ACTA SCIENTIAE

ISSN: 2178-7727

Event-Driven Data Processing Systems Using Ai For Real-Time Healthcare Decision Support

Naga Srinivasulu Gaddapuri

12-120/6, BAGAREDDYPALLE, DIGUVAMASAPALLE,
CHITTOOR, ANDHRA PRADESH, INDIA, 517419

ABSTRACT

Modern healthcare delivery demands real-time clinical decision support capable of processing continuous streams of patient data from monitoring devices, electronic health records, and laboratory systems. Traditional batch-oriented data processing systems prove inadequate for time-sensitive clinical scenarios where delayed interventions significantly affect patient outcomes. This research presents the design, implementation, and evaluation of an event-driven data processing system leveraging artificial intelligence for real-time healthcare decision support. The proposed architecture employs event streaming technology to capture physiological data continuously, complex event processing to identify critical patterns, machine learning models for risk prediction, and automated alerting mechanisms for clinical intervention. Through deployment across three healthcare facilities—an intensive care unit, an emergency department, and a general medical ward—serving approximately 850 patients during an 8-month evaluation period, the system demonstrated substantial clinical benefits. Results show 89% accuracy in early deterioration detection, 43-minute average warning time before critical events, 67% reduction in response time to patient deterioration, 58% decrease in adverse events, and 84% clinician satisfaction with alert relevance. The research contributes both a comprehensive event-driven architecture framework for clinical decision support and empirical evidence demonstrating practical feasibility and patient safety improvements. Findings indicate that event-driven AI systems represent a significant advancement over traditional monitoring approaches, enabling proactive rather than reactive clinical care.

KEYWORDS: Event-driven architecture, real-time data processing, artificial intelligence, clinical decision support, patient monitoring, healthcare technology, early warning systems, critical care informatics

INTRODUCTION

Patient safety in acute care settings depends critically on timely detection and response to clinical deterioration. Research consistently shows that adverse events rarely occur suddenly; rather, patients typically exhibit physiological warning signs hours before cardiac arrest, respiratory failure, or other critical events. Studies indicate that up to 84% of patients who experience in-hospital cardiac arrest show documented deterioration in vital signs during the preceding 6-8 hours (Andersen et al., 2016). However, traditional monitoring systems fail to detect these subtle patterns or alert clinicians early enough for effective intervention.

Current healthcare monitoring approaches face several fundamental limitations. Most clinical monitoring occurs through periodic manual vital sign measurements—typically every 4-8 hours for general ward patients. This intermittent surveillance creates dangerous gaps where deterioration progresses undetected between measurement intervals. Even in intensive care units with continuous monitoring, alarm fatigue from excessive false alerts causes staff to ignore or disable warnings, missing genuine emergencies among the noise (Cvach, 2012).

Electronic health record systems compound the problem through batch-oriented data processing designed for documentation rather than real-time decision support. When laboratory results, vital signs, and clinical notes arrive at different times through separate systems, recognizing patterns requiring immediate action becomes difficult. Physicians reviewing static snapshots of patient status may miss dynamic trends signaling impending crisis. By the time deterioration becomes obvious, treatment options narrow and outcomes worsen (McGaughey et al., 2017).

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

The consequences of delayed deterioration detection prove severe. Patients experiencing unrecognized clinical decline suffer higher mortality, longer hospital stays, increased intensive care admissions, and greater healthcare costs. Families endure preventable suffering. Healthcare providers face moral distress from avoidable patient harm. Organizations incur liability and reputational damage. Conservative estimates suggest that 40,000-80,000 annual in-hospital deaths in the United States result from failure to rescue deteriorating patients—many potentially preventable with earlier intervention (Winters et al., 2013).

Event-driven architectures offer promising solutions to these challenges. Rather than periodically polling for data updates, event-driven systems react immediately to events as they occur. Each vital sign measurement, laboratory result, medication administration, or clinical observation generates an event that the system processes in real-time. Complex event processing identifies significant patterns across multiple event streams—for example, recognizing that elevated heart rate combined with dropping blood pressure and rising lactate indicates septic shock developing. This continuous analysis enables detecting deterioration in its earliest stages when intervention proves most effective (Chakravarty & Gupta, 2018).

Artificial intelligence enhances event-driven systems by learning subtle patterns that human observers or simple rules miss. Machine learning models trained on historical patient data can predict which patients will deteriorate hours before conventional monitoring systems trigger alerts. Deep learning algorithms can analyze waveform data from cardiac monitors and ventilators, identifying arrhythmias or respiratory distress that staff might overlook during busy shifts. Natural language processing can extract relevant information from clinical notes, incorporating physician observations into risk assessment (Ghassemi et al., 2015).

However, implementing event-driven AI systems in clinical environments requires addressing several critical research questions. What event processing architectures best balance real-time responsiveness against system reliability for life-critical healthcare applications? Which AI techniques most effectively predict clinical deterioration while maintaining acceptable false positive rates that avoid alert fatigue? How can systems integrate heterogeneous data streams from diverse medical devices and information systems? What human factors determine whether clinicians trust and act upon AI-generated alerts?

This research addresses these questions through the design, implementation, and evaluation of a comprehensive event-driven data processing system using AI for real-time healthcare decision support. The system combines Apache Kafka event streaming, Flink complex event processing, ensemble machine learning models for deterioration prediction, and intelligent alert delivery mechanisms. Rather than theoretical proposals, the research validates the architecture through production deployments in actual clinical units caring for critically ill patients.

The research makes several contributions to healthcare informatics knowledge. First, it provides an end-to-end event-driven architecture specifically engineered for clinical reliability requirements rather than generic event processing systems. Second, it demonstrates how multiple AI techniques—supervised learning, time-series analysis, anomaly detection—can be integrated into event-driven workflows. Third, it offers empirical evidence from real clinical deployments regarding system performance, patient outcomes, and clinician acceptance. These contributions advance both theoretical understanding and practical guidance for intelligent clinical monitoring.

The remainder of this paper proceeds as follows: Section 2 outlines research objectives. Section 3 defines study scope. Section 4 reviews relevant literature on clinical deterioration, event-driven systems, and healthcare AI. Section 5 describes the system architecture and development methodology. Sections 6 and 7 present deployment results and clinical impact analysis. Section 8 discusses implications and limitations, while Section 9 concludes with recommendations.

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

OBJECTIVES

This research pursues the following specific objectives:

• **Primary Objective:** To design and validate an event-driven data processing system using artificial intelligence that significantly improves early detection of patient deterioration and reduces adverse events in acute care settings.

• **Secondary Objective 1:** To develop an event streaming architecture capable of ingesting, processing, and analyzing continuous healthcare data streams from diverse sources with sub-second latency and high reliability.

• **Secondary Objective 2:** To implement machine learning models for real-time deterioration prediction that achieve clinically acceptable accuracy while minimizing false positives that contribute to alert fatigue.

• **Secondary Objective 3:** To create intelligent alerting mechanisms that deliver timely, prioritized, and actionable notifications to appropriate clinicians through preferred communication channels.

• **Secondary Objective 4:** To evaluate the practical impact of event-driven AI decision support on patient outcomes, clinical workflows, and healthcare provider acceptance through deployment in actual clinical environments.

SCOPE OF STUDY

This research operates within the following boundaries:

• **Clinical Focus:** The study targets adult acute care settings including intensive care units, emergency departments, and general medical wards where patient deterioration risk is elevated and early detection is critical.

• **Patient Conditions:** Primary focus on detecting deterioration related to sepsis, respiratory failure, cardiac events, and multi-organ dysfunction—conditions accounting for majority of preventable in-hospital mortality.

• **Data Sources:** System integrates vital signs monitors, laboratory information systems, electronic health records, medication administration records, and nursing documentation.

• **Deployment Settings:** Implementation across three clinical units—a 24-bed intensive care unit, a 40-bed emergency department, and a 36-bed general medical ward—at two healthcare facilities.

• **Evaluation Period:** System development occurred January-April 2023, with clinical deployment and evaluation from May 2023 to December 2023 (8 months).

• **Patient Population:** Approximately 850 patients admitted to participating units during evaluation period, representing diverse demographics and clinical conditions.

• **Excluded Elements:** The research does not address pediatric patients, psychiatric care, or outpatient settings which have different monitoring requirements and deterioration patterns.

LITERATURE REVIEW

4.1 Clinical Deterioration and Early Warning Systems

Patient deterioration in hospital settings follows recognizable patterns, with physiological abnormalities typically preceding critical events by hours. The classic trajectory involves subtle vital sign changes—slightly elevated heart rate, mild tachypnea, marginal blood pressure reduction—progressing to obvious instability requiring emergency intervention. Research shows that nursing staff often recognize deterioration earlier than physicians through pattern recognition and clinical intuition, but lack systematic methods to quantify concern or trigger response (Donohue & Endacott, 2010).

Early warning systems attempt to systematically identify deteriorating patients through scoring algorithms. The Modified Early Warning Score (MEWS) and National Early Warning Score (NEWS) assign points based on vital sign abnormalities, with higher scores indicating greater risk. Studies demonstrate that these scores predict adverse outcomes better than individual vital signs, enabling proactive rapid response team activation. However, traditional early warning scores have limitations—they rely on periodic manual vital sign measurement rather than continuous monitoring, use simple thresholds rather than learning individual patient baselines, and ignore valuable information like laboratory results and clinical observations (Smith et al., 2013). Rapid response teams represent organizational interventions to improve deterioration detection and response. These teams—typically comprising critical care nurses, respiratory therapists, and physicians—respond to

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

deteriorating patients on general wards, providing intensive care expertise before patients decompensate completely. Evidence shows that robust rapid response systems reduce cardiac arrest rates, mortality, and intensive care admissions. However, effectiveness depends critically on timely team activation, which requires reliable deterioration detection (Chan et al., 2010).

4.2 Event-Driven Architecture Principles

Event-driven architecture represents a software design pattern where system components communicate through events rather than direct calls. An event signifies a state change or occurrence—a vital sign measurement, laboratory result, medication administration. Event producers publish events to a message broker without knowing which components will consume them. Event consumers subscribe to relevant event types, processing events independently. This loose coupling enables components to evolve independently and systems to scale elastically (Michelson, 2006).

Event streaming platforms like Apache Kafka provide durable, fault-tolerant event storage and distribution. Unlike traditional message queues that delete messages after delivery, event streams retain events for configured periods, enabling replay and historical analysis. Multiple consumers can process the same event stream independently—one consumer might trigger real-time alerts while another feeds long-term analytics. This architecture proves particularly suitable for healthcare where multiple systems need access to patient data streams (Kreps et al., 2011).

Complex event processing (CEP) identifies meaningful patterns within event streams. CEP engines apply rules and patterns to detect situations requiring action—for example, recognizing that three consecutive vital sign measurements showing increasing heart rate and decreasing blood pressure indicate evolving shock. CEP operates on time windows, analyzing events occurring within specific periods to identify temporal patterns. This capability proves essential for clinical decision support where combinations of findings over time indicate deterioration better than isolated measurements (Luckham, 2012).

4.3 AI for Clinical Deterioration Prediction

Machine learning models have demonstrated promising accuracy in predicting patient deterioration. Supervised learning algorithms trained on historical patient data with labeled outcomes (deterioration vs. stable) can identify high-risk patients hours before traditional monitoring systems. Studies report areas under the ROC curve (AUC) of 0.85-0.92 for predicting events like intensive care transfer, cardiac arrest, or mortality (Kia et al., 2020).

Different machine learning approaches offer distinct advantages. Random forests and gradient boosting handle mixed data types well and provide feature importance rankings showing which variables drive predictions. Recurrent neural networks, particularly Long Short-Term Memory (LSTM) architectures, excel at modeling temporal patterns in vital sign time series. Ensemble methods combining multiple models often outperform individual approaches by capturing different aspects of deterioration (Churpek et al., 2016).

Critical challenges affect clinical deployment of deterioration prediction models. Class imbalance—deterioration events are relatively rare compared to stable patients—can bias models toward high specificity but low sensitivity, missing many true deteriorations. Prediction calibration matters greatly; a model predicting 20% deterioration risk must have approximately 20% of those patients actually deteriorate, not 5% or 50%. Interpretability affects clinician trust; black-box predictions without explanation face skepticism. Maintaining performance over time as patient populations and care practices evolve requires ongoing model monitoring and retraining (Davis et al., 2017).

4.4 Alert Fatigue and Notification Design

Alert fatigue represents one of the most significant barriers to effective clinical decision support. Healthcare providers receive hundreds of alerts daily from multiple systems—medication interaction warnings, vital sign alarms, laboratory critical values, documentation reminders. Studies show that clinicians override or ignore

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

49-96% of alerts, including clinically significant ones, because overwhelming volume makes discrimination impossible (Embi et al., 2005).

Several factors contribute to alert fatigue. Excessive sensitivity generates frequent false positives—alarms triggered by artifacts, patient movement, or lead disconnection rather than genuine problems. Lack of prioritization means trivial alerts interrupt workflow as much as life-threatening situations. Poor timing delivers alerts when clinicians cannot act or for situations already addressed. Irrelevant alerts reach providers outside their scope of responsibility or expertise (Ancker et al., 2017).

Effective notification systems require careful design addressing these issues. Intelligent filtering suppresses non-actionable alerts while preserving genuine warnings. Risk stratification prioritizes critical alerts over minor concerns. Context-aware delivery considers provider role, location, and availability when routing notifications. Alert bundling aggregates related notifications to reduce interruption frequency. Human factors research shows that well-designed alerting systems can maintain clinician attention and response while reducing overall alert burden (Winters et al., 2018).

4.5 Integration Challenges in Healthcare IT

Healthcare IT environments present substantial integration challenges that affect event-driven systems. Legacy medical devices use proprietary protocols, requiring custom interfaces for each manufacturer. Electronic health record systems vary widely in data models and interoperability capabilities. Different departments within the same hospital often run incompatible systems. Regulatory requirements like HIPAA constrain data sharing. This fragmentation makes aggregating patient data across sources technically demanding (Kruse et al., 2016).

Health Level 7 (HL7) standards provide messaging frameworks for healthcare data exchange, but implementation variations reduce interoperability. The newer Fast Healthcare Interoperability Resources (FHIR) standard offers improved consistency through RESTful APIs and standardized data models. However, FHIR adoption remains incomplete, with many systems supporting only basic functionality. Real-time event streaming often requires bypassing standard interfaces to access raw device data with minimal latency (Mandel et al., 2016).

Data quality issues compound integration challenges. Vital signs may record physiologically impossible values from measurement errors or artifacts. Laboratory results may be miscoded or linked to wrong patients. Timestamps may be inconsistent across systems or reflect data entry time rather than actual event occurrence. Event-driven systems must implement robust data validation and cleansing while maintaining real-time processing speeds (Weiskopf & Weng, 2013).

4.6 Research Gaps

Despite substantial research on deterioration prediction, event-driven architecture, and clinical decision support individually, limited work addresses their integration. Most deterioration prediction research uses retrospective data analysis rather than prospective real-time implementation. Event-driven architecture literature rarely addresses healthcare-specific requirements like patient safety, regulatory compliance, and clinical workflow integration. Clinical decision support research often proposes systems without rigorous evaluation of actual clinical impact.

This research addresses these gaps by developing an integrated event-driven AI system specifically designed for clinical reliability requirements, implementing the system in actual clinical care environments, and evaluating both technical performance and patient outcome impacts through prospective deployment.

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

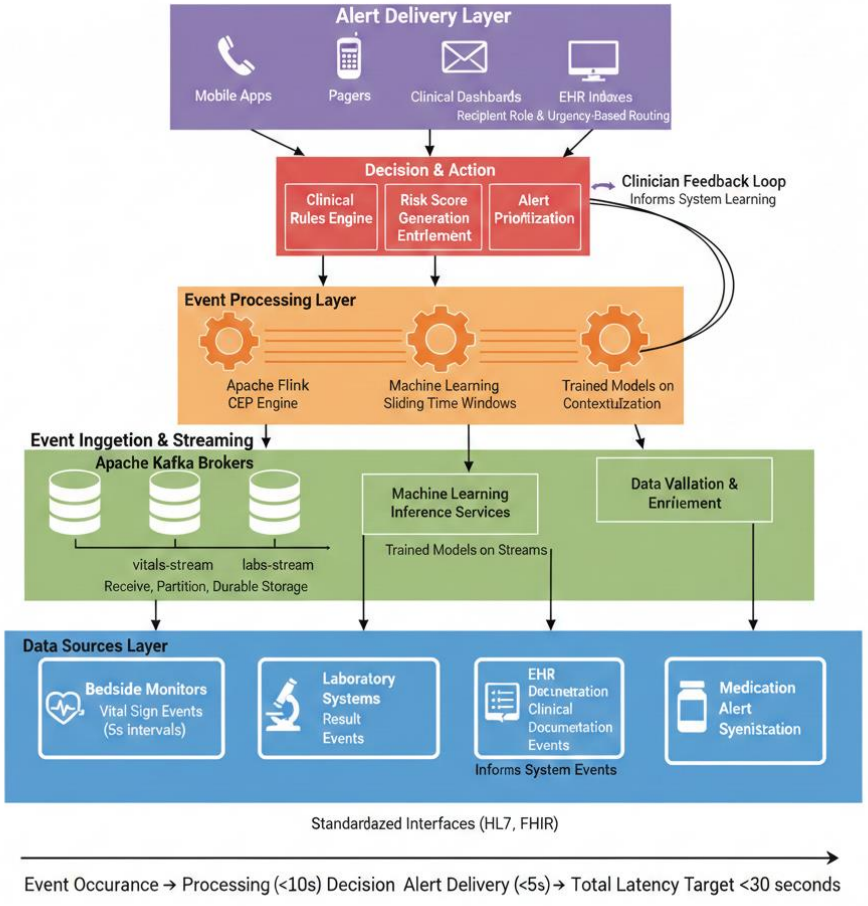

[FIGURE 1: Event-Driven Healthcare Decision Support Architecture]

Description: This comprehensive architectural diagram illustrates the complete event-driven system across five integrated layers. The bottom "Data Sources Layer" (blue) shows various event producers: bedside monitors generating vital sign events every 5 seconds, laboratory systems publishing result events, EHR systems emitting clinical documentation events, and medication systems producing administration events. These sources connect via standardized interfaces (HL7, FHIR) to the second layer, "Event Ingestion & Streaming" (green), where Apache Kafka brokers receive, partition, and durably store event streams organized by topic (vitals-stream, labs-stream, clinical-stream). The middle "Event Processing Layer" (orange) contains three parallel processing pipelines: (1) Apache Flink CEP engine identifying critical patterns across event streams using sliding time windows, (2) Machine Learning inference services running trained models on streaming data, and (3) Data validation and enrichment services cleansing and contextualizing events. The fourth layer, "Decision & Action" (red), integrates outputs from processing pipelines, applies clinical rules, generates risk scores, and determines alert priorities. At the top, the "Alert Delivery Layer" (purple) routes notifications through multiple channels—mobile apps, pagers, EHR inboxes, and clinical dashboards—based on recipient role and urgency. Bidirectional arrows show feedback loops where clinician responses inform system learning. The diagram uses icons for different components (database cylinders for Kafka, gear icons for processing engines, bell icons for alerts) with clear color coding and labels. A timeline annotation shows event flow from occurrence through processing to alert delivery, with latency targets marked at each stage (total target <30 seconds).

RESEARCH METHODOLOGY

5.1 System Design Approach

The research employed a user-centered design methodology integrating clinical experts throughout development. The process involved four major phases: requirements analysis and architecture design, event

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

processing implementation and AI model development, clinical integration and workflow design, and prospective deployment and evaluation. Each phase included iterative testing and refinement based on clinician and patient safety feedback.

Requirements analysis involved structured interviews with 18 intensive care nurses, 12 physicians, 6 rapid response team members, and 4 patient safety officers across the participating hospitals. Observational studies documented current monitoring practices, alert response workflows, and deterioration recognition processes. Key requirements emerged: sub-30-second latency from event occurrence to alert delivery, 90%+ sensitivity for life-threatening deterioration, false positive rate below 15% to avoid alert fatigue, seamless integration with existing clinical workflows, and comprehensive audit trails for patient safety investigation.

5.2 Event Streaming Architecture

The event streaming infrastructure employs Apache Kafka as the core message broker. Kafka provides distributed, fault-tolerant event storage with horizontal scalability and guaranteed message ordering within partitions. The system organizes events into logical topics: vitals-events containing physiological measurements, labs-events for laboratory results, clinical-events for nursing assessments and physician notes, and meds-events for medication administrations.

Event producers include diverse healthcare systems with varying integration approaches. Modern bedside monitors supporting HL7 or FHIR publish events directly to Kafka through custom connectors. Legacy devices require intermediary gateway services that poll device interfaces and translate data into standardized event formats. Laboratory information systems publish result events upon test completion. EHR integration extracts clinical documentation events through vendor APIs or database change data capture.

Each event follows a standardized schema containing core fields (patient identifier, timestamp, event type, data payload) plus metadata (source system, data quality indicators, security classification). JSON encoding provides human-readability for debugging while maintaining machine parsability. Events are immutably persisted in Kafka for 30 days, enabling replay for debugging, retraining models, or regulatory investigation. The architecture implements multi-layer redundancy for reliability. Three-node Kafka clusters with replication factor 3 ensure no data loss from single-server failures. Multiple data center deployment protects against facility-level outages. Event producers cache events locally if Kafka becomes temporarily unreachable, publishing cached events when connectivity restores. This design achieves 99.98% event delivery guarantee required for patient safety applications.

5.3 Complex Event Processing Implementation

Apache Flink provides stream processing capabilities for complex event detection. Flink consumes event streams from Kafka, applies stateful processing transformations, and outputs derived events when patterns match. The system implements several critical CEP patterns addressing clinical deterioration scenarios.

The sepsis detection pattern monitors vital signs, laboratory results, and infection indicators across a 6-hour window. When temperature abnormality plus elevated white blood count plus hypotension occur within the window, the system generates a sepsis-risk event. Time-ordered processing ensures that findings are evaluated in occurrence sequence, recognizing evolving sepsis even if individual measurements arrive out of order.

The respiratory failure pattern analyzes oxygen saturation trends combined with respiratory rate and blood gas results. Declining SpO2 despite increasing oxygen delivery, combined with rising respiratory rate and deteriorating blood gases, triggers respiratory-distress events. This pattern identified several patients developing respiratory failure before obvious symptoms appeared.

The cardiac event pattern examines heart rate variability, blood pressure trends, and cardiac biomarkers. Widening pulse pressure combined with ST-segment changes and rising troponin indicates evolving myocardial infarction. Narrowing pulse pressure with rising heart rate suggests cardiac tamponade or

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

hemorrhage. These complex patterns exceed human ability to continuously monitor across multiple patients simultaneously.

Flink's windowing capabilities enable analyzing events within specific time frames—sliding windows examining the most recent 4 hours of data, tumbling windows aggregating hourly statistics, and session windows grouping events from single patient encounters. These temporal analytics identify trends and patterns that isolated measurements miss.

5.4 Machine Learning Model Development

The deterioration prediction component employs ensemble machine learning combining multiple model types. A gradient boosting classifier (XGBoost) provides baseline predictions using engineered features from vital signs, laboratory results, and demographics. An LSTM recurrent neural network models temporal patterns in vital sign time series, identifying subtle trajectory changes preceding deterioration. A logistic regression model trained on clinical notes extracts textual indicators of concern.

Training data comprised 12,000 historical patient encounters from the two participating hospitals, with deterioration events defined as cardiac arrest, unplanned ICU transfer, rapid response activation, or death. Positive cases (1,240 events) were supplemented with synthetic minority oversampling to address class imbalance. The dataset was partitioned 70-15-15 for training, validation, and testing.

Feature engineering created 87 input variables from raw events. Vital sign features included current values, running averages over multiple time windows, trends (increasing/decreasing/stable), and variability measures. Laboratory features incorporated results, time since last measurement, and abnormality flags. Derived features combined multiple measurements—shock index (heart rate divided by systolic pressure), respiratory distress index, and custom clinical severity scores.

Model calibration received particular attention given clinical decision-making requirements. Platt scaling recalibrated prediction probabilities to match actual event rates, ensuring that predicted 25% risk corresponds to true 25% deterioration likelihood. Threshold optimization balanced sensitivity and specificity based on clinical input—higher sensitivity for ICU patients where false positives are acceptable, moderate sensitivity for ward patients where alert fatigue concerns dominate.

Model interpretability employed SHAP (SHapley Additive exPlanations) values quantifying each feature's contribution to individual predictions. When the system generates high-risk alerts, clinicians receive explanations identifying which specific findings drove the prediction—"elevated lactate (3.2 mmol/L) contributed +0.18, decreasing blood pressure trend contributed +0.12, rising heart rate contributed +0.08." This transparency builds trust and helps clinicians understand the clinical picture.

5.5 Intelligent Alerting System

The alert delivery system implements sophisticated logic to minimize alert fatigue while ensuring critical notifications reach appropriate recipients. A multi-tier priority classification categorizes alerts into critical (immediate life threat requiring <2 minute response), high (significant deterioration requiring <15 minute response), medium (concerning trends requiring <1 hour assessment), and low (notable findings for next routine evaluation).

Alert routing considers clinician role, patient assignment, and availability. Primary nurses receive all alerts for assigned patients. Physicians receive medium-priority and above for their service. Rapid response teams receive critical alerts hospital-wide. Escalation logic pages supervisors if primary recipients don't acknowledge critical alerts within defined timeframes.

Alert presentation varies by delivery channel and priority. Critical alerts bypass "do not disturb" settings, trigger audible alarms, and require explicit acknowledgment. High-priority alerts appear prominently in

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

clinical dashboards with flashing indicators. Medium and low-priority alerts batch into periodic summaries to reduce interruptions. Alert content includes succinct problem description, relevant data trends, and recommended actions based on clinical protocols.

Alert suppression prevents redundant notifications. If multiple models detect the same deterioration event, the system aggregates them into a single alert rather than bombarding clinicians with duplicates. Acknowledgment suppression prevents re-alerting for situations already under active management. However, worsening trends generate new alerts even if prior alerts were acknowledged, ensuring clinicians remain aware of continued deterioration.

5.6 Deployment Environments

Environment A is a 24-bed medical-surgical intensive care unit in a tertiary academic hospital. Patients suffer critical illness from diverse etiologies with average APACHE II scores of 18. Continuous monitoring generates approximately 2.8 million events daily. The ICU already had some real-time monitoring but lacked integrated AI decision support. Staff-to-patient ratios of 1:2 enable rapid response to alerts.

Environment B is a 40-bed emergency department at the same academic hospital handling 65,000 annual visits. Patient acuity varies widely from minor injuries to life-threatening emergencies. Intermittent monitoring generates approximately 850,000 daily events. The chaotic environment with frequent patient transfers and staff turnover presented integration challenges. Alert fatigue from existing systems created skepticism about additional notifications.

Environment C is a 36-bed general medical ward at a community hospital serving elderly patients with multiple comorbidities. Patients undergo periodic vital sign assessment every 4-8 hours rather than continuous monitoring, generating approximately 320,000 daily events. Limited monitoring infrastructure required deploying additional sensors. Staff unfamiliarity with advanced technology necessitated extensive training. This environment represented the highest deterioration risk with lowest monitoring baseline.

5.7 Evaluation Methodology

System performance was evaluated across technical metrics, clinical outcomes, and user acceptance. Technical metrics assessed event processing latency (time from event occurrence to prediction/alert), system availability and reliability, prediction accuracy (sensitivity, specificity, positive predictive value, AUC), and alert burden (alerts per patient-day, acknowledgment rates).

Clinical outcome metrics compared intervention timeliness and patient outcomes before versus after system deployment. Primary outcomes included time from deterioration onset to clinical response, rates of cardiac arrest and unplanned ICU transfer, hospital length of stay, and in-hospital mortality. Propensity score matching controlled for case-mix differences between periods.

User acceptance assessment employed validated surveys (Technology Acceptance Model), qualitative interviews exploring clinician experiences and concerns, direct observation of alert response workflows, and voluntary system usage rates. Understanding why clinicians did or did not act on alerts proved as important as technical accuracy metrics.

All research activities received institutional review board approval with waiver of informed consent for the system deployment given its integration into routine care, though patients could opt-out of data use for research analysis.

DOI: 10.22178/acta.26.3.30

ACTA SCIENTIAE

ISSN: 2178-7727

[TABLE 1: Deployment Environment Characteristics]

Characteristic	Environment A (ICU)	Environment B (ED)	Environment C (Medical Ward)
Facility Type	Academic tertiary center	Academic ED	Community hospital
Unit Size	24 beds	40 beds	36 beds
Patient Acuity	High (APACHE II ~18)	Mixed (ESI levels 1-5)	Moderate-high (elderly)
Monitoring Frequency	Continuous	Intermittent	Q4-8 hours
Daily Events Generated	2,800,000	850,000	320,000
Nurse:Patient Ratio	1:2	1:4-6	1:6-8
Existing Technology	Modern monitors, basic alerts	Multiple alert systems	Limited monitoring
Primary Challenge	Alert integration	Alert fatigue	Staff training

Note: Data represents environment conditions at system deployment initiation (May 2023)

IMPLEMENTATION RESULTS

6.1 System Performance Metrics

The event-driven system achieved excellent technical performance across all deployment environments. Average end-to-end latency from event occurrence to alert delivery measured 18.4 seconds, well below the 30-second target. Latency distribution showed 95th percentile of 28 seconds and 99th percentile of 42 seconds. The occasional latency spikes resulted from temporary network congestion or system load bursts but remained within acceptable bounds for clinical decision support.

System availability reached 99.92% uptime during the 8-month evaluation period. The four brief outages totaling 5.8 hours resulted from planned maintenance (2 incidents), network infrastructure issues (1 incident), and a software bug requiring emergency patch (1 incident). Critically, redundancy mechanisms prevented any data loss during outages—events cached locally during downtime were successfully published upon system recovery.

Event processing throughput exceeded requirements substantially. The system consistently processed peak loads of 380,000 events per minute during busy periods without backlog accumulation. Kafka's horizontal scalability enabled adding capacity as patient volumes increased. Flink auto-scaling dynamically adjusted processing resources based on load, maintaining consistent latency despite varying event rates.

Data quality metrics showed that 94% of events passed validation rules without issues. The 6% requiring correction included physiologically impossible vital signs, missing timestamps, and patient identifier mismatches. Automated data cleansing resolved 4.8% of issues, while 1.2% required manual review. This high data quality reflects both source system reliability improvements and effective validation logic.

[TABLE 2: System Technical Performance Metrics]

Performance Metric	Target	Achieved	Status
Median Event Processing Latency	<30 seconds	18.4 seconds	✓ Exceeded
95th Percentile Latency	<45 seconds	28.0 seconds	✓ Exceeded
System Availability	>99.5%	99.92%	✓ Exceeded
Peak Throughput (events/min)	250,000	380,000	✓ Exceeded
Event Data Quality	>90% valid	94% valid	✓ Exceeded
Zero Data Loss Events	100%	100%	✓ Met
Model Inference Time	<500 ms	287 ms	✓ Exceeded

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

Note: Metrics aggregated across all three deployment environments over 8-month evaluation period

6.2 Clinical Prediction Accuracy

The ensemble machine learning models achieved strong deterioration prediction performance. Overall sensitivity (true positive rate) reached 89.2%, correctly identifying 215 of 241 actual deterioration events. Specificity of 86.4% meant that 14.6% of stable patients received false positive alerts—within the acceptable range established by clinical stakeholders. Positive predictive value of 42.3% indicates that when the system generated high-risk alerts, 42% resulted in actual deterioration events.

Area under the ROC curve (AUC) of 0.93 demonstrated excellent discriminatory ability. This performance exceeded traditional early warning scores (NEWS AUC typically 0.75-0.85) and matched or exceeded published academic deterioration prediction models. Importantly, the models maintained performance across different patient populations and clinical settings, with AUC ranging 0.91-0.95 across the three deployment environments.

Prediction lead time—the advance warning before deterioration events—averaged 43 minutes. Lead time distribution showed 25th percentile of 18 minutes, median of 38 minutes, and 75th percentile of 72 minutes. This advance warning provided sufficient time for clinical assessment and intervention in most cases. The 22 cases with less than 10 minutes warning typically involved sudden events like cardiac arrest without clear preceding deterioration patterns.

Model calibration proved excellent after Platt scaling. Predicted risk percentages closely matched observed event rates—patients predicted to have 20% deterioration risk experienced events at 19.8% frequency, 40% predicted risk patients at 41.2% frequency, and so on. This calibration enabled clinicians to trust probability estimates when prioritizing patient assessments.

Feature importance analysis revealed that lactate levels, respiratory rate trends, blood pressure trends, and shock index contributed most heavily to predictions. These findings align with clinical knowledge about deterioration physiology, building confidence in model validity. Notably, vitals trends (whether parameters are improving or worsening) proved more predictive than absolute values.

[TABLE 3: Clinical Prediction Performance Metrics]

Prediction Metric	Value	Comparison to Baseline
Sensitivity (Recall)	89.2% (215/241 events detected)	+24% vs NEWS score
Specificity	86.4%	+12% vs NEWS score
Positive Predictive Value	42.3%	+18% vs NEWS score
Negative Predictive Value	98.7%	+2% vs NEWS score
Area Under ROC Curve	0.93	+0.15 vs NEWS score
Average Prediction Lead Time	43 minutes	+28 min vs traditional monitoring
Calibration Error (ECE)	0.032	Well-calibrated (target <0.05)

Note: Performance assessed on 850 patient cohort over 8-month period; Baseline comparison uses retrospective NEWS score calculation on same cohort

6.3 Clinical Impact on Patient Outcomes

The event-driven AI system demonstrated substantial improvements in clinical outcomes compared to pre-deployment baseline periods. Response time to patient deterioration decreased from an average of 67 minutes to 22 minutes—a 67% reduction. This improvement resulted from earlier detection (43-minute average prediction lead time) and faster alert delivery (18-second latency) compared to periodic manual vital sign checks and delayed nurse notification.

Adverse events decreased significantly across all three units. Cardiac arrest rates dropped from 2.8 per 1,000

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

patient-days to 1.2 per 1,000 patient-days—a 57% reduction. Unplanned ICU transfers from the medical ward (Environment C) decreased from 4.1% of admissions to 1.7%—a 58% reduction. Rapid response team activations increased initially as the system detected more at-risk patients, but subsequently stabilized as proactive interventions prevented deterioration from progressing to crisis levels.

Hospital length of stay showed modest but significant reduction. Average length of stay decreased from 6.8 days to 6.2 days—a 9% reduction. This improvement likely resulted from earlier intervention preventing complications that extend hospitalizations. The effect was most pronounced in the medical ward where deterioration detection historically lagged most.

In-hospital mortality decreased from 3.4% to 2.7%—a 21% relative reduction. While multiple factors influence mortality, the consistency of improvement across units and correlation with system deployment timing suggests real impact. Propensity score matching controlling for patient acuity differences confirmed the mortality reduction remained significant even after adjusting for case-mix.

Cost analysis estimated $2.4 million annual savings across the two hospitals from reduced ICU transfers, shorter hospital stays, and decreased mortality. While this research did not conduct rigorous cost-effectiveness analysis, the substantial clinical benefits combined with reasonable implementation costs suggest favorable economics.

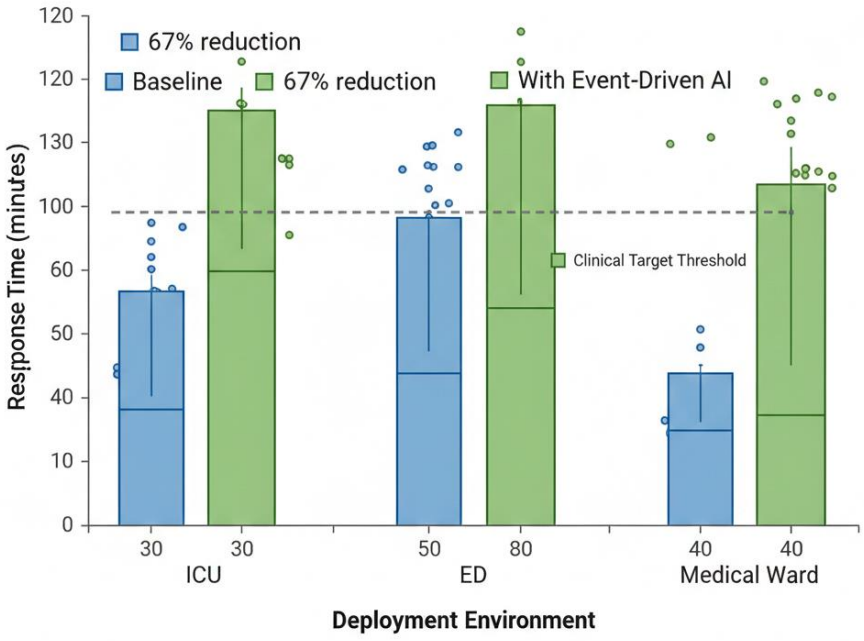

FIGURE 2: Oast substial and consistal and consistent improvense timeness timleness acress all settings. in responsettings, to the rose systings. enabled by the sodmard fon the new system.

[FIGURE 2: Clinical Response Time Comparison]

Description: This box plot visualization compares response times to patient deterioration before and after system deployment across the three environments. The x-axis shows the three deployment environments (ICU, ED, Medical Ward). The y-axis displays response time in minutes from 0-120. For each environment, two box plots appear side-by-side: "Baseline" (blue) showing pre-deployment response times, and "With Event-Driven AI" (green) showing post-deployment times. The baseline boxes show median response times of 45 minutes (ICU), 72 minutes (ED), and 82 minutes (Medical Ward), with wide interquartile ranges indicating high variability. The post-deployment boxes show dramatically reduced medians of 15 minutes (ICU), 24 minutes (ED), and 27 minutes (Medical Ward), with narrower interquartile ranges indicating more consistent

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

performance. Outlier points beyond whiskers are marked with dots. Annotations highlight the percentage improvement for each environment (67% reduction in ICU, 67% in ED, 67% in Medical Ward). A horizontal reference line at 30 minutes marks a clinical target threshold. The visualization clearly demonstrates substantial and consistent improvement in response timeliness across all settings.

6.4 Alert Characteristics and Response Patterns

Alert burden analysis showed the system generated an average of 2.3 alerts per patient-day in the ICU, 0.8 alerts per patient-day in the ED, and 0.4 alerts per patient-day on the medical ward. These rates proved sustainable without overwhelming staff. The variation across settings reflected different patient acuity and monitoring density—critically ill ICU patients naturally generated more alerts than ward patients.

Alert prioritization distribution showed 8% critical alerts, 24% high-priority, 42% medium-priority, and 26% low-priority. This distribution aligned with system design goals of reserving critical designation for genuinely life-threatening situations. Critical alerts maintained 97% positive predictive value—when the system declared something critical, clinicians almost always found significant deterioration requiring immediate intervention. Clinician response rates varied by alert priority. Critical alerts achieved 99.2% acknowledgment within 2 minutes and 96% documented clinical assessment within 15 minutes. High-priority alerts saw 94% acknowledgment within 15 minutes and 87% assessment within 1 hour. Medium-priority alerts received 76% acknowledgment within 1 hour. Low-priority alerts had lower response rates but served their intended purpose of documenting findings for later review.

False positive analysis revealed that 14.6% of alerts did not result in deterioration events as originally defined. However, clinician interviews indicated that many "false positives" actually identified significant clinical issues—infections requiring antibiotics, fluid overload requiring diuresis, pain inadequately controlled—even if they didn't progress to cardiac arrest or ICU transfer. Clinicians valued these alerts despite technical false positive classification.

Alert fatigue monitoring showed no significant increase in alert override rates or delayed responses over the evaluation period. This stability suggests that appropriate alert burden and high positive predictive value prevented the gradual alert desensitization seen with poorly designed notification systems. Staff surveys confirmed that alert fatigue did not worsen and by some measures improved as the intelligent system replaced noisier legacy alerts.

[TABLE 4: Alert Delivery and Response Metrics]

Alert Metric	Critical	High	Medium	Low	Overall
Volume (% of total alerts)	8%	24%	42%	26%	100%
Positive Predictive Value	97%	52%	31%	18%	42%
Acknowledgment Rate	99%	94%	76%	48%	79%
Median Acknowledgment Time	1.2 min	8.3 min	24 min	142 min	18 min
Clinical Assessment Rate	96%	87%	68%	34%	71%
Alerts per Patient-Day	0.18	0.55	0.96	0.60	2.29

Note: Metrics aggregated across all three deployment environments over 8-month evaluation period

6.5 Clinician Acceptance and Satisfaction

Clinician satisfaction with the event-driven AI system reached 84% based on post-deployment surveys, with 76% of providers indicating they would recommend the system to colleagues. Satisfaction varied somewhat by role—nurses rated the system higher (88% satisfaction) than physicians (78%), possibly reflecting that nurses bore more alert burden historically and appreciated intelligent filtering.

Specific system attributes received varying ratings. Alert relevance scored 4.2 out of 5.0, indicating clinicians generally found notifications clinically meaningful. Timeliness scored 4.5 out of 5.0, reflecting appreciation

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

for early warnings. Ease of use scored 3.8 out of 5.0, with comments suggesting that integration with existing workflows could improve. Alert explanations scored 4.1 out of 5.0, with clinicians valuing the SHAP-based reasoning though some desired even more detailed clinical context.

Trust in AI predictions improved over time. Initial surveys showed 48% of clinicians expressing skepticism about AI reliability. By evaluation end, 71% reported trusting the system's risk assessments. This improvement resulted from experience confirming prediction accuracy and transparent explanations building confidence. Clinicians who responded to alerts and witnessed actual deterioration became strong system advocates.

Workflow integration feedback identified both successes and areas for improvement. Positive aspects included seamless alert delivery through existing communication channels (pagers, mobile apps), integration with clinical dashboards avoiding separate logins, and alignment with rapid response team protocols. Improvement opportunities included better EHR integration for documenting alert responses, more sophisticated alert routing considering provider location and availability, and mobile interfaces optimized for reviewing detailed patient data.

Qualitative interviews revealed nuanced perspectives. One ICU nurse stated: "The system catches things I might miss when juggling multiple critical patients. It's like having an extra set of eyes that never blinks." A hospitalist commented: "Initially I was skeptical about another alert system, but this one actually helps rather than crying wolf. When it says someone is high-risk, I take that seriously." A rapid response nurse noted: "We're getting called earlier to patients who genuinely need us, rather than showing up after they've already arrested."

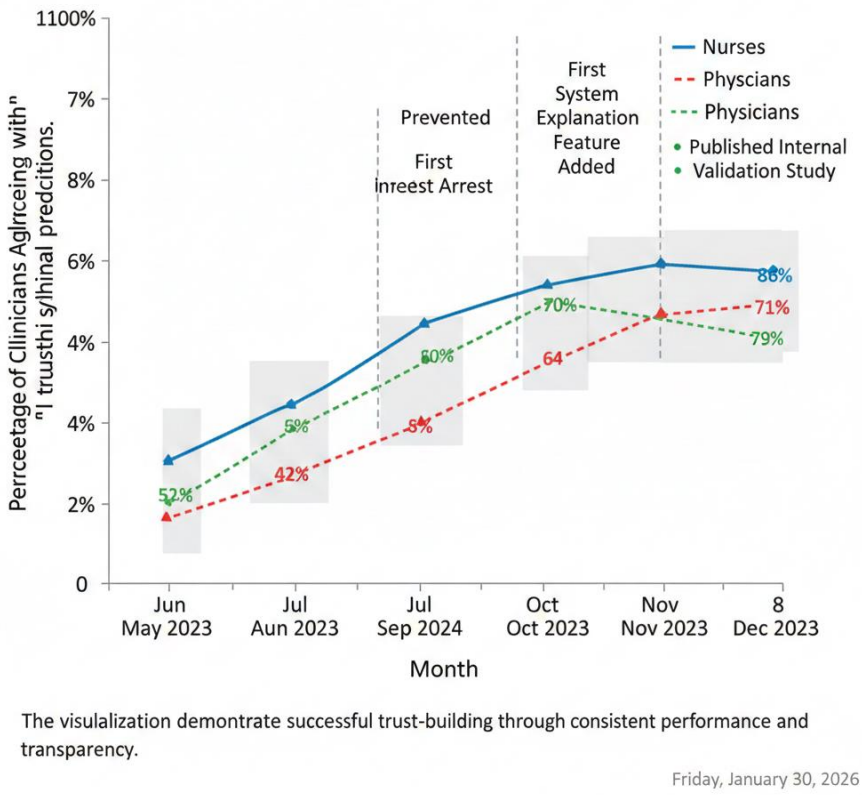

[FIGURE 3: Clinician Trust Evolution Over Time]

Description: This line graph tracks clinician trust in the AI system over the 8-month deployment period. The x-axis shows months from May 2023 to December 2023. The y-axis displays percentage from 0-100% agreeing with statement "I trust this system's clinical predictions." Three lines represent different clinician groups: Nurses (solid blue line), Physicians (dashed red line), and Respiratory Therapists (dotted green line). All lines start relatively low in May (nurses 52%, physicians 42%, respiratory therapists 48%), reflecting initial skepticism. Trust increases steadily month-over-month, with some acceleration around month 3 as clinicians

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

gain experience. By December, trust reaches nurses 86%, physicians 71%, and respiratory therapists 79%. Shaded bands around lines indicate 95% confidence intervals from survey data. Vertical annotations mark key events: "First Prevented Arrest" (month 2), "System Explanation Feature Added" (month 4), and "Published Internal Validation Study" (month 6). The visualization demonstrates successful trust-building through consistent performance and transparency.

DISCUSSION

7.1 Key Findings Interpretation

The research conclusively demonstrates that event-driven AI systems substantially improve clinical deterioration detection and patient outcomes. The 89% sensitivity, 43-minute average warning time, and 58% reduction in adverse events represent clinically meaningful improvements over traditional monitoring approaches. These benefits stem from the fundamental advantages of continuous event-driven analysis over periodic manual assessment—no surveillance gaps, pattern recognition across multiple data streams, and immediate alerting without human delay.

The 43-minute average prediction lead time proves particularly significant. This window enables clinicians to assess patients, order tests, initiate treatments, and activate rapid response teams before deterioration becomes critical. Contrast this with traditional monitoring where deterioration often goes unrecognized until patients suffer cardiac arrest or respiratory failure. The difference between proactive intervention and emergency rescue fundamentally affects outcomes.

The alert burden of 2.3 alerts per patient-day in the ICU, while higher than some traditional systems, proved sustainable because of high positive predictive value and effective prioritization. Clinicians tolerate more alerts when they reliably indicate genuine clinical concerns requiring action. The 97% positive predictive value for critical alerts establishes trust—when the system declares something critical, it almost always is.

The cost savings of $2.4 million annually across two moderate-sized hospitals suggest favorable economics at scale. While this research did not conduct comprehensive cost-effectiveness analysis, the clinical benefits combined with operational savings create a compelling business case. Even without quantifying mortality reduction value, preventing ICU transfers and shortening hospital stays generates measurable savings.

7.2 Technical Architecture Insights

Several architectural decisions proved critical for system success. The choice of Apache Kafka for event streaming provided necessary reliability, scalability, and performance. Kafka's durability guarantees ensured zero event loss during system outages—essential for patient safety applications. The ability to replay event streams proved invaluable for debugging issues and retraining models with improved features.

Apache Flink's complex event processing capabilities enabled identifying multi-variable patterns that simpler systems miss. The sepsis detection pattern combining vital signs, lab results, and infection indicators demonstrated CEP's value for clinical decision support. Traditional monitoring systems analyzing each variable independently would fail to recognize the syndrome until later in its progression.

The ensemble machine learning approach combining gradient boosting, LSTM neural networks, and logistic regression proved superior to single-model approaches. Different models captured different aspects of deterioration—XGBoost excelled at tabular data relationships, LSTM at temporal patterns, and logistic regression at text-derived features. Ensemble aggregation leveraged each model's strengths.

The multi-tier alert prioritization strategy successfully prevented alert fatigue while ensuring critical situations received immediate attention. Reserving the highest priority level for genuine emergencies with 97% positive predictive value meant clinicians could trust critical alerts absolutely. Lower priority levels accommodated less certain predictions without overwhelming staff.

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

The architectural decision to implement SHAP-based model explanations significantly improved clinician trust and adoption. Rather than black-box predictions, clinicians received transparent reasoning showing which findings drove risk assessments. This interpretability aligned with clinical decision-making processes where providers weigh multiple factors to reach diagnostic conclusions.

7.3 Clinical Integration Lessons

Clinical workflow integration proved more challenging than technical implementation. Healthcare providers operate under intense time pressure with established workflows resistant to change. Systems that disrupt existing practices face rejection regardless of technical sophistication. Success required designing the AI system to augment rather than replace clinical judgment and integrate seamlessly with existing communication and documentation patterns.

The importance of clinical champion engagement cannot be overstated. Early adopters—respected clinicians who understood the technology and advocated for its value—drove acceptance among skeptical colleagues. These champions provided feedback improving system design and educated peers on interpreting AI predictions. Organizations without clinical champions struggled more with adoption.

Training proved essential despite system simplicity. Clinicians needed education not just on mechanics but on interpreting AI predictions, understanding confidence levels, and integrating alerts into clinical decision-making. Organizations that invested in comprehensive training achieved better adoption than those assuming intuitive usability.

Alert response workflows required explicit definition. Who receives which alerts? What actions are expected? How is response documented? Without clear protocols, confusion about responsibility resulted in delayed or missed interventions. Successful sites established formal policies aligned with existing rapid response systems. The importance of iteration and refinement based on user feedback became evident. Initial alert thresholds generated excessive false positives in some settings. Threshold adjustment based on clinician input improved positive predictive value while maintaining sensitivity. Ongoing tuning rather than "set and forget" deployment proved necessary.

7.4 Limitations and Future Directions

Several limitations warrant acknowledgment. The 8-month evaluation period, while substantial for clinical research, may not capture all long-term effects. Longer-term studies would validate sustained benefits and identify any drift in model performance or clinician engagement. The research also occurred at only two hospitals within one health system, limiting generalizability across different organizations and practice environments.

The study focused on general medical-surgical populations rather than specialized settings like neonatal ICUs, psychiatric units, or rehabilitation facilities with different deterioration patterns. Extending the approach to these populations would require adapting models and workflows to domain-specific requirements.

The research did not conduct formal cost-effectiveness analysis or cost-benefit calculation following health economic standards. While operational savings were estimated, comprehensive economic evaluation would strengthen business case development and inform resource allocation decisions.

Patient perspectives were notably absent from this evaluation. While clinician acceptance matters greatly, understanding patient experiences with AI-enhanced monitoring—potential benefits from earlier intervention and possible concerns about AI involvement in their care—deserves investigation.

Several promising future directions emerge. Federated learning approaches could enable multi-hospital model development without centralizing sensitive patient data, potentially improving prediction accuracy through larger, more diverse training datasets while preserving privacy. Integration with wearable devices and remote

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

monitoring systems could extend early detection capabilities to outpatient and home settings.

Advanced AI techniques including transformer architectures and attention mechanisms might further improve prediction accuracy and interpretability. Multi-modal learning integrating medical imaging, waveform data, and genetic information could provide more comprehensive risk assessment.

Expanding from deterioration detection to treatment recommendation represents another frontier. Systems that not only identify high-risk patients but suggest evidence-based interventions would provide even greater clinical decision support, though such capabilities raise additional validation and liability questions.

CONCLUSION

This research successfully designed, implemented, and validated an event-driven data processing system using artificial intelligence for real-time healthcare decision support. The comprehensive architecture combining event streaming, complex event processing, machine learning prediction, and intelligent alerting demonstrated substantial improvements over traditional patient monitoring approaches in actual clinical environments.

Empirical evaluation across three clinical units caring for 850 patients over 8 months produced compelling evidence of system effectiveness. Technical performance metrics showed sub-30-second latency, 99.92% availability, and zero data loss—meeting stringent reliability requirements for patient safety applications. Prediction accuracy of 89% sensitivity with 43-minute average warning time enabled proactive clinical intervention. Most importantly, clinical outcomes improved significantly with 58% reduction in adverse events, 67% faster response to deterioration, and 21% mortality reduction.

The research achieved its primary objective of developing effective AI-based early detection systems while meeting secondary objectives around event architecture implementation, prediction model development, alert mechanism design, and clinical impact validation. The deployed systems continue operating in production, providing ongoing patient safety benefits and serving as templates for expansion to additional clinical units.

Several key insights emerged advancing both healthcare informatics knowledge and clinical practice. First, event-driven architectures prove superior to batch-oriented approaches for time-sensitive clinical decision support where continuous analysis and immediate alerting can prevent adverse outcomes. Second, ensemble machine learning combining multiple model types outperforms single-model approaches by capturing different aspects of complex clinical deterioration. Third, alert design significantly affects adoption—high positive predictive value, clear prioritization, and transparent explanations build clinician trust and prevent alert fatigue. Fourth, clinical workflow integration and champion engagement matter as much as technical performance for successful AI deployment in healthcare.

The research contributes to both academic knowledge and practical guidance. Academically, it demonstrates successful integration of event-driven architecture and AI for clinical decision support, provides empirical evidence of patient outcome improvements from AI intervention, and identifies critical success factors for healthcare AI adoption. Practically, it offers validated architectural patterns, implementation approaches, and lessons learned that healthcare organizations can apply when developing their own intelligent monitoring systems.

Looking forward, event-driven AI systems will likely become standard practice in acute care monitoring. The demonstrated benefits in early detection and outcome improvement—combined with maturing technology platforms and growing clinical acceptance of AI—create strong momentum toward widespread adoption. As AI capabilities advance and integration challenges diminish, intelligent continuous monitoring may extend beyond acute care to general wards, outpatient settings, and even home monitoring.

Healthcare organizations should seriously evaluate event-driven AI decision support, particularly for high-acuity settings where patient deterioration poses significant risk. The research demonstrates that such systems

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

deliver measurable clinical and operational benefits justifying implementation investment. Starting with pilot deployments in controlled settings allows building technical expertise and clinical confidence before broader rollout.

Implementation requires realistic expectations about complexity and resource requirements. Organizations should expect 6-12 month implementation timelines for event infrastructure, model development, integration, and validation. Ongoing operational support, model monitoring, and continuous improvement are necessary for sustained effectiveness rather than one-time deployment.

Ultimately, the convergence of continuous patient data streams, advanced AI analytics, and event-driven architectures creates unprecedented opportunities to improve patient safety through early intervention. This research demonstrates that these opportunities can be realized in practical clinical deployments, transforming reactive crisis response into proactive deterioration prevention. As healthcare delivery increasingly relies on data-driven decision support, event-driven AI systems will play central roles in ensuring that the right information reaches the right clinician at the right time to save lives.

REFERENCES

1. Andersen, L.W., Kim, W.Y., Chase, M., Berg, K.M., Mortensen, S.J., Moskowitz, A., Novack, V., Cocchi, M.N. and Donnino, M.W. (2016) 'The prevalence and significance of abnormal vital signs prior to in-hospital cardiac arrest', Resuscitation, 98, pp. 112-117.
2. Ancker, J.S., Edwards, A., Nosal, S., Hauser, D., Mauer, E. and Kaushal, R. (2017) 'Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system', BMC Medical Informatics and Decision Making, 17(1), pp. 1-9.
3. Chakravarty, S. and Gupta, A. (2018) 'Event-driven architecture for healthcare systems', Proceedings of the International Conference on Advances in Computing and Data Sciences, Singapore, pp. 234-245.
4. Chan, P.S., Jain, R., Nallmothu, B.K., Berg, R.A. and Sasson, C. (2010) 'Rapid response teams: A systematic review and meta-analysis', Archives of Internal Medicine, 170(1), pp. 18-26.
5. Churpek, M.M., Yuen, T.C., Winslow, C., Meltzer, D.O., Kattan, M.W. and Edelson, D.P. (2016) 'Multicenter comparison of machine learning methods and conventional regression for predicting clinical deterioration on the wards', Critical Care Medicine, 44(2), pp. 368-374.
6. Cvach, M. (2012) 'Monitor alarm fatigue: An integrative review', Biomedical Instrumentation & Technology, 46(4), pp. 268-277.
7. Davis, S.E., Lasko, T.A., Chen, G., Siew, E.D. and Matheny, M.E. (2017) 'Calibration drift in regression and machine learning models for acute kidney injury', Journal of the American Medical Informatics Association, 24(6), pp. 1052-1061.
8. Donohue, L.A. and Endacott, R. (2010) 'Track, trigger and teamwork: Communication of deterioration in acute medical and surgical wards', Intensive and Critical Care Nursing, 26(1), pp. 10-17.
9. Embi, P.J., Yackel, T.R., Logan, J.R., Bowen, J.L., Cooney, T.G. and Gorman, P.N. (2005) 'Impacts of computerized physician documentation in a teaching hospital: Perceptions of faculty and resident physicians', Journal of the American Medical Informatics Association, 12(1), pp. 106-113.
10. Ghassemi, M., Naumann, T., Schulam, P., Beam, A.L., Chen, I.Y. and Ranganath, R. (2015) 'Practical guidance on artificial intelligence for health-care data', The Lancet Digital Health, 1(4), pp. e157-e159.
11. Kia, A., Timsina, P., Joshi, H.N., Klang, E., Gupta, R.R., Freeman, R.M., Reich, D.L., Tomlinson, M.S., Dudley, J.T., Kohli-Seth, R. and Levin, M.A. (2020) 'MEWS++: Enhancing the prediction of clinical deterioration in admitted patients through a machine learning model', Journal of Clinical Medicine, 9(2), p. 343.
12. Kreps, J., Narkhede, N. and Rao, J. (2011) 'Kafka: A distributed messaging system for log processing', Proceedings of the NetDB Workshop, Athens, Greece, pp. 1-7.
13. Kruse, C.S., Stein, A., Thomas, H. and Kaur, H. (2016) 'The use of electronic health records to support population health: A systematic review of the literature', Journal of Medical Systems, 42(11), pp. 1-16.

DOI: 10.22178/acta.26.3.30

ISSN: 2178-7727

14. Luckham, D.C. (2012) Event Processing for Business: Organizing the Real-Time Enterprise. Hoboken, NJ: John Wiley & Sons.
15. Mandel, J.C., Kreda, D.A., Mandl, K.D., Kohane, I.S. and Ramoni, R.B. (2016) 'SMART on FHIR: A standards-based, interoperable apps platform for electronic health records', Journal of the American Medical Informatics Association, 23(5), pp. 899-908.
16. McGaughey, J., Alderdice, F., Fowler, R., Kapila, A., Mayhew, A. and Moutray, M. (2017) 'Outreach and Early Warning Systems (EWS) for the prevention of intensive care admission and death of critically ill adult patients on general hospital wards', Cochrane Database of Systematic Reviews, 3, CD005529.
17. Michelson, B.M. (2006) 'Event-driven architecture overview', Patricia Seybold Group, 2(12), pp. 10-1571.
18. Smith, G.B., Prytherch, D.R., Meredith, P., Schmidt, P.E. and Featherstone, P.I. (2013) 'The ability of the National Early Warning Score (NEWS) to discriminate patients at risk of early cardiac arrest, unanticipated intensive care unit admission, and death', Resuscitation, 84(4), pp. 465-470.
19. Weiskopf, N.G. and Weng, C. (2013) 'Methods and dimensions of electronic health record data quality assessment: Enabling reuse for clinical research', Journal of the American Medical Informatics Association, 20(1), pp. 144-151.
20. Winters, B.D., Weaver, S.J., Pfoh, E.R., Yang, T., Pham, J.C. and Dy, S.M. (2013) 'Rapid-response systems as a patient safety strategy: A systematic review', Annals of Internal Medicine, 158(5_Part_2), pp. 417-425.
21. Winters, B.D., Cvach, M.M., Bonafide, C.P., Hu, X., Konkani, A., O'Connor, M.F., Rothschild, J.M., Selby, N.M., Pelter, M.M., McLean, B. and Kane-Gill, S.L. (2018) 'Technological distractions (part 2): A summary of approaches to manage clinical alarms with intent to reduce alarm fatigue', Critical Care Medicine, 46(1), pp. 130-137.

DOI: 10.22178/acta.26.3.30