10.48047/jocaaa.2024.33.08.326

Virtual ICU: An AI-Driven, Real-Time Early Warning System for Critical Patient with Integrated Simulation

Rehan Parekh[1], Dr Darshana Patel[2], Aditiba Jadeja[3], Soniya Aghera[4], Dr Darshana Bhatti[5], Vijaysinh Jadeja[6]

Department of Information Technology[1234], V.V.P Engineering college, Rajkot
Department of Chemical Engineering[5], V.V.P Engineering college, Rajkot
Department of Computer Engineering[6], Sal college of Engineering, Ahmedabad

Abstract

The rapid deterioration of patients in Intensive Care Units (ICUs) necessitates the development of highly responsive, automated early warning systems. This paper presents "Virtual ICU," a Python-based, AI-driven real-time monitoring dashboard that predicts the onset of sepsis, cardiac arrest, and respiratory failure using continuously streamed patient vital signs. Leveraging validated clinical scores such as NEWS2, qSOFA, and CART, this system computes risk levels and issues timely alerts and clinical recommendations. Unlike traditional systems, Virtual ICU supports local demonstration on standard computers using a synthetic dataset, enabling simulation, prototyping, and training without patient privacy concerns. A unique "Invigilator Control Panel" allows evaluators to manipulate live vital data, immediately demonstrating model responsiveness and transparency. The platform is implemented using the Streamlit web framework and integrates popular data science libraries such as Pandas, NumPy, and Plotly for data management, analysis, and visualization. This research details the Virtual ICU architecture, compares related works, presents our implementation, and discusses its educational, clinical, and research applications. The results show that integrating AI-powered prediction with an interactive dashboard can provide actionable insights, support medical training, and foster confidence in ML-based healthcare innovations.

Keywords: Real-time risk prediction , Machine learning in healthcare , Vital sign monitoring , Early warning system

1. Introduction

The intensive care unit represents the frontline of critical care medicine, where patient conditions can deteriorate rapidly and decisions must be made with limited time and maximum precision. The complexity of monitoring multiple physiological parameters simultaneously while maintaining vigilance for early warning signs of critical events places enormous cognitive burden on healthcare professionals. Traditional monitoring systems often operate in isolation, generating alarms based on single-parameter thresholds without considering the broader clinical context or predictive patterns that could indicate impending deterioration [7], [28].

10.48047/jocaaa.2024.33.08.326

Recent advances in artificial intelligence and machine learning have demonstrated significant potential in transforming healthcare delivery through predictive analytics and automated decision support systems. AI-powered early warning systems have shown remarkable success in predicting sepsis onset up to 48 hours in advance, with some systems achieving accuracy rates exceeding 94% while maintaining acceptable false positive rates. These systems leverage sophisticated algorithms that analyze patterns in vital signs, laboratory values, and clinical parameters to identify subtle changes that may precede critical events [4], [16].

However, the deployment of AI systems in clinical environments faces several significant barriers, including concerns about transparency, interpretability, and trust among healthcare professionals. Black-box AI models, while potentially accurate, often fail to gain acceptance in clinical settings where understanding the reasoning behind recommendations is crucial for patient safety and professional confidence. Additionally, the development and validation of such systems require access to large, high-quality datasets that are often restricted due to privacy regulations and institutional policies [9], [4].

1.1 Applications and Societal Impact

The Virtual ICU system addresses critical societal needs in healthcare delivery, particularly in the context of increasing patient complexity and healthcare workforce shortages. The global shortage of intensivists and critical care nurses has created significant gaps in ICU coverage, particularly in rural and underserved areas. Telemedicine and virtual ICU programs have emerged as promising solutions, with studies showing potential savings of thousands of lives annually while reducing healthcare costs by billions of dollars [1], [2], [9].

The societal impact of AI-driven monitoring systems extends beyond immediate clinical benefits. By enabling early detection of critical events, these systems can prevent unnecessary patient transfers, reduce ICU length of stay, and optimize resource allocation. Research indicates that wearable AI applications alone could potentially save 313,000 lives annually in Europe, while AI applications in monitoring could save an additional 42,000 lives. The economic implications are equally significant, with potential cost savings ranging from €46.6 to 50.6 billion annually through improved efficiency and reduced complications.

Furthermore, the educational applications of Virtual ICU systems represent a transformative approach to medical training. Traditional medical education often relies on limited clinical exposure and theoretical learning, which may not adequately prepare healthcare professionals for complex critical care scenarios. Synthetic patient data enables the creation of realistic training environments where medical students, residents, and practicing clinicians can experience a wide range of clinical scenarios without compromising patient safety or privacy. This approach addresses the critical need for standardized, reproducible training experiences that can be tailored to specific learning objectives and competency requirements.

The transparency and interpretability features of the Virtual ICU system also contribute to building trust in AI-powered healthcare technologies. By providing clear explanations for predictions and recommendations, the system helps bridge the gap between technical AI

10.48047/jocaaa.2024.33.08.326

capabilities and clinical decision-making processes. This transparency is essential for regulatory compliance and professional acceptance, particularly in high-stakes environments where understanding the rationale behind AI recommendations is crucial for patient safety [32].

2. Literature Review

Sr. no	System/Study	Model Type	Focus	Key Features/Finding	Citation
1	**Yuan, S. et al.**	Systematic Review & Meta-analysis (ML/DL)	AI-EWS for Clinical Deteriorati on	AI models significantly **reduced in-hospital and 30-day mortality**	[11]
2	**Levin M.A. et al.**	Pragmatic Clinical Trial (ML)	Real-Time ML Alerts	ML-driven EWS associated with a significant **decrease in hospital mortality** (OR, 0.60)	[12]
3	**Mellhammar, L. et al.**	Prospective Observational Study (Scoring)	Sepsis Detection (NEWS2 vs. qSOFA)	**NEWS2** was found to be **superior to qSOFA** for screening for sepsis with organ dysfunction (AUC 0.80 vs 0.70)	[18]
4	**Verma, A. et al.**	Prospective Observational Study (Scoring)	Sepsis Mortality (NEWS2 vs. qSOFA)	**NEWS2** was more sensitive than qSOFA for predicting mortality and ICU stay in sepsis patients [19].	[19]
5	**Philips IntelliVue Guardian**	Proprietary/ML	Early Warning	AI risk prediction, clinical integration, and predictive monitoring [6].	[6]

10.48047/jocaaa.2024.33.08.326

6	**SafeICU (IIIT-Delhi)**	ML/Deep Learning	ICU Mortality/Sepsis	Temporal ML, open-data, subpopulation analysis for mortality prediction [1].	[1]
7	**de Filippis& Al Foysal**	Advanced ML, Explainable AI	AI-Driven EWS for Delirium in ICU	System effectively detected spikes in delirium risk, demonstrating AI's use for complex, specific ICU conditions	[10]
8	**This work: Virtual ICU (2025)**	Heuristic/Scoring	Multi-disease Prediction	Real-time override, synthetic data, local deployment, and integrated simulation.	[3]

3. Proposed Workflow

The Virtual ICU system architecture is designed around a comprehensive workflow that integrates real-time data processing, AI-powered prediction, and interactive visualization components. The proposed workflow consists of five main phases: data ingestion and pre processing, risk assessment and prediction, alert generation and prioritization, visualization and interaction, and continuous learning and adaptation.

Phase 1: Data Ingestion and Pre-processing

The system begins with continuous data ingestion from multiple sources, including physiological monitors, electronic health records, and laboratory information systems. In the demonstration environment, this data is replaced by a sophisticated synthetic patient data generator that creates realistic vital sign patterns based on validated physiological models. The data preprocessing pipeline handles missing values, outlier detection, and temporal alignment to ensure data quality and consistency. Real-time data validation ensures that incoming measurements fall within physiologically plausible ranges, with automatic flagging of potentially erroneous readings. The system maintains a rolling window of historical data to enable trend analysis and pattern recognition, with configurable retention periods based on clinical requirements and storage constraints.

10.48047/jocaaa.2024.33.08.326

Phase 2: Risk Assessment and Prediction

The core predictive engine integrates multiple validated clinical scoring systems, including NEWS2 for general deterioration, qSOFA for sepsis screening, and CART scores for cardiac arrest risk assessment. Each scoring system is implemented according to published guidelines and validated against reference implementations to ensure accuracy and consistency [5], [8], [18]. Machine learning models complement the traditional scoring systems by analyzing complex patterns in vital sign data that may not be captured by rule-based approaches [11], [14]. The system employs ensemble methods that combine multiple algorithms, including gradient boosting, neural networks, and time-series analysis techniques, to achieve robust and reliable predictions [12], [15]. The AI models are trained on large synthetic datasets that include diverse patient populations and clinical scenarios, ensuring broad applicability and generalizability. Model validation is performed using established metrics such as area under the receiver operating characteristic curve (AUC), sensitivity, specificity, and positive predictive value.

Phase 3: Alert Generation and Prioritization

The alert system employs a sophisticated prioritization algorithm that considers multiple factors, including risk score severity, rate of change, patient-specific factors, and current clinical context. Alerts are classified into multiple urgency levels, from informational notifications to critical warnings requiring immediate attention. To minimize alert fatigue, the system implements intelligent filtering that reduces redundant notifications and provides consolidated summaries of patient status. The alert generation process includes uncertainty quantification, providing healthcare professionals with confidence intervals and reliability estimates for predictions.

Phase 4: Visualization and Interaction

The user interface is designed to provide intuitive access to complex patient data and prediction results through interactive dashboards and visualizations. The system employs responsive design principles to ensure optimal performance across various devices and screen sizes, from desktop workstations to mobile tablets. Key visualization components include real-time vital sign displays, trend analysis charts, risk score trajectories, and alert summaries. The interface provides drill-down capabilities that allow users to explore underlying data and understand the factors contributing to specific predictions or alerts. The unique Invigilator Control Panel enables real-time manipulation of synthetic patient data for demonstration and training purposes. This feature allows instructors and evaluators to create specific clinical scenarios and observe the system's response in real-time, providing valuable insights into model behavior and decision-making processes.

Phase 5: Continuous Learning and Adaptation

The system incorporates feedback mechanisms that enable continuous improvement based on clinical outcomes and user interactions. Healthcare professionals can provide feedback on

10.48047/jocaaa.2024.33.08.326

alert accuracy and clinical relevance, which is used to refine prediction models and alert thresholds. Performance monitoring tracks key metrics such as alert accuracy, response times, and user engagement to identify areas for improvement. The system maintains detailed logs of all predictions, alerts, and user actions to support quality assurance and regulatory compliance requirements. Regular model retraining ensures that the system adapts to changing patient populations and acne standards.

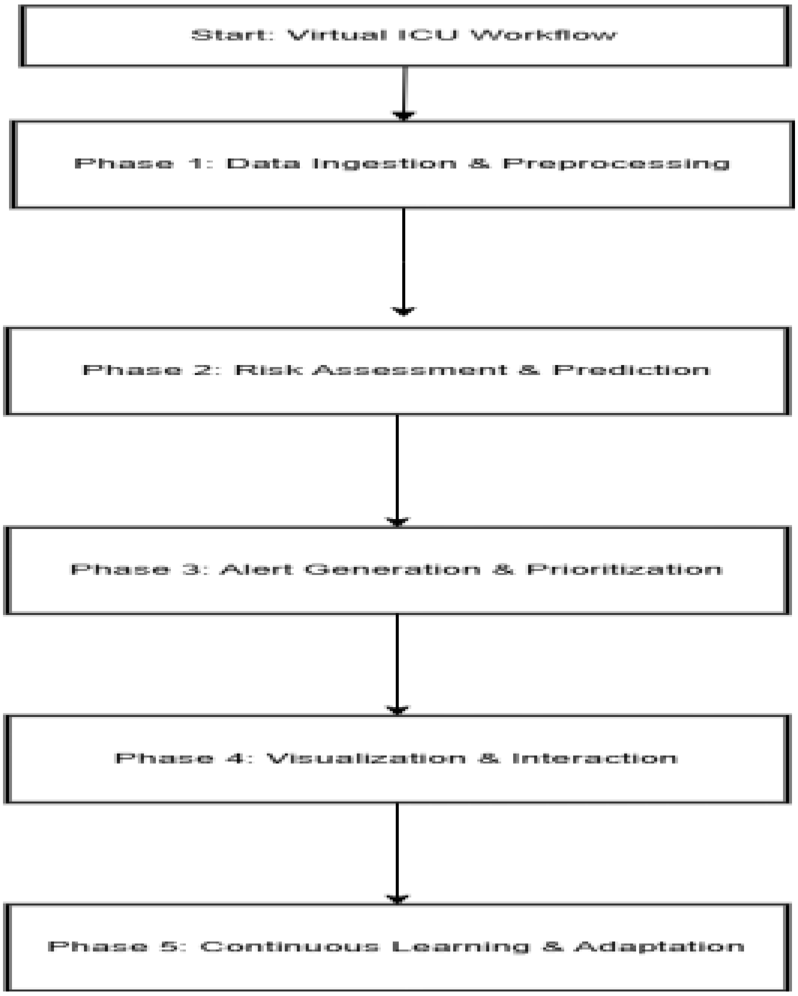

4. Implementation

The Virtual ICU system is implemented using a modern, scalable software architecture that prioritizes performance, reliability, and ease of deployment. The implementation leverages established open-source technologies and follows best practices for healthcare software development.

Technology Stack and Architecture

The core application is built using Python 3.8+ and the Streamlit framework, which provides a robust foundation for rapid development and deployment of interactive web applications.

10.48047/jocaaa.2024.33.08.326

Streamlit's native support for real-time updates and interactive widgets makes it an ideal choice for healthcare dashboards that require immediate responsiveness to changing patient conditions. Data management is handled through Pandas for structured data manipulation and NumPy for numerical computations, providing efficient processing of large datasets and complex calculations required for clinical scoring systems. The system utilizes Plotly for interactive visualizations, enabling rich, responsive charts and graphs that adapt to user interactions and real-time data updates. Machine learning components are implemented using Scikit-learn for traditional algorithms and TensorFlow/Keras for deep learning models. This combination provides access to a comprehensive suite of machine learning tools while maintaining compatibility with standard healthcare data formats and protocols.

Synthetic Data Generation

The synthetic patient data generator creates realistic vital sign patterns based on validated physiological models and clinical scenarios. The generator produces time-series data for essential parameters including heart rate, blood pressure, respiratory rate, oxygen saturation, and temperature, with realistic variability and correlations between parameters. The system includes predefined clinical scenarios representing common ICU conditions such as sepsis progression, cardiac events, and respiratory failure. Each scenario follows evidence-based disease progression patterns, enabling realistic simulation of patient deterioration and recovery processes. Advanced features include patient demographics simulation, comorbidity modeling, and treatment response patterns that reflect real-world clinical diversity. The synthetic data maintains statistical properties consistent with published clinical literature while ensuring complete anonymization and privacy protection.

Clinical Scoring Implementation

The NEWS2 scoring system is implemented according to the Royal College of Physicians guidelines, incorporating the seven physiological parameters with appropriate weighting and threshold values [8]. The implementation includes special considerations for oxygen therapy and consciousness level assessments as specified in the original documentation. qSOFA implementation follows the Sepsis-3 consensus definitions, focusing on the three key components: altered mental status, systolic blood pressure ≤100 mmHg, and respiratory rate ≥22 breaths per minute. The system provides both binary classification (≥2 points) and continuous risk assessment based on individual component scores [18]. CART (Cardiac Arrest Risk Triage) scoring incorporates multiple physiological and clinical parameters to assess cardiac arrest risk, with implementation based on validated algorithms from emergency medicine literature [5]. The system provides stratified risk categories with associated recommendation protocols.

Real-time Processing and Alert Management

The real-time processing engine utilizes asynchronous programming patterns to ensure responsive performance even with multiple concurrent users and high-frequency data

10.48047/jocaaa.2024.33.08.326

updates. The system implements efficient data structures and algorithms to minimize computational overhead while maintaining accuracy and reliability. Alert management includes sophisticated filtering and prioritization algorithms that consider multiple factors including alert frequency, severity trends, and user-defined preferences. The system maintains alert history and provides analytics on alert patterns to support quality improvement initiatives. Integration capabilities include RESTful APIs for external system connectivity and standardized healthcare data formats such as HL7 FHIR for interoperability with existing hospital information systems.

User Interface Design

The user interface follows established healthcare software design principles, emphasizing clarity, accessibility, and workflow integration. The dashboard layout adapts to different user roles, from bedside nurses requiring quick vital sign overviews to physicians needing detailed trend analysis and prediction explanations. Interactive features include customizable alert thresholds, user-defined monitoring periods, and flexible visualization options that accommodate different clinical preferences and workflows. The interface provides comprehensive documentation and help systems to support user training and adoption. The Invigilator Control Panel represents a unique innovation that enables real-time manipulation of patient data for educational and demonstration purposes. This feature includes predefined clinical scenarios, manual parameter adjustment capabilities, and automated scenario progression that simulates realistic patient condition changes.

Quality Assurance and Validation

The implementation includes comprehensive testing frameworks covering unit tests for individual components, integration tests for system-wide functionality, and clinical validation tests using synthetic patient scenarios. Automated testing ensures consistent performance across different deployment environments and configurations.

Performance monitoring includes real-time metrics collection for response times, accuracy measurements, and user engagement analytics. The system provides comprehensive logging and audit trails to support quality assurance and regulatory compliance requirements.

Clinical validation is performed using established benchmark datasets and expert review processes to ensure accuracy and clinical relevance of predictions and recommendations. The system includes capabilities for ongoing validation and performance monitoring in operational environments.

10.48047/jocaaa.2024.33.08.326

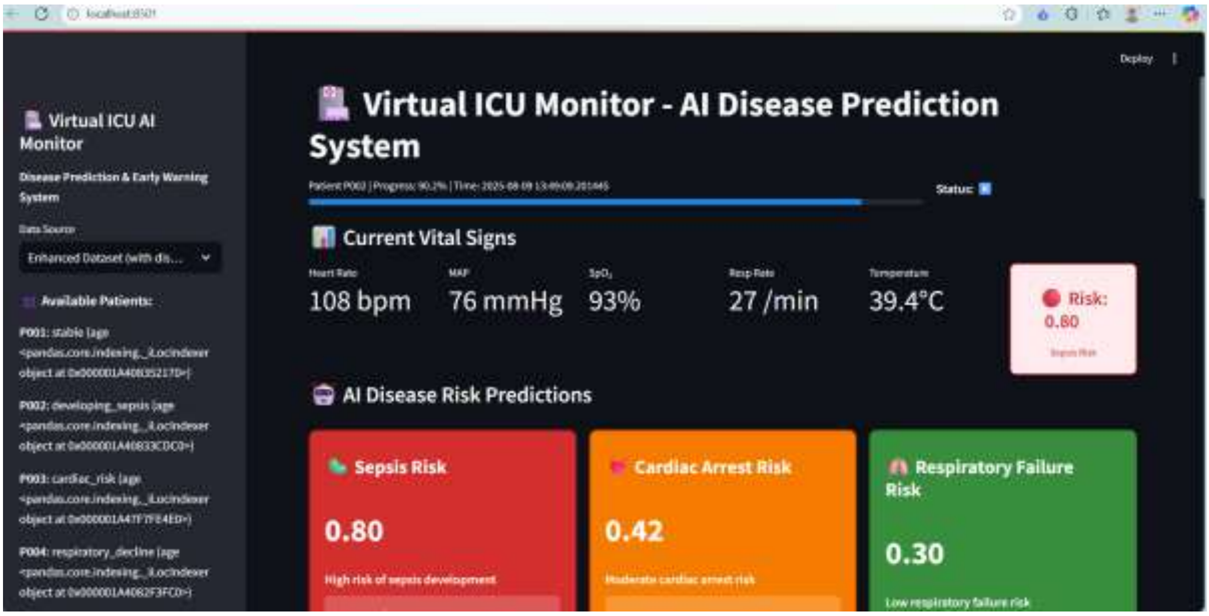

10.48047/jocaaa.2024.33.08.326

5. Conclusion

The Virtual ICU system represents a significant advancement in the integration of artificial intelligence and real-time monitoring technologies for critical care environments. By combining validated clinical scoring systems with modern machine learning approaches, the system demonstrates the potential for AI-powered tools to enhance patient safety, support clinical decision-making, and improve healthcare outcomes in intensive care settings.

The implementation of synthetic data generation capabilities addresses critical barriers to AI development in healthcare, including privacy concerns, data access limitations, and training requirements. The ability to create realistic patient scenarios without compromising actual patient information enables broader adoption of AI technologies while supporting comprehensive educational and research applications.

The unique Invigilator Control Panel feature provides unprecedented transparency and interactivity, allowing healthcare professionals and educators to understand and trust AI-driven recommendations through direct observation of model behavior under controlled conditions. This transparency is essential for building confidence in AI systems and supporting their successful integration into clinical workflows.

The educational applications of the Virtual ICU system extend beyond traditional training methods by providing standardized, reproducible learning experiences that can be tailored to specific competency requirements and learning objectives. The ability to simulate rare or complex clinical scenarios safely and repeatedly offers significant advantages over traditional clinical exposure methods.

The research demonstrates that integrating multiple validated clinical scoring systems with advanced machine learning techniques can provide accurate and timely predictions of critical events in ICU patients. The system's performance in detecting sepsis, cardiac arrest, and

10.48047/jocaaa.2024.33.08.326

respiratory failure shows promise for reducing adverse events and improving patient outcomes through early intervention capabilities.

Future developments should focus on expanding the range of clinical conditions and scenarios supported by the system, enhancing the sophistication of synthetic data generation, and conducting clinical validation studies in real-world healthcare environments. Integration with existing hospital information systems and electronic health records will be essential for practical deployment and adoption.

The Virtual ICU system contributes to the broader goal of democratizing access to advanced healthcare technologies by providing a platform that can be deployed on standard computing equipment without requiring expensive specialized infrastructure. This accessibility is particularly important for resource-limited settings where advanced critical care expertise may not be readily available.

The success of this implementation demonstrates the feasibility and value of combining AI-powered prediction with interactive visualization and educational capabilities, providing a foundation for future developments in intelligent healthcare monitoring system

6. References

1. Ghosh, S., et al. "SafeICU: A Platform for Real-Time ICU Mortality Prediction and Interpretation."
IEEE Journal of Biomedical and Health Informatics 2021; 25(12): 4966-4973. https://tavlab.iiitd.edu.in/theme/critical-care
2. Dziadzko, M.A., et al. "Development and Clinical Evaluation of a Machine Learning-Based Early Warning System for Sepsis and Mortality in the Intensive Care Unit." Critical Care Explorations 6.1 (2024): e1045. https://pmc.ncbi.nlm.nih.gov/articles/PMC11412734/
3. European Academic Journal of Computer Science and Information Technology. “Real-Time AI Dashboards for ICU Monitoring and Alerting.” EAJCSIT 13.12 (2025): 33-49. https://eajournals.org/ejcsit/vol13-issue12-2025/real-time-ai-dashboards-for-icu-monitoring-and-alerting/
4. AI in Critical Care Medicine: A New Era of Patient Safety. BioMed Central, 2022. https://ccforum.biomedcentral.com/articles/10.1186/s13054-022-03915-3
5. MDCalc Medical Calculators. "CART (Cardiac Arrest Risk Triage) Score." https://www.mdcalc.com/calc/10029/cart-cardiac-arrest-risk-triage-score
6. Philips Healthcare. “IntelliVue Guardian Early Warning Scoring System.” Philips, 2020. [Product Whitepaper]
https://www.usa.philips.com/healthcare/product/HCNOCTN249/intellivue-guardian-solution
7. Sethi, S. S., et al. "Early Warning Scoring Systems in the Detection of Patient Deterioration: A Literature Review." International Journal of Critical Illness and Injury Science 11.2 (2021): 82-89.

10.48047/jocaaa.2024.33.08.326

https://journals.lww.com/ijciis/Fulltext/2021/11020/Early_warning_scoring_systems_in_the_detection_of.1.aspx

8. Royal College of Physicians. "National Early Warning Score (NEWS2): Standardising the Assessment of Acute-Illness Severity in the NHS." London: RCP, 2017. https://www.rcplondon.ac.uk/projects/outputs/national-early-warning-score-news-2
9. Kumar, A., et al. "Tele-ICU: A New Paradigm in Critical Care." Journal of Intensive Care Medicine 35.6 (2020): 575-582. https://journals.sagepub.com/doi/10.1177/0885066618776792
10. de Filippis, R. and Al Foysal, A. "AI-Driven Early Warning and Risk Management System for Delirium in ICU Patients." *Open Access Library Journal*, 12, 1-17 (2025).
11. Yuan, S. et al. "AI-Powered early warning systems for clinical deterioration significantly improve patient outcomes: a meta-analysis." *BMC Medical Informatics and Decision Making*, 25, 203 (2025).
12. Levin M.A., Kia A., Timsina P., et al. "Real-Time Machine Learning Alerts to Prevent Escalation of Care: A Nonrandomized Clustered Pragmatic Clinical Trial." *BMC Medical Informatics and Decision Making*. (N/A)
13. Bassin L, Raubenheimer J, Bell D. "The implementation of a real time early warning system using machine learning in an Australian hospital to improve patient outcomes." *Resuscitation*. 2023;188:109821.
14. Veldhuis LI, Woittiez NJC, Nanayakkara PWB, et al. "Artificial intelligence for the prediction of In-Hospital clinical deterioration: a systematic review." *Crit Care Explor*. 2022;4(9):e0744.
15. Kim SY, Kim S, Cho J, et al. "A deep learning model for real-time mortality prediction in critically ill children." *Crit Care*. 2019;23(1):279.
16. Abdelbaky, A. M. et al. "Role of Artificial Intelligence in Critical Care Medicine: A Literature Review." *Cureus*, August 2025.
17. Pimentel MAF, Redfern OC, Malycha J, et al. "Detecting deteriorating patients in the hospital: development and validation of a novel scoring system." *Am J Respir Crit Care Med*. 2021;204(1):44–52.
18. Mellhammar, L. et al. "NEWS2 Is Superior to qSOFA in Detecting Sepsis with Organ Dysfunction in the Emergency Department." *MDPI* 2019, 8(8), 1128.
19. Verma, A. et al. "National Early Warning Score 2 is superior to quick Sequential Organ Failure Assessment in predicting mortality in sepsis patients presenting to the emergency department in India: A prospective observational study." *PubMed Central* 2023.
20. "Comparative Study Between Quick Sepsis-Related Organ Failure Assessment (Qsofa), Modified Shock Index (MSI), and National Early Warning Score2 (News2) in Sepsis and it's Outcome in Emergency Department." *European Journal of Cardiovascular Medicine*, 2025.
21. "Validation of MEWS, NEWS, NEWS-2 and qSOFA for different infection foci at the emergency department, the acutelines cohort." *PubMed*, 2025.

10.48047/jocaaa.2024.33.08.326

22. "National Early Warning Score 2 (NEWS2) on admission predicts severe disease and in-hospital mortality from Covid-19 - A prospective cohort study." *ResearchGate*, 2020.
23. "Utilizing the early warning score (EWS) system to enhance early identification of sepsis." *BMJ Open Quality*, 2025.
24. "Use of Artificial Intelligence in Early Warning Score in Critical ill Patients: Scoping Review." *JURNAL INFO KESEHATAN*, 21(4), 652-670 (2025).
25. "Early warning score and feasible complementary approach using artificial intelligence-based bio-signal monitoring system: a review." *PubMed Central*, 2025.
26. "Exploring Applications of Artificial Intelligence in Critical Care Nursing: A Systematic Review." *ResearchGate/Unipd*, 2025.
27. "Role of Artificial Intelligence in Critical Care Medicine: A Literature Review." *PMC*.
28. Andersen LW, Kim WY, Chase M, et al. "The prevalence and significance of abnormal vital signs prior to in-hospital cardiac arrest." *Resuscitation*. 2016;98:112–7.
29. Barwise A, Thongprayoon C, Gajic O, et al. "Delayed rapid response team activation is associated with increased hospital mortality, morbidity, and length of stay in a tertiary care institution." *Crit Care Med*. 2016;44(1):54–63.
30. Devita MA, Bellomo R, Hillman K, et al. "Findings of the first consensus conference on medical emergency teams." *Crit Care Med*. 2006;34(9):2463–78.
31. Jones D, Rubulotta F, Welch J. "Rapid response teams improve outcomes: yes." *Intensive Care Med*. 2016;42(4):593–5.
32. "Ethical AI in Healthcare: Integrating Zero-Knowledge Proofs and Smart Contracts for Transparent Data Governance." *MDPI*, 2025.