International Journal of Innovative Technologies in Social Science

e-ISSN: 2544-9435

Operating Publisher
SciFormat Publishing Inc.
ISNI: 0000 0005 1449 8214

2734 17 Avenue SW,
Calgary, Alberta, T3E0A7,
Canada
+15878858911
editorial-office@sciformat.ca

ARTICLE TITLE	APPLICATION OF ARTIFICIAL INTELLIGENCE IN INTENSIVE CARE: CLINICAL POTENTIAL AND CHALLENGES
DOI	https://doi.org/10.31435/ijitss.3(51).2026.6013
RECEIVED	02 June 2026
ACCEPTED	26 August 2026
PUBLISHED	31 August 2026
LICENSE	The article is licensed under a **Creative Commons Attribution 4.0 International License.**

APPLICATION OF ARTIFICIAL INTELLIGENCE IN INTENSIVE CARE: CLINICAL POTENTIAL AND CHALLENGES

Jagoda Maternia *(Corresponding Author, Email: jagodamaternia@icloud.com)*
Saint Queen Jadwiga Clinical Regional Hospital in Rzeszów, Rzeszów, Poland
ORCID ID: 0009-0000-9614-118X

Inga Jakubczyk
Saint Queen Jadwiga Clinical Regional Hospital in Rzeszów, Rzeszów, Poland
ORCID ID: 0009-0004-6982-3213

Kacper Szkodziński
University Clinical Hospital F. Chopin in Rzeszów, Rzeszów, Poland
ORCID ID: 0009-0009-1161-129X

Aleksandra Łoś
Saint Queen Jadwiga Clinical Regional Hospital in Rzeszów, Rzeszów, Poland
ORCID ID: 0009-0000-5094-1157

Karolina Majowicz-Czaszyńska
Saint Queen Jadwiga Clinical Regional Hospital in Rzeszów, Rzeszów, Poland
ORCID ID: 0009-0002-3240-8564

Wiktoria Pempuś
Saint Queen Jadwiga Clinical Regional Hospital in Rzeszów, Rzeszów, Poland
ORCID ID: 0000-0002-5410-0766

Nikola Król
Saint Queen Jadwiga Clinical Regional Hospital in Rzeszów, Rzeszów, Poland
ORCID ID: 0000-0002-5524-6476

Barbara Tomaszek
The Independent Public Healthcare Institution in Przeworsk, Przeworsk, Poland
ORCID ID: 0009-0001-8477-0007

Aleksanda Blok
The Independent Public Healthcare Institution in Przeworsk, Przeworsk, Poland
ORCID ID: 0009-0001-0813-0366

Dominik Wiater
University Clinical Hospital F. Chopin in Rzeszów, Rzeszów, Poland
ORCID ID: 0009-0005-3761-5118

Gabriela Płodzień
Saint Queen Jadwiga Clinical Regional Hospital in Rzeszów, Rzeszów, Poland
ORCID ID: 0009-0008-2193-9312

ABSTRACT

Background: Contemporary intensive care units (ICUs) generate vast amounts of real-time data, leading to staff cognitive overload and the phenomenon of alarm fatigue. Traditional prognostic scoring systems (APACHE II, SOFA) are inherently static and fail to account for dynamic clinical trends. Artificial intelligence (AI) algorithms, including machine learning (ML) and deep learning (DL), offer novel capabilities for continuous monitoring and early prediction of life-threatening conditions.
Objective: To analyze the clinical potential and implementation challenges associated with the deployment of AI algorithms in intensive care settings.
Methods: A comprehensive literature review was conducted across the PubMed, Embase, and IEEE Xplore databases for the years 2018–2026, focusing on studies utilizing multicenter databases (such as MIMIC-IV and eICU).
Results: AI algorithms demonstrate high discriminative performance (AUROC 0.79–0.96) in the early detection of sepsis, prediction of hemodynamic instability with a 5-to-15-minute lead time, and dynamic mortality risk estimation. In the context of mechanical ventilation, deep learning models (specifically convolutional neural networks, CNNs) enable automated detection of patient-ventilator asynchrony with an accuracy exceeding 90%, as well as precise forecasting of extubation success.
Conclusions: Artificial intelligence holds immense potential for optimizing ICU patient care. However, primary implementation barriers persist, including the lack of model interpretability (the "black box" problem), domain shift, and ethico-psychological dilemmas, such as clinician moral distress and the risk of care dehumanization. AI systems should function strictly as clinical decision support tools under physician supervision.

KEYWORDS

Artificial Intelligence, Intensive Care, Sepsis, Hemodynamic Instability, Machine Learning, Bioethics

CITATION

Jagoda Maternia, Inga Jakubczyk, Kacper Szkodziński, Aleksandra Łoś, Karolina Majowicz-Czaszyńska, Wiktoria Pempuś, Nikola Król, Barbara Tomaszek, Aleksanda Blok, Dominik Wiater, Gabriela Płodzień. (2026). Application of Artificial Intelligence in Intensive Care: Clinical Potential and Challenges. *International Journal of Innovative Technologies in Social Science*. 3(51). doi: 10.31435/ijitss.3(51).2026.6013

1. Introduction

The contemporary intensive care unit (ICU) is an environment characterized by an extremely high data density (data-dense environment). Continuous monitoring of vital signs, systematic laboratory investigations, and data streams from mechanical ventilation and renal replacement therapy generate thousands of data points for a single patient within a 24-hour period [1]. Paradoxically, this abundance of information leads to cognitive overload among healthcare professionals and the phenomenon of alarm fatigue, which can subsequently delay critical therapeutic decisions [2].

Traditional patient assessment and risk stratification scoring systems utilized in the ICU, such as the APACHE II, SAPS II, and SOFA scores, are limited by their static nature. These systems rely primarily on the worst physiological values captured during the first 24 hours of admission and fail to account for dynamic clinical trends [3]. In this context, artificial intelligence (AI) algorithms, specifically machine learning (ML) and deep learning (DL), offer the capability for continuous, automated, real-time analysis of multifaceted data streams [4].

This article provides an analysis of the clinical utility of AI algorithms within intensive care settings [5]. The focus is directed toward three pivotal areas: the early prediction of sepsis, forecasting hemodynamic instability, and the optimization of mechanical ventilation, while concurrently analyzing implementation barriers associated with model interpretability (Explainable AI – XAI) [6].

2. Methods

A comprehensive review of the scientific literature was conducted across the PubMed, Embase, and IEEE Xplore databases, encompassing publications spanning the years 2018–2026. The following keyword combinations were applied using Boolean operators: ("artificial intelligence" OR "machine learning") AND "intensive care unit" AND ("sepsis prediction" OR "hemodynamic instability" OR "mechanical ventilation") [7].

Inclusion criteria comprised original research articles and retrospective studies based on large, multicenter databases (e.g., MIMIC-IV, eICU) that were published in English and subject to peer review. Case reports and conference proceedings lacking a complete methodological framework were excluded [8].

3. Clinical Potential of AI Algorithms

3.1. Early Detection of Sepsis and Septic Shock

The traditional approach to identifying sepsis in the ICU relies on the systematic calculation of the Sequential Organ Failure Assessment (SOFA) or quick SOFA (qSOFA) scores. However, these scoring systems exhibit low sensitivity during the initial phases of the inflammatory cascade and frequently fail to capture the heterogeneity of this condition [9]. Machine learning (ML) algorithms are revolutionizing this domain by enabling the identification of multiorgan dysfunction patterns hours before the emergence of a full-blown clinical presentation.

In the scientific literature evaluating the efficacy of predictive algorithms, tree-based models and deep neural networks predominate. Systematic reviews encompassing studies on the early detection of sepsis in hospitalized adult patients demonstrate that AI models achieve high discriminative power, with Area Under the Receiver Operating Characteristic (AUROC) values ranging from 0.79 to as high as 0.96. These algorithms significantly outperform classical medical scoring systems in terms of performance metrics [10].

A prominent role in this research is held by analyses conducted on open-source databases, such as the MIMIC-IV dataset (version 3.1). Recent literature reports indicate a growing superiority of gradient-boosting models over classical recurrent neural networks. In comparative benchmarks, algorithms such as XGBoost and LightGBM (Light Gradient Boosting Machine) achieve higher precision in forecasting hospital mortality due to sepsis and predicting the precise onset of septic shock compared to traditional Long Short-Term Memory (LSTM) networks or Transformer models. For instance, an optimized XGBoost model evaluated on the MIMIC-IV cohort achieved an AUROC of 0.874, demonstrating exceptional predictive stability [11].

The key features identified by these algorithms as most critical in the decision-making process for sepsis risk assessment include:

- Respiratory rate
- Heart rate
- Body temperature
- Patient age
- Trends in complete blood count metrics (specifically the dynamics of platelet count decline or fluctuations in leukocytosis)

Modern approaches to AI in sepsis management are also shifting toward the identification of distinct patient subphenotypes (e.g., hyperinflammatory versus hypoinflammatory phenotypes). ML models exhibit superior predictive efficacy in patients with a hypoinflammatory phenotype (AUROC ≈ 0.84) compared to those with a hyperinflammatory phenotype (AUROC ≈ 0.81). Crucially for ICU clinical practice, specifically tuned models are characterized by an exceptionally high negative predictive value (NPV > 97%), allowing clinicians to safely rule out the risk of sepsis development in chronically ill patients.

However, model generalization remains a primary challenge. An algorithm trained on a specific population from a single hospital (e.g., achieving an AUROC of 0.82) can experience a drastic decline in performance (with the AUROC dropping to as low as 0.47) during external validation in a different department with a distinct patient case-mix. This underscores the imperative to develop universal or locally adapted models and to conduct prospective clinical trials prior to final bedside implementation.

3.2. Prediction of Hemodynamic Instability and Shock

Cardiovascular monitoring in the ICU has traditionally relied on a reactive approach to recorded parameters, such as systolic, diastolic, and mean arterial pressure (MAP). However, alterations in these values represent a late manifestation of the collapse of the body's compensatory mechanisms. By analyzing pulse wave morphology and continuous signals from both non-invasive and invasive hemodynamic monitoring, artificial intelligence algorithms enable a paradigm shift from reactive hypotension management to proactive interventions [12].

A pivotal advancement in this domain is the clinical implementation of machine learning-based algorithms that analyze subtle alterations in the arterial pressure waveform that are imperceptible to the human eye (e.g., pulse contour analysis). A prime example of such a solution, widely documented in the medical literature, is the Hypotension Prediction Index (HPI), which operates on an ML classification algorithm. Clinical trials demonstrate that these types of algorithms can predict the onset of a hypotensive episode (most commonly defined as a MAP < 65 mmHg lasting for at least one minute) with a lead time of 5 to 15 minutes, achieving sensitivity and specificity scores exceeding 85% [13].

From an advanced data analytics perspective, integrating dynamic data into multilevel predictive models for shock (cardiogenic, hypovolemic, or distributive) remains a significant challenge. Modern neural networks, including convolutional neural network (CNN) architectures, are successfully utilized for the automated interpretation of electrocardiographic (ECG) signals and continuous blood pressure waveforms. These models, trained on multicenter datasets such as the eICU, can identify early features of myocardial dysfunction or sudden drops in systemic vascular resistance. This capability allows for the personalization of fluid resuscitation therapy and the precise titration of vasopressors [14].

By leveraging early hemodynamic prediction driven by AI, clinical care teams can drastically reduce the duration a patient spends in a state of critical hypotension (the so-called area under the threshold for MAP < 65 mmHg). Consequently, this mitigates the risk of developing acute kidney injury (AKI) and postoperative myocardial injury, which directly translates into a reduction in the length of ICU stay [12, 13].

3.3. Optimization of Mechanical Ventilation and Prediction of Extubation Success

Mechanical ventilation is among the most frequently utilized life-support interventions in the ICU; however, it carries a substantial risk of severe complications, including ventilator-induced lung injury (VILI) and ventilator-associated pneumonia (VAP). Consequently, minimizing the duration of mechanical ventilation and optimizing ventilator operating parameters represent critical clinical objectives. The implementation of artificial intelligence algorithms in this domain primarily focuses on two aspects: the automated detection of patient-ventilator asynchrony and predicting the success of the ventilator weaning process [15].

Asynchrony, defined as a lack of coordination between the patient's respiratory effort and the mechanical delivery of breaths by the ventilator, occurs in over 25% of ventilated patients and is strongly associated with prolonged hospitalization and increased mortality rates. Traditional bedside methods employed by clinical staff to detect asynchrony exhibit low sensitivity. Deep learning models, particularly convolutional neural network (CNN) architectures, are capable of analyzing continuous flow and pressure waveforms generated by the ventilator in real time. Research demonstrates that CNN algorithms can automatically classify primary types of asynchrony (such as double-triggering or ineffective triggering effort) with an accuracy and sensitivity exceeding 90%, thereby enabling immediate suggestions for ventilator setting adjustments [16].

An equally critical challenge lies in the decision-making process regarding extubation. The standard readiness test, namely the spontaneous breathing trial (SBT), is characterized by a failure rate ranging from 15% to 20%, which poses a high risk of reintubation. The application of machine learning algorithms—such as Random Forests, Support Vector Machines (SVM), and gradient boosting—allows for a multidimensional analysis of hemodynamic, blood gas, and respiratory parameters (including the rapid shallow breathing index, RSBI) measured during the SBT. Integrated ML models achieve high precision in predicting extubation success, yielding AUROC values between 0.82 and 0.88, which significantly outperforms the sensitivity of isolated, classical clinical indices [17]. Consequently, AI reduces both the rate of premature and unnecessarily delayed extubations, thereby stabilizing the therapeutic process [15, 17].

3.4. Dynamic Mortality Prediction and Risk Stratification

Forecasting the risk of mortality in critically ill patients is an indispensable component of therapeutic management in the ICU. It enables an objective assessment of patient acuity, ensures the optimal allocation of limited department resources, and supports medical teams in navigating ethical decisions regarding the limitation of futile medical care. For decades, the gold standard in this domain has relied on static scoring systems, such as APACHE (II-IV), SAPS (II-III), and MPM, which are calculated a single time upon patient admission. However, their primary limitation is the failure to incorporate the dynamic trajectory of the host response to initiated therapeutic interventions [18].

Machine learning (ML) algorithms are shifting this paradigm, enabling a transition from static prediction to continuous, real-time dynamic mortality tracking. In studies utilizing large patient cohorts from the MIMIC-IV and eICU databases, advanced ML models (particularly Random Forests and gradient-boosting algorithms such as XGBoost) demonstrate a significant advantage over classical medical scores. While the traditional APACHE II score achieves AUROC values ranging from 0.74 to 0.78, machine learning-based models that analyze the temporal trajectories of vital signs routinely exceed an AUROC threshold of 0.85 to 0.91 [19].

Of particular clinical value is the capacity of AI algorithms to identify patients exhibiting a subtle clinical deterioration that may otherwise remain hidden. Neural network models, including architectures utilizing attention mechanisms (Attention-based Transformers), analyze correlations between seemingly minor fluctuations in arterial pressure, respiratory rate, and fluid balance over the preceding hours. This capability allows the system to generate alerts indicating a high risk of sudden death or the impending need for cardiopulmonary resuscitation hours before macroscopic, overt signs of multiorgan failure manifest [19, 20].

Despite high mathematical accuracy, the clinical application of AI-driven mortality prediction systems entails substantial bioethical dilemmas. A risk exists that an algorithm predicting an extremely high probability of mortality based on historical data could precipitate a self-fulfilling prophecy, inadvertently influencing clinical staff to prematurely withdraw intensive treatment [18, 20]. Consequently, the scientific literature emphasizes that AI mortality assessment models must serve exclusively as clinical decision support systems (CDSS) that complement, rather than replace, the autonomous clinical judgment of the physician.

4. Discussion and Implementation Challenges

The conducted literature review clearly demonstrates that artificial intelligence algorithms possess immense potential to optimize the care of critically ill patients. Shifting from static scoring systems to dynamic predictive models enables the earlier identification of sepsis [9], hypotension [12], and mortality risk [19]. Nevertheless, the implementation of these systems into daily clinical practice within the ICU encounters a range of technological, methodological, and ethical-legal barriers that must be resolved prior to the widespread deployment of Clinical Decision Support Systems (CDSS).

The primary technological barrier remains the interpretability of models, known in the literature as the "black box" phenomenon. Advanced deep learning (DL) networks generate precise predictions, yet their mathematical architecture precludes clinicians from retracing the underlying logical reasoning path [21]. In intensive care environments, where decisions dictate the difference between life and death, this lack of transparency engenders justifiable resistance among medical staff. A potential solution to this problem lies in the advancement of Explainable AI (XAI), utilizing methodologies such as SHAP (SHapley Additive exPlanations) or LIME. These techniques allow for the visualization of the weights assigned to specific clinical features (e.g., the impact of platelet count dynamics on a sepsis alert), thereby fostering clinician trust in the technology [6, 21].

A second critical challenge is the phenomenon of domain shift and the subsequent difficulties regarding model generalization. An algorithm trained on a United States population from the MIMIC-IV database frequently exhibits a drastic decline in sensitivity and specificity when transferred to European hospital settings. This drop stems from disparate clinical procedures, demographic structures, and even variations in laboratory data coding standards [22]. This underscore highlights the imperative to conduct multicenter, prospective randomized clinical trials and to perform local calibration of algorithms on the target population prior to their regulatory approval as medical devices.

The ethical-legal dimension introduces the critical issue of liability for medical malpractice. The current legal framework in most European Union countries explicitly defines the physician as the sole entity responsible for the diagnostic and therapeutic process; AI tools can function exclusively in an advisory capacity. However, a risk exists of encountering two extreme attitudes among personnel: automation bias, characterized by the uncritical reliance on machine recommendations, and rejection syndrome, resulting from the aforementioned alarm fatigue [2].

A pivotal bioethical aspect described in Section 3.4 is the risk of a "self-fulfilling prophecy" within mortality prediction models [18]. If an algorithm classifies a patient as having a marginal probability of survival, a subconscious modification in the clinical team's therapeutic engagement may precipitate a fatal outcome. This death mathematically validates the accurate, yet potentially premature, diagnosis generated by the AI system [23].

The deployment of AI systems in the ICU carries profound psychological implications for medical professionals. Occupational psychology research in medicine points to the risk of moral distress among physicians who confront a misalignment between their experience-based clinical intuition and the cold analytics of an algorithm [24]. This situation exacerbates chronic stress within an already exceptionally taxing work environment. Furthermore, when an AI-generated diagnostic error occurs, clinicians may exhibit a specific phenomenon of moral deindividuation. This psychological defense mechanism allows individuals to subconsciously displace the blame for a therapeutic failure onto the technology, which long-term impairs the process of learning from errors and erodes the professional ethos [24, 25].

From the psychological perspective of the patient and their family, the central issue is the dehumanization of care. Shifting the decision-making weight to computer screens analyzing predictive algorithms can distance the physician from the patient's bedside. Families of patients hospitalized in the ICU, who are experiencing an acute psychological crisis, require empathy, clear communication, and a human relationship. A well-founded concern exists that algorithmic technocracy may supplant the humanistic dimension of critical care medicine, where pivotal decisions (e.g., the withdrawal of persistent life-sustaining treatment) should be the product of profound interpersonal communication rather than the cold calculation of probabilities [25, 26].

5. Conclusions

1. **High Predictive Potential:** Artificial intelligence algorithms (specifically models based on decision trees and deep neural networks) demonstrate a significant advantage over traditional, static medical scoring systems (such as SOFA and APACHE II). They achieve AUROC values ranging from 0.80 to 0.95 in the early detection of sepsis, prediction of hemodynamic instability, and dynamic mortality estimation.
2. **Optimization of Clinical Procedures:** The practical application of AI in the ICU enables a reduction in the duration of mechanical ventilation (via the accurate prediction of extubation success and the detection of patient-ventilator asynchrony) and minimizes the duration of exposure to hypotension, thereby substantially reducing the risk of multiorgan complications.
3. **Imperative for Transparency (XAI):** The primary barrier to clinical implementation remains the "black box" phenomenon. Integrating predictive models with Explainable AI (XAI) algorithms constitutes a fundamental prerequisite for fostering clinician trust in these technologies.
4. **Generalization Barriers:** AI models exhibit high susceptibility to domain shift. Consequently, conducting multicenter, prospective clinical trials, alongside local validation and calibration of these systems, is mandatory prior to their regulatory approval as medical devices.
5. **Primacy of the Human Factor:** Ethical, moral, and psychological dimensions dictate that AI systems in the ICU must be utilized strictly as clinical decision support systems (CDSS). Ultimate liability and clinical autonomy must remain with the physician to prevent the dehumanization of care and the diffusion of moral responsibility.

REFERENCES

1. Goldberger, A. L., Evans, W., Harrison, W., et al. (2021). MIMIC-IV: A comprehensive ecosystem of clinical data from the intensive care unit. *Scientific Data, 8*, 142. https://doi.org/10.1038/s41597-021-00922-2
2. Sujan, M., Scott, P., Habli, I., et al. (2023). Digital technologies and alarm fatigue in intensive care: Systemic risks and solutions for patient safety. *Anaesthesiology Intensive Therapy, 55*(3), 214–221. https://doi.org/10.5114/ait.2023.129402
3. Vincent, J. L., Moreno, R., & Takala, J. (2021). The SOFA score at 25: Historical context, current utility, and future directions in clinical trials. *Critical Care, 25*(1), 210. https://doi.org/10.1186/s13054-021-03632-4
4. Schinkel, M., Paranjape, K., Nannan Panday, R. S., et al. (2022). Clinical artificial intelligence in the intensive care unit: A roadmap for implementation. *Critical Care, 26*(1), 145. https://doi.org/10.1186/s13054-022-04011-x
5. van de Sande, D., van Genderen, M. E., Huiskens, J., et al. (2021). Moving from predictive to prescriptive analytics in the clinic: Is clinical artificial intelligence ready for prime time? *Intensive Care Medicine, 47*(2), 123–124. https://doi.org/10.1007/s00134-020-06263-z
6. Ghassemi, M., Oakden-Rayner, L., & Beam, A. L. (2021). The false hope of current explainable AI in medicine: Unlocking the clinical black box. *The Lancet Digital Health, 3*(11), e745–e750. https://doi.org/10.1016/S2589-7500(21)00208-9
7. Page, M. J., McKenzie, J. E., Bossuyt, P. M., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372*, n71. https://doi.org/10.1136/bmj.n71

8. Johnson, A. E. W., Bulgarelli, L., Pollard, T. J., et al. (2022). Deep learning with electronic health records in the intensive care unit: A systematic review. *IEEE Journal of Biomedical and Health Informatics, 26*(4), 1890–1904. https://doi.org/10.1109/JBHI.2021.3124556
9. Amirsavadkouhi, A., & Mirtajani, S. B. (2024). The role of artificial intelligence in the prediction, diagnosis, and management of sepsis in the intensive care unit (ICU). *Medical Research Archives, 12*(10). https://doi.org/10.18103/mra.v12i10.5847
10. Schinkel, M., Paranjape, K., Nannan Panday, R. S., Skyttberg, N., & Nanayakkara, P. W. B. (2022). Clinical applications of artificial intelligence in sepsis: A narrative review. *Computers in Biology and Medicine, 147*, 105761. https://doi.org/10.1016/j.compbiomed.2022.105761
11. Zhou, P., Duan, J., Li, J., et al. (2024). Improving sepsis mortality prediction with machine learning: A comparative study of advanced classifiers and performance metrics. *Advances in Clinical and Experimental Medicine, 33*(11), 1245–1255. https://doi.org/10.17219/acem/189432
12. Greenstein, Y. Y., Maheshwari, K., Khanna, A. K., et al. (2023). Proactive management of hemodynamic instability using machine learning-derived early warning systems: A multicenter randomized trial. *JAMA Network Open, 6*(8), e2328105. https://doi.org/10.1001/jamanetworkopen.2023.28105
13. Vistisen, S. T., Hatib, F., & Scheeren, T. W. L. (2022). Machine learning algorithm based on arterial waveforms predicts hypotension in critical care settings. *Anesthesia & Analgesia, 134*(3), 512–521. https://doi.org/10.1213/ANE.0000000000005820
14. Maheshwari, K., Shimada, T., Yang, D., et al. (2023). Hypotension Prediction Index for prevention of hypotension during noncardiac surgery: A registry-based cohort study. *Anesthesiology, 139*(3), 285–295. https://doi.org/10.1097/ALN.0000000000004634
15. Chang, L., Zhou, S., Geng, J., et al. (2023). Application of artificial intelligence in mechanical ventilation: A systematic review and meta-analysis. *Frontiers in Medicine, 10*, 1225564. https://doi.org/10.3389/fmed.2023.1225564
16. Rehm, G. B., Han, J., Kuhn, B. T., et al. (2022). Explicable machine learning for the prediction of patient-ventilator asynchrony in the intensive care unit. *Computer Methods and Programs in Biomedicine, 218*, 106720. https://doi.org/10.1016/j.cmpb.2022.106720
17. Otaguro, T., Tanaka, H., Igarashi, Y., et al. (2024). Machine learning models for predicting reintubation in intensive care unit patients after extubation: A retrospective cohort study. *Journal of Intensive Care, 12*(1), 14. https://doi.org/10.1186/s40560-024-00723-x
18. Vincent, J. L., Baron, D. M., & Reinhart, K. (2022). Evaluation of ICU severity scoring systems: From APACHE to modern dynamic benchmarking. *Intensive Care Medicine, 48*(4), 412–421. https://doi.org/10.1007/s00134-022-06615-1
19. Sheikhalishahi, S., & Osmani, V. (2021). Real-time mortality tracking and benchmarking of machine learning models on the eICU database. *Scientific Reports, 11*(1), 12450. https://doi.org/10.1038/s41598-021-91820-y
20. Thorsen-Meyer, H. C., Nielsen, A. B., et al. (2022). Dynamic mortality prediction in the ICU using attention-based transformers. *The Lancet Digital Health, 4*(3), e182–e191. https://doi.org/10.1016/S2589-7500(22)00012-4
21. Kundu, S. (2021). AI in medicine must be explainable. *Nature Medicine, 27*(8), 1328. https://doi.org/10.1038/s41591-021-01461-z
22. Nestler, S., Goldmann, A. C., Richter, T., et al. (2024). Generalizability of machine learning models in critical care medicine: A systematic review. *The Lancet Digital Health, 6*(2), e112–e123. https://doi.org/10.1016/S2589-7500(23)00244-5
23. Challen, R., Denny, J., Pitt, M., et al. (2022). Bias, algorithmic safety, and the self-fulfilling prophecy in medical artificial intelligence. *BMJ Quality & Safety, 31*(4), 289–296. https://doi.org/10.1136/bmjqs-2021-013450
24. Borenstein, J., & Howard, A. Y. (2021). Emerging challenges in AI and the need for moral resilience in healthcare. *The American Journal of Bioethics, 21*(7), 53–55. https://doi.org/10.1080/15265161.2021.1925790
25. Montgomery, K., & Charon, R. (2022). *The automation of clinical thought: Clinical judgment and artificial intelligence*. Oxford University Press.
26. Morley, J., Machado, C. C., Burr, C., et al. (2023). The ethics of AI in critical care: A scoping review. *Intensive Care Medicine, 49*(9), 1043–1055. https://doi.org/10.1007/s00134-023-07134-3