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	ARTIFICIAL INTELLIGENCE IN INTENSIVE CARE UNITS: CLINICAL APPLICATIONS, ETHICAL CHALLENGES AND FUTURE DIRECTIONS - A NARRATIVE REVIEW
DOI	https://doi.org/10.31435/ijitss.2(50).2026.5480
RECEIVED	26 March 2026
ACCEPTED	11 June 2026
PUBLISHED	30 June 2026
LICENSE	The article is licensed under a **Creative Commons Attribution 4.0 International License.**

ARTIFICIAL INTELLIGENCE IN INTENSIVE CARE UNITS: CLINICAL APPLICATIONS, ETHICAL CHALLENGES AND FUTURE DIRECTIONS - A NARRATIVE REVIEW

Jakub Świech *(Corresponding Author, Email: jasza.swiech@gmail.com)*
University Clinical Hospital No. 2 of Pomeranian Medical University in Szczecin, Poland
ORCID ID: 0009-0001-1732-4162

Marta Kamrowska
University Clinical Hospital No. 2 of Pomeranian Medical University in Szczecin, Poland
ORCID ID: 0009-0001-1570-0144

Katarzyna Apanasewicz
J. Strus Multidisciplinary Municipal Hospital in Poznań, Poland
ORCID ID: 0009-0001-0765-5398

Filip Kopacki
University Clinical Hospital No. 2 of Pomeranian Medical University in Szczecin, Poland
ORCID ID: 0009-0005-7897-8222

Weronika Szczeblewska
Independent Public Provincial Integrated Hospital in Szczecin, Poland
ORCID ID: 0009-0007-3956-622X

Daniel Jaskot
Independent Public Provincial Integrated Hospital in Szczecin, Poland
ORCID ID: 0009-0004-8678-9838

Jakub Marzec
University Clinical Hospital No. 1 of Pomeranian Medical University in Szczecin, Poland
ORCID ID: 0009-0007-8336-7260

Natalia Kita
University Clinical Hospital No. 2 of Pomeranian Medical University in Szczecin, Poland
ORCID ID: 0009-0001-2277-6198

Natalia Sztenc
Pomeranian Medical University, Szczecin, Poland
ORCID ID: 0009-0006-4307-7958

Patrycja Broniszewska
University Clinical Hospital No. 2 of Pomeranian Medical University in Szczecin, Poland
ORCID ID: 0009-0000-6577-307X

ABSTRACT

Artificial intelligence (AI) is increasingly recognized as a transformative tool in intensive care units (ICUs), where large volumes of complex patient data present both challenges and opportunities for clinical decision-making. This narrative review aims to synthesize current evidence on the applications of AI in ICU settings, focusing on predictive analytics, clinical decision support systems, therapeutic optimization, and ethical and implementation challenges.
A targeted literature search was conducted across PubMed, Scopus, Web of Science, and Google Scholar, prioritizing peer-reviewed studies published between 2015 and 2025. Findings were synthesized thematically. The results indicate that AI models can outperform traditional clinical scoring systems in predicting sepsis, mortality, and clinical deterioration. Reinforcement learning approaches further demonstrate potential for personalized treatment strategies.
However, significant barriers to implementation remain, including limited external validation, lack of interpretability, ethical concerns, and infrastructural limitations. AI holds substantial promise, but its successful integration requires careful consideration of clinical, ethical, and organizational factors.
Future research should focus on explainable AI, prospective validation, and seamless integration into clinical workflows.

KEYWORDS

Artificial Intelligence; Intensive Care Unit; Machine Learning; Clinical Decision Support; Critical Care; Digital Health

CITATION

Jakub Świech, Marta Kamrowska, Katarzyna Apanasewicz, Filip Kopacki, Weronika Szczeblewska, Daniel Jaskot, Jakub Marzec, Natalia Kita, Natalia Sztenc, Patrycja Broniszewska. (2026) Artificial Intelligence in Intensive Care Units: Clinical Applications, Ethical Challenges and Future Directions - a Narrative Review. *International Journal of Innovative Technologies in Social Science*. 2(50). doi: 10.31435/ijitss.2(50).2026.5480

INTRODUCTION

The intensive care unit (ICU) represents one of the most complex and data-intensive environments in modern healthcare. Critically ill patients require continuous monitoring, rapid clinical decision-making, and individualized therapeutic strategies. Clinicians must interpret vast amounts of heterogeneous data, including physiological signals, laboratory results, imaging findings, and electronic health records (EHRs), often under significant time pressure. Despite the availability of established clinical scoring systems such as the SOFA score and APACHE II, these tools are limited by their static nature and reliance on a restricted set of variables (Vincent et al., 1996).

Over the past decade, the rapid digitization of healthcare systems has enabled the collection of large-scale clinical datasets, creating new opportunities for data-driven approaches to critical care. One of the most influential developments in this field has been the introduction of publicly available ICU datasets such as the MIMIC-III database, which contains detailed clinical data from over 40,000 patients (Johnson et al., 2016). The subsequent release of MIMIC-IV (Johnson et al., 2021) further expanded access to high-quality ICU data, facilitating large-scale research and reproducibility.

These advances have coincided with significant progress in artificial intelligence (AI), particularly in machine learning (ML) and deep learning. AI systems are capable of analyzing complex, high-dimensional datasets and identifying patterns that may not be apparent through traditional statistical methods (LeCun et al., 2015; Esteva et al., 2019). In healthcare, AI has been increasingly applied to tasks such as diagnosis, prognosis, and treatment optimization, with promising results across multiple domains.

In the context of critical care, AI applications have focused primarily on predictive analytics, clinical decision support, and therapeutic optimization. For example, machine learning models have demonstrated improved performance in early detection of sepsis—a life-threatening condition associated with high mortality—compared to traditional diagnostic approaches (Nemati et al., 2018; Singer et al., 2016). Similarly, deep learning models trained on EHR data have shown superior accuracy in predicting in-hospital mortality and clinical deterioration (Rajkomar et al., 2018; Tomašev et al., 2019).

Reinforcement learning, a subset of AI that focuses on sequential decision-making, has also been explored in ICU settings. Komorowski et al. (2018) developed a reinforcement learning model capable of recommending optimal treatment strategies for sepsis, demonstrating the potential of AI to guide personalized clinical interventions. Subsequent studies have further investigated the use of reinforcement learning for dynamic treatment regimes, highlighting its applicability in complex and evolving clinical environments (Raghu et al., 2017).

In addition to predictive and decision-support applications, AI has been applied to the analysis of continuous monitoring data in ICUs. Machine learning models, including recurrent neural networks and long short-term memory (LSTM) architectures, have demonstrated the ability to process time-series data and detect early signs of clinical deterioration (Lipton et al., 2016; Ghassemi et al., 2015). These capabilities are particularly valuable in ICU settings, where early intervention can significantly improve patient outcomes.

Despite these promising developments, the integration of AI into routine ICU practice remains limited. A growing body of literature highlights the challenges associated with translating AI models from research settings to clinical environments. These challenges include lack of external validation, limited generalizability, and difficulties in integrating AI systems into existing clinical workflows (Wynants et al., 2020; Sendak et al., 2020). Furthermore, many AI models rely on retrospective datasets, which may not accurately reflect real-world clinical conditions.

Ethical and societal considerations also play a critical role in the adoption of AI in healthcare. Concerns regarding algorithmic bias, transparency, accountability, and patient autonomy have been widely discussed in the literature. Obermeyer et al. (2019) demonstrated that widely used healthcare algorithms may exhibit racial bias, highlighting the importance of ensuring fairness in AI systems. Similarly, Char et al. (2018) emphasized the ethical challenges associated with implementing machine learning in clinical practice, particularly in high-stakes environments such as ICUs.

The concept of "high-performance medicine," introduced by Topol (2019), emphasizes the integration of AI with human clinical expertise rather than replacement. In this model, AI systems serve as tools to augment clinician decision-making, enabling more precise and efficient care while preserving the central role of human judgment.

Given the rapid expansion of research in this field and the complexity of its clinical, ethical, and organizational implications, there is a need for a focused and integrative synthesis of existing evidence. While numerous studies have explored specific applications of AI in critical care, fewer reviews have combined predictive analytics, clinical decision support systems, therapeutic optimization, and implementation challenges within one structured discussion. This narrative review addresses that gap by providing a clinically oriented overview of AI applications in ICUs, with particular emphasis on translational barriers, ethical concerns, and the conditions necessary for safe implementation in real-world practice. This narrative review aims to synthesize and critically discuss current literature on artificial intelligence in intensive care units. Specifically, the objectives of this study are to: (1) examine the main clinical applications of AI in ICU settings, including predictive analytics, decision support systems, and therapeutic optimization; (2) assess the benefits and limitations of AI-based approaches reported in the literature; (3) discuss ethical, organizational, and implementation-related challenges associated with the use of AI in critical care; and (4) identify priorities for future research and safe clinical integration. By integrating findings from a wide range of studies, this review seeks to provide a comprehensive and critical overview of the role of AI in intensive care, contributing to ongoing discussions on its safe and effective implementation in clinical practice.

METHODOLOGY

This study was conducted as a narrative review aimed at synthesizing current evidence on the application of artificial intelligence (AI) in intensive care units (ICUs), with particular focus on predictive analytics, clinical decision support systems, therapeutic optimization, and the ethical and organizational implications of AI-assisted care.

A targeted, non-systematic literature search was performed across major electronic databases, including PubMed/MEDLINE, Scopus, Web of Science, and Google Scholar. The search strategy used combinations of keywords such as "artificial intelligence," "machine learning," "deep learning," "reinforcement learning," "intensive care unit," "critical care," "clinical decision support," "sepsis prediction," and "mortality prediction." The search was designed to identify influential and clinically relevant publications rather than to retrieve all available studies on the topic.

Priority was given to peer-reviewed articles published between 2015 and 2025 in order to capture recent advances in AI methodology and critical care applications following the expansion of deep learning research (LeCun et al., 2015). Earlier landmark publications were also included when necessary to provide clinical or methodological context, such as foundational work on ICU scoring systems (Vincent et al., 1996) and the development of large-scale clinical databases such as the MIMIC-III database (Johnson et al., 2016).

Eligible sources included clinical trials, observational studies, retrospective database analyses, machine learning model development studies, narrative reviews, and other influential papers relevant to AI use in intensive care. Particular attention was given to publications evaluating AI applications in key ICU domains, such as sepsis detection (Nemati et al., 2018), mortality prediction (Rajkomar et al., 2018), and reinforcement learning-based treatment strategies (Komorowski et al., 2018).

Because of the narrative character of this review, the identified literature was synthesized thematically rather than quantitatively. Studies were grouped into major domains, including predictive analytics, clinical decision support systems, automation and therapeutic optimization, big data and patient monitoring, ethical and legal challenges, and implementation barriers. This approach was intended to provide an integrative overview of a rapidly evolving field rather than a formal systematic assessment. Accordingly, this review should be interpreted as a critical narrative synthesis of influential and clinically relevant literature rather than as an exhaustive systematic review of all available evidence.

No formal risk-of-bias assessment tool was applied. Instead, emphasis was placed on the clinical relevance, methodological rigor, and reproducibility of included studies. Particular consideration was given to commonly discussed limitations in the literature, including limited external validation and the risk of bias in AI models (Wynants et al., 2020; Obermeyer et al., 2019).

This methodological approach allows for a comprehensive and flexible synthesis of a rapidly evolving field while maintaining a focus on clinically meaningful outcomes and the broader societal implications of AI in critical care.

This narrative review follows a structured approach to literature identification and thematic synthesis, aiming to balance comprehensiveness with clinical relevance. While not systematic, the review process was guided by clearly defined search terms and selection principles to enhance transparency and consistency.

RESULTS

Predictive Analytics in Intensive Care

Predictive analytics represents the most extensively studied and clinically advanced application of artificial intelligence (AI) in intensive care units (ICUs). The primary objective of these models is to identify patterns in complex clinical data that enable early detection of critical conditions, risk stratification, and outcome prediction.

Traditional clinical scoring systems, such as the SOFA score and APACHE II, have long been used to assess patient severity and predict outcomes. However, these models are inherently limited by their reliance on static variables and predefined thresholds (Vincent et al., 1996). In contrast, AI-based models can process large volumes of heterogeneous and time-dependent data, enabling more dynamic and individualized predictions (Rajkomar et al., 2018; Shickel et al., 2018).

Early Detection of Sepsis

Sepsis remains one of the leading causes of mortality in ICU settings, and early detection is critical for improving patient outcomes (Singer et al., 2016). AI-based predictive models have demonstrated substantial improvements over traditional diagnostic approaches.

Nemati et al. (2018) developed an interpretable machine learning model capable of predicting sepsis onset several hours before clinical recognition. The study highlighted the importance of incorporating temporal physiological data and demonstrated improved predictive accuracy compared to conventional methods.

Similarly, Henry et al. (2015) introduced the Targeted Real-time Early Warning Score (TREWS), which uses machine learning to continuously analyze patient data. A subsequent prospective study by Adams et al. (2022) demonstrated that TREWS implementation was associated with earlier treatment initiation and reduced mortality, providing one of the strongest pieces of evidence supporting real-world clinical impact of AI systems.

Other approaches have focused on deep learning architectures. Futoma et al. (2017) compared multiple machine learning models for sepsis detection and showed that sequential models outperform traditional classifiers by effectively capturing temporal dynamics. Moor et al. (2019) further demonstrated that temporal convolutional networks can significantly improve predictive performance in early sepsis detection.

Despite these advances, heterogeneity in sepsis definitions and variability in clinical practice remain challenges for model generalization (Wynants et al., 2020).

Mortality Prediction and Risk Stratification

Accurate prediction of mortality is essential for clinical decision-making, resource allocation, and communication with patients' families. AI models have shown considerable promise in this area.

Rajkomar et al. (2018) developed deep learning models using large-scale EHR data and demonstrated superior performance compared to traditional statistical methods in predicting in-hospital mortality. Their work emphasized the ability of AI systems to integrate diverse data sources, including structured and unstructured clinical data.

Tomašev et al. (2019) developed a deep learning model for predicting acute kidney injury (AKI), which is strongly associated with increased mortality in ICU patients. The model was able to predict AKI up to 48 hours in advance, highlighting the potential of AI for early intervention.

Harutyunyan et al. (2019) contributed to the field by providing benchmark datasets and standardized evaluation methods for ICU predictive modeling. Their work facilitates comparison between different machine learning approaches and supports reproducibility.

Shickel et al. (2018), in a comprehensive review, highlighted the superiority of deep learning models in handling complex clinical data, particularly in ICU settings where multiple variables interact dynamically.

However, concerns remain regarding model overfitting and lack of external validation. Wynants et al. (2020) reported that many predictive models demonstrate high performance in development datasets but fail to generalize to new populations.

Prediction of Clinical Deterioration

Beyond specific diagnoses such as sepsis, AI models have been widely applied to predict general clinical deterioration, including respiratory failure, cardiac arrest, and organ dysfunction.

Churpek et al. (2016) compared machine learning-based early warning scores with traditional scoring systems and found that AI models significantly improved prediction of clinical deterioration. Their findings support the use of machine learning for continuous patient monitoring and early intervention.

Ghassemi et al. (2015) emphasized the importance of analyzing time-series data in ICU settings, demonstrating that machine learning models can identify subtle changes in physiological parameters that precede clinical deterioration.

Recurrent neural networks, particularly long short-term memory (LSTM) models, have shown strong performance in modeling temporal data. Lipton et al. (2016) demonstrated that LSTM-based models can effectively predict clinical events using multivariate time-series data.

These findings highlight the potential of AI to enhance early warning systems and improve patient outcomes through timely intervention.

Comparative Performance of AI Models and Traditional Methods

A consistent finding across studies is that AI-based models outperform traditional statistical approaches in predictive accuracy.

Machine learning models can incorporate a larger number of variables and capture nonlinear relationships between clinical parameters, which are often missed by traditional models (Rajkomar et al., 2018; Shillan et al., 2019).

However, this increased accuracy often comes at the cost of reduced interpretability. While traditional scoring systems are transparent and easy to use, AI models—particularly deep learning systems—are often perceived as "black boxes" (Tonekaboni et al., 2019).

Wiens et al. (2019) emphasize that both accuracy and interpretability are essential for clinical adoption, suggesting that future research should focus on balancing these factors.

Methodological Limitations in Predictive Modeling

Despite promising results, several methodological limitations are consistently identified in the literature.

Data Quality and Missing Data

ICU datasets frequently contain missing or inconsistent data, which can affect model performance. Approaches such as imputation may introduce bias (Ghassemi et al., 2015).

Bias and Fairness

Bias in training datasets can lead to unequal model performance across patient populations. Obermeyer et al. (2019) demonstrated that healthcare algorithms may systematically disadvantage certain groups.

Lack of External Validation

Many models are developed using datasets such as MIMIC-III (Johnson et al., 2016), limiting generalizability. Wynants et al. (2020) identified lack of external validation as a major weakness in predictive modeling studies.

Overfitting

Complex models may perform well on training data but fail in real-world settings, highlighting the need for robust validation.

Synthesis of Findings

Overall, the reviewed studies suggest that AI-based predictive models may outperform traditional approaches in ICU settings, particularly in sepsis detection, mortality prediction, and early warning tasks. At the same time, their practical adoption remains limited by poor external validation, reduced interpretability, and integration challenges.

Clinical Decision Support Systems in Intensive Care

Clinical decision support systems (CDSS) represent one of the most clinically relevant and rapidly developing applications of artificial intelligence (AI) in intensive care units (ICUs). Unlike purely predictive models, CDSS aim to translate data-driven insights into actionable recommendations that can directly influence clinical decision-making.

In ICU environments—characterized by high patient acuity, complex treatment protocols, and time-critical decisions—CDSS have the potential to significantly improve both efficiency and quality of care (Topol, 2019). By integrating real-time patient data with advanced analytical models, AI-driven CDSS can assist clinicians in diagnosis, risk assessment, and treatment planning.

AI-Driven Decision-Making Models

AI-based CDSS typically rely on machine learning and reinforcement learning algorithms to generate clinical recommendations.

One of the most influential contributions in this field is the study by Komorowski et al. (2018), which introduced a reinforcement learning model—often referred to as the “AI Clinician”—for optimizing sepsis treatment. The model analyzed retrospective ICU data and generated treatment policies for fluid administration and vasopressor use. Importantly, the study demonstrated that patients whose treatment aligned more closely with AI-generated recommendations had lower mortality rates.

Similarly, Raghu et al. (2017) explored deep reinforcement learning approaches to sepsis management and showed that AI systems could identify treatment strategies that differ from standard clinical practice while maintaining comparable or improved outcomes.

These studies collectively suggest that AI-driven CDSS can move beyond passive prediction toward active clinical guidance, representing a significant shift in the role of AI in healthcare.

Real-Time Decision Support and Early Warning Systems

A key advantage of AI-based CDSS is their ability to operate in real time.

The Targeted Real-time Early Warning Score (TREWS), developed by Henry et al. (2015), is one of the most widely cited examples of a real-time AI system for sepsis detection. TREWS continuously analyzes patient data from electronic health records and generates alerts when patients are at risk of developing sepsis.

A large prospective study by Adams et al. (2022) demonstrated that the implementation of TREWS was associated with earlier clinical intervention and improved patient outcomes, including reduced mortality. This represents one of the strongest examples of successful translation of AI into clinical practice.

Churpek et al. (2016) also demonstrated that machine learning-based early warning systems outperform traditional scoring systems such as MEWS in predicting clinical deterioration.

However, the effectiveness of real-time CDSS depends heavily on alert accuracy. Excessive false positives can lead to alarm fatigue, reducing clinician responsiveness and potentially compromising patient safety (Sendak et al., 2020).

Integration with Clinical Workflow

Despite their potential, the integration of AI-based CDSS into clinical workflows remains a major challenge.

Sendak et al. (2020) emphasize that many AI systems fail not because of poor predictive performance but because they are not designed with clinical usability in mind. Common issues include:

- lack of interoperability with electronic health record systems
- poorly designed user interfaces
- disruption of established clinical routines

Wiens et al. (2019) argue that successful implementation requires a human-centered design approach, ensuring that AI tools align with clinician workflows and support rather than hinder decision-making.

Without proper integration, even highly accurate systems may be underutilized or ignored in practice.

Interpretability and Trust in AI Systems

A critical factor influencing the adoption of AI-based CDSS is interpretability.

Many AI models, particularly deep learning systems, are perceived as "black boxes," meaning that their internal decision-making processes are not easily understood. This lack of transparency can reduce clinician trust, especially in high-risk environments such as ICUs.

Tonekaboni et al. (2019) found that clinicians are more likely to adopt AI systems that provide clear explanations for their recommendations. Even when less accurate, interpretable models may be preferred over opaque high-performance models.

Wiens et al. (2019) similarly emphasize that interpretability is essential for ensuring accountability and enabling clinicians to validate AI outputs.

Taken together, these findings suggest the importance of explainable AI (XAI) in the future development of CDSS.

Human–AI Collaboration

Rather than replacing clinicians, AI-based CDSS are increasingly viewed as tools for enhancing human decision-making.

Topol (2019) introduced the concept of "high-performance medicine," in which AI systems and clinicians work collaboratively to improve patient care. In this model:

- AI provides data-driven insights
- clinicians apply contextual judgment and ethical reasoning

This collaborative approach is particularly important in ICU settings, where decisions often involve uncertainty and complex trade-offs.

The literature consistently suggests that the most effective CDSS are those that support, rather than replace, clinician expertise.

Ethical Considerations in CDSS

The use of AI-based CDSS introduces several ethical challenges.

Algorithmic bias, as demonstrated by Obermeyer et al. (2019), may result in unequal treatment recommendations across patient populations. This is particularly concerning in ICU settings, where decisions directly impact survival.

Char et al. (2018) emphasize the importance of transparency and accountability in AI systems, particularly when they influence clinical decisions.

In addition, reliance on AI systems may raise concerns about reduced clinician autonomy and overdependence on automated recommendations.

These ethical considerations must be addressed to ensure responsible implementation of CDSS in clinical practice.

Limitations and Challenges

Despite promising results, AI-based CDSS face several limitations:

- reliance on retrospective datasets such as MIMIC-III (Johnson et al., 2016)
- limited external validation (Wynants et al., 2020)
- challenges in generalizability across healthcare systems
- regulatory uncertainty

These challenges highlight the need for rigorous validation and cautious implementation.

Synthesis of Findings

The literature indicates that AI-based CDSS have the potential to significantly improve clinical decision-making in ICU settings.

Key advantages include:

- enhanced diagnostic accuracy
- real-time decision support
- personalized treatment recommendations

However, a central insight is that technical performance alone is insufficient for successful implementation. Factors such as usability, interpretability, and clinician trust are equally critical.

Future research should focus on:

- developing explainable models
- improving integration into clinical workflows
- conducting prospective validation studies

Automation and Therapeutic Optimization in Intensive Care

Automation and therapeutic optimization represent a more advanced stage of artificial intelligence (AI) application in intensive care units (ICUs), extending beyond prediction and decision support toward active participation in clinical management. These approaches focus on optimizing treatment strategies in real time, adapting interventions to individual patient trajectories, and reducing variability in clinical practice.

In contrast to traditional guideline-based care, which often relies on standardized protocols, AI-driven systems can dynamically adjust therapeutic decisions based on continuous patient data. This shift reflects a broader transition toward precision medicine and individualized care in critical care settings (Topol, 2019).

Reinforcement Learning for Treatment Optimization

Reinforcement learning (RL) has emerged as a particularly promising approach for therapeutic optimization in ICU environments, where decision-making is inherently sequential and time-dependent.

Komorowski et al. (2018) developed a landmark reinforcement learning model for sepsis management, often referred to as the "AI Clinician." By analyzing large-scale ICU data, the model learned optimal policies for fluid resuscitation and vasopressor administration. Importantly, the study demonstrated that deviations from AI-recommended treatment strategies were associated with increased mortality, suggesting potential clinical relevance.

Raghu et al. (2017) further explored deep reinforcement learning approaches for sepsis treatment and found that AI systems could identify treatment strategies that differed from standard clinical practice while maintaining or improving patient outcomes.

These studies support the potential to move beyond static protocols and enable adaptive, patient-specific treatment strategies.

Mechanical Ventilation and Respiratory Management

Mechanical ventilation is a cornerstone of ICU care, yet its management is complex and requires continuous adjustment of parameters such as tidal volume, respiratory rate, and positive end-expiratory pressure (PEEP).

AI-based approaches have been explored to optimize ventilator management and reduce complications such as ventilator-induced lung injury.

Prasad et al. (2017) proposed a reinforcement learning model to optimize ventilator weaning strategies. Their findings suggest that AI can identify optimal timing for weaning, potentially reducing the duration of mechanical ventilation and associated risks.

These applications highlight the potential of AI to enhance respiratory management through data-driven decision-making.

Hemodynamic Management and Drug Optimization

Hemodynamic instability is a common and life-threatening condition in ICU patients, particularly in cases of sepsis and septic shock. Optimal management requires careful titration of fluids, vasopressors, and other medications.

AI-based models offer the potential to optimize these interventions by analyzing patient-specific responses over time.

Komorowski et al. (2018) demonstrated that reinforcement learning can be used to recommend individualized dosing strategies for vasopressors and fluids. Similarly, Futoma et al. (2017) emphasized the importance of temporal modeling in predicting patient responses to interventions, enabling dynamic adjustment of treatment strategies.

Ghassemi et al. (2015) highlighted the potential of machine learning to identify complex patterns in physiological data that may inform treatment decisions in unstable patients.

These approaches represent a significant shift toward personalized medicine, where treatment is tailored to the unique characteristics of each patient rather than based solely on standardized guidelines.

Closed-Loop and Semi-Autonomous Systems

Closed-loop systems represent the most advanced form of automation in ICU care. These systems continuously monitor patient data and automatically adjust treatment parameters without direct human intervention.

Examples of such systems include:

- automated insulin delivery systems
- closed-loop anesthesia systems
- adaptive ventilator control systems

While these systems are already used in limited clinical contexts, the integration of AI algorithms has the potential to significantly enhance their capabilities.

However, fully autonomous systems raise important safety and ethical concerns. Wiens et al. (2019) emphasize that human oversight is essential, particularly in high-risk environments such as ICUs.

The concept of “human-in-the-loop” systems—where clinicians retain ultimate control over AI-driven decisions—is increasingly recognized as a necessary safeguard.

Benefits of AI-Driven Automation

The potential benefits of AI-driven automation in ICU settings include:

- improved precision in treatment decisions
- reduction in clinician workload
- faster response to patient deterioration
- increased consistency and standardization of care

Topol (2019) argues that automation can free clinicians from routine tasks, allowing them to focus more on patient-centered care.

In addition, AI-driven systems may help reduce variability in clinical practice, which is a known contributor to differences in patient outcomes.

Risks and Limitations of Automation

Despite its potential, AI-driven automation introduces several risks and limitations.

Safety and Reliability

Errors in automated systems may have immediate and severe consequences in critically ill patients. Robust validation and fail-safe mechanisms are therefore essential.

Overreliance on AI

There is a risk that clinicians may become overly dependent on AI systems, potentially reducing critical thinking and clinical judgment.

Generalizability

Many models are trained on datasets such as MIMIC-III (Johnson et al., 2016), which may limit their applicability across different healthcare settings.

Regulatory Challenges

The regulation of AI-driven therapeutic systems remains underdeveloped, particularly for adaptive models that evolve over time.

Wynants et al. (2020) highlight that many AI models lack sufficient validation for clinical deployment, raising concerns about safety and reliability.

Synthesis of Findings

Available evidence suggests that, although most data remain retrospective, AI-driven automation has the potential to fundamentally transform ICU care by enabling adaptive, personalized treatment strategies.

Key insights include:

- reinforcement learning enables dynamic treatment optimization
- AI can improve management of ventilation and hemodynamics
- automation may reduce variability and improve efficiency
- human oversight remains essential for safety and ethical reasons

A central conclusion is that automation should augment, rather than replace, clinician decision-making.

Future research should prioritize:

- prospective validation of reinforcement learning models
- development of safe and interpretable automation systems
- integration of AI into clinical workflows

Big Data and Patient Monitoring in Intensive Care

The intensive care unit (ICU) is one of the most data-rich environments in healthcare, characterized by continuous monitoring and the generation of high-frequency physiological data. The emergence of "big data" in critical care has played a central role in enabling the development and application of artificial intelligence (AI), providing the foundation for predictive modeling, decision support, and therapeutic optimization.

Sources and Nature of ICU Data

ICU data are inherently complex and heterogeneous, encompassing continuous physiological signals, laboratory findings, medication records, imaging data, and electronic health records. A major milestone in this domain was the introduction of large, publicly available datasets such as the MIMIC-III database (Johnson et al., 2016), which enabled large-scale machine learning research in ICU populations. The subsequent release of MIMIC-IV (Johnson et al., 2021) further expanded access to high-quality, longitudinal patient data.

These datasets have significantly improved reproducibility and benchmarking in AI research, allowing for standardized evaluation of predictive models (Harutyunyan et al., 2019).

Time-Series Data and Temporal Modeling

A defining characteristic of ICU data is its temporal structure. Unlike traditional datasets, ICU data consist of continuous time-series measurements that reflect rapidly changing patient conditions.

Machine learning models capable of handling temporal data—such as recurrent neural networks (RNNs) and long short-term memory (LSTM) networks—have demonstrated strong performance in this context (Lipton et al., 2016).

Moor et al. (2019) showed that incorporating temporal dynamics significantly improves predictive performance in sepsis detection. Similarly, Ghassemi et al. (2015) emphasized that time-series analysis enables early identification of clinical deterioration by capturing subtle physiological trends.

These findings highlight the importance of temporal modeling in ICU AI applications, particularly for real-time monitoring and early warning systems.

Data Integration and Interoperability

Despite the abundance of data, one of the major challenges in ICU environments is the integration of heterogeneous data sources.

Different systems—such as bedside monitors, laboratory systems, and electronic health records—often operate independently and use incompatible formats. This fragmentation limits the ability to develop comprehensive AI models.

Rajkomar et al. (2018) demonstrated that integrating diverse data types significantly improves predictive performance. However, achieving such integration in real-world clinical environments remains challenging.

Wiens et al. (2019) highlight that lack of interoperability is one of the key barriers to the effective implementation of machine learning in healthcare.

Alarm Systems and Alert Fatigue

Modern ICU monitoring systems rely heavily on alarms to alert clinicians to abnormal patient conditions. However, these systems often generate excessive numbers of alerts, many of which are false positives.

Alarm fatigue—defined as desensitization to frequent alarms—is a well-documented issue in critical care and has been associated with delayed responses and increased risk of adverse events.

AI has been proposed as a solution to this problem. By analyzing contextual patterns in patient data, machine learning models can generate more accurate and clinically relevant alerts.

Sendak et al. (2020) emphasize that improving alert specificity is critical for maintaining clinician trust and ensuring effective use of monitoring systems.

However, poorly calibrated AI systems may exacerbate alarm fatigue if they generate excessive or non-actionable alerts.

Data Quality, Missing Data, and Bias

The performance of AI models is highly dependent on data quality. ICU datasets often contain:

- missing values
- measurement errors
- inconsistencies across systems

Handling missing data remains a significant methodological challenge. Techniques such as imputation are commonly used but may introduce bias.

Wynants et al. (2020) identified poor data quality as a major contributor to bias in predictive models. Similarly, Obermeyer et al. (2019) demonstrated that biases in healthcare data can lead to unequal outcomes across patient populations.

Ensuring high-quality, representative datasets is therefore essential for developing reliable and equitable AI systems.

Data Privacy and Security

The use of large-scale patient data raises important concerns regarding privacy and data security.

Datasets such as the MIMIC-III database (Johnson et al., 2016) demonstrate that data sharing is possible when appropriate de-identification procedures are applied. However, the increasing complexity of AI models raises concerns about potential re-identification risks.

Regulatory frameworks, such as the General Data Protection Regulation (GDPR), impose strict requirements on data handling, including transparency, data minimization, and accountability.

Failure to address privacy concerns may undermine public trust and hinder the adoption of AI technologies in healthcare.

Infrastructure and Computational Requirements

The implementation of AI in ICU settings requires substantial computational infrastructure.

Key requirements include:

- real-time data processing capabilities
- large-scale data storage
- integration with hospital IT systems

Rajkomar et al. (2018) emphasize that scalable infrastructure is essential for deploying AI models in clinical practice.

In many healthcare systems, limited infrastructure and outdated technologies represent significant barriers to AI adoption.

Synthesis of Findings

The literature indicates that big data and continuous patient monitoring are fundamental to the development of AI in intensive care.

Key insights include:

- ICU environments generate high-frequency, multidimensional data suitable for AI analysis
- time-series modeling significantly enhances predictive capabilities
- data integration and interoperability remain major challenges
- AI has the potential to reduce alarm fatigue through improved alert systems
- data quality, bias, and privacy are critical concerns

A central conclusion is that data availability alone is insufficient—effective utilization requires robust infrastructure, integration, and governance.

Future research should focus on:

- improving data standardization and interoperability
- developing robust methods for handling missing data
- ensuring ethical and secure data use

Table 1. Major applications of artificial intelligence in intensive care units

Domain	Example methods	Main clinical use	Main limitations
Predictive analytics	Machine learning, deep learning	Sepsis prediction, mortality prediction, deterioration detection	Limited external validation, risk of overfitting
Clinical decision support systems	Machine learning, reinforcement learning	Diagnostic support, treatment recommendations, early warning alerts	Alert fatigue, poor workflow integration, limited trust
Therapeutic optimization	Reinforcement learning	Fluid therapy, vasopressor dosing, ventilator weaning	Mostly retrospective evidence, safety concerns
Patient monitoring and big data	RNNs, LSTMs, temporal convolutional models	Continuous monitoring, time-series analysis, risk stratification	Data fragmentation, missing data, interoperability problems
Ethical and implementation issues	Explainable AI, governance frameworks	Fairness, accountability, transparency, safe deployment	Bias, regulatory uncertainty, infrastructure limitations

The main domains of AI application in intensive care, together with their potential benefits and limitations, are summarized in Table 1.

DISCUSSION

The present review highlights the rapidly expanding role of artificial intelligence (AI) in intensive care units (ICUs), demonstrating its potential to significantly enhance predictive accuracy, support clinical decision-making, and enable personalized therapeutic strategies. From a practical perspective, this review may be useful for clinicians, hospital decision-makers, and researchers seeking to understand not only where AI performs well in ICU settings, but also why successful implementation remains difficult. By combining clinical, technical, and ethical perspectives, the review aims to support a more realistic understanding of AI readiness in critical care. However, despite promising results across multiple domains, the translation of AI into routine clinical practice remains limited, reflecting a complex interplay of technical, clinical, ethical, and organizational challenges. Taken together, the findings of this review suggest that the future of intensive care will likely depend on the successful integration of data-driven technologies with human-centered clinical practice.

Clinical Impact of AI in Intensive Care

A consistent theme across the reviewed literature is that AI-based models may outperform traditional clinical tools under controlled study conditions; however, this apparent advantage does not automatically translate into routine clinical usefulness. Machine learning and deep learning models have demonstrated improved accuracy in predicting sepsis, mortality, and clinical deterioration (Nemati et al., 2018; Rajkomar et al., 2018; Tomašev et al., 2019).

These improvements can be attributed to the ability of AI systems to process large-scale, high-dimensional data and capture complex nonlinear relationships between clinical variables (Shickel et al., 2018). In contrast, traditional scoring systems such as SOFA are limited by their static structure and reliance on a small number of predefined parameters (Vincent et al., 1996).

Importantly, some studies have demonstrated not only improved predictive performance but also potential clinical benefits. The implementation of the TREWS system, for example, was associated with earlier treatment initiation and reduced mortality in patients with sepsis (Adams et al., 2022).

However, it should be noted that most studies remain retrospective, and evidence from prospective clinical trials is still limited. This raises important questions about the real-world effectiveness of AI systems.

From Prediction to Decision-Making and Automation

A key development in the field is the transition from predictive analytics to decision support and therapeutic optimization.

Reinforcement learning approaches, such as those proposed by Komorowski et al. (2018), represent a paradigm shift by enabling AI systems to recommend treatment strategies rather than simply predict outcomes. These models demonstrate the potential for dynamic, patient-specific decision-making, particularly in complex conditions such as sepsis.

Similarly, AI-based clinical decision support systems (CDSS) have shown promise in improving diagnostic accuracy and guiding treatment decisions (Sendak et al., 2020). Real-time systems such as TREWS illustrate the potential for integrating AI into clinical workflows.

Nevertheless, adoption remains limited, suggesting that the main barrier is no longer algorithm development itself but the ability to implement, validate, and govern these tools within real clinical environments.

The Role of Big Data in Enabling AI

The development of AI in intensive care is closely linked to the availability of large-scale clinical datasets.

Resources such as the MIMIC-III database (Johnson et al., 2016) have enabled significant advances in predictive modeling and reproducibility. Time-series analysis of high-frequency ICU data has further improved model performance (Ghassemi et al., 2015; Lipton et al., 2016).

However, the reliance on specific datasets raises concerns about generalizability. Models trained on single-center datasets may not perform well in different clinical environments, particularly those with different patient populations or healthcare systems (Wynants et al., 2020).

This highlights the need for multicenter datasets and standardized data formats to support robust model development.

Key Barriers to Implementation

Despite strong performance in research settings, several barriers limit the implementation of AI in clinical practice.

Lack of External Validation

Many models lack validation across diverse populations, reducing confidence in their reliability (Wynants et al., 2020).

Integration Challenges

AI systems are often not well integrated into clinical workflows, limiting usability and adoption (Sendak et al., 2020).

Interpretability and Trust

The "black box" nature of many AI models reduces clinician trust, particularly in high-risk environments (Tonekaboni et al., 2019).

Data Quality and Infrastructure

Fragmented data systems and inconsistent data quality further complicate implementation (Rajkomar et al., 2018).

These findings suggest that technical performance alone is insufficient for successful adoption. Instead, a holistic approach is required.

Ethical and Societal Implications

The integration of AI into ICU care raises important ethical considerations.

Algorithmic bias represents a major concern. Obermeyer et al. (2019) demonstrated that healthcare algorithms may disproportionately disadvantage certain patient groups, raising questions about fairness and equity.

Transparency and accountability are also critical issues. Clinicians must be able to understand and justify AI-driven decisions, particularly when they affect patient outcomes (Char et al., 2018).

In addition, the use of AI in critical care raises questions about patient autonomy, particularly in situations where patients are unable to participate in decision-making.

These ethical challenges underscore the need for robust governance frameworks and careful consideration of the societal impact of AI.

Human–AI Collaboration in Critical Care

Rather than replacing clinicians, AI is increasingly viewed as a tool for augmenting human expertise.

Topol (2019) introduced the concept of "high-performance medicine," in which AI systems and clinicians work collaboratively to improve patient care. In this model, AI provides data-driven insights, while clinicians contribute contextual understanding, ethical reasoning, and clinical judgment.

This collaborative approach is particularly important in ICU settings, where decisions often involve uncertainty and complex trade-offs.

The literature suggests that the most effective AI systems are those that support, rather than replace, clinician decision-making.

Future Perspectives and Research Priorities

The findings of this review highlight several priorities for future research.

First, prospective and multicenter validation studies are needed to establish the clinical effectiveness of AI systems.

Second, the development of explainable AI models is essential for improving transparency and clinician trust (Tonekaboni et al., 2019).

Third, greater emphasis should be placed on integration into clinical workflows, including user-centered design and interoperability.

Finally, ethical and regulatory frameworks must evolve to address the unique challenges posed by AI in healthcare.

Overall Synthesis

Overall, available evidence suggests that AI has the potential to transform intensive care by improving prediction, decision-making, and treatment optimization.

However, a key conclusion of this review is that **the main barriers to implementation are not technical but systemic**. Addressing issues related to validation, integration, ethics, and infrastructure will be critical for realizing the full potential of AI in critical care.

LIMITATIONS

This review has several limitations that should be considered when interpreting the findings.

First, as a narrative review, the study does not follow a fully systematic methodology. While this approach allows for a comprehensive and flexible synthesis of a rapidly evolving field, it may introduce selection bias and limit reproducibility. The absence of formal quantitative synthesis also precludes direct comparison of model performance across studies.

Second, the literature on artificial intelligence in intensive care is highly heterogeneous. Included studies differ in terms of patient populations, data sources, model architectures, and evaluation metrics, which complicates the interpretation and generalization of findings. This variability reflects the interdisciplinary nature of the field but also limits the ability to draw definitive conclusions.

Third, a substantial proportion of the available evidence is based on retrospective analyses, frequently using datasets such as the MIMIC-III database (Johnson et al., 2016). While these datasets have significantly advanced research, they may not fully represent real-world clinical variability. As noted by Wynants et al. (2020), lack of external validation remains a major limitation in many AI studies, raising concerns regarding generalizability.

Fourth, publication bias may influence the overall interpretation of results, as studies reporting positive outcomes are more likely to be published. This may lead to an overestimation of the effectiveness and readiness of AI systems for clinical implementation.

Finally, although ethical and societal aspects were included in this review, these dimensions remain underrepresented in the broader literature, which continues to focus predominantly on technical performance (Obermeyer et al., 2019; Char et al., 2018). This imbalance highlights an important gap that warrants further investigation.

Despite these limitations, the review provides a structured and critical overview of current evidence and identifies key challenges and directions for future research. In addition, because this review was designed as a narrative synthesis, the selection of studies may have been influenced by the authors' judgment regarding relevance and influence, which should be considered when interpreting the conclusions.

CONCLUSIONS

Artificial intelligence (AI) has emerged as a powerful and rapidly evolving tool in intensive care medicine, offering significant potential to improve patient outcomes, enhance clinical decision-making, and optimize resource utilization. This review demonstrates that AI applications in intensive care units (ICUs) are most advanced in the domains of predictive analytics, clinical decision support systems, and therapeutic optimization.

The analyzed literature consistently shows that AI-based models can outperform traditional clinical scoring systems in predicting sepsis, mortality, and clinical deterioration (Nemati et al., 2018; Rajkomar et al., 2018; Tomašev et al., 2019). In addition, reinforcement learning approaches highlight the potential for personalized, dynamic treatment strategies tailored to individual patient needs (Komorowski et al., 2018). These developments reflect a broader shift toward data-driven and precision medicine in critical care.

Despite these promising advances, the implementation of AI in routine clinical practice remains limited. As highlighted in this review, key barriers include lack of external validation, limited generalizability, challenges in integration with clinical workflows, and concerns regarding interpretability and trust (Wynants et al., 2020; Sendak et al., 2020; Tonekaboni et al., 2019). Furthermore, ethical issues such as algorithmic bias, transparency, and accountability continue to pose significant challenges (Obermeyer et al., 2019; Char et al., 2018).

A central conclusion of this review is that the successful adoption of AI in intensive care depends not only on technological performance but also on clinical, organizational, and societal factors. The concept of human–AI collaboration, rather than replacement, appears to be the most promising model for future integration (Topol, 2019).

Future research should focus on prospective and multicenter validation of AI models, development of explainable and transparent systems, and improved integration into clinical workflows. In parallel, there is a need for robust ethical and regulatory frameworks to ensure safe and equitable implementation.

In conclusion, artificial intelligence has clear potential to improve prediction, decision support, and therapeutic personalization in intensive care. However, the literature reviewed here suggests that clinical readiness still lags behind technical innovation. Future progress will depend not only on better algorithms, but also on prospective validation, explainability, workflow integration, and robust ethical oversight. For AI to become a safe and effective component of ICU practice, implementation science and governance must advance alongside model development.

CONFLICT OF INTEREST: The authors declare no conflicts of interest.

REFERENCES

1. Adams, R., Henry, K. E., Sridharan, A., Huang, J., Pattharanitima, P., Chou, J., … Saria, S. (2022). Prospective, multi-site study of patient outcomes after implementation of the TREWS machine learning-based sepsis early warning system. *Nature Medicine, 28*(7), 1455–1460. https://doi.org/10.1038/s41591-022-01894-0
2. Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing machine learning in health care—Addressing ethical challenges. *The New England Journal of Medicine, 378*(11), 981–983. https://doi.org/10.1056/NEJMp1714229
3. Churpek, M. M., Yuen, T. C., Winslow, C., Meltzer, D. O., Kattan, M. W., & 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), 368–374. https://doi.org/10.1097/CCM.0000000000001571
4. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., … Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine, 25*(1), 24–29. https://doi.org/10.1038/s41591-018-0316-z
5. Futoma, J., Hariharan, S., & Heller, K. (2017). Learning to detect sepsis with a multitask Gaussian process RNN classifier. *Proceedings of Machine Learning for Healthcare*, 1–17. https://doi.org/10.48550/arXiv.1706.04152
6. Ghassemi, M., Pimentel, M. A. F., Naumann, T., Brennan, T., Clifton, D. A., Szolovits, P., & Feng, M. (2015). A multivariate timeseries modeling approach to severity of illness assessment and forecasting in ICU with sparse, heterogeneous clinical data. *Proceedings of the AAAI Conference on Artificial Intelligence, 29*(1), 446–453. https://ojs.aaai.org/index.php/AAAI/article/view/9448
7. Harutyunyan, H., Khachatrian, H., Kale, D. C., Ver Steeg, G., & Galstyan, A. (2019). Multitask learning and benchmarking with clinical time series data. *Scientific Data, 6*, Article 96. https://doi.org/10.1038/s41597-019-0103-9
8. Henry, K. E., Hager, D. N., Pronovost, P. J., & Saria, S. (2015). A targeted real-time early warning score (TREWS) for septic shock. *Science Translational Medicine, 7*(299), Article 299ra122. https://doi.org/10.1126/scitranslmed.aab3719

9. Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L. W. H., Feng, M., Ghassemi, M., … Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data, 3*, Article 160035. https://doi.org/10.1038/sdata.2016.35
10. Johnson, A. E. W., Bulgarelli, L., Pollard, T. J., Horng, S., Celi, L. A., & Mark, R. G. (2021). MIMIC-IV, a freely accessible electronic health record dataset. *Scientific Data, 8*, Article 190. https://doi.org/10.1038/s41597-022-01899-x
11. Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., & Faisal, A. A. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine, 24*(11), 1716–1720. https://doi.org/10.1038/s41591-018-0213-5
12. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436–444. https://doi.org/10.1038/nature14539
13. Lipton, Z. C., Kale, D. C., Elkan, C., & Wetzel, R. (2016). *Learning to diagnose with LSTM recurrent neural networks* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.1511.03677
14. Moor, M., Horn, M., Rieck, B., Roqueiro, D., & Borgwardt, K. (2019). *Early recognition of sepsis with Gaussian process temporal convolutional networks and dynamic time warping* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.1902.01659
15. Nemati, S., Holder, A., Razmi, F., Stanley, M. D., Clifford, G. D., & Buchman, T. G. (2018). An interpretable machine learning model for accurate prediction of sepsis in the ICU. *Critical Care Medicine, 46*(4), 547–553. https://doi.org/10.1097/CCM.0000000000002936
16. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science, 366*(6464), 447–453. https://doi.org/10.1126/science.aax2342
17. Prasad, N., Cheng, L. F., Chivers, C., Draugelis, M., & Engelhardt, B. E. (2017). *A reinforcement learning approach to weaning of mechanical ventilation* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.1704.06300
18. Raghu, A., Komorowski, M., Ahmed, I., Celi, L. A., Szolovits, P., & Ghassemi, M. (2017). *Deep reinforcement learning for sepsis treatment* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.1711.09602
19. Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Liu, P. J., … Dean, J. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine, 1*, Article 18. https://doi.org/10.1038/s41746-018-0029-1
20. Sendak, M. P., D'Arcy, J., Kashyap, S., Gao, M., Nichols, M., Corey, K., & Ratliff, W. (2020). A path for translation of machine learning products into healthcare delivery. *European Medical Journal Innovations*, 1–9. https://doi.org/10.33590/EMJINNOV/19-00172
21. Shickel, B., Tighe, P. J., Bihorac, A., & Rashidi, P. (2018). Deep EHR: A survey of recent advances in deep learning techniques for electronic health record analysis. *IEEE Journal of Biomedical and Health Informatics, 22*(5), 1589–1604. https://doi.org/10.1109/JBHI.2017.2767063
22. Shillan, D., Sterne, J. A. C., Champneys, A., & Gibbison, B. (2019). Use of machine learning to analyse routinely collected intensive care unit data: A systematic review. *Critical Care, 23*, Article 284. https://doi.org/10.1186/s13054-019-2564-9
23. Singer, M., Deutschman, C. S., Seymour, C. W., Shankar-Hari, M., Annane, D., Bauer, M., & Angus, D. C. (2016). The third international consensus definitions for sepsis and septic shock (Sepsis-3). *JAMA, 315*(8), 801–810. https://doi.org/10.1001/jama.2016.0287
24. Tomašev, N., Glorot, X., Rae, J. W., Zielinski, M., Askham, H., Saraiva, A., … Hassabis, D. (2019). A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature, 572*(7767), 116–119. https://doi.org/10.1038/s41586-019-1390-1
25. Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). *What clinicians want: Contextualizing explainable machine learning for clinical end use* [Preprint]. arXiv. https://doi.org/10.48550/arXiv.1905.05134
26. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine, 25*(1), 44–56. https://doi.org/10.1038/s41591-018-0300-7
27. Vincent, J. L., Moreno, R., Takala, J., Willatts, S., De Mendonça, A., Bruining, H., … Thijs, L. G. (1996). The SOFA (Sepsis-related Organ Failure Assessment) score to describe organ dysfunction/failure. *Intensive Care Medicine, 22*(7), 707–710. https://doi.org/10.1007/BF01709751
28. Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., … Goldenberg, A. (2019). Do no harm: A roadmap for responsible machine learning for health care. *Nature Medicine, 25*(9), 1337–1340. https://doi.org/10.1038/s41591-019-0548-6
29. Wynants, L., Van Calster, B., Collins, G. S., Riley, R. D., Heinze, G., Schuit, E., … van Smeden, M. (2020). Prediction models for diagnosis and prognosis of COVID-19 infection: Systematic review and critical appraisal. *BMJ, 369*, Article m1328. https://doi.org/10.1136/bmj.m1328