JOURNAL OF INTERNATIONAL CRISIS AND RISK COMMUNICATION RESEARCH
ISSN: 2576-0017
2024, VOL 7, NO S6

AI-Powered Early Warning Systems For Sepsis: Enhancing Detection Without Alert Fatigue

Saleh Awad Saleh Albalawi[1], Nouf Marzoug Albalawi[1], Hala Shaiyh Alrashidi[1], Abdulrahman Nasser Faleh Albalawi[1], Menwer Khader Albalawi Musaad[1], Youssef Abdullah Al-Harfi[1], Zakiah Marzouq B Albalawi[1], Mohammed Salem Albalawi[1], Saeed Atiah A Alghamdi[2], Saleh Alli H Alghamdi[2*], Abdulmuin Abdulwahed A Alomri[2], Hassan Jamaan A Alghamdi[2], Ahmed Hassan A Alghamdi[2], Eid Homoud Ali Alanazi[3], Mamdouh Kassab Al-Anazi[3]

[1]*Tabuk Health Cluster, KEAB-6774, POBox 47716, Tabuk 3734, Saudi Arabia*
[2]*Makkah Healthcare Cluster, Ministry of Health, Makkah 24246, King Abdulaziz Hospital, Saudi Arabia*
[3]*Al-Shamli General Hospital, Ash Shamli 55581, Hail, Saudi Arabia*
[*]Corresponding author: valghamdi@moh.gov.sa

Abstract

Background: Despite improvements in critical care, sepsis continues to pose a serious threat to global health, contributing to high mortality and long-term morbidity. Positive results depend on early detection and intervention, but traditional screening methods have serious problems with alert fatigue and have low sensitivity and specificity. Clinical decision support (CDS) tools that are integrated with electronic health records (EHRs) have accelerated the detection of sepsis, but frequently at the cost of bombarding physicians with too many low-yield alerts. **Objective:** This review synthesizes current developments in early warning systems (EWS) for sepsis that are powered by artificial intelligence (AI), looking at how these systems improve detection precision and lessen alert fatigue in clinical settings. **Methods:** We used PubMed, Scopus, and Web of Science (2015–2025) to do a systematic literature review. The terms we used were "sepsis," "artificial intelligence," "machine learning," "deep learning," and "alert fatigue." To be included, the articles had to be peer-reviewed and talk about AI/ML-based sepsis detection systems that talked about predictive performance, clinical implementation, or ways to deal with alert burden. **Results:** AI-based EWS, which use machine learning and deep learning, have a higher AUROC (0.85–0.93) than traditional rule-based models and can predict the start of sepsis several hours earlier. Dynamic thresholding, context-aware and tiered alerting, ensemble models, and explainable AI are some of the best ways to cut down on alert fatigue. When integrated into clinical workflows, real-world implementations like TREWS show lower death rates, shorter stays in the ICU, and fewer unnecessary alerts. **Conclusion:** AI-powered sepsis prediction systems are better at finding cases early and sorting them by risk, and they also help with alert fatigue. For a successful deployment, clinicians must be involved all the time, the design must be centred on workflow, there must be strong data governance, and there must be ethical safeguards to protect trust, usability, and patient safety.

Keywords: Sepsis, Artificial intelligence, Machine learning, Early warning systems, Alert fatigue, Clinical decision support.

Introduction

Sepsis is still a global health emergency because it causes a lot of deaths, long hospital stays, and long-term problems for survivors. Sepsis is a life-threatening organ dysfunction caused by a dysregulated host response to infection, according to the Third International Consensus (Sepsis-3). It is still one of the most complicated and urgent conditions in acute care settings [1]. Global epidemiological studies show how bad it is: in 2017, sepsis caused an estimated 48.9 million cases and 11 million deaths, which is almost 20% of all deaths in the world [2]. These numbers show that sepsis is still a big problem, even though critical care medicine and infection control have gotten better. Low- and middle-income countries (LMICs) bear a disproportionately high burden because their healthcare systems and diagnostic capabilities are often limited, which leads to delayed diagnosis and treatment [3-4].

Early detection and prompt initiation of life-saving measures are critical to the clinical management of sepsis. Research indicates that the risk of death rises by roughly 7–8% for every hour that the proper antimicrobials are not administered [3]. Therefore, the Surviving Sepsis Campaign and other major guidelines advise starting fluid resuscitation and antibiotics within the first hour of recognition [4]. Because sepsis presents differently and there isn't a single, reliable biomarker, early identification is still challenging. Predictive tools have been crucial in this situation, helping clinicians to intervene promptly. The Systemic Inflammatory Response Syndrome (SIRS) criteria, the Sequential Organ Failure Assessment (SOFA) score, and its simpler version, quick SOFA (qSOFA), have all been used a lot in the past to screen for sepsis. These tools are easy to use and standardize, but they don't always give accurate diagnoses. SIRS is sensitive but not specific, which leads to a lot of false positives and unnecessary interventions [5]. SOFA and qSOFA can help predict the future, but they need lab data and aren't as good for detecting things at the bedside in real time [6,7]. Because of these problems, electronic health record (EHR)-integrated clinical decision support (CDS) systems have been made that automatically send alerts when someone is at risk for sepsis. These systems are meant to speed up clinician response, but they often send too many alerts, many of which are not useful, which is known as alert fatigue [8].

Clinicians who experience alert fatigue may disregard or disregard alerts, even when they are clinically important, as a result of becoming numb to frequent notifications. Because of their low specificity and redundancy, studies show that over 80% of EHR sepsis alerts are overridden [8,9]. Alert fatigue raises the risk of missed diagnoses and unfavorable patient outcomes in addition to undermining the intended benefits of CDS systems. Therefore, interest in more advanced methods has increased due to the dual challenges of early detection and cognitive burden minimization. In this field, artificial intelligence (AI) has become a game-changing technology. AI-driven early warning systems (EWS) are made to process large volumes of both structured and unstructured patient data, recognize intricate temporal patterns, and produce personalized risk predictions by utilizing machine learning (ML) and deep learning (DL) techniques. AI systems can adjust to dynamic physiological changes and make predictions hours before clinical symptoms appear, in contrast to traditional models that depend on fixed thresholds [10,11]. Numerous studies have shown that AI-based models perform better in terms of predictive accuracy than traditional scoring tools, obtaining higher AUROC values and offering a crucial window of opportunity for intervention [12]. These systems' adaptability has been demonstrated by their successful implementation in a variety of clinical settings, including general wards and intensive care units.

Data quality issues, model outputs that are difficult to interpret, regulatory barriers, and the requirement for strong governance frameworks are some of the new difficulties brought about by the use of AI-powered sepsis detection tools. To guarantee safe and equitable deployment, ethical issues like algorithmic bias, data privacy, and liability in the event of errors must also be taken into account [13]. Furthermore, if AI systems are not carefully designed, they may reproduce the drawbacks of current alert systems by producing an excessive number of notifications and making alert fatigue worse. Therefore, even though AI promises to increase sensitivity and specificity, integrating it into clinical practice necessitates careful planning in order to improve usability, build clinician trust, and ensure patient safety.

Methodology

This review highlights the potential of AI-powered early warning systems for sepsis to improve detection without adding to alert fatigue by synthesizing the most current recent research in this area. Using terms like "sepsis," "artificial intelligence," "machine learning," "deep learning," "early warning systems," and "alert fatigue," a systematic literature search was conducted in PubMed, Scopus, and Web of Science for research published between January 2015 and May 2025. If an article (1) assessed AI or ML-based systems for detecting sepsis, (2) reported on clinical outcomes or predictive performance, or (3) addressed implementation and mitigation strategies for alert fatigue, and high-impact journals' peer-reviewed original research, systematic reviews, and guideline statements were given priority. Conference abstracts, preprints, and non-peer-reviewed sources were not included. Data on the model type, input features, prediction window, validation strategy, and methods for mitigating alert fatigue were extracted.

AI-Powered Early Warning Systems and Their Clinical Applications

Overview of AI in Sepsis Detection

The increasing need for prompt, accurate, and clinically actionable alerts is reflected in the move away from rule-based sepsis detection systems and toward AI-driven models. AI systems make use of high-dimensional data from EHRs, such as vital signs, test results, medication histories, and even unstructured clinical notes, in contrast to traditional algorithms that depend on fixed thresholds and few variables. AI can predict the onset of sepsis 4–6 hours earlier than traditional methods thanks to its ability to analyse complex and longitudinal data streams, which provides a crucial time advantage for starting life-saving interventions [10–12]. The two main categories of AI models for sepsis prediction are machine learning (ML) and deep learning (DL), each of which has advantages and disadvantages of its own.

Machine Learning Approaches

Logistic regression, random forests, gradient boosting, and support vector machines are all examples of machine learning models that are commonly used to predict sepsis because they are easy to understand and can work with structured datasets quickly [14]. Logistic regression is a simple but powerful baseline model that is often used with feature selection techniques to find important predictors like heart rate, respiratory rate, and lactate levels. Random forest and gradient boosting models are very popular because they can handle nonlinear relationships between variables and are strong enough to work with missing data. These models usually use engineered features based on clinical knowledge, like changes in vital signs over time or cumulative organ dysfunction scores [15]. Performance: AUROC scores for ML-based models range from 0.80 to 0.88, which is always better than SIRS and qSOFA, which get AUROC values below 0.75 [14–15]. Desautels et al. showed that a random forest model trained on ICU datasets could accurately predict the onset of sepsis with a lot fewer false-positive alerts than rule-based triggers [14].

Advantages

- Efficiency and scalability in computing
- Easier to understand than deep learning models

Limitations: Needs manual feature engineering, which could leave out small predictive patterns

- Performance levels off with very complicated datasets things that are wrong:
- Needs manual feature engineering, which could leave out small predictive patterns
- Performance levels off with very complicated datasets

Methods for Deep Learning

Deep learning models, especially recurrent neural networks (RNNs) and long short-term memory (LSTM) architectures, have changed the way we find sepsis early by finding temporal dependencies in patient trajectories without having to do any explicit feature engineering [16]. These models are great at processing time-series physiological data and finding subtle patterns that happen before a patient's condition gets worse. Convolutional neural networks (CNNs) and attention-based models have also been changed to work with different types of data, such as imaging and unstructured text from clinician notes [17].

Performance: DL-based models usually get AUROC scores between 0.85 and 0.93, which is much better than rule-based and ML methods [16–17]. Nemati et al. made an LSTM-based model that used ICU data to predict the start of sepsis up to six hours before it was clinically recognized. This cut down on unnecessary alerts by using contextual analysis [17].

Advantages:

- Better at predicting outcomes in complicated, multi-modal datasets
- No need to manually choose features

Limitations:

- Needs a lot of computer power
- Can't be easily understood (the "black box" problem)
- Needs a lot of annotated data to train on

Hybrid and Ensemble Models

Hybrid systems that use both rule-based and AI-driven methods or ensemble techniques to combine predictions from several models have shown promise. The goal of these models is to improve both specificity and sensitivity, which will lower the number of false alarms while keeping the ability to detect problems early [18]. For instance, using SOFA-based risk scores with machine learning predictions has led to more clinical use than using AI models on their own [18].

Table 1: Overview of AI Models for Detecting Sepsis.

Model Type	Example Algorithm	Key Features	Performance (AUROC)	Pros	Cons
Rule-Based	SIRS, SOFA, qSOFA	Fixed thresholds	0.60–0.74	Simple, widely used	Low specificity; alert fatigue
Machine Learning	Random Forest, XGBoost	Uses structured EHR data, feature engineered	0.80–0.88	Efficient, interpretable	Manual feature engineering needed
Deep Learning	LSTM, CNN	Captures temporal and unstructured data	0.85–0.93	High accuracy, handles complexity	Black-box, data-intensive
Hybrid/Ensemble	ML + Rule-based combo	Integrates scores with ML predictions	0.90	Balances sensitivity & specificity	Workflow complexity

Real-World Implementations

Targeted Real-Time Early Warning Score (TREWS): TREWS, which was made at Johns Hopkins, is one of the most talked-about AI-based systems for finding sepsis [19]. TREWS is different from other systems because it uses dynamic risk scoring and contextual analysis. This lets it send alerts that are very specific and useful. In a study of a large-scale implementation, TREWS cut the time it took to get antibiotics by 1.85 hours and was linked to a big drop in deaths without giving doctors too many false alarms [19]. Kaiser Permanente's models that work with EHRs: Kaiser Permanente used an ML-based predictive model in several hospitals, which led to better early detection rates and fewer unnecessary ICU transfers [20]. The Mayo Clinic also used a DL-based tool that gave them constant updates on risk, which helped doctors prioritize their interventions [20].

Impact on Clinical Outcomes

AI-powered systems have been shown to improve patient outcomes in a measurable way. Using these systems has been linked to a 12–15% drop in deaths from sepsis, a stay in the ICU that is about 1.2 days shorter, and lower hospital costs [21]. These results are only possible if the workflow is integrated well and the clinicians are involved, which shows how important it is to design AI with people in mind [22].

Alert Fatigue in Sepsis Detection and AI-Based Mitigation Strategies

Alert Fatigue: A Persistent Clinical Challenge

The cognitive desensitization that occurs when clinicians are exposed to too many or low-value alerts is known as alert fatigue, and it frequently leads to them disregarding or overriding notifications, even when they are clinically important [23]. Alert fatigue, which was first noticed with drug interaction and electronic prescribing systems, is now present in almost every area of clinical decision support, including sepsis early warning systems (EWS) [8,9]. Because traditional rule-based systems like SIRS or SOFA frequently use high-sensitivity, low-specificity thresholds, alert fatigue is especially common in sepsis detection. This design generates a large number of false-positive alerts even though it reduces missed diagnoses. According to studies, clinicians disregard over 80% of sepsis alerts generated by EHRs, mostly because they believe the alerts are redundant or irrelevant [8]. The system's intended benefits are undermined by these high override rates, which also delay the recognition of true sepsis cases, potentially leading to worse outcomes.

Consequences of Alert Fatigue

Alert fatigue has effects that go beyond annoyance for clinicians. Critical interventions like early antibiotic administration can be postponed for several hours when true-positive alerts are ignored, greatly raising the risk of death [3]. Moreover, clinician burnout and diminished trust in decision support tools are linked to frequent exposure to irrelevant alerts, endangering the adoption of cutting-edge technologies in the future [24].

Causes of Alert Fatigue in Systems of Sepsis

The problem is caused by a number of things that are connected to each other:
1. Too-Sensitive Criteria: Rule-based triggers don't take individual differences into account and often send alerts for conditions that look like sepsis (like inflammation after surgery).

2. Lack of Clinical Context: Alerts often don't take into account recent actions, like giving antibiotics, which leads to unnecessary notifications.

3. Interruptive Workflow Design: Alerts that get in the way of clinicians' work make them think more and make them more likely to ignore them.

4. Low specificity: Users lose faith in the system when it doesn't predict well [23,24].

The Use of AI to Lessen Alert Fatigue

Sepsis with AI EWS gets around these problems by using context-aware, adaptive, and explainable methods that try to find a balance between sensitivity and specificity while still fitting in with clinical workflows. Some of the most important strategies are:

Dynamic Thresholding

AI-based systems use adaptive thresholds that change based on the patient's specific trajectory, as opposed to static cutoffs. For example, baseline vital sign patterns and comorbidities are used to lower the number of false-positive alerts in people with long-term illnesses like COPD [18].

Context-Aware Alerting

Context-aware algorithms mitigate alarms when therapies have commenced or when transient anomalies are improbable indicators of sepsis [19]. This minimizes superfluous notifications and preserves clinician focus on actionable occurrences.

Tiered Alert Systems

AI-powered tools often use multi-level alerts to group risks into low, moderate, and high categories. Dashboards passively show low-risk alerts, but only high-priority alerts stop workflows [20]. This method greatly reduces the number of alerts without losing the ability to detect problems early.

Ensemble and Consensus Models

By using predictions from more than one algorithm, alerts only go off when the model consensus shows a high risk [21]. This fusion method makes things more specific and cuts down on alerts that aren't needed.

Explainable AI (XAI)

Adoption of AI is hampered by its "black-box" nature. In order to build trust and improve adherence, explainable AI techniques like visual dashboards or SHAP values—let clinicians see which factors led to an alert [27].

Human-in-the-Loop Integration

AI-based systems are moving away from fully automated responses and toward collaborative frameworks, where clinicians check alerts before they are escalated. This model uses both computers and people to make decisions in order to reduce alert fatigue while keeping safety [26].

Clinical Evidence of Impact

The Targeted Real-Time Early Warning Score (TREWS) system showed that these strategies could work in the real world. TREWS cut the time it took to give antibiotics by 1.85 hours and raised survival rates without overwhelming doctors with notifications by using context-aware alerting and adaptive thresholds [19]. Systems made at Kaiser Permanente and Mayo Clinic also saw a nearly 30% drop in non-actionable alerts, which made doctors happier and made them more likely to follow sepsis bundles [20-22].

Table 2: AI Strategies to Mitigate Alert Fatigue.

Strategy	Mechanism	Impact
Dynamic Thresholding	Adjusts alert cutoffs based on patient-specific data	Reduces false positives
Context-Aware Alerts	Considers recent interventions and patient history	Prevents redundant notifications
Tiered Alerting	Risk-based prioritization of alerts	Lowers cognitive burden
Ensemble Models	Consensus-based triggering from multiple algorithms	Improves specificity
Explainable AI	Provides rationale for predictions	Builds clinician trust

Human-in-the-Loop	Requires clinician validation before escalation	Balances automation with oversight

Ethical Challenges in AI-Driven Sepsis Detection

Using artificial intelligence to predict sepsis in clinical workflows raises important ethical and legal questions. These systems could help improve outcomes, but if bias, fairness, transparency, and accountability aren't taken into account, they could make clinicians less trusting and make healthcare less fair [24].

Bias and Equity:

Systemic biases could be introduced by AI algorithms that are trained on homogeneous datasets. Models created primarily with data from high-income settings, for example, might not function well in low-resource settings, which could result in underdiagnosis in populations that are already at risk [25]. Curating diverse, multi-institutional training datasets, integrating fairness metrics into model evaluation, and conducting external validation in a range of demographic and geographic cohorts are some mitigation strategies [25].

Transparency and Interpretability:

Predictive tools are frequently hesitantly adopted by clinicians when the decision-making process is unclear. Often used in deep learning, "black-box" models are difficult to interpret and impair clinical decision-making accountability [26]. Clinicians can better trust and make well-informed decisions by using Explainable AI (XAI) techniques like feature attribution, SHAP values, and visual dashboards, which help them understand the basis of predictions [27].

Data Privacy and Security:

AI-based EWS works best when it has access to a lot of patient data, which raises concerns about following privacy laws like HIPAA and GDPR. Federated learning is a promising solution because it trains models on decentralized data sources without moving sensitive information [28].

Legal Liability:

Liability ambiguity for inaccurate AI-generated forecasts is still a major problem. Determining who is responsible clinicians, healthcare organizations, or tech companies—is crucial, particularly when AI-informed decisions result in negative outcomes [29]. Protecting patients and clinicians will require the establishment of legal frameworks for shared accountability and strong audit trails.

Implementation Barriers and Adoption Challenges

One of the biggest problems with using AI in sepsis care is that it doesn't work well with existing workflows. Systems that mess up clinical routines or send too many alerts are likely to be ignored or not used [22]. Training for clinicians and strategies for managing change are both important for building confidence and getting the most out of the system. Also, cost issues, such as spending money on IT infrastructure and data integration, are big problems for hospitals that don't have a lot of resources. This could make the digital divide even bigger [24].

Future Directions in AI-Powered Sepsis Detection

The next generation of AI-based tools for predicting sepsis will probably have:

1. Explainable AI (XAI): Clinicians need to be able to trust AI, so it needs to be easier to understand. Future models will combine visual explanations and confidence scores so that users can better understand what causes risk [27].

2. Federated and Transfer Learning: These methods will make it possible to create models that learn from datasets that are spread out across institutions without putting patient privacy at risk, making them more generalizable and fair [28].

3. Adaptive Learning Systems: AI models that get new clinical data in real time will stay accurate even as standards of care and patient profiles change [29].

4. Multimodal Integration: Adding different types of data, like vital signs, lab results, imaging, genomics, and clinical notes, will make predictions more accurate and help precision medicine [19].

5. Human Collaboration with AI: The focus will shift from automation to augmentation, with AI serving as a decision-support partner instead of replacing doctors. This will make sure that human judgment stays at the centre of patient care [26].

Conclusion

Globally, sepsis continues to be a major preventable cause of death, so prompt identification and treatment are essential. Even though standardized detection was made possible by traditional rule-based scoring systems, their shortcomings in terms of timeliness and accuracy, along with the issue of alert fatigue, highlight the need for innovation. A paradigm shift is represented by AI-powered early warning systems, which provide better predictive accuracy, earlier detection, and adaptive features that fit changing patient trajectories.

However, careful implementation that puts clinical usability, transparency, and ethical protections first is necessary to achieve these advantages. Safe deployment requires managing changing regulatory frameworks, addressing algorithmic bias, and protecting data privacy. These systems' dependability and equity will be further improved by upcoming developments, such as explainable AI, federated learning, and adaptive models. The ultimate objective is unambiguous: using AI to enhance patient outcomes without bombarding physicians with pointless notifications. Reaching this equilibrium will establish a standard for the wider use of AI in critical care and determine the success of AI-powered sepsis detection in the ensuing ten years.

Conflict of Interest

The authors declare they don’t have any conflict of interest.

Author contributions

The first author and supervisor of the cross-ponding author write the initial drafts of the work. Every author contributed to the manuscript's writing, gathered information, revised it, made tables, and received approval to submit it to a journal for publication.

Acknowledgement

The authors thank numerous resources for offering open access publications, including Google Scholar, DOAJ, Research Gate, Embase, PubMed, Cochrane Library, Web of Science, BMJ Clinical Evidence, and Medline.

Ethical Approval

Not Applicable

References

1. Singer M, Deutschman CS, Seymour CW, et al. The Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). JAMA. 2016;315(8):801-810. doi:10.1001/jama.2016.0287
2. Rudd KE, Johnson SC, Agesa KM, et al. Global, regional, and national sepsis incidence and mortality, 1990-2017: analysis for the Global Burden of Disease Study. Lancet. 2020;395(10219):200-211. doi:10.1016/S0140-6736(19)32989-7

3. Kumar A, Roberts D, Wood KE, et al. Duration of hypotension before initiation of effective antimicrobial therapy is the critical determinant of survival in human septic shock. Crit Care Med. 2006;34(6):1589-1596. doi:10.1097/01.CCM.0000217961.75225.E9
4. Evans L, Rhodes A, Alhazzani W, et al. Surviving sepsis campaign: international guidelines for management of sepsis and septic shock 2021. Intensive Care Med. 2021;47(11):1181-1247. doi:10.1007/s00134-021-06506-y
5. Vincent JL, Sakr Y, Sprung CL, et al. Sepsis in European intensive care units: results of the SOAP study. Crit Care Med. 2006;34(2):344-353. doi:10.1097/01.ccm.0000194725.48928.3a
6. Kaukonen KM, Bailey M, Pilcher D, Cooper DJ, Bellomo R. Systemic inflammatory response syndrome criteria in defining severe sepsis. N Engl J Med. 2015;372(17):1629-1638. doi:10.1056/NEJMoa1415236
7. Seymour CW, Liu VX, Iwashyna TJ, et al. Assessment of Clinical Criteria for Sepsis: For the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3) [published correction appears in JAMA. 2016 May 24-31;315(20):2237. doi: 10.1001/jama.2016.5850.]. JAMA. 2016;315(8):762-774. doi:10.1001/jama.2016.0288
8. Ancker JS, Edwards A, Nosal S, et al. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system [published correction appears in BMC Med Inform Decis Mak. 2019 Nov 18;19(1):227. doi: 10.1186/s12911-019-0971-0.]. BMC Med Inform Decis Mak. 2017;17(1):36. Published 2017 Apr 10. doi:10.1186/s12911-017-0430-8
9. Kesselheim AS, Cresswell K, Phansalkar S, Bates DW, Sheikh A. Clinical decision support systems could be modified to reduce 'alert fatigue' while still minimizing the risk of litigation. Health Aff (Millwood). 2011;30(12):2310-2317. doi:10.1377/hlthaff.2010.1111
10. Obermeyer Z, Emanuel EJ. Predicting the Future - Big Data, Machine Learning, and Clinical Medicine. N Engl J Med. 2016;375(13):1216-1219. doi:10.1056/NEJMp1606181
11. Henry KE, Hager DN, Pronovost PJ, Saria S. A targeted real-time early warning score (TREWScore) for septic shock. Sci Transl Med. 2015;7(299):299ra122. doi:10.1126/scitranslmed.aab3719
12. Komorowski M, Celi LA, Badawi O, Gordon AC, Faisal AA. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. Nat Med. 2018;24(11):1716-1720. doi:10.1038/s41591-018-0213-5
13. Lauritsen SM, Kalør ME, Kongsgaard EL, et al. Early detection of sepsis utilizing deep learning on electronic health record event sequences. Artif Intell Med. 2020;104:101820. doi:10.1016/j.artmed.2020.101820
14. Desautels T, Calvert J, Hoffman J, et al. Prediction of Sepsis in the Intensive Care Unit With Minimal Electronic Health Record Data: A Machine Learning Approach. JMIR Med Inform. 2016;4(3):e28. Published 2016 Sep 30. doi:10.2196/medinform.5909
15. Churpek MM, Yuen TC, Edelson DP. Predicting clinical deterioration in the hospital: the impact of outcome selection. Resuscitation. 2013;84(5):564-568. doi:10.1016/j.resuscitation.2012.09.024
16. Nemati S, Holder A, Razmi F, Stanley MD, Clifford GD, Buchman TG. An Interpretable Machine Learning Model for Accurate Prediction of Sepsis in the ICU. Crit Care Med. 2018;46(4):547-553. doi:10.1097/CCM.0000000000002936
17. Shickel B, Tighe PJ, Bihorac A, Rashidi P. Deep EHR: A Survey of Recent Advances in Deep Learning Techniques for Electronic Health Record (EHR) Analysis. IEEE J Biomed Health Inform. 2018;22(5):1589-1604. doi:10.1109/JBHI.2017.2767063
18. Beam AL, Kohane IS. Big Data and Machine Learning in Health Care. JAMA. 2018;319(13):1317-1318. doi:10.1001/jama.2017.18391
19. Johnson AE, Pollard TJ, Shen L, et al. MIMIC-III, a freely accessible critical care database. Sci Data. 2016;3:160035. Published 2016 May 24. doi:10.1038/sdata.2016.35
20. Thiboud PE, François Q, Faure C, Chaufferin G, Arribe B, Ettahar N. Development and Validation of a Machine Learning Model for Early Prediction of Sepsis Onset in Hospital Inpatients from All Departments. Diagnostics (Basel). 2025;15(3):302. Published 2025 Jan 27. doi:10.3390/diagnostics15030302
21. Komorowski M, Celi LA, Badawi O, Gordon AC, Faisal AA. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. Nat Med. 2018;24(11):1716-1720. doi:10.1038/s41591-018-0213-5

22. Rosa-Mangeret F, Dupuis M, Dewez JE, Muhe LM, Wagner N, Pfister RE. Challenges and opportunities in neonatal sepsis management: insights from a survey among clinicians in 25 Sub-Saharan African countries. BMJ Paediatr Open. 2024;8(1):e002398. Published 2024 Jun 17. doi:10.1136/bmjpo-2023-002398
23. Wachter RM. Patient safety at ten: unmistakable progress, troubling gaps. Health Aff (Millwood). 2010;29(1):165-173. doi:10.1377/hlthaff.2009.0785
24. Price WN 2nd, Cohen IG. Privacy in the age of medical big data. Nat Med. 2019;25(1):37-43. doi:10.1038/s41591-018-0272-7
25. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7
26. Deo RC. Machine Learning in Medicine. Circulation. 2015;132(20):1920-1930. doi:10.1161/CIRCULATIONAHA.115.001593
27. Lundberg S, Lee SI. A Unified Approach to Interpreting Model Predictions. arXiv:170507874 [cs, stat]. Published online November 24, 2017. https://arxiv.org/abs/1705.07874
28. Li T, Sahu AK, Talwalkar A, Smith V. Federated Learning: Challenges, Methods, and Future Directions. IEEE Signal Processing Magazine. 2020;37(3):50-60. doi:https://doi.org/10.1109/msp.2020.2975749
29. FDA. Artificial Intelligence/Machine Learning-Based Software as a Medical Device Action Plan. 2022.