

# Human-Centered Sepsis Management in Clinical Work Systems: A Socio-Technical AI Framework for Patient Safety

Firda Rahmadani<sup>1</sup>, Mecit Can Emre Simsekler<sup>1</sup>, and Siddiq Anwar<sup>2</sup>

<sup>1</sup>Department of Management Science & Engineering, Khalifa University of Science & Technology, Abu Dhabi, UAE

<sup>2</sup>Department of Nephrology, King's College Hospital London, Dubai, UAE

## ABSTRACT

Sepsis is a time-critical condition associated with substantial morbidity and mortality, where delays in recognition and treatment markedly worsen outcomes. Although machine learning models show promise for early detection, their clinical impact has been constrained by poor integration into workflows, limited interpretability, and insufficient support for coordinated action. This study introduces a human-centered, systems-based agentic AI architecture for sepsis risk modeling and proactive clinical management. Rather than generating static risk scores, the system continuously interprets evolving patient data, situates risk within the clinical workflow, and supports timely, clinician-supervised interventions. Grounded in systems engineering and guided by the Systems Engineering Initiative for Patient Safety (SEIPS) framework, the architecture embeds predictive intelligence within the broader socio-technical work system, enabling closed-loop monitoring, coordination of safety-critical tasks, and feedback-driven adaptation. By reframing sepsis prediction as an adaptive, workflow-aware safety intervention, this approach advances AI from passive decision support toward an accountable, action-oriented partner in care delivery while preserving clinician oversight.

**Keywords:** Human-AI interaction, Patient safety, Human factors, Risk management, Sepsis, Systems thinking, Agentic AI

## INTRODUCTION

Sepsis continues to be a major contributor to global death, posing an ongoing threat to patient safety within healthcare systems worldwide (Rudd et al., 2020). Its clinical trajectory is highly time-sensitive, with delays in detection and management strongly associated with increased mortality, prolonged hospital stays, and higher resource utilization (Ashqar et al., 2025). Despite advances in clinical knowledge and standardized treatment bundles, early identification of sepsis remains inconsistent, reflecting variability in patient presentation and differences in clinical workflows across care settings (Draeger et al., 2025).

Traditional approaches to sepsis detection have relied heavily on rule-based screening tools and early warning scores, such as threshold-based vital sign criteria or aggregated risk indices (Al-Juhani et al., 2025). While these tools provide some standardization, they often lack sensitivity and specificity in

real-world settings and may fail to account for temporal trends (Papareddy et al., 2025). As a result, clinicians may experience frequent false alarms or delayed alerts, undermining trust in automated systems and limiting their effectiveness in supporting timely clinical decision-making (Ayvaci et al., 2026).

Importantly, sepsis is not solely a clinical diagnostic challenge but a socio-technical systems problem. Recognition and response depend on complex interactions among clinicians, information systems, organizational policies, workload pressures, and communication pathways (Rahmadani et al., 2025). Failures in sepsis care often arise not from a lack of medical knowledge but from breakdowns in coordination, situational awareness, and system-level responsiveness. Addressing sepsis-related harm, therefore, requires solutions that extend beyond clinical prediction to encompass workflow integration, human factors, and system resilience (Taylor & Kowalkowski, 2024).

In recent years, there has been a growing awareness in applying machine learning models to sepsis risk prediction, with many studies demonstrating improved predictive performance compared to traditional scoring systems. However, most existing AI-based approaches remain limited by their static and observational nature, generating risk scores at discrete time points without meaningful integration into clinical workflows. These models often function as passive decision-support tools, leaving clinicians to interpret and act on predictions amid competing demands and time pressures (Abbas et al., 2025).

In practice, the deployment of AI-generated risk scores has frequently contributed to alert fatigue and inconsistent adoption (Rahmadani & Simsekler, 2024). Poor interpretability, lack of transparency regarding contributing factors, and insufficient alignment with clinical roles can reduce clinician trust and engagement (Abdelwanis et al., 2026). Moreover, risk predictions are rarely accompanied by actionable guidance or system-level support, diminishing their potential impact on patient outcomes (Giddings et al., 2024). Conventional AI models typically operate in an open-loop manner: they predict risk but do not monitor downstream actions, adapt to clinician responses, or learn from system-level outcomes (Oei et al., 2025). Without mechanisms for feedback, escalation, or coordination of care processes, these models cannot compensate for delays, omissions, or workflow disruptions. As a result, high predictive accuracy alone has not translated reliably into improved sepsis care or patient safety (O'Reilly et al., 2024).

To meaningfully improve sepsis outcomes, AI systems must move beyond prediction toward active participation within the clinical system. This requires AI that not only estimates risk but also understands context, supports timely action, and adapts to evolving patient and system states while preserving human authority over clinical decisions (Persson et al., 2021). Agentic AI offers a promising paradigm for achieving this shift. Agentic AI is conceptualized as a goal-oriented, human-supervised system that operationalizes sepsis risk prediction through continuous monitoring, contextual reasoning, and bounded clinical support (Sapkota et al., 2026). Rather than replacing clinician judgment, the agent functions as a safety partner by translating predictive insights into workflow-aware prompts, coordinating care processes, and supporting shared situational awareness (Li et al., 2025). By embedding intelligence within the care delivery system,

agentic AI enables a closed-loop approach to sepsis safety that aligns with principles of systems engineering and resilience (Choi & Yoo, 2026).

To overcome these challenges, this study aims to develop and evaluate a human-centered, systems-based agentic AI framework for sepsis risk modeling. Specifically, we aim to (1) integrate machine learning-based sepsis risk prediction within an agentic architecture, (2) apply human factors and systems engineering principles to ensure safe and effective interaction with clinicians, and (3) enhance patient safety, resilience, and quality of care through proactive, context-aware clinical support. This work demonstrates how agentic AI can bridge the gap between predictive analytics and real-world clinical impact in sepsis care.

## **CONCEPTUAL FRAMEWORK: AGENTIC AI FOR SEPSIS SAFETY**

### **Definition of Agentic AI in Healthcare**

Agentic AI in healthcare is a goal-oriented, context-aware, and human-supervised artificial intelligence system that operates over time to support clinical objectives in complex care environments (Hinojosa Fuentes et al., 2025). Unlike conventional predictive models that generate isolated risk estimates, agentic AI continuously perceives patient data, reasons about evolving clinical and system states, and takes bounded actions to pursue safety and quality goals (Acharya et al., 2025). In the context of sepsis care, this includes maintaining situational awareness of patient deterioration, care processes, and workflow constraints while adapting one's behaviour in response to clinician actions and patient outcomes.

A key distinction exists between predictive AI and agentic AI. Predictive AI addresses the question of what might happen by estimating the likelihood of adverse events, such as the onset of sepsis. In contrast, agentic AI addresses how the system should respond over time by translating predictions into context-sensitive support actions, monitoring downstream effects, and adjusting interventions as conditions change. This temporal, interactive capability enables agentic AI to function as an active component of the clinical system rather than a passive source of information.

### **Human-AI Shared Control Model**

The proposed framework is grounded in a human–AI shared control model, in which clinicians retain full authority over diagnosis, treatment, and care decisions, while AI agents provide continuous assistance, coordination, and monitoring. Within this model, agentic AI supports clinicians by synthesizing complex data streams, highlighting emerging risks, and facilitating the timely execution of evidence-based care processes without superseding human judgment.

Agent autonomy is explicitly bounded through predefined action spaces, safety constraints, and oversight mechanisms. Agents may generate alerts, recommend actions, track care bundle completion, and provide explanations, but they do not independently initiate treatment or alter clinical orders. All agent outputs are transparent, reviewable, and interruptible, allowing clinicians to accept, modify, or disregard recommendations. This shared

control approach is designed to mitigate risks such as automation bias and alert fatigue while fostering trust, accountability, and effective collaboration between humans and AI systems.

### **Systems Engineering Perspective**

The Systems Engineering Initiative for Patient Safety (SEIPS) is a human-centered systems engineering framework that explains patient safety as the result of interactions within a complex healthcare work system (Carayon et al., 2020). SEIPS conceptualizes this system as comprising five interrelated elements: people (e.g., clinicians and patients), tasks, tools and technologies, the physical environment, and organizational factors. Rather than attributing safety events to individual failures, the SEIPS perspective focuses on how poorly designed systems, misaligned workflows, or inadequate support structures contribute to errors and inefficiencies. By emphasizing the design, evaluation, and continuous improvement of the overall work system, SEIPS provides a structured approach to identifying hazards, improving care processes, and enhancing patient safety and quality of care (Park et al., 2026).

From a systems engineering standpoint, sepsis care constitutes a dynamic, high-risk socio-technical system characterized by uncertainty, time pressure, and integration between human decision-making and technological infrastructure. Adverse outcomes often arise from delayed recognition, coordination failures, or system bottlenecks rather than isolated clinical errors. Addressing these challenges requires solutions that support not only accurate prediction but also reliable system performance under variable conditions.

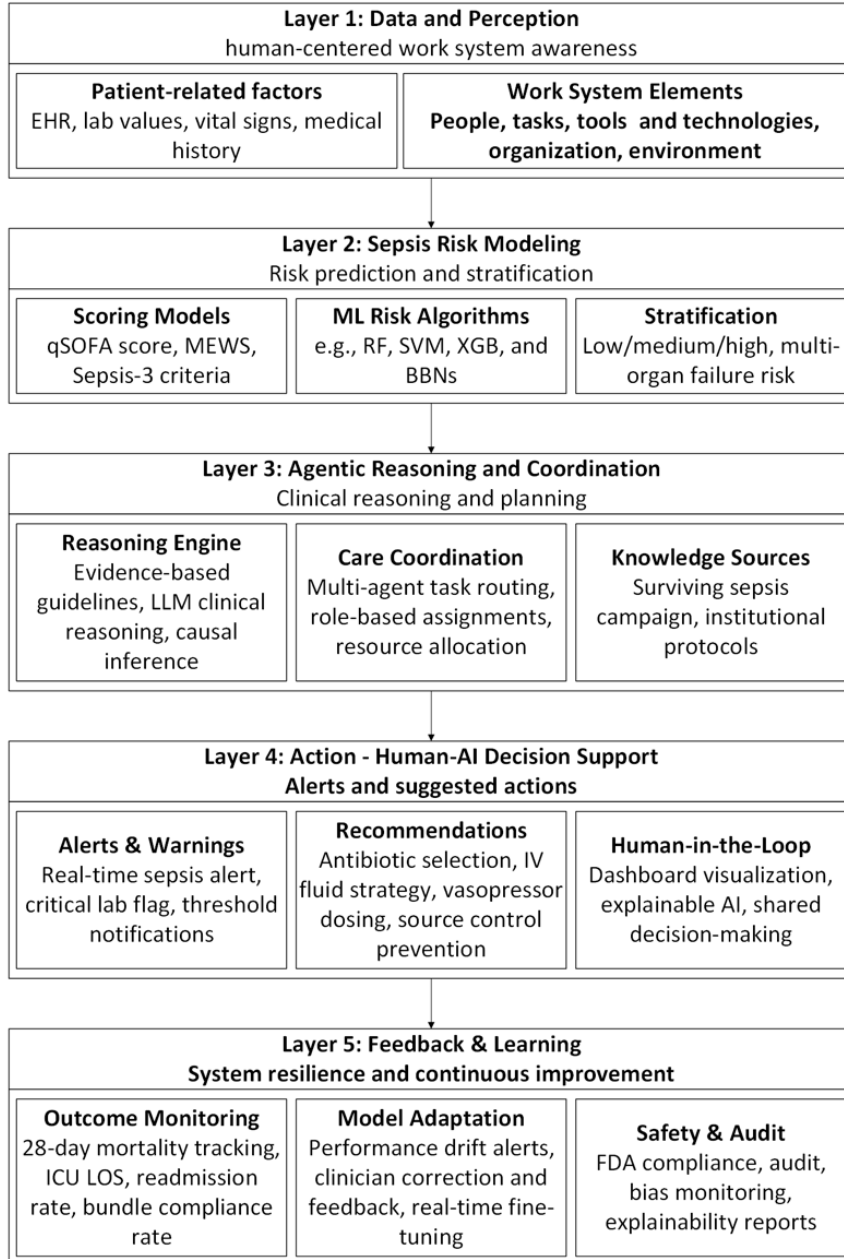
Within this framework, agentic AI functions as a closed-loop safety controller embedded in the sepsis care process. The agent continuously monitors patient status and system responses, supports timely intervention, and assesses whether safety-critical actions have occurred. By incorporating feedback and adapting to evolving conditions, agentic AI contributes to system resilience by helping the healthcare system anticipate, respond to, and recover from disruptions. This approach aligns with principles of resilience engineering and contemporary patient safety theory, reframing AI from a predictive tool into an active participant in maintaining safe and effective care delivery.

## **AGENTIC AI SYSTEM ARCHITECTURE**

### **Overview of the SEIPS-Aligned Architecture**

The proposed agentic AI system architecture is designed as an embedded component of the SEIPS work system, integrating people, tasks, tools, and technologies, organizational context, and care environments to support safe and timely sepsis care. Rather than operating as an isolated analytic module, the architecture positions agentic AI as a coordinating layer that links predictive intelligence with clinical workflows and system-level feedback. At a high level, the architecture consists of four interacting layers: (1) data and perception, (2) sepsis risk modeling, (3) agentic reasoning and coordination, (4) action, and (5) feedback and learning. These layers collectively support a

closed-loop safety function that aligns with SEIPS' emphasis on care processes as the mediators between work system design and patient outcomes. The agentic AI architecture is shown in Figure 1.



**Figure 1:** AI agent for human-centered, systems-based risk modeling for sepsis.

Table 1 shows the SEIPS components mapped into the agentic AI architecture in sepsis safety context.

**Table 1:** Mapping of SEIPS components to the agentic AI architecture for sepsis management.

SEIPS Component	Definition	Agentic AI Role	Operational in Sepsis Care
People	Clinicians, patients, care teams	Supports clinician situational awareness and decision-making without replacing human judgment	Risk explanations tailored to nurses vs physicians; role-specific alerts
Tasks	Clinical and cognitive activities required for care	Monitors and coordinates safety-critical sepsis tasks over time	Tracking screening, lactate measurement, antibiotic timing, fluid resuscitation
Tools and technologies	Technology that supports the detection and care process	Act as an intelligent, adaptive tool embedded in the EHR	Risk modeling, alerts, task reminders, interfaces
Organization	Policies, protocols, staffing, and escalation pathways	Align actions with institutional sepsis protocols and workflows	Escalation based on unit policies, alignment with sepsis bundles
Environment	Care setting and physical context	Adapts behavior to care environment constraints	The ward layout to accommodate the sepsis care process

### Data and Perception Layer: Work System Awareness

The data and perception layer provides the agentic AI with continuous awareness of the clinical work system. This layer takes real-time and near-real-time inputs from the electronic health record, including physiological observations, laboratory data, medication administration information, clinical notes, and the patient's location within the care setting (e.g., emergency department, ward, or ICU).

From a SEIPS perspective, this layer captures signals across multiple domains:

- Person-related data, reflecting patient physiology and clinician documentation;
- Task-related data, such as timing and completion of sepsis screening and bundle elements;
- Environmental and organizational context, including care setting, shift changes, and escalation pathways.

By integrating heterogeneous data streams, the perception layer enables the agent to detect not only physiological deterioration but also emerging system-level risks, such as delays in care or breakdowns in task execution.

### Sepsis Risk Modeling Layer

The sepsis risk modeling layer serves as the core predictive component of the architecture. Using machine learning techniques applied to longitudinal patient data, this layer generates dynamic estimates of sepsis risk that evolve

over time rather than relying on single-time-point assessments. Risk outputs include both probabilistic estimates and trajectories indicating whether a patient's risk is increasing, stable, or decreasing.

Within the SEIPS framework, this layer supports care processes by enhancing clinicians' ability to anticipate patient deterioration. However, consistent with the limitations of conventional AI models, this layer alone does not determine clinical action. Instead, it functions as an input to the agentic reasoning layer, where predictions are interpreted in light of human, task, and organizational factors.

### **Agentic Reasoning and Coordination Layer**

The agentic reasoning layer represents the defining element of the proposed architecture. This layer integrates sepsis risk predictions with contextual information about the work system to determine when and how to support clinical action. Reasoning processes account for temporal risk trends, recent clinician responses, workload considerations, and organizational protocols, enabling the agent to adapt its behaviour to evolving conditions.

From a SEIPS standpoint, this layer facilitates coordination across people and tasks by maintaining situational awareness of who is responsible for care, which safety-critical tasks are pending, and where bottlenecks may arise. The agent does not replace human coordination but augments it by tracking task status, identifying delays, and supporting shared awareness among team members. This capability is particularly important in sepsis care, where time pressure and distributed responsibility increase the likelihood of system-level failures.

### **Action Layer: Human-Centered Decision Support**

The action layer translates agentic reasoning into human-centered interventions that align with clinical workflows and safety principles. Actions are intentionally limited to supportive functions, including notifications based on sepsis risk and urgency, reminders related to sepsis screening or bundle completion, and suggesting next steps aligned with evidence-based guidelines.

In alignment with SEIPS, these actions are designed to support human performance. Alerts and recommendations are tailored to clinical roles and settings, minimizing unnecessary interruptions and cognitive overload. Clinicians retain full authority to accept, modify, or ignore agent outputs, and all actions are transparent and auditable. This design reinforces trust and accountability while mitigating risks, including alert fatigue and automation bias.

### **Feedback and Learning Layer: Supporting System Resilience**

The final architectural component is a feedback and learning layer that monitors downstream effects of agent actions and clinician responses. This layer captures data on alert acknowledgment, task completion, delays, and patient outcomes, enabling retrospective analysis of system performance.

Consistent with SEIPS and resilience engineering principles, this feedback supports organizational learning rather than autonomous adaptation. Insights

derived from agent performance inform human-led quality improvement efforts, such as refining alert thresholds, adjusting escalation pathways, or redesigning workflows. By closing the loop between prediction, action, and outcome, the architecture supports continuous improvement of sepsis care processes and reinforces the resilience of the healthcare system.

## DISCUSSION

This study employed a systems-based evaluation of an agentic AI framework for sepsis risk modeling, grounded in the SEIPS model of work system design. The framework was evaluated using retrospective and/or prospective clinical data from adult patients admitted to acute care settings, including the ED, inpatient wards, and ICUs. The study design focused on assessing both predictive performance and system-level impacts on sepsis care processes and patient safety.

System design was guided explicitly by the SEIPS framework to ensure alignment with socio-technical principles. Stakeholder analysis was conducted to identify key clinical roles, safety-critical tasks, and organizational constraints relevant to sepsis care. These findings informed the configuration of agentic AI functions, including alert thresholds, escalation pathways, and role-specific interfaces. The agentic AI was implemented as a non-authoritative clinical decision support layer integrated with the electronic health record. All agent actions were restricted to supportive functions, consistent with SEIPS principles that emphasize human-centered system design and clinician control.

A central contribution of this work is the reframing of sepsis prediction from a static, model-centric task to an adaptive, workflow-aware safety intervention. Traditional machine learning-based sepsis tools have often been evaluated primarily on calibration metrics, with limited consideration of how predictions are translated into action within complex clinical environments. By embedding the risk model within an agentic AI architecture, this study demonstrates how continuous interpretation of evolving patient data can be operationalized as a series of context-sensitive prompts, reminders, and coordination aligned with existing work practices. In doing so, the framework helps to bridge the persistent gap between algorithmic performance and meaningful improvements in patient safety.

The systems-based perspective also highlighted that the value of agentic AI is tightly integrated within the broader socio-technical work system. The SEIPS analysis underscored that sepsis care depends on coordinated activities across multiple actors: nurses, physicians, pharmacists, operating under time pressure and resource constraints. By explicitly mapping these interdependencies, the study positioned the agent not as an isolated “intelligent” component but as a supporting actor that facilitates information flow, prioritization, and task execution. For example, the configuration of role-specific interfaces and escalation pathways was driven by observed communication patterns and handoff practices, rather than by technical considerations alone. This alignment likely contributed to the perceived acceptability and potential sustainability of the system in practice.

Importantly, the findings also strengthen the idea that maintaining clinician authority and accountability is not merely a regulatory or ethical requirement but a design feature that shapes how the agent is used and trusted. By constraining the agentic AI to non-authoritative actions, such as surfacing risk, suggesting next steps, and tracking completion of safety-critical tasks, the framework supports a human-centered model of automation. This design choice mitigates risks of overreliance and automation bias while still enabling the system to act proactively through closed-loop monitoring and follow-up. The ability of clinicians to override, ignore, or adapt recommendations serves as a dynamic feedback signal that can inform ongoing tuning of alert logic, thresholds, and workflows.

From a patient safety perspective, the agentic AI framework can be interpreted as a layered defense against delays in sepsis recognition and treatment. At the end of care, it assists frontline clinicians in detecting subtle deteriorations, recalling time-sensitive bundle elements, and coordinating actions across disciplines. It generates process-level data on timeliness, adherence, and variation to identify system weaknesses and inform organizational learning. This dual function aligns with contemporary safety paradigms that emphasize harm prevention and the support of resilient performance, positioning agentic AI as a mechanism for strengthening resilience in the face of clinical uncertainty and workload variability.

This study highlights several challenges that must be addressed before broader implementation. The effectiveness of the agentic framework depends on the quality, availability, and timeliness of EHR data, as missing, delayed, or inaccurate documentation can degrade risk estimates and recommended actions. In addition, the evaluation was conducted within specific organizational contexts, limiting generalizability across institutions, workflows, and resource settings. Successful adaptation will require local stakeholder engagement, model recalibration, and configuration of agent behaviors to align with site-specific policies and resources.

Although the system aims to reduce cognitive burden, poorly calibrated behaviors may contribute to alert fatigue, underscoring the need for continuous monitoring of alert volume, override rates, and workflow effects. Clinician feedback on trust, and usability should be treated as core outcomes alongside performance metrics to ensure sustained engagement and alignment with evolving clinical practice.

## **CONCLUSION AND FUTURE WORK**

By embedding agentic AI within the SEIPS work system, this architecture reframes sepsis risk modeling as a system-level safety intervention, aligning human-centered design, workflow coordination, and feedback to support adaptive care delivery. This integration helps ensure that gains in model performance translate into meaningful improvements in patient safety and system resilience.

Future work should focus on a deeper evaluation of human–AI interaction in sepsis care, examining how clinicians interpret, respond to, and learn from agentic recommendations in real time. Mixed-methods studies that combine

system logs, observation, and clinician feedback can inform refinements to autonomy, interaction design, and interpretability, thereby enabling a clearer understanding of why risk is identified and which data drive recommendations.

More broadly, this framework may serve as a template for AI-enabled safety interventions in other acute conditions. It advances a shift from static prediction tools toward adaptive, workflow-aware AI partners, a shift that will require sustained multidisciplinary collaboration and rigorous, multi-site evaluation to ensure safe, equitable, and scalable impact.

## ACKNOWLEDGMENT

This research was funded by Khalifa University of Science and Technology through the Research & Innovation Grant Program under Project ID: KU-INT-RIG-2025-8471000044.

## REFERENCES

- Abbas, G. H., Sen, P., Giri, O. A., & Khan, N. H. (2025). Artificial Intelligence-Based Predictive Modeling for Early Detection of Sepsis in Hospitalized Patients: A Systematic Review and Meta-Analysis. *Critical Care Explorations*, 7(12), e1360. <https://doi.org/10.1097/CCE.0000000000001360>
- Abdelwanis, M., Simsekler, M. C. E., Gabor, A. F., Sleptchenko, A., & Omar, M. (2026). Artificial intelligence adoption challenges from healthcare providers' perspectives: A comprehensive review and future directions. *Safety Science*, 193, 107028. <https://doi.org/10.1016/j.ssci.2025.107028>
- Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey. *IEEE Access*, 13, 18912–18936. <https://doi.org/10.1109/ACCESS.2025.3532853>
- Al-Juhani, A., Desoky, R., Iskander, Z., Alshehri, K. T., Alshehri, A. A., Almuhaimeid, A., Alharbi, N. L., Mominah, R. H., Al-humoud, F. M., & Desoky, A. (2025). Advances in Data-Driven Early Warning Systems for Sepsis Recognition and Intervention in Emergency Care: A Systematic Review of Diagnostic Performance and Clinical Outcomes. *Cureus*. <https://doi.org/10.7759/cureus.89882>
- Ashqar, E., Alkaissi, A., Ateeq, O., & Shawahna, R. (2025). Factors predicting discharge outcomes of sepsis patients admitted to intensive care unit in a major tertiary care hospital: A retrospective study from Palestine. *PLOS Global Public Health*, 5(12), e0005643. <https://doi.org/10.1371/journal.pgph.0005643>
- Ayvaci, M. U. S., Mobini, Z., & Özer, Ö. (2026). It Takes Two to Make It Right: How Nurses' Response to Sepsis Alerts Impacts Physicians' Process Compliance. *Manufacturing & Service Operations Management*, 28(1), 20–41. <https://doi.org/10.1287/msom.2022.0242>
- Carayon, P., Wooldridge, A., Hoonakker, P., Hundt, A. S., & Kelly, M. M. (2020). SEIPS 3.0: Human-centered design of the patient journey for patient safety. *Applied Ergonomics*, 84, 103033. <https://doi.org/10.1016/j.apergo.2019.103033>
- Choi, J. Y., & Yoo, T. K. (2026). Transforming clinical medicine with multimodal artificial intelligence, agentic systems, and the model-context protocol: A perspective on future directions. *Discover Health Systems*, 5(1), 6. <https://doi.org/10.1007/s44250-026-00343-w>
- Draeger, L., Fleischmann-Struzek, C., Gehrke-Beck, S., Heintze, C., Thomas-Rueddel, D. O., & Schmidt, K. (2025). Barriers and facilitators to optimal sepsis care – a systematized review of healthcare professionals' perspectives. *BMC Health Services Research*, 25(1), 591. <https://doi.org/10.1186/s12913-025-12777-8>

- Giddings, R., Joseph, A., Callender, T., Janes, S. M., Van Der Schaar, M., Sheringham, J., & Navani, N. (2024). Factors influencing clinician and patient interaction with machine learning-based risk prediction models: A systematic review. *The Lancet Digital Health*, 6(2), e131–e144. [https://doi.org/10.1016/S2589-7500\(23\)00241-8](https://doi.org/10.1016/S2589-7500(23)00241-8)
- Hinostroza Fuentes, V. G., Karim, H. A., Tan, M. J. T., & AlDahoul, N. (2025). AI with agency: A vision for adaptive, efficient, and ethical healthcare. *Frontiers in Digital Health*, 7, 1600216. <https://doi.org/10.3389/fdgth.2025.1600216>
- Li, H., Cheng, X., & Zhang, X. (2025). Accurate Insights, Trustworthy Interactions: Designing a Collaborative AI-Human Multi-Agent System with Knowledge Graph for Diagnosis Prediction. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–15. <https://doi.org/10.1145/3706598.3713526>
- Oei, S. P., Bakkes, T. H. G. F., Mischi, M., Bouwman, R. A., Van Sloun, R. J. G., & Turco, S. (2025). Artificial intelligence in clinical decision support and the prediction of adverse events. *Frontiers in Digital Health*, 7, 1403047. <https://doi.org/10.3389/fdgth.2025.1403047>
- O'Reilly, D., McGrath, J., & Martin-Loeches, I. (2024). Optimizing artificial intelligence in sepsis management: Opportunities in the present and looking closely to the future. *Journal of Intensive Medicine*, 4(1), 34–45. <https://doi.org/10.1016/j.jointm.2023.10.001>
- Papareddy, P., Lobo, T. J., Holub, M., Bouma, H., Maca, J., Strodthoff, N., & Herwald, H. (2025). Transforming sepsis management: AI-driven innovations in early detection and tailored therapies. *Critical Care*, 29(1), 366. <https://doi.org/10.1186/s13054-025-05588-0>
- Park, Y.-E., Ock, M., Lee, J.-H., Ko, D.-H., Lee, H.-J., Park, T., Yoo, J., & Lee, Y. (2026). Assessing Health Care Professionals' Perceptions of a New System in Clinical Workflows: Systems Engineering Initiative for Patient Safety–Based Consensual Qualitative Research. *Journal of Medical Internet Research*, 28, e86166. <https://doi.org/10.2196/86166>
- Persson, I., Östling, A., Arlbrandt, M., Söderberg, J., & Becedas, D. (2021). A Machine Learning Sepsis Prediction Algorithm for Intended Intensive Care Unit Use (NAVVOY Sepsis): Proof-of-Concept Study. *JMIR Formative Research*, 5(9), e28000. <https://doi.org/10.2196/28000>
- Rahmadani, F., & Simsekler, M. C. E. (2024). A Path Towards Human-AI Decision-Making in Sepsis Care through Human-Centered Systems-Based Design Approach. *2024 IEEE International Conference on Technology Management, Operations and Decisions (ICTMOD)*, 1–6. <https://doi.org/10.1109/ICTMOD63116.2024.10878191>
- Rahmadani, F., Simsekler, M. C. E., & Kim, N. (2025). Human-AI Decision Support for Enhanced Healthcare Management: The Case of A New Sepsis Prediction and Treatment System. *IEEE Engineering Management Review*, 1–37. <https://doi.org/10.1109/EMR.2025.3639504>
- Rudd, K. E., Johnson, S. C., Agesa, K. M., Shackelford, K. A., Tsoi, D., Kievlan, D. R., Colombara, D. V., Ikuta, K. S., Kissoon, N., Finfer, S., Fleischmann-Struzek, C., Machado, F. R., Reinhart, K. K., Rowan, K., Seymour, C. W., Watson, R. S., West, T. E., Marinho, F., Hay, S. I., ... Naghavi, M. (2020). Global, regional, and national sepsis incidence and mortality, 1990–2017: Analysis for the Global Burden of Disease Study. *The Lancet*, 395(10219), 200–211. [https://doi.org/10.1016/S0140-6736\(19\)32989-7](https://doi.org/10.1016/S0140-6736(19)32989-7)
- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2026). AI Agents vs. Agentic AI: A Conceptual taxonomy, applications and challenges. *Information Fusion*, 126, 103599. <https://doi.org/10.1016/j.inffus.2025.103599>
- Taylor, S. P., & Kowalkowski, M. (2024). Failure to Rescue as a Quality Measure in Sepsis. *JAMA*, 332(1), 11. <https://doi.org/10.1001/jama.2024.6771>