Journal of Translational Critical Care Medicine **2026** vol. 8(2) | e00022
DOI: 10.14218/JTCCM.2025.00022

Review Article

Generative Artificial Intelligence in Critical Care Medicine: A Narrative Review of Applications, Predictive Analytics, Documentation, and Ethical Imperatives

Wisit Cheungpasitporn[1], Charat Thongprayoon[1] and Kianoush Kashani[1,2*]

[1]*Division of Nephrology and Hypertension, Department of Medicine, Mayo Clinic, Rochester, MN, USA;* [2]*Division of Pulmonary and Critical Care Medicine, Department of Medicine, Mayo Clinic, Rochester, MN, USA*

Received: September 02, 2025 | **Revised:** September 21, 2025 | **Accepted:** October 28, 2025 | **Published online:** June 26, 2026

Abstract

Generative artificial intelligence (AI), particularly large language models (LLMs) and multimodal systems, is emerging as a potentially important innovation in intensive care medicine. The intensive care unit (ICU) is a data-dense, high-acuity setting where rapid and accurate decisions are critical. These models can translate complex multimodal data into interpretable and clinically actionable insights across diagnostic, prognostic, and documentation workflows. This review outlines six key domains in which generative AI is currently being explored for its potential to reshape critical care: clinical decision support; clinical documentation automation (AI scribe, voice-to-note); predictive analytics, including sepsis and acute respiratory distress syndrome prediction, acute kidney injury management, ventilator liberation readiness, delirium monitoring, and continuous renal replacement therapy optimization; ICU data summarization and multimodal monitoring; synthetic data generation; and legal and ethical governance. In clinical decision support, hybrid models that integrate time-series monitoring data with LLMs can contextualize alerts, generate diagnostic suggestions, and offer treatment plans with explainable reasoning. Documentation tools that leverage ambient listening and voice-to-note AI can streamline progress notes and discharge summaries, thereby reducing clinician workload. In predictive analytics, LLMs enhance model performance by augmenting sparse electronic health record data and translating outputs into interpretable narratives. Synthetic data generation enables algorithm development and training, particularly for rare events, while protecting patient privacy. However, the realism and ethical deployment of such data require rigorous validation. Widespread implementation of generative AI will require careful attention to challenges related to trust, validation, bias, liability, and regulatory compliance. The use of these tools must remain under clinician supervision to ensure transparency and accountability. With responsible deployment, generative AI may augment ICU workflows, improve outcomes, and reduce clinician burden, potentially becoming an indispensable component of critical care delivery.

Introduction

Generative artificial intelligence (AI), especially large language models (LLMs) and multimodal foundation models, is rapidly emerging as a promising tool with the potential to reshape aspects of intensive care medicine.[1,2] The intensive care unit (ICU) presents a uniquely challenging environment defined by its constant high acuity, overwhelming data volume, and the need for rapid, high-consequence clinical decisions.[3] In this setting, generative AI offers capabilities that extend far beyond conventional data analysis. These systems can integrate diverse data streams, including clinical narratives, laboratory values, continuous physiologic monitoring, imaging studies, and even genomic information.[4] The outputs they generate are human-readable and clinically actionable, providing diagnostic suggestions, prognostic assessments, treatment recommendations, and structured documentation that support real-time care delivery.[5–8]

Advances in computational capabilities have accelerated the incorporation of generative AI into emerging clinical workflows.[3,9–11] However, as these systems demonstrate early utility, they also intro-

---

Keywords: Generative artificial intelligence; Large language models (LLMs); Intensive care medicine; ICU; Clinical decision support systems (CDSS); Electronic health record automation (EHR automation); Predictive analytics; Early warning systems; Multimodal data integration; Synthetic data generation; AI ethics and governance.

[*]**Correspondence to**: Kianoush Kashani, Division of Pulmonary and Critical Care Medicine, Mayo Clinic, Rochester, MN 59005, USA. ORCID: https://orcid.org/0000-0003-2184-3683. Tel: +1-507-266-7960, E-mail: kashani.kianoush@mayo.edu

How to cite this article: Cheungpasitporn W, Thongprayoon C, Kashani K. Generative Artificial Intelligence in Critical Care Medicine: A Narrative Review of Applications, Predictive Analytics, Documentation, and Ethical Imperatives. *J Transl Crit Care Med* 2026;8(2):e00022. doi: 10.14218/JTCCM.2025.00022.

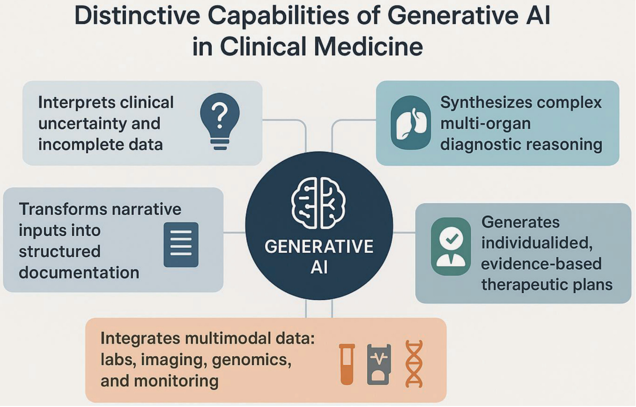

Fig. 1. Distinctive capabilities of generative AI in clinical medicine and critical care. Illustrates unique functions of LLMs including multimodal data integration, chain-of-thought reasoning, predictive analytics, and clinical documentation automation. AI, artificial intelligence; LLM, large language model.

duce challenges that require careful governance.[12,13] Key considerations include safety, bias, transparency, regulatory compliance, and optimal models of human–AI collaboration.[14–18] This review summarizes the current applications of generative AI in intensive care across multiple domains, including clinical decision support, knowledge management, education, documentation automation, predictive analytics, synthetic data generation, and legal and ethical oversight.[6] This synthesis draws on evolving evidence and initial clinical implementation to provide a comprehensive perspective on integrating generative AI into critical care practice.

Generative AI and LLMs in medicine

Generative AI represents a significant leap forward in the broader evolution of AI.[19,20] Earlier phases focused on symbolic reasoning and deep learning.[3,9,10] The current era, driven by foundation models such as GPT-4o (OpenAI), Claude 3.5 Sonnet (Anthropic), Gemini 1.5 Pro (Google DeepMind), and Meta Llama 3, introduces models capable of few-shot learning, chain-of-thought reasoning, and seamless multimodal integration.[20] These models combine large-scale natural language processing, unsupervised pretraining, reinforcement learning from human feedback, and retrieval-augmented generation (RAG) to perform complex clinical reasoning and generate meaningful outputs directly from unstructured data.[21,22]

In clinical medicine, several features distinguish generative AI from earlier approaches (Fig. 1).[19,20] These include: 1) the ability to interpret ambiguous or incomplete clinical scenarios that often challenge rule-based systems, 2) generation of structured documentation directly from narrative inputs, 3) assistance with complex diagnostic reasoning across multiple organ systems, 4) production of individualized, evidence-informed therapeutic recommendations, and 5) integration of multimodal data sources, such as text, laboratory data, imaging, physiologic monitoring, and genomics, into coherent clinical insights.

Advanced prompt engineering strategies for critical care applications

Effective clinical use of LLMs in critical care depends on domain-specific prompting strategies that optimize output quality, safety, and clinical relevance.[23] Several prompting frameworks are critical for ICU applications (Table 1). These include: (1) zero-shot and few-shot prompting generate structured clinical summaries without prior fine-tuning (Supplementary Table 1); (2) chain-of-thought prompting supports stepwise diagnostic reasoning, such as in hypoxemia evaluation, shock classification, or acid-base interpretation; (3) role prompting simulates expert consultant recommendations tailored to specific subspecialties such as intensivist, infectious disease, or nephrology input; (4) multimodal prompting enables the interpretation of integrated data sources, including laboratory data, physiologic waveforms, ventilator parameters, and imaging findings; and (5) context-enriched prompting incorporates longitudinal patient trajectories, ventilator settings, hemodynamic trends, sedation levels, and patient-specific goals of care.

These advanced prompting strategies provide essential support that enables generative AI to operate effectively in the complex, data-intensive ICU environment while also helping to minimize clinical risk.

Potential applications of generative AI in critical care

The following sections highlight key domains where generative AI may enhance critical care delivery. These applications encompass both bedside clinical decision-making and broader system-level functions, showcasing the unique ability of generative models to integrate diverse data sources, synthesize evolving information, and support high-acuity care. For each domain, we summarize current capabilities, emerging use cases, and potential clinical implications within the intensive care setting.

Clinical decision support in the ICU

Clinical decision support is one of the most transformative applications of generative AI in the ICU. Traditional rule-based systems rely on predefined algorithms and limited structured data inputs.[24,25] In contrast, generative AI models, particularly LLMs, are capable of processing unstructured and semi-structured clinical data, including free-text progress notes, laboratory trends, consultant narratives, imaging reports, and longitudinal patient records.[25] These models synthesize heterogeneous information sources to generate patient-specific, context-aware clinical recommendations that are both interpretable and actionable. In this role, generative AI serves as a form of AI medical consultant, augmenting clinical reasoning by integrating multimodal data and producing diagnostic suggestions, risk stratifications, and treatment considerations in complex or uncertain cases.

Several studies have demonstrated that advanced generative AI models, such as GPT-4, can match or exceed clinician-level performance when applied to complex diagnostic vignettes, particularly in rare, atypical, or diagnostically challenging scenarios.[24,26,27] This capability goes beyond isolated data retrieval, enabling integrated clinical synthesis that informs both diagnostic formulation and management decisions.

A key advancement enabled by generative AI is the integration of multimodal data streams (Table 2). Generative multimodal AI systems have the potential to support a broad range of critical care applications by integrating diverse structured and unstructured data sources. For instance, in sepsis detection, these models may combine vital signs, laboratory values, medication orders, nursing notes, and clinical documentation using RAG to produce individualized risk assessments and accompanying narrative explanations.[9,28] By contrast, proposed applications such as ventila-

Table 1. Prompting techniques and example prompts for generative AI applications in intensive care medicine

Prompting technique	Description	Example of prompt
Zero-shot prompting	Provides instructions without examples or prior context	"List possible causes of acute kidney injury in a 65-year-old ICU patient with hypotension and rising creatinine"
Few-shot prompting	Uses a few examples to guide model reasoning	"Given these cases of prerenal AKI and ATN, classify the following patient's AKI phenotype based on labs and volume status"
Chain-of-thought prompting	Enables stepwise logical reasoning	"Step through the differential for a patient with metabolic acidosis: pH 7.22, bicarbonate 12, AG 20, lactate 5"
Role prompting	Assigns an expert identity to the model	"You are a nephrologist. Recommend CRRT prescription for a hemodynamically unstable patient with hyperkalemia and volume overload"
Multimodal prompting	Integrates data beyond text (labs, imaging, waveforms)	"Analyze this patient's ultrasound, CVP, and MAP trends to suggest fluid management during CRRT"
Context-enriched prompting	Incorporates longitudinal and real-time patient data	"Given 24-hour labs, hourly urine output, and vasopressor doses, recommend CRRT adjustments for ongoing hemodynamic instability"
Guideline-constrained prompting	Embeds clinical protocols into outputs	"Provide hyperkalemia treatment following KDIGO guidelines for a patient with K+ 6.8 and ECG changes"
Safety-constrained prompting	Applies safety boundaries to outputs	"Avoid correcting serum sodium by more than 8 mmol/L per 24 hours in hyponatremia management"

AG, anion gap; AKI, acute kidney injury; ATN, acute tubular necrosis; CRRT, continuous renal replacement therapy; CVP, central venous pressure; ECG, electrocardiogram; ICU, intensive care unit; KDIGO, Kidney Disease: Improving Global Outcomes; MAP, mean arterial pressure.

tor weaning, delirium monitoring, and acute respiratory distress syndrome (ARDS) phenotyping remain conceptual only, with no clinical trial validation to date. These concepts envision multimodal systems, synthesizing ventilator waveforms, chest imaging, arterial blood gas measurements, progress notes, and nursing documentation to generate readiness assessments, delirium status reports, or ARDS subphenotype classifications. For acute kidney injury, generative AI models may integrate laboratory data, urine output, medication exposures, hemodynamic variables, and imaging findings to generate individualized risk estimates and suggest management options. In the evaluation of code status and goals of care, LLMs may synthesize prior clinical documentation, advance directives, and family discussions to generate patient-centered summaries. Collectively, these generative AI systems represent a potential advance in real-time, data-driven clinical reasoning, offering the possibility of enhancing diagnostic accuracy, improving care consistency, and reducing cognitive workload in the complex ICU environment.[9,28]

Table 2. Integration of generative multimodal AI mechanisms across key critical care scenarios with clinical translation

Clinical scenario	Generative AI multimodal mechanism
Sepsis detection	Multimodal generative AI integrates structured data (vital signs, labs, orders) with unstructured data (nursing notes, progress notes) using LLMs with RAG. The model generates patient-specific risk scores combined with narrative explanations (e.g., "Sepsis risk elevated due to rising lactate and hypotension")
Ventilator weaning	Multimodal generative AI integrates ventilator waveform data, radiographic imaging (CXR), and blood gases, combined with textual clinical narratives. Generative AI produces weaning-readiness summaries, such as "Patient meeting ventilator liberation criteria; minimal secretions; ABG normalized"
Delirium monitoring	Multimodal generative AI integrates audio inputs (patient voice), nursing documentation, medication records, and clinical notes. The system synthesizes real-time narrative assessments of delirium presence, subtype, and potential contributing factors to assist in clinical management and handoff documentation
AKI prediction and management	Multimodal generative AI integrates labs (creatinine, electrolytes), urine output, medications, hemodynamics, and imaging (ultrasound). Generative AI produces both predictive risk scores and actionable management narratives ("Recommend reduction in nephrotoxic exposure and fluid re-evaluation")
Code status and goals-of-care assessment	LLMs with RAG can synthesize prior notes, advance directives, family discussions, and demographics to auto-generate goals-of-care summaries ("Patient previously expressed preference for comfort-oriented care if mechanical ventilation required")
ARDS detection and phenotyping	Multimodal generative AI integrates SpO_2 waveform analysis, ABG trends, chest X-rays (processed via vision transformers), and ventilator data. The generative system classifies ARDS subphenotypes and produces management suggestions ("ARDS with hyperinflammatory phenotype; recommend conservative fluid management and prone positioning")

AI, artificial intelligence; ABG, arterial blood gas; AKI, acute kidney injury; ARDS, acute respiratory distress syndrome; CXR, chest radiograph; LLM, large language model; RAG, retrieval-augmented generation; SpO_2, peripheral capillary oxygen saturation.

Knowledge management and evidence synthesis in critical care

In parallel with patient-specific decision support, generative AI also enables advanced knowledge management functions that support clinical reasoning, research, and hypothesis generation in the ICU. Tools such as OpenEvidence employ RAG architectures to synthesize large bodies of biomedical literature into focused, evidence-based summaries relevant to specific clinical questions.[29] These systems integrate structured databases (e.g., PubMed, clinical trial registries, guideline repositories) with unstructured text sources, allowing clinicians to rapidly access up-to-date evidence while formulating differential diagnoses, evaluating complex cases, or considering novel therapeutic approaches.[29]

In the context of rare or diagnostically challenging ICU presentations, knowledge management platforms powered by generative AI can dynamically generate literature syntheses that highlight potential etiologies, diagnostic strategies, and therapeutic options drawn from current medical knowledge.[20] Moreover, these systems facilitate hypothesis generation for clinical research by identifying knowledge gaps, aggregating related studies, and summarizing emerging evidence across diverse domains. Such AI-augmented knowledge curation reduces the time burden of manual literature review and supports continuous learning at the point of care, particularly in rapidly evolving fields like critical care.

The integration of patient-specific data synthesis and global knowledge management represents a dual capability of generative AI, simultaneously enhancing bedside decision-making while informing clinician education and research development within the ICU environment.

Education and training

Generative AI offers novel capabilities to support education, training, and cognitive skill development for intensive care providers.[30,31] In contrast to traditional didactic learning modalities, AI-enabled educational tools offer dynamic, personalized, and context-aware learning experiences tailored to the evolving needs of ICU clinicians and trainees.[20]

One application involves automated curriculum development and educational content generation.[30] Generative AI models can synthesize recent literature and clinical guidelines to develop up-to-date learning modules, journal club summaries, and continuing education materials for critical care fellows and staff. These AI-generated resources facilitate the rapid dissemination of evolving knowledge, enabling clinicians to stay current with emerging evidence while minimizing the time burden of manual literature review.[31]

Generative models also support scenario-based training through the creation of AI-generated clinical vignettes and simulation cases. By utilizing structured and unstructured ICU data, AI systems can produce diverse, high-fidelity patient scenarios that replicate complex clinical situations for trainees to manage.[31] These synthetic cases allow for the safe practice of diagnostic reasoning, therapeutic decision-making, and team-based communication in a controlled learning environment. Furthermore, adaptive learning platforms may use real-time performance feedback to tailor the difficulty, complexity, and instructional focus based on each learner's progress, thereby optimizing educational outcomes. In addition, generative AI-powered virtual teaching assistants may support asynchronous learning by answering trainee questions, providing evidence-based clarifications, and guiding clinical reasoning exercises. This AI-augmented mentorship can complement traditional faculty supervision, particularly in resource-limited or high-volume training environments.

Documentation and electronic health record (EHR) automation

Clinical documentation remains a substantial contributor to clinician workload in the ICU.[32] The complexity of ICU care requires frequent documentation of progress notes, multidisciplinary care plans, handoffs, and discharge summaries. Generative AI offers the potential to reduce this burden by automating aspects of documentation and streamlining data synthesis.[33–35]

In outpatient and general inpatient settings, ambient AI scribe technologies, such as Nuance Dragon Ambient eXperience and Abridge,[33–35] have demonstrated the ability to transcribe clinician-patient conversations and automatically populate structured fields in the electronic health record (EHR). These systems allow clinicians to focus on patient care while improving note completeness and timeliness. Early reports suggest that such tools can reduce administrative workload and improve provider satisfaction.[33–35]

However, their application in the ICU is more limited. Many critically ill patients are sedated, mechanically ventilated, or otherwise unable to participate in verbal interactions.[32] As a result, the use of real-time ambient voice capture is inherently restricted in this setting.[36] Nevertheless, generative AI remains applicable to other components of ICU documentation (Supplementary Table 2). Models can synthesize structured and unstructured data sources, including laboratory values, physiologic monitoring, imaging reports, and consult notes, to generate daily progress notes, multi-day ICU summaries, and discharge documentation.[9] Automated generation of structured handoff reports, such as SBAR (Situation, Background, Assessment, Recommendation) and I-PASS templates, may also improve transitions of care. In addition, generative AI models may assist with clinical coding by extracting relevant diagnoses, procedures, and billing elements from clinical documentation, potentially improving revenue capture and reducing the need for coder queries. Summaries can also be tailored; for example, an AI might generate a focused handoff note highlighting changes in the last 12 hours or a patient and family communication note translating medical jargon into lay language.[9] Importantly, these AI-generated outputs remain subject to clinician review and final approval.

Generative AI for EHR integration and critical care workflow optimization

In critical care, the application of generative AI extends beyond documentation into active clinical workflow integration.[13] LLM-embedded EHR interfaces enable clinicians to query complex patient data in real time using natural language, facilitating rapid clinical decision-making.

One principal application is context-aware data retrieval. An intensivist may, for example, query: "*Has this patient developed hypotension over the past 24 hours?*" The LLM analyzes structured and unstructured data, including nursing flowsheets, physiologic monitoring trends, laboratory results, ventilator settings, and clinical narratives, to generate a targeted response. This represents a notable departure from traditional keyword-based search methods to multimodal, conversational data synthesis. Such tools may be particularly valuable in time-sensitive scenarios, including acute hemodynamic deterioration or pre-rounding clinical assessments, where the rapid integration of heterogeneous data is essential (Supplementary Fig. 1).

In parallel, generative AI is being evaluated for prospective task automation within intensive care workflows. Ambient listening systems may capture multidisciplinary team discussions (e.g., "*If hypotension persists, initiate vasopressin and repeat chest radiography in the morning*"), with the AI autonomously generating corresponding draft orders, imaging requests, and documentation for clinician review. This represents an evolution from passive transcription toward actionable documentation, in which expressed clinical intent is translated into structured orders under physician oversight. While still investigational, these systems have the potential to reduce clinician cognitive load, streamline task execution, and enhance workflow efficiency in high-acuity environments.[9,36]

Generative AI in sepsis prediction and early warning systems

Predictive analytics have long been employed in ICUs to anticipate clinical deterioration, using scoring systems and machine learning algorithms to identify early signs of sepsis, shock, respiratory failure, and other critical events.[37] Generative AI is advancing these applications through two primary mechanisms: 1) by enabling models that learn from complex, high-dimensional data to generate individualized risk estimates and 2) by enhancing interpretability through narrative-based, context-aware outputs that augment clinical decision-making.[19,20]

An emerging application involves real-time physiologic monitoring systems that not only detect abnormal trends but also translate these into actionable clinical narratives. These platforms integrate streaming data from bedside monitors, EHRs, and laboratory systems to generate timely, interpretable alerts that facilitate early intervention. For example, generative AI systems can continuously analyze temporal trends in vital signs, laboratory results, and medication administrations to identify subtle patterns indicative of impending clinical deterioration. Rather than issuing generic threshold-based alarms, these systems provide narrative, context-sensitive recommendations, such as "*Progressive increase in lactate levels and heart rate since 04:00; suggest evaluating for evolving shock state*."

Unlike conventional alert systems that rely on fixed thresholds and generate generic alarms, these narrative, context-sensitive recommendations provide clinicians with an interpretable synthesis of complex data streams, facilitating earlier evaluation and targeted intervention. The integration of predictive analytics with generative AI-enabled interpretability represents a significant advance in the development of next-generation decision support systems for critical care.[38,39]

Specific high-risk conditions such as sepsis and ARDS have become major targets for generative AI-based predictive modeling.[40,41] Generative models are capable of capturing the complex temporal interplay among physiologic signals, laboratory markers, and clinical context that often precede the onset of sepsis. For example, diffusion-based generative models have been applied to anomaly detection in time-series data by reconstructing expected physiologic trajectories and identifying subtle deviations that may signal early clinical deterioration.[40] In parallel, LLMs may augment conventional risk scores by synthesizing quantitative predictions with clinical narratives. An LLM may receive output from a sepsis risk model and generate a context-enriched alert: "*The model predicts high sepsis risk in 12 hours; note that the patient's wound has been oozing and WBC count is trending upward, consistent with early infection*" (Fig. 2). This hybrid framework, which combines predictive modeling with generative narrative synthesis, may offer clinicians more actionable and interpretable insights than numerical risk estimates alone, thereby enhancing situational awareness and supporting timely clinical intervention.

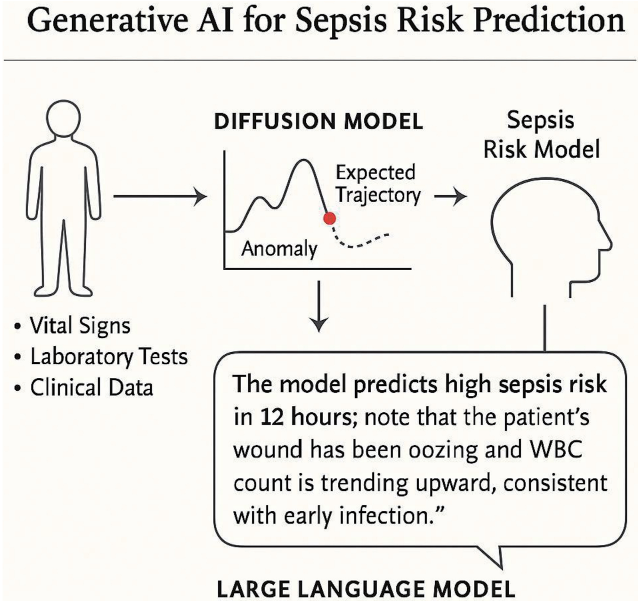

Fig. 2. Hybrid LLM and predictive analytics system for early sepsis detection and ICU outcome prediction with narrative synthesis. Demonstrates integration of structured physiologic data, laboratory values, and clinical notes for sepsis risk prediction and interpretability. AI, artificial intelligence; ICU, intensive care unit; LLM, large language model; WBC, white blood cell.

Real-world barriers to implementation

Despite promising applications, the widespread deployment of generative AI in ICUs faces several real-world barriers. Technically, integration into fragmented EHR systems remains challenging, with interoperability gaps, limited infrastructure for real-time multimodal data streaming, and concerns about reliability during downtime or system errors. From the human perspective, clinician acceptance is influenced by trust in AI outputs, workload implications, and medico-legal accountability. Without adequate training and demonstration of clear clinical value, adoption may remain limited. Financially, the costs of procuring, maintaining, and updating AI systems, particularly those requiring secure high-performance computing environments, can be prohibitive for many institutions. Reimbursement pathways for AI-augmented care are still evolving, which further complicates sustainability. Together, these barriers highlight the need for careful pilot testing, transparent cost-benefit evaluations, and robust implementation science frameworks before large-scale adoption.

Generative AI for data augmentation and predictive model optimization in critical care

An emerging application of generative AI in critical care involves augmenting the underlying data inputs used for predictive modeling. Traditional machine learning models that rely on EHR data are frequently constrained by data sparsity, irregular documentation intervals, and missing information, particularly among underserved populations with limited clinical documentation. LLMs offer novel capabilities to address these challenges by generating synthetic clinical content that enhances longitudinal data continuity and improves predictive model performance.[42]

One approach employs LLMs to impute missing data across temporal gaps in patient records. In many intensive care settings, formal documentation occurs intermittently, leaving extended periods without explicit narrative updates on clinical status. Recent studies have demonstrated that prompting LLMs to generate plausible interim progress notes, based on prior clinical context, can construct more continuous patient trajectories.[43] Incorporating these synthetic narratives into one-year mortality prediction models has been associated with improved model discrimination. Importantly, the inclusion of generated text narrowed performance gaps between patients with extensive versus sparse documentation, suggesting that such generative augmentation may mitigate biases arising from differential documentation density and improve fairness across patient subgroups.[43]

In addition to temporal imputation, LLMs have been applied as feature engineering tools to distill complex unstructured narratives into concise, prognostically relevant summaries.[44] In one investigation of ICU readmission and length-of-stay prediction,[43] daily LLM-generated summaries capturing key complications and clinical events were used as structured inputs for downstream predictive models.[43] The incorporation of these generative features resulted in substantial improvement in both discrimination and precision-recall performance, outperforming models trained solely on conventional structured EHR data. In this capacity, generative models convert unstructured documentation into clinically salient features that may be otherwise inaccessible through standard preprocessing methods.

These data augmentation strategies represent a distinct paradigm in critical care data science. Rather than functioning solely as documentation tools or real-time decision aids, generative models actively reconstruct incomplete patient trajectories, thereby improving both the completeness and equity of predictive analytics.[20] This evolving role positions generative AI as a critical component in dataset optimization, model calibration, and bias reduction, particularly as predictive algorithms are increasingly deployed to inform high-stakes clinical decision-making in the ICU.

Synthetic data generation and simulation for critical care research and development

Generative AI offers a novel capability to create fully synthetic ICU datasets that resemble real-world patient data but do not correspond to actual individuals.[9] This approach has particular relevance in critical care, where acquiring large, high-quality datasets for algorithm development, clinical research, or systems testing is often limited by privacy constraints, institutional data access, and the rarity of certain clinical events. Synthetic ICU data may include artificial time-series signals, laboratory trajectories, waveforms, imaging data, and even entire simulated patient records that approximate the complexity of true clinical cases.

By generating fully artificial datasets, researchers can overcome data scarcity and enrich model training in domains where limited real-world examples exist.[45] For instance, uncommon ICU events such as ARDS, cardiac tamponade, or specific drug-related adverse events may not occur with sufficient frequency at a single institution to support robust model development (Supplementary Fig. 2). Generative adversarial networks (GANs), diffusion models, and variational autoencoders have been applied to create synthetic patient cohorts for these scenarios. In one study, a GAN was trained to generate synthetic ICU time-series data, which, when added to limited real-world training sets, improved model performance in predicting vasopressor requirements and the need for mechanical ventilation. This synthetic expansion allows models to learn nuanced patterns of disease progression that would otherwise be underrepresented in small datasets.

Beyond predictive model development, synthetic data enables simulation-based evaluation of clinical interventions and care policies in a no-risk environment (Supplementary Fig. 2).[46] Virtual patient populations can be generated to simulate the application of novel therapeutic strategies or care pathways, facilitating hypothesis generation and informing the design of prospective trials.[45] Bayesian network-based models and autoencoders have been used to simulate multi-organ ICU trajectories, enabling researchers to explore the potential benefits and harms of specific management strategies before clinical deployment.[47]

Synthetic EHR datasets also play an important role in the development, validation, and security testing of health information technology infrastructure.[48] Before integrating new ICU software platforms or clinical decision support systems into live environments, engineers can utilize fully synthetic patient records to conduct functionality and safety testing without compromising patient privacy.[48]

Multiple generative architectures contribute to the production of synthetic ICU data.[48] GANs and variational autoencoders are commonly used for structured and time-series data, whereas LLMs have shown notable capabilities in generating both structured clinical data and synthetic clinical text. In few-shot learning scenarios, LLMs presented with limited sample patient records have produced synthetic datasets that closely mirror real-world distributions of laboratory values, vital signs, and comorbidity profiles. Furthermore, LLMs can generate simulated clinical narratives such as progress notes, consultation summaries, and discharge documentation. These synthetic text corpora are valuable for developing natural language processing algorithms, enabling systems to train and validate on a broader range of writing styles, terminologies, and clinical scenarios than typically found in limited institutional datasets.[48] While synthetic data generation does not replace prospective clinical studies, these fully artificial datasets provide valuable complementary resources for model development, systems testing, educational simulation, and exploratory research in critical care (Supplementary Fig. 3).[48]

Generative AI for ICU data summarization and real-time monitoring

Continuous, high-dimensional data across multiple domains, including physiologic waveforms, vital signs, ventilator parameters, laboratory results, and medication administration records, are generated in the ICU setting.[20] Synthesizing these complex data streams into actionable clinical insights in real time remains a persistent cognitive challenge for ICU clinicians. Generative AI, particularly LLMs and emerging multimodal architectures, offers novel capabilities to transform raw physiologic and laboratory data into coherent, human-readable narratives that may enhance situational awareness, reduce cognitive burden, and support timely clinical decision-making.[20]

In contrast to clinical decision support systems, where generative AI functions as an AI medical agent that offers diagnostic or therapeutic recommendations, ICU data summarization serves a complementary role as an AI clinical reporter. In this role, generative models continuously monitor evolving physiologic data, integrating multiple structured inputs to produce concise, temporally contextualized narrative assessments of patient status. For example, an LLM-based system monitoring hemodynamics may generate: "*Over the past six hours, mean arterial pressure decreased from 75 to 60 mm Hg, requiring norepinephrine escalation from 5 to 15 µg/min. Heart rate increased from 90 to 120 beats per minute, while urine output declined, and lactate levels rose, consistent with worsening hypoperfusion*" (Fig. 3). These synthesized outputs transform complex temporal trajectories into concise, interpretable summaries, enabling clinicians to track evolving physiological patterns rather than interpret isolated data points.

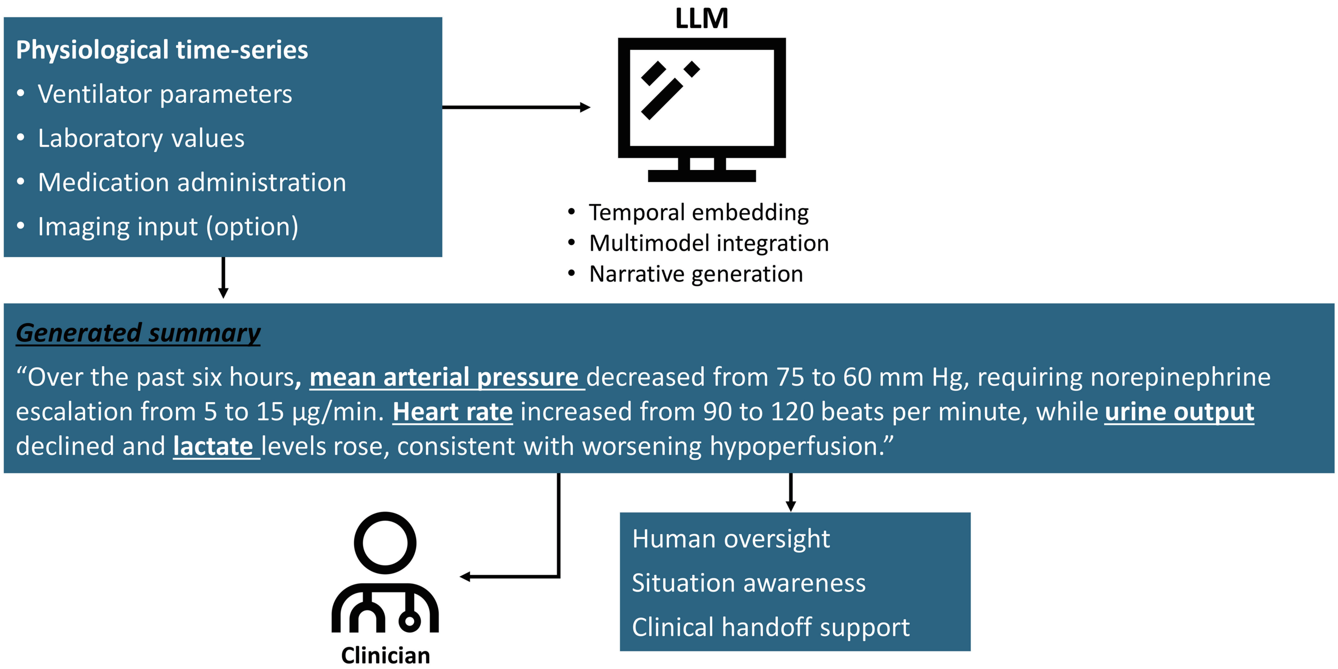

Fig. 3. Generative AI system for ICU data summarization and real-time multimodal monitoring. Example of transforming high-dimensional vital signs, laboratory trends, and ventilator parameters into human-readable clinical narratives. AI, artificial intelligence; ICU, intensive care unit; LLM, large language model.

Recent studies have demonstrated the feasibility of using transformer-based architectures to encode time-series data for LLM-driven interpretation.[49] These approaches convert streaming physiologic measurements into temporally ordered text or embeddings that can be processed by language models, enabling the generation of narrative summaries that reflect both current values and their clinical trajectories.[49] Some prototypes incorporate waveform segmentation and anomaly detection, allowing real-time interpretation of respiratory patterns, cardiac arrhythmias, and ventilator waveforms within a unified generative framework.[50,51]

Legal, ethical, and regulatory considerations

The integration of generative AI into intensive care introduces complex legal, ethical, and regulatory challenges.[19,20] Despite early promise, these opportunities coexist with substantial risks that must be addressed before broad adoption. In a setting where clinical decisions carry immediate life-or-death consequences, issues related to transparency, safety, accountability, and fairness are paramount. Several domains require particular attention.

Explainability and transparency

AI systems influencing clinical decisions must provide transparent and interpretable outputs. Clinicians need to understand the rationale behind AI recommendations to independently assess their validity. Generative AI models, particularly LLMs, offer improved explainability by producing chain-of-thought reasoning and citing sources using RAG techniques. For example, an AI system may recommend initiating hydrocortisone for septic shock while providing the underlying clinical criteria and citing relevant guidelines. Both the U.S. Food and Drug Administration (FDA) and the European Union (EU) AI Act emphasize transparency as a prerequisite for clinical use, requiring that clinicians be trained to understand and evaluate AI-generated outputs.[9]

Safety and hallucination mitigation

AI hallucinations, instances in which systems generate fabricated or erroneous information with unwarranted confidence, pose significant risks in medicine. Hallucinations may lead to incorrect diagnoses, inappropriate treatments, or the presentation of false evidence. Recognizing this risk does not negate AI's potential utility; rather, it underscores the need for safeguards. Developers are employing techniques such as RAG, external knowledge integration, and post-processing validation. Continuous performance monitoring, version control upon model updates, and independent verification of AI outputs are critical safeguards prior to widespread clinical deployment.[9]

Bias and fairness

Bias in AI systems remains a significant ethical and legal concern, particularly in critical care settings where vulnerable populations may be disproportionately affected. Generative AI models trained on historical data may inadvertently perpetuate or amplify existing disparities related to race, gender, or socioeconomic status. Empirical studies have documented that LLMs can generate biased differential diagnoses and management plans based on demographic cues. Addressing these risks is essential to realizing AI's potential to improve equity. Mitigation strategies include diverse and representative training datasets, bias auditing, subgroup performance evaluations, and application of fairness constraints during model fine-tuning. Both the EU AI Act and proposed U.S. regulatory frameworks emphasize bias detection and correction as regulatory requirements.

Liability and accountability

Determining legal responsibility when AI-generated recommendations contribute to adverse outcomes remains unresolved. Current legal frameworks generally hold clinicians accountable as final decision-makers. However, as AI tools evolve toward greater autonomy and complexity, the question of shared liability involving clinicians, institutions, and AI developers is under active legal debate. Regulatory clarity regarding safe harbors, manufacturer obligations, and clinician protections will be necessary as generative AI adoption expands. Hospitals are also advised to update informed consent processes to disclose AI use where appropriate.[52]

Regulatory and governance perspectives

Regulatory bodies worldwide are developing frameworks to guide the safe integration of AI into clinical practice, with notable regional variations.[53–55] In the United States, the FDA regulates many clinical AI applications under its medical device framework, emphasizing safety, effectiveness, and transparency. The Joint Commission and other accreditation organizations have begun incorporating AI oversight into institutional standards. In the EU, the forthcoming AI Act classifies most healthcare AI systems as "high-risk," requiring robust measures for risk management, transparency, and human oversight.[53]

Asian jurisdictions are also advancing distinct approaches.[53] China mandates algorithmic registration and transparency reporting, while Singapore has piloted governance sandboxes and implemented a Model AI Governance Framework emphasizing accountability and human-in-the-loop safeguards. India has identified healthcare as a priority in its National AI Strategy and is building a regulatory structure that seeks to balance innovation with equity. These initiatives illustrate Asia's growing role in shaping the global regulatory landscape.[53]

In low- and middle-income regions, particularly in Africa and Southeast Asia, the feasibility of deploying AI depends not only on regulation but also on infrastructure.[53] Limited internet connectivity, low rates of health record digitization, workforce shortages, and underdeveloped data governance present significant challenges. Yet these settings also offer opportunities for AI to bridge gaps in access, support limited clinical capacity, and leapfrog traditional infrastructure. Success will require partnerships that build local capacity and tailor solutions to resource-limited environments.

Regional emphases differ[53,55]: the United States prioritizes safety and validation, the EU stresses transparency and proportional risk, Asian regulators highlight accountability and algorithm registration, and resource-limited settings focus on scalability and affordability. At the institutional level, health systems are forming multidisciplinary governance committees that include clinicians, ethicists, data scientists, legal experts, and patient representatives. These committees oversee performance, audit fairness across patient subgroups, and support clinician education.[53] Together, these regional and institutional efforts underscore the importance of adaptable regulatory principles and international collaboration to balance innovation, equity, and accountability in critical care AI.

Patient consent and ethical deployment

While generative AI may improve care quality and efficiency, it raises concerns regarding patient autonomy and trust. Transparent communication with patients and families about AI's role in care decisions may be warranted, particularly for high-stakes recommendations. Informed consent remains challenging in critical care, where many patients lack decisional capacity. Institutions may adopt policies that ensure ethical oversight and protect patient interests while enabling responsible innovation.[9,52]

Institutional oversight and deployment

Most institutions are adopting staged deployment approaches, often beginning with pilot studies, research protocols, or quality improvement initiatives. Establishing multidisciplinary AI oversight committees that include clinicians, ethicists, data scientists, legal experts, and patient representatives is increasingly recommended. These bodies monitor AI performance, evaluate safety and fairness, guide institutional policy, and support clinician education.

Patient and family perspectives: Trust, consent, and communication

Critical care is uniquely stressful for patients and families, where uncertainty and information overload are common. Generative AI tools can augment communication by transforming complex medical findings into clear, accessible, and culturally tailored narratives, thereby reducing misunderstanding and supporting shared decision-making. For example, AI-generated family updates or lay-language discharge summaries may improve the comprehension of ICU care plans. However, patient trust depends on transparency. Families must be informed when AI contributes to diagnostic reasoning or prognostic estimates, and clinicians should emphasize that ultimate responsibility rests with the care team. Consent processes may need to adapt, particularly for AI tools that directly inform diagnosis, treatment, or prognosis, by explicitly disclosing the role of AI in care decisions. In addition, institutions should consider incorporating patient and family representatives into AI oversight committees, ensuring that deployment strategies reflect the perspectives and values of those most affected. By centering patients and families, generative AI can evolve from being primarily a clinician support tool to a bridge that enhances trust, communication, and autonomy in critical care.

Future perspectives

Generative AI may, with further validation, help redefine multiple dimensions of intensive care delivery as it advances toward deeper clinical integration (Supplementary Table 3). AI agents may eventually execute complex, multistep workflows such as sepsis bundle initiation, antimicrobial stewardship, and ventilator liberation protocols by autonomously monitoring physiologic data, initiating early interventions, and coordinating care pathways under clinician oversight.[56] Embedded within EHRs, generative AI models may serve as clinical co-pilots, providing real-time assistance with documentation, clinical ordering, and administrative tasks while continuously synthesizing patient data to support clinical decision-making. Institution-specific models trained on local protocols and care standards may further enhance decision support relevance and consistency across diverse ICU environments. Meanwhile, precision critical care applications that integrate continuous physiologic monitoring with genomic, multi-omic, and biomarker data may enable individualized prognostication and therapeutic tailoring for critically ill patients.

Beyond bedside care, generative AI may advance scientific discovery by generating hypotheses, identifying biomarkers, optimizing clinical trial design, and synthesizing emerging literature to accelerate critical care research.[19,20] Future integration of diffusion models, GANs, and RAG-enhanced multimodal AI may enable real-time continuous renal replacement therapy (CRRT) adjustments, precision ARDS phenotyping and individualized ventilator strategies, and bias-mitigated sepsis prediction. However, the safe and effective deployment of these technologies will require ongoing evaluation of transparency, safety, bias mitigation, and regulatory compliance. Sustained, multidisciplinary collaboration among intensivists, data scientists, regulators, ethicists, and patient representatives will be essential to ensure that generative AI augments, rather than supplants, human clinical judgment. In the high-acuity ICU setting, where complex decisions must be made rapidly and under conditions of data overload, responsibly deployed generative AI has the potential to become a foundational element of next-generation critical care practice.

Limitations

This review has several limitations. First, as a narrative review, it does not use a formal systematic search strategy or quantitative evidence synthesis; therefore, relevant studies may have been missed, and the strength of evidence across applications could not be formally graded. Second, the field of generative AI is evolving rapidly, and some cited models, platforms, and regulatory frameworks may change soon after publication. Third, many proposed applications in critical care remain conceptual or are supported primarily by retrospective studies, prototype systems, or early implementation reports rather than prospective clinical trials demonstrating improvements in patient-centered outcomes. Fourth, the performance, safety, and generalizability of generative AI tools may vary across institutions, EHR systems, patient populations, languages, and resource settings. Finally, ethical, legal, regulatory, and workflow considerations remain incompletely resolved. Accordingly, the applications discussed in this review should be interpreted as emerging opportunities that require rigorous validation, prospective evaluation, and ongoing governance before widespread clinical implementation.

Conclusions

Generative AI represents a potential paradigm shift pending rigorous validation and integration, with current evidence supporting roles in documentation assistance, knowledge synthesis, and prototype EHR integration. Emerging applications such as multimodal physiologic monitoring, ventilator liberation support, and synthetic ICU data simulation remain largely speculative and require further study. With careful oversight, governance, and ethical safeguards, generative AI may augment clinician efficiency, improve patient safety, and enable precision care. Future progress is likely to leverage diffusion models, GANs, and RAG-enhanced multimodal AI, enabling real-time CRRT adjustments, precision ARDS phenotyping, and bias-mitigated sepsis prediction. Generative AI will not replace intensivists but will likely become a foundational co-pilot in high-acuity care, provided rigorous validation and ethical safeguards are in place.

Supporting information

Supplementary material for this article is available at https://doi.org/10.14218/JTCCM.2025.00022.

Acknowledgments

During the preparation of this work, the authors used ChatGPT 5.5 to edit the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Funding

This research received no external funding.

Conflict of interest

KK is an Associate Editor-in-Chief of J*ournal of Translational Critical Care Medicine*. The article was subject to the journal's standard procedures, with peer review handled independently of the editor and the affiliated research groups.

Author contributions

Conceptualization (WC, CT, KK), writing-original draft preparation (WC, CT), writing-editing (KK). All authors have read and approved the final version.

References

[1] Idan D, Einav S. Primer on large language models: an educational overview for intensivists. Crit Care 2025;29(1):238. doi:10.1186/s13054-025-05479-4, PMID:40506762.

[2] Hwai H, Ho YJ, Wang CH, Huang CH. Large language model application in emergency medicine and critical care. J Formos Med Assoc 2025;124(8):696–699. doi:10.1016/j.jfma.2024.08.032, PMID:39198112.

[3] Biesheuvel LA, Dongelmans DA, Elbers PWG. Artificial intelligence to advance acute and intensive care medicine. Curr Opin Crit Care 2024;30(3):246–250. doi:10.1097/MCC.0000000000001150, PMID:38525882.

[4] Moor M, Banerjee O, Abad ZSH, Krumholz HM, Leskovec J, Topol EJ, *et al*. Foundation models for generalist medical artificial intelligence. Nature 2023;616(7956):259–265. doi:10.1038/s41586-023-05881-4, PMID:37045921.

[5] Wu X, Huang Y, He Q. A large language model improves clinicians' diagnostic performance in complex critical illness cases. Crit Care 2025;29(1):230. doi:10.1186/s13054-025-05468-7, PMID:40481529.

[6] Meskó B. The Impact of Multimodal Large Language Models on Health Care's Future. J Med Internet Res 2023;25:e52865. doi:10.2196/52865, PMID:37917126.

[7] Bedi S, Liu Y, Orr-Ewing L, Dash D, Koyejo S, Callahan A, *et al*. Testing and Evaluation of Health Care Applications of Large Language Models: A Systematic Review. JAMA 2025;333(4):319–328. doi:10.1001/jama.2024.21700, PMID:39405325.

[8] Kunze KN, Nwachukwu BU, Cote MP, Ramkumar PN. Large Language Models Applied to Health Care Tasks May Improve Clinical Efficiency, Value of Care Rendered, Research, and Medical Education. Arthroscopy 2025;41(3):547–556. doi:10.1016/j.arthro.2024.12.010, PMID:39694303.

[9] Biesheuvel LA, Workum JD, Reuland M, van Genderen ME, Thoral P, Dongelmans D, *et al*. Large language models in critical care. J Intensive Med 2025;5(2):113–118. doi:10.1016/j.jointm.2024.12.001, PMID:40241839.

[10] Lu Y, Wu H, Qi S, Cheng K. Artificial Intelligence in Intensive Care Medicine: Toward a ChatGPT/GPT-4 Way? Ann Biomed Eng 2023;51(9):1898–1903. doi:10.1007/s10439-023-03234-w, PMID:37179277.

[11] Sheikh MS, Thongprayoon C, Suppadungsuk S, Miao J, Qureshi F, Kashani K, *et al*. The Future of AI-Assisted Patient Education in Critical Care Nephrology. Blood Purif 2025;54(3):156–159. doi:10.1159/000543350, PMID:39733768.

[12] Moulaei K, Yadegari A, Baharestani M, Farzanbakhsh S, Sabet B, Reza Afrash M. Generative artificial intelligence in healthcare: A scoping review on benefits, challenges and applications. Int J Med Inform 2024;188:105474. doi:10.1016/j.ijmedinf.2024.105474, PMID:38733640.

[13] Reddy S. Generative AI in healthcare: an implementation science informed translational path on application, integration and governance. Implement Sci 2024;19(1):27. doi:10.1186/s13012-024-01357-9, PMID:38491544.

[14] Pinsky MR, Bedoya A, Bihorac A, Celi L, Churpek M, Economou-Zavlanos NJ, *et al*. Use of artificial intelligence in critical care: opportunities and obstacles. Crit Care 2024;28(1):113. doi:10.1186/s13054-024-04860-z, PMID:38589940.

[15] Soroush A, Giuffrè M, Chung S, Shung DL. Generative Artificial Intelligence in Clinical Medicine and Impact on Gastroenterology. Gastroenterology 2025;169(3):502–517.e1. doi:10.1053/j.gastro.2025.03.038, PMID:40245953.

[16] Sun QW, Miller J, Hull SC. Charting the ethical landscape of generative AI-augmented clinical documentation. J Med Ethics 2026;52(e1):e4–e7. doi:10.1136/jme-2024-110656, PMID:40436441.

[17] Howell MD. Generative artificial intelligence, patient safety and healthcare quality: a review. BMJ Qual Saf 2024;33(11):748–754. doi:10.1136/bmjqs-2023-016690, PMID:39048292.

[18] Sheth S, Baker HP, Prescher H, Strelzow JA. Ethical Considerations of Artificial Intelligence in Health Care: Examining the Role of Generative Pretrained Transformer-4. J Am Acad Orthop Surg 2024;32(5):205–210. doi:10.5435/JAAOS-D-23-00787, PMID:38175996.

[19] Cheungpasitporn W, Thongprayoon C, Kashani KB. Advances in critical care nephrology through artificial intelligence. Curr Opin Crit Care 2024;30(6):533–541. doi:10.1097/MCC.0000000000001202, PMID:39248074.

[20] Cheungpasitporn W, Thongprayoon C, Ronco C, Kashani KB. Generative AI in Critical Care Nephrology: Applications and Future Prospects. Blood Purif 2024;53(11-12):871–883. doi:10.1159/000541168, PMID:39217985.

[21] Miao J, Thongprayoon C, Craici IM, Cheungpasitporn W. How to improve ChatGPT performance for nephrologists: a technique guide. J Nephrol 2024;37(5):1397–1403. doi:10.1007/s40620-024-01974-z, PMID:38771519.

[22] Miao J, Thongprayoon C, Suppadungsuk S, Garcia Valencia OA, Cheungpasitporn W. Integrating Retrieval-Augmented Generation with Large Language Models in Nephrology: Advancing Practical Applications. Medicina (Kaunas) 2024;60(3):445. doi:10.3390/medicina60030445, PMID:38541171.

[23] Miao J, Thongprayoon C, Suppadungsuk S, Krisanapan P, Radhakrishnan Y, Cheungpasitporn W. Chain of Thought Utilization in Large Language Models and Application in Nephrology. Medicina (Kaunas) 2024;60(1):148. doi:10.3390/medicina60010148, PMID:38256408.

[24] Cabral S, Restrepo D, Kanjee Z, Wilson P, Crowe B, Abdulnour RE, *et al*. Clinical Reasoning of a Generative Artificial Intelligence Model Compared With Physicians. JAMA Intern Med 2024;184(5):581–583. doi:10.1001/jamainternmed.2024.0295, PMID:38557971.

[25] Abdullahi T, Singh R, Eickhoff C. Learning to Make Rare and Complex Diagnoses With Generative AI Assistance: Qualitative Study of Popular Large Language Models. JMIR Med Educ 2024;10:e51391. doi:10.2196/51391, PMID:38349725.

[26] Barabucci G, Shia V, Chu E, Harack B, Laskowski K, Fu N. Combining Multiple Large Language Models Improves Diagnostic Accuracy. NEJM AI 2024;1(11):AIcs2400502. doi:10.1056/AIcs2400502.

[27] Eriksen AV, Möller S, Ryg J. Use of GPT-4 to Diagnose Complex Clinical Cases. NEJM AI 2023;1(1):AIp2300031. doi:10.1056/AIp2300031.

[28] Tan S. Advancing Clinical Decision Support with Reliable, Transparent Large Language Models. Yale School of Medicine; 2024. Available from: https://medicine.yale.edu/news-article/advancing-clinical-decision-support-with-reliable-transparent-large-language-models/.

[29] Low YS, Jackson ML, Hyde RJ, Brown RE, Sanghavi NM, Baldwin JD, *et al*. Answering real-world clinical questions using large language model, retrieval-augmented generation, and agentic systems. Digit Health 2025;11:20552076251348850. doi:10.1177/20552076251348850, PMID:40510193.

[30] Sheikh MS, Kashani KB, Gregoire JR, Thongprayoon C, Miao J, Craici IM, *et al*. Digital transformation of nephrology POCUS education-Integrating a multiagent, artificial intelligence, and human collaboration-enhanced curriculum with expert feedback. Digit Health 2025;11:20552076251328807. doi:10.1177/20552076251328807, PMID:40162173.

[31] Miao J, Thongprayoon C, Craici IM, Cheungpasitporn W. How to incorporate generative artificial intelligence in nephrology fellowship education. J Nephrol 2024;37(9):2491–2497. doi:10.1007/s40620-024-02165-6, PMID:39621255.

[32] Sanderson AL, Burns JP. Clinical Documentation for Intensivists: The Impact of Diagnosis Documentation. Crit Care Med 2020;48(4):579–587. doi:10.1097/CCM.0000000000004200, PMID:32205605.

[33] Cao DY, Silkey JR, Decker MC, Wanat KA. Artificial intelligence-driven digital scribes in clinical documentation: Pilot study assessing the impact on dermatologist workflow and patient encounters. JAAD Int 2024;15:149–151. doi:10.1016/j.jdin.2024.02.009, PMID:38571698.

[34] Ma SP, Liang AS, Shah SJ, Smith M, Jeong Y, Devon-Sand A, *et al*. Ambient artificial intelligence scribes: utilization and impact on documentation time. J Am Med Inform Assoc 2025;32(2):381–385. doi:10.1093/jamia/ocae304, PMID:39688515.

[35] Hassan H, Zipursky AR, Rabbani N, You JG, Tse G, Orenstein E, *et al*. Clinical Implementation of Artificial Intelligence Scribes in Health Care: A Systematic Review. Appl Clin Inform 2025;16(4):1121–1135. doi:10.1055/a-2597-2017, PMID:40306686.

[36] Diaz N. Epic, Nuance to bring GPT-4-powered documentation to EHRs. Becker's Hospital Review; 2023. Available from: https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/epic-nuance-to-bring-gpt-4-powered-documentation-to-ehrs/.

[37] Rayan Z, Alfonse M, Salem ABM. Intensive Care Unit (ICU) Data Analytics using Machine Learning Techniques. Int J Inf Theor Appl 2019;26(1):69–82.

[38] Ramos G, Gjini E, Coelho L, Silveira M. Unsupervised learning approach for predicting sepsis onset in ICU patients. Annu Int Conf IEEE Eng Med Biol Soc 2021;2021:1916–1919. doi:10.1109/EMBC46164.2021.9629559, PMID:34891661.

[39] Yao J, Ong ML, Mun KK, Liu S, Motani M. Hybrid feature learning using autoencoders for early prediction of sepsis. Computing in Cardiology 2019;46:1–4. doi:10.22489/CinC.2019.243.

[40] Liu J, Ma Z, Wang Z, Zou C, Ren J, Wang Z, *et al*. A survey on diffusion models for anomaly detection. arXiv 2025. doi:10.48550/arXiv.2501.11430.

[41] Reuter E. FDA grants de novo nod to AI tool for detecting sepsis. MedTech Dive; 2024. Available from: https://www.medtechdive.com/news/fda-de-novo-clearance-ai-tool-sepsis-prenosis/712279/.

[42] Choi I, Long Q, Getzen E. Filling the gaps: leveraging large language models for temporal harmonization of clinical text across multiple medical visits for clinical prediction. medRxiv 2024. doi:10.1101/2024.05.06.24306959, PMID:38765975.

[43] Choudhuri A, Polgreen P, Segre A, Adhikari B. Summarizing clinical notes using LLMs for ICU bounceback and length-of-stay prediction. medRxiv 2025. doi:10.1101/2025.01.19.25320797.

[44] Hadi MU, Al-Tashi Q, Qureshi R, Shah A, Muneer A, Irfan M, *et al*. LLMs: a comprehensive survey of applications, challenges, datasets, models, limitations, and future prospects. TechRxiv 2023. doi:10.36227/techrxiv.23589741.v2.

[45] Gangwal A, Ansari A, Ahmad I, Azad AK, Wan Sulaiman WMA. Current strategies to address data scarcity in artificial intelligence-based drug discovery: A comprehensive review. Comput Biol Med 2024;179:108734. doi:10.1016/j.compbiomed.2024.108734, PMID:38964243.

[46] Abdalla HB, Kumar Y, Marchena J, Guzman S, Awlla A, Gheisari M, *et al*. The Future of Artificial Intelligence in the Face of Data Scarcity. Comput Mater Contin 2025;84(1):1073–1099. doi:10.32604/cmc.2025.063551.

[47] Waddell TH. Applying Bayesian-networks to model multi-organ pathologies [Dissertation]. OxfordUniversity of Oxford; 2023. doi:10.5287/ora-7qxk1jq52.

[48] Yoon J, Mizrahi M, Ghalaty NF, Jarvinen T, Ravi AS, Brune P, *et al*. EHR-Safe: generating high-fidelity and privacy-preserving synthetic electronic health records. NPJ Digit Med 2023;6(1):141. doi:10.1038/s41746-023-00888-7, PMID:37567968.

[49] Antonesi G, Cioara T, Anghel I, Michalakopoulos V, Sarmas E, Toderean L. A systematic review of transformers and large language models in the energy sector: towards agentic digital twins. Appl Energy 2025;401(Pt A):126670. doi:10.1016/j.apenergy.2025.126670.

[50] Chan N, Parker F, Bennett W, Wu T, Jia MY, Fackler J, *et al*. Medtsllm: Leveraging llms for multimodal medical time series analysis. arXiv 2024. doi:10.48550/arXiv.2408.07773.

[51] Qiu J, Yuan W, Lam K. The application of multimodal large language models in medicine. Lancet Reg Health West Pac 2024;45:101048. doi:10.1016/j.lanwpc.2024.101048, PMID:38524685.

[52] Shumway DO, Hartman HJ. Medical malpractice liability in large language model artificial intelligence: legal review and policy recommendations. J Osteopath Med 2024;124(7):287–290. doi:10.1515/jom-2023-0229, PMID:38295300.

[53] Zaidan E, Ibrahim IA. AI Governance in a Complex and Rapidly Changing Regulatory Landscape: A Global Perspective. Humanit Soc Sci Commun 2024;11:1121. doi:10.1057/s41599-024-03560-x.

[54] De Almeida PGR, Dos Santos CD, Farias JS. Artificial intelligence regulation: a framework for governance. Ethics Inf Technol 2021;23(3):505–525. doi:10.1007/s10676-021-09593-z.

[55] Taeihagh A. Governance of artificial intelligence. Policy Soc 2021;40(2):137–157. doi:10.1080/14494035.2021.1928377.

[56] Karunanayake N. Next-generation agentic AI for transforming healthcare. Inform Health 2025;2(2):73–83. doi:10.1016/j.infoh.2025.03.001.