International Journal of Innovative Technologies in Social Science

e-ISSN: 2544-9435

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

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

ARTICLE TITLE	ARTIFICIAL INTELLIGENCE IN HEALTHCARE: DIAGNOSTIC SUPPORT AND ADMINISTRATIVE AUTOMATION

DOI	https://doi.org/10.31435/ijitss.2(50).2026.5250
RECEIVED	27 February 2026
ACCEPTED	07 April 2026
PUBLISHED	12 May 2026

ARTIFICIAL INTELLIGENCE IN HEALTHCARE: DIAGNOSTIC SUPPORT AND ADMINISTRATIVE AUTOMATION

Julia Żak *(Corresponding Author, Email: juliaza13@gmai.com)*
Czerniakowski Hospital, Warsaw, Poland
ORCID ID: 0009-0008-3383-8025

Monika Stępińska
University Clinical Center of the Medical University of Warsaw, Warsaw, Poland
ORCID ID: 0009-0008-0347-2704

Marta Omiecińska
Międzyleski Specialist Hospital of Warsaw, Warsaw, Poland
ORCID ID: 0009-0002-3134-8141

Alicja Maciejewska
Independent Public Health Care Complex in Pruszków, Pruszków, Poland
ORCID ID: 0009-0009-6903-2434

Maja Kaczor
National Medical Institute of the Ministry of the Interior and Administration, Warsaw, Poland
ORCID ID: 0009-0001-8218-5429

Weronika Trynkiewicz
Ludwik Rydygier Specialist Hospital, Kraków, Poland
ORCID ID: 0009-0008-2267-7166

Zuzanna Rybka
University Clinical Center, Medical University of Warsaw, Warsaw, Poland
ORCID ID: 0009-0001-4056-8800

Emilia Lenkiewicz
National Medical Institute of the Ministry of the Interior and Administration, Warsaw, Poland
ORCID ID: 0009-0001-4822-8157

Karolina Dąbrowska
Wojewódzki Szpital w Łomży, Łomża, Poland
ORCID ID: 0009-0002-9053-2002

Jakub Winiarczyk
Independent Public Health Care Complex in Pruszków, Pruszków, Poland
ORCID ID: 0009-0008-8163-5580

ABSTRACT

Background: Artificial intelligence (AI) is rapidly moving from theoretical research into daily hospital operations. While AI promises to improve diagnostic speed and reduce administrative burnout, its real-world success depends on how well it integrates into complex clinical workflows. This review evaluates the practical impact of both discriminative and generative AI on modern healthcare delivery.
Methods: We conducted a targeted literature search using PubMed and Semantic Scholar (via Consensus AI) for research published between January 2018 and March 2026. We structured this work as a critical narrative review to focus specifically on the practical challenges and successes of AI in real-world hospital settings. After screening for clinical relevance and the presence of deployment data, we selected 21 peer-reviewed articles for final inclusion.
Results: The evidence shows that discriminative AI significantly improves speed in acute care, particularly in sepsis forecasting and stroke triage. Newer models are also showing success in evaluating surgical risks from ECGs and tracking traumatic brain injuries. However, these tools often face challenges like "alert fatigue" and data bias. On the administrative side, generative AI—such as ambient clinical scribes—is successfully reducing documentation time and improving billing accuracy. We found that the success of these tools depends on "Human-in-the-Loop" models and mathematical explainability frameworks, like SHAP values, which help doctors understand and trust AI decisions.
Conclusion: AI is becoming a necessary requirement for modern medicine, not just an optional upgrade. By offloading high-volume data tasks to machines, AI can help return the physician's focus to direct patient care. Future success depends on using protocols that ensure algorithmic transparency and prevent automation bias, ensuring the doctor remains the final authority in clinical decisions.

KEYWORDS

Artificial Intelligence, Clinical Workflow, Generative AI, Machine Learning, Physician Burnout, Explainable AI (XAI), Implementation Science

CITATION

Julia Żak, Monika Stępińska, Marta Omiecińska, Alicja Maciejewska, Maja Kaczor, Weronika Trynkiewicz, Zuzanna Rybka, Emilia Lenkiewicz, Karolina Dąbrowska, Jakub Winiarczyk. (2026) Artificial Intelligence in Healthcare: Diagnostic Support and Administrative Automation. *International Journal of Innovative Technologies in Social Science*. 2(50). doi: 10.31435/ijitss.2(50).2026.5250

1. Introduction

Modern healthcare faces a structural crisis. Diagnostic technology advances rapidly, yet the daily clinical workflow is collapsing under extreme administrative overload (Micek et al., 2020; Sinsky et al., 2016). While artificial intelligence promises to fix this burden, a gap exists between theoretical algorithmic performance and actual hospital use. Many models that perform perfectly in a laboratory fail completely when put into a chaotic, real-world clinical environment (Challen et al., 2019; Chen et al., 2023). Because hospital administrators and clinicians are currently overwhelmed by aggressive tech marketing, the medical community urgently needs an objective guide that separates theoretical hype from practical clinical reality.

Institutionalized documentation forces highly trained physicians to act as data clerks. Direct observational time-and-motion studies reveal a stark reality. During a standard office day, physicians spend only 27.0% of their total time on direct clinical face time with patients. They spend 49.2% on Electronic Health Records and desk work (Sinsky et al., 2016). For every hour of face-to-face patient time, administrative tasks consume nearly two additional hours (Sinsky et al., 2016). This boils down to 1 to 2 hours of unpaid after-hours work each night (Sinsky et al., 2016).

This administrative weight correlates heavily with the global physician burnout crisis (Micek et al., 2020). When paperwork eclipses patient interaction, clinical empathy and trust decline (Espinoza-Vinces et al., 2025). To preserve the humanistic core of healthcare, AI has emerged as a functional necessity. AI is highly autonomous, yet, it requires constant human oversight. The algorithms process data instantly, but they lack the clinical context required to apply that data safely. To maintain clinical safety, hospitals must put in place a collaborative human-in-the-loop (HITL) strategy. In this framework, artificial intelligence processes high-volume data and the healthcare provider has diagnostic authority, validating the algorithm's outputs (Sezgin, 2023).

AI transitioned from theoretical computer science to daily diagnostic care. Radiology uses tools like Convolutional Neural Networks (CNNs) to predict cancer outcomes directly from imaging (Soganciogluet al., 2021; Bera et al., 2019). Emergency departments use machine learning to rapidly categorize patient acuity (Berlyand et al., 2018). Stroke algorithms autonomously detect large vessel occlusions to speed up mechanical thrombectomy (Heeralal et al., 2025). Intensive care units deploy predictive analytics to identify impending septic shock (Adams et al., 2022).

Existing literature catalogues the theoretical and mathematical performance of AI models in isolated settings. Still, there is a need for a better understanding of how these tools actually impact daily hospital operations. The aim of this narrative review is to discuss the operational integration of clinical AI. Specifically, this paper assesses the impact of Generative AI on reducing the clinical documentation burden and the efficacy of Discriminative AI in acute diagnostics. Finally, we evaluate the boundaries of algorithmic safety, demographic bias, and physician-computer interaction.

2. Methodology

Search Strategy and Selection Criteria To evaluate how artificial intelligence works in everyday healthcare, we conducted a targeted literature search using PubMed and Semantic Scholar (via Consensus AI). To focus on the most modern AI tools—including both deep learning and large language models—we limited our search to articles published between January 2018 and March 2026. Our search terms included combinations of "artificial intelligence," "clinical workflow," "generative AI," "ambient clinical scribing," "predictive modeling," "sepsis forecasting," and "algorithmic bias."

We structured this paper as a critical narrative review rather than a strict systematic review. This allowed us to focus specifically on the practical, everyday challenges of deploying AI in real hospital settings. We strictly excluded articles if they were not peer-reviewed, if they only discussed theoretical AI models without clinical testing, or if they lacked real-world hospital data. We selected a total of 21 peer-reviewed articles for final inclusion.

Study Selection and Synthesis We screened article titles and abstracts for clinical relevance and reviewed the full texts when necessary. We also checked the reference lists of included clinical trials and reviews to find any additional eligible studies. Finally, we grouped the evidence by key themes: diagnostic support versus administrative automation, acute care triage, sepsis forecasting, ambient AI scribes, revenue cycle management, algorithmic bias, clinical safety, and regulatory rules.

3. AI Architecture

Evaluating artificial intelligence for hospital use requires understanding the architectural differences in these systems. Clinical AI functions either as an analytical diagnostic tool (discriminative AI) or as an administrative assistant (generative AI).

Discriminative Models and Convolutional Neural Networks (CNNs) Diagnostic algorithms rely on discriminative architectures to classify data. In computational radiology, this happens through Convolutional Neural Networks (CNNs) (Yasaka & Abe, 2018). These are a type of artificial intelligence that act like a digital eye, learning to recognize diseases in medical images by automatically piecing together simple visual cues—like lines and edges—into complex diagnostic patterns (Bera et al., 2019; Yasaka & Abe, 2018). They extract hierarchical mathematical features directly from raw pixel arrays (Bera et al., 2019). Initial convolutional layers are simple edge detectors. Deeper network layers synthesize these maps into complex representations. This architecture allows the algorithm to independently recognize complex morphological patterns without explicit human instruction. Discriminative models excel at image-level prediction and are the majority of diagnostic support tools (Sogancioglu et al., 2021). For example, in active clinical practice, discriminative models detect acute large vessel occlusions on stroke CT scans (Heeralal et al., 2025) and predict septic shock by analyzing continuous ICU vital signs (Adams et al., 2022).

Generative Artificial Intelligence Administrative automation relies on Generative AI, which produces entirely new content based on learned patterns. This capability origins from "foundation models" which are massive models trained on digital general information that are subsequently adapted for specific clinical uses. Large Language Models (LLMs) operate under this framework. Highly adapted LLMs like Med-PaLM encode professional clinical knowledge. They successfully change complex human queries to do medical question answering and reasoning tasks (Singhal et al., 2023). Large Language Models understand how medical professionals naturally speak and write. Rather than categorizing a symptom, these models actively draft the entire clinical note (Bhuyan et al., 2025; Duggan et al., 2025).

To safely integrate Generative AI into the clinical environment, the underlying Large Language Models (LLMs) must accurately reflect professional medical consensus (Singhal et al., 2023). A significant challenge in AI development is the occurrence of hallucinations, where models generate false information that has no basis in the input data. To address it, researchers developed the MultiMedQA benchmark. This framework combines datasets including: professional licensing exams, complex research queries, and highly searched consumer medical questions (Singhal et al., 2023). This rigorous benchmark provided the reference for Med-PaLM which is a highly adapted, domain-specific foundation model engineered to encode professional clinical knowledge. When tested against the MultiMedQA standard, the model's outputs were assessed by licensed clinicians (Singhal et al., 2023). They judged the AI across multiple axes: factuality, reading comprehension, reasoning capability, demographic bias, and the likelihood of inducing clinical harm (Singhal et al., 2023). Evaluated under this multi-dimensional scrutiny, Med-PaLM achieved a 92.6% alignment with scientific consensus (Singhal et al., 2023). This score shows that AI tools are ready for everyday administrative tasks and direct patient care (Bhuyan et al., 2025; Singhal et al., 2023).

Hospitals are already using similar generative foundation models to automate administrative labor. In live clinical environments, Memorial Sloan Kettering Cancer Center deployed generative ambient AI scribes to autonomously draft clinical notes during complex outpatient oncology visits (Baldwin et al., 2025). Similarly, an academic health system in Philadelphia deployed acoustic foundation models to successfully reduce documentation burdens across their outpatient clinics (Duggan et al., 2025). Healthcare systems are also rapidly expanding the use of AI directly into their financial infrastructures to extract diagnoses from physician notes and pair them with medical billing codes for effective revenue cycle management (Bhuyan et al., 2025).

Table 1. Architectural and Operational Divergence of Clinical AI Modalities

System Feature	Discriminative AI	Generative AI
Mechanism	Analyzes existing physiological and anatomical data patterns.	Synthesizes novel, unstructured clinical data into structured textual outputs.
Clinical Use	Acute stroke triage, radiomic tumor microenvironment mapping, sepsis deterioration forecasting.	Ambient clinical scribing, automated revenue cycle management (RCM), medical question answering.
Algorithmics	CNNs, Random Forests, High-frequency vital sign classifiers.	Domain-specific foundation models (e.g., Med-PaLM), transformer networks, high-fidelity acoustic natural language processors.
Vulnerabilities	Data leakage (shortcut learning), demographic bias, diagnostic alert fatigue, automation bias.	Algorithmic hallucination, semantic omission, severe cognitive editing friction in complex cases.
Clinician’s Function	Diagnostic Arbiter: Mandates human verification of flagged radiological anomalies prior to clinical intervention.	Data Editor: Demands clinician supervision to meticulously correct and validate auto-generated clinical documentation.
Literature References	Adams et al. (2022); Bera et al. (2019); Challen et al. (2019); Chen et al. (2023); Heeralal et al. (2025)	Baldwin et al. (2025); Bhuyan et al. (2025); Duggan et al. (2025); Singhal et al. (2023)

4. Diagnostic AI in Radiology and Emergency Medicine, Acute Care and Pathology

Discriminative artificial intelligence yields measurable improvements in diagnostic speed and accuracy (Berlyand et al., 2018; Heeralal et al., 2025). However, real-world use reveals a sharp contrast between theoretical precision and practical limitations.

Stroke Triage and Brain Injuries Radiology and emergency medicine serve as the primary testing grounds for diagnostic AI. The greatest benefit is the acceleration of clinical workflows. Machine learning-driven triage models in the emergency department prove superior to traditional human methods in categorizing patients with complex diagnosis (Berlyand et al., 2018). In acute stroke management, this acceleration saves lives (Berlyand et al., 2018). AI systems autonomously analyze non-contrast CTs to detect large vessel occlusions. Across commercial platforms, AI detected the occlusions with sensitivity ranging from 80% to 96 % (Heeralal et al., 2025). Integrating these tools reduces the door-in-door-out time for transfer cases by an average of 66 minutes (Heeralal et al., 2025).

However, false positives also occur in the acute care workflow. AI systems frequently encounter anatomical mimickers—dense calcifications or benign vascular stenoses—which the algorithm incorrectly flags as acute thrombi (Heeralal et al., 2025). To prevent unnecessary surgeries, these tools should function as notification systems rather than final diagnostic decision makers. Effective management relies on immediate human verification of the source images.

In the management of traumatic brain injury (TBI), AI systems are increasingly being used to process complex, multimodal monitoring data—including advanced neuroimaging and biomarker profiles—to identify key prognostic indicators. This computational approach assists clinicians in adjusting treatment strategies that can dynamically respond to acute neuroinflammation and blood-brain barrier disruption (Lu et al., 2024).

Thoracic Imaging Beyond neuroimaging, deep learning speeds up the analysis of chest X-rays. CNNs identify thoracic pathologies simultaneously, ranging from pneumonia to subtle lung nodules (Sogancioglu et al., 2021). Yet, clinical applicability is constrained by the diversity of training data. Algorithms optimized on specific hospital datasets experience reduced accuracy when deployed in new environments due to differing patient demographics (Sogancioglu et al., 2021). Despite these generalizability challenges, rigorously validated AI models are proving invaluable for lung cancer triage, particularly in resource-limited settings where standard screening tools like low-dose computed tomography are not accessible. For example, the prospective, multicenter CREATE study evaluated an AI-based lung nodule malignancy score on incidental findings from chest X-rays, achieving an overall negative predictive value of 93.5% (Koksal et al., 2026). This high negative predictive performance was maintained across diverse subgroups, including younger patients and non-smokers (Koksal et al., 2026).

Cardiological Triage AI-enhanced electrocardiography models are transforming perioperative risk assessments by extracting subtle predictive features that outperform traditional scoring systems. For instance, the AI-derived QCG-Critical score, generated from a standard preoperative ECG, demonstrated strong discriminative performance for predicting 30-day postoperative mortality and other adverse acute outcomes, offering a rapid triage mechanism prior to non-cardiac surgery (Choi et al., 2026).

Sepsis Forecasting Sepsis remains a leading cause of in-hospital mortality. Advanced machine learning models improve upon traditional scoring systems by analyzing vital signs. In a prospective study involving nearly 600,000 patients, the Targeted Real-time Early Warning System (TREWS) was associated with a relative reduction in in-hospital mortality (Adams et al., 2022). The median lead time before the onset of overt organ failure was 5.7 hours, representing a great clinical advantage (Adams et al., 2022). It gives physicians time to review the chart, verify the alert, order broad-spectrum IV antibiotics, and administer IV fluids to prevent organ failure. The efficacy of predictive models is inherently constrained by the clinician's action. The 18.7% mortality reduction was achieved when clinicians acknowledged and acted upon the alert within 3 hours (Adams et al., 2022). The primary barrier to this compliance is alert fatigue. When an AI system triggers frequent false alarms, the physician can dismiss the notification and the algorithmic advantage vanishes (Adams et al., 2022).

To overcome the AI nature of early warning systems and build clinical trust, FDA-authorized AI sepsis diagnostic tools have begun incorporating Shapley Additive exPlanations (SHAP) values to provide patient-specific feature importance. Derived from cooperative game theory, SHAP quantifies the exact mathematical contribution of individual clinical variables—such as a sudden spike in heart rate or an abnormal white blood cell count—to the algorithm's final prediction. The algorithmic transparency has been shown to improve clinicians' understanding of the underlying risk calculations, successfully achieving a 98% interpretability rate among providers in acute care settings (Watson et al., 2026).

Digital Pathology Deep learning models extract sub-visual features that are inaccessible to the human eye. Precision oncology tools discriminate benign treatment confounders from true disease progression, with AI architectures successfully predicting non-small-cell lung cancer survival and treatment response with high discriminative accuracy (Bera et al., 2019). Yet, these deep learning models lack inherent transparency. Because it is challenging to determine the specific factors driving an AI's prediction, clinicians must rely on post-hoc explainability approaches to gain insight into the model's decision-making process (Bera et al., 2019).

5. Automating the Clinical Work

Ambient Clinical Intelligence and History Taking Modern Ambient Clinical Intelligence (ACI) overhauls clinical documentation. Following the widespread adoption of Electronic Health Records (EHRs) the daily clinical workflow collapsed under an extreme typing burden (Micek et al., 2020; Sinsky et al., 2016). To cope, some hospitals introduced legacy speech-to-text dictation software. While dictation sped up data entry compared to manual typing, it introduced severe flaws. It forced doctors to hold a microphone and dictate like a robot, explicitly reciting punctuation commands out loud (e.g., *"Patient denies chest pain period"*). Furthermore, these older dictation tools struggled with background noise and complex accents, requiring heavy manual correction that kept the physician tethered to the computer screen (Duggan et al., 2025). ACI removes both the keyboard and the dictation microphone. These systems use high-fidelity acoustic sensors and Generative AI to capture the rhythm of a natural medical conversation. By synthesizing raw dialogue into a formal clinical note, ambient scribes fix the disproportionate time physicians spend navigating the EHR (Baldwin et al., 2025; Duggan et al., 2025). During a 2024 evaluation within an academic health system, investigators recorded a 20.4% drop in documentation time per appointment (Duggan et al., 2025). The software pushes the clinician out of the role of data creator and turns them into a supervisory data editor. This shift helps collect documentation faster, allowing physicians to finally focus on the patient rather than the screen.

ACI must also prove safe in complex specialty environments where clinical detail is key. A prospective pilot study at Cancer Center put ambient scribes into outpatient oncology clinics. Almost half the participants saw a drop in the work hours. The algorithm handled a lengthy, unstructured patient history well. Doctors used the software more often for comprehensive new patient visits (21%) than for routine follow-ups (12%) (Baldwin et al., 2025). In the context of preoperative workflow, generative AI-assisted scribes are also being used to ease the acute administrative burden. Recent simulation trials show that integrating AI into preanesthetic consultations reduces total documentation duration by 18% and significantly decreases the time physicians spend fixated on computer screens (Rahrisch et al., 2026).

Algorithms still fail, however, to capture highly complex clinical scenarios. They miss intricate oncology treatment logic or specific clinical trial parameters. While many providers saved time, 17% actually experienced an increase in documentation hours (Baldwin et al., 2025). The effort of editing the machine's complex outputs replaced the physical fatigue of typing. While ambient AI shows clear promise for collecting histories and drafting notes, the technology requires deeper clinical adaptation before doctors can truly reclaim that saved time for direct patient interaction.

Clinical Finances Generative AI does more than draft clinical narratives. It assists in Revenue Cycle Management (RCM) (Bhuyan et al., 2025). RCM is the financial process hospitals use to track a patient from the first appointment to the final payment, ensuring every medical service is accurately documented. Manual process of assigning International Classification of Diseases and Current Procedural Terminology tags creates a backlog (Bhuyan et al., 2025). AI models can map the semantic structure of clinical language and bypass rudimentary keyword matching entirely (Bhuyan et al., 2025). The algorithm reads the doctor's note, follows the complex medical logic, and retrieves the important data on its own. It then puts this information directly into the correct billing and matches medical codes with the complex treatments, which reduces human mistakes that lead to rejected insurance claims (Bhuyan et al., 2025). Using AI as a coding tool speeds up how quickly the hospital gets paid and allows it to move money away from paperwork and put it back into treating patients.

6. Social, Technical, Ethical, and Safety Implications

As artificial intelligence transitions from an experimental novelty to a deeply integrated clinical tool, its use introduces some risks that extend beyond algorithmic accuracy. Putting these systems in place without rigorous oversight threatens to embed diagnostic biases and introduces new clinical errors.

Algorithmic Bias and Health Disparities A fundamental limitation of current clinical AI is its vulnerability to algorithmic bias, which can strengthen systemic healthcare inequalities. Machine learning models rely entirely on the data used to train them; models trained on homogeneous datasets from well-funded, tertiary academic centers frequently fail to generalize across diverse demographic groups (Chen et al., 2023). Assessments of AI models across clinical sub-populations have revealed inequalities in how patients are diagnosed and treated. Diagnostic imaging algorithms have demonstrated less diagnostic accuracy for Hispanic patients and socioeconomically disadvantaged populations (such as those insured by Medicaid) compared to privately insured, predominantly White populations (Chen et al., 2023). Furthermore, deep learning algorithms are highly susceptible to "data leakage." In computational pathology, image intensity on digitized slides has been shown to accidentally leak ethnicity data or hospital-specific acquisition protocols, causing the algorithm to base its predictions on demographic confounders rather than true tumor biology (Chen et al., 2023).

Clinical Safety Integrating autonomous systems into medical environments introduces unique safety vulnerabilities. One of the most documented risks is "automation bias," a psychological phenomenon where clinicians become overly reliant on automated systems, occasionally not following their own clinical judgment or failing to seek out contradictory evidence when an AI provides a diagnosis (Challen et al., 2019). If an AI system flags a chest radiograph as normal, a fatigued clinician may resign from their visual search and miss a subtle pathology.

Additionally, AI systems using reinforcement learning face the unique risk of "reward hacking" (Challen et al., 2019). In these scenarios, the algorithm discovers a mathematical shortcut to maximize its programmed objective and inadvertently cause clinical harm. For example, an automatic ICU tool meant to keep vital signs steady might try to make its digital score better by waiting too long to start a life-saving treatment. This happens because the treatment makes the patient's vitals look shaky for a moment, which the AI sees as a failure (Challen et al., 2019).

Regulatory Frameworks, Software as a Medical Device As AI algorithms assume more active roles in clinical triage and diagnosis, the legal regulations must develop to ensure patient safety. The U.S. Food and Drug Administration has begun categorizing these advanced algorithms under the regulatory framework of "Software as a Medical Device" (Espinoza-Vinces et al., 2025). Regulating AI presents a unique challenge: unlike traditional medical devices, machine learning models dynamically change their performance in response to new data. Consequently, traditional, static approval pathways are insufficient. Regulatory agencies must put in place modernized evaluation frameworks that monitor algorithmic safety throughout the software's clinical use (Espinoza-Vinces et al., 2025).

Preserving Empathy and the Human-in-the-Loop Necessity The rapid evolution of these technologies frequently sparks anxiety regarding whether AI will ultimately replace the physician. However, the literature contraindicates this hypothesis, showing that the therapeutic alliance between doctor and patient—built on empathy and compassion—cannot be replicated by algorithms (Espinoza-Vinces et al., 2025). The solution to decrease both algorithmic bias and clinical safety risks is the strict maintenance of a human-in-the-loop (HITL) framework, where "human" is a doctor validating AI outcomes. HITL requires a health care system to involve multidisciplinary teams—including clinicians, data scientists, and ethicists—operating within an integrated loop to continually evaluate AI outputs and refine workflows (Sezgin, 2023). Hospitals can fundamentally transform patient treatment by integrating the expertise of clinical teams with the analytical power of AI. This approach ensures a safer and higher standard of care (Sezgin, 2023).

7. Conclusion

The integration of artificial intelligence marks a fundamental change in healthcare. AI has moved from a theoretical concept into a likely future requirement because it saves time, reduces costs, and improves diagnostic accuracy. By dividing the computational labor, healthcare systems use discriminative algorithms to speed up clinical tasks and generative AI to handle heavy administrative work. Discriminative AI tools—such as automated large vessel occlusion detectors for acute stroke and the Targeted Real-time Early Warning System for sepsis—improve speed in acute care and measurably reduce in-hospital mortality. Moving beyond basic triage, these predictive models also show great value in evaluating surgical risks from ECGs, tracking traumatic brain injuries, and detecting incidental lung cancer.

The research highlights that mathematical precision is not enough without human context. Real-world use reveals persistent flaws. Highly sensitive stroke algorithms often trigger false positives by mistaking benign anatomical structures for blockages. Similarly, chest X-ray models remain limited by biases in their training data, and predictive triage systems can cause severe alert fatigue if not carefully integrated into the workflow. Therefore, the collaborative "Human-in-the-Loop" model is not just a temporary safety measure; it is a strict necessity for safely incorporating medical AI.

Ultimately, the goal of AI in medicine is to empower the provider, not to replace them. By offloading high-volume data tasks and billing work to machines, hospitals can remove the current administrative obstacles that block the doctor-patient relationship. Success depends on using rigorous protocols that stop automation bias and ensure fairness for all patients. This ensures the physician retains final authority, preserving the human core of clinical care.

REFERENCES

1. Adams, R., Henry, K. E., Sridharan, A., Soleimani, H., Zhan, A., Nielson, J. A., ... Saria, S. (2022). Prospective, multi-site study of patient outcomes after implementation of the TREWS machine learning-based early warning system for sepsis. *Nature Medicine, 28*(7), 1455–1460. https://doi.org/10.1038/s41591-022-01894-0
2. Baldwin, A., Lipitz-Snyderman, A., Goyal, S., Kuperman, G., Stetson, P. D., Chimonas, S., ... Damstrom, M. (2025). Early insights into the impact of an ambient AI scribe solution on clinical documentation to reduce clinician burnout in oncology. *Journal of Clinical Oncology, 43*(16_suppl), Article e13722. https://doi.org/10.1200/JCO.2025.43.16_suppl.e13722
3. Bera, K., Braman, N., Gupta, A., Velcheti, V., & Madabhushi, A. (2019). Predicting cancer outcomes with radiomics and artificial intelligence in radiology. *Nature Reviews Clinical Oncology, 16*(11), 703–715. https://doi.org/10.1038/s41571-021-00560-7
4. Berlyand, Y., Raja, A. S., Dorner, S. C., Prabhakar, A. M., Sonis, J. D., ... Yun, B. J. (2018). How artificial intelligence could transform emergency department operations. *The American Journal of Emergency Medicine, 36*(8), 1515–1517. https://doi.org/10.1016/j.ajem.2018.01.017
5. Bhuyan, S. S., Sateesh, V., Mukul, N., Galvankar, A., Mahmood, A., Nauman, M., Rai, A., Bordoloi, K., Basu, U., & Samuel, J. (2025). Generative AI in healthcare: Opportunities for clinical excellence and revenue cycle management. *Journal of Medical Systems, 49*, Article 10. https://doi.org/10.1007/s10916-024-02136-1
6. Challen, R., Denny, J., Pitt, M., Gompels, L., Edwards, T., & Tsaneva-Atanasova, K. (2019). Artificial intelligence, bias and clinical safety. *BMJ Quality & Safety, 28*(3), 231–237. https://doi.org/10.1136/bmjqs-2018-008370
7. Chen, R. J., Wang, J. J., Williamson, D. F., Chen, T. Y., Lipkova, J., Lu, M. Y., ... Mahmood, F. (2023). Algorithm fairness in artificial intelligence for medicine and healthcare. *Nature Biomedical Engineering, 7*(6), 719–742. https://doi.org/10.1038/s41551-023-01056-8
8. Choi, H. M., Kim, Y., Kim, J., Park, J., Lee, J. H., Yoon, Y. E., Oh, I. Y., Song, I. A., & Cho, Y. (2026). Artificial intelligence-enhanced ECG score for perioperative risk assessment in non-cardiac surgery. *European Heart Journal - Digital Health, 7*, 1–10. https://doi.org/10.1093/ehjdh/ztag006
9. Duggan, M. J., Gervase, J., Schoenbaum, A., Hanson, W., Howell, J. T., III, Sheinberg, M., & Johnson, K. B. (2025). Clinician experiences with ambient scribe technology to assist with documentation burden and efficiency. *JAMA Network Open, 8*(2), Article e2460637. https://doi.org/10.1001/jamanetworkopen.2024.60637
10. Espinoza-Vinces, C., Martínez, M. C., & Irimia, P. (2025). The impact of artificial intelligence on the person-centred, doctor-patient relationship. *The Journal of Headache and Pain, 26*, Article 192. https://doi.org/10.1186/s10194-025-02143-8
11. Heeralal, V. T., Chadee, S. E., Ilyaev, B., Ilyaev, R., & Ilyayeva, S. (2025). Artificial intelligence in stroke care: A narrative review of diagnostic, predictive, and workflow applications. *Cureus, 17*(9), Article e93430. https://doi.org/10.7759/cureus.93430
12. Koksal, D., Govindarajan, A., Gonuguntla, H. K., Nayci, S., Cordova, R., Zidan, M. H., McCutcheon, S., Saha, A., Kantharaju, P., Sen, S., Agrawal, R., & Wulandari, L. (2026). Evaluation of an artificial intelligence defined lung nodule malignancy score in incidental pulmonary nodules: The CREATE study. *Mayo Clinic Proceedings: Digital Health, 4*(1), Article 100335. https://doi.org/10.1016/j.mcpdig.2026.100335
13. Lu, Y., Jin, J., Zhang, H., Lu, Q., Zhang, Y., Liu, C., Liang, Y., Tian, S., Zhao, Y., & Fan, H. (2024). Traumatic brain injury: Bridging pathophysiological insights and precision treatment strategies. *Neural Regeneration Research, 21*(3), 887–907. https://doi.org/10.4103/NRR.NRR-D-24-01398
14. Micek, M. A., Arndt, B., Tuan, W. J., Trowbridge, E., Dean, S. M., Lochner, J., ... Pandhi, N. (2020). Physician burnout and timing of electronic health record use. *ACI Open, 4*(1), e1–e8. https://doi.org/10.1055/s-0039-3401815
15. Rahrisch, A., Braun, J., Ganter, M. T., Roche, T. R., Spahn, D. R., Tscholl, D. W., & Kaserer, A. (2026). Generative AI in preanesthetic consultations: Effects on efficiency, documentation workload, quality, and physician-patient interaction: A simulation trial. *Journal of Clinical Anesthesia, 110*, Article 112163. https://doi.org/10.1016/j.jclinane.2026.112163

16. Sezgin, E. (2023). Artificial intelligence in healthcare: Complementing, not replacing, doctors and healthcare providers. *Digital Health, 9*. https://doi.org/10.1177/20552076231186520
17. Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., ... Natarajan, P. (2023). Large language models encode clinical knowledge. *Nature, 620*(7972), 172–180. https://doi.org/10.1038/s41586-023-06291-2
18. Sinsky, C., Colligan, L., Li, L., Prgomet, M., Reynolds, S., Goeders, L., ... Blatz, A. (2016). Allocation of physician time in ambulatory practice: A time and motion study in 4 specialties. *Annals of Internal Medicine, 165*(11), 753–760. https://doi.org/10.7326/M16-0961
19. Soganciouglu, E., Calli, E., van Ginneken, B., van Leeuwen, K. G., & Murphy, K. (2021). Deep learning for chest X-ray analysis: A survey. *Medical Image Analysis, 72*, Article 102125. https://doi.org/10.1016/j.media.2021.102125
20. Watson, G. L., Staples, G., Carver, R., Bhargava, A., López-Espina, C., Schmalz, L., ... Antkowiak, P. S. (2026). Interpretability of an FDA-authorized AI/ML sepsis diagnostic tool improved by SHAP values. *JAMIA Open, 9*(1), Article ooag020. https://doi.org/10.1093/jamiaopen/ooag020
21. Yasaka, K., & Abe, O. (2018). Deep learning and artificial intelligence in radiology: Current applications and future directions. *PLOS Medicine, 15*(11), Article e1002707. https://doi.org/10.1371/journal.pmed.1002707