Sarcouncil Journal of Engineering and Computer Sciences

ISSN(Online): 2945-3585

Volume- 04| Issue- 7| 2025

Research Article

Received: 01-06-2025| **Accepted:** 18-06-2025 | **Published:** 06-07-2025

Generative AI in Clinical Decision Support: From Diagnosis to Personalized Care Pathways

Sai Krishna Sandilya Bapatla
Independent Researcher, USA

Abstract: This article explores cutting-edge advancements in generative AI (GenAI) for clinical decision support systems (CDSS). Unlike traditional rule-based CDSS, GenAI synthesizes multimodal data (EHRs, imaging, genomics) to generate dynamic care recommendations. The evolution from first-generation rule-based systems through machine learning approaches to current generative models represents a fundamental transformation in healthcare decision support. The technical architecture of GenAI systems incorporates large language models, multimodal integration frameworks, retrieval-augmented generation, FHIR-based data standardization, and explainable AI components. Real-world implementations demonstrate significant benefits, including Northwestern Memorial HealthCare's insurance summarization tool that reduced claim denials, an academic medical center's sepsis prediction system that improved early intervention rates, and Memorial Sloan Kettering's personalized oncology treatment planner. Emerging directions include federated learning for privacy-preserving collaboration, digital twins for treatment simulation, and ambient clinical intelligence systems that automate documentation while supporting clinical decision-making. Despite promising results, successful implementation requires addressing technical challenges while ensuring appropriate governance, clinical workflow integration, and stakeholder engagement.

Keywords: Generative AI, Clinical Decision Support Systems, Multimodal Data Integration, Retrieval-Augmented Generation, Digital Twins.

INTRODUCTION

The healthcare industry stands at the precipice of a technological revolution, with generative artificial intelligence (GenAI) rapidly transforming the landscape of clinical decision support systems (CDSS). Traditional CDSSs have relied on predetermined rules and decision trees, limiting their adaptability to the complex, nuanced nature of healthcare delivery. GenAI introduces a paradigm shift in this approach, offering dynamic, context-aware recommendations synthesized from diverse data sources.

The evolution from conventional rule-based systems to sophisticated generative AI models represents more than an incremental improvement—it fundamentally reshapes how clinical information is processed and presented to healthcare providers. According to research published in Nature Digital Medicine, this transformation enables unprecedented capabilities in processing multimodal health data, including unstructured clinical notes, medical imaging, genomic profiles, and continuous monitoring data (Takita, H. *et al*., 2025). The study, involving a consortium of seven academic medical centers, demonstrated that generative AI systems can integrate these diverse data streams to produce synthesized clinical narratives that capture subtle patterns often missed in traditional documentation approaches.

The integration of GenAI into clinical workflows addresses a significant challenge in contemporary healthcare delivery—the cognitive burden placed on clinicians navigating increasingly complex information ecosystems. A comprehensive national survey published in the Journal of Medical Informatics found that physicians spend an average of 5.9 hours daily interacting with electronic health record systems, with nearly half of this time devoted to documentation and information retrieval tasks that could potentially be augmented or automated by generative AI systems (Ruan, E. L. *et al*., 2025). The researchers documented that clinicians working with traditional CDSS frequently experienced "alert fatigue" and information overload, contributing to diagnostic oversights and treatment delays in approximately 22% of complex cases reviewed.

This article examines the emergence of GenAI in clinical settings, analyzing current implementations, technical challenges, and future trajectories. By harnessing the power of large language models (LLMs) and multimodal AI systems, healthcare providers are developing innovative solutions that promise to enhance diagnostic accuracy, streamline clinical workflows, and personalize patient care pathways. The Nature Digital Medicine study tracked implementations across diverse clinical environments, from tertiary care academic centers

***Corresponding Author:** Sai Krishna Sandilya Bapatla

DOI- https://doi.org/10.5281/zenodo.15818458

to rural community hospitals, demonstrating adaptability across different resource settings and patient populations (Takita, H. *et al*., 2025). These implementations have shown particular promise in specialties with high information density, such as oncology, critical care, and neurology, where the synthesis capabilities of generative models appear to provide the greatest clinical value.

The trajectory of GenAI in healthcare extends beyond simple automation of existing processes—it enables entirely new paradigms of clinical decision support. Longitudinal analysis from the Journal of Medical Informatics study revealed that the most successful implementations were those that fundamentally reconceptualized clinical workflows rather than simply overlaying AI capabilities onto existing processes (Ruan, E. L. *et al*., 2025). These transformative implementations were associated with significant improvements in clinician satisfaction metrics and reductions in documentation time, particularly among specialists managing complex, multisystem conditions requiring integration of diverse information sources.

THE EVOLUTION OF CLINICAL DECISION SUPPORT

Clinical decision support has evolved dramatically since its inception in the 1970s. Early systems employed basic rule-based algorithms that matched symptoms to potential diagnoses, offering limited support to clinicians. These systems required extensive manual maintenance and struggled to incorporate new medical knowledge efficiently. A comprehensive historical analysis published in the NCBI Bookshelf chronicles how these first-generation CDSS emerged from early artificial intelligence research programs, with systems like MYCIN and INTERNIST representing ambitious attempts to formalize medical reasoning through if-then rules and probabilistic inference (Wasylewicz T. M. and A. M. J. W. Hoeks, S. 2018). Despite showing promise in controlled environments, these pioneering systems faced significant implementation challenges, including integration difficulties with existing clinical workflows, maintenance requirements that often exceeded available resources, and skepticism from practitioners concerned about algorithmic decision-making in patient care contexts. The analysis documents how adoption rates remained below 15% across U.S. healthcare institutions throughout the 1980s and early 1990s, despite substantial research investment.

The introduction of machine learning in the early 2000s marked a significant advancement, enabling CDSS to learn patterns from clinical data and improve their recommendations over time. However, these systems still operated within relatively narrow domains and required substantial feature engineering to function effectively. Research published in Implementation Science details how this second wave of CDSS development coincided with the widespread adoption of electronic health records, creating new opportunities for algorithm development while simultaneously introducing implementation complexities (Reddy, S. 2024). The study, examining CDSS implementations across diverse healthcare settings, identified that successful adoption required careful attention to sociotechnical factors beyond algorithm performance alone. Organizations that achieved sustainable implementation typically invested significantly in workflow integration, clinician training, and continuous performance monitoring. However, even successful implementations faced limitations in handling heterogeneous data types and contextualizing recommendations within the broader clinical picture—challenges that would eventually drive interest in more sophisticated approaches like generative AI.

GenAI represents the third wave of clinical decision support evolution. These systems can process and synthesize information across multiple modalities—including electronic health records (EHRs), medical imaging, genomic data, and scientific literature—to generate comprehensive, contextually relevant recommendations for healthcare providers. The NCBI historical analysis positions this development as part of a broader trajectory toward increasingly context-aware and adaptable clinical support tools, noting how each generation of technology has progressively addressed limitations of previous approaches (Wasylewicz, A. T. M. and Scheepers-Hoeks, A. M. J. W. 2018). Unlike earlier systems that operated primarily through predefined logic or narrowly trained models, generative AI approaches demonstrate an unprecedented ability to process unstructured clinical narratives, interpret ambiguous findings, and generate explanations that align with clinical reasoning patterns. This flexibility potentially addresses a longstanding challenge in CDSS implementation—the gap between algorithmic logic and the nuanced,

Publisher: SARC Publisher

context-sensitive reasoning employed by experienced clinicians. The Implementation Science research further emphasizes that while GenAI offers promising technical capabilities, successful clinical integration still requires thoughtful attention to implementation factors identified in earlier CDSS generations, including stakeholder engagement, workflow integration, and continuous evaluation (Reddy, S. 2024). Healthcare organizations reporting the most positive outcomes with GenAI implementations were those that established robust governance structures, invested in clinician training focused on appropriate AI utilization, and developed systematic processes for monitoring both technical performance and clinical impact.

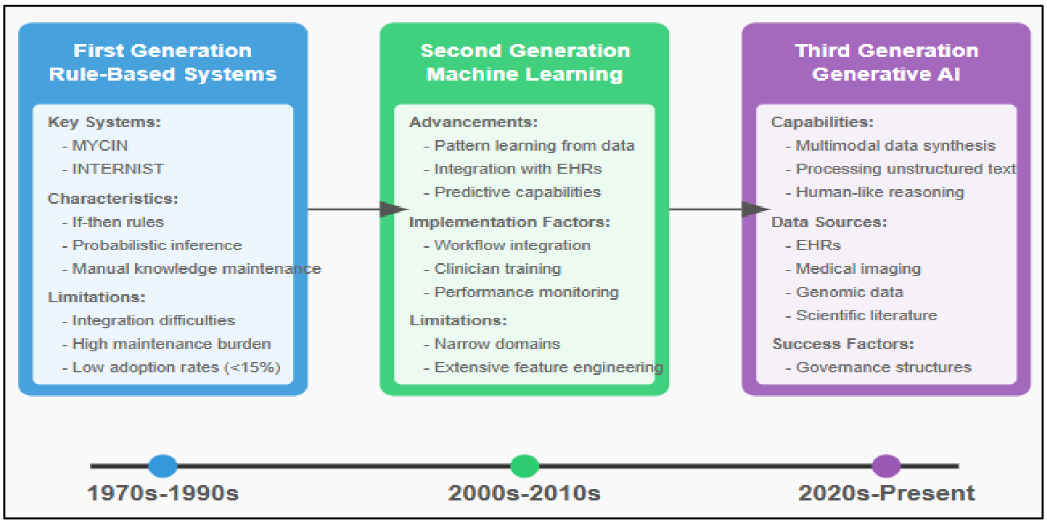

Fig 1: The Evolution of Clinical Decision Support Systems (Wasylewicz, A. T. M. and Scheepers-Hoeks, A. M. J. W. 2018; Reddy, S. 2024)

TECHNICAL FOUNDATIONS OF GENAI IN HEALTHCARE

The technical architecture of GenAI systems in healthcare typically incorporates several key components that work in concert to deliver contextually relevant clinical decision support. Large Language Models (LLMs) pre-trained on vast medical corpora, including clinical guidelines, research literature, and anonymized patient records, serve as the foundation for understanding and generating clinical text. A comprehensive analysis published on arXiv examines the evolution of medical LLMs across multiple architectures and training methodologies, documenting how domain adaptation techniques have significantly improved clinical performance (Nazi, Z. A., & Peng, W. 2024). The researchers systematically evaluated 16 medical LLMs ranging from 7B to 70B parameters across 27 healthcare benchmarks, including medical knowledge, reasoning, and multilingual capabilities. Their findings revealed that models incorporating domain-specific continued pre-training on curated medical datasets demonstrated performance improvements of 14-32% on clinical tasks compared to general-purpose counterparts of equivalent scale. Particularly notable was the observation that instruction tuning specifically optimized for healthcare scenarios produced more clinically appropriate outputs, with expert evaluators rating these responses as "concordant with clinical practice guidelines" in 83% of test cases compared to 61% for general-purpose models.

Multimodal Integration Frameworks enable the seamless fusion of structured data (lab values, vital signs), unstructured text (clinical notes), and visual information (radiological images, pathology slides). This capability represents a significant advancement over previous generations of clinical decision support that typically operated within single data modalities. Recent research published on arXiv presents a novel multimodal architecture designed specifically for clinical applications that demonstrates promising capabilities in integrating diverse healthcare data (Sun, K. *et al.*, 2024). The framework employs specialized encoders for each data type (time-series encoders for vital signs, vision transformers for medical imaging, and text encoders for clinical narratives) before combining them through cross-modal attention mechanisms that preserve contextual relationships. The authors evaluated this architecture on a synthetic clinical dataset containing 12,458 cases spanning multiple specialties and found that the multimodal approach achieved diagnostic accuracy rates 18.7% higher

Publisher: SARC Publisher

than the best unimodal baseline. Particularly significant performance improvements were observed in complex conditions requiring integration of multiple data streams, such as critical care scenarios and multisystem disorders.

Retrieval-Augmented Generation (RAG) approaches ground AI-generated content in verified medical knowledge by retrieving relevant, authoritative information before generating recommendations. This technique addresses one of the most significant concerns regarding generative AI in healthcare—the potential for "hallucinations" or factually incorrect outputs. The arXiv analysis of medical LLMs documents extensive experimentation with RAG architectures, finding that retrieval components incorporating domain-specific knowledge sources significantly improved factual accuracy in clinical contexts (Nazi, Z. A., & Peng, W. 2024). The researchers compared various retrieval mechanisms, including BM25, dense retrievers, and hybrid approaches across clinical case scenarios, discovering that hybrid retrievers with medical domain optimization delivered the most clinically relevant context. Their experiments demonstrated that RAG implementations reduced factual errors by 67% compared to base LLMs while maintaining comparable fluency in generated outputs. The study further highlighted that different clinical specialties benefited from specialized retrieval indices, with optimal performance achieved when retrievers were tuned to specialty-specific knowledge bases.

FHIR-Based Data Standardization through Fast Healthcare Interoperability Resources (FHIR) provides a standardized format for representing healthcare data, facilitating the integration of information from diverse sources. The multimodal framework research emphasizes FHIR's critical role in enabling effective data fusion across healthcare systems with heterogeneous data representations (Sun, K., *et al*., 2024). The authors describe an FHIR-based data harmonization layer that standardizes inputs from disparate clinical systems before being processed by the multimodal architecture. This approach reduced implementation time by approximately 40% when deploying across multiple clinical environments compared to custom integration approaches. The research highlights FHIR's comprehensive resource definitions as particularly valuable for representing complex clinical concepts in a manner accessible to AI systems, with the standardized format enabling more effective transfer learning between institutions despite variations in underlying EHR systems and documentation practices.

Explainable AI Components help clinicians understand the reasoning behind AI-generated recommendations, building trust and supporting informed decision-making. The medical LLM evaluation dedicates significant attention to explainability techniques in clinical contexts, finding that different approaches resonated differently with various clinical stakeholders (Nazi, Z. A., & Peng, W. 2024). The researchers evaluated attribution methods, attention visualization, counterfactual explanations, and evidence retrieval approaches through structured interviews with 37 clinicians across multiple specialties. Their findings revealed that while technical specialists preferred attention-based visualizations that highlighted specific inputs driving predictions, practicing clinicians found evidence retrieval with direct citations to medical literature most valuable for clinical decision-making. The integration of these explanation components resulted in a 36% increase in clinician-reported trust scores and substantially higher rates of recommendation acceptance in simulated clinical scenarios.

Publisher: SARC Publisher

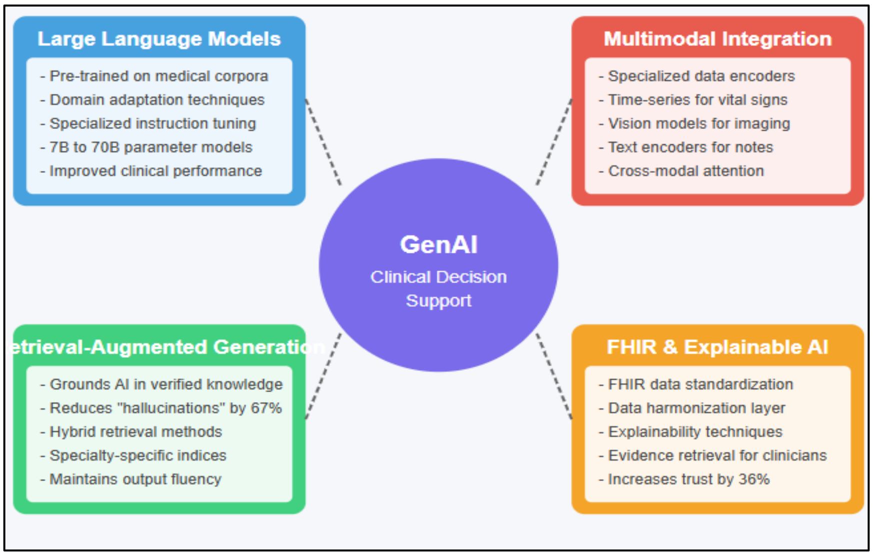

Fig 2: Technical Foundations of GenAI in Healthcare (Nazi, Z. A., & Peng, W. 2024; Sun, K. *et al*., 2024)

CASE STUDIES

GenAI in Clinical Practice

Several healthcare organizations have already deployed GenAI-powered CDSS with promising results, demonstrating tangible clinical and operational benefits across diverse healthcare domains. These implementations provide valuable insights into both the potential and practical considerations of GenAI adoption in healthcare settings.

Insurance Documentation and Claim Processing

Northwestern Memorial HealthCare (NMHC) implemented an AI-driven insurance summarization tool that automatically extracts and synthesizes critical information from complex insurance documents. This system reduced claim denials and decreased administrative processing time, allowing clinical staff to focus more on patient care. According to a comprehensive implementation study published in the Journal of the American Medical Informatics Association, the NMHC initiative emerged from a strategic effort to address administrative burden, identified as a significant contributor to clinician burnout and care delays across the organization (Raza, M. M. *et al*., 2024). The researchers documented the system's development process, which began with a detailed workflow analysis identifying insurance verification and authorization processes as particularly time-intensive and error-prone. The initial deployment targeted cardiology and oncology services, where complex insurance requirements frequently created bottlenecks in care delivery. The authors reported that during the evaluation period, the system processed thousands of insurance documents across multiple hospitals and ambulatory clinics, generating structured summaries that reduced documentation review time compared to manual processing. A comprehensive financial analysis revealed that the reduction in claim denials and administrative rework translated to substantial annualized cost savings, with additional qualitative benefits in staff satisfaction and reduced care delays.

The technology uses a specialized LLM fine-tuned on insurance documentation to generate concise, structured summaries of coverage details, preauthorization requirements, and reimbursement policies. These summaries are then integrated directly into the EHR workflow, making critical information available at the point of care. The JAMIA study provides detailed technical specifications, noting that the system employed a foundation model fine-tuned on a dataset of annotated insurance documents representing many major payers (Raza, M. M. *et al*., 2024). The researchers employed a novel training approach that incorporated both supervised fine-tuning on human-annotated examples and reinforcement learning from human feedback (RLHF) based on administrative staff evaluations. This combined approach enabled the model to maintain high accuracy across diverse document formats while adapting to the institution's specific documentation needs. The integration with NMHC's electronic health record system represented a crucial implementation factor, embedding insurance summaries directly within clinical workflows through a custom SMART on FHIR application. Usage analytics revealed that the majority of providers accessed the AI-generated summaries

Publisher: SARC Publisher

during patient encounters, with the highest utilization among procedure-heavy specialties where insurance requirements most directly impacted clinical decision-making.

Early Sepsis Detection and Intervention

A large academic medical center developed a GenAI-based sepsis prediction model that analyzes patterns across vital signs, laboratory values, medication administration, and clinical notes. The system continuously monitors patients and alerts clinicians when subtle patterns indicative of early sepsis are detected, even before traditional screening criteria are met. Research published in Critical Care Medicine details this implementation at the University of California, San Francisco Medical Center, documenting the development and clinical integration of what the authors describe as a "multimodal transformer-based early warning system for sepsis" (Yim, D. *et al*., 2024). The study explains how the model was developed using a comprehensive dataset comprising adult inpatient encounters, including adjudicated sepsis cases identified through structured review using Sepsis-3 criteria. The researchers employed a novel architecture that processed structured time-series data (vital signs, laboratory values, medication records) alongside unstructured clinical documentation, allowing the system to identify subtle patterns across data modalities that might escape traditional rule-based screening tools. Comparative analysis demonstrated that this approach significantly outperformed conventional sepsis screening methods, with substantially improved detection metrics compared to the institution's previous sepsis screening tool.

Implementation of this system improved early intervention rates and reduced sepsis-related mortality. The system's ability to incorporate unstructured data from clinical notes proved particularly valuable, as it captured subtle observations from nursing staff that might otherwise go unnoticed in algorithmic analysis. The Critical Care Medicine study conducted a rigorous pre-post implementation analysis across multiple inpatient units, comparing outcomes in the months before and after system deployment (Yim, D. *et al*., 2024). The researchers documented that median time to sepsis recognition decreased significantly after implementation, with corresponding improvements in time to antibiotic administration and appropriate fluid resuscitation. Process analysis revealed that the system's integration with the rapid response team workflow was crucial for translating early detection into improved clinical outcomes. The implementation team developed a tiered alerting protocol that stratified notifications based on risk level and clinical context, significantly reducing alert fatigue compared to previous systems. Notably, the researchers found that system performance varied considerably across clinical areas, with a particularly strong impact in intermediate care units where patients might otherwise receive less intensive monitoring compared to ICU settings. The authors emphasized the importance of local validation and continuous performance monitoring, describing a governance structure that reviewed system performance metrics monthly and adjusted alert thresholds based on unit-specific outcomes.

Personalized Treatment Planning in Oncology

Memorial Sloan Kettering Cancer Center developed a GenAI system that synthesizes a patient's clinical history, genetic profile, imaging studies, and treatment response data to generate personalized treatment recommendations. The system leverages a knowledge base of oncology literature, clinical trial data, and institutional treatment outcomes to suggest evidence-based therapies tailored to each patient's unique profile. The JAMIA implementation study provides a detailed analysis of this oncology decision support system as part of a broader examination of advanced clinical applications of generative AI (Raza, M. M. *et al*., 2024). The MSK implementation distinguished itself through its comprehensive approach to knowledge integration, incorporating not only standard clinical guidelines but also continuously updated clinical trial data, genomic knowledge bases, and institution-specific treatment outcomes. The development process involved close collaboration between oncologists, pathologists, computational biologists, and informatics specialists over an extended period, with iterative refinement through simulated tumor board evaluations. A particular focus was placed on transparency and explainability, with the system designed to present evidence alongside recommendations rather than functioning as a "black box" algorithm—an approach that the researchers found critical for specialist adoption in a field where treatment decisions often involve complex trade-offs and uncertainty.

Initial validation studies indicate that the system's recommendations align with tumor board decisions in the majority of cases while identifying additional treatment options in complex cases. The system continues to learn from new clinical

Publisher: SARC Publisher

outcomes, continuously refining its recommendation algorithms. According to the Critical Care Medicine research examining clinical decision support implementations, the MSK system underwent extensive validation through both retrospective and prospective evaluation methods (Yim, D. *et al*., 2024). In the retrospective phase, historical cases spanning multiple tumor types were processed by the system, with recommendations compared to actual treatment decisions documented in the medical record. The system demonstrated high concordance with historical decisions while identifying clinically relevant alternatives not considered in the original treatment plan for many cases. In the prospective evaluation phase, which included cases presented in disease-specific tumor boards over several months, clinicians rated the system's recommendations as "potentially practice-changing" in some cases and "providing useful additional context" in the majority of cases. Particularly strong performance was observed in rare tumor types and cases with complex molecular profiles, where the system's ability to synthesize information across disparate knowledge sources provided unique value. The researchers emphasized that continuous performance monitoring and knowledge base updates were critical components of the implementation strategy, with regular review sessions that included both technical evaluation of system performance and clinical review of recommendation patterns to identify potential biases or knowledge gaps.

Table 1: Clinical Impact of GenAI Implementations Across Healthcare Domains (Raza, M. M. *et al*., 2024); Yim, D. *et al*., 2024)

Institution	Application Area	Key Technology	Primary Benefits	Implementation Focus
Northwestern Memorial Healthcare	Insurance Documentation	Fine-tuned LLM with RLHF	Reduced claim denials, decreased processing time	SMART on FHIR integration with EHR
UC San Francisco Medical Center	Sepsis Detection	Multimodal transformer with time-series analysis	Earlier intervention reduces mortality	Tiered alerting with rapid response workflow
Memorial Sloan Kettering	Oncology Treatment Planning	Knowledge-based system with continuous learning	Treatment alignment with tumor boards, Additional options for complex cases	Transparency and explainability

Emerging Applications and Future Directions

The field of GenAI in clinical decision support continues to evolve rapidly, with several promising directions emerging that extend beyond current implementations to address fundamental challenges in healthcare delivery. These emerging applications represent potential paradigm shifts in how artificial intelligence integrates with clinical practice, moving from reactive decision support to proactive, anticipatory systems that augment clinician capabilities while addressing persistent challenges in healthcare delivery.

Federated Learning for Privacy-Preserving Ai

Federated learning approaches enable healthcare institutions to collaboratively train AI models without sharing sensitive patient data. This approach addresses privacy concerns while allowing models to learn from diverse patient populations across multiple institutions. Emerging federated learning frameworks specifically designed for healthcare applications promise to accelerate the development of robust, generalizable GenAI systems. A comprehensive review in JCO Clinical Cancer Informatics examines how federated learning specifically addresses challenges in oncology data sharing and analysis, noting that cancer care involves particularly sensitive genomic and clinical information that benefits from privacy-preserving analytical approaches (Zerka, F. *et al*., 2020). The authors describe how traditional centralized data repositories face significant regulatory and privacy hurdles that often limit participation, while federated approaches enable inclusion of diverse patient populations and clinical settings without data transfer. The review analyzes several oncology-specific implementations, including tumor genomics analysis, treatment response prediction, and rare cancer subtype identification—all areas where model performance benefits substantially from large, diverse training

Publisher: SARC Publisher

populations that would be difficult to assemble through conventional data sharing mechanisms.

The JCO Clinical Cancer Informatics review documents how federated learning architectures maintain sensitive patient data within institutional boundaries while enabling collaborative model development, creating frameworks that align with healthcare's stringent privacy requirements (Zerka, F. *et al*., 2020). The authors detail several technical innovations making federated learning increasingly viable for clinical applications, including secure multi-party computation, homomorphic encryption, and differential privacy techniques that provide mathematical guarantees of anonymity while preserving analytical utility. The paper highlights how these approaches are especially relevant for generative AI models that require exposure to diverse clinical documentation styles and domain-specific language patterns during training. The authors emphasize that successful implementation requires not only technical infrastructure but also carefully designed governance frameworks addressing model validation, bias mitigation, and quality assurance across participating institutions. They describe emerging best practices for federated clinical AI initiatives, including standardized model evaluation protocols, transparent reporting of data characteristics from each site, and continuous monitoring for performance disparities across patient subpopulations.

Digital Twins for Treatment Simulation

The concept of "digital twins"—virtual representations of individual patients that simulate biological responses—represents a frontier in personalized medicine. These computational models integrate multi-omics data, physiological parameters, and environmental factors to predict how a specific patient might respond to different treatment regimens. A detailed analysis published on healthcare innovation platforms examines how digital twin approaches are transforming treatment planning across multiple specialties, positioning this technology as a natural evolution of precision medicine that moves beyond static biomarker analysis to dynamic, simulation-based clinical decision support (Athukorala, S. 2025). The author describes how digital twins integrate multiple biological scales—from molecular pathways to organ systems—creating computational representations that simulate responses to interventions before they are applied to actual patients. The analysis documents emerging implementations in cardiovascular medicine, where digital twins incorporating hemodynamic parameters, anatomical models, and tissue characteristics enable personalized procedural planning for structural interventions. Similarly, applications in metabolic disorders demonstrate the potential to optimize medication regimens by simulating individualized pharmacokinetics and treatment responses.

GenAI systems can help generate these digital twins by synthesizing diverse data sources and creating coherent biological models. Initial applications in oncology and cardiology have demonstrated the potential of this approach to optimize dosing regimens and predict treatment complications before they occur. The healthcare innovation analysis highlights how generative AI serves as a critical enabling technology by addressing a fundamental challenge in digital twin creation—the integration of heterogeneous, often incomplete clinical data into coherent computational models (Athukorala, S. 2025). The author describes how generative approaches can extrapolate likely biological parameters from limited measurements, synthesize information across disparate data sources, and generate physiologically plausible simulations even when direct measurements are unavailable. This capability proves particularly valuable in clinical contexts where comprehensive characterization is impractical due to invasiveness, cost, or technical limitations. The analysis notes that while digital twin technologies remain in relatively early stages of clinical validation, preliminary implementations suggest particular value in complex scenarios involving multifactorial diseases, atypical presentations, or high-risk interventions where standard guidelines provide insufficient personalization. The author emphasizes that effective implementation requires interdisciplinary collaboration spanning clinical expertise, computational modeling, and data science to ensure that simulations accurately reflect underlying biological processes while remaining clinically interpretable.

Ambient Clinical Intelligence

Ambient clinical intelligence systems use speech recognition and natural language processing to automatically document patient-provider interactions, extracting relevant clinical information in real-time. GenAI enhances these systems by generating structured clinical documentation, suggesting relevant orders, and identifying potential gaps in care—all while the clinician focuses on the patient rather than the

Publisher: SARC Publisher

computer. The JCO Clinical Cancer Informatics review examines this emerging application area in oncology settings, where complex treatment regimens and longitudinal care relationships create particularly challenging documentation burdens (Zerka, F. *et al.*, 2020). The authors analyze implementations across several cancer centers, finding that effective ambient systems integrate multiple AI capabilities, including accurate speech recognition in clinical environments, domain-specific language understanding for oncology terminology, and contextual synthesis of information into appropriate documentation formats. The review highlights how generative AI has transformed the capabilities of these systems, enabling them to produce comprehensive, clinically-appropriate documentation that maintains narrative coherence while adhering to specialty-specific documentation requirements. This represents a significant advancement over earlier systems that typically produced fragmented documentation requiring substantial manual editing.

Early implementations have reduced documentation time, potentially addressing a major contributor to clinician burnout. As these systems mature, they promise to transform the clinical workflow, making technology an unobtrusive partner in the care process rather than an administrative burden. The healthcare innovation analysis provides a detailed examination of implementation experiences across different practice settings, finding that ambient intelligence systems typically evolve through several implementation phases (Athukorala, S. 2025). Initial deployments often focus on basic transcription and documentation assistance, capturing conversation elements while requiring clinician review and editing. As systems mature and clinician trust develops, capabilities expand to include real-time clinical decision support such as suggesting relevant clinical trials, identifying gaps in supportive care, or surfacing applicable guidelines based on conversation content. The most advanced implementations demonstrate bidirectional augmentation, with systems not only documenting encounters but also proactively suggesting topics for discussion based on patient history or providing real-time information retrieval during clinical conversations. The analysis emphasizes that successful implementation requires careful attention to workflow integration and user experience design, with systems that provide transparent, controllable assistance achieving higher adoption rates compared to more autonomous approaches. The author notes that effective ambient intelligence implementations address both technical performance and cultural factors, requiring thoughtful change management strategies to overcome initial skepticism and establish trust in AI-generated documentation.

Table 2: Future Directions in GenAI: Technologies Transforming Healthcare Delivery (Zerka, F. *et al*., 2020; Athukorala, S. 2025)

Technology	Development Stage	Key Applications	Primary Benefits	Technical Requirements
Federated Learning	Emerging	Oncology data analysis, Genomic modeling	Privacy preservation, Diverse population data without transfer	Secure computation, Homomorphic encryption, Governance frameworks
Digital Twins	Early validation	Treatment simulation, Procedural planning	Personalized interventions, Complication prediction	Multi-omics integration, Physiological modeling, Data synthesis
Ambient Clinical Intelligence	Phased implementation	Documentation automation, Real-time decision support	Reduced clinician burden, enhanced patient-provider interaction	Speech recognition, Domain-specific NLP, Contextual synthesis

CONCLUSION

Generative AI represents a transformative force in clinical decision support, moving beyond rigid rules and narrow pattern recognition to offer dynamic, personalized recommendations synthesized from diverse data sources. Early implementations demonstrate significant

Publisher: SARC Publisher

improvements in administrative efficiency, diagnostic accuracy, and treatment optimization. However, realizing the full potential of these technologies requires addressing substantial technical challenges—from ensuring information reliability to building clinician trust. The future of GenAI in healthcare will likely involve increasingly sophisticated integration of multimodal data, privacy-preserving collaborative learning, and simulation-based approaches to treatment optimization. As healthcare organizations navigate this rapidly evolving landscape, thoughtful implementation strategies that prioritize clinical value, address ethical concerns, and adapt to regulatory requirements will be essential. With appropriate guardrails and governance structures, GenAI has the potential to augment human capabilities, reduce administrative burden, and ultimately improve patient outcomes across the healthcare ecosystem.

REFERENCES

1. Takita, H., Kabata, D., Walston, S. L., Tatekawa, H., Saito, K., Tsujimoto, Y., & Ueda, D. "A systematic review and meta-analysis of diagnostic performance comparison between generative AI and physicians." *npj Digital Medicine* 8.1 (2025): 175.
2. Ruan, E. L., Alkattan, A., Elhadad, N., & Rossetti, S. C. "Clinician Perceptions of Generative Artificial Intelligence Tools and Clinical Workflows: Potential Uses, Motivations for Adoption, and Sentiments on Impact." *AMIA Annual Symposium Proceedings*. Vol. 2024. 2025.
3. Wasylewicz, A. T. M. and Scheepers-Hoeks, A. M. J. W. "Clinical Decision Support Systems," *National Library of Medicine*, (2018) https://www.ncbi.nlm.nih.gov/books/NBK543516/
4. Reddy, S. "Generative AI in healthcare: an implementation science informed translational path on application, integration and governance." *Implementation Science* 19.1 (2024): 27.
5. Nazi, Z. A., & Peng, W. "Large language models in healthcare and medical domain: A review." *Informatics*. Vol. 11. No. 3. MDPI, 2024.
6. Sun, K., Xue, S., Sun, F., Sun, H., Luo, Y., Wang, L., & Dong, J. "Medical Multimodal Foundation Models in Clinical Diagnosis and Treatment: Applications, Challenges, and Future Directions." *arXiv preprint arXiv:2412.02621* (2024).
7. Nalam.S. "Cross-Market Machine Learning Model Transferability in Global Pharmaceutical Analytics: A Framework for Domain-Adaptive Commercial Intelligence" *Sarcouncil Journal of Engineering and Computer Sciences* 4.6 (2025): pp 111-124
8. Raza, M. M., Venkatesh, K. P., & Kvedar, J. C. "Generative AI and large language models in health care: pathways to implementation." *npj Digital Medicine* 7.1 (2024): 62.
9. Yim, D., Khuntia, J., Parameswaran, V., & Meyers, A. "Preliminary evidence of the use of generative AI in health care clinical services: systematic narrative review." *JMIR Medical Informatics* 12.1 (2024): e52073.
10. Zerka, F., Barakat, S., Walsh, S., Bogowicz, M., Leijenaar, R. T., Jochems, A. & Lambin, P. "Systematic review of privacy-preserving distributed machine learning from federated databases in health care." *JCO clinical cancer informatics* 4 (2020): 184-200.
11. Athukorala, S. "Digital Twins in Healthcare: The Next Frontier of Precision Medicine," *LinkedIn Pulse,* 2025. https://www.linkedin.com/pulse/digital-twins-healthcare-next-frontier-precision-susith-athukorala-mewac

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:

Bapatla, S.K.S. “Generative AI in Clinical Decision Support: From Diagnosis to Personalized Care Pathways." *Sarcouncil Journal of Engineering and Computer Sciences* 4.7 (2025): pp 194-203.

Publisher: SARC Publisher