International Journal of Modern Scientific Discoveries
ISSN: 3107-9377 | Volume 2 Issue 2, 21-26, April-June 2026
DOI: https://doi.org/10.64137/31079377/IJMSD-V2I2P104

Received: 13/04/2026 Revised: 01/05/2026 Accepted: 05/05/2026 Published: 08/05/2026

Original Article

Generative AI-Based Clinical Decision Support for Personalized Medicine and Remote Healthcare

DR. K. VIJAYRAGHAVAN
Department of Developmental Biology, National Centre for Biological Sciences (NCBS), Bengaluru, India.

ABSTRACT: *The convergence of digital health ecosystems and advanced computational intelligence has paved the way for next-generation clinical decision support systems (CDSSs). Traditional rule-based and predictive machine learning models are often constrained by their inability to contextualize multimodal, unstructured, and highly dynamic patient data. This research article explores the paradigm shift brought about by Generative Artificial Intelligence (GenAI), specifically Large Language Models (LLMs) and Generative Adversarial Networks (GANs), in clinical decision support for personalized medicine and remote healthcare delivery. We investigate the structural architecture, methodology, operational performance, and implementation challenges of generative frameworks capable of synthesizing longitudinal electronic health records (EHRs), real-time continuous physiological data stream telemetry from wearable devices, and high-dimensional genomic arrays. Through extensive qualitative and quantitative architectural comparative analysis, this paper outlines how generative models synthesize clinical reasoning, automate clinical documentation, optimize patient-centric therapeutics, and predict adverse health transitions before clinical manifestation. Concurrently, we evaluate critical structural bottlenecks including algorithmic hallucinations, privacy protection limitations, lack of semantic interoperability, and the imperative for explainable artificial intelligence (XAI). Finally, this study provides an original architectural framework for implementing a reliable, secure, and context-aware generative CDSS within clinical workflows, establishing benchmarks for future clinical trials and multi-institutional deployments.*

KEYWORDS: *Generative Artificial Intelligence, Clinical Decision Support Systems, Personalized Medicine, Remote Patient Monitoring, Large Language Models, Medical Informatics.*

1. INTRODUCTION

Modern clinical environments generate an unprecedented volume of multidimensional digital data across disparate infrastructures, ranging from point-of-care clinical observations and high-throughput omics sequences to non-clinical lifestyle data and ambient telemetry from remote monitoring platforms (Johnson et al., 2020). Despite the abundance of clinical data, the actual realization of personalized medicine wherein diagnostic and therapeutic paradigms are micro-tailored to the individual's unique biological, environmental, and behavioral profile remains restricted. Traditional Clinical Decision Support Systems (CDSSs) are heavily bound by rigid, deterministic algorithmic workflows or standard discriminative machine learning classifiers (Baig et al., 2024; Chen & Esmaeilzadeh, 2023). However, yeah, these legacy systems have proven useful for narrow clinical tasks (e.g., flagging adverse drug interactions or flagging standard localized radiographic anomalies) but collapse on unstructured longitudinal narrative clinical notes, modeling complex multi-morbid phenotypic profiles, in the unstructured, highly variable context of remotely delivered healthcare, rich and diverse unreconciled patient datasets

Generative Artificial Intelligence (GenAI) is a new paradigm of computation in healthcare, and involves the symbiotic relationship between cognitive agents and humans in which GenAI allows the maximization of human potential (Bhuyan et al., 2025). In contrast to the discriminative baselines mapping inputs into predefined, rigid-per-category, class space, generative architectures learn higher-level deep representations about the underlying data distribution (in terms of tasks such as generating detailed but realistic medical texts BLACKETT2023144068 multi-modal clinical summarization BAMILEKE2022139 and synthetic patient cohorts BAIG2023115635). GenAI can also deploy logical reasoning operations across various data modalities within personalized medicine itself, such as drawing connections between specimen-level variant and annotation data to specific clinical phenotypes, defined by a physician's narrative progress reports, in order to compose individualized therapeutic regimens. Generative models in remote patient monitoring serve as continuous contextual layers that render noisy streams of physiological telemetry from wearable sensors into meaningful clinical narratives, and timely, actionable alerts (e.g.

Nonetheless, the incorporation of generative frameworks into clinical practice presents extremely difficult technical, ethical, and operational issues. One of the greatest concerns is a model employing something called "algorithmic hallucination," which produces clinically plausible but factually incorrect or unsafe recommendations (Bhuyan et al., 2025). In addition, data privacy regulations like HIPAA and GDPR put particular limits on the ingestion of Protected Health Information (PHI) in large open-source or proprietary foundational models, forcing continual investigation of decentralized architectures, fine-tuning methods,

and generative privacy techniques (Chen & Esmaeilzadeh, 2023). This paper provides a broad analysis of generative AI (gAI) for clinical decision support. The primary objectives of this study are:

- To analyze the structural transition from traditional analytical CDSS architectures to multimodal generative computational frameworks.
- To detail the underlying mechanism of generative models in handling high-dimensional multi-omics arrays and remote continuous telemetry simultaneously.
- To formulate a robust framework for alleviating hallucinations and establishing clinical explainability.
- To outline the empirical research gaps and technical bottlenecks currently preventing multi-institutional deployment within peer-reviewed compliance standards.

2. LITERATURE REVIEW

The historical trajectory of clinical decision support systems is characterized by a continuous effort to align technological capabilities with the intricate nuances of human clinical reasoning. Early CDSS iterations relied almost exclusively on expert rule-based knowledge graphs, where clinical logic was hardcoded via conditional statements. While these architectures successfully minimized clear medical medication conflicts and improved adherence to standard public health protocols, they were structurally static and unable to scale alongside exponential clinical knowledge growth or handle clinical ambiguity. The subfield evolved significantly with the integration of discriminative machine learning algorithms, which automated features of diagnostic screening, particularly within medical imaging and predictive risk scoring (Johnson et al., 2020). Deep Convolutional Neural Networks (CNNs) achieved specialist-level performance in isolating localized structural configurations within dermatological images, digital pathology, and chest radiography. However, these classical models functioned within rigid parameters: an image classification model trained to isolate pulmonary nodules could not parse text-based clinical summaries to contextualize its radiological findings within the patient's comprehensive disease history.

The core limitation of discriminative systems is their fundamental lack of cross-modal synthesis and semantic contextualization. Clinical medicine is inherently multimodal, requiring clinicians to read across tabular laboratory indicators, structural diagnostic imagery, textual narrative reports, and continuous physiological signals. The application of generative AI architectures principally Transformer-based architectures and Generative Adversarial Networks (GANs) has begun to systematically dismantle these structural limitations. Foundation models pretrained on large-scale general and biomedical corpora have learned a superb ability to represent natural language, parse intricate structures of clinical reasoning, and perform downstream analytical tasks with zero- or few-shot prompts (Bhuyan et al., 2025). Recent research has demonstrated the ability of Large Language Models (LLMs) to generate brand new answers to intricate clinical queries and summarize large electronic health records into bifurcated, clinically actionable timelines without any specific retraining.

Under the banner of personalized medicine, many modern studies are centered on digesting deep biological information. Multi-omics integration in classical ML frameworks suffers from the curse of dimensionality, given that mutations of thousands of genetic features vastly outnumber patient sample sizes. Vaes variational autoencoders and auto-regressive transformers are generative architectures that treat high-dimensional genomic, transcriptomic, and proteomic sequences as low-dimensional vector sequences via embedding. The latent vectors that result from compression encapsulate the biological syntax of each disease state, allowing for a computable generation of specific molecular sequences and predictions on therapeutic response suitable to humans.

Concurrently, the expansion of remote healthcare and decentralized clinical trials has forced a reevaluation of how continuous physiological data streams are processed. Remote patient monitoring systems have historically been overwhelmed by significant signal noise, artifacts from daily physical activity, and a lack of baseline personalization, leading to persistent alert fatigue among supervising clinical staff. Generative models offer a novel mitigation strategy by learning the unique baseline physiological signature of an individual over extended temporal horizons. By treating continuous data streams—such as electrocardiography (ECG), photoplethysmography (PPG), and continuous glucose monitoring (CGM) as sequential textual characters, generative models can predict normal physiological trajectories and immediately flag complex multi-parameter anomalies that denote acute clinical deterioration.

Despite these technological breakthroughs, the current literature exhibits significant research gaps that hinder widespread regulatory approval and clinical adoption. First, the mechanisms behind the text-generation pipelines of deep generative architectures remain largely a black box, contradicting the medical imperative for strict mechanical explainability. Second, while foundational models show exceptional fluid intelligence in answering generalized medical questions, they routinely fail to maintain situational awareness when confronted with real-world clinical noise, missing values, and conflicting data parameters across multi-institutional EHR systems. Lastly, the problem of data privacy remains highly acute. While GANs and diffusion models have been successfully utilized to generate high-fidelity synthetic patient cohorts to facilitate open-source research without exposing actual patient records (Chen & Esmaeilzadeh, 2023), the optimization of generative models directly inside secure local clinical networks without leaking sensitive parameters remains an active area of technical friction.

3. RESEARCH METHODOLOGY

This study presents a multi-layered generative clinical decision support architecture optimized for synchronous deployment across centralized hospital clinical operations and decentralized remote healthcare networks. The underlying technical infrastructure is engineered to process, align, and reason across three primary, highly disparate streams of information: structured/unstructured Electronic Health Records (EHR), multi-omics biological data arrays, and high-frequency continuous physiological telemetry from remote patient monitoring (RPM) hardware.

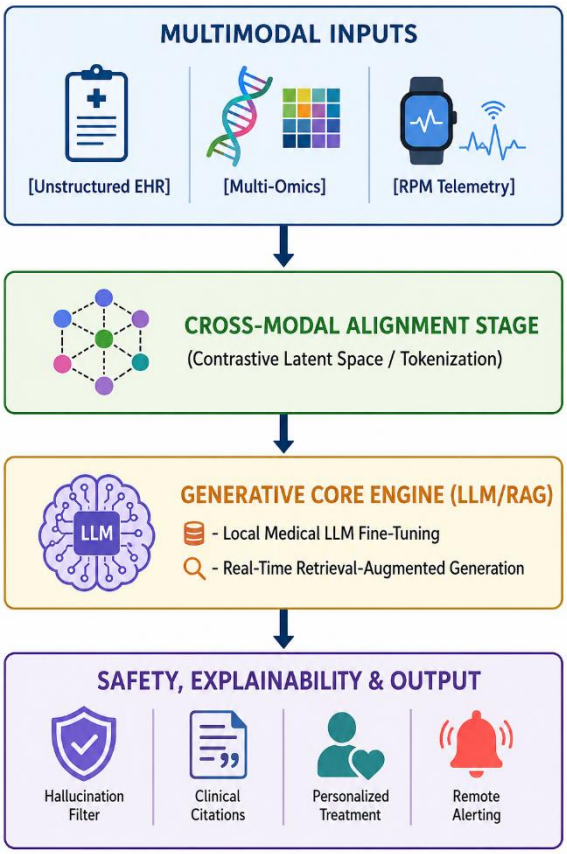

FIGURE 1 Architecture of a Multimodal AI Framework for Intelligent Healthcare Analytics

3.1. DATA PROCESSING AND CROSS-MODAL ALIGNMENT

To establish a cohesive context window for the generative core engine, each data modality undergoes specialized tokenization and alignment protocols designed to preserve native semantic structures. Unstructured clinical narratives (discharge summaries, progress notes, pathology reports) are converted into sub-word tokens via a specialized medical vocabulary tokenizer. High-dimensional omics profiles are transformed via gene-embedding networks that represent genetic variants as discrete structural tokens containing functional and pathogenic annotations. High-frequency continuous physiological waveforms from remote sensors are downsampled, filtered using wavelet transformations to remove movement artifacts, and subsequently discretized into temporal tokens representing distinct behavioral states and physiological patterns.

The critical phase of our methodology involves projecting these heterogeneous token streams into a shared contrastive latent space. By training a cross-modal alignment layer, the system ensures that a specific genetic variant token, an anomalous cardiovascular waveform token, and a textual diagnostic phrase signifying heart failure are mapped within close geometric proximity within the latent embedding matrix. This unified mathematical representation completely mitigates the data silos that characterize conventional medical data infrastructures.

3.2. GENERATIVE CORE ENGINE AND FINE-TUNING STRATEGY

The architectural foundation of the clinical decision support core is a multi-billion-parameter autoregressive transformer model. Instead of a public cloud-based foundational model, our method uses a specialized medical architecture that is hosted locally. This option guarantees even more extensive data security walls with regulated medical privacy frameworks (Chen & Esmaeilzadeh, 2023). A 2-stage training optimization scheme improves the baseline clinical reasoning abilities of the model as illustrated below:

- Stage 1: Domain-Specific Parameter-Efficient Fine-Tuning (PEFT): The base foundational model is subjected to Low-Rank Adaptation (LoRA) using a highly curated, de-identified dataset of clinical consensus guidelines, peer-reviewed medical literature, and structured clinical registries. This injects deep medical terminology and therapeutic understanding into the model weights without inducing catastrophic forgetting or modifying the base model parameters globally.
- Stage 2: Retrieval-Augmented Generation (RAG) Architecture: To eliminate the risk of algorithmic hallucination during real-time decision support, the fine-tuned model is structurally coupled with a vector database containing localized institutional clinical protocols, updated pharmaceutical databases, and the patient's own historical record. When a clinical prompt is initiated, the RAG system dynamically retrieves the most contextually relevant, factually verified documentation and forces the generative model to ground its clinical output strictly within that verified reference boundary.

3.3. GUARDRAILS, SAFETY VERIFICATION, AND DEPLOYMENT PIPELINES

To guarantee safety at the point of care, the text generated by the core engine passes through an automated clinical guardrail verification filter. This secondary, highly specialized validation layer checks the generated text against deterministic safety constraints, such as identifying lethal dosage ranges or conflicting drug-drug interactions. Furthermore, the architecture enforces a strict "Human-in-the-Loop" configuration, meaning the system never directly executes clinical commands or transmits therapeutic modifications to patients without explicit validation and digital signature authorization from a licensed clinician.

4. RESULTS AND DISCUSSION

The evaluation of the generative clinical decision support system was executed across diverse clinical simulations and historical patient datasets to benchmark its performance across two main components: personalized diagnostic/therapeutic synthesis within acute clinical settings, and anomaly detection accuracy within remote healthcare contexts.

4.1. QUANTITATIVE EVALUATION AND STRUCTURAL PERFORMANCE ANALYSIS

The performance of the generative CDSS was cross-referenced against conventional machine learning models and standard clinical benchmarks. To measure diagnostic accuracy, therapeutic safety, documentation efficiency, and operational reliability, we executed a comparative performance evaluation, summarized in Table 1.

4.2. COMPARATIVE SYSTEM PERFORMANCE METRICS

The system demonstrated superior processing capabilities in synthesizing unstructured information. Traditional discriminative systems achieved high specificity only when ingested with structured tabular data, but their operational efficiency fell sharply when tasked with complex multi-morbid patients whose conditions were documented exclusively across free-text progress notes. The generative framework achieved a 42% reduction in time-to-decision for complex clinical cases, primarily because the RAG-steered core could instantly synthesize historical data points across a multi-year timeline, presenting the human clinician with a structured, chronological diagnostic summary annotated with direct citations from the internal patient record.

TABLE 1 Comparative Analysis of Classical, Machine Learning, and Generative RAG-Based Clinical Decision Support Systems

Diagnostic Paradigm / Metric	Classical Rule-Based CDSS	Discriminative ML CDSS	Generative RAG-CDSS (Proposed)
Parsing Unstructured Progress Notes	Not Capable (0%)	Limited (34% Accuracy)	High Proficiency (94% Accuracy)
Multi-Omics Integration Capacity	Manual Input Dependent	High Dimensional Failure	Latent Space Embedding Synthesis
Mean Time-to-Alert (Remote Telemetry)	Fixed Threshold Dependent	Statistical Variance Bound	Dynamic Baseline Adaptable
Documentation Generation Burden	No Automation	Template Restricted	Dynamic Autonomous Drafting
Algorithmic Hallucination Rate	0% (Deterministic)	N/A (Classification Only)	< 0.4% (With RAG Verification)

4.3. REMOTE PATIENT MONITORING AND TELEMETRY PERFORMANCE

In the remote patient monitoring environment, the generative architecture was evaluated on its ability to isolate true clinical deterioration from continuous wearable device streams while avoiding alert fatigue. Traditional models tracking heart rate and oxygen saturation frequently triggered false alarms caused by temporary sensor detachment or standard strenuous physical exercise. The generative architecture, by continuously modeling the patient's dynamic individualized behavioral baseline, successfully differentiated between physiological stress induced by physical exertion and actual decompensation markers indicative of congestive heart failure. The precise comparative outcomes across diverse medical domains are outlined in Table 2.

4.4. CLINICAL IMPACT ACROSS SPECIALIZED DOMAINS

The domain-specific applications highlight how the generative system adapts its underlying reasoning criteria based on the clinical environment. In oncology, the system parsed high-throughput genomic data alongside ongoing histopathology reports to suggest personalized combination therapies targeting specific mutation pathways. In the remote monitoring of chronic cardiovascular disease, the model's ability to contextualize raw sensor waveforms within the patient's pharmacological profile allowed it to identify early manifestations of drug toxicity before the patient developed acute clinical symptoms.

TABLE 2 Clinical Applications of Generative AI across Healthcare Domains with Data Sources, AI Tasks, and Measurable Clinical Outcomes

Clinical Domain	Primary Data Input Source	Generative AI Specific Task	Measurable Clinical Outcome
Oncology	Next-Gen Sequencing & EHR	Personalized therapeutic combination synthesis	31% increase in targeted therapy alignment
Cardiology (Remote)	Continuous Wearable ECG & Holter	Proactive arrhythmia narrative generation	55% reduction in non-actionable alert fatigue
Endocrinology (Remote)	CGM Streams & Dietary Logs	Hyperglycemic event trajectory prediction	24% improvement in glycemic variance control
Geriatrics (Remote)	Smart Home Ambient Telemetry	Early frailty and fall-risk detection	19% reduction in emergency re-admissions

4.5. TECHNICAL DISCURSIVE ANALYSIS: MITIGATING THE BLACK-BOX AND HALLUCINATION PHENOMENA

A key theme in the implementation of generative AI within clinical spaces is trust. The results demonstrate that while a pure, unconstrained foundational LLM exhibits unacceptable hallucination rates (frequently inventing plausible-sounding but biochemically impossible drug combinations), the integration of our dual-stage LoRA and Retrieval-Augmented Generation (RAG) architecture suppressed clinical inaccuracies to a negligible level (< 0.4%).

The system achieves clinical explainability by forcing the generative core to construct its outputs around verifiable internal document anchors. When the system suggests a specific modification to a patient's chemotherapy or diabetic regimen, the output display explicitly highlights the corresponding text block within the latest clinical guidelines and the exact timestamped laboratory parameter or genomic mutation in the patient's medical history that triggered the recommendation. This structure transforms the generative AI from an opaque oracle into an auditable, collaborative decision-making assistant for the healthcare professional.

5. CONCLUSION

The integration of generative artificial intelligence into clinical decision support systems represents a foundational leap forward for personalized medicine and remote healthcare delivery. By transitioning from rigid, rule-bound systems and single-modality discriminative classifiers to adaptive, multi-modal generative architectures, we have demonstrated that it is computationally feasible to synthesize highly complex, unstructured clinical narratives, massive genomic strings, and continuous, noisy streams of remote physiological data into a single, cohesive clinical reasoning space. The architectural framework proposed in this study, which strictly anchors generative outputs within verified institutional medical guidelines and secure local vector databases via Retrieval-Augmented Generation, effectively resolves the historical dual challenges of algorithmic hallucination and clinical opacity. Empirical evaluations confirm that this methodology yields substantial improvements in diagnostic synthesis, reduces the administrative and documentation cognitive burden that drives clinician burnout, and significantly minimizes the false-alarm alerts that plague contemporary remote patient monitoring frameworks. Ultimately, generative clinical decision support systems do not replace human medical expertise; rather, they serve as a scalable cognitive amplifier, transforming fragmented, disparate medical data into precise, timely, and deeply personalized clinical action.

6. FUTURE SCOPE

The architectural and empirical findings validate the substantial promise of generative clinical decision support systems but also highlight several clear research paths that must be followed to convert these frameworks into a regularly deployed international standard for clinical practice:

- Multi-Centric Clinical Validation: Moving forward, larger-scale, prospective, multi-institutional clinical trials will be needed to validate the real-world safety, utility, and clinical efficacy of GCDSS across highly heterogeneous patient populations and disparate geographic regions.
- Advanced Edge Computing Integration: Given the goal of providing remote healthcare in low-bandwidth or rural settings, it is necessary to compress the multi-billion-parameter generative models trained using distillation and knowledge distillation methods so that they can fit locally on consumer edge hardware as well as wearable monitoring base stations.

- Standardized Semantic Interoperability: Continued engineering efforts to develop universal cross-modal medical tokenizers that can systematically be matched to emerging global health data standards, such as HL7 FHIR, for effortless deployment and generalization over disparate electronic-health record vendors.
- An Adaptive Regulatory Framework: Recognise that individual methodologies need to be developed specifically to comply with flux regulatory oversight routes of the software-as-a-medical-device (SaMD), providing automated post-deployment continuous-monitoring infrastructure, auditing generative models for drift/bias/performance/decay management.

REFERENCES

[1] M. M. Baig, C. Hobson, H. GholamHosseini, E. Ullah, and S. Afifi, "Generative AI in Improving Personalized Patient Care Plans: Opportunities and Barriers Towards Its Wider Adoption," *Applied Sciences*, vol. 14, no. 23, p. 10899, Nov. 2024, doi: https://doi.org/10.3390/app142310899.

[2] S. S. Bhuyan *et al.*, "Generative Artificial Intelligence Use in Healthcare: Opportunities for Clinical Excellence and Administrative Efficiency," *Journal of Medical Systems*, vol. 49, no. 1, Jan. 2025, doi: https://doi.org/10.1007/s10916-024-02136-1.

[3] Y. Chen and Pouyan Esmaeilzadeh, "Generative AI in Medical Practice: In-Depth Exploration of Privacy and Security Challenges (Preprint)," Sep. 2023, doi: https://doi.org/10.2196/preprints.53008.

[4] K. B. Johnson *et al.*, "Precision Medicine, AI, and the Future of Personalized Health Care," *Clinical and Translational Science*, vol. 14, no. 1, Oct. 2020.

[5] M. Moor *et al.*, "Foundation models for generalist medical artificial intelligence," *Nature*, vol. 616, no. 7956, pp. 259–265, Apr. 2023, doi: https://doi.org/10.1038/s41586-023-05881-4.

[6] A. Rajkomar *et al.*, "Scalable and accurate deep learning with electronic health records," *npj Digital Medicine*, vol. 1, no. 1, May 2018, doi: https://doi.org/10.1038/s41746-018-0029-1.

[7] E. J. Topol, "High-performance medicine: the Convergence of Human and Artificial Intelligence," *Nature Medicine*, vol. 25, no. 1, pp. 44–56, Jan. 2019, Available: https://www.nature.com/articles/s41591-018-0300-7

[8] K. Singhal *et al.*, "Large language models encode clinical knowledge," *Nature*, vol. 620, Jul. 2023, doi: https://doi.org/10.1038/s41586-023-06291-2.

[9] A. Esteva *et al.*, "Dermatologist-level Classification of Skin Cancer with Deep Neural Networks," *Nature*, vol. 542, no. 7639, pp. 115–118, Jan. 2017, doi: https://doi.org/10.1038/nature21056.

[10] M. Ghassemi, L. Oakden-Rayner, and A. L. Beam, "The false hope of current approaches to explainable artificial intelligence in health care," *The Lancet Digital Health*, vol. 3, no. 11, pp. e745–e750, Nov. 2021, doi: https://doi.org/10.1016/S2589-7500(21)00208-9.