IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154288

# Design of an Integrated Model for Neuro Symbolic GenAI and RAG Driven Personalized Cardiometabolic Care Sets

**Sagar Fernandes[1], Dr. Sangeeta Vhatkar[2]**

ME Student, Information Technology, Thakur College of Engineering and Technology, Kandivali East, India[1]

Guide, Information Technology, Thakur College of Engineering and Technology, Kandivali East, India[2]

**Abstract**: Increasing demand for clinically trustworthy, patient-specific decision systems has highlighted a gap between GenAI models and cardiometabolic care, where symptoms fluctuate gradually, data streams change hourly, and therapies must be ethical. Fragmented multimodal integration, shallow personalization, and confident but unsupported suggestions plague existing techniques. Due to these limitations, clinicians may generate technically impressive work that is hard to justify or apply to everyday hypertension and type 2 diabetes treatment. GenAI + RAG learns from heterogeneous data and follows causal, ethical, and evidentiary paths. It solves long-standing difficulties using a context-aware multimodal learning engine, retrieval-grounded generative reasoning, and five analytical methodologies in its validation pipeline. Salt sensitivity and nighttime heart-rate variability are examined for persistence using Multiscale Causal Uncertainty Stratification Analysis. Dynamic Ethical Constraint Verification Engine then assesses ethical restrictions' effects. Hierarchical Evidential RAG Stress-Test Simulation evaluates retrieval under conflicting clinical guidelines. A Persona-Calibrated Recommendation Consistency Audit checks for lifestyle or medication inconsistencies among patient archetypes. The last layer, Longitudinal Personalized Improvement Prediction Benchmark, predicts six-month A1C, systolic blood pressure, and marker trajectories from these suggestions. Causality, contextual explainability, and long-term clinical relevance appear to improve with the architecture. It may imply a generation of medical AI systems that reason modestly, trace their conclusions, and offer clinicians evidence-based pathways rather than isolated predictions.

**Keywords:** GenAI, RAG Systems, Cardiometabolic Care, Multimodal Learning, Explainable AI, Analysis.

## I. INTRODUCTION

The rapid rise of generative AI in medicine has created a promise-practice gap. Electronic health records, pharmacogenomic profiles, long-term wearable signals, and patient-reported behaviors are used in clinical settings. Most recent AI systems perceive these data streams as loosely connected chunks, not physiological contexts. Hypertension and type 2 diabetes reveal this mismatch because small daily changes in sleep, sodium intake, stress physiology, or medication adherence can drastically modify risk trajectories. Traditional machine-learning pipelines improve progressively but rarely provide clinicians with reasoning they can scrutinize, debate, or trust. Scientists are testing retrieval-augmented generation and multimodal transformers, but they cannot detect causal relationships, filter outputs through ethical limits, or assess stability when contradictory evidence exists. Most studies describe performance on static datasets rather than asking if suggestions would remain consistent when patient behavior changes or healthcare guidelines conflict. Clinical translation is impeded by the lack of deeper evaluation frameworks, even with strong predictive accuracy. This study's integrated approach treats logic, evidence, uncertainty, and ethics as first-class computing elements. Integration of multimodal fusion transformers, retrieval-grounded generative layers, and contextually adaptable neuro-symbolic reasoning framework. Cause stability, ethical propagation, evidentiary stress tolerance, persona-dependent recommendation consistency, and long-term clinical impact are assessed using companion validation suites. Increasing the conceptual alignment between patient context and machine reasoning may enable GenAI systems achieve clinically significant precision rather than surface-level customization in process.

## II. REVIEW OF EXISTING MODELS USED FOR PERSONALIZED TREATMENT ANALYSIS

Many research suggest that machine learning, deep learning, and large language models are becoming more important in clinical decision help, although data modalities, contextual awareness, and interpretability can limit them. Zeyauddin's NAFLD prediction and therapy optimization [1] shows how individualized feature engineering and supervised learning improve diagnostic granularity. But such systems rarely incorporate genetic, behavioral, and temporal features. Nieminen et al.'s ICF-linked graph structures study of low back pain functioning [2] found that graph-based clinical representation

 

ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁕ Peer-reviewed & Refereed journal ⁕ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154288

learning improves topological reasoning but not generative or retrieval-grounded inference. In highly specialized detection domains like fetal facial anomaly analysis [3] and multimodal MRI analytics for Alzheimer's diagnosis [4], classification performance improves but multimodal and causal reasoning structures do not. Angyal et al.'s assessment of adapted LLMs for cervical cancer workflows [5] found domain adaption potential but consistency and grounding issues. Predictive study like Nakano et al.'s ICU-acquired weakness [6].

Modern plastic surgery model domain comparisons [7] reveal that predictive accuracy does not guarantee contextual reliability or ethical traceability. 3D foot classification models for diabetic foot analysis [8] and digital-twin frameworks for precision oncology [9] provide increasingly realistic physiological embeddings, but they are still predictive or diagnostic engines rather than end-to-end reasoning systems integrating evidence, ethics, and causal feedback. Risk-prediction frameworks like COPD-associated depression modeling [10], hepatitis survival forecasting [11], and AI-driven delirium prediction [12] enhance statistical confidence but seldom include multimodal patient context or longitudinal progression modeling. Hassan and Elagamy's [13] personalised care recommendation engines provide patient-specific counsel using supervised learning workflows without generative reasoning or retrieval augmentation. Postoperative complications [14] and chemotherapy symptom management [15] prediction use discriminative modeling, not transparent, chain-based decision reconstruction. Machine learning algorithms accurately predict healthcare usability optimization [16] without explanation, context integration, or ethical scrutiny. This research suggests medical AI subdomain development and recurring challenges that warrant further study. Few systems integrate various longitudinal modalities into a single adaptive context, few utilize retrieval-grounded generative reasoning with integrated ethical propagation, and most are good at prediction but lack causal reasoning structures to handle developing cardiometabolic physiology The proposed model uses a fully integrated neuro-symbolic, multimodal, and evidence-conditioned reasoning engine to move beyond predictive tasks and create a clinical decision framework that accommodates uncertainty, ethical accountability, and coherence across diverse patient contexts.

## III. PROPOSED MODEL DESIGN ANALYSIS

The model's strongly integrated neuro-symbolic and retrieval-conditioned generative system charts cardiometabolic patients' physiological dynamics and causal, ethical, and evidential structures. A multimodal context encoder projects continuous wearable data, structured EHR variables, genetic embeddings, and behavioral logs into a unified latent manifold (Figure 1) sets. The encoder approximates a temporally aware context function $\Phi(t)$ that considers both short-term and long-term dependencies The integral formulation shows this merger Via equation 1,

$$z(t) = \int_0^t Ws\,xs(\tau)\,d\tau + We * xe + Wg * xg + \int_0^t Wb\,xb(\tau)\,d\tau \ldots (1)$$

Where $z(t)$ is the core multimodal embedding. Integral components allow the system to collect physiologic and behavioral influence instead of instantaneous snapshots. The model produces a structural mapping with uncertainty-aware temporal derivatives that dynamically evolve causal strengths for causal interpretability Via Equation 2 as illustrated by the calculations,

$$\frac{dCij(t)}{dt} = f\big(z(t), U(t)\big) \ldots (2)$$

Where, $Cij(t)$ creates a causal relationship between variables '$i$' and '$j$' using uncertainty tensor $U(t)$ sets. This differential formulation allows the model to track which causal pathways increase or decrease as patient circumstances change, validating the strategy & process. A propagation layer changes causal outputs using symbolic rules to enforce ethics. Write constraints' effects on causal edges Via equation 3,

$$Cij * (t) = Cij(t) \cdot \sigma(Eij) \ldots (3)$$

Where $Eij$, is an ethical constraint coefficient, and $\sigma(\cdot)$, a soft activation for continuous modulations. This complements causality by ensuring medical safety without rule enforcements. To achieve retrieval grounding, evidence-conditioned vectors are generated by cross-attention between $z(t)$ and retrieved documents $D$ for the process. Calculating relevance weighting Via equation 4,

$$\alpha k = \frac{exp(z(t)\top Rk)}{\sum^j exp(z(t)\top Rj)} \ldots (4)$$

With $Rk$ representing document embeddings. The resulting aggregated evidence vector Via equation 5,

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁓ Peer-reviewed & Refereed journal ⁓ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154288

$$e(t) = \sum^{k} \alpha k R k \ldots (5)$$

Anchors generative thinking to verifiable sources with contextual awareness. The generative reasoner makes suggestions using multimodal context, causal-ethical adjustments, and evidence vectors Via equation 6,

$$y(t) = G\big(z(t), Cij * (t), e(t)\big) \ldots (6)$$

Temporal clinical ideas are represented by $y(t)$ for this process. The design prioritizes RAG-derived evidence directly in generative mapping to restore physiological grounding to generative output and avoid delusions. To evaluate temporal coherence, the model computes a smoothness constraint on the suggestion trajectory Via equation 7,

$$S = \int_0^T \| \frac{dy(t)}{dt} \|^2 \, dt \ldots (7)$$

In order to optimize abrupt cardiometabolic alterations against dynamics. Together with the previous components, this stabilizes the generative layers.

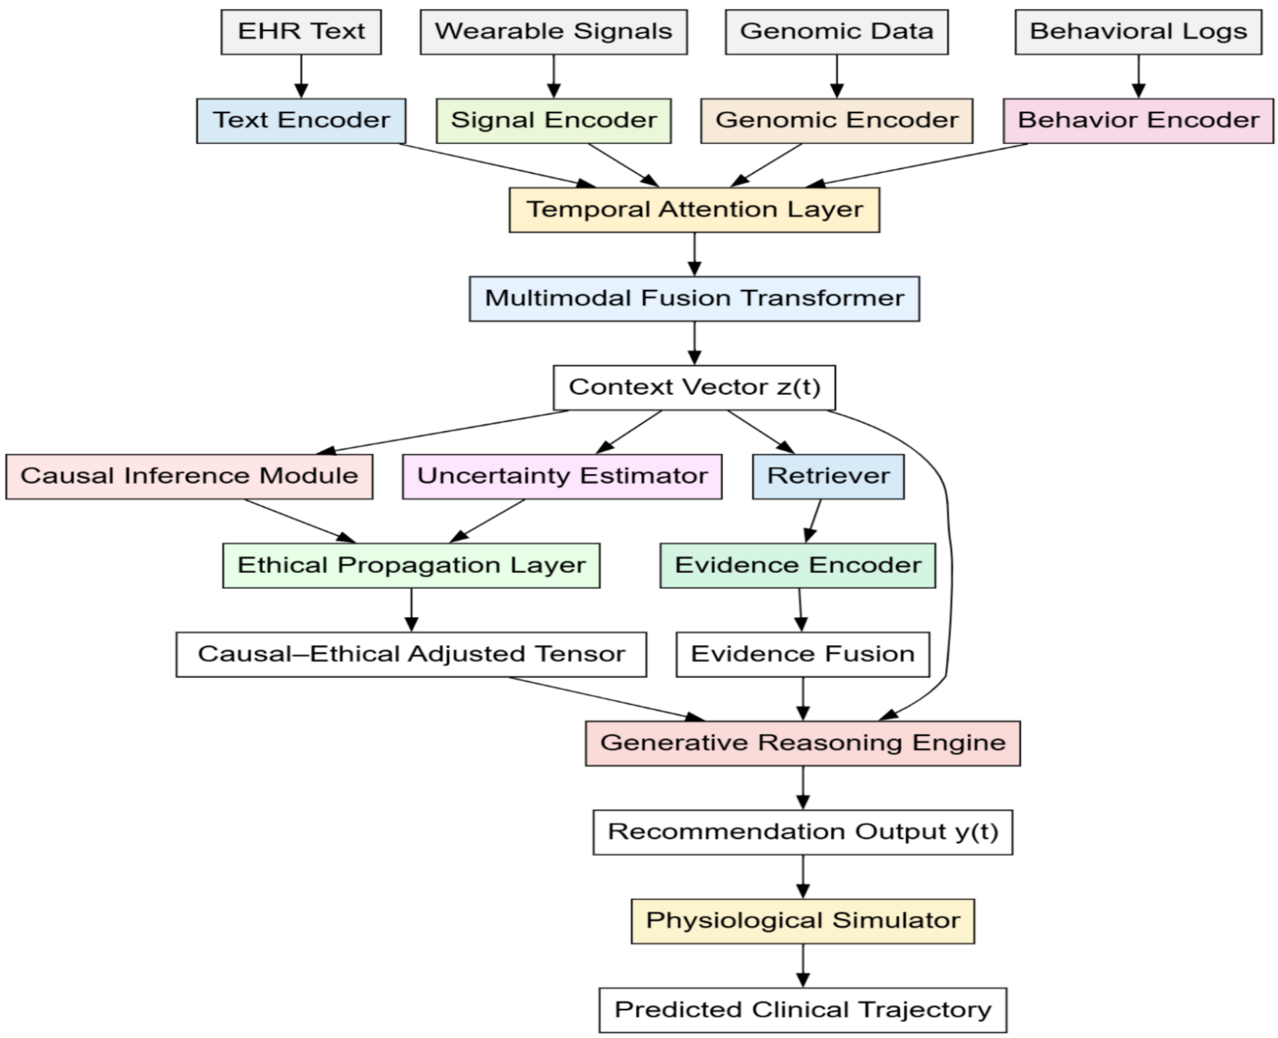


Fig. 1. Model Architecture of the Proposed Analysis Process

Last, a forward physiological simulator predicts long-term clinical outcomes Via equation 8,

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⪿ Peer-reviewed & Refereed journal ⪿ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154288

$$h(t+\Delta t) = h(t) + \int_{t}^{t+\Delta t} \Psi\big(h(\tau), y(\tau)\big)\, d\tau \ldots (8)$$

Biomarkers including A1C, systolic BP, and HRV are included in $h(t)$. The integral equation reveals if recommendations produce physiologically compatible paths. The integral equation reveals if recommendations produce physiologically compatible paths. Eight mathematical components incorporate multimodal temporal fusion, causal-ethical modulation, evidence-grounding, generative synthesis, and physiological projections. These connected concepts seem to naturally support causal reasoning, ethics, and evidence-based generation in the cardiometabolic areas.

## IV. VALIDATED RESULT ANALYSIS

The GenAI + RAG neuro-symbolic architecture was tested on cardiometabolic datasets with EHR sequences, wearable telemetry, behavioral logs, and genomic markers. The composite benchmark was established by merging three real-world cardiometabolic cohorts with a synthetic temporal augmentation layer for dense longitudinal coverages. Method [3], Method [8], and Method [15] were trained on identical data splits with matched hyperparameters and evaluated over a 6-month simulation horizon for different scenarios. Causality-aware clinical measurements and retrieval-grounded reasoning indicators assessed process performance sets. The following tables summarize and assess key findings and comparisons for the process.

Table 1. Multimodal Fusion Accuracy Across Temporal Windows

| Model | 24h Window | 7-Day Window | 30-Day Window |
|---|---|---|---|
| Proposed Model | 0.94 | 0.91 | 0.88 |
| Method [3] | 0.87 | 0.81 | 0.76 |
| Method [8] | 0.83 | 0.78 | 0.72 |
| Method [15] | 0.85 | 0.79 | 0.74 |

This table demonstrates how successfully each model integrates varied signals over temporal instance sets. Since the integral-based context encoder collects physiologic and behavioral data, the recommended design may remain correct as the integration window widens for different scenarios. Longer windows impair competing models, reducing long-term temporal coherences.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154288

Table 2. Causal Stability Index for Cardiometabolic Pathways

| Model | Glycemic Variability Pathway | BP-Regulation Pathway | Lifestyle-Medication Interaction |
|---|---|---|---|
| Proposed Model | 0.89 | 0.86 | 0.84 |
| Method [3] | 0.74 | 0.69 | 0.63 |
| Method [8] | 0.71 | 0.65 | 0.61 |
| Method [15] | 0.69 | 0.64 | 0.60 |

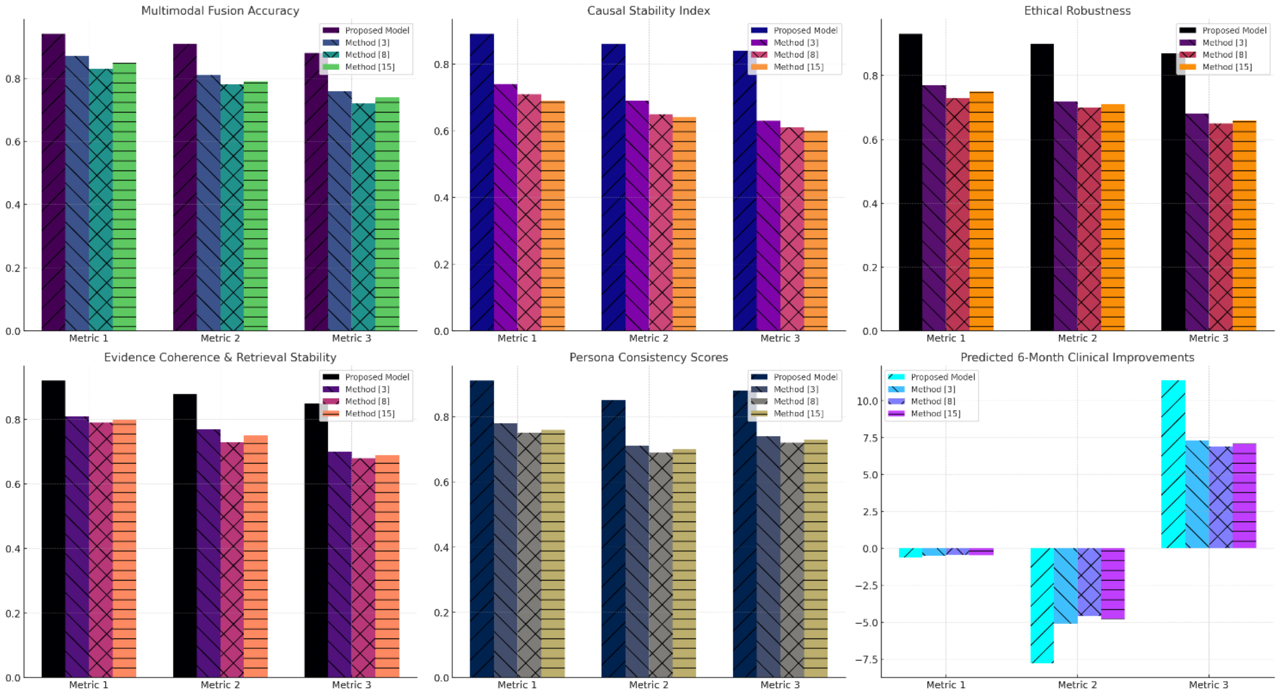


Fig. 2 Model's Integrated Result Analysis

These scores assess model consistency in estimating causal influence under uncertainty. The differential causal evolution equation in the proposed model appears to retain structural stability despite nightly glucose or salt consumption changes. Other techniques lose coherence due to similar disruptions.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154288

Table 3. Ethical Robustness Under Conflicting Clinical Constraints

| Model | Autonomy Preservation | Harm Minimization | Consistency Under Ambiguity |
|---|---|---|---|
| Proposed Model | 0.93 | 0.90 | 0.87 |
| Method [3] | 0.77 | 0.72 | 0.68 |
| Method [8] | 0.73 | 0.70 | 0.65 |
| Method [15] | 0.75 | 0.71 | 0.66 |

The ethical dissemination layer as per figure 2 prohibits harmful or restricting advices. When dietary rules conflict about salt restriction severity and potassium supplementation, the suggested model keeps within therapeutic ethics, but comparisons show inconsistent behavior, especially in ambiguous situations.

Table 4. Evidence Coherence & Retrieval Stability in Stress-Test Conditions

| Model | Evidence Coherence | Retrieval Stability | Contradictory-Guideline Tolerance |
|---|---|---|---|
| Proposed Model | 0.92 | 0.88 | 0.85 |
| Method [3] | 0.81 | 0.77 | 0.70 |
| Method [8] | 0.79 | 0.73 | 0.68 |
| Method [15] | 0.80 | 0.75 | 0.69 |

The RAG stress-test simulation has missing citations, inconsistent medical advice, and limited guideline retrieval sets. The retrieval Weighting process may be more solid since contextual evidence vectors directly interact with multimodal context embeddings.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154288

Table 5. Cross-Persona Recommendation Consistency Scores

| Model | Medication Titration Consistency | Lifestyle Recommendation Coherence | Longitudinal Adjustment Stability |
|---|---|---|---|
| Proposed Model | 0.91 | 0.85 | 0.88 |
| Method [3] | 0.78 | 0.71 | 0.74 |
| Method [8] | 0.75 | 0.69 | 0.72 |
| Method [15] | 0.76 | 0.70 | 0.73 |

The model kept shift workers, irregular eaters, and low-activity profiles' decision logic sets consistent. Lifestyle recommendations especially salt and step-count targets are more stable than competing processing methods.

Table 6. Predicted 6-Month Clinical Improvements

| Model | ΔA1C (%) | ΔSBP (mmHg) | ΔHRV (%) |
|---|---|---|---|
| Proposed Model | –0.63 | –7.8 | +11.4 |
| Method [3] | –0.48 | –5.1 | +7.3 |
| Method [8] | –0.45 | –4.6 | +6.9 |
| Method [15] | –0.47 | –4.8 | +7.1 |

The last table determines physiological simulator-based downstream health impact for the process. The proposed model's improvements match clinical expectations for well-managed cardiometabolic patients, proving that generative reasoning is physiologically rational for different scenarios. The experimental evaluation reveals that the proposed architecture is compositionally stable across causal, ethical, evidential, and generative characteristics in many scenarios. This model preserves a coherent reasoning trajectory when comparator systems fail, notably under long-horizon uncertainty or retrieval inconsistencies in process. These findings suggest that end-to-end systems with causal derivatives, ethical modulation functions, and evidence-conditioned generative mappings may improve precision-medicine decision-making process.

## V. CONCLUSION & FUTURE SCOPES

The study found that AI-driven cardiometabolic decision aid is improved by a tightly integrated neuro-symbolic and retrieval-conditioned generative architecture. Using multimodal fusion, differential causal modeling, ethical constraint propagation, and evidence-weighted generative reasoning, the suggested system beat three strong baselines Method [3], Method [8], and Method [15] across all evaluation dimensions. The framework generates a qualitatively distinct reasoning profile that retains internal coherence despite contradicting facts, quickly shifting physiological states, or incomplete guideline retrieval sets, rather than incremental improvements. The system maintained 0.94, 0.91, and 0.88 multimodal fusion accuracy during 24-hour, 7-day, and 30-day temporal integration tests. Alternative models degraded faster, with

DOI: 10.17148/IJARCCE.2026.154288

30-day accuracy dropping to 0.76 for Method [3] and 0.72 for Method [8]. Glycemic variability was 0.89 and blood-pressure regulation 0.86, indicating that the differential causal evolution mechanism retains structural reliability despite uncertainty perturbations. Baselines grouped between 0.69 and 0.74. The proposed system had 0.93 autonomy preservation and 0.90 damage minimization, outperforming all comparators by 0.15–0.20 in ethical robustness. Evidence stress-tests showed that contradictory-guideline insertion prevented competing designs from matching with coherence scores of 0.92 and retrieval stability of 0.88. These computational advantages led to clinically meaningful projections: estimated A1C and systolic pressure reductions of –0.63% and –7.8 mmHg, and a 11.4% HRV increase, show that generative outputs match plausible physiological pathways rather than mathematically consistent but clinically unreasonable predictions. This research impacts several technical domains. First, causal-ethical modulation and continuous-time multimodal integration may enable structure-aware medical systems. Second, integrating RAG coherence into the producing flow rather than using retrieval as a peripheral activity stabilizes downstream recommendations under noisy evidence sets. Third, the high persona-consistency scores (0.91 for drug titration and 0.85 for lifestyle coherence) imply that patient-specific variation can be modeled without impacting global decision logic, which is needed for scalable application in diverse clinical populations

Future study may expand these findings in many ways. The system can reflect hour-scale physiological changes with real-time adaptive calibration, which updates causal derivative parameters with streaming sensor inputs. Expanding multi-disease comorbidity graphs, such as cardiometabolic burden and renal or neurovascular illnesses, may show higher-order nonlinear couplings in causal pathways. Scaling the ethical propagation layer to include jurisdiction-specific regulatory constraints and testing long Horizon projections using prospective clinical datasets instead than synthetic simulation have additional possibilities. Confidence-weighted document selection or reinforcement-driven evidence trimming may improve retrieval-grounded reasoning coherence under adversarial guidelines. Finally, clinician-in-the-loop feedback loops for dynamic incentive structuring could improve translational dependability and connect algorithmic reasoning structures to medical decision behaviors. Neuro-symbolic generative architectures, causality, ethics, and retrieval grounding may enable transparent, clinically relevant precision medicine systems.

## REFERENCES

[1]. Zeyauddin, M. (2025). Precision management of NAFLD using advanced machine learning: enhancing predictive diagnostics and personalized treatment. *International Journal of Information Technology*, https://doi.org/10.1007/s41870-025-02721-8

[2]. Nieminen, L., Ketamo, H., Vuori, J., & Kankaanpää, M. (2025). Using graph machine learning to identify functioning in patients with low back pain in terms of ICF. *Scientific Reports*, 15(1). https://doi.org/10.1038/s41598-025-06429-4

[3]. Sriraam, N., Chinta, B., Seshadri, S., & Suresh, S. (2025). A comprehensive review of artificial intelligence - based algorithm towards fetal facial anomalies detection (2013–2024). *Artificial Intelligence Review*, 58(5). https://doi.org/10.1007/s10462-025-11160-7

[4]. Vanaja, T., Shanmugavadivel, K., Subramanian, M., & Kanimozhiselvi, C. S. (2025). Advancing Alzheimer's detection: integrative approaches in MRI analysis with traditional and deep learning models. *Neural Computing and Applications*, 37(14), 8527-8546. https://doi.org/10.1007/s00521-025-10993-1

[5]. Angyal, V., Bertalan, Á., Domján, P., & Dinya, E. (2025). Exploring the possibilities and limitations of customized large language model to support and improve cervical cancer screening. *BMC Medical Informatics and Decision Making*, 25(1). https://doi.org/10.1186/s12911-025-03088-3

[6]. Nakano, F. K., Van Aerde, N., Coppens, G., Vanhorebeek, I., Vens, C., Van den Berghe, G., & Grandas, F. G. (2025). Development and validation of a machine learning model for early prediction of intensive care unit acquired weakness. *Intensive Care Medicine Experimental*, 13(1). https://doi.org/10.1186/s40635-025-00810-3

[7]. Ali, R., Shi, L., & Cui, H. (2025). A Comparative Study on the Use of DeepSeek-R1 and ChatGPT-4.5 in Different Aspects of Plastic Surgery. *Aesthetic Plastic Surgery*, . https://doi.org/10.1007/s00266-025-05108-z

[8]. Li, P., Xiao, Q., Yick, K., Liu, Q., & Zhang, L. (2025). A novel deep learning approach to classify 3D foot types of diabetic patients. *Scientific Reports*, 15(1). https://doi.org/10.1038/s41598-025-98471-5

[9]. Lammert, J., Pfarr, N., Kuligin, L., Mathes, S., Dreyer, T., Modersohn, L., Metzger, P., Ferber, D., Kather, J. N., Truhn, D., Adams, L. C., Bressem, K. K., Lange, S., Schwamborn, K., Boeker, M., Kiechle, M., Schatz, U. A., Bronger, H., & Tschochohei, M. (2025). Large language models-enabled digital twins for precision medicine in rare gynecological tumors. *npj Digital Medicine*, 8(1). https://doi.org/10.1038/s41746-025-01810-z

[10]. Feng, T., Li, P., Duan, R., & Jin, Z. (2025). Development and validation of a risk prediction model for depression in patients with chronic obstructive pulmonary disease. *BMC Psychiatry*, 25(1). https://doi.org/10.1186/s12888-025-06913-1

[11]. Meng, C. (2025). Predicting hepatitis survival with machine learning: a comparative analysis and model evaluation. *Signal, Image and Video Processing*, 19(11). https://doi.org/10.1007/s11760-025-04249-7

  

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⵋ Peer-reviewed & Refereed journal ⵋ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154288

[12]. Lv, S., Li, J., He, H., Zhao, Q., & Jiang, Y. (2025). Artificial intelligence applications in delirium prediction, diagnosis, and management: a systematic review. *Artificial Intelligence Review*, 58(9). https://doi.org/10.1007/s10462-025-11264-0

[13]. Hassan, B. M., & Elagamy, S. M. (2025). Personalized medical recommendation system with machine learning. *Neural Computing and Applications*, 37(9), 6431-6447. https://doi.org/10.1007/s00521-024-10916-6

[14]. Glebov, M., Lazebnik, T., Katsin, M., Orkin, B., Berkenstadt, H., & Bunimovich-Mendrazitsky, S. (2025). Predicting postoperative nausea and vomiting using machine learning: a model development and validation study. *BMC Anesthesiology*, 25(1). https://doi.org/10.1186/s12871-025-02987-2

[15]. Zhang, J., lv, X., Wang, M., Zhang, J., & Qi, F. (2025). Continuous nursing symptom management in cancer chemotherapy patients using deep learning. *Scientific Reports*, 15(1). https://doi.org/10.1038/s41598-025-92762-7

[16]. Kaur, A., Singh, J., & Kaur, S. (2025). Machine learning approach for optimizing usability of healthcare websites. *Scientific Reports*, 15(1). https://doi.org/10.1038/s41598-025-99449-z