

MedVersa: A Review of NLP and GAN Frameworks for Patient–Provider Communication

Madhuri Zavar¹, Sanmati Kumar Jain², Pradnya Vikhar³

¹Research Scholar & Assistant Professor, Department of AIML GHRCEM Pune, India,

²Associate Professor, Department of CSE Vikrant Group of Institutions Indore, India

³Associate Professor, Department of CSE RSCOE Pune, India

Abstract- Telemedicine is now an essential care delivery mode, but patients (especially in rural and resource-restricted environments) continue to experience challenges associated with access, symptom interpretation, and communication with providers. The development of Artificial Intelligence (AI), Natural Language Processing (NLP) and Generative Adversarial Networks (GANs) presents the opportunities to address these shortcomings. This review summarizes the latest papers on AI-enabled telemedicine with the focus on: (i) telehealth applications augmented by AI analytics, (ii) NLP models of clinical records and chatbots, and (iii) GAN-based models of synthetic medical data. Recent literature was located by searching large digital libraries on the keyword's telemedicine, NLP, GANs, and patient-provider communication and filtered on remote-care applications. The reviewed articles present a positive change in predicting risks, remote monitoring, imaging, and dialogue support, whereas patient-centered communication outcomes (e.g., comprehension, satisfaction, trust) are mentioned less frequently, and most studies use accuracy-oriented indicators, which are not statistically validated at the time. Judging by these gaps, we propose the concept of MedVersa as a conceptual framework, which includes GAN-based data generation with large language model (LLM) dialogue features to increase the dependability, fairness, and contextual sensitivity of teleconsultations and guide the implementation and evaluations in the future. This article contains a literature review and conceptual architecture, and does not indicate implemented prototype or new empirical findings.

Keywords: Telemedicine, Natural Language Processing, Generative Adversarial Networks, Patient–Provider Communication, Healthcare AI.

I. INTRODUCTION

Healthcare delivery has always been transformed by technology and innovation, and lately digital health technologies in the form of telemedicine, mobile applications, and artificial intelligence (AI) have changed access and personalization of care [1]. However, even alongside a rapid uptake, telemedicine is known to have enduring weaknesses: patients in rural or resource-constrained regions usually have issues with access, and symptom misunderstanding in remote consultations is not rare, and lack of communication between patients and providers is what is likely to destroy trust and clinical precision [2]. These problems highlight the idea that even though telehealth has been facilitated by infrastructure and policy, the quality of patient-provider conversation is a bottleneck in the middle.

This review discusses how methods based on AI-enabled methods might be used to overcome these

shortcomings by enhancing remote-care systems and facilitating more detailed and contextual communications. Recent development focuses on high capacity, low-latency connectivity and AI-based analytics as an important enabler of large-scale telemedicine and continuous monitoring [3], and newer mobile health communication architectures that can transmit clinical signals dependably in resource-constrained environments [4].

However, a significant part of the literature currently is still disjointed, in the sense of it viewing infrastructure, sensing/monitoring, and clinical AI as independent layers, without considering the integration of generative data augmentation with conversational intelligence. To fill this gap, the MedVersa is introduced as a theoretical framework that involves the synthesis of synthetic medical data using GANs and the adaptation of dialogue-based LLM as the means of enhancing the reliability and

contextual sensitivity of patient-provider communication in telemedicine [8], [12].

This research has four contributions. It starts by considering AI-based telemedicine studies in the domain of clinical language understanding, dialogue support, and data generation [1], to which it includes work suggesting how deep learning can learn clinically informative structure in unstructured medical text to perform downstream prediction tasks [5]. Second, it examines measurement measures and methodological shortcomings of the current literature, such as case of low statistical validation and consideration of equity factors [1], [2]. Third, it constructs the main technical, clinical and ethical issues in terms of patient-provider communication in telemedicine [1], [2]. Lastly, it proposes conceptual framework MedVersa as an idea combining GANs and LLMs as a guideline for future studies on strong contextual teleconsultation instead of a system.

II. BACKGROUND

The next phase of digital healthcare transformation is being driven by Artificial intelligence (AI) and six-generation (6G) networks. AI complements diagnostics, predictive models and personalization of treatments by analyzing large and complex data sets, and 6G supports the delivery of ultra-fast, low-latency, and reliable communication needed in real-time applications, including telesurgery, remote monitoring, and large-scale telemedicine [3]. Collectively, these technologies are in favor of Smart Healthcare 4.0, in which Internet of Medical Things (IoMT) devices constantly gather and provide patient information to be early diagnosed and receive timely interventions, even in underserved areas.

Generalized frequency division multiplexing (GFDM), Multi-input Multi-output (MIMO), electroencephalography (EEG) communication systems have ultra-low power consumption, high throughput, and very low error rates and are thus highly applicable to mHealth applications, IoMT applications and telemedicine applications [4]. The fact that they can transmit the EEG signals reliably even in erroneous channel Estimation demonstrates

the possibility of incorporating neuro-technology into the future healthcare. Figure 1 shows integration of AI and 6G in healthcare services. Next-generation systems enhance mobile health as demonstrated by the emerging communication architecture.




Fig. 1: Integration of AI and 6G in Healthcare Services.

Collectively, AI-based analytics, 6G connectivity, and advanced donation models compose the enabling layer of such structures as MedVersa, where synthetic data generation and dialogue models based on NLP rely on strong and real-time infrastructures. These developments form the enabling infrastructure for frame-works such as MedVersa, which depend on real-time connectivity and secure data exchange [3], [4].

III. REVIEW METHODOLOGY

The recent literature was found by conducting searches in major digital libraries where such keywords as telemedicine, patient-provider communication, NLP, clinical dialogue, GAN, and synthetic medical data were used. It was also required that studies had to be related to remote-care or telemedicine contexts and use AI techniques related to analytics, clinical NLP/dialogue, or synthetic data generation. Research that was specifically done in the context of non-telemedicine or with no relevance to healthcare was filtered out. We also managed to extract (i) the type of application (e.g., prediction, monitoring, imaging, dialogue) of the study, (ii) data modality, (iii) the metrics used to evaluate it (e.g., statistical validation confidence intervals, repeated runs, significance

testing), and (iv) the presence of statistical validation like confidence intervals, significance testing and patient-centered communication outcomes.

IV. REVIEW OF LITERATURE

The increasing overlap between telemedicine and artificial intelligence (AI) and sophisticated computational models has created a variety of methods to enhance healthcare delivery. New research points to applications in predictive modeling and bio signal monitoring to conversational NLP and medical image synthesis which can serve as a useful guide to the opportunities and limitations of AI-based telemedicine.

Telemedicine and AI Integration Overview

AI is growing telemedicine in the clinical text analysis, community health surveillance, medical imaging, and remote monitoring. Supposedly, deep learning and NLP are capable of forecasting the recurrence of atrial fibrillation using unstructured clinical reports [5]. Telemedicine interventions can enhance the results of underserved populations, including digital cardiovascular risk management at the population level, but the absence of access and low health literacy sabotage them [6]. Prenatal and maternal care is another area that has seen improvements in patients in geographically far regions in obstetrics and gynecology through telemedicine [7]. Overall, AI can be used to take telemedicine to greater scales because disconnected clinical and social data are converted into practical insights, but the non-uniform adoption and inappropriate integration into the healthcare process is an unsolved problem.

Healthcare and NLP Role

NLP has become one of the main pillars of AI-based healthcare capable of transforming unstructured medical data into functional knowledge [12]. Contextual information, such as clinical narratives, discharge reports, imaging reports, and patient-provider communications, are lost in structured datasets [11]. By using modern NLP models, healthcare systems can not only improve predictive accuracy, minimize errors in coding and support

more meaningful communication in both clinical and telemedicine settings [11], [12].

NLP has advanced healthcare through imaging annotation [8], [9], [10], EHR structuring [11], and conversational AI [12], all of which improve patient-provider communication. However, persistent gaps remain in data quality, interpretability, and validation for real-world deployment. These limitations in Table I highlight the need for integrated frameworks like MedVersa that combine GAN-based data augmentation with dialogue-specific NLP models. However, few of these systems are tested in real-world telemedicine workflows, and only a minority report patient-centered communication outcomes, such as understanding, satisfaction, or trust [8], [9], [10], [11], [12].

Healthcare based Generative Adversarial Networks

GANs find more applications in healthcare to create artificial data to overcome the medical data underrepresentation and sensitivity. Their adversarial training is able to create realistic images to augment, reconstruct, and translate, but more recent studies prove that visual realism does not necessarily translate to clinical reliability since medically meaningful statistics may not be maintained [13]. Consequently, GAN research is commonly measured based on image-level measures (SSIM, PSNR, FID, or subjective visual rating) [8], [9], [10], [13] and comparisons between model variants and the existence of clinically-relevant distributions are less frequently reported.

Table I: Overview of NLP in Healthcare

Citation	Advances	Impact on Communication	Research Gap
S. Islam et al [8]	GANs + text in imaging	Links annotations with diagnostics	Limited datasets; weak NLP-GAN integration
Bruno et al [9]	GAN-augmented endoscopy	Free-text guides synthetic data	Needs broader validation, domain alignment
Shobayo et al [10]	Multimodal DL + NLP	Improves accuracy via	Experimental; untested in telemedicine

		text–image synergy	
S. Daphne et al [11]	EHR graphs with BERTs	Makes records interpretable, aids prediction	EHR noise, bias; domain adaptation needed
Mohamed et al [12]	Transformers in healthcare + conversational NLP	Automates reporting, supports triage, education, dialogue	Interpretability/fairness issues; data scarcity; trust concerns




Fig. 3. Role of GANs in Healthcare.

This reduces trust in the GAN-generated data, which can be applied in downstream diagnostic or dialogue systems, which encourages domain-specific assessment that entails statistical tests to real clinical data and indicators of effects on downstream tasks.

GANs have also been applied to radiology, pathology, and endoscopy to enhance the quality of images and increase training data even when faced with low-resource scenarios [13]. Figure 3 shows that GANs can facilitate multimodal learning by matching synthesized results with clinical labels and enhancing downstream NLP pipelines. As the use of LLMs gets more embedded in the sphere of healthcare communication, synthetic datasets generated with the help of GANs can also aid in the reduction of bias and in improving adaptation in low-resource languages [14]. The latest surveys also indicate that data generation by GANs combined with dialogue models by means of NLPs can enhance the telemedicine level of robustness and communication stability [15].

In summary, GANs mitigate data scarcity and enhance AI readiness in healthcare. However, clinical trust requires evaluation methods that prioritize medical significance. Their convergence with NLP and LLMs presents an opportunity to build hybrid systems, such as MedVersa, that combine synthetic data with dialogue intelligence to improve patient-provider communication.

Key Challenges in Enhancement of Patient-Provider Communication

Despite advances in AI-driven telemedicine, patient-provider communication continues to face critical challenges.

Another problem is the possibility of wrong diagnosis in the case of remote consultations, where it is possible to forget about sufficient visual and contextual information, especially in cardiovascular, maternal, and post-acute treatment [6], [7].

Technically, unstructured clinical data is also a significant challenge. Electronic health records (EHRs) are noisy, fragmented, and prone to variability in style, abbreviations, and errors, complicating their direct use in predictive models [11]. While transformer-based NLP models and embeddings such as ClinicalBERT and BioBERT improve interpretability, they still operate as “black boxes,” limiting physician trust and clinical adoption [12].

GANs are also applicable in the field of imaging, where they assist in alleviating data paucity, but the reliability and clinical validity of generated data have not been properly addressed yet [13]. Models can generate mediocre realistic percepts, but unrealistic from the medical perspective, or they may do harm when not used critically. The possibility of integrating GANs with NLP and LLMs exists, yet it has not been incorporated yet [14], [15].

The obstacle is developing patient rapport. Conversation based AI can triage symptoms and deliver health education, but is likely to yield false or misinformed results because probabilistic generation and the established weaknesses of large language models in safety-critical tasks [14], [15].

This brings up the issue of the authenticity of AI-supported telemedicine and inspires interpretable, patient-centered models with human-in-the-loop supervision [12], [14].

In brief, the barriers to communication are the result of clinical, technical, ethical limitations: incomplete data, information disjunction, data absence, bias, and patient distrust. The resolution of these problems promotes the MedVersa who integrates the GAN-generated synthetic data with the dialogue-inspired NLP to produce more reliable, more equitable and more context-aware telemedicine.

Evaluation Metrics and Statistical Validation in Existing Studies

The reviewed studies provide diverse evaluation metrics-accuracy/precision/recall/F1 and AUC to predict and classify as they summarize discriminative performance on clinical risk and triage tasks, BLEU-like to generate dialogue as they provide a few words a quick measure of lexical overlap of generated and reference responses, and SSIM/FID to generate synthetic imaging since they approximate visual similarity and distributional realism of the generated samples [8], [9], [12], [13].

$$\text{Precision} = \frac{TP}{TP+FP} \quad \text{Recall} = \frac{TP}{TP+FN}$$

$$F1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

Nevertheless, these convenient metrics are not entirely indicative of the requirements of telemedicine: AUC/F1 can determine errors of clinical significance without calibration and subgroup analysis, BLEU-like metrics might not predict safety, accuracy, compassion, and clear communication in patient-physician communication, and SSIM/FID are not guaranteed to signal medically significant fidelity and diagnostic validity of artificial data [13]. Further, a significant number of the studies do not contain confidence intervals, hypothesis testing, or repeated-run statistics, and hardly ever provide patient-centered outcomes, including comprehension, satisfaction, trust, and workload of clinicians, which makes it challenging to understand

whether AI has a meaningful impact on enhancing telemedicine encounters [12], [13].

The future implementations of MedVersa thus expect statistically significant, communication-based assessment where repeated runs are used alongside confidence intervals/significance test over patient-reported outcomes and task measures.

V. CONCEPTUAL FRAMEWORK: MEDVERSA

The conceptual framework of MedVersa is introduced as the theory of work but no actual implementation or empirical research and future research directions in telemedicine platforms that are AI-enabled are derived based on the architecture.




Fig. 4. Conceptual MedVersa Framework.

The suggested MedVersa framework will combine the Generative Adversarial Networks (GANs) and the Large Language Models (LLMs) to improve telemedicine communication. MedVersa solves two major issues, namely, scanty medical data and untrustworthy patient-provider communication. The system is intended to produce synthetic medical data and optimize dialogue-specific LLM more accurately, context-sensitive, and reliable responses during the teleconsultation as depicted in Figure 4.

Architecture of the Medical GAN Component

The Medical GAN Component in MedVersa is tasked with generating synthetic medical records that closely resemble real-world patient data. This is an essential step since the size of healthcare datasets is normally constrained by privacy issues, high annotation, and regulatory limitations.

Data Augmentation and Synthesis: MedVersa takes the form of a generator-based discriminator cycle in which synthetic samples are generated and their authenticity is tested by the discriminator. An input random noise vector leads to a synthetic patient record produced by the generator $Gen(z)$, and it is expected that the generator does not minimize any error, rather the discriminator $Dis(x)$ is trained to differentiate between real data x and synthetic data. The adversarial training process is a recurrent process: the generator is also being enhanced to generate more realistic data, and the discriminator is being enhanced to identify fabricated entries better.

$$\min_G \max_D E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} [\log (1 - D(G(z)))]$$

Unique Contribution: In contrast to typical GAN pipelines, which produce purely image-based data, MedVersa uses the architecture on multi-modal health data, that is, structured EHR fields, narrative medical notes, and conversation transcripts. This makes sure that the numerical and linguistic aspects of the healthcare data are captured. The output is a wide range of synthetic medical samples with the ability to increase the training corpus of downstream NLP tasks.

Impact: The aim of the GAN module is that it will help lessen overfitting of the LLM by enlarging and diversifying the training set, enhance the coverage of rare clinical instances, and more accurately produce data to be used in the downstream dialogue modeling, as seen in previous studies of GAN-based augmentation in medical fields [8], [9], [13]. This design purpose is to improve telemedicine dialogue models, but it will have to be confirmed in the future implementation.

Enhancing Communication through LLM

The generated enriched dataset stems from the Medical GAN and is used as the input to the LLM NLP Component. MedVersa is a transformer-based model (e.g. Clinical BERT, GPT styles), fine-tuned on mixed real and synthetic data.

Contextual Consideration: The LLM takes the tokenized consideration of inputs from patient queries. LLM subsumes it into a high-dimensional semantic space. It takes the contextual dependencies of the conversation through self-attention mechanisms. It enables the interpretation of nuances to the speech of the patient like a negation (e.g. Regarding chest pain (e.g. "I do not feel chest pain anymore") or development of the symptoms (e.g. the cough worsened at night).

Response Generation: The LLM translates the context into medical dialogue responses after being trained in the natural language. The responses generated are grammatically correct and clinically meaningful to be relevant to the telemedicine environment.

Advantage Over Existing Systems: Most telemedicine chatbots are either retrieval based (selecting pre-existing answers), or domain generic LLMs with superficial adaptation to healthcare. The MedVersa, its device, makes use of artificial GAN-enhanced training data to allow the LLM to handle expert clinical language more effectively.

Contextual Understanding and Response Generation
MedVersa addresses the problem by introducing contextual integrity in response generation mechanism. The system is able to analyze intent on

a finer granularity using tokenized and vectorized patient queries. Attention mechanisms are implemented in the form of transformers for identifying the inter token relationships.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

It makes sure conversations are understood as a whole, and not in individual statements. This contextualization makes the model to be able to rebuild the embeddings into fluent and medically applicable sentences for representation of the clinical accuracy and the patient understanding. As an example, a patient may give a question, asking, is my headache connected to my treatment? MedVersa provides the response by acknowledging the symptom and, taking into account the medical context. It provides a patient friendly explanation. In this way, breaking the dialogue loop, MedVersa generates context-sensitive results, which can increase patient trust, continuity of health care and the overall telemedicine experience.

Workflow of the MedVersa Framework

MedVersa has a systematic and end-to-end workflow in which every aspect is incorporated within a telemedicine platform.

Data Collection and Preprocessing: The patient's, clinical notes and the dialogue transcripts are collected. Preprocessing is the stage, which involves cleaning and tokenization of data, data vectorization and normalization to standardize input.

Medical GAN Training: GANs produce synthetic data in cycles. It will make the data more diverse and richer in infrequent cases of clinical observations.

LLM Fine-Tuning: The LLM based on transformers are trained on the biased dataset (real + synthetic) to reduce their bias. The generalizability is provided by validation and testing.

System Integration and Deployment: The GAN and the LLM modules are combined to create one telemedicine interface which is deployed on a load balanced cloud infrastructure.

Real-Time Interaction: The queries on the patients are processed by the LLDM. The augmented training set of GAN involves real time generation of the reactions.

Monitoring and Maintenance: It is powerful and meets the needs of the healthcare because of the constant control of the performance, the security inspection, and the upgrades of the model.

Steps of algorithm for MedVersa Framework as depicted in Table II.

Implementation and Evaluation Considerations
The MedVersa as it is now would describe a blueprint, as opposed to a prototype implemented. An actual implementation would necessitate (1) a secure integration with the current EHR and telehealth systems, (2) strict deidentification and governance of both real and synthetic data, and (3) human-in-the-loop systems in order to control clinicians. Assessment would involve not only technical measures (e.g., task accuracy, F1-score, latency) but also telemedicine-specific measures, including patient comprehension, patient satisfaction, length of consultation

Table II: Pseudocode for MedVersa Framework

Steps of MedVersa Algorithm
Initialize MedVersa Framework
Load Medical Data for Training
Preprocess Data: Tokenization and Vectorization
Train Medical GAN:
Initialize Generator and Discriminator Networks
Repeat until convergence:
Generate synthetic data using Generator
Evaluate synthetic data using Discriminator
Update Generator and Discriminator weights
Generate augmented medical data using trained Medical GAN
Integrate synthetic data with real data
Fine-tune LLM on combined dataset:
Initialize LLM model
Train on combined dataset using disentangled attention
Fine-tune for medical dialogue generation
Deploy LLM in Telemedicine System:
Receive patient queries
Process queries using LLM
Generate and return contextually relevant responses
Monitor system performance and update model as needed

and workload in clinicians. Before such a large-scale adoption, randomized or stepped-wedge trials of MedVersa-enhanced teleconsultations versus usual care would be needed. These considerations are presented as guidelines on what will be implemented in future work as opposed to reporting them as successful experiments.

VI. CONCLUSION

This is a review of recent research on AI-enabled telemedicine and clinical dialogue NLP, as well as GAN-powered synthetic medical data, particularly patient-provider communication. Whereas AI promises to be effective in risk prediction, remote monitoring, and documentation, communication quality, equity, and trust are less commonly studied, and most studies are based on accuracy-focused measures with limited statistical validation. We aim to fill these gaps by introducing a proposed conceptual framework named MedVersa, which involves the synthesis of synthetic data and dialogue modeling by GAN and LLM, respectively, to facilitate more contexts and inclusive teleconsultations. The presented paper is a narrative review and conceptual framework and does not feature an implemented prototype or a new empirical discovery.

Synthesis is based on what the literature in the existing studies tells us, most of them lack communication-focused results and statistically sound validation. Consequently, MedVersa architecture and evaluation approach can be viewed as a guideline which must be prototyped in the future, assessed under controlled conditions, and tested in clinical conditions to determine its real impact. The next step will be development of the prototype of MedVersa into real telemedicine systems and pilot studies to establish the impacts of the technology on the communication, safety, and availability of MedVersa particularly in small resource and rural settings.

REFERENCES

1. E. Cano-Marin, S. Sánchez-Alonso and M. Mora-Cantalops, "Unleashing Competitive Intelligence: News Mining Analysis on

- Technology Trends and Digital Health Driving Healthcare Innovation," in *IEEE Transactions on Engineering Management*, vol. 71, pp. 12311-12325, 2024.
2. W. M. Glinkowski, T. Cedro, A. Wołk, R. Doniec, K. Wołk, and S. Wilk, "Telemedicine, eHealth, and Digital Transformation in Poland (2014–2024): Trends, Specializations, and Systemic Implications," *Applied Sciences*, vol. 15, no. 16, Art. no. 8793, 2025.
3. Ahmad Khan et al., "Big Data Analytics Model Using Artificial Intelligence (AI) and 6G Technologies for Healthcare," in *IEEE Access*, vol. 12, pp. 97924-97937, 2024.
4. C.-F. Lin and K.-Y. Chen, "Generalized Frequency Division Multiplexing–Based Direct Mapping–Multiple-Input Multiple-Output Mobile Electroencephalography Communication Technique," *Applied Sciences*, vol. 15, no. 17, Art. no. 9451, 2025.
5. García-Olea, A. G. Domingo-Aldama, M. Merino, K. Gojenola, J. Goikoetxea, A. Atutxa, and J. M. Ormaetxe, "The Application of Deep Learning Tools on Medical Reports to Optimize the Input of an Atrial-Fibrillation-Recurrence Predictive Model," *Journal of Clinical Medicine*, vol. 14, no. 7, Art. no. 2297, 2025.
6. Gherman, C. M. Levai, O. A. Hațegan, C. M. Popoiu, E. R. Stoicescu, and A. L. Maghiari, "Telemedicine-Supported Intervention Versus Standard Care for Managing Cardiovascular Risk Factors in a Socially Deprived Urban Population: A Prospective Study," *Healthcare*, vol. 13, no. 17, Art. no. 2202, 2025.
7. E. Medani, A. M. Hakami, U. H. Chourasia, B. Rahamtalla, N. M. Adawi, M. Fadailu, A. Salih, A. Abdelmola, K. N. Hashim, A. M. Dawelbait, et al., "Telemedicine in Obstetrics and Gynecology: A Scoping Review of Enhancing Access and Outcomes in Modern Healthcare," *Healthcare*, vol. 13, no. 16, Art. no. 2036, 2025.
8. S. Islam et al., "Generative Adversarial Networks (GANs) in Medical Imaging: Advancements, Applications, and Challenges," in *IEEE Access*, vol. 12, pp. 35728-35753, 2024.
9. Magalhães, A. Neto and A. Cunha, "Generative Adversarial Networks for Augmenting Endoscopy Image Datasets of Stomach

- Precancerous Lesions: A Review," in IEEE Access, vol. 11, pp. 136292-136307, 2023.
10. O. Shobayo and R. Saatchi, "Developments in Deep Learning Artificial Neural Network Techniques for Medical Image Analysis and Interpretation," *Diagnostics*, vol. 15, no. 9, Art. no. 1072, 2025.
 11. S. Daphne, V. M. A. Rajam, P. Hemanth, and S. Dinesh, "An Ensemble Patient Graph Framework for Predictive Modelling from Electronic Health Records and Medical Notes," *Diagnostics*, vol. 15, no. 6, Art. no. 756, 2025.
 12. Mohamed, R. AlAleeli, and K. Shaalan, "Advancing Predictive Healthcare: A Systematic Review of Transformer Models in Electronic Health Records," *Computers*, vol. 14, no. 4, Art. no. 148, 2025.
 13. V. A. Kelkar et al., "Assessing the Ability of Generative Adversarial Networks to Learn Canonical Medical Image Statistics," in *IEEE Transactions on Medical Imaging*, vol. 42, no. 6, pp. 1799-1808, June 2023.
 14. P. Peykani, F. Ramezanlou, C. Tanasescu, and S. Ghanidel, "Large Language Models: A Structured Taxonomy and Review of Challenges, Limitations, Solutions, and Future Directions," *Applied Sciences*, vol. 15, no. 14, Art. no. 8103, 2025.
 15. M. A. K. Raiaan et al., "A Review on Large Language Models: Architectures, Applications, Taxonomies, Open Issues and Challenges," in *IEEE Access*, vol. 12, pp. 26839-26874, 2024.