

THE IMPACT OF ARTIFICIAL INTELLIGENCE ON HEALTHCARE

IULIA-CRISTINA STĂNICĂ¹, RADU-GIORGIAN CONSTANTIN², IONUT-SEBASTIAN LUTAN²,
ANDREI-ALEXANDRU CHITU², COSTIN-ANTON BOIANGIU²

Keywords: Artificial Intelligence; Healthcare; Large language models; Explainable Artificial Intelligence; Generative Artificial Intelligence.

The advanced development of artificial intelligence (AI) in healthcare has exhibited such promise hence it revolutionizes diagnostic precision, facilitates personalized treatments, and streamlines procedures. AI has progressed from primitive rule-based systems such as MYCIN to today's last generation of AI models such as GPT-4 and Med-Gemini. AI has been quite consistent in enhancing its capabilities and bringing continuing innovation in evidence-based medicine from medical-image predictive analytics to patient monitoring. Explainable AI (XAI) helps address the main problems of trust in AI because it allows transparent and understandable machine-learning predictions on which the physician can rely. Our paper investigates the potential of AI to change the delivery of care and define the cost, access, and outcome of care for different populations within the integration of human expertise and machine efficiency. Future research would need to combine innovation with ethical considerations and achieve full adoption, regarding its potential and quality operation in the healthcare sector.

1. INTRODUCTION

There have been significant advances in artificial intelligence (AI) and rapidly growing healthcare systems, which are producing remarkable efficiencies and innovations that are paving the way for a better and healthier future [1]. These technological transitions also make an impact on the processes of improving health services.

Generative Models and large language models (LLMs) are the main forces driving change and innovation, and they are transforming themselves in many applications such as medical image processing, analysis, and automatic diagnosis, while personalized therapy and drug discovery can be done at an exponentially higher speed [2]. LLMs offer the best programs for clinical decision support, delivering evidence-based insights, automating manual processes, and enhancing precision and effectiveness in healthcare services.

Despite all this, the full transition of AI in the health sector would necessitate a balanced approach that would allow for overcoming operational bottlenecks, fine-tuning the delivery of care, and bringing all public trust into the capability of AI for ethical governance. Since the evolution of AI from the early stages to the current methods of application, it has emerged as a swift influence upon the medical sector, setting down challenges offered and future directional focus for AI integration into healthcare systems.

Health systems around the world are being stretched to their limit, battling between increasing demand and spiraling costs while medicine itself becomes more complex and necessitates more personalized health services. To this end, most of the developed economies are interested in innovation; artificial intelligence has made a leap into the power that uses groundbreaking ways to schedule immense data volumes and patterns. But as soon as this work turns into practice, traditional tools would appear less competitive, and deficiencies in diagnostic precision, planned therapies, and operational reform would be detailed all along. This lag underlines the absolute necessity of integrating innovative strategies into the medical industry while touching upon new issues arising in terms of ethical considerations and public concerns over AI. Furthermore, peoples skepticism should be taken into account when using such novel approaches that

would implicate considerable risks.

Our current paper presents an analysis of the potential impact AI can have in healthcare, starting with the presentation in chapter 2 of its historical evolution. Chapter 3 delves into the “hows” and “whys” of the technology presenting the concept of explainable artificial intelligence, while section 4 underlines ethical and social considerations on the matter. Finally, chapter 5 presents innovations in medicine using generative AI including comparison of various AI models, and chapter 6 draws the conclusions and possible future directions of the subject. Some of the research questions that we are planning to answer in this article include the following - RQ1: How has artificial intelligence evolved in healthcare from early rule-based systems (such as MYCIN) to contemporary generative models like GPT-4 and Med-Gemini? RQ2: In what ways do contemporary AI models outperform earlier systems in medical applications? RQ3: How does AI contribute to increasing physician trust and adoption of AI-supported diagnostic tools? RQ4: What ethical concerns arise with the implementation of AI in healthcare, particularly regarding data privacy, bias, and consent? RQ5: How can policymakers and developers align AI innovation with ethical standards to ensure responsible deployment? RQ6: How do generative AI models differ in their medical applications, and what are their comparative strengths and weaknesses? RQ7: What are the most promising current and emerging applications of generative AI in healthcare?

2. HISTORICAL CONTEXT OF AI IN HEALTHCARE

To address RQ1 and RQ2, we have researched the evolution of artificial intelligence in healthcare over time. It can be said to have originated from the application of rule-based expert systems in the 1950s to simulate human decision-making using pre-set rules. The most important event on the matter was undoubtedly MYCIN, a system designed in the 1970s to help diagnose bacterial infections and recommend their treatment with antibiotics. This system, while never implemented in clinical practice, powerfully indicated the potential of AI for transformational change in medicine. It used backward chaining reasoning and a

¹ Faculty of Engineering in Foreign Languages, National University of Science and Technology POLITEHNICA Bucharest, Romania.

² Faculty of Automatic Control and Computer Science, National University of Science and Technology POLITEHNICA Bucharest, Romania.

E-mails: iulia.stanica@upb.ro, radu.constantin1005@stud.acs.upb.ro, ionut.lutan@stud.acs.upb.ro, andrei.chitu1312@stud.acs.upb.ro, costin.boiangiu@cs.pub.ro (Corresponding author)

knowledge base of approximately 600 production rules for recommending therapies in infections such as bacteremia and meningitis [3]. Such reliance on static rules limited the system's scalability and transferability but highlighted the possibility of using artificial systems to supplement human expertise in clinical decision support.

MYCIN's contributions were extensively evaluated, showing its performance to be similar to that of infectious disease experts. In one study, its recommendations were compared to those of 10 specialists using data from 15 patients with positive blood cultures [4]. It comprised three main elements: a consultation module for processing input data and producing recommendations, an explanation system to clarify its reasoning, and a knowledge acquisition system that allowed domain experts to update rules without programming skills. Designed to provide clinically actionable advice while explaining its decision-making process, MYCIN presented a novel integration of symbolic reasoning into medicine. However, it relied on extensive computing resources which were inaccessible to most hospitals at the time. Despite its limitations, MYCIN laid the foundations for AI-based decision support systems, and consequently, enhancements were proposed in the following years for networking technology to broaden access and adapt it for smaller, portable machines.

In the 1990s, machine learning was adopted in the field of healthcare, which brought changes towards an entirely different patient care paradigm, from strictly rule-based systems to something more flexible and driven by data. These machine learning algorithms could learn and improve from data, creating predictive modeling, early warning systems, and basic clinical decision-support tools. This revolution coincided with the proliferation of Electronic Health Records (EHR), which supplied the large datasets necessary for training those models. This enabled high-accuracy identification of high-risk patients, hence massive proactive interventions and fewer complications. During this period, artificial neural networks (ANNs) were introduced to the medical field, attaining a central role in the diagnosis and classification of diseases. The early applications positioned ANNs as multirole actors, targeting classification, prediction, and diagnosis. Meso-level applications focused on strategy-oriented decisions such as cost forecasting or technology adoption, while macro-level applications addressed system-wide models like risk adjustment and revenue generation. (Nida Shahid, 2019). A review of over 3000 articles [5] categorized ANN applications into three levels: micro (patient diagnostics), meso (intra-organizational decisions), and macro (system-wide behaviors). ANN models were associated with accuracy ranging from 50% to 100% and built with architectures like standard ANNs, feed-forward networks, and hybrid models. Findings confirm ANNs' increasing role in healthcare, from resource allocation to better patient flow and improved care quality.

The 2010s were the decade that saw the emergence of deep learning (DL) with a specific emphasis on convolutional neural networks (CNNs). Such technologies revolutionized medical imaging, resulting in human-level accuracy, for instance, when detecting conditions such as diabetic retinopathy, breast cancer, or lung diseases. One of the most significant developments in this field is Google DeepMind [6], which designed an AI tool with the same

precise understanding as an ophthalmologist for diagnosing retinal diseases. The initiatives were timely, considering that in Europe, age-related macular degeneration (AMD) affected almost 25% of people above 60 years, with about 15% moving toward exudative AMD (exAMD), the advanced and blinding form of the disease [7]. Curating a dataset on retinal images in partnership with companies and hospitals for training an AI system at DeepMind was established. The objective of developing this system is to provide exAMD occurrence prediction using two deep convolutional neural networks. Clinically, the system predicted by 90% specificity the high-risk exAMD patients which was compared to 6 retinal specialists (three ophthalmologists and three optometrists), each with at least a decade of experience. These findings emphasized that the variability in assessments among experts was minimized by a standardized analysis to which the AI system defaulted. The segmentation of eye scans anatomically also provided doctors with a visual representation of the changes in retinal tissues over time. Hence, during the observation of a patient over 13 months, the system could follow changes in anatomy and forecast exAMD transience much before visible signs of evidence were collected. This capability not only improves accuracy for other predictions but also gives clinicians evidence that can be acted upon to better understand disease progression.

Last but not least, the fusion of artificial intelligence with the Internet of Things (IoT) has become a growing application area in medicine, enabling real-time monitoring and data analysis [8]. Wearable devices like smartwatches and implanted sensors collect health data continuously, which AI algorithms process for anomaly detection and health event prediction. Key platforms advancing this field include the Apple Watch, AliveCor's KardiaMobile, and Android Wear, the most widely used platform in healthcare studies. Smartwatch applications provide diversified healthcare services, including health monitoring for elderly individuals (25%), Parkinson's management (21%), and drug adherence for chronic conditions (13%). Android Wear led health research (46%), with Samsung Galaxy Gear being the most used brand (25%). The most common sensors were accelerometers and gyroscopes, used in 67% of the studies. Some studies also used smartwatches as aid devices with screens or voice inputs rather than relying solely on sensors [9]. These applications are especially significant given the projected growth in the elderly population in the U.S., where a majority of older adults manage multiple chronic conditions. These advancements represent a shift from reactive to preventive care, allowing earlier interventions and reducing the burden on healthcare systems.

3. EXPLAINABLE AI: THE "PERFECT" ASSISTANT

Based on the previous studies on the historical evolution and impact of AI, one common problem emerges: How does AI contribute to increasing physician trust and adoption of AI-supported diagnostic tools? (RQ3). The answer lies in Explainable Artificial Intelligence (XAI), a set of techniques and tools that make AI model predictions interpretable and understandable for humans. Unlike classical "black-box" models, which offer limited insight into their operations, XAI provides interpretation, allowing users to understand the "how and why" behind an AI's predictions or decisions [10].

As the stakes in healthcare are extremely high, AI can have enormous consequences on patient outcomes, and clinicians need to understand its use, especially given the trust and understanding already established among a broader consumer audience. Some of the possible advantages provided by Explainable AI (XAI) include:

- **Strengthened Trust:** XAI enables medical professionals and patients to understand the reasoning behind AI-based predictions, fostering confidence in the new information presented.
- **Follow-Up Accountability:** XAI ensures models are transparent, allowing them to be held accountable for their decisions by end users.
- **Error Detection:** XAI helps identify biases or mistakes in AI systems, enabling healthcare practitioners to mitigate potential harm.
- **Regulatory Compliance:** XAI supports healthcare compliance regulations, like the GDPR, by improving interpretability and making automated decision-making more understandable.

The "black-box" nature of many AI models remains a significant hurdle to their adoption in healthcare, as clinicians are reluctant to rely on opaque decision-making processes, especially in high-stakes environments [11]. Transparency in AI models is essential for fostering trust among healthcare professionals, who are more likely to embrace tools that provide clear and interpretable outputs. Transparent systems are also better suited to meet regulatory standards and address ethical concerns related to accountability and fairness. By making AI recommendations interpretable, clinicians can validate these outputs against established medical knowledge, ensuring alignment with clinical standards and boosting confidence in AI-driven solutions.

Further studies utilized a publicly available diabetes dataset from Kaggle, consisting of 768 samples with 8 features, such as glucose level, blood pressure, and age. A Random Forest Classifier trained on this dataset achieved good results: precision = 0.74, recall = 0.70, and F1 score = 0.72. With SHAP inclusion, it was found that glucose level is the most significant predictor of diabetes, followed by age and body mass index (BMI). SHAP dependency plots revealed that patients younger than 30 years had a lower likelihood of diabetes, while those older than 30 showed a higher probability [12].

Another notable contribution to the medical industry during the COVID-19 pandemic was the use of XAI, which developed a proof of concept in public health management by providing transparency and actionable insights for critical decision-making. Transparent AI models capable of justifying predictive analytics for outbreak predictions empowered policymakers to allocate medical resources effectively and prioritize interventions wisely. XAI systems optimized resources by analyzing trends in infection and hospital bed usage, assisting in the strategic allocation of ventilators and personal protective equipment (PPE) to areas of highest need. Additionally, XAI fostered public trust by clarifying the rationale behind recommendations, reducing resistance to policies like lockdowns and vaccine rollouts, and promoting collaboration during the global crisis [13].

4. ETHICAL AND SOCIAL CONSIDERATIONS

As previously highlighted through the promising potential of XAI, the reason for its existence is to resolve ethical and social considerations (RQ4) that take center stage in the heavy dominance of the integration of AI in medical engineering. Applications span fields such as radiology, surgery, pathology, dermatology, ophthalmology, and general practice, using technologies for purposes like diagnostics, surgical assistance, and patient monitoring. This integration has prompted individuals from diverse backgrounds to address critical issues, including privacy concerns, mitigating biases, ensuring accountability, and fostering trust between stakeholders.

A scientific paper introduces the term Technology Availability Level (TAL) scale to gauge AI readiness and accessibility, ranging from TAL 0 (unknown / not feasible) to TAL 9 (publicly available) [14]. Based on this new scale, algorithms for computer-aided diagnosis and structured reports for eHealth (TAL 8, 9) show high social impact through applications in clinical decision support and improving workflow efficiency. AR/VR tools for advanced imaging and navigation (TAL 6, 7, 9) provide transformative capabilities, particularly in image-guided surgery and automated analysis. Companion robots for elderly care and big data analysis (TAL 2-9) enhance patient care and epidemiological research. Controversially, technologies like brain-machine interfaces (TAL 5-8), gene editing for superhumans (TAL 2, 6), and human-animal embryos (TAL 2, 4, 5) raise ethical concerns. At lower TAL levels, advancements such as autonomous AI systems for surgery (TAL 2-5) and the quest for immortality (TAL 1-3) spark debates about their societal risks. Bioterrorism applications, rated TAL 1-2, represent the most negative potential, emphasizing the wide spectrum of benefits and hazards associated with these technologies.

Furthermore, among World Health Organization priorities we encounter promoting fairness in healthcare, leveraging emerging technologies, and building public trust. However, it also raises significant concerns about patient privacy, data security, and the risk of disrupting traditional doctor-patient relationships. Potential pitfalls are starting to take a toll on the industry, beginning with the fragmentation of medicine into "fake-based," "patient-generated," and "scientifically tailored" practices, stressing the importance of upholding scientific and ethical standards. This led to an abrupt rise in intrigue for the ethics of AI when the focus switched to the zero-shot learning capabilities of LLMs in tasks such as diagnostic assistance, drug discovery, and personalized medicine [15]. Different adaptation strategies were taken into consideration, emphasizing fine-tuning methods for uni-modal and multi-modal LLMs to address challenges like medical question answering and processing biomedical literature. Identified challenges include limited model interpretability, dataset quality issues, and ethical implications in healthcare deployment.

One other key aspect of the rapid evolution of AI is the legal implications and the way policymakers can ensure the development of AI according to ethical standards to ensure responsible deployment (RQ5). Attention started growing towards frameworks like the GDPR in Europe and HIPAA in the United States which enforce stringent rules to safeguard patient information and outline protocols for addressing breaches. Additional regulations, such as Canada's PIPEDA and the UK's Data Protection Act, impose similar

requirements, ensuring global alignment in data privacy standards. Despite these measures, healthcare organizations remain highly vulnerable to cyberattacks, which highlights the urgent need for advanced security practices and innovative solutions like federated learning, which enable AI models to train on decentralized data without transferring sensitive information, effectively balancing privacy with utility while minimizing the risk of exposure [16].

As advanced AI technologies are often available only in well-resourced systems, this widens health discrepancies, leaving underserved populations behind—a concept known as the Matthew Principle [17]. Another concern is the costs of AI development and maintenance, which burden healthcare budgets, diverting resources from essential services. Resource dislocation is another issue, where jobs in healthcare are disrupted by AI automation, leading to employment challenges and a changing work environment.

Nonetheless, the advent of personalized medicine has brought transformative advancements in healthcare, sparking debates about the privacy and accessibility of patient data. Building upon traditional personalized medicine, the concept of Extended Personalized Medicine incorporates diverse data sources beyond genetics to tailor healthcare more dynamically and holistically. While traditional approaches emphasize genetic profiling, Extended Personalized Medicine includes biological, demographic, social, environmental, and lifestyle data, offering a more comprehensive framework for health management. Key data sources include biophysical metrics like bioelectromagnetic fields and biomarkers, alongside insights into brain function and connections. Social and demographic data, combined with lifestyle parameters such as sleep, stress, and physical activity, are collected via IoT devices and wearables. Behavioral sensors capture mood and physiological changes through smartphones and smart home technologies, while environmental factors like pollution and weather conditions complement traditional clinical data, including imaging scans and genomics [14]. Despite its transformative potential, implementing Extended Personalized Medicine poses technical challenges. The integration of heterogeneous datasets requires advanced AI tools and interactive visualization techniques, such as augmented and virtual reality, to enable practical applications. Additionally, the accumulation of intimate data raises privacy and security concerns, necessitating robust governance frameworks to maintain patient trust. Addressing these challenges while leveraging benefits like individualized treatment and proactive health management will shape the future of this innovative paradigm.

Thus, establishing trust in the public's perception of AI through techniques like timely communication, ethical robustness, and patient-centered approaches is crucial for healthcare success. Transparency about AI applications helps patients understand how data is collected, processed, and utilized in diagnosis and treatment plans. Ethical practices involve creating guidelines for data use that safeguard patients and respect their consent. Research shows that 78% of patients express concerns about data usage, while over 60% feel more comfortable with healthcare systems that transparently discuss AI policies [18]. Informed consent processes must ensure patients comprehend AI's benefits, limitations, and risks. Public awareness campaigns also play a vital role in reducing skepticism and improving acceptance. For instance, Britain's NHS initiative demonstrated that educating people on AI's

capabilities and ensuring data protection increased public acceptability by over 50%. Additionally, 74% of patients feel reassured when informed about how AI improves diagnostic accuracy, which research shows may rise by 15%-20% for conditions like diabetic retinopathy or breast cancer detection. These efforts emphasize open communication among stakeholders—patients, healthcare providers, and AI developers—to foster trust and the responsible adoption of AI in medicine [19].

5. IMPACTFUL AI INNOVATIONS IN MEDICINE

The current chapter focuses on tackling RQ6 and comparing different generative AI models, including their comparative strengths and weaknesses in medical applications. Recent increase in generative AI is impacting the medical world, contributing to various innovations. OpenAI's world-renowned GPT (Generative Pre-trained Transformer) family comprises advanced language models capable of producing coherent and human-like text for various applications. These models have evolved from GPT-2 through GPT-3 and GPT-3.5, culminating in GPT-4, with each version boasting improved complexity, accuracy, and natural language understanding. Similarly, Google's advancements include MedPaLM 1 and MedPaLM 2, which adapt large language models specifically for clinical knowledge and reasoning. MedPaLM 2, for instance, excels at answering specific medical questions, summarizing patient notes, and supporting diagnostic decision-making using datasets tailored for tasks like medical question answering, clinical note summarization, and diagnostic decision support. Tuned for health-related tasks, MedPaLM 2 has demonstrated remarkable performance in medical reasoning and diagnostics, while GPT-4 scores well in the generation and interpretation of medical content. These models differ from specialized LLMs like Microsoft BioGPT and PubMedBERT, which exhibit varying accuracy, contextual understanding, and clinical applicability.

Medical question-answering datasets and diagnostic accuracy tests serve as benchmarks to evaluate these models, providing insights into their potential for improving healthcare delivery and research. A subsequent table highlights the performance dominance of models developed by OpenAI and Google (Table 1).

Even though it is the one with the least accuracy, amazing advancements have been made by GPT-4, the state-of-the-art large language model, particularly in medical competency assessments [20]. This has to be mentioned because the GPT-4 is not a model created specifically for medical assessments. Thus, the results shown by this model are much more impressive than the others that were specifically pre-trained to perform better on medical datasets. Qualitatively, GPT-4 excelled in explaining medical reasoning, personalizing explanations for learners, and offering counterfactual scenarios during interactive case-based discussions. It demonstrated robustness in processing both text and image-based questions, minimizing content memorization, which further highlights its potential in medical education and assessment. These advancements render GPT-4 a powerful tool for education and evaluation, with promising applications in clinical decision-making, provided challenges around accuracy and safety are resolved before field deployment.

Table 1
Accurate LLMs used in Medical Sciences

Model	Description	Performance metrics	Acc.
GPT-4	LLM by OpenAI, GPT-4 excels in various medical tasks, passing the U.S. Medical Licensing Examination (USMLE).	Exceeded the passing score on USMLE by over 20 points, showing high competency in medical reasoning and interpretation.	>80% (USMLE)
PaLM	Developed by Google, PaLM is a 540-billion-parameter transformer-based LLM capable of handling diverse tasks, including medical scenarios.	Achieved state-of-the-art performance on various benchmarks, including 85% accuracy on medical question-answering datasets and clinical text summarization.	85%
Med - PaLM 2	Developed by Google, Med-PaLM 2 is fine-tuned on medical data and achieves state-of-the-art performance on medical benchmarks.	Scored 86.5% on the MedQA dataset, improving upon its predecessor by over 19%.	86.5%
ChatDoctor	Fine-tuned on a large dataset of patient-doctor dialogues, ChatDoctor provides accurate and context-aware medical advice.	Significant improvement in understanding patient inquiries and precise medical responses with up to 88% accuracy in dialogs.	88%
ClinicalGPT	Designed for clinical scenarios, ClinicalGPT integrates diverse real-world medical data to perform effectively across clinical tasks.	Outperformed others in tasks such as medical knowledge question-answering and diagnostic analysis with high accuracy in real-world medical cases.	~88%
Med-Gemini	A state-of-the-art multimodal AI model designed for medicine, excelling in reasoning, multimodal data analysis, and clinical knowledge.	Recorded the highest accuracy in medical AI, processing data 4x faster, and reducing diagnostic errors by 30% setting a new standard.	92%

Text mining and knowledge discovery in biomedical literature have gained prominence due to the rapidly growing pool of biomedical data, with over 30 million articles on PubMed alone. This wealth of information necessitates automated approaches for efficient knowledge extraction. General-purpose pre-trained language models such as BERT-like for understanding tasks and GPT-like for text generation have demonstrated success in general natural language processing (NLP) applications [20]. However, these models often underperform in biomedical applications due to the domain shift. Domain-specific models like

BioBERT and PubMedBERT have shown improved performance in understanding tasks, but effective modeling for generation tasks has remained underdeveloped.

The model that stood out the most is Med-Gemini [21], achieving state-of-the-art performance with an impressive accuracy of 92%. Med-Gemini is the newest member of a family of multimodal models designed for medically targeted applications, which includes innovations like web search capabilities and customizable encoders, allowing for seamless adaptation to new data modalities. Med-Gemini delivers exceptionally high performance, achieving SoTA on 10 out of 14 medical benchmarks, consistently outperforming the GPT-4 family on similar tasks [20]. Notably, it achieved 91.1% accuracy on the MedQA (USMLE) benchmark, a 4.6% improvement over Med-PaLM 2, using an innovative uncertainty-guided search strategy. Advanced multimodal diagnostic tasks, such as those from the New England Journal of Medicine and GeneTuring benchmarks, further demonstrated its reliability, maintaining baseline test results while excelling in long-context processing tasks like medical video question answering and retrieving large health records. Medical professionals using Med-Gemini surpassed human expert-level performance in summarizing medical texts, drafting reference letters, and communicating in simplified layman language across various fields of medicine. The integration of self-training, web search, and customizable encoders enhances its potential for clinical reasoning and versatility across diverse data types. This represents a major milestone in AI-driven medicine, showcasing strong reasoning, multimodal integration, and long-context processing. Such groundbreaking advances reiterate Med-Gemini's potential to transform medical research, education, and clinical practice.

6. CONCLUSIONS

Artificial intelligence ignited a transformational spark across systems, stakeholders, and operations. At the systems level, AI perfects diagnostics, streamlines workflows, enables early disease detection, removes bottlenecks, and reduces costs while expanding care access. For healthcare professionals, AI places robust tools in their hands, automating routine tasks and enhancing decision-making in complex cases, while strengthening their connection to society. Additionally, wearables and personalized therapies empower patients to actively manage their health (RQ7). However, these advancements demand robust policies and regulations, ensuring ethical implementation, data protection, and equity in healthcare delivery.

AI signals the advent of a data-driven and consumer-oriented age in healthcare. Emerging technologies such as generative AI, digital twins, and predictive analytics promise to transform healthcare systems comprehensively. The future success of AI hinges on finding a balance between trade-offs and challenges, enabling its adaptive application to create an all-encompassing, efficient system.

A potential bright future lies in merging human expertise with the immense processing power of machines, delivering more effective, accessible, and affordable healthcare. This must be achieved in an ethical and accountable manner, paving the way for a healthier and sustainable future, as emphasized by the critical changes required for seamless AI

integration in healthcare environments.

CREDIT AUTHORSHIP CONTRIBUTION STATEMENT

Iulia-Cristina Stanica: writing original draft preparation, review, and supervision, ethical considerations analysis

Radu Georgian Constantin: writing original draft preparation, state-of-the-art analysis

Ionut-Sebastian Lutan: writing original draft preparation, comparison analysis

Andrei-Alexandru Chitu: writing original draft preparation, research of medical applications

Costin-Anton Boiangiu: supervision, project administration

Received on March 26, 2025

REFERENCES

1. S. Palanisamy, T. Ramasamy, *Medi-Net: Cloud-based framework for medical data retrieval system using deep learning*, Rev. Roum. Des Sci. Tech. — Série Électrotechnique Énergétique, **69**, 2, pp. 255–260 (2024).
2. A. Ramaiah, P.D. Balasubramanian, A. Appathurai, N.A. Muthukumar, *Detection Of Parkinson's disease via Clifford gradient-based recurrent neural network using multi-dimensional data*, Rev. Roum. Des Sci. Tech. — Série Électrotechnique Énergétique, **69**, 1, pp. 103–108 (2024).
3. E.H. Shortliffe, *Mycin: A Knowledge-Based Computer Program Applied To Infectious Diseases*, Proc Annu Symp Comput Appl Med Care (1977).
4. V.L. Yu, L.M. Fagan, S.W. Bennett, W.J. Clancey, A.C. Scott, J.F. Hannigan, R.L. Blum, B.G. Buchanan, S.N. Cohen, *An Evaluation Of Mycin's Advice*, 1979.
5. N. Shahid, T. Rappon, W. Berta, *Applications Of Artificial Neural Networks In Health Care Organizational Decision-Making: A Scoping Review*, Plos One, **14**, 2, P. E0212356 (2019).
6. ***Google DeepMind.
7. J. Yim, et al., *Predicting conversion to wet age-related macular degeneration using Deep Learning*, Nat. Med. 2020, **26**, 6, pp. 892–899 (2020).
8. R.A. Craciun, R.N. Pietraru, M.A. Moisescu, *Internet Of Things Platform Benchmark: An Artificial Intelligence Assessment*, Rev. Roum. Des Sci. Tech. — Série Électrotechnique Énergétique, **69**, 1, pp. 97–102 (2024).
9. T.C. Lu, C.M. Fu, M.H.M. Ma, C.C. Fang, A.M. Turner, *Healthcare Applications Of Smart Watches: A Systematic Review*, Appl. Clin. Inform., **7**, 3, pp. 850 (2016).
10. A. Varghese, J. Varghese, J. Biju, R. Thomas, M. Thomas, K. Jyothi, *Explainable Ai In Healthcare Applications*, Int. Res. J. Adv. Eng. Manag., **2**, 12, pp. 3671–3679 (2024).
11. A. Rai, *Explainable Ai: From Black Box To Glass Box*, J. Acad. Mark. Sci., **48**, 1, pp. 137–141 (2020).
12. M. Panda, S.R. Mahanta, *Explainable Artificial intelligence for healthcare applications using random forest classifier with lime and shap*, Arxiv (2023).
13. O. Lohaj, J. Paralič, P. Bednár, Z. Paraličová, M. Huba, *Unraveling Covid-19 dynamics via machine learning and Xai: Investigating variant influence and prognostic classification*, Mach. Learn. Knowl. Extr. 2023, **5**, 4, pp. 1266–1281 (2023).
14. E. Gómez-González, E. Gomez, J. Marquez-Rivas, M. Guerrero-Claro, I. Fernandez-Lizarran, M.I. Relimpio-Lopez, M.E. Dorado, M.J. Mayorga-Buiza, G. Izquierdo-Ayuso, L. Capitan-Morales, *Artificial intelligence in medicine and healthcare: A review and classification of current and near-future applications and their ethical and social impact*, preprint, (2025).
15. C. Wang, M. Li, J. He, Z. Wang, E. Darzi, Z. Chen, J. Ye, T. Li, Y. Su, J. Ke, K. Qu, S. Li, Y. Yu, P. Lio, T. Wang, Y. G. Wang, Y. Shen, *A survey for large language models in biomedicine*, Preprint (2024).
16. A.K. Momani, *Implications of artificial intelligence on health data privacy and confidentiality*, Deep Learning Monitor (2025).
17. X. Wang, N.X. Zhang, H. He, R. Nguyen, K.H. Yu, H. Deng, C. Brandt, D.S. Bitterman, L. Pan, C.Y. Cheng, J. Zou, D. Liu, *Safety challenges of Ai in medicine in the era of large language models*, preprint (2025).
18. V. Sivaraman, L.A. Bukowski, J. Levin, J.M. Kahn, A. Perer, *Ignore, trust, or negotiate: Understanding clinician acceptance of AI-based treatment recommendations in health care*; CHI Conference on Human Factors in Computing Systems, pp. 18 (2023).
19. N. Scharowski, S.A.C. Perrig, F. Brühlmann, *Trust and reliance in Xai-Distinguishing between attitudinal and behavioral measures*, arXiv (2022).
20. H. Nori, N. King, S.M. McKinney, D. Carignan, E. Horvitz, *Capabilities of GPT-4 on medical challenge problems*, Arxiv (2023).
21. K. Saab, et al., *Capabilities of Gemini models in medicine*, Deep. Google Research (2024).