Design and Implementation of an Intelligent Virtual Medical Agent for Health Risk Assessment

Anita Petreska[1], Igor Nedelkovski[2], Andrijana Bocevska[3], Blagoj Ristevski[4]

[1,2,3,4] *Faculty of Information and Communication Technologies - Bitola, University "St Kliment Ohridski" – Bitola Partizanska bb, Bitola, Republic of North Macedonia*

petreska.anita@uklo.edu.mk; igor.nedelkovski@uklo.edu.mk; andrijana.bocevska@uklo.edu.mk; blagoj.ristevski@uklo.edu.mk

Abstract:
The development of artificial intelligence (AI) has opened new opportunities for the digital transformation of the healthcare system, enabling automation of initial symptom assessment, intelligent decision support, and improved communication between patients and healthcare professionals. This paper presents the design and conceptual framework of an intelligent virtual medical agent (IVMA) that combines machine learning (ML) and natural language processing (NLP) for interactive, contextual, and explainable healthcare assessment. Methodologically, the research is based on a Design Science Research (DSR) approach and includes architectural modeling, system logic design, and development of a demonstration prototype (Python/Gradio) with local processing, visual BMI analysis, and automatic PDF report generation. The architecture applies the principles of Explainable AI (XAI) and privacy-by-design for transparency, confidentiality, and compliance with GDPR and HIPAA. The proposed model demonstrates technical feasibility, interoperability, and potential for application in preventive medicine, digital triage, and clinical decision-making, especially in resource-limited settings.

Keywords:
Artificial Intelligence, Machine Learning, Natural Language Processing, Virtual Medical Agent, Explainable AI

1. Introduction

Modern healthcare systems face increasing demand, staff shortages, and growing administrative burdens, leading to limited service availability and prolonged waiting times [1] [2]. In response to these challenges, digital transformation, driven by artificial intelligence (AI), machine learning (ML), and natural language processing (NLP), is becoming a key direction for clinical process optimization, automation, and medical decision support [3][8][9][10].

Virtual engineering, as an integrative discipline that unites engineering principles, cognitive sciences, and ethical standards, enables the development of intelligent digital agents that mediate between patients and healthcare institutions. These agents do not replace human expertise, but rather complement it through digital triage, automated symptom collection, and referral to appropriate medical services [4] [5] [11] [12]. This approach enables more efficient use of healthcare resources and the creation of new models for preventive and personalized medicine.

The COVID-19 pandemic has highlighted the importance of scalable digital solutions and accelerated the implementation of virtual health assistants. Platforms such as Infermedica, Ada Health, and Babylon Health have demonstrated significant potential for initial symptom assessment and clinical decision support, confirming the applicability of such systems even in resource-constrained healthcare settings [6].

In this context, this paper presents a conceptual and architectural model of an intelligent medical agent that combines NLP and machine learning for automated initial symptom assessment and clinical judgment support. The main goal is to develop an integrated architectural framework that enables cognitive, technological, and ethical alignment in the design of such systems [13] [14] [15]. In addition, the research proposes methodological and ethical guidelines for their design and presents a simulation framework that can be used as a basis for future pilot implementations in a clinical context.

Methodologically, the paper is based on the Design Science Research (DSR) approach, which enables integrated modeling of technological, cognitive, and ethical aspects in the development of digital solutions [7], [16]. Rather than full implementation, the emphasis is placed on defining the systemic, architectural, and logical principles that support the engineering development and future integration of such systems into healthcare infrastructures.

The rest of the paper is structured as follows. The second section focuses on the theoretical foundations of AI, ML, NLP, and medical ontologies, while the next section presents the concept of virtual engineering and its application in healthcare, with a review of the theoretical foundations, architecture, and design of empathic virtual medical agents as intelligent decision support systems. The subsequent section concerns the practical implementation through Gradio. The fifth section addresses ethical and regulatory aspects, and the last section provides concluding remarks and directions for future research and integration into real healthcare contexts.

2. Theoretical foundations

The design of intelligent digital agents in healthcare is an interdisciplinary process that brings together the principles of AI, ML, NLP, and medical semantics. This section lays the foundation for understanding the cognitive, technological, and communication behavior of such systems.

2.1 Artificial intelligence, machine and deep learning in healthcare

AI is now a key driver of the digital transformation in healthcare, enabling analysis, prediction, and automation of clinical processes. ML enables models that learn from data and recognize relationships between symptoms and diagnoses, using algorithms such as Support Vector Machine (SVM), Random Forest, and XGBoost [17]. Deep learning (DL), through architectures such as Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), and transformer models, enables the analysis of complex data, such as electronic health records, clinical notes, and medical dialogues [18].

Current trends include Federated Learning (FL), which enables modeling of decentralized data sources without compromising privacy, while complying with GDPR and HIPAA regulations [19] [20]. In parallel, Explainable AI (XAI) is developing techniques such as SHAP and LIME, which enable visualization and interpretation of model decisions, increasing confidence in clinical applications [21] [22].

In this context, AI and ML do not represent a replacement for medical staff, but rather an intelligent cognitive layer that collects, structures, and interprets data, contributing to increased accuracy, better efficiency, and more informed clinical decision-making.

2.2 Natural Language Processing and Medical Ontologies

Natural language processing (NLP) plays a central role in the design of systems that understand, interpret, and answer medical questions. Models such as BioBERT, ClinicalBERT, and Med-PaLM have been trained on large biomedical corpora (e.g., PubMed, MIMIC-III) and are used for medical entity recognition, symptom analysis, and triage through dialogue [23] [24] [25] [26].

The integration of medical ontologies such as SNOMED CT, ICD, and UMLS enables semantic interoperability and standardization of clinical concepts. In this way, NLP agents connect linguistic analysis with formalized medical knowledge, which is a prerequisite for reliable, explainable, and ethical healthcare recommendations [27] [28].

Modern solutions, such as the Infermedica API, combine these principles and provide interactive triage based on ontological knowledge, NLP inference, and structured decision support.

Within this theoretical framework, AI, ML, NLP, and medical ontologies together form a cognitive and semantic foundation that enables medical agents to understand context, analyze symptoms, and communicate in natural language. These technologies provide a foundation for developing ethical, transparent, and interoperable systems that support the modern model of digital health triage and personalized medical support.

3. Introduction to virtual engineering

Virtual engineering is a modern interdisciplinary paradigm that brings together simulation, digital prototyping, and cognitive virtualization to design and evaluate complex systems in digital environments. In the healthcare context, this means that clinical workflows can be modeled and tested through intelligent virtual agents—from conversational interfaces to cognitive decision support systems—that operate based on AI, NLP, and medical reasoning. These agents serve as the first digital point of contact, structuring symptoms, initiating triage, and communicating recommendations in an ethically and regulatory-compliant manner.

Virtual engineering enables the creation of simulated healthcare environments where technological, cognitive, and ethical principles are integrated before practical implementation. In this way, a bridge is created between academic theory and real-world clinical transformation to reduce risk and ensure methodological rigor [29] [30] [31].

3.1 Theoretical foundations of the concept

The concept is based on four pillars: digital representation of clinical entities (including digital twins) for realistic simulation; a modular architecture that allows for independent development and integration of NLU, inference, and interfaces; simulation before real implementation for greater stability; and a human-centric design that allows for dynamic adaptation of the agent according to the context and user profile [32]. These principles are applied in this research to design an intelligent medical agent that links symptom narratives to standardized clinical concepts (SNOMED CT, UMLS) and generates explainable recommendations [33].

3.2 Virtual medical avatars: role and functionalities

A virtual medical avatar is defined as an intelligent communicative intermediary between the patient and medical knowledge. Its NLU layer converts text into structured input, bound to ontologies (SNOMED CT, ICD-11), while reasoning logic initiates triage or diagnostic recommendations. The interface can be textual, voice, or visual (XR), but the principles of semantic interoperability, ethical transparency, and GDPR/HIPAA compliance remain central.

Psychological factors play a key role in the acceptance of these systems. Empathetic expression, clear communication about limitations, and data confidentiality significantly increase user satisfaction and trust, especially in domains such as mental health and preventive medicine.

Table 1 presents an overview of the key design principles for empathic interaction with a virtual medical agent, with their functional role and expected psychological effect on users.

Table 9:
Design principles for empathic interaction with a virtual medical agent

Principle	Function	Expected effect
Empathic verbalization	Emotionally sensitive communication	Greater confidence and comfort
Adaptive dialogue logic	Contextual Personalization	Perception of intelligence and attention
Multimodal empathy	Visual/audio support	Reduced psychological distance
Transparency (XAI)	Explanatory recommendations	Cognitive trust
Privacy and non-judgment	Identity and data protection	Greater honesty in responses
Cultural adaptability	Language and visual localization	Inclusion without alienation

3.3 Engineering architecture

The proposed system is organized into several functional layers: an interaction layer for text, voice, and XR communication; an NLU layer for intent and symptom identification; a dialog management layer for contextual coherence; a clinical reasoning core with ontological knowledge and probabilistic logic (or integration with the Infermedica API); and an XAI security layer for explainability, encryption, and regulatory compliance (GDPR/HIPAA). This modular structure ensures interoperability (HL7 FHIR and OMOP) and the possibility of multilingual expansion [34] [35].

4. Practical Implementation: Virtual Agent for Health Assessment with Gradio

Within the framework of this research, a demonstration prototype of an intelligent virtual health agent, designed for interactive assessment of chronic disease risk, has been developed. The prototype serves as a proof of concept for the technical feasibility and cognitive coherence of the proposed architecture. The system is built in Python, using the Gradio library to create an intuitive user interface, scikit-learn for machine learning, and matplotlib for visualization of the results. Through this approach, the practical application of ML and NLP principles in the context of digital triage, self-assessment, and preventive medicine is demonstrated.

During the interaction process, the user enters biometric and clinical parameters such as age, gender, height, weight, blood pressure, sugar level, symptoms, and chronic conditions. Based on this data, the system calculates a body mass index (BMI) and, through a trained Random Forest model, performs a risk assessment for the development of a chronic disease. The results are presented through a visual graphic display with colored zones, indicating different levels of risk (low, medium, high), supplemented with a textual interpretation and a recommendation for further health care.

In addition to the interactive assessment, the application allows for tracking interaction history, with each session being saved as an individual PDF report. The report includes a graphical representation, descriptive assessment, and recommendations, which increase the traceability and transparency of the results. All data is processed locally, without external storage or transmission, ensuring GDPR compliance and respecting the principles of privacy-by-design.

Table 2 shows the main functional components of the system, their role, and the effects of their application, which illustrates the logic of architectural integration and the cognitive flow of the application.

Table 2:
Main functional components of the system

Component	Function	Result / Effect
Data entry	Reception of individual parameters via an interactive interface	A personalized basis for analysis
Machine learning model	Risk classification with the Random Forest model	Risk calculation (low, medium, high)
Visualization	Graphical display of BMI with color zones	Improved perception and understanding of results
Explanatory Logic (XAI)	Textual explanation of the result	Greater trust, transparency, and interpretability
History and report	View and generate PDF reports	Traceability and documentation of sessions

The visual design of the interface is structured into three functional units: data entry, Figure 1, results display, and history review with report. The user enters data via a simple form, and the system automatically generates a visual assessment with a graphical display of BMI and accompanying textual recommendation. The history allows for a review of previous sessions and downloading of an individualized report, thus ensuring continuity and documentation of health self-evaluation.

Although the prototype is of a demonstration nature, its architecture allows for real expansion with additional functionalities, such as processing of textual descriptive symptoms via NLP, support for multimodal data (images, audio), and integration of XAI techniques, such as SHAP and LIME, for better interpretability and transparency of results [36].

Future research phases include a methodological assessment of the accuracy, interpretability, and user experience of the system. This plan is not a formal evaluation, but rather an orientation framework that defines the future direction of the experimental phase. The evaluation will be based on synthetic scenarios and sample dialogues to test the cognitive coherence and stability of the agent in a simulated environment. Key performance indicators include accuracy, precision, recall, F1-score, and Top-N, which will allow for a quantitative assessment of the model's classification performance. Additionally, a subjective UX survey will be conducted with a small group of test users to gain insight into their perception of the system's intuitiveness, trustworthiness and understandability. In this way, the assessment will ensure a balance between technical accuracy and user experience, which is an essential aspect for the successful integration of intelligent agents into healthcare practice. Such an evaluation approach is directional and will serve as a basis for future pilot testing in a clinical context, following the development of a fully functional prototype and the provision of real-world health data. In doing so, the system not only demonstrates technical feasibility but also creates a platform for sustainable ethical, regulatory, and clinical integration in future phases of research.

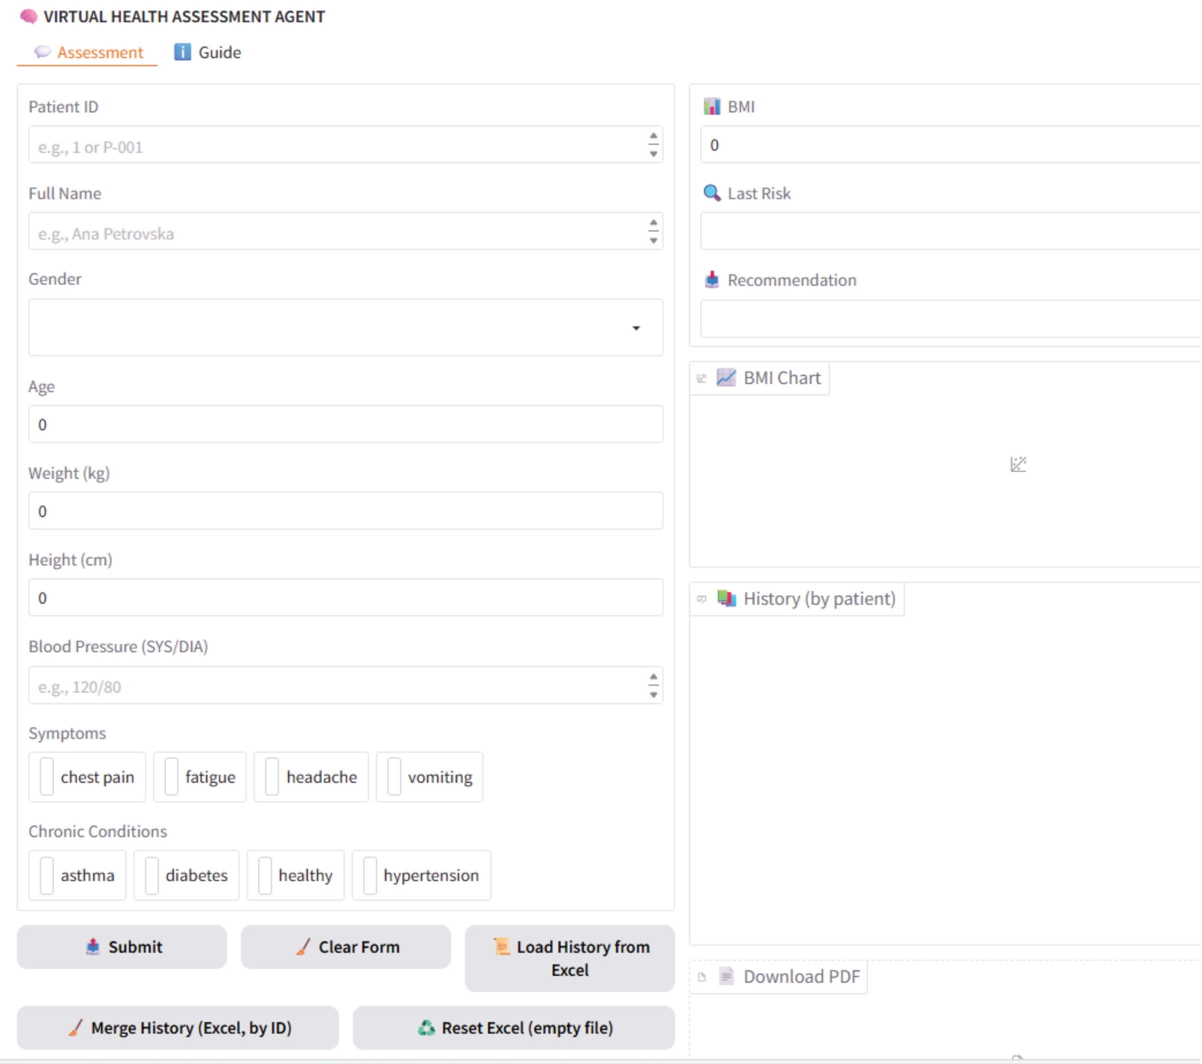

Figure 1: Data entry interface of Virtual Agent for Health Assessment

5. Ethical, regulatory, and security aspects in the design of intelligent medical agents

The integration of intelligent medical agents into clinical processes represents not only a technological advance but also a significant ethical and regulatory challenge. To ensure trust, safety, and social acceptability, such systems must be based on the principles of non-harm, autonomy, fairness, accountability, and transparency, especially when mediating initial (pre-diagnostic) interactions with patients.

Ethical operability implies a design that is not only functionally correct but also compliant with European and international regulations, such as GDPR, HIPAA, and Software as a Medical Device (SaMD), with built-in mechanisms for fair algorithms, auditability, and traceability of decisions. In this way, the system demonstrates ethical confidentiality, prevents discriminatory effects, and enables responsible use of technology in clinical practice.

5.1 Compliance and data protection

AI-based systems in healthcare process highly sensitive personal data, which requires the application of the principles of privacy-by-design and compliance-by-design. In the proposed model, all data is processed locally and sessionally, without storing personal identifiers, which minimizes regulatory risk and eliminates the possibility of unauthorized access [37].

Mechanisms such as TLS encryption, pseudonymization, access control, audit, and data revocation are implemented, in accordance with GDPR, HIPAA, and HL7/FHIR standards; privacy-by-design and compliance-by-design approaches are built into the architectural phase. This approach strengthens patient trust, as it ensures transparency in processing, interoperability with existing healthcare platforms, and consistent compliance with regulatory requirements of the European Union and the United States [38] [39].

5.2 Ethical AI and explainability

The ethical application of AI in healthcare implies that systems should be explainable – XAI, especially when providing automated recommendations or diagnostic guidance. The proposed model applies micro-XAI layers that offer concise, understandable interpretations of the decisions made [39]. For example: "Symptoms contributing to the assessment: abdominal pain, nausea, fever (2 days). Possible condition: gastroenteritis."

These micro-explanations enable cognitive transparency and increase user trust, as the system is not a “black box” but a partner in understanding symptoms. Potential algorithmic bias (linguistic, demographic, or cultural) is addressed through localized corpora, Named Entity Recognition (NER), and metadata analysis for fairness [40].

The human-in-the-loop principle is applied, whereby human intervention remains essential in situations of uncertainty, ambiguity, or contradictory data. This ensures a balance between automation and ethical control, which is especially important when applied in clinical contexts [41].

A systematic review of the key ethical principles and corresponding mechanisms implemented within the proposed model, as well as the expected effects of their implementation in a real healthcare context, is presented in Table 3. This review confirms the integration of technological and ethical components in the system design.

Table 3: Ethical principles and applied mechanisms in the proposed system

Ethical principle	Applied mechanism	Expected effect
Privacy and confidentiality	Pseudonymization, TLS encryption, and offline processing	Personal data protection and risk reduction
Explainability (XAI)	SHAP/LIME interpretations and micro-explanations	Transparency and greater user trust
Algorithmic fairness	Federated and fair learning, metadata analysis	Reducing bias towards certain groups
Data control and traceability	Audit trail, user consent, audit	Compliance with GDPR and HIPAA regulations
Human intervention	Involving a clinician in decision-making	Security and control in situations of

Ethical principle	Applied mechanism	Expected effect
(Human-in-the-loop)		uncertainty

5.3 Limitations and future directions

Although the research provides a solid conceptual and architectural framework, based on technological and ethical alignment, empirical clinical validation is lacking at this stage. The main limitations stem from limited local datasets, the unavailability of national NLP resources, and insufficient standardization of medical semantics in the Macedonian language.

In the future phases, a functional prototype (MVP) is planned to be developed with the integration of Rasa Framework and Infermedica API, which will enable the simulation of real clinical scenarios and testing of the cognitive consistency of the system. Pilot testing in healthcare institutions is also planned, with the involvement of medical professionals to validate accuracy, explainability, and user experience. Additionally, the application of large language models (LLMs) for multilingual communication and improved contextual understanding of symptoms will be investigated, which will increase the applicability of the system in various healthcare settings.

These activities represent a logical extension of this research and will enable the development of ethically responsible, interpretable, and reliable medical agents, which combine safety, transparency, and clinical applicability – prerequisites for their real integration into modern healthcare systems [41] [42].

6. Conclusion

This research establishes a conceptual, architectural, and methodological framework for the design of an intelligent medical agent that integrates the principles of NLP, dialogic logic, and clinical reasoning. Through a synthesis of machine learning, semantic ontologies, and explainable algorithms, a system is modeled that represents an innovative engineering concept for simulated health triage and interactive symptom assessment in a digital health environment.

The proposed architecture is modular and hybrid, with clearly defined functional layers that ensure interoperability, scalability, and transparency. A special value is the separation of cognitive reasoning from dialogic logic, which allows integration with reasoning-as-a-service platforms and the creation of multilingual, contextually aware interaction. The architecture incorporates XAI mechanisms for interpretability and privacy-by-design principles for full compliance with GDPR and HIPAA, which guarantees the ethical and regulatory sustainability of the system.

Although the paper is theoretical and conceptual in nature, without clinical validation at this stage, it lays a solid foundation for future prototyping, testing, and evaluation in a real healthcare environment. The identified challenges - such as the limited NLP resources in the Macedonian language, the dependence on commercial API services, and the need for local medical corpora - indicate the main directions for further development. They also highlight the need for national initiatives to create open health data infrastructures that will encourage the application of artificial intelligence in the domestic healthcare system.

At its core, this research represents an initial step towards creating ethically responsible, explainable, and interoperable digital medical agents that combine technical feasibility with clinical and ethical relevance. Such systems have the potential to transform initial health assessment, support preventive and personalized medicine, and create more accessible, transparent, and equitable healthcare, especially in resource-limited settings.

References

[1] Siciliani, L., Lafortune, G., Fujisawa, R., Vuik, S., & Hewlett, E. (2023). A review of policies to reduce waiting times for health services across OECD countries: Waiting times for health services. Nordic Journal of Health Economics, 11(1), 162-181.

[2] World Health Organization. (2022). Health and care workforce in Europe: time to act. In *Health and care workforce in Europe: time to act*.

[3] Busch, F., Hoffmann, L., Rueger, C., van Dijk, E. H., Kader, R., Ortiz-Prado, E., ... & Bressem, K. K. (2025). Current applications and challenges in large language models for patient care: a systematic review. *Communications Medicine*, *5*(1), 26.

[4] Ding, H., Simmich, J., Vaezipour, A., Andrews, N., & Russell, T. (2024). Evaluation framework for conversational agents with artificial intelligence in health interventions: a systematic scoping review. *Journal of the American Medical Informatics Association*, *31*(3), 746-761.

[5] Porto, B. M. (2024). Improving triage performance in emergency departments using machine learning and natural language processing: a systematic review. *BMC Emergency Medicine*, *24*(1), 219.

[6] Sutaria, S., Devakumar, D., Mallinson, P., Kinra, S., Malak, T. T., & Meczner, A. (2025). Evaluating the use of red flags by online symptom checkers. *BMC Health Services Research*, *25*(1), 1263.

[7] Staehelin, D., Dolata, M., Stöckli, L., & Schwabe, G. (2024). How patient-generated data enhance patient-provider communication in chronic care: field study in design science research. *JMIR Medical Informatics*, *12*(1), e57406.

[8] Bedi, S., Liu, Y., Orr-Ewing, L., Dash, D., Koyejo, S., Callahan, A., ... & Shah, N. H. (2025). Testing and evaluation of health care applications of large language models: a systematic review. *Jama*.

[9] Omiye, J. A., Gui, H., Rezaei, S. J., Zou, J., & Daneshjou, R. (2024). Large language models in medicine: the potentials and pitfalls: a narrative review. *Annals of internal medicine*, *177*(2), 210-220.

[10] Hager, P., Jungmann, F., Holland, R., Bhagat, K., Hubrecht, I., Knauer, M., ... & Rueckert, D. (2024). Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nature medicine*, *30*(9), 2613-2622.

[11] Kopka, M., von Kalckreuth, N., & Feufel, M. A. (2025). Accuracy of online symptom assessment applications, large language models, and laypeople for self–triage decisions. *npj Digital Medicine*, *8*(1), 178.

[12] Knitza, J., Hasanaj, R., Beyer, J., Ganzer, F., Slagman, A., Bolanaki, M., ... & Dechant, K. (2024). Comparison of two symptom checkers (Ada and Symptoma) in the emergency department: randomized, crossover, head-to-head, double-blinded study. *Journal of Medical Internet Research*, *26*, e56514.

[13] World Health Organization. (2024). *Ethics and governance of artificial intelligence for health: large multi-modal models. WHO guidance*. World Health Organization.

[14] Van Kolfschooten, H., & Van Oirschot, J. (2024). The EU artificial intelligence act (2024): implications for healthcare. *Health Policy*, *149*, 105152.

[15] Zahra, F. A., & Kate, R. J. (2024). Obtaining clinical term embeddings from SNOMED CT ontology. *Journal of Biomedical Informatics*, *149*, 104560.

[16] Katsoulakis, E., Wang, Q., Wu, H., Shahriyari, L., Fletcher, R., Liu, J., ... & Xiao, Y. (2024). Digital twins for health: A scoping review. npj Digital Med.

[17] Tuama, M. H. (2025). A Comparative Evaluation of Random Forest and XGBoost Models for Disease Detection Using Medical Indicators. *International Journal of Professional Studies*, *19*.

[18] Nerella, S., Bandyopadhyay, S., Zhang, J., Contreras, M., Siegel, S., Bumin, A., ... & Rashidi, P. (2024). Transformers and large language models in healthcare: A review. *Artificial intelligence in medicine*, *154*, 102900.

[19] Teo, Z. L., Jin, L., Liu, N., Li, S., Miao, D., Zhang, X., ... & Ting, D. S. W. (2024). Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture. *Cell Reports Medicine*, *5*(2).

[20] Pati, S., Kumar, S., Varma, A., Edwards, B., Lu, C., Qu, L., ... & Bakas, S. (2024). Privacy preservation for federated learning in health care. *Patterns*, *5*(7).

[21] Rosenbacke, R., Melhus, Å., McKee, M., & Stuckler, D. (2024). How explainable artificial intelligence can increase or decrease clinicians' trust in AI applications in health care: systematic review. *Jmir Ai*, *3*, e53207.

[22] Sadeghi, Z., Alizadehsani, R., Cifci, M. A., Kausar, S., Rehman, R., Mahanta, P., ... & Pardalos, P. M. (2024). A review of Explainable Artificial Intelligence in healthcare. *Computers and Electrical Engineering*, *118*, 109370.

[23] Singhal, K., Tu, T., Gottweis, J., Sayres, R., Wulczyn, E., Amin, M., ... & Natarajan, V. (2025). Toward expert-level medical question answering with large language models. *Nature Medicine*, *31*(3), 943-950.

[24] Jin, Q., Leaman, R., & Lu, Z. (2024). *PubMed and beyond: biomedical literature search in the age of artificial intelligence. eBioMedicine. 2024; 100: 104988*.

[25] Johnson, A. E., Bulgarelli, L., Shen, L., Gayles, A., Shammout, A., Horng, S., ... & Mark, R. G. (2023). MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data*, *10*(1), 1.

[26] Harada, Y., Sakamoto, T., Sugimoto, S., & Shimizu, T. (2024). Longitudinal changes in diagnostic accuracy of a differential diagnosis list developed by an AI-Based symptom checker: retrospective observational study. *JMIR Formative Research*, *8*, e53985.

[27] Fung, K. W., Xu, J., Brear, H., Lane, A., Lau, M., Wong, A., & D'Havé, A. (2024). Promoting interoperability between SNOMED CT and ICD-11: lessons learned from the pilot project mapping between SNOMED CT and the ICD-11 Foundation. *Journal of the American Medical Informatics Association*, *31*(8), 1631-1637.

[28] Vuokko, R., Vakkuri, A., & Palojoki, S. (2023). Systematized nomenclature of medicine–clinical terminology (SNOMED CT) clinical use cases in the context of electronic health record systems: systematic literature review. *JMIR medical informatics*, *11*, e43750.

[29] Wang, F., & Beecy, A. (2025). Implementing AI models in clinical workflows: a roadmap. *BMJ Evidence-Based Medicine*, *30*(5), 285-287.

[30] Fąferek, J., Cariou, P. L., Hege, I., Mayer, A., Morin, L., Rodriguez-Molina, D., ... & Kononowicz, A. A. (2024). Integrating virtual patients into undergraduate health professions curricula: a framework synthesis of stakeholders' opinions based on a systematic literature review. *BMC Medical Education*, *24*(1), 727.

[31] García-Torres, D., Vicente Ripoll, M. A., Fernández Peris, C., & Mira Solves, J. J. (2024, November). Enhancing clinical reasoning with virtual patients: a hybrid systematic review combining human reviewers and ChatGPT. In *Healthcare* (Vol. 12, No. 22, p. 2241). MDPI.

[32] Cevasco, K. E., Morrison Brown, R. E., Woldeselassie, R., & Kaplan, S. (2024). Patient Engagement with Conversational Agents in Health Applications 2016–2022: A Systematic Review and Meta-Analysis. *Journal of medical systems*, *48*(1), 40.

[33] Faizan Siddiqui, M., Jabeen, S., Alwazzan, A., Vacca, S., Dalal, L., Al-Haddad, B., ... & Kalmatov, R. (2025). Integration of Augmented Reality, Virtual Reality, and Extended Reality in Healthcare and Medical Education: A Glimpse into the Emerging Horizon in LMICs—A Systematic Review. *Journal of Medical Education and Curricular Development*, *12*, 23821205251342315.

[34] Fruchart, M., Quindroit, P., Jacquemont, C., Beuscart, J. B., Calafiore, M., & Lamer, A. (2024). Transforming primary care data into the observational medical outcomes partnership common data model: Development and usability study. *JMIR Medical Informatics*, *12*(1), e49542.

[35] Yim, W. W., Abacha, A. B., Fu, Y., Sun, Z., Xia, F., Yetisgen-Yildiz, M., & Krallinger, M. (2024, June). Overview of the MEDIQA-M3G 2024 Shared Task on Multilingual Multimodal Medical Answer Generation. In *Proceedings of the 6th Clinical Natural Language Processing Workshop* (pp. 581-589).

[36] Nzenwata, I., OO, I., Tai-Ojuolape, E. O., Aderogba, T. A., Durodola, O. F., Kesinro, P. O., ... & Adesuyan, M. A. (2024). Explainable AI: A systematic literature review focusing on healthcare. *en, J. Comput. Sci. Appl*, *12*(1), 10-16.

[37] Moses, M. E., & Gipson Rankin, S. M. (2024). Medical artificial intelligence should do no harm. *Nature Reviews Electrical Engineering*, *1*(5), 280-281.

[38] Stalla-Bourdillon, S. (2025). The state of pseudonymisation in the EU: Where do we stand today?. *Privacy and Data Protection*, *18*(2).

[39] Dolin, R. H., Heale, B. S., Alterovitz, G., Gupta, R., Aronson, J., Boxwala, A., ... & Chamala, S. (2023). Introducing HL7 FHIR Genomics Operations: a developer-friendly approach to genomics-EHR integration. Journal of the American Medical Informatics Association, 30(3), 485-493.

[40] Sadeghi, Z., Alizadehsani, R., Cifci, M. A., Kausar, S., Rehman, R., Mahanta, P., ... & Pardalos, P. M. (2024). A review of Explainable Artificial Intelligence in healthcare. *Computers and Electrical Engineering*, *118*, 109370.

[41] Chen, R. J., Wang, J. J., Williamson, D. F., Chen, T. Y., Lipkova, J., Lu, M. Y., ... & Mahmood, F. (2023). Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nature biomedical engineering*, *7*(6), 719-742.

[42] Fink, M. (2025). Human Oversight under Article 14 of the EU AI Act. *Available at SSRN*.

[43] Ben Abacha, A., Yim, W., Fu, Y., Sun, Z., Xia, F., Yetisgen, M., & Krallinger, M. (2024). Overview of the MEDIQA-M3G 2024 shared tasks on multilingual multimodal medical answer generation.