SICURA: A Symbiotic AI System to Support Pharmacological Management

Raffaele Manna[1,2,*,†], Giulia Speranza[1,2,†], Maria Pia di Buono[1,2,†], Gianmarco Troia[3,†] and Giovanni Vizzini[3,†]

[1]Dahlia srl, Spin-off University of Naples "L'Orientale", Via Duomo 219, Naples, 80139, Italy
[2]University of Naples "L'Orientale", Via Chiatamone 61/62, Naples, 80121, Italy
[3]Qwince Innovation srl, Via Principe di Belmonte, 102, Palermo, 90139, Italy

Abstract

This paper describes the development of an AI-driven decision-making support system designed to assist healthcare professionals in managing pharmacological prescriptions and drug interactions, with a focus on cardiovascular diseases. The system utilizes Tiny Language Models (TLMs) to improve patient safety, treatment effectiveness, and reduce medication errors. The main objective is to symbiotically support the healthcare professionals by leveraging the artificial intelligence system's ability to rapidly and precisely retrieve and extract information.

Keywords

Decision-making Support System, Symbiotic AI, TLM, Drug Interactions

1. Introduction

This article aims to describe the methodological framework underlying the development of a task-oriented decision-making support system based on a Language Model (LM), designed to assist healthcare professionals in the analysis and management of prescriptions and drug interactions, with the goal of improving patient medical care.

The application domain of the case study —namely, cardiovascular implications of cancer therapy (cardiotoxicity and heart failure)— was selected due to its significance and widespread impact. In fact, its prevalence in Italy is considerable, representing one of the leading causes of morbidity and mortality. Indeed, the Italian National Institute of Health (ISS) reports that cardiovascular diseases account for 35% of all deaths in the country. Furthermore, the risk of drug interactions is directly proportional to the number of medications taken by the patient, and the most frequent interactions involve precisely those drugs that are most commonly used, such as medications for cardiovascular diseases.

The decision-making support system is currently under development within the framework of the SICURA project (SymbIotIC AI for drUg inteRactions Assessment), proposed by Dahlia Srl and Qwince Innovation Srl. The project was selected as a winner of the Cascade Call under Spoke 6 "Symbiotic AI", coordinated by the University of Bari "Aldo Moro", within the "Future Artificial Intelligence Research" (FAIR) initiative.

Among the general objectives of the SICURA project is the goal of improving patient safety by reducing the risk of side effects, enhancing treatment effectiveness, and mitigating potential risks associated with errors in the management of drug interactions. The main objective is to symbiotically support the healthcare professionals by leveraging the artificial intelligence system's ability to rapidly and precisely retrieve and extract information, thus reducing search times and improving the therapeutic decision-making process.

Specifically, the project aims to:

Ital-IA 2025: 5th National Conference on Artificial Intelligence, organized by CINI, June 23-24, 2025, Trieste, Italy
*Corresponding author.
†These authors contributed equally.
rmanna@dahliasrl.it (R. Manna); gsperanza@dahliasrl.it (G. Speranza); mpdibuono@unior.it (M. P. di Buono); gianmarco@qwince.com (G. Troia); giovanni.vizzini@qwince.com (G. Vizzini)
0009-0006-6285-8557 (R. Manna); 0000-0002-9249-492X (G. Speranza); 0000-0001-7284-6556 (M. P. di Buono)

- Improve pharmacological management: The symbiotic AI system is intended to support and enhance the pharmacological management of patients with cardiovascular problems, assisting healthcare professionals in identifying more effective and safer therapies to reduce the risk of complications.
- Reduce prescription errors: The system aims to minimize prescription errors through data analysis and artificial intelligence, ensuring that patients receive the correct therapy and the appropriate dosage.
- Optimize therapy: The system should be able to optimize drug therapy for each patient, taking into account individual medical and personal needs, comorbidities, and personal responses to medications.
- Improve quality of life: Optimized pharmacological management combined with effective prevention of cardiovascular complications is expected to improve patients' quality of life by reducing disease burden and enhancing overall well-being.
- Reduce healthcare costs: The system will eventually lead to long-term reductions in healthcare costs, decreasing the need for expensive medical interventions and hospitalizations.

2. Methodology

The initial phase of the SICURA project includes studies and activities focused on modeling the human agent and the virtual patient, aimed at better defining the medical sub-domain for experimentation and the informational needs of healthcare professionals, as well as the parameters of the medical profile and the therapeutic needs of patients.

Additionally, the prototyping of the LM-based decision-making support system and its optimization are planned. Finally, a testing phase is foreseen. However, the testing of the medical prescription support model will not take place in a real clinical environment with actual patients. Instead, it will be conducted using laboratory data based on electronic health records and virtual patients.

2.1. Modeling the Human Agent

In the human agent modeling phase —which represents the decision-making support system in charge of supporting the healthcare professional in the pharmacological management— an analysis is conducted on the requirements, needs, and activities performed by healthcare workers during their daily drug management. This analysis is carried out through interview techniques and workflow analysis. Specifically, modeling a human agent that represents a healthcare professional managing a patient's medication requires an approach that captures the complexity of decision-making processes. Furthermore, it has to manage pharmaceutical domain knowledge and adherence to national and international clinical protocols and guidelines.

For the purposes of this project, a hybrid methodological approach has been chosen by combining neural architectures with symbolic representation models such as ontologies and knowledge graphs. This approach includes the process of knowledge injection of information from domain-specific databases and Italian-language curated linguistic resources, such as those developed by Manna et al. [1], to perform the fine-tuning of the LM for the identified subdomain. Additionally, the model's capacity to manage drug interaction information and identify potential correlations will be preliminarily optimized through targeted training. This training is aimed at teaching the model to appropriately allocate computational reasoning abilities without direct user feedback, using techniques such as Policy Optimization (PO) [2, 3]. This methodology ensures the enhancement of the system's capabilities through an approach that integrates Retrieval-Augmented Generation (RAG), Chain-of-Thought (CoT) reasoning, and function-calling strategies. The adoption of this unified framework allows for the integration of reasoning capabilities by leveraging the symbolic representation of knowledge and domain-specific expertise required to support the human healthcare professional's decisions.

2.2. Modeling of the Virtual Patient

The modeling of virtual patients aims to establish a structured framework for representing clinical profiles by leveraging synthetic data derived from real patient records. This approach enables the creation of realistic clinical scenarios, producing virtual patients with coherent medical histories and characteristics closely resembling those of actual individuals [4, 5].

This information includes:

- Age, sex, primary disease, and any associated conditions or comorbidities,
- Allergies (e.g. to food or medication),
- Smoking and alcohol use,
- Relevant medical history (past illnesses, previous surgeries, family history of cardiovascular disease and diabetes), and medication history,
- Clinical presentation at the time of medical evaluation (general condition, sensorium status, significant signs and symptoms),
- Nutritional status, weight and height, BMI,
- Ongoing therapy (drug, dosage, schedule, duration, reason, evaluation of response),
- Laboratory data available in the medical records (renal function, liver function, coagulation parameters).

The types of patients and the subsequent experimentation are quite heterogeneous, while the clinical problem to be addressed is restricted to the sub-domain of cardiovascular system diseases. The prototyping phase focuses on the cardiovascular implications of cancer therapy (cardiotoxicity and heart failure). Representative data of typical clinical and therapeutic situations for patients treated with Bevacizumab and Trastuzumab are therefore being created (based on real patient profiles). Indeed, patients undergoing chemotherapy with these drugs often have complex treatment regimens involving multiple active substances. Subsequently, the interaction of the chemotherapeutic drugs with other medications by the patient will be examined, highlighting any potential risk areas. Following this, synthetic data will be generated to validate the model. During the prototyping phase, prototypical virtual patients will be generated using LM that leverages the information from the previously defined clinical profiles. An ethical approach will be applied, incorporating de-identification and anonymization techniques when necessary [6, 7], in order to prevent personal data leakage from the models.

2.3. System Prototyping

The results of the human agent and virtual patient modeling are used to create a dynamic list of all the requirements, functionalities, modifications, and improvements that need to be developed for the decision support system. The virtual patient model and the human agent model/decision-making support system will be integrated in order to provide personalized outputs tailored to the human healthcare professional queries, based on the needs of the virtual profiles, also leveraging information retrieved from external resources. The creation of the decision-making support system is based on a LM trained on domain-specific knowledge bases, leveraging recent methods such as RAG [8, 9, 10].

For the tuning of the decision support model, the experimentation will primarily focus on Tiny Language Models (TLMs) such as LLaMA, Mistral, as well as Minerva, the first Italian language model. These models offer advantages in terms of energy efficiency and easier integration into other platforms. In this specific case, various TLMs will be considered to simulate the role of patients during medical consultations in order to assess the predictive capabilities of the system. The TLM, interfacing with a drug database, must identify potential interactions between prescribed drugs and other medications or supplements taken by the patient, generating an alert for the healthcare professional.

TLMs offer several specific advantages when applied to a virtual human agent (chatbot or support system) for managing drug prescriptions, particularly in contexts where computational resources are limited or when privacy is a priority. Their reduced hardware requirements —due to their lower complexity— enable faster response times and higher inference speeds. This is crucial in this domain, as it

ensures smooth and responsive interaction, which is essential when patients or healthcare professionals need immediate access to information. Additionally, TLMs are also suitable for scalability. In terms of data privacy and security, TLMs allow for on-premise deployment, since they can easily operate within the infrastructure of healthcare organizations, which may also be not very powerful. They also support federated learning processes, which enable the model to learn from data distributed across different devices or servers without ever centralizing raw data. Moreover, another important aspect is that TLMs can be trained on smaller datasets with fair and satisfactory results.

Finally, the adoption of TLMs is also driven by the need for efficient fine-tuning with limited resources, the necessity of integration with systems managing healthcare profiles, and the importance of making the model's reasoning criteria understandable and interpretable. As such, the quality of the training set emerges as a critical factor for the project's success.

3. Conclusion and Future Work

The symbiotic AI-based decision-making support system represents a potential solution to assist healthcare professionals in pharmacological management, since it aims at improving the quality, safety, and efficiency of healthcare services. This is primarily due to its ability to access different sources of information and its capabilities in continuous learning. The implementation of such an agent requires a multidisciplinary approach involving clinical, technological and ethical expertise. Therefore, a prudent and collaborative approach must be adopted, also with the active involvement of healthcare professionals and patients. Considering that the project includes a validation phase, further steps will be taken into account to more precisely deepen the phases of agent modeling, in order to refine clinical validation procedures. For the testing phase, the accuracy of the decision support model's outputs regarding drug interactions mentioned in the virtual profiles will be assessed using reliable external domain resources. Performance measurements will be based on scales that evaluate both the type of interaction (e.g., synergy-antagonism) and the severity of the interaction (ranging from mild to severe). Additionally, the interaction between the virtual patient model and the decision-making support system will be evaluated. This should enable a dynamic, interpretable, and personalized interaction during the simulation of specific clinical scenarios. Furthermore, to address hallucinations and erroneous information, a second phase of evaluation will involve human feedback on the reliability of the AI prototype.

Acknowledgments

SICURA Project is funded by Spoke 6 "Symbiotic AI" Project "Future of Artificial Intelligence Research" (FAIR) PE00000013, CUP H97G22000210007 funded by the Italian National Recovery and Resilience Plan (PNRR), Mission 4, Component 2, Investment 1.3 and by European Union – NextGenerationEU.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT and Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] R. Manna, M. P. di Buono, L. Giordano, et al., Dimmi-drug information mining in italian: A calamita challenge, in: Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-it 2024), 2024.

[2] S. Sane, Hybrid group relative policy optimization: A multi-sample approach to enhancing policy optimization, arXiv preprint arXiv:2502.01652 (2025).

[3] H. Zhao, W. Tang, D. Yao, Policy optimization for continuous reinforcement learning, Advances in Neural Information Processing Systems 36 (2023) 13637–13663.
[4] B. Draghi, Z. Wang, P. Myles, A. Tucker, Identifying and handling data bias within primary healthcare data using synthetic data generators, Heliyon 10 (2024).
[5] E. K. Sichani, A. Smith, K. El Emam, L. Mosquera, et al., Creating high-quality synthetic health data: Framework for model development and validation, JMIR formative research 8 (2024) e53241.
[6] Y. Liu, U. R. Acharya, J. H. Tan, Preserving privacy in healthcare: A systematic review of deep learning approaches for synthetic data generation, Computer Methods and Programs in Biomedicine (2024) 108571.
[7] A. R. Sarkar, Y.-S. Chuang, N. Mohammed, X. Jiang, De-identification is not enough: a comparison between de-identified and synthetic clinical notes, Scientific Reports 14 (2024) 29669.
[8] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, H. Wang, H. Wang, Retrieval-augmented generation for large language models: A survey, arXiv preprint arXiv:2312.10997 2 (2023) 1.
[9] C. S. Ong, N. T. Obey, Y. Zheng, A. Cohan, E. B. Schneider, Surgeryllm: a retrieval-augmented generation large language model framework for surgical decision support and workflow enhancement, npj Digital Medicine 7 (2024) 364.
[10] K. K. Y. Ng, I. Matsuba, P. C. Zhang, Rag in health care: a novel framework for improving communication and decision-making by addressing llm limitations, NEJM AI 2 (2025) AIra2400380.