I’d Like to Argue: Medical Decision Support as Arguments instead of Advice

Bachelor/Master project "Information Systems"

Table of Contents

Prof. Dr. Mario Nadj
University of Duisburg-Essen
Faculty of Computer Science

Chair and Website
Information Systems and AI in Marketing
ai.ris.uni-due.de

Address Campus Essen
Universitätsstraße 9
D-45141 Essen

www.uni-due.de

1: Project Overview

Background: This project combines two vital developments in usage of artificial intelligence (AI) systems for decision making in healthcare. First, it aims to utilize large language models (LLM) as part of a retrieval-augmented generation (RAG) system to obtain trustworthy answers based on actual medical studies. Second, due to rising issues of overreliance on AI systems, the system does not present a single definitive treatment recommendation, but requires that doctors enter their own assessment first, and then receive both supporting and contradicting evidence, based on the RAG system.

Project Goal: Develop and implement RAG system for supporting treatment decisions for breast cancer. The outputs of the RAG system should be presented as supporting and contrasting arguments for a doctor's initial treatment choice.

Highlights: Develop the RAG system from scratch; engage with a highly relevant medical context; freedom to pursue individual solutions; close cooperation with the chair team

2: Project Background

Doctors work in a demanding, time-constrained environment where decision making has severe consequences for their patients. Thus, it is highly important that they receive every possible aid there is, and that they apply it correctly. Today, AI, and especially LLMs show potential to increase healthcare quality by synthesizing the wealth of existing medical knowledge and making it available through simple prompting (Xiong et al., 2024). However, this promise does not come without cost. If doctors receive direct recommendations for how to treat their patients, they may fail to consider relevant alternatives, or may not be able to responsibly communicate such a decision to patients (Liu et al., 2025).

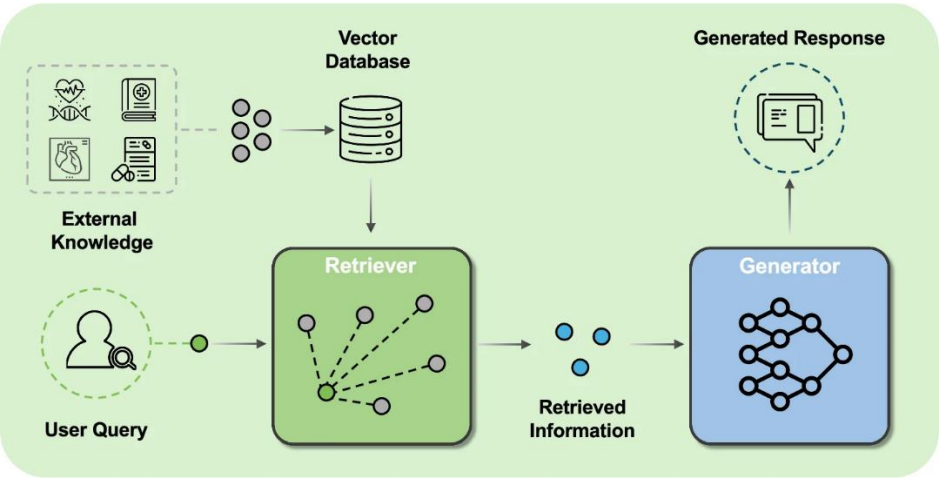

Figure 1: Illustration of RAG framework by Yang et al. (2025).

These risks stem from cognitive biases such as anchoring or automation bias, that could lead to doctors not recognizing when an AI system has made a mistake, thus leading to unsafe or inappropriate treatments. Even if an AI system on average outperforms human doctors, it may not capture all possible variables that a human doctor could perceive in direct patient interaction. A previous project has explored approaches to mitigating cognitive bias in medical decision-making. It has shown that providing both supporting and contradicting or cautioning arguments once the doctor has entered an initial assessment may help to keep treatment decisions safe and under human responsibility, while still profiting from modern LLMs' ability to collect and present medical knowledge. RAG could be used to implement such a system and provide hallucination-free sources (Lewis et al., 2020).

3: Project Objective

The objective of this specific Bachelor/Master project is to develop a RAG system to support medical treatment decision making. It will focus on breast cancer treatments in order to establish a manageable scope.

The main task will be to build the RAG system and (1) establish a working knowledge source, e.g., through a vector database, (2) implement a working retrieval component, and (3) use it to generate both supporting and contradicting arguments for a doctor's initial treatment suggestion. These arguments should be integrated into a user interface to demonstrate the results to interview partners for evaluation.

The objectives of this specific project are the following:

- **Building** a database of relevant oncological publications on breast cancer treatment
- **Implementing** the RAG system and testing different configurations
- **Implementing** a user interface for presenting the generated answers
- **Evaluating** the system against clinical guidelines and through interviews with medical students and expert oncologists

Depending on group size, adjustments to the project scope can be made. This will be announced at the kickoff and in subsequent milestone meetings.

Each finished artifact has to be properly documented to facilitate understanding and further development by students in follow-up projects.

4: Project Registration and further Information

All Bachelor- and Master students interested in this IS-project may apply via our online application form. The link to the application will be **available here** (Link hinterlegt) as of the start of the application period on **Monday, 28.09.2026.** The deadline for the application is **Sunday, 11.10.2026, 23:59.**

For successful completion of the project, Master students receive 12 ECTS (corresponding to 360 working hours) and Bachelor students receive 6 ECTS (corresponding to 180 working hours)

- Master students will invest 18 hours per week into the project.
- Bachelor students will invest 9 hours per week into the project.

In addition to project-specific deliverables during the process, we also require some fixed deliverables from each group. On the 04.12.2026, the project team needs to submit a one-page management summary on the current state of the project. A template will be provided by the Chair. As final deliverables, the project team has to submit four items by 28.02.2027. These are (1) a 5-minute video presenting the results (video for orientation), (2) a scientific poster summarizing the results (whether the poster will be considered for presentation will be determined by our Chair, and a final centralized vote among the Information Systems Chairs), and (3) a project report along with the (4) project code. The project report will contain the final results on 20 to 50 pages (Bachelor) or 30 to 60 pages (Master). The exact number of pages depends on group size and the number of bachelor and master students and will be communicated at kickoff. A template is provided on the WIPRO page. The (4) code needs to be accessible as Open Source via GitHub with appropriate documentation. We want to put up the best posters of each semester (2) at the Chair's premises to highlight your achievements. For outstanding projects, the final video (1) and a link to your code (4) will be shown on the Chair's page.

5: Project Schedule (Dates are to be announced)

Deliverables for each milestone are marked in blue.

6: Grading

The final grade will be calculated as follows based on the sum of the milestones:

- Milestone 1: 5%
- Milestone 2: 30%
- Milestone 3: 10%
- Milestone 4: 5%
- Milestone 5: 10%
- Milestone 6: 40%

7: Sources

- Xiong, G., Jin, Q., Lu, Z., & Zhang, A. (2024). Benchmarking retrieval-augmented generation for medicine. In Findings of the Association for Computational Linguistics: ACL 2024 (pp. 6233-6251).
- Liu, P., Zhang, J., Chen, S., & Chen, S. (2025). Human-AI teaming in healthcare: 1+ 1> 2?. npj Artificial Intelligence, 1(1), 47.
- Yang, R., Ning, Y., Keppo, E., Liu, M., Hong, C., Bitterman, D. S., Ong, J., Ting, D. & Liu, N. (2025). Retrieval-augmented generation for generative artificial intelligence in health care. Npj health systems, 2(1),
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S. & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. Advances in neural information processing systems, 33, 9459-9474.