CliniGuide - Foundational Framework for Medication Recommendation based on Reasoning

Vidushraj Selvaraj

A dissertation submitted in partial fulfilment of the requirements for

Bachelor of Engineer (Honours) degree in Software Engineering

School of Computing

Informatics Institute of Technology, Sri Lanka

in collaboration with

University of Westminster, UK

2025

ABSTRACT

The healthcare industry increasingly demands precise, context-aware medication recommendations to support both clinical decision-making and patient self-management. However, many existing artificial intelligence (AI) systems rely on generic models that often provide insufficient details, such as precise dosages and usage instructions, which can undermine trust among healthcare professionals and patients alike. CliniGuide is developed to address these challenges by leveraging advanced large language models (LLMs) combined with knowledge distillation (KD) and Retrieval-Augmented Generation (RAG) techniques. This innovative approach aims to produce explicit, personalized medication suggestions and to enhance the system's overall transparency through chain-of-thought prompt engineering, thereby bridging the gap between raw AI outputs and clinical best practices.

To solve the identified issues, the project implements a two-phase framework. In the first phase, a teacher LLM generates distilled data, which is then used to fine-tune a smaller, more efficient student model. This knowledge distillation process ensures that critical medication information is preserved while reducing computational overhead. In the second phase, RAG is applied: the system retrieves relevant, domain-specific information from a vector store based on user input, and then combines this knowledge with the student LLM's reasoning through chain-of-thought prompt engineering. This integration not only improves the specificity of the medication recommendations (including brand and dosage details) but also ensures that the decision-making process is transparent, interpretable, and aligned with regional and clinical requirements.

Preliminary evaluations of CliniGuide demonstrate that the proposed framework enhances both the accuracy and clarity of medication recommendations. Quantitative metrics, including BLEU (0.0196), ROUGE-1 (0.2674), ROUGE-2 (0.0519), and ROUGE-L (0.1348), reflect that while the model effectively captures key concepts and maintains relevance, there remains potential for improvement in generating more precise textual outputs. Overall, the evaluation indicates that the integration of KD, RAG, and chain-of-thought reasoning brings CliniGuide closer to the desired goal of delivering trustworthy, contextually grounded, and actionable healthcare insights—a significant step toward realizing AI-driven, patient-specific medication recommendation systems.