

LIMICS at ArchEHR-QA 2025

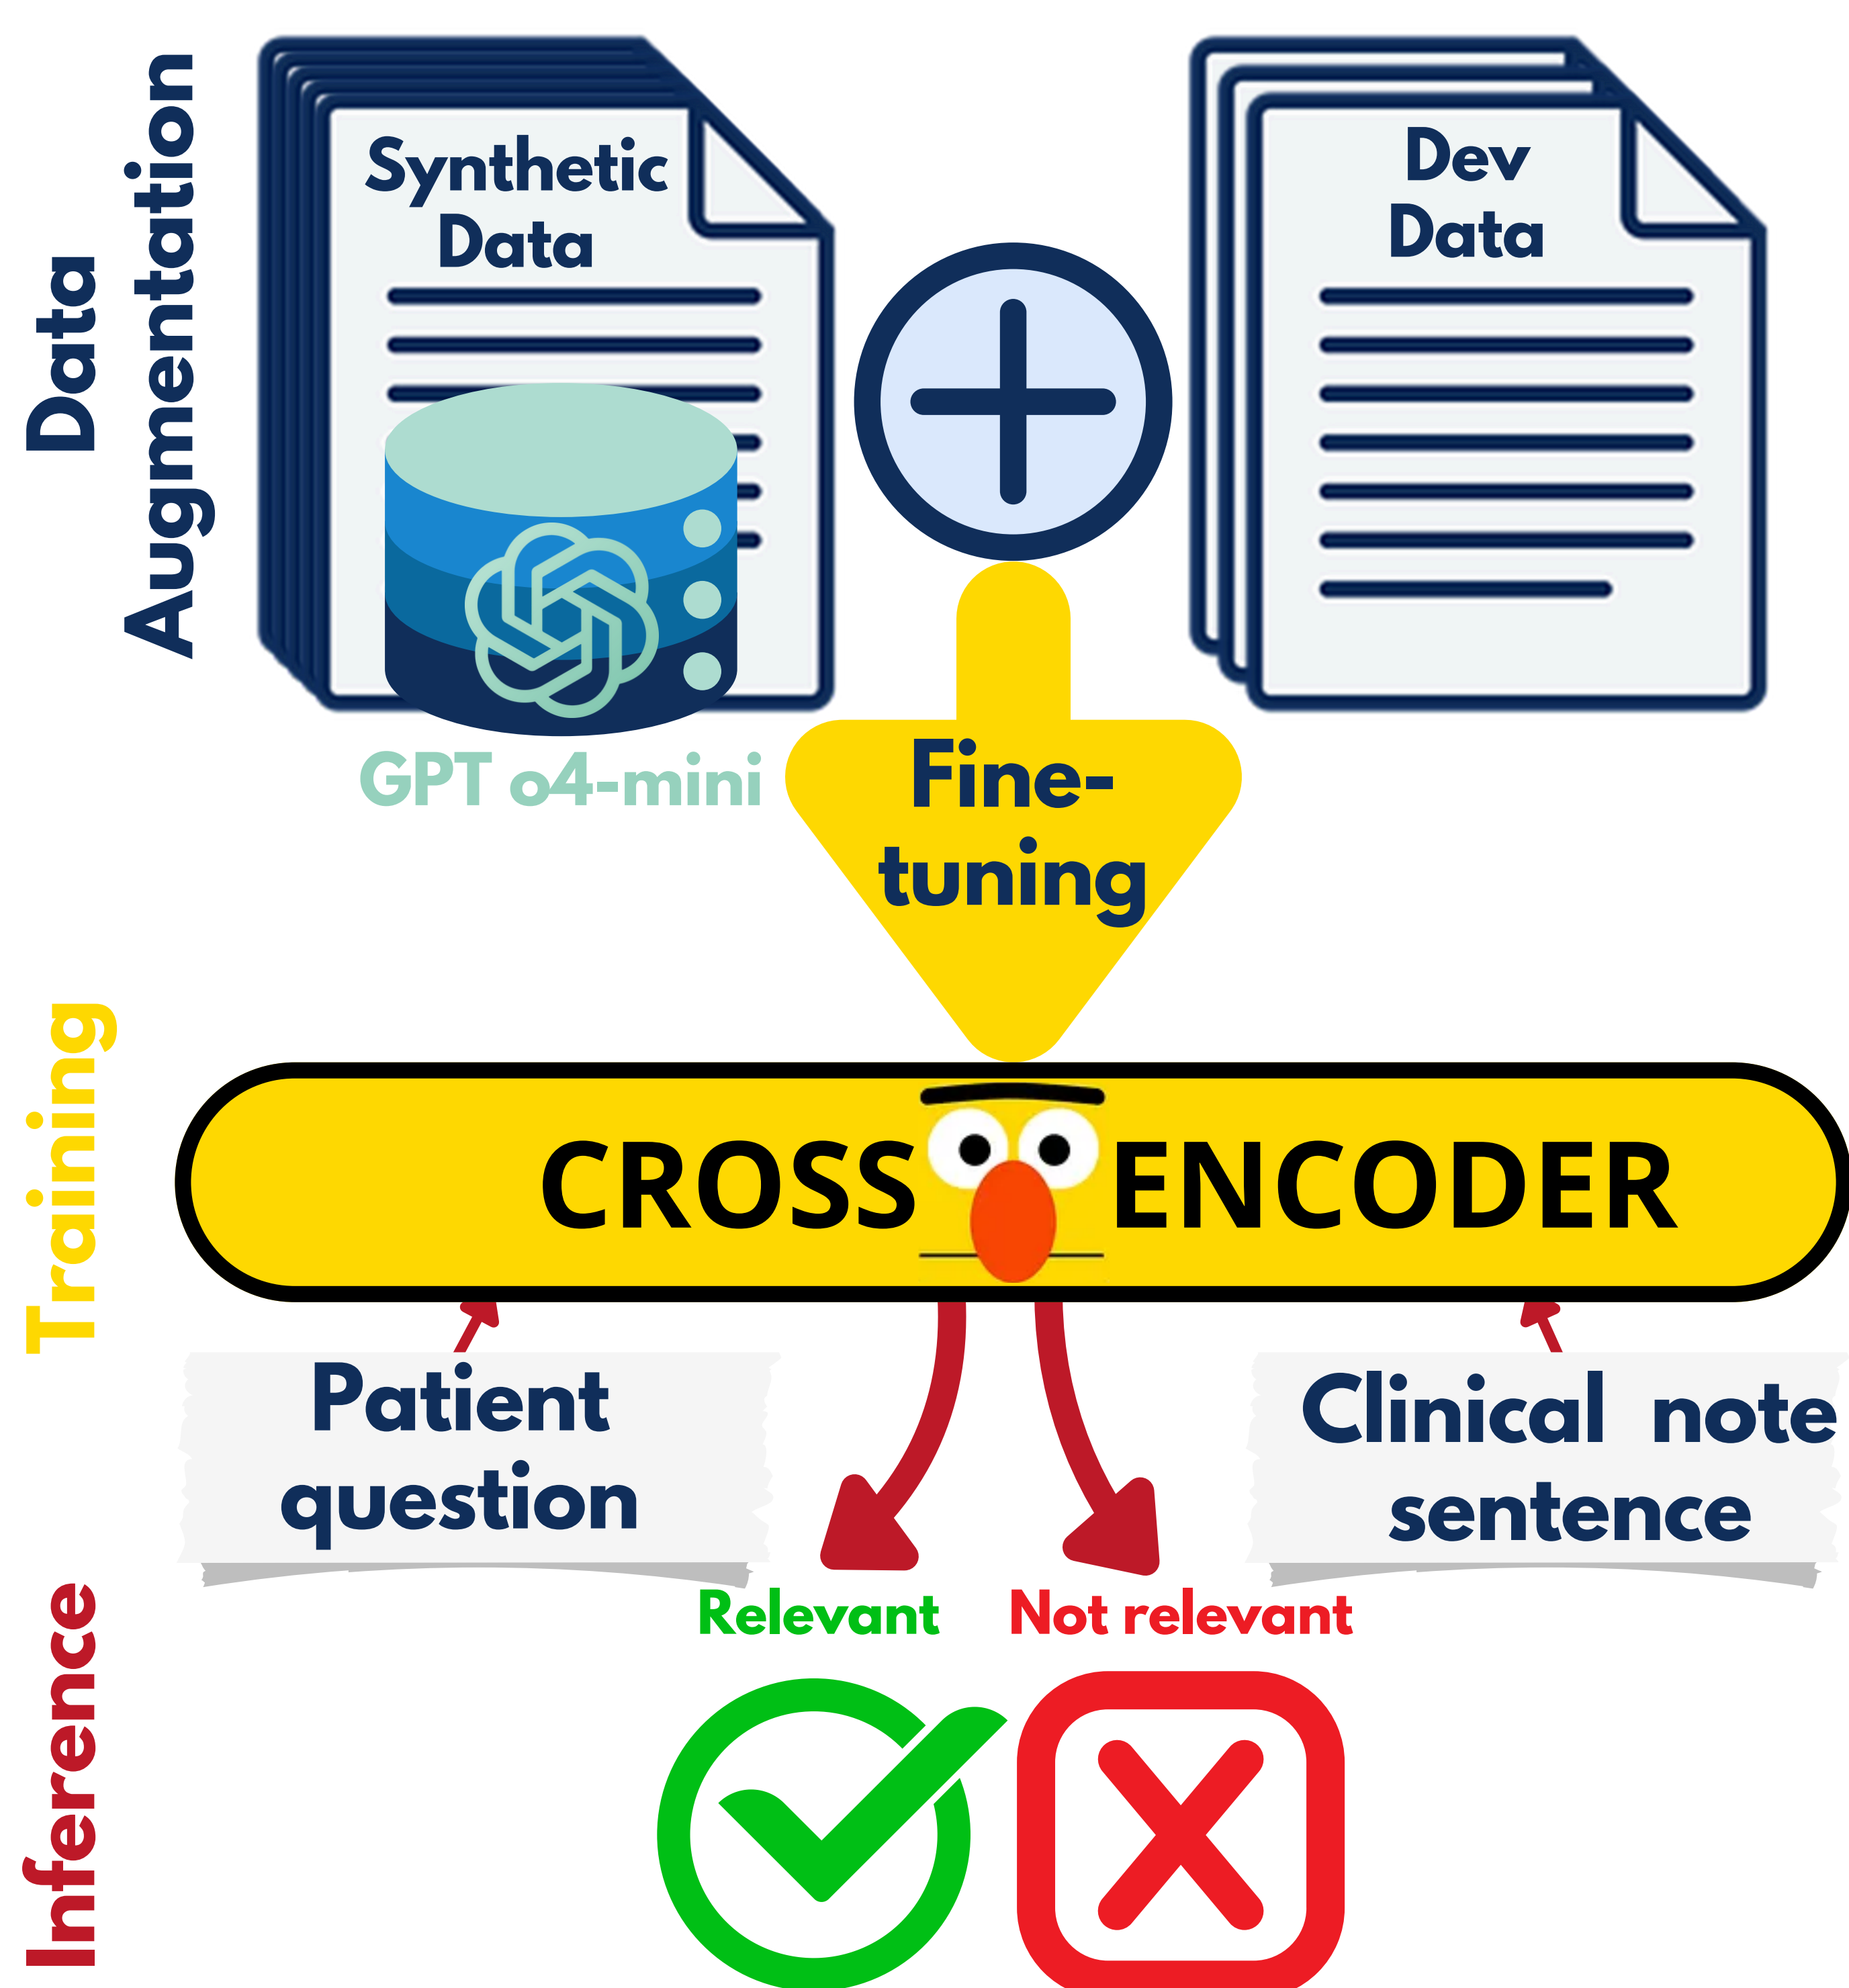
Prompting LLMs Beats Fine-Tuned Embeddings



Background: Clinicians are increasingly overwhelmed by the volume of patient messages received through online portals. The ArchEHR-QA 2025 shared task addresses this challenge by aiming to automatically generate **evidence-grounded responses to patient questions** using their electronic health records.

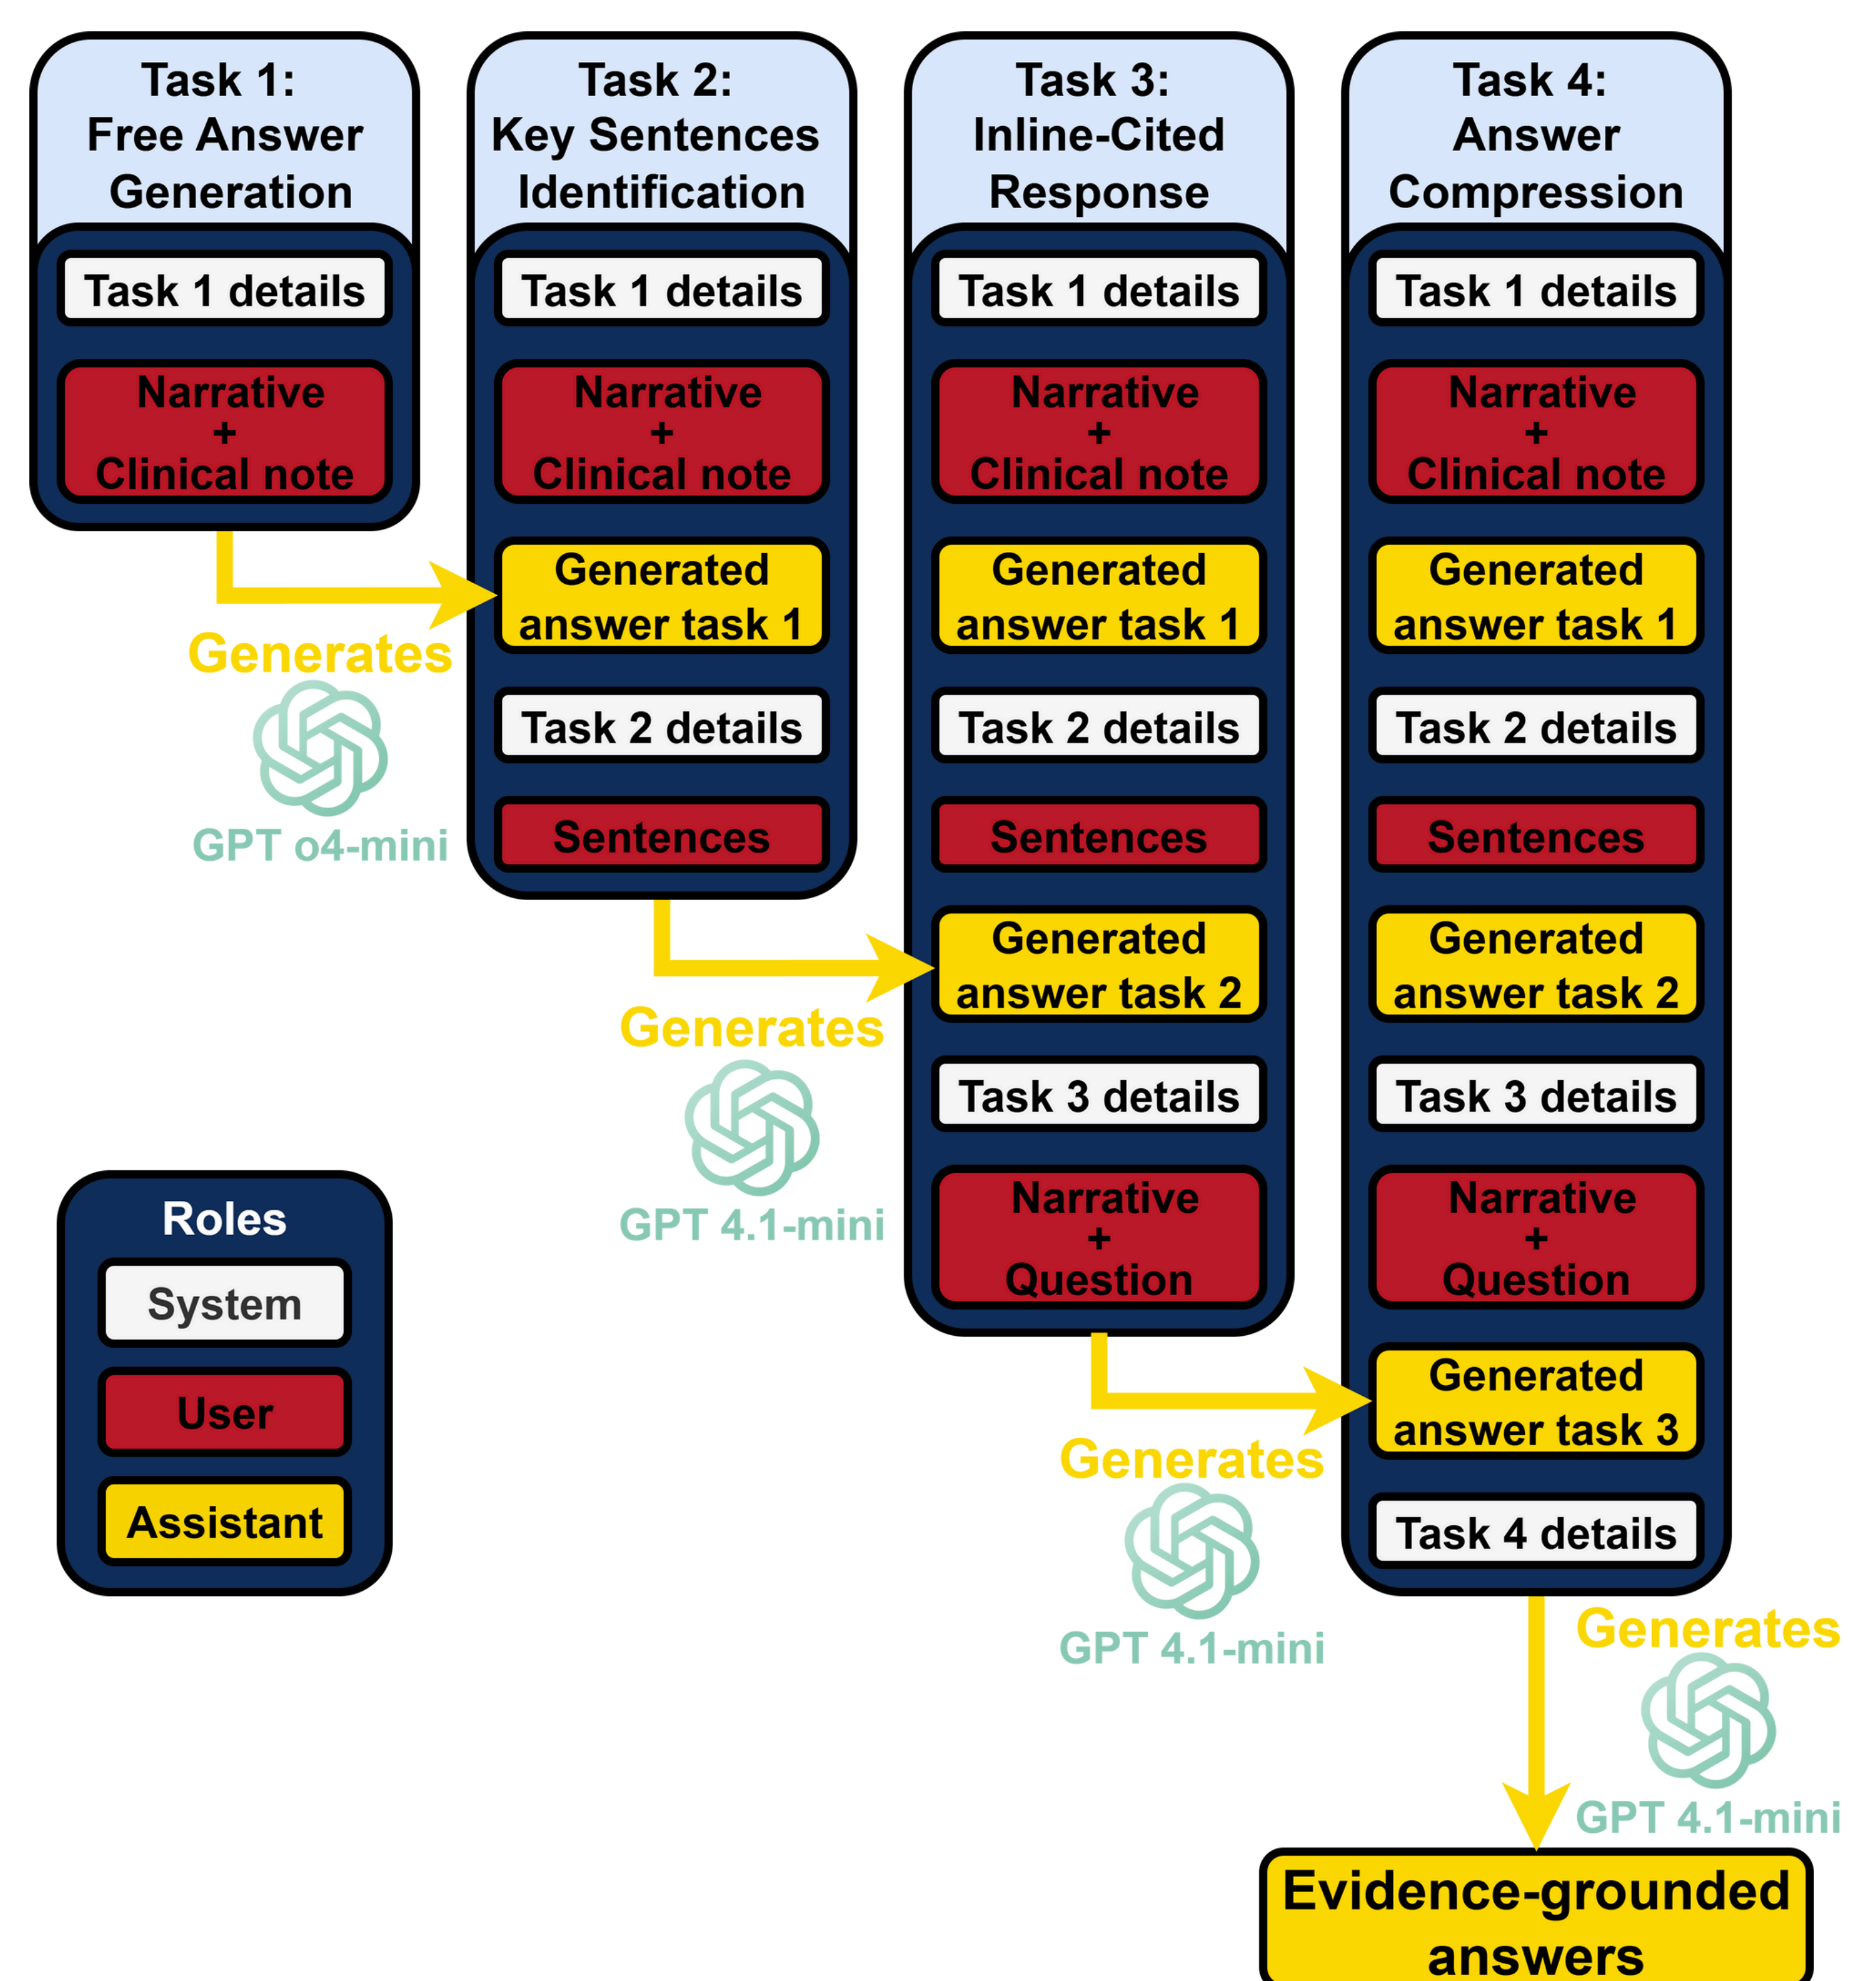
Method 1: Fine-Tuned Cross-Encoder.

We used a BERT-based cross-encoder model **fine-tuned** on the development dataset, **supplemented with synthetic data** from three external corpora using GPT-4o-mini. Its task is to identify which sentences in a clinical note are relevant to a given patient question.



Method 2: Prompt Chaining with LLM.

This method uses large language models in a structured **prompt chaining approach**. The task is broken down, guiding the model to first select relevant information from the clinical notes, and then generate evidence-grounded answers based on that information.



Results: Comparison of F1-scores for both methods in identifying sentences relevant to the question.



Authors: Adam Remaki, Armand Violle, Vikram Natraj, Étienne Guével, Akram Redjda

