INSTITUTIONEN FÖR DATAVETENSKAP | LUNDS TEKNISKA HÖGSKOLA | PRESENTERAD 2016-04-28

EXAMENSARBETE Transforming Project Documentation with AI STUDENTER Edoardo Vaira, Georgios Chavales HANDLEDARE Pierre Nugues (LTH) EXAMINATOR Jacek Malek (LTH)

## Transforming Project Documentation: An AI Chatbot That Talks Your Code

## POPULÄRVETENSKAPLIG SAMMANFATTNING Edoardo Vaira, Georgios Chavales

You are a software developer trying to find information. Imagine instantly finding the exact piece you need buried deep in hundreds of project files, without manually wading through repos, wikis, READMEs and code comments. We've built a smart assistant that automatically ingests all your GitHub repositories and answers natural-language questions about your project, even generating fresh documentation on demand.

Modern software teams often waste precious time hunting for buried details across wikis, README files and code comments. Protocols forgotten in an old PDF or legacy decisions hidden in a comment can stall development, introduce costly bugs and erode institutional memory. These problems are especially urgent in Agile and DevOps-driven environments, where rapid iteration and reliable knowledge access directly impact feature delivery, team productivity and software quality.

To address these needs, we designed a retrievalaugmented generation chatbot that automatically ingests every file in a GitHub repository (Markdown, PDFs, source code and more ), splits them into context-rich segments and indexes them with a hybrid search engine combining deep semantic embeddings and keyword matching. A neural reranker filters out irrelevant passages, and a generative model synthesizes concise, accurate answers strictly grounded in the retrieved content. This interface lets developers pose questions in plain language, like "How do I configure the payment gateway?" or "What's the API endpoint for user authentication?", and receive instant, context-aware responses without ever leaving their chat window.

benchmarks, scaling In controlled our Qwen-2.5-Instruct models from 3 billion to 14 billion parameters increased answer faithfulness on the MSMARCO dataset from 86.2% to 89.8% (a 3.6pp gain) and improved contextual precision by over 2 percentage points. Enabling hypothetical document embeddings further boosted contextual recall by nearly one point. When evaluated on the TechQA dataset with GPT-4o-mini, our pipeline achieved an 86.6% faithfulness score and matching semantic similarity, outperforming larger baselines across key metrics. In qualitative user studies, developers praised the assistant's ability to trace configurations across multiple files and generate coherent documentation from entire chat sessions.

By converting static repositories into living, interactive knowledge bases, our collaborative solution slashes search time from minutes to seconds, preserves institutional memory and prevents costly misunderstandings, empowering teams to focus on innovation rather than information retrieval.