

Text2SQL with „small” Large Language Models

This article conducts a comprehensive comparison of Mistral7B and SQLCoder 2.0, two „small” large language models, in the context of the Text2SQL task using the „Spider” dataset. Despite its modest scale, Mistral7B achieves a notable 33% accuracy without any query-answer examples in the prompt, showcasing promising prospects for compact large language models in the Text2SQL domain. The portability of Mistral7B, capable of running on mobile devices, suggests novel applications for natural language processing in portable services. Challenges persist in handling complex queries and diverse SQL dialects. The article emphasizes the significance of understanding nuanced insights into Mistral7B’s effectiveness, especially when varying the number of provided examples in the prompt.

The study explores models performance with different prompt configurations, revealing unexpected trends in its effectiveness. The findings prompt a detailed discussion on potential reasons, such as model confusion or interference from provided examples.

Despite the results show promise, it is noteworthy to acknowledge that the observed trends may pose limitations for practical production use cases. While the findings provide valuable insights, further refinement and exploration of small large language models are imperative to fully unlock their potential across diverse applications in the Text2SQL domain.

Index Terms—Text2SQL, Large Language Models, LLM, Mistral7B, SQLCoder 2.0, Natural Language Processing, NLP, Fine-tuning, Database Schema, Query Generation, Spider Dataset, SQL

Primary author: JANCZAREK, Damian

Presenter: JANCZAREK, Damian

Session Classification: Session C (Poster)