JUORBIT-AI:
A Retrieval-Augmented AI Assistant for Orbitronics Literature

Amal Aldarawsheh[1], Ahmed Alia[2], and Yuriy Mokrousov[1,3]

[1]Peter Grünberg Institute (PGI-1), Forschungszentrum Jülich, Jülich, Germany;
[2] Institute for Advanced Simulation (IAS-7), Forschungszentrum Jülich, Jülich, Germany;
[3] Institute of Physics, Johannes Gutenberg University Mainz, Germany.

Orbitronics, which studies orbital currents and the transport of orbital angular momentum in quantum materials, is an emerging field with strong potential for next-generation electronic and spintronic devices. However, the rapid growth of the orbitronics literature poses a significant challenge for researchers, as efficiently accessing relevant domain-specific knowledge is difficult. General-purpose large language models (LLMs) can assist with scientific question answering, but they often generate hallucinated or weakly grounded responses when applied to specialized fields such as orbitronics. To address this limitation, we developed JUORBIT-AI, a domain-specific question-answering system for orbitronics literature based on retrieval-augmented generation (RAG) (Figure 1). The system combines large language models with retrieval from peer-reviewed publications to generate context-aware, literature-grounded answers. To evaluate the system, we constructed a benchmark of 50 domain-specific questions derived from several orbitronics review papers, with reference answers manually extracted from the literature and validated by domain experts.

We evaluated open- and closed-weight LLMs, including GPT[1], Llama[2], Qwen[3], DeepSeek[4], and Mistral, using correctness and faithfulness as evaluation metrics. Retrieval augmentation substantially improved performance: GPT-5.2 increased from 49% correctness without RAG to 87% with RAG, showing that retrieval grounding reduces hallucinations. Among the open-weight models, Mistral-Small-3.1-Instruct-2503 performed best, achieving 80% correctness and 89% faithfulness. These findings support the development of an open-source, LLM-based assistant for orbitronics, which will be hosted on Blablador at Forschungszentrum Jülich.

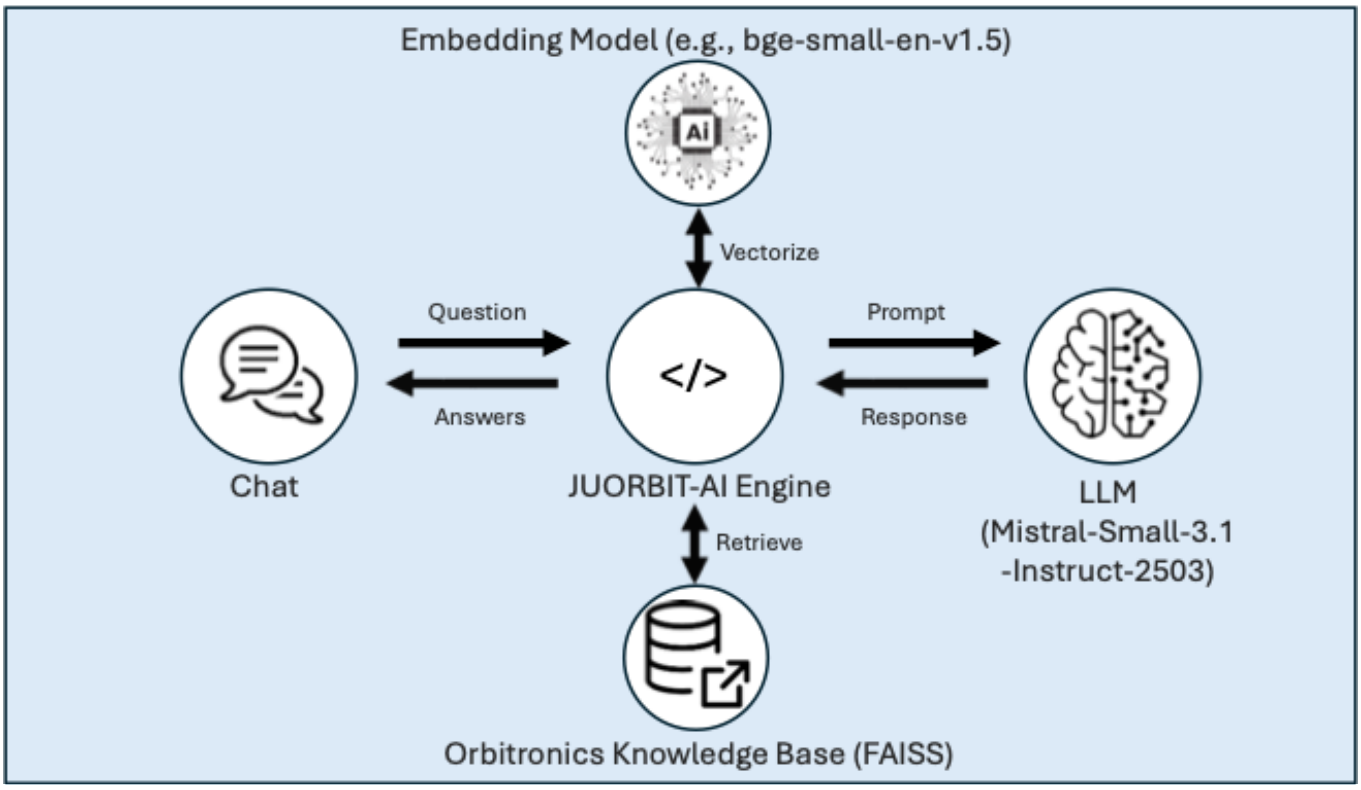

Figure 1: JUORBIT-AI Architecture

[1] OpenAI. *GPT-4 Technical Report*. arXiv:2303.08774, 2023.
[2] Dubey, A., et al. *The Llama 3 Herd of Models*. arXiv:2407.21783, 2024.
[3]Qwen Team. *Qwen2.5 Technical Report.* arXiv:2412.15115, 2024.
[4] DeepSeek-AI et al. *DeepSeek-R1*.arXiv:2501.12948, 2025.