

Graph Based Rag For Legal Lms Research Proposal

Research Proposal: Graph-Based Retrieval-Augmented Generation for Legal LLMs

April 28, 2025

Research Proposal: Graph-Based Retrieval-Augmented Generation for Legal LLMs

Introduction

Large Language Models (LLMs) have demonstrated impressive capabilities but often struggle with domain-specific knowledge and factual accuracy, especially in specialized fields like law. Retrieval-Augmented Generation (RAG) addresses these gaps by grounding LLM outputs in external knowledge bases^{[1][24]}. Traditional RAG methods use flat text stores, which may overlook complex relationships in legal data. Graph-based RAG frameworks (e.g., GraphRAG and PathRAG) instead structure knowledge as interconnected entities, better capturing dependencies across statutes and concepts^{[5][13]}.

GraphRAG, developed by Microsoft researchers, builds knowledge graphs from text by extracting entities and relationships, then uses network analysis and LLM prompting to summarize and retrieve information^{[10][2]}. PathRAG further refines this idea by identifying key relational paths in the graph and pruning redundant information, guiding the LLM with more coherent, low-overlap context^[13]. While powerful, these approaches can be computationally heavy.

LightRAG is a recent, lightweight graph-based RAG system that offers a promising alternative^{[1][2]}. It incorporates dual-level retrieval — combining low-level (entity-level) and high-level (topic-level) retrieval — with incremental graph updates, enabling efficient exploration of large corpora. Early results show LightRAG outperforms GraphRAG on complex datasets, delivering higher retrieval accuracy and faster response times^{[2][1]}.

This proposal plans to leverage LightRAG for a legal knowledge graph built from the *Vade Mecum* (a comprehensive 888-page compendium of Brazilian laws) to enhance LLM reasoning on legal exam questions. We will compare LightRAG to GraphRAG and PathRAG, hypothesizing that LightRAG's simpler, scalable design is better suited for dense legal knowledge. To support multi-step legal reasoning, we will also incorporate the Graph-of-Thoughts framework^[18].

Objectives

- **Construct a Legal Knowledge Graph:** Extract entities and relationships from the *Vade Mecum*.
- **Implement LightRAG on Legal KG:** Use dual-level retrieval and incremental updates^{[1][26]}.
- **Integrate Graph-of-Thoughts Reasoning:** Enable multi-hop inference using Graph-of-Thoughts^[18].
- **Evaluate on Legal Exam Data:** Test on Brazilian OAB and public service exams (2020–2024).
- **Compare with GraphRAG and PathRAG:** Empirical and conceptual analysis.
- **Experiment with State-of-the-Art LLMs:** Test using GPT-4o Mini, DeepSeek-V2, and Gemini 1.5 Pro.

Methodology

1. **Knowledge Graph Construction:** Preprocess *Vade Mecum*, extract legal entities/relationships, form a graph^[26].
2. **LightRAG Implementation:** Apply LightRAG's fast dual-level retrieval and incremental updates^[1].
3. **Graph-of-Thoughts Reasoning:** Iterative multi-hop reasoning using GoT approach^[18].

4. **Baseline Comparisons:** Implement comparable GraphRAG and PathRAG pipelines^{[10][13]}.
5. **Evaluation and Metrics:** Measure accuracy, retrieval efficiency, token usage, and latency.
6. **Experimental Setup:** Query LLMs with exam questions post-training cutoff; ensure no data leakage.

Related Work

- ***Retrieval-Augmented Generation (RAG):*** Widely used for grounding LLMs in external knowledge^{[1][24]}.
- ***Graph-Based RAG:***
 - *GraphRAG*: Builds and traverses knowledge graphs via network analysis^{[5][10][2]}.
 - *LightRAG*: Uses entity/topic key-value retrieval, avoiding heavy summarization^{[1][2]}.
 - *PathRAG*: Extracts minimal relational paths to guide retrieval, improving context quality^[13].
- ***Graph-of-Thoughts (GoT):*** Represents multi-step reasoning as a graph of intermediate thoughts^[18].

This work extends these ideas, proposing that LightRAG's efficient structure particularly suits legal knowledge's complex, interconnected nature.

Expected Contributions

- **LightRAG-based Legal QA System** using the *Vade Mecum*.
- **Comparative Analysis** of LightRAG vs. GraphRAG vs. PathRAG.

- **Demonstration of Multi-step Legal Reasoning** with Graph-of-Thoughts.
- **New Benchmark Dataset** for Brazilian legal exams.
- **Insights for Future RAG Systems** emphasizing scalability and speed.

References

[^1]: Guo, Z., Xia, L., Yu, Y., Ao, T., & Huang, C. (2024). *LightRAG: Simple and Fast Retrieval-Augmented Generation*. arXiv:2410.05779.

[^2]: LightRAG Paper Details. (2024).

[^5]: Peng, B., Zhu, Y., Liu, Y., Bo, X., Shi, H., Hong, C., Zhang, Y., & Tang, S. (2024). *Graph Retrieval-Augmented Generation: A Survey*. arXiv:2408.08921.

[^10]: Edge, D., Trinh, H., Truitt, S., Larson, J., et al. (2024). *Project GraphRAG* (Microsoft Research blog).

[^13]: Chen, B., Guo, Z., Yang, Z., Chen, Y., Chen, J., Liu, Z., Shi, C., & Yang, C. (2025). *PathRAG: Pruning Graph-based Retrieval Augmented Generation with Relational Paths*. arXiv:2502.14902.

[^18]: Besta, M., Blach, N., Kubicek, A., et al. (2024). *Graph of Thoughts: Solving Elaborate Problems with Large Language Models*. AAAI Conference on Artificial Intelligence.

[^24]: Barron, R. C., Eren, M. E., Serafimova, O. M., Matuszek, C., & Alexandrov, B. S. (2025). *Bridging Legal Knowledge and AI: RAG with Vector Stores, KGs, and NMF*. arXiv:2502.20364.

[^26]: Graph indexing methods and LightRAG enhancements.