

Fact Finder - Enhancing Domain Expertise of Large Language Models by Incorporating Knowledge Graphs

Daniel Steinigen¹
Roman Teucher¹
Timm Heine Ruland¹
Max Rudat¹
Nicolas Flores-Herr
Fraunhofer IAIS
Sankt Augustin, Germany

<first-name>.<last-name>@iaais.fraunhofer.de

¹All authors contributed equally to this work

Peter Fischer²
Nikola Milosevic²
Christopher Schymura²
Angelo Ziletti²
Bayer AG

<first-name>.<last-name>@bayer.com

²Shared senior authorship

Abstract

Recent advancements in Large Language Models (LLMs) have showcased their proficiency in answering natural language queries. However, their effectiveness is hindered by limited domain-specific knowledge, raising concerns about the reliability of their responses. We introduce a hybrid system that augments LLMs with domain-specific knowledge graphs (KGs), thereby aiming to enhance factual correctness using a KG-based retrieval approach. We focus on a medical KG to demonstrate our methodology, which includes (1) pre-processing, (2) Cypher query generation, (3) Cypher query processing, (4) KG retrieval, and (5) LLM-enhanced response generation. We evaluate our system on a curated dataset of 69 samples, achieving a precision of 78% in retrieving correct KG nodes. Our findings indicate that the hybrid system surpasses a standalone LLM in accuracy and completeness, as verified by an LLM-as-a-Judge evaluation method. This positions the system as a promising tool for applications that demand factual correctness and completeness, such as target identification — a critical process in pinpointing biological entities for disease treatment or crop enhancement. Moreover, its intuitive search interface and ability to provide accurate responses within seconds make it well-suited for time-sensitive, precision-focused research contexts. We publish the source code together with the dataset and the prompt templates used¹.

1 Introduction

Recently, Large Language Models (LLMs) have enabled sophisticated question-answering systems, revolutionizing the landscape of natural language processing (OpenAI, 2023; Jiang et al., 2023a;

¹<https://github.com/chrschy/fact-finder>

Bubeck et al., 2023; Park et al., 2023; Touvron et al., 2023; Katz et al., 2023). These advanced models, with their ability to understand and generate human-like text, have shown great potential in various domains, including life sciences (Nori et al., 2023; Waisberg et al., 2023; Bašaragin et al., 2024). However, LLMs are limited by the timeframe of their training data and can produce incorrect statements, known as hallucinations (Ji et al., 2023), or incomplete answers by providing only a few relevant entities while missing others not included in their internal knowledge.

In domains such as life sciences, obtaining answers with current and factual information is paramount for many use cases (Malaviya et al., 2023; Ljajić et al., 2024). Factually correct AI-generated reviews can aid researchers in information retrieval and hypothesis building. For example, target identification requires up-to-date knowledge of the latest literature. Target identification involves pinpointing a biological entity, such as a gene or protein, that can be manipulated to achieve a desired effect, like treating a disease in humans (pharmaceuticals) or improving crop resilience in plants (crop sciences). Similarly, designing effective field or clinical trials requires considering up-to-date information. For field trials in crop sciences, this may include environmental and climatic conditions, market developments, and regulatory requirements. For clinical trials and medical writing, relevant information includes details about the drug under development, market conditions, planned sites, and regulatory requirements. Current and timely information is crucial for competitive intelligence, including insights on competitor products, disease epidemiology, and market size. AI-based solutions can assist in identifying concurrent disease occurrences and help researchers develop hypotheses

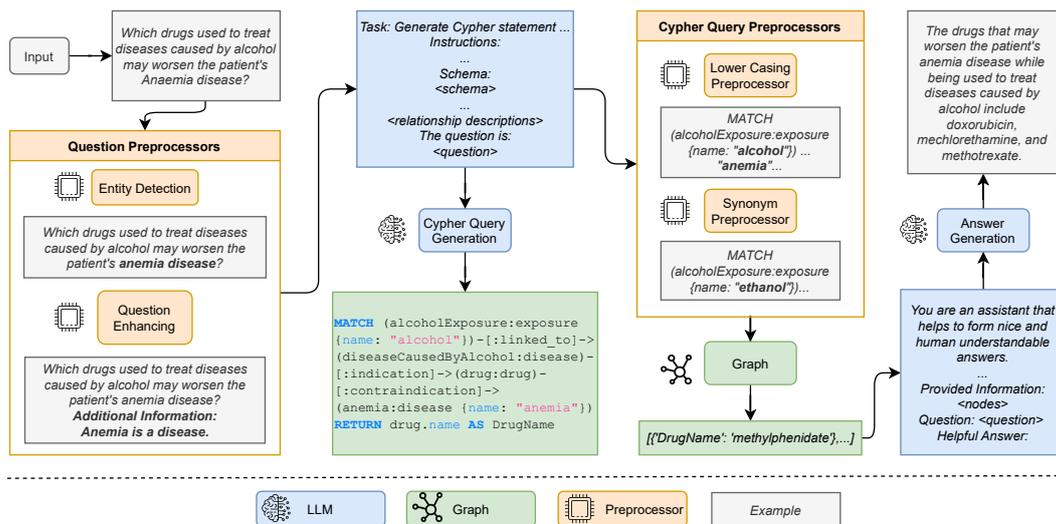


Figure 1: Overview of the FactFinder pipeline using large language models and knowledge graphs to answer scientific questions.

using real-world data.

Knowledge Graphs (KGs) represent a promising strategy for improving factual correctness in LLMs (Jiang et al., 2023b; Baek et al., 2023; Sen et al., 2023), including the life science domain (Feng et al., 2024). By organizing entities such as drugs, diseases, and genes, along with their relationships, into a structured network, KGs provide useful additional context for LLMs for precise and relevant information retrieval (Chandak et al., 2023; Milošević and Thielemann, 2023; Badenes-Olmedo and Corcho, 2023). This organized data framework allows LLMs to produce more factually accurate and comprehensive responses (Pan et al., 2024). In addition, KGs enable systems to leverage current and comprehensive information, including recent data not available during the LLMs’ training phase. From an organizational standpoint, integrating KGs enhances top-tier LLMs with proprietary or specialized knowledge. This integration facilitates the inclusion of unique organizational data sources, such as historical and ongoing lab experiments or licensed datasets.

In this paper, we present FactFinder - a hybrid question answering (QA) system - which leverages both KG and LLM to provide answers to scientific questions. We use the Neo4j graph database² to provide the KG and Cypher³ as query language, and build our pipeline borrowing components from

Langchain⁴. Fig. 1 depicts the system’s architecture, which is structured as a pipeline with several subcomponents. Our main contributions are:

- We provide an easy-to-use system that answers scientific questions combining LLMs and KGs, with automatic subgraph visualizations that enhance interpretability by showing graph-native evidence for each answer.
- We release a dataset of manually annotated text-to-Cypher query pairs, which could serve as benchmark for validating text-to-Cypher conversion system.
- We present a methodology showing that current state-of-the-art LLMs are able to generate satisfactory Cypher queries for the life science domain.
- We share our dataset, source code and prompt templates⁵.

2 Data

Knowledge graph. We use PrimeKG (Chandak et al., 2023) as our source of fact-based background knowledge. PrimeKG integrates 20 high-quality resources to describe 17,080 diseases with 4,050,249 relationships, including over 100,000 nodes and 29 types of edges that densely connect disease nodes with drugs, genes, exposures, and phenotypes. We preprocess the graph data by mapping names to

²<https://neo4j.com/>

³<https://opencypher.org/>

⁴<https://www.langchain.com/>

⁵<https://github.com/chrschy/fact-finder>

their preferred terms, as described in Section 3.1, and converting all entries to lowercase.

Text-to-Cypher dataset. We manually generated a ground-truth dataset containing 69 text-to-Cypher query pairs specifically designed for medical questions. These queries are complex, often involving multiple hops in the graph, aggregation, and boolean question structures. They require deep knowledge of Cypher and the KG. Each entry includes a natural language question, the corresponding Cypher query, the expected answer, and relevant nodes and relationships. This dataset provides a benchmark to evaluate the ability of text-to-Cypher systems to interpret and execute complex queries. While specialized to the PrimeKG graph, the dataset leverages PrimeKG’s extensive applicability, making it a valuable resource for various medical information retrieval tasks. Examples from the dataset include simpler questions like *Which drugs have pterygium as a side effect?* and more complex ones such as *Which medications have more off-label uses than approved indications?* and *Which diseases have only treatments that have no side effects at all?*

3 System Description

3.1 Cypher Query Generation

Generating code to query structured databases from natural language inquiries used to be a complex process, involving steps such as entity and relation extraction, entity and relation linking, query type classification, template-based or compositional query generation (Srivastava et al., 2021; Chakraborty et al., 2021). Retrieval-based approaches like Graph-RAG aim to simplify this by partitioning the graph into communities of nodes and edges, which are then retrieved and summarized using a LLM to generate an answer (Edge et al., 2024). However, Graph-RAG struggles with complex queries that span multiple graph communities and require exact graph operations.

With the advent of LLMs, however, QA systems can now understand domain-specific questions and generate valid queries directly, allowing for more flexible approaches. Much of the research has been centered on text-to-SQL generation, where LLMs have demonstrated considerable effectiveness (Gao et al., 2023; Chang and Fosler-Lussier, 2023), including in the medical domain (Ziletti and D’Ambrosi, 2024), and have often performed better than specialized models (Pourreza and Rafiei,

2023). Conversely, the area of text-to-Cypher query generation remains relatively under-explored, with prior research primarily focused on sequence-to-sequence models (Zhao et al., 2024; Guo et al., 2022). Only recently has the application of LLMs to this task begun to emerge (Feng et al., 2023). To bridge this gap, our work evaluates the capabilities of LLMs to produce robust Cypher queries for scientific QA in the medical domain (see Sec. 4.1).

We prompt LLMs with questions and graph schemas, including node and relationship types and their properties, to generate Cypher queries. When a graph relation is not self-explanatory, we add its natural language description to the prompt. For instance, for the relation *ppi*, we add *"Temporary, non-covalent binding between protein molecules. Protein-protein interactions occur..."*. During instruction prompting, we also identify questions that cannot be answered by the given graph schema. In such cases, the LLM returns the string `SCHEMA_ERROR` along with a brief explanation of why it could not generate an answer. FactFinder detects this marker using a regex and returns the explanation to the user.

We include an entity extraction model to align entity names in questions with those in the KG, based on Linnaeus (Gerner et al., 2010) and developed set of vocabularies for entity types. This step ensures consistency by replacing detected entities with their preferred KG terms (e.g., *alcohol* to *ethanol*) and generating sentences linking each entity to its category (see Fig. 1 left), reducing the LLM’s reliance on domain-specific knowledge.

3.2 Query Pre-Processors

Before querying the graph with the generated Cypher query, we preprocess it to increase the system’s robustness. This leverages the structured natural language understanding provided by the translation of questions to Cypher queries. Various regular expression-based methods target specific query elements.

Formatting. First, we format the query to improve readability and consistency. This includes adding indentation, line breaks, and ensuring consistent naming conventions, which simplifies the application of regular expressions in subsequent steps.

Lowercasing Property Values. We convert property values in the Cypher query to lowercase, matching the pre-lowered graph properties.

Synonym Selection. We map entities in the Cypher query to preferred terms used in the graph. If no

mapping is available, we use external tools (e.g., `skos:altLabel` queries against Wikidata) to find synonyms and match them to graph terms.

Deprecated Code Handling. We correct deprecated code generated by the LLM, such as replacing the obsolete `SIZE()` keyword with the current `COUNT()` keyword.

Child to Parent Node Mapping. In PrimeKG (Chandak et al., 2023), some node types are connected by parent-child relationships. We replace child nodes with their parent nodes in the Cypher query to ensure completeness.

3.3 Graph Question Answering and Verbalization

The pre-processed Cypher query is executed on the graph, returning a unique set of nodes that may include names, properties, IDs, and other elements. Next, the question and graph results are incorporated into a prompt template and sent to an LLM. The prompt instructs the LLM to rely solely on the graph information to formulate an answer, which is then provided as the final natural language output.

3.4 Explainability through Evidence

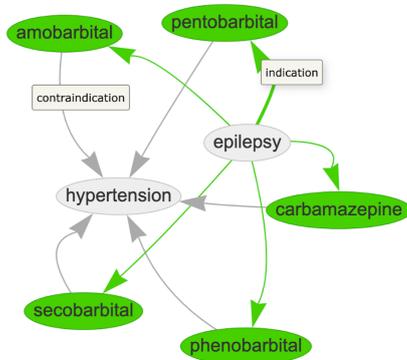


Figure 2: Example of the evidence subgraph for *Which drugs against epilepsy should not be used by patients with hypertension?*

To ensure transparency and explainability, the system provides various forms of evidence alongside the natural language answer. These include intermediate results, such as Cypher generation prompts and graph results, as well as explicitly created information like the underlying subgraph for a query. These evidences are displayed in the frontend to track the system’s behavior.

Cypher Query Evidence. The primary evidence is the Cypher query generated by the model, demonstrating how the given question maps to the graph

structure. This enables expert users to evaluate the model’s understanding of the question and the quality of the generated query.

Graph Response. The system also provides the actual response from the graph, consisting of the nodes and relationships returned by the executed query.

Subgraph Visualization. To enhance interpretability, we provide a subgraph as part of the evidence. This subgraph visually displays (via Pyvis⁶) the relevant nodes and edges, illustrating the subset of the main graph that contributed to the specific answer, as shown in Fig. 2.

Sub-Graph Generator. A notable challenge is that the original Cypher query usually returns only the nodes required to answer the question, omitting the edges that connect the question’s entity to the answer nodes. To address this, we generate a new Cypher query that fetches both the answer nodes and the connecting edges. This is achieved by submitting the original Cypher query to a LLM and instructing it to return all nodes and edges present in the query.

For example, if our original Cypher query is:

```
MATCH (g:gene_or_protein name:"pink1")-
[:associated_with]->(d:disease)
RETURN d.id AS ID, d.name AS Name
```

our subgraph Cypher query is:

```
MATCH (g:gene_or_protein name:"pink1")-
[a:associated_with]->(d:disease)
RETURN g, d, a
```

Note that this last subgraph Cypher query returns all relevant nodes and edges.

3.5 User Interface and Example of Usage

Our target audience includes researchers in the life sciences, such as those working in medical research and crop science, who are interested in exploring connections between drugs, genes, proteins, and other biological entities to discover new research directions. To make our pipelines accessible, we developed a graphical user interface using Streamlit⁷. This interface allows users to input questions and view generated answers. Users can select different pipelines, such as LLM-only or those incorporating KGs or documents. For instance, a medical researcher might ask, *Which drugs against epilepsy*

⁶<https://pyvis.readthedocs.io/>

⁷<https://streamlit.io/>

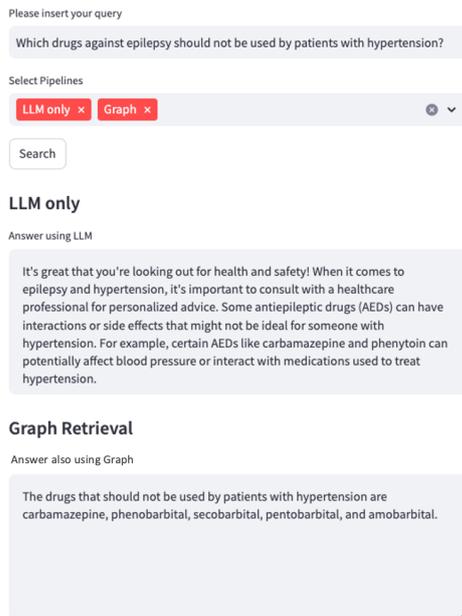


Figure 3: User Interface with question and answers of the standalone LLM and our graph-based hybrid system.

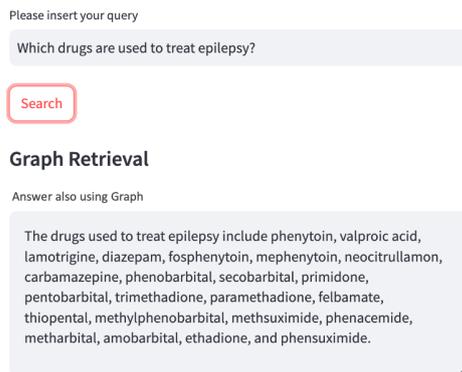


Figure 4: Answer for exploring drugs used to treat epilepsy.

should not be used by patients with hypertension? (Fig. 3). To delve deeper, they could follow up with *Which drugs are used to treat epilepsy?* (Fig. 4). Going further, with the question *Which genes are targeted by amobarbital but not lamotrigine?*, they could take a first step towards understanding the genetic interactions that differentiate drug effects, resulting in the identification of four genes exclusively targeted by amobarbital (Fig. 5). Additionally, users can interactively visualize the relevant subgraph, generated Cypher query, and graph response, enabling them to verify response accuracy and understand the underlying data (Fig. 6). These features (see Section 3.4) enhance transparency and foster trust in the system.

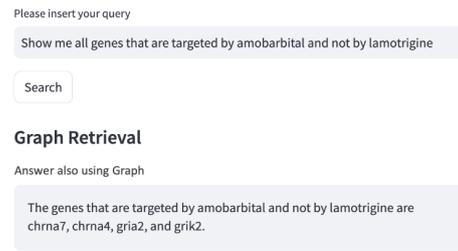


Figure 5: Answer for exploring genes targeted by amobarbital but not lamotrigine.

4 Evaluation

4.1 Graph Retrieval Evaluation

To quantify the graph retrieval step, we assess the returned nodes from graph queries by comparing result nodes from executing the ground truth queries with those from the generated queries. This enables a quantitative evaluation of the text-to-Cypher step. We use the ground truth text-to-Cypher dataset described in Sec. 2 for this evaluation. We compute intersection over union (IoU), precision, and recall for the expected and generated graph result sets, as shown in Table 1.

The results demonstrate robust performance across models, with the highest accuracy exceeding 75%. GPT-4o demonstrates superior performance across all evaluation metrics. Contrary to expectations, Entity Enhancement (EE, Fig.1 left and Sec.3.1) decreased performance in all models except GPT-4-Turbo. This decline stems from two main factors: (1) PrimeKG’s merged gene-protein nodes create ambiguity when EE is applied, leading to incorrect relations, and (2) a terminology mismatch where the synonym endpoint returns formal medical terms while PrimeKG uses non-medical terminology. The results suggest that leveraging LLMs’ internal knowledge is more effective for Cypher query generation than external entity enhancement. The following more detailed analysis focuses only on the most strongly performing models, GPT-4o and GPT-4-turbo.

4.2 Evaluating Correctness and Completeness

To assess the quality of LLM-generated answers, we conduct two evaluations using the LLM-as-a-Judge approach (Zheng et al., 2024): (1) comparing the answers from our KG-LLM-based system to those from an LLM-only system, and (2) evaluating the reliability of LLM verbalization of information provided by the KG. In both cases, correctness



Figure 6: User Interface with the generated Cypher query and the graph response as evidence.

Model	EE	IoU	Precision	Recall
gpt-4-turbo	True	71.3	73.6	71.6
gpt-4-turbo	False	62.7	65.3	64.9
gpt-4o	True	74.9	77.0	77.6
gpt-4o	False	75.2	77.5	77.8
Llama-3.1-70B	True	46.6	46.9	53.3
Llama-3.1-70B	False	52.8	55.4	59.7
Mixtral-8x7B	True	31.4	32.3	33.6
Mixtral-8x7B	False	37.6	39.5	40.1

Table 1: Results for the graph retrieval evaluation (metrics in %). IoU stands for intersection over union and EE for Entity Enhancement, see Section 4.1.

is defined as the inclusion of only facts from the graph nodes, and completeness as the inclusion of all such facts.

Hybrid system vs. LLM-only. We compare the hybrid KG-based system against a standalone LLM. The hybrid system (GPT-4o without entity enhancement, Sec. 4.1) is evaluated to produce more correct (complete) answers in 94.12% (96.08%) of cases, demonstrating its superior performance in providing accurate and complete responses.

LLM verbalization. We evaluate the verbalization of natural language answers from graph results. In this evaluation, 89.13% of answers are deemed correct, and 80.43% complete, indicating high accuracy in verbalization.

4.3 Handling Incorrect Graph Responses

Finally, we evaluate FactFinder’s ability to handle incorrect or incomplete information in graph responses. The system should be able to refuse to answer if the Cypher query generation step produces a wrong query, thus retrieving the correct data from the KG. We test this by disabling Cypher query generation and supplying incorrect Cypher queries for each question, resulting in incorrect graph results.

The results in Table 2 show that both GPT-4-turbo and GPT-4-o can detect irrelevant informa-

(in %)	gpt-4o	gpt-4-turbo
Answer Denied	65/69 (94.2)	63/69 (91.3)
Uncertain Answer	1/69 (1.5)	1/69 (1.5)
Full Answer	3/69 (4.3)	5/69 (7.3)

Table 2: Handling irrelevant information in graph responses.

tion and correctly respond with "I don’t know" in over 90% of cases, demonstrating that the LLMs can reason and understand when the knowledge passed to them is not relevant. This highlights FactFinder’s ability to enhance reliability by leveraging both structured and world knowledge.

Manual analysis revealed that in one case, the LLM expressed uncertainty with the phrase "The provided information does not mention ..." indicating potential inaccuracies. In a few cases, the models provided full answers despite unrelated KG results, especially for count, boolean, and lengthy responses, highlighting detection challenges in these scenarios.

5 Conclusion

This work demonstrates the value of integrating structured, factual knowledge into a user-friendly chat system, providing researchers with a reliable tool for answering scientific questions while minimizing hallucinations. We show that LLMs can generate valid Cypher queries to retrieve relevant data from a KG, informing accurate answers. By grounding every answer in the underlying knowledge graph, FactFinder provides an interpretable and reproducible response pathway that pure LLM systems cannot match. The creation of such a robust system is crucial for enhancing research capabilities. Future work will focus on expanding the evaluation dataset, quantifying system uncertainty, and enabling access to multiple KGs, possibly through agent-based retrieval.

References

- Carlos Badenes-Olmedo and Oscar Corcho. 2023. Lessons learned to enable question answering on knowledge graphs extracted from scientific publications: A case study on the coronavirus literature. *Journal of Biomedical Informatics*, 142:104382.
- Jinheon Baek, Alham Fikri Aji, and Amir Saffari. 2023. Knowledge-augmented language model prompting for zero-shot knowledge graph question answering. In *Proceedings of the 1st Workshop on Natural Language Reasoning and Structured Explanations (NLRSE)*, pages 78–106, Toronto, Canada. Association for Computational Linguistics.
- Bojana Bašaragin, Adela Ljajić, Darija Medvecki, Lorenzo Cassano, Miloš Košprdić, and Nikola Milošević. 2024. How do you know that? teaching generative language models to reference answers to biomedical questions. *BioNLP Workshop 2024*.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Nilesh Chakraborty, Denis Lukovnikov, Gaurav Maheshwari, Priyansh Trivedi, Jens Lehmann, and Asja Fischer. 2021. Introduction to neural network-based question answering over knowledge graphs. *WIREs Data Mining and Knowledge Discovery*, 11(3):e1389.
- Payal Chandak, Kexin Huang, and Marinka Zitnik. 2023. Building a knowledge graph to enable precision medicine. *Scientific Data*, 10(1):67.
- Shuaichen Chang and Eric Fosler-Lussier. 2023. How to prompt llms for text-to-sql: A study in zero-shot, single-domain, and cross-domain settings. *Preprint*, arXiv:2305.11853.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, and Jonathan Larson. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Guandong Feng, Guoliang Zhu, Shengze Shi, Yue Sun, Zhongyi Fan, Sulin Gao, and Jun Hu. 2023. Robust nl-to-cypher translation for kbqa: Harnessing large language model with chain of prompts. In *Knowledge Graph and Semantic Computing: Knowledge Graph Empowers Artificial General Intelligence*, pages 317–326, Singapore. Springer Nature Singapore.
- Yichun Feng, Lu Zhou, Yikai Zheng, Ruikun He, Chao Ma, and Yixue Li. 2024. Knowledge graph-based thought: a knowledge graph enhanced llms framework for pan-cancer question answering. *bioRxiv*.
- Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou. 2023. Text-to-sql empowered by large language models: A benchmark evaluation. *Preprint*, arXiv:2308.15363.
- Martin Gerner, G. Nenadic, and Casey M. Bergman. 2010. Linnaeus: A species name identification system for biomedical literature. *BMC Bioinformatics*, 11:85 – 85.
- Aibo Guo, Xinyi Li, Guanchen Xiao, Zhen Tan, and Xiang Zhao. 2022. Spcql: A semantic parsing dataset for converting natural language into cypher. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, CIKM '22*, page 3973–3977, New York, NY, USA. Association for Computing Machinery.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Comput. Surv.*, 55(12).
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023a. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Xin Zhao, and Ji-Rong Wen. 2023b. StructGPT: A general framework for large language model to reason over structured data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9237–9251, Singapore. Association for Computational Linguistics.
- Daniel Martin Katz, Michael James Bommarito, Shang Gao, and Pablo Arredondo. 2023. Gpt-4 passes the bar exam. *Available at SSRN 4389233*.
- Adela Ljajić, Miloš Košprdić, Bojana Bašaragin, Darija Medvecki, Lorenzo Cassano, and Nikola Milošević. 2024. Scientific qa system with verifiable answers. *The 6th International Open Search Symposium*.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2023. Expertqa: Expert-curated questions and attributed answers. *arXiv preprint arXiv:2309.07852*.
- Nikola Milošević and Wolfgang Thielemann. 2023. Comparison of biomedical relationship extraction methods and models for knowledge graph creation. *Journal of Web Semantics*, 75:100756.
- Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. 2023. Capabilities of gpt-4 on medical challenge problems. *arXiv preprint arXiv:2303.13375*.
- OpenAI. 2023. Gpt-4 technical report. *Preprint*, arXiv:2303.08774.

A Appendix

- Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. 2024. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22.
- Mohammadreza Pourreza and Davood Rafiei. 2023. [DIN-SQL: Decomposed in-context learning of text-to-SQL with self-correction](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Priyanka Sen, Sandeep Mavadia, and Amir Saffari. 2023. [Knowledge graph-augmented language models for complex question answering](#). In *Proceedings of the 1st Workshop on Natural Language Reasoning and Structured Explanations (NLRSE)*, pages 1–8, Toronto, Canada. Association for Computational Linguistics.
- Saurabh Srivastava, Mayur Patidar, Sudip Chowdhury, Puneet Agarwal, Indrajit Bhattacharya, and Gautam Shroff. 2021. [Complex question answering on knowledge graphs using machine translation and multi-task learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3428–3439, Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Ethan Waisberg, Joshua Ong, Mouayad Masalkhi, Sharif Amit Kamran, Nasif Zaman, Prithul Sarker, Andrew G Lee, and Alireza Tavakkoli. 2023. Gpt-4: a new era of artificial intelligence in medicine. *Irish Journal of Medical Science (1971-)*, 192(6):3197–3200.
- Ziyu Zhao, Wei Liu, Tim French, and Michael Stewart. 2024. Cyspider: A neural semantic parsing corpus with baseline models for property graphs. In *AI 2023: Advances in Artificial Intelligence*, pages 120–132, Singapore. Springer Nature Singapore.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.
- Angelo Ziletti and Leonardo D’Ambrosi. 2024. [Retrieval augmented text-to-sql generation for epidemiological question answering using electronic health records](#). *Preprint*, arXiv:2403.09226.



Fact Finder

Factual Question-Answering

Select Example

- Insert Query What are the phenotypes associated with cardiofacial dysplasia? What are the genes responsible for psoriasis? Which diseases involve PINK1? How many drugs against epilepsy are available?
 Which medications have more off-label uses than approved indications?

Please insert your query

Which drugs are used to treat ocular hypertension?

Select Pipelines

LLM only Graph

Search

LLM only

Answer using LLM

Ocular hypertension is a condition where the pressure inside the eye (intraocular pressure) is higher than normal, which can lead to glaucoma if not managed properly. There are several types of medications that are commonly used to treat ocular hypertension:

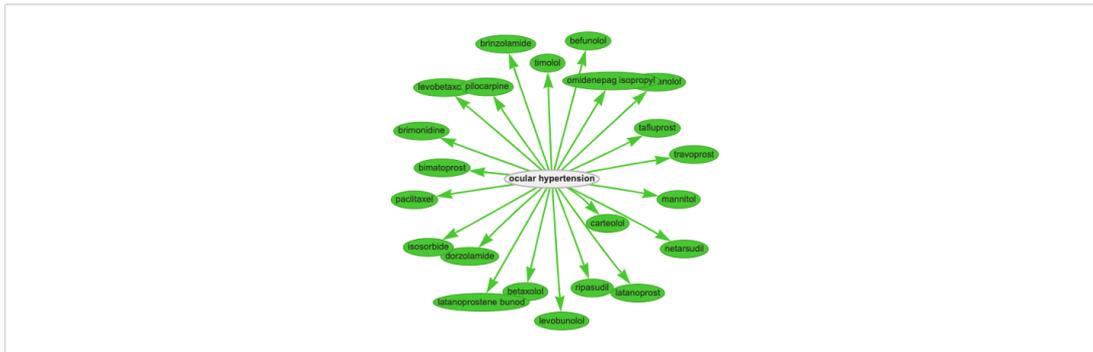
1. **Prostaglandin Analogues**: These help to increase the outflow of fluid from the eye. Examples include latanoprost (Xalatan), bimatoprost (Lumigan), and travoprost (Travatan).
2. **Beta Blockers**: These reduce the production of fluid within the eye. Examples include timolol (Timoptic) and betaxolol (Betoptol).

Graph Retrieval

Answer using Graph

The drugs used to treat ocular hypertension include tafiprost, paclitaxel, bimatoprost, betaxolol, timolol, carteolol, levobunolol, metipranolol, befunolol, levobetaxolol, pilocarpine, brinzolamide, dorzolamide, brimonidine, isosorbide, netarsudil, travoprost, latanoprost, latanoprostene bunod, ripasudil, mannitol, and omdenepag isopropyl.

Relevant Subgraph:



Show Evidence

Cypher Query:

```
MATCH (d:disease (name: "ocular hypertension"))-[:indication]->(drug:drug)
RETURN DISTINCT drug.name AS drug_name
```

Cypher Response:

```
[{'drug_name': 'tafluprost'}, {'drug_name': 'paclitaxel'}, {'drug_name': 'bimatoprost'}, {'drug_name': 'betaxolol'}, {'drug_name': 'timolol'}, {'drug_name': 'carteolol'}, {'drug_name': 'le
```

Cypher Prompt

Task: Generate Cypher statement to query a graph database described in the following schema.

Instructions:

Use only the provided relationship types and properties in the schema.

Do not use any other relationship types or properties that are not provided.

If there is no sensible Cypher statement for the given question and schema, state so and prepend SCHEMA_ERROR to your answer.

Any variables that are returned by the query must have readable names.

Remove modifying adjectives from the entities queried to the graph.

Answer Prompt

You are an assistant that helps to form nice and human understandable answers.

The information part contains the provided information that you must use to construct an answer.

The provided information is authoritative, you must never doubt it or try to use your internal knowledge to correct it.

Make the answer sound as a response to the question. Do not mention that you based the result on the given information.

If the provided information is a list, include all entries in your response.

If the provided information is empty, say that you don't know the answer.

Provided information:

Figure 7: User interface of Fact Finder for the question *Which drugs are used to treat ocular hypertension?*. The answers of the standalone LLM and our graph-based hybrid system are compared as output. In addition, the relevant subgraph is displayed as evidence together with the generated Cypher query and the answer from the graph.



Fact Finder

Factual Question-Answering

Select Example

- Insert Query What are the phenotypes associated with cardiofacrofacial dysplasia? What are the genes responsible for psoriasis? Which diseases involve PINK1? How many drugs against epilepsy are available?
 Which medications have more off-label uses than approved indications?

Please insert your query

Which drugs against epilepsy should not be used by patients with hypertension?

Select Pipelines

LLM only x Graph x

Search

LLM only

Answer using LLM

It's great that you're looking out for health and safety! When it comes to epilepsy and hypertension, it's important to consult with a healthcare professional for personalized advice. Some antiepileptic drugs (AEDs) can have interactions or side effects that might not be ideal for someone with hypertension. For example, certain AEDs like carbamazepine and phenytoin can potentially affect blood pressure or interact with medications used to treat hypertension.

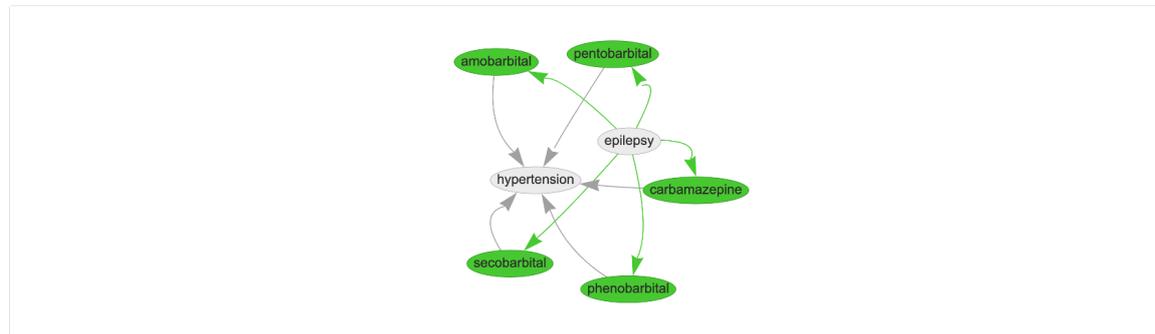
However, the best course of action is always to discuss your specific situation with your doctor or a pharmacist. They can provide guidance tailored to your individual health needs. Is there anything else you're curious about or any other topic you'd like to chat about?

Graph Retrieval

Answer using Graph

The drugs that should not be used by patients with hypertension are carbamazepine, phenobarbital, secobarbital, pentobarbital, and amobarbital.

Relevant Subgraph:



Show Evidence

Cypher Query:

```
MATCH (epilepsy:disease {name: "epilepsy"})-[:indication]->(drug:drug)-[:contraindication]->(hypertension:disease {name: "hypertension"})
RETURN DISTINCT drug_name AS drug_name
```

Cypher Response:

```
[{'drug_name': 'carbamazepine'}, {'drug_name': 'phenobarbital'}, {'drug_name': 'secobarbital'}, {'drug_name': 'pentobarbital'}, {'drug_name': 'amobarbital'}]
```

Cypher Prompt

Task: Generate Cypher statement to query a graph database described in the following schema.
Instructions:
Use only the provided relationship types and properties in the schema.
Do not use any other relationship types or properties that are not provided.
If there is no sensible Cypher statement for the given question and schema, state so and prepend SCHEMA_ERROR to your answer.
Any variables that are returned by the query must have readable names.
Remove modifying adjectives from the entities queried to the graph.

Answer Prompt

You are an assistant that helps to form nice and human understandable answers.
The information part contains the provided information that you must use to construct an answer.
The provided information is authoritative, you must never doubt it or try to use your internal knowledge to correct it.
Make the answer sound as a response to the question. Do not mention that you based the result on the given information.
If the provided information is a list, include all entries in your response.
If the provided information is empty, say that you don't know the answer.
Provided Information:

Figure 8: User interface of Fact Finder for the question *Which drugs against epilepsy should not be used by patients with hypertension?*.