

# Enhancing Text2Cypher with Schema Filtering

Makbule Gulcin Ozsoy<sup>1,\*</sup>

<sup>1</sup>Neo4j, London, UK

## Abstract

Knowledge graphs represent complex data using nodes, relationships, and properties. Cypher, a powerful query language for graph databases, enables efficient modeling and querying. Recent advancements in large language models allow translation of natural language questions into Cypher queries—Text2Cypher. A common approach is incorporating database schema into prompts. However, complex schemas can introduce noise, increase hallucinations, and raise computational costs. Schema filtering addresses these challenges by including only relevant schema elements, improving query generation while reducing token costs. This work explores various schema filtering methods for Text2Cypher task and analyzes their impact on token length, performance, and cost. Results show that schema filtering effectively optimizes Text2Cypher, especially for smaller models. Consistent with prior research, we find that larger models benefit less from schema filtering due to their longer context capabilities. However, schema filtering remains valuable for both larger and smaller models in cost reduction.

## Keywords

Text2Cypher, Large Language Model, Schema Filtering

## 1. Introduction

Databases are an essential part of modern computer systems for storing and managing data. They are typically accessed via query languages like SQL (for relational databases), SPARQL (for RDF graphs) or Cypher (for graph databases) which allow users to store and query data for insight [1]. Advancements in LLMs have enabled the translation of natural language questions into database queries (Text2SQL, Text2SPARQL, Text2Cypher), allowing non-expert users to query data models on their own terms.

To help contextualize an LLM when generating database queries from natural language, a common practice is to incorporate database schema information. Figure 1 shows an example schema where nodes (e.g., Organization, Person) connect through relations (e.g., Has\_CEO, Has\_Investor) with their properties (e.g., name, age). Schemas can be provided to LLMs via prompting, but complex schemas introduce noise, increase hallucinations, and raise costs [2, 3]. Schema filtering addresses these challenges by selecting only relevant elements, improving query generation while reducing token costs.

In this paper, we apply five schema linking and filtering approaches that improve Text2Cypher: Two static methods that extract the full database schema in different formats and three dynamic methods that prune the schema based on the input question. We evaluate their impact on a Text2Cypher dataset, analyzing token distribution, Cypher generation performance, and cost. Our main contributions are:

- We propose new schema filtering techniques. The two static methods use the full database schema in different formats, while our three dynamic methods prune it based on the input question.
- We analyze their impact on Text2Cypher task, specifically on prompt token length distribution, query generation performance, and computational cost.
- Our results show that schema filtering improves Text2Cypher efficiency. While larger models benefit less due to their extended context windows, smaller models perform better with shorter prompts. Nevertheless, schema filtering remains a cost-effective strategy for all models.

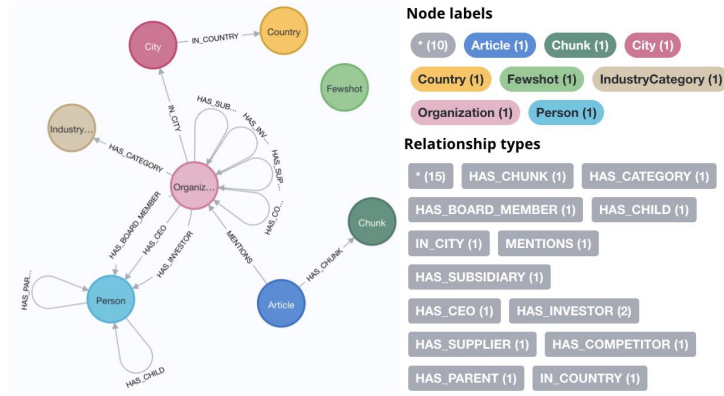
The paper is structured as follows: Section 2 covers related work, and Section 3 details our schema-filtering approaches for the Text2Cypher task. Section 4 presents our experiments and results, and Section 5 concludes the paper.

*LLM-TEXT2KG 2025: 4th International Workshop on LLM-Integrated Knowledge Graph Generation from Text (Text2KG), June 1 - June 5, 2025, co-located with the Extended Semantic Web Conference (ESWC 2025), Portorož, Slovenia*

✉ makbule.ozsoy@neo4j.com (M. G. Ozsoy)



© 2025 Copyright © 2025 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



**Figure 1:** Overview of an Example Database

**Table 1**

Instructions used

Type	Instruction prompt
System Instruction	Task: Generate Cypher statement to query a graph database. 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 in the schema. Do not include any explanations or apologies in your responses. Do not respond to any questions that might ask anything else than for you to construct a Cypher statement. Do not include any text except the generated Cypher statement.
User Instruction	Generate Cypher statement to query a graph database. Use only the provided relationship types and properties in the schema. <b>Schema:</b> {schema} <b>Question:</b> {question} <b>Cypher output:</b>

## 2. Related Work

### 2.1. Natural Language to Database Query Language

Recent advances in large language models (LLMs) have significantly improved the ability to translate natural language into database query languages. For instance, there has been extensive research on the Text2SQL and Text2SPARQL tasks, which translates natural language queries to SQL or SPARQL, respectively [4, 5, 6, 7, 8, 9, 10, 11, 12]. Until recently, the Text2Cypher task, which translates natural language into Cypher; the query language used by Neo4j and other graph database systems; had received less attention. However, with advancements in the integration of large language models (LLMs) and knowledge graphs, text-to-graph query language (GQL) tasks, particularly Text2Cypher, have gained increasing interest. Several datasets have been developed to support Text2Cypher research, including Opitz and Hochgeschwender [13], S2CTrans [14], CySpider [15], Rel2Graph [16], SyntheT2C [17], and Text2Cypher [18]. Additionally, studies have explored benchmarking and fine-tuning models for this task, with contributions such as GPT4Graph [19], TopoChat [20], Baraki et al. [21], FCAV [22], Liang et al. [23] and Text2Cypher [18]. In most cases, the baseline model is fine-tuned using prompts that include natural language questions, database schema information, and ground-truth Cypher queries.

### 2.2. Schema Filtering in Query Generation

Schema information is essential for accurate query generation, ensuring correct linking of query terms to database structures [2, 3]. This process, known as schema linking, plays a key role in Text2SQL and Text2Cypher tasks by mapping query words to relevant database elements [24, 2]. While providing the full schema in the prompt is possible, schema filtering is often preferred to reduce noise, computational cost, and hallucinations [2, 25]. However, we must remain aware that excessive filtering can remove essential components, harming accuracy [26]. Early Text2SQL schema filtering relied on heuristics like string matching, as seen in IRNet [5] and TypeSQL [27]. Later, learning-based methods such as Dong et

```

"Node properties:\n-
**Person**\n - `name`: STRING Example: \"Julie Spellman
Sweet\"\n - `id`: STRING Example:
\"Eaf0bpz6NNqLVUCqNZPAew\"\n - `summary`: STRING Example:
\"CEO at Accenture\"\n-
...
**Chunk**\n - `text`: STRING Example: \"Boston and
Mumbai-based consumer behavior analyses\"\n-
**Fewshot**\n - `Question`: STRING Available options:
['How is Emil Eifrem connected to Michael Hunger? ', 'What
are the latest news regarding Google? '... - `Cypher`:
STRING Available options: ['MATCH (p1:Person {name:\"Emil
Eifrem\"})), (p2:Pers', 'MATCH (o:Organization
{name:\"Google\"})<-[:MENTIO'... \n
Relationship properties:\n\n
The relationships:\n
(:Person)-[:HAS_PARENT]->(:Person)\n
(:Person)-[:HAS_CHILD]->(:Person)\n
(:Organization)-[:IN_CITY]->(:City)\n
(:Organization)-[:HAS_CEO]->(:Person)\n
...

```

(a) Enhanced Schema

```

"Node properties:\n
Person {name: STRING, id: STRING, summary: STRING}\n
...
Chunk {text: STRING, embedding: LIST, embedding_google:
LIST, embedding_google_004: LIST}\n
Fewshot {Question: STRING, Cypher: STRING, id: INTEGER,
embedding: LIST}\n
Relationship properties:\n\n
The relationships:\n
(:Person)-[:HAS_PARENT]->(:Person)\n
(:Person)-[:HAS_CHILD]->(:Person)\n
(:Organization)-[:IN_CITY]->(:City)\n
(:Organization)-[:HAS_CATEGORY]->(:IndustryCategory)\n
...

```

(b) Base Schema

**Figure 2:** Example static schemas for the example database presented in Figure 1

al. [28], Bogin et al. [29], and RAT-SQL [30] were proposed. Recent approaches utilize LLMs through prompting, fine-tuning, or agent-based techniques, such as DIN-SQL [31], RESDSQL [6], CHESS [32], E-SQL [2], ExSL [33] and KaSLA [34]. While schema filtering is common, studies suggest it is less necessary for LLMs with long context windows but remains valuable for smaller models [26, 2, 3]. The trade-off is, however, that larger context sizes increase latency and computational cost for complex databases, making filtering highly beneficial [3].

Research on schema filtering for Text2Cypher or other graph query languages (Text2GQL) is presently limited compared to Text2SQL. Liang et al. [23] explored aligning LLMs for a Text2GQL task in Chinese, using a schema filtering module that executes: (i) extraction of the database schema as a dictionary, (ii) extraction of the named entities from the query, and (iii) mapping these entities to the schema dictionary. For queries requiring multiple nodes and relations, they used A\* algorithm [35] to find the shortest path. NAT-NL2GQL [36] includes a module for preprocessing inputs and executing schema extraction, following a similar approach to Liang et al. [23]. Additionally, they use an LLM for filtering multiple matched schema items before proceeding with the Text2GQL task. In this work, we examine the impact of schema filtering on the Text2Cypher task, focusing on both performance and cost.

### 3. Schema Filtering for Text2Cypher

We now present schema filtering for Text2Cypher using a template from [18] (Table 1), focusing on the schema field with two static and three dynamic formats.

#### 3.1. Static Schemas

Cypher is the query language for Neo4j, a graph database. Neo4j offers various tools for retrieving database schema information, based on the database structure rather than the input query. While this allows efficient caching, it leads to longer schema representations, increasing token length and context requirements for LLMs. We utilized two static schema formats provided by Neo4j frameworks:

- **Enhanced Schema:** This is one of the default schema types provided by Neo4j. It provides an enhanced view of the database schema, including list of nodes, relationships and their properties. It additionally provides example values for the fields. For instance, if the property is the 'name' of the 'Actor' node, examples might include: ['Tom Hanks', 'Julia Roberts', ...]. An example enhanced schema is presented in Figure 2a.

```
"Node properties:\n
Organization {name: STRING, nbrEmployees: INTEGER,
isDissolved: BOOLEAN, id: STRING, motto: STRING, summary:
STRING, isPublic: BOOLEAN, diffbotId: STRING, revenue:
FLOAT}\n
Article {id: STRING, sentiment: FLOAT, siteName: STRING,
summary: STRING, date: DATE_TIME, title: STRING, author:
STRING}\n
Relationship properties:\n\n
The relationships:\n
(:Organization)-[:HAS_SUPPLIER]->(:Organization)\n
(:Organization)-[:IN_CITY]->(:City)\n
(:Organization)-[:HAS_CEO]->(:Person)\n
(:Organization)-[:HAS_BOARD_MEMBER]->(:Person)\n
(:Organization)-[:HAS_CATEGORY]->(:IndustryCategory)\n
(:Organization)-[:HAS_INVESTOR]->(:Organization)\n
(:Organization)-[:HAS_INVESTOR]->(:Person)\n
(:Organization)-[:HAS_SUBSIDIARY]->(:Organization)\n
(:Organization)-[:HAS_COMPETITOR]->(:Organization)\n
(:Article)-[:HAS_CHUNK]->(:Chunk)\n
(:Article)-[:MENTIONS]->(:Organization)"
```

(a) Pruned By Exact-Match Schema

```
"Node properties:\n
Organization {name: STRING, nbrEmployees: INTEGER,
isDissolved: BOOLEAN, id: STRING, motto: STRING, summary:
STRING, isPublic: BOOLEAN, diffbotId: STRING, revenue:
FLOAT}\n
Article {id: STRING, sentiment: FLOAT, author: STRING,
siteName: STRING, summary: STRING, date: DATE_TIME, title:
STRING}\n
Relationship properties:\n\n
The relationships:\n
(:Organization)-[:IN_CITY]->(:City)\n
(:Organization)-[:HAS_CEO]->(:Person)\n
(:Organization)-[:HAS_CATEGORY]->(:IndustryCategory)\n
(:Organization)-[:HAS_SUBSIDIARY]->(:Organization)\n
(:Organization)-[:HAS_SUPPLIER]->(:Organization)\n
(:Organization)-[:HAS_INVESTOR]->(:Organization)\n
(:Organization)-[:HAS_INVESTOR]->(:Person)\n
(:Organization)-[:HAS_BOARD_MEMBER]->(:Person)\n
(:Organization)-[:HAS_COMPETITOR]->(:Organization)\n
(:Article)-[:HAS_CHUNK]->(:Chunk)\n
(:Article)-[:MENTIONS]->(:Organization)"
```

(b) NER Masked & Pruned by Exact-Match Schema

```
"Node properties:\n
Person {name: STRING, id: STRING, summary: STRING}\n
Organization {isPublic: BOOLEAN, motto: STRING,
nbrEmployees: INTEGER, isDissolved: BOOLEAN, id: STRING,
name: STRING, diffbotId: STRING, summary: STRING, revenue:
FLOAT}\n
IndustryCategory {name: STRING, id: STRING}\n
City {id: STRING, summary: STRING, name: STRING}\n
Country {name: STRING, id: STRING, summary: STRING}\n
Article {id: STRING, sentiment: FLOAT, siteName: STRING,
summary: STRING, date: DATE_TIME, title: STRING, author:
STRING}\n
Fewshot {Question: STRING, Cypher: STRING, id: INTEGER,
embedding: LIST}\n
Relationship properties:\n\n
The relationships:\n
(:Person)-[:HAS_PARENT]->(:Person)\n
(:Person)-[:HAS_CHILD]->(:Person)\n
(:Organization)-[:HAS_SUPPLIER]->(:Organization)\n
(:Organization)-[:IN_CITY]->(:City)\n
(:Organization)-[:HAS_CATEGORY]->(:IndustryCategory)\n
(:Organization)-[:HAS_CEO]->(:Person)\n
(:Organization)-[:HAS_BOARD_MEMBER]->(:Person)\n
(:Organization)-[:HAS_INVESTOR]->(:Organization)\n
(:Organization)-[:HAS_INVESTOR]->(:Person)\n
(:Organization)-[:HAS_SUBSIDIARY]->(:Organization)\n
(:Organization)-[:HAS_COMPETITOR]->(:Organization)\n
(:Article)-[:HAS_CHUNK]->(:Chunk)\n
(:Article)-[:MENTIONS]->(:Organization)"
```

(c) Pruned by Similarity Schema

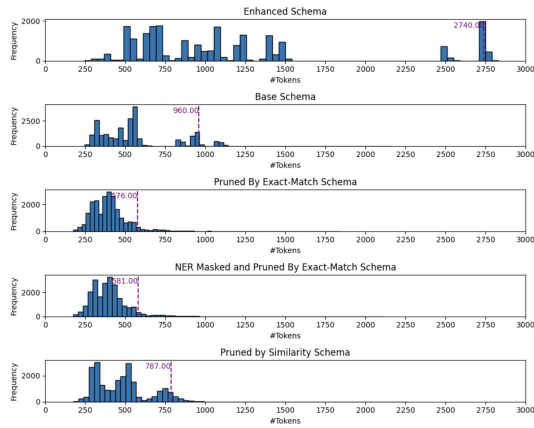
**Figure 3:** Dynamic schemas for the question: "List the articles that mention the organization 'Acme Group'"

- **Base Schema:** This is another default schema types provided by Neo4j. It provides similar information as the Enhanced Schema, except it does not include examples of properties, and the formatting is different. An example for this schema format is presented in Figure 2b.

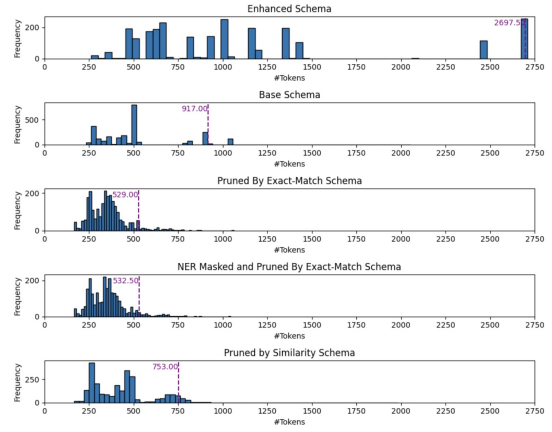
### 3.2. Dynamically Pruned Schemas

We implement three dynamic schema filtering approaches, which prune the baseline schemas based on the input natural language question.

- **Pruned By Exact-Match:** This approach compares node labels, relationship types, and properties to words in the input question. Similar to Liang et al. [23] and NAT-NL2GQL [36], if an exact case-insensitive match is found, the corresponding schema elements are retained; otherwise, they are removed. Our method also considers properties as well as labels, and we retain multiple matching elements (e.g., synonyms) to prevent excessive pruning. See Figure 3a for an example.
- **NER Masked & Pruned By Exact-Match:** This approach replaces named entities with their entity types before applying exact-match filtering. NER-masking prevents irrelevant matches. For example, in the query "List the articles that mention the organization 'Acme Energy'", it avoids incorrect matches, such as retaining properties of a node labeled 'Energy,' which is unrelated. See Figure 3b for an example.
- **Pruned by Similarity:** This approach extends exact-match pruning by incorporating similarity-based filtering. Instead of requiring an exact match, it computes similarity scores between query



(a) Token Distribution on Training Set



(b) Token Distribution on Test Set

**Figure 4:** Token distributions of the training and test sets, based on the tokenizer from the Llama-3.1-8B model. The 95th percentile (p95) value is marked with a purple line.

Schema Type	min	max	median	p95
Enhanced Schema	263	2710	921	2697
Base Schema	233	2052	494	917
Pruned By Exact-Match	166	1063	344	529
NER Masked & Pruned By Exact-Match	166	1045	344	532
Pruned by Similarity	166	1982	422	753

**Table 2**

Token distribution statistics of the test set

terms and schema elements, retaining only those above a predefined threshold. While various similarity measures could be used, we rely on embedding-based similarity. An example of this schema filtering approach is shown in Figure 3c.

## 4. Experimental Setup and Results

### 4.1. Experimental Setup and Evaluation Metrics

We conducted experiments using a publicly available Text2Cypher dataset [18], focusing on a subset with accessible databases for query execution, resulting in 22,093 training and 2,471 test samples. Schema filtering was assessed using the 'unsloth/Meta-Llama-3.1-8B-Instruct-bnb-4bit', 'unsloth/Qwen2.5-7B-Instruct-bnb-4bit' and 'GoogleAIStudio/Gemini-1.5-Flash' models, referred as Llama-3.1-8B, Qwen2.5-7B and Gemini-1.5-Flash, respectively, in the remainder of the paper. For Cypher generation, after utilizing the LLMs, an additional post-processing step is executed to remove unwanted text, such as 'cypher:' suffix. Furthermore, the spaCy framework is used for named entity extraction and similarity computations. To compute evaluation metrics, we used the Hugging Face Evaluate library [37]. We employed two evaluation procedures: (i) Translation-based (Lexical) evaluation: Compares generated Cypher queries with reference queries based on text content. We used Google-BLEU score while presenting the results. (ii) Execution-based evaluation: Executes both generated and reference queries on target databases and compares their outputs (sorted lexicographically) using the same metrics as the translation-based evaluation. We used ExactMatch score while presenting the results.

### 4.2. Evaluation Results

We evaluate the proposed schema formats based on (i) token distribution and cost, and (ii) performance.



**Table 3**

Schema type vs. Cost. Notes: Median token counts: Table 2, Number of instances: 20K, Prices-Feb. 27, 2025: (i) GoogleAIStudio/Gemini-2.0-Flash: \$0.15 / 1M tokens (ii) Anthropic/Claude 3.5 Haiku: \$0.80 / 1M tokens (iii) Self hosted/LLaMA-3.1-8B: Assumed A40 48GB machine on RunPod (\$0.44 per hour) with 20 tokens/sec

Schema Type	# Total tokens	Gemini-2.0 Flash	Claude-3.5 Haiku	LLaMA-3.1-8B (Self Hosted)
<b>Enhanced Schema</b>	20K * 921=~18.5M	\$2.7	\$14.7	\$112.5
<b>Base Schema</b>	20K * 494=~10M	\$1.5	\$7.9	\$60.3
<b>Pruned By Exact-Match</b>	20K * 344=~7M	\$1.0	\$5.5	\$42.0
<b>NER Masked &amp;</b>	20K * 344=~7M	\$1.0	\$5.5	\$42.0
<b>Pruned By Exact-Match</b>				
<b>Pruned by Similarity</b>	20K * 422=~8.5M	\$1.3	\$6.7	\$51.5

#### 4.2.1. Impact on Token Distribution & Cost

Schema format impacts both prompt length and token count. For example, with the Llama-3.1-8B tokenizer, the base prompt is about 150 tokens, but adding schema information increases it to over 2,700 tokens. Figure 4 shows token distributions for training and test sets. Table 2 provides additional token details for the test set. Results show that the Enhanced Schema leads to the longest prompts, while switching to the Base Schema reduces the P95 token length by one-third. Exact-match pruning (with or without NER masking) further reduces the P95 token length to 1/6th of the original. Similarity-based pruning increases schema length but reduces the P95 token length to about 1/4th of the original.


Reducing the token count reduces costs, whether for LLM vendor payments or infrastructure expenses for self-hosted models (e.g., storage and GPU access). In a scenario with 20,000 instances, where input token length aligns with the median in Table 2, we compare costs across models (Table 3). In the table, we assume output lengths remain constant and only input tokens contribute to the cost. The results show that cost scales linearly with token usage, but factors like output token count, caching, and batch processing can affect this. Shorter prompts lead to significant cost reductions.

While dynamic pruning reduces token length and costs, it may introduce computational overhead as a side-effect. Unlike Enhanced or Base Schema (which are cached), dynamic pruning is performed for each query, which might increase latency. However, we observe this overhead is minimal, especially for methods like ‘Pruned by Exact-Match,’ which uses regular expression matching.

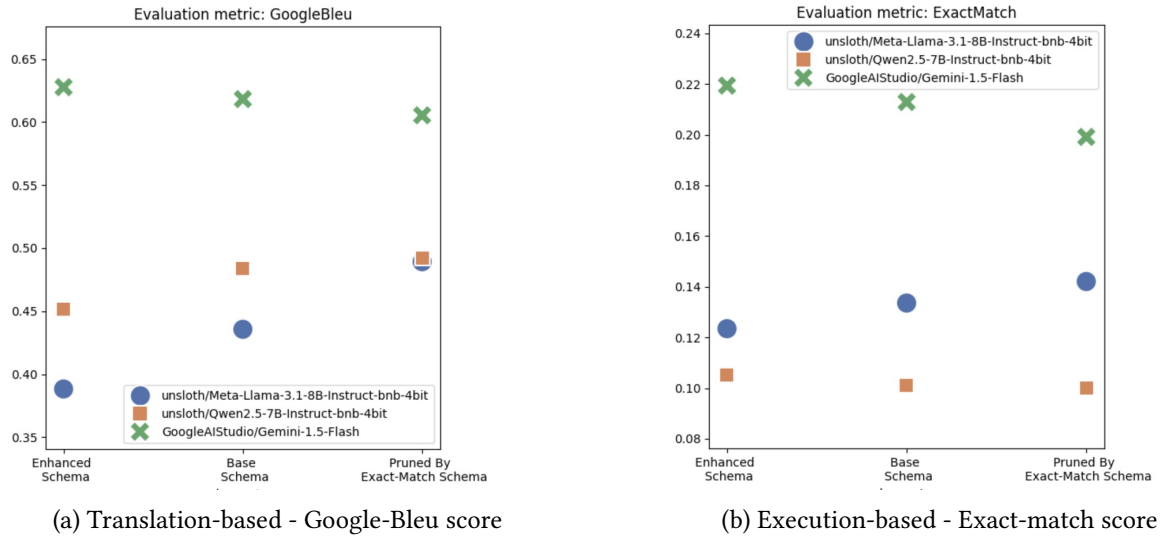
#### 4.2.2. Impact on Performance

We evaluate the impact of proposed schema formats on Text2Cypher performance using the Llama-3.1-8B model. Figure 5 presents the results, showing that longer prompts lead to lower performance. The highest accuracy is achieved with the ‘Pruned by Exact-Match Schema.’ NER masking and Similarity-Based Matching did not improve performance but may be beneficial for other datasets.

We further compared the performance of different LLMs on a selected subset of schema formats. In addition to Llama-3.1-8B, we evaluated Qwen2.5-7B and Gemini-1.5-Flash. While Llama-3.1-8B and Qwen2.5-7B are similar in size, they differ in multiple ways, such as tokenization strategies. Gemini-1.5-Flash, in contrast, has a larger model size and a significantly longer context window. For comparison, we used three schema formats—Enhanced, Base, and Pruned by Exact-Match—with decreasing token lengths. Figure 6 presents the results, highlighting key trends: (i) In terms of lexical (translation-based) comparison, performance of Llama-3.1-8B and Qwen2.5-7B models are improved as prompt length decreased. However, Gemini-1.5-Flash had the opposite trend, performing better with longer prompts. The drop in Gemini-1.5-Flash for shorter prompts was minor, remaining below 5%. (ii) In terms of execution-based evaluation, Llama-3.1-8B model showed improved performance with shorter prompts, while Qwen2.5-7B and Gemini-1.5-Flash experienced slight declines, both around 2%. These findings align with observation made by previous research [2, 3]: The impact of schema length varies across models, with Gemini-1.5-Flash potentially benefiting from longer context while the other smaller models perform better with shorter inputs.



**Figure 5:** Performance across various schema formats



**Figure 6:** Performance comparison of various models

## 5. Conclusion

We presented schema filtering for Text2Cypher and analyzed its effects on token length, performance, and cost. We found that reducing schema size improved performance for most models, and reduced cost for all of those we tested. Comparison of various models revealed that smaller models performed better with shorter prompts, while larger models benefited from longer contexts. Dynamically pruning schemas reduced both token counts and cost, introducing slightly more latency but remained the most efficient overall. This work has two main limitations. First, experiments used a subset of a public dataset [18], selecting instances with accessible databases for schema extraction. However, these demo-oriented databases often have simpler schemas than real-world ones. The longest schema had around 2700 tokens, whereas real databases likely have more complex structures, making schema filtering more critical. Second, our filtering methods are heuristic-based. More advanced techniques, like those in Text2SQL (see Section 2), may yield better results and require further exploration. In the future, we will explore adaptive schema selection based on model characteristics, as well as the impact of schema filtering on the fine-tuning process and its effects on fine-tuned models.

## Declaration on Generative AI

During the preparation of this work, the author(s) used Chat-GPT in order to: 'Improve writing style' and 'Paraphrase and reword'. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

## References

- [1] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, et al., Knowledge graphs, *ACM Computing Surveys (Csur)* 54 (2021) 1–37.
- [2] H. A. Caferoğlu, Ö. Ulusoy, E-sql: Direct schema linking via question enrichment in text-to-sql, *arXiv preprint arXiv:2409.16751* (2024).
- [3] Y. Chung, G. T. Kakkar, Y. Gan, B. Milne, F. Ozcan, Is long context all you need? leveraging llm's extended context for nl2sql, *arXiv preprint arXiv:2501.12372* (2025).
- [4] T. Yu, R. Zhang, K. Yang, M. Yasunaga, D. Wang, Z. Li, J. Ma, I. Li, Q. Yao, S. Roman, et al., Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task, *arXiv preprint arXiv:1809.08887* (2018).
- [5] J. Guo, Z. Zhan, Y. Gao, Y. Xiao, J.-G. Lou, T. Liu, D. Zhang, Towards complex text-to-sql in cross-domain database with intermediate representation, *arXiv preprint arXiv:1905.08205* (2019).
- [6] H. Li, J. Zhang, C. Li, H. Chen, Resdsq: Decoupling schema linking and skeleton parsing for text-to-sql, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2023, pp. 13067–13075.
- [7] Y. Fan, Z. He, T. Ren, C. Huang, Y. Jing, K. Zhang, X. S. Wang, Metasql: A generate-then-rank framework for natural language to sql translation, *arXiv preprint arXiv:2402.17144* (2024).
- [8] K. Shen, M. Kejriwal, Select-sql: Self-correcting ensemble chain-of-thought for text-to-sql, *arXiv preprint arXiv:2409.10007* (2024).
- [9] J. Lee, I. Baek, B. Kim, H. Lee, Safe-sql: Self-augmented in-context learning with fine-grained example selection for text-to-sql, *arXiv preprint arXiv:2502.11438* (2025).
- [10] F. Brei, J. Frey, L.-P. Meyer, Leveraging small language models for text2sparql tasks to improve the resilience of ai assistance, *arXiv preprint arXiv:2405.17076* (2024).
- [11] L.-P. Meyer, J. Frey, F. Brei, N. Arndt, Assessing sparql capabilities of large language models, *arXiv preprint arXiv:2409.05925* (2024).
- [12] V. Emonet, J. Bolleman, S. Duvaud, T. M. de Farias, A. C. Sima, Llm-based sparql query generation from natural language over federated knowledge graphs, *arXiv preprint arXiv:2410.06062* (2024).
- [13] D. Opitz, N. Hochgeschwender, From zero to hero: generating training data for question-to-cypher models, in: *Proceedings of the 1st International Workshop on Natural Language-based Software Engineering*, 2022, pp. 17–20.
- [14] Z. Zhao, X. Ge, Z. Shen, C. Hu, H. Wang, S2ctrans: Building a bridge from sparql to cypher, in: *International Conference on Database and Expert Systems Applications*, Springer, 2023, pp. 424–430.
- [15] Z. Zhao, W. Liu, T. French, M. Stewart, Cyspider: A neural semantic parsing corpus with baseline models for property graphs, in: *Australasian Joint Conference on Artificial Intelligence*, Springer, 2023, pp. 120–132.
- [16] Z. Zhao, W. Liu, T. French, M. Stewart, Rel2graph: Automated mapping from relational databases to a unified property knowledge graph, *arXiv preprint arXiv:2310.01080* (2023).
- [17] Z. Zhong, L. Zhong, Z. Sun, Q. Jin, Z. Qin, X. Zhang, Synthet2c: Generating synthetic data for fine-tuning large language models on the text2cypher task, *arXiv preprint arXiv:2406.10710* (2024).
- [18] M. G. Ozsoy, L. Messallem, J. Besga, G. Minneci, Text2cypher: Bridging natural language and graph databases, in: *Proceedings of the Workshop on Generative AI and Knowledge Graphs (GenAIK)*, 2025, pp. 100–108.
- [19] J. Guo, L. Du, H. Liu, M. Zhou, X. He, S. Han, Gpt4graph: Can large language models understand graph structured data? an empirical evaluation and benchmarking, *arXiv preprint arXiv:2305.15066* (2023).
- [20] H. Xu, B. Zhang, Z. Jin, T. Zhu, Q. Wu, H. Weng, Topochat: Enhancing topological materials retrieval with large language model and multi-source knowledge, *arXiv preprint arXiv:2409.13732* (2024).
- [21] W. W. Baraki, Leveraging large language models for accurate Cypher query generation: Natural language query to Cypher statements, Master degree project, University of Skövde, 2024. <https://www.scribd.com/document/800000000/Natural-language-query-to-Cypher-statements>



[//www.diva-portal.org/smash/get/diva2:1881385/FULLTEXT01.pdf](http://www.diva-portal.org/smash/get/diva2:1881385/FULLTEXT01.pdf).

- [22] Y. Liu, X. Wang, J. Ge, H. Wang, D. Xu, Y. Jia, Text to graph query using filter condition attributes, *Proceedings of the VLDB Endowment*. ISSN 2150 (2024) 8097.
- [23] Y. Liang, K. Tan, T. Xie, W. Tao, S. Wang, Y. Lan, W. Qian, Aligning large language models to a domain-specific graph database for nl2gql, in: *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 1367–1377.
- [24] W. Lei, W. Wang, Z. Ma, T. Gan, W. Lu, M.-Y. Kan, T.-S. Chua, Re-examining the role of schema linking in text-to-sql, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 6943–6954.
- [25] Z. Cao, Y. Zheng, Z. Fan, X. Zhang, W. Chen, X. Bai, Rsl-sql: Robust schema linking in text-to-sql generation, *arXiv preprint arXiv:2411.00073* (2024).
- [26] K. Maamari, F. Abubaker, D. Jaroslawicz, A. Mhedhbi, The death of schema linking? text-to-sql in the age of well-reasoned language models, *arXiv preprint arXiv:2408.07702* (2024).
- [27] T. Yu, Z. Li, Z. Zhang, R. Zhang, D. Radev, Typesql: Knowledge-based type-aware neural text-to-sql generation, *arXiv preprint arXiv:1804.09769* (2018).
- [28] Z. Dong, S. Sun, H. Liu, J.-G. Lou, D. Zhang, Data-anonymous encoding for text-to-sql generation, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 5405–5414.
- [29] B. Bogin, M. Gardner, J. Berant, Global reasoning over database structures for text-to-sql parsing, *arXiv preprint arXiv:1908.11214* (2019).
- [30] B. Wang, R. Shin, X. Liu, O. Polozov, M. Richardson, Rat-sql: Relation-aware schema encoding and linking for text-to-sql parsers, *arXiv preprint arXiv:1911.04942* (2019).
- [31] M. Pourreza, D. Rafiei, Din-sql: Decomposed in-context learning of text-to-sql with self-correction, *Advances in Neural Information Processing Systems* 36 (2023) 36339–36348.
- [32] S. Talaei, M. Pourreza, Y.-C. Chang, A. Mirhoseini, A. Saberi, Chess: Contextual harnessing for efficient sql synthesis, *arXiv preprint arXiv:2405.16755* (2024).
- [33] M. Glass, M. Eyceoz, D. Subramanian, G. Rossiello, L. Vu, A. Gliozzo, Extractive schema linking for text-to-sql, *arXiv preprint arXiv:2501.17174* (2025).
- [34] Z. Yuan, H. Chen, Z. Hong, Q. Zhang, F. Huang, X. Huang, Knapsack optimization-based schema linking for llm-based text-to-sql generation, *arXiv preprint arXiv:2502.12911* (2025).
- [35] P. E. Hart, N. J. Nilsson, B. Raphael, A formal basis for the heuristic determination of minimum cost paths, *IEEE transactions on Systems Science and Cybernetics* 4 (1968) 100–107.
- [36] Y. Liang, T. Xie, G. Peng, Z. Huang, Y. Lan, W. Qian, Nat-nl2gql: A novel multi-agent framework for translating natural language to graph query language, *arXiv preprint arXiv:2412.10434* (2024).
- [37] HuggingFace, Huggingface evaluate, 2024. <https://huggingface.co/evaluate-metric>.