

Constrained Entity Selection under Partial Knowledge for LLM-Based Knowledge Graph QA^{*}

Emanuel Kitzelmann¹[0000–0001–5614–3463]

Brandenburg University of Applied Sciences, Brandenburg a.d.H., Germany
 emanuel.kitzelmann@th-brandenburg.de

Abstract. Large language models are increasingly used for knowledge graph question answering (KGQA), but can fail to correctly ground answers in the underlying graph. Current approaches to LLM-based KGQA either rely on full semantic parsing into executable queries such as SPARQL, which is brittle in practice due to complex schemas or incompleteness of real-world KGs, or on LLM-reasoning and answer generation over KGs, which can be more robust but lacks formal guarantees.

In this work, we study a complementary setting in which *candidate* answers are generated by an LLM-based system and subsequently verified using lightweight symbolic constraints derived from the question. We introduce *Constrained Entity Selection under Partial Knowledge (CES-PK)*, a problem formulation that focuses on eliminating invalid answers and providing symbolic support for valid ones without requiring construction of executable logical forms. To account for incomplete KGs, we employ a three-valued constraint semantics (*satisfied*, *violated*, *unknown*) that avoids incorrect rejections under open-world assumptions.

To demonstrate the effects of our method, we instantiate this framework over the Hetionet biomedical knowledge graph and evaluate the impact of type, relation, and exclusion constraints. Experiments show that precision improves by filtering invalid candidates, while recall is preserved due to retaining candidates whose constraints are not explicitly violated. Satisfied constraints provide additional positive symbolic evidence to rank remaining candidates.

Keywords: Knowledge Graph Question Answering · Large Language Models · Constraint-Based Verification · Incomplete Knowledge Graphs

1 Introduction

Large language models (LLMs) are increasingly used for knowledge graph question answering (KGQA) [10], enabling natural language access to structured

^{*} This preprint has not undergone peer review or any post-submission improvements or corrections. The Version of Record of this contribution is published in KI 2026: Advances in Artificial Intelligence, KI 2026, LNCS, vol 16830, and is available online at https://doi.org/10.1007/978-3-032-32335-4_21

knowledge. However, LLMs can fail to correctly ground answers in the underlying graph and may violate basic structural or semantic constraints.

Current LLM approaches to KGQA typically follow one of two paradigms. In semantic parsing methods [4,3], questions are translated into executable logical forms such as SPARQL queries. While this enables precise and verifiable answers, generating correct queries remains challenging in practice, especially for complex schemas and incomplete KGs [5]. In contrast, retrieval-augmented or agentic approaches such as [11] use LLMs to retrieve and reason over KGs. These methods are more flexible and robust to missing information but lack formal guarantees and may produce hallucinated or non-grounded answers.

In this work, we propose a complementary perspective that avoids both full semantic parsing and unconstrained LLM reasoning. Instead of deriving complete logical forms, we extract lightweight necessary constraints from the question and use them to verify candidate answers produced by an upstream system. This yields a verification layer that filters invalid candidates, retains candidates that cannot be disproven under incomplete knowledge, and provides symbolic support for valid candidates.

We formalize this setting as *Constrained Entity Selection under Partial Knowledge* (CES-PK). To account for the open-world nature of real-world knowledge graphs, we employ a three-valued constraint semantics that distinguishes between satisfied, violated, and unknown constraints. This allows us to reject candidates only in the presence of explicit violations, thereby avoiding false negatives caused by missing information.

We instantiate this framework for entity-centric queries over the Hetionet¹ biomedical knowledge graph [7] and evaluate its effect on simulated candidate answer sets. Using controlled candidate sets that simulate LLM-generated answers, we show that constraint-based verification improves precision substantially while preserving recall and yielding support for KG-grounded answers. Furthermore, we analyze the complementary roles of constraint types: exclusion constraints contribute to filtering, whereas positive relation constraints contribute to verification through support.

Contributions. This paper makes the following contributions:

- We introduce Constrained Entity Selection under Partial Knowledge (CES-PK), a problem formulation for post-hoc verification of KGQA answers without requiring full semantic parsing.
- We analyze the complementary roles of constraint types for filtering and verification with a three-valued constraint semantics for incomplete KGs.
- We provide an empirical study demonstrating that lightweight constraint-based verification improves precision while preserving recall and additionally provides support scores for retained answers.

¹ <https://het.io/>

2 Problem and Method

2.1 Problem Setting

We consider knowledge graph question answering (KGQA) in a setting where candidate answers are produced by an upstream component (e.g., an agentic LLM pipeline) and may contain incorrect or unsupported entities.

Let $G = (V, E)$ be a knowledge graph [8] with entities V and typed relations $E \subseteq V \times R \times V$. Given a question q , we assume a candidate set $C(q) \subseteq V$ and an (unknown) set of correct answers $A(q) \subseteq V$. Rather than deriving $A(q)$ directly, we aim to identify candidates consistent with necessary conditions implied by the question.

Constrained Entity Selection under Partial Knowledge (CES-PK). Given $C(q)$ and a set of constraints $\mathcal{C}(q)$ derived from q , compute a subset $C^*(q) \subseteq C(q)$ such that no retained candidate violates any constraint under incomplete knowledge and associate remaining answers in $C^*(q)$ with support scores reflecting the degree of constraint satisfaction.

2.2 Assumptions on the Knowledge Graph

We assume that the KG is factually correct but may be incomplete, i.e., all represented triples are true, while some true facts may be missing (open-world assumption). Furthermore, all candidate entities are assumed to be present in the graph.

2.3 Constraints

We consider lightweight necessary constraints extracted from the question with respect to a known schema (this set of constraint types can be extended, e.g., to capture joins or cardinality constraints):

- Type constraints: $\text{type}(x, T)$
- Relation constraints: $R(x, y)$
- Exclusion constraints: $\neg R(x, y)$

2.4 Three-Valued Constraint Semantics

Due to incompleteness, constraints $c(x)$ for candidate answers x over graph G are evaluated under a three-valued semantics: $\llbracket c(x) \rrbracket_G \in \{\text{sat}, \text{viol}, \text{unk}\}$.

Relation constraints

$$\llbracket R(x, y) \rrbracket_G = \begin{cases} \text{sat} & \text{if } (x, R, y) \in E \\ \text{unk} & \text{otherwise} \end{cases}$$

Exclusion constraints

$$\llbracket \neg R(x, y) \rrbracket_G = \begin{cases} \text{viol} & \text{if } (x, R, y) \in E \\ \text{unk} & \text{otherwise} \end{cases}$$

Type constraints

$$\llbracket \text{type}(x, T) \rrbracket_G = \begin{cases} \text{sat} & \text{if } \text{type}(x) = T \\ \text{viol} & \text{otherwise} \end{cases}$$

2.5 Constraint-Based Filtering and Support-Based Verification

We retain only candidates that do not violate any constraint:

$$C^*(q) = \{x \in C(q) \mid \forall c \in \mathcal{C}(q) : \llbracket c(x) \rrbracket_G \neq \text{viol}\}.$$

This strategy relies solely on explicit violations and is therefore robust under incomplete knowledge.

Further, let

$$\mathcal{C}^+(q) = \{c \in \mathcal{C}(q) \mid \text{sat} \in \text{range}(c)\}$$

denote the subset of constraints that can provide positive evidence by being satisfied. The support of a candidate x is defined as the fraction of constraints in $\mathcal{C}^+(q)$ that evaluate to **sat**:

$$\text{support}(x, q) = \frac{|\{c \in \mathcal{C}^+(q) \mid \llbracket c(x) \rrbracket_G = \text{sat}\}|}{|\mathcal{C}^+(q)|}.$$

3 Experiments**3.1 Experimental Setup**

We evaluate the proposed constraint-based verification approach on entity-centric KGQA tasks using the Hetionet biomedical knowledge graph [7].² This evaluation is designed as a controlled study to isolate the effect of constraint-based verification independently of upstream candidate generation. Therefore, we simulate answer generation by sampling queries and gold answers from the Hetionet KG and then generate candidate answers from gold answers by systematically adding incorrect answers (false positives). Since our method only removes candidates with explicit constraint violations, recall is unaffected and we do not simulate false negatives.

We consider two query types that fit our constraints (these can be extended by more complex forms including more complex constraint types):

- Q1: Which compounds treat disease D ?

² Code and experimental artifacts: <https://github.com/ekitzelmann/ces-pk>

Table 1. Performance of constraint-based filtering and support-based verification.

Query	Prec. (before)	Prec. (after)	Δ	Recall	Supp. (gold)	Supp. (non-gold)
Q1	0.66	0.70	+0.04	1.00	1.00	0.42
Q2	0.17	0.62	+0.45	1.00	1.00	0.74

Table 2. Effect of simulated KG incompleteness (10% edge removal).

Query	Prec. (after)	Δ	Recall	Supp. (gold)	Supp. (non-gold)
Q1	0.70	+0.04	1.00	0.95	0.42
Q2	0.53	+0.37	1.00	0.95	0.69

- Q2: Which compounds treat disease D but do not cause side effect S ?

For each query type, we derive queries by sampling diseases and side effects from the KG and construct gold answer sets from the graph. We construct candidate sets that simulate LLM-generated answers by combining gold answers with sampled distractors, including type-incorrect entities and candidates that violate the exclusion constraint (for Q2).

Constraints include type (C1), relation (C2), and exclusion (C3) constraints. Evaluation is based on precision before and after filtering and support scores after filtering. Since all gold answers are included in the candidate sets, recall remains constant at 1.0. In the present experiments, $\mathcal{C}^+(q)$ consists of the type and positive relation constraints, such that $|\mathcal{C}^+(q)| = 2$ for both Q1 and Q2.

3.2 Results

Table 1 summarizes filtering performance and support-based verification evidence on the Hetionet KG. For Q1, constraint-based filtering yields a modest precision improvement (+0.04) due to removed type-inconsistent candidates. For Q2, we achieve a substantial precision increase from 0.17 to 0.62 (+0.45) due to exclusion constraints that eliminate candidates violating the negated condition. Overall, precision improves from 0.41 to 0.66 while recall remains unaffected.

The results show distinct roles of constraint types. Type constraints (C1) provide consistent filtering signals, while exclusion constraints (C3) are the primary driver of filtering in Q2. Relation constraints (C2) do not lead to violations under OWA but contribute to verification by providing support if satisfied.

Support scores provide additional discrimination between candidates. Gold answers achieve a mean support of 1.0 compared to 0.42 (Q1) and 0.74 (Q2) for non-gold candidates. This indicates that positive constraints contribute useful verification signals even when filtering is limited.

To further analyze the impact of KG incompleteness, in a second experiment we randomly removed 10% of the relevant *treats* and *causes* edges after constructing the gold answer sets. Results are shown in Table 2. Precision and

support scores decrease, as missing edges reduce the number of detectable violations and verifications, while the method continues to preserve recall.

3.3 Discussion

The results show that lightweight constraints improve precision without requiring full semantic parsing. Negative constraints provide filtering signals when explicit violations are present, while positive constraints contribute evidence for verification. The three-valued semantics preserves recall under incomplete knowledge.

Limitations. The evaluation is based on controlled candidate sets rather than capturing end-to-end KGQA performance. The KG is assumed to be factually correct but incomplete. In the presence of incorrect triples or missing entities, constraint evaluation may lead to false rejections and support.

4 Related Work

KGQA with LLMs [1,11,13] is approached in different ways: In semantic parsing approaches [2], an LLM translates natural language questions into executable logical forms [6,4]. Such methods enable precise and verifiable answers, but remain brittle in practice as generating correct queries over complex schemas is challenging. They typically involve fine-tuning, require complete KGs and leave LLM knowledge unexploited that could help to reason more robustly over incomplete KGs.

Another line of work uses LLMs for reasoning over knowledge graphs and generating answers. In [1], relevant triples are retrieved from the KG by entity linking which are then verbalized and injected as additional context into the prompt. Think-on-Graph [11] allows LLMs to iteratively explore graph neighborhoods and construct reasoning paths. These methods typically work without additional training and are more robust but can produce hallucinated or non-grounded answers. Moreover, incomplete KGs remain challenging. To address the issue of unreliable LLM-based reasoning, [9] introduce constraints constructed from the KG that guide the LLM decoding to generate well-formed reasoning chains. In contrast, the method proposed in this paper uses constraints to post-hoc verify answer candidates. In [5], integration of link prediction into LLM-based reasoning over KGs is suggested to address incomplete KGs.

Our work is also related to research on generation followed by verification [12], where it is shown that LLMs are effective at generating solution candidates but unreliable at verifying them.

5 Conclusions

We introduced Constrained Entity Selection under Partial Knowledge (CES-PK), a framework for post-hoc verification of candidate answers in knowledge

graph question answering. Instead of relying on full semantic parsing or unconstrained neural generation, our approach uses lightweight necessary constraints together with a three-valued semantics to filter invalid candidates and provide symbolic support for valid ones under incomplete knowledge.

Our experimental results demonstrate that constraint-based verification can substantially improve precision while preserving recall in a controlled setting. In particular, we show that exclusion constraints provide strong negative signals for filtering, whereas positive relation constraints contribute to verification through support.

The proposed approach assumes that the underlying knowledge graph is factually correct but incomplete and that all candidate entities are present in the graph. Furthermore, our evaluation is based on controlled candidate sets and does not yet capture end-to-end KGQA performance under noisy knowledge graphs.

Future work includes integrating the proposed verification layer into full KGQA pipelines, studying robustness under incomplete and noisy graphs, and extending constraint extraction to more complex query types and schemas.

References

1. Baek, J., Aji, A.F., Saffari, A.: Knowledge-Augmented Language Model Prompting for Zero-Shot Knowledge Graph Question Answering. In: Dalvi Mishra, B., Durrett, G., Jansen, P., Neves Ribeiro, D., Wei, J. (eds.) *Proceedings of the 1st Workshop on Natural Language Reasoning and Structured Explanations (NLRSE)*. pp. 78–106. Association for Computational Linguistics, Toronto, Canada (Jun 2023). <https://doi.org/10.18653/v1/2023.nlrse-1.7>
2. Berant, J., Chou, A., Frostig, R., Liang, P.: Semantic Parsing on Freebase from Question-Answer Pairs. In: Yarowsky, D., Baldwin, T., Korhonen, A., Livescu, K., Bethard, S. (eds.) *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. pp. 1533–1544. Association for Computational Linguistics, Seattle, Washington, USA (Oct 2013)
3. Bhandiwad, D., Gattogi, P., Kanger, A., Basaldella, M., Ferré, S., Vahdati, S., Lehmann, J.: Bridging language models and knowledge graphs with controlled natural languages. *Knowledge-Based Systems* **337**, 115405 (Mar 2026). <https://doi.org/10.1016/j.knosys.2026.115405>
4. D’Abramo, J., Zugarini, A., Torroni, P.: Investigating Large Language Models for Text-to-SPARQL Generation. In: Shi, W., Yu, W., Asai, A., Jiang, M., Durrett, G., Hajishirzi, H., Zettlemoyer, L. (eds.) *Proceedings of the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing*. pp. 66–80. Association for Computational Linguistics, Albuquerque, New Mexico, USA (May 2025). <https://doi.org/10.18653/v1/2025.knowledgenlp-1.5>
5. El Khatib, O.: Reasoning over Incomplete Knowledge Graphs. In: *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*. pp. 6785–6788. CIKM ’25, Association for Computing Machinery, New York, NY, USA (Nov 2025). <https://doi.org/10.1145/3746252.3761657>
6. Feng, T., He, L.: RGR-KBQA: Generating Logical Forms for Question Answering Using Knowledge-Graph-Enhanced Large Language Model. In: Rambow, O., Wanner, L., Apidianaki, M., Al-Khalifa, H., Eugenio, B.D., Schockaert, S. (eds.)

- Proceedings of the 31st International Conference on Computational Linguistics. pp. 3057–3070. Association for Computational Linguistics, Abu Dhabi, UAE (Jan 2025)
7. Himmelstein, D.S., Lizée, A., Hessler, C., Brueggeman, L., Chen, S.L., Hadley, D., Green, A., Khankhanian, P., Baranzini, S.E.: Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *eLife* **6**, e26726 (Sep 2017). <https://doi.org/10.7554/eLife.26726>
 8. Hogan, A., Blomqvist, E., Cochez, M., D’amato, C., Melo, G.D., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., Ngomo, A.C.N., Polleres, A., Rashid, S.M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., Zimmermann, A.: Knowledge Graphs. *ACM Comput. Surv.* **54**(4), 71:1–71:37 (Jul 2021). <https://doi.org/10.1145/3447772>
 9. Li, K., Zhang, T., Wu, X., Luo, H., Glass, J.R., Meng, H.M.: Decoding on Graphs: Faithful and Sound Reasoning on Knowledge Graphs through Generation of Well-Formed Chains. In: Che, W., Nabende, J., Shutova, E., Pilehvar, M.T. (eds.) *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pp. 24349–24364. Association for Computational Linguistics, Vienna, Austria (Jul 2025). <https://doi.org/10.18653/v1/2025.acl-long.1186>
 10. Ma, C., Chen, Y., Wu, T., Khan, A., Wang, H.: Large Language Models Meet Knowledge Graphs for Question Answering: Synthesis and Opportunities. In: Christodoulopoulos, C., Chakraborty, T., Rose, C., Peng, V. (eds.) *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. pp. 24578–24597. Association for Computational Linguistics, Suzhou, China (Nov 2025). <https://doi.org/10.18653/v1/2025.emnlp-main.1249>
 11. Sun, J., Xu, C., Tang, L., Wang, S., Lin, C., Gong, Y., Ni, L., Shum, H.Y., Guo, J.: Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph. In: *The Twelfth International Conference on Learning Representations* (Oct 2023)
 12. Valmeekam, K., Marquez, M., Sreedharan, S., Kambhampati, S.: On the Planning Abilities of Large Language Models - A Critical Investigation. In: *Advances in Neural Information Processing Systems*. vol. 36, pp. 75993–76005 (Dec 2023)
 13. Wagner, R., Kitzelmann, E., Boersch, I.: Mitigating Hallucination by Integrating Knowledge Graphs into LLM Inference – a Systematic Literature Review. In: Zhao, J., Wang, M., Liu, Z. (eds.) *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*. pp. 795–805. Association for Computational Linguistics, Vienna, Austria (Jul 2025). <https://doi.org/10.18653/v1/2025.acl-srw.53>