

## Responsible NLP Checklist

Paper title: *From Query to Logic: Ontology-Driven Multi-Hop Reasoning in LLMs*

Authors: *Haonan Bian, Yutao Qi, Rui Yang, Yuanxi Che, Jiaqian Wang, Heming Xia, Ranran Zhen*

How to read the checklist symbols:

- the authors responded 'yes'
- the authors responded 'no'
- the authors indicated that the question does not apply to their work
- the authors did not respond to the checkbox question

For background on the checklist and guidance provided to the authors, see the [Responsible NLP Checklist](#) page at ACL Rolling Review.

---

### A. Questions mandatory for all submissions.

- A1. Did you describe the limitations of your work?

*This paper has a Limitations section.*

- A2. Did you discuss any potential risks of your work?

*This work presents a training-free reasoning framework using public datasets and does not introduce significant societal risks.*

### B. Did you use or create scientific artifacts? (e.g. code, datasets, models)

- B4. Did you discuss the steps taken to check whether the data that was collected/used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect/anonymize it?

*We use publicly available benchmark datasets (HotPotQA, 2WikiMQA, MuSiQue) that do not contain personally identifying information or offensive content.*

- B6. Did you report relevant statistics like the number of examples, details of train/test/dev splits, etc. for the data that you used/created?

*We report dataset statistics including the number of examples for each benchmark (HotPotQA, 2WikiMQA, MuSiQue) in Section 5.1.1 and Appendix A.2.1.*

### C. Did you run computational experiments?

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*The experimental setup, including model configurations and implementation details, is described in Section 5.1.3 and Appendix A.2.2 and A.2.3.*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*We report main results as single-run performance scores (EM and F1) on standard benchmarks in Section 5.2 (Table 1), with ablation results in Section 6.1 (Table 2) and cross-model results in Section 6.4 (Table 3).*

*The Responsible NLP Checklist used at ACL Rolling Review is adopted from NAACL 2022, with the addition of ACL 2023 question on AI writing assistance and further refinements based on ARR practice. ACL 2026 used a subset of ARR checklist form.*

**D. Did you use human annotators (e.g., crowdworkers) or research with human subjects?**

D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*This work does not involve human annotators or participants.*

D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*(left blank)*

D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating (e.g., did your instructions explain how the data would be used)?

*(left blank)*

D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*(left blank)*

**E. Did you use AI assistants (e.g., ChatGPT, Copilot) in your research, coding, or writing?**

E1. If you used AI assistants, did you include information about their use?

*We used AI assistants for writing polish and coding assistance, but did not include a dedicated statement in the paper.*