

Responsible NLP Checklist

Paper title: *Graph Explorer: Training Faithful KG Agents with Visibility-Grounded Supervision*

Authors: *Yifeng Chen, Sicheng Wan, Tianyi Zhang, Xuezhou Zhang*

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?

Not applicable. This work does not involve risks requiring elaboration.

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?

No personally identifying or offensive data is used. All experiments are conducted on publicly available knowledge graph datasets (CWQ, WebQSP).

- 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?

Section 5.1 (Experimental Setups) and Table 1 report dataset statistics, including dataset sizes, splits, and filtering coverage.

C. Did you run computational experiments?

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

Section 5.1 and Appendix D describe the experimental setup, including model configuration, training procedure, and hyperparameters.

- 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?

Section 5.3 and Tables 23 report evaluation results, including Hit@1 under both best-effort and finish-or-fail protocols.

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.?

No human participants or annotators were involved in this work.

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)?

No human participants were recruited or compensated.

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)?

No human data requiring consent was used.

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

This work does not involve human subjects and does not require ethics review board approval.

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?

AI assistants (e.g., ChatGPT) were used for minor writing assistance, including grammar checking and phrasing improvements. All technical content, experimental design, and results were produced and verified by the authors.