

## Responsible NLP Checklist

Paper title: *Enhancing LLM-based Search Agents via Contribution Weighted Group Relative Policy Optimization*

Authors: *Junzhe Wang, Zhiheng Xi, yajie yang, Hao Luo, Shihan Dou, Tao Gui, Qi Zhang*

How to read the checklist symbols:

- the authors responded 'yes'
- the authors responded 'no'
- <sup>N/A</sup> 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.*

<sup>N/A</sup> A2. Did you discuss any potential risks of your work?  
*(left blank)*

### 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 only use general and multi-hop question-and-answer datasets, and there is no sensitive information.*

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 4.1 and 4.3*

### C. Did you run computational experiments?

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?  
*Section 4 and 5.2*

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?  
*Due to significant computational requirements, we report results from single-run training experiments using a fixed seed. Consequently, descriptive statistics such as error bars are not provided, as is common practice in LLM agent training.*

---

*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.?

*We did not provide a formal written instruction protocol, as the annotations were conducted by a single author solely for validating LLM judge accuracy on a small dataset.*

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 external annotators were recruited or paid.*

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)?

*Annotation was performed solely by a single author; no external participants data were collected, so separate consent procedures were not applicable.*

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

*No ethics review was required. The annotations were performed solely by a single author, involved no external participants, collected no personal or sensitive information, and posed no participant risk.*

**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?

*Acknowledgements*