

Responsible NLP Checklist

Paper title: *LoRE: Enhancing Search Relevance with Progressive Chain-of-Thought and Preference Alignment*

Authors: *Chenji Lu, Zhuo Chen, Hui Zhao, Zhiyuan Zeng, Gang Zhao, Junjie Ren, Lihaoran, Songyan Liu, Pengjie Wang, Chuan Yu, Jian Xu, Bo Zheng*

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?

No. All experiments were conducted offline using publicly available open-source datasets, presenting no compliance or privacy 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?

no. The datasets we use contain no personally identifiable information (PII) 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?

Yes. Section 5 reports the statistics of the benchmark datasets used in our experiments. Additional details on the training and validation sets are provided in Appendix C.

C. Did you run computational experiments?

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

no. Our experiments introduce no additional hyperparameters beyond basic ones such as learning rate.

- 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?

Yes. Section 5 reports all statistical results and explicitly states that they are from single runs of the experiments.

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. No human subjects were involved in this study.

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. Our data was not annotated using crowdsourcing platforms or student annotators.

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. We use publicly available benchmark and cite them in the paper.

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

no. Our research uses only publicly available datasets and does not involve human subjects, thus raising no ethical concerns.

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?

GPT-5 was used only for text polishing during the writing phase.