

Responsible NLP Checklist

Paper title: *Mind Reader: Latent User Demand-Guided Content Optimization for Generative Search Engine*

Authors: *Tong Chen, JiaWei Guo, Yuxi Li, Baiming Chen, Houxing Ren, Zhang Zhiwei, Yunxiang Zhang, Hanyang Xia, Kun Liang, Zhaoran Fan*

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. Given that our work neither collects nor processes identifiable or private information, and operates with no capacity for autonomous decision-making or user manipulation, we think that it presents no societal, ethical, or safety 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 only consist of public web content and do not contain real user data, personal identifiers, or user-generated content, and therefore pose no privacy or safety risks.

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 B.1

C. Did you run computational experiments?

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

Section A.1

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 4.1

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.?

(left blank)

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)?

Ethics Statement

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

Ethics Statement

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 only for proofreading, grammar checking, and improving clarity of exposition, but not for generating core scientific content, experimental design, data analysis, or result interpretation. The final manuscript was fully reviewed and validated by the human authors, who take full responsibility for all technical claims and reported results.