

Responsible NLP Checklist

Paper title: *Reasoning While Asking: Transforming Reasoning Large Language Models from Passive Solvers to Proactive Inquirers*

Authors: *Xin Chen, Feng Jiang, Yiqian Zhang, Hardy Chen, Shuo Yan, Wenya Xie, Min Yang, Shujian Huang*

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?

The paper discuss broader societal risks or ethical concerns in Ethical Considerations Section.

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?

Ethical Considerations Section. The data used includes MathChat, BigCodeBench-Chat, and DocEdit-Chat, which are academic dataset and do not typically contain PII or offensive conten

- 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?

See Appendix A.2 provides statistics for the Reasoning-While-Asking SFT Dataset, Math-CHAT, BigCodeBench-Chat, and MediumDocEdit-Chat (Train/Test splits).

C. Did you run computational experiments?

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

See Section 4.1 and Appendix A.2 which lists hyperparameters for SFT and US-GRPO (Batch size, LR, Epochs, GPU, etc.).

- 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 the high computational costs associated with reinforcement learning training for Large Language Models, we performed single runs for our main experiments and reported the results on the standard test sets. For the training dynamics analysis (e.g., Figure 6), we utilized exponential moving averages to visualize the trends and stability.

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.?

The user study involved straightforward preference comparison between model outputs. Instructions were minimal and did not involve any risk to 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)?

Not applicable. The data was provided by co-authors or members of the research group, and no external 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)?

All participants were informed that their annotations would be used for research evaluation purposes. As the study only involved preference judgments on model-generated outputs with no personal or sensitive data collected, formal written consent was not obtained

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

The study involved only preference comparisons of model outputs by volunteer graduate students, with no collection of personal or sensitive data and no risk to participants. Formal ethics review board approval was not required.

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?

Ethical Considerations Section. We used AI assistants for minor text editing and proofreading.