

Responsible NLP Checklist

Paper title: *RiT: Rubrics-in-Thinking Reinforcement Learning for Improved Reasoning in Large Language Models*

Authors: *Xiaobin Tian, Shuai Yuan, Muyun Ding, Haonan Chen, xiaoxi jiang*

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?

Our work is foundational research. In this paper, we aim to use thinking-rubric-guided reinforcement learning to improve reasoning in large language models. We do not anticipate any potential risks arising from our work.

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?

The datasets used in this work do not contain any personally identifiable information (PII) or offensive content. All data are publicly available benchmark datasets or synthetic data, and no private or sensitive information about individuals is included. We did not need to perform additional anonymization steps, as the data are already sanitized and curated for research purposes.

- 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?

Relevant statistics, including the number of examples and details of train/dev/test splits, are provided in Appendix A.2.

C. Did you run computational experiments?

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

The experimental setup is described in Section 5.1, Appendix A.1 and Appendix A.2. While we did not perform a dedicated hyperparameter search, our choice of hyperparameters is discussed in Appendix A.4.

- 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?

Descriptive statistics for our results are reported in Section 5.1 and Appendix A.3.

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.?

This work does not involve human participants or annotators. Therefore, no instructions or disclaimers were provided.

- 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)?

This work does not involve human participants or annotators. Therefore, no recruitment or payment was applicable.

- 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)?

This work uses only publicly available datasets or anonymized data. No additional consent was required from individuals, as all data are already curated for research purposes.

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

This work does not involve the collection of new data from human participants. As only publicly available datasets were used, no ethics review board approval was 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?

We report the use of AI assistants in Sections 5.1, 5.5, and Appendix A.1, where they serve as judge models for both training supervision and evaluation.