

Responsible NLP Checklist

Paper title: *Harnessing Negative Signals: Reinforcement Distillation from Teacher Data for LLM Reasoning*

Authors: *Shuyao Xu, Cheng Peng, Jiangxuan Long, Weidi Xu, Wei Chu, Yuan Qi*

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?

Section I (Potential Harms). the appendix explicitly discusses misuse risks, bias propagation, and risks in high-stakes deployment.

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 paper does not describe a dedicated procedure for screening the data for personally identifying information or offensive content, nor an anonymization pipeline. The work uses public math/problem-solving datasets and model-generated reasoning traces rather than human personal data, so this was not a central concern of the study.

- 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?

Sections 3.1 and 3.3. Reason: the paper reports dataset sizes such as 78k positive pairs and 53k preference triplets, and describes evaluation protocol/sample counts.

C. Did you run computational experiments?

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

Sections 3.2, 3.3, and Appendix C.4. Reason: the paper describes training stages, optimizer, batch sizes, learning-rate schedule, decoding setup, and says detailed hyperparameter ranges/chosen values are in Appendix C.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?

Section 3.3, Tables 1 and 2, and Appendix C.5. Reason: the paper reports pass@1/pass@16, states how pass@k is computed, includes SEM, and tables show means with error bars such as values.

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)?

(left blank)

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

(left blank)

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?

Section H (Statement on the Use of AI Assistants). Reason: that section states the assistants were used only for grammar/style/clarity and not for the core content, analysis, or conclusions.