

Responsible NLP Checklist

Paper title: *Verifier-Free RL for LLMs via Intrinsic Gradient-Norm Reward*

Authors: *Xuexiang Wen, Hang Yu, Linchao Zhu, Gaoang Wang*

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 work does not involve any potentially harmful applications, deployment risks, or data that could cause ethical concerns. The experiments are limited to benchmarking training methods for NLP models on standard public datasets.

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 are standard publicly available NLP benchmark datasets. They do not contain any personally identifiable information or offensive content, and our work does not involve collecting or using any sensitive personal data.

- 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 (Training Data) and Appendix B (Table 10) report dataset statistics including the number of training problems and evaluation examples for each benchmark.

C. Did you run computational experiments?

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

We did not perform hyperparameter search for this work, as we leverage established default settings from prior work.

- 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?

Appendix B (Evaluation Setting) reports how results are computed: pass@1 with greedy decoding (temperature=0) for most benchmarks, and avg@8 (average over 8 independent runs) for AMC to reduce variance given its smaller size (83 problems).

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 any annotation tasks requiring instructions. All experiments use standard benchmark datasets and automatic evaluation.

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 human participants were recruited or compensated in this study. The work uses only publicly available benchmarks and automatic evaluation metrics.

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 benchmark datasets that do not involve personal data collected from individuals; therefore, data consent from subjects is not applicable.

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

This work does not involve any data collection from human participants or sensitive personal data; therefore, no ethics review board approval is applicable.

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 ChatGPT during the writing process to assist with wording and phrasing, but this usage was not included in the manuscript itself. We disclose it here for transparency. All content generated was reviewed and edited by authors.