Responsible NLP Checklist

Paper title: *G^2RPO-A: Guided Group Relative Policy Optimization with Adaptive Guidance*
Authors: *Yongxin Guo, Wenbo Deng, Zhenglin Cheng, Xiaoying Tang*

How to read the checklist symbols:

☑ the authors responded 'yes'

☒ the authors responded 'no'

N/A 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?
We acknowledge that, in a broad sense, the research on language models does carry potential social risks, such as being maliciously used to generate misleading information or amplifying the inherent biases in the training data. However, we believe that our work does not introduce new risks, and the level of existing risks is extremely low. Our contribution is essentially At the fundamental algorithmic level, it aims to improve the training methods of small language models (SLMs) in specific domains such as mathematics and code generation. Since our research does not involve user-oriented deployment or handling sensitive data, the potential risks are indirect. Moreover, we focus on enhancing the performance of smaller and more accessible models, which may have a positive impact on society by lowering the threshold and reducing resource consumption of AI technology.

☑ **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?
A discussion on this topic was omitted as the datasets used consist of mathematical problems and coding challenges from the curated Open-R1 collection. By their nature, these datasets are not expected to contain personally identifiable information (PII) or offensive content, making an explicit anonymization or filtering process unnecessary.

☑ 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?
4.4

☑ **C. Did you run computational experiments?**

☑ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
4.2, 4.4

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.

☒ 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?
We report point estimates from a single experimental run for each result, rather than descriptive statistics (e.g., mean/std over multiple runs). This decision was made due to the substantial computational resources required to conduct multiple full training and evaluation cycles for each experiment. We have provided a detailed description of the experimental setup and hyperparameters in Section 4.4 to ensure the reproducibility of our reported results.

☒ **D. Did you use human annotators (e.g., crowdworkers) or research with human subjects?**

N/A 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 question is not applicable as our research did not involve human participants or annotators. All experiments were conducted using pre-existing, publicly available datasets.

N/A 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 question is not applicable as our research did not involve human participants or annotators. All experiments were conducted using pre-existing, publicly available datasets.

N/A 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 question is not applicable. Our research utilizes publicly available datasets of mathematical and coding problems, not personal data collected from human participants. Therefore, the concept of obtaining data consent from individuals is not relevant to our work.

N/A D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
This question is not applicable. Our research did not involve human subjects or any new data collection from individuals, which is the type of work that requires ethics review board approval. All our experiments were conducted on pre-existing, publicly available technical datasets.

☑ **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 did not include a statement in the paper acknowledging the use of AI assistants. We confirm that AI-based tools were utilized to assist with improving the grammar, clarity, and style of the language in the manuscript. The role of these tools was limited to that of a writing aid, similar to a grammar checker. All the scientific contributions, including the methodology, experiments, analysis, and conclusions, were solely the work of the authors. The authors have reviewed, edited, and take full responsibility for the final content of the paper.