

## Responsible NLP Checklist

Paper title: *CSRP: Chain-of-Thought Reasoning for Chinese Text Correction via Reinforcement Learning with Efficiency-Aware Rewards*

Authors: *Wei Tian, Yuhao Zhou, Man Lan*

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 6 (Limitations)*

### 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 (Lang8, HSK, NACGEC, CSCD, etc.) are established public benchmarks for Chinese text correction and do not contain personally identifying information.*

- 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*

### C. Did you run computational experiments?

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

*Section 4.1 and Appendix D*

- 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 cost of the three-stage training pipeline (CPT + SFT + GRPO), each configuration was run once. Reporting multiple runs with error bars was not feasible within our resource constraints. Results are from a single run.*

*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.?

*Appendix E*

- 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)?

*The three annotators were domain experts from our research group who participated voluntarily. No external payment was involved as the annotation was conducted as part of the research collaboration.*

- 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 datasets used are publicly available benchmarks. No new data was collected from human subjects.*

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

*No new data collection from human subjects was conducted. The human annotation study involved internal domain experts evaluating model-generated rationales, which does not require ethics board approval.*

**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?

*Qwen-Plus was used as a teacher model to distill Chain-of-Thought rationales for the SFT stage, as described in Section 3.2. This use is central to our methodology and fully disclosed in the paper.*