

Responsible NLP Checklist

Paper title: *Self-SoftCoT: A Self-Consistent Framework via Position-Aware Latent Space Reinforcement Learning*

Authors: *Liangliang Dong, Lianlei Shan, Shuaimin Li*

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?

Limitations section. The paper discusses practical reliability and deployment risks, including reduced auditability and the need for caution in high-stakes settings.

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?

We used established benchmark datasets and did not collect or curate 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?

Appendix C.1, Table 6.

C. Did you run computational experiments?

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

Section 4.3 and Appendix D.1D.3, including Tables 7 and 8.

- 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 5.15.3, Tables 15. Results are reported with mean standard deviation, and some evaluations explicitly use 5 random seeds.

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.?

No human participants or annotators were involved.

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.

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 or annotators were involved.

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)?

No human-subject data were collected by the authors.

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

No human-subject data collection or participant study was conducted.

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 have disclosed this in the "Acknowledgments" section on page 9. We only used generative AI tools for limited language assistance, including minor wording refinement and polishing of the authors' original text. All technical content, experimental design, results, and final writing decisions were verified and determined by the authors.