

Responsible NLP Checklist

Paper title: *Act-Adaptive Margin: Dynamically Calibrating Reward Models for Subjective Ambiguity*

Authors: *Feiteng Fang, Dingwei Chen, Xiang Huang, Ting-En Lin, Yuchuan Wu, Xiong Liu, Jing Ye, Ziqiang Liu, Haonan Zhang, Liang Zhu, Hamid Alinejad-Rokny, Min Yang, Yongbin Li*

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?
There are no potential risks of this work.

B. Did you use or create scientific artifacts? (e.g. code, datasets, models)

^{N/A} 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?
There are no personally identifying info or offensive content.

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 5

C. Did you run computational experiments?

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 5

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 5.1 and Section 5.2

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.?
We described the key evaluation criteria in the Appendix A.

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)?

We recruited annotators via crowdsourcing and they were compensated. However, we did not disclose specific payment rates in the paper due to privacy considerations.

- 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)?

Participants implied consent by voluntarily accepting the annotation task, which included a description of the data usage. Formal written consent forms were not included in the paper text.

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

This research does not involve clinical trials, or private data collection. It relies solely on public benchmarks and synthetic data, so ethics review board approval is not 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 utilized GPT-4o for generating synthetic narrative data as described in Section 5. Additionally, we used AI assistants to help with grammatical error checking and polishing of the manuscript.