

## Responsible NLP Checklist

Paper title: *Learning Temporally-Aware Sample Weights for Preference Optimization*

Authors: *Mengyang Li, Xudong Zhou, Pinlong Zhao*

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?

*This work focuses on sample weighting for LLM alignment and does not introduce new risks beyond those of standard preference optimization methods.*

### 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 use publicly available datasets (UltraFeedback, HH-RLHF, Llama-3-UltraFeedback-ArmoRM) that have been previously released and vetted by their original authors.*

- 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 reports dataset sizes (64k, 61k, 161k samples) and validation set construction (3% stratified split).*

### 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 A describe hyperparameters including EMA decay =0.9, =0.1, K=100, and all hardware/library details.*

- 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?

*All results in Tables 15 report meanstd over 3 runs with different seeds; statistical significance verified via paired t-tests ( $p < 0.05$ ).*

*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.?

*No human subjects or annotators were involved in this study.*

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 subjects or annotators were involved in this study.*

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 subjects or annotators were involved in this study.*

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

*No human subjects or annotators were involved in this study.*

**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?

*Disclosed in the Generative AI Disclosure statement at the end of Section 5: Claude was used to improve language and grammar; all content was reviewed and edited by the authors.*