

Responsible NLP Checklist

Paper title: *Min-k Sampling: Decoupling Truncation from Temperature Scaling via Relative Logit Dynamics*

Authors: *Yuanhao Ding, Meimingwei Li, Esteban Garces Arias, Matthias Aenmacher, Christian Heumann, Chongsheng ZHANG*

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?

Discussed in the Ethics Statement (Section 7 / page 9)

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?

Addressed in the Ethics Statement (Section 7 / page 9). We explicitly state that all datasets are publicly available and do not contain personally identifiable 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?

Reported in Section 4 (Experimental Setup), where dataset sizes and compositions are explicitly specified (e.g., AQuA: 254, GSM8K: 1319, etc.).

C. Did you run computational experiments?

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

Detailed in Section 4 (Experimental Setup), including model configurations, baseline methods, and hyperparameter settings (e.g., Top-p = 0.9, = 3.0).

- 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?

Reported in Section 5 (Results), including aggregate performance metrics (Tables 13), as well as statistical analyses such as binomial tests and inter-rater agreement (Fleiss).

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.?

Included in the Appendix

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)?

Reported in Section 5.3 (Human Evaluation) and the Ethics Statement (page 9), including number of annotators and compensation (\$20/hour).

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)?

(left blank)

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

(left blank)

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?

(left blank)