

Responsible NLP Checklist

Paper title: *CrowdSelect: Synthetic Instruction Data Selection with Multi-LLM Wisdom*

Authors: *Yisen Li, Lingfeng Yang, Wenxuan Shen, Pan Zhou, Yao Wan, Weiwei Lin, Dongping Chen*

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?

Yes, in Limitation

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

- B1. Did you cite the creators of artifacts you used?

Yes, in Reference

- B2. Did you discuss the license or terms for use and/or distribution of any artifacts?

No, I did not explicitly discuss the license or terms for use/distribution of the artifacts in the paper. However, all artifacts (including code, data, etc.) associated with this work follow the standard licensing agreements commonly adopted by the relevant research community (e.g., the MIT License, which is widely used in the field for open-access resources)

- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?

Yes, in 4. Experiment

- 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?

No, because my explanation: The current paper presents a selection method based on the Magpie dataset, and the Magpie dataset itself does not contain any personal identifiable information.

- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?

Yes, in Reference

- 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?

Yes, in 4, Experiment

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.

C. Did you run computational experiments?

C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

Yes, in appendix B

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

Yes, in 4, Experiment

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?

Yes, in 4, Experiment

C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation, such as NLTK, SpaCy, ROUGE, etc.), did you report the implementation, model, and parameter settings used?

No, This detail is not elaborated on in the main body of the paper due to space constraints, but it is explicitly specified and documented in the open-source code repository associated with this work.

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.?

in appendix C.9

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)?

Due to cost constraints (e.g., funding and time limitations for recruiting and compensating a large-scale human participant pool), the main experimental conclusions of this study are primarily derived from LLM-as-a-Judge evaluations

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 artifacts (including code, data, etc.) associated with this work follow the standard licensing agreements commonly adopted by the relevant research community (e.g., the MIT License, which is widely used in the field for open-access resources)

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

All artifacts (including code, data, etc.) associated with this work follow the standard licensing agreements commonly adopted by the relevant research community (e.g., the MIT License, which is widely used in the field for open-access resources)

D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Due to cost constraints (e.g., funding and time limitations for recruiting and compensating a large-scale human participant pool), the main experimental conclusions of this study are primarily derived from LLM-as-a-Judge evaluations

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 only used AI writing assistants for paper polishing and formatting checks, which complies with community guidelines.