

Responsible NLP Checklist

Paper title: *Do Morals Guide How LLMs Think? The Role of Ethical Perspectives in General Problem Solving*

Authors: *Iseo Kim, Eunjin Hong, Juae Kim*

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?

Limitations

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 verified that all datasets used in our study (MMLU and DIT) do not contain any personally identifiable information or offensive content. Both datasets were either publicly released for research or obtained through formal purchase for academic use. Although these checks were performed during the research process, this information is not explicitly stated in the paper.

- 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 3, Section 4, and Appendix G

C. Did you run computational experiments?

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

Appendix C.

- ^{N/A} 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 3, Section 4, Appendices C, D, F, and G.

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.?

N/A. No human participants or annotators were involved in this research. All experiments used existing datasets.

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)?

N/A. No participants were recruited or paid. The study relies exclusively on publicly available datasets, with no new human data collection.

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)?

N/A. The study uses only publicly available and licensed datasets (MMLU and DIT). We did not collect any new data from individuals, and the original data sources ensured consent as part of their data collection processes.

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

N/A. We did not collect any new data involving human subjects. All datasets used were either publicly available or licensed, and the original data providers are responsible for any ethical approvals related to their collection.

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?

Appendix I