

Responsible NLP Checklist

Paper title: *LLEOT: A Privacy-Enhancing Offsite Tuning Framework via Loss Landscape Elevation*

Authors: *Jin Zhong, Jinglin Liang, Tongtong Yang, Zijian Xie, Shuangping Huang, Hanlin Gu*

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 proposes a fine-tuning framework and does not introduce any additional potential risks beyond those inherent in the base models.

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?

*This work exclusively utilizes well-established, publicly available datasets (*The Pile-uncopyrighted*, *WikiText-2*, *OpenBookQA*, *ARC-Challenge*, *Social IQA*, and *WebQuestions*). We rely on the data curation and anonymization procedures described in the original publications of these benchmarks.*

- 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?

Appendix A.3

C. Did you run computational experiments?

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

Appendix A.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?

Appendix A.5

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 annotators or participants were involved in this work.

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)?

No human annotators or participants were involved in this work.

- 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)?

This work utilizes existing public datasets and did not involve data collection from human subjects requiring new informed consent.

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

As the study relies solely on publicly available datasets and does not involve human subjects, ethics review board approval was not required.

- 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 C