

Responsible NLP Checklist

Paper title: *SOS-LoRA: Static Orthogonal-Subspace Low-Rank Adaptation with Fixed Multi-Scale Scaling*

Authors: *Yupeng Chang, Yuan Wu, Yi Chang*

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?

No. This paper studies a parameter-efficient adaptation method and does not involve new data collection, human-subject experiments, or high-risk downstream deployment. General limitations are discussed in the Limitations section.

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?

N/A. We do not collect or create new datasets. We only use widely used public benchmarks and pretrained checkpoints released by their original providers.

- 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?

We describe the datasets, official splits, and evaluation metrics in Appendix A.4, and report key experimental settings and protocols in Section 4 and Appendix A.3.

C. Did you run computational experiments?

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

The experimental setup and key hyperparameters are reported in Section 4 and Appendix A.3, including adaptation scope, ranks, number of experts, scaling, regularization, learning rates, batch sizes, sequence lengths, epochs, and optimizer settings.

- 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?

Unless otherwise stated, results are reported as mean standard deviation over three independent runs with different random seeds; see Section 4 and Appendix A.3.

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.?

This work does not involve human participants or annotators.

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)?

This work does not involve recruiting or paying participants or annotators.

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)?

We do not collect or curate new human data in this work.

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

This work does not involve human-subject data collection or experimentation requiring ethics review board approval.

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 include an LLM Usage Statement in Appendix A.10, stating that an LLM was used only for language polishing and proofreading, not for method design, algorithm development, experimental results, or scientific decisions.