

Responsible NLP Checklist

Paper title: *MoE Adapter for Large Audio Language Models: Sparsity, Disentanglement, and Gradient-Conflict-Free*

Authors: *Yishu Lei, Shuwei He, Hu Jing, Dan Zhang, Xianlong Luo, Danxiang Zhu, shikun feng, Rui Liu, Jingzhou HE, Yu Sun, Hua Wu, Haifeng Wang*

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 is pure fundamental technical research focusing on resolving gradient conflicts in adapters for large audio language models. No specific productization or deployment scenarios are involved. All experiments are conducted on publicly available academic benchmarks, and no direct malicious application risks have been identified. Potential risks and mitigation strategies in real-world applications will be further explored in future work.

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 study exclusively uses publicly available academic benchmark datasets (MMAU, MMSU, OBQA). All datasets have been fully cleaned, anonymized, and content-audited by their original authors, and do not contain any personally identifiable information (PII) or inappropriate content.

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?

All datasets used in this study are standard public benchmarks in the NLP and audio domains. Their detailed statistics (including train/validation/test splits, sample sizes, and data distributions) have been fully disclosed in the original dataset papers, which have been properly cited in this work.

C. Did you run computational experiments?

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

Section 4.1 and Appendix 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,

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.

etc. or just a single run?

All experiments fixed random seeds to ensure full reproducibility, and the results reported in the paper are deterministic results from a single run. To further verify the stability of the results, we have supplemented the mean and standard deviation of 3 runs with different random seeds in the Appendix.

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 study does not involve any human subjects or human annotation work. Not applicable.

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 study does not involve any human subjects or human annotation work. Not applicable.

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 study does not involve any human subjects or human annotation work. Not applicable.

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

This study does not involve any human subjects or human annotation work. Not applicable.

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?

In this study, we used Claude to assist in writing partial code comments and repetitive code snippets, and Gemini for English grammar checking and proofreading. All AI-generated content has been carefully reviewed, revised, and verified by all authors to ensure its accuracy, originality, and academic rigor. AI assistants did not participate in the formulation of core ideas, experimental design, or result analysis of the paper.