

Responsible NLP Checklist

Paper title: *HiSVD: Principled Low-Rank Approximation of LLMs via Hierarchical Modeling of Information Capacity and Spectral Structure*

Authors: *Zhuo Chen, Minghao Li, Xiaoqian Ma, Siqi Fan, Xiusheng Huang, ZhangLiujie, Weihang 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.

- N/A* A2. Did you discuss any potential risks of your work?

Not applicable (N/A). This work focuses on theoretical model compression of existing open-source models and does not introduce new applications or datasets that pose specific societal risks.

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?

No. We use standard, publicly available benchmarks (e.g., WikiText-2, C4) and did not collect any new data containing PII.

- 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?

See Section 4 (Experiments and Analysis), where we detail the datasets and evaluation metrics used.

C. Did you run computational experiments?

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

See Section 4 (Experiments and Analysis) and Appendix A, where we detail the model architectures, hyperparameters, and hardware configurations.

- 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?

We report results based on single runs due to the deterministic nature of SVD and the computational cost of LLMs. See Section 4.

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 annotators or participants; no instructions to report.

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 recruitment or payment of participants/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)?

N/A: We use public datasets only; we do not curate data from identifiable individuals requiring a new consent discussion for this work.

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

N/A: No new human-subjects data collection protocol; thus no IRB/ethics approval applies to this study's data pipeline.

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?

Yes. We have included a declaration in Appendix A.6 (AI Assistant Declaration). We used AI assistants solely for grammatical error correction and polishing the text; no new scientific content was generated.