

Responsible NLP Checklist

Paper title: *WeightLoRA: Keep Only Necessary Adapters*

Authors: *Andrey Veprikov, Vladimir Solodkin, Zyl Alexander, Andrey Savchenko, Aleksandr Beznosikov*

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?

Potential risks are discussed in Ethical Considerations. We state that the proposed method does not introduce new risks beyond those associated with the underlying pretrained models and standard benchmarks, and we note that improved parameter efficiency may lower the barrier to deployment, which can increase misuse potential.

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 do not use sensitive personal data. All evaluations are conducted on established, publicly available benchmarks, and any data-related considerations are inherited from these existing resources. See Ethical Considerations and Appendix E for datasets descriptions.

- 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 use standard benchmarks with their commonly adopted train/dev/test splits. For all experiments, we report results averaged over five random seeds to reduce variance and improve robustness. Details of the experimental setup are provided in Section 3 and Appendix E.

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 hyperparameters are described in Section 3, with additional implementation and training details provided in Appendix E. This includes model architectures, adapter configurations, training schedules, optimizer settings, and the values of key hyperparameters used in all experiments.

- 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?

We report results averaged over five random seeds to account for variance across runs. This is stated explicitly in the experimental section, and all reported metrics correspond to mean performance across seeds rather than single runs. See captions for all tables in Section 3 and Appendix E.

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.?

The study does not involve human participants or annotators. All experiments are conducted using existing pretrained models and publicly available benchmark 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)?

No human participants were recruited or compensated, as the work does not involve human subjects or annotation.

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)?

The paper does not involve the collection or curation of new human data. All datasets used are established public benchmarks with existing consent and usage terms determined by their original creators.

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

Ethics review board approval was not required, as the study does not involve human subjects, human annotation, or the collection of new personal data.

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?

The use of AI assistants is disclosed in Appendix F. AI-based tools were used to improve the clarity and presentation of the manuscript text and to assist with code development. All research ideas, experimental design, results, and conclusions are the responsibility of the authors.