Responsible NLP Checklist

Paper title: *Small Data, Big Noise: Adversarial Training for Robust ParameterEfficient Fine-Tuning*
Authors: *Eitan Cohen, Idan Simai, Uri Shaham*

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?
No. This work focuses on methodological advances in robust parameter-efficient fine-tuning and does not involve deployment, human subjects, or release of harmful content. We do not identify direct or immediate risks beyond the standard dual-use considerations that apply broadly to NLP robustness research.

☒ 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 used existing public benchmark datasets and did not collect new data. The paper does not include a dedicated discussion of whether these datasets contain personally identifying or offensive content, nor additional anonymization steps beyond the original dataset releases.

☑ 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?
Yes Section 5. We report the datasets used and relevant experimental data settings, including the benchmark datasets, training-data regimes (5% to 100% of the training data), and concrete subset sizes used in experiments (e.g., 1,000 samples for Banking77 and 200 samples for SQuAD/TweetQA).

☒ C. Did you run computational experiments?

☑ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Yes Section 5 and Appendix C/D. We describe the experimental setup, including the models, datasets, PEFT methods, training protocol (e.g., warm-up and adversarial-training epochs), perturbation settings, and the selected hyperparameter values such as the perturbation radius . Additional implementation and hyperparameter details are provided in the appendix.

☑ 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?
Yes Section 5 and Appendix C/D. We report descriptive statistics for results, including mean standard deviation in several tables, making clear that the reported values summarize multiple runs rather than a single run.

☒ **D. Did you use human annotators (e.g., crowdworkers) or research with human subjects?**

N/A 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.?
(left blank)

N/A 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)?
(left blank)

N/A 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)?
(left blank)

N/A D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
(left blank)

☑ **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?
No. AI assistants were used in a limited way for parts of the research and/or writing process, but this was not explicitly disclosed in the manuscript.