

Responsible NLP Checklist

Paper title: *Context-Conditioned Masked LoRA: Dynamic Rank Routing for Compute-Efficient Parameter-Efficient Fine-Tuning*

Authors: *Rifat Rafiuddin, Rafae Abdullah*

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?

Limitations section discusses potential risks, including differential treatment from input-conditioned routing across languages/dialects/domains and transparency risks if router overhead or kernel constraints are not reported.

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 use established public benchmarks (GLUE/SuperGLUE, XNLI, TyDi QA) and do not collect or release new datasets. Therefore, we did not perform additional PII/offensive-content screening or anonymization beyond the documentation and protocols of the original datasets, which are cited in 3.2.

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?

Section 3.2 describes the benchmarks used and the relevant split/evaluation protocols. Canonical train/dev splits for GLUE/SuperGLUE, English training with multilingual dev/test evaluation for XNLI, and official evaluation for TyDi QA.

C. Did you run computational experiments?

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

Section 3 describes the experimental setup, including models, datasets, baselines, metrics, and profiling protocol. Appendix B.3 reports the training setup and hyperparameter search/selection protocol, including optimizer, learning-rate grid, weight decay, warmup ratio, LoRA scaling, target rank, budget penalty, entropy weight, seed protocol, and dev-set model selection.

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.

- 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?

Section 4 reports result tables with mean std over seeds. Appendix G further clarifies the aggregation protocol, including the number of seeds, 95% confidence intervals, paired tests, and paired bootstrap statistics.

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.?

N/A the work uses existing public benchmarks and does not involve new human-subject studies, participants, or human annotation. See Section 3.2.

- 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)?

N/A no participants or annotators were recruited or paid because the work uses existing public benchmarks and does not involve new human-subject studies or human annotation. See Section 3.2.

- 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)?

N/A the work uses existing publicly released benchmarks and does not collect or curate new human-subject data. Consent and data-use terms are governed by the original dataset protocols/licenses, which are cited. See section 3.2.

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

N/A no new data collection or human-subject research was conducted. The work relies on existing public benchmarks, so ethics review board approval/exemption is not applicable. See section 3.2.

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 used AI assistants for language polishing, LaTeX/table formatting, and drafting submission metadata/checklist text. All research contributions, method design, experiments, analysis, reported results, and final claims were produced and verified by the authors.