

Responsible NLP Checklist

Paper title: *Relaxing the Constraints: A Dual-Importance Projection Mechanism for Lifelong Model Editing*

Authors: *Zhenghai Chen, Senbin Xu, Jiayi Tan, Xinhua Wu, Yan Zhang, Xiawu Zheng, Shengchuan Zhang, Ke Li, Sicheng Zhao, Liujuan Cao, Rongrong Ji*

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?

The proposed method focuses on model editing efficiency and does not introduce new safety risks beyond those inherent in LLMs.

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?

The ZsRE and CounterFact benchmarks are standard, publicly available datasets that do not contain personally identifying information.

- 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?

Data statistics, including the number of edits and categories, are reported in Sections 4.1 and 4.2.

C. Did you run computational experiments?

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

Detailed experimental setup and hyperparameters are provided in 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, etc. or just a single run?

We report mean results and performance improvements (e.g., 10.36% average gain) across multiple metrics in Table 1.

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.?

No human annotators or human subject research were involved.

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 annotators or human subject research were involved.

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)?

No human annotators or human subject research were involved.

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

No human annotators or human subject research were involved.

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 ChatGPT solely for the purpose of refining the linguistic clarity and improving the grammatical correctness of the manuscript. The AI assistant was not used for any core research, data analysis, or original content generation of the paper, 'Relaxing the Constraints: A Dual-Importance Projection Mechanism for Lifelong Model Editing'.