

Responsible NLP Checklist

Paper title: *Harmonizing the Past, Present, and Future: A Null-Space Constrained Region-Specific Method for Continual Learning in LLMs*

Authors: *Jinhui Chen, Shizhu He, Xingchang Yang, Huanxuan Liao, Yequan Wang, Xiangwen Liao, Wenhao Teng, Kang Liu, Jun Zhao*

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?

We did not discuss potential risks because this is a foundational algorithmic study focusing on continual learning methodology. It does not introduce new direct societal risks or biases beyond those already inherent in the pre-trained open-source LLMs utilized.

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 did not discuss anonymization steps because we exclusively utilized established, publicly available benchmarks (e.g., Standard CL, Long Sequence, TRACE, MMLU, GSM8K) that are widely adopted by the research community. We did not collect any new data containing personally identifying information (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?

Detailed dataset statistics, including train/test splits and the number of instances sampled per task (e.g., 1,000 for training, 500 for testing), are explicitly reported in the "Datasets and Benchmarks" section of Appendix D.

C. Did you run computational experiments?

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

Comprehensive details of the experimental setup, including specific learning rates, maximum training epochs, early stopping mechanisms, and sparsity ratios for different benchmarks, are provided in the "Implementation Details" section of Appendix D.

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?

Yes. We report the mean and standard deviation of our main experimental results across three independent runs initialized with different random seeds. This is explicitly stated in Section 5.2 ("Main Results") and in the captions of the relevant tables.

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?

During the preparation of this work, the authors utilized AI-based writing assistants primarily for grammatical error correction, sentence polishing, and LaTeX formatting optimization to enhance readability. The authors refined and validated all AI-generated suggestions. We explicitly state that all scientific concepts, experimental designs, data analyses, and conclusions presented in this paper are the original work of the authors, who take full responsibility for the content.