

## Responsible NLP Checklist

Paper title: *Learning What to Remember: Adaptive Probabilistic Memory Retention for Memory-Efficient Language Models*

Authors: *S M Rafiuddin, Muntaha Nujat Khan*

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 paper does not include an explicit discussion of potential societal risks (e.g., misuse, privacy, fairness, or safety). It only reports technical limitations (Sec. 5, Limitations) such as lack of evaluation on autoregressive decoding, overhead profile, scale, and hyperparameter sensitivity.*

### B. Did you use or create scientific artifacts? (e.g. code, datasets, models)

- B1. Did you cite the creators of artifacts you used?

*Creators of all used artifacts (datasets and pretrained/baseline models) are cited in Section 4.1.*

- B2. Did you discuss the license or terms for use and/or distribution of any artifacts?

*Licenses/terms for artifacts are not discussed in the paper; licensing information will be provided in the released repository.*

- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?

*All external artifacts were used in accordance with their stated access and intended-use restrictions, and our released artifacts specify academic research-only use (Section 4.1 and Section 5).*

- 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 datasets used are publicly available and anonymized; the paper does not report an explicit PII/offensive-content analysis.*

- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?

*Documentation of artifacts, including dataset domains and dataset characteristics (e.g., avg. lengths), is provided in Section 4.1; implementation and release details appear in Section 4.2 and Appendix A.*

---

*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.*

- 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 did not report dataset statistics (number of examples or train/dev/test splits) in the paper.*

**C. Did you run computational experiments?**

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

*Model parameter counts are reported in Section 4.1, training and experimental setup are in Section 4.2, and compute budget/infrastructure details are provided in Appendix A.*

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

*Experimental setup, including hyperparameter search procedures and the best-found hyperparameter values, is described in Section 4.2.*

- 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?

*Results are reported as single-run values without descriptive statistics (e.g., mean/std or error bars).*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation, such as NLTK, SpaCy, ROUGE, etc.), did you report the implementation, model, and parameter settings used?

*Implementation, model, and parameter settings for external packages (preprocessing, normalization, and evaluation) are detailed in Section 4.2 and Appendix A.*

**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.?

*(left blank)*

- 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)*

- 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)*

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

*(left blank)*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*(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?

*AI assistants were not used in the research, coding, or writing of this paper.*