

Responsible NLP Checklist

Paper title: *AURORA: Neuro-Symbolic Continual Indexing for Evolving RAG Systems*

Authors: *Manoj Saravanan, Rohit Kumar Salla, Ramya Manasa Amancherla*

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.

- A2. Did you discuss any potential risks of your work?

10 (Ethical Considerations); see also 9 (Limitations)

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 standard public research benchmarks (AG News, SQuAD, NRCorpus) and a synthetic vector benchmark (SIFT1M). We did not collect, curate, annotate, or release new human-subject data in this work, and the paper does not include a separate PII/offensive-content screening or anonymization procedure beyond the original benchmark curation.

- 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?

5.1 and App. B.4. These sections describe the datasets and drift protocols used in evaluation, including the 50/50 D_old/D_new partition, SIFT1M (N=10^6, d=128), and the sequential update protocol sizes used in the continual-learning experiments.

C. Did you run computational experiments?

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

5.3 and App. B.1. We describe the experimental setup in 5.3 and report the best-found hyperparameters in App. B.1 (Hyperparameter Configuration), including Reptile settings, router configuration, dense quantizer settings, and LoRA ranks.

- 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 single-run point estimates in the main tables and figures and do not report multi-seed meanstd, confidence intervals, or error bars.

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.?

(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)

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?

(left blank)