

Responsible NLP Checklist

Paper title: *PanoramaRAG: Enabling Consistent Global Topic Awareness in Graph-Based RAG*

Authors: *Ding Deng, Xiang Li, Yaqing Zhang, Meng Li, Xiting Wang*

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?

We do not identify significant potential risks associated with our work, as it focuses on improving retrieval-augmented generation methods on publicly available datasets and does not involve deployment in safety-critical or user-facing decision-making scenarios.

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 used publicly available benchmark datasets that do not contain personally identifying information. We also performed standard preprocessing to remove malformed or irrelevant content. Therefore, no additional anonymization procedures were required.

- 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 4

C. Did you run computational experiments?

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

section 4.5

- 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 4.3

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.?

section 4

- 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)?

A small amount of manual annotation was performed by internal annotators for evaluation purposes. No external participants were recruited through crowdsourcing platforms, and no payment or compensation was provided beyond regular research involvement within the project.

- 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)?

We use publicly available benchmark datasets that do not contain personally identifiable information, and no new data was collected from human subjects. Therefore, no additional consent procedures were required.

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

We use publicly available benchmark datasets that do not contain personally identifiable information, and no new data was collected from human subjects. Therefore, no additional consent procedures were required.

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?

appendix