

Responsible NLP Checklist

Paper title: *HiKEY: Hierarchical Multimodal Retrieval for Open-Domain Document Question Answering*

Authors: *Joongmin Shin, Gyuho Shim, Jeongbae Park, Jaehyung Seo, Heuiseok Lim*

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?

Limitations (unnumbered section after Sec. 5). In particular, the paper discusses risks and failure conditions related to preprocessing/indexing overhead, dependence on upstream OCR/layout parsing, structural assumptions, and post-retrieval evidence-state.

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?

N/A. The paper uses publicly available research benchmarks/corpora only, and FRAMES-PDF is constructed by rendering publicly available Wikipedia articles into multi-page PDFs; it does not collect new personal data or recruit participants.

- 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?

Appendix A.1 and Table 9. The paper reports the number of documents/questions, evidence-type counts, and hop distributions for M3DocVQA and FRAMES.

C. Did you run computational experiments?

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

Sec. 4.1, Appendix A.2, and Appendix F (especially F.1F.3), plus Table 12. These sections describe the preprocessing pipeline, retrievers, reader model, context limit, decoding setup, and evaluation configuration.

- 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?

Sec. 4.34.9, Appendix B, and Tables 16 and 9. The paper is explicit about reporting averages

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?
We used an AI assistant only for translation and language polishing due to the authors being non-native English speakers. It was not used for research design, experiments, coding, or result interpretation; all content was reviewed and verified by the authors.