

## Responsible NLP Checklist

Paper title: *Evidence-Aligned Entity Verification for Hallucination Detection in Retrieval-Augmented Generation*

Authors: *Runsong Jia, Zhen Fang, Mengjia Wu, Jie Lu, Yi Zhang*

How to read the checklist symbols:

- the authors responded 'yes'
- the authors responded 'no'
- <sup>N/A</sup> 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.*

- <sup>N/A</sup> A2. Did you discuss any potential risks of your work?

*Our study poses minimal risk since all datasets and baseline methods used are publicly available and widely adopted in prior research.*

### 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 publicly available datasets that have been used in prior work. We discussed the details of the dataset in A.2 Experimental Details and the dataset does not allow identification of individuals beyond what is already publicly accessible.*

- 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?

*A.2 Experimental Details*

### C. Did you run computational experiments?

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

*A.2 Experimental Details, A.3 Additional Results*

- 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?

*4 Experiment*

*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.?

*We did not involve any human annotators or research with human subjects in this study.*

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)?

*As no human subjects or annotators were used, this item is not applicable.*

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)?

*All datasets used in this work are publicly available.*

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

*All datasets used in this work are publicly available.*

**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?

*A.6 Usage Claim of AI assistants. AI assistants were used only to refine and revise phrasing in the written content.*