

Responsible NLP Checklist

Paper title: *FactVerse: A Benchmark for Factual Consistency in Interleaved ImageText Generation*

Authors: *Yubo Shan, Kun Zhang, Qiming Xu, Liping Cao, Yingying Cao, Jian Zhang, Yu Wang, Jingyuan Li, Yuanzhuo 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 did not discuss potential risks because this work focuses on an evaluation benchmark designed to measure and improve the factual consistency of MLLMs. The dataset contains publicly available, non-sensitive information, and the framework is intended to mitigate the risks of hallucinations rather than introducing new generative risks.

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?

AppendixC.3

- 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?

AppendixA.1

C. Did you run computational experiments?

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

AppendixD.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?

5 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.?
(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)?
AppendixC.1

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 did not explicitly discuss consent in the paper because the collected data consists of objective factual knowledge and model generations, rather than personal data from human subjects. The annotators were recruited as paid workers for dataset construction and were informed of the task instructions.

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 utilized AI assistants solely for improving the linguistic quality of the manuscript (grammar checking and rephrasing) and generating auxiliary code snippets for data formatting. The AI tools contributed neither to the scientific conceptualization nor to the drafting of the core intellectual content, hence we did not include an explicit disclosure in the main text.