

Responsible NLP Checklist

Paper title: *Empowering Reliable Visual-Centric Instruction Following in MLLMs*

Authors: *Weilei He, Feng Ju, Zhiyuan Fan, Rui Min, Minhao Cheng*

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?

There is no potential risks of our work.

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?

Our work uses only publicly available datasets that have been previously released for research purposes (e.g., OpenAssistant, Visual Genome, LLaVA-Instruct). We did not collect any new human data. Therefore, no additional steps for anonymization or offensive content filtering were necessary.

- 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?

3

C. Did you run computational experiments?

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

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?

5

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.?

The annotators were provided with general task descriptions and examples consistent with standard evaluation practices in multimodal instruction-following benchmarks. No sensitive, risky, or ethically concerning information was involved in the annotation process.

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.

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)?

We are not paid.

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)?

The data used in this work were either publicly available or generated for research purposes under standard academic settings.

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

No ethics review board approval was required for this work. The study involved only publicly available data and voluntary human annotation for non-sensitive, research-only purposes, which are exempt from formal ethics review under standard academic policies.

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 use open-source frameworks like LLaMA Factory, they are not hard to use.