

## Responsible NLP Checklist

Paper title: *Interleaved Latent Visual Reasoning with Selective Perceptual Modeling*

Authors: *Shuai Dong, Siyuan Wang, Xingyu Liu, Chenglin Li, Haowen Hou, zhongyu wei*

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?

*This work proposes an architectural framework (ILVR) to improve MLLM efficiency. It does not introduce new datasets or high-risk capabilities; thus, we foresee no specific ethical risks beyond those inherent to general MLLMs.*

### 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 established public benchmarks and relied on their original anonymization. No new data involving human subjects or uncurated sources was collected.*

- 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 4.1, Appendix D*

### C. Did you run computational experiments?

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

*Section 4.1, Appendix B*

- 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?

*Due to the high computational cost of MLLMs, we report reproducible single-run results using a fixed seed (42) without error bars.*

*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.?

*Not applicable. This research did not involve human subjects, crowdworkers, or human annotators.*

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)?

*Not applicable as no human participants or annotators were involved in this study.*

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 collect new data from individuals, so obtaining direct consent was not applicable to this work.*

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

*This computational research on public datasets involves no human subjects and is exempt from IRB approval.*

**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?

*See the Use of AI Assistants section at the end of the main paper (before References).*