

Responsible NLP Checklist

Paper title: *AICA-Bench: Holistically Examining the Capabilities of VLMs in Affective Image Content Analysis*

Authors: *Dong She, Xianrong Yao, Liqun Chen, Jinghe Yu, Yang Gao, Zhanpeng Jin*

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.

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

(left blank)

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?

Section 3.2

- 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 3.2 and Appendix A.1 (Table 5) for AICA-Bench (number of images from each source and overall counts), and in Appendix B.1 for the human-annotated data used to train the scoring model.

C. Did you run computational experiments?

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

Section 4.1 (Experiments Experimental Setup), Appendix B.2 (Training Details, Table 7).

- 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?

Section 4.2 (Main Results and Analysis, Table 3), regression metrics in Table 2, GAT vs. baseline comparisons in Table 4, the full leaderboard in Appendix B.3B.4 (Table 8), and mean human ratings from 10 annotators in Table 9.

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 describe annotator recruitment and compensation in Appendix A.3.

- 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 describe annotator recruitment and compensation in Appendix A.3.

- 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)?

see Section 3.2 and Appendix A.3 for consent, privacy, and sensitive-content filtering.

- 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 AI assistants for minor grammar correction, wording refinement, coding, and simple icon designs.