

Responsible NLP Checklist

Paper title: *MIND Your Reasoning: A Meta-Cognitive Intuitive-Reflective Network for Dual-Reasoning in Multimodal Stance Detection*

Authors: *Bingbing Wang, Zhengda Jin, Bin Liang, Wenjie Li, Jing Li, Ruifeng Xu, Min Zhang*

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 include a dedicated discussion of potential risks in the current manuscript. The paper focuses on a model and experiments for multimodal stance detection on a public benchmark, and the existing Limitations section discusses computational cost and inference latency rather than broader societal or misuse 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?

We did not include a dedicated discussion of whether the data contains personally identifying information or offensive content, nor did we describe any anonymization or protection procedures. Our experiments were conducted on the public MMSD benchmark rather than on a newly collected dataset, and the manuscript does not provide a section discussing data screening or anonymization steps.

- 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?

We used the public MMSD benchmark and listed its five domains in Section 4.1, but we did not restate detailed dataset statistics such as the number of examples and train/dev/test splits in the current manuscript. Since our work adopts an existing public benchmark rather than introducing a new dataset, we focused on the model design and evaluation results and referred to the original dataset paper for dataset construction details.

C. Did you run computational experiments?

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

Section 4.1

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.

- 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 4.3

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)?

(left blank)

- 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)?

(left blank)

- 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?

(left blank)