

Responsible NLP Checklist

Paper title: *EmoMAS: Emotion-Aware Multi-Agent System for High-Stakes Edge-Deployable Negotiation with Bayesian Orchestration*

Authors: *Yunbo Long, Yuhan Liu, Liming Xu*

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?

Yes, discussed in Section 7 (Limitations) and Section 8 (Ethical Considerations).

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?

All data is fully synthetic, generated by LLMs with fictional scenarios; no real persons or personally identifying information are involved.

- 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?

Dataset sizes and per-domain statistics are reported in Section 4 and Appendix (Dataset Details).

C. Did you run computational experiments?

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

Experimental setup is described in Section 4, with full hyperparameter values in Appendix (Hyperparameters).

- 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?

Results are reported as means with 95% bootstrap confidence intervals over 100 scenarios per configuration in Section 5 (Tables 13).

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.?

No human participants or annotators were involved; all evaluation is agent-to-agent with synthetic LLM-generated scenarios.

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)?

No human participants were recruited; all experiments are agent-to-agent with synthetic LLM-generated data

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)?

No human data was collected; all scenarios are synthetically generated by LLMs.

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

No human subjects; all scenarios are synthetic LLM-generated, exempt from ethics review.

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?

AI assistants (ChatGPT) were used for language polishing and LaTeX formatting; all technical contributions and experimental design are the authors' own work. Disclosed in the Appendix (AI Assistant Disclosure).