

## Responsible NLP Checklist

Paper title: *Generalization or Memorization? Multi-Agent vs. Baseline LLMs and AutoML Models for Tabular Classification*

Authors: *Aida Sanatizadeh, Sorouralsadat Fatemi, Reza Mousavi, Ahmed Abbasi*

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?

*Sections 8 (Limitations) and 9 (Ethical Considerations) discuss risks related to memorization, poor calibration on novel data, misuse in high-stakes settings, sensitive mental health attributes, lack of fairness analysis, and deployment 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?

*Sections 8 and 9 discuss the presence of sensitive attributes in mental health and workplace datasets, risks of sending data to external LLM APIs, and the need for safeguards such as IRB approval and informed consent for real-world use.*

- 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 and Table 1 report the number of instances, features, class distributions, and dataset sources, with additional details in Appendix A.*

### C. Did you run computational experiments?

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

*Section 4 describes all models and configurations. Appendix F provides detailed hyperparameters*

- N/A 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 5 reports confidence intervals and error analysis across models and datasets, quantifying uncertainty, variance, and failure modes. Appendix H reports mean and standard deviation across 5 runs.*

---

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.?

*(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?

*We used AI assistants for writing and debugging code, and assistance with manuscript writing, in line with ACL policies.*