

Responsible NLP Checklist

Paper title: *Fair-CCD: Mitigating Bias in Large Language Models for Tabular Classification Through Context-Contrastive Decoding*

Authors: *Donghan Liu, Han Sun, Zhaohui Wang, Qin Li, 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?

Section 6 (Conclusions) and Limitations

B. Did you use or create scientific artifacts? (e.g. code, datasets, models)

- N/A 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?

Not applicable. The datasets used in this work (Adult and COMPAS) are publicly available and have been preprocessed and anonymized. No personally identifiable information is directly used or collected in this study.

- 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?

No. This work uses publicly available benchmark datasets (Adult and COMPAS), and we follow standard evaluation protocols without introducing new data splits or reporting additional dataset statistics.

C. Did you run computational experiments?

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

No. The proposed method does not rely on explicit hyperparameter search. The scaling factors are determined dynamically via the attention-guided mechanism during inference.

- 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?

Yes. We report aggregated results across multiple experiments, including averaged performance over sampled demonstrations and summary statistics presented in tables and figures.

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 work does not involve human subjects or human annotation.

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. This work does not involve recruiting or compensating human participants.

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)?

Not applicable. This work does not involve collecting data from human participants.

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

Not applicable. This work does not involve human subjects or data collection requiring ethical 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?

No. The use of AI assistants was limited to minor language editing and does not affect the scientific content of this work.