

Responsible NLP Checklist

Paper title: *From Numbers to Narratives: Efficient Language Model-Based Detection for Safety-Critical Minority Classes*

Authors: *Ahatsham Hayat, Hunter Tridle, Mohammad Rashedul Hasan*

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

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

- B1. Did you cite the creators of artifacts you used?

Section 3.2 and Appendix A.4

- B2. Did you discuss the license or terms for use and/or distribution of any artifacts?

We did not discuss artifact licenses, as the focus was on methodological contributions, and datasets/models are publicly available with licensing details provided by their original sources.

- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?

Datasets (e.g., A4H 2020, Sklearn, SECOM) and models like ChaGPT4, intended for numerical classification and general text generation respectively, are repurposed for narrative-based classification through our "numbers to narratives" framework (Section 3.1), which transforms tabular data into textual representations, an unconventional application not aligned with their original design.

- 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?

The datasets (e.g., AI4I, Glass Identification, Student Performance, Gas Detection, Mammographic, SECOM) utilized in the paper, as described in Section 3 and Table 1, consist of numerical or categorical features relevant to specific domains (e.g., machine failure, student performance, semiconductor manufacturing) and do not include personally identifying information such as names, addresses, or sensitive personal attributes. Additionally, the data processing and transformation into narratives (Section 3.1 and Section A.4.1) focus on abstract feature representations without generating or incorporating offensive content, ensuring compliance with ethical data usage standards.

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.

- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 3.3
- 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 4
- C. Did you run computational experiments?**
- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4
- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Appendix A.4.2, A.4.3, and A.4.4
- 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 4.1, 4.2, and 4.3
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation, such as NLTK, SpaCy, ROUGE, etc.), did you report the implementation, model, and parameter settings used?
Section 4
- 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)
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
(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 did not include information about AI assistant use in the paper itself. While we used AI assistants during the writing and revision process to help refine language, structure arguments, and ensure clarity of presentation, we did not explicitly acknowledge this assistance in any section of the paper.