

Responsible NLP Checklist

Paper title: *Towards Automated Error Discovery: A Study in Conversational AI*

Authors: *Dominic Petrak, Thy Thy Tran, Iryna Gurevych*

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?

Section 7 (Limitations)

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

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

Section 4, 5, Appendix B

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

The artifacts we use have all been published at representative venues and use established open source licences. We will make our artifacts available under an established open-source license that is compatible.

- 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)?

The artifacts we use have all been published for scientific purposes. We will publish our artifacts on GitHub and clearly state their intended scientific use in the repository.

- 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 collect new data and the datasets we use already discuss these aspects in their accompanying peer-reviewed publications.

- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?

We provide a documented simplification of our proposed sampling mechanism (Section 4.2 (Algorithm 1)), the prompts used (Appendix A), and the error type distributions of the datasets considered in our experiments (Appendix C.1 (Datasets Statistics)).

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.

- 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?

Appendix C.1

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?

Appendix B.2

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

Appendix B

- 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, Appendix D

- 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 5

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.?

We did not involve crowdworkers or other external participants.

- 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)?

We did not involve crowdworkers or other external 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)?

We did not involve crowdworkers or other external participants.

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

We did not involve crowdworkers or other external participants.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

We did not involve crowdworkers or other external participants.

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 only used AI assistance for improving spelling and grammar, which is permitted under the ACL policy (<https://2023.aclweb.org/blog/ACL-2023-policy/>).