

Responsible NLP Checklist

Paper title: *On Finding Inconsistencies in Documents*

Authors: *Charles Lovering, Seth Ebner, Brandon Smock, Michael Krumdick, Muhammad Saad Rabbani, Ahmed Muhammad, Varshini Reddy, Chris Tanner*

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?

Limitations section.

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?

No personal information.

- 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?

4 Dataset

C. Did you run computational experiments?

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

H.1; Overall, because the task is expensive we did not extensively tune the tested models. We discuss development dataset results in the appendix.

- 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?

6 Results and Analysis; we report mean and additionally report the top equivalence class of model performances. In the appendix we report the standard errors (table H.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.?

We include the links to the documentation we provide to the annotators (who are collaborators and authors on this paper).

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.

- 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)?

See "Housekeeping", start of appendix. We do include minimal information here because it is private. We list that all annotators are paid members of our organization.

- 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)?

Yes, see 4. Dataset (and G. Dataset in the appendix) for permission information on the dataset.

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

See "Housekeeping", start of appendix.

- 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?

See "Housekeeping", start of appendix.