

## Responsible NLP Checklist

Paper title: *Just Pass Twice: Efficient Token Classification with LLMs for Zero-Shot NER*

Authors: *Ahmed Ewais, Ahmed Hashish, Amr Ali*

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?

*The work presents a method for named entity recognition and does not introduce risks beyond those inherent to any NER system (e.g., misclassification). No sensitive applications are proposed.*

### 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?

*Training data is derived from public Wikipedia articles. Evaluation benchmarks are established public datasets. No private or personally identifying information was collected.*

- 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 B, Section B.1 (Table 7) reports dataset statistics including number of examples, tokens, entity mentions, and unique entity types. Section B.2 reports train/test split sizes (17,489 training and 2,500 test examples).*

### C. Did you run computational experiments?

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

*Section 4.1 describes the experimental setup. Appendix A, Section A.3 (Table 6) reports all training hyperparameters.*

- 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?

*We report single-run F1 scores consistent with prior work in this area (UniNER, GoLLIE, GLiNER, SaM). Error bars are not reported as the evaluation is deterministic (no sampling) on fixed test sets.*

---

*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?  
*Claude Sonnet 4.5 and Claude Opus 4.5 were used to generate and validate training data annotations, as described in Appendix B.2 (Dataset Construction).*