

Responsible NLP Checklist

Paper title: *Understanding or Memorizing? A Case Study of German Definite Articles in Language Models*

Authors: *Jonathan Drechsel, Erisa Bytyqi, Steffen Herbold*

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?

We omit a risks discussion because this is a non-deployment interpretability study of German grammar that introduces no new model or sensitive data and is unlikely to add meaningful risks beyond those of standard language models.

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?

We use only publicly available, widely used corpora (German Wikipedia and Leipzig news) and restrict to filtered, non-user-generated text, so we did not perform additional PII/offensive-content anonymization beyond this filtering.

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 A

C. Did you run computational experiments?

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

We provide details on experimental setup including chosen hyperparameters, but we do not perform hyperparameter search, as our experiments build up on existing frameworks (GRADIEND, SuperGLEBer).

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 C

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.?

Appendix 1.4 provides full details for data annotation used to assess spaCy quality. The annotation task - gender-case identification - is straightforward and does not require annotation guidelines. All annotations were carried out by one of the authors; no external participants were involved.

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?

We did not include a paper statement because AI assistance was limited to routine coding and writing support, while all experiments, analysis, and final text were produced and verified by the authors.