

Responsible NLP Checklist

Paper title: *DOS: Dependency-Oriented Sampler for Masked Diffusion Language Models*

Authors: *Xueyu Zhou, Yangrong Hu, Jian Huang*

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?

We discuss potential risks and ethical considerations in the Ethics Statement on Page 9.

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?

This work does not involve the collection or creation of new datasets. We only use existing publicly available benchmarks, which do not contain personally identifying 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?

We do not report dataset statistics such as the number of examples or train/dev/test splits because this work does not involve collecting, creating, or modifying any datasets. All experiments are conducted on standard, publicly available benchmark datasets.

C. Did you run computational experiments?

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

We report the experimental setup and hyperparameters in Section 6 and Appendix A.

- 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 do not report descriptive statistics such as error bars or variance across multiple runs because, following prior work, all results are reported from a single run with a fixed random seed.

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.?

This work does not involve human subjects or human annotators.

- 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)?

This work does not involve human subjects or human annotators.

- 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)?

This work does not involve the collection or curation of human data. We only use existing publicly available datasets.

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

This work does not involve the collection of data from human subjects.

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?

AI assistants were used solely for minor language polishing and improving the clarity of presentation (e.g., grammar and wording).