

Responsible NLP Checklist

Paper title: *Arctic-Text2SQL-RI: Simple Rewards, Strong Reasoning in Text-to-SQL*

Authors: *Zhewei Yao, Guoheng Sun, ukasz Borchmann, Zheyu Shen, Minghang Deng, Bohan Zhai, Hao Zhang, Ang Li, Yuxiong He*

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?

Yes. See the 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. The datasets used in this work (e.g., BIRD, SPIDER) are publicly available benchmarks that do not contain personally identifying information, and we did not collect any new human-related data.

- 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?

Yes. Dataset statistics, including the number of examples and splits, are reported in Appendix B (Table B.1).

C. Did you run computational experiments?

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

Yes. We describe experimental setup and hyperparameters in Section 3 and Appendix B, including training configurations such as temperature, batch size, KL penalty, and clipping parameters.

- 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?

No. We primarily report single-run performance and benchmark accuracy, and do not include variance estimates or multiple-run statistics such as standard deviations.

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.?

N/A. This work does not involve human participants or 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)?

N/A. This work does not involve human participants or 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)?

N/A. This work does not involve human participants or annotators.

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

N/A. This work does not involve human participants or annotators.

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?

Yes. We disclose the use of LLMs for language polishing in Appendix A (LLM Usage Statement), including the scope and limitations of their use.