

Responsible NLP Checklist

Paper title: *CARL: Constraint-Aware Reinforcement Learning for Planning with LLMs*

Authors: *Qiuyi Qi, Jinjian Zhang, Mutian Bao, Tian Liang, Guocong Li, Dongnan Liu, Wei Zhou, Jie Liu, Ming Kong, Linjian Mo, Feng Zhang, Qiang Zhu*

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.

- ^{N/A} A2. Did you discuss any potential risks of your work?

The proposed method focuses on improving the planning capabilities of LLMs on standard academic benchmarks (BlocksWorld, TravelPlanner, T-Eval). We do not foresee significant ethical risks or negative societal impacts associated with this specific technical contribution.

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 utilized established public benchmarks (BlocksWorld, TravelPlanner, T-Eval) which are widely used in the research community and do not contain personally identifying information or offensive content.

- 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?

5

C. Did you run computational experiments?

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

5

- 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?

5

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.?

Not applicable (no human subjects).

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)?

Not applicable (no human subjects).

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)?

Not applicable (no human subjects).

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

Not applicable (no 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?

During the preparation of this manuscript, we utilized AI assistants (e.g., large language models) for the following purposes: 1) to assist with literature searches; and 2) to improve the grammar, phrasing, and overall clarity of our writing. All AI-generated text and suggestions were critically reviewed, edited, and verified by the authors. The authors retain full responsibility for the final content of the manuscript, including its accuracy, originality, and any potential errors.