

Responsible NLP Checklist

Paper title: *ProcWorld: Benchmarking Large Model Planning in Reachability-Constrained Environments*
Authors: *Dong Wang, Xinghang Li, Zhengshen Zhang, Jirong Liu, Xiao Ma, Hanbo Zhang, Tao Kong, Huaping Liu*

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?

(left blank)

B. Did you use or create scientific artifacts? (e.g. code, datasets, models)

- B1. Did you cite the creators of artifacts you used?

Section 2 (Related Work) and Section 3.2 (Environment Setups) cite creators of artifacts like TextWorld (Ct et al., 2019), ProcTHOR (Deitke et al., 2022), and others.

- B2. Did you discuss the license or terms for use and/or distribution of any artifacts?

While the paper states that the code and dataset will be open-sourced upon publication (Footnote 1), it does not explicitly discuss licenses or terms for existing artifacts (e.g., TextWorld, ProcTHOR).

- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?

Section 3.2 explains how ProcWORLD modifies TextWorld and ProcTHOR to align with LLM/VLM evaluation goals, preserving compatibility while introducing partial observability and reachability constraints.

- 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?

The dataset is procedurally generated using synthetic environments (Section 3.3). The paper does not mention personally identifiable information or offensive content, nor steps to anonymize data.

- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?

Section 3.3 (Benchmark Statistics) documents scene diversity (5,000 rooms), task types (16 categories), and evaluation trajectories (10 million).

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.

- 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?
Section 3.3 and Appendix C.1 detail dataset statistics, including task distributions, object categories, and validation/test splits.

C. Did you run computational experiments?

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4.5 reports model sizes (e.g., Qwen2.5 from 0.5B to 72B) and total evaluation runs (1,080,000 episodes). However, GPU hours and infrastructure details are not specified.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 4 describes the evaluation protocol, including observation modes (text/vision), reachability constraints, and in-context learning methods (ReAct, Reflexion).

- 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 4 reports success rates, average steps (Tables 214), and performance trends (Figures 58). However, error bars or variance metrics are not included.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation, such as NLTK, SpaCy, ROUGE, etc.), did you report the implementation, model, and parameter settings used?
The paper uses PDDL (Section 3.2) and ProcTHOR but does not specify implementation parameters for these tools.

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)

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
(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?
(left blank)