

Responsible NLP Checklist

Paper title: *Consolidation or Adaptation? PRISM: Disentangling SFT and RL Data via Gradient Concentration*

Authors: *Yang Zhao, Yangou Ouyang, Xiao Ding, Hepeng Wang, Bibo Cai, Kai Xiong, Jinglong Gao, Zhouhao Sun, Li Du, Bing Qin, Ting 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.

- A2. Did you discuss any potential risks of your work?

The work focuses on algorithmic efficiency and data routing for standard LLM agent benchmarks (WebShop, ALFWorld) using open-source models. It does not introduce new datasets, release models with hazardous capabilities, or involve applications with direct societal risks. Technical limitations are discussed in 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?

The experiments utilize established, publicly available benchmarks (WebShop and ALFWorld) and standard open-source models (Qwen, Llama). No new datasets containing personally identifying information or offensive content were collected or created for this study.

- 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 3.1 ("Experimental Setup") and Table 1/Table 2 (referencing "Seen" and "Unseen" splits and data allocation ratios). Further details on context length and steps are in Appendix C ("Implementation Details").

C. Did you run computational experiments?

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

Appendix C ("Implementation Details") lists specific hyperparameters for Gradient Probing, SFT (learning rate, epochs, batch size), and RL (GRPO settings, reward functions).

- 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,

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.

etc. or just a single run?

Section 3.1 ("Baselines") explicitly states: "Crucially, to ensure robustness, all reported results represent the mean across three random seeds."

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)

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?

Appendix H ("AI Assistance Disclosure") states: "We acknowledge the use of AI tools solely for language polishing and grammatical editing"