

Responsible NLP Checklist

Paper title: *DR-Arena: an Automated Evaluation Framework for Deep Research Agents*

Authors: *Yiwen Gao, Ruo Chen Zhao, Yang Deng, Wenxuan Zhang*

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?

The proposed DR-Arena is an evaluation framework for research agents using publicly available web data and does not introduce significant ethical or security risks.

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 personally identifiable information was collected during automated data construction or human audit.

- 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?

Detailed statistics on source domains and topic distributions are provided in Appendix F (Table 8 and 9).

C. Did you run computational experiments?

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

Model choices, tournament settings, and scoring mechanisms are described in Section 4.1.

- 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 report Pearson/Spearman correlations, Elo ratings, and error distributions in Section 4.2 and Appendix I.

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.?

The exact questionnaires and criteria given to participants are provided in Appendix H.1.

- 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)?

Details are provided in Appendix H.

- 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)?

Details are provided in Appendix H.

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

No explicit ethics review was sought as the study involved auditing non-sensitive, automated system outputs by university students, posing no significant ethical risks to participants.

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 tools were used for linguistic refinement and proofreading to enhance the clarity of the text.