

Responsible NLP Checklist

Paper title: *Is Grokking Worthwhile? Functional Analysis and Transferability of Generalization Circuits in Transformers*

Authors: *Kaiyu He, Mian Zhang, Peilin Wu, Xinya Du, Zhiyu Chen*

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?

This research is a mechanistic study of Transformer architectures using synthetic datasets to understand internal reasoning paths. It does not involve real-world data, human subjects, or applications that pose foreseeable ethical or social 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?

The study utilizes a synthetically generated knowledge graph consisting of abstract entities and relations. It contains no 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?

Section 2 and Appendix B provide the total number of entities (2,000), relations (200), and the 95%/5% split for In-Distribution and Out-of-Distribution facts.

C. Did you run computational experiments?

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

Section 2 and Appendix B.3

- 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?

Sections 3 and 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.?

The research relies entirely on algorithmically generated synthetic data and automated mechanistic analysis; no human participants or annotators were involved.

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)?

The research relies entirely on algorithmically generated synthetic data and automated mechanistic analysis; no human participants or annotators were involved.

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)?

The research relies entirely on algorithmically generated synthetic data and automated mechanistic analysis; no human participants or annotators were involved.

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

The research relies entirely on algorithmically generated synthetic data and automated mechanistic analysis; no human participants or annotators were involved.

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 D