

## Responsible NLP Checklist

Paper title: *A Layer-wise Analysis of Supervised Fine-Tuning*

Authors: *Qinghua Zhao, Xueling Gong, Xinyu Chen, Zhongfeng Kang, Xinlu Li*

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 work focuses on analyzing internal representations of language models rather than deploying them in real-world applications. We do not identify specific potential risks associated with this work.*

### 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 data used in this work consists of publicly available datasets (MMLU, GSM8K, WikiText, IFEval, HumanEval, MT-Bench, ToxiGen) and model-generated representations. These datasets do not contain personally identifying information. ToxiGen contains synthetic toxic content for safety evaluation purposes, but no real personal data.*

- 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 4 (Experimental Setup) describes the datasets used, including sample sizes (100 instances for token-level analysis, 1,000 instances for pooled-level analysis) and the benchmarks evaluated.*

### C. Did you run computational experiments?

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

*Section 4 (Experimental Setup) describes the experimental setup in detail, including model configurations (OLMo2-1B/7B/13B/32B, Mistral-7B), LoRA hyperparameters (rank  $r$  {8, 64}), dataset sampling strategies (100 instances for token-level, 1,000 for pooled-level analysis), and implementation details such as random seed fixing and streaming computation for memory efficiency.*

- 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 4.2 and Section 4.4 report mean performance across runs and clearly state the reported*

---

*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.*

statistics. Robustness checks using different random seeds are conducted, and statistical significance is verified via  $t$ -test ( $p < 0.01$ ) as described in Section 4.2.

**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?

*(left blank)*