

Responsible NLP Checklist

Paper title: *On the Cultural Anachronism and Temporal Reasoning in Vision Language Models*

Authors: *Mukul Ranjan, Prince Jha, Khushboo Kumari, Zhiqiang Shen*

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?

This work presents a diagnostic evaluation benchmark using publicly available museum artifacts to identify limitations in existing models rather than introducing new capabilities or deployable systems, thus presenting minimal risks beyond standard academic research.

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 dataset consists exclusively of inanimate historical artifacts (pottery, sculptures, manuscripts, textiles) from public museum collections, which contain no personally identifying information or depictions of identifiable individuals.

- 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 3.2

C. Did you run computational experiments?

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

Section 4

- 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, Table 1

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.?

Annotators were authors of the paper; the task was a binary judgment (pipeline scoring correct/incorrect) requiring no external instructions.

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)?

Annotators were authors of the paper; no external recruitment or payment was 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)?

Annotation was performed on model outputs, not on data from human subjects; no personal data was collected.

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

No human-subjects data was collected; annotation of model outputs by the authors does not require IRB review.

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?

Section 8