Responsible NLP Checklist

Paper title: *Generating Attribution Reports for Manipulated Facial Images: A Dataset and Baseline*
Authors: *Jingchun Lian, Lingyu Liu, Yaxiong Wang, Yujiao Wu, Lianwei Wu, Li Zhu, Zhedong Zheng*

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?
See Enthical Consideration

☑ **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 datasets utilized in this study are publicly available and widely used benchmarks. Our work involves modifications based on these existing datasets and does not involve the collection of any new personally identifiable information (PII). Therefore, additional privacy protection steps were not discussed.

☑ 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 Multi-Modal Tamper Tracing Dataset, Supplementary

☑ **C. Did you run computational experiments?**

☑ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Supplementary D Experimental Setup and F ablation study

☒ 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?
Due to the high computational cost of training large multimodal models (approx. 48 GPU hours on H100), we follow previous works and only report the results of a single run with fixed random seeds for reproducibility, which is standard practice in this domain.

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.?
Section 2 Multi-Modal Tamper Tracing Dataset

☑ 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)?
section 2, we employ volunteer experts for dataset annotation

☑ 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)?
section 2 dataset source collection and Supplementary I Statements

☑ D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
section 2, this research primarily utilizes existing, publicly available datasets. We did not engage in active data collection from human subjects that would typically mandate an institutional ethics review. The small-scale human evaluation was conducted internally for validation purposes only.

☑ **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?
Supplementary I.3 LLM usage statements