

Responsible NLP Checklist

Paper title: *XLSR-MamBo: Scaling the Hybrid Mamba-Attention Backbone for Audio Deepfake Detection*

Authors: *Kwok-Ho Ng, Tingting Song, Yongdong WU, Zhihua Xia*

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?
Section "Ethical Considerations"

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 used are publicly available academic benchmark datasets that contain only anonymized speaker IDs, with no real names or personally identifiable information, and no offensive content. As they were released with privacy compliance already ensured by the providers, no additional steps for anonymization or removal of offensive material were required.

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 4.1

C. Did you run computational experiments?

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 4.2

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 5.1, 5.2

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.?
Our work does not involve human annotators (e.g., crowdworkers) or research with human subjects.

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.

All experiments are conducted on publicly available audio datasets with binary bona-fide/spoofed labels. Therefore, no instructions, screenshots, disclaimers, or related information were provided to participants.

- N/A 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)?

Our work does not involve human participants, annotators, or crowdworkers. All experiments are performed on publicly available audio datasets with pre-existing binary labels. Therefore, no recruitment, payment, or related demographic considerations apply.

- N/A 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)?

We do not collect, curate, or annotate any new data from individuals. Our experiments rely solely on publicly available audio datasets standard in the audio deepfake detection field. These datasets were originally collected with consent for research purposes, as stated in their respective publications and access agreements (cited in the paper). We use them strictly in accordance with their licenses and terms, which permit academic research use. Therefore, no additional consent was obtained or discussed.

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

This work does not involve any data collection from human subjects. All experiments are conducted using pre-existing, publicly available audio datasets. Therefore, no ethics review board approval or exemption was required.

- 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 "Acknowledgments"