

## Responsible NLP Checklist

Paper title: *UniSonate: A Unified Model for Speech, Music, and Sound Effect Generation with Text Instructions*

Authors: *Chunyu Qiang, Xiaopeng Wang, Kang Yin, Yuzhe Liang, Yuxin Guo, Teng Ma, Ziyu Zhang, Tianrui Wang, Cheng Gong, Yushen CHEN, Ruibo Fu, Longbiao Wang, Jianwu Dang*

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 7(Ethical Considerations) discusses potential risks such as misuse for deepfakes and dataset biases.*

### 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?

*Section 7 (Ethical Considerations) addresses the risks of impersonation and bias in internet-collected data. Section 4.1 describes the data collection.*

- 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 reports the scale (50K hrs speech, 20K hrs music, 1.5M SFX clips) and distribution of the data.*

### C. Did you run computational experiments?

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

*Section 4.2 detail the model parameters (1.34B), hardware (32 A800 GPUs), and training hyperparameters.*

- 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?

*Tables 2 and 4 report subjective metrics*

*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.?

*Appendix A.3 Subjective Evaluation Details*

- 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)?

*We clarified that participants were volunteers recruited from our QA team.*

- 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 7 (Ethical Considerations)*

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

*Standard subjective evaluation of audio quality does not typically require IRB approval.*

**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?

*AI assistants were used to polish the writing and structure of the paper.*