

## Responsible NLP Checklist

Paper title: *Learning Dynamic Representations and Policies from Multimodal Clinical Time-Series with Informative Missingness*

Authors: *Zihan Liang, Ziwen Pan, Ruoxuan Xiong*

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 6 (Conclusion) and Section 7 (Limitations)*

### 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?

*No. All experiments use publicly available, fully de-identified clinical datasets (MIMIC-III, MIMIC-IV and eICU). The data do not contain personally identifying information, and no new personal data were collected or released as part of this work.*

- 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 (Experimental Setup), including dataset sizes, mortality rates, observation windows, modality coverage, and train/validation/test splits. Additional dataset details are provided in Appendix B.*

### C. Did you run computational experiments?

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

*Section 3.3, Section 4 and Appendix H/I/N/R*

- 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?

*Yes. Descriptive statistics are reported in Section 4 and Tables 13.*

*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.?  
*(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)?  
*N/A. This work uses existing, publicly available, de-identified clinical datasets (MIMIC-III, MIMIC-IV and eICU). Consent procedures were handled by the original data collection protocols and are described in the respective dataset documentation.*

D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?  
*N/A. No new data collection involving human subjects was conducted in this study. All analyses were performed on existing, de-identified datasets that were previously approved or deemed exempt by the original ethics review boards.*

**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?  
*Appendix, Section Use of AI Assistants.*