

Responsible NLP Checklist

Paper title: *If Only My CGM Could Speak: A Privacy-Preserving Agent for Question Answering over Continuous Glucose Data*

Authors: *Yanjun Cui, Ali Emami, Temiloluwa Prioleau, Nikhil Singh*

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?

Limitations and Ethical Considerations sections

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 3.3 (Real User Questions) and Appendix F. We did de-identification, normalization, and instantiation.

- 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 Dataset and Appendix J. We reported public dataset used, and generated QA benchmark statistics.

C. Did you run computational experiments?

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

Section 4 (discussed agent layer setup) and Section 5.1 (models setup with temperature for each model).

- 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, in Table 3 - 7; Appendix, in Table 9-11, 14

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 I (Task Instructions for participants)

- 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)?

Appendix I (Participant Requirements and Recruitment, and Compensation)

- 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 3.3 (Real User Questions), Appendix I (Ethics and Protocol, Task Instructions), Appendix F (Instantiation)

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

Appendix I (Ethics and Protocol)

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?

We used AI-based tools to assist with language editing and grammar checking of the manuscript. The AI tools were not used to generate scientific content, experimental results, or claims. All content was reviewed and verified by the authors.