

## Responsible NLP Checklist

Paper title: *Anonpsy: A Graph-Based Framework for Structure-Preserving De-identification of Psychiatric Narratives*

Authors: *Kyungho Lim, Byung-Hoon Kim*

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?

*Potential risks are discussed in the Limitations section*

### B. Did you use or create scientific artifacts? (e.g. code, datasets, models)

- N/A* 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?

*(left blank)*

- 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?

*Yes. Dataset size, exclusions, and evaluation subsets are reported in Section 3.3 (Dataset) and Section 4 (Experiments), with additional statistics and full GPT-based evaluation results provided in Appendix D.5D.8.*

### C. Did you run computational experiments?

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

*Yes. The experimental setup, including datasets, baselines, evaluation protocols, and model configurations, is described in Sections 3 and 4 of the main paper and in the Appendix (Section D, Experimental Details), with decoding hyperparameters reported in Appendix C (Hyperparameter Settings).*

- 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 including mean performance metrics (e.g., soft-F1, cosine similarity, and risk scores) are reported throughout the experimental evaluation (Section 4).*

---

*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?**

N/A 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)?  
*Details of recruitment and payment are reported in the expert evaluation section (Section 4.1).*

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)?  
*(left blank)*

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

**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?  
*Yes. Information about the use of AI assistants is provided in Section 3.2 (LLM backbone), which describes how large language models are used as controlled components of the proposed pipeline (graph conversion, perturbation, and generation) and as tools for automated evaluation.*