

Responsible NLP Checklist

Paper title: *GraphDx: A Cost-Aware Knowledge-Enhanced Multi-Agent Framework for Sequential Diagnosis*

Authors: *Shaoting Tan, Ning Liu, Yuntao Du, Shuyue Wei, Wu Shuai, Qian Li, Yanyu Xu, Wei Zhang, Lizhen Cui, Haitao Yuan*

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 "Ethics Statement" (Page 9).

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 "Ethics Statement" confirms that the datasets (MedQA, MIMIC-IV) are publicly available and de-identified.

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 "Implementation Details" and Appendix D.2 "Datasets" (Page 12).

C. Did you run computational experiments?

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 4 "Implementation Details" and Appendix D.3 "Parameter Settings" (Page 12).

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 4 "Implementation Details" states that results are Mean Standard Deviation over 3 independent runs (Tables 2 & 3).

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.?

Section 4.1.3 describes the evaluation criteria (Diagnostic Score, Test Cost) used by the experts. Full instruction text was not included as the review relied on their professional medical judgment.

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)?

Detailed recruitment and compensation information was omitted for anonymity. The evaluators were qualified medical professionals.

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)?

The task involved reviewing model outputs, not collecting personal data from the experts.

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?

We used LLMs (e.g., ChatGPT) solely for grammatical error correction and code debugging. All scientific claims, experimental designs, and writing logic are original.