ClaraNP: Generative Nursing Interface for Nursing Education with Hallucination Mitigation
Response to Reviewers

Dear AI4ED Workshop Committee,

Thank you for giving us the opportunity to submit a revised version of the in-progress paper "ClaraNP: Generative Nursing Interface for Nursing Education and Hallucination Mitigation" for poster presentation during the AI4ED workshop. We would like to express our gratitude and thanks to your committee and the reviewers for accepting and reviewing our manuscript. Your comments and suggestions made the revisions in our present version possible. We have considered the comments and done our best to address each one. It is our hope that the changes made help our paper meet the standards of AAAI and satisfy any questions that arose during initial review. Below we provide point-by-point responses to the comments and concerns of the reviewers. Responses are written in blue, while all changes to the manuscript are highlighted in yellow. All page and location numbers refer to the revised version of our paper.

Response to Reviewer 1

[General Comment] Overall, the paper suffers from lack of details and significant results. However, the concept seems promising; the planned future work would certainly improve the strength of the paper.

[Response] Thank you for noticing the potential behind our initial efforts. Although ClaraNP is still in development, we have made an effort to introduce additional technical depth to the methodological sections of this paper in our revisions.

[Minor Comment] The paper well-motivates the need to reduce hallucination rates in LLMs when used in education settings.

[Response] Thank you for your recognition!

[Comment 1] Although bias mitigation is stated as a primary contribution of the model, only one sentence is dedicated to it.

[Response] While Reviewer 1 is correct that bias mitigation is an integral part of the goal for ClaraNP, this portion of the architecture is still in progress. We placed mentions

of bias mitigation under the "Future Work" section to demonstrate ongoing development of fairness capabilities. However in response to the concerns of Reviewer 1, we worked to more specifically outline the next steps in developing an anti-bias mechanism.

[Comment 2] The focus of the paper seems to be on the preparation of the training architecture. But few details are discussed regarding the machine learning/fine tuning methods. It is unclear to me what keyword ranking model is used, how the SNN+ cosine similarity scores are weighted, etc.

[Response] Reviewer 1 is correct in their assessment that our initial description of the learning and fine tuning methods was not adequately explained. We have elaborated on the technical aspects of our methodology in the latter half of the methodology section.

The remainder of the comments were reiterations of the first responses. We thank you for your time and effort and appreciate your feedback.