

An Automated Explainable Educational Assessment System Built on LLMs

Jiazheng Li^{1*}, Artem Bobrov^{1*}, David West², Cesare Aloisi², Yulan He^{1,3}

¹King's College London,

²AQA,

³The Alan Turing Institute

{jiazheng.li, artem.bobrov, yulan.he}@kcl.ac.uk, {dwest, caloisi}@aqa.org.uk

Abstract

In this demo, we present AERA Chat, an automated and explainable educational assessment system designed for interactive and visual evaluations of student responses. This system leverages large language models (LLMs) to generate automated marking and rationale explanations, addressing the challenge of limited explainability in automated educational assessment and the high costs associated with annotation. Our system allows users to input questions and student answers, providing educators and researchers with insights into assessment accuracy and the quality of LLM-assessed rationales. Additionally, it offers advanced visualization and robust evaluation tools, enhancing the usability for educational assessment and facilitating efficient rationale verification.

Introduction

Automated Student Answer Scoring (ASAS) systems are vital educational NLP applications that streamline the manual grading process, offering a swift and consistent evaluation of student performance (Larkey 1998; Alikaniotis, Yannakoudakis, and Rei 2016; Dong, Zhang, and Yang 2017). Traditional ASAS systems typically utilize text classifiers built on pre-trained language models (Mayfield and Black 2020; Xie et al. 2022). These models process inputs such as question, key answer elements (e.g., “*student should specify the materials to be tested in their response*”), marking rubric (e.g., “*2 points for accurately describing two additional pieces of information*”), and student responses to generate marks. However, concerns about the lack of transparency of these systems have raised questions about their reliability in real-world assessments. Various approaches have been developed to address this issue by interpreting the model marking processes. These include feature analysis (Tornqvist et al. 2023; Vanga et al. 2023; Li et al. 2023b) and visualizations of internal mechanisms, such as weights and attention (Alikaniotis, Yannakoudakis, and Rei 2016; Yang et al. 2020). However, these interpretations often require significant NLP expertise, posing challenges for educators without technical backgrounds.

The recent development of large language models (LLMs) has introduced a new approach that leverages in-context learning and reasoning capabilities (Brown et al. 2020; Wei

et al. 2022; Zhou et al. 2024) to generate natural language rationales that justify model decisions (Camburu et al. 2018; Marasovic et al. 2022; Gurrapu et al. 2023). This approach improves explainability, making it more accessible to both educators and students. However, acquiring rationale annotations is costly. Most ground truth labels used for training rationale generation models rely on noisy rationales generated by LLMs, often without human verification.

To *facilitate easy access to the latest LLM-based assessment rationales generation models* and *simplify the rationale evaluation and annotation process*, we introduce an interactive platform for Automated Explainable student Response Assessment, called **AERA Chat**. Our platform features an interactive user interface, leveraging multiple LLMs as the backend for automated assessment and rationale generation. AERA Chat is designed to advance the practical use of explainable assessment methods in educational scenarios and support research by providing a unified platform for benchmarking. To the best of our knowledge, AERA Chat is the first interactive platform designed to utilize LLMs for explainable student answer scoring. Our demo video can be found at <https://youtu.be/qUSjz-sxlBc>.

Overview of AERA Chat Functionality

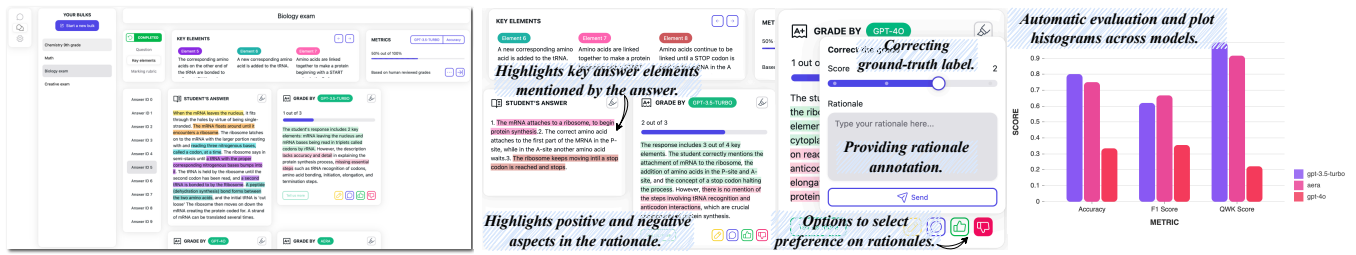
Bulk Marking Interface

As shown in Fig.1a, the bulk marking interface is our main interface to provide automated assessment and integrate with our innovative tools. Our interface has been implemented and equipped with the following key functionalities:

Bulk Question Initialization Users can set up the question, key answer elements, and a point-based marking rubric in the system. Our system automatically compiles the question information into a template prompt, which is then provided to LLMs for assessment. Users can upload a batch of student answers they want to assess in a file, which may or may not include the ground-truth marks.

Automated Scoring and Rationale Generation Once the question information and student answers have been uploaded to the system, our platform can automatically assess student answers and generate rationales with the user-selected LLMs. In this system demonstration, we used two OpenAI models: GPT-3.5-turbo and GPT-4o, as well as our developed rationale generation model – the AERA

*These authors contributed equally.



(a) Overview of Bulk Marking Interface.

(b) Key Functionalities Provided in the Bulk Marking Interface.

Figure 1: The AERA Chat interface.

model (Li et al. 2023a). Our system can easily extend to include other models for student answer assessment and rationale generation. Our database system will automatically monitor the assessment status. As shown in Fig. 1a, our platform will display rationales and marks assessed by each LLM as a card view once the assessment is completed.

Explainable Highlighting Verifying the faithfulness of the assessment rationales for student answers is a complex task that requires a detailed understanding of context. Our platform aims to enhance user experience by providing clear, high-contrast visual cues within student answers and assessment rationales. As shown in the leftmost picture in Fig. 1b, users can choose to visualize the key answer elements mentioned in the student answers or visualize the positive aspects (reasons for awarding points) and negative aspects (reasons for deducting points) within rationales. This functionality is powered by GPT-4o, which performs word-level tagging and generates context highlights in a JSON format, ensuring effective and efficient context visualization (Li et al. 2023c).

Annotation Toolkit To enable easy access for rationale evaluation and annotation, we have implemented three functionalities, as demonstrated in the middle picture in Fig. 1b. **(1) Ground-truth Label Correction:** As noted in (Li et al. 2023a), some ground-truth labels provided by publicly available ASAS dataset (Hamner et al. 2012) could be wrong. Therefore, we offer users an option to correct wrong ground-truth labels. **(2) Human Preference Annotation:** Evaluating the quality of rationales often models the qualitative evaluation task as binary preferences (Li et al. 2024b,a), where the factually correct and more detailed rationale is preferred. Consequently, our platform includes options for users to select “preferred” or “not preferred” rationales. These selections are automatically recorded in the system’s database, allowing researchers to construct preference data that can be used to leverage the reinforcement learning from human feedback techniques (Rafailov et al. 2023; Lu et al. 2024; Ouyang et al. 2022) for training LLMs for rationale generation. **(3) Rationale Annotation:** If none of the candidate rationales generated by LLMs is correct, our platform allows users to submit their assessment rationales as annotated data for future supervised fine-tuning LLMs.

Assessment Performance Evaluation As shown in the rightmost picture of Fig. 1b, our platform can automatically evaluate LLMs’ assessment performance if the user uploads their pre-assessed marks. This evaluation is visualized in a

histogram displaying Accuracy, macro F1 Scores, and QWK (Quadratic Weighted Kappa) scores.

Chat Interface

To leverage the advanced chat capabilities of LLMs and their potential to support explainable student answer scoring, our platform integrates a chat interface. This feature allows users to bring question details and rationales from the bulk marking system into interactive discussion with LLMs. Educators can use this functionality to request more detailed explanations for unclear assessment rationales, while researchers can use it to analyze incorrect assessment rationales and re-generate assessment decisions for further refinement.

System Implementation Details

AERA Chat is designed using a microservices architecture that integrates multiple LLMs through a unified web interface. We use Docker to modularize services. **Frontend:** The frontend provides a responsive web interface tailored for educators and researchers to engage with the system. It was developed using the Remix framework, which is built on React. **Backend:** The backend layer serves as the backbone of the AERA Chat platform, integrating all other layers. It operates based on a REST API, handling requests over HTTP connections without maintaining the state of each connection. This layer is implemented using the Flask framework. **LLM Services:** Users of our system can opt to employ publicly available API-based LLMs, or they may choose to develop and utilize privately trained, customized LLMs to enhance question-specific performance or ensure better data privacy. We developed a local API module allowing users to deploy privately trained models using the HuggingFace package as their private API service. **Database:** We utilize a PostgreSQL relational database to manage various data, including user profiles, assessment records (both for active tracking and background processing), and chat histories.

Conclusion and Future Work

In conclusion, we have developed an interactive platform, AERA Chat, for explainable student answer assessment via rationale generation. Our platform supports public and private LLM assessment, uses highlighting techniques for better visualization, and includes annotation functions to help educators and researchers generate more accurate data for developing trustworthy automated assessment LLMs. With a potential for wider applications to enable trustworthy automated student answer scoring, AERA Chat paves the way for future advancements of LLMs in educational settings.

Acknowledgments

This work was supported in part by the UK Engineering and Physical Sciences Research Council through KCL's Impact Acceleration Account (grant no. EP/X525571/1) and a Turing AI Fellowship (grant no. EP/V020579/1, EP/V020579/2). JL is funded by a PhD scholarship provided by AQA.

References

- Alikaniotis, D.; Yannakoudakis, H.; and Rei, M. 2016. Automatic Text Scoring Using Neural Networks. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*.
- Camburu, O.-M.; Rocktäschel, T.; Lukasiewicz, T.; and Blunsom, P. 2018. e-SNLI: Natural Language Inference with Natural Language Explanations. In *Advances in Neural Information Processing Systems*.
- Dong, F.; Zhang, Y.; and Yang, J. 2017. Attention-based Recurrent Convolutional Neural Network for Automatic Essay Scoring. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*.
- Gurrapu, S.; Kulkarni, A.; Huang, L.; Lourentzou, I.; and Batareseh, F. A. 2023. Rationalization for explainable NLP: a survey. *Frontiers in Artificial Intelligence*.
- Hamner, B.; Morgan, J.; Lynnvandev, M. S.; and Ark, T. V. 2012. The Hewlett Foundation: Automated Essay Scoring.
- Larkey, L. S. 1998. Automatic essay grading using text categorization techniques. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '98*. Association for Computing Machinery.
- Li, J.; Bobrov, A.; West, D.; Aloisi, C.; and He, Y. 2024a. AERA Chat: An Interactive Platform for Automated Explainable Student Answer Assessment. *ArXiv*, abs/2410.09507.
- Li, J.; Gui, L.; Zhou, Y.; West, D.; Aloisi, C.; and He, Y. 2023a. Distilling ChatGPT for Explainable Automated Student Answer Assessment. In *Findings of the Association for Computational Linguistics: EMNLP 2023*.
- Li, J.; Sun, Z.; Liang, B.; Gui, L.; and He, Y. 2023b. CUE: An Uncertainty Interpretation Framework for Text Classifiers Built on Pre-Trained Language Models. In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*.
- Li, J.; Xu, H.; Sun, Z.; Zhou, Y.; West, D.; Aloisi, C.; and He, Y. 2024b. Calibrating LLMs with Preference Optimization on Thought Trees for Generating Rationale in Science Question Scoring. In *Findings of the Association for Computational Linguistics: EMNLP 2024*.
- Li, J.; Zhao, R.; Yang, Y.; He, Y.; and Gui, L. 2023c. Over-Prompt: Enhancing ChatGPT through Efficient In-Context Learning. In *RO-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models*.
- Lu, J.; Li, J.; An, S.; Zhao, M.; He, Y.; Yin, D.; and Sun, X. 2024. Eliminating Biased Length Reliance of Direct Preference Optimization via Down-Sampled KL Divergence. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.
- Marasovic, A.; Beltagy, I.; Downey, D.; and Peters, M. 2022. Few-Shot Self-Rationalization with Natural Language Prompts. In *Findings of the Association for Computational Linguistics: NAACL 2022*. Association for Computational Linguistics.
- Mayfield, E.; and Black, A. W. 2020. Should You Fine-Tune BERT for Automated Essay Scoring? In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Gray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Tornqvist, M.; Mahamud, M.; Mendez Guzman, E.; and Farazouli, A. 2023. ExASAG: Explainable Framework for Automatic Short Answer Grading. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*. Association for Computational Linguistics.
- Vanga, R. R.; Sindhu, C.; Bharath, M. S.; Reddy, T. C.; and Kanneganti, M. 2023. Autograder: A Feature-Based Quantitative Essay Grading System Using BERT. In *ICT Infrastructure and Computing*.
- Wei, J.; Bosma, M.; Zhao, V.; Guu, K.; Yu, A. W.; Lester, B.; Du, N.; Dai, A. M.; and Le, Q. V. 2022. Finetuned Language Models are Zero-Shot Learners. In *International Conference on Learning Representations*.
- Xie, J.; Cai, K.; Kong, L.; Zhou, J.; and Qu, W. 2022. Automated Essay Scoring via Pairwise Contrastive Regression. In *Proceedings of the 29th International Conference on Computational Linguistics*.
- Yang, R.; Cao, J.; Wen, Z.; Wu, Y.; and He, X. 2020. Enhancing Automated Essay Scoring Performance via Fine-tuning Pre-trained Language Models with Combination of Regression and Ranking. In *Findings of the Association for Computational Linguistics: EMNLP 2020*.
- Zhou, Y.; Li, J.; Xiang, Y.; Yan, H.; Gui, L.; and He, Y. 2024. The Mystery of In-Context Learning: A Comprehensive Survey on Interpretation and Analysis. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.