

Beyond Surface Human-Likeness: AI-Mentor Feedback Alignment, Pedagogical Adaptation, and Student Engagement in Longitudinal Learning Data

Jing Fan¹, Jiseon Kim², Kimia Abedini¹, Daniel Koch-Truhponen¹, Charles Koutcheme¹ and Juho Leinonen¹

¹Aalto University, Espoo, Finland

²KAIST Global Institute for Talented Education, Daejeon, Republic of Korea

Abstract

Automated feedback is often evaluated by comparing it with human feedback, but surface resemblance may be an insufficient proxy for pedagogical quality. This extended abstract examines AI-mentor feedback alignment in the *KAIST Cyberbridge* program, a 10-week online science program for underserved gifted elementary students. The dataset includes 920 student-week observations from 92 students, with AI-generated feedback, AI next-session guidance, mentor-refined feedback, task scores, and participation records. We compared AI and mentor feedback using string similarity, token overlap, TF-IDF cosine similarity, sentence-embedding semantic similarity, and traditional lexicon-based sentiment analysis with TextBlob and VADER. Results show low surface similarity but moderate-to-high semantic similarity, suggesting that mentors often rewrote AI feedback while preserving its core meaning. TextBlob indicated mentor feedback was more positive than AI feedback, while VADER suggested that both AI and mentor feedback were highly positive. Adaptation also differed by instructional group: Basic group students received feedback that was less similar to AI output and more positively reframed than Advanced group students. These findings suggest that automated feedback evaluation should move beyond surface human-likeness and consider semantic preservation, pedagogical adaptation, affective tone, learner level, and participation-based engagement variables.

Keywords

Automated Feedback, Feedback Evaluation, Human-AI Collaboration, Pedagogical Adaptation, Student Engagement, Generative AI, Large Language Models

1. Introduction

Automated feedback is increasingly used in education, but its pedagogical quality cannot be evaluated only by fluency, correctness, or resemblance to human feedback. Formative feedback research emphasizes that useful feedback should be actionable, understandable, supportive, and usable by learners over time [1, 2, 3]. Recent work on LLM-based educational feedback similarly highlights the need to evaluate automated feedback against pedagogical criteria rather than surface fluency alone [4, 5]. This is especially important for underserved learners, for whom feedback may also need to support confidence, belonging, and continued participation.

This extended abstract examines a human-AI feedback workflow in the *KAIST Cyberbridge* program, an online science program for underserved gifted elementary students. In this program, AI-generated feedback was first produced from student submissions, after which human mentors could keep, modify, or replace it before delivery. This setting allows us to examine automated feedback not as a standalone AI output, but as a resource that human mentors adapt for pedagogical use. Because mentors reviewed AI-generated feedback before writing or refining the final feedback, mentor feedback should not be interpreted as an independent human gold standard. Rather, our analyses examine alignment between AI-generated feedback and mentor-refined feedback within a human-AI feedback workflow. Our research questions for this work are:

PEAF 2026: The First International Workshop on Pedagogical Evaluation of Automated Feedback, June 28 2026, Seoul, South Korea
✉ jing.fan@aalto.fi (J. Fan); jskim315@kaist.ac.kr (J. Kim); kimia.abedini@aalto.fi (K. Abedini); daniel.koch-truhponen@aalto.fi (D. Koch-Truhponen); charles.koutcheme@aalto.fi (C. Koutcheme); juho.2.leinonen@aalto.fi (J. Leinonen)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

- RQ1: How closely is mentor feedback aligned with AI-generated feedback at surface, lexical, and semantic levels?
- RQ2: How does mentor feedback differ from AI feedback in affective tone?
- RQ3: How do alignment and affective adjustment vary by instructional group and mentor, and how are they associated with participation-based engagement variables?

2. Context and Data

We use anonymized¹ data from the *KAIST Cyberbridge* program which integrates AI-assisted feedback with human mentoring to support underserved gifted students in science education. In this program, AI-generated feedback is first produced based on student reflection journals, after which human mentors review the AI output and may keep, modify, or replace the feedback before delivering it to students. When generating the feedback, the AI was given information about the concept students were learning that week and instructed to write a single paragraph of feedback related to the task (what the student did), process (feedback on their approach or strategy), and self-regulation based on the student's reflection journal. The system used OpenAI GPT-4 to generate the feedback. The exact prompts are not provided due to space limitations. The dataset contains 92 students across 10 instructional sessions, yielding 920 possible student-week rows. For each session, the data include student reflections, AI-generated feedback, AI next-session guidance, mentor-refined feedback, task scores, and submission records. The original wide-format dataset was reshaped into a student-week format, and AI feedback and AI next-session guidance were combined into one AI text field for comparison with mentor feedback. Across the 920 rows, 670 contained reflections, 656 contained AI feedback, 761 contained mentor feedback, and 652 contained both AI and mentor feedback. Because task scores were generally high and showed limited variability, we treated them as secondary descriptive indicators and focused on participation-based engagement variables, including assignment submission rate, reflection participation rate, full-submission rate, and next-week reflection participation.

Table 1

Terminology and calculation of engagement variables used in the analysis.

Term used in paper	Definition and calculation
Assignment submission indicator	Student-week binary indicator for whether the student submitted the assignment in a given session.
Assignment submission rate	Mean of the assignment submission indicator, aggregated within students, groups, mentors, or the full dataset.
Reflection participation indicator	Student-week binary indicator for whether the student submitted a reflection journal in a given session.
Reflection participation rate	Mean of the reflection participation indicator, aggregated within students, groups, mentors, or the full dataset.
Assignment submission count	Student-level number of submitted assignments across the 10 sessions, ranging from 0 to 10.
Reflection count	Student-level number of submitted reflection journals across the 10 sessions, ranging from 0 to 10.
Full-submission indicator	Student-level binary indicator for whether the student submitted all 10 assignments.
Full-submission rate	Proportion of students with full-submission indicator equal to 1.
Next-week reflection participation	Student-week binary indicator for whether the student submitted a reflection journal in the following session.

¹All identifiers were removed and the text was translated from Korean to English.

3. Method

We compared the combined AI text with mentor feedback at the student-week level. Surface alignment was measured using sequence-based string similarity and token-level Jaccard overlap. Lexical alignment was measured using TF-IDF cosine similarity [6], and semantic alignment was measured using cosine similarity between multilingual sentence embeddings generated with paraphrase-multilingual-MiniLM-L12-v2, following Sentence-BERT and multilingual sentence-embedding methods [7, 8]. Affective tone was analyzed using TextBlob polarity [9] and VADER compound sentiment [10]. For each sentiment measure, we computed mentor-minus-AI differences.

Because task scores showed limited variability, we summarized participation-based engagement variables at the student, mentor, and instructional-group levels. We also compared Basic and Advanced groups to examine whether feedback adaptation varied by learner level. To examine mentor-level differences, we used Kruskal–Wallis tests [11] for similarity metrics, sentiment-adjustment metrics, assignment submission rate, and reflection participation rate because these variables were not assumed to be normally distributed across mentors. To explore associations between feedback metrics and participation-based engagement variables, we conducted exploratory correlation analyses at two levels. At the student level, we computed Pearson and Spearman correlations between student-level average feedback metrics and participation-based engagement variables, including assignment submission count, reflection count, assignment submission rate, and reflection participation rate. Because of space constraints, we report the clearest statistically significant student-level pattern in the Findings and treat the remaining correlations as exploratory checks. At the row level, we operationalized feedback modification as one minus the corresponding similarity score, so that higher values indicated greater departure from the AI-generated text, and examined whether greater modification was associated with next-week reflection participation. These analyses were exploratory, and statistical significance was evaluated using an alpha threshold of .05.

4. Findings

Table 2

Main quantitative patterns by instructional group. Sentiment values are mentor-minus-AI differences; similarity and sentiment values are computed only for student-week rows where both the combined AI text and mentor feedback were available, whereas engagement rates are computed over the relevant student-week rows.

Measure	Overall	Basic	Advanced
String similarity	.149	.078	.217
TF-IDF cosine	.394	.227	.553
Semantic similarity	.690	.561	.813
TextBlob diff.	.089	.161	.020
VADER diff.	.003	.013	-.006
Assignment submission rate	.903	.865	.947
Reflection participation rate	.728	.684	.779

Table 3

Exploratory student-level association between affective adjustment and reflection count.

Association	Coef.	<i>p</i>
Pearson: VADER diff. vs. reflection count	.243	.021
Spearman: VADER diff. vs. reflection count	.248	.018

Low surface similarity but moderate semantic similarity. Across the 652 student-week rows with both combined AI text and mentor feedback, mentor feedback showed low surface similarity to AI feedback: mean string similarity was 0.149, token Jaccard similarity was 0.435, and TF-IDF cosine

similarity was 0.394. However, mean sentence-embedding semantic similarity was 0.690. This contrast suggests that mentors often changed the wording substantially while preserving the core meaning of AI-generated feedback.

Affective tone showed mixed but informative patterns. TextBlob indicated that mentor feedback was more positive than AI feedback: mean AI polarity was 0.185, mean mentor polarity was 0.277, and the mentor-minus-AI difference was 0.089; this paired difference was statistically significant ($p < .001$). VADER produced a more conservative pattern: both AI and mentor feedback had high compound scores, and the paired mentor-minus-AI difference was only 0.003. This suggests a sentiment ceiling, where AI feedback was already highly positive and mentor adaptation may have involved subtle reframing rather than a simple shift from negative to positive.

Adaptation varied by learner level and mentor. As shown in Table 2, Basic group feedback had lower semantic similarity to AI output than Advanced group feedback (0.561 vs. 0.813) and stronger positive TextBlob adjustment (0.161 vs. 0.020). This suggests that mentors adapted AI feedback more strongly for Basic group students, whereas Advanced group feedback remained closer to the AI-generated version. Kruskal–Wallis tests indicated mentor-level differences in string similarity, TF-IDF similarity, semantic similarity, TextBlob sentiment adjustment, VADER sentiment adjustment, assignment submission rate, and reflection participation rate (all $p < .001$, $\alpha = .05$). This suggests that human intervention varied across mentors.

Engagement associations were exploratory. Task scores were high on average (mean = 95.16), suggesting a ceiling effect. Participation-based engagement variables were more informative: the overall assignment submission rate was 90.3%, the reflection participation rate was 72.8%, and the full-submission rate was 68.5%. Student-level association analyses showed limited evidence that feedback metrics were related to participation-based engagement variables. The clearest pattern was a weak positive association between mean VADER mentor-minus-AI sentiment difference and reflection count (Pearson $r = .243$, $p = .021$; Spearman $\rho = .248$, $p = .018$; Table 3).

Row-level analyses of next-week reflection participation showed only weak and inconsistent associations with the extent of AI-feedback modification. When modification was operationalized as 1–similarity, next-week reflection participation was weakly associated with TF-IDF-based modification extent (Pearson $r = .082$, $p = .046$), but this association was not robust in the Spearman correlation ($\rho = .061$, $p = .135$). Modification extent based on string similarity (Pearson $r = .056$, $p = .170$) and semantic similarity (Pearson $r = -.049$, $p = .236$) was not significantly associated with next-week reflection participation. Overall, the results do not support a simple directional claim that either closer AI–mentor alignment or greater modification predicts better participation-based engagement.

5. Discussion

These findings suggest that human-likeness is a multi-level construct. Low surface similarity should not be interpreted as evidence that AI feedback was irrelevant or unusable: mentors may preserve the instructional meaning of AI feedback while changing wording, tone, emphasis, or readability. Automated feedback evaluation should therefore distinguish surface resemblance from semantic preservation and pedagogical adaptability.

The results also suggest that feedback evaluation should account for learner context. Basic group students received feedback that was less similar to AI output and more positively adjusted, whereas Advanced group feedback remained closer to the AI-generated version. Engagement analyses were exploratory: student-level reflection count showed the clearest association with affective adjustment, while row-level next-week reflection participation showed only weak and inconsistent associations with feedback modification. Overall, these findings suggest that evaluation frameworks should consider

at least three dimensions: semantic preservation, pedagogical adaptation, and learner engagement as reflected in participation-based variables.

6. Limitations and Future Work

This study is observational and should not be interpreted causally. Task scores showed limited variability, limiting their usefulness as learning outcome measures. TextBlob and VADER are coarse sentiment indicators and may not capture culturally specific encouragement or pedagogical warmth. In addition, the original Korean feedback was translated into English before analysis, which may have affected similarity and sentiment scores. Combining AI feedback and AI next-session guidance into one AI text field may also partly reflect differences in text scope rather than mentor adaptation alone if mentors decided not to include next-session guidance in their feedback. Finally, instructional group and mentor identity were confounded because mentors worked with only one instructional group, so Basic/Advanced differences may partly reflect mentor-specific feedback style rather than learner level alone. Future work should examine feedback examples in depth and develop richer outcome measures such as reflection quality, feedback uptake, help-seeking, and sustained participation.

7. Conclusion

This extended abstract shows that AI–mentor feedback alignment is not well captured by surface human-likeness alone. Mentor feedback was lexically different from AI feedback but moderately aligned at the semantic level, suggesting that mentors used AI feedback as a resource for pedagogical adaptation rather than as a final message. For automated feedback systems embedded in human mentoring workflows, the key evaluation question is not only whether AI feedback resembles the final mentor-refined feedback, but whether it can support meaningful human adaptation and learner-centered use.

Acknowledgments

This work was supported by Research Council of Finland grants #356114 and #367787.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT for language polishing. The authors reviewed and edited the content as needed and take full responsibility for the publication's content. The study also analyzes AI-generated feedback produced by GPT-4 as part of the educational feedback workflow described in the paper.

References

- [1] V. J. Shute, Focus on formative feedback, *Review of Educational Research* 78 (2008) 153–189. doi:10.3102/0034654307313795.
- [2] J. Hattie, H. Timperley, The power of feedback, *Review of Educational Research* 77 (2007) 81–112. doi:10.3102/003465430298487.
- [3] D. Carless, D. Boud, The development of student feedback literacy: Enabling uptake of feedback, *Assessment & Evaluation in Higher Education* 43 (2018) 1315–1325. doi:10.1080/02602938.2018.1463354.
- [4] E. Kasneci, K. Seßler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, U. Gasser, G. Groh, S. Günemann, E. Hüllermeier, S. Krusche, G. Kutyniok, T. Michaeli, C. Nerdel, J. Pfeffer, O. Poquet, M. Sailer, A. Schmidt, T. Seidel, M. Stadler, J. Weller, J. Kuhn, G. Kasneci, Chatgpt for good? on

- opportunities and challenges of large language models for education, *Learning and Individual Differences* 103 (2023) 102274. doi:10.1016/j.lindif.2023.102274.
- [5] J. Steiss, T. Tate, S. Graham, J. Cruz, M. Hebert, J. Wang, Y. Moon, W. Tseng, M. Warschauer, C. B. Olson, Comparing the quality of human and chatgpt feedback on students' writing, *Learning and Instruction* 91 (2024) 101894. doi:10.1016/j.learninstruc.2024.101894.
 - [6] G. Salton, C. Buckley, Term-weighting approaches in automatic text retrieval, *Information Processing & Management* 24 (1988) 513–523. doi:10.1016/0306-4573(88)90021-0.
 - [7] N. Reimers, I. Gurevych, Sentence-bert: Sentence embeddings using siamese bert-networks, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 3982–3992. URL: <https://aclanthology.org/D19-1410/>. doi:10.18653/v1/D19-1410.
 - [8] N. Reimers, I. Gurevych, Making monolingual sentence embeddings multilingual using knowledge distillation, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 4512–4525. URL: <https://aclanthology.org/2020.emnlp-main.365/>. doi:10.18653/v1/2020.emnlp-main.365.
 - [9] S. Loria, Textblob: Simplified text processing, <https://textblob.readthedocs.io/>, 2025. Python library documentation, accessed 2026-05-15.
 - [10] C. J. Hutto, E. Gilbert, Vader: A parsimonious rule-based model for sentiment analysis of social media text, *Proceedings of the International AAAI Conference on Web and Social Media* 8 (2014) 216–225. URL: <https://ojs.aaai.org/index.php/ICWSM/article/view/14550>. doi:10.1609/icwsm.v8i1.14550.
 - [11] W. H. Kruskal, W. A. Wallis, Use of ranks in one-criterion variance analysis, *Journal of the American Statistical Association* 47 (1952) 583–621. doi:10.1080/01621459.1952.10483441.