

“I understand why I got this grade”: Automatic Short Answer Grading (ASAG) with Feedback

Dishank Aggarwal^{*1}[0009–0006–8961–2768], Pritam Sil²[0000–0002–0992–6892],
Bhaskaran Raman²[0000–0002–2234–7123], and Pushpak
Bhattacharyya²[0000–0001–5319–5508]

¹ Fujitsu Research of India Private Limited, Bangalore, India

² Dept. of Computer Science and Engineering, IIT Bombay, Mumbai, India
dishank.aggarwal@fujitsu.com {pritamsil,pb,br}@cse.iitb.ac.in

Abstract. In recent years, there has been a growing interest in using Artificial Intelligence (AI) to automate student assessment in education. Among different types of assessments, summative assessments play a crucial role in evaluating a student’s understanding level of a course. Such examinations often involve short-answer questions. However, grading these responses and providing meaningful feedback manually at scale is both time-consuming and labor-intensive. Feedback is particularly important, as it helps students recognize their strengths and areas for improvement. Despite the importance of this task, there is a significant lack of publicly available datasets that support automatic short-answer grading with feedback generation. To address this gap, we introduce Engineering Short Answer Feedback (EngSAF), a dataset designed for automatic short-answer grading with feedback. The dataset covers a diverse range of subjects, questions, and answer patterns from multiple engineering domains and contains $\sim 5.8k$ data points. We incorporate feedback into our dataset by leveraging the generative capabilities of state-of-the-art large language models (LLMs) using our Label-Aware Synthetic Feedback Generation (LASFG) strategy. This paper underscores the importance of enhanced feedback in practical educational settings, outlines dataset annotation and feedback generation processes, conducts a thorough EngSAF analysis, and provides different LLMs-based zero-shot and finetuned baselines for future comparison³. The best-performing model (Mistral-7B) achieves an overall accuracy of 75.4% and 58.7% on unseen answers and unseen question test sets, respectively. Additionally, we demonstrate the efficiency and effectiveness of our ASAG system through its deployment in a real-world end-semester exam at a reputed institute. There, we achieve an output label accuracy of 92.5% along with feedback correctness and emotional impact scores above 4.5 (out of 5) on human evaluation, thus showcasing its practical viability and potential for broader implementation in educational institutions.

^{*} Work done while at IITB

³ Dataset and code are open-sourced at: <https://github.com/dishankaggarwal/EngSAF>

Keywords: LLM · Automatic Short Answer Grading · EngSAF · Feedback

1 Introduction

Technology integration in education, particularly artificial intelligence (AI), has resulted in transformative changes that have redefined traditional pedagogical approaches and assessment methodologies. Effective education relies on feedback and explanations provided during assessments to ensure quality learning outcomes [30]. In the context of education, educators use different kinds of assessments, including formative, summative, and diagnostic⁴, each of which serves a vital role in a student’s learning journey. This work deals with summative assessments, which are essential in determining whether the student has achieved the learning goals of a course once it is over. Grading summative assessments often involve grading short answers and essays, which are more complicated in nature due to the flexibility and natural language in the response. Hence, automating the grading process for such questions becomes crucial, especially in countries with extremely high student-to-teacher ratios, as it can significantly reduce instructor’s workloads and improve the assessment process. For instance, during the COVID-19 pandemic, educators faced increased stress and exhaustion due to overloaded grading burden [15,25].

This challenge can be approached as a machine learning problem, where the objective is to grade a student’s response based on how similar it is to the reference answers. However, simply assigning a score or label to a learner’s response is often inadequate in practical educational contexts [13,31]. Ahea et al. [5] highlights the value and effectiveness of feedback in improving student’s learning and standardizing teaching in higher education. Although earlier work formulated this problem as the Semantic Textual Similarity (STS) task [3,4,2,1], recent works acknowledge the importance of feedback generation along with the grading task [12]. Hence, together with the grading task and the feedback generation task, researchers have formulated the problem of automatic short answer grading (ASAG) with feedback.

1.1 Automatic Short Answer Grading (ASAG) with Feedback

The ASAG with feedback problem is as follows: given a question, a reference answer, and a student’s answer, the aim is to assign a label indicating the degree of correctness in the student’s answer compared to the reference answer and provide content-focused elaborated feedback/explanation for the same. Note that the degree of correctness is limited to only three labels here, namely correct, partially correct, and incorrect. In this problem, we only focus on questions where the answers are short, varying between a sentence and a short paragraph.

⁴ <https://www.niu.edu/citl/resources/guides/instructional-guide/formative-and-summative-assessment.shtml>

This task involves evaluating the alignment between the student’s answer and the reference answers which is expressed via the degree of correctness along with proper reasoning as to why that label is assigned.

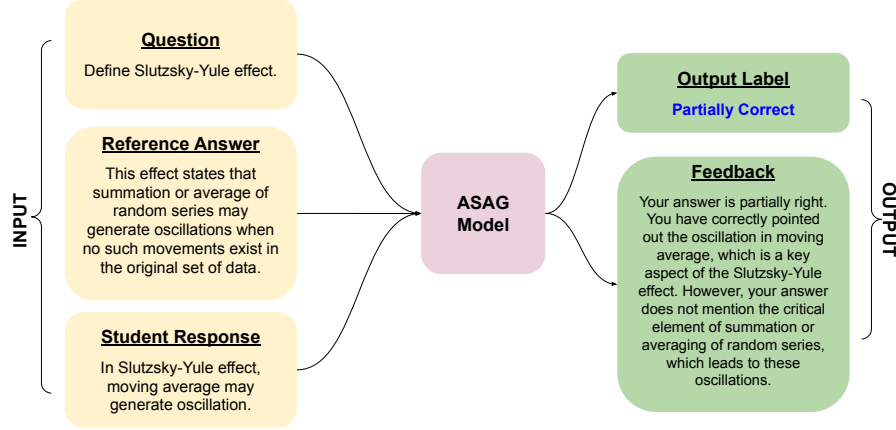


Fig. 1. Illustrative example of the ASAG with Feedback pipeline that classifies a student response as ‘correct’, ‘partially correct’, or ‘incorrect’ and provides feedback.

Figure 1 illustrates the problem statement, showing the input and output using an example from the EngSAF dataset.

1.2 Motivation

With recent advancements in state-of-the-art (SOTA) Natural Language Processing (NLP) techniques, there is an increasing interest in applying such methodologies in the field of education, particularly in short-answer evaluations. The benefits of designing such AI-assisted systems that solve the ASAG with feedback problem are two-fold:

1. The learners⁵ benefit by understanding their strengths and shortcomings from the content-elaborated feedback provided by the system.
2. The educators benefit by having prior feedback on the student’s degree of correctness, along with a proper explanation for the same. This can reduce an educator’s cognitive load while grading such short answer questions.

However, researchers face challenges when designing such systems due to the lack of public, content-centered, elaborated feedback datasets in different domains. These datasets are crucial for training and developing automated grading

⁵ We often use the terms student and learner synonymously. Also, we do the same for the terms graders, teachers, and educators.

systems that can provide feedback. Thus, this calls for a more efficient and effective dataset spanning multiple engineering domains along with proper grades and content-elaborated feedback. This is where we introduce the Engineering Short Answer Feedback (EngSAF) dataset.

Our contributions are:

1. A Label Aware Synthetic Feedback Generation (LASFG) strategy to convert conventional automatic short answer grading datasets to the ones containing feedback. (Section 3.2)
2. **EngSAF** dataset containing around **5.8K** student responses to 119 questions from multiple engineering domains taken from real-life examinations at a reputed institute along with synthetically generated feedback using LASFG strategy explaining the assigned grade for the task of ASAG. To the best of our knowledge, this is the first dataset containing questions and responses from multiple engineering domains. *The overarching goal of this work is to provide helpful feedback to students automatically.* (Section 3)
3. Benchmark scores on the EngSAF dataset using different Large Language Models (LLMs) for future comparison and research. (Section 5)
4. Real-world deployment of the best-performing fine-tuned ASAG model in an end-semester exam at a reputed institute. (Section 6)

2 Related work

This section provides a comprehensive literature review, beginning with previous research on the ASAG problem, followed by an in-depth survey of existing short-answer grading datasets.

2.1 Automatic Short Answer Grading (ASAG)

ASAG is an essential area of research that has garnered significant attention in recent years. Several approaches have been proposed for traditional ASAG techniques, ranging from rule-based methods to more sophisticated machine-learning techniques. One early approach for solving the ASAG problem was based on keyword or pattern matching, where the presence or absence of certain keywords in the student’s answer was used to determine its accuracy [20,32,23]. To overcome these limitations, researchers have developed more sophisticated methods that use NLP-based techniques. One such method is based on Latent Semantic Analysis (LSA), which represents texts as high-dimensional vectors and compares them to the reference answers using cosine similarity [17]. In a related study, the task of ASAG is addressed by incorporating features such as answer length, grammatical correctness, and semantic similarity in comparison to reference answers [33].

More recently, deep learning models such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs) have been applied to ASAG’s

task. These models are trained on large amounts of annotated data and can capture the semantic relationships between words in a student’s answer and the reference answers [34,36]. Pre-trained Language Models (PLMs), such as BERT [8], GPT [27], RoBERTa [18], DistillBERT [29], ALBERT [16], and SBERT[28] have performed exceptionally well for various NLP tasks, are widely used even for the task of ASAG on traditional ASAG datasets [7]. A more recent work by Fateen et al. [11] extends this problem by generating feedback while solving the ASAG problem using a RAG-based approach. However, a common limitation of all the prior works is that only grades are assigned, and none of the datasets provide a ground truth for the problem content-elaborated feedback generation.

2.2 Short Answer Grading Datasets

Again, when it comes to short answer grading datasets, researchers have carefully curated several publicly accessible datasets to facilitate research and benchmarking of different solutions on the ASAG task. In 2009, Mohler et al. [22] published a short answer dataset on a data structures course that contained 630 records. In 2011, the same authors [21] published an extended dataset on the same course that contained 2273 records and named it the Texas Extended dataset. Apart from this, Nielsen et al. [24] created the Beetle corpus, which contains fine-grained annotations of entailment relationships in the context of student answers to science assessment questions. Another dataset that is publicly available on Kaggle and could be used for the task of ASAG is provided by Hewlett Foundation named ASAP-AES⁶ (Automated Assessment Prize Competition for Essay Scoring).

However, as mentioned above, all datasets are conventional ASAG datasets that solely consist of grades or scores. Our extensive literature survey includes only one short-answer grading dataset with content-focused elaborated feedback, i.e., [12], which introduces an inaugural dataset for this problem and comprises bilingual responses in English and German. Human annotations were meticulously collected and refined to uphold the quality of the feedback. However, this dataset is constrained by its limited number of student responses ($\sim 2k$) in English and its exclusive focus on questions from a single domain, specifically computer network queries, thereby lacking diversity across different engineering domains. To address these limitations, we curate the EngSAF dataset, which contains $\sim 5.8k$ data points and questions from multiple engineering domains.

3 Engineering Short Answer Feedback (EngSAF) Dataset

This dataset contains 119 questions drawn from different undergraduate and graduate engineering courses, accompanied by approximately 5.8k student responses. An instructor provided correct/reference answer also accompanies each question. These questions and answers have been taken from real-life examinations at a reputed institute, thus covering a diverse range of 25 courses. The

⁶ <https://www.kaggle.com/c/asap-aes/data>

questions and responses span across multiple domains, including image processing, water quality management, and operating systems, to name a few. The dataset has been curated using a well-known e-learning platform that has been extensively used and developed at a reputed institute. Based on the instructor’s assigned marks, each response has been categorized as “*correct*,” “*partially correct*,” or “*incorrect*.”

Table 1. An example showing a question, reference answer, and three student answers (Student#1, Student#2, and Student#3) alongside their corresponding labels and synthetically generated Feedback/Explanation for the assigned label from the EngSAF Dataset

Question	What is the difference between basin order and channel order?
Reference Answer	Basin order is highest order of any stream in that basin whereas channel order is order of stream which denotes that in what order of streams has joined the channel."
Student Answer 1	Highest order channel is the basin order whereas channel order is the order of channel from tributaries to reaches to main river stream.
Label	2 (Correct response)
Feedback	Excellent! You have a clear understanding of the distinction between basin order and channel order.
Student Answer 2	Channel order reflects to the number of streams coming together to form a channel.
Label	1 (Partially Correct response)
Feedback	Your answer includes a part of the distinction. Channel order indeed indicates the number of streams joining together to make a channel, but the difference between basin order and channel order is not mentioned.
Student Answer 3	Channel order is the order of the highest order streams. For example, two first order streams (or more) will make a second order stream and similarly for highest orders.
Label	0 (Incorrect response)
Feedback	The student answer confuses basin order with channel order. Basin order refers to the highest order of streams within a basin, while channel order refers to the order of streams based on the sequence of junctions.

Table 1 provides examples from the EngSAF dataset showcasing a question, its corresponding reference answer, and three different student responses alongside their associated output labels and synthetically generated feedback. The dataset is applicable for both traditional automatic short answer grading and the generation of elaborated feedback.

3.1 Dataset Construction

The EngSAF dataset contains questions and student responses from different undergraduate and graduate engineering courses from a reputed institute. The data was collected using a well-known e-learning platform that has been developed and is extensively used at this institute. Using the e-learning platform, the instructor provided a correct answer/ reference answer to each question. Additionally, the instructors and teaching assistants (TAs) have assigned marks to each student’s answer from their respective courses. The student’s responses and grades were then obtained from the e-learning platform. The quality assessment of the proposed dataset is covered in subsection 3.4.

Table 2. Labelling scheme used for the annotation. Here, x is the marks given by the instructor/TA, and k is the maximum marks for that question.

Condition	Output Label
$x = k$	2 (Correct Response)
$0 < x < k - 1$	1 (Partially Correct Response)
$x = 0$	0 (Incorrect Response)

Each data point includes a question, student answer, and reference answer, along with a label indicating the level of correctness of the student’s response and detailed, content-specific feedback. If the maximum mark for a question is k and a student’s answer is graded x marks by the instructor, then the output label is annotated according to the scheme in Table 2.

3.2 Label Aware Synthetic Feedback Generation (LASFG)

Apart from the question, the student’s answer, and the reference answer, each data point contains a content-elaborated feedback part. To obtain this, we leverage the advanced language generative and reasoning capabilities of SOTA LLMs like Gemini⁷ and ChatGPT. The approach involves utilizing Gemini’s ability to comprehend input prompts consisting of a question, a student’s answer, a reference answer, and the corresponding grading label provided by an instructor or teaching assistant to generate content-focused elaborated feedback. The generated feedback covers the reasoning or explanation of the Gold output label. Based on initial observations of the generated feedback quality, we use `gemini-1.5-flash` model for feedback generation. To ensure the quality of the generated feedback, we performed quality estimation as described in Section 3.4.

3.3 Corpus Statistics

Following Dzikovska et al.’s approach [10], we partitioned the data into training sets, comprising 70% of the dataset, as well as unseen answers (16%) and

⁷ <https://gemini.google.com/>

unseen questions (14%) for test sets as shown in Table 3. The test split with unseen answers (**UA**) includes fresh responses to the questions used in training. In contrast, the test split with unseen questions (**UQ**) comprises entirely new questions for testing the model’s ability to generalize to new questions without prior exposure. The average text length (in tokens) for different fields in our dataset is as follows: Questions have an average length of **17.50** tokens, student answers average **25.85** tokens, reference answers are **26.47** tokens long on average, and feedback is the longest, with an average of **41.95** tokens.

Table 3. Distribution of gold label outputs across train and test split on EngSAF dataset. The test set is further subdivided into Unseen Answers (UA) and Unseen Questions (UQ).

Label	Train	UA	UQ	Total
Correct	1716	403	321	2440
Partially Correct	1412	344	278	2034
Incorrect	941	233	166	1340
Total	4069	980	765	5814

3.4 Quality Estimation

To show the reliability and credibility of our dataset and synthetically generated feedback, we randomly sampled **300 data points** and equally distributed them across the output label. This sampled data is a good representation of the EngSAF dataset. Each synthetically generated feedback was evaluated on three aspects by three proficient English-speaking annotators, all Master’s students. They were provided with detailed annotation guidelines and examples to ensure consistency.

Each aspect is scored on a scale (1-5), with a high score indicating a better response. Each feedback is analyzed on the following aspects.

1. **Fluency and Grammatical Correctness (FGC):** This aspect checks if the feedback is fluent and grammatically correct in English.
2. **Feedback Correctness/Accuracy (FC):** This aspect assesses the overall quality of the generated feedback regarding content, relevance, quality, and explanation for the assigned grade.
3. **Emotional Impact (EI):** This aspect assesses how feedback affects the learner’s emotional state. Annotators are tasked with rating whether the feedback avoids triggering negative emotions or impacts by refraining from using words such as “fail” that may evoke feelings of discouragement or distress in the learner.

Table 4. Average human annotation scores across Fluency & Grammatical Correctness (FGC), Feedback Correctness/Accuracy (FC), and Emotional Impact (EI)

Aspect	Avg. Score
Fluency & Grammatical Correctness (FGC)	4.73
Feedback Correctness/Accuracy (FC)	4.55
Emotional Impact (EI)	4.61

Three human annotators evaluate the correctness of each output label for sampled data points, achieving an accuracy of **98%** and pair-wise average Cohen’s Kappa (κ) score of **0.65** (substantial agreement), showcasing the high reliability of the assigned output label. Table 4 shows the average score over all the designed aspects to measure the reliability of the proposed dataset. Human evaluation across various designed aspects consistently yields an average score greater than **4.5** out of 5, which underscores the reliability and quality of the EngSAF dataset.

As part of our ablation studies, we have also performed an LLM-based evaluation by following the LLM-as-a-judge paradigm [38]. We conducted such an evaluation using the Gemini model to assess the quality of the generated feedback. Specifically, we instructed Gemini to evaluate feedback quality on the EngSAF test set, and we obtained average scores of **4.23** for feedback quality on unseen answers (UA) and **3.91** on unseen questions (UQ) out of 5. These scores reflect the model’s assessment of the feedback quality and provide an additional layer of validation to the reliability and quality of the dataset.

4 Experiments

Following Filighera et al.’s [12] methodology, our primary objective is to establish a baseline for our EngSAF dataset. We also assess the impact of incorporating questions on the generated feedback and the assigned labels, by performing experiments in two different settings, namely **with_question** and **without_question**. Traditionally, in Automatic Short Answer Grading (ASAG), assessments have focused solely on evaluating reference answers and student responses. However, Lv et al. [19] challenge this convention by demonstrating that integrating questions into the evaluation process improves the performance of traditional ASAG tasks.

4.1 Experimental Setting

To establish baselines for EngSAF, we have conducted experiments in two settings.

- 1. Fine-tuning Large Language Models (LLMs):** We use the Llama-2 model [35], Llama-3.1 model [9], Mistral 7b model⁸, which are fine-tuned to

⁸ <https://mistral.ai/>

predict the output label for student responses, categorizing them as correct, partially correct, or incorrect and jointly providing feedback explaining the assigned output label. Furthermore, we have conducted this experiment using two distinct methodologies.

- (a) **Without Question:** Student answer and Reference answer are passed as input.
- (b) **With Question:** Question, Student answer, and Reference answer are passed as input.

The models used in this experiment are `llama-3.1-8B-Instruct-chat`, `llama-2-13B` and `Mistral-7B-Instruct-v0.1`.

2. **Zero-shot experimentation:** For this experiment, we prompt ChatGPT⁹ and DeepSeek [14] to assign an output label along with feedback to each student’s response by evaluating its correctness compared to the reference answer for the ASAG task. The following models are used: `gpt-4o` and `DeepSeek-R1`

All experiments are performed on 2 **NVIDIA A100-SXM GPU** with **80 GB** of memory. Fine-tuning takes around 2 hours per epoch for training.

5 Results

Table 5. Fine-tuning and zero-shot experiment results on the EngSAG unseen answers and unseen questions test splits. `w_quest` models additionally received the questions as input, while `wo_quest` did not. Please note that the text similarity measures, accuracy, and F1 scores are in percent. Zero-shot experiments were conducted solely on the unseen answers test set, as both unseen questions and answers are identical in zero-shot setting.

Model	Unseen Answers						Unseen Questions					
	Acc.	F1	BLEU	MET.	ROU.	BERT	Acc.	F1	BLEU	MET.	ROU.	BERT
Majority	43.3	26.1	1.2	8.6	12.7	16.2	40.5	23.4	0.1	8.64	2.78	12.32
Fine-tuning experiments												
<i>Llama – 2_{wo_quest}</i>	74.4	73.0	13.3	34.1	16.8	31.9	55.6	53.6	12.5	31.4	16.2	28.9
<i>Llama – 2_{w_quest}</i>	73.9	73.7	11.7	35.9	16.9	35.1	56.3	54.9	9.2	31.4	13.9	32.6
<i>Llama – 3.1_{wo_quest}</i>	70.1	69.4	13.2	34.7	17.7	42.4	53.5	52.0	10.8	30.6	15.1	40.1
<i>Llama – 3.1_{w_quest}</i>	72.4	72.9	13.6	35.7	17.6	40.4	54.9	56.1	13.4	34.4	17.5	39.1
<i>Mistral_{wo_quest}</i>	72.8	73.1	11.3	33.3	14.82	37.4	54.7	55.4	11.3	32.0	15.2	36.2
<i>Mistral_{w_quest}</i>	75.4	75.7	13.9	38.3	19.5	41.6	58.7	57.9	14.9	36.7	19.7	40.9
Zero-shot experiments												
ChatGPT-4o	48.9	46.9	3.64	22.3	5.2	20.8	-	-	-	-	-	-
DeepSeek-R1	47.4	46.7	4.1	23.6	6.5	18.0	-	-	-	-	-	-

Table 5 shows a majority baseline, fine-tuning experiments, and zero-shot experiments results. The majority baseline contains the most occurring label

⁹ <https://openai.com/>

and feedback from the EngSAF train set. The most common label is “correct response,” and the most common feedback is “Well done! You have answered the question correctly, covering all the required aspects.”

5.1 Analysis

In this section, we delve into the insights drawn from the performance results of experiments on the EngSAF dataset for the ASAG task. Following Filighera et al.’s [12] evaluation methodology, we measure the accuracy and macro-averaged F1 for output labels and the ROUGE-2 [26], SACREBLEU¹⁰, METEOR [6] and BERTScore [37] for the feedback part. For all metrics, a higher score indicates better results.

As per Table 5, Mistral-7b significantly outperforms the majority baseline in both output label and feedback metrics. However, we observe a significant performance gap between unseen answers and unseen questions, with accuracy dropping by approximately **23%** for unseen questions and by **4%** in BERTScore, suggesting the necessity for new evaluation metrics that comprehensively assess text on the context level instead of the lexical level. The above two observations highlight the challenge that even fine-tuned models face when generalizing to new questions or domains (UQ). In the case of unseen answers (UA), other responses to the same question are present in the training data, making generalization easier compared to entirely new questions. A similar pattern can be observed for other fine-tuned models. Notably, including questions in the input significantly improves accuracy, F1 score, and BERTScore. For example, in our experiments with Llama-2 and Mistral-7B, incorporating the question led to a **10%** improvement in BERTScore for unseen answers and a **12%** improvement for unseen questions.

Our zero-shot experiments with ChatGPT and DeepSeek-R1 yield an accuracy of **48.9%** and **47.4%**, respectively, falling below the fine-tuning baselines, highlighting the complexity of the ASAG task. We observe a similar performance gap in the feedback evaluation part, suggesting that fine-tuning plays a crucial role in capturing complex grading patterns. Another contributing factor to this performance disparity of zero-shot experiments in comparison to fine-tuned experiments could be the inherent bias present in teacher grading, which models can only learn when trained on such data.

6 Real-World Deployment

For real-world deployment, we integrated our best-performing fine-tuned ASAG model for automatic evaluation of an end-semester exam of a course at a reputed institute. The end-semester examination included two short-answer questions, each accompanied by the instructor’s correct/reference answer. We randomly sampled 25 student answers for each question and used our fine-tuned

¹⁰ <https://pypi.org/project/sacrebleu/>

Mistral_{w_quest} model to predict the output label and feedback/ explanation for the predicted output. To ascertain the quality of the generated output, we employed **three** Subject Matter Experts (SMEs) to evaluate the predicted output and feedback. The SMEs were PhD students who were teaching assistants (TAs) of that course. Each predicted output label was evaluated based on its correctness, where the subject matter expert checks whether the predicted output label is correct or not. Each predicted feedback is analyzed in terms of **Feedback Quality/Correctness (FC)** and **Emotional Impact (EI)** as discussed in the Quality Estimate section 3.4. Each aspect is scored on a **scale (1-5)**, with a high score indicating a better response.

Table 6. Subject Matter Experts (SME) scores evaluating the real-world deployment across label accuracy, Feedback Quality/ Correctness (FC), and Emotional Impact (EI)

Aspect	Avg. Score
Output Label Accuracy	92.5%
Feedback Quality/Correctness (FC)	4.5
Emotional Impact (EI)	4.9

Table 6 presents the average scores across all the aspects used to measure the reliability of our ASAG model. Upon examining the table, it is evident that the subject matter evaluation scores are greater than **4.5** for both FC and the EI aspect. This demonstrates the reliability and effectiveness of the ASAG model in real-world scenarios. Additionally, the model’s predicted output label achieved an accuracy of **92.5%**, further showcasing its performance and reliability. By providing accurate labels and feedback at scale, our approach helps reduce the manual workload for teachers and instructors, making the grading process more efficient and scalable.

7 Conclusion

This work introduces a novel and extensive dataset for automatic short answer grading (ASAG) across multiple engineering domains. Addressing the need for a standardized evaluation platform, our dataset encompasses diverse subjects and a wide range of short-answer responses. We meticulously curated and annotated the dataset to ensure its quality, consistency, and applicability to real-world grading scenarios. We propose Label-Aware Synthetic Feedback Generation (LASFG)—a method that enables the transformation of any traditional ASAG dataset into one that incorporates meaningful synthetic feedback. Additionally, we benchmark our dataset using fine-tuned large language models (LLMs) and conduct zero-shot evaluations, providing a comprehensive assessment of its effectiveness.

Furthermore, we deployed our ASAG system in a real-world course exam at a prestigious institution, showcasing its practical utility in educational settings.

Human evaluation showcases the effectiveness of the ASAG system by achieving scores more than **4.5 out of 5** on Feedback Quality/Correctness (FC) and Emotional Impact (EI). Additionally, the model demonstrates its ability to correctly detect an answer with an accuracy of **92.5%**. This work represents a significant milestone in educational technology, offering educators and researchers a valuable tool to drive innovation in automated assessment methods.

Future extensions of this work could focus on expanding the dataset to include more complex short-answer responses, enabling broader grading scenarios. Additionally, incorporating the evaluation of diagrams and images within student responses could further enhance the system’s capabilities, paving the way for more comprehensive automated assessment solutions.

References

1. Agirre, E., Banea, C., Cardie, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Guo, W., Lopez-Gazpio, I., Maritxalar, M., Mihalcea, R., et al.: Semeval-2015 task 2: Semantic textual similarity, english, spanish and pilot on interpretability. In: Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015). pp. 252–263 (2015)
2. Agirre, E., Banea, C., Cardie, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Guo, W., Mihalcea, R., Rigau, G., Wiebe, J.: Semeval-2014 task 10: Multilingual semantic textual similarity. In: Proceedings of the 8th international workshop on semantic evaluation (SemEval 2014). pp. 81–91 (2014)
3. Agirre, E., Cer, D., Diab, M., Gonzalez-Agirre, A.: Semeval-2012 task 6: A pilot on semantic textual similarity. In: * SEM 2012: The First Joint Conference on Lexical and Computational Semantics–Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012). pp. 385–393 (2012)
4. Agirre, E., Cer, D., Diab, M., Gonzalez-Agirre, A., Guo, W.: * sem 2013 shared task: Semantic textual similarity. In: Second joint conference on lexical and computational semantics (* SEM), volume 1: proceedings of the Main conference and the shared task: semantic textual similarity. pp. 32–43 (2013)
5. Ahea, M.M.A.B., Ahea, M.R.K., Rahman, I.: The value and effectiveness of feedback in improving students’ learning and professionalizing teaching in higher education. *Journal of Education and Practice* **7**(16), 38–41 (2016)
6. Banerjee, S., Lavie, A.: Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization. pp. 65–72 (2005)
7. Condor, A., Litster, M., Pardos, Z.: Automatic short answer grading with sbert on out-of-sample questions. *International Educational Data Mining Society* (2021)
8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
9. Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al.: The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024)

10. Dzikovska, M.O., Nielsen, R., Brew, C., Leacock, C., Giampiccolo, D., Bentivogli, L., Clark, P., Dagan, I., Dang, H.T.: Semeval-2013 task 7: The joint student response analysis and 8th recognizing textual entailment challenge. In: Second Joint Conference on Lexical and Computational Semantics (* SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013). pp. 263–274 (2013)
11. Fateen, M., Wang, B., Mine, T.: Beyond scores: A modular rag-based system for automatic short answer scoring with feedback (2024), <https://arxiv.org/abs/2409.20042>
12. Filighera, A., Parihar, S., Steuer, T., Meuser, T., Ochs, S.: Your answer is incorrect... would you like to know why? introducing a bilingual short answer feedback dataset. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 8577–8591 (2022)
13. Franz, D.J., Richter, T., Lenhard, W., Marx, P., Stein, R., Ratz, C.: The influence of diagnostic labels on the evaluation of students: A multilevel meta-analysis. *Educational Psychology Review* **35**(1), 17 (2023)
14. Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al.: Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025)
15. Klapproth, F., Federkeil, L., Heinschke, F., Jungmann, T.: Teachers’ experiences of stress and their coping strategies during covid-19 induced distance teaching. *Journal of Pedagogical Research* **4**(4), 444–452 (2020)
16. Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., Soricut, R.: Albert: A lite bert for self-supervised learning of language representations. arXiv preprint arXiv:1909.11942 (2019)
17. LaVoie, N., Parker, J., Legree, P.J., Ardison, S., Kilcullen, R.N.: Using latent semantic analysis to score short answer constructed responses: Automated scoring of the consequences test. *Educational and Psychological Measurement* **80**(2), 399–414 (2020)
18. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
19. Lv, G., Song, W., Cheng, M., Liu, L.: Exploring the effectiveness of question for neural short answer scoring system. In: 2021 IEEE 11th International Conference on Electronics Information and Emergency Communication (ICEIEC) 2021 IEEE 11th International Conference on Electronics Information and Emergency Communication (ICEIEC). pp. 1–4. IEEE (2021)
20. Mitchell, T., Russell, T., Broomhead, P., Aldridge, N.: Towards robust computerised marking of free-text responses (2002)
21. Mohler, M., Bunescu, R., Mihalcea, R.: Learning to grade short answer questions using semantic similarity measures and dependency graph alignments. In: Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies. pp. 752–762 (2011)
22. Mohler, M., Mihalcea, R.: Text-to-text semantic similarity for automatic short answer grading. In: Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009). pp. 567–575 (2009)
23. Nielsen, R.D., Ward, W., Martin, J.H.: Recognizing entailment in intelligent tutoring systems. *Natural Language Engineering* **15**(4), 479–501 (2009)
24. Nielsen, R.D., Ward, W.H., Martin, J.H., Palmer, M.: Annotating students’ understanding of science concepts. In: LREC. Citeseer (2008)

25. Oliveira, G., Grenha Teixeira, J., Torres, A., Morais, C.: An exploratory study on the emergency remote education experience of higher education students and teachers during the covid-19 pandemic. *British Journal of Educational Technology* **52**(4), 1357–1376 (2021)
26. Post, M.: A call for clarity in reporting bleu scores. *arXiv preprint arXiv:1804.08771* (2018)
27. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners (2019)
28. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019)
29. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019)
30. Shute, V.J.: Focus on formative feedback. *Review of educational research* **78**(1), 153–189 (2008)
31. Sree Lakshmi, P., Simha, J., Ranjan, R.: Empowering educators: Automated short answer grading with inconsistency check and feedback integration using machine learning. *SN Computer Science* **5**(5), 653 (2024)
32. Sukkariéh, J.Z., Pulman, S.G., Raikes, N.: Auto-marking 2: An update on the ucles-oxford university research into using computational linguistics to score short, free text responses. *International Association of Educational Assessment*, Philadelphia (2004)
33. Sultan, M.A., Salazar, C., Sumner, T.: Fast and easy short answer grading with high accuracy. In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 1070–1075 (2016)
34. Surya, K., Gayakwad, E., Nallakaruppan, M.: Deep learning for short answer scoring. *Int. J. Recent. Technol. Eng.(IJRTE)* **7**(6) (2019)
35. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023)
36. Zhang, L., Huang, Y., Yang, X., Yu, S., Zhuang, F.: An automatic short-answer grading model for semi-open-ended questions. *Interactive learning environments* **30**(1), 177–190 (2022)
37. Zhang, T., Kishore, V., Wu, F., Weinberger, K.Q., Artzi, Y.: Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675* (2019)
38. Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E.P., Zhang, H., Gonzalez, J.E., Stoica, I.: Judging llm-as-a-judge with mt-bench and chatbot arena. In: *Proceedings of the 37th International Conference on Neural Information Processing Systems. NIPS '23*, Curran Associates Inc., Red Hook, NY, USA (2023)