

A LLM-Powered Automatic Grading Framework with Human-Level Guidelines Optimization

Yucheng Chu^{*}
Michigan State University
chuyuch2@msu.edu

Harry Shomer
Michigan State University
shomerha@msu.edu

Kevin Haudek
Michigan State University
haudekke@msu.edu

Hui Liu
Michigan State University
liuhui7@msu.edu

Hang Li
Michigan State University
lihang4@msu.edu

Yasemin Copur-Gencturk
University of Southern
California
copurgen@usc.edu

Joseph Krajcik
Michigan State University
krajcik@msu.edu

Jiliang Tang
Michigan State University
tangjili@msu.edu

Kaiqi Yang
Michigan State University
kqyang@msu.edu

Leonora Kaldaras
Texas Tech University
leonora.kaldaras@ttu.edu

Namsoo Shin
Michigan State University
namsoo@msu.edu

ABSTRACT

Open-text responses provide researchers and educators with rich, nuanced insights that multiple-choice questions cannot capture. When reliably assessed, such responses have the potential to enhance teaching and learning. However, scaling and consistently capturing these nuances remain significant challenges, limiting the widespread use of open-text questions in educational research and assessments. In this paper, we introduce and evaluate *GradeOpt*, a unified multi-agent automatic short-answer grading (ASAG) framework that leverages large language models (LLMs) as graders for short-answer responses. More importantly, *GradeOpt* incorporates two additional LLM-based agents—the *reflector* and the *refiner*—into the multi-agent system. This enables *GradeOpt* to automatically optimize the original grading guidelines by performing self-reflection on its errors. To assess *GradeOpt*'s effectiveness, we conducted experiments on two representative ASAG datasets, which include items designed to capture key aspects of teachers' pedagogical knowledge and students' learning progress. Our results demonstrate that *GradeOpt* consistently outperforms representative baselines in both grading accuracy and alignment with human evaluators across different knowledge domains. Finally, comprehensive ablation studies validate the contributions of *GradeOpt*'s individual components, confirming their impact on overall performance.

^{*}Both authors contribute equally.

Yucheng Chu, Hang Li, Kaiqi Yang, Harry Shomer, Yasemin Copur-Gencturk, Leonora Kaldaras, Kevin Haudek, Joseph Krajcik, Namsoo Shin, Hui Liu, and Jiliang Tang. A LLM-Powered Automatic Grading Framework with Human-Level Guidelines Optimization. In Caitlin Mills, Giora Alexandron, Davide Taibi, Giosuè Lo Bosco, and Luc Paquette (eds.) Proceedings of the 18th International Conference on Educational Data Mining, Palermo, Italy, July, 2025, pp. 31–41. International Educational Data Mining Society (2025).

© 2025 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.

<https://doi.org/10.5281/zenodo.15870201>

Keywords

Large Language Models, Automatic Grading, Mathematics Education, Teacher Knowledge, Assessments, Large-Scale Testing, Mathematical Knowledge for Teaching, Physical Science, Learning Assessments, Content Knowledge, Pedagogical Content Knowledge

1. INTRODUCTION

Accurate evaluation of assignments and examinations in a timely manner is vital to learning due to the significance of performance measurement in the learning process [38]. Traditionally, multiple-choice questions (MCQs), which asks students to select the correct answer from distracting options, dominated learning assessment studies. While this approach makes data available promptly [29, 2], it falls short in providing insights into learners' thinking. Open-ended short-answer questions (SAQs) can provide deeper insights into students' answering rationale and knowledge concepts. This is because they are known to elicit the thinking path that describes how a student arrives at their conclusion [23]. Unfortunately, grading open-ended textual answers is tedious as substantial resources and time are needed to train raters to accurately and consistently code responses [35]. More importantly, the inconsistent or unfair assessments, caused by diverged interpretations, biases, or mistakes create another challenge to SAQs grading in practice [38]. To mitigate these issues and provide timely and consistent evaluation, automatic short-answer grading (ASAG) [2] systems have become appealing. ASAG, which can be traced back to the 1960s, has bloomed in recent years due to advancements in natural language processing (NLP) [23, 41]. Early ASAG systems often used pattern-matching techniques and hand-crafted features [23]. Thus, those systems required intensive human labor to build and were limited to a few specific grading tasks. The rise of deep learning (DL) has lessened the amount of burdensome feature designs needed for early ASAG systems. DL provides an end-to-end so-

lution that automatically learns to output grading scores from a large number of graded answer samples [13]. Due to the strong data-fitting capability of DL models, DL-based ASAG systems are able to be extended to different tasks if a large number of annotated samples are available. However, when the annotated sample size is limited, DL-based ASAG systems often face serious over-fitting issues. Beyond that, as DL is a black-box model whose results lack interpretation, the application of DL-based ASAG systems is still limited [7].

The emergence of pre-trained language models (PLMs) and the more advanced Large Language Models (LLMs) have recently revolutionized the design of ASAG systems due to their human-like language ability and human-interpretable intermediate textual results. Therefore, many recent studies have attempted to build ASAG systems with LLMs. Promising results have been demonstrated that using fine-tuning [22] and prompting techniques such as Chain-of-Thought (CoT) [5] and in-context learning [24]. Yet these recent techniques are still limited due to LLMs’ inherent limitations such as sensitivity to prompts, context window restriction, etc., making the complex ASAG task challenging for the LLM grader. In reality, accurate, standardized, and unambiguous guidelines are critical to help human graders formulate a precise interpretation of scoring criteria. For LLM-based ASAG systems, those guidelines also serve as the principal instructions. They teach LLMs to perform the grading task following a similar standard as human graders. However, using guidelines composed by pedagogical experts directly for LLMs is sub-optimal since the general-purposed LLMs lack domain-specific knowledge and can misinterpret the guidelines [23]. Meanwhile, LLMs are often sensitive to various facets of prompts [18] where minor changes could lead to great differences in LLM’s performance. Optimizing the guidelines manually for LLMs can further take a lot of trial and error. Thus, recent works propose to conduct guideline modification with LLMs to offload human burden [5]. While the modified guidelines yield performance improvement, the prompt search space in these methods is relatively limited. Because of this, the modified guidelines are not necessarily optimal. Additionally, abundant human efforts such as timely feedback or a large amount of labeling are required. Therefore, methods to optimize grading guidelines automatically and effectively are still desired.

In this paper, we propose a unified **multi-agent** ASAG framework that **automatically optimizes** grading guidelines. Specifically, it employs an iterative reflection mechanism to generate task prompts (guidelines) that effectively capture learners’ thinking and knowledge from a small dataset of short answers. To achieve this, we innovatively introduce prompt optimization in ASAG, framing grading guideline refinement as an optimization problem aimed at maximizing accuracy. Inspired by APO [33], we develop novel techniques such as *misconfidence*-based selection, iterative optimization, and log-probability-based robustness to enhance the framework’s **stability** in producing accurate and **trustworthy** score predictions on unseen datasets. To minimize human labeling effort, our mechanism intelligently selects short-answer samples that contribute to optimal guideline refinement. Additionally, the framework supports assessments across varying levels of complexity, offering interpretable evaluations for

each learning objective while improving overall scoring accuracy. To validate our approach, we conducted experiments on two real-world grading datasets. The first dataset comprises responses from high school students within a physical sciences curriculum, while the second consists of a national sample of teachers answering questions designed to assess content-specific knowledge required for teaching [8]. Experimental results demonstrate that *GradeOpt* outperforms representative baselines in both accuracy and alignment. Further analysis highlights consistent improvements in test accuracy across iterations, showcasing the framework’s ability to continuously enhance grading guidelines. To the best of our knowledge, we are the first to apply prompt optimization in ASAG by refining grading guidelines akin to generating an optimal task prompt. We believe that our multi-agent reflective mechanism can unlock the full potential of LLMs in learning analytics by providing detailed and accurate assessments while significantly reducing educators’ grading workload.

2. RELATED WORK

2.1 Automatic Short Answer Grading

Automatic Short Answer Grading (ASAG) is often treated as a text classification or regression problem in NLP studies. Here we mainly focus on classification due to its relevancy to our setting. Traditional ASAG models mainly rely on text similarity and employ classic ML classifiers. They use lexical features such as bag-of-words (BOW) [29] and TF-IDF [12], or syntactic features indicating the structure of sentences [23]. However, these methods require significant manual design, which makes them hard to be applied to new datasets. To reduce the burden of feature engineering, Deep Neural Networks (DNNs) such as Long-Short-Term-Memory (LSTM) are utilized [13], which produce superior results but suffer from limited generalizability. Pre-trained BERT-based models provide enhanced versatility through *transfer learning* including on ASAG datasets [3]. To further enhance grading accuracy, researchers have made attempts to ensemble BERT with statistics-based methods [10] and data augmentation [26]. LLMs are increasingly utilized in ASAG and similar assessment tasks [44, 5, 4]. However, their prompts are mostly always manually-crafted and thus are unable to properly adapt to new datasets. To solve this issue, several works have shifted attention to assisting educators with guideline creation [41, 5].

2.2 LLM Prompt Optimization and Reflection

Prompts are critical to the success of LLMs [45]. To tailor LLMs to challenging tasks, manually crafted prompts are adopted to enhance the performance [40]. To automate the generation and optimization of prompts, prompt optimization emerges as a promising method for input prompt refinement.

Using these techniques, LLMs have demonstrated superior performance in many down-stream tasks, particularly in instruction following and reasoning [33, 45, 43]. However, such automatic methods are risky when directly applied to ASAG tasks considering the limitations of LLMs such as hallucination [17] and misalignment [21]. To enhance both accuracy and trustworthiness, we adopt the idea of state-of-the-art prompt optimization APO [33] and implement novel tech-

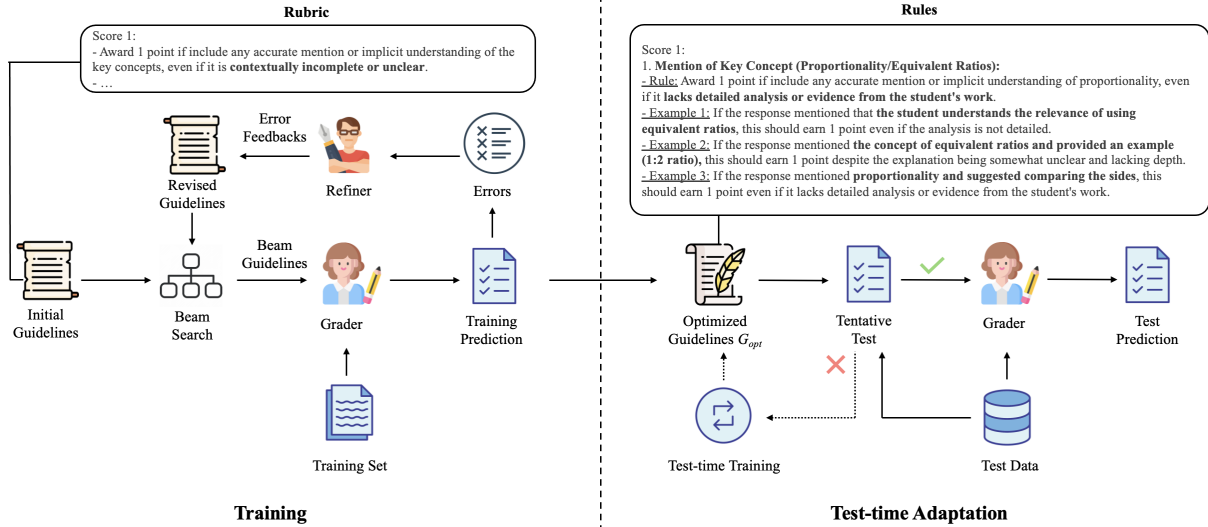


Figure 1: An Illustration of the proposed framework.

niques for reliability. Similar to how humans gather knowledge from failures, experience-and-reflect [31] is an important technique for improving LLMs’ alignment with task specifications. By reflection, LLMs learn through failure, which enriches its knowledge base and provides valuable reference in similar scenarios. Self-reflection has demonstrated promising results in improving LLM reasoning [37, 28]. However, LLMs’ reflection ability is relatively limited when it comes to self-correction without human feedback or true labels [16]. A recent work [39] divides the task of self-correction into two steps: mistake finding and output correction. They empirically show that while LLMs struggle to find errors, their correction ability is robust when given ground-truth labels. This provides grounding support for our proposed framework due to the similar use of true labels in guiding LLM reflection.

3. PROBLEM STATEMENT

We define ASAG as a text classification task, which grades the short answer text x_i by classifying it into the discrete score categories $\{\hat{y}_i = \mathcal{F}(x_i) = s_j | j = 1, \dots, C\}$, where \mathcal{F} is a ASAG system, \hat{y}_i is the score prediction, s_j is the score category, and C is size of the score category set. When \mathcal{F} is an LLM, the grading guideline text G will be concatenated at the front of x_i as an instructional prompt, and the grading process can be expressed as $\hat{y}_i = \mathcal{F}(G, x_i) = s_j$. In this work, we focuses on leveraging the reflection and refining capabilities of LLMs to automatically generate an optimized grading guideline G^* based on a small amount of graded short answer text $\mathcal{D} = \{(x_i, y_i) | i = 1, \dots, N\}$, where N is the number of graded samples. The goal of our framework can be expressed as: $G^* = \underset{G}{\operatorname{argmax}} \frac{\sum_{i=1}^N \mathbb{1}_{y_i = \hat{y}_i}}{N}$, $G \in \mathcal{G}$, where \mathcal{G} is

the potential grading space, $\mathbb{1}_{\{\cdot\}}$ is an indicator function that is 1 if $y_i = \hat{y}_i$ and 0 otherwise. Once the optimization process is finished, our framework will concatenate the optimized guidelines G^* at the front of unlabeled short answer text and generate the grading results, $\hat{y} = \mathcal{F}(G^*, x_i)$.

4. METHOD

In this section, we introduce our unified multi-agent ASAG framework *GradeOpt*. It can automatically optimize the grading guidelines and achieve better grading alignment with human experts. Next, we first give an overview of *GradeOpt*. Then, we detail the LLM-based agent design, and implementation details.

4.1 An Overview

As demonstrated in Figure 1, *GradeOpt* consists of two stages: **training** and **test-time adaptation**. The training stage is supported by three LLM-based agents: **Grader**, **Reflector**, and **Refiner**. They synergically enhance the grading guidelines by optimizing the score classification accuracy using the graded answers to the SAQs \mathcal{D}_{train} (i.e., the training data). In the test-time adaptation stage, the system first performs an out-of-distribution (OOD) test over a small amount of unlabeled answers sampled from the test data. To be specific, by checking the log likelihood score of the predicted grading results, *GradeOpt* decides whether the optimized guidelines G_{opt} can be applied to the test data directly. If the test failed, the current guideline is not optimal for the test data. Therefore, our framework will improve G_{opt} via test-time training. If the test successes, *GradeOpt* will perform the auto-grading over the whole test data automatically.

4.2 Training Stage

The training stage is to optimize the guideline for the *Grader* agent to achieve the optimal grading performance over the training dataset \mathcal{D}_{train} . *GradeOpt* leverages a multi-agent framework powered by three agents which collaboratively predict scores for \mathcal{D}_{train} , identify errors, and suggest rule modifications to mitigate errors.

Before diving into the details of this stage, we first give a brief introduction to the three key components of a common grading guideline: Question Stem (G_{qs}), Key Concept (G_{kc}) and Scoring Rubric (G_{sr}). Specifically, G_{qs} contains the complete question contents, G_{kc} describes the test knowl-

Task Description: You are GradeGPT. You assess teachers' knowledge of students' mathematical thinking by grading their responses to a pedagogical content knowledge question.

Question Stem: Based on the student's work, what is student likely to understand about the relationship between the length of the shadow and the height of the object?

Key Concept: Teachers should infer that the student possibly understands that there is a proportional relationship between the height of an object and the length of its shadow. However, because the concept of halving/doubling is natural to students, it is unclear if the student understands the relationship between object height and shadow length is proportional or if they understand equivalent ratios.

Scoring Rubrics:

- Award 0 points if it does not address key concept ...
- Award 1 point if the response includes an accurate mention or implicit understanding of the key concept ...
- Award 2 points if the response offers a clear and explicit analysis of the proportional relationship between the objects ...

Adaptation Rules:

1. Mention of Key Concept:
 - Rule: Award 1 point if the teacher response includes any accurate mention or implicit understanding of the Key Concept (proportionality/equivalent ratios), even if it lacks detailed analysis or evidence from the student's work.
 - Example 1: If the response mentions that the student understands the relevance of using equivalent ratios, this should earn 1 point even if the analysis is not detailed.

Figure 2: An example of the optimized guidelines, G^* .

edge concepts, and G_{sr} is the operational guidance instructing human graders how to score responses. As we previously mentioned, directly using $G_0 = \{G_{qs}||G_{kc}||G_{sr}\}$ as grading guideline for *Grader* is sub-optimal since the human-based scoring rubrics commonly lack detailed explanations to some concepts. As a result, LLM-based grading methods could provide ambiguous judgments. To solve this issue, *GradeOpt* focuses on optimizing G by appending new Adaptation Rules (G_{ar}) that provides the detailed explanations regarding reflections from failed predictions and identified errors. In Figure 2, we present an example of optimized grading guideline G_{opt} . Specifically, when given the expert-designed input containing "Task Description", "Question Stem", "Key Concept" and "Scoring Rubrics", *GradeOpt* automatically generates the additional descriptions in "Adaptation Rules". These new rules help describe how to assign a grade based on answer patterns and details.

The training procedure is shown on the left sub-figure of Figure 1. During training, the optimization is conducted in an iterative manner. In the t -th round, *GradeOpt* first draws a batch of samples b from \mathcal{D}_{train} and sends them to the grader agent for grading. *GradeOpt* compares the grades outputted by LLMs with human-annotated scores, then identifies error samples. These samples are then sent to the reflector agent for error reflections. Based on the reflections generated from those error samples, the reflector agent proposes a series of suggestions for improving G_{t-1} , represented by ΔG_t . ΔG_t is then sent to the refiner agent, which fuses G_{t-1} with ΔG_t and generates G_t for the next iteration of optimization. Next, we will introduce detailed designs of the three agents in *GradeOpt*. Then, we will present the implementation details of the iterative optimization process.

4.2.1 Agent Configurations

Grader. The *Grader* focuses on mapping x_i to s_j based on the given G . In *GradeOpt*, we leverage the exceptional instruction-following capability of LLMs by using a prompt to instruct LLMs to simulate the grading process of human graders. To fully exploit the potential of LLMs, we incorporate the prompt engineering strategy Chain-of-Thought [40]. This encourages LLMs to provide both judgment and intermediate reasoning steps in their outputs. With such design, the *Grader* becomes better aligned with the human-like grading process. Meanwhile, the intermediate reasoning steps provide support for the *Reflector* to discover the potential improvements to the given guideline. The prompt for the *Grader* agent is shown in Figure 3.

Grader Prompt

Task Description: In this task, you perform the task of assessing teachers' knowledge of students' mathematical thinking by grading teacher's response to a math teaching question.
Question Stem: <question stem>
Key Concept: <key concept>
Scoring Rubrics: <scoring rubrics>
Adaptation Rules: <adaptation rules>
Output format
 <score>
 Reasoning: <reasoning>
Output Rules
 1. Replace <score> with only one integer from 0, 1, or 2.
 2. Replace <reasoning> with your reasoning.
 Let's think step by step!

Figure 3: An example of the prompt to *Grader*.

Reflector. The role of *Reflector* is to propose ways to improve the current guideline G_{t-1} by reflecting over the error samples returned by *Grader*. To be specific, we design a two-step instruction prompt for LLMs to achieve this goal. In the first step, LLM is instructed to analyze the individual and shared failure reasons for a set of error samples. Then, in the second step, we ask LLMs to propose suggestions that can help resolve those issues. In general, the two-step improving process is analogous to the gradient descent algorithm used by parameter optimization for machine learning algorithm [36]. In our case, the guideline G serves as the parameter of *Grader* and identifying the error reason is similar to the "gradient". Finally, proposing improving suggestions based on discovered reasons is similar to making a descent down the "gradient" and thus optimizing G_{t-1} . The prompt for the *Reflector* agent is shown in Figure 4.

Refiner. The role of *Refiner* is to generate a new guideline G_t based on the suggestions from *Reflector*. Specifically, *Refiner* is asked to make modifications to the examples and illustrations to the content in G_{ar} . Such edits include adding, removing, or editing. Note that we keep the other components, i.e., G_{qs} , G_{kc} , G_{sc} unchanged since they are composed by human experts, and any small change may distort the scoring logic away from its original design. The refined guideline can be expressed as $G_t = \{G_{qs}||G_{kc}||G_{sc}||G_{ar}\}$, where $||$ is the text concatenation operator. The prompt for the *Refiner* agent is given in Figure 5.

4.2.2 Iterative Optimization Designs

Nested Iteration . The high complexity of test questions and grading guidelines makes it nontrivial to implement the optimization directly. Beyond that, the constraint over the input context window size of LLMs forbids it to accept all examples in \mathcal{D} for processing at once. To resolve that, we propose a nested iterative optimization approach, i.e., inner and outer loop, in *GradeOpt*. Specifically, during the t -th outer loop, *GradeOpt* selects a batch of samples b_{out} from \mathcal{D}_{train} and sends them with G_{t-1} to *Grader* for grading. Then, the wrongly graded answers e_t are filtered for reflections. However, due to the input context window size limitation, all errors in e_t cannot be entirely processed by *reflector* and *refiner* simultaneously. Thus, we introduce the inner loop, which samples an inner batch b_{in} from e_t , and updates G_{t-1} with the iterative procedure.

Reflector Prompt

You are ReflectorGPT, a helpful AI agent capable of reflecting on [adaptation rules] that is used by a classifier for a grading task. Your task is to reflect and give reasons for why [adaptation rules] have gotten the given examples in [failed examples] wrong.

The prompt contains two components: 1. [question stem], [key concept] and [scoring rubrics] (these three are given by experts and should not be modified); 2. [adaptation rules] (your task to modify).

Important Steps For Devising Rules: Read [failed examples]. For each one of the errors, perform the following steps:
 - Step 1: Explain why the classifier made the mistakes, and provide detailed, explanative analyses for why this teacher response should not be interpreted in that wrong way.
 - Step 2: Devise or modify [adaptation rules] for each mistake to help classifier effectively avoid the mistake and classify the teacher response into the correct category (label). Make sure the devised rule is explanative, straightforward, detailed, concise, and in 1 to 3 sentences.

Question Stem: <question stem>
Key Concept: <key concept>
Scoring Rubrics: <scoring rubrics>
Adaptation Rules: <adaptation rules>

But [adaptation rules] gets the following examples wrong:
Failed Examples: <errors>

Give reasons for why [adaptation rules] could have gotten the examples wrong. *Let's think step by step!*

Figure 4: An example of the prompt to *Reflector*.

To accelerate the optimization process and encourage a wider exploration of all possible combinations of error samples in b_{out} , we integrate the beam searching strategy [11] within both inner and outer loops. The algorithm of the nested iteration is shown in Algorithm 1. To be specific, in the w -th inner loop of the t -th outer iteration, *GradeOpt* accepts guidelines beam $G_{t,w-1} = \{g_{t,w-1}^{(k)} \mid 1 \leq k \leq K\}$ from $(w-1)$ -th inner iteration instead of a single guideline for refining (line 5). Then, during the inner iteration, each $g_{t,w-1}^{(k)}$ will be sent for reflection and refinement with L independently sampled inner batches b_{in} in a parallel manner (line 9). After all refined guidelines for the w -th inner loop are finished, $G_{t,w} = \{g_{t,w}^{(l,k)} \mid 1 \leq l \leq L, 1 \leq k \leq K\}$, each new guideline $g_{t,w}^{(l,k)}$ will be tested over a hold-out validation set

\mathcal{D}_{val} (line 14). Meanwhile, the top- K performing guidelines will be kept as $G_{t,w}$ and passed to the $(w+1)$ -th inner loop. Finally, the beam output of the last iteration of inner loop $G_{t,W}$ will be sent to the $(t+1)$ -th outer iteration (line 4).

While this procedure helps increase the accuracy and reliability, blindly increasing the iteration could lead to overfitting and higher computational overheads[19]. This is particularly true for smaller datasets. To help address these challenges, we introduce an early-stopping criteria. Specifically, during the selection for top- K performed $G_{t,w}$ in the w -th inner loop, we record the performance metric m_w of the best performed guideline. Then, in the next $(w+1)$ -th inner iteration, we check if m_{w+1} is improved. If m_w stops improving for two consecutive inner iterations, it indicates that the current guideline is facing risks to be over-fitted, thus following inner iterations are skipped. Similarly, during the t -th outer iterations, if its inner iteration is terminated due to the early-stopping and $(t-1)$ -th outer iteration's inner iteration is also terminated by early-stopping, the following outer iterations will also be skipped.

Batch Sampling Strategies . Using self-reflective approaches of LLMs to refine grading guidelines requires the exposure of similar errors in consecutive optimization iterations due to LLMs' lack of ability in generating appropriate modifications with one attempt [27]. This is especially true for complicated cases involving nuance differences between score categories. However, the randomness of batch sampling in the outer loop fails to guarantee this pre-requisite, which limits the performance of *GradeOpt*. To solve this, we develop a novel sampling strategy, which leverages the *misconfidence* metric (ψ) [42] to find challenging examples in \mathcal{D}_{train} . To be specific, given x_i as an input to *Grader* and y_i as its human grading result, we calculate $\psi_i = \frac{\max_{\hat{y}_i \neq y_i} \log P_{LLM}(\hat{y}_i | G, x_i)}{\log P_{LLM}(y_i | G, x_i)}$, where \hat{y}_i is the prediction of *Grader*. The *misconfidence* quantifies the discrepancy between the highest log probability of *Grader*'s incorrect prediction \hat{y}_i and the log probability of correct prediction y_i . Intuitively, the larger ψ indicates that the *Grader* is giving the wrong judgment with a relatively high confidence over the correct one, thereby implying that the sample is more challenging. However, calculating ψ over all $x_i \in \mathcal{D}_{train}$ is computationally expensive and cannot be directly done in each iteration. To avoid introducing the additional computing cost to the current algorithm, we only calculate ψ_i for samples in current iteration batch $x_i \in b_{out}$ and select the top- C samples as seeds to query similar samples from \mathcal{D}_{train} through embedding similarities. In this way, we simplify the selection process and ensure the consecutive appearance of the similar challenging examples between iterations. At last, to avoid the optimization being operated over the same portion of samples from \mathcal{D}_{train} all the time, we only select half batch based on *misconfidence*, and keep the another half as random samples. The detailed comparisons between the batch sampling strategies are presented in Section 5.7.

4.3 Test-time Adaptation Stage

In this stage, *GradeOpt* begins to perform the automatic grading to the large scaled unlabeled responses in test data. However, due to the diversity of language expressions existing in open-ended answers and other influence factors such as geography and time that change users' expression

Algorithm 1: Nested Iterative Prompt Optimization Algorithm

Data: training split of Dataset \mathcal{D}_{train} , validation split of Dataset \mathcal{D}_{val} , initial guidelines \mathcal{G} , outer loop iteration number T , inner loop iteration number W , parallel inner batch number L , guidelines beam size K .

Result: Optimized guidelines G_{opt} .

Initialize $G_{0,W} = \{g_{0,W}^{(k)}\} = \{\mathcal{G}\}$;

for $t \leftarrow 1$ to T do

$b_{out} \leftarrow$ sample an outer iteration batch from \mathcal{D}_{train} ;

Initialize $G_{t,0} = G_{t-1,W}$;

for $w \leftarrow 1$ to W do

for $k \leftarrow 1$ to K do

$\bar{y}_{out} \leftarrow$ generate grading results for b_{out} by *Grader*

with guideline $g_{t,w}^{(k)}$;

$e_{t,k} \leftarrow$ find error graded samples from b_{out} caused

by guideline $g_{t,w}^{(k)}$;

do $l \leftarrow 1$ to L in parallel

$b_{in}^{(l)} \leftarrow$ randomly sample an inner batch from

$e_{t,k}$;

$g_{t,w}^{(k,l)} \leftarrow$ generate optimized guideline with

inputting $b_{in}^{(l)}$ and $g_{t,w-1}^{(k)}$ to *Reflector* and

Refiner ;

$G_{t,w} = \{g_{t,w}^{(k)} \mid 1 \leq k \leq K\} \leftarrow$ select top- K performed

$g_{t,w}^{(k,l)}$ based on grading performance over \mathcal{D}_{val} ;

styles, the performance of the auto-graded is not always guaranteed to be the same as during training. Such phenomenon is well-recognized as the out-of-distribution (OOD) issue in many machine learning problems [14]. Prior work [14] has shown that capturing prediction probability statistics about correct or in-sample examples is often sufficient for detecting whether an example is in error or abnormal. Inspired by this, we compose a confidence indicator $\zeta = \frac{1}{|\mathcal{D}_{test}|} \sum x_i \in \mathcal{D}_{test} \max_j (\log P_{LLM}(s_j | G, x_i))$, where $\log P_{LLM}(\cdot)$ denotes the log likelihood probability given by the LLM. Intuitively, the log probability reflects the confidence that *Grader* gives to its graded results. By comparing ζ with the average LLM confidence scores μ on samples in \mathcal{D}_{train} , we can know how serious the OOD phenomenon is. Specifically, when $\zeta > \mu$, it indicates that G is well-applicable to \mathcal{D}' . When $\zeta < \mu$, it suggests that the guideline is facing serious OOD influences, which suggests that grader may struggle to produce reliable and accurate predictions for \mathcal{D}_{test} .

If the test samples are deemed to be OOD, a common solution is to first compose an adaption dataset from the testing scenario. Using this adaption dataset, we then perform test-time training on the existing model. To be specific, test-time training leverages the annotation samples from \mathcal{D}_{test} and fine-tunes the optimized guideline G_{opt} with the same training process introduced in Section 4.2. Unfortunately, in the ASAG scenario, the annotation is usually expensive. Besides, it is challenging to ask pedagogical experts to provide a large amount of annotation samples to help the existing system adapt to any changes in a timely manner. To solve this issue, we propose an incremental labeling approach which checks the marginal performance changes brought by gradually increasing the size of annotation samples. By selecting the size with highest marginal gains in metrics like accuracy and Kappa, *GradeOpt* only asks pedagogical experts for necessary annotations. This not only reduces the annotation work loads but also increases the adaption effi-

ciency of the framework. Finally, when the G_{opt} passes the OOD test, *GradeOpt* will be leveraged to finish ASAG over all samples in \mathcal{D}_{test} .

Refiner Prompt

I'm trying to write a classifier for a grading task. You are RefinerGPT, a helpful AI agent capable of refining [adaptation rules] to be used by the classifier.

The [adaptation rules] must contain patterns learned from failed examples, explaining why the predicted score is wrong comparing to the correct label. The set of rules must strictly abide by [scoring rubrics] and must clearly use patterns/details from examples to clearly illustrate and explain.

Question Stem: <question stem>

Key Concept: <key concept>

Scoring Rubrics: <scoring rubrics>

Adaptation Rules: <adaptation rules>

But [adaptation rules] have gotten several examples wrong, with the reasons of the problems examined as follows:

Failed Examples: <errors>

Error Feedbacks: <error feedbacks>

Based on the above information, I wrote **one** different improved set of rules in replacement of [adaptation rules] for instructing the classifier to learn patterns from examples for avoiding such errors.

Let's think step by step!

Figure 5: An example of the prompt to *Refiner*.

5. EXPERIMENT

In this section, we conduct experiments to validate the effectiveness of *GradeOpt*. Through the experiments, we aim to answer the following research questions. **RQ1:** Whether the refined guidelines based on prompt optimization match or exceed the performance of human-crafted guidelines? **RQ2:** Are the optimized guidelines applicable to new datasets of the same or similar questions? **RQ3:** How does each component contribute to the overall effectiveness of the guideline optimization system?

5.1 Datasets

To address the research questions above, we conduct experiments using two representative datasets for SAQ grading. Unlike existing ASAG studies [9, 29], which focus solely on student responses, our work extends ASAG to the grading of pedagogical answers from both students and teachers. The first dataset, \mathcal{D}_1 , consists of teachers' responses to questions designed to assess the knowledge and skills essential for teaching mathematics [8]. Since grading pedagogical answers requires a more nuanced interpretation to capture the underlying thought process, evaluating *GradeOpt* on this dataset allows us to examine its performance on more complex ASAG tasks. Specifically, \mathcal{D}_1 includes six questions addressing different aspects of teacher knowledge, with responses labeled on a three-point scale: Bad (0), Fair (1), and Good (2). The second dataset, \mathcal{D}_2 , evaluates *GradeOpt* on student responses, aligning with prior studies. It comprises 252 high school student responses to 11 assessment items within a physical sciences curriculum. These assessments measure Learning Progress (LP)-aligned scientific text-based explanations, reflecting students' ability to apply knowledge of electrical interactions in high school Physical Science [20]. Responses in \mathcal{D}_2 are graded on a binary scale: Fail (0) or

Pass (1). All grading labels in both \mathcal{D}_1 and \mathcal{D}_2 were assigned by at least two human raters. In cases of disagreement, a third rater provided the final judgment. Detailed statistics for both datasets are presented in Table 1. For our experiments, we split both datasets into training, validation, and test sets using a 7:1:2 ratio.

Table 1: Detailed statistics of different questions in both datasets \mathcal{D}_1 and \mathcal{D}_2 . The number of samples in each label category is shown as C_i .

\mathcal{D}_1			\mathcal{D}_2		
Question	Total	$C_1 / C_2 / C_3$	Question	Total	C_1 / C_2
Q_1	261	36 / 104 / 121	Q_1	252	43 / 209
Q_2	265	78 / 47 / 140	Q_2	252	123 / 129
Q_3	236	132 / 66 / 38	Q_3	252	113 / 139
Q_4	231	180 / 44 / 7	Q_4	252	183 / 69
Q_5	232	83 / 112 / 37	Q_5	252	242 / 10
Q_6	229	74 / 43 / 112	Q_6	252	244 / 8
Q_7	230	64 / 114 / 52	Q_7	252	245 / 7
Q_8	231	108 / 24 / 99	Q_8	252	227 / 25
-	-	-	Q_9	252	243 / 9
-	-	-	Q_{10}	252	210 / 42
-	-	-	Q_{11}	252	241 / 11

5.2 Baselines

We compared our model with several representative ASAG baselines. Firstly, we choose two popular non-LLM methods, i.e., SBERT [34] with Logistic Regression and RoBERTa [25] with Fine-tuning. Both of them have demonstrated strong performance in prior studies [6, 32]. In addition, we adopt GPT-4o with zero-shot prompting, referred to as GPT-4o, as another baseline. Compared with non-LLM methods, LLM’s exceptional instruction and human-like reasoning capabilities make it a powerful method when facing complicated grading cases [15]. To mitigate the manual burden of revising the guidelines in the GPT-4o setting, we implement and compare *GradeOpt* with APO [33], which is a state-of-the-art method for automatic prompt optimization tasks.

5.3 Implementations

To implement the nested iterative optimization, we set the outer batch size $|b_{out}| = 64$ and inner batch size $|b_{in}| = 8$. The outer loop iteration number $T = 5$ and the inner loop iteration number $W = 3$. We implement the beam search selection mechanism with Upper Confidence Bound (UCB) [1], where the guideline beam size $K = 4$. The evaluation metric for UCB is Cohen’s Kappa as it empirically works better than other metrics. The agents in our framework are all powered by GPT-4o [30] with zero-shot prompting. The temperature for *Grader* is set to 0.0 to decrease the randomness of the result. The temperatures for both *Reflector* and *Refiner* are set to 0.5, since we want to encourage the LLMs to be more open in exploring the error reasons and propose the improving suggestions. For each question, we run the algorithm 3 times and report the average results.

5.4 Evaluation Metrics

In this work, we use Accuracy (Acc) and Cohen’s Kappa (κ_c) as the evaluation metrics to compare the performance of different models. To be specific, accuracy measures the percentage of correct predictions across all cases, while Cohen’s Kappa measures the inter-rater alignment between model’s predictions and expert annotations, accounting for agreement by chance. For the \mathcal{D}_1 dataset, which involves

multi-class classification, we additionally utilize Quadratic Weighted Kappa (κ_w), which is particularly suitable for ordinal data as it assigns different weights to disagreements based on their magnitude.

5.5 Main Results

In this section, we address **RQ1** by comparing baseline models with *GradeOpt* on both datasets, \mathcal{D}_1 and \mathcal{D}_2 . Table 2 presents the performance of baseline models and *GradeOpt* on \mathcal{D}_1 . The results reveal several key observations. While all models achieve relatively high accuracy across questions, the Cohen’s kappa values for baseline models such as RoBERTa and SBERT on some questions are notably low, often close to zero. This indicates a poor alignment between automated and manual grading. A deeper analysis reveals that non-LLM-based models exhibit a uniform majority classification phenomenon, leading to skewed grading patterns. This suggests that LLM-based models provide more reliable grading results compared to their non-LLM counterparts. Comparing GPT-4o with prompt-optimized methods further highlights the importance of optimization. Optimized prompts consistently enhance grading performance while reducing variance across different questions. This finding confirms that directly applying raw human-provided rubrics is suboptimal, and prompt optimization is necessary to fully leverage LLMs in automatic grading. Lastly, *GradeOpt* outperforms the state-of-the-art (SOTA) automatic prompt optimization method, APO, across all questions. This result demonstrates the superior effectiveness of *GradeOpt* in improving grading performance, reinforcing its advantage over existing methods.

Table 3 presents the performance of baseline models and *GradeOpt* on dataset \mathcal{D}_2 . Consistent with the findings in Table 2, LLM-based ASAG methods consistently outperform non-LLM-based models, reaffirming the advantages of leveraging LLMs for ASAG tasks. Moreover, *GradeOpt* consistently achieves the best performance among all baselines in Table 3, demonstrating its effectiveness across different datasets. This confirms that *GradeOpt* is a generalizable framework suitable for various ASAG tasks. Comparing model performance across the two datasets reveals an expected trend: grading responses related to teacher knowledge yields lower performance metrics. This aligns with our expectations, as evaluating pedagogical knowledge is inherently more complex, requiring deeper expertise and more intricate logical reasoning. Beyond performance improvements, LLM-based ASAG methods enhance grading transparency and comprehensibility. Educators can easily interpret the LLM’s scoring rationale, which can be further broken down at the expectation level. Each key concept or criterion specified in the rubric or learning objectives is individually assessed, allowing educators to evaluate how well student responses align with specific learning goals. This level of explainability fosters greater trust in automated grading and facilitates more informed instructional decisions.

5.6 Adaptation Results

To address **RQ2**, we collect an external dataset, \mathcal{D}'_1 , containing responses to two questions, Q_1 and Q_2 , from new teachers across the nation. Specifically, we gather 1,352 responses for Q_1 and 1,364 responses for Q_2 . From these, we randomly select 100 responses per question to serve as the test-time training dataset, while the remaining responses are used for evaluation. By applying the optimized guidelines

Table 2: Performance of *GradeOpt* and Baseline Models on \mathcal{D}_1 . The best performed model of each metric is marked with bold, the second best one is marked with underline.

Question	RoBERTa	SBERT	GPT-4o	APO	GradeOpt
Accuracy (Acc)					
Q_1	0.80	0.61	0.85	<u>0.90</u>	0.92
Q_2	0.81	0.70	0.72	<u>0.89</u>	0.91
Q_3	0.76	0.76	0.75	<u>0.80</u>	0.86
Q_4	0.79	<u>0.74</u>	0.51	0.67	0.70
Q_5	<u>0.79</u>	0.69	0.64	<u>0.79</u>	0.80
Q_6	0.49	0.76	0.70	<u>0.81</u>	0.84
Q_7	0.55	<u>0.68</u>	0.51	<u>0.68</u>	0.73
Q_8	0.66	0.62	0.66	<u>0.85</u>	0.89
Cohen's Kappa (κ_c)					
Q_1	0.65	0.32	0.76	<u>0.85</u>	0.88
Q_2	0.66	0.42	0.56	<u>0.80</u>	0.85
Q_3	0.00	0.00	0.38	<u>0.51</u>	0.68
Q_4	0.00	0.00	0.09	<u>0.35</u>	0.36
Q_5	0.58	0.44	0.30	<u>0.60</u>	0.63
Q_6	0.00	0.17	0.55	<u>0.69</u>	0.70
Q_7	0.00	0.41	0.33	<u>0.50</u>	0.52
Q_8	0.37	0.29	0.48	<u>0.75</u>	0.80
Quadratic Weighted Kappa (κ_w)					
Q_1	0.70	0.31	0.82	<u>0.88</u>	0.89
Q_2	0.81	0.52	0.75	<u>0.93</u>	0.94
Q_3	0.00	0.00	0.61	<u>0.67</u>	0.76
Q_4	0.00	0.00	0.24	<u>0.41</u>	0.54
Q_5	0.59	0.32	0.56	<u>0.68</u>	0.71
Q_6	0.00	0.17	0.62	<u>0.77</u>	0.80
Q_7	0.00	0.48	0.58	<u>0.64</u>	0.70
Q_8	0.44	0.33	0.68	<u>0.84</u>	0.87

Table 3: Performance of *GradeOpt* and Baseline Models on \mathcal{D}_2 . The best performed model of each metric is marked with bold, the second best one is marked with underline.

Question	RoBERTa	SBERT	GPT-4o	APO	GradeOPT
Accuracy (Acc)					
Q_1	0.84	0.84	0.84	<u>0.96</u>	0.98
Q_2	0.57	0.55	0.76	0.86	0.86
Q_3	0.69	0.67	0.88	0.92	0.92
Q_4	0.65	0.65	0.78	<u>0.80</u>	0.82
Q_5	<u>0.94</u>	<u>0.94</u>	<u>0.94</u>	<u>0.94</u>	0.98
Q_6	0.98	0.98	0.88	0.92	0.94
Q_7	0.98	0.98	1.00	1.00	1.00
Q_8	<u>0.92</u>	<u>0.92</u>	0.47	0.76	0.78
Q_9	0.98	0.98	0.96	0.96	0.98
Q_{10}	0.88	0.88	0.73	0.86	0.88
Q_{11}	0.94	0.94	0.90	0.96	0.96
Cohen's Kappa (κ_c)					
Q_1	0.00	0.00	0.55	<u>0.83</u>	0.92
Q_2	0.05	0.00	0.52	0.72	0.72
Q_3	0.08	0.00	0.74	0.81	0.81
Q_4	0.00	0.00	0.59	<u>0.62</u>	0.65
Q_5	0.00	0.00	0.64	0.64	0.64
Q_6	0.00	0.00	0.22	<u>0.31</u>	0.38
Q_7	0.00	0.00	1.00	1.00	1.00
Q_8	0.00	0.00	0.10	<u>0.15</u>	0.17
Q_9	0.00	0.00	<u>0.48</u>	<u>0.48</u>	0.66
Q_{10}	0.00	0.00	0.34	<u>0.56</u>	0.60
Q_{11}	0.00	0.00	0.50	0.73	0.73

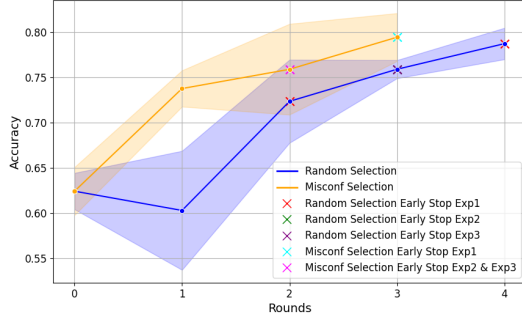
Table 4: Result for Q_1 and Q_2 over new dataset \mathcal{D}'_1 before (G_{opt}) and after (G'_{opt}) test-time training. The average confidence indicator (CI) on Pilot Dataset \mathcal{D}_1 is $\mu = -0.2$ and questions with CI ($\zeta < -0.2$) are marked with OOD.

Question	Metric	G_0	G_{opt}	G'_{opt}	Δ
Q_1	Acc	0.67	0.7	0.78	+0.08
	Kappa	0.49	0.52	0.64	+0.12
	CI (ζ)	-	-0.22	-0.17	+0.05
	OOD	-	✓	×	-
Q_2	Acc	0.75	0.82	-	-
	Kappa	0.59	0.7	-	-
	CI (ζ)	-	-0.16	-	-
	OOD	-	×	-	-


G_{opt} learned from \mathcal{D}_1 in Section 5.5, we explore its grading performance over \mathcal{D}'_1 . In addition, if the G_{opt} fails to pass the OOD test, where the confidence indicator $\zeta < \mu$, we will implement the test-time training with the train split of national dataset \mathcal{D}'_{train} . Then, the adapted guideline will be tested over the same \mathcal{D}'_{test} again. Overall, the results are reported in Table 4. From the table, we can find that Q_1 is marked with OOD Flag as its confidence indicator $\zeta = -0.22 < -0.2 = \mu$. By comparing its performance between Table 2 and Table 4, we can confirm that it suffers great performance drops. Meanwhile, the confidence indicator of Q_2 , $\zeta = -0.16 > -0.2 = \mu$, and its performance gap between \mathcal{D}_2 and \mathcal{D}'_2 is relatively smaller. These two observations indicate that the proposed confidence indicator is a valid indicator for the OOD detection purpose. On the other hand, by comparing G_{opt} with the raw guidelines provided by G_0 , we find that G_{opt} consistently outperforms G_0 . This observation indicates that even the automatically optimized guideline suffers from the OOD issue, it is still better than the raw guideline. Finally, by calculating the performance change between the guideline before and after the test-time training, we can find that *GradeOpt* is able to get adapted to the new examples with only limited available annotated examples. In addition, the performance of guidelines after the test-time training is restored back to the acceptable grading range, e.g., Kappa > 0.6 , which indicates that the test-time training is an effective solution to help a high-performed G_{opt} to quickly get applied to different datasets.

5.7 Ablation Studies

To answer **RQ3**, we conduct ablation studies to the nested iteration introduced in Section 4.2. We choose to experiment with Q_5 from \mathcal{D}_1 as its relatively simple rubric design of Q_5 makes the shortest guideline prompt, leaving room for *GradeOpt* to add in its reflective experience as it iteratively learns from \mathcal{D}_2 . Thus, experimenting with Q_5 can better showcase *GradeOpt*'s optimization power. With the experimental results shown in the following sections, we demonstrate the effectiveness of each component. First, we demonstrate the superiority of our *misconfidence*-based batch sampling strategy by comparing it with the random-based one. From Figure 6, we can observe that *misconfidence*-based batch sampling results are more consistent and accurate. While random selection generates optimal guidelines in 2 to 4 rounds, *misconfidence*-based selection consistently optimizes guidelines in 3 rounds. This, together with high predictive accuracy and alignment brought




(a) Accuracy Comparison




(b) Cohen's Kappa Comparison

Figure 6: Performance comparison between *GradeOpt* with *misconfidence*-based and random-based outer batch selection strategies.



(a) Outer Batch Size



(b) Inner Batch Size

Figure 7: Performance of *GradeOpt* with different outer and inner iteration batch sizes.

by *misconfidence*-based selection, makes the system reliable in practical educational scenarios as the required training round number is coherent. Then, we conduct experiments over the sizes of outer batch $|b_{out}| \in \{20, 32, 64\}$ and inner batch $|b_{in}| \in \{4, 6, 8\}$, targeting at exploring the influence of those two hyper-parameters on the performance of *GradeOpt*. From the results in Figure 5.7, we observe an increasing trend in accuracy and kappa as its outer batch size increases. This observation suggests that increasing the number of examples is always beneficial to the final performance. Similarly, from the Figure 5.7, we find an consistent increasing trend of performance as error number increases. Based on these two findings, we can conclude that the larger batch size is likely to bring performance gains to the *GradeOpt*. At last, we study how the iterations number impacts accuracy performance. In our experiment, we explore different iteration numbers, ranging from 1 to 5. Whilst increasing iteration on a minibatch, we utilize the early-stopping signal introduced in Section 4.2.2 to carefully monitor overfitting. As shown in Figure 8, increasing iteration with the help of early-stopping signal can effectively lead to higher test accuracy as well as more stable performance. While five iterations produce higher accuracy, due to limited computational resource, we use three iterations as our default setting.




Figure 8: Performance of *GradeOpt* with different numbers of iterations.

6. CASE STUDY

In this section, we conduct case study by going through the refinement history and presenting the extraordinary performance of *GradeOpt* in optimizing the guidelines to match with task specifications and training dataset. We showcase example responses of *Reflector* and *Refiner* separately.

6.1 Reflector

The *Reflector* reflects and gives feedback when provided with a small selection of error samples. For example, as shown in Table 5, given a selection of errors in the prompt, *Reflector* follows a two-step reflect-and-suggest approach. It first provides an explanation to each error, illustrating the reason for why the mistake is made and analyzing why the response should be assessed differently. Based on this reflection, *Reflector* next discusses its suggestion of rule modification, based on the evidence from error sample.

6.2 Refiner

As introduced in Section 4.2.1, *Refiner* is tasked with guideline modification given the feedback from *Reflector*. *Refiner* can edit by adding, changing, removing, etc. We first

Table 5: Case Study for *Reflector*

Example Error	Teacher Response: “They are equal proportions because they are situations during the same time and the same angle of the sun.” - Expert Label: 2; AI Prediction: 1
Reflector’s Error Feedback	<p>Step 1: Explanation of Mistake - The classifier predicted a score of 1 because the response mentions equal proportions and references conditions affecting proportionality (same time and same angle of the sun). However, it lacks specific evidence from the student’s work and a thorough interpretation of the student’s mathematical thinking. The correct label is 2 because the response does provide a clear explanation of the proportional relationship and the conditions under which it holds, which is a nuanced understanding of proportionality.</p> <p>Step 2: Rule Modification - Rule: If the response explains equal proportions and references conditions affecting proportionality, classify as 2, even if it lacks specific evidence from the student’s work.</p>

showcase an example of *Refiner* adding a rule in Table 6, in response to *Reflector*’s feedback in Section 6.1. We observe that *Refiner* takes *Reflector*’s rules suggested in “*Step 2: Rule Modification*”. The new rule explains two components for grading a certain type of answer: first, the answer pattern, by saying “*if the teacher’s response explains equal proportions and references conditions affecting proportionality, ..., even if it lacks specific evidence from the student’s work*”; second, the score assignment, by saying “*classify as 2*”. It additionally adds in details from the example by noting “(e.g., same time, same angle of the sun)” and includes the whole answer to give an elucidative grasping of to what answers this rule can be applied. To conclude, when adding rules *Refiner* is to include new types of answers, concluding their patterns and explaining word or phrase details. Next, we show an example of *Refiner* editing the guidelines. This happens when *Refiner* determines that a defective rule can be adjusted to give a better explanation or more complete details. As the example in Table 6 shows, part of the given rule is misleading, so *Refiner* revises this part. The changed rule first redefines the scenarios when this rule can be utilized, then it provides a detailed explanation by citing the answer and illustrating how the answer falls into the pattern category defined in the new rule to receive the correct score.

7. CONCLUSION

This paper explores fully automating guideline optimization to leverage LLM techniques including reflection and prompt engineering to solve ASAG tasks. We innovatively decompose the ASAG procedure into two steps: guideline optimization and grading. Specifically, we set our focus on automatic guideline optimization to avoid the manual efforts of composing a task-optimal guideline. To further prevent labeling a large amount of data, we propose a two-phase “*train and test-adapt*” procedure to maximally tune a guideline on a small training set and securely ensure this optimized output is reliable for large-scale grading. The proposed *GradeOpt* is a multi-agent guideline optimization system that iteratively leads the LLM to reflect on mistakes, learn answer patterns, and make improving modifications. Empirical experiments on two pedagogical datasets have demonstrated the effectiveness of *GradeOpt*.

Table 6: Case Study for *Refiner*

Action	<i>Refiner</i> ’s generated rules
Add	<p>“1. <i>Equal Proportions and Conditions</i>:</p> <p>- Rule: <i>If the teacher’s response explains equal proportions and references conditions affecting proportionality (e.g., same time, same angle of the sun), classify as 2, even if it lacks specific evidence from the student’s work.</i></p> <p>- Example: <i>They are equal proportions because they are situations during the same time and the same angle of the sun.</i>”</p>
Edit	<p>Before Edit: “- Award 1 point if ... <i>explicit evidence. For instance, if the response mentions that the student understands the unit rate, which is related to the concept of proportionality.</i>” (misleading statement)</p> <p>After Edit: “- Award 1 point if ... <i>explicit evidence. For example, if the response states that the student might have a limited understanding of proportionality, it should be awarded 1 point. For instance, in the response “their answer makes sense only if there is a proportional relationship between the height of the object and the length of the shadow,” teacher mentioned “limited/partial understanding of the proportional relationship” but lacked depth.</i> (ambiguity resolved)</p>

8. ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under Grant No. 1813760, No. 2405483, No. 2200757 and No. 2234015. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We thank Dr. Clare Carlson for help in revising one of the scoring rubrics.

References

- [1] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *JMLR*, 2003.
- [2] S. Burrows, I. Gurevych, and B. Stein. The eras and trends of automatic short answer grading. *IJAIED*, 2015.
- [3] L. Camus and A. Filighera. Investigating transformers for automatic short answer grading. In *AIED*, 2020.
- [4] Y. Chu et al. Enhancing llm-based short answer grading with retrieval-augmented generation. *arXiv preprint arXiv:2504.05276*, 2025.
- [5] C. Cohn, N. Hutchins, T. Le, and G. Biswas. A chain-of-thought prompting approach with llms for evaluating students’ formative assessment responses in science. In *AAAI*, 2024.
- [6] A. Condor, M. Litster, and Z. A. Pardos. Automatic short answer grading with sbert on out-of-sample questions. In *EDM*, 2021.
- [7] A. Condor and Z. Pardos. Explainable automatic grading with neural additive models. In *AIED*, 2024.
- [8] Y. Copur-Gencturk and T. Tolar. Mathematics teaching expertise: A study of the dimensionality of content knowledge, pedagogical content knowledge, and content-specific noticing skills. *Teaching and Teacher Education*, 114:103696, 2022.

- [9] M. O. Dzikovska et al. Semeval-2013 task 7: The joint student response analysis and 8th recognizing textual entailment challenge. In *SemEval*, 2013.
- [10] J. A. Erickson et al. The automated grading of student open responses in mathematics. In *LAK*, 2020.
- [11] M. Freitag and Y. Al-Onaizan. Beam search strategies for neural machine translation. *arXiv preprint arXiv:1702.01806*, 2017.
- [12] D. Gobbo et al. Gradeaid: a framework for automatic short answers grading in educational contexts—design, implementation and evaluation. *KAIS*, 2023.
- [13] S. Hassan, A. A. Fahmy, and M. El-Ramly. Automatic short answer scoring based on paragraph embeddings. *IJACSA*, 2018.
- [14] D. Hendrycks and K. Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- [15] O. Henkel et al. Can large language models make the grade? an empirical study evaluating llms ability to mark short answer questions in k-12 education. In *Learning@ Scale*, 2024.
- [16] J. Huang et al. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*, 2023.
- [17] L. Huang et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*, 2023.
- [18] Z. Jiang et al. How can we know what language models know? *arXiv preprint arXiv:1911.12543*, 2020.
- [19] G. Juneja et al. Task facet learning: A structured approach to prompt optimization. *arXiv preprint arXiv:2406.10504*, 2024.
- [20] L. Kaldaras, H. Akaeze, and J. Krajcik. Developing and validating next generation science standards-aligned learning progression to track three-dimensional learning of electrical interactions in high school physical science. *Journal of Research in Science Teaching*, 2021.
- [21] Z. Kenton et al. Alignment of language agents. *arXiv preprint arXiv:2103.14659*, 2021.
- [22] E. Latif and X. Zhai. Fine-tuning chatgpt for automatic scoring. *arXiv preprint arXiv:2310.10072*, 2023.
- [23] C. Leacock and M. Chodorow. C-rater: Automated scoring of short-answer questions. *Computers and the Humanities*, 2003.
- [24] G.-G. Lee et al. Applying large language models and chain-of-thought for automatic scoring. *arXiv preprint arXiv:2312.03748*, 2024.
- [25] Y. Liu et al. Roberta: A robustly optimized bert pre-training approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [26] J. Lun et al. Multiple data augmentation strategies for improving performance on automatic short answer scoring. In *AAAI*, 2020.
- [27] R. Ma et al. Are large language models good prompt optimizers? *arXiv preprint arXiv:2402.02101*, 2024.
- [28] A. Madaan et al. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*, 2023.
- [29] M. Mohler, R. Bunescu, and R. Mihalcea. Learning to grade short answer questions using semantic similarity measures and dependency graph alignments. In *ACL*, 2011.
- [30] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2024.
- [31] L. Pan et al. Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies. *arXiv preprint arXiv:2308.03188*, 2023.
- [32] A. Poulton and S. Eliens. Explaining transformer-based models for automatic short answer grading. In *ICDTE*, 2021.
- [33] R. Pryzant et al. Automatic prompt optimization with "gradient descent" and beam search. *arXiv preprint arXiv:2305.03495*, 2023.
- [34] N. Reimers and I. Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*, 2019.
- [35] S. Roy, Y. Narahari, and O. D. Deshmukh. A perspective on computer assisted assessment techniques for short free-text answers. In *CAA*, 2015.
- [36] S. Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2017.
- [37] N. Shinn et al. Reflexion: Language agents with verbal reinforcement learning. *arXiv preprint arXiv:2303.11366*, 2023.
- [38] N. Stützen et al. Automatic short answer grading and feedback using text mining methods. *Procedia Computer Science*, 2020.
- [39] G. Tyen et al. LLMs cannot find reasoning errors, but can correct them given the error location. In *ACL*, 2024.
- [40] J. Wei et al. Chain-of-thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2023.
- [41] W. Xie et al. Grade like a human: Rethinking automated assessment with large language models. *arXiv preprint arXiv:2405.19694*, 2024.
- [42] S. Xu and C. Zhang. Misconfidence-based demonstration selection for llm in-context learning. *arXiv preprint arXiv:2401.06301*, 2024.
- [43] C. Yang et al. Large language models as optimizers. *arXiv preprint arXiv:2309.03409*, 2024.
- [44] K. Yang et al. Content knowledge identification with multi-agent large language models (llms). In *AIED*, 2024.
- [45] Y. Zhou et al. Large language models are human-level prompt engineers. *arXiv preprint arXiv:2211.01910*, 2023.