

# How Can I Improve? Using GPT to Highlight the Desired and Undesired Parts of Open-ended Responses

Jionghao Lin  
Carnegie Mellon University  
jionghal@cs.cmu.edu

Ashish Gurung  
Carnegie Mellon University  
agurung@andrew.cmu.edu

Ngoc Dang Nguyen  
Monash University  
dan.nguyen2@monash.edu

Eason Chen  
Carnegie Mellon University  
easonc13@cmu.edu

Danielle R. Thomas  
Carnegie Mellon University  
drthomas@cmu.edu

Kenneth R. Koedinger  
Carnegie Mellon University  
koedinger@cmu.edu

Zifei Han  
University of Toronto  
feifei.han@mail.utoronto.ca

Wei Tan  
Monash University  
wei.tan2@monash.edu

## ABSTRACT

Automated explanatory feedback systems play a crucial role in facilitating learning for a large cohort of learners by offering feedback that incorporates explanations, significantly enhancing the learning process. However, delivering such explanatory feedback in real-time poses challenges, particularly when high classification accuracy for domain-specific, nuanced responses is essential. Our study leverages the capabilities of large language models, specifically Generative Pre-Trained Transformers (GPT), to explore a sequence labeling approach focused on identifying components of desired and undesired praise for providing explanatory feedback within a tutor training dataset. Our aim is to equip tutors with actionable, explanatory feedback during online training lessons. To investigate the potential of GPT models for providing the explanatory feedback, we employed two commonly-used approaches: *prompting* and *fine-tuning*. To quantify the quality of highlighted praise components identified by GPT models, we introduced a Modified Intersection over Union (M-IoU) score. Our findings demonstrate that: (1) the M-IoU score effectively correlates with human judgment in evaluating sequence quality; (2) using two-shot prompting on GPT-3.5 resulted in decent performance in recognizing effort-based (M-IoU of 0.46) and outcome-based praise (M-IoU of 0.68); and (3) our optimally fine-tuned GPT-3.5 model achieved M-IoU scores of 0.64 for effort-based praise and 0.84 for outcome-based praise, aligning with the satisfaction levels evaluated by human coders. Our results show promise for using GPT models to provide feedback that focuses on specific elements in their open-ended responses that are desirable or could use improvement.

J. Lin, E. Chen, Z. Han, A. Gurung, D. R. Thomas, W. Tan, N. D. Nguyen, and K. R. Koedinger. How can i improve? using gpt to highlight the desired and undesired parts of open-ended responses. In B. Paaßen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 236–250, Atlanta, Georgia, USA, July 2024. International Educational Data Mining Society.

© 2024 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.12729804>

## Keywords

Generative Artificial Intelligence, Large Language Models, Tutor Training, Feedback, GPT, Sequence Labeling

## 1. INTRODUCTION

Tutoring is an important instructional method that can be highly effective in supporting students. Tutors utilize various tutoring strategies to effectively facilitate learning opportunities [32, 48, 39]. While the effectiveness of tutoring is widely recognized, various logistical challenges have restricted its widespread implementation. Specifically, recruiting, training, and retention of tutors have presented major hurdles [55]. Training tutors can be highly resource-intensive and often requires hands-on training from experienced tutors. A key component of effective tutor training involves helping novice tutors learn effective tutoring strategies [40, 55]. For instance, instead of simply pointing out an incorrect answer, effective tutors often engage with the student to identify the underlying misconceptions or gaps in knowledge that can provide additional context to the incorrect answer. This contextual insight can then assist the tutor in providing more effective support. Traditionally, these types of nuanced insights have been facilitated through hands-on training from more experienced tutors. However, the scalability of this hands-on approach remains a well-recognized limitation [40, 24, 27, 38], necessitating innovative solutions to extend this model of training tutors without compromising the quality of feedback.

In response to the growing need for scalable hands-on support in tutor training, researchers have increasingly turned to automated feedback systems. The integration of such systems to enhance feedback is well-established within Educational Data Mining (EDM), with numerous studies demonstrating their efficacy [37, 2, 17, 50]. While many implementations have employed AI algorithms to generate automated feedback [5], the specific application to tutor training remains underexplored. In this emerging field, the development of automated explanatory feedback systems designed for tutors presents a promising avenue. An illustrative example includes work by [40], which utilized the BERT language model [12] to enhance tutor training. Although the re-

sults showed potential, a significant challenge emerged: The BERT model was hampered by a lack of access to extensive datasets, limiting its ability to offer precise, context-specific feedback. This challenge is similarly problematic for other traditional models such as Conditional Random Fields (CRF) and Hidden Markov Models (HMM), which also require adequate domain-specific training data [47, 43, 46]. Recent advances in large language models (LLMs) present a viable solution to these challenges. LLMs, such as Generative Pre-trained Transformers (GPT) developed by OpenAI, are pre-trained on vast and diverse datasets, enabling them to generalize effectively across different domains without extensive task-specific data. The inherent adaptability of GPT models to dynamically adjust to specific contextual scenarios makes them well-suited for developing real-time, tailored feedback systems for tutor training—offering the adaptive, hands-on support that models like BERT could not.

By referencing recent LLM literature [26, 51, 28], we explore two approaches to leverage the potentials of GPT models in educational contexts: *prompting* and *fine-tuning*. Prompting involves designing input queries that guide the GPT model to generate desired outputs by leveraging its pre-existing knowledge and capabilities [26, 28]. This approach is particularly useful for tasks requiring immediate, context-specific responses without the need for extensive model re-training. In comparison, fine-tuning adjusts the model’s parameters on a targeted dataset, thereby optimizing its performance for specific tasks or domains [26, 28]. The fine-tuning approach allows for a more tailored response generation, closely aligned with the nuances of the given context. Both approaches exhibit significant promise in text comprehension and generation, suggesting their potential effectiveness in producing nuanced, explanatory feedback. Thus, our study aims to harness these approaches to unveil the full capacity of GPT models in automating the generation of high-quality explanatory feedback, thereby addressing a critical need in educational feedback systems. Driven by this, our study proposed two **Research Questions (RQs)**:

**RQ1:** To what extent can we prompt the GPT models to enhance the prediction accuracy of providing explanatory feedback?

**RQ2:** To what extent can the fine-tuned GPT models enhance the prediction accuracy of providing explanatory feedback?

Through this work, we aim to offer a scalable solution that enhances tutor training programs and, ultimately, the learning experience for students. Our study developed an automated explanatory feedback system to highlight the correct and incorrect components of praise from novice tutor attempts, as illustrated in Figure 1. We implemented sequence labeling method for highlighting the correct and incorrect components by using the approaches of prompting and fine-tuning GPT models. To evaluate the quality of highlighted praise components from tutor responses by GPT models, we introduced the Modified Intersection over Union (M-IoU) score, a metric designed for our task. Our results indicate a strong correlation between the M-IoU score and human evaluators’ judgments regarding the quality of highlights, affirming the metric’s reliability.

In addressing **RQ1**, we employed a two-shot prompting method to prompt GPT-3.5 and GPT-4 models to highlight desired and undesired parts of praise in tutor responses. Notably, the GPT-3.5 model demonstrated performance on par with that of the GPT-4 model, exhibiting commendable accuracy in identifying effort-based (M-IoU of 0.46) and outcome-based praise (M-IoU of 0.68). These levels of accuracy are considered decent by human coders, highlighting the effectiveness of our prompting strategies. For **RQ2**, we delved into fine-tuning the GPT models across a set of training sample sizes—from 10% of our dataset (13 samples) to 50% (65 samples)—to gauge how fine-tuning influences the model’s ability to enhance the precision of explanatory feedback. Due to limitations in accessing the fine-tuning GPT-4 model, our investigation focused on fine-tuning the GPT-3.5 model. The optimal fine-tuned GPT-3.5 model achieved M-IoU scores of 0.64 for effort-based praise and 0.84 for outcome-based praise, aligning with the satisfaction levels observed by human coders. Motivated by the effectiveness of our fine-tuned GPT model, we have built a demo<sup>1</sup> of our automated explanatory feedback system.

### Enhancing Automatic Explanatory Feedback Using Generative Pre-trained Transformers (Demo)

You’re tutoring a student named Kevin. He is struggling to understand a math problem. After trying several different approaches, Kevin gets the problem correct. As Kevin’s tutor, you want him to continue working through solving more problems on his math assignment.

**Tutor 1:** You're **doing great**, let's see what the next step is.

🗨️: Saying "**doing great**" is praising students for the outcome. You should focus on praising the students for their effort and process towards learning. Do you want to try responding again?

**Tutor 2:** I appreciate how hard you worked to solve the problem, and I believe you are now ready to do more. Lets give another a try.

🗨️: Saying "**I appreciate how hard you worked to solve the problem**" is a nice example of effort-based praise, which praises students for their effort.

Figure 1: Automated explanatory feedback using sequence labeling method facilitated by fine-tuned GPT-3.5 model

## 2. BACKGROUND

### 2.1 Effective Tutoring Practice

Effective tutoring plays an important role in enhancing student learning by integrating academic knowledge with the capability to address students’ socio-motivational needs [13, 19, 48, 39]. However, equipping tutors with these skills proves challenging, given the limited active learning opportunities that bring situational, scenario-based experiences to the professional development of tutors [7]. Thus, current tutor training for tutors in addressing the social-emotional and motivational aspects of student learning remain under-developed[7, 52].

<sup>1</sup>You can experience the promising features and limitations of our existing system in this demo <https://edm24-effort-outcome.vercel.app/>

Our study focuses on a particular aspect of tutoring practice: the delivery of effective praise. Praise is a fundamental tutoring practice during the human tutoring process, consistently shown to have a positive impact on student motivation, engagement, and learning outcomes [25, 29, 55]. Research highlights that for praise to be effective, it should be: (1) sincere, earned, and truthful; (2) specific by giving details of a student’s strengths; (3) immediate, with praise given right after the student’s action; (4) authentic, avoiding repetitive phrases like “*great job*” which diminishes meaning and becomes predictable, and (5) focused on the learning process rather than innate ability [55]. Existing literature categorizes praise into three types: effort-based *Effort*, outcome-based *Outcome*, and person-based *Person* [29, 55, 8, 7]. Effort-based praise emphasizes the student’s learning process (e.g., “*I love your effort that you put into the writing...*”). Outcome-based praise highlights a student’s achievements, like scoring high on an assignment or solving a problem correctly, and it’s sometimes linked to less effective praise strategies such as “*Good job!*”. Person-based praise attributes success to innate qualities beyond the student’s control such as “*You are smart!*” and is often, similar to outcome-focused praise, considered less effective [29].

Training novice tutors to provide more effective praise (i.e., effort-based praise) requires a comprehensive understanding of both the desirable and undesired parts of their praise responses. For tutors to refine their skills effectively, they should engage in a feedback process to know how well their responses align with the effective praise in tutoring [7, 55]. However, manual feedback generation by expert tutors poses significant challenges due to its time-consuming and labor-intensive nature. This underscores the necessity for exploring automated feedback systems within tutor training programs. Such systems could offer scalable and timely feedback, thereby enhancing tutors’ ability to effectively address student motivation issues.

## 2.2 Feedback for Tutor Training

Feedback in the learning process is universally recognized for its significant impact on learning outcomes [50, 17, 18, 37, 23, 21], with effects ranging from significantly positive [50, 17] to occasionally negative [18], depending on the content and method of delivery. The effectiveness of feedback, as highlighted by Hattie and Timperley [21], is intricately linked to its relevance to the learning context, its timing following initial instruction, and its focus on addressing misconceptions or incorrect reasoning [21]. In particular, immediate, explanatory feedback, which clarifies why certain responses are correct or incorrect, plays a crucial role in promoting active engagement and thoughtful practice among learners [53, 37, 21, 23, 17, 50]. The growing importance of feedback has motivated the adoption of automated feedback systems in educational settings, such as OnTask, which allows educators to provide scalable feedback based on conditional rules related to students’ academic activities and performance [49]. Yet, the application of such systems in tutor training remains under-explored.

An important method of deploying automated feedback in tutor training involves the use of templated feedback. The templated feedback, including specific references to desired and undesired elements of the tutor responses, is informed

by earlier results on the effectiveness of having a rich, data-driven error diagnosis taxonomy driving template-based feedback [1]. Our study aims to employ natural language processing techniques to automate the identification of desirable and undesirable elements within tutor responses, facilitating the generation of templated explanatory feedback.

## 2.3 Sequence Labeling for Feedback Generation

Sequence labeling, a fundamental task in natural language processing (NLP), plays a pivotal role in identifying and categorizing key segments of text according to predefined labels [26]. To elucidate the mechanism of sequence labeling, we consider Named Entity Recognition (NER) as a representative subtask, which is closed to the task in our study. NER seeks to automatically detect and classify named entities—words or phrases with specific attributes—into categories such as person, organization, and location [26, 36]. For instance, in the sentence “*John said that Pittsburgh is wonderful in the winter,*” the terms “*John*”, “*Pittsburgh*”, and “*winter*” would be labeled as **Person**, **Location**, and **Time**, respectively, showcasing how NER distinguishes and categorizes entities within a textual context.

Our study extends the application of sequence labeling to identify and highlight components related to different types of praise within tutor responses. This process involves discerning specific words or phrases that signify the kind of praise being used, thereby offering tutors insight into their feedback practices. For example, “*You are doing great.*”, the phrase “*doing great*” in this context is identified as an outcome-based praise. Leveraging sequence labeling allows our AI model to spotlight such instances of praise, enabling the provision of nuanced, explanatory feedback. An example of such feedback might be, “*Saying “doing great” is praising the student for the outcome. You should focus on praising the students for their effort and process towards learning. Do you want to try responding again?*” This approach facilitates the generation of targeted, template-based feedback for tutors.

Notably, while previous research has explored sequence labeling techniques for similar purposes [40], the accuracy of their proposed models in precisely identifying and categorizing feedback elements remains a challenge. This limitation underscores the need for leveraging more advanced models to provide accurate, informative feedback to tutors.

## 2.4 Large Language Models in Education

Recent advancements in natural language processing have seen the evaluation of large language models (LLMs) like GPT models in various educational tasks, leveraging techniques such as prompting and fine-tuning [30]. The GPT models (e.g., GPT-3.5 or GPT-4) have demonstrated significant potential in enhancing many educational tasks (e.g., feedback generation and learning content generation) [30]. Motivated by these developments, our study aims to investigate the applicability of prompting and fine-tuning GPT models to identify both desirable and undesirable aspects of tutoring responses. We intend to evaluate the effectiveness of these approaches in developing an automated system for providing explanatory feedback.

### 2.4.1 Prompting large language models

Prompting, which involves the use of specific queries or statements to guide a large language model’s (LLM) output, has been identified as a significant technique for leveraging the capabilities of LLMs in education [30]. The prompting strategy plays a pivotal role in effectively guiding the models, such as GPT-3 and GPT-4, to produce responses that are more aligned with the context and requirements of the tasks. Research by Dai *et al.* [10] on the GPT-3.5 and GPT-4 model highlighted GPT models’ ability to generate student feedback that surpassed human instructors in readability. Furthermore, Hirunyasiri *et al.* [24] demonstrated the superiority of the GPT-4 model over human expert tutors in assessing specific tutoring practices. [34] used GPT-4 model to generate high-quality answer responses for middle school math questions. [45] providing feedback for multiple-choice questions at the middle-school math level. Given that GPT models has shown remarkable performance on various educational tasks [9, 24, 34, 45], our study also leveraged the GPT models to further explore its capabilities in automatically generating explanatory feedback since the exploration of prompting GPT models for providing explanatory feedback in response to open-ended questions remains limited.

### 2.4.2 Fine-tuning large language models

In addition to prompting the GPT models, the fine-tuning of GPT models has also shown considerable promise in various educational tasks [30]. The fine-tuning method adjusts the model’s neural network to better suit particular domain, thereby enhancing its performance in relevant contexts [26]. Latif and Zhai [33] employed fine-tuned GPT-3.5 model for the purpose of automatic scoring in science education. Their findings indicate that GPT-3.5, once fine-tuned with domain-specific data, not only surpassed the performance of the established BERT model [12] but also demonstrated superior accuracy in assessing a variety of science education tasks. Such advancements underscore the value of fine-tuning GPT models for educational applications, showcasing their ability to provide precise, scalable solutions across diverse educational settings. Bhat *et al.* [3] introduced a method for generating assessment questions from text-based learning materials using a fine-tuned GPT-3 model. The generated questions was further assessed regard to their usefulness to the learning outcome by human experts, with the findings revealing a favorable reception among human experts. Inspired by these pioneering research, our study aims to extend the application of fine-tuning method to GPT models within the context of generating explanatory feedback. While the aforementioned studies [12, 33] have not directly addressed the generation of explanatory feedback, their success in applying fine-tuned LLMs within educational domains suggests a promising avenue for our investigation. By customizing GPT models to the nuances of educational feedback, we anticipate uncovering new potentials in automating and enhancing the feedback process. These efforts will contribute to the growing body of evidence supporting the integration of fine-tuned LLMs in educational technology, potentially revolutionizing the way feedback is generated and applied in learning environments.

## 3. METHOD

### 3.1 Dataset

Our study received IRB ethical approval with the protocol number: STUDY2018\_00000287 from Carnegie Mellon University. The study utilized a dataset comprising responses from 65 volunteer tutors who participated in the *Giving Effective Praise* lesson. The demographic breakdown of these tutors was as follows: 52% were White, 18% Asian, and 52% male, with over half being 50 years or older. The objective of *Giving Effective Praise* is to equip tutors with skills to boost student motivation through the delivery of effective praise. We collected 129 responses from the tutors who completed the lesson, and these responses are sorted according to the type of praise (i.e., effort-based praise, outcome-based praise, and person-based praise). Notably, the dataset contained only one instance of person-based praise (“*You are very smart*”), leading to its exclusion from the analysis. Thus, our study mainly focused on the analysis of effort-based and outcome-based praise.

### 3.2 Sequence Labeling

We aim to provide explanatory feedback that can highlight components of effort-based and outcome-based praise within the tutor responses. Thus, we decided to use a sequence labeling method. By doing so, we created the annotation guideline as well as specific examples of *Effort* and *Outcome* based on the studies [55, 8, 7]. To carry out the annotation work, we hired two expert educators who first completed the lesson of *Giving Effective Praise* from our platform and then started annotating the praise tags representing attributes associated with *Effort* and *Outcome*, for 129 tutor responses.

In the pursuit of advancing our understanding of effective praise within tutoring dialogues, our study leverages the **Inside-Outside (IO)** labeling scheme [31] in our study. The **IO** scheme can capture the necessary information for our analysis, allowing us to maintain focus on the core aspects of praise within tutor responses without the need for differentiating between the beginning or end of entities, which aligns with our needs. The **IO** scheme, characterized by its simplicity and efficiency, labels tokens as an inside tag (**I**) and an outside tag (**O**). The **I** tag is for the praise components, whereas the **O** tag is for non-praise words. For example, when annotating praise components for a tutor’s praise “*You are making a great effort*”, the words “*great*” and “*effort*” are identified as part of the *Effort* (i.e., **I<sub>Effort</sub>**) and the remaining text in the response is identified as the outside (i.e., ‘**O**’) of the praise components. By annotating the praise components for each tutor response, we can obtain a list of tokens as shown in Figure 2.

[You, are, making, a, great, effort]  
[ O, O, O, O, I<sub>Effort</sub>, I<sub>Effort</sub>]

Figure 2: Labeling the praise components using IO scheme

In assessing inter-rater reliability for our study, we note that while Cohen’s Kappa is considered the standard measure of inter-annotator agreement for most annotation tasks [44], its suitability for sequence labeling tasks—Named Entity Recognition or similar tasks where labels are assigned to specific words or tokens within a sequence—is limited [4, 16]. Specifically, sequence labeling may result in partial agreements between annotators (e.g., consensus on token label

type but not on exact boundaries), which may not be effectively captured by Cohen’s Kappa as it does not account for partial agreement [4]. Additionally, in sequence labeling, a large proportion of tokens are typically labeled as ‘O’ (the distribution of token labels in our study is shown in Table 2), leading to an imbalanced label distribution. Since Cohen’s Kappa assumes an equal likelihood of each category being chosen, it may not provide a meaningful measure of agreement in situations where the vast majority of labels belong to a single category, making the metric less informative or even misleading [4]. Given these limitations, F1 score is often preferred for evaluating inter-rater reliability in sequence labeling tasks as suggested in previous studies [4, 11]. As the token level Cohen’s Kappa scores can also provide some insight, we provide both Cohen’s Kappa and F1 scores to provide a comprehensive view of annotator agreement in our study. Our results—0.49 for Cohen’s Kappa and 0.79 for the F1 score—were deemed acceptable for the purposes of our task as suggested by [4, 15]. To address discrepancies between two annotators, a third expert was invited to resolve inconsistencies. The distribution of annotated praise in our dataset is as follows: 59 responses with only effort-based praise, 22 with only outcome-based praise, 31 containing both types, and 17 lacking mentions of either, illustrating the varied nature of praise within the responses.

### 3.3 GPT Facilitated Sequence Labeling

As discussed, our study employed two widely used approaches for adapting GPT models to sequence labeling tasks: prompt engineering and fine-tuning. Each method offers unique advantages and impacts the process of creating automated explanatory feedback in different ways.

#### 3.3.1 Prompt engineering for identifying praise components

To answer **RQ 1**, we conducted prompt engineering to design certain prompting strategies to enable GPT models to identify the praise components within the tutor responses. Prompting engineering approach involves designing and structuring input prompts to guide the GPT model in generating desired outputs [42, 56]. The art of prompt engineering lies in crafting prompts that can effectively communicate the context and requirements of the task to the model [42, 56]. In our study, given the presence of tutor responses that exemplify both effort-based and outcome-based praise in the tutor training lesson, we employed a two-shot prompting strategy to guide GPT-3.5 (`gpt-3.5-turbo-0125`) and GPT-4 (`gpt-4-0125-preview`) models to highlight praise components within tutor responses. Our prompt is shown in Table 1. The following explains our prompt design, aimed at extracting specific elements from tutor responses related to praising student effort and outcomes.

- **{Lesson Principle}**: This segment provides the guiding principles for desired tutor responses. It includes key aspects of effective praise in educational settings, such as sincerity, specificity, immediacy, authenticity, and focus on the learning process. This principle acts as a reference for evaluating the tutor responses. The lesson principle is detailed in Appendix A
- **{Tutor Response}**: This part simulates an interactive

**Table 1: Prompt for identifying praise from tutor responses**

Role	Content
<b>System</b>	<i>You are a response evaluator designed to output JSON. Your task is to analyze tutor responses based on the principles of effective praise focusing on ‘effort’ and ‘outcome’. Extract words or phrases that represent praise for the student’s effort and outcome, and output the results in JSON format with keys titled ‘Effort’ and ‘Outcome’.</i>
<b>User</b>	<b>Lesson Principle</b>
<b>Assistant</b>	<i>Sure, can you provide a tutor response for analysis</i>
<b>User</b>	<i>An example of outcome-based praise is: “Great job! You are a genius!”</i>
<b>Assistant</b>	<i>An output json format is: {“effort”: [], “outcome”: [“Great job”]}</i>
<b>User</b>	<i>Nice, let’s do it again.</i>
<b>Assistant</b>	<i>Sure, can you provide a tutor response for analysis?</i>
<b>User</b>	<i>An example of effort-based praise is: “You are almost there! I am proud of how you are persevering through and striving to solve the problem. Keep going!”</i>
<b>Assistant</b>	<i>An output json format is: {“effort”: [“persevering through and striving to solve the problem”, “Keep going”], “outcome”: []}</i>
<b>User</b>	<i>Nice, let’s do it again.</i>
<b>Assistant</b>	<i>Sure, can you provide a tutor response for analysis</i>
<b>User</b>	<b>Tutor Response</b>

environment where the model identify the praise components from the input of tutor responses.

#### 3.3.2 Fine-tuned GPT models for identifying praise components

Given limited access to fine-tuning capabilities for the GPT-4 model, we focused on optimizing the use of GPT-3.5 (`gpt-3.5-turbo-1106`) to answer the **RQ2**, particularly within the constraints of a modestly sized training dataset. The model fine-tuning approach was implemented to train the GPT-3.5 model to recognize and understand the patterns associated with identifying praise components in tutor responses. To prepare our data for the fine-tuning process, we converted tutor responses and their associated tags into JSON format. This format facilitated the structured representation of our data, mirroring the input style typically expected by the GPT model. The structure of our input data closely resembled the prompts used in GPT model training, with a key distinction: instead of prompting the model to generate text containing praise, we supplied it with annotated outcomes and effort-based praises. Due to the page limit and avoid repetitive content appearing in the paper, we decided to put the details of the fine-tuning input in the Appendix B.

Our approach aimed to investigate the extent to which fine-tuned model can accurately classify and label praise components under limited training dataset, thereby enhancing its performance on our task. By doing so, we first divided our dataset evenly, allocating 50% (65 responses) for training and the remaining 50% (64 responses) for testing. The dis-

tribution of annotation is shown in Table 2 which presents **O** as the major tag in our dataset. Then, we subdivided our training set into five distinct partitions: 13, 26, 39, 52, and 65 responses. For each partition, the training process was repeated five times using different random seeds. These partitions represented 10%, 20%, 30%, 40%, and 50% of our original dataset, respectively. This stratified approach allowed us to simulate different training conditions, thereby enabling a comprehensive analysis of the model’s adaptability and learning efficiency as the amount of available training data varied.

**Table 2: Distribution of token labels.**

	% Annotation (I/O)		
	<b>O</b>	<b>I<sub>Effort</sub></b>	<b>I<sub>Outcome</sub></b>
<b>Full</b>	2415 (77.72%)	562 (18.09%)	130 (4.18%)
<b>Training</b>	1241 (78.99%)	282 (17.95%)	48 (3.06%)
<b>Testing</b>	1174 (76.43%)	280 (18.23%)	82 (5.34%)

### 3.4 Metrics

#### 3.4.1 Modified Inter Section over Union Scores

In sequence labeling tasks, traditional metrics like the F1 score, as depicted in Equation 1, are commonly used to assess model performance [4]. In the context of our study, True Positives (TP) represent the number of tokens correctly identified as praise by the model, False Positives (FP) refer to tokens incorrectly classified as praise, often resulting from the model predicting additional words as part of the praise. False Negatives (FN), on the other hand, are tokens that were part of praise but were overlooked by the model, indicative of missed praise components. Previous research [40] has highlighted that certain additional entities identified by the model can still contribute meaningfully to human tutors’ understanding of response correctness. For instance, as illustrated in Table 3, while the first row shows expert annotations of effort-based praise, subsequent examples (rows 2-5) might be model-generated. Notably, rows 2 to 4, despite including additional words for effort-based praise (i.e., FP), offer valuable insights that could assist tutors, contrasting with row 5 where the model’s highlighting of merely “*great*” (i.e., FN) fails to clearly convey the type of praise intended. This observation suggests the a need for a metric that accommodates the evaluation of additional identified praise tokens more flexibly. However, the F1 score, as shown in Equation 1, applies the same weight to both FP and FN, a treatment that diverges from our requirement for a more nuanced metric. Consequently, we propose adopting the Intersection over Union (IoU) concept, commonly utilized in the computer vision domain, to better suit our evaluation needs in our task.

$$F1\ score = \frac{TP}{TP + \frac{1}{2}(FP + FN)} \quad (1)$$

The Intersection over Union (IoU) metric, frequently applied in object detection and segmentation tasks as depicted in

**Table 3: Original training instance and different types of augmented instances. We highlighted outcome-based praise using yellow color and effort-based praise using blue.**

	<b>Instance</b>	<b>Label</b>
1	<i>John, you are making a really <b>great effort</b>.</i>	True
2	<i>John, you are making a <b>really great effort</b>.</i>	Pred
3	<i>John, you are making <b>a really great effort</b>.</i>	Pred
4	<i>John, you are <b>making a really great effort</b>.</i>	Pred
5	<i>John, you are making a really <b>great effort</b>.</i>	Pred

Equation 2, quantifies the extent of overlap between predicted and actual annotations, offering a balanced approach to assess model performance [41, 6]. In the context of sequence labeling, the ‘*Area of Overlap*’ (i.e., TP) corresponds to the tokens the model accurately identifies as praise, whereas the ‘*Area of Union*’ includes all tokens labeled as praise by the model (TP and FP) along with all actual praise tokens in the ground truth (TP and FN). Since we recognize the significance of additionally detected words in our study, we propose a Modified Intersection over Union (M-IOU) metric (shown in Equation 3) to refine IoU metric further. This modification incorporates a weight coefficient,  $\alpha$ , which aims to reduce the influence of FPs on the overall performance score, thus introducing a measure of flexibility towards additional identified words without neglecting the potential for inaccuracies. The coefficient  $\alpha$  is introduced as a real number set at or above zero, enabling users to adjust the tolerance level for additional words identified. A higher  $\alpha$  value enforces a stricter penalty on FPs, while a lower value indicates a more lenient approach. In our analysis,  $\alpha$  is set to 0.2 based on our observation of expert annotations. Notably, in situations where a response lacks praise and the model’s prediction concurs (i.e.,  $TP + FP + FN$  equals 0), indicating a perfect match between model and ground truth in identifying no relevant praise tokens, we encounter a scenario reflective of novice tutors possibly providing irrelevant responses (e.g., “*Can I see any of your writing*”). Such irrelevant responses underscore the necessity of our explanatory feedback system in guiding tutors on giving effective praise. For this case, we adjust the M-IOU formula to directly assign a score of 1 to reflect perfect agreement and underscore the adaptability of our M-IOU in accurately evaluating model precision, particularly in the absence of praise, thus showing its effectiveness in practical applications.

$$IoU = \frac{\text{Area of Overlap}}{\text{Area of Union}} = \frac{TP}{TP + FP + FN} \quad (2)$$

$$M\text{-IOU} = \frac{TP}{TP + \alpha \times FP + FN} \quad (3)$$

#### 3.4.2 Human annotation and correlation analysis

To assess the efficacy of our proposed M-IOU score, we undertook a rigorous process involving human annotation to rate the quality of identified components of praise within tutor responses. The human rating scores are further used to compare with our proposed M-IoU score to ensure that

M-IoU not only holds computational validity but also aligns with human judgments regarding the praise components in the tutoring responses. Recognizing the importance of human judgment in our study, we hired two additional human coders to scrutinize the highlighted components of praise in tutor responses. These coders attended a detailed annotation training session and completed the lesson of *Giving Effective Praise*, equipping them with the necessary background to perform their evaluations effectively.

Before beginning their rating tasks, we randomized the presentation order of highlighted texts generated by both GPT models and expert annotations for each tutor response. This approach ensured the unpredictability of expert annotation sequence, aiming to mitigate any potential bias in the coders’ evaluations. Inspired by the study [54], we guided the coders to assess each highlighted response based on two questions: **Question 1:** “Do you think the highlighted text provides enough information to identify praise on effort?” and **Question 2:** “Do you think the highlighted text provides enough information to identify praise on the outcome?”. These questions were designed to capture the coders’ assessments of the highlighted texts’ adequacy in conveying praise, either for the student’s effort or the outcome of their work. The coders were instructed to use a five-point Likert scale for their annotations, with the options being: 1 - *Strongly Disagree*, 2 - *Disagree*, 3 - *Neutral*, 4 - *Agree*, 5 - *Strongly Agree*.

Upon completing the annotations, we calculated the average score for each response, providing a quantitative measure of the consensus between the coders regarding the effectiveness of the highlighted praise text. To determine the effectiveness of the M-IoU score as a metric for evaluating model predictions, we conducted a correlation analysis using Pearson’s  $r$  to understand the strength and direction of the linear relationship between the M-IoU scores and the human coders’ ratings. The correlation analysis help us understand how well the M-IoU score aligns with human judgment and its potential as a surrogate metric for evaluating the model performance in identifying praise components.

## 4. RESULTS

### 4.1 Results on RQ1

To answer **RQ1**, we aimed to evaluate the extent to which the models’ highlighted elements by prompting approach could adequately convey information necessary for identifying the type of praise being expressed. By doing so, we first conducted correlation analysis (as described in the Section 3.4.2) to validate the efficacy of Modified Intersection over Union (M-IoU) score. Due to the page limit,<sup>2</sup> we only present the results of correlation analysis for effort-based praise between the M-IoU scores and the coders’ ratings (shown in Figure 3). The findings revealed a significant positive correlation ( $p < 0.01$ ) between the M-IoU scores and the ratings from both individual coders (Coder 1 and Coder 2) as well as the averaged scores of the coders (Coder AVG.) for the identification of effort-based praise. This significant correlation underscores the reliability and effectiveness of our proposed M-IoU metric in evaluating the quality of highlighted components from GPT models.

<sup>2</sup>The correlation analysis for outcome-based praise can be found in the Appendix C.

Then, we examined the quality of highlighted elements by prompting GPT-3.5 and GPT-4. In Table 4, we presented the descriptive statistics of the scores rated by two human coders and measured by M-IoU scores. Since the M-IoU score ranging from 0 to 1, to facilitate a direct comparison between human scores and M-IoU scores, we normalized the human coders’ rating scores (originally on a scale from 1 to 5) to the same 0 to 1 range. It is important to note that the calculation of M-IoU scores was based on the overlap of highlighted text between the GPT models and expert annotations; consequently, assigning an M-IoU score for expert annotation was not applicable and is thus indicated as N/A in Table 4. The results in Table 4 revealed that the highlighted text for outcome-based praise consistently received higher human rating scores and M-IoU scores than that for effort-based praise across both GPT models. This finding aligns with our intuition, considering that outcome-based praise, characterized by expressions such as “*Good job*” and “*Well done,*” tends to be more structured and straightforward to identify than the more nuanced effort-based praise. Interestingly, the difference in M-IoU scores between GPT-3.5 and GPT-4 for both types of praise was marginal, despite the reputed superiority of the GPT-4 model in numerous educational tasks.

**Table 4: Descriptive statistics of the scores rated by two human coders and measured by M-IoU scores**

	GPT-3.5 turbo		GPT 4 turbo		Expert Annotation	
	Effort	Outcome	Effort	Outcome	Effort	Outcome
<b>Comparison between the normalized human ratings and M-IoU scores</b>						
<b>Coder 1</b>	0.68 <sub>0.38</sub>	0.79 <sub>0.40</sub>	0.63 <sub>0.36</sub>	0.75 <sub>0.39</sub>	0.77 <sub>0.35</sub>	0.89 <sub>0.30</sub>
<b>Coder 2</b>	0.60 <sub>0.43</sub>	0.76 <sub>0.40</sub>	0.57 <sub>0.40</sub>	0.74 <sub>0.40</sub>	0.77 <sub>0.35</sub>	0.84 <sub>0.35</sub>
<b>Avg.</b>	0.64 <sub>0.35</sub>	0.77 <sub>0.37</sub>	0.60 <sub>0.33</sub>	0.75 <sub>0.38</sub>	0.77 <sub>0.29</sub>	0.87 <sub>0.29</sub>
<b>M-IoU</b>	0.46 <sub>0.36</sub>	0.68 <sub>0.44</sub>	0.47 <sub>0.38</sub>	0.64 <sub>0.46</sub>	N/A	N/A
<b>Proportion of human rating ‘Agree’ or higher on our scale*</b>						
<b>Coder 1</b>	64.06%	76.56%	53.13%	75.00%	73.44%	89.06%
<b>Coder2</b>	53.13%	73.44%	46.88%	71.88%	75.00%	84.38%
<b>Avg.</b>	56.25%	73.44%	53.13%	73.44%	75.00%	85.94%

\*Note: Proportion of rating greater than or equal to ‘Agree’ (i.e., agree with the highlight text provides enough information to identify the praise on effort or outcome)

We further investigated the proportion of highlights that achieved a rating of 4 or above (corresponding to ‘Agree’ or higher on our scale), termed here as ‘satisfied highlighted text’. The proportion serves as an indicator of the highlights’ utility in facilitating the identification of the accurate type of praise. In Table 4, our analysis disclosed that over 50% of the effort-based praise highlights generated by prompting the GPT-3.5 model were deemed effective by the coders in identifying effort-based praise, whereas for outcome-based praise, the proportion exceeded 70%. Interestingly, the satisfaction proportion for GPT-3.5’s highlights surpassed those of GPT-4, suggesting a nuanced difference in their performance. Moreover, expert annotations were observed to yield the highest satisfaction rates, with over 70% for effort-based praise and 80% for outcome-based praise highlights considered satisfactory by the coders. The coders’ ratings for expert-annotated text were generally higher, reflecting the expert annotations’ authenticity and precision



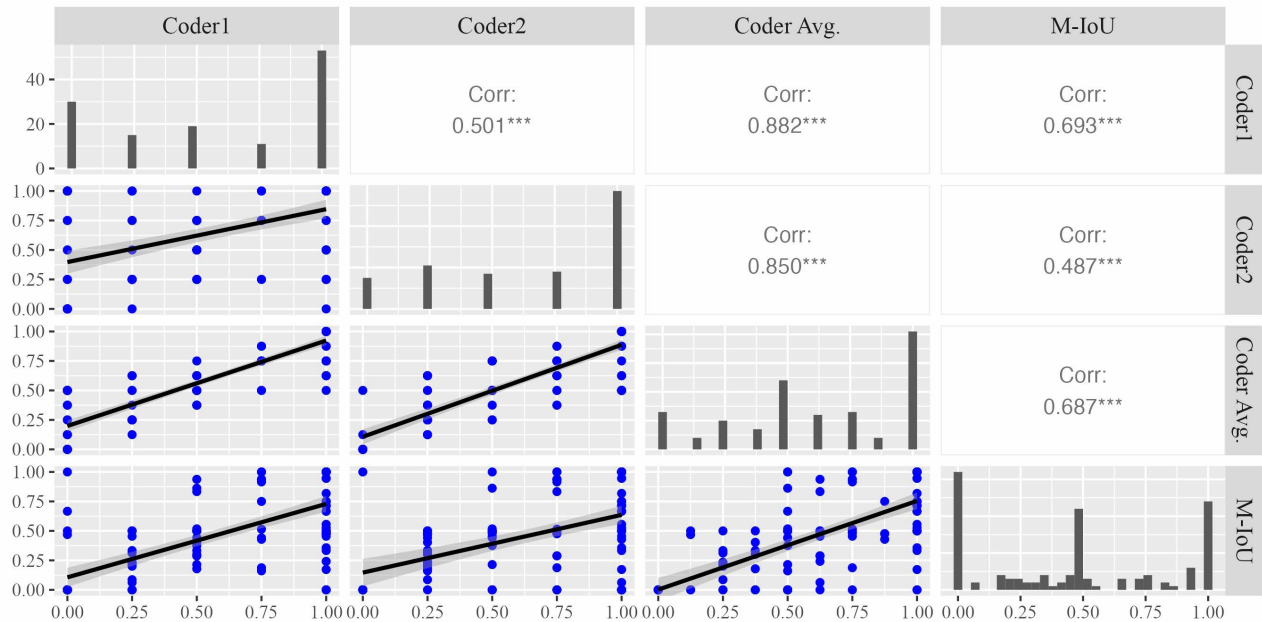


Figure 3: Scatter matrix of the correlation on the effort-based praise

in capturing the essential elements of praise, which indicates the potential limitations of relying solely on Cohen’s Kappa for evaluating agreement in sequence labeling tasks. The coders’ perceptions affirm the significance of the highlighted texts’ quality over mere statistical agreement, indicating that expert annotations, despite a lower Cohen’s Kappa, effectively convey the essential attributes of praise within the tutor responses.

## 4.2 Results on RQ2

Building upon the insights gained from the results on RQ1, where the M-IoU was established as a viable proxy for assessing the quality of text highlighted by GPT models, we delved into the potential of fine-tuning the GPT-3.5 model to enhance its performance in identifying praise within tutor responses. Notably, our ability to fine-tune the GPT-4 model was constrained due to access limitations. Consequently, our efforts were concentrated on the GPT-3.5 model, the performance of which is depicted in Figure 4 and the detailed results of model performance was shown in the Appendix D

In Figure 4, we present the model’s performance, quantified by averaging the M-IoU scores derived from five distinct random seeds across various training partitions (from 13 to 65 training sample size). The inclusion of error bars in our analysis offers a visual representation of the model’s performance variability, ranging from its maximum to its minimum across these partitions. We simulated a low-resource scenario—characterized by a limited training dataset—and observe the fine-tuned GPT-3.5 model’s capability to maintain satisfactory performance under such constraints. Starting with a mere 13 training samples (10% of the full dataset), the model demonstrated an approximate M-IoU score of 0.5 for effort-based praise and 0.65 for outcome-based praise, showcasing performance on par with that achieved through the prompting method applied to the GPT models. As the




Figure 4: Performance of fine-tuning GPT-3.5 model on highlighting correct types of praise in different training size

training sample size increased, we generally observed an improvement in model performance, with a peculiar exception observed in the outcome-based praise performance when utilizing 52 training samples. Expanding the dataset to 65 samples resulted in the model attaining an M-IoU score of roughly 0.6 for effort-based praise—surpassing the efficacy of the prompting method. Correspondingly, the performance on outcome-based praise reached an M-IoU score of 0.75, rivaling that of expert annotations given that an M-IoU of 0.68 equates to a human rating score of 0.77.

Motivated by these promising outcomes, we elected to adopt the model exhibiting optimal performance in highlighting effort-based praise as the foundation for our automated feedback system within tutor training programs. This decision



is underpinned by the pivotal role of effort-based praise in educational feedback; it is the essence of effective praise, determining the appropriateness and impact of the tutor’s feedback on student motivation and learning. The ability to accurately identify and underscore effort-based praise in tutors’ responses is thus deemed crucial for enhancing the quality of educational feedback. In support of this initiative, a demo of our automated explanatory feedback system is accessible via the provided link<sup>3</sup>, showcasing the application’s potential to transform tutor training by emphasizing the significance of effort-based praise.

## 5. DISCUSSION

Our study examined the potential of GPT models to highlight the desired and undesired parts of praise from trainee responses and further integrated the highlighted parts into the feedback for tutor training. By employing the Modified Intersection over Union (M-IoU) as a novel metric, we measured the quality of the praise highlighted by GPT models. The M-IoU metric, validated through correlation with human coders, underscores the potential of combining human intuition with algorithmic metrics to enhance the specificity and relevance of educational feedback. The findings from our investigation confirmed the considerable promise of employing techniques such as *prompting* and *fine-tuning* within GPT models to generate automated, explanatory feedback tailored for tutor training programs. By leveraging a fine-tuned GPT model, we developed an automated feedback system specifically designed for tutor training, with the objective of delivering immediate and explanatory feedback. This innovation presented a viable, scalable solution to the pressing challenge of delivering personalized feedback to learners (trainee tutors as learners in our study). The implementation of an automated explanatory feedback system in our study exemplifies how such technology can be leveraged to identify specific elements in tutors’ open-ended responses that are either desirable or in need of enhancement.

**Prompting GPT model to highlight key components.** Upon evaluating the highlighted praise components from GPT models (prompting) and expert annotations, we observe that the quality of highlighted praise components by experts typically outperform those highlighted by GPT models. For instance, as indicated in the first row of Table 5, there is unanimous agreement among human coders that the highlighted praise components by expert annotations are better than those highlighted by both GPT models. Specifically, while the GPT-3.5 model accurately identified phrases such as *“doing a great job”* and *“Stick with this”* as forms of praise, it erroneously categorized *“doing a great job”* as effort-based praise, contrary to the established praise principle which classifies it as outcome-based praise [55]. Conversely, the GPT-4 model correctly classified *“doing a great job”* as outcome-based praise but included additional words like *“We can finish it”* in its identification of effort-based praise. The comparison of additional words included in effort-based praise annotations between GPT-3.5 and GPT-4 models resulted in identical scores from the coders (0.5 in Table 5, reflecting a neutral stance equivalent to a score of 3 on the Likert Scale). The identical scoring stems from the

<sup>3</sup>The demo of our automated explanatory feedback can be found here <https://edm24-effort-outcome.vercel.app/>

equal number of additional words identified by both models. Furthermore, the M-IoU score aligns with the coders’ assessments, underscoring the metric’s utility in capturing the accuracy of the models’ annotations. In another observation, detailed in the second row of Table 5, both coders concurred that in certain instances, prompting GPT-3.5’s identification of praise components was superior to that of expert annotations. Additionally, the third row of Table 5 presents a scenario with significant discrepancies in the ratings assigned to the highlighted praise components by GPT-3.5 and GPT-4 between two coders. Here, the M-IoU score proved instrumental in mitigating the variances in individual assessments, effectively approximating the average score derived from both coders’ ratings.

**Fine-tuning GPT model to highlight key components.** Then, our study assessed the impact of fine-tuning the GPT-3.5 model with different amount of training data to determine the optimal dataset size required to achieve satisfactory performance in generating explanatory feedback. This insight is important for researchers and educational practitioners seeking to use LLMs effectively, especially when faced with constraints on data availability. Our findings highlight the critical role of task-specific optimization for LLMs, illustrating how strategic modifications to the quantity of training data can markedly enhance the performance of automated feedback systems. By identifying the minimum dataset requirements for fine-tuning GPT models, our study provides valuable guidelines for developing effective explanatory feedback. Furthermore, by integrating our proposed prompting strategies alongside a certain number of training datasets, we found that the fine-tuned GPT-3.5 model generally outperforming prompting models (both GPT-3.5 and GPT-4) in identifying the praise elements (including effort- and outcome-based praise). It suggest that, despite the general advancements represented by newer models like GPT-4, fine-tuning earlier versions such as GPT-3.5 can achieve comparable or even superior performance in specific applications. This insight is important for educational practitioners and researchers, particularly those constrained by financial limitations, as fine-tuning GPT-3.5 proves to be a more cost-effective option than prompting GPT-4. Moreover, the fine-tuning approach offers a solution to challenges related to accessing the latest models or dealing with limited resources, such as a restricted number of training datasets.

**Comparison of prompting and fine-tuning approaches.** In our study, we employed both prompting and fine-tuning approaches to adapt large language models, specifically GPT models, for providing highlighting the desired and undesired parts of trainee responses. Prompting enables rapid model adaptation to highlight the components of effort- and outcome-based praise without extensive retraining, thus conserving computational resources and time. However, since the model parameters are not updated, prompting might not capture deeper insights from annotated data, potentially limiting performance on highlighting key components from complex responses. For example, consider the tutor response *“Great job figuring out that problem! Would you like help with anything else?”* Because of the use of *“figuring out”*, this response is categorized as effort-based praise, however, GPT-4 without fine-tuning mistakenly classified it as outcome-based, whereas GPT-3.5 with fine-tuning cor-

Table 5: Examples of evaluation on highlighted praise components from two human coders’ rating scores and M-IoU scores.

Row	Categories	Responsese	Coder 1		Coder 2		M-IoU	
			Effort	Outcome	Effort	Outcome	Effort	Outcome
1	GPT-3.5	Carla you are <i>doing a great job!</i> <i>Stick with this.</i> <i>We can finish it.</i>	0.50	0	0.50	0	0.50	0
	GPT-4	Carla you are <i>doing a great job!</i> <i>Stick with this.</i> <i>We can finish it.</i>	0.50	0.75	0.50	0.75	0.50	0.83
	Expert	Carla you are <i>doing a great job!</i> <i>Stick with this.</i> <i>We can finish it.</i>	1.00	1.00	1.00	1.00	N/A	N/A
2	GPT-3.5	<i>Great job, Kevin!</i> <i>I can tell how hard you worked to get there.</i>	1.00	1.00	1.00	1.00	0.53	1.00
	GPT-4	<i>Great job, Kevin!</i> <i>I can tell how hard you worked to get there.</i>	0.75	1.00	0.75	1.00	1.00	1.00
	Expert	<i>Great job, Kevin!</i> <i>I can tell how hard you worked to get there.</i>	0.75	1.00	0.75	1.00	N/A	N/A
3	GPT-3.5	<i>Great job Kevin!</i> <i>Your determination is really admirable!</i> <i>Pretty sure you can complete it with this determination!</i>	0.25	1.00	1.00	1.00	0.48	1.00
	GPT-4	<i>Great job Kevin!</i> <i>Your determination is really admirable!</i> <i>Pretty sure you can complete it with this determination!</i>	0.25	1.00	1.00	1.00	0.48	1.00
	Expert	<i>Great job Kevin!</i> <i>Your determination is really admirable!</i> <i>Pretty sure you can complete it with this determination!</i>	1.00	1.00	0.25	1.00	N/A	N/A

rectly classified it as effort-based. This error likely occurred because the model over-weighted the generic phrase “Great job”. Additionally, while the prompting approach offers flexibility in testing different prompts to quickly gauge the model’s capabilities on our task, its effectiveness heavily depends on the quality of the prompt design. As observed during our prompt engineering phase, inadequate prompts can lead to misleading outputs.

On the other hand, fine-tuning allows for deeper model customization by adjusting internal parameters to closely align with our task in identifying the components of praises from tutor responses, often resulting in superior performance measured by M-IoU scores, as observed in our study. Fine-tuning enables the GPT model to deeply integrate new knowledge and adjust its existing knowledge, better fitting the task requirements of identifying components of effort- and outcome-based praise. Despite these advantages, fine-tuning requires a substantial amount of relevant and high-quality data and significant computational resources. The data must be carefully annotated to guide the model effectively toward the desired behavior, which present a significant limitation if such data is scarce or difficult to collect. Additionally, fine-tuning involves updating the weights of a neural network based on a specific dataset, a process that can be resource-intensive and requires access to powerful hardware, especially for larger models.

To address some of these challenges and further enhance our highlighted feedback system, we are considering the integration of Retrieval-Augmented Generation (RAG). RAG combines the strengths of both retrieval and generation models to improve the performance of language models on specific tasks [35]. RAG could enhance the performance of prompting LLMs by dynamically incorporating relevant external information into responses, providing more informed and contextually accurate outputs (e.g., [20]). Additionally, RAG can be integrated with the fine-tuning approach for providing highlighted feedback, potentially improving the model’s accuracy in highlighting components of praise. This integration aims to create a model that not only leverages external data through RAG but also adapts more finely to specialized tasks through fine-tuning, demonstrating superior performance in contextually rich and dynamic environments.

## 6. LIMITATION AND FUTURE WORKS

**Measuring the impact of the proposed feedback system.** While the current study demonstrates the potential of using GPT-based models for providing explanatory feedback in a novice tutor training context, we acknowledge the necessity of validating the effectiveness of feedback with highlighted components through empirical research involving actual users. To this end, we propose a comprehensive study aimed at assessing the real-world effectiveness and impact of our feedback system on novice tutors. The planned study will involve a group of novice tutors who will use our automated feedback system during their training sessions. The study will be designed to capture both qualitative and quantitative data to provide a holistic evaluation of the feedback system’s performance. Quantitative data will be collected through pre- and post-tests to measure the learning gains of tutors, while qualitative data will be gathered from surveys and interviews to assess tutors’ perceptions and experiences with the feedback.

**Expanding the scope of the proposed feedback systems for diverse tutoring scenarios.** We aim to empower novice tutors through automated explanatory feedback, enabling them to grasp effective tutoring strategies within our training programs. While the fine-tuned GPT-3.5 model has shown promising results in delivering explanatory feedback for giving effective praise, its applicability and effectiveness across a broader range of tutoring scenarios, such as responding to student errors and assessing student understanding, have yet to be explored. This gap highlights the necessity of broadening the scope of our proposed method. Expanding and rigorously evaluating our approach to encompass diverse educational contexts and lesson types is essential for building a more versatile and universally applicable automated feedback system.

**Enhancing the proposed feedback system with data augmentation.** We also recognized the inherent challenges associated with sequence labeling for highlighting key components of tutoring practice (e.g., praise components in our study). To achieve satisfactory performance, our study required the use of 50% of the total dataset, equivalent to 65 training samples. This substantial annotation workload raises concerns, particularly when considering the extension of fine-tuning

GPT models to more tutor training lessons (e.g. our tutor training platform has designed 20 lessons for different tutoring strategies). To address this issue and reduce the reliance on extensive manual annotation, we are exploring the implementation of data augmentation techniques, such as random swap and synonym replacement [14]. By applying these data augmentation techniques to merely 10% of the dataset or 13 training samples, we aim can reduce the dependency on extensive manual annotation efforts.

**Examining the applicability of the proposed feedback system across different platforms.** In our future work, we aim to apply sequence labeling methods to analyze real-world tutoring transcripts and diverse datasets, such as teacher comments from educational platforms like ASSISTments [22]. Leveraging fine-tuned GPT models on highlighting the key components of instructional strategies (e.g., effective praise, response to student errors, and engaging with difficult students), we plan to generate comprehensive reports that highlight the desired and undesired components from the teacher feedback or comments and provide targeted feedback with suggestions for improvements. This initiative will potentially offer actionable insights to tutors on enhancing their pedagogical approaches in future sessions.

## 7. CONCLUSION

In this study, we investigated the enhancement of automated feedback systems through the application of GPT models, employing a multifaceted approach that included the utilization of prompting GPT-3.5 and GPT-4 models and fine-tuning GPT-3.5 models for improved performance. Prompting GPT models demonstrated their potential in guiding models to identify specific components of praise, emphasizing the critical role of prompt design in optimizing model outputs. In comparison, fine-tuning the GPT-3.5 model, in particular, significantly enhanced the system’s ability to accurately highlight key components from tutor responses. This led to the development of an automated feedback system aimed at delivering immediate and explanatory feedback for tutor training, addressing the crucial need for scalable and effective feedback. Our implementation showcases the potential of leveraging advanced large language models to provide highlighting explanatory feedback on tutors’ open-ended responses, offering insights for future research in the development of automated feedback systems.

## 8. ACKNOWLEDGMENTS

This work is supported by funding from the Richard King Mellon Foundation (Grant #10851) and the Learning Engineering Virtual Institute (<https://learning-engineering-virtual-institute.org/>). Any opinions, findings, and conclusions expressed in this paper are those of the authors. We also wish to express our gratitude to Dr. Ralph Abboud and Dr. Carolyn P. Rosé for their invaluable guidance and recommendations, and to Ms. Yiyang Zhao and Ms. Yuting Wang for their assistance in verifying the rating scheme.

## 9. REFERENCES

- [1] V. Aleven, O. Popescu, and K. R. Koedinger. Towards tutorial dialog to support self-explanation: Adding natural language understanding to a cognitive tutor. In *Proceedings of Artificial Intelligence in Education*, pages 246–255, 2001.
- [2] J. E. Beck, K.-m. Chang, J. Mostow, and A. Corbett. Does help help? introducing the bayesian evaluation and assessment methodology. In *Intelligent Tutoring Systems: 9th International Conference, ITS 2008, Montreal, Canada, June 23-27, 2008 Proceedings 9*, pages 383–394. Springer, 2008.
- [3] S. Bhat, H. A. Nguyen, S. Moore, J. Stamper, M. Sakr, and E. Nyberg. Towards automated generation and evaluation of questions in educational domains. In *Proceedings of the 15th International Conference on Educational Data Mining*, volume 701, 2022.
- [4] A. Brandsen, S. Verberne, M. Wansleben, and K. Lambers. Creating a dataset for named entity recognition in the archaeology domain. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4573–4577, 2020.
- [5] A. P. Cavalcanti, A. Barbosa, R. Carvalho, F. Freitas, Y.-S. Tsai, D. Gašević, and R. F. Mello. Automatic feedback in online learning environments: A systematic literature review. *Computers and Education: Artificial Intelligence*, 2:100027, 2021.
- [6] Y. Chen, L. Wu, Q. Zheng, R. Huang, J. Liu, L. Deng, J. Yu, Y. Qing, B. Dong, and P. Chen. A boundary regression model for nested named entity recognition. *Cognitive Computation*, 15(2):534–551, 2023.
- [7] D. R. Chine, P. Chhabra, A. Adeniran, S. Gupta, and K. R. Koedinger. Development of scenario-based mentor lessons: an iterative design process for training at scale. In *Proceedings of the Ninth ACM Conference on Learning@ Scale*, pages 469–471, 2022.
- [8] D. R. Chine, P. Chhabra, A. Adeniran, J. Kopko, C. Tipper, S. Gupta, and K. R. Koedinger. Scenario-based training and on-the-job support for equitable mentoring. In *The Learning Ideas Conference*, pages 581–592. Springer, 2022.
- [9] W. Dai, J. Lin, H. Jin, T. Li, Y.-S. Tsai, D. Gašević, and G. Chen. Can large language models provide feedback to students? a case study on chatgpt. In *2023 IEEE International Conference on Advanced Learning Technologies (ICALT)*, pages 323–325. IEEE, 2023.
- [10] W. Dai, Y.-S. Tsai, J. Lin, A. Aldino, H. Jin, T. Li, D. Gašević, and G. Chen. Assessing the proficiency of large language models in automatic feedback generation: An evaluation study. 2024.
- [11] L. Deleger, Q. Li, T. Lingren, M. Kaiser, K. Molnar, L. Stoutenborough, M. Kouril, K. Marsolo, I. Solti, et al. Building gold standard corpora for medical natural language processing tasks. In *AMIA Annual Symposium Proceedings*, volume 2012, page 144. American Medical Informatics Association, 2012.
- [12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the NAACL-HLT, Volume 1*, pages 4171–4186, 2019.
- [13] J. Dietrichson, M. Bøg, T. Filges, and A.-M. Klinton Jørgensen. Academic interventions for elementary and middle school students with low socioeconomic status: A systematic review and meta-analysis. *Review of educational research*, 87(2):243–282, 2017.

- [14] S. Y. Feng, V. Gangal, J. Wei, S. Chandar, S. Vosoughi, T. Mitamura, and E. Hovy. A survey of data augmentation approaches for nlp. In *Findings of the ACL: ACL-IJCNLP 2021*, pages 968–988, 2021.
- [15] N. Gisev, J. S. Bell, and T. F. Chen. Interrater agreement and interrater reliability: key concepts, approaches, and applications. *Research in Social and Administrative Pharmacy*, 9(3):330–338, 2013.
- [16] C. Grouin, S. Rosset, P. Zweigenbaum, K. Fort, O. Galibert, and L. Quintard. Proposal for an extension of traditional named entities: From guidelines to evaluation, an overview. In *Proceedings of the 5th linguistic annotation workshop*, pages 92–100, 2011.
- [17] A. Gurung, S. Baral, M. P. Lee, A. C. Sales, A. Haim, K. P. Vanacore, A. A. McReynolds, H. Kreisberg, C. Heffernan, and N. T. Heffernan. How common are common wrong answers? crowdsourcing remediation at scale. In *Proceedings of the Tenth ACM Conference on Learning@ Scale*, pages 70–80, 2023.
- [18] A. Gurung, S. Baral, K. P. Vanacore, A. A. McReynolds, H. Kreisberg, A. F. Botelho, S. T. Shaw, and N. T. Hefferna. Identification, exploration, and remediation: Can teachers predict common wrong answers? In *LAK23: 13th International Learning Analytics and Knowledge Conference*, pages 399–410, 2023.
- [19] J. Guryan, J. Ludwig, M. P. Bhatt, P. J. Cook, J. M. Davis, K. Dodge, G. Farkas, R. G. Fryer Jr, S. Mayer, H. Pollack, et al. Not too late: Improving academic outcomes among adolescents. *American Economic Review*, 113(3):738–765, 2023.
- [20] Z. F. Han, J. Lin, A. Gurung, D. R. Thomas, E. Chen, C. Borchers, S. Gupta, and K. R. Koedinger. Improving assessment of tutoring practices using retrieval-augmented generation, 2024.
- [21] J. Hattie and H. Timperley. The power of feedback. *Review of Educational Research*, 2007.
- [22] N. T. Heffernan and C. L. Heffernan. The assistments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education*, 24:470–497, 2014.
- [23] M. Henderson, R. Ajjawi, D. Boud, and E. Molloy. *The Impact of Feedback in Higher Education: Improving assessment outcomes for learners*. Springer Nature, 2019.
- [24] D. Hirunyasiri, D. R. Thomas, J. Lin, K. R. Koedinger, and V. Alevan. Comparative analysis of gpt-4 and human graders in evaluating praise given to students in synthetic dialogues. *arXiv preprint arXiv:2307.02018*, 2023.
- [25] L. N. Jenkins, M. T. Floress, and W. Reinke. Rates and types of teacher praise: A review and future directions. *Psychology in the Schools*, 52(5):463–476, 2015.
- [26] D. Jurafsky and J. H. Martin. *Speech and language processing*. 3rd, 2022.
- [27] S. Kakarla, D. Thomas, J. Lin, S. Gupta, and K. R. Koedinger. Using large language models to assess tutors’ performance in reacting to students making math errors, 2024.
- [28] K. S. Kalyan. A survey of gpt-3 family large language models including chatgpt and gpt-4. *Natural Language Processing Journal*, page 100048, 2023.
- [29] M. L. Kamins and C. S. Dweck. Person versus process praise and criticism: implications for contingent self-worth and coping. *Developmental psychology*, 35(3):835, 1999.
- [30] E. Kasneci, K. Seßler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, U. Gasser, G. Groh, S. Günemann, E. Hüllermeier, et al. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274, 2023.
- [31] M. Konkol and M. Konopík. Segment representations in named entity recognition. In *International conference on text, speech, and dialogue*, pages 61–70. Springer, 2015.
- [32] M. A. Kraft and G. T. Falken. A blueprint for scaling tutoring and mentoring across public schools. *AERA Open*, 7:23328584211042858, 2021.
- [33] E. Latif and X. Zhai. Fine-tuning chatgpt for automatic scoring. *Computers and Education: Artificial Intelligence*, page 100210, 2024.
- [34] Z. Levonian, C. Li, W. Zhu, A. Gade, O. Henkel, M.-E. Postle, and W. Xing. Retrieval-augmented generation to improve math question-answering: Trade-offs between groundedness and human preference, 2023.
- [35] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.
- [36] J. Li, A. Sun, J. Han, and C. Li. A survey on deep learning for named entity recognition. *IEEE TKDE*, 34(1):50–70, 2020.
- [37] J. Lin, W. Dai, L.-A. Lim, Y.-S. Tsai, R. F. Mello, H. Khosravi, D. Gašević, and G. Chen. Learner-centred analytics of feedback content in higher education. In *LAK23*, pages 100–110, 2023.
- [38] J. Lin, Z. Han, D. R. Thomas, A. Gurung, S. Gupta, V. Alevan, and K. R. Koedinger. Leveraging large language models to enhance feedback provision in tutor training program. *International Journal of Artificial Intelligence in Education*, 2024.
- [39] J. Lin, S. Singh, L. Sha, W. Tan, D. Lang, D. Gašević, and G. Chen. Is it a good move? mining effective tutoring strategies from human–human tutorial dialogues. *Future Generation Computer Systems*, 127:194–207, 2022.
- [40] J. Lin, D. R. Thomas, F. Han, S. Gupta, W. Tan, N. D. Nguyen, and K. R. Koedinger. Using large language models to provide explanatory feedback to human tutors. *arXiv preprint arXiv:2306.15498*, 2023.
- [41] C. Liu, H. Fan, and J. Liu. Span-based nested named entity recognition with pretrained language model. In *Database Systems for Advanced Applications: 26th International Conference, DASFAA 2021, Taipei, Taiwan, April 11–14, 2021, Proceedings, Part II 26*, pages 620–628. Springer, 2021.

- [42] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35, 2023.
- [43] H. Luo, W. Tan, N. D. Nguyen, and L. Du. Re-weighting tokens: A simple and effective active learning strategy for named entity recognition. *arXiv preprint arXiv:2311.00906*, 2023.
- [44] M. L. McHugh. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3):276–282, 2012.
- [45] H. McNichols, W. Feng, J. Lee, A. Scarlatos, D. Smith, S. Woodhead, and A. Lan. Exploring automated distractor and feedback generation for math multiple-choice questions via in-context learning. *arXiv preprint arXiv:2308.03234*, 2023.
- [46] N. D. Nguyen, W. Tan, L. Du, W. Buntine, R. Beare, and C. Chen. Auc maximization for low-resource named entity recognition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):13389–13399, Jun. 2023.
- [47] N. D. Nguyen, W. Tan, L. Du, W. Buntine, R. Beare, and C. Chen. Low-resource named entity recognition: Can one-vs-all auc maximization help? In *2023 IEEE International Conference on Data Mining (ICDM)*, pages 1241–1246. IEEE, 2023.
- [48] A. Nickow, P. Oreopoulos, and V. Quan. The impressive effects of tutoring on prek-12 learning: A systematic review and meta-analysis of the experimental evidence. 2020.
- [49] A. Pardo, K. Bartimote, S. B. Shum, S. Dawson, J. Gao, D. Gašević, S. Leichtweis, D. Liu, R. Martínez-Maldonado, N. Mirriahi, et al. Ontask: Delivering data-informed, personalized learning support actions. *Journal of Learning Analytics*, 5(3):235–249, 2018.
- [50] T. Patikorn and N. T. Heffernan. Effectiveness of crowd-sourcing on-demand assistance from teachers in online learning platforms. In *Proceedings of the Seventh ACM Conference on Learning@ Scale*, pages 115–124, 2020.
- [51] C. Pornprasit and C. Tantithamthavorn. Gpt-3.5 for code review automation: How do few-shot learning, prompt design, and model fine-tuning impact their performance? *arXiv preprint arXiv:2402.00905*, 2024.
- [52] J. Reich. Teaching drills: Advancing practice-based teacher education through short, low-stakes, high-frequency practice. *Journal of Technology and Teacher Education*, 30(2):217–228, 2022.
- [53] T. Ryan, M. Henderson, K. Ryan, and G. Kennedy. Designing learner-centred text-based feedback: a rapid review and qualitative synthesis. *Assessment & Evaluation in Higher Education*, 46(6):894–912, 2021.
- [54] A. Shrivastava and J. Heer. Iseql: Interactive sequence learning. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*, pages 43–54, 2020.
- [55] D. Thomas, X. Yang, S. Gupta, A. Adeniran, E. Mclaughlin, and K. Koedinger. When the tutor becomes the student: Design and evaluation of efficient scenario-based lessons for tutors. In *LAK23: 13th International Learning Analytics and Knowledge Conference*, pages 250–261, 2023.
- [56] Y. Zhou, A. I. Muresanu, Z. Han, K. Paster, S. Pitis, H. Chan, and J. Ba. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*, 2022.

## APPENDIX

### A. LESSON PRINCIPLES

The following is the principle that a correct response should follow:

Praising students for working hard and putting forth effort is a great way to increase student motivation. When the learning gets tough, giving correct praise is a powerful strategy to encourage students to keep going.

The correct response should be :

- perceived as sincere, earned, and truthful.
- specific by giving details of what the student did well.
- immediate with praise given right after the student action.
- authentic and is not repeated often, such as “great job” which loses meaning and becomes predictable.
- focused on the learning process, not ability (AJTutoring.com, 2022)

Correct responses must follow some, but not all the above.

There are two types of praise responses: Effort and Outcome praise

- Effort praise focuses on the learning process. Effort praise recognizes students for putting forth effort and persevering through the learning process instead of focusing on whether a student got the problem correct or pure ability.

- Outcome praise showcases student’s achievements, such as getting a grade A on an assignment or getting a problem correct, and is often, but not always, associated with unproductive praise.

To receive full credit of correct praise, tutors cannot just say “great job” and praise with no specific reasoning. Tutors need to praise for effort AND be positive and encouraging.

### B. INPUT FOR FINE-TUNING GPT-3.5

Table 6: Input for fine-tuning GPT 3.5 model

Role	Content
<b>System</b>	<i>You are a response evaluator designed to output JSON. Your task is to analyze tutor responses based on the principles of effective praise focusing on ‘effort’ and ‘outcome’. Extract words or phrases that represent praise for the student’s effort and outcome, and output the results in JSON format with keys titled ‘Effort’ and ‘Outcome’.</i>
<b>User</b>	<b>Lesson Principle</b>
<b>Assistant</b>	<i>Sure, can you provide a tutor response for analysis</i>
<b>User</b>	<i>An example of outcome-based praise is: “Great job! You are a genius!”</i>
<b>Assistant</b>	<i>An output json format is: {“effort”: [], “outcome”: [“Great job”]}</i>
<b>User</b>	<i>Nice, let’s do it again.</i>
<b>Assistant</b>	<i>Sure, can you provide a tutor response for analysis?</i>
<b>User</b>	<i>An example of effort-based praise is: “You are almost there! I am proud of how you are persevering through and striving to solve the problem. Keep going!”</i>
<b>Assistant</b>	<i>An output json format is: {“effort”: [“persevering through and striving to solve the problem”, “Keep going”], “outcome”: []}</i>
<b>User</b>	<i>Nice, let’s do it again.</i>
<b>Assistant</b>	<i>Sure, can you provide a tutor response for analysis</i>
<b>User</b>	<b>Tutor Response</b>
<b>Assistant</b>	<b>Praise Type and Content</b>

*Note: Praise Type and Content:* This part simulates an interactive environment where the model plays the role of a response evaluator. The conversation flow is designed to mimic a real-world interaction, with system and user roles alternately providing context, instruction, and input (the tutor response) for processing.



### C. SCATTER MATRIX OF THE CORRELATION ON THE OUTCOME-BASED PRAISE

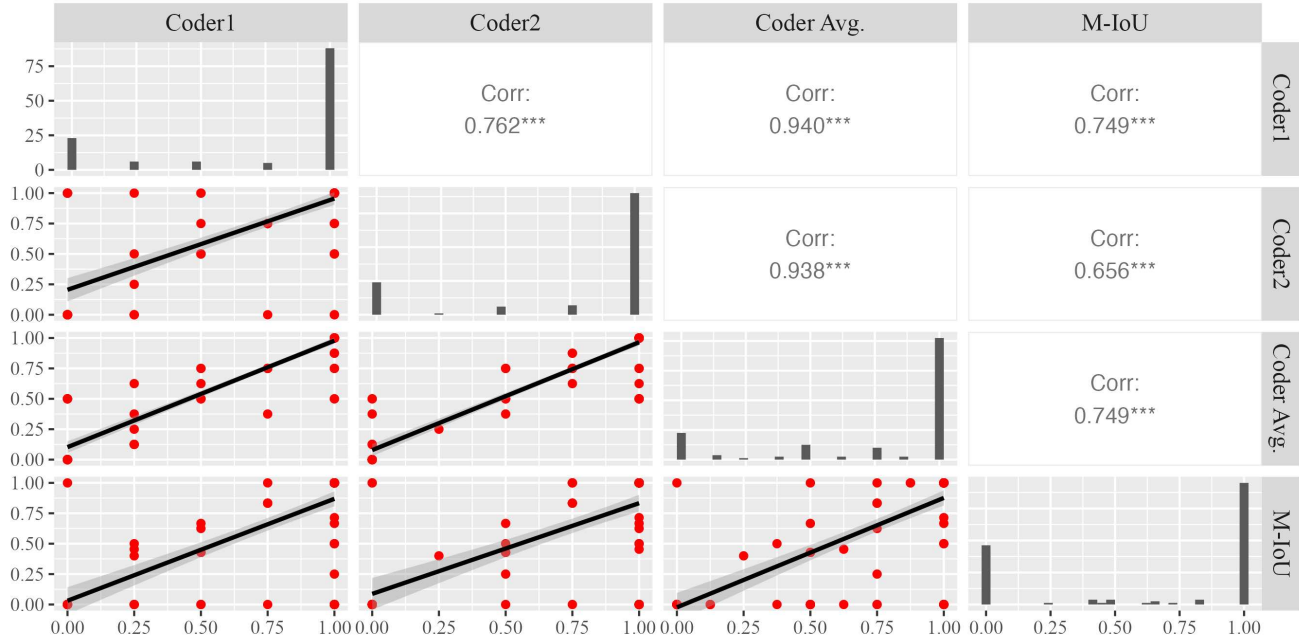


Figure 5: Scatter matrix of the correlation on the outcome-based praise

### D. DETAILED RESULTS OF FINE-TUNED GPT-3.5 MODELS' PERFORMANCE

Table 7: Detailed results of fine-tuned GPT-3.5 models on identifying praise components

Training size	Effort				Outcome			
	Mean	Std.	Min.	Max.	Mean	Std.	Min.	Max.
<b>13 (10%)</b>	0.51	0.05	0.44	0.58	0.62	0.05	0.56	0.69
<b>26 (20%)</b>	0.55	0.03	0.51	0.58	0.64	0.04	0.59	0.68
<b>39 (30%)</b>	0.56	0.05	0.46	0.59	0.69	0.06	0.59	0.75
<b>52 (40%)</b>	0.58	0.04	0.54	0.63	0.66	0.05	0.60	0.73
<b>65 (50%)</b>	0.60	0.04	0.54	0.64	0.76	0.07	0.66	0.84