

Beyond Agreement: Diagnosing the Rationale Alignment of Automated Essay Scoring Methods based on Linguistically-informed Counterfactuals

Yupei Wang[†] Renfen Hu^{†*} Zhe Zhao[‡]

[†] Beijing Normal University

[‡] Tencent AI Lab

{wangyupei, irishu}@mail.bnu.edu.cn, nlpzhezhaotencent.com

Abstract

While current Automated Essay Scoring (AES) methods demonstrate high scoring agreement with human raters, their decision-making mechanisms are not fully understood. Our proposed method, using counterfactual intervention assisted by Large Language Models (LLMs), reveals that BERT-like models primarily focus on sentence-level features, whereas LLMs such as GPT-3.5, GPT-4 and Llama-3 are sensitive to conventions & accuracy, language complexity, and organization, indicating a more comprehensive rationale alignment with scoring rubrics. Moreover, LLMs can discern counterfactual interventions when giving feedback on essays. Our approach improves understanding of neural AES methods and can also apply to other domains seeking transparency in model-driven decisions¹.

1 Introduction

In recent years, neural approaches to Automated Essay Scoring (AES) have demonstrated remarkable performance (Ke and Ng, 2019; Ramesh and Sanampudi, 2022). The advent of Large Language Models (LLMs) has shifted focus not only towards their scoring capabilities but also towards the potential for providing feedback (Mizumoto and Eguchi, 2023; Caines et al., 2023; Han et al., 2023; Xiao et al., 2024), enabling a better understanding of the models' rationale. However, current model evaluations mainly use metrics such as Quadratic Weighted Kappa (QWK) to measure agreement with human ratings. This approach leaves the models' underlying reasoning opaque, thereby raising risks and questioning the validity of their use in high-stakes educational tests (Fiacco et al., 2023).

A series of studies have found that neural models can be *right for the wrong reasons*, a concern that

persists into the era of LLMs (McCoy et al., 2020; Turpin et al., 2023). To understand the decision-making basis of neural models, researchers have primarily adopted two primary avenues: *what* knowledge a model encodes and *why* a model makes certain predictions (Lyu et al., 2024). Both paradigms have garnered attention in the field of AES. Fiacco et al. (2023) addresses the *what* question by extracting meaningful functional groups from the representations of transformer models and aligning them with human-understandable features. However, a model encodes a myriad of features does not mean that the features are utilized in decision-making (Lyu et al., 2024). To tackle the *why* question, Singla et al. (2023) employed integrated gradients (Sundararajan et al., 2017) to analyze token importance, and discovered that for BERT-based model, most of the attributions are over non-linguistic tokens and stop words. It can be seen that the gradients-based methods only target lower-level token features, thus failing to reveal whether models leverage higher-level linguistic features. Moreover, both Kabra et al. (2022) and Singla et al. (2023) employed adversarial modifications to assess models, but these interventions did not target the linguistic features critical to the AES task, and they did not control for other variables that could affect essay scores during modification. Therefore, even for traditional AES models, reliable explanations of their inner workings remain elusive. Additionally, the explainability of scoring in LLMs is largely unexplored, indicating considerable work is needed to advance our understanding of model reasoning within this domain.

In this paper, we systematically investigate whether models adhere to scoring rubrics when producing scores and feedback on essays, thereby assessing their alignment with human rationale. To quantify this alignment, we introduce the concept of *rationale alignment*, measured by the difference in scores and feedback between original essays and

*Corresponding author.

¹We release our code at <https://github.com/YpLarryWang/beyond-agreement-aes-2024>.

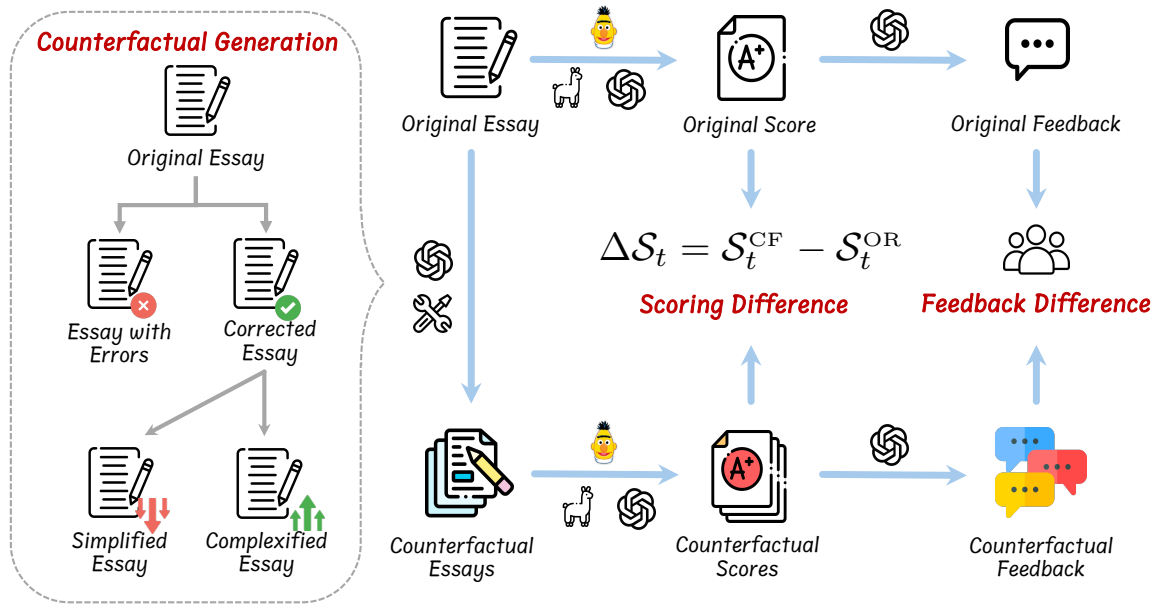


Figure 1: The pipeline of our proposed method.

their counterfactuals. As shown in Figure 1, we propose a model-agnostic diagnosis method that uses linguistically-informed counterfactuals to scrutinize the scoring behavior of BERT-like models and LLMs. The diagnostic approach closely integrates linguistic knowledge from scoring rubrics, such as conventions, accuracy, vocabulary, syntax, and coherence, with LLMs employed for fine-grained and controllable counterfactual generation.

Our investigation reveals that: (1) BERT-like models can discern differences in conventions and language complexity but struggle to grasp the logical structure and coherence of essays; and (2) LLMs, although have lower score agreement than traditional models, display a superior alignment with human experts’ reasoning during scoring and can also address counterfactual interventions in their feedback. Through few-shot learning or fine-tuning, LLMs can achieve both high scoring agreement and rationale alignment.

2 Related Work

2.1 AES based on Neural Language Models

Pre-trained neural language models have made significant progress in the field of AES. After fine-tuning on specific datasets, these models can achieve high levels of agreement with human raters (Rodriguez et al., 2019; Yang et al., 2020; Ormerod et al., 2021; Wang et al., 2022). Since the emergence of ChatGPT, the scoring performance of LLMs has garnered considerable attention. Leveraging their powerful language understanding capa-

bilities and in-context learning abilities, LLMs can evaluate essays and assign overall scores or scores for specific dimensions (Naismith et al., 2023). However, research has shown that zero-shot and few-shot LLMs fail to achieve state-of-the-art scoring performance (Mizumoto and Eguchi, 2023), while fine-tuned LLM models exhibit notable superiority (Xiao et al., 2024).

Although the scoring ability of LLMs without fine-tuning is not particularly remarkable, they can provide explainable feedback in natural language. Previously, essay feedback was primarily provided through trait scores (e.g., vocabulary) (Carlile et al., 2018; Hussein et al., 2020; Lee et al., 2023). With the emergence of LLMs, researchers discovered that it is possible to elicit explanations about assessment decisions from the models (Caines et al., 2023). Han et al. (2023) assessed the feedback generated by GPT-3.5 on level of detail, accuracy, relevance, and helpfulness, while Xiao et al. (2024) found that GPT-4 feedback could elevate novice raters to expert levels.

2.2 Interpretability and Robustness of AES Models

In terms of model interpretability in AES research, Fiacco et al. (2023) analyzed the features encoded by transformer models, but this approach provides limited insight into the decision-making rationale of the models. Singla et al. (2023) employed the integrated gradients (IG) method (Sundararajan et al., 2017) on neural models to analyze token-level feature importance and discovered that BERT-based

models frequently assign substantial importance to stopwords and non-linguistic tokens. This counterintuitive result may stem from the fact that the IG method does not address interactions between tokens, thereby failing to capture abstract linguistic concepts such as cohesion and syntax. Moreover, these methods cannot be directly applied to closed-source models like GPT-3.5 and GPT-4.

Additionally, a line of works have utilized adversarial modifications to diagnose model robustness. Powers et al. (2002) invited human writers to compose essays that would "trick" the AES system and found that repeating, rewording, and reordering were effective strategies. Bejar et al. (2014) employed the substitution of words with less frequent and longer synonyms. Kabra et al. (2022) used methods such as the addition of irrelevant lines, the introduction of grammatical errors, and the deletion of lines from the responses. Myers and Wilson (2023) evaluated models using a sentence-level randomization approach. It is important to note that these studies aim to expose model vulnerabilities by introducing input perturbations rather than exploring the interpretability of model decisions.

2.3 Counterfactual Analysis

Counterfactuals are hypothetical scenarios created to understand the causal effects of specific interventions in a given situation (Feder et al., 2022). Existing counterfactual generation methods utilize keyword replacement (Garg et al., 2019), sentence rewriting (Ross et al., 2022; Wu et al., 2021), and manual editing (Gardner et al., 2020). However, these approaches are often limited to simple local interventions or require costly manual annotation, which hinders the practical estimation of the causal effects of high-level concepts on NLP models. While recent efforts have leveraged LLMs for generating more natural and diverse counterfactuals (Dixit et al., 2022; Chen et al., 2023), most have only exploited LLMs' powerful language generation capabilities without tapping into their potential to understand and manipulate abstract concepts within texts. Gat et al. (2023) found that LLMs can produce high-quality counterfactuals, which assist in providing strong black-box model explanations. Li et al. (2024) prompted LLMs to identify and modify *causal terms* to generate counterfactuals. Inspired by these works, we decided to combine LLMs with rule-based methods to achieve controlled sample generation in AES.

3 Method

We employed counterfactual interventions to establish causality between target concepts and predicted scores. Typically, counterfactual intervention involves manipulating a specific feature or concept while controlling for others and observing the subsequent effect on the model's prediction. We firstly extracted target concepts from scoring rubrics for intervention, and then generated counterfactual samples for different concepts using LLMs and heuristic rules.

3.1 Concepts for Intervention

To identify the target concepts for AES scenarios, we reviewed scoring rubrics from major standardized English tests (IELTS, TOEFL iBT, TOEIC, PTE Academic) and the ELLIPSE dataset (Crossley et al., 2023), which is based on various state and industrial English language proficiency assessments. We conducted a detailed annotation process to identify common linguistic features across the five rubrics. See Appendix A for more information. Through this analysis, we discovered that all the scoring criteria consistently emphasize three key aspects:

Conventions and Accuracy: An essay is considered to adhere to conventions and demonstrate accuracy when it is free from mechanical (spelling, capitalization, and punctuation) mistakes and grammatical inaccuracies.

Language Complexity: An essay demonstrates lexical and syntactic complexity through the use of a broad vocabulary, sophisticated lexical control, and varied sentence structures.

Organization and Development: An essay exhibits effective organization and development by presenting a logical structure with skillful paragraphing and the use of cohesive devices to ensure unity, progression, and seamless connection of thoughts.

3.2 Measurement of Rationale Alignment

We introduce two methods—*score differences* and *feedback differences*—to measure rationale alignment, assessing how well models adhere to human-defined scoring rubrics.

Score Differences: For each concept C_i , let $O_i = \{o_1^{(i)}, o_2^{(i)}, \dots, o_k^{(i)}\}$ be the set of possible types of interventions. Each $o_j^{(i)}$ represents a specific way to alter the value v_i of concept C_i in essay t , resulting in a modified value $v'_i = o_j^{(i)}(v_i)$.

Concept	Intervention	Description
Conventions & Accuracy	Error Correction	Prompt GPT-4 Turbo to correct spelling, punctuation, and grammar errors.
	Spelling Errors Introduction	Use nlpaug (Ma, 2019) to misspell 30% of words in 50% of sentences.
	Agreement Errors Introduction	Introduce subject-verb agreement (SVA) errors in 50% of sentences.
Language Complexity	Word Order Swapping (WOS)	Use nlpaug to swap 30% of words in 50% of sentences.
	Complexification	Prompt GPT-4 Turbo to enhance vocabulary and sentence structure.
Organization & Development	Simplification	Prompt GPT-4 Turbo to simplify vocabulary and sentence structure.
	Intra-paragraph Shuffling	Shuffle sentence order <i>within paragraphs</i> to disrupt <i>local</i> cohesion.
	Inter-text Shuffling	Shuffle sentence order <i>across the entire essay</i> to disrupt <i>global</i> cohesion.

Table 1: Overview of [positive](#) and [negative](#) counterfactual intervention methods used.

For each essay t in the test set \mathcal{T} , we generate multiple counterfactual versions by altering the values of different concepts $\{C_i\}$ using various interventions $\{o_j^{(i)}\}$. We use a model \mathcal{M} to predict scores for both the original essay and its counterfactuals. The effect of a specific counterfactual intervention on model \mathcal{M} with respect to essay t is calculated by subtracting the original score from the counterfactual score:

$$\Delta\mathcal{S}_t^{\mathcal{M}}(C_i, o_j^{(i)}) = \mathcal{S}_t^{\mathcal{M}}(C_i = v'_i) - \mathcal{S}_t^{\mathcal{M}}(C_i = v_i) \quad (1)$$

We then compute the mean effect of this specific intervention on model \mathcal{M} across all essays in \mathcal{T} :

$$\Delta\mathcal{S}_{\mathcal{T}}^{\mathcal{M}}(C_i, o_j^{(i)}) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \Delta\mathcal{S}_t^{\mathcal{M}}(C_i, o_j^{(i)}) \quad (2)$$

Feedback Differences: We manually compare the feedback provided for both the original and counterfactual essays. This helps us understand shifts in the model’s reasoning and justification.

3.3 Counterfactual Generation

We employ a hybrid approach combining rule-based and LLM-based methods to generate eight types of linguistically informed counterfactuals for diagnostic purposes, as detailed in Table 1. These interventions derive from three aforementioned linguistic concepts and are implemented in both positive and negative directions for conventions and language complexity. As shown in Figure 1, for conventions and accuracy, we introduce errors such as spelling, subject-verb agreement, and word order for negative impacts, and use LLMs to correct all errors for positive impacts. Regarding the language complexity, we leverage LLMs to increase and decrease the language complexity along both vocabulary and syntax dimensions, building upon the basis of error correction. For the organizational aspect, negative interventions include disrupting the sentence order within paragraphs to affect local coherence and across the entire article to impact global coherence. See Appendix B.2 for LLM prompts used to generate counterfactuals.

3.4 The Validity of LLM Generated Counterfactuals

As shown in Table 1, we prompted LLMs to correct errors, complexify, and simplify essays to manipulate their conventions and language complexity. To evaluate counterfactual essays generated by LLMs, we introduced seven linguistic metrics that measure the essay length, lexical diversity, lexical sophistication, syntactic complexity and writing error density. The descriptions of these metrics can be seen in Table 2. We compute Cohen’s \mathcal{D} (Cohen, 2013) effect size for each metric as follows:

$$\mathcal{D} = \frac{\bar{x}_{CF} - \bar{x}_{OR}}{s} \quad (3)$$

where \bar{x}_{CF} and \bar{x}_{OR} are the mean values of a metric for the counterfactual and original samples, and the pooled standard deviation s is defined as:

$$s = \sqrt{\frac{(n_{OR} - 1)s_{OR}^2 + (n_{CF} - 1)s_{CF}^2}{n_{OR} + n_{CF} - 2}} \quad (4)$$

where n_{OR} and n_{CF} are the sample sizes, and s_{OR}^2 and s_{CF}^2 are the variances of the original and counterfactual samples respectively.

Moreover, we assess content preservation during interventions by calculating the average cosine similarity between text embeddings of "original-counterfactual" essay pairs for each of the three LLM-based interventions.

4 Experiments

4.1 Settings

Our study utilized TOEFL11 (Blanchard et al., 2013) and ELLIPSE datasets. TOEFL11 includes 12,100 essays from the 2006-2007 TOEFL exams, divided into 9,900 for training, 1,100 for validation, and 1,100 for testing, with essays categorized into low, medium, or high proficiency by human raters. We assessed performance using weighted F1 and quadratic weighted kappa (QWK). The ELLIPSE dataset contains 6,482 essays from 8th to 12th-grade English learners, with 2,568 reserved for testing. Essays were rated on a 1 to 5 scale (with

Metric	Description
WordNum	The number of words in an essay.
SentNum	The number of sentences in an essay.
MLS	Mean length of sentences. The length of each sentence is the number of words it has.
ADDT	Average depth of dependency tree for all sentences in an essay.
LemmaTTR	A <i>lexical diversity</i> measure based on the Type-Token Ratio (TTR) of an essay, where each word is lemmatized.
LexSoph	A <i>lexical sophistication</i> measure based on word frequency statistics from the 1980s-2010s COHA corpus (Davies, 2010). For an essay with N words, let w_1, w_2, \dots, w_N be the individual words (including repetitions), ℓ_i be the lemma of w_i , and $\text{Freq}(\ell_i)$ be the frequency of ℓ_i in the selected COHA subset. LexSoph is defined as: $\frac{1}{N} \sum_{i=1}^N \frac{1}{\log(\text{Freq}(\ell_i) + 1)}$
ErrorDensity	Density of writing errors in an essay with N words, defined as $\#\text{error}/N$. Writing error analyses are implemented using LanguageTool (Naber et al., 2003).

Table 2: The linguistics metrics used for the evaluation of counterfactual samples.

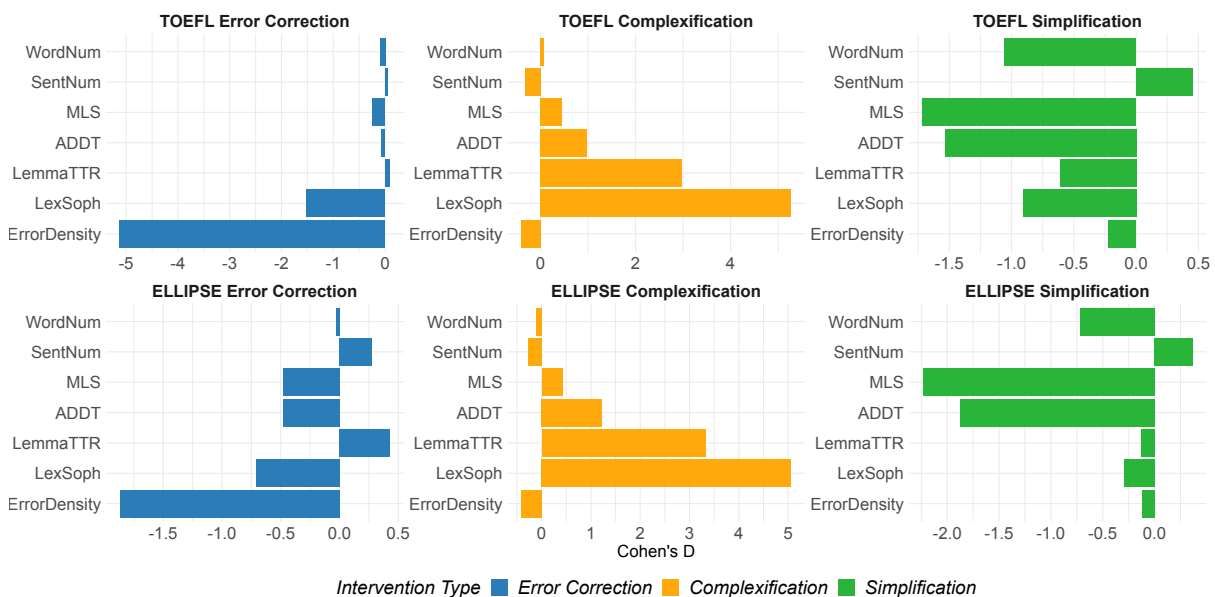


Figure 2: Cohen's \mathcal{D} measured for seven linguistic metrics on three interventions.

0.5 increments), adjusted to the nearest 0.5 for QWK calculations, alongside Root Mean Square Error (RMSE) evaluation.

The counterfactuals were generated on the test set using GPT-4 Turbo and Llama-3-70b-Instruct models. Comparative analysis revealed that both models successfully completed the task, but the GPT-4 Turbo model exhibited more stable performance in the aforementioned measures across both datasets (see Appendix B.3 for detailed comparisons). Consequently, we employed the counterfactual essays generated by the GPT-4 Turbo model for subsequent analyses.

For automated scoring, we fine-tuned BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) and DeBERTa (He et al., 2021) on the training set. For LLMs, we utilized GPT-3.5 Turbo, GPT-4 Turbo and Llama 3 instruction-fine-tuned models (8B & 70B) in zero-shot learning (ZSL) and few-shot

learning (FSL) scenarios, and performed supervised fine-tuning (SFT) on GPT-3.5 Turbo. Detailed fine-tuning and inference settings are provided in Appendix C.

Since score levels of TOEFL11 essays are labeled as three discrete categories²: low, medium and high, we calculated the score delta of BERT-like models after converting the classifier's output probabilities into 1-5 scores for consistency with TOEFL scoring rubrics scale:

$$\mathcal{S}_t^{\mathcal{M}} = \sum_{k=1}^3 (2k - 1) \Pr^{\mathcal{M}}(\hat{y}_t = k) \quad (5)$$

where \hat{y}_t is the predicted class of essay t .

²When few-shot prompting LLMs to rate TOEFL11 essays, score levels of example essays in prompts can only be one of the three categories. Therefore it is not possible to expect numerical predicted scores, which is the reason why in Table 6 we do not show intervention effects of FSL on TOEFL11.

Original: In the **up coming** years **i wanna** save up and work hard to **get** money, so that **i** can open my own personal store.. \n\n in the store **i going to be** selling high **branded** shoes,clothes, and **outer collectives**. and **Im** going to **learning** how to invest my money as **i** sell things before opening up my dream store. \n\n **i wanna** do this as a job because **ive awlays had a passhine** for shoes and high **branded** things. and **if i** can pull this off **i think ill be veary succes full** in life **i will be getting and veary good salary ill be able to pay off anything and have to money** to support my family.

Error Correction: "In the **upcoming** years, **I want to** save up and work hard to **earn** money, so that **I** can open my own personal store. \n\n In the store, **I am going to sell** high-branded shoes, clothes, and **other collectibles**. **I'm also** going to **learn** how to invest my money as **I** sell things before opening up my dream store. \n\n **I want to** do this as a job because **I've always had a passion** for shoes and high-branded **items**. **If I** can pull this off, **I think I'll be very successful** in life. **I will be earning a very good salary, I'll be able to pay off anything, and have the money** to support my family."

Complexification: In the forthcoming years, my ambition is to accumulate sufficient capital and diligently strive to generate income, thereby enabling me to inaugurate my own boutique. Within this emporium, I intend to offer an array of premium-branded footwear, apparel, and assorted collectibles. Concurrently, I shall endeavor to acquire financial acumen, particularly in investment strategies, as I engage in commerce prior to the launch of my envisioned enterprise. Pursuing this career path is fueled by an enduring zeal for exclusive footwear and luxury goods. Should I succeed in actualizing this venture, I am confident that it will herald substantial success in my life. Not only will I secure a lucrative income, but I will also possess the financial capacity to discharge any obligations and provide for my family's needs.

Simplification: In the next few years, I plan to save money and work hard. I want to make enough money to open my own store. In my store, I will sell popular shoes, clothes, and other special items. I will also learn how to use my money well as I sell things. I want to open my dream store one day. I have always loved shoes and popular brands. If I can do this, I think I will do very well in life. I will make a lot of money. I will be able to pay for everything and have enough to take care of my family.

Table 3: Example of a 3-point ELLIPSE essay and its counterfactual counterparts generated by GPT-4 Turbo.

4.2 Counterfactual Validation Results

Figure 2 shows the effect size of three GPT-4-based interventions on seven linguistic metrics across two datasets. Both datasets show similar patterns, with Error Correction notably reducing error density and leaving lexical and syntactic complexity almost untouched. Complexification significantly boosts lexical diversity and sophistication while moderately enhancing syntactic complexity, without substantially changing overall text length and error density. Similarly, Simplification effectively reduces sentence length and complexity while also making corresponding changes to lexical properties.

Table 4 presents the embedding similarities between counterfactuals and original essays. It can be seen that Error Correction almost completely retains the original meaning, and Complexification and Simplification, although inevitably making more changes to the original text, still retain most of the original meaning. For better clarification, Table 3 shows counterfactual examples of a 3-point ELLIPSE essay generated by GPT-4 Turbo. Examples of its rule-based counterfactuals can be seen in Appendix B.1.

Intervention	TOEFL11	ELLIPSE
Error Correction	0.935	0.942
Complexification	0.760	0.749
Simplification	0.816	0.849

Table 4: Content preservation for GPT-4-based interventions: text cosine similarities computed by OpenAI text-embedding-3-large.

4.3 Scoring Results

Table 5 displays the performance of scoring agreement between models and human on the test sets of both datasets. Note that the GPT-3.5-SFT-100 set-

ting is fine-tuned on 100 essays randomly selected from the overall training set, ensuring a stratified distribution based on essay scores. Table 6 shows intervention effects of different types of counterfactual interventions. Based on these results, our findings are as follows:

Setting	TOEFL11		ELLIPSE	
	F1 ↑	QWK ↑	RMSE ↓	QWK ↑
BERT	0.783	0.736	0.437	0.680
ROBERTA	0.795	0.739	0.430	0.695
DEBERTA	0.790	0.741	0.422	0.720
GPT-3.5-ZSL	0.599	0.408	0.701	0.399
GPT-3.5-FSL	0.546	0.314	<u>0.570</u>	0.378
GPT-3.5-SFT-100	0.710	0.592	0.550	0.629
GPT-4-ZSL	0.368	0.380	0.960	0.261
GPT-4-FSL	0.490	0.477	0.680	0.466
LLAMA-3-8B-ZSL	0.558	0.297	0.628	0.345
LLAMA-3-8B-FSL	0.435	0.441	1.039	0.054
LLAMA-3-70B-ZSL	0.524	0.390	0.903	0.182
LLAMA-3-70B-FSL	<u>0.609</u>	<u>0.562</u>	0.589	<u>0.503</u>

Table 5: The scoring agreement performance on both test sets: **best setting** in bold, **fine-tuned GPT-3.5** with a green shadow, **best off-the-shelf LLMs** underlined. Metrics with ↑ indicate that higher values are better, while the one with ↓ indicates that lower values are better.

Firstly, BERT-like models show higher scoring agreement with human raters than LLMs. These models can discern complex concepts, i.e., they are highly sensitive to features in the conventions and language complexity categories. After interventions on these concepts, the absolute value of the score difference predicted by BERT-like models often exceeds that of LLMs. This differs from the phenomenon observed by Singla et al. (2023), in which BERT-based models were found to function like *bag-of-words* models when scoring essays. However, BERT-like models struggle to dis-

Dataset	Setting	Conventions & Accuracy				Language Complexity		Organization & Development	
		Error Correction (+)	Error Introduction (-)		Complexification (+)	Simplification (-)	InParaShuffle (-)	InTextShuffle (-)	
			Spelling	SVA					WOS
	BERT	1.03 ^{+0.43}	-0.92 ^{+0.32}	-0.22 ^{+0.13}	-1.26 ^{+0.33}	0.42 ^{+0.35}	-0.69 ^{+0.33}	-0.01 ^{+0.06}	-0.01 ^{+0.06}
	RoBERTA	0.99 ^{+0.44}	-0.79 ^{+0.32}	-0.45 ^{+0.14}	-1.13 ^{+0.33}	0.24 ^{+0.33}	-0.35 ^{+0.33}	-0.19 ^{+0.06}	-0.02 ^{+0.05}
	DeBERTA	1.19 ^{+0.45}	-0.92 ^{+0.31}	-0.35 ^{+0.16}	-1.24 ^{+0.32}	0.33 ^{+0.32}	-0.27 ^{+0.26}	-0.06 ^{+0.05}	-0.06 ^{+0.05}
TOEFL11	GPT-3.5-ZSL	0.64 ^{+0.32}	-0.76 ^{+0.33}	-0.20 ^{+0.26}	-0.59 ^{+0.32}	0.27 ^{+0.25}	0.01 ^{+0.19}	-0.31 ^{+0.30}	-0.42 ^{+0.32}
	GPT-4-ZSL	0.92 ^{+0.31}	-0.80 ^{+0.34}	-0.35 ^{+0.21}	-0.80 ^{+0.30}	0.66 ^{+0.24}	-0.24 ^{+0.21}	-0.24 ^{+0.17}	-0.29 ^{+0.19}
	LLAMA-3-8B-ZSL	0.58 ^{+0.27}	-0.37 ^{+0.29}	-0.07 ^{+0.18}	-0.17 ^{+0.24}	0.57 ^{+0.26}	-0.11 ^{+0.23}	-0.15 ^{+0.24}	-0.23 ^{+0.26}
	LLAMA-3-70B-ZSL	0.64 ^{+0.26}	-0.56 ^{+0.25}	-0.24 ^{+0.21}	-0.41 ^{+0.23}	1.19 ^{+0.32}	-0.17 ^{+0.24}	-0.15 ^{+0.19}	-0.19 ^{+0.21}
	BERT	0.84 ^{+0.14}	-0.57 ^{+0.11}	-0.09 ^{+0.03}	-0.57 ^{+0.11}	0.31 ^{+0.09}	-0.11 ^{+0.08}	-0.01 ^{+0.02}	-0.02 ^{+0.02}
ELLIPSE	RoBERTA	0.92 ^{+0.14}	-0.50 ^{+0.09}	-0.11 ^{+0.03}	-0.54 ^{+0.09}	0.25 ^{+0.08}	-0.05 ^{+0.07}	-0.01 ^{+0.02}	-0.10 ^{+0.03}
	DeBERTA	1.06 ^{+0.16}	-0.64 ^{+0.13}	-0.20 ^{+0.06}	-0.64 ^{+0.13}	-0.08 ^{+0.07}	0.01 ^{+0.05}	-0.02 ^{+0.01}	-0.07 ^{+0.02}
	GPT-3.5-ZSL	0.77 ^{+0.19}	-0.60 ^{+0.19}	-0.19 ^{+0.15}	-0.35 ^{+0.18}	0.48 ^{+0.16}	0.08 ^{+0.14}	-0.15 ^{+0.15}	-0.18 ^{+0.16}
	GPT-3.5-FSL	0.35 ^{+0.14}	-0.46 ^{+0.15}	-0.15 ^{+0.12}	-0.31 ^{+0.14}	0.36 ^{+0.14}	-0.04 ^{+0.12}	-0.11 ^{+0.13}	-0.16 ^{+0.14}
	GPT-4-ZSL*	0.87 ^{+0.06}	-0.64 ^{+0.47}	-0.30 ^{+0.45}	-0.56 ^{+0.45}	0.96 ^{+0.65}	-0.05 ^{+0.58}	-0.10 ^{+0.33}	-0.19 ^{+0.37}
ELLIPSE	GPT-4-FSL*	0.61 ^{+0.52}	-0.71 ^{+0.60}	-0.27 ^{+0.50}	-0.56 ^{+0.48}	0.67 ^{+0.55}	-0.09 ^{+0.45}	-0.14 ^{+0.32}	-0.23 ^{+0.42}
	LLAMA-3-8B-ZSL	0.32 ^{+0.19}	-0.31 ^{+0.18}	-0.06 ^{+0.11}	-0.11 ^{+0.13}	0.70 ^{+0.13}	0.01 ^{+0.10}	-0.06 ^{+0.11}	-0.10 ^{+0.14}
	LLAMA-3-8B-FSL	0.06 ^{+0.11}	-0.11 ^{+0.16}	-0.02 ^{+0.08}	-0.06 ^{+0.11}	0.07 ^{+0.16}	-0.00 ^{+0.07}	-0.02 ^{+0.10}	-0.02 ^{+0.11}
	LLAMA-3-70B-ZSL*	0.51 ^{+0.18}	-0.41 ^{+0.11}	-0.11 ^{+0.09}	-0.19 ^{+0.10}	1.63 ^{+0.19}	0.03 ^{+0.18}	-0.03 ^{+0.07}	-0.06 ^{+0.08}
	LLAMA-3-70B-FSL*	0.51 ^{+0.17}	-0.54 ^{+0.05}	-0.12 ^{+0.33}	-0.24 ^{+0.50}	1.08 ^{+0.55}	-0.04 ^{+0.40}	-0.11 ^{+0.40}	-0.13 ^{+0.45}
ELLIPSE	GPT-3.5-SFT-50*	0.83 ^{+0.75}	-0.64 ^{+0.77}	-0.14 ^{+0.45}	-0.34 ^{+0.65}	0.96 ^{+0.60}	0.08 ^{+0.55}	-0.09 ^{+0.45}	-0.10 ^{+0.47}
	GPT-3.5-SFT-100*	1.12 ^{+0.81}	-0.95 ^{+0.80}	-0.26 ^{+0.52}	-0.58 ^{+0.57}	0.88 ^{+0.58}	0.05 ^{+0.48}	-0.18 ^{+0.50}	-0.19 ^{+0.60}
	GPT-3.5-SFT-200*	1.03 ^{+0.92}	-0.57 ^{+0.87}	-0.01 ^{+0.58}	-0.32 ^{+0.72}	0.79 ^{+0.52}	-0.02 ^{+0.37}	0.06 ^{+0.60}	0.02 ^{+0.62}
	GPT-3.5-SFT-400*	1.11 ^{+0.90}	-0.95 ^{+0.75}	-0.30 ^{+0.60}	-0.66 ^{+0.68}	0.76 ^{+0.55}	-0.03 ^{+0.45}	-0.18 ^{+0.52}	-0.23 ^{+0.60}
	GPT-3.5-SFT-800*	1.02 ^{+0.89}	-0.83 ^{+0.80}	-0.23 ^{+0.67}	-0.55 ^{+0.70}	0.94 ^{+0.55}	-0.03 ^{+0.50}	-0.14 ^{+0.52}	-0.23 ^{+0.62}
	BERT	0.84 ^{+0.14}	-0.57 ^{+0.11}	-0.09 ^{+0.03}	-0.57 ^{+0.11}	0.31 ^{+0.09}	-0.11 ^{+0.08}	-0.01 ^{+0.02}	-0.02 ^{+0.02}
	RoBERTA	0.92 ^{+0.14}	-0.50 ^{+0.09}	-0.11 ^{+0.03}	-0.54 ^{+0.09}	0.25 ^{+0.08}	-0.05 ^{+0.07}	-0.01 ^{+0.02}	-0.10 ^{+0.03}

Table 6: Mean score differences between original and counterfactual groups, with scores ranging from 1 to 5. Results are shown for both full and stratified subsets (stratified results are marked with *). Subscripts and superscripts indicate confidence intervals, obtained through 10,000 bootstrap iterations. Gray values indicate non-significant differences ($p > 0.01$), while coral values represent significant differences contrary to the expected intervention trend. (+) and (-) denote the expected direction of intervention effect.

tinguish interventions on organization and development. Their predicted score differences are often an order of magnitude smaller than those of LLMs (especially GPT-series models), indicating insensitivity to logical structures and coherence within essays.

Secondly, LLMs respond more comprehensively to our interventions than BERT-like models do, indicating a closer alignment with the criteria specified in scoring rubrics. As shown in Table 5, LLMs have lower agreement with human ratings when scoring essays in the ZSL setting. However, introducing FSL and SFT considerably improves their performance while maintaining the strength of their rationale alignment, as demonstrated in Table 6.

To better understand the impact of SFT on LLM performance, we fine-tuned GPT-3.5 Turbo with varying data sizes (50, 200, 400, and 800 samples), building on previous results with GPT-3.5-SFT-100. Figure 3 shows that scoring performance improves as the number of training essays increases. Additionally, with 400 essays for TOEFL11 and 200 for ELLIPSE, the performance nearly stabilizes, reaching levels comparable to BERT, while consistently maintaining rationale alignment capability.

4.4 Feedback Analysis

Han et al. (2023) and Xiao et al. (2024) have proposed that LLMs can provide helpful essay feed-

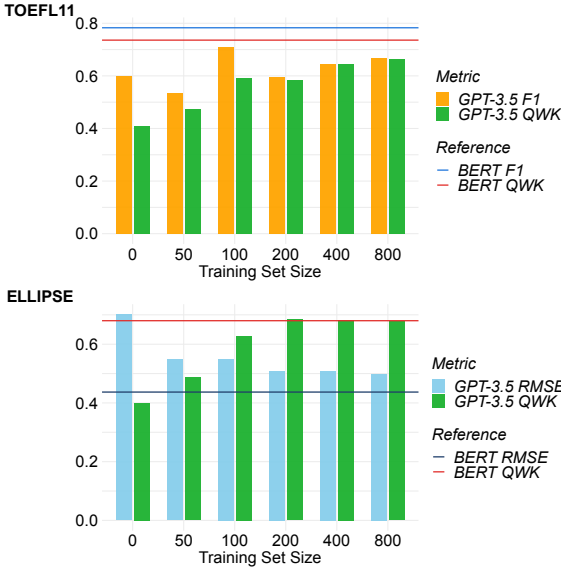


Figure 3: Scoring performance of GPT-3.5 SFT models with varying size of training data. The models’ performance improves as the number of training samples increases, reaching comparable or equivalent levels to BERT-like models.

back for both writers and evaluators. More importantly, but still largely overlooked, this feedback offers an opportunity to assess the construct validity of models. Therefore we further investigated the feedback differences provided by LLMs regarding the interventions. As Table 6 shows that GPT-4-FSL exhibits comprehensive sensitivity in all types of our interventions, while other models typically show inadequate sensitivity in one way or another,

we generated feedback using GPT-4 Turbo for further analysis of the model’s faithfulness.

As shown in Figure 4, after few-shot prompting on scoring task, we continued to ask GPT-4 Turbo to generate feedback based on the scoring rubrics, explaining the scores in terms of aforementioned three concepts. In this way, we obtained feedback for each essay and its counterfactual counterparts³. Then, three trained annotators were hired to evaluate the feedback differences within each feedback pair, determining whether counterfactual interventions can be detected without accessing essay content. See the detailed evaluation procedures in Appendix D.

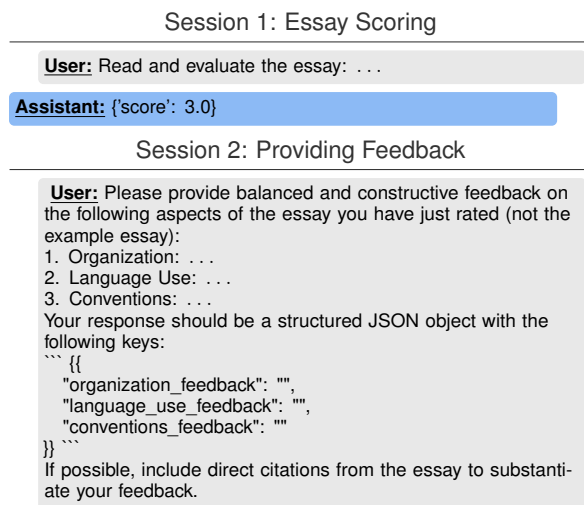


Figure 4: An Example of Feedback Generation

Table 7 presents the annotator-voted results, demonstrating that a large proportion of counterfactual interventions can be identified simply from the feedback given by the GPT-4 Turbo, especially for complexification, error correction and error introduction except for SVA. On the other hand, simplification and organization interventions are hard to be detected simply from feedback pairs, which is consistent with their relatively smaller absolute effect as shown in Table 6. One possible reason is that the ELLIPSE essays, written by 8th to 12th grade English learners, tend to be simple in vocabulary and syntax, contain some spelling and SVA errors, and exhibit imperfect logic flow and coherence. Consequently, the model frequently identified SVA issues and offered numerous organizational and develop-

³We conducted stratified sampling on the ELLIPSE dataset to obtain 200 essay samples and, through two rounds of dialogues, acquired 200 "original-counterfactual" feedback pairs for human evaluation. For the evaluation process, we categorized these pairs based on eight counterfactual interventions and assessed each category of cases accordingly.

Category	Counterfactual Type	Detection Rate%
Conventions	Error Correction	72
	Spelling	68
	SVA	48
	WOS	80
Language	Complexification	100
Complexity	Simplification	32
Organization	InParaShuffle	40
	InTextShuffle	20

Table 7: Voting-Based Detection Rates of Original vs. Counterfactual Feedback.

mental suggestions both in feedback of original and counterfactual essays, leading to less distinct differences.

5 Conclusion

We generated linguistically-informed counterfactuals with an integrated approach combining LLM and rule-based methods, analyzing their impact on essay scoring results of BERT-like models and LLMs. Our findings emphasize that a higher scoring agreement with human raters does not necessarily indicate a better alignment with scoring rubrics, suggesting that a more holistic evaluation approach should consider both aspects. Moreover, our study highlights LLMs’ considerable potential in AES domain: while zero-shot prompted LLMs show lower scoring agreement compared to BERT-like models, a major reason for this is that they tend to be conservative or strict when evaluating the essay. FSL and SFT could significantly increase the agreement level with annotated essays serve as anchors to neutralize the conservatism. In the mean time, LLMs demonstrate comprehensive rationale alignment with scoring rubrics. This ability is stably maintained in ZSL, FSL and SFT settings. Lastly, LLMs are not only sensitive to counterfactual interventions when scoring but can also reflect a large part of these differences in their feedback, an advantage beyond the reach of traditional AES methods.

This study sheds light on *why* a neural model assigns specific scores to essays. It unveils how modifying domain-specific concepts in texts to craft counterfactuals enhances transparency in model decisions—a method applicable across multiple fields. With LLMs, counterfactual generation has been greatly empowered, boosting transparency and accountability in machine learning applications.

6 Limitations

In addition to conventions, language complexity, and organization, TOEFL independent writing rubrics also emphasize content-related evaluations—namely, assessing relevance to the prompt and fulfillment of task requirements. These aspects, being beyond mere linguistic concepts, were not included in the current scope of our study. This is because counterfactual interventions require modifying a specific aspect while keeping others constant. This is also because we can adjust linguistic features without affecting content, but altering content inevitably impacts the linguistic aspect. However, we acknowledge that task and topic relevance, as important scoring dimensions, warrant future in-depth exploration.

Our experiment demonstrated that LLMs have significant potential in providing feedback. In this paper, we focus on the feedback differences between original and counterfactual samples. A comprehensive evaluation of the LLM-generated feedback is a crucial step for future research.

Acknowledgements

This research was supported by a grant from the Center for Language Education and Cooperation at the Ministry of Education of China (No. 22YH04ZW), a grant from the National Language Commission of China (No. ZDA145-9), and a grant from the Beijing Federation of Social Science Circles (No. 21DTR037). It is also supported by the Fundamental Research Funds for the Central Universities of China.

References

- Isaac I Bejar, Michael Flor, Yoko Futagi, and Chaintanya Ramineni. 2014. On the vulnerability of automated scoring to construct-irrelevant response strategies (cirs): An illustration. *Assessing Writing*, 22:48–59.
- Daniel Blanchard, Joel Tetreault, Derrick Higgins, Aoife Cahill, and Martin Chodorow. 2013. Toefl11: A corpus of non-native english. *ETS Research Report Series*, 2013(2):i–15.
- Andrew Caines, Luca Benedetto, Shiva Taslimipoor, Christopher Davis, Yuan Gao, Oeistein Andersen, Zheng Yuan, Mark Elliott, Russell Moore, Christopher Bryant, Marek Rei, Helen Yannakoudakis, Andrew Mullooly, Diane Nicholls, and Paula Buttery. 2023. [On the application of large language models for language teaching and assessment technology](#).
- Winston Carlile, Nishant Gurrupadi, Zixuan Ke, and Vincent Ng. 2018. [Give me more feedback: Annotating argument persuasiveness and related attributes in student essays](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 621–631, Melbourne, Australia. Association for Computational Linguistics.
- Zeming Chen, Qiyue Gao, Antoine Bosselut, Ashish Sabharwal, and Kyle Richardson. 2023. [Disco: distilling counterfactuals with large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5514–5528.
- Jacob Cohen. 2013. *Statistical power analysis for the behavioral sciences*. Academic press.
- Scott Crossley, Yu Tian, Perpetual Baffour, Alex Franklin, Youngmeen Kim, Wesley Morris, Meg Benner, Aigner Picou, and Ulrich Boser. 2023. [The english language learner insight, proficiency and skills evaluation \(ellipse\) corpus](#). *International Journal of Learner Corpus Research*, 9(2):248–269.
- Mark Davies. 2010. *The corpus of historical American English: COHA*. BYE, Brigham Young University.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Tanay Dixit, Bhargavi Paranjape, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Core: A retrieve-then-edit framework for counterfactual data generation](#). *arXiv preprint arXiv:2210.04873*.
- Amir Feder, Katherine A. Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E. Roberts, Brandon M. Stewart, Victor Veitch, and Diyi Yang. 2022. [Causal inference in natural language processing: Estimation, prediction, interpretation and beyond](#). *Transactions of the Association for Computational Linguistics*, 10:1138–1158.
- James Fiacco, David Adamson, and Carolyn Ros. 2023. [Towards extracting and understanding the implicit rubrics of transformer based automatic essay scoring models](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 232–241, Toronto, Canada. Association for Computational Linguistics.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, et al. 2020. [Evaluating models’ local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323.

- Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H Chi, and Alex Beutel. 2019. Counterfactual fairness in text classification through robustness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 219–226.
- Yair Gat, Nitay Calderon, Amir Feder, Alexander Charanin, Amit Sharma, and Roi Reichart. 2023. [Faithful explanations of black-box nlp models using llm-generated counterfactuals](#).
- Jieun Han, Haneul Yoo, Junho Myung, Minsun Kim, Hyunseung Lim, Yoonsu Kim, Tak Yeon Lee, Hwajung Hong, Juho Kim, So-Yeon Ahn, and Alice Oh. 2023. [Fabric: Automated scoring and feedback generation for essays](#).
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#).
- Mohamed A. Hussein, Hesham A., and Mohammad Nassef. 2020. [A trait-based deep learning automated essay scoring system with adaptive feedback](#). *International Journal of Advanced Computer Science and Applications*, 11(5).
- Anubha Kabra, Mehar Bhatia, Yaman Kumar Singla, Junyi Jessy Li, and Rajiv Ratn Shah. 2022. [Evaluation toolkit for robustness testing of automatic essay scoring systems](#). In *Proceedings of the 5th Joint International Conference on Data Science & Management of Data (9th ACM IKDD CODS and 27th COMAD), CODS-COMAD '22*, pages 90–99, New York, NY, USA. Association for Computing Machinery.
- Zixuan Ke and Vincent Ng. 2019. Automated essay scoring: A survey of the state of the art. In *IJCAI*, volume 19, pages 6300–6308.
- Alwyn Vwen Yen Lee, Andrés Carlos Luco, and Seng Chee Tan. 2023. [A human-centric automated essay scoring and feedback system for the development of ethical reasoning](#). *Educational Technology & Society*, 26(1):147–159.
- Yongqi Li, Mayi Xu, Xin Miao, Shen Zhou, and Tiejun Qian. 2024. Prompting large language models for counterfactual generation: An empirical study. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13201–13221.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. 2024. [Towards faithful model explanation in NLP: A survey](#). *Computational Linguistics*, 50(2):657–723.
- Edward Ma. 2019. [NLP augmentation](#).
- R Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2020. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *57th Annual Meeting of the Association for Computational Linguistics, ACL 2019*, pages 3428–3448. Association for Computational Linguistics (ACL).
- Atsushi Mizumoto and Masaki Eguchi. 2023. [Exploring the potential of using an ai language model for automated essay scoring](#). *Research Methods in Applied Linguistics*, 2(2):100050.
- Matthew C Myers and Joshua Wilson. 2023. Evaluating the construct validity of an automated writing evaluation system with a randomization algorithm. *International Journal of Artificial Intelligence in Education*, 33(3):609–634.
- Daniel Naber et al. 2003. *A rule-based style and grammar checker*. GRIN Verlag Munich, Germany.
- Ben Naismith, Phoebe Mulcaire, and Jill Burstein. 2023. [Automated evaluation of written discourse coherence using gpt-4](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 394–403, Toronto, Canada. Association for Computational Linguistics.
- Christopher M. Ormerod, Akanksha Malhotra, and Amir Jafari. 2021. [Automated essay scoring using efficient transformer-based language models](#).
- Donald E Powers, Jill C Burstein, Martin Chodorow, Mary E Fowles, and Karen Kukich. 2002. Stumping e-rater: challenging the validity of automated essay scoring. *Computers in Human Behavior*, 18(2):103–134.
- Dadi Ramesh and Suresh Kumar Sanampudi. 2022. [An automated essay scoring systems: A systematic literature review](#). *Artificial Intelligence Review*, 55(3):2495–2527.
- Pedro Uria Rodriguez, Amir Jafari, and Christopher M. Ormerod. 2019. [Language models and automated essay scoring](#).
- Alexis Ross, Tongshuang Wu, Hao Peng, Matthew Peters, and Matt Gardner. 2022. [Tailor: Generating and perturbing text with semantic controls](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3194–3213, Dublin, Ireland. Association for Computational Linguistics.
- Yaman Kumar Singla, Swapnil Parekh, Somesh Singh, Junyi Jessy Li, Rajiv Ratn Shah, and Changyou Chen. 2023. Automatic essay scoring systems are both overstable and oversensitive: explaining why and proposing defenses. *Dialogue & Discourse*, 14(1):1–33.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. [Axiomatic attribution for deep networks](#).

- Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. 2023. [Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 74952–74965. Curran Associates, Inc.
- Yongjie Wang, Chuan Wang, Ruobing Li, and Hui Lin. 2022. [On the use of bert for automated essay scoring: Joint learning of multi-scale essay representation](#).
- Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel Weld. 2021. [Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6707–6723, Online. Association for Computational Linguistics.
- Changrong Xiao, Wenxing Ma, Qingping Song, Sean Xin Xu, Kunpeng Zhang, Yufang Wang, and Qi Fu. 2024. [Human-ai collaborative essay scoring: A dual-process framework with llms](#).
- Ruosong Yang, Jiannong Cao, Zhiyuan Wen, Youzheng Wu, and Xiaodong He. 2020. [Enhancing automated essay scoring performance via fine-tuning pre-trained language models with combination of regression and ranking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1560–1569, Online. Association for Computational Linguistics.

Appendix

A Rubrics

To identify the core concepts for intervention, we reviewed five scoring rubrics from IELTS Writing, TOEFL iBT Independent Writing, TOEIC Writing, PTE Academic Writing and the ELLIPSE dataset. We aimed to uncover commonalities across the five rubrics and found that they could be categorized into three dimensions: (1) conventions and accuracy; (2) language complexity; and (3) organization and development. For clarity, this section will present the descriptors for the highest score in each rubric, with **color-coded highlights** to indicate the corresponding dimensions.

A.1 IELTS Writing

- Task achievement: fully satisfies all the requirements of the task; clearly presents a **fully developed response**.
- Coherence and cohesion: uses **cohesion** in such a way that it attracts no attention; skillfully **manages paragraphing**.
- Lexical resource: uses **a wide range of vocabulary with very natural and sophisticated control of lexical features**; **rare minor errors occur only as 'slips'**.
- Grammatical range and accuracy: uses **a wide range of structures with full flexibility and accuracy**; **rare minor errors occur only as 'slips'**.

A.2 TOEFL Independent Writing

- Effectively addresses the topic and task.
- Is **well organized and well developed**, using clearly appropriate explanations, exemplifications and/or details.
- Displays **unity, progression and coherence**.
- Displays consistent facility in the use of language, demonstrating **syntactic variety, appropriate word choice and idiomaticity**, though it may have **minor lexical or grammatical errors**.

A.3 TOEIC Writing

- Typically, test takers at level 9 can communicate straightforward information effectively and use reasons, examples, or explanations to support an opinion.

- When using reasons, examples, or explanations to support an opinion, their writing is **well-organized and well-developed**.
- The use of English is natural, with **a variety of sentence structures and appropriate word choice**, and is **grammatically accurate**.
- When giving straightforward information, asking questions, giving instructions, or making requests, their writing is clear, **coherent**, and effective.

A.4 PTE Academic Writing

- Content: Adequately deals with the prompt.
- Form: Length is between 200 and 300 words.
- Development, Structure & Coherence: Shows **good development and logical structure**.
- Grammar: Shows consistent **grammatical control of complex language**. **Errors are rare and difficult to spot**.
- General Linguistic Range: Exhibits mastery of **a wide range of language** to formulate thoughts precisely, give emphasis, differentiate and eliminate ambiguity. No sign that the test taker is restricted in what they want to communicate.
- Vocabulary: Good command of **a broad lexical repertoire, idiomatic expressions and colloquialisms**.
- Spelling: **Correct spelling**.

A.5 ELLIPSE Dataset

- Overall: Native-like facility in the use of language with **syntactic variety, Appropriate word choice and phrases**; **well-controlled text organization**; **precise use of grammar and conventions**; **rare language inaccuracies** that do not impede communication.
- Cohesion: **Text organization consistently well controlled using a variety of effective linguistic features such as reference and transitional words and phrases to connect ideas across sentences and paragraphs**; **appropriate overlap of ideas**.

- Syntax: Flexible and effective use of a **full range of syntactic structures including simple, compound, and complex sentences**; There may be **rare minor and negligible errors** in sentence formation.
- Vocabulary: **Wide range of vocabulary** flexibly and effectively used to convey precise meanings; skillful use of **topic-related terms and less common words**; **rare negligible inaccuracies in word use**.
- Phraseology: Flexible and effective use of a **variety of phrases, such as idioms, collocations, and lexical bundles**, to convey precise and subtle meanings; **rare minor inaccuracies that are negligible**.
- Grammar: Command of grammar and **usage with few or no errors**.
- Conventions: **Consistent use of appropriate conventions to convey meaning**; **spelling, capitalization, and punctuation errors nonexistent or negligible**.

B Detail of Counterfactual Generation

In this section, we present the details of our counterfactual generation experiment. This includes examples of rule-based counterfactuals, information on the models used, the prompts provided to the LLMs, and a comparative analysis of various aspects of interest in the counterfactuals generated by GPT-4 Turbo and Llama-3-70b-Instruct.

B.1 Examples of Rule-based Counterfactuals

In this study, all the interventions designed to introduce errors into essays and decrease organization are rule-based. In this subsection, we provide counterfactual examples for each of the rule-based interventions. See Table 8.

B.2 Prompts for Counterfactual Generation

The counterfactual samples of text correction, complexification and simplification are generated by the gpt-4-1106-preview model. When calling OpenAI’s APIs, we turn on JSON mode to get easier parsing results. For reproducibility, we set the temperature parameter to 0 and the seed to 42.

B.2.1 Prompt for Error Correction

System: You are an experienced writing tutor.

User: Please fix the spelling, punctuation and grammatical errors in the given essay. Ensure the main idea, the words used, the sentence structure, and the length of the text remain consistent with the original text.

Input Essay:
"{"

Please return the output essay in JSON format as below:

```
```
{"output_essay": "..."}
```
```

Output:

B.2.2 Prompt for Complexification

System: You are an experienced writing tutor.

User: Modify the provided essay to enhance its lexical sophistication and syntactic variety following the instructions below:

1. Expand lexical range: Vary word choice and replace common words with advanced vocabulary when suitable without compromising clarity or meaning. Avoid repeating the same words and capture subtle differences in meaning.
2. Increase syntactic complexity: Incorporate a wider range of sentence structures including compound-complex sentences, varied clause types, subordination and coordination. Use advanced constructions such as non-finite clauses, adverbials, conditionals, inversion and passives where appropriate.
3. Maintain meaning, length and clarity: The revised text should retain the original ideas and conform to the initial length while remaining clear and understandable.

Input Essay:
"{"

Please return the output essay in JSON format as below:

Original: In the up coming years i wanna save up and work hard to get money, so that i can open my own personal store.. \n\n in the store i going to be selling high branded shoes,clothes, and outer collectives. and Im going to learning how to invest my money as i sell things before opening up my dream store. \n\n i wanna do this as a job because ive awlays had a passhine for shoes and high branded things. and if i can pull this off i think ill be veary succes full in life i will be getting and veary good salary ill be able to pay off anything and have to money to support my family.

Spelling Error Introduction: In the up coming years i wanna save up and work hard to get money, so that i can open my own personal store. \n\n ia the store ye goind to be salling high branded shoes, couthes, anso outer collectives. and Im going to learning how to invest my money as i sell things before opening up my dream store. \n\n i wanna do these as a job because live awlays had e passhine foi shooes and highest branded things. and if i can pull this off i think ill be veary succes full in life i will be getting and veary good salary ill be able to pay off anything and have to money to support my family.

Subject-verb Agreement Error Introduction: In the up coming years i wanna save up and work hard to get money, so that i can open my own personal store . \n\n in the store i going to be selling high branded shoes,clothes, and outer collectives. and Im going to learning how to invest my money as i sell things before opening up my dream store. \n\n i wanna does this as a job because ives awlays had a passhine for shoes and high branded things. and if i can pull this off i thinks ill is veary succes full in life i will be getting and veary good salary ill be able to pay off anything and have to money to support my family.

Word Order Swapping: In the up coming years i wanna save up and work hard to get money, so that i can open my own personal store . \n\n the in store i be going to selling branded high, shoes clothes, and collectives outer. and Im going to learning how to invest my money as i sell things before opening up my dream store. \n\n i wanna do this as a job because ive awlays had a passhine for shoes and high branded things. and if i can pull off this i ill think be veary succes full in i life be will getting and veary salary good ill be able to pay off anything and to have money to support my family.

Intra-paragraph Shuffling: In the up coming years i wanna save up and work hard to get money, so that i can open my own personal store . \n\n and Im going to learning how to invest my money as i sell things before opening up my dream store. in the store i going to be selling high branded shoes,clothes, and outer collectives. \n\n and if i can pull this off i think ill be veary succes full in life i will be getting and veary good salary ill be able to pay off anything and have to money to support my family. i wanna do this as a job because ive awlays had a passhine for shoes and high branded things.

Inter-text Shuffling: In the up coming years i wanna save up and work hard to get money, so that i can open my own personal store . \n\n i wanna do this as a job because ive awlays had a passhine for shoes and high branded things.and Im going to learning how to invest my money as i sell things before opening up my dream store. \n\n in the store i going to be selling high branded shoes,clothes, and outer collectives.and if i can pull this off i think ill be veary succes full in life i will be getting and veary good salary ill be able to pay off anything and have to money to support my family.

Table 8: Example of the medium-level TOEFL11 essay used in Table 3 and its rule-based counterfactual counterparts.

```

...
{"output_essay": "..."}
...

Output:

```

```

Please return the output essay in JSON for-
mat as below:
...
{"output_essay": "..."}
...

Output:

```

B.2.3 Prompt for Simplification

System: You are an experienced writing tutor.

User: Modify the provided essay to simplify its vocabulary and sentence structure following the instructions below:

1. Simplify vocabulary: Replace advanced words with common everyday equivalents for clear understanding. Limit synonyms to favor those most commonly used.
2. Simplify sentence structure: Break down complex sentences and avoid clauses, conjunctions, and nesting where possible. Favor short, simple subject-verb-object sentences.
3. Maintain meaning, length and clarity: The revised text should retain the original ideas and conform to the initial length while remaining clear and understandable.

Input Essay:
{"}

B.3 Comparative Performance of GPT-4 Turbo and Llama-3-70b-Instruct in Counterfactual Generation

Table 9 shows the effect size of three types of interventions performed by both GPT-4 Turbo and Llama-3-70b-Instruct on seven linguistic metrics across two datasets. It can be seen that the impact of the two models on the original essay, across various language metrics of interest during counterfactual interventions, aligns with expectations, albeit with slight variations in degree. In terms of error correction, GPT-4 significantly reduces error density. Meanwhile, for complexification and simplification, GPT-4 intervenes more comprehensively in vocabulary and syntax, with generally smaller changes in length.

Table 4 presents the embedding similarities between counterfactuals and original essays given by both LLMs. Although Llama-3-70b-Instruct retains a higher degree of the original text’s meaning than GPT-4 Turbo in most cases, it shows a signif-

icant drop when simplifying the ELLIPSE essay, indicating its potential lack of stability.

C The Implementation of AES methods

C.1 Fine-tuning BERT-like Models

We fine-tuned three commonly used pre-trained transformer-based encoder models, specifically bert-base-uncased, roberta-base, and deberta-v3-base.

C.1.1 Basic Settings

As the essays in the TOEFL11 dataset are categorized into low, medium, and high categories, we developed a three-class classifier using the cross-entropy loss. We use the `AutoModelForSequenceClassification` class from Hugging Face transformer, setting `num_labels=3` to load the pre-training checkpoints. For the ELLIPSE dataset, with scores ranging from 1.0 to 5.0, we model it as a regression problem by setting `num_labels=1` and using the mean squared error (MSE) loss function.

C.1.2 Hyperparameters

In our model fine-tuning process, we experimented with four distinct learning rates: $1e-5$, $2e-5$, $3e-5$, and $5e-5$, using Hugging Face's Trainer. We identify the best learning rate that led to the lowest loss on the validation set (results see Table 11). We used a linear learning rate scheduler that includes a 50-step warm-up phase, where the learning rate initially increases from a lower value to a specified maximum (chosen from the four rates: $1e-5$, $2e-5$, $3e-5$, and $5e-5$) and then decreases linearly. This method ensures gradual adaptation of the model's weights, with the peak learning rates being reached at the end of the warm-up.

For other parameters, we used a seed of 42 and a batch size of 16 for both training and evaluation. We aimed for a maximum of 10 epochs, with the actual duration potentially reduced by early stopping, triggered if loss value fails to improve after 5 checks. The approach included a weight decay of 0.01 for overfitting prevention and FP16 for efficient training. Input lengths were adjusted to 512 tokens through padding and truncation to ensure uniformity across all samples.

C.2 Prompting LLMs to Score Essays

As introduced in Section 4, we also used LLMs for essay scoring, including gpt-3.5-turbo-1106 and gpt-4-1106-preview based on OpenAI's

API. We turned on JSON mode to get easier parsing results, and set the temperature parameter to 0 and the seed parameter to 42 for reproducibility.

C.2.1 Prompts for Scoring TOEFL11 Essays with Zero-shot Learning

Below is the scoring template for TOEFL11 essays. In the zero-shot setting, we provide the LLMs with the essay prompt, the essay itself, and the scoring rubrics. Notably, while the TOEFL11 dataset only provides low, medium, and high score levels for the essays without specific scores, the TOEFL rating rubric is actually based on a 1 to 5 scale. Consequently, even in zero-shot scenarios without examples or training data, we can still prompt LLMs to assess and score TOEFL11 essays.

System: You are a TOEFL rater specializing in the evaluation of the Independent Writing section.

User: Read and evaluate the essay written in response to the prompt: "{}"

Essay: "{}"

Please assign it a score from 1 to 5 (in increments of 0.5 points) based on rubric below: "{TOEFL11_RUBRICS}"

Your response should be a JSON object containing only one key: 'score', which should be a numeric value representing the score you gave.

TOEFL11 Rubrics

- A 5-point essay effectively addresses all aspects of the topic and task. Well organized and developed with clearly appropriate explanations and details. Displays strong unity, progression and coherence. Shows consistent language facility with syntactic variety, appropriate word choice and idiomaticity. May have minor lexical or grammatical errors.
- A 4-point essay addresses the topic and task well, though some points may not be fully elaborated. Generally well organized and developed with appropriate and sufficient explanations, exemplifications and details. Displays unity, progression and coher-

Metrics	Model	TOEFL11			ELLIPSE		
		Error Correction	Complexification	Simplification	Error Correction	Complexification	Simplification
WordNum	LLAMA-3-70B-IT	-0.045	-0.170	-1.325	0.065	-0.332	-0.981
	GPT-4 TURBO	-0.098	0.078	-1.060	-0.027	-0.103	-0.714
SentNum	LLAMA-3-70B-IT	0.037	-0.508	-0.037	0.215	-0.406	-0.074
	GPT-4 TURBO	0.047	-0.323	0.454	0.280	-0.264	0.367
MLS	LLAMA-3-70B-IT	-0.176	0.385	-1.473	-0.354	0.156	-1.816
	GPT-4 TURBO	-0.245	0.449	-1.714	-0.481	0.423	-2.237
ADDT	LLAMA-3-70B-IT	-0.030	0.734	-1.359	-0.353	0.817	-1.628
	GPT-4 TURBO	-0.066	0.982	-1.535	-0.481	1.220	-1.875
LemmaTTR	LLAMA-3-70B-IT	-0.074	2.130	-0.467	0.020	2.647	-0.009
	GPT-4 TURBO	0.094	2.985	-0.611	0.429	3.323	-0.128
LexSoph	LLAMA-3-70B-IT	-1.538	3.596	-0.186	-0.799	3.710	0.301
	GPT-4 TURBO	-1.514	5.277	-0.909	-0.711	5.063	-0.291
ErrorDensity	LLAMA-3-70B-IT	-5.015	-0.616	-0.535	-1.887	-0.628	-0.331
	GPT-4 TURBO	-5.122	-0.407	-0.219	-1.869	-0.412	-0.123

Table 9: Cohen’s D for seven linguistic metrics on three interventions of GPT-4 Turbo and Llama-3-70b-Instruct.

Model	Intervention	TOEFL11	ELLIPSE
GPT-4 TURBO	Error Correction	0.935	0.942
	Complexification	0.760	0.749
	Simplification	0.816	0.849
LLAMA-3-70B-IT	Error Correction	0.944	0.957
	Complexification	0.817	0.813
	Simplification	0.853	0.610

Table 10: Mean cosine similarity between original and counterfactual essays for GPT-4 and Llama-3-70b-instruct given by text-embedding-3-large.

ence, though may contain occasional redundancy, digression or unclear connections. Demonstrates syntactic variety and vocabulary range. May have occasional minor errors that do not interfere with meaning.

- A 3-point essay addresses the topic and task with somewhat developed explanations, exemplifications and details. Displays unity, progression and coherence, though connection of ideas may be occasionally obscured. May demonstrate inconsistent language facility resulting in lack of clarity and obscured meaning. May display accurate but limited structures and vocabulary.
- A 2-point essay shows limited development in response to the topic and task. Inadequate organization or connection of ideas. Insufficient or inappropriate exemplifications, explanations or details to support generalizations. Noticeable inappropriate word choices or word forms. An accumulation of errors in sentence structure and/or usage.

Dataset	Model	Learning Rate	EarlyStop@Step	Validation Loss ↓
TOEFL11	BERT	1e-5	450	.443
		2e-5	550	.453
		3e-5	350	.462
		5e-5	150	.482
	ROBERTA	1e-5	450	.403
		2e-5	450	.424
		3e-5	400	.442
		5e-5	500	.467
	DEBERTA	1e-5	500	.398
		2e-5	400	.400
		3e-5	250	.416
		5e-5	250	.427
ELLIPSE	BERT	1e-5	500	.173
		2e-5	200	.172
		3e-5	300	.179
		5e-5	150	.185
	ROBERTA	1e-5	250	.196
		2e-5	100	.199
		3e-5	500	.171
		5e-5	300	.176
DEBERTA	1e-5	200	.157	
	2e-5	150	.167	
	3e-5	200	.160	
	5e-5	150	.181	

Table 11: Performance of the three models on the validation set after fine-tuning using different learning rates on both TOEFL11 and ELLIPSE datasets. Learning rates for achieving minimum loss in each model for both datasets are **bolded**.

- A 1-point essay is seriously flawed due to disorganization, underdevelopment, little or no supporting detail, and unresponsiveness to the task. Contains serious and frequent errors in sentence structure or usage.

C.2.2 Prompts for Scoring ELLIPSE Essays with Zero-shot Learning

Below is the scoring template for ELLIPSE essays. Since the ELLIPSE’s rubrics do not require adherence to a specific prompt or fulfillment of task requirements. We only provide the LLMs with the essay to be rated and the scoring rubrics.

System: You are an essay rater specializing in the evaluation of essays written by students from 8th to 12th grade who are learning English as a second language.

User: Read and evaluate the essay: "{}"

Assign it a score from 1 to 5, in increments of 0.5, based on this rubric:
"{ELLIPSE_RUBRICS}"

Your response should be a JSON object containing only one key: 'score', which should be a numeric value representing the score you gave.

ELLIPSE Rubrics

- A 5-point essay demonstrates native-like facility in the use of language with syntactic variety, appropriate word choice and phrases; well-controlled text organization; precise use of grammar and conventions; rare language inaccuracies that do not impede communication.
- A 4-point essay demonstrates facility in the use of language with syntactic variety and range of words and phrases; controlled organization; accuracy in grammar and conventions; occasional language inaccuracies that rarely impede communication.
- A 3-point essay demonstrates facility limited to the use of common structures and generic vocabulary; organization generally controlled although connection sometimes absent or unsuccessful; errors in grammar and syntax and usage. Communication is impeded by language inaccuracies in some cases.
- A 2-point essay demonstrates inconsistent facility in sentence formation, word choice, and mechanics; organization partially developed but may be missing or unsuccessful. Communication impeded in many instances by language inaccuracies.
- A 1-point essay demonstrates a limited range of familiar words or phrases loosely strung together; frequent errors in grammar (including syntax) and usage. Communication impeded in most cases by language inaccuracies.

C.2.3 Prompts for Scoring TOEFL11 Essays with Few-shot Learning

For few-shot learning on TOEFL11 dataset, we gave three examples from the low, medium and high categories, and asked the models to return the score level as well. See the prompt below.

System: You are a TOEFL rater specializing in the evaluation of the Independent Writing section.

User: Read and evaluate the essay written in response to the prompt: "{}"

Example essay 1 of score level "High":
"{A_REPRESENTATIVE_HIGH_SCORE_ESSAY}"

Example Essay 2 of score level "Medium":
"{A_REPRESENTATIVE_MEDIUM_SCORE_ESSAY}"

Example Essay 3 of score level "Low":
"{A_REPRESENTATIVE_LOW_SCORE_ESSAY}"

Essay to score:
"{}"

Please note:

- Low corresponds to scores of 1.0 - 2.0
- Medium corresponds to scores of 2.5 - 3.5
- High corresponds to scores of 4.0 - 5.0

Assign the essay a score level of Low, Medium or High based on the criteria in the rubric below:
"{TOEFL11_RUBRICS}"

Your response should be a JSON object with the key "score_level" set to either "Low", "Medium", or "High" representing the level you determined for this essay.

C.2.4 Prompts for Scoring ELLIPSE Essays with Few-shot Learning

To align with the process of rating TOEFL11 essays, we also provide three example essays from the ELLIPSE dataset, representing low, medium, and high score levels. However, we give the specific scores of these examples and require the model to return numerical scores as well. Refer to the following prompt. For information on how to select samples, see the next section.

System: You are an essay rater specializing in the evaluation of essays written by students from 8th to 12th grade who are learning English as a second language.

User: Read and evaluate the essay:

Example essay 1 of score "4.0":
"{A_REPRESENTATIVE_HIGH_SCORE_ESSAY}"

Example Essay 2 of score "3.0":
"{A_REPRESENTATIVE_MEDIUM_SCORE_ESSAY}"

Example Essay 3 of score "2.0":
"{A_REPRESENTATIVE_LOW_SCORE_ESSAY}"

Essay to score:
"{}"

Assign it a score from 1 to 5, in increments of 0.5, based on this rubric:
"{ELLIPSE_RUBRICS}"

Your response should be a JSON object containing only one key: 'score', which should be a numeric value representing the score you gave.

C.2.5 Few-shot Example Selection

We use a linguistic-based approach to select the representative examples for few-shot learning by following the steps:

1. **Calculate Metrics:** Calculate and normalize

the seven linguistic metrics mentioned in Section 3.4 for training sets of both TOEFL11 and ELLIPSE datasets.

2. **Process Data:** Apply Principal Component Analysis (PCA) to identify the top five components that explain 95% of the variance, representing essential linguistic metrics.
3. **Represent Samples:** Utilize these principal components to represent the linguistic metrics of all training samples.
4. **Determine Medoids:** Categorize samples into proficiency levels (low, medium, high) and find the medoid of each group using Euclidean distance.

Note that a medoid is an object within a dataset that minimally differs from all other objects in the dataset, according to a given distance metric. It is similar to the concept of a centroid, but while a centroid may not be an actual data point, a medoid is always a member of the dataset.

C.3 Fine-tuning GPT-3.5 Turbo

We fine-tuned GPT-3.5 Turbo model using the OpenAI API⁴ with the following hyperparameters: 3 epochs, a batch size of 1, and a learning rate multiplier of 2. These are the default settings provided by OpenAI, as the size and weight of GPT-3.5 Turbo model are not accessible, a systematic parameter search would be very costly and even impossible.

D Details for Feedback Generation and Evaluation

D.1 Feedback Generation

Given the stable performance of few-shot GPT-4 in handling a variety of counterfactual interventions, we conducted the manual evaluations on this model. As shown in Figure 4, we prompted the few-shot GPT-4 to provide writing feedback to the essay it just scored. The experiments were conducted on a stratified subset of ELLIPSE. For 200 samples in the subset, we requested GPT-4 to provide feedback respectively on each of the original samples and their specific counterfactual counterparts.

⁴<https://platform.openai.com/docs/guides/fine-tuning>

Full Prompt Instructing GPT-4 to Provide Feedback

User: Please provide balanced and constructive feedback on the following aspects of the essay you have just rated (not the example essay):

1. Organization:

- Evaluate how effectively ideas are communicated and organized. Identify any issues with the logical flow, transitions between ideas, and clarity in conveying concepts. Comment on the introduction's setup, idea development throughout the body, and the conclusiveness of the ending.

2. Language Use:

- Morphology: Identify errors in word formation and structure, focusing on verb tenses, irregular verbs, plurals, possessives, affixes, agreement, and gerund/participle usage.

- Syntax: Comment on the arrangement of words and phrases to create well-formed sentences, coherence in sentence construction, and the complexity and variety of sentence types.

- Vocabulary: Assess the appropriateness of word choice, the diversity and sophistication of vocabulary employed, and note any imprecise use of words where more accurate or specific terms could be used.

3. Conventions:

- Highlight any errors in spelling, capitalization, and punctuation.

Your response should be a structured JSON object with the following keys:

```
```json
{{
 "organization_feedback": "",
 "language_use_feedback": "",
 "conventions_feedback": ""
}}
```

If possible, include direct citations from the essay to substantiate your feedback.

## D.2 Feedback Evaluation

After collecting 200 "original-counterfactual" feedback pairs, we divided them into 8 equal portions, with each portion corresponding to one counterfactual intervention for manual evaluation. Three grad-

uate students majoring in Linguistics were tasked with identifying feedback differences within each pair by deciding whether one of the feedback better aligned with specific features. For instance, for interventions introducing *spelling errors*, we asked,

*Which feedback ( I or II ) more clearly or frequently points out spelling errors?*

Annotators could choose from four choices:

- (a) *Feedback I;*
- (b) *Feedback II;*
- (c) *Both are similar;*
- (d) *Uncertain.*

Correct identification of the counterfactual feedback was recorded as "correct", incorrect identification was recorded as "incorrect", and choosing options (c) or (d) was recorded as "Indeterminate".

Three graduate students majoring in Linguistics were tasked with identifying the feedback differences within each pair. The final classification for each pair was determined by a majority vote among the annotators. In cases where no majority was reached due to each annotator choosing a different outcome, the pair was labeled as "indeterminate". Results are in Table 12. Table 13 gives an example of a feedback pair where the counterfactual feedback corresponds to a sample obtained by introducing spelling errors to the original sample.

## D.3 Ethical Considerations

The three annotators involved in this project were graduate students in linguistics. Prior to assigning them the annotation task, we provided a comprehensive introduction to the content, purpose, and significance of the project. Each annotator was responsible for reviewing 200 feedback pairs and received compensation of \$0.42 per annotated pair.

Category	Counterfactual Type	Correct%	Incorrect%	Indeterminate%
Conventions	Error Correction	72	4	24
	Spelling	68	0	32
	SVA	48	0	52
	WOS	80	0	20
Language	Complexification	100	0	0
Complexity	Simplification	32	0	68
Orgnization	InParaShuffle	40	4	56
	InTextShuffle	20	8	72

Table 12: Feedback Type Voting Results by Annotators

Feedback for the Original	Feedback for the Counterfactual
<p><b>**organization_feedback**</b> The essay begins with a clear topic sentence that introduces the importance of parks for improving mood, but the development of ideas throughout the body is somewhat disjointed. The writer provides personal anecdotes and general observations to support the argument, but transitions between ideas are abrupt, and there is a lack of clear logical flow. The conclusion does not effectively summarize the main points or provide a strong final thought on the topic. The essay would benefit from a more structured approach, with each paragraph focusing on a single aspect of the importance of parks and using transitions to guide the reader through the argument.</p> <p><b>**language_use_feedback**</b> Morphology: There are several errors in verb tense consistency ('i use to just go outside' should be 'I used to just go outside') and subject-verb agreement ('kids finds parks activity enjoyable' should be 'kids find park activities enjoyable'). Syntax: Sentence structures are often simplistic and repetitive, lacking variety and complexity. There are also instances of awkward phrasing ('my mom always straggle with tell my brothers to go outside' should be 'my mom always struggles to tell my brothers to go outside'). Vocabulary: The vocabulary is generally appropriate but lacks diversity and sophistication. There are instances of imprecise language ('be open minded' could be 'become open-minded', 'enjoy life' could be 'enjoy the outdoors').</p> <p><b>**conventions_feedback**</b> The essay contains several spelling errors ('straggle' should be 'struggle', 'rob' should be 'rope'), incorrect capitalization ('iPod', 'iPhone', 'I', and the first word of sentences should be capitalized), and punctuation mistakes (missing commas in compound sentences, incorrect spacing around punctuation marks). Attention to these details would improve the readability and professionalism of the essay.</p>	<p><b>**organization_feedback**</b> The essay's organization needs improvement. Ideas about the importance of parks and their impact on mood and childhood experiences are presented, but the logical flow is disjointed. Transitions between ideas are abrupt or missing, making it difficult to follow the argument. The introduction lacks a clear thesis statement, and the body paragraphs do not consistently develop the initial idea. The conclusion is not clearly defined, and the essay ends abruptly without summarizing the main points or providing a final thought on the importance of parks for children.</p> <p><b>**language_use_feedback**</b> Morphology: There are several errors in verb tenses and subject-verb agreement (e.g., 'my mm takes him ty the park and ha play' should be 'my mom takes him to the park and he plays'). Syntax: Sentence structures are often incomplete or incoherent, which disrupts the flow of information (e.g., 'Know i barely even go outside i go ti school and then i come back is just' is difficult to understand). Vocabulary: Word choice is sometimes inaccurate or too simplistic, and there is a lack of diversity in vocabulary. More precise language could enhance the clarity of the essay (e.g., 'bouring' should be 'boring', 'activity enjoyable' should be 'activities enjoyable').</p> <p><b>**conventions_feedback**</b> The essay contains numerous spelling errors (e.g., 'iMy', 'btother', 'ty', 'ha play', 'lf', 'alls', 'jn', 'tie', 'suing', 'activity', 'mint', 'whi offen', 'leke', 'ext', 'litter', 'bi on m'i phone', 'utside', 'tou', 'weill', 'offen mindend', 'thingh', 'dayl', 'field', 'wak ta giong tte park', 'nown', 'mer', 'estay', 'gou'). Capitalization is inconsistent, and punctuation is often missing or used incorrectly, making it challenging to understand the intended meaning of sentences.</p>

Table 13: Example feedback pair of original and counterfactual obtained by introducing spelling errors.