

# Double-Edged Sword or Sharp Tool? Designing and Evaluating Triadic LLM-Teacher Collaboration for K-12 Writing at Scale

Canran WANG<sup>1#</sup>, Yuwen Yang<sup>1#</sup>, Zhen Wang<sup>2</sup>, Ming MA<sup>3</sup>,  
Chentai Wang<sup>3</sup>, Ding Yu<sup>3</sup>, Keman Huang<sup>1\*</sup>, Xiaoyong Du<sup>1</sup>

<sup>1</sup>Renmin University of China, Beijing, China,

<sup>2</sup>Central University of Finance and Economics, Beijing, China,

<sup>3</sup>Beijing HQ Intelligent Technology, Ltd., Beijing, China

## Abstract

The double-edged sword of integrating Large Language Models (LLMs) requires an effective triadic collaboration mechanism among LLMs, teachers and students, especially for K-12 education. By developing a triadic collaboration system to support K-12 writing learning, a multidimensional evaluation framework grounded in Systemic Functional Linguistics and the suggestion trajectory tracing pipeline, this paper contributes a large-scale empirical dataset involving 57,954 essays from 10,195 students across 120 schools over two years. Our findings confirm the efficacy of this system in improving writing quality through a strategic labor division: the LLM serves as a generative engine to mitigate teacher burnout, and the teacher acts as a pedagogical gatekeeper and bridge to guarantee feedback quality. While both LLM and teacher are critical for skill improvement, we uncover a ceiling effect where excessive linguistic expansion yields diminishing marginal utility. These suggest a dynamically adaptive LLM-teacher collaboration as student proficiency increases.

## 1 Introduction

The Large Language Models (LLMs) have fundamentally transformed educational practices (Florencio and Prieto, 2025), supporting the entire learning process from curriculum design (Shojaei et al., 2025) to adaptive tutoring (Wang et al., 2024), particularly within higher education (AlGhamdi et al., 2025). However, a significant research gap persists in K-12 pedagogy: while LLMs lower the entry barrier for novice writers and enhance mnemonic retention through personalized interaction (Park et al., 2024), current studies predominantly target adult populations in controlled laboratory settings (Yuan et al., 2022; Jakesch et al., 2023), often overlooking

the developmental nuances and socio-cognitive constraints of the K-12 classroom.

Crucially, LLM adoption presents a “double-edged sword” for developing learners: while these models can encourage divergent thinking (Wang et al., 2025), they risk homogenizing collective output (Doshi and Hauser, 2024) and diminishing student ownership (Dhillon et al., 2024; Kadoma et al., 2024). Without human mediation, a reliance on algorithmic norms (Sundaresan and Guler, 2025) threatens to stifle unique cultural voices and cognitive agency (Agarwal et al., 2025). This suggests that LLMs alone are insufficient as partners for K-12 learners. Consequently, determining how LLMs and teachers can collaborate effectively to foster learning has emerged as a critical challenge for the application of NLP in pedagogy.

To bridge this gap, using K-12 writing as the research context, we design and implement a hybrid system featuring a **triadic collaboration** among students, LLMs, and teachers. This system synergizes the scalable generative power of LLMs with the pedagogical expertise of human educators, guiding students to transform their initial drafts into high-quality revisions. To evaluate this intervention, we introduce a **multidimensional evaluation framework** grounded in Systemic Functional Linguistics (SFL) (Halliday and Matthiessen, 2013), capturing learning outcomes across the *Ideational, Textual, and Interpersonal* metafunctions. Furthermore, we develop an analytical pipeline to **trace the trajectory** of how students incorporate specific suggestions into their final revisions. By deploying this system in real-world K-12 writing environments, this study asks:

**RQ1:** To what extent does the triadic collaboration system enhance student writing across the three SFL metafunctions?

**RQ2:** How is the labor divided between LLMs and teachers during feedback generation, and how are these distinct sources adopted by students?

\*The first two authors contributed equally to this work. Keman Huang is the corresponding author (keman@ruc.edu.cn).

**RQ3:** What are the underlying mechanisms through which hybrid LLM-teacher collaboration impacts writing performance?

To address these questions, we contribute **the first large-scale empirical study** centered on the K-12 iterative writing revision process for real-world K-12 writing environments. Our dataset spans over two years and comprises  $N = 57,954$  essays from 10,195 students across 120 schools. To the best of our knowledge, this is the most extensive empirical dataset in the domain of LLM-assisted K-12 education.

Furthermore, our findings provide systematic evidence for **the efficacy of triadic collaboration**, yielding a significant 5%+ improvement in rewriting quality. Beyond holistic gains cross all dimensions, our data reveals a pronounced growth in the interpersonal function, characterized by a distinct pro-social shift. Moreover, there exists **a strategic functional division of labor**: the LLM acts as a scalable generative engine to mitigate instructor burnout, and the teacher serves as a pedagogical gatekeeper and bridge to ensure feedback quality. Finally, while **both the LLM and teacher contribute significantly to performance, a critical “ceiling effect”** emerges: although linguistic expansion initially proxies performance improvement, its marginal utility for the final grade decreases as proficiency rises, with certain metrics eventually turning negative. Overall, these findings highlight that effective LLM integration in K-12 requires a dynamic triadic design to ensure technological power is successfully converted into meaningful learning outcomes.

## 2 Related Work

**LLMs in Education** LLMs have been deeply embedded in educational ecosystems (Florencio and Prieto, 2025), facilitating tasks ranging from mathematical reasoning (Wu et al., 2023) to the entire teaching-and-learning lifecycle (Chu et al., 2025). For educators, domain-specific agents serve as reflective planning tools that simulate classroom dynamics (Xu et al., 2025); for learners, they provide personalized feedback that bolsters motivation (Park et al., 2024). Despite these affordances, the proliferation of LLMs introduces critical challenges beyond academic integrity (Tang et al., 2024). Recent literature warns of cognitive deskilling, where outsourcing synthesis to algorithms leads to surface-level learning (Yang et al.,

2026). Moreover, over-reliance on LLMs may overdraw problem-solving capacities (Sundaresan and Guler, 2025), necessitating rigorous evaluation of AI’s long-term pedagogical impact.

Critically, these conflicting outcomes remain unresolved due to a methodological reliance on adult cohorts in controlled settings (such as (Agarwal et al., 2025; Jakesch et al., 2023), etc.). This demographic bias limits the generalizability of findings to K-12 contexts, necessitating empirical evidence that specifically targets the multidimensional developmental trajectories of developing writers.

**Efficacy of LLM-Assisted Writing** Current literature primarily frames LLM-assisted writing as a dyadic collaboration between a single user and an agent. While this partnership lowers entry barriers (Wang et al., 2024) and boosts efficiency (Yuan et al., 2022; Agarwal et al., 2025; Dhillon et al., 2024), it often results in a “double-edged sword” effect: quality gains often come at the expense of distinctiveness and psychological ownership (Dhillon et al., 2024). Critically, unmediated dyadic interaction risks linguistic homogenization (Guo et al., 2024; Padmakumar and He, 2023) and convergence toward Western-centric norms, potentially compelling minority users to abandon unique expressive styles (Jakesch et al., 2023; Agarwal et al., 2025; Kadoma et al., 2024). Furthermore, these interventions may exacerbate inequalities, with outcomes varying significantly based on gender and cognitive capabilities (Lu et al., 2025).

Notably, existing research prioritizes dyadic “Writer-LLM” pairs in controlled settings, failing to capture the socially mentored nature of K-12 writing: efficacy is moderated by a **triadic collaboration** where teachers act as pedagogical mediators, filtering and contextualizing LLM outputs to align with learning objectives. Consequently, this study shifts the lens to this “LLM-Teacher-Student” collaboration to evaluate its impact on multidimensional linguistic development.

## 3 Methodology

### 3.1 The LLM-Teacher Collaboration System for K-12 Student Writing

As illustrated in Figure 1, we developed a hybrid system featuring a triadic collaboration between students, LLMs, and teachers. This architecture synergizes the scalable analytical power of LLMs with the pedagogical expertise of human educators

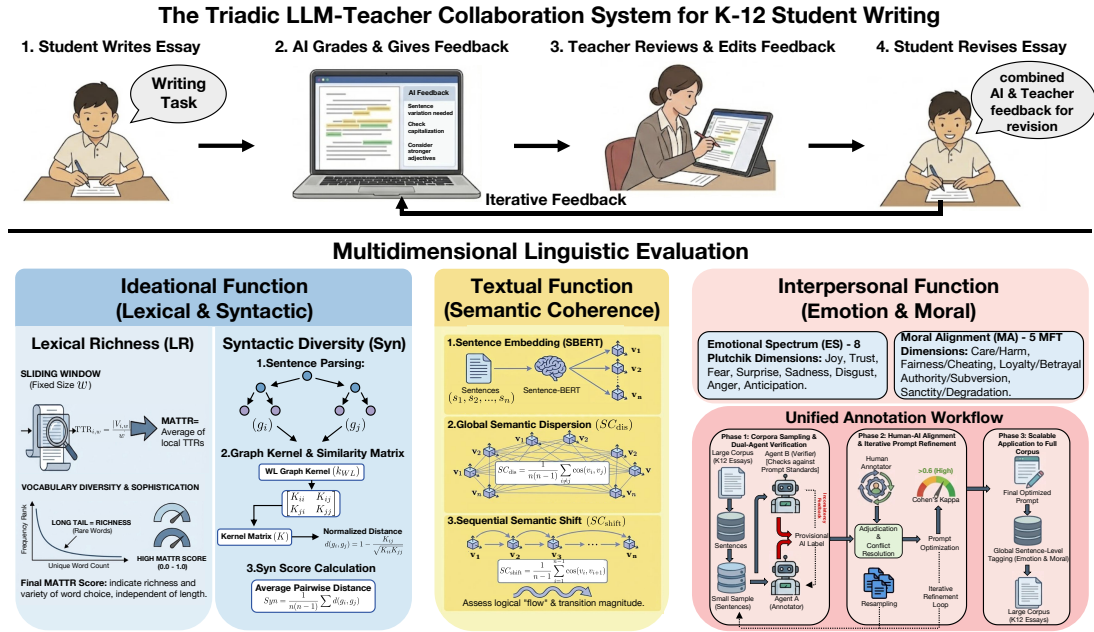


Figure 1: **The LLM-Teacher Collaboration System for K-12 Student Writing.** The upper panel illustrates the iterative workflow where teacher expertise refines LLM-generated suggestions ( $S_{initial}$ ) into actionable feedback ( $S_{final}$ ) for student revisions. The lower panel displays the evaluation metrics based on the Halliday’s Systemic Functional Linguistics (SFL), covering Ideational ( $Lex, Syn$ ), Textual ( $SC_{dis}, SC_{shift}$ ), and Interpersonal ( $ES, MA$ ) linguistic dimensions.

to augment K-12 student writing skills. It transforms initial student drafts into enhanced compositions through a four-stage iterative process:

### Stage 1: Writing Task and Draft Generation.

The process initiates with students engaging in a standardized writing task tailored to K-12 curriculum standards. Students compose an initial draft, denoted as  $D_{pre}$ , which reflects their baseline linguistic and thematic proficiency.

**Stage 2: LLM-Generated Feedback.** Upon submission, the initial draft ( $D_{pre}$ ) is processed by the LLM, acting as an expert writing assistant. It rates the baseline writing quality ( $grade_{LLM}^{pre}$ ) and generates a set of preliminary, constructive suggestions ( $S_{initial}$ ) across multiple dimensions.

**Stage 3: Teacher Review and Refinement.** Rather than passing the LLM output directly to the student, the system incorporates a teacher who accesses the raw LLM suggestions ( $S_{initial}$ ) alongside the student’s draft ( $D_{pre}$ ). The teacher screens and refines the suggestions, including validating accurate insights, adjusting the tone, removing inappropriate or redundant outputs, adding missing pedagogical scaffolding to finalize the suggestions ( $S_{final}$ ) and providing the rating ( $grade_{teacher}^{pre}$ ).

**Stage 4: Feedback Adoption and Student Revision.** Upon receiving the teacher-mediated feed-

back ( $S_{final}$ ), students review and revise their original draft ( $D_{pre}$ ) to produce a finalized version ( $D_{post}$ ). To foster continuous improvement, this new draft cycles back for additional assessment, generating both the LLM evaluation ( $grade_{LLM}^{post}$ ) and the teacher assessment ( $grade_{teacher}^{post}$ ).

## 3.2 Multidimensional Linguistic Evaluation

To provide a holistic assessment of student writing development, we ground our evaluation framework in Halliday’s Systemic Functional Linguistics (SFL) (Halliday and Matthiessen, 2013), which treats language as a social semiotic system where meaning-making occurs across three simultaneous functions (Ideational, Textual, and Interpersonal).

### 3.2.1 Ideational Function

The ideational function, measured by Lexical Richness ( $Lex$ ) and Syntactic Diversity ( $Syn$ ), quantifies the student’s ability to deploy diverse vocabulary and complex grammatical structures.

**Lexical Richness ( $Lex$ ):** evaluates a student’s expressive capacity and vocabulary breadth. For a given essay, we utilize the Moving-Average Type-Token Ratio (MATTR) (Agarwal et al., 2025) to calculate the Type-Token Ratio within a sliding window of fixed size  $w$ , which moves incremen-

tally through the whole essay.

$$\text{MATTR} = \frac{1}{N - w + 1} \sum_{i=1}^{N-w+1} \frac{|V_{i,w}|}{w} \quad (1)$$

where  $|V_{i,w}|$  is the number of unique tokens in the window starting at position  $i$ ,  $w$  is the window size and  $N$  is the total number of tokens. A higher lexical richness score indicates higher richness and variety of word choice, typically associated with more advanced writing proficiency.

**Syntactic Diversity ( $Syn$ ):** We utilize a graph-based kernel approach (Guo et al., 2024) to capture the structural variation within each essay. Specifically, each sentence is parsed into a dependency graph  $g$ , where nodes represent words labeled with their part-of-speech (POS) tags and edges represent dependency relations (Che et al., 2021). For a set of  $n$  sentence graphs, we calculate the Weisfeiler-Lehman (WL) graph kernel (Sherashidze et al., 2011) matrix  $\mathbf{K}$  to derive the average pairwise distance between all sentences:

$$Syn = \frac{1}{n(n-1)} \sum_{i \neq j} \left(1 - \frac{k_{ij}}{\sqrt{k_{ii}k_{jj}}}\right) \quad (2)$$

where  $k_{ij}$  is the structural similarity between the dependency graph for sentence  $s_i$  and  $s_j$  calculated by the WL kernel. A higher  $Syn$  score indicates a broader range of grammatical structures and reduced syntactic monotonicity, typically representing a higher writing proficiency.

### 3.2.2 Textual Function

The textual function is operationalized through Semantic Dispersion ( $SC_{dis}$ ) and Semantic Shift ( $SC_{shift}$ ) to capture semantic coherence, which reflects how individual linguistic choices are woven into a logically and conceptually consistent student essay. In particular, we represent each sentence  $s_i$  as a high-dimensional vector  $\mathbf{v}_i$  using a pre-trained Sentence-BERT (SBERT) model (Reimers and Gurevych, 2019) and then quantify semantic coherence through two distinct measures:

**Semantic Dispersion ( $SC_{dis}$ ):** We measure the overall breadth of ideas within an essay by calculating the average pairwise cosine distance between all sentence embeddings:

$$SC_{dis} = \frac{1}{n(n-1)} \sum_{i \neq j} \cos(\mathbf{v}_i, \mathbf{v}_j) \quad (3)$$

where  $n$  is the number of sentences and a higher  $SC_{dis}$  indicates a broader semantic coverage and a more diverse range of topics within the essay.

**Semantic Shift ( $SC_{shift}$ ):** To assess the logical flow of the essay, we calculate the average distance between contiguous sentences:

$$SC_{shift} = \frac{1}{n-1} \sum_{i=1}^{n-1} \cos(\mathbf{v}_i, \mathbf{v}_{i+1}) \quad (4)$$

This sequential metric captures the magnitude of topical transitions. In the context of student writing, a very low  $SC_{shift}$  score may indicate repetitive or redundant transitions.

### 3.2.3 Interpersonal Function

The interpersonal function is measured using the Emotional Spectrum ( $ES$ ) and Moral Alignment ( $MA$ ) to evaluate students' capacity to engage readers through pathos and ethical reasoning.

**Emotional Spectrum ( $ES$ ):** We ground our analysis in Plutchik's wheel of emotions (Plazadel Arco et al., 2024), which identifies eight primary emotional dimensions in two sets: approach-oriented emotions including *Anticipation*, *Joy*, *Trust*, *Fear* and avoidance-oriented emotions including *Surprise*, *Sadness*, *Disgust*, *Anger*. To facilitate large-scale annotation, we developed a two-stage multi-agent pipeline utilizing the social-scientific labeling capabilities of LLMs. The process begins with iterative multi-agent discussions and human verification to optimize labeling prompts; our pilot validation yielded a Cohen's Kappa of 0.61, indicating strong alignment between the system and human experts. This optimized prompt was then deployed for large-scale automated annotation, where the Fleiss's Kappa for the final annotations reaches 0.981, confirming the high consistency for our labels. We provide the complete technical architecture, the finalized prompts, and the details of the consistency verification process in Appendix A.

Through this process, the emotional state of an essay is represented as an 8-dimensional vector:  $\mathbf{e} = [e_1, \dots, e_8]$  where  $e_i$  denotes the proportion of sentences categorized within a specific emotional dimension relative to the total sentence count  $n$  of the essay. To quantify the balance of these emotions, we calculate the Emotional Spectrum ( $ES$ ) using Shannon entropy:

$$ES = - \sum_{i=1}^8 \frac{e_i}{n} \log \frac{e_i}{n} \quad (5)$$

A higher  $ES$  score indicates a wider range of emotional registers, whereas a lower score suggests an emotionally flat or monotonic essay.

**Moral Alignment (MA).** Our analysis is based on the Moral Foundations Theory (MFT) (Abdulhai et al., 2024), which identifies five positive ethics (*Care, Fairness, Loyalty, Authority, Sanctity*) and corresponding negative ethics (*Harm, Cheating, Betrayal, Subversion, Degradation*). We follow the same pipeline to reconfigure the multi-agent system to classify the moral alignment of each related sentence. The pilot validation yielded a Cohen’s Kappa of 0.67 and the agent consistency score reaches a Fleiss’s Kappa of 0.989, confirming the effectiveness of the system.

Through this process, the moral alignment for each essay is represented as a 10-dimensional value vector:  $\mathbf{m} = [m_1, \dots, m_{10}]$ , where  $m_i$  denotes the normalized frequency of a specific moral category within the text. Finally, we compute the moral alignment (MA) Shannon entropy to quantify the moral diversity of an essay:

$$MA = - \sum_{i=1}^{10} \frac{m_i}{n} \log \frac{m_i}{n} \quad (6)$$

A higher MA score signifies that the student has engaged with a diverse set of moral perspectives.

### 3.3 Feedback Uptake Quantification

To unfold how students incorporate the feedback from the hybrid LLM-teacher system into their revision, we develop a four-step pipeline as follows:

**Suggestion Extraction and Classification.** For each essay, we collect the raw LLM-generated suggestions  $\mathcal{S}_{\text{initial}}$  from stage 2 and the instructor-mediated finalized suggestions  $\mathcal{S}_{\text{final}}$  from stage 3, and then map all suggestions into a shared semantic space using Sentence-BERT (SBERT).

For each finalized suggestion  $s_j \in \mathcal{S}_{\text{final}}$ , we calculate its maximum similarity against the raw LLM set  $s_i \in \mathcal{S}_{\text{initial}}$ :  $\max_{s_i \in \mathcal{S}_{\text{initial}}} (\cos(v_j, v_i))$  where  $v_i$  and  $v_j$  are the vector embeddings of  $s_i$  and  $s_j$ . If it is larger than the threshold  $\delta_m = 0.75$ , the suggestion is categorized as  $\mathcal{S}_L$  representing teacher retention of LLM feedback. Otherwise, it is categorized as  $\mathcal{S}_T$ , meaning the teacher significantly revises the LLM output or injects novel contents.

Finally, we categorize each suggestion  $s_j \in \mathcal{S}_{\text{final}}$  based on its fine-grained linguistic focus into one of six dimensions. We provide the categorization method in Appendix B.

**Revision Candidate Identification** To isolate the student’s response to feedback, we identify

the subset of sentences in the revised draft  $D_{\text{post}}$  that underwent meaningful modification. Using the difflib sequence matching algorithm, we compute the similarity between sentences in  $D_{\text{pre}}$  and  $D_{\text{post}}$ . We define a set of revision candidates  $\mathcal{C} = \{c_1, c_2, \dots, c_m\}$ , where a sentence  $s \in D_{\text{post}}$  is included in  $\mathcal{C}$  if its maximum similarity to any sentence in  $D_{\text{pre}}$   $\max_{s' \in D_{\text{pre}}} \text{sim}(s, s')$  is lower than the threshold  $\delta_r = 0.95$ . This ensures that  $\mathcal{C}$  represents only the active revisions made by the student in the second iteration while excluding those remain unchanged.

**Attention-based Semantic Matching** We develop an attention mechanism to determine if a suggestion  $s_j \in \mathcal{S}_{\text{final}}$  is adopted by the student. In particular, given a suggestion  $s_j$  and the revision candidates  $\mathcal{C}$ , we compute attention weights  $\alpha_{ij}$  using a softmax function:

$$\alpha_{ji} = \frac{\exp(\cos(v_j, v_i)/\tau)}{\sum_{i=1}^m \exp(\cos(v_j, v_i)/\tau)} \quad (7)$$

where  $v_j$  and  $v_i$  are the vector embeddings of  $s_j$  in the suggestion set and  $c_i$  in the revision candidate set, while  $\tau = 0.1$  is the temperature parameter. The matching score  $Match_j$  for suggestion  $s_j$  is the weighted sum of similarities:

$$Match_j = \sum_{i=1}^m \alpha_{ji} \cdot \cos(v_j, v_i) \quad (8)$$

A finalized suggestion  $s_j \in \mathcal{S}_{\text{final}}$  is considered adopted if its matching score  $Match_j$  exceeds a specific adoption threshold  $\delta_a = 0.5$ .

**Feedback Uptake Ratio Calculation** The Feedback Uptake Rate (FUR) is the proportion of the total count of suggestions adopted by the student, Feedback Uptake Amount (FUA), relative to the total number of provided suggestions  $|\mathcal{S}_{\text{final}}|$ :

$$FUA = \sum_{j=1}^{|\mathcal{S}_{\text{final}}|} \mathbb{I}(\text{Match}_j > \delta_a) \quad FUR = \frac{FUA}{|\mathcal{S}_{\text{final}}|} \quad (9)$$

where  $\mathbb{I}(\cdot)$  denotes the indicator function. To evaluate the specific impact of teacher intervention, we decompose these metrics into  $FUA_L$  and  $FUR_L$  for suggestions from the LLM subset ( $\mathcal{S}_L$ ), and  $FUA_T$  and  $FUR_T$  for those refined by the teacher ( $\mathcal{S}_T$ ).

## 4 Experiments and Results

### 4.1 System Implementation and Dataset

The hybrid LLM-teacher system was implemented and deployed in a real-world K-12 writing platform

Table 1: Statistics of the constructed dataset.

Statistic	Value
Time Span	2023.05 – 2025.07
# Schools	120
# Students	10,195
# Writing Tasks	1,602
<b>Essay</b>	
# Total Essays	57,954
# Average Sentences	14.56
# Average Tokens	129.01
# Student-Task Pairs	28,977
<b>Grade (Mean)</b>	
$\text{grade}_{\text{LLM}}^{\text{pre}}$ vs $\text{grade}_{\text{LLM}}^{\text{post}}$	80.74 vs 85.75 ***
$\text{grade}_{\text{Teacher}}^{\text{pre}}$ vs $\text{grade}_{\text{Teacher}}^{\text{post}}$	75.17 vs 80.34 ***

starting in 2023. For this study, we focus on the process where a student submitted a revised draft following the receipt of the feedback, resulting in a large-scale dataset centered on this iterative revision process. To ensure data integrity, we excluded invalid submissions lacking corresponding LLM or teacher grades. Then, all scores were normalized to a standard  $[0, 100]$  scale. Furthermore, we applied a 2.5% trimming strategy based on sentence and token counts to remove outliers.

As reported in Table 1, the dataset spanned over two years, involving 10,195 students across 120 distinct schools. The resulting corpus comprises 57,954 essays derived from 1,602 unique writing tasks, yielding a total of 28,977 student-task pairs. On average, each essay contains 129.01 tokens and 14.56 sentences. Each pair captures a complete cycle of the writing process, consisting of the initial draft ( $D_{\text{pre}}$ ), the raw LLM-generated suggestion ( $S_{\text{initial}}$ ), the instructor-mediated finalized suggestion ( $S_{\text{final}}$ ) and the revised version ( $D_{\text{post}}$ ).

Notably, on average, **student’s performance significantly<sup>1</sup> improved across both evaluative lenses**: teacher-assigned grades increased by 5.17 points, from  $\text{grade}_{\text{Teacher}}^{\text{pre}}$  to  $\text{grade}_{\text{Teacher}}^{\text{post}}$ , and LLM-assigned grades by 5.01 points, from  $\text{grade}_{\text{LLM}}^{\text{pre}}$  to  $\text{grade}_{\text{LLM}}^{\text{post}}$  in the revised versions. Notably, the increase in teacher-assigned grades was significantly greater than the increase in LLM-assigned grades. These confirm the efficacy of our triadic system in fostering measurable writing skill improvement.

<sup>1</sup>Given the non-normal distribution of the data (Shapiro-Wilk test,  $p < 0.001$ ), we employ the Wilcoxon Signed-Rank Test for paired comparisons and the Mann-Whitney U Test for independent samples. In this study, statistical significance levels are denoted as follows: \*\*\* :  $p < 0.001$ ; \*\* :  $p < 0.01$ ; \* :  $p < 0.05$ ; † :  $p < 0.1$ ;  $^{ns}$  :  $p > 0.1$ ;

## 4.2 Results

### 4.2.1 Multidimensional Linguistic Growth

**Holistic Writing Skill Growth cross ALL Linguistic Dimensions with Pronounced Gains in the Interpersonal Function.** As detailed in Table 2, we observe a statistically significant improvement across all six dimensions, revealing a holistic writing skill improvement. Notably, the interpersonal function demonstrates the most substantial growth, with the emotional spectrum expanding by 7.719% and moral alignment by 5.735%. This suggests that our system transcends mere formal correction: **beyond refining students’ writing mechanics and structural cohesion, the hybrid LLM-teacher intervention can foster deeper ideological maturity and cognitive complexity**, enabling students to navigate more nuanced emotional and moral landscapes in their revisions.

Table 2: Linguistic growth between the initial drafts ( $D_{\text{pre}}$ ) and the refined versions ( $D_{\text{post}}$ ).

Dimension	Mean <sub>pre</sub>	Mean <sub>post</sub>	$\Delta\%$
Lexical Richness	0.850	0.875	+3.313***
Syntactic Diversity	0.696	0.701	+1.334***
Semantic Dispersion	0.483	0.486	+1.556***
Semantic Shift	0.524	0.528	+1.641***
Emotional Spectrum	1.115	1.231	+7.719***
Moral Alignment	0.480	0.536	+5.735***

**Shift Toward Pro-social Emotional and Moral Narratives.** Table 3 reveals a distinct polarity shift for the interpersonal function improvement: in approach-oriented emotions, including *Anticipation, Joy, Trust and Fear*, all exhibiting significant growth. However, avoidance-oriented emotions, specifically Anger and Disgust, are significantly reduced. A similar pro-social effect appears in moral alignment. The positive foundations, including *Sanctity, Care, and Loyalty*, increase significantly, while negative foundations (*Harm, Cheating, Subversion, and Degradation*) are consistently suppressed. Notably, decreases in *Authority* and *Fairness* suggest a shift toward more egalitarian narrative styles. **These indicate that the intervention steers students away from conflict-based toward more pro-social narratives.**

### 4.2.2 Labor Division and Feedback Adoption

**Hybridized Team Collaboration with Teacher’s Pedagogical Bridge Role.** As reported in Table 4, the collaborative synergy between LLM and teacher results in a hybridized feedback landscape where the final suggestion pool ( $S_{\text{final}}$ ) consists

Table 3: Shifts in emotional and moral dimensions

Emotion	$\Delta\%$	Moral	$\Delta\%$
<b>I. Approach-oriented</b>		<b>I. Positive</b>	
Anticipation	+2.498***	Sanctity	+1.513***
Joy	+1.789***	Care	+0.767***
Trust	+0.975***	Loyalty	+0.509***
Fear	+0.305***	Authority	-0.127***
		Fairness	-0.021**
<b>II. Avoidance-oriented</b>		<b>II. Negative</b>	
Surprise	+0.235***	Degradation	-0.017**
Sadness	+0.212 <sup>ns</sup>	Harm	-0.028***
Disgust	-0.073***	Betrayal	-0.003 <sup>ns</sup>
Anger	-0.141***	Subversion	-0.029***
		Cheating	-0.052***

of 58.2% raw LLM suggestions ( $\mathcal{S}_L$ ) and 41.8% teacher-mediated refinements ( $\mathcal{S}_T$ ). The LLM provides the majority of the volume, particularly in the moral alignment category, which is almost entirely LLM-driven (98.912%), while teachers provide a higher relative contribution to the semantic shift dimension to ensure logical cohesion.

However, teachers demonstrate a significantly higher feedback adoption ratio overall and across nearly all individual dimensions, with the sole exception of non-significant difference in semantic shift. Notably, our design ensured that students could not distinguish the source of the provided feedback directly. This “trust premium” for teacher feedback confirms that the teacher’s refinements, acting as a “Pedagogical Bridge”, possess an inherent qualitative alignment with student needs that raw LLM output lacks. Hence, **while the LLM acts as the primary generative engine, teacher intervention remains the critical bridge for actual student implementation.**

Table 4: Finalized suggestions contributed by LLM and teacher, and their adoption rate

Dimension	Feed. Share (%)			FUR (%)		
	$\mathcal{S}_L$	$\mathcal{S}_T$	Sig.	FUR <sub>L</sub>	FUR <sub>T</sub>	Sig.
Lexical Richness	53.037	46.963	***	85.368	91.761	***
Syntactic Diversity	62.847	37.153	***	79.369	93.539	***
Semantic Dispersion	76.280	23.720	***	84.573	91.311	***
Semantic Shift	39.140	60.860	***	92.621	93.701	<sup>ns</sup>
Emotional Spectrum	54.075	45.925	***	88.926	94.026	***
Moral Alignment	98.912	1.088	***	82.699	84.463	*
Overall	58.158	41.842	***	84.688	91.895	***

**Strategic Labor Division with the LLM as Generative Engine and Teacher as Pedagogical Gatekeeper.** We quantify *Teacher Effort* as the volume of modifications to LLM output plus manual additions, calculated as  $|\mathcal{S}_{\text{initial}} - \mathcal{S}_L| + \mathcal{S}_T$ . We distinguish between two scenarios: *Creation*, where the teacher drafts feedback from scratch,

and *Modification*, where the teacher refines LLM-generated content.

Table 5: Teacher effort under creation and modification

Dimension	Creation	Modification	Sig. ( $\Delta$ )
Lexical Richness	1.400	0.026	*** ↓
Syntactic Diversity	0.290	0.017	*** ↓
Semantic Dispersion	0.232	0.014	*** ↓
Semantic Shift	0.520	0.024	*** ↓
Emotional Spectrum	0.270	0.022	*** ↓
Moral Alignment	0.011	0.026	*** ↑
Overall	3.951	0.095	*** ↓

As reported in Table 5, **LLM integration significantly reduces the teacher’s burden.** In the creation scenario, total teacher effort is 3.951, dominated by lexical richness (1.400), followed by semantic shift (0.520) and syntactic diversity (0.290). In contrast, in the modification scenario, total teacher effort is remarkably low (0.095), **representing a nearly 40-fold decrease in workload.** Beyond lexical richness, interventions are redistributed to moral alignment (0.026) and the emotional spectrum (0.024). Notably, when transitioning from creation to modification, moral alignment is the only dimension where teacher effort increases. This highlights the role of the teacher not just as an editor, but as a moral gatekeeper, ensuring that automated suggestions align with the specific sensitive requirements of K-12 educational standards. In other words, **the triadic collaboration does not simply replace the teacher, but reorients human expertise toward high-stakes qualitative gatekeeping, mitigating the instructor burnout common in K-12 writing programs.**

#### 4.2.3 The Triadic Collaboration Impact

**Two-channel Feedback Efficacy with Linguistic Dimensional Mediation.** As shown in Model 1 (Table 6), the volume of adopted suggestions from both the LLM (FUA<sub>L</sub>) and the teacher (FUA<sub>T</sub>) significantly predicts improvements in writing quality. **This confirms the baseline efficacy of both channels, LLMs and teachers, in the revision process. Notably, teacher-mediated suggestions yield a higher marginal coefficient than raw LLM suggestions.** This suggests that human intervention effectively “aligns” student compositions with high-level evaluative criteria, confirming the teacher’s role as a linguistic anchor that grounds generative output in pedagogical reality.

Model 2 further elucidates how specific SFL dimensions mediate these gains. Both the interpersonal metafunction (encompassing emotional spec-

trum and moral alignment) and the ideational metafunction (specifically lexical richness) emerge as potent positive mediators. While syntactic diversity contributes positively to LLM evaluations, it demonstrates a non-significant impact on final teacher grades, suggesting a discrepancy between algorithmic fluency and human-perceived quality. A critical evaluative divergence is observed within the textual metafunction: while semantic shift (local logical flow) contributes positively to the grade, semantic dispersion (global thematic breadth) actually penalizes the final score. As demonstrated in our functional division of labor analysis, LLMs intend to prioritize semantic dispersion whereas teachers prioritize semantic shift, suggesting **teacher’s crucial role in mitigating LLM’s potential over-expansion**.

Table 6: Fixed effects regression results. All models include student, teacher, and task fixed effects. Standard errors are reported in parentheses.

Predictors	Model 1		Model 2	
	Tchr	LLM	Tchr	LLM
FUA <sub>L</sub>	0.005** (0.002)	0.008*** (0.001)	0.004* (0.002)	0.007*** (0.001)
FUA <sub>T</sub>	0.005*** (0.001)	0.011*** (0.001)	0.005*** (0.001)	0.010*** (0.002)
Lexical Richness	–	–	0.003*** (0.001)	0.008*** (0.001)
Syntactic Diversity	–	–	-0.001 (0.001)	0.002* (0.001)
Semantic Dispersion	–	–	-0.001 (0.001)	-0.004** (0.001)
Semantic Shift	–	–	0.002† (0.001)	0.004*** (0.001)
Emotion Spectrum	–	–	0.003*** (0.001)	0.008*** (0.001)
Moral Alignment	–	–	0.003*** (0.001)	0.007*** (0.001)
Baseline Score	-0.009*** (0.000)	-0.007*** (0.000)	-0.009*** (0.000)	-0.006*** (0.000)
<b>R-squared</b>	0.880	0.566	0.880	0.578
<b>Observations</b>	23,647	23,647	23,467	23,467

**Diminishing Marginal Utility and the Proficiency-Based Ceiling Effect.** To unfold the mediation contribution from the linguistic growth, we analyze Spearman correlations between dimensional gains and grade improvements by partitioning the sample into quartiles based on initial linguistic proficiency.

As illustrated in Figure 2, a consistent pattern emerges across the proficiency spectrum: correlation coefficients steadily decrease from the lower to the higher quartiles. For dimensions such as syntactic diversity, semantic dispersion, and semantic

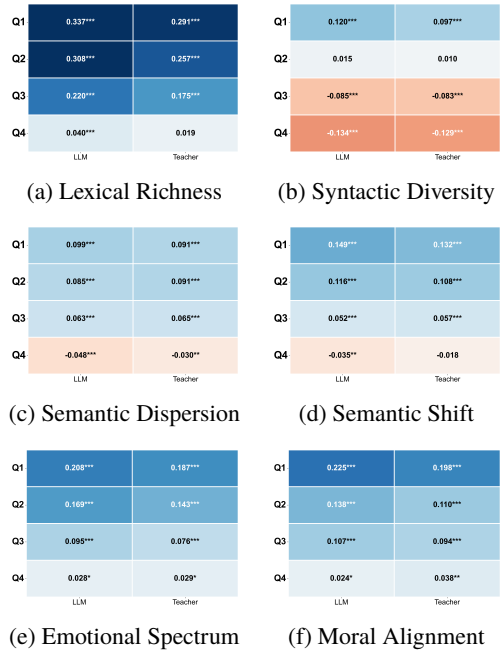


Figure 2: Spearman correlation between linguistics growth and grade split into baseline quartiles.

shift, these correlations eventually turn negative in the highest quartiles (Q4). This trend indicates that while **lower-performing students derive significant pedagogical value from linguistic expansion, higher-performing students encounter a pronounced “ceiling effect”**: further increases in structural or thematic complexity may no longer translate into higher scores which can even become counterproductive. This suggests a dynamic, proficiency-aware triadic system that adjusts its strategy (i.e. adjusting LLM’s system prompts) according to student’s linguistic proficiency.

## 5 Conclusion

This study implements and deploys a triadic collaboration model, integrating LLMs, teachers and students, within real-world K-12 writing environment. By conducting the first large-scale empirical dataset regarding the iterative process and quantifying the linguistic skill improvement from this collaboration, this study provides the empirical evidence demonstrating the efficacy of the designed triadic collaboration model.

More importantly, the system achieves an effective labor division: the LLM functions as a highly efficient generative engine to mitigate instructor burnout while teacher intervention shifts to a strategic gatekeeper and bridge for suggestion adoption. Moreover, while both LLM and teacher contribute to writing skill improvement, we identify a proficiency-based ceiling effect: while lin-

guistic expansion significantly aids novice writers, higher-proficiency students encounter diminishing marginal utility where excessive complexity can become counterproductive. These findings advocate for a triadic collaboration system making LLMs sharp tools for K-12 education.

## Limitations

While this study provides empirical evidence for the triadic collaboration model when implementing LLMs for K-12 students, several limitations warrant consideration:

First, the generalizability of the SFL framework used to quantify linguistic growth may vary across writing genres. While the SFL framework has been widely adopted in linguistic evaluation, the weight of specific linguistic dimensions might manifest differently. Exploring how the impact of these dimensions varies across genres (e.g., argumentative, expository, narrative, analytical) would be highly valuable for refining AI feedback strategies.

Second, this research focuses on the short-term impact of our triadic collaboration model. We did not fully consider the long-term longitudinal effects, such as the potential emergence of automation bias. For example, teachers might become increasingly over-reliant on the LLM's generative engine, potentially degrading the "strategic gatekeeping" and "pedagogical bridge" roles we identified.

Additionally, our study is bounded by the specific capabilities of the underlying LLMs. In real-world K-12 environments, it is often technically or financially unfeasible to frequently update or switch underlying models. Future work should investigate the implementation of the suggested dynamically adaptive triadic system to test its efficacy and robustness.

Finally, several external factors (such as the AI literacy of teachers and the socio-economic diversity of school environments) may influence the results. While we employ the fixed effect to control these factors' impact, assessing the equity of our system across these variables would be highly valuable.

## Ethical Consideration and Design

The responsible deployment of Large Language Models (LLMs) in K-12 educational settings is paramount. We identify and address three primary ethical domains critical to our triadic collaboration model:

First, a central ethical concern in educational NLP is the potential for LLMs to displace human educators. Our study provides strong empirical evidence against this "automation" narrative by demonstrating that the triadic model redefines, rather than removes, the teacher's role. By offloading repetitive linguistic drafting to the LLM (the "Generative Engine"), the system creates space for teachers to focus on high-stakes pedagogical gatekeeping and bridging. Our results show that both teacher and LLM interventions are indispensable for student development; however, teachers maintain a higher adoption rate and a higher marginal coefficient for improvement. This confirms that LLM implementation must be designed to mitigate teacher burnout while maintaining human accountability at the center of the learning process.

Second, to prevent cognitive disengagement, our model avoids generating "raw" complete writing. Instead, the LLM acts as a teaching assistant, providing granular suggestions that require student evaluation and revision. The teacher acts as a pedagogical bridge, filtering these suggestions to ensure alignment with specific educational objectives. Furthermore, our discovery of the proficiency-based ceiling effect serves as an ethical safeguard: it warns against the risks of "overscaffolding" advanced learners with excessive linguistic complexity. By advocating for a dynamically adaptive system, we protect student agency and ensure that LLM intervention supports genuine skill acquisition rather than mere automated completion.

Third, this study was conducted with the approval of participating school administrations and was deemed exempt from Institutional Review Board (IRB) review as it did not involve direct intervention on students. Nevertheless, we strictly adhered to standard IRB ethics and privacy requirements. We utilized exclusively de-identified textual data, ensuring no Personally Identifiable Information (PII) was processed. All analysis was conducted within a secure, access-restricted environment, effectively controlling for any potential risks of data leakage.

## Declaration of Generative AI in the Writing Process.

During the preparation of this work the authors used LLMs like ChatGPT, Gemini, Claude and Deepseek, in order to improve language and read-

ability. After using the generative AI service, the authors reviewed and edited the content and take full responsibility for the content of the submission.

## Acknowledgments

The work was supported by the National Natural Science Foundation of China (62441230, 62172425), the Scientific Research Innovation Capability Support Project for Young Faculty (Grant No. SRICSPYF-ZY2025001), and the Fundamental Research Funds for the Central Universities and the Research Funds of Renmin University of China (22XNKJ04).

## References

- Marwa Abdulhai, Gregory Serapio-Garcia, Clément Crepy, Daria Valter, John Canny, and Natasha Jaques. 2024. Moral foundations of large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17737–17752.
- Dhruv Agarwal, Mor Naaman, and Aditya Vashista. 2025. Ai suggestions homogenize writing toward western styles and diminish cultural nuances. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–21.
- Eman AlGhamdi, Yuheng Li, Dragan Gašević, and Guanliang Chen. 2025. Leveraging prompt-based llms for automated scoring and feedback generation in higher education. *Computers & Education*, page 105511.
- Wanxiang Che, Yunlong Feng, Libo Qin, and Ting Liu. 2021. N-ltp: An open-source neural language technology platform for chinese. In *Proceedings of the 2021 conference on empirical methods in natural language processing: System demonstrations*, pages 42–49.
- Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jinheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S Yu, and 1 others. 2025. Llm agents for education: Advances and applications. *arXiv preprint arXiv:2503.11733*.
- Paramveer S Dhillon, Somayeh Molaei, Jiaqi Li, Maximilian Golub, Shaochun Zheng, and Lionel Peter Robert. 2024. Shaping human-ai collaboration: Varied scaffolding levels in co-writing with language models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–18.
- Anil R Doshi and Oliver P Hauser. 2024. Generative ai enhances individual creativity but reduces the collective diversity of novel content. *Science advances*, 10(28):eadn5290.
- Marcos Florencio and Francielle Prieto. 2025. Vibe learning: Education in the age of ai. *arXiv preprint arXiv:2511.01956*.
- Yanzhu Guo, Guokan Shang, Michalis Vazirgiannis, and Chloé Clavel. 2024. The curious decline of linguistic diversity: Training language models on synthetic text. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3589–3604.
- Michael Alexander Kirkwood Halliday and Christian MIM Matthiessen. 2013. *Halliday's introduction to functional grammar*. Routledge.
- Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-writing with opinionated language models affects users' views. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, pages 1–15.
- Kowe Kadoma, Marianne Aubin Le Quere, Xiyu Jenny Fu, Christin Munsch, Danaë Metaxa, and Mor Naaman. 2024. The role of inclusion, control, and ownership in workplace ai-mediated communication. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–10.
- Jinwei Lu, Yikuan Yan, Keman Huang, Ming Yin, and Fang Zhang. 2025. Do we learn from each other: Understanding the human-ai co-learning process embedded in human-ai collaboration. *Group Decision and Negotiation*, 34(2):235–271.
- Vishakh Padmakumar and He He. 2023. Does writing with language models reduce content diversity? *arXiv preprint arXiv:2309.05196*.
- Minju Park, Sojung Kim, Seunghyun Lee, Soonwoo Kwon, and Kyuseok Kim. 2024. Empowering personalized learning through a conversation-based tutoring system with student modeling. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–10.
- Flor Miriam Plaza-del Arco, Alba A Cercas Curry, Amanda Cercas Curry, and Dirk Hovy. 2024. Emotion analysis in nlp: Trends, gaps and roadmap for future directions. In *Proceedings of the 2024 joint international conference on computational linguistics, language resources and evaluation (LREC-COLING 2024)*, pages 5696–5710.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992.
- Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. 2011. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9).

- Mostafa Faghieh Shojaei, Rahul Gulati, Benjamin A Jaspersen, Shangshang Wang, Simone Cimolato, Dangli Cao, Willie Neiswanger, and Krishna Garikipati. 2025. Ai-university: An llm-based platform for instructional alignment to scientific classrooms. *arXiv preprint arXiv:2504.08846*.
- Shirish Sundaresan and Isin Guler. 2025. Algorithmic recommendation tools and experiential learning in clinical care. *Organization Science*.
- Ruixiang Tang, Yu-Neng Chuang, and Xia Hu. 2024. The science of detecting llm-generated text. *Communications of the ACM*, 67(4):50–59.
- Dawei Wang, Difang Huang, Haipeng Shen, and Brian Uzzi. 2025. A large-scale comparison of divergent creativity in humans and large language models. *Nature Human Behaviour*, pages 1–10.
- Shen Wang, Tianlong Xu, Hang Li, Chaoli Zhang, Joleen Liang, Jiliang Tang, Philip S Yu, and Qingsong Wen. 2024. Large language models for education: A survey and outlook. *arXiv preprint arXiv:2403.18105*.
- Yiran Wu, Feiran Jia, Shaokun Zhang, Hangyu Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, Qingyun Wu, and Chi Wang. 2023. Mathchat: Converse to tackle challenging math problems with llm agents. *arXiv preprint arXiv:2306.01337*.
- Songlin Xu, Hao-Ning Wen, Hongyi Pan, Dallas Dominguez, Dongyin Hu, and Xinyu Zhang. 2025. Classroom simulacra: Building contextual student generative agents in online education for learning behavioral simulation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–26.
- Bo Yang, Yongqiang Sun, Zihan Zeng, and Qinwei Li. 2026. Deskillling, reskillling, or upskilling? unpacking the pathways of student adaptation to generative artificial intelligence. *International Journal of Information Management*, 87:103002.
- Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: story writing with large language models. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*, pages 841–852.

## A Automatic Emotion and Moral Annotation Pipeline

Figure 3: **Schematic overview of the automatic emotion and moral annotation pipeline.** Essays are segmented into sentences and filtered for emotion or moral relevance. Relevant sentences are annotated using an iterative LLM-based framework with critical feedback (A–B–A). Annotation prompts are calibrated based on human–model agreement before large-scale annotation. Sentence-level labels are finally aggregated at the writing level for distributional analysis.

Figure 3 provides an overview of the complete annotation workflow. We describe each stage in detail below. In this study, both emotion and moral labels are generated through an automated annotation pipeline designed for large-scale text processing. The pipeline consists of four main stages: sentence-level preprocessing and relevance filtering; an iterative large language model (LLM) annotation framework guided by critical feedback (A–B–A); prompt calibration based on human–model agreement; and large-scale full annotation. To ensure the usability and reproducibility of the full annotation results, the full-run stage incorporates strict output format validation and completeness control. Missing samples are eliminated through set-difference detection and repeated re-runs until coverage is complete. Finally, sentence-level outputs are aggregated at the writing level to produce vectorized indicators that reflect the content structure of each entire essay.

### A.1 Sentence-Level Preprocessing and Relevance Filtering

The annotation pipeline starts with essay-level texts. The raw data are organized at the essay level, and each essay is assigned a stable identifier (`writing_id`). This identifier is kept consistent across all intermediate files throughout the pipeline to enable cross-step alignment and traceability. Given each `writing_id` and its corresponding text, the essays are segmented into sentence-level analysis units using Chinese punctuation marks (e.g., “。”, “?”, “! ”) and newline characters (`\n`). Within each essay, sentences are assigned a primary sentence-level key (`sentence_id`) according to their order of appearance in the original text.

Before entering LLM-based relevance judgment, a rule-based coarse filtering step is applied to remove fragments that clearly lack semantic content, such as symbol-only strings or extremely short texts of no more than four Chinese characters. Subsequently, a large language model (GPT-4) is used to determine whether the remaining sentences are related to emotion or moral content. This stage outputs only a binary decision, classifying sentences into filter pass and filter reject (AI-based). Only sentences that pass the filter proceed to fine-grained classification.

## **A.2 Iterative Sentence-Level Classification with Critical Feedback Guidance**

To improve the stability and consistency of automatic annotations, we adopt an iterative LLM-based annotation framework guided by critical feedback, structured as an A–B–A process. The workflow is as follows.

First, the primary annotation model (Agent A; GPT-4.1, temperature = 0) generates an initial classification result (A1), including the predicted label, a confidence score, and a brief justification. Next, a review model (Agent B; Claude Sonnet 4, temperature = 0) evaluates A1 and produces an agree/disagree judgment, a critical critique, and a suggested label. Finally, Agent A revises its initial decision based on the reviewer’s feedback and outputs the final classification result (A2), along with a final confidence score and justification.

The sentence-level final decision is taken from the A2 output. The same A–B–A mechanism and validation framework are used for both emotion and moral annotation tasks, with task-specific label sets fixed as Plutchik’s eight basic emotions for emotion classification and a 10-class Moral Foundations Theory label set for moral classification.

## **A.3 Iterative Annotation with Critical Feedback and Large-Scale Full Annotation**

### **A.3.1 Prompt Calibration Based on Human–Model Agreement**

Under the A–B–A sentence-level classification framework, model judgments are highly dependent on explicit constraints in the prompts regarding task definitions, label semantics, and output formats. To ensure that model classification logic aligns statistically with human understanding before large-scale deployment, we conduct systematic iterative prompt calibration for both emotion and moral classification tasks.

In each calibration round, 20 sentences are randomly sampled from the relevance-filtered sentence pool to form a batch. Human annotators first label these sentences with emotion or moral categories. In parallel, using the current version of the classification prompts, Agent A and Agent B perform a single-round automatic classification on the same batch. Model outputs are then aligned sentence by sentence with human annotations, and Cohen’s Kappa is computed to measure human–model agreement.

Based on the observed agreement levels and error patterns in each round, prompts are revised and refined before entering the next iteration. Calibration is conducted independently for the emotion and moral tasks. The process terminates when the average Cohen’s Kappa across three consecutive iterations exceeds 0.6, indicating substantial consistency between system and human annotations.

Final agreement scores reach 0.61 for the emotion task and 0.67 for the moral task. The corresponding prompts are then fixed as the unified configuration for subsequent large-scale automatic annotation. For transparency and reproducibility, we report the complete system and task prompts for both emotion and moral annotation below.

Table 7: Emotion annotation prompts used in the A–B–A framework

Prompt Type	Content
Agent A1 (task)	<p><b>Task Description</b> Read each sentence carefully and identify the single most dominant emotion expressed. <b>Emotion definitions and label examples are provided in Table 8.</b></p> <p><b>Decision Rules</b> – Select exactly one emotion label. – Output a confidence score in <math>[0, 1]</math>, rounded to two decimals. – Provide a brief reason (1–3 sentences) grounded in explicit textual cues.</p> <p><b>Output Example (TSV)</b> sentence_id&lt;TAB&gt;emotion&lt;TAB&gt;confidence&lt;TAB&gt;reason S001&lt;TAB&gt;anger&lt;TAB&gt;0.80&lt;TAB&gt;The sentence contains explicit blame and irritation cues, indicating anger.</p>
Agent B (critic)	<p><b>Task Description</b> Review Agent A1’s initial judgment and assess whether it is reasonable. Identify possible confusion and recommend a more appropriate emotion if needed. <b>Emotion definitions and label examples are provided in Table 8.</b></p> <p><b>Decision Rules</b> – Output y if you agree with Agent A1, otherwise n. – Provide a short critique (1–3 sentences). – Recommend exactly one emotion label.</p> <p><b>Output Example (TSV)</b> sentence_id&lt;TAB&gt;agree&lt;TAB&gt;critique&lt;TAB&gt;suggested_emotion S001&lt;TAB&gt;n&lt;TAB&gt;The sentence emphasizes uncertainty and threat rather than confrontation.&lt;TAB&gt;fear</p>
Agent A2 (final decision)	<p><b>Decision Task</b> Determine the final emotion label after considering the critic’s feedback. <b>Emotion definitions and label examples are provided in Table 8.</b></p> <p><b>Decision Rules</b> – Output exactly one final emotion label. – Output a final confidence score in <math>[0, 1]</math>, rounded to two decimals. – Explain in 1–3 sentences why the initial judgment is kept or revised.</p> <p><b>Output Example (TSV)</b> sentence_id&lt;TAB&gt;final_emotion&lt;TAB&gt;final_confidence&lt;TAB&gt;final_reason S001&lt;TAB&gt;fear&lt;TAB&gt;0.74&lt;TAB&gt;The sentence highlights uncertainty and perceived threat, aligning more with fear.</p>

Table 8: Emotion label definitions and examples (Plutchik-8)

Label	Definition / Typical Cues
anger	<p>Core meaning: Anger is an antagonistic emotional state that arises when an individual perceives that their goals are being obstructed, their rights are being violated, or they are being treated unfairly. Its core lies in blame toward an external target, hostility, or an impulse to retaliate. This emotion is typically accompanied by a clear attribution of responsibility, emphasizing conflict, opposition, and a motivation to change the status quo. It may manifest as overt indignation or exist in an implicit form such as dissatisfaction or suppression.</p> <p>Typical cues: getting angry, anger, complaining, accusing, dissatisfaction.....</p>
anticipation	<p>Core meaning: Anticipation is a future-oriented emotional state, reflecting an individual's psychological preparation and anticipatory attention toward upcoming events, potential outcomes, or consequences of actions. Its core is not an immediate experience of pleasure or pain, but a cognitive orientation toward future possibilities and motivational activation. It is often associated with planning, hope, vigilance, or moral orientation, and plays an important role in learning, growth, and the formation of goals.</p> <p>Typical cues: anticipation, looking forward to, hope, preparation, planning.....</p>
disgust	<p>Core meaning: Disgust is an emotional reaction that stems from a strong sense of repulsion toward certain objects, behaviors, or situations. Its function is to prompt individuals to distance themselves from things perceived as harmful, unclean, or violating moral norms. This emotion can be grounded in physiological discomfort, and it can also extend to social and moral aversion and contempt, emphasizing a tendency to reject and sever contact.</p> <p>Typical cues: nausea, aversion, dislike, being despised.....</p>
fear	<p>Core meaning: Fear is an emotional response that arises when an individual perceives threats, uncertainty, or potential danger. Its core lies in the anticipation of negative consequences and the resulting defensive psychological state. This emotion is typically accompanied by tension, anxiety, and vigilance, emphasizing a motivation to avoid risk or reduce losses.</p> <p>Typical cues: fear, worry, tension, unease.....</p>
joy	<p>Core meaning: Joy is an emotional state triggered by satisfaction, success, or positive experiences, reflecting an individual's positive evaluation of the current situation and pleasant feelings. Its core lies in the immediate emotional experience itself, often expressed as relaxation, delight, or happiness, and accompanied by acceptance and enjoyment of the current state, rather than being primarily characterized by future expectations or risk assessment.</p> <p>Typical cues: happiness, joy, being glad, satisfaction, excitement, enjoyment, warmth.....</p>
sadness	<p>Core meaning: Sadness is a low emotional state that arises when an individual experiences loss, failure, setbacks, or powerlessness to change reality. Its core lies in an affective response to negative outcomes and inward-oriented regulation.</p> <p>Typical cues: sadness, sorrow, loss, depression, regret.....</p>
surprise	<p>Core meaning: Surprise is an instantaneous emotional response triggered by unexpected events or information. Its core lies in a rapid cognitive reaction to broken expectations. This emotion itself does not have a stable positive or negative affective polarity; rather, as a transitional state, it prompts individuals to quickly adjust their cognitive framework to cope with a new situation.</p> <p>Typical cues: surprise, didn't expect, actually, unexpectedly.....</p>
trust	<p>Core meaning: Trust is an emotional experience based on the perceived reliability and safety of others, relationships, or the environment. Its core lies in reducing defensiveness, accepting dependence, and being willing to bear uncertainty in interactions. This emotion emphasizes stability and continuity, often accompanied by feelings of reassurance, identification, and reliance, and plays a role in social relationships and cooperative behavior.</p> <p>Typical cues: believe, feel assured, trust, reassurance.....</p>

*(Continued on next page)*

Label	Definition / Typical Cues
<b>[Label examples]</b>	<p>“My mother also understood how I felt at this moment, so she quietly walked out. I watched my mother’s departing back, and I suddenly felt a trace of guilt.”: sadness</p> <p>“I still remember before—its leaves were almost withered, but today it actually bloomed. How amazing.”: surprise</p> <p>“Dark clouds shrouded my heart; the light seemed out of reach.”: fear</p> <p>“This kind of spirit is worth learning from.”: anticipation</p> <p>“This time my essay score actually reached as high as ninety-three!”: surprise</p> <p>“It felt as if I was pulled back into the long river of memories; when I was 7 years old, a warm recollection surged in my heart.”: joy</p> <p>“From fairy tales to novels, the types of books have changed, but the feelings and inspiration they bring me have never changed.”: anticipation</p> <p>“By reading <i>Special Forces School</i>, I not only experienced the excitement and thrills of adventure, but also learned the importance of courage and wisdom when facing difficulties.”: anticipation</p> <p>“We praised her for a long time.”: trust</p> <p>“This is truly awful!”: disgust</p>

Table 9: Moral annotation prompts used in the A–B–A framework

Prompt Type	Content
Agent A1 (task)	<p><b>Task Description</b> Read each sentence carefully and identify the single most salient moral expressed in the text. <b>moral definitions, neutral-sentence rules, and reference examples are provided in Table 10.</b></p> <p><b>Decision Rules</b> – Select exactly one moral label. – Output a confidence score in <math>[0, 1]</math>, rounded to two decimals. – Provide a brief reason (1–3 sentences) grounded in explicit textual evidence.</p> <p><b>Output Example (TSV)</b> sentence_id&lt;TAB&gt;moral&lt;TAB&gt;confidence&lt;TAB&gt;reason S001&lt;TAB&gt;care_pos&lt;TAB&gt;0.83&lt;TAB&gt;The sentence explicitly evaluates helping others as morally appropriate, reflecting a care-oriented moral.</p>
Agent B (critic)	<p><b>Task Description</b> Review Agent A1’s moral judgment and assess whether it is appropriate given the sentence focus. Identify potential confusion and recommend a more suitable moral if necessary. <b>moral definitions, decision standards, and examples are provided in Table 10.</b></p> <p><b>Decision Rules</b> – Output y if you agree with Agent A1, otherwise n. – Provide a short critique (1–3 sentences). – Recommend exactly one moral label.</p> <p><b>Output Example (TSV)</b> sentence_id&lt;TAB&gt;agree&lt;TAB&gt;critique&lt;TAB&gt;suggested_moral S001&lt;TAB&gt;n&lt;TAB&gt;The sentence emphasizes responsibility and rule adherence more than concern for others.&lt;TAB&gt;authority_pos</p>
Agent A2 (final decision)	<p><b>Decision Task</b> Make the final moral assignment after considering the critic’s feedback. <b>moral definitions, neutral-sentence rules, and examples are provided in Table 10.</b></p> <p><b>Decision Rules</b> – Output exactly one final moral label. – Output a final confidence score in <math>[0, 1]</math>, rounded to two decimals. – Explain in 1–3 sentences why the initial judgment is retained or revised.</p> <p><b>Output Example (TSV)</b> sentence_id&lt;TAB&gt;final_moral&lt;TAB&gt;final_confidence&lt;TAB&gt;final_reason S001&lt;TAB&gt;authority_pos&lt;TAB&gt;0.78&lt;TAB&gt;After considering the critique, rule-following and responsibility are the dominant evaluative focus.</p>

Table 10: Moral label definitions and guidelines (MFT-10)

Label	Definition / Typical Scenarios / Targets / Expressions
Care	<p><b>Core definition:</b> Care refers to the significant emotional resonance and discomfort that arise when an individual observes others or animals in a state of suffering, together with a tendency to take action to reduce harm, provide help, and offer protection, thereby improving the target’s situation as much as possible.</p> <p><b>Typical triggering scenarios:</b> For example, when seeing someone who is ill but unattended, an elderly person with limited mobility, a neglected child, or a stray animal suffering from hunger and cold, people often naturally develop a judgmental tendency such as “Can I offer help?” or “How can their suffering be reduced?”</p> <p><b>Typical targets:</b> The targets of Care include not only humans but also animals; it may also be directed at “states requiring care” themselves, such as injury, hunger, cold, loneliness, and helplessness.</p> <p><b>Common expressions:</b> These include “That’s so pitiful,” “They shouldn’t have to endure this,” and “We should take care of them,” often followed by concrete supportive actions such as delivering medicine, providing companionship, offering food, reducing waste of resources, reducing environmental damage, or arranging care.</p>
Harm	<p><b>Core definition:</b> Harm refers to inflicting physical or psychological damage on others (not limited to humans) despite knowing that the target may experience pain or discomfort, or showing indifference and disregard toward their suffering.</p> <p><b>Typical triggering scenarios:</b> For example, someone deliberately humiliates others with harsh language, mocks others in public to create embarrassment, deliberately frightens someone for amusement despite knowing their fear, abuses animals, or damages the environment.</p> <p><b>Typical targets:</b> The targets of Harm may be individuals or animals, as well as the dignity and psychological feelings of others (e.g., causing someone to lose face, feel fear, or feel wronged can also constitute Harm).</p> <p><b>Common expressions:</b> These include “It’s his fault for being so stupid,” “I just want to make him feel bad,” and “Let him cry, I don’t care,” or continuing to use harmful language under the pretext of joking, such as “I’m just kidding—why are you so sensitive,” as well as expressions involving “damaging the environment.”</p>
Fairness	<p><b>Core definition:</b> Fairness refers to applying the same standards to different targets under the same circumstances, without granting differential treatment based on identity differences.</p> <p><b>Typical triggering scenarios:</b> For example, cutting in line while shopping is considered inappropriate because first-come-first-served should be followed; similarly, in competitions, awards, and admissions, relevant rules should apply equally to all participants.</p> <p><b>Typical targets:</b> The targets of Fairness may include individuals and groups, as well as the rules and procedures themselves, such as queueing rules, competition rules, selection procedures, and allocation mechanisms.</p> <p><b>Common expressions:</b> These include “Why does he get to cut in line,” “Everyone should follow the rules,” and “We should treat everyone equally,” emphasizing that “the same situations should be handled in the same way.”</p>
Cheating	<p><b>Core definition:</b> Cheating refers to circumventing rules through covert means in the presence of explicit regulations in order to obtain unjust benefits, while attempting to appear as though nothing has happened.</p> <p><b>Typical triggering scenarios:</b> For example, peeking at answers or using cheat sheets during an exam; or deliberately tampering with an inaccurate scale or concealing expired goods in a transaction to avoid detection by others.</p> <p><b>Typical targets:</b> The targets of Cheating may include deceived individuals, as well as the exploited rules and transactional systems themselves, such as exam regulations, commercial integrity, or contractual agreements.</p> <p><b>Common expressions:</b> These include “No one will find out anyway,” “If I don’t do this, I’ll lose out,” and “Everyone does it,” with the core characteristics being “covert rule evasion, concern about being discovered, and gaining advantage through such actions.”</p>

(Continued on next page)

Label	Definition / Typical Scenarios / Targets / Expressions
Loyalty	<p><b>Core definition:</b> Loyalty refers to treating “one’s own side” as in-group members and partners, being willing to stand by them and jointly bear responsibilities and risks, rather than easily abandoning or withdrawing.</p> <p><b>Typical triggering scenarios:</b> For example, being willing to support and share burdens when a family encounters hardship; or not leaving immediately when a team faces difficulties, but instead making efforts to overcome them together.</p> <p><b>Typical targets:</b> The targets of Loyalty may include family, friends, groups, or organizations, as well as “shared reputation and commitments,” such as “what we promised to complete together” or “this family/team must not fall apart.”</p> <p><b>Common expressions:</b> These include “Our own people should look out for each other,” “Since we agreed to do it together, we should do it together,” and “You can’t quit at the critical moment,” emphasizing “maintaining the in-group and not leaving lightly.”</p>
Betrayal	<p><b>Core definition:</b> Betrayal refers to an individual, as a member of an in-group, turning to sell out the interests of that group, or using resources, information, or responsibilities entrusted by the group to harm it.</p> <p><b>Typical triggering scenarios:</b> For example, revealing family or team secrets to outsiders and causing harm to the group; or accepting benefits from others and then sabotaging one’s own group in critical matters.</p> <p><b>Typical targets:</b> The targets of Betrayal primarily include the side to which the individual belongs (such as family, team, or organization), as well as “the trust established by the group” and “the secrets and responsibilities entrusted to the individual.”</p> <p><b>Common expressions:</b> These include “selling out one’s own people,” “betraying your own at the critical moment,” and “you dragged all of us down,” emphasizing “exploiting internal trust to harm one’s own group.”</p>
Authority	<p><b>Core definition:</b> Authority refers to recognizing and following established rules and role-based orders, believing that family rules, school rules, or organizational regulations help maintain order and ensure the stability of collective life and operation, and therefore should be respected and implemented.</p> <p><b>Typical triggering scenarios:</b> For example, in classroom settings, students should listen to teachers rather than interrupting at will; in organizational contexts, members should act according to division of labor and procedures to avoid operational disorder.</p> <p><b>Typical targets:</b> The targets of Authority may include specific individuals (such as parents, teachers, and leaders), as well as roles and rules themselves (such as family rules, school rules, work processes, or norms of etiquette). In addition to institutional regulations, this also includes “executable methodological norms/procedural requirements” (such as writing conventions, learning methods, operational steps, and classroom requirements). The core lies in being “learnable, executable, and usable to constrain behavior or improve order and effectiveness.”</p> <p><b>Common expressions:</b> These include “You should listen to the teacher,” “Follow the rules,” and “Everyone has their own role,” emphasizing “observing rules, respecting roles, and avoiding disruption of order.”</p>
Subversion	<p><b>Core definition:</b> Subversion refers to the belief that certain rules, role arrangements, or authority structures themselves lack legitimacy, leading individuals to explicitly refuse compliance with related demands and potentially advocate overturning, modifying, or restructuring such arrangements; the key characteristic lies in denying the premise that “authority ought to be obeyed.”</p> <p><b>Typical triggering scenarios:</b> For example, a long-standing institutional regulation is considered unreasonable or discriminatory, and someone publicly states their refusal to accept its constraints, declines to act according to its requirements, and may call on or organize others to jointly oppose, resist, or change the rule.</p> <p><b>Typical targets:</b> The targets of Subversion mainly include authority structures themselves, including institutional rules, role divisions, hierarchical arrangements, and their underlying legitimacy; the focus is not on a specific enforcer, but on “whether this arrangement should be obeyed at all.”</p> <p><b>Common expressions:</b> These include “I don’t recognize the legitimacy of this set of rules,” “I refuse to follow this arrangement,” “This authority itself should not be obeyed,” and “We can’t keep doing things your way,” emphasizing the denial of authority legitimacy and the resulting refusal to comply or acts of challenge.</p>

(Continued on next page)

Label	Definition / Typical Scenarios / Targets / Expressions
Sanctity	<p><b>Core definition:</b> Sanctity/Purity refers to viewing certain things, morals, or qualities as possessing elevated nobility and purity that transcend everyday utilitarian significance, believing they should be cherished, protected, and kept from contamination, and therefore showing strong moral aversion to trivialization, profanation, or secularization.</p> <p><b>Typical triggering scenarios:</b> For example, maintaining solemnity in religious rituals or commemorative occasions; or when discussing ideals, art, beliefs, personal integrity, or personal aspirations, emphasizing that they should remain pure, pursue higher standards, and not be reduced to utilitarian calculation, transaction, or ridicule.</p> <p><b>Typical targets:</b> The targets of Sanctity/Purity may include rituals, beliefs, and commemorative activities, as well as morals and qualities endowed with elevated meaning, such as ideals, beliefs, personal dignity, artistic pursuits, moral aspirations, life goals, spiritual pursuits, and the boundaries they symbolize.</p> <p><b>Common expressions:</b> These include “Show respect,” “It must not be profaned,” “This is a bottom line,” “It should not be traded or mocked,” “Pursue beauty,” and “Pursue dreams,” emphasizing reverence for elevated morals and the protection and maintenance of boundaries.</p>
Degradation	<p><b>Core definition:</b> Degradation/Impurity refers to deliberately handling certain things, morals, or qualities in a vulgarized, utilitarianized, joking, or transactional manner, despite knowing that they are widely regarded as requiring nobility and bottom-line significance, thereby dissolving their proper solemnity and spiritual elevation.</p> <p><b>Typical triggering scenarios:</b> For example, engaging in vulgar performances in front of solemn memorials; or treating ideals, beliefs, personal integrity, or artistic pursuits as tools for profit, jokes, or cheap performance materials, passively devaluing one’s life meaning and rejecting progress, thus weakening their intended seriousness.</p> <p><b>Typical targets:</b> The targets of Degradation/Impurity are typically those things, morals, or qualities that should be respected and kept noble, including rituals, symbols, ideals, beliefs, personal bottom lines, and the spiritual boundaries they represent.</p> <p><b>Common expressions:</b> These include disrespecting heroes, “It’s all just a show,” “Don’t take it so seriously,” and “It’s just a gimmick,” emphasizing deliberately lowering nobility and denying that such things deserve reverence.</p>

*(Continued on next page)*

Label	Definition / Typical Scenarios / Targets / Expressions
-------	--

**[Criteria for classifying sentences with no explicit moral polarity]**

**1. Default rules for neutral statements**

When a sentence does not directly express a moral judgment, it may be pre-classified according to its primary focus:

- If it focuses on vulnerability or harmed states (e.g., illness, loneliness), pre-classify as Care/Harm;
- If it focuses on allocation ratios, procedures, or transparency, pre-classify as Fairness/Cheating;
- If it focuses on in-group versus out-group relations, belonging, or reputation, pre-classify as Loyalty/Betrayal;
- If it focuses on tradition, roles, or stability of order, pre-classify as Authority/Subversion;
- If it focuses on cleanliness, solemnity, or boundaries against vulgarity, pre-classify as Sanctity/Degradation.

**2. Priority principles when moral dimensions overlap**

When a sentence involves multiple moral dimensions, classification should be based on the “explicit focus of the text,” prioritizing the object that is most directly evaluated, required, or protected in the sentence. When necessary, adjudication should rely only on observable signals in the sentence (normative words, role/rule nouns, victims/beneficiaries, group boundaries, bottom-line/profanation terms):

- Care/Harm has the highest priority: if the direct consequence of an action is significant suffering of a specific individual or the avoidance of such suffering, then even when rule conflicts are involved, Care/Harm should be prioritized;
- Sanctity/Degradation is used to judge bottom lines of personality and dignity: if the text involves trampling on such bottom lines, priority should be given to the Sanctity dimension;
- Fairness has priority over Authority: if the emphasis of a rule lies in equal treatment rather than obedience, it should be labeled as Fairness;
- Loyalty, when the motivation purely points to “protecting our own side,” has priority over other dimensions, even if the behavior violates external rules.

**3. Rules for assigning moral polarity**

- Attention should be paid to pragmatic stance and rhetorical devices. When encountering rhetorical questions, irony, sarcasm, exaggerated negation, or derogatory metaphors, the moral orientation of the sentence should be carefully analyzed;
- If a sentence primarily expresses dissatisfaction or negation toward a behavior, phenomenon, or state through criticism, denial, ridicule, or accusation, it should be labeled with the corresponding opposite moral label;
- When a sentence directly expresses support, endorsement, advocacy, or calls for protecting a certain moral dimension or behavior, it should be labeled with the corresponding positive moral label.

**[Reference examples]**

“Do not bring dogs into public places.”: Authority (Note: this leans toward a recommendation and also constitutes a social rule.)

“I walked up to him and sincerely said, ‘I’m sorry, it was my fault. I shouldn’t have broken your toy, and I shouldn’t have argued with you just to save face.’ ”: Care (Note: the apology expresses concern for others’ states and emotions.)

“He never complained about hardship, but instead smiled and said, ‘While I can still run, I’ll run a few more deliveries.’ ”: Sanctity (Note: this tends toward perseverance and effort, falling under the Sanctity dimension.)

“The tree planter diligently cultivated the saplings, while the man with glasses brazenly sat under the tree waiting to enjoy the shade. This kind of gaining without contributing is truly disgraceful.”: Cheating (Note: the gains are disproportionate to the contributions.)

“The entire scene uses exaggeration to present the desolation of the environment and the helplessness of the woodpecker, metaphorically pointing to the heavy blow of ecological crisis to the survival of living beings.”: Care (Note: the exaggerated depiction reflects concern for nature.)

“The giant ants arrived at the construction site. The giant ants voluntarily took on the task of carrying materials, and soon the building was completed.”: Loyalty (Note: working together in unity.)

“Under Teacher Lin’s guidance, we learned to feel art with our hearts and depict the world with colors.”: Sanctity (Note: the trigger lies in ‘feeling art with our hearts’ —the elevation of moral meaning—rather than ‘teacher guidance’ (Authority).)

“Caution is a member of our life; we must always stay alert!”: Authority (Note: the emphasis is on advocating caution and carries a prescriptive tone.)

“Maintain up to ten hours of sleep every day.”: Care (Note: the focus is on health and sleep.)

### A.3.2 Batch Processing and Parallel Configuration

Full annotation is executed via batch processing with multi-threaded parallelism to improve throughput. We set the batch size to 30 sentences and run 5 parallel threads at the batch level. Each model call is allowed up to three retries. To reduce the risk of data loss caused by long-running task interruptions, outputs are appended to the streaming result file immediately after each batch returns successfully and passes validation. We also maintain runtime logs that record format-validation failures, invocation exceptions, retry-limit events, and re-run statistics, enabling traceability and auditability.

### A.3.3 Strict Output Constraints and Validation

To ensure that batch outputs are automatically parsable and strictly aligned with inputs, we impose strict output constraints and validate them automatically. At the sentence level, model outputs must be in TSV (tab-separated values) format and satisfy the following conditions: (i) the number of output rows matches the number of input sentences; (ii) each row contains a fixed number of fields consistent with the predefined schema; (iii) the first column (`sentence_id`) exactly matches the input batch’s `sentence_id` order row by row; (iv) output labels belong to a predefined closed set; (v) confidence scores fall within  $[0, 1]$  and are rounded to two decimal places. If any condition is violated, the batch output is deemed invalid and triggers retries and subsequent re-run procedures (Section A.3.4).

### A.3.4 Result Completeness Control: Set-Difference Detection and Iterative Re-Runs Until Zero Missing

After the full run completes, we execute a completeness-control procedure to ensure that all sentences requiring annotation receive valid outputs. Using (`writing_id`, `sentence_id`) as a composite primary key, we define the target set as all sentence keys in the filter-pass file, and the done set as the keys already written to the streaming output file. The missing set is computed via set difference.

Re-runs are conducted iteratively. In each round, sentences in the current missing set are reprocessed using the A–B–A framework. Successfully returned and validated outputs are appended to the main output file immediately, and the done set is

updated, after which the missing set is recomputed. This loop continues until `missing = 0`.

To improve re-run success rates and reduce format failures caused by output truncation, we apply an automatic batch-splitting strategy during re-runs. If a batch fails validation after exhausting retries, it is recursively split into smaller sub-batches down to the single-sentence level. If a single sentence still fails to produce valid output, it is logged as a final failure and skipped in subsequent rounds, maximizing sentence-level completeness and reproducibility under the defined constraints.

Consistency analysis of the full annotation results shows that, for both emotion and moral tasks, Fleiss’  $\kappa$  among Agent A’s initial labels, Agent B’s suggested labels, and Agent A’s final labels reaches 0.9814 and 0.9895, respectively, indicating extremely high agreement.

### A.4 Writing-Level Aggregation and Entropy Computation

After obtaining sentence-level final outputs, we aggregate labels to the writing level. The numerator consists of sentences that pass relevance filtering and receive valid final labels, and the denominator is the total number of sentences under the same `writing_id`. For the emotion task, we count sentences assigned to each of the eight emotion categories; for the moral task, we count sentences assigned to each of the ten moral categories.

We then compute category proportion vectors by dividing category counts by the total number of sentences. Vector dimensionality and ordering are fixed and stored as vector fields. Finally, we compute Shannon entropy (base-2) over the emotion or moral distribution for each essay as an indicator of distributional diversity.

## B Dimensional Classification Logic

The classification maps specific pedagogical directions to six linguistic dimensions based on the following dictionary-based heuristic:

**1. Extraction and Mapping** Feedback is segmented into discrete suggestions. Each segment’s pedagogical direction is mapped to a dimension using the exact logic defined in Table 11.

### 2. Processing Logic

- **Direction Identification:** The system prioritizes leading tags (e.g., “Language:”) but

Table 11: Exact Mapping from Pedagogical Directions to Dimensions

Linguistic Dimension	Mapped Pedagogical Directions (Keywords)
Lexical Richness	Language, Description, Details
Syntactic Diversity	Expression, Transition
Semantic Dispersion	Content, Structure, Persona
Semantic Shift	Plot, Opening, Ending
Emotion Spectrum	Psychology, Emotion, Environment, Action
Moral Alignment	Morals, Material Selection, Title, Off-topic

scans the full segment if no explicit tag is present.

- **Multi-Labeling:** Suggestions containing multiple directions are assigned up to two distinct dimensions based on the first two keywords identified.
- **Model Exclusion:** Suggestions that cannot be mapped to these specific categories (labeled as “Other”) contribute to the total uptake count (*FUA*) but are excluded from the dimensional regression analysis to ensure thematic consistency.