

Unsupervised Text Style Transfer for Controllable Intensity

Shuhuan Gu¹, Wenbiao Tao¹, Xinchen Ma¹, Kangkang He², Ye Guo²,
Xiang Li¹, Yunshi Lan^{1*}

¹East China Normal University, Shanghai, China

²Baowu Group, Shanghai, China

treeegu2003@gmail.com

yslan@dase.ecnu.edu.cn

Abstract

Unsupervised Text Style Transfer (UTST) aims to build a system to transfer the stylistic properties of a given text without parallel text pairs. Compared with text transfer between style polarities, UTST for controllable intensity is more challenging due to the subtle differences in stylistic features across different intensity levels. Faced with the challenges posed by the lack of parallel data and the indistinguishability between adjacent intensity levels, we propose a SFT-then-PPO paradigm to fine-tune an LLM. We first fine-tune the LLM with synthesized parallel data. Then, we further train the LLM with PPO, where the rewards are elaborately designed for distinguishing the stylistic intensity in hierarchical levels. Both the global and local stylistic features are considered to formulate the reward functions. The experiments on two UTST benchmarks showcase that both rewards have their advantages and applying them to LLM fine-tuning can effectively improve the performance of an LLM backbone based on various evaluation metrics. Even for adjacent levels of intensity, we can still observe a noticeable stylistic difference among the generated text across these levels¹.

1 Introduction

Text Style Transfer (TST) aims to automatically modify the stylistic properties of a given text, such as readability (North et al., 2023), sentiment (Zhang et al., 2015), or formality (Briakou et al., 2021), while preserving its original semantic content. This task, translating the original text into the desired style, has a wide range of applications in education (Tan et al., 2024), legal (Li et al., 2021), and social media (Prabhumoye et al., 2018) fields. Besides the transfer between style polarity, recent trends start to evaluate the style transfer on

* Corresponding author

¹The code is available at <https://github.com/TreeGu2003/UTST-CI>.

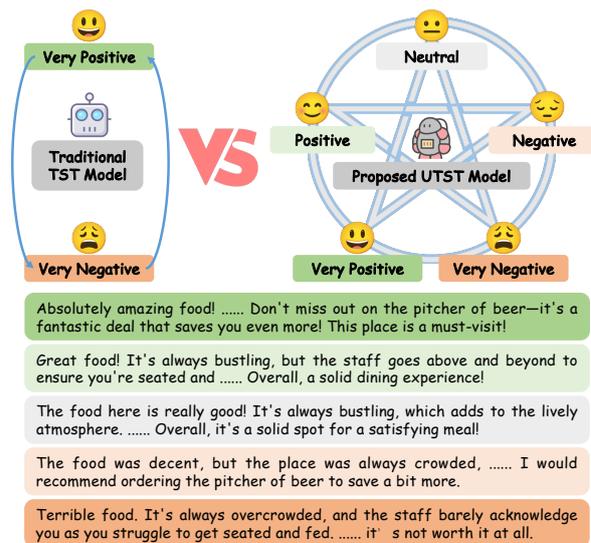


Figure 1: A comparison between the text transfer for style polarity and across levels of intensity.

a scale of intensity, which is more challenging due to the subtle difference of stylistic attributions and features for different levels of intensity (Mir et al., 2019; Mukherjee et al., 2025). Moreover, due to the lack of large high-quality parallel corpora, it brings in new difficulties in text style transfer without supervision, namely Unsupervised Text Style Transfer (UTST). We display an example of UTST for controllable intensity in Figure 1.

Noticeable efforts have been made in existing studies. We classify them into two categories: (1) *Attribute substitution*. These methods replace the style-specific phrases in the text with the ones in the target style. Hence, extra corpora (Li et al., 2018; Sudhakar et al., 2019), classifier model (Lee, 2020; Fu et al., 2018) and mapping model (Luo et al., 2019b) are introduced into the transfer model. However, due to the direct substitution of attributes, the transferred sentences may suffer from semantic incoherence. (2) *Sentence generation*. Thanks to the emergence of LLMs, text styles can be trans-

ferred token by token. Prompt engineering (Reif et al., 2022; Suzgun et al., 2022), instruction fine-tuning (Mukherjee and Dušek, 2023; Zong et al., 2024) and reinforcement learning (RL) (Deng et al., 2022; Liu et al., 2024) are applied so that language models learn to generate target style text. Nevertheless, faced with the need for controllable intensity generation, these methods frequently suffer from style deficiencies.

The controllable intensity of text transformation forces a model to fully understand different levels of stylistic features and conduct style transfer. In summary, we need to focus on two challenges:

- **(C1) Absence of the high-quality TST parallel data.** Paired text style transfer data allows the model to learn the precise transfer pattern. While in the absence of high-quality parallel data, the knowledge source of the model is limited, which impairs the TST for controllable intensity.
- **(C2) Transfer across indistinguishable intensities.** Widely publicized bipolar style transformations have distinctive features for identification. However, there are more subtle differences between texts of stepped intensities, as shown in Figure 1, which clearly demonstrates the UTST task.

To address these challenges, we propose a SFT-then-PPO paradigm for UTST. Our approach begins by constructing a pseudo parallel corpus using an LLM to rewrite source texts into target styles at varying intensity levels. The generated outputs are then filtered using a pretrained style classifier to ensure alignment with the desired stylistic strength, resulting in high-quality pseudo-parallel corpora. Based on the synthesized data, we fine-tune an existing LLM to establish foundational TST capabilities. To further improve stylistic consistency and enable fine-grained control beyond supervised learning, we introduce proximal policy optimization with hierarchical rewards. Specifically, we design sentence-level rewards to capture global stylistic alignment, and lexicon-level rewards to guide fine-grained lexical choices. This dual-level reward structure allows the model to produce outputs with more accurate and consistent stylistic intensity.

To sum up, our contributions are as follows:

- We introduce an intensity-controlled UTST paradigm that first fine-tunes a model for baseline style transfer via supervised fine-tuning

and then refines it with reinforcement learning to further handle varying style intensities.

- We construct pseudo parallel corpus to solve the problem of scarcity of parallel data for TST model fine-tuning, as well as design sentence-level, lexicon-level and semantic-level rewards to allow the model to adapt to UTST for controllable intensity from multiple granularities.
- We construct two benchmarks for our task. The experimental results on readability and sentiment style transfer demonstrate that our method can effectively generate text with noticeable stylistic difference even for close levels of intensity.

2 Related Work

Text Style Transfer. Text Style Transfer modifies text style while preserving meaning, typically via attribute substitution or sentence generation.

Attribute substitution methods replace style-specific phrases using external corpora (Li et al., 2018; Sudhakar et al., 2019), classifiers (Lee, 2020; Fu et al., 2018) or mapping models (Luo et al., 2019b). While effective for coarse-grained polarity transfer, these methods often suffer from semantic incoherence and cannot handle fine-grained intensity—directly replacing phrases (e.g., ‘good’ → ‘bad’) fails to capture intensity differences (e.g., ‘great’ vs. ‘amazing’) (Mir et al., 2019).

Sentence generation methods based on LLMs are classified into non-disentangled and disentangled approaches. Non-disentangled methods (e.g., prompt engineering (Reif et al., 2022; Suzgun et al., 2022), instruction fine-tuning (Mukherjee and Dušek, 2023; Zong et al., 2024)) adopt intertwined style-content representations, which simplifies the training process yet lacks the capability to regulate subtle stylistic intensity. RL extensions (Deng et al., 2022; Liu et al., 2024) only optimize broad style alignment rather than nuanced gradations of style. Disentangled methods (Han et al., 2024) decouple style and content through auxiliary tasks to achieve precise control, albeit requiring extensive manual annotation.

Our paradigm follows the line of non-disentangled methods, enabling robust UTST by generating pseudo-parallel data with LLMs and applying hierarchical PPO-based RL for fine-grained style control, without needing parallel data.

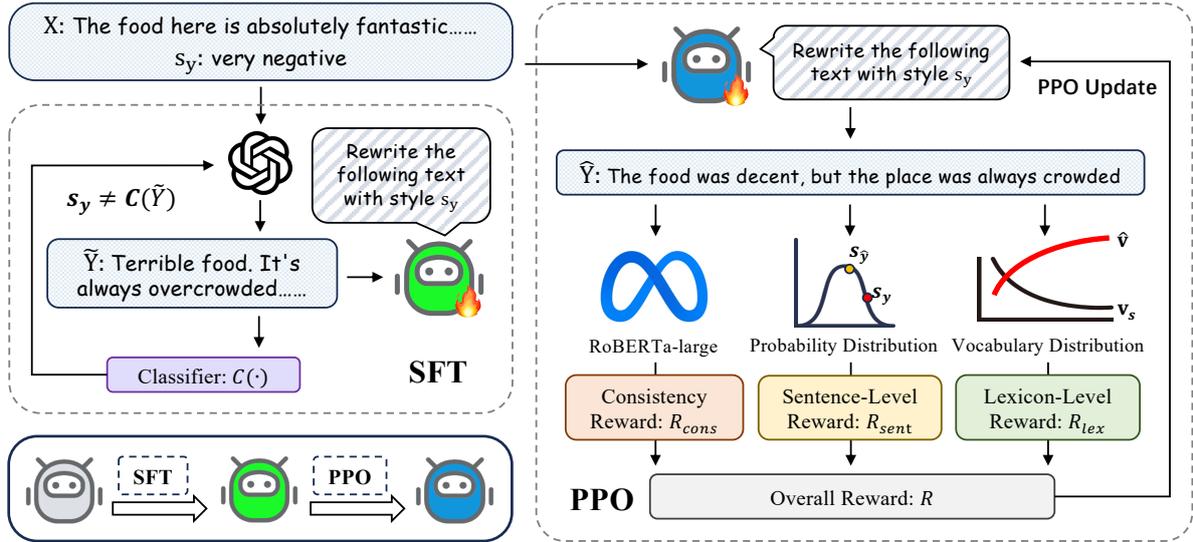


Figure 2: UTST with SFT-then-PPO paradigm.

Controllable Text Generation. Controllable text generation (CTG) produces text with desired attributes. Transformer-based LLMs have advanced CTG through prompt-based methods (Suzgun et al., 2022), allowing flexible control via natural language instructions. Fine-tuning (Mukherjee and Dušek, 2023) and reinforcement learning (RL) (Deng et al., 2022; Liu et al., 2024) further improve adaptability for tasks like sentiment modification (Zhang et al., 2015) and readability adjustment (North et al., 2023). However, CTG for UTST with controllable intensity remains difficult due to the lack of parallel data and the challenge of capturing nuanced stylistic differences (Mukherjee et al., 2025). Our approach fills these gaps by combining pseudo-parallel data generation with hierarchical rewards in a PPO method, enabling fine-grained style control and strong UTST performance without supervision.

3 Task Statement

UTST is an NLP task aiming to convert the style of the given text without parallel text pairs. Specifically, given k sets of non-parallel corpora, each of them can be represented as $\mathcal{D}^s = \{X_i\}_{i=1}^{|\mathcal{D}^s|}$. Here, s is annotated with a pre-defined stylistic intensity set $\mathcal{S} = \{s_1, s_2, \dots, s_k\}$, which is the categorical scale of the stylistic transfer. Given a text $X = (x_1, x_2, \dots, x_n)$ in the *source* stylistic intensity $s_x \in \mathcal{S}$, the key of UTST is to develop a model which is able to transfer it into the text $Y = (y_1, y_2, \dots, y_m)$ in the *target* stylistic intensity $s_y \in \mathcal{S}$ while preserving the original content.

Regarding UTST for controllable intensity, the generation process can be modeled as:

$$\hat{Y} \leftarrow \mathcal{G}(X, s_y),$$

where $\mathcal{G}(\cdot)$ is a UTST model and \hat{Y} should be the transferred text in s_y stylistic intensity.

4 Paradigm for UTST

To train such a model $\mathcal{G}(\cdot)$ without parallel corpus, we take full advantage of the LLMs for UTST with controllable intensity and introduce a SFT-then-PPO paradigm to train an LLM with surrogate rewards in hierarchical levels. Particularly, we first generate high-quality pseudo parallel corpus with an LLM-based pipeline. Then, we fine-tune an LLM as $\mathcal{G}(\cdot)$ with parallel stylistic text. Next, we further train the LLM with Proximal Policy Optimization (PPO), where the rewards are elaborately designed for distinguishing the fine-grained stylistic intensity. The overview of our method is illustrated in Figure 2. We will discuss the steps of obtaining $\mathcal{G}(\cdot)$ in detail.

4.1 SFT with Synthesized Parallel Text

SFT is critical for enhancing the text style transfer capabilities of LLMs. While LLMs exhibit impressive zero-shot and few-shot performance in TST tasks, their outputs often lack consistency in adhering to specific stylistic constraints, such as targeted readability levels or sentiment tones (Suzgun et al., 2022; Hu et al., 2020). SFT addresses this by accurately aligning the LLMs' outputs with the attributes of the target text, especially when the

style is annotated with the fine-grained intensity. To this end, we employ instruction fine-tuning (Zhang et al., 2023), where LLMs are fine-tuned on elaborately designed prompts that specify the attributes of the target style.

Since SFT requires a high-quality dataset of parallel text pairs as the strong supervision, where each pair consists of an input text and its stylistically transformed counterpart. Due to the absence of the parallel text pairs for UTST task, we propose to synthetically generate the data for fine-tuning with the following pipeline.

- For each source text X sampled from the corpus, we pair it with a variety of intensities in \mathcal{S} to form a set of inputs (X, s_y) , where $s_y \in \mathcal{S}$ is the target style that we would like to transfer X into and the transferred style $s_y \neq s_x$.
- Given an input, we utilize GPT-4o-mini (OpenAI et al., 2024) to transfer X into the text with the target stylistic intensity \tilde{Y}^2 .
- We train a classifier $\mathcal{C}(\cdot)$ using the non-parallel corpus as the judge to determine if the synthesized target text satisfies the requested intensity. If the synthesized text \tilde{Y} is predicted as the target intensity, we add the parallel data (X, s_y, \tilde{Y}) into the data set. Otherwise, we repeat the above step till the LLMs generate the desired text. In the case where no suitable output is generated after 10 iterations, we discard X .

As a result, we obtain a set of parallel data $\mathcal{D} = \{X_i, s_y, \tilde{Y}_i\}_{i=1}^N$ which contains aligned text pairs for style transfer. We consider an LLM as $\mathcal{G}(\cdot)$ and fine-tune it on \mathcal{D} with the following objective:

$$\mathcal{L}_{\text{SFT}} = - \sum_{i=1}^N \sum_{j=1}^k \log P(\tilde{Y}_i | X_i, s_j),$$

where the probability is modeled via the LLM. After fine-tuning, the model $\mathcal{G}(\cdot)$ is able to generate text \hat{Y} in the desired style.

4.2 Reward Function for PPO

While SFT effectively guides the model to capture stylistic features through strong supervision, it often falls short in ensuring consistent adherence to subtle or nuanced style attributes across diverse inputs. To address this, we introduce an RL-based

approach to refine the model’s text generation, inspired by prior RL methods (Gong et al., 2019; Luo et al., 2019a). Unlike prior RL-based methods that use single-scale rewards to judge polarity compliance (e.g., positive vs. negative), a contribution of our method lies in its hierarchical reward optimization quantifying intensity compliance (e.g., very positive vs. positive), which explicitly evaluates the global and local attributes of the stylistic text. We introduce two-level rewards: a sentence-level reward that measures the stylistic intensity of the generated text in two different formats, and a lexicon-level reward that aligns the generated text to the fine-grained linguistic features by analyzing lexical distributions. We integrate the above rewards into our final reward function.

4.2.1 Hierarchical Stylistic Rewards

Sentence-level Reward. The sentence-level reward, derived from classifier-based evaluations, captures global stylistic attributes. Since the stylistic intensity can be either formulated as the fine-scale classification (e.g., sentiment analysis (Zhang et al., 2015)) or regression (e.g., readability assessment (Feng et al., 2010)), we formulate the sentence-level rewards for both regression and classification.

- *Regression-based:* This type of reward is designed for the case that a regression model $\mathcal{C}(\cdot)$ is used to measure the intensity of the generated text \hat{Y} . A large value reward is provided when the generated text has the low deviation from the target intensity. Otherwise, the reward value is small. Inspired by existing study (Ribeiro et al., 2023), we model the reward for regression-based classifier using a normalized Gaussian distribution centered at the target intensity $f \sim \mathcal{N}(s_y, \sigma^2)$:

$$\hat{s} = \mathcal{C}(\hat{Y});$$

$$\mathcal{R}_{\text{sent}}(\hat{Y}) = f(\hat{s})/f(s_y).$$

The reward will assign a maximum value 1 if the observed intensity is equal to the target intensity and decrease exponentially as \hat{s} deviates from s_y .

- *Classification-based:* This type of reward is designed for the case that a classification model $\mathcal{C}(\cdot)$ is used to measure the intensity of the generated text. The reward is formulated as:

$$\mathcal{R}_{\text{sent}}(\hat{Y}) = \mathcal{C}(\hat{Y}),$$

²We display the prompts in Appendix A.

where $\mathcal{C}(\hat{Y})$ outputs the probability of the target intensity given the generated text. The reward will assign a maximum value 1 to the generated text if it is predicted to have the target style.

As we can see, both rewards encourage the model to generate text in the target stylistic intensity.

Lexicon-level Reward. The sentence-level reward focuses on overall characteristics while lexicon-level reward captures vocabulary intensity from a style-specific vocabulary and lexical distribution. Since stylistic distinctions often manifest through different word usage patterns (Mir et al., 2019), identifying and rewarding style-specific vocabulary can enhance the LLMs’ ability to generate text that conforms to the linguistic features of the target style.

To identify the style-specific vocabulary, we employ TF-IDF to measure the importance of a word to a text in a collection or corpus. As a result, for the UTST task, the style-specific words are usually those having high TF-IDF values. We select the words with top-1000 TF-IDF values per level of intensity to form the style-specific vocabulary. We display the representative words in a corpus for stylistic intensities in Appendix E. Then for each level of stylistic intensity, we concatenate all the text in the corpus and form a TF-IDF vector demonstrating a pivot of the lexical distribution. For each generated text, we also form a TF-IDF vector showing its lexical distribution:

$$\begin{aligned} \mathbf{v}_s &= \text{TF-IDF}(\mathcal{D}^s); \\ \hat{\mathbf{v}} &= \text{TF-IDF}(\hat{Y}), \end{aligned}$$

where \mathbf{v}_s indicates the pivot TF-IDF vector for the requested intensity s and $\hat{\mathbf{v}}$ means the TF-IDF vector for the generated text. By computing the similarity of the vectors of the generated text and the pivot vector, the lexicon-level reward can be measured as follows:

$$\mathcal{R}_{\text{lex}} = \frac{\exp(\text{sim}(\mathbf{v}, \mathbf{v}_{s_y})/T)}{\sum_{j=1}^k \exp(\text{sim}(\hat{\mathbf{v}}, \mathbf{v}_{s_j})/T)},$$

where \mathbf{v}_{s_y} is the pivot TF-IDF vector of the target intensity, $\text{sim}(\cdot, \cdot)$ denotes the cosine similarity function, and T is a temperature parameter that controls the sharpness of the distribution.

This reward encourages the model to emphasize vocabulary patterns that are characteristic of the target style, thereby improving stylistic fidelity.

4.2.2 Consistency Reward

To preserve the semantic information of the input during style transfer, we follow the existing study (Zhang* et al., 2020) proposing a consistency reward:

$$\mathcal{R}_{\text{cons}} = \frac{1}{|\hat{Y}|} \sum_{y_j \in \hat{Y}} \max_{x_i \in X} (\text{Emb}(y_j) \cdot \text{Emb}(x_i)),$$

where $|\hat{Y}|$ is the number of tokens in \hat{Y} , $\text{Emb}(\cdot)$ denotes the RoBERTa (Liu et al., 2019) embeddings of token. This score is normalized to $[0, 1]$, which indicates the semantic alignment between the source and target text.

4.2.3 Total Reward

The total reward is formulated as the weighted sum of the above rewards:

$$\mathcal{R} = \lambda_1 \cdot \mathcal{R}_{\text{sent}} + \lambda_2 \cdot \mathcal{R}_{\text{lex}} + \lambda_3 \cdot \mathcal{R}_{\text{cons}},$$

where λ_1 , λ_2 and λ_3 are three hyper-parameters to balance the rewards. Our final model is obtained via the fine-tuning in Section 4.1 and PPO optimization. As a result, our model $\mathcal{G}(\cdot)$ can perform text style transfer for controllable intensity without any parallel data.

5 Experimental Setup

5.1 Datasets

Differently from most previous work, which only focused on the text style transfer between style polarity (e.g., positive \rightarrow negative), we conduct experiments on two datasets³ of UTST for controllable intensity:

- **CNN/DM:** This is a dataset for readability transfer which is collected from CNN Daily mails (Hermann et al., 2015) containing the original text and its summaries with varying readability levels. To construct a dataset for UTST task, we calculate the Flesch Reading Ease (FRE) scores (Flesch, 1948) of the summaries and divide them into four levels of readability, namely $\{\textit{elementary school (1)} \rightarrow \textit{middle school (2)} \rightarrow \textit{high school (3)} \rightarrow \textit{college (4)}\}$. The numbers of text in these sets are 42824, 176262, 61099 and 6459, respectively.
- **YELP:** This is a dataset initially for sentiment analysis which is collected from Yelp reviews

³The details of the datasets are shown in the Appendix D

(Zhang et al., 2015). It contains the ratings ranging from 1 to 5 stars, which corresponds to $\{\text{very negative (1)} \rightarrow \text{negative (2)} \rightarrow \text{neutral (3)} \rightarrow \text{positive (4)} \rightarrow \text{very positive (5)}\}$. We consider it as the UTST dataset and partition it into five sets of non-parallel corpus for training. The number of texts in each set is 130k.

5.2 Evaluation Metrics

Following the standard evaluation of TST tasks, we employ metrics that assess style alignment and content preservation. Regarding style alignment, we use **FRE** to measure the readability intensity of the prediction and **FRE $_{\Delta}$** to denote the deviation of intensity between the prediction and the target. We use star rating (**STAR**) to measure the sentiment intensity of the prediction and **STAR $_{\Delta}$** to denote the deviation of intensity between the prediction and the target. We also employ the sum of the hierarchical reward **H-Re**, illustrated in Section 4.2 to measure the stylistic alignment of the test data. For content preservation, ROUGE-L (**RG-L**) is used to measure the overlap between the prediction and the input text. Detailed description can be found in Appendix C.

5.3 Comparable Methods

To evaluate the performance of our proposed UTST method with controllable intensity, we comprehensively compare a variety of our SFT-then-PPO paradigms for comparison: **SFT**: We employ a flan-t5-large model as the backbone LLM and fine-tune it via the SFT with the synthesized parallel text. **SFT + PPO ($\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{cons}}$)**: This is our method with SFT-then-PPO paradigm and we employ sentence-level reward as the style alignment judge. **SFT + PPO ($\mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$)**: This is our method with SFT-then-PPO paradigm and we employ lexicon-level reward as the style alignment judge. **SFT + PPO ($\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$)**: This is our method with both hierarchical rewards as the style alignment judge.

Due to the inapplicability of the prior TST methods for polarity transfer, we employ **GPT-4o-mini** (OpenAI et al., 2024), which is an advanced LLM, as the strong baseline for UTST tasks. We simply prompt it with instructions to generate text with varying style intensities under the zero-shot setting, as detailed in Appendix B.

5.4 Implementation Details

For YELP, we employ a DistilBERT⁴ fine-tuned on non-parallel corpus as the classification-based model $\mathcal{C}(\cdot)$. For CNN/DM, we employ FRE function, which is a regression model, as $\mathcal{C}(\cdot)$.

We adopt the TRLX library (Havrilla et al., 2023) with PPO for the RL phase. Reward weights are set to $\lambda_1 = 0.5$, $\lambda_2 = 0.3$, and $\lambda_3 = 0.2$ for readability and sentiment, explained in Appendix G.1, with a softmax temperature $T = 0.01$ for the lexicon-level reward. Detailed parameter settings can be seen in Table 7.

6 Results

6.1 Main Results

Readability Transfer. Table 1 compares the UTST paradigm with baselines on the CNN/DM dataset across four target readability levels. Without task-specific training, GPT-4o-mini underperforms in zero-shot settings, highlighting the limitations of prompt-based approaches—even with a strong LLM. In contrast, our UTST paradigm achieves significant gains, outperforming baselines by up to 64.05% in average FRE $_{\Delta}$ and 40.74% in H-Re, demonstrating the effectiveness of the SFT-then-PPO pipeline.

Sentiment Transfer. We further evaluate the UTST method on the YELP dataset for sentiment transfer, with results presented in Table 2. Consistent with the readability transfer task, UTST significantly outperforms GPT-4o-mini, achieving gains of 26.85% in STAR $_{\Delta}$ and 36.36% in H-Re, reinforcing its effectiveness in text style transfer.

Overall Analysis. **H-Re** serves as the primary metric for UTST tasks, as it comprehensively assesses performance at both the sentence and lexical levels. Overall, the SFT-then-PPO paradigm outperforms SFT alone (H-Re: 0.38 vs. 0.34 in Table 1; 0.45 vs. 0.43 in Table 2). SFT is only capable of learning to mimic the superficial stylistic characteristics of the dataset, and fails to proactively align with complex or continuous style control objectives. In contrast, PPO facilitates the direct optimization of non-differentiable style metrics via reward signals, thus improving control over fine-grained styles.

The optimal reward combination exhibits variation across different style types: for readability, $\mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$ is deemed optimal, as readability relies on lexical selections (captured by \mathcal{R}_{lex})

⁴<https://huggingface.co/distilbert/distilbert-base-uncased>

Readability (s_y)	FRE	FRE $_{\Delta}$ ↓	RG-L ↑	H-Re ↑
GPT-4o-mini				
Elementary school	70.96	19.04	54.12	0.24
Middle school	65.41	4.59	59.64	0.40
High school	60.73	10.73	62.54	0.30
College	47.42	27.42	60.92	0.14
Average	-	15.44	<u>59.31</u>	0.27
SFT				
Elementary school	73.97	16.03	61.63	0.29
Middle school	64.96	5.04	63.38	0.40
High school	51.89	1.89	61.91	0.38
College	32.27	12.27	55.22	0.29
Average	-	8.81	60.54	0.34
SFT + PPO ($\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{cons}}$)				
Elementary school	78.77	11.23	58.62	0.38
Middle school	65.25	4.75	63.11	0.40
High school	50.70	0.70	61.27	0.38
College	28.87	8.87	53.06	0.32
Average	-	6.39	59.02	<u>0.37</u>
SFT + PPO ($\mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$)				
Elementary school	79.96	10.04	56.73	0.40
Middle school	64.65	5.35	61.65	0.40
High school	48.92	1.08	59.50	0.38
College	25.71	5.71	49.99	0.34
Average	-	5.55	56.97	0.38
SFT + PPO ($\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$)				
Elementary school	78.87	11.13	58.12	0.38
Middle school	65.25	4.75	62.69	0.40
High school	50.89	0.89	60.90	0.38
College	26.67	6.67	51.49	0.33
Average	-	<u>5.86</u>	58.30	<u>0.37</u>

Table 1: The results of CNN/DM. Both breakdown results and average results are displayed.

and content preservation (captured by $\mathcal{R}_{\text{cons}}$)—sentential features are redundant here, given that word difficulty exhibits a strong correlation with sentence structure. For sentiment, $\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$ is considered optimal, as emotional meaning is context-dependent (e.g. 'busy' can convey positive or negative connotations), necessitating $\mathcal{R}_{\text{sent}}$ to ascertain the emotional polarity of words within their contextual framework. This adaptability highlights the generalizability of the proposed model across diverse style dimensions.

However, **RG-L** exhibits a slight reduction, attributed to the intrinsic calculation mechanism. A detailed explication is provided in Appendix F.

6.2 Further Analysis

6.2.1 Factual Consistency Analysis

The heatmap in Figure 3a reveals strong diagonal alignment in predictions, particularly for Very Positive (4743 cases) and Very Negative (3996 cases),

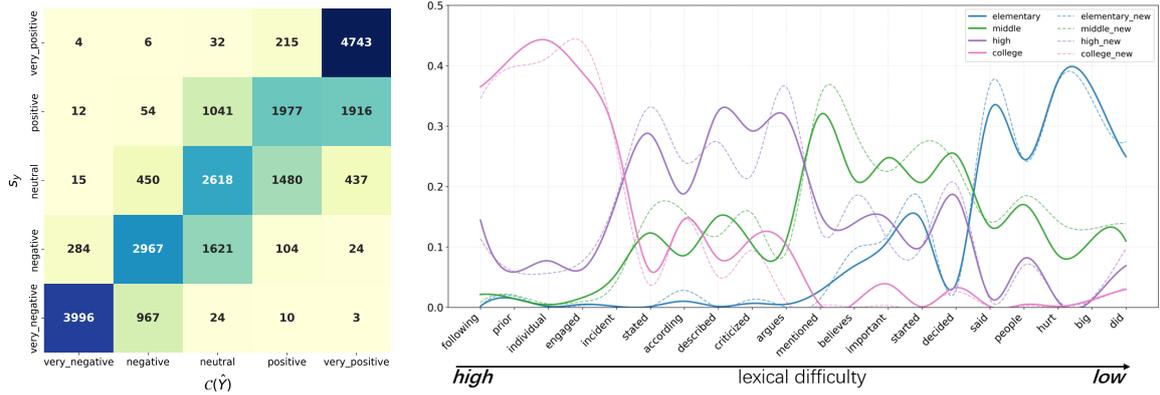
Sentiment (s_y)	STAR	STAR $_{\Delta}$ ↓	RG-L ↑	H-Re ↑
GPT-4o-mini				
Very Negative	1.5190	0.5190	45.39	0.35
Negative	1.8018	0.1982	47.11	0.38
Neutral	2.9476	0.0524	46.58	0.23
Positive	4.3652	0.3652	49.78	0.30
Very Positive	4.6032	0.3968	48.98	0.41
Average	-	0.3063	47.57	0.33
SFT				
Very Negative	1.2012	0.2012	43.68	0.50
Negative	2.3338	0.3338	45.88	0.38
Neutral	3.3854	0.3854	46.24	0.37
Positive	4.1382	0.1382	46.20	0.32
Very Positive	4.9416	0.0584	43.92	0.60
Average	-	0.2234	45.18	0.43
SFT + PPO ($\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{cons}}$)				
Very Negative	1.1954	0.1954	46.16	0.50
Negative	2.3254	0.3254	48.55	0.38
Neutral	3.3760	0.3760	48.95	0.38
Positive	4.1330	0.1330	49.21	0.32
Very Positive	4.9374	0.0626	47.88	0.60
Average	-	0.2185	48.15	<u>0.44</u>
SFT + PPO ($\mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$)				
Very Negative	1.2004	0.2004	46.19	0.50
Negative	2.3146	0.3146	48.80	0.38
Neutral	3.3904	0.3904	49.14	0.38
Positive	4.1370	0.1370	49.21	0.32
Very Positive	4.9410	0.0590	46.53	0.60
Average	-	0.2203	<u>47.97</u>	<u>0.44</u>
SFT + PPO ($\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$)				
Very Negative	1.2114	0.2114	46.18	0.51
Negative	2.3234	0.3234	48.49	0.40
Neutral	3.3748	0.3748	48.89	0.39
Positive	4.1462	0.1462	49.04	0.33
Very Positive	4.9374	0.0374	46.39	0.62
Average	-	<u>0.2186</u>	47.80	0.45

Table 2: The results of YELP. Both breakdown results and average results are displayed.

demonstrating the model’s accuracy in identifying emotionally charged expressions. Most misclassifications occur between adjacent categories. For example, Positive as Very Positive (1916 cases), underscoring the challenge of distinguishing subtle emotional shifts. Neutral samples exhibit high ambiguity, often confused with neighboring categories (4985 instances total), likely due to overlapping lexical patterns. Despite these challenges, the model shows strong performance, effectively capturing both clear and nuanced emotional signals.

6.2.2 Visualization of the Lexicon Distribution

Figure 3b compares lexicon distribution across readability intensity levels regarding both reference and generated texts. Results show high lexical consistency, especially at style polarity—vocabulary



(a) Heatmap of factual consistency on the YELP dataset. Each cell indicates the number of instances. The sentiment of the transferred text is predicted via a fine-tuned DistilBERT classifier.

(b) Visualization of the lexicon distribution. We showcase some style-specific words and their TF-IDF values in the pivot vectors and the generated text. The solid line denotes the distribution of the annotated corpus and the imaginary line denotes the distribution of the generated text. The words from left to right sides showcase a decreasing lexical difficulty.

Figure 3: Factual consistency and lexicon distribution analysis of SFT + PPO ($\mathcal{R}_{\text{sent}} + \mathcal{R}_{\text{lex}} + \mathcal{R}_{\text{cons}}$).

Category	Text
Original	The food here is absolutely fantastic. It’s always busy, which speaks to its popularity, but the staff are efficient and make sure you’re seated and served quickly. Be sure to order the pitcher of beer—it’s a great deal and will save you a good amount.
Elementary	The food here is really good . It’s always busy, which is good because many people come here. The staff is very helpful . They make sure you get seated quickly and get your food. Don’t forget to order a pitcher of beer. It’s a great deal and will save you a lot.
Middle	The food at this restaurant is always great , and it’s easy to see why it’s so popular . The staff is always friendly and ready to help, making sure your meal comes out quickly . It’s also a good idea to order a pitcher of beer—it’s a great deal and can save you a lot of money.
High	The food at this establishment is consistently excellent , and it’s clear why it’s so popular . The service is prompt and attentive , ensuring that your meal is served promptly . Additionally , it’s a good idea to order a pitcher of beer—it’s a great deal and can save you quite a bit of money.
College	The dining experience at this establishment epitomizes a remarkable level of quality and service. Despite its perpetually bustling atmosphere , which reflects the restaurant’s considerable appeal, the staff exhibits a remarkable level of efficiency , ensuring prompt service and prompt attention to patrons . Moreover , it is imperative to reserve a pitcher of lager—it offers an exceptional value and significantly reduces your overall expenditure .

Table 3: Transferred text of a restaurant review with different levels of readability intensities. Highlighted words are those in style-specific vocabulary and the darker color indicates a high level of intensity.

in generated texts closely mirrors that of references. Even at intermediate levels, differences remain minimal, highlighting the model’s ability to reproduce readability-specific vocabulary distributions. This strong lexical alignment, supported by the hierarchical reward design, further confirms the effectiveness of our paradigm in capturing fine-grained stylistic attributes.

6.2.3 Case Study

As shown in Table 3, we transform a restaurant review in YELP dataset into four versions with different readability levels. The text length increases progressively with higher reading difficulty. Fur-

thermore, the number of highlighted words indicates that lower-readability versions tend to include more advanced vocabulary, such as “*epitomizes*” and “*imperative*”. This example demonstrates the effectiveness and sensitivity of our method in generating text with varying stylistic intensities.

7 Conclusion

We propose an SFT-then-PPO paradigm for UTST with controllable intensity, using filtered pseudo-parallel corpora for SFT and hierarchical PPO for fine-grained style control. Experiments on CN-N/DM and Yelp show improved readability and

sentiment transfer, with accurate style alignment, strong content preservation, and enhanced lexical precision over GPT-4o-mini. This approach advances UTST and enables broader applications in text transformation. Future work will explore extending intensity control to additional stylistic dimensions for greater real-world applicability.

Limitations

Our UTST paradigm exhibits promising performance in controllable intensity text style transfer, one potential limitation concerns the linguistic scope of the evaluation. The current experiments are conducted on English-language datasets (CNN/DM and Yelp), which, while representative of common style transfer benchmarks, may not fully capture the method's effectiveness across diverse linguistic and cultural contexts. Given that the proposed approach is inherently language-agnostic, future work could benefit from extending the evaluation to multilingual or non-English datasets, thereby providing a more comprehensive assessment of the model's generalizability and practical utility in broader real-world scenarios.

Acknowledgments

The authors would like to thank the anonymous reviewers for their insightful comments. This work is supported by the Chenguang Program of Shanghai Education Development Foundation and Shanghai Municipal Education Commission under Grant 24CGA26.

References

- Eleftheria Briakou, Sweta Agrawal, Joel Tetreault, and Marine Carpuat. 2021. Evaluating the evaluation metrics for style transfer: A case study in multilingual formality transfer. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1321–1336.
- Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric Xing, and Zhiting Hu. 2022. Rlprompt: Optimizing discrete text prompts with reinforcement learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3369–3391.
- Lijun Feng, Martin Jansche, Matt Huenerfauth, and Noémie Elhadad. 2010. A comparison of features for automatic readability assessment. In *Coling 2010: Posters*, pages 276–284.
- Rudolf Franz Flesch. 1948. A new readability yardstick. *Journal of Applied Psychology*, 32(3):221–233.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2018. Style transfer in text: Exploration and evaluation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Hongyu Gong, Suma Bhat, Lingfei Wu, JinJun Xiong, and Wen-mei Hwu. 2019. Reinforcement learning based text style transfer without parallel training corpus. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3168–3180, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jingxuan Han, Quan Wang, Zikang Guo, Benfeng Xu, Licheng Zhang, and Zhendong Mao. 2024. Disentangled learning with synthetic parallel data for text style transfer. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15187–15201, Bangkok, Thailand. Association for Computational Linguistics.
- Alexander Havrilla, Maksym Zhuravinskyi, Duy Phung, Aman Tiwari, Jonathan Tow, Stella Biderman, Quentin Anthony, and Louis Castricato. 2023. trIX: A framework for large scale reinforcement learning from human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8578–8595, Singapore. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomáš Kociský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *NIPS*, pages 1693–1701.
- Zhiqiang Hu, Roy Ka-Wei Lee, and Charu C. Aggarwal. 2020. Text style transfer: A review and experiment evaluation. *CoRR*, abs/2010.12742.
- Joosung Lee. 2020. Stable style transformer: Delete and generate approach with encoder-decoder for text style transfer. In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 195–204.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. Delete, retrieve, generate: a simple approach to sentiment and style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1865–1874.
- Xiaolin Li, Lei Huang, Yifan Zhou, and Changcheng Shao. 2021. Tst-gan: A legal document generation model based on text style transfer. In *2021 4th International Conference on Robotics, Control and Automation Engineering (RCAE)*, pages 90–93. IEEE.

- Pusheng Liu, Lianwei Wu, Linyong Wang, Sensen Guo, and Yang Liu. 2024. Step-by-step: Controlling arbitrary style in text with large language models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15285–15295.
- 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](#). *CoRR*, abs/1907.11692.
- Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Zhifang Sui, and Xu Sun. 2019a. A dual reinforcement learning framework for unsupervised text style transfer. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI 2019*.
- Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Xu Sun, and Zhifang Sui. 2019b. [A dual reinforcement learning framework for unsupervised text style transfer](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5116–5122. International Joint Conferences on Artificial Intelligence Organization.
- Remi Mir, Bjarke Felbo, Nick Obradovich, and Iyad Rahwan. 2019. [Evaluating style transfer for text](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 495–504, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sourabrata Mukherjee and Ondřej Dušek. 2023. Leveraging low-resource parallel data for text style transfer. In *Proceedings of the 16th International Natural Language Generation Conference*, pages 388–395.
- Sourabrata Mukherjee, Atul Kr. Ojha, John Philip McCrae, and Ondřej Dusek. 2025. [Evaluating text style transfer evaluation: Are there any reliable metrics?](#) In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*, pages 418–434, Albuquerque, USA. Association for Computational Linguistics.
- Kai North, Marcos Zampieri, and Matthew Shardlow. 2023. Lexical complexity prediction: An overview. *ACM Computing Surveys*, 55(9):1–42.
- OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, and 401 others. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. [Style transfer through back-translation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 866–876, Melbourne, Australia. Association for Computational Linguistics.
- Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. A recipe for arbitrary text style transfer with large language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848.
- Leonardo F. R. Ribeiro, Mohit Bansal, and Markus Dreyer. 2023. [Generating summaries with controllable readability levels](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11669–11687, Singapore. Association for Computational Linguistics.
- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. “transforming” delete, retrieve, generate approach for controlled text style transfer. 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 3269–3279.
- Mirac Suzgun, Luke Melas-Kyriazi, and Dan Jurafsky. 2022. Prompt-and-rerank: A method for zero-shot and few-shot arbitrary textual style transfer with small language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2195–2222.
- Keren Tan, Kangyang Luo, Yunshi Lan, Zheng Yuan, and Jinlong Shu. 2024. An llm-enhanced adversarial editing system for lexical simplification. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1136–1146.
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, and 1 others. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). *CoRR*, abs/1509.01626.
- Chang Zong, Yuyan Chen, Weiming Lu, Jian Shao, Yongfeng Huang, Heng Chang, and Yueting Zhuang. 2024. Proswitch: Knowledge-guided instruction tuning to switch between professional and non-professional responses. *arXiv preprint arXiv:2403.09131*.

A Prompts for Parallel Data Generation

We present prompts for generating readability and sentiment parallel data.

A.1 readability

Rewrite the following text to match the {outputDesc} readability level (Flesch Reading Ease score {FREScore}).

Adjust vocabulary, sentence length, and complexity to strongly reflect the target readability level.

For Elementary, use very simple words and short sentences.

For Middle School, use moderately simple words and slightly longer sentences.

For High School, use more complex words and varied sentence structures.

For College, use advanced vocabulary and complex sentence structures.

Keep the core meaning and content consistent, but adapt the style to be natural and concise.

Input: {text}

Output:

A.2 sentiment

Rewrite the following {inputRating} Yelp review into an {outputRating} review.

Transform the tone and wording to reflect the target rating's sentiment strongly.

For 5 Stars, use highly positive words like "amazing" or "outstanding";

for 1 Star, use strongly negative words like "terrible" or "awful".

Keep core aspects (food, service, atmosphere) consistent, but adjust their descriptions to match the target rating.

Ensure the output is natural, realistic, concise (similar length to the input), and avoids vague terms like "okay" or "fine".

Input: {text}

Output:

B Prompts for Controllable Text Generation

We present prompts for generating readability and sentiment text with controllable intensity.

B.1 readability

Rewrite the following text for {readability level}

Input: {text}

Output:

B.2 sentiment

Rewrite the following text for {sentiment level}

Input: {text}

Output:

C Evaluation Metrics

- **FRE**: Measures readability using the Flesch Reading Ease score, assessing text complexity for readability transfer:

$$\text{FRE} = 206.835 - 1.015 \cdot \frac{\text{total words}}{\text{total sentences}} - 84.6 \cdot \frac{\text{total syllables}}{\text{total words}}.$$

FRE_{Δ} quantifies deviation from the target range's midpoint:

$$\text{FRE}_{\Delta} = |\text{FRE}_{\hat{y}} - \text{FRE}_{s_y}|,$$

where $\text{FRE}_{\hat{y}}$ is the FRE of the generated text, and FRE_{s_y} is the target range's midpoint (Table 4).

- **STAR**: Represents the predicted style class (1–4 for readability, 1–5 for sentiment) using classifiers (e.g., DistilBERT for sentiment, FRE-to-level mapping for readability). STAR_Δ measures deviation from the target class:

$$\text{STAR}_\Delta = |\text{STAR}_{\hat{y}} - \text{STAR}_{s_y}|,$$

where $\text{STAR}_{\hat{y}}$ is the predicted class, and STAR_{s_y} is the target class.

- **RG-L**: Assesses content preservation via ROUGE-L, measuring the longest common subsequence between input and generated texts:

$$\text{RG-L} = \frac{\text{LCS}(X, \hat{Y})}{|\hat{Y}|},$$

where $\text{LCS}(X, \hat{Y})$ is the length of the longest common subsequence between input X and generated text \hat{Y} , and $|\hat{Y}|$ is the length of \hat{Y} .

- **H-Re**: Evaluates overall transfer quality, equivalent to the reward functions for readability and sentiment in Section 4.2.3, but with the BERTScore coefficient set to 0:

$$\text{H-Re} = 0.5 \cdot \mathcal{R}_{\text{sent}} + 0.5 \cdot \mathcal{R}_{\text{lex}},$$

where $\mathcal{R}_{\text{sent}}$ and \mathcal{R}_{lex} are defined in Section 4.2.1.

D DataSet Overview

Table 4 demonstrates the details of the readability and sentiment datasets, involving level scores, intensity attributes, and data distribution.

Style	Dataset	SCORES	ATTRIBUTES	Train	Test
Readability	CNN/DM	$80 \leq \text{FRE} \leq 100$	Elementary School Students	4,000	4,000
		$60 \leq \text{FRE} < 80$	Middle School Students	4,000	
		$40 \leq \text{FRE} < 60$	High School Students	4,000	
		$0 \leq \text{FRE} < 40$	College Students	4,000	
Sentiment	Yelp	5 stars	Very Positive Tone	4,000	5,000
		4 stars	Positive Tone	4,000	
		3 stars	Neutral Tone	4,000	
		2 stars	Negative Tone	4,000	
		1 star	Very Negative Tone	4,000	

Table 4: Attributes for Style, categorized by Readability and Sentiment. The number of training sets here refers to the number of constructed pseudo-parallel datasets.

E Representative words for stylistic intensities

As shown in Table 5, some representative words in the readability and sentiment corpora for stylistic intensities are listed, forming the vocabulary distribution for the lexicon-level reward.

F Explanation of the decrease in RG-L

On the Readability dataset, the RG-L metric only exhibits a slight decrease. This phenomenon can be attributed to the core calculation logic of the metric: RG-L incorporates the Longest Common Subsequence (LCS) of the text before and after style transformation for analysis. The design of this metric focuses on measuring the degree of textual content preservation between the pre- and post-transformation texts. However, the requirement for such content preservation inevitably leads to changes in the lexical hierarchy,

Scores	feature words
Readability	
elementary school	say, people, big, hurt, did, help, get, year, call, bad
middle school	mention, just, believe, like, old, decide, share, know, right, important
high school	state, according, express, additionally, currently, describe, recently, claim
college	regarding, significant, individual, following, subsequently, notably, furthermore
Sentiment	
very positive	outstanding, amazing, absolutely, incredible, truly, fantastic, gem, incredibly
positive	solid, great, overall, definitely, enjoyable, satisfying, delightful, tasty, really
neutral	bit, decent, quite, improvement, potential, overall, nice, consider, appreciate
negative	disappointing, unfortunately, left, desired, frustrating, better, expected, lacking
very negative	terrible, awful, bland, disaster, disappointment, avoid, worst, completely, barely

Table 5: Representative words in readability and sentiment corpora for stylistic intensities.

making the RG-L metric irrelevant to the core goal of "evaluating style transformation intensity". Thus, the effectiveness of style transformation methods cannot be comprehensively assessed solely based on this metric.

F.1 Evidence from Tabular Case

Table 6 presents a specific empirical case that supports the aforementioned conclusion.

Text Category	FRE	RG-L	Text Content
Original Text	66.10	—	The scientist did careful tests to get correct results.
Rewritten Text 1	33.07	66.67	The scientist did careful experiments to obtain precise results.
Rewritten Text 2	15.64	52.63	The scientist carefully implemented the experiments to ensure precise results.

Table 6: Example of Text Style Transformation with Different Intensity Levels

F.2 Analysis of Metric-Goal Adaptability

The initial objective of this study was to rewrite the original text into a version that meets the "college-level and above reading standard", which corresponds to a FRE score range of 0–40. As shown in the table, both Rewritten Text 1 and Rewritten Text 2 satisfy this requirement.

Further analysis of style transformation intensity reveals that the FRE score of Rewritten Text 2 (15.64) is lower than that of Rewritten Text 1 (33.07). This numerical difference indicates that Rewritten Text 2 exhibits a more significant stylistic shift. Nevertheless, although the RG-L value of Rewritten Text 2 (52.63) is lower than that of Rewritten Text 1 (66.67), this difference in RG-L values does not diminish its superiority in style transformation effectiveness—this further confirms the conclusion that the RG-L metric cannot effectively reflect style transformation intensity.

G Hyperparameter Settings

Table 7 shows some settings for UTST with the SFT-then-PPO paradigm, including computing infrastructure, GPT-4o-mini settings, Deepspeed settings, SFT settings and PPO settings.

G.1 Setting Strategy of Hyperparameters for the Combined Reward Function

For the combined reward function, λ_1 , λ_2 , and λ_3 are configured via a fixed-weight strategy, excluding resource-intensive grid search. The weights adhere to an order-preserving principle, sorted by their significance: the sentence-level reward (R_{sent}) is assigned the highest weight, with the objective of prioritizing the alignment between the global style of generated text and the target intensity; the word-level reward (R_{lex}) follows, employed to enhance the matching between fine-grained lexical features and the target style; the consistency reward (R_{cons}) is assigned the lowest priority, as it is not the core focus of

Parameter	Value
<i>Computational Settings</i>	
GPU	NVIDIA A6000 48GB (2 units)
flan-t5-large	783M parameters
Roberta-large	355M parameters
DistilBERT base model	67M parameters
<i>GPT-4o-mini Settings</i>	
temperature	0.7
<i>Deepspeed Settings</i>	
ZeRO Optimization Stage	2
Mixed Precision	bf16
Gradient Accumulation Steps	8
Gradient Clipping	1.0
Communication Window Size	512
<i>SFT Settings</i>	
Optimizer	AdamW
Optimizer Parameters	$\beta = (0.9, 0.999), \epsilon = 10^{-8}$
Learning Rate	10^{-4}
Weight Decay	0.0
Dropout Rate	0.1
Batch Size	16
Epochs	20
Random Seed	42
<i>PPO Settings</i>	
Optimizer	AdamW
Learning Rate	10^{-5}
Weight Decay	10^{-6}
Batch Size	2
Epochs	10
Total Steps	100000
Num Rollouts	512
Initial KL Coefficient	0.05
Target KL Divergence	6
Horizon	10000
Discount Factor (γ)	0.99
GAE Lambda (λ)	0.95
Policy Clip Range	0.2
Value Clip Range	0.2
Reward Clip Range	10
Max New Tokens (Generation)	256
Top-k Sampling	50
Top-p Sampling	0.95
Random Seed	42
<i>Sentiment Classifier</i>	
Optimization	AdamW
Learning rate	$2e-5$
Epochs	6
Batch size	16

Table 7: Hyperparameter settings

this study. This reward merely requires the preservation of core semantics during style transfer, and no ablation experiments were performed for this component.