

Cognitive Enhancement Chain-of-Thought Towards Enhancing Style Learning and Content Preservation for Long Style Transfer

Lianwei Wu^{1,2}, Botao Wang¹, Wenbo An¹, Tieqiao Li¹, Xianghua Li^{3*}

¹School of Computer Science, Northwestern Polytechnical University, China

²Research & Development Institute of Northwestern Polytechnical University in Shenzhen, China

³School of Artificial Intelligence, Optics and Electronics (iOPEN), Northwest Polytechnical University, China
wlw@nwpu.edu.cn, {wangbotao, 2074137878, litieqiao}@mail.nwpu.edu.cn, li_xianghua@163.com

Abstract

Current text style transfer task mainly focuses on short texts, while the field has not been fully developed for long texts. Considering the richer semantics and more complex sentence structures in long text sequences, existing methods that employ traditional style-content disentanglement ways and learn the target style to generate target sequences face two key issues: 1) During disentanglement, they usually directly separate style words or fragments, such as coarse-grained disentanglement risks losing original semantics and hinder the model’s content preservation. 2) During target style learning, they often focus on the transfer of certain style attributes or aspects, which makes it difficult to grasp the holistic style of target objects. To this end, we propose **Cognitive enhancement Chain-of-Thought (CeCoT)** towards enhancing style learning and content preservation for long style transfer. CeCoT first constructs progressive CoT to facilitate LLMs to gradually rewrite source content and separate source styles, thereby enhancing the retention of original content. Then, we propose cognitive CoT, which comprehensively considers hierarchical cognitive content (i.e., shallower-deeper-normal level) and cognitive behavior (i.e., prompt order of CoT) to learn the overall target style. To enhance the robustness of our model, we also propose two constraint losses in a dual validation way towards content preservation enhancing and style consistency learning. Extensive experiments on two competitive datasets demonstrate the superiority of our CeCoT.

Introduction

Text style transfer task is the process of changing style attributes of a source text sequence while keeping its content unchanged, to generate a new sequence (a.k.a. the target sequence). Transfer types of this task primarily include emotional transfer (Zhu et al. 2024b), formal transfer (Hu et al. 2022), humor transfer (Christ et al. 2024), and imitation of specific characters, etc., which have also been fully applied in the fields (Jin et al. 2022; Wu et al. 2023b,c) of poetry creation, advertising generation, e-commerce customer service, etc. Appropriate text style can promote interpersonal communication, product promotion, and even political relations. Recently, with the advent of the era of large language models (LLMs), this field has received increasing attention to enhance the personalization and fun of generated content (Jing et al. 2023).

Text style transfer could be categorized into short-text and long-text approaches. For **short text transfer**, methods mainly adopt explicit and implicit style-content disentanglement. *Explicit-based methods* separate content and style in source text sequences through editing or replacement, such as frequency-based (Liu et al. 2024) or attention-based methods (Lee et al. 2021; Lin et al. 2024). *Implicit-based methods* capture implicit representations of style and transfer them to flexible expressions of target style, which include back-translation (Wei et al. 2023), reinforcement learning (Deng et al. 2022), and disentanglement (Zhu et al. 2022, 2024a) strategies. Owing to its relatively limited semantics in short text, these methods exhibit low complexity and achieve relatively straightforward transfer. **Long text style transfer**, by contrast, has broader applications such as literary adaptation, imitation of specific characters, etc. However, these methods are relatively scarce and largely extend short-text methods. They mainly learn complex hidden features in long texts through implicit disentanglement, and rely on autoregression (Xiao et al. 2023), variational autoencoder (Sharma, Dhiman, and Kumar 2024), and generative adversarial networks (Chen, Tennent, and Hsu 2024) to generate new sequences with target styles.

Nevertheless, there are two key challenges in long text style transfer: **1) Easy to lose original semantics during style-content disentanglement.** Long text has richer semantics and more complex discourse structures compared to short text. Existing methods typically leverage explicit or implicit disentanglement to separate style words or fragments from source sequences to obtain pure original semantics. Considering that these fragments still contain key original semantics, such as coarse-grained separating methods easily lose original semantics and interfere with the preservation of content. **2) Difficult to learn holistic style during target style learning.** In long text, target style is not as concise and clear as in short text. Like in short text emotion transfer, only emotional words in source sequences need to be replaced (e.g. positive \rightarrow negative). For the target style of long texts, such as the style of a specific author, it has rich style features such as word usage habits, sentence structure, expression framework, etc. Current methods often focus on word-level imitation, failing to capture the target’s complex global style for holistic style transfer.

To address these issues, we propose **Cognitive enhancement Chain-of-Thought (CeCoT)** to enhance content preservation and style learning for long style transfer. Specifically, to obtain pure

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

content semantics from the source sequence while removing the style features, leveraging the powerful understanding and reasoning capabilities of LLMs, we design progressive CoT to facilitate LLMs to gradually learn deep pure semantics from source sequences. To learn global style features of the target object, we construct cognitive CoT that models cognitive content around “shallower-deeper-normal” prompt strategies and cognitive behavior of “prompt order of CoT” to enhance overall target style learning. Finally, we propose two constraint losses to strengthen original content preservation and target style consistency learning, thereby improving long text transfer. Experiments validate on two public datasets confirm the effectiveness of our model. Contributions of our paper are summarized as follows:

- To the best of our knowledge, we propose the first framework that integrates cognition into LLMs for long text style transfer, which constructs progressive and cognitive CoTs to promote LLMs to enhance the preservation of original semantics while capturing the overall target style from multiple perspectives.
- We devise CCoT around cognitive content and cognitive behavior, which establishes “shallower-deeper-normal” multi-perspective style rewriting strategies in cognitive content and considers the influence of prompt order in cognitive behavior, so as to enhance the overall target style learning.
- Two constraint losses respectively boost the preservation of source semantics and target style learning. Extensive experiments verify the superiority of CeCoT.

Related Work

Text Style Transfer

Text style transfer methods are classified as supervised, semi-supervised, or unsupervised based on parallel corpus requirements (see Figure 1). **Supervised learning methods** require parallel corpora and mainly include two types: *data augmentation* (Zhang, Ge, and Sun 2020) and *vocabulary constraint decoding* (Zhang, Ge, and Sun 2020). Data augmentation focuses more on joint learning based on multi-task learning (Jin et al. 2022). **Semi-supervised methods** prefer to learn *hidden space cross-projection* (Shang et al. 2019) and employ *pseudo parallel corpus* (He et al. 2020). The first relies on Seq2Seq and projects alignment of hidden layers between encoders to improve content preservation. The second (Horvitz et al. 2024) mainly leverages retrieval and generative-based strategies to enhance the consistency between target style and its content.

Unsupervised methods that do not require parallel corpora have become the mainstream of research (Han et al. 2023). They are divided into *disentanglement-based* (Xu et al. 2024) and *non-disentanglement methods* (Hong et al. 2024) according to whether style and content need to be separated. The former mainly includes explicit (Shi et al. 2023) and implicit (Chen, Tennent, and Hsu 2024) content-style separation. The focus of **explicit separation** is on how to identify and replace style words, which mainly includes frequency-based (Narasimhan et al. 2023), attention-based (Wu et al. 2023a), and a combination of both (Ge et al. 2023; Liu et al. 2024) ways. This type of method has the advantage of strong interpretability, but it is difficult to meet the requirement of complex style transfer tasks. The goal of **implicit separation** is to indirectly separate the content and style

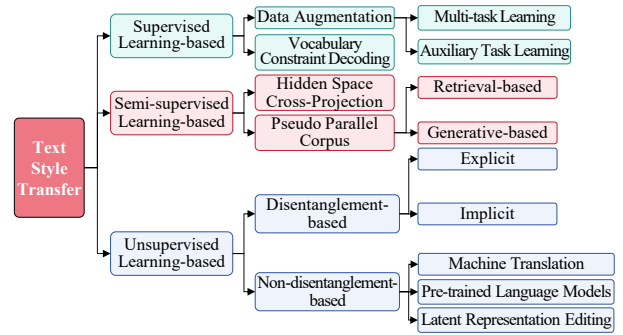


Figure 1: Classification of text style transfer methods.

of source sequences based on autoencoders (Pour et al. 2023), variational autoencoders (Sharma, Dhiman, and Kumar 2024), and adversarial networks (Liu et al. 2023), which has the advantage of better capturing original semantics. Considering that the content and style of most text sequences are often intertwined, it is difficult to separate them into two independent parts (Lv et al. 2024). Therefore, **non-disentanglement methods** have emerged, which generally include machine translation (Wei et al. 2023), pre-trained language models (Narasimhan et al. 2023), latent representation editing (Li, Li, and Hoi 2024; Shu et al. 2024), etc. Compared with short text, long text has richer semantics and more complex discourse structures. Current methods often suffer from drawbacks of losing semantics from source text and difficulty in learning target styles. Thus, we strive to stimulate reasoning abilities of LLMs to enhance style transfer in long texts.

Chain-of-Thought of LLMs

The key to the powerful reasoning ability of LLMs lies in the invention of CoT prompts (Wei et al. 2022). CoT prompts simulate the process of human problem-solving, guiding the model step by step to demonstrate its reasoning ideas and logical relationships, thereby achieving complex thinking processes. The current CoT approaches mainly focus on how to design effective prompts to explore more valuable knowledge in LLMs, including self-consistent prompts (Lin et al. 2025), generated knowledge prompts (Zhang et al. 2024), and minimum to maximum prompts (Maia Polo et al. 2024). Current methods do not fully consider the deep-level cognition of human, i.e., people have the characteristics of viewing things from multiple perspectives and thinking leaps. Therefore, our model designs cognitive CoTs from both cognitive content (around shallower-deeper-normal level) and cognitive behavior (around prompt order consideration) to enhance the ability to transfer long text styles.

The Proposed Approach

We propose cognitive enhancement chain-of-thought method (CeCoT) for long text style transfer, consisting of five core components, as shown in Figure 2. Next, we will describe each level of our module in detail.

Task Definition

This task could be defined as: Given a long text sequence (a.k.a. a source sequence) $x = \{x_1, x_2, \dots, x_n\}$ with n tokens and its

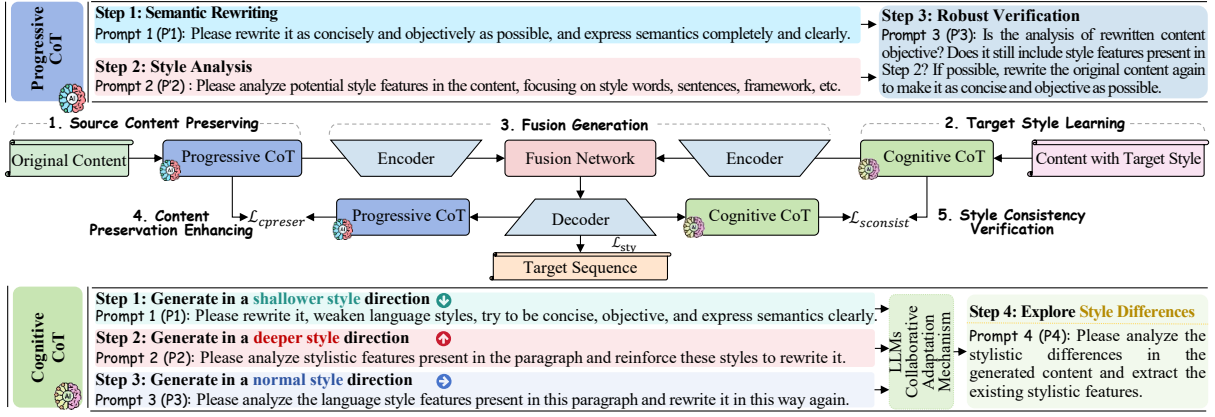


Figure 2: An overview of our CeCoT. Source content preserving layer designs progressive CoT to apply LLM for capturing pure content semantics, and target style learning layer explores cognitive CoT to promote LLM to learn overall styles.

style label $s \in \mathcal{S}$ (\mathcal{S} is the set of long text styles). The goal is to generate a new long text sequence (a.k.a. the target sequence) that contains multiple sentences, which possesses significant target style features $s' \in \mathcal{S}$ ($s' \neq s$) while preserving original semantics.

Source Content Preserving

To remove style features from the source sequence and preserve content semantics, considering the powerful understanding and reasoning ability of LLMs, we design progressive CoT prompts to facilitate LLMs to gradually disentangle content and style features from the source sequence, thereby learning pure original content semantics. Progressive CoT prompts mainly involve the following three steps:

Step 1: Semantic Rewriting Firstly, to ensure that the model focuses on understanding semantic content of the source text, we design *Semantic Rewriting Prompt* to guide the model to rewrite the semantics of the source text in a more concise and objective way. This step can be formalized as:

$$c^{re} = \text{LLM}(x, \mathbf{P}'1) \quad (1)$$

Step 2: Style Analysis Next, to separate the original style from source text, we instruct LLMs to fully explore the style of the source text from multiple perspectives of word-level, sentence-level, and framework-level, which are formalized as:

$$c^{sa} = \text{LLM}(x, \mathbf{P}'2) \quad (2)$$

Step 3: Robust Verification Finally, to verify whether content semantics of source text obtained in step 1 are still mixed with style information, we design robust verification prompt that combines the style information in step 2 to check the content in step 1 and further promote LLMs rewriting, thereby acquiring a more robust and objective content of source text. This step can be formalized as:

$$c^{rv} = \text{LLM}(x, c^{re}, c^{sa}, \mathbf{P}'3) \quad (3)$$

where $\mathbf{P}'1$, $\mathbf{P}'2$, and $\mathbf{P}'3$ are shown in the progressive CoT module of Figure 2.

Target Style Learning

Styles in long text generally imply multiple complex structures throughout the entire text, making it difficult to clearly describe through a single or a few token labels. For this purpose, we develop target style learning layer that incorporates cognitive

CoT (CCoT) with dual-view of cognitive content and behaviors to facilitate LLMs in learning more precise target style features. CCoT includes two aspects:

Content Cognition We know that people have different sensitivities to colors, and similarly, readers have different sensitivities to text style. Some people can keenly perceive the style of the text, while others may not feel the style features (they may only perceive the semantic information of the text). Therefore, we build prompt strategies by simulating people’s different levels of recognition of styles, designing three types of gradual style rewriting prompts “shallower-deeper-normal”, and exploring their differences to comprehensively enhance target style learning in the target sample y' . Content cognition includes four steps:

Step 1: Generate in a shallower style direction When people look at the style of long text, some may not be sensitive to style features and may not perceive their existence. To this end, we model the cognition of this group of people and rewrite the long text in a ‘shallower’ direction, formalized as:

$$s^{ws} = \text{LLM}(y', \mathbf{P}1) \quad (4)$$

Step 2: Generate in a deeper style direction On the contrary, some people are very sensitive to language style and can easily capture style information in long text. Thus, we rewrite this long text in a ‘deeper’ direction. This step can be formalized as:

$$s^{ss} = \text{LLM}(y', \mathbf{P}2) \quad (5)$$

Step 3: Generate in a normal style direction Certainly, most people have a normal perception of language style and can effectively discover style information in long text. Therefore, we rewrite this long text in a ‘normal’ direction, formalized as:

$$s^{ps} = \text{LLM}(y', \mathbf{P}3) \quad (6)$$

Step 4: Explore style differences To highlight overall styles, we construct style differentiation prompt to facilitate LLM to analyze styles in three steps mentioned above, formalized as:

$$s^s = \text{LLM}(S, s^{ws}, s^{ss}, s^{ps}, \mathbf{P}4) \quad (7)$$

where S represents original target style, $\mathbf{P}1$, $\mathbf{P}2$, $\mathbf{P}3$, and $\mathbf{P}4$ are shown in cognitive CoT of Figure 2. In this way, we can obtain the highlighted target style s^s .

Behavioral Cognition Experiments have found that different prompt orders can leave a preconceived impression on LLMs and generate responses of different qualities (Li et al. 2023), just

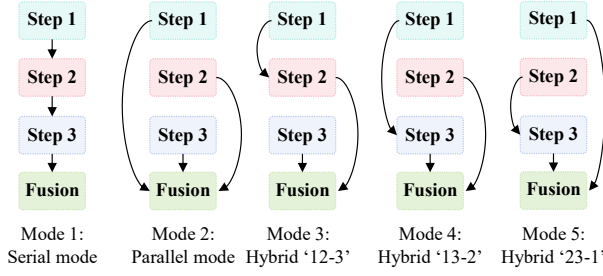


Figure 3: The behavioral cognition patterns of prompts.

like our human behavioral cognition. Considering that “different prompt orders of agents affect results”, we devise prompt combination mechanism that explores different prompt execution orders in LLMs, thereby further boosting the learning of target styles. It includes five types of combination modes:

Mode 1: Serial mode To learn influence of different prompt execution orders, three prompts are first input into the same LLM to obtain their rewritten content and further are integrated together.

$$S_{m1} = [\text{LLM}(\text{LLM}(\text{LLM}(y', \mathbf{P1}), \mathbf{P2}), \mathbf{P3})] \quad (8)$$

where $;$ is the concatenation operation.

Mode 2: Parallel mode To ensure that different prompts do not affect each other, we input three prompts into different LLMs (which can be copies of the same LLMs) to obtain their rewritten content, and then integrate them together. It can be formalized as:

$$S_{m2} = [\text{LLM}_1(y', \mathbf{P1}); \text{LLM}_2(y', \mathbf{P2}); \text{LLM}_3(y', \mathbf{P3})] \quad (9)$$

Furthermore, we also learn the influence between the two prompts and then integrate the content obtained from the third prompt. Thus, three hybrid modes are designed: **Hybrid ‘12-3’**, **Hybrid ‘13-2’**, and **Hybrid ‘23-1’**, where Hybrid ‘12-3’ means inputting prompt 1 and prompt 2 into the same LLM, and prompt 3 into another LLM. Similarly, Hybrid ‘13-2’ and Hybrid ‘23-1’ will not be further elaborated.

$$S_{m3} = [\text{LLM}_1(\text{LLM}_1(y', \mathbf{P1}), \mathbf{P2}); \text{LLM}_2(y', \mathbf{P3})] \quad (10)$$

$$S_{m4} = [\text{LLM}_1(\text{LLM}_1(y', \mathbf{P1}), \mathbf{P3}); \text{LLM}_2(y', \mathbf{P2})] \quad (11)$$

$$S_{m5} = [\text{LLM}_1(\text{LLM}_1(y', \mathbf{P2}), \mathbf{P3}); \text{LLM}_2(y', \mathbf{P1})] \quad (12)$$

The formula for prompt 4 can be optimized as:

$$s^s = \text{LLM}([S, S_{m1}; S_{m2}; S_{m3}; S_{m4}; S_{m5}], \mathbf{P4}) \quad (13)$$

In the end, we obtain precise target style information to be imitated.

Fusion Generation

After obtaining semantic content of the source sequence and target style, we design fusion generation to facilitate interaction and integration of the two and gradually generate new sequences with the target style.

Encoder Network We first encode source content semantics and the target style to obtain their hidden representations. In our model, Transformer encoder serves as the encoder.

As the core of Transformer encoder, multi-head self-attention network can capture deep dependencies of a sequence from multiple perspectives. For each self attention, it focuses on learning internal semantic dependencies between two tokens in a sequence:

$$H = \text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (14)$$

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V) \quad (15)$$

$$\begin{aligned} M &= \text{Attention}(Q, K, V) \\ &= \text{Concat}(\text{head}_1, \text{head}_2, \dots, \text{head}_m)W^o \end{aligned} \quad (16)$$

$$E = \text{Feed-Forward}(M) \quad (17)$$

where Q , K , and V are queries, keys, and value matrices, respectively. d_k is feature dimension size of K . all W are trainable parameters. In our settings, in the encoding of source text, $Q = K = V = C^{rv}$, C^{rv} is the embeddings of source sequence c^{rv} . In the encoding, S^s is the embeddings of target sequence s^s . In this way, we gain encoding features of source text and target style as E_c and E_s , respectively.

Cross-fusion Network To effectively associate and fuse the source text and target style, we design cross fusion network, which first concatenates the two types of information and then uses multiple Transformer layers as fusion network to learn target style-aware representations.

$$F_{cs} = \text{MTL}(Q = K = V = [E_c; E_s]) \quad (18)$$

where F_{cs} contains fusion features of the association between source text and target style.

Decoder Network To generate smooth sequences with target style, we use an autoregressive (AR) decoder to generate target content. The decoder absorbs style-aware associated fusion features F_{cs} by the cross-attention layer.

$$P(y|F_{cs}) = \prod_{t=1} P(y_t|y_{>t}, F_{cs}) \quad (19)$$

Training Objectives

In addition, to further strengthen source content preservation and target style consistency in the training phase, we propose two constraint losses:

Content Preservation Enhancing Loss To enhance the retention of original content, we construct content preservation enhancing loss to compare content semantics of the source sequence obtained in Section with intrinsic content in the generated sequence. Here, we still employ LLMs to mine the intrinsic semantics $c^{rv'}$ in the generated sequence y .

$$c^{rv'} = \text{LLM}(y, c^{re'}, c^{sa'}, \mathbf{P3}) \quad (20)$$

where $c^{re'}$ and $c^{sa'}$ are learned by steps 1 and 2 in section, respectively.

$$\mathcal{L}_{cpre.ser} = \sum_{i=1}^{|E_c|} \sum_{j=1}^{|E_{c'}|} \|E_{ci} - E_{c'j}\|_2^2 \quad (21)$$

where $E_{c'j}$ is the encoded j -th hidden state of intrinsic content semantics $c^{rv'}$.

Style Consistency Verification Loss To enhance the consistency of the generated target style, we construct style consistency verification loss to compare the style of the generated sequence with the style we imitated. For the style learning, we also utilize LLMs described in Section to obtain:

$$s^{s'} = \text{LLM}([y, S'_{m1}; S'_{m2}; S'_{m3}; S'_{m4}; S'_{m5}], \mathbf{P4}) \quad (22)$$

S'_{m1} to S'_{m5} are obtained by steps 1 to 5, where S'_{m1} comes from the output of Mode 1.

$$\mathcal{L}_{sconsist} = \sum_{i=1}^{E_{s'}} \sum_{j=1}^{E_{s'}} \|E_{s'i} - E_{s'j}\|_2^2 \quad (23)$$

Datasets	Chinese			English		
	Style	Size	Avg. Len.	Style	Size	Avg. Len.
Training	JY	2964	344	SP	1161	71
	LX	3036	168	RS	1161	49
	FT	1456	175	-	-	-
Validation	FT	242	175	RS	290	48
Test	FT	729	176	RS	290	50

Table 1: Statistics of two style transfer datasets. Avg. Len. denotes the average length of tokens in each sample.

where E_{s^s} and $E_{s'}$ are the encoded hidden state of sequences s^s and s' , respectively.

Style Classifier Loss We expect generated sequences possessing target styles. Thus, we train a style classifier to construct style loss:

$$\mathcal{L}_{sty} = -\mathbb{E}_{y \sim Decoder} [\log P(s^s | y)] \quad (24)$$

In style classifier, we employ cross-entropy loss to train it in whole training set and leverage soft sampling to allow gradient back-propagation (Lee et al. 2021).

On the whole, the overall loss \mathcal{L} of our CeCoT is constructed by the following balanced way:

$$\mathcal{L} = \mathcal{L}_{sty} + \lambda_1 \mathcal{L}_{cpreser} + \lambda_2 \mathcal{L}_{sconsist} \quad (25)$$

where λ_1 and λ_2 are trainable parameters.

Experiments

We design different CoTs to improve style transfer and try to answer the following questions:

- **RQ1:** Could CeCoT achieve optimal performance compared to the latest baselines?
- **RQ2:** Are different modules of model beneficial for the task?
- **RQ3:** Are different types of CoTs progressiveness?
- **RQ4:** What is impact of CoTs behavior on model performance?
- **RQ5:** Could the impact of CoT on model performance be visually demonstrated?

Datasets

To evaluate the effectiveness of our model in long text style transfer task, we utilize competitive stylized story datasets (Zhu et al. 2023) consisting of both Chinese and English. English dataset mainly includes two types of text styles, everyday essays (ER) from ROCStories (Mostafazadeh et al. 2016) and scripts from Shakespeare’s plays (SP). Chinese dataset mainly includes three types of text styles, namely Fairy Tale (FT) for multiple novel fragments, Lu Xun’s (LX) for realism novels and Jin Yong’s (JY) for martial arts novels. In our experiments, we carry out multiple transfer tasks, i.e., FT \rightarrow LX, FT \rightarrow JY, and ER \rightarrow SP, where FT \rightarrow LX represents transferring Fairy Tale style into Lu Xun’s style.

Settings and Evaluation Metrics

In our model, we utilize *LongLM_{small}* (Guan et al. 2022) and *T5_{small}* (Raffel et al. 2020) as the backbone model of the encoder-decoder architecture for Chinese and English datasets, respectively. For hyperparameters, λ_1 and λ_2 are set to 1.0 and 0.8, respectively. the mini-batch size is 4, and the learning rate of Adam optimizer to 2×10^{-5} .

Evaluation Metrics. Inspired by previous methods (Zhu et al. 2023; Zhao et al. 2024), we evaluate our CeCoT by both automatic and manual evaluation ways. **1) Automatic evaluation.** *Style Transfer Accuracy.* To evaluate style transfer accuracy in target sequences, we apply our model’s encoder fusing a classification layer as style classifier to evaluate models. Here, we utilize document-/sentence-level style classifiers, i.e., Acc_d and Acc_s . *Content Preservation.* We adopt shallow and deep evaluation to measure content preservation. For shallow evaluation, we adopt BLEU₁ and BLEU₂ to measure N-gram overlaps. For deep evaluation, we rely on BERT Score (BS) to gauge semantic similarity between two sequences. Here, we employ BS_p (Precision), BS_r (Recall), and BS_{F1} (F1-score) as deep metrics. We use geometric mean of Acc_d and BLEU₂/BS_{F1} (i.e., Overall₁ and Overall₂) to evaluate entire models’ performance. **2) Manual evaluation.** To enhance user readability, we also randomly sampled 200 samples from Chinese dataset and generated target sequences with LX and JY styles and invite three Chinese native users to score style features (Sty.), content preservation (Cont.), and fluency (Flu.). We set scoring range from poor to high as (Zhu et al. 2023) and utilize average scores as final results.

Performance Evaluation (RQ1)

Baselines Style Transformer (Dai et al. 2019) employs Transformer to learn context dependencies of text sequences for style transfer. **StyleLM** (Syed et al. 2020) generates target text by fine-tuning pre-trained language models (PLMs). **Reverse Attention** (Lee et al. 2021) designs reverse attention in encoder, which evaluates each token’s contribution for style classification, so as to filtering styles in original content. **AugZero-Shot** (Reif et al. 2022) constructs augmented zero-shot prompts to improve PLMs for style transfer. **StoryTrans** (Zhu et al. 2023) devises discourse representations and content enhancing for long story author-style transfer. **SC2** (Zhao et al. 2024) improves larger PLMs by a joint style-content weighted module and style consistent loss for long text transfer.

Results of Comparison

Automatic Evaluation Table 2 shows the results of all models on two datasets, and we observe that:

- In baseline models, Style Transformer and StyleLM show weaker performance compared to other methods, especially in achieving a maximum accuracy of 6.1% in style transfer. However, they show the highest performance on BLEU (up to 87%), demonstrating that they copy a large amount of source content and find it difficult to disentangle source styles from original content.
- Furthermore, Reverse Attention and AugZero-Shot present more compromise performance than the previous two methods, with a significant improvement (obtaining up to 42.4% boost in Acc_d and Acc_s), revealing the usefulness of pre-trained language models for text style transfer.
- SC2 gains the best among all baseline models, showing improvements of up to 5.4% and 3.5% than StoryTrans on Overall₁ and Overall₂, respectively, confirming the effectiveness of jointly optimizing style and content and designing consistency loss.
- Overall, our model fulfills superior results than the latest baseline (SC2) with the overall boost of up to 2.4%, which

Task	Methods	Acc _d	Acc _s	BLEU ₁	BLEU ₂	BS _p	BS _R	BS _{F1}	Overall ₁	Overall ₂
FT → LX	Style Transformer	0.1	1.7	71.7	64.4	91.7	92.6	92.1	2.5	3.0
	StyleLM	0.1	6.1	71.2	61.1	90.6	92.2	91.4	2.5	3.0
	Reverse Attention	6.4	28.0	17.3	4.8	61.4	58.8	60.0	5.4	19.6
	AugZero-Shot	5.9	38.5	11.7	4.1	56.5	59.8	58.0	4.9	18.5
	StoryTrans	39.6	52.0	25.7	9.6	63.0	63.9	63.4	19.5	50.1
	SC2	41.6	70.5	27.5	11.9	61.6	65.5	63.5	22.2	51.4
	Ours	43.3	75.6	29.1	13.2	64.1	67.6	65.8	23.9	53.4
FT → JY	Style Transformer	0.8	3.4	71.8	64.5	91.6	92.6	92.1	7.2	8.7
	StyleLM	1.2	0.9	70.6	60.3	90.3	92.1	91.1	8.5	10.5
	Reverse Attention	63.8	29.1	14.8	3.9	59.1	60.1	59.6	15.8	61.7
	AugZero-Shot	56.0	5.0	12.0	4.2	56.4	59.5	57.8	15.3	56.9
	StoryTrans	70.6	25.1	22.3	7.9	61.2	65.5	63.2	23.6	66.8
	SC2	74.9	41.5	26.2	12.5	64.2	67.7	65.9	30.6	70.3
	Ours	75.8	44.3	28.1	13.5	65.7	68.9	67.3	32.0	71.4
ER → SP	Style Transformer	1.4	5.9	87.0	82.4	97.1	97.5	97.3	10.7	11.7
	StyleLM	0.3	2.7	83.6	80.0	96.0	97.3	96.7	4.9	5.4
	Reverse Attention	43.8	45.7	8.7	2.6	73.6	81.9	77.4	10.7	58.2
	AugZero-Shot	7.6	28.5	19.7	14.1	81.7	85.7	83.7	10.4	25.2
	StoryTrans	55.9	47.4	22.0	7.1	82.1	83.7	82.9	19.9	68.1
	SC2	60.3	54.4	27.3	10.6	82.7	84.7	83.6	25.3	71.0
	Ours	62.1	56.7	29.1	15.4	84.3	85.6	84.9	30.9	72.6

Table 2: Automatic evaluation of our model with baselines on two datasets, where metrics of **Overall** are more valuable.

Methods	FT → LX			FT → JY		
	Sty.	Cont.	Flu.	Sty.	Cont.	Flu.
StyleLM	0.89	1.90	1.60	1.45	1.87	1.75
Reverse Attention	0.95	1.52	1.69	1.56	1.67	1.80
AugZero-Shot	1.35	1.20	2.46	1.43	1.63	2.65
StoryTrans	2.26	2.01	1.86	2.46	2.00	1.97
SC2	2.45	2.25	2.20	2.55	2.24	2.20
Ours	2.66	2.42	2.51	2.78	2.48	2.53

Table 3: Human evaluation of our model with baselines.

Methods	Acc _d	Acc _s	BLEU ₁	BLEU ₂	BS _p	BS _R	BS _{F1}	O ₁	O ₂
-SCP	46.5	48.9	10.6	6.5	74.5	80.4	77.3	17.4	60.0
-TSL	30.2	34.7	32.5	20.7	75.7	81.2	78.4	25.0	48.7
-FG	58.3	53.4	26.6	12.1	83.3	83.9	83.6	26.6	69.8
-loss ₁	56.4	50.4	24.3	10.2	82.6	83.3	82.9	24.0	68.4
-loss ₂	35.6	38.9	32.0	18.6	77.3	82.4	79.8	25.7	53.3
CeCoT	62.1	56.7	29.1	15.4	84.3	85.6	84.9	30.9	72.6

Table 4: Ablation analysis of our CeCoT.

expresses the superiority of our elaborately designed CoTs in promoting LLMs to improve long text style transfer.

Manual Evaluation We also conduct human evaluation to verify style transfer effect based on user experience. As shown in Table 3, we observe that: Although AugZero-Shot gains a slight advantage (0.12 improvement in FT → JY) in terms of fluency, our model obtains the best score in content retention and style transfer, showing an improvement of up to 0.41 than SC2 and StoryTrans. Therefore, our model has achieved significant advantages compared to other baseline methods.

Discussions

Ablation Study (RQ2) To verify the effectiveness of each module, we separate CeCoT: **-SCP**, **-TSL**, and **-FG** represent

the removal of LLMs in source content preserving, LLMs in target style learning module, and fusion network in CeCoT. $-\text{loss}_1$ and $-\text{loss}_2$ indicate that our model removes content preservation enhancing loss and style consistency verification loss, respectively. As shown in Table 4, we observe that:

- Effectiveness of constituent modules of our model. After removing LLMs (-SCP and -TSL), our model experiences varying degrees of performance degradation, showing a decrease of up to 23.9%. Meanwhile, -FG underperforms our model with at least 2.5% in all metrics. These fully demonstrate the effectiveness of each modules of our model.
- Effectiveness of designed losses. The removal of two loss modules has had a certain impact on the model performance, reflecting at least a 4.2% weakening in overall metrics. This illustrates that strengthening the constraints of content retention and style consistency has positive effects on improving the ability to transfer text styles.

Evaluation of Two Types of CoTs (RQ3) To verify the superiority of the two types of CoTs designed in CeCoT, we evaluate specific prompt strategies in TSL and SCP modules: **StdPrompt(TSL)** refers to modifying prompts of TSL module to standard prompts, i.e. “remove style and mine content”, while **StdPrompt(SCP)** denotes that replacing prompts of SCL module with standard prompts, i.e., “remove content semantics and explore style”. **SCP(-P1)** and **SCP(-P2)** indicate that we remove P’1 and P’2 from the source content preserving. **SCP(-P1)** and **SCP(-P2)** indicate that we will remove Prompt1 and Prompt2 from the source content preserving. **TSL(-Mode 1)**, **TSL(-Mode 2)**, **TSL(-Mode 3)**, **TSL(-Mode 4)**, and **TSL(-Mode 5)** respectively represent the removal of the corresponding different modes. The results are shown in Figure 4:

- When we replace prompts in TSL and SCL modules with standard prompts, it shows at least 3% performance weakening,




Figure 4: Comparison results of two types of CoTs (i.e., the impact of different steps in CoTs on model performance).




Figure 5: Performance comparison of different steps in CoTs.

which confirms the superiority of the prompt strategy constructed by our model. Furthermore, SCP (-P1) and SCP (-P2) achieve a performance degradation of at least 10.3% across all metrics, confirming the usefulness of the prompt strategy constructed by SCP for preserving source content.

- The removal of different prompt modes in TSL results in the model degradation, with a decline from 1.2% to 5.5% compared to our model. This not only confirms that prompt order behaviors affect the output quality of large models, but also reflects the effectiveness of considering different prompt behaviors to improve our model.

Evaluation of Each Step in CoTs (RQ4)

To evaluate the sustained effectiveness of each step in CoTs, we assess the changes in model performance (i.e., *Overall₂*) under different sample sizes, with a random sampling range from 100 to 800 samples. As shown in Figure 5, we can observe that: Overall, the performance of each step in the design thinking chain continues to improve with the increase of sample size, ranging from 2.2% to 5.6% in terms of overall performance, which confirms the effectiveness of our proposed steps in CoTs; In different modes, different modes will bring different model performance, among which mode 5 can promote the model to achieve greater performance improvement, while mode 1 only achieves the minimum performance. The influence of the execution order of these different thinking steps fully reflects the impact of human behavior on the results.

Case Study (RQ5)

To more intuitively show the outputs of CeCoT, we choose a specific case from Chinese dataset for verification. Figure 6 shows comparison of our model and the one without using CoTs (i.e., StdPrompt(TSL) and StdPrompt(SCP)), and we observe that:

- In StdPrompt(TSL), there was less text style rewriting (only transferring “slick” and “fabricated sophistry”) and more

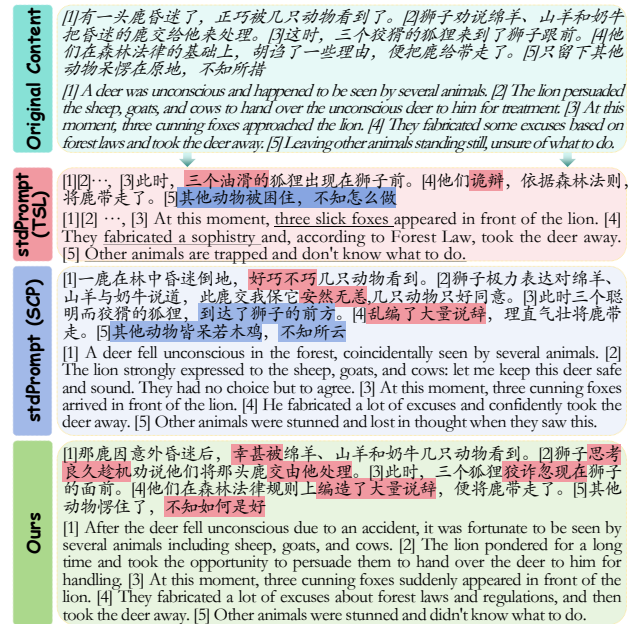


Figure 6: The visualization of sequences generated by CeCoT.

repetitive content (such as “taking the deer away according to the forest rule”), which is probably due to the lack of well-designed prompts, making it difficult to accurately identify the specific style information of the sequence.

- While transferring target text styles, StdPrompt(SCP) also loses some vital source content, such as entities “forest law” and “standing still”, resulting in semantic deviation. Our model outperforms the other two methods in preserving original content and style polarity. The red fragments show a more outstanding transfer effect, which further verifies the superiority of our model in long text style transfer from a visual perspective.

Conclusion

We constructed a novel long text style transfer framework, CeCoT, which designed two types of CoTs around cognition strategies into LLMs and promoted mutual combination between LLMs and small language models for improving source content retention and target style learning. Experimental results confirmed the superiority of CeCoT on two publicly competitive datasets. In the future, we will further design better prompts around the knowledge of cognitive psychology to improve this task.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grants (62572403, 62202381, and U22B2036), in part by Guang dong Basic and Applied Basic Research Foundation (2024A1515010087), and General Program of Chongqing Natural Science Foundation (No. CSTB2022NSCQ-MSX1284). We would like to thank the anonymous reviewers for their constructive comments.

References

- Chen, D.-Y.; Tennent, H.; and Hsu, C.-W. 2024. ArtAdapter: Text-to-Image Style Transfer using Multi-Level Style Encoder and Explicit Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8619–8628.
- Christ, L.; Amiriparian, S.; Kathan, A.; Müller, N.; König, A.; and Schuller, B. W. 2024. Towards Multimodal Prediction of Spontaneous Humor: A Novel Dataset and First Results. *IEEE Transactions on Affective Computing*.
- Dai, N.; Liang, J.; Qiu, X.; and Huang, X.-J. 2019. Style Transformer: Unpaired Text Style Transfer without Disentangled Latent Representation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5997–6007.
- Deng, M.; Wang, J.; Hsieh, C.-P.; Wang, Y.; Guo, H.; Shu, T.; Song, M.; Xing, E.; and Hu, Z. 2022. RLPrompt: Optimizing Discrete Text Prompts with Reinforcement Learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 3369–3391.
- Ge, S.; Park, T.; Zhu, J.-Y.; and Huang, J.-B. 2023. Expressive text-to-image generation with rich text. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7545–7556.
- Guan, J.; Feng, Z.; Chen, Y.; He, R.; Mao, X.; Fan, C.; and Huang, M. 2022. LOT: A story-centric benchmark for evaluating Chinese long text understanding and generation. *Transactions of the Association for Computational Linguistics*, 10: 434–451.
- Han, J.; Wang, Q.; Zhang, L.; Chen, W.; Song, Y.; and Mao, Z. 2023. Text style transfer with contrastive transfer pattern mining. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7914–7927.
- He, J.; Wang, X.; Neubig, G.; and Berg-Kirkpatrick, T. 2020. A Probabilistic Formulation of Unsupervised Text Style Transfer. In *International Conference on Learning Representations*.
- Hong, Z.; Wang, Z.; Shen, L.; Yao, Y.; Huang, Z.; Chen, S.; Yang, C.; Gong, M.; and Liu, T. 2024. Improving non-transferable representation learning by harnessing content and style. In *The Twelfth International Conference on Learning Representations*.
- Horvitz, Z.; Patel, A.; Callison-Burch, C.; Yu, Z.; and McKeown, K. 2024. ParaGuide: Guided Diffusion Paraphrasers for Plug-and-Play Textual Style Transfer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18216–18224.
- Hu, Z.; Lee, R. K.-W.; Aggarwal, C. C.; and Zhang, A. 2022. Text style transfer: A review and experimental evaluation. *ACM SIGKDD Explorations Newsletter*, 24(1): 14–45.
- Jin, D.; Jin, Z.; Hu, Z.; Vechtomova, O.; and Mihalcea, R. 2022. Deep learning for text style transfer: A survey. *Computational Linguistics*, 48(1): 155–205.
- Jing, L.; Song, X.; Lin, X.; Zhao, Z.; Zhou, W.; and Nie, L. 2023. Stylized data-to-text generation: A case study in the e-commerce domain. *ACM Transactions on Information Systems*, 42(1): 1–24.
- Lee, D.; Tian, Z.; Xue, L.; and Zhang, N. L. 2021. Enhancing Content Preservation in Text Style Transfer Using Reverse Attention and Conditional Layer Normalization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 93–102.
- Li, D.; Li, J.; and Hoi, S. 2024. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Advances in Neural Information Processing Systems*, 36.
- Li, Y.; Lin, Z.; Zhang, S.; Fu, Q.; Chen, B.; Lou, J.-G.; and Chen, W. 2023. Making Language Models Better Reasoners with Step-Aware Verifier. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5315–5333. Toronto, Canada: Association for Computational Linguistics.
- Lin, F.; Song, Y.; Tian, Z.; Chen, W.; Dong, D.; and Liu, B. 2024. Memory-enhanced text style transfer with dynamic style learning and calibration. *Science China Information Sciences*, 67(4): 142105.
- Lin, Z.; Tao, J.; Yuan, Y.; and Yao, A. C.-C. 2025. Existing LLMs Are Not Self-Consistent For Simple Tasks. *arXiv preprint arXiv:2506.18781*.
- Liu, M.; Zhu, A. H.; Maiti, P.; Thomopoulos, S. I.; Gadewar, S.; Chai, Y.; Kim, H.; Jahanshad, N.; and Initiative, A. D. N. 2023. Style transfer generative adversarial networks to harmonize multisite MRI to a single reference image to avoid overcorrection. *Human Brain Mapping*, 44(14): 4875–4892.
- Liu, P.; Wu, L.; Wang, L.; Guo, S.; and Liu, Y. 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)*, 15285–15295.
- Lv, Z.; Huang, Z.; Lu, M.; Yang, Y.; Tian, Z.; Niu, X.; and Li, D. 2024. Meta Learning Based Rumor Detection with Awareness of Social Bot. In *International Conference on Knowledge Science, Engineering and Management*, 135–151. Springer.
- Maia Polo, F.; Xu, R.; Weber, L.; Silva, M.; Bhardwaj, O.; Choshen, L.; de Oliveira, A.; Sun, Y.; and Yurochkin, M. 2024. Efficient multi-prompt evaluation of LLMs. *Advances in Neural Information Processing Systems*, 37: 22483–22512.

- Mostafazadeh, N.; Chambers, N.; He, X.; Parikh, D.; Batra, D.; Vanderwende, L.; Kohli, P.; and Allen, J. 2016. A corpus and cloze evaluation for deeper understanding of common-sense stories. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 839–849.
- Narasimhan, S.; Pooja, H.; Dey, S.; and Desarkar, M. S. 2023. On Text Style Transfer via Style-Aware Masked Language Models. In *Proceedings of the 16th International Natural Language Generation Conference*, 362–374.
- Pour, M. M. A.; Farinneya, P.; Bharadwaj, M.; Verma, N.; Pesaranghader, A.; and Sanner, S. 2023. COUNT: COntRastive UNlikelihoOd text style transfer for text detoxification. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 8658–8666.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Reif, E.; Ippolito, D.; Yuan, A.; Coenen, A.; Callison-Burch, C.; and Wei, J. 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)*, 837–848.
- Shang, M.; Li, P.; Fu, Z.; Bing, L.; Zhao, D.; Shi, S.; and Yan, R. 2019. Semi-supervised Text Style Transfer: Cross Projection in Latent Space. 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)*, 4937–4946.
- Sharma, D.; Dhiman, C.; and Kumar, D. 2024. Control With Style: Style Embedding-based Variational Autoencoder for Controlled Stylized Caption Generation Framework. *IEEE Transactions on Cognitive and Developmental Systems*.
- Shi, K.; Sun, X.; He, L.; Wang, D.; Li, Q.; and Xu, G. 2023. AMR-TST: Abstract meaning representation-based text style transfer. In *Findings of the Association for Computational Linguistics: ACL 2023*, 4231–4243.
- Shu, L.; Luo, L.; Hoskore, J.; Zhu, Y.; Liu, Y.; Tong, S.; Chen, J.; and Meng, L. 2024. RewritelM: An instruction-tuned large language model for text rewriting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18970–18980.
- Syed, B.; Verma, G.; Srinivasan, B. V.; Natarajan, A.; and Varma, V. 2020. Adapting language models for non-parallel author-stylized rewriting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 9008–9015.
- Wei, D.; Wu, Z.; Shang, H.; Li, Z.; Wang, M.; Jiabin, G.; Chen, X.; Zhengzhe, Y.; and Yang, H. 2023. Text Style Transfer Back-Translation. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wu, L.; Liu, P.; Yuan, Y.; Liu, S.; and Zhang, Y. 2023a. Context-aware style learning and content recovery networks for neural style transfer. *Information Processing & Management*, 60(3): 103265.
- Wu, L.; Liu, P.; Zhao, Y.; Wang, P.; and Zhang, Y. 2023b. Human cognition-based consistency inference networks for multi-modal fake news detection. *IEEE Transactions on Knowledge and Data Engineering*, 36(1): 211–225.
- Wu, L.; Long, Y.; Gao, C.; Wang, Z.; and Zhang, Y. 2023c. MFIR: Multimodal fusion and inconsistency reasoning for explainable fake news detection. *Information Fusion*, 100: 101944.
- Xiao, Y.; Wu, L.; Guo, J.; Li, J.; Zhang, M.; Qin, T.; and Liu, T.-y. 2023. A survey on non-autoregressive generation for neural machine translation and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10): 11407–11427.
- Xu, C.; Xu, Y.; Zhang, H.; Xu, X.; and He, S. 2024. DreamAnime: Learning Style-Identity Textual Disentanglement for Anime and Beyond. *IEEE Transactions on Visualization and Computer Graphics*.
- Zhang, Q.; Dong, J.; Chen, H.; Zha, D.; Yu, Z.; and Huang, X. 2024. Knowgpt: Knowledge graph based prompting for large language models. *Advances in Neural Information Processing Systems*, 37: 6052–6080.
- Zhang, Y.; Ge, T.; and Sun, X. 2020. Parallel Data Augmentation for Formality Style Transfer. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 3221–3228.
- Zhao, J.; Guan, Z.; Xu, C.; Zhao, W.; and Jiang, Y. 2024. SC2: Towards Enhancing Content Preservation and Style Consistency in Long Text Style Transfer. *arXiv preprint arXiv:2406.04578*.
- Zhu, A.; Yin, Z.; Iwana, B. K.; Zhou, X.; and Xiong, S. 2022. Text style transfer based on multi-factor disentanglement and mixture. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2430–2440.
- Zhu, K.; Tian, Z.; Wei, J.; Luo, R.; Song, Y.; and Mao, X. 2024a. StyleFlow: Disentangle latent representations via normalizing flow for unsupervised text style transfer. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 15384–15397.
- Zhu, X.; Guan, J.; Huang, M.; and Liu, J. 2023. StoryTrans: Non-Parallel Story Author-Style Transfer with Discourse Representations and Content Enhancing. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 14803–14819.
- Zhu, X.; Lei, Y.; Li, T.; Zhang, Y.; Zhou, H.; Lu, H.; and Xie, L. 2024b. Metts: Multilingual emotional text-to-speech by cross-speaker and cross-lingual emotion transfer. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.