

SC2: Towards Enhancing Content Preservation and Style Consistency in Long Text Style Transfer

Jie Zhao, Ziyu Guan*, Cai Xu, Wei Zhao and Yue Jiang

School of Computer Science and Technology, Xidian University, Xi'an, 710126, Chian
{jzhao1992@stu., zyguan@, cxu@, ywzhao@mail., 22031212489@stu.}xidian.edu.cn

Abstract

Text style transfer (TST) aims to vary the style polarity of text while preserving the semantic content. Although recent advancements have demonstrated remarkable progress in short TST, it remains a relatively straightforward task with limited practical applications. The more comprehensive long TST task presents two challenges: (1) existing methods encounter difficulties in accurately evaluating content attributes in multiple words, leading to content degradation; (2) the conventional vanilla style classifier loss encounters obstacles in maintaining consistent style across multiple generated sentences.

In this paper, we propose a novel method SC2, where a multilayer Joint Style-Content Weighed (JSCW) module and a Style Consistency loss are designed to address the two issues. The JSCW simultaneously assesses the amounts of style and content attributes within a token, aiming to acquire a lossless content representation and thereby enhancing content preservation. The multiple JSCW layers further progressively refine content representations. We design a style consistency loss to ensure the generated multiple sentences consistently reflect the target style polarity. Moreover, we incorporate a denoising non-autoregressive decoder to accelerate the training. We conduct plentiful experiments and the results show significant improvements of SC2 over competitive baselines. Our code: <https://github.com/jiezhao6/SC2>.

1 Introduction

Text style transfer (TST) aims to generate a text exhibiting a desired style based on the source text (e.g., negative \rightarrow positive), while endeavoring to faithfully preserve the semantic content. The applications of TST cover a wide range of user-centric natural language generation tasks, such as personalized dialogue systems (Gao et al., 2019), educa-

*Corresponding author.




Figure 1: Comparisons of content learning between existing approaches and the proposed method: (a) evaluating the relevance between text x and its style; (b) evaluating the relevance between source text x_s and target text x_t ; and (c) joint evaluating the relevance between text x and its style as well as content.

tional platforms (Wang et al., 2019), and writing assistants (Syed et al., 2020).

Given the scarcity of parallel data (i.e., text pairs conveying the same content but differing in styles) and the labor-intensive nature of annotating such pairs, existing research has predominantly focused on unsupervised TST. Recent contributions in this domain, including studies by (Zhou et al., 2020; Lee et al., 2021; Huang et al., 2021; Suzgun et al., 2022; Ramesh Kashyap et al., 2022; Han et al., 2023), have demonstrated significant progress. Despite notable success, these works primarily concentrate on the transfer of a single sentence, which we call short TST. This is a relatively simple task and is difficult to apply to complex scenarios, i.e., long TST such as transferring news articles and novels. It is challenging to achieve desirable content preservation and style consistency across multiple sentences for these methods. In a very recent study, Zhu et al. (2023) first concentrated on the long TST task and proposed StoryTrans, which learns content representations and fills stylistic tokens in separate stages. While somewhat effective, challenges persist in preserving content and ensuring style consistency.

Content Preservation: The critical factor for preserving content lies in accurately assessing the amount of content attribute (CA) within a token to improve the content representation. In tradi-

tional approaches, the content learning primarily involves explicitly (Wu et al., 2019; Huang et al., 2021) or implicitly (Zhou et al., 2020; Lee et al., 2021) removing style tokens. In these processes, they evaluate the relevance between text and style, which solely consider the amount of style attribute (SA) within a token and neglect the CA amount (Figure 1 (a)). This results in tokens with both strong style and content attributes, such as “euphonious”¹, potentially receiving higher SA scores, making them more drastic to be removed and consequently leading to content degradation. On the other hand, in StoryTrans, the disentanglement of content from style is achieved by encouraging texts with distinct styles to be close together in the content space (Figure 1 (b)). However, owing to the non-parallel nature of the data, this unavoidably results in a loss of content information.

Style Consistency: To control the style polarity of generated text, existing methods employ a style discriminator that operates on the entire output text. However, in the context of long TST, it becomes challenging for the discriminator to ensure that the style of the generated multiple sentences is consistent. As a result, some generated sentences might exhibit strong target style polarity while others remain weak, which creates a less reader-friendly outcome. Therefore, we argue that maintaining style (polarity) consistency across multiple sentences is crucial.

To tackle the above issues, we propose a novel approach aimed at enhancing content preservation and maintaining style consistency in long TST. Our approach achieves these objectives by carefully designing a multilayer Joint Style-Content Weigher (JSCW) module and a Style Consistency loss, thus we call it SC2. (1) The JSCW utilizes the convolution operation to measure the SA amount within the center token. Simultaneously, it assesses the CA amount by computing and integrating the content relevance of a token across all sentences. Then by normalizing these two amounts and weighting the CA amount to the tokens’ representations, we obtain preliminary content representation. Furthermore, we employ multiple JSCW layers to progressively refine content representations. Finally, we fuse the target style information to content representations and feed them to the decoder to generate target text. (2) For the other challenge, we design

a contrastive learning-based style consistency loss. It brings each generated sentence closer to the previously generated sentences or target sentences in the corpus, while farther away from the sentences in the source text.

Additionally, within the unsupervised long TST setting, directly employing an autoregressive (AR) decoder substantially slows down the training since the masked self-attention technique (Vaswani et al., 2017) cannot be exploited. Hence, drawing inspiration from the research on non-AR (NAR) text generation (Gu et al., 2017; Huang et al., 2021), we incorporate an auxiliary denoising NAR decoder. It parallelly generates pseudo-target texts, which are then fed into the AR decoder to accelerate the training process. Our main contributions are summarized as follows:

- We propose to explicitly and simultaneously assess the SA and CA amounts within tokens to learn lossless content representations and show that this idea can significantly improve content preservation in long TST.
- We first propose the concept of style consistency for long TST, and design a corresponding loss to encourage the generated text to consistently maintain style polarity across multiple sentences.
- Extensive experiments are conducted on both Chinese and English datasets to verify the effectiveness of the proposed SC2. The results demonstrate significant improvements over competitive baselines.

2 Related Work

Text Style Transfer. Recently, there has been considerable research attention on TST. The predominant of current efforts lies in unsupervised TST since acquiring parallel data is labor-intensive, even impossible to crowdsource for some styles (Jin et al., 2022). One line explicitly disentangles text into content and style representations, and then incorporates the target style information into the content-based generation process. Li et al. (2018); Wu et al. (2019); Lee et al. (2021) proposed using frequency ratio or attention scores to measure the SAs within tokens, and then entirely or partially removing the stylistic tokens. Tokpo and Caldera (2022) introduced the LIME explainer (Ribeiro et al., 2016) to identify biased (style) tokens and then masked them for non-biased text generation. John et al. (2019) and Yi et al. (2020) employed

¹“Euphonious” is strongly associated with a positive sentiment style and concurrently indicates content related to music.

the adversarial loss to learn the style-agnostic content representation. Krishna et al. (2020) employed a pretrained paraphrase model to strip away the style information and generated target text conditioning on the style-dependent inverse paraphrase model. The other line uses entangled representations to perform TST. Jin et al. (2019) and Xu et al. (2023) created pseudo-parallel data and then trained a *Seq2Seq* model. Syed et al. (2020) trained a generator for each target style. Riley et al. (2021) extracted a style vector from the adjacent sentence and used it to condition the decoder. Our work falls into the first line. Nevertheless, what distinguishes our work is the explicit and simultaneous measurement of SA and CA amounts within tokens.

The aforementioned works all handle the task of Single-Sentence (SS) transfer, conforming to an SS-input and SS-output paradigm. Differing from these approaches, Cheng et al. (2020) introduced the use of context sentences as additional information to preserve topical coherence between the generated sentence and its surrounding context. Although it employs multiple sentences as input, fundamentally, this work still falls under the setting of SS transfer. Zhu et al. (2023) introduced the task of long TST and proposed a two-stage training and inference framework to learn content representations by aligning source texts and target texts. We follow this task, aiming to enhance content preservation and style consistency.

Non-AutoRegressive (NAR) Text Generation. Gu et al. (2017) first introduced the NAR concept in machine translation which produces output tokens in parallel. Subsequent works introduce the NAR decoder into other domains, such as dialogue generation (Han et al., 2020), speech synthesis (Peng et al., 2020), and short TST (Huang et al., 2021). Inspired by these works, we incorporate an auxiliary denoising NAR decoder to accelerate the training of SC2.

3 Approach

3.1 Problem Definition

Let $x = (x_1, x_2, \dots, x_{T_x})$ be a source text with multiple sentences and $s \in \mathcal{S}$ be its style label, where \mathcal{S} is the style label set. Our goal is to develop a model f_θ that can generate a text $y = (y_1, y_2, \dots, y_{T_y})$ exhibiting a desired target style \hat{s} ($\hat{s} \in \mathcal{S}$ and $\hat{s} \neq s$) for a given text x . The generated text y shares the same semantic content

with x . Subsequently, we use bold face lower/upper case letters to denote vectors/matrices respectively.

3.2 Model Architecture

As shown in Figure 2, the proposed SC2 fits within the encoder-decoder framework, and mainly consists of multilayer JSCW module, style fusion module, and denoising NAR decoder. After the encoding of original tokens, the JSCW evaluates the SA and CA amounts simultaneously and learns elementary content representations. And we propose to progressively refine content representations by multiple JSCW layers. Subsequently, the style fusion module combines the style information and style-agnostic content representations to provide signals of transfer direction. And the results are fed into the decoder. Additionally, the denoising NAR decoder generates pseudo-target text to accelerate the training. For simplicity, our discussion in the following section focuses on the transfer direction of $s \rightarrow \hat{s}$. Nevertheless, our approach can be easily extended to handle transfers among multiple styles by incorporating multiple style embeddings.

3.2.1 Multilayer Joint Style-Content Weigher

Assuming that the length of x is n and it consists of m sentences, we insert a special token $\langle \text{SEN} \rangle$ at the end of each sentence, inspired by (Lee et al., 2020; Zhu et al., 2023). Consequently, we feed x into the encoder, obtaining $n + m$ hidden states (i.e., $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^{n+m}$). For convenience, we use \mathbf{x}_i ($i \in [1, n]$) and \mathbf{e}_j ($j \in [1, m]$) to denote the representations of i -th word token and j -th sentence token respectively. Inspired by the work of (Lee et al., 2021), we assume that the representation of a word token is comprised of style information $\alpha_i \mathbf{x}_i$ and content information $\beta_i \mathbf{x}_i$, where α_i is the SA score (amount), β_i is the CA score (amount), and $\alpha_i + \beta_i = 1$. The essence of the JSCW lies in discerning the values of α_i and β_i simultaneously to effectively disentangle the content from style for a word token.

Style Weigher. We introduce a learnable style embedding $\mathbf{s} \in \mathbb{R}^d$ to assist the computation of SA score, which will also serve as a style control signal. A straightforward idea is to directly measure the SA score based on the current word token (e.g., $\mathbf{x}_i \mathbf{W}_s \mathbf{s}$, where \mathbf{W}_s is a bilinear term enabling flexible estimation). However, individual words sometimes do not reflect their association with a specific style. Therefore, we employ a con-

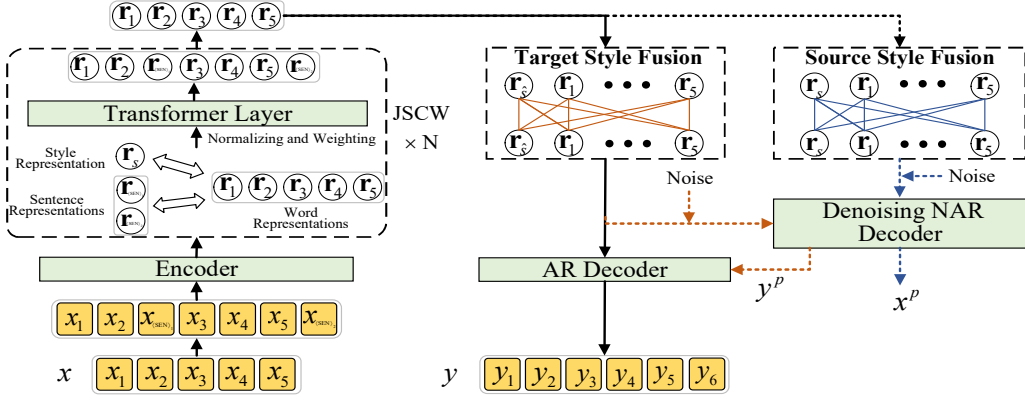


Figure 2: The framework of SC2. We show an example of transferring x (long source text) to y (long generated text). Dashed lines with arrows indicate that these processes occur only during the training phase.

volution operation applied to a window of $2h + 1$ tokens, generating a new feature used to assess the SA score of the center token. Formally, the non-normalizing SA score of a word token is computed as follows:

$$\tilde{\beta}_i = \sigma(\mathbf{W}_{conv}\mathbf{x}_{i-h:i+h} + \mathbf{b}_{conv})\mathbf{W}_s\mathbf{s}, \quad (1)$$

where $\sigma(x) = \max(x, 0)$ is the rectifier activation function, and \mathbf{W}_{conv} and \mathbf{b}_{conv} are convolution matrix and bias vector respectively.

Content Weigher. The content of a word token varies across diverse contexts, exemplified by the polysemy. Hence, employing a single representation akin to style embedding is inadvisable for assessing a word’s CA score. On the other hand, in a multi-sentence scenario, a word’s semantics may be interconnected with or influenced by sentences beyond its immediate sentence. To address these challenges, we compute the content-relevance scores between a word and all sentences, subsequently attentionally aggregating them as the non-normalizing CA score. Specifically, we employ the following equation to achieve this objective:

$$\mu_{i,j} = \frac{\exp(\mathbf{x}_i\mathbf{W}_c\mathbf{e}_j)}{\sum_{j=1}^m \exp(\mathbf{x}_i\mathbf{W}_c\mathbf{e}_j)}, \quad (2)$$

$$\tilde{\alpha}_i = \sum_{j=1}^m \mu_{i,j}(\mathbf{x}_i\mathbf{W}_c\mathbf{e}_j),$$

where \mathbf{W}_c is a bilinear term.

After obtaining the individual SA score and CA score, we normalize them as: $[\alpha_i, \beta_i] = \text{softmax}([\tilde{\alpha}_i, \tilde{\beta}_i])$. Finally, the content representation of a word token is the weighted \mathbf{x}_i :

$$\mathbf{x}_i^c = \beta_i\mathbf{x}_i. \quad (3)$$

We define the above disentanglement process (Eqs. (1)~(3)) as a JSCW layer. However, the assessment of the CA score might introduce bias due to the potential inclusion of style information in the sentence representation \mathbf{e}_j . To address this concern, we modify the JSCW layer by incorporating a subsequent Transformer layer (Vaswani et al., 2017). Leveraging multi-head self-attention mechanisms, this Transformer layer utilizes elementary content representations obtained from the original JSCW layer to learn new sentence representations. And we employ multiple JSCW layers, which progressively diminishes the style information within sentence representations and enables a more precise assessment of CA score.

3.2.2 Style Fusion Module

After extracting the words’ content representations, we integrate the style information into them to provide transfer signals. We concatenate the target style embedding with content representations and employ multiple Transformer layers (MTL) as the fusion network to learn style-aware representations. Formally, the process is formulated as:

$$\mathbf{Z}_{\hat{s}} = \text{MTL}(Q = K = V = \tilde{\mathbf{Z}}_{\hat{s}}), \quad (4)$$

where $\tilde{\mathbf{Z}}_{\hat{s}} = [\hat{s}, \mathbf{x}_1^c, \mathbf{x}_1^c, \dots, \mathbf{x}_n^c] \in \mathbb{R}^{d \times (n+1)}$, and Q , K and V respectively represent the query, key, and value within the Transformer layer. In summary, we denote the encoder, multiple JSCW layers, and the style fusion module as a network F , i.e., $\mathbf{Z}_{\hat{s}} = F(x, \hat{s})$.

3.2.3 Decoder

We employ an autoregressive (AR) decoder to generate target text. The decoder accesses style-

aware words' representations $\mathbf{Z}_{\hat{s}}$ through the cross-attention layer. The AR decoder is defined as:

$$P_{D_{AR}}(y | \mathbf{Z}_{\hat{s}}) = \prod_{t=1}^{T_y} P_{AR}(y_t | y_{<t}, \mathbf{Z}_{\hat{s}}) \quad (5)$$

However, a practical issue arises in terms of time consumption during training. In the unsupervised setting, the decoder cannot leverage the masked self-attention technique (Vaswani et al., 2017) to efficiently train the network and can only generate target tokens one by one. For long TST, it significantly increases the training time. Hence, inspired by recent research on non-AR (NAR) text generation (Gu et al., 2017; Huang et al., 2021), we propose a simple but effective denoising NAR decoder², which generates pseudo-target tokens in parallel, to accelerate the training. The generated text y^p will be fed into the AR decoder to exploit the masked self-attention to generate target-styled text. The denoising NAR decoder is defined as:

$$P_{D_{NAR}}(y^p | \mathbf{Z}_{\hat{s}}, \mathbf{Z}_{\hat{s}}^n) = \prod_{t=1}^{T_y} P_{D_{NAR}}(y_t^p | \mathbf{Z}_{\hat{s}}, \mathbf{Z}_{\hat{s}}^n), \quad (6)$$

where $\mathbf{Z}_{\hat{s}}^n$ is a corrupted version of $\mathbf{Z}_{\hat{s}}$. Specifically, we randomly swap a token with one of its $2k$ neighboring tokens, where the swapping probability is p . We also employ a source style embedding in the fusion module and use the NAR to obtain x^p . During training, the NAR decoder will learn to acquire the capability of reconstructing source texts of various styles. It serves as providing pseudo-labels to the AR decoder, which is somewhat similar to some pseudo-parallel data construction approaches (Jin et al., 2019). The key difference lies in the end-to-end training for both denoising NAR and AR decoders.

3.3 Training Objectives

In this section, we elaborate on the training objectives employed in SC2.

Style-Oriented Objectives. To constrain the style polarity of the generated text, we employ a document-level objective and design a sentence-level objective. We pre-train two style classifiers C_{docu} and C_{sent} using standard cross-entropy loss on the entire long text and segmented sentences, respectively. The weights of these two classifiers are frozen during the training of SC2.

²The NAR decoder will be removed during inference since a simple NAR generator may not perform very well.

We first use the widely employed vanilla style transfer loss based on the document-level classifier:

$$\mathcal{L}_{sty_d} = -\mathbb{E}_{y \sim D_{AR}} [\log P_{C_{docu}}(\hat{s} | y)].$$

Secondly, we consider maintaining the consistency of style polarity across multiple sentences. It is challenging for the document-level style classifier to achieve this goal. Therefore, we derive the following contrastive style consistency loss:

$$\mathcal{L}_{sty_s} = -\frac{1}{\hat{m}} \sum_{j=1}^{\hat{m}} \sum_{j'=1}^{j-1} \log \left(\frac{\exp((\mathbf{o}_j^y \cdot \mathbf{o}_{j'}^y) / \tau)}{\sum_{j''=1}^{\hat{m}} \exp((\mathbf{o}_j^y \cdot \mathbf{o}_{j''}^x) / \tau)} \right),$$

where τ is a temperature parameter, \hat{m} (m) is the number of sentences in y (x), j and j' (j'') are the sentence indexes in y (x), and \mathbf{o} is the output of the last layer of the sentence-level classifier. When $j = 1$, the \mathbf{o}_j^y is obtained by randomly sampling a target sentence from the training corpus. This loss makes the j -th generated sentence maintain stylistic similarity with sentences $1 \sim j - 1$ and target sentences in corpus, while diverging from the source sentences. This ensures both the style consistency and polarity of the generated text.

The overall loss for controlling style polarity is defined as:

$$\mathcal{L}_{sty} = \mathcal{L}_{sty_d} + \mathcal{L}_{sty_s}. \quad (7)$$

Content-Oriented Objectives. We employ a reconstruction loss containing self- and cycle-reconstruction, which is formulated as:

$$\mathcal{L}_{rec} = -\mathbb{E}_{x \sim P_x} [\log P_{D_{AR}}(x | F(x, s))] - \mathbb{E}_{x \sim P_x} [\log P_{D_{AR}}(x | F(y, s))]. \quad (8)$$

NAR Decoder-Oriented Objective. The denoising NAR decoder is trained to reconstruct the source text, and we use the following loss:

$$\mathcal{L}_{NAR} = -\mathbb{E}_{x \sim P_x} [\log P_{D_{NAR}}(x^p | \mathbf{Z}_s, \mathbf{Z}_s^n)], \quad (9)$$

where $\mathbf{Z}_s = F(x, s)$.

Disentanglement-Oriented Objective. We empirically observe that during the early stage of training, the distribution formed by $[\alpha_i, \beta_i]$ tends to align with the Bernoulli distribution. This undermines the disentanglement learning. We design a penalizing term to mitigate this issue.

$$\mathcal{L}_{dis} = \sum_{i=1}^n \max(0, \epsilon - \mathcal{H}([\alpha_i, \beta_i])), \quad (10)$$

where $\mathcal{H}(\cdot)$ represents the entropy of a distribution, and ϵ is a threshold.

In summary, the overall loss of SC2 is defined to balance the above losses:

$$\mathcal{L} = \mathcal{L}_{rec} + \lambda_1 \mathcal{L}_{sty} + \lambda_2 \mathcal{L}_{NAR} + \lambda_3 \mathcal{L}_{dis}. \quad (11)$$

An training issue arises with Eqs. (7)~(8), where discrete tokens in y prevent the back-propagation of gradients. We address this issue by employing the soft-sampling strategy (Dai et al., 2019).

4 Experiments

4.1 Dataset

Following the work of Zhu et al. (2023), we employ the stylized story dataset to evaluate the performance of SC2. The dataset encompasses sub-datasets in both Chinese and English. Important statistics are summarized in Table 1.

The **Chinese** sub-dataset comprises three styles of texts: LuXun (LX)’s novels, JinYong (JY)’s novels³, and Fairy Tale (FT), totaling approximately 7.5k samples for training. The corresponding tasks involve transforming an FT into text with LX’s style (FT \rightarrow LX) or JY’s style (FT \rightarrow JY). The **English** sub-dataset contains everyday stories and fragments from Shakespeare (SP)’s plays. There are approximately 2.3k training samples. The aim is to transfer an everyday story (ER) into the style of Shakespeare’s plays (ER \rightarrow SP).

4.2 Baselines

We employ **StoryTrans** (Zhu et al., 2023), a two-stage framework aiming to learn content representations and filling stylistic tokens for long TST, as the strongest baseline in our experiments. Additionally, we compare SC2 with the following baselines designed for short TST:

Style Transformer (Dai et al., 2019) equips the power of attention mechanism within the Transformer (Vaswani et al., 2017) architecture to learn entangled text representations.

StyleLM (Syed et al., 2020) fine-tunes a pre-trained language model (PLM) on a corrupted version of the style-specific corpus for each style.

Reverse Attention (Lee et al., 2021) employs reverse attention scores, which assess each token’s contribution to the style classification, to implicitly remove the style information in tokens.

³LuXun was a realism novelist and JinYong was known for his martial arts novels.

Dataset	Chinese			English		
	Style	Size	Avg L	Style	Size	Avg L
Training	JY	2,964	344	SP	1,161	71
	LX	3,036	168	RS	1,161	49
	FT	1,456	175	-	-	-
Val.	FT	242	175	RS	290	48
Test	FT	729	176	RS	290	50

Table 1: Statistics of the employed dataset.

AugZero-Shot (Reif et al., 2022) is an augmented zero-shot prompting method, which prompts the PLM with diverse sentence-rewriting examples.

4.3 Evaluation Metrics

Following previous studies of (Syed et al., 2020; Lee et al., 2021; Reif et al., 2022; Zhu et al., 2023), we employ both automatic and human evaluations.

4.3.1 Automatic Evaluation

Style Transfer Accuracy: To evaluate the style polarity of the generated text, we separately pre-trained document- and sentence-level style classifiers for each language. The style classifier utilizes the same architecture as the encoder of SC2 and incorporates a mean-pooling layer and a classification layer. In the case of Chinese, our style classifier achieves accuracies of 100% and 98% for document- and sentence-level test corpus, respectively. For English, these accuracies stand at 100% and 100%. Consequently, we employ two metrics, denoted as Acc_d and Acc_s , to measure document- and sentence-level style transfer accuracy, respectively.

Content Preservation: Firstly, we employ $BLEU_1$ and $BLEU_2$ (Papineni et al., 2002) to evaluate the content preservation through n-gram overlap. Secondly, we employ the BERT Score (Zhang et al., 2020) metric. It leverages PLMs to assess the semantic similarity of tokens and provides a more human-like evaluation. Precision (BS_p), recall (BS_r), and F1 (BS_{F1}) scores are reported.

Overall: Recognizing that content preservation and style transfer as two major challenges within the TST task, we report two metrics for the overall performance (Lee et al., 2021; Zhu et al., 2023) and use them as the main metrics. Specifically, we employ the geometric means of Acc_d and $BLEU(G-BL_d)$, and Acc_d and $BS_{F1}(G-BS_d)$.

Zhu et al. (2023) have noted that the perplexity based on PLMs is unreliable for fluency evaluation on this dataset. Consequently, we employ the man-

Task	Methods	Acc _d	Acc _s	BLEU ₁	BLEU ₂	BS _p	BS _R	BS _{F1}	G-BL _d	G-BS _d
FT → LX	Style Transformer	0.1	1.7	71.7	64.4	91.7	92.6	92.1	3.1	3.6
	StyleLM	0.1	6.1	71.2	61.1	90.6	92.2	91.4	3.0	3.5
	Reverse Attention	6.4	28.0	17.3	4.8	61.4	58.8	60.0	8.4	19.7
	AugZero-Shot	5.9	38.5	11.7	4.1	56.5	59.8	58.0	6.8	18.5
	StoryTrans	39.6	52.0	25.7	9.6	63.0	63.9	63.4	26.4	50.1
	Ours	41.6	70.5	27.5	11.9	61.6	65.5	63.5	28.6	51.4
FT → JY	Style Transformer	0.8	3.4	71.8	64.5	91.6	92.6	92.1	7.5	8.7
	StyleLM	1.2	0.9	70.6	60.3	90.3	92.1	91.1	9.0	10.6
	Reverse Attention	63.8	29.1	14.8	3.9	59.1	60.1	59.6	24.4	61.6
	AugZero-Shot	56.0	5.0	12.0	4.2	56.4	59.5	57.8	21.2	56.9
	StoryTrans	70.6	25.1	22.3	7.9	61.2	65.5	63.2	32.7	66.8
	Ours	74.9	41.5	26.2	12.5	64.2	67.7	65.9	38.1	70.2
ER → SP	Style Transformer	1.4	5.9	87.0	82.4	97.1	97.5	97.3	10.8	11.6
	StyleLM	0.3	2.7	83.6	80.0	96.0	97.3	96.7	5.3	5.8
	Reverse Attention	43.8	45.7	8.7	2.6	73.6	81.9	77.4	15.8	58.2
	AugZero-Shot	7.6	28.5	19.7	14.1	81.7	85.7	83.7	11.3	25.2
	StoryTrans	55.9	47.4	22.0	7.1	82.1	83.7	82.9	28.5	68.0
	Ours	60.3	54.4	27.3	10.6	82.7	84.7	83.6	33.8	71.0

Table 2: Automatic evaluation results. Higher values are considered desirable for all metrics.

ual evaluation to assess the fluency of the generated text.

4.3.2 Human Evaluation

We randomly select 100 samples from the Chinese test set and generate target texts with LX style and JY style for baselines and proposed SC2. Given the source text and the target style, three Chinese native speakers are assigned the task of scoring for style polarity (Sty), content preservation (Con), and fluency (Flu), using a scale of 1 (very bad) to 3 (very good). We report the average scores across the three annotators as our final results.

4.4 Implementation Details

We employ LongLM_{SMALL} (Guan et al., 2022) and T5_{SMALL} (Raffel et al., 2020) as backbone models for Chinese and English data respectively. For the NAR decoder, we use 6 Transformer layers. The number of JSCW layers is set to 3. For Chinese data, we set $\lambda_1/\lambda_2/\lambda_3$ to 0.05/1/1, p/k for the NAR decoder to 0.3/2, the threshold ϵ to 0.15, the batch size to 2, and the learning rate to 5×10^{-5} (Adam optimizer). For English data, the hyperparameters are the same except that $\lambda_1/p/k$ are set to 0.001/0.1/1. The experiments are conducted on one RTX 4060Ti GPU (16G) and the training of SC2 approximately takes 7~9 hours.

4.5 Results

Table 2 shows the experimental results using automatic metrics. We can obtain the following observations: (1) the methods (StoryTrans, SC2) that are

carefully designed for long TST usually perform better than traditional methods (Style Transformer, StyleLM, Reverse Attention, and AugZero-Shot). This demonstrates a specialized model is essential for long TST. (2) The Style Transformer and StyleLM exhibit significant shortcomings in controlling style polarity, while the high BLEU scores and Bert scores suggest substantial copying of the source text. This may be attributed to the difficulty of transferring entangled text representations to target tokens. (3) Compared to StoryTrans, the proposed SC2 demonstrates significant advantages in both style polarity control and content preservation, particularly concerning the control of sentence style. (4) The difficulties of imitating the writing styles of different authors vary. The similar BLEU metrics for SC2 suggest consistent surface-level content preservation across three sub-tasks. However, significant differences are observed in the control of document- and sentence-level style polarities for three sub-tasks. (5) For the overall metrics, we use t-test with significance level 0.05 to test the significance of performance difference. Results show that SC2 significantly outperforms all the baselines.

As for the human evaluations, the results are shown in Table 3. Here we neglect the baselines of Style Transformer and StyleLM, which show poor performance in automatic evaluation. For content preservation and style polarity, the results from manual evaluation align with those from automatic evaluation, with our proposed SC2 achieving the highest scores. Additionally, annotators provided

Source Text	<p>1一天晚上，一只觅食的大灰狼溜进了树林里，一不留神踩到了串喇叭花。2喇叭花告诉了树林里的小动物们大灰狼来了的消息。3小动物们做足了准备，他们各显神通赶跑了大灰狼。4大灰狼再也不敢踏进树林觅食了，树林又恢复了平日的欢乐与宁静。</p> <p>1 One evening, a foraging gray wolf sneaked into the woods and accidentally stepped on a cluster of trumpet flowers. 2 The trumpet flowers informed the small animals in the forest about the arrival of the gray wolf. 3 The small animals made thorough preparations, showcased their unique abilities, and successfully drove away the gray wolf. 4 The gray wolf no longer dared to step into the woods to look for food, and the woods returned to usual joy and tranquility.</p>
StoryTrans	<p>1待狼群来到树林，瞧见一大灰狼在宁静的一头燕时，一不留神拔出了刺花的喇叭。2喇叭喇的骂道：“到了半年，还是没见到牛人家大灰狗。大伙儿的小心在谷底溜一圈再回去，别给狼吞没食的动物觅来。”3说罢大狼转身欲逃，林前树下的乌云慢慢散去，待狼们到达树边，宁静告知罗刹消息的动向。4小黑狼大人等了大半个时辰，才道：“狼大的足见与众不同，上来送吃的了。大人留心好生留心便是。我们大人走不了的，去瞧瞧动物的回头再说。”</p> <p>1 When the wolfpack arrived in the woods, they saw a gray wolf (at a tranquil swallow) and inadvertently pulled out a thorny trumpet flower. 2 The trumpet flower scolded, “Half a year has passed, and still no sign of the gray dog from the cow family. Everyone, be cautious, take a round in the valley before going back, and don’t let the wolves come.” 3 After saying this, the big wolf turned to escape. The dark clouds under the trees slowly dispersed. When wolves reached the trees, the “tranquility” told Rakshasa of the news. 4 The adult black wolf waited for nearly half an hour, and said, “The wolf leader’s behavior is extraordinary; he’s coming up to offer food. Pay attention, and be vigilant. We adults cannot leave, let’s go and see animals.”</p>
Ours	<p>1傍晚时分，一只疲累的狼来不及回避，一不小心摔进了树林里。2只听得叮叮当当两声响，兔儿们大哭起来。松鼠赶紧拉住马缰，急奔而去。3灰狼叫道：“大家干甚么去了？咱们做过手脚和神通广大的动物们呀！”众狼听道这里，都是兴高采烈、欢喜与宁静。4他们用火把烤熟了，再也不敢移动半步，更不用提防自己修炼内功，立即跃出树林，一溜烟逃走了。</p> <p>1 In the evening, a weary wolf, unable to evade in time, accidentally fell into the woods. 2 Only hearing two clinking sounds, the rabbits burst into tears. The squirrel quickly grabbed the reins and rushed away. 3 The gray wolf shouted, “What is everyone doing? We have manipulated and outwitted the animals with extraordinary abilities!” Upon hearing this, the wolves were all jubilant and filled with joy and tranquility. 4 They roasted the torches and no longer dared to move a step, let alone practice their internal power. Immediately, they leaped out of the woods and fled in a flash.</p>

Table 4: Generated texts with the style of JY. The number preceding each sentence in the generated texts corresponds to the respective sentence in the source text in terms of semantics. Underlined sentences or phrases denote inserted contents tailored to match the target style. We use the corresponding colors of texts between the source and generated texts to emphasize the rewritten content.

Task	Methods	Sty	Con	Flu
FT → LX	Reverse Attention	0.72	1.40	1.49
	AugZero-Shot	0.85	1.03	2.35
	StoryTrans	1.87	1.73	1.50
	Ours	2.29	1.96	1.85
FT → JY	Reverse Attention	1.42	1.50	1.55
	AugZero-Shot	1.30	1.49	2.56
	StoryTrans	2.28	1.70	1.61
	Ours	2.63	2.15	2.02

Table 3: Human evaluation results.

evaluations on fluency, where our proposed method significantly outperformed the StoryTrans. While the AugZero-Shot excels in fluency, its style polarity and content preservation are both too weak. Results indicate the superiority of the proposed SC2 across various metrics.

4.6 Analysis

Case Study: Table 4 illustrates examples of generated texts with the target style of JY by the best baseline StoryTrans and SC2. In the text generated by StoryTrans, a substantial portion exhibits low relevance with the source text, indicating poor content retention. For style polarity, we observe that SC2 rewrites more of the source text with the target style and inserts more target-style phrases. In terms of fluency, StoryTrans performs less satisfactorily, as exemplified by the generation of peculiar phrases such as “在宁静的一头燕时” (“at a tranquil swallow”). It is evident that, in terms




Figure 3: Visualization of Chinese and English test dataset on original space (blue dots), content space (red dots), and fused space (green dots) using t-SNE.

of content preservation, style polarity, and fluency, SC2 outperforms StoryTrans significantly.

Ablation Study: We conduct ablation experiments on the English dataset to demonstrate the effectiveness of two crucial components in SC2. The results are presented in Table 5, where “-” indicates that we remove the corresponding component. Particularly, “-JSCW” refers to the removal of the calculation of CA score, forming a new computational process as $\beta_i = \text{sigmoid}(\beta_i)$, $\alpha_i = 1 - \beta_i$. we can see that the performance of these variants drops apparently, which confirms the effectiveness of these proposed components in SC2.

Visualization of the Latent Space: In Figure 3, we employ the t-SNE to visualize the original, content and fused latent space of Chinese and English test datasets. For the Chinese dataset, we use “JY” as the target style. The representation of an input text x is the average of all tokens. It is evident

Methods	Acc _d	Acc _s	BLEU ₁	BLEU ₂	BS _P	BS _R	BS _{F1}	G-BL _d	G-BS _d
Ours	60.3	54.4	27.3	10.6	82.7	84.7	83.6	33.8	71.0
- JSCW	57.6	58.1	20.5	6.0	81.8	84.1	82.9	27.6	69.1
- \mathcal{L}_{sty_s}	41.7	46.2	30.2	12.7	83.6	85.3	84.4	29.9	59.4

Table 5: Ablation study results on English dataset.

that the entangled features, disentangled content features, and fused stylistic features exhibit three distinct distributions. This demonstrates that our approach can effectively learn content information from the raw data and integrate the target style information into it.

4.7 Extrinsic Experiment

We further conduct an extrinsic experiment to show that the proposed SC2 can function as a data aug-menter, improving the performance of data-scarce natural language processing tasks. Specifically, we concentrate on a charge prediction task within the legal domain. Based on the factual details of legal cases, it aims at predicting the ultimate charges convicted by courts for suspects. Traditional methods, which utilize professional legal-linguistic style (PLLS) data, typically benefit only legal experts because the domain discrepancy between PLLS and non-PLLS degrades the models’ performance on non-PLLS texts. However, unprofessional users also show an increasing demand on such a prediction service. The scarcity of non-PLLS data poses significant challenges in developing effective models that benefit ordinary users (Zhao et al., 2023).

We address this problem by training SC2 to transfer the PLLS fact descriptions to non-PLLS texts, and then developing a charge prediction model based on the hybrid data (we mix PLLS and non-PLLS texts). The utilized datasets are CAIL (Xiao et al., 2018) and NCCP (Zhao et al., 2023) for PLLS and non-PLLS respectively. We employ three charge prediction methods (FewShot (Hu et al., 2018), CECP (Zhao et al., 2022) and DLCCP (Zhao et al., 2023)) and three domain adaptation methods (FADA (Motiian et al., 2017), *d*-SNE (Xu et al., 2019) and DAGE (Morsing et al., 2021)) as baselines, following the work of (Zhao et al., 2023). We also employ the long TST method StoryTrans (Zhu et al., 2023) as a data aug-menter. The settings for training SC2 and StoryTrans are consistent with that of the Chinese sub-dataset in this paper. And the settings for training baselines and charge prediction model are consistent with (Zhao et al., 2023). We employ Accuracy (Acc.), Macro-

Metrics	Acc.	MP	MR	F1
FewShot	0.5527	0.5883	0.5677	0.5056
CECP	0.5170	0.5854	0.5461	0.4962
FADA	0.6230	0.5584	0.6464	0.5482
<i>d</i> -SNE	0.6185	0.5778	0.6214	0.5528
DAGE	0.5959	<u>0.5863</u>	0.6145	0.5473
DLCCP	<u>0.6390</u>	0.5810	0.6596	0.5619
StoryTrans	0.6094	0.4840	0.5106	0.4665
SC2	0.6881	0.6417	0.6259	0.5860

Table 6: Extrinsic experiment results. Underlined values denote the optimal results of baselines.

Precision (MP), Macro-Recall (MR), and Macro-F1 (F1) as the evaluation metrics. The results are shown in Table 6. It is evident that the proposed SC2 effectively improves the performance of non-PLLS charge prediction compared to all baselines.

5 Conclusion

In this paper, we propose a novel method SC2 for long TST. We design multiple JSCW layers to progressively refine content representations of tokens, thus enhancing content preservation. And we design a style consistency loss to ensure the generated sentences share similarities in style polarity. Experimental results with automatic and human evaluations confirm the effectiveness of SC2 compared to competitive baselines.

Limitations

We have not yet evaluated the performance of SC2 with larger pre-trained language models since they are not mainstream methods for TST and they exceed the capacity of our available GPU resources. Another limitation is that, for both Chinese and English, the data used in our experiments only involves a single type of long TST, specifically imitating the writing style of a particular author. We also find the styles to be quite distinct because of the high accuracy scores obtained by style classifiers. This limitation arises from the fact that there is only one available dataset for the long TST in the existing research. We leave the construction of diverse data (e.g., formality or humor transfer, and

texts with closer styles) and the comprehensive validation of SC2’s generalizability and adaptability to future work.

Ethics Statement

Long TST involves generating target text based on both the original text and target style signals. The nature of text generation and the diversified stylistic corpus raise potential ethical concerns, including the risk of generating text that may be inappropriate, offensive, or biased. On the other hand, the model carries the risk of being maliciously exploited, such as generating fake news and fabricating political statements. Future research on how to mitigate these issues is in crucial need.

We hired three Chinese native speakers as annotators to manually evaluate the performance of the proposed method and baselines. Considering the wage standards of China, annotators will get 2.0 yuan (RMB) for each sample.

Acknowledgements

This research was supported by the National Natural Science Foundation of China (Grant Nos. 62133012, 61936006, 62103314, 62073255, 62303366), the Key Research and Development Program of Shaanxi (Program No. 2020ZDLGY04-07), Innovation Capability Support Program of Shaanxi (Program No. 2021TD-05) and Natural Science Basic Research Program of Shaanxi under Grant No.2023-JC-QN-0648.

References

- Yu Cheng, Zhe Gan, Yizhe Zhang, Oussama Elachqar, Dianqi Li, and Jingjing Liu. 2020. [Contextual text style transfer](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2915–2924, Online. Association for Computational Linguistics.
- Ning Dai, Jianze Liang, Xipeng Qiu, and Xuanjing Huang. 2019. [Style transformer: Unpaired text style transfer without disentangled latent representation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5997–6007, Florence, Italy. Association for Computational Linguistics.
- Xiang Gao, Yizhe Zhang, Sungjin Lee, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2019. [Structuring latent spaces for stylized response generation](#). 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 1814–1823, Hong Kong, China. Association for Computational Linguistics.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2017. [Non-autoregressive neural machine translation](#). *arXiv preprint arXiv:1711.02281*.
- Jian Guan, Zhuoer Feng, Yamei Chen, Ruilin He, Xiaoxi Mao, Changjie Fan, and Minlie Huang. 2022. [LOT: A story-centric benchmark for evaluating Chinese long text understanding and generation](#). *Transactions of the Association for Computational Linguistics*, 10:434–451.
- Jingxuan Han, Quan Wang, Licheng Zhang, Weidong Chen, Yan Song, and Zhendong Mao. 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)*, pages 7914–7927, Toronto, Canada. Association for Computational Linguistics.
- Qinghong Han, Yuxian Meng, Fei Wu, and Jiwei Li. 2020. [Non-autoregressive neural dialogue generation](#). *arXiv preprint arXiv:2002.04250*.
- Zikun Hu, Xiang Li, Cunchao Tu, Zhiyuan Liu, and Maosong Sun. 2018. [Few-shot charge prediction with discriminative legal attributes](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 487–498.
- Fei Huang, Zikai Chen, Chen Henry Wu, Qihan Guo, Xiaoyan Zhu, and Minlie Huang. 2021. [NAST: A non-autoregressive generator with word alignment for unsupervised text style transfer](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1577–1590, Online. Association for Computational Linguistics.
- Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2022. [Deep learning for text style transfer: A survey](#). *Computational Linguistics*, 48(1):155–205.
- Zhijing Jin, Di Jin, Jonas Mueller, Nicholas Matthews, and Enrico Santus. 2019. [IMaT: Unsupervised text attribute transfer via iterative matching and translation](#). 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 3097–3109, Hong Kong, China. Association for Computational Linguistics.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. [Disentangled representation learning for non-parallel text style transfer](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 424–434, Florence, Italy. Association for Computational Linguistics.

- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Dongkyu Lee, Zhiliang Tian, Lanqing Xue, and Nevin L. Zhang. 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)*, pages 93–102, Online. Association for Computational Linguistics.
- Haejun Lee, Drew A. Hudson, Kangwook Lee, and Christopher D. Manning. 2020. [SLM: Learning a discourse language representation with sentence unshuffling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1551–1562, Online. Association for Computational Linguistics.
- 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, New Orleans, Louisiana. Association for Computational Linguistics.
- Lukas Hedegaard Morsing, Omar Ali Sheikh-Omar, and Alexandros Iosifidis. 2021. Supervised domain adaptation using graph embedding. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7841–7847. IEEE.
- Saeid Motiian, Quinn Jones, Seyed Iranmanesh, and Gianfranco Doretto. 2017. Few-shot adversarial domain adaptation. *Advances in neural information processing systems*, 30.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Kainan Peng, Wei Ping, Zhao Song, and Kexin Zhao. 2020. Non-autoregressive neural text-to-speech. In *International conference on machine learning*, pages 7586–7598. PMLR.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Abhinav Ramesh Kashyap, Devamanyu Hazarika, Min-Yen Kan, Roger Zimmermann, and Soujanya Poria. 2022. [So different yet so alike! constrained unsupervised text style transfer](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 416–431, Dublin, Ireland. 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, Dublin, Ireland. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ["why should i trust you?": Explaining the predictions of any classifier](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Parker Riley, Noah Constant, Mandy Guo, Girish Kumar, David Uthus, and Zarana Parekh. 2021. [TextSETTR: Few-shot text style extraction and tunable targeted restyling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3786–3800, Online. Association for Computational Linguistics.
- 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, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Bakhtiyar Syed, Gaurav Verma, Balaji Vasan Srinivasan, Anandhavelu Natarajan, and Vasudeva Varma. 2020. Adapting language models for non-parallel author-stylized rewriting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9008–9015.
- Ewoenam Kwaku Tokpo and Toon Calders. 2022. [Text style transfer for bias mitigation using masked language modeling](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*, pages 163–171, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

- Yunli Wang, Yu Wu, Lili Mou, Zhoujun Li, and Wenhan Chao. 2019. [Harnessing pre-trained neural networks with rules for formality 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 3573–3578, Hong Kong, China. Association for Computational Linguistics.
- Xing Wu, Tao Zhang, Liangjun Zang, Jizhong Han, and Songlin Hu. 2019. [Mask and infill: Applying masked language model for sentiment transfer](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5271–5277. International Joint Conferences on Artificial Intelligence Organization.
- Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Yansong Feng, Xianpei Han, Zhen Hu, Heng Wang, and Jianfeng Xu. 2018. CAIL2018: A large-scale legal dataset for judgment prediction.
- Minzhang Xu, Min Peng, and Fang Liu. 2023. Text style transfer between classical and modern chinese through prompt-based reinforcement learning. *World Wide Web*, 26(2):733–750.
- Xiang Xu, Xiong Zhou, Ragav Venkatesan, Gurumurthy Swaminathan, and Orchid Majumder. 2019. d-sne: Domain adaptation using stochastic neighborhood embedding. In *CVPR*, pages 2497–2506.
- Xiaoyuan Yi, Zhenghao Liu, Wenhao Li, and Maosong Sun. 2020. [Text style transfer via learning style instance supported latent space](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3801–3807. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Jie Zhao, Ziyu Guan, Cai Xu, Wei Zhao, and Enze Chen. 2022. Charge prediction by constitutive elements matching of crimes. In *IJCAI*, pages 4517–4523. IJCAI Organization.
- Jie Zhao, Ziyu Guan, Wei Zhao, Yue Jiang, and Xiaofei He. 2023. Few-shot domain adaptation for charge prediction on unprofessional descriptions. *arXiv preprint arXiv:2309.17313*.
- Chulun Zhou, Liangyu Chen, Jiachen Liu, Xinyan Xiao, Jinsong Su, Sheng Guo, and Hua Wu. 2020. [Exploring contextual word-level style relevance for unsupervised style transfer](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7135–7144, Online. Association for Computational Linguistics.
- Xuekai Zhu, Jian Guan, Minlie Huang, and Juan Liu. 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)*, pages 14803–14819, Toronto, Canada. Association for Computational Linguistics.