Positive Text Reframing under Multi-strategy Optimization

Anonymous ACL submission

Abstract

Differing from sentiment transfer, positive reframing seeks to substitute negative perspectives with positive expressions while preserving the original meaning. With the emergence of pre-trained language models (PLMs), it is possible to achieve acceptable results by fine-tuning PLMs. Nevertheless, generating fluent, diverse and task-constrained reframing text remains a significant challenge. To tackle this issue, a **m**ulti-**s**trategy **o**ptimization **f**ramework (MSOF) is proposed in this paper. MSOF starts with the objective of positive reframing, introducing positive sentiment reward and content preservation reward to encourage the model to transform the negative expressions of the original text while ensuring the integrity and consistency of the semantics. Then, different decoding optimization approaches are introduced to improve the quality of text generation. Finally, based on the modeling formula of positive reframing, the candidate sentences are further selected from three dimensions: strategy consistency, text similarity and fluency. Extensive experiments on two Seq2Seq PLMs, BART and T5, demonstrate our framework achieves significant improvements on unconstrained and controlled positive reframing tasks.

1 Introduction

The concept of style transfer initially emerges within the domain of computer vision (CV) with the objective of accomplishing image style transfer (Gatys et al., 2016). Inspired by this, Hu et al. (2017) proposed text style transfer (TST), whose main purpose is to automatically control the text style and preserve the style-independent content. There also have been some related research before this, such as paraphrase (Xu et al., 2012). In recent years, there has been an increasing focus on TST, which has gradually evolved into a significant subfield within the domain of natural language generation. Many corresponding task variants also have been proposed, such as text form transfer (Briakou et al., 2021), topic transfer (Huang et al., 2020), text simplification (Cao et al., 2020), and sentiment transfer (Mueller et al., 2017), etc.

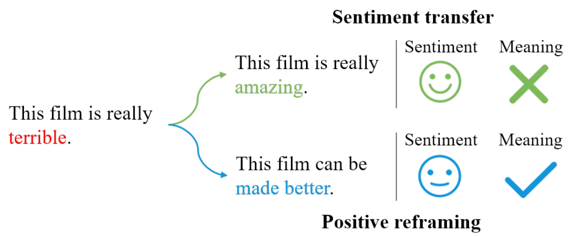

Figure 1: The difference between sentiment transfer and positive reframing.

Among them, sentiment transfer primarily focuses on reversing the sentiment polarity of the original text. However, it relies on the straightforward replacement of opinion words, such as substituting negative opinion words with their positive counterparts of the opposite meaning. On the one hand, it retains the content irrelevant to style to some extent, such as the invariance of described object entities. On the other hand, it also inherently alters the meaning of the original text (Liao et al., 2018; Li et al., 2018). To this end, Ziems et al. (2022) proposed positive reframing. In contrast to sentiment transfer, positive reframing adopts principles from psychology to reframe negative text by introducing a complementary positive viewpoint while simultaneously maintaining the underlying meaning conveyed in the original text. A toy example of their difference can be seen in Figure 1.

More specifically, positive reframing encompasses various tasks, including unconstrained positive reframing, controlled positive reframing, and derivative tasks such as reframe strategy classification. The unconstrained positive reframing task focuses on generating reframed text without explicit guidance of the corresponding reframe strategy. In contrast, the controlled positive reframing task involves reframing text based on the given strategy. And the reframe strategy classification task entails determining the specific strategy employed in reframing text. Ziems et al. (2022) gives six positive reframing strategies, namely growth mindset, impermanence, neutralization, optimism, self-affirmation and thankfulness.

However, most of the existing methods only fine-tune PLMs on the corresponding dataset, ignoring the consistency requirement between the model training objective and the target of positive reframing, and also failing to fully utilize the known condition of the reframing strategy under the controlled setting, making it difficult to ensure that the generated text meets the task requirements. Therefore, this paper proposes a multi-strategy optimization framework (MSOF) for positive reframing and our contributions are as follows:

- Firstly, from the target of positive reframing, we design and implement the positive sentiment reward and content preservation reward to optimize the sequence-level training objective, and then apply various decoding improvement approaches to alleviate text degeneration and elevate the quality and diversity of the generated text.
- Secondly, we propose a multi-dimensional re-ranking approach based on the modeling formula of positive reframing, which comprehensively evaluates the quality of the candidate text based on strategy consistency, text similarity and fluency.
- Extensive experimental results demonstrate that our proposed multi-strategy optimization framework achieves significant improvement on both unconstrained and controlled positive reframing task. And we would release our code to encourage future research[1].

2 Related Work

Early research on **text style transfer** mostly relied on artificial design features such as syntax (Zhu et al., 2010) and phrase (Xu et al., 2012) modeling, etc. Similar to other tasks in NLP, the advent of deep learning has resulted in the growing application of neural network models to TST. For example, Jhamtani et al. (2017) investigated the utilization of the Seq2Seq model for transforming modern English into Shakespearean-style English. Wang et al. (2019) applied GPT-2 to accomplish the formal-informal transfer. Sancheti et al. (2020) extended the work of Jhamtani et al. (2017) by incorporating a reinforcement learning framework. Lai et al. (2021) further applied this framework to PLMs. Above studies are mainly based on parallel corpora. Although satisfactory results can be achieved, the cost of constructing parallel corpora is expensive. Therefore, semi-supervised learning and unsupervised learning are widely used in TST. The main methods include data augmentation or text retrieval (Zhang et al., 2020; Jin et al., 2019), adversarial learning (Hu et al., 2017; Fu et al., 2018), back-translation (Prabhumoye et al., 2018; Wei et al., 2023), and reinforcement learning (Luo et al., 2019; Gong et al., 2019).

Specific to **sentiment transfer**, the early goal is to extract sentiment words that describe the corresponding entities, and then replace them with expressions of the opposite sentiment attribute. The representative one is the "Delete, Retrieve, Generate" strategy (Li et al., 2018). Furthermore, Sudhakar et al. (2019) applied the transformer architecture to the above strategy. To better distinguish content and style, Kim and Sohn (2020) divided the model into sentence reconstruction module and style module to complete their respective task. Han et al. (2023) introduced the adaptive clustering and contrastive learning modules to better explore sentence transmission patterns to main and utilize the latent transfer patterns.

Although sentiment transfer preserves attribute-independent content, the intrinsic meaning of the original text expression is also changed. To this end, Ziems et al. (2022) introduced **positive reframing**, aiming to preserve the original meaning by substituting negative viewpoints with complementary positive expressions, and constructed the corresponding parallel dataset. For unconstrained positive reframing, Xu et al. (2023) decoupled the sentiment and style of the text to complete the positive reframing. Then, Sheng et al. (2023) further decomposed positive reframing into paraphrase generation and sentiment transfer and constructed corresponding pseudo datasets to fuse generation capabilities through multi-task learning, but also led to the inability to apply their method under the controlled setting.

3 Methodology

3.1 Problem Definition

Let (x, y, ψ_x) be a triple in the positive reframing task, where $x = \{x_1, x_2, \ldots, x_n\}$ is the original text with negative sentiment, and $y = \{y_1, y_2, \ldots, y_m\}$

[1] https://anonymous.4open.science/r/code-for-paper-B875/

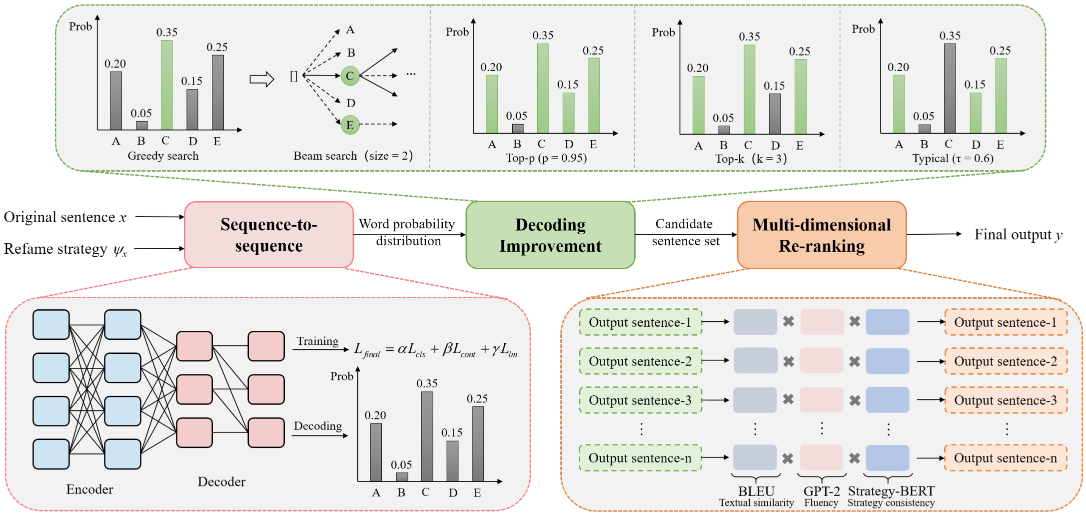

Figure 2: The overall architecture of MSOF. We respectively use BART and T5 as the basic model for positive reframing. The positive sentiment reward and content preservation reward are applied to optimize the model training process. Then, we adopt various decoding improvement approaches (e.g. beam search, random sampling) during the decoding stage to improve the quality of text generation. Finally, multi-dimensional re-ranking is used to comprehensively evaluate candidate sentences and select the candidate with the highest score as the final output.

is the target sentence with complementary positive expressions corresponding to x, m and n represent the sentence length. $\psi_x \subseteq$ {Growth Mindset, Impermanence, Neutralizing, Optimism, Self-affirmation, Thankfulness} is the positive reframing strategy used to reframe the negative text x, which can use multiple strategies simultaneously. This paper researches the following three tasks and ultimately focuses on controlled positive reframing task.

The target of unconstrained positive reframing is to generate the target sentence y from the original text x **without** any reframe strategy guidance. This task can be modeled as follows:

$$p(y|x) = \prod_{t=1}^{m} p(y_t|x, y_{<t}) \quad (1)$$

where $y_{<t}$ represents what has been generated before time t.

Regarding reframe strategy classification, its requirement is to predict the positive reframing strategy ψ_x used to reframe the original sentence x.

For controlled positive reframing, the primary objective is to generate the target sentence y from the original text x **under** given strategy ψ_x, This problem can be modeled as the following formula.

$$p(y|x, \psi_x) = \prod_{t=1}^{m} p(y_t|x, \psi_x, y_{<t}) \quad (2)$$

3.2 Framework

As shown in Figure 2, our proposed framework mainly consists of four modules, namely sequence-to-sequence, reinforcement training, decoding improvement and multi-dimensional re-ranking.

3.2.1 Sequence-to-sequence

Consistent with Ziems et al. (2022), we also use T5 (Raffel et al., 2019) and BART (Lewis et al., 2020) as the basic text generation model, which are both mainly composed of two components, namely encoder and decoder.

Encoder This part is to encode original sentence x and reframe strategy ψ_x into hidden vector H. We use T5 and BART as the basic generation model, and the encoder part is as follows:

$$H = \text{Encoder}([x_1, x_2, \ldots, x_n], \psi_x) \quad (3)$$

where $H \in \mathbb{R}^{l \times d}$, l is the length of sequence, and d is the hidden dimension.

Decoder The output y_t of the decoder part takes the hidden vector output of the encoder and the output $y_{<t}$ of the decoder before time t as input, the equation is as follows.

$$y_t = \text{Decoder}(H; y_{<t}) \quad (4)$$

3.2.2 Reinforcement Training

As shown in Figure 3, based on the objective of positive reframing, the generated text should transform the negative sentiment of the original text and keep the semantics unchanged. Therefore, we design and implement positive sentiment reward and content preservation reward to optimize the overall training process.

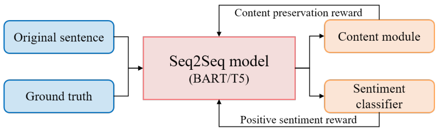

Figure 3: The reinforcement training procedure of the Seq2Seq-based model.

Positive sentiment reward We first design the positive sentiment reward loss based on binary cross entropy (BCE). Specifically, we fine-tune the binary sentiment classifier RoBERTa (Liu et al., 2019) and utilize it to determine the sentiment change degree of the generated sentence relative to the original text. The positive sentiment reward loss function is formulated as follows:

$$p(s_t|y', x) = \mathrm{Sigmoid}(\mathrm{RoBERTa}(y', x)) \quad (5)$$

$$L_{cls} = -\log(p(s_t|y', x)) \quad (6)$$

where s_t represents the target style, and y' is the generated sentence.

Content preservation reward Inspired by Lai et al. (2021), we use BLEU score as the reward for content preservation and leverage SCST (Self-Critic Sequence Training) approach (Rennie et al., 2017) as the optimization method. The corresponding loss function is as follows:

$$L_{cont} = \sum_i log(p(y_i^s|y_{1:i-1}^s, x))(bleu(y', y) - bleu(y^s, y)) \quad (7)$$

where y^s is sampled from the distribution of model outputs at each time step, and y' is the greedy generation from the model.

The overall loss is a weighted sum of the positive sentiment reward loss L_{cls}, content preservation reward loss L_{cont}, and language modeling loss L_{lm}.

$$L_{lm} = \sum_i log(p(y_i|y_{1:i-1}, x)) \quad (8)$$

$$L_{final} = \alpha L_{cls} + \beta L_{cont} + \gamma L_{lm} \quad (9)$$

3.2.3 Decoding Improvement

Although T5 and BART have demonstrated their superiority in the field of NLG, the sentences generated by default greedy search often result in text degeneration (i.e., empty or repeated sequences) during the decoding stage (Fan et al., 2018; Holtzman et al., 2019). Therefore, in this paper, various decoding improvement ways such as Beam search (Wiseman and Rush, 2016), Top-k sampling (Fan et al., 2018), Top-p sampling (Holtzman et al., 2019) and Typical sampling (Meister et al., 2023) are applied to the decoding stage of the Seq2Seq model to improve the quality of text generation. And Eq. 4 is changed as follows.

$$y_t = \text{Post-Processing}(\text{Decoder}(H; y_{<t})) \quad (10)$$

3.2.4 Multi-dimensional Re-ranking

According to Bayes Rule, we can decompose Eq. 2 into the product of three probabilities:

$$p(y|x, \psi_x) = p(\psi_x|y, x) \times p(x|y) \times p(y) \quad (11)$$

The first term $p(\psi_x|y, x)$ can be seen as the consistency of original-to-generative sentence transformation with given reframe strategy[2]. The second term $p(x|y)$ represents the textual similarity. And the last term $p(y)$ can be regarded as the overall fluency of the output.

Strategy consistency For this term, we propose Strategy-BERT to evaluate the consistency between text reframing and the given strategy, which draws on the idea of "breaking the whole into pieces" and prompt learning to transform the multi-label problem into multiple binary classification tasks, i.e. training the corresponding model for each reframing strategy. For one thing, this approach enables each model to concentrate on its specific aspect and thus not affect each other. For another thing, it facilitates context semantic enhancement by constructing an auxiliary sentence that incorporates supplementary task prompt to effectively mine the implicit task-specific knowledge contained in PLMs and alleviate the task awareness challenge.

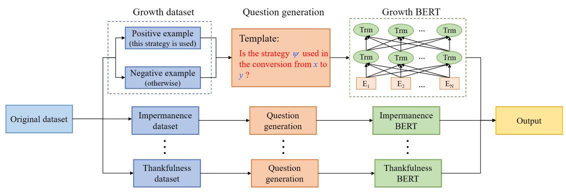

Figure 4: The overall procedure of reframe strategy classification.

As shown in Figure 4, the original dataset is firstly divided according to the different strategies used in reframing, that is, if the strategy ψ is used

[2]For unconstrained setting, there is no this term.

in the original-reframed text transfer, this sentence pair will be regarded as a positive sample of corresponding strategy dataset, otherwise, it will be a negative sample. The dataset division results are shown in Table 2.

For different reframe strategies, this paper uses the following way to construct auxiliary question:

"Is the strategy + strategy type + used in the conversion from + original + to + reframe + ?" where the artificially added tokens are marked in red, and the reframe strategy, original sentence and reframed sentence are marked in blue. In this way, context semantic enhancement can be achieved by constructing auxiliary question.

Then, we fine-tune BERT on above dataset and propose Strategy-BERT specific to each reframe strategy, which is used to evaluate the strategy consistency score of candidate sentences. We also implement the above thought on RoBERTa, and propose Strategy-RoBERTa.

Textual similarity Regarding this item, we still use BLEU to calculate this term because it measures the overlap between the ground truth and the generated text (Sancheti et al., 2020).

Fluency Recent works suggest that the probability of output generated from PLM is an appropriate automatic and referenceless measure of fluency (Suzgun et al., 2022; Ramirez et al., 2023). Therefore, we use GPT-2_{large} (Radford et al., 2019) to calculate the overall fluency of each candidate.

4 Experiment

4.1 Dataset

Positive reframing For unconstrained positive reframing and controlled positive reframing, we adopt the dataset provided by Ziems et al. (2022) . and the specific statistics are given in Table 1.

Label	Train	Dev	Test
Growth	1683	216	221
Impermanence	1296	172	157
Neutralizing	2410	303	302
Optimism	3295	373	400
Self-affirmation	673	92	76
Thankfulness	882	94	109

Table 1: The statistics of the positive reframing dataset (unconstrained & controlled).

Reframe strategy classification To verify the effectiveness of Strategy-BERT, we conduct experiments on reframe strategy classification task. Since this paper converts the multi-label classification problem into multiple binary classification tasks, the dataset is also divided accordingly, and the division results are presented in Table 2.

Label	Train		Dev		Test	
	POS	NEG	POS	NEG	POS	NEG
Growth	1683	4996	216	619	221	614
Impermenence	1296	5383	172	663	157	678
Neutralizing	2410	4269	303	532	302	533
Optimism	3295	3383	373	462	400	435
Self-affirmation	673	6006	92	743	76	759
Thankfulness	882	5797	94	741	109	726

Table 2: The statistics of the reframe strategy classification dataset.

4.2 Evaluating Metrics

Regarding classification task, following Ziems et al. (2022), we use F1 score as the evaluation metric.

For generation task, the following automatic metrics are used: (1) BLEU (Papineni et al., 2002) is used to evaluate the overall quality of the generated sentences, specifically, this paper adopts the implementation of Post (2018). (2) ROUGE (Lin, 2004) is used for evaluating the degree of overlap between generated and reference text, specifically, this paper uses ROUGE-1 (R-1), ROUGE-2 (R-2) and ROUGE-L (R-L). (3) BERTScore (BScore) (Zhang et al., 2019) utilizes the contextual embedding representation of text to measure the similarity between generated and reference text. (4) ΔTextBlob (ΔTB) (Loria, 2018) is used to report the average change in sentiment. (5) **R**eframing **T**ext **Q**uality **E**valuation (RTQE) is used for evaluating the degree of positive text reframing (i.e. style strength), inspired by Lai et al. (2021), we fine-tune RoBERTa$_{\text{large}}$ (Liu et al., 2019) to evaluate reframing degree and we regard the probability from the model prediction as the degree of positive reframing between the original and generated sentence; on the human reference it has the F1 score of 95.98% and accuracy of 97.41%, (6) Perplexity (PPL) is an indicator of text fluency, and we use GPT-2_{large} as the evaluation model.

Finally, following Ziems et al. (2022), We randomly selected 50 samples from each generated file and assigned them to 3 well-educated raters with relevant professional backgrounds to score Meaning Preservation (Meaning), Positivity and Fluency of reframed sentences on a scale of 1 to 5. Since the main research of this paper falls on controlled positive reframing task, we only conducted human evaluation on this task.

4.3 Implementation Details

Reframe strategy classification $\text{BERT}_{\text{base}}$ (Devlin et al., 2019) and $\text{RoBERTa}_{\text{base}}$ (Liu et al., 2019) are used as the backbone model in this task respectively. The maximum text embedding length is set to 110. AdamW is used as the optimizer, and the batch size is 16. In addition, all models in this paper are implemented through HuggingFace (Wolf et al., 2020) and PyTorch (Paszke et al., 2019) on TITAN Xp GPU.

Positive reframing Following Ziems et al. (2022), we use T5 (Raffel et al., 2019) and BART (Lewis et al., 2020) with 6 layers in each of the encoder and decoder, and the hidden size of 768. The value of the learning rate is from 3e-5 to 3e-4, the batch size processed by each device is 6, and the text maximum input length is 80. α, β, γ are respectively set to 1, 0.2, 1.

4.4 Main Results

4.4.1 Reframe Strategy Classification

For this task, this paper selects the Multi-label-BERT and Multi-label-RoBERTa proposed by Ziems et al. (2022) as baselines to compare with the Strategy-BERT and Strategy-RoBERTa proposed in this paper. For fairness, we directly adopt the results reported by Ziems et al. (2022). Since they only report F1 score of their models, we only use it as the evaluation metric in this task. The detailed performance of our proposed models on other metrics can be found in Table 12 in Appendix D.1.

Label	Multi-label-BERT	Multi-label-RoBERTa	Strategy-BERT	Strategy-RoBERTa
Thankfulness	0.71	0.69	**0.73**	0.72
Neutralizing	0.59	0.61	**0.61**	0.61
Optimism	0.71	0.71	0.71	**0.73**
Impermanence	0.55	0.55	**0.57**	0.57
Growth	0.63	0.63	0.67	**0.69**
Self-affirmation	0.43	0.44	**0.48**	0.46

Table 3: The experimental results of reframe strategy classification on F1 score. And the best results in each label are in **bold**.

It can be seen from Table 3 that our models are able to outperform baselines on all labels, significantly on the Growth (Growth Mindset) label, the two models proposed in this paper have increased by 4 points and 6 points respectively. Furthermore, in terms of the Self-affirmation label, Strategy-BERT demonstrates a noteworthy improvement of 5 points compared to the corresponding baseline. Additionally, our method consistently achieves approximately 1 point of improvement on other labels, further affirming the effectiveness and superiority of our approach. Since the performance of Strategy-BERT and Strategy-RoBERTa are similar, we only use Strategy-BERT as the evaluation model to measure the strategy consistency of each candidate.

Label	Strategy-BERT w/o auxiliary	Strategy-BERT	Strategy-RoBERTa w/o auxiliary	Strategy-RoBERTa
Thankfulness	0.71	**0.73**	0.69	0.72
Neutralizing	0.59	**0.61**	0.60	0.61
Optimism	0.71	0.71	0.71	**0.73**
Impermanence	0.55	**0.57**	0.55	0.57
Growth	0.61	0.67	0.65	**0.69**
Self-affirmation	0.44	**0.48**	0.44	0.46

Table 4: The experimental results of different input ways on F1 score. The best results in each label are in **bold** and w/o auxiliary means without using auxiliary sentence.

In addition, the performance of the input approach of directly connecting the original and generated sentence is also tested to demonstrate the effectiveness of the contextual semantic enhancement strategy (i.e., the construction of auxiliary question) used in this paper. And the experimental results are given in Table 4. As can be seen, the F1 score on each label is greatly reduced without context enhancement strategy, but our models still achieve comparable performance with the multi-label classification models which once again proves the effectiveness of our method.

4.4.2 Unconstrained Positive Reframing

As shown in Table 5, our proposed framework MSOF achieves significant improvements compared to the baselines. When combining positive sentiment reward and content preservation reward only during the training process, i.e. $\text{MSOF}_{\text{Greedy}}$, already outperforms the baselines on almost all metrics, especially ROUGE, BScore, RTQE, and PPL. When incorporating decoding optimization and multi-dimensional re-ranking, the performance of the model will be further improved. From the perspective of the model, the T5-based models achieve the best results on metrics such as ΔTB, RTQE and PPL, while the BART-based models reach SOTA on content preservation-related metrics such as ROUGE, BLEU, and BScore. This may be because BART prioritizes semantic preservation rather than sentiment change when reframing the negative text. Among different decoding methods, both beam search and random sampling-based

Model	R-1	R-2	R-L	BLEU	BScore	ΔTB	RTQE	PPL
T5 (Ziems et al., 2022)	27.4	9.8	23.8	8.7	88.7	0.38	84.8	42.7
FDSC (Xu et al., 2023)	30.4	10.9	25.2	8.1	88.8	0.39	93.1	30.0
PG2ST (Sheng et al., 2023)	31.1	11.2	25.5	8.9	88.7	0.35	85.4	41.0
ST2PG (Sheng et al., 2023)	30.8	11.3	25.5	8.8	88.7	0.33	84.6	43.2
MSOF$_{\text{Greedy}}$	32.9	13.0	26.0	8.8	89.1	0.37	86.2	36.8
MSOF$_{\text{Beam}}$	34.1	14.0	27.1	9.7	89.2	0.37	89.0	35.4
MSOF$_{\text{Top-k}}$	**34.8**	**14.7**	**27.7**	**10.1**	**89.5**	**0.44**	93.5	22.3
MSOF$_{\text{Top-p}}$	34.4	14.6	27.6	10.1	89.4	0.43	93.5	**22.2**
MSOF$_{\text{Typical}}$	32.9	13.5	26.2	9.1	89.3	0.39	**94.5**	22.6
BART (Ziems et al., 2022)	27.7	10.8	24.3	10.3	89.3	0.23	63.8	86.0
FDSC (Xu et al., 2023)	32.7	13.4	27.0	10.4	88.5	0.21	60.1	77.5
PG2ST (Sheng et al., 2023)	32.6	13.5	26.9	10.3	88.4	0.19	60.9	86.2
ST2PG (Sheng et al., 2023)	32.9	13.6	27.1	10.9	88.4	0.20	61.5	78.9
MSOF$_{\text{Greedy}}$	32.3	13.2	26.9	10.4	89.4	0.24	80.1	47.0
MSOF$_{\text{Beam}}$	34.2	14.2	28.1	10.9	89.5	0.24	87.3	33.6
MSOF$_{\text{Top-k}}$	**34.8**	**14.9**	**29.3**	**12.0**	**89.9**	**0.31**	87.3	**25.8**
MSOF$_{\text{Top-p}}$	34.8	14.9	29.2	12.0	89.8	0.30	87.2	27.3
MSOF$_{\text{Typical}}$	32.5	12.8	26.9	10.4	89.5	0.30	**88.5**	29.6

Table 5: The experimental results of **unconstrained positive reframing**. The best in-category performance is **bolded** and the best overall performance is **highlighted**. And expect for PPL, all other metrics are better when they are higher.

Model	R-1	R-2	R-L	BLEU	BScore	ΔTB	RTQE	PPL
T5 (Ziems et al., 2022)	27.7	10.0	23.9	8.8	88.8	0.36	86.2	62.1
MSOF$_{\text{Greedy}}$	33.6	13.6	26.7	8.8	89.2	0.37	94.6	34.6
MSOF$_{\text{Beam}}$	34.6	14.4	27.5	9.5	89.3	0.36	96.2	34.5
MSOF$_{\text{Top-k}}$	**34.8**	**15.0**	**28.0**	**9.9**	**89.5**	**0.43**	**97.7**	23.1
MSOF$_{\text{Top-p}}$	34.1	14.2	27.6	9.3	89.5	0.42	96.6	**23.0**
MSOF$_{\text{Typical}}$	33.2	13.4	26.5	8.6	89.3	0.42	97.0	23.8
BART (Ziems et al., 2022)	28.8	10.9	25.1	10.1	89.6	0.27	69.5	89.1
MSOF$_{\text{Greedy}}$	33.0	13.3	27.2	10.0	89.6	0.31	89.1	44.4
MSOF$_{\text{Beam}}$	34.6	14.2	28.2	10.5	89.7	0.34	**94.8**	31.8
MSOF$_{\text{Top-k}}$	**34.8**	**14.7**	**29.0**	**11.4**	**90.1**	**0.36**	94.0	**29.4**
MSOF$_{\text{Top-p}}$	34.6	14.4	28.8	11.3	90.0	0.36	94.0	30.8
MSOF$_{\text{Typical}}$	33.2	13.2	27.5	10.1	89.8	0.36	94.0	29.8

Table 6: The experimental results of **controlled positive reframing**.

methods are superior to greedy search. Specifically, Top-k sampling has the best overall performance, achieving the best or sub-optimal results on almost all metrics. Top-p sampling performs slightly lower than Top-k sampling. Compared to the above two decoding methods, beam search and Typical sampling are not satisfactory but still superior to the baseline method. Ultimately, regardless of whether T5 or BART is used as the basic generation model, MSOF$_{\text{Top-k}}$ achieves the best results among all variant models, basically achieving at least 7% improvement on each metric compared to baselines, which strongly proves the effectiveness of our proposed framework.

4.4.3 Controlled Positive Reframing

Since only Ziems et al. (2022) have studied controlled positive reframing, we use T5 and BART (Ziems et al., 2022) that are fine-tuned on the corresponding dataset as baselines for comparison. The primary experimental results are given in Table 6. It can be concluded that the performance of models under constraints is generally better than unconstrained, which proves that the reframe strategy plays a role in assisting model inference to a certain extent. Consistent with the experimental results under the unconstrained setting, MSOF$_{\text{Top-k}}$ still achieves the best results among all variant models. Compared with the baselines, MSOF$_{\text{Top-k}}$ achieves an average improvement of 5 points on ROUGE, 1 point in BLEU, more than 10 points on both RTQE and PPL, and an improvement of about 20% on ΔTB. Moreover, it can be found that although Typical sampling does not perform as well as other decoding approaches on content preservation-related metrics such as ROUGE,

	Model	R-1	R-2	R-L	BLEU	BScore	ΔTB	RTQE	PPL
T5	$\text{MSOF}_{\text{Top}-k}$	34.8	**15.0**	**28.0**	**9.9**	89.5	**0.43**	**97.7**	**23.1**
	w.o Cls	34.5	14.5	27.5	9.4	89.4	0.41	96.7	25.3
	w.o Cont	**35.0**	14.8	27.7	9.6	**89.6**	0.37	95.7	24.2
	w.o Re-ranking	32.1	12.0	25.2	7.6	89.1	0.43	96.1	28.3
BART	$\text{MSOF}_{\text{Top}-k}$	**34.8**	**14.7**	**29.0**	**11.4**	**90.1**	0.36	**94.0**	**29.4**
	w.o Cls	33.6	13.7	28.2	10.8	90.0	0.35	86.9	31.3
	w.o Cont	33.1	13.7	27.5	10.9	89.7	**0.38**	86.2	34.6
	w.o Re-ranking	31.9	11.9	26.2	9.4	89.6	0.35	92.9	38.8

Table 7: The ablation experimental results of MSOF under controlled setting. w.o Cls means without positive sentiment reward, w.o Cont represents without content preservation reward, w.o Re-ranking represents not using multi-dimensional re-ranking.

BLEU, and BScore, it still achieves impressive results on ΔTB, RTQE and PPL, suggesting that its corresponding output is consistent with task requirements to some extend, even though there is less overlap with human reference.

4.4.4 Ablation Experiment

In addition, from the ablation experimental results shown in Table 7, we can conclude that only applying content preservation reward helps the model perform well on ROUGE, BLEU and BScore, but hinders the model from transferring text style. When using only positive sentiment reward, although the model performs well on ΔTB and RTQE, it is not satisfactory in terms of content preservation. However, when the two are combined, the model can achieve a better balance between sentiment change and content preservation, exhibiting a more comprehensive performance. Furthermore, it can be observed that the multi-dimensional re-ranking significantly improves the model's performance on multiple metrics. This demonstrates that it can effectively select the sentence from the candidate that better meets the requirements of positive reframing. Based on the above experimental results and analysis, the validity and rationality of each component of MSOF can be effectively proved. For more ablation experiments, please refer to Table 13 in Appendix D.2 and Table 14 in Appendix D.3.

4.4.5 Human Evaluation

Finally, we adopt human evaluation to manually judge the quality of the reframed text. As can be seen from Table 8, our method is more applicable to T5, but for BART, its performance on Positivity is not satisfactory, which can also be reflected by ΔTB and RTQE. Combining the relevant experimental results in Table 6, we speculate this is because the BART-based models prioritize content preservation over sentiment change. In general, consistent with the results and conclusion of automatic metrics, our method can effectively improve the model's performance, where the T5-based models perform better on Positivity and have a slightly higher score on Fluency, while BART-based models are better on Meaning.

Model	Meaning	Positivity	Fluency
T5 (Ziems et al., 2022)	4.13	3.89	4.07
$\text{MSOF}_{\text{Top}-k}$	**4.38**	**4.22**	**4.58**
BART (Ziems et al., 2022)	4.23	4.07	4.27
$\text{MSOF}_{\text{Top}-k}$	**4.42**	**4.10**	**4.54**

Table 8: The human evaluation results of controlled positive reframing.

5 Conclusion

We propose an original multi-strategy optimization framework (MSOF), which consists of reinforcement training, decoding improvement, and multi-dimensional re-ranking, to enhance the performance of PLMs on positive reframing. By conducting extensive experiments on T5-based and BART-based models separately, our framework achieves significant improvements over the baselines on various metrics. Future work includes further cleaning and expansion of the existing dataset to improve the quality and alleviate the imbalanced distribution of different reframe strategy labels, then exploring how the thought of controlled text generation can be applied to this task, followed by trying different approaches of context enhancement, and finally exploring how to apply large language models (LLMs) to positive reframing.

Limitations

Firstly, the multi-strategy optimization framework proposed in this paper introduces reinforced reward in the model training stage and the multi-dimensional re-ranking to select the candidate text generated by the model. Therefore, compared with the baselines, our proposed framework needs more memory space and time during training and prediction. Then, this paper finds that the dataset provided by Ziems et al. (2022) has certain noise and label imbalance issues that may hinder the training of the model and there are currently no corresponding datasets in other languages. Finally, we also suggest that if PLMs could be further trained in a rich psychological corpus, the performance would be improved more.

Ethics Statement

Similar to sentiment transfer, positive reframing has two sides, that is, our method can also be used to generate negative text and cause possible harmful effects on society. However, we still make our code public and hope others will be aware of the possible risks. We welcome any discussion and suggestions to minimize such risks.

References

Eleftheria Briakou, Di Lu, Ke Zhang, and Joel Tetreault. 2021. Olá, bonjour, salve! XFORMAL: A benchmark for multilingual formality style transfer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3199–3216, Online. Association for Computational Linguistics.

Yixin Cao, Ruihao Shui, Liangming Pan, Min-Yen Kan, Zhiyuan Liu, and Tat-Seng Chua. 2020. Expertise style transfer: A new task towards better communication between experts and laymen. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1061–1071, Online. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.

Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2018. Style transfer in text: Exploration and evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Hongyu Gong, Suma Bhat, Lingfei Wu, JinJun Xiong, and Wen-mei Hwu. 2019. Reinforcement learning based text style transfer without parallel training corpus. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3168–3180, Minneapolis, Minnesota. Association for Computational Linguistics.

Jingxuan Han, Quan Wang, 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.

Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *CoRR*, abs/1904.09751.

Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. Toward controlled generation of text. *CoRR*, abs/1703.00955.

Yufang Huang, Wentao Zhu, Deyi Xiong, Yiye Zhang, Changjian Hu, and Feiyu Xu. 2020. Cycle-consistent adversarial autoencoders for unsupervised text style transfer. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2213–2223, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Harsh Jhamtani, Varun Gangal, Eduard Hovy, and Eric Nyberg. 2017. Shakespearizing modern language using copy-enriched sequence to sequence models. In *Proceedings of the Workshop on Stylistic Variation*, pages 10–19, Copenhagen, Denmark. Association for Computational Linguistics.

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.

Heejin Kim and Kyung-Ah Sohn. 2020. How positive are you: Text style transfer using adaptive style embedding. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2115–2125, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Huiyuan Lai, Antonio Toral, and Malvina Nissim. 2021. Thank you BART! rewarding pre-trained models improves formality style transfer. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 484–494, Online. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 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.

Yi Liao, Lidong Bing, Piji Li, Shuming Shi, Wai Lam, and Tong Zhang. 2018. QuaSE: Sequence editing under quantifiable guidance. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3855–3864, Brussels, Belgium. Association for Computational Linguistics.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.

Steven Loria. 2018. textblob documentation. *Release 0.15, 2:269.*

Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Zhifang Sui, and Xu Sun. 2019. A dual reinforcement learning framework for unsupervised text style transfer. *CoRR*, abs/1905.10060.

Clara Meister, Tiago Pimentel, Gian Wiher, and Ryan Cotterell. 2023. Locally typical sampling. *Transactions of the Association for Computational Linguistics*, 11:102–121.

Jonas Mueller, David Gifford, and Tommi Jaakkola. 2017. Sequence to better sequence: continuous revision of combinatorial structures. In *International Conference on Machine Learning*, pages 2536–2544. PMLR.

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*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.

Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.

Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. Style transfer through back-translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 866–876, Melbourne, Australia. Association for Computational Linguistics.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *CoRR*, abs/1910.10683.

Angela Ramirez, Kartik Agarwal, Juraj Juraska, Utkarsh Garg, and Marilyn Walker. 2023. Controllable generation of dialogue acts for dialogue systems via few-shot response generation and ranking. In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 355–369, Prague, Czechia. Association for Computational Linguistics.

Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Abhilasha Sancheti, Kundan Krishna, Balaji Vasan Srinivasan, and Anandhavelu Natarajan. 2020. Reinforced rewards framework for text style transfer. In *Advances in Information Retrieval*, pages 545–560, Cham. Springer International Publishing.

Xu Sheng, Fumiyo Fukumoto, Jiyi Li, Go Kentaro, and Yoshimi Suzuki. 2023. Learning disentangled meaning and style representations for positive text reframing. In *Proceedings of the 16th International Natural Language Generation Conference*, pages 424–430, Prague, Czechia. Association for Computational Linguistics.

Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. "transforming" delete, retrieve, generate approach for controlled text style transfer. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3269–3279, Hong Kong, China. 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.

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.

Daimeng Wei, Zhanglin Wu, Hengchao Shang, Zongyao Li, Minghan Wang, Jiaxin Guo, Xiaoyu Chen, Zhengzhe Yu, and Hao Yang. 2023. Text style transfer back-translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7944–7959, Toronto, Canada. Association for Computational Linguistics.

Sam Wiseman and Alexander M. Rush. 2016. Sequence-to-sequence learning as beam-search optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1296–1306, Austin, Texas. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Sheng Xu, Yoshimi Suzuki, Jiyi Li, and Fumiyo Fukumoto. 2023. Decoupling style from contents for positive text reframing. In *Neural Information Processing*, pages 73–84, Singapore. Springer Nature Singapore.

Wei Xu, Alan Ritter, Bill Dolan, Ralph Grishman, and Colin Cherry. 2012. Paraphrasing for style. In *Proceedings of COLING 2012*, pages 2899–2914, Mumbai, India. The COLING 2012 Organizing Committee.

Boliang Zhang, Ajay Nagesh, and Kevin Knight. 2020. Parallel corpus filtering via pre-trained language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8545–8554, Online. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with BERT. *CoRR*, abs/1904.09675.

Zhemin Zhu, Delphine Bernhard, and Iryna Gurevych. 2010. A monolingual tree-based translation model for sentence simplification. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 1353–1361, Beijing, China. Coling 2010 Organizing Committee.

Caleb Ziems, Minzhi Li, Anthony Zhang, and Diyi Yang. 2022. Inducing positive perspectives with text reframing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3682–3700, Dublin, Ireland. Association for Computational Linguistics.

A Reframing Text Quality Evaluation

A.1 Problem Statement

The essence of existing TST metrics such as ROUGE and BLEU is to evaluate the similarity between the generated and reference sentence, so a simple copy can lead to a high score (Fan et al., 2018; Holtzman et al., 2019). And for an original sentence, there may be multiple corresponding reframed sentences, especially in the unconstrained case. Furthermore, existing metrics also cannot directly measure the degree of positive reframing. Therefore, this paper proposes a new metric RTQE (**R**eframing **T**ext **Q**uality **E**valuation), which aims to evaluate the degree of positive reframing relationship between the generated and original text that can avoid the limitation of only compared with human reference given in the dataset.

A.2 Evaluation Model

Taking the inspiration from Lai et al. (2021), the above problem is simplified into a binary classifi-

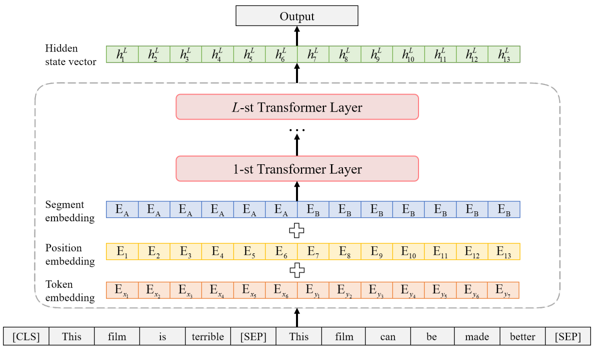

Figure 5: The model for RTQE.

cation task, i.e., judging whether there is a positive reframing relationship between two sentences. In practical evaluation, we regard the probability from the model prediction as the degree of positive reframing between the original and generated sentence. And the RTQE evaluation model established in this paper is shown in Figure 5. Given the original sentence x and the corresponding sentence y, we firstly concatenate them and input into the auto-encoding models such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) (without segment embedding). The encoder part is as follows:

$$H^e = \text{Encoder}([\text{CLS}], x, [\text{SEP}], y, [\text{SEP}]) \quad (12)$$

where [CLS] and [SEP] are special tokens.

The feature vector can be refined through L-layer transformer and the representation of H^l at the l-th layer ($l \in [1, L]$) is calculated as below:

$$H^l = \text{Transformer}_l(H^{l-1}), H^0 = H^e \quad (13)$$

We regard the hidden vector $H^{[\text{CLS}]}$ corresponding to [CLS] at the last layer as the contextualized representation of the whole sequence. And the prediction is obtained through the following equation:

$$\text{Output} = \text{Sigmoid}(W_o H^{[\text{CLS}]} + b_o) \quad (14)$$

where $W_o \in \mathbb{R}^{\text{dim}_H \times |y|}$ is the learnable parameter of the linear layer and b_o is the bias.

A.3 Dataset

As we simplified the RTQE task as a binary classification question, which determines whether two sentences constitute the positive reframing relationship. Therefore, this paper reconstructs the positive reframing dataset (Ziems et al., 2022) in the following way: for each original sentence, we consider its corresponding reframing sentence as a positive sample, and we pair the original sentence with itself or randomly select other reframing sentences to create negative samples, aiming to enhance the learning depth and generalization ability of the model. The specific statistics are presented in Table 9.

Set	Positive	Negative
Train	6679	13358
Dev	835	1670
Test	835	1670

Table 9: The statistics of the RTQE dataset.

A.4 Implementation Details

We use BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) as the backbone model respectively. For the base version, the model has 12 transformer encoder layers, and the hidden size is 768. For the large version, the model has 24 transformer encoder layers, and the hidden size is 1024. In this paper, the maximum text embedding length is set to 100 tokens, AdamW with an initial learning rate 1e-5 is used as the optimizer, and batch size is 32.

A.5 Experiment Results

This paper mainly tests the performance of four models: $\text{BERT}_{\text{base}}$, $\text{BERT}_{\text{large}}$, $\text{RoBERTa}_{\text{base}}$ and $\text{RoBERTa}_{\text{large}}$. And the experimental results are shown in Table 10.

Model	P(%)	R(%)	F1(%)	Acc(%)	Ref(%)
$\text{BERT}_{\text{base}}$	94.49	92.09	93.41	96.37	93.36
$\text{BERT}_{\text{large}}$	95.65	94.85	95.25	96.85	93.49
$\text{RoBERTa}_{\text{base}}$	94.52	94.97	94.74	96.48	94.59
$\text{RoBERTa}_{\text{large}}$	**96.16**	**96.05**	**96.11**	**97.41**	**95.98**

Table 10: The experimental results of RTQE task. The column of Ref refers to the average degree of positive reframing relationship between the human reference and original text in the test set obtained by our models. The best results are in **bold**.

It can be seen from Table 10 that the performance of RoBERTa is generally better than BERT on all metrics, and the large version is better than the base, which proves that the more parameters and training corpus the model has, the better its performance will be. In the end, $\text{RoBERTa}_{\text{large}}$ basically achieves the best results in all metrics and also reaches the F1 score of 95.98% and accuracy of 97.41% in the test of evaluating human reference, so finally this paper uses it as the evaluation model for RTQE.

Finally, we present the Pearson correlation between RTQE and manual evaluation in Table 11. It can be inferred that both the results of the T5-based models and BART-based models show a positive correlation with the three human evaluation metrics, particularly in terms of meaning preservation. This demonstrates that the introduction of the RTQE metric aligns with the task requirements, that is, positive reframing needs to prioritize maintaining the original meaning intact.

	Meaning	Positivity	Fluency
T5-based models	0.78	0.22	0.91
BART-based models	0.85	0.62	0.43

Table 11: Pearson correlation between RTQE and human evaluation.

B The Approach of Obtaining the Candidate Sentence

The approach of obtaining the candidate sentence set is as follows: when beam search is used, the number of candidate sentences with the same beam size can be returned directly, and beam size of 4, 5, and 6 are experimented in this paper; for Top-k sampling, the generated sentences of k = 30, 40, 50 and 60 are composed of candidate sentence set; for Top-p sampling, the generated sentences of p = 0.80, 0.85, 0.90 and 0.95 are selected to be composed the candidate sentence set; for Typical sampling, the sentences generated by τ = 0.20 and 0.95 are selected according to the settings recommended by Meister et al. (2023) to form the candidate sentence set.

C The Instruction for Human Evaluation

The specific instruction for human evaluation is as follows.

Give the **original sentence** with negative viewpoint and **reframed sentence** generated by our models. You need to score the Meaning Preservation (Meaning), Positivity and Fluency of the reframed sentence on a scale of 1 to 5.

Meaning: Indicate whether the reframed sentence preserves the original meaning.

1: Completely changed the original meaning.

3: Meaning related but with slight inconsistency or contradiction.

5: Faithful to the original meaning.

Choose **2** or **4** when you are hesitant.

Positivity: Indicate how positive the reframed sentence is.

1: As negative as the original sentence.

3: Neutral Sentiment, i.e. neither negative nor positive.

5: Very positive compared to the original sentence.

Choose **2** or **4** when you are hesitant.

Fluency: Indicate the fluency of the reframed sentence.

1: The reframed sentence does not make sense and it is unreadable.

3: The reframed sentence contains some minor grammatical errors, but does not affect reading.

5: The reframed sentence is human-like, without any grammatical errors.

Choose **2** or **4** when you are hesitant.

D Additional Results

D.1 Reframe Strategy Classification

We provide the detailed scores of our models on all classification evaluation metrics (i.e., accuracy, precision, recall, and F1 score) for others to compare and refer to, which can be found in Table 12.

D.2 Unconstrained Positive Reframing

For this task, we provide additional ablation results of unconstrained positive reframing in Table 13. It can be seen that when the positive sentiment reward is not used, the model's score on metrics such as ΔTB and RTQE decrease. And when the content preservation reward is not used, the model's performance on metrics such as ROUGE and BLEU may decline. In addition, multi-dimensional re-ranking can effectively improve the model's performance on content preservation-related metrics.

D.3 Controlled Positive Reframing

Here, we present the ablation results of multi-dimensional re-ranking under controlled setting in Table 14. It can be observed that when the strategy consistency evaluation is not used, the scores of $\mathrm{MSOF}_{\mathrm{Top-k}}$ on RTQE and PPL will decrease significantly, but it has better performance on ROUGE and BLEU. When the text similarity evaluation is not used, the performance of $\mathrm{MSOF}_{\mathrm{Top-k}}$ would significantly lower on content preservation-related metrics, but achieves best or sub-optimal results on ΔTB and RTQE. And when the fluency evaluation is not used, the model scores significantly lower on PPL, but still achieves sub-optimal results on RTQE and content preservation-related metrics. This paper suggests that the reason for the above phenomenon may be that the strategy consistency evaluation considers excessive content preservation as indicating incomplete reframing, and thus interacts with the text similarity evaluation. In addition, as can be seen from the results in the table, a decrease in text fluency (high PPL) is often accompanied by a decrease on ΔTB and RTQE. Therefore, there may be some positive correlation among them. Finally, although the overall framework does not achieve optimal results on all metrics, considering the performance of each variant model on each metric, choosing this way is the best trade-off at present.

D.4 Case Study

We provide the generated examples of unconstrained and controlled experiments in Tables 15 and 16. A comparative analysis reveals that our models generate outputs that are more diverse and comprehensive, while effectively preserving the underlying meaning of the original text. Specifically, the outputs of the BART-based models are mostly similar, except for the sentences generated by Typical sampling. On the other hand, the T5-based models outperform the BART-based models and baselines by providing the benefits of weekends consistent with human reference. Additionally, although the text in the dataset may contain colloquialisms and even grammatical errors, our models can generate more formal sentences that avoid these issues. Therefore, we speculate that further cleaning and filtering of the data in the dataset can further improve the model's performance. By comparing the results generated by the model in the unconstrained and controlled settings, it can be inferred that without reframe strategy, the reframing performance of the models will decrease, which proves that the reframing strategy plays an auxiliary role in helping the model generate results that better meet task requirements.

Finally, to further explore whether different reframe strategy will affect the generation results of the model, Table 17 shows the generation result of using different strategy to reframe the same negative text. It is obvious from the results that the model can generate reframing text with corresponding characteristics under the guidance of different reframe strategy, especially "Self-affirmation", "Thankfulness" and "Growth Mindset". This proves that the model can learn some information from the reframe strategy and it also shows that the research on controlled positive reframing is valuable.

Label	**Strategy-BERT**				**Strategy-RoBERTa**			
	P(%)	R(%)	F1(%)	Acc(%)	P(%)	R(%)	F1(%)	Acc(%)
Thankfulness	77.55	69.72	**73.43**	**93.41**	76.84	66.97	71.57	93.05
Neutralizing	52.75	72.84	**61.20**	66.59	58.70	62.58	60.58	**70.54**
Optimism	61.04	85.00	71.06	66.83	63.57	84.50	**72.69**	**69.58**
Impermanence	56.10	58.60	**57.32**	83.59	49.76	65.61	56.59	81.08
Growth Mindset	58.70	77.82	66.92	79.64	65.04	72.40	**68.52**	**82.40**
Self-affirmation	50.72	46.05	**48.28**	**91.02**	47.22	44.74	45.94	90.42

Table 12: The detailed experimental results of reframe strategy classification. We provide detailed experimental results of our models on all classification metrics here for analysis and comparison. And the best results in each label are in **bold**.

	Model	**R-1**	**R-2**	**R-L**	**BLEU**	**BScore**	**ΔTB**	**RTQE**	**PPL**
T5	$\text{MSOF}_{\text{Top}-\text{k}}$	**34.8**	14.7	27.7	10.1	**89.5**	**0.44**	93.5	**22.3**
	w.o Cls	34.6	**14.9**	**27.8**	**10.2**	89.5	0.42	93.5	22.6
	w.o Cont	34.0	14.5	27.4	9.6	89.4	0.39	**94.1**	23.6
	w.o Re-ranking	31.9	11.7	25.1	7.7	89.1	0.42	92.7	27.0
BART	$\text{MSOF}_{\text{Top}-\text{k}}$	34.8	14.9	**29.3**	12.0	**89.9**	**0.31**	**87.3**	**25.8**
	w.o Cls	**34.9**	**15.1**	29.1	**12.2**	89.8	0.31	85.6	30.2
	w.o Cont	34.7	15.0	29.0	12.2	89.8	0.27	84.1	30.5
	w.o Re-ranking	31.6	11.7	26.0	9.4	89.4	0.28	84.8	38.9

Table 13: The ablation experimental results of **unconstrained positive reframing**

	Model	**R-1**	**R-2**	**R-L**	**BLEU**	**BScore**	**ΔTB**	**RTQE**	**PPL**
T5	$\text{MSOF}_{\text{Top}-\text{k}}$	34.8	15.0	28.0	9.9	**89.5**	0.43	**97.7**	23.1
	w.o Strategy	**35.6**	**15.8**	**28.8**	**10.7**	89.5	0.41	95.0	30.0
	w.o Similar	32.2	12.1	25.4	7.6	89.2	**0.44**	97.5	**21.3**
	w.o Fluency	35.0	15.3	28.1	10.1	89.5	0.41	97.1	28.6
BART	$\text{MSOF}_{\text{Top}-\text{k}}$	**34.8**	**14.7**	**29.0**	11.4	**90.1**	0.36	**94.0**	**29.4**
	w.o Strategy	34.0	14.6	28.4	**11.8**	89.7	0.37	84.3	34.0
	w.o Similar	29.6	10.6	24.4	8.3	89.3	**0.41**	85.8	32.3
	w.o Fluency	33.9	14.3	28.2	11.6	89.7	0.35	86.2	46.9

Table 14: The ablation experimental results of **multi-dimensional re-ranking**. w.o Strategy means without strategy consistency evaluation, w.o Similarity represents without textual similarity evaluation and w.o Fluency represents not using fluency evaluation.

Original text	So glad that tomorrow is Friday. This has seriously been the longest week of my life
Reference	I'm glad the weekend is coming up, so I can rest .
T5 (Ziems et al., 2022)	This week has been a long one, but I'm sure it will be over soon.
FDSC (Xu et al., 2023)	I'm so glad that tomorrow is Friday. This week has been a long one.
ST2PG (Sheng et al., 2023)	I'm glad that tomorrow is Friday. This has been the longest week of my life.
PG2ST (Sheng et al., 2023)	I'm glad that tomorrow is Friday. This has been the longest week of my life.
MSOF$_{\text{Beam}}$	I'm glad that tomorrow is Friday. It's been a long week, but it's going to be a good one .
MSOF$_{\text{Top-k}}$	This week has been a long week, but I'm glad it's Friday. I'll be able to relax and enjoy the weekend .
MSOF$_{\text{Top-p}}$	It's been a long week, but it's a good chance to get some rest .
MSOF$_{\text{Typical}}$	I'm glad that tomorrow is Friday. This week has been challenging, but I'm going to get through it.
BART (Ziems et al., 2022)	I'm glad that tomorrow is Friday. This has been the longest week of my life, but I'm sure I'll get through it.
FDSC (Xu et al., 2023)	So glad that tomorrow is Friday. This has been the longest week of my life. I'm tired, but I'm sure I can get through it.
ST2PG (Sheng et al., 2023)	I'm glad that tomorrow is Friday. This has been the longest week of my life, but I'm sure it will be over soon.
PG2ST (Sheng et al., 2023)	I'm glad that tomorrow is Friday. This has been the longest week of my life, but I'm sure it will be over soon.
MSOF$_{\text{Beam}}$	I'm glad that tomorrow is Friday. This week has been very challenging.
MSOF$_{\text{Top-k}}$	I'm glad that tomorrow is Friday. This week has been very challenging.
MSOF$_{\text{Top-p}}$	I'm glad that tomorrow is Friday. This week has been so long.
MSOF$_{\text{Typical}}$	I'm glad that tomorrow is Friday. This week has been challenging, but I'm going to get through it.

Table 15: The reframing examples of **unconstrained positive reframing**. In order to better compare with the constrained settings. The pink text shows the positive perspective.

Original text	So glad that tomorrow is Friday. This has seriously been the longest week of my life! tired
Reference	I'm glad the weekend is coming up, so I can rest .
T5 (Ziems et al., 2022)	So glad that tomorrow is Friday. This has seriously been the longest week of my life. I'm tired, but I know I'll get through it.
MSOF$_{\text{Beam}}$	I'm so glad that tomorrow is Friday. This has been the longest week of my life, but I know that tomorrow will be a better day .
MSOF$_{\text{Top-k}}$	I'm glad that tomorrow is Friday. This has been the longest week of my life, and I've had a lot of fun .
MSOF$_{\text{Top-p}}$	I'm glad that tomorrow is Friday. This has been the longest week of my life, but I know it's going to be a great day.
MSOF$_{\text{Typical}}$	Tomorrow is Friday. This has been the longest week of my life, but I know I will make it to the end of the week. It will be great.
BART (Ziems et al., 2022)	I'm glad that tomorrow is Friday. This week has been long, but I'm looking forward to the weekend.
MSOF$_{\text{Beam}}$	I'm so glad that tomorrow is Friday. This has been the longest week of my life! I'm tired but I'm sure it will be good .
MSOF$_{\text{Top-k}}$	I'm really looking forward to Friday, so I can relax a bit .
MSOF$_{\text{Top-p}}$	I'm glad that tomorrow is Friday. I'm going to feel so much better .
MSOF$_{\text{Typical}}$	Even though Friday is the longest week in my life, I'm happy to have the chance to rest for a few days .

Table 16: The model comparison for reframing the same text and the reframing strategy is optimism. And we selected the same example as Table 15 to better compare the output of models under different settings.

Original text		I hate that I stress my self out so much that I can't fall asleep!
MSOF$_{\text{Top-k}}$(T5)	Growth Mindset	I need to take better care of myself so that I can fall asleep in no time! I'm going to try to reduce my stress and improve my sleep.
	Impermanence	I don't like that I stress myself out so much that I can't fall asleep, but I'm sure I'll get better soon.
	Neutralizing	I am stressed out so much that I can't fall asleep, but I'm going to take a nap and sleep better so I can sleep better.
	Optimism	I don't like to stress myself out so much that I can't fall asleep, but I'm sure I'll fall asleep soon.
	Self-affirmation	I don't like that I stress my self out so much that I can't fall asleep, but I'm a strong person, and I know I can do it.
	Thankfulness	I'm glad I have a bed to sleep in after a long day of stressing myself out, I can't sleep.
MSOF$_{\text{Top-k}}$(BART)	Growth Mindset	I'm going to stop stressing out about things so that I can fall asleep.
	Impermanence	I'm going to take some time to myself to clear my head.
	Neutralizing	Stress is part of life, and I can't fall asleep, but I'm sure I'll feel better soon.
	Optimism	I'm going to have to stay up all night tonight so that I can get some peace of mind.
	Self-affirmation	I am not able to sleep because of my stress. But I am a strong person, and I know I can get through this.
	Thankfulness	I'm thankful that I have a bed to sleep in when I'm stressed.

Table 17: A model comparison for reframing the same text using different reframe strategy