

Reinforced Rewards Framework for Text Style Transfer

Abhilasha Sancheti^{1(⊠)}, Kundan Krishna², Balaji Vasan Srinivasan³, and Anandhavelu Natarajan³

¹ University of Maryland, College Park, USA sancheti@cs.umd.edu
² Language Technologies Institute, Carnegie Mellon University, Pittsburgh, USA kundank@andrew.cmu.edu
³ Adobe Research, Bangalore, India {balsrini.anandyn}@adobe.com

Abstract. Style transfer deals with the algorithms to transfer the stylistic properties of a piece of text into that of another while ensuring that the core content is preserved. There has been a lot of interest in the field of text style transfer due to its wide application to tailored text generation. Existing works evaluate the style transfer models based on content preservation and transfer strength. In this work, we propose a reinforcement learning based framework that directly rewards the framework on these target metrics yielding a better transfer of the target style. We show the improved performance of our proposed framework based on automatic and human evaluation on three independent tasks: wherein we transfer the style of text from formal to informal, high excitement to low excitement, modern English to Shakespearean English, and vice-versa in all the three cases. Improved performance of the proposed framework over existing state-of-the-art frameworks indicates the viability of the approach.

Keywords: Style transfer \cdot Rewards \cdot Content preservation \cdot Transfer strength

1 Introduction

Text style transfer deals with transforming a given piece of text in such a way that the stylistic properties change to that of the target text while preserving the core content of the given text. This is an active area of research because of its wide applicability in the field of content creation including news rewriting, generating messages with a particular style to maintain the personality of a brand, etc. The stylistic properties may denote various linguistic phenomenon,

© Springer Nature Switzerland AG 2020

A. Sancheti and K. Krishna—This work was done while the authors were working at Adobe Research, Bangalore, India.

J. M. Jose et al. (Eds.): ECIR 2020, LNCS 12035, pp. 545–560, 2020. https://doi.org/10.1007/978-3-030-45439-5_36

from syntactic changes [7,23] to sentiment modifications [4,10,18] or extent of formality in a sentence [16].

Most of the existing works in this area either use copy-enriched sequenceto-sequence models [7] or employ an adversarial [4,15,18] or much simpler generative approaches [10] based on the disentanglement of style and content in text. On the other hand, more recent works like [19] and [3] perform the task of style transfer without disentangling style and content, as practically this condition cannot always be met. However, all of these works use word-level objective function (eg. cross-entropy) while training which is inconsistent with the desired metrics (content preservation and transfer strength) to be optimized in style transfer tasks. These metrics are generally calculated at a sentence-level and use of word level objective functions is not sufficient. Moreover, discreteness of these metrics makes it even harder to directly optimize the model over these metrics.

Recent advancements in Reinforcement Learning and its effectiveness in various NLP tasks like sequence modelling [8], abstractive summarization [14], and a related one machine translation [21] have motivated us to leverage reinforcement learning approaches in style transfer tasks.

In this paper, we propose a reinforcement learning (RL) based framework which adopts to optimize sequence-level objectives to perform text style transfer. Our reinforced rewards framework is based on a sequence-to-sequence model with attention [1,12] and copy-mechanism [7] to perform the task of text style transfer. The sentence generated by this model along with the ground truth sentence is passed to a content module and a style classifier which calculates the metric scores to finally obtain the reward values. These rewards are then propagated back to the sequence-to-sequence model in the form of loss terms.

The rest of our paper is organized as follows: we discuss related work on text style transfer in Sect. 2. The proposed reinforced rewards framework is introduced in Sect. 3. We evaluate our framework and report the results on formality transfer task in Sect. 4, on affective dimension like excitement in Sect. 5 and on Shakespearean-Modern English corpus in Sect. 6. In Sect. 7, we discuss few qualitative sample outputs. Finally, we conclude the paper in Sect. 8.

2 Related Work

Style transfer approaches can be broadly categorized as style transfer with parallel corpus and style transfer with non-parallel corpus.

Parallel corpus consists of input-output sentence pairs with mapping. Since such corpora are not readily available and difficult to curate, efforts here are limited. [23] introduced a parallel corpus of 30K sentence pairs to transfer Shakespearean English to modern English and benchmark various phrase-based machine translation methods for this task. [7] use a copy-enriched sequence-tosequence approach for Shakespearizing modern English and show that it outperforms the previous benchmarks by [23]. Recently, [16] introduced a parallel corpus of formal and informal sentences and benchmark various neural frameworks to transfer sentences across different formality levels. Our approach contributes in this field of parallel style transfer and extends the work by [7] by directly optimizing the metrics used for evaluating the style transfer tasks.

Another class of explorations are in the area of non-parallel text style transfer [4, 10, 15, 18] which does not require mapping between the input and output sentences. [4] compose a non-parallel dataset for paper-news titles and propose models to learn separate representations for style and content using adversarial frameworks. [18] assume a shared latent content distribution across a given corpora and propose a method that leverages refined alignment of latent representations to perform style transfer. [10] define style in terms of attributes (such as, sentiment) localized to parts of the sentence and learn to disentangle style from content in an unsupervised setting. Although these approaches perform well on the transfer task, content preservation is generally observed to be low due to the non-parallel nature of the data. Along this line, parallel style transfer approaches have shown better performance in benchmarks despite the data curation challenges [16].

Style transfer models are primarily evaluated on **content preservation** and **transfer strength**. But the existing approaches do not optimize on these metrics and rather teach the model to generate sentences to match the ground truth. This is partly because of the reliance on a differentiable training objective and discreteness of these metrics makes it challenging to differentiate the objective. Leveraging recent advancements in reinforcement learning approaches, we propose a reinforcement learning based text style transfer framework which directly optimizes the model on the desired evaluation metrics. Though there exists some prior work on reinforcement learning for machine translation [21], sequence modelling [8] and abstractive summarization [14] dealing model optimization for qualitative metrics like Rouge [11], they do not consider style aspects which is one of the main requirements of style transfer tasks. More recently, efforts [5,22] have been made to incorporate RL in style transfer tasks in a non-parallel setup. However, our work is in the field of parallel text style transfer which is not much explored.

Our work is different from these related works in the sense that we take care of content preservation and transfer strength with the use of a content module (to ensure content preservation) and cooperative style discriminator (style classifier) without explicitly separating content and style. We illustrate the improvement in the performance of the framework on the task of transferring text between different levels of formality [16]. Furthermore, we present the generalizability of the proposed approach by evaluating it on a self-curated excitement corpus as well as modern English to Shakespearean corpus [7].

3 Reinforced Rewards Framework

The proposed approach takes an input sentence $x = x_1 \dots x_l$ from source style s_1 and translates it to sentence $y = y_1 \dots y_m$ with style s_2 , where x and y are represented as a sequence of words. If x is given by (c_1, s_1) where c_1 represents the content and s_1 the style of the source, our objective is to generate $y = (c_1, s_2)$ which has same content as the source but with the target style.

Our approach is based on a copy-enriched sequence-to-sequence framework [7] which allows the model to retain factual parts of the text while changing the style specific text using an attention mechanism. At the time of training, the framework takes in the source style and the target style sentence as input to the attention based sequence-to-sequence encoder-decoder model. The words in the input sentence are mapped into an embedding space and the sentence is encoded into a latent space by the LSTM encoder. The network learns to pay attention to the words in the source sentence and creates a context vector based on the attention. The decoder model is a mixture of RNN and pointer (PTR) network where the RNN predicts the probability distribution over the vocabulary and the pointer network predicts the probability over the words in the input sentence based on the context vector. A weighted average of the two probabilities yields the final probability distribution at time step t given by,

$$P_t(w) = \delta P_t^{RNN}(w) + (1 - \delta) P_t^{PTR}(w),$$

where δ is computed based on encoder outputs and previous decoder hidden states. The decoder generates the transferred sentence by selecting the most probable word at each time step. This model is trained to minimize cross entropy loss given by

$$L_{ml} = -\sum_{t=1}^{m} \log(p(P_t(y_t^*))),$$

where m is the maximum length of the output sentence and y_t^* is the ground truth word at time t in the transferred sentence. While this framework optimizes for generating sentences close to the ground truth, it does not explicitly teach the network to preserve the content and generate sentences in target style. To achieve this, we introduce a style classifier and a content module which takes in the generated sentence from the sequence-to-sequence model along with the ground truth target sentence to provide reward to the sentence, as shown in Fig. 1. We leverage BLEU [13] score to measure the reward for preserving content and because of the lack of any formal score for transfer strength, we use a cooperative discriminator to provide score to the generated sentence. This score from the discriminator is used as a measure to reward for transfer strength. These rewards are then back propagated as explicit loss terms to penalize the network for incorrect generation.



Fig. 1. Model overview

3.1 Content Module: Rewarding Content Preservation

To preserve the content while transferring the style, we leverage Self-Critic Sequence Training (SCST) [17] approach and optimize the framework with BLEU scores as the reward. SCST is a policy gradient method for reinforcement learning and is used to train end-to-end models directly on non-differentiable metrics. We use BLEU score as reward for content preservation because it measures the overlap between the ground truth and the generated sentences. Teaching the network to favor this would result in high overlap with the ground truth and subsequently preserve the content of the source sentence since ground truth ensures this preservation.

We produce two output sentences y^s and y', where y^s is sampled from the distribution $p(y_t^s|y_{1:t-1}^s, x)$ at each decoding time step and y' (baseline output) is obtained by greedily maximizing the output distribution at each time step. The BLEU score between the sampled and greedy sequences is computed as the reward and the corresponding content-preservation loss is given by,

$$L_{cp} = (r(y') - r(y^s)) \sum_{t=1}^{m} \log(p(y_t^s | y_{1:t-1}^s, x)),$$

where the log term is the log likelihood on sampled sequence and the difference term is the difference between the reward (BLEU score) for the greedily sampled y' and multinomially sampled y^s sentences. Note that our formulation is flexible and does not require the metric to be differentiable because rewards are used as weights to the log-likelihood loss. Minimizing L_{cp} is equivalent to encouraging the model to generate sentences which have higher reward as compared to the baseline y' and thus increasing the reward expectation of the model. The framework can now be trained end to end by using this loss function along with the cross entropy loss to preserve the content of the source sentence in the transferred sentence.

3.2 Style Classifier: Rewarding Transfer Strength

To optimize the model to generate sentences which belong to the target style, it is possible to use a similar loss function as above and use it with the SCST framework [17]. However, that will require a formal measure for the target style aspect. Here, we present an alternate framework where such a formal measure is not readily available. We train a convolutional neural network based style classifier as proposed by [9] on the training dataset. This style classifier predicts the likelihood that an input sentence is in the target style, and the likelihood is taken as a proxy to the reward for style of a sentence and appended to a discriminatorbased loss function extended from [6]. Based on the transfer direction, we add the following term to the cross-entropy loss,

$$L_{ts} = \begin{cases} -\log(1 - s(y')), & \text{high to low level} \\ -\log(s(y')), & \text{low to high level} \end{cases}$$

In this formulation, y' is the greedily generated output from the decoder and s(y') is the likelihood score predicted by the classifier for y'. When transfer is done from high to low level of style, minimization of L_{ts} will encourage generation of sentences such that the classifier score is as low as possible. When the sentences are transferred from low to high level of style then the formulation ensures that the generated sentences have a score as high as possible. The framework is trained end-to-end using this loss function to generate the sentences which belong to the target style.

3.3 Training and Inference

The overall loss function thus can be written as a combination of the 3 loss functions,

$$Loss = \alpha L_{ml} + \beta L_{cp} + \gamma L_{ts}$$

We train various models using this loss function and different training methodologies (setting $\alpha = 1.0$, $\beta = 0.125$, $\gamma = 1.0$ after hyper-parameter tuning) as described in the next section. During the inference phase, the model predicts a probability distribution over the vocabulary based on the sentence generated so far and the word having the highest probability is chosen as the next word till the maximum length of the output sentence is reached. Note that unlike training phase in which case both the input and ground truth transferred sentences are available to the model, only the input sentence is made available to the model.

4 Experiments: Reinforcing Formality (GYAFC Dataset)

We evaluate the proposed approach on the GYAFC [16] dataset which is a parallel corpus for formal-informal text. We present the transfer task results in both the directions - formal to informal and vice-versa. This dataset (from Entertainment and Music domain) consists of \sim 56K informal-formal sentence pairs: \sim 52K in train, \sim 1.5K in test and \sim 2.5K in validation split.

We use both human and automatic evaluation measures for content preservation and transfer strength to illustrate the performance of the proposed approach.

Content preservation measures the degree to which the target style model outputs have the same meaning as the input style sentence. Following [16], we measure preservation of content using BLEU [13] score between the ground truth and the generated sentence since the ground truth ensures that content of the source style sentence is preserved in it. For human evaluation, we presented 50 randomly selected model outputs to the Mechanical turk annotators and requested them to rate the outputs on a Likert [2] scale of 6 as described in [16].

Transfer strength measures the degree to which style transfer was carried out. We reuse the classifiers that we built to provide rewards to the generated sentences (Sect. 3.2). A score above 0.5 from the classifier represents that the generated sentence belongs to the target style and to the source style otherwise. We define accuracy as the fraction of generated sentences which are classified to be in the target style. The higher the accuracy, higher is the transfer strength. For human evaluation, we ask the Mechanical turk annotators to rate the generated sentence on a Likert scale of 5 as described in [16].

Following [4] who illustrate the trade-off between the two metrics - content preservation and transfer strength, we combine the two evaluation measures and present an **overall score** for the transfer task since both the measures are central to different aspects of text style transfer task. The trade-off arises because the best content preservation can be achieved by simply copying the source sentence. However, the transfer strength in such scenario will be the worst. We compute overall score in the following way

$$Overall = \frac{BLEU \times Accuracy}{BLEU + Accuracy}$$

which is similar to F1-score since content preservation can be considered as measuring recall of the amount of source content retained in the target style sentence and transfer strength acts as a measure of precision with which the transfer task is carried out. In the above formulation, both BLEU and accuracy scores are normalized to be between 0 and 1.

We first ran an **ablation study** to demonstrate the improvement in performance of the model with introduction of the two loss terms in the various settings differing in the way training is being carried out. Below we provide details about each of the settings.

CopyNMT: Trained with L_{ml} **TS:** Trained with L_{ml} followed by $\alpha L_{ml} + \gamma L_{ts}$ **CP:** Trained with L_{ml} followed by $\alpha L_{ml} + \beta L_{cp}$ **TS+CP:** Trained with L_{ml} followed by $\alpha L_{ml} + \beta L_{cp} + \gamma L_{ts}$ **TS** \rightarrow **CP:** Trained with L_{ml} followed by $\alpha L_{ml} + \gamma L_{ts}$ and finally with $\alpha L_{ml} + \beta L_{cp}$ **CP** \rightarrow **TS:** Trained with L_{ml} followed by $\alpha L_{ml} + \beta L_{cp}$ and finally with $\alpha L_{ml} + \gamma L_{ts}$

 Table 1. Ablation study to demonstrate the improvement of the addition of the loss terms on formality transfer task.

| Models | Informal to Formal | | Formal to Informal | | | |
|---------------------------------------|--------------------|------------------------------|--------------------|-------------------------|------------------------------|--------------------|
| | $BLEU\uparrow$ | $\mathrm{Accuracy} \uparrow$ | $Overall\uparrow$ | $\mathrm{BLEU}\uparrow$ | $\mathrm{Accuracy} \uparrow$ | $Overall \uparrow$ |
| CopyNMT | 0.263 | 0.774 | 0.196 | 0.280 | 0.503 | 0.180 |
| TS | 0.240 | 0.801 | 0.184 | 0.271 | 0.527 | 0.179 |
| CP | 0.272 | 0.749 | 0.199 | 0.281 | 0.487 | 0.178 |
| TS+CP | 0.259 | 0.772 | 0.194 | 0.271 | 0.527 | 0.179 |
| $CP \rightarrow TS$ | 0.227 | 0.817 | 0.178 | 0.259 | 0.5441 | 0.175 |
| $\mathrm{TS}{\rightarrow}\mathrm{CP}$ | 0.286 | 0.723 | 0.205 | 0.298 | 0.516 | 0.189 |

Training with L_{ml} alone in all the above settings is done for 10 epochs with all the hyper-parameters set as default in the off-the-shelf implementation of [7]. Each of the iterative model training is done using the model with the best performance on validation set for 5 more epochs. We can observe from Table 1 that L_{ts} and L_{cp} helps in improving the accuracy which measures transfer strength (TS) and BLEU score which measures content preservation (CP) respectively as compared to CopyNMT. When all the three loss terms are used simultaneously (TS+CP)the resulting performance lies between TS and CP, indicating that there is a trade-off between the two metrics and improvement in one metric is at the cost of another as observed by [4]. This phenomenon is evident from the results of $TS \rightarrow CP$ and $CP \rightarrow TS$ where the network gets a bit biased towards the latter optimization. Moreover, improvement in $CP \rightarrow TS$ and $TS \rightarrow CP$ as compared to TS and CP respectively suggests that incremental training better helps in teaching the framework. Since the performance on both transfer strength and content preservation metrics plays an important role in text style transfer task. we chose $TS \rightarrow CP$, which has the maximum overall score, over the other models for further analysis.

Baselines: We compare the proposed approach $TS \rightarrow CP$ against the state-of-the-art cross-aligned autoencoder style transfer approach (Cross-Aligned) by

| Models | Informal to Formal | | | Formal to Informal | | |
|--------------------------------|--------------------------|-------------------|-------------------------------|--------------------------|-------------------|-------------------------------|
| | $\mathbf{BLEU}\uparrow$ | Accuracy ↑ | $\mathbf{Overall}^{\uparrow}$ | $\mathbf{BLEU}\uparrow$ | Accuracy ↑ | $\mathbf{Overall}^{\uparrow}$ |
| Transformer [20] | 0.125 | 0.933 | 0.110 | 0.099 | 0.894 | 0.089 |
| Cross-Aligned [18] | 0.116 | 0.670 | 0.098 | 0.117 | 0.766 | 0.101 |
| CopyNMT [7] | 0.263 | 0.774 | 0.196 | 0.280 | 0.503 | 0.180 |
| $TS \rightarrow CP (Proposed)$ | 0.286 | 0.723 | 0.205 | 0.298 | 0.516 | 0.189 |
| | Exciting to Non-exciting | | | Non-exciting to Exciting | | |
| Transformer [20] | 0.077 | 0.922 | 0.071 | 0.069 | 0.605 | 0.062 |
| Cross-Aligned [18] | 0.059 | 0.818 | 0.055 | 0.061 | 0.547 | 0.054 |
| CopyNMT [7] | 0.143 | 0.919 | 0.124 | 0.071 | 0.813 | 0.065 |
| $TS \rightarrow CP (Proposed)$ | 0.153 | 0.922 | 0.131 | 0.088 | 0.744 | 0.078 |
| | Modern to Shakespearean | | Shakesp | bearean to | Modern | |
| Transformer [20] | 0.027 | 0.736 | 0.026 | 0.046 | 0.915 | 0.043 |
| Cross-Aligned [18] | 0.044 | 0.614 | 0.041 | 0.049 | 0.537 | 0.044 |
| CopyNMT [7] | 0.104 | 0.495 | 0.085 | 0.111 | 0.596 | 0.093 |
| $TS \rightarrow CP (Proposed)$ | 0.127 | 0.489 | 0.100 | 0.137 | 0.567 | 0.110 |

Table 2. Comparison of $TS \rightarrow CP$ with the baselines on the three transfer tasks in both the directions. All the scores are normalized to be between 0 and 1.

 $[18]^1$, parallel style transfer approach (CopyNMT) by $[7]^2$ and neural encoderdecoder based transformer model $[20]^3$.

Results: It can be seen from Table 2 that even though the transformer model has the best accuracy, it fails in preserving the content. Closer look at the outputs (formal to informal transfer task in Table 4) reveal that it generates sentences in target style but the sentences do not preserve the meaning of the input and sometimes are out of context (discussed in the Sect. 7). Cross-Aligned performs the worst in informal to formal transfer task among all the other approaches because it is generating a lot of unknowns and is not able to preserve content. $TS \rightarrow CP$, on the other hand, has the highest overall score and performs the best in preserving the content. We also observed that the dataset had many sentences containing proper nouns like name of the songs, person or artists. In such cases, copy mechanism helps in retaining the proper nouns whereas other models are not able to do so. This is evident from the higher BLEU scores for our proposed model. Table 3 presents the human evaluation results aggregated over three annotators per sample. It can be seen that in at least 70% of the cases, annotators rated model outputs from $TS \rightarrow CP$ as better than the three baselines on both the evaluated metrics except for the content preservation as compared to CopyNMT in formal to informal task wherein, both the models perform equally good. One reason behind this is that both the models use copy-mechanism.

Table 3. Human evaluation results of 50 randomly selected model outputs. The values represent the % of times annotators rated model outputs from $TS \rightarrow CP$ (R) as better than the baseline CopyNMT (C), Transformer (T) and Cross-Aligned (S) over the metrics. I-F (E-NE) refers to informal to formal (exciting to non-exciting) task.

| Task | Transfer strength | | | Content preservation | | |
|------|---------------------------|-------|-------------|---------------------------|---------------|-------|
| | $\mathbf{R} > \mathbf{C}$ | R > T | $\rm R{>}S$ | $\mathbf{R} > \mathbf{C}$ | $\rm R {>} T$ | R > S |
| I-F | 88.67 | 81.34 | 70.00 | 70.00 | 72.67 | 83.67 |
| F-I | 73.34 | 88.67 | 61.22 | 59.34 | 79.34 | 91.80 |
| E-NE | 64.00 | 79.34 | 68.00 | 60.67 | 71.34 | 71.73 |
| NE-E | 76.67 | 70.67 | 68.00 | 69.34 | 74.00 | 70.00 |

¹ We use the off-the-shelf implementation provided by the authors at https://github. com/shentianxiao/language-style-transfer.

² https://github.com/harsh19/Shakespearizing-Modern-English.

³ https://github.com/pytorch/fairseq We also tried using the model proposed by [5] to compare against out proposed approach but we couldn't get stable performance on our datasets.

5 Experiments: Beyond Formality (Excitement Dataset)

In order to demonstrate the generalizability of our approach on an affective style dimension like excitement (the feeling of enthusiasm and eagerness), we curated our own dataset using reviews from Yelp dataset⁴ which is a subset of Yelp's businesses, reviews, and user data. We request human annotators to provide rewrites for given exciting sentences such that they sound as non-exciting/boring as possible. Reviews with rating greater than or equal to 3 were filtered out and considered as exciting to get the non-exciting/boring rewrites. We also asked the annotators to rate the given and transferred sentences on a Likert scale of 1 (No Excitement at all) to 5 (Very high Excitement). The dataset thus curated was split into train (~36K), test (1K) and validation (2K) sets. We evaluate the transfer quality on content preservation and transfer strength metrics as defined in Sect. 4.

For measuring the transfer strength we train a classifier as described in Sect. 3.2. We use the annotations provided by the human annotators on these sentences to get the labels for the two styles. Sentences with a rating greater than or equal to 3 were considered as exciting and non-exciting otherwise.

Results: The transfer task in this case is to convert the input sentence with high excitement (exciting) to a sentence with low excitement (non-exciting) and vice-versa. We can observe from Table 2 that model performance in the case of excitement transfer task is similar to what we observed in the formality transfer task. However, CopyNMT performs the best in transferring style in case of non-exciting to exciting transfer task because the model has picked up on expressive words ('awesome', 'great', and 'amazing') which helps in boosting the transfer strength. TS \rightarrow CP (with highest overall score) consistently outperforms Cross-Aligned in all the metrics and both the directions. Table 3 presents the human evaluation results on this transfer task. We notice that humans preferred outputs from our proposed model at least 60% of the times on both the measures as compared to the other three baselines. This provides an evidence that the proposed RL-based framework indeed helps in improving generation of more content preserving sentences which align with the target style.

6 Experiments: Beyond Affective Elements (English Dataset)

Besides affective style dimensions, our approach can also be extended to other style transfer tasks like converting modern English to Shakespearean English. To illustrate the performance of our model on this task we experimented with the corpus used in [7]. The dataset consists of ~ 21 K modern-Shakespearean English sentence pairs with ~ 18 K in train, ~ 1.5 K in test and ~ 1.2 K in validation split.

⁴ https://www.yelp.com/dataset.

We use the same evaluation measures as in the previous two tasks for illustrating the model performance and generalizability of the approach. For this task we present only the automatic evaluation results because manual evaluation of this task is not easy since it requires an understanding of Shakespearean english and finding such population is a difficult task due to limited availability.

Results: We can observe from Table 2 that model performance in the case of this transfer task is also similar to what we have observed in the earlier two transfer tasks. Although Cross-Aligned has better accuracy than $TS \rightarrow CP$, it fails to preserve the content (sample 3 of Table 6). Similar is the case with transformer which outperforms others in accuracy but is not able to retain the content (sample 1 of Table 6). TS \rightarrow CP outperforms the three baselines in preserving the content with the highest overall score. This establishes the viability of our approach to various types of text style transfer tasks.

These experiments further indicate that our proposed reinforcement learning framework improves the transfer strength and content preservation of parallel style transfer frameworks and is also generalizable across various stylistic expression.

7 Discussion

In this section, we provide few qualitative samples from the baselines and the proposed reinforcement learning based model. We can observe from the transformer model output for Input 1 and 2 in formal to informal column of Table 4 that it generates sentences with correct target style but does not preserve the content. It either adds random content or deletes the required content ('band' instead of 'better' in 1 and 'hot' instead of 'talented' in 2). As mentioned earlier, in sample output 3 of Table 4, Cross-Aligned is unable to retain the content and tend to generate unknown tokens. CopyNMT, even though is able to preserve content, tend to generate repeated token like 'please' in sample input 2 (informal to formal task) which results in lower BLEU score than our proposed approach. Transformer model outputs for exciting to non-exciting task in samples 1 and 2 of Table 5, miss specific content words like 'environment' and 'alisha' respectively. However, it is able to generate the sentences in target style. Similary, Cross-Aligned and CopyNMT are also not able to retain the name of the server in sample 2 of Table 5. Sample 2 of Shakespearean to Modern English and 1 of Modern to Shakespearean English task in Table 6 provide evidence for high accuracy and lower BLEU scores for transformer model. From sample 2 of Shakespearean to modern English transfer task, we can observe that Cross-Aligned although can generate the sentence in the target style is not able to preserve the entities like 'father' and 'child'. On the other hand, $TS \rightarrow CP$ can not only generate the sentences in the target style but is also able to retain the entities. There are few cases when CopyNMT is better in preserving the content **Table 4.** Sample model outputs and target style reference for Informal to Formal and Formal to Informal style transfer task. The first line is the source style sentence (input), second line is the reference output and the following lines correspond to the outputs from the baselines and the RL-based model.

| | Model | Informal to Formal | Formal to Informal |
|---|---------------------|---------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| 1 | Input | I want to be on TV! | I do not understand what that has to do with who's better looking? |
| | Reference | I would like to be on television | I don't know what the hell that has to do with who's better looking but OKAY! |
| | Transformer | I want to be on TV | I don't know what that's better looking with the band that do u? |
| | Cross-Aligned | I want to be on TV! | I do n't know that that do to have to talk of more better? |
| | CopyNMT | I would like to be on TV | I don't understand what that has to do with who's better looking for? |
| | $TS \rightarrow CP$ | I would like to be on TV | I don't understand what that has to do with who better? |
| 2 | Input | When you find out please let me know | I think that she is so talented, if she does not win, I am going to be really disappointed |
| | Reference | Please let me know when you find out | He is so talented, if she didn't win, I'd be really disappointed! |
| | Transformer | Keep me informed as soon as you know anything | I don't think she's hot, but i'm going to win so she'll win |
| | Cross-Aligned | If you find out please let me know | I think she is so funny, she doesn't win, I'm not sure to be gonna be cute |
| | CopyNMT | When you find out please please please me know? | I think she's so talented, she's not that i'm going to be really disappointed |
| | $TS \rightarrow CP$ | Please inform me if you find out | I think she is so talented, if she doesn't win, I'm gonna be really disappointed |
| 3 | Input | I dono I think that is the DUMBEST show EVER!!!!!! | Our mother is so unintelligent that she was hit & by a cop and told the police that she was mugged |
| | Reference | I don't think it's a very intelligent show | Your mama is so stupid, she got hit by a cop and told the police that she got mugged |
| | Transformer | I do not think that the show is appropriate | Your mama is so stupid that she sat on the ocean and said she was a bus |
| | Cross-Aligned | I think that I am $\langle unk \rangle$ the show $\langle unk \rangle \langle unk \rangle!$ | Yo mama is so fat that she had a $\langle unk \rangle$ and got a bunch of that's and she was $\langle unk \rangle$ |
| | CopyNMT | I am not sure that is the DUMBEST show EVER! | Your mama is so unintelligent she she hit hit cop and told the police that she was |
| | $TS \rightarrow CP$ | I think that is the DUMBEST show EVER! | Your mama is so unintelligent she got hit by a cop and told that she was so |

as compared to other models, for instance, sample 1 of formal to informal transfer task and sample 3 of non-exciting to exciting transfer task since it leverages copy-mechanism.

Another point to notice is the lexical level changes made to reflect the target style. For example, the use of 'would', 'don't' and 'inform' instead of 'want', 'dono' and 'let me know' respectively for transforming informal sentences into formal ones. Use of colloquial words like 'u', 'gonna' and 'mama' for converting the formal sentences to informal can be observed from the sample outputs. Not **Table 5.** Sample model outputs and target style reference for Exciting to Non-exciting and Non-exciting to Exciting style transfer task. The first line is the source style sentence (input), second line is the reference output and the following lines correspond to the outputs from the baselines and RL-based model.

| _ | Model | Exciting to Non-exciting | Non-exciting to Exciting |
|---|---------------------|-------------------------------------|----------------------------------------------------------------------------------------|
| 1 | Input | Delicious food and good environment | A good choice if you are in the phoenix area |
| | Reference | Good food and environment | A must visit if in the phoenix area |
| | Transformer | I recommend this food | If you're in the phoenix area, this is the place to |
| | | | go |
| | Cross-Aligned | Good food and good drinks | A great spot if you're in the area area |
| | CopyNMT | The food was good | This is a great choice of if you are in the phoenix area |
| | $TS \rightarrow CP$ | Good food and atmosphere | If you're in the phoenix area, this is a great choice if you're in the phoenix area |
| 2 | Input | Our server alisha was amazing | The food menu is reasonable and happy hour specials are good |
| | Reference | Our server alisha did a good job | Reasonable food menu and great happy hour specials |
| | Transformer | Our server was good | They have a great happy hour menu and the food is very good |
| | Cross-Aligned | Our server server was good | The food is great and happy hour prices are awesome |
| | CopyNMT | Our server was good | The food menu is great and the food is amazing |
| | $TS \rightarrow CP$ | Our server alisha was very good | The food menu is reasonable and happy hour specials are great |
| 3 | Input | The patio is amazing too | Acceptable food and beers with live music sometimes |
| | Reference | I like the patio also | Good food and great beers with occasional live music |
| | Transformer | The patio great | Live bands, good food and great beer |
| | Cross-Aligned | The patio is pretty good | Awesome food and great selection of music and music |
| | CopyNMT | The patio is good | Great food and great drinks and live music |
| | $TS \rightarrow CP$ | The patio is good | Great food, great beers, and great music |

only lexical level changes but structural transformations can also be observed as in 'Please inform me if you find out'. In case of excitement transfer task, use of strong expressive words like 'amazing' and 'great' makes the sentence sound more exciting while less expressive words such as 'okay' and 'good' makes the sentence less exciting. Use of 'thou' for you and 'hither' for here are more frequently used in Shakespearean English than in modern English. These sample outputs indeed provide an evidence that our model is able to learn these lexical or structural level differences in various transfer tasks, be it formality, beyond formality or beyond affective dimensions. **Table 6.** Sample model outputs and target style reference for Modern to Shakespearean English and Shakespearean to Modern English transfer task. The first line is the source style sentence (input), second line is the reference output and the following lines correspond to the outputs from the baselines and the RL-based model.

| | Model | Modern to Shakespearean | Shakespearean to Modern |
|---|---------------------|-----------------------------------------|-----------------------------------------------|
| 1 | Input | Don't you see that I'm out of breath? | Good morrow to you both |
| | Reference | Do you not see that I am out of breath? | Good morning to you both |
| | Transformer | Do you not hear me? | Good morning to you |
| | Cross-Aligned | Do you not think I had out of breath? | Good morrow to you |
| | CopyNMT | Do not see see I breath of breath? | Good morning, you both |
| | $TS \rightarrow CP$ | Do you not see that I am out of breath? | Good morning to you both |
| 2 | Input | Do you love me? | Well, well, thou hast a careful father, child |
| | Reference | Dost thou love me? | Well, well, you have a careful father, child |
| | Transformer | Do you love me? | Well, good luck |
| | Cross-Aligned | Dost thou love me? | Well, sir, be a man, Give it this |
| | CopyNMT | Do you love? | Well, well, you hast a father father, child |
| | $TS \rightarrow CP$ | Dost thou love me? | Well, well, you have a careful father, child |
| 3 | Input | Come here, man | Thou know'st my daughter's of a pretty age |
| | Reference | Come hither, man | You know how young my daughter is |
| | Transformer | Come, man | You are my daughter |
| | Cross-Aligned | Come hither, Iago | You know how noble my name is |
| | CopyNMT | Come hither, man | You know'st my daughter's age |
| | $TS \rightarrow CP$ | Come hither, man | You're know'st my daughter's of a pretty age |

8 Conclusion and Future Work

The primary contribution of this work is a reinforce rewards based sequence-tosequence model which explicitly optimizes over content preservation and transfer strength metrics for style transfer with parallel corpus. Initial results are promising and generalize to other stylistic characteristics as illustrated in our experimental sections. Leveraging this approach for simultaneously changing multiple stylistic properties (for e.g. high excitement and low formality) is a subject of further research.

References

- Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014)
- 2. Bertram, D.: Likert scales (2007). Accessed 2 Nov 2013
- Dai, N., Liang, J., Qiu, X., Huang, X.: Style transformer: unpaired text style transfer without disentangled latent representation. arXiv preprint arXiv:1905.05621 (2019)
- Fu, Z., Tan, X., Peng, N., Zhao, D., Yan, R.: Style transfer in text: exploration and evaluation. arXiv preprint arXiv:1711.06861 (2017)

- 5. Gong, H., Bhat, S., Wu, L., Xiong, J., Hwu, W.-H.: Reinforcement learning based text style transfer without parallel training corpus. arXiv preprint arXiv:1903.10671 (2019)
- Holtzman, A., Buys, J., Forbes, M., Bosselut, A., Golub, D., Choi, Y.: Learning to write with cooperative discriminators. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1638–1649. Association for Computational Linguistics (2018). http://aclweb.org/ anthology/P18-1152
- Jhamtani, H., Gangal, V., Hovy, E., Nyberg, E.: Shakespearizing modern language using copy-enriched sequence to sequence models. In: Proceedings of the Workshop on Stylistic Variation, pp. 10–19. Association for Computational Linguistics (2017). https://doi.org/10.18653/v1/W17-4902. http://aclweb.org/anthology/W17-4902
- Keneshloo, Y., Shi, T., Reddy, C.K., Ramakrishnan, N.: Deep reinforcement learning for sequence to sequence models. arXiv preprint arXiv:1805.09461 (2018)
- 9. Kim, Y.: Convolutional neural networks for sentence classification. arXiv preprint arXiv:1408.5882 (2014)
- Li, J., Jia, R., He, H., Liang, P.: Delete, retrieve, generate: a simple approach to sentiment and style transfer. arXiv preprint arXiv:1804.06437 (2018)
- 11. Lin, C.Y.: ROUGE: a package for automatic evaluation of summaries. Text Summarization Branches Out (2004)
- Luong, M.T., Pham, H., Manning, C.D.: Effective approaches to attention-based neural machine translation. arXiv preprint arXiv:1508.04025 (2015)
- Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, pp. 311–318. Association for Computational Linguistics (2002)
- Paulus, R., Xiong, C., Socher, R.: A deep reinforced model for abstractive summarization. arXiv preprint arXiv:1705.04304 (2017)
- Prabhumoye, S., Tsvetkov, Y., Salakhutdinov, R., Black, A.W.: Style transfer through back-translation. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 866–876. Association for Computational Linguistics (2018). http://aclweb.org/anthology/P18-1080
- 16. Rao, S., Tetreault, J.: Dear sir or madam, may i introduce the GYAFC dataset: corpus, benchmarks and metrics for formality style transfer. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Long Papers), vol. 1, pp. 129–140 (2018)
- Rennie, S.J., Marcheret, E., Mroueh, Y., Ross, J., Goel, V.: Self-critical sequence training for image captioning. In: CVPR, vol. 1, p. 3 (2017)
- Shen, T., Lei, T., Barzilay, R., Jaakkola, T.: Style transfer from non-parallel text by cross-alignment. In: Advances in Neural Information Processing Systems, pp. 6830–6841 (2017)
- Subramanian, S., Lample, G., Smith, E.M., Denoyer, L., Ranzato, M., Boureau, Y.L.: Multiple-attribute text style transfer. arXiv preprint arXiv:1811.00552 (2018)
- Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems, pp. 5998–6008 (2017)

- Wu, L., Tian, F., Qin, T., Lai, J., Liu, T.Y.: A study of reinforcement learning for neural machine translation. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 3612–3621. Association for Computational Linguistics (2018). http://aclweb.org/anthology/D18-1397
- Xu, J., et al.: Unpaired sentiment-to-sentiment translation: a cycled reinforcement learning approach. arXiv preprint arXiv:1805.05181 (2018)
- Xu, W., Ritter, A., Dolan, B., Grishman, R., Cherry, C.: Paraphrasing for style. In: Proceedings of COLING 2012, pp. 2899–2914 (2012)