

# Generic resources are what you need: Style transfer tasks without task-specific parallel training data

Huiyuan Lai, Antonio Toral, Malvina Nissim  
CLCG, University of Groningen / The Netherlands  
{h.lai, a.toral.ruiz, m.nissim}@rug.nl

## Abstract

Style transfer aims to rewrite a source text in a different target style while preserving its content. We propose a novel approach to this task that leverages generic resources, and without using any task-specific parallel (source–target) data outperforms existing unsupervised approaches on the two most popular style transfer tasks: formality transfer and polarity swap. In practice, we adopt a multi-step procedure which builds on a generic pre-trained sequence-to-sequence model (BART). First, we strengthen the model’s ability to *rewrite* by further pre-training BART on both an existing collection of generic paraphrases, as well as on synthetic pairs created using a general-purpose lexical resource. Second, through an iterative back-translation approach, we train two models, each in a transfer direction, so that they can provide each other with synthetically generated pairs, dynamically in the training process. Lastly, we let our best resulting model generate static synthetic pairs to be used in a supervised training regime. Besides methodology and state-of-the-art results, a core contribution of this work is a reflection on the nature of the two tasks we address, and how their differences are highlighted by their response to our approach.

## 1 Introduction

Text style transfer is, broadly put, the task converting a text of one style into another while preserving its content. In its recent tradition within Natural Language Generation (NLG), two tasks and their corresponding datasets have been commonly used (Zhang et al., 2018; Luo et al., 2019; Wu et al., 2019; Yi et al., 2020; Zhou et al., 2020). One dataset was specifically created for *formality transfer* and contains parallel data (GYAFC (Rao and Tetreault, 2018)), while the other one contains a large amount of non-parallel sentiment labelled texts (YELP (Li et al., 2018)), with parallel pairs

for test, and is used for the task of *polarity swap*. Examples from these datasets are shown in Table 1.

The two tasks are usually conflated in the literature under the general *style transfer* label and addressed with the same methods, but we find this an oversimplification. Formality transfer implies rewriting a formal sentence into its informal counterpart (or viceversa) while preserving its meaning. Polarity swap, instead, aims to change a positive text into a negative one (or viceversa); and while the general theme must be preserved, the meaning is by definition not maintained (e.g. “I hated that film” → “I loved that film”). In line with previous work, we also address both tasks in a similar way, but this is actually to unveil how their different nature affects modelling and evaluation.

Due to the general scarcity of parallel data, previous works mainly adopted unsupervised approaches, dubbed *unpaired* methods (Dai et al., 2019) since they do not rely on labelled training pairs. However, it has also been shown that best results, unsurprisingly, can be achieved if parallel training data (such as the formality dataset (Rao and Tetreault, 2018)) is available (Sancheti et al., 2020; Lai et al., 2021). For this reason, substantial work has gone into the creation of artificial training pairs through various methods (see Section 2); approaches using synthetic pairs are thus still considered unsupervised in the style transfer literature, since they do not use manually labelled data.

We explore how parallel data can best be derived and integrated in a general style transfer framework. To do so, we create pairs in a variety of ways and use them in different stages of our framework. A core aspect of our approach is leveraging generic resources to derive training pairs, both natural and synthetic. On the natural front, we use abundant data from a generic rewriting task: paraphrasing. As for synthetic data, we leverage a general-purpose computational lexicon using its antonymy relation to generate polarity pairs.

In practice, we propose a framework that adopts a multi-step procedure which builds upon a general-purpose pre-trained sequence-to-sequence (seq2seq) model. First, we strengthen the model’s ability to rewrite by conducting a second phase of pre-training on natural pairs derived from an existing collection of generic paraphrases, as well as on synthetic pairs created using a general-purpose lexical resource. Second, through an iterative back-translation (Hoang et al., 2018) approach, we train two models, each in a transfer direction, so that they can provide each other with synthetically generated pairs on-the-fly. Lastly, we use our best resulting model to generate static synthetic pairs, which are then used offline as parallel training data.

**Contributions** Using a large pre-trained seq2seq model (1) we achieve state-of-the-art results for the two most popular style transfer tasks without task-specific parallel data. We show that (2) generic resources can be leveraged to derive parallel data for additional model pre-training, which boosts performance substantially and that (3) an iterative back-translation setting where models in the two transfer directions are trained simultaneously is successful, especially if enriched with a reward strategy. We also offer (4) a theoretical contribution over the nature of the two tasks: while they are usually treated as the same task, our results suggest that they could possibly be treated separately.<sup>1</sup>

## 2 Related Work

Style transfer is most successful if task-specific parallel data is available, as in the case of formality transfer (Rao and Tetreault, 2018). Like in most NLP tasks, large pre-trained models have been shown to provide an excellent base for fine-tuning in a supervised setting (Chawla and Yang, 2020; Lai et al., 2021).

Since parallel data for fine-tuning such large models for style transfer is scarce, a substantial amount of work has gone into methods for creating artificial sentence pairs so that models can be trained in a supervised regime.

One way to do this is to artificially generate parallel data via back-translation, so that training pairs are created on-the-fly during the training process itself (Zhang et al., 2018; Lample et al., 2019; Prabhunoye et al., 2018; Luo et al., 2019). In these systems, one direction’s outputs and its inputs can

be used as pairs to train the model of the opposite transfer direction.

Another common strategy is to use style-word-editing (Li et al., 2018; Xu et al., 2018; Wu et al., 2019; Lee, 2020) to explicitly separate content and style. These approaches first detect relevant words in the source and then do operations like deleting, inserting and combining to create the pair’s target. Back-transferring is generally used to reconstruct the source sentence for training, so that pairs are also made on-the-fly.

Lample et al. (2019) provide evidence that disentangling style and content to learn distinct representations (Shen et al., 2017; Fu et al., 2018; John et al., 2019; Yi et al., 2020) is not necessary. Reconstructing the source, instead, appears beneficial: it is used by Dai et al. (2019) who pre-train a model on style transfer data with the Transformer architecture (Vaswani et al., 2017); and by Zhou et al. (2020), who use an attentional seq2seq model that pre-trains the model to reconstruct the source sentence and re-predict its word-level style relevance.

Luo et al. (2019) pre-train a LSTM-based seq2seq model (Bahdanau et al., 2015) using sentence pairs generated by a template-based baseline. More recently, Li et al. (2020a) proposed a two-stage strategy of search and learning for formality transfer where they perform a simulated annealing search (Liu et al., 2020) to obtain output sentences as pseudo-references, and then fine-tune GPT-2 (Radford et al., 2019) with the resulting pairs.

The methods above create task-specific artificial pairs, some using pre-crafted manual rules or templates. We aim to overcome this by exploiting generic resources. Additionally, it is not evident which strategy works best for creating parallel data, whether offline or on-the-fly, and the simultaneous advantage of both strategies has not been fully explored. Lastly, Chawla and Yang (2020) develop a semi-supervised model based on sequence-to-sequence pre-trained model (BART, Lewis et al. (2020)) using parallel training data and large amounts of non-parallel data, which achieves a significant performance. In previous work, we have also shown that a sequence-to-sequence pre-trained model (BART) outperforms a language model (GPT-2) in content preservation and overall performance when task-specific parallel training data is available (Lai et al., 2021).

Therefore, we use BART as generic base model; we enrich it with iterative back-translation to cre-

<sup>1</sup>All code at <https://github.com/laihuiyuan/Generic-resources-for-TST>.

DATASET	STYLE	SENTENCE-PAIR	
GYAFC	Informal	that is just my gut feeling.	no different between ages if the mind is near to eachother
	Formal	That is my personal opinion.	There is no difference between ages if the intellect is similar.
YELP	Negative	this branch is getting worse and worse.	bad service in these areas and really ruined our visit.
	Positive	this branch is getting better and better.	good service in these areas and really made our visit.
PARABANK 2	Source	The bank is coming up on your left.	I guess I've always been pretty good with words.
	Target	You have the bank on the left side.	I think narrating has always been my strong suit.

Table 1: Samples of each dataset.

DATASET	STYLE	PAIRED			UNPAIRED	
		Train	Valid	Test	Train	Valid
GYAFC [F&R]	Informal	51,967	2,788	1,332	N/A	N/A
	Formal	51,967	2,247	1,019	N/A	N/A
YELP	Negative	N/A	N/A	500	177,218	2,000
	Positive	N/A	N/A	500	266,041	2,000
PARABANK 2	Source	1,132,289	N/A	N/A	N/A	N/A
	Target	1,132,289	N/A	N/A	N/A	N/A

Table 2: Dataset Statistics.

ate training pairs on-the-fly. We also explore the advantage of further pre-training by creating pairs through generic resources, as well as the benefits of a final training using generated pairs.

### 3 Tasks, Datasets, and Evaluation

#### 3.1 Tasks and Datasets

The task of style transfer is generally defined as the conversion of a text written in a given style to approximately the same text in a different style: style should be changed while preserving the original “content”. We focus on the two most popular tasks, namely *formality transfer* and *polarity swap*, and use the two standard available datasets. Example pairs are shown in Table 1; statistics are in Table 2.

**Formality Transfer Dataset** Grammarly’s Yahoo Answers Formality Corpus (GYAFC) (Rao and Tetreault, 2018) is a dataset containing aligned formal and informal sentences from two domains: Entertainment & Music (E&M) and Family & Relationships (F&R). Parallel pairs are provided for training, validation, and test, with four human references for every test sentence. In the experiments we report in this paper we use data from the F&R domain, which is the one more commonly used.

**Polarity Swap Dataset** YELP is a dataset of business reviews on Yelp (with scores 1–5) processed by Li et al. (2018). Samples with a score greater than 3 are considered as positive otherwise they are negative. The dataset comes in the form of large amounts of non-parallel data for training and development, while parallel pairs are provided for

evaluation. For each test sentence, Li et al. (2018) provide one human reference; three additional human references are released by Luo et al. (2019).

Although these two tasks have been conflated in previous work as “style transfer”, they are not exactly the same, which we hypothesise affects both their modelling and evaluation. More specifically, in polarity swap the actual content is not exactly preserved (the message is actually the opposite), rather it’s the general “theme/topic” that needs to be preserved. In formality transfer, instead, the “translation” happens really more at style level, and content needs to stay the same. This is evident if we look at examples in Table 1 (top two blocks). In YELP, we can see that the theme-related words are expected to stay while changing the polarity words. Therefore, although the two sentences refer to the same event/concept, they convey opposite meanings. On the contrary, in formality transfer, an informal text should be changed into a formal one, but the overall meaning should be preserved. In this sense, formality transfer can be seen much more as rewriting than polarity swap and can be conceived akin to the more general task of paraphrasing.

Leveraging this observation, we explore if paraphrase pairs can be used to make the model learn the basic task of “rewriting” in a first stage. The advantage of using paraphrases is the large amount of parallel data available. Specifically, we use PARABANK 2, a large-scale, diverse, collection of paraphrases (Hu et al., 2019). Given the different nature of the two tasks, we expect this strategy to help more formality transfer than polarity swap, since the latter is much less of a rewriting task than the former. In spite of the differences highlighted above, we approach both tasks within the same framework for two reasons: (i) to compare to previous works, which have treated the tasks as manifestations of the same “style transfer” task; but also (ii) to observe if and how the tasks respond differently to modelling and evaluation metrics.




Figure 1: General overview of our pipeline.

### 3.2 Task Evaluation

The performance of text style transfer is commonly assessed on style strength and content preservation. For style strength, using a pre-trained style classifier is the most popular automatic evaluation strategy. For content preservation,  $n$ -gram-based matching metrics such as BLEU (Papineni et al., 2002) are most commonly used. However, these metrics usually fail to recognise information beyond the lexical level. Since word embeddings (Mikolov et al., 2013; Pennington et al., 2014) have become the prime alternative to  $n$ -gram-based matching to capture similarity, embeddings-based metrics have also been developed (Fu et al., 2018). However, embedding-based metrics like cosine similarity still work at the token-level, and might fail to capture the overall semantics of a sentence.

To overcome such limitations, recent work has developed *learnable metrics*, which attempt to directly optimize correlation with human judgments. These metrics, with the prime examples of BLEURT (Sellam et al., 2020) and COMET (Rei et al., 2020), have recently shown promising results in machine translation evaluation. To the best of our knowledge, only our previous work used BLEURT in the evaluation of formality style transfer models (Lai et al., 2021); we are now proposing to use it also for the evaluation of polarity swap, and to add COMET to the pool of evaluation metrics to be systematically adopted in the evaluation of text style transfer tasks.

Therefore, in addition to BLEU, which allows us to compare to previous work, we also use BLEURT and COMET. Let us bear in mind that “content preservation” does not mean exactly the same thing for the two tasks that we consider (cf. Section 3.1), so that we might observe different reactions to different evaluation measures for the two tasks.

## 4 Approach

We propose a framework that adopts a multi-step procedure on top of the large pre-trained seq2seq

model BART (Lewis et al., 2020).

Given a source sentence  $x = \{x_1, \dots, x_n\}$  of length  $n$  with style  $s_1$ , the goal of text style transfer is to generate a sentence  $y$  with style  $s_2$ , preserving the source sentence’s meaning in formality transfer or the source sentence’s theme in polarity swap.<sup>2</sup> Formally, the objective is to minimize the following negative log likelihood:

$$L(\phi) = -\sum_i \log(p(y_i | y_{1:i-1}, x; \phi)) \quad (1)$$

where  $\phi$  are the parameters of BART.

Our framework can be conceived as a pipeline, visualised in Figure 1. At the core of the framework are two BART models (model A and model B), one for each transfer direction. Since the main challenge in unpaired style transfer is that we cannot directly employ supervision (i.e. task-specific parallel training pairs), we explore and evaluate different ways of creating and using sentence pairs at different stages of the pipeline.

First, we strengthen the model’s ability to rewrite by conducting a second phase of pre-training on natural pairs derived from an existing collection of generic paraphrases, as well as on synthetic pairs created using a general-purpose lexical resource (Step 1, Section 4.1).

Second, we use iterative back-translation with several reward strategies to train the two models in both transfer directions simultaneously; sentence pairs are created on the fly (Step 2, Section 4.2).

Third, we create high-quality synthetic pairs using our best systems from the previous step, to create a static resource of parallel data that can be used to train new transfer models (Step 3, Section 4.3).

### 4.1 Further Pre-training: Learning to Rewrite

As hinted at in Section 3, style transfer can be seen as a specific way of *paraphrasing*. On the basis

<sup>2</sup>In what follows, we use the term “content” in a more general way to refer to both cases.

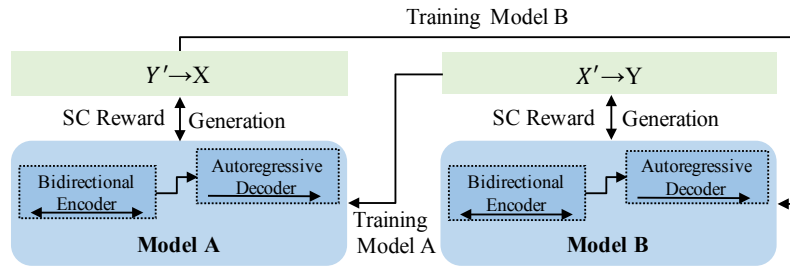


Figure 2: General overview of IBT training.

of this intuition, we hypothesise that generic paraphrase data, which already exists in much larger amounts than task-specific style transfer data, can be useful for text style transfer in terms of teaching the models the more generic task of “rewriting”. For polarity swap, which is less of a rewriting task than formality transfer, as the meaning is reversed rather than preserved, we also create synthetic pairs using a general-purpose lexical resource.

Using the natural and the synthetic pairs we conduct a second phase of pre-training. We expect this strategy to help specifically with content preservation, which is known to be the most difficult part of style transfer, especially in an unsupervised setting (Sancheti et al., 2020; Lai et al., 2021).

**Generic Training Pairs** We use data from PARABANK 2 to make the model learn the basic task of “rewriting”. We use this dataset in its entirety or filtered (models M1.1 and M1.2 in Table 3). In the first case, the whole of the paraphrase pairs from PARABANK 2 are used to further pre-train the model. In the second case, we follow the rationale that not all pairs are equally relevant for our tasks, and selecting task-specific ones could be beneficial. For instance, while both PARABANK 2 pairs in Table 1 are good examples of rewriting, the one on the right is more meaningful in terms of formality transfer. Therefore, we train two binary style classifiers, one for formality and one for polarity, using TextCNN (Kim, 2014) on the training sets of GYAFC and YELP. These classifiers are then used to automatically select more strongly style-opposed pairs. The resulting filtered paraphrase subset  $D_p$  is such a set of pairs:

$$D_p = \{(\mathbf{x}, \mathbf{y}) | (p(s_1|\mathbf{x}) + p(s_2|\mathbf{y}))/2 > \sigma\} \quad (2)$$

where  $p(s_i|*)$  is the probability of a sentence being a style  $s_i$ , predicted by the style classifier, and  $\sigma$  is the threshold for data selection<sup>3</sup>;  $\mathbf{x}$  and  $\mathbf{y}$  constitute

<sup>3</sup> $\sigma = 0.85$  in our experiments.

the sentence pair.

**Synthetic Pairs for Polarity Swap** Due to the nature of polarity swap, we expect that even filtered paraphrases might not benefit polarity swap as much as formality. We therefore add another strategy to enhance polarity swap rewriting and create pairs for further pre-training exploiting a general-purpose lexical resource (model M1.3 in Table 3). Specifically, we use SentiWordNet (Baccianella et al., 2010) to obtain words’ sentiment scores to detect the polarity of each word in the sentence. To maximise the quality of synthetic pairs, we select sentences that contain one polarity word only, and swap that one with its WordNet antonym (Miller, 1995). The new synthetic sentence is regarded as the target sentence corresponding to the original sentence.

The generic/filtered/synthetic pairs are used for a second phase of seq2seq pre-training for BART. Examples of these pairs are in Appendix A.5.

#### 4.2 Iterative Back-translation and Rewards: Pairs on-the-fly

After further pre-training BART, we use iterative back-translation to train two models, each in a transfer direction, so that they can provide one another with synthetically generated pairs on-the-fly. We obtain pseudo-parallel data via back-transfer: the outputs of one direction are used to provide the supervision to train the model of the opposite direction (Figure 2). To explicitly guide the model to preserve the content and to apply the target style, we add content and style rewards in a reinforcement learning fashion (models M2.\* in Table 3).

**Rewarding Style Strength** To provide an explicit signal to teach the model to change the sentence’s style, a style classifier (SC) based reward is used to push the model to change the sentence into the target style. For this SC reward, which evaluates

how well the transferred sentence  $\mathbf{y}'$  matches the target style, we reuse the style classifier trained for selecting paraphrase data (Section 4.1). The SC’s confidence in each transfer direction is

$$p(s_i|\mathbf{y}') = \text{softmax}_i(\text{TextCNN}(\mathbf{y}', \theta)) \quad (3)$$

where  $i = \{1,2\}$  and  $\theta$  are the parameters of the style classifier, fixed during training transfer models. Formally, the reward is

$$R_{sc} = \lambda_{sc}[p(s_2|\mathbf{y}') - p(s_1|\mathbf{y}')] \quad (4)$$

where  $s_1$  and  $s_2$  are source style and target style, respectively.  $\mathbf{y}'$  is the generated target sentence sampled from the distribution of model outputs at each decoding time step.

We apply the SC reward in two ways: in the supervised training process using pseudo-parallel data (SC0); and in the process of *generating* pseudo-parallel data itself (SC1). For the latter, we generate text in the target style by sampling the distribution of model outputs, while at the same time use the SC reward to feed back its corresponding style signals to the model.

**Rewarding Content Preservation** Following Sancheti et al. (2020), we use a BLEU-based reward, formulated as follows:

$$R_{\text{bleu}} = \lambda_{\text{bleu}}[\text{BLEU}(\mathbf{y}'_{s_i}, \mathbf{x}) - \text{BLEU}(\mathbf{y}_{s_i}^s, \mathbf{x})] \quad (5)$$

where  $\mathbf{y}_{s_i}^s$  is the generated sentence in target style  $s_i$  sampled from the distribution of model outputs at each time step in decoding, and  $\mathbf{y}'_{s_i}$  is obtained by greedily maximizing the distribution.

Since new-generation metrics show promising results in evaluation (Section 3), we use BLEURT also as an alternative metric to BLEU in the reward strategy, expecting it might be better at measuring semantics at the sentence level. Formally, we formulate the BLEURT-based reward as

$$R_{\text{bleurt}} = \lambda_{\text{bleurt}}[\text{BLEURT}(\mathbf{y}_{s_i}^s, \mathbf{x})] \quad (6)$$

where  $\mathbf{y}_{s_i}^s$  is the generated sentence in target style  $s_i$  sampled from the distribution of model outputs.

**Gradients and Objectives** We use the policy gradient algorithm (Williams, 1992) to maximize the expected reward of the generated sentence  $\mathbf{y}^s$ , whose gradient with respect to the parameters  $\phi$  of the neural network model is estimated by sampling:

$$\nabla_{\phi} J(\phi) = E[R \cdot \nabla_{\phi} \log(P(\mathbf{y}^s|\mathbf{x}; \phi))] \quad (7)$$

where  $\nabla_{\phi} J(\cdot)$  is the gradient of objective function  $J(\cdot)$  with respect to model parameters  $\phi$ ,  $E(\cdot)$  is the expectation,  $R$  is the reward of the sequence  $\mathbf{y}^s$  that is sampled from the distribution of model outputs at each decoding time step. The overall objectives are the combination of the base model’s loss (Eq. 1) and the policy gradient of rewards (Eq. 7) which are used to train our framework end-to-end.

### 4.3 Final Training: High-quality Pairs

As a final step, we let our best models generate pairs to create a static resource of parallel data. We feed the system source sentences randomly picked from the training sets and generate the corresponding sentences in the target style. We then select high-quality pairs using BLEURT and our style classifier. The resulting dataset  $\mathbf{D}_h$  is a set of pairs:

$$\mathbf{D}_h = \{(\mathbf{x}, \mathbf{y}') | \text{BLEURT}(\mathbf{x}, \mathbf{y}') > \sigma_c \text{ and } (p(s_1|\mathbf{x}) + p(s_2|\mathbf{y}'))/2 > \sigma_s\} \quad (8)$$

where  $\mathbf{x}$  and  $\mathbf{y}'$  are the source sentence and generated sentence, respectively.  $p(s_i|*)$  is the probability of a sentence being of style  $s_i$  as predicted by the style classifier, and  $\sigma_*$  is the threshold for data selection regarding content and style.<sup>4</sup>

Finally, these pairs are used to fine-tune the original BART with all reward strategies, so as to train new transfer models in a supervised way (model M3.1 in Table 3).

## 5 Experiments

All experiments are implemented atop Huggingface Transformers (Wolf et al., 2020), taking the BART base model (139M parameters) for our experiments. We train our framework using the Adam optimiser (Kingma and Ba, 2015) with the initial learning rate  $1e^{-5}$ . The batch size is set to 32. The final values  $\lambda$  for style and content rewards are both set to 1 based on validation results. Both WordNet and SentiWordNet are used from NLTK<sup>5</sup>.

### 5.1 Evaluation Metrics

To assess style and content we use common metrics for this task. For content preservation we add two learnable metrics, which we hope will be adopted from now on, to glean better insights into the systems’ behaviour in the two tasks (Section 3.2).

We measure **style strength** automatically by evaluating the target style accuracy of transferred

<sup>4</sup> $\sigma_c = 0.15$  and  $\sigma_s = 0.9$  in our experiments.

<sup>5</sup><https://www.nltk.org/>

MODEL	DATASET	GYAFC (FORMALITY TRANSFER)					YELP (POLARITY SWAP)				
		BLEURT	COMET	BLEU	ACC	HM	BLEURT	COMET	BLEU	ACC	HM
M0: Original BART		-0.116	0.242	0.414	0.333	0.369	-0.388	-0.146	0.309	0.022	0.041
STEP 1: Further pre-training											
M1.1: Further pre-training using whole dataset		0.012	0.209	0.420	0.357	0.386	-0.412	-0.282	0.179	0.040	0.065
M1.2: Further pre-training using subset		0.011	0.225	0.441	0.693	0.539	-0.347	-0.178	0.247	0.166	0.199
M1.3: Further pre-training using synthetic data		-	-	-	-	-	-0.321	-0.074	0.326	0.189	0.239
STEP 2: IBT + Rewards											
M2.1: IBT + all rewards with M0		-0.010	0.292	0.507	0.836	0.631	-0.229	-0.017	0.298	0.826	0.438
M2.2: IBT + all rewards with M1.2		<b>0.041</b>	0.318	0.553	<b>0.932</b>	<b>0.694</b>	-0.176	0.026	0.295	0.853	0.438
M2.3: IBT + all rewards with M1.3		-	-	-	-	-	-0.246	-0.035	0.302	0.884	0.450
M2.4: M2.2 except BLEURT		0.033	0.313	0.552	0.929	0.693	-0.187	0.001	0.285	0.860	0.428
M2.5: M2.2 except BLEU		<b>0.041</b>	0.320	0.551	0.925	0.691	<b>-0.149</b>	0.031	0.295	0.784	0.429
M2.6: M2.2 except SC0		0.024	<b>0.321</b>	0.544	0.928	0.686	-0.195	-0.016	0.286	0.881	0.432
M2.7: M2.2 except SC1		0.039	0.318	0.555	0.873	0.679	-0.176	0.039	<b>0.331</b>	0.500	0.398
STEP 3: Offline training (Model used: original BART + Rewards)											
M3.1: training pairs generated with M2.2 (GYAFC) / M2.3 (YELP)		0.030	<b>0.321</b>	<b>0.560</b>	0.904	0.692	-0.183	<b>0.046</b>	0.316	<b>0.887</b>	<b>0.466</b>
M3.2: training pairs are subset of paraphrase data (same as in M1.2)		0.012	0.229	0.455	0.783	0.576	-0.338	-0.221	0.215	0.457	0.292

Table 3: Results for the different steps of the pipeline. SC0 is the SC reward used in the supervised training process using pseudo-parallel data, and SC1 is used in the process of generating pseudo-parallel data.

TASKS	MODEL	SENTENCE	BLEU	BLEURT	COMET	ACC
Informal → Formal	Source	So if you're set on that, that's the way to go!!			-	
	M0	so if you're set on that, that's the way to go!!	0.417	0.175	0.568	0.000
	M1.1	so if you want to do this, this is the way to go!	0.301	0.204	0.354	0.003
	M1.2	If you want to do this, this is the way to go.	0.416	0.339	0.433	0.855
	M2.1	So if you're set on that, that is the way to go.	0.763	0.525	0.689	0.179
	M2.2	So, if you are set on that, then that is the way to go.	<b>0.884</b>	0.456	0.722	<b>0.880</b>
	M3.1	So if you are set on that, that is the way to go.	0.541	<b>0.941</b>	<b>0.734</b>	0.617
	M3.2	If you're on board, that's the way to go.	0.352	0.200	0.311	0.552
Positive → Negative	Source	the staff are all super friendly and on top of there jobs.			-	
	M0	the staff are all super friendly and on top of there jobs.	0.163	-0.561	-0.169	0.000
	M1.1	all the staff are very friendly and they're doing their jobs well.	0.107	-0.571	-0.301	0.003
	M1.2	the staff are all super friendly and on top of each same jobs.	0.149	-0.662	-0.507	0.000
	M1.3	the staff are all super unfriendly and on top of there jobs.	0.151	-0.239	0.095	<b>1.000</b>
	M2.1	the staff are all super rude and on top of there jobs.	0.151	-0.513	0.048	<b>1.000</b>
	M2.2	the staff are all super rude and on top of there jobs.	0.151	-0.513	0.048	<b>1.000</b>
	M2.3	the staff are all super rude and on top of there jobs.	0.151	-0.513	0.048	<b>1.000</b>
	M3.1	the staff are not super friendly or on top of there jobs.	<b>0.320</b>	<b>0.322</b>	<b>0.621</b>	<b>1.000</b>
M3.2	the staff are so friendly and they're doing their jobs.	0.148	-0.663	-0.326	0.001	

Table 4: Example outputs for the different steps of the pipeline and their corresponding evaluation results. Note that ACC represents style confidence here.

sentences. We use the style classifiers trained for selecting paraphrase data (Section 4.1). The classifiers have an accuracy of 92.6% and 98.1% on the test sets of F&R and YELP, respectively.

To assess **content preservation**, we follow previous work and calculate BLEU<sup>6</sup> between the generated sentence and the human reference(s). Additionally, we compute BLEURT and COMET<sup>7</sup>. As the human references for YELP are released from different researchers and appear to differ quite a lot in nature (see Appendix A.6 for examples), we provide two evaluation results: one using the first human reference only (Table 3), and the other using all four (Appendix A.2).

As **overall score**, for a direct comparison to pre-

vious work (Luo et al., 2019; Zhou et al., 2020; Lai et al., 2021) we compute the harmonic mean (HM) of style accuracy and BLEU.

## 5.2 Results

Table 3 reports results for each step.<sup>8</sup>

Results of Step 1 show that using paraphrase data benefits more formality transfer than polarity swap, confirming the latter is much less of a rewriting task than the former. Filtering paraphrases to a subset closer to the task (M1.2) substantially helps formality and yields some improvement in polarity. WordNet-derived synthetic pairs (M1.3) are definitely a better strategy for polarity.<sup>9</sup>

The first block of Step 2 confirms that further

<sup>6</sup>We use `multi-bleu.perl` with default settings.

<sup>7</sup>COMET is designed to also take input sentences into account, but our evaluations including them yielded lower correlations with human judgements. This might be because in COMET training input and output are different languages.

<sup>8</sup>Results for more models per step are in Appendix A.1.

<sup>9</sup>The WordNet-based strategy could in principle be used on its own to solve the polarity swap task with no learning involved, but results prove it insufficient: BLEURT: -0.475; COMET: -0.221; BLEU: 0.296; ACC: 0.206; HM: 0.243.

GYAFC (FORMALITY TRANSFER)						YELP (POLARITY SWAP)					
MODEL	BLEURT	COMET	BLEU	ACC	HM	MODEL	BLEURT	COMET	BLEU	ACC	HM
Input Copy	-0.114	0.272	0.474	0.120	0.192	Input Copy	-0.383	-0.139	0.312	0.019	0.036
UnsuperMT (Zhang et al., 2018)	-0.665	-0.446	0.327	0.670	0.439	Style-Transformer (Dai et al., 2019)	-0.469	-0.269	0.282	0.857	0.424
DualRL (Luo et al., 2019)	-0.589	-0.451	0.404	0.654	0.499	DualRL (Luo et al., 2019)	-0.385	-0.202	0.278	0.894	0.424
StyIns (Yi et al., 2020)	-0.395	-0.112	0.458	0.761	0.573	StyIns (Yi et al., 2020)	-0.576	-0.390	0.250	<b>0.924</b>	0.394
Zhou’s (Zhou et al., 2020)	-0.454	-0.203	0.447	0.799	0.573	Zhou’s (Zhou et al., 2020)	-0.270	-0.051	0.302	0.865	0.448
*TGLS (Li et al. (2020a); 0 → 1)	-	-	0.603	-	-	DGST (Li et al., 2020b)	-0.421	-0.240	0.268	0.781	0.399
Ours (M2.2; lowercase)	<b>0.009</b>	<b>0.328</b>	<b>0.563</b>	<b>0.866</b>	<b>0.682</b>	Ours (M3.1)	<b>-0.183</b>	<b>0.046</b>	<b>0.316</b>	0.887	<b>0.466</b>
Ours (M2.2; lowercase; 0 → 1)	-	-	<b>0.741</b>	-	-	-	-	-	-	-	-

Table 5: Comparison with other systems. Notes: (i) we lowercase the GYAFC texts for a fairer comparison to previous works, as they do so; (ii) if the output of previous work is available, we re-calculate the scores using our metrics. Otherwise we take the scores from the paper and mark this with a (\*); (iii) we report our results on informal-to-formal (0 → 1) alone to compare with Li et al. (2020a), who only transfer in this direction.

pre-training significantly improves performance on formality transfer (compare M2.2 with M2.1). This results in the best model for formality transfer. For polarity, instead, we see improvement from further pre-training only when using WordNet-based synthetic pairs (compare M2.2 with M2.3). Overall, in Step 2 we see that combining SC rewards and content-related rewards results in the best balance regarding content preservation and style strength.

In Step 3, we see that the model trained with high-quality synthetic pairs (M3.1) achieves the best overall performance on polarity swap. For comparison, we use the subset of paraphrase data as training pairs in place of the generated pairs, and see that performance is lower (M3.2).

Table 4 shows example outputs of each step and their evaluation results.<sup>10</sup> It is interesting to see the impact of paraphrase-based pre-training: for formality, in M1.1 and M1.2, the phrase “if you want to do this” is used in place of “if you’re set on that”. This rewriting ability can also be observed on the polarity swap (“on top of there jobs” → “they’re doing their jobs well”; note also that using paraphrases seems to prompt better writing: “there” → “their”, M1.1/M3.2, though this is not consistent throughout the models). For formality, the quality of the output gradually improves in Step 2, with M2.2 achieving the best performance on BLEU and style confidence (M2.2); the model trained with high-quality synthetic pairs (M3.1) has the highest BLEURT and COMET. In M3.2, trained on paraphrase pairs, we find nice variability again (“if you’re on board”). For polarity, M1.3 (using WordNet-based synthetic pairs), swaps a polarity word with its antonym (“friendly” → “unfriendly”). In Step 2, the models are indeed changing the polarity of the sentence; finally, the model trained with high-quality pairs (M3.1) nicely changes “and”

into “or” to get the right semantics (though it loses the correct form “their”) and is scored best. Further exploration of combining generic and task-specific rewriting appears very promising for these tasks.

As an additional curiosity-driven qualitative assessment of the behaviour of our models, we probed the polarity swap models with *neutral* sentences.<sup>11</sup> As a first example, we use “the earth revolves around the sun.” as the source sentence, and observe that the models in both transfer directions generate the same sentences as the input. With as input the neutral sentence “there is a grocery store near my house.”, the model which transforms negative sentences into positive ones generates “there is a great grocery store near my house.” while into the other direction it generates “there is no grocery store near my house.” It is worth mentioning that all the training data comes from business reviews on YELP, and the first example is clearly outside that domain. For the second example, closer to the domain of YELP, the transformation proposed by the model is rather reasonable in terms of obtaining a positive (“great grocery store”) or negative (“no grocery store”) output. It is left to future research to investigate what it should mean to transform a neutral sentence into a positive/negative one, and how such a test can help to better understand the models’ behaviour and the task itself.

**Comparison to other systems** To put our results in perspective, we compare our best system (M2.2 for formality and M3.1 for polarity in Table 3) against the most recent and best performing unpaired systems. For formality: UnsuperMT (Zhang et al., 2018); DualRL (Luo et al., 2019); StyIns (Yi et al., 2020); Zhou’s (Zhou et al., 2020); TGLS (Li et al., 2020a). For polarity: Style-Transformer (Dai

<sup>10</sup>References and more examples are in Appendix A.3.

<sup>11</sup>This was a suggestion of a reviewer, and we found indeed that this perspective could provide helpful insights in the models’ behaviour to be further studied in future work.



TASKS	N	BLEURT	COMET	BLEU
		COMET	BLEU	BLEURT
Formality Transfer	21	0.980 ( $p < 0.01$ )	0.775 ( $p < 0.01$ )	0.761 ( $p < 0.01$ )
Polarity Swap	21	0.968 ( $p < 0.01$ )	0.671 ( $p < 0.01$ )	0.479 ( $p = 0.03$ )

Table 6: Pearson correlation between evaluation metrics for content preservation over  $N$  systems.

et al., 2019); DualRL (Luo et al., 2019); StyIns (Yi et al., 2020); Zhou’s (Zhou et al., 2020); DGST (Li et al., 2020b).<sup>12</sup> We also add a simple baseline that just copies the input as output.

As visible in Table 5, our models achieve the best overall performance on both tasks. For formality transfer, this is true in all evaluation metrics. For polarity swap, StyIns has the highest style accuracy, while our model is better on all other metrics.<sup>13</sup>

### 5.3 Reflections on Tasks and Evaluation

The strategy of making the model learn the basic task of “rewriting” in a first stage clearly benefits more formality transfer than polarity swap. This is not surprising, since the latter is not simply “rewriting a sentence in a different style”; rather, the task involves changing the meaning of a sentence to obtain its opposite polarity, and thus, broadly put, its meaning. The fact that polarity swap cannot be regarded as a “style change” task is also evident from evaluation. Rather than only using BLEU, we suggested to also use BLEURT and COMET, and this provides us with additional evidence. Specifically, from Table 6 we observe that BLEU has a high correlation with BLEURT/COMET for formality transfer but not for polarity swap.

To glean further insights into this difference, we leverage human judgments released by Li et al. (2018) for YELP and see how they correlate with the used metrics. We calculate system-level Pearson correlation between the automatic evaluations and human judgment.

Results show that while COMET and BLEURT highly correlate with human judgments, BLEU does so to a lesser extent, suggesting this might be a less strong measure to assess the goodness of polarity swap.<sup>14</sup> Intuitively, if a system does not change the polarity it may still have a high  $n$ -gram overlap (high BLEU) while new-generation met-

<sup>12</sup>See Section 2 for details on these models.

<sup>13</sup>A sample comparison of outputs is in Appendix A.4.

<sup>14</sup>Pearson’s  $r = .922$  for BLEURT,  $r = .941$  for COMET, and  $r = .901$  for BLEU. All  $p < .001$ ,  $N = 7$ .

rics do not have this problem. For formality this limitation of BLEU is not much of an issue, since meaning is not altered. Nevertheless, we suggest that the evaluation of style transfer and related tasks should use learned metrics whenever possible.

## 6 Conclusions

We proposed an unpaired approach that adopts a multi-step procedure based on the general-purpose pre-trained seq2seq model BART.

Achieving state-of-the-art results on the two most popular “style transfer” tasks, we have shown the benefit of further pre-training using data derived from generic resources as well as the advantage of back-translation, paired with rewards, especially towards content preservation. We have also seen how leveraging paraphrases can enhance both variability and naturalness in the generated text.

Through experimental settings as well as the introduction of BLEURT and COMET as metrics, we have also highlighted how the two tasks we addressed differ, and should probably not be conflated into a single “style transfer” label. Indeed, we show that they benefit from partially different modelling, and react differently to evaluation metrics, both key aspects to improve future modelling of these tasks.

## Acknowledgments

This work was partly funded by the China Scholarship Council (CSC). The anonymous EMNLP reviewers provided us with useful comments which contributed to improving this paper and its presentation, so we’re grateful to them. We would also like to thank the Center for Information Technology of the University of Groningen for their support and for providing access to the Peregrine high performance computing cluster.

## Ethics Statement

All work that automatically generates and/or alters natural text could unfortunately be used maliciously. While we cannot fully prevent such uses once our models are made public, we do hope that writing about risks explicitly and also raising awareness of this possibility in the general public are ways to contain the effects of potential harmful uses. We are open to any discussion and suggestions to minimise such risks.

## References

- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. volume 10.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). *arXiv preprint, arXiv: 1409.0473*.
- Kunal Chawla and Diyi Yang. 2020. Semi-supervised formality style transfer using language model discriminator and mutual information maximization. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2340–2354.
- Ning Dai, Jianze Liang, Xipeng Qiu, and Xuanjing Huang. 2019. [Style transformer: Unpaired text style transfer without disentangled latent representation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5997–6007, Florence, Italy. Association for Computational Linguistics.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2018. Style transfer in text: Exploration and evaluation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 663–670.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. [Iterative back-translation for neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 18–24, Melbourne, Australia. Association for Computational Linguistics.
- J. Edward Hu, Abhinav Singh, Nils Holzenberger, Matt Post, and Benjamin Van Durme. 2019. Large-scale, diverse, paraphrastic bitexts via sampling and clustering. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 44–54.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. Disentangled representation learning for non-parallel text style transfer. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 424–434.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). *arXiv preprint, arXiv: 1412.6980*.
- 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.
- Guillaume Lample, Sandeep Subramanian, Eric Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2019. Multiple-attribute text rewriting. In *International Conference on Learning Representations*.
- Joosung Lee. 2020. [Stable style transformer: Delete and generate approach with encoder-decoder for text style transfer](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 195–204, Dublin, Ireland. 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.
- Jingjing Li, Zichao Li, Lili Mou, Xin Jiang, Michael R. Lyu, and Irwin King. 2020a. Unsupervised text generation by learning from search. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- 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.
- Xiao Li, Guanyi Chen, Chenghua Lin, and Ruizhe Li. 2020b. DGST: a dual-generator network for text style transfer. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7131–7136.
- Xianggen Liu, Lili Mou, Fandong Meng, Hao Zhou, Jie Zhou, and Sen Song. 2020. Unsupervised paraphrasing by simulated annealing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 302–312.
- 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. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 5116–5122.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems*, volume 26, page 3111–3119.
- George A. Miller. 1995. [Wordnet: A lexical database for english](#). *Commun. ACM*, 38(11):39–41.
- 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.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. 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.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Sudha Rao and Joel Tetreault. 2018. [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*, pages 129–140.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Abhilasha Sancheti, Kundan Krishna, Balaji Vasanth Srinivasan, and Anandhavelu Natarajan. 2020. [Reinforced rewards framework for text style transfer](#). In *Advances in Information Retrieval*, pages 545–560.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. [Style transfer from non-parallel text by cross-alignment](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6833–6844.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Machine Learning*, 8:229–256.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. 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.
- Chen Wu, Xuancheng Ren, Fuli Luo, and Xu Sun. 2019. [A hierarchical reinforced sequence operation method for unsupervised text style transfer](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4873–4883, Florence, Italy. Association for Computational Linguistics.
- Jingjing Xu, Xu Sun, Qi Zeng, Xiaodong Zhang, Xuancheng Ren, Houfeng Wang, and Wenjie Li. 2018. [Unpaired sentiment-to-sentiment translation: A cycled reinforcement learning approach](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 979–988, Melbourne, Australia. Association for Computational Linguistics.
- Xiaoyuan Yi, Zhenghao Liu, Wenhao Li, and Maosong Sun. 2020. [Text style transfer via learning style instance supported latent space](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3801–3807.
- Zhirui Zhang, Shuo Ren, Shujie Liu, Jianyong Wang, Peng Chen, Mu Li, Ming Zhou, and Enhong Chen. 2018. [Style transfer as unsupervised machine translation](#). *arXiv preprint, arXiv: 1808.07894*.
- Chulun Zhou, Liangyu Chen, Jiachen Liu, Xinyan Xiao, Jinsong Su, Sheng Guo, and Hua Wu. 2020. [Exploring contextual word-level style relevance for unsupervised style transfer](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7135–7144.

## A Appendices:

This appendices include: 1) detailed results for the different steps of the pipeline (A.1); 2) detailed evaluation results of using four human references on YELP (A.2); 3) example outputs for the different steps of the pipeline (A.3); 4) example outputs for existing systems we compare to, and our best models (A.4); 5) sample examples for further pre-training (A.5); 6) sample examples of human reference on YELP (A.6).

### A.1 Detailed results for the different steps of the pipeline

MODEL	DATASET	GYAFC (FORMALITY TRANSFER)					YELP (POLARITY SWAP)				
		BLEURT	COMET	BLEU	ACC	HM	BLEURT	COMET	BLEU	ACC	HM
Original BART		-0.116	0.242	0.414	0.333	0.369	-0.388	-0.146	0.309	0.022	0.041
STEP 1: Further pre-training											
Further pre-trained BART using whole dataset		0.012	0.209	0.420	0.357	0.386	-0.412	-0.282	0.179	0.040	0.065
Further pre-trained BART using subset		0.011	0.225	0.441	0.693	0.539	-0.347	-0.178	0.247	0.166	0.199
Further pre-trained BART using synthetic data		-	-	-	-	-	-0.321	-0.074	0.326	0.189	0.239
STEP 2: IBT + Rewards											
IBT (original BART)		-0.010	0.292	0.507	0.836	0.631	-0.229	-0.017	0.298	0.826	0.438
IBT (Further pre-trained BART using whole dataset)		<b>0.048</b>	0.319	0.550	0.907	0.685	-0.192	-0.041	0.252	0.854	0.389
IBT (Further pre-trained BART using subset)		0.041	0.318	0.553	<b>0.932</b>	<b>0.694</b>	-0.176	0.026	0.295	0.853	0.438
IBT (Further pre-trained BART using synthetic data)		-	-	-	-	-	-0.246	-0.035	0.302	0.884	0.450
IBT + SC0 + SC1 + BLEU		0.033	0.313	0.552	0.929	0.693	-0.187	0.001	0.285	0.860	0.428
IBT + SC0 + SC1 + BLEURT		0.041	0.320	0.551	0.925	0.691	<b>-0.149</b>	0.031	0.295	0.784	0.429
IBT + SC1 + BLEU + BLEURT		0.024	<b>0.321</b>	0.544	0.928	0.686	-0.195	-0.016	0.286	0.881	0.432
IBT + SC0 + BLEU + BLEURT		0.039	0.318	0.555	0.873	0.679	-0.176	0.039	<b>0.331</b>	0.500	0.398
IBT + SC0 + SC1		0.027	0.314	0.550	<b>0.932</b>	0.692	-0.208	-0.028	0.279	0.859	0.421
IBT + BLEU + BLEURT		0.036	0.318	0.552	0.857	0.671	-0.204	0.017	<b>0.331</b>	0.413	0.367
IBT without reward		0.032	0.319	0.551	0.849	0.668	-0.181	0.037	<b>0.331</b>	0.489	0.395
STEP 3: Offline training (Model used: original BART + Rewards)											
Trained with high-quality pairs		0.030	<b>0.321</b>	<b>0.560</b>	0.904	0.692	-0.183	<b>0.046</b>	0.316	<b>0.887</b>	<b>0.466</b>
Trained with subset of paraphrase data		0.012	0.229	0.455	0.783	0.576	-0.338	-0.221	0.215	0.457	0.292

Table A.1: Detailed results for the different steps of the pipeline. Note that (i) SC0 represents the SC reward is used in the supervised training process using pseudo-parallel data, and SC1 is in the process of generating pseudo-parallel data.

### A.2 Detailed evaluation results of using four human references on YELP

MODEL	BLEURT	COMET	BLEU	ACC	HM
Input Copy	-0.337	-0.033	<b>0.640</b>	0.019	0.037
DualRL (Luo et al., 2019)	-0.281	-0.080	0.550	0.894	0.681
Style-Transformer (Dai et al., 2019)	-0.390	-0.158	0.553	0.857	0.672
DGST (Li et al., 2020b)	-0.337	-0.131	0.520	0.781	0.624
StyIns (Yi et al., 2020)	-0.487	-0.280	0.489	<b>0.924</b>	0.640
Zhou’s (Zhou et al., 2020)	-0.162	0.090	0.608	0.865	0.714
Ours	<b>-0.053</b>	<b>0.192</b>	0.610	0.887	<b>0.723</b>

Table A.2: Automatic evaluation results using four human references on YELP.

### A.3 Example outputs for the different steps of the pipeline

MODEL	INFORMAL → FORMAL	BLEU	BLEURT	COMET	Style Confidence
Source	So if you're set on that, that's the way to go!!			-	
Reference 1	If you are set on that, that is the way to go.			-	
Reference 2	If that is your decision, then that is what you should do.			-	
Reference 3	So that is the way to go if you are set on that.			-	
Reference 4	If you are set on that, then that is the way to go.			-	
M0	so if you're set on that, that's the way to go!!	0.417	0.175	0.568	0.000
M1.1	so if you want to do this, this is the way to go!	0.301	0.204	0.354	0.003
M1.2	If you want to do this, this is the way to go.	0.416	0.339	0.433	0.855
M2.1	So if you're set on that, that is the way to go.	0.763	0.525	0.689	0.179
M2.2	So, if you are set on that, then that is the way to go.	<b>0.884</b>	0.456	0.722	<b>0.880</b>
M3.1	So if you are set on that, that is the way to go.	0.541	<b>0.941</b>	<b>0.734</b>	0.617
M3.2	If you're on board, that's the way to go.	0.352	0.200	0.311	0.552
Source	(saying sorry to him without committing a mistake is humiliation to ur self...AND DON'T EVER GET WEAK!!)			-	
Reference 1	Do not apologize if you have not made a mistake.			-	
Reference 2	Don't say sorry to him unless you've actually made a mistake.			-	
Reference 3	You will just humiliate yourself if you apologize to him without committing a mistake. Do not get weak!			-	
Reference 4	Avoid being weak by not saying sorry for something you did not do.			-	
M0	(saying sorry to him without committing a mistake is humiliation to ur self...and don't ever get weak!!)	0.100	-0.985	-0.452	0.000
M1.1	(saying sorry to him without apologizing for a mistake is humiliation to your self... and don't ever get weak	0.095	-0.697	-0.414	0.002
M1.2	I'm sorry, but to apologize without making a mistake is humiliation to my self, and don't ever get weak	0.031	<b>-0.648</b>	-0.608	0.694
M2.1	(Saying sorry to him without committing a mistake is humiliation to your self. Please don't ever get weak.	0.089	-0.661	<b>-0.247</b>	<b>0.980</b>
M2.2	(Saying sorry to him without committing a mistake is humiliation to your self and do not ever get weak.	0.099	-0.655	-0.377	0.758
M3.1	Saying sorry to him without committing a mistake is humiliation to your self and do not ever get weak.	0.099	-0.655	-0.377	0.758
M3.2	Saying sorry to him without making a mistake is humiliation to your self, and don't ever get weak!	<b>0.108</b>	-0.652	-0.400	0.083
MODEL	POSITIVE → NEGATIVE	BLEU	BLEURT	COMET	Style Confidence
Source	the staff are all super friendly and on top of there jobs.			-	
Reference	the staff are not friendly and not on top of their jobs.			-	
M0	the staff are all super friendly and on top of there jobs.	0.163	-0.561	-0.169	0.000
M1.1	all the staff are very friendly and they're doing their jobs well.	0.107	-0.571	-0.301	0.003
M1.2	the staff are all super friendly and on top of each same jobs.	0.149	-0.662	-0.507	0.000
M1.3	the staff are all super unfriendly and on top of there jobs.	0.151	-0.239	0.095	<b>1.000</b>
M2.1	the staff are all super rude and on top of there jobs.	0.151	-0.513	0.048	<b>1.000</b>
M2.2	the staff are all super rude and on top of there jobs.	0.151	-0.513	0.048	<b>1.000</b>
M2.3	the staff are all super rude and on top of there jobs.	0.151	-0.513	0.048	<b>1.000</b>
M3.1	the staff are not super friendly or on top of there jobs.	<b>0.320</b>	<b>0.322</b>	<b>0.621</b>	<b>1.000</b>
M3.2	the staff are so friendly and they're doing their jobs.	0.148	-0.663	-0.326	0.001
Source	very good brunch, was impressed with selection and quality.			-	
Reference	the brunch was bad, with little selection			-	
M0	very good brunch, was impressed with selection and quality.	0.028	-1.134	-0.400	0.000
M1.1	it was a very good brunch, i was impressed by the selection and the quality.	0.017	-0.714	-0.497	0.000
M1.2	very good brunch, was impressed with the selection and quality.	0.027	-1.101	-0.441	0.000
M1.3	very bad brunch, was impressed with selection and quality.	0.030	-1.100	-0.118	0.778
M2.1	very mediocre brunch, was disappointed with selection and quality.	0.028	<b>-0.367</b>	0.160	<b>1.000</b>
M2.2	very disappointing brunch, was disappointed with selection and service.	0.028	-0.532	-0.020	<b>1.000</b>
M2.3	very mediocre brunch, was disappointed with selection and quality.	0.028	<b>-0.367</b>	-0.160	<b>1.000</b>
M3.1	very bad brunch, was disappointed with selection and quality.	<b>0.030</b>	-0.495	<b>0.236</b>	<b>1.000</b>
M3.2	it was a very good brunch, i was impressed with the food and the service.	0.017	-0.865	-0.649	0.000

Table A.3: Example outputs for the different steps of the pipeline.

### A.4 Example outputs for existing systems we compare to, and our best models

SYSTEMS	INFORMAL → FORMAL	NEGATIVE → POSITIVE
Input	i hardly everrr see him in school either usually I see hima t my brothers basketball games.	so, no treatment and no medication to help me deal with my condition.
DualRL	<b>recognize him</b> see him in school either I usually <b>see my brothers</b> .	so, great treatment and <b>great</b> help me deal with my condition.
StyIns	<b>I would not</b> see him in school either <b>because</b> I see <b>to profess</b> my brothers basketball games.	so, great <b>service</b> and <b>great location</b> to help me deal with my condition.
Zhou's	I hardly <b>everrr</b> see him in school either. see I <b>""</b> my brothers games.	so, great treatment and <b>no</b> medication to help me deal with my condition.
Ours	I hardly ever see him in school, but usually I see him at my brothers basketball games.	so great treatment and great medication to help me deal with my condition.
Human	I hardly ever see him in school, usually I see him when I go to my brother 's basketball games.	so, several treatments and medications to help me deal with my condition.

Table A.4: Example outputs for existing systems we compare to, and our best models. Improperly generated words/phrases are in red. We can observe that: 1) there are still some informal/negative expressions in the generated sentences of previous systems like Zhou et al. (2020)'s. 2) Some systems introduce noise in the generated sentences and fail in preserving content like DualRAL (Luo et al., 2019)'s and StyIns (Yi et al., 2020)'s. On the contrary, our proposed approach is better at changing input sentences into the target style while preserving most style-independent parts. Furthermore, generated sentences of our system are more fluent than previous systems.

## A.5 Sample examples for further pre-training

RESOURCE	TASK	INFORMAL/NEGATIVE	FORMAL/*POSITIVE
Paraphrase	Formality	Now I... I like the interactive side of this job. Y'all got five minutes to finish your smoke. Yeah, well, there's a bridge right here. If they go.....they leave the source of power behind. Ain't many of us can face it out there sober all the time.	I like the interactivity on our work. All of you have five minutes to finish your long smoke. Here's one bridge, ten kilometers from there. If they leave, they'll leave their source of strength. Not too many of us could face this outside as sober.
	Polarity	he makes me feel wired. there's a room on the empty floor. it's only part of work, you know – routine clearance. it was the only thing i liked to buy here. wherever it went, it was followed by an admiring crowd of small lassans.	i... it gives me a funny feeling. there's plenty of free space on the next floor. great. yeah, it's just part of the job, you know... a routine imprint. that is the one thing i actually enjoy buying at this store. wherever he moved, he was followed by an astonished mob of little lassans.
WordNet	Polarity	some of the worst pizza i've ever had. also the inside is dirty as heck. the guy never really even apologized for the mistake. wake up or you are going to lose your business. absolutely the worst care in all my experience with vets!	some of the best pizza i've ever had. also the inside is clean as heck. the guy ever really even apologized for the mistake. wake up or you are going to find your business. absolutely the best care in all my experience with vets!

Table A.5: Sample examples for further pre-training. \* indicates that the sentences are synthetic.

## A.6 Sample examples of human reference on YELP

	NEGATIVE → POSITIVE	POSITIVE → NEGATIVE
Source	ever since joes has changed hands it's just gotten worse and worse.	it's small yet they make you feel right at home.
Reference 1	ever since joes has changed hands it's gotten better and better.	it's small yet they make you feel like a stranger.
Reference 2	ever since joes has changed hands it's gotten better and better.	it's small and make you feel as small office cabin
Reference 3	since joe changed hands, it has become a better place.	it's small and not friendly at all
Reference 4	ever since joes has changed hands it is getting better	it's small and they make you feel like a stranger
Source	there is definitely not enough room in that part of the venue.	i will be going back and enjoying this great place!
Reference 1	there is so much room in that part of the venue	i won't be going back and suffering at this terrible place!
Reference 2	there is definitely enough room in that part of the venue.	i will not be going back to this terrible place!
Reference 3	there is enough space in that oart of the venue.	i will never come back to this bad place!
Reference 4	there is many room on that venue	i will not be returning to this place and it was unenjoyable
Source	so basically tasted watered down.	the drinks were affordable and a good pour.
Reference 1	it didn't taste watered down at all.	the drinks were expensive and half full.
Reference 2	so it's fine because it is not watered down.	the drinks were expensive and a less pour
Reference 3	so basically not tasted watered down.	the drinks were very expensive and a less pour
Reference 4	so basically did not taste watered down.	the drinks were not affordable and a not good pour.
Source	she said she'd be back and disappeared for a few minutes.	my husband got a ruben sandwich, he loved it.
Reference 1	she said she'd be back, and didn't disappear at all.	my husband got a reuben sandwich, he hated it.
Reference 2	she said she'd be back and enjoy herself	my husband got a ruben sandwich, he hate it very much.
Reference 3	she said she'd be back and will not disappeared	my husband got a ruben sandwich, he hated it.
Reference 4	she said she'd be back and have a good time	my husband got a ruben sandwich , he did not love it
Source	i can't believe how inconsiderate this pharmacy is.	i signed up for their email and got a coupon.
Reference 1	this pharmacy is really considerate.	i signed up for their email and got spam.
Reference 2	i can not imagine how considerate this pharmacy is.	i signed up for their email and got nothing.
Reference 3	the pharmacy was so considerate of me	i signed up for their email and didnt even get offered a deal or anything.
Reference 4	i can not believe how considerate this pharmacy is	i wrote an email and did not obtube anything

Table A.6: Sample examples of human reference on YELP. The first human reference is provided by [Li et al. \(2018\)](#), and the 3 additional references are released by [Luo et al. \(2019\)](#).