

Contrastive Representation Learning for Exemplar-Guided Paraphrase Generation *

Haoran Yang¹, Wai Lam¹, Piji Li²

¹The Chinese University of Hong Kong

²Tencent AI Lab

{hryang, wlam}@se.cuhk.edu.hk

lipiji.pz@gmail.com

Abstract

Exemplar-Guided Paraphrase Generation (EGPG) aims to generate a target sentence which conforms to the style of the given exemplar while encapsulating the content information of the source sentence. In this paper, we propose a new method with the goal of learning a better representation of the style and the content. This method is mainly motivated by the recent success of contrastive learning which has demonstrated its power in *unsupervised* feature extraction tasks. The idea is to design two contrastive losses with respect to the content and the style by considering two problem characteristics during training. One characteristic is that the target sentence shares the same content with the source sentence, and the second characteristic is that the target sentence shares the same style with the exemplar. These two contrastive losses are incorporated into the general encoder-decoder paradigm. Experiments on two datasets, namely QQP-Pos and ParaNMT, demonstrate the effectiveness of our proposed contrastive losses. The code is available at https://github.com/LHRYANG/CRL_EGPG.

1 Introduction

Paraphrase generation (Gupta et al., 2017; Li et al., 2019), aiming to generate a sentence with the same semantic meaning of the source sentence, has achieved a great success in recent years. To obtain a paraphrase sentence with a particular style, Exemplar-Guided Paraphrase Generation (EGPG) (Chen et al., 2019) has attracted considerable attention. Different from other controllable text generation tasks whose constraints are taken from a finite set, e.g., binary sentiment or political slant (Yang et al., 2018; Prabhumoye et al., 2018), multiple personas (Kang et al., 2019), over

The work described in this paper is substantially supported by a grant from the Research Grant Council of the Hong Kong Special Administrative Region, China (Project Code: 14200620).

source (X)	<i>what is the easiest way to get followers on quora ?</i>
exemplar (Z)	<i>how do i avoid plagiarism in my article ?</i>
target (Y)	<i>how do i get more followers for my quora ?</i>
retrieved (Y')	<i>what are the better ways to ask questions on quora ?</i>

Table 1: An example of EGPG

which a classifier can be trained to guide the disentanglement process, the constraints of EGPG are exemplar sentences that can be arbitrarily provided, making it more challenging to learn a good representation for the style and the content. For example, as shown in Table 1, when using the content embedding of X to retrieve a sentence with the most similar content in the target sentences list, we observe that the ordinary model, which is described in Section 4.5, can often match the sentence Y' whose content differs from X instead of the correct target sentence Y . This reveals that the content encoder cannot encode the content information of a sentence appropriately which can result in inconsistent content of the generated sentence. The same problem also exists in the style encoder.

To learn a better content and style representations, we explore the incorporation of *contrastive learning* in EGPG to design an end-to-end encoder-decoder paradigm with multiple losses. Contrastive learning originates from computer vision area (Chen et al., 2020a; Khosla et al., 2020) and now, it also shows its powerfulness in natural language processing area. For instance, Iyer et al. (2020) employ contrastive learning to improve the quality of discourse-level sentence representations. In our proposed model, besides the basic encoder-decoder generation task, a content contrastive loss is designed to force the content encoder to distinguish features of the same content from features of different content. Similarly, a style contrastive loss is also employed to obtain a similar distinguishing effect for the style features. Experimental results on two benchmark datasets, namely QQP-Pos and ParaNMT, show that superior performance can be

achieved with the help of the contrastive losses.

2 Related Work

Paraphrase Generation Researches on paraphrase generation has been for a long time. Traditional methods solve this problem mainly through statistical machine translation (Quirk et al., 2004) or rule-based word substitution (Wubben et al., 2010). In the era of deep learning, approaches based on the encoder-decoder framework have emerged in large numbers (Prakash et al., 2016; Chen et al., 2020b). In addition to basic seq2seq model, Li et al. (2018) add a pair-wise discriminator to judge whether the input sentence and generated sentence are paraphrases of each other, with the help of reinforcement learning. To generate diverse paraphrases, i.e., one to many mapping, Gupta et al. (2017) combine the power of RNN-based sequence-to-sequence model and the variational autoencoder. At decoding time, a noise sampled from the Gaussian distribution are appended to input to generate a diverse output. Qian et al. (2019) propose a approach which use multiple generators to generate diverse paraphrases without sacrificing quality.

Exemplar-Guided Paraphrase Generation Making the generated paraphrases satisfy the style of an exemplar sentence is recently a hot research topic. EGPH is similar to other controlled text generation tasks whose constraints are sentiment (Yang et al., 2018; Xu et al., 2021), gender (Prabhumoye et al., 2018), topics (Wang et al., 2021). These tasks are highly related to Disentangled Representation Learning (DRL) which maps different aspects of the input data to independent low-dimensional spaces (Cheng et al., 2020). Iyyer et al. (2018) and Kumar et al. (2020) directly utilise the parse tree information of the exemplar as the style information without separating style from sentences. Chen et al. (2019) propose a model which can directly extract style features from a modified target sentence. Goyal and Durrett (2020) provide a way to generate paraphrase which is a component rearrangement of the original input through manipulating the parse tree.

Contrastive Learning In the past few years, many unsupervised feature extraction algorithms have emerged, for instance, variational autoencoder (Kingma and Welling, 2014; Xu et al., 2020; Gao et al., 2019b,a), generalised language models (Brown et al., 2020; Devlin et al., 2019). All the

above methods obtain the feature of input by reconstructing the original input or predicting masked words and so on which do not take the relationships between the inputs into consideration. Therefore, contrastive learning, whose loss is designed to narrow down the distance between features of similar inputs and to enlarge the distance of dissimilar inputs, has been proposed and achieved a great success in both unsupervised (Chen et al., 2020a) and supervised (Khosla et al., 2020) image feature extraction. There are also some works trying to apply contrastive learning into natural language processing domain. For instance, Iyer et al. (2020) propose a pretraining method for sentence representation which employs contrastive learning to improve the quality of discourse-level representations. Giorgi et al. (2020) utilise it to pretrain the transformer and obtains state-of-the-art results on SentEval (Conneau and Kiela, 2018). All of the above successes spur us to test whether contrastive learning is helpful on EGPG.

3 Proposed Model

Given a source sentence X_i and an exemplar sentence Z_i , our goal is to generate a sentence Y_i that has the same style (syntax) with Z_i and retains the content (semantics) of X_i . As shown in Figure 1, we design the encoders E_s and E_c for style and content respectively. The decoder D generates the output. Our model is trained by optimizing three losses simultaneously: (1) generation loss; (2) content contrastive loss; (3) style contrastive loss.

Generation Task For X_i and Z_i , we firstly obtain their corresponding content features c_{X_i} and style features s_{Z_i} :

$$c_{X_i} = E_c(X_i) \quad (1)$$

$$s_{Z_i} = E_s(Z_i) \quad (2)$$

Then c_{X_i} and s_{Z_i} are concatenated and inputted into the decoder as the initial hidden state to generate a sequence of probabilities over vocabulary. At the step t , the predicted probability p_t of the t -th target word is obtained as follows

$$p_t = \text{softmax}(Wh_t) \quad (3)$$

$$h_t = \text{GRU}(h_{t-1}, e(y_{t-1})) \quad (4)$$

where h_0 is initialized as $[c_{X_i}, s_{Z_i}]$ and W is a parameter matrix. y_{t-1} is the word in the previous step $t-1$ and y_0 is the special symbol [SOS] which represents the start of the sentence. $e(y_{t-1})$ is the embedding of the word y_{t-1} .

Negative log-likelihood loss (NLL) is employed

as the basic optimization objective

$$\mathcal{L}_i^{nll} = -\frac{1}{|Y_i|} \sum_{t=1}^{|Y_i|} I(y_t)^T \log p_t \quad (5)$$

where $I(y_t)$ represents the one-hot encoding of the word y_t in the vocabulary.

Content Contrastive Learning (CCL) Considering that X_i and Y_i share the same content, their content features should be close with each other in the content feature space. Contrastive Learning which is designed to minimize the distance between positive pairs and maximize the distance between negative pairs can help model this relationship. Formally, during training, given a batch $\{(X_i, Y_i, Z_i)\}_{i=1}^n$ where n is the batch size, we firstly obtain the corresponding content features of X_i and Y_i , denoted by $\{(c_{X_i}, c_{Y_i})\}_{i=1}^n$. For c_{X_i} , the positive pair is (c_{X_i}, c_{Y_i}) and c_{X_i} with the other remaining features in this batch form $2n - 2$ negative pairs. For c_{Y_i} , the definition of positive/negative pairs is the same as c_{X_i} . Then the contrastive loss is employed, giving

$$\mathcal{L}_{X_i}^{ccl} = -\log \frac{\exp(c_{X_i} \cdot c_{Y_i} / \tau)}{\exp(\frac{c_{Y_i} \cdot c_{X_i}}{\tau}) + \sum_{\substack{j \neq i \\ T \in \{X, Y\}}} \exp(\frac{c_{X_i} \cdot c_{T_j}}{\tau})} \quad (6)$$

$$\mathcal{L}_{Y_i}^{ccl} = -\log \frac{\exp(c_{Y_i} \cdot c_{X_i} / \tau)}{\exp(\frac{c_{Y_i} \cdot c_{X_i}}{\tau}) + \sum_{\substack{j \neq i \\ T \in \{X, Y\}}} \exp(\frac{c_{Y_i} \cdot c_{T_j}}{\tau})} \quad (7)$$

$$\mathcal{L}^{ccl} = \sum_{i=1}^n (\mathcal{L}_{X_i}^{ccl} + \mathcal{L}_{Y_i}^{ccl}) \quad (8)$$

where \cdot represents the dot product between two vectors and τ denotes a temperature parameter.

Style Contrastive Learning (SCL) aims to help E_s learn a better style representation by considering that Z_i and Y_i share the same style. Similar to CCL, we firstly obtain the style features $\{(s_{Z_i}, s_{Y_i})\}_{i=1}^n$ and then apply the contrastive loss to these features

$$\mathcal{L}_{Y_i}^{scl} = -\log \frac{\exp(s_{Y_i} \cdot s_{Z_i} / \tau)}{\exp(\frac{s_{Y_i} \cdot s_{Z_i}}{\tau}) + \sum_{\substack{j \neq i \\ T \in \{Z, Y\}}} \exp(\frac{s_{Y_i} \cdot s_{T_j}}{\tau})} \quad (9)$$

$$\mathcal{L}_{Z_i}^{scl} = -\log \frac{\exp(s_{Z_i} \cdot s_{Y_i} / \tau)}{\exp(\frac{s_{Z_i} \cdot s_{Y_i}}{\tau}) + \sum_{\substack{j \neq i \\ T \in \{Z, Y\}}} \exp(\frac{s_{Z_i} \cdot s_{T_j}}{\tau})} \quad (10)$$




Figure 1: An Overview of Our Model

$$\mathcal{L}^{scl} = \sum_{i=1}^n (\mathcal{L}_{Y_i}^{scl} + \mathcal{L}_{Z_i}^{scl}) \quad (11)$$

As a result, the total loss for a batch is as follows

$$\mathcal{L} = \sum_{i=1}^n \mathcal{L}_i^{nll} + \lambda_1 \mathcal{L}^{ccl} + \lambda_2 \mathcal{L}^{scl} \quad (12)$$

4 Experiments

4.1 Datasets

We conduct experiments on two benchmark datasets, namely ParaNMT (Chen et al., 2019) and QQP-Pos (Kumar et al., 2020). ParaNMT consists of about 500k training, 800 testing and 500 validation sentence pairs which are automatically generated through backtranslation of the original English sentences. QQP-Pos consists of 130k training, 3k testing and 3k validation quora question pairs which are more formal than ParaNMT. The split size is the same as previous works to have a fair comparison. Since the exemplar sentences are not provided in both datasets, we adopt a method similar with Kumar et al. (2020) to search an exemplar Z_i for each source-target pair (X_i, Y_i) based on the POS tag sequence¹ similarity (refer to Appendix A).

4.2 Baselines & Metrics

We compare our model with (1) SCPN (Iyyer et al., 2018) which employs a parse generator to output the full linearized parse tree as the style by inputting a parse template; (2) SGCP (Kumar et al., 2020) which extracts the style information directly from the parse tree of the exemplar sentence; (3) CGEN (Chen et al., 2019), an approach based on variational inference (Kingma and Welling, 2014).

The evaluation metrics are BLEU (Papineni et al., 2002), METEOR (Lavie and Agarwal, 2007) and ROUGE (R) (Lin, 2004). We also conduct human evaluation to investigate the quality of the generated sentences. Moreover, we propose **Content**

¹We use NLTK for POS tagging.

Matching Accuracy (CMA) to gauge the quality of the generated embeddings for the content. CMA will be introduced in Section 4.5.

4.3 Implementation Details

Each sentence is trimmed with a maximum length 15. The word embedding is initialized with 300- d pretrained GloVe (Pennington et al., 2014). We use a BERT-based (Devlin et al., 2019) architecture for the style encoder E_s and the dimension of style features is 768. For content encoder E_c , we use GRU (Chung et al., 2014) with hidden state size 512. During training, the teacher forcing technique is applied with the rate 1.0. The balancing parameters λ_1 and λ_2 are both set to 0.1. The temperature parameter τ is set to 0.5. We train our model using Adam optimizer with the learning rate $1e-4$ and the training epochs are set to 30 and 45 for ParaNMT and QQP-Pos respectively.

4.4 Results

As summarized in Table 2, our model outperforms SCPN, SGCP and CGEN by a large margin on automatic evaluation metrics. We also conduct human evaluation to investigate the holistic quality of the generated sentences. For each dataset, we firstly choose two source sentences and then randomly select 25 exemplars for each source sentence to generate a total of 50 sentences. Table 5 shows the results of human assessment. It can be seen that our model obtains a higher score than SGCP and CGEN, which is consistent with the automatic evaluation results. These results are expected because SCPN and SGCP use a parse tree as the style which is lack of the lexical information and very unstable. Moreover, CEGN is VAE-based which is intrinsically harder to train (Bowman et al., 2016).

4.5 Ablation Study

We conduct ablation study with three variants, namely ours without SCL (Ours-w/o-SCL), ours without CCL (Ours-w/o-CCL), ours without both SCL and CCL (Ours-w/o-both). We show that our model is better than the three variants to demonstrate the effectiveness of the contrastive losses.

As presented in Table 2, we can see that Ours can achieve better results than the three variants on all automatic metrics for QQP-Pos. Particularly, Ours, Ours-w/o-CCL and Ours-w/o-SCL outperform Ours-w/o-both a lot, demonstrating the usefulness of the contrastive losses. For ParaNMT, Ours obtains the highest score on BLEU, ROUGE-2,

QQP-Pos					
Model	BLEU	R-1	R-2	R-L	METEOR
SCPN	15.6	40.6	20.5	44.6	19.6
SGCP	36.7	66.9	45.0	69.6	39.8
CGEN	34.9	62.6	42.7	65.4	37.4
Ours	45.8	71.0	52.8	73.3	45.8
Ours-w/o-CCL	43.1	70.0	50.6	72.3	43.5
Ours-w/o-SCL	42.7	69.7	49.9	71.8	43.6
Ours-w/o-both	40.8	68.4	48.4	70.8	41.6
ParaNMT					
SCPN	6.4	30.3	11.2	34.6	14.6
SGCP	15.3	46.6	21.8	49.7	25.9
CGEN	13.6	44.8	21.0	48.3	24.8
Ours	16.2	50.6	25.3	52.1	28.4
Ours-w/o-CCL	15.3	50.8	25.2	52.4	28.0
Ours-w/o-SCL	15.2	50.2	24.4	51.5	28.0
Ours-w/o-both	15.3	50.2	24.9	51.6	27.7

Table 2: Automatic Evaluation Results.

Model	Ours	Ours-w/o-CCL	Ours-w/o-SCL	Ours-w/o-both
QQP-Pos				
ED-E	2.49	2.42	2.56	2.57
ED-R	2.64	2.65	2.78	2.82
ParaNMT				
ED-E	4.36	4.47	4.49	4.49
ED-R	4.22	4.24	4.28	4.25

Table 3: Style Evaluation

MENTOR while it lags behind Ours-w/o-CCL on ROUGE-1 and ROUGE-L. This phenomena may be caused by the poor quality of the dataset. The human evaluation results are also listed in Table 5. Our model can generally generate fluent sentences on QQP-Pos. But the overall quality of sentences generated by all models on ParaNMT is unsatisfactory which shows that a high-quality dataset is necessary for training a good model.

We also provide the style evaluation of **ED-E** (edit distance between the POS tag sequence of the generated paraphrase and the exemplar) and **ED-R** (edit distance between the POS tag sequence of the generated paraphrase and the ground truth reference) in Table 3. We can see that models with style contrastive losses have smaller edit distance.

To directly assess the quality of the generated embeddings for the content we propose **Content Matching Accuracy**. To calculate CMA, firstly we input all the source and target sentences into E_c to get the content representations $\mathbf{A}, \mathbf{B} \in \mathbf{R}^{m \times k}$, where m is the size of the test dataset and k is the dimension of content feature vectors. Then, we calculate the similarity matrix $\mathbf{S} = \mathbf{AB}^T$. In ideal situation, each diagonal element $\mathbf{S}_{[i,i]}$ should be the largest value in row i since the content embedding \mathbf{A}_i of the i_{th} source sentence should have the greatest similarity with the content embedding \mathbf{B}_i of its corresponding target sentence. Therefore, the

SOURCE: how do i develop good project management skills ?	
EXEMPLAR	GENERATION
which is the best laptop model to buy within 30k ?	which is the best way to develop project management skills ?
how many cups of coffee should i consume in a day ?	from what skills can i start in a project management ?
which subjects are important to become a chartered accountant ?	what skills are necessary to develop a project management ?

Table 4: An Example Generated by Our Model

Model	Ours	SGCP	CGEN	Ours-w/o-CCL	Ours-w/o-SCL	Ours-w/o-both
QQP-Pos	3.71	3.12	2.97	3.59	3.45	3.33
ParaNMT	2.53	1.9	2.05	2.51	2.53	2.38

Table 5: Human evaluation results. Each sentence is given a score ranging from one to five to assess the holistic quality. We report the average value of two annotators. Higher score is better. The Spearman’s correlation coefficients of these two annotators are 0.707 for QQP-Pos and 0.37 for ParaNMT.

X	can you view a private facebook profile?
Z	how would you learn a new programming language?
Y	how do you view a private facebook profile?
Y'	can you see who visited your instagram last?
Ours	how do you see a private facebook profile?
Ours-w/o-SCL	how can you view a private facebook profile?
Ours-w/o-both	how can you view a private instagram profile?

Table 6: A failed example without CCL

content matching accuracy (CMA) is defined as:

$$CMA = \frac{\sum_{i=1}^m \mathbb{1}(\arg\max(\mathbf{S}_i) = i)}{m} \quad (13)$$

where $\mathbb{1}(s)$ equals 1 if s is true, otherwise 0. The results are illustrated in Figure 2. We notice that models with CCL can achieve higher accuracy than Ours-w/o-both. It signifies that the content encoder E_c is improved with the help of CCL. We provide a failed example in Table 6. Y' is the sentence retrieved given X under the model Ours-w/o-both. Y is the target sentence and it is also the sentence retrieved given X under the model Ours or Ours-w/o-SCL. We can see that the content generated by Ours-w/o-both is incorrect (changing facebook to instagram) and instagram exists in Y' . This illustrates that the poor-quality content embedding of X can cause the incorrect content of the generated sentence. In general, matching accuracy can also be calculated for the style. However, the exemplar selection process has a high probability of dropping the sentence with the most similar style of Y . Therefore, style matching accuracy is not provided here. Instead, we list some retrieved sentences based on the style embedding in Appendix B.

4.6 Case Study

Some examples generated by our model are shown in Table 4. It can be observed that our model can generally generate high-quality sentences which have similar style with the exemplar and retain the semantic meaning of the source sentence. More-

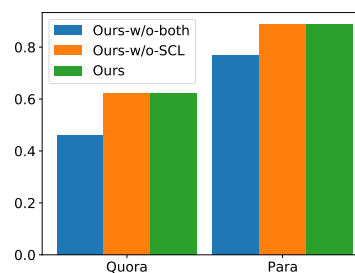


Figure 2: Content Matching Accuracy

over, our model does not directly copy the style words from the exemplar, but instead adopts the overall structure of the exemplar to generate sentences, for example, the second one. More examples are provided in Appendix C.

4.7 Explanations

We attempt to provide some possible explanations about why the model with these two contrastive losses can achieve better performance. The first reason is that adding additional losses on the output of encoders can alleviate gradient vanishing which is a serious issue when training the encoder-decoder model. The second reason is that the overfitting issue may be prevented, since the contrastive losses restrict the free adjustment of parameters in the model by forcing the encoder and decoder to focus on their own tasks, i.e., feature extraction and sentence generation.

5 Conclusion

We introduce the content contrastive loss and the style contrastive loss into EGPG to design a multi-losses scheme without requiring additional labeled data. This scheme can obtain better results compared with the baseline and ablative models, which demonstrates the effectiveness of contrastive learning for learning better representations. Moreover, the proposed framework is general and may benefit other similar NLP tasks.

References

- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. [Generating sentences from a continuous space](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *ArXiv*.
- Mingda Chen, Qingming Tang, Sam Wiseman, and Kevin Gimpel. 2019. Controllable paraphrase generation with a syntactic exemplar. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5972–5984.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020a. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*.
- Yen-Chun Chen, Zhe Gan, Yu Cheng, Jingzhou Liu, and Jingjing Liu. 2020b. Distilling knowledge learned in bert for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Pengyu Cheng, Martin Renqiang Min, Dinghan Shen, Christopher Malon, Yizhe Zhang, Yitong Li, and Lawrence Carin. 2020. Improving disentangled text representation learning with information-theoretic guidance. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7530–7541, Online.
- Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning*.
- Alexis Conneau and Douwe Kiela. 2018. [SentEval: An evaluation toolkit for universal sentence representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186.
- Jun Gao, Wei Bi, Xiaojiang Liu, Junhui Li, and Shuming Shi. 2019a. Generating multiple diverse responses for short-text conversation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6383–6390.
- Jun Gao, Wei Bi, Xiaojiang Liu, Junhui Li, Guodong Zhou, and Shuming Shi. 2019b. A discrete cvae for response generation on short-text conversation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1898–1908.
- John M Giorgi, Osvald Nitski, Gary D. Bader, and Bo Wang. 2020. Decludr: Deep contrastive learning for unsupervised textual representations. *ArXiv*, abs/2006.03659.
- Tanya Goyal and Greg Durrett. 2020. Neural syntactic preordering for controlled paraphrase generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 238–252.
- Ankush Gupta, Arvind Agarwal, Prawaan Singh, and Piyush Rai. 2017. A deep generative framework for paraphrase generation. In *Proceedings of the 32th AAAI Conference on Artificial Intelligence*.
- Dan Iter, Kelvin Guu, Larry Lansing, and Dan Jurafsky. 2020. Pretraining with contrastive sentence objectives improves discourse performance of language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial example generation with syntactically controlled paraphrase networks. In *16th Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1875–1885.
- Dongyeop Kang, Varun Gangal, and Eduard Hovy. 2019. (male, bachelor) and (female, ph.d) have different connotations: Parallely annotated stylistic language dataset with multiple personas. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *ArXiv*.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations*.
- Ashutosh Kumar, Kabir Ahuja, Raghuram Vadapalli, and Partha Talukdar. 2020. Syntax-guided controlled generation of paraphrases. *Transactions of the Association for Computational Linguistics*.

- Alon Lavie and Abhaya Agarwal. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. In *Proceedings of the Second Workshop on Statistical Machine Translation*, page 228–231.
- 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 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1865–1874.
- Zichao Li, Xin Jiang, Lifeng Shang, and Qun Liu. 2019. Decomposable neural paraphrase generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3403–3414.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Proceedings of the ACL Workshop: Text Summarization*, page 10.
- 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 on Association for Computational Linguistics*, page 311–318.
- 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*, pages 1532–1543.
- Shrimai Prabhunoye, 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*, pages 866–876.
- Aaditya Prakash, Sadid A. Hasan, Kathy Lee, Vivek Datla, Ashequl Qadir, Joey Liu, and Oladimeji Farri. 2016. Neural paraphrase generation with stacked residual LSTM networks. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2923–2934.
- Lihua Qian, Lin Qiu, Weinan Zhang, Xin Jiang, and Yong Yu. 2019. Exploring diverse expressions for paraphrase generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3173–3182.
- Chris Quirk, Chris Brockett, and William Dolan. 2004. Monolingual machine translation for paraphrase generation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 142–149.
- Wei Wang, Piji Li, and Hai-Tao Zheng. 2021. Generating diversified comments via reader-aware topic modeling and saliency detection. *CoRR*, abs/2102.06856.
- Sander Wubben, Antal van den Bosch, and Emiel Krahmer. 2010. Paraphrase generation as monolingual translation: Data and evaluation. In *Proceedings of the 6th International Natural Language Generation Conference*.
- Chen Xu, Jianyu Zhao, Rang Li, Changjian Hu, and Chuangbai Xiao. 2021. Change or not: A simple approach for plug and play language models on sentiment control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 15935–15936.
- Minghong Xu, Piji Li, Haoran Yang, Pengjie Ren, Zhaochun Ren, Zhumin Chen, and Jun Ma. 2020. A neural topical expansion framework for unstructured persona-oriented dialogue generation. *arXiv preprint arXiv:2002.02153*.
- Zichao Yang, Zhiting Hu, Chris Dyer, Eric P Xing, and Taylor Berg-Kirkpatrick. 2018. Unsupervised text style transfer using language models as discriminators. In *Proceedings of the 32th Annual Conference on Advances in Neural Information Processing Systems*, pages 7287–7298.

A Exemplar Searching Algorithm

Algorithm 1 Searching Exemplar Sentences

Require: dataset $\mathbb{D} = (\mathbb{D}_X, \mathbb{D}_Y)$

- 1: **for** Y in \mathbb{D}_Y **do**
- 2: find the sentence set $\mathbb{C}_1 \subseteq \mathbb{D}_X$ that each $C \in \mathbb{C}_1$ satisfies $|\text{len}(C) - \text{len}(Y)| \leq 2$
- 3: find the sentence set $\mathbb{C}_2 \subseteq \mathbb{C}_1$ that for each $C \in \mathbb{C}_2$, the number of shared words between C and Y , denoted by c , satisfies $c + 2 \leq \text{len}(Y)$
- 4: find the exemplar $Z \in \mathbb{C}_2$ which has the smallest POS tag sequence editdistance with Y
- 5: **end for**

The detailed steps for exemplar sentence searching are described in Algorithm 1. Step 2 is done to accelerate the searching procedure since sentences with similar style tend to have similar token lengths. Step 3 guarantees that the content information of Y and the selected Z does not overlap much.

B Style Embedding Quality

We list some sentences retrieved by ablative models given the style embedding of a sentence S in Table 7. For each model, we obtain the top-5 sentences which are most similar to S . We can see that the sentences retrieved by Ours and Ours-w/o-CCL are more similar to S in style dimension than Ours-w/o-both on the whole. For example, in the second case, the fifth sentence of Ours-w/o-both lacks the adverbial modifier compared with S_2 .

S_1	what are newton 's laws of motion?
Ours	what are the after effects of masturbation ? what are the health benefits of coffee ? what are the safety precautions on handling shotguns ? what are some interesting facts about bengaluru ? what are some unknown facts about football ?
Ours-w/o-CCL	what are the health benefits of coffee ? what are the good things about pakistan ? what are some interesting facts about bengaluru ? what are some unknown facts about football ? what are considered abiotic factors of grasslands ?
Ours-w/o-both	what were nelson mandela 's greatest accomplishments ? what are craig good 's qualifications to talk about nutrition ? what are reasons of china 's success ? what are president obama 's greatest accomplishments and failures ? what is newton 's third law of motion with examples ?
S_2	how do i impress a girl on chat ?
Ours	how do i become an engineer in robotics ? how do i get a job in europe countries ? how do i get the crown on musical ly ? how do i make a website responsive without bootstrap ? how do i find the best seo company in delhi ncr ?
Ours-w/o-CCL	how do i become an engineer in robotics ? how do i get a job in europe countries ? how do i leave a girl without hurting her feelings ? how do i get the crown on musical ly ? how do i root a galaxy s550 at t ?
Ours-w/o-both	how do i become an engineer in robotics ? how do i get a job in europe countries ? how do i buy a suit online ? how do i get the crown on musical ly ? how can i help a friend get off drugs ?

Table 7: Sentences Retrieved by Style Embedding

C Multiple Paraphrase Sentences Generation

X_1	which is the best anime to watch ?
Z_{11}	can you jailbreak an ios 8 3 ?
Y_{11}	can you recommend the best anime ?
Z_{12}	which are the best mba colleges in gwaliior ?
Y_{12}	what are the best anime films of all time ?
Z_{13}	how can i earn from online ?
Y_{13}	what anime should i watch now ?
X_2	what type of music do you listen ?
Z_{21}	is equatorial guinea really rich ?
Y_{21}	which music is really good ?
Z_{22}	what tv series have you watched and why did you like them ? explain
Y_{22}	what type of music do you like ? and how do you recommend ?
Z_{23}	what is the best way to reduce weight ?
Y_{23}	what is the best music to listen to ?

Table 8: Generate different sentences given different exemplars.