

Multi-Pair Text Style Transfer for Unbalanced Data via Task-Adaptive Meta-Learning

Xing Han

University of Texas at Austin
Austin, TX, 78712
aaronhan223@utexas.edu

Jessica Lundin

Salesforce
San Francisco, CA, 94105
jlundin@salesforce.com

Abstract

Text-style transfer aims to convert text given in one domain into another by paraphrasing the sentence or substituting the keywords without altering the content. By necessity, state-of-the-art methods have evolved to accommodate nonparallel training data, as it is frequently the case there are multiple data sources of unequal size, with a mixture of labeled and unlabeled sentences. Moreover, the inherent style defined within each source might be distinct. A generic bidirectional (e.g., formal \leftrightarrow informal) style transfer regardless of different groups may not generalize well to different applications. In this work, we developed a task adaptive meta-learning framework that can simultaneously perform a multi-pair text-style transfer using a single model. The proposed method can adaptively balance the difference of meta-knowledge across multiple tasks. Results show that our method leads to better quantitative performance as well as coherent style variations. Common challenges of unbalanced data and mismatched domains are handled well by this method.

1 Introduction

Text-style transfer is a fundamental challenge in natural language processing. Applications include non-native speaker assistants, child education, personalization and generative design (Fu et al., 2017; Zhou et al., 2017; Yang et al., 2018a; Gatys et al., 2016b,a; Zhu et al., 2017; Li et al., 2017). Figure 1 shows a prominent example on applying style transfer into a hypothetical online shopping platform, where the generated style variations can be used for personalized recommendations. However, compared with other domains, the lack of parallel corpus and quality training data is currently an obstacle for text-style transfer research. For example, assume one supports a multi-tenant service platform including tenant-specific

text data, but there is no guarantee that each tenant will provide sufficient amount of data for model training. To build a multi-task language model that matches the text-style of each tenant is more practical and efficient than training individual models. This single-model approach might also have relatively favorable empirical performance.

Existing works on text style transfer have addressed different applications such as sentiment transfer (Shen et al., 2017), word decipherment (Knight et al., 2006), and author imitation (Xu et al., 2012). If parallel training data is available, a wide range of supervised techniques in machine translation (e.g., Seq2Seq models (Bahdanau et al., 2014) and Transformers (Vaswani et al., 2017)) can also be applied to style transfer problems. For non-parallel data, He et al. (2020) proposed a probabilistic formulation that models non-parallel data from two domains as a partially observed parallel corpus, and learn the style transfer model in a completely unsupervised fashion. Unsupervised machine translation method has also been adapted to this setting (Zhang et al., 2018). In recent research focused on learning disentangled content and style representations using adversarial training (John et al., 2018; Yang et al., 2018b; Shen et al., 2017), models are designed for non-parallel data while preserving content. Lample et al. (2018) argued that the adversarial models are not really doing disentanglement, and proposed a denoising auto-encoding approach instead. Another way to approach this problem is through identifying and substituting style-related sub-sentences (Li et al., 2018; Sudhakar et al., 2019), where the unchanged part guarantees consistency over content. Additionally, state-of-the-art language models (BERT (Devlin et al., 2018), GPT-2 (Radford et al., 2019), CTRL (Keskar et al., 2019), etc.) and text-to-text models (Raffel et al., 2019) achieve good performance generating text

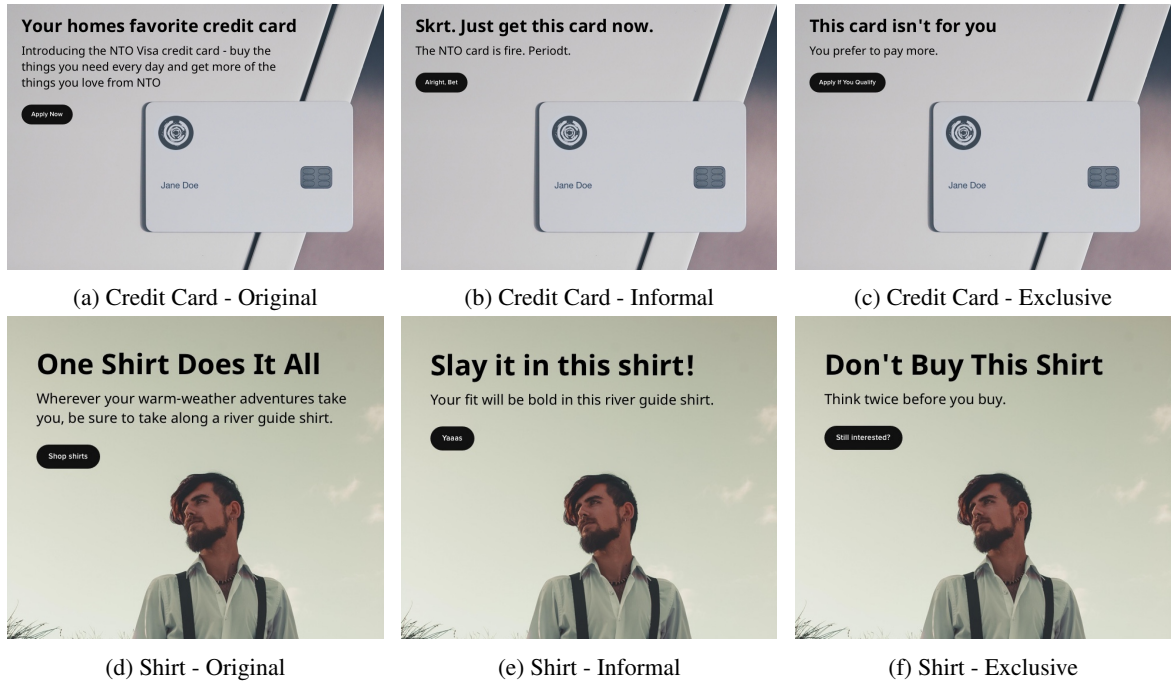


Figure 1: Text style transfer examples in generative design: the original text is meta-data from e-commerce websites. Two target style variations are predefined for certain groups of customers.

in different styles on multiple tasks (Dathathri et al., 2019; Wolf et al., 2019).

Building upon previous work, we aim to bridge real applications while accounting for the aforementioned data problems. Specifically, we wish to design an efficient training method for a style transfer model that 1. quickly learns and adapts to different style domains with limited data; 2. handling class-imbalance and out-of-distribution tasks. To achieve this, we introduce meta-learning into the style-transfer problem.

Meta-learning (Schmidhuber, 1987) is a method to enable generalization ability to a model over a distribution of tasks. We focus on optimization-based meta-learning for our applications. MAML (Finn et al., 2017) learns a common initialization parameter for each task using a few gradient steps. This standard MAML approach has been applied to text style transfer problems with low resources (Chen and Zhu, 2020) and achieved better performance in this situation. However, this method did not take into account the internal variations between tasks. A similar algorithm called Reptile (Nichol et al., 2018) achieves better performance by maximizing the inner product between gradient of different mini-batches from the same task in its update. Recent works (Qiao et al., 2018; Lee and Choi, 2018) improved a single meta-learner to task-adaptive meta-learning models, which in-

cludes task-specific parameters to help generalize better between tasks. Bayesian meta-learning is another active area of research: Finn et al. (2018) proposed a probabilistic version of MAML, where the variational inference framework utilizes a task-specific gradient update. More recently, Lee et al. (2019) incorporated a Bayesian framework into task-adaptive meta-learning. Specifically, they introduce balancing variables for task and class-specific learning and leverage the uncertainties of these parameters derived from training data statistics. In this paper, we will adapt the Bayesian task adaptive meta-learning (TAML) for our application shown in Figure 2 overview.

2 Balancing Variations between Tasks

A common challenge in aforementioned real application is that data from multiple sources may suffer from different problems, such as insufficient training samples, unbalanced class labels, or domain mismatch. However, simply ignore these differences and concatenate all tenants' data for model training will not lead to ideal results.

Meta-learning is one of the most relevant approaches for generalized learning from few samples of different tasks. Assume a task distribution $p(\tau)$ that randomly generates task τ consisting of training set $\mathcal{D}^\tau = \{X^\tau, \bar{X}^\tau\}$ and a test set




Figure 2: An overview of our multi-pair style transfer method: assume learning from each tenant’s data is a task, and the training data available for each task varies. The style transfer model can adaptively learn tasks using our method and the resulting model performs style transfer across multiple domains.

$\mathcal{D}_{\text{test}}^\tau = \{X_{\text{test}}^\tau, \bar{X}_{\text{test}}^\tau\}$. If parallel training data is not available, then we only have $\mathcal{D}^\tau = X^\tau$ and $\mathcal{D}_{\text{test}}^\tau = X_{\text{test}}^\tau$. The MAML algorithm initialize task-specific parameter θ^τ using a few gradient steps on a small amount of data. In this case, the optimized parameters can generalize to new tasks. Specifically, we have the loss minimization

$$\min_{\theta} \sum_{\tau \sim p(\tau)} \mathcal{L}(\theta - \alpha \nabla_{\theta} \mathcal{L}(\theta; \mathcal{D}^\tau); \mathcal{D}_{\text{test}}^\tau), \quad (1)$$

where α is the step size when learning each task. The initial parameter of each task then becomes $\theta^\tau = \theta - \alpha \nabla_{\theta} \mathcal{L}(\theta; \mathcal{D}^\tau)$, which has been proved to minimize the test loss $\mathcal{L}(\theta^\tau; \mathcal{D}_{\text{test}}^\tau)$. The training set \mathcal{D}^τ may consist of only a few samples.

Eq (1) is effective in numerous applications, yet insufficient in addressing our data problems, as it treats the initialization and learning parameters with equal importance for each task. Inspired by Lee et al. (2019), we now introduce three balancing variables: $z^\tau, \gamma^\tau, \omega^\tau$ for every task τ .

Let $\omega^\tau = (\omega_1^\tau, \dots, \omega_C^\tau) \in [0, 1]^C$ be the multiplier of each of the class specific gradients to vary the learning rate for each class. In real applications, we often have a style transfer problem with unbalanced training data. For instance, when training formality style transfer models, the number of formal/positive sentences is normally much larger than the number of informal/exclusive sentences. Also, denote $\gamma^\tau = (\gamma_1^\tau, \dots, \gamma_L^\tau) \in [0, \infty)^L$ to be the multipliers of the original learning-rate α , where the new learning rate becomes $\gamma_1^\tau \alpha, \gamma_2^\tau \alpha, \dots, \gamma_L^\tau \alpha$. Note that the value of γ is task-dependent (e.g., sample size of the training data from each task), and is meant to deal with the small data problem in multi-pair text style transfer. Moreover, since the text data collected from every source or tenant is very hard to be aligned, it is common to have training data with significantly different context. We can treat this as an out of dis-

tribution problem and this can be reflected on the value of initial parameters. We use z^τ to modulate the initial parameter θ for each task. Specifically, z^τ relocates the initial θ to a task-dependent starting point prior to the learning process. We unify these properties as the learning framework below:

$$\theta_0 = \theta * z^\tau, \quad \text{and for } k = 1, \dots, K :$$

$$\theta_k = \theta_{k-1} - \gamma^\tau \circ \alpha \circ \sum_{c=1}^C \omega_c^\tau \nabla_{\theta_{k-1}} \mathcal{L}(\theta_{k-1}; \mathcal{D}_c^\tau), \quad (2)$$

where ω_c and \mathcal{D}_c are class-specific parameters and data; K is the total number of iterations for updating parameters. We currently assume $C = 2$ in the following discussions of this paper, since pair-wise style transfer is the primary problem of interest so far.

3 Learning the Balancing Variables through Variational Inference

We now discuss how to find the most suitable value of each balancing variable. We employ the variational inference framework from probabilistic MAML (He et al., 2020) and TAML (Lee et al., 2019) to extract the task-specific information. The variational inference framework is used to compute posterior distributions for the balancing variables $z^\tau, \gamma^\tau, \omega^\tau$. Assume the training data $X^\tau = \{x_n^\tau\}_{n=1}^{N_\tau}$, $\bar{X}^\tau = \{\bar{x}_n^\tau\}_{n=1}^{N_\tau}$; test data $X_{\text{test}}^\tau = \{x_m^\tau\}_{m=1}^{M_\tau}$, $\bar{X}_{\text{test}}^\tau = \{\bar{x}_m^\tau\}_{m=1}^{M_\tau}$, and $\phi^\tau = \{\tilde{\omega}^\tau, \tilde{\gamma}^\tau, \tilde{z}^\tau\}$ to be a collection of three balancing variables. The goal of learning for each task τ is to maximize the conditional log-likelihood of the joint dataset $\mathcal{D}_{\text{test}}^\tau$ and \mathcal{D}^τ : $\log p(\bar{X}_{\text{test}}^\tau, \bar{X}^\tau | X_{\text{test}}^\tau, X^\tau; \theta)$. To solve the optimization problem requires determining the true posterior $p(\phi^\tau | \mathcal{D}^\tau, \mathcal{D}_{\text{test}}^\tau)$, which is intractable. We resort to variational inference with a tractable




Figure 3: The inference network for generating posterior distribution of balancing variables γ , z , and ω for task i .

form of approximate posterior $q(\phi^\tau | \mathcal{D}^\tau, \mathcal{D}_{\text{test}}^\tau; \psi)$ parameterized by ψ . In order to make the inference network of meta training and meta testing consistent, we drop the dependency of $\mathcal{D}_{\text{test}}^\tau$ since the test labels are unknown in meta-testing. Hence the approximate posterior becomes $q(\phi^\tau | \mathcal{D}^\tau; \psi)$. We now have the approximated lower bound for task adaptive meta learning:

$$\mathcal{L}_{\theta, \psi}^\tau = \frac{N_\tau + M_\tau}{M_\tau} \sum_{m=1}^{M_\tau} E_{q(\phi^\tau | \mathcal{D}^\tau; \psi)} [\log p(\tilde{y}_m^\tau | \tilde{x}_m^\tau, \phi^\tau; \theta)] - \text{KL}[q(\phi^\tau | \mathcal{D}^\tau; \psi) \| p(\phi^\tau)]. \quad (3)$$

Given that each balancing variable is independent, $q(\phi^\tau | \mathcal{D}^\tau; \psi)$ can therefore be fully factorized

$$q(\phi^\tau | \mathcal{D}^\tau; \psi) = \prod_c q(\omega_c^\tau | \mathcal{D}^\tau; \psi) \prod_l q(\gamma_l^\tau | \mathcal{D}^\tau; \psi) \prod_i q(z_i^\tau | \mathcal{D}^\tau; \psi).$$

We assume each single dimension of $q(\phi^\tau | \mathcal{D}^\tau; \psi)$ follows a uni-variate Gaussian distribution with trainable mean and variance. Given $\phi_s^\tau \sim q(\phi^\tau | \mathcal{D}^\tau; \psi)$, we then use the Monte-Carlo approximation on Eq (3) as a new objective:

$$\min_{\theta, \psi} \frac{1}{M_\tau} \sum_{m=1}^{M_\tau} \frac{1}{S} \sum_{s=1}^S -\log p(\tilde{y}_m^\tau | \tilde{x}_m^\tau, \phi_s^\tau; \theta) + \frac{1}{N_\tau + M_\tau} \text{KL}[q(\phi^\tau | \mathcal{D}^\tau; \psi) \| p(\phi^\tau)]. \quad (4)$$

To better model the variational distribution $q(\phi^\tau | \mathcal{D}^\tau; \psi)$, an informative representation encoded from the training dataset \mathcal{D}^τ is necessary. In this case, the inference network can capture all useful statistical information in \mathcal{D}^τ to recognize its imbalances. We use a two-stage hierarchical set encoder, for a given text style transfer task, we first

encodes each class, and then encodes the whole set of classes. Define the encoder $\text{StatisticsPooling}(\cdot)$ that generates concatenation of the class statistics such as mean, variance and cardinality. The two-stage encoder first encodes all text sentences of each class into s_c , followed by encoding representations of the whole set of classes:

$$v^\tau = \text{StatisticsPooling}(\{\text{NN}_2(s_c)\}_{c=1}^C), \\ s_c = \text{StatisticsPooling}(\{\text{NN}_1(x)\}_{x \in X_c^\tau}),$$

where $c = 1, \dots, C$ represents classes; X_c^τ is the collection of class c examples in task τ ; NN_1 and NN_2 are some neural networks parameterized by ψ . Therefore, the summarized feature vectors of \mathcal{D}^τ can be used to infer the Gaussian distribution parameters of balancing variables ω^τ , z^τ and γ^τ to be further applied in the update of meta-learning. Note that since the balancing variable ω is class-specific, inference its distributional parameters does not need to go through the second stage of encoding. The overall structure of the inference network is shown in Figure 3.

4 Task-Adaptive Style Transfer

We discuss formulation of multi-pair text-style transfer problem using the TAML framework. An overview of our method is shown in Figure 2. We assume training data in each task could either be parallel (task 1 and 4) or non-parallel (task 2 and 3). The number of training samples in task i is represented by N_i , which is not necessarily equal for each task. In addition, the class distribution in non-parallel training data is heavily skewed.

We now formulate our problem as follows. Given a distribution of similar tasks $p(\tau)$, each task represents performing text style transfer on a certain dataset \mathcal{D}^τ . Define a generic loss function \mathcal{L} and shared parameters θ within tasks, the

Algorithm 1 Multi-Pair Text Style Transfer via TAML

```
1: Input: style pair for each task  $\tau$ ,  $\{(s^\tau, \bar{s}^\tau)\}_{\tau=1}^T$ , parameters  $\alpha, \beta$ ,
2: Meta-training procedure:
3: while not done do
4:   for each style pair  $(s^\tau, \bar{s}^\tau)$  do
5:     Train inference network  $q(\phi^\tau | \mathcal{D}^\tau; \psi)$  by minimizing objective (4)
6:     Obtain balancing variables  $\{z^\tau, \gamma^\tau, \omega^\tau\} \sim q(\phi^\tau | \mathcal{D}^\tau; \psi)$ 
7:     Initialize sub-learner with  $\theta_0^\tau = \theta * z^\tau$ 
8:     for step in  $1, \dots, K$  do
9:       Sample batch data from  $\mathcal{D}_s^\tau$ 
10:      Update parameters for task  $\tau$  using  $\theta_k^\tau = \theta_{k-1}^\tau - \gamma^\tau \circ \alpha \circ \sum_{c=1}^2 \omega_c^\tau \nabla_{\theta_{k-1}^\tau} \mathcal{L}(\theta_{k-1}^\tau, \mathcal{D}_s^\tau)$ 
11:    end for
12:    Sample batch data from  $\mathcal{D}_t^\tau$ 
13:    Evaluate  $\mathcal{L}(\theta_K^\tau, \mathcal{D}_t^\tau)$ 
14:  end for
15:  Update meta-learner  $f_\theta$  with  $\theta = \theta - \beta \nabla_\theta \sum_{\tau=1}^T \mathcal{L}(\theta^\tau, \mathcal{D}_t^\tau)$ 
16: end while
17: Meta testing:  $Y^\tau \leftarrow f_\theta(X_{\text{test}}^\tau, \mathcal{S}_{\text{test}}^\tau)$ 
```

goal is to jointly learn a task-agnostic model $f_\theta : (X^\tau, \mathcal{S}^\tau) \mapsto Y^\tau$, where for each τ , \mathcal{S}^τ is the corresponding set of style labels of original text X^τ , and Y^τ is the resulting style transformed text. Ideally, Y^τ should be consistent with \bar{X}^τ , the corresponding input text sentence in another style domain which may or may not be available in model training. In fine-tuning with a new task, the parameters are initialized accounting for the imperfect nature of the given dataset. Similar to the standard meta-learning approach, the training data of task τ is divided into a support set \mathcal{D}_s^τ and a query set \mathcal{D}_t^τ , where \mathcal{D}_s^τ is used to update each sub-task and \mathcal{D}_t^τ is used to evaluate the loss, and later used for meta-learner updates. A detailed description can be found in Algorithm 1.

5 Experiments

We conduct experiments on multiple style-transfer datasets: Shakespeare (Xu et al., 2012), Yelp reviews (Shen et al., 2017) and an internal dataset from a company that contains formal/informal text sentences. Performing style transfer on each of the above dataset defines a unique task. The Shakespeare dataset contains 21k parallel sentences, which includes original text style and Shakespeare’s style. The maximum length of the sentences is 20. The Yelp dataset contains around 252k sentences of positive and negative restaurant reviews, where we use a maximum length of 15 to conduct the experiment. We evaluate our method using state-of-the-art transformers including BERT, GPT-2, T5, and VAE (John et al., 2018) designed for style transfer by learning disentangled representations. Our baseline method in-

cludes regular model training without distinguishing the difference between tasks, and the MAML method in Eq (1) to fine-tune the style transfer models on multiple distinct tasks, which has also been proposed by Chen and Zhu (2020). We then employ Algorithm 1 to adaptively fine-tune the style transfer models for each task.

The unbalanced training data is created by sampling from each class at different rate (75% positive class, 25% negative class). We use the pre-trained transformers in Huggingface library (Wolf et al., 2020) as our initial style transfer models. Specifically, we build a two-head model (Figure 4) on top of the decoders where each head is composed of multiple dense layers. We do not perform end-to-end training for the entire transformer but only train the two-head model. The model input is the sentence and style pairs $(X^\tau, \mathcal{S}^\tau)$ while the forward propagation of transformer’s output to each model head is dependent on the style labels. The resulting output sentences are style dependent, and one can perform text-style transfer by flipping the style labels during the inference phase. Similarly, we use both baseline and TAML to train VAE and obtain disentangled style and content representations, and replace the style embedding during the inference stage to get style transferred sentences. Note that we focus on improving the fine-tuning part of text style transfer models, while we do not modify the model structure themselves. In terms of content preservation, the objective function of the VAE model proposed by (John et al., 2018) contains a content-oriented loss, while for other transformer-based models, we designed the loss \mathcal{L} to be the cross entropy loss




Figure 4: Two-head architecture

Model	Shakespeare			Yelp		
	BLEU [↑]	PPL [↓]	ACC [↑]	BLEU [↑]	PPL [↓]	ACC [↑]
BERT	12.04	26.43	78.77	9.56	15.31	74.68
GPT-2	2.83	38.47	74.45	4.81	45.49	76.67
T5	3.65	59.39	82.58	5.22	39.41	75.14
VAE	14.36	22.29	81.92	10.81	10.65	77.27
MAML-BERT	16.31	21.09	79.34	10.87	15.02	74.98
MAML-GPT-2	7.01	36.94	75.25	5.04	41.76	77.06
MAML-T5	4.77	50.44	83.02	6.46	33.72	75.86
MAML-VAE	15.52	21.45	81.96	11.74	11.04	77.24
TAML-BERT	17.56	19.36	79.34	11.02	16.82	75.22
TAML-GPT-2	7.42	36.67	76.02	5.63	37.66	77.18
TAML-T5	4.81	47.23	83.45	6.92	32.30	75.64
TAML-VAE	17.98	20.14	82.61	12.31	10.59	77.33

Table 1: Evaluations of multiple text style transfer models on testing set of the listed data. TAML-based model training methods achieve better performance on multi-task text style transfer.

between Y^τ and \bar{X}^τ , or between Y^τ and X^τ in non-parallel situations.

For BERT, GPT-2, and T5, we use the built-in vocabulary within the transformers library. Adam optimizer is used with learning of 5×10^{-4} to train the model. The batch size is set to 16 and the model is trained for 100 epochs. We build the two-head model by using 6 fully connected layers with hidden size of 256 and ReLU activation function. The parameters are chosen empirically with the best performance. For VAE approach, we use the same parameter settings as reported in (John et al., 2018). As for NN_1 in inference network, we used two consecutive blocks of 3×3 convolution layer followed by 2×2 max pooling layer, the output is then fed into one fully connected layer for statistics pooling. We then use two fully connected layers for NN_2 . All the activation functions are ReLU.

We evaluate competing methods on quality and accuracy of style transfer. The adopted metrics are common choices among recent works.

BLEU: We use BLEU (Papineni et al., 2002) score to evaluate the content preservation, the scores are calculated using ScareBLEU (Post, 2018). When parallel sentences are available, we

compute the BLEU score between the style transferred sentences Y^τ and the ground truth sentences \bar{X}^τ . Otherwise, we use the original sentences X^τ instead.

PPL: We implemented a bigram language model (Kneser and Ney, 1995) to quantitatively evaluate the fluency of a sentence. The language model is trained on the target-style domain, and we report the PPL of the generated sentences.

Accuracy: We also trained a TextCNN classifier (Rakhlina, 2016) simultaneously while training style transfer models. The trained classifier is then used to evaluate the classification accuracy on the generated sentences.

Table 1 shows our results for each method. By applying task-adaptive meta learning on each style-transfer model, the performance with respect to every metric is generally improved on the datasets we evaluated. We observe that the VAE method performs better in style transfer, as other models are not explicitly designed for this goal.

6 Conclusion

In this paper, we investigated meta-learning approaches applied to text-style transfer, for situations with multiple data sources. Given the distinct context and total amount of data, we propose a task-adaptive meta-learning approach to fine-tune style-transfer models. The proposed method introduces three balancing variables with probabilistic distributions, which can be encoded from training data. These balancing variables are then used to solve class and task imbalance problems. Empirically, we found that TAML improves the style-transfer performance on multiple models. In the future, we wish to explore generating style variations in more fine-grained levels (for $C > 2$) with the help of meta-learning.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Xiwen Chen and Kenny Q Zhu. 2020. St 2: Small-data text style transfer via multi-task meta-learning. *arXiv preprint arXiv:2004.11742*.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2019. Plug and play language mod-

- els: a simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*.
- Chelsea Finn, Kelvin Xu, and Sergey Levine. 2018. Probabilistic model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, pages 9516–9527.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2017. Style transfer in text: Exploration and evaluation. *arXiv preprint arXiv:1711.06861*.
- Leon A Gatys, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. 2016a. Preserving color in neural artistic style transfer. *arXiv preprint arXiv:1606.05897*.
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2016b. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423.
- Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. 2020. A probabilistic formulation of unsupervised text style transfer. *arXiv preprint arXiv:2002.03912*.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2018. Disentangled representation learning for non-parallel text style transfer. *arXiv preprint arXiv:1808.04339*.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Reinhard Kneser and Hermann Ney. 1995. Improved backing-off for m-gram language modeling. In *1995 international conference on acoustics, speech, and signal processing*, volume 1, pages 181–184. IEEE.
- Kevin Knight, Anish Nair, Nishit Rathod, and Kenji Yamada. 2006. Unsupervised analysis for decipherment problems. In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 499–506.
- Guillaume Lample, Sandeep Subramanian, Eric Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2018. Multiple-attribute text rewriting. In *International Conference on Learning Representations*.
- Hae Beom Lee, Hayeon Lee, Donghyun Na, Sae-hoon Kim, Minseop Park, Eunho Yang, and Sung Ju Hwang. 2019. Learning to balance: Bayesian meta-learning for imbalanced and out-of-distribution tasks. *arXiv preprint arXiv:1905.12917*.
- Yoonho Lee and Seungjin Choi. 2018. Gradient-based meta-learning with learned layerwise metric and subspace. *arXiv preprint arXiv:1801.05558*.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. Delete, retrieve, generate: A simple approach to sentiment and style transfer. *arXiv preprint arXiv:1804.06437*.
- Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. 2017. Demystifying neural style transfer. *arXiv preprint arXiv:1701.01036*.
- Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*.
- 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.
- Matt Post. 2018. **A call for clarity in reporting BLEU scores**. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L Yuille. 2018. Few-shot image recognition by predicting parameters from activations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7229–7238.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- A Rakhlin. 2016. Convolutional neural networks for sentence classification. *GitHub*.
- Jürgen Schmidhuber. 1987. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. Ph.D. thesis, Technische Universität München.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Advances in neural information processing systems*, pages 6830–6841.

- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. Transforming delete, retrieve, generate approach for controlled text style transfer. *arXiv preprint arXiv:1908.09368*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- 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. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- 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.
- Wei Xu, Alan Ritter, Bill Dolan, Ralph Grishman, and Colin Cherry. 2012. Paraphrasing for style. In *Proceedings of COLING 2012*, pages 2899–2914.
- Cheng Yang, Maosong Sun, Xiaoyuan Yi, and Wenhao Li. 2018a. Stylistic chinese poetry generation via unsupervised style disentanglement. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3960–3969.
- Zichao Yang, Zhiting Hu, Chris Dyer, Eric P Xing, and Taylor Berg-Kirkpatrick. 2018b. Unsupervised text style transfer using language models as discriminators. In *Advances in Neural Information Processing Systems*, pages 7287–7298.
- 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*.
- Ganbin Zhou, Ping Luo, Rongyu Cao, Fen Lin, Bo Chen, and Qing He. 2017. Mechanism-aware neural machine for dialogue response generation. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232.