

Improving Text Generation with Student-Forcing Optimal Transport

Jianqiao Li^{*1}, Chunyuan Li^{*2}, Guoyin Wang^{*3}, Hao Fu¹, Yuh-Chen Lin³

Liqun Chen¹, Yizhe Zhang², Chenyang Tao¹, Ruiyi Zhang¹

Wenlin Wang¹, Dinghan Shen⁴, Qian Yang¹ and Lawrence Carin¹

¹Duke University ²Microsoft Research ³Amazon ⁴Microsoft Dynamics 365 AI

jianqiao.li@duke.edu, chunyl@microsoft.com, guoyiwan@amazon.com

Abstract

Neural language models are often trained with maximum likelihood estimation (MLE), where the next word is generated conditioned on the ground-truth word tokens. During testing, however, the model is instead conditioned on previously generated tokens, resulting in what is termed *exposure bias*. To reduce this gap between training and testing, we propose using optimal transport (OT) to match the sequences generated in these two modes. An extension is further proposed to improve the OT learning, based on the structural and contextual information of the text sequences. The effectiveness of the proposed method is validated on machine translation, text summarization, and text generation tasks.

1 Introduction

Natural language generation is an essential component of many NLP applications, such as machine translation (Bahdanau et al., 2015), image captioning (You et al., 2016), text summarization (See et al., 2017), dialogue systems (Vinyals and Le, 2015), and machine comprehension (Nguyen et al., 2016). Generating human-like natural language is typically cast as predicting a sequence of consecutive words in a recurrent manner. Maximum likelihood estimation (MLE) is commonly employed as the objective to train such text-generation models, maximizing the log-likelihood of producing the ground-truth tokens within a sentence or paragraph (Salakhutdinov, 2015). In Recurrent Neural Network (RNN) models, this is also known as Teacher-Forcing (TF) (Williams and Zipser, 1989), due to the use of ground-truth tokens for next-token prediction.

However, in the maximum likelihood paradigm, previous observed tokens are usually provided during training, giving rise to an issue termed *exposure bias* (Bengio et al., 2015): the ground-truth

tokens seen by the model during training are not available at inference time. During inference, the model is required to use outputs from the last step instead of the unseen ground-truth, which is often referred to as Student-Forcing (SF). As a result, there is a discrepancy between training and inference, accumulating errors along the sequence-generation trajectory (Ranzato et al., 2016a).

This challenge has been addressed by incorporating model-generated text at training time. Bengio et al. (2015) proposed *scheduled sampling* (SS), where the training samples are systematically tampered with by replacing some of the ground-truth tokens with model-predicted tokens. Other works regard text generation as a sequential decision making problem, applying reinforcement learning (RL) techniques (Ranzato et al., 2016b; Bahdanau et al., 2017). In particular, quantitative evaluation metrics such as BLEU and ROGUE are used as sequence-level rewards for model-generated texts.

Despite the encouraging results reported, concerns have been raised w.r.t. the above strategies. Scheduled sampling is known to be statistically inconsistent and fails to address the fundamental issues (Huszár, 2015). RL solutions usually suffer from slow and unstable training due to the high variance of REINFORCE-based policy gradients; consequently, specific training techniques are often needed (Rennie et al., 2017; Liu et al., 2018). Additionally, the results from RL models often do not correlate well with human evaluations, as the rewards used are typically biased towards specific aspects of a language model (Wang et al., 2018).

On the other hand, recent developments in likelihood-free modeling techniques, most prominently the *generative adversarial networks* (GANs), tackle exposure bias in a more principled fashion (Lamb et al., 2016; Yu et al., 2017; Zhang et al., 2017; Li et al., 2017). GAN-based text models, however, suffer from a number of severe dif-

* Equal contribution

difficulties, including mode collapse, where generated text looks real but lacks necessary diversity (Zhu et al., 2018; Caccia et al., 2018). Additionally, the training of GAN-based models is often unstable, and model learning easily breaks down in the event of vanishing or exploding gradients (Arjovsky et al., 2017; Zhang et al., 2017). Therefore, existing adversarial methods may not be able to match sentences generated by student-forcing with ground-truth sentences.

To mitigate the challenges from the adversarial methods, we utilize a sequence-matching loss based on Optimal Transport (OT), which avoids a neural discriminator and delivers a smoother gradient for the generator. Recently, Chen et al. (2019) leverage OT loss based on a Teacher-Forcing scheme, however, it degenerates to word-level matching, making it difficult to capture temporal semantic information. In this work, in order to enable sequence-level matching, we instead propose an OT-based sequence-level training scheme to directly optimize the discrepancy loss between the ground-truth and free-running text samples. Further, we introduced various OT cost functions for loss calculation. The significance of this work is two-fold: firstly, this approach alleviates the exposure bias, boosting model performance at the inference stage by using a sequence-level objective between free-running output and reference. Secondly, with the use of OT, this approach provides a direct objective that is easy and robust to optimize, without biasing towards a specific, manually-defined metric.

Our work provides the following contributions: *i)* We introduce a novel method for text generation called Student-Forcing OT (SFOT), leveraging OT loss to improve long-term sequence sampling. *ii)* A new context-preserving OT approach is proposed to effectively match a text sequence with order information. *iii)* We examine the necessity of integrating OT with Student-Forcing via Imitation Learning. *iv)* The proposed models are robust demonstrated by extensive empirical evaluations on Neural Machine Translation (NMT), Text Summarization, and Neural Text Generation (NLG).

2 Student Forcing Optimal Transport

To reduce exposure bias, the output sequences of the generator in the teacher-forcing (training) and student-forcing (inference) stages should be indis-

tinguishable. Therefore, we propose to use OT loss to measure sequence matching distance between the two stages, in conjunction with the maximum likelihood estimate.

2.1 Maximum Likelihood Estimate

We denote the training dataset as N sequence pairs $\mathcal{D} = \{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$, with output sequence $\mathbf{x} = [x_1, \dots, x_T]$ and input sequence $\mathbf{y} = [y_1, \dots, y_{T'}]$. Depending on the specific task, \mathbf{y} may have different definitions. For *seq2seq* models like neural machine translation, \mathbf{y} represents the source sequence, conditioned on which the target sequence \mathbf{x} is generated. For language modeling tasks, \mathbf{y} is empty, and \mathbf{x} becomes the unconditionally generated sequence.

To generate a text sequence, neural language models (Mikolov et al., 2010) generate every token x_t conditioned on the previous tokens in an auto-regressive manner:

$$\log p_{\theta}(\mathbf{x}|\mathbf{y}) = \sum_{t=1}^T \log p_{\theta}(x_t|x_{<t}, \mathbf{y}) \quad (1)$$

where θ are model parameters, and $x_{<t}$ indicates all tokens before t . Learning θ is often performed with maximum likelihood estimation (MLE):

$$\mathcal{L}_{\text{MLE}} = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\log p_{\theta}(\mathbf{x}|\mathbf{y})] \quad (2)$$

To facilitate MLE training, the *teacher-forcing* scheme is considered. The word tokens $x_{<t}$ from the ground-truth sequence are fed into (1) to generate the next token:

$$\tilde{\mathbf{x}} = [\tilde{x}_1, \dots, \tilde{x}_T] \text{ with } \tilde{x}_t \sim p_{\theta}(\tilde{x}_t|x_{<t}, \mathbf{y}) \quad (3)$$

Learned neural language models are often evaluated using the *student-forcing* scheme. The previously generated word tokens of the model are conditioned to generate the next token:

$$\hat{\mathbf{x}} = [\hat{x}_1, \dots, \hat{x}_T] \text{ with } \hat{x}_t \sim p_{\theta}(\hat{x}_t|\hat{x}_{<t}, \mathbf{y}) \quad (4)$$

The difference between (3) and (4) reveals a gap between training and evaluation in the MLE method. To reduce the gap, a natural idea is to define a tractable function to measure the discrepancy between ground-truth and SF-generated text sequences, to regularize TF-based MLE learning.

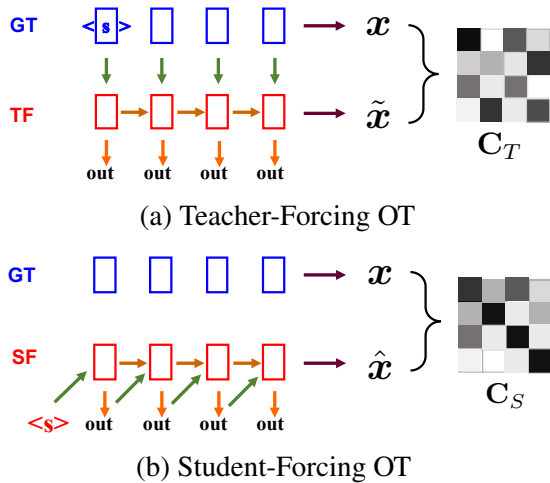


Figure 1: Comparison of (a) teacher-forcing (TF) and (b) student-forcing (SF), where $\langle s \rangle$ is the start token. In (a), the ground-truth (GT) sequence \mathbf{x} is compared with TF-generated sequence $\tilde{\mathbf{x}}$ to produce the cost matrix \mathbf{C}_T ; In (b), the GT sequence \mathbf{x} is compared with SF-generated sequence $\hat{\mathbf{x}}$ to produce the cost matrix \mathbf{C}_S .

2.2 Student Forcing Optimal Transport

We propose to use optimal transport (OT) to measure the discrepancy between the student-forcing generated sequence $\hat{\mathbf{x}}$ and the ground-truth sequence \mathbf{x} . Assuming there is an oracle/target distribution $\mu(\mathbf{x})$ to generate \mathbf{x} , our goal is to learn θ such that the model distribution $p_\theta(\hat{\mathbf{x}})$ matches $\mu(\mathbf{x})$. Formally, OT provides a distance metric between the two probability measures μ and p on a domain \mathbb{X} (the sequence of word tokens):

$$\mathcal{R}_{\text{OT}}(\mathbf{x}, \hat{\mathbf{x}}) = \inf_{\gamma \in \Pi(\mu, p)} \mathbb{E}_{(\mathbf{x}, \hat{\mathbf{x}}) \sim \gamma} [c(\mathbf{x}, \hat{\mathbf{x}})], \quad (5)$$

where $\Pi(\mu, p)$ denotes the set of all joint distributions $\gamma(\mathbf{x}, \hat{\mathbf{x}})$ with marginals $\mu(\mathbf{x})$ and $p(\hat{\mathbf{x}})$. The function $c(\mathbf{x}, \hat{\mathbf{x}}) : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ defines the cost of moving $\hat{\mathbf{x}}$ to \mathbf{x} . Intuitively, OT provides a method of matching the sequence $\hat{\mathbf{x}}$ to \mathbf{x} with the minimum cost, given μ, p and $c(\cdot, \cdot)$.

OT distance on discrete domains For discrete distributions μ, p on \mathbb{X} , we have $\mu = \sum_{i=1}^T u_i \delta_{\mathbf{x}_i}$ and $p = \sum_{j=1}^T p_j \delta_{\hat{\mathbf{x}}_j}$ with $\delta_{\mathbf{x}}$ the Dirac function centered on \mathbf{x} . The weight vectors $\mathbf{u} = \{u_i\}_{i=1}^T \in \Delta^T$ and $\mathbf{p} = \{p_j\}_{j=1}^T \in \Delta^T$ belong to the T -dimensional simplex, *i.e.*, $\sum_{i=1}^T u_i = \sum_{j=1}^T p_j = 1$, as both μ and p are probability distributions. Under such a setting, computing the OT distance is equivalent to solving the following network-flow problem (Luise et al., 2018):

$$\begin{aligned} \mathcal{R}_{\text{OT}}(\mathbf{x}, \hat{\mathbf{x}}) &= \min_{\mathbf{M} \in \Pi(\mathbf{u}, \mathbf{p})} \sum_{i=1}^T \sum_{j=1}^T \mathbf{M}_{ij} \cdot c(\mathbf{x}_i, \hat{\mathbf{x}}_j) \\ &= \min_{\mathbf{M} \in \Pi(\mathbf{u}, \mathbf{p})} \langle \mathbf{M}, \mathbf{C} \rangle \end{aligned} \quad (6)$$

where $\Pi(\mathbf{u}, \mathbf{p}) = \{\mathbf{M} \in \mathbb{R}_+^{T \times T} | \mathbf{M} \mathbf{1}_T = \mathbf{u}, \mathbf{M}^\top \mathbf{1}_T = \mathbf{p}\}$, $\mathbf{1}_T$ denotes a T -dimensional all-one vector, \mathbf{C} is the cost matrix given by $\mathbf{C}_{ij} = c(\mathbf{x}_i, \hat{\mathbf{x}}_j)$, and $\langle \mathbf{M}, \mathbf{C} \rangle = \text{Tr}(\mathbf{M}^\top \mathbf{C})$ represents the Frobenius dot-product. We refer to the minimizer \mathbf{M}^* of (6) as *OT matching*.

Summarizing, our student-forcing optimal transport (SFOT) objective is:

$$\begin{aligned} \mathcal{L}_{\text{SFOT}} &= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\log p_\theta(\mathbf{x} | \mathbf{y}) + \\ &\quad \lambda \mathbb{E}_{\hat{\mathbf{x}} \sim p_\theta(\hat{\mathbf{x}} | \mathbf{y})} \mathcal{R}_{\text{OT}}(\mathbf{x}, \hat{\mathbf{x}})] \end{aligned} \quad (7)$$

where λ is the weighting hyper-parameter that balances the MLE and OT losses. In practice, we only take one sample in student-forcing. Note that our SFOT objective considers $\mathcal{R}_{\text{OT}}(\mathbf{x}, \hat{\mathbf{x}})$. This is a key difference from Chen et al. (2019), where teacher-forcing is used in OT $\mathcal{R}_{\text{OT}}(\mathbf{x}, \tilde{\mathbf{x}})$. To note the difference, we refer to the method in Chen et al. (2019) as TFOT.

The exact minimization over \mathbf{M} is generally computationally intractable (Arjovsky et al., 2017; Genevay et al., 2018). Hence we use the recently introduced Inexact Proximal point method for Optimal Transport (IPOT) (Xie et al., 2018) algorithm to approximate \mathbf{M}^* . The details of the IPOT algorithm are shown in Appendix A.1.

2.3 Cost Functions in SFOT

OT-matching quality largely depends on the cost function $c(\cdot, \cdot)$. In particular, there is flexibility in how we represent the elements to be transported, for which we outline two alternatives below.

Vanilla OT A natural choice for the cost distance is to use the word embeddings, denoted $\{\mathbf{h}_i^0\}_{i=1}^T$, as used by previous works:

$$c(\mathbf{x}_i, \hat{\mathbf{x}}_j) = 1 - \frac{\mathbf{h}_i^0 \top \mathbf{h}_j^0}{\|\mathbf{h}_i^0\| \|\mathbf{h}_j^0\|}, \quad (8)$$

However, a word-embedding-based cost function only captures the token-to-token similarity. On the other hand, the semantics of words can be different in different positions or contexts. This inspires the proposal of two novel cost functions to improve text sequence matching in OT.




Figure 2: Illustration of the intuition of the proposed OT extensions. The goal is to match the sequence ABACA to the sequence AABAC. Left: traditional OT, where only the token information is considered; thus letter A in any positions of the two sequences can be aligned. Right: the proposed extensions of OT, where the *context* and *ordering* information help to eliminate undesirable alignment. In terms of the context, letter A will be matched to those that share a similar context $\{A, B\}$ and $\{A, C\}$. In terms of ordering, letter A will be aligned to the letters with similar positions

Contextualized OT with Order-Preserving Regularizer Because the same word in different linguistic contexts may have different meanings, a cost function that cannot capture such variability may lead to undesirable matching results. While word embeddings $\{\mathbf{h}_t^0\}_{t=1}^T$ may be myopic, hidden representations $\{\mathbf{h}_t^\ell\}_{t=1}^T$ at higher layers ($\ell > 0$) of deep language models (e.g. LSTM (Hochreiter and Schmidhuber, 1997) or the Transformer (Vaswani et al., 2017)) often capture contextualized representations of the word in the sequence. Inspired by works on deep contextualized word representations (Peters et al., 2018; Devlin et al., 2018; Radford et al., 2018), we replace the word embeddings with $\{\mathbf{h}_t^\ell\}_{t=1}^T$ to represent the meaning of words inside the sequence. Then the cost function can be defined

$$c(\mathbf{x}_i, \hat{\mathbf{x}}_j) = 1 - \frac{\mathbf{h}_i^{\ell \top} \mathbf{h}_j^\ell}{\|\mathbf{h}_i^\ell\| \|\mathbf{h}_j^\ell\|} \quad (9)$$

Figure 2 shows that the meaning of letter $\{A\}$ can be less ambiguous when considering its local context $\{B, A\}$ and $\{C, A\}$. This context information further helps eliminate the undesirable matching configurations. Otherwise, letter A can match to A of any position in the other sequence. Note that contextual information may implicitly capture relatively long-term dependency information compared with vanilla OT, and it can be perceived as a “soft” n -gram matching.

We also consider an order-preserving regularizer for the contextualized OT, motivated by the

Algorithm 1 Student Forcing Optimal Transport

- 1: **Input:** Ground truth paired sequences $\mathcal{D} = \{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$
 - 2: Initialize MLE model parameters θ
 - 3: **while** training **do**
 - 4: Draw samples $\{\mathbf{x}_n, \mathbf{y}_n\} \sim \mathcal{D}$
 - 5: Compute the outputs \mathbf{h} of the model via teacher-forcing and $\hat{\mathbf{h}}$ via student-forcing
 - 6: Compute the cost matrix \mathbf{C} based on the choice of cost function in (8), (10)
 - 7: Compute the OT Loss $\mathcal{R}_{\text{OT}}(\mathbf{x}, \hat{\mathbf{x}})$ defined in (6) via IPOT algorithm
 - 8: Update θ by optimizing $\mathcal{L}_{\text{SFOT}}$ defined in (7)
 - 9: **end while**
-

fact that positional information of a token can be crucial in natural language understanding. For example, two sentences may have opposite meanings when the word order is changed: “*He hated it but then started loving it*” versus “*He loved it but then started hating it*”. Hence, it is desirable to have the transportation matrix concentrate to diagonal entries, by transporting the neighboring elements in one sequence into some other neighboring elements in another sequence with nearby temporal position (Su and Hua, 2018). Inspired by Albreghsen et al. (2008), we penalize the contextual cost function with *inverse difference moment* as

$$c_c(\mathbf{x}_i, \hat{\mathbf{x}}_j) = c(\mathbf{x}_i, \hat{\mathbf{x}}_j) - \frac{\beta}{\left(\frac{i}{T} - \frac{j}{T}\right)^2 + 1} \quad (10)$$

where $\beta \geq 0$ is the weighting hyper-parameter for the order-preserving penalty. Figure 2 shows that when only considering the token-to-token similarity, the first letter A from the left sequence can be moved to any A in the right sequence. However, when considering the ordering penalty, letter A is aligned to the letters with similar position. SFOT is summarized in Algorithm 1.

Note that order-preserving regularization is only applied to contextualized OT, as requiring position-wise matching for the generated and target sentence may be too restrictive. Empirically, we observed adding the order-preserving regularizer to vanilla OT gives marginal improvement. However, since contextualized OT operates on a feature space, the position-wise matching is softened and can be naturally coupled with contextual cost function.

2.4 An Imitation Learning Interpretation of Student-Forcing

The sequential text generation process can be reformulated using Markov decision processes

(MDPs). Formally, an MDP $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P_s, r \rangle$ ¹ defines a transition probability $P_s(s'|s, \mathbf{a})$ from state $s \in \mathcal{S}$ to the next state $s' \in \mathcal{S}$ after the agent takes an action $\mathbf{a} \in \mathcal{A}$; $r(s, \mathbf{a})$ is an unknown reward function. To cast text generation as an MDP, we may consider the action as the selection of the next word x_t from the vocabulary, conditioned on the states of observed words $x_{<t}$, *i.e.*, $\mathbf{s}_t = x_{<t}$ and $\mathbf{a}_t = x_t$. At each time step t , the agent takes an action \mathbf{a}_t at state \mathbf{s}_t according to policy $\pi(\mathbf{a}_t|\mathbf{s}_t)$, and receives a reward $r(\mathbf{s}_t, \mathbf{a}_t)$. The conjunction of the generated word and the previous methods constitute the next state \mathbf{s}_{t+1} .

Imitation learning seeks to learn an optimal policy from demonstrations of an expert policy π_E . In language models, the training text plays the role of expert trajectories. This objective is formally:

$$\max_{r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}} \min_{\pi \in \Pi} -H(\pi) + \mathcal{R}(\pi_E) - \mathcal{R}(\pi) \quad (11)$$

with $\mathcal{R}(\pi) = \sum_{\mathbf{s}, \mathbf{a}} \rho_\pi(\mathbf{s}, \mathbf{a}) r(\mathbf{s}, \mathbf{a})$.

where $\rho_\pi(\mathbf{s}, \mathbf{a})$ is the stationary joint distribution of (\mathbf{s}, \mathbf{a}) induced by the learning policy π ; $H(\pi)$ is the entropy. Intuitively, the objective encourages that higher rewards are assigned to the expert policy π_E , while π is trained to mimic π_E .

Importantly, (11) suggests that each individual word of sentences that induces the distribution ρ_π should be fully generated by the learned language model π . In other words, at each time step, π generates a word using its own previously generated words. This is exactly the student-forcing scheme we employ for the proposed SFOT algorithm. This reveals the key difference with Chen et al. (2019), where teaching-forcing is employed. Since such an agent takes actions based on the partial expert trajectories, it will induce a biased occupancy measure. This results in a sub-optimal policy for text generation, even when the imitation learning objective (11) is optimized.

3 Related Work

Text Generation Natural Language Generation (NLG) is a challenging NLP task. Neural language models parameterized by autoregressive architectures are widely used for NLG. To improve the global control ability of generated sentences, variational auto-encoders are considered for language generation (Bowman et al., 2016; Fu et al.,

2019; Fang et al., 2019; Li et al., 2020a). Recently, GPT-2 (Radford et al., 2019) and GPT-3 (Brown et al., 2020) improve the generation fluency via pre-training on massive text corpus. All of them are trained with MLE using Teacher-Forcing, which are known to suffer from *exposure bias* in principle (Bengio et al., 2015). Several methods have been proposed to solve the problem, including (Shao et al., 2018; Zhang et al., 2019). Adversarial training techniques were also proposed (Yu et al., 2017; Zhu et al., 2018; Che et al., 2017; Lin et al., 2017; Guo et al., 2018; Chen et al., 2018; Li et al., 2020b; Yang et al., 2019; Zhang et al., 2018; Liang et al., 2018). However, adversarial-based NLG models can suffer from gradient vanishing and unstable training. Indeed, (Caccia et al., 2018) argues that a temperature sweeping approach on MLE can outperform GAN-based models. Our model further improves this work by adopting a principled sequence-matching loss via optimal transport and achieve state-of-the-art results on NLG tasks.

Optimal Transport Optimal transport is widely employed for a variety of NLP tasks, including document classification (Kusner et al., 2015), word embedding space alignment (Alvarez-Melis and Jaakkola, 2018), and generative adversarial networks (Chen et al., 2018). The most related work to ours is TFOT Chen et al. (2019). We discuss the difference between the proposed SFOT and TFOT as follows:

- *Strong empirical evidence on long sentences.* While the overall performance of SFOT is superior to TFOT only by a decent margin on standard datasets, we emphasize that the main advantage of SFOT is for long sentences. As shown in the break-down analysis (Figure 4), SFOT is significantly better than TFOT when sentence length is larger than 60. This highlights the critical contribution of this paper to addressing the exposure bias problem, which longer sentences usually suffer from.
- *Methodology novelty.* Besides the difference in SF decoding and TF decoding in two methods, we also propose a technique on ‘‘Contextualized OT with Order-Preserving Regularizer’’, which improves both SFOT and TFOT, as shown in Table 4. Note that, there is no order information used in TFOT, which de-

¹The discount factor is set as one for simplicity

generates it to word embedding alignment instead of sequence matching. In contrast, the proposed technique can utilize sequence level, thus improving the performance.

- *Theoretical difference.* In Section 2.3, we provide theoretical justification on why SFOT can reduce exposure bias, while TFOT still suffers from it: TFOT is based on partial expert trajectories and induces a bias occupancy measure, while our proposed method SFOT uses previous self-generated words and can obtain an optimal policy.

Sequence matching Direct sequence matching has been explored widely in various machine learning tasks. Jaccard distance has been used to retrieve prototypes for sentence generation (Guu et al., 2018). Chernoff distance has been applied to image classification (Su et al., 2015). However, these distances assume each instance in the sequence is independent, ignoring temporal information. These distances consequently measure sequence alignment poorly, as they miss semantic relationships inside a sentence (e.g. cause and effect). Dynamic time warping (DTW) (Sakoe et al., 1990) and Connectionist Temporal Classification (CTC) (Graves and Jaitly, 2014) consider temporal information, and have been adopted widely in speech recognition. However, these losses preserve strict ordered alignment, and hence cannot be directly applied to text sequences.

4 Experiments

We perform experiments on neural machine translation (NMT), abstractive text summarization and unconditional natural language generation (NLG) tasks. Algorithms are implemented in Tensorflow and trained on an NVIDIA TITAN X GPU.

4.1 Neural Machine Translation

Two standard datasets are tested for NMT tasks: a small-scale English-Vietnamese corpus from the IWSLT 2015 Evaluation Campaign (Cettolo et al., 2015) and a large-scale English-German corpus from the WMT16 Evaluate Campaign². Further details of the datasets and the experimental setup are shown in Appendix A.2.

We compare SFOT with a variety of methods: MLE (Luong et al., 2017), Scheudle Sampling (SS), TFOT and several RL-based models,

²<http://statmt.org/wmt16>

Task	Algorithm	NT2012	NT2013
VI-EN	MLE	21.8	24.5
	SS	21.8	24.6
	RAML	22.0	25.0
	MIXER	21.9	24.7
	SPG	22.0	25.1
	TFOT	21.9	25.5
	SFOT	22.3	25.8
EN-VI	MLE	23.8	26.1
	SS	23.9	26.2
	RAML	24.2	26.6
	MIXER	24.0	26.3
	SPG	24.3	26.7
	TFOT	24.5	26.9
	SFOT	24.9	27.4

Table 1: VI-EN and EN-VI translation BLEU scores.




Figure 3: Convergence on NT2012 EN-VI.

i.e. RAML (Norouzi et al., 2016), SPG (Ding and Soricut, 2017), and MIXER (Ranzato et al., 2016b). The results are summarized in Tables 1 and 2. The proposed SFOT approach consistently improves upon MLE training and outperforms other models in all experimental setups. Besides the quantitative results, we observe that SFOT correctly maintains the information from the source side to make correct translations (Table 3). As can be seen from these examples, our model can better preserve information from the source and is less likely to transfer words incorrectly. Notice that most errors in the baseline models occur in the latter part of sequences, due to error accumulation from *exposure-bias*, which SFOT addresses by matching the free-running outputs to the ground-truth. In conjunction with the quantitative results presented above, these qualitative observations confirm that our model can generate more reliable translation for long sentences and address the *exposure-bias* problem.

We further investigate the performance on the English-Vietnamese dataset. We first check the convergence of the BLEU score on the validation set, shown in Figure 3. SS converges similarly to the MLE model and only shows marginal improvement. The models with OT converge much

Task	Algorithm	NT2013	NT2015
DE-EN	MLE	29.0	29.9
	SS	29.0	29.9
	RAML	29.1	30.1
	MIXER	29.0	30.0
	SPG	29.1	30.1
	TFOT	29.2	30.1
	SFOT	29.3	30.3
EN-DE	MLE	24.3	26.5
	SS	24.3	26.5
	RAML	24.5	26.7
	MIXER	24.4	26.6
	SPG	24.5	26.7
	TFOT	24.6	26.8
	SFOT	24.8	27.0

Table 2: DE-EN and EN-DE translation BLEU scores.




Figure 4: Translation quality as sentences become longer on NT2013 EN-VI.

faster in the early stages and finish with a higher final performance. However, TFOT has an unstable convergence trajectory, as it can degrade MLE training performance. SFOT by contrast is consistently better than MLE training, achieving the best BLEU score among the three models.

To further demonstrate SFOT’s ability to address *exposure bias*, we follow Bahdanau et al. (2015) and group sentences of similar lengths together, computing a BLEU score per group on the test set. As errors accumulate in generation, longer sentences suffer more from *exposure bias* and have lower quality. Figure 4 shows that our model is more effective in handling long sentences: compared with other methods, SFOT is more robust for longer sentence lengths, indicating that matching the generated and ground-truth sentence in the sequence level may alleviate *exposure bias* for long-text generation.

Additionally, we compare different OT cost function variants in Table 4. We denote contextualized OT with order-preserving regularizer as -c. Contextualized OT improves translation quality when used with TFOT and SFOT. The improvements on both models indicate that contextualized representations with order preservation is capable

of capturing more sentence semantic information.

4.2 Abstractive Text Summarization

We use a widely considered English Gigawords corpus (Graff et al., 2003) for the text-summarization task. Similar to NMT experiments, we use MLE as our baseline model and further compare SFOT with SS, TFOT, and several RL-based methods, *i.e.* RAML, SPG and MIXER. We evaluate the model performance using ROUGE (including -1, -2, -L) score (Lin, 2004), the most popular metric for summarization. Details of the datasets and the experimental setup are shown in Appendix A.3.

Summarization results are provided in Table 5. Consistent with our NMT results, SFOT outperforms all other methods, showing that the contextualized matching is capable of capturing semantic information essential for high-quality generation. Moreover, the superiority of SFOT over RL-based models demonstrates that the OT (sentence) matching is more robust than word/phrase matching in RL rewards.

4.3 Neural Language Generation

Following recent unconditional long text generation work (Caccia et al., 2018), we perform experiments on EMNLP2017 WMT News dataset³. All sentences are longer than 20, making the dataset appropriate for testing the *exposure bias* problem.

To evaluate the effectiveness of our model, we consider various baseline methods, including recent GAN-based text generation approaches, such as SeqGAN (Ramachandran et al., 2017), RankGAN (Lin et al., 2017), MaliGAN (Che et al., 2017), and LeakGAN (Guo et al., 2018), as well as an MLE-trained model using temperature sweep (Caccia et al., 2018). We apply SFOT with contextualized OT to improve the temperature-sweep MLE model. SS does not show significantly different results compared to the MLE model. For the evaluation metric, we follow the current protocol for NLG evaluation (Zhu et al., 2018) w.r.t. both quality and diversity. Specifically, the quality of the generation is measured with BLEU score (Papineni et al., 2002) and the diversity is evaluated with Self-BLEU (Zhu et al., 2018). Human evaluation is further considered to measure the quality of generation⁴. More details

³<http://www.statmt.org/wmt17/>

⁴We perform human evaluation using Amazon Mechanical Turk. 100 generated sentences are sampled from each

Reference:	India’s new prime minister, Narendra Modi, is meeting his Japanese counterpart, Shinzo Abe, in Tokyo to discuss economic and security ties, on his first major foreign visit since winning May’s election .
MLE:	India’s new prime minister, Narendra Modi, meets his Japanese counterpart, Shinzo Abe, in Tokyo, during his first major foreign visit in May to discuss economic and security relations.
TFOT:	India’s new prime minister Narendra Modi meets his Japanese counterpart, Shinzo Abe, in Tokyo at his first major foreign visit since his election in May in order to discuss economic and security relations.
SFOT:	India’s new prime minister, Narendra Modi, is meeting his Japanese counterpart Shinzo Abe in Tokyo in his first major foreign visit since his election victory in May to discuss economic and security relations.

Reference:	Chinese leaders presented the Sunday ruling as a democratic breakthrough because it gives Hong Kongers a direct vote, but the decision also makes clear that Chinese leaders would retain a firm hold on the process through a nominating committee tightly controlled by Beijing .
MLE:	The Chinese leadership presented the decision of Sunday as a democratic breakthrough, because it gives Hong Kong citizens a direct right to vote, but the decision also makes it clear that the Chinese leadership maintains the expiration of a nomination committee closely controlled by Beijing .
TFOT:	The Chinese leadership presented Sunday’s decision as a democratic breakthrough because it gives the citizens of Hong Kong a direct right to vote, but the decision also makes it clear that the Chinese leadership keeps the process firmly in the hands of a government -controlled Nomination Committee.
SFOT:	The Chinese leadership presented the decision on Sunday as a democratic breakthrough, because Hong Kong citizens have a direct electoral right, but the decision also makes it clear that the Chinese leadership remains firmly in hand with a nominating committee controlled by Beijing .

Table 3: Comparison of German-to-English translation examples. For each example, we show the human translation (reference) and the translation from MLE, TFOT, and SFOT. We highlight the key phrase differences between reference and translation outputs in blue and red, and annotate translation errors in bold. In the first example, SFOT correctly maintains all the information in “*since winning in May election*” by translating to “*since his election victory in May*”, whereas MLE only generates “*in May*” and TFOT also misses “*winning*” in the reference. In the second example, SFOT successfully keeps the information “*Beijing*”, whereas MLE generates wrong words “*expiration of*” and TFOT changes “*Beijing*” to “*government*”.

Task	Algorithm	NT2012	NT2013
VI-EN	MLE	21.8	24.5
	TFOT	22.2	25.1
	TFOT-c	22.2	25.4
	SFOT	22.2	25.6
	SFOT-c	22.3	25.8
EN-VI	MLE	23.8	26.1
	TFOT	24.5	26.9
	TFOT-c	24.8	27.2
	SFOT	24.8	27.4
	SFOT-c	24.9	27.4

Table 4: BLEU scores for VI-EN and EN-VI ablation study.

Method	ROUGE-1	ROUGE-2	ROUGE-L
MLE	36.1	16.4	32.3
SS	36.6	16.8	32.7
RAML	36.3	16.7	32.5
SPG	36.5	16.8	32.8
MIXER	36.3	16.6	32.6
TFOT	36.8	17.2	33.5
SFOT	37.0	17.5	33.9

Table 5: Results of text summarization on English Gigawords dataset.

of experiment setup are shown in Appendix A.4.

To reasonably select the best model along the temperature sweep, we are motivated by (Gu et al., 2019) and propose the BLEU-F1 score to evaluate model. Ten native speakers are asked to rate each sentence in the scale 1 to 5 in terms of readability and meaningfulness

the trade-off between the quality and diversity simultaneously, defined as

$$\text{BLEU-F1} = \frac{2 \times \text{BLEU} \times (1 - \text{Self-BLEU})}{\text{BLEU} + (1 - \text{Self-BLEU})}. \quad (12)$$

Figure 5 shows the BLEU-F1 score versus reverse temperature on MLE and SFOT. We observed that the best temperature for MLE model is 1/1.5 and for SFOT is 1/1.4. We further conduct analysis under these temperatures. Figure 5 also indicates that the SFOT model consistently improves the MLE model on the BLEU-F1 score.

We compare SFOT with the proposed strong baselines in Figure 6 and report human evaluation of generated quality in Table 6. We observe that SFOT has the highest generation quality in human evaluation. With better guided sequence-level semantics information, SFOT generates high-quality sentences at higher temperatures, compared with the MLE model. The MLE model decreases temperature to concentrate on generating *safe* words (with high probability); this avoids error accumulation by avoiding risky words. However, the model loses generation diversity when increasing temperature, as the model only focuses on *safe* words. SFOT obtains better quality at higher temperature, indicating SFOT can generate reasonable sequences on more risky words and hence

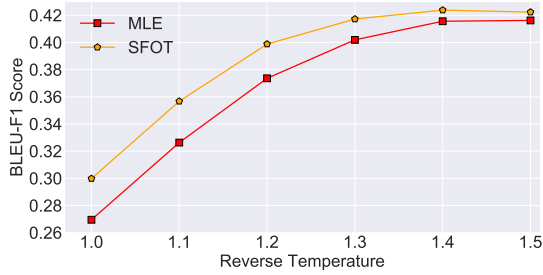


Figure 5: BLEU-5 F1 score plotted against reverse temperature α on EMNLP2017 News test set.

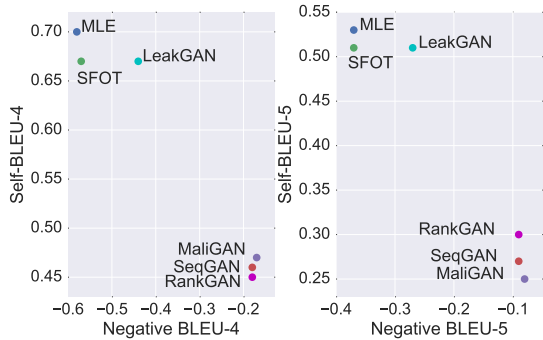


Figure 6: Negative BLEU (lower is better) plotted against Self-BLEU (lower is better) for EMNLP2017 News test set, for BLEU-4 (left) and BLEU-5 (right).

Method	SeqGAN	RankGAN	MaliGAN
Score	2.56 ± 0.49	2.89 ± 0.56	2.50 ± 0.46
Method	LeakGAN	MLE	SFOT
Score	3.45 ± 0.47	3.46 ± 0.42	3.54 ± 0.37

Table 6: Human evaluation of NLG on EMNLP news 2017 dataset. 100 generated sentences from each model are rated 1-5, with means and standard deviations reported. Real sentences were rated 4.21 ± 0.44 .

can address *exposure bias* and gain more diversity in generation. We also observe that SFOT outperforms all text GANs in terms of quality-diversity trade-off and human evaluation. Under similar Self-BLEU score, SFOT significantly improves the quality of LeakGAN (Guo et al., 2018), the best GAN by BLEU metric.

5 Conclusions

We have introduced SFOT to mitigate exposure bias in text generation. The proposed model captures positional and contextual information of word tokens in OT matching. Experiments on neural machine translation, text summarization, and text generation have demonstrated the effectiveness of our SFOT algorithm, yielding improved performance over strong baselines on these tasks.

References

- Fritz Albregtsen et al. 2008. Statistical texture measures computed from gray level cooccurrence matrices. *Image processing laboratory, department of informatics, university of oslo*, 5.
- David Alvarez-Melis and Tommi S Jaakkola. 2018. Gromov-wasserstein alignment of word embedding spaces. In *EMNLP*.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *ICML*.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. An actor-critic algorithm for sequence prediction. In *ICLR*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. *CONLL*.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Massimo Caccia, Lucas Caccia, William Fedus, Hugo Larochelle, Joelle Pineau, and Laurent Charlin. 2018. Language gans falling short. *arXiv preprint arXiv:1811.02549*.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2015. The iwslt 2015 evaluation campaign. In *IWSLT*.
- Tong Che, Yanran Li, Ruixiang Zhang, R Devon Hjelm, Wenjie Li, Yangqiu Song, and Yoshua Bengio. 2017. Maximum-likelihood augmented discrete generative adversarial networks. In *CoRR*.
- Liqun Chen, Shuyang Dai, Chenyang Tao, Haichao Zhang, Zhe Gan, Dinghan Shen, Yizhe Zhang, Guoyin Wang, Ruiyi Zhang, and Lawrence Carin. 2018. Adversarial text generation via feature-mover’s distance. In *NIPS*.
- Liqun Chen, Yizhe Zhang, Ruiyi Zhang, Chenyang Tao, Zhe Gan, Haichao Zhang, Bai Li, Dinghan Shen, Changyou Chen, and Lawrence Carin. 2019. Improving sequence-to-sequence learning via optimal transport. In *ICLR*.

- 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*.
- Nan Ding and Radu Soricut. 2017. Cold-start reinforcement learning with softmax policy gradient. In *Advances in Neural Information Processing Systems*, pages 2817–2826.
- Le Fang, Chunyuan Li, Jianfeng Gao, Wen Dong, and Changyou Chen. 2019. Implicit deep latent variable models for text generation. *EMNLP*.
- Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, Lawrence Carin, et al. 2019. Cyclical annealing schedule: A simple approach to mitigating KL vanishing. *NAACL*.
- Aude Genevay, Gabriel Peyré, and Marco Cuturi. 2018. Learning generative models with sinkhorn divergences. *AISTATS*.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.
- Alex Graves and Navdeep Jaitly. 2014. Towards end-to-end speech recognition with recurrent neural networks. In *ICML*.
- Xiaodong Gu, Kyunghyun Cho, Jungwoo Ha, and Sunghun Kim. 2019. Dialogwae: Multimodal response generation with conditional wasserstein auto-encoder. *ICLR*.
- Jiaxian Guo, Sidi Lu, Han Cai, Weinan Zhang, Yong Yu, and Jun Wang. 2018. Long text generation via adversarial training with leaked information. In *AAAI*.
- Kelvin Guu, Tatsunori B Hashimoto, Yonatan Oren, and Percy Liang. 2018. Generating sentences by editing prototypes. *TACL*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*.
- Ferenc Huszár. 2015. How (not) to train your generative model: Scheduled sampling, likelihood, adversary? In *CoRR*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. OpenNMT: Open-source toolkit for neural machine translation. In *ACL*.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *ICML*.
- Alex M Lamb, Anirudh Goyal, Parth Goyal, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. 2016. Professor forcing: A new algorithm for training recurrent networks. In *NIPS*.
- Chunyu Li, Xiang Gao, Yuan Li, Xiujun Li, Baolin Peng, Yizhe Zhang, and Jianfeng Gao. 2020a. Optimus: Organizing sentences via pre-trained modeling of a latent space. *arXiv preprint arXiv:2004.04092*.
- Dianqi Li, Yizhe Zhang, Hao Peng, Liqun Chen, Chris Brockett, Ming-Ting Sun, and Bill Dolan. 2020b. Contextualized perturbation for textual adversarial attack. *arXiv preprint arXiv:2009.07502*.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial learning for neural dialogue generation. In *EMNLP*.
- Kevin J Liang, Chunyuan Li, Guoyin Wang, and Lawrence Carin. 2018. Generative adversarial network training is a continual learning problem. *arXiv preprint arXiv:1811.11083*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.
- Kevin Lin, Dianqi Li, Xiaodong He, Zhengyou Zhang, and Ming-Ting Sun. 2017. Adversarial ranking for language generation. In *NIPS*.
- Hao Liu, Yihao Feng, Yi Mao, Dengyong Zhou, Jian Peng, and Qiang Liu. 2018. Action-dependent control variates for policy optimization via stein’s identity. In *ICLR*.
- Giulia Luise, Alessandro Rudi, Massimiliano Pontil, and Carlo Ciliberto. 2018. Differential properties of sinkhorn approximation for learning with wasserstein distance. *arXiv:1805.11897*.
- Minh-Thang Luong, Eugene Brevdo, and Rui Zhao. 2017. Neural machine translation (seq2seq) tutorial. <https://github.com/tensorflow/nmt>.
- Minh-Thang Luong and Christopher D Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *IWSLT*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015a. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- Minh-Thang Luong, Ilya Sutskever, Quoc V Le, Oriol Vinyals, and Wojciech Zaremba. 2015b. Addressing the rare word problem in neural machine translation. In *ACL*.
- Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *ISCA*.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human generated machine reading comprehension dataset. In *NIPS*.
- Mohammad Norouzi, Samy Bengio, Navdeep Jaitly, Mike Schuster, Yonghui Wu, Dale Schuurmans, et al. 2016. Reward augmented maximum likelihood for neural structured prediction. In *Advances*

- In *Neural Information Processing Systems*, pages 1723–1731.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/languageunsupervised/language_understanding_paper.pdf.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*.
- Prajit Ramachandran, Peter J Liu, and Quoc V Le. 2017. Unsupervised pretraining for sequence to sequence learning. In *EMNLP*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016a. Sequence level training with recurrent neural networks. *CoRR*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016b. Sequence level training with recurrent neural networks. In *ICLR*.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *CVPR*.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685*.
- Hiroaki Sakoe, Seibi Chiba, A Waibel, and KF Lee. 1990. Dynamic programming algorithm optimization for spoken word recognition. *Readings in speech recognition*.
- Ruslan Salakhutdinov. 2015. Learning deep generative models. *Annual Review of Statistics and Its Application*, 2:361–385.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *ACL*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. In *ACL*.
- Chenze Shao, Yang Feng, and Xilin Chen. 2018. Greedy search with probabilistic n-gram matching for neural machine translation. *arXiv preprint arXiv:1809.03132*.
- Bing Su, Xiaoqing Ding, Changsong Liu, and Ying Wu. 2015. Heteroscedastic max-min distance analysis. In *CVPR*.
- Bing Su and Gang Hua. 2018. Order-preserving optimal transport for distances between sequences. *IEEE transactions on pattern analysis and machine intelligence*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- 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 *NIPS*.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. In *ICML workshop*.
- Xin Wang, Wenhua Chen, Yuan-Fang Wang, and William Yang Wang. 2018. No metrics are perfect: Adversarial reward learning for visual storytelling. In *ACL*.
- Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Yujia Xie, Xiangfeng Wang, Ruijia Wang, and Hongyuan Zha. 2018. A fast proximal point method for Wasserstein distance. In *arXiv:1802.04307*.
- Qian Yang, Dinghan Shen, Yong Cheng, Wenlin Wang, Guoyin Wang, Lawrence Carin, et al. 2019. An end-to-end generative architecture 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 3123–3133.
- Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image captioning with semantic attention. In *CVPR*.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*.

Ruiyi Zhang, Changyou Chen, Zhe Gan, Wenlin Wang, Liqun Chen, Dinghan Shen, Guoyin Wang, and Lawrence Carin. 2018. Sequence generation with guider network. *arXiv preprint arXiv:1811.00696*.

Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. Bridging the gap between training and inference for neural machine translation. *arXiv preprint arXiv:1906.02448*.

Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ricardo Henao, Dinghan Shen, and Lawrence Carin. 2017. Adversarial feature matching for text generation. In *ICML*.

Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Texygen: a benchmarking platform for text generation models. In *SIGIR*.

A Appendix

A.1 Inexact Proximal point method for Optimal Transport (IPOT) Algorithm

The IPOT algorithm to approximate M^* is shown in Algorithm 2

Algorithm 2 IPOT algorithm

```

1: Input: two sequence  $\mathbf{x}, \hat{\mathbf{x}}$  with length  $T$ , and generalized
   step size  $1/\varepsilon, \sigma$ 
2:  $\boldsymbol{\sigma} = \frac{1}{T^r} \mathbf{1}_T, M^{(1)} = \mathbf{1}_T \mathbf{1}_T^\top,$ 
3:  $C_{ij} = c(\mathbf{x}_i, \hat{\mathbf{x}}_j), A_{ij} = e^{-\frac{C_{ij}}{\varepsilon}}$ 
4: for  $t = 1, 2, 3 \dots$  do
5:    $\mathbf{Q} = \mathbf{A} \odot \mathbf{T}^{(t)}$  //  $\odot$  is Hadamard product
6:   for  $k = 1, \dots, K$  do //  $K = 1$  in practice
7:      $\boldsymbol{\delta} = \frac{1}{T\mathbf{Q}\boldsymbol{\sigma}}, \boldsymbol{\sigma} = \frac{1}{T\mathbf{Q}^\top\boldsymbol{\delta}}$ 
8:   end for
9:    $M^{(t+1)} = \text{diag}(\boldsymbol{\delta})\mathbf{Q}\text{diag}(\boldsymbol{\sigma})$ 
10: end for
11: Return  $(M, C)$ 

```

A.2 Neural Machine Translation Experiments

Dataset Two standard datasets are tested for NMT tasks. The first one is a small-scale English-Vietnamese corpus from the IWSLT 2015 Evaluation Campaign (Cettolo et al., 2015), which is a parallel corpus of TED-talks and contains 133K sentence pairs. We follow the pre-processing procedure in (Luong and Manning, 2015) by replacing words with frequencies less than 5 with $\langle \text{unk} \rangle$. As a result, our vocabulary reduces to 17K for English and 7.7K for Vietnamese. We use TED tst2012 as development set and TED tst2013 as the test set. For a large-scale dataset, we select an

English-German corpus from the WMT16 Evaluate Campaign⁵, which contains 4.5M sentence pairs. Newstest 2013 is used as the development set and Newstest 2015 is used as the test set. We conduct the sub-word tokenization on the corpus using the Byte Pair Encoding (BPE) method (Sennrich et al., 2015). Following Klein et al. (2017), we set the vocabulary size of both English and German to 32K.

Setup We use Google’s Neural Machine Translation (GNMT) system (Wu et al., 2016) as our baseline MLE model, which follows the standard architecture and hyper-parameters⁶ for fair comparison. All other models are built on top of with same network structure. We evaluate the model performance using BLEU scores (Papineni et al., 2002). We set OT weighting parameter $\lambda = 0.1$ and order-preserving penalty weighting parameter $\beta = 0.1$.

For English-Vietnamese translation tasks (*i.e.*, EN-VI or VI-EN), we follow the setup in (Sutskever et al., 2014; Luong et al., 2015b,a). We use one bidirectional LSTM layer with 512 hidden units as encoder and two-layer LSTM with 512 hidden units at each layer as decoder. The embedding dimension is set as 512. We follow the attention method described in (Luong et al., 2015a) and use dropout with probability 0.2 as suggested by (Zaremba et al., 2014). All parameters are initialized uniformly between $[-0.1, 0.1]$. We train the model for 12 epochs with 12 epochs using Stochastic Gradient Decent (SGD). For the first 8 epochs, we set learning rate as 1.0. After that, we anneal the learning rate at half at every epoch.

For English-German translation tasks (*i.e.*, EN-GE or GE-EN), we adopt a stacked LSTM with a 2-layer bidirectional of 1024 units as encoder and 4-layer LSTM with units 1024 as decoder. The embedding dimension is set to 1024. We adopt the attention used in (Wu et al., 2016). We train the model for 10 epochs. For the first 5 epochs, we set the learning rate as 1 and then halving the learning rate every half epoch.

A.3 Abstractive Text Summarization Experiments

We use a widely accepted English Gigawords corpus (Graff et al., 2003) for the text summariza-

⁵<http://statmt.org/wmt16>

⁶<https://github.com/tensorflow/nmt>

tion task. We follow the pre-process in (Rush et al., 2015). The dataset is sampled and split into train/dev/test set with size 200K/8K/2K.

A.4 Natural Language Generation Experiment

In the NLG experiment, 200K sentences are collected as the training set and 10K sentences as the test set. In the NLG experiment, we set OT weighting parameter $\lambda = 1$ and the order-preserving penalty weighting parameter is $\beta = 0.1$. Since input sequence \mathbf{y} is empty in a language model, to better guide the student forcing output, we adopt schedule sampling with ratio 0.3 in our experiments. The samples generated by SFOT are presented in Table 7 .

So , this is a great way , but I'm not sure how to do that , he said .

When made an emergency landing , the driver who was also injured in the blast was arrested on suspicion of causing death .

The result is that the company's economic growth rate is rising by a substantial margin in November, which is even higher than a year ago .

It's really a big deal for us and we're going to get ready for the second game .

We feel like it's very hard to say that it's really going to be the next generation .

You don't want to be a kid , and there are a lot of things that you can do .

The government's decision to extend its coal policy vote will be announced in the first half of 2017 .

I'm not able to do it , but I think it's pretty important for him to be the best player .

But I'm not sure what the supporters can do in this election , he said , referring to the Sanders campaign .

Table 7: Examples generated by SFOT in NLG experiments