

Generative Bridging Network for Neural Sequence Prediction

Wenhu Chen,¹ Guanlin Li,² Shuo Ren,⁵ Shujie Liu,³ Zhirui Zhang,⁴ Mu Li,³ Ming Zhou³

University of California, Santa Barbara¹

Harbin Institute of Technology²

Microsoft Research Asia³

University of Science and Technology of China⁴

Beijing University of Aeronautics and Astronautics⁵

wenhuchen@cs.ucsb.edu epsilonlee.green@gmail.com {v-shure, shujliu, v-zhirzh, muli, mingzhou}@microsoft.com

Abstract

In order to alleviate data sparsity and overfitting problems in maximum likelihood estimation (MLE) for sequence prediction tasks, we propose the Generative Bridging Network (GBN), in which a novel bridge module is introduced to assist the training of the sequence prediction model (the generator network). Unlike MLE directly maximizing the conditional likelihood, the bridge extends the point-wise ground truth to a bridge distribution conditioned on it, and the generator is optimized to minimize their KL-divergence. Three different GBNs, namely uniform GBN, language-model GBN and coaching GBN, are proposed to penalize confidence, enhance language smoothness and relieve learning burden. Experiments conducted on two recognized sequence prediction tasks (machine translation and abstractive text summarization) show that our proposed GBNs can yield significant improvements over strong baselines. Furthermore, by analyzing samples drawn from different bridges, expected influences on the generator are verified.

1 Introduction

Sequence prediction has been widely used in tasks where the outputs are sequentially structured and mutually dependent. Recently, massive explorations in this area have been made to solve practical problems, such as machine translation (Bahdanau et al., 2014; Ma et al., 2017; Norouzi et al., 2016), syntactic parsing (Vinyals et al., 2015), spelling correction (Bahdanau et al., 2014), image captioning (Xu et al., 2015) and speech recognition (Chorowski et al., 2015). Armed with modern computation power, deep LSTM (Hochreiter and Schmidhuber, 1997) or GRU (Chung et al., 2014) based neural sequence prediction models have achieved the state-of-the-art performance.

The typical training algorithm for sequence prediction is Maximum Likelihood Estimation

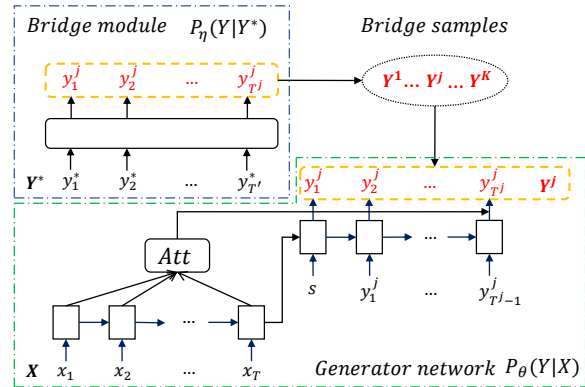


Figure 1: The overall architecture of our novel Generative Bridging Network (GBN). Two main components, namely the generator network and the bridge module, are connected through samples ($Y^1 \dots Y^K$ in red) from the bridge module during training time. (We sometimes call them generator and bridge in brief respectively in the following discussion.) The generator is implemented through an attentive encoder-decoder, where in the figure *Att* represents the attention module.

(MLE), which maximizes the likelihood of the target sequences conditioned on the source ones:

$$\theta^* = \operatorname{argmax}_{\theta} \mathbb{E}_{(X, Y^*) \sim D} \log p_{\theta}(Y^* | X) \quad (1)$$

Despite the popularity of MLE or teacher forcing (Doya, 1992) in neural sequence prediction tasks, two general issues are always haunting: 1). data sparsity and 2). tendency for overfitting, with which can both harm model generalization.

To combat data sparsity, different strategies have been proposed. Most of them try to take advantage of monolingual data (Sennrich et al., 2015; Zhang and Zong, 2016; Cheng et al., 2016). Others try to modify the ground truth target based on derived rules to get more similar examples for training (Norouzi et al., 2016; Ma et al., 2017). To alleviate overfitting, regularization techniques,

such as confidence penalization (Pereyra et al., 2017) and posterior regularization (Zhang et al., 2017), are proposed recently.

As shown in Figure 1, we propose a novel learning architecture, titled Generative Bridging Network (GBN), to combine both of the benefits from synthetic data and regularization. Within the architecture, the bridge module (bridge) first transforms the point-wise ground truth into a bridge distribution, which can be viewed as a target proposer from whom more target examples are drawn to train the generator. By introducing different constraints, the bridge can be set or trained to possess specific property, with which the drawn samples can augment target-side data (alleviate data sparsity) while regularizing the training (avoid overfitting) of the generator network (generator).

In this paper, we introduce three different constraints to build three bridge modules. Together with the generator network, three GBN systems are constructed: 1). a uniform GBN, instantiating the constraint as a uniform distribution to penalize confidence; 2). a language-model GBN, instantiating the constraint as a pre-trained neural language model to increase language smoothness; 3). a coaching GBN, instantiating the constraint as the generator’s output distribution to seek a close-to-generator distribution, which enables the bridge to draw easy-to-learn samples for the generator to learn. Without any constraint, our GBN degrades to MLE. The uniform GBN is proved to minimize KL-divergence with a so-called payoff distribution as in reward augmented maximum likelihood or RAML (Norouzi et al., 2016).

Experiments are conducted on two sequence prediction tasks, namely machine translation and abstractive text summarization. On both of them, our proposed GBNs can significantly improve task performance, compared with strong baselines. Among them, the coaching GBN achieves the best. Samples from these three different bridges are demonstrated to confirm the expected impacts they have on the training of the generator. In summary, our contributions are:

- A novel GBN architecture is proposed for sequence prediction to alleviate the data sparsity and overfitting problems, where the bridge module and the generator network are integrated and jointly trained.
- Different constraints are introduced to build GBN variants: uniform GBN, language-

model GBN and coaching GBN. Our GBN architecture is proved to be a generalized form of both MLE and RAML.

- All proposed GBN variants outperform the MLE baselines on machine translation and abstractive text summarization. Similar relative improvements are achieved compared to recent state-of-the-art methods in the translation task. We also demonstrate the advantage of our GBNs qualitatively by comparing ground truth and samples from bridges.




Figure 2: Conceptual interpretation of our Generative Bridging Network (GBN). See detailed discussion in the beginning of Sec. 2.

2 Generative Bridging Network

In this section, we first give a conceptual interpretation of our novel learning architecture which is sketched in Figure 2. Since data augmentation and regularization are two golden solutions for tackling data sparsity and overfitting issues. We are willing to design an architecture which can integrate both of their benefits. The basic idea is to use a so-called bridge which transforms Y^* to an easy-to-sample distribution, and then use this distribution (samples) to train and meanwhile regularize the sequence prediction model (the generator).

The bridge is viewed as a conditional distribution¹ $p_\eta(Y|Y^*)$ to get more target Y s given Y^* so as to construct more training pairs (X, Y) . In the meantime, we could inject (empirical) prior knowledge into the bridge through its optimization objective which is inspired by the design of the payoff distribution in RAML. We formulate the optimization objective with two parts in Equation (2): a) an expected similarity score computed through a similarity score function $S(\cdot, Y^*)$ interpolated with b) a knowledge injection constraint² $\mathcal{C}(p_\eta(Y|Y^*), p_c(Y))$ where α controls the

¹ η should be treated as an index of the bridge distribution, so it is not necessarily the parameters to be learned.

²Note that, in our paper, we specify \mathcal{C} to be KL-divergence between the bridge distribution p_η and certain constraint dis-

strength of the regularization, formally, we write the objective function $L_B(\eta)$ as follows:

$$L_B(\eta) = \mathbb{E}_{Y \sim p_\eta(Y|Y^*)} [-S(Y, Y^*)] + \alpha \mathcal{C}(p_\eta(Y|Y^*), p_c(Y)) \quad (2)$$

Minimizing it empowers the bridge distribution not only to concentrate its mass around the ground truth Y^* but also to adopt certain hope property from $p_c(Y)$. With the constructed bridge distribution, we optimize the generator network $P_\theta(Y|X)$ to match its output distribution towards the bridge distribution by minimizing their KL-divergence:

$$L_G(\theta) = KL(p_\eta(Y|Y^*) || p_\theta(Y|X)) \quad (3)$$

In practice, the KL-divergence is approximated through sampling process detailed in Sec. 2.3. As a matter of fact, the bridge is the crux of the integration: it synthesizes new targets to alleviate data sparsity and then uses the synthetic data as regularization to overcome overfitting. Thus a regularization-by-synthetic-example approach, which is very similar to the prior-incorporation-by-virtual-example method (Niyogi et al., 1998).

2.1 Generator Network

Our generator network is parameterized with the commonly used encoder-decoder architecture (Bahdanau et al., 2014; Cho et al., 2014). The encoder is used to encode the input sequence X to a sequence of hidden states, based on which an attention mechanism is leveraged to compute context vectors at the decoding stage. The context vector together with previous decoder’s hidden state and previously predicted label are used, at each time step, to compute the next hidden state and predict an output label.

As claimed in Equation (3), the generator network is not trained to maximize the likelihood of the ground truth but tries best to match the bridge distribution, which is a delegate of the ground truth. We use gradient descent to optimize the KL-divergence with respect to the generator:

$$\nabla L_G(\theta) = \mathbb{E}_{Y \sim p_\eta(Y|Y^*)} \log \nabla p_\theta(Y|X) \quad (4)$$

The optimization process can be viewed as the generator maximizing the likelihood of samples distribution p_c , however, we believe mathematical form of \mathcal{C} is not restricted, which could motivate further development.

drawn from the bridge. This may alleviate data sparsity and overfitting by posing more unseen scenarios to the generator and may help the generator generalize better in test time.

2.2 Bridge Module³

Our bridge module is designed to transform a single target example Y^* to a bridge distribution $p_\eta(Y|Y^*)$. We design its optimization target in Equation (2) to consist of two terms, namely, a concentration requirement and a constraint. The constraint is instantiated as KL-divergence between the bridge and a constraint distribution $p_c(Y)$. We transform Equation (2) as follows, which is convenient for mathematical manipulation later:

$$L_B(\eta) = \mathbb{E}_{Y \sim p_\eta} \left[-\frac{S(Y, Y^*)}{\tau} \right] + KL(p_\eta(Y|Y^*) || p_c(Y)) \quad (5)$$

$S(Y, Y^*)$ is a predefined score function which measures similarity between Y and Y^* and peaks when $Y = Y^*$, while $p_c(Y)$ reshapes the bridge distribution. More specifically, the first term ensures that the bridge should concentrate around the ground truth Y^* , and the second introduces willing property which can help regularize the generator. The hyperparameter τ can be interpreted as a temperature which scales the score function. In the following bridge specifications, the score function $S(Y, Y^*)$ is instantiated according to Sec. 3.1.

1. Delta Bridge The delta bridge can be seen as the simplest case where $\alpha = 0$ or no constraint is imposed. The bridge seeks to minimize $\mathbb{E}_{Y \sim p_\eta(Y|Y^*)} \left[-\frac{S(Y, Y^*)}{\tau} \right]$. The optimal solution is when the bridge only samples Y^* , thus the Dirac delta distribution is described as follows:

$$p_\eta(Y|Y^*) = \delta_{Y^*}(Y) \quad (6)$$

This exactly corresponds to MLE, where only examples in the dataset are used to train the generator. We regard this case as our baseline.

2. Uniform Bridge The uniform bridge adopts a uniform distribution $U(Y)$ as constraint. This

³Although we name it bridge module, we explicitly learn it with the generator when a closed-form static solution exists in terms of Equation (5). Otherwise, we will adopt an encoder-decoder to construct a dynamic bridge network.

bridge motivates to include noise into target example, which is similar to label smoothing (Szegedy et al., 2016). The loss function can be written as:

$$L_B(\eta) = \mathbb{E}_{Y \sim p_\eta} \left[-\frac{S(Y, Y^*)}{\tau} \right] + KL(p_\eta(Y|Y^*) || U(Y)) \quad (7)$$

We can re-write it as follows by adding a constant to not change the optimization result:

$$L_B(\eta) + C = KL(p_\eta(Y|Y^*) || \frac{\exp \frac{S(Y, Y^*)}{\tau}}{Z}) \quad (8)$$

This bridge is static for having a closed-form solution:

$$p_\eta(Y|Y^*) = \frac{\exp \frac{S(Y, Y^*)}{\tau}}{Z} \quad (9)$$

where Z is the partition function. Note that our uniform bridge corresponds to the payoff distribution described in RAML (Norouzi et al., 2016).

3. Language-model (LM) Bridge The LM bridge utilizes a pretrained neural language model $p_{LM}(Y)$ as constraint, which motivates to propose target examples conforming to language fluency.

$$L_B(\eta) = \mathbb{E}_{Y \sim p_\eta(Y|Y^*)} \left[-\frac{S(Y, Y^*)}{\tau} \right] + KL(p_\eta(Y|Y^*) || p_{LM}) \quad (10)$$

Similar to the uniform bridge case, we can re-write the loss function to a KL-divergence:

$$L_B(\eta) + C = KL(p_\eta(Y|Y^*) || \frac{p_{LM}(Y) \cdot \exp \frac{S(Y, Y^*)}{\tau}}{Z}) \quad (11)$$

Thus, the LM bridge is also static and can be seen as an extension of the uniform bridge, where the exponentiated similarity score is re-weighted by a pretrained LM score, and renormalized:

$$p(Y|Y^*) = \frac{p_{LM}(Y) \exp \frac{S(Y, Y^*)}{\tau}}{Z} \quad (12)$$

where Z is the partition function. The above equation looks just like the payoff distribution, whereas an additional factor is considered.

4. Coaching Bridge The coaching bridge utilizes the generator’s output distribution as constraint, which motivates to generate training samples which are easy to be understood by the generator, so as to relieve its learning burden. The coaching bridge follows the same spirit as the coach proposed in Imitation-via-Coaching (He et al., 2012), which, in reinforcement learning vocabulary, advocates to guide the policy (generator) with easy-to-learn action trajectories and let it gradually approach the oracle when the optimal action is hard to achieve.

$$L_B(\eta) = \mathbb{E}_{Y \sim p_\eta} \left[-\frac{S(Y, Y^*)}{\tau} \right] + KL(p_\theta(Y|X) || p_\eta(Y|Y^*)) \quad (13)$$

Since the KL constraint is a moving target when the generator is updated, the coaching bridge should not remain static. Therefore, we perform iterative optimization to train the bridge and the generator jointly. Formally, the derivatives for the coaching bridge are written as follows:

$$\begin{aligned} \nabla L_B(\eta) = & \mathbb{E}_{Y \sim p_\eta} \left[-\frac{S(Y, Y^*)}{\tau} \nabla \log p_\eta(Y|Y^*) \right] \\ & + \mathbb{E}_{Y \sim p_\theta} \nabla \log p_\eta(Y|Y^*) \end{aligned} \quad (14)$$

The first term corresponds to the policy gradient algorithm described in REINFORCE (Williams, 1992), where the coefficient $-S(Y, Y^*)/\tau$ corresponds to reward function. Due to the mutual dependence between bridge module and generator network, we design an iterative training strategy, i.e. the two networks take turns to update their own parameters treating the other as fixed.

2.3 Training

The training of the above three variants is illustrated in Figure 3. Since the proposed bridges can be divided into static ones, which only require pre-training, and dynamic ones, which require continual training with the generator, we describe their training process in details respectively.

2.3.1 Stratified-Sampled Training

Since closed-formed optimal distributions can be found for uniform/LM GBNs, we only need to draw samples from the static bridge distributions to train our sequence generator. Unfortunately,

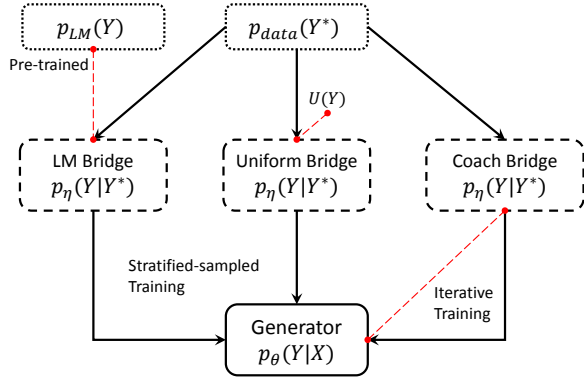


Figure 3: The training processes of the three different variants of our GBN architecture (Sec. 2.3).

due to the intractability of these bridge distributions, direct sampling is infeasible. Therefore, we follow Norouzi et al. (2016); Ma et al. (2017) and adopt stratified sampling to approximate the direct sampling process. Given a sentence Y^* , we first sample an edit distance m , and then randomly select m positions to replace the original tokens. The difference between the uniform and the LM bridge lies in that the uniform bridge replaces labels by drawing substitutions from a uniform distribution, while LM bridge takes the history as condition and draws substitutions from its step-wise distribution.

2.3.2 Iterative Training

Since the KL-constraint is a moving target for the coaching bridge, an iterative training strategy is designed to alternately update both the generator and the bridge (Algorithm 1). We first pre-train both the generator and the bridge and then start to alternately update their parameters. Figure 4 intuitively demonstrates the intertwined optimization effects over the coaching bridge and the generator. We hypothesize that iterative training with easy-to-learn guidance could benefit gradient update, thus result in better local minimum.

3 Experiment

We select machine translation and abstractive text summarization as benchmarks to verify our GBN framework.

3.1 Similarity Score Function

In our experiments, instead of directly using BLEU or ROUGE as reward to guide the bridge network’s policy search, we design a simple sur-

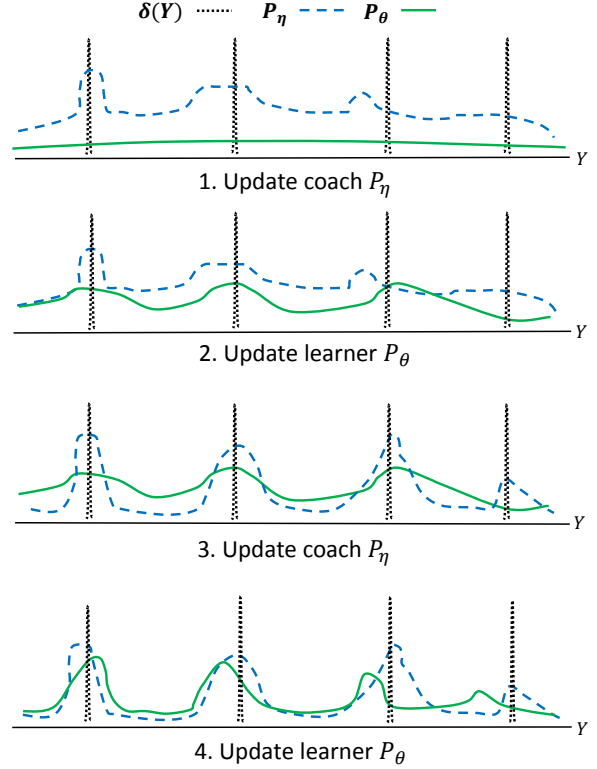


Figure 4: Four iterative updates of the coaching bridge and the generator. In an early stage, the pre-trained generator P_{θ} may not put mass on some ground truth target points within the output space, shown by $\delta(Y)$. The coaching bridge is first updated with Equation (14) to locate in between the Dirac delta distribution and the generator’s output distribution. Then, by sampling from the coaching bridge for approximating Equation (4), target samples which demonstrate easy-to-learn sequence segments facilitate the generator to be optimized to achieve closeness with the coaching bridge. Then this process repeats until the generator converges.

rogate n-gram matching reward as follows:

$$S(Y, Y^*) = 0.4 * N_4 + 0.3 * N_3 + 0.2 * N_2 + 0.1 * N_1 \quad (15)$$

N_n represents the n-gram matching score between Y and Y^* . In order to alleviate reward sparsity at sequence level, we further decompose the global reward $S(Y, Y^*)$ as a series of local rewards at every time step. Formally, we write the step-wise reward $s(y_t|y_{1:t-1}, Y^*)$ as follows:

$$s(y_t|y_{1:t-1}, Y^*) = \begin{cases} 1.0; & N(y_{1:t}, y_{t-3:t}) \leq N(Y^*, y_{t-3:t}) \\ 0.6; & N(y_{1:t}, y_{t-2:t}) \leq N(Y^*, y_{t-2:t}) \\ 0.3; & N(y_{1:t}, y_{t-1:t}) \leq N(Y^*, y_{t-1:t}) \\ 0.1; & N(y_{1:t}, y_t) \leq N(Y^*, y_t) \\ 0.0; & otherwise \end{cases} \quad (16)$$

where $N(Y, \tilde{Y})$ represents the occurrence of sub-sequence \tilde{Y} in whole sequence Y . Specifically, if

Algorithm 1 Training Coaching GBN**procedure** PRE-TRAINING

Initialize $p_\theta(Y|X)$ and $p_\eta(Y|Y^*)$ with random weights θ and η

Pre-train $p_\theta(Y|X)$ to predict Y^* given X

Use pre-trained $p_\theta(Y|X)$ to generate \hat{Y} given X

Pre-train $p_\eta(Y|Y^*)$ to predict \hat{Y} given Y^*

end procedure**procedure** ITERATIVE-TRAINING**while** Not Converged **do**

Receive a random example (X, Y^*)

if Bridge-step **then**

Draw samples Y from $p_\theta(Y|X)$

Update bridge via Equation (14)

else if Generator-step **then**

Draw samples Y from $p_\eta(Y|Y^*)$

Update generator via Equation (4)

end if**end while****end procedure**

a certain sub-sequence $y_{t-n+1:t}$ from Y appears less times than in the reference Y^* , y_t receives reward. Formally, we rewrite the step-level gradient for each sampled Y as follows:

$$\begin{aligned} & -\frac{S(Y, Y^*)}{\tau} \nabla \log p_\eta(Y|Y^*) \\ &= \sum_t -\frac{s(y_t|y_{1:t-1}, Y^*)}{\tau} \cdot \nabla \log p_\eta(y_t|y_{1:t-1}, Y^*) \end{aligned} \quad (17)$$

3.2 Machine Translation

Dataset We follow Ranzato et al. (2015); Bahdanau et al. (2016) and select German-English machine translation track of the IWSLT 2014 evaluation campaign. The corpus contains sentence-wise aligned subtitles of TED and TEDx talks. We use Moses toolkit (Koehn et al., 2007) and remove sentences longer than 50 words as well as lower-casing. The evaluation metric is BLEU (Papineni et al., 2002) computed via the multi-bleu.perl.

System Setting We use a unified GRU-based RNN (Chung et al., 2014) for both the generator and the coaching bridge. In order to compare with existing papers, we use a similar system setting with 512 RNN hidden units and 256 as embedding size. We use attentive encoder-decoder to build our system (Bahdanau et al., 2014). During training, we apply ADADELTA (Zeiler, 2012)

Methods	Baseline	Model
MIXER	20.10	21.81 +1.71
BSO	24.03	26.36 +2.33
AC	27.56	28.53 +0.97
Softmax-Q	27.66	28.77 +1.11
Uniform GBN ($\tau = 0.8$)	29.10	29.80 +0.70
LM GBN ($\tau = 0.8$)		29.90 +0.80
Coaching GBN ($\tau = 0.8$)		29.98 +0.88
Coaching GBN ($\tau = 1.2$)		30.15 +1.05
Coaching GBN ($\tau = 1.0$)		30.18 +1.08

Table 1: Comparison with existing works on IWSLT-2014 German-English Machine Translation Task.




Figure 5: Coaching GBN’s learning curve on IWSLT German-English Dev set.

with $\epsilon = 10^{-6}$ and $\rho = 0.95$ to optimize parameters of the generator and the coaching bridge. During decoding, a beam size of 8 is used to approximate the full search space. An important hyper-parameter for our experiments is the temperature τ . For the uniform/LM bridge, we follow Norouzi et al. (2016) to adopt an optimal temperature $\tau = 0.8$. And for the coaching bridge, we test hyper-parameters from $\tau \in \{0.8, 1.0, 1.2\}$. Besides comparing with our fine-tuned baseline, other systems for comparison of relative BLEU improvement are: MIXER (Ranzato et al., 2015), BSO (Wiseman and Rush, 2016), AC (Bahdanau et al., 2016), Softmax-Q (Ma et al., 2017).

Results The experimental results are summarized in Table 1. We can observe that our fine-tuned MLE baseline (29.10) is already over-

Methods	RG-1	RG-2	RG-L
ABS	29.55	11.32	26.42
ABS+	29.76	11.88	26.96
Luong-NMT	33.10	14.45	30.71
SAEASS	36.15	17.54	33.63
seq2seq+att	34.04	15.95	31.68
Uniform GBN ($\tau = 0.8$)	34.10	16.70	31.75
LM GBN ($\tau = 0.8$)	34.32	16.88	31.89
Coaching GBN ($\tau = 0.8$)	34.49	16.70	31.95
Coaching GBN ($\tau = 1.2$)	34.83	16.83	32.25
Coaching GBN ($\tau = 1.0$)	35.26	17.22	32.67

Table 2: Full length ROUGE F1 evaluation results on the English Gigaword test set used by (Rush et al., 2015). RG in the Table denotes ROUGE. Results for comparison are taken from SAEASS (Zhou et al., 2017).

competing other systems and our proposed GBN can yield a further improvement. We also observe that LM GBN and coaching GBN have both achieved better performance than Uniform GBN, which confirms that better regularization effects are achieved, and the generators become more robust and generalize better. We draw the learning curve of both the bridge and the generator in Figure 5 to demonstrate how they cooperate during training. We can easily observe the interaction between them: as the generator makes progress, the coaching bridge also improves itself to propose harsher targets for the generator to learn.

3.3 Abstractive Text Summarization

Dataset We follow the previous works by Rush et al. (2015); Zhou et al. (2017) and use the same corpus from Annotated English Gigaword dataset (Napoles et al., 2012). In order to be comparable, we use the same script⁴ released by Rush et al. (2015) to pre-process and extract the training and validation sets. For the test set, we use the English Gigaword, released by Rush et al. (2015), and evaluate our system through ROUGE (Lin, 2004). Following previous works, we employ ROUGE-1, ROUGE-2, and ROUGE-L as the evaluation metrics in the reported experimental results.

⁴<https://github.com/facebookarchive/NAMAS>




Figure 6: Coaching GBN’s learning curve on Abstractive Text Summarization Dev set.

System Setting We follow Zhou et al. (2017); Rush et al. (2015) to set input and output vocabularies to 119,504 and 68,883 respectively, and we also set the word embedding size to 300 and all GRU hidden state size to 512. Then we adopt dropout (Srivastava et al., 2014) with probability $p = 0.5$ strategy in our output layer. We use attention-based sequence-to-sequence model (Bahdanau et al., 2014; Cho et al., 2014) as our baseline and reproduce the results of the baseline reported in Zhou et al. (2017). As stated, the attentive encoder-decode architecture can already outperform existing ABS/ABS+ systems (Rush et al., 2015). In coaching GBN, due to the fact that the input of abstractive summarization X contains more information than the summary target Y^* , directly training the bridge $p_\eta(Y|Y^*)$ to understand the generator $p_\theta(Y|X)$ is infeasible. Therefore, we re-design the coaching bridge to receive both source and target input X, Y and we enlarge its vocabulary size to 88,883 to encompass more information about the source side. In Uniform/LM GBN experiments, we also fix the hyper-parameter $\tau = 0.8$ as the optimal setting.

Results The experimental results are summarized in Table 2. We can observe a significant improvement via our GBN systems. Similarly, the coaching GBN system achieves the strongest performance among all, which again reflects our assumption that more sophisticated regularization can benefit generator’s training. We draw the learning curve of the coaching GBN in Figure 6 to demonstrate how the bridge and the generator promote each other.

4 Analysis

By introducing different constraints into the bridge module, the bridge distribution will propose different training samples for the generator to learn. From Table 3, we can observe that most samples still reserve their original meaning. The uniform bridge simply performs random replacement without considering any linguistic constraint. The LM bridge strives to smooth reference sentence with high-frequent words. And the coaching bridge simplifies difficult expressions to relieve generator’s learning burden. From our experimental results, the more rational and aggressive diversification from the coaching GBN clearly benefits generator the most and helps the generator generalize to more unseen scenarios.

5 Related Literature

5.1 Data Augmentation and Self-training

In order to resolve the data sparsity problem in Neural Machine Translation (NMT), many works have been conducted to augment the dataset. The most popular strategy is via self-learning, which incorporates the self-generated data directly into training. Zhang and Zong (2016) and Sennrich et al. (2015) both use self-learning to leverage massive monolingual data for NMT training. Our bridge can take advantage of the parallel training data only, instead of external monolingual ones to synthesize new training data.

5.2 Reward Augmented Maximum Likelihood

Reward augmented maximum likelihood or RAML (Norouzi et al., 2016) proposes to integrate task-level reward into MLE training by using an exponentiated payoff distribution. KL divergence between the payoff distribution and the generator’s output distribution are minimized to achieve an optimal task-level reward. Following this work, Ma et al. (2017) introduces softmax Q-Distribution to interpret RAML and reveals its relation with Bayesian decision theory. These two works both alleviate data sparsity problem by augmenting target examples based on the ground truth. Our method draws inspiration from them but seeks to propose the more general Generative Bridging Network, which can transform the ground truth into different bridge distributions, from where samples are drawn will account for different interpretable factors.

System	Uniform GBN
Property	Random Replacement
Reference	the question <i>is</i> , is it worth it ?
Bridge	the question <i>lemon</i> , was it worth it ?
System	Language-model GBN
Property	Word Replacement
Reference	<i>now</i> how can this help us ?
Bridge	<i>so</i> how can this help us ?
System	Coaching GBN
Property	Reordering
Reference	i need <i>to have a health care lexicon</i> .
Bridge	i need <i>a lexicon for health care</i> .
Property	Simplification
Reference	<i>this is the way that</i> most of us were taught to <i>tie</i> our shoes .
Bridge	most of us learned to <i>bind</i> our shoes .

Table 3: Qualitative analysis for three different bridge distributions.

5.3 Coaching

Our coaching GBN system is inspired by imitation learning by coaching (He et al., 2012). Instead of directly behavior cloning the oracle, they advocate learning hope actions as targets from a coach which is interpolated between learner’s policy and the environment loss. As the learner makes progress, the targets provided by the coach will become harsher to gradually improve the learner. Similarly, our proposed coaching GBN is motivated to construct an easy-to-learn bridge distribution which lies in between the ground truth and the generator. Our experimental results confirm its effectiveness to relieve the learning burden.

6 Conclusion

In this paper, we present the Generative Bridging Network (GBN) to overcome data sparsity and overfitting issues with Maximum Likelihood Estimation in neural sequence prediction. Our implemented systems prove to significantly improve the performance, compared with strong baselines. We believe the concept of bridge distribution can be applicable to a wide range of distribution matching tasks in probabilistic learning. In the future, we intend to explore more about GBN’s applications as well as its provable computational and statistical guarantees.

References

- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. *arXiv preprint arXiv:1607.07086* .
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* .
- Yong Cheng, Wei Xu, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Semi-supervised learning for neural machine translation. *arXiv preprint arXiv:1606.04596* .
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* .
- Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. 2015. Attention-based models for speech recognition. In *Advances in Neural Information Processing Systems*. pages 577–585.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* .
- Kenji Doya. 1992. Bifurcations in the learning of recurrent neural networks. In *Circuits and Systems, 1992. ISCAS'92. Proceedings., 1992 IEEE International Symposium on*. IEEE, volume 6, pages 2777–2780.
- He He, Jason Eisner, and Hal Daume. 2012. Imitation learning by coaching. In *Advances in Neural Information Processing Systems*. pages 3149–3157.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th annual meeting of the ACL on interactive poster and demonstration sessions*. Association for Computational Linguistics, pages 177–180.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL-04 workshop*. Barcelona, Spain, volume 8.
- Xuezhe Ma, Pengcheng Yin, Jingzhou Liu, Graham Neubig, and Eduard Hovy. 2017. Softmax q-distribution estimation for structured prediction: A theoretical interpretation for raml. *arXiv preprint arXiv:1705.07136* .
- Courtney Napoles, Matthew Gormley, and Benjamin Van Durme. 2012. Annotated gigaword. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction*. Association for Computational Linguistics, pages 95–100.
- Partha Niyogi, Federico Girosi, and Tomaso Poggio. 1998. Incorporating prior information in machine learning by creating virtual examples. *Proceedings of the IEEE* 86(11):2196–2209.
- 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 *Proceedings of the 40th annual meeting on association for computational linguistics*. Association for Computational Linguistics, pages 311–318.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. 2017. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548* .
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732* .
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685* .
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709* .
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15(1):1929–1958.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 2818–2826.

- Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015. Grammar as a foreign language. In *Advances in Neural Information Processing Systems*. pages 2773–2781.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4):229–256.
- Sam Wiseman and Alexander M Rush. 2016. Sequence-to-sequence learning as beam-search optimization. *arXiv preprint arXiv:1606.02960* .
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*. pages 2048–2057.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701* .
- Jiacheng Zhang, Yang Liu, Huanbo Luan, Jingfang Xu, and Maosong Sun. 2017. Prior knowledge integration for neural machine translation using posterior regularization. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. volume 1, pages 1514–1523.
- Jiajun Zhang and Chengqing Zong. 2016. Exploiting source-side monolingual data in neural machine translation. In *EMNLP*. pages 1535–1545.
- Qingyu Zhou, Nan Yang, Furu Wei, and Ming Zhou. 2017. Selective encoding for abstractive sentence summarization. *arXiv preprint arXiv:1704.07073* .