Generative Bridging Network for Neural Sequence Prediction

Wenhu Chen, Guanlin Li, Shuo Ren, Shujie Liu, Zhirui Zhang, Mu Li, Ming Zhou
Microsoft Research Asia & University of California, Santa Barbara
Contact Email: wenhuchen@cs.ucsb.edu

Introduction

- We propose a new training algorithm to resolve the data sparsity problem in neural sequence prediction.
- We introduce a bridging network to assist us in training sequence prediction models by extending point-wise ground truth to a bridge distribution

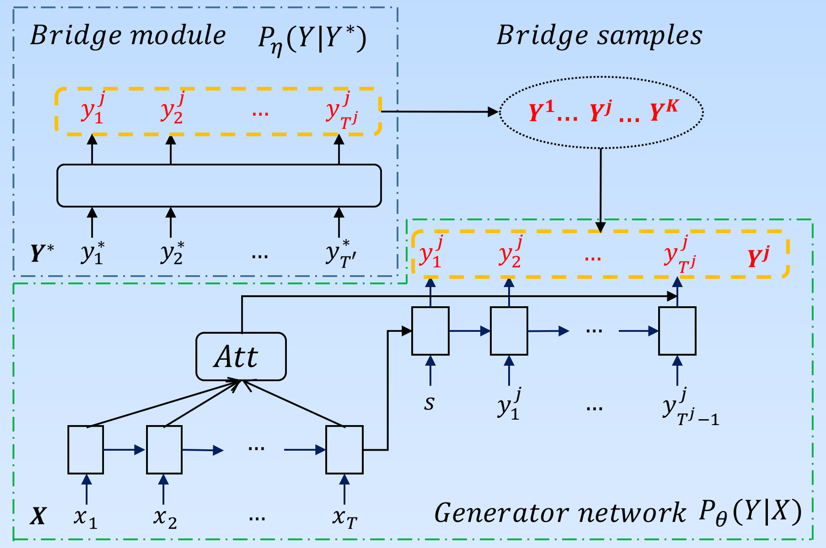

Model

Background: The typical training algorithm is Maximum Likelihood Estimation

$$\theta^* = argmax_\theta \mathbb{E}_{(X^*,Y^*)} \log p_\theta(Y^*|X^*)$$

where D is the dataset, X^*, Y^* are the input sequence and output sequence. MLE is known to suffer from **sparsity problem**.

Goal: We use Bridge Network to transform the point-wise ground truth into target-side distribution $p_\eta(Y|Y^*)$ by imposing prior constraints over the target sequence. The target distribution draws samples for the sequence model to learn.

Objective: We design our Bridge Network based on the following criterion:

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

where $S(Y, Y^*)$ denotes the similarity function (BLEU, METEOR, etc) between Y and Y^*, α is the balancing factor between similarity and regularization, $p_c(Y)$ is our prior knowledge over sequence Y.

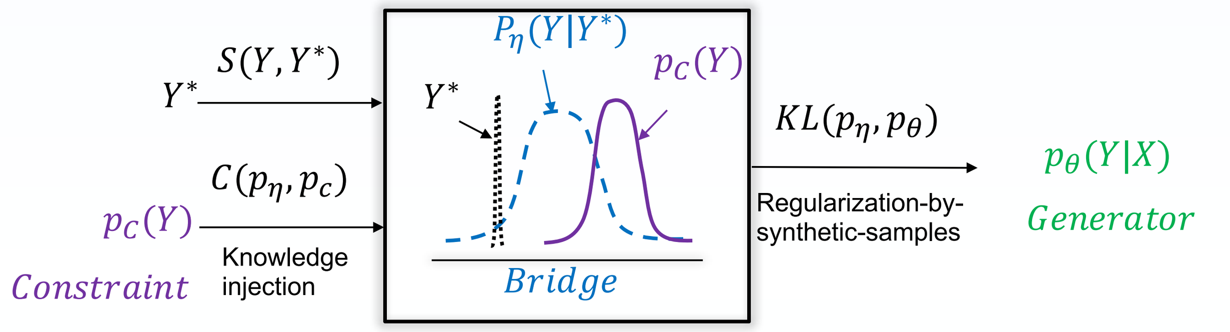

Algorithm

Three types of Bridge Networks:

1) Uniform Bridge: we hope the bridge will be as diverse as possible, we set $p_c(Y) = U(Y)$, we arrive the solution $p_\eta(Y|Y^*) = \frac{e^{\frac{S(Y,Y^*)}{\tau}}}{Z}$
2) Language-Model Bridge: we hope the bridge will follow the language model, we set $p_c(Y) = p_{LM}(Y)$, we arrive the solution $p_\eta(Y|Y^*) = \frac{p_{LM}(Y|Y^*) e^{\frac{S(Y,Y^*)}{\tau}}}{Z}$
3) Coaching Bridge: we hope the bridge lies in the middle of ground truth and model prediction, we set $p_c(Y) = p_\theta(Y|X^*)$. Coaching bridge can ease the training and guide the sequence model to gradually approach the oracle. We parametrize $p_\eta(Y|Y^*)$ with RNN model, which takes ground truth as input.

Sequence model:

$$p_\theta(Y|X^*) = argmin_{p_\theta(Y|X^*)} KL(p_\theta(Y|X^*)||p_\eta(Y|Y^*))$$

Generative Bridging Networks (GBN):

1) For uniform bridge and LM bridge, sequence model and bridge is not interleaved, we use closed-form bridge distribution to draw samples.
2) For coaching bridge, the sequence model and bridge is interleaved, we adopt SGD algorithm to alternately update them.

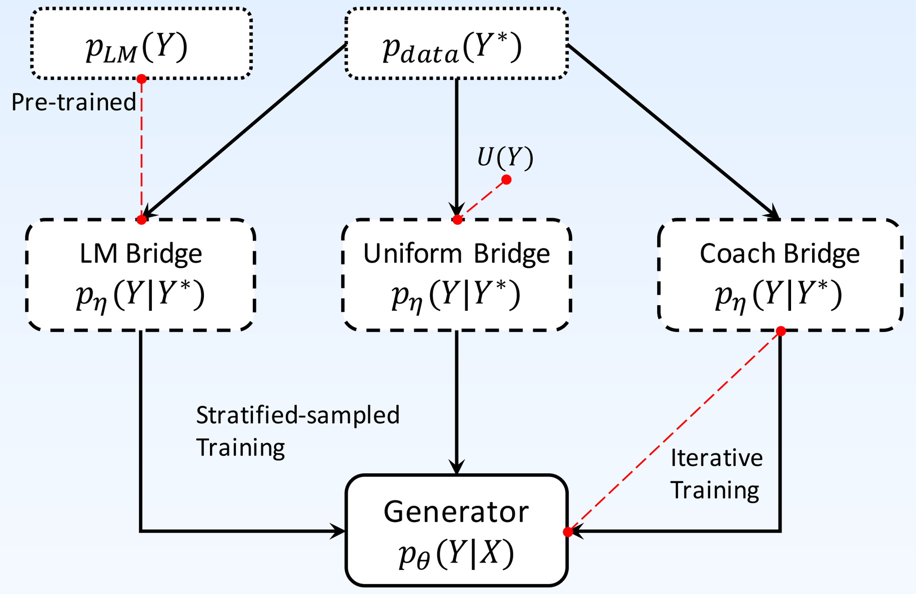

Iterative Training Procedure for Coaching GBN:

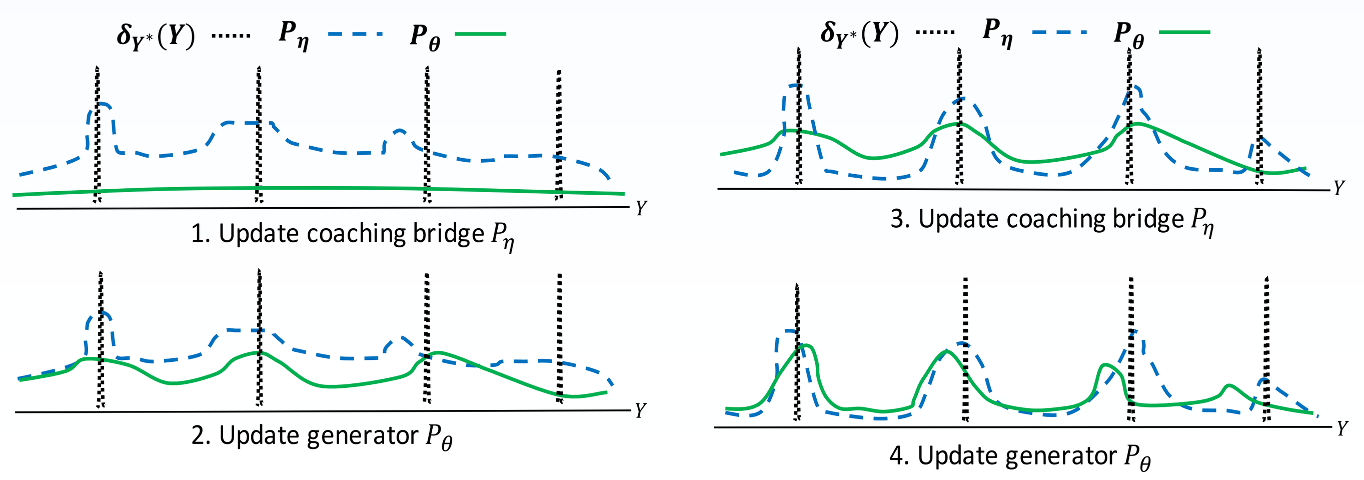

Machine Translation Results

Dataset:

1) We 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 test our models on newstest corpus.
2) We use BLEU to evaluate the performance of our model

Methods	Baselines	BLEU
MIXER (Marc'Aurelio et al. 2015)	20.10	21.81 (+ 1.71)
BSO (Sam Wiseman et al. 2016)	24.03	26.36 (+2.33)
Actor-critic (Bahdanau et al. 2017)	27.53	28.53 (+0.93)
Softmax-Q (Ma et al. 2017)	27.66	28.77 (+1.11)
Uniform GBN	29.10	29.88 (+0.70)
LM GBN	29.10	29.98 (+0.88)
Coaching GBN	29.10	**30.18** (+1.08)

Abstractive Summarization Results

Dataset:

1) We use the same corpus from Annotated English Gigaword dataset (Napoles et al., 2012). we use the same script released by Rush et al. (2015) to pre-process and extract the training and validation sets
2) We use ROUGE to evaluate the performance of our model

Methods	RG-1	RG-2	RG-L
ABS (Rush et al. 2015)	29.55	11.32	26.42
ABS+ (Rush et al. 2015)	29.77	11.88	26.96
Luong-NMT (Luong et al. 2016)	33.10	14.45	30.71
SAEASS (Zhou et al. 2017)	**36.15**	**17.54**	**33.63**
Seq2seq-Att (Bahdanau et al. 2014)	34.04	15.95	31.68
Uniform GBN	34.10	16.70	31.75
LM GBN	34.32	16.88	31.89
Coaching GBN	**35.26**	**17.22**	**32.67**

Conclusion

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.

Result Analysis

We can observe that most samples still preserve their original meanings. The uniform bridge simply performs random replacement without considering any linguistic constraint. The LM bridge strives to smooth reference sentence with high-frequency words, and the coaching bridge simplifies difficult expressions to relieve generator's learning burden.

System	Uniform GBN
Property	Random Replacement
Reference	the question is , **is** it worth it ?
Bridge Sample	the question **lemon** , **was** it worth it ?

System	Language-model GBN
Property	Word Replacement
Reference	**now** how can this help us ?
Bridge Sample	**so** how can this help us ?

System	Coaching GBN
Property	Reordering
Reference	I need to **have a health care lexicon** .
Bridge Sample	I need a **lexicon for health care** .
Property	Simplification
Reference	**this is the way that** most of us **were taught to** tie our shoes .
Bridge Sample	most of us **learned** to bind our shoes .

References

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.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 .

Sam Wiseman and Alexander M Rush. 2016. Sequence-to-sequence learning as beam-search optimization. arXiv preprint arXiv:1606.02960 .

Marc'Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. arXiv preprint arXiv:1511.06732 .

Qingyu Zhou, Nan Yang, Furu Wei, and Ming Zhou. 2017. Selective encoding for abstractive sentence summarization. arXiv preprint arXiv:1704.07073

He He, Jason Eisner, and Hal Daume. 2012. Imitation learning by coaching. In Advances in Neural Information Processing Systems. pages 3149–3157.

Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. arXiv preprint arXiv:1509.00685 .

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 .

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.