Adversarial Ranking for Language Generation

Kevin Lin*†, Dianqi Li*†, Xiaodong He‡, Zhengyou Zhang‡, Ming-Ting Sun†
{kvlin,dianqili,mts}@uw.edu, {xiaohe,zhang}@microsoft.com
†University of Washington, Seattle, ‡Microsoft Research, Redmond

Introduction

- Generate natural language descriptions similar to human-written.
- Existing GANs restrict the discriminator to be a binary classifier.
- The proposed RankGAN learns from relative ranking between machine- and human-written sentences in an adversarial framework.

Adversarial Ranking

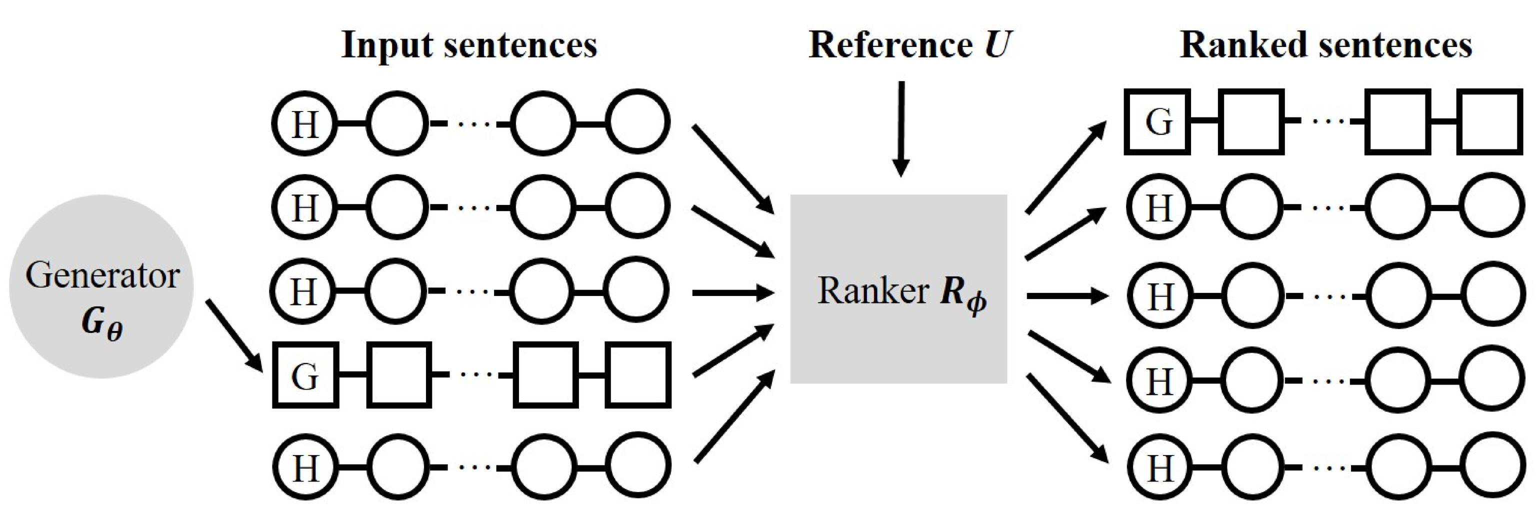

Figure 1: **H** denotes the sentence sampled from the human-written sentences. **G** is the sentence generated by the generator G_θ. It is illustrated that the generator tries to fool the ranker and let the synthetic sentence to be ranked at the top with respect to the reference sentence.

MinMax game:

$$\min_\theta \max_\phi \mathfrak{L}(G_\theta, R_\phi) = \mathop{\mathbb{E}}_{s\sim\mathcal{P}_h}\left[\log R_\phi(s|U,\mathcal{C}^-)\right] + \mathop{\mathbb{E}}_{s\sim G_\theta}\left[\log(1 - R_\phi(s|U,\mathcal{C}^+))\right]$$

where U is the reference set used for estimating relative ranks, and $\mathcal{C}^+, \mathcal{C}^-$ are the comparison set with regard to input sentences s.

Ranking score:

$$\log R_\phi(s|U,\mathcal{C}) = \mathop{\mathbb{E}}_{u\in U} \log\left[\frac{exp(\gamma\alpha(s|u))}{\sum_{s'\in\mathcal{C}'} exp(\gamma\alpha(s'|u))}\right]$$

$\alpha(s|u)$ computes cosine similarity between reference u and input sequence s. $\mathcal{C}'$ denotes all sentences to be ranked. γ is a hyperparameter.

Ranking score for partial sentence:

$$V_{\theta,\phi}(s_{1:t-1}, U) = \mathop{\mathbb{E}}_{s_r\sim G_\theta}\left[\log R_\phi(s_r|U,\mathcal{C}^+, s_{1:t-1})\right]$$

where s_r is a simulated sentence by Markov rollout method.

Policy gradient for generator G_θ:

$$\nabla_\theta \mathfrak{L}_\theta(s_0) = \mathop{\mathbb{E}}_{s_{1:T}\sim G_\theta}\left[\sum_{t=1}^{T}\sum_{w_t\in V} \nabla_\theta \pi_\theta(w_t|s_{1:t-1}) V_{\theta,\phi}(s_{1:t}, U)\right]$$

where w_t is the next token sampled by policy π_θ.

Synthetic Data Experiments

Method	MLE	PG-BLEU	SeqGAN	RankGAN
NLL	9.038	8.946	8.736	**8.247**

Table 1: The performance comparison of different methods on the synthetic data [SeqGAN, AAA17] in terms of the negative log-likelihood (NLL) scores. The lower the NLL score is, the higher chance the synthesized sentences will pass Turing test.

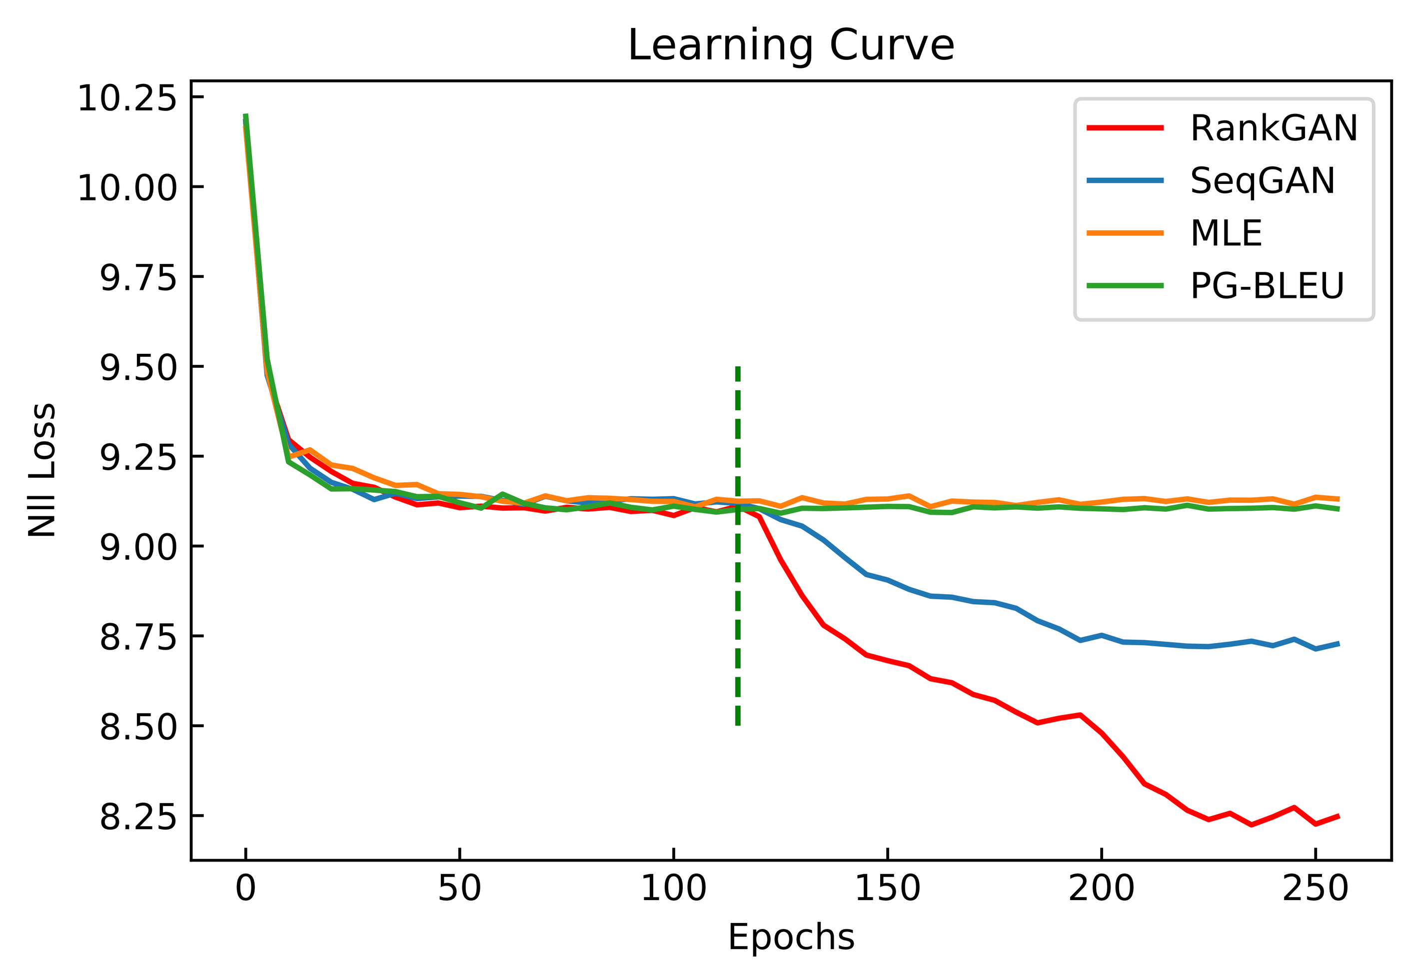

Figure 2: Learning curves of different methods on the simulation of synthetic data with respect to different training epochs. Note that the vertical dashed line indicates the end of the pre-training of PG-BLEU, SeqGAN and RankGAN.

Results on Chinese Poems Composition

Method	BLEU-2
MLE	0.667
SeqGAN	0.738
RankGAN	**0.812**

Method	Human score
SeqGAN	3.58
RankGAN	4.52
Human-written	**6.69**

Table 2: The performance comparison of different methods on the Chinese poem generation in terms of the BLEU scores and human evaluation scores.

Results on COCO image captions

Method	BLEU-2	BLEU-3	BLEU-4
MLE	0.781	0.624	0.589
SeqGAN	0.815	0.636	0.587
RankGAN	**0.845**	**0.668**	**0.614**

Method	Human score
SeqGAN	3.44
RankGAN	4.61
Human-written	**6.42**

Table 3: The performance comparison of different methods on the COCO captions in terms of the BLEU scores and human evaluation scores.

Human-written
Two men happily working on a plastic computer. The toilet in the bathroom is filled with a bunch of ice. A bottle of wine near stacks of dishes and food.
SeqGAN (Baseline)
A baked mother cake sits on a street with a rear of it. A tennis player who is in the ocean. A highly many fried scissors sits next to the older.
RankGAN (Ours)
Three people standing in front of some kind of boats. A bedroom has silver photograph desk. The bears standing in front of a palm state park.

Table 4: Example of the generated descriptions with different methods. Note that the language models are trained on COCO caption dataset without the images.

Results on Shakespear's Plays

Method	BLEU-2	BLEU-3	BLEU-4
MLE	0.796	0.695	0.635
SeqGAN	0.887	0.842	0.815
RankGAN	**0.914**	**0.878**	**0.856**

Table 5: The performance comparison of different methods on Shakespeare's play *Romeo and Juliet* in terms of the BLEU scores.