The Thirty-Second AAAI Conference
on Artificial Intelligence (AAAI-18)

Dialogue Generation with GAN

Hui Su,[1*] Xiaoyu Shen,[2] Pengwei Hu,[1] Wenjie Li,[1] Yun Chen[3]
[1]The Hong Kong Polytechnic University, Hong Kong
[2]Max Planck Institute Informatics, Saarbrücken, Germany
[3]The University of Hong Kong, Hong Kong

Abstract

This paper presents a Generative Adversarial Network (GAN) to model multi-turn dialogue generation, which trains a latent hierarchical recurrent encoder-decoder simultaneously with a discriminative classifier that make the prior approximate to the posterior. Experiments show that our model achieves better results.

Introduction

Generative Adversarial Nets (GANs) (Goodfellow et al. 2014) have made significant progress in learning smooth latent variable representations of continuous data. However, directly applying GAN on discrete text data is difficult. Current approaches usually combine it with reinforcement learning or continuous approximations (Kim et al. 2017), but none of them have demonstrated convincing advantages. Using GAN with reinforcement learning (RL) (Li et al. 2017) has been made to investigate dialogue generation, while the model variance is too high.

Recently (Serban et al. 2017) proposed the Variational Hierarchical Recurrent Encoder-Decoder (VHRED) model which borrows the idea from conditional variational autoencoders (CVAE) (Sohn, Lee, and Yan 2015) to model multi-turn dialogue generation. However, the approximated posterior distribution needs to be specified and the prior distribution is not guaranteed to be the same as the marginal posterior distribution in the global optimum.

Inspired by the recent success of Adversarial Autoencoder (Makhzani et al. 2015), in this paper, we propose an end-to-end differentiable GAN framework (D-GAN) for dialogue generation, which learns a discriminator that measures the differences between the continuous latent representation of prior and posterior distribution. Specifically, in a conversation, a response x is generated as follows: A latent variable z is sampled from a prior distribution $p_\theta(z|c)$ based on the current dialogue context c, then the response is generated from the distribution $p_\theta(x|z,c)$. In training phrase, we generate response from posterior $p_\theta(z|c,x)$ and the discriminator is trained to distinguish between samples from the prior and posterior distribution, thereby pushing the prior distribution to match the posterior. The posterior take the input x and a random noise η with a fixed distributionn (e.g., Gaussian), which can learn any arbitrary distribution and no longer constrained to be Gaussian.

We demonstrate that the GAN framework can be effectively used for common dialogue generation and generate more consistent latent variables than VHRED model.

[*]Correspondence to H. Su (suhui15@mails.ucas.ac.cn).

Related work

VHRED The variational hierarchical recurrent encoder decoder (VHRED) model has previously been proposed for dialogue modeling. The response x is generated from the distribution $p_\theta(x|z,c)$. In contrast to calculating the exact log-likelihood, it can be efficiently trained by optimizing a valid lower bound. The objective takes the following form:

$$\begin{aligned}&-\log p_\theta(x|c) = -\log \int_z p_\theta(z|c)p_\theta(x|z,c)dz\\ &\leq -\mathbb{E}_{q_\phi(z|x,c)}[\log p(x|z,c)] + \mathrm{KL}(q_\phi(z|x,c)||p(z)|c)))\end{aligned} \tag{1}$$

Adversarial autoencoders CVAE model typically use a factorial Gaussian as the prior, which enables closed-form optimization while restricting the expressive power of the model. Adversarial autoencoders replace the KL divergence with an adversarial training criterion to allow richer families of priors. Our work differs in that we do not sample from a fixed prior distribution, both prior and posterior are instead parameterized through the neural network.

Model

We decompose a dialogue into two levels: sequences of utterances and sub-sequences of words, as in (Serban et al. 2017). Let $\mathbf{w}_1, \ldots, \mathbf{w}_N$ be a dialogue with N utterances, where $\mathbf{w}_n = (w_{n,1}, \ldots, w_{n,M_n})$ is the n-th utterance. The probability distribution of the utterance sequence factorizes as:

$$\prod_{n=1}^{N}\prod_{m=1}^{M_n} P_\theta(\mathbf{w}_{m,n}|\mathbf{w}_{m,<n}, \mathbf{w}_{<n}) \tag{2}$$

where θ represents the model parameters and $\mathbf{w}_{<n}$ encodes the dialog context until step n. A word encoder and context encoder is implemented to separately model the dynamics of these two levels.

In the training phase, we have access to the golden response and generate latent variables with $z_1 = f_\phi(c, h, \eta)$, where c is the last hidden state of the context encoder, h is the last hidden state of the word encoder run on the golden response and η is a Gaussian noise sampled from $\mathcal{N}(0, 1)$. For dialogue modeling, the true posterior distribution should be complex and multimodal rather than a simple Gaussian, so we use neural networks to appropriately generating such complex distribution. A sample is drawn from this distribution, then this sample is given as input to the decoder RNN, which then computes the output probabilities of the words in the next utterance. In the testing phase, we have only the context information and latent variables are generated with $z_2 = \mathcal{G}(c, \eta)$. We use GAN to close the divergence between z_1 and z_2. $f_\phi(c, h)$ defines an aggregated distribution of z_1 given a dialogue context as follows:

$$q_\phi(z_1|c) = \int_h q_\phi(z_1|c, h, \eta) p_d(h) dh \tag{3}$$

$p_d(h)$ is the real distribution of the response. $\mathcal{G}(c, \eta)$ also defines a distribution $p_\theta(z_2|c)$. We apply an adversarial network on top of them to match the distribution $q_\phi(z_1|c)$ and $p_\phi(z_2|c)$. The discriminator is implemented as a feedforward neural network conditioned on the dialogue context c, judging whether a latent variable comes from $q_\phi(z_1|c)$ or not. The objective function is:

$$\mathcal{L}(G) = \min_D \max_{\phi, \mathcal{G}} \mathbb{E}_{h,\eta} D(f_\phi(c, h) + \eta) - \mathbb{E}_\eta D(\mathcal{G}(c, \eta)) \tag{4}$$

Here we use the WGAN objective (details can be found in the supplementary material) to replace the original log-likelihood. We update both the z_1 generator $\mathcal{G}$ and the z_2 generator $f_\phi(c, h)$ to prevent the overfitting of z_2 generator. It also functions as a regularizer to force $f_\phi(c, h)$ to extract information from both c and h.When optimizing with $\mathcal{L}(\mathcal{G})$, we first train the discriminator $\mathcal{D}$ for several steps then train the z_2 generator to receive stable gradient information from the discriminator.

Experiment

Datasets and Baseline

We conduct our experiments on the multi-turn dialogue datasets Switchboard (Godfrey and Holliman 1997). This dataset is randomly separated into training/validation/test sets with the ratio of 10:1:1. We compare our approach with two baseline methods, including HRED (Serban et al. 2016) and VHRED (Serban et al. 2017).

Implementation Details

For every model, the word embeddings are initialized with the Word2Vec embeddings trained on the Google News Corpus[1]. The vocabulary set is defined as the most frequent 20,000 words on every corpus, the left words are mapped to an unknown token. The batch size is set to 128, truncated backpropagation and gradient clipping are used. To decide the stopping point, we first train an independent GRU language model on the training corpus, then apply it to test the perplexity of the decoded responses. The weight clipping threshold ϵ for the discriminator training in GAN is set as 0.01.

[1]https://code.google.com/archive/p/word2vec/

Table 1: Evaluation results

Model	Greedy	Average	Extrema	BLEU-1
HRED	0.482	0.343	0.290	0.254
VHRED	0.502	0.358	0.335	0.267
D-GAN	**0.537**	**0.372**	**0.461**	**0.286**

Experimental results

In our experiments, the results are evaluated using e embedding-based (Liu et al. 2016) metrics and BLEU-1 score. We summarize the experiment results in Table 1. Compared to previous methods, our approaches achieve better results on the experimental dataset.

Conclusion

In this paper, we proposed a novel variant GAN framework to improve the performance of dialogue generation. Experiments showed that our method achieves better results.

References

Godfrey, J., and Holliman, E. 1997. Switchboard-1 release 2, 1997. *Linguistic Data Consortium, Philadelphia*.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.

Kim, Y.; Zhang, K.; Rush, A. M.; LeCun, Y.; et al. 2017. Adversarially regularized autoencoders for generating discrete structures. *arXiv preprint arXiv:1706.04223*.

Li, J.; Monroe, W.; Shi, T.; Ritter, A.; and Jurafsky, D. 2017. Adversarial learning for neural dialogue generation. *arXiv preprint arXiv:1701.06547*.

Liu, C.-W.; Lowe, R.; Serban, I. V.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. *arXiv preprint arXiv:1603.08023*.

Makhzani, A.; Shlens, J.; Jaitly, N.; and Goodfellow, I. J. 2015. Adversarial autoencoders. *CoRR*.

Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A. C.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models.

Serban, I. V.; Sordoni, A.; Lowe, R.; Charlin, L.; Pineau, J.; Courville, A.; and Bengio, Y. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. *AAAI*.

Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, 3483–3491.