Generative Adversarial Networks (GANs) for Dialog Response Generation

Background about dialog response generation
Dialogue systems, also known as interactive conversational agents, virtual agents and sometimes chatbots, are used in a wide set of applications ranging from technical support services to language learning tools and entertainment. In this project we are going to focus on applying a very new method, called generative adversarial networks(GANs), to automatically generate responses given an input query and context.

Background about GANs
In the deep learning community, a recently proposed *adversarial nets* framework [1] haven been proved to be good at image generation tasks. In GANs the generative model is pitted against an adversary: a discriminative model that learns to determine whether a sample is realistic. The generative model can be thought of as analogous to a team of counterfeiters, trying to produce fake currency and use it without detection, while the discriminative model is analogous to the police, trying to detect the counterfeit currency. Competition in this game drives both teams to improve their methods until the counterfeits are indistinguishable from the genuine articles, that is, the resulting generative model is good.

Your Task
Despite the prosperity and fast development of GANs, GANs have not been applied successfully to NLP tasks because regular GANs are only defined for real-valued data. It works by training a generator network that outputs synthetic data, then running a discriminator network on the synthetic data. The gradient of the output of the discriminator network with respect to the synthetic data tells us how to slightly change the synthetic data to make it more realistic. We can make slight changes to the synthetic data only if it is based on continuous numbers. If it is based on discrete numbers, there is no way to make a slight change. Besides, the network will not be differentiable from end to end with discrete hidden layer.

Several potential approaches are worthy trying to be studied under the framework of GANs to generate discrete tokens like in dialog setting. Several potential directions are elaborated below. In the project, you are going to conducting research from one of the following perspective and try it out on some real datasets (released by Facebook, Mircrosoft, etc.) for response generation task.
Depending on the quality of the results we might attempt to publish it as a scientific publication.

Potential perspectives

1. Reinforcement learning

Recently, [2] models the data generator as a stochastic policy in reinforcement learning, called SeqGAN. It bypasses the generator differentiation problem by directly performing policy gradient update. Policy gradient method is a type of reinforcement learning techniques that relies upon optimizing parametrized policies (generator in GANs) with respect to the expected return (long-term cumulative reward) by gradient descent.

2. Gradient approximation

Some other works are very recently proposed to deal with discrete latent variables for neural networks by estimating gradient. [3] present an efficient gradient estimator that replaces the non-differentiable sample from a categorical distribution with a differentiable sample from a *Gumbel-Softmax* distribution which can be smoothly annealed into a categorical distribution. Another distribution, called *Concrete* distribution [4], can also be considered to approximate the gradients of discrete latent variables.

Ideal candidate

• Experienced with Python or other script languages.
• Familiar with machine learning, especially deep learning.
• Quick leaner and passionate about the topic.

Contact

Fei Mi (fei.mi@epfl.ch)

References

[1] Goodfellow, Ian, et al. "Generative adversarial nets." *Advances in Neural Information Processing Systems*. 2014.

[2] Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. Seqgan: sequence generative adversarial nets with policy gradient. arXiv preprint arXiv:1609.05473, 2016.

[3] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. arXiv preprint arXiv:1611.01144, 2016.

[4] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. arXiv preprint arXiv:1611.00712, 2016.