

# SELF-ADVERSARIAL LEARNING WITH COMPARATIVE DISCRIMINATION FOR TEXT GENERATION

Wangchunshu Zhou<sup>1\*</sup> Tao Ge<sup>2</sup> Ke Xu<sup>1</sup> Furu Wei<sup>2</sup> Ming Zhou<sup>2</sup>

<sup>1</sup>Beihang University, Beijing, China

<sup>2</sup>Microsoft Research Asia, Beijing, China

zhouwangchunshu@buaa.edu.cn, kexu@nlsde.buaa.edu.cn  
{tage, fuwei, mingzhou}@microsoft.com

## ABSTRACT

Conventional Generative Adversarial Networks (GANs) for text generation tend to have issues of reward sparsity and mode collapse that affect the quality and diversity of generated samples. To address the issues, we propose a novel self-adversarial learning (SAL) paradigm for improving GANs’ performance in text generation. In contrast to standard GANs that use a binary classifier as its discriminator to predict whether a sample is real or generated, SAL employs a comparative discriminator which is a pairwise classifier for comparing the text quality between a pair of samples. During training, SAL rewards the generator when its currently generated sentence is found to be better than its previously generated samples. This self-improvement reward mechanism allows the model to receive credits more easily and avoid collapsing towards the limited number of real samples, which not only helps alleviate the reward sparsity issue but also reduces the risk of mode collapse. Experiments on text generation benchmark datasets show that our proposed approach substantially improves both the quality and the diversity, and yields more stable performance compared to the previous GANs for text generation.

## 1 INTRODUCTION

Generative Adversarial Networks (Goodfellow et al., 2014) (GANs) have achieved tremendous success for image generation and received much attention in computer vision. For text generation, however, the performance of GANs is severely limited due to reward sparsity and mode collapse: reward sparsity refers to the difficulty for the generator to receive reward signals when its generated samples can hardly fool the discriminator that is much easier to train; while mode collapse refers to the phenomenon that the generator only learns limited patterns from the real data. As a result, both the quality and diversity of generated text samples are limited.

To address the above issues, we propose a novel self-adversarial learning (SAL) paradigm for improving adversarial text generation. In contrast to standard GANs (Figure 1(a)) that use a binary classifier as its discriminator to predict whether a sample is real or generated, SAL employs a comparative discriminator which is a pairwise classifier assessing whether the currently generated sample is better than its previously generated one, as shown in Figure 1(b). During training, SAL rewards the generator when its currently generated samples are found to be better than its previously generated samples. In the earlier training stage when the quality of generated samples is far below the real data, this self-improvement reward mechanism makes it easier for the generator to receive non-sparse rewards with informative learning signals, effectively alleviating the reward sparsity issue; while in the later training stage, SAL can prevent a sample from keeping receiving high reward as the self-improvement for a popular mode will become more and more difficult, and therefore help the generator avoid collapsing toward the limited patterns of real data.

We comprehensively evaluate the proposed self-adversarial learning paradigm in both synthetic data and real data on the text generation benchmark platform (Zhu et al., 2018). Compared to the previous approaches for adversarial text generation (Yu et al., 2017; Che et al., 2017; Lin et al., 2017), our

\* This work was done during the first author’s internship at Microsoft Research Asia.