
Improving GAN Training with Probability Ratio Clipping and Sample Reweighting

Yue Wu¹, Pan Zhou³, Andrew Gordon Wilson⁴, Eric P. Xing^{1,2}, Zhiting Hu^{1,2}

¹Carnegie Mellon University, ²Petuum Inc., ³National University of Singapore, ⁴New York University

Abstract

Despite success on a wide range of problems related to vision, generative adversarial networks (GANs) can suffer from inferior performance due to unstable training, especially for text generation. We propose a new variational GAN training framework which enjoys superior training stability. Our approach is inspired by a connection of GANs and reinforcement learning under a variational perspective. The connection leads to (1) probability ratio clipping that regularizes generator training to prevent excessively large updates, and (2) a sample re-weighting mechanism that stabilizes discriminator training by downplaying bad-quality fake samples. We provide theoretical analysis on the convergence of our approach. By plugging the training approach in diverse state-of-the-art GAN architectures, we obtain significantly improved performance over a range of tasks, including text generation, text style transfer, and image generation¹

1 Introduction

Generative adversarial networks (GANs) [16] have drawn great research interests and achieved remarkable success in image synthesis [4, 40], video generation [33], and others. However, it is usually hard to train a GAN well, because the training process is commonly unstable, subject to disturbances and even collapses. To alleviate this issue, substantial efforts have been paid to improve the training stability from different perspectives, e.g., divergence minimization [37, 38], Wasserstein distance with Lipschitz continuity of the discriminator [2, 18, 51], energy-based models [3, 56], etc.

In spite of the above progresses, the instability in training has not been well resolved [9], since it is difficult to well balance the strength of the generator and the discriminator. What is worse, such an instability issue is exacerbated in text generation due to the sequential and discrete nature of text [6, 14, 36]. Specifically, the high sensitivity of text generation to noise and the underlying errors caused by sparse discriminator signals in the generated text can often result in destructive updates to both generator and discriminator, enlarging the instability in GANs.

In this work, we develop a novel variational GAN training framework to improve the training stability, which is broadly applicable to GANs of a variety of architectures for image and text generation. This training framework is derived from a variational perspective of GANs and the resulting connections to reinforcement learning (in particular, RL-as-inference) [1, 30, 45] and other rich literature [5, 17, 25]. Specifically, our approach consists of two stabilization techniques, namely, probability ratio clipping and sample re-weighting, for stabilizing the generator and discriminator respectively. (1) Under the new variational perspective, the generator update is subject to a KL penalty on the change of the generator distribution. This KL penalty closely resembles that in the popular Trust-Region Policy Optimization (TRPO) [44] and its related Proximal Policy Optimization (PPO) [45]. This connection motivates a simple surrogate objective with a clipped probability ratio between the new generator and the old one. The probability ratio clipping discourages excessively large generator updates, and has

¹Code available at: github.com/Holmeswww/PPOGAN

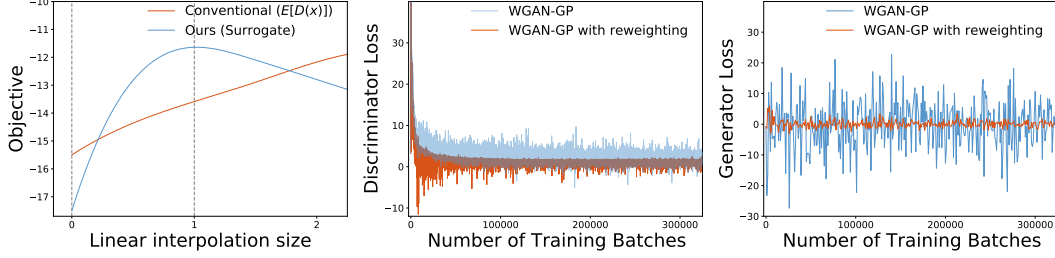


Figure 1: Illustration of the proposed approach for stabilizing GAN training. Results are from the CIFAR-10 experiment in Sec. 4.1. **Left:** The conventional and surrogate objectives for generator training, as we interpolate between the initial generator parameters θ_{old} and the updated generator parameters θ_{new} which we compute after one iteration of training. The updated θ_{new} obtains maximal surrogate objective. The surrogate objective imposes a penalty for having too large of a generator update, since the curve starts decreasing after $x = 1$. In contrast, the conventional objective (for WGAN-GP) keeps increasing with larger generator updates. **Middle and right:** Discriminator and generator losses w/ and w/o sample re-weighting. WGAN-GP with our re-weighting plugged in shows lower variance in both discriminator and generator losses throughout training (and achieves better final performance as shown in Sec. 4.1).

shown to be effective in the context of stabilizing policy optimization [45]. Figure 1 (left) shows the intuition about the surrogate objective, where we can observe the objective decreases with an overly large generator change and thus imposes regularization on the updates.

(2) When updating the discriminator, the new perspective induces an importance sampling mechanism, which effectively re-weights fake samples by their discriminator scores. Since low-quality samples tend to receive smaller weights, the discriminator trained on the re-weighted samples is more likely to maintain stable performance, and in turn provide informative gradients for subsequent generator updates. Figure 1 (middle/right) demonstrates the effect of the re-weighting in reducing the variance of both discriminator and generator losses. Similar importance weighting methods have recently been used in other contexts, such as de-biasing generative models [17] and sampling from energy-based models [12]. Our derivations can be seen as a variant for the new application of improving GANs.

We give theoretical convergence analysis showing the generator receives informative gradient [57] from the generator. Empirically, we conduct extensive experiments on a range of tasks, including text generation, text style transfer, and image generation. Our approach shows significant improvement over state-of-the-art methods, demonstrating its broad applicability and efficacy.

2 Related Work

Wasserstein distance, WGAN, and Lipschitz continuity. The GAN framework [16] features two components: a generator G_θ that synthesizes samples x given some noise source z , namely $x = G_\theta(z)$ with $z \sim p_z(z)$, and a discriminator that distinguishes generator’s output and real data, which provides gradient feedback to improve the generator’s performance. WGAN [2] improves the training stability of GANs by minimizing the Wasserstein distance $W(p_r, p_\theta)$ between the generation distribution p_θ (induced from G_θ) and the real data distribution p_r . Its training loss is formulated as:

$$\min_{\theta} \max_{f \in \mathcal{D}} \mathbb{E}_{x \sim p_r} [f(x)] - \mathbb{E}_{x \sim p_\theta} [f(x)], \quad (1)$$

where \mathcal{D} is the set of 1-Lipschitz functions; f acts as the discriminator and is usually implemented by a neural network f_ϕ . The original resort to enforce the Lipschitz constraint is through weight clipping [2]. WGAN-GP [18] later improves it by replacing it with a gradient penalty on the discriminator. CT-GAN [51] further imposes the Lipschitz continuity constraint on the manifold of the real data $x \sim p_r$. Our approach is orthogonal to these prior works and can serve as a drop-in replacement to stabilize generator and discriminator in various kinds of GANs, such as WGAN-GP and CT-GAN.

Research on the Lipschitz continuity of GAN discriminators have resulted in the theory of “informative gradients” [57, 58]. Under certain mild conditions, a Lipschitz discriminator can provide informative gradient to the generator in a GAN framework: when p_θ and p_r are disjoint, the gradient $\nabla f^*(x)$ of optimal discriminator f^* w.r.t each sample $x \sim p_\theta$ points to a sample $x^* \sim p_r$, which guarantees that the generation distribution p_θ is moving towards p_r . We extend the informative gradient theory to our new case and show convergence of our approach.

Reinforcement Learning as Inference. Casting RL as probabilistic inference has a long history of research [1, 10, 11, 30, 41]. For example, Abdolmaleki et al. [1] introduced maximum a-posteriori policy optimization from a variational perspective. TRPO [44] is closely related to this line by using a KL divergence regularizer to stabilize standard RL objectives. PPO [45] further proposed a practical clipped surrogate objective that emulates the regularization. Our approach draws on the connections to the research, particularly the variational perspective and PPO, to improve GAN training.

Other related work. Importance re-weighting has been adopted in different problems, such as improving variational auto-encoders [5], de-biasing generative models [17], learning knowledge constraints [25], etc. We derive from the variational perspective which leads to re-weighting in the new context of discriminator stabilization. Our approach is orthogonal to and can be combined with other stabilization techniques such as large-batch training [4] and parameter averaging (EMA) [4, 53].

3 Improving GAN Training

3.1 Motivations

Our approach is motivated by connecting GAN training with the well-established RL-as-inference methods [1, 30] under a variational perspective. The connections enable us to augment GAN training with existing powerful probabilistic inference tools as well as draw inspirations from the rich RL literature for stable training. In particular, the connection to the popular TRPO [44] and PPO [45] yields the probability ratio clipping in generator training that avoids destructive updates (Sec. 3.2), and the application of importance sampling estimation gives rise to sample re-weighting for adaptive discriminator updates (Sec. 3.3). The full training procedure is then summarized in Alg.1.

Specifically, as described in Sec. 2, the conventional formulation e.g., WGAN [2] for updating the generator $p_\theta(\mathbf{x})$ maximizes the expected discriminator score: $\mathbb{E}_{p_\theta}[f_\phi(\mathbf{x})]$, where f_ϕ is the Lipschitz-continuous discriminator parameterized with ϕ . The objective straightforwardly relates to policy optimization in RL by seeing p_θ as a policy and f_ϕ as a reward function. Thus, inspired by the probabilistic inference formulations of policy optimization [1, 13, 25], here we transform the conventional objective by introducing a non-parametric auxiliary distribution $q(\mathbf{x})$ and defining a new variational objective:

$$\mathcal{L}(\theta, q) = \mathbb{E}_q[f_\phi(\mathbf{x})] - \text{KL}(q(\mathbf{x})\|p_\theta(\mathbf{x})), \quad (2)$$

where KL is the KL divergence. Intuitively, we are maximizing the expected discriminator score of the auxiliary q (instead of generator p_θ), and meanwhile encouraging the generator to stay close to q .

As we shall see in more details shortly, the new formulation allows us to take advantage of off-the-shelf inference methods, which naturally leads to new components to improve the GAN training. In particular, maximizing the above objective is solved by the expectation maximization (EM) algorithm [35] which alternatingly optimizes q at E-step and optimizes θ at M-step. More specifically, at each iteration t , given the current status of $\theta = \theta^{(t)}$, the E-step that maximizes $\mathcal{L}(\theta^{(t)}, q)$ w.r.t q has a closed-form solution:

$$q^{(t)}(\mathbf{x}) = \frac{p_{\theta^{(t)}}(\mathbf{x}) \exp\{f_\phi(\mathbf{x})\}}{Z_\phi}, \quad (3)$$

where $Z_\phi = \int_{\mathbf{x}} p_{\theta^{(t)}}(\mathbf{x}) \exp\{f_\phi(\mathbf{x})\} d\mathbf{x}$ is the normalization term depending on the discriminator parameters ϕ . We elaborate on the M-step in the following, where we continue to develop the practical procedures for updating the generator and the discriminator, respectively.

3.2 Generator Training with Probability Ratio Clipping

The M-step optimizes $\mathcal{L}(\theta, q^{(t)})$ w.r.t θ , which is equivalent to minimizing the KL divergence term in Eq.(2). However, since the generator p_θ in GANs is often an *implicit* distribution that does not permit evaluating likelihood, the above KL term (which involves evaluating the likelihood of samples from q) is not applicable. We adopt an approximation, which has also been used in the classical wake-sleep algorithm [21] and recent work [25], by minimizing the *reverse* KL divergence as below. With Eq.(3) plugged in, we have:

$$\min_{\theta} \text{KL}(p_\theta(\mathbf{x})\|q^{(t)}(\mathbf{x})) = \min_{\theta} -\mathbb{E}_{p_\theta}[f_\phi(\mathbf{x})] + \text{KL}(p_\theta(\mathbf{x})\|p_{\theta^{(t)}}(\mathbf{x})). \quad (4)$$

As proven in the appendix, approximating with the reverse KL does not change the optimization problem. The first term on the right-hand side of the equation recovers the conventional objective of updating the generator. Of particular interest is the second term, which is a new KL regularizer between the generator p_θ and its “old” state $p_{\theta^{(t)}}$ from the previous iteration. The regularizer discourages the generator from changing too much between updates, which is useful to stabilize the stochastic optimization procedure. The regularization closely resembles to that of TRPO/PPO, where a similar KL regularizer is imposed to prevent uncontrolled policy updates and make policy gradient robust to noises. Sec. 3.4 gives convergence analysis on the KL-regularized generator updates.

In practice, directly optimizing with the KL regularizer can be infeasible due to the same difficulty with the implicit distribution as above. Fortunately, PPO [45] has presented a simplified solution that emulates the regularized updates using a clipped surrogate objective, which is widely-used in RL. We adapt the solution to our context, leading to the following practical procedure of generator updates.

Probability Ratio Clipping. Let r_t denote the probability ratio $r_t(\theta) = \frac{p_\theta(\mathbf{x})}{p_{\theta^{(t)}}(\mathbf{x})}$ which measures the difference between the new and old distributions. For instance, $r_t(\theta^{(t)}) = 1$. The clipped surrogate objective for updating the generator, as adapted from PPO, is:

$$\mathcal{L}^{CLIP}(\theta) = \mathbb{E}_{p_\theta} \left[\min \left(r_t(\theta) f_\phi(\mathbf{x}), r_t^{clip}(\theta) f_\phi(\mathbf{x}) \right) \right], \quad (5)$$

where $r_t^{clip}(\theta) = \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)$ clips the probability ratio, so that moving $r_t(\theta)$ outside of the interval $[1 - \epsilon, 1 + \epsilon]$ is discouraged. Taking the minimum puts a ceiling on the increase of the objective. Thus the generator does not benefit by going far away from the old generator.

Finally, to estimate the probability ratio $r_t(\theta)$ when p_θ is implicit, we use an efficient approximation similar to [7, 17] by introducing a binary classifier C trained to distinguish real and generated samples. Assuming an optimal classifier C [7, 16] which has $p_\theta(\mathbf{x}) = \frac{1-C(\mathbf{x})}{C(\mathbf{x})} p_r(\mathbf{x})$, we then approximate r_t by keeping a historical version (C) of $C^{(t)}$ at time-step t :

$$r_t(\theta) = \frac{p_\theta(\mathbf{x})}{p_{\theta^{(t)}}(\mathbf{x})} \approx \frac{(1 - C(\mathbf{x})) \cdot C^{(t)}(\mathbf{x})}{(1 - C^{(t)}(\mathbf{x})) \cdot C(\mathbf{x})}. \quad (6)$$

Note that after plugging the rightmost expression into Eq.(5), gradient can propagate through C to θ since $\mathbf{x} = G_\theta(\mathbf{z})$. In practice, during the phase of generator training, we maintain C by fine-tuning it for one iteration every time after θ is updated (Alg.1). Thus the maintenance of C is cheap. We give more details of C in the appendix.

3.3 Discriminator Training with Sample Re-weighting

We next discuss the training of the discriminator f_ϕ , where we augment the conventional training with an importance weighting mechanism for adaptive updates. Concretely, given the form of the auxiliary distribution solution $q^{(t)}$ in Eq.(3), we first draw from the recent energy-based modeling work [8, 12, 25, 27] and propose to maximize the data log-likelihood of $q^{(t)}$ w.r.t ϕ . By taking the gradient, we derive the loss ($\mathcal{L}(\phi)$) for the discriminator (f_ϕ):

$$\nabla_\phi \mathcal{L}(\phi) = \nabla_\phi \left(\mathbb{E}_{p_r} [f_\phi(\mathbf{x})] - \log Z_\phi \right) = \mathbb{E}_{p_r} [\nabla_\phi f_\phi(\mathbf{x})] - \mathbb{E}_{q^{(t)}} [\nabla_\phi f_\phi(\mathbf{x})]. \quad (7)$$

We can observe that the resulting form resembles the conventional one (Sec. 2) as we are essentially maximizing f_ϕ on real data while minimizing f_ϕ on fake samples. An important difference is that here fake samples are drawn from the auxiliary distribution $q^{(t)}$ instead of the generator p_θ . This difference leads to the new sample re-weighting component as below. Note that, as in WGAN (Sec.2), we maintain f_ϕ to be from the class of 1-Lipschitz functions, which is necessary for the convergence analysis in Sec.3.4. In practice, we can use gradient penalty following [18, 51].

Sample Re-weighting. We use the tool of importance sampling to estimate the expectation under $q^{(t)}$ in Eq.(7). Given the multiplicative form of $q^{(t)}$ in Eq.(3), similar to [1, 12, 25], we use the generator $p_{\theta^{(t)}}$ as the proposal distribution. This leads to

$$\mathbb{E}_{q^{(t)}} [\nabla_\phi f_\phi(\mathbf{x})] = \mathbb{E}_{p_{\theta^{(t)}}} [\exp\{f_\phi(\mathbf{x})\} \cdot \nabla_\phi f_\phi(\mathbf{x})] / Z_\phi. \quad (8)$$

Note that Z_ϕ is the normalization factor defined in Eq.(3). That is, fake samples from the generator are weighted by the exponentiated discriminator score when used to update the discriminator. Intuitively,

the mechanism assigns higher weights to samples that can fool the discriminator better, while low-quality samples are downplayed to avoid destructing the discriminator performance. It is worth mentioning that similar importance weighting scheme has been used in [7, 24] for generator training in GANs, and [5] for improving variational auto-encoders. Our work instead results in a re-weighting scheme in the new context of discriminator training.

Alg.1 summarizes the proposed training procedure for the generator and discriminator.

Algorithm 1 GAN Training with Probability Ratio Clipping and Sampling Re-weighting

```

1: Initialize the generator  $p_\theta$ , the discriminator  $f_\phi$ , and the auxiliary classifier  $C$ 
2: for  $t \leftarrow 1$  to  $T$  do
3:   for certain number of steps do
4:     Update the discriminator  $f_\phi$  with sample re-weighting through Eqs.(7)-(8), and maintain  $f_\phi$ 
       to have upper-bounded Lipschitz constant through, e.g., gradient penalty [18].
5:   end for
6:   for certain number of steps do
7:     Finetune the real/fake classifier  $C$  (for 1 step)
8:     Estimate probability ratio  $r_t(\theta)$  using  $C$  through Eq.(6)
9:     Update the generator  $p_\theta$  with probability ratio clipping through Eq.(5)
10:  end for
11: end for

```

3.4 Theoretical Analysis

In this section, we show that our framework produces informative gradient [57] that points the generator distribution p_θ toward the real data distribution p_r assuming an optimal discriminator. The analysis is based on the reverse KL updates for the generator (Eq.4), while the probability ratio clipping serves as a practical emulation for the updates. We begin by adapting Proposition 1 in Gulrajani et al. [18] to our problem:

Proposition 3.1. *Let p_r and p_θ be two distributions in X , a compact metric space. Then, there is a 1-Lipschitz function f^* which is the optimal solution of*

$$\max_{\|f_\phi\|_L \leq 1} \mathbb{E}_{\mathbf{x} \sim p_r} [f_\phi(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim q} [f_\phi(\mathbf{x})]$$

Let π be the optimal coupling between p_r and q , defined as the minimizer of: $W(p_r, q) = \inf_{\pi \in \Pi(p_r, q)} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \pi} [\|\mathbf{x} - \mathbf{y}\|]$ where $\Pi(p_r, p_\theta)$ is the set of joint distributions $\pi(\mathbf{x}, \mathbf{y})$ whose marginals are p_r and q , respectively. Then, if f^ is differentiable, $\pi(\mathbf{x} = \mathbf{y}) = 0$, and $\mathbf{x}_\tau = \tau \mathbf{x} + (1 - \tau) \mathbf{y}$ with $0 \leq \tau \leq 1$, it holds that $p_{(\mathbf{x}, \mathbf{y}) \sim \pi} \left[\nabla f^*(\mathbf{x}_\tau) = \frac{\mathbf{y} - \mathbf{x}_\tau}{\|\mathbf{y} - \mathbf{x}_\tau\|} \right] = 1$.*

The proposition shows that the optimal f^* provides informative gradient [58] from q towards p_r . We then generalize the conclusion to p_θ by considering correlation between q and p_θ .

By the definition of q with respect to p_θ in Equation (3), the support of p_θ and q are the same; namely, given $\mathbf{x} \sim p_\theta, \mathbf{y} \sim p_r$, we also have $q(\mathbf{x}) \neq 0$. Therefore, for all $\mathbf{x} \sim p_\theta$, \mathbf{x} is also a valid sample from q , the f^* in Proposition 3.1 provides informative gradient with respect to $\mathbf{x}_\tau = \tau \mathbf{x} + (1 - \tau) \mathbf{y}, \forall \tau \in [0, 1]$:

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \pi} \left[\nabla f^*(\mathbf{x}_\tau) = \frac{\mathbf{y} - \mathbf{x}_\tau}{\|\mathbf{y} - \mathbf{x}_\tau\|} \right] = 1$$

Therefore, assuming f_ϕ^* is the optimal discriminator to (7), optimizing Eq.(4) can provide informative gradient that points the generator toward p_r .

4 Experiments

We conduct extensive experiments on three unsupervised generation tasks, including image generation, text generation, and text style transfer. The three tasks apply GANs to model different data modalities, namely image, text, and neural hidden representations. Our approach consistently offers improvement over the state-of-the-arts on all tasks. See appendix for experimental details.

Method	IS (\uparrow)	FID (\downarrow)
Real data	11.24 \pm .12	7.8
WGAN-GP (2017)	7.86 \pm .08	-
CT-GAN (2018)	8.12 \pm .12	-
SN-GANs (2018)	8.22 \pm .05	21.7 \pm .21
WGAN-ALP (2020)	8.34 \pm .06	12.96 \pm .35
SRNGAN (2020)	8.53 \pm .04	19.83
AutoGAN (2019)	8.55 \pm .10	12.42
Ours (re-weighting only)	8.45 \pm .14	13.21 \pm .60
Ours (full)	8.69\pm.13	10.70\pm.10

Table 1: CIFAR-10 results. Our method is run 3 times for average and standard deviation.



Figure 2: Generated samples by WGAN-GP (top-left), CT-GAN (bottom-left), and ours (right).

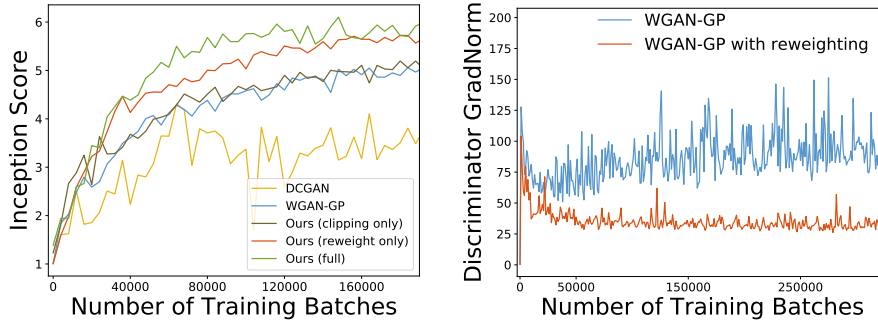


Figure 3: **Left:** Inception score on CIFAR-10 v.s. training iterations. The DCGAN [40] architecture is used. **Right:** The gradient norms of discriminators on fake samples.

4.1 Image generation

We first use the popular CIFAR-10 benchmark for evaluation and in-depth analysis of our approach.

Setup. CIFAR-10 [29] contains 50K images of sizes 32×32 . Following the setup in CT-GAN [51] and its public implementation², we use a residual architecture to implement both generator and discriminator, and also impose a Lipschitz constraint on the discriminator. For each iteration, we update both generator and discriminator for 5 times. We use Inception Score (IS) [42] and Frechet Inception Distance (FID) [20] as our evaluation metrics. Among them, IS evaluates the quality and diversity of generated images, and FID captures model issues, e.g., mode collapse [52].

Results. Table 1 reports the results on CIFAR-10. For the four latest methods, SN-GANs [34] introduced spectral normalization to stabilize the discriminator training; WGAN-ALP [48] developed an explicit Lipschitz penalty; WGAN-ALP [43] introduced a weight-normalization scheme for generalization; and AutoGAN [15] incorporated neural architecture search for the generator architecture. From Table 1, one can observe that our full approach (CT-GAN + discriminator sample re-weighting + generator probability ratio clipping) achieves the best performance, with both IS and FID significantly surpassing the baselines. These results accord with the visual results in Figure 2 where our generated samples show higher visual quality than those of the baselines. Moreover, comparison between CT-GAN and our approach with only discriminator re-weighting shows significant improvement. By further adding the probability ratio clipping to arrive our full approach, the performance in terms of both IS and FID is further improved with a large margin. The results demonstrate the effectiveness of the two components in our approach.

Figure 1 in Sec. 1 has shown the effects of the proposed approach in stabilizing the generator and discriminator training. Here we further analyze these two components. Figure 3 (left) shows the convergence curves of different GAN methods. For a fair comparison, all models use the same DCGAN architecture [40], and both our approach and WGAN-GP [18] enforce the same discriminator Lipschitz constraint. From Figure 3 (left), one can observe that our full approach surpasses our approach with only sample re-weighting, and they both converge faster and achieve a higher IS score than clipping-only, WGAN-GP, and DCGAN. Note that clipping-only does not offer a performance improvement but does not slow down training. Figure 3 (right) further looks into how the re-weighting

²https://github.com/igul222/improved_wgan_training

Length	MLE	SeqGAN [54]	LeakGAN [19]	RelGAN [36]	WGAN-GP [18]	Ours	Real
20	9.038	8.736	7.038	6.680	6.89	5.67	5.750
40	10.411	10.310	7.191	6.765	6.78	6.14	4.071

Table 2: Oracle negative log-likelihood scores (\downarrow) on synthetic data.

Method	BLEU-2 (\uparrow)	BLEU-3 (\uparrow)	BLEU-4 (\uparrow)	BLEU-5 (\uparrow)	NLL _{gen} (\downarrow)
MLE	0.768	0.473	0.240	0.126	2.382
SeqGAN [54]	0.777	0.491	0.261	0.138	2.773
LeakGAN [19]	0.826	0.645	0.437	0.272	2.356
RelGAN (100) [36]	0.881	0.705	0.501	0.319	2.482
RelGAN (1000) [36]	0.837	0.654	0.435	0.265	2.285
WGAN-GP [18]	0.872	0.636	0.379	0.220	2.209
Ours	0.905	0.692	0.470	0.322	2.265

Table 3: Results on EMNLP2017 WMT News. BLEU measures text quality and NLL_{gen} evaluates sample diversity. We copied the results of previous text GAN models from [36], where RelGAN (100) and RelGAN (1000) use different hyper-parameter for gumbel-softmax.

on fake samples can affect the discriminator training. It shows that by injecting sample re-weighting into WGAN-GP, its gradients on fake samples become more stable and show much lower variance, which also partially explains the higher training stability of discriminator in Figure 1.

4.2 Text Generation

In this section, we evaluate our approach on text generation, a task that is known to be notoriously difficult for GANs due to the discrete and sequential nature of text.

Setup. We implement our approach based on the RelGAN [36] architecture, a state-of-the-art GAN model for text generation. Specifically, we replace the generator and discriminator objectives in RelGAN with ours. We follow WGAN-GP [18] and impose discriminator Lipschitz constraint with gradient penalty. Same as [36], we use Gumbel-softmax approximation [26, 32] on the discrete text to enable gradient backpropagation, and the generator is initialized with maximum likelihood (MLE) pre-training. Our implementation is based on the public PyTorch code of RelGAN³. Same as previous studies [19, 36, 54], we evaluate our approach on both synthetic and real text datasets.

Results on Synthetic Data. The synthetic data consists of 10K discrete sequences generated by an oracle-LSTM with fixed parameters [54]. This setup facilitates evaluation, as the quality of generated samples can be directly measured by the negative log-likelihood (NLL) of the oracle on the samples. We use synthetic data with sequence lengths 20 and 40, respectively. Table 2 reports the results. MLE is the baseline with maximum likelihood training, whose output model is used to initialize the generators of GANs. Besides the previous text generation GANs [19, 36, 54], we also compare with WGAN-GP which uses the same neural architecture as RelGAN and ours. From Table 2, one can observe that our approach significantly outperforms all other approaches on both synthetic sets. Our improvement over RelGAN and WGAN-GP demonstrates that our proposed generator and discriminator objectives are more effective than the previous ones.

Results on Real Data. We then evaluate our method on the EMNLP2017 WMT News, a large real text data used for text GAN studies [19, 36]. The dataset consists of 270K/10K training/test sentences with a maximum length of 51 and a vocabulary size of 5,255. To measure the generation quality, we use the popular BLEU- n metric which measures n -gram overlap between generated and real text ($n \in \{2, 3, 4, 5\}$). To evaluate the diversity of generation, we use the negative log-likelihood of the generator on the real test set (NLL_{gen}) [19, 36]. From the results in Table 3, one can see that our approach shows comparable performance with the previous best model RelGAN (100) in terms of text quality (BLEU), but has better sample diversity. We use the same gumbel-softmax parameter as RelGAN (1000). Our model also achieves much higher BLEU scores than WGAN-GP (e.g., 0.322 v.s. 0.220 on BLEU-5), demonstrating its ability of generating higher-quality samples.

³<https://github.com/williamSYSU/TextGAN-PyTorch>

Method	BLEU
Zhang et al. [55]	24.48
Tian et al. [49]	24.90
Subramanian et al. [47]	31.20
Tikhonov et al. [50]	32.82
Ours	33.45\pm.95

Table 4: BLEU scores between model generations and human-written text on the Yelp data. We run our method for 5 times and report the average and standard deviation.

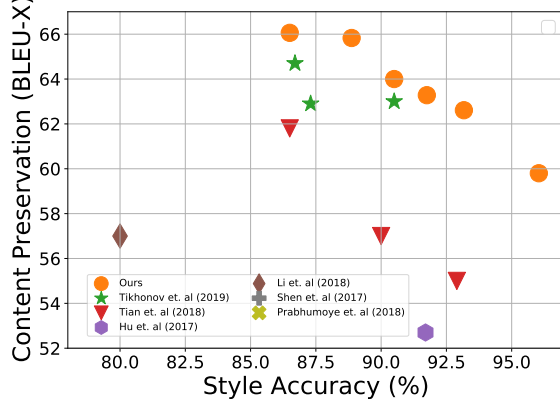


Figure 4: Trade-off between style accuracy and content preservation. The orange circles denote our results using varying values for an objective weight [50] which manages the trade-off.

4.3 Text Style Transfer

We further apply our approach on the text style transfer task which is gaining increasing attention in NLP [23, 46]. The task aims at rewriting a sentence to modify its style (e.g., sentiment) while preserving the content. Previous work applies GANs on neural hidden states to learn disentangled representations [46, 50]. The task thus can serve as a good benchmark task for GANs, as hidden state modeling provides a new modality that differs from image and text modeling as studied above.

Setup. We follow the same experimental setting and use the same model architecture in the latest work [50]. In particular, Tikhonov et al. [50] extended the variational autoencoder based model [23, 28] by adding a latent code discriminator which eliminates stylistic information in the latent code. We replace their adversarial objectives with our proposed ones, and impose discriminator Lipschitz constraint with gradient penalty [18]. Our implementation is based on the public code⁴ released in [50]. We test our approach on sentiment transfer, in which the sentiment (positive or negative) is treated as the style of the text. We use the standard Yelp review dataset⁵, and the human written output text provided by [31] as the ground truth for evaluation.

Results. Following the previous work [50], we first report the BLEU score that measures the similarity of the generated samples against the human written text. Table 4 shows that our approach achieves best performance, improving the state-of-the-art result [50] from BLEU 32.82 to 33.45.

The second widely used evaluation method is to measure (1) the style accuracy by applying a pre-trained style classifier on generated text, and (2) the content preservation by computing the BLEU score between the generated text and the original input text (BLEU-X). There is often a trade-off between the two metrics. Figure 4 displays the trade-off by different models. Our results locate on the top-right corner, indicating that our approach achieves the best overall style-content trade-off.

5 Conclusion

We have presented a new training framework of GANs derived from a new variational perspective and draws on rich connections with RL-as-inference. This results in probably ratio clipping for generator updates to discourage overly large changes, and fake sample re-weighting for discriminator updates to stabilize training. Experiments show our approach achieves better results than previous best methods on image generation, text generation, and text style transfer. Our approach also shows more stable training. We are interested in exploring more connections between GANs and other learning paradigms to inspire more techniques for improved GAN training.

⁴https://github.com/VAShibae/text_style_transfer

⁵www.yelp.com/dataset

Broader Impacts

This work offers a unique viewpoint on two promising fields with lots of applications and impacts: Generative Adversarial Networks and Reinforcement Learning. The improvement to image generation results may be adapted to speed up photo editing, improve scene rendering, and create more realistic simulation for robot training. Furthermore, the contribution to text generation and text style transfer can be adopted to improve the quality of machine translation, and automated news-summaries.

Nevertheless, GANs can also be applied to faking images of people and jeopardize personal identities (i.e. Deepfake). We hope that future works can counter this issue through deep-fake detection.

References

- [1] Abdolmaleki, A., J. T. Springenberg, Y. Tassa, R. Munos, N. Heess, and M. Riedmiller (2018). Maximum a posteriori policy optimisation. In *ICLR*.
- [2] Arjovsky, M., S. Chintala, and L. Bottou (2017). Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pp. 214–223.
- [3] Berthelot, D., T. Schumm, and L. Metz (2017). Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*.
- [4] Brock, A., J. Donahue, and K. Simonyan (2018). Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.
- [5] Burda, Y., R. Grosse, and R. Salakhutdinov (2015). Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*.
- [6] Caccia, M., L. Caccia, W. Fedus, H. Larochelle, J. Pineau, and L. Charlin (2020). Language gans falling short. In *ICLR*.
- [7] Che, T., Y. Li, R. Zhang, R. D. Hjelm, W. Li, Y. Song, and Y. Bengio (2017). Maximum-likelihood augmented discrete generative adversarial networks. *arXiv preprint arXiv:1702.07983*.
- [8] Che, T., R. Zhang, J. Sohl-Dickstein, H. Larochelle, L. Paull, Y. Cao, and Y. Bengio (2020). Your GAN is secretly an energy-based model and you should use discriminator driven latent sampling. *arXiv preprint arXiv:2003.06060*.
- [9] Chu, C., K. Minami, and K. Fukumizu (2020). Smoothness and stability in gans. In *ICLR*.
- [10] Dayan, P. and G. E. Hinton (1997). Using expectation-maximization for reinforcement learning. *Neural Computation* 9(2), 271–278.
- [11] Deisenroth, M. P., G. Neumann, J. Peters, et al. (2013). A survey on policy search for robotics. *Foundations and Trends® in Robotics* 2(1–2), 1–142.
- [12] Deng, Y., A. Bakhtin, M. Ott, A. Szlam, and M. Ranzato (2020). Residual energy-based models for text generation. In *ICLR*.
- [13] Ding, N. and R. Soricut (2017). Cold-start reinforcement learning with softmax policy gradient. In *NeurIPS*.
- [14] Fedus, W., I. Goodfellow, and A. M. Dai (2018). MaskGAN: better text generation via filling in the_. In *ICLR*.
- [15] Gong, X., S. Chang, Y. Jiang, and Z. Wang (2019). Autogan: Neural architecture search for generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3224–3234.
- [16] Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio (2014). Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680.

- [17] Grover, A., J. Song, A. Kapoor, K. Tran, A. Agarwal, E. J. Horvitz, and S. Ermon (2019). Bias correction of learned generative models using likelihood-free importance weighting. In *NeurIPS*.
- [18] Gulrajani, I., F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville (2017). Improved training of wasserstein gans. In *Advances in neural information processing systems*, pp. 5767–5777.
- [19] Guo, J., S. Lu, H. Cai, W. Zhang, Y. Yu, and J. Wang (2018). Long text generation via adversarial training with leaked information. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [20] Heusel, M., H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter (2017). Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pp. 6626–6637.
- [21] Hinton, G. E., P. Dayan, B. J. Frey, and R. M. Neal (1995). The "wake-sleep" algorithm for unsupervised neural networks. *Science* 268(5214), 1158–1161.
- [22] Hu, Z., H. Shi, B. Tan, W. Wang, Z. Yang, T. Zhao, J. He, L. Qin, D. Wang, X. Ma, et al. (2019). Texar: A modularized, versatile, and extensible toolkit for text generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 159–164.
- [23] Hu, Z., Z. Yang, X. Liang, R. Salakhutdinov, and E. P. Xing (2017). Toward controlled generation of text. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1587–1596. JMLR. org.
- [24] Hu, Z., Z. Yang, R. Salakhutdinov, and E. P. Xing (2018). On unifying deep generative models. In *ICLR*.
- [25] Hu, Z., Z. Yang, R. R. Salakhutdinov, L. Qin, X. Liang, H. Dong, and E. P. Xing (2018). Deep generative models with learnable knowledge constraints. In *Advances in Neural Information Processing Systems*, pp. 10501–10512.
- [26] Jang, E., S. Gu, and B. Poole (2017). Categorical reparameterization with gumbel-softmax. In *ICLR*.
- [27] Kim, T. and Y. Bengio (2016). Deep directed generative models with energy-based probability estimation. *arXiv preprint arXiv:1606.03439*.
- [28] Kingma, D. P. and M. Welling (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- [29] Krizhevsky, A. and G. Hinton (2010). Convolutional deep belief networks on cifar-10. *Unpublished manuscript* 40(7), 1–9.
- [30] Levine, S. (2018). Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*.
- [31] Li, J., R. Jia, H. He, and P. Liang (2018). Delete, retrieve, generate: A simple approach to sentiment and style transfer. In *Proceedings of NAACL-HLT*, pp. 1865–1874.
- [32] Maddison, C. J., A. Mnih, and Y. W. Teh (2017). The concrete distribution: A continuous relaxation of discrete random variables.
- [33] Mathieu, M., C. Couprie, and Y. LeCun (2015). Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*.
- [34] Miyato, T., T. Kataoka, M. Koyama, and Y. Yoshida (2018). Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*.
- [35] Neal, R. M. and G. E. Hinton (1998). A view of the em algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pp. 355–368. Springer.
- [36] Nie, W., N. Narodytska, and A. Patel (2018). Relgan: Relational generative adversarial networks for text generation.

- [37] Nock, R., Z. Cranko, A. K. Menon, L. Qu, and R. C. Williamson (2017). f-GANs in an information geometric nutshell. In *Advances in Neural Information Processing Systems*, pp. 456–464.
- [38] Nowozin, S., B. Cseke, and R. Tomioka (2016). f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pp. 271–279.
- [39] Odena, A., C. Olah, and J. Shlens (2017). Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2642–2651. JMLR. org.
- [40] Radford, A., L. Metz, and S. Chintala (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- [41] Rawlik, K., M. Toussaint, and S. Vijayakumar (2013). On stochastic optimal control and reinforcement learning by approximate inference. In *IJCAI*.
- [42] Salimans, T., I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen (2016). Improved techniques for training gans. In *Advances in neural information processing systems*, pp. 2234–2242.
- [43] Sanyal, A., P. H. Torr, and P. K. Dokania (2020). Stable rank normalization for improved generalization in neural networks and gans. In *ICLR*.
- [44] Schulman, J., S. Levine, P. Abbeel, M. Jordan, and P. Moritz (2015). Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897.
- [45] Schulman, J., F. Wolski, P. Dhariwal, A. Radford, and O. Klimov (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- [46] Shen, T., T. Lei, R. Barzilay, and T. Jaakkola (2017). Style transfer from non-parallel text by cross-alignment. In *Advances in neural information processing systems*, pp. 6830–6841.
- [47] Subramanian, S., G. Lample, E. M. Smith, L. Denoyer, M. Ranzato, and Y.-L. Boureau (2018). Multiple-attribute text style transfer. *arXiv preprint arXiv:1811.00552*.
- [48] Terjék, D. (2020). Adversarial lipschitz regularization. In *ICLR*.
- [49] Tian, Y., Z. Hu, and Z. Yu (2018). Structured content preservation for unsupervised text style transfer. *arXiv preprint arXiv:1810.06526*.
- [50] Tikhonov, A., V. Shibaev, A. Nagaev, A. Nugmanova, and I. P. Yamshchikov (2019). Style transfer for texts: to err is human, but error margins matter. In *EMNLP*.
- [51] Wei, X., B. Gong, Z. Liu, W. Lu, and L. Wang (2018). Improving the improved training of wasserstein GANs: A consistency term and its dual effect. *arXiv preprint arXiv:1803.01541*.
- [52] Xu, Q., G. Huang, Y. Yuan, C. Guo, Y. Sun, F. Wu, and K. Weinberger (2018). An empirical study on evaluation metrics of generative adversarial networks. *arXiv preprint arXiv:1806.07755*.
- [53] Yaz, Y., C.-S. Foo, S. Winkler, K.-H. Yap, G. Piliouras, V. Chandrasekhar, et al. (2018). The unusual effectiveness of averaging in gan training. In *International Conference on Learning Representations*.
- [54] Yu, L., W. Zhang, J. Wang, and Y. Yu (2017). Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [55] Zhang, Z., S. Ren, S. Liu, J. Wang, P. Chen, M. Li, M. Zhou, and E. Chen (2018). Style transfer as unsupervised machine translation. *arXiv preprint arXiv:1808.07894*.
- [56] Zhao, J., M. Mathieu, and Y. LeCun (2016). Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*.

- [57] Zhou, Z., J. Liang, Y. Song, L. Yu, H. Wang, W. Zhang, Y. Yu, and Z. Zhang (2019). Lipschitz generative adversarial nets. *arXiv preprint arXiv:1902.05687*.
- [58] Zhou, Z., Y. Song, L. Yu, H. Wang, J. Liang, W. Zhang, Z. Zhang, and Y. Yu (2018). Understanding the effectiveness of lipschitz-continuity in generative adversarial nets. *arXiv preprint arXiv:1807.00751*.

6 Appendix

6.1 Proof on the equivalence between Reverse KL Divergence and KL Divergence

We prove that optimizing $\mathbf{KL}(p_\theta||q)$ are equivalent to optimizing $\mathbf{KL}(q||p_\theta)$. This provides guarantee for the approximation that leads to (4).

Claim: Under the assumption that f_ϕ Lipschitz, f_ϕ is bounded because the input \mathbf{x} is bounded. Let K be the Lipschitz constant of f_ϕ , and let $c = f_\phi(0)$

$$|f_\phi(\mathbf{x}) - c| \leq K|\mathbf{x} - 0| = K|\mathbf{x}| \quad (9)$$

We then show that $\mathbf{KL}(p_\theta||q)$ differ $\mathbf{KL}(q||p_\theta)$ by at most a constant. Since the function $f_\phi(\mathbf{x})$ is lower and upper-bounded. There exists a, b , such that $-a \leq f_\phi(\mathbf{x}) \leq b$ for any \mathbf{x} bounded.

$$\begin{aligned} & \mathbf{KL}(q||p_\theta) - \mathbf{KL}(p_\theta||q) \\ &= \int_{\mathbf{x}} \left[q(\mathbf{x}) \log \left(\frac{q(\mathbf{x})}{p_\theta(\mathbf{x})} \right) - p_\theta(\mathbf{x}) \log \left(\frac{p_\theta(\mathbf{x})}{q(\mathbf{x})} \right) \right] d\mathbf{x} \\ &= \int_{\mathbf{x}} [q(\mathbf{x}) + p_\theta(\mathbf{x})] \log \left(\frac{q(\mathbf{x})}{p_\theta(\mathbf{x})} \right) d\mathbf{x} \\ &\stackrel{\textcircled{1}}{=} \int_{\mathbf{x}} p_\theta(\mathbf{x}) \left[1 + \frac{\exp(\alpha f_\phi(\mathbf{x}))}{Z} \right] \log \left(\frac{\exp(\alpha f_\phi(\mathbf{x}))}{Z} \right) d\mathbf{x} \\ &\stackrel{\textcircled{2}}{\leq} \alpha(a+b) \int_{\mathbf{x}} p_\theta(\mathbf{x}) \left[1 + \frac{\exp(\alpha f_\phi(\mathbf{x}))}{Z} \right] d\mathbf{x} \\ &\stackrel{\textcircled{3}}{=} 2\alpha(a+b), \end{aligned} \quad (10)$$

where $\textcircled{1}$ plugs $q^*(\mathbf{x}) = \frac{p_\theta(\mathbf{x}) \exp(\alpha f_\phi(\mathbf{x}))}{Z}$; $\textcircled{2}$ uses the fact $\log \left(\frac{\exp(\alpha f_\phi(\mathbf{x}))}{Z} \right) = \log \left(\frac{\exp(\alpha f_\phi(\mathbf{x}))}{\int_{\mathbf{x}} p_\theta(\mathbf{x}) \exp(\alpha f_\phi(\mathbf{x})) d\mathbf{x}} \right) \leq \log \left(\frac{\exp(\alpha b)}{\int_{\mathbf{x}} p_\theta(\mathbf{x}) \exp(-\alpha a) d\mathbf{x}} \right) = \alpha(a+b)$; $\textcircled{3}$ uses $\int_{\mathbf{x}} p_\theta(\mathbf{x}) \frac{\exp(\alpha f_\phi(\mathbf{x}))}{Z} d\mathbf{x} = 1$. The above claim completes the theoretical guarantee on the reverse-KL approximation in (4).

6.2 Proof on the necessity of Lipschitz constraint on the discriminator

Although [25] shows preliminary connections between PR and GAN, the proposed PR framework does not provide informative gradient to the generator when treated as a GAN loss. Following [58], we consider the training problem when the discriminator (i.e. $f_\phi(\mathbf{x})$ here) is optimal: when discriminator $f_\phi^*(\mathbf{x})$ is optimal, then the gradient of generator $g(f_\phi(\mathbf{x}))$ is $\nabla_{f_\phi^*(\mathbf{x})} g(f_\phi^*(\mathbf{x})) \cdot \nabla_{\mathbf{x}} f_\phi^*(\mathbf{x})$ which could be very small due to vanished $\nabla_{\mathbf{x}} f_\phi^*(\mathbf{x})$. In this way, it is hard to push the generated data distribution p_θ towards the targeted real distribution p_r . This problem also exists in ?? because

$$f_\phi^*(\mathbf{x}) = \arg \min_{f_\phi(\mathbf{x})} \alpha [p_r(\mathbf{x}) f_\phi(\mathbf{x}) - q(\mathbf{x}) f_\phi(\mathbf{x})]. \quad (11)$$

So if p_r and q are disjoint, we have

$$\begin{aligned} f_\phi^*(\mathbf{x}) &= \arg \min_{f_\phi(\mathbf{x})} \alpha [p_r(\mathbf{x}) f_\phi(\mathbf{x}) - q(\mathbf{x}) f_\phi(\mathbf{x})] \\ &= \begin{cases} \arg \min_{f_\phi(\mathbf{x})} p_r(\mathbf{x}) f_\phi(\mathbf{x}), & \text{if } \mathbf{x} \sim p_r \\ \arg \min_{f_\phi(\mathbf{x})} -q(\mathbf{x}) f_\phi(\mathbf{x}), & \text{if } \mathbf{x} \sim q. \end{cases} \end{aligned} \quad (12)$$

Note that for any $\mathbf{x} \sim p_r$, $f_\phi^*(\mathbf{x})$ is not related to q and thus its gradient $\nabla_{f_\phi^*(\mathbf{x})}$ also does not relate to q . Similarly, for any $\mathbf{x} \sim q$, $\nabla_{f_\phi^*(\mathbf{x})}$ does not provide any information of p_r . Therefore, the proposed loss in [25] cannot guarantee informative gradient [58] that pushes q or p_θ towards to p_r .

6.3 Experiments: More Details and Results

6.3.1 Binary classifier for probability ratio clipping

For the image generation and text generation, the binary classifier C in Eq.(6) has the same architecture as the discriminator except an additional Sigmoid activation at the output layer. The binary

classifier is trained with real and fake mini-batches alongside the generator, and requires no additional loops. We select the clipping parameter ϵ from $\{0.2, 0.4\}$, as they are typically used in PPO.

In addition in the task of image generation, we observe similar overall performance between training on raw inputs from the generator/dataset and training on input features from the first residual block of the discriminator (D), thus further reducing the computational overhead of the binary classifier.

6.3.2 Image Generation on CIFAR-10

We translate the code⁶ provided by Wei et al. [51] into Pytorch to conduct our experiments. We use the same architecture: a residual architecture for both generator and discriminator, and enforcing Lipschitz constraint on the discriminator in the same way as CT-GAN [51]. During training, we interleave 5 generator iterations with 5 discriminator iterations. We optimize the generator and discriminators with Adam (Generator lr: $5e-5$, Discriminator lr: $1e-4$, betas: (0.0, 0.9)). We set the clipping threshold $\epsilon := 0.4$ for the surrogate loss and we linearly anneal the learning rate with respect to the number of training epochs.

Discriminator sample re-weighting stabilizes DCGAN We quantitatively evaluate the effect of discriminator re-weighted sampling by comparing DCGAN [40] against DCGAN with discriminator re-weighting. Starting from the DCGAN architecture and hyper-parameters, we run 200 random configurations of learning rate, batch size, non-linearity (ReLU/LeakyReLU), and base filter count (32, 64). Results are summarized in Table 5. DCGANs trained with re-weighted sampling has significantly less collapse rate, and achieves better overall performance in terms of Inception Score. These results well demonstrate the effectiveness of the proposed discriminator re-weighted sampling mechanism.

Method	Collapse rate	Avg IS	Best IS
DCGAN	52.4%	4.2	6.1
DCGAN + Re-weighting	30.2%	5.1	6.7

Table 5: Outcomes of 200 trials with random configurations. The performance of the models are measured through Inception score. We identify training collapse when the average discriminator loss over 2000 batches is below $1e^{-20}$ or above $1 - 1e^{-20}$. DCGAN re-weighted with our loss has lower collapse rate and higher average performance.

Discriminator re-weighted samples To provide an illustration of how discriminator weights can help the discriminator concentrate on the fake samples of better quality during the training phase, in Figure 5 we plot the fake samples of a trained ResNet model alongside their corresponding discriminator weights.

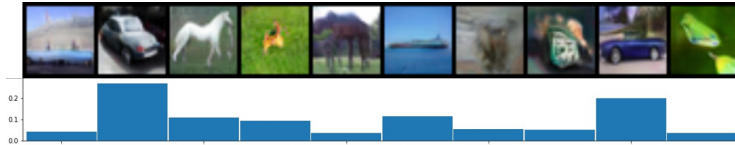


Figure 5: One batch of generated images together with their corresponding softmax discriminator weights. The more photo-realistic images (columns 2, 3, 5, 8) receive higher discriminator weights. In this batch, the generator will be influenced more by gradients from the better-quality samples above.

Clipped surrogate objective One unique benefit of the clipped surrogate objective is that it allows our model to obtain an estimate of the effectiveness of the discriminator, which then enables us to follow a curriculum that takes more than one (n_g) generator steps per (n_c) critic steps. In practice, setting $n_g = n_c = 5$ achieves good quality, which also allows us to take 5 times more generator steps than prior works [2, 18, 34, 51] with the same number of discriminator iterations. Table 1 shows the improvement enabled by applying the surrogate objective.

Generated samples Figure 6 shows more image samples by our model.

⁶github.com/biuyq/CT-GAN

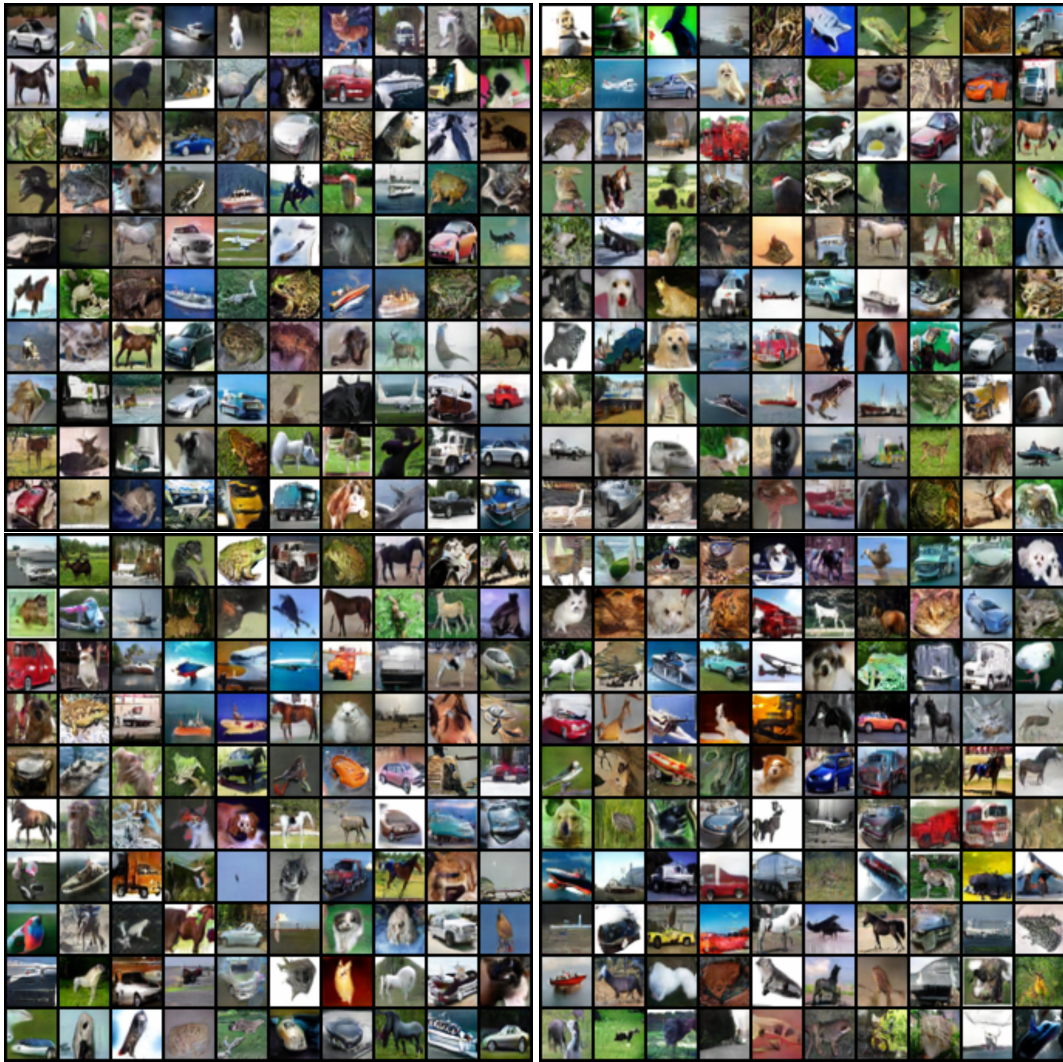


Figure 6: More samples from our generator on CIFAR-10

6.3.3 Text Generation

We build upon the Pytorch implementation⁷ of RelGAN. We use the exact same model architecture as provided in the code, and enforce Lipschitz constraint on the discriminator in the same way as in WGAN-GP [2].

During training, we interleave 5 generator iterations with 5 discriminator iterations. We use Adam optimizer (generator lr: 1e-4, discriminator lr: 3e-4). We set the clipping threshold $\epsilon = 0.2$ for the surrogate loss and we linearly anneal the learning rate with respect to the number of training epochs.

6.3.4 Text Style Transfer

We build upon the Texar-TensorFlow [22] style-transfer model by Tikhonov et al. [50]⁸. We use the exact same model architecture and hyper-parameters as provided in the code, and enforce Lipschitz constraint on the discriminator in the same way as WGAN-GP [2]. In addition, we replace the discriminator D in Figure 7, by our loss with an auxiliary linear style classifier as in Odena et al. [39].

⁷github.com/williamSYSU/TextGAN-PyTorch

⁸https://github.com/VAShibaev/text_style_transfer

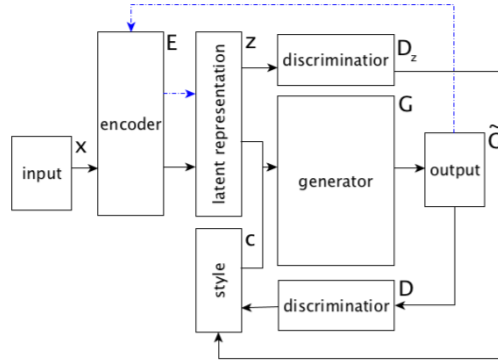


Figure 7: Model architecture from [50], where the style discriminator (D) is a structured constraint the generator optimize against. A latent code discriminator ensure the independence between semantic part of the latent representation and the style of the text. **Blue** dashed arrows denote additional independence constraints of latent representation and controlled attribute, see [50] for the details.

We did not apply the surrogate loss to approximate the KL divergence, but relied on gradient clipping on the generator.