

Improved Consistency Regularization for GANs

Zhengli Zhao,^{1,2} Sameer Singh,¹ Honglak Lee,² Zizhao Zhang,² Augustus Odena,² Han Zhang²

¹ University of California, Irvine

² Google Research

zhengliz@uci.edu, zhanghan@google.com

Abstract

Recent work has increased the performance of Generative Adversarial Networks (GANs) by enforcing a consistency cost on the discriminator. We improve on this technique in several ways. We first show that consistency regularization can introduce artifacts into the GAN samples and explain how to fix this issue. We then propose several modifications to the consistency regularization procedure designed to improve its performance. We carry out extensive experiments quantifying the benefit of our improvements. For unconditional image synthesis on CIFAR-10 and CelebA, our modifications yield the best known FID scores on various GAN architectures. For conditional image synthesis on CIFAR-10, we improve the state-of-the-art FID score from 11.48 to 9.21. Finally, on ImageNet-2012, we apply our technique to the original BigGAN model and improve the FID from 6.66 to 5.38, which is the best score at that model size.

1 Introduction

Generative Adversarial Networks (GANs; Goodfellow et al. 2014) are a powerful class of deep generative models, but are known for training difficulties (Salimans et al. 2016). Many approaches have been introduced to improve GAN performance (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017; Miyato et al. 2018a; Sinha et al. 2020). Recent work (Wei et al. 2018; Zhang et al. 2020) suggests that the performance of generative models can be improved by introducing consistency regularization techniques – which are popular in the semi-supervised learning literature (Oliver et al. 2018). In particular, Zhang et al. (2020) show that Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) augmented with consistency regularization can achieve state-of-the-art image-synthesis results. In CR-GAN, real images and their corresponding augmented counterparts are fed into the discriminator. The discriminator is then encouraged — via an auxiliary loss term — to produce similar outputs for an image and its corresponding augmentation.

Though the consistency regularization in CR-GAN is effective, the augmentations are only applied to the real images and not to generated samples, making the whole procedure somewhat imbalanced. In particular, the generator can learn these artificial augmentation features and introduce them into

generated samples as undesirable artifacts.¹ Further, by regularizing only the discriminator, and by only using augmentations in image space, the regularizations in Wei et al. (2018) and Zhang et al. (2020) do not act directly on the generator. By constraining the mapping from the prior to the generated samples, we can achieve further performance gains on top of those yielded by performing consistency regularization on the discriminator in the first place.

In this work, we introduce *Improved Consistency Regularization* (ICR) which applies forms of consistency regularization to the generated images, the latent vector space, and the generator. First, we address the lack of regularization on the generated samples by introducing *balanced consistency regularization* (bCR), where a consistency term on the discriminator is applied to both real images and samples coming from the generator. Second, we introduce *latent consistency regularization* (zCR), which incorporates regularization terms modulating the sensitivity of both the generator and discriminator changes in the prior. In particular, given augmented/perturbed latent vectors, we show that it is helpful to encourage the generator to be *sensitive* to the perturbations and the discriminator to be *insensitive*. We combine bCR and zCR, and call it Improved Consistency Regularization (ICR).

ICR yields state-of-the-art image synthesis results. For unconditional image synthesis on CIFAR-10 and CelebA, our method yields the best known FID scores on various GAN architectures. For conditional image synthesis on CIFAR-10, we improve the state-of-the-art FID score from 11.48 to 9.21. Finally, on ImageNet-2012, we apply our technique to the original BigGAN (Brock, Donahue, and Simonyan 2019) model and improve the FID from 6.66 to 5.38, which is the best score at that model size.

2 Improved Consistency Regularization

For semi-supervised or unsupervised learning, consistency regularization techniques are effective and have become broadly used recently (Sajjadi, Javanmardi, and Tasdizen 2016; Laine and Aila 2016; Zhai et al. 2019; Xie et al. 2019; Berthelot et al. 2019). The intuition behind these techniques is to encode into model training some prior knowledge: that the model should produce consistent predictions given input instances and their semantics-preserving augmentations.

¹We show examples in Fig. 5 and discuss further in Section 4.1.




Figure 1: Illustrations comparing our methods to the baseline. (1) CR-GAN (Zhang et al. 2020) is the baseline, with consistency regularization applied only between real images and their augmentations. (2) In Balanced Consistency Regularization (bCR-GAN), we also introduce consistency regularization between generated fake images and their augmentations. With consistency regularization on both real and fake images, the discriminator is trained in a balanced way and less augmentation artifacts are generated. (3) Furthermore, we propose Latent Consistency Regularization (zCR-GAN), where latent z is augmented with noise of small magnitude. Then for the discriminator, we regularize the consistency between corresponding pairs; while for the generator we encourage the corresponding generated images to be more diverse. Note that $\{\rightarrow\leftarrow\}$ indicates a loss term encouraging pairs to be closer together, while $\{\leftarrow\rightarrow\}$ indicates a loss term pushing pairs apart.

The augmentations (or transformations) can take many forms, such as image flipping and rotating, sentence back-translating, or even adversarial attacks. Penalizing the inconsistency can be easily achieved by minimizing L_2 loss (Sajjadi, Javanmardi, and Tasdizen 2016; Laine and Aila 2016) between instance pairs, or KL-divergence loss (Xie et al. 2019; Miyato et al. 2018b) between distributions. In the GAN literature, Wei et al. (2018) propose a consistency term derived from Lipschitz continuity considerations to improve the training of WGAN. Recently, CR-GAN (Zhang et al. 2020) applies consistency regularization to the discriminator and achieves substantial improvements.

Below we start by introducing our two new techniques, abbreviated as bCR and zCR, to improve and generalize CR for GANs. We denote the combination of both of these techniques as ICR, and we will later show that ICR yields state-of-the-art image synthesis results in a variety of settings. Figure 1 shows illustrations comparing our methods to the baseline CR-GAN Zhang et al. (2020).

2.1 Balanced Consistency Regularization (bCR)

Figure 1(1) illustrates the baseline CR-GAN, in which a term is added to the discriminator loss function that penalizes its sensitivity to the difference between the original image x and the augmented image $T(x)$. One key problem with the original CR-GAN is that the discriminator might ‘mistakenly believe’ that the augmentations are actual features of the target data set, since these augmentations are only performed on the real images. This phenomenon, which we refer to as consistency imbalance, is not easy to notice for certain types

Algorithm 1 Balanced Consistency Regularization (bCR)

Input: parameters of generator θ_G and discriminator θ_D , consistency regularization coefficient for real images λ_{real} and fake images λ_{fake} , augmentation transform T (for images, e.g. shift, flip, cutout, etc).

for number of training iterations **do**

 Sample batch $z \sim p(z)$, $x \sim p_{\text{real}}(x)$

 Augment both real $T(x)$ and fake $T(G(z))$ images

$L_D \leftarrow D(G(z)) - D(x)$

$L_{\text{real}} \leftarrow \|D(x) - D(T(x))\|^2$

$L_{\text{fake}} \leftarrow \|D(G(z)) - D(T(G(z)))\|^2$

$\theta_D \leftarrow \text{AdamOptimizer}(L_D + \lambda_{\text{real}}L_{\text{real}} + \lambda_{\text{fake}}L_{\text{fake}})$

$L_G \leftarrow -D(G(z))$

$\theta_G \leftarrow \text{AdamOptimizer}(L_G)$

end for

of augmentation (e.g. image shifting and flipping). However, it can result in generated samples with explicit augmentation artifacts when augmented samples contain visual artifacts not belonging to real images. For example, we can easily observe this effect for CR-GAN with cutout augmentation: see the second column in Figure 5. This undesirable effect greatly limits the choice of advanced augmentations we could use.

In order to correct this issue, we propose to also augment generated samples before they are fed into the discriminator, so that the discriminator will be evenly regularized with respect to both real and fake augmentations and thereby be encouraged to focus on meaningful visual information.

Specifically, a gradient update step will involve four batches, a batch of real images x , augmentations of these

real images $T(x)$, a batch of generated samples $G(z)$, and that same batch with augmentations $T(G(z))$. The discriminator will have terms that penalize its sensitivity between corresponding $\{x, T(x)\}$ and also $\{G(z), T(G(z))\}$, while the generator cost remains unmodified.

This technique is described in more detail in Algorithm 1 and visualized in Figure 1(2). We abuse the notation a little in the sense that $D(x)$ denotes the output vector before activation of the last layer of the discriminator given input x . $T(x)$ denotes an augmentation transform, here for images (e.g. shift, flip, cutout, etc). The consistency regularization can be balanced by adjusting the strength of λ_{real} and λ_{fake} . This proposed bCR technique not only removes augmentation artifacts (see third column of Figure 5), but also brings substantial performance improvement (see Section 3 and 4).

2.2 Latent Consistency Regularization (zCR)

Algorithm 2 Latent Consistency Regularization (zCR)

Input: parameters of generator θ_G and discriminator θ_D , consistency regularization coefficient for generator λ_{gen} and discriminator λ_{dis} , augmentation transform T (for latent vectors, e.g. adding small perturbation noise $\sim \mathcal{N}(0, \sigma_{\text{noise}})$).

for number of training iterations **do**

Sample batch $z \sim p(z)$, $x \sim p_{\text{real}}(x)$
 Sample perturbation noise $\Delta z \sim \mathcal{N}(0, \sigma_{\text{noise}})$
 Augment latent vectors $T(z) \leftarrow z + \Delta z$
 $L_D \leftarrow D(G(z)) - D(x)$
 $L_{\text{dis}} \leftarrow \|D(G(z)) - D(G(T(z)))\|^2$
 $\theta_D \leftarrow \text{AdamOptimizer}(L_D + \lambda_{\text{dis}}L_{\text{dis}})$
 $L_G \leftarrow -D(G(z))$
 $L_{\text{gen}} \leftarrow -\|G(z) - G(T(z))\|^2$
 $\theta_G \leftarrow \text{AdamOptimizer}(L_G + \lambda_{\text{gen}}L_{\text{gen}})$

end for

In Section 2.1, we focus on consistency regularization with respect to augmentations in image space on the inputs to the discriminator. In this section, we consider a different question: Would it help if we enforce consistency regularization on augmentations in latent space (Zhao, Dua, and Singh 2018)? Given that a GAN model consists of both a generator and a discriminator, it seems reasonable to ask if techniques that can be applied to the discriminator can also be effectively applied to the generator in certain analogous way.

Towards this end, we propose to augment inputs to the generator by slightly perturbing draws z from the prior to yield $T(z) = z + \Delta z$, $\Delta z \sim \mathcal{N}(0, \sigma_{\text{noise}})$. Assuming the perturbations Δz are small enough, we expect that output of the discriminator ought not to change much with respect to this perturbation and modify the discriminator loss by enforcing $\|D(G(z)) - D(G(T(z)))\|^2$ is small.

However, with only this term added onto the GAN loss, the generator would be prone to collapse to generating specific samples for any latent z , since that would easily satisfy the constraint above. To avoid this, we also modify the loss function for the generator with a term that maximizes the difference between $G(z)$ and $G(T(z))$, which also encourages generations from similar latent vectors to be diverse. Though motivated differently, this can be seen as related to

the Jacobian Clamping technique from Odena et al. (2018) and diversity increase technique in Yang et al. (2019).

This method is described in more detail in Algorithm 2 and visualized in Figure 1(3). $G(z)$ denotes the output images of the generator given input z . $T(x)$ denotes an augmentation transform, here for latent vectors (e.g. adding small perturbation noise). The strength of consistency regularization for the discriminator can be adjusted via λ_{dis} . From the view of the generator, intuitively, the term $L_{\text{gen}} = -\|G(z) - G(T(z))\|^2$ encourages $\{G(z), G(T(z))\}$ to be diverse. We have conducted analysis on the effect of λ_{gen} with experiments in Section 4.3. This technique substantially improves the performance of GANs, as measured by FID. We present experimental results in Section 3 and 4.

2.3 Putting it All Together (ICR)

Though both Balanced Consistency Regularization and Latent Consistency Regularization improve GAN performance (see Section 3), it is not obvious that they would work when ‘stacked on top’ of each other. That is, maybe they are accomplishing the same thing in different ways, and we cannot add up their benefits. However, validated with extensive experiments, we achieve the best experimental results when combining Algorithm 1 and Algorithm 2 together. We call this combination Improved Consistency Regularization (ICR). Note that in ICR, we augment inputs in both image and latent spaces, and add regularization terms to both the discriminator and the generator. We regularize the discriminator’s consistency between corresponding pairs of $\{D(x), D(T(x))\}$, $\{D(G(z)), D(T(G(z)))\}$, and $\{D(G(z)), D(G(T(z)))\}$; For the generator, we encourage diversity between $\{G(z), G(T(z))\}$.

3 Experiments

In this section, we validate our methods on different data sets, model architectures, and GAN loss functions. We compare both Balanced Consistency Regularization (Algorithm 1) and Latent Consistency Regularization (Algorithm 2) with several baseline methods. We also combine both techniques (we abbreviate this combination as ICR) and show that this yields state-of-the-art FID numbers. We follow the best experimental practices established in Kurach et al. (2019), aggregating all runs and reporting the FID distribution of the top 15% of trained models. We provide both quantitative and qualitative results (with more in the appendix).

3.1 Baseline Methods

We compare our methods with four GAN regularization techniques: Gradient Penalty (GP) (Gulrajani et al. 2017), DRAGAN (DR) (Kodali et al. 2017), Jensen-Shannon Regularizer (JSR) (Roth et al. 2017), and vanilla Consistency Regularization (CR) (Zhang et al. 2020). The regularization strength λ is set to 0.1 for JSR, and 10 for all others.

Following the procedures from Lucic et al. (2018); Kurach et al. (2019), we evaluate these methods across different data sets, neural architectures, and loss functions. For optimization, we use the Adam optimizer with batch size of 64 for all

experiments. By default, spectral normalization (SN) (Miyato et al. 2018a) is used in the discriminator, as it is the most effective normalization method for GANs (Kurach et al. 2019) and is becoming the standard for recent GANs (Brock, Donahue, and Simonyan 2019; Wu et al. 2019).

3.2 Data Sets and Evaluation

We carry out extensive experiments comparing our methods against the above baselines on three commonly used data sets in the GAN literature: CIFAR-10 (Krizhevsky, Hinton et al. 2009), CelebA-HQ-128 (Karras et al. 2018), and ImageNet-2012 (Russakovsky et al. 2015).

For data set preparation, we follow the detailed procedures in Kurach et al. (2019). CIFAR-10 contains 60K 32×32 images with 10 labels, out of which 50K are used for training and 10K are used for testing. CelebA-HQ-128 (CelebA) consists of 30K 128×128 facial images, out of which we use 3K images for testing and train models with the rest. ImageNet-2012 has approximately 1.2M images with 1000 labels, and we down-sample the images to 128×128 . We stop training after 200k generator update steps for CIFAR-10, 100k steps for CelebA, and 250k for ImageNet.

We use the Fréchet Inception Distance (FID) (Heusel et al. 2017) as the primary metric for quantitative evaluation. FID has been shown to correlate well with human evaluation of image quality and to be helpful in detecting intra-class mode collapse. We calculate FID between generated samples and real test images, using 10K images on CIFAR-10, 3K on CelebA, and 50K on ImageNet. We also report Inception Scores (Salimans et al. 2016) in the appendix.

By default, the augmentation transform T on latent vectors z is adding Gaussian noise $\Delta z \sim \mathcal{N}(0, \sigma_{\text{noise}})$. The augmentation transform T on images is a combination of randomly flipping horizontally and shifting by multiple pixels (up to 4 for CIFAR-10 and CelebA, and up to 16 for ImageNet). This transform combination results in better performance than alternatives (see Zhang et al. (2020)). Though we outperform CRGAN for different augmentation strategies, we use the same image augmentation strategies as the best one (random flip and shift) in CRGAN for comparison.

There are many different GAN loss functions and we elaborate on several of them in the Appendix. Following Zhang et al. (2020), for each data set and model architecture combination, we conduct experiments using the loss function that achieves the best performance on baselines.

3.3 Unconditional GAN Models

We first test out techniques on unconditional image generation, which is to model images from an object-recognition data set without any reference to the underlying classes. We conduct experiments on the CIFAR-10 and CelebA data sets, and use both DCGAN (Radford, Metz, and Chintala 2015) and ResNet (He et al. 2016) GAN architectures.

DCGAN on CIFAR-10 Figure 2 presents the results of DCGAN on CIFAR-10 with the hinge loss. Vanilla Consistency Regularization (CR) (Zhang et al. 2020) outperforms all other baselines. Our Balanced Consistency Regularization (bCR) technique improves on CR by more than 3.0 FID




Figure 2: FID scores for DCGAN trained on CIFAR-10 with the hinge loss, for a variety of regularization techniques. Consistency regularization significantly outperforms non-consistency regularizations. Adding Balanced Consistency Regularization causes a larger improvement than Latent Consistency Regularization, but both yield improvements much larger than measurement variances.

points. Our Latent Consistency Regularization (zCR) technique improves scores less than bCR, but the improvement is still significant compared to the measurement variance. We set $\lambda_{\text{real}} = \lambda_{\text{fake}} = 10$ for bCR, while using $\sigma_{\text{noise}} = 0.03$, $\lambda_{\text{gen}} = 0.5$, and $\lambda_{\text{dis}} = 5$ for zCR.

ResNet on CIFAR-10 DCGAN-type models are well-known and it is encouraging that our techniques increase performance for those models, but they have been substantially surpassed in performance by newer techniques. We then validate our methods on more recent architectures that use residual connections (He et al. 2016). Figure 3 shows unconditional image synthesis results on CIFAR-10 using a GAN model with residual connections and the non-saturating loss. Though both of our proposed modifications still outperform all baselines, Latent Consistency Regularization works better than Balanced Consistency Regularization, contrary to the results in Figure 2. For hyper-parameters, we set $\lambda_{\text{real}} = 10$ and $\lambda_{\text{fake}} = 5$ for bCR, while using $\sigma_{\text{noise}} = 0.07$, $\lambda_{\text{gen}} = 0.5$, and $\lambda_{\text{dis}} = 20$ for zCR.

DCGAN on CelebA We also conduct experiments on the CelebA data set. The baseline model we use in this case is a DCGAN model with the non-saturating loss. We set $\lambda_{\text{real}} = \lambda_{\text{fake}} = 10$ for bCR, while using $\sigma_{\text{noise}} = 0.1$, $\lambda_{\text{gen}} = 1$, and $\lambda_{\text{dis}} = 10$ for zCR. The results are shown in Figure 4 and are overall similar to those in Figure 2. The improvements in performance for CelebA are not as large as those for CIFAR-10, but they are still substantial, suggesting that our methods generalize across data sets.

Improved Consistency Regularization As alluded to above, we observe experimentally that combining Balanced Consistency regularization (bCR) and Latent Consistency Regularization (zCR) (into Improved Consistency Regularization (ICR)) yields results that are better than those given by either method alone. Using the above experimental results,




Figure 3: FID scores for a ResNet-style GAN trained on CIFAR-10 with the non-saturating loss, for a variety of regularization techniques. Contrary to the results in Figure 2, Latent Consistency Regularization outperforms Balanced Consistency Regularization, though they both substantially surpass all baselines.

Methods	CIFAR-10 (DCGAN)	CIFAR-10 (ResNet)	CelebA (DCGAN)
W/O	24.73	19.00	25.95
GP	25.83	19.74	22.57
DR	25.08	18.94	21.91
JSR	25.17	19.59	22.17
CR	18.72	14.56	16.97
ICR (ours)	15.87	13.36	15.43

Table 1: FID scores for Unconditional Image Synthesis. ICR achieves the best performance overall. Baselines are: not using regularization (W/O), Gradient Penalty (GP) (Gulrajani et al. 2017), DRAGAN (DR) (Kodali et al. 2017), Jensen-Shannon Regularizer (JSR) (Roth et al. 2017), and vanilla Consistency Regularization (CR) (Zhang et al. 2020).

we choose the best-performing hyper-parameters to carry out experiments for ICR, regularizing with both bCR and zCR. Table 1 shows that ICR yields the best results for all three unconditional synthesis settings we study. Moreover, the results of the ResNet model on CIFAR-10 are, to the best of our knowledge, the best reported results for unconditional CIFAR-10 synthesis.

3.4 Conditional GAN Models

In this section, we apply our consistency regularization techniques to the publicly available implementation of BigGAN (Brock, Donahue, and Simonyan 2019) from Kurach et al. (2019). We compare it to baselines from Brock, Donahue, and Simonyan (2019); Miyato et al. (2018a); Zhang et al. (2020). Note that the FID numbers from Wu et al. (2019) are based on a larger version of BigGAN called BigGAN-Deep with substantially more parameters than the original BigGAN, and are thus not comparable to the numbers we report here.

On CIFAR-10, our techniques yield the best known FID




Figure 4: FID scores for DCGAN trained on CelebA with the non-saturating loss, for a variety of regularization techniques. Consistency regularization significantly outperforms all other baselines. Balanced Consistency Regularization further improves on Consistency Regularization by more than 2.0 in terms of FID, while Latent Consistency Regularization improves by around 1.0.

Models	CIFAR-10	ImageNet
SNGAN	17.50	27.62
BigGAN	14.73	8.73
CR-BigGAN	11.48	6.66
bCR-BigGAN	10.54	6.24
zCR-BigGAN	10.19	5.87
ICR-BigGAN	9.21	5.38

Table 2: FID scores for class conditional image generation on CIFAR-10 and ImageNet. We compare our ICR technique with state-of-the-art GAN models including SNGAN (Miyato et al. 2018a), BigGAN (Brock, Donahue, and Simonyan 2019), and CR-GAN (Zhang et al. 2020). The BigGAN implementation we use is from Kurach et al. (2019). (*)-BigGAN has the exactly same architecture as the publicly available BigGAN and is trained with the same settings, but with our consistency regularization techniques added to GAN losses. On CIFAR-10 and ImageNet, we improve the FID numbers to 9.21 and 5.38 correspondingly, which are the best known scores at that model size.

score for conditional synthesis with CIFAR-10²: 9.21. On conditional Image Synthesis on the ImageNet data set, our technique yields FID of 5.38. This is the best known score using the same number of parameters as in the original BigGAN model, though the much larger model from Wu et al. (2019) achieves a better score. For both setups, we set

²There are a few papers that report lower scores using the PyTorch implementation of the FID. That implementation outputs numbers that are much lower, which are not comparable to numbers from the official TF implementation, as explained at <https://github.com/ajbrock/BigGAN-PyTorch/#an-important-note-on-inception-metrics>




Figure 5: Illustration of resolving generation artifacts by Balanced Consistency Regularization. The first column shows CIFAR-10 training images augmented with cutout of different sizes. The second column demonstrates that the vanilla CR-GAN (Zhang et al. 2020) can cause augmentation artifacts to appear in generated samples. This is because CR-GAN only has consistency regularization on real images passed into the discriminator. In the last column (our Balanced Consistency Regularization: bCR in Algorithm 1) this issue is fixed with both real and generated fake images augmented before being fed into the discriminator.

$\lambda_{\text{real}} = \lambda_{\text{fake}} = 10$, together with $\sigma_{\text{noise}} = 0.05$, $\lambda_{\text{gen}} = 0.5$, and $\lambda_{\text{dis}} = 20$.

4 Ablation Studies

To better understand how the various hyper-parameters introduced by our new techniques affect performance, we conduct a series of ablation studies. We include both quantitative and qualitative results.

4.1 Examining Artifacts Resulting from ‘Vanilla’ Consistency Regularization

To understand the augmentation artifacts resulting from using vanilla CR-GAN (Zhang et al. 2020), and to validate that Balanced Consistency Regularization removes those artifacts, we carry out a series of qualitative experiments using varying sizes for the cutout (DeVries and Taylor 2017) augmentation. We experiment with cutouts of size 8×8 , 16×16 , and 32×32 , training both vanilla CR-GANs and GANs with

λ_{fake}	8x8 cutout	16x16 cutout	32x32 cutout
0	0.07 ± 0.03	0.12 ± 0.02	0.66 ± 0.03
2	0.03 ± 0.02	0.08 ± 0.01	0.09 ± 0.02
5	0.01 ± 0.01	0.05 ± 0.01	0.03 ± 0.01
10	0.01 ± 0.01	0.01 ± 0.01	0.01 ± 0.01

Table 3: Fraction of Artifacts: bCR alleviates generation artifacts the more it is enforced (higher λ_{fake}).

Balanced Consistency Regularization. The results are shown in Figure 5. The first column shows CIFAR-10 training images augmented with cutout of different sizes. The second column demonstrates that the vanilla CR-GAN (Zhang et al. 2020) can cause augmentation artifacts to appear in generated samples. This is because CR-GAN only has consistency regularization on real images passed into the discriminator. In the last column (our Balanced Consistency Regularization: bCR in Algorithm 1) this issue is fixed with both real and generated fake images augmented before being fed into the discriminator. Broadly speaking, we observe more substantial cutout artifacts (black rectangles) in samples from CR-GANs with larger cutout augmentations, and essentially no such artifacts for GANs trained with Balanced Consistency Regularization with $\lambda_{\text{fake}} \geq \lambda_{\text{real}}$.

To quantify how much bCR alleviates generation artifacts, we vary cutout sizes and the strength of CR for generated images. We examine 600 of generated images from 3 random runs, and report the fraction of images that contain artifacts of cutouts in Table 3. The strength of CR for real images is fixed at $\lambda_{\text{real}} = 10$. We do observe a few artifacts when $0 < \lambda_{\text{fake}} \ll \lambda_{\text{real}}$, but much less than those from the vanilla CR-GAN. We believe that this phenomenon of introducing augmentation artifacts into generations likely holds for other types of augmentation, but it is much more difficult to confirm for less visible transforms, and sometimes it may not actually be harmful (e.g. flipping of images in most contexts).

4.2 Effect of Hyper-Parameters on Balanced Consistency Regularization’s Performance

In Balanced Consistency Regularization (Algorithm 1), the cost associated with sensitivity to augmentations of the real images is weighted by λ_{real} and the cost associated with sensitivity to augmentations of the generated samples is weighted by λ_{fake} . In order to better understand the interplay between these parameters, we train a DCGAN-type model with spectral normalization on the CIFAR-10 data set with the hinge loss, for many different values of $\lambda_{\text{fake}}, \lambda_{\text{real}}$. The heat map in the appendix shows that it never pays to set either of the parameters to zero: this means that Balanced Consistency Regularization always outperforms vanilla consistency regularization (the baseline CR-GAN). Generally speaking, setting λ_{real} and λ_{fake} similar in magnitude works well. This is encouraging, since it means that the performance of bCR is relatively insensitive to hyper-parameters.




Figure 6: Analysis on the hyper-parameters of Latent Consistency Regularization. We conduct experiments using a ResNet-style GAN on CIFAR-10 with non-saturating loss in order to better understand the interplay between σ_{noise} , λ_{gen} and λ_{dis} . The results show that a moderate value for the generator coefficient (e.g. $\lambda_{\text{gen}} = 0.5$) works the best. With the added term $L_{\text{gen}} = -\|G(z) - G(T(z))\|^2$, the generator is encouraged to be sensitive to perturbations in latent space. For this set of experiments, we observe the best performance adding perturbations with standard deviation of $\sigma_{\text{noise}} = 0.07$, and higher (but not extremely high) values for the discriminator coefficient λ_{dis} also improve further.

4.3 Effect of Hyper-Parameters on Latent Consistency Regularization’s Performance

Latent Consistency Regularization (Algorithm 2) has three hyper-parameters: σ_{noise} , λ_{gen} and λ_{dis} , which respectively govern the magnitude of the perturbation made to the draw from the prior, the weight of the sensitivity of the generator to that perturbation, and the weight of the sensitivity of the discriminator to that perturbation. From the view of the generator, intuitively, the extra loss term added $L_{\text{gen}} = -\|G(z) - G(T(z))\|^2$ encourages $G(z)$ and $G(T(z))$ to be far away from each other.

We conduct experiments using a ResNet-style GAN on the CIFAR-10 data set with the non-saturating loss in order to better understand the interplay between these hyper-parameters. The results in Figure 6 show that a moderate value for the generator coefficient (e.g. $\lambda_{\text{gen}} = 0.5$) works the best (as measured by FID). This corresponds to encouraging the generator to be sensitive to perturbations of samples from the prior. For this experimental setup, perturbations with standard deviation of $\sigma_{\text{noise}} = 0.07$ are the best, and higher (but not extremely high) values for the discriminator coefficient λ_{dis} also perform better.

5 Related Work

There is so much related work on GANs (Goodfellow et al. 2014) that it is impossible to do it justice (see Odena (2019); Kurach et al. (2019) for different overviews of the field), but here we sketch out a few different threads. There is a several-year-long thread of work on scaling GANs up to do

conditional image synthesis on the ImageNet-2012 data set beginning with Odena, Olah, and Shlens (2017), extending through Miyato et al. (2018a); Zhang et al. (2019); Brock, Donahue, and Simonyan (2019); Daras et al. (2019) and most recently culminating in Wu et al. (2019) and Zhang et al. (2020), which presently represent the state-of-the-art models at this task (Wu et al. (2019) uses a larger model size than Zhang et al. (2020) and correspondingly report better scores).

Zhou and Krähenbühl (2019) try to make the discriminator robust to adversarial attacks to the generated images. Our zCR is different in two aspects: zCR enforces the robustness of the compound function $D(G(*))$ to make $D(G(z))$ and $D(G(z + \Delta z))$ consistent, while Zhou and Krähenbühl (2019) only encourage the robustness in the generated image space as they regularize between $D(G(z))$ and $D(G(z) + v)$ where v is a fast normalized gradient attack vector; Instead of only regularizing D , zCR also regularizes G to make $G(z)$ and $G(z + \Delta z)$ different to avoid mode collapse.

Most related work on consistency regularization is from the semi-supervised learning literature, and focuses on regularizing model predictions to be invariant to small perturbations (Bachman, Alsharif, and Precup 2014; Sajjadi, Javanmardi, and Tasdizen 2016; Laine and Aila 2016; Miyato et al. 2018b; Xie et al. 2019) for the purpose of learning from limited labeled data. Wei et al. (2018); Zhang et al. (2020) apply related ideas to training GAN models and observe initial gains, which motivates this work.

There are also several concurrent work related to this paper, indicating an emerging direction of GAN training with augmentations. Zhao et al. (2020a) and Karras et al. (2020) research on how to train GANs with limited data; while Zhao et al. (2020b) mainly focus on thoroughly investigating the effectiveness of different types of augmentations.

6 Conclusion

Extending the recent success of consistency regularization in GANs (Wei et al. 2018; Zhang et al. 2020), we present two novel improvements: Balanced Consistency Regularization, in which generator samples are also augmented along with training data, and Latent Consistency Regularization, in which draws from the prior are perturbed, and the sensitivity to those perturbations is discouraged and encouraged for the discriminator and the generator, respectively.

In addition to fixing a new issue we observe with the vanilla Consistency Regularization (augmentation artifacts in samples), our techniques yield the best known FID numbers for both unconditional and conditional image synthesis on the CIFAR-10 data set. They also achieve the best FID numbers (with the fixed number of parameters used in the original BigGAN (Brock, Donahue, and Simonyan 2019) model) for conditional image synthesis on ImageNet.

These techniques are simple to implement, not particularly computationally burdensome, and relatively insensitive to hyper-parameters. We hope they become a standard part of the GAN training toolkit, and that their use allows more interesting usage of GANs to many sorts of applications.

Acknowledgements

We would like to thank Pouya Pezeshkpour and Colin Raffel for helpful discussions. This work is funded in part by the National Science Foundation grant IIS-1756023.

References

- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein gan. *arXiv preprint arXiv:1701.07875* .
- Bachman, P.; Alsharif, O.; and Precup, D. 2014. Learning with pseudo-ensembles. In *Advances in neural information processing systems*, 3365–3373.
- Berthelot, D.; Carlini, N.; Goodfellow, I. J.; Papernot, N.; Oliver, A.; and Raffel, C. 2019. MixMatch: A Holistic Approach to Semi-Supervised Learning. *NeurIPS* .
- Brock, A.; Donahue, J.; and Simonyan, K. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *ICLR*.
- Daras, G.; Odena, A.; Zhang, H.; and Dimakis, A. G. 2019. Your Local GAN: Designing Two Dimensional Local Attention Mechanisms for Generative Models. *arXiv preprint arXiv:1911.12287* .
- DeVries, T.; and Taylor, G. W. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552* .
- 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.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. C. 2017. Improved training of wasserstein gans. In *Advances in neural information processing systems*, 5767–5777.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Advances in Neural Information Processing Systems*.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*.
- Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020. Training generative adversarial networks with limited data. *arXiv preprint arXiv:2006.06676* .
- Kodali, N.; Abernethy, J.; Hays, J.; and Kira, Z. 2017. On convergence and stability of gans. *arXiv preprint arXiv:1705.07215* .
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. *Citeseer* .
- Kurach, K.; Lucic, M.; Zhai, X.; Michalski, M.; and Gelly, S. 2019. A large-scale study on regularization and normalization in GANs. In *ICML*.
- Laine, S.; and Aila, T. 2016. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242* .
- Lucic, M.; Kurach, K.; Michalski, M.; Gelly, S.; and Bousquet, O. 2018. Are gans created equal? a large-scale study. In *Advances in neural information processing systems*, 700–709.
- Miyato, T.; Kataoka, T.; Koyama, M.; and Yoshida, Y. 2018a. Spectral normalization for generative adversarial networks. In *ICLR*.
- Miyato, T.; Maeda, S.-i.; Ishii, S.; and Koyama, M. 2018b. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* .
- Odena, A. 2019. Open questions about generative adversarial networks. *Distill* 4(4): e18.
- Odena, A.; Buckman, J.; Olsson, C.; Brown, T. B.; Olah, C.; Raffel, C.; and Goodfellow, I. 2018. Is generator conditioning causally related to gan performance? *arXiv preprint arXiv:1802.08768* .
- Odena, A.; Olah, C.; and Shlens, J. 2017. Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2642–2651. JMLR. org.
- Oliver, A.; Odena, A.; Raffel, C. A.; Cubuk, E. D.; and Goodfellow, I. 2018. Realistic evaluation of deep semi-supervised learning algorithms. In *NeurIPS*, 3235–3246.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* .
- Roth, K.; Lucchi, A.; Nowozin, S.; and Hofmann, T. 2017. Stabilizing training of generative adversarial networks through regularization. In *Advances in neural information processing systems*, 2018–2028.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115(3): 211–252. doi: 10.1007/s11263-015-0816-y.
- Sajjadi, M.; Javanmardi, M.; and Tasdizen, T. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NeurIPS*.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. In *Advances in neural information processing systems*, 2234–2242.
- Sinha, S.; Zhao, Z.; Goyal, A.; Raffel, C.; and Odena, A. 2020. Top-k Training of GANs: Improving GAN Performance by Throwing Away Bad Samples. In *NeurIPS*.
- Wei, X.; Gong, B.; Liu, Z.; Lu, W.; and Wang, L. 2018. Improving the improved training of wasserstein gans: A consistency term and its dual effect. *arXiv preprint arXiv:1803.01541* .

Wu, Y.; Donahue, J.; Balduzzi, D.; Simonyan, K.; and Lillcrap, T. 2019. LOGAN: Latent Optimisation for Generative Adversarial Networks. *arXiv preprint arXiv:1912.00953* .

Xie, Q.; Dai, Z.; Hovy, E.; Luong, M.-T.; and Le, Q. V. 2019. Unsupervised Data Augmentation for Consistency Training. *arXiv preprint arXiv:1904.12848* .

Yang, D.; Hong, S.; Jang, Y.; Zhao, T.; and Lee, H. 2019. Diversity-sensitive conditional generative adversarial networks. In *ICLR*.

Zhai, X.; Oliver, A.; Kolesnikov, A.; and Beyer, L. 2019. S⁴L: Self-Supervised Semi-Supervised Learning. *arXiv preprint arXiv:1905.03670* .

Zhang, H.; Goodfellow, I.; Metaxas, D.; and Odena, A. 2019. Self-attention generative adversarial networks. In *ICML*.

Zhang, H.; Zhang, Z.; Odena, A.; and Lee, H. 2020. Consistency Regularization for Generative Adversarial Networks. In *ICLR*.

Zhao, S.; Liu, Z.; Lin, J.; Zhu, J.-Y.; and Han, S. 2020a. Differentiable augmentation for data-efficient gan training. *arXiv preprint arXiv:2006.10738* .

Zhao, Z.; Dua, D.; and Singh, S. 2018. Generating Natural Adversarial Examples. In *ICLR*.

Zhao, Z.; Zhang, Z.; Chen, T.; Singh, S.; and Zhang, H. 2020b. Image Augmentations for GAN Training. *arXiv preprint arXiv:2006.02595* .

Zhou, B.; and Krähenbühl, P. 2019. Don't let your Discriminator be fooled. In *ICLR*.