

---

# Self-Supervised GAN to Counter Forgetting

---

**Ting Chen**<sup>\*†</sup>  
University of California, Los Angeles  
tingchen@cs.ucla.edu

**Xiaohua Zhai**<sup>†</sup>  
Google Brain  
xzhai@google.com

**Neil Houlsby**  
Google Brain  
neilhoulby@google.com

## Abstract

GANs involve training two networks in an adversarial game, where each network’s task depends on its adversary. Recently, several works have framed GAN training as an online or continual learning problem [1–6]. We focus on the discriminator, which must perform classification under an (adversarially) shifting data distribution. When trained on sequential tasks, neural networks exhibit *forgetting*. For GANs, discriminator forgetting leads to training instability [1]. To counter forgetting, we encourage the discriminator to maintain useful representations by adding a self-supervision. Conditional GANs have a similar effect using labels. However, our self-supervised GAN does not require labels, and closes the performance gap between conditional and unconditional models. We show that, in doing so, the self-supervised discriminator learns better representations than regular GANs.<sup>1</sup>

## 1 Introduction

Generative Adversarial Networks (GANs) are a class of unsupervised generative models [7]. GANs involve training a *generator* and *discriminator* network in an adversarial game, such that the generator learns to produce samples from a desired data distribution. Training GANs is challenging because one searches for a Nash equilibrium of a non-convex game in a high-dimensional parameter space. GANs are typically trained with alternating stochastic gradient descent. However, this training procedure is unstable and lacks guarantees [8]. Training may exhibit instability, divergence, cyclic behaviour, or mode collapse [9]. As a result, many works propose techniques to stabilize GAN training [10–16].

Recently, several papers frame GAN training as a continual, or online, learning problem [1–6]. GANs may be viewed as online learning because each network, the generator and discriminator, must learn in the context of a non-stationary environment. The discriminator, for example, is a binary classifier where one class (the fake samples) is non-stationary. With this view, Thanh-Tung et al. [1] study catastrophic forgetting in the discriminator and mode collapse, relating these to training instability. Anonymous [2] counter discriminator forgetting by leveraging techniques from continual learning directly (Elastic Weight Sharing [17] and Intelligent Synapses [18]). Other works address forgetting in GANs by retraining on old data, or adding an explicit memory [4, 5]. Finally, Seff et al. [3] extend continual learning for GANs by training on nonstationary true-data distributions also.

We follow this line of work, but propose an alternative approach. Instead of adding explicit replay or memorization strategies, we apply representation learning. We train the discriminator to learn stable representations that are useful for identifying real images. We propose to add self-supervision to the discriminator, in particular, we apply a rotation-based loss [19].

Conditional GANs, which use side information such as class labels, are state-of-the-art in generating high fidelity complex images [20–22]. A major contributor to their success is augmentation of

---

<sup>\*</sup>Work done at Google

<sup>†</sup>Equal contribution

<sup>1</sup>Full version of this paper available at <https://arxiv.org/pdf/1811.11212.pdf>

the discriminator with supervised information, which encourages it to learn stable and diverse representations. However, labels are not always available, and unconditional GANs perform much worse. Our self-supervised GAN requires no additional information, and closes the performance gap for image generation between conditional and unconditional GANs. As intended, we observe that the self-supervised GAN learns better discriminator representations than other GANs for the task of image classification. We hope that this work encourages further investigation into representation and continual learning to improve GANs.

## 2 The Self-supervised GAN

We first discuss discriminator forgetting and motivate representation learning for GANs. We then present our solution using rotation based self-supervision.

### 2.1 Discriminator Forgetting

The original value function for GAN training is:

$$V(G, D) = \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}(\mathbf{x})} [\log P_D(S = \text{real}|\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim P_G(\mathbf{x})} [\log(1 - P_D(S = \text{real}|\mathbf{x}))] \quad (1)$$

where  $P_{\text{data}}$  is the true data distribution, and  $P_G$  is the data distribution induced by feeding noise drawn from a simple distribution  $\mathbf{z} \sim P(\mathbf{z})$  through the generator,  $\mathbf{x} = G(\mathbf{z})$ .  $P_D(S|\mathbf{x})$  is the discriminator’s Bernoulli distribution over the sources (real or fake).

The generator maximizes Equation 1, while the discriminator minimizes it. Training is typically performed via alternating stochastic gradient descent. The discriminator, therefore, classifies between two data distributions,  $P_{\text{data}}(\mathbf{x})$ , and  $P_G^{(t)}(\mathbf{x})$ , where  $t$  indexes training iterations. The latter, however, is non-stationary because the generator’s parameters are updated over time. This is an online learning problem for which explicit temporal dependencies have been proposed to improve training [2, 4, 6, 8].

In online learning, neural networks tend to forget previous tasks [17, 23, 24]. We demonstrate forgetting in a toy setting, Figure 1(a). Here, a classifier performs sequential 1-vs.-all classification task on each of the ten classes in CIFAR10. The classifier is trained for 1k iterations on each task before switching to the next. Figure 1(a) shows substantial forgetting, despite the tasks being similar. The classifier returns almost to random accuracy, 0.5, each time the task switches. The accuracy even drops substantially when returning to the original task at iteration 10k. This demonstrates that the model does not retain generalizable representations during online learning.

Consideration of convergence provides further indication that discriminators may forget. Goodfellow et al. [7] show that the optimal discriminator estimates the likelihood ratio between the generated and real data distributions. Therefore, given a perfect generator, where  $P_G = P_D$ , the optimal discriminator simply outputs 0.5, and has no requirement to retain any meaningful representations.

Discriminator forgetting may cause training difficulties because (1) it does not learn meaningful representations to guide the generator, and (2) the generator can revert to generating old images to fool it [1]. This is particularly problematic if the data is diverse and complex and the generator exhibits mode collapse. Therefore, we add self-supervision to encourage the discriminator to retain useful representations.

### 2.2 The Self-Supervised GAN

Self-supervised learning is a family of methods for building representations from unsupervised data. Self-supervision works by creating artificial supervised tasks from unsupervised data, training on these tasks, and then extracting representations from the resulting networks [25–27]. Here, we apply the successful image rotation self-supervision method [19]. In this method, the self-supervised network predicts the angle of rotation of an image. In Figure 1(b) we motivate this loss using our toy problem. When we add the self-supervised loss, the network learns features that transfer across tasks; performance continually improves, and does not drop to 0.5 when the distribution shifts.




Figure 1: Image classification accuracy when the underlying class distribution shifts every 1k iterations. The vertical dashed line indicates the end of an entire cycle, and return to the original classification task at  $t = 0$ . *Left*: vanilla classifier. *Right*: classifier with additional self-supervised loss. This example demonstrates that a classifier/discriminator may fail to learn generalizable representations in a non-stationary environment, but an auxiliary loss mitigates this problem.

For the self-supervised GAN, the specific losses we use for the generator and discriminator are:

$$L_G = -V(G, D) - \alpha \mathbb{E}_{\mathbf{x} \sim P_G} \left[ \sum_{r \in \mathcal{R}} \log P'_D(R = r | \mathbf{x}^r) \right], \quad (2)$$

$$L_D = V(G, D) - \beta \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} \left[ \sum_{r \in \mathcal{R}} \log P'_D(R = r | \mathbf{x}^r) \right], \quad (3)$$

where  $V(G, D)$  is the original GAN loss in Equation 1.  $r \in \mathcal{R}$  is a rotation selected from a set of possible rotations. We use  $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$  as in Gidaris et al. [19].  $P'_D(R | \mathbf{x}^r)$  is the discriminator’s distribution over rotations, and  $\mathbf{x}^r$  is the image  $\mathbf{x}$  transformed by rotation  $r$ . Architecturally, we use a single discriminator network with two heads to compute  $P_D$  and  $P'_D$ . See supplementary Figure 4 for an overview.

**A note on convergence** With  $\alpha > 0$  convergence, even under optimal conditions, to the true data distribution  $P_G = P_{\text{data}}$  is not guaranteed. This may not be a concern because current GANs are far from attaining the optimal solution. If it is, one could anneal  $\alpha$  to zero during training. Our intuition is that the proposed loss encourages the discriminator to learn and retain meaningful representations that allow it to distinguish rotations as well as true/fake images. The generator is then trained to match distributions in this feature space which encourages the generation of realistic objects.

### 3 Experiments

We show that self-supervision improves GAN training, even matching the performance of conditional GANs. We also demonstrate, by way of representation evaluation, that the features learned by the self-supervised GAN are more meaningful than those of a vanilla GAN.

**Datasets and Settings** Our main result uses the IMAGENET dataset. For the self-supervised GAN (SsGAN), we adopt the architectures of, and compare to, the unconditional GAN in Miyato et al. [12] (SN-GAN) and conditional GAN in Miyato and Koyama [22] (P-cGAN). Both the generator and discriminator consist of ResNet architectures. For the conditional generator in P-cGAN, we apply label-conditional batch norm. SsGAN does not use labels, so by default the generator does not use condition batch-norm, as in SN-GAN. Therefore, we also consider a variant of the SsGAN with self-modulated batch norm (sBN) [13] on the generator to replace label-conditional batch norm. We compare image generation quality using the Frchet Inception Distance (FID) [28]. We train the model for 1M steps using a single P100 GPU. Further details and results can be found in the supplementary material.

We also test three smaller datasets, CIFAR10, LSUN-BEDROOM, and CELEBA-HQ. Labels for Pc-GAN are only available for CIFAR10. For all datasets we set  $\beta = 1$  and  $\alpha = 0.2$ , chosen using a small hyperparameter sweep. For all other hyperparameters, we use the values in Miyato et al. [12] and Miyato and Koyama [22]. We train on these datasets for 100k steps on a single P100 GPU.

**Image generation** Figure 2 shows FID training curves on IMAGENET. The unconditional GAN is unstable, and sometimes diverges. The conditional P-cGAN substantially outperforms the unconditional SN-GAN. The SsGAN is stable, and even performs as well as the P-cGAN.




Figure 2: Mean training curves across three random seeds.

Dataset	Method	FID
CIFAR10	SN-GAN	19.725
	P-cGAN	<b>15.604</b>
	SsGAN	17.114
	SsGAN (sBN)	<b>15.653</b>
IMAGENET	SN-GAN	56.670
	P-cGAN	<b>42.071</b>
	SsGAN	47.555
	SsGAN (sBN)	43.869
BEDROOM	SN-GAN	16.021
	SsGAN	13.664
	SsGAN (sBN)	<b>13.298</b>
CELEBA-HQ	SN-GAN	<b>23.765</b>
	SsGAN	26.113
	SsGAN (sBN)	24.355

Table 1: Best FID attained across three random seeds.




Figure 3: Top 1 accuracy (mean across three seeds) to predict labels from discriminator representations from the final ResNet block. x-axis gives the number of GAN training iterations.

Table 1 shows FID of the best run across three random seeds for each dataset/model combination. Again, for CIFAR the SsGAN (with sBN) improves unconditional performance, and even equals P-cGAN. For BEDROOM, we also see some improvement. Self supervision appears not to help in CELEBA-HQ, we believe that this data is too simple and face orientations do not provide useful signal.

**Representation Quality** We compare each model’s discriminator representations. We also ablate the GAN loss from SsGAN to assess the representation learned from self-supervision alone. We followed the evaluation protocol of [27]. We train a logistic regression to predict the class labels from the feature maps of the discriminator and report top-1 classification accuracy.

Figure 3 shows the classification performance using the final ResNet block as a function of GAN training iterations. SsGAN produces better representations than SN-GAN. The same is true for all the other blocks (see supplementary, Figures 8 and 9). Interestingly, during training the SsGAN overtakes P-cGAN, which sees the class labels. This indicates that P-cGAN is overfitting the training data. Using rotation loss alone substantially decreases representation quality of the SsGAN. Note that our model uses hyperparameters designed for image generation; therefore, we do not seek state-of-the-art representations. Instead, we use representation evaluation to understand and motivate

our strategy. Indeed it seems that self-supervision improves both discriminator representations and image generation.

**Acknowledgments** We would like to thank Mario Lucic, Marvin Ritter, Sylvain Gelly, Ilya Tolstikhin, Alexander Kolesnikov and Lucas Beyer for help with and discussions on this project.

## References

- [1] Hoang Thanh-Tung, Truyen Tran, and Svetha Venkatesh. On catastrophic forgetting and mode collapse in generative adversarial networks. *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018.
- [2] Anonymous. Generative adversarial network training is a continual learning problem. In *Submitted to International Conference on Learning Representations (ICLR)*, 2019. URL <https://openreview.net/forum?id=SJzuHia9tQ>. under review.
- [3] Ari Seff, Alex Beatson, Daniel Suo, and Han Liu. Continual learning in generative adversarial nets. *arXiv preprint arXiv:1705.08395*, 2017.
- [4] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [5] Youngjin Kim, Minjung Kim, and Gunhee Kim. Memorization precedes generation: Learning unsupervised gans with memory networks. *International Conference on Learning Representations (ICLR)*, 2018.
- [6] Paulina Grnarova, Kfir Y Levy, Aurelien Lucchi, Thomas Hofmann, and Andreas Krause. An online learning approach to generative adversarial networks. *arXiv preprint arXiv:1706.03269*, 2017.
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [8] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [9] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International Conference on Machine Learning (ICML)*, pages 3478–3487, 2018.
- [10] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. *International Conference on Computer Vision (ICCV)*, 2016.
- [11] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of Wasserstein GANs. *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [12] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *International Conference on Learning Representations (ICLR)*, 2018.
- [13] Ting Chen, Mario Lucic, Neil Houlsby, and Sylvain Gelly. On self modulation for generative adversarial networks. *arXiv preprint arXiv:1810.01365*, 2018.
- [14] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *International Conference on Learning Representations (ICLR)*, 2016.
- [15] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*, 2018.
- [16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [17] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 2017.
- [18] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. *International Conference on Machine Learning (ICML)*, 2017.

- [19] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations (ICLR)*, 2018.
- [20] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [21] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier GANs. In *International Conference on Machine Learning (ICML)*, 2017.
- [22] Takeru Miyato and Masanori Koyama. cGANs with projection discriminator. *International Conference on Learning Representations (ICLR)*, 2018.
- [23] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*. 1989.
- [24] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 1999.
- [25] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [26] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *International Conference on Computer Vision (ICCV)*, 2015.
- [27] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European Conference on Computer Vision (ECCV)*, 2016.
- [28] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Günter Klambauer, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a Nash equilibrium. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [29] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- [30] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *To Appear in Advances in Neural Information Processing Systems (NIPS)*, 2018.
- [31] Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, and Olivier Bousquet. Are GANs Created Equal? A Large-scale Study. In *To appear in Advances in Neural Information Processing Systems (NIPS)*, 2018.
- [32] Zhiming Zhou, Yuxuan Song, Lantao Yu, and Yong Yu. Understanding the effectiveness of lipschitz constraint in training of gans via gradient analysis. *arXiv preprint arXiv:1807.00751*, 2018.
- [33] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *International Conference on Learning Representations (ICLR)*, 2017.

## A Additional Details of The Proposed Method

Figure 4 depicts the proposed method. Real and fake/generated images are all rotated four directions. For the shared discriminator  $D$ , its goal is to predict true/fake labels of the *non-rotated* images, and detect the rotation of the *real* images. For the generator, its goal is to generate images to fool discriminator, at the same time trying to generate images that are easier for discriminator to detect rotations (adjusted by  $\alpha$ ).




Figure 4: Rotation feature learning discriminator. The dotted line arrow uses only the non-rotated images, for real/fake classification task. The solid line arrow uses all the four rotated images, for the rotation classification task.

## B Further Experiments and Details

**Datasets** We consider four datasets: CIFAR10, CELEBA-HQ, LSUN-BEDROOM, and IMAGENET. The LSUN-BEDROOM dataset [29] contains around 3M images. We partition the images randomly into a test set containing 30588 images and a train set containing the rest. CELEBA-HQ contains 30k images [16]. We use the  $128 \times 128 \times 3$  version obtained by running the code provided by the authors<sup>3</sup>. We use 3000 examples as the test set and the remaining examples as the training set. CIFAR10 contains 70K images ( $32 \times 32 \times 3$ ), partitioned into 60000 training instances and 10000 testing instances. Finally, we evaluate our method on IMAGENET, which contains 1.3M training images and 50K test images. We re-size the images to  $128 \times 128 \times 3$  as done in Miyato and Koyama [22] and Zhang et al. [15].

**Metrics** To quantitatively evaluate generated samples from different methods, we mainly *Frechet Inception Distance* (FID) [28]. In FID, the true data and generated samples are first embedded in some feature space (a specific layer of Inception). Then, a multivariate Gaussian is fit to the data and the distance computed as  $FID(x, g) = \|\mu_x - \mu_g\|_2^2 + \text{Tr}(\Sigma_x + \Sigma_g - 2(\Sigma_x \Sigma_g)^{\frac{1}{2}})$ , where  $\mu$  and  $\Sigma$  denote the empirical mean and covariance and subscripts  $x$  and  $g$  denote the true and generated data respectively. As shown in Sajjadi et al. [30], Lucic et al. [31], FID is shown to be sensitive to both the addition of spurious modes and to mode dropping.

**Additional experimental details** When the rotation loss is added, we take a quarter of images in the batch, rotate for all four directions to synthetic images of known rotation orientations. This subset

<sup>3</sup>Available at [https://github.com/tkarras/progressive\\_growing\\_of\\_gans](https://github.com/tkarras/progressive_growing_of_gans).

rotation trick could reduce and keep the extra computation small. It is worth noting that other than the extra loss term, other parts (architectures, hyper-parameters) are all kept the same for our method and baselines (both unconditional and conditional GANs). For simplicity, we fix  $\beta = 1$  and compare  $\alpha$  in  $\{0.2, 0.5, 1\}$ ;  $\alpha = 0.2$  is used by default.

We use the Adam optimizer. The learning rate is fixed for both datastes as 0.0002. All models are trained for 100k steps<sup>4</sup>. All models are trained with batch size of 64 in a single Nvidia P100 GPU. Reported results by default are the mean of from 3 independent runs with different random seeds.

**Robustness test** We test the effects of training stabilization brought by the SsGAN. We consider two types of hyper-parameter settings: First, controlling the Lipschitz constant of the discriminator, which is a central quantity analyzed in the GAN literature [12, 32]. We consider two state-of-the-art techniques: Gradient Penalty [11], and Spectral Normalization [12]. For the Gradient Penalty regularizer we consider regularization strength  $\lambda \in \{1, 10\}$ . Second, we consider different Adam optimizer hyper-parameters. We test two popular settings  $(\beta_1, \beta_2)$ :  $(0.5, 0.999)$  and  $(0, 0.9)$ . Previous studies find that multiple discriminator steps per generator step help training [7, 8], thus we also consider both 1 and 2 discriminator steps per generator step<sup>5</sup>. In total, this amounts to three different sets of hyper-parameters for  $(\beta_1, \beta_2, \text{disc\_iter})$ :  $(0, 0.9, 1)$ ,  $(0, 0.9, 2)$ ,  $(0.5, 0.999, 1)$ .

### B.1 Quantitative Comparisons of the Generative Samples

**Convergence curve** Figure 5 shows the convergence curves of different methods. We can clearly see the proposed SsGAN significantly improves the performance over the other unconditional GAN, SN-GAN, and even matches the conditional P-cGAN.




Figure 5: FID convergence curves for different methods with 100k training steps. The self supervision close the gap between the conditional and unconditional GANs, sometimes even surpass the class label conditional GAN.

**Robustness improvement** The use of self-supervised loss not only improves the performance of the default hyper-parameters we use, but also improves under multiple types of hyper-parameters. Table 2 shows the performance comparison with Gradient Penalty with various hyper-parameters, and Table 3 shows the performance comparison with Spectral Normalization with various hyper-parameters. Both figures demonstrate SsGAN reduces the sensitivity of GAN training to hyper-parameters, hence the training is more robust.

**Selection of the hyper-parameters** We further compare different choices of  $\alpha$  (while fixing  $\beta = 1$ ). As discussed, an inappropriate value of  $\alpha$  could bias the generator, while a reasonable one would help. As shown in Table 6, we indeed see the effectiveness of  $\alpha$ . In the values compared, the optimal  $\alpha$  is 1 for CIFAR10, and 0.2 for IMAGENET. By default we use 0.2 for all datasets for simplicity.

<sup>4</sup>Results in the main paper for IMAGENET are obtained from training with 1M steps.

<sup>5</sup>We also experimented with 5 steps which didn't outperform the 2 step setting.



$\lambda$	$\beta_1$	$\beta_2$	Disc_iters	CIFAR10		IMAGENET	
				SN-GAN	SsGAN	SN-GAN	SsGAN
1	0.0	0.900	1	121.047	<b>25.798</b>	183.364	<b>80.669</b>
			2	28.107	<b>26.984</b>	85.130	<b>83.081</b>
	0.5	0.999	1	78.542	<b>25.890</b>	104.731	<b>91.628</b>
10	0.0	0.900	1	188.523	<b>28.481</b>	227.041	<b>85.376</b>
			2	29.110	<b>27.735</b>	227.738	<b>80.822</b>
	0.5	0.999	1	117.671	<b>25.216</b>	242.705	<b>144.353</b>

Table 2: FID comparisons for unconditional GANs under various hyper-parameter settings with gradient penalty. Self-supervision loss reduces the sensitivity of GAN training to hyper-parameters.

$\beta_1$	$\beta_2$	Disc_iters	CIFAR10		IMAGENET	
			SN-GAN	SsGAN	SN-GAN	SsGAN
0.0	0.900	1	87.864	<b>19.645</b>	129.964	<b>86.095</b>
		2	20.245	<b>17.885</b>	80.047	<b>70.637</b>
0.5	0.999	1	86.869	<b>18.235</b>	201.936	<b>99.968</b>

Table 3: FID comparisons for unconditional GANs under various hyper-parameter settings with spectral normalization. Self-supervision loss reduces the sensitivity of GAN training to hyper-parameters.

## B.2 Qualitative Comparisons of the Generative Samples

Figure 7 shows the generated examples from different GANs for IMAGENET. Although the baseline unconditional GAN generates sharp images, majority of the generated objects e.g. dogs are distorted. With the supervised labels, P-cGAN’s samples are improved, e.g. we see the dog faces and diverse objects. SsGAN generates similar quality images to the conditional P-cGAN.

## B.3 Comparison of Representation Quality

We compare the representation qualities of the discriminator’s intermediate layers. The comparison is performed by training logistic regression classifiers on top of the feature maps from each ResNet block to perform the 1000-way IMAGENET classification task, as proposed in Zhang et al. [27]. The representation features are spatially resized to around 9000 dimension with adaptive max pooling. We report top-1 classification accuracy.

We compare the SN-GAN, P-cGAN and SsGAN (SBN) models with the best FID score from Section B.1. We also consider ablating the GAN loss from SsGAN, to yield a rotation only model (Rotation). Rotation uses the same architecture and hyper parameters as the SsGAN discriminator. Evaluation is performed on three independent models with different random seeds, where the mean accuracy and the standard deviation are reported. We tune the learning rate for linear regression on a validation set created by holding out 10% of the training set. The linear regression results are insensitive to the hyperparameters. On CIFAR10, we use SGD with batch size 256, momentum 0.9 and learning rate 0.1. We drop the learning rates by a factor of 10 after every 30 epochs. We train in total for 100 epochs, where the learning rates are dropped by a factor of 10 after every 30 epochs. On IMAGENET, we use SGD with batch size 64, momentum 0.9 and learning rate 0.01. We train in total for 30 epochs, where the learning rates are dropped by a factor of 10 after every 10 epochs.

Table 4 shows the top-1 classification accuracy on CIFAR10. Figure 8 shows the results as a function of training steps of the original model. We plot the quality of representations extracted from each of the 4 blocks for ResNet CIFAR10 architecture. The conditional model, P-cGAN, produces similar results to the other unsupervised model GAN and Rotation model. The SsGAN outperforms the other models under most conditions.

Table 5 and Figure 9 show results on IMAGENET. There are six blocks in the IMAGENET Resnet architecture. For SN-GAN and Rotation, block3 performs best, for P-cGAN and SsGAN, block 5 performs best. It is expected that P-cGAN would benefit from representations closer to the classification layer because it sees the actual labels being used in the downstream representation




Figure 6: Performance under different  $\alpha$  values.

evaluation. In block5, we observe that SsGAN and P-cGAN representation quality is not improving after 200K steps, while SsGAN improves the representation quality for all IM training steps. Overall, representations are improved with self-supervised GAN and conditional GAN, which correlated with their improved samples.

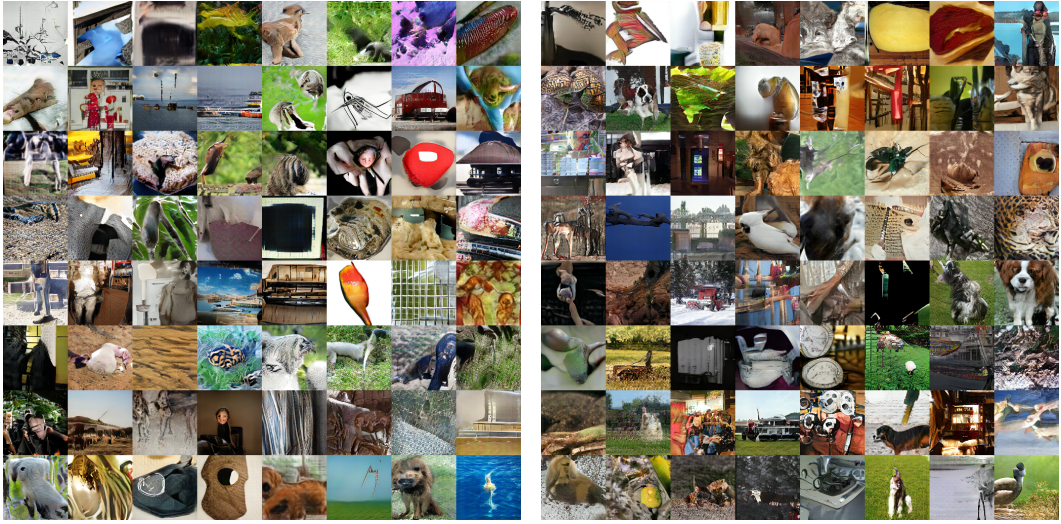
Table 6 shows the comparison of our method to other self supervised learning methods. We select the best run according to the validation results. Overall, the SsGAN model achieves comparable results on IMAGENET. Note that our method is not tuned for representation, but rather for image generation. Among those methods, BiGAN [33] learns the representation from GANs as well. BiGAN learns the representation with an additional encoder network, while our method is arguably simpler because it extracts the representation directly from the discriminator network.

Method	SN-GAN	P-cGAN	Rotation	SsGAN (sBN)
Block0	$0.719 \pm 0.002$	$0.719 \pm 0.003$	$0.710 \pm 0.002$	$0.721 \pm 0.002$
Block1	$0.762 \pm 0.001$	$0.759 \pm 0.003$	$0.749 \pm 0.003$	$0.774 \pm 0.003$
Block2	<b><math>0.778 \pm 0.001</math></b>	$0.776 \pm 0.005$	<b><math>0.762 \pm 0.003</math></b>	$0.796 \pm 0.005$
Block3	$0.776 \pm 0.005$	<b><math>0.780 \pm 0.006</math></b>	$0.752 \pm 0.006$	<b><math>0.799 \pm 0.003</math></b>

Table 4: Top-1 accuracy on CIFAR10.


Method	SN-GAN	P-cGAN	Rotation	SsGAN (sBN)
Block0	$0.084 \pm 0.076$	$0.169 \pm 0.001$	$0.172 \pm 0.002$	$0.170 \pm 0.001$
Block1	$0.064 \pm 0.107$	$0.190 \pm 0.004$	$0.169 \pm 0.005$	$0.198 \pm 0.003$
Block2	$0.078 \pm 0.134$	$0.232 \pm 0.002$	$0.202 \pm 0.004$	$0.250 \pm 0.001$
Block3	<b><math>0.087 \pm 0.149</math></b>	$0.271 \pm 0.003$	<b><math>0.239 \pm 0.005</math></b>	$0.284 \pm 0.002$
Block4	$0.0744 \pm 0.127$	$0.273 \pm 0.019$	$0.180 \pm 0.003$	$0.345 \pm 0.003$
Block5	$0.075 \pm 0.128$	<b><math>0.324 \pm 0.007</math></b>	$0.176 \pm 0.033$	<b><math>0.353 \pm 0.006</math></b>

Table 5: Top-1 accuracy on IMAGENET for all Resnet blocks.



(a) SN-GAN (FID=59.3)

(b) P-cGAN (FID=43.0)



(c) SsGAN (FID=47.8)

(d) SsGAN (sBN) (FID=43.9)

Figure 7: Random samples generated by GANs described in this paper on IMAGENET (at 1M steps).

Method	Context [26]	BiGAN [33]	Colorization [27]	RotNet [19]	SsGAN (sBN)
Accuracy	0.317	0.310	0.315	0.387	0.359

Table 6: Comparison with other self-supervised representation learning methods by top-1 accuracy on IMAGENET.




Figure 8: Top 1 accuracy on CIFAR10 with training steps from 10k to 100k.




Figure 9: Top 1 accuracy on IMAGENET validation set with training steps from 10k to 1M.