DuelGAN: A Duel Between Two Discriminators Stabilizes the GAN Training

Jiaheng Wei*[0000-0003-3573-1711], Minghao Liu*[0000-0002-1629-2429], Jiahao Luo[0000-0003-3374-6135], Andrew Zhu[0000-0002-6664-3215], James Davis[0000-0002-1413-2616], and Yang Liu[0000-0001-8420-6011]

University of California, Santa Cruz, Santa Cruz CA 95060, USA *: equal contributions Correspondence to yangliu@ucsc.edu

Abstract. In this paper, we introduce DuelGAN, a generative adversarial network (GAN) solution to improve the stability of the generated samples and to mitigate mode collapse. Built upon the Vanilla GAN's two-player game between the discriminator D_1 and the generator G, we introduce a peer discriminator D_2 to the min-max game. Similar to previous work using two discriminators, the first role of both D_1 , D_2 is to distinguish between generated samples and real ones, while the generator tries to generate high-quality samples which are able to fool both discriminators. Different from existing methods, we introduce a duel between D_1 and D_2 to discourage their agreement and therefore increase the level of diversity of the generated samples. This property alleviates the issue of early mode collapse by preventing D_1 and D_2 from converging too fast. We provide theoretical analysis for the equilibrium of the min-max game formed among G, D_1, D_2 . We offer convergence behavior of DuelGAN as well as stability of the min-max game. It's worth mentioning that DuelGAN operates in the unsupervised setting, and the duel between D_1 and D_2 does not need any label supervision. Experiments results on a synthetic dataset and on real-world image datasets (MNIST, Fashion MNIST, CIFAR-10, STL-10, CelebA, VGG) demonstrate that DuelGAN outperforms competitive baseline work in generating diverse and highquality samples, while only introduces negligible computation cost. Our code is publicly available at https://github.com/UCSC-REAL/DuelGAN.

Keywords: GAN, Image Generation, Peer Discriminator, Stability.

1 Introduction

Vanilla GAN (Generative Adversarial Nets [16]) proposed a data generating framework through an adversarial process which has achieved great success in image generation [16,23,51,15,6,63,49,38,47,4,18,12], image translation [64,62,11,54], and other real-life applications [22,35,48,43,3,58,29,42,61,31,53]. However, training Vanilla GAN is usually accompanied with a number of common problems, for example, vanishing gradients, mode collapse and failure to converge. Unfortunately, none of these issues have been completely addressed. There is a large

amount of follow up work on Vanilla GAN. Due to space limitations, we only discuss the two most related lines of works.

Stable and Diverse GAN Training. Several stabilization techniques have been implemented in GAN variants. Modifying architectures is the most extensively explored category. Radford et al. [45] make use of convolutional and convolutional-transpose layer in training the discriminator and generator. Karras et al. [23] adopt a hierarchical architecture and trains the discriminator and generator with progressively increasing size. Huang et al. [21] proposed a generative model which consists of a top-down stack of GANs. Chen et al. [8] split the generator into the noise prior and also latent variables. The optimization task includes maximizing the mutual information between latent variables and the observation. Designing suitable loss functions is another favored technique. Successful designs include f-divergence based GAN [41,36] (these two approaches replace loss functions of GAN by estimated variational f-divergence or leastsquare loss respectively), introducing auxiliary terms in the loss function [37] and integral probability metric based GAN [5,18,26,44]. A detailed survey of methods for stabilizing GANs exists [57].

Multi-Player GANs. Multi-player GANs explore the situation where there are multiple generators or multiple discriminators. The first published work to introduce multiple discriminators to GANs is multi-adversarial networks, in which discriminators can range from an unfavorable adversary to a forgiving teacher [13]. Nguyen et al. [40] formulate D2GAN, a three-player min-max game which utilizes a combination of Kullback-Leibler (KL) and reverse KL divergences in the objective function and is the most closely related to our work. Albuquerque et al. [1] show that training GAN variants with multiple discriminators is a practical approach even though extra capacity and computational cost are needed. Employing multiple generators and one discriminator to overcome the mode collapse issue and encourages diverse images has also been proposed [20,14].

In contrast to the above existing work, we demonstrate the possibility of improving GAN training with a computationally light modification by adding only one competing discriminator. We introduce a duel game among two discriminators and demonstrate the benefits of doing so in stabilizing and diversifying the training. Our main contributions summarize as follows:

- We introduce a duel between two discriminators to encourage diverse predictions and avoid early failure. The intuition is that predictions with high consensus will be discouraged, and effectively both discriminators are rewarded for having diverse predictions. The introduced game between the two discriminators results in a different convergence pattern for the generator.
- Theoretically, we derive the equilibrium for discriminators and the generator. We show how DuelGAN alleviates the vanishing gradient issue and mode collapse intuitively and empirically. We derive evidence for how the peer discriminator helps the dynamics of the learning. In addition, we demonstrate

that if the peer discriminator is better than a random guess classifier, the intermediate game and the objective function in DuelGAN are stable/robust to a bad peer discriminator.

- Experimental results on a synthetic dataset validate that DuelGAN addresses mode collapse. Results on real datasets demonstrate that DuelGAN generates high-quality image samples compared with baseline works. Besides, the introduced duel-game could also be viewed as a regularizer which complements well with existing methods and further improves the performance.

2 Background

We first review Vanilla GAN and D2GAN, which are the most relevant to understanding our proposed DuelGAN.

Vanilla GAN [16]. Let $\{x_i\}_{i=1}^n \subseteq \mathcal{X}$ denote the given training dataset drawn from the unknown distribution p_{data} . Traditional GAN formulates a two-player game: a discriminator D and a generator G. To learn the generator G's distribution over \mathcal{X} , G maps a prior noise distribution $p_z(z)$ to the data space. $\forall x \in \mathcal{X}, D(x)$ returns the probability that x belongs to p_{data} rather than p_g , where p_g denotes the distribution of G(z) implicitly defined by G. GAN trains D to maximize the probability of assigning the correct label to both training samples and those from the generator G. Meanwhile, GAN trains G to minimize $\log(1 - D(G(z)))$.

$$\min_{G} \max_{D} V(D,G) = \mathbb{E}_{x \sim p_{\text{data}}}[\log D(x)] + \mathbb{E}_{z \sim p_{z}}\left[\log\left(1 - D(G(z))\right)\right].$$
(1)

D2GAN [40]. D2GAN is the most closely related method to DuelGAN. This three-player game aims to solve the mode collapse issue and the optimization task is equivalent to minimize both KL divergence and Reverse-KL divergence between p_{data} and p_g . The formulation of D2GAN comes as follows:

$$\min_{G} \max_{D_1, D_2} V(D_1, D_2, G) = \alpha \cdot \mathbb{E}_{x \sim p_{\text{data}}} [\log D_1(x)] + \mathbb{E}_{z \sim p_z} \left[-D_1(G(z)) \right] \\
+ \mathbb{E}_{x \sim p_{\text{data}}} \left[-D_2(x) \right] + \beta \cdot \mathbb{E}_{z \sim p_z} \left[\log D_2(G(z)) \right], \quad (2)$$

with $0 \leq \alpha, \beta \leq 1$. Given a sample x in data space, $D_1(x)$ rewards a high score if x is drawn from p_{data} , and gives a low score if generated from the generator distribution p_g . In contrast, $D_2(x)$ returns a high score for x generated from p_g and gives a low score for a sample drawn from p_{data} . A large α induces multiple modes of p_g , but may include undesirable samples, while a large β induces the single mode. However, it is non-trivial to make a balance of these two terms in practice; if D_i fails/crashes (unstable), the overall training is influenced. And D_i won't be corrected since there is no interaction between discriminators!

Our work is similar to D2GAN in containing a pair of discriminators, instead of discriminators with different goals, we use identical discriminators and introduce a duel/competition between them.



Fig. 1. Illustration of the proposed DuelGAN. Compared with Vanilla GAN, DuelGAN has one more identical discriminator and a Duel Game between two discriminators. The introduced Duel Game induces diversified generated samples by discouraging the agreement between D_1 and D_2 . In D2GAN, although both discriminators are trained with different loss functions, they do not interfere with each other in the training.

3 DuelGAN: A Duel Between Two Discriminators

In this section, we first give the formulation and intuition of DuelGAN. Then we will present the equilibrium strategy of the generator and the discriminators.

3.1 Formulation

Similar to related works, we assume that the data follows the distribution p_{data} , our ultimate goal is to achieve $p_g = p_{\text{data}}$ where p_g is the generator's distribution. DuelGAN formulates a three-player game which consists of two discriminators D_1 , D_2 and one generator G. Denote by p_{duel} an equal mixture of p_{data} , p_g , $\forall x$: $p_{\text{duel}}(x) = [p_{\text{data}}(x) + p_g(x)]/2$. Recall that p_z denotes the prior noise distribution, now we are ready to formulate the min-max game of DuelGAN as follows:

$$\min_{G} \max_{D_1, D_2} \mathcal{L}(D_1, D_2, G)
= \min_{G} \max_{D_1, D_2} \mathbb{E}_{x \sim p_{\text{data}}} \left[\log D_1(x) \right] + \underline{\mathbb{E}_{x \sim p_{\text{data}}}} \left[\log D_2(x) \right]
+ \beta \cdot \underline{\text{Duel-D}} + \underline{\mathbb{E}_{z \sim p_z}} \left[\log \left(1 - D_1 \left(G(z) \right) \right) \right] + \underline{\mathbb{E}_{z \sim p_z}} \left[\log \left(1 - D_2 \left(G(z) \right) \right) \right], \quad (3)$$

where Duel-D introduces the duel (a peer competition game) among D_1, D_2 , defined as:

$$Duel-D = \mathbb{E}_{x \sim p_{duel}} \left[\underbrace{\ell \left(D_1(x), \mathbb{1} \left(D_2(x) > \frac{1}{2} \right) \right)}_{\text{Term 1a}} - \alpha \cdot \ell \left(D_1(x_{p_1}), \mathbb{1} \left(D_2(x_{p_2}) > \frac{1}{2} \right) \right) \right]}_{\text{Term 1b}} + \mathbb{E}_{x \sim p_{duel}} \left[\underbrace{\ell \left(D_2(x), \mathbb{1} \left(D_1(x) > \frac{1}{2} \right) \right)}_{\text{Term 2a}} - \alpha \cdot \underbrace{\ell \left(D_2(x_{p_1}), \mathbb{1} \left(D_1(x_{p_2}) > \frac{1}{2} \right) \right)}_{\text{Term 2b}} \right]_{\text{Term 2b}} \right]$$

$$(4)$$

In Duel-D, x_{p_1} and x_{p_2} are drawn randomly from p_{duel} and that \underline{x}, x_{p_1} and x_{p_2} are independent with each other. $\mathbb{1}(\cdot)$ is the indicator function, $\alpha, \beta \in [0, 1]$ are hyper-parameters controlling the disagreement level and the weight of the competition game between two discriminators, respectively. ℓ is an evaluation function, for simplicity, we adopt $\ell = \log(\cdot)$, as commonly used in other terms in the min-max game. Thus, for $i \in \{0, 1\}$, we have:

$$\ell(D_i(x), y) = \begin{cases} \log(D_i(x)) & \text{if } y = 1; \\ \log(1 - D_i(x)) & \text{if } y = 0. \end{cases}$$
(5)

To clarify the differences among Vanilla GAN [16], D2GAN [40] and DuelGAN, we use an workflow to illustrate in Figure 1. The key differences in DuelGAN's formulation can be summarized as follows:

- Compared with Vanilla GAN (see Eqn.(1)), DuelGAN (see Eqn.(3)) introduces a peer discriminator D_2 which has the same objective function as D appeared in Eqn.(1). An intermediate duel game Duel-D is added which will be explained below.
- The differences between D2GAN (see Eqn.(2)) and DuelGAN are highlighted with the underscores in red. Primarily, there is no interaction between discriminators in D2GAN, while our Duel-D term introduces another duel game between the discriminators, which we explain below. In addition to Duel-D, the objective function in DuelGAN encourages both discriminators to fit perfectly on both training samples and generated samples. While in D2GAN, one discriminator fits overly on training samples, the other fits overly on generated samples.

Competition Introduced by Duel-D. Duel-D bridges D_1 and D_2 by introducing 4 terms specified in Eqn.(4). Since we do not expect arbitrarily different discriminators, and both D_i s should play against the generator G, Term 1a and Term 2a encourage agreements between D_1 and D_2 . With only these two terms, D_1 and D_2 will eventually be encouraged to converge to agree with each other. Mode collapse issue remains a possibility. DuelGAN introduces Term 1b and Term 2b to the objective function which punish D_1 and D_2 from over-agreeing with each other (where the duel happens), especially at the early phase of training. Particularly, the Term 1b and 2b are evaluating the agreements of D_1 and D_2 on two entirely independent samples x_{p_1}, x_{p_2} . Because of the independence, the two discriminators' predictions should not match with high probability. Note that the calculation of Duel-D does not need label supervisions, which distinguishes our work from other works that introduces multiple discriminators but would require additional label supervisions [11].

We provide more details of our intuition as well as theoretical evidences of this property in Section 4.

3.2 The Max Game of Discriminators

Denote the true label of x as y = 1 if x comes from p_{data} , otherwise, y = 0. For any given generator G, let us first analyze the best responding/optimal discriminator $D_{i,G}^*(x)$, $i \in 1, 2$. We define the following quantities:

$$r_{i,G}(x) := \mathbb{P}_{x \sim p_{\text{duel}}} \left(\mathbb{1} \left(D_i(x) > \frac{1}{2} \right) = 1 \right), \quad p_{i,G} := \mathbb{E}_{x \sim p_{\text{duel}}} [r_{i,G}(x)], \quad (6)$$

where $r_{i,G}(x)$ represents the probability/confidence of x being categorized as the real data by D_i and $p_{i,G}$ is the expectation of $r_{i,G}(x)$ for $x \sim p_{duel}$. Let $\hat{r}^*_{i,G}(x) := r_{i,G}(x) - \alpha \cdot p_{i,G}$. Given discriminator D_i , when there is no confusion, we use D_j to denote the peer discriminator without telling $j \neq i$ in later sections.

Proposition 1 For G fixed, denote by $w := \beta \cdot (1-\alpha)$, the optimal discriminators D_1, D_2 are given by:

$$D_{i,G}^{*}(x) = \frac{p_{data}(x) + \beta \cdot \hat{r}_{j,G}^{*}(x) \cdot p_{duel}(x)}{p_{data}(x) + p_{a}(x) + w \cdot p_{duel}(x)}, \quad i = 1, 2.$$
(7)

3.3 The Min Game of the Generator

Remember that the training objective for D_i can be interpreted as maximizing the log-likelihood for estimating the conditional probability $\mathbb{P}(Y = y|x)$ where Yindicates whether x comes from p_{data} (with y = 1) or from p_g (with y = 0). With the introduce of Duel Game, the distributions p_{data} and p_g in the Vanilla GAN got changed due to the appearance of p_{duel} . Thus, we define the corresponding updated distributions in DuelGAN w.r.t. discriminator D_i as p_{data_i} and p_{g_i} , respectively. For a clean presentation, we defer the exact form of p_{data_i} , p_{g_i} in Appendix (Eqn.(22)).

Denote $C(G) := \max_D \mathcal{L}(G, D_1, D_2)$, the inner-max game (C(G)) can be rewritten as (straightforward in the proof of Proposition 1 which is available in the Appendix B.1):

$$C(G) = \mathbb{E}_{x \sim p_{\text{data}_1}} [\log D^*_{1,G}(x)] + \mathbb{E}_{x \sim p_{g_1}} [\log (1 - D^*_{1,G}(x))] \\ + \mathbb{E}_{x \sim p_{\text{data}_2}} [\log D^*_{2,G}(x)] + \mathbb{E}_{x \sim p_{g_2}} [\log (1 - D^*_{2,G}(x))].$$
(8)

Theorem 1. When $\alpha = 0$, $r_{j,G}(x) = 1/2$, the global minimum of the virtual training criterion C(G) is achieved if and only if $p_{data} = p_g$. At this point, C(G) achieves the value of $-\log 16$.

When $r_{j,G}(x) = 1/2$? Note that $r_{j,G}(x)$ is merely representing the probability that D_j classifies x to be real samples, $p_{j,G}$ is the probability that D_j classifies a random sample as the real one. Without loss of generality, we assume real and generated samples are of uniform/equal prior. At the very beginning of the training process, the discriminator can do well in distinguishing real or generated

samples, since the generator at this time generates low-quality samples. In this case, $r_{j,G}(x)$ is supposed to approach its max/min value, for example, $r_{j,G}(x) \to 0$ if x is from generated samples, and otherwise, $r_{j,G}(x) \to 1$. During the training process, the generator progressively tries to mislead the predictions made by discriminators, which means the discriminator can not decide whether the sample is being fake or real. Thus, $r_{j,G}(x) \to 1/2$. At this time, for $\alpha = 0, i = 1, 2$, we have:

$$D_{i,G}^*(x) = \frac{p_{\text{data}}(x) + \beta \cdot \hat{r}_{j,G}^*(x) \cdot p_{\text{duel}}(x)}{p_{\text{data}}(x) + p_g(x) + \beta \cdot p_{\text{duel}}(x)} \to \frac{p_{\text{data}}(x) + \frac{\beta}{2} \cdot p_{\text{duel}}(x)}{p_{\text{data}}(x) + p_g(x) + \beta \cdot p_{\text{duel}}(x)}.$$
(9)

This allows us to rewrite C(G)/2 as: $\mathbb{E}_{x \sim p_{\text{data}_i}} \left[\log \frac{p_{\text{data}}(x) + \frac{\beta}{2} \cdot p_{\text{duel}}(x)}{p_{\text{data}}(x) + p_g(x) + \beta \cdot p_{\text{duel}}(x)} \right] + \mathbb{E}_{x \sim p_{g_i}} \left[\log \frac{p_g(x) + \frac{\beta}{2} \cdot p_{\text{duel}}(x)}{p_{\text{data}}(x) + p_g(x) + \beta \cdot p_{\text{duel}}(x)} \right]$. Our subsequent proof is then based on the above reformulation. Due to space limits, we defer the formal statement of DuelGAN algorithm to the beginning of Appendix.

4 Properties of DuelGAN

In this section, we first illustrate how DuelGAN alleviates common issues in GAN training, for example, the vanishing gradients issue and the mode collapse issue. Then we present properties of DuelGAN including its stability guarantee and converging behavior.

4.1 DuelGAN and Common Issues in GAN Training

Vanishing Gradients Issue. In training GAN, discriminators might be too good for the generator to fool with and to improve progressively. When training with neural networks with back-propagation or gradient-based learning approaches, a vanishing small gradient only results in minor changes even with a large weight. As a result, the generator training may fail due to the vanishing gradients issue.

Mode Collapse Issue. Mode collapse refers to the phenomenon that the generator will rotate through a small set of output types. For the given fixed discriminator, the generator over-optimizes in each iteration. Thus, the corresponding discriminator fails to learn its way out of the trap.

How DuelGAN Alleviates the Vanish Gradient and Mode Collapse. DuelGAN alleviates the above two issues by preventing discriminators from "colluding" on its discrimination ability. In DuelGAN, for either discriminator D_i , recall that x_{p_1} and x_{p_2} are randomly drawn from p_{duel} which are independent from each other. Then the max game of D_i , given its peer discriminator D_j , is to perform the following task:

$$\max_{D_{i}} \mathcal{L}(D_{i},G)|_{D_{j}}$$

$$= \max_{D_{i}} \mathbb{E}_{x \sim p_{\text{data}}}[\log D_{i}(x)] + \mathbb{E}_{z \sim p_{z}}\left[\log\left(1 - D_{i}(G(z))\right)\right]$$

$$+ \beta \cdot \mathbb{E}_{x \sim p_{\text{duel}}}\left[\underbrace{\ell\left(D_{i}(x), \mathbb{1}\left(D_{j}(x) > \frac{1}{2}\right)\right)}_{\text{Term (b)}} \underbrace{-\alpha \cdot \ell\left(D_{i}(x_{p_{1}}), \mathbb{1}\left(D_{j}(x_{p_{2}}) > \frac{1}{2}\right)\right)}_{\text{Term (c)}}\right]. \quad (10)$$

Term (a) maximizes the probability of assigning the correct label to both real samples and generated samples. Term (b) maximizes the probability of matching predicted label with peer discriminator predicted ones. In other words, Term (b) controls the agreement level of D_i with respect to its peer discriminator D_j . However, note that Term (c) checks on the predictions of D_j on two different tasks x_{p_1}, x_{p_2} . When D_i agrees/fits overly on D_j , Term (c) returns a lower value if D_j 's predictions on these two different tasks are matching, mathematically, $\mathbb{1}(D_j(x_{p_1}) > 1/2) = \mathbb{1}(D_j(x_{p_2}) > 1/2)$. And Term (c) will return a high value if D_j 's predictions on these two different tasks are indeed different $\mathbb{1}(D_j(x_{p_1}) > 1/2) \neq \mathbb{1}(D_j(x_{p_2}) > 1/2)$. The weight α controls this disagreement level compared with Term (b) by referring to the fact that a larger α encourages more disagreement/diverse predictions from discriminators.

Based on the above intuitions, when two discriminators are of a high disagreement level, there exists a set S_{dis} such that $\mathbb{1}(D_i(x) > 1/2) \neq \mathbb{1}(D_j(x) > 1/2)$ for $x \in S_{\text{dis}}$ and S_{dis} is non-negligible. Therefore, there exists at least one discriminator D_i that can't perfectly predict labels (real/generated) of given data samples. The generator will then be provided with sufficient information, e.g., information or features that can be extracted from S_{dis} , to progress. This property helps us address the vanishing gradients issue. As for the mode collapse issue, suppose the over-optimized generator is able to find plausible outputs for both discriminators in the next generation. However, note that optimization is implemented on mini-batches in practice, the randomly selected samples x_{p_1}, x_{p_2} in Duel-D as well as the dynamically changing weights α, β can bring a certain degree of randomness in the next generation. Thus, rotating through this subset of the generator's output types could not force Term (\mathbf{c} to remain unchanged, so that the discriminators won't maintain a constant disagreement level and they unlikely get stuck in a local optimum.

To theoretically demonstrate how DuelGAN alleviates the mode collapse issue, we borrow from [32] and formalize the mode collapse issue as:

Definition 1. Data distribution p_{data} and a generator p_g in vanilla GAN exhibit (ϵ, δ) -mode collapse for some $0 \leq \epsilon < \delta \leq 1$ if \exists a set $S \subseteq \mathcal{X}$ such that $p_{data}(S) \geq \delta$ and $p_g(S) \leq \epsilon$.

A generator with a small δ or a large ϵ indicates a mild mode collapse issue (the better mode coverage). For D_i in DuelGAN, we could replace the distribution

 $p_{\text{data}}, p_{\text{g}}$ by $p_{\text{data},i} := c_{i,1} \cdot p_{\text{data}} + (1 - c_{i,1}) \cdot p_{\text{g}}$ and $p_{\text{g},i} := (1 - c_{i,2}) \cdot p_{\text{data}} + c_{i,2} \cdot p_{\text{g}}$, respectively, for some distributions $c_{i,1}(x), c_{i,2}(x) \geq \frac{1}{2}$, which implicitly encode the information of β , $\hat{r}^*_{i,G}$ (appeared in Eqn. (7)). We have:

Theorem 2. Suppose $\exists S \subseteq \mathcal{X}$ such that the Vanilla GAN has the (ϵ, δ) -mode collapse, DuelGAN has (ϵ_d, δ_d) -mode collapse given S, where:

$$\begin{split} \delta_d &:= \mathbb{E}_{x \in S}[c_{i,1}(x)] \cdot \mathbb{E}_{x \in S}[p_{data}(x)] + \mathbb{E}_{x \in S}[(1 - c_{i,1}(x))] \cdot \mathbb{E}_{x \in S}[p_g(x)] + Cov(c_{i,1}, p_{data} - p_g), \\ \epsilon_d &:= \mathbb{E}_{x \in S}[c_{i,2}(x)] \cdot \mathbb{E}_{x \in S}[p_g(x)] + \mathbb{E}_{x \in S}[(1 - c_{i,2}(x))] \cdot \mathbb{E}_{x \in S}[p_{data}(x)] + Cov(c_{i,2}, p_g - p_{data}). \end{split}$$

In the mode collapse scenario where D_j is fooled by G easily, we have $c_{i,1}, c_{i,2} \rightarrow (1 + \frac{1}{2})/2 = \frac{3}{4}$ and Covariance terms become 0. Note that $0 \leq \mathbb{E}_{x \in S}[p_g(x)] \leq \epsilon < \delta \leq \mathbb{E}_{x \in S}[p_{\text{data}}(x)] \leq 1$, we then have: $\delta_d < \delta$ and $\epsilon_d > \epsilon$, the mode collapse issue is alleviated with the introduce of duel game.

In Section 5.1, we use synthetic experiments to show that DuelGAN addresses mode collapse issues. And we include more empirical observations of the competition introduced by Duel-D in the Appendix C.5, i.e., the stability of the DuelGAN training, and the visualization of agreement levels between D_1 and D_2 due to the introduce of the duel game.

4.2 Stability and Convergence Behavior

In Section 4.1, we discussed the significant role of the introduced intermediate duel game. Now we discuss the potential downsides of introducing a second discriminator. Particularly, we are interested in understanding if the introduce of a peer discriminator D_j will disrupt the training and make the competition game with D_i unstable. Suppose D_j diverges from the optimum in the max game, in other words, the diverged peer discriminator \widetilde{D}_j fails to provide qualified verification label Y_i^* (given by $D_{i,G}^*$), and provides \widetilde{Y}_j instead. Denote by:

$$e_{\text{data},j} := \mathbb{P}(\widetilde{Y}_j = 0 | Y_j^* = 1), \quad e_{g,j} := \mathbb{P}(\widetilde{Y}_j = 1 | Y_j^* = 0).$$
 (11)

For any peer discriminator D_j , D_j may be a diverged peer discriminator D_j or an optimal one $D_{i,G}^*$, we denote the Duel Game of D_i given D_j as:

$$\operatorname{Duel}(D_i)|_{D_j} := \mathbb{E}_{x \sim p_{\operatorname{duel}}} \left[\ell \left(D_i(x), \mathbb{1} \left(D_j(x) > \frac{1}{2} \right) \right) - \alpha \cdot \ell \left(D_i(x_{p_1}), \mathbb{1} \left(D_j(x_{p_2}) > \frac{1}{2} \right) \right) \right]. \tag{12}$$

Theorem 3 explains the condition of stability (for D_i) when its peer discriminator in DuelGAN diverges from the corresponding optimum.

Theorem 3. Given G, suppose D_i has enough capacity, and at one step of Algorithm 1 (Appendix), if $e_{data,j} + e_{g,j} < 1$, $\alpha = 1$, the duel term of discriminator D_i is stable/robust with diverged peer discriminator \widetilde{D}_j . Mathematically,

$$\max_{D_i} Duel(D_i)|_{\widetilde{D}_j} \quad is \ equivalent \ with \quad \max_{D_i} Duel(D_i)|_{D_{j,G}^*}. \tag{13}$$

The above theorem implies that a diverging and degrading peer discriminator D_i will not disrupt the training of D_i .

Remark 1. Note that assuming uniform prior of real and generated samples, the condition to be stable is merely requiring that the proportion of false/wrong D_j 's prediction is less than a half (random guessing). This condition can be easily satisfied in practice. Thus, Theorem 3 provides the stability/robustness guarantee when the peer discriminator diverged from its optimum.

Build upon Theorem 1, with sufficiently small updates, Theorem 4 presents when p_q converges to p_{data} .

Theorem 4. If G and D_is have enough capacity, and at each step of Algorithm 1, D_is are allowed to reach its optimum given G, D_i is updated so as to improve the criterion in Eqn.(10), and p_g is updated so as to improve:

$$C(G) = \mathbb{E}_{x \sim p_{data_1}} [\log D_{1,G}^*(x)] + \mathbb{E}_{x \sim p_{g_1}} [\log (1 - D_{1,G}^*(x))] \\ + \mathbb{E}_{x \sim p_{data_2}} [\log D_{2,G}^*(x)] + \mathbb{E}_{x \sim p_{g_2}} [\log (1 - D_{2,G}^*(x))].$$
(14)

If $\beta = 0$ (remove the duel game to end the training and lead to convergence), we have $D_{1,G}^* = D_{2,G}^*$, p_g converges to p_{data} .

5 Experiments

In this section, we empirically validate the properties of DuelGAN through a set of datasets, including a synthetic task and several real world datasets ranging from hand-written digits to human faces.

5.1 Experiment Results on Synthetic Data

We apply the experiment and model structures proposed in UnrolledGAN [37] to investigate whether the DuelGAN design can prevent mode collapse. This experiment aims to generate eight 2D Gaussian distributions with a covariance matrix 0.02*I*, arranged around the same centroid with radius 2.0. Vanilla GAN fails on this example. D2GAN has been shown to outperform UnrolledGAN, so we include it as an alternate method which performs well.

Figure 2 shows symmetric KL-divergence, Wasserstein distance, and a visualization of results with Vanilla GAN, D2GAN, and DuelGAN. Knowing the target distribution p_{data} , we can employ symmetric KL divergence and Wasserstein distance, which calculate the distance between the true p_{data} and the normalized histogram of 10,000 generated points. On the left of Figure 2, the plots for symmetric KL-divergence and Wasserstein distance show that DuelGAN has a much better score than Vanilla GAN and slightly better than D2GAN.

On the right side of Figure 2 is a visualization of 512 generated blue samples points, together with red data points drawn from the true distribution. Vanilla GAN generates data points around only a single valid mode of the data distribution. D2GAN and DuelGAN distribute data around all eight mixture components, demonstrating the ability to resolve modal collapse in this case.



Fig. 2. Comparison of Vanilla GAN, D2GAN, and proposed DuelGAN on 2D synthesized data. The top-left graph shows the symmetric KL divergence over the training iterations, while the bottom left graph shows the Wasserstein distance. Both metrics compare the generated data points to data points drawn from the true target distribution. DuelGAN has the best performance. The right side visualizes generated blue data points and true red p_{data} data points. Note that Vanilla GAN has a clear mode collapse which both D2GAN and DuelGAN avoid.

5.2 Experiments on Real Image Datasets

We tested the proposed DuelGAN and baseline methods on MNIST [28], FashionMNIST [59], CIFAR-10 [27], STL-10 [10], CelebA [34] and VGGFace2 [7]. For quantitative evaluation, we adopt Fréchet Inception Distance (FID) [19] and Inception score(IS) [46] as the evaluation metric. FID summarizes the distance between the Inception features of the generated images and the real images. A lower FID indicates both better accuracy and higher diversity, so that a batch of generated images with good accuracy but identical to each other will have a poor FID score. A higher IS score indicates a higher generated image quality.

Baseline Methods. We reproduce/report the performance of a list of existing baseline methods, including: DCGAN [45], D2GAN [40], WGAN [18], DRAGAN [26], LSGAN [44], MicroBatchGAN [39], Dist-GAN [50], PresGAN [12], and QSNGAN [17]. We used the same generator and discriminator backbone for all the comparison methods in each dataset unless specified by the original author. We recorded the best performing checkpoints when evaluating each method.

Grey-Scale Images. MNIST [28] and FashionMNIST [59] are small grey-scale image datasets including 60,000 training and 10,000 testing 28×28 gray-scale images of hand-written digits and clothing. Since they are of small-scale, we adopt the shallow version of the generator and discriminators to generate the grey-scale images. We include a comprehensive comparison via FID score in Table 1. And the first two columns in Table 1 show our method has the best FID score among all tested methods. Figure 3 (left) shows FashionMNIST image results.

Table 1. Experiment FID score results of grey-scale image dataset: MNIST and FashionMNIST; natural scene image dataset: CIFAR-10 and STL-10; human face image dataset: CelebA and VGGFace2. Baseline results denoted with (*) were extracted from the original paper report, not independently run in our experiments.

	MNIST	FasionMNIST	CIFAR10	STL-10	CelebA	VGG
DCGAN [45]	19.86	24.78	27.45	59.79	17.38	49.99
WGAN* [18]	14.07	28.24	35.37	60.21	15.23	39.24
DRAGAN [26]	66.96	62.64	36.49	91.07	14.57	50.20
D2GAN [40]	22.20	29.33	27.38	54.12	17.30	20.67
Dist-GAN $*$ [50]	_	_	22.95	36.19	23.70	_
PresGAN [*] [12]	42.02	_	52.20	_	—	_
LSGAN [44]	23.80	43.00	51.42	70.37	15.35	55.96
MicroBatchGAN* [39]	17.10	_	77.70	_	34.50	_
QSNGAN* [17]	_	_	31.97	59.61	_	_
DuelGAN (ours)	7.87	21.73	21.55	51.37	13.95	19.05

Natural Scene Images. CIFAR-10 [27] and STL-10 [10] are natural scene RGB image datasets. CIFAR-10 includes 50K training and 10K testing 32×32 images with ten unique categories: airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck. STL-10 is sub-sampled from ImageNet, and has more diverse samples than CIFAR-10, containing about 100K 96×96 images. We adopt the deep version of the generator and discriminator to generate 32×32 RGB images. Table 1 middle two columns show FID score results. Note that the introduce of competitive Duel Game in two discriminator GAN setup, brings performance boost in all the experiments. Figure 3 (middle) shows STL-10 image results.



Fig. 3. Image results generated by proposed DuelGAN. Left: FashionMNIST, grey-scale clothing images; Middle: STL-10, natural scene images; Right: CelebA, large-scale celebrate face images.

Human Face Images. CelebA [34] and VGGFace2 [7] are large-scale face datasets. CelebA includes 162,770 training and 19,962 testing images of celebrity faces. VGGFace2 contains more than 3.3 million face images of celebrities caught

in the 'wild'. There are different lighting conditions, emotions, and viewing angles. We randomly choose 200 categories from VGGFace2 and trained on the reduced dataset. We adopt the deep version of the generator and discriminators to generate 32×32 RGB images on CelebA and 64×64 RGB images on VGGFace2. Table 1 last two columns show our method has the best FID score among tested methods. Figure 3 (right) shows CelebA image results.

Implementation Details Our model architecture adopts the same generator and discriminator backbone as DCGAN [45]. In DuelGAN, the newly introduced discriminator is a duplicate of the first one. DuelGAN achieves low FID scores and high IS scores when α and β are simply set to constant values. However we found that we could obtain an approximately 10% improvement through dynamic tuning. The parameter β controls the overall weight of Duel-D, while α punishes the condition when D_1 over-agrees with D_2 . In the early training phase, when we have an unstable generator and discriminator, we set α and β to 0. As training progresses, we gradually increase these parameters to a max value, which helps with vanishing gradients. After the midpoint of training we decrease these parameters to help the discriminators converge, until the parameters reach approximately 0 at the end of the training process. We adopt 0.3, 0.5 as the max value for α and β , respectively.

5.3 Duel Game as a Regularizer

Intuitively, the introduced duel game could be well applied to a large family of GAN variants defined w.r.t a single discriminator D_1 and a generator G. This is due to the fact that Eqn.(3) could be denoted by:

$$\min_{G} \max_{D_1, D_2} \mathcal{L}(D_1, D_2, G) = \min_{G} \max_{D_1, D_2} \left[\operatorname{GAN}(D_1) + \beta \cdot \operatorname{Duel-D} + \operatorname{GAN}(D_2) \right],$$
(15)

where $\operatorname{GAN}(D_i) := \mathbb{E}_{x \sim p_{\text{data}}} [\log D_i(x)] + \mathbb{E}_{z \sim p_z} [\log (1 - D_i(G(z)))]$. Thus, if we substitute the GAN loss $\operatorname{GAN}(D_i)$ by a state-of-the-art GAN variant, i.e., StyleGAN-ADA [24], one could view the duel game Duel-D as a regularizer.

We take the higher resolution version $(256 \times 256 \text{ RGB images})$ of CelebA [34] for illustration. Clearly in Table 2, StyleGAN-ADA reaches the state-of-the-art result on this task. And the introduced Duel-D regularizer could further improve its performance. Figure 4 shows the corresponding generated images.

We train StyleGAN-ADA with and w/o Dual-D regularizer on three datasets. Table 3 shows StyleGAN-ADA has higher recall score by introducing the Dual-D regularizer. This indicates the duel game regularizer could alleviate mode collapse issue and have a better mode coverage.

More Experiment Results. Due to space limits, we defer more experiment results to the Appendix C, including: an ablation study of hyper-parameters tuning; experiment validations about the stability of training; the visualization of the duel game between D_1 and D_2 .

Table 2. Experiment FID score results of CelebA (256×256 RGB images). Baseline results denoted with (*) were obtained from the original paper report.

Method	GLF* [60]	MSP* [30]	NCP-VAE* [2]	LSGM* [52]	StyleGAN-ADA [24]	StyleGAN-ADA+Duel-D
FID	41.80	35.00	24.79	7.22	4.85	4.32

Table 3. Comparisons of recall (\uparrow) on 3 datasets (StyleGAN-ADA w/o a duel game).

Methods	FFHQ	CelebA	CIFAR-10
StyleGAN-ADA	0.61	0.75	0.49
StyleGAN-ADA + Dual GAME	0.67	0.83	0.60



Fig. 4. Image results generated by proposed DuelGAN. (Trained on CelebA 256×256 RGB images. More generated images are deferred to the Appendix C.)

6 Conclusion

We propose DuelGAN which introduces a peer discriminator to Vanilla GAN. The role of the peer discriminator is to allow an intermediate game (duel game) between discriminators. Theoretical analysis demonstrates that the introduced duel game incentivizes incremental improvement, addresses vanishing gradients and mode collapse issues, punishes over-agreements among discriminators and is stable with diverged peer discriminator. Experimental results on a synthetic dataset and multiple real world datasets validate that DuelGAN produces high quality images, with lower error than competing techniques.

Acknowledgements. JHW and YL are partially supported by the National Science Foundation (NSF) under grants IIS-2007951, IIS-2143895, and CCF-2023495. MHL, JHL, and JD are supported in part by WISEautomotive through a ATC+ Program award from the Korean Ministry of Trade, Industry and Energy (MOTIE).

References

- Albuquerque, I., Monteiro, J., Doan, T., Considine, B., Falk, T., Mitliagkas, I.: Multiobjective training of generative adversarial networks with multiple discriminators. arXiv preprint arXiv:1901.08680 (2019) 2
- Aneja, J., Schwing, A., Kautz, J., Vahdat, A.: A contrastive learning approach for training variational autoencoder priors. Advances in Neural Information Processing Systems 34 (2021) 14
- Antipov, G., Baccouche, M., Dugelay, J.L.: Face aging with conditional generative adversarial networks. In: 2017 IEEE international conference on image processing (ICIP). pp. 2089–2093. IEEE (2017) 1
- Arbel, M., Sutherland, D., Bińkowski, M., Gretton, A.: On gradient regularizers for mmd gans. In: Advances in neural information processing systems. pp. 6700–6710 (2018) 1
- 5. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein gan. arXiv preprint arXiv:1701.07875 (2017) 2
- Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096 (2018) 1
- Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: International Conference on Automatic Face and Gesture Recognition (2018) 11, 12
- Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., Abbeel, P.: Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In: Advances in neural information processing systems. pp. 2172–2180 (2016) 2
- Cheng, H., Zhu, Z., Li, X., Gong, Y., Sun, X., Liu, Y.: Learning with instancedependent label noise: A sample sieve approach. In: International Conference on Learning Representations (2021), https://openreview.net/forum?id= 2VXyy9mIyU3 1
- Coates, A., Ng, A., Lee, H.: An Analysis of Single Layer Networks in Unsupervised Feature Learning. In: AISTATS (2011), https://cs.stanford.edu/~acoates/ papers/coatesleeng_aistats_2011.pdf 11, 12
- Dash, A., Gamboa, J.C.B., Ahmed, S., Liwicki, M., Afzal, M.Z.: Tac-gan-text conditioned auxiliary classifier generative adversarial network. arXiv preprint arXiv:1703.06412 (2017) 1, 5
- Dieng, A.B., Ruiz, F.J., Blei, D.M., Titsias, M.K.: Prescribed generative adversarial networks. arXiv preprint arXiv:1910.04302 (2019) 1, 11, 12
- Durugkar, I., Gemp, I., Mahadevan, S.: Generative multi-adversarial networks. arXiv preprint arXiv:1611.01673 (2016) 2
- Ghosh, A., Kulharia, V., Namboodiri, V.P., Torr, P.H., Dokania, P.K.: Multi-agent diverse generative adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8513–8521 (2018) 2
- Gong, X., Chang, S., Jiang, Y., Wang, Z.: Autogan: Neural architecture search for generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3224–3234 (2019) 1
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. pp. 2672–2680 (2014) 1, 3, 5, 7
- 17. Grassucci, E., Cicero, E., Comminiello, D.: Quaternion generative adversarial networks. arXiv preprint arXiv:2104.09630 (2021) 11, 12

- 16 Wei, Liu. et al.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. In: Advances in neural information processing systems. pp. 5767–5777 (2017) 1, 2, 11, 12
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: 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 (2017) 11
- Hoang, Q., Nguyen, T.D., Le, T., Phung, D.: Multi-generator generative adversarial nets. arXiv preprint arXiv:1708.02556 (2017) 2
- Huang, X., Li, Y., Poursaeed, O., Hopcroft, J., Belongie, S.: Stacked generative adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5077–5086 (2017) 2
- 22. Jin, Y., Zhang, J., Li, M., Tian, Y., Zhu, H., Fang, Z.: Towards the automatic anime characters creation with generative adversarial networks. arXiv preprint arXiv:1708.05509 (2017) 1
- 23. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196 (2017) 1, 2
- Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. Advances in Neural Information Processing Systems 33, 12104–12114 (2020) 13, 14, 8
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019) 8
- Kodali, N., Abernethy, J., Hays, J., Kira, Z.: On convergence and stability of gans. arXiv preprint arXiv:1705.07215 (2017) 2, 11, 12
- Krizhevsky, A., Hinton, G.: Convolutional deep belief networks on cifar-10. Unpublished manuscript 40(7), 1–9 (2010) 11, 12
- LeCun, Y., Cortes, C.: MNIST handwritten digit database (2010), http://yann. lecun.com/exdb/mnist/ 11
- Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4681–4690 (2017) 1
- Li, X., Lin, C., Li, R., Wang, C., Guerin, F.: Latent space factorisation and manipulation via matrix subspace projection. In: International Conference on Machine Learning, pp. 5916–5926. PMLR (2020) 14
- Li, Y., Liu, S., Yang, J., Yang, M.H.: Generative face completion. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3911–3919 (2017) 1
- Lin, Z., Khetan, A., Fanti, G., Oh, S.P.: The power of two samples in generative adversarial networks. arxiv 2017. arXiv preprint arXiv:1712.04086
- Liu, Y., Guo, H.: Peer loss functions: Learning from noisy labels without knowing noise rates. In: International conference on machine learning. pp. 6226–6236. PMLR (2020) 1
- Liu, Z., Luo, P., Wang, X., Tang, X.: Large-scale celebfaces attributes (celeba) dataset. Retrieved August 15(2018), 11 (2018) 11, 12, 13, 8
- Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. In: Advances in neural information processing systems. pp. 406–416 (2017) 1
- Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2794–2802 (2017) 2

- 37. Metz, L., Poole, B., Pfau, D., Sohl-Dickstein, J.: Unrolled generative adversarial networks. arXiv preprint arXiv:1611.02163 (2016) 2, 10
- Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. arXiv preprint arXiv:1802.05957 (2018) 1
- Mordido, G., Yang, H., Meinel, C.: microbatchgan: Stimulating diversity with multiadversarial discrimination. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3061–3070 (2020) 11, 12
- Nguyen, T., Le, T., Vu, H., Phung, D.: Dual discriminator generative adversarial nets. In: Advances in Neural Information Processing Systems. pp. 2670–2680 (2017) 2, 3, 5, 11, 12
- Nowozin, S., Cseke, B., Tomioka, R.: f-gan: Training generative neural samplers using variational divergence minimization. In: Advances in neural information processing systems. pp. 271–279 (2016) 2
- Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2536–2544 (2016) 1
- Perarnau, G., Van De Weijer, J., Raducanu, B., Álvarez, J.M.: Invertible conditional gans for image editing. arXiv preprint arXiv:1611.06355 (2016) 1
- Qi, G.J.: Loss-sensitive generative adversarial networks on lipschitz densities. International Journal of Computer Vision 128(5), 1118–1140 (2020) 2, 11, 12
- Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015) 2, 11, 12, 13, 7, 8
- 46. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. arXiv preprint arXiv:1606.03498 (2016) 11
- Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: Advances in Neural Information Processing Systems. pp. 11918– 11930 (2019) 1
- Taigman, Y., Polyak, A., Wolf, L.: Unsupervised cross-domain image generation. arXiv preprint arXiv:1611.02200 (2016) 1
- Tran, N.T., Bui, T.A., Cheung, N.M.: Dist-gan: An improved gan using distance constraints. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 370–385 (2018) 1
- Tran, N.T., Bui, T.A., Cheung, N.M.: Dist-gan: An improved gan using distance constraints. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 370–385 (2018) 11, 12
- Tran, N.T., Tran, V.H., Nguyen, B.N., Yang, L., Cheung, N.M.M.: Self-supervised gan: Analysis and improvement with multi-class minimax game. Advances in Neural Information Processing Systems 32, 13253–13264 (2019) 1
- Vahdat, A., Kreis, K., Kautz, J.: Score-based generative modeling in latent space. Advances in Neural Information Processing Systems 34 (2021) 14
- Vondrick, C., Pirsiavash, H., Torralba, A.: Generating videos with scene dynamics. In: Advances in neural information processing systems. pp. 613–621 (2016) 1
- Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018) 1
- 55. Wei, J., Liu, H., Liu, T., Niu, G., Liu, Y.: Understanding generalized label smoothing when learning with noisy labels. arXiv preprint arXiv:2106.04149 (2021) 1
- 56. Wei, J., Liu, Y.: When optimizing f-divergence is robust with label noise. arXiv preprint arXiv:2011.03687 (2020) 1

- 18 Wei, Liu. et al.
- Wiatrak, M., Albrecht, S.V., Nystrom, A.: Stabilizing generative adversarial networks: A survey. arXiv preprint arXiv:1910.00927 (2019) 2
- Wu, H., Zheng, S., Zhang, J., Huang, K.: Gp-gan: Towards realistic high-resolution image blending. In: Proceedings of the 27th ACM International Conference on Multimedia. pp. 2487–2495 (2019) 1
- 59. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747 (2017) 11
- Xiao, Z., Yan, Q., Amit, Y.: Generative latent flow. arXiv preprint arXiv:1905.10485 (2019) 14
- Yeh, R.A., Chen, C., Yian Lim, T., Schwing, A.G., Hasegawa-Johnson, M., Do, M.N.: Semantic image inpainting with deep generative models. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5485–5493 (2017) 1
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 5907–5915 (2017) 1
- Zhang, H., Zhang, Z., Odena, A., Lee, H.: Consistency regularization for generative adversarial networks. arXiv preprint arXiv:1910.12027 (2019) 1
- 64. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017) 1