On Characterizing GAN Convergence Through Proximal Duality Gap

Sahil Sidheekh [1] **Aroof Aimen** [1] **Narayanan C. Krishnan** [1]

Abstract

Despite the accomplishments of Generative Adversarial Networks (GANs) in modeling data distributions, training them remains a challenging task. A contributing factor to this difficulty is the non-intuitive nature of the GAN loss curves, which necessitates a subjective evaluation of the generated output to infer training progress. Recently, motivated by game theory, duality gap has been proposed as a domain agnostic measure to monitor GAN training. However, it is restricted to the setting when the GAN converges to a Nash equilibrium. But GANs need not always converge to a Nash equilibrium to model the data distribution. In this work, we extend the notion of duality gap to proximal duality gap that is applicable to the general context of training GANs where Nash equilibria may not exist. We show theoretically that the proximal duality gap is capable of monitoring the convergence of GANs to a wider spectrum of equilibria that subsumes Nash equilibria. We also theoretically establish the relationship between the proximal duality gap and the divergence between the real and generated data distributions for different GAN formulations. Our results provide new insights into the nature of GAN convergence. Finally, we validate experimentally the usefulness of proximal duality gap for monitoring and influencing GAN training.

1. Introduction

Generative modeling is an important machine learning paradigm, aiming to learn data distributions. The ability to parametrically model the true underlying distribution of real-world data from a given empirical distribution brings with it the power to generate new and unseen instances. Generative adversarial networks (GANs) are perhaps the most popular and successful of innovations for learning data distributions. A GAN formulates the generative modeling problem as a zero-sum game between two agents - a Discriminator (D) and a Generator (G). The discriminator aims to differentiate the fake samples produced by the generator from samples belonging to the true data distribution. On the other hand, the generator seeks to fool the discriminator by learning a mapping from an input noise space to the data space. The generator can also be viewed as performing adversarial attacks on the discriminator, exploiting the information leak through the discriminator and learning the real data distribution as the game proceeds to an equilibrium.

Formally, the GAN game is defined as :

$$\min_{\theta_g \in \Theta_G} \max_{\theta_d \in \Theta_D} V(D_{\theta_d}, G_{\theta_g}), \tag{1}$$

where the generator (parametrized by θ_g) and discriminator (parametrized by θ_d) are neural networks and V is the objective function that the agents seek to optimize. Different GAN formulations yield different expressions for V, each minimizing a unique divergence between the real and generated data distributions. The classic GAN formulation (Goodfellow et al., 2014) minimizes the JS divergence and is defined by :

$$V = \mathbb{E}_{\mathbf{x} \sim P_r}[\log(D(\mathbf{x}))] + \mathbb{E}_{\mathbf{x} \sim P_{\theta_g}}[\log(1 - D(\mathbf{x}))]$$

where P_r denotes the real data distribution and P_{θ_g} denotes the generated data distribution.

In any learning problem, the trajectory of the loss functions should indicate the goodness of the trained model. However, such intuitive inferences cannot be drawn from the loss curves of a GAN. This is because classical training of a GAN involves alternate gradient descent optimization of the objective function w.r.t the individual agents. Each optimization step of an agent alters its adversary's loss surface, resulting in non-intuitive loss curves for both the agents over time. Figure 1 shows discriminator and generator loss curves for a GAN when it (a) converges and (b) diverges. Ideally, losses should decrease during model convergence and increase during divergence. However, we observe a diminishing generator loss and an increasing discriminator loss when the GAN converges. When it diverges, there is an interplay between both the losses. These loss curves do not

[1]Department of Computer Science, Indian Institute of Technology, Ropar, India. Correspondence to: Sahil Sidheekh <2017csb1104@iitrpr.ac.in>, Narayanan C Krishnan <ckn@iitrpr.ac.in>.

Proceedings of the 38^{th} International Conference on Machine Learning, PMLR 139, 2021.

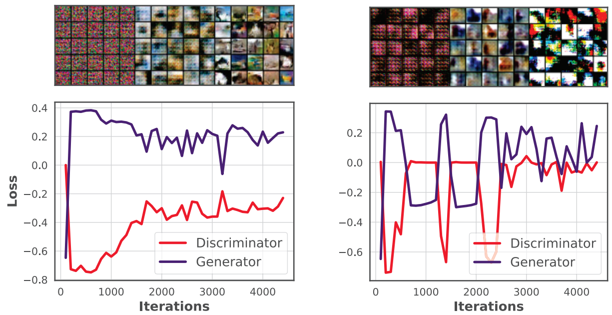

(a) Convergence (b) Divergence

Figure 1. Loss curves throughout the training progress of WGAN on the CIFAR-10 dataset.

give any insight into the gradual improvement or degradation in the GAN's performance. Thus, monitoring GANs often requires a subjective evaluation of the generated output. As a result, an exhaustive search over the architecture and hyperparameter space to find the delicate balance demanded by the GAN game becomes infeasible. This increases the complexity of GAN training that is already challenging due to the instabilities posed by the min-max gradient optimization. Objective measures capable of quantifying GAN training progress can reduce the training complexity.

Duality Gap (DG) (Grnarova et al., 2019) for GANs, motivated by principles of game theory, is a recently proposed objective measure for monitoring GAN training. The DG quantifies a GAN configuration's goodness in terms of the agents' ability to deviate from it in search of better optima. When a GAN converges to a Nash equilibrium, no agent can unilaterally deviate to find a better optima, and hence the DG would be zero. The ability to quantify convergence as well as the domain agnostic nature that requires no pre-trained models nor labeled data, makes DG a potentially powerful tool to monitor GAN training.

However, DG relies on the notion that GANs converge to Nash equilibria. On the contrary, recent studies (Farnia & Ozdaglar, 2020; Berard et al., 2020) suggest that a Nash equilibrium need not always exist for a GAN, especially when trained under regularized environments that most modern GAN formulations employ. GANs can converge to stable stationary points that are not Nash equilibria, all the while producing realistic data samples with high fidelity. This weakens the foundation upon which the notion of DG as a performance monitoring tool for GANs is built, eliciting the following questions: Can GANs capture the real data distribution even at non-Nash critical points? If so, would the DG at such stationary points be zero? If not, how do we monitor the GAN training in such situations?

In this work, we study the above questions by introducing the notion of proximal duality gap for GANs that is generalizable to scenarios where Nash equilibria may not exist. Our work is motivated by the notion of proximal equilibria for GANs that serves as a general concept for characterizing GAN optimality (Farnia & Ozdaglar, 2020). We define the DG in terms of the agents ability to optimize the proximal GAN objective (see Eq. 9) and call it as the Proximal Duality Gap (DG^λ). A proximal equilibrium for the GAN game (Eq. 1) is a Nash equilibrium w.r.t the proximal objective. Thus, whenever the GAN game attains a proximal equilibrium, DG^λ will tend to zero, indicating model convergence. As all Nash equilibria form a subset of proximal equilibria, DG^λ serves as a generic and robust measure that can quantify GAN convergence in the wild.

Overall, we make the following contributions:

- We present an acute limitation of DG for monitoring GAN training.
- We propose a theoretically grounded and robust extension - DG^λ, that overcomes this limitation and is also applicable to the broader context, when GANs converge to a non-Nash equilibrium.
- Using DG^λ, we derive insights into the nature of GAN convergence. Specifically, we study the relationship between the quality of the learned data distribution and the game equilibria. We show that for various GANs, a configuration (θ_d, θ_g) where $P_{\theta_g} = P_r$ corresponds to a Stackelberg equilibrium.
- We demonstrate through experiments, the proficiency of DG^λ for monitoring and influencing GAN training.

2. Related Work

Motivated by the non-inferrable nature of GAN loss curves, developing extrinsic measures to monitor GAN training has emerged as an active research area (Borji, 2019; Lucic et al., 2018; Olsson et al., 2018). Existing measures such as average log-likelihood (Goodfellow et al., 2014; Theis et al., 2016), Inception Score (IS) (Salimans et al., 2016), Frechet Inception Distance (FID) (Heusel et al., 2017) evaluate the output of GANs, but require pre-trained models. Further, these measures, including the more recent ones such as precision and recall (Sajjadi et al., 2018; Kynkäänniemi et al., 2019), density and coverage (Tolstikhin et al., 2017; Naeem et al., 2020) do not monitor the training progress nor characterize the equilibria of the GAN game.

Duality Gap (DG) (Grnarova et al., 2019) is a recently proposed domain agnostic and computationally feasible metric for monitoring and evaluating GAN training. DG is zero when the GAN converges to a Nash equilibrium making it an attractive metric for objectively monitoring GAN training.

However, DG has a fundamental limitation with its estimation process due to vanishing gradients. Adding perturbations to the GAN configuration before estimating the duality gap (perturbed duality gap) helps to overcome this issue (Sidheekh et al., 2020). But both these approaches assume the convergence of GANs to Nash equilibria, which may not always be the case, especially for high dimensional datasets (as we demonstrate in the next section). Thus, limiting the applicability of DG and perturbed DG for monitoring GAN training.

3. Background

3.1. A Brief Overview of GAN formulations

Classic GAN : The min-max objective in the classic GAN (Goodfellow et al., 2014) formulation is :

$$V_c = \frac{1}{2}\mathbb{E}_{\mathbf{x}\sim P_r}[\log D(\mathbf{x})] + \frac{1}{2}\mathbb{E}_{\mathbf{x}\sim P_{\theta_g}}[\log(1 - D(\mathbf{x}))] \quad (2)$$

where the probabilistic discriminator D outputs the likelihood of the input data point belonging to the real data distribution. The discriminator's objective is to maximize the log-likelihood to learn the conditional probability $P(y|\mathbf{x})$, where $y = 0$ and 1 indicate a fake and real data point respectively. Minimizing the above objective w.r.t the generator for the optimal discriminator is equivalent to minimizing the Jenson Shannon divergence (JSD) between P_{θ_g} and P_r.

F-GAN: F-GAN (Nowozin et al., 2016) is the generalization of the classic GAN to minimize arbitrary $f-$ divergences by incorporating an extension of the variational divergence estimation framework (Nguyen et al., 2010). For a convex, lower semi-continuous function $f : \mathbb{R}_+ \to \mathbb{R}$ that satisfies $f(1) = 0$, the $f-$ divergence between distributions P, and Q is,

$$D_f(P||Q) = \int p(x) f\left(\frac{q(x)}{p(x)}\right) dx \quad (3)$$

The F-GAN objective that minimizes the $f-$ divergence (D_f) between P_{θ_g} and P_r is defined as

$$V_f = \mathbb{E}_{\mathbf{x}\sim P_r}[D(\mathbf{x})] - \mathbb{E}_{\mathbf{x}\sim P_{\theta_g}}[f^*(D(\mathbf{x}))] \quad (4)$$

where $f^*(x) = \sup_{t\in \mathrm{Dom}(f)} \{xt - f(t)\}$ is the Fenchel-conjugate of f. The classic GAN is a special case of F-GAN when $f(t) = t\log(t) - (t+1)\log\left(\frac{t+1}{2}\right)$.

Wasserstein GAN (WGAN): WGAN formulates the GAN game as a minimization of the optimal transport cost, the Wasserstein distance, between P_r and P_{θ_g}, a more efficient cost function to learn data distributions having support on low dimensional manifolds (Arjovsky et al., 2017). Specifically, the Kantorovich-Rubinstein duality is used to arrive at the Wasserstein-1 (Earth Movers) distance between the distributions defined as: $\sup_{||D||_{L\leq 1}} \mathbb{E}_{\mathbf{x}\sim P_r}[D(\mathbf{x})] - \mathbb{E}_{\mathbf{x}\sim P_{\theta_g}}[D(\mathbf{x})]$, where the supremum is over all 1-Lipschitz discriminators. In practice, the Lipschitz constraint is enforced through weight clipping resulting in the following WGAN objective:

$$V_{w_1} = \mathbb{E}_{\mathbf{x}\sim P_r}[D(\mathbf{x})] - \mathbb{E}_{\mathbf{x}\sim P_{\theta_g}}[D(\mathbf{x})] \quad (5)$$

The WGAN formulation is also extended to a general transport cost (Farnia & Tse, 2018) $c(\mathbf{x}, \mathbf{y})$ by constraining the discriminators to be c-concave as

$$V_w = \mathbb{E}_{\mathbf{x}\sim P_r}[D(\mathbf{x})] - \mathbb{E}_{\mathbf{x}\sim P_{\theta_g}}[D^c(\mathbf{x})], \quad (6)$$

where D^c is the c-transform of the discriminator D, i.e.,

$$D^c(\mathbf{x}) = \sup_{\mathbf{y}} \{D(\mathbf{y}) - c(\mathbf{x},\mathbf{y})\} \quad (7)$$

and the Wasserstein distance between P_{θ_g} and P_r is :

$$W_c(P_{\theta_g}||P_r) = \sup_{D-c-concave} \mathbb{E}_{\mathbf{x}\sim P_r}[D(\mathbf{x})] - \mathbb{E}_{\mathbf{x}\sim P_{\theta_g}}[D^c(\mathbf{x})] \quad (8)$$

Despite the added stability of Wasserstein distance over other divergences, training GANs remains an arduous task. This has motivated efforts towards understanding the nature of GAN convergence.

3.2. Understanding GAN convergence

Classical Notion of GAN Equilibrium

Traditionally, the GAN game was expected to converge to a pure Nash equilibrium, a configuration (θ_d^*, θ_g^*) that is optimal for both the players i.e.

$$\max_{\theta_d} V(D_{\theta_d}, G_{\theta_g^*}) = \min_{\theta_g} V(D_{\theta_d^*}, G_{\theta_g}) = V(D_{\theta_d^*}, G_{\theta_g^*})$$

A GAN having unbounded capacity learns the true data distribution at such a solution (Goodfellow et al., 2014). However, a pure Nash equilibrium need not always exist for a zero sum game (Nash, 1950). Only an extended notion - the mixed strategy Nash equilibrium (MNE) is guaranteed to exist. Recent GAN formulations explicitly seek the MNE (Arora et al., 2017; Hsieh et al., 2019). As a mixed strategy gives a distribution over the model parameters, the stationary point to which the GAN converges need not be individually optimal for both the players.

GANs Need Not Converge to Nash Equilibria

GAN convergence has also been well studied from an optimization perspective using the notion of stability (Daskalakis et al., 2018; Fiez et al., 2019; Nouiehed et al., 2019; Zhang et al., 2019b; Mazumdar et al., 2019; Mokhtari et al., 2020; Lin et al., 2020). As GAN formulations are non-convex, only local surrogates of equilibria may be attainable while employing gradient based optimization. A (differential) local Nash equilibrium (LNE) (θ_d^*, θ_g^*) satisfies two properties: 1) $\nabla_{\theta_d} V(\theta_d^*, \theta_g^*) = \nabla_{\theta_g} V(\theta_d^*, \theta_g^*) = 0$ and 2) $\nabla^2_{\theta_d} V(\theta_d^*, \theta_g^*) \prec 0$, $\nabla^2_{\theta_g} V(\theta_d^*, \theta_g^*) \succ 0$. However, recent literature (Mazumdar & Ratliff, 2018; Adolphs et al., 2019) suggests the existence of many stable attractors that are not LNE for the alternate gradient descent optimization, nor produce realistic data and proposes methods to escape these attractors. (Fiez et al., 2019; Jin et al., 2020) study the gradient dynamics of sequential games and establish that their only stable attractors are Stackelberg equilibria. However, the relationship between the learned data distribution and the game configuration has not been well studied. Further, recent empirical studies (Berard et al., 2020; Farnia & Ozdaglar, 2020) also suggest that *GANs need not attain an LNE to produce realistic data*.

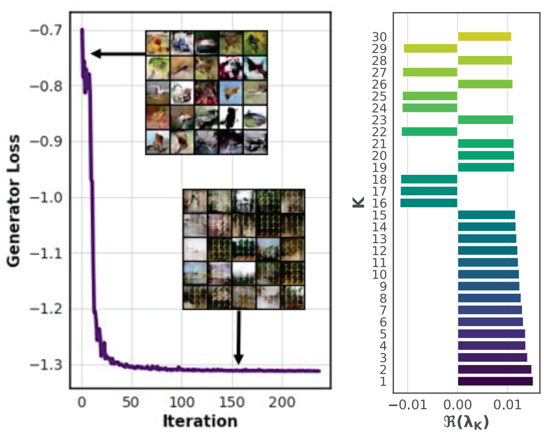

Figure 2. The high fidelity images outputted by a converged GAN deteriorates on optimizing only w.r.t the generator while attaining a lower loss (left), indicating that the GAN has not converged to a Nash equilibrium; confirmed by the positive and negative eigenvalues of the Hessian (right).

We verify the above hypothesis by training a spectral normalized GAN (SNGAN) on the CIFAR-10 dataset for 100 epochs ensuring that the models have converged producing high fidelity samples. As a local Nash equilibrium is locally optimal for the agents, optimizing the objective function w.r.t any individual player should not facilitate departure of the player from such a point. However, as demonstrated in Figure 2, the generator deviates from the attained equilibrium on further optimizing the objective function w.r.t only the generator. While the generator attains a lower cost, there is a clear deterioration in the generated samples' quality. This implies that the GAN has not converged to an LNE. We further strengthen the claim by verifying the top-K eigenvalues (λ_K) (by magnitude) of the Hessian of the objective w.r.t the generator's parameters. However, as shown in Figure 2, the Hessian has both positive as well as negative eigen values violating the second property of an LNE, thus emphasizing that GANs can converge to non-Nash attractors, all the while producing high fidelity samples. Similar results on GANs trained for MNIST and CELEB-A datasets are discussed in the supplementary material.

Proximal Equilibria for GANs

Majority of studies attempting to characterize the equilibria for GANs assume unbounded capacity for the models (realizable setting) (Goodfellow et al., 2014; Arora et al., 2017; Hsieh et al., 2019). However, in practice, most GAN architectures (Brock et al., 2019; Zhang et al., 2019a; Miyato et al., 2018; Gulrajani et al., 2017; Arjovsky et al., 2017; Radford et al., 2016) employ normalization and regularization to achieve state of the art performance. A recent study (Farnia & Ozdaglar, 2020) on GAN convergence under the non-realizable setting proposes a more generic notion of equilibira - the Proximal Equilibria (PE). This notion of equilibira is derived from a sequential game-play perspective for a GAN, for which a Stackelberg equilibrium is guaranteed to exist under mild continuity assumptions (Jin et al., 2020; Fiez et al., 2019). Farnia and Ozdaglar define a proximal operator over the original GAN objective, allowing the discriminator to be optimal in a neighbourhood (controlled by λ),

$$V^{\lambda}(D_{\theta_d}, G_{\theta_g}) = \max_{\tilde{\theta_d} \in \Theta_D} V(D_{\tilde{\theta_d}}, G_{\theta_g}) - \lambda ||D_{\tilde{\theta_d}} - D_{\theta_d}||^2 \tag{9}$$

A proximal equilibrium is defined as the Nash equilibrium for the objective V^{λ}, which is guaranteed to exist (Farnia & Ozdaglar, 2020). Formally, a configuration (θ_d^*, θ_g^*) of the GAN game (Eq 1) is called a $\lambda-$proximal equilibrium if and only if $\forall\ \theta_d, \theta_g$,

$$\begin{aligned} V(D_{\theta_d}, G_{\theta_g^*}) &\leq V(D_{\theta_d^*}, G_{\theta_g^*}) \\ &\leq \max_{\tilde{\theta_d} \in \Theta_D} V(D_{\tilde{\theta_d}}, G_{\theta_g}) - \lambda ||D_{\tilde{\theta_d}} - D_{\theta_d^*}||^2 \end{aligned} \tag{10}$$

As the extreme cases $\lambda \to \infty(0)$ recreate a Nash (Stackelberg) equilibrium, the $\lambda-$proximal equilibrium explores the spectrum of equilibria between the two and thus serves as a generic notion for GAN convergence. (Farnia & Ozdaglar, 2020) also suggest proximal training to explicitly enforce convergence to a proximal equilibrium. Our work, in contrast, is aimed towards evaluating GAN convergence and quantifying its goodness, irrespective of how it was trained.

4. Proximal Duality Gap

We first define the classical duality gap (Grnarova et al., 2019) for GANs before moving on to the proposed measure.

Definition 1. *Consider the GAN game presented in Eq.1. Then, for a configuration (θ_d, θ_g) of the game, the duality gap (DG) is defined as :*

$$DG(\theta_d, \theta_g) = \max_{\tilde{\theta_d} \in \Theta_D} V(D_{\tilde{\theta_d}}, G_{\theta_g}) - \min_{\tilde{\theta_g} \in \Theta_G} V(D_{\theta_d}, G_{\tilde{\theta_g}})$$

At a pure Nash equilibrium (θ_d^*, θ_g^*), $DG(\theta_d^*, \theta_g^*) = 0$. DG has some interesting properties. First, it is lower bounded by the JSD between P_r and P_{θ_g} (for $V = V_c$). Second, as DG is applicable to any GAN objective $V(D, G)$ and does not require pre-trained classifier nor labeled data, it is domain agnostic and potentially better equipped to monitor GAN training over other prevalent evaluation measures. However, as discussed previously, GANs can converge to non Nash attractors, where P_r and P_{θ_g} are aligned well. DG at such is not very well understood, limiting its practicality for monitoring GAN training.

We extend the notion of Duality Gap to the general context of training GANs where Nash equilibria need not be attainable, utilizing the proximal operator. We define the Proximal Duality Gap (DG^λ) as below.

Definition 2. *The proximal duality gap (DG^λ) at (θ_d, θ_g) for the GAN game presented in Eq.1 is defined as*

$$\begin{aligned} DG^\lambda(\theta_d, \theta_g) &= V_{D_w}(\theta_g) - V^\lambda_{G_w}(\theta_d), \text{ where} \\ V_{D_w}(\theta_g) &= \max_{\theta'_d \in \Theta_D} V(D_{\theta'_d}, G_{\theta_g}) \\ V^\lambda_{G_w}(\theta_d) &= \min_{\theta'_g \in \Theta_G} V^\lambda(D_{\theta_d}, G_{\theta'_g}) \end{aligned}$$

The terms D_w(or G_w) indicate the worst adversary that the generator (or discriminator) might face. Note that $\max_{\theta'_d \in \Theta_D} V^\lambda(D_{\theta'_d}, G_{\theta_g}) = \max_{\theta'_d \in \Theta_D} V(D_{\theta'_d}, G_{\theta_g})$. Thus, for a GAN configuration attained using V, the DG^λ measures the ability of the agents to deviate from it w.r.t the proximal objective V^λ.

Remark. *For all proximal equilibria (θ_d^*, θ_g^*) of the GAN game defined by Eq. 1, $V_{D_w}(\theta_g^*) = V^\lambda_{G_w}(\theta_d^*) = V^\lambda(\theta_d^*, \theta_g^*)$, thus $DG^\lambda(\theta_d^*, \theta_g^*) = 0$.*

This remark directly follows from the definition of proximal equilibria (Eq. 10). Thus DG^λ tending to zero implies that the GAN game has converged to a proximal equilibrium. However, as GANs are used for learning data distributions, it is important that measures to quantify GAN convergence should also give insights into the nature of the learned data distribution. Thus, to establish the applicability of DG^λ, we study how it relates to the divergence between the real and generated data distributions for various GAN formulations.

4.1. Theoretical Analysis

The definition of DG^λ has two terms - V_{D_w} and $V^\lambda_{G_w}$. We first establish the relationship between V_{D_w} and the divergences (DIV) used in various GAN formulations - the JS divergence ($JSD(P_{\theta_g}||P_r)$) for classical GAN objective V_c, the Wasserstein distance ($W_c(P_{\theta_g}||P_r)$) for the WGAN objective V_w, and the $f-$divergence ($D_f(P_{\theta_g}||P_r)$) for the F-GAN objective V_f. In our analysis, following (Farnia & Ozdaglar, 2020; Grnarova et al., 2019; Arjovsky et al., 2017), we assume that for a fixed generator, an optimal discriminator (D_w) that maximizes V exists.

Lemma 1. *Given a generator θ_g, V_{D_w} is related to the divergences between P_r and P_{θ_g} in the various GAN objectives as follows*

$$V_{D_w}(\theta_g) = \begin{cases} JSD(P_{\theta_g}||P_r) - \log 2, & \text{if } V = V_c \\ W_c(P_{\theta_g}||P_r), & \text{if } V = V_w \\ D_f(P_{\theta_g}||P_r), & \text{if } V = V_f \end{cases}$$

Proof. Deferred to the supplementary material. □

Thus, $V_{D_w}(\theta_g)$ measures the quality of the generator G_{θ_g}. If θ_g is optimal such that P_{θ_g} covers the real distribution P_r, then $V_{D_w}(\theta_g)$ achieves the minimum value. In case of a mismatch between P_{θ_g} and P_r, either due to insufficient support or poor sample quality, $V_{D_w}(\theta_g)$ will increase, thus making it a potentially useful metric in itself to monitor GAN training. However, V_{D_w} does not incorporate the ability of the discriminator to deviate from the current game configuration. Thus, it cannot identify if the game has converged to an equilibrium. Further, as the minimum value for V_{D_w} will vary depending upon the GAN formulation, it cannot serve as a domain agnostic measure for monitoring GAN training. DG^λ on the other hand will always tend to zero on attaining a proximal equilibrium irrespective of the GAN formulation. Thus, in the next result we analyze the behavior of DG^λ for the three GAN formulations. Specifically, we show that both in the realizable setting (G is of unbounded capacity; $\exists\, \theta_g$ such that $P_{\theta_g} = P_r$) and the more practical non-realizable setting (where G is of bounded capacity), the proximal duality gap is positive and lower bounded closely by the divergence between P_r and P_{θ_g}.

Theorem 1. *Consider a GAN game governed by an objective function V. Then the proximal duality gap (DG^λ) at a configuration (θ_d, θ_g) is related to the divergence between the real (P_r) and generated (P_{θ_g}) data distributions as follows.*

$$DG^\lambda(\theta_d, \theta_g) \geq DIV(P_{\theta_g}||P_r) - \kappa$$

where,

$$DIV(P_{\theta_g}||P_r) = \begin{cases} JSD(P_{\theta_g}||P_r), & \text{if } V = V_c \\ W_c(P_{\theta_g}||P_r), & \text{if } V = V_w \\ D_f(P_{\theta_g}||P_r), & \text{if } V = V_f \end{cases}$$

and κ (≥ 0) denotes the minimum divergence that the considered class of generator functions can achieve with the real data distribution.

Proof. Deferred to the supplementary material. □

Corollary. *Under the realizable setting, since $\exists\, \theta_g$ such that $P_{\theta_g} = P_r$, $\kappa = 0$ and hence $DG^{\lambda}(\theta_d, \theta_g) \geq DIV(P_{\theta_g}||P_r)$.*

This theorem non-trivially extends the prior result on the bound of DG only for classic GAN (Grnarova et al., 2019) under the realizable setting.

4.2. Implications

As DG^{λ} is lower bounded by the divergence between the real and generated data distributions, $DG^{\lambda} \to 0$ not only implies that the GAN has reached an equilibrium, but also that generated distribution is close to the real data distribution.

Corollary. *For GAN formulations defined by V_c, V_w or V_f, the generator attains the minimum possible divergence with the real data distribution at a proximal equilibrium (θ_d^*, θ_g^*), as $DG^{\lambda}(\theta_d^*, \theta_g^*) = 0$.*

This furthers the adeptness of proximal equilibria to serve as a general optimality notion for GANs. As a proximal equilibrium need not be a Nash equilibrium, it also facilitates the following interesting observation:

Remark. *GANs can capture the real data distribution even at non Nash game configurations.*

Thus, the empirical observation (Berard et al., 2020) that GANs can produce realistic data samples having high fidelity despite converging to a non-Nash attractor of the gradient dynamics is theoretically justified.

On similar lines, a natural question that arises concerning the behaviour of DG^{λ} is whether proximal equilibria constitute an exhaustive notion of equilibria at which GANs can capture the real data distribution. Precisely, we ask the question: Does GAN converging to a solution such that $P_{\theta_g} \to P_r \implies DG^{\lambda} \to 0$? Our answer begins with the following proposition concerning the extreme case for DG^{λ} as $\lambda \to 0$.

Theorem 2. *The proximal duality gap (DG^{λ}) at a configuration (θ_d^*, θ_g^*) for the GAN game defined by V_c, V_w, or V_f is equal to zero for $\lambda = 0$, when the generator learns the real data distribution; $P_{\theta_g^*} = P_r \implies DG^{\lambda=0}(\theta_d^*, \theta_g^*) = 0$.*

Proof. Deferred to the supplementary material. □

Corollary. *For the GAN formulations defined by V_c, V_w or V_f, the generator learns the real data distribution at a configuration (θ_d^*, θ_g^*) if and only if (θ_d^*, θ_g^*) constitutes a Stackelberg equilibrium.*

Proof. Deferred to the supplementary material. □

Thus $DG^{\lambda=0}(\theta_d, \theta_g) \to 0$ whenever $P_{\theta_g} \to P_r$. The value of λ restricts the discriminator in the proximal objective (V^{λ}) to be optimal within a neighbourhood. As $\lambda \to 0$, V^{λ} considers the optimal discriminator over the entire parameter space (Θ_D). Thus $DG^{\lambda=0}(\theta_d, \theta_g) = 0$ implies that (θ_d, θ_g) is a Stackelberg equilibrium. As all proximal equilibria form a subset of Stackelberg equilibria, $DG^{\lambda=0}$ would thus be an ideal choice to monitor GAN convergence. However, as λ decreases, the complexity of computing V^{λ} increases rapidly and becomes infeasible as $\lambda \to 0$. Hence, it is only practical to check if a GAN configuration is a $\lambda(>0)-$proximal equilibrium. But can DG^{λ}, for a fixed value of $\lambda(>0)$ monitor convergence of GANs to all $\lambda^{'}-$proximal equilibria? To address this question, let us study the two cases - (i) $\lambda^{'} \geq \lambda$ and (ii) $\lambda^{'} < \lambda$ separately. The following theorem addresses case (i) utilizing the hierarchical property of proximal equilibria.

Theorem 3. *Consider a GAN configuration (θ_d, θ_g). Then, $\forall \lambda^{'} \geq \lambda_0$,*

$$DG^{\lambda=\lambda^{'}}(\theta_d, \theta_g) = 0 \implies DG^{\lambda=\lambda_0}(\theta_d, \theta_g) = 0$$

Proof. Deferred to Supplementary material □

$DG^{\lambda^{'}}(\theta_d, \theta_g) = 0$ is a sufficient condition for (θ_d, θ_g) being a $\lambda^{'}-$proximal equilibrium. Thus, it follows from theorem 3 that DG^{λ} is adept to monitor convergence of GANs to all $\lambda^{'}(\geq \lambda)-$proximal equilibria. However, when a GAN converges to a $\lambda^{'}(<\lambda)-$proximal equilibrium, DG^{λ} can be prone to error. The following theorem addresses this issue by upper bounding the difference between DG^{λ} and the divergence between real and generated data distributions.

Theorem 4. *Consider a GAN game governed by an objective function V. For $\lambda > 0$, let V^{λ} denote the proximal objective defined by $V^{\lambda}(\theta_d, \theta_g) = max_{\tilde{\theta_d}} V(\tilde{\theta}_d, \theta_g) - \lambda||D_{\tilde{\theta_d}} - D_{\theta_d}||^2$. Then, $\forall\, \epsilon > 0$, $\exists\, \delta > 0$ such that if $||D_{\theta_d} - D_{\tilde{\theta_d}}|| < \delta$, then $DG^{\lambda}(\theta_d, \theta_g) - DIV(P_{\theta_g}||P_r) < \epsilon$ where,*

$$DIV(P_{\theta_g}||P_r) = \begin{cases} JSD(P_{\theta_g}||P_r), & \textit{if } V = V_c \\ W_c(P_{\theta_g}||P_r), & \textit{if } V = V_w \\ D_f(P_{\theta_g}||P_r), & \textit{if } V = V_f \end{cases}$$

Proof. Deferred to Supplementary material □

Corollary. *For a GAN configuration (θ_d^*, θ_g^*) such that $P_{\theta_g^*} = P_r, DG^{\lambda}(\theta_d^*, \theta_g^*) < \epsilon$*

Thus even when the GAN converges to a $\lambda^{'}(<\lambda)-$proximal equilibrium, the error that DG^{λ} can incur is bounded. Previously, the absence of such an upper bound as implied by theorem 4 for DG meant that DG need not necessarily be close to zero when P_{θ_g} is close to P_r. DG^{λ} however, rules out this possibility and is thus a theoretically grounded and robust measure that can serve as a tool for monitoring convergence of GANs in the wild.

Figure 3. Monitoring Convergence of WGAN (Top Row) and SNGAN (Bottom Row) over the datasets MNIST (Col 1), CIFAR-10 (Col 2) and CELEB-A (Col 3) using duality gap. DG is not reflective of the GAN convergence. DG^λ, on the other hand is indicative of convergence and saturates close to zero. The shaded region indicates the standard deviation over 5 independent trials.

4.3. Estimating Proximal Duality Gap

Computing the true DG^λ for a GAN configuration is a hard task as it involves finding the optima of non convex functions. We approximate DG^λ by employing gradient descent to estimate V_{D_w} and $V^\lambda_{G_w}$. Following (Farnia & Ozdaglar, 2020), we use the Sobolev norm in the proximal objective V^λ, given by $||D|| = \sqrt{\mathbb{E}_{x\sim P_r}[||\nabla_x D(x)||_2^2]}$. For a GAN game governed by V, estimating V_{D_w} involves optimizing $V(\theta_d, \theta_g)$ w.r.t θ_d using gradient descent. However, estimating $V^\lambda_{G_w}$ requires gradient computation over the proximal operator. As shown in (Farnia & Ozdaglar, 2020), for a GAN objective function V that is smooth w.r.t θ_d, the gradient of the proximal objective (V^λ) w.r.t θ_g can be obtained in terms of V as : $\nabla_{\theta_g} V^\lambda(D_{\theta_d}, G_{\theta_g}) = \nabla_{\theta_g} V(D_{\theta^*_d}, G_{\theta_g})$, where θ^*_d represents the optimal discriminator implied by the proximal objective. Thus, to estimate $V^\lambda_{G_w}$, at every iteration we use gradient descent to obtain θ^*_d for the corresponding θ_g and update θ_g to minimize $V(\theta^*_d, \theta_g)$. The algorithm for the overall estimation process and the associated computational complexity are discussed in the supplementary material. To ensure that we obtain an unbiased estimate for DG^λ, following (Grnarova et al., 2019), we split the dataset into 3 disjoint sets - S_A, S_B and S_C. We train the GAN using S_A, we use S_B to find the worst case counter parts D_w and G_w via gradient descent, and S_C to evaluate the objective function at the obtained worst case configurations.

5. Experimentation

To experimentally establish the proficiency of DG^λ, we consider a WGAN with weight-clipping (that optimizes V_w) (Arjovsky et al., 2017) and a Spectral Normalized GAN (SNGAN) (that optimizes V_c) (Miyato et al., 2018) over 3 datasets - MNIST (Deng, 2012), CIFAR-10 (Krizhevsky et al., 2014) and CELEB-A (Liu et al., 2015). For all the experiments, we use the 4-layer DCGAN (Radford et al., 2016) architecture for both the generator and the discriminator networks, and an Adam optimizer (Kingma & Ba, 2015) to train the models. To compute DG^λ, we use λ=0.1 and 20 optimization steps for approximating the proximal objective. We used the torchgan framework (Pal & Das, 2019) to train and evaluate all GAN models. Further implementation details for each experiment are provided in the supplementary material and the source code is publicly available [1].

[1] https://github.com/proximal-dg/proximal_dg

	Pearson Correlation Coefficient (r)			
	$\mathrm{r}_{DG,IS}$	$\mathrm{r}_{DG^\lambda,IS}$	$\mathrm{r}_{DG,FID}$	$\mathrm{r}_{DG^\lambda,FID}$
MNIST	-0.752	$\mathbf{-0.892}$	0.845	$\mathbf{0.957}$
CIFAR-10	-0.368	$\mathbf{-0.854}$	0.213	$\mathbf{0.738}$
CELEB-A	-0.213	$\mathbf{-0.524}$	0.263	$\mathbf{0.699}$

Table 1. Comparing the correlation of DG and DG^λ with IS and FID computed during the training of WGAN over the 3 datasets.

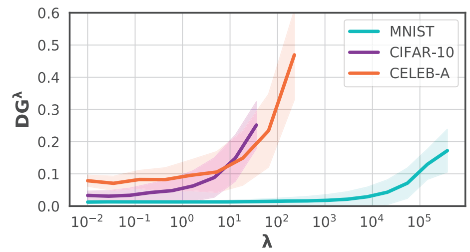

Figure 4. The behaviour of DG^λ for increasing λ

Monitoring GAN training using DG$^\lambda$

Our first experiment aims to establish that DG^λ is better equipped over DG to monitor GAN convergence in practice. To this end, we train a WGAN and SNGAN over the 3 datasets till the models have converged. We compute DG and DG^λ throughout the training process, in addition to the (image) domain specific evaluation measures - IS and FID. Figure 3 demonstrates the training progress of each GAN, qualitatively through visualization of samples from the learned data distribution and quantitatively in terms of DG and DG^λ. The high fidelity of the learned data samples indicates that the models have converged. However, DG is not reflective of the training progress. This suggests that the GANs have not attained a Nash equilibrium - the behaviour of DG at non-Nash critical points is not well understood. DG^λ, on the other hand, captures the trend in the training progress, and eventually saturates close to zero. Thus, it is able to better characterize convergence. We also quantitatively validate the above observation by examining the correlation between DG and DG^λ against popular measures - IS and FID that quantify the quality of P_{θ_g}. As shown in Table 1, DG^λ has a higher positive correlation with FID and a higher negative correlation with IS as compared to DG. The duality gap is negatively correlated with IS because the latter increases as the GAN learns the real data distribution, whereas the former decreases. A larger (or smaller) IS (or FID) implies better fidelity of the learned data distribution. The higher correlation of DG^λ with IS and FID over the training process thus validates that DG^λ is adept to monitor not only the convergence of GANs to an equilibrium but also the goodness of P_{θ_g}.

Visualizing the effect of λ

λ is a critical hyperparameter that determines the proficiency of DG^λ. We observed from the theoretical analysis that, while DG^λ is adept to monitor convergence of GANs to all $\lambda^{'}(\geq \lambda)$-proximal equilibria, it is prone to error as λ increases. As $\lambda \to \infty$, DG^λ becomes equivalent to DG. We thus experimentally study the behaviour of DG^λ for increasing values of λ. We compute DG^λ at the converged WGAN configurations (as shown in Fig 3) for each of the three datasets by varying λ in the range $[10^{-2}, 10^6]$. We observe (Figure 4) that for all the datasets, DG^λ remains close to zero and unaffected for λ in range $[10^{-2}, 1]$. Interestingly, for the MNIST dataset, DG^λ remains unaffected even for larger values of $\lambda (\approx 10^4)$. This suggests that the WGAN configuration for MNIST is closer to a Nash equilibrium, also explaining why DG and DG^λ are closer for the same in Figure 3. As λ crosses a threshold (10^1 for CELEB-A, CIFAR-10 and 10^4 for MNIST), DG^λ increases sharply and behaves similar to DG. Thus, for a small value for λ (< 1) DG^λ is a robust tool for monitoring GAN convergence.

Influencing GAN training using DG$^\lambda$

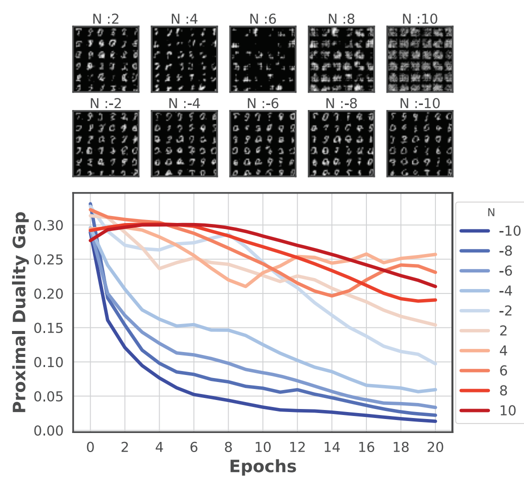

Figure 5. Tuning GAN Hyperparameter using DG^λ

A quantitative measure to monitor GAN training would enable easier tuning of hyperparameters. In this section, we explore DG^λ as an effective tool for influencing GAN training. A decisive hyperparameter that governs the delicate balance of the GAN game and hence its convergence is the update ratio of the agents. Let us denote by N, the number of discriminator updates per generator update, where a negative value for N indicates a larger number of generator updates. We train WGAN over the MNIST dataset by performing a grid search over N in the range -10 to 10 and computing DG^λ. Figure 5 depicts the qualitative output at the end of 20 epochs and DG^{λ} across training for each value of N. We observe that as N increases, the quality of the generated data samples diminishes and the learned data distribution eventually diverges as $N \to 10$. Correspondingly, we observe that DG^{λ} is close to zero for lower values of N(colored blue) and increases with N(colored red), suggesting that the GAN diverges for larger values of N. DG^{λ} thus enables us to quantitatively identify the optimal range of values for the hyperparameters of a GAN.

6. Summary and Future Work

GANs have pushed the boundaries of learning complex data distributions. However, the non-intuitive nature of GAN loss curves makes training a challenging task. We propose Proximal Duality Gap (DG^{λ}) as a generic and quantitative tool to monitor GAN training and understand GAN convergence. DG^{λ} characterizes GAN convergence as the game attaining a $\lambda-$proximal equilibrium. It also helps derive insights into the nature of GAN convergence - a GAN learns the real data distribution if and only if it attains a Stackelberg equilibrium. The ability of DG^{λ} to objectively quantify GAN convergence makes it a useful measure to tune the hyperparameters of a GAN. A couple of open questions that can improve the utility of DG^{λ}, if addressed, include identifying an optimal λ and making the range of values for DG^{λ} invariant across GAN formulations. The characterization of GAN convergence through proximal duality gap opens up new avenues for effortless GAN training.

Acknowledgements

The resources provided by the *PARAM Shivay* Facility under the National Super-computing Mission, Government of India at the Indian Institute of Technology, Varanasi and under Google Tensorflow Research award are gratefully acknowledged.

References

Adolphs, L., Daneshmand, H., Lucchi, A., and Hofmann, T. Local saddle point optimization: A curvature exploitation approach. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pp. 486–495, 2019.

Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 214–223, 2017.

Arora, S., Ge, R., Liang, Y., Ma, T., and Zhang, Y. Generalization and equilibrium in generative adversarial nets (gans). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 224–232, 2017.

Berard, H., Gidel, G., Almahairi, A., Vincent, P., and Lacoste-Julien, S. A closer look at the optimization landscapes of generative adversarial networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.

Borji, A. Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding*, 179:41–65, 2019. ISSN 1077-3142.

Brock, A., Donahue, J., and Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.

Daskalakis, C., Ilyas, A., Syrgkanis, V., and Zeng, H. Training gans with optimism. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.

Deng, L. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 2012.

Farnia, F. and Ozdaglar, A. E. Do gans always have nash equilibria? In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3029–3039, 2020.

Farnia, F. and Tse, D. A convex duality framework for gans. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 5254–5263, 2018.

Fiez, T., Chasnov, B., and Ratliff, L. J. Convergence of learning dynamics in stackelberg games. *Computing Research Repository (CoRR)*, 2019.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. C., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 2672–2680, 2014.

Grnarova, P., Levy, K. Y., Lucchi, A., Perraudin, N., Goodfellow, I., Hofmann, T., and Krause, A. A domain agnostic measure for monitoring and evaluating gans. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 12069–12079, 2019.

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5767–5777, 2017.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6626–6637, 2017.

Hsieh, Y., Liu, C., and Cevher, V. Finding mixed nash equilibria of generative adversarial networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2810–2819, 2019.

Jin, C., Netrapalli, P., and Jordan, M. I. What is local optimality in nonconvex-nonconcave minimax optimization? In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 4880–4889, 2020.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

Krizhevsky, A., Nair, V., and Hinton, G. Cifar-10 (canadian institute for advanced research). 2014.

Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., and Aila, T. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 3929–3938, 2019.

Lin, T., Jin, C., and Jordan, M. I. On gradient descent ascent for nonconvex-concave minimax problems. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6083–6093, 2020.

Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 3730–3738, 2015.

Lucic, M., Kurach, K., Michalski, M., Gelly, S., and Bousquet, O. Are gans created equal? A large-scale study. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 698–707, 2018.

Mazumdar, E. and Ratliff, L. J. On the convergence of competitive, multi-agent gradient-based learning. *Computing Research Repository (CoRR)*, 2018.

Mazumdar, E. V., Jordan, M. I., and Sastry, S. S. On finding local nash equilibria (and only local nash equilibria) in zero-sum games. *Computing Research Repository (CoRR)*, 2019.

Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.

Mokhtari, A., Ozdaglar, A. E., and Pattathil, S. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1497–1507, 2020.

Naeem, M. F., Oh, S. J., Uh, Y., Choi, Y., and Yoo, J. Reliable fidelity and diversity metrics for generative models. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7176–7185, 2020.

Nash, J. F. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 1950.

Nguyen, X., Wainwright, M. J., and Jordan, M. I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 2010.

Nouiehed, M., Sanjabi, M., Huang, T., Lee, J. D., and Razaviyayn, M. Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 14905–14916, 2019.

Nowozin, S., Cseke, B., and Tomioka, R. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 271–279, 2016.

Olsson, C., Bhupatiraju, S., Brown, T. B., Odena, A., and Goodfellow, I. J. Skill rating for generative models. *Computing Research Repository (CoRR)*, 2018.

Pal, A. and Das, A. Torchgan: A flexible framework for gan training and evaluation. *Computing Research Repository (CoRR)*, 2019.

Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

Sajjadi, M. S. M., Bachem, O., Lucic, M., Bousquet, O., and Gelly, S. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 5234–5243, 2018.

Salimans, T., Goodfellow, I. J., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 2226–2234, 2016.

Sidheekh, S., Aimen, A., Madan, V., and Krishnan, N. C. On duality gap as a measure for monitoring gan training. *Computing Research Repository (CoRR)*, 2020.

Theis, L., van den Oord, A., and Bethge, M. A note on the evaluation of generative models. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

Tolstikhin, I. O., Gelly, S., Bousquet, O., Simon-Gabriel, C., and Schölkopf, B. Adagan: Boosting generative models. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5424–5433, 2017.

Zhang, H., Goodfellow, I. J., Metaxas, D. N., and Odena, A. Self-attention generative adversarial networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7354–7363, 2019a.

Zhang, K., Yang, Z., and Basar, T. Policy optimization provably converges to nash equilibria in zero-sum linear quadratic games. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 11598–11610, 2019b.