
Lipschitz Generative Adversarial Nets

Zhiming Zhou¹ Jiadong Liang² Yuxuan Song¹ Lantao Yu³ Hongwei Wang³ Weinan Zhang¹ Yong Yu¹
Zhihua Zhang²

Abstract

In this paper we show that generative adversarial networks (GANs) without restriction on the discriminative function space commonly suffer from the problem that the gradient produced by the discriminator is uninformative to guide the generator. By contrast, Wasserstein GAN (WGAN), where the discriminative function is restricted to 1-Lipschitz, does not suffer from such a gradient uninformative problem. We further show in the paper that the model with a compact dual form of Wasserstein distance, where the Lipschitz condition is relaxed, may also theoretically suffer from this issue. This implies the importance of Lipschitz condition and motivates us to study the general formulation of GANs with Lipschitz constraint, which leads to a new family of GANs that we call Lipschitz GANs (LGANs). We show that LGANs guarantee the existence and uniqueness of the optimal discriminative function as well as the existence of a unique Nash equilibrium. We prove that LGANs are generally capable of eliminating the gradient uninformative problem. According to our empirical analysis, LGANs are more stable and generate consistently higher quality samples compared with WGAN.

1. Introduction

Generative adversarial networks (GANs) (Goodfellow et al., 2014), as one of the most successful generative models, have shown promising results in various challenging tasks. GANs are popular and widely used, but they are notoriously hard to train (Goodfellow, 2016). The underlying obstacles, though have been heavily studied (Arjovsky & Bottou, 2017; Lucic et al., 2017; Heusel et al., 2017; Mescheder et al., 2017;

¹Shanghai Jiao Tong University ²Peking University ³Stanford University. Correspondence to: Zhiming Zhou <heyohai@apex.sjtu.edu.cn>.

2018; Yadav et al., 2017), are still not fully understood.

The objective of GAN is usually defined as a distance metric between the real distribution \mathcal{P}_r and the generative distribution \mathcal{P}_g , which implies that $\mathcal{P}_r = \mathcal{P}_g$ is the unique global optimum. The nonconvergence of traditional GANs has been considered as a result of ill-behaving distance metric (Arjovsky & Bottou, 2017), i.e., the distance between \mathcal{P}_r and \mathcal{P}_g keeps constant when their supports are disjoint. Arjovsky et al. (2017) accordingly suggested using the Wasserstein distance, which can properly measure the distance between two distributions no matter whether their supports are disjoint.

In this paper, we conduct a further study on the convergence of GANs from the perspective of the informativeness of the gradient of the optimal discriminative function f^* . We show that for GANs that have no restriction on the discriminative function space, e.g., the vanilla GAN and its most variants, $f^*(x)$ is only related to the densities of the local point x and does not reflect any information about other points in the distributions. We demonstrate that under these circumstances, the gradient of the optimal discriminative function with respect to its input, on which the generator updates generated samples, usually tells nothing about the real distribution. We refer to this phenomenon as the *gradient uninformative*, which is substantially different from the gradient vanishing and is a fundamental cause of nonconvergence of GANs.

According to the analysis of Gulrajani et al. (2017), Wasserstein GAN can avoid the gradient uninformative problem. Meanwhile, we show in the paper that the Lipschitz constraint in the Kantorovich-Rubinstein dual of the Wasserstein distance can be relaxed, leading to a new equivalent dual; and with the new dual form, the gradient may also not reflect any information about how to refine \mathcal{P}_g towards \mathcal{P}_r . It suggests that Lipschitz condition would be a vital element for resolving the gradient uninformative problem.

Motivated by the above analysis, we investigate the general formulation of GANs with Lipschitz constraint. We show that under a mild condition, penalizing Lipschitz constant guarantees the existence and uniqueness of the optimal discriminative function as well as the existence of the unique Nash equilibrium between f^* and \mathcal{P}_g where $\mathcal{P}_r = \mathcal{P}_g$. It leads to a new family of GANs that we call Lipschitz GANs

Table 1: Comparison of different objectives in GANs.

	ϕ	φ	\mathcal{F}	$f^*(x)$	Gradient Vanishing	Gradient Uninformative	$f^*(x)$ Uniqueness
Vanilla GAN	$-\log(\sigma(-x))$	$-\log(\sigma(x))$	$\{f: \mathbb{R}^n \rightarrow \mathbb{R}\}$	$\log \frac{\mathcal{P}_r(x)}{\mathcal{P}_g(x)}$	Yes	Yes	Yes
Least-Squares GAN	$(x - \alpha)^2$	$(x - \beta)^2$	$\{f: \mathbb{R}^n \rightarrow \mathbb{R}\}$	$\frac{\alpha \cdot \mathcal{P}_g(x) + \beta \cdot \mathcal{P}_r(x)}{\mathcal{P}_r(x) + \mathcal{P}_g(x)}$	No	Yes	Yes
μ -Fisher GAN	x	$-x$	$\{f: \mathbb{R}^n \rightarrow \mathbb{R}, \mathbb{E}_{x \sim \mu} f(x) ^2 \leq 1\}$	$\frac{1}{\mathcal{F}_\mu(\mathcal{P}_r, \mathcal{P}_g)} \frac{\mathcal{P}_r(x) - \mathcal{P}_g(x)}{\mu(x)}$	No	Yes	Yes
Wasserstein GAN	x	$-x$	$\{f: \mathbb{R}^n \rightarrow \mathbb{R}, k(f) \leq 1\}$	N/A	No	No	No
Lipschitz GAN	any ϕ and φ satisfying Eq. (11)		$\{f: \mathbb{R}^n \rightarrow \mathbb{R}; k(f) \text{ is penalized}\}$	N/A	No	No	Yes

(LGANs). We show that LGANs are generally capable of eliminating the gradient uninformative in the manner that with the optimal discriminative function, the gradient for each generated sample, if nonzero, will point towards some real sample. This process continues until the Nash equilibrium $\mathcal{P}_r = \mathcal{P}_g$ is reached.

The remainder of this paper is organized as follows. In Section 2, we provide some preliminaries that will be used in this paper. In Section 3, we study the gradient uninformative issue in detail. In Section 4, we present LGANs and their theoretical analysis. We conduct the empirical analysis in Section 5. Finally, we discuss related work in Section 6 and conclude the paper in Section 7.

2. Preliminaries

In this section we first give some notions and then present a general formulation for generative adversarial networks.

2.1. Notation and Notions

Given two metric spaces (X, d_X) and (Y, d_Y) , a function $f: X \rightarrow Y$ is said to be Lipschitz continuous if there exists a constant $k \geq 0$ such that

$$d_Y(f(x_1), f(x_2)) \leq k \cdot d_X(x_1, x_2), \forall x_1, x_2 \in X. \quad (1)$$

In this paper and in most existing GANs, the metrics d_X and d_Y are by default Euclidean distance which we also denote by $\|\cdot\|$. The smallest constant k is called the (best) Lipschitz constant of f , denoted by $k(f)$.

The first-order Wasserstein distance W_1 between two probability distributions is defined as

$$W_1(\mathcal{P}_r, \mathcal{P}_g) = \inf_{\pi \in \Pi(\mathcal{P}_r, \mathcal{P}_g)} \mathbb{E}_{(x,y) \sim \pi} [d(x, y)], \quad (2)$$

where $\Pi(\mathcal{P}_r, \mathcal{P}_g)$ denotes the set of all probability measures with marginals \mathcal{P}_r and \mathcal{P}_g . It can be interpreted as the minimum cost of transporting the distribution \mathcal{P}_g to the distribution \mathcal{P}_r . We use π^* to denote the optimal transport plan, and let \mathcal{S}_r and \mathcal{S}_g denote the supports of \mathcal{P}_r and \mathcal{P}_g , respectively. We say two distributions are disjoint if their supports are disjoint.

The Kantorovich-Rubinstein (KR) duality (Villani, 2008) provides a way of more efficiently computing of Wasserstein

distance. The duality states that

$$W_1(\mathcal{P}_r, \mathcal{P}_g) = \sup_f \mathbb{E}_{x \sim \mathcal{P}_r} [f(x)] - \mathbb{E}_{x \sim \mathcal{P}_g} [f(x)], \quad (3)$$

s.t. $f(x) - f(y) \leq d(x, y), \forall x, y$.

The constraint in Eq. (3) implies that f is Lipschitz continuous with $k(f) \leq 1$. Interestingly, we have a more compact dual form of the Wasserstein distance. That is,

$$W_1(\mathcal{P}_r, \mathcal{P}_g) = \sup_f \mathbb{E}_{x \sim \mathcal{P}_r} [f(x)] - \mathbb{E}_{x \sim \mathcal{P}_g} [f(x)], \quad (4)$$

s.t. $f(x) - f(y) \leq d(x, y), \forall x \in \mathcal{S}_r, \forall y \in \mathcal{S}_g$.

The proof for this dual form is given in Appendix A.5. We see that this new dual relaxes the Lipschitz continuity condition of the dual form in Eq. (3).

2.2. Generative Adversarial Networks (GANs)

Typically, GANs can be formulated as

$$\min_{f \in \mathcal{F}} J_D \triangleq \mathbb{E}_{z \sim \mathcal{P}_z} [\phi(f(g(z)))] + \mathbb{E}_{x \sim \mathcal{P}_r} [\varphi(f(x))], \quad (5)$$

$$\min_{g \in \mathcal{G}} J_G \triangleq \mathbb{E}_{z \sim \mathcal{P}_z} [\psi(f(g(z)))],$$

where \mathcal{P}_z is the source distribution of the generator in \mathbb{R}^m and \mathcal{P}_r is the target (real) distribution in \mathbb{R}^n . The generative function $g: \mathbb{R}^m \rightarrow \mathbb{R}^n$ learns to output samples that share the same dimension as samples in \mathcal{P}_r , while the discriminative function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ learns to output a score indicating the authenticity of a given sample. Here \mathcal{F} and \mathcal{G} denote discriminative and generative function spaces, respectively; and $\phi, \varphi, \psi: \mathbb{R} \rightarrow \mathbb{R}$ are loss metrics. We denote the implicit distribution of the generated samples by \mathcal{P}_g .

We list the choices of \mathcal{F} , ϕ and φ in some representative GAN models in Table 1. In these GANs, the gradient that the generator receives from the discriminator with respect to (w.r.t.) a generated sample $x \in \mathcal{S}_g$ is

$$\nabla_x J_G(x) \triangleq \nabla_x \psi(f(x)) = \nabla_{f(x)} \psi(f(x)) \cdot \nabla_x f(x), \quad (6)$$

where the first term $\nabla_{f(x)} \psi(f(x))$ is a step-related scalar, and the second term $\nabla_x f(x)$ is a vector with the same dimension as x which indicates the direction that the generator should follow for optimizing the generated sample x .

We use f^* to denote the optimal discriminative function, i.e., $f^* \triangleq \arg \min_{f \in \mathcal{F}} J_D$. For further notation, we let $\dot{J}_D(x) \triangleq \mathcal{P}_g(x) \phi(f(x)) + \mathcal{P}_r(x) \varphi(f(x))$. It has $J_D = \int \dot{J}_D(x) dx$.

2.3. The Gradient Vanishing

The gradient vanishing problem has been typically thought as a key factor for causing the nonconvergence of GANs, i.e., the gradient becomes zero when the discriminator is perfectly trained.

Goodfellow et al. (2014) addressed this problem by using an alternative objective for the generator. Actually, only the scalar $\nabla_{f(x)}\psi(f(x))$ is changed. The Least-Squares GAN (Mao et al., 2016), which aims at addressing the gradient vanishing problem, also focused on $\nabla_{f(x)}\psi(f(x))$.

Arjovsky & Bottou (2017) provided a new perspective for understanding the gradient vanishing. They argued that \mathcal{S}_r and \mathcal{S}_g are usually disjoint and the gradient vanishing stems from the ill-behaving of traditional distance metrics, i.e., the distance between \mathcal{P}_r and \mathcal{P}_g remains constant when they are disjoint. The Wasserstein distance was thus used (Arjovsky et al., 2017) as an alternative metric, which can properly measure the distance between two distributions no matter whether they are disjoint.

3. The Gradient Uninformativeness

In this paper we pay our main attention on the gradient direction of the optimal discriminative function, i.e., $\nabla_x f^*(x)$, along which the generated sample x is updated. We show that for many distance metrics, such a gradient may fail to bring any useful information about \mathcal{P}_r . Consequently, \mathcal{P}_g is not guaranteed to converge to \mathcal{P}_r . We name this phenomenon as the *gradient uninformativeness* and argue that it is a fundamental factor of resulting in nonconvergence and instability in the training of traditional GANs.

The gradient uninformativeness is substantially different from the gradient vanishing. The gradient vanishing is about the scalar term $\nabla_{f(x)}\psi(f(x))$ in $\nabla_x J_G(x)$ or the overall scale of $\nabla_x J_G(x)$, while the gradient uninformativeness is about the direction of $\nabla_x J_G(x)$, which is defined by $\nabla_x f^*(x)$. The two issues are orthogonal, though they sometimes exist simultaneously. See Table 1 for a summary of issues for representative GANs.

Next, we discuss the gradient uninformativeness in the taxonomy of restrictions on the discriminative function space \mathcal{F} . We will show that for unrestricted GANs, gradient uninformativeness commonly exists; for restricted GANs, such an issue might still exist; and with Lipschitz condition, it generally does not exist.

3.1. Unrestricted GANs

For many GAN models, there is no restriction on \mathcal{F} . Typical cases include f -divergence based GANs, such as the vanilla GAN (Goodfellow et al., 2014), Least-Squares GAN (Mao et al., 2016) and f -GAN (Nowozin et al., 2016).

In these GANs, the value of the optimal discriminative function at each point $f^*(x)$ is independent of other points and only reflects the local densities $\mathcal{P}_r(x)$ and $\mathcal{P}_g(x)$:

$$f^*(x) = \arg \min_{f(x) \in \mathbb{R}} \mathcal{P}_g(x)\phi(f(x)) + \mathcal{P}_r(x)\varphi(f(x)), \quad \forall x.$$

Hence, for each generated sample x which is not surrounded by real samples (there exists $\epsilon > 0$ such that for all y with $0 < \|y - x\| < \epsilon$, it holds that $y \notin \mathcal{S}_r$), $f^*(x)$ in the surrounding of x would contain no information about \mathcal{P}_r . Thus $\nabla_x f^*(x)$, the gradient that x receives from the optimal discriminative function, does not reflect any information about \mathcal{P}_r .

Typical situation is that \mathcal{S}_r and \mathcal{S}_g are disjoint, which is common in practice according to (Arjovsky & Bottou, 2017). To further distinguish the gradient uninformativeness from the gradient vanishing, we consider an ideal case: \mathcal{S}_r and \mathcal{S}_g are totally overlapped and both consist of n discrete points, but their probability masses over these points are different. In this case, $\nabla_x f^*(x)$ for each generated sample is still uninformative, but the gradient does not vanish.

3.2. Restricted GANs: Fisher GAN as an Instance

Some GANs impose restrictions on \mathcal{F} . Typical instances are the Integral Probability Metric (IPM) based GANs (Mroueh & Sercu, 2017; Mroueh et al., 2017; Bellemare et al., 2017) and the Wasserstein GAN (Arjovsky et al., 2017). We next show that GANs with restriction on \mathcal{F} might also suffer from the gradient uninformativeness.

The optimal discriminative function of μ -Fisher IPM $\mathcal{F}_\mu(\mathcal{P}_r, \mathcal{P}_g)$, the generalized objective of the Fisher GAN (Mroueh et al., 2017), has the following form:

$$f^*(x) = \frac{1}{\mathcal{F}_\mu(\mathcal{P}_r, \mathcal{P}_g)} \frac{\mathcal{P}_r(x) - \mathcal{P}_g(x)}{\mu(x)}, \quad (7)$$

where μ is a distribution whose support covers \mathcal{S}_r and \mathcal{S}_g , and $\frac{1}{\mathcal{F}_\mu(\mathcal{P}_r, \mathcal{P}_g)}$ is a constant. It can be observed that μ -Fisher IPM also defines $f^*(x)$ at each point according to the local densities and does not reflect information of other locations. Similar as above, we can conclude that for each generated sample that is not surrounded by real samples, $\nabla_x f^*(x)$ is uninformative.

3.3. The Wasserstein GAN

As shown by Gulrajani et al. (2017), the gradient of the optimal discriminative function in the KR dual form of the Wasserstein distance has the following property:

Proposition 1. *Let π^* be the optimal transport plan in Eq. (2) and $x_t = tx + (1-t)y$ with $0 \leq t \leq 1$. If the optimal discriminative function f^* in Eq. (3) is differentiable and $\pi^*(x, x) = 0$ for all x , then it holds that*

$$\mathbb{P}_{(x,y) \sim \pi^*} \left[\nabla_{x_t} f^*(x_t) = \frac{y-x}{\|y-x\|} \right] = 1. \quad (8)$$

This proposition indicates: (i) for each generated sample x , there exists a real sample y such that $\nabla_{x_t} f^*(x_t) = \frac{y-x}{\|y-x\|}$ for all linear interpolations x_t between x and y , i.e., the gradient at any x_t is pointing towards the real sample y ; (ii) these (x, y) pairs match the optimal coupling π^* in the optimal transport perspective. It implies that WGAN is able to overcome the gradient uninformative as well as the gradient vanishing.

Our concern turns to the reason why WGAN can avoid gradient uninformative. To address this question, we alternatively apply the compact dual of the Wasserstein distance in Eq. (4) and study the optimal discriminative function.

Since there is generally no closed-form solution for f^* in Eq. (4), we take an illustrative example, but the conclusion is general. Let $Z \sim U[0, 1]$ be a uniform variable on interval $[0, 1]$, \mathcal{P}_r be the distribution of $(1, Z)$ in \mathbb{R}^2 , and \mathcal{P}_g be the distribution of $(0, Z)$ in \mathbb{R}^2 . According to Eq. (4), we have an optimal f^* as follows

$$f^*(x) = \begin{cases} 1, & \forall x \in \mathcal{S}_r; \\ 0, & \forall x \in \mathcal{S}_g. \end{cases} \quad (9)$$

Though having the constraint “ $f(x) - f(y) \leq d(x, y)$, $\forall x \in \mathcal{S}_r, \forall y \in \mathcal{S}_g$,” the Wasserstein distance in this dual form also only defines the values of $f^*(x)$ on \mathcal{S}_r and \mathcal{S}_g . For each generated sample x which is isolated or at the boundary (there does not exist $\epsilon > 0$ such that it holds $y \in \mathcal{S}_r \cup \mathcal{S}_g$ for all y with $0 < \|y - x\| < \epsilon$), the gradient of $f^*(x)$ is theoretically undefined and thus cannot provide useful information about \mathcal{P}_r . We can consider the more extreme case where \mathcal{S}_g are isolated points to make it clearer.

These examples imply that Lipschitz condition would be critical for resolving the gradient uninformative problem. Motivated by this, we study the general formulation of GANs with Lipschitz constraint, which leads to a family of more general GANs that we call Lipschitz GANs. We will see that in Lipschitz GANs, the similarity measure between \mathcal{P}_r and \mathcal{P}_g might not be some Wasserstein distance, but they still perform very well.

4. Lipschitz GANs

Lipschitz continuity recently becomes popular in GANs. It was observed that introducing Lipschitz continuity as a regularization of the discriminator leads to improved stability and sample quality (Arjovsky et al., 2017; Kodali et al., 2017; Fedus et al., 2017; Miyato et al., 2018; Qi, 2017).

In this paper, we investigate the general formulation of GANs with Lipschitz constraint, where the Lipschitz constant of discriminative function is penalized via a quadratic loss, to theoretically analyze the properties of such GANs. In particular, we define the Lipschitz Generative Adversarial

Nets (LGANs) as:

$$\begin{aligned} & \min_{f \in \mathcal{F}} \mathbb{E}_{z \sim \mathcal{P}_z} [\phi(f(g(z)))] + \mathbb{E}_{x \sim \mathcal{P}_r} [\varphi(f(x))] + \lambda \cdot k(f)^2, \\ & \min_{g \in \mathcal{G}} \mathbb{E}_{z \sim \mathcal{P}_z} [\psi(f(g(z)))] \end{aligned} \quad (10)$$

In this work, we further assume that the loss functions ϕ and φ satisfy the following conditions:

$$\begin{cases} \phi'(x) > 0, \varphi'(x) < 0, \\ \phi''(x) \geq 0, \varphi''(x) \geq 0, \\ \exists a, \phi'(a) + \varphi'(a) = 0. \end{cases} \quad (11)$$

The assumptions for the losses ϕ and φ are very mild. Note that in WGAN $\phi(x) = \varphi(-x) = x$ is used, which satisfies Eq. (11). There are many other instances, such as $\phi(x) = \varphi(-x) = -\log(\sigma(-x))$, $\phi(x) = \varphi(-x) = x + \sqrt{x^2 + 1}$ and $\phi(x) = \varphi(-x) = \exp(x)$. Meanwhile, there also exist losses used in GANs that do not satisfy Eq. (11), e.g., the quadratic loss (Mao et al., 2016) and the hinge loss (Zhao et al., 2016; Lim & Ye, 2017; Miyato et al., 2018).

To devise a loss in LGANs, it is practical to let ϕ be an increasing function with non-decreasing derivative and set $\phi(x) = \varphi(-x)$. Moreover, the linear combinations of such losses still satisfy Eq. (11). Figure 13 illustrates some of these loss metrics.

Note that $\phi(x) = \varphi(-x) = -\log(\sigma(-x))$ is the objective of vanilla GAN. As we have shown, the vanilla GAN suffers from the gradient uninformative problem. However, as we will show next, when imposing the Lipschitz regularization, the resulting model as a specific case of LGANs behaves very well.

4.1. Theoretical Analysis

We now present the theoretical analysis of LGANs. First, we consider the existence and uniqueness of the optimal discriminative function.

Theorem 1. *Under Assumption (11) and if ϕ or φ is strictly convex, the optimal discriminative function f^* of Eq. (10) exists and is unique.*

Note that although WGAN does not satisfy the condition in Theorem 1, its solution still exists but is not unique. Specifically, if f^* is an optimal solution then $f^* + \alpha$ for any $\alpha \in \mathbb{R}$ is also an optimal solution. The following theorems can be regarded as a generalization of Proposition 1 to LGANs.

Theorem 2. *Assume $\phi'(x) > 0$, $\varphi'(x) < 0$, and the optimal discriminator f^* exists and is smooth. We have*

- For all $x \in \mathcal{S}_r \cup \mathcal{S}_g$, if it holds that $\nabla_{f^*(x)} J_D(x) \neq 0$, then there exists $y \in \mathcal{S}_r \cup \mathcal{S}_g$ with $y \neq x$ such that $|f^*(y) - f^*(x)| = k(f^*) \cdot \|y - x\|$;
- For all $x \in \mathcal{S}_r \cup \mathcal{S}_g - \mathcal{S}_r \cap \mathcal{S}_g$, there exists $y \in \mathcal{S}_r \cup \mathcal{S}_g$ with $y \neq x$ such that $|f^*(y) - f^*(x)| = k(f^*) \cdot \|y - x\|$;

- (c) If $\mathcal{S}_r = \mathcal{S}_g$ and $\mathcal{P}_r \neq \mathcal{P}_g$, then there exists (x, y) pair with both points in $\mathcal{S}_r \cup \mathcal{S}_g$ and $y \neq x$ such that $|f^*(y) - f^*(x)| = k(f^*) \cdot \|y - x\|$ and $\nabla_{f^*(x)} \mathring{J}_D(x) \neq 0$;
- (d) There is a unique Nash equilibrium between \mathcal{P}_g and f^* under the objective $J_D + \lambda \cdot k(f)^2$, where it holds that $\mathcal{P}_r = \mathcal{P}_g$ and $k(f^*) = 0$.

The proof is given in Appendix A.2. This theorem states the basic properties of LGANs, including the existence of unique Nash equilibrium where $\mathcal{P}_r = \mathcal{P}_g$ and the existence of *bounding relationships* in the optimal discriminative function (i.e., $\exists y \neq x$ such that $|f^*(y) - f^*(x)| = k(f^*) \cdot \|y - x\|$). The former ensures that the objective is a well-defined distance metric, and the latter, as we will show next, eliminates the gradient uninformative problem.

It is worth noticing that the penalty $k(f)$ is in fact necessary for Property-(c) and Property-(d). The reason is due to the existence of the case that $\nabla_{f^*(x)} \mathring{J}_D(x) = 0$ for $\mathcal{P}_r(x) \neq \mathcal{P}_g(x)$. Minimizing $k(f)$ guarantees that the only Nash equilibrium is achieved when $\mathcal{P}_r = \mathcal{P}_g$. In WGAN, minimizing $k(f)$ is not necessary. However, if $k(f)$ is not minimized towards zero, $\nabla_x f^*(x)$ is not guaranteed to be zero at the convergence state $\mathcal{P}_r = \mathcal{P}_g$ where any function subject to 1-Lipschitz constraint is an optimal f^* in WGAN. It implies that minimizing $k(f)$ also benefits WGAN.

4.2. Refining the Bounding Relationship

From Theorem 2, we know that for any point x , as long as $\mathring{J}_D(x)$ does not hold a zero gradient with respect to $f^*(x)$, $f^*(x)$ must be bounded by another point y such that $|f^*(y) - f^*(x)| = k(f^*) \cdot \|y - x\|$. We further clarify that when there is a bounding relationship, it must involve both real sample(s) and fake sample(s). More formally, we have

Theorem 3. *Under the conditions in Theorem 2, we have*

- 1) For any $x \in \mathcal{S}_g$, if $\nabla_{f^*(x)} \mathring{J}_D(x) > 0$, then there must exist some $y \in \mathcal{S}_r$ with $y \neq x$ such that $f^*(y) - f^*(x) = k(f^*) \cdot \|y - x\|$ and $\nabla_{f^*(y)} \mathring{J}_D(y) < 0$;
- 2) For any $y \in \mathcal{S}_r$, if $\nabla_{f^*(y)} \mathring{J}_D(y) < 0$, then there must exist some $x \in \mathcal{S}_g$ with $y \neq x$ such that $f^*(y) - f^*(x) = k(f^*) \cdot \|y - x\|$ and $\nabla_{f^*(x)} \mathring{J}_D(x) > 0$.

The intuition behind the above theorem is that samples from the same distribution (e.g., the fake samples) will not bound each other to violate the optimality of $\mathring{J}_D(x)$. So, when there is strict bounding relationship (i.e., it involves points that hold $\nabla_{f^*(x)} \mathring{J}_D(x) \neq 0$), it must involve both real and fake samples. It is worth noticing that if only it is not the overlapping case, all fake samples hold $\nabla_{f^*(x)} \mathring{J}_D(x) > 0$, while all real samples hold $\nabla_{f^*(y)} \mathring{J}_D(y) < 0$.

Note that there might exist a dozen real and fake samples that bound each other. Under the Lipschitz continuity condition, the bounding relationship on the value surface of f^* is

the basic building block that connects \mathcal{P}_r and \mathcal{P}_g , and each fake sample with $\nabla_{f^*(x)} \mathring{J}_D(x) \neq 0$ lies in at least one of these bounded relationships. Next we will further interpret the implication of bounding relationship and show that it guarantees meaningful $\nabla_x f^*(x)$ for all involved points.

4.3. The Implication of Bounding Relationship

Recall that the Proposition 1 states that $\nabla_{x_t} f^*(x_t) = \frac{y-x}{\|y-x\|}$. We next show that it is actually a direct consequence of bounding relationship between x and y . We formally state it as follows:

Theorem 4. *Assume function f is differentiable and its Lipschitz constant is k , then for all x and y which satisfy $y \neq x$ and $f(y) - f(x) = k \cdot \|y - x\|$, we have $\nabla_{x_t} f(x_t) = k \cdot \frac{y-x}{\|y-x\|}$ for all $x_t = tx + (1-t)y$ with $0 \leq t \leq 1$.*

In other words, if two points x and y bound each other in terms of $f(y) - f(x) = k \cdot \|y - x\|$, there is a straight line between x and y on the value surface of f . Any point in this line holds the maximum gradient slope k , and the gradient direction at any point in this line is pointing towards the $x \rightarrow y$ direction. The proof is provided in Appendix A.4.

Combining Theorems 2 and 3, we can conclude that when \mathcal{S}_r and \mathcal{S}_g are disjoint, the gradient $\nabla_x f^*(x)$ for each generated sample $x \in \mathcal{S}_g$ points towards some real sample $y \in \mathcal{S}_r$, which guarantees that $\nabla_x f^*(x)$ -based updating would pull \mathcal{P}_g towards \mathcal{P}_r at every step.

In fact, Theorem 2 provides further guarantee on the convergence. Property-(b) implies that for any generated sample $x \in \mathcal{S}_g$ that does not lie in \mathcal{S}_r , its gradient $\nabla_x f^*(x)$ must point towards some real sample $y \in \mathcal{S}_r$. And in the fully overlapped case, according to Property-(c), unless $\mathcal{P}_r = \mathcal{P}_g$, there must exist at least one pair of (x, y) in strict bounding relationship and $\nabla_x f^*(x)$ pulls x towards y . Finally, Property-(d) guarantees that the only Nash equilibrium is $\mathcal{P}_r = \mathcal{P}_g$ where $\nabla_x f^*(x) = 0$ for all generated samples.

5. Empirical Analysis

In this section, we empirically study the gradient uninformative problem and the performance of various objectives of Lipschitz GANs. The anonymous code is provided in the supplemental material.

5.1. Gradient Uninformativeness in Practice

According to our analysis, $\nabla_x f^*(x)$ for most traditional GANs is uninformative. Here we investigate the practical behaviors of the gradient uninformative. Note that the behaviors of GANs without restriction on \mathcal{F} are essentially identical. We choose the Least-Squares GAN whose f^* is relatively simple as the representative and study it with a set

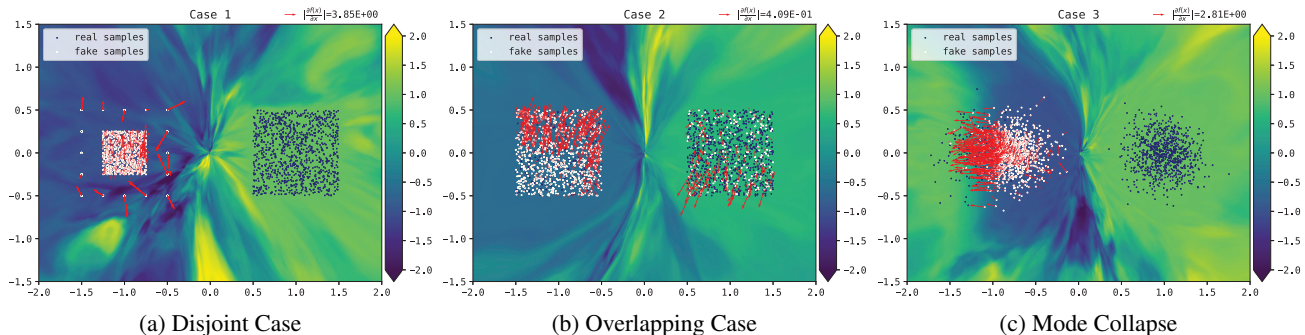


Figure 1: Practical behaviors of gradient uninformative: noisy gradient. Local greedy gradient leads to mode collapse.

of synthetic experiments which benefits the visualization.

The results are shown in Figure 1. We find that the gradient is very random, which we believe is the typical practical behavior of the gradient uninformative. Given the non-deterministic property of $f^*(x)$ for points out of $\mathcal{S}_r \cup \mathcal{S}_g$, $\nabla_x f^*(x)$ is highly sensitive to the hyper-parameters. We actually conduct the same experiments with a set of different hyper-parameters. The rest is provided in Appendix B.

In Section 3, we discussed the gradient uninformative under the circumstances that the fake sample is not surrounded by real samples. Actually, the problem of $\nabla_x f^*(x)$ in traditional GANs is more general, which can also be regarded as the gradient uninformative. For example, in the case of Figure 1b where the real and fake samples are both evenly distributed in the two regions with different densities, $f^*(x)$ is constant in each region and undefined outside. It theoretically has zero $\nabla_x f^*(x)$ for inner points and undefined $\nabla_x f^*(x)$ for boundary points. They in practice also behave as noisy gradient. We note that in the totally overlapping and continuous case, $\nabla_x f^*(x)$ is also ill-behaving, which seems to be an intrinsic cause of mode collapse, as illustrated in Figure 1c where \mathcal{P}_r and \mathcal{P}_g are both devised to be Gaussian(s).

5.2. Verifying $\nabla_x f^*(x)$ of LGANs

One important theoretical benefit of LGANs is that $\nabla_x f^*(x)$ for each generated sample is guaranteed to point towards some real sample. We here verify the gradient direction of $\nabla_x f^*(x)$ with a set of ϕ and φ that satisfy Eq. (11).

The tested objectives include: (a) $\phi(x) = \varphi(-x) = x$; (b) $\phi(x) = \varphi(-x) = -\log(\sigma(-x))$; (c) $\phi(x) = \varphi(-x) = x + \sqrt{x^2 + 1}$; (d) $\phi(x) = \varphi(-x) = \exp(x)$. And they are tested in two scenarios: two-dimensional toy data and real-world high-dimensional data. In the two-dimensional case, \mathcal{P}_r consists of two Gaussians and \mathcal{P}_g is fixed as one Gaussian which is close to one of the two real Gaussians, as illustrated in Figure 2. For the latter case, we use the CIFAR-10 training set. To make solving f^* feasible, we use ten CIFAR-10 images as \mathcal{P}_r , and ten fixed noise images as

\mathcal{P}_g . Note that we fix \mathcal{P}_g on purpose because to verify the direction of $\nabla_x f^*(x)$, learning \mathcal{P}_g is not necessary.

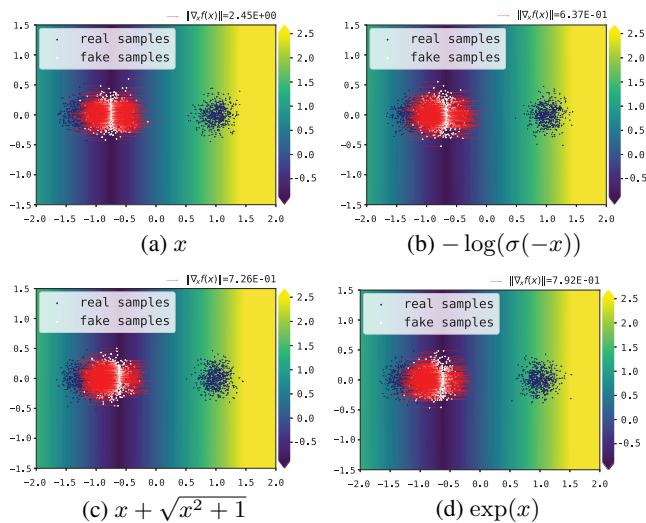
The results are shown in Figures 2 and 3, respectively. In Figure 2, we can see that the gradient of each generated sample is pointing towards some real sample. For the high dimensional case, visualizing the gradient direction is non-trivial. Hence, we plot the gradient and corresponding increments. In Figure 3, the leftmost in each row is a sample x from \mathcal{P}_g and the second is its gradient $\nabla_x f(x)$. The interiors are $x + \epsilon \cdot \nabla_x f(x)$ with increasing ϵ and the rightmost is the nearest real sample y from \mathcal{P}_r . This result visually demonstrates that the gradient of a generated sample is towards a real sample. Note that the final results of Figure 3 keep almost identical when varying the loss metric ϕ and φ in the family of LGANs.

5.3. Stabilized Discriminative Functions

The Wasserstein distance is a very special case that has solution under Lipschitz constraint. It is the only case where both ϕ and φ have constant derivative. As a result, f^* under the Wasserstein distance has a free offset, i.e., given some f^* , $f^* + \alpha$ with any $\alpha \in \mathbb{R}$ is also an optimal. In practice, it behaves as oscillations in $f(x)$ during training. The oscillations affect the practical performance of WGAN; Karras et al. (2017) and Adler & Lutz (2018) introduced regularization to the discriminative function to prevent $f(x)$ drifting during the training. By contrast, any other instance of LGANs does not have this problem. We illustrate the practical difference in Figure 5.

5.4. Max Gradient Penalty (MaxGP)

LGANs impose penalty on the Lipschitz constant of the discriminative function. There are works that investigate different implementations of Lipschitz continuity in GANs, such as gradient penalty (GP) (Gulrajani et al., 2017), Lipschitz penalty (LP) (Petzka et al., 2017) and spectral normalization (SN) (Miyato et al., 2018). However, the existing regularization methods do not directly penalize the Lipschitz constant. According to (Adler & Lutz, 2018), Lipschitz constant $k(f)$


 Figure 2: $\nabla_x f^*(x)$ in LGANs point towards real samples.

is equivalent to the maximum scale of $\|\nabla_x f(x)\|$. Both GP and LP penalize all gradients whose scales are larger than the given target Lipschitz constant k_0 . SN directly restricts the Lipschitz constant via normalizing the network weights by their largest eigenvalues. However, it is currently unclear how to effectively penalize the Lipschitz constant with SN.

To directly penalize Lipschitz constant, we approximate $k(f)$ in Eq. (10) with the maximum sampled gradient scale:

$$k(f) \simeq \max_x \|\nabla_x f(x)\|. \quad (12)$$


Practically, we follow (Gulrajani et al., 2017) and sample x as random interpolation of real and fake samples. We provide more details of this algorithm (MaxGP) in Appendix C.

According to our experiments, MaxGP in practice is usually comparable with GP and LP. However, in some of our synthetic experiments, we find that MaxGP is able to achieve the optimal discriminative function while GP and LP fail, e.g., the problem of solving f^* in Figure 3. Also, in some real data experiments, we find the training with GP or LP diverges and it is able to converge if we switch to MaxGP, e.g., the training with metric $\phi(x) = \varphi(-x) = \exp(x)$.

5.5. Benchmark with Unsupervised Image Generation

To quantitatively compare the performance of different objectives under Lipschitz constraint, we test them with unsupervised image generation tasks. In this part of experiments, we also include the hinge loss $\phi(x) = \varphi(-x) = \max(0, x + \alpha)$ and quadratic loss (Mao et al., 2016), which do not fit the assumption of strict monotonicity. For the quadratic loss, we set $\phi(x) = \varphi(-x) = (x + \alpha)^2$. To make the comparison simple, we fix $\psi(x)$ in the objective of generator as $-x$. We set $\alpha = 1.0$ in the experiment.

The strict monotonicity assumption of ϕ and φ is critical


 Figure 3: $\nabla_x f^*(x)$ gradation with CIFAR-10.

in Theorem 2 to theoretically guarantee the existences of bounding relationships for *arbitrary datas*. But if we further assume S_r and S_g are limited, it is possible that there exists a suitable λ such that all real and fake samples lie in a strict monotone region of ϕ and φ : for the hinge loss, it would mean $2\alpha < k(f) \cdot \|y - x\|$ for all $y \in S_r$ and $x \in S_g$.

The results in terms of Inception Score (IS) (Salimans et al., 2016) and Frechet Inception Distance (FID) (Heusel et al., 2017) are presented in Table 2. For all experiments, we adopt the network structures and hyper-parameter setting from (Gulrajani et al., 2017), where WGAN-GP in our implementation achieves $IS 7.71 \pm 0.03$ and $FID 18.86 \pm 0.13$ on CIFAR-10. We use MaxGP for all experiments and search the best λ in $[0.01, 0.1, 1.0, 10.0]$. We use 200,000 iterations for better convergence and use 500k samples to evaluate IS and FID for preferable stability. We note that IS is remarkably unstable during training and among different initializations. By contrast, FID is fairly stable.


From Table 2, we can see that LGANs generally work better than WGAN. Different LGANs have relatively similar final results, while the objectives $\phi(x) = \varphi(-x) = \exp(x)$ and $\phi(x) = \varphi(-x) = x + \sqrt{x^2 + 1}$ achieve the best performances. The hinge loss and quadratic loss with a suitable λ turn out to also work pretty good. We plot the training curves in terms of FID in Figures 4 and 6. Due to page limitation, we leave more results and details in Appendix D.

6. Related Work

WGAN (Arjovsky et al., 2017) based on the KR dual does not suffer from the gradient uninformative problem. We have shown that the Lipschitz constraint in the KR dual of the Wasserstein distance can be relaxed. With the new dual form, the resulting model suffers from the gradient

Table 2: Quantitative comparisons with unsupervised image generation.

Objective	CIFAR-10		Tiny ImageNet	
	IS	FID	IS	FID
x	7.68 ± 0.03	18.35 ± 0.12	8.66 ± 0.04	16.47 ± 0.04
$\exp(x)$	8.03 ± 0.03	15.64 ± 0.07	8.67 ± 0.04	14.90 ± 0.07
$-\log(\sigma(-x))$	7.95 ± 0.04	16.47 ± 0.11	8.70 ± 0.04	15.05 ± 0.07
$x + \sqrt{x^2 + 1}$	7.97 ± 0.03	16.03 ± 0.09	8.82 ± 0.03	15.11 ± 0.06
$(x + 1)^2$	7.97 ± 0.04	15.90 ± 0.09	8.53 ± 0.04	15.72 ± 0.11
$\max(0, x + 1)$	7.91 ± 0.04	16.52 ± 0.12	8.63 ± 0.04	15.75 ± 0.06

Figure 5: $f^*(x)$ in LGANs is more stable. Left: WGAN. Right: LGANs.

uninformativeness problem.

We have shown that Lipschitz constraint is able to ensure the convergence for a family of GAN objectives, which is not limited to the Wasserstein distance. For example, Lipschitz continuity is also introduced to the vanilla GAN (Miyato et al., 2018; Kodali et al., 2017; Fedus et al., 2017), achieving improvements in the quality of generated samples. As a matter of fact, the vanilla GAN objective $\phi(x) = \varphi(-x) = -\log(\sigma(-x))$ is a special case of our LGANs. Thus our analysis explains why and how it works. (Farnia & Tse, 2018) also provide some analysis on how f -divergence behaviors when combined with Lipschitz. However, their analysis is limited to the symmetric f -divergence.

Fedus et al. (2017) also argued that divergence is not the primary guide of the training of GANs. However, they thought that the vanilla GAN with a non-saturating generator objective somehow works. According to our analysis, given the optimal f^* , the vanilla GAN has no guarantee on its convergence. Unterthiner et al. (2017) provided some arguments on the unreliability of $\nabla_x f^*(x)$ in traditional GANs, which motivates their proposal of Coulomb GAN. However, the arguments there are not thorough. By contrast, we identify the gradient uninformativeness problem and link it to the restrictions on \mathcal{F} . Moreover, we have accordingly proposed a new solution, i.e., the Lipschitz GANs.

Some work studies the suboptimal convergence of GANs (Mescheder et al., 2017; 2018; Arora et al., 2017; Liu et al., 2017; Farnia & Tse, 2018), which is another important direction for theoretically understanding GANs. Despite the fact




Figure 4: Training curves on CIFAR.




Figure 6: Training curves on Tiny.

that the behaviors of suboptimal can be different, we think the optimal should well-behave in the first place, e.g., informative gradient and stable Nash equilibrium. Researchers found that applying Lipschitz continuity condition to the generator also benefits the quality of generated samples (Zhang et al., 2018; Odena et al., 2018). And (Qi, 2017) studied the Lipschitz condition from the perspective of loss-sensitive with a Lipschitz data density assumption.

7. Conclusion

In this paper we have studied one fundamental cause of failure in the training of GANs, i.e., the gradient uninformativeness issue. In particular, for generated samples which are not surrounded by real samples, the gradients of the optimal discriminative function $\nabla_x f^*(x)$ tell nothing about \mathcal{P}_r . That is, in a sense, there is no guarantee that \mathcal{P}_g will converge to \mathcal{P}_r . Typical case is that \mathcal{P}_r and \mathcal{P}_g are disjoint, which is common in practice. The gradient uninformativeness is common for unrestricted GANs and also appears in restricted GANs.

To address the nonconvergence problem caused by uninformative $\nabla_x f^*(x)$, we have proposed LGANs and shown that it makes $\nabla_x f^*(x)$ informative in the way that the gradient for each generated sample points towards some real sample. We have also shown that in LGANs, the optimal discriminative function exists and is unique, and the only Nash equilibrium is achieved when $\mathcal{P}_r = \mathcal{P}_g$ where $k(f^*) = 0$. Our experiments shown LGANs lead to more stable discriminative functions and achieve higher sample qualities.

Acknowledgements

This work is sponsored by APEX-YITU Joint Research Program. The corresponding authors Zhiming Zhou and Weinan Zhang thank the support of National Natural Science Foundation of China (61702327, 61772333, 61632017), Shanghai Sailing Program (17YF1428200) and the helpful discussions with Dachao Lin. Jiadong Liang and Zhihua Zhang have been supported by Beijing Municipal Commission of Science and Technology under Grant No. 181100008918005, and by Beijing Academy of Artificial Intelligence (BAAI).

References

- Adler, J. and Lunz, S. Banach Wasserstein GAN. *arXiv preprint arXiv:1806.06621*, 2018.
- Arjovsky, M. and Bottou, L. Towards principled methods for training generative adversarial networks. In *ICLR*, 2017.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- Arora, S., Ge, R., Liang, Y., Ma, T., and Zhang, Y. Generalization and equilibrium in generative adversarial nets (GANs). *arXiv preprint arXiv:1703.00573*, 2017.
- Bellemare, M. G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., and Munos, R. The Cramer distance as a solution to biased Wasserstein gradients. *arXiv preprint arXiv:1705.10743*, 2017.
- Farnia, F. and Tse, D. A convex duality framework for GANs. In *Advances in Neural Information Processing Systems 31*. 2018.
- Fedus, W., Rosca, M., Lakshminarayanan, B., Dai, A. M., Mohamed, S., and Goodfellow, I. Many paths to equilibrium: GANs do not need to decrease divergence at every step. *arXiv preprint arXiv:1710.08446*, 2017.
- Goodfellow, I. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of Wasserstein GANs. *arXiv preprint arXiv:1704.00028*, 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*, pp. 6626–6637, 2017.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Kodali, N., Abernethy, J., Hays, J., and Kira, Z. On convergence and stability of GANs. *arXiv preprint arXiv:1705.07215*, 2017.
- Lim, J. H. and Ye, J. C. Geometric GAN. *arXiv preprint arXiv:1705.02894*, 2017.
- Liu, S., Bousquet, O., and Chaudhuri, K. Approximation and convergence properties of generative adversarial learning. In *Advances in Neural Information Processing Systems*, pp. 5545–5553, 2017.
- Lucic, M., Kurach, K., Michalski, M., Gelly, S., and Bousquet, O. Are GANs created equal? a large-scale study. *arXiv preprint arXiv:1711.10337*, 2017.
- Mao, X., Li, Q., Xie, H., Lau, R. Y., Wang, Z., and Smolley, S. P. Least squares generative adversarial networks. *arXiv preprint ArXiv:1611.04076*, 2016.
- Mescheder, L., Nowozin, S., and Geiger, A. The numerics of GANs. In *Advances in Neural Information Processing Systems*, pp. 1825–1835, 2017.
- Mescheder, L., Geiger, A., and Nowozin, S. Which training methods for GANs do actually converge? In *International Conference on Machine Learning*, pp. 3478–3487, 2018.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Mroueh, Y. and Sercu, T. Fisher GAN. In *Advances in Neural Information Processing Systems*, pp. 2510–2520, 2017.
- Mroueh, Y., Li, C., Sercu, T., Raj, A., and Cheng, Y. Sobolev GAN. *arXiv preprint arXiv:1711.04894*, 2017.
- Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pp. 271–279, 2016.
- Odena, A., Buckman, J., Olsson, C., Brown, T. B., Olah, C., Raffel, C., and Goodfellow, I. Is generator conditioning causally related to GAN performance? *arXiv preprint arXiv:1802.08768*, 2018.
- Petzka, H., Fischer, A., and Lukovnicov, D. On the regularization of Wasserstein GANs. *arXiv preprint arXiv:1709.08894*, 2017.

- Qi, G.-J. Loss-sensitive generative adversarial networks on lipschitz densities. *arXiv preprint arXiv:1701.06264*, 2017.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*, pp. 2226–2234, 2016.
- Unterthiner, T., Nessler, B., Klambauer, G., Heusel, M., Ramsauer, H., and Hochreiter, S. Coulomb GANs: Provably optimal nash equilibria via potential fields. *arXiv preprint arXiv:1708.08819*, 2017.
- Villani, C. *Optimal Transport: Old and New*, volume 338. Springer Science & Business Media, 2008.
- Yadav, A., Shah, S., Xu, Z., Jacobs, D., and Goldstein, T. Stabilizing adversarial nets with prediction methods. *arXiv preprint arXiv:1705.07364*, 2017.
- Zhang, H., Goodfellow, I., Metaxas, D., and Odena, A. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*, 2018.
- Zhao, J., Mathieu, M., and LeCun, Y. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016.