
Bridging the Gap Between f -GANs and Wasserstein GANs

Jiaming Song¹ Stefano Ermon¹

Abstract

Generative adversarial networks (GANs) variants approximately minimize divergences between the model and the data distribution using a discriminator. Wasserstein GANs (WGANs) enjoy superior empirical performance, however, unlike in f -GANs, the discriminator does not provide an estimate for the ratio between model and data densities, which is useful in applications such as inverse reinforcement learning. To overcome this limitation, we propose a new training objective where we additionally optimize over a set of importance weights over the generated samples. By suitably constraining the feasible set of importance weights, we obtain a family of objectives which includes and generalizes the original f -GAN and WGAN objectives. We show that a natural extension outperforms WGANs while providing density ratios as in f -GAN, and demonstrate empirical success on distribution modeling, density ratio estimation and image generation.

1. Introduction

Learning generative models to sample from complex, high-dimensional distributions is an important task in machine learning with many important applications, such as image generation (Kingma & Welling, 2013), imitation learning (Ho & Ermon, 2016) and representation learning (Chen et al., 2016). Generative adversarial networks (GANs, Goodfellow et al. (2014)) are likelihood-free deep generative models (Mohamed & Lakshminarayanan, 2016) based on finding the equilibrium of a two-player minimax game between a generator and a critic (discriminator). Assuming the optimal critic is obtained, one can cast the GAN learning procedure as minimizing a discrepancy measure between the distribution induced by the generator and the training data distribution.

¹Stanford University. Correspondence to: Jimaing Song <tsong@cs.stanford.edu>.

Various GAN learning procedures have been proposed for different discrepancy measures. f -GANs (Nowozin et al., 2016) minimize a variational approximation of the f -divergence between two distributions (Csiszár, 1964; Nguyen et al., 2008). In this case, the critic acts as a density ratio estimator (Uehara et al., 2016; Grover & Ermon, 2017), i.e., it estimates if points are more likely to be generated by the data or the generator distribution. This includes the original GAN approach (Goodfellow et al., 2014) which can be seen as minimizing a variational approximation to the Jensen-Shannon divergence. Knowledge of the density ratio between two distributions can be used for importance sampling and in a range of practical applications such as mutual information estimation (Hjelm et al., 2018), off-policy policy evaluation (Liu et al., 2018), and de-biasing of generative models (Grover et al., 2019).

Another family of GAN approaches are developed based on Integral Probability Metrics (IPMs, Müller (1997)), where the critic (discriminator) is restricted to particular function families. For the family of Lipschitz-1 functions, the IPM reduces to the Wasserstein-1 or earth mover’s distance (Rubner et al., 2000), which motivates the Wasserstein GAN (WGAN, Arjovsky et al. (2017)) setting. Various approaches have been applied to enforce Lipschitzness, including weight clipping (Arjovsky et al., 2017), gradient penalty (Gulrajani et al., 2017) and spectral normalization (Miyato et al., 2018). Despite its strong empirical success in image generation (Karras et al., 2017; Brock et al., 2018), the learned critic cannot be interpreted as a density ratio estimator, which limits its usefulness for importance sampling or other GAN-related applications such as inverse reinforcement learning (Yu et al., 2019).

In this paper, we address this problem via a generalized view of f -GANs and WGANs. The generalized view introduces importance weights over the generated samples in the critic objective, allowing prioritization over the training of different samples. The algorithm designer can select suitable feasible sets to constrain the importance weights; we show that both f -GAN and WGAN are special cases to this generalization when specific feasible sets are considered. We further discuss cases that select alternative feasible sets where divergences other than f -divergence and IPMs can be obtained.

To derive concrete algorithms, we turn to a case where the importance weights belong to the set of valid density ratios over the generated distribution. In certain cases, the optimal importance weights can be obtained via closed-form solutions, bypassing the need to perform an additional inner-loop optimization. We discuss one such approach, named KL-Wasserstein GAN (KL-WGAN), that is easy to implement from existing WGAN approaches, and is compatible with state-of-the-art GAN architectures. We evaluate KL-WGAN empirically on distribution modeling, density estimation and image generation tasks. Empirical results demonstrate that KL-WGAN enjoys superior quantitative performance compared to its WGAN counterparts on several benchmarks.

2. Preliminaries

Notations Let X denote a random variable with separable sample space \mathcal{X} and let $\mathcal{P}(\mathcal{X})$ denote the set of all probability measures over the Borel σ -algebra on \mathcal{X} . We use P, Q to denote probability measures, and $P \ll Q$ to denote P is absolutely continuous with respect to Q , i.e. the Radon-Nikodym derivative dP/dQ exists. Under $Q \in \mathcal{P}(\mathcal{X})$, the p -norm of a function $r : \mathcal{X} \rightarrow \mathbb{R}$ is defined as

$$\|r\|_p := \left(\int |r(\mathbf{x})|^p dQ(\mathbf{x}) \right)^{1/p}, \quad (1)$$

with $\|r\|_\infty = \lim_{p \rightarrow \infty} \|r\|_p$. The set of locally p -integrable functions is defined as

$$L^p(Q) := \{r : \mathcal{X} \rightarrow \mathbb{R} : \|r\|_p < \infty\}, \quad (2)$$

i.e. its norm with respect to Q is finite. We denote $L^p_{\geq 0}(Q) := \{r \in L^p(Q) : \forall \mathbf{x} \in \mathcal{X}, r(\mathbf{x}) \geq 0\}$ which considers non-negative functions in $L^p(Q)$. The space of probability measures wrt. Q is defined as

$$\Delta(Q) := \{r \in L^1_{\geq 0}(Q) : \|r\|_1 = 1\}. \quad (3)$$

For example, for any $P \ll Q$, $dP/dQ \in \Delta(Q)$ because $\int (dP/dQ) dQ = 1$. We define $\mathbb{1}$ such that $\forall \mathbf{x} \in \mathcal{X}$, $\mathbb{1}(\mathbf{x}) = 1$, and define $\text{im}(\cdot)$ and $\text{dom}(\cdot)$ as image and domain of a function respectively.

Fenchel duality For functions $g : \mathcal{X} \rightarrow \mathbb{R}$ defined over a Banach space \mathcal{X} , the Fenchel dual of $g, g^* : \mathcal{X}^* \rightarrow \mathbb{R}$ is defined over the dual space \mathcal{X}^* by:

$$g^*(\mathbf{x}^*) := \sup_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{x}^*, \mathbf{x} \rangle - g(\mathbf{x}), \quad (4)$$

where $\langle \cdot, \cdot \rangle$ is the duality pairing. For example, the dual space of \mathbb{R}^d is also \mathbb{R}^d and $\langle \cdot, \cdot \rangle$ is the usual inner product (Rockafellar, 1970).

Generative adversarial networks In generative adversarial networks (GANs, Goodfellow et al. (2014)), the goal is to fit an (empirical) data distribution P_{data} with an implicit generative model over \mathcal{X} , denoted as $Q_\theta \in \mathcal{P}(\mathcal{X})$. Q_θ is defined implicitly via the process $X = G_\theta(Z)$, where Z is a random variable with a fixed prior distribution. Assuming access to *i.i.d.* samples from P_{data} and Q_θ , a discriminator $T_\phi : \mathcal{X} \rightarrow [0, 1]$ is used to classify samples from the two distributions, leading to the following objective:

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} [\log T_\phi(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim Q_\theta} [\log(1 - T_\phi(\mathbf{x}))].$$

If we have infinite samples from P_{data} , and T_ϕ and Q_θ are sufficiently expressive, then the above minimax objective will reach an equilibrium where $Q_\theta = P_{\text{data}}$ and $T_\phi(\mathbf{x}) = 1/2$ for all $\mathbf{x} \in \mathcal{X}$.

2.1. Variational Representation of f -Divergences

For any convex and semi-continuous function $f : [0, \infty) \rightarrow \mathbb{R}$ satisfying $f(1) = 0$, the f -divergence (Csiszár, 1964; Ali & Silvey, 1966) between two probabilistic measures $P, Q \in \mathcal{P}(\mathcal{X})$ is defined as:

$$D_f(P\|Q) := \mathbb{E}_Q \left[f \left(\frac{dP}{dQ} \right) \right] \quad (5)$$

$$= \int_{\mathcal{X}} f \left(\frac{dP}{dQ}(\mathbf{x}) \right) dQ(\mathbf{x}), \quad (6)$$

if $P \ll Q$ and $+\infty$ otherwise. Nguyen et al. (2010) derive a general variational method to estimate f -divergences given only samples from P and Q .

Lemma 1 (Nguyen et al. (2010)). $\forall P, Q \in \mathcal{P}(\mathcal{X})$ such that $P \ll Q$, and differentiable f :

$$D_f(P\|Q) = \sup_{T \in L^\infty(Q)} I_f(T; P, Q), \quad (7)$$

$$\text{where } I_f(T; P, Q) := \mathbb{E}_P[T(\mathbf{x})] - \mathbb{E}_Q[f^*(T(\mathbf{x}))] \quad (8)$$

and the supremum is achieved when $T = f'(dP/dQ)$.

In the context of GANs, Nowozin et al. (2016) proposed variational f -divergence minimization where one estimates $D_f(P_{\text{data}}\|Q_\theta)$ with the variational lower bound in Eq.(7) while minimizing over θ the estimated divergence. This leads to the f -GAN objective:

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} [T_\phi(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim Q_\theta} [f^*(T_\phi(\mathbf{x}))], \quad (9)$$

where the original GAN objective is a special case for $f(u) = u \log u - (u + 1) \log(u + 1) + 2 \log 2$.

2.2. Integral Probability Metrics and Wasserstein GANs

For a fixed class of real-valued bounded Borel measurable functions \mathcal{F} on \mathcal{X} , the integral probability metric (IPM)

based on \mathcal{F} and between $P, Q \in \mathcal{P}(\mathcal{X})$ is defined as:

$$\text{IPM}_{\mathcal{F}}(P, Q) := \sup_{T \in \mathcal{F}} \left| \int T(\mathbf{x}) dP(\mathbf{x}) - \int T(\mathbf{x}) dQ(\mathbf{x}) \right|.$$

If for all $T \in \mathcal{F}$, $-T \in \mathcal{F}$ then $\text{IPM}_{\mathcal{F}}$ forms a metric over $\mathcal{P}(\mathcal{X})$ (Müller, 1997); we assume this is always true for \mathcal{F} in this paper (so we can remove the absolute values). In particular, if \mathcal{F} is the set of all bounded 1-Lipschitz functions with respect to the metric over \mathcal{X} , then the corresponding IPM becomes the Wasserstein distance between P and Q (Villani, 2008). This motivates the Wasserstein GAN objective (Arjovsky et al., 2017):

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} [T_{\phi}(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim Q_{\theta}} [T_{\phi}(\mathbf{x})], \quad (10)$$

where T_{ϕ} is regularized to be approximately k -Lipschitz for some k . Various approaches have been applied to enforce Lipschitzness of neural networks, including weight clipping (Arjovsky et al., 2017), gradient penalty (Gulrajani et al., 2017), and spectral normalization over the weights (Miyato et al., 2018).

Despite its strong empirical performance, WGAN has two drawbacks. First, unlike f -GAN (Lemma 1), it does not naturally recover a density ratio estimator from the critic. Granted, the WGAN objective corresponds to an f -GAN one (Sriperumbudur et al., 2009) when $f(x) = 0$ if $x = 1$ and $f(x) = +\infty$ otherwise, so that $f^*(x) = x$; however, we can no longer use Lemma 1 to recover density ratios given an optimal critic T , because the derivative $f'(x)$ does not exist. Second, WGAN places the same weight on the objective for each generated sample, which could be sub-optimal when the generated samples are of different qualities.

3. A Generalization of f -GANs and WGANs

In order to achieve the best of both worlds, we propose an alternative generalization to the critic objectives to *both* f -GANs and WGANs. Consider the following functional:

$$\begin{aligned} \ell_f(T, r; P, Q) \\ := \mathbb{E}_{\mathbf{x} \sim Q} [f(r(\mathbf{x}))] + \mathbb{E}_{\mathbf{x} \sim P} [T(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim Q} [r(\mathbf{x}) \cdot T(\mathbf{x})] \end{aligned} \quad (11)$$

which depends on the distributions P and Q , the critic function $T : \mathcal{X} \rightarrow \mathbb{R}$, and an additional function $r : \mathcal{X} \rightarrow \mathbb{R}$. For conciseness, we remove the dependency on the argument \mathbf{x} for T, r, P, Q in the remainder of the paper.

The function $r : \mathcal{X} \rightarrow \mathbb{R}$ here plays the role of ‘‘importance weights’’, as they changes the weights to the critic objective over the generator samples. When $r = dP/dQ$, the objective above simplifies to $\mathbb{E}_Q[f(dP/dQ)]$ which is exactly the definition of the f -divergence between P and Q (Eq. 6).

To recover an objective over only the critic T , we minimize ℓ_f as a function of r over a suitable set $\mathcal{R} \subseteq L_{\geq 0}^{\infty}(Q)$, thus

eliminating the dependence over r :

$$\mathcal{L}_f^{\mathcal{R}}(T; P, Q) := \inf_{r \in \mathcal{R}} \ell_f(T, r; P, Q) \quad (12)$$

We note that the minimization step is performed within a particular set $\mathcal{R} \subseteq L^{\infty}(Q)$, which can be selected by the algorithm designer. The choice of the set \mathcal{R} naturally gives rise to different critic objectives. As we demonstrate below (and in Figure 1), we can obtain critic objectives for f -GAN as well as WGANs as special cases via different choices of \mathcal{R} in $\mathcal{L}_f^{\mathcal{R}}(T; P, Q)$.

3.1. Recovering the f -GAN Critic Objective

First, we can recover the critic in the f -GAN objective by setting $\mathcal{R} = L_{\geq 0}^{\infty}(Q)$, which is the set of all non-negative functions in $L^{\infty}(Q)$. Recall from Lemma 1 the f -GAN objective:

$$D_f(P||Q) = \sup_{T \in L^{\infty}(Q)} I_f(T; P, Q) \quad (13)$$

where $I_f(T; P, Q) := \mathbb{E}_P[T] - \mathbb{E}_Q[f^*(T)]$ as defined in Lemma 1. The following proposition shows that when $\mathcal{R} = L_{\geq 0}^{\infty}(Q)$, we recover $I_f = \mathcal{L}_f^{\mathcal{R}}$.

Proposition 1. *Assume that f is differentiable at $[0, \infty)$. $\forall P, Q \in \mathcal{P}(\mathcal{X})$ such that $P \ll Q$, and $\forall T \in \mathcal{F} \subseteq L^{\infty}(Q)$ such that $\text{im}(T) \subseteq \text{dom}((f')^{-1})$,*

$$I_f(T; P, Q) = \inf_{r \in L_{\geq 0}^{\infty}(Q)} \ell_f(T, r; P, Q). \quad (14)$$

where $I_f(T; P, Q) := \mathbb{E}_P[T] - \mathbb{E}_Q[f^*(T)]$.

Proof. From Fenchel’s inequality we have for convex $f : \mathbb{R} \rightarrow \mathbb{R}$, $\forall T(\mathbf{x}) \in \mathbb{R}$ and $\forall r(\mathbf{x}) \geq 0$, $f(r(\mathbf{x})) + f^*(T(\mathbf{x})) \geq r(\mathbf{x})T(\mathbf{x})$ where equality holds when $T(\mathbf{x}) = f'(r(\mathbf{x}))$. Taking the expectation over Q , we have

$$\mathbb{E}_Q[f(r)] - \mathbb{E}_Q[rT] \geq -\mathbb{E}_Q[f^*(T)]; \quad (15)$$

applying this to the definition of $\ell_f(T, r; P, Q)$, we have:

$$\begin{aligned} \ell_f(T, r; P, Q) &:= \mathbb{E}_Q[f(r)] + \mathbb{E}_P[T] - \mathbb{E}_Q[rT] \\ &\geq \mathbb{E}_P[T] - \mathbb{E}_Q[f^*(T)] = I_f(T; P, Q). \end{aligned} \quad (16)$$

where the inequality comes from Equation 15. The inequality becomes an equality when $r(\mathbf{x}) = (f')^{-1}(T(\mathbf{x}))$ for all $\mathbf{x} \in \mathcal{X}$. We note that such a case can be achieved, i.e., $(f')^{-1}(T) \in L_{\geq 0}^{\infty}(Q)$, because $\forall \mathbf{x} \in \mathcal{X}$, $(f')^{-1}(T(\mathbf{x})) \in \text{dom}(f) = [0, \infty)$ from the assumption over $\text{im}(T)$. Therefore, taking the infimum over $r \in L_{\geq 0}^{\infty}(Q)$, we have:

$$I_f(T; P, Q) = \inf_{r \in L_{\geq 0}^{\infty}(Q)} \ell_f(T, r; P, Q), \quad (17)$$

which completes the proof. \square




Figure 1. (Left) Minimization over different \mathcal{R} in $\mathcal{L}_f^{\mathcal{R}}$ gives different critic objectives. Minimizing over $L_{\geq 0}^{\infty}(Q)$ recovers f -GAN (blue set), minimizing over $\{\mathbb{1}\}$ recovers WGAN (orange set), and minimizing over $\Delta(Q)$ recovers f -WGAN (green set). (Right) Naturally, as we consider smaller sets \mathcal{R} to minimize over, the critic objective becomes larger for the same T .

3.2. Recovering the WGAN Critic Objective

Next, we recover the WGAN critic objective (IPM) by setting $\mathcal{R} = \{\mathbb{1}\}$, where $\mathbb{1}(x) = 1$ is a constant function. First, we can equivalently rewrite the definition of an IPM using the following notation:

$$\text{IPM}_{\mathcal{F}}(P, Q) = \sup_{T \in \mathcal{F}} I_W(T; P, Q) \quad (18)$$

where I_W represents the critic objective. We show that $I_W = \mathcal{L}_f^{\mathcal{R}}$ when $\mathcal{R} = \{\mathbb{1}\}$ as follows.

Proposition 2. $\forall P, Q \in \mathcal{P}(\mathcal{X})$ such that $P \ll Q$, and $\forall T \in \mathcal{F} \subseteq L^{\infty}(Q)$:

$$I_W(T; P, Q) = \inf_{r \in \{\mathbb{1}\}} \ell_f(T, r; P, Q) \quad (19)$$

where $I_W(T; P, Q) := \mathbb{E}_P[T] - \mathbb{E}_Q[T]$.

Proof. As $\{\mathbb{1}\}$ has only one element, the infimum is:

$$\ell_f(T, \mathbb{1}; P, Q) = \mathbb{E}_Q[f(1)] + \mathbb{E}_P[T] - \mathbb{E}_Q[T] \quad (20)$$

$$= I_W(T; P, Q) \quad (21)$$

where we used $f(1) = 0$ for the second equality. \square

The above propositions show that $\mathcal{L}_f^{\mathcal{R}}$ generalizes both f -GAN and WGANs critic objectives by setting $\mathcal{R} = L_{\geq 0}^{\infty}(Q)$ and $\mathcal{R} = \{\mathbb{1}\}$ respectively.

3.3. Extensions to Alternative Constraints

The generalization with $\mathcal{L}_f^{\mathcal{R}}$ allows us to introduce new objectives when we consider alternative choices for the constraint set \mathcal{R} . We consider sets \mathcal{R} such that $\{\mathbb{1}\} \subseteq \mathcal{R} \subseteq L_{\geq 0}^{\infty}(Q)$. The following proposition shows for some fixed T , the corresponding objective with \mathcal{R} is bounded between the f -GAN objective (where $\mathcal{R} = L_{\geq 0}^{\infty}(Q)$) and the WGAN objective (where $\mathcal{R} = \{\mathbb{1}\}$).

Proposition 3. $\forall P, Q \in \mathcal{P}(\mathcal{X})$ such that $P \ll Q$, $\forall T \in L^{\infty}(Q)$ such that $\text{im}(T) \subseteq \text{dom}((f')^{-1})$, and $\forall \mathcal{R} \subseteq L_{\geq 0}^{\infty}(Q)$ such that $\{\mathbb{1}\} \subseteq \mathcal{R}$ we have:

$$I_f(T; P, Q) \leq \mathcal{L}_f^{\mathcal{R}}(T; P, Q) \leq I_W(T; P, Q). \quad (22)$$

Proof. In Appendix A. \square

We visualize this in Figure 1. Selecting the set \mathcal{R} allows us to control the critic objective in a more flexible manner, interpolating between the f -GAN critic and the IPM critic objective and finding suitable trade-offs. Moreover, if we additionally take the supremum of $\mathcal{L}_f^{\mathcal{R}}(T; P, Q)$ over T , the result will be bounded between the supremum of I_f over T (corresponding to the f -divergence) and the supremum of I_W over T , as stated in the following theorem.

Theorem 1. For $\{\mathbb{1}\} \subseteq \mathcal{R} \subseteq L_{\geq 0}^{\infty}(Q)$, define

$$D_{f, \mathcal{R}}(P \| Q) := \sup_{T \in \mathcal{F}} \mathcal{L}_f^{\mathcal{R}}(T; P, Q) \quad (23)$$

where $\mathcal{F} := \{T : \mathcal{X} \rightarrow \text{dom}((f')^{-1}), T \in L^{\infty}(Q)\}$. Then

$$D_f(P \| Q) \leq D_{f, \mathcal{R}}(P \| Q) \leq \sup_{T \in \mathcal{F}} I_W(T; P, Q). \quad (24)$$

Proof. In Appendix A. \square

A natural corollary is that $D_{f, \mathcal{R}}$ defines a divergence between two distributions.

Corollary 1. $D_{f, \mathcal{R}}(P \| Q)$ defines a divergence between P and Q : $D_{f, \mathcal{R}}(P \| Q) \geq 0$ for all $P, Q \in \mathcal{P}(\mathcal{X})$, and $D_{f, \mathcal{R}}(P \| Q) = 0$ if and only if $P = Q$.

This allows us to interpret the corresponding GAN algorithm as variational minimization of a certain divergence bounded between the corresponding f -divergence and IPM.

4. Practical f -Wasserstein GANs

As a concrete example, we consider the set $\mathcal{R} = \Delta(Q)$, which is the set of all valid density ratios over Q . We note that $\{\mathbb{1}\} \subset \Delta(Q) \subset L_{\geq 0}^{\infty}(Q)$ (see Figure 1), so the corresponding objective is a divergence (from Corollary 1). We can then consider the variational divergence minimization objective over $\mathcal{L}_f^{\Delta(Q)}(T; P, Q)$:

$$\inf_{Q \in \mathcal{P}(\mathcal{X})} \sup_{T \in \mathcal{F}} \inf_{r \in \Delta(Q)} \ell_f(T, r; P, Q), \quad (25)$$

We name this the “ f -Wasserstein GAN” (f -WGAN) objective, since it provides an interpolation between f -GAN and Wasserstein GANs while recovering a density ratio estimate between two distributions.

4.1. KL-Wasserstein GANs

For the f -WGAN objective in Eq.(25), the trivial algorithm would have to perform iterative updates to three quantities Q , T and r , which involves three nested optimizations. While this seems impractical, we show that for certain choices of f -divergences, we can obtain closed-form solutions for the

optimal $r \in \Delta(Q)$ in the innermost minimization; this bypasses the need to perform an inner-loop optimization over $r \in \Delta(Q)$, as we can simply assign the optimal solution from the close-form expression.

Theorem 2. *Let $f(u) = u \log u$ and \mathcal{F} a set of real-valued bounded measurable functions on \mathcal{X} . For any fixed choice of P, Q , and $T \in \mathcal{F}$, we have*

$$\arg \min_{r \in \Delta(Q)} \mathbb{E}_Q[f(r)] + \mathbb{E}_P[T] - \mathbb{E}_Q[r \cdot T] = \frac{e^T}{\mathbb{E}_Q[e^T]} \quad (26)$$

Proof. In Appendix A. □

The above theorem shows that if the f -divergence of interest is the KL divergence, we can directly obtain the optimal $r \in \Delta(Q)$ using Eq.(26) for any fixed critic T . Then, we can apply this r to the f -WGAN objective, and perform gradient descent updates on Q and T only. Avoiding the optimization procedure over r allows us to propose practical algorithms that are similar to existing WGAN procedures. In Appendix C, we show a similar argument with χ^2 -divergence, another f -divergence admitting a closed-form solution, and discuss its connections with the χ^2 -GAN approach (Tao et al., 2018).

4.2. Implementation Details

In Algorithm 1, we describe KL-Wasserstein GAN (KL-WGAN), a practical algorithm motivated by the f -WGAN objectives based on the observations in Theorem 2. We note that r_0 corresponds to selecting the optimal value for r from Theorem 2; once r_0 is selected, we ignore the effect of $E_Q[f(r_0)]$ to the objective and optimize the networks with the remaining terms, which corresponds to weighting the generated samples with r_0 ; the critic will be updated as if the generated samples are reweighted. In particular, $\nabla_\phi(D_0 - D_1)$ corresponds to the critic gradient (T , which is parameterized by ϕ) and $\nabla_\theta D_1$ corresponds to the generator gradient (Q , parameterized by θ).

In terms of implementation, the only differences between KL-WGAN and WGAN are between lines 8 and 11, where WGAN will assign $r_0(\mathbf{x}) = 1$ for all $\mathbf{x} \sim Q_m$. In contrast, KL-WGAN ‘‘importance weights’’ the samples using the critic, in the sense that it will assign higher weights to samples that have large $T_\phi(\mathbf{x})$ and lower weights to samples that have low $T_\phi(\mathbf{x})$. This will encourage the generator $Q_\theta(\mathbf{x})$ to put more emphasis on samples that have high critic scores. It is relatively easy to implement the KL-WGAN algorithm from an existing WGAN implementation, as we only need to modify the loss function. We present an implementation of KL-WGAN losses (in PyTorch) in Appendix B.

While the mini-batch estimation for $r_0(\mathbf{x})$ provides a biased estimate to the optimal $r \in \Delta(Q)$ (which according to The-

Algorithm 1 Pseudo-code for KL-Wasserstein GAN

- 1: **Input:** the (empirical) data distribution P_{data} ;
 - 2: **Output:** implicit generative model Q_θ .
 - 3: Initialize generator Q_θ and discriminator T_ϕ .
 - 4: **repeat**
 - 5: Draw $P_m := m$ i.i.d. samples from P_{data} ;
 - 6: Draw $Q_m := m$ i.i.d. samples from $Q_\theta(\mathbf{x})$.
 - 7: Compute $D_1 := \mathbb{E}_{P_m}[T_\phi(\mathbf{x})]$ (real samples)
 - 8: **for all** $\mathbf{x} \in Q_m$ (fake samples) **do**
 - 9: Compute $r_0(\mathbf{x}) := e^{T_\phi(\mathbf{x})} / \mathbb{E}_{Q_m}[e^{T_\phi(\mathbf{x})}]$
 - 10: **end for**
 - 11: Compute $D_0 := \mathbb{E}_{Q_m}[r_0(\mathbf{x})T_\phi(\mathbf{x})]$.
 - 12: Perform SGD over θ with $-\nabla_\theta D_0$;
 - 13: Perform SGD over ϕ with $\nabla_\phi(D_0 - D_1)$.
 - 14: Regularize T_ϕ to satisfy k -Lipschitzness.
 - 15: **until** Stopping criterion
 - 16: **return** learned implicit generative model Q_θ .
-

orem 2 is $e^{T_\theta(\mathbf{x})} / \mathbb{E}_Q[e^{T_\theta(\mathbf{x})}]$, i.e., normalized with respect to Q instead of over a minibatch of m samples as done in line 8), we found that this does not affect performance significantly. We further note that computing $r_0(\mathbf{x})$ does not require additional network evaluations, so the computational cost for each iteration is nearly identical between WGAN and KL-WGAN. To promote reproducible research, we include code in the supplementary material.

5. Related Work

5.1. f -divergences, IPMs and GANs

Variational f -divergence minimization and IPM minimization paradigms are widely adopted in GANs. A non-exhaustive list includes f -GAN (Nowozin et al., 2016), Wasserstein GAN (Arjovsky et al., 2017), MMD-GAN (Li et al., 2017), WGAN-GP (Gulrajani et al., 2017), SNGAN (Miyato et al., 2018), LSGAN (Mao et al., 2017), etc. The f -divergence paradigms enjoy better interpretations over the role of learned discriminator (in terms of density ratio estimation), whereas IPM-based paradigms enjoy better training stability and empirical performance. Prior work have connected IPMs with χ^2 divergences between mixtures of data and model distributions (Mao et al., 2017; Tao et al., 2018; Mroueh & Sercu, 2017); our approach can be applied to χ^2 divergences as well, and we discuss its connections with χ^2 -GAN in Appendix C.

Several works (Liu et al., 2017; Farnia & Tse, 2018) considered restricting function classes directly over the f -GAN objective; Husain et al. (2019) show that restricted f -GAN objectives are lower bounds to Wasserstein autoencoder (Tolstikhin et al., 2017) objectives, aligning with our argument for f -GAN and WGAN (Figure 1).

Our approach is most related to regularized variational f -divergence estimators (Nguyen et al., 2010; Ruderman et al., 2012) and linear f -GANs (Liu et al., 2017; Liu & Chaudhuri, 2018) where the function family \mathcal{F} is a RKHS with fixed “feature maps”. Different from these approaches, ours naturally allows the “feature maps” to be learned. Moreover, considering both restrictions allows us to bypass inner-loop optimization via closed-form solutions in certain cases (such as KL or χ^2 divergences); this leads to our KL-WGAN approach which is easy to implement from existing WGAN implementations, and also have similar computational cost per iteration.

5.2. Reweighting of Generated Samples

The learned discriminators in GANs can further be used to perform reweighting over the generated samples (Tao et al., 2018); these include rejection sampling (Azadi et al., 2018), importance sampling (Grover et al., 2019; Tao et al., 2018), and Markov chain monte carlo (Turner et al., 2018). These approaches can only be performed after training has finished, unlike our KL-WGAN case where discriminator-based reweighting are performed during training.

Moreover, prior reweighting approaches assume that the discriminator learns to approximate some (fixed) function of the density ratio $dP_{\text{data}}/dQ_{\theta}$, which does not apply directly to general IPM-based GAN objectives (such as WGAN); in KL-WGAN, we interpret the discriminator outputs as (un-normalized, regularized) log density ratios, introducing the density ratio interpretation to the IPM paradigm. We note that post-training discriminator-based reweighting can also be applied to our approach, and is orthogonal to our contributions; we leave this as future work.

6. Experiments

We release code for our experiments (implemented in PyTorch) in <https://github.com/ermongroup/f-wgan>.

6.1. Synthetic and UCI Benchmark Datasets

We first demonstrate the effectiveness of KL-WGAN on synthetic and UCI benchmark datasets (Asuncion & Newman, 2007) considered in (Wenliang et al., 2018). The 2-d synthetic datasets include Mixture of Gaussians (MoG), Banana, Ring, Square, Cosine and Funnel; these datasets cover different modalities and geometries. We use Red-Wine, WhiteWine and Parkinsons from the UCI datasets. We use the same SNGAN (Miyato et al., 2018) architectures for WGAN and KL-WGANs, which uses spectral normalization to enforce Lipschitzness (detailed in Appendix D).

After training, we draw 5,000 samples from the generator and then evaluate two metrics over a fixed validation set. One is the negative log-likelihood (NLL) of the validation

samples on a kernel density estimator fitted over the generated samples; the other is the maximum mean discrepancy (MMD, Borgwardt et al. (2006)) between the generated samples and validation samples. To ensure a fair comparison, we use identical kernel bandwidths for all cases.

Distribution modeling We report the mean and standard error for the NLL and MMD results in Tables 1 and 2 (with 5 random seeds in each case) for the synthetic datasets and UCI datasets respectively. The results demonstrate that our KL-WGAN approach outperforms its WGAN counterpart on all but the Cosine dataset. From the histograms of samples in Figure 2, we can visually observe where our KL-WGAN performs significantly better than WGAN. For example, WGAN fails to place enough probability mass in the center of the Gaussians in MoG and fails to learn a proper square in Square, unlike our KL-WGAN approaches.

Density ratio estimation We demonstrate that adding the constraint $r \in \Delta(Q)$ leads to effective density ratio estimators. We consider measuring the density ratio from synthetic datasets, and compare them with the original f -GAN with KL divergence. We evaluate the density ratio estimation quality by multiplying dQ with the estimated density ratios, and compare that with the density of P ; ideally the two quantities should be identical. We demonstrate empirical results in Figure 3, where we plot the samples used for training, the ground truth density of P and the two estimates given by two methods. In terms of estimating density ratios, our proposed approach is comparable to the f -GAN one.

Stability of critic objectives For the MoG, Square and Cosine datasets, we further show the estimated divergences over a batch of 256 samples in Figure 4, where WGAN uses I_W and KL-WGAN uses the proposed $\mathcal{L}_f^{\Delta(Q)}$. While both estimated divergences decrease over the course of training, our KL-WGAN divergence is more stable on all three cases. In addition, we evaluate the number of occurrences when a negative estimate of the divergences was produced for an epoch (which contradicts the fact that divergences should be non-negative); over 500 batches, WGAN has 46, 181 and 55 occurrences on MoG, Square and Cosine respectively, while KL-WGAN only has 29, 100 and 7 occurrences. This suggests that the proposed objective is easier to estimate and optimize, and is more stable across different iterations.

6.2. Image Generation

We further evaluate our KL-WGAN’s practical on image generation tasks on CIFAR10 and CelebA datasets. Our experiments are based on the BigGAN (Brock et al., 2018) PyTorch implementation¹. We use a smaller network than

¹<https://github.com/ajbrock/BigGAN-PyTorch>

Table 1. Negative Log-likelihood (NLL) and Maximum mean discrepancy (MMD, multiplied by 10^3) results on six 2-d synthetic datasets. Lower is better. W denotes the original WGAN objective, and KL-W denotes the proposed KL-WGAN objective.

Metric	GAN	MoG	Banana	Rings	Square	Cosine	Funnel
NLL	W	2.65 ± 0.00	3.61 ± 0.02	4.25 ± 0.01	3.73 ± 0.01	3.98 ± 0.00	3.60 ± 0.01
	KL-W	2.54 ± 0.00	3.57 ± 0.00	4.25 ± 0.00	3.72 ± 0.00	4.00 ± 0.01	3.57 ± 0.00
MMD	W	25.45 ± 7.78	3.33 ± 0.59	2.05 ± 0.47	2.42 ± 0.24	1.24 ± 0.40	1.71 ± 0.65
	KL-W	6.51 ± 3.16	1.45 ± 0.12	1.20 ± 0.10	1.10 ± 0.23	1.33 ± 0.23	1.08 ± 0.23

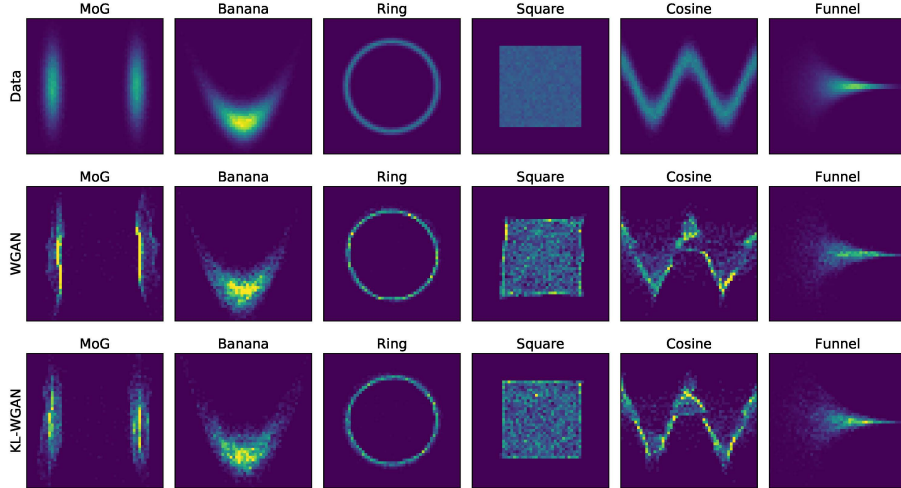


Figure 2. Histograms of samples from the data distribution (top), WGAN (middle) and our KL-WGAN (bottom).




Figure 3. Estimating density ratios. The first column contains the samples used for training, the second column is the ground truth density of P , the third and fourth columns are the density of Q times the estimated density ratios from original f -GAN (third column) and our KL-WGAN (fourth column).

Table 2. Negative Log-likelihood (NLL, top two rows) and Maximum mean discrepancy (MMD, multiplied by 10^3 , bottom two rows) results on real-world datasets. Lower is better for both evaluation metrics. W denotes the original WGAN objective, and KL denotes the proposed KL-WGAN objective.

	RedWine	WhiteWine	Parkinsons
W	14.55 ± 0.04	14.12 ± 0.02	20.24 ± 0.08
KL	14.41 ± 0.03	14.08 ± 0.02	20.16 ± 0.05
W	2.61 ± 0.37	1.32 ± 0.10	1.30 ± 0.09
KL	2.55 ± 0.11	1.23 ± 0.17	0.84 ± 0.04

the one reported in Brock et al. (2018) (implemented on TensorFlow), using the default architecture in the PyTorch implementation.

We compare training a BigGAN network with its original objective and training same network with our proposed KL-WGAN algorithm, where we add steps 8 to 11 in Algorithm 1. In addition, we also experimented with the original f -GAN with KL divergence; this failed to train properly due to numerical issues where exponents of very large critic




Figure 4. Estimated divergence with respect to training epochs (smoothed with a window of 10).

Table 3. Inception and FID scores for CIFAR10 image generation. We list comparisons with results reported by WGAN-GP (Gulrajani et al., 2017), Fisher GAN (Mroueh & Sercu, 2017), χ^2 GAN (Tao et al., 2018), MoLM (Ravuri et al., 2018), SNGAN (Miyato et al., 2018), NCSN (Song & Ermon, 2019), BigGAN (Brock et al., 2018) and Sphere GAN (Park & Kwon, 2019). (*) denotes our experiments with the PyTorch BigGAN implementation.

Method	Inception score	FID score
CIFAR10 Unconditional		
WGAN-GP	$7.86 \pm .07$	-
Fisher GAN	$7.90 \pm .05$	-
MoLM	$7.90 \pm .10$	18.9
SNGAN	$8.22 \pm .05$	21.7
Sphere GAN	$8.39 \pm .08$	17.1
NCSN	8.91	25.32
BigGAN*	$8.60 \pm .10$	16.38
KL-BigGAN*	$8.66 \pm .09$	15.23
CIFAR10 Conditional		
Fisher GAN	$8.16 \pm .12$	-
WGAN-GP	$8.42 \pm .10$	-
χ^2 -GAN	$8.44 \pm .10$	-
SNGAN	$8.60 \pm .08$	17.5
BigGAN	9.22	14.73
BigGAN*	$9.08 \pm .11$	9.51
KL-BigGAN*	$9.20 \pm .09$	9.17

Table 4. FID scores for CelebA image generation. The mean and standard deviation are obtained from 4 instances trained with different random seeds.

Method	Image Size	FID score
BigGAN	64×64	18.07 ± 0.47
KL-BigGAN		17.70 ± 0.32

values gives infinity values in the objective.

We report two common benchmarks for image generation, Inception scores (Salimans et al., 2016) and Fréchet Inception Distance (FID) (Heusel et al., 2017)² in Table 3 (CIFAR10) and Table 4 (CelebA). We do not report inception score on CelebA since the real dataset only has a score of less than 3, so the score is not very indicative of generation performance (Heusel et al., 2017). We show generated samples from the model in Appendix E.

Despite the strong performance of BigGAN, our method is able to consistently achieve superior inception scores and FID scores consistently on all the datasets and across different random seeds. This demonstrates that the KL-WGAN algorithm is practically useful, and can serve as a viable drop-in replacement for the existing WGAN objective even on state-of-the-art GAN models, such as BigGAN.

7. Conclusions

In this paper, we introduce a generalization of f -GANs and WGANs based on optimizing a (regularized) objective over importance weighted samples. This perspective allows us to recover both f -GANs and WGANs when different sets to optimize for the importance weights are considered. In addition, we show that this generalization leads to alternative practical objectives for training GANs and demonstrate its effectiveness on several different applications, such as distribution modeling, density ratio estimation and image generation. The proposed method only requires a small change in the original training algorithm and is easy to implement in practice.

In future work, we are interested in considering other constraints that could lead to alternative objectives and/or inequalities and their practical performances. It would also be interesting to investigate the KL-WGAN approaches on high-dimensional density ratio estimation tasks such as off-policy policy evaluation, inverse reinforcement learning and contrastive representation learning.

²Based on <https://github.com/mseitzer/pytorch-fid>

Acknowledgements

The authors would like to thank Lantao Yu, Yang Song, Abhishek Sinha, Yilun Xu and Shengjia Zhao for helpful discussions about the idea, proofreading a draft, and details about the image generation experiments. This research was supported by AFOSR (FA9550-19-1-0024), NSF (#1651565, #1522054, #1733686), ONR, and FLI.

References

- Ali, S. M. and Silvey, S. D. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society: Series B (Methodological)*, 28(1):131–142, 1966.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, January 2017.
- Asuncion, A. and Newman, D. UCI machine learning repository, 2007.
- Azadi, S., Olsson, C., Darrell, T., Goodfellow, I., and Odena, A. Discriminator rejection sampling. *arXiv preprint arXiv:1810.06758*, October 2018.
- Borgwardt, K. M., Gretton, A., Rasch, M. J., Kriegel, H.-P., Schölkopf, B., and Smola, A. J. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57, 2006.
- Brock, A., Donahue, J., and Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, September 2018.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 2172–2180. Curran Associates, Inc., 2016.
- Csiszár, I. Eine informationstheoretische ungleichung und ihre anwendung auf beweis der ergodizitaet von markoff-schen ketten. *Magyer Tud. Akad. Mat. Kutato Int. Koezl.*, 8:85–108, 1964.
- Farnia, F. and Tse, D. A convex duality framework for GANs. *arXiv preprint arXiv:1810.11740*, October 2018.
- 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.
- Grover, A. and Ermon, S. Boosted generative models. *arXiv preprint arXiv:1702.08484*, February 2017.
- Grover, A., Song, J., Agarwal, A., Tran, K., Kapoor, A., Horvitz, E., and Ermon, S. Bias correction of learned generative models using Likelihood-Free importance weighting. *arXiv preprint arXiv:1906.09531*, June 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*, pp. 5769–5779, 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. *arXiv preprint arXiv:1706.08500*, June 2017.
- Hjelm, D. R., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, August 2018.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pp. 4565–4573, 2016.
- Husain, H., Nock, R., and Williamson, R. C. Adversarial networks and autoencoders: The Primal-Dual relationship and generalization bounds. *arXiv preprint arXiv:1902.00985*, February 2019.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, October 2017.
- Kingma, D. P. and Welling, M. Auto-Encoding variational bayes. *arXiv preprint arXiv:1312.6114v10*, December 2013.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. MMD GAN: Towards deeper understanding of moment matching network. *arXiv preprint arXiv:1705.08584*, May 2017.
- Liu, Q., Li, L., Tang, Z., and Zhou, D. Breaking the curse of horizon: Infinite-Horizon Off-Policy estimation. *arXiv preprint arXiv:1810.12429*, October 2018.
- Liu, S. and Chaudhuri, K. The inductive bias of restricted f-GANs. *arXiv preprint arXiv:1809.04542*, September 2018.
- Liu, S., Bousquet, O., and Chaudhuri, K. Approximation and convergence properties of generative adversarial learning. *arXiv preprint arXiv:1705.08991*, May 2017.
- Mao, X., Li, Q., Xie, H., Lau, R. Y. K., Wang, Z., and Paul Smolley, S. Least squares generative adversarial

-
- networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2794–2802. openaccess.thecvf.com, 2017.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, February 2018.
- Mohamed, S. and Lakshminarayanan, B. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, October 2016.
- Mroueh, Y. and Sercu, T. Fisher GAN. *arXiv preprint arXiv:1705.09675*, May 2017.
- Müller, A. Integral probability metrics and their generating classes of functions. *Advances in applied probability*, 29(2):429–443, June 1997. ISSN 0001-8678, 1475-6064. doi: 10.2307/1428011.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *arXiv preprint arXiv:0809.0853*, (11): 5847–5861, September 2008. doi: 10.1109/TIT.2010.2068870.
- 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*, 56(11):5847–5861, 2010.
- Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: Training generative neural samplers using variational divergence minimization. *arXiv preprint arXiv:1606.00709*, June 2016.
- Park, S. W. and Kwon, J. Sphere generative adversarial network based on geometric moment matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4292–4301, 2019.
- Ravuri, S., Mohamed, S., Rosca, M., and Vinyals, O. Learning implicit generative models with the method of learned moments. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4311–4320, Stockholmsmässan, Stockholm Sweden, 2018. PMLR.
- Rockafellar, R. T. *Convex analysis*, volume 28. Princeton university press, 1970.
- Rubner, Y., Tomasi, C., and Guibas, L. J. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000.
- Ruderman, A., Reid, M., Garcia-Garcia, D., and Petterson, J. Tighter variational representations of f-divergences via restriction to probability measures. *arXiv preprint arXiv:1206.4664*, June 2012.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training GANs. *arXiv preprint arXiv:1606.03498*, June 2016.
- Song, Y. and Ermon, S. Generative modeling by estimating gradients of the data distribution. *arXiv preprint arXiv:1907.05600*, July 2019.
- Sriperumbudur, B. K., Fukumizu, K., Gretton, A., Schölkopf, B., and Lanckriet, G. R. G. On integral probability metrics, φ -divergences and binary classification. *arXiv preprint arXiv:0901.2698*, January 2009.
- Tao, C., Chen, L., Henao, R., Feng, J., and others. Chi-square generative adversarial network. *on Machine Learning*, 2018.
- Tolstikhin, I., Bousquet, O., Gelly, S., and Schoelkopf, B. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.
- Turner, R., Hung, J., Frank, E., Saatci, Y., and Yosinski, J. Metropolis-Hastings generative adversarial networks. *arXiv preprint arXiv:1811.11357*, November 2018.
- Uehara, M., Sato, I., Suzuki, M., Nakayama, K., and Matsuo, Y. Generative adversarial nets from a density ratio estimation perspective. *arXiv preprint arXiv:1610.02920*, 2016.
- Villani, C. *Optimal Transport: Old and New*. Springer Science & Business Media, October 2008. ISBN 9783540710509.
- Wenliang, L., Sutherland, D., Strathmann, H., and Gretton, A. Learning deep kernels for exponential family densities. *arXiv preprint arXiv:1811.08357*, November 2018.
- Yu, L., Song, J., and Ermon, S. Multi-Agent adversarial inverse reinforcement learning. *arXiv preprint arXiv:1907.13220*, July 2019.