MINIMAX OPTIMALITY (PROBABLY) DOESN'T IMPLY DISTRIBUTION LEARNING FOR GANS

Anonymous authors
Paper under double-blind review

ABSTRACT

Arguably the most fundamental question in the theory of generative adversarial networks (GANs) is to understand when GANs can actually learn the underlying distribution. Theoretical and empirical evidence (see e.g. (Arora et al., 2018)) suggests local optimality of the empirical training objective is insufficient. Yet, it does not rule out the possibility that achieving a true population minimax optimal solution might imply distribution learning. In this paper, we show that standard cryptographic assumptions imply that this stronger condition is still insufficient. Namely, we show that if local pseudorandom generators (PRGs) exist, then for a large family of natural target distributions, there are ReLU network generators of constant depth and poly size which take Gaussian random seeds so that (i) the output is far in Wasserstein distance from the target distribution, but (ii) no polynomially large Lipschitz discriminator ReLU network can detect this. This implies that even achieving a population minimax optimal solution to the Wasserstein GAN objective is likely insufficient for distribution learning in the usual statistical sense. Our techniques reveal a deep connection between GANs and PRGs, which we believe will lead to further insights into the computational landscape of GANs.

1 INTRODUCTION

When will a generative adversarial network (GAN) trained with samples from a distribution $\mathcal{D}$ actually output samples from a distribution that is close to $\mathcal{D}$? This question is one of the most foundational questions in GAN theory—indeed, it was raised since the original paper introducing GANs. However, despite significant interest, this question still remains to be fully understood for general classes of generators and discriminators.

A significant literature has developed discussing the role of the training dynamics (Liu et al., 2017; Li et al., 2018; Arora et al., 2018; Berard et al., 2019; Wiatrak et al., 2019; Thanh-Tung & Tran, 2020; Allen-Zhu & Li, 2021), as well as the generalization error of the GAN objective (Zhang et al., 2017; Arora et al., 2017; Thanh-Tung et al., 2019). In most cases, researchers have demonstrated that given sufficient training data, GANs are able to learn some specific form of distributions after successful training. Underlying these works appears to be a tacit belief that if we are able to achieve the minimax optimal solution to the population-level GAN objective, then the GAN should be able to learn the target distribution. In this work, we take a closer look at this assumption.

What does it mean to learn the target distribution? As a starting point, we must first formally define what we mean by learning a distribution; more concretely, what do we mean when we say that two distributions are close? The original paper of (Goodfellow et al., 2020) proposed to measure closeness with KL divergence. However, learning the target distribution in KL divergence is quite unlikely to be satisfied for real-world distributions. This is because learning distributions in KL divergence also requires us to exactly recover the support of the target distribution, which we cannot really hope to do if the distribution lies in an unknown (complicated) low-dimensional manifold. To rectify this, one may instead consider learning in Wasserstein distance, as introduced in the context of GANs by (Arjovsky et al., 2017), which has no such "trivial" barriers. Recall that the Wasserstein distance between two distributions $\mathcal{D}_1, \mathcal{D}_2$ over $\mathbb{R}^d$ is given by

$$W_1(\mathcal{D}_1, \mathcal{D}_2) = \sup_{\mathrm{Lip}(f) \leq 1} \mathbb{E}_{\mathcal{D}_1}[f] - \mathbb{E}_{\mathcal{D}_2}[f] \ ,$$

where for any $f : \mathbb{R}^d \to R$, we let $\mathrm{Lip}(f)$ denote the Lipschitz constant of f. That is, two densities are close in Wasserstein distance if no Lipschitz function can distinguish between them. In this work we will focus on Wasserstein distance as it is the most standard notion of distance between probability distributions considered in the context of GANs.

Note that if the class of discriminators contains sufficiently large neural networks, then minimax optimality of the GAN objective does imply learning in Wasserstein distance. This is because we can approximate any Lipschitz function arbitrarily well, with an exponentially large network with one hidden layer (see e.g. (Poggio et al., 2017)). Thus, in this case, minimizing the population GAN objective is actually equivalent to learning in Wasserstein distance. Of course in practice, however, we are limited to polynomially large networks for both the generator and the discriminator. This raises the natural question:

Does achieving small error against all poly-size neural network discriminators imply that the poly-size generator has learned the distribution in Wasserstein distance?

One might conjecture that this claim is true, since the generator is only of poly-size. Thus, using a (larger) poly-size discriminator (as opposite to the class of all 1-Lipschitz functions) might still be sufficient to minimize the actual Wasserstein distance. In this paper, however, we provide strong evidence to the contrary. We demonstrate that widely accepted cryptographic assumptions imply that this is false, *even if the generator is of constant depth*:

Theorem 1.1 (Informal, see Theorem 3.1)**.** *For any $n \in \mathbb{N}$, let γ_n be the standard Gaussian measure over $\mathbb{R}^n$. Assuming local pseudorandom generators exist, the following holds for any sufficiently large $m \in \mathbb{Z}$, $d, r \leq \mathrm{poly}(m)$, and any diverse*[1] *target distribution $\mathcal{D}^*$ over $[0,1]^d$ given by the pushforward of the uniform distribution I_r on $[0,1]^r$ by a constant depth ReLU network of polynomial size/Lipschitzness.*[2]

There exist generators $G : \mathbb{R}^m \to \mathbb{R}^d$ computed by (deterministic) ReLU networks of constant depth *and polynomial size for which no ReLU network of polynomial depth/size/Lipschitzness can tell apart the distributions $G(\gamma_m)$ and $\mathcal{D}^*$, yet $G(\gamma_m)$ and $\mathcal{D}^*$ are $\Omega(1)$-far in Wasserstein distance.*

While Theorem 1.1 pertains to the practically relevant setting of *continuous* seed and output distributions, we also give guarantees for the discrete setting. In fact, if we replace $\mathcal{D}^*$ and γ_m by the uniform distributions over $\{\pm 1\}^d$ and $\{\pm 1\}^m$, we show this holds for generators whose output coordinates are given by *constant-size* networks (see Theorem 3.2).

We defer the formal definition of local pseudorandom generators (PRGs) to Section 2.2. We pause to make a number of remarks about this theorem.

First, our theorem talks about the population loss of the GAN objective; namely, it says that the true population GAN objective is small for this generator G, meaning that for every ReLU network discriminator f of polynomial depth/size/Lipschitzness, we have that

$$|\mathbb{E}[f(\mathcal{D}^*)] - \mathbb{E}[f(G(\gamma_m))]| \leq \frac{1}{d^{\omega(1)}} \ .$$

In other words, our theorem states that even optimizing the true population minimax objective is insufficient for distribution learning. In fact, we show this even when the target distribution can be represented *perfectly* by some other generative model.

Second, notice that our generator is extremely simple: notably, it is only constant depth. On the other hand, the discriminator is allowed to be much more complex, namely any ReLU network of polynomial complexity. This discriminator class thus constitutes the most powerful family of functions we could hope to use in practice. Despite this, we show that the discriminators are still not powerful enough to distinguish the output of the (much simpler) generator from the target distribution.

Third, our conclusions hold both for $d \geq m$ and $d \leq m$, so long as the input and output dimensions are related by polynomial factors.

[1] See Definition 6. In the discussion proceeding this definition, we give a number of examples making clear that this is a mild and practically relevant assumption to make.

[2] When we say "polynomial," we are implicitly referring to the dependence on the parameter m, though because d, r are bounded by $\mathrm{poly}(m)$, "polynomial" could equivalently refer to the dependence on those parameters if they exceeded m.

Finally, we formally define the class of "diverse" target distributions for which our conclusions hold in Section 2.3. We note that this class is quite general: for instance, it includes pushforwards of the uniform distribution under random leaky ReLU networks (see Lemma 2.3).

Empirical results. To complement these theoretical results, we also perform some empirical validations of our findings (see Section 4). Our theorem is constructive; that is, given a local PRG, we give an explicit generator which satisfies the theorem. We instantiate this construction with Goldreich's PRG with the "Tri-Sum-And" (TSA) predicate (Goldreich, 2011), which is an explicit function which is believed to satisfy the local PRG property. We then demonstrate that a neural network discriminator trained via standard methods empirically cannot distinguish between the output of this generator and the uniform distribution. While of course we cannot guarantee that we achieve the truly optimal discriminator using these methods, this still demonstrates that our construction leads to a function which does appear to be hard to distinguish in practice.

GANs, PRGs, and circuit lower bounds. At a high level, our results and techniques demonstrate surprising and deep connections between GANs and more "classical" problems in cryptography and complexity theory. Theorem 1.1 already shows that cryptographic assumptions may pose a fundamental barrier to the most basic question in GAN theory. In addition to this, we also show a connection between this question and circuit lower bounds. In the supplementary material, we show that if we are able to unconditionally exhibit generators which can fool polynomially large ReLU network discriminators, then we would obtain breakthrough circuit lower bounds against TC^0 (see Theorem C.2 and Remark C.3). This complements Theorem 1.1, as it says that if we can unconditionally construct generators which fool realistic discriminators, then we make progress on long-standing questions in circuit complexity. We believe that exploring these connections may be crucial to achieving a deeper understanding of what GANs can and cannot learn.

1.1 Related Work

GANs and Distribution Learning The literature on GAN theory is vast and we cannot hope to do it full justice here. For a more extensive review, see e.g. (Gui et al., 2020). Besides the previously mentioned work on understanding GAN dynamics and generalization, we only mention the most relevant papers here. One closely related line of work derives concrete bounds on when minimax optimality of the GAN objective implies distribution learning (Bai et al., 2018; Liang, 2018; Singh et al., 2018; Uppal et al., 2019; Chen et al., 2020a; Schreuder et al., 2021). However, the rates they achieve scale poorly with the dimensionality of the data, and/or require strong assumptions on the class of generators and discriminators, such as invertability. Another line of work has demonstrated that first order methods can learn very simple GAN architectures in polynomial time (Feizi et al., 2017; Daskalakis et al., 2017; Gidel et al., 2019; Lei et al., 2020). However, these results do not cover many of the generators used in practice, such as ReLU networks with > 1 hidden layers.

Local PRGs and Learning PRGs have had a rich history of study in cryptography and complexity theory (see e.g. (Vadhan, 2012)). From this literature, the object most relevant to the present work is the notion of a *local PRG*. These are part of a broader research program of building constant parallel-time cryptography (Applebaum et al., 2006). One popular local PRG candidate was suggested in (Goldreich, 2011). By now there is compelling evidence that this candidate is a valid PRG, as a rich family of algorithms including sum-of-squares (ODonnell & Witmer, 2014) and statistical query algorithms (Feldman et al., 2018) provably cannot break it.

Finally, we remark that Goldreich's PRG and, more generally, hardness of refuting random CSPs have been used in a number of works showing hardness for various supervised learning problems (Daniely & Vardi, 2021; Daniely et al., 2014; Daniely, 2016; Daniely & Shalev-Shwartz, 2016; Applebaum et al., 2006; Applebaum & Raykov, 2016). We consider a very different setting, and our techniques are very different from the aforementioned papers.

2 Technical Preliminaries

Notation Denote by U_n the uniform distribution over $\{\pm 1\}^n$, by γ_n the standard n-dimensional Gaussian measure, and by I_n the uniform measure on $[0,1]^n$. Given distribution $\mathcal{D}$ over $\mathbb{R}^m$ and measurable function $G : \mathbb{R}^m \to \mathbb{R}^d$, let $G(\mathcal{D})$ denote the distribution over $\mathbb{R}^d$ given by the *push-*

forward of $\mathcal{D}$ under G—to sample from the pushforward $G(\mathcal{D})$, sample x from $\mathcal{D}$ and output $G(x)$. Given $\sigma : \mathbb{R} \to \mathbb{R}$ and vector $v \in \mathbb{R}^d$, denote by $\sigma(v) \in \mathbb{R}^d$ the result of applying σ entrywise to v. To avoid dealing with issues of real-valued computation, let $\mathbb{R}_\tau \subset \mathbb{R}$ be the set of multiples of $2^{-\tau}$ bounded in magnitude by 2^τ.

2.1 GANs and Pseudorandom Generators

In this section we review basic notions about generative models and PRGs.

Definition 1 (ReLU Networks)**.** *Let $\mathcal{C}_{L,S,d}$ denote the family of ReLU networks $F : \mathbb{R}^d \to \mathbb{R}$ of depth L and size S. Formally, $F \in \mathcal{C}_{L,S,d}$ if there exist weight matrices $\mathbf{W}_1 \in \mathbb{R}^{k_1 \times d}, \mathbf{W}_2 \in \mathbb{R}^{k_2 \times k_1}, \ldots, \mathbf{W}_L \in \mathbb{R}^{1 \times k_{L-1}}$ and biases $b_1 \in \mathbb{R}^{k_1}, b_2 \in \mathbb{R}^{k_2} \ldots, b_L \in \mathbb{R}$ such that*

$$F(x) \triangleq \mathbf{W}_L \phi\left(\mathbf{W}_{L-1}\phi\left(\cdots \phi(\mathbf{W}_1 x + b_1)\cdots\right) + b_{L-1}\right) + b_L,$$

and $\sum_{i=1}^{L-1} k_i = S$, where $\phi(z) \triangleq \max(0, z)$ is the ReLU activation. Additionally, we let $\mathcal{C}^{\tau,\Lambda}_{L,S,d}$ be the subset of such networks which are additionally Λ-Lipschitz and whose weight matrices and biases have entries in $\mathbb{R}_\tau$– we will refer to τ as the bit complexity *of the network.*

The following allows us to control the complexity of compositions of ReLU networks:

Lemma 2.1. *Let $J : \mathbb{R}^s \to \mathbb{R}^r$ be a function each of whose output coordinates is computed by some network in $\mathcal{C}^{\tau_1,\Lambda_1}_{L_1,S_1,s}$, and let $f \in \mathcal{C}^{\tau_2,\Lambda_2}_{L_2,S_2,r}$. Then $f \circ J \in \mathcal{C}^{\tau,\Lambda}_{L,S,s}$ for $\tau = \max(\tau_1, \tau_2)$, $\Lambda = \Lambda_1\Lambda_2\sqrt{r}$, $L = L_1 + L_2$, and $S = (S_1 + 1)r + S_2$. Furthermore, for the network in $\mathcal{C}^{\tau,\Lambda}_{L,S,s}$ realizing $f \circ J$, the bias and weight vector entries in the output layer lie in $\mathbb{R}_{\tau_2}$.*

Next, we formalize the probability metric we will work with.

Definition 2 (IPM)**.** *Given a family $\mathcal{F}$ of functions, define the $\mathcal{F}$-*integral probability metric *between two distributions p, q by $W_{\mathcal{F}}(p, q) = \sup_{f\in\mathcal{F}} |\mathbb{E}_{y\sim p}[f(y)] - \mathbb{E}_{y\sim q}[f(y)]|$. When $\mathcal{F}$ consists of the family of 1-Lipschitz functions, this is the standard* Wasserstein-1 *metric, which we denote by W_1.*

In the context of GANs, we will focus on discriminators given by ReLU networks of polynomial size, depth, Lipschitzness, and bit complexity:

Definition 3 (Discriminators)**.** *$\mathcal{F}^*$ denotes the set of all sequences of discriminators $f_d : \mathbb{R}^d \to \mathbb{R}$, indexed by $d \in \mathbb{N}$, whose size, depth, Lipschitzness, bit complexity grow at most polynomially in d.*

We now formalize the definition of GANs, which closely parallels the definition of PRGs.

Definition 4 (GANs/PRGs)**.** *Let $\epsilon : \mathbb{N} \to [0, 1]$ be an arbitrary function, and let $\{d(m)\}_{m\in\mathbb{N}}$ be some sequence of positive integers. Given a sequence of* seed distributions *$\{\mathcal{D}_m\}_m$ over $\mathbb{R}^m$, a sequence of* target distributions *$\{\mathcal{D}^*_{d(m)}\}$ over $\mathbb{R}^{d(m)}$, a family $\mathcal{F}$ of* discriminators *$f : \mathbb{R}^{d(m)} \to \mathbb{R}$, and a sequence of* generators *$G_m : \mathbb{R}^m \to \mathbb{R}^{d(m)}$, we say that $\{G_m\}$ ϵ-*fools *$\mathcal{F}$* relative to *$\{\mathcal{D}^*_{d(m)}\}$* with seed *$\{\mathcal{D}_m\}$ if for all sufficiently large m,*

$$\left|\mathbb{E}[f(G_m(\mathcal{D}_m))] - \mathbb{E}[f(\mathcal{D}^*_{d(m)})]\right| \leq \epsilon(m) \quad \forall f \in \mathcal{F}, f : \mathbb{R}^{d(m)} \to \mathbb{R}. \tag{1}$$

Note that in the notation of Definition 2, (1) is equivalent to $W_{\mathcal{F}}(G_m(\mathcal{D}_m), \mathcal{D}^*_{d(m)}) \leq \epsilon(m)$.[3]

In this definition, if the discriminators and generators were instead Boolean functions, we would refer to $\{G_m\}$ as pseudorandom generators.

Remark 2.2. It will often be cumbersome to refer to sequences of target/seed distributions and discriminators/generators as in Definitions 3 and 4, so occasionally we will refer to a single choice of m and d even though we implicitly mean that m and d are parameters that increase towards infinity. In this vein, we will often say that a single network f is in $\mathcal{F}^*$, though we really mean that f belongs to a sequence of networks which lies in $\mathcal{F}^*$. And for distributions p, q which implicitly belong to sequences $\{p_d\}, \{q_d\}$, when we refer to bounds on $W_{\mathcal{F}^*}(p, q)$ we really mean that for any sequence of discriminators $f_d \in \mathcal{F}^*$, $|\mathbb{E}[f_d(p_d)] - \mathbb{E}[f_d(q_d)]|$ is bounded.

[3]Here we are slightly abusing notation as $\mathcal{F}$ contains functions which have domain not equal to $\mathbb{R}^{d(m)}$, and we ignore these functions in the supremum.

2.2 LOCAL PSEUDORANDOM GENERATORS

In complexity theory, the role of neural network discriminators is played by *Boolean circuits*, which one can roughly think of as networks which take in Boolean strings as input and whose activations are logical operations (e.g. AND/OR/NOT) rather than ReLUs (see Section A.2 in the supplement).

It is widely believed that there exist so-called *local PRGs* capable of fooling all polynomial-sized Boolean circuits and all of whose output coordinates are functions of a constant number of input coordinates (Applebaum et al., 2006). One prominent candidate is Goldreich's PRG:

Definition 5 ((Goldreich, 2011)). *Let H be a collection of d subsets $S_1, \dots, S_d$ of $\{1, \dots, m\}$, each of size k and each sampled independently from the uniform distribution over subsets of $\{1, \dots, m\}$ of size k. Let $P : \{\pm 1\}^k \to \{\pm 1\}$ be a Boolean function; P is often referred to as a* predicate.

Let $G_{P,H} : \{\pm 1\}^m \to \{\pm 1\}^d$ denote the Boolean function whose ℓ-th output coordinate is computed by evaluating P on the coordinates of the input indexed by subset S_ℓ in H.

The following is a standard assumption in cryptography (see (Applebaum, 2016)), namely that the incredibly simple functions in Definition 5 fool all Boolean circuits of polynomial size:

Assumption 1. *For any $c > 1$, there exists $k \in \mathbb{N}$ and $P : \{\pm 1\}^k \to \{\pm 1\}$ such that for m sufficiently large, with probability $1 - o_m(1)$ over the randomness of H in Definition 5, the function $G = G_{P,H}$ $\mathsf{negl}(m)$-fools all polynomial-size Boolean circuits relative to U_{m^c} with seed U_m for some negligible function $\mathsf{negl} : \mathbb{N} \to [0, 1]$.*[4]

2.3 DIVERSE DISTRIBUTIONS

Recall that the main result of this paper is to construct generators that look indistinguishable from natural target distributions $\mathcal{D}^*$ according to any poly-sized neural network, but which are far from $\mathcal{D}^*$ in Wasserstein. In this section we describe in greater detail the properties that these $\mathcal{D}^*$ satisfy.

Definition 6. *A distribution μ over $\mathbb{R}^d$ is (N, β)-diverse if for any discrete distribution ν on $\mathbb{R}^d$ supported on at most N points, $W_1(\mu, \nu) \geq \beta$.*

The main fact we use about (N, β)-diverse distributions is that they cannot be approximated by pushforwards of $U_{\log_2 N}$ (see Lemma A.12 in the supplement for a formal statement).

Note that Definition 6 is a very mild assumption that simply requires that the distribution not be tightly concentrated around a few points. Distributions that satisfy Definition 6 are both practically relevant and highly expressive. For starters, any reasonable real-world image distribution will be diverse as it will not be concentrated around a few unique images. One can also show that distributions like uniform distributions over well-separated discrete point sets and over $[0, 1]^d$ are diverse (Lemma A.13 and Lemma A.15 in supplement). Definition 6 also captures random expansive neural networks with leaky ReLU activations:

Lemma 2.3 (Random expansive leaky ReLU networks)**.** *For $k_0, \dots, k_L \in \mathbb{N}$ satisfying $k_i \geq 1.1k_{i-1}$ for all $i \in [L]$, let $\mathbf{W}_1 \in \mathbb{R}^{k_1 \times k_0}, \mathbf{W}_2 \in \mathbb{R}^{k_2 \times k_1}, \dots, \mathbf{W}_L \in \mathbb{R}^{k_L \times k_{L-1}}$ be random weight matrices, where every entry of $\mathbf{W}_i$ is an independent draw from $\mathcal{N}(0, 1/k_i)$. For the function $F : \mathbb{R}^{k_0} \to \mathbb{R}^{k_L}$ given by $F(x) \triangleq \mathbf{W}_L \psi_\lambda(\mathbf{W}_{L-1} \psi_\lambda(\cdots \psi_\lambda(\mathbf{W}_1 x) \cdots))$, where $\psi_\lambda(z) = \psi_\lambda(z) = z/2 + (1/2 - \lambda)|z|$ is the leaky ReLU activation, $F(I_{k_0})$ is $(2^m, \beta)$-diverse for $m = (k_0/2)\log(k_0/2) - k_L/1.1 - 1$ and $\beta = \Theta(\lambda)^L$ with probability at least $1 - \exp(-\Omega(d))$.*

For example, if $\lambda, L = \Theta(1)$, then $F(I_{k_0})$ is $(2^{\Omega(k_0 \log k_0)}, \Omega(1))$-diverse for k_0 sufficiently large.

It is easy to see that the networks in Lemma 2.3 can be implemented as ReLU networks, so our main theorem applies to target distributions given by pushing the uniform distribution over the continuous cube through a random expansive leaky ReLU network of constant depth.

3 FOOLING RELU NETWORK DISCRIMINATORS DOES NOT SUFFICE

In this section we will show that even though a generative model looks indistinguishable from some target distribution $\mathcal{D}^*$ according to any ReLU network in $\mathcal{F}^*$, it can be quite far from $\mathcal{D}^*$ in Wasser-

[4] $g : \mathbb{N} \to \mathbb{R}_{\geq 0}$ is negligible if for every polynomial p, $g(n) < |1/p(n)|$ for all sufficiently large n.

stein. We begin by describing a simple version of this result over discrete domains in Section 3.1. In Section 3.2 we extend this to target distributions over continuous domains, but where the generator still takes in a discrete-valued seed. Finally, in Section 3.3 we give a simple reduction that extends these results to give generators that take in a continuous-valued random (Gaussian) seed, culminating in the following main result:

Theorem 3.1. *Let $\{H_m\}_m$ be a sequence of generators $H_m : \mathbb{R}^{r(m)} \to \mathbb{R}^{d(m)}$ for $r(m), d(m) \le \mathrm{poly}(m)$ whose output coordinates are computable by networks in $\mathcal{C}^{\tau'(m),\Lambda'(m)}_{L'(m),S'(m),r(m)}$ for $\tau'(m), \Lambda'(m) \le \mathrm{poly}(m)$. Suppose that $H_m(I_{r(m)})$ is $(2^m, \Omega(1))$-diverse.*

Fix any $\epsilon : \mathbb{N} \to [0,1]$ satisfying $\epsilon(m) \ge \max(\mathsf{negl}(m), \exp(-O(m))$. Under Assumption 1, there is a sequence of generators $G_m : \mathbb{R}^m \to \mathbb{R}^{d(m)}$ such that for all m sufficiently large:

1. *Every output coordinate of G_m is computable by a network in $\mathcal{C}^{\tau,\Lambda}_{L,S,m}$ for*

$$\tau = \max(O(\log(\Lambda'(m) \cdot m \cdot d(m)/\epsilon(m))), \tau'(m), O(1)), \qquad \Lambda = O(\Lambda'(m)^2 \mathrm{poly}(m)/\epsilon(m)),$$

$$L = L'(m) + O(1), \qquad S = O(r(m)\log(1/\epsilon(m))) + 3m + S'(m).$$

2. $W_{\mathcal{F}^*}(G_m(\gamma_m), H_m(I_{r(m)})) \le \epsilon(m) \cdot \mathrm{poly}(m)$,
3. $W_1(G_m(\gamma_m), H_m(I_{r(m)})) \ge \Omega(1)$.

Note that a natural choice of parameters for H_m would be

$$\tau'(m) \le O(\log m), \;\; \Lambda'(m), S'(m), 1/\epsilon(m) \le \mathrm{poly}(m), \;\; L'(m) \le O(1).$$

(In fact, the $\mathrm{poly}(m)$ factor in bullet point 2 simply comes from the Lipschitzness of H_m, so $\epsilon(m)$ only needs to scale inversely in this quantity for $W_{\mathcal{F}^*}$ to be small.) Altogether, we conclude that G_m's output coordinates are computable by *constant-depth ReLU networks* with polynomial size and Lipschitzness and logarithmic bit complexity τ.

3.1 Stretching Bits to Bits

As a warmup, in this subsection we prove the following special case of Theorem 3.1 when the target distribution and seed distribution are discrete.

Theorem 3.2. *Under Assumption 1, for any constant $c > 1$, there is a sequence of generators $G_m : \mathbb{R}^m \to \mathbb{R}^{d(m)}$ for $d(m) \ge m^c$ such that for all m,*

1. *Every output coordinate of G_m is computable by a network in $\mathcal{C}^{\tau,\Lambda}_{L,S,m}$ for $\tau, \Lambda, L, S = O_c(1)$.*
2. $W_{\mathcal{F}^*}(G_m(U_m), U_{d(m)}) \le \mathsf{negl}(m)$,
3. $W_1(G_m(U_m), U_{d(m)}) \ge \Omega(1)$.

We emphasize that in this discrete setting, our quantitative guarantees are even stronger: *all parameters τ, Λ, L, S of the generator are constant*, and no polynomial-sized ReLU network can distinguish between $G_m(U_m)$ and $U_{d(m)}$ with even *non-negligible* advantage.

As discussed in the introduction, a basic but important building block in the proof of Theorem 3.2 is the connection between Goldreich's PRG and generative models computed by neural networks of constant depth/size/Lipschitzness. We begin by elaborating on this connection and showing that any predicate $\{\pm 1\}^k \to \{\pm 1\}$ can be implemented as a network in $\mathcal{C}^{\tau,\Lambda}_{L,S,d}$ where $\tau, \Lambda, L, S = O_k(1)$. As a consequence of known constructions for implementing Boolean functions with ReLU networks (see e.g. Lemma A.2 of (Chen et al., 2020b)), we can show:

Corollary 3.3. *For any $k \in \mathbb{N}$ and any H and P as in Definition 5, the coordinates of the output of $G_{P,H}$ are computed by networks in $\mathcal{C}^{\tau,\Lambda}_{L,S,m}$ for $\tau = O(k), \Lambda = \exp(O(k)), L = k, S = O(2^k k)$.*

Before we use this to prove Theorem 3.2, we need an extra technical ingredient to formalize the fact that a discriminator given by a ReLU network of polynomially bounded complexity yields a discriminator computable by a polynomial-sized Boolean circuit. The idea is that if $W_{\mathcal{F}^*}$ is large so that there exists some ReLU network discriminator, then because the input to the discriminator is sufficiently well concentrated, some affine threshold of the ReLU network can distinguish between the two distributions:

Lemma 3.4. *Given independent X and Y such that $\mathbb{E}[Y]-\mathbb{E}[X]=\alpha$ and for which $X-\mathbb{E}[X]$ and $Y-\mathbb{E}[Y]$ are σ^2-sub-Gaussian, there exists a threshold $t \in [\mathbb{E}[X]-O(\sigma\sqrt{\log(\sigma/|\alpha|)}), \mathbb{E}[Y]+O(\sigma\sqrt{\log(\sigma/|\alpha|)})]$ for which $|\,\mathbb{P}[X>t]-\mathbb{P}[Y>t]| \geq \min\left(1/2, \widetilde{\Omega}(|\alpha|/\sigma)\right)$.*

We can now sketch the proof of Theorem 3.2, deferring a formal argument to the supplement.

Proof sketch of Theorem 3.2. The parameter m will be clear from context in the following discussion, so for convenience we will refer to $d(m)$ and G_m as d and G. Let $k, P, G, \mathsf{negl}(\cdot)$ be such that the outcome of Assumption 1 holds. By Corollary 3.3, every output coordinate of G is computable by a network in $\mathcal{C}^{\tau,\Lambda}_{L,S,m}$ for $\tau = O(k), \Lambda = \exp(O(k)), L = k, S = O(2^k k)$.

The fact that $W_1(G(U_m), U_d) > 1/3$ follows from the fact that $G(U_m)$ is uniform over 2^m points on the hypercube (with multiplicity), and U_d is $(2^m, 1/3)$-diverse.

It remains to check that G fools $\mathcal{F}^*$ relative to U_d. Suppose to the contrary that there exists some $f : \mathbb{R}^d \to \mathbb{R}$ in $\mathcal{F}^*$ for which $|\mathbb{E}[f(G(U_m))] - \mathbb{E}[f(U_d)]| > d^{-a}$ for some constant $a > 0$. Note that for any threshold $t \in \mathbb{R}_\tau$, there is a Turing machine $\mathcal{M}_\tau : \{\pm 1\}^d \to \{\pm 1\}$ that computes $y \mapsto \mathrm{sgn}(f(y) - t)$ using τ bits of advice. By applying Lemma 3.4 to random variables $X = f(G(U_m))$ and $Y = f(U_d)$, we get that there exists a threshold $t \in \mathbb{R}_{\mathrm{poly}(d)}$ for which $|\mathbb{E}[\mathcal{M}_\tau(G(U_m))] - \mathbb{E}[\mathcal{M}_\tau(U_d)]| > 1/\mathrm{poly}(d)$, contradicting Assumption 1. □

3.2 From Binary Outputs to Continuous Outputs

In this section we show how to extend Theorem 3.2 to the setting where the target distribution $\mathcal{D}^*$ is a pushforward of the uniform distribution on $[0,1]^r$. At a high level, the idea will be to *post-process* the output of the generator constructed in Theorem 3.2. Roughly speaking, we take weighted averages of clusters of output coordinates from the generator in Theorem 3.2 and pass these averages through the pushforward map defining $\mathcal{D}^*$. Formally, we show:

Theorem 3.5. *Let $\{H_m\}_m$ be a sequence of generators $H_m : \mathbb{R}^{r(m)} \to \mathbb{R}^{d(m)}$ for $r(m), d(m) \leq \mathrm{poly}(m)$ whose output coordinates are computable by networks in $\mathcal{C}^{\tau'(m),\Lambda'(m)}_{L'(m),S'(m),r(m)}$ for $\tau'(m), \Lambda'(m) \leq \mathrm{poly}(m)$. Suppose that $H_m(I_{r(m)})$ is $(2^m, \Omega(1))$-diverse.*

Fix any $\epsilon : \mathbb{N} \to [0,1]$ satisfying $\epsilon(m) \geq \max(\mathsf{negl}(m), \exp(-O(m)))$. Under Assumption 1, there is a sequence of generators $G_m : \mathbb{R}^m \to \mathbb{R}^{d(m)}$ such that for all m sufficiently large:

1. *Every output coordinate of G_m is computable by a network in $\mathcal{C}^{\tau,\Lambda}_{L,S,m}$ for*
$$\tau = \max(O(\log(1/\epsilon(m))), \tau'(m), O(1)), \qquad \Lambda = O(\Lambda'(m)) \cdot \mathrm{poly}(m),$$
$$L = L'(m) + O(1), \qquad S = O(r(m) \cdot \log(1/\epsilon(m))) + S'(m).$$
2. $W_{\mathcal{F}^*}(G_m(U_m), H_m(I_{r(m)})) \leq \epsilon(m) \cdot \mathrm{poly}(m)$,
3. $W_1(G_m(U_m), H_m(I_{r(m)})) \geq \Omega(1)$.

Theorem 3.5 retains many of the nice properties of Theorem 3.2, e.g. it can tolerate distinguishing advantage $\epsilon(m)$ which is negligible, at the mild cost of an extra logarithmic dependence on $1/\epsilon(m)$ in the bit complexity τ. And as with Theorem 3.1, if the networks H_m are of constant depth, the resulting generators G_m are also of constant depth.

To prove Theorem 3.5, we begin by showing that for any pair of distributions which are close under the $W_{\mathcal{F}}$ metric for some family of neural networks $\mathcal{F}$, their pushforwards under a simple generative model will still be close under the $W_{\mathcal{F}'}$ metric for some slightly weaker family of networks $\mathcal{F}'$.

Lemma 3.6. *Fix parameters $\Lambda, \Lambda' > 1$. Let $\mathcal{F} = \mathcal{C}^{\tau,\Lambda}_{L,S,s}$. If $W_{\mathcal{F}}(p,q) \leq \epsilon$ for some distributions p, q on $\mathbb{R}^s$, then for any $J : \mathbb{R}^s \to \mathbb{R}^r$ each of whose output coordinates is computed by a function in $\mathcal{C}^{\tau,\Lambda'}_{L',S',s}$ for some $L' < L$ and $S' \leq \frac{S-r-1}{r}$, we have $W_{\mathcal{F}'}(J(p), J(q)) \leq 2\epsilon\Lambda'\sqrt{r}$ for $\mathcal{F}' = \mathcal{C}^{\tau',\Lambda}_{L-L',S'',r}$ where $\tau' = \tau - \lceil \log_2 \Lambda\sqrt{r} \rceil$ and $S'' = S - r(S'+1)$.*

Now recall that in Theorem 3.2, we exhibited a GAN which is close in $W_{\mathcal{F}^*}$ to U_s. Using Lemma 3.6, we can show that a certain simple pushforward of this GAN will be close in $W_{\mathcal{F}^*}$ to *the uniform distribution over* $[0,1]^r$ for r slightly smaller than s. The starting point is the following:

Fact 3.7. *For any $0 < \epsilon < 1$ and $n \geq \log_2(1/\epsilon)$, let $h : \mathbb{R}^n \to \mathbb{R}$ be given by $h(x) = \langle w, x + \mathbf{1} \rangle$ for $w = \left(1/4, 1/8, \ldots, (1/2)^{n+1}\right)$ and $\mathbf{1}$ the all-1's vector. Then $W_1(h(U_n), I_1) \leq \epsilon$.*

By leveraging Lemma 3.6 and Fact 3.7, we can get an approximation to the uniform distribution over $[0,1]^r$ out of the uniform distribution over $\{\pm 1\}^s$:

Lemma 3.8. *Suppose $\epsilon > 0$ satisfies $\log(1/\epsilon) \leq \mathrm{poly}(s)$. If a distribution $\widetilde{\mathcal{D}}$ over $\mathbb{R}^s$ satisfies $W_{\mathcal{F}^*}(\widetilde{\mathcal{D}}, U_s) \leq \epsilon$, then for $r \triangleq s/\lceil \log(1/\epsilon) \rceil$,[5] there is a function $J : \mathbb{R}^s \to \mathbb{R}^r$ each of whose output coordinates is computed by a function in $\mathcal{C}_{1,0,s}^{O(\log(1/\epsilon)),O(1)}$ such that $W_{\mathcal{F}^*}(J(\widetilde{\mathcal{D}}), I_r) \leq \epsilon \cdot \mathrm{poly}(r)$.*

By further combining Lemma 3.6 and Lemma 3.8, we can thus extend the latter from the uniform distribution on $[0,1]^r$ to simple pushforwards thereof.

Lemma 3.9. *Under the hypotheses of Lemma 3.8, for any $d \leq \mathrm{poly}(s)$ and any function $H : \mathbb{R}^r \to \mathbb{R}^d$ each of whose output coordinates is computed by a function in $\mathcal{C}_{L',S',r}^{\tau',\Lambda'}$ for $\tau' \leq \mathrm{poly}(s)$, there is a function $J' : \mathbb{R}^s \to \mathbb{R}^d$ each of whose output coordinates is computed by a function in $\mathcal{C}_{L'+1,r+S',s}^{\max(O(\log(1/\epsilon)),\tau'),O(\Lambda'\sqrt{d})}$ such that $W_{\mathcal{F}^*}(J'(\widetilde{\mathcal{D}}), H(I_r)) \leq \epsilon\Lambda' \cdot \mathrm{poly}(s)$.*

So for any pushforward $H(I_r)$ of the uniform distribution on $[0,1]^r$, Lemma 3.9 lets us take the GAN given by Theorem 3.2 and slightly post-process its output so that it is close in $W_{\mathcal{F}^*}$ to $H(I_r)$. We can now sketch Theorem 3.5's proof, deferring the full argument to the supplement.

Proof of Theorem 3.5. Condition 3 follows because $G(U_m)$ is a uniform distribution on 2^m points (w/ multiplicity) and $H(I_r)$ is $(2^m, \Omega(1))$-diverse. Next, let $s = r \cdot \lceil \log(1/\epsilon) \rceil$. As we are assuming $\epsilon \geq \exp(-O(m))$, $s \leq r \cdot m = m^c$ for some constant $c > 1$. Take G' to be the generator $G : \mathbb{R}^m \to \mathbb{R}^s$ constructed in Theorem 3.2. By applying Lemma 3.9 to $\widetilde{\mathcal{D}} = G'(U_m)$, we get a function $J' : \mathbb{R}^s \to \mathbb{R}^d$ each of whose output coordinates is computed by a function in $\mathcal{C}_{L'+1,r+S',s}^{\max(O(\log(1/\epsilon)),\tau'),O(\Lambda'\sqrt{d})}$ such that $W_{F_d}(J'(G'(U_m)), H(I_r)) \leq \epsilon\Lambda' \cdot \mathrm{poly}(m) \leq \epsilon \cdot \mathrm{poly}(m)$, so we get condition 2 of the theorem for $G \triangleq J' \circ G'$. Finally, by Lemma 2.1, every output coordinate of G can be realized by a network in $\mathcal{C}_{L,S,m}^{\tau,\Lambda}$ for $\tau = \max(O(\log(1/\epsilon)), \tau', O(1))$, $\Lambda = O(\Lambda'\sqrt{ds}) = O(\Lambda') \cdot \mathrm{poly}(m)$, $L = L' + O(1)$, and $S = O(s) + r + S' = O(s) + S'$ (where we used the fact that $r = s/\lceil \log(1/\epsilon) \rceil < s$). This establishes condition 1 of the theorem. □

3.3 From Binary Inputs to Continuous Inputs

In Theorem 3.5, we have shown how to go from U_m to any simple pushforward $H_m(I_{r(m)})$ of the uniform distribution over $[0,1]^{r(m)}$. Here we complete the proof of our main result, Theorem 3.1, by giving a simple reduction showing how to use *Gaussian seed* γ_m instead of U_m. At a high level, the idea will be to *pre-process* the inputs to the generator constructed in Theorem 3.5 by appending appropriate activations at the input layer. We will need the following elementary construction.

Lemma 3.10. *For any ξ for which $1/\xi \in \mathbb{R}_\tau$, the piecewise linear function h_ξ which outputs -1 for inputs $x \leq -\xi$, 1 for $x \geq \xi$, and x/ξ otherwise, lies in $\mathcal{C}_{2,2,1}^{\tau,1/\xi}$.*

The function h_ξ will let us approximately convert from γ_m to U_m. Specifically, the following says that if we want to approximate the output distribution of the generator in Theorem 3.5 using Gaussians instead of bits as seed, it suffices to attach entrywise applications of h_ξ at the input layer:

Lemma 3.11. *Let $G_0 : \mathbb{R}^m \to \mathbb{R}^d$ have output coordinates computable by networks in $\mathcal{C}_{L'',S'',m''}^{\tau'',\Lambda''}$. For any $\epsilon > 0$, let $\xi' \triangleq \epsilon / \left(\Lambda''\sqrt{(2/\pi)m^3 d}\right)$ and let ξ bethe multiplicative inverse of $\lceil 1/\xi' \rceil$.*

The function $G : \mathbb{R}^m \to \mathbb{R}^d$ given by $G(x) = G_0(h_\xi(x))$, where $h_\xi(x)$ denotes entrywise application of h_ξ defined in Lemma 3.10, satisfies that 1) each of the output coordinates of G is computable by a network in $\mathcal{C}_{L,S,m}^{\tau,\Lambda}$ for the parameters $\tau = \max(\tau'', O(\log(\Lambda'' md/\epsilon)))$, $\Lambda = O(\Lambda''^2 m^2\sqrt{d}/\epsilon)$, $L = L'' + 2$, and $S = 3m + S''$, and 2) $W_1(G_0(U_m), G(\gamma_m)) \leq \epsilon$.

[5] We will assume for simplicity that this is an integer, though it is not hard to handle the case where $\lceil \log(1/\epsilon) \rceil$ does not divide s.

Remark 3.12. Note the only fact we use about γ_m in the proof of Lemma 3.11 is that outside of an event with probability $O(m\xi)$, $h_\xi(\gamma_m)$ is uniform over $\{\pm 1\}^m$. In particular, the same would hold for any product measure each of whose coordinates is symmetric and anticoncentrated around zero. For instance, up to a constant factor in ξ', Lemma 3.11 also holds with γ_m replaced by I_m.

Proof of Theorem 3.1. Substitute the generator constructed in Theorem 3.5, call it G_0, into Lemma 3.11; let G be the generator resulting from the lemma. Recall from Theorem 3.5 that each output coordinate of G_0 lies in $\mathcal{C}^{\tau'',\Lambda''}_{L'',S'',m}$ for $\tau'' = \max(O(\log(1/\epsilon)), \tau', O(1))$, $\Lambda'' = O(\Lambda' \mathrm{poly}(m))$, $L'' = L' + O(1)$, $S'' = O(r\log(1/\epsilon)) + S'$. So by Lemma 3.11, every output coordinate of G is computable by a network in $\mathcal{C}^{\tau,\Lambda}_{L,S,m}$ for $\tau = \max(\tau'', O(\log(\Lambda' md/\epsilon))) = \max(O(\log(\Lambda md/\epsilon)), \tau', O(1))$, $\Lambda = O(\Lambda''^2 m^2 \sqrt{d}/\epsilon) = O(\Lambda'^2 \mathrm{poly}(m)/\epsilon)$, $L = L'' + 2 = L' + O(1)$, and $S = 3m + S'' = O(r\log(1/\epsilon)) + 3m + S''$. □

4 Experimental Results

To empirically demonstrate the existence of a constant depth generator that can fool polynomially-bounded discriminators, we evaluated the generator G given by Goldreich's PRG (Goldreich, 2011) with input dimension $m = 50$, output dimension $d = 200$, and predicate $P : \{\pm 1\}^5 \to \{\pm 1\}$ given by the popular TSA predicate, namely $P(x_1, \ldots, x_5) = x_1 \cdot x_2 \cdot x_3 \cdot (x_4 \wedge x_5)$. This is the smallest predicate under which Goldreich's candidate construction is believed to be secure.

The target distribution $\mathcal{D}^*$ is the uniform distribution U_{200} over $\{\pm 1\}^{200}$. As we prove in Lemma A.13, U_d is sufficiently diverse that $W_1(G(U_m), U_d) \geq \Omega(1)$. We trained four different discriminators given respectively by $1, 2, 3, 4$ hidden-layer ReLU networks, where each hidden layer is fully connected with dimensions 200×200, to discriminate the output of the generator $G(U_m)$ from the target distribution U_d. We used the Adam optimizer with step size 0.001 over the DCGAN training objective, with batch-size 128. As we can see in Figure 1, the test loss $\mathbb{E}[-\log(D(X))] + \mathbb{E}[-\log(1 - D(G(z)))] - 2\log(2)$ stays consistently above zero, indicating that the discriminator can not discriminate the true distribution from the generator output, even though the Wasserstein distance between these two distributions is provably large.

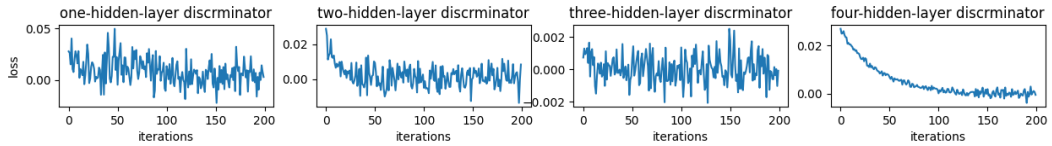

Figure 1: Test loss ($\mathbb{E}[-\log(D(X))] + \mathbb{E}[-\log(1 - D(G(z)))] - 2\log(2)$) over course of training. Discriminators cannot distinguish generator output from true distribution, though Wasserstein provably large.

5 Conclusions

In light of the obstructions presented in this paper, what are natural next steps for the theory of GANs? Here, we offer a couple of thoughts and possible future directions.

One limitation of our lower bound is that it holds for a specific generator. Of course, it is quite unlikely that we will ever encounter such a generator through natural GAN training. One way to circumvent our lower bound is to argue that the training dynamics of the generator may have some regularization effect which allows us to avoid these troublesome generators, and which allows GANs to learn distributions in polynomial time.

Another orthogonal perspective is that our results suggest that perhaps statistical learning is too strong of a goal. If our GAN is indeed indistinguishable from the target distribution to all polynomial time algorithms, then not only should the output of the GAN be sufficient for humans, but it should also be sufficient for all downstream applications, which presumably run in polynomial time. This raises the intriguing possibility that the correct metric for measuring closeness between distances in the context of GANs should inherently involve some computational component (e.g. in the sense of Dwork et al. (2021)) as opposed to the purely statistical metrics generally considered in the literature.

REFERENCES

Zeyuan Allen-Zhu and Yuanzhi Li. Forward super-resolution: How can gans learn hierarchical generative models for real-world distributions. *arXiv preprint arXiv:2106.02619*, 2021.

Benny Applebaum. Cryptographic hardness of random local functions. *Computational complexity*, 25(3):667–722, 2016.

Benny Applebaum and Pavel Raykov. Fast pseudorandom functions based on expander graphs. In *Theory of Cryptography Conference*, pp. 27–56. Springer, 2016.

Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. *SIAM Journal on Computing*, 36(4):845–888, 2006.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.

Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.

Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, and Yi Zhang. Generalization and equilibrium in generative adversarial nets (gans). In *International Conference on Machine Learning*, pp. 224–232. PMLR, 2017.

Sanjeev Arora, Andrej Risteski, and Yi Zhang. Do gans learn the distribution? some theory and empirics. In *International Conference on Learning Representations*, 2018.

Yu Bai, Tengyu Ma, and Andrej Risteski. Approximability of discriminators implies diversity in gans. In *International Conference on Learning Representations*, 2018.

Hugo Berard, Gauthier Gidel, Amjad Almahairi, Pascal Vincent, and Simon Lacoste-Julien. A closer look at the optimization landscapes of generative adversarial networks. *arXiv preprint arXiv:1906.04848*, 2019.

Lijie Chen and Roei Tell. Bootstrapping results for threshold circuits "just beyond" known lower bounds. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 34–41, 2019.

Minshuo Chen, Wenjing Liao, Hongyuan Zha, and Tuo Zhao. Statistical guarantees of generative adversarial networks for distribution estimation. *arXiv preprint arXiv:2002.03938*, 2020a.

Ruiwen Chen, Rahul Santhanam, and Srikanth Srinivasan. Average-Case Lower Bounds and Satisfiability Algorithms for Small Threshold Circuits. In Ran Raz (ed.), *31st Conference on Computational Complexity (CCC 2016)*, volume 50 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 1:1–1:35, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-008-8. doi: 10.4230/LIPIcs.CCC.2016.1. URL http://drops.dagstuhl.de/opus/volltexte/2016/5844.

Sitan Chen, Adam R Klivans, and Raghu Meka. Learning deep relu networks is fixed-parameter tractable. *arXiv preprint arXiv:2009.13512*, 2020b.

Amit Daniely. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 105–117, 2016.

Amit Daniely and Shai Shalev-Shwartz. Complexity theoretic limitations on learning dnf's. In *Conference on Learning Theory*, pp. 815–830. PMLR, 2016.

Amit Daniely and Gal Vardi. From local pseudorandom generators to hardness of learning. *arXiv preprint arXiv:2101.08303*, 2021.

Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. From average case complexity to improper learning complexity. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pp. 441–448, 2014.

Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. *arXiv preprint arXiv:1711.00141*, 2017.

Cynthia Dwork, Michael P Kim, Omer Reingold, Guy N Rothblum, and Gal Yona. Outcome indistinguishability. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1095–1108, 2021.

Soheil Feizi, Farzan Farnia, Tony Ginart, and David Tse. Understanding gans: the lqg setting. *arXiv preprint arXiv:1710.10793*, 2017.

Vitaly Feldman, Will Perkins, and Santosh Vempala. On the complexity of random satisfiability problems with planted solutions. *SIAM Journal on Computing*, 47(4):1294–1338, 2018.

Gauthier Gidel, Reyhane Askari Hemmat, Mohammad Pezeshki, Rémi Le Priol, Gabriel Huang, Simon Lacoste-Julien, and Ioannis Mitliagkas. Negative momentum for improved game dynamics. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1802–1811. PMLR, 2019.

Mikael Goldmann and Marek Karpinski. Simulating threshold circuits by majority circuits. *SIAM Journal on Computing*, 27(1):230–246, 1998.

Mikael Goldmann, Johan Håstad, and Alexander Razborov. Majority gates vs. general weighted threshold gates. *Computational Complexity*, 2(4):277–300, 1992.

Oded Goldreich. Candidate one-way functions based on expander graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pp. 76–87. Springer, 2011.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

Jie Gui, Zhenan Sun, Yonggang Wen, Dacheng Tao, and Jieping Ye. A review on generative adversarial networks: Algorithms, theory, and applications. *arXiv preprint arXiv:2001.06937*, 2020.

Russell Impagliazzo, Ramamohan Paturi, and Michael E Saks. Size–depth tradeoffs for threshold circuits. *SIAM Journal on Computing*, 26(3):693–707, 1997.

Qi Lei, Jason Lee, Alex Dimakis, and Constantinos Daskalakis. SGD learns one-layer networks in wgans. In *International Conference on Machine Learning*, pp. 5799–5808. PMLR, 2020.

Jerry Li, Aleksander Madry, John Peebles, and Ludwig Schmidt. On the limitations of first-order approximation in gan dynamics. In *International Conference on Machine Learning*, pp. 3005–3013. PMLR, 2018.

Tengyuan Liang. How well generative adversarial networks learn distributions. *arXiv preprint arXiv:1811.03179*, 2018.

Shuang Liu, Olivier Bousquet, and Kamalika Chaudhuri. Approximation and convergence properties of generative adversarial learning. *arXiv preprint arXiv:1705.08991*, 2017.

Ester Mariucci and Markus Reiß. Wasserstein and total variation distance between marginals of lévy processes. *Electronic Journal of Statistics*, 12(2):2482–2514, 2018.

Ryan ODonnell and David Witmer. Goldreich's prg: evidence for near-optimal polynomial stretch. In *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pp. 1–12. IEEE, 2014.

Tomaso Poggio, Hrushikesh Mhaskar, Lorenzo Rosasco, Brando Miranda, and Qianli Liao. Why and when can deep-but not shallow-networks avoid the curse of dimensionality: a review. *International Journal of Automation and Computing*, 14(5):503–519, 2017.

Mark Rudelson and Roman Vershynin. Smallest singular value of a random rectangular matrix. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 62(12):1707–1739, 2009.

Nicolas Schreuder, Victor-Emmanuel Brunel, and Arnak Dalalyan. Statistical guarantees for generative models without domination. In *Algorithmic Learning Theory*, pp. 1051–1071. PMLR, 2021.

Shashank Singh, Ananya Uppal, Boyue Li, Chun-Liang Li, Manzil Zaheer, and Barnabás Póczos. Nonparametric density estimation under adversarial losses. In *NeurIPS*, 2018.

Michael Sipser. Introduction to the theory of computation. *ACM Sigact News*, 27(1):27–29, 1996.

Hoang Thanh-Tung and Truyen Tran. Catastrophic forgetting and mode collapse in gans. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–10. IEEE, 2020.

Hoang Thanh-Tung, Truyen Tran, and Svetha Venkatesh. Improving generalization and stability of generative adversarial networks. *arXiv preprint arXiv:1902.03984*, 2019.

Ananya Uppal, Shashank Singh, and Barnabas Poczos. Nonparametric density estimation & convergence rates for gans under besov ipm losses. *Advances in Neural Information Processing Systems*, 32:9089–9100, 2019.

Salil P Vadhan. *Pseudorandomness*, volume 7. Now Delft, 2012.

Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d. *Computational Complexity*, 18(2):209–217, 2009.

Ingo Wegener. *The complexity of Boolean functions*. John Wiley & Sons, Inc., 1987.

Maciej Wiatrak, Stefano V Albrecht, and Andrew Nystrom. Stabilizing generative adversarial networks: A survey. *arXiv preprint arXiv:1910.00927*, 2019.

Pengchuan Zhang, Qiang Liu, Dengyong Zhou, Tao Xu, and Xiaodong He. On the discrimination-generalization tradeoff in gans. *arXiv preprint arXiv:1711.02771*, 2017.