

Convergence and Sample Complexity of SGD in GANs

Vasilis Kontonis

University of Wisconsin-Madison
kontonis@wisc.edu

Sihan Liu

University of Wisconsin-Madison
sliu556@wisc.edu

Christos Tzamos

University of Wisconsin-Madison
tzamos@wisc.edu

Abstract

We provide theoretical convergence guarantees on training Generative Adversarial Networks (GANs) via SGD. We consider learning a target distribution modeled by a 1-layer Generator network with a non-linear activation function $\phi(\cdot)$ parametrized by a $d \times d$ weight matrix \mathbf{W}_* , i.e., $f_*(\mathbf{x}) = \phi(\mathbf{W}_* \mathbf{x})$.

Our main result is that by training the Generator together with a Discriminator according to the Stochastic Gradient Descent-Ascent iteration proposed by Goodfellow et al. yields a Generator distribution that approaches the target distribution of f_* . Specifically, we can learn the target distribution within total-variation distance ϵ using $\tilde{O}(d^2/\epsilon^2)$ samples which is (near-)information theoretically optimal.

Our results apply to a broad class of non-linear activation functions ϕ , including ReLUs and is enabled by a connection with truncated statistics and an appropriate design of the Discriminator network. Our approach relies on a bilevel optimization framework to show that vanilla SGDA works.

1 Introduction

1.1 Background and Motivation

Since the influential work of [GPAM⁺14], Generative Adversarial Networks (GANs) have seen enormous success in diverse applications, see, for example, [ACB17, RMC15, AB17, JWS⁺20, ZXY18, HTP⁺17]. Despite their success in practice, very little is currently known about their theoretical guarantees in terms of generalization properties and the number of samples they require for training. In comparison, supervised models based on neural networks for classification, are much better understood through the theory of VC dimension and Rademacher complexity. One of the main reasons for the limited understanding of GANs is the fact that their training dynamics are quite complex as they correspond to a min-max game between two neural networks, the Generator and Discriminator. Analyzing such min-max games even in simple settings can be quite challenging, [MPPS16, KAHK17] as the natural methods commonly used for training based on stochastic gradient descent ascent (SGDA) fail to converge.

While there are countless versions of GANs proposed in the literature that are often times domain specific, we focus on the original GAN formulation proposed in [GPAM⁺14]. In particular, we consider the following min-max game

$$\min_{\mathcal{G}} \max_{\mathcal{D}} \mathbf{E}_{\mathbf{x} \sim \mathcal{T}} \log \mathcal{D}(\mathbf{x}) + \mathbf{E}_{\mathbf{x} \sim \mathcal{G}} \log(1 - \mathcal{D}(\mathbf{x})), \quad (1)$$

where \mathcal{T} is the true and unknown distribution, \mathcal{G} is the Generator distribution, and \mathcal{D} is the Discriminator. The above game is a *zero-sum* game between the Discriminator \mathcal{D} and the Generator \mathcal{G} . Our goal is to provide theoretical convergence guarantees on training such Generative Adversarial Networks (GANs) via SGD.

1.2 Our Contribution

In this work, we take a learning theoretic approach to understand the convergence and sample complexity of GANs for learning distributions corresponding to one layer neural nets. More formally we consider the following model.

Model. We assume that the underlying target distribution has the following form:

Samples are generated by drawing a random variable \mathbf{z} from a standard d -dimensional Gaussian distribution which are then transformed through a one-layer neural network. That is, for some unknown $d \times d$ parameter matrix \mathbf{W}_* and a known function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$, the output is equal to $\phi(\mathbf{W}_* \mathbf{z})$. We denote by $p(\mathbf{W}_*, \phi)$ the distribution of the random variable $\phi(\mathbf{W}_* \mathbf{z})$, where \mathbf{z} is drawn from a standard normal distribution.

This class of distributions corresponds to one layer neural networks with standard separable activation functions, like ReLU and sigmoid applied on each coordinate of the output. Moreover, it also captures much more complex non-linearities as it allows for arbitrary functions $\mathbb{R}^d \rightarrow \mathbb{R}^d$ that are given as input. For instance, it can capture multi-layer neural networks as long as only the parameters \mathbf{W}_* of the first layer are unknown while all others are fixed in advance.

Without any additional assumptions on the transformation function ϕ the problem is information theoretically intractable. We identify a natural property of the transformation ϕ that makes it possible to learn the underlying distribution using a GAN architecture without restricting the expressiveness of our model.

We require the transformation $\phi : \mathbb{R}^d \mapsto \mathbb{R}^d$ to be invertible with non-trivial probability over the samples of the true distribution $p(\mathbf{W}_, \phi)$.*

We note that, commonly used activation functions in neural networks are either fully invertible (e.g. sigmoid) or partially invertible like ReLU that is invertible when the coordinates of \mathbf{x} are positive.

More precisely we define the following class of transformed distributions.

Definition 1 (Partially Invertible Network). *A pair (\mathbf{W}_*, ϕ) composed of a weight matrix \mathbf{W}_* and an activation ϕ is denoted as a one layer partially invertible network if there exists some set $T \subseteq \mathbb{R}^d$ such that ϕ is invertible on T and $\mathcal{N}(T; \mathbf{W}_*) \geq \alpha > 0$ ¹, where $\alpha = \Omega(1)$ is some absolute constant.*

Our main result is that Generative Adversarial Networks with Partially Invertible Generator Networks converges to the true distribution when trained by stochastic gradient descent ascent. In particular, simultaneous training of the Generator with an appropriately designed Discriminator succeeds in learning the target distribution in polynomially many iterations and near-optimal sample complexity.

Theorem 1. *Consider samples generated by a Partially Invertible Network (\mathbf{W}_*, ϕ) for some unknown \mathbf{W}_* with bounded distance to \mathbf{I} :*

$$\max (\|\mathbf{W}_* \mathbf{W}_*^T - \mathbf{I}\|_F, \|(\mathbf{W}_* \mathbf{W}_*^T)^{-1} - \mathbf{I}\|_F) < c.$$

¹We use $\mathcal{N}(T; \mathbf{W}_*)$ to denote the mass of the set T under the normal distribution with covariance $\mathbf{W}_* \mathbf{W}_*^T$. See Section 2 for details.

Then, Nested Stochastic Gradient Descent-Ascent (Algorithm 1) uses $\tilde{O}_c(d^2/\epsilon^2)$ samples from $p(\mathbf{W}_*, \phi)$, performs $\tilde{O}_c(d^4/\epsilon^6)$ gradient updates, and converges to a matrix $\widetilde{\mathbf{W}}$ where $d_{\text{TV}}(p(\mathbf{W}_*, \phi), p(\widetilde{\mathbf{W}}, \phi)) \leq \epsilon$ with probability 99%.

Theorem 1 shows that Nested Stochastic Gradient Descent Ascent recovers a parameter matrix $\widetilde{\mathbf{W}}$ such that the Generator distribution $p(\widetilde{\mathbf{W}}, \phi)$ is close in total variation distance to the underlying distribution $p(\mathbf{W}_*, \phi)$. We note that as this is a gradient descent method the number of iterations naturally grows larger with the distance of the target matrix \mathbf{W}_* to the initial weights \mathbf{W} which are assumed to be \mathbf{I}^2 .

We also remark that the sample-complexity of our result is information theoretically optimal up to polylogarithmic factors as even the case where the transformation ϕ is the identity transformation (i.e., $\phi(\mathbf{x}) = \mathbf{x}$) that corresponds to learning the covariance matrix of a Gaussian distribution is well known to require $\Omega(d^2/\epsilon^2)$ samples in order to learn (the covariance of) a Gaussian within total variation distance ϵ .

A key challenge in showing Theorem 1 is to construct an appropriate Discriminator network that is powerful enough to distinguish between the true distribution and fake samples while, at the same time, simple enough to have few parameters and be efficiently trainable.

For any fixed Generator distribution there always exists a Discriminator that optimally distinguishes samples from the Generator and the target distribution (see Proposition 1 of [GPAM⁺14]). Unfortunately, this Discriminator may be arbitrarily complex. Even for a single layer neural network with ReLU activations, it requires treating samples differently according to their non-zero patterns, which is challenging to express directly as a simple low-depth neural network.

Instead, we focus on simple Discriminator networks that only discriminate samples that fall in the invertible region T of the transformation ϕ . In particular, our Discriminator first checks whether the received sample belongs to (the image of) set T , then performs the inverse transformation ϕ^{-1} followed by a quadratic layer and a sigmoid activation, see Figure 1b. For the full GAN architecture see Figure 1a. We train both Generator and Discriminator with Nested Stochastic Gradient Descent Ascent, that is we perform multiple iterations for the Discriminator per Generator update.

Our choice of Discriminator allows us to use techniques and ideas from truncated statistics, an area of statistics that deals with estimating the parameters of a distribution given only conditional samples from a subset of the distribution. The Discriminator essentially performs such a truncation operation to the data, see Figure 1b, as a result of the non-invertible function ϕ .

Learning from high-dimensional truncated datasets is a notoriously challenging task and a computationally efficient algorithm for Gaussian data was only recently obtained in [DGTZ18a] through maximum-likelihood estimation. As a byproduct of our analysis, we show that the min-max GAN iteration is an alternative computationally efficient and near-sample optimal approach for the task of learning a truncated Gaussian considered in [DGTZ18a].

1.3 Related Work

Our work is inspired and motivated mainly by the success of Generative Adversarial Neural nets in practice and aims to provide provable guarantees for their convergence and sample complexity. There are other works in the computer science and optimization communities that try to theoretically analyze the behavior of GANs. One such work related to ours is [FFGT17]. The authors consider the problem of learning a Gaussian distribution using a Wasserstein GAN which corresponds to the special case of our model without a non-linearity, i.e., $\phi(\mathbf{x}) = \mathbf{x}$. Another related work is [GM18].

²The dependence on the distance to \mathbf{I} can be eliminated by an additional preconditioning step using a few samples from the target distribution (See Remark 2).

The authors analyze a similar setting where the Generator is again linear (learning a Gaussian distribution) and the Discriminator is quadratic. They train their W-GAN using a custom method that they denote as “Crossing-the-curl”. Interestingly, they show that simultaneous alternating SGDA diverges in their setting. We view this as strong evidence that *Nested* SGDA is indeed required in order to have convergence for GANs. A third work in this direction is [MGN18], where the authors study local convergence of different GAN architectures. In particular, the results show that GAN training diverges when the underlying distribution is not absolutely continuous.

Prior work also studied how well GANs generalize - does minimization of GAN’s objective function offer any guarantee on the statistical distance between the Generator distribution and the target distribution? One such work is [LBC17], where the authors address the problems by giving a new notion of statistical distance (called *adversarial divergence*) that captures a wide range of GAN objectives frequently used in practice. They show that for objectives falling in the category, successfully optimizing the objectives implies weak convergence of the output distribution to the target. Instead of treating it in a black-box manner, in our work we focus on the optimization process of specific GAN instances and show that the output distribution converges to the target.

A more recent work related to ours is [LLDD19]. The authors prove that Wasserstein GANs can be trained via SGD to learn one layer neural networks. They assume that the activation function has a separable form, i.e., $\phi(\mathbf{x}) = (q(\mathbf{x}_1), \dots, q(\mathbf{x}_d))$, for some univariate function $q : \mathbb{R} \rightarrow \mathbb{R}$ that has a simple form, e.g. is a Lipschitz and odd function. Our result instead focuses on the standard GAN and shows that SGD converges for a much broader class of activation functions including non-invertible ones like ReLUs.

The interplay between min-max dynamics and GAN dynamics is already a very active field of research. An interesting recent work, that focuses on the negative side of min-max games is [FVGP19]. The authors there show that for a general class of non-convex, non-concave zero sum games Stochastic Gradient Descent Ascent may not converge to fixed points that are meaningful within game theoretical settings.

On the positive side, in [AGL⁺17], the authors studied the existence of pure equilibrium under various min-max game formulations of the Generator/Discriminator training dynamics. In [RLLY18], a class of non-convex concave optimization problem is studied, where the minimizer’s objective function is weakly convex and the maximizer’s objective is strongly concave. Moreover, in [LJJ19], the performance of the Gradient Ascent Descent (GDA) under similar setting is studied. In [NK17], Nagarajan et al. studied GAN’s stability around the local Nash equilibrium of the min-max game. Finally, in [DP18b, DP18a] the authors use optimism to show convergence of gradient based methods in min-max optimization.

2 Preliminaries and Notation

We use small bold letters \mathbf{x} to refer to real vectors in \mathbb{R}^d and capital bold letters \mathbf{A} to refer to matrices in $\mathbb{R}^{d \times \ell}$. We define $\mathbf{1}\{\mathbf{x} \in S\} = \mathbf{1}_S(\mathbf{x})$ to be the 0 – 1 indicator of a set. The *Frobenius norm* of a matrix \mathbf{A} is defined as $\|\mathbf{A}\|_F = \sqrt{\sum_{ij} \mathbf{A}_{ij}^2}$.

For distributions Q, P we denote by $d_{TV}(Q, P) = (1/2) \int |Q(x) - P(x)| dx$ their statistical or total variation distance. Let $\mathcal{N}(\mathbf{W})$ be the normal distribution with mean $\mathbf{0} \in \mathbb{R}^d$ and covariance matrix $\mathbf{W}\mathbf{W}^T \in \mathbb{R}^{d \times d}$, with the following probability density function

$$\mathcal{N}(\mathbf{x}; \mathbf{W}) = \frac{1}{\sqrt{\det(2\pi(\mathbf{W}\mathbf{W}^T))}} \exp\left(-\frac{1}{2}\mathbf{x}^T(\mathbf{W}\mathbf{W}^T)^{-1}\mathbf{x}\right). \tag{2}$$

Also, let $\mathcal{N}(S; \mathbf{W})$ denote the *probability mass of a measurable set S* under this Gaussian measure. We shall also denote by \mathcal{N} the standard Gaussian; whether it is single or multidimensional will be clear from the context.

In a multi-variable function, we often want to focus on just a subset of variables. We then use semi-colon to separate the primary variables from the secondary. Usually, the secondary variables are treated as constants which parametrize the function.

3 Technical Overview

3.1 Discriminator Design

We study GAN Dynamics for learning one layer non-linear Neural Networks. In particular we consider the following GAN architecture where the Generator is a one layer neural net with a fully connected linear layer parametrized by some matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$ followed by some general non-linear activation function $\phi: \mathbb{R}^d \mapsto \mathbb{R}^d$. Furthermore, we will use $\mathbf{W}_* \in \mathbb{R}^{d \times d}$ to denote the parameters of the target Generator network. If we denote the density functions of Generator and the target distributions as p_g and p_d respectively, it is known [GPAM⁺14] that the optimal Discriminator for this problem is $D_*(\mathbf{x}; p_g) = \frac{p_d(\mathbf{x})}{p_d(\mathbf{x}) + p_g(\mathbf{x})}$. Denote $\Sigma_* = \mathbf{W}_* \mathbf{W}_*^T$. When the activation function ϕ is invertible over its whole domain \mathbb{R}^d , this optimal Discriminator takes the following form

$$\begin{aligned} D_*(\mathbf{x}; \mathbf{W}) &= \frac{\mathcal{N}(\phi^{-1}(\mathbf{x}); \Sigma_*^{1/2})}{\mathcal{N}(\phi^{-1}(\mathbf{x}); \Sigma_*^{1/2}) + \mathcal{N}(\phi^{-1}(\mathbf{x}); \mathbf{W})} \\ &= \sigma(\phi^{-1}(\mathbf{x})^T \mathbf{A}_* \phi^{-1}(\mathbf{x}) + b_*), \end{aligned}$$

where $\mathbf{A}_* = \frac{1}{2}((\mathbf{W}\mathbf{W}^T)^{-1} - \Sigma_*^{-1})$, $b_* = \log \det(\mathbf{W}\Sigma_*^{-1/2})$, σ is the sigmoid function $\sigma(x) = \frac{1}{1 + \exp(-x)}$.

Unfortunately, many popular activation functions are not invertible over their whole domain but usually only on a subset of \mathbb{R}^d . For example the well-known ReLU activation i.e., $\text{ReLU}(\mathbf{x}) = \mathbb{1}\{\mathbf{x} \geq 0\}\mathbf{x}$ is invertible only when every coordinate of \mathbf{x} is positive. In order to capture these important activation functions, we relax the invertibility assumption to hold only on a subset of \mathbb{R}^d (see Definition 1). Recall that we denote by $p(\mathbf{W}, \phi)$ the output distribution of the network. In this general setting, the optimal Discriminator may not have a simple form. Even for ReLU (which is simply the identity function restricted on the set T) the optimal Discriminator is a complicated piecewise function consisting of 2^d different cases (these cases correspond to all possible subsets of coordinates that may be negative):

$$\begin{aligned} D(\mathbf{x}) &= \frac{\int_{\mathbf{z} \in \mathcal{Y}_{\mathbf{x}}} \mathcal{N}(\mathbf{z}; \mathbf{W}_*) d\mathbf{z}}{\int_{\mathbf{z} \in \mathcal{Y}_{\mathbf{x}}} \mathcal{N}(\mathbf{z}; \mathbf{W}) d\mathbf{z} + \int_{\mathbf{z} \in \mathcal{Y}_{\mathbf{x}}} \mathcal{N}(\mathbf{z}; \mathbf{W}_*) d\mathbf{z}}, \\ \text{where } \mathcal{Y}_{\mathbf{x}} &:= \{\mathbf{z} \text{ such that } \mathbf{x}_i = \text{ReLU}(\mathbf{z}_i)\} \end{aligned}$$

The optimal Discriminator is therefore a very complicated neural network and implementing such a network is infeasible even in rather low-dimensional scenarios. We take advantage of the fact that the activation function is known and invertible on some subset $T \in \mathbb{R}^d$ and design the following Discriminator architecture that balances simplicity and expressiveness. We denote by S the image of T under ϕ , i.e., $S = \phi(T)$ and define

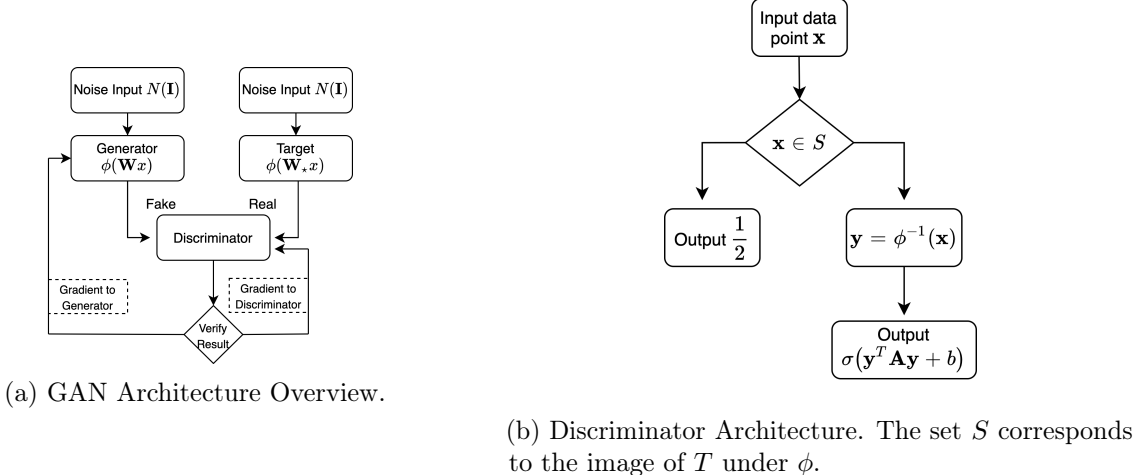


Figure 1: Our GAN Architecture.

$$D(\mathbf{x}; \mathbf{A}, b) = \mathbb{1}_S(\mathbf{x}) \sigma(\phi^{-1}(\mathbf{x})^T \mathbf{A} \phi^{-1}(\mathbf{x}) + b) + \frac{\mathbb{1}_{S^c}(\mathbf{x})}{2},$$

where σ is the sigmoid function (see Figure 1b). Then, the one layer Generator is paired with the above Discriminator to form the full architecture shown in Figure 1a. We then show that Nested Stochastic Gradient Descent Ascent (Algorithm 1) on this pair of neural nets provably converges, enabling Generator to recover the target distribution.

Algorithm 1 Nested Stochastic Gradient Descent Ascent on Standard GAN

- 1: Set $k = \tilde{O}(d^2/\epsilon^2)$
 - 2: Initialize $\mathbf{W} = \mathbf{I}$.
 - 3: Sample $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k)}$ from $p(\mathbf{W}_*, \phi)$
 - 4: **for** $i = 1$ to $\tilde{O}(d^2/\epsilon^4)$ **do**
 - 5: \triangleright Discriminator Training
 - 6: Sample $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k)}$ from $p(\mathbf{W}, \phi)$
 - 7: **for** $j = 1$ to k **do**
 - 8: Update \mathbf{A} and b with stochastic gradient
$$\nabla_{\mathbf{A}, b} [\log(D(\mathbf{x}^{(j)}; \mathbf{A}, b)) + \log(1 - D(\mathbf{y}^{(j)}; \mathbf{A}, b))]$$
 - 9: **end for**
 - 10: \triangleright Generator Training
 - 11: Sample $\mathbf{z} \sim \mathcal{N}(\mathbf{I})$
 - 12: Update \mathbf{W} with stochastic gradient
$$\nabla_{\mathbf{W}} \log(D(\phi(\mathbf{W}\mathbf{z}); \mathbf{A}, b)) \tag{3}$$
 - 13: **end for**
-

We show that this Projected Nested Stochastic Gradient Descent-Ascent (NSGDA) algorithm converges.

3.2 Roadmap of the Proof

The zero sum game used in the GAN formulation corresponds to the min-max optimization problem with loss function

$$\mathcal{L}(\mathbf{W}, \mathbf{A}, b) = \mathbf{E}_{\mathbf{x} \sim p(\mathbf{W}_*, \phi)} \log(D(\mathbf{x}; \mathbf{A}, b)) + \mathbf{E}_{\mathbf{x} \sim p(\mathbf{W}, \phi)} \log(1 - D(\mathbf{x}; \mathbf{A}, b)) \quad (4)$$

We use Nested Stochastic Gradient Descent-Ascent to solve the problem $\min_{\mathbf{W}} \max_{\mathbf{A}, b} \mathcal{L}(\mathbf{W}, \mathbf{A}, b)$. The Nested SGDA solves this problem by trying to fully optimize the inner maximization optimization for a given Generator parameter \mathbf{W} . In other words, in the inner loop of Algorithm 1, the Discriminator player is maximizing over \mathbf{A}, b the objective function $\mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W}) = \mathcal{L}(\mathbf{W}, \mathbf{A}, b)$; we stress that for $\mathcal{L}_{\mathcal{D}}$ the weight matrix \mathbf{W} is a fixed parameter.

We first show that by doing Stochastic Gradient Ascend we can train Discriminator’s parameters to being almost optimal. Using the structure of Discriminator we are able to show that $\mathcal{L}_{\mathcal{D}}$ is *strongly concave* with respect to the Discriminator parameters \mathbf{A}, b , see Lemma 2. Since the Discriminator is using samples from the underlying model $p(\mathbf{W}_*, \phi)$, from a learning theoretic point of view, we want to make its optimization as efficient as possible in order to get tight sample complexity results. Strong concavity is crucial in that sense: we are able to depend optimally not only on the dimension d but also on ϵ ; simple concavity would give us a substantially sub-optimal dependence on ϵ . The full discussion and detailed versions of the corresponding lemmas can be found in Subsection 4.2.

Showing convergence of Generator is more involved. With Discriminator’s parameters \mathbf{A}, b fixed, in expectation, Generator in Algorithm 1 receives training gradients from the objective function

$$\mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b) = \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} \log(1 - D(\phi(\mathbf{W}\mathbf{x}); \mathbf{A}, b)). \quad (5)$$

By Danskin’s Theorem [Dan12], when \mathbf{A}, b are fully optimized ($\mathbf{A} = \mathbf{A}_*, b = b_*$), the training gradients in expectation will be equal to the gradient of the function $\mathcal{V}(\mathbf{W}) = \max_{\mathbf{A}, b} \mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W})$, which is known as the *Virtual Training Criteria* of Generator in the work of [GPAM⁺14]. In contrast with the Discriminator objective function, minimizing $\mathcal{V}(\mathbf{W})$ is a *non-convex* minimization problem. In fact, any factorization of the covariance matrix $\Sigma_* = \bar{\mathbf{W}}\bar{\mathbf{W}}^T$ corresponds to a minimizer of this problem: the Gaussian distribution is invariant under orthogonal transformations, and therefore these matrices are indeed indistinguishable since they all produce the same distribution. Our main structural result shows that *finding approximate stationary points* of the virtual training criteria $\mathcal{V}(\mathbf{W})$ is sufficient to recover a matrix \mathbf{W} whose corresponding distribution $p(\mathbf{W}, \phi)$ is close in total variation distance to the true underlying distribution $p(\mathbf{W}^*, \phi)$. The proof of this statement relies on Gaussian anti-concentration of polynomials, see Lemma 3. At a high level, we first argue that the norm of gradient of the Generator is proportional to the probability that a specific quadratic form takes large values with respect to the standard normal distribution. Then using anti-concentration we show that this probability cannot be too small unless the distributions $p(\mathbf{W}, \phi)$ and $p(\mathbf{W}^*, \phi)$ are close. For the formal statement of the above discussion see Lemma 4.

A final complication that we face is that with finitely many samples, it is impossible to recover the optimal Discriminator parameters \mathbf{A}_*, b_* exactly. This introduces *biases* in the gradients used to train Generator. To overcome the difficulty, we use a Biased PSGD lemma, which guarantees convergence of SGD to first-order stationary points of the underlying objective function even when some bias are added to the gradient oracle used (See Lemma 7). In particular, the framework requires the bias to be bounded. We control the bias by showing that the training gradients $\nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}, \mathbf{A}, b)$ are Lipchitz continuous with respect to the Discriminator parameters \mathbf{A}, b (see Lemma 5). Thus, as long as we train the Discriminator enough to ensure that \mathbf{A}, b are close to the optimal \mathbf{A}_*, b_* , the bias $\|\nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b) - \nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W})\|$ will be small.

4 Convergence of GANs

In this section, we prove our main result and show that the GAN iteration converges and learns the one-layer Generator network (\mathbf{W}_*, ϕ) . Denote $\Sigma_* = \mathbf{W}_* \mathbf{W}_*^T$. Without loss of generality, we can assume the underlying target network has the form $(\Sigma_*^{1/2}, \phi)$ as we have already seen that this does not affect the corresponding distribution. The *Generator* is a one-layer neural network of the form (\mathbf{W}, ϕ) . The Generator will be paired with the *Discriminator* that tries to discern samples from $p(\Sigma_*^{1/2}, \phi)$ and $p(\mathbf{W}, \phi)$.

If the Generator’s parameter \mathbf{W} is initialized very far from the target distribution, most of its samples will fall outside the truncation set S of the Discriminator, leading to “vanishing gradients”. We thus make a closeness assumption that our initialization is close to the true covariance matrix. Assuming without loss of generality that we initialize the generator with $\mathbf{W} = \mathbf{I}$, we require that:

Assumption 1 (Initialization). *We assume that the matrix Σ_* satisfies*

$$\max(\|\Sigma_*^{-1} - \mathbf{I}\|_F, \|\Sigma_* - \mathbf{I}\|_F) \leq c.$$

Remark 2. *As shown in Corollary 3 of [DGTZ18b], we can initialize the algorithm with the empirical covariance matrix computed using $O_\alpha(d^2)$ samples from the truncated normal distribution $\mathcal{N}(\Sigma_*^{1/2}, T)$ and then transform the space so that $\mathbf{W}^{(0)} \rightarrow \mathbf{I}$. Then constant c in Assumption 1 depends only on the mass of the set α , i.e., $c = \text{poly}(1/\alpha) = O(1)$ under our assumption that $\alpha = \Omega(1)$.*

4.1 Projection Set

In order to avoid moving towards regions where the gradients vanish we will use the following convex projection set for the Generator parameters \mathbf{W} .

$$\mathcal{Q}_G = \left\{ \|\mathbf{W} - \mathbf{I}\|_F \leq \text{poly}(c), \text{poly}(1/c) \leq \mathbf{x}^T \mathbf{W} \mathbf{x} \leq \text{poly}(c), \text{ for all } \|\mathbf{x}\|_2 = 1 \right\}, \quad (6)$$

The important property of the above projection set is that the set T (the set where ϕ is invertible) has non-trivial mass under any matrix $\mathbf{W} \in \mathcal{Q}_G$. Interpreting the set T as a truncation set and using tools developed in [DGTZ18b], we can show that the set T always has non-trivial mass with respect to the Gaussian distribution $\mathcal{N}(\mathbf{W})$. This is a crucial property because our Discriminator relies on seeing samples that fall inside the set S (recall that S is the image of T under ϕ). In order for the Discriminator to produce non-trivial gradients we need to ensure that the mass of the invertible set T is not-trivial with respect to the parameter of the Generator.

Lemma 1 (Non-trivial mass). *Under Assumption 1, if we have $\mathbf{W} \in \mathcal{Q}_G$, it holds that $\mathcal{N}(T; \mathbf{W}) = \Omega_c(1)$.*

As we discussed previously, to obtain the optimal sample complexity we require the loss function of the Discriminator to be strongly concave. Unfortunately, strong concavity does not hold globally for the objective function \mathcal{L} . Hence, we shall define the following projection set for Discriminator’s parameters that ensures (see Lemma 2) this desired property.

$$\mathcal{Q}_D = \left\{ \|\mathbf{A}\|_F \leq \text{poly}(c), |b| \leq \text{poly}(c) \right\}. \quad (7)$$

We remark both sets \mathcal{Q}_G and \mathcal{Q}_D are convex and their projections can be efficiently computed, see, for example, Algorithm 3 in [DGTZ18b].

4.2 Training the Discriminator

In this section, we show convergence property of Discriminator training given in the following proposition.

Proposition 3 (Convergence of Discriminator Training). *Fix $\mathbf{W} \in \mathcal{Q}_G$ and assume that Assumption 1 is true. Setting the inner loop for $M_{\mathcal{D}} = \tilde{O}(d^2/\epsilon^2 \log^2(1/\delta))$. Then, with probability at least $1 - \delta$, it holds when Algorithm 1 exits the inner loop, the parameters \mathbf{A}, b satisfy $\|\mathbf{A} - \mathbf{A}_*\|_F + |b - b_*| \leq \epsilon$, where $\mathbf{A}_* = \frac{1}{2}((\mathbf{W}\mathbf{W}^T)^{-1} - \boldsymbol{\Sigma}_*^{-1})$, $b_* = \log \det(\mathbf{W}\boldsymbol{\Sigma}_*^{-1/2})$.*

The main step of the proof is the following Lemma which shows that the loss function \mathcal{L} as specified in Equation (4) is strongly concave with respect to \mathbf{A}, b . Its proof relies on the anti-concentration of polynomials under the Gaussian measure (Lemma 3) and can be found in the Appendix C.1.

Lemma 2 (Strong concavity for Discriminator). *Fix the Target Network (\mathbf{W}_*, ϕ) and Generator Network (\mathbf{W}, ϕ) , it holds that $\mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W})$ is at least $\Omega_c(1)$ -strongly concave when $\mathbf{A}, b \in \mathcal{Q}_D$ described in Equation (7), $\mathbf{W} \in \mathcal{Q}_G$ described in Equation (6) and Assumption 1 is satisfied.*

4.3 Training the Generator

As we discussed previously, the Generator tries to optimize the loss function

$$\mathcal{L}_G(\mathbf{W}; \mathbf{A}, b) = \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \log(1 - D(\phi(\mathbf{W}\mathbf{x}); \mathbf{A}, b)) \quad (8)$$

where $D(\mathbf{x}; \mathbf{A}, b) = \mathbb{1}_S(\mathbf{x}) / (1 + \exp(-\mathbf{x}^T \mathbf{A} \mathbf{x} - b)) + \mathbb{1}_{S^c}(\mathbf{x})/2$. By Danskin's Theorem [Dan12], when we use the optimal Discriminator parameters, namely $\mathbf{A}_* = \frac{1}{2}((\mathbf{W}\mathbf{W}^T)^{-1} - \boldsymbol{\Sigma}_*^{-1})$, $b_* = \log \det \mathbf{W} - \log \det \boldsymbol{\Sigma}_*^{1/2}$, we essentially optimize the function

$$\begin{aligned} \mathcal{V}(\mathbf{W}) &= \max_{\mathbf{A}, b} \mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W}) \\ &= \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{W})} [\log(1 - D(\phi(\mathbf{x}); \mathbf{A}_*, b_*))] + \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\boldsymbol{\Sigma}_*^{1/2})} [\log(D(\phi(\mathbf{x}); \mathbf{A}_*, b_*))] \end{aligned} \quad (9)$$

When the Discriminator is not fully optimized, the training gradients can still be treated as *biased* estimators of the true gradients of $\mathcal{V}(\mathbf{W})$. We first ignore the bias introduced from the sub-optimal Discriminator and prove our main structural result, showing that finding any stationary point of the Virtual Training Criteria $\mathcal{V}(\mathbf{W})$ suffices to learn the underlying distribution. Since $\mathcal{V}(\cdot)$ is not convex and we have a projection set, there are many obstacles in optimizing this objective function. Firstly, we need to make sure that stationary points in the interior of \mathcal{Q}_G are close to being optimal. Secondly, we need to make sure that the projection set does not introduce new "bad" stationary points (that is matrices \mathbf{W} whose corresponding distribution $p(\mathbf{W}, \phi)$ is far from $p(\boldsymbol{\Sigma}_*, \phi)$ lying on the boundary. To do so, we will employ the anti-concentration property of polynomials under Gaussian measure, which is stated in the following lemma.

Lemma 3 (Theorem 8 of [CW01]). *Let $k, \gamma \in \mathbb{R}_+$, $\mathbf{m} \in \mathbb{R}^d$, $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ such that $\boldsymbol{\Sigma}$ is positive semidefinite and $p : \mathbb{R}^d \rightarrow \mathbb{R}$ be a multivariate polynomial of degree at most ℓ , we define $\bar{\mathcal{Q}} = \{\mathbf{x} \in \mathbb{R}^d \mid |p(\mathbf{x})| \leq \gamma\}$, then there exists an absolute constant C such that*

$$\mathcal{N}(\bar{\mathcal{Q}}; \boldsymbol{\Sigma}^{1/2}) \leq \frac{Ck\gamma^{1/\ell}}{\left(\mathbf{E}_{\mathbf{z} \sim \mathcal{N}(\boldsymbol{\Sigma}^{1/2})} [|p(\mathbf{z})|^{k/\ell}]\right)^{1/k}}.$$

We are now ready to show the optimality of stationary points with respect to the learning problem.

Definition 2. A point $\mathbf{w} \in \mathcal{Q}$ is an ϵ -approximate first order stationary point of the function $f : \mathbb{R}^d \mapsto \mathbb{R}$ (ϵ -FOSP) if and only if for all $\mathbf{u} \in \mathcal{Q}$ the following holds

$$\frac{1}{\|\mathbf{w} - \mathbf{u}\|_2} \langle \nabla_{\mathbf{w}} f(\mathbf{w}), \mathbf{w} - \mathbf{u} \rangle \leq \epsilon.$$

Lemma 4 (Stationary Points Suffice). Let \mathbf{W} be an ϵ -first order stationary point (ϵ -FOSP) of $\mathcal{V}(\mathbf{W})$ in \mathcal{Q}_G . Then it holds $d_{\text{TV}}(p(\mathbf{W}, \phi), p(\mathbf{W}_*, \phi)) \leq O_c(\epsilon)$.

Proof Sketch. Here we only deal with the case when \mathbf{W} is an interior point of \mathcal{Q}_G . The rest of the proof which considers the case when \mathbf{W} lies on the boundary of \mathcal{Q}_G and can be found in Appendix D.1. For convenience, we define the expressions

$$h(\mathbf{x}; \mathbf{W}) = \frac{1}{2} \mathbf{x}^T \left((\mathbf{W}\mathbf{W}^T)^{-1} - \Sigma_*^{-1} \right) \mathbf{x} + \log \det \mathbf{W} \Sigma_*^{-1/2} \quad (10)$$

$$f(y) = \log(1 + \exp(y)) \quad (11)$$

Then, the gradient of the Virtual Training Criteria $\mathcal{V}(\mathbf{W})$ is given by

$$\nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W}) = \left(\Sigma_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right) \cdot \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) \mathbf{x} \mathbf{x}^T \mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} \right] \quad (12)$$

Given two matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{d \times d}$ where \mathbf{B} is a symmetric positive definite matrix, we always have $\|\mathbf{A} \mathbf{B}\|_F \geq \|\mathbf{A}\|_F \min_{\|z\|_2=1} z^T \mathbf{B} z$. Hence, the frobenius norm of the gradient can be lower bounded by

$$\|\nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W})\|_F \geq \left\| \left(\Sigma_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right) \right\|_F \cdot \min_{\|z\|_2=1} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (z^T \mathbf{x})^2 \right]$$

We now bound from below $\min_{\|z\|_2=1} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (z^T \mathbf{x})^2 \right]$. Notice that $f'(y) = \sigma(y) = 1/(\exp(-y) + 1)$ is the sigmoid function. Using the property that $f'(\cdot)$ is positive, and non-decreasing and also that $\mathbf{x} \mathbf{x}^T$ is positive semi-definite, we get the following inequality

$$\begin{aligned} & \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (z^T \mathbf{x})^2 \right] \\ & \geq f'(r) \gamma \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} [\mathbb{1}\{h(\mathbf{W}\mathbf{x}; \mathbf{W}) \geq r\} \mathbb{1}\{(z^T \mathbf{x})^2 \geq \gamma\} \cdot \mathbb{1}\{\mathbf{W}\mathbf{x} \in T\}]. \end{aligned}$$

Since we know that $\mathbf{W} \in \mathcal{Q}_G$, we can prove (see Appendix for details) that $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |h(\mathbf{W}\mathbf{x}; \mathbf{W})| \leq \text{poly}(c)$. Thus, if we choose $r < 0$, by Markov's inequality, we have

$$\mathcal{N}(h(\mathbf{W}\mathbf{x}; \mathbf{W}) \leq r) \leq \mathcal{N}(|h(\mathbf{W}\mathbf{x}; \mathbf{W})| \geq |r|) \leq \frac{1}{|r|} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |h(\mathbf{W}\mathbf{x}; \mathbf{W})|$$

By Lemma 1, the mass $\mathcal{N}(T; \mathbf{W})$ is always lower bounded by some absolute constant k_c that depends only on c . Hence, by setting $r = -\frac{4}{k_c} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |h(\mathbf{W}\mathbf{x}; \mathbf{W})|$, we have $\mathcal{N}(h(\mathbf{W}\mathbf{x}; \mathbf{W}) \leq r) \leq k_c/4$. On the other hand, we have $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[(z^T \mathbf{x})^2 \right] = 1$ given that $\|z\|_2 = 1$.

Now we can use the Gaussian anti-concentration of polynomials, Lemma 3, for the degree 2 polynomial $(z^T \mathbf{x})^2$. We choose

$$\gamma = \frac{1}{2} \left(\frac{k_c}{8C} \right)^2 \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[(z^T \mathbf{x})^2 \right],$$

and therefore, we have $\mathcal{N}\left(\left(\mathbf{z}^T \mathbf{x}\right)^2 \leq \gamma\right) \leq k_c/4$. Thus, by union bound, we conclude

$$\begin{aligned} & \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{h(\mathbf{W}\mathbf{x}; \mathbf{W}) \geq r\} \mathbb{1}\{(\mathbf{z}^T \mathbf{x})^2 \geq \gamma\} \mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} \right] \\ & \geq k_c - k_c/4 - k_c/4 \geq k_c/2. \end{aligned}$$

Using the inequality $f'(y) \geq e^y/2$ when $y < 0$, we obtain the bound

$$\min_{\|\mathbf{z}\|_2=1} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (\mathbf{z}^T \mathbf{x})^2 \right] \geq \Omega_c(1)$$

Therefore, given $\|\nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W})\|_F \leq \epsilon$, it holds

$$\left\| \Sigma_*^{-1/2} (\mathbf{W}\mathbf{W}^T) \Sigma_*^{-1/2} - \mathbf{I} \right\|_F \leq \left\| \left(\Sigma_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right) \right\|_F c^2 \leq O_c(\epsilon)$$

Using Pinsker's inequality (and the exact expression of Kullback-Leibler divergence for normal distributions) we have

$$d_{\text{TV}}(\mathcal{N}(\mathbf{W}), \mathcal{N}(\Sigma_*^{1/2})) \leq \left\| \Sigma_*^{-1/2} (\mathbf{W}\mathbf{W}^T) \Sigma_*^{-1/2} - \mathbf{I} \right\|_F \leq O_c(\epsilon).$$

Using the data processing inequality it follows that the total variation distance between the transformed distributions $p(\mathbf{W}, \phi)$, $p(\mathbf{W}_*, \phi)$ is small, i.e., $d_{\text{TV}}(p(\mathbf{W}, \phi), p(\mathbf{W}_*, \phi)) \leq O_c(\epsilon)$. \square

We have seen that finding stationary points of the non-convex objective suffices in order to compute a good parameter matrix \mathbf{W} . However, as we have already discussed we cannot optimize the Discriminator exactly and this leads to biased gradients when we train the Generator, that is gradients that do not exactly match the stochastic gradients of the function $\mathcal{V}(\mathbf{W})$. We now show how to overcome this obstacle. If we compute the gradient given in Equation (3), we get

$$\mathbf{g}_{\mathcal{G}} = D(\phi(\mathbf{W}\mathbf{z}); \mathbf{A}, b) \mathbf{A}\mathbf{W}\mathbf{z}\mathbf{z}^T, \quad (13)$$

where $\mathbf{z} \sim \mathcal{N}(\mathbf{I})$. In the following lemma, we show that the bias can be controlled as long as the parameters of the Discriminator are approximately optimal. In particular, we prove that the gradients $\nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W})$ (namely the gradient oracle $\mathbf{g}_{\mathcal{G}}$ in expectation) are Lipschitz with respect to Discriminator's parameters.

Lemma 5. $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{I})} [\mathbf{g}_{\mathcal{G}}] = \nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b)$ is $O_c(1)$ -Lipchitz with respect to \mathbf{A} and b when $\mathbf{W} \in \mathcal{Q}_G$, $\mathbf{A}, b \in \mathcal{Q}_D$ and Assumption 1 is satisfied.

Apart from that, we also need that the variance of the gradient oracle is bounded. We show the following lemma (see Appendix D.3).

Lemma 6. Let $\mathbf{g}_{\mathcal{G}} = D(\phi(\mathbf{W}\mathbf{z}); \mathbf{A}, b) \mathbf{A}\mathbf{W}\mathbf{z}\mathbf{z}^T$ be the gradient update of the Generator network and assume that $\mathbf{A}, b \in \mathcal{Q}_D$ described in Equation (7), $\mathbf{W} \in \mathcal{Q}_G$ described in Equation (6), and that Assumption 1 is satisfied. Then it holds that $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{I})} \left[\|\mathbf{g}_{\mathcal{G}}\|_2^2 \right] \leq O_c(d^2)$

Finally, we prove the Biased SGD Lemma which shows that the properties guaranteed by Lemmas 5 and 6 are essentially enough for us to optimize the Virtual Training Criteria. Technically, its proof is similar to the work of [GLZ16]; we adapt it so that it handle biased gradients (see Appendix D.4).

Lemma 7 (Biased Nonconvex PSGD). *Let f be an l -Lipschitz and L -smooth function, such that $\max_{\mathbf{x}, \mathbf{y} \in \mathcal{Q}} \|f(\mathbf{x}) - f(\mathbf{y})\|_2 \leq R$ on a convex domain \mathcal{Q} . At step t of the SGD we are given a biased gradient $\boldsymbol{\xi}^{(t)}$ such that $\left\| \mathbf{E}[\boldsymbol{\xi}^{(t)} | \boldsymbol{\xi}^{(1)}, \dots, \boldsymbol{\xi}^{(t-1)}] - \nabla f(\mathbf{x}) \right\|_2 \leq \alpha$ and $\mathbf{E}[\|\boldsymbol{\xi}^{(t)}\|_2^2] \leq B$. Set $M = O(BLR/\epsilon^4)$ and sample the stopping time m uniformly at random from $\{1, \dots, M\}$. Then, with step size $\beta = \sqrt{2R/(LBM)}$ and the update rule $\mathbf{w}^{(t+1)} = \operatorname{argmin}_{\mathbf{w} \in \mathcal{Q}} \left\| \mathbf{w} - (\mathbf{w}^{(t)} - \beta \boldsymbol{\xi}^{(t)}) \right\|_2$, we have that with probability at least 99%, the last iteration $\mathbf{w}^{(m)}$ of PSGD is an $O(\epsilon + \sqrt{l\alpha})$ -stationary point of f .*

Finally, we will combine the lemmas together with the Biased-SGD framework to obtain the sample complexity and number of iterations needed of Algorithm 1.

4.3.1 Proof of Theorem 1

Using Proposition 3, if we run the inner loop for $M_{\mathcal{D}} = O_c(d^2/\epsilon^2 \log^2(1/\delta))$ iterations, with probability at least $1 - \delta$, when Algorithm 1 exits the inner loop, the parameters \mathbf{A}, b satisfy $\|\mathbf{A} - \mathbf{A}_*\|_F + |b - b_*| \leq O_c(\epsilon)$. Using Lemma 5, we know $\nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}, \mathbf{A}, b)$ is $\tau := \Omega_c(1)$ -Lipchitz with respect to Discriminator's parameters \mathbf{A}, b . Furthermore, $\nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}, \mathbf{A}_*, b_*) = \nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W})$. Hence, it holds

$$\|\nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b) - \nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W})\|_F \leq \tau (\|\mathbf{A} - \mathbf{A}_*\|_F + |b - b_*|) \leq O_c(\epsilon).$$

This implies that the gradient oracle $\mathbf{g}_{\mathcal{G}}$ used by Algorithm 1 satisfies

$$\left\| \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} [\mathbf{g}_{\mathcal{G}}(\mathbf{x}; \mathbf{W}, \mathbf{A}, b)] - \nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W}) \right\|_F \leq O_c(\epsilon). \quad (14)$$

We have that the virtual objective function $\mathcal{V}(\mathbf{W})$ is $L := O_c(1)$ -smooth and $l := O_c(1)$ -Lipchitz continuous (see Appendix for a proof). Moreover, using Lemma 6, we have that the variance of the gradient oracle, namely $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} [\|\mathbf{g}_{\mathcal{G}}(\mathbf{x})\|_2^2]$, is bounded by $B = O_c(d^2)$. By the definition of the Projection Set \mathcal{Q}_G in Equation (6) and the fact that $\mathcal{V}(\mathbf{W})$ is $O_c(1)$ -Lipchitz continuous, it holds

$$R := \max_{\mathbf{W}_1, \mathbf{W}_2 \in \mathcal{Q}_G} |\mathcal{V}(\mathbf{W}_1) - \mathcal{V}(\mathbf{W}_2)| \leq O_c(1).$$

Conditioning on the event that the guarantee in Equation (14) is met, by Lemma 7, if we run the outer loop of Algorithm 1 for $M_{\mathcal{G}} = O_c(LBR/\epsilon^4) = \tilde{O}_c(d^2/\epsilon^4)$ rounds with step size $\eta_{\mathcal{G}} = \sqrt{\frac{LB}{RM_{\mathcal{G}}}} = O_c(\epsilon^2)$, it holds that the last iteration Generator parameters $\widetilde{\mathbf{W}}$ are an $O_c(\epsilon)$ - first order stationary point of $\mathcal{V}(\mathbf{W})$. If we set $\delta = \frac{1}{100M_{\mathcal{G}}}$, by the union bound, the probability that Equation (14) fails to hold in any iteration is less than 1%. Finally, by Lemma 4, we can transform the guarantee into bounds on the total variation distance between Generator distribution and target distribution and conclude that with probability at least 99% in the last iteration $d_{\text{TV}}(p(\widetilde{\mathbf{W}}, \phi), p(\boldsymbol{\Sigma}_*^{1/2}, \phi)) \leq \epsilon$.

References

- [AB17] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862*, 2017.
- [ACB17] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017.
- [AGL⁺17] Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, and Yi Zhang. Generalization and equilibrium in generative adversarial nets (gans). In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, page 224–232. JMLR.org, 2017.
- [CW01] Anthony Carbery and James Wright. Distributional and ℓ_1 norm inequalities for polynomials over convex bodies in \mathbb{R}^n . *Mathematical research letters*, 8(3):233–248, 2001.
- [Dan12] John M Danskin. *The theory of max-min and its application to weapons allocation problems*, volume 5. Springer Science & Business Media, 2012.
- [DGTZ18a] Constantinos Daskalakis, Themis Gouleakis, Christos Tzamos, and Manolis Zampetakis. Efficient statistics, in high dimensions, from truncated samples. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 639–649. IEEE, 2018.
- [DGTZ18b] Constantinos Daskalakis, Themis Gouleakis, Christos Tzamos, and Manolis Zampetakis. Efficient statistics, in high dimensions, from truncated samples. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 639–649, 2018.
- [DP18a] Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. *arXiv preprint arXiv:1807.04252*, 2018.
- [DP18b] Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 9256–9266, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [FFGT17] Soheil Feizi, Farzan Farnia, Tony Ginart, and David Tse. Understanding gans: the lqg setting. *arXiv preprint arXiv:1710.10793*, 2017.
- [FVGP19] Lampros Flokas, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Georgios Piliouras. Poincaré recurrence, cycles and spurious equilibria in gradient-descent-ascent for non-convex non-concave zero-sum games, 2019.
- [GLZ16] Saeed Ghadimi, Guanghui Lan, and Hongchao Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2):267–305, 2016.
- [GM18] Ian Gemp and Sridhar Mahadevan. Global convergence to the equilibrium of gans using variational inequalities. *arXiv preprint arXiv:1808.01531*, 2018.

- [GPAM⁺14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [HLPR19] Nicholas JA Harvey, Christopher Liaw, Yaniv Plan, and Sikander Randhawa. Tight analyses for non-smooth stochastic gradient descent. In *Conference on Learning Theory*, pages 1579–1613. PMLR, 2019.
- [HLR19] Nicholas JA Harvey, Christopher Liaw, and Sikander Randhawa. Simple and optimal high-probability bounds for strongly-convex stochastic gradient descent. *arXiv preprint arXiv:1909.00843*, 2019.
- [HTP⁺17] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. *CoRR*, abs/1711.03213, 2017.
- [JWS⁺20] Yongjun Jing, Hao Wang, Kun Shao, Xing Huo, and Yangyang Zhang. Unsupervised graph representation learning with variable heat kernel. *IEEE Access*, 8:15800–15811, 2020.
- [KAHK17] Naveen Kodali, Jacob D. Abernethy, James Hays, and Zolt Kira. How to train your DRAGAN. *CoRR*, abs/1705.07215, 2017.
- [LBC17] Shuang Liu, Olivier Bousquet, and Kamalika Chaudhuri. Approximation and convergence properties of generative adversarial learning. In *Advances in Neural Information Processing Systems*, pages 5545–5553, 2017.
- [LJJ19] Tianyi Lin, Chi Jin, and Michael I Jordan. On gradient descent ascent for nonconvex-concave minimax problems. *arXiv preprint arXiv:1906.00331*, 2019.
- [LLDD19] Qi Lei, Jason D. Lee, Alexandros G. Dimakis, and Constantinos Daskalakis. Sgd learns one-layer networks in wgans, 2019.
- [MGN18] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? *arXiv preprint arXiv:1801.04406*, 2018.
- [MPPS16] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *CoRR*, abs/1611.02163, 2016.
- [Nes13] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
- [NK17] Vaishnavh Nagarajan and J. Zico Kolter. Gradient descent gan optimization is locally stable. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5585–5595. Curran Associates, Inc., 2017.
- [O’D14] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [RLLY18] Hassan Rafique, Mingrui Liu, Qihang Lin, and Tianbao Yang. Non-convex min-max optimization: Provable algorithms and applications in machine learning. *arXiv preprint arXiv:1810.02060*, 2018.

- [RMC15] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [ZXY18] Zizhao Zhang, Yuanpu Xie, and Lin Yang. Photographic text-to-image synthesis with a hierarchically-nested adversarial network. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6199–6208. IEEE Computer Society, 2018.

A Additional Notation

In this section we define some additional notation used in the following sections of the appendix. We denote by $\mathbf{A} \otimes \mathbf{B}$ the Kronecker product between two matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{B} \in \mathbb{R}^{k \times \ell}$, is the block matrix

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & \cdots & a_{1n}\mathbf{B} \\ \vdots & \ddots & \vdots \\ a_{m1}\mathbf{B} & \cdots & a_{mn}\mathbf{B} \end{bmatrix} \in \mathbb{R}^{mk \times n\ell}$$

We also define the symmetrization of a matrix $\mathcal{S}(\mathbf{A}) = \mathbf{A} + \mathbf{A}^T$.

B Projection Set

Recall that the projection sets for Generator and Discriminator are given by

$$\mathcal{Q}_G = \left\{ \|\mathbf{W} - \mathbf{I}\|_F \leq \text{poly}(c), \text{poly}(1/c) \leq \mathbf{x}^T \mathbf{W} \mathbf{x} \leq \text{poly}(c), \text{ for all } \|\mathbf{x}\|_2 = 1 \right\},$$

$$\mathcal{Q}_D = \left\{ \|\mathbf{A}\|_F \leq \text{poly}(c), |b| \leq \text{poly}(c) \right\}.$$

The following lemma shows that the above projection sets always contain valid solutions of the problem. Moreover, when the parameters of the Discriminator and the Generator lie inside their corresponding projection sets we have the following bounds that will be useful throughout our analysis.

Lemma 8 (Projection Sets). *Under Assumption 1, we have that the convex set \mathcal{Q}_G contains some matrix \mathbf{W} such that the corresponding distribution $p(\mathbf{W}, \phi)$ is equal to the true underlying distribution $p(\mathbf{W}_*, \phi)$. Moreover, for any $\mathbf{W} \in \mathcal{Q}_G$ we have that the optimal parameters for the Discriminator, $\mathbf{A}_* = \frac{1}{2} \left((\mathbf{W}\mathbf{W}^T)^{-1} - \Sigma_*^{-1} \right)$, $b_* = \log \det \left(\mathbf{W}\Sigma_*^{-1/2} \right)$ lie in the Discriminator projection set \mathcal{Q}_D . Finally, for all $\mathbf{W} \in \mathcal{Q}_G$ we have that the following bounds hold $\|\mathbf{W}\mathbf{W}^T - \Sigma_*\|_F, \left\| (\mathbf{W}\mathbf{W}^T)^{-1} - \Sigma_*^{-1} \right\|_F, \|\mathbf{W}^T \Sigma_*^{-1} \mathbf{W} - \mathbf{I}\|_F, \left\| \Sigma_*^{-1/2} \mathbf{W}\mathbf{W}^T \Sigma_*^{-1/2} - \mathbf{I} \right\|_F, \left\| \Sigma_*^{1/2} (\mathbf{W}\mathbf{W}^T)^{-1} \Sigma_*^{1/2} - \mathbf{I} \right\|_F, \left| \log \det \left(\mathbf{W}\Sigma_*^{-1/2} \right) \right| \leq \text{poly}(c)$*

Proof. First, we consider the the projection set \mathcal{Q}_G . By Assumption 1, we have $\|\Sigma_* - \mathbf{I}\|_F \leq c$. Assume that we have the following eigenvalue decomposition $\Sigma_* = \mathbf{U}\Lambda^2\mathbf{U}^T$. Then the inequality assumed can be rewritten as $\|\mathbf{U}(\Lambda^2 - \mathbf{I})\mathbf{U}^T\|_F \leq c$. Since the Frobenious norm is invariant under unitary transformations, this gives $\|\Lambda^2 - \mathbf{I}\|_F = \|(\Lambda - \mathbf{I})(\Lambda + \mathbf{I})\|_F \leq c$. As $\Lambda + \mathbf{I}$ clearly has its eigenvalues lower bounded by 1, it implies $\left\| \Sigma_*^{1/2} - \mathbf{I} \right\|_F = \|\mathbf{U}(\Lambda - \mathbf{I})\mathbf{U}^T\|_F = \|\Lambda - \mathbf{I}\|_F \leq c$.

On the other hand, since $\left\| \boldsymbol{\Sigma}_*^{1/2} - \mathbf{I} \right\|_2 \leq \left\| \boldsymbol{\Sigma}_*^{1/2} - \mathbf{I} \right\|_F \leq c$, it holds $\left\| \boldsymbol{\Sigma}_*^{1/2} \right\|_2 \leq 1 + c$. Similarly, $\left\| \boldsymbol{\Sigma}_*^{-1/2} \right\|_2 \leq 1 + c$. Thus, the eigenvalues of $\boldsymbol{\Sigma}_*^{1/2}$ lie in the interval $[1/(1+c), 1+c]$. Hence, we have shown that the projection set \mathcal{Q}_G contains some matrix \mathbf{W}_* that is essentially optimal (up to orthogonal transformations). On the other hand, the six expressions in the statement are all poly(c)-Lipchitz with respect to \mathbf{W} as the l2-norms of \mathbf{W} , \mathbf{W}^{-1} , $\boldsymbol{\Sigma}_*^{1/2}$, $\boldsymbol{\Sigma}_*^{-1/2}$ are all bounded by poly(c). The upper bounds of these expressions then follow from their Lipchitzness, the diameters of the projection set (poly(c)) and the fact that they all evaluate to 0 when $\mathbf{W} = \boldsymbol{\Sigma}_*^{1/2}$.

Next, we consider the Discriminator projection set. Recall that after fixing the Generator parameters \mathbf{W} , the optimal Discriminator parameters are given by $\mathbf{A}_* = \frac{1}{2} \left((\mathbf{W}\mathbf{W}^T)^{-1} - \boldsymbol{\Sigma}_*^{-1} \right)$, $b_* = \log \det \left(\mathbf{W}\boldsymbol{\Sigma}_*^{-1/2} \right)$. From our discussion of Generator Projection Set, we know both expressions are bounded by poly(c) when $\mathbf{W} \in \mathcal{Q}_G$. Hence, for any $\mathbf{W} \in \mathcal{Q}_G$ we have that the corresponding optimal Discriminator parameters lie in the projection set \mathcal{Q}_D . \square

B.1 Proof of Lemma 1

We will use the following lemma from the work of [DGTZ18b] which relates the probability mass that two different normal distributions assign to the same set.

Lemma 9 (Lemma 7 of [DGTZ18b]). *Consider two normal distributions $\mathcal{N}(\boldsymbol{\Sigma}_1^{1/2})$, $\mathcal{N}(\boldsymbol{\Sigma}_2^{1/2})$ and a set S satisfying $\mathcal{N}(S; \boldsymbol{\Sigma}_1^{1/2}) \geq \alpha$. Suppose the parameters satisfy $\left\| \boldsymbol{\Sigma}_1^{1/2} \boldsymbol{\Sigma}_2^{-1} \boldsymbol{\Sigma}_1^{1/2} - \mathbf{I} \right\|_F \leq B$. Then, it holds $\mathcal{N}(S; \boldsymbol{\Sigma}_2^{1/2}) \geq k_{B,\alpha}$ for some constant $k_{B,\alpha}$ that depends only on B and α .*

From Lemma 8, we know that $\left\| \boldsymbol{\Sigma}_*^{1/2} (\mathbf{W}\mathbf{W}^T)^{-1} \boldsymbol{\Sigma}_*^{1/2} - \mathbf{I} \right\|_F \leq \text{poly}(c)$. Recall that by Definition 1, we have $\mathcal{N}(T; \boldsymbol{\Sigma}_*^{1/2}) \geq \alpha = \Omega(1)$. From Lemma 9 it follows that $\mathcal{N}(T; \mathbf{W}) \geq \Omega_c(1)$.

C Training the Discriminator

For convenience, we define the following expressions which commonly appear in the formulas of the training gradients. Let $h(\mathbf{x}; \mathbf{A}, b) = \mathbf{x}^T \mathbf{A} \mathbf{x} + b$, $h(\mathbf{x}; \mathbf{W}) = \frac{1}{2} \mathbf{x}^T \left((\mathbf{W}\mathbf{W}^T)^{-1} - \boldsymbol{\Sigma}_*^{-1} \right) \mathbf{x} + \log \det \mathbf{W} \boldsymbol{\Sigma}_*^{-1/2}$, $\sigma(y) = 1/(1 + \exp(-y))$, $f(y) = \log(1 + \exp(y))$. Notice that the second expression is equivalent to the first expression when \mathbf{A} , b are exactly the optimal Discriminator parameters.

C.1 Proof of Lemma 2

We will use the following facts.

Fact 4. *For any symmetric matrix $\mathbf{X} \in \mathbb{R}^{d \times d}$, there exist two semidefinite matrices \mathbf{Y} and \mathbf{Z} such that $\mathbf{X} = \mathbf{Y} + \mathbf{Z}$, $\|\mathbf{Y}\|_F + \|\mathbf{Z}\|_F \leq \sqrt{2} \|\mathbf{X}\|_F$ and $\|\mathbf{Y}\|_2 + \|\mathbf{Z}\|_2 \leq 2 \|\mathbf{X}\|_2$.*

Proof. Since \mathbf{X} is symmetric, we can always diagonalize it as $\mathbf{X} = \mathbf{Q} \boldsymbol{\Lambda} \mathbf{Q}^T$. Then, we rewrite $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}^+ + \boldsymbol{\Lambda}^-$ where $\boldsymbol{\Lambda}^+$ contains only positive diagonal elements and $\boldsymbol{\Lambda}^-$ contains only negative diagonal elements. Set $\mathbf{Y} = \mathbf{Q} \boldsymbol{\Lambda}^+ \mathbf{Q}^T$ and $\mathbf{Z} = \mathbf{Q} \boldsymbol{\Lambda}^- \mathbf{Q}^T$. For ℓ_2 norm, as all the eigenvalues of \mathbf{Y} and \mathbf{Z} comes from the eigenvalues of \mathbf{X} , the inequality is obvious. Then, for the Frobenius norm, it is easy to see that $\|\mathbf{Y}\|_F^2 + \|\mathbf{Z}\|_F^2 = \|\mathbf{X}\|_F^2$. Since $(\|\mathbf{Y}\|_F + \|\mathbf{Z}\|_F)^2 \leq 2 \|\mathbf{Y}\|_F^2 + 2 \|\mathbf{Z}\|_F^2$, the fact follows. \square

Fact 5. For any matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ it holds $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |\mathbf{x}^T \mathbf{A} \mathbf{x}| \leq \sqrt{2} \|\mathbf{A}\|_F$

Proof. Recall that we defined the symmetrization of matrix \mathbf{A} as $\mathcal{S}(\mathbf{A}) = \mathbf{A} + \mathbf{A}^T$. We first replace \mathbf{A} with $\frac{1}{2}\mathcal{S}(\mathbf{A})$. Then, use Fact 4 to rewrite $\mathcal{S}(\mathbf{A})$ as the sum of two definite matrices \mathbf{A}_1 and \mathbf{A}_2 , where $\|\mathbf{A}_1\|_F + \|\mathbf{A}_2\|_F \leq \sqrt{2} \|\mathbf{A}\|_F$. Then, it holds

$$\begin{aligned} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |\mathbf{x}^T \mathbf{A} \mathbf{x}| &= \frac{1}{2} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |\mathbf{x}^T \mathcal{S}(\mathbf{A}) \mathbf{x}| \leq \frac{1}{2} \left(\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \mathbf{x}^T \mathbf{A}_1 \mathbf{x} + \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \mathbf{x}^T (-\mathbf{A}_2) \mathbf{x} \right) \\ &\leq \frac{1}{2} (\|\mathbf{A}_1\|_F + \|\mathbf{A}_2\|_F) \leq \sqrt{2} \|\mathbf{A}\|_F \end{aligned}$$

□

We first compute the first and second order derivatives of Discriminator's objective function $\mathcal{L}_{\mathcal{D}}$ with respect to \mathbf{A}, b . Given that the activation function ϕ is invertible on the set T , recall that the Discriminator objective function is given by

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(\mathbf{A}, b, \mathbf{W}) &= \mathbf{E}_{\mathbf{x} \sim p(\boldsymbol{\Sigma}_*^{1/2}, \phi)} \log(D(\mathbf{x}; \mathbf{A}, b)) + \mathbf{E}_{\mathbf{x} \sim p(\mathbf{W}, \phi)} \log(1 - D(\mathbf{x}; \mathbf{A}, b)) \\ &= \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\boldsymbol{\Sigma}_*^{1/2})} \log(D(\phi(\mathbf{x}); \mathbf{A}, b)) + \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{W})} \log(1 - D(\phi(\mathbf{x}); \mathbf{A}, b)), \end{aligned}$$

where $D(\mathbf{x}; \mathbf{A}, b) = \mathbb{1}_S(\mathbf{x}) \sigma(\phi^{-1}(\mathbf{x})^T \mathbf{A} \phi^{-1}(\mathbf{x}) + b) + \mathbb{1}_{S^c}(\mathbf{x})/2$ and $S = \phi(T)$.

Applying the chain rule then gives us

$$\begin{aligned} \nabla_{\mathbf{A}, b} \mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W}) &= - \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{W})} \left[f'(h(\mathbf{x}; \mathbf{A}, b)) \begin{pmatrix} (\mathbf{x} \mathbf{x}^T)^b \\ 1 \end{pmatrix} \mathbb{1}\{\mathbf{x} \in T\} \right] \\ &\quad - \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\boldsymbol{\Sigma}_*^{1/2})} \left[f'(h(\mathbf{x}; \mathbf{A}, b)) \begin{pmatrix} (\mathbf{x} \mathbf{x}^T)^b \\ 1 \end{pmatrix} \mathbb{1}\{\mathbf{x} \in T\} \right] \\ &\quad + \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\boldsymbol{\Sigma}_*^{1/2})} \left[\begin{pmatrix} (\mathbf{x} \mathbf{x}^T)^b \\ 1 \end{pmatrix} \mathbb{1}\{\mathbf{x} \in T\} \right] \end{aligned} \quad (15)$$

$$\begin{aligned} \nabla_{\mathbf{A}, b}^2 \mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W}) &= - \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f''(h(\mathbf{W} \mathbf{x}; \mathbf{A}, b)) q(\mathbf{W} \mathbf{x}) \mathbb{1}\{\mathbf{W} \mathbf{x} \in T\} \right] \\ &\quad - \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f''(h(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}; \mathbf{A}, b)) q(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}) \mathbb{1}\{\boldsymbol{\Sigma}_*^{1/2} \mathbf{x} \in T\} \right], \end{aligned} \quad (16)$$

where for convenience we denote $q(\mathbf{x}) = \begin{pmatrix} (\mathbf{x} \mathbf{x}^T)^b \\ 1 \end{pmatrix} \otimes \begin{pmatrix} (\mathbf{x} \mathbf{x}^T)^b \\ 1 \end{pmatrix}$.

As the two terms of the Hessian above are similar, we will show how to handle the second term containing $\boldsymbol{\Sigma}_*^{1/2}$. The other case follows similarly. More specifically, for any $\mathbf{z} \in \mathbb{R}^{d^2}$ such that $\|\mathbf{z}\|_2 = 1$, we will show that

$$\mathbf{z}^T \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f''(h(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}; \mathbf{A}, b)) q(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}) \mathbb{1}\{\boldsymbol{\Sigma}_*^{1/2} \mathbf{x} \in T\} \right] \mathbf{z}$$

is bounded from below by some constant $\Omega_c(1)$ (that depends only on c). First, notice that $f''(y) = e^y / (e^y + 1)^2$ is a positive, even function and is strictly decreasing when $y > 0$. Thus, we must have $f''(h(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}; \mathbf{A}, b)) \geq f''(r)$ when $|h(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}; \mathbf{A}, b)| \leq r$. Besides, as $q(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x})$ gives rise to a positive semi-definite matrix, we always have $\mathbf{z}^T q(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}) \mathbf{z} = |\mathbf{z}^T q(\boldsymbol{\Sigma}_*^{1/2} \mathbf{x}) \mathbf{z}|$. Hence,

we proceed by defining two sets P and Q where the values of $f''(\cdot)$ and $q(\cdot)$ are lower bounded respectively. Let $P = \{\mathbf{x} \in \mathbb{R}^d : |h(\mathbf{\Sigma}_*^{1/2}\mathbf{x}; \mathbf{A}, b)| \leq r\}$ and $Q = \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{z}^T q(\mathbf{\Sigma}_*^{1/2}\mathbf{x})\mathbf{z}| \geq \gamma\}$. It then holds

$$\begin{aligned} & \mathbf{z}^T \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f''(h(\mathbf{\Sigma}_*^{1/2}\mathbf{x}; \mathbf{A}, b)) q(\mathbf{\Sigma}_*^{1/2}\mathbf{x}) \mathbb{1}\{\mathbf{\Sigma}_*^{1/2}\mathbf{x} \in T\} \right] \mathbf{z} \\ & \geq f''(r) \gamma \mathbf{E}_{\mathbf{x} \in \mathcal{N}} \left[\mathbb{1}\{\mathbf{x} \in Q\} \mathbb{1}\{\mathbf{x} \in P\} \mathbb{1}\{\mathbf{\Sigma}_*^{1/2}\mathbf{x} \in T\} \right] \end{aligned}$$

Then, we lower bound the mass of each set. For set Q , we can use the Gaussian anti-concentration of polynomials, Lemma 3, for the degree 4 polynomial $\mathbf{z}^T q(\mathbf{\Sigma}_*^{1/2}\mathbf{x})\mathbf{z}$ with respect to \mathbf{x} . We choose

$$\gamma = \alpha^4 \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\left| \mathbf{z}^T q(\mathbf{\Sigma}_*^{1/2}\mathbf{x})\mathbf{z} \right| \right] / (16C)^4,$$

where $\alpha := \mathcal{N}(T; \mathbf{\Sigma}_*^{1/2}) = \Omega(1)$ as defined in Definition 1 and $C = O(1)$ is the absolute constant defined in Lemma 3. It then holds $\mathcal{N}(Q; \mathbf{I}) = Pr_{\mathbf{x} \sim \mathcal{N}} \left[\mathbf{z}^T q(\mathbf{\Sigma}_*^{1/2}\mathbf{x})\mathbf{z} \leq \gamma \right] \leq \frac{\alpha}{4}$. Moreover, for this specific choice of γ , it holds

$$\begin{aligned} \gamma &= \frac{\alpha^4}{(16C)^4} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbf{z}^T \begin{pmatrix} (\mathbf{\Sigma}_*^{1/2}\mathbf{x}\mathbf{x}^T\mathbf{\Sigma}_*^{1/2})^b \\ 1 \end{pmatrix} \otimes \begin{pmatrix} (\mathbf{\Sigma}_*^{1/2}\mathbf{x}\mathbf{x}^T\mathbf{\Sigma}_*^{1/2})^b \\ 1 \end{pmatrix} \mathbf{z} \right] \\ &\geq \lambda_{\min}(\mathbf{\Sigma}_*)^2 \frac{\alpha^4}{(16C)^4} \geq \Omega_c(1), \end{aligned} \tag{17}$$

since $\|\mathbf{\Sigma}_*^{-1}\|_2 \leq c + 1$ by Assumption 1.

Next, we will use Markov's Inequality to lower bound $Pr_{\mathbf{x} \sim \mathcal{N}} \left[|h(\mathbf{\Sigma}_*^{1/2}\mathbf{x}; \mathbf{A}, b)| > r \right]$. We will first derive an upper bound for the expected value $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[|h(\mathbf{\Sigma}_*^{1/2}\mathbf{x}; \mathbf{A}, b)| \right]$. In particular, by the definition of Discriminator's projection set in Equation (7) ($\|\mathbf{A}\|_F \leq \text{poly}(c)$) and the constraint $\|\mathbf{\Sigma}_*\|_2 \leq 1 + c$ as implied in Assumption 1, it follows from Fact 5 that

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[|h(\mathbf{\Sigma}_*^{1/2}\mathbf{x}; \mathbf{A}, b)| \right] \leq \sqrt{2} \left\| \mathbf{\Sigma}_*^{1/2} \mathbf{A} \mathbf{\Sigma}_*^{1/2} \right\|_F + |b| \leq \text{poly}(c).$$

By Markov's inequality, we have $Pr_{\mathbf{x} \sim \mathcal{N}} \left[|h(\mathbf{\Sigma}_*^{1/2}\mathbf{x}; \mathbf{A}, b)| > r \right] < \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[|h(\mathbf{\Sigma}_*^{1/2}\mathbf{x}; \mathbf{A}, b)| \right] \cdot \frac{1}{r}$. By setting $r = \Omega_c(1/\alpha)$, we obtain $\mathcal{N}(|h(\mathbf{x}; \mathbf{A}, b)| > r) < \frac{\alpha}{4}$. Using the union bound, we then have

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} (\mathbb{1}\{\mathbf{x} \in P\} \mathbb{1}\{\mathbf{x} \in Q\} \mathbb{1}\{\mathbf{\Sigma}_*^{1/2}\mathbf{x} \in T\}) \geq \alpha - \mathcal{N}(\bar{P}) - \mathcal{N}(\bar{Q}) \geq \alpha/2.$$

Overall, we get

$$\mathbf{z}^T \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f''(h(\mathbf{W}\mathbf{x}; \mathbf{A}, b)) q(\mathbf{W}\mathbf{x}) \mathbb{1}\{\mathbf{\Sigma}_*^{1/2}\mathbf{x} \in T\} \right] \mathbf{z} \geq \frac{\alpha}{2} f''(r) \gamma.$$

Recall that we have set $r = \Omega_c(1/\alpha)$, $\gamma = \Omega_c(1)$ as shown in Equation (17). Since $f''(r) = \frac{e^r}{(e^r+1)^2} \geq e^{-r}/4$ and α is an absolute constant, we conclude $\mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W})$ is at least $\Omega_c(1)$ strongly concave.

C.2 High Probability Projected Stochastic Gradient Descent

The main tool used is Theorem 3 in the work of [HLR19], which gives tight convergence rates of Projected Stochastic Gradient Descent in the high probability regime. In their work, the theorem is

proved for gradient oracle with sub-gaussian noise but it is not hard to see that the statement holds for sub-exponential noise as well. We state the more general version of the theorem and provide its proof sketch.

To be consistent with the notation used in the original proof, here we will use subscript to denote iterations rather than index of elements from vector/matrix. Also, since we won't use the notion of "invertible region" in the section, we will use T to denote the total number of iterations just in this section.

Lemma 10 (High probability Projected Stochastic Gradient Descent, [HLR19], Theorem C.12). *Let \mathcal{Q} be a convex set and $f : \mathcal{Q} \mapsto \mathbb{R}$ be μ -strongly convex and L -Lipchitz function with minimizer \mathbf{x}_* . Moreover, let \mathbf{g}_t be an unbiased stochastic gradient oracle of f . Define $\eta_t = \frac{2}{\mu(t+1)}$. Let $\mathbf{x}_t = \text{proj}_{\mathcal{Q}}(\mathbf{x}_{t-1} - \eta_t \mathbf{g}_t)$. Assume the following holds.*

- (a) *There exists a constant τ such that $\sum_{t=1}^T \|\mathbf{g}_t\|_2^2 \leq \tilde{O}(\tau T \log^2(1/\delta))$ with probability $1 - \delta$.*
- (b) *There exists a pair of constants κ and ζ such that for any $\lambda \in (0, 1/\zeta)$, we have $\mathbf{E}[\exp(\lambda \langle \mathbf{g}_t - \mathbf{E}[\mathbf{g}_t], \mathbf{x}_t - \mathbf{x}_* \rangle)] \leq \exp\left(\lambda^2 \kappa \|\mathbf{x}_t - \mathbf{x}_*\|_2^2\right)$, where the expectation is conditional on $\mathbf{g}_{t-1}, \dots, \mathbf{g}_1$.*

Let $\gamma_t = \frac{t}{T(T+1)/2}$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds

$$f\left(\sum_{t=1}^T \gamma_t \mathbf{x}_t\right) - f(\mathbf{x}_*) \leq \tilde{O}\left(\frac{(L^2 + \tau + \kappa + \zeta) \log^2(1/\delta)}{\mu} \frac{1}{T}\right).$$

We can follow the original proof of Theorem C.12 in [HLR19] until we reach the inequality

$$f\left(\sum_{t=1}^T \gamma_t \mathbf{x}_t\right) - f(\mathbf{x}_*) \leq \frac{2}{T(T+1)} \sum_{t=1}^T t \cdot \langle \mathbf{g}_t - \mathbf{E}[\mathbf{g}_t], \mathbf{x}_t - \mathbf{x}_* \rangle + \frac{2}{\mu T(T+1)} \sum_{t=1}^T \|\mathbf{g}_t\|_2^2. \quad (18)$$

The original proof exploits the properties of sub-gaussian noise to bound the two summation sequences respectively (Lemma C.4 and C.5 in the original work). We state the two supporting lemmas and show they still hold in our scenario.

Lemma 11 (Lemma C.4). *For any $\delta \in (0, 1)$, $\sum_{t=1}^T \|\mathbf{g}_t\|_2^2 = \tilde{O}(\tau \cdot T \cdot \log^2(1/\delta))$ with probability at least $1 - \delta$.*

Proof. The lemma is trivially true due to property (a). □

Lemma 12 (Lemma C.5). *Let $Z_T = \sum_{t=1}^T t \cdot \langle \mathbf{g}_t - \mathbf{E}[\mathbf{g}_t], \mathbf{x}_t - \mathbf{x}_* \rangle$. Then, for any $\delta \in (0, 1)$, we have $Z_T = \tilde{O}\left(\frac{L^2 + \tau + \kappa + \zeta}{\mu} \cdot T \cdot \log(1/\delta)\right)$ with probability at least $1 - \delta$.*

The proof of Lemma 12 relies on the following claims and lemma. We will follow the notation of the original proof and define $d_t := t \cdot \langle \mathbf{g}_t - \mathbf{E}[\mathbf{g}_t], \mathbf{x}_t - \mathbf{x}_* \rangle$, $v_{t-1} := 2 \cdot \kappa \cdot t^2 \cdot \|\mathbf{x}_t - \mathbf{x}_*\|_2^2$, and $V_T := \sum_{t=1}^T v_{t-1}$.

Claim 6 (Claim C.9). *For any $\lambda \in (0, 1/(\zeta \cdot T))$, $\mathbf{E}[\exp(\lambda \cdot t \cdot d_t)] \leq \exp(\lambda^2 v_{t-1}/2)$, where the expectation is conditional on $\mathbf{g}_{t-1}, \dots, \mathbf{g}_1$.*

Proof. The claim follows by scaling both sides of property (b) by a factor of $t \leq T$. □

Claim 7 (Lemma C.11). *There exists non-negative constants $\alpha_1, \dots, \alpha_T = O\left(\kappa \cdot \frac{T}{\mu}\right)$, and $\beta = \tilde{O}\left(\kappa \cdot (L^2 + \tau) \cdot \frac{T^2}{\mu^2}\right)$ such that for every $\delta \in (0, 1)$, $V_T \leq \sum_{t=1}^T \alpha_t d_t + \beta \log^2(1/\delta)$ with probability at least $1 - \delta$.*

Proof. We follow the original proof until we reach the inequality

$$V_T \leq \sum_{t=1}^T \alpha_t d_t + \kappa \cdot \sum_{t=1}^T O(T) \cdot \|\mathbf{g}_t\|_2^2 + O(\kappa L^2 / \mu^2),$$

where $\alpha_t = O(\kappa \cdot T / \mu)$. By property (a), we have $\sum_{t=1}^T \|\mathbf{g}_t\|_2^2 \leq \tilde{O}(\tau T \log^2(1/\delta))$ with probability $1 - \delta$. Hence, overall,

$$\beta := \frac{1}{\log^2(1/\delta)} \cdot \left(\kappa \cdot \sum_{t=1}^T O(T) \cdot \|\mathbf{g}_t\|_2^2 + O(\kappa L^2 / \mu^2) \right) \leq \tilde{O}\left(\kappa \cdot (L^2 + \tau) \cdot \frac{T^2}{\mu^2}\right).$$

□

Lemma 13 (Generalized Freedman, [HLPR19], Theorem 3.3). *Let $\{d_t, \mathcal{F}_t\}_{t=1}^T$ be a martingale difference sequence. Suppose that, for $t \in [T]$, v_{t-1} are non-negative \mathcal{F}_{t-1} -measurable random variables satisfying $\mathbf{E}[\exp(\lambda d_t) | \mathcal{F}_{t-1}] \leq \exp(\lambda^2 v_{t-1} / 2)$, for any $\lambda \in (0, 1/\zeta')$. Let $Z_T = \sum_{t=1}^T d_t$ and $V_T = \sum_{t=1}^T v_{t-1}$. Suppose there exists $\alpha_1, \dots, \alpha_T, \beta \geq 0$ such that for every $\delta \in (0, 1)$, $V_T \leq \sum_{t=1}^T \alpha_t d_t + \beta \log^2(1/\delta)$ with probability at least $1 - \delta$. Let $\alpha \geq \max_{t \in [T]} \alpha_t$. Then*

$$\Pr[Z_T \geq x] \leq \exp\left(-\frac{x^2}{4\alpha x + 2\zeta' x + 8\beta \log^2(1/\delta)}\right) + \delta.$$

Proof Sketch. The original theorem does not have the constrain on λ . Nevertheless, the proof only requires $\mathbf{E}[\exp(\lambda d_t) | \mathcal{F}_{t-1}] \leq \exp(\lambda^2 v_{t-1} / 2)$ to hold for λ upper bounded by some constant C . In fact, λ is fixed to be in the interval $[0, 1/(2\alpha)]$ at the beginning of the proof. Hence, if $\zeta' \leq 2\alpha$, the original proof holds. In general, we can fix $\lambda \leq \min(1/(2\alpha), 1/\zeta')$. Then, we can follow the original proof until we reach the inequality

$$\Pr\left[Z_T \geq x \text{ and } V_T \leq \sum_{i=1}^T \alpha_i d_i + \beta \log^2(1/\delta)\right] \leq \exp(-\lambda(x - 2\lambda\beta \log^2(1/\delta))).$$

Instead of picking $\lambda = 1/(2\alpha + 4\beta \cdot \log^2(1/\delta)/x) \leq \frac{1}{2\alpha}$, we pick

$$\lambda = 1/(2\alpha + \zeta' + 4\beta \cdot \log^2(1/\delta)/x) \leq \min\left(\frac{1}{2\alpha}, \frac{1}{\zeta'}\right).$$

For this specific choice of λ , we conclude

$$\Pr\left[Z_T \geq x \text{ and } V_T \leq \sum_{i=1}^T \alpha_i d_i + \beta \log^2(1/\delta)\right] \leq \exp\left(-\frac{x^2}{4\alpha x + 2\zeta' x + 8\beta \log^2(1/\delta)}\right).$$

Then, the result follows by applying union bounds on the events $\{Z_t \geq x\}$ and $\{V_T \leq \sum_{t=1}^T \alpha_i d_i + \beta \cdot \log^2(1/\delta)\}$. □

Proof of Lemma 12. Claim 6 shows that $\mathbf{E}[\exp(\lambda d_t) | \mathcal{F}_{t-1}] \leq \exp(\lambda^2 v_{t-1}/2)$, for all $\lambda \in (0, 1/(\zeta T))$. By Claim 7, we have for every $\delta \in (0, 1)$, $V_T \leq \sum_{t=1}^T \alpha_t d_t + \beta \log^2(1/\delta)$. Hence, we can plugin $\alpha := \max_t \alpha_t = O\left(\kappa \cdot \frac{T}{\mu}\right)$, $\beta := \tilde{O}\left(\kappa \cdot (L^2 + \tau) \cdot \frac{T^2}{\mu^2}\right)$ and $\zeta' := \zeta \cdot T$ into Lemma 13. Furthermore, we can set

$$\begin{aligned} x &= \tilde{\Theta}\left(\alpha + \zeta' + \sqrt{\beta \log^2(1/\delta)}\right) \\ &\leq \tilde{\Theta}\left(\kappa \cdot \frac{T}{\mu} + \zeta \cdot T + (\kappa + L^2 + \tau) \cdot \frac{T}{\mu} \log(1/\delta)\right) \\ &\leq \tilde{\Theta}\left((\kappa + L^2 + \tau + \zeta) \cdot \frac{T}{\mu} \log(1/\delta)\right), \end{aligned}$$

where in the first inequality we use the fact $\sqrt{ab+ac} \leq O(a+b+c)$ for all $a, b, c > 0$. This then gives $\Pr[Z_T \geq x] \leq 2 \cdot \delta$. We then get the statement in Lemma 12 by rescaling the failing probability. \square

Proof of Lemma 10. Substituting the bounds obtained from Lemmas 11 and 12 into Equation (18) then gives

$$\begin{aligned} f\left(\sum_{t=1}^t \gamma_t \mathbf{x}_t\right) - f(\mathbf{x}_*) &\leq \tilde{O}\left(\frac{L^2 + \tau + \kappa + \zeta}{T(T+1)} \cdot \frac{T}{\mu} \cdot \log(1/\delta)\right) + \tilde{O}(\tau \cdot T \cdot \log^2(1/\delta)) \\ &\leq \tilde{O}\left(\frac{(L^2 + \tau + \kappa + \zeta)}{\mu} \cdot \frac{\log^2(1/\delta)}{T}\right). \end{aligned}$$

\square

C.3 Proof of Proposition 3

We then proceed to show that our gradient oracle for Discriminator does satisfy conditions (a) and (b) in Lemma 10. We need the following well-known result on concentration of polynomials of independent Gaussian random variables. See, e.g., [O'D14].

Lemma 14 (Gaussian Hypercontractivity). *Let $h(\mathbf{x}) : \mathbb{R}^d \mapsto \mathbb{R}$ be a degree- m polynomial. Then,*

$$\Pr_{\mathbf{x} \sim \mathcal{N}} \left[\left| h(\mathbf{x}) - \mathbf{E}_{\mathbf{y} \sim \mathcal{N}}[h(\mathbf{y})] \right| \geq l \right] \leq e^2 \exp\left(-\left(\frac{l^2}{C \text{Var}_{\mathbf{x} \sim \mathcal{N}}[h(\mathbf{x})]}\right)^{1/m}\right),$$

where $C > 0$ is an absolute constant.

For our case, following Algorithm 1, the gradient oracle used by the Discriminator is given by

$$\mathbf{g}^{(t)} = \mathbb{1}\{\mathbf{x}^{(t)} \in T\} \frac{1}{1+e^k} \cdot \begin{bmatrix} \left(\mathbf{x}^{(t)} \mathbf{x}^{(t)T}\right)^b \\ 1 \end{bmatrix} + \mathbb{1}\{\mathbf{y}^{(t)} \in T\} \frac{1}{1+e^{-q}} \cdot \begin{bmatrix} \left(\mathbf{y}^{(t)} \mathbf{y}^{(t)T}\right)^b \\ 1 \end{bmatrix}, \quad (19)$$

where $k = \mathbf{x}^{(t)T} \mathbf{A}^{(t)} \mathbf{x}^{(t)} + b^{(t)}$, $q = \mathbf{y}^{(t)T} \mathbf{A} \mathbf{y}^{(t)} + b^{(t)}$, $\mathbf{x}^{(t)} \sim \mathcal{N}(\mathbf{W})$, $\mathbf{y}^{(t)} \sim \mathcal{N}(\boldsymbol{\Sigma}_*^{1/2})$. The next lemma shows that this specific gradient oracle satisfies condition (a) required by Lemma 10.

Lemma 15. *For any $\delta \in (0, 1)$, we have $\sum_{t=1}^{M_D} \|\mathbf{g}^{(t)}\|_2^2 = \tilde{O}_c(d^2 M_D \log^2(1/\delta))$ with probability at least $1 - \delta$.*

Proof. First, notice that the sum is upper bounded by

$$\sum_{t=1}^{M_{\mathcal{D}}} \left\| \mathbf{g}^{(t)} \right\|_2^2 \leq \sum_{t=1}^{M_{\mathcal{D}}} \left\| \mathbf{x}^{(t)} \mathbf{x}^{(t)T} \right\|_F^2 + \sum_{t=1}^{M_{\mathcal{D}}} \left\| \mathbf{y}^{(t)} \mathbf{y}^{(t)T} \right\|_F^2 + 2M_{\mathcal{D}}.$$

Next, we will upper bound the expectation and variance of $\left\| \mathbf{x}^{(t)} \mathbf{x}^{(t)T} \right\|_F^2$ respectively (the bounds for $\left\| \mathbf{y}^{(t)} \mathbf{y}^{(t)T} \right\|_F^2$ can be obtained similarly). Recall that $\mathbf{x}^{(t)} \sim \mathcal{N}(\mathbf{W})$. As $\|\mathbf{W}\|_2 \leq \text{poly}(c)$ by definition of the projection set, we can assume that the data are generated by a standard normal, i.e., $\mathbf{x} \sim \mathcal{N}(\mathbf{I})$ by only losing a $\text{poly}(c)$ factor in the upper bound. Hence, we have

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} \left[\|\mathbf{x}\mathbf{x}\|_F^2 \right] = \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} \left[\sum_{i,j} \mathbf{x}_i^2 \mathbf{x}_j^2 \right] \leq O(d^2),$$

where the last inequality follows from standard bounds of moments of normal variables. On the other hand,

$$\begin{aligned} \text{Var}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} \left[\|\mathbf{x}\mathbf{x}^T\|_F^2 \right] &= \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} \left[\|\mathbf{x}\mathbf{x}^T\|_F^4 \right] - \left(\mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} \left[\|\mathbf{x}\mathbf{x}^T\|_F^2 \right] \right)^2 \\ &\leq \mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} \left[\left(\sum_{i,j} \mathbf{x}_i^2 \mathbf{x}_j^2 \right)^2 \right] \leq O(d^4), \end{aligned}$$

where the last inequality comes from the fact that $\left(\sum_{i,j} \mathbf{x}_i^2 \mathbf{x}_j^2 \right)^2$ equals the sum of d^4 eighth moments of normal variables.

We can then apply Lemma 14 with $l = \Theta_c(d^2 \log^2(1/\delta'))$ on the degree-4 polynomial $\left\| \mathbf{x}^{(t)} \mathbf{x}^{(t)T} \right\|_F^2$ and obtain

$$\Pr \left[\left| \left\| \mathbf{x}^{(t)} \mathbf{x}^{(t)T} \right\|_F^2 - \mathbf{E} \left[\left\| \mathbf{x}^{(t)} \mathbf{x}^{(t)T} \right\|_F^2 \right] \right| \geq l \right] \leq e^2 \exp \left(- \left(\frac{l^2}{O_c(d^4)} \right)^{1/4} \right) \leq \delta'.$$

This implies with probability at least $1 - \delta'$,

$$\left\| \mathbf{g}^{(t)} \right\|_2^2 \leq \left\| \mathbf{x}^{(t)} \mathbf{x}^{(t)T} \right\|_F^2 \leq O_c(d^2 \log^2(1/\delta)). \quad (20)$$

If we choose $\delta' = \delta/M_{\mathcal{D}}$ and takes the union bound over the events that Equation (20) fails to hold for some $t \in [M_{\mathcal{D}}]$, we have

$$\sum_{t=1}^{M_{\mathcal{D}}} \left\| \mathbf{x}^{(t)} \mathbf{x}^{(t)T} \right\|_F^2 \leq O_c(d^2 T \log^2(1/\delta))$$

with probability at least $1 - \delta$. With a similar argument on $\sum_{t=1}^{M_{\mathcal{D}}} \left\| \mathbf{y}^{(t)} \mathbf{y}^{(t)T} \right\|_F^2$, the statement follows. \square

Define the noise of the gradient oracle as $\mathbf{z}^{(t)} = \mathbf{g}^{(t)} - \mathbf{E}[\mathbf{g}^{(t)}]$ conditioned on $\mathbf{g}_{t-1}, \dots, \mathbf{g}_1$. Denote $\boldsymbol{\theta}^{(t)} = \begin{bmatrix} (\mathbf{A}^{(t)})^b \\ b^{(t)} \end{bmatrix}$ and $\boldsymbol{\theta}_* = \begin{bmatrix} (\mathbf{A}_*)^b \\ b_* \end{bmatrix}$. The next lemma ensures that the noise satisfy condition (b) required by Lemma 10.

Lemma 16. For all $\lambda \leq \frac{1}{O_c(1)}$, $t \in [M_{\mathcal{D}}]$, it holds

$$\mathbf{E} \left[\exp \left(\lambda \cdot \langle \mathbf{z}^{(t)}, \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \rangle \right) \right] \leq \exp \left(\lambda^2 O_c(d) \left\| \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \right\|_2^2 \right)$$

Proof. By definition of $\mathbf{z}^{(t)}$, we have

$$\begin{aligned} \langle \mathbf{z}^{(t)}, \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \rangle &\leq \left| \langle \mathbf{g}^{(t)}, \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \rangle \right| + \left| \langle \mathbf{E} \left[\mathbf{g}^{(t)} \right], \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \rangle \right| \leq \left| \langle \mathbf{g}^{(t)}, \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \rangle \right| + O_c \left(\left\| \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \right\|_2 \right) \\ &\leq \left| \mathbf{x}^{(t)T} \left(\mathbf{A}^{(t)} - \mathbf{A}_* \right) \mathbf{x}^{(t)} \right| + \left| \mathbf{y}^{(t)T} \left(\mathbf{A}^{(t)} - \mathbf{A}_* \right) \mathbf{y}^{(t)} \right| + O_c \left(\left\| \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \right\|_2 \right) \\ &\leq O_c \left(\left\| \mathbf{x}^{(t)} \right\|_2^2 + \left\| \mathbf{y}^{(t)} \right\|_2^2 \right) \cdot \left\| \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \right\|_2 \\ &\leq O_c \left(\left\| \tilde{\mathbf{x}} \right\|_2^2 + \left\| \tilde{\mathbf{y}} \right\|_2^2 \right) \cdot \left\| \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \right\|_2, \end{aligned} \quad (21)$$

where $\tilde{\mathbf{x}} := \mathbf{W}^{-1} \mathbf{x}^{(t)}$, $\tilde{\mathbf{y}} := \boldsymbol{\Sigma}_*^{-1/2} \mathbf{y}^{(t)}$ and the last inequality follows from the bounded ℓ_2 norm of \mathbf{W}^{-1} and $\boldsymbol{\Sigma}_*^{-1/2}$. Notice that $\left\| \tilde{\mathbf{x}} \right\|_2^2, \left\| \tilde{\mathbf{y}} \right\|_2^2$ are the sum of $2d$ independent random variables following the chi-squared distribution. Hence,

$$\mathbf{E} \left[\exp \left(\lambda \left(\left\| \tilde{\mathbf{x}} \right\|_2^2 + \left\| \tilde{\mathbf{y}} \right\|_2^2 \right) \right) \right] \leq \exp(\lambda^2 O(d)) \quad (22)$$

for all $\lambda \leq 1/O(1)$. By Equation (21), we have for all $\lambda > 0$

$$\mathbf{E} \left[\exp \left(\lambda \cdot \langle \mathbf{z}^{(t)}, \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \rangle \right) \right] \leq \mathbf{E} \left[\exp \left(\lambda \cdot \left\| \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \right\|_2 \cdot O_c \left(\left\| \tilde{\mathbf{x}} \right\|_2^2 + \left\| \tilde{\mathbf{y}} \right\|_2^2 \right) \right) \right].$$

Recall that for any $\boldsymbol{\theta} \in \mathcal{Q}_D$, we have $\left\| \boldsymbol{\theta}_* - \boldsymbol{\theta} \right\|_2 \leq O_c(1)$. Hence, we can scale Equation (22) by a factor of $\left\| \boldsymbol{\theta} - \boldsymbol{\theta}_* \right\|_2^2$ and conclude for all $\lambda \in (0, 1/O_c(1))$

$$\mathbf{E} \left[\exp \left(\lambda \cdot \langle \mathbf{z}^{(t)}, \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \rangle \right) \right] \leq \exp \left(\lambda^2 O_c(d) \left\| \boldsymbol{\theta}_* - \boldsymbol{\theta}^{(t)} \right\|_2^2 \right).$$

□

Now, we can substitute the bounds obtained into Lemma 10 to finish the proof of Proposition 3. By Lemma 2, the objective function $\mathcal{L}_{\mathcal{D}}$ is $\mu = \Omega_c(1)$ strongly concave; by Lemma 15, we have $\tau = O_c(d^2)$; by Lemma 16, we have $\kappa = O_c(d)$ and $\zeta = O_c(1)$. Hence, fixing the Generator parameter $\mathbf{W} \in \mathcal{Q}_G$, if we run the inner loop for $M_{\mathcal{D}}$ iterations, with probability at least $1 - \delta$, when the algorithm exits the inner loop, the parameters \mathbf{A}, b satisfy

$$\left| \mathcal{L}_{\mathcal{D}}(\mathbf{A}_*, b_*; \mathbf{W}) - \mathcal{L}_{\mathcal{D}}(\mathbf{A}, b; \mathbf{W}) \right| \leq O_c \left(d^2 \log^2(1/\delta) / M_{\mathcal{D}} \right).$$

Again, by strong concavity, it holds $\left\| \mathbf{A} - \mathbf{A}_* \right\|_F + |b - b_*| \leq O_c \left(\sqrt{d^2 \log^2(1/\delta) / M_{\mathcal{D}}} \right)$. Setting $M_{\mathcal{D}} = O_c \left(d^2 \log^2(1/\delta) / \epsilon^2 \right)$, we then obtain the statement.

D Training the Generator

For convenience, we will use the same notations of $h(\mathbf{x}; \mathbf{A}, b)$, $h(\mathbf{x}; \mathbf{W})$ and $f(y)$ as used in Section C.

D.1 Proof of Lemma 4

The gradient of the Virtual Training Criteria $\mathcal{V}(\mathbf{W})$ is given as

$$\nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W}) = \left(\boldsymbol{\Sigma}_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right) \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) \mathbf{x} \mathbf{x}^T \mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} \right]. \quad (23)$$

Thus, the Frobenius Norm of the gradient can be lower bounded by

$$\|\nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W})\|_F \geq \left\| \left(\boldsymbol{\Sigma}_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right) \right\|_F \min_{\|\mathbf{z}\|_2=1} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (\mathbf{z}^T \mathbf{x})^2 \right].$$

We now try to lower bound $\min_{\|\mathbf{z}\|_2=1} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (\mathbf{z}^T \mathbf{x})^2 \right]$. Using the property that $f'(y) = \sigma(y)$ is positive and monotonically increasing, we get the following inequality.

$$\begin{aligned} & \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (\mathbf{z}^T \mathbf{x})^2 \right] \\ & \geq f'(r) \gamma \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{h(\mathbf{W}\mathbf{x}; \mathbf{W}) \geq r\} \mathbb{1}\{(\mathbf{z}^T \mathbf{x})^2 \geq \gamma\} \mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} \right]. \end{aligned}$$

By Fact 5 and Lemma 8, we can upper bound $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |h(\mathbf{W}\mathbf{x}; \mathbf{W})|$ by

$$\begin{aligned} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |h(\mathbf{W}\mathbf{x}; \mathbf{W})| & \leq \frac{1}{2} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} [|\mathbf{x}^T (\mathbf{I} - \mathbf{W}^T \boldsymbol{\Sigma}_*^{-1} \mathbf{W}) \mathbf{x}|] + \frac{1}{2} |\log \det (\mathbf{W}^T \boldsymbol{\Sigma}_*^{-1} \mathbf{W})| \\ & \leq \frac{1}{\sqrt{2}} \|\mathbf{I} - \mathbf{W}^T \boldsymbol{\Sigma}_*^{-1} \mathbf{W}\|_F + \left| \log \det (\mathbf{W}^T \boldsymbol{\Sigma}_*^{-1/2}) \right| \\ & \leq \text{poly}(c). \end{aligned}$$

Thus, if we choose $r < 0$, by Markov's Inequality, we have

$$\mathcal{N}(h(\mathbf{W}\mathbf{x}; \mathbf{W}) \leq r) \leq \mathcal{N}(|h(\mathbf{W}\mathbf{x}; \mathbf{W})| \geq |r|) \leq \frac{1}{|r|} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |h(\mathbf{W}\mathbf{x}; \mathbf{W})|.$$

By Lemma 1, the Generator's mass in the set T is always lower bounded by some absolute constant k_c that depends only on c . By setting $r = -\frac{4}{k_c} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} |h(\mathbf{W}\mathbf{x}; \mathbf{W})|$, we have $\mathcal{N}(h(\mathbf{W}\mathbf{x}; \mathbf{W}) \leq r) \leq k_c/4$. On the other hand, for the degree 2 polynomial $(\mathbf{z}^T \mathbf{x})^2$, we can again use the Gaussian anti-concentration of polynomials (Lemma 3). We choose

$$\gamma = \frac{1}{2} \left(\frac{k_c}{8C} \right)^2 \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} [(\mathbf{z}^T \mathbf{x})^2],$$

and therefore, we have $\mathcal{N}((\mathbf{z}^T \mathbf{x})^2 \leq \gamma) \leq k_c/4$. Thus, by Union Bound, we conclude

$$\begin{aligned} & \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbb{1}\{h(\mathbf{W}\mathbf{x}; \mathbf{W}) \geq r\} \mathbb{1}\{(\mathbf{z}^T \mathbf{x})^2 \geq \gamma\} \mathbb{1}\{\mathbf{W}\mathbf{x} \in T\} \right] \\ & \geq k_c - k_c/4 - k_c/4 \geq k_c/2. \end{aligned}$$

Using the inequality $f'(y) = \frac{1}{1+\exp(-y)} \geq e^y/2$ when $y < 0$, we obtain the bound

$$\min_{\|\mathbf{z}\|_2=1} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) (\mathbf{z}^T \mathbf{x})^2 \right] \geq \Omega_c(1).$$

Therefore, given $\|\nabla_{\mathbf{W}}\mathcal{V}(\mathbf{W})\|_F \leq \epsilon$, it holds

$$\left\| \boldsymbol{\Sigma}_*^{-1/2} (\mathbf{W}\mathbf{W}^T) \boldsymbol{\Sigma}_*^{-1/2} - \mathbf{I} \right\|_F \leq \left\| \left(\boldsymbol{\Sigma}_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right) \right\|_F c^2 \leq O_c(\epsilon).$$

Using Pinsker's inequality (and the exact expression of Kullback-Leibler divergence for normal distributions) we have

$$d_{\text{TV}}(\mathcal{N}(\mathbf{W}), \mathcal{N}(\boldsymbol{\Sigma}_*^{1/2})) \leq \left\| \boldsymbol{\Sigma}_*^{-1/2} (\mathbf{W}\mathbf{W}^T) \boldsymbol{\Sigma}_*^{-1/2} - \mathbf{I} \right\|_F \leq O_c(\epsilon).$$

Using the data processing inequality it follows that the total variation distance between the transformed distributions $p(\mathbf{W}, \phi)$, $p(\mathbf{W}_*, \phi)$ is small, i.e., $d_{\text{TV}}(p(\mathbf{W}, \phi), p(\mathbf{W}_*, \phi)) \leq O_c(\epsilon)$.

Next, we consider the case when \mathbf{W} lies on the boundary of \mathcal{Q}_G . For convenience, denote $\mathbf{X} = \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbf{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) \mathbf{x}\mathbf{x}^T \right]$. Then, the gradient can be written as

$$\nabla_{\mathbf{W}}\mathcal{V}(\mathbf{W}) = (\boldsymbol{\Sigma}_*^{-1} \mathbf{W} - \mathbf{W}^{-1}) \mathbf{X}. \quad (24)$$

Consider the Singular Value Decomposition $\mathbf{W}^T \boldsymbol{\Sigma}_*^{-1/2} = \mathbf{U} \boldsymbol{\Lambda} \mathbf{V}$. Since \mathbf{W} is a first order stationary point, it holds

$$\langle \nabla \mathcal{V}(\mathbf{W}), \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \rangle \leq \epsilon \left\| \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \right\|_F. \quad (25)$$

By expanding the inner product, we obtain the lower bound

$$\begin{aligned} \langle \nabla \mathcal{V}(\mathbf{W}), \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \rangle &= \text{Tr} \left(\mathbf{X} (\mathbf{W}^T \boldsymbol{\Sigma}_*^{-1} - \mathbf{W}^{-1}) (\mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T) \right) \\ &= \text{Tr} \left(\mathbf{X} (\mathbf{W}^T \boldsymbol{\Sigma}_*^{-1/2} - \mathbf{W}^{-1} \boldsymbol{\Sigma}_*^{1/2}) (\boldsymbol{\Sigma}_*^{-1/2} \mathbf{W} - \mathbf{V}^T \mathbf{U}^T) \right) \\ &= \text{Tr} (\mathbf{X} \mathbf{U} (\boldsymbol{\Lambda} - \boldsymbol{\Lambda}^{-1}) \mathbf{V} \mathbf{V}^T (\boldsymbol{\Lambda} - \mathbf{I}) \mathbf{U}^T) \\ &\geq \lambda_{\min}(\mathbf{X}) \text{Tr} (\boldsymbol{\Lambda}^2 - \boldsymbol{\Lambda} - \mathbf{I} + \boldsymbol{\Lambda}^{-1}). \end{aligned} \quad (26)$$

From the discussion of the previous case where \mathbf{W} is an interior point, we have $\lambda_{\min}(\mathbf{X}) \geq \Omega_c(1)$. Hence, combining Equation (26) with Equation (25) gives

$$\text{Tr} (\boldsymbol{\Lambda}^2 - \boldsymbol{\Lambda} - \mathbf{I} + \boldsymbol{\Lambda}^{-1}) \leq O_c(\epsilon) \left\| \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \right\|_F.$$

Notice that $\text{Tr} (\boldsymbol{\Lambda}^2 - \boldsymbol{\Lambda} - \mathbf{I} + \boldsymbol{\Lambda}^{-1})$ is 2-strongly convex with respect to $\boldsymbol{\Lambda}$ and minimizes at $\boldsymbol{\Lambda} = \mathbf{I}$. Hence, using convexity, we get

$$\|\boldsymbol{\Lambda} - \mathbf{I}\|_F^2 \leq O_c(\epsilon) \left\| \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \right\|_F. \quad (27)$$

On the other hand, it holds

$$\left\| \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \right\|_F = \left\| (\mathbf{W}^T \boldsymbol{\Sigma}_*^{-1/2} - \mathbf{U} \mathbf{V}) \boldsymbol{\Sigma}_*^{1/2} \right\|_F \leq \|\boldsymbol{\Lambda} - \mathbf{I}\|_F \left\| \boldsymbol{\Sigma}_*^{1/2} \right\|_2.$$

By Assumption 1, we have $\|\boldsymbol{\Sigma}_*\|_2 \leq (1+c)$. Hence,

$$\|\boldsymbol{\Lambda} - \mathbf{I}\|_F \geq \Omega_c(1) \left\| \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \right\|_F. \quad (28)$$

Combining Equation (27) and (28), we then have

$$\left\| \mathbf{W} - \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T \right\|_F \leq O_c(\epsilon).$$

This therefore implies

$$\left\| \boldsymbol{\Sigma}_*^{-1/2} (\mathbf{W} \mathbf{W}^T) \boldsymbol{\Sigma}_*^{-1/2} - \mathbf{I} \right\|_F \leq O_c(\epsilon)$$

since the expression is $O_c(1)$ -Lipchitz with respect to \mathbf{W} and the expressions evaluates to 0 when $\mathbf{W} = \boldsymbol{\Sigma}_*^{1/2} \mathbf{V}^T \mathbf{U}^T$. The rest of the proof is then identical to the case when \mathbf{W} is an interior point.

D.2 Proof of Lemma 5

The actual gradient used in training takes a similar form as Equation (23). The difference is that the Discriminator has now parameters \mathbf{A}, b instead of the optimal \mathbf{A}_*, b_* . In particular, the expected value of the training gradients are given by

$$\nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b) = 2 \mathbf{A} \mathbf{W} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f'(h(\mathbf{W} \mathbf{x}; \mathbf{A}, b)) \mathbb{1}\{\mathbf{W} \mathbf{x} \in T\} \mathbf{x} \mathbf{x}^T \right].$$

We proceed to compute the expression's derivatives with respect to \mathbf{A} and b . For \mathbf{A} , we have

$$\begin{aligned} \nabla_{\mathbf{A}} \nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b) &= 2 \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f'(h(\mathbf{W} \mathbf{x}; \mathbf{A}, b)) \mathbb{1}\{\mathbf{W} \mathbf{x} \in T\} (\mathbf{W} \mathbf{x} \mathbf{x}^T \mathbf{W}^T) \otimes (\mathbf{A} \mathbf{W} \mathbf{x} \mathbf{x}^T) \right] \\ &\quad + 2 \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f''(h(\mathbf{W} \mathbf{x}; \mathbf{A}, b)) \mathbb{1}\{\mathbf{W} \mathbf{x} \in T\} (\mathbf{x} \mathbf{x}^T \mathbf{W}^T) \otimes \mathbf{I} \right]. \end{aligned}$$

Notice that \mathbf{A} is a symmetric matrix but it is not necessarily positive semi-definite. Nevertheless, using Fact 4, we can write it as $\mathbf{A} = \mathbf{A}^- + \mathbf{A}^+$, where \mathbf{A}^- is negative semi-definite, \mathbf{A}^+ is positive semi-definite and both have their l2-norms bounded by $\|\mathbf{A}\|_2$. Then, by splitting the expressions with triangle inequality, we can without loss of generality assume \mathbf{A} is positive semi-definite by losing a constant factor in the upper bound. Then, we replace all the non-negative scalar-valued $f'(\cdot)$ and $f''(\cdot)$ functions with their upper bound 1. Lastly, using linearity of expectation, we take expectation over terms involving \mathbf{x} , which gives $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}(I)} [\mathbf{x} \mathbf{x}^T] = \mathbf{I}$. Hence, we obtain

$$\|\nabla_{\mathbf{A}} \nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b)\|_2 \leq 4 \|(\mathbf{W} \mathbf{W}^T) \otimes (\mathbf{A} \mathbf{W})\|_2 + 2 \|\mathbf{W}^T \otimes \mathbf{I}\|_2.$$

Since $\|\mathbf{A}\|_2, \|\mathbf{W}\|_2 \leq \text{poly}(c)$, the l2-norm is bounded above by $\text{poly}(c)$.

For b , we have $\frac{\partial}{\partial b} \nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{G}}(\mathbf{W}; \mathbf{A}, b) = 2 \mathbf{A} \mathbf{W} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f''(h(\mathbf{W} \mathbf{x}; \mathbf{A}, b)) \mathbb{1}\{\mathbf{W} \mathbf{x} \in T\} \mathbf{x} \mathbf{x}^T \right]$. Similarly, the norm can be upper bounded by $\text{poly}(c)$. Hence, overall, the training gradient is $O_c(1)$ -Lipchitz with respect to \mathbf{A}, b .

D.3 Proof of Lemma 6

For Generator, it is easy to see that the gradient oracle takes the form

$$\mathbf{g}_{\mathcal{G}} = \nabla_{\mathbf{W}} \log(1 - D(\phi(\mathbf{W} \mathbf{x}); \mathbf{A}, b)) = \mathbb{1}\{\mathbf{W} \mathbf{x} \in T\} f'(h(\mathbf{W} \mathbf{x}; \mathbf{A}, b)) 2 \mathbf{A} \mathbf{W} \mathbf{x} \mathbf{x}^T,$$

where $\mathbf{x} \sim \mathcal{N}(I)$. Hence, we can upper bound the square ℓ_2 -norm of it by

$$\begin{aligned} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\|\mathbf{g}_{\mathcal{G}}\|_2^2 \right] &= \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\left\| \mathbb{1}\{\mathbf{W} \mathbf{x} \in T\} f'(h(\mathbf{W} \mathbf{x}; \mathbf{A}, b)) 2 \mathbf{A} \mathbf{W} \mathbf{x} \mathbf{x}^T \right\|_F^2 \right] \\ &\leq 4 \|\mathbf{A} \mathbf{W}\|_2^2 \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \|\mathbf{x} \mathbf{x}^T\|_F^2 \leq 4 \|\mathbf{A} \mathbf{W}\|_2^2 (d^2 + 2d). \end{aligned}$$

By definition of the projection sets \mathcal{Q}_G and \mathcal{Q}_D , we have $\|\mathbf{A}\|_2, \|\mathbf{W}\|_2 \leq \text{poly}(c)$. Hence, it holds $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\|\mathbf{g}_{\mathcal{G}}\|_2^2 \right] \leq O_c(d^2)$.

D.4 Proof of Lemma 7

In this section, we prove our main optimization tool: a lemma stating the convergence of Biased Stochastic Gradient Descent (BSGD). Technically its proof is standard (similar to the work of [GLZ16]) and we provide it here for completeness. In the setting, we try to optimize a function $f(\mathbf{x})$ when only a biased gradient estimator $\xi(\mathbf{x})$ of $\nabla f(\mathbf{x})$ is provided. In particular, we study the following Projected Stochastic Gradient Descent Algorithm under a convex set \mathcal{Q} .

Algorithm 2 Biased PSGD for $f(\mathbf{w})$

Procedure: BPSGD(f, M, β)

- 1: Sample the stopping time m uniformly from $\{1, \dots, M\}$
 - 2: **for** $i = 1, \dots, m$ **do**
 - 3: Sample $\xi^{(i)}$. $\triangleright \|\mathbf{E}[\xi] - \nabla_{\mathbf{w}} f(\mathbf{w}^{(i)})\|_2 < \alpha$
 - 4: $\mathbf{w}^{(i+1/2)} \leftarrow \mathbf{w}^{(i)} - \beta \xi^{(i)}$.
 - 5: $\mathbf{w}^{(i+1)} \leftarrow \operatorname{argmin}_{\mathbf{w} \in \mathcal{Q}} \|\mathbf{w} - \mathbf{w}^{(i+1/2)}\|_2$.
 - 6: **end for**
-

Notice that, since we do not require the objective function $f(\mathbf{x})$ to be convex, we can only guarantee convergence to stationary points of the objective function.

We will use the following Lemma which is standard for non-convex projected gradient descent.

Lemma 17. *Assume function f is L -smooth. Consider the gradient mapping $g_{\mathcal{Q}}^{\eta}(\mathbf{w}, \nabla f(\mathbf{w})) = \frac{1}{\eta}(\mathbf{w} - \operatorname{proj}_{\mathcal{Q}}(\mathbf{w} - \eta \nabla f(\mathbf{w})))$. It holds $g_{\mathcal{Q}}^{\eta}(\mathbf{w}, \nabla f(\mathbf{w})) \geq \epsilon_{\frac{1}{1+L\eta}}$ if there exists $\mathbf{u} \in \mathcal{Q}$ such that $\frac{1}{\|\bar{\mathbf{w}} - \mathbf{u}\|_2} \langle \nabla f(\bar{\mathbf{w}}), \bar{\mathbf{w}} - \mathbf{u} \rangle \geq \epsilon$, where $\bar{\mathbf{w}} = \operatorname{proj}_{\mathcal{Q}}(\mathbf{w} - \eta \nabla f(\mathbf{w}))$.*

Proof of Lemma 7. Consider the update before the projection step $\mathbf{w}^{(i+1/2)} = \mathbf{w}^{(i)} - \beta \xi^{(i)}$. After the projection step we have $\mathbf{w}^{(i+1)} = \operatorname{argmin}_{\mathbf{x} \in \mathcal{Q}} \|\mathbf{w}^{(i+1/2)} - \mathbf{x}\|_2^2$. Denote the projection operator as $\mathbf{p}_{\mathcal{Q}}(\mathbf{w}, \xi) = \operatorname{argmin}_{\mathbf{x} \in \mathcal{Q}} \|\mathbf{x} - \mathbf{w} - \beta \xi\|_2$. Besides, we define the gradient mapping on the convex set \mathcal{Q} of point \mathbf{x} to be $g_{\mathcal{Q}}(\mathbf{w}, \xi) = (1/\beta)(\mathbf{w} - \mathbf{p}_{\mathcal{Q}}(\mathbf{w}, \xi))$. It follows from standard arguments (for example, Theorem 1.2.3 of [Nes13]) that

$$\begin{aligned} f(\mathbf{w}^{(i+1)}) - f(\mathbf{w}^{(i)}) &\leq \nabla f(\mathbf{w}^{(i)})^T (\mathbf{w}^{(i+1)} - \mathbf{w}^{(i)}) + \frac{L}{2} \|\mathbf{w}^{(i+1)} - \mathbf{w}^{(i)}\|_2^2 \\ &\leq -\beta \nabla f(\mathbf{w}^{(i)})^T g_{\mathcal{Q}}(\mathbf{w}^{(i)}, \xi) + \frac{L\beta^2}{2} \|g_{\mathcal{Q}}(\mathbf{w}^{(i)}, \xi)\|_2^2 \\ &\leq -\beta \nabla f(\mathbf{w}^{(i)})^T g_{\mathcal{Q}}(\mathbf{w}^{(i)}, \xi) + \frac{LB\beta^2}{2}. \end{aligned}$$

Notice that since ξ is a biased estimate of the gradient we have $\mathbf{E}[g_{\mathcal{Q}}(\mathbf{w}, \xi)] = g_{\mathcal{Q}}(\mathbf{w}, \nabla f(\mathbf{w})) + \mathbf{e}$, for some error vector \mathbf{e} with $\|\mathbf{e}\|_2 \leq \alpha$. This is true because in expectation the minimizer of $\|\mathbf{w}^{(i+1/2)} - \mathbf{x}\|_2$ only changes by \mathbf{e} . Additionally, f is l -Lipchitz. Therefore, after taking the expectation conditional on $\mathbf{w}^{(i)}$, we have

$$\mathbf{E}[f(\mathbf{w}^{(i+1)}) - f(\mathbf{w}^{(i)}) | \mathbf{w}^{(i)}] \leq -\beta \nabla f(\mathbf{w}^{(i)})^T g_{\mathcal{Q}}(\mathbf{w}^{(i)}, \nabla f(\mathbf{w}^{(i)})) + \frac{LB\beta^2}{2} + l\beta\alpha.$$

Since we project onto a convex set \mathcal{Q} we have that

$$\nabla f(\mathbf{w}^{(i)})^T g_{\mathcal{Q}}(\mathbf{w}^{(i)}, \nabla f(\mathbf{w}^{(i)})) \geq \|g_{\mathcal{Q}}(\mathbf{w}^{(i)}, \nabla f(\mathbf{w}^{(i)}))\|_2^2.$$

Therefore putting everything together we obtain,

$$\mathbf{E}[f(\mathbf{w}^{(i+1)}) - f(\mathbf{w}^{(i)}) | \mathbf{w}^{(i)}] \leq -\beta \left\| g_{\mathcal{Q}}(\mathbf{w}^{(i)}, \nabla f(\mathbf{w}^{(i)})) \right\|_2^2 + \frac{LB\beta^2}{2} + l\beta\alpha.$$

Rearranging, summing over $i = 1, \dots, M$, and using the law of total expectation, we obtain

$$\sum_{i=1}^M \left\| g_{\mathcal{Q}}^{(i)} \right\|_2^2 \leq \frac{R}{\beta} + \beta \frac{LB}{2} M + l\alpha M.$$

Picking step size $\beta = \sqrt{2R/(LBM)}$ we obtain that $\sum_{i=1}^M \left\| g_{\mathcal{Q}}^{(i)} \right\|_2^2 \leq \sqrt{2RLBM} + l\alpha M$. Next we choose a random stopping time m uniformly in $\{1, \dots, M\}$, where $M = O(RLB/\epsilon^4)$. We then have

$$E \left[\left\| g_{\mathcal{Q}}^{(m)} \right\|_2^2 \right] = \frac{1}{M} \sum_{i=1}^M \left\| g_{\mathcal{Q}}^{(i)} \right\|_2^2 \leq \sqrt{2RLB/M} + l\alpha \leq O(\epsilon^2) + l\alpha$$

From Markov's inequality we get that with probability at least 99% it holds that when the SGD stops we have $\left\| g_{\mathcal{Q}}^{(m)} \right\|_2 \leq \sqrt{O(\epsilon^2 + l\alpha)} \leq O(\epsilon + \sqrt{l\alpha})$. Then, applying Lemma 17, it holds for any $\mathbf{u} \in \mathcal{Q}$,

$$\frac{1}{\left\| \mathbf{w}^{(m)} - \mathbf{u} \right\|_2} \langle \nabla_{\mathbf{w}} f(\mathbf{w}^{(m)}), \mathbf{w}^{(m)} - \mathbf{u} \rangle \leq O(\epsilon + \sqrt{l\alpha})(1 + L\beta) \leq O(\epsilon + \sqrt{l\alpha}),$$

where the last inequality follows from the fact that $L\beta = L\sqrt{\frac{R}{LBM}} = O(\epsilon^2/B^2) = o(1)$. \square

D.5 Some additional properties required by Biased SGD

As a standard requirement for Gradient Descent, we show that $\mathcal{V}(\mathbf{W})$ is locally smooth and Lipchitz-continuous.

Lemma 18. $\mathcal{V}(\mathbf{W})$ is poly(c)-smooth and poly(c)-Lipchitz continuous with respect to \mathbf{W} in the projection set \mathcal{Q}_G when Assumption 1 is satisfied.

Proof. Recall that the gradient of the Virtual Training Criteria $\mathcal{V}(\mathbf{W})$ is given as

$$\nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W}) = \left(\boldsymbol{\Sigma}_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right) \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) \mathbf{x} \mathbf{x}^T \mathbf{1}\{\mathbf{W}\mathbf{x} \in T\} \right]$$

Since $f'(\cdot)$ is a positive function upper bounded by 1, it holds

$$\left\| \nabla_{\mathbf{W}} \mathcal{V}(\mathbf{W}) \right\|_2 \leq \left\| \boldsymbol{\Sigma}_*^{-1} \mathbf{W} - (\mathbf{W}^{-1})^T \right\|_2$$

Since $\left\| \mathbf{W} \right\|_2, \left\| \boldsymbol{\Sigma}_* \right\|_2 \leq \text{poly}(c)$ by definition of the projection set and Assumption 1, we conclude $\mathcal{V}(\mathbf{W})$ is poly(c)-Lipchitz continuous.

Next, we compute and upper bound the l2-norm of the hessian.

$$\begin{aligned} \nabla_{\mathbf{W}}^2 \mathcal{V}(\mathbf{W}) &= \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbf{1}\{\mathbf{W}\mathbf{x} \in T\} f''(h(\mathbf{W}\mathbf{x}; \mathbf{W})) \left((\mathbf{W}^{-1})^T - \boldsymbol{\Sigma}_*^{-1} \mathbf{W} \mathbf{x} \mathbf{x}^T \right) \right. \\ &\quad \left. \otimes \left(\boldsymbol{\Sigma}_*^{-1} \mathbf{W} \mathbf{x} \mathbf{x}^T - (\mathbf{W}^{-1})^T \mathbf{x} \mathbf{x}^T \right) \right] \\ &\quad + \mathbf{E}_{\mathbf{x} \sim \mathcal{N}} \left[\mathbf{1}\{\mathbf{W}\mathbf{x} \in T\} f'(h(\mathbf{W}\mathbf{x}; \mathbf{W})) \left(\mathbf{x} \mathbf{x}^T \otimes \boldsymbol{\Sigma}_*^{-1} + (\mathbf{x} \mathbf{x}^T (\mathbf{W}^{-1})^T) \otimes \mathbf{W}^{-1} \right) \right]. \end{aligned}$$

We could then use triangle inequality to split the expressions into sum of positive definite matrices (in a sense that $\mathbf{z}^T \mathbf{M} \mathbf{z} \geq 0$ for any \mathbf{z}). Then, we replace all the non-negative scalar-valued $f'(\cdot)$ and $f''(\cdot)$ with their upper bound 1. Lastly, by linearity of expectation, we take expectation over terms involving \mathbf{x} , which gives $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{I})} [\mathbf{x} \mathbf{x}^T] = \mathbf{I}$. Hence, we obtain the following upper bound

$$\begin{aligned} \|\nabla_{\mathbf{W}}^2 \mathcal{V}(\mathbf{W})\|_2 &\leq \|(\mathbf{W}^{-1})^T \otimes (\boldsymbol{\Sigma}_*^{-1} \mathbf{W})\|_2 + \|(\mathbf{W}^{-1})^T \otimes (\mathbf{W}^{-1})^T\|_2 + \|(\boldsymbol{\Sigma}_*^{-1} \mathbf{W}) \otimes (\boldsymbol{\Sigma}_*^{-1} \mathbf{W})\|_2 \\ &\quad + \|(\boldsymbol{\Sigma}_*^{-1} \mathbf{W}) \otimes (\mathbf{W}^{-1})^T\|_2 + \|\mathbf{I} \otimes \boldsymbol{\Sigma}_*^{-1}\|_2 + \|(\mathbf{W}^{-1})^T \otimes \mathbf{W}^{-1}\|_2 \end{aligned}$$

Since the l2-norm of \mathbf{W} , \mathbf{W}^{-1} , $\boldsymbol{\Sigma}_*^{1/2}$ and $\boldsymbol{\Sigma}_*^{-1/2}$ are all bounded by $\text{poly}(c)$, it then follows $\|\nabla_{\mathbf{W}}^2 \mathcal{V}(\mathbf{W})\|_2 \leq \text{poly}(c)$. \square