
On Convergence of Gradient Descent Ascent: A Tight Local Analysis

Haochuan Li¹ Farzan Farnia² Subhro Das³ Ali Jadbabaie⁴

Abstract

Gradient Descent Ascent (GDA) methods are the mainstream algorithms for minimax optimization in generative adversarial networks (GANs). Convergence properties of GDA have drawn significant interest in the recent literature. Specifically, for $\min_{\mathbf{x}} \max_{\mathbf{y}} f(\mathbf{x}; \mathbf{y})$ where f is strongly-concave in \mathbf{y} and possibly nonconvex in \mathbf{x} , (Lin et al., 2020) proved the convergence of GDA with a stepsize ratio $\eta_{\mathbf{y}}/\eta_{\mathbf{x}} = \Theta(\kappa^2)$ where $\eta_{\mathbf{x}}$ and $\eta_{\mathbf{y}}$ are the stepsizes for \mathbf{x} and \mathbf{y} and κ is the condition number for \mathbf{y} . While this stepsize ratio suggests a slow training of the min player, practical GAN algorithms typically adopt similar stepsizes for both variables, indicating a wide gap between theoretical and empirical results. In this paper, we aim to bridge this gap by analyzing the *local convergence* of general *nonconvex-nonconcave* minimax problems. We demonstrate that a stepsize ratio of $\Theta(\kappa)$ is necessary and sufficient for local convergence of GDA to a Stackelberg Equilibrium, where κ is the local condition number for \mathbf{y} . We prove a nearly tight convergence rate with a matching lower bound. We further extend the convergence guarantees to stochastic GDA and extra-gradient methods (EG). Finally, we conduct several numerical experiments to support our theoretical findings.

1. Introduction

Minimax learning frameworks including generative adversarial networks (GANs) (Goodfellow et al., 2014) and adversarial training (Madry et al., 2017) have achieved great success in various machine learning tasks. According to these frameworks, the underlying machine learning prob-

lem is formulated as a zero-sum game between two min and max players trying to optimize a learning objective in the opposite directions. Therefore, minimax learning tasks are commonly formulated through a minimax optimization problem of the following form:

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}; \mathbf{y}). \quad (1)$$

In the above problem, \mathbf{x} and \mathbf{y} denote the optimization variables for the minimization and maximization subproblems, respectively, and $f(\mathbf{x}; \mathbf{y})$ represents the minimax learning objective estimated over observed training data. For example, in GAN settings, \mathbf{x} contains the parameters of a generator machine producing real-like samples, and \mathbf{y} represents the parameters of a discriminator machine distinguishing the generated samples from real training data.

Since the machine players of modern minimax learning frameworks are typically chosen as deep neural networks, these frameworks lead to difficult nonconvex-nonconcave minimax optimization problems where $f(\mathbf{x}; \mathbf{y})$ is possibly nonconvex in \mathbf{x} and nonconcave in \mathbf{y} . In such optimization problems, standard gradient-based algorithms are not guaranteed to find a locally optimal solution. However, numerous empirical studies on GANs suggest that a simultaneous-update gradient-based optimization algorithm such as Gradient Descent Ascent (GDA) can successfully train a satisfactory generator function. The wide gap between standard convergence guarantees and empirical results in GAN experiments has inspired several recent studies on the convergence behavior of GDA-type methods in general nonconvex-nonconcave minimax problems.

A recent line of works in the minimax optimization literature focuses on the convergence properties of gradient-based optimization algorithms in the *nonconvex-concave minimax settings* where the objective $f(\mathbf{x}; \mathbf{y})$ is further assumed to be a concave function of \mathbf{y} . Regarding this class of minimax optimization problems, a well-known result proved by (Lin et al., 2020a) demonstrates a convergence guarantee for the GDA and stochastic GDA (SGDA) algorithms in the smooth nonconvex-strongly-concave (NC-SC) minimax settings where $f(\mathbf{x}; \mathbf{y})$ is assumed to be L -smooth and μ -strongly-concave in \mathbf{y} . Under these assumptions, Lin et al. (2020a) analyze the convergence of two-time-scale GDA with different stepsizes $\eta_{\mathbf{x}} = \Theta(\frac{1}{\kappa^2 L})$ and $\eta_{\mathbf{y}} = \Theta(\frac{1}{L})$ which achieves a gradient complexity bound of $\mathcal{O}(\kappa^2/\epsilon^2)$

¹Department of EECS, Massachusetts Institute of Technology

²Department of CSE, The Chinese University of Hong Kong ³MIT-IBM Watson AI Lab, IBM Research ⁴Department of CEE, Massachusetts Institute of Technology. Correspondence to: Haochuan Li <haochuan@mit.edu>.



Figure 1. Generated images of the learned generator on MNIST (a) and CIFAR10 (b). For both MNIST and CIFAR10, we train WGAN-GP models (Gulrajani et al., 2017) using simultaneous GDA with $\eta_x = \eta_y = 0.001$.

for $\kappa := L/\mu$ defined as the condition number of the minimax objective and ϵ denoting the stationarity degree of the desired solution.

While the above result sheds light on the convergence properties of two-time-scale GDA methods for general minimax settings, it suggests choosing a significantly smaller stepsize value for the minimization subproblem compared to that of the maximization subproblem. While the recommended stepsize ratio $\eta_y/\eta_x = \Theta(\kappa^2)$ would tend to infinity for general nonconvex-nonconcave problems, practical GAN experiments often choose similar stepsize values for the minimization and maximization tasks. For example, Figure 1 shows that choosing a $\Theta(1)$ stepsize ratio may already suffice to converge to a good solution for GANs trained over standard MNIST and CIFAR-10 datasets. In addition, Farnia & Ozdaglar (2021) provide theoretical and numerical evidence that a smaller GDA stepsize ratio can result in better generalization performance for the gradient-based minimax learners. Such empirical and theoretical results question the necessity of a large stepsize ratio for minimax learning algorithms, and further motivate the following question:

What is the best stepsize ratio that ensures convergence of GDA and what is the corresponding convergence rate?

To answer this question, we develop a tight local convergence analysis near a Stackelberg Equilibrium of smooth minimax optimization problems. Specifically, we show that a stepsize ratio of $\Theta(\kappa)$, where κ is the local condition number of the max player, is necessary and sufficient to ensure the convergence of GDA. When applied to the NC-SC setting, it is much smaller than that in (Lin et al., 2020a) and enables a faster local convergence rate. We also show that our stepsize ratio and the corresponding convergence rate are nearly optimal by providing matching lower bounds.

Different from standard analysis methods in optimization, we leverage the tools in linear dynamical systems and control theory to derive necessary and sufficient conditions under which GDA is guaranteed to converge. Specifically, note that a smooth function is approximately quadratic NC-

SC locally around its Stackelberg Equilibrium and thus the dynamics of GDA is nearly a linear time invariant (LTI) dynamical system. The convergence is then determined by the spectral radius ρ , the largest modulus of the complex eigenvalues, of the transition matrix of this system. For example, $\rho < 1$ implies convergence and $\rho > 1$ implies divergence. However, the transition matrix is a non-symmetric and pretty general block matrix whose eigenvalues could be complex numbers and have no closed-form solution. In Lemma 5.2, we provide a novel and fine-grained analysis of its spectral properties, which is then used to obtain our convergence results. Lemma 5.2 is also our main technical contribution which in itself can be interesting to the optimization community.

1.1. Summary of main contributions

- We provide a precise answer to the above question for local convergence to a Stackelberg Equilibrium of locally L smooth functions. In this setting, we prove that choosing $\eta_x = \Theta(1/rL)$ and $\eta_y = \Theta(1/L)$ with a stepsize ratio $r = \Omega(\kappa)$ suffices to guarantee convergence. We then show the necessity of the $\Omega(\kappa)$ stepsize ratio to guarantee convergence. Therefore the optimal choice is $r = \Theta(\kappa)$.
- For any $r = \Omega(\kappa)$, we provide a nearly tight convergence rate with matching lower bounds in this setting. Our result is locally much faster than that under the conservative stepsize choice in (Lin et al., 2020a).
- We further extend our local convergence results to Stochastic Gradient Descent Ascent (SGDA) and extragradient methods (EG). For EG, the convergence rate is also optimal for any $r = \Omega(\kappa)$.
- We conduct experiments on quadratic NC-SC functions to illustrate the convergence behaviors under different stepsize ratios. We also run the GDA algorithm to optimize non-quadratic functions and show that such behaviors further exist globally for several non-quadratic functions. Finally, we conduct experiments on GANs

to show that simultaneous GDA with $\eta_x \approx \eta_y$ enjoys fast convergence to a desired solution.

2. Related work

GDA-based minimax optimization algorithms. Gradient Descent Ascent (GDA) naturally extends gradient descent (GD) to the minimax optimization problem (1). The algorithm takes one step of gradient descent over the minimization variable \mathbf{x} with stepsize η_x and another step of gradient ascent over the maximization variable \mathbf{y} with stepsize η_y . A lot of convergence results have been established for the average iterates generated by GDA when the objective function is convex in \mathbf{x} and concave in \mathbf{y} (Chen & Rockafellar, 1997; Nedić & Ozdaglar, 2009; Nemirovski, 2004; Du & Hu, 2019).

However, in terms of last iterate convergence, some works (Benaïm & Hirsch, 1999; Hommes & Ochea, 2012; Mertikopoulos et al., 2018b) show that GDA with equal stepsizes ($\eta_x = \eta_y$) can converge to limit cycles or diverge even for convex-concave functions. Two-time-scale GDA with unequal stepsizes ($\eta_x \neq \eta_y$) is shown empirically to alleviate the issues of limit cycles and its local asymptotic convergence to Nash equilibria is theoretically guaranteed (Heusel et al., 2017). Lin et al. (2020a) established non-asymptotic convergence guarantees in the nonconvex-concave setting.

In a closely-related work, Fiez & Ratliff (2021) demonstrate how to construct a ratio r^* that ensures GDA’s local convergence. However, this work does not focus on answering what the optimal choice of r^* will be. Moreover, the questions of whether the proposed r^* is the optimal choice and how large r^* and their convergence rates are remain unanswered in this paper. In our work, we provide a precise answer by showing what the optimal r^* and the convergence rates are for local convergence.

Other algorithms. Many works study algorithms other than GDA and provide strong theoretical results in the convex-concave setting. The popular extra-gradient method (EG), which takes two gradient updates per iteration, is one of them (Korpelevich, 1976; Tseng, 1995; Nemirovski, 2004; Chambolle & Pock, 2011; Yadav et al., 2017). Another line studies the optimistic gradient method (Rakhlin & Sridharan, 2013; Daskalakis et al., 2017; Gidel et al., 2018; Mertikopoulos et al., 2018a; Hsieh et al., 2019). Mokhtari et al. (2020) provided a unified analysis of both methods as approximations of the classical proximal point method. (Lin et al., 2020b) focused on accelerating the known rates in terms of the conditional number dependence and their rates match the lower bounds (Ouyang & Xu, 2019; Ibrahim et al., 2019; Zhang et al., 2019).

Recently, several convergence results under the nonconvex-concave setup are established. Rafique et al. (2018) pro-

posed proximally guided gradient methods and analyzed its convergence. Thekumparampil et al. (2019) proposed an inexact proximal point method for the nonconvex-concave setting. Many other works also studied different variations of convergence (Nouiehed et al., 2019; Kong & Monteiro, 2019; Lu et al., 2020; Ostrovskii et al., 2020). The algorithm of (Lin et al., 2020b) was shown to be nearly optimal in the nonconvex-strongly-concave setting later in (Li et al., 2021; Zhang et al., 2021; Han et al., 2021)

The general nonconvex-nonconcave setting is even more challenging. First, a saddle-point may not exist (Jin et al., 2020; Farnia & Ozdaglar, 2020). Moreover, determining its existence is known to be NP-hard (Daskalakis et al., 2020), and finding an approximate local saddle point is PPAD-complete (Daskalakis et al., 2020). It remains an open problem how to define a nontrivial and tractable notion of minimax points.

3. Preliminaries

Notation. We use bold lower-case letters to denote vectors, upper-case letters to denote matrices, and calligraphic upper-case letters to denote sets. Let $\|\mathbf{x}\|_2$ be the ℓ_2 norm of vector \mathbf{x} and $\|A\|_2$ be the spectral norm of matrix A . For a matrix A , we use A^\top , $\det(A)$, and $\text{Trace}(A)$ to denote its transpose, determinant, and trace respectively. Given a symmetric real matrix A , we use $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ to denote its smallest and largest eigenvalues. Given another symmetric real matrix B , we write $A < B$ or $B > A$ if $B - A$ is positive definite. We also write $A \leq B$ or $B \geq A$ if $B - A$ is positive semi-definite. We use I_d to denote the identity matrix in $\mathbb{R}^{d \times d}$ and omit the subscript when d is clear from the context. Finally, we use the standard $\mathcal{O}(\cdot)$, $\Theta(\cdot)$, and $\Omega(\cdot)$ notation, with $\tilde{\mathcal{O}}(\cdot)$, $\tilde{\Theta}(\cdot)$, and $\tilde{\Omega}(\cdot)$ further hiding logarithmic factors.

3.1. Problem setup

We study the minimax optimization problem (1) where f is twice-continuously differentiable. This assumption guarantees that $\nabla^2 f(\mathbf{x}; \mathbf{y})$ is a symmetric Hessian matrix for any $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$ and $\mathbf{y} \in \mathcal{Y} \subseteq \mathbb{R}^m$. We will often denote $\mathbf{z} = (\mathbf{x}, \mathbf{y})$ for simplicity.

We further assume that there exists a differential Stackelberg Equilibrium $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$. Formally, denote the Hessian at \mathbf{z}^* as

$$\begin{pmatrix} \nabla_{\mathbf{x}\mathbf{x}}^2 f(\mathbf{z}^*) & \nabla_{\mathbf{x}\mathbf{y}}^2 f(\mathbf{z}^*) \\ \nabla_{\mathbf{y}\mathbf{x}}^2 f(\mathbf{z}^*) & \nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*) \end{pmatrix} = \begin{pmatrix} C & B \\ B^\top & -A \end{pmatrix}. \quad (2)$$

Then differential Stackelberg Equilibrium is defined as follows.

Definition 3.1 ((Fiez et al., 2020)). A point \mathbf{z}^* satisfying (2) is a differential Stackelberg Equilibrium if $\nabla_{\mathbf{x}} f(\mathbf{z}^*) = \mathbf{0}$,

$$\nabla_{\mathbf{y}} f(\mathbf{z}^*) = \mathbf{0}, A > 0, C + BA^{-1}B^\top > 0.$$

Here $C + BA^{-1}B^\top$ is the Schur complement with respect to the top left $n \times n$ block in the block matrix $\nabla^2 f(\mathbf{z}^*)$. We will see later that it is also the Hessian of the primal function $\Phi(\cdot) = \max_{\mathbf{y} \in \mathcal{Y}} f(\cdot; \mathbf{y})$ at \mathbf{x}^* . The condition $A > 0$ implies the function is locally strictly concave in \mathbf{y} around \mathbf{z}^* . The condition $C + BA^{-1}B^\top$ further suggests Φ is locally strictly convex.

We assume that f is locally L smooth around \mathbf{z}^* . Specifically, assume¹

$$\max\{\|A\|_2, \|B\|_2, \|C\|_2\} \leq L.$$

Denote

$$\mu = \lambda_{\min}(A), \quad \mu_x = \min\{L, \lambda_{\min}(C + BA^{-1}B^\top)\}$$

and the corresponding condition numbers

$$\kappa = L/\mu, \quad \kappa_x = L/\mu_x.$$

We will call \mathbf{z}^* a Stackelberg Equilibrium with parameters (L, μ, μ_x) . Also remark that Φ is locally $(\kappa + 1)L$ smooth and μ_x strongly convex, which implies the local condition number of Φ is actually $\kappa_x(\kappa + 1)$.

Learning Dynamics In this paper, we focus on the two-time-scale GDA algorithm and also extend the results to two-time-scale EG. Let η_x and η_y be the stepsizes for \mathbf{x} and \mathbf{y} and $r = \eta_y/\eta_x$ be the stepsize ratio.

At each iterate $\mathbf{z}^k = (\mathbf{x}^k, \mathbf{y}^k)$, we are given an exact or stochastic oracle that returns gradient $(g_x(\mathbf{z}^k), g_y(\mathbf{z}^k))$. If it is an exact oracle, we have

$$g_x(\mathbf{z}^k) = \nabla_{\mathbf{x}} f(\mathbf{z}^k), \quad g_y(\mathbf{z}^k) = \nabla_{\mathbf{y}} f(\mathbf{z}^k).$$

In the stochastic setting, we hide the randomness in the random functions g_x and g_y . Denote the noise of gradient as

$$\begin{aligned} \xi_{\mathbf{x}}^k &= g_x(\mathbf{z}^k) - \nabla_{\mathbf{x}} f(\mathbf{z}^k), \\ \xi_{\mathbf{y}}^k &= g_y(\mathbf{z}^k) - \nabla_{\mathbf{y}} f(\mathbf{z}^k). \end{aligned} \quad (3)$$

We assume the gradient is estimated with a mini-batch of samples of size S and is unbiased with bounded variance. Formally, we assume

Assumption 3.2. $\mathbb{E}[\xi_{\mathbf{x}}^k] = \mathbf{0}$, $\mathbb{E}[\xi_{\mathbf{y}}^k] = \mathbf{0}$, $\mathbb{E}[\|\xi_{\mathbf{x}}^k\|_2^2] \leq \sigma^2/S$, and $\mathbb{E}[\|\xi_{\mathbf{y}}^k\|_2^2] \leq \sigma^2/S$ for some $\sigma \geq 0$. Also, the noises for different k are independent.

Remark that in the stochastic setting, it takes S gradient evaluations to estimate $(g_x(\mathbf{z}^k), g_y(\mathbf{z}^k))$.

¹Our definition of smoothness is slightly different from the usual one up to a factor of 2.

The learning dynamics of (stochastic) GDA is given by

$$\begin{aligned} \mathbf{x}^{k+1} &= \mathbf{x}^k - \eta_x g_x(\mathbf{x}^k; \mathbf{y}^k), \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + \eta_y g_y(\mathbf{x}^k; \mathbf{y}^k). \end{aligned} \quad (4)$$

The learning dynamics of (stochastic) EG is given by

$$\begin{aligned} \mathbf{x}^{k+1/2} &= \mathbf{x}^k - \eta_x g_x(\mathbf{x}^k; \mathbf{y}^k), \\ \mathbf{y}^{k+1/2} &= \mathbf{y}^k + \eta_y g_y(\mathbf{x}^k; \mathbf{y}^k); \\ \mathbf{x}^{k+1} &= \mathbf{x}^k - \eta_x g_x(\mathbf{x}^{k+1/2}; \mathbf{y}^{k+1/2}), \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + \eta_y g_y(\mathbf{x}^{k+1/2}; \mathbf{y}^{k+1/2}). \end{aligned} \quad (5)$$

From Definition 3.1, we can see that f is approximately a quadratic NC-SC function near \mathbf{z}^* . The dynamics of GDA or EG is a linear time invariant system. Specifically, define matrix

$$M = \begin{pmatrix} -C & -B \\ rB^\top & -rA \end{pmatrix}. \quad (6)$$

Let $\mathbf{z}^k = (\mathbf{x}^k, \mathbf{y}^k)$. We will see that with an exact gradient oracle, the dynamics of GDA is approximately

$$\mathbf{z}^{k+1} - \mathbf{z}^* \approx (1 + \eta_x M) \cdot (\mathbf{z}^k - \mathbf{z}^*). \quad (7)$$

Similarly, the dynamics of EG is approximately

$$\mathbf{z}^{k+1} - \mathbf{z}^* \approx (1 + \eta_x M + \eta_x^2 M^2) \cdot (\mathbf{z}^k - \mathbf{z}^*). \quad (8)$$

4. Main results

In this section, we present our main results on the local convergence of GDA and EG.

4.1. A lower bound of the stepsize ratio for GDA

To start with, we show that for GDA, the stepsize ratio r must be at least $\Omega(\kappa)$ to ensure local convergence to a Stackelberg Equilibrium. Formally, we have the following theorem.

Theorem 4.1. *For any given $L, \mu > 0$ and $0 \leq \mu_x \leq L$ such that $\kappa = L/\mu \geq 2$, there exists a function f that has a differential Stackelberg Equilibrium with parameters (L, μ, μ_x) such that if choosing a stepsize ratio $r \leq \kappa$, the GDA algorithm does not locally converge for all initial points irrespective of any choice of positive stepsize η_x .*

The hard function f in Theorem 4.1 is constructed as $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ with the following expression

$$f(x; y) = -\frac{\mu}{2}y^2 + Lxy - \frac{L}{2}x^2.$$

Note that since the above f is a quadratic NC-SC function and global convergence implies local convergence, Theorem 4.1 can be also viewed as a lower bound of the stepsize

ratio of GDA for the general NC-SC minimax optimization problem studied in (Lin et al., 2020a).

For such a simple function f , we can directly solve the two eigenvalues of its corresponding M matrix and show that one of them, say λ_1 , has a positive real part. Then $|1 + \eta_x \lambda_1| > 1$ for any positive η_x which implies divergence of GDA. However, it does not directly apply to EG because it is much harder to lower bound $|1 + \eta_x \lambda_1 + \eta_x^2 \lambda_1^2|$ unless assuming η_x is very small (which is not desirable). We leave the optimal choice of the stepsize ratio for EG as an interesting open problem.

4.2. Convergence rate of GDA and EG

Now we show that $r = \Omega(\kappa)$ is sufficient for both GDA and EG to converge. We also provide a tight local convergence rate for them.

Theorem 4.2. *Given a function f and its Stackelberg Equilibrium \mathbf{z}^* satisfying all the assumptions in Section 3. Choose $\eta_x = \frac{1}{4rL}$ and $\eta_y = \frac{1}{4L}$ with $r \geq 2\kappa$. Suppose that M defined in (6) is diagonalizable and its eigenvalue decomposition is $M = P^{-1}\Lambda P$. Let $C_P = \|P\|_2 \|P^{-1}\|_2$. There exists some $\delta > 0$ such that if $\|\mathbf{z}^0 - \mathbf{z}^*\|_2 \leq \delta$, we have the iterates of GDA and EG with an exact gradient oracle satisfy*

$$\|\mathbf{z}^T - \mathbf{z}^*\|_2 \leq C_P \|\mathbf{z}^0 - \mathbf{z}^*\|_2 \left(1 - \frac{c_0}{r\kappa_x}\right)^T$$

for all $T \geq 0$, where c_0 is some numerical constant. It means the gradient complexity of achieving an error less than ϵ is

$$\mathcal{O}\left(r\kappa_x \log\left(\frac{C_P \|\mathbf{z}^0 - \mathbf{z}^*\|_2}{\epsilon}\right)\right).$$

Remark 1. The assumption that M is diagonalizable is actually without loss of generality. One can easily obtain similar convergence result without this assumption by using Lemma 5.3 and (Polyak, 2020, Chapter 2.1.2, Theorem 1). However, we still keep this assumption because (1) it is a mild assumption as the set of non-diagonalizable matrices has zero measure and (2) it enables us to explicitly derive the dependence on all involved problem-dependent constants like the condition number C_P .

Remark 2. The factor 2 in the assumption $r \geq 2\kappa$ is also arbitrary. Following our analysis, we can replace it by $r \geq c\kappa$ for any numerical constant $c > 1$ and obtain nearly the same convergence result up to constant factors. However, we will always choose $c = 2$ in this paper to make expressions as simple as possible. Even if $c > 1$ is not a numerical constant and depends on κ or κ_x , local convergence can still be guaranteed although the rate might

be worse. Therefore, combined with the negative result in Theorem 4.1, we know $r = \kappa$ is the exact *phase transition point* between convergence and non-convergence.

The above complexity is $\tilde{\mathcal{O}}(r\kappa_x)$ which scales linearly in r when $r \geq 2\kappa$. So choosing $r = 2\kappa$ results in the best complexity bound of $\tilde{\mathcal{O}}(\kappa\kappa_x)$. This rate significantly improves upon the $\tilde{\mathcal{O}}(\kappa^2\kappa_x)$ complexity under the choice of $r = \Theta(\kappa^2)$ as in (Lin et al., 2020a).

Note that the condition number of the local primal function Φ is $\kappa_x(\kappa + 1)$ as Φ is $(\kappa + 1)L$ smooth. Therefore, the $\tilde{\mathcal{O}}(\kappa\kappa_x)$ gradient complexity for $r = 2\kappa$ is nearly the same as the iteration complexity of running gradient descent directly on Φ although we do not really have access to its value and gradient!

We also want to point out we fix $\eta_y = \Theta(1/L)$ and vary η_x for different r because it is well known that $\Theta(1/L)$ is the largest possible step size for L smooth functions to ensure convergence.

To prove Theorem 4.2, we note that the dynamics of GDA or EG is approximately a linear time invariant system as in (7,8). By a careful analysis of the spectral properties of M , we can show that the spectral radius of the transition matrix $I + \eta_x M$ or $I + \eta_x M + \eta_x^2 M^2$ is strictly upper bounded by 1. In Section 5, we will introduce our analysis methods in detail through simpler quadratic NC-SC functions. In this case, the dynamical systems (7,8) are exactly LTIs. Then bounding the spectral radius of the transition matrix is sufficient to prove convergence. We will show how to use the results for quadratic functions to prove Theorem 4.2 in Appendix 12.

It is also worth pointing out that when f is quadratic, it is actually nonconvex-strongly-concave (NC-SC). Also, the results for quadratic NC-SC functions are slightly stronger than Theorem 4.2. For example, M no longer needs to satisfy the diagonalizability assumption. As the dynamics of the quadratic functions are globally LTI, a global convergence result is feasible to obtain in this setting. It means that all our local convergence results can be viewed as global results for quadratic NC-SC functions.

Now we provide a matching lower bound to show the convergence rate in Theorem 4.2 is nearly tight for GDA.

Theorem 4.3. *Suppose we run GDA or EG with the stepsize choices in Theorem 4.2 with $r \geq 2\kappa$. There exists a function f and its Stackelberg Equilibrium \mathbf{z}^* satisfying all the assumptions in Section 3 with a diagonalizable M matrix, such that for any there exists a \mathbf{z}^0 that is arbitrarily close to \mathbf{z}^* and the iterates satisfy*

$$\|\mathbf{z}^T - \mathbf{z}^*\|_2 \geq \left(1 - \frac{1}{r\kappa_x}\right)^T \|\mathbf{z}^0 - \mathbf{z}^*\|_2.$$

The proof of Theorem 4.3 is deferred in Appendix 9. The following theorem establishes the convergence rate of GDA and EG when the gradient oracle is stochastic.

Theorem 4.4. *With the same assumptions and notation as in Theorem 4.2. There exists some $\delta > 0$ such that if $\|\mathbf{z}^0 - \mathbf{z}^*\|_2 \leq \delta$, we have with probability at least $1 - p$ for arbitrarily small constant p , the iterates of mini-batch SGDA and stochastic EG under Assumption 3.2 satisfy $\|\mathbf{z}^T - \mathbf{z}^*\|_2 \leq \epsilon$ if choosing*

$$T = O\left(r\kappa_x \log\left(\frac{C_P \|\tilde{\mathbf{z}}^0\|_2}{\epsilon}\right)\right)$$

$$S = O\left(\frac{r^2 \kappa_x^2 C_P^2 \sigma^2}{L^2 \epsilon^2} \log^2\left(\frac{C_P \|\tilde{\mathbf{z}}^0\|_2}{\epsilon}\right)\right).$$

This implies a gradient complexity of

$$T \cdot S = O\left(\frac{r^3 \kappa_x^3 C_P^2 \sigma^2}{L^2 \epsilon^2} \log^3\left(\frac{C_P \|\tilde{\mathbf{z}}^0\|_2}{\epsilon}\right)\right).$$

We are not sure whether the above $\tilde{O}\left(\frac{r^3 \kappa_x^3}{\epsilon^2}\right)$ complexity is tight or not. Actually, as we will show in next section, the convergence rate for quadratic NC-SC functions in the stochastic setting is $\tilde{O}\left(\frac{r^2 \kappa_x^2}{\epsilon^2}\right)$, strictly better than the above theorem. Here we get a worse rate because the deviation of the local dynamics of GDA/EG from the corresponding LTI can depend on the gradient noise, which makes the analysis harder. It is an interesting future direction to see whether the above rate can be improved.

5. Analysis through quadratic NC-SC functions

In this section, we use the quadratic NC-SC function class as an example to illustrate our analysis.

5.1. Problem setup for quadratic functions

We first formulate the problem in the quadratic setting. Consider the quadratic function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ of the following form:

$$f(\mathbf{x}; \mathbf{y}) = -\frac{1}{2}(\mathbf{y} - \mathbf{y}^*)^\top A(\mathbf{y} - \mathbf{y}^*) + (\mathbf{x} - \mathbf{x}^*)^\top B(\mathbf{y} - \mathbf{y}^*) + \frac{1}{2}(\mathbf{x} - \mathbf{x}^*)^\top C(\mathbf{x} - \mathbf{x}^*) \quad (9)$$

for some matrices $A \in \mathbb{R}^{m \times m}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{n \times n}$ where A and C are symmetric matrices. $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$ is the differential Stackelberg Equilibrium. In this setting, the assumptions in Section 3 are reduced as follows.

Assumption 5.1. The function f defined in (9) satisfies

1. $\mu I \leq A \leq LI$ and $\|B\|_2, \|C\|_2 \leq L$.
2. $C + BA^{-1}B^\top \geq \mu_x I > 0$.

Then one can check that the primal function is

$$\Phi(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}^*)^\top (C + BA^{-1}B^\top)(\mathbf{x} - \mathbf{x}^*).$$

Also, we have $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} \Phi(\mathbf{x})$ and $\mathbf{y}^* = \operatorname{argmax}_{\mathbf{y}} f(\mathbf{x}^*; \mathbf{y})$, i.e., \mathbf{z}^* is a global minimax point for this quadratic NC-SC minimax optimization problem.

5.2. Convergence in the deterministic setting

For quadratic functions, (7,8) are exactly Linear Time-Invariant (LTI) systems with transition matrices $I + \eta_x M$ and $I + \eta_x M + \eta_x^2 M^2$ respectively. Note that from the definition of M in (6), these two matrices are non-symmetric and pretty general block matrices. Their eigenvalues could be complex numbers. Although their spectral norms are easy to bound, they could be strictly greater than 1. Therefore, we have to bound a more fine-grained quantity, the spectral radius, which could be strictly less than the spectral norm. To bound the spectral radius, we provide a careful analysis of the spectral properties of M in the following lemma which is also our main technical contribution.

Lemma 5.2. *Suppose $\lambda = \lambda_0 + i\lambda_1$ is an eigenvalue of M where $\lambda_0, \lambda_1 \in \mathbb{R}$ and $i = \sqrt{-1}$. Then*

1. $|\lambda_1| \leq \sqrt{r}L$.
2. If $\lambda_1 \neq 0$, we have $\lambda_0 \leq -\frac{\mu(r-\kappa)}{2}$.
3. If $r \geq 1$, the spectral norm of M is bounded: $\lambda_0^2 + \lambda_1^2 \leq \|M\|_2^2 \leq 4r^2L^2$.
4. Choosing $r > \kappa$, if $\lambda_1 = 0$, we have $\lambda_0 \leq -\mu_x$.
5. If $r > \kappa$ and $\mu_x > 0$, we must have $\lambda_0 < 0$.

Remark that Lemma 5.2.1 and 5.2.5 together suggest that if $\mu_x > 0$ and $r > \kappa$, every eigenvalue λ of M has a strictly negative real part and a bounded imaginary part. Then if fixing $r = \eta_y/\eta_x > \kappa$ and choosing a small enough η_x (or η_y), one can guarantee $|1 + \eta_x \lambda| < 1$, i.e., the spectral radius of the transition matrix of GDA is less than 1 which implies convergence. Similarly, the convergence of EG can also be guaranteed.

Lemma 5.2 also suggests that the angle between the complex number λ and the real line in the complex plane is bounded by some $\theta < \pi/2$. If further assuming $r \geq c\kappa$ for some numerical constant $c > 1$, we can get a tight bound of θ which allows us to choose η_x to make $|1 + \eta_x \lambda|$ as small as possible. Then we can obtain a quantitative and tight bound of the spectral radius of the transition matrices in the following lemma.

Lemma 5.3. Choose $\eta_x = \frac{1}{4rL}$ and $\eta_y = \frac{1}{4L}$ where $r \geq 2\kappa$. Let ρ_1 be the spectral radius of $I + \eta_x M$ and ρ_2 be that of $I + \eta_x M + \eta_x^2 M^2$. We have

$$\max\{\rho_1, \rho_2\} \leq 1 - \frac{1}{64r\kappa_x}.$$

The proofs of Lemma 5.2 and Lemma 5.3 are deferred in Appendix 10. With Lemma 5.3, the convergence result of GDA and EG for quadratic NC-SC functions follows naturally.

Theorem 5.4. Let $M = P^{-1}JP$ be its Jordan decomposition. Define $C_P = \|P\|_2 \|P^{-1}\|_2$ as the condition number of P and s be the size of the largest Jordan block in J . Under Assumption 5.1, choosing $\eta_x = \frac{1}{4rL}$ and $\eta_y = \frac{1}{4L}$ with $r \geq 2\kappa$, we have the iterates of GDA or EG with an exact gradient oracle satisfy

$$\|z^T - z^*\|_2 \leq sC_P \|z^0 - z^*\|_2 T^{s-1} \left(1 - \frac{c_0}{r\kappa_x}\right)^{T-s+1}$$

for all $T \geq 0$, where c_0 is some numerical constant. It means the gradient complexity of achieving $\|z^T - z^*\|_2 \leq \epsilon$ is

$$\mathcal{O}\left(sr\kappa_x \log\left(\frac{sr\kappa_x C_P \|z^0 - z^*\|_2}{\epsilon}\right)\right).$$

Note that Theorem 5.4 is a global convergence result and does not assume M is diagonalizable. However, the gradient complexity has a factor of s which could be proportional to the dimension $d = m + n$ in the worst case when M is not diagonalizable but is usually very small.

5.3. Convergence in the stochastic setting

The following theorem establishes a convergence result for mini-batch SGDA.

Theorem 5.5. We use the same assumptions and parameters as in Theorem 5.4 except that now we consider stochastic gradient oracles. The iterates of SGDA and stochastic EG satisfy

$$\mathbb{E}[\|z^T - z^*\|_2^2] \leq C_P^2 \left(1 - \frac{c_1}{r\kappa_x}\right)^{2T} \|z^0 - z^*\|_2^2 + \frac{c_2 r\kappa_x C_P^2 \sigma^2}{L^2 S}$$

for all $T \geq 0$, where c_1 and c_2 are numerical constants. It means the gradient complexity of achieving $\mathbb{E}[\|z^T - z^*\|_2] \leq \epsilon$ is

$$\mathcal{O}\left(r\kappa_x \log\left(\frac{C_P \|z^0 - z^*\|_2}{\epsilon}\right) \max\left\{1, \frac{r\kappa_x C_P^2 \sigma^2}{L^2 \epsilon^2}\right\}\right).$$

The proofs of Theorem 5.4 and Theorem 5.5 can be found in Appendix 11. As we mentioned before, the above $\tilde{O}\left(\frac{r^2 \kappa_x^2}{\epsilon^2}\right)$

rate is better than the stochastic local convergence counterpart in Theorem 4.4.

Again note that when $r = 2\kappa$, the rate is $\tilde{O}\left(\frac{(\kappa\kappa_x)^2}{\epsilon^2}\right)$ which is the gradient complexity of stochastic gradient descent directly on Φ although we do not have access to the stochastic oracle of Φ !

6. Experimental results

6.1. Quadratic functions

In this section, we empirically show the convergence behavior of GDA with different stepsize ratios on quadratic functions. We set $z^* = 0$ w.l.o.g. as it does not affect the convergence behavior. The matrices $A, B, C \in \mathbb{R}^{4 \times 4}$ are randomly generated and processed to satisfy Assumption 5.1. We choose $L = 100$, $\mu = 1$ in the beginning and compute μ_x after the matrices are sampled. We keep $\eta_y = \frac{1}{2L}$ and change the stepsize ratio by varying η_x .

Figure 2 shows the training curves under two independent samples for four different stepsize ratios. The orange curve uses the best stepsize ratio, 2κ , as suggest by our theorems. The red curve uses the stepsizes as in (Lin et al., 2020a) (note $200\kappa = 2\kappa^2$). According to Figure 2a, as the stepsize ratio increases from 2κ , the convergence becomes slower. For the red curve, since η_x is much smaller than η_y , x updates much more slowly than y . As a result, the training can be roughly splitted into two phases. During the first phase (first 1000 iterations in Figure 2a), the change of x can be ignored and y is updated to maximize f . Therefore the red curve goes up. During the second phase (after 1000 iterations in Figure 2a), y is updated to keep maximizing f for each x . Therefore essentially, x is trained to minimize the primal function $\Phi(x)$ and thus the curve goes down. However, such phenomenon does not exist for the orange curve, as both variables are basically simultaneously trained to the optimal point.

For the blue curve which uses a stepsize ratio even much smaller than 2κ , although it converges even faster than the orange curve as in Figure 2a, it sometimes does diverge as in Figure 2b. This is consistent with our lower bound in Theorem 4.1.

6.2. A non-quadratic function

Although our theory only applies to quadratic functions or local convergence, the empirical training behavior we observed above may exist for more general functions. In this section, we empirically run GDA on a non-quadratic function defined as follows. We first generate an instance of quadratic functions, f_0 , as in Section 6.1. Then we add a non-quadratic function of x only. Formally, we run GDA




Figure 2. Training curves of GDA on two independently generated quadratic functions satisfying Assumption 5.1 with four different stepsize ratios. The corresponding μ_x value is shown in the title of each plot.




Figure 3. Training curves of GDA on two independently generated non-quadratic functions defined in (10) with four different stepsize ratios. The corresponding μ_x value is shown in the title of each plot.

on the following function:

$$f(\mathbf{x}; \mathbf{y}) := f_0(\mathbf{x}; \mathbf{y}) + \frac{L}{n} \sum_{i=1}^n [\log(1 + \exp(a(x_i - b_i))) + \log(1 + \exp(-a(x_i - b_i)))]. \quad (10)$$

The training curves for different stepsize ratios are shown in Figure 3, with similar behaviors to those in Figure 2.

6.3. GAN training

In this section, we empirically show the convergence behavior of GANs trained using GDA with different stepsize ratios. The WGAN-GP model in (Gulrajani et al., 2017) was trained using ADAM, a variant of GDA, with $\eta_x = \eta_y = 0.0001$. Moreover, they ran multiple steps of inner maximization each outer minimization step. Our experiment uses the same WGAN-GP model. However, different from their algorithm, we use simultaneous GDA with the same number of gradient steps for both variables. Figure 4 shows the generated images of the learned generator and the evolution of the Wasserstein Distance which is also the primal function Φ . As we can see, simultaneous GDA with $\eta_x = \eta_y = 0.001$ is able to converge with a

high speed for both MNIST and CIFAR10. However, if we decrease η_x , i.e., increase the stepsize ratio, it converges with a lower speed. The convergent value is also larger. We will show more experimental results for different models under different stepsize ratios in the Appendix.

7. Conclusion

There exists a wide gap between existing theory on gradient-based minimax optimization which suggests a significantly large max/min stepsize ratio and state-of-the-art GAN training algorithms where similar learning rates are used for the min and max players. In this paper, we take one step towards understanding and bridging this gap by providing a tight local convergence analysis of GDA near a strict Stackelberg Equilibrium. In particular, we characterize the optimal choices of the stepsize values and resulting stepsize ratio and further provide a local convergence guarantee under these stepsize selection. We also extend the analysis to extra-gradient and stochastic GDA methods.

There remain several open problems and future directions




Figure 4. Evolution of Wasserstein Distance on MNIST (a) and CIFAR10 (b). For both MNIST and CIFAR10, we train WGAN-GP models using under different stepsize ratios $(\eta_x, \eta_y) = (0.001, 0.001), (0.0001, 0.001), (1e-05, 0.001)$.

following our paper. First, in the stochastic setting, the complexity bounds of GDA and EG for local convergence in Theorem 4.4 are larger than the bounds under quadratic objective functions in Theorem 5.5. Hence, an interesting future direction is to improve the bounds in Theorem 4.4. In addition, the extension of our results to optimistic GDA (OGDA) methods will be another interesting problem. Finally, our analysis focuses on local convergence and quadratic NC-SC objective functions. The problem of characterizing the best stepsize ratio for general NC-SC functions and NC-NC functions still remains open.

8. Acknowledgement

This work was supported, in part, by the MIT-IBM Watson AI Lab, ONR Grant N00014-20-1-2394, a Vannevar Bush fellowship from Office of the Secretary of Defense, and a CUHK Direct Grant for Research. The authors thank the anonymous reviewers for their constructive feedback and suggestions.

References

- Benaïm, M. and Hirsch, M. Mixed equilibria and dynamical systems arising from fictitious play in perturbed games. *Games and Economic Behavior*, 29:36–72, 1999.
- Chambolle, A. and Pock, T. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- Chen, G. H. and Rockafellar, R. T. Convergence rates in forward-backward splitting. *SIAM J. Optim.*, 7:421–444, 1997.
- Daskalakis, C., Ilyas, A., Syrgkanis, V., and Zeng, H. Training gans with optimism. *arXiv preprint arXiv:1711.00141*, 2017.
- Daskalakis, C., Skoulakis, S., and Zampetakis, M. The complexity of constrained min-max optimization. *arXiv preprint arXiv:2009.09623*, 2020.
- Du, S. and Hu, W. Linear convergence of the primal-dual gradient method for convex-concave saddle point problems without strong convexity. *ArXiv*, abs/1802.01504, 2019.
- Farnia, F. and Ozdaglar, A. Do gans always have nash equilibria? In *International Conference on Machine Learning*, pp. 3029–3039. PMLR, 2020.
- Farnia, F. and Ozdaglar, A. Train simultaneously, generalize better: Stability of gradient-based minimax learners. In *International Conference on Machine Learning*, pp. 3174–3185. PMLR, 2021.
- Fiez, T. and Ratliff, L. J. Local convergence analysis of gradient descent ascent with finite timescale separation. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=AWOSz_mMAPx.
- Fiez, T., Chasnov, B., and Ratliff, L. Implicit learning dynamics in stackelberg games: Equilibria characterization, convergence analysis, and empirical study. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3133–3144. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/fiez20a.html>.
- Gidel, G., Berard, H., Vignoud, G., Vincent, P., and Lacoste-Julien, S. A variational inequality perspec-

- tive on generative adversarial networks. *arXiv preprint arXiv:1802.10551*, 2018.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of wasserstein gans. In *NIPS*, 2017.
- Han, Y., Xie, G., and Zhang, Z. Lower complexity bounds of finite-sum optimization problems: The results and construction. *ArXiv*, abs/2103.08280, 2021.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017.
- Hommes, C. and Ochea, M. Multiple equilibria and limit cycles in evolutionary games with logit dynamics. *Games Econ. Behav.*, 74:434–441, 2012.
- Hsieh, Y.-G., Iutzeler, F., Malick, J., and Mertikopoulos, P. On the convergence of single-call stochastic extragradient methods. *arXiv preprint arXiv:1908.08465*, 2019.
- Ibrahim, A., Azizian, W., Gidel, G., and Mitliagkas, I. Lower bounds and conditioning of differentiable games. *ArXiv*, abs/1906.07300, 2019.
- Jin, C., Netrapalli, P., and Jordan, M. What is local optimality in nonconvex-nonconcave minimax optimization? In *International Conference on Machine Learning*, pp. 4880–4889. PMLR, 2020.
- Kong, W. and Monteiro, R. D. An accelerated inexact proximal point method for solving nonconvex-concave min-max problems. *arXiv preprint arXiv:1905.13433*, 2019.
- Korpelevich, G. M. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- Li, H., Tian, Y., Zhang, J., and Jadbabaie, A. Complexity lower bounds for nonconvex-strongly-concave min-max optimization. *ArXiv*, abs/2104.08708, 2021.
- Lin, T., Jin, C., and Jordan, M. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pp. 6083–6093. PMLR, 2020a.
- Lin, T., Jin, C., and Jordan, M. I. Near-optimal algorithms for minimax optimization, 2020b.
- Lu, S., Tsaknakis, I., Hong, M., and Chen, Y. Hybrid block successive approximation for one-sided non-convex min-max problems: algorithms and applications. *IEEE Transactions on Signal Processing*, 68:3676–3691, 2020.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Mertikopoulos, P., Lecouat, B., Zenati, H., Foo, C.-S., Chandrasekhar, V., and Piliouras, G. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. *arXiv preprint arXiv:1807.02629*, 2018a.
- Mertikopoulos, P., Papadimitriou, C., and Piliouras, G. Cycles in adversarial regularized learning. In *SODA*, 2018b.
- Mokhtari, A., Ozdaglar, A. E., and Pattathil, S. Convergence rate of $\mathcal{O}(1/k)$ for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 30(4):3230–3251, 2020.
- Nedić, A. and Ozdaglar, A. Subgradient methods for saddle-point problems. *Journal of Optimization Theory and Applications*, 142:205–228, 2009.
- Nemirovski, A. Prox-method with rate of convergence $\mathcal{O}(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1): 229–251, 2004.
- Nouiehed, M., Sanjabi, M., Huang, T., Lee, J. D., and Razaviyayn, M. Solving a class of non-convex min-max games using iterative first order methods. *arXiv preprint arXiv:1902.08297*, 2019.
- Ostrovskii, D. M., Lowy, A., and Razaviyayn, M. Efficient search of first-order nash equilibria in nonconvex-concave smooth min-max problems. *arXiv preprint arXiv:2002.07919*, 2020.
- Ouyang, Y. and Xu, Y. Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems. *Mathematical Programming*, pp. 1–35, 2019.
- Polyak, B. *Introduction to Optimization*. 07 2020.
- Rafique, H., Liu, M., Lin, Q., and Yang, T. Non-convex minimax optimization: Provable algorithms and applications in machine learning. *arXiv preprint arXiv:1810.02060*, 2018.
- Rakhlin, A. and Sridharan, K. Optimization, learning, and games with predictable sequences. *arXiv preprint arXiv:1311.1869*, 2013.

Thekumparampil, K. K., Jain, P., Netrapalli, P., and Oh, S. Efficient algorithms for smooth minimax optimization. *arXiv preprint arXiv:1907.01543*, 2019.

Tseng, P. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2):237–252, 1995.

Yadav, A., Shah, S., Xu, Z., Jacobs, D., and Goldstein, T. Stabilizing adversarial nets with prediction methods. *arXiv preprint arXiv:1705.07364*, 2017.

Zhang, J., Hong, M., and Zhang, S. On lower iteration complexity bounds for the convex concave saddle point problems. *arXiv: Optimization and Control*, 2019.

Zhang, S., Yang, J., Guzmán, C., Kiyavash, N., and He, N. The complexity of nonconvex-strongly-concave minimax optimization. *arXiv preprint arXiv:2103.15888*, 2021.

9. Proofs for lower bounds

In this section, we provide the proofs for all the lower bounds. Specifically, we prove Theorem 4.1 and Theorem 4.3.

9.1. Proof of Theorem 4.1

Proof of Theorem 4.1. Consider the following quadratic NC-SC function $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$:

$$\begin{aligned} f(x; y) &= -\frac{\mu}{2}(y - \kappa x)^2 + \frac{\mu}{2}(\kappa^2 - \kappa)x^2 \\ &= -\frac{\mu}{2}y^2 + Lxy - \frac{L}{2}x^2 \\ &= \frac{\mu}{2} \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} -\kappa & \kappa \\ \kappa & -\mu \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}. \end{aligned}$$

Note that $\Phi(x) = \max_{y \in \mathbb{R}} f(x; y) = \frac{\mu}{2}(\kappa^2 - \kappa)x^2 \geq \frac{\mu_x}{2}x^2$. It is easy to verify that f satisfies Assumption 5.1 and $(0, 0)$ is its unique differential Stackelberg Equilibrium with parameters (L, μ, μ_x) . Indeed, it is also the global minimax point. The M matrix for f is

$$M = \begin{pmatrix} \kappa & -\kappa \\ r\kappa & -r \end{pmatrix}.$$

Let λ_1 and λ_2 be the eigenvalues of M , we have

$$\text{Trace}(M) = \lambda_1 + \lambda_2 = \kappa - r.$$

There are two possible cases:

- If $r < \kappa$, we know $\lambda_1 + \lambda_2 > 0$ which implies at least one of them has a positive real part. Suppose $\text{Re}[\lambda_1] > 0$ without loss of generality. We know $|1 + \eta_x \lambda_1| > 1$ for every $\eta_x > 0$.
- If $r = \kappa$, $M = \begin{pmatrix} \kappa & -\kappa \\ \kappa^2 & -\kappa \end{pmatrix}$. We can directly compute its eigenvalues $\lambda_1 = i\sqrt{\kappa^3 - \kappa^2}$ and $\lambda_2 = -i\sqrt{\kappa^3 - \kappa^2}$ where $i = \sqrt{-1}$. We also have $|1 + \eta_x \lambda_1| = \sqrt{1 + \eta_x^2(\kappa^3 - \kappa^2)} > 1$ for every $\eta_x > 0$.

Then the power sequence $\{(I + \eta_x M)^k\}_{k=0}^\infty$ will diverge and we complete the proof. \square

9.2. Proof of Theorem 4.3

Proof of Theorem 4.3. Consider the following quadratic NC-SC function $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$

$$f(x, y) := -\frac{1}{2}Lx^2 + bxy - \frac{1}{2}\mu y^2,$$

where $b = \sqrt{\mu(L + \mu_x)}$. The primal function is $\Phi(x) := \max_y f(x, y) = \frac{1}{2}\mu_x x^2$. Therefore, it is straight-forward to verify that f satisfies Assumption 5.1 and $(0, 0)$ is its unique differential Stackelberg Equilibrium with parameters (L, μ, μ_x) . Assume $\kappa = L/\mu \geq 2$. We first show the lower bound of GDA to achieve $(0, 0)$. The M matrix is

$$M = \begin{pmatrix} L & -b \\ rb & -\mu r \end{pmatrix}.$$

Let λ_1 and λ_2 be the two eigenvalues of M , we have

$$\lambda_{1,2} = -\frac{1}{2}(\mu r - L) \pm \frac{1}{2}\sqrt{(\mu r - L)^2 - 4r\mu\mu_x}.$$

$$\begin{aligned} 0 \geq \lambda_1 &= -\frac{1}{2}(\mu r - L) \left(1 - \sqrt{1 - \frac{4r\mu\mu_x}{(\mu r - L)^2}} \right) \\ &\geq -\frac{2r\mu\mu_x}{r\mu - L} \geq -4\mu_x. \end{aligned}$$

Let s_1 be the corresponding eigenvalue of $I + \eta_x M$, it satisfies

$$0 \leq 1 - \frac{\mu_x}{rL} \leq s_1 = 1 + \eta_x \lambda_1 \leq 1.$$

We adversarially choose the initial point z^0 such that it is parallel to the eigenvector of $I + \eta_x M$ corresponding to s_1 . Note that for any $\delta > 0$, we can choose such a z^0 such that $\|z^0 - z^*\|_2 \leq \delta$. Then we have

$$z^T = (I + \eta_x M)^T z^0 = s_1^T z^0,$$

which implies

$$\|z^T\|_2 = s_1^T \|z^0\|_2 \geq \left(1 - \frac{1}{r\kappa_x}\right)^T \|z^0\|_2.$$

For EG, note that $1 + \eta_x \lambda_1^2 + \eta_x^2 \lambda_2^2 \geq 1 + \eta_x \lambda_1^2 = s_1$. The rest of proof is the same as that of GDA. \square

10. Analysis of the spectral properties

In this section, we prove the spectral properties of M (Lemma 5.2) and the spectral radius bound of GDA and GD (Lemma 5.3). We first present the following lemma and its corollary useful for analyzing spectral properties of a matrix.

Lemma 10.1. *Suppose $G = A + iB \in \mathbb{C}^{d \times d}$ where $A, B \in \mathbb{R}^{d \times d}$ are both symmetric real matrices and $i = \sqrt{-1}$. If there exist $a, b \in \mathbb{R}$ such that $aA + bB$ is strictly positive or negative definite, then G is invertible, i.e., $\det(G) \neq 0$.*

Proof of Lemma 10.1. If $\det(G) = 0$, we know there exists $z = z_0 + iz_1 \neq 0$ for some $z_0, z_1 \in \mathbb{R}^d$ such that $Gz = 0$. Then we have

$$z^\dagger Gz = z_0^\top A z_0 + z_1^\top A z_1 + i(z_0^\top B z_0 + z_1^\top B z_1) = 0,$$

where z^\dagger is the conjugate transpose of z . It means both the real and imaginary parts of $z^\dagger Gz$ are zero. So is any linear combination of them. Therefore

$$z_0^\top (aA + bB) z_0 + z_1^\top (aA + bB) z_1 = 0$$

which contradicts with the strict definiteness of $aA + bB$. Therefore $\det(G) \neq 0$. \square

Choosing $a = 1, b = 0$ or $a = 0, b = 1$, we obtain the following corollary immediately.

Corollary 10.2. *Suppose $G = A + iB \in \mathbb{C}^{d \times d}$ where $A, B \in \mathbb{R}^{d \times d}$ are both symmetric real matrices and $i = \sqrt{-1}$. If at least one of A and B is strictly positive or negative definite, then G is invertible, i.e., $\det(G) \neq 0$.*

Now we are ready to prove Lemma 5.2, the spectral properties of M .

Proof of Lemma 5.2. Before proving the five parts separately, we will first prove that $\lambda_0 \leq 0$ if $r > \kappa$, which is a slightly weaker version of part 5 (where we claim $\lambda_0 < 0$ if $r > \kappa$). We will prove it by contradiction. Assume that $r > \kappa$ but $\lambda_0 > 0$. The eigenvalue satisfies $\det(\lambda I - M) = 0$. Note by the theory of Schur Complement,

$$\begin{aligned} \det(\lambda I - M) &= \det \begin{pmatrix} C + \lambda I & B \\ -rB^\top & rA + \lambda I \end{pmatrix} \\ &= \det(rA + \lambda I) \det \left(C + \lambda I + B \left(A + \frac{\lambda}{r} I \right)^{-1} B^\top \right), \end{aligned}$$

where in the second equality we use the fact that $rA + \lambda I$ is invertible due to Corollary 10.2. Let $A = U\Lambda U^\top$ where $U \in \mathbb{R}^{m \times m}$ is orthonormal and $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_m)$ is diagonal. We know $\mu \leq \sigma_i \leq L$ for every $i \in [m]$. Let $\tilde{B} = BU$. Define

$$\begin{aligned} H(\lambda) &:= C + \lambda I + B \left(A + \frac{\lambda}{r} I \right)^{-1} B^\top \\ &= C + \lambda I + \tilde{B} D \tilde{B}^\top \end{aligned}$$

where $D = \text{diag}(d_1, \dots, d_m)$ is a diagonal matrix with

$$d_i = \frac{\sigma_i + \frac{\lambda_0}{r} - i \left(\frac{\lambda_1}{r}\right)}{\left(\sigma_i + \frac{\lambda_0}{r}\right)^2 + \left(\frac{\lambda_1}{r}\right)^2}$$

whose real part is strictly positive. We know $\det(\lambda I - M) = 0$ if and only if $\det(H(\lambda)) = 0$ as $\det(rA + \lambda I) \neq 0$. The real part of $H(\lambda)$ is

$$\text{Re}[H(\lambda)] = C + \lambda_0 I + \tilde{B} \text{Re}[D] \tilde{B}^\top.$$

Since $C \geq -LI$ and $\text{Re}[D] \geq 0$, we know that if $\lambda_0 > L$, then $\text{Re}[H(\lambda)] > 0$. Then by Corollary 10.2, $\det(H(\lambda)) \neq 0$. Therefore we must have $\lambda_0 \leq L$.

Now suppose $\lambda_1 \neq 0$. Then note the imaginary part of $H(\lambda)$ is

$$\text{Im}[H(\lambda)] = \lambda_1 \left(I - \frac{1}{r} \tilde{B} E \tilde{B}^\top \right),$$

where $E = \text{diag}(e_1, \dots, e_m)$ with

$$e_i = \frac{1}{\left(\sigma_i + \frac{\lambda_0}{r}\right)^2 + \left(\frac{\lambda_1}{r}\right)^2}.$$

Note

$$\text{Re}[H(\lambda)] + \frac{r\mu}{\lambda_1} \text{Im}[H(\lambda)] = C + \lambda_0 I + r\mu I + \tilde{B} F \tilde{B}^\top,$$

where $F = \text{diag}(f_1, \dots, f_m)$ with

$$f_i = \frac{\sigma_i + \frac{\lambda_0}{r} - \mu}{\left(\sigma_i + \frac{\lambda_0}{r}\right)^2 + \left(\frac{\lambda_1}{r}\right)^2} > 0.$$

Then we have $\text{Re}[H(\lambda)] + \frac{r\mu}{\lambda_1} \text{Im}[H(\lambda)] > 0$ which contradicts with Lemma 10.1. Therefore we must have $\lambda_1 = 0$, i.e., λ is real.

Then we can calculate

$$\begin{aligned} H(\lambda) &= C + \lambda_0 I + B \left(A + \frac{\lambda_0}{r} I \right)^{-1} B^\top \\ &= C + \lambda_0 I + B \left(A + \frac{\lambda_0}{r} I \right)^{-1} \left[\left(A + \frac{\lambda_0}{r} I \right) - \frac{\lambda_0}{r} I \right] A^{-1} B^\top \\ &= C + \lambda_0 I + B A^{-1} B^\top - \frac{\lambda_0}{r} B \left(A + \frac{\lambda_0}{r} I \right)^{-1} A^{-1} B^\top \\ &\geq C + B A^{-1} B^\top + \lambda_0 I - \frac{\lambda_0}{r\mu} B A^{-1} B^\top \\ &= \left(1 - \frac{\lambda_0}{r\mu} \right) (C + B A^{-1} B^\top) + \lambda_0 \left(I + \frac{C}{r\mu} \right) \\ &> 0, \end{aligned}$$

which again contradicts with Corollary 10.2, where in the last inequality we use the fact $1 - \lambda_0/r\mu > 0$ and $I + C/r\mu > 0$ if $r > \kappa$ and $0 < \lambda_0 \leq L$. Therefore we must have $\lambda_0 \leq 0$.

Now we are ready to prove the five parts one by one.

1. First note that if $|\lambda_1| > \sqrt{r}L$, we know $rA + \lambda I$ is still invertible without assuming $\lambda_0 > 0$ as $\lambda_1 \neq 0$, which means we are still able to define $H(\lambda)$ and obtain

$$\text{Im}[H(\lambda)] = \lambda_1 \left(I - \frac{1}{r} \tilde{B}E\tilde{B}^\top \right).$$

It is easy to see $\|E\|_2 \leq \left(\frac{r}{\lambda_1}\right)^2$. Noting $\|\tilde{B}\|_2 = \|B\|_2 \leq L$, we have

$$I - \frac{1}{r} \tilde{B}E\tilde{B}^\top > 0.$$

Then by Corollary 10.2, $\det(H(\lambda)) \neq 0$ which leads to contradiction. Therefore we must have $|\lambda_1| \leq \sqrt{r}L$.

2. Note if $\lambda_1 \neq 0$, we have

$$\text{Re}[H(\lambda)] + \frac{\alpha r \mu}{\lambda_1} \text{Im}[H(\lambda)] = C + \lambda_0 I + \alpha \mu r I + \tilde{B}F^{(\alpha)}\tilde{B}^\top,$$

where we choose $\alpha = \frac{r+\kappa}{2r}$ and $F^{(\alpha)} = \text{diag}(f_1^{(\alpha)}, \dots, f_m^{(\alpha)})$ with

$$f_i^{(\alpha)} = \frac{\sigma_i + \frac{\lambda_0}{r} - \alpha \mu}{\left(\sigma_i + \frac{\lambda_0}{r}\right)^2 + \left(\frac{\lambda_1}{r}\right)^2}.$$

Suppose $\lambda_0 > -(\alpha \mu r - L) = -\frac{\mu(r-\kappa)}{2}$. We have $C + \lambda_0 I + \alpha \mu r I > 0$. Also note that

$$\sigma_i + \frac{\lambda_0}{r} - \alpha \mu \geq \sigma_i - 2\alpha \mu + L/r \geq (1 - 2\alpha + \kappa/r)\mu = 0$$

which implies $f_i^{(\alpha)} \geq 0$ and thus $F^{(\alpha)}$ is positive semi-definite. Therefore $\text{Re}[H(\lambda)] + \frac{\alpha r \mu}{\lambda_1} \text{Im}[H(\lambda)] > 0$ which leads to contradiction by Lemma 10.1. Therefore $\lambda_0 \leq -\frac{\mu(r-\kappa)}{2}$.

3. Note that the spectral radius of a matrix is always bounded by its operator norm, i.e., $\lambda_0^2 + \lambda_1^2 \leq \|M\|_2^2$. Note

$$\begin{aligned} \left\| M \begin{pmatrix} x \\ y \end{pmatrix} \right\|_2 &= \left\| \begin{pmatrix} -Cx - By \\ rB^\top x - rAy \end{pmatrix} \right\|_2 \\ &= \sqrt{\| -Cx - By \|_2^2 + \| rB^\top x - rAy \|_2^2} \\ &\leq \sqrt{2(\|Cx\|_2^2 + \|By\|_2^2) + 2r^2(\|B^\top x\|_2^2 + \|Ay\|_2^2)} \\ &\leq \sqrt{2(r^2 + 1)L^2(\|x\|_2^2 + \|y\|_2^2)} \\ &\leq 2rL \left\| \begin{pmatrix} x \\ y \end{pmatrix} \right\|_2. \end{aligned}$$

Therefore $\|M\|_2 \leq 2rL$ and $\lambda_0^2 + \lambda_1^2 \leq 4r^2L^2$.

4. If $\lambda_1 = 0$ and $-\min\{L, \lambda_{\min}(C + BA^{-1}B^\top)\} < \lambda_0 \leq 0$, we have

$$\begin{aligned} H(\lambda) &= C + \lambda_0 I + B \left(A + \frac{\lambda_0}{r} I \right)^{-1} B^\top \\ &> \lambda_0 I + C + BA^{-1}B^\top \\ &> 0, \end{aligned}$$

which leads to contradiction.

5. This part can be directly shown by combining part 2 and part 4.

□

With the properties of M , we proceed to prove Lemma 5.3.

Proof of Lemma 5.3. Let $\lambda = \lambda_0 + i\lambda_1$ be an eigenvalue of M , we only need to show $\rho_{1,\lambda} = |1 + \eta_{\mathbf{x}}\lambda| \leq 1 - \frac{1}{16r\kappa_{\mathbf{x}}}$ and $\rho_{2,\lambda} = |1 + \eta_{\mathbf{x}}\lambda + \eta_{\mathbf{x}}^2\lambda^2| \leq 1 - \frac{1}{16r\kappa_{\mathbf{x}}}$, for any eigenvalue λ .

We first bound $\rho_{1,\lambda}$. There are two possible cases.

- If $\lambda_1 = 0$, by Lemma 5.2, we know $-2rL \leq \lambda_0 \leq -\mu_{\mathbf{x}}$. Therefore $\rho_{1,\lambda} = |1 + \eta_{\mathbf{x}}\lambda_0| \leq 1 - \frac{1}{4r\kappa_{\mathbf{x}}} \leq 1 - \frac{1}{16r\kappa_{\mathbf{x}}}$.
- If $\lambda_1 \neq 0$, we know $-2rL \leq \lambda_0 \leq -\frac{\mu(r-\kappa)}{2}$ and $|\lambda_1| \leq \sqrt{r}L$. Then we have

$$\begin{aligned} \rho_{1,\lambda}^2 &= (1 + \eta_{\mathbf{x}}\lambda_0)^2 + \eta_{\mathbf{x}}^2\lambda_1^2 \\ &\leq \left(1 + \frac{-\mu(r-\kappa)}{8rL}\right)^2 + \frac{1}{16r} \\ &\leq \left(1 - \frac{1}{16\kappa}\right)^2 + \frac{1}{32\kappa} \leq 1 - \frac{1}{16\kappa}. \end{aligned}$$

Therefore

$$\rho_{1,\lambda} < \sqrt{1 - \frac{1}{16\kappa}} \leq 1 - \frac{1}{32\kappa} \leq 1 - \frac{1}{16r\kappa_{\mathbf{x}}}.$$

Next, we bound $\rho_{2,\lambda}$. Similarly, there are also two possible cases.

- If $\lambda_1 = 0$, noting that $-2rL \leq \lambda_0 \leq -\mu_{\mathbf{x}}$ and $|\eta_{\mathbf{x}}\lambda_0| \leq 1/2$, we have

$$\rho_{2,\lambda} \triangleq |1 + \eta_{\mathbf{x}}\lambda_0 + \eta_{\mathbf{x}}^2\lambda_0^2| \leq 1 + \frac{1}{2}\eta_{\mathbf{x}}\lambda_0 \leq 1 - \frac{1}{8r\kappa_{\mathbf{x}}}.$$

- If $\lambda_1 \neq 0$,

$$\begin{aligned} \rho_{2,\lambda}^2 &= |1 + \eta_{\mathbf{x}}\lambda + \eta_{\mathbf{x}}^2\lambda^2|^2 \\ &= |1 + \eta_{\mathbf{x}}\lambda_0 + \eta_{\mathbf{x}}^2(\lambda_0^2 - \lambda_1^2) + i\eta_{\mathbf{x}}\lambda_1(1 + 2\eta_{\mathbf{x}}\lambda_0)|^2 \\ &= |1 + \eta_{\mathbf{x}}\lambda_0 + \eta_{\mathbf{x}}^2(\lambda_0^2 - \lambda_1^2)|^2 + \eta_{\mathbf{x}}^2\lambda_1^2(1 + 2\eta_{\mathbf{x}}\lambda_0)^2 \\ &\leq \left(1 + \frac{1}{2}\eta_{\mathbf{x}}\lambda_0\right)^2 + \eta_{\mathbf{x}}^2\lambda_1^2 \\ &\leq \left(1 - \frac{\mu}{4}\eta_{\mathbf{x}}(r - \kappa)\right)^2 + \frac{1}{16r} \\ &\leq 1 - \frac{1}{64\kappa}. \end{aligned}$$

Therefore

$$\rho_{2,\lambda} < \sqrt{1 - \frac{1}{64\kappa}} \leq 1 - \frac{1}{128\kappa} \leq 1 - \frac{1}{64r\kappa_{\mathbf{x}}}.$$

Noting that $\rho_1 = \max_{\lambda} \rho_{1,\lambda}$ and $\rho_2 = \max_{\lambda} \rho_{2,\lambda}$, we complete the proof. □

11. Proofs of the convergence results for quadratic NC-SC functions

11.1. Proof of Theorem 5.4

In this section, we show the linear convergence results for quadratic functions. Given Lemma 5.3, if M is diagonalizable, it is straight-forward to obtain a linear convergence rate. However, in some rare cases, M may not be diagonalizable and things become trickier. Therefore, we further need the following lemma.

Lemma 11.1. *Let $A \in \mathbb{C}^{n \times n}$ be a (possibly nondiagonalizable) matrix and its spectral radius is $\rho < 1$. Suppose its Jordan decomposition is $A = P^{-1}JP$, then we have*

$$\|A^k\|_2 \leq sC_P k^{s-1} \rho^{k-s+1},$$

where s is the size of the largest Jordan block in J and $C_P = \|P^{-1}\|_2 \|P\|_2$ is the condition number of P .

Proof of Lemma 11.1. Note that

$$\|A^k\|_2 = \|P^{-1}J^kP\|_2 \leq C_P \|J^k\|_2.$$

So it suffices to bound $\|J^k\|_2$. Use $J(\lambda, m)$ to denote the Jordan block with eigenvalue λ and size m . We know $[J(\lambda, m)]^k = (a_{i,j})$ where $a_{i,j} = \binom{k}{j-i} \lambda^{k-(j-i)}$ (here we define $\binom{N}{n} = 0$ if $n < 0$ or $n > N$). Then since $|\lambda| \leq \rho < 1$, we can bound $\|[J(\lambda, m)]^k\|_2 \leq mk^{m-1} \rho^{k-m+1} \leq sk^{s-1} \rho^{k-s+1}$. Then we have $\|J^k\|_2 \leq sk^{s-1} \rho^{k-s+1}$ and complete the proof. \square

With Lemma 11.1, it is straight forward to show Theorem 5.4 and we omit the proof.

11.2. SGDA on quadratic functions

Now we show Theorem 5.5, the converge result for SGDA on quadratic functions.

Proof of Theorem 5.5. We first prove the convergence for GDA. Denote $\tilde{z} = z - z^*$ and $\xi^k = (\xi_x^k, \xi_y^k)$ where ξ_x^k, ξ_y^k are defined in (3). The dynamics of GDA is

$$\tilde{z}^{k+1} = (I + \eta_x M) \tilde{z}^k + \eta_x \xi^k,$$

where we can bound $\mathbb{E} \left[\|\xi^k\|_2^2 \right] \leq 2r^2 \sigma^2 / S$. Note that we have

$$\left\| \prod_{k \in S} (I + \eta_x M) \right\|_2 \leq C_P \left(1 - \frac{1}{64r\kappa_x} \right)^{|S|}.$$

Then we can bound

$$\begin{aligned} \mathbb{E} \left[\|\tilde{z}^T\|_2^2 \right] &\leq \left\| (I + \eta_x M)^T \right\|_2^2 \|\tilde{z}^0\|_2^2 + \eta_x^2 \sum_{i=0}^{T-1} \left\| (I + \eta_x M)^{T-i-1} \right\|_2^2 \mathbb{E}[\|\xi^i\|_2^2] \\ &\leq C_P^2 \left(1 - \frac{1}{64r\kappa_x} \right)^{2T} \|\tilde{z}^0\|_2^2 + \eta_x^2 C_P^2 \frac{2r^2 \sigma^2}{S} \sum_{i=0}^{T-1} \left(1 - \frac{1}{64r\kappa_x} \right)^{T-i-1} \\ &\leq C_P^2 \left(1 - \frac{1}{64r\kappa_x} \right)^{2T} \|\tilde{z}^0\|_2^2 + \frac{8r\kappa_x C_P^2 \sigma^2}{L^2 S}, \end{aligned}$$

where the first inequality is due to the independence of $\{\xi^k\}_{k \geq 0}$. If we require $\mathbb{E} \left[\|\tilde{z}^T\|_2^2 \right] \leq \epsilon^2$, we only need to have

$$\begin{aligned} T &= O \left(r\kappa_x \log \left(\frac{C_P \|\tilde{z}^0\|_2}{\epsilon^2} \right) \right), \\ S &= O \left(\frac{r\kappa_x C_P^2 \sigma^2}{L^2 \epsilon^2} \right). \end{aligned}$$

This implies a gradient complexity bound of

$$T \cdot S = O\left(\frac{r^2 \kappa_{\mathbf{x}}^2 C_P^2 \sigma^2}{L^2 \epsilon^2} \log\left(\frac{C_P \|\tilde{\mathbf{z}}^0\|_2}{\epsilon^2}\right)\right).$$

Now we prove the convergence result for EG. Note that the dynamics of EG is

$$\begin{aligned} \tilde{\mathbf{z}}^{k+1} &= \tilde{\mathbf{z}}^k + \eta_{\mathbf{x}} M \tilde{\mathbf{z}}^{k+1/2} + \eta_{\mathbf{x}} \boldsymbol{\xi}^{k+1/2} \\ &= \tilde{\mathbf{z}}^k + \eta_{\mathbf{x}} M \left(\tilde{\mathbf{z}}^k + \eta_{\mathbf{x}} M \tilde{\mathbf{z}}^k + \eta_{\mathbf{x}} \boldsymbol{\xi}^k \right) + \eta_{\mathbf{x}} \boldsymbol{\xi}^{k+1/2} \\ &= (I + \eta_{\mathbf{x}} M + \eta_{\mathbf{x}}^2 M^2) \tilde{\mathbf{z}}^k + \eta_{\mathbf{x}} \boldsymbol{\xi}^{k+1/2} + \eta_{\mathbf{x}}^2 M \boldsymbol{\xi}^k. \end{aligned}$$

Since $\eta_{\mathbf{x}} \|M\|_2 \leq 1$ by Lemma 5.2, we can bound $\mathbb{E} \left[\left\| \boldsymbol{\xi}^{k+1/2} + \eta_{\mathbf{x}} M \boldsymbol{\xi}^k \right\|_2^2 \right] \leq 4r^2 \sigma^2 / S$. Also note that by Lemma 5.3, the transition matrix of EG also satisfies $\|I + \eta_{\mathbf{x}} M + \eta_{\mathbf{x}}^2 M^2\|_2 \leq 1 - \frac{1}{64r\kappa_{\mathbf{x}}}$. Then the rest of proof is the same as that for GDA. \square

12. Proof of Theorem 4.2

In this section, we prove our main theorem (Theorem 4.2). Since the function f is locally approximately quadratic, we are essentially extending the results of quadratic functions to approximately quadratic functions. So the proofs in this section build upon the lemmas and proofs for quadratic functions in Appendix 10 and Appendix 11.

Note that since the Hessian of f is continuous under our assumption, for any $\Delta > 0$, there exist some $\delta_m > 0$ such that if $\|\mathbf{z} - \mathbf{z}^*\|_2 \leq \delta_m$, we have

$$\max \{ \|\nabla_{\mathbf{x}\mathbf{x}}^2 f(\mathbf{x}; \mathbf{y}) - C\|_2, \|\nabla_{\mathbf{x}\mathbf{y}}^2 f(\mathbf{x}; \mathbf{y}) - B\|_2, \|\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{x}; \mathbf{y}) + A\|_2 \} \leq \Delta.$$

Before proving the theorem, we first present the following useful lemma.

Lemma 12.1. *For any two points \mathbf{z}_1 and \mathbf{z}_2 such that $\max\{\|\mathbf{z}_1 - \mathbf{z}^*\|_2, \|\mathbf{z}_2 - \mathbf{z}^*\|_2\} \leq \delta_m$, there exists a symmetric matrix $\tilde{H} = \begin{pmatrix} \tilde{C} & \tilde{B} \\ \tilde{B}^\top & -\tilde{A} \end{pmatrix}$ such that*

$$\nabla f(\mathbf{z}_2) - \nabla f(\mathbf{z}_1) = \tilde{H} \cdot (\mathbf{z}_2 - \mathbf{z}_1),$$

where we also have

$$\max \left\{ \|C - \tilde{C}\|_2, \|B - \tilde{B}\|_2, \|A - \tilde{A}\|_2 \right\} \leq \Delta.$$

Proof of Lemma 12.1. For any $0 \leq t \leq 1$, denote $\mathbf{z}_t = (1-t)\mathbf{z}_1 + t\mathbf{z}_2$. We can denote the hessian at \mathbf{z}_t as

$$\nabla^2 f(\mathbf{z}_t) = \begin{pmatrix} C_t & B_t \\ B_t^\top & -A_t \end{pmatrix}.$$

By Jensen's inequality, we know $\|\mathbf{z}_t - \mathbf{z}^*\|_2 \leq \delta_m$. Therefore

$$\max \{ \|C - C_t\|_2, \|B - B_t\|_2, \|A - A_t\|_2 \} \leq \Delta. \quad (11)$$

Note that

$$\nabla f(\mathbf{z}_2) - \nabla f(\mathbf{z}_1) = \int_0^1 \nabla^2 f(\mathbf{z}_t) \cdot (\mathbf{z}_2 - \mathbf{z}_1) dt =: \tilde{H} \cdot (\mathbf{z}_2 - \mathbf{z}_1),$$

where we define

$$\tilde{H} := \int_0^1 \nabla^2 f(z_t) dt = \mathbb{E}_{t \sim \mathcal{U}(0,1)} [\nabla^2 f(z_t)],$$

where $\mathcal{U}(0,1)$ denotes the uniform distribution over $[0,1]$. Then we have

$$\tilde{A} = \mathbb{E}_{t \sim \mathcal{U}(0,1)}[A_t], \quad \tilde{B} = \mathbb{E}_{t \sim \mathcal{U}(0,1)}[B_t], \quad \tilde{C} = \mathbb{E}_{t \sim \mathcal{U}(0,1)}[C_t].$$

Taking expectation of the (11) and applying Jensen's inequality, we complete the proof. \square

With Lemma 12.1, we are ready to prove Theorem 4.2.

Proof of Theorem 4.2. Since $\nabla f(z^*) = \mathbf{0}$, by Lemma 12.1, for any point z such that $\|z - z^*\|_2 \leq \delta_m$, there exists a matrix $\tilde{H} = \begin{pmatrix} \tilde{C} & \tilde{B} \\ \tilde{B}^\top & -\tilde{A} \end{pmatrix}$ such that

$$\nabla f(z) = \tilde{H} \cdot (z - z^*).$$

We will use $\tilde{H}_k = \begin{pmatrix} \tilde{C}_k & \tilde{B}_k \\ \tilde{B}_k^\top & -\tilde{A}_k \end{pmatrix}$ to such a matrix for z^k , i.e., $\nabla f(z^k) = \tilde{H}_k \cdot (z^k - z^*)$.

Let us first prove the convergence rate of GDA. Choosing $\Delta \leq \frac{c_0 L}{r \kappa_{\mathbf{x}} C_P}$ for some small enough numerical constant c_0 and initialize z^0 such that $\|z^0 - z^*\|_2 \leq \delta_m / C_P$. We claim that $\|z^t - z^*\|_2 \leq \delta_m$ for all $t \geq 0$. We will prove this claim by induction. First note that it is trivially true when $t = 0$. Let us assume it is true for any $t \leq T$ and try to show it also holds for $t = T + 1$.

Denote $\tilde{z} = z - z^*$. We have for any $k \leq T$,

$$\tilde{z}^{k+1} = (I + \eta_{\mathbf{x}} \tilde{H}_k) \tilde{z}^k = (I + \eta_{\mathbf{x}} M + E_k) \tilde{z}^k,$$

where

$$E_k := \eta_{\mathbf{x}} (\tilde{H}_k - M) = \eta_{\mathbf{x}} \begin{pmatrix} -(\tilde{C}_k - C) & -(\tilde{B}_k - B) \\ r(\tilde{B}_k - B)^\top & r(\tilde{A}_k - A) \end{pmatrix}.$$

It is easy to bound $\|E_k\|_2 \leq 2\eta_{\mathbf{x}} r \Delta$. Suppose $I + \eta_{\mathbf{x}} M = P^{-1} \Lambda P$ is diagonalizable, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_{m+n})$ and we have shown $|\lambda_i| \leq 1 - \frac{1}{64r\kappa_{\mathbf{x}}}$. Then we can bound that

$$\begin{aligned} \left\| \prod_{k=0}^T (I + \eta_{\mathbf{x}} M + E_k) \right\|_2 &\leq \sum_{i=0}^{T+1} \binom{T+1}{i} C_P^{i+1} \left(1 - \frac{1}{64r\kappa_{\mathbf{x}}}\right)^{T+1-i} (2\eta_{\mathbf{x}} r \Delta)^i \\ &\leq C_P \left(1 - \frac{1}{64r\kappa_{\mathbf{x}}} + 2\eta_{\mathbf{x}} r \Delta C_P\right)^{T+1} \\ &\leq C_P \left(1 - \frac{1}{128r\kappa_{\mathbf{x}}}\right)^{T+1}. \end{aligned}$$

Then we have

$$\|z^{T+1} - z^*\|_2 \leq C_P \left(1 - \frac{1}{128r\kappa_{\mathbf{x}}}\right)^{T+1} \|z^0 - z^*\|_2 \leq C_P \cdot \frac{\delta_m}{C_P} = \delta_m.$$

Therefore we have proved the claim as well as obtained the convergence rate for GDA.

Now we prove the convergence rate for EG. Note that the dynamics of EG can be written as

$$\begin{aligned}\tilde{\mathbf{z}}^{k+1} &= \tilde{\mathbf{z}}^k + \eta_{\mathbf{x}} \tilde{H}_{k+1/2} \cdot \tilde{\mathbf{z}}^{k+1/2} \\ &= \left(1 + \eta_{\mathbf{x}} \tilde{H}_{k+1/2} + \eta_{\mathbf{x}}^2 \tilde{H}_{k+1/2} \tilde{H}_k\right) \tilde{\mathbf{z}}^k \\ &=: \left(1 + \eta_{\mathbf{x}} M + \eta_{\mathbf{x}}^2 M^2 + F_k\right) \tilde{\mathbf{z}}^k,\end{aligned}$$

where we have

$$\|F_k\|_2 = \left\| \left(1 + \eta_{\mathbf{x}} M\right) E_{k+1/2} + E_k \left(\eta_{\mathbf{x}} M + E_{k+1/2}\right) \right\|_2 \leq 4 \|E_k\|_2,$$

where in the inequality we use the fact $\|E_{k+1/2}\|_2 \leq 1$ and $\eta_{\mathbf{x}} \|M\|_2 \leq 1$ due to Lemma 5.2. Also note that by Lemma 5.3, the transition matrix of EG also satisfies $\|I + \eta_{\mathbf{x}} M + \eta_{\mathbf{x}}^2 M^2\|_2 \leq 1 - \frac{1}{64r\kappa_{\mathbf{x}}}$. The rest of proof for EG is similar to that of GDA. □

13. Proof of Theorem 4.4

Proof of Theorem 4.4. We use the same notation as in the proof of Theorem 5.5. We first prove the convergence for GDA. Similar to the proof of Theorem 4.2, we choose $\Delta \leq \frac{c_0 L}{r\kappa_{\mathbf{x}} C_P}$ for some small enough numerical constant c_0 and initialize \mathbf{z}^0 such that $\|\mathbf{z}^0 - \mathbf{z}^*\|_2 \leq \delta_m / 2C_P$. We claim that with probability at least $1 - p$ where p is an arbitrarily small constant, we have $\|\mathbf{z}^t - \mathbf{z}^*\|_2 \leq \delta_m$ for any $t \leq T$ where T is the maximum number of iterations. Denote the event

$$\mathcal{E} = \left\{ \boldsymbol{\xi} \left| \eta_{\mathbf{x}} C_P \sum_{k=0}^{T-1} \|\boldsymbol{\xi}^k\|_2 \leq \epsilon/2 \right. \right\},$$

where $\boldsymbol{\xi} = \{\boldsymbol{\xi}^k\}_{k=0}^{T-1}$ and we can bound $\mathbb{E} \left[\|\boldsymbol{\xi}^k\|_2^2 \right] \leq 2r^2 \sigma^2 / S$. Choosing

$$S = \frac{CT^2 C_P^2 \sigma^2}{L^2 \epsilon^2}$$

for some large enough constant C , by Markov inequality, we have

$$\begin{aligned}\mathbb{P}(\boldsymbol{\xi} \notin \mathcal{E}) &\leq \frac{2\eta_{\mathbf{x}} C_P}{\epsilon} \sum_{k=0}^{T-1} \mathbb{E}[\|\boldsymbol{\xi}^k\|_2] \\ &\leq \frac{2C_P}{4rL\epsilon} \cdot T \cdot \sqrt{2r\sigma} / \sqrt{S} \\ &\leq p.\end{aligned}$$

In other words, with probability at least $1 - p$, we have $\boldsymbol{\xi} \in \mathcal{E}$. We will condition on this event in the following proof.

Then we can use the induction technique similar to the proof of Theorem 4.2. Suppose that $\|\tilde{\mathbf{z}}^k\| \leq \delta_m$ for all $k \leq K$ where $K < T$. The dynamics of GDA is

$$\tilde{\mathbf{z}}^{k+1} = (I + \eta_{\mathbf{x}} M + E_k) \tilde{\mathbf{z}}^k + \eta_{\mathbf{x}} \boldsymbol{\xi}^k,$$

where E_k is defined in the proof of Theorem 4.2. Note that in this case, following the proof of Theorem 4.2, we can bound

$$\left\| \prod_{k \in \mathcal{I}} (I + \eta_{\mathbf{x}} M + E_k) \right\|_2 \leq C_P \left(1 - \frac{1}{128r\kappa_{\mathbf{x}}}\right)^{|\mathcal{I}|}, \text{ for any index set } \mathcal{I}.$$

The problem is more difficult than that in Theorem 5.5 because E_k could potentially depend on the gradient noise $\{\xi^i\}_{i < k}$. Conditioned on the event $\xi \in \mathcal{E}$, we can bound that

$$\begin{aligned} \|\tilde{z}^{K+1}\|_2 &\leq \left\| \prod_{k=0}^K (I + \eta_{\mathbf{x}} M + E_k) \right\|_2 \|\tilde{z}^0\|_2 + \eta_{\mathbf{x}} \sum_{i=0}^K \left\| \prod_{k=i}^K (I + \eta_{\mathbf{x}} M + E_k) \right\|_2 \|\xi^i\|_2 \\ &\leq C_P \left(1 - \frac{1}{128r\kappa_{\mathbf{x}}}\right)^{K+1} \|\tilde{z}^0\|_2 + \eta_{\mathbf{x}} C_P \sum_{i=0}^{T-1} \|\xi^i\|_2 \\ &\leq C_P \|\tilde{z}^0\|_2 + \epsilon/2 \leq \delta_m, \end{aligned}$$

where in the last inequality we use $\epsilon \leq \delta_m$ which holds because otherwise the initialization point is already ϵ -approximately optimal. Therefore the induction hypothesis also holds for $K + 1$. Then we can conclude that $\|\tilde{z}^k\|_2 \leq \delta_m$ for all $0 \leq k \leq T - 1$, i.e., the iterates never escape from the neighborhood around z^* under the event $\xi \in \mathcal{E}$.

From the above inequalities, when $K = T - 1$, we get

$$\|\tilde{z}^T\|_2 \leq C_P \left(1 - \frac{1}{32r\kappa_{\mathbf{x}}}\right)^T \|\tilde{z}^0\|_2 + \epsilon/2.$$

If we require $\|\tilde{z}^T\|_2 \leq \epsilon$, we only need to have

$$T = O\left(r\kappa_{\mathbf{x}} \log\left(\frac{C_P \|\tilde{z}^0\|_2}{\epsilon}\right)\right),$$

which implies

$$S = O\left(\frac{r^2 \kappa_{\mathbf{x}}^2 C_P^2 \sigma^2}{L^2 \epsilon^2} \log^2\left(\frac{C_P \|\tilde{z}^0\|_2}{\epsilon}\right)\right).$$

This implies a gradient complexity bound of

$$T \cdot S = O\left(\frac{r^3 \kappa_{\mathbf{x}}^3 C_P^2 \sigma^2}{L^2 \epsilon^2} \log^3\left(\frac{C_P \|\tilde{z}^0\|_2}{\epsilon}\right)\right).$$

Now we prove the convergence result for EG. Note that the dynamics of EG is

$$\begin{aligned} \tilde{z}^{k+1} &= \tilde{z}^k + (\eta_{\mathbf{x}} M + E_{k+1/2}) \tilde{z}^{k+1/2} + \eta_{\mathbf{x}} \xi^{k+1/2} \\ &= \tilde{z}^k + (\eta_{\mathbf{x}} M + E_{k+1/2}) \left(\tilde{z}^k + (\eta_{\mathbf{x}} M + E_k) \tilde{z}^k + \eta_{\mathbf{x}} \xi^k \right) + \eta_{\mathbf{x}} \xi^{k+1/2} \\ &= (I + \eta_{\mathbf{x}} M + \eta_{\mathbf{x}}^2 M^2 + F_k) \tilde{z}^k + \eta_{\mathbf{x}} \xi^{k+1/2} + \eta_{\mathbf{x}} (\eta_{\mathbf{x}} M + E_{k+1/2}) \xi^k. \end{aligned}$$

Note that it is straight forward to bound $\|\eta_{\mathbf{x}} M + E_{k+1/2}\|_2 \leq 1$. Then we can bound $\mathbb{E} \left[\left\| \xi^{k+1/2} + (\eta_{\mathbf{x}} M + E_{k+1/2}) \xi^k \right\|_2 \right] \leq 4r\sigma/\sqrt{S}$. Then the rest of proof is similar to that for GDA. In particular, we can define the event

$$\mathcal{E} = \left\{ \xi \mid \eta_{\mathbf{x}} C_P \sum_{k=0}^{T-1} (\|\xi^k\|_2 + \|\xi^{k+1/2}\|_2) \leq \epsilon/2 \right\}$$

and show it holds with high probability if choosing a large enough S . Following the proof of Theorem 4.2 for EG, we can also show that

$$\left\| \prod_{k \in \mathcal{I}} (I + \eta_{\mathbf{x}} M + \eta_{\mathbf{x}}^2 M^2 + F_k) \right\|_2 \leq C_P \left(1 - \frac{c_0}{r\kappa_{\mathbf{x}}}\right)^{|\mathcal{I}|},$$

for some small enough numerical constant c_0 . Then the the following steps are straight forward. □



Figure 5. Generated images of the convergent generators for MNIST (a) and CIFAR10 (b). All the generators are trained using the WGAN-GP model with SGDA. The stepsize choices (η_x, η_y) are set to be $(0.001, 0.001)$, $(0.0001, 0.001)$, and $(1e-5, 0.001)$ from the first row to the third row.

14. Additional experimental results

We present additional experimental results in this section. In Figure 4, we showed the training curves of SGDA on MNIST and CIFAR10 with different stepsize ratios. Now we show the generated images of the trained generator in Figure 5. As we can clearly see, a larger stepsize ratio not only makes it converging more slowly, but also results in a worse generator for both MNIST and CIFAR10 datasets. This is also shown in by the evolution FID score during training for CIFAR10 in Figure 6. All the experimental results so far use the WGAN-GP model. We will show that another GAN model, DCGAN, has similar training behaviors. Figure 7 shows the evolution of the Wasserstein Distance and FID score under different stepsize choices for DCGAN and CIFAR10. Figure 8 shows the generated images in this setting. As we can see, the results are similar to those for WGAN-GP.






Figure 6. Evolution of FID score on CIFAR10. We train WGAN-GP models under different stepsize ratios $(\eta_x, \eta_y) = (0.001, 0.001), (0.0001, 0.001), (1e-05, 0.001)$.



(a)



(b)

Figure 7. Evolution of Wasserstein Distance (a) and FID score (b) for CIFAR10. All the generators are trained using the DCGAN with SGDA under different stepsize choices $(\eta_x, \eta_y) = (0.001, 0.001), (0.0001, 0.001), (1e-5, 0.001)$.



Figure 8. Generated images of the convergent generators for DCGAN and CIFAR10. The stepsize choices (η_x, η_y) are set to be $(0.001, 0.001)$, $(0.0001, 0.001)$, and $(1e - 5, 0.001)$ from the first row to the third row.