PROXIMAL GRADIENT DESCENT-ASCENT: VARIABLE CONVERGENCE UNDER KŁ GEOMETRY

Anonymous authors
Paper under double-blind review

ABSTRACT

The gradient descent-ascent (GDA) algorithm has been widely applied to solve minimax optimization problems. In order to achieve convergent policy parameters for minimax optimization, it is important that GDA generates convergent variable sequences rather than convergent sequences of function value or gradient norm. However, the variable convergence of GDA has been proved only under convexity geometries, and it lacks understanding in general nonconvex minimax optimization. This paper fills such a gap by studying the convergence of a more general proximal-GDA for regularized nonconvex-strongly-concave minimax optimization. Specifically, we show that proximal-GDA admits a novel Lyapunov function, which monotonically decreases in the minimax optimization process and drives the variable sequences to a critical point. By leveraging this Lyapunov function and the KŁ geometry that parameterizes the local geometries of general nonconvex functions, we formally establish the variable convergence of proximal-GDA to a certain critical point x^*, i.e., $x_t \to x^*, y_t \to y^*(x^*)$. Furthermore, over the full spectrum of the KŁ-parameterized geometry, we show that proximal-GDA achieves different types of convergence rates ranging from sublinear convergence up to finite-step convergence, depending on the geometry associated with the KŁ parameter. This is the first theoretical result on the variable convergence for nonconvex minimax optimization.

1 INTRODUCTION

Minimax optimization is a classical optimization framework that has been widely applied in various modern machine learning applications, including game theory Ferreira et al. (2012), generative adversarial networks (GANs) Goodfellow et al. (2014), adversarial training Sinha et al. (2017), reinforcement learning Qiu et al. (2020), imitation learning Ho and Ermon (2016); Song et al. (2018), etc. A typical minimax optimization problem is shown below, where f is a differentiable function.

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

A popular algorithm for solving the above minimax problem is gradient descent-ascent (GDA), which performs a gradient descent update on the variable x and a gradient ascent update on the variable y alternatively in each iteration. Under the alternation between descent and ascent updates, it is much desired that GDA generates sequences of variables that *converge* to a certain optimal point, i.e., the minimax players obtain convergent optimal policies. In the existing literature, many studies have established the convergence of GDA-type algorithms under various global geometries of the objective function, e.g., convex-concave geometry (f is convex in x and concave in y) Nedić and Ozdaglar (2009), bi-linear geometry Neumann (1928); Robinson (1951) and Polyak-Łojasiewicz (PŁ) geometry Nouiehed et al. (2019); Yang et al. (2020). Some other work studied GDA under stronger global geometric conditions of f such as convex-strongly-concave geometry Du and Hu (2019) and strongly-convex-strongly-concave geometry Mokhtari et al. (2020); Zhang and Wang (2020), under which GDA is shown to generate convergent variable sequences. However, these special global function geometries do not cover modern machine learning problems that usually have complex models and nonconvex geometry.

Recently, many studies characterized the convergence of GDA in nonconvex minimax optimization, where the objective function is nonconvex in x. Specifically, Lin et al. (2020); Nouiehed et al.

(2019); Xu et al. (2020b) studied the convergence of GDA in the nonconvex-concave setting whereas Lin et al. (2020); Xu et al. (2020b) focused on the nonconvex-strongly-concave setting. In these general nonconvex settings, it is shown that GDA converges to a certain stationary point at a sublinear rate, i.e., $\|G(x_t)\| \leq t^{-\alpha}$ for some $\alpha > 0$, where $G(x_t)$ corresponds to a certain notion of gradient. Although such a gradient convergence result implies the stability of the algorithm, namely, $\lim_{t\to\infty} \|x_{t+1} - x_t\| = 0$, it does not guarantee the convergence of the variable sequences $\{x_t\}_t, \{y_t\}_t$ generated by GDA. So far, the variable convergence of GDA has not been established for nonconvex problems, but only under (strongly) convex function geometries that are mentioned previously Du and Hu (2019); Mokhtari et al. (2020); Zhang and Wang (2020). Therefore, we want to ask the following fundamental question:

- *Q1: Does GDA have guaranteed* **variable convergence** *in nonconvex minimax optimization? If so, where do they converge to?*

In fact, proving the variable convergence of GDA in the nonconvex setting is highly nontrivial due to the following reasons: 1) the algorithm alternates between a minimization step and a maximization step; 2) It is well understood that strong global function geometry leads to the convergence of GDA. However, in general nonconvex setting, the objective functions typically do not have an amenable global geometry. Instead, they may satisfy different types of local geometries around the critical points. Hence, it is natural and much desired to exploit the local geometries of functions in analyzing the convergence of GDA. The Kurdyka-Łojasiewicz (KŁ) geometry provides a broad characterization of such a local geometry for nonconvex functions.

The Kurdyka-Łojasiewicz (KŁ) geometry (see Section 2 for details) Bolte et al. (2007; 2014) parameterizes a broad spectrum of the *local nonconvex geometries* and has been shown to hold for a broad class of practical functions. Moreover, it also generalizes other global geometries such as strong convexity and PŁ geometry. In the existing literature, the KŁ geometry has been exploited extensively to analyze the convergence rate of various gradient-based algorithms in nonconvex optimization, e.g., gradient descent Attouch and Bolte (2009); Li et al. (2017) and its accelerated version Zhou et al. (2020) as well as the distributed version Zhou et al. (2016a). Hence, we are highly motivated to study the convergence rate of **variable convergence** of GDA in nonconvex minimax optimization under the KŁ geometry. In particular, we want to address the following question:

- *Q2: How does the local function geometry captured by the KŁ parameter affects the variable convergence rate of GDA?*

In this paper, we provide comprehensive answers to these questions. We develop a new analysis framework to study the variable convergence of GDA in nonconvex-strongly-concave minimax optimization under the KŁ geometry. We also characterize the convergence rates of GDA in the full spectrum of the parameterization of the KŁ geometry.

1.1 Our contributions

We consider the following regularized nonconvex-strongly-concave minimax optimization problem

$$\min_{x\in\mathbb{R}^m} \max_{y\in\mathcal{Y}} \; f(x,y) + g(x) - h(y), \tag{P}$$

where f is a differentiable and nonconvex-strongly concave function, g is a general nonconvex regularizer and h is a convex regularizer. Both g and h can be possibly nonsmooth. To solve the above regularized minimax problem, we study a proximal-GDA algorithm that leverages the forward-backward splitting update Lions and Mercier (1979); Attouch et al. (2013) .

We study the variable convergence property of proximal-GDA in solving the minimax problem (P). Specifically, we show that proximal-GDA admits a novel Lyapunov function $H(x, y)$ (see Proposition 2), which is monotonically decreasing along the trajectory of proximal GDA, i.e., $H(x_{t+1}, y_{t+1}) < H(x_t, y_t)$. Based on the monotonicity of this Lyapunov function, we show that every limit point of the variable sequences generated by proximal-GDA is a critical point of the objective function.

Moreover, by exploiting the ubiquitous KŁ geometry of the Lyapunov function, we prove that the entire variable sequence of proximal-GDA has a unique limit point, or equivalently speaking, it converges to a certain critical point x^*, i.e., $x_t \to x^*, y_t \to y^*(x^*)$ (see the definition of y^* in Section

2). To the best of our knowledge, this is the first variable convergence result of GDA-type algorithms in nonconvex minimax optimization.

Furthermore, we characterize the asymptotic convergence rates of both the variable sequences and function value of proximal-GDA in different parameterization regimes of the KŁ geometry. Depending on the value of the KŁ parameter θ, we show that proximal-GDA achieves different types of convergence rates ranging from sublinear convergence up to finite-step convergence, as we summarize in Table 1 below.

Table 1: Convergence rates of proximal-GDA under different parameterizations of KŁ geometry. Note that t_0 denotes a sufficiently large positive integer.

KŁ parameter	Function value convergence rate	Variable convergence rate
$\theta = 1$	Finite-step convergence	Finite-step convergence
$\theta \in (\frac{1}{2}, 1)$	$\mathcal{O}\big(\exp\big(-[2(1-\theta)]^{t_0-t}\big)\big)$ Super-linear convergence	$\mathcal{O}\big(\exp\big(-[2(1-\theta)]^{t_0-t}\big)\big)$ Super-linear convergence
$\theta = \frac{1}{2}$	$\mathcal{O}\big((1+\rho)^{t_0-t}\big), \rho > 0$ Linear convergence	$\mathcal{O}\big(\big(\min\{2, 1+\rho\}\big)^{(t_0-t)/2}\big), \rho > 0$ Linear convergence
$\theta \in (0, \frac{1}{2})$	$\mathcal{O}\big((t-t_0)^{-\frac{1}{1-2\theta}}\big)$ Sub-linear convergence	$\mathcal{O}\big((t-t_0)^{-\frac{\theta}{1-2\theta}}\big)$ Sub-linear convergence

1.2 Related work

Deterministic GDA algorithms: Yang et al. (2020) studied an alternating gradient descent-ascent (AGDA) algorithm in which the gradient ascent step uses the current variable x_{t+1} instead of x_t. Xu et al. (2020b) studied an alternating gradient projection algorithm which applies ℓ_2 regularizer to the local objective function of GDA followed by projection onto the constraint sets. Daskalakis and Panageas (2018); Mokhtari et al. (2020); Zhang and Wang (2020) analyzed optimistic gradient descent-ascent (OGDA) which applies negative momentum to accelerate GDA. Mokhtari et al. (2020) also studied an extra-gradient algorithm which applies two-step GDA in each iteration. Nouiehed et al. (2019) studied multi-step GDA where multiple gradient ascent steps are performed, and they also studied the momentum-accelerated version. Cherukuri et al. (2017); Daskalakis and Panageas (2018); Jin et al. (2020) studied GDA in continuous time dynamics using differential equations. Adolphs et al. (2019) analyzed a second-order variant of the GDA algorithm.

Stochastic GDA algorithms: Lin et al. (2020); Yang et al. (2020) analyzed stochastic GDA and stochastic AGDA, which are direct extension of GDA and AGDA to the stochastic setting. Variance reduction techniques have been applied to stochastic minimax optimization, including SVRG-based Du and Hu (2019); Yang et al. (2020), SPIDER-based Xu et al. (2020a), STORM Qiu et al. (2020) and its gradient free version Huang et al. (2020). Xie et al. (2020) studied the complexity lower bound of first-order stochastic algorithms for finite-sum minimax problem.

KŁ geometry: The KŁ geometry was defined in Bolte et al. (2007). The KŁ geometry has been exploited to study the convergence of various first-order algorithms for solving minimization problems, including gradient descent Attouch and Bolte (2009), alternating gradient descent Bolte et al. (2014), distributed gradient descent Zhou et al. (2016a; 2018a), accelerated gradient descent Li et al. (2017). It has also been exploited to study the convergence of second-order algorithms such as Newton's method Noll and Rondepierre (2013); Frankel et al. (2015) and cubic regularization method Zhou et al. (2018b).

2 Problem Formulation and KŁ Geometry

In this section, we introduce the problem formulation, technical assumptions and the Kurdyka-Łojasiewicz (KŁ) geometry. We consider the following regularized minimax optimization problem.

$$\min_{x \in \mathbb{R}^m} \max_{y \in \mathcal{Y}} \; f(x, y) + g(x) - h(y), \tag{P}$$

where $f : \mathbb{R}^m \times \mathbb{R}^n \to \mathbb{R}$ is a differentiable and nonconvex-strongly concave loss function, $\mathcal{Y} \subset \mathbb{R}^n$ is a compact and convex set, and g, h are the regularizers that are possibly non-smooth. In particular, define $\Phi(x) := \max_{y\in\mathcal{Y}} f(x, y) - h(y)$, and then the problem (P) is equivalent to the minimization problem $\min_{x\in\mathbb{R}^m} \Phi(x) + g(x)$.

Throughout the paper, we adopt the following standard assumptions on the problem (P).

Assumption 1. *The objective function of the problem* (P) *satisfies:*

1. *Function $f(\cdot, \cdot)$ is L-smooth and function $f(x, \cdot)$ is μ-strongly concave;*
2. *Function $(\Phi + g)(x)$ is bounded below, i.e., $\inf_{\mathbf{x}\in\mathbb{R}^m}(\Phi + g)(x) > -\infty$;*
3. *For any $\alpha \in \mathbb{R}$, the sub-level set $\{x : (\Phi + g)(x) \leq \alpha\}$ is compact;*
4. *Function h is proper and convex, and function g is proper and lower semi-continuous.*

To elaborate, item 1 considers the class of nonconvex-strongly-concave functions f that has been widely studied in the existing literature Lin et al. (2020); Jin et al. (2020); Xu et al. (2020b;a); Lu et al. (2020). Items 2 and 3 guarantee that the minimax problem (P) has at least one solution, and the variable sequences generated by the proximal-GDA algorithm (See Algorithm 1) are bounded. Item 4 requires the regularizer h to be convex (possibly non-smooth), which includes many norm-based popular regularizers such as ℓ_p $(p \geq 1)$, elastic net, nuclear norm, spectral norm, etc. On the other hand, the other regularizer g can be nonconvex but lower semi-continuous, which includes all the aforementioned convex regularizers, ℓ_p $(0 \leq p < 1)$, Schatten-p norm, rank, etc. Hence, our formulation of the problem (P) covers a rich class of nonconvex objective functions and regularizers and is more general than the existing nonconvex minimax formulation in Lin et al. (2020), which does not consider any regularizer.

Remark 1. *We note that the strong concavity of $f(x, \cdot)$ in item 1 can be relaxed to concavity, provided that the regularizer $h(y)$ is μ -strongly convex. In this case, we can add $-\frac{\mu}{2}\|y\|^2$ to both $f(x, y)$ and $h(y)$ such that Assumption 1 still holds. For simplicity, we will omit the discussion on this case.*

By strong concavity of $f(x, \cdot)$, it is clear that the mapping $y^*(x) := \arg\max_{y\in\mathcal{Y}} f(x, y) - h(y)$ is uniquely defined for every $x \in \mathbb{R}^m$. In particular, if x^* is the desired minimizer of $\Phi(x)$, then $(x^*, y^*(x^*))$ is the desired solution of the minimax problem (P).

Next, we present some important properties regarding the function $\Phi(x)$ and the mapping $y^*(x)$. The following proposition generalizes the Lemma 4.3 of Lin et al. (2020) to the regularized setting. The proof can be found in Appendix A. Throughout, we denote $\kappa = L/\mu$ as the condition number and denote $\nabla_1 f(x, y), \nabla_2 f(x, y)$ as the gradients with respect to the first and the second input argument, respectively. For example, with this notation, $\nabla_1 f(x, y^*(x))$ denotes the gradient of $f(x, y^*(x))$ with respect to only the first input argument x, and the x in the second input argument $y^*(x)$ is treated as a constant.

Proposition 1 (Lipschitz continuity of $y^*(x)$ and $\nabla\Phi(x)$)**.** *Let Assumption 1 hold. Then, the mapping $y^*(x)$ and the function $\Phi(x)$ satisfy*

1. *Mapping $y^*(x)$ is κ-Lipschitz continuous;*
2. *Function $\Phi(x)$ is $L(1 + \kappa)$-smooth with $\nabla\Phi(x) = \nabla_1 f(x, y^*(x))$.*

As an intuitive explanation of Proposition 1, since the function $f(x, y) - h(y)$ is L-smooth with respect to x, both the maximizer $y^*(x)$ and the corresponding maximum function value $\Phi(x)$ should not change substantially with regard to a small change of x.

Recall that the minimax problem (P) is equivalent to the standard minimization problem $\min_{x\in\mathbb{R}^m} \Phi(x) + g(x)$, which, according to item 2 of Proposition 1, includes a smooth nonconvex function $\Phi(x)$ and a lower semi-continuous regularizer $g(x)$. Hence, we can define the optimization goal of the minimax problem (P) as **finding a critical point x^* of the nonconvex function** $\Phi(x) + g(x)$ that satisfies the necessary optimality condition $\mathbf{0} \in \partial(\Phi + g)(x^*)$ for minimizing nonconvex functions. Here, ∂ denotes the notion of subdifferential as we elaborate below.

Definition 1. *(Subdifferential and critical point, Rockafellar and Wets (2009)) The Frechét subdifferential $\widehat{\partial} h$ of function h at $x \in \operatorname{dom} h$ is the set of $u \in \mathbb{R}^d$ defined as*

$$\widehat{\partial} h(x) := \Big\{u : \liminf_{z\neq x, z\to x} \frac{h(z) - h(x) - u^\intercal(z - x)}{\|z - x\|} \geq 0\Big\},$$

and the limiting subdifferential ∂h at $x \in \operatorname{dom} h$ is the graphical closure of $\widehat{\partial} h$ defined as:

$$\partial h(x) := \{u : \exists x_k \to x, h(x_k) \to h(x), u_k \in \widehat{\partial} h(x_k), u_k \to u\}.$$

The set of ***critical points*** *of h is defined as* $\operatorname{crit} h := \{x : \mathbf{0} \in \partial h(x)\}$.

Throughout, we refer to the limiting subdifferential as subdifferential. We note that subdifferential is a generalization of gradient (when h is differentiable) and subgradient (when h is convex) to the nonconvex setting. In particular, any local minimizer of h must be a critical point.

Next, we introduce the Kurdyka-Łojasiewicz (KŁ) geometry of a function h. Throughout, the point-to-set distance is denoted as $\operatorname{dist}_\Omega(x) := \inf_{u \in \Omega} \|x - u\|$.

Definition 2 (KŁ geometry, Bolte et al. (2014))**.** *A proper and lower semi-continuous function h is said to have the KŁ geometry if for every compact set $\Omega \subset \operatorname{dom} h$ on which h takes a constant value $h_\Omega \in \mathbb{R}$, there exist $\varepsilon, \lambda > 0$ such that for all $\bar{x} \in \Omega$ and all $x \in \{z \in \mathbb{R}^m : \operatorname{dist}_\Omega(z) < \varepsilon, h_\Omega < h(z) < h_\Omega + \lambda\}$, the following condition holds:*

$$\varphi'\left(h(x) - h_\Omega\right) \cdot \operatorname{dist}_{\partial h(x)}(\mathbf{0}) \geq 1, \tag{1}$$

where φ' is the derivative of function $\varphi : [0, \lambda) \to \mathbb{R}_+$, which takes the form $\varphi(t) = \frac{c}{\theta} t^\theta$ for certain universal constant $c > 0$ and KŁ parameter $\theta \in (0, 1]$.

The KŁ geometry characterizes the local geometry of a nonconvex function around the set of critical points. To explain, consider the case where h is a differentiable function so that $\partial h(x) = \nabla h(x)$. Then, the KŁ inequality in eq. (1) becomes $h(x) - h_\Omega \leq \mathcal{O}(\|\nabla h(x)\|^{\frac{1}{1-\theta}})$, which generalizes the Polyak-Łojasiewicz (PL) condition $h(x) - h_\Omega \leq \mathcal{O}(\|\nabla h(x)\|^2)$ Łojasiewicz (1963); Karimi et al. (2016) (i.e., KŁ parameter $\theta = \frac{1}{2}$). Moreover, the KŁ geometry has been shown to hold for a large class of functions including sub-analytic functions, logarithm and exponential functions and semi-algebraic functions. These function classes cover most of the nonconvex objective functions encountered in practical machine learning applications Zhou et al. (2016b); Yue et al. (2018); Zhou and Liang (2017); Zhou et al. (2018b).

The KŁ geometry has been exploited extensively to analyze the convergence of various first-order algorithms, e.g., gradient descent Attouch and Bolte (2009); Li et al. (2017), alternating minimization Bolte et al. (2014) and distributed gradient methods Zhou et al. (2016a). It has also been exploited to study the convergence of second-order algorithms such cubic regularization Zhou et al. (2018b). In these works, it has been shown that the variable sequences generated by these algorithms converge to a desired critical point in nonconvex optimization, and the convergence rates critically depend on the parameterization θ of the KŁ geometry. In the subsequent sections, we provide a comprehensive understanding of the convergence and convergence rate of proximal-GDA under the KŁ geometry.

3 Proximal-GDA and Global Convergence Analysis

In this section, we study the following proximal-GDA algorithm that leverages the forward-backward splitting updates Lions and Mercier (1979); Attouch et al. (2013) to solve the regularized minimax problem (P) and analyze its global convergence properties. In particular, the proximal-GDA algorithm is a generalization of the GDA Du and Hu (2019) and projected GDA Nedić and Ozdaglar (2009) algorithms. The algorithm update rule is specified in Algorithm 1, where the two proximal gradient steps are formally defined as

$$\operatorname{prox}_{\eta_x g}\big(x_t - \eta_x \nabla_1 f(x_t, y_t)\big) :\in \operatorname*{argmin}_{u \in \mathbb{R}^m} \Big\{ g(u) + \frac{1}{2\eta_x} \|u - x_t + \eta_x \nabla_1 f(x_t, y_t)\|^2 \Big\}, \tag{2}$$

$$\operatorname{prox}_{\eta_y h}\big(y_t + \eta_y \nabla_2 f(x_t, y_t)\big) := \operatorname*{argmin}_{v \in \mathcal{Y}} \Big\{ h(v) + \frac{1}{2\eta_y} \|v - y_t - \eta_y \nabla_2 f(x_t, y_t)\|^2 \Big\}, \tag{3}$$

Recall that our goal is to obtain a critical point of the minimization problem $\min_{x \in \mathbb{R}^m} \Phi(x) + g(x)$. Unlike the gradient descent algorithm which generates a sequence of monotonically decreasing function values, the function value $(\Phi + g)(x_k)$ along the variable sequence generated by proximal-GDA is generally oscillating due to the alternation between the gradient descent and gradient ascent steps. Hence, it seems that proximal-GDA is less stable than gradient descent. However, our next result shows that, for the problem (P), the proximal-GDA admits a special Lyapunov function that monotonically decreases in the optimization process. The proof of Proposition 2 is in Appendix B.

Algorithm 1 Proximal-GDA

Input: Initialization x_0, y_0, learning rates η_x, η_y.
for $t = 0, 1, 2, \ldots, T-1$ **do**

$$x_{t+1} \in \mathrm{prox}_{\eta_x g}\big(x_t - \eta_x \nabla_1 f(x_t, y_t)\big),$$
$$y_{t+1} = \mathrm{prox}_{\eta_y h}\big(y_t + \eta_y \nabla_2 f(x_t, y_t)\big).$$

end
Output: x_T, y_T.

Proposition 2. *Let Assumption 1 hold and define the Lyapunov function* $H(z) := \Phi(x) + g(x) + (1 - \frac{1}{4\kappa^2})\|y - y^*(x)\|^2$ *with* $z := (x, y)$. *Choose the learning rates such that* $\eta_x \leq \frac{1}{\kappa^3(L+3)^2}$, $\eta_y \leq \frac{1}{L}$. *Then, the variables* $z_t = (x_t, y_t)$ *generated by proximal-GDA satisfy, for all* $t = 0, 1, 2, \ldots$

$$H(z_{t+1}) \leq H(z_t) - 2\|x_{t+1} - x_t\|^2 - \frac{1}{4\kappa^2}\big(\|y_{t+1} - y^*(x_{t+1})\|^2 + \|y_t - y^*(x_t)\|^2\big). \quad (4)$$

We first explain how this Lyapunov function is introduced in the proof. By eq. (19) in the supplementary material, we established a recursive inequality on the objective function $(\Phi + g)(x_{t+1})$. One can see that the right hand side of eq. (19) contains a negative term $-\|x_{t+1} - x_t\|^2$ and an undesired positive term $\|y^*(x_t) - y_t\|^2$. Hence, the objective function $(\Phi + g)(x_{t+1})$ may be oscillating and cannot serve as a proper Lyapunov function. In the subsequent analysis, we break this positive term into a difference of two terms $\|y^*(x_t) - y_t\|^2 - \|y^*(x_{t+1}) - y_{t+1}\|^2$, by leveraging the update of y_{t+1} for solving the strongly-concave maximization problem. After proper rearranging, this difference term contributes to the quadratic term in the Lyapunov function.

We note that the Lyapunov function $H(z)$ is the objective function $\Phi(x) + g(x)$ regularized by the additional quadratic term $(1 - \frac{1}{4\kappa^2})\|y - y^*(x)\|^2$, and such a Lyapunov function clearly characterizes our optimization goal. To elaborate, consider a desired case where the sequence x_t converges to a certain critical point x^* and the sequence y_t converges to the corresponding point $y^*(x^*)$. In this case, it can be seen that the Lyapunov function $H(z_t)$ converges to the desired function value $(\Phi + g)(x^*)$. Hence, solving the minimax problem (P) is equivalent to minimizing the Lyapunov function. More importantly, Proposition 2 shows that the Lyapunov function value sequence $\{H(z_t)\}_t$ is monotonically decreasing in the optimization process of proximal-GDA, implying that the algorithm continuously makes optimization progress. We also note that the coefficient $(1 - \frac{1}{4\kappa^2})$ in the Lyapunov function is chosen in a way so that eq. (4) can be proven to be strictly decreasing. This monotonic property is the core of our analysis of proximal-GDA.

Based on Proposition 2, we obtain the following asymptotic properties of the variable sequences generated by proximal-GDA. The proof can be found in Appendix C.

Corollary 1. *Based on Proposition 2, the sequences* $\{x_t, y_t\}_t$ *generated by proximal-GDA satisfy*

$$\lim_{t\to\infty} \|x_{t+1} - x_t\| = 0, \quad \lim_{t\to\infty} \|y_{t+1} - y_t\| = 0, \quad \lim_{t\to\infty} \|y_t - y^*(x_t)\| = 0.$$

The above result shows that the variable sequences generated by proximal-GDA in solving the problem (P) are asymptotically stable. In particular, the last two equations show that y_t asymptotically approaches the corresponding maximizer $y^*(x_t)$ of the objective function $f(x_t, y) + g(x_t) - h(y)$. Hence, if x_t converges to a certain critical point, y_t will converge to the corresponding maximizer.

Discussion: We note that the monotonicity property in Proposition 2 further implies the convergence rate result $\min_{0\leq k\leq t} \|x_{k+1} - x_k\| \leq \mathcal{O}(t^{-1/2})$ (by telescoping over t). When there is no regularizer, this convergence rate result can be shown to further imply that $\min_{0\leq k\leq t} \|\nabla\Phi(x_k)\| \leq \mathcal{O}(t^{-1/2})$, which reduces to the Theorem 4.4 of Lin et al. (2020). However, such a convergence rate result does not imply the convergence of the variable sequences $\{x_t\}_t, \{y_t\}_t$. To explain, we can apply the convergence rate result $\|x_{t+1} - x_t\| \leq \mathcal{O}(t^{-1/2})$ to bound the trajectory norm as $\|x_T\| \leq \|x_0\| + \sum_{t=0}^{T-1} \|x_{t+1} - x_t\| \approx \sqrt{T}$, which diverges to $+\infty$ as $T \to \infty$. Therefore, such a type of convergence rate does not even imply the boundedness of the trajectory. In this paper, our focus is to establish the convergence of the variable sequences generated by proximal-GDA.

All the results in Corollary 1 imply that the alternating proximal gradient descent & ascent updates of proximal-GDA can achieve stationary points, which we show below to be critical points.

Theorem 1 (Global convergence)**.** *Let Assumption 1 hold and choose the learning rates* $\eta_x \leq \frac{1}{\kappa^3(L+3)^2}$, $\eta_y \leq \frac{1}{L}$. *Then, proximal-GDA satisfies the following properties.*

1. *The function value sequence* $\{(\Phi+g)(x_t)\}_t$ *converges to a finite limit* $H^* > -\infty$*;*
2. *The sequences* $\{x_t\}_t, \{y_t\}_t$ *are bounded and have compact sets of limit points. Moreover,* $(\Phi+g)(x^*) \equiv H^*$ *for any limit point* x^* *of* $\{x_t\}_t$*;*
3. *Every limit point of* $\{x_t\}_t$ *is a critical point of* $(\Phi+g)(x)$.

The proof of Theorem 1 is presented in Appendix D. The above theorem establishes the global convergence property of proximal-GDA. Specifically, item 1 shows that the function value sequence $\{(\Phi+g)(x_t)\}_t$ converges to a finite limit H^*, which is also the limit of the Lyapunov function sequence $\{H(z_t)\}_t$. Moreover, items 2 & 3 further show that all the converging subsequences of $\{x_t\}_t$ converge to critical points of the problem, at which the function $\Phi+g$ achieves the constant value H^*. These results show that proximal-GDA can properly find critical points of the minimax problem (P). Furthermore, based on these results, the variable sequences generated by proximal-GDA are guaranteed to enter a local parameter region where the Kurdyka-Łojasiewicz geometry holds, which we exploit in the next section to establish stronger convergence results of the algorithm.

4 Variable Convergence of Proximal-GDA under KŁ Geometry

We note that Theorem 1 only shows that every limit point of $\{x_t\}_t$ is a critical point, and the sequences $\{x_t, y_t\}_t$ may not necessarily be convergent. In this section, we exploit the local KŁ geometry of the Lyapunov function to formally prove the convergence of these sequences. Throughout this section, we adopt the following assumption.

Assumption 2. *Regarding the mapping* $y^*(x)$*, the function* $\|y^*(x)-y\|^2$ *has a non-empty subdifferential, i.e.,* $\partial_x(\|y^*(x)-y\|^2) \neq \emptyset$.

Note that in many practical scenarios $y^*(x)$ is sub-differentiable. In addition, Assumption 2 ensures the sub-differentiability of the Lyapunov function $H(z) := \Phi(x) + g(x) + (1-\frac{1}{4\kappa^2})\|y-y^*(x)\|^2$. We obtain the following variable convergence result of proximal-GDA under the KŁ geometry. The proof is presented in Appendix E.

Theorem 2 (Variable convergence)**.** *Let Assumption 1 & 2 hold and assume that* H *has the KŁ geometry. Choose the learning rates* $\eta_x \leq \frac{1}{\kappa^3(L+3)^2}$ *and* $\eta_y \leq \frac{1}{L}$. *Then, the sequence* $\{(x_t, y_t)\}_t$ *generated by proximal-GDA converges to a certain critical point* $(x^*, y^*(x^*))$ *of* $(\Phi+g)(x)$*, i.e.,*

$$x_t \xrightarrow{t} x^*, \quad y_t \xrightarrow{t} y^*(x^*).$$

Theorem 2 formally shows that proximal-GDA is guaranteed to converge to a certain critical point $(x^*, y^*(x^*))$ of the minimax problem (P), provided that the Lyapunov function belongs to the large class of KŁ functions. To the best of our knowledge, this is the first variable convergence result of GDA-type algorithms in nonconvex minimax optimization. The proof logic of Theorem 2 can be summarized as the following two key steps.

Step 1: By leveraging the monotonicity property of the Lyapunov function in Proposition 2, we first show that the variable sequences of proximal-GDA eventually enter a local region where the KŁ geometry holds;

Step 2: Then, combining the KŁ inequality in eq. (1) and the monotonicity property of the Lyapunov function in eq. (4), we show that the variable sequences of proximal-GDA are Cauchy sequences and hence converge to a certain critical point.

5 Convergence Rate of Proximal-GDA under KŁ Geometry

In this section, we exploit the parameterization of the KŁ geometry to establish various types of asymptotic convergence rates of proximal-GDA.

We obtain the following asymptotic convergence rates of proximal-GDA under different parameter regimes of the KŁ geometry. The proof is presented in Appendix F. In the sequel, we denote t_0 as a sufficiently large positive integer, denote $c > 0$ as the constant in Definition 2 and also define

$$M := \max\Big\{\frac{1}{2}\Big(\frac{1}{\eta_x} + (L + 4\kappa^2)(1+\kappa)\Big)^2, 4\kappa^2(L+4\kappa)^2\Big\}. \tag{5}$$

Theorem 3 (Funtion value convergence rate)**.** *Under the same conditions as those of Theorem 2, the Lyapunov function value sequence $\{H(z_t)\}_t$ converges to the limit H^* at the following rates.*

1. *If KŁ geometry holds with $\theta = 1$, then $H(z_t) \downarrow H^*$ within* **finite number of iterations***;*
2. *If KŁ geometry holds with $\theta \in (\frac{1}{2}, 1)$, then $H(z_t) \downarrow H^*$* **super-linearly** *as*

$$H(z_t) - H^* \leq (2Mc^2)^{-\frac{1}{2\theta-1}} \exp\Big(-\Big(\frac{1}{2(1-\theta)}\Big)^{t-t_0}\Big), \quad \forall t \geq t_0; \tag{6}$$

3. *If KŁ geometry holds with $\theta = \frac{1}{2}$, then $H(z_t) \downarrow H^*$* **linearly** *as*

$$H(z_t) - H^* \leq \Big(1 + \frac{1}{2Mc^2}\Big)^{t_0 - t}(H(z_{t_0}) - H^*), \quad \forall t \geq t_0; \tag{7}$$

4. *If KŁ geometry holds with $\theta \in (0, \frac{1}{2})$, then $H(z_t) \downarrow H^*$* **sub-linearly** *as*

$$H(z_t) - H^* \leq \big[C(t - t_0)\big]^{-\frac{1}{1-2\theta}}, \quad \forall t \geq t_0. \tag{8}$$

where $C = \min\Big[\frac{1-2\theta}{8Mc^2}, d_{t_0}^{-(1-2\theta)}\big(1 - 2^{-(1-2\theta)}\big)\Big] > 0$.

It can be seen from the above theorem that the convergence rate of the Lyapunov function of proximal-GDA is determined by the KŁ parameter θ. A larger θ implies that the local geometry of H is 'sharper', and hence the corresponding convergence rate is orderwise faster. In particular, the algorithm converges at a linear rate when the KŁ geometry holds with $\theta = \frac{1}{2}$ (see the item 3), which is a generalization of the Polyak-Łojasiewicz (PL) geometry. As a comparison, in the existing analysis of GDA, such a linear convergence result is established under stronger geometries, e.g., convex-strongly concave Du and Hu (2019), strongly convex-strongly concave Mokhtari et al. (2020); Zhang and Wang (2020) and two-sided PL condition Yang et al. (2020). In summary, the above theorem provides a full characterization of the fast convergence rates of proximal-GDA in the full spectrum of the KŁ geometry.

Moreover, we also obtain the following asymptotic convergence rates of the variable sequences that are generated by proximal-GDA under different parameterization of the KŁ geometry. The proof is presented in Appendix G.

Theorem 4 (Variable convergence rate)**.** *Under the same conditions as those of Theorem 2, the sequences $\{x_t, y_t\}_t$ converge to their limits $x^*, y^*(x^*)$ respectively at the following rates.*

1. *If KŁ geometry holds with $\theta = 1$, then $(x_t, y_t) \to (x^*, y^*(x^*))$ within* **finite number of iterations***;*
2. *If KŁ geometry holds with $\theta \in (\frac{1}{2}, 1)$, then $(x_t, y_t) \to (x^*, y^*(x^*))$* **super-linearly** *as*

$$\max\big\{\|x_t - x^*\|, \|y_t - y^*(x^*)\|\big\} \leq \mathcal{O}\Big(\exp\Big(-\big(\frac{1}{2(1-\theta)}\big)^{t-t_0}\Big)\Big), \quad \forall t \geq t_0; \tag{9}$$

3. *If KŁ geometry holds with $\theta = \frac{1}{2}$, then $(x_t, y_t) \to (x^*, y^*(x^*))$* **linearly** *as*

$$\max\big\{\|x_t - x^*\|, \|y_t - y^*(x^*)\|\big\} \leq \mathcal{O}\Big(\Big(\min\Big\{2, 1 + \frac{1}{2Mc^2}\Big\}\Big)^{(t_0-t)/2}\Big), \quad \forall t \geq t_0; \tag{10}$$

4. *If KŁ geometry holds with $\theta \in (0, \frac{1}{2})$, then $(x_t, y_t) \to (x^*, y^*(x^*))$* **sub-linearly** *as*

$$\max\big\{\|x_t - x^*\|, \|y_t - y^*(x^*)\|\big\} \leq \mathcal{O}\Big((t - t_0)^{-\frac{\theta}{1-2\theta}}\Big), \quad \forall t \geq t_0. \tag{11}$$

To the best of our knowledge, this is the first characterization of the variable convergence rates of proximal-GDA in the full spectrum of the KŁ geometry. It can be seen that, similar to the convergence rate results of the function value sequence, the convergence rate of the variable sequences is also affected by the parameterization of the KŁ geometry.

6 CONCLUSION

In this paper, we develop a new analysis framework for the proximal-GDA algorithm in nonconvex-strongly-concave optimization. Our key observation is that proximal-GDA has a intrinsic Lyapunov function that monotonically decreases in the minimax optimization process. Such a property demonstrates the stability of the algorithm. Moreover, we establish the formal variable convergence of proximal-GDA to a critical point of the objective function under the ubiquitous KŁ geometry. Our results fully characterize the impact of the parameterization of the KŁ geometry on the convergence rate of the algorithm. In the future study, we will leverage such an analysis framework to explore the convergence of stochastic GDA algorithms and their variance-reduced variants.

REFERENCES

Adolphs, L., Daneshmand, H., Lucchi, A., and Hofmann, T. (2019). Local saddle point optimization: A curvature exploitation approach. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 486–495.

Attouch, H. and Bolte, J. (2009). On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Mathematical Programming*, 116(1-2):5–16.

Attouch, H., Bolte, J., and Svaiter, B. F. (2013). Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized gauss–seidel methods. *Mathematical Programming*, 137(1-2):91–129.

Bernhard, P. and Rapaport, A. (1995). On a theorem of danskin with an application to a theorem of von neumann-sion. *Nonlinear Analysis: Theory, Methods & Applications*, 24(8):1163–1181.

Bolte, J., Daniilidis, A., and Lewis, A. (2007). The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM Journal on Optimization*, 17:1205–1223.

Bolte, J., Sabach, S., and Teboulle, M. (2014). Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494.

Boţ, R. I. and Böhm, A. (2020). Alternating proximal-gradient steps for (stochastic) nonconvex-concave minimax problems. *ArXiv:2007.13605*.

Cherukuri, A., Gharesifard, B., and Cortes, J. (2017). Saddle-point dynamics: conditions for asymptotic stability of saddle points. *SIAM Journal on Control and Optimization*, 55(1):486–511.

Daskalakis, C. and Panageas, I. (2018). The limit points of (optimistic) gradient descent in min-max optimization. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 9236–9246.

Du, S. S. and Hu, W. (2019). Linear convergence of the primal-dual gradient method for convex-concave saddle point problems without strong convexity. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 196–205.

Ferreira, M. A. M., Andrade, M., Matos, M. C. P., Filipe, J. A., and Coelho, M. P. (2012). Minimax theorem and nash equilibrium.

Frankel, P., Garrigos, G., and Peypouquet, J. (2015). Splitting methods with variable metric for Kurdyka–Łojasiewicz functions and general convergence rates. *Journal of Optimization Theory and Applications*, 165(3):874–900.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 2672–2680.

Ho, J. and Ermon, S. (2016). Generative adversarial imitation learning. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 4565–4573.

Huang, F., Gao, S., Pei, J., and Huang, H. (2020). Accelerated zeroth-order momentum methods from mini to minimax optimization. *ArXiv:2008.08170*.

Jin, C., Netrapalli, P., and Jordan, M. I. (2020). What is local optimality in nonconvex-nonconcave minimax optimization?

Karimi, H., Nutini, J., and Schmidt, M. (2016). *Linear Convergence of Gradient and Proximal-Gradient Methods Under the Polyak-Łojasiewicz Condition*, pages 795–811.

Kruger, A. Y. (2003). On fréchet subdifferentials. *Journal of Mathematical Sciences*, 116(3):3325–3358.

Li, Q., Zhou, Y., Liang, Y., and Varshney, P. K. (2017). Convergence analysis of proximal gradient with momentum for nonconvex optimization. In *Proc. International Conference on Machine Learning (ICML*, volume 70, pages 2111–2119.

Lin, T., Jin, C., and Jordan, M. I. (2020). On gradient descent ascent for nonconvex-concave minimax problems.

Lions, P.-L. and Mercier, B. (1979). Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979.

Łojasiewicz, S. (1963). A topological property of real analytic subsets. *Coll. du CNRS, Les equations aux derivees partielles*, page 87–89.

Lu, S., Tsaknakis, I., Hong, M., and Chen, Y. (2020). Hybrid block successive approximation for one-sided non-convex min-max problems: algorithms and applications. *IEEE Transactions on Signal Processing*.

Mokhtari, A., Ozdaglar, A., and Pattathil, S. (2020). A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1497–1507.

Nedić, A. and Ozdaglar, A. (2009). Subgradient methods for saddle-point problems. *Journal of optimization theory and applications*, 142(1):205–228.

Neumann, J. v. (1928). Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320.

Noll, D. and Rondepierre, A. (2013). Convergence of linesearch and trust-region methods using the Kurdyka–Łojasiewicz inequality. In *Proc. Computational and Analytical Mathematics*, pages 593–611.

Nouiehed, M., Sanjabi, M., Huang, T., Lee, J. D., and Razaviyayn, M. (2019). Solving a class of non-convex min-max games using iterative first order methods. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 14934–14942.

Qiu, S., Yang, Z., Wei, X., Ye, J., and Wang, Z. (2020). Single-timescale stochastic nonconvex-concave optimization for smooth nonlinear td learning. *ArXiv:2008.10103*.

Robinson, J. (1951). An iterative method of solving a game. *Annals of mathematics*, 54(2):296–301.

Rockafellar, R. T. (1970). *Convex analysis*. Number 28. Princeton university press.

Rockafellar, R. T. and Wets, R. J.-B. (2009). *Variational analysis*, volume 317. Springer Science & Business Media.

Sinha, A., Namkoong, H., and Duchi, J. C. (2017). Certifying some distributional robustness with principled adversarial training. In *Proc. International Conference on Learning Representations (ICLR)*.

Song, J., Ren, H., Sadigh, D., and Ermon, S. (2018). Multi-agent generative adversarial imitation learning. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 7461–7472.

Xie, G., Luo, L., Lian, Y., and Zhang, Z. (2020). Lower complexity bounds for finite-sum convex-concave minimax optimization problems.

Xu, T., Wang, Z., Liang, Y., and Poor, H. V. (2020a). Enhanced first and zeroth order variance reduced algorithms for min-max optimization. *arXiv:2006.09361*.

Xu, Z., Zhang, H., Xu, Y., and Lan, G. (2020b). A unified single-loop alternating gradient projection algorithm for nonconvex-concave and convex-nonconcave minimax problems. *ArXiv:2006.02032*.

Yang, J., Kiyavash, N., and He, N. (2020). Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems.

Yue, M., Zhou, Z., and So, M. (2018). On the Quadratic Convergence of the Cubic Regularization Method under a Local Error Bound Condition. *ArXiv:1801.09387v1*.

Zhang, G. and Wang, Y. (2020). On the suboptimality of negative momentum for minimax optimization. *ArXiv:2008.07459*.

Zhou, Y. and Liang, Y. (2017). Characterization of Gradient Dominance and Regularity Conditions for Neural Networks. *ArXiv:1710.06910v2*.

Zhou, Y., Liang, Y., Yu, Y., Dai, W., and Xing, E. P. (2018a). Distributed Proximal Gradient Algorithm for Partially Asynchronous Computer Clusters. *Journal of Machine Learning Research (JMLR)*, 19(19):1–32.

Zhou, Y., Wang, Z., Ji, K., Liang, Y., and Tarokh, V. (2020). Proximal gradient algorithm with momentum and flexible parameter restart for nonconvex optimization. In *Proc. International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 1445–1451.

Zhou, Y., Wang, Z., and Liang, Y. (2018b). Convergence of cubic regularization for nonconvex optimization under kl property. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 3760–3769.

Zhou, Y., Yu, Y., Dai, W., Liang, Y., and Xing, P. (2016a). On convergence of model parallel proximal gradient algorithm for stale synchronous parallel system. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS*, volume 51, pages 713–722.

Zhou, Y., Zhang, H., and Liang, Y. (2016b). Geometrical properties and accelerated gradient solvers of non-convex phase retrieval. *In Proc. 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 331–335.