

Two-Timescale Gradient Descent Ascent Algorithms for Nonconvex Minimax Optimization

Tianyi Lin

*Department of Industrial Engineering and Operations Research
Columbia University
New York, NY 10027, USA*

TL3335@COLUMBIA.EDU

Chi Jin

*Department of Electrical and Computer Engineering
Princeton University
Princeton, NJ 08544, USA*

CHIJ@PRINCETON.EDU

Michael I. Jordan

*Department of Electrical Engineering and Computer Science and Department of Statistics
University of California, Berkeley
Berkeley, CA 94720-1776, USA*

JORDAN@CS.BERKELEY.EDU

Editor: Po-Ling Loh

Abstract

We provide a unified analysis of *two-timescale* gradient descent ascent (TTGDA) for solving structured nonconvex minimax optimization problems in the form of $\min_{\mathbf{x}} \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y})$, where the objective function $f(\mathbf{x}, \mathbf{y})$ is nonconvex in \mathbf{x} and concave in \mathbf{y} , and the constraint set $\mathcal{Y} \subseteq \mathbb{R}^n$ is convex and bounded. In the convex-concave setting, the single-timescale gradient descent ascent (GDA) algorithm is widely used in applications and has been shown to have strong convergence guarantees. In more general settings, however, it can fail to converge. Our contribution is to design TTGDA algorithms that are effective beyond the convex-concave setting, efficiently finding a stationary point of the function $\Phi(\cdot) := \max_{\mathbf{y} \in \mathcal{Y}} f(\cdot, \mathbf{y})$. We also establish theoretical bounds on the complexity of solving both *smooth* and *nonsmooth* nonconvex-concave minimax optimization problems. To the best of our knowledge, this is the first systematic analysis of TTGDA for nonconvex minimax optimization, shedding light on its superior performance in training generative adversarial networks (GANs) and in other real-world application problems.

Keywords: Structured nonconvex minimax optimization, two-timescale gradient descent ascent, iteration complexity analysis

1. Introduction

In the wake of von Neumann’s seminal work (Neumann, 1928), the problem of finding a saddle-point solution to minimax optimization problems, under various assumptions on the objective functions, has been a major focus of research in many areas, including economics, control theory, operations research and computer science (Basar and Olsder, 1999; Nisan et al., 2007; Von Neumann and Morgenstern, 2007). Recently, minimax optimization theory has begun to see applications in machine learning, with examples including generative adversarial networks (GANs) (Goodfellow et al., 2014), robust statistics (Xu et al., 2009;

Algorithm 1 TTGDA

Input: $(\mathbf{x}_0, \mathbf{y}_0), (\eta_{\mathbf{x}}, \eta_{\mathbf{y}})_{t \geq 0}$ with $\eta_{\mathbf{x}}^t \ll \eta_{\mathbf{y}}^t$.
for $t = 1, 2, \dots, T$ **do**
 $(\mathbf{g}_{\mathbf{x}}^{t-1}, \mathbf{g}_{\mathbf{y}}^{t-1}) \in \partial f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$.
 $\mathbf{x}_t \leftarrow \mathbf{x}_{t-1} - \eta_{\mathbf{x}}^{t-1} \mathbf{g}_{\mathbf{x}}^{t-1}$,
 $\mathbf{y}_t \leftarrow \mathcal{P}_{\mathcal{Y}}(\mathbf{y}_{t-1} + \eta_{\mathbf{y}}^{t-1} \mathbf{g}_{\mathbf{y}}^{t-1})$.
end for
 Randomly draw $\hat{\mathbf{x}}$ from $\{\mathbf{x}_t\}_{t=0}^T$ at uniform.
Return: $\hat{\mathbf{x}}$.

Algorithm 2 TTSGDA

Input: $(\mathbf{x}_0, \mathbf{y}_0), M, (\eta_{\mathbf{x}}, \eta_{\mathbf{y}})_{t \geq 0}$ with $\eta_{\mathbf{x}}^t \ll \eta_{\mathbf{y}}^t$.
for $t = 1, 2, \dots, T$ **do**
 $(\hat{\mathbf{g}}_{\mathbf{x}}^{t-1}, \hat{\mathbf{g}}_{\mathbf{y}}^{t-1}) = \text{SG}(G, \{\xi_i^{t-1}\}_{i=1}^M, \mathbf{x}_{t-1}, \mathbf{y}_{t-1})$.
 $\mathbf{x}_t \leftarrow \mathbf{x}_{t-1} - \eta_{\mathbf{x}}^{t-1} \hat{\mathbf{g}}_{\mathbf{x}}^{t-1}$.
 $\mathbf{y}_t \leftarrow \mathcal{P}_{\mathcal{Y}}(\mathbf{y}_{t-1} + \eta_{\mathbf{y}}^{t-1} \hat{\mathbf{g}}_{\mathbf{y}}^{t-1})$.
end for
 Randomly draw $\hat{\mathbf{x}}$ from $\{\mathbf{x}_t\}_{t=0}^T$ at uniform.
Return: $\hat{\mathbf{x}}$.

Abadeh et al., 2015), online adversarial learning (Cesa-Bianchi and Lugosi, 2006), robust training of deep neural networks (Sinha et al., 2018) and distributed computing over networks (Shamma, 2008; Mateos et al., 2010). Current trends in multi-agent learning with large-scale economic systems seem likely to further increase the need for efficient algorithms for computing minimax optima and related equilibrium solutions.

Formally, we study the following minimax optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^m} \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y}), \tag{1.1}$$

where $f : \mathbb{R}^m \times \mathbb{R}^n \mapsto \mathbb{R}$ is the objective function and \mathcal{Y} is a constraint set. One of simplest methods for solving the minimax optimization problems is to generalize gradient descent (GD) so that each loop comprises both a descent step and an ascent step. The algorithm, referred to as *gradient descent ascent* (GDA), performs one-step gradient descent over the min variable \mathbf{x} with a learning rate $\eta_{\mathbf{x}} > 0$ and one-step gradient ascent over the max variable \mathbf{y} with another learning rate $\eta_{\mathbf{y}} > 0$.

On the positive side, when f is convex in \mathbf{x} and concave in \mathbf{y} , there is a vast literature establishing asymptotic and nonasymptotic convergence for average iterates generated by classical single-timescale GDA (i.e., $\eta_{\mathbf{x}} = \eta_{\mathbf{y}}$) (see, e.g., Korpelevich, 1976; Chen and Rockafellar, 1997; Nemirovski, 2004; Auslender and Teboulle, 2009; Nedić and Ozdaglar, 2009). In addition, local linear convergence can be established under an additional assumption that f is locally strongly convex in \mathbf{x} and locally strongly concave in \mathbf{y} (Cherukuri et al., 2017; Adolphs et al., 2019). In contrast, there has been no shortage of negative results that highlight the fact that in the general settings, the single-timescale GDA algorithm might converge to limit cycles or even diverge (Mertikopoulos et al., 2018).

A recent line of research has focused on alternative gradient-based algorithms that have guarantees beyond the convex-concave setting (Daskalakis et al., 2018; Heusel et al., 2017; Mertikopoulos et al., 2019). Notably, the two-timescale GDA (TTGDA), which incorporates unequal learning rates ($\eta_{\mathbf{x}} \neq \eta_{\mathbf{y}}$), has been shown empirically to alleviate problems with limit circles; indeed, it attains theoretical support in terms of local asymptotic convergence (see Heusel et al. (2017, Theorem 2) for the formal statements and proofs). For sequential problems such as robust learning, where the natural order of min-max is important (i.e., min-max is not equal to max-min), practitioners often prefer the faster convergence for the max problem. As such, there can be certain practical motivations for choosing $\eta_{\mathbf{x}} \ll \eta_{\mathbf{y}}$ in addition to those relating to nonasymptotic convergence guarantees.

We provide detailed pseudocode for TTGDA and its stochastic counterpart (TTSGDA) in Algorithms 1 and 2. Note that Algorithm 2 repeatedly samples while remaining at the current iterate and the subroutine $(\hat{\mathbf{g}}_{\mathbf{x}}^t, \hat{\mathbf{g}}_{\mathbf{y}}^t) = \text{SG}(G, \{\xi_i^t\}_{i=1}^M, \mathbf{x}_t, \mathbf{y}_t)$ can be implemented in general by the updates

$$\hat{\mathbf{g}}_{\mathbf{x}}^t = \frac{1}{M} \left(\sum_{i=1}^M G_{\mathbf{x}}(\mathbf{x}_t, \mathbf{y}_t, \xi_i^t) \right), \quad \hat{\mathbf{g}}_{\mathbf{y}}^t = \frac{1}{M} \left(\sum_{i=1}^M G_{\mathbf{y}}(\mathbf{x}_t, \mathbf{y}_t, \xi_i^t) \right), \quad (1.2)$$

where $G = (G_{\mathbf{x}}, G_{\mathbf{y}})$ is unbiased and has bounded variance such that

$$\mathbb{E}[G(\mathbf{x}, \mathbf{y}, \xi) \mid \mathbf{x}, \mathbf{y}] \in \partial f(\mathbf{x}, \mathbf{y}), \quad \mathbb{E}[\|G(\mathbf{x}, \mathbf{y}, \xi) - \mathbb{E}[G(\mathbf{x}, \mathbf{y}, \xi)]\|^2 \mid \mathbf{x}, \mathbf{y}] \leq \sigma^2. \quad (1.3)$$

In the current paper, we aim to provide a theoretical analysis of these algorithms, extending the asymptotic analysis in Heusel et al. (2017, Theorem 2) to a nonasymptotic analysis that characterizes algorithmic efficiency. Notably, we pursue our analysis in a general setting in which $f(\mathbf{x}, \cdot)$ is concave for any \mathbf{x} and \mathcal{Y} is a convex and bounded set. This setting arises in many practical applications, including robust training of a deep neural network and robust learning of a classifier from multiple distributions. Both of these models can be reformulated as nonconvex-concave minimax optimization problems.

Our main results in this paper are presented in four theorems. In particular, the first and second theorems (Theorems 17 and 19) provide a complete characterization of TTGDA and TTSGDA when applied to solve *smooth* minimax optimization problems. In the nonconvex-strongly-concave setting, TTGDA and TTSGDA require $O(\kappa^2 \epsilon^{-2})$ and $O(\kappa^3 \epsilon^{-4})$ (stochastic) gradient evaluations to find an ϵ -stationary point (cf. Definition 3) where κ is a condition number. In the nonconvex-concave setting, TTGDA and TTSGDA require $O(\epsilon^{-6})$ and $O(\epsilon^{-8})$ (stochastic) gradient evaluations to find an ϵ -stationary point (cf. Definition 8). While these results have appeared in our preliminary work (Lin et al., 2020a), the analysis that we provide here is simpler and more direct.

Another two theorems (Theorems 24 and 26) are the more substantive contributions of this paper. These theorems show that TTGDA and TTSGDA can efficiently find an ϵ -stationary point (cf. Definition 8) even when applied to solve *nonsmooth* minimax optimization problems. The complexity bounds of $\tilde{O}(\epsilon^{-6})$ and $O(\epsilon^{-8})$ in terms of (stochastic) gradient evaluations can be established for the nonconvex-strongly-concave and nonconvex-concave settings. These results are useful since the aforementioned machine learning applications often have the nonsmooth objective functions (e.g., the use of ReLU neural network in GANs). Our theorems demonstrate that, when lack of smoothness occurs, TTGDA and TTSGDA may slow down, but the nonasymptotic complexity bounds will still be attained as long as a certain Lipschitz continuity property is retained and adaptive stepsizes are chosen carefully (i.e., our algorithms use $(\eta_{\mathbf{x}}^t, \eta_{\mathbf{y}}^t)_{t \geq 0}$ instead of $(\eta_{\mathbf{x}}, \eta_{\mathbf{y}})$).

Our proof technique in this paper is worth being highlighted. It is different from existing techniques for analyzing nested-loop algorithms as exemplified in Nouiehed et al. (2019) and Jin et al. (2020). Indeed, the previous frameworks are based on inexact gradient descent algorithms for minimizing a function $\Phi(\cdot) = \max_{\mathbf{y} \in \mathcal{Y}} f(\cdot, \mathbf{y})$, where a subroutine is designed for the inner problem of $\max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y})$ and which requires \mathbf{y}_t to be close to $\mathbf{y}^*(\mathbf{x}_t) = \arg\max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}_t, \mathbf{y})$ at each iteration. Applying this approach to analyze TTGDA and TTSGDA is challenging since \mathbf{y}_t might not be close enough to $\mathbf{y}^*(\mathbf{x}_t)$, so that $\nabla_{\mathbf{x}} f(\mathbf{x}_t, \mathbf{y}_t)$

does not approximate a descent direction well. In order to overcome this issue, we have developed a new proof technique that directly analyzes the concave optimization problem with a slowly changing objective function.

Subsequent to our work in Lin et al. (2020a), several new algorithms were developed for solving smooth and nonconvex-(strongly)-concave minimax optimization problems (Lin et al., 2020b; Yang et al., 2020b; Ostrovskii et al., 2021b; Yan et al., 2020; Zhang et al., 2020b, 2021; Li et al., 2021; Yang et al., 2022b; Zhao, 2023; Boş and Böhm, 2023; Xu et al., 2023). In the smooth and nonconvex-strongly-concave setting, a lower bound has been derived for gradient-based algorithms (Zhang et al., 2021; Li et al., 2021) and there exist various nested-loop *near-optimal* deterministic algorithms (Lin et al., 2020b; Ostrovskii et al., 2021b; Kong and Monteiro, 2021; Zhang et al., 2021). In the smooth and nonconvex-concave setting, no lower bound is known, and the best-known upper bound has been achieved by various nested-loop algorithms (Thekumparampil et al., 2019; Lin et al., 2020b; Yang et al., 2020b; Ostrovskii et al., 2021b; Kong and Monteiro, 2021; Zhao, 2023). Compared to deterministic counterparts, investigations of stochastic algorithms are scarce; indeed, there exists a gap between the best-known lower bound (Li et al., 2021) and upper bound (Yan et al., 2020).

A growing body of research has focused on *single-loop* algorithms for nonconvex minimax optimization problems (Zhang et al., 2020b; Yang et al., 2022b; Boş and Böhm, 2023; Xu et al., 2023). For example, Xu et al. (2023) proposed an alternating GDA which converges in both nonconvex-concave and convex-nonconcave settings and showed that the required number of (stochastic) gradient evaluations to find an ϵ -stationary point of f is the best among all single-loop algorithms. Boş and Böhm (2023) studied the alternating TTGDA and TTSGDA in a more general setting where $f(\cdot, \mathbf{y})$ is Lipschitz and weakly convex but $f(\mathbf{x}, \cdot)$ is smooth. They showed that the algorithms can achieve the same complexity bound in this setting as in the smooth setting. These two papers extended our proof techniques to analyze the alternating TTGDA and TTSGDA but still assume the smoothness of $f(\mathbf{x}, \cdot)$. By contrast, the results in the current paper are the first to demonstrate the efficiency of TTGDA and TTSGDA in the nonsmooth setting. Both Zhang et al. (2020b) and Yang et al. (2022b) proposed single-loop algorithms for solving other structured nonconvex-concave minimax optimization problems with theoretical guarantee.

Notation. For $n \geq 2$, we let $[n]$ denote the set $\{1, 2, \dots, n\}$ and \mathbb{R}_+^n denote the set of all vectors in \mathbb{R}^n with nonnegative coordinates. We also use bold lower-case letters to denote vectors (i.e., \mathbf{x} and \mathbf{y}), and calligraphic upper-case letters to denote sets (i.e., \mathcal{X} and \mathcal{Y}). For a Lipschitz function f , we let $\partial f(\cdot)$ denote its Clarke subdifferential (Clarke, 1990). For a differentiable function f , we let $\nabla f(\cdot)$ denote its gradient. For a Lipschitz function $f(\cdot, \cdot)$ of two variables, we let $\partial_{\mathbf{x}} f(\cdot, \cdot)$ and $\partial_{\mathbf{y}} f(\cdot, \cdot)$ denote its partial Clarke subdifferentials with respect to the first and second variables. For a differentiable function $f(\cdot, \cdot)$, we let $\nabla_{\mathbf{x}} f(\cdot, \cdot)$ and $\nabla_{\mathbf{y}} f(\cdot, \cdot)$ denote the partial gradients of f with respect to the first and second variables. For a vector \mathbf{x} and a matrix X , we let $\|\mathbf{x}\|$ and $\|X\|$ denote the ℓ_2 -norm of \mathbf{x} and the spectral norm of X . For a set \mathcal{X} , we let $D = \max_{\mathbf{x}, \mathbf{x}' \in \mathcal{X}} \|\mathbf{x} - \mathbf{x}'\|$ denote its diameter and $\mathcal{P}_{\mathcal{X}}$ denote its orthogonal projection. Lastly, we use $O(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ to hide absolute constants which do not depend on any problem parameter, and $\tilde{O}(\cdot)$, $\tilde{\Omega}(\cdot)$ and $\tilde{\Theta}(\cdot)$ to hide both absolute constants and log factors.

2. Related Work

Historically, an early concrete instantiation of minimax optimization involved computing a pair of probability vectors (\mathbf{x}, \mathbf{y}) for solving the problem $\min_{\mathbf{x} \in \Delta^m} \max_{\mathbf{y} \in \Delta^n} \mathbf{x}^\top \mathbf{A} \mathbf{y}$ for $A \in \mathbb{R}^{m \times n}$ and simplices $\Delta^m \subseteq \mathbb{R}^m$ and $\Delta^n \subseteq \mathbb{R}^n$. This bilinear minimax problem together with von Neumann’s minimax theorem (Neumann, 1928) were a cornerstone in the development of game theory. A simple and generic algorithm was developed for solving this problem where the min and max players implemented a learning procedure in tandem (Robinson, 1951). Subsequently, Sion (1958) generalized von Neumann’s result from bilinear games to general convex-concave games, i.e., $\min_{\mathbf{x}} \max_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}) = \max_{\mathbf{y}} \min_{\mathbf{x}} f(\mathbf{x}, \mathbf{y})$, and triggered a line of research on algorithmic design for convex-concave minimax optimization in continuous time (Kose, 1956; Cherukuri et al., 2017) and discrete time (Uzawa, 1958; Golshtein, 1974; Korpelevich, 1976; Nemirovski, 2004; Auslender and Teboulle, 2009; Nedić and Ozdaglar, 2009; Liang and Stokes, 2019; Mokhtari et al., 2020a,b; Azizian et al., 2020). This line of research has yielded two important conclusions: single-timescale GDA, with decreasing learning rates, can find an ϵ -approximate saddle point within $O(\epsilon^{-2})$ iterations in the convex-concave setting (Auslender and Teboulle, 2009), and $O(\kappa^2 \log(1/\epsilon))$ iterations in the strongly-convex-strongly-concave setting (Liang and Stokes, 2019).

Nonconvex-concave setting. Smooth and nonconvex-concave minimax optimization problems are a class of computationally tractable problems in the form of Eq. (1.1). They have emerged as a focus in the optimization and machine learning communities (Sinha et al., 2018; Rafique et al., 2021; Sanjabi et al., 2018; Grnarova et al., 2018; Lu et al., 2020; Nouiehed et al., 2019; Thekumparampil et al., 2019; Ostrovskii et al., 2021b; Kong and Monteiro, 2021; Zhao, 2023). Among the existing works, we highlight the work of Grnarova et al. (2018), who proposed a variant of GDA for this class of nonconvex minimax optimization problems and the works of Sinha et al. (2018) and Sanjabi et al. (2018), who studied inexact SGD algorithms (i.e., variants of SGDmax) and established a convergence guarantee in terms of iteration numbers. Subsequently, Jin et al. (2020) proved the refined convergence results for GDmax and SGDmax in terms of (stochastic) gradient evaluations.

There are some other algorithms developed for solving smooth and nonconvex-concave minimax optimization problems. For example, Rafique et al. (2021) proposed two stochastic algorithms and prove that they efficiently find a stationary point of $\Phi(\cdot) = \max_{\mathbf{y} \in \mathcal{Y}} f(\cdot, \mathbf{y})$. However, these algorithms are nested-loop algorithms and the convergence guarantees hold only when $f(\mathbf{x}, \cdot)$ is linear and $\mathcal{Y} = \mathbb{R}^n$ (see Rafique et al. (2021, Assumption 2)). Nouiehed et al. (2019) have designed a multistep variant of GDA by incorporating momentum terms into the scheme.¹ In this context, the best-known gradient complexity is $\tilde{O}(\epsilon^{-3})$ and has been achieved by the ProxDIAG algorithm (Thekumparampil et al., 2019). This complexity result was also achieved by an accelerated inexact proximal point algorithm. Indeed, Kong and Monteiro (2021) showed that their algorithm finds one ϵ -stationary point of Φ within $O(\epsilon^{-3})$ iterations.² Solving the subproblem at each iteration is equivalent to minimizing a smooth and strongly convex function and can be approximately done using $\tilde{O}(1)$ gradient

1. Nouiehed et al. (2019) considered a setting where (i) \mathbf{x} is constrained to a convex set \mathcal{X} ; (ii) $\mathcal{Y} = \mathbb{R}^n$; and (iii) f is smooth in \mathbf{x} over \mathcal{X} and satisfies the Polyak-Lojasiewicz condition in \mathbf{y} .

2. Kong and Monteiro (2021) considered a more general setting than ours, where \mathbf{x} is constrained to a convex set \mathcal{X} and f is smooth in \mathbf{x} over \mathcal{X} . Their analysis and results are valid in our setting.

evaluations. Combining these two facts yields the desired claim. Subsequent to Lin et al. (2020a), this complexity bound was also achieved by other algorithms (Lin et al., 2020b; Yang et al., 2020b; Ostrovskii et al., 2021b; Zhao, 2023). Despite the appealing theoretical convergence guarantee, the above algorithms are all nested-loop algorithms and thus can be complex to implement. One would like to understand whether the nested-loop structure is necessary or whether a single-loop algorithm provably converges in smooth and nonconvex-concave settings. An example of such an investigation is Lu et al. (2020), where a single-loop algorithm was developed for solving smooth and nonconvex-concave minimax optimization problems. However, their analysis requires some restrictive conditions; e.g., that $f(\cdot, \cdot)$ is lower bounded. In contrast, we only require that $\max_{\mathbf{y} \in \mathcal{Y}} f(\cdot, \mathbf{y})$ is lower bounded. An example which meets our conditions rather than theirs is $\min_{\mathbf{x} \in \mathbb{R}} \max_{\mathbf{y} \in [-1, 1]} \mathbf{x}^\top \mathbf{y}$. Our less-restrictive conditions make the problem more challenging and our analysis accordingly differs significantly from theirs.

To our knowledge, we are the first to analyze TTGDA and TTSGDA in the nonsmooth and nonconvex-concave setting where f is Lipschitz and $f(\cdot, \mathbf{y})$ is weakly convex. As already noted, Rafique et al. (2021) considered the nonsmooth setting but focused on nested-loop algorithms. Barazandeh and Razaviyayn (2020) and Huang et al. (2021) proposed to study the composite nonconvex-concave setting with nonsmooth regularization terms which are however structured such that the associated proximal mappings can be computed efficiently. The closest work to ours is Boş and Böhm (2023) who studied the alternating TTGDA and TTSGDA in the nonsmooth and nonconvex-concave setting. However, their setting is different from ours—they assume that $f(\cdot, \mathbf{y})$ is Lipschitz and weakly convex but $f(\mathbf{x}, \cdot)$ is smooth. Their results demonstrate that the algorithms can achieve the same complexity bound in this setting as in the smooth setting. In contrast, as we will show in this paper, relaxing the smoothness of f makes the problem substantially more difficult, especially for deterministic algorithms.

Nonconvex-nonconcave setting. The investigation of nonconvex-nonconcave minimax optimization problems has become a central topic in the machine learning community, inspired by the advent of generative adversarial networks (GANs) (Goodfellow et al., 2014) and adversarial learning (Madry et al., 2018; Sinha et al., 2018). In this context, most works focus on defining a notion of goodness or the development of new procedures for reducing oscillations (Daskalakis and Panageas, 2018; Adolphs et al., 2019; Jin et al., 2020) and speeding up the convergence of gradient dynamics (Heusel et al., 2017; Balduzzi et al., 2018; Mertikopoulos et al., 2019). For example, Daskalakis and Panageas (2018) showed that the stable fixed points of GDA are not necessarily saddle points, while Adolphs et al. (2019) proposed Hessian-based algorithms where their stable fixed points are saddle points. In this context, Balduzzi et al. (2018) also developed a symplectic gradient adjustment algorithm for finding stable fixed points in potential games and Hamiltonian games. Notably, Heusel et al. (2017) have proposed TTGDA and TTSGDA formally for training GANs and showed that the saddle points are also stable fixed points of a continuous limit of these algorithms under strong conditions. All of these convergence results are either local or asymptotic and thus can not cover our nonasymptotic results for TTGDA and TTSGDA in the nonconvex-concave settings. Mertikopoulos et al. (2019) has provided nonasymptotic guarantees for a special class of nonconvex-nonconcave minimax optimization problems under variational

stability and/or the so-called Minty condition. However, while these conditions must hold in convex-concave setting, they do not necessarily hold in the nonconvex-concave setting. Subsequent to our earlier work (Lin et al., 2020a), we are aware of significant effort that has been devoted to understanding the nonconvex-nonconcave minimax optimization problem (Yang et al., 2020a; Diakonikolas et al., 2021; Liu et al., 2021; Lee and Kim, 2021; Ostrovskii et al., 2021a; Fiez et al., 2022; Jiang and Chen, 2023; Grimmer et al., 2023; Daskalakis et al., 2023; Zheng et al., 2023; Cai et al., 2024). Some of their results are valid for the nonconvex-concave setting but can not cover our results in this paper. This is an active research area and we refer the interested reader to the aforementioned articles for more details.

Other settings. Analysis in the online learning perspective aims at understanding whether an algorithm enjoys the no-regret property. For example, the optimistic GDA from Daskalakis and Panageas (2019) is a no-regret learning algorithm, while the extragradient algorithm from Mertikopoulos et al. (2019) is not. Moreover, by carefully comparing limit behavior of algorithm dynamics, Bailey and Piliouras (2018) have shown that the multiplicative weights update behaves similarly to single-timescale GDA, thus justifying the need for introducing the optimistic update for achieving the last-iterate convergence guarantee. There are also several works on finite-sum minimax optimization and variance reduction (Luo et al., 2020; Alacaoglu and Malitsky, 2022; Huang et al., 2022; Han et al., 2024).

3. Preliminaries and Technical Background

We present standard definitions for Lipschitz functions and smooth functions.

Definition 1. f is L -Lipschitz if for all \mathbf{x}, \mathbf{x}' , we have $\|f(\mathbf{x}) - f(\mathbf{x}')\| \leq L\|\mathbf{x} - \mathbf{x}'\|$.

Definition 2. f is ℓ -smooth if for all \mathbf{x}, \mathbf{x}' , we have $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}')\| \leq \ell\|\mathbf{x} - \mathbf{x}'\|$.

The minimax optimization problem in Eq. (1.1) can be rewritten as $\min_{\mathbf{x} \in \mathbb{R}^m} \Phi(\mathbf{x})$ where $\Phi(\cdot) = \max_{\mathbf{y} \in \mathcal{Y}} f(\cdot, \mathbf{y})$. If $f(\mathbf{x}, \cdot)$ is concave for each $\mathbf{x} \in \mathbb{R}^m$, the problem of $\max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y})$ can be approximately solved efficiently, thus providing useful information about Φ . However, it is impossible to find any reasonable approximation of the global minimum of Φ in general since Φ can be nonconvex and nonsmooth (Murty and Kabadi, 1987; Zhang et al., 2020a). This motivates us to pursue a local surrogate for the global minimum of Φ .

A common surrogate in nonconvex optimization is the notion of stationarity, which is appropriate if Φ is differentiable.

Definition 3. A point \mathbf{x} is an ϵ -stationary point of a differentiable function Φ for some tolerance $\epsilon \geq 0$ if $\|\nabla \Phi(\mathbf{x})\| \leq \epsilon$. If $\epsilon = 0$, we have that \mathbf{x} is a stationary point.

Definition 3 is suitable in the smooth and nonconvex-strongly-concave setting since Φ is differentiable. However, it is worth remarking that Φ is not necessarily differentiable in the other three settings: (i) smooth and nonconvex-concave; (ii) nonsmooth and nonconvex-strongly-concave; and (iii) nonsmooth and nonconvex-concave. Thus, we need to introduce a weaker condition as follows.

Definition 4. A function Φ is ρ -weakly convex if the function $\Phi(\cdot) + 0.5\rho\|\cdot\|^2$ is convex.

For a ρ -weakly convex function Φ , we know that its Clarke subdifferential $\partial\Phi$ is uniquely determined by the subdifferential of the convex function $\Phi(\cdot) + 0.5\rho\|\cdot\|^2$. This then implies that a naive measure of ϵ -approximate stationarity can be defined as a point \mathbf{x} satisfying $\min_{\xi \in \partial\Phi(\mathbf{x})} \|\xi\| \leq \epsilon$ (i.e., at least one subgradient is small). However, this notion is restrictive for nonsmooth optimization. For example, when $\Phi(\cdot) = |\cdot|$ is a one-dimensional function, an ϵ -stationary point is 0 for all $\epsilon \in [0, 1)$. Therefore, we see that finding an ϵ -approximate stationary point of Φ under this notion is as hard as finding an exact stationary point even when $\epsilon > 0$ is not very small.

In respond to this issue, Davis and Drusvyatskiy (2019) employed an alternative notion of stationarity based on the Moreau envelope. Their notion can be suitable for our settings for two reasons: (i) the problem of $\max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y})$ can be approximately solved efficiently, thus providing useful information about Φ and its Moreau envelope; (ii) a stationary point based on the Moreau envelope has practical implications in real machine learning applications. For example, \mathbf{x} and \mathbf{y} stand for a classifier and an adversarial example arising from adversarial machine learning models. In general, practitioners are interested in finding a robust classifier \mathbf{x} rather than recovering the adversarial example \mathbf{y} .

Definition 5. A function Φ_λ is the Moreau envelope of Φ with the parameter $\lambda > 0$ if $\Phi_\lambda(\cdot) = \min_{\mathbf{w}} \Phi(\mathbf{w}) + \frac{1}{2\lambda}\|\mathbf{w} - \cdot\|^2$.

In what follows, we prove two lemmas for the nonconvex-concave settings. For simplicity, we use $\Phi(\cdot) = \max_{\mathbf{y} \in \mathcal{Y}} f(\cdot, \mathbf{y})$ throughout in the subsequent statements.

Lemma 6. If $f(\mathbf{x}, \mathbf{y})$ is ℓ -smooth, concave in \mathbf{y} , and \mathcal{Y} is convex and bounded, we have that Φ is ℓ -weakly convex and the following statements hold true,

1. $\Phi(\text{prox}_{\Phi/2\ell}(\mathbf{x})) \leq \Phi(\mathbf{x})$ where $\text{prox}_{\Phi/2\ell}(\cdot) = \text{argmin}_{\mathbf{w}} \Phi(\mathbf{w}) + \ell\|\mathbf{w} - \cdot\|^2$.
2. $\Phi_{1/2\ell}$ is ℓ -smooth with $\nabla\Phi_{1/2\ell}(\mathbf{x}) = 2\ell(\mathbf{x} - \text{prox}_{\Phi/2\ell}(\mathbf{x}))$.

Lemma 7. If $f(\mathbf{x}, \mathbf{y})$ is ρ -weakly convex in \mathbf{x} , concave in \mathbf{y} , and \mathcal{Y} is convex and bounded, we have that Φ is ρ -weakly convex and the following statements hold true,

1. $\Phi(\text{prox}_{\Phi/2\rho}(\mathbf{x})) \leq \Phi(\mathbf{x})$ where $\text{prox}_{\Phi/2\rho}(\cdot) = \text{argmin}_{\mathbf{w}} \Phi(\mathbf{w}) + \rho\|\mathbf{w} - \cdot\|^2$.
2. $\Phi_{1/2\rho}$ is ρ -smooth with $\nabla\Phi_{1/2\rho}(\mathbf{x}) = 2\rho(\mathbf{x} - \text{prox}_{\Phi/2\rho}(\mathbf{x}))$.

Based on Lemmas 6 and 7, an alternative measure of approximate stationarity of a ρ -weakly convex function Φ can be defined as a point \mathbf{x} satisfying $\|\nabla\Phi_{1/2\rho}(\mathbf{x})\| \leq \epsilon$.

Definition 8. A point \mathbf{x} is an ϵ -stationary point of a ρ -weakly convex function Φ for some tolerance $\epsilon \geq 0$ if $\|\nabla\Phi_{1/2\rho}(\mathbf{x})\| \leq \epsilon$. If $\epsilon = 0$, we have that \mathbf{x} is a stationary point.

Although Definition 8 uses the Moreau envelope $\Phi_{1/2\rho}$ to define stationarity, it connects to approximate stationarity of Φ in terms of small subgradients as follows.

Lemma 9. If \mathbf{x} is an ϵ -stationary point of a ρ -weakly convex function Φ , there exists $\hat{\mathbf{x}}$ such that $\min_{\xi \in \partial\Phi(\hat{\mathbf{x}})} \|\xi\| \leq \epsilon$ and $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq \frac{\epsilon}{2\rho}$.

Lemma 9 shows that Definition 8 gives an intuitive surrogate for the naive measure of ϵ -approximate stationarity. Indeed, if \mathbf{x} is an ϵ -stationary point of a ρ -weakly convex function Φ , it must be $\frac{\epsilon}{2\rho}$ -close to a point $\hat{\mathbf{x}}$ which has at least one small subgradient.

There is another notion of stationarity based on ∇f in smooth and nonconvex-concave setting (Nouiehed et al., 2019; Lu et al., 2020; Lin et al., 2020b; Xu et al., 2023). A pair of points $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^m \times \mathcal{Y}$ is an ϵ -stationary point of f if $\|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y})\| \leq \epsilon$ and $\|\mathbf{y}^+ - \mathbf{y}\| \leq \frac{\epsilon}{\ell}$ where \mathbf{y}^+ is defined by

$$\mathbf{y}^+ = \mathcal{P}_{\mathcal{Y}}(\mathbf{y} + \frac{1}{\ell}\nabla_{\mathbf{y}}f(\mathbf{x}, \mathbf{y})).$$

Intuitively, the term $\ell\|\mathbf{y}^+ - \mathbf{y}\|$ stands for the distance between $\mathbf{y} \in \mathcal{Y}$ and \mathbf{y}^+ obtained by performing a one-step projected partial gradient ascent at (\mathbf{x}, \mathbf{y}) starting from \mathbf{y} . This is akin to the norm of the gradient mapping (Nesterov, 2013). As highlighted by Thekumparampil et al. (2019), this notion is weaker than Definitions 3 and 8. For the sake of completeness, we provide a complete characterization of the relationship between these notions by showing that they can be translated in both directions with the extra computational cost. To pursue this goal, we introduce a new notion of stationarity based on ∇f as follows.

Definition 10. *A pair of points $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^m \times \mathcal{Y}$ is an ϵ -stationary point of a function f for some tolerance $\epsilon \geq 0$ if $\|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}^+)\| \leq \epsilon$ and $\|\mathbf{y}^+ - \mathbf{y}\| \leq \frac{\epsilon}{\ell}$.*

The stationarity of f is in fact equivalent to Definition 10 up to a constant. Indeed, if $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^m \times \mathcal{Y}$ is an ϵ -stationary point as per Definition 10, we have $\|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y})\| \leq \|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}^+)\| + \ell\|\mathbf{y}^+ - \mathbf{y}\| \leq 2\epsilon$. Conversely, if $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^m \times \mathcal{Y}$ satisfies $\|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y})\| \leq \epsilon$ and $\|\mathbf{y}^+ - \mathbf{y}\| \leq \frac{\epsilon}{\ell}$, we have $\|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}^+)\| \leq \|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y})\| + \ell\|\mathbf{y}^+ - \mathbf{y}\| \leq 2\epsilon$.

We focus on the notion in Definition 10 and relate it to the stationarity of Φ (cf. Definition 3 and 8). The following propositions summarize the details.

Proposition 11. *Suppose that f is ℓ -smooth, μ -strongly concave in \mathbf{y} , and \mathcal{Y} is convex and bounded. Given that $\hat{\mathbf{x}}$ is an ϵ -stationary point as per Definition 3, we can obtain a 2ϵ -stationary point $(\mathbf{x}', \mathbf{y}')$ as per Definition 10 using $O(\kappa \log(1/\epsilon))$ gradient evaluations or $O(\epsilon^{-2})$ stochastic gradient evaluations. Given that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a $\frac{\epsilon}{2\kappa}$ -stationary point as per Definition 10, we have that $\hat{\mathbf{x}}$ is an 2ϵ -stationary point as per Definition 3.*

Proposition 12. *Suppose that f is ℓ -smooth, concave in \mathbf{y} , and \mathcal{Y} is convex and bounded. Given that $\hat{\mathbf{x}}$ is an ϵ -stationary point as per Definition 8, we can obtain a 4ϵ -stationary point $(\mathbf{x}', \mathbf{y}')$ as per Definition 10 using $O(\epsilon^{-2})$ gradient evaluations or $O(\epsilon^{-4})$ stochastic gradient evaluations. Given that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a $\frac{\epsilon^2}{\max\{1, \ell D\}}$ -stationary point as per Definition 10, we have that $\hat{\mathbf{x}}$ is a $(4\epsilon^2 + 4\epsilon)$ -stationary point as per Definition 8.*

Propositions 11 and 12 show that one needs to pay an additional $O(\kappa \log(1/\epsilon))$ or $O(\epsilon^{-2})$ gradient evaluations for translating the stationarity of f to that of Φ in nonconvex-strongly-concave and nonconvex-concave settings. In this sense, the notions of stationarity as per Definitions 3 and 8 are stronger than that as per Definition 10.

Remark 13. *Ostrovskii et al. (2021b) identified an error in the proof of Lin et al. (2020a, Proposition 4.12) and proposed a new notion of stationarity based on ∇f which is stronger than the existing stationarity of f . After communicating with the authors of Ostrovskii et al. (2021b), we agree that Lin et al. (2020a, Proposition 4.12) (i.e., the precursor of*

Proposition 12) is still valid and the error can be fixed by using the notion in Definition 10. We also acknowledge that their notion is strictly stronger than ours in Definition 10 when $\mathcal{Y} \neq \mathbb{R}^n$ while being tractable from a computational viewpoint.

4. Smooth Minimax Optimization

We prove the bounds on the complexity of solving smooth and nonconvex-concave minimax optimization problems. Our focus is TTGDA and TTSGDA (see Algorithms 1 and 2) with a proper choice of $(\eta_{\mathbf{x}}^t, \eta_{\mathbf{y}}^t)_{t \geq 0}$; recall that these methods were shown to converge locally and are useful in practice (Heusel et al., 2017). Intuitively, the reason for this phenomenon is that the choice of $\eta_{\mathbf{x}}^t \ll \eta_{\mathbf{y}}^t$ matches the asymmetric structure of nonconvex-concave minimax optimization problems.

4.1 Main results

We make the following assumptions throughout this subsection.

Assumption 14. *The function $f(\mathbf{x}, \mathbf{y})$ is ℓ -smooth and μ -strongly concave in \mathbf{y} , and the constraint set \mathcal{Y} is convex and bounded with a diameter $D > 0$.*

Assumption 15. *The function $f(\mathbf{x}, \mathbf{y})$ is ℓ -smooth, L -Lipschitz in \mathbf{x} and concave in \mathbf{y} , and the constraint set \mathcal{Y} is convex and bounded with a diameter $D > 0$.*

In addition to smoothness condition, the function $f(\mathbf{x}, \mathbf{y})$ is assumed to be concave in \mathbf{y} but nonconvex in \mathbf{x} . Imposing such one-side concavity is necessary since Daskalakis et al. (2021) have shown that it is NP-hard to determine if an approximate min-max point exists in a general smooth and nonconvex-nonconcave setting. Imposing the boundedness condition for \mathcal{Y} is an assumption of convenience that reflects the fact that the ambiguity sets are generally bounded in real applications (Sinha et al., 2018). Relaxing this condition requires more complicated arguments and we leave it to future work.

For the ease of presentation, we define

$$\Phi(\mathbf{x}) = \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y}), \quad \mathcal{Y}^*(\mathbf{x}) = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y}), \quad \text{for each } \mathbf{x} \in \mathbb{R}^m.$$

We present a key lemma on the structure of the function $\Phi(\cdot)$ and the set function $\mathcal{Y}^*(\cdot)$.

Lemma 16. *The following statements hold true,*

1. *Under Assumption 14, we have that $\mathcal{Y}^*(\mathbf{x})$ is a singleton for each \mathbf{x} such that $\mathcal{Y}^*(\mathbf{x}) = \{\mathbf{y}^*(\mathbf{x})\}$ where $\mathbf{y}^*(\cdot)$ is κ -Lipschitz. Also, $\Phi(\cdot)$ is $(2\kappa\ell)$ -smooth and $\nabla\Phi(\cdot) = \nabla_{\mathbf{x}}f(\cdot, \mathbf{y}^*(\cdot))$.*
2. *Under Assumption 15, we have that $\mathcal{Y}^*(\mathbf{x})$ contains many entries for each \mathbf{x} such that $\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}) \in \partial\Phi(\mathbf{x})$ where $\mathbf{y} \in \mathcal{Y}^*(\mathbf{x})$. Also, $\Phi(\cdot)$ is ℓ -weakly convex and L -Lipschitz.*

Lemma 16 shows that the function $\Phi(\cdot)$ is smooth under Assumption 14, implying that the stationarity notion in Definition 3 is our target. Denoting $\Delta_{\Phi} = \Phi(\mathbf{x}_0) - \min_{\mathbf{x}} \Phi(\mathbf{x})$, we summarize our results for Algorithm 1 and 2 in the following theorem.

Theorem 17. *Under Assumption 14, the following statements hold true,*

1. If we set

$$\eta_{\mathbf{x}}^t = \frac{1}{16(\kappa+1)^2\bar{\ell}}, \quad \eta_{\mathbf{y}}^t = \frac{1}{\bar{\ell}}, \quad (4.1)$$

in the TTGDA scheme (see Algorithm 1), the required number of gradient evaluations to return an ϵ -stationary point is

$$O\left(\frac{\kappa^2\ell\Delta_{\Phi} + \kappa\ell^2D^2}{\epsilon^2}\right).$$

2. If we set $\eta_{\mathbf{x}}^t$ and $\eta_{\mathbf{y}}^t$ as per Eq. (4.1), $M = \max\{1, 48\kappa\sigma^2\epsilon^{-2}\}$ and G as per Eq. (1.3) in the TTSGDA scheme (see Algorithm 2), the required number of stochastic gradient evaluations to return an ϵ -stationary point is

$$O\left(\frac{\kappa^2\ell\Delta_{\Phi} + \kappa\ell^2D^2}{\epsilon^2} \max\left\{1, \frac{\kappa\sigma^2}{\epsilon^2}\right\}\right).$$

Remark 18. Theorem 17 provides the complexity results for TTGDA and TTSGDA when applied to solve the smooth and nonconvex-strongly-concave minimax optimization problems. The ratio $\frac{\eta_{\mathbf{y}}^t}{\eta_{\mathbf{x}}^t} = \Theta(\kappa^2)$ arises due to asymmetric structure and the order of $O(\kappa^2)$ reflects an algorithmic efficiency trade-off. Furthermore, our algorithms are only guaranteed to visit an ϵ -stationary point within a certain number of iterations. Such the convergence guarantees are standard in nonconvex optimization and practitioners usually return an iterate at which the learning curve stops changing significantly. Finally, the batch size $M = \Theta(\kappa\sigma^2\epsilon^{-2})$ is necessary for proving Theorem 17. While our proof technique can be extended to the case of $M = 1$, the complexity bound becomes $O(\kappa^3\epsilon^{-5})$. This gap has been closed using other new single-loop algorithms (Yang et al., 2022b) but remains open for TTSGDA.

Lemma 16 shows that the function $\Phi(\cdot)$ is weakly convex under Assumption 15, implying that the stationarity notion in Definition 8 is our target. We summarize our results for Algorithm 1 and 2 in the following theorem.

Theorem 19. Under Assumption 15, the following statements hold true,

1. If we set

$$\eta_{\mathbf{x}}^t = \min\left\{\frac{\epsilon^2}{80\ell L^2}, \frac{\epsilon^4}{4096\ell^3 L^2 D^2}\right\}, \quad \eta_{\mathbf{y}}^t = \frac{1}{\bar{\ell}},$$

in the TTGDA scheme (see Algorithm 1), the required number of gradient evaluations to return an ϵ -stationary point is

$$O\left(\left(\frac{\ell L^2 \Delta_{\Phi}}{\epsilon^4} + \frac{\ell \Delta_0}{\epsilon^2}\right) \max\left\{1, \frac{\ell^2 D^2}{\epsilon^2}\right\}\right).$$

2. If we set $M = 1$, G as per Eq. (1.3) and

$$\eta_{\mathbf{x}}^t = \min\left\{\frac{\epsilon^2}{80\ell(L^2 + \sigma^2)}, \frac{\epsilon^4}{8192\ell^3(L^2 + \sigma^2)D^2}, \frac{\epsilon^6}{131072\ell^3(L^2 + \sigma^2)D^2\sigma^2}\right\}, \quad \eta_{\mathbf{y}}^t = \min\left\{\frac{1}{2\bar{\ell}}, \frac{\epsilon^2}{32\ell\sigma^2}\right\},$$

in the TTSGDA scheme (see Algorithm 2), the required number of stochastic gradient evaluations to return an ϵ -stationary point is

$$O\left(\left(\frac{\ell(L^2 + \sigma^2)\Delta_{\Phi}}{\epsilon^4} + \frac{\ell\Delta_0}{\epsilon^2}\right) \max\left\{1, \frac{\ell^2 D^2}{\epsilon^2}, \frac{\ell^2 D^2 \sigma^2}{\epsilon^4}\right\}\right).$$

Remark 20. *Theorem 19 is the analog of Theorem 17 in the nonconvex-concave setting but with two different results. First, the order of the ratio becomes $\eta_{\mathbf{y}}/\eta_{\mathbf{x}} = \Theta(\epsilon^4)$ which highlights the highly asymmetric structure here. Second, the theoretical results are valid for the case of $M = 1$ and in that case they are different from that derived in Theorem 17.*

4.2 Discussions

We remark that we have not shown that TTGDA and TTSGDA are optimal in any sense, nor that acceleration might be achieved by using momentum for updating the variable \mathbf{y} .

In smooth and nonconvex-strongly-concave setting, the lower bound is $\Omega(\sqrt{\kappa}\epsilon^{-2})$ in terms of gradient evaluations (Zhang et al., 2021; Li et al., 2021) and was already achieved by several nested-loop algorithms (Kong and Monteiro, 2021; Lin et al., 2020b; Ostrovskii et al., 2021b; Zhao, 2023) up to logarithmic factors. In the context of single-loop algorithms, Yang et al. (2022b) has extended the smoothed GDA (Zhang et al., 2020b) to more general settings where the function $f(\mathbf{x}, \mathbf{y})$ is nonconvex in \mathbf{x} and satisfies the Polyak-Łojasiewicz condition in \mathbf{y} and proved a bound of $O(\kappa\epsilon^{-2})$. This result demonstrates that the proper adjustments of TTGDA leads to tighter bounds in smooth and nonconvex-strongly-concave setting. In addition, Li et al. (2022) provided a local analysis for TTGDA and showed that the stepsize ratio of $\Theta(\kappa)$ is necessary and sufficient for local convergence to a stationary point.

In smooth and nonconvex-concave setting, the lower bound remains open and the best-known upper bound is presented by recent works (Thekumparampil et al., 2019; Kong and Monteiro, 2021; Lin et al., 2020b; Ostrovskii et al., 2021b; Zhao, 2023). All of the analyses require $\tilde{O}(\epsilon^{-3})$ gradient evaluations for returning an ϵ -stationary point but the proposed algorithmic schemes are complex. The general question of constructing lower bounds and optimal algorithms in this setting is an important topic for future research.

Our complexity results do not contradict the classical results on the divergence of GDA with *fixed stepsize* in the convex-concave setting. Indeed, there are a few key distinctions: (1) our results guarantee that the proposed algorithm visits an ϵ -stationary point at some iterate, not the last iterate; (2) our results guarantee the stationarity of some iterate \mathbf{x}_t rather than $(\mathbf{x}_t, \mathbf{y}_t)$. As such, our proof permits the possibility of significant changes in \mathbf{y}_t even when \mathbf{x}_t is very close to stationarity. This, together with our choice $\eta_{\mathbf{x}}^t \ll \eta_{\mathbf{y}}^t$, removes the contradiction. Based on the above fact, we highlight that our algorithms can be used to obtain an approximate saddle point in the smooth and convex-concave setting (i.e., optimality for both \mathbf{x} and \mathbf{y}). Instead of averaging, we can run two passes of TTGDA for minimax and maximin optimization problems separately. We use $\eta_{\mathbf{x}}^t \ll \eta_{\mathbf{y}}^t$ in the first pass while we use $\eta_{\mathbf{x}}^t \gg \eta_{\mathbf{y}}^t$ in the second pass. Either pass returns an approximate stationary point for each variable \mathbf{x} and/or \mathbf{y} , which is defined using the corresponding Moreau envelope and which jointly forms an approximate saddle point.

4.3 Proof sketch

We sketch the proofs for TTGDA by highlighting the ideas for proving Theorems 17 and 19.

Nonconvex-strongly-concave setting. We have that the function $\mathbf{y}^*(\cdot)$ is κ -Lipschitz (cf. Lemma 16), which guarantees that $\{\mathbf{y}^*(\mathbf{x}_t)\}_{t \geq 1}$ moves slowly if $\{\mathbf{x}_t\}_{t \geq 1}$ moves slowly. This implies that one-step gradient ascent over \mathbf{y} at each iteration is sufficient for minimizing

this slowly changing sequence of strongly concave functions $\{f(\mathbf{x}_t, \cdot)\}_{t>1}$. Indeed, we define $\delta_t = \|\mathbf{y}^*(\mathbf{x}_t) - \mathbf{y}_t\|^2$ and prove that $\sum_{t=0}^T \delta_t$ can be controlled using the relations as follows,

$$\delta_t \leq \left(1 - \frac{1}{2\kappa} + 4\kappa^3 \ell^2 (\eta_{\mathbf{x}}^t)^2\right) \delta_{t-1} + 4\kappa^3 (\eta_{\mathbf{x}}^t)^2 \|\nabla \Phi(\mathbf{x}_{t-1})\|^2,$$

where $\eta_{\mathbf{x}}^t$ can be chosen such that $\sum_{t=0}^T \delta_t$ is controlled by $\sum_{t=0}^T \|\nabla \Phi(\mathbf{x}_t)\|^2$.

By viewing TTGDA as inexact gradient descent for minimizing the function $\Phi(\cdot)$, we derive the key descent inequality as follows,

$$\Phi(\mathbf{x}_{T+1}) - \Phi(\mathbf{x}_0) \leq -\frac{7\eta_{\mathbf{x}}^t}{16} \left(\sum_{t=0}^T \|\nabla \Phi(\mathbf{x}_t)\|^2 \right) + \frac{9(\eta_{\mathbf{x}}^t)\ell^2}{16} \left(\sum_{t=0}^T \delta_t \right).$$

Putting these pieces together yields the desired results in Theorem 17.

Nonconvex-concave setting. We note that $\mathcal{Y}^*(\mathbf{x})$ will not be a singleton for each \mathbf{x} and $\mathcal{Y}^*(\mathbf{x}_1)$ can be dramatically different from $\mathcal{Y}^*(\mathbf{x}_2)$ even when $\mathbf{x}_1, \mathbf{x}_2$ are close to each other. This implies that $\min_{\mathbf{y} \in \mathcal{Y}^*(\mathbf{x}_t)} \|\mathbf{y}_t - \mathbf{y}\|$ is no longer a viable error to control.

However, we still have that one-step gradient ascent over \mathbf{y} at each iteration is sufficient for minimizing this slowly changing sequence of concave functions $\{f(\mathbf{x}_t, \cdot)\}_{t>1}$. Indeed, we have that the function $\Phi(\cdot)$ is L -Lipschitz (cf. Lemma 23), which guarantees that $\{\Phi(\mathbf{x}_t)\}_{t \geq 1}$ moves slowly if $\{\mathbf{x}_t\}_{t \geq 1}$ moves slowly. We can define $\Delta_t = \Phi(\mathbf{x}_t) - f(\mathbf{x}_t, \mathbf{y}_t)$ and prove that $\frac{1}{T+1}(\sum_{t=0}^T \Delta_t)$ can be controlled using the relations as follows,

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}}^t (B+1)L^2 + \frac{\ell D^2}{2B} + \frac{\Delta_0}{T+1},$$

where $B = \left\lceil \frac{D}{2L} \sqrt{\frac{\ell}{\eta_{\mathbf{x}}^t}} \right\rceil$ and $\eta_{\mathbf{x}}^t$ can be chosen such that $\frac{1}{T+1}(\sum_{t=0}^T \Delta_t)$ is well controlled.

By viewing TTGDA as an inexact subgradient descent for minimizing the function $\Phi(\cdot)$, we follow Davis and Drusvyatskiy (2019) and derive the key descent inequality as follows,

$$\Phi_{1/2\ell}(\mathbf{x}_{T+1}) - \Phi_{1/2\ell}(\mathbf{x}_0) \leq 2\eta_{\mathbf{x}}\ell \left(\sum_{t=0}^T \Delta_t \right) + \eta_{\mathbf{x}}^2 \ell L^2 (T+1) - \frac{\eta_{\mathbf{x}}}{4} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2 \right).$$

Putting these pieces together yields the desired results in Theorem 19.

5. Nonsmooth Minimax Optimization

We prove the bounds on the complexity of solving nonsmooth and nonconvex-concave minimax optimization problems. Indeed, we assume that $f(\mathbf{x}, \mathbf{y})$ is L -Lipschitz and ρ -weakly convex in \mathbf{x} . Our focus is TTGDA and TTSGDA with a proper choice of $(\eta_{\mathbf{x}}^t, \eta_{\mathbf{y}}^t)_{t \geq 0}$.

5.1 Main results

We make the following assumptions throughout this subsection.

Assumption 21. *The function $f(\mathbf{x}, \mathbf{y})$ is L -Lipschitz, ρ -weakly convex in \mathbf{x} and μ -strongly concave in \mathbf{y} , and the constraint set \mathcal{Y} is convex and bounded with a diameter $D > 0$.*

Assumption 22. *The function $f(\mathbf{x}, \mathbf{y})$ is L -Lipschitz, ρ -weakly convex in \mathbf{x} and concave in \mathbf{y} , and the constraint set \mathcal{Y} is convex and bounded with a diameter $D > 0$.*

The setting defined by Assumptions 21 and 22 generalizes the nonsmooth optimization and are intuitively more challenging than smooth minimax optimization studied in Section 4. The following lemma is an analog of Lemma 16.

Lemma 23. *The following statements hold true,*

1. *Under Assumption 21, we have that $\mathcal{Y}^*(\mathbf{x})$ is a singleton for each \mathbf{x} such that $\mathcal{Y}^*(\mathbf{x}) = \{\mathbf{y}^*(\mathbf{x})\}$ and $\partial_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}^*(\mathbf{x})) \subseteq \partial\Phi(\mathbf{x})$. Also, $\Phi(\cdot)$ is L -Lipschitz and ρ -weakly convex.*
2. *Under Assumption 22, we have that $\mathcal{Y}^*(\mathbf{x})$ contains many entries for each \mathbf{x} such that $\partial_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}) \subseteq \partial\Phi(\mathbf{x})$ where $\mathbf{y} \in \mathcal{Y}^*(\mathbf{x})$. Also, $\Phi(\cdot)$ is L -Lipschitz and ρ -weakly convex.*

Lemma 23 shows that the function $\Phi(\cdot)$ is weakly convex in both of settings, implying that the stationarity notion in Definition 8 is our target. We summarize our results for Algorithm 1 and 2 in the following theorem.

Theorem 24. *Under Assumption 21, the following statements hold true,*

1. *If we set*

$$\eta_{\mathbf{x}}^t = \min \left\{ \frac{\epsilon^2}{48\rho L^2}, \frac{\mu\epsilon^4}{4096\rho^2 L^4}, \frac{\mu\epsilon^4}{4096\rho^2 L^4 \log^2(1+4096\rho^2 L^4 \mu^{-2}\epsilon^{-4})} \right\},$$

and

$$\eta_{\mathbf{y}}^t = \begin{cases} \frac{1}{\mu^t}, & \text{if } 1 \leq t \leq B, \\ \frac{1}{\mu^{(t-B)}}, & \text{else if } B+1 \leq t \leq 2B, \\ \vdots & \vdots \\ \frac{1}{\mu^{(t-jB)}}, & \text{else if } jB+1 \leq t \leq (j+1)B, \\ \vdots & \vdots \end{cases} \quad \text{for } B = \left\lfloor \sqrt{\frac{1}{\mu\eta_{\mathbf{x}}}} \right\rfloor + 1, \quad (5.1)$$

in the TTGDA scheme (see Algorithm 1), the required number of gradient evaluations to return an ϵ -stationary point is

$$O \left(\frac{\rho L^2 \Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho L^2}{\mu\epsilon^2}, \frac{\rho L^2}{\mu\epsilon^2} \log^2 \left(1 + \frac{\rho^2 L^4}{\mu^2 \epsilon^4} \right) \right\} \right).$$

2. *If we set $M = 1$, G as per Eq. (1.3), $\eta_{\mathbf{y}}^t$ as per Eq. (5.1) and*

$$\eta_{\mathbf{x}}^t = \min \left\{ \frac{\epsilon^2}{48\rho(L^2+\sigma^2)}, \frac{\mu\epsilon^4}{4096\rho^2(L^2+\sigma^2)^2}, \frac{\mu\epsilon^4}{4096\rho^2(L^2+\sigma^2)^2 \log^2(1+4096\rho^2(L^2+\sigma^2)^2 \mu^{-2}\epsilon^{-4})} \right\},$$

in the TTSFDA scheme (see Algorithm 2), the required number of stochastic gradient evaluations to return an ϵ -stationary point is

$$O \left(\frac{\rho(L^2+\sigma^2)\Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho^2(L^2+\sigma^2)}{\mu\epsilon^2}, \frac{\rho^2(L^2+\sigma^2)}{\mu\epsilon^2} \log^2 \left(1 + \frac{\rho^2(L^2+\sigma^2)^2}{\mu^2\epsilon^4} \right) \right\} \right).$$

Remark 25. *Theorem 24 shows that the adaptive choice of $\eta_{\mathbf{y}}^t$ better exploits the nonconvex-strongly-concave structure in the nonsmooth setting; indeed, $\eta_{\mathbf{y}}^t = 1/(\mu t)$ for $t \in \{1, 2, \dots, B\}$ which corresponds to the first epoch. Then, for each subsequent epoch, we set the same sequence of stepsizes. Such a finding is consistent with that for (stochastic) subgradient-type methods for minimizing a strongly convex function (Nemirovski and Yudin, 1983; Rakhlin et al., 2012; Shamir and Zhang, 2013). We remark that the complexity bounds for TTGDA and TTSGDA are the same in terms of $1/\epsilon$, which is also consistent with the classical results for nonsmooth optimization (Nemirovski and Yudin, 1983; Davis and Drusvyatskiy, 2019).*

Theorem 26. *Under Assumption 22, the following statements hold true,*

1. *If we set*

$$\eta_{\mathbf{x}}^t = \min \left\{ \frac{\epsilon^2}{48\rho L^2}, \frac{\epsilon^6}{65536\rho^3 L^4 D^2} \right\}, \quad \eta_{\mathbf{y}}^t = \frac{\epsilon^2}{16\rho L^2},$$

in the TTGDA scheme (see Algorithm 1), the required number of gradient evaluations to return an ϵ -stationary point is

$$O \left(\frac{\rho L^2 \Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho^2 L^2 D^2}{\epsilon^4} \right\} \right).$$

2. *If we set $M = 1$, G as per Eq. (1.3) and*

$$\eta_{\mathbf{x}}^t = \min \left\{ \frac{\epsilon^2}{48\rho(L^2 + \sigma^2)}, \frac{\epsilon^6}{131072\rho^3(L^2 + \sigma^2)^2 D^2} \right\}, \quad \eta_{\mathbf{y}}^t = \frac{\epsilon^2}{32\rho(L^2 + \sigma^2)},$$

in the TTSGDA scheme (see Algorithm 2), the required number of stochastic gradient evaluations to return an ϵ -stationary point is

$$O \left(\frac{\rho(L^2 + \sigma^2) \Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho^2(L^2 + \sigma^2) D^2}{\epsilon^4} \right\} \right).$$

Remark 27. *Theorems 24 and 26 provide the gradient complexity results for TTGDA and TTSGDA when applied to solve the nonsmooth and nonconvex-(strongly)-concave minimax optimization problems. The orders of $O(\epsilon^{-4})$ and $\Theta(\epsilon^{-4})$ reflect the effect of nonsmoothness when compared to that presented in Theorem 17 and 19. Furthermore, we guarantee that an ϵ -stationary point is visited within a certain number of iterations rather than guaranteeing last-iterate convergence. Finally, Theorem 24 is valid for the case of $M = 1$ and such a result is different from the theoretical results in Theorem 17.*

6. Applications and Empirical Results

We discuss our evaluation of TTGDA and TTSGDA in two domains—robust regression and Wasserstein generative adversarial networks (WGANs). All of algorithms were implemented on a MacBook Pro with an Intel Core i9 2.4GHz and 16GB memory.

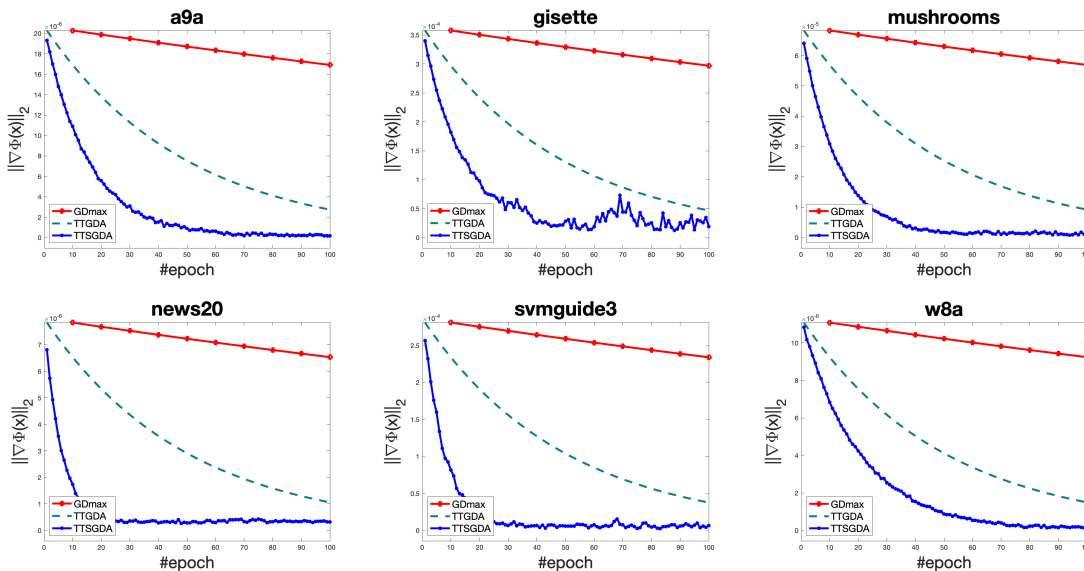


Figure 1: Performance of all the algorithms with 6 LIBSVM datasets. The numerical results are presented in terms of epoch count where the evaluation metric is the gradient norm of the function $\Phi(\cdot) = \max_{\mathbf{y} \in \mathcal{Y}} f(\cdot, \mathbf{y})$.

6.1 Robust logistic regression

We will consider the problem of robust logistic regression with nonconvex penalty functions. Given a set of samples $\{(\mathbf{a}_i, b_i)\}_{i=1}^N$ where $\mathbf{a}_i \in \mathbb{R}^d$ is the feature of the i^{th} sample and $b_i \in \{1, -1\}$ is the label, the problem can be written in the form of Eq. (1.1) with

$$f(\mathbf{x}, \mathbf{y}) = \frac{1}{N} \left(\sum_{i=1}^N y_i \log(1 + \exp(-b_i \mathbf{a}_i^\top \mathbf{x})) \right) - \frac{1}{2} \lambda_1 \|N\mathbf{y} - \mathbf{1}_n\|^2 + \lambda_2 \left(\sum_{i=1}^d \frac{\alpha x_i^2}{1 + \alpha x_i^2} \right),$$

and $\mathcal{Y} = \{\mathbf{y} \in \mathbb{R}_+^n : \mathbf{y}^\top \mathbf{1}_n = 1\}$. Here, we compare TTGDA and TTSGDA with GDmax (Jin et al., 2020) on 6 LIBSVM datasets³ and follow the setup of Kohler and Lucchi (2017) to set $\lambda_1 = \frac{1}{n^2}$, $\lambda_2 = 10^{-2}$ and $\alpha = 10$ for our experiments.

The algorithmic parameters are tuned as follows: we consider different pairs of $(\eta_{\mathbf{x}}, \eta_{\mathbf{y}})$ where $\eta_{\mathbf{y}} \in \{10^{-1}, 1\}$ and the ratio $\frac{\eta_{\mathbf{y}}}{\eta_{\mathbf{x}}} \in \{10, 10^2, 10^3\}$, and different sizes $M \in \{10, 100, 200\}$ for TTSGDA. Figure 1 presents the performance of each algorithm with its corresponding fine-tuning parameters, demonstrating that TTGDA and TTSGDA consistently converge faster than GDmax on 6 LIBSVM datasets that we consider in this paper.

6.2 WGANs with linear generators

We will consider the same setting as Loizou et al. (2020) which uses the WGAN (Arjovsky et al., 2017) to approximate a one-dimensional Gaussian distribution.

3. <https://www.csie.ntu.edu.tw/~cjlin/libsvm/>




Figure 2: Performance of all the algorithms for training WGANs with linear generators. The numerical results are presented in terms of iteration count where the evaluation metric is the gradient norm of the function $f(\cdot, \cdot)$.

The problem setup is as follows: We have the real data a^{real} generated from a normal distribution with $\hat{\mu} = 0$ and $\hat{\sigma} = 0.1$, and the latent variable z drawn from a standard normal distribution. We define the generator as $G_{\mathbf{x}}(z) = x_1 + x_2 z$ and the discriminator as $D_{\mathbf{y}}(a) = y_1 a + y_2 a^2$, where a is either real data or fake data generated from the generator. The problem can be written in the form of Eq. (1.1) with

$$f(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{(a^{\text{real}}, z)} \left[D_{\mathbf{y}}(a^{\text{real}}) - D_{\mathbf{y}}(G_{\mathbf{x}}(z)) - \lambda \|\mathbf{y}\|^2 \right],$$

where $\lambda = 10^{-3}$ is chosen to make the function $f(\mathbf{x}, \mathbf{y})$ concave in \mathbf{y} .

We fix the batch size $M = 100$ in TTSGDA and tune the parameters $(\eta_{\mathbf{x}}, \eta_{\mathbf{y}})$ as we have done in Section 6.1. Figure 2 shows that TTSGDA outperforms other adaptive algorithms, including Adam (Kingma and Ba, 2015) and RMSprop (Tieleman and Hinton, 2012). Here, the possible reason is that the linear generators result in a relatively simple structure such that the two-timescale update with tuning parameters is sufficient.

6.3 WGANs with nonlinear generators

Following the setup of Lei et al. (2020) and Yang et al. (2022b), we consider the WGAN with a ReLU neural network as generator. For ease of comparison, we retain all the problem settings as in the previous paragraph and view \mathbf{x} as the parameters of a small neural network (one hidden layer with five neurons and ReLU activations).

Figure 3 shows that TTSGDA is worse with Adam and RMSprop in this complex setting. This makes sense since the nonlinear generators result in a richer structure and the adaptive algorithms are highly optimized for handling such structure. The vanilla implementation of TTSGDA does not achieve a satisfactory level of performance and can be enhanced by exploiting the advanced tools of regularization and adaptiveness (Yang et al., 2022a,b; Li et al., 2023). For more details on the implementation heuristics and empirical performance for TTSGDA in training GANs and its variants, we refer to Heusel et al. (2017).




Figure 3: Performance of all the algorithms for training WGANs with nonlinear generators. The numerical results are presented in terms of iteration count where the evaluation metric is the gradient norm of the function $f(\cdot, \cdot)$.

7. Conclusion

We proposed two-timescale GDA and SGDA algorithms (TTGDA and TTSGDA for short) for solving a class of nonconvex minimax optimization problems. We established complexity bounds for TTGDA and TTSGDA in terms of the number of (stochastic) gradient evaluations, demonstrating the efficiency of these two algorithms in theory. Our theoretical analysis is of interest to both optimization and machine learning communities as it not only provides a new combination of a two-timescale update rule and GDA/SGDA but supplies a novel proof technique that applies to concave optimization problems with slowly changing objective functions. Future directions of research include the investigation of lower bounds for gradient-based algorithms in the nonconvex-concave settings and the extension of upper bounds for TTGDA and TTSGDA in the structured nonconvex-nonconcave settings.

References

- S. S. Abadeh, P. M. M. Esfahani, and D. Kuhn. Distributionally robust logistic regression. In *NeurIPS*, pages 1576–1584, 2015.
- L. Adolphs, H. Daneshmand, A. Lucchi, and T. Hofmann. Local saddle point optimization: A curvature exploitation approach. In *AISTATS*, pages 486–495. PMLR, 2019.
- A. Alacaoglu and Y. Malitsky. Stochastic variance reduction for variational inequality methods. In *COLT*, pages 778–816. PMLR, 2022.
- M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, pages 214–223, 2017.
- A. Auslender and M. Teboulle. Projected subgradient methods with non-euclidean distances for non-differentiable convex minimization and variational inequalities. *Mathematical Programming*, 120:27–48, 2009.

- W. Azizian, I. Mitliagkas, S. Lacoste-Julien, and G. Gidel. A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *AISTATS*, pages 2863–2873. PMLR, 2020.
- J. P. Bailey and G. Piliouras. Multiplicative weights update in zero-sum games. In *EC*, pages 321–338, 2018.
- D. Balduzzi, S. Racaniere, J. Martens, J. Foerster, K. Tuyls, and T. Graepel. The mechanics of N-player differentiable games. In *ICML*, pages 354–363. PMLR, 2018.
- B. Barazandeh and M. Razaviyayn. Solving non-convex non-differentiable min-max games using proximal gradient method. In *ICASSP*, pages 3162–3166. IEEE, 2020.
- T. Basar and G. J. Olsder. *Dynamic Noncooperative Game Theory*, volume 23. SIAM, 1999.
- D. P. Bertsekas. *Control of uncertain systems with a set-membership description of the uncertainty*. PhD thesis, Massachusetts Institute of Technology, 1971.
- R. I. Boç and A. Böhm. Alternating proximal-gradient steps for (stochastic) nonconvex-concave minimax problems. *SIAM Journal on Optimization*, 33(3):1884–1913, 2023.
- Y. Cai, A. Oikonomou, and W. Zheng. Finite-time last-iterate convergence for learning in multi-player games. In *NeurIPS*, pages 33904–33919, 2022.
- Y. Cai, A. Oikonomou, and W. Zheng. Accelerated algorithms for constrained nonconvex-nonconcave min-max optimization and comonotone inclusion. In *ICML*, 2024. URL <https://openreview.net/forum?id=EK7fuAMNoI>.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- G. H. G. Chen and R. T. Rockafellar. Convergence rates in forward–backward splitting. *SIAM Journal on Optimization*, 7(2):421–444, 1997.
- A. Cherukuri, B. Ghahsifard, and J. Cortes. Saddle-point dynamics: conditions for asymptotic stability of saddle points. *SIAM Journal on Control and Optimization*, 55(1):486–511, 2017.
- F. H. Clarke. *Optimization and Nonsmooth Analysis*. SIAM, 1990.
- J. M. Danskin. *The Theory of Max-min and Its Applications to Weapons Allocation Problems*. Springer, 1967.
- C. Daskalakis and I. Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *NeurIPS*, pages 9236–9246, 2018.
- C. Daskalakis and I. Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *ITCS*, pages 27:1–27:18, 2019.

- C. Daskalakis, A. Ilyas, V. Syrgkanis, and H. Zeng. Training GANs with optimism. In *ICLR*, 2018. URL <https://openreview.net/forum?id=SJJySbbAZ>.
- C. Daskalakis, S. Skoulakis, and M. Zampetakis. The complexity of constrained min-max optimization. In *STOC*, pages 1466–1478, 2021.
- C. Daskalakis, N. Golowich, S. Skoulakis, and M. Zampetakis. Stay-on-the-ridge: Guaranteed convergence to local minimax equilibrium in nonconvex-nonconcave games. In *COLT*, pages 5146–5198. PMLR, 2023.
- D. Davis and D. Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019.
- J. Diakonikolas, C. Daskalakis, and M. I. Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *AISTATS*, pages 2746–2754. PMLR, 2021.
- T. Fiez, C. Jin, P. Netrapalli, and L. J. Ratliff. Minimax optimization with smooth algorithmic adversaries. In *ICLR*, 2022. URL <https://openreview.net/forum?id=UdxJ2fJx7N0>.
- E. G. Golshtein. Generalized gradient method for finding saddle points. *Matekon*, 10(3):36–52, 1974.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014.
- E. Gorbunov, N. Loizou, and G. Gidel. Extragradients method: $O(1/k)$ last-iterate convergence for monotone variational inequalities and connections with cocoercivity. In *AISTATS*, pages 366–402. PMLR, 2022.
- B. Grimmer, H. Lu, P. Worah, and V. Mirrokni. The landscape of the proximal point method for nonconvex-nonconcave minimax optimization. *Mathematical Programming*, 201(1):373–407, 2023.
- P. Grnarova, K. Y. Levy, A. Lucchi, T. Hofmann, and A. Krause. An online learning approach to generative adversarial networks. In *ICLR*, 2018. URL <https://openreview.net/forum?id=H1Yp-j1Cb>.
- Y. Han, G. Xie, and Z. Zhang. Lower complexity bounds of finite-sum optimization problems: The results and construction. *Journal of Machine Learning Research*, 25(2):1–86, 2024.
- M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *NeurIPS*, pages 6626–6637, 2017.
- F. Huang, X. Wu, and H. Huang. Efficient mirror descent ascent methods for nonsmooth minimax problems. In *NeurIPS*, pages 10431–10443, 2021.

- F. Huang, S. Gao, J. Pei, and H. Huang. Accelerated zeroth-order and first-order momentum methods from mini to minimax optimization. *Journal of Machine Learning Research*, 23(36):1–70, 2022.
- J. Jiang and X. Chen. Optimality conditions for nonsmooth nonconvex-nonconcave min-max problems and generative adversarial networks. *SIAM Journal on Mathematics of Data Science*, 5(3):693–722, 2023.
- C. Jin, P. Netrapalli, and M. Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *ICML*, pages 4880–4889. PMLR, 2020.
- A. Juditsky, A. Nemirovski, and C. Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- J. M. Kohler and A. Lucchi. Sub-sampled cubic regularization for non-convex optimization. In *ICML*, pages 1895–1904, 2017.
- W. Kong and R. D. C. Monteiro. An accelerated inexact proximal point method for solving nonconvex-concave min-max problems. *SIAM Journal on Optimization*, 31(4):2558–2585, 2021.
- G. M. Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- T. Kose. Solutions of saddle value problems by differential equations. *Econometrica, Journal of the Econometric Society*, pages 59–70, 1956.
- S. Lee and D. Kim. Semi-anchored multi-step gradient descent ascent method for structured nonconvex-nonconcave composite minimax problems. *ArXiv Preprint: 2105.15042*, 2021.
- Q. Lei, J. D. Lee, A. G. Dimakis, and C. Daskalakis. SGD learns one-layer networks in WGANs. In *ICML*, pages 5799–5808, 2020.
- H. Li, Y. Tian, J. Zhang, and A. Jadbabaie. Complexity lower bounds for nonconvex-strongly-concave min-max optimization. In *NeurIPS*, volume 34, 2021.
- H. Li, F. Farnia, S. Das, and A. Jadbabaie. On convergence of gradient descent ascent: A tight local analysis. In *ICML*, pages 12717–12740. PMLR, 2022.
- X. Li, J. Yang, and N. He. TiAda: A time-scale adaptive algorithm for nonconvex minimax optimization. In *ICLR*, 2023. URL <https://openreview.net/forum?id=zClyiz5V6sL>.
- T. Liang and James Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *AISTATS*, pages 907–915. PMLR, 2019.
- T. Lin, C. Jin, and M. I. Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *ICML*, pages 6083–6093. PMLR, 2020a.

- T. Lin, C. Jin, and M. I. Jordan. Near-optimal algorithms for minimax optimization. In *COLT*, pages 2738–2779. PMLR, 2020b.
- M. Liu, H. Rafique, Q. Lin, and T. Yang. First-order convergence theory for weakly-convex-weakly-concave min-max problems. *Journal of Machine Learning Research*, 22(169):1–34, 2021.
- N. Loizou, H. Berard, A. Jolicoeur-Martineau, P. Vincent, S. Lacoste-Julien, and I. Mitliagkas. Stochastic Hamiltonian gradient methods for smooth games. In *ICML*, pages 6370–6381, 2020.
- S. Lu, I. Tsaknakis, M. Hong, and Y. Chen. Hybrid block successive approximation for one-sided non-convex min-max problems: Algorithms and applications. *IEEE Transactions on Signal Processing*, 68:3676–3691, 2020.
- L. Luo, H. Ye, Z. Huang, and T. Zhang. Stochastic recursive gradient descent ascent for stochastic nonconvex-strongly-concave minimax problems. In *NeurIPS*, pages 20566–20577, 2020.
- A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. URL <https://openreview.net/forum?id=rJzIBfZAb>.
- G. Mateos, J. A. Bazerque, and G. B. Giannakis. Distributed sparse linear regression. *IEEE Transactions on Signal Processing*, 58(10):5262–5276, 2010.
- P. Mertikopoulos, C. Papadimitriou, and G. Piliouras. Cycles in adversarial regularized learning. In *SODA*, pages 2703–2717. SIAM, 2018.
- P. Mertikopoulos, B. Lecouat, H. Zenati, C-S Foo, V. Chandrasekhar, and G. Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra(-gradient) mile. In *ICLR*, 2019. URL <https://openreview.net/forum?id=Bkg8jjC9KQ>.
- A. Mokhtari, A. Ozdaglar, and S. Pattathil. Convergence rate of $o(1/k)$ for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 30(4):3230–3251, 2020a.
- A. Mokhtari, A. Ozdaglar, and S. Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *AISTATS*, pages 1497–1507. PMLR, 2020b.
- K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39:117–129, 1987.
- A. Nedić and A. Ozdaglar. Subgradient methods for saddle-point problems. *Journal of Optimization Theory and Applications*, 142(1):205–228, 2009.
- A. Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

- A. Nemirovski and D. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley-Interscience, 1983.
- Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- Y. Nesterov. *Lectures on Convex Optimization*, volume 137. Springer, 2018.
- J. V. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.
- N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- M. Nouiehed, M. Sanjabi, T. Huang, J. D. Lee, and M. Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. In *NeurIPS*, pages 14905–14916, 2019.
- D. M. Ostrovskii, B. Barazandeh, and M. Razaviyayn. Nonconvex-nonconcave min-max optimization with a small maximization domain. *ArXiv Preprint: 2110.03950*, 2021a.
- D. M. Ostrovskii, A. Lowy, and M. Razaviyayn. Efficient search of first-order nash equilibria in nonconvex-concave smooth min-max problems. *SIAM Journal on Optimization*, 31(4): 2508–2538, 2021b.
- H. Rafique, M. Liu, Q. Lin, and T. Yang. Weakly-convex-concave min-max optimization: provable algorithms and applications in machine learning. *Optimization Methods and Software*, pages 1–35, 2021.
- A. Rakhlin, O. Shamir, and K. Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *ICML*, pages 1571–1578. Omnipress, 2012.
- J. Robinson. An iterative method of solving a game. *Annals of Mathematics*, pages 296–301, 1951.
- M. Sanjabi, J. Ba, M. Razaviyayn, and J. D. Lee. On the convergence and robustness of training GANs with regularized optimal transport. In *NeurIPS*, pages 7091–7101, 2018.
- O. Shamir and T. Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *ICML*, pages 71–79. PMLR, 2013.
- J. Shamma. *Cooperative Control of Distributed Multi-agent Systems*. John Wiley & Sons, 2008.
- A. Sinha, H. Namkoong, and J. Duchi. Certifiable distributional robustness with principled adversarial training. In *ICLR*, 2018. URL <https://openreview.net/forum?id=Hk6kPgZA->.
- M. Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.

- K. K. Thekumparampil, P. Jain, P. Netrapalli, and S. Oh. Efficient algorithms for smooth minimax optimization. In *NeurIPS*, pages 12659–12670, 2019.
- T. Tieleman and G. Hinton. Lecture 6.5-RMSprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4(2):26, 2012.
- H. Uzawa. Iterative methods for concave programming. *Studies in Linear and Nonlinear Programming*, 6:154–165, 1958.
- J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior (Commemorative Edition)*. Princeton University Press, 2007.
- H. Xu, C. Caramanis, and S. Mannor. Robustness and regularization of support vector machines. *Journal of Machine Learning Research*, 10(Jul):1485–1510, 2009.
- Z. Xu, H. Zhang, Y. Xu, and G. Lan. A unified single-loop alternating gradient projection algorithm for nonconvex-concave and convex-nonconcave minimax problems. *Mathematical Programming*, 201(1):635–706, 2023.
- Y. Yan, Y. Xu, Q. Lin, W. Liu, and T. Yang. Optimal epoch stochastic gradient descent ascent methods for min-max optimization. In *NeurIPS*, pages 5789–5800, 2020.
- J. Yang, N. Kiyavash, and N. He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. In *NeurIPS*, pages 1153–1165, 2020a.
- J. Yang, S. Zhang, N. Kiyavash, and N. He. A catalyst framework for minimax optimization. In *NeurIPS*, pages 5667–5678, 2020b.
- J. Yang, X. Li, and N. He. Nest your adaptive algorithm for parameter-agnostic nonconvex minimax optimization. In *NeurIPS*, pages 11202–11216, 2022a.
- J. Yang, A. Orvieto, A. Lucchi, and N. He. Faster single-loop algorithms for minimax optimization without strong concavity. In *AISTATS*, pages 5485–5517. PMLR, 2022b.
- J. Zhang, H. Lin, S. Jegelka, S. Sra, and A. Jadbabaie. Complexity of finding stationary points of nonsmooth nonconvex functions. In *ICML*, pages 11173–11182, 2020a.
- J. Zhang, Peijun Xiao, Ruoyu Sun, and Zhiqian Luo. A single-loop smoothed gradient descent-ascent algorithm for nonconvex-concave min-max problems. In *NeurIPS*, pages 7377–7389, 2020b.
- S. Zhang, J. Yang, C. Guzmán, N. Kiyavash, and N. He. The complexity of nonconvex-strongly-concave minimax optimization. In *UAI*, pages 482–492. PMLR, 2021.
- R. Zhao. A primal-dual smoothing framework for max-structured non-convex optimization. *Mathematics of Operations Research*, To appear, 2023.
- T. Zheng, L. Zhu, A. M-C. So, J. Blanchet, and J. Li. Universal gradient descent ascent method for nonconvex-nonconcave minimax optimization. In *NeurIPS*, pages 54075–54110, 2023.

Appendix A. Proofs for Technical Lemmas

We provide the detailed proofs for Lemmas 6, 7, 9, 16, 23, Propositions 11, 12, and auxiliary results on stochastic gradients.

A.1 Proof of Lemma 6

We first prove that the Moreau envelope $\Phi_{1/2\ell}(\cdot)$ is well defined. Fixing \mathbf{x} and using that f is ℓ -smooth, the function $f(\cdot, \mathbf{y}) + 0.5\ell\|\cdot - \mathbf{x}\|^2$ is convex for any $\mathbf{y} \in \mathcal{Y}$. By definition, we have $\Phi(\cdot) + 0.5\ell\|\cdot - \mathbf{x}\|^2 = \max_{\mathbf{y} \in \mathcal{Y}} \{f(\cdot, \mathbf{y}) + 0.5\ell\|\cdot - \mathbf{x}\|^2\}$. Since the supremum of many convex functions is convex, the function $\Phi(\cdot) + 0.5\ell\|\cdot - \mathbf{x}\|^2$ is convex. As such, the function $\Phi(\cdot) + \ell\|\cdot - \mathbf{x}\|^2$ is 0.5ℓ -strongly convex and the Moreau envelope $\Phi_{1/2\ell}(\cdot)$ is well defined.

As a byproduct of the above argument, $\text{prox}_{\Phi/2\ell}(\cdot) = \text{argmin}_{\mathbf{w}} \Phi(\mathbf{w}) + \ell\|\mathbf{w} - \cdot\|^2$ is well defined. By definition, we have

$$\Phi(\text{prox}_{\Phi/2\ell}(\mathbf{x})) \leq \Phi(\text{prox}_{\Phi/2\ell}(\mathbf{x})) + \ell\|\text{prox}_{\Phi/2\ell}(\mathbf{x}) - \mathbf{x}\|^2 \leq \Phi(\mathbf{x}), \quad \text{for all } \mathbf{x} \in \mathbb{R}^m.$$

Finally, Davis and Drusvyatskiy (2019, Lemma 2.2) implies that the Moreau envelope $\Phi_{1/2\ell}$ is ℓ -smooth with $\nabla\Phi_{1/2\ell}(\mathbf{x}) = 2\ell(\mathbf{x} - \text{prox}_{\Phi/2\ell}(\mathbf{x}))$.

A.2 Proof of Lemma 7

Fixing \mathbf{x} and using that $f(\mathbf{x}, \mathbf{y})$ is ρ -weakly convex in \mathbf{x} , the function $f(\cdot, \mathbf{y}) + 0.5\rho\|\cdot - \mathbf{x}\|^2$ is convex for any $\mathbf{y} \in \mathcal{Y}$. By applying the same argument used for proving Lemma 6, we have that $\Phi_{1/2\rho}$ and $\text{prox}_{\Phi/2\rho}(\cdot)$ are well defined, $\Phi(\text{prox}_{\Phi/2\rho}(\mathbf{x})) \leq \Phi(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^m$, and $\Phi_{1/2\rho}$ is ρ -smooth with $\nabla\Phi_{1/2\rho}(\mathbf{x}) = 2\rho(\mathbf{x} - \text{prox}_{\Phi/2\rho}(\mathbf{x}))$.

A.3 Proof of Lemma 9

It suffices to show that $\hat{\mathbf{x}} = \text{prox}_{\Phi/2\ell}(\mathbf{x})$ satisfies $\min_{\xi \in \partial\Phi(\hat{\mathbf{x}})} \|\xi\| \leq \epsilon$ and $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq \epsilon/2\rho$. Indeed, Davis and Drusvyatskiy (2019, Lemma 2.2) guarantees that $\nabla\Phi_{1/2\rho}(\mathbf{x}) = 2\rho(\mathbf{x} - \hat{\mathbf{x}})$ where $\hat{\mathbf{x}} = \text{prox}_{\Phi/2\rho}(\mathbf{x})$ and $\|\hat{\mathbf{x}} - \mathbf{x}\| = \|\nabla\Phi_{1/2\rho}(\mathbf{x})\|/2\rho$. By the definition of $\text{prox}_{\Phi/2\rho}(\cdot)$, we have $2\rho(\mathbf{x} - \hat{\mathbf{x}}) \in \partial\Phi(\hat{\mathbf{x}})$. Putting these pieces together yields $\min_{\xi \in \partial\Phi(\hat{\mathbf{x}})} \|\xi\| \leq \|\nabla\Phi_{1/2\rho}(\mathbf{x})\|$. Since $\|\nabla\Phi_{1/2\rho}(\mathbf{x})\| \leq \epsilon$, we achieve the desired results.

A.4 Proof of Proposition 11

Given that $\hat{\mathbf{x}}$ satisfies $\|\nabla\Phi(\hat{\mathbf{x}})\| \leq \epsilon$, the function $f(\hat{\mathbf{x}}, \mathbf{y})$ is ℓ -smooth and μ -strongly concave in \mathbf{y} , and $\mathbf{y}^*(\hat{\mathbf{x}}) = \text{argmax}_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y})$ is uniquely defined. By applying the standard gradient ascent algorithm for the optimization problem of $\max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \cdot)$, we obtain a point $\mathbf{y}' \in \mathcal{Y}$ satisfying that, for $\mathbf{y}^+ = \mathcal{P}_{\mathcal{Y}}(\mathbf{y}' + \frac{1}{\ell}\nabla_{\mathbf{y}}f(\hat{\mathbf{x}}, \mathbf{y}'))$, the following inequalities hold true,

$$\|\mathbf{y}^+ - \mathbf{y}'\| \leq \frac{\epsilon}{\ell}, \quad \|\mathbf{y}^+ - \mathbf{y}^*(\hat{\mathbf{x}})\| \leq \frac{\epsilon}{\ell}. \quad (\text{A.1})$$

Since $\|\nabla\Phi(\hat{\mathbf{x}})\| \leq \epsilon$ and $\nabla\Phi(\hat{\mathbf{x}}) = \nabla_{\mathbf{x}}f(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}}))$, we have

$$\|\nabla_{\mathbf{x}}f(\hat{\mathbf{x}}, \mathbf{y}^+)\| \leq \|\nabla_{\mathbf{x}}f(\hat{\mathbf{x}}, \mathbf{y}^+) - \nabla\Phi(\hat{\mathbf{x}})\| + \|\nabla\Phi(\hat{\mathbf{x}})\| = \|\nabla_{\mathbf{x}}f(\hat{\mathbf{x}}, \mathbf{y}^+) - \nabla_{\mathbf{x}}f(\hat{\mathbf{x}}, \mathbf{y}^*(\hat{\mathbf{x}}))\| + \epsilon.$$

Since f is ℓ -smooth, we have $\|\nabla_{\mathbf{x}}f(\hat{\mathbf{x}}, \mathbf{y}^+)\| \leq \ell\|\mathbf{y}^+ - \mathbf{y}^*(\hat{\mathbf{x}})\| + \epsilon \leq 2\epsilon$ and the convergence guarantee of standard gradient descent method implies that the required number of gradient

evaluations to output a point $\mathbf{y}' \in \mathcal{Y}$ satisfying Eq. (A.1) is $\mathcal{O}(\kappa \log(1/\epsilon))$ (Nesterov, 2018). The similar argument holds for applying stochastic gradient descent with proper stepsize and the required number of stochastic gradient evaluations is $\mathcal{O}(\epsilon^{-2})$.

Conversely, given that $(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \in \mathbb{R}^m \times \mathcal{Y}$ is a $\frac{\epsilon}{2\kappa}$ -stationary point satisfying

$$\|\nabla_{\mathbf{x}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+)\| \leq \frac{\epsilon}{2\kappa}, \quad \|\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}\| \leq \frac{\epsilon}{2\kappa\ell},$$

where $\hat{\mathbf{y}}^+ = \mathcal{P}_{\mathcal{Y}}(\hat{\mathbf{y}} + (1/\ell)\nabla_{\mathbf{y}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}))$, we have

$$\|\nabla\Phi(\hat{\mathbf{x}})\| \leq \|\nabla\Phi(\hat{\mathbf{x}}) - \nabla_{\mathbf{x}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+)\| + \|\nabla_{\mathbf{x}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+)\| \leq \ell\|\hat{\mathbf{y}}^+ - \mathbf{y}^*(\hat{\mathbf{x}})\| + \frac{\epsilon}{2\kappa}.$$

By the definition of $\hat{\mathbf{y}}^+$ and $\mathbf{y}^*(\hat{\mathbf{x}})$, we have

$$(\mathbf{y}^*(\hat{\mathbf{x}}) - \hat{\mathbf{y}}^+)^\top (\hat{\mathbf{y}}^+ - \hat{\mathbf{y}} - \frac{1}{\ell}\nabla_{\mathbf{y}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}})) \geq 0,$$

which implies

$$\ell(\mathbf{y}^*(\hat{\mathbf{x}}) - \hat{\mathbf{y}}^+)^\top (\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}) \geq (\mathbf{y}^*(\hat{\mathbf{x}}) - \hat{\mathbf{y}}^+)^\top \nabla_{\mathbf{y}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}).$$

Since the function $f(\hat{\mathbf{x}}, \cdot)$ is ℓ -smooth and μ -strongly concave, we have

$$\ell(\mathbf{y}^*(\hat{\mathbf{x}}) - \hat{\mathbf{y}}^+)^\top (\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}) \geq \frac{\mu}{2}\|\mathbf{y}^*(\hat{\mathbf{x}}) - \hat{\mathbf{y}}\|^2 - \frac{\ell}{2}\|\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}\|^2.$$

In addition, we have

$$(\mathbf{y}^*(\hat{\mathbf{x}}) - \hat{\mathbf{y}}^+)^\top (\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}) \leq \|\mathbf{y}^*(\hat{\mathbf{x}}) - \hat{\mathbf{y}}\| \|\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}\| - \|\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}\|^2.$$

Putting these pieces together yields $\|\hat{\mathbf{y}}^+ - \mathbf{y}^*(\hat{\mathbf{x}})\| \leq 2\kappa\|\hat{\mathbf{y}} - \hat{\mathbf{y}}^+\| \leq \frac{\epsilon}{\ell}$. This together with $\|\hat{\mathbf{y}}^+ - \mathbf{y}^*(\hat{\mathbf{x}})\| \leq \|\hat{\mathbf{y}} - \mathbf{y}^*(\hat{\mathbf{x}})\|$ yields $\|\nabla\Phi(\hat{\mathbf{x}})\| \leq \epsilon + \frac{\epsilon}{2\kappa} \leq 2\epsilon$ which implies the desired result.

A.5 Proof of Proposition 12

Given that $\hat{\mathbf{x}}$ satisfies $\|\nabla\Phi_{1/2\ell}(\hat{\mathbf{x}})\| \leq \epsilon$, the function $f(\mathbf{x}, \mathbf{y}) + \ell\|\mathbf{x} - \hat{\mathbf{x}}\|^2$ is 0.5ℓ -strongly convex in \mathbf{x} and concave in \mathbf{y} , and $\mathbf{x}^*(\hat{\mathbf{x}}) = \operatorname{argmin}_{\mathbf{w}} \Phi(\mathbf{w}) + \ell\|\mathbf{w} - \hat{\mathbf{x}}\|^2$ is uniquely defined. By applying the standard extragradient algorithm for solving the optimization problem of $\min_{\mathbf{x}} \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y}) + \ell\|\mathbf{x} - \hat{\mathbf{x}}\|^2$, we obtain a pair of points $(\mathbf{x}', \mathbf{y}') \in \mathbb{R}^m \times \mathcal{Y}$ satisfying that, for $\mathbf{y}^+ = \mathcal{P}_{\mathcal{Y}}(\mathbf{y}' + \frac{1}{\ell}\nabla_{\mathbf{y}} f(\mathbf{x}', \mathbf{y}'))$, the following inequalities hold true,

$$\|\nabla_{\mathbf{x}} f(\mathbf{x}', \mathbf{y}^+) + 2\ell(\mathbf{x}' - \hat{\mathbf{x}})\| \leq \epsilon, \quad \|\mathbf{x}' - \mathbf{x}^*(\hat{\mathbf{x}})\| \leq \frac{\epsilon}{2\ell}, \quad \|\mathbf{y}^+ - \mathbf{y}'\| \leq \frac{\epsilon}{\ell}. \quad (\text{A.2})$$

Since $\|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\| = \frac{1}{2\ell}\|\nabla\Phi_{1/2\ell}(\hat{\mathbf{x}})\| \leq \frac{\epsilon}{2\ell}$, we have

$$\|\nabla_{\mathbf{x}} f(\mathbf{x}', \mathbf{y}^+)\| \leq \|\nabla_{\mathbf{x}} f(\mathbf{x}', \mathbf{y}^+) + 2\ell(\mathbf{x}' - \hat{\mathbf{x}})\| + 2\ell\|\mathbf{x}' - \mathbf{x}^*(\hat{\mathbf{x}})\| + 2\ell\|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\| \leq 4\epsilon.$$

To facilitate the reader's understanding, we give a brief description of the extragradient algorithm. It was introduced by Korpelevich (1976) and has been recognized as one of the most useful methods for solving minimax optimization problems. In the convex-concave setting, its global convergence rate guarantees were established for both averaged iterates (Nemirovski, 2004; Mokhtari et al., 2020a,b) and last iterates (Gorbunov et al., 2022; Cai et al.,

2022). When applied to solve the optimization problem of $\min_{\mathbf{x}} \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y}) + \ell \|\mathbf{x} - \hat{\mathbf{x}}\|^2$, one iteration of the extragradient algorithm is summarized as follows,

$$\begin{aligned} \tilde{\mathbf{x}}_t &\leftarrow \mathbf{x}_{t-1} - \eta(\nabla_{\mathbf{x}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + 2\ell(\mathbf{x}_{t-1} - \hat{\mathbf{x}})), & \tilde{\mathbf{y}}_t &\leftarrow \mathbf{y}_{t-1} + \eta \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}), \\ \mathbf{x}_t &\leftarrow \mathbf{x}_{t-1} - \eta(\nabla_{\mathbf{x}} f(\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t) + 2\ell(\tilde{\mathbf{x}}_t - \hat{\mathbf{x}})), & \mathbf{y}_t &\leftarrow \mathbf{y}_{t-1} + \eta \nabla_{\mathbf{y}} f(\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t). \end{aligned}$$

The convergence guarantee of extragradient algorithm implies that the required number of gradient evaluations to output a pair of points $(\mathbf{x}', \mathbf{y}') \in \mathbb{R}^m \times \mathcal{Y}$ satisfying Eq. (A.2) is $O(\epsilon^{-2})$. The similar argument holds for applying stochastic mirror-prox algorithm and the required number of stochastic gradient evaluations is $O(\epsilon^{-4})$ (Juditsky et al., 2011).

Conversely, given that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a $\frac{\epsilon^2}{\max\{1, \ell D\}}$ -stationary point satisfying

$$\|\nabla_{\mathbf{x}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+)\| \leq \frac{\epsilon^2}{\max\{1, \ell D\}}, \quad \|\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}\| \leq \frac{\epsilon^2}{\ell \max\{1, \ell D\}}, \quad (\text{A.3})$$

where $\hat{\mathbf{y}}^+ = \mathcal{P}_{\mathcal{Y}}(\hat{\mathbf{y}} + (1/\ell)\nabla_{\mathbf{y}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}))$, we have

$$\|\nabla \Phi_{1/2\ell}(\hat{\mathbf{x}})\|^2 = 4\ell^2 \|\hat{\mathbf{x}} - \mathbf{x}^*(\hat{\mathbf{x}})\|^2.$$

Since $\Phi(\cdot) + \ell \|\cdot - \hat{\mathbf{x}}\|^2$ is 0.5ℓ -strongly-convex, we have

$$\begin{aligned} &\max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}^*(\hat{\mathbf{x}}), \mathbf{y}) - \ell \|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\|^2 \\ &= \Phi(\hat{\mathbf{x}}) - \Phi(\mathbf{x}^*(\hat{\mathbf{x}})) - \ell \|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\|^2 \geq \frac{\ell}{4} \|\hat{\mathbf{x}} - \mathbf{x}^*(\hat{\mathbf{x}})\|^2 = \frac{1}{16\ell} \|\nabla \Phi_{1/2\ell}(\hat{\mathbf{x}})\|^2. \end{aligned} \quad (\text{A.4})$$

Further, we have

$$\begin{aligned} &\max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}^*(\hat{\mathbf{x}}), \mathbf{y}) - \ell \|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\|^2 \\ &= \max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+) + f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+) - \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}^*(\hat{\mathbf{x}}), \mathbf{y}) - \ell \|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\|^2 \\ &\leq \max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+) + (f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+) - f(\mathbf{x}^*(\hat{\mathbf{x}}), \hat{\mathbf{y}}^+) - \ell \|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\|^2) \\ &\leq \max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+) + (\|\hat{\mathbf{x}} - \mathbf{x}^*(\hat{\mathbf{x}})\| \|\nabla_{\mathbf{x}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+)\| - \frac{\ell}{2} \|\hat{\mathbf{x}} - \mathbf{x}^*(\hat{\mathbf{x}})\|^2) \\ &\leq \max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+) + \frac{1}{2\ell} \|\nabla_{\mathbf{x}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+)\|^2. \end{aligned}$$

By the definition of $\hat{\mathbf{y}}^+$, we have $(\mathbf{y} - \hat{\mathbf{y}}^+)^\top (\hat{\mathbf{y}}^+ - \hat{\mathbf{y}} - (1/\ell)\nabla_{\mathbf{y}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}})) \geq 0$ for all $\mathbf{y} \in \mathcal{Y}$. Since the function $f(\hat{\mathbf{x}}, \cdot)$ is ℓ -smooth and the diameter of \mathcal{Y} is D , we have

$$f(\hat{\mathbf{x}}, \mathbf{y}) - f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+) \leq \ell D \|\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}\| \text{ for all } \mathbf{y} \in \mathcal{Y}.$$

Putting these pieces together yields

$$\max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}^*(\hat{\mathbf{x}}), \mathbf{y}) - \ell \|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\|^2 \leq \ell D \|\hat{\mathbf{y}}^+ - \hat{\mathbf{y}}\| + \frac{1}{2\ell} \|\nabla_{\mathbf{x}} f(\hat{\mathbf{x}}, \hat{\mathbf{y}}^+)\|^2.$$

Plugging Eq. (A.3) into the above inequality yields

$$\max_{\mathbf{y} \in \mathcal{Y}} f(\hat{\mathbf{x}}, \mathbf{y}) - \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}^*(\hat{\mathbf{x}}), \mathbf{y}) - \ell \|\mathbf{x}^*(\hat{\mathbf{x}}) - \hat{\mathbf{x}}\|^2 \leq \frac{\epsilon^2}{\ell} + \frac{\epsilon^4}{2\ell}. \quad (\text{A.5})$$

Combining Eq. (A.4) with Eq. (A.5) yields $\|\nabla \Phi_{1/2\ell}(\hat{\mathbf{x}})\| \leq 4\epsilon^2 + 4\epsilon$ and the desired result.

A.6 Proof of Lemma 16

We prove the results under Assumption 14. Indeed, we have that $\mathcal{Y}^*(\mathbf{x}) = \{\mathbf{y}^*(\mathbf{x})\}$ since the function $f(\mathbf{x}, \cdot)$ is strongly concave for each \mathbf{x} . Then, let $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^m$, the definition of $\mathbf{y}^*(\mathbf{x}_1)$ and the definition of $\mathbf{y}^*(\mathbf{x}_2)$ imply

$$(\mathbf{y} - \mathbf{y}^*(\mathbf{x}_1))^\top \nabla_{\mathbf{y}} f(\mathbf{x}_1, \mathbf{y}^*(\mathbf{x}_1)) \leq 0, \quad \text{for all } \mathbf{y} \in \mathcal{Y}, \quad (\text{A.6})$$

$$(\mathbf{y} - \mathbf{y}^*(\mathbf{x}_2))^\top \nabla_{\mathbf{y}} f(\mathbf{x}_2, \mathbf{y}^*(\mathbf{x}_2)) \leq 0, \quad \text{for all } \mathbf{y} \in \mathcal{Y}. \quad (\text{A.7})$$

Letting $\mathbf{y} = \mathbf{y}^*(\mathbf{x}_2)$ in Eq. (A.6) and $\mathbf{y} = \mathbf{y}^*(\mathbf{x}_1)$ in Eq. (A.7) and adding them yields

$$(\mathbf{y}^*(\mathbf{x}_2) - \mathbf{y}^*(\mathbf{x}_1))^\top (\nabla_{\mathbf{y}} f(\mathbf{x}_1, \mathbf{y}^*(\mathbf{x}_1)) - \nabla_{\mathbf{y}} f(\mathbf{x}_2, \mathbf{y}^*(\mathbf{x}_2))) \leq 0. \quad (\text{A.8})$$

Since the function $f(\mathbf{x}_1, \cdot)$ is μ -strongly concave, we have

$$(\mathbf{y}^*(\mathbf{x}_2) - \mathbf{y}^*(\mathbf{x}_1))^\top (\nabla_{\mathbf{y}} f(\mathbf{x}_1, \mathbf{y}^*(\mathbf{x}_2)) - \nabla_{\mathbf{y}} f(\mathbf{x}_1, \mathbf{y}^*(\mathbf{x}_1))) \leq -\mu \|\mathbf{y}^*(\mathbf{x}_2) - \mathbf{y}^*(\mathbf{x}_1)\|^2. \quad (\text{A.9})$$

Combining Eq. (A.8) and Eq. (A.9) with the fact that f is ℓ -smooth yields

$$\mu \|\mathbf{y}^*(\mathbf{x}_2) - \mathbf{y}^*(\mathbf{x}_1)\|^2 \leq \ell \|\mathbf{y}^*(\mathbf{x}_2) - \mathbf{y}^*(\mathbf{x}_1)\| \|\mathbf{x}_2 - \mathbf{x}_1\|.$$

which implies that the function $\mathbf{y}^*(\cdot)$ is κ -Lipschitz.

Since $\mathbf{y}^*(\cdot)$ is a well-defined function and \mathcal{Y} is a convex and bounded set, the celebrated Danskin theorem (see Danskin (1967) and Bertsekas (1971, Proposition A.22)) implies that $\nabla \Phi(\cdot) = \nabla_{\mathbf{x}} f(\cdot, \mathbf{y}^*(\cdot))$. Since $\mathbf{y}^*(\cdot)$ is κ -Lipschitz and $\kappa \geq 1$, we have

$$\|\nabla \Phi(\mathbf{x}) - \nabla \Phi(\mathbf{x}')\| = \|\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}^*(\mathbf{x})) - \nabla_{\mathbf{x}} f(\mathbf{x}', \mathbf{y}^*(\mathbf{x}'))\| \leq \ell (\|\mathbf{x} - \mathbf{x}'\| + \kappa \|\mathbf{x} - \mathbf{x}'\|) \leq 2\kappa\ell \|\mathbf{x} - \mathbf{x}'\|.$$

We then prove the results under Assumption 15. Indeed, Lemma 6 implies that the function $\Phi(\cdot)$ is ℓ -weakly convex. Since \mathcal{Y} is a convex and bounded set, there exist $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{Y}$ such that

$$f(\mathbf{x}_1, \mathbf{y}_1) = \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}_1, \mathbf{y}), \quad f(\mathbf{x}_2, \mathbf{y}_2) = \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}_2, \mathbf{y}), \quad \text{for any given } \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^m.$$

Thus, we have

$$\begin{aligned} \Phi(\mathbf{x}_1) - \Phi(\mathbf{x}_2) &= f(\mathbf{x}_1, \mathbf{y}_1) - f(\mathbf{x}_2, \mathbf{y}_2) \leq f(\mathbf{x}_1, \mathbf{y}_1) - f(\mathbf{x}_2, \mathbf{y}_1), \\ \Phi(\mathbf{x}_2) - \Phi(\mathbf{x}_1) &= f(\mathbf{x}_2, \mathbf{y}_2) - f(\mathbf{x}_1, \mathbf{y}_1) \leq f(\mathbf{x}_2, \mathbf{y}_2) - f(\mathbf{x}_1, \mathbf{y}_2). \end{aligned}$$

Since the function $f(\cdot, \mathbf{y})$ is L -Lipschitz for each $\mathbf{y} \in \mathcal{Y}$, we have

$$|\Phi(\mathbf{x}_1) - \Phi(\mathbf{x}_2)| = \max \{\Phi(\mathbf{x}_1) - \Phi(\mathbf{x}_2), \Phi(\mathbf{x}_2) - \Phi(\mathbf{x}_1)\} \leq L \|\mathbf{x}_1 - \mathbf{x}_2\|.$$

Further, we have $\partial \Phi(\mathbf{x}) = \partial \Psi(\mathbf{x}) - \ell \mathbf{x}$ for each \mathbf{x} where $\Psi(\mathbf{x}) = \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y}) + 0.5\ell \|\mathbf{x}\|^2$. Since $f(\cdot, \mathbf{y}) + 0.5\ell \|\cdot\|^2$ is convex for each $\mathbf{y} \in \mathcal{Y}$ where \mathcal{Y} is convex and bounded, we have that $\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) + \ell \mathbf{x} \in \partial \Psi(\mathbf{x})$ where $\mathbf{y} \in \mathcal{Y}^*(\mathbf{x})$. This yields the desired result.

A.7 Proof of Lemma 23

We prove the results under Assumption 22. Indeed, Lemma 7 implies that the function $\Phi(\cdot)$ is ρ -weakly convex. Given any $\mathbf{x}_1, \mathbf{x}_2$, we have that there exist $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{Y}$ such that

$$\begin{aligned}\Phi(\mathbf{x}_1) - \Phi(\mathbf{x}_2) &= f(\mathbf{x}_1, \mathbf{y}_1) - f(\mathbf{x}_2, \mathbf{y}_2) \leq f(\mathbf{x}_1, \mathbf{y}_1) - f(\mathbf{x}_2, \mathbf{y}_1), \\ \Phi(\mathbf{x}_2) - \Phi(\mathbf{x}_1) &= f(\mathbf{x}_2, \mathbf{y}_2) - f(\mathbf{x}_1, \mathbf{y}_1) \leq f(\mathbf{x}_2, \mathbf{y}_2) - f(\mathbf{x}_1, \mathbf{y}_2).\end{aligned}$$

Since f is L -Lipschitz, we have

$$|\Phi(\mathbf{x}_1) - \Phi(\mathbf{x}_2)| \leq L\|\mathbf{x}_1 - \mathbf{x}_2\|.$$

Applying the same argument used for proving Lemma 16, we have that $\partial_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}) \subseteq \partial\Phi(\mathbf{x})$ for each \mathbf{x} where $\mathbf{y} \in \mathcal{Y}^*(\mathbf{x})$.

We then prove the results under Assumption 21. Since the function $f(\mathbf{x}, \cdot)$ is μ -strongly concave for each \mathbf{x} , we have that $\mathcal{Y}^*(\mathbf{x}) = \{\mathbf{y}^*(\mathbf{x})\}$. Since Assumption 21 is stronger than Assumption 22, all other desired results follow from the previous analysis.

A.8 Auxiliary Results on Stochastic Gradient

We establish the properties of stochastic gradients sampled at each iteration.

Lemma 28. *If Eq. (1.3) holds true, the estimators $(\hat{\mathbf{g}}_{\mathbf{x}}^t, \hat{\mathbf{g}}_{\mathbf{y}}^t) = \text{SG}(G, \{\xi_i^t\}_{i=1}^M, \mathbf{x}_t, \mathbf{y}_t)$ satisfy*

$$\begin{aligned}\mathbb{E}[\hat{\mathbf{g}}_{\mathbf{x}}^t \mid \mathbf{x}_t, \mathbf{y}_t] &\in \partial_{\mathbf{x}}f(\mathbf{x}_t, \mathbf{y}_t), & \mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{x}}^t - \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{x}}^t \mid \mathbf{x}_t, \mathbf{y}_t]\|^2 \mid \mathbf{x}_t, \mathbf{y}_t] &\leq \frac{\sigma^2}{M}, \\ \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^t \mid \mathbf{x}_t, \mathbf{y}_t] &\in \partial_{\mathbf{y}}f(\mathbf{x}_t, \mathbf{y}_t), & \mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{y}}^t - \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^t \mid \mathbf{x}_t, \mathbf{y}_t]\|^2 \mid \mathbf{x}_t, \mathbf{y}_t] &\leq \frac{\sigma^2}{M}.\end{aligned}$$

Proof. Since $G = (G_{\mathbf{x}}, G_{\mathbf{y}})$ is unbiased and

$$\hat{\mathbf{g}}_{\mathbf{x}}^t = \frac{1}{M} \left(\sum_{i=1}^M G_{\mathbf{x}}(\mathbf{x}_t, \mathbf{y}_t, \xi_i^t) \right), \quad \hat{\mathbf{g}}_{\mathbf{y}}^t = \frac{1}{M} \left(\sum_{i=1}^M G_{\mathbf{y}}(\mathbf{x}_t, \mathbf{y}_t, \xi_i^t) \right),$$

we have

$$\mathbb{E}[\hat{\mathbf{g}}_{\mathbf{x}}^t \mid \mathbf{x}_t, \mathbf{y}_t] \in \partial_{\mathbf{x}}f(\mathbf{x}_t, \mathbf{y}_t), \quad \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^t \mid \mathbf{x}_t, \mathbf{y}_t] \in \partial_{\mathbf{y}}f(\mathbf{x}_t, \mathbf{y}_t).$$

Furthermore, we have

$$\mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{x}}^t - \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{x}}^t \mid \mathbf{x}_t, \mathbf{y}_t]\|^2 \mid \mathbf{x}_t, \mathbf{y}_t] \leq \frac{\sigma^2}{M}, \quad \mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{y}}^t - \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^t \mid \mathbf{x}_t, \mathbf{y}_t]\|^2 \mid \mathbf{x}_t, \mathbf{y}_t] \leq \frac{\sigma^2}{M}.$$

This completes the proof. \square

Appendix B. Smooth and Nonconvex-Strongly-Concave Setting

We provide the proof for Theorem 17. The trivial case is $\kappa = 1$ since it is as easy as smooth nonconvex optimization, which has been investigated extensively. For simplicity, we assume $\kappa > 1$ in the subsequent analysis. Throughout this subsection, we set

$$\eta_{\mathbf{x}}^t \equiv \eta_{\mathbf{x}} = \frac{1}{16(\kappa+1)^2\ell}, \quad \eta_{\mathbf{y}}^t \equiv \eta_{\mathbf{y}} = \frac{1}{\ell}, \quad \text{for both Algorithms 1 and 2.}$$

Lemma 29. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\Phi(\mathbf{x}_t) \leq \Phi(\mathbf{x}_{t-1}) - \left(\frac{\eta_{\mathbf{x}}}{2} - 2\eta_{\mathbf{x}}^2\kappa\ell\right)\|\nabla\Phi(\mathbf{x}_{t-1})\|^2 + \left(\frac{\eta_{\mathbf{x}}}{2} + 2\eta_{\mathbf{x}}^2\kappa\ell\right)\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\begin{aligned} \mathbb{E}[\Phi(\mathbf{x}_t)] &\leq \mathbb{E}[\Phi(\mathbf{x}_{t-1})] - \left(\frac{\eta_{\mathbf{x}}}{2} - 2\eta_{\mathbf{x}}^2\kappa\ell\right)\mathbb{E}[\|\nabla\Phi(\mathbf{x}_{t-1})\|^2] \\ &\quad + \left(\frac{\eta_{\mathbf{x}}}{2} + 2\eta_{\mathbf{x}}^2\kappa\ell\right)\mathbb{E}[\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2] + \frac{\eta_{\mathbf{x}}^2\kappa\ell\sigma^2}{M}. \end{aligned}$$

Proof. We first prove the inequality for Algorithm 1. Since $\Phi(\cdot)$ is $(2\kappa\ell)$ -smooth, we have

$$\Phi(\mathbf{x}_t) - \Phi(\mathbf{x}_{t-1}) - (\mathbf{x}_t - \mathbf{x}_{t-1})^\top \nabla\Phi(\mathbf{x}_{t-1}) \leq \kappa\ell\|\mathbf{x}_t - \mathbf{x}_{t-1}\|^2.$$

Since $\mathbf{x}_t - \mathbf{x}_{t-1} = -\eta_{\mathbf{x}}\nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we have

$$\Phi(\mathbf{x}_t) \leq \Phi(\mathbf{x}_{t-1}) - \eta_{\mathbf{x}}(\nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}))^\top \nabla\Phi(\mathbf{x}_{t-1}) + \eta_{\mathbf{x}}^2\kappa\ell\|\nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2$$

Using Young's inequality yields

$$\Phi(\mathbf{x}_t) \leq \Phi(\mathbf{x}_{t-1}) + \eta_{\mathbf{x}}^2\kappa\ell\|\nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 + \frac{\eta_{\mathbf{x}}}{2}(\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 - \|\nabla\Phi(\mathbf{x}_{t-1})\|^2).$$

By the Cauchy-Schwartz inequality, we have

$$\|\nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 \leq 2(\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 + \|\nabla\Phi(\mathbf{x}_{t-1})\|^2).$$

Putting these pieces yields the inequality for Algorithm 1.

We then prove the inequality for Algorithm 2. By using the same argument, we have

$$\Phi(\mathbf{x}_t) \leq \Phi(\mathbf{x}_{t-1}) - \eta_{\mathbf{x}}\|\nabla\Phi(\mathbf{x}_{t-1})\|^2 + \eta_{\mathbf{x}}^2\kappa\ell\|\hat{\mathbf{g}}_{\mathbf{x}}^{t-1}\|^2 + \eta_{\mathbf{x}}(\nabla\Phi(\mathbf{x}_{t-1}) - \hat{\mathbf{g}}_{\mathbf{x}}^{t-1})^\top \nabla\Phi(\mathbf{x}_{t-1}).$$

Conditioned on the iterate $(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of samples $\{\xi_i^{t-1}\}_{i=1}^M$) and use Lemma 28 to yield

$$\begin{aligned} \mathbb{E}[\Phi(\mathbf{x}_t) \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] &\leq \Phi(\mathbf{x}_{t-1}) - \eta_{\mathbf{x}}\|\nabla\Phi(\mathbf{x}_{t-1})\|^2 + \frac{\eta_{\mathbf{x}}^2\kappa\ell\sigma^2}{M} \\ &\quad + \eta_{\mathbf{x}}^2\kappa\ell\|\nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 + \eta_{\mathbf{x}}(\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}))^\top \nabla\Phi(\mathbf{x}_{t-1}). \end{aligned}$$

By using the same argument, we have

$$\begin{aligned} \mathbb{E}[\Phi(\mathbf{x}_t) \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] &\leq \Phi(\mathbf{x}_{t-1}) - \left(\frac{\eta_{\mathbf{x}}}{2} - 2\eta_{\mathbf{x}}^2\kappa\ell\right)\|\nabla\Phi(\mathbf{x}_{t-1})\|^2 \\ &\quad + \left(\frac{\eta_{\mathbf{x}}}{2} + 2\eta_{\mathbf{x}}^2\kappa\ell\right)\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 + \frac{\eta_{\mathbf{x}}^2\kappa\ell\sigma^2}{M}. \end{aligned}$$

We take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of all previous samples). This yields the inequality for Algorithm 2. \square

In the subsequent analysis, we define $\delta_t = \|\mathbf{y}^*(\mathbf{x}_t) - \mathbf{y}_t\|^2$ and $\delta_t = \mathbb{E}[\|\mathbf{y}^*(\mathbf{x}_t) - \mathbf{y}_t\|^2]$ for the iterates generated by Algorithms 1 and 2 respectively.

Lemma 30. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\delta_t \leq \left(1 - \frac{1}{2\kappa} + 4\kappa^3 \ell^2 \eta_{\mathbf{x}}^2\right) \delta_{t-1} + 4\kappa^3 \eta_{\mathbf{x}}^2 \|\nabla \Phi(\mathbf{x}_{t-1})\|^2.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\delta_t \leq \left(1 - \frac{1}{2\kappa} + 4\kappa^3 \ell^2 \eta_{\mathbf{x}}^2\right) \delta_{t-1} + 4\kappa^3 \eta_{\mathbf{x}}^2 \mathbb{E}[\|\nabla \Phi(\mathbf{x}_{t-1})\|^2] + \frac{2\eta_{\mathbf{x}}^2 \sigma^2 \kappa^3}{M} + \frac{\sigma^2}{\ell^2 M}.$$

Proof. We first prove the inequality for Algorithm 1. Since $\eta_{\mathbf{y}} = \frac{1}{\ell}$, we have

$$(\mathbf{y} - \mathbf{y}_t)^\top (\mathbf{y}_t - \mathbf{y}_{t-1} - \frac{1}{\ell} \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})) \geq 0, \quad \text{for all } \mathbf{y} \in \mathcal{Y}.$$

Rearranging the above inequality with $\mathbf{y} = \mathbf{y}^*(\mathbf{x}_{t-1})$ yields

$$\ell(\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t)^\top (\mathbf{y}_t - \mathbf{y}_{t-1}) \geq (\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}).$$

Since the function $f(\mathbf{x}_{t-1}, \cdot)$ is ℓ -smooth and μ -strongly concave, we have

$$(\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \geq \frac{\mu}{2} \|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_{t-1}\|^2 - \frac{\ell}{2} \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2.$$

In addition, we have

$$(\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t)^\top (\mathbf{y}_t - \mathbf{y}_{t-1}) = \frac{1}{2} (\|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t\|^2 - \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2).$$

Putting these pieces together yields

$$\|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t\|^2 \leq \left(1 - \frac{1}{\kappa}\right) \delta_{t-1}. \quad (\text{B.1})$$

Using Young's inequality and the fact that the function $\mathbf{y}^*(\cdot)$ is κ -Lipschitz, we have

$$\begin{aligned} \delta_t &\leq \left(1 + \frac{1}{2\kappa-2}\right) \|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t\|^2 + (2\kappa - 1) \|\mathbf{y}^*(\mathbf{x}_t) - \mathbf{y}^*(\mathbf{x}_{t-1})\|^2 \\ &\leq \frac{2\kappa-1}{2\kappa-2} \|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t\|^2 + 2\kappa \|\mathbf{y}^*(\mathbf{x}_t) - \mathbf{y}^*(\mathbf{x}_{t-1})\|^2 \\ &\stackrel{\text{Eq. (B.1)}}{\leq} \left(1 - \frac{1}{2\kappa}\right) \delta_{t-1} + 2\kappa \|\mathbf{y}^*(\mathbf{x}_t) - \mathbf{y}^*(\mathbf{x}_{t-1})\|^2 \\ &\leq \left(1 - \frac{1}{2\kappa}\right) \delta_{t-1} + 2\kappa^3 \|\mathbf{x}_t - \mathbf{x}_{t-1}\|^2. \end{aligned}$$

Furthermore, we have

$$\|\mathbf{x}_t - \mathbf{x}_{t-1}\|^2 = \eta_{\mathbf{x}}^2 \|\nabla_{\mathbf{x}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 \leq 2\eta_{\mathbf{x}}^2 (\ell^2 \delta_{t-1} + \|\nabla \Phi(\mathbf{x}_{t-1})\|^2).$$

Putting these pieces together yields the inequality for Algorithm 1.

We now prove the results for Algorithm 2. Applying the same argument used for proving Eq. (B.1), we have

$$\mathbb{E}[\|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_t\|^2] \leq \left(1 - \frac{1}{\kappa}\right) \delta_{t-1} + \frac{\sigma^2}{\ell^2 M},$$

which further implies

$$\delta_t \leq \left(1 - \frac{1}{2\kappa}\right) \delta_{t-1} + 2\kappa^3 \mathbb{E}[\|\mathbf{x}_t - \mathbf{x}_{t-1}\|^2] + \frac{\sigma^2}{\ell^2 M}.$$

Furthermore, we have

$$\mathbb{E}[\|\mathbf{x}_t - \mathbf{x}_{t-1}\|^2] = \eta_{\mathbf{x}}^2 \mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{x}}^{t-1}\|^2] \leq 2\eta_{\mathbf{x}}^2 (\ell^2 \delta_{t-1} + \mathbb{E}[\|\nabla \Phi(\mathbf{x}_{t-1})\|^2]) + \frac{\eta_{\mathbf{x}}^2 \sigma^2}{M}.$$

Putting these pieces together yields the inequality for Algorithm 2. \square

Lemma 31. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\Phi(\mathbf{x}_t) \leq \Phi(\mathbf{x}_{t-1}) - \frac{7\eta_{\mathbf{x}}}{16} \|\nabla\Phi(\mathbf{x}_{t-1})\|^2 + \frac{9\eta_{\mathbf{x}}\ell^2\delta_{t-1}}{16}.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\mathbb{E}[\Phi(\mathbf{x}_t)] \leq \mathbb{E}[\Phi(\mathbf{x}_{t-1})] - \frac{7\eta_{\mathbf{x}}}{16} \mathbb{E}[\|\nabla\Phi(\mathbf{x}_{t-1})\|^2] + \frac{9\eta_{\mathbf{x}}\ell^2\delta_{t-1}}{16} + \frac{\eta_{\mathbf{x}}^2\kappa\ell\sigma^2}{M}.$$

Proof. Since $\eta_{\mathbf{x}} = \frac{1}{16(\kappa+1)^2\bar{\ell}}$ in Algorithms 1 and 2, we have

$$\frac{7\eta_{\mathbf{x}}}{16} \leq \frac{\eta_{\mathbf{x}}}{2} - 2\eta_{\mathbf{x}}^2\kappa\ell \leq \frac{\eta_{\mathbf{x}}}{2} + 2\eta_{\mathbf{x}}^2\kappa\ell \leq \frac{9\eta_{\mathbf{x}}}{16}. \quad (\text{B.2})$$

We first prove the results for Algorithm 1. By combining Eq. (B.2) with the first inequality in Lemma 29, we have

$$\Phi(\mathbf{x}_t) \leq \Phi(\mathbf{x}_{t-1}) - \frac{7\eta_{\mathbf{x}}}{16} \|\nabla\Phi(\mathbf{x}_{t-1})\|^2 + \frac{9\eta_{\mathbf{x}}}{16} \|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2.$$

Since $\nabla\Phi(\mathbf{x}_{t-1}) = \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1}))$, we have

$$\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 \leq \ell^2 \|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_{t-1}\|^2 = \ell^2\delta_{t-1}.$$

Putting these pieces together yields the inequality for Algorithm 1.

We then prove the results for Algorithm 2. By combining Eq. (B.2) with the second inequality in Lemma 29, we have

$$\mathbb{E}[\Phi(\mathbf{x}_t)] \leq \mathbb{E}[\Phi(\mathbf{x}_{t-1})] - \frac{7\eta_{\mathbf{x}}}{16} \mathbb{E}[\|\nabla\Phi(\mathbf{x}_{t-1})\|^2] + \frac{9\eta_{\mathbf{x}}}{16} \mathbb{E}[\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2] + \frac{\eta_{\mathbf{x}}^2\kappa\ell\sigma^2}{M}.$$

Since $\nabla\Phi(\mathbf{x}_{t-1}) = \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1}))$, we have

$$\mathbb{E}[\|\nabla\Phi(\mathbf{x}_{t-1}) - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2] \leq \ell^2 \mathbb{E}[\|\mathbf{y}^*(\mathbf{x}_{t-1}) - \mathbf{y}_{t-1}\|^2] = \ell^2\delta_{t-1}.$$

Putting these pieces together yields the inequality for Algorithm 2. \square

Proof of Theorem 17. For simplicity, we define $\gamma = 1 - \frac{1}{2\kappa} + 4\kappa^3\ell^2\eta_{\mathbf{x}}^2$.

We first prove the results for Algorithm 1. Since $\delta_t = \|\mathbf{y}^*(\mathbf{x}_t) - \mathbf{y}_t\|^2$, we have $\delta_t \leq D^2$. Performing the first inequality in Lemma 30 recursively and using $\delta_0 \leq D^2$ yields

$$\delta_t \leq \gamma^t D^2 + 4\kappa^3\eta_{\mathbf{x}}^2 \left(\sum_{j=0}^{t-1} \gamma^{t-1-j} \|\nabla\Phi(\mathbf{x}_j)\|^2 \right). \quad (\text{B.3})$$

Combining Eq. (B.3) with the first inequality in Lemma 31 yields

$$\Phi(\mathbf{x}_t) \leq \Phi(\mathbf{x}_{t-1}) - \frac{7\eta_{\mathbf{x}}}{16} \|\nabla\Phi(\mathbf{x}_{t-1})\|^2 + \frac{9\eta_{\mathbf{x}}\ell^2\gamma^{t-1}D^2}{16} + \frac{9\eta_{\mathbf{x}}^3\ell^2\kappa^3}{4} \left(\sum_{j=0}^{t-2} \gamma^{t-2-j} \|\nabla\Phi(\mathbf{x}_j)\|^2 \right). \quad (\text{B.4})$$

Summing Eq. (B.4) over $t = 1, 2, \dots, T+1$ and rearranging the terms yields

$$\Phi(\mathbf{x}_{T+1}) \leq \Phi(\mathbf{x}_0) - \frac{7\eta_{\mathbf{x}}}{16} \left(\sum_{t=0}^T \|\nabla\Phi(\mathbf{x}_t)\|^2 \right) + \frac{9\eta_{\mathbf{x}}\ell^2D^2}{16} \left(\sum_{t=0}^T \gamma^t \right) + \frac{9\eta_{\mathbf{x}}^3\ell^2\kappa^3}{4} \left(\sum_{t=1}^{T+1} \sum_{j=0}^{t-2} \gamma^{t-2-j} \|\nabla\Phi(\mathbf{x}_j)\|^2 \right).$$

Since $\eta_{\mathbf{x}} = \frac{1}{16(\kappa+1)^2\ell}$, we have $\gamma \leq 1 - \frac{1}{4\kappa}$ and $\frac{9\eta_{\mathbf{x}}^3\ell^2\kappa^3}{4} \leq \frac{9\eta_{\mathbf{x}}}{1024\kappa}$. This yields $\sum_{t=0}^T \gamma^t \leq 4\kappa$ and

$$\sum_{t=1}^{T+1} \sum_{j=0}^{t-2} \gamma^{t-2-j} \|\nabla\Phi(\mathbf{x}_j)\|^2 \leq 4\kappa \left(\sum_{t=0}^T \|\nabla\Phi(\mathbf{x}_t)\|^2 \right)$$

Putting these pieces together yields

$$\Phi(\mathbf{x}_{T+1}) \leq \Phi(\mathbf{x}_0) - \frac{103\eta_{\mathbf{x}}}{256} \left(\sum_{t=0}^T \|\nabla\Phi(\mathbf{x}_t)\|^2 \right) + \frac{9\eta_{\mathbf{x}}\kappa\ell^2 D^2}{4}.$$

By the definition of Δ_{Φ} , we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla\Phi(\mathbf{x}_t)\|^2 \right) \leq \frac{256(\Phi(\mathbf{x}_0) - \Phi(\mathbf{x}_{T+1}))}{103\eta_{\mathbf{x}}(T+1)} + \frac{576\kappa\ell^2 D^2}{103(T+1)} \leq \frac{128\kappa^2\ell\Delta_{\Phi} + 5\kappa\ell^2 D^2}{T+1}.$$

This implies that the number of gradient evaluations required by Algorithm 1 to return an ϵ -stationary point is

$$O\left(\frac{\kappa^2\ell\Delta_{\Phi} + \kappa\ell^2 D^2}{\epsilon^2}\right).$$

We then prove the results for Algorithm 2. Applying the same argument used for analyzing Algorithm 1 but with the second inequalities in Lemmas 30 and 31, we have

$$\begin{aligned} \frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla\Phi(\mathbf{x}_t)\|^2] \right) &\leq \frac{256(\Phi(\mathbf{x}_0) - \mathbb{E}[\Phi(\mathbf{x}_{T+1})])}{103\eta_{\mathbf{x}}(T+1)} + \frac{576\kappa\ell^2 D^2}{103(T+1)} + \frac{2304\kappa\sigma^2}{103M} \\ &\leq \frac{128\kappa^2\ell\Delta_{\Phi} + 5\kappa\ell^2 D^2}{T+1} + \frac{24\kappa\sigma^2}{M}. \end{aligned}$$

This implies that the number of stochastic gradient evaluations required by Algorithm 2 to return an ϵ -stationary point is

$$O\left(\frac{\kappa^2\ell\Delta_{\Phi} + \kappa\ell^2 D^2}{\epsilon^2} \max\left\{1, \frac{\kappa\sigma^2}{\epsilon^2}\right\}\right).$$

This completes the proof.

Appendix C. Smooth and Nonconvex-Concave Setting

We provide the proof for Theorem 19. Throughout this subsection, we set

$$\eta_{\mathbf{x}}^t \equiv \eta_{\mathbf{x}} = \min\left\{\frac{\epsilon^2}{80\ell L^2}, \frac{\epsilon^4}{4096\ell^3 L^2 D^2}\right\}, \quad \eta_{\mathbf{y}}^t \equiv \eta_{\mathbf{y}} = \frac{1}{\ell},$$

for Algorithm 1, and

$$\eta_{\mathbf{x}}^t \equiv \eta_{\mathbf{x}} = \min\left\{\frac{\epsilon^2}{80\ell(L^2 + \sigma^2)}, \frac{\epsilon^4}{8192\ell^3(L^2 + \sigma^2)D^2}, \frac{\epsilon^6}{131072\ell^3(L^2 + \sigma^2)D^2\sigma^2}\right\}, \quad \eta_{\mathbf{y}}^t \equiv \eta_{\mathbf{y}} = \min\left\{\frac{1}{2\ell}, \frac{\epsilon^2}{32\ell\sigma^2}\right\},$$

for Algorithm 2. We define $\Delta_t = \Phi(\mathbf{x}_t) - f(\mathbf{x}_t, \mathbf{y}_t)$ and $\bar{\Delta}_t = \mathbb{E}[\Phi(\mathbf{x}_t) - f(\mathbf{x}_t, \mathbf{y}_t)]$ for the iterates generated by Algorithms 1 and 2 respectively.

Lemma 32. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\Phi_{1/2\ell}(\mathbf{x}_t) \leq \Phi_{1/2\ell}(\mathbf{x}_{t-1}) + 2\eta_{\mathbf{x}}\ell\Delta_{t-1} - \frac{1}{4}\eta_{\mathbf{x}}\|\nabla\Phi_{1/2\ell}(\mathbf{x}_{t-1})\|^2 + \eta_{\mathbf{x}}^2\ell L^2.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\mathbb{E}[\Phi_{1/2\ell}(\mathbf{x}_t)] \leq \mathbb{E}[\Phi_{1/2\ell}(\mathbf{x}_{t-1})] + 2\eta_{\mathbf{x}}\ell\Delta_{t-1} - \frac{1}{4}\eta_{\mathbf{x}}\mathbb{E}[\|\nabla\Phi_{1/2\ell}(\mathbf{x}_{t-1})\|^2] + \eta_{\mathbf{x}}^2\ell(L^2 + \sigma^2).$$

Proof. We first prove the results for Algorithm 1. Letting $\hat{\mathbf{x}}_{t-1} = \text{prox}_{\Phi/2\ell}(\mathbf{x}_{t-1})$, we have

$$\Phi_{1/2\ell}(\mathbf{x}_t) \leq \Phi(\hat{\mathbf{x}}_{t-1}) + \ell\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t\|^2.$$

Since the function $f(\cdot, \mathbf{y})$ is L -Lipschitz for each $\mathbf{y} \in \mathcal{Y}$, we have $\|\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y})\| \leq L$ for each pair of \mathbf{x} and $\mathbf{y} \in \mathcal{Y}$. This yields

$$\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t\|^2 \leq \|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2 + 2\eta_{\mathbf{x}}\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \rangle + \eta_{\mathbf{x}}^2L^2.$$

Combining the above two inequalities yields

$$\Phi_{1/2\ell}(\mathbf{x}_t) \leq \Phi_{1/2\ell}(\mathbf{x}_{t-1}) + 2\eta_{\mathbf{x}}\ell\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \rangle + \eta_{\mathbf{x}}^2\ell L^2. \quad (\text{C.1})$$

Since the function $f(\cdot, \mathbf{y}_{t-1})$ is ℓ -smooth, we have

$$\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \rangle \leq f(\hat{\mathbf{x}}_{t-1}, \mathbf{y}_{t-1}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + \frac{\ell}{2}\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2. \quad (\text{C.2})$$

Since $\Phi(\mathbf{x}) = \max_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y})$ and $\hat{\mathbf{x}}_{t-1} = \text{prox}_{\Phi/2\ell}(\mathbf{x}_{t-1})$, we have

$$f(\hat{\mathbf{x}}_{t-1}, \mathbf{y}_{t-1}) + \ell\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2 \leq \Phi(\hat{\mathbf{x}}_{t-1}) + \ell\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2 \leq \Phi(\mathbf{x}_{t-1}).$$

Since $\|\nabla\Phi_{1/2\ell}(\mathbf{x}_{t-1})\| = 2\ell\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|$, we have

$$f(\hat{\mathbf{x}}_{t-1}, \mathbf{y}_{t-1}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + \frac{\ell}{2}\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2 \leq \Phi(\mathbf{x}_{t-1}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) - \frac{1}{8\ell}\|\nabla\Phi_{1/2\ell}(\mathbf{x}_{t-1})\|^2. \quad (\text{C.3})$$

Plugging Eq. (C.2) and Eq. (C.3) into Eq. (C.1) yields the inequality for Algorithm 1.

We then prove the results for Algorithm 2. By using the same argument, we have

$$\begin{aligned} \|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t\|^2 &= \|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2 + 2\eta_{\mathbf{x}}\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \hat{\mathbf{g}}_{\mathbf{x}}^{t-1} \rangle + \eta_{\mathbf{x}}^2\|\hat{\mathbf{g}}_{\mathbf{x}}^{t-1} - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2 \\ &\quad + 2\eta_{\mathbf{x}}^2\langle \hat{\mathbf{g}}_{\mathbf{x}}^{t-1} - \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}), \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \rangle + \eta_{\mathbf{x}}^2L^2. \end{aligned}$$

Conditioned on the iterate $(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of samples $\{\zeta_i^{t-1}\}_{i=1}^M$) and use Lemma 28 with $M = 1$ to yield

$$\mathbb{E}[\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t\|^2 \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] \leq \|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2 + 2\eta_{\mathbf{x}}\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \rangle + \eta_{\mathbf{x}}^2(L^2 + \sigma^2).$$

We take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of all previous samples). This yields

$$\mathbb{E}[\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t\|^2] \leq \mathbb{E}[\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2] + 2\eta_{\mathbf{x}}\mathbb{E}[\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \rangle] + \eta_{\mathbf{x}}^2(L^2 + \sigma^2).$$

Applying the same argument used for proving Eq. (C.1), we have

$$\mathbb{E}[\Phi_{1/2\ell}(\mathbf{x}_t)] \leq \mathbb{E}[\Phi_{1/2\ell}(\mathbf{x}_{t-1})] + 2\eta_{\mathbf{x}}\ell\mathbb{E}[\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \nabla_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \rangle] + \eta_{\mathbf{x}}^2\ell(L^2 + \sigma^2). \quad (\text{C.4})$$

Plugging Eq. (C.2) and Eq. (C.3) into Eq. (C.4) yields the inequality for Algorithm 2. \square

Lemma 33. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy for $\forall s \leq t-1$ that*

$$\Delta_{t-1} \leq \eta_{\mathbf{x}} L^2 (2t - 2s - 1) + \frac{\ell}{2} (\|\mathbf{y}_{t-1} - \mathbf{y}^*(\mathbf{x}_s)\|^2 - \|\mathbf{y}_t - \mathbf{y}^*(\mathbf{x}_s)\|^2) + (f(\mathbf{x}_t, \mathbf{y}_t) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})).$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy for $\forall s \leq t-1$ that

$$\begin{aligned} \Delta_{t-1} &\leq \eta_{\mathbf{x}} (L^2 + \sigma^2) (2t - 2s - 1) + \frac{1}{2\eta_{\mathbf{y}}} (\mathbb{E}[\|\mathbf{y}_{t-1} - \mathbf{y}^*(\mathbf{x}_s)\|^2] - \mathbb{E}[\|\mathbf{y}_t - \mathbf{y}^*(\mathbf{x}_s)\|^2]) \\ &\quad + \mathbb{E}[f(\mathbf{x}_t, \mathbf{y}_t) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})] + \eta_{\mathbf{y}} \sigma^2. \end{aligned}$$

Proof. We first prove the results for Algorithm 1. By the definition of Δ_{t-1} , we have

$$\Delta_{t-1} \leq (f(\mathbf{x}_t, \mathbf{y}_t) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})) + \mathbf{A} + \mathbf{B} + \mathbf{C},$$

where

$$\begin{aligned} \mathbf{A} &= f(\mathbf{x}_{t-1}, \mathbf{y}_t) - f(\mathbf{x}_t, \mathbf{y}_t), \\ \mathbf{B} &= f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)) - f(\mathbf{x}_{t-1}, \mathbf{y}_t), \\ \mathbf{C} &= f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1})) - f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)). \end{aligned}$$

Since $f(\cdot, \mathbf{y})$ is L -Lipschitz for each $\mathbf{y} \in \mathcal{Y}$, we have $\|\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y})\| \leq L$ for each pair of \mathbf{x} and $\mathbf{y} \in \mathcal{Y}$. This yields

$$\mathbf{A} \leq L \|\mathbf{x}_{t-1} - \mathbf{x}_t\| = \eta_{\mathbf{x}} L \|\nabla_{\mathbf{x}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\| \leq \eta_{\mathbf{x}} L^2.$$

By using the same argument and that $f(\mathbf{x}_s, \mathbf{y}^*(\mathbf{x}_s)) \geq f(\mathbf{x}_s, \mathbf{y})$, we have

$$\begin{aligned} \mathbf{C} &= f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1})) - f(\mathbf{x}_s, \mathbf{y}^*(\mathbf{x}_{t-1})) + f(\mathbf{x}_s, \mathbf{y}^*(\mathbf{x}_{t-1})) - f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)) \\ &\leq f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1})) - f(\mathbf{x}_s, \mathbf{y}^*(\mathbf{x}_{t-1})) + f(\mathbf{x}_s, \mathbf{y}^*(\mathbf{x}_s)) - f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)) \\ &\leq 2L \|\mathbf{x}_{t-1} - \mathbf{x}_s\| \leq 2\eta_{\mathbf{x}} L \left(\sum_{k=s}^{t-1} \|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\| \right) \leq 2\eta_{\mathbf{x}} L^2 (t - s - 1). \end{aligned}$$

It suffices to bound the term \mathbf{B} . Since $\eta_{\mathbf{y}} = \frac{1}{\ell}$, we have

$$(\mathbf{y} - \mathbf{y}_t)^\top (\mathbf{y}_t - \mathbf{y}_{t-1} - \frac{1}{\ell} \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})) \geq 0, \quad \text{for all } \mathbf{y} \in \mathcal{Y}.$$

Rearranging the above inequality yields

$$(\mathbf{y} - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \leq \frac{\ell}{2} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y} - \mathbf{y}_t\|^2 - \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2).$$

Since the function $f(\mathbf{x}_{t-1}, \cdot)$ is concave and ℓ -smooth, we have

$$\begin{aligned} (\mathbf{y} - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) &= (\mathbf{y} - \mathbf{y}_{t-1})^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + (\mathbf{y}_{t-1} - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \\ &\geq f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) - f(\mathbf{x}_{t-1}, \mathbf{y}_t) - \frac{\ell}{2} \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2. \end{aligned}$$

Putting these pieces together yields

$$f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_t) \leq \frac{\ell}{2} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y} - \mathbf{y}_t\|^2). \quad (\text{C.5})$$

Since $\mathbf{y}^*(\mathbf{x}_s) \in \mathcal{Y}$ for $s \leq t-1$, Eq. (C.5) implies

$$\mathbf{B} \leq \frac{\ell}{2} (\|\mathbf{y}^*(\mathbf{x}_s) - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y}^*(\mathbf{x}_s) - \mathbf{y}_t\|^2).$$

Putting these pieces together yields the inequality for Algorithm 1.

We then prove the results for Algorithm 2. By the definition of Δ_{t-1} , we have

$$\Delta_{t-1} \leq \mathbb{E}[f(\mathbf{x}_t, \mathbf{y}_t) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})] + \mathbf{A} + \mathbf{B} + \mathbf{C},$$

where

$$\begin{aligned} \mathbf{A} &= \mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}_t) - f(\mathbf{x}_t, \mathbf{y}_t)], \\ \mathbf{B} &= \mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)) - f(\mathbf{x}_{t-1}, \mathbf{y}_t)], \\ \mathbf{C} &= \mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1})) - f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s))]. \end{aligned}$$

Applying the same argument used for analyzing Algorithm 1, we have

$$\mathbf{A} + \mathbf{C} \leq \eta_{\mathbf{x}} L \left(\mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{x}}^{t-1}\|] + 2 \left(\sum_{k=s}^{t-1} \mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{x}}^k\|] \right) \right).$$

For $\forall k \in \{s, s+1, \dots, t-1\}$, we have

$$\|\hat{\mathbf{g}}_{\mathbf{x}}^k\|^2 = \|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 + 2 \langle \hat{\mathbf{g}}_{\mathbf{x}}^k - \nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k), \nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k) \rangle + \|\hat{\mathbf{g}}_{\mathbf{x}}^k - \nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2.$$

Conditioned on the iterate $(\mathbf{x}_k, \mathbf{y}_k)$, we take the expectation of both sides of this equation (where the expectation is taken over the randomness in the selection of samples $\{\xi_i^k\}_{i=1}^M$) and use Lemma 28 with $M = 1$ to yield

$$\mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{x}}^k\|^2 \mid \mathbf{x}_k, \mathbf{y}_k] \leq \|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 + \sigma^2 \leq L^2 + \sigma^2.$$

We take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of all previous samples). This yields

$$\mathbb{E}[\|\hat{\mathbf{g}}_{\mathbf{x}}^k\|^2] \leq L^2 + \sigma^2.$$

Putting these pieces together yields

$$\mathbf{A} + \mathbf{C} \leq \eta_{\mathbf{x}} (L^2 + \sigma^2) (2t - 2s - 1).$$

It suffices to bound the term \mathbf{B} . By using the same argument, we have

$$\eta_{\mathbf{y}} (\mathbf{y} - \mathbf{y}_t)^\top \hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \leq \frac{1}{2} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y} - \mathbf{y}_t\|^2 - \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2).$$

Using the Young's inequality, we have

$$\begin{aligned} \eta_{\mathbf{y}} (\mathbf{y} - \mathbf{y}_t)^\top \hat{\mathbf{g}}_{\mathbf{y}}^{t-1} &\geq \eta_{\mathbf{y}} (\mathbf{y} - \mathbf{y}_{t-1})^\top \hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \\ &\quad + \eta_{\mathbf{y}} (\mathbf{y}_{t-1} - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) - \frac{1}{4} \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2 - \eta_{\mathbf{y}}^2 \|\hat{\mathbf{g}}_{\mathbf{y}}^{t-1} - \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})\|^2. \end{aligned}$$

Conditioned on the iterate $(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of samples $\{\xi_i^k\}_{i=1}^M$) and use Lemma 28 with $M = 1$ to yield

$$\begin{aligned} & (\mathbf{y} - \mathbf{y}_{t-1})^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + \mathbb{E}[(\mathbf{y}_{t-1} - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] \\ & \leq \frac{1}{2\eta_{\mathbf{y}}} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \mathbb{E}[\|\mathbf{y} - \mathbf{y}_t\|^2 \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}]) - \frac{1}{4\eta_{\mathbf{y}}} \mathbb{E}[\|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2 \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] + \eta_{\mathbf{y}} \sigma^2. \end{aligned}$$

We take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of all previous samples). This yields

$$\mathbb{E}[(\mathbf{y} - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})] \leq \frac{1}{2\eta_{\mathbf{y}}} (\mathbb{E}[\|\mathbf{y} - \mathbf{y}_{t-1}\|^2] - \mathbb{E}[\|\mathbf{y} - \mathbf{y}_t\|^2]) - \frac{1}{4\eta_{\mathbf{y}}} \mathbb{E}[\|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2] + \eta_{\mathbf{y}} \sigma^2.$$

Since the function $f(\mathbf{x}_{t-1}, \cdot)$ is concave and ℓ -smooth, we have

$$(\mathbf{y} - \mathbf{y}_t)^\top \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \geq f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_t) - \frac{\ell}{2} \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2.$$

Since $\eta_{\mathbf{y}} \leq \frac{1}{2\ell}$, we have

$$\mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_t)] \leq \frac{1}{2\eta_{\mathbf{y}}} (\mathbb{E}[\|\mathbf{y} - \mathbf{y}_{t-1}\|^2] - \mathbb{E}[\|\mathbf{y} - \mathbf{y}_t\|^2]) + \eta_{\mathbf{y}} \sigma^2. \quad (\text{C.6})$$

Since $\mathbf{y}^*(\mathbf{x}_s) \in \mathcal{Y}$ for $s \leq t-1$, Eq. (C.6) implies

$$\mathbf{B} \leq \frac{1}{2\eta_{\mathbf{y}}} (\mathbb{E}[\|\mathbf{y}^*(\mathbf{x}_s) - \mathbf{y}_{t-1}\|^2] - \mathbb{E}[\|\mathbf{y}^*(\mathbf{x}_s) - \mathbf{y}_t\|^2]) + \eta_{\mathbf{y}} \sigma^2.$$

Putting these pieces together yields the inequality for Algorithm 2. \square

Without loss of generality, we let $B \leq T+1$ satisfy that $\frac{T+1}{B}$ is an integer. The following lemma bounds $\frac{1}{T+1} (\sum_{t=0}^T \Delta_t)$ for Algorithm 1 and 2 using a novel localization technique.

Lemma 34. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}} (B+1) L^2 + \frac{\ell D^2}{2B} + \frac{\Delta_0}{T+1}.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}} (B+1) (L^2 + \sigma^2) + \frac{D^2}{2B\eta_{\mathbf{y}}} + \eta_{\mathbf{y}} \sigma^2 + \frac{\Delta_0}{T+1}.$$

Proof. We first prove the results for Algorithm 1. Indeed, we can divide $\{\Delta_t\}_{t=0}^T$ into several blocks in which each block contains at most B terms, given by

$$\{\Delta_t\}_{t=0}^{B-1}, \{\Delta_t\}_{t=B}^{2B-1}, \dots, \{\Delta_t\}_{t=T-B+1}^T.$$

Then, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \frac{B}{T+1} \left(\sum_{j=0}^{\frac{T+1}{B}-1} \left(\frac{1}{B} \sum_{t=jB}^{(j+1)B-1} \Delta_t \right) \right). \quad (\text{C.7})$$

Further, letting $s = 0$ in the first inequality in Lemma 33 yields

$$\sum_{t=0}^{B-1} \Delta_t \leq \eta_{\mathbf{x}} L^2 B^2 + \frac{\ell D^2}{2} + (f(\mathbf{x}_B, \mathbf{y}_B) - f(\mathbf{x}_0, \mathbf{y}_0)).$$

Similarly, letting $s = jB$ yields, for $1 \leq j \leq \frac{T+1}{B} - 1$,

$$\sum_{t=jB}^{(j+1)B-1} \Delta_t \leq \eta_{\mathbf{x}} L^2 B^2 + \frac{\ell D^2}{2} + (f(\mathbf{x}_{jB+B}, \mathbf{y}_{jB+B}) - f(\mathbf{x}_{jB}, \mathbf{y}_{jB})).$$

Plugging the above two inequalities into Eq. (C.7) yields

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}} L^2 B + \frac{\ell D^2}{2B} + \frac{f(\mathbf{x}_{T+1}, \mathbf{y}_{T+1}) - f(\mathbf{x}_0, \mathbf{y}_0)}{T+1}. \quad (\text{C.8})$$

Since $f(\cdot, \mathbf{y})$ is L -Lipschitz for each $\mathbf{y} \in \mathcal{Y}$, we have $\|\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y})\| \leq L$ for each pair of \mathbf{x} and $\mathbf{y} \in \mathcal{Y}$. This yields

$$\begin{aligned} f(\mathbf{x}_{T+1}, \mathbf{y}_{T+1}) - f(\mathbf{x}_0, \mathbf{y}_0) &= f(\mathbf{x}_{T+1}, \mathbf{y}_{T+1}) - f(\mathbf{x}_0, \mathbf{y}_{T+1}) + f(\mathbf{x}_0, \mathbf{y}_{T+1}) - f(\mathbf{x}_0, \mathbf{y}_0) \\ &\leq \eta_{\mathbf{x}} L^2 (T+1) + \Delta_0. \end{aligned}$$

Plugging this inequality into Eq. (C.8) yields the inequality for Algorithm 1.

We then prove the results for Algorithm 2. Applying the same argument used for analyzing Algorithm 1 but with the second inequality in Lemma 33, we can obtain the inequality for Algorithm 2. \square

Proof of Theorem 19. We first prove the results for Algorithm 1. Summing the first inequality in Lemma 32 over $t = 1, \dots, T+1$ yields

$$\Phi_{1/2\ell}(\mathbf{x}_{T+1}) \leq \Phi_{1/2\ell}(\mathbf{x}_0) + 2\eta_{\mathbf{x}} \ell \left(\sum_{t=0}^T \Delta_t \right) - \frac{\eta_{\mathbf{x}}}{4} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2 \right) + \eta_{\mathbf{x}}^2 \ell L^2 (T+1).$$

Combining the above inequality with the first inequality in Lemma 34 yields

$$\begin{aligned} \Phi_{1/2\ell}(\mathbf{x}_{T+1}) &\leq \Phi_{1/2\ell}(\mathbf{x}_0) + 2\eta_{\mathbf{x}} \ell (T+1) (\eta_{\mathbf{x}} L^2 (B+1) + \frac{\ell D^2}{2B}) + 2\eta_{\mathbf{x}} \ell \Delta_0 \\ &\quad - \frac{\eta_{\mathbf{x}}}{4} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2 \right) + \eta_{\mathbf{x}}^2 \ell L^2 (T+1). \end{aligned}$$

By the definition of the function $\Phi_{1/2\ell}(\cdot)$, we have

$$\begin{aligned} \Phi_{1/2\ell}(\mathbf{x}_0) &= \min_{\mathbf{w}} \Phi(\mathbf{w}) + \ell \|\mathbf{w} - \mathbf{x}_0\|^2 \leq \Phi(\mathbf{x}_0), \\ \Phi_{1/2\ell}(\mathbf{x}_{T+1}) &= \min_{\mathbf{w}} \Phi(\mathbf{w}) + \ell \|\mathbf{w} - \mathbf{x}_{T+1}\|^2 \geq \min_{\mathbf{x}} \Phi(\mathbf{x}), \end{aligned}$$

which implies $\Phi_{1/2\ell}(\mathbf{x}_0) - \Phi_{1/2\ell}(\mathbf{x}_{T+1}) \leq \Delta_{\Phi}$. Putting these pieces together yields

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 8\ell(\eta_{\mathbf{x}}(B+1)L^2 + \frac{\ell D^2}{2B}) + \frac{8\ell\Delta_0}{T+1} + 4\eta_{\mathbf{x}}\ell L^2.$$

Letting $B = \left\lfloor \frac{D}{2L} \sqrt{\frac{\ell}{\eta_{\mathbf{x}}}} \right\rfloor + 1$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{8\ell\Delta_0}{T+1} + 16\ell LD\sqrt{\ell\eta_{\mathbf{x}}} + 20\eta_{\mathbf{x}}\ell L^2.$$

By the definition of $\eta_{\mathbf{x}}$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{8\ell\Delta_0}{T+1} + \frac{\epsilon^2}{2}.$$

This implies that the number of gradient evaluations required by Algorithm 1 to return an ϵ -stationary point is

$$O\left(\left(\frac{\ell L^2 \Delta_{\Phi}}{\epsilon^4} + \frac{\ell \Delta_0}{\epsilon^2}\right) \max\left\{1, \frac{\ell^2 D^2}{\epsilon^2}\right\}\right).$$

We then prove the results for Algorithm 2. Applying the same argument used for analyzing Algorithm 1 but with the second inequalities in Lemmas 32 and 34, we have

$$\begin{aligned} \frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2] \right) &\leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 8\ell(\eta_{\mathbf{x}}(L^2 + \sigma^2)(B+1) + \frac{D^2}{2B\eta_{\mathbf{y}}}) + \eta_{\mathbf{y}}\sigma^2 \\ &\quad + \frac{8\ell\Delta_0}{T+1} + 4\eta_{\mathbf{x}}\ell(L^2 + \sigma^2). \end{aligned}$$

Letting $B = \left\lfloor \frac{D}{2} \sqrt{\frac{1}{\eta_{\mathbf{x}}\eta_{\mathbf{y}}(L^2 + \sigma^2)}} \right\rfloor + 1$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2] \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{8\ell\Delta_0}{T+1} + 16\ell D\sqrt{\frac{\eta_{\mathbf{x}}(L^2 + \sigma^2)}{\eta_{\mathbf{y}}}} + 8\eta_{\mathbf{y}}\ell\sigma^2 + 20\eta_{\mathbf{x}}\ell(L^2 + \sigma^2).$$

By the definition of $\eta_{\mathbf{x}}$ and $\eta_{\mathbf{y}}$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla \Phi_{1/2\ell}(\mathbf{x}_t)\|^2] \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{8\ell\Delta_0}{T+1} + \frac{3\epsilon^2}{4}.$$

This implies that the number of stochastic gradient evaluations required by Algorithm 2 to return an ϵ -stationary point is

$$O\left(\left(\frac{\ell(L^2 + \sigma^2)\widehat{\Delta}_{\Phi}}{\epsilon^4} + \frac{\ell\Delta_0}{\epsilon^2}\right) \max\left\{1, \frac{\ell^2 D^2}{\epsilon^2}, \frac{\ell^2 D^2 \sigma^2}{\epsilon^4}\right\}\right).$$

This completes the proof.

Appendix D. Nonsmooth and Nonconvex-Strongly-Concave Setting

We provide the proof for Theorem 24. Throughout this subsection, we set

$$\eta_{\mathbf{x}}^t \equiv \eta_{\mathbf{x}} = \min\left\{\frac{\epsilon^2}{48\rho L^2}, \frac{\mu\epsilon^4}{4096\rho^2 L^4}, \frac{\mu\epsilon^4}{4096\rho^2 L^4 \log^2(1+4096\rho^2 L^4 \mu^{-2}\epsilon^{-4})}\right\},$$

and

$$\eta_{\mathbf{y}}^t = \begin{cases} \frac{1}{\mu^t}, & \text{if } 1 \leq t \leq B, \\ \frac{1}{\mu^{(t-B)}}, & \text{else if } B+1 \leq t \leq 2B, \\ \vdots & \vdots \\ \frac{1}{\mu^{(t-jB)}}, & \text{else if } jB+1 \leq t \leq (j+1)B, \\ \vdots & \vdots \end{cases} \quad \text{for } B = \left\lfloor \sqrt{\frac{1}{\mu\eta_{\mathbf{x}}}} \right\rfloor + 1, \quad (\text{D.1})$$

for Algorithm 1, and

$$\eta_{\mathbf{x}}^t \equiv \eta_{\mathbf{x}} = \min \left\{ \frac{\epsilon^2}{48\rho(L^2+\sigma^2)}, \frac{\mu\epsilon^4}{4096\rho^2(L^2+\sigma^2)^2}, \frac{\mu\epsilon^4}{4096\rho^2(L^2+\sigma^2)^2 \log^2(1+4096\rho^2(L^2+\sigma^2)^2\mu^{-2}\epsilon^{-4})} \right\},$$

and $\eta_{\mathbf{y}}^t$ in Eq. (D.1) for Algorithm 2. We define $\Delta_t = \Phi(\mathbf{x}_t) - f(\mathbf{x}_t, \mathbf{y}_t)$ and $\Delta_t = \mathbb{E}[\Phi(\mathbf{x}_t) - f(\mathbf{x}_t, \mathbf{y}_t)]$ for the iterates generated by Algorithms 1 and 2 respectively.

Lemma 35. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\Phi_{1/2\rho}(\mathbf{x}_t) \leq \Phi_{1/2\rho}(\mathbf{x}_{t-1}) + 2\eta_{\mathbf{x}}\rho\Delta_{t-1} - \frac{1}{4}\eta_{\mathbf{x}}\|\nabla\Phi_{1/2\rho}(\mathbf{x}_{t-1})\|^2 + \eta_{\mathbf{x}}^2\rho L^2.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\mathbb{E}[\Phi_{1/2\rho}(\mathbf{x}_t)] \leq \mathbb{E}[\Phi_{1/2\rho}(\mathbf{x}_{t-1})] + 2\eta_{\mathbf{x}}\rho\Delta_{t-1} - \frac{1}{4}\eta_{\mathbf{x}}\mathbb{E}[\|\nabla\Phi_{1/2\rho}(\mathbf{x}_{t-1})\|^2] + \eta_{\mathbf{x}}^2\rho(L^2 + \sigma^2).$$

Proof. We first prove the results for Algorithm 1. Applying the same argument used for proving Lemma 32, we have

$$\Phi_{1/2\rho}(\mathbf{x}_t) \leq \Phi_{1/2\rho}(\mathbf{x}_{t-1}) + 2\eta_{\mathbf{x}}\rho\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \mathbf{g}_{\mathbf{x}}^{t-1} \rangle + \eta_{\mathbf{x}}^2\rho L^2. \quad (\text{D.2})$$

Since the function $f(\cdot, \mathbf{y}_{t-1})$ is ρ -weakly convex and $\mathbf{g}_{\mathbf{x}}^{t-1} \in \partial_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we have

$$\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \mathbf{g}_{\mathbf{x}}^{t-1} \rangle \leq f(\hat{\mathbf{x}}_{t-1}, \mathbf{y}_{t-1}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + \frac{\rho}{2}\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2. \quad (\text{D.3})$$

Applying the same argument used for proving Lemma 32, we have

$$f(\hat{\mathbf{x}}_{t-1}, \mathbf{y}_{t-1}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + \frac{\rho}{2}\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2 \leq \Phi(\mathbf{x}_{t-1}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) - \frac{1}{8\rho}\|\nabla\Phi_{1/2\rho}(\mathbf{x}_{t-1})\|^2. \quad (\text{D.4})$$

Plugging Eq. (D.3) and Eq. (D.4) into Eq. (D.2) yields the inequality for Algorithm 1.

We then prove the results for Algorithm 2. Applying the same argument used for proving Lemma 32, we have

$$\mathbb{E}[\Phi_{1/2\rho}(\mathbf{x}_t)] \leq \mathbb{E}[\Phi_{1/2\rho}(\mathbf{x}_{t-1})] + 2\eta_{\mathbf{x}}\rho\mathbb{E}[\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{x}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] \rangle] + \eta_{\mathbf{x}}^2\rho(L^2 + \sigma^2).$$

Since $\mathbb{E}[\hat{\mathbf{g}}_{\mathbf{x}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] \in \partial_{\mathbf{x}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we have

$$\mathbb{E}[\langle \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}, \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{x}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] \rangle] \leq \mathbb{E}[f(\hat{\mathbf{x}}_{t-1}, \mathbf{y}_{t-1})] - \mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})] + \frac{\rho}{2}\mathbb{E}[\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\|^2].$$

Combining two above inequalities with Eq. (D.4) yields the inequality for Algorithm 2. \square

Lemma 36. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy for $\forall s \leq t-1$ that*

$$\Delta_{t-1} \leq 2\eta_{\mathbf{x}}L^2(t-s-1) + \left(\frac{1}{2\eta_{\mathbf{y}}^{t-1}} - \frac{\mu}{2}\right)\|\mathbf{y}_{t-1} - \mathbf{y}^*(\mathbf{x}_s)\|^2 - \frac{1}{2\eta_{\mathbf{y}}^{t-1}}\|\mathbf{y}_t - \mathbf{y}^*(\mathbf{x}_s)\|^2 + \frac{1}{2}\eta_{\mathbf{y}}^{t-1}L^2.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy for $\forall s \leq t-1$ that

$$\Delta_{t-1} \leq 2\eta_{\mathbf{x}}(L^2 + \sigma^2)(t-s-1) + \left(\frac{1}{2\eta_{\mathbf{y}}^{t-1}} - \frac{\mu}{2}\right)\mathbb{E}[\|\mathbf{y}_{t-1} - \mathbf{y}^*(\mathbf{x}_s)\|^2] - \frac{1}{2\eta_{\mathbf{y}}^{t-1}}\mathbb{E}[\|\mathbf{y}_t - \mathbf{y}^*(\mathbf{x}_s)\|^2] + \eta_{\mathbf{y}}^{t-1}(L^2 + \sigma^2).$$

Proof. We first prove the results for Algorithm 1. By the definition of Δ_{t-1} , we have

$$\Delta_{t-1} \leq \mathbf{A} + \mathbf{B},$$

where

$$\begin{aligned} \mathbf{A} &= f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1})) - f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)), \\ \mathbf{B} &= f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}). \end{aligned}$$

Since $f(\cdot, \mathbf{y})$ is L -Lipschitz for each $\mathbf{y} \in \mathcal{Y}$, we have $\|\mathbf{g}\| \leq L$ for all $\mathbf{g} \in \partial_{\mathbf{x}}f(\mathbf{x}, \mathbf{y})$ and each pair of \mathbf{x} and $\mathbf{y} \in \mathcal{Y}$. Applying the same argument used for proving Lemma 33 yields

$$\mathbf{A} \leq 2\eta_{\mathbf{x}}L^2(t-s-1).$$

It suffices to bound the term \mathbf{B} . Indeed, we have

$$(\mathbf{y} - \mathbf{y}_t)^\top \mathbf{g}_{\mathbf{y}}^{t-1} \leq \frac{1}{2\eta_{\mathbf{y}}^{t-1}} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y} - \mathbf{y}_t\|^2 - \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2), \quad \text{for all } \mathbf{y} \in \mathcal{Y}.$$

Since the function $f(\mathbf{x}_{t-1}, \cdot)$ is μ -strongly concave and $\mathbf{g}_{\mathbf{y}}^{t-1} \in \partial_{\mathbf{y}}f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we have

$$(\mathbf{y} - \mathbf{y}_{t-1})^\top \mathbf{g}_{\mathbf{y}}^{t-1} \geq f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + \frac{\mu}{2}\|\mathbf{y} - \mathbf{y}_{t-1}\|^2.$$

Since the function $f(\mathbf{x}_{t-1}, \cdot)$ is L -Lipschitz, we have $\|\mathbf{g}_{\mathbf{y}}^{t-1}\| \leq L$. Combining this fact with Young's inequality yields

$$(\mathbf{y}_{t-1} - \mathbf{y}_t)^\top \mathbf{g}_{\mathbf{y}}^{t-1} \geq -\frac{1}{2\eta_{\mathbf{y}}^{t-1}}\|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2 - \frac{1}{2}\eta_{\mathbf{y}}^{t-1}L^2.$$

Putting these pieces together yields

$$f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) \leq \frac{1}{2\eta_{\mathbf{y}}^{t-1}}(\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y} - \mathbf{y}_t\|^2) - \frac{\mu}{2}\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 + \frac{1}{2}\eta_{\mathbf{y}}^{t-1}L^2. \quad (\text{D.5})$$

Since $\mathbf{y}^*(\mathbf{x}_s) \in \mathcal{Y}$ for $s \leq t-1$, Eq. (D.5) implies

$$\mathbf{B} \leq \frac{1}{2\eta_{\mathbf{y}}^{t-1}}(\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y} - \mathbf{y}_t\|^2) - \frac{\mu}{2}\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 + \frac{1}{2}\eta_{\mathbf{y}}^{t-1}L^2.$$

Putting these pieces together yields the inequality for Algorithm 1.

We then prove the results for Algorithm 2. By the definition of Δ_{t-1} , we have

$$\Delta_{t-1} \leq \mathbf{A} + \mathbf{B},$$

where

$$\begin{aligned}\mathbf{A} &= \mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_{t-1})) - f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s))], \\ \mathbf{B} &= \mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}^*(\mathbf{x}_s)) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})].\end{aligned}$$

Applying the same argument used for proving Lemma 33 yields

$$\mathbf{A} \leq 2\eta_{\mathbf{x}}(L^2 + \sigma^2)(t - s - 1).$$

It suffices to bound the term \mathbf{B} . Indeed, we have

$$(\mathbf{y} - \mathbf{y}_t)^\top \hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \leq \frac{1}{2\eta_{\mathbf{y}}^{t-1}} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \|\mathbf{y} - \mathbf{y}_t\|^2 - \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2), \quad \text{for all } \mathbf{y} \in \mathcal{Y}.$$

Using Young's inequality, we have

$$\begin{aligned}(\mathbf{y} - \mathbf{y}_t)^\top \hat{\mathbf{g}}_{\mathbf{y}}^{t-1} &\geq (\mathbf{y} - \mathbf{y}_{t-1})^\top \hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \\ &\quad - \frac{1}{2\eta_{\mathbf{y}}^{t-1}} \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2 - \eta_{\mathbf{y}}^{t-1} \|\mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}]\|^2 - \eta_{\mathbf{y}}^{t-1} \|\hat{\mathbf{g}}_{\mathbf{y}}^{t-1} - \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}]\|^2.\end{aligned}$$

Conditioned on the iterate $(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$, we take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of samples $\{\xi_i^k\}_{i=1}^M$) and use Lemma 28 with $M = 1$ to yield

$$\begin{aligned}(\mathbf{y} - \mathbf{y}_{t-1})^\top \mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] \\ \leq \frac{1}{2\eta_{\mathbf{y}}^{t-1}} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \mathbb{E}[\|\mathbf{y} - \mathbf{y}_t\|^2 \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}]) + \eta_{\mathbf{y}}^{t-1} \|\mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}]\|^2 + \eta_{\mathbf{y}}^{t-1} \sigma^2.\end{aligned}$$

Note that $\mathbb{E}[\hat{\mathbf{g}}_{\mathbf{y}}^{t-1} \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}] \in \partial_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$ (cf. Lemma 28). Thus, by using the same argument used for analyzing Algorithm 1, we have

$$\begin{aligned}f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) &\leq \frac{1}{2\eta_{\mathbf{y}}^{t-1}} (\|\mathbf{y} - \mathbf{y}_{t-1}\|^2 - \mathbb{E}[\|\mathbf{y} - \mathbf{y}_t\|^2 \mid \mathbf{x}_{t-1}, \mathbf{y}_{t-1}]) \\ &\quad - \frac{\mu}{2} \|\mathbf{y} - \mathbf{y}_{t-1}\|^2 + \eta_{\mathbf{y}}^{t-1} (L^2 + \sigma^2).\end{aligned}$$

We take the expectation of both sides of the above inequality (where the expectation is taken over the randomness in the selection of all previous samples). This yields

$$\mathbb{E}[f(\mathbf{x}_{t-1}, \mathbf{y}) - f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})] \leq \left(\frac{1}{2\eta_{\mathbf{y}}^{t-1}} - \frac{\mu}{2}\right) \mathbb{E}[\|\mathbf{y} - \mathbf{y}_{t-1}\|^2] - \frac{1}{2\eta_{\mathbf{y}}^{t-1}} \mathbb{E}[\|\mathbf{y} - \mathbf{y}_t\|^2] + \eta_{\mathbf{y}}^{t-1} (L^2 + \sigma^2). \quad (\text{D.6})$$

Since $\mathbf{y}^*(\mathbf{x}_s) \in \mathcal{Y}$ for $s \leq t - 1$, Eq. (D.6) implies

$$\mathbf{B} \leq \left(\frac{1}{2\eta_{\mathbf{y}}^{t-1}} - \frac{\mu}{2}\right) \mathbb{E}[\|\mathbf{y}^*(\mathbf{x}_s) - \mathbf{y}_{t-1}\|^2] - \frac{1}{2\eta_{\mathbf{y}}^{t-1}} \mathbb{E}[\|\mathbf{y}^*(\mathbf{x}_s) - \mathbf{y}_t\|^2] + \eta_{\mathbf{y}}^{t-1} (L^2 + \sigma^2).$$

Putting these pieces together yields the inequality for Algorithm 2. \square

Without loss of generality, we let $B \leq T + 1$ satisfy that $\frac{T+1}{B}$ is an integer. The following lemma bounds $\frac{1}{T+1} (\sum_{t=0}^T \Delta_t)$ for Algorithm 1 and 2 using the same localization technique.

Lemma 37. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}} B L^2 + \frac{L^2(1+\log(B))}{2\mu B}.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}} B (L^2 + \sigma^2) + \frac{(L^2 + \sigma^2)(1+\log(B))}{2\mu B}.$$

Proof. Applying the same argument used for proving Lemma 34 but with the two inequalities in Lemma 36, we can obtain the inequalities for Algorithm 1 and 2. \square

Proof of Theorem 24. We first prove the results for Algorithm 1. Applying the same argument used for proving Theorem 19 but with the first inequalities in Lemmas 35 and 37, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 8\rho L^2 \left(\eta_{\mathbf{x}} B + \frac{1+\log(B)}{2\mu B} \right) + 4\eta_{\mathbf{x}} \rho L^2.$$

Letting $B = \lfloor \sqrt{\frac{1}{\mu\eta_{\mathbf{x}}}} \rfloor + 1$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 16\rho L^2 \sqrt{\frac{\eta_{\mathbf{x}}}{\mu}} + 4\rho L^2 \sqrt{\frac{\eta_{\mathbf{x}}}{\mu}} \log \left(1 + \frac{1}{\mu\eta_{\mathbf{x}}} \right) + 12\eta_{\mathbf{x}} \rho L^2.$$

By the definition of $\eta_{\mathbf{x}}$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{3\epsilon^2}{4}.$$

This implies that the number of gradient evaluations required by Algorithm 1 to return an ϵ -stationary point is

$$O \left(\frac{\rho L^2 \Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho L^2}{\mu \epsilon^2}, \frac{\rho L^2}{\mu \epsilon^2} \log^2 \left(1 + \frac{\rho^2 L^4}{\mu^2 \epsilon^4} \right) \right\} \right).$$

We then prove the results for Algorithm 2. Applying the same argument used for analyzing Algorithm 1 but with the second inequalities in Lemmas 35 and 37 and using $B = \lfloor \sqrt{\frac{1}{\mu\eta_{\mathbf{x}}}} \rfloor + 1$ and the definition of $\eta_{\mathbf{x}}$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2] \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{3\epsilon^2}{4}.$$

This implies that the number of stochastic gradient evaluations required by Algorithm 2 to return an ϵ -stationary point is

$$O \left(\frac{\rho(L^2 + \sigma^2)\Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho(L^2 + \sigma^2)}{\mu \epsilon^2}, \frac{\rho(L^2 + \sigma^2)}{\mu \epsilon^2} \log^2 \left(1 + \frac{\rho^2(L^2 + \sigma^2)^2}{\mu^2 \epsilon^4} \right) \right\} \right).$$

This completes the proof.

Appendix E. Nonsmooth and Nonconvex-Concave Setting

We provide the proof for Theorem 26. Throughout this subsection, we set

$$\eta_{\mathbf{x}}^t \equiv \eta_{\mathbf{x}} = \min \left\{ \frac{\epsilon^2}{48\rho L^2}, \frac{\epsilon^6}{65536\rho^3 L^4 D^2} \right\}, \quad \eta_{\mathbf{y}}^t \equiv \eta_{\mathbf{y}} = \frac{\epsilon^2}{16\rho L^2},$$

for Algorithm 1, and

$$\eta_{\mathbf{x}}^t \equiv \eta_{\mathbf{x}} = \min \left\{ \frac{\epsilon^2}{48\rho(L^2 + \sigma^2)}, \frac{\epsilon^6}{131072\rho^3(L^2 + \sigma^2)^2 D^2} \right\}, \quad \eta_{\mathbf{y}}^t \equiv \eta_{\mathbf{y}} = \frac{\epsilon^2}{32\rho(L^2 + \sigma^2)},$$

for Algorithm 2. We define $\Delta_t = \Phi(\mathbf{x}_t) - f(\mathbf{x}_t, \mathbf{y}_t)$ and $\bar{\Delta}_t = \mathbb{E}[\Phi(\mathbf{x}_t) - f(\mathbf{x}_t, \mathbf{y}_t)]$ for the iterates generated by Algorithms 1 and 2 respectively.

Lemma 38. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy for $\forall s \leq t - 1$ that*

$$\Delta_{t-1} \leq 2\eta_{\mathbf{x}} L^2 (t - s - 1) + \frac{1}{2\eta_{\mathbf{y}}} (\|\mathbf{y}_{t-1} - \mathbf{y}^*(\mathbf{x}_s)\|^2 - \|\mathbf{y}_t - \mathbf{y}^*(\mathbf{x}_s)\|^2) + \frac{1}{2}\eta_{\mathbf{y}} L^2.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy for $\forall s \leq t - 1$ that

$$\Delta_{t-1} \leq 2\eta_{\mathbf{x}}(L^2 + \sigma^2)(t - s - 1) + \frac{1}{2\eta_{\mathbf{y}}} (\mathbb{E}[\|\mathbf{y}_{t-1} - \mathbf{y}^*(\mathbf{x}_s)\|^2] - \mathbb{E}[\|\mathbf{y}_t - \mathbf{y}^*(\mathbf{x}_s)\|^2]) + \eta_{\mathbf{y}}(L^2 + \sigma^2).$$

Proof. The proof is the same as that of Lemma 36 with $\mu = 0$. \square

Lemma 39. *The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 1 satisfy*

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}} B L^2 + \frac{D^2}{2B\eta_{\mathbf{y}}} + \frac{\eta_{\mathbf{y}} L^2}{2}.$$

The iterates $\{\mathbf{x}_t\}_{t \geq 1}$ generated by Algorithm 2 satisfy

$$\frac{1}{T+1} \left(\sum_{t=0}^T \Delta_t \right) \leq \eta_{\mathbf{x}} B (L^2 + \sigma^2) + \frac{D^2}{2B\eta_{\mathbf{y}}} + \eta_{\mathbf{y}} (L^2 + \sigma^2).$$

Proof. Applying the same argument used for proving Lemma 34 but with the two inequalities in Lemma 38, we can obtain the inequalities for Algorithm 1 and 2. \square

Proof of Theorem 26. We first prove the results for Algorithm 1. Applying the same argument used for proving Theorem 19 but with the first inequalities in Lemmas 35 and 39, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 8\rho(\eta_{\mathbf{x}} B L^2 + \frac{D^2}{2B\eta_{\mathbf{y}}} + \frac{\eta_{\mathbf{y}} L^2}{2}) + 4\eta_{\mathbf{x}} \rho L^2.$$

Letting $B = \left\lfloor \frac{D}{2L} \sqrt{\frac{1}{\eta_{\mathbf{x}} \eta_{\mathbf{y}}}} \right\rfloor + 1$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 16\rho L D \sqrt{\frac{\eta_{\mathbf{x}}}{\eta_{\mathbf{y}}}} + 4\eta_{\mathbf{y}} \rho L^2 + 12\eta_{\mathbf{x}} \rho L^2.$$

By the definition of $\eta_{\mathbf{x}}$ and $\eta_{\mathbf{y}}$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2 \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{3\epsilon^2}{4}.$$

This implies that the number of gradient evaluations required by Algorithm 1 to return an ϵ -stationary point is

$$O \left(\frac{\rho L^2 \Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho^2 L^2 D^2}{\epsilon^4} \right\} \right).$$

We now prove the results for Algorithm 2. Applying the same argument used for analyzing Algorithm 1 but with the second inequalities in Lemmas 35 and 39, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2] \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 8\rho(\eta_{\mathbf{x}}B(L^2 + \sigma^2) + \frac{D^2}{2B\eta_{\mathbf{y}}} + \eta_{\mathbf{y}}(L^2 + \sigma^2)) + 4\eta_{\mathbf{x}}\rho(L^2 + \sigma^2).$$

Letting $B = \left\lfloor \frac{D}{2} \sqrt{\frac{1}{\eta_{\mathbf{x}}\eta_{\mathbf{y}}(L^2 + \sigma^2)}} \right\rfloor + 1$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2] \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + 16\rho D \sqrt{\frac{\eta_{\mathbf{x}}(L^2 + \sigma^2)}{\eta_{\mathbf{y}}}} + 8\eta_{\mathbf{y}}\rho(L^2 + \sigma^2) + 12\eta_{\mathbf{x}}\rho(L^2 + \sigma^2).$$

By the definition of $\eta_{\mathbf{x}}$ and $\eta_{\mathbf{y}}$, we have

$$\frac{1}{T+1} \left(\sum_{t=0}^T \mathbb{E}[\|\nabla \Phi_{1/2\rho}(\mathbf{x}_t)\|^2] \right) \leq \frac{4\Delta_{\Phi}}{\eta_{\mathbf{x}}(T+1)} + \frac{3\epsilon^2}{4}.$$

This implies that the number of stochastic gradient evaluations required by Algorithm 2 to return an ϵ -stationary point is

$$O \left(\frac{\rho(L^2 + \sigma^2)\Delta_{\Phi}}{\epsilon^4} \max \left\{ 1, \frac{\rho^2(L^2 + \sigma^2)D^2}{\epsilon^4} \right\} \right).$$

This completes the proof.