

# ALTERNATING PROXIMAL-GRADIENT STEPS FOR (STOCHASTIC) NONCONVEX-CONCAVE MINIMAX PROBLEMS\*

RADU IOAN BOT<sup>†</sup> AND AXEL BÖHM<sup>‡</sup>

**Abstract.** Minimax problems of the form  $\min_x \max_y \Psi(x, y)$  have attracted increased interest largely due to advances in machine learning, in particular generative adversarial networks and adversarial learning. These are typically trained using variants of stochastic gradient descent for the two players. Although convex-concave problems are well understood with many efficient solution methods to choose from, theoretical guarantees outside of this setting are sometimes lacking even for the simplest algorithms. In particular, this is the case for alternating gradient descent ascent, where the two agents take turns updating their strategies. To partially close this gap in the literature we prove a novel global convergence rate for the stochastic version of this method for finding a critical point of  $\psi(\cdot) := \max_y \Psi(\cdot, y)$  in a setting which is not convex-concave.

**Key words.** Minimax, Saddle point, Nonconvex-concave, Complexity, Prox-gradient method, Stochastic gradient descent

**AMS subject classifications.** 90C47, 90C15, 90C25

**1. Introduction.** We investigate the *alternating* variant of gradient descent ascent (GDA) with proximal steps for weakly convex-(strongly) concave saddle point problems, given by

$$(1.1) \quad \min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^n} \left\{ \Psi(x, y) := f(x) + \Phi(x, y) - h(y) \right\}$$

for a weakly convex-concave coupling function  $\Phi : \mathbb{R}^d \times \mathbb{R}^n \rightarrow \mathbb{R}$  and proper, convex and lower semicontinuous (l.s.c.) regularizers  $h$  and  $f$ , see Assumption 3, 5 and 7 for details.

Nonconvex-concave saddle point problems have received a great deal of attention recently due to their application in adversarial learning [47], learning with nondecomposable losses [15, 53], learning with uncertain data [9] and generative adversarial imitation learning of linear quadratic regulators [23]. Additionally, albeit typically resulting in nonconvex-nonconcave objectives, the large interest in *generative adversarial networks (GANs)* [19, 2] has led to the studying of saddle point problems under different simplifying assumptions [3, 16, 6, 10, 30].

In the nonconvex-concave setting *inner loop* methods have received much of the attention [40, 49, 30, 25, 44] with them obtaining the best complexity results in this class, see Table 1. Despite superior theoretical performance these methods have not been as popular in practice, especially in the training of GANs where *single loop* methods are still state-of-the-art [16, 3, 10, 19, 17, 23, 32]. The simplest approach is given by *simultaneous* GDA, which, for a smooth coupling function  $\Phi$  and stepsizes  $\eta_x, \eta_y > 0$ , reads as:

$$\text{(simultaneous)} \quad \begin{cases} x^+ = x - \eta_x \nabla_x \Phi(x, y) \\ y^+ = y + \eta_y \nabla_y \Phi(x, y). \end{cases}$$

---

\*Submitted to the editors DATE.

**Funding:** Research of the second author was supported by the doctoral programme *Vienna Graduate School on Computational Optimization (VGSCO)*, FWF (Austrian Science Fund), project W 1260.

<sup>†</sup>Faculty of Mathematics, University of Vienna ([radu.bot@univie.ac.at](mailto:radu.bot@univie.ac.at))

<sup>‡</sup>Faculty of Mathematics, University of Vienna ([axel.boehm@univie.ac.at](mailto:axel.boehm@univie.ac.at))

After the first step of this method, however, more information is already available, which can be used in the update of the second variable, resulting in

$$\text{(alternating)} \quad \begin{cases} x^+ = x - \eta_x \nabla_x \Phi(x, y) \\ y^+ = y + \eta_y \nabla_y \Phi(x^+, y). \end{cases}$$

It has been widely known that the alternating version of GDA has many favorable convergence properties of the simultaneous one [3, 17, 52, 54]. It has been long known that for bilinear problems the iterates of simultaneous GDA may diverge while those of the alternating version at least remain bounded. Furthermore, [17] showed that the alternating version can be made convergent for this simple setting if negative momentum is used, while the same is false for simultaneous GDA. In another special setting [54] was able to show better local dependence on the condition number for strongly convex-strongly-concave quadratic problems. We are naturally interested in — and will give an affirmative answer to the question:

**Does stochastic alternating GDA have nonasymptotic convergence guarantees for nonconvex minimax problems?**

This might seem surprising as it has been sufficiently demonstrated [18, 16, 36, 6] that both versions of GDA fail to converge for simple bilinear problems if equal stepsizes are used. We therefore want to point out the importance of the *two-time-scale* approach which was also emphasized in [30, 22]. However, this alone is also not enough as shown in [17]. The seeming contradiction is resolved through the observation that our convergence guarantees only concern the  $x$ -component of the objective function.

*Optimality.* For convex-concave minimax problems, the notion of solution is simple. We aim to find a so-called saddle point  $(x^*, y^*) \in \mathbb{R}^d \times \mathbb{R}^n$  satisfying

$$\Psi(x^*, y) \leq \Psi(x^*, y^*) \leq \Psi(x, y^*) \quad \forall (x, y) \in \mathbb{R}^d \times \mathbb{R}^n.$$

For convex-concave problems this is equivalent to the first order optimality condition

$$(1.2) \quad \begin{pmatrix} 0 \\ 0 \end{pmatrix} \in \begin{pmatrix} \nabla_x \Phi(x^*, y^*) \\ -\nabla_y \Phi(x^*, y^*) \end{pmatrix} + \begin{pmatrix} \partial f(x^*) \\ \partial h(y^*) \end{pmatrix}.$$

Similarly to the nonconvex single objective optimization where one cannot expect to find global minima, if the minimax problem is not convex-concave the notion of saddle point is too strong. So one natural approach is to focus on conditions such as (1.2), as done in [52, 43, 33]. However, treating the two components in such a symmetric fashion might not seem fitting since in contrast to the convex-concave problem  $\min_x \max_y \neq \max_y \min_x$ . Instead we will focus, in the spirit of [30, 49, 46], on the stationarity of what we will refer to as the *max function* given by

$$(1.3) \quad \varphi(x) := \max_{y \in \mathbb{R}^n} \Phi(x, y) - h(y), \quad \text{where } \varphi : \mathbb{R}^d \rightarrow \mathbb{R}.$$

This makes sense from the point of view of many practical applications. Problems arising from adversarial learning can be formulated as minimax, but typically only  $x$ , which corresponds to the classifier is relevant as  $y$  is adversarial noise. Similarly, for GANs, one is typically only interested in the generator and not the discriminator. See Table 1 for a comparison of other methods using the same notion of optimality. Note that it is possible to move from one notion of optimality to the other [30], but as both directions are typically associated with additional computational effort a comparison is not trivial and out of scope of this work.

Table 1: The gradient complexity of algorithms for nonconvex-(strongly) concave minimax problems for computing  $\epsilon$ -stationary points of the max function.  $\kappa > 0$  is the condition number. The notation  $\tilde{\mathcal{O}}$  hides logarithmic terms.

	Noncvx-Strongly Concave		Noncvx-Concave		single
	det.	stoch.	det.	stoch.	loop
[46]	$\tilde{\mathcal{O}}(\kappa^2\epsilon^{-2})$	$\tilde{\mathcal{O}}(\kappa^3\epsilon^{-4})$	$\tilde{\mathcal{O}}(\epsilon^{-6})$	$\tilde{\mathcal{O}}(\epsilon^{-6})$	$\times$
[55, 49]	–	–	$\tilde{\mathcal{O}}(\epsilon^{-3})$	–	$\times$
[29, 44]	$\tilde{\mathcal{O}}(\sqrt{\kappa}\epsilon^{-2})$	–	$\tilde{\mathcal{O}}(\epsilon^{-3})$	–	$\times$
[30]	$\mathcal{O}(\kappa^2\epsilon^{-2})$	$\mathcal{O}(\kappa^3\epsilon^{-4})$	$\mathcal{O}(\epsilon^{-6})$	$\mathcal{O}(\epsilon^{-8})^*$	$\checkmark$
this work	$\mathcal{O}(\kappa^2\epsilon^{-2})$	$\mathcal{O}(\kappa^3\epsilon^{-4})$	$\mathcal{O}(\epsilon^{-6})$	$\mathcal{O}(\epsilon^{-8})^*$	$\checkmark$

*Contributions.* We prove novel convergence rates for *alternating* prox-gradient descent ascent for nonconvex-(strongly) concave minimax problems in a deterministic and stochastic setting. For deterministic problems, [52] has proved convergence rates for alternating GDA in terms of the criticality of  $\Phi$  while we use the max function  $\varphi$ , see (1.3), instead. Our results are also more general than e.g. [29, 55, 30] in the sense that they require  $\Phi$  to be smooth in the first component whereas we only require weak convexity, similar to [46]. Furthermore, we allow for our method to include possibly nonsmooth regularizers, similar to [55, 46], by passing from a regular projected-gradient to more the more general *proximal-gradient* steps which captures and extends the common constraint setting, necessitating us to prove a more general version of Danskins theorem in the process.

*Roadmap.* In the remainder of this section we discuss related literature and some real-world applications resulting in nonconvex-concave problems. In Section 2 we discuss the mathematical preliminaries as well as our main assumptions about the involved functions. Section 3 and Section 4 are devoted to the setting where the objective function is assumed to be convex and strongly convex, respectively. Both times we treat the deterministic problem first and then the scenario where we are only given a stochastic gradient oracle. Finally, in Section 5 we discussed numerical experiments in adversarial learning. For the interested reader we highlighted the improvements in the analysis of alternating GDA over its simultaneous counterpart in Sections 3.5 and 4.4.

**1.1. Related literature.** For the purpose of this paper we separate the quantitative study of minimax problems into the following domains.

*Convex-concave.* For convex-concave problems historically the *extra-gradient* and the *forward-backward-forward* method have been known to converge. For the former even a rate of  $\mathcal{O}(\epsilon^{-1})$  has been proven in [41] under the name of *mirror-prox*. Both of these methods suffer from the drawback of requiring two gradient evaluations per iteration. This has led to the development of methods such as *optimistic GDA* [10, 11] or [21, 6, 16, 35] which use past gradients to reduce the need of gradient evaluations to one per iteration. In all of these cases, however, convergence guarantees typically do not go beyond the convex-concave setting. Nevertheless, these methods have been employed successfully in the GAN setting [16, 6, 10].

\* using large batchsizes of order  $\mathcal{O}(\epsilon^{-2})$

*Nonconvex-concave with inner loops.* Approximating the max function by running multiple iterations of a solver on the second component or convexifying the problem by adding a quadratic term and then solving the convex-concave problem constitute natural approaches [49, 29, 55, 46, 43]. Such methods achieve the best known rates [49, 29, 55, 44] in this class. However, they are usually quite involved and have for the most part not been used in deep learning applications.

*Nonconvex-concave with single loop.* While these methods have received some attention in the training of GANs [10, 16, 6] most of the theoretical statement are for convex-concave problems. In the nonconvex setting only two methods have been studied. Previous research, see [30, 33], has focused on the *simultaneous* version of the gradient descent ascent algorithm where both components are updated at the same time. The only other work which focuses on *alternating* GDA is [52]. Their results are in terms of stationarity of  $\Phi$  and they do not treat the stochastic case. Note that our work is most similar to [30] where the same notion of optimality is used and similar rates to our are obtained for *simultaneous* GDA.

*Others.* Clearly the above categories do not cover the entire field. However, other settings have not received as much attention. Only [52] treats (strongly) convex-nonconcave problems and proves convergence rates similar to the nonconvex-(strongly) concave setting. In [51] a special stochastic nonconvex-linear problem with regularizers is solved via a variance reduced single loop method with a significantly improved rate over the general nonconvex-concave problem.

The most general setting out of all the aforementioned ones is discussed in [31, 32, 48], namely the weakly convex-weakly concave setting. They use however, a weaker notion of optimality related to the Minty variational inequality formulation. We also only mentioned (sub)gradient methods, but the restrictive assumption that the proximal operator of a component can be evaluated has been considered as well [25].

## 1.2. Nonconvex-concave applications.

**1.2.1. Adversarial learning.** Such problems often use an attack model [34] that allows for every pixel to be perturbed up to given threshold  $\epsilon$ :

$$\min_{\theta} \max_{\|z - z_0\|_{\infty} \leq \epsilon} \ell(\theta, z),$$

where  $z_0$  denotes the “true” training examples, and  $z$  the adversarial attack. However, this typically leads to nonconvex-nonconcave formulation. So [40] proposed a distributionally robust model, making use of the Wasserstein distance  $W$

$$\min_{\theta} \max_{P: W(P, P_0) \leq \rho} \mathbb{E}_P[\ell(\theta, Z)],$$

where  $Z \sim P_0$ , which they reformulated via a Lagrangian penalty approach to

$$(1.4) \quad \min_{\theta} \max_z \ell(\theta, z) - \gamma \|z - z_0\|^2.$$

While a larger  $\gamma$  corresponds to smaller robustness  $\rho$ , the model can be made nonconvex-strongly-concave if it is set big enough.

**1.2.2. Generative adversarial imitation learning of linear quadratic regulators.** In imitation learning the objective is to learn from an expert’s demonstration of performing a given task. In this case the minimization is performed over the policies with the goal of reducing the discrepancy between the reward of the expert’s policy

and the proposed one. The maximization is over the parameters of the reward function, see [23]. If the underlying dynamic and the reward function come from a linear quadratic regulator, see [8], this can be expressed as a nonconvex-strongly-concave minimax problem

$$(1.5) \quad \min_K \max_{\theta} m(K, \theta),$$

where  $K$  represents the choice of policy and  $\theta$  the parameters of the dynamic and reward functions.

**1.2.3. Fair learning.** The work [38] observed that a logistic regression model trained on the Fashion-MNIST dataset (comprised of  $n = 10$  classes) can lead to a bias against certain classes. In order to remove this bias, they proposed to minimize the maximal loss of the different categories, i.e.

$$(1.6) \quad \min_{\theta} \max_{1 \leq i \leq n} \ell_i(\theta),$$

where  $\ell_i$  denotes the loss incurred by all examples of class  $i$ . A similar approach was taken in [43], but for a more sophisticated CNN model. For the purpose of implementation (1.6) can be rewritten as

$$(1.7) \quad \min_{\theta} \max_{(t_1, \dots, t_n) \in \Delta} \sum_{i=1}^n t_i \ell_i(\theta)$$

where  $\Delta := \{(t_1, \dots, t_n) : t_i \geq 0, \sum_{i=1}^n t_i = 1\}$  denotes the unit simplex. Due to the linearity of (1.7) in the second variable  $(t_1, \dots, t_n)$ , the inner maximization problem is in particular concave.

**2. Preliminaries.** As mentioned in the earlier we will consider optimality in terms of the max function for any  $x \in \mathbb{R}^d$  given by  $\varphi(x) := \max_{y \in \mathbb{R}^n} \Phi(x, y) - h(y)$ . Similarly, we also need the regularized max function

$$\psi := \varphi + f, \quad \text{where } \psi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}.$$

In the remainder of the section we will focus on the necessary preliminaries connected to the weak convexity of the max function in the nonconvex-concave setting, see Section 3.

**2.1. Weak convexity.** In the nonconvex-concave setting of Section 3 the max function  $\varphi$  will in general be nonsmooth, which makes it nonobvious how to define near stationarity. The max function  $\varphi$  will, however, turn out to be *weakly convex*, see Proposition 3.1. For some  $\rho \geq 0$ , we say that

$$\psi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\} \text{ is } \rho\text{-weakly convex if } \psi + (\rho/2)\|\cdot\|^2 \text{ is convex.}$$

An example of a weakly convex function is one which is differentiable and the gradient is uniformly Lipschitz continuous with constant  $L$  (we call such a function  $L$ -smooth). In this case, the weak convexity parameter  $\rho$  is given by the Lipschitz constant.

Following [12, 49, 14, 7], we make use of a smooth approximation of  $\psi$  known as the *Moreau envelope*  $\psi_{\lambda}$ , parametrized  $\lambda > 0$ . For a proper,  $\rho$ -weakly convex and l.s.c. function  $\psi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ , the Moreau envelope of  $\psi$  with the parameter  $\lambda \in (0, \rho^{-1})$  is the function from  $\mathbb{R}^d$  to  $\mathbb{R}$  defined by

$$\psi_{\lambda}(x) := \inf_{z \in \mathbb{R}^d} \left\{ \psi(z) + \frac{1}{2\lambda} \|z - x\|^2 \right\}.$$

The proximal operator of the function  $\lambda\psi$  is the arg min of the right-hand side in this definition, that is,

$$(2.1) \quad \text{prox}_{\lambda\psi}(x) := \arg \min_{z \in \mathbb{R}^d} \left\{ \psi(z) + \frac{1}{2\lambda} \|z - x\|^2 \right\}.$$

Note that  $\text{prox}_{\lambda\psi}(x)$  is uniquely defined by (2.1) because the function being minimized is proper, l.s.c. and strongly convex. For a function  $\psi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$  and a point  $\bar{x}$  such that  $\psi(\bar{x})$  is finite, the *Fréchet subdifferential* of  $\psi$  at  $\bar{x}$ , denoted by  $\partial\psi(\bar{x})$ , is the set of all vectors  $v \in \mathbb{R}^d$  such that

$$\psi(x) \geq \psi(\bar{x}) + \langle v, x - \bar{x} \rangle + o(\|x - \bar{x}\|) \quad \text{as } x \rightarrow \bar{x}.$$

For weakly convex function the Fréchet subdifferential can simply be expressed in terms of the convex subdifferential of the (convex) function  $\psi + (\rho/2)\|\cdot\|^2$ .

While the next result is standard for the gradient and convex subgradients we explicitly mention the general case.

LEMMA 2.1 (see [39, Theorem 3.52]). *For an  $L_\psi$ -Lipschitz continuous function  $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$  every Fréchet subgradient is bounded in norm by  $L_\psi$ .*

Now, we provide a useful characterization of the gradient of the Moreau envelope.

LEMMA 2.2 (see [13, Lemma 2.2]). *Let  $\psi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$  be a proper,  $\rho$ -weakly convex, and l.s.c. function, and let  $\lambda \in (0, \rho^{-1})$ . Then the Moreau envelope  $\psi_\lambda$  is continuously differentiable on  $\mathbb{R}^d$  with gradient*

$$\nabla\psi_\lambda(x) = \frac{1}{\lambda} (x - \text{prox}_{\lambda\psi}(x)) \quad \text{for all } x \in \mathbb{R}^d,$$

and this gradient is Lipschitz continuous.

In particular, a gradient step with respect to the Moreau envelope corresponds to a proximal step, that is,

$$(2.2) \quad x - \lambda\nabla\psi_\lambda(x) = \text{prox}_{\lambda\psi}(x) \quad \text{for all } x \in \mathbb{R}^d.$$

*Stationarity.* The Moreau envelope allows us to naturally define a notion of *near stationarity* even for nonsmooth and  $\rho$ -weakly convex functions. We say that for an  $\epsilon > 0$  and a  $\lambda \in (0, \rho^{-1})$

$$(2.3) \quad \text{a point } x \text{ is } \epsilon\text{-stationary for } \psi \text{ if } \|\nabla\psi_\lambda(x)\| \leq \epsilon.$$

Canonically, we call a point stationary if the above holds for  $\epsilon = 0$ . This notion of near stationarity can also be expressed in terms the original function  $\psi$ .

LEMMA 2.3. *Let  $x$  be  $\epsilon$ -stationary for the proper,  $\rho$ -weakly convex and l.s.c. function  $\psi$ , i.e.  $\|\nabla\psi_\lambda(x)\| \leq \epsilon$  with  $\lambda \in (0, \rho^{-1})$ . Then there exist a point  $\hat{x}$  such that  $\|x - \hat{x}\| \leq \epsilon\lambda$  and  $\text{dist}(0, \partial\psi(\hat{x})) \leq \epsilon$ .*

*Proof.* From the definition of the Moreau envelope, we have that

$$0 \in \partial\psi(\text{prox}_{\lambda\psi}(x)) + \frac{1}{\lambda}(\text{prox}_{\lambda\psi}(x) - x),$$

from which  $\nabla\psi_\lambda(x) \in \partial\psi(\text{prox}_{\lambda\psi}(x))$  follows by using (2.2). It is easy to see that  $\hat{x} = \text{prox}_{\lambda\psi}(x)$  fulfills the required conditions.  $\square$

**2.2. About the stochastic setting.** We discuss the stochastic version of problem (1.1) where the coupling function  $\Phi$  is actually given as an expectation,

$$\Phi(x, y) = \mathbb{E}_{\xi \sim \mathcal{D}} [\Phi(x, y; \xi)] \quad \forall (x, y) \in \mathbb{R}^d \times \mathbb{R}^n$$

and we can only access independent samples of the gradient  $\nabla_x \Phi(x, y; \xi)$  (or subgradient) and  $\nabla_y \Phi(x, y; \zeta)$ , where  $\xi$  and  $\zeta$  are drawn from the (in general unknown) distribution  $\mathcal{D}$ .

We require the following standard assumption with respect to these stochastic gradient estimators.

ASSUMPTION 1 (unbiased). *The stochastic gradient estimator is unbiased, i.e.*

$$\mathbb{E}[\nabla \Phi(x, y; \xi)] = \nabla \Phi(x, y) \quad \forall (x, y) \in \mathbb{R}^d \times \mathbb{R}^n,$$

or in the case of subgradients

$$\mathbb{E}[g^\xi] \in \partial[\Phi(\cdot, y)](x), \quad \text{where } g^\xi \in \partial[\Phi(\cdot, y; \xi)](x).$$

ASSUMPTION 2 (bounded variance). *The variance of the estimator is uniformly bounded, i.e. for all  $(x, y) \in \mathbb{R}^d \times \mathbb{R}^n$  and a variance  $\sigma^2 \geq 0$  we have*

$$(2.4) \quad \mathbb{E}[\|\nabla \Phi(x, y; \xi) - \nabla \Phi(x, y)\|^2] \leq \sigma^2.$$

In the setting of Section 3 where  $\Phi$  is not necessarily smooth in the first component, we make the analogous assumption for subgradients, i.e.

$$(2.5) \quad \mathbb{E}[\|g^\xi - \mathbb{E}[g^\xi]\|^2] \leq \sigma^2,$$

for a stochastic subgradient  $g^\xi \in \partial[\Phi(\cdot, y; \xi)](x)$ .

**2.3. The algorithm.** Since we cover different settings such as smooth or not, deterministic and stochastic we try to formulate a unifying scheme.

ALGORITHM 1 (proximal alternating GDA). *Let  $(x_0, y_0) \in \mathbb{R}^d \times \mathbb{R}^n$  and stepsizes  $\eta_x, \eta_y > 0$ . Consider the following iterative scheme*

$$(\forall k \geq 0) \quad \begin{cases} x_{k+1} = \text{prox}_{\eta_x f}(x_k - \eta_x G_x(x_k, y_k)) \\ y_{k+1} = \text{prox}_{\eta_y h}(y_k + \eta_y G_y(x_{k+1}, y_k)), \end{cases}$$

where  $G_x$  and  $G_y$  will be replaced by the appropriate (sub)gradient and its estimator in the deterministic and stochastic setting, respectively.

**2.4. Notation.** We collect different symbols used through this manuscript.

Object	Definition
Coupling function	$\Phi(x, y)$
Objective function	$\Psi(x, y) := f(x) + \Phi(x, y) - h(y)$
Regularized coupling function	$\Gamma(x, y) := \Phi(x, y) - h(y)$
Max function	$\varphi(x) := \max_y \Phi(x, y) - h(y)$
Regularized max function	$\psi(x) := f(x) + \varphi(x)$

Note that the regularized coupling function  $\Gamma$  is only needed in proofs and some technical lemmata. The remaining functions confirm to the logic that small letters denote functions maximized in the second component (and thus only depend on  $x$ ). On the other hand (no matter if capital or not) the letter *psi* indicates the presence of regularizers and *phi* their absence.

**3. Nonconvex-concave objective.** In this section we treat the case where the objective function is weakly convex and Lipschitz in  $x$ , but not necessarily smooth, and concave and smooth in  $y$ . This will result in a weakly convex and Lipschitz max function whose Moreau envelope we will study for criticality, see (2.3).

**3.1. Assumptions.** While the first assumption concerns general setting of this section, i.e. weakly convex-concave, the latter ones are more of a technical nature.

ASSUMPTION 3. *The coupling function  $\Phi$  is*

(i) *concave and  $L_{\nabla\Phi}$ -smooth in the second component uniformly in  $x$ ,*

$$\|\nabla_y\Phi(x, y) - \nabla_y\Phi(x, y')\| \leq L_{\nabla\Phi}\|y - y'\| \quad \forall x \in \mathbb{R}^d \forall y, y' \in \mathbb{R}^n.$$

(ii)  *$\rho$ -weakly convex in the first component uniformly in the second one, i.e.*

$$\Phi(\cdot, y) + \frac{\rho}{2}\|\cdot\|^2 \quad \text{is convex for all } y \in \mathbb{R}^n.$$

Assumption 3 is fulfilled if e.g.  $\Phi$  is  $L_{\nabla\Phi}$ -smooth jointly in both components, i.e.

$$\|\nabla\Phi(x, y) - \nabla\Phi(x', y')\| \leq L_{\nabla\Phi}\|(x, y) - (x', y')\| \quad \forall x, x' \in \mathbb{R}^d \forall y, y' \in \mathbb{R}^n,$$

in which case (ii) holds with  $\rho = L_{\nabla\Phi}$ .

The next assumption is classical in nonconvex optimization.

ASSUMPTION 4. *The function  $\psi$  is lower bounded, i.e.  $\inf_{x \in \mathbb{R}^d} \psi(x) > -\infty$ .*

In Section 3 we will actually need to bound the Moreau envelope  $\psi_\lambda$ , but these two conditions are in fact equivalent as for all  $x \in \mathbb{R}^d$  and any  $\lambda \in (0, \rho^{-1})$

$$\psi(x) \geq \psi_\lambda(x) \geq \inf_{u \in \mathbb{R}^d} \psi(u).$$

We also want to point out that this is weaker than the lower boundedness of  $\Psi$ , which is usually required if stationary points of the type (1.2) are used, see for example [33].

ASSUMPTION 5.  *$\Phi$  is  $L$ -Lipschitz in the first component uniformly over  $\text{dom } h$  in the second one, i.e.*

$$\|\Phi(x, y) - \Phi(x', y)\| \leq L\|x - x'\| \quad \forall x, x' \in \mathbb{R}^d \forall y \in \text{dom } h.$$

ASSUMPTION 6. *The regularizers  $f$  and  $h$  are proper, l.s.c. and convex*

(i) *Additionally,  $f$  is either  $L_f$ -Lipschitz continuous on its domain, which is assumed to be open, or the indicator of a nonempty, convex and closed set. Either of those assumptions guarantees for any  $\gamma > 0$  the bound*

$$(3.1) \quad \|\text{prox}_{\gamma f}(x) - x\| \leq \gamma L_f \quad \forall x \in \text{dom } f.$$

(ii) *Furthermore,  $h$  has a bounded domain  $\text{dom } h$  such that the diameter of  $\text{dom } h$  is bounded by  $D_h$ .*



**3.2. Properties of the max function.** Previous research, when concluding the weak convexity of the max function, has relied on the compactness of the domain over which to maximize. This is done so that the classical Danskin Theorem can be applied. This assumption is e.g. fulfilled in the context of *Wasserstein GANs* [2] with weight clipping, but not in other formulations such as [20]. We provide an extension of the classical Danskin Theorem, which only relies on the concavity and l.s.c. of the objective in the second component and the boundedness of  $\text{dom } h$ , see Assumption 3 and 6. This implies that for every  $x \in \mathbb{R}^d$  the set

$$(3.2) \quad Y(x) := \left\{ y^* \in \mathbb{R}^n : \varphi(x) = \Phi(x, y^*) - h(y^*) = \max_{y \in \mathbb{R}^n} \{\Phi(x, y) - h(y)\} \right\}$$

is nonempty. For brevity we denote arbitrary elements of  $Y(x_k)$  by  $y_k^*$  for all  $k \geq 0$ .

PROPOSITION 3.1 (Subgradient characterization of the max function). *Let Assumption 3 and 6 hold true. Then, the function  $\varphi$ , see (1.3), fulfills*

$$\partial[\Phi(\cdot, y^*)](x) \subseteq \partial\varphi(x) \quad \forall y^* \in Y(x), \forall x \in \mathbb{R}^d.$$

In particular,  $\varphi$  is  $\rho$ -weakly convex.

*Proof.* From the  $\rho$ -weak convexity of  $\Phi(\cdot, y)$ , we have that  $\Phi(\cdot, y) + \frac{\rho}{2}\|\cdot\|^2$  is convex for all  $y \in \mathbb{R}^n$ . We define  $\tilde{\Phi}(x, y) = \Phi(x, y) + \frac{\rho}{2}\|x\|^2$  and  $\tilde{\Gamma}(x, y) = \Gamma(x, y) + \frac{\rho}{2}\|x\|^2$  for  $(x, y) \in \mathbb{R}^d \times \mathbb{R}^n$  as well as

$$\tilde{\varphi}(x) = \max_{y \in \mathbb{R}^n} \tilde{\Gamma}(x, y) = \varphi(x) + \frac{\rho}{2}\|x\|^2.$$

Notice that  $\tilde{\Gamma}(x, \cdot)$  is concave for any  $x \in \mathbb{R}^d$  and  $\tilde{\Gamma}(\cdot, y)$  is convex for any  $y \in \mathbb{R}^n$ . Thus, the function  $\tilde{\varphi}$  is convex and  $\text{dom } \tilde{\varphi} = \text{dom } \varphi = \mathbb{R}^d$ . Therefore  $\varphi$  is continuous, which implies that  $\partial\varphi(x) \neq \emptyset$  for any  $x \in \mathbb{R}^d$ . Let  $x \in \mathbb{R}^d$ ,  $y \in Y(x)$  and  $v \in \mathbb{R}^d$ . For any  $\alpha > 0$  it holds

$$\frac{\tilde{\varphi}(x + \alpha v) - \tilde{\varphi}(x)}{\alpha} \geq \frac{\tilde{\Gamma}(x + \alpha v, y) - \tilde{\Gamma}(x, y)}{\alpha} = \frac{\tilde{\Phi}(x + \alpha v, y) - \tilde{\Phi}(x, y)}{\alpha},$$

thus

$$(3.3) \quad \tilde{\varphi}'(x; v) = \inf_{\alpha > 0} \frac{\tilde{\varphi}(x + \alpha v) - \tilde{\varphi}(x)}{\alpha} \geq \inf_{\alpha > 0} \frac{\tilde{\Phi}(x + \alpha v, y) - \tilde{\Phi}(x, y)}{\alpha} = [\tilde{\Phi}(\cdot, y)]'(x; v),$$

where  $[\tilde{\Phi}(\cdot, y)]'(x; v)$  denotes the directional derivative of  $\tilde{\Phi}$  in the first component at  $x$  in the direction  $v$ . In conclusion,

$$(3.4) \quad \tilde{\varphi}'(x; v) \geq \sup_{y \in Y(x)} [\tilde{\Phi}(\cdot, y)]'(x; v) \quad \forall v \in \mathbb{R}^d$$

and for  $y \in Y(x)$  we therefore conclude  $\partial[\tilde{\Phi}(\cdot, y)](x) \subseteq \partial\tilde{\varphi}(x)$ . The first statement is obtained by subtracting  $\rho x$  on both sides of the inclusion.  $\square$

LEMMA 3.2 (Lipschitz continuity of the max function). *The Lipschitz continuity of  $\Phi$  in its first component implies that  $\varphi$  is Lipschitz with the same constant.*

*Proof.* Let  $x, x' \in \mathbb{R}^d$  and  $y^* \in Y(x)$ . On the one hand

$$\begin{aligned} \varphi(x) - \varphi(x') &= \Phi(x, y^*) - h(y^*) - \varphi(x') \\ &\leq \Phi(x, y^*) - h(y^*) - \Phi(x', y^*) + h(y^*) \leq L\|x - x'\|. \end{aligned}$$

The reverse direction  $\varphi(x') - \varphi(x) \leq L\|x - x'\|$  follows analogously.  $\square$

**3.3. Deterministic setting.** For initial values  $(x_0, y_0) \in \text{dom } f \times \text{dom } h$  the deterministic version of alternating GDA, for  $g_k \in \partial[\Phi(\cdot, y_k)](x_k)$ , reads as

$$(3.5) \quad (\forall k \geq 0) \begin{cases} x_{k+1} = \text{prox}_{\eta_x f}(x_k - \eta_x g_k) \\ y_{k+1} = \text{prox}_{\eta_y h}(y_k + \eta_y \nabla_y \Phi(x_{k+1}, y_k)). \end{cases}$$

**THEOREM 3.3.** *Let Assumption 3, 4, 5 and 6 hold true. For algorithm (3.5) with the stepsizes*

$$(3.6) \quad \eta_x = \min \left\{ \frac{\epsilon^4}{L_{\nabla\Phi} \rho^2 D_h^2 (L + L_f)^2}, \frac{\epsilon^2}{\rho L^2} \right\}, \eta_y = \frac{1}{L_{\nabla\Phi}} \text{ and } \lambda = \frac{1}{2\rho}$$

*the number of gradient evaluations  $K$  required is*

$$\mathcal{O} \left( \frac{\Delta^* \rho (L + L_f)^2}{\epsilon^4} \max \left\{ 1, \frac{\rho L_{\nabla\Phi} D_h^2}{\epsilon^2} \right\} + \frac{\rho \Delta_0}{\epsilon^2} \right),$$

*to visit an  $\epsilon$ -stationary given by  $\min_{1 \leq k \leq K} \|\nabla \psi_\lambda(x_k)\| \leq \epsilon$ , where  $\Delta^* := \psi(x_0) - \inf_{x \in \mathbb{R}^d} \psi(x)$  and  $\Delta_0 := \psi(x_0) - \Psi(x_0, y_0)$ .*

Similarly to the proofs in [12, 30] and others, the main descent statement makes use of the quantity  $\text{prox}_{\lambda\psi}(x_k)$  for a  $\lambda > 0$ . This is somewhat surprising as this point does not appear in the algorithm and can in general not be computed.

But first, we need to establish the fact that  $\hat{x}_k := \text{prox}_{\lambda\psi}(x_k)$  can also be written as the proximal operator of  $f$  evaluated at an auxiliary point.

**LEMMA 3.4.** *For any  $\lambda \in (0, \rho^{-1})$  and all  $k \geq 0$  the point  $\hat{x}_k := \text{prox}_{\lambda\psi}(x_k)$  can also be written for some  $v_k \in \partial\varphi(\hat{x}_k)$  as*

$$\hat{x}_k = \text{prox}_{\eta_x f}(\eta_x \lambda^{-1} x_k - \eta_x v_k + (1 - \eta_x \lambda^{-1}) \hat{x}_k).$$

*Proof.* Let  $k \geq 0$  be arbitrary but fixed and recall that  $g = f + \varphi$ . By the definition of  $\hat{x}_k$  we have that

$$0 \in \partial\psi(\hat{x}_k) + \frac{1}{\lambda}(\hat{x}_k - x_k) = \partial(\varphi + f)(\hat{x}_k) + \frac{1}{\lambda}(\hat{x}_k - x_k).$$

We can estimate through the continuity of  $\varphi$  and subdifferential calculus

$$\frac{1}{\lambda}(x_k - \hat{x}_k) \in \partial(\varphi + f)(\hat{x}_k) \subseteq \partial\varphi(\hat{x}_k) + \partial f(\hat{x}_k).$$

Thus, there exists  $v_k \in \partial\varphi(\hat{x}_k)$  such that  $\frac{1}{\lambda}(x_k - \hat{x}_k) \in v_k + \partial f(\hat{x}_k)$ . Also,  $\square$

$$\begin{aligned} \frac{1}{\lambda}(x_k - \hat{x}_k) \in \partial f(\hat{x}_k) + v_k &\Leftrightarrow \frac{\eta_x}{\lambda} x_k - \eta_x v_k + (1 - \frac{\eta_x}{\lambda}) \hat{x}_k \in \hat{x}_k + \eta_x \partial f(\hat{x}_k) \\ &\Leftrightarrow \hat{x}_k = \text{prox}_{\eta_x f} \left( \frac{\eta_x}{\lambda} x_k - \eta_x v_k + (1 - \frac{\eta_x}{\lambda}) \hat{x}_k \right). \end{aligned}$$

With the previous lemma in place we can now turn our attention to the first step of the actual convergence proof.

**LEMMA 3.5.** *With  $\lambda = 1/2\rho$  and  $\eta_x \geq 0$  we have for all  $k \geq 0$  that*

$$\psi_\lambda(x_{k+1}) \leq \psi_\lambda(x_k) + 2\rho\eta_x \Delta_k - \frac{1}{2}\eta_x \|\nabla \psi_\lambda(x_k)\|^2 + 4\rho\eta_x^2 L^2,$$

*where  $\Delta_k := \psi(x_k) - \Psi(x_k, y_k) \geq 0$ .*

*Proof.* Let  $k \geq 0$  be fixed. As before we denote  $\hat{x}_k = \text{prox}_{\lambda\psi}(x_k)$ . From the definition of the Moreau envelope we have that

$$(3.7) \quad \psi_\lambda(x_{k+1}) = \min_{x \in \mathbb{R}^d} \left\{ \psi(x) + \frac{1}{2\lambda} \|x - x_{k+1}\|^2 \right\} \leq \psi(\hat{x}_k) + \frac{1}{2\lambda} \|\hat{x}_k - x_{k+1}\|^2.$$

Let now  $v_k \in \partial\varphi(\hat{x}_k)$  as in Lemma 3.4. We successively deduce for  $\beta := 1 - \eta_x \lambda^{-1}$

$$(3.8) \quad \|\hat{x}_k - x_{k+1}\|^2 = \|\text{prox}_{\eta_x f}(\eta_x \lambda^{-1} x_k - \eta_x v_k + \beta \hat{x}_k) - \text{prox}_{\eta_x f}(x_k - \eta_x g_k)\|^2$$

$$(3.9) \quad \begin{aligned} &\leq \|\beta(\hat{x}_k - x_k) + \eta_x(g_k - v_k)\|^2 \\ &= \beta^2 \|\hat{x}_k - x_k\|^2 + 2\eta_x \beta \langle g_k - v_k, \hat{x}_k - x_k \rangle + \eta_x^2 \|g_k - v_k\|^2 \end{aligned}$$

$$(3.10) \quad \leq \beta^2 \|\hat{x}_k - x_k\|^2 + 2\eta_x \beta \langle g_k - v_k, \hat{x}_k - x_k \rangle + 4\eta_x^2 L^2$$

where (3.8) uses Lemma 3.4 and the definition of  $x_{k+1}$ , inequality (3.9) holds because of the nonexpansiveness of the proximal operator, and (3.10) follows from the Lipschitz continuity of  $\Phi$  and  $\varphi$  (see Lemma 3.2) and the fact that Lipschitz continuity implies bounded subgradients. We are left with estimating the inner product in the above inequality and we do so by splitting it into two: first of all, from the weak convexity of  $\Phi$  in  $x$  we have that

$$\begin{aligned} \langle g_k, \hat{x}_k - x_k \rangle &\leq \Phi(\hat{x}_k, y_k) - \Phi(x_k, y_k) + \frac{\rho}{2} \|\hat{x}_k - x_k\|^2 \\ &\leq \varphi(\hat{x}_k) - \Gamma(x_k, y_k) + \frac{\rho}{2} \|\hat{x}_k - x_k\|^2. \end{aligned}$$

Secondly, by the  $\rho$ -weak convexity of  $\varphi$

$$-\langle v_k, \hat{x}_k - x_k \rangle \leq \varphi(x_k) - \varphi(\hat{x}_k) + \frac{\rho}{2} \|\hat{x}_k - x_k\|^2.$$

Combining the last two inequalities we get that

$$(3.11) \quad \langle g_k - v_k, \hat{x}_k - x_k \rangle \leq \varphi(x_k) - \Gamma(x_k, y_k) + \rho \|\hat{x}_k - x_k\|^2.$$

Plugging (3.11) into (3.10) we deduce

$$(3.12) \quad \|\hat{x}_k - x_{k+1}\|^2 \leq \underbrace{[(1 - \eta_x \lambda^{-1})^2 + 2\eta_x (1 - \eta_x \lambda^{-1}) \rho]}_{= (*)} \|\hat{x}_k - x_k\|^2 + 2\eta_x \Delta_k + 4\eta_x^2 L^2,$$

where we used the fact that  $1 - \beta \leq 1$  in the factor of  $\Delta_k$ . Now note that

$$(3.13) \quad \begin{aligned} (*) &= 1 - 2\eta_x \lambda^{-1} + \eta_x^2 \lambda^{-2} + 2\eta_x \rho - 2\eta_x^2 \lambda^{-1} \rho \\ &= 1 - 4\eta_x \rho + 4\eta_x^2 \rho^2 + 2\eta_x \rho - 4\eta_x^2 \rho^2 = 1 - 2\eta_x \rho. \end{aligned}$$

Combining (3.7), (3.12) and (3.13) we deduce, using  $\lambda = 1/2\rho$ , □

$$\begin{aligned} \psi_\lambda(x_{k+1}) &\leq \psi(\hat{x}_k) + \frac{1}{2\lambda} (\|\hat{x}_k - x_k\|^2 + 2\eta_x \Delta_k - 2\eta_x \rho \|\hat{x}_k - x_k\|^2 + 4\eta_x^2 L^2) \\ &= \psi_\lambda(x_k) + 2\rho \eta_x \Delta_k - \frac{1}{2} \eta_x \|\nabla \psi_\lambda(x_k)\|^2 + 4\rho \eta_x^2 L^2. \end{aligned}$$

Naturally, we want to telescope the inequality established by the previous lemma. We are left with estimating  $\Delta_k$ , preferably even in a summable way. But first we need the following technical, yet standard lemma, estimating the amount of increase obtained by a single iteration of gradient ascent.

LEMMA 3.6. *It holds for all  $y \in \mathbb{R}^n$  and  $k \geq 0$  that*

$$(3.14) \quad \Psi(x_{k+1}, y) - \Psi(x_{k+1}, y_{k+1}) \leq \frac{1}{2\eta_y} \left( \|y - y_k\|^2 - \|y - y_{k+1}\|^2 \right).$$

*Proof.* This is a standard estimate on the improvement made by a single prox-gradient step for a convex (in this case concave) function, see for example [5, Lemma 2.3].  $\square$

We can now use the previous lemma to estimate  $\Delta_k$ . Recall also that  $y_k^*$  denotes a maximizer of  $\Psi(x_k, \cdot)$  for all  $k \geq 0$ .

LEMMA 3.7. *We have that for all  $1 \leq m \leq k$ ,*

$$(3.15) \quad \Delta_k \leq 2\eta_x L(L + L_f)(k - m) + \frac{1}{2\eta_y} \left( \|y_{k-1} - y_m^*\|^2 - \|y_k - y_m^*\|^2 \right).$$

*Proof.* Plugging  $y = y_m^*$  into (3.14) we deduce that

$$(3.16) \quad 0 \leq \Psi(x_k, y_k) - \Psi(x_k, y_m^*) + \frac{1}{2\eta_y} \left( \|y_m^* - y_{k-1}\|^2 - \|y_m^* - y_k\|^2 \right).$$

Starting from the definition of  $\Delta_k = \Psi(x_k, y_k^*) - \Psi(x_k, y_k)$ , we add (3.16) to obtain

$$(3.17) \quad \Delta_k \leq \Psi(x_k, y_k^*) - \Psi(x_k, y_m^*) + \frac{1}{2\eta_y} \left( \|y_m^* - y_{k-1}\|^2 - \|y_m^* - y_k\|^2 \right).$$

Due to the Lipschitz continuity of  $\Phi$ , terms which only differ in their first argument will be easy to estimate. Therefore, we insert and subtract  $\Phi(x_m, y_k^*)$  to deduce

$$(3.18) \quad \begin{aligned} & \Psi(x_k, y_k^*) - \Psi(x_k, y_m^*) \\ &= \Phi(x_k, y_k^*) - \Phi(x_m, y_k^*) + \Phi(x_m, y_k^*) - h(y_k^*) - \Phi(x_k, y_m^*) + h(y_m^*) \\ &\leq \Phi(x_k, y_k^*) - \Phi(x_m, y_k^*) + \Phi(x_m, y_m^*) - h(y_m^*) - \Phi(x_k, y_m^*) + h(y_m^*) \\ &= \Phi(x_k, y_k^*) - \Phi(x_m, y_k^*) + \Phi(x_m, y_m^*) - \Phi(x_k, y_m^*). \end{aligned}$$

We estimate the above expression for  $k > m$  by making use of the Lipschitz continuity of  $\Phi(\cdot, y)$  and (3.1) deducing

$$(3.19) \quad \begin{aligned} \Phi(x_k, y_k^*) - \Phi(x_m, y_k^*) &\leq L \|x_k - x_m\| \leq L \sum_{l=m}^{k-1} \|x_{l+1} - x_l\| \\ &\leq L \sum_{l=m}^{k-1} \left( \|\text{prox}_{\eta_x f}(x_l - \eta_x g_l) - \text{prox}_{\eta_x f}(x_l)\| + \|\text{prox}_{\eta_x f}(x_l) - x_l\| \right) \\ &\leq \eta_x L(L + L_f)(k - m). \end{aligned}$$

For  $k = m$  the inequality follows trivially. Analogously, we deduce

$$(3.20) \quad \Phi(x_m, y_m^*) - \Phi(x_k, y_m^*) \leq \eta_x L(L + L_f)(k - m).$$

Plugging (3.18), (3.19) and (3.20) into (3.17) gives the statement of the lemma.  $\square$

In order to estimate the summation of  $\Delta_k$  we will use a trick to sum over it in blocks, where the size  $B$  of these blocks will depend on the total number of iterations  $K$ . Note that w.l.o.g. we assume that the block size  $B \leq K$  divides  $K$  without remainder.

LEMMA 3.8. *It holds that for all  $K \geq 1$*

$$(3.21) \quad \frac{1}{K} \sum_{k=0}^{K-1} \Delta_k \leq \eta_x L(L + L_f)B + \frac{L_{\nabla\Phi} D_h^2}{2B} + \frac{\Delta_0}{K}.$$

*Proof.* By splitting the summation into blocks we get that

$$(3.22) \quad \sum_{k=0}^{K-1} \Delta_k = \sum_{j=0}^{K/B-1} \sum_{k=jB}^{(j+1)B-1} \Delta_k.$$

By using (3.15) from Lemma 3.7 with  $j > 0$  and  $m = jB$  and the fact that  $\sum_{k=1}^{B-1} k \leq B^2/2$  we have

$$(3.23) \quad \sum_{k=jB}^{(j+1)B-1} \Delta_k \leq \eta_x L(L + L_f)B^2 + \frac{1}{2\eta_y} \|y_{jB-1} - y_{jB}^*\|^2$$

where the last term can be bounded by  $D_h^2$ , which was defined in Assumption 6 and denotes the diameter of  $\text{dom } h$ . We do the same for the case  $j = 0$  but choose here  $m = 1$  and have separate out the first summand of  $\sum_{k=0}^{B-1} \Delta_k$  as Lemma 3.7 does not hold for  $k = 0$  and therefore get an extra  $\Delta_0$  summand. where  $D_h$  was defined in Assumption 6 and denotes the diameter of  $\text{dom } h$ . Plugging (3.23) into (3.22) gives

$$\frac{1}{K} \sum_{k=0}^{K-1} \Delta_k \leq \eta_x L(L + L_f)B + \frac{1}{2\eta_y B} D_h^2 + \frac{\Delta_0}{K}.$$

The desired statement is obtained by using the stepsize  $\eta_y = 1/L_{\nabla\Phi}$ .  $\square$

*Proof of Theorem 3.3.* From Lemma 3.5 we deduce by summing up

$$\psi_\lambda(x_K) \leq \psi_\lambda(x_0) + 2\eta_x \rho \sum_{k=0}^{K-1} \Delta_k - \frac{1}{2} \eta_x \sum_{k=0}^{K-1} \|\nabla \psi_\lambda(x_k)\|^2 + 4K\rho\eta_x^2 L^2.$$

Next, we divide by  $K$  and obtain that

$$\frac{1}{K} \sum_{k=0}^{K-1} \|\nabla \psi_\lambda(x_k)\|^2 \leq 2 \frac{\Delta^*}{\eta_x K} + \frac{4\rho}{K} \sum_{k=0}^{K-1} \Delta_k + 8\rho\eta_x L^2.$$

Now, we plug in (3.21) to deduce that

$$\frac{1}{K} \sum_{k=0}^{K-1} \|\nabla \psi_\lambda(x_k)\|^2 \leq 2 \frac{\Delta^*}{\eta_x K} + 4\rho \left( \eta_x L(L + L_f)B + \frac{L_{\nabla\Phi} D_h^2}{2B} \right) + \frac{4\rho\Delta_0}{K} + 8\eta_x \rho L^2.$$

With  $B = \frac{D_h}{L} \sqrt{\frac{L_{\nabla\Phi}}{\eta_x}}$ , we have that

$$\frac{1}{K} \sum_{k=0}^{K-1} \|\nabla \psi_\lambda(x_k)\|^2 \leq 2 \frac{\Delta^*}{\eta_x K} + 6\rho \sqrt{L_{\nabla\Phi} \eta_x} D_h (L + L_f) + \frac{4\rho\Delta_0}{K} + 8\eta_x \rho L^2.$$

By plugging in the stepsize described in the statement of the theorem we obtain

$$\frac{1}{K} \sum_{k=0}^{K-1} \|\nabla \psi_\lambda(x_k)\|^2 \leq 2 \frac{\Delta^*}{K} \max \left\{ \frac{L^2 \rho}{\epsilon^2}, \frac{L_{\nabla\Phi} \rho^2 D_h^2 (L + L_f)^2}{\epsilon^4} \right\} + \frac{4\rho\Delta_0}{K} + 14\epsilon^2.$$

proving the desired complexity result.  $\square$

**3.4. Stochastic setting.** For initial values  $(x_0, y_0) \in \text{dom } f \times \text{dom } h$  the stochastic version of alternating GDA is given by

$$(3.24) \quad (\forall k \geq 0) \begin{cases} x_{k+1} = \text{prox}_{\eta_x f} \left( x_k - \eta_x g_k^\xi \right) \\ y_{k+1} = \text{prox}_{\eta_y h} \left( y_k + \eta_y \nabla_y \Phi(x_{k+1}, y_k; \zeta_k) \right), \end{cases}$$

for  $g_k^\xi \in \partial[\Phi(\cdot, y_k; \xi_k)](x_k)$  for  $\xi_k, \zeta_k \sim \mathcal{D}$  independent from all previous iterates.

**THEOREM 3.9.** *Let in addition to the assumptions of Theorem 3.3 also Assumption 1 and 2 hold true. For algorithm (3.24) with stepsizes*

$$\eta_x = \min \left\{ \frac{\epsilon^2}{\rho(L^2 + \sigma^2)}, \frac{\epsilon^4}{\rho^2 L(L + L_f + \sigma) D_h^2 L_{\nabla \Phi}}, \frac{\epsilon^6}{\rho^3 L(L + L_f + \sigma) \sigma^2 D_h^2} \right\},$$

$\eta_y = \min \left\{ \frac{1}{2L_{\nabla \Phi}}, \frac{\epsilon^2}{\rho \sigma^2} \right\}$  and  $\lambda = \frac{1}{2\rho}$  the number of stochastic gradient evaluations  $K$  required is

$$\mathcal{O} \left( \frac{\Delta^*(L^2 + L_f^2 + \sigma^2)\rho}{\epsilon^4} \max \left\{ 1, \frac{\rho L_{\nabla \Phi} L D_h^2}{\epsilon^2}, \frac{\rho^2 D_h^2 \sigma^2}{\epsilon^4} \right\} + \frac{\Delta_0 \rho}{\epsilon^2} \right),$$

where  $\Delta^* = \psi(x_0) - \inf_{x \in \mathbb{R}^d} \psi(x)$  and  $\Delta_0 = \psi(x_0) - \Psi(x_0, y_0)$ , to visit an  $\epsilon$ -stationary point in expectation such that  $\min_{1 \leq k \leq K} \mathbb{E}[\|\nabla \psi_\lambda(x_k)\|] \leq \epsilon$ .

The proof proceeds along the same lines of the deterministic case. Similarly we show an adapted version of Lemma 3.5.

**LEMMA 3.10.** *With  $\lambda = 1/2\rho$  we have for all  $k \geq 0$  that*

$$\mathbb{E}[\psi_\lambda(x_{k+1})] \leq \mathbb{E}[\psi_\lambda(x_k)] + 2\rho\eta_x \hat{\Delta}_k - \frac{\eta_x}{2} \mathbb{E}[\|\nabla \psi_\lambda(x_k)\|^2] + 4\rho\eta_x^2(L^2 + \sigma^2)$$

where  $\hat{\Delta}_k := \mathbb{E}[\psi(x_k) - \Psi(x_k, y_k)]$ .

*Proof.* Let  $k \geq 0$  be arbitrary but fixed. It follows easily from (2.5) that

$$(3.25) \quad \mathbb{E}[\|g_k^\xi\|^2] \leq \mathbb{E}[\|g_k\|^2] + \sigma^2 \leq L^2 + \sigma^2,$$

where  $\mathbb{E}[g_k^\xi] = g_k \in \partial_x \Phi(x_k, y_k)$ . The definition of the Moreau envelope yields

$$(3.26) \quad \mathbb{E}[g_\lambda(x_{k+1})] \leq \mathbb{E}[\psi(\hat{x}_k)] + \frac{1}{2\lambda} \mathbb{E}[\|\hat{x}_k - x_{k+1}\|^2].$$

Similarly to Lemma 3.5 we deduce that for  $v_k \in \partial\varphi(\hat{x}_k)$  (as given in Lemma 3.4) and  $\beta = 1 - \eta_x \lambda^{-1}$

$$\|\hat{x}_k - x_{k+1}\|^2 = \beta^2 \|\hat{x}_k - x_k\|^2 + 2\eta_x \beta \langle g_k^\xi - v_k, \hat{x}_k - x_k \rangle + \eta_x^2 \|g_k^\xi - v_k\|^2.$$

By applying the conditional expectation  $\mathbb{E}[\cdot | x_k, y_k]$ , then the unconditional one and using (3.25), we get that

$$\mathbb{E}[\|\hat{x}_k - x_{k+1}\|^2] \leq \mathbb{E}[\|\hat{x}_k - x_k\|^2 + 2\eta_x (1 - \frac{\eta_x}{\lambda}) \langle g_k - v_k, \hat{x}_k - x_k \rangle] + 4\eta_x^2 (L^2 + \sigma^2).$$

where  $g_k = \mathbb{E}[g_k^\xi]$ . Lastly, we combine the above inequality with (3.26) and the estimate for the inner product (3.11) as in Lemma 3.5 to deduce the statement of the lemma.  $\square$

Next, we discuss the stochastic version of Lemma 3.6. It is clear that we cannot expect the same amount of function value increase by a single iteration of gradient ascent if we do not use the exact gradient.

LEMMA 3.11. *With  $\eta_y \leq 1/2L_{\nabla\Phi}$  we have for all  $k \geq 0$  and all  $y \in \mathbb{R}^n$*

$$(3.27) \quad \mathbb{E}[\Psi(x_{k+1}, y) - \Psi(x_{k+1}, y_{k+1})] \leq \frac{1}{2\eta_y} \left( \mathbb{E}[\|y_m^* - y_k\|^2 - \|y_m^* - y_{k+1}\|^2] \right) + \eta_y \sigma^2.$$

*Proof.* Let  $k \geq 0$  and  $y \in \mathbb{R}^n$  be arbitrary but fixed. As in Lemma 3.6, we deduce

$$\begin{aligned} h(y_{k+1}) - \langle \nabla_y \Phi(x_{k+1}, y_k; \zeta_k), y_{k+1} - y_k \rangle + \frac{1}{2\eta_y} \left( \|y_{k+1} - y_k\|^2 + \|y - y_{k+1}\|^2 \right) \\ \leq h(y) - \langle \nabla_y \Phi(x_{k+1}, y_k; \zeta_k), y - y_k \rangle + \frac{1}{2\eta_y} \|y - y_k\|^2. \end{aligned}$$

The term  $\langle \nabla_y \Phi(x_{k+1}, y_k; \zeta_k), y_{k+1} - y_k \rangle$  is problematic, because the right hand side of the inner product is not measurable with respect to the sigma algebra generated by past iterates  $\mathcal{F}_k := \sigma\{x_{k+1}, \dots, x_1, y_k, \dots, y_1\}$ , so we insert and subtract  $\langle \nabla_y \Phi(x_{k+1}, y_k), y_{k+1} - y_k \rangle$ . Now, using Young's inequality we estimate the resulting inner product

$$\begin{aligned} \langle \nabla_y \Phi(x_{k+1}, y_k; \zeta_k) - \nabla_y \Phi(x_{k+1}, y_k), y_{k+1} - y_k \rangle \\ \leq \eta_y \|\nabla_y \Phi(x_{k+1}, y_k; \zeta_k) - \nabla_y \Phi(x_{k+1}, y_k)\|^2 + \frac{1}{4\eta_y} \|y_{k+1} - y_k\|^2. \end{aligned}$$

Combining the above two inequalities for  $y = y_m^*$  with  $1 \leq m \leq k$  and taking the expectation together with the bounded variance assumption (2.4) gives

$$(3.28) \quad \begin{aligned} & \mathbb{E}[\mathbb{E}[\langle \nabla_y \Phi(x_{k+1}, y_k; \zeta_k), y_m^* - y_k \rangle \mid \mathcal{F}_k]] + \mathbb{E} \left[ h(y_{k+1}) + \frac{1}{2\eta_y} \|y_m^* - y_{k+1}\|^2 \right] \\ & \leq \mathbb{E} \left[ \langle \nabla_y \Phi(x_{k+1}, y_k), y_{k+1} - y_k \rangle - \frac{1}{4\eta_y} \|y_{k+1} - y_k\|^2 \right] \\ & \quad + \eta_y \sigma^2 + \mathbb{E} \left[ h(y_m^*) + \frac{1}{2\eta_y} \|y_m^* - y_k\|^2 \right]. \end{aligned}$$

From the descent lemma (in ascent form) and the fact that  $\eta_y \leq 1/2L_{\nabla\Phi}$  we have

$$\Phi(x_{k+1}, y_k) + \langle y_{k+1} - y_k, \nabla_y \Phi(x_{k+1}, y_k) \rangle - \frac{1}{4\eta_y} \|y_{k+1} - y_k\|^2 \leq \Phi(x_{k+1}, y_{k+1}).$$

We plug the above inequality into (3.28), make use of the concavity and add  $f(x_{k+1})$  on both sides to deduce the statement of the lemma.  $\square$

We can now use the previous lemma to estimate  $\hat{\Delta}_k$ .

LEMMA 3.12. *For all  $1 \leq m \leq k$ , we have that*

$$(3.29) \quad \hat{\Delta}_k \leq 2\eta_x L(L_f + L + \sigma)(k - m) + \frac{1}{2\eta_y} \mathbb{E}[\|y_{k-1} - y_m^*\|^2 - \|y_k - y_m^*\|^2] + \eta_y \sigma^2.$$

*Proof.* Let the numbers  $1 \leq m \leq k$  be fixed. Starting from the definition of  $\hat{\Delta}_k$ , we add (3.27) to obtain

$$(3.30) \quad \begin{aligned} \hat{\Delta}_k &= \mathbb{E}[\psi(x_k) - \Psi(x_k, y_k)] \\ &\leq \mathbb{E}[\psi(x_k) - \Psi(x_k, y_m^*)] + \frac{1}{2\eta_y} \left( \mathbb{E}[\|y_m^* - y_{k-1}\|^2 - \|y_m^* - y_k\|^2] \right) + \eta_y \sigma^2. \end{aligned}$$

As in (3.18) we deduce that

$$\Psi(x_k, y_k^*) - \Psi(x_k, y_m^*) \leq \Phi(x_k, y_k^*) - \Phi(x_m, y_k^*) + \Phi(x_m, y_m^*) - \Phi(x_k, y_m^*).$$

Together with the  $L$ -Lipschitz continuity of  $\Phi(\cdot, y)$  and (3.1) we estimate for  $k \geq m$  as in (3.19), but using  $\mathbb{E}[\|g_l^\xi\|] \leq L + \sigma$ ,

$$\mathbb{E}[\Phi(x_k, y_k^*) - \Phi(x_m, y_k^*)] \leq \eta_x L \left( L_f + \sqrt{L^2 + \sigma^2} \right) (k - m).$$

Plugging all of these into (3.30) gives the statement of the lemma.  $\square$

In order to estimate the summation of  $\hat{\Delta}_k$  we will use the same trick as in the deterministic setting and sum over it in blocks, where the size  $B$  of these blocks will divide the total number of iterations  $K$ .

LEMMA 3.13. *We have that for all  $K \geq 1$*

$$(3.31) \quad \frac{1}{K} \sum_{k=0}^{K-1} \hat{\Delta}_k \leq \eta_x L(L + L_f + \sigma)B + \frac{D_h^2}{2\eta_y B} + \eta_y \sigma^2 + \frac{\Delta_0}{K}.$$

*Proof.* We proceed as in Lemma 3.8. By using Lemma 3.12 we obtain  $j > 0$  and  $m = jB$  we have that

$$(3.32) \quad \sum_{k=jB}^{(j+1)B-1} \hat{\Delta}_k \leq \eta_x L(L + L_f + \sigma)B^2 + \frac{1}{2\eta_y} D_h^2 + B\eta_y \sigma^2.$$

For  $j = 0$  we use  $m = 1$  and do not estimate  $\Delta_0$  but leave it there. Plugging (3.32) into (3.22) gives the statement of the lemma.  $\square$

Now we can prove the convergence result for the stochastic algorithm.

*Proof of Theorem 3.9.* We sum up the inequality of Lemma 3.10 to deduce that

$$\mathbb{E}[\psi_\lambda(x_K)] \leq \psi_\lambda(x_0) + 2\eta_x \rho \sum_{k=0}^{K-1} \hat{\Delta}_k - \frac{\eta_x}{2} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla \psi_\lambda(x_k)\|^2] + 4K\rho\eta_x^2(L^2 + \sigma^2).$$

Thus, by dividing by  $K$  and  $\eta_x$  yields

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla \psi_\lambda(x_k)\|^2] \leq \frac{2\Delta^*}{\eta_x K} + \frac{4\rho}{K} \sum_{k=0}^{K-1} \hat{\Delta}_k + 8\rho\eta_x(L^2 + \sigma^2).$$

Now we plug in (3.31) to obtain

$$\begin{aligned} &\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla \psi_\lambda(x_k)\|^2] \\ &\leq \frac{2\Delta^*}{\eta_x K} + 4\rho \left( \eta_x L(L + L_f + \sigma)B + \frac{D_h^2}{2\eta_y B} + \eta_y \sigma^2 \right) + 4\rho \frac{\Delta_0}{K} + 8\rho\eta_x(L^2 + \sigma^2). \end{aligned}$$



With the block size  $B = D_h \sqrt{1/(\eta_x \eta_y L(L + L_f + \sigma))}$  we have that

$$\eta_x L(L + L_f + \sigma)B + \frac{D_h^2}{2\eta_y B} + \eta_y \sigma^2 = \sqrt{\frac{\eta_x}{\eta_y}} \sqrt{L(L + L_f + \sigma)} D_h + \eta_y \sigma^2$$

Via the stepsize choice presented in the theorem we obtain the desired complexity.  $\square$

**3.5. Alternating vs simultaneous.** Although we are not able to show improved rates for the alternating version of GDA in this setting, we would still like to point out some improvements in the constants which otherwise might go unnoticed since the statements are quite technical.

In the nonconvex-concave setting the main descent type property we focus on can be seen in Lemma 3.5 and is given by

$$g_\lambda(x_{k+1}) \leq g_\lambda(x_k) + 2\rho\eta_x \Delta_k - \frac{1}{2}\eta_x \|\nabla g_\lambda(x_k)\|^2 + 4\rho\eta_x^2 L^2,$$

From this it clear the only troublesome part is the estimation of  $\Delta_k = \psi(x_k) - \Psi(x_k, y_k)$  (to be precise we, we need to estimate the sum of  $\Delta_k$  after telescoping). While we obtain

$$\Delta_k \leq \Phi(x_k, y_k) - \Phi(x_m, y_k^*) + \Phi(x_m, y_m^*) - \Phi(x_k, y_m^*) + \frac{1}{2\eta_y} (\|y - y_{k-1}\|^2 - \|y - y_k\|^2)$$

in [30] the same estimate is obtained for simultaneous GDA plus an additional term

$$\Phi(x_k, y_k) - \Phi(x_{k-1}, y_{k-1}),$$

see [30, Lemma D.4]. While we do not have information about the sign of this term it is clear that its absence is preferable as it needs to be estimated after telescoping and averaging via

$$\frac{1}{K} (\Phi(x_K, y_K) - \Phi(x_0, y_0)) \leq \eta_x L^2 + \Delta_0.$$

Looking at the statement of Lemma 3.7 we see, however, that factors of both of these terms already appear in the final statement due to other estimations which is why their appearance gets lost in the big O notation.

**4. Nonconvex-strongly concave objective.** By requiring in addition to the assumptions of Section 3 *strong* convexity in the second component and smoothness of the coupling function in  $x$ , we can drop any assumption about Lipschitz continuity and will be able to deduce the max function  $\varphi$  is smooth with Lipschitz continuous gradient (making it weakly convex). For the technical details see the following assumptions.

**ASSUMPTION 7.** *Let  $\Phi$  be  $L_{\nabla\Phi}$ -smooth uniformly in both components and concave in the second one. The regularizers  $f$  and  $-h$  are proper, l.s.c. and convex. Additionally, either  $\Phi$  is  $\mu$ -strongly concave in the second component, uniformly in the first one, or  $-h$  is  $\mu$ -strongly concave.*

*Notation.* In Proposition 4.2 we will show that under the above assumptions  $\varphi = \max_{y \in \mathbb{R}^n} \{\Phi(\cdot, y) - h(y)\}$  is  $L_{\nabla\varphi}$ -smooth, with  $L_{\nabla\varphi} = (1 + \kappa)L_{\nabla\Phi}$ , for  $\kappa := \max\{L_{\nabla\Phi}/\mu, 1\}$  denoting the *condition number*. In the setting without regularizers, where the strong concavity arises from  $\Phi$  it is well known that  $\mu \leq L_{\nabla\Phi}$  and therefore  $1 \leq L_{\nabla\Phi}/\mu$  (the standard definition of the condition number). If the strong concavity

stems from the regularizers  $h$  this is no longer true and  $L_{\nabla\Phi}/\mu$  might be smaller than 1 which would lead to tedious case distinctions, which is why we adapt the definition of the condition number in order to provide a unified analysis. Additionally, the solution set  $Y(x)$  defined in (3.2) consists only of a single element which we will denote by  $y^*(x)$ . We denote the quantity  $\delta_k := \|y_k - y_k^*\|^2$ , measuring the distance between the current strategy of the second player and her best response according to the current strategy of the first player.

**4.1. Properties of the max function.** In the following we will show the smoothness of  $\varphi$ , as well as the fact that the solution map fulfills a strong Lipschitz property.

LEMMA 4.1 (Lipschitz continuity of the solution mapping). *The solution map  $y^* : \mathbb{R}^d \rightarrow \mathbb{R}^n$  which fulfills  $\Gamma(x, y^*(x)) = \max_{y \in \mathbb{R}^n} \Gamma(x, y)$  for all  $x \in \mathbb{R}^d$  is well defined and  $\kappa$ -Lipschitz where  $\kappa = \max\{L_{\nabla\Phi}/\mu, 1\}$ .*

*Proof.* Let  $x, x' \in \mathbb{R}^d$  be fixed. From the optimality condition we deduce that

$$\nabla_y \Phi(x', y^*(x')) - \nabla_y \Phi(x, y^*(x')) \in \partial h(y^*(x')) - \nabla_y \Phi(x, y^*(x')).$$

Thus by the strong monotonicity of  $\partial h - \nabla_y \Phi(x, \cdot)$  we obtain

$$\begin{aligned} \mu \|y^*(x) - y^*(x')\|^2 &\leq \langle y^*(x) - y^*(x'), \nabla_y \Phi(x, y^*(x')) - \nabla_y \Phi(x', y^*(x')) \rangle \\ &\leq \|y^*(x) - y^*(x')\| L_{\nabla\Phi} \|x - x'\|. \end{aligned}$$

The statement of the lemma follows.  $\square$

PROPOSITION 4.2 (Smoothness of the max function). *Let Assumption 7 hold true. Then,  $\varphi$  is smooth and its gradient is given by*

$$\nabla \varphi(x) = \nabla_x \Phi(x, y^*(x))$$

and is therefore  $L_{\nabla\Phi}(1 + \kappa)$ -Lipschitz.

*Proof.* Following Proposition 3.1, we define the quantities  $\tilde{\varphi}, \tilde{\Phi}$  and  $\tilde{\Gamma}$  as there using  $\rho = L_{\nabla\Phi}$ . Let  $x, v \in \mathbb{R}^d$ ,  $\alpha_k \downarrow 0$  and  $x^k := x + \alpha_k v$  for any  $k \geq 0$ . Further, let be  $y^k = y^*(x^k)$  for any  $k \geq 0$ . Then, by the Lipschitz continuity of  $y^*(\cdot)$ , see Lemma 4.1,  $\lim_{k \rightarrow \infty} y^k = y^*(x)$ . In addition, for any  $v \in \mathbb{R}^d$  and all  $k \geq 0$ ,

$$\begin{aligned} \tilde{\varphi}'(x; v) &\leq \frac{\tilde{\varphi}(x^k) - \tilde{\varphi}(x)}{\alpha_k} \leq \frac{\tilde{\Gamma}(x^k, y^k) - \tilde{\Gamma}(x, y^k)}{\alpha_k} \\ &= \frac{\tilde{\Phi}(x^k, y^k) - \tilde{\Phi}(x, y^k)}{\alpha_k} \leq [\tilde{\Phi}(\cdot, y^k)]'(x + \alpha_k v; v) = \langle \nabla_x \tilde{\Phi}(x^k, y^k), v \rangle. \end{aligned}$$

Since the gradient of  $\tilde{\Phi}$  is continuous, this implies by letting  $k \rightarrow +\infty$  that

$$\tilde{\varphi}'(x; v) \leq \langle \nabla_x \tilde{\Phi}(x, y^*(x)), v \rangle, \quad \forall v \in \mathbb{R}^d,$$

which, together with (3.4), yields that (16) holds with equality. The fact that the gradient of  $\varphi$  is Lipschitz continuous follows, with  $y^* = y^*(x)$  and  $\bar{y}^* = y^*(\bar{x})$ , from

$$\begin{aligned} \|\nabla \varphi(x) - \nabla \varphi(\bar{x})\| &\leq \|\nabla_x \Phi(x, y^*) - \nabla_x \Phi(\bar{x}, y^*)\| + \|\nabla_x \Phi(\bar{x}, y^*) - \nabla_x \Phi(\bar{x}, \bar{y}^*)\| \\ &\leq L_{\nabla\Phi} \|x - \bar{x}\| + L_{\nabla\Phi} \|y^* - \bar{y}^*\| \leq (L_{\nabla\Phi} + L_{\nabla\Phi} \kappa) \|x - \bar{x}\|, \end{aligned}$$

together with the claimed constant.  $\square$

**4.2. Deterministic setting.** For this section Algorithm 1 reads as

$$(4.1) \quad (\forall k \geq 0) \begin{cases} x_{k+1} = \text{prox}_{\eta_x f}(x_k - \eta_x \nabla_x \Phi(x_k, y_k)) \\ y_{k+1} = \text{prox}_{\eta_y h}(y_k + \eta_y \nabla_y \Phi(x_{k+1}, y_k)). \end{cases}$$

We start with the main convergence result of this section.

**THEOREM 4.3.** *Let Assumption 7 and 4 hold. For algorithm 4.1 with stepsize  $\eta_y = 1/L_{\nabla\Phi}$  and  $\eta_x = 1/(3(\kappa+1)^2 L_{\nabla\Phi})$  the number of gradient evaluations  $K$  required, using the notation  $\Delta^* = \psi(x_0) - \inf_{x \in \mathbb{R}^d} \psi(x)$ , is*

$$\mathcal{O}\left(\frac{\kappa L_{\nabla\Phi}}{\epsilon^2} \max\{\kappa \Delta^*, L_{\nabla\Phi} \|y^*(x_0) - y_0\|^2\}\right),$$

to visit an  $\epsilon$ -stationary point such that  $\min_{1 \leq k \leq K} \text{dist}(-\nabla\varphi(x_k), \partial f(x_k)) \leq \epsilon$ .

Before we start with the first lemma, let us recall that  $\delta_k = \|y_k - y_k^*\|^2$  denotes for all  $k \geq 0$  the squared distance between the current iterate  $y_k$  and the maximizing argument  $y_k^* = \arg \max_y \Psi(x_k, y)$ .

**LEMMA 4.4.** *There exists a sequence  $(w_k)_{k \geq 1}$  such that  $w_k \in (\partial f + \nabla\varphi)(x_k)$  and its norm can be bounded for all  $k \geq 0$  by*

$$\frac{1}{2} \eta_x \|w_{k+1}\|^2 \leq \psi(x_k) - \psi(x_{k+1}) + \frac{1}{2} \left( L_{\nabla\varphi} + 2L_{\nabla\varphi}^2 \eta_x - \frac{1}{\eta_x} \right) \|x_k - x_{k+1}\|^2 + \eta_x L_{\nabla\Phi}^2 \delta_k.$$

*Proof.* Let  $k \geq 0$  be arbitrary but fixed. From the optimality condition of the proximal operator we deduce by adding  $\nabla\varphi(x_{k+1})$  on both sides

$$w_{k+1} := \frac{1}{\eta_x} (x_k - x_{k+1}) + \nabla\varphi(x_{k+1}) - \nabla_x \Phi(x_k, y_k) \in \partial f(x_{k+1}) + \nabla\varphi(x_{k+1}),$$

as claimed. In order to prove the bound on  $\|w_{k+1}\|$  we proceed as follows:

$$(4.2) \quad \|w_{k+1}\|^2 = \eta_x^{-2} \|x_k - x_{k+1}\|^2 + 2\eta_x^{-1} \langle x_k - x_{k+1}, \nabla\varphi(x_{k+1}) - \nabla_x \Phi(x_k, y_k) \rangle + \|\nabla\varphi(x_{k+1}) - \nabla_x \Phi(x_k, y_k)\|^2.$$

The smoothness of  $\varphi$  implies via the descent lemma that

$$(4.3) \quad \varphi(x_{k+1}) + \langle \nabla\varphi(x_{k+1}), x_k - x_{k+1} \rangle - \frac{L_{\nabla\varphi}}{2} \|x_{k+1} - x_k\|^2 \leq \varphi(x_k).$$

Since the proximal operator minimizes a  $1/\eta_x$ -strongly convex function we have that

$$f(x_{k+1}) + \langle \nabla_x \Phi(x_k, y_k), x_{k+1} - x_k \rangle + \frac{1}{\eta_x} \|x_{k+1} - x_k\|^2 \leq f(x_k)$$

Adding this inequality to (4.3) we deduce that

$$(4.4) \quad \begin{aligned} \langle \nabla\varphi(x_{k+1}) - \nabla_x \Phi(x_k, y_k), x_k - x_{k+1} \rangle &\leq \psi(x_k) - \psi(x_{k+1}) \\ &+ \frac{1}{2} \left( L_{\nabla\varphi} - \frac{2}{\eta_x} \right) \|x_{k+1} - x_k\|^2. \end{aligned}$$

Lastly, by the Young inequality

$$\begin{aligned} \|\nabla\varphi(x_{k+1}) - \nabla_x \Phi(x_k, y_k)\|^2 &= \|\nabla\varphi(x_{k+1}) - \nabla\varphi(x_k) + \nabla\varphi(x_k) - \nabla_x \Phi(x_k, y_k)\|^2 \\ &\leq 2L_{\nabla\varphi}^2 \|x_{k+1} - x_k\|^2 + 2L_{\nabla\Phi}^2 \delta_k. \end{aligned}$$

Plugging (4.4) and (4.3) into (4.2) yields the desired statement.  $\square$

In the next lemma it remains to bound the gap between the current iterate and the maximizing argument of the second component  $\delta_k = \|y_k^* - y_k\|^2$ .

LEMMA 4.5. *We have that for all  $k \geq 0$ , then*

$$\delta_{k+1} \leq \left(1 - \frac{1}{2\kappa}\right) \delta_k + \kappa^3 \|x_{k+1} - x_k\|^2.$$

*Proof.* Let  $k \geq 0$  be fixed. From the definition of  $y_{k+1}$ , see (4.1), and the fact that  $y_{k+1}^*$  is a fixed point of the proximal-gradient operator, we deduce

$$\delta_{k+1} = \|\text{prox}_{\eta_y h}(y_{k+1}^* + \eta_y \nabla_y \Phi(x_{k+1}, y_{k+1}^*)) - \text{prox}_{\eta_y h}(y_k + \eta_y \nabla_y \Phi(x_{k+1}, y_k))\|^2.$$

If  $\Phi$  is strongly concave in its second component we can use the nonexpansiveness of the proximal operator and [42, Theorem 2.1.11], which states

$$(4.5) \quad \begin{aligned} & \langle \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k), y_{k+1}^* - y_k \rangle \\ & \leq -\frac{\mu L_{\nabla \Phi}}{\mu + L_{\nabla \Phi}} \|y_{k+1}^* - y_k\|^2 - \frac{1}{\mu + L_{\nabla \Phi}} \|\nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k)\|^2 \end{aligned}$$

to conclude

$$\begin{aligned} \delta_{k+1} & \leq \|y_{k+1}^* + \eta_y \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - y_k - \eta_y \nabla_y \Phi(x_{k+1}, y_k)\|^2 \\ & = \|y_{k+1}^* - y_k\|^2 + 2\eta_y \langle y_{k+1}^* - y_k, \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k) \rangle \\ & \quad + \eta_y^2 \|\nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k)\|^2 \\ & \stackrel{(4.5)}{\leq} \left(\frac{\kappa - 1}{\kappa + 1}\right) \|y_{k+1}^* - y_k\|^2 \leq q \|y_{k+1}^* - y_k\|^2 \end{aligned}$$

with  $q := \left(\frac{\kappa}{\kappa+1}\right)^2$ , where we used that  $\eta_y = 1/L_{\nabla \Phi}$ . If on the other hand  $-h$  is strongly concave we can use the fact that the proximal operator (of  $h$ ) is even a contraction, see [4, Proposition 25.9 (i)], to deduce that  $\delta_{k+1} \leq q \|y_{k+1}^* - y_k\|^2$ . Therefore, in either case  $\delta_{k+1} \leq q \|y_{k+1}^* - y_k\|^2$ . Using this, the triangle inequality and Young's inequality, we have

$$(4.6) \quad \begin{aligned} \delta_{k+1} & \leq q \|y_{k+1}^* - y_k\|^2 \leq q \left( \|y_{k+1}^* - y_k^*\| + \|y_k^* - y_k\| \right)^2 \\ & \leq q \left( 1 + \frac{3\kappa^2 - 1}{2\kappa^3} \right) \underbrace{\|y_k^* - y_k\|^2}_{=\delta_k} + q \left( 1 + \frac{2\kappa^3}{3\kappa^2 - 1} \right) \|y_{k+1}^* - y_k^*\|^2 \\ & \leq \left( 1 - \frac{1}{2\kappa} \right) \delta_k + \kappa \|y_{k+1}^* - y_k^*\|^2. \end{aligned}$$

Due to the  $\kappa$ -Lipschitz continuity of  $y^*(\cdot)$  we have that  $\|y_{k+1}^* - y_k^*\| \leq \kappa \|x_{k+1} - x_k\|$ , which finishes the proof.  $\square$

Now we can bound the sum of  $\delta_k$ .

LEMMA 4.6. *We have that, for all  $K \geq 1$*

$$\sum_{k=0}^{K-1} \delta_k \leq 2\kappa \delta_0 + 2\kappa^4 \sum_{k=0}^{K-1} \|x_{k+1} - x_k\|^2.$$

*Proof.* By recursively applying the previous lemma we obtain for  $k \geq 1$

$$\delta_k \leq \left(1 - \frac{1}{2\kappa}\right)^k \delta_0 + \kappa^3 \sum_{j=0}^{k-1} \left(1 - \frac{1}{2\kappa}\right)^{k-j-1} \|x_{j+1} - x_j\|^2.$$

Now we sum this inequality from  $k = 1$  to  $K - 1$  and add  $\delta_0$  on both sides to deduce

$$\sum_{k=0}^{K-1} \delta_k \leq 2\kappa\delta_0 + 2\kappa^4 \sum_{k=0}^{K-1} \|x_{k+1} - x_k\|^2,$$

where we used that

$$(4.7) \quad \sum_{k=1}^{K-1} \sum_{j=0}^{k-1} \left(1 - \frac{1}{2\kappa}\right)^{k-1-j} \|x_{j+1} - x_j\|^2 \leq \sum_{j=0}^{K-1} \left(1 - \frac{1}{2\kappa}\right)^j \left(\sum_{k=0}^{K-1} \|x_{k+1} - x_k\|^2\right),$$

and  $\sum_{j=0}^{\infty} (1 - (2\kappa)^{-1})^j = 2\kappa$ .  $\square$

We can now put the pieces together.

*Proof of Theorem 4.3.* Summing up the inequality of Lemma 4.4 from  $k = 0$  to  $K - 1$  and applying Lemma 4.6 we deduce that

$$\begin{aligned} \frac{1}{2}\eta_x \sum_{k=1}^K \|w_k\|^2 &\leq \psi(x_0) - \psi(x_K) + 2L_{\nabla\Phi}^2 \eta_x \kappa \delta_0 \\ &\quad + \frac{1}{2} \left( L_{\nabla\varphi} + 2L_{\nabla\varphi}^2 \eta_x - \frac{1}{\eta_x} + 2\kappa^4 L_{\nabla\Phi}^2 \eta_x \right) \sum_{k=0}^{K-1} \|x_{k+1} - x_k\|^2. \end{aligned}$$

With the stepsize  $\eta_x = 1/(3(\kappa + 1)^2 L_{\nabla\Phi})$  it follows that

$$L_{\nabla\varphi} + 2L_{\nabla\varphi}^2 \eta_x - \frac{1}{\eta_x} + 2\kappa^4 L_{\nabla\Phi}^2 \eta_x \leq -\frac{2}{3}(\kappa + 1)^2 L_{\nabla\Phi} \leq 0,$$

which concludes the proof.  $\square$

**4.3. Stochastic setting.** For the purpose of this section Algorithm 1 reads

$$(4.8) \quad (\forall k \geq 0) \begin{cases} x_{k+1} = \text{prox}_{\eta_x f}(x_k - \eta_x G_x) \\ y_{k+1} = \text{prox}_{\eta_y h}(y_k + \eta_y G_y), \end{cases}$$

for the minibatch gradient estimators, given by  $G_x = \frac{1}{M} \sum_{i=1}^M \nabla_x \Phi(x_k, y_k; \xi_k^i)$  and  $G_y = \frac{1}{M} \sum_{i=1}^M \nabla_y \Phi(x_{k+1}, y_k; \zeta_k^i)$ .

**THEOREM 4.7.** *Let in addition to the assumptions of Theorem 4.3 also the two properties of the gradient estimator Assumption 1 and 2 hold true. For algorithm (4.8) with stepsize  $\eta_y = 1/L_{\nabla\Phi}$  and  $\eta_x = 1/(4(1+\kappa)^2 L_{\nabla\Phi})$  and batch size  $M = \mathcal{O}(\kappa\sigma^2\epsilon^{-2})$  the number of stochastic gradient evaluations  $K$  required is*

$$\mathcal{O} \left( \frac{\sigma^2 \kappa^2 L_{\nabla\Phi}}{\epsilon^4} \max\{\kappa\Delta^*, L_{\nabla\Phi} \|y^*(x_0) - y_0\|^2\} \right),$$

such that  $\min_{1 \leq k \leq K} \mathbb{E}[\text{dist}(-\nabla\varphi(x_k), \partial f(x_k))] \leq \epsilon$ , i.e. to visit an  $\epsilon$ -stationary point in expectation, where  $\Delta^* = \psi(x_0) - \inf_{x \in \mathbb{R}^d} \psi(x)$ .

LEMMA 4.8. *There exists a sequence  $(w_k)_{k \geq 1}$  such that  $w_k \in (\partial f + \nabla \varphi)(x_k)$  and its norm can be bounded for all  $k \geq 0$  by*

$$\begin{aligned} \frac{1}{2} \eta_x \mathbb{E}[\|w_{k+1}\|^2] &\leq \mathbb{E}[\psi(x_k) - \psi(x_{k+1})] + \eta_x L_{\nabla \Phi}^2 \mathbb{E}[\delta_k] + \eta_x \frac{\sigma^2}{M} \\ &\quad + \frac{1}{2} \left( L_{\nabla \varphi} + 3L_{\nabla \varphi}^2 \eta_x - \frac{1}{\eta_x} \right) \mathbb{E}[\|x_k - x_{k+1}\|^2]. \end{aligned}$$

*Proof.* Let  $k \geq 0$ . From the proximal operator we deduce that

$$0 \in \partial f(x_{k+1}) + G_x + \frac{1}{\eta_x} (x_{k+1} - x_k).$$

Thus, we define  $w_{k+1}$  such that we immediately obtain the desired inclusion

$$w_{k+1} := \frac{1}{\eta_x} (x_k - x_{k+1}) + \nabla \varphi(x_{k+1}) - G_x \in \partial f(x_{k+1}) + \nabla \varphi(x_{k+1}).$$

In order to bound  $w_{k+1}$  we consider

$$(4.9) \quad \begin{aligned} &\|w_{k+1}\|^2 \\ &= \frac{1}{\eta_x^2} \|x_k - x_{k+1}\|^2 + \frac{2}{\eta_x} \langle x_k - x_{k+1}, \nabla \varphi(x_{k+1}) - G_x \rangle + \|\nabla \varphi(x_{k+1}) - G_x\|^2. \end{aligned}$$

Using the analogous statement to (4.4)

$$\begin{aligned} \mathbb{E}[\|\nabla \varphi(x_{k+1}) - G_x\|^2] &= \mathbb{E}[\|\nabla \varphi(x_{k+1}) - \nabla_x \Phi(x_k, y_k) + \nabla_x \Phi(x_k, y_k) - G_x\|^2] \\ &= \mathbb{E}[\|\nabla \varphi(x_{k+1}) - \nabla \varphi(x_k) + \nabla \varphi(x_k) - \nabla_x \Phi(x_k, y_k)\|^2] \\ &\quad + 2\mathbb{E}[\langle \nabla \varphi(x_{k+1}) - \nabla \varphi(x_k), \nabla_x \Phi(x_k, y_k) - G_x \rangle] + \mathbb{E}[\|\nabla_x \Phi(x_k, y_k) - G_x\|^2] \\ &\quad + 2\mathbb{E}[\underbrace{\langle \nabla \varphi(x_k) - \nabla_x \Phi(x_k, y_k), \nabla_x \Phi(x_k, y_k) - G_x \rangle}_{=0}] \\ &\leq 3L_{\nabla \varphi}^2 \mathbb{E}[\|x_{k+1} - x_k\|^2] + 2L_{\nabla \Phi}^2 \mathbb{E}[\delta_k] + 2\frac{\sigma^2}{M} \end{aligned}$$

in (4.9) yields the desired statement.  $\square$

In the next lemma it remains to bound  $\delta_k$ .

LEMMA 4.9. *We have that for all  $k \geq 0$*

$$\mathbb{E}[\delta_{k+1}] \leq \left(1 - \frac{1}{2\kappa}\right) \mathbb{E}[\delta_k] + \kappa^3 \|x_{k+1} - x_k\|^2 + \frac{\sigma^2}{ML_{\nabla \Phi}^2}.$$

*Proof.* Let  $k \geq 0$  be fixed. We first consider the case where  $\Phi$  is strongly concave in its second component from the definition of  $y_{k+1}$  (see (4.8)) we deduce that

$$\begin{aligned} \delta_{k+1} &= \|\text{prox}_{\eta_y h}(y_{k+1}^* + \eta_y \nabla_y \Phi(x_{k+1}, y_{k+1}^*)) - \text{prox}_{\eta_y h}(y_k + \eta_y G_y)\|^2 \\ &\leq \|y_{k+1}^* + \eta_y \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - y_k - \eta_y G_y\|^2 \end{aligned}$$

and

$$\begin{aligned}
(4.10) \quad & \|y_{k+1}^* + \eta_y \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - y_k - \eta_y G_y\|^2 \\
&= \|y_{k+1}^* - y_k\|^2 + 2\eta_y \langle y_{k+1}^* - y_k, \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k) \rangle \\
&\quad + 2\eta_y \underbrace{\langle y_{k+1}^* - y_k, \nabla_y \Phi(x_{k+1}, y_k) - G_y \rangle}_{(\square)} \\
&\quad + \eta_y^2 \|\nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k)\|^2 + \eta_y^2 \|\nabla_y \Phi(x_{k+1}, y_k) - G_y\|^2 \\
&\quad + \eta_y^2 \underbrace{\langle \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k), \nabla_y \Phi(x_{k+1}, y_k) - G_y \rangle}_{(*)}.
\end{aligned}$$

Some of the terms vanish after taking the expectation such as

$$\mathbb{E}[(*)] = \mathbb{E}[\mathbb{E}[(*) \mid y_k, x_{k+1}]] = \mathbb{E}[0] = 0 = \mathbb{E}[\mathbb{E}[(\square) \mid y_k, x_{k+1}]] = \mathbb{E}[(\square)].$$

Using furthermore [42, Theorem 2.1.11] which states that

$$\begin{aligned}
(4.11) \quad & \langle \nabla_y \Phi(x_{k+1}, y_k) - \nabla_y \Phi(x_{k+1}, y_{k+1}^*), y_{k+1}^* - y_k \rangle \\
&\geq \frac{\mu L_{\nabla \Phi}}{\mu + L_{\nabla \Phi}} \|y_{k+1}^* - y_k\|^2 + \frac{1}{\mu + L_{\nabla \Phi}} \|\nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k)\|^2
\end{aligned}$$

results in

$$(4.12) \quad \mathbb{E}[\delta_{k+1}] \leq \left( \frac{\kappa - 1}{\kappa + 1} \right) \mathbb{E}[\|y_{k+1}^* - y_k\|^2] + \frac{\sigma^2}{ML_{\nabla \Phi}^2} \leq q \mathbb{E}[\|y_{k+1}^* - y_k\|^2] + \frac{\sigma^2}{ML_{\nabla \Phi}^2},$$

with  $q = \left( \frac{\kappa}{\kappa + 1} \right)^2$ , where we used that  $\eta_y = 1/L_{\nabla \Phi}$ . If  $h$  is strongly concave then we use the fact that the proximal operator is a contraction, see [4, Proposition 23.11], to deduce that

$$\begin{aligned}
\mathbb{E}\delta_{k+1} &= \mathbb{E} \left[ \left\| \text{prox}_{\eta_y h} (y_{k+1}^* + \eta_y \nabla_y \Phi(x_{k+1}, y_{k+1}^*)) - \text{prox}_{\eta_y h} (y_k + \eta_y G_y) \right\|^2 \right] \\
&= q \mathbb{E} \left[ \|y_{k+1}^* - y_k + \eta_y \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \eta_y G_y\|^2 \right] \\
(4.10) \quad &= q \mathbb{E} \left[ \|y_{k+1}^* - y_k\|^2 + 2q\eta_y \langle y_{k+1}^* - y_k, \nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k) \rangle \right. \\
&\quad \left. + q\eta_y^2 \|\nabla_y \Phi(x_{k+1}, y_{k+1}^*) - \nabla_y \Phi(x_{k+1}, y_k)\|^2 + q\eta_y^2 \|\nabla_y \Phi(x_{k+1}, y_k) - G_y\|^2 \right].
\end{aligned}$$

Using now (4.11) with  $\mu = 0$ , i.e. the cocoercivity of the gradient, we deduce that

$$\mathbb{E}[\delta_{k+1}] = \mathbb{E}[\|y_{k+1}^* - y_{k+1}\|^2] \leq q \|y_{k+1}^* - y_k\|^2 + q \frac{\sigma^2}{ML_{\nabla \Phi}^2},$$

meaning that we concluded (4.12) in both cases. Next, using (4.12) and the considerations made in (4.6) we deduce that

$$\mathbb{E}[\delta_{k+1}] \leq \left( 1 - \frac{1}{2\kappa} \right) \mathbb{E}[\delta_k] + \kappa \mathbb{E}[\|y_{k+1}^* - y_k^*\|^2] + \frac{\sigma^2}{ML_{\nabla \Phi}^2}.$$

Again, due to the  $\kappa$ -Lipschitz continuity of  $y^*(\cdot)$  we have that  $\|y_{k+1}^* - y_k^*\| \leq \kappa \|x_{k+1} - x_k\|$ , which finishes the proof.  $\square$

Now we can bound the sum of  $\delta_k$ .

LEMMA 4.10. *We have that, for all  $K \geq 1$*

$$\sum_{k=0}^{K-1} \mathbb{E}[\delta_k] \leq 2\kappa\delta_0 + 2\kappa^4 \sum_{k=0}^{K-1} \mathbb{E}[\|x_{k+1} - x_k\|^2] + 2K \frac{\kappa\sigma^2}{ML_{\nabla\Phi}^2}.$$

*Proof.* By recursively applying the previous lemma we obtain for  $k \geq 1$

$$\mathbb{E}[\delta_k] \leq \left(1 - \frac{1}{2\kappa}\right)^k \delta_0 + \sum_{j=0}^{k-1} \left(1 - \frac{1}{2\kappa}\right)^{k-j-1} \left(\kappa^3 \mathbb{E}[\|x_{j+1} - x_j\|^2] + \frac{\sigma^2}{ML_{\nabla\Phi}^2}\right).$$

Now we sum this inequality from  $k = 1$  to  $K - 1$  and add  $\delta_0$  on both sides to deduce

$$\sum_{k=0}^{K-1} \mathbb{E}[\delta_k] \leq 2\kappa\delta_0 + 2K \frac{\kappa\sigma^2}{ML_{\nabla\Phi}^2} + 2\kappa^4 \sum_{k=0}^{K-1} \mathbb{E}[\|x_{k+1} - x_k\|^2]$$

using the considerations made in (4.7).  $\square$

We can now put the pieces together.

*Proof of Theorem 4.7.* We sum up the inequality of Lemma 4.8 from  $k = 0$  to  $K - 1$  and applying Lemma 4.10 we deduce that

$$\begin{aligned} \eta_x \sum_{k=1}^K \mathbb{E}[\|w_k\|^2] &\leq 2\mathbb{E}[\psi(x_0) - \psi(x_K)] + 4\eta_x \kappa L_{\nabla\Phi}^2 \delta_0 + 4\eta_x \kappa K \frac{\sigma^2}{M} + 2\eta_x K \frac{\sigma^2}{M} \\ &\quad + \underbrace{\left(L_{\nabla\varphi} + 3L_{\nabla\varphi}^2 \eta_x - \frac{1}{\eta_x} + 2\kappa^4 L_{\nabla\Phi}^2 \eta_x\right)}_{= (*)} \sum_{k=0}^{K-1} \mathbb{E}[\|x_{k+1} - x_k\|^2]. \end{aligned}$$

Applying the stepsize  $\eta_x = 1/(3(1+\kappa)^2 L_{\nabla\Phi})$  it follows that

$$(*) \leq 2(\kappa + 1)L_{\nabla\Phi} - 3(\kappa + 1)^2 L_{\nabla\Phi} + \frac{2\kappa^2 L_{\nabla\Phi}}{3} \leq -\frac{1}{3}(\kappa + 1)^2 L_{\nabla\Phi} \leq 0$$

which concludes the proof.  $\square$

**4.4. Alternating vs simultaneous.** Similarly to Section 3.5 we want to highlight here the difference in the analysis between the two versions of GDA. Again, in this nonconvex-strongly-concave setting the task is to estimate  $\delta_k := \|y_k^* - y_k\|^2$ . While in the simultaneous version [30] obtains for all  $k \geq 0$  the following inequality

$$\begin{aligned} \delta_{k+1} &\leq (1 + \epsilon) \|y_k^* - y_{k+1}\|^2 + (1 + \epsilon^{-1}) \|y_{k+1}^* - y_k^*\|^2 \\ &\leq q(1 + \epsilon) \delta_k + (1 + \epsilon^{-1}) \|y_{k+1}^* - y_k^*\|^2, \end{aligned}$$

where  $q$  is the contraction constant derived from the gradient ascent step and is roughly  $1 - \frac{1}{\kappa}$ . In the alternating version however, we estimate for all  $k \geq 0$

$$\delta_{k+1} \leq q \|y_{k+1}^* - y_k\|^2 \leq q(1 + \epsilon) \delta_k + q(1 + \epsilon^{-1}) \|y_{k+1}^* - y_k^*\|^2.$$

Evidently, in the alternating version the contraction property is applied before the triangle inequality, which leads to the second term being multiplied by  $q < 1$  as well, influencing the final complexity bound favorably, albeit only slightly.



**5. Numerical Experiments.** In this section, we present several experiments outlining the empirical benefits of alternating GDA over its simultaneous counterpart.

**5.1. Toy example.** The recent paper [54] showed an improved convergence rate of alternating GDA over the simultaneous version in the strongly convex-strongly concave quadratic setting from  $\mathcal{O}(\kappa^2)$  to  $\mathcal{O}(\kappa)$ . Inspired by these results we study a nonconvex-strongly-concave toy example

$$(5.1) \quad \min_{x \in \mathbb{R}} \max_{y \in \mathbb{R}} -\frac{1}{4}x^2 + xy - \frac{1}{2}y^2,$$

where the resulting max function happens to be a strongly convex quadratic

$$(5.2) \quad \max_{y \in \mathbb{R}} \left\{ -\frac{1}{4}x^2 + xy - \frac{1}{2}y^2 \right\} = \frac{1}{4}x^2.$$

As we can see from Figure 1, alternating GDA outperforms not only its simultaneous counterpart but also the extra-gradient method (EG) [26] and the multistep method GDmax [24] (employing 10 ascent steps per descent step). For the Minimax-PPA method [29] we only counted iterations but did not account for the computational cost of the double inner procedure which required over 100 solves of a proximally regularized subproblem per iteration. We suspect that for a significantly more ill-conditioned problem the Minimax-PPA methods might have performed more competitively.

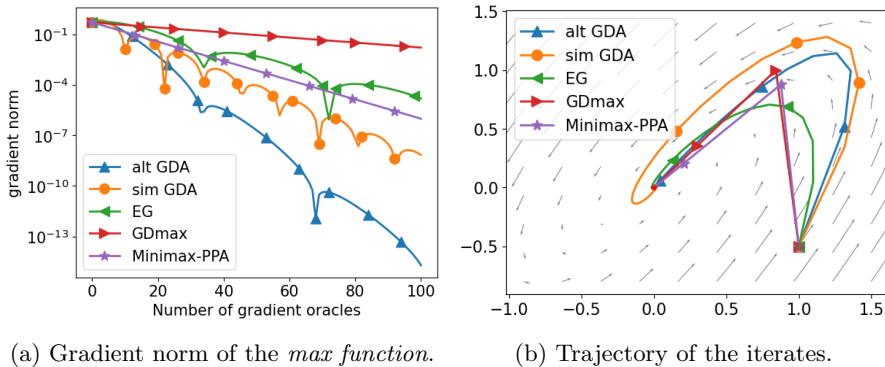


Figure 1: Comparison of different methods on the 2-d toy problem (5.1). The starting point is given by  $(x_0, y_0) = (1, -0.5)$ , step sizes for the first 4 methods are given by  $\eta_x = 1/(\kappa L_{\nabla\Phi})$ ,  $\eta_y = 1/L_{\nabla\Phi}$  where  $\mu = 1$  and  $L_{\nabla\Phi} \approx 1.78$  and the unique stationary point of the *max function* is  $x^* = 0$  as evident from (5.2). EG denotes the extra-gradient method [26], GDmax a simple multistep method [24] employing 10 ascent steps for each descent step and Minimax-PPA is the method from [29] for which all parameters chosen according to theory.

In order to account for differences in stepsizes we do a grid search across possible stepsizes for the two components and plot the number of iteration required to reach a target accuracy, see Figure 2. Since the Minimax-PPA method has many inner step sizes and number of inner loop calls to tune it cannot easily be compared here, so we excluded it. We can see that alternating GDA is not only convergent for




Figure 2: Gradient oracle calls needed to achieve an iterate gradient norm of the *max* function, see (1.3), smaller than  $10^{-4}$  using different methods.

many different combinations of stepsizes but also consistently outperforms the other methods in terms of the required gradient oracle calls.

We also want to point out that, it might appear from the convergence behavior of the different methods that our toy example (5.1) is easier than the more classical bilinear problem of  $\min_x \max_y xy$ , see [17, 16, 37], where neither version of GDA converges. While for the latter problem the vector field is always perpendicular to the direction of solution, our new problem exhibits areas where the vector field points away from the solution (see the upper right corner of the Figure 1 (b)) and thus exhibits a novel (and challenging behavior) which cannot be captured by bilinear problems.

**5.2. Adversarially robust learning.** We now highlight the performance of alternating GDA for adversarial learning on the MNIST [28], Fashion MNIST [1] and CIFAR10 [27] datasets respectively. We focus on the adversarial learning formulation described in (1.4), which originated in [47] and results in a nonconvex-strongly-concave minimax formulation for large enough  $\gamma$ . We use standard convolutional networks (CNN) for all three datasets. For MNIST and Fashion-MNIST we use the architecture proposed by [47] of three convolutional layers followed by a dense layer and softmax output. For CIFAR10 we follow the default architecture in the tutorial of [45] with seven convolutional layers where the third, fifth and seventh are followed by an average pooling operation.

Since the robust training did not significantly impact the performance of the CNNs on the clean test examples on any of the data sets, we only report the performance on the adversarial examples.

In contrast to multistep methods, as proposed in [47, 24, 43] which aim to (approximately) solve the inner maximization problem and typically start in every iteration




Figure 3: Test accuracy (percentage of correctly classified) of different adversarial training methods on multiple datasets averaged over 5 random seeds. ERM denotes empirical risk minimization and represents a standard network trained purely for classification without any consideration of the adversarial examples. For the minimization step we actually used the ADAM optimizer with a learning rate of 0.001 as done in the code of [47, 45]. For the Fashion-MNIST and CIFAR10 datasets we set the Lagrange parameter  $\gamma$  to be  $1/\gamma = 0.4$ , whereas for MNIST we use  $1/\gamma = 1.3$ .

from either the “clean” training example (or a randomly perturbed point [50]), GDA can be interpreted as a warm starting procedure which stores the adversarial example computed in the previous epoch. This also emphasizes the advantage of alternating GDA (and why it can be considered the more natural approach) as this method uses computed adversarial examples right away whereas the simultaneous version stores them to only use them in the next epoch.

Although this is not the main focus this work, we also contrast the single step methods with a method approximately solving the maximization problem (GDmax, see [30]) and observe that while the multistep method slightly outperforms alternating GDA, see Figure 3, this comes at the cost of an  $\approx 8$  times higher computation time (based on  $\approx 15$  gradient ascent steps for the multistep method).

Nevertheless, all results clearly show that alternating GDA consistently outperforms simultaneous GDA without any additional computational cost.

**6. Conclusion.** We show novel complexity results for the alternating gradient descent ascent method for nonconvex-(strongly) concave minimax problems using stochastic or deterministic gradient evaluations. Since these bounds are only a first step into the analysis of alternating GDA in this sophisticated setting, they do not explain theoretically the benefit of the alternating version. However, we provide empirical evidence that this method is favorable.

## REFERENCES

- [1] A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter, and R. Vollgraf. In *Proceedings of the 2019 Conference of the North. Association for Computational Linguistics*, 2019.
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *arXiv:1701.07875*, 2017.
- [3] J. P. Bailey, G. Gidel, and G. Piliouras. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *Proceedings of the 33rd Conference on Learning Theory*, pages 391–407. PMLR, 2020.
- [4] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, volume 408. Springer, 2011.
- [5] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse

- problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [6] A. Böhm, M. Sedlmayer, E. R. Csetnek, and R. I. Boç. Two steps at a time – taking GAN training in stride with Tseng’s method. *SIAM Journal on Mathematics of Data Science*, 4(2):750–771, 2022.
- [7] A. Böhm and S. J. Wright. Variable smoothing for weakly convex composite functions. *Journal of Optimization Theory and Applications*, 188(3):628–649, 2021.
- [8] Q. Cai, M. Hong, Y. Chen, and Z. Wang. On the global convergence of imitation learning: A case for linear quadratic regulator. *arXiv preprint arXiv:1901.03674*, 2019.
- [9] R. S. Chen, B. Lucier, Y. Singer, and V. Syrgkanis. Robust optimization for non-convex objectives. In *Advances in Neural Information Processing Systems*, pages 4705–4714, 2017.
- [10] C. Daskalakis, A. Ilyas, V. Syrgkanis, and H. Zeng. Training GANs with optimism. In *International Conference on Learning Representations*, 2018.
- [11] C. Daskalakis and I. Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *Advances in Neural Information Processing Systems*, pages 9236–9246, 2018.
- [12] D. Davis and D. Drusvyatskiy. Stochastic subgradient method converges at the rate  $O(k^{-1/4})$  on weakly convex functions. *arXiv:1802.02988*, 2018.
- [13] D. Davis and D. Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019.
- [14] D. Drusvyatskiy and C. Paquette. Efficiency of minimizing compositions of convex functions and smooth maps. *Mathematical Programming*, 178(1-2):503–558, 2019.
- [15] Y. Fan, S. Lyu, Y. Ying, and B. Hu. Learning with average top-k loss. In *Advances in Neural Information Processing Systems*, pages 497–505, 2017.
- [16] G. Gidel, H. Berard, G. Vignoud, P. Vincent, and S. Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *International Conference on Learning Representations*, 2019.
- [17] G. Gidel, R. A. Hemmat, M. Pezeshki, R. Le Priol, G. Huang, S. Lacoste-Julien, and I. Mitliagkas. Negative momentum for improved game dynamics. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1802–1811, 2019.
- [18] I. Goodfellow. Nips 2016 tutorial: Generative adversarial networks. *arXiv:1701.00160*, 2016.
- [19] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [20] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.
- [21] E. Y. Hamedani and N. S. Aybat. A primal-dual algorithm for general convex-concave saddle point problems. *SIAM Journal on Optimization*, 31(2):1299–1329, 2021.
- [22] 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 *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017.
- [23] J. Ho and S. Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pages 4565–4573, 2016.
- [24] C. Jin, P. Netrapalli, and M. I. Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? pages 4880–4889, 2020.
- [25] W. Kong and R. D. Monteiro. An accelerated inexact proximal point method for solving nonconvex-concave min-max problems. *SIAM Journal on Optimization*, 31(4):2558–2585, 2021.
- [26] G. Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- [27] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Master’s thesis, University of Toronto, Canada*, 2009.
- [28] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [29] T. Lin, C. Jin, and M. I. Jordan. Near-optimal algorithms for minimax optimization. In *Conference on Learning Theory*, pages 2738–2779. PMLR, 2020.
- [30] T. Lin, C. Jin, and M. I. Jordan. On gradient descent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pages 6083–6093. PMLR, 2020.
- [31] M. Liu, H. Rafique, Q. Lin, and T. Yang. First-order convergence theory for weakly-convex-weakly-concave min-max problems. *The Journal of Machine Learning Research*, 2021.
- [32] M. Liu, W. Zhang, Y. Mroueh, X. Cui, J. Ross, T. Yang, and P. Das. A decentralized par-

- allel algorithm for training generative adversarial nets. *Advances in Neural Information Processing Systems*, 33:11056–11070, 2020.
- [33] 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*, 2020.
- [34] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- [35] Y. Malitsky and M. K. Tam. A forward-backward splitting method for monotone inclusions without cocoercivity. *SIAM Journal on Optimization*, 30(2):1451–1472, 2020.
- [36] L. Mescheder, A. Geiger, and S. Nowozin. Which training methods for GANs do actually converge? In *International Conference on Machine Learning*, 2018.
- [37] L. Mescheder, S. Nowozin, and A. Geiger. The numerics of GANs. In *Advances in neural information processing systems*, volume 30, 2017.
- [38] M. Mohri, G. Sivek, and A. T. Suresh. Agnostic federated learning. In *International Conference on Machine Learning*, pages 4615–4625. PMLR, 2019.
- [39] B. S. Mordukhovich. *Variational Analysis and Generalized Differentiation I: Basic Theory*, volume 330. Springer Science & Business Media, 2006.
- [40] H. Namkoong and J. C. Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *Advances in Neural Information Processing Systems*, pages 2208–2216, 2016.
- [41] 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.
- [42] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media, 2013.
- [43] 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 *Advances in Neural Information Processing Systems*, pages 14934–14942, 2019.
- [44] 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, 2021.
- [45] N. Papernot, F. Faghri, N. Carlini, I. Goodfellow, R. Feinman, A. Kurakin, C. Xie, Y. Sharma, T. Brown, A. Roy, A. Matyasko, V. Behzadan, K. Hambarzumyan, Z. Zhang, Y.-L. Juang, Z. Li, R. Sheatsley, A. Garg, J. Uesato, W. Gierke, Y. Dong, D. Berthelot, P. Hendricks, J. Rauber, and R. Long. Technical report on the cleverhans v2.1.0 adversarial examples library. *arXiv preprint arXiv:1610.00768*, 2018.
- [46] H. Rafique, M. Liu, Q. Lin, and T. Yang. Non-convex min-max optimization: Provable algorithms and applications in machine learning. *arXiv:1810.02060*, 2018.
- [47] A. Sinha, H. Namkoong, and J. Duchi. Certifiable distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018.
- [48] C. Song, Z. Zhou, Y. Zhou, Y. Jiang, and Y. Ma. Optimistic dual extrapolation for coherent non-monotone variational inequalities. In *Advances in Neural Information Processing Systems*, volume 33, pages 14303–14314, 2020.
- [49] K. K. Thekumparampil, P. Jain, P. Netrapalli, and S. Oh. Efficient algorithms for smooth minimax optimization. In *Advances in Neural Information Processing Systems*, pages 12680–12691, 2019.
- [50] F. Tramèr, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel. The space of transferable adversarial examples. *arXiv preprint arXiv:1704.03453*, 2017.
- [51] Q. Tran-Dinh, D. Liu, and L. M. Nguyen. Hybrid variance-reduced SGD algorithms for nonconvex-concave minimax problems. *arXiv:2006.15266*, 2020.
- [52] 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*, pages 1–72, 2023.
- [53] Y. Ying, L. Wen, and S. Lyu. Stochastic online AUC maximization. In *Advances in Neural Information Processing Systems*, pages 451–459, 2016.
- [54] G. Zhang, Y. Wang, L. Lessard, and R. B. Grosse. Near-optimal local convergence of alternating gradient descent-ascent for minimax optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 7659–7679. PMLR, 2022.
- [55] R. Zhao. A primal dual smoothing framework for max-structured nonconvex optimization. *arXiv:2003.04375*, 2020.