

SOLVING CONSTRAINED VARIATIONAL INEQUALITIES VIA A FIRST-ORDER INTERIOR POINT-BASED METHOD

Tong Yang*
University of California, Berkeley
pptmiao@berkeley.edu

Michael I. Jordan*
University of California, Berkeley
jordan@cs.berkeley.edu

Tatjana Chavdarova*
University of California, Berkeley
tatjana.chavdarova@berkeley.edu

ABSTRACT

We develop an interior-point approach to solve constrained variational inequality (cVI) problems. Inspired by the efficacy of the *alternating direction method of multipliers* (ADMM) method in the single-objective context, we generalize ADMM to derive a first-order method for cVIs, that we refer to as ADMM-based interior-point method for constrained VIs (ACVI). We provide convergence guarantees for ACVI in two general classes of problems: (i) when the operator is ξ -monotone, and (ii) when it is monotone, some constraints are active and the game is not purely rotational. When the operator is, in addition, L -Lipschitz for the latter case, we match known lower bounds on rates for the gap function of $\mathcal{O}(1/\sqrt{K})$ and $\mathcal{O}(1/K)$ for the last and average iterate, respectively. To the best of our knowledge, this is the first presentation of a *first-order* interior-point method for the general cVI problem that has a global convergence guarantee. Moreover, unlike previous work in this setting, ACVI provides a means to solve cVIs when the constraints are nontrivial. Empirical analyses demonstrate clear advantages of ACVI over common first-order methods. In particular, (i) cyclical behavior is notably reduced as our methods approach the solution from the analytic center, and (ii) unlike projection-based methods that zigzag when near a constraint, ACVI efficiently handles the constraints.

1 INTRODUCTION

We are interested in the *constrained variational inequality* problem (Stampacchia, 1964):

$$\text{find } \mathbf{x}^* \in \mathcal{X} \text{ s.t. } \langle \mathbf{x} - \mathbf{x}^*, F(\mathbf{x}^*) \rangle \geq 0, \quad \forall \mathbf{x} \in \mathcal{X}, \quad (\text{cVI})$$

where \mathcal{X} is a subset of the Euclidean n -dimensional space \mathbb{R}^n , and where $F : \mathcal{X} \mapsto \mathbb{R}^n$ is a continuous map. Finding (an element of) the solution set $\mathcal{S}_{\mathcal{X}, F}^*$ of cVI is a key problem in multiple fields such as economics and game theory. More pertinent to machine learning, CVIs generalize standard single-objective optimization, complementarity problems (Cottle & Dantzig, 1968), zero-sum games (von Neumann & Morgenstern, 1947; Rockafellar, 1970) and multi-player games. For example, solving cVI is the optimization problem underlying *reinforcement learning* (e.g., Omidshafiei et al., 2017)—and *generative adversarial networks* (Goodfellow et al., 2014). Moreover, even when training one set of parameters with one loss f , that is $F(\mathbf{x}) \equiv \nabla_{\mathbf{x}} f(\mathbf{x})$, a natural way to improve the model’s robustness in some regard is to introduce an adversary to perturb the objective or the input, or to consider the worst sample distribution of the empirical objective. As has been noted in many problem domains, including robust classification (Mazuelas et al., 2020), adversarial training (Szegedy et al., 2014), causal inference (Christiansen et al., 2020), and robust objectives (e.g., Rothenhäusler et al., 2018), this leads to a min-max structure, which is an instance of the cVI problem. To see this, consider two sets of parameters (agents), $\mathbf{x}_1 \in \mathcal{X}_1$ and $\mathbf{x}_2 \in \mathcal{X}_2$, that share a loss/utility function, $f : \mathcal{X}_1 \times \mathcal{X}_2 \rightarrow \mathbb{R}$, which the first agent aims to minimize

*All authors contributed equally. Link to source code: <https://github.com/Chavdarova/ACVI>.




Figure 1: ACVI (Algorithm 1) and EG iterates—depicted in *red* and *green*, resp.—on the game:

$$\min_{x_1 \in \mathbb{R}_+} \max_{x_2 \in \mathbb{R}_+} 0.05 \cdot x_1^2 + x_1 x_2 - 0.05 \cdot x_2^2.$$

The constraints are depicted with dashed lines and the iterates with circles. ACVI gets close to the Nash Equilibrium (\star) in a single step, whereas EG zigzags when hitting a constraint. The remaining commonly used methods—GDA, OGD, and LA-GDA—perform similarly to EG, see App. E.

and the second agent aims to maximize. Then the problem is to find a *saddle point* of f , i.e., a point (x_1^*, x_2^*) such that $f(x_1^*, x_2) \leq f(x_1^*, x_2^*) \leq f(x_1, x_2^*)$. This corresponds to a *cVI* with $F(x) \equiv [\nabla_{x_1} f(x_1, x_2) \quad -\nabla_{x_2} f(x_1, x_2)]^\top$.

Solving cVIs is significantly more challenging than single-objective optimization problems, due to the fact that F is a general vector field, leading to “rotational” trajectories in parameter space (App. A). In response, the development of efficient algorithms with provable convergence has recently been the focus of interest in machine learning and optimization, particularly in the unconstrained setting, where $\mathcal{X} \equiv \mathbb{R}^n$ (e.g., Tseng, 1995; Daskalakis et al., 2018; Mokhtari et al., 2019; 2020; Golowich et al., 2020b; Azizian et al., 2020; Chavdarova et al., 2021a; Gorbunov et al., 2022; Bot et al., 2022).

In many applications, however, we have *constraints* on (part of) the decision variable x , that is, \mathcal{X} is often a strict subset of \mathbb{R}^n . As an example, let us revisit the aforementioned distributionally robust prediction problem: consider a linear setting (cf. Eq. 1 in Rothenhäusler et al., 2018) and class of parametrized distributions $\Delta \equiv \{w \in \mathbb{R}^d | w \geq 0, e^\top w = 1\}$, where $e \in \mathbb{R}^d$ is a vector of all ones. Thus, the robust problem is: $\min_{x \in \mathbb{R}^n} \max_{w \in \mathbb{R}^d} w^\top (y - Dx)$, *subject to* $w \geq 0, e^\top w = 1$, where $D \in \mathbb{R}^{d \times n}$ contains d samples of an n -dimensional covariate vector, and $y \in \mathbb{R}^d$ is the vector of target variables (the constraint $w \leq 1$ is implied). This illustrates that given a standard minimization problem, its robustification immediately leads to an instance of the *cVI* problem; see further examples in § 5. Additional example applications include (i) machine learning applications in business, finance, and economics where often the sum of the decision variables—representing, for example, *resources*—cannot exceed a specific value, (ii) contract theory (e.g. §2.3.2 in (Bates et al., 2022) where one player is the parameters of a probability distribution as above), and (iii) solving optimal control problems numerically, among others.

Significantly fewer works address the convergence of first-order optimization methods in the constrained setting; see § 2 for an overview. Recently, Cai et al. (2022) established a convergence rate for the *projected extragradient method* (Korpelevich, 1976), when F is monotone and Lipschitz (see § 3 for definitions). However, (i) the proof that the authors presented is computer-assisted, which makes it hard to interpret and of limited usefulness for inspiring novel (e.g., accelerated) methods, and (ii) the considered setting assumes the projection is fast to compute and thus ignores the projection in the rate. The latter assumption only holds in rare cases when the constraints are relatively simple so that operations such as clipping suffice. However, when the inequality and/or equality constraints are of a general form, each EG update requires *two* projections (see App. A.4). Each projection requires solving a new/separate *constrained* optimization problem, which if given general constraints implies the need for a *second-order method* as explained next.

Interior point (IP) methods are the de facto family of iterative algorithms for constrained optimization. These methods enjoy well-established guarantees and theoretical understanding in the context of single-objective optimization [see, e.g., Boyd & Vandenberghe (2004), Ch.11, Megiddo (1989), Wright (1997)], and have extensions to a wide range of problem settings (e.g., Tseng, 1993; Nesterov & Nemirovski, 1994; Nesterov & Todd, 1998; Renegar, 2001; Wright, 2001). They build on a natural idea of solving a simplified homotopic problem that makes it possible to “smoothly” transition to the original complex problem; see § 3.1. Several works extend IP methods to *cVI*, by applying the second-order Newton method to a modified *Karush-Kuhn-Tucker* (KKT) system appropriate for the *cVI* (Ralph & Wright, 2000; Qi & Sun, 2002; Fan & Yan, 2010; Monteiro & Pang, 1996; Chen et al., 1998). Many of these approaches, however, rely on strong assumptions—see § 2. Moreover, although these methods enjoy fast convergence in terms of the number of iterations, each iteration involves the computation of the Jacobian of F (or Hessian when $F \equiv \nabla f(x)$) which quickly becomes prohibitive for large dimension n of x . Hence first-order methods are preferred in practice.

We are currently missing a *first-order* optimization method for solving *cVI* with *general* constraints. Accordingly, in this paper, we focus on the following open question:

*Can we derive **first-order** algorithms for the *cVI* problem that (i) can be applied when **general** constraints are given, and that (ii) have **global** convergence guarantees?*

In this paper, we develop precisely such a method. To mitigate the computational burden of second-derivative computation, we replace the Newton step of the traditional IP methods with the *alternating direction method of multipliers* (ADMM) method. ADMM was designed with a different purpose: it is applicable only when the objective is separable into two or more different functions whose arguments are non disjoint—see § 3.1 for full description—and can be seen as equivalent to *Douglas–Rachford* operator splitting (Douglas & Rachford, 1956) applied in the *dual* space (see e.g. Lin et al., 2022). ADMM owes its popularity primarily to its computational efficiency (Boyd et al., 2011) for large-scale machine learning problems and its fast convergence in some machine-learning settings (e.g., Nishihara et al., 2015). The core idea of our approach is to reformulate the original *cVI* problem with equality and inequality constraints via the KKT conditions, so as to apply ADMM in such a way that the subproblems of the resulting algorithm have desirable properties (see § 4.1). That is, by generalizing the technique underlying ADMM, we derive a novel first-order algorithm for solving monotone VIs with very general constraints. Furthermore, this framework can be used to design novel algorithms for solving *cVIs*; see App. C.

Our contributions can be summarized as follows:

- Based on the KKT system for the constrained VI problem and the ADMM technique, we derive an algorithm that we refer to as the *ADMM-based Interior Point Method for Constrained VIs* (ACVI)—see § 4.1 and Algorithm 1.
- We prove the global convergence of ACVI given two sets of assumptions: (i) when F is ξ -monotone, and (ii) when it is monotone, the constraints are active at the solution, and the game is not purely rotational. By further assuming F is a Lipschitz operator, we upper bound the rate of decrease of the gap function and we match the known lower bound for the gap function of $\mathcal{O}(1/\sqrt{K})$ for the last iterate—see § 4.2.
- Empirically, we document two notable advantages of ACVI over popular projection-based saddle-point methods: (i) the ACVI iterates exhibit significantly reduced rotations, as they approach the solution from the analytic center, and (ii) while projection-based methods show extensive zigzagging when hitting a constraint, ACVI avoids this, resulting in more efficient updates—§ 5.

Our convergence guarantees are parameter-free, meaning these do not require a priori knowledge of the constants of the problem (such as the Lipschitz constant), and, interestingly, the convergence guarantee does not require that F is Lipschitz. This assumption is solely used to express the rate of decrease of the gap function (in contrast to the extragradient method (Korpelevich, 1976) where such an assumption is necessary to show convergence. To the best of our knowledge, the proposed ACVI method is the first first-order IP algorithm for VIs with a global convergence proof.

2 RELATED WORK

Unconstrained VIs: methods and guarantees. Apart from the standard gradient descent ascent (GDA) method, among the most commonly used methods for VI optimization are the extragradient method (EG, Korpelevich, 1976), optimistic GDA (OGDA, Popov, 1980), and the lookahead method (LA, Zhang et al., 2019; Chavdarova et al., 2021b). (See App. A for a full description). In contrast to gradient fields (as in a single-objective setting), when F is a general vector field, the last iterate can be far from the solution even though the average iterate converges to it (Daskalakis et al., 2018; Chavdarova et al., 2019). This is problematic since it implies that the average convergence guarantee is weaker in the sense that it may not extend to more general setups where we can no longer rely on the convexity of \mathcal{X} . Golowich et al. (2020b;a) provided a last-iterate lower bound of $\mathcal{O}(\frac{1}{\tilde{p}\sqrt{K}})$ for the broad class of \tilde{p} -stationary canonical linear iterative (\tilde{p} -SCLI) first-order methods (Arjevani et al., 2016). An extensive line of further work has provided guarantees for the last iterate for other problem classes. For the general monotone VI (MVI) class, the following \tilde{p} -SCLI methods come with guarantees that match the lower bound: (i) Golowich et al. (2020b) obtained a rate in terms of

the gap function relying on first- and second-order smoothness of F , and Gorbunov et al. (2022) obtained a rate of $\mathcal{O}(\frac{1}{K})$ in terms of reducing the squared norm of the operator relying on first-order smoothness of F (Assumption 1), using a computer-assisted proof, and (ii) Golowich et al. (2020b) and Chavdarova et al. (2021a) provided the best iterate rate for OGDA.

Constrained zero-sum and VI classes of problems. Gidel et al. (2017b) extended the Frank-Wolfe (Frank & Wolfe, 1956; Jaggi, 2013; Lacoste-Julien & Jaggi, 2015) method—also known as the *conditional gradient* (V.F Demyanov, 1970)—to solve a subclass of cVI, specifically constrained zero-sum problems. This extension was carried out under a strong convex-concavity assumption and also under the assumption that the constraint set is *strongly* convex; that is, it has sublevel sets that are strongly convex functions (Vial, 1983). Daskalakis & Panageas (2019) provided an asymptotic proof for the last iterate for zero-sum convex-concave constrained problems for the *optimistic multiplicative weights update* (OMWU) method. Wei et al. (2021) focused on OGDA and OMWU in the constrained setting and provided convergence rates for bilinear games over the simplex. In her seminal work, Korpelevich (1976) proposed the classical (projected) extragradient method (EG)—see App. A—and proved its convergence for monotone (c)VIs with an L -Lipschitz operator, and Cai et al. (2022) established a rate with respect to the gap function using a computer-aided proof. Tseng (1995) built on (Pang, 1987) and provided a linear convergence rate for EG in the setting of *strongly* monotone F , whereas Malitsky (2015) focused on the same constrained setting but on the projected reflected gradient method. Diakonikolas (2020) obtained parameter-free guarantee for Halpern iteration (Halpern, 1967) for cocoercive operators. Goffin et al. (1997) described a second-order cutting-plane method for solving pseudomonotone VIs with linear inequalities.

Interior point (IP) methods in single-objective and VI settings. Traditionally IP methods primarily express the inequality constraints by augmenting the objective with a log-barrier penalty (see § 3.1), and then use Newton’s method to solve the subproblem (Boyd & Vandenberghe, 2004). The latter involves computing either the inverse of a large matrix or a Cholesky decomposition and yet it can be highly efficient in low dimensions as it requires only a few iterations to converge. When the dimensionality of the variable is large, however, the computation becomes infeasible. Among other IP variants that address this issue, Lin et al. (2018) replaced the Newton step with the ADMM method, which is known to be highly scalable in terms of the dimension (Boyd et al., 2011). In the context of cVIs, a few works apply IP methods, mostly Newton-based (e.g., Nesterov & Nemirovski, 1994, Chapter 7). Monteiro & Pang (1996) analyze path-following IP methods for complementarity problems, which are a subclass of cVI, using local homeomorphic maps. Chen et al. (1998) provided a superlinear global convergence rate of the smoothing Newton method when F is semi-smooth for *box constrained* VIs. Similarly, Qi & Sun (2002); Qi et al. (2000) focused on the smoothing Newton method and provided the rate for the outer loop. Ralph & Wright (2000) showed superlinear convergence for MVI problems under inequality constraints, under the following set of assumptions: (i) existence of a *strictly* complementary solution, (ii) full rank of the Jacobian of the active constraints at the solution, and (iii) twice differentiable constraints. They provided a *local* convergence rate. Fan & Yan (2010) considered inequality constraints and proposed a second-order Newton-based method that has global convergence guarantees under certain conditions. A rate was not provided.

3 PRELIMINARIES

Notation. Bold small and capital letters denote vectors and matrices, respectively. Sets are denoted with curly capital letters, e.g., S . The Euclidean norm of v is denoted by $\|v\|$, and the inner product in Euclidean space with $\langle \cdot, \cdot \rangle$. With \odot we denote element-wise product. We let $[n]$ denote $\{1, \dots, n\}$ and let e denote vector of all 1’s. We let $x \perp y$ denote x and y are perpendicular.

In the remainder of the paper, we consider a general setting in which the constraint set $\mathcal{C} \subseteq \mathcal{X}$ is defined as an intersection of finitely many inequalities and linear equalities:

$$\mathcal{C} = \{x \in \mathbb{R}^n \mid \varphi_i(x) \leq 0, i \in [m], \mathbf{C}x = \mathbf{d}\}, \quad (\text{CS})$$

where each $\varphi_i : \mathbb{R}^n \mapsto \mathbb{R}$, $\mathbf{C} \in \mathbb{R}^{p \times n}$, $\mathbf{d} \in \mathbb{R}^p$, where we assume $\text{rank}(\mathbf{C}) = p$. For brevity, with φ we denote the concatenated $\varphi_i(\cdot), i \in [m]$, and in the remainder of the paper, each $\varphi_i \in C^1(\mathbb{R}^n), i \in [m]$ and is convex. For convenience we denote:

$$\mathcal{C}_{\leq} \triangleq \{x \in \mathbb{R}^n \mid \varphi(x) \leq \mathbf{0}\}, \quad \mathcal{C}_{<} \triangleq \{x \in \mathbb{R}^n \mid \varphi(x) < \mathbf{0}\}, \quad \text{and} \quad \mathcal{C}_{=} \triangleq \{y \in \mathbb{R}^n \mid \mathbf{C}y = \mathbf{d}\};$$

thus the *relative interior* of \mathcal{C} is $\text{int } \mathcal{C} \triangleq \mathcal{C}_{<} \cap \mathcal{C}_{=}$, and we consider $\text{int } \mathcal{C} \neq \emptyset$ and \mathcal{C} is compact.

In the following, we list the definitions and assumptions we refer to later on.

Definition 1 ((strong/ ξ) monotonicity). *An operator $F : \mathcal{X} \supseteq \mathcal{S} \rightarrow \mathbb{R}^n$ is monotone on \mathcal{S} if: $\langle \mathbf{x} - \mathbf{x}', F(\mathbf{x}) - F(\mathbf{x}') \rangle \geq 0, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{S}$. F is said to be ξ -monotone on \mathcal{S} iff there exist $c > 0$ and $\xi > 1$ such that $\langle \mathbf{x} - \mathbf{x}', F(\mathbf{x}) - F(\mathbf{x}') \rangle \geq c \|\mathbf{x} - \mathbf{x}'\|^\xi$, for all $\mathbf{x}, \mathbf{x}' \in \mathcal{S}$. Finally, F is μ -strongly monotone on \mathcal{S} if there exists $\mu > 0$, such that $\langle \mathbf{x} - \mathbf{x}', F(\mathbf{x}) - F(\mathbf{x}') \rangle \geq \mu \|\mathbf{x} - \mathbf{x}'\|^2$, for all $\mathbf{x}, \mathbf{x}' \in \mathcal{S}$. Moreover, we say that an operator F is star-monotone, star- ξ -monotone or star-strongly-monotone (on \mathcal{S}) if the respective definition holds for $\mathbf{x}' \equiv \mathbf{x}^*$, where $\mathbf{x}^* \in \mathcal{S}_{\mathcal{S}, F}^*$.*

Note that the “star-” definitions are weaker relative to their respective non-star counterparts. The above definition holds similarly for unconstrained VIs, by setting $\mathcal{S} \equiv \mathbb{R}^n$. The analog for cVI of the function values used as a performance measure for convergence rates in convex optimization is the *gap function* (a.k.a., the *optimality gap* or *primal gap*), defined next.

Definition 2 (gap function). *Given a candidate point $\mathbf{x}' \in \mathcal{X}$ and a map $F : \mathcal{X} \supseteq \mathcal{S} \rightarrow \mathbb{R}^n$ where \mathcal{S} is compact, the gap function $\mathcal{G} : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as $\mathcal{G}(\mathbf{x}', \mathcal{S}) \triangleq \max_{\mathbf{x} \in \mathcal{S}} \langle F(\mathbf{x}'), \mathbf{x}' - \mathbf{x} \rangle$.*

Note that the gap function requires \mathcal{S} to be compact in order to be defined (as otherwise, it can be infinite). We will rely on the following assumption to express our rates in terms of the gap function.

Assumption 1 (first-order smoothness). *Let $F : \mathcal{X} \supseteq \mathcal{S} \rightarrow \mathbb{R}^n$ be an operator, we say that F satisfies L -first-order smoothness on \mathcal{S} , or L -smoothness, if F is an L -Lipschitz map; that is, there exists $L > 0$ such that $\|F(\mathbf{x}) - F(\mathbf{x}')\| \leq L \|\mathbf{x} - \mathbf{x}'\|$, for all $\mathbf{x}, \mathbf{x}' \in \mathcal{S}$.*

As an informal summary, a solution existence guarantee follows when \mathcal{X} is compact; see Chapter 2.2 of (Facchinei & Pang, 2003), and App. A.2.

3.1 RELEVANT PATH-FOLLOWING INTERIOR-POINT METHODS AND ADMM

In this section, we overview the interior-point approach to single-objective optimization, focusing on aspects that are most relevant to our proposed method. Consider the following problem:

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{s.t.} \quad \varphi(\mathbf{x}) \leq \mathbf{0} \quad \text{and} \quad \mathbf{C}\mathbf{x} = \mathbf{d}, \quad (\text{cCVX})$$

where $f, \varphi_i : \mathbb{R}^n \rightarrow \mathbb{R}$ are convex and continuously differentiable, $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{C} \in \mathbb{R}^{p \times n}$, and $\mathbf{d} \in \mathbb{R}^p$. IP methods solve problem (cCVX) by reducing it to a sequence of linear equality-constrained problems via a logarithmic barrier (see, e.g., Boyd & Vandenberghe, 2004, Chapter 11):

$$\min_{\mathbf{x}} f(\mathbf{x}) - \mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{x})) \quad \text{s.t.} \quad \mathbf{C}\mathbf{x} = \mathbf{d}, \quad \text{with} \quad \mu > 0. \quad (\text{l-cCVX})$$

Assume that (l-cCVX) has a solution for each $\mu > 0$, and let \mathbf{x}^μ denote the solution of (l-cCVX) for a given μ . The *central path* of (l-cCVX) is defined as the set of points $\mathbf{x}^\mu, \mu > 0$. Note that $\mathbf{x}^\mu \in \mathbb{R}^n$ is a strictly feasible point of (cCVX) as it satisfies $\varphi(\mathbf{x}^\mu) < \mathbf{0}$ and $\mathbf{C}\mathbf{x}^\mu = \mathbf{d}$.

Alternating direction method of multipliers (ADMM) method. ADMM (Glowinski & Marroco, 1975; Gabay & Mercier, 1976; Lions & Mercier, 1979; Glowinski & Le Tallec, 1989) is a gradient-based algorithm for convex optimization problems that splits the objective into subproblems each of which is easier to solve. Its popularity is due to its computational scalability (Boyd et al., 2011). Consider a problem of the following form:

$$\min_{\mathbf{x}, \mathbf{y}} f(\mathbf{x}) + g(\mathbf{y}) \quad \text{s.t.} \quad \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y} = \mathbf{b}, \quad (\text{ADMM-Pr})$$

where $f, g : \mathbb{R}^n \rightarrow \mathbb{R}$ are convex, $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n' \times n}$, and $\mathbf{b} \in \mathbb{R}^{n'}$. The augmented Lagrangian function, $\mathcal{L}_\beta(\cdot)$, of the (ADMM-Pr) problem is:

$$\mathcal{L}_\beta(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) = f(\mathbf{x}) + g(\mathbf{y}) + \langle \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y} - \mathbf{b}, \boldsymbol{\lambda} \rangle + \frac{\beta}{2} \|\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y} - \mathbf{b}\|^2, \quad (\text{AL-CVX})$$

where $\beta > 0$ is referred to as the *penalty parameter*. If the augmented Lagrangian method is used to solve (AL-CVX), at each step k we have: $\mathbf{x}_{k+1}, \mathbf{y}_{k+1} = \underset{\mathbf{x}, \mathbf{y}}{\operatorname{argmin}} \mathcal{L}_\beta(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}_k)$ and $\boldsymbol{\lambda}_{k+1} = \boldsymbol{\lambda}_k + \beta(\mathbf{A}\mathbf{x}_{k+1} + \mathbf{B}\mathbf{y}_{k+1} - \mathbf{b})$, where the latter step is gradient ascent on the dual. In contrast, ADMM updates \mathbf{x} and \mathbf{y} in an alternating way as follows:

$$\begin{aligned} \mathbf{x}_{k+1} &= \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{L}_\beta(\mathbf{x}, \mathbf{y}_k, \boldsymbol{\lambda}_k), \\ \mathbf{y}_{k+1} &= \underset{\mathbf{y}}{\operatorname{argmin}} \mathcal{L}_\beta(\mathbf{x}_{k+1}, \mathbf{y}, \boldsymbol{\lambda}_k), \\ \boldsymbol{\lambda}_{k+1} &= \boldsymbol{\lambda}_k + \beta(\mathbf{A}\mathbf{x}_{k+1} + \mathbf{B}\mathbf{y}_{k+1} - \mathbf{b}). \end{aligned} \quad (\text{ADMM})$$

4 ACVI: FIRST-ORDER ADMM-BASED IP METHOD FOR CONSTRAINED VIs

4.1 DERIVING THE ACVI ALGORITHM

In this section, we derive an interior-point method for the **cVI** problem that we refer to as ACVI (ADMM-based interior problem for constrained VIs). We first restate the **cVI** problem in a form that will allow us to derive an interior-point procedure. By the definition of **cVI** it follows (see §1.3 in [Facchinei & Pang, 2003](#)) that:

$$\mathbf{x} \in \mathcal{S}_{\mathcal{C},F}^* \Leftrightarrow \begin{cases} \mathbf{w} = \mathbf{x} \\ \mathbf{x} = \underset{\mathbf{z}}{\operatorname{argmin}} F(\mathbf{w})^\top \mathbf{z} \\ \text{s.t. } \varphi(\mathbf{z}) \leq \mathbf{0} \\ \mathbf{C}\mathbf{z} = \mathbf{d} \end{cases} \Leftrightarrow \begin{cases} F(\mathbf{x}) + \nabla\varphi^\top(\mathbf{x})\boldsymbol{\lambda} + \mathbf{C}^\top\boldsymbol{\nu} = \mathbf{0} \\ \mathbf{C}\mathbf{x} = \mathbf{d} \\ \mathbf{0} \leq \boldsymbol{\lambda} \perp \varphi(\mathbf{x}) \leq \mathbf{0}, \end{cases} \quad (\text{KKT})$$

where $\boldsymbol{\lambda} \in \mathbb{R}^m$ and $\boldsymbol{\nu} \in \mathbb{R}^p$ are dual variables, and \perp denotes perpendicular. Recall that we assume that $\operatorname{int} \mathcal{C} \neq \emptyset$, thus, by the Slater condition (using the fact that $\varphi_i(\mathbf{x}), i \in [m]$ are convex) and the KKT conditions, the second equivalence holds, yielding the KKT system of **cVI**. Note that the above equivalence also guarantees the two solutions coincide; see [Facchinei & Pang \(2003, Prop. 1.3.4 \(b\)\)](#). Analogous to the method described in §3, we add a log-barrier term to the objective to remove the inequality constraints and obtain the following modified version of (KKT):

$$\begin{cases} \mathbf{w} = \mathbf{x} \\ \mathbf{x} = \underset{\mathbf{z}}{\operatorname{argmin}} F(\mathbf{w})^\top \mathbf{z} - \mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{z})) \\ \text{s.t. } \mathbf{C}\mathbf{z} = \mathbf{d} \end{cases} \Leftrightarrow \begin{cases} F(\mathbf{x}) + \nabla\varphi^\top(\mathbf{x})\boldsymbol{\lambda} + \mathbf{C}^\top\boldsymbol{\nu} = \mathbf{0} \\ \boldsymbol{\lambda} \odot \varphi(\mathbf{x}) + \mu\mathbf{e} = \mathbf{0} \\ \mathbf{C}\mathbf{x} - \mathbf{d} = \mathbf{0} \\ \varphi(\mathbf{x}) < \mathbf{0}, \boldsymbol{\lambda} > \mathbf{0}, \end{cases} \quad (\text{KKT-2})$$

with $\mu > 0$, $\mathbf{e} \triangleq [1, \dots, 1]^\top \in \mathbb{R}^m$. Again, the equivalence holds by the KKT and the Slater condition. We derive the update rule at step k via the following subproblem: $\min_{\mathbf{x}} F(\mathbf{w}_k)^\top \mathbf{x} - \mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{x}))$, s.t. $\mathbf{C}\mathbf{x} = \mathbf{d}$, where we fix $\mathbf{w} = \mathbf{w}_k$. Directly projecting on the equality constraint may cause the vectors to fall out of the domain of the log term. On the other hand, (i) \mathbf{w}_k is a constant vector in this subproblem, and (ii) the objective is split, making ADMM a natural choice to solve the subproblem. Hence, we introduce a new variable $\mathbf{y} \in \mathbb{R}^n$ yielding:

$$\begin{cases} \min_{\mathbf{x}, \mathbf{y}} F(\mathbf{w}_k)^\top \mathbf{x} + \mathbb{1}[\mathbf{C}\mathbf{x} = \mathbf{d}] - \mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{y})) \\ \text{s.t. } \mathbf{x} = \mathbf{y} \end{cases}, \quad \mathbb{1}[\mathbf{C}\mathbf{x} = \mathbf{d}] \triangleq \begin{cases} 0, & \text{if } \mathbf{C}\mathbf{x} = \mathbf{d} \\ +\infty, & \text{if } \mathbf{C}\mathbf{x} \neq \mathbf{d}. \end{cases} \quad (1)$$

Note that $\mathbb{1}[\mathbf{C}\mathbf{x} = \mathbf{d}]$ is a generalized real-valued convex function of \mathbf{x} . We introduce the following:

$$\mathbf{P}_c \triangleq \mathbf{I} - \mathbf{C}^\top(\mathbf{C}\mathbf{C}^\top)^{-1}\mathbf{C}, \quad (\mathbf{P}_c) \quad \text{and} \quad \mathbf{d}_c \triangleq \mathbf{C}^\top(\mathbf{C}\mathbf{C}^\top)^{-1}\mathbf{d}, \quad (\mathbf{d}_c\text{-EQ})$$

where $\mathbf{P}_c \in \mathbb{R}^{n \times n}$ and $\mathbf{d}_c \in \mathbb{R}^n$. The augmented Lagrangian of (1) is thus:

$$\mathcal{L}_\beta(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) = F(\mathbf{w}_k)^\top \mathbf{x} + \mathbb{1}(\mathbf{C}\mathbf{x} = \mathbf{d}) - \mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{y})) + \langle \boldsymbol{\lambda}, \mathbf{x} - \mathbf{y} \rangle + \frac{\beta}{2} \|\mathbf{x} - \mathbf{y}\|^2, \quad (\text{AL})$$

where $\beta > 0$ is the penalty parameter. Finally, using ADMM, we have the following update rule for \mathbf{x} at step k :

$$\mathbf{x}_{k+1} = \arg \min_{\mathbf{x} \in \mathcal{C}_=} \mathcal{L}_\beta(\mathbf{x}, \mathbf{y}_k, \boldsymbol{\lambda}_k) = \arg \min_{\mathbf{x} \in \mathcal{C}_=} \frac{\beta}{2} \left\| \mathbf{x} - \mathbf{y}_k + \frac{1}{\beta} (F(\mathbf{w}_k) + \boldsymbol{\lambda}_k) \right\|^2. \quad (2)$$

This yields the following update for \mathbf{x} :

$$\mathbf{x}_{k+1} = \mathbf{P}_c \left(\mathbf{y}_k - \frac{1}{\beta} (F(\mathbf{w}_k) + \boldsymbol{\lambda}_k) \right) + \mathbf{d}_c. \quad (\text{X-EQ})$$

For \mathbf{y} and the dual variable $\boldsymbol{\lambda}$, we have:

$$\mathbf{y}_{k+1} = \underset{\mathbf{y}}{\operatorname{argmin}} \mathcal{L}_\beta(\mathbf{x}_{k+1}, \mathbf{y}, \boldsymbol{\lambda}_k) = \underset{\mathbf{y}}{\operatorname{argmin}} \left(-\mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{y})) + \frac{\beta}{2} \left\| \mathbf{y} - \mathbf{x}_{k+1} - \frac{1}{\beta} \boldsymbol{\lambda}_k \right\|^2 \right), \quad (\text{Y-EQ})$$

$$\boldsymbol{\lambda}_{k+1} = \boldsymbol{\lambda}_k + \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1}). \quad (\lambda\text{-EQ})$$

Next, we derive the update rule for w . We set w_k to be the solution of the following equation:

$$w + \frac{1}{\beta} P_c F(w) - P_c y_k + \frac{1}{\beta} P_c \lambda_k - d_c = 0. \quad (\text{W-EQ})$$

The following theorem ensures the solution of (W-EQ) exists and is unique, see App. B.1 for proof.

Theorem 1 (W-EQ: solution uniqueness). *If F is monotone on $\mathcal{C}_=$, the following statements hold true for the solution of (W-EQ): (i) it always exists, (ii) it is unique, and (iii) it is contained in $\mathcal{C}_=$.*

Remark 1. *Note that when there are no equality constraints, $\mathcal{C}_=$ becomes the entire space \mathbb{R}^n . Further notice that $w_k = x_{k+1}$, thus it is redundant to state it in the algorithm, and we remove w .*

We summarize the full algorithm as Algorithm 1. For problems such as affine or low-dimensional VIs, or optimization over the probability simplex, (W-EQ) can be solved analytically, such that step 8 is fast to compute. For problems where (W-EQ) is cumbersome to solve analytically—e.g., in GANs—one could use optimization methods for the *unconstrained* case, e.g., EG and GDA, among others. See App. B.5 for further discussion. In the remaining discussion, where clear from context, we drop the superscript from the iterate $x_k^{(t)}$.

Algorithm 1 ACVI pseudocode.

- 1: **Input:** operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$, constraints $Cx = d$ and $\varphi_i(x) \leq 0, i = [m]$, hyperparameters $\mu_{-1}, \beta > 0, \delta \in (0, 1)$, number of outer and inner loop iterations T and K , resp.
 - 2: **Initialize:** $y_0^{(0)} \in \mathbb{R}^n, \lambda_0^{(0)} \in \mathbb{R}^n$
 - 3: $P_c \triangleq I - C^\top(CC^\top)^{-1}C$ where $P_c \in \mathbb{R}^{n \times n}$
 - 4: $d_c \triangleq C^\top(CC^\top)^{-1}d$ where $d_c \in \mathbb{R}^n$
 - 5: **for** $t = 0, \dots, T - 1$ **do**
 - 6: $\mu_t = \delta\mu_{t-1}$
 - 7: **for** $k = 0, \dots, K - 1$ **do**
 - 8: Set $x_{k+1}^{(t)}$ to be the solution of: $x + \frac{1}{\beta} P_c F(x) - P_c y_k^{(t)} + \frac{1}{\beta} P_c \lambda_k^{(t)} - d_c = 0$ (w.r.t. x)
 - 9: $y_{k+1}^{(t)} = \underset{y}{\operatorname{argmin}} - \mu_t \sum_{i=1}^m \log(-\varphi_i(y)) + \frac{\beta}{2} \left\| y - x_{k+1}^{(t)} - \frac{1}{\beta} \lambda_k^{(t)} \right\|^2$
 - 10: $\lambda_{k+1}^{(t)} = \lambda_k^{(t)} + \beta(x_{k+1}^{(t)} - y_{k+1}^{(t)})$
 - 11: **end for**
 - 12: $(y_0^{(t+1)}, \lambda_0^{(t+1)}) \triangleq (y_K^{(t)}, \lambda_K^{(t)})$
 - 13: **end for**
-

4.2 CONVERGENCE ANALYSIS

We consider two broad classes of problems. The first class assumes that F is ξ -monotone on $\mathcal{C}_=$ —a stronger assumption than monotonicity, yet weaker than strong monotonicity. The second setup requires that (i) F is monotone, (ii) the constraints are active at the solution, and (iii) F is not purely rotational. Note that (iii) is weaker than requiring that the active constraints at the solution form an acute angle with the operator; in other words, given the latter, the former holds due to monotonicity of F . (See App. B). Note that (iii) is not strong, as purely rotational games occur “almost never” in a Baire category sense (Kupka, 1963; Smale, 1963; Balduzzi et al., 2018; Hsieh et al., 2021). The proofs of the main theorems use the following lemma.

Lemma 1 (Upper bound for $\mathcal{G}(\cdot)$). *When F is L -Lipschitz on $\mathcal{C}_=$ —as per Assumption 1—we have that any iterate x_k produced by Algorithm 1 satisfies $\mathcal{G}(x_k, \mathcal{C}) \leq M_0 \|x_k - x^*\|$, where $M_0 > 0$ depends linearly on L , and $x^* \in \mathcal{S}_{\mathcal{C}, F}^*$.*

To state the results we define the following sets. For $r, s > 0$, let $\hat{\mathcal{C}}_r \triangleq \{x \in \mathbb{R}^n | Cx = d, \varphi(x) \leq re\}$, and similarly let $\tilde{\mathcal{C}}_s \triangleq \{x \in \mathbb{R}^n | \|Cx - d\| \leq s, \varphi(x) \leq 0\}$. We have the following.

Theorem 2 (Last and average iterate convergence for star- ξ -monotone operator). *Given an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$ monotone on $\mathcal{C}_=$ (Def. 1), assume that either F is strictly monotone on \mathcal{C} or one of φ_i is strictly convex. Assume there exists $r > 0$ or $s > 0$ such that F is star- ξ -monotone on either $\hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$. Let $x_K^{(t)}$ and $\hat{x}_K^{(t)} \triangleq \frac{1}{K} \sum_{k=1}^K x_k^{(t)}$ denote the last and average iterate of Algorithm 1, respectively, run with sufficiently small μ_{-1} . Then for all $t \in [T]$, we have that:*

1. $\|\mathbf{x}_K^{(t)} - \mathbf{x}^*\| \leq \mathcal{O}(\frac{1}{K^{1/(2\xi)}})$.
2. If in addition F is ξ -monotone on $\mathcal{C}_=$, we have $\|\hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^*\| \leq \mathcal{O}(\frac{1}{K^{1/\xi}})$.
3. Moreover, if F is L -Lipschitz on $\mathcal{C}_=$ —as per Assumption 1—the same corresponding upper bounds hold for $\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C})$ and $\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C})$; that is, $\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}(\frac{L}{K^{1/(2\xi)}})$ and $\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}(\frac{L}{K^{1/\xi}})$.

Remark 2. Note that the convergence guarantee does not rely on Assumption 1, and it is solely used to relate the rate to the gap function. Also, note that μ_{-1} does not impact the convergence rate. Moreover, for simplicity we state the result with sufficiently small μ_{-1} , however, the proof extends to any $\mu_{-1} > 0$. That is, the above result can be made parameter-free; see App. B.4.

Theorem 3 (Last and average iterate convergence for monotone operator). Given an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$, assume (i) F is monotone on $\mathcal{C}_=$, and (ii) either F is strictly monotone on \mathcal{C} or one of φ_i is strictly convex, and (iii) $\inf_{\mathbf{x} \in \mathcal{S} \setminus \{\mathbf{x}^*\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0$, where $\mathcal{S} \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$. Let $\mathbf{x}_K^{(t)}$

and $\hat{\mathbf{x}}_K^{(t)} \triangleq \frac{1}{K} \sum_{k=1}^K \mathbf{x}_k^{(t)}$ denote the last and average iterate of Algorithm 1, respectively, run with sufficiently small μ_{-1} . Then for all $t \in [T]$, we have that:

1. $\|\mathbf{x}_K^{(t)} - \mathbf{x}^*\| \leq \mathcal{O}(\frac{1}{\sqrt{K}})$.
2. If in addition $\inf_{\mathbf{x} \in \mathcal{S} \setminus \{\mathbf{x}^*\}} F(\mathbf{x}^*)^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0$ (with $\mathcal{S} \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$), then $\|\hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^*\| \leq \mathcal{O}(\frac{1}{K})$.
3. Moreover, if F is L -Lipschitz on $\mathcal{C}_=$ —as per Assumption 1—the same corresponding upper bounds hold for $\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C})$ and $\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C})$, that is, $\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}(\frac{L}{\sqrt{K}})$ and $\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}(\frac{L}{K})$.

Assumption (iii) in Theorem 3 requires the angle of $F(\mathbf{x})$ and $\mathbf{x} - \mathbf{x}^*$ to be acute on $\mathcal{S} \setminus \{\mathbf{x}^*\}$, where $\mathcal{S} = \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$. For example, when there are no equality constraints, Assumption (iii) becomes $\inf_{\mathbf{x} \in \mathcal{C} \setminus \{\mathbf{x}^*\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0$. From (cVI) and by the monotonicity of F , we can see that for any

point $\mathbf{x} \in \mathcal{C} \setminus \{\mathbf{x}^*\}$, the angle between $F(\mathbf{x})$ and $\mathbf{x} - \mathbf{x}^*$ is always less than or equal to $\pi/2$. And assumption (iii) requires that $F(\mathbf{x}^*) \neq \mathbf{0}$, which means some constraints are active at \mathbf{x}^* , and $\exists \theta \in (0, \pi/2)$ s.t. for any $\mathbf{x} \in \mathcal{C} \setminus \{\mathbf{x}^*\}$, the angle between $F(\mathbf{x})$ and $\mathbf{x} - \mathbf{x}^*$ is upper bounded by θ .

Remark 3. Our proofs rely on the existence of the central path—see Appendix A. Note that since \mathcal{C} is compact, it suffices that either: (i) F is strictly monotone on \mathcal{C} , or that (ii) one of the inequality constraints φ_i is strictly convex for the central path to exist (Facchinei & Pang, 2003, Corollary 11.4.24). Thus, if F is ξ -monotone on \mathcal{C} , then the central path exists. However, to relax the former assumption, notice that—by the compactness of \mathcal{C} —there exists a sufficiently large M such that for any $\mathbf{x} \in \mathcal{C}$, $\mathbf{x}^\top \mathbf{x} \leq M$. Thus, one can add a strictly convex inequality constraint $\varphi_{m+1}(\mathbf{x})$ —e.g., $\mathbf{x}^\top \mathbf{x} - M \leq 0$ —and the solution set remains intact. That is, as μ tends to 0 the original problem is recovered. This ensures the existence of the central path without changing the original problem.

5 EXPERIMENTS

Problems. To study the empirical performance of ACVI we use the following 2D problems: (i) *cBG*: the common bilinear game, constrained on \mathbb{R}_+ for the two players, stated in Fig. 1, (ii) *Von Neumann’s ratio game* (Von Neumann, 1971; Daskalakis et al., 2020; Diakonikolas et al., 2021), (iii) *Forsaken game* (Hsieh et al., 2021)—which exhibits *limit cycles*, as well as (iv) *toy GAN*—used in (Daskalakis et al., 2018; Antonakopoulos et al., 2021). Note that these are known to be challenging problems in the literature, and interestingly the latter three are non-monotone, going beyond the assumptions that we made in our theoretical results. We also consider the following higher-dimension bilinear game on the probability simplex, with $\eta \in (0, 1)$, $n = 1000$:

$$\min_{\mathbf{x}_1 \in \Delta} \max_{\mathbf{x}_2 \in \Delta} \eta \mathbf{x}_1^\top \mathbf{x}_1 + (1 - \eta) \mathbf{x}_1^\top \mathbf{x}_2 - \eta \mathbf{x}_2^\top \mathbf{x}_2; \quad \Delta = \{\mathbf{x}_i \in \mathbb{R}^{500} \mid \mathbf{x}_i \geq \mathbf{0}, \text{ and } \mathbf{e}^\top \mathbf{x}_i = 1\}. \quad (\text{HBG})$$

As GANs on MNIST (Lecun & Cortes, 1998) enjoy well-established metrics, we use this setup and augment it solely with linear inequalities. We implement the baselines with the greedy projection algorithm—see App. D.3 for details—hence these baselines will be slower when equality constraints are also given. App. E lists additional experiments, including on Fashion-MNIST (Xiao et al., 2017).

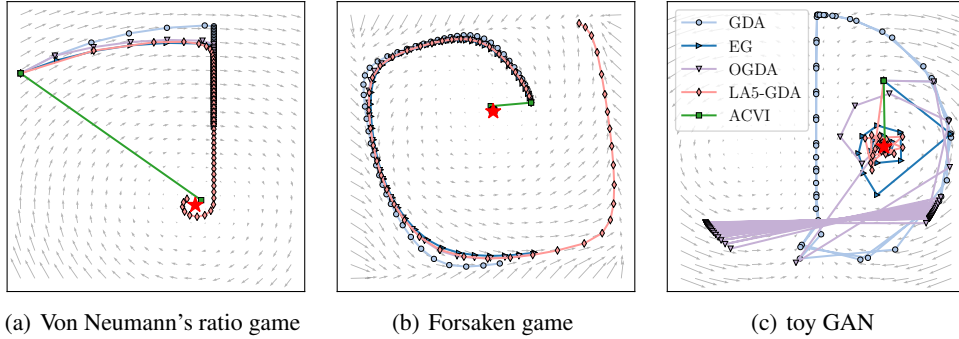


Figure 2: **Convergence of GDA, EG, OGDA, LA-GDA, and ACVI on three different 2d problems,** for a *fixed* number of *total* iterations, where markers denote the *iterates* of the respective method.

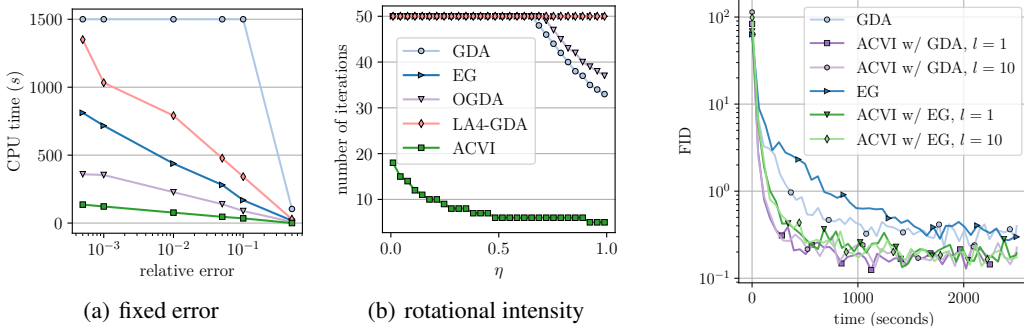


Figure 3: **Comparison on the (HBG) problem:** (a)–CPU time given fixed error, (b)–number of iterations needed to reach ϵ -distance to solution for varying intensity of the rotational component $(1 - \eta)$. For both plots, we set maximum iterations/time to run. See § 5 for discussion.




Figure 4: **FID (lower is better) on MNIST with added constraints,** over wall-clock time; averaged over 3 seeds. See § 5 and App. D for discussion and implementation, resp.

Methods. We compare ACVI with the projected variants of the common saddle point optimizers (fully described in App. A.4): (i) **GDA**, (ii) **EG** (Korpelevich, 1976), (iii) **OGDA** (Popov, 1980), and (iv) **LA \tilde{k} -GDA** (Chavdarova et al., 2021b; Zhang et al., 2019), where \tilde{k} is the hyperparameter of LA. For ACVI on MNIST, l denotes the number of steps to solve the subproblems; see Algorithm 4.

Results. From Fig. 1, we observe that projection-based algorithms may zigzag when hitting a constraint due to the rotational nature of F , behavior that ACVI avoids because it incorporates the constraints in its update rule; see Fig. 5 for the remaining baselines. Fig. 2 shows that even with problems that go beyond our theoretical assumptions, a single step of ACVI significantly reduces the distance to the solution. Moreover, from Fig. 2(b), we observe that ACVI escapes the limit cycles; see also Fig. 6. Fig. 3 shows results for (HBG), indicating that ACVI is time efficient, and that ACVI performs well relative to projection-based methods for varying rotational intensity $(1 - \eta)$. Fig. 4 summarizes the experiments on MNIST with linear inequality constraints; we observe that ACVI converges significantly faster than the corresponding baseline.

6 CONCLUSION

Motivated by the lack of a *first-order* method to solve constrained VI (cVI) problems with general constraints, we proposed a framework that combines (i) *interior-point* methods—needed to be able to handle general constraints—with (ii) the ADMM method—designed to deal with *separable* objectives. The combination yields ACVI—a first-order ADMM-based interior point method for cVIs. We proved convergence for two broad classes of problems and derived the corresponding convergence rates. Numerical experiments showed that while projection-based methods zigzag when hitting a constraint due to the rotational vector field, ACVI avoids this by incorporating the constraints in the update rule.

ACKNOWLEDGMENTS

TC thanks the support of the Swiss National Science Foundation (SNSF), grant P2ELP2_199740. The authors thank Matteo Pagliardini and Tianyi Lin for insightful discussions and feedback.

REFERENCES

- Kimon Antonakopoulos, E. Veronica Belmega, and Panayotis Mertikopoulos. Adaptive extra-gradient methods for min-max optimization and games. In *ICLR*, 2021.
- Yossi Arjevani, Shai Shalev-Shwartz, and Ohad Shamir. On lower and upper bounds for smooth and strongly convex optimization problems. In *JMLR*, 2016.
- Waïss Azizian, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *AISTATS*, pp. 2863–2873, 2020.
- David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In *ICML*, 2018.
- Stephen Bates, Michael I. Jordan, Michael Sklar, and Jake A. Soloff. Principal-agent hypothesis testing. *arXiv:2205.06812*, 2022.
- Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2nd edition, 2017. ISBN 1441994661.
- Amir Beck. *First-Order Methods in Optimization*. SIAM, 2017.
- Radu Ioan Bot, Ernő Robert Csetnek, and Dang-Khoa Nguyen. Fast OGDA in continuous and discrete time. *arXiv preprint arXiv:2203.10947*, 2022.
- Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge university press, 2004.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3, 2011. ISSN 1935-8237. doi: 10.1561/22000000016.
- Yang Cai, Argyris Oikonomou, and Weiqiang Zheng. Tight last-iterate convergence of the extragradient method for constrained monotone variational inequalities. *arXiv:2204.09228*, 2022.
- Tatjana Chavdarova, Gauthier Gidel, François Fleuret, and Simon Lacoste-Julien. Reducing noise in GAN training with variance reduced extragradient. In *NeurIPS*, 2019.
- Tatjana Chavdarova, Michael I. Jordan, and Manolis Zampetakis. Last-iterate convergence of saddle point optimizers via high-resolution differential equations. *NeurIPS Workshop on Optimization for Machine Learning*, 2021a.
- Tatjana Chavdarova, Matteo Pagliardini, Sebastian U Stich, François Fleuret, and Martin Jaggi. Taming GANs with Lookahead-Minmax. In *ICLR*, 2021b.
- Xiaojun Chen, Liquan Qi, and Defeng Sun. Global and superlinear convergence of the smoothing newton method and its application to general box constrained variational inequalities. *Mathematics of Computation*, 67(222):519–540, 1998.
- Rune Christiansen, Niklas Pfister, Martin Emil Jakobsen, Nicola Gnecco, and Jonas Peters. A causal framework for distribution generalization. *arXiv:2006.07433*, 2020.
- Richard W. Cottle and George B. Dantzig. Complementary pivot theory of mathematical programming. *Linear Algebra and its Applications*, 1(1):103–125, 1968. ISSN 0024-3795.
- Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *ITCS*, 2019.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training GANs with optimism. In *ICLR*, 2018.

- Constantinos Daskalakis, Dylan J. Foster, and Noah Golowich. Independent policy gradient methods for competitive reinforcement learning. In *NeurIPS*, 2020.
- Damek Davis and Wotao Yin. Convergence rate analysis of several splitting schemes. In *Splitting Methods in Communication, Imaging, Science, and Engineering*, pp. 115–163. Springer, 2016.
- Jelena Diakonikolas. Halpern iteration for near-optimal and parameter-free monotone inclusion and strong solutions to variational inequalities. *COLT*, 125, 2020.
- Jelena Diakonikolas, Constantinos Daskalakis, and Michael I. Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *AISTATS*, 2021.
- Jim Douglas and H. H. Jr. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82:421–439, 1956.
- Jonathan Eckstein and Dimitri P Bertsekas. On the Douglas—Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1): 293–318, 1992.
- Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional Variational Inequalities and Complementarity Problems*. Springer, 2003.
- Xiaona Fan and Qinglun Yan. An interior point algorithm for variational inequality problems. *International Journal of Contemporary Mathematical Sciences*, 5(52):2595–2604, 2010.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956. doi: <https://doi.org/10.1002/nav.3800030109>.
- Daniel Gabay. Applications of the method of multipliers to variational inequalities. In *Studies in Mathematics and its Applications*, volume 15, pp. 299–331. Elsevier, 1983.
- Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2:17–40, 1976. ISSN 0898-1221.
- Gauthier Gidel, Tony Jebara, and Simon Lacoste-Julien. Frank-wolfe algorithms for saddle point problems. In *Artificial Intelligence and Statistics*, pp. 362–371. PMLR, 2017a.
- Gauthier Gidel, Tony Jebara, and Simon Lacoste-Julien. Frank-Wolfe algorithms for saddle point problems. In *AISTATS*, 2017b.
- R. Glowinski and A. Marroco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de Dirichlet non linéaires. *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique*, 9(R2): 41–76, 1975.
- Roland Glowinski and Patrick Le Tallec. *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*. Society for Industrial and Applied Mathematics, 1989. doi: 10.1137/1.9781611970838.
- Jean-Louis Goffin, Patrice Marcotte, and Daoli Zhu. An analytic center cutting plane method for pseudomonotone variational inequalities. *Operations Research Letters*, 20(1):1–6, 1997. ISSN 0167-6377.
- Noah Golowich, Sarath Pattathil, and Constantinos Daskalakis. Tight last-iterate convergence rates for no-regret learning in multi-player games. In *NeurIPS*, 2020a.
- Noah Golowich, Sarath Pattathil, Constantinos Daskalakis, and Asuman Ozdaglar. Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems. In *COLT*, pp. 1758–1784, 2020b.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.

- Eduard Gorbunov, Nicolas Loizou, and Gauthier Gidel. Extragradient method: $\mathcal{O}(1/K)$ last-iterate convergence for monotone variational inequalities and connections with cocoercivity. In *AISTATS*, 2022.
- Benjamin Halpern. Fixed points of nonexpanding maps. *Bulletin of the American Mathematical Society*, 73:957–961, 1967.
- Bingsheng He and Xiaoming Yuan. On the $\mathcal{O}(1/n)$ convergence rate of the Douglas-Rachford alternating direction method. *SIAM Journal on Numerical Analysis*, 50(2):700–709, 2012.
- Bingsheng He and Xiaoming Yuan. On non-ergodic convergence rate of douglas–rachford alternating direction method of multipliers. *Numerische Mathematik*, 130(3):567–577, 2015.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017.
- Ya-Ping Hsieh, Panayotis Mertikopoulos, and Volkan Cevher. The limits of min-max optimization algorithms: convergence to spurious non-critical sets. In *ICML*, 2021.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- Martin Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *ICML*, volume 28, pp. 427–435. PMLR, 2013.
- Galina Michailovna Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 1976.
- Ivan Kupka. Contribution à la théorie des champs génériques. *Contributions to Differential Equations*, 2:457–484, 1963.
- Simon Lacoste-Julien and Martin Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. In *NIPS*, 2015.
- Yann Lecun and Corinna Cortes. The MNIST database of handwritten digits. 1998. URL <http://yann.lecun.com/exdb/mnist/>.
- Tianyi Lin, Shiqian Ma, Yinyu Ye, and Shuzhong Zhang. An ADMM-based interior-point method for large-scale linear programming. *arXiv:1805.12344*, 2018.
- Zhouchen Lin, Huan Li, and Cong Fang. *Alternating Direction Method of Multipliers for Machine Learning*. Springer, 2022.
- P. L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.
- Yu. Malitsky. Projected reflected gradient methods for monotone variational inequalities. *SIAM Journal on Optimization*, 25:502–520, 2015.
- Santiago Mazuelas, Andrea Zanoni, and Aritz Pérez. Minimax classification with 0-1 loss and performance guarantees. In *NeurIPS*, volume 33, pp. 302–312. Curran Associates, Inc., 2020.
- Nimrod Megiddo. *Progress in Mathematical Programming: Interior-Point and Related Methods*. Springer-Verlag, 1989. ISBN 0387968474.
- Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. *arXiv:1901.08511*, 2019.
- Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. Convergence rate of $\mathcal{O}(1/k)$ for optimistic gradient and extra-gradient methods in smooth convex-concave saddle point problems. In *SIAM Journal on Optimization*, 2020.
- Renato DC Monteiro and Jong-Shi Pang. Properties of an interior-point mapping for mixed complementarity problems. *Mathematics of Operations Research*, 21(3):629–654, 1996.

- Yurii Nesterov and Arkadi Nemirovski. Interior-point polynomial algorithms in convex programming. In *Siam Studies in Applied Mathematics*, 1994.
- Yurii Nesterov and Michael J. Todd. Primal-dual interior-point methods for self-scaled cones. *SIAM Journal on Optimization*, 8:324–364, 1998.
- Robert Nishihara, Laurent Lessard, Ben Recht, Andrew Packard, and Michael Jordan. A general analysis of the convergence of admm. In *ICML*, volume 37 of *Proceedings of Machine Learning Research*, pp. 343–352, 2015.
- Shayegan Omidshafiei, Jason Pazis, Chris Amato, Jonathan P. How, and John Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *ICML*, 2017.
- Jong-Shi Pang. A posteriori error bounds for the linearly-constrained variational inequality problem. *Mathematics of Operations Research*, 12:474–484, 1987.
- Leonid Denisovich Popov. A modification of the arrow–hurwicz method for search of saddle points. *Mathematical Notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980.
- Liquan Qi and Defeng Sun. Smoothing functions and smoothing newton method for complementarity and variational inequality problems. *Journal of Optimization Theory and Applications*, 113(1): 121–147, 2002.
- Liquan Qi, Defeng Sun, and Guanglu Zhou. A new look at smoothing newton methods for non-linear complementarity problems and box constrained variational inequalities. *Mathematical Programming*, 87(1):1–35, 2000.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016.
- Daniel Ralph and Stephen J Wright. Superlinear convergence of an interior-point method despite dependent constraints. *Mathematics of Operations Research*, 25(2):179–194, 2000.
- James Renegar. *A Mathematical View of Interior-Point Methods in Convex Optimization*. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, 2001. ISBN 9780898715026.
- Ralph Tyrrell Rockafellar. Monotone operators associated with saddle-functions and minimax problems. *Nonlinear functional analysis*, 18(part 1):397–407, 1970.
- Dominik Rothenhäusler, Nicolai Meinshausen, Peter Bühlmann, and Jonas Peters. Anchor regression: heterogeneous data meets causality. *ArXiv:1801.06229*, 2018.
- Ernest K. Ryu and Wotao Yin. *Large-Scale Convex Optimization via Monotone Operators*. Springer Publishing Company, Incorporated, 2022.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *NIPS*, 2016.
- Stephen Smale. Stable manifolds for differential equations and diffeomorphisms. *Annali della Scuola Normale Superiore di Pisa-Classe di Scienze*, 17(1-2):97–116, 1963.
- Guido Stampacchia. Formes bilinéaires coercitives sur les ensembles convexes. *Académie des Sciences de Paris*, 258:4413–4416, 1964.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv:1312.6199*, 2014.
- Paul Tseng. A path-following algorithm for linear programming using quadratic and logarithmic penalty functions. *SIAM Journal on Control and Optimization*, 31:1578–1598, 1993.
- Paul Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60:237–252, 1995. ISSN 0377-0427.

- A.M Rubinov V.F Demyanov. *Approximate Methods in Optimization Problems*, volume 9. American Elsevier, New York, 1970.
- Jean-Philippe Vial. Strong and weak convexity of sets and functions. *Mathematics of Operations Research*, 8:231–259, 1983.
- John Von Neumann. *A Model of General Economic Equilibrium*. Palgrave Macmillan UK, 1971. ISBN 978-1-349-15430-2.
- John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947.
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. In *ICLR*, 2021.
- Stephen J. Wright. *Primal-Dual Interior-Points Methods*. SIAM, 1997.
- Stephen J. Wright. *Lectures on Modern Convex Optimization. Analysis, Algorithms, and Engineering Applications*. SIAM, 2001.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Michael Zhang, James Lucas, Jimmy Ba, and Geoffrey E Hinton. Lookahead optimizer: k steps forward, 1 step back. In *NeurIPS*, 2019.

A BACKGROUND: ADDITIONAL DETAILS

This section lists additional background such as omitted definitions and a description of the used baseline methods.

A.1 ADDITIONAL VI DEFINITIONS & EQUIVALENT FORMULATIONS

Seeing an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$ as the graph $GrF = \{(\mathbf{x}, \mathbf{y}) | \mathbf{x} \in \mathcal{X}, \mathbf{y} = F(\mathbf{x})\}$, its inverse F^{-1} is defined as $GrF^{-1} = \{(\mathbf{y}, \mathbf{x}) | (\mathbf{x}, \mathbf{y}) \in GrF\}$. See, for example, (Ryu & Yin, 2022) for further discussion. We denote the projection to the set \mathcal{X} with $\Pi_{\mathcal{X}}$.

Definition 3 ($\frac{1}{\mu}$ -cocoercive operator). *An operator $F : \mathcal{X} \supseteq \mathcal{S} \rightarrow \mathbb{R}^n$ is $\frac{1}{\mu}$ -cocoercive (or $\frac{1}{\mu}$ -inverse strongly monotone) on \mathcal{S} if its inverse (graph) F^{-1} is μ -strongly monotone on \mathcal{S} , that is,*

$$\exists \mu > 0, \quad \text{s.t.} \quad \langle \mathbf{x} - \mathbf{x}', F(\mathbf{x}) - F(\mathbf{x}') \rangle \geq \mu \|F(\mathbf{x}) - F(\mathbf{x}')\|^2, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{S}.$$

It is star $\frac{1}{\mu}$ -cocoercive if the above holds when setting $\mathbf{x}' \equiv \mathbf{x}^$ where \mathbf{x}^* denotes a solution, that is:*

$$\exists \mu > 0, \quad \text{s.t.} \quad \langle \mathbf{x} - \mathbf{x}^*, F(\mathbf{x}) - F(\mathbf{x}^*) \rangle \geq \mu \|F(\mathbf{x}) - F(\mathbf{x}^*)\|^2, \forall \mathbf{x} \in \mathcal{S}, \mathbf{x}^* \in \mathcal{S}_{\mathcal{X}, F}^*.$$

Note from Def. 3 that cocoercivity is a strict subclass of monotone and L-Lipschitz operators, thus it is a stronger assumption. See Chapter 4.2 of (Bauschke & Combettes, 2017) for further relations of cocoercivity with other properties of operators.

In the following, we will make use of the natural and normal mappings of an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$, where $\mathcal{X} \subset \mathbb{R}^n$. Following the notation of (Facchinei & Pang, 2003), the natural map $F_{\mathcal{X}}^{\text{NAT}} : \mathcal{X} \rightarrow \mathbb{R}^n$ is:

$$F_{\mathcal{X}}^{\text{NAT}} \triangleq \mathbf{x} - \Pi_{\mathcal{X}}(\mathbf{x} - F(\mathbf{x})), \quad \forall \mathbf{x} \in \mathcal{X}, \quad (\text{F-NAT})$$

whereas the normal map $F_{\mathcal{X}}^{\text{NOR}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is:

$$F_{\mathcal{X}}^{\text{NOR}} \triangleq F(\Pi_{\mathcal{X}}(\mathbf{x})) + \mathbf{x} - \Pi_{\mathcal{X}}(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{R}^n. \quad (\text{F-NOR})$$

Moreover, we have the following solution characterizations:

- (i) $\mathbf{x}^* \in \mathcal{S}_{\mathcal{X}, F}^*$ iff $F_{\mathcal{X}}^{\text{NAT}}(\mathbf{x}^*) = \mathbf{0}$, and
- (ii) $\mathbf{x}^* \in \mathcal{S}_{\mathcal{X}, F}^*$ iff $\exists \mathbf{x}' \in \mathbb{R}^n$ s.t. $\mathbf{x}^* = \Pi_{\mathcal{X}}(\mathbf{x}')$ and $F_{\mathcal{X}}^{\text{NOR}}(\mathbf{x}') = \mathbf{0}$.

Remark 4 (“rotational” component of the vector field). *The rotational trajectories in parameter space are induced by the fact that the eigenvalues of the Jacobian of the vector field F (the second-order derivative matrix) belong to the complex set \mathbb{C} ; that is $\lambda_i \in \mathbb{C}$. In contrast, when $F \equiv \nabla f$ the second-order derivative matrix known as the Hessian is always symmetric, and thus the eigenvalues are real.*

A.2 EXISTENCE OF SOLUTION

We provide brief informal summary of some sufficient conditions for solution existence, that $\mathcal{S}_{\mathcal{X}, F}^* \neq \emptyset$. See Chapter 2 of (Facchinei & Pang, 2003) for a full treatment of the topic.

The common underlying tool to establish that a solution to the VI(\mathcal{X} , F) problem exists is using *topological degree*. The topological degree tool is *designed* so as to satisfy the so-called *homotopy invariance* axiom, which in turn allows for reducing a solution existence question of a complicated map to a simpler one (which is homotopy-invariant to the original one) for which we can more easily show that it has a solution (e.g., the identity map on a closed domain). It can be used (as one way) to prove the celebrated *Brouwer fixed-point theorem*, which states that any continuous map $\Phi : \mathcal{S} \rightarrow \mathcal{S}$, where \mathcal{S} is a nonempty convex compact set, has a fixed point in \mathcal{S} .

We have the following sufficient condition, see (Cor. 2.2.5, Facchinei & Pang, 2003)

Theorem 4 (sufficient condition for existence of the solution, Cor. 2.2.5, (Facchinei & Pang, 2003)). *If $\mathcal{X} \subseteq \mathbb{R}^n$ is compact and convex, and $F : \mathcal{X} \rightarrow \mathbb{R}^n$ is continuous, then the solution set is nonempty and compact.*

It can also be shown that when \mathcal{X} is closed convex (and F continuous), if one can find $\mathbf{x}' \in \mathcal{X}$ s.t. $\langle F(\mathbf{x}), \mathbf{x} - \mathbf{x}' \rangle \geq 0, \forall \mathbf{x} \in \mathcal{X}$, then $\mathcal{S}_{\mathcal{X}, F}^* \neq \emptyset$. The same conclusion follows if one can show that $\exists \mathbf{x}' \in \mathcal{X}$ and the set $\{\mathbf{x} \in \mathcal{X} | \langle F(\mathbf{x}), \mathbf{x} - \mathbf{x}' \rangle < 0\}$ is bounded (possibly empty, see Prop. 2.2.3 in (Facchinei & Pang, 2003)).

Sufficient conditions can also be established via the natural and the normal map due to the above solution characterizations. In this case, we require that F is continuous on an open set \mathcal{S} and we are interested if $\mathcal{S}_{\mathcal{X}, F}^* \neq \emptyset$, where \mathcal{X} is assumed closed and convex and subset of \mathcal{S} , $\mathcal{X} \subseteq \mathcal{S}$. If one establishes that a solution exists for $F_{\mathcal{X}}^{\text{NAT}}$ on a bounded open set \mathcal{U} , and if $\text{cl } \mathcal{U} \subseteq \mathcal{S}$, then it follows that $\mathcal{S}_{\mathcal{X}, F}^* \neq \emptyset$. A similar implication holds when we have such a guarantee for $F_{\mathcal{X}}^{\text{NOR}}$. See Theorem 2.2.1 of (Facchinei & Pang, 2003).

In summary, the solution existence guarantee follows from the boundness of some set which includes \mathcal{X} , the boundness of the set of potential solutions (if we can construct such set), or the compactness of \mathcal{X} itself.

A.3 EXISTENCE OF CENTRAL PATH

In this section, we discuss the results that establish guarantees of the existence of the central path. Let $L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \triangleq F(\mathbf{x}) + \nabla \varphi^\top(\mathbf{x})\boldsymbol{\lambda} + \mathbf{C}^\top \boldsymbol{\nu}$, $h(\mathbf{x}) = \mathbf{C}^\top \mathbf{x} - \mathbf{d}$. For $(\boldsymbol{\lambda}, \mathbf{w}, \mathbf{x}, \boldsymbol{\nu}) \in \mathbb{R}^{2m+n+p}$, let

$$G(\boldsymbol{\lambda}, \mathbf{w}, \mathbf{x}, \boldsymbol{\nu}) \triangleq \begin{pmatrix} \mathbf{w} \circ \boldsymbol{\lambda} \\ \mathbf{w} + \varphi(\mathbf{x}) \\ L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \\ h(\mathbf{x}) \end{pmatrix} \in \mathbb{R}^{2m+n+p},$$

and

$$H(\boldsymbol{\lambda}, \mathbf{w}, \mathbf{x}, \boldsymbol{\nu}) \triangleq \begin{pmatrix} \mathbf{w} + \varphi(\mathbf{x}) \\ L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \\ h(\mathbf{x}) \end{pmatrix} \in \mathbb{R}^{m+n+p}.$$

Let $H_{++} \triangleq H(\mathbb{R}_{++}^{2m} \times \mathbb{R}^n \times \mathbb{R}^p)$. By (Corollary 11.4.24, Facchinei & Pang, 2003) we have the following proposition.

Proposition 1 (sufficient condition for the existence of the central path.). *If F is monotone, either F is strictly monotone or one of φ_i is strictly convex, and \mathcal{C} is bounded. The following four statements hold for the functions G and H :*

1. G maps $\mathbb{R}_{++}^{2m} \times \mathbb{R}^{n+p}$ homeomorphically onto $\mathbb{R}_{++}^m \times H_{++}$;
2. $\mathbb{R}_{++}^m \times H_{++} \subseteq G(\mathbb{R}_{++}^{2m} \times \mathbb{R}^{n+p})$;
3. for every vector $\mathbf{a} \in \mathbb{R}_{++}^m$, the system

$$H(\boldsymbol{\lambda}, \mathbf{w}, \mathbf{x}, \boldsymbol{\nu}) = \mathbf{0}, \quad \mathbf{w} \circ \boldsymbol{\lambda} = \mathbf{a}$$

has a solution $(\boldsymbol{\lambda}, \mathbf{w}, \mathbf{x}, \boldsymbol{\nu}) \in \mathbb{R}_{++}^{2m} \times \mathbb{R}^{n+p}$; and

4. the set H_{++} is convex.

A.4 SADDLE-POINT OPTIMIZATION METHODS

In this section, we describe in detail the saddle point methods that we compare within the main part (in § 5). We denote the projection to the set \mathcal{X} with $\Pi_{\mathcal{X}}$, and when the method is applied in the unconstrained setting $\Pi_{\mathcal{X}} \equiv \mathbf{I}$.

For an example of the associated vector field and its Jacobian, consider the following constrained zero-sum game:

$$\min_{\mathbf{x}_1 \in \mathcal{X}_1} \max_{\mathbf{x}_2 \in \mathcal{X}_2} f(\mathbf{x}_1, \mathbf{x}_2), \quad (\text{ZS-G})$$

where $f : \mathcal{X}_1 \times \mathcal{X}_2 \rightarrow \mathbb{R}$ is smooth and convex in \mathbf{x}_1 and concave in \mathbf{x}_2 . As in the main paper, we write $\mathbf{x} \triangleq (\mathbf{x}_1, \mathbf{x}_2) \in \mathbb{R}^n$. The vector field $F : \mathcal{X} \rightarrow \mathbb{R}^n$ and its Jacobian J are defined as:

$$F(\mathbf{x}) = \begin{bmatrix} \nabla_{\mathbf{x}_1} f(\mathbf{x}) \\ -\nabla_{\mathbf{x}_2} f(\mathbf{x}) \end{bmatrix}, \quad J(\mathbf{x}) = \begin{bmatrix} \nabla_{\mathbf{x}_1}^2 f(\mathbf{x}) & \nabla_{\mathbf{x}_2} \nabla_{\mathbf{x}_1} f(\mathbf{x}) \\ -\nabla_{\mathbf{x}_1} \nabla_{\mathbf{x}_2} f(\mathbf{x}) & -\nabla_{\mathbf{x}_2}^2 f(\mathbf{x}) \end{bmatrix}.$$

In the remainder of this section, we will only refer to the joint variable \mathbf{x} , and (with abuse of notation) the subscript will denote the step. Let $\gamma \in [0, 1]$ denote the step size.

(Projected) Gradient Descent Ascent (GDA). The extension of gradient descent for the cVI problem is *gradient descent ascent* (GDA). The GDA update at step k is then:

$$\mathbf{x}_{k+1} = \Pi_{\mathcal{X}}(\mathbf{x}_k - \gamma F(\mathbf{x}_k)). \quad (\text{GDA})$$

(Projected) Extragradient (EG). EG (Korpelevich, 1976) uses a ‘‘prediction’’ step to obtain an extrapolated point $\mathbf{x}_{k+\frac{1}{2}}$ using GDA: $\mathbf{x}_{k+\frac{1}{2}} = \Pi_{\mathcal{X}}(\mathbf{x}_k - \gamma F(\mathbf{x}_k))$, and the gradients at the *extrapolated* point are then applied to the *current* iterate \mathbf{x}_k :

$$\mathbf{x}_{k+1} = \Pi_{\mathcal{X}}\left(\mathbf{x}_k - \gamma F\left(\Pi_{\mathcal{X}}(\mathbf{x}_k - \gamma F(\mathbf{x}_k))\right)\right). \quad (\text{EG})$$

In the original EG paper, (Korpelevich, 1976) proved that the EG method (with a fixed step size) converges for monotone VIs, as follows.

Theorem 5 (Korpelevich (1976)). *Given a map $F : \mathcal{X} \mapsto \mathbb{R}^n$, if the following is satisfied:*

1. *the set \mathcal{X} is closed and convex,*
2. *F is single-valued, definite, and monotone on \mathcal{X} —as per Def. 1,*
3. *F is L -Lipschitz—as per Asm. 1.*

then there exists a solution $\mathbf{x}^ \in \mathcal{X}$, such that the iterates \mathbf{x}_k produced by the EG update rule with a fixed step size $\gamma \in (0, \frac{1}{L})$ converge to it, that is $\mathbf{x}_k \rightarrow \mathbf{x}^*$, as $k \rightarrow \infty$.*

Facchinei & Pang (2003) also show that for any *convex-concave* function f and any closed convex sets $\mathbf{x}_1 \in \mathcal{X}_1$ and $\mathbf{x}_2 \in \mathcal{X}_2$, the EG method converges (Facchinei & Pang, 2003, Theorem 12.1.11).

(Projected) Optimistic Gradient Descent Ascent (OGDA). The update rule of Optimistic Gradient Descent Ascent OGDA ((OGDA) Popov, 1980) is:

$$\mathbf{x}_{n+1} = \Pi_{\mathcal{X}}(\mathbf{x}_n - 2\gamma F(\mathbf{x}_n) + \gamma F(\mathbf{x}_{n-1})). \quad (\text{OGDA})$$

(Projected) Lookahead–Minmax (LA). The LA algorithm for min-max optimization (Chavdarova et al., 2021b), originally proposed for minimization by Zhang et al. (2019), is a general wrapper of a ‘‘base’’ optimizer where, at every step t : (i) a copy of the current iterate $\tilde{\mathbf{x}}_n$ is made: $\tilde{\mathbf{x}}_n \leftarrow \mathbf{x}_n$, (ii) $\tilde{\mathbf{x}}_n$ is updated $k \geq 1$ times, yielding $\tilde{\omega}_{n+k}$, and finally (iii) the actual update \mathbf{x}_{n+1} is obtained as a point that lies on a line between the current \mathbf{x}_n iterate and the predicted one $\tilde{\mathbf{x}}_{n+k}$:

$$\mathbf{x}_{n+1} \leftarrow \mathbf{x}_n + \alpha(\tilde{\mathbf{x}}_{n+k} - \mathbf{x}_n), \quad \alpha \in [0, 1]. \quad (\text{LA})$$

In this work, we use solely GDA as a base optimizer for LA and thus denote it with *LAk-GDA*.

The projection-free Frank-Wolfe (FW). FW (Frank & Wolfe, 1956) is an IP-type method for solving constrained smooth zero-sum games (ZS-G). It avoids the projection operator by ensuring we never leave the constraint set. To do so, it finds the intersection points of the inequality constraints—hence, it requires that the inequality constraints satisfy certain structures (such as linear) in order for this operation to be computationally cheap. We state FW for zero-sum games in Algorithm 2 as proposed by Gidel et al. (2017a, Alg. 2) for completeness. It requires access to a linear minimization oracle (LMO) over the constraint set—to minimize the linear function in line 5 in Algorithm 2. Currently, FW-style algorithms only have convergence guarantees when the constraint set is a polytope ((Gidel et al., 2017a)) with additional assumptions (which we listed in § 2.

Algorithm 2 Frank-Wolfe algorithm for zero-sum games.

```
1: Input:  $C, \nu > 0$ 
2: Initialize:  $\mathbf{z}^{(0)} = (\mathbf{x}_1^{(0)}, \mathbf{x}_2^{(0)}) \in \mathcal{X}_1 \times \mathcal{X}_2$ 
3: for  $t = 0, \dots, T$  do
4:   Compute  $\mathbf{r}^{(t)} \triangleq \begin{bmatrix} \nabla_{\mathbf{x}_1} f(\mathbf{x}_1^{(t)}, \mathbf{x}_2^{(t)}) \\ -\nabla_{\mathbf{x}_2} f(\mathbf{x}_1^{(t)}, \mathbf{x}_2^{(t)}) \end{bmatrix}$ 
5:    $\mathbf{s}^{(t)} \triangleq \arg \min_{\mathbf{z} \in \mathcal{X}_1 \times \mathcal{X}_2} \langle \mathbf{z}, \mathbf{r}^{(t)} \rangle$ 
6:   Compute  $g_t \triangleq \langle \mathbf{z}^{(t)} - \mathbf{s}^{(t)}, \mathbf{r}^{(t)} \rangle$ 
7:   if  $g_t \leq \varepsilon$  then
8:     return  $\mathbf{z}^{(t)}$ 
9:   end if
10:  Let  $\gamma = \min(1, \frac{\nu}{2C} g_t)$  or  $\gamma = \frac{2}{2+t}$ 
11:  Update  $\mathbf{z}^{(t+1)} = (1 - \gamma)\mathbf{z}^{(t)} + \gamma\mathbf{s}^{(t)}$ 
12: end for
```

B OMITTED PROOFS AND DISCUSSIONS CONCERNING ALGORITHM 1

This section provides the proofs of the theoretical results in § 4.2.

B.1 PROOF OF THEOREM 1: UNIQUENESS OF THE SOLUTION OF EQ. W-EQ

Recall first that Eq. W-EQ is as follows:

$$\mathbf{x} + \frac{1}{\beta} \mathbf{P}_c F(\mathbf{x}) - \mathbf{P}_c \mathbf{y}_k + \frac{1}{\beta} \mathbf{P}_c \boldsymbol{\lambda}_k - \mathbf{d}_c = \mathbf{0},$$

since $\mathbf{w} = \mathbf{x}$.

Proof of Theorem 1: uniqueness of the solution of (W-EQ). Let $G(\mathbf{x})$ denote the LHS of (W-EQ), that is:

$$G(\mathbf{x}) \triangleq \mathbf{x} + \frac{1}{\beta} \mathbf{P}_c F(\mathbf{x}) - \mathbf{P}_c \mathbf{y}_k + \frac{1}{\beta} \mathbf{P}_c \boldsymbol{\lambda}_k - \mathbf{d}_c \quad (3)$$

We claim that $G(\mathbf{x})$ is strongly monotone on $\mathcal{C}_=$. In fact, $\forall \mathbf{x}, \mathbf{y} \in \mathcal{C}_=$, $\mathbf{P}_c(\mathbf{x} - \mathbf{y}) = \mathbf{x} - \mathbf{y}$. Note that \mathbf{P}_c is symmetric, thus we have:

$$\begin{aligned} \langle G(\mathbf{x}) - G(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle &= \|\mathbf{x} - \mathbf{y}\|^2 + \frac{1}{\beta} \langle \mathbf{P}_c F(\mathbf{x}) - \mathbf{P}_c F(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \\ &= \|\mathbf{x} - \mathbf{y}\|^2 + \frac{1}{\beta} \langle \mathbf{x} - \mathbf{y}, F(\mathbf{x}) - F(\mathbf{y}) \rangle \\ &\geq \|\mathbf{x} - \mathbf{y}\|^2. \end{aligned}$$

Therefore, according to Theorem 2.3.3 (b) in (Facchinei & Pang, 2003), $\mathcal{S}_{\mathcal{C}_=, G}^*$ has a unique solution $\tilde{\mathbf{x}} \in \mathcal{C}_=$. Thus, we have:

$$G(\tilde{\mathbf{x}})^\top (\mathbf{x} - \tilde{\mathbf{x}}) = 0, \quad \forall \mathbf{x} \in \mathcal{C}_=.$$

From the above, we deduce that $G(\mathbf{x}) \in \text{Span}\{\mathbf{c}_1, \dots, \mathbf{c}_p\}$, where \mathbf{c}_i is the row vectors of \mathbf{C} , $i \in [p]$.

Suppose that $G(\tilde{\mathbf{x}}) = \sum_{i=1}^p \alpha_i \mathbf{c}_i$. Notice that $\mathbf{C}G(\mathbf{x}) = \mathbf{0}$, $\forall \mathbf{x} \in \mathcal{C}_=$. Thus, we have that:

$$\mathbf{c}_j^\top G(\tilde{\mathbf{x}}) = \mathbf{c}_j^\top \sum_{i=1}^p \alpha_i \mathbf{c}_i = 0, \quad \forall j \in [p].$$

Hence,

$$\langle G(\tilde{\mathbf{x}}), G(\tilde{\mathbf{x}}) \rangle = \left\langle \sum_{i=1}^p \alpha_i \mathbf{c}_i, \sum_{i=1}^p \alpha_i \mathbf{c}_i \right\rangle = 0,$$

which indicates that $G(\tilde{\mathbf{x}}) = \mathbf{0}$. Hence, $\tilde{\mathbf{x}}$ is a solution of (W-EQ) in $\mathcal{C}_=$.

On the other hand, $\forall \mathbf{x} \in \mathbb{R}^n$, if \mathbf{x} is a solution of (W-EQ), i.e. $G(\mathbf{x}) = \mathbf{0}$, then $\mathbf{x} \in \mathcal{C}_=$. By the uniqueness of $\tilde{\mathbf{x}}$ in $\mathcal{C}_=$ we have that $\mathbf{x} = \tilde{\mathbf{x}}$, which means $\tilde{\mathbf{x}}$ is unique in \mathbb{R}^n . \square

B.2 PROOF OF LEMMA 1: UPPER BOUND ON THE GAP FUNCTION

Proof of Lemma 1: Upper bound on the gap function. Let \mathbf{x}_k denote an iterate produced by Algorithm 1, and let $\mathbf{x} \in \mathcal{C}$. Note that we always have $\mathbf{x}_k \in \mathcal{C}_=$. We have that:

$$\begin{aligned} \langle F(\mathbf{x}_k), \mathbf{x}_k - \mathbf{x} \rangle &= \langle F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle + \langle F(\mathbf{x}_k), \mathbf{x}_k - \mathbf{x} \rangle - \langle F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \\ &= \langle F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle + \langle F(\mathbf{x}_k), \mathbf{x}_k - \mathbf{x}^* + \mathbf{x}^* - \mathbf{x} \rangle - \langle F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle. \end{aligned}$$

From the proof of Theorem 6 we know that \mathbf{x}_k is bounded, which gives: $\langle F(\mathbf{x}_k), \mathbf{x}_k - \mathbf{x}^* \rangle \leq M \|\mathbf{x}_k - \mathbf{x}^*\|$, with $M > 0$, as well as that $\langle F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \leq 0$. Thus, for the above we get:

$$\begin{aligned} \langle F(\mathbf{x}_k), \mathbf{x}_k - \mathbf{x} \rangle &\leq \langle F(\mathbf{x}_k) - F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle + M \|\mathbf{x}_k - \mathbf{x}^*\| \\ &\leq D \|F(\mathbf{x}_k) - F(\mathbf{x}^*)\| + M \|\mathbf{x}_k - \mathbf{x}^*\| \\ &\leq (DL + M) \|\mathbf{x}_k - \mathbf{x}^*\|, \end{aligned}$$

where for the second row we used $\mathbf{x}^* - \mathbf{x} \leq D$ where $D \triangleq \max_{\mathbf{x}' \in \mathcal{C}} \|\mathbf{x}^* - \mathbf{x}'\|$ is the largest distance between any point in \mathcal{C} and \mathbf{x}^* . For the last row we used that F is L -Lipschitz—Assumption 1—which concludes the proof. \square

The proof is analogous for the $\mathbf{y}_k \in \mathcal{C}_{\leq}$ iterates produced by Algorithm 1.

B.3 PROOFS OF THE CONVERGENCE RATE: THEOREMS 2 AND 3

Let

$$\begin{aligned} f^{(t)}(\mathbf{x}) &\triangleq F(\mathbf{x}^{\mu_t})^\top \mathbf{x} + \mathbb{1}(\mathbf{C}\mathbf{x} = \mathbf{d}), \\ f_k^{(t)}(\mathbf{x}) &\triangleq F(\mathbf{x}_{k+1}^{(t)})^\top \mathbf{x} + \mathbb{1}(\mathbf{C}\mathbf{x} = \mathbf{d}), \\ g^{(t)}(\mathbf{y}) &\triangleq -\mu_t \sum_{i=1}^m \log(-\varphi_i(\mathbf{y})), \end{aligned}$$

where \mathbf{x}^{μ_t} is a solution of (KKT-2) when $\mu = \mu_t$. Note that the existence of \mathbf{x}^{μ_t} is guaranteed by the existence of the central path—see App. A, and that $f^{(t)}$, $f_k^{(t)}$ and $g^{(t)}$ are all convex.

In the following proofs, unless causing confusion, we drop the subscript t to simplify notations.

Let $\mathbf{y}^\mu = \mathbf{x}^\mu$, from (KKT-2) we can see that $(\mathbf{x}^\mu, \mathbf{y}^\mu)$ is an optimal solution of

$$\begin{cases} \min f(\mathbf{x}) + g(\mathbf{y}) \\ \text{s.t. } \mathbf{x} = \mathbf{y} \end{cases}. \quad (4)$$

There exists $\boldsymbol{\lambda}^\mu \in \mathbb{R}^n$ such that $(\mathbf{x}^\mu, \mathbf{y}^\mu, \boldsymbol{\lambda}^\mu)$ is a KKT point of (4). We give the following proposition which we will repeatedly use in the proofs:

Proposition 2. *If F is monotone, then $\forall k \in \mathbb{N}$,*

$$f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}^\mu) \geq f(\mathbf{x}_{k+1}) - f(\mathbf{x}^\mu).$$

Furthermore, if F is ξ -monotone, as per Def. 1

$$f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}^\mu) \geq f(\mathbf{x}_{k+1}) - f(\mathbf{x}^\mu) + c\|\mathbf{x}_{k+1} - \mathbf{x}^\mu\|_2^\xi.$$

Proof of Proposition 2. It suffices to note that:

$$f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}^\mu) - (f(\mathbf{x}_{k+1}) - f(\mathbf{x}^\mu)) = (F(\mathbf{x}_{k+1}) - F(\mathbf{x}^\mu))^\top (\mathbf{x}_{k+1} - \mathbf{x}^\mu).$$

\square

Some of our proofs that follow are inspired by some convergence proofs in ADMM (Gabay, 1983; Eckstein & Bertsekas, 1992; Davis & Yin, 2016; He & Yuan, 2012; 2015; Lin et al., 2022). However, although Algorithm 1 adopts the high level idea of ADMM, we can not directly refer to the convergence proofs of ADMM, but need to substantially modify these.

We will use the following lemma.

Lemma 2. $f(\mathbf{x}) + g(\mathbf{y}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x} - \mathbf{y} \rangle \geq 0, \forall \mathbf{x}, \mathbf{y}$.

Proof. The Lagrange function of (4) is

$$L(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) = f(\mathbf{x}) + g(\mathbf{y}) + \boldsymbol{\lambda}^\top (\mathbf{x} - \mathbf{y}).$$

And by the property of KKT point, we have

$$L(\mathbf{x}^\mu, \mathbf{y}^\mu, \boldsymbol{\lambda}) \leq L(\mathbf{x}^\mu, \mathbf{y}^\mu, \boldsymbol{\lambda}^\mu) \leq L(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}^\mu), \quad \forall (\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}),$$

from which the conclusion follows. \square

The following lemma is straightforward to verify:

Lemma 3. *If*

$$\begin{aligned} f(\mathbf{x}) + g(\mathbf{y}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x} - \mathbf{y} \rangle &\leq \alpha_1, \\ \|\mathbf{x} - \mathbf{y}\| &\leq \alpha_2 \end{aligned}$$

then we have

$$-\|\boldsymbol{\lambda}^\mu\|_{\alpha_2} \leq f(\mathbf{x}) + g(\mathbf{y}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) \leq \|\boldsymbol{\lambda}^\mu\|_{\alpha_2} + \alpha_1.$$

The following lemma lists some simple but useful facts that we will use in the following proofs.

Lemma 4. *For (4) and Algorithm 1, we have*

$$\mathbf{0} \in \partial f_k(\mathbf{x}_{k+1}) + \boldsymbol{\lambda}_k + \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1}) \quad (5)$$

$$\mathbf{0} \in \partial g(\mathbf{y}_{k+1}) - \boldsymbol{\lambda}_k - \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1}), \quad (6)$$

$$\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k = \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1}), \quad (7)$$

$$\mathbf{0} \in \partial f(\mathbf{x}^\mu) + \boldsymbol{\lambda}^\mu, \quad (8)$$

$$\mathbf{0} \in \partial g(\mathbf{y}^\mu) - \boldsymbol{\lambda}^\mu, \quad (9)$$

$$\mathbf{x}^\mu = \mathbf{y}^\mu. \quad (10)$$

We define:

$$\hat{\nabla} f_k(\mathbf{x}_{k+1}) \triangleq -\boldsymbol{\lambda}_k - \beta(\mathbf{x}_{k+1} - \mathbf{y}_k), \quad (11)$$

$$\hat{\nabla} g(\mathbf{y}_{k+1}) \triangleq \boldsymbol{\lambda}_k + \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1}). \quad (12)$$

Then from (5) and (6) we can see that

$$\hat{\nabla} f_k(\mathbf{x}_{k+1}) \in \partial f_k(\mathbf{x}_{k+1}) \text{ and } \hat{\nabla} g(\mathbf{y}_{k+1}) \in \partial g(\mathbf{y}_{k+1}). \quad (13)$$

Lemma 5. *For Algorithm 1, we have*

$$\langle \hat{\nabla} g(\mathbf{y}_{k+1}), \mathbf{y}_{k+1} - \mathbf{y} \rangle = -\langle \boldsymbol{\lambda}_{k+1}, \mathbf{y} - \mathbf{y}_{k+1} \rangle, \quad (14)$$

and

$$\begin{aligned} &\langle \hat{\nabla} f_k(\mathbf{x}_{k+1}), \mathbf{x}_{k+1} - \mathbf{x} \rangle + \langle \hat{\nabla} g(\mathbf{y}_{k+1}), \mathbf{y}_{k+1} - \mathbf{y} \rangle \\ &= -\langle \boldsymbol{\lambda}_{k+1}, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} - \mathbf{x} + \mathbf{y} \rangle + \beta \langle -\mathbf{y}_{k+1} + \mathbf{y}_k, \mathbf{x}_{k+1} - \mathbf{x} \rangle. \end{aligned} \quad (15)$$

Proof of Lemma 5. From (7), (11) and (12) we have:

$$\begin{aligned} &\langle \hat{\nabla} f_k(\mathbf{x}_{k+1}), \mathbf{x}_{k+1} - \mathbf{x} \rangle \\ &= -\langle \boldsymbol{\lambda}_k + \beta(\mathbf{x}_{k+1} - \mathbf{y}_k), \mathbf{x}_{k+1} - \mathbf{x} \rangle \\ &= -\langle \boldsymbol{\lambda}_{k+1}, \mathbf{x}_{k+1} - \mathbf{x} \rangle + \beta \langle -\mathbf{y}_{k+1} + \mathbf{y}_k, \mathbf{x}_{k+1} - \mathbf{x} \rangle, \end{aligned}$$

and

$$\langle \hat{\nabla} g(\mathbf{y}_{k+1}), \mathbf{y}_{k+1} - \mathbf{y} \rangle = -\langle \boldsymbol{\lambda}_{k+1}, \mathbf{y} - \mathbf{y}_{k+1} \rangle.$$

Adding these together yields (15). \square

Lemma 6. *For Algorithm 1, we have*

$$\begin{aligned} &\langle \hat{\nabla} f_k(\mathbf{x}_{k+1}), \mathbf{x}_{k+1} - \mathbf{x}^\mu \rangle + \langle \hat{\nabla} g(\mathbf{y}_{k+1}), \mathbf{y}_{k+1} - \mathbf{y}^\mu \rangle + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \\ &\leq \frac{1}{2\beta} \|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 \\ &\quad + \frac{\beta}{2} \|\mathbf{y}^\mu - \mathbf{y}_k\|^2 - \frac{\beta}{2} \|\mathbf{y}^\mu - \mathbf{y}_{k+1}\|^2 \\ &\quad - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 - \frac{\beta}{2} \|\mathbf{y}_k - \mathbf{y}_{k+1}\|^2 \end{aligned}$$

Proof of Lemma 6. Letting $(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) = (\mathbf{x}^\mu, \mathbf{y}^\mu, \boldsymbol{\lambda}^\mu)$ in (15), adding $\langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle$ to both sides, and using (7) and (10), we have:

$$\begin{aligned} & \langle \hat{\nabla} f_k(\mathbf{x}_{k+1}), \mathbf{x}_{k+1} - \mathbf{x}^\mu \rangle + \langle \hat{\nabla} g(\mathbf{y}_{k+1}), \mathbf{y}_{k+1} - \mathbf{y}^\mu \rangle + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \\ &= -\langle \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle + \beta \langle -\mathbf{y}_{k+1} + \mathbf{y}_k, \mathbf{x}_{k+1} - \mathbf{x}^\mu \rangle \\ &= -\frac{1}{\beta} \langle \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu, \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k \rangle + \langle -\mathbf{y}_{k+1} + \mathbf{y}_k, \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k \rangle \\ &\quad - \beta \langle -\mathbf{y}_{k+1} + \mathbf{y}_k, -\mathbf{y}_{k+1} + \mathbf{y}^\mu \rangle \end{aligned} \tag{16}$$

$$\begin{aligned} &= \frac{1}{2\beta} \|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 \\ &\quad + \frac{\beta}{2} \|\mathbf{y}_k - \mathbf{y}^\mu\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}^\mu\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}_k\|^2 \\ &\quad + \langle -\mathbf{y}_{k+1} + \mathbf{y}_k, \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k \rangle. \end{aligned} \tag{17}$$

On the other hand, (14) gives

$$\langle \hat{\nabla} g(\mathbf{y}_k), \mathbf{y}_k - \mathbf{y} \rangle + \langle \boldsymbol{\lambda}_k, -\mathbf{y}_k + \mathbf{y} \rangle = 0. \tag{18}$$

Letting $\mathbf{y} = \mathbf{y}_k$ in (14) and $\mathbf{y} = \mathbf{y}_{k+1}$ in (18), and adding them together, we have:

$$\langle \hat{\nabla} g(\mathbf{y}_{k+1}) - \hat{\nabla} g(\mathbf{y}_k), \mathbf{y}_{k+1} - \mathbf{y}_k \rangle + \langle \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k, -\mathbf{y}_{k+1} + \mathbf{y}_k \rangle = 0.$$

By the monotonicity of ∂g we know that the first term of the above equality is non-negative. Thus, we have:

$$\langle \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k, -\mathbf{y}_{k+1} + \mathbf{y}_k \rangle \leq 0. \tag{19}$$

Plugging it into (17), we have the conclusion. \square

Lemma 7. For Algorithm 1, we have

$$\begin{aligned} & f(\mathbf{x}_{k+1}) + g(\mathbf{y}_{k+1}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \\ &\leq \frac{1}{2\beta} \|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 \\ &\quad + \frac{\beta}{2} \|\mathbf{y}_k - \mathbf{y}^\mu\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}^\mu\|^2 \\ &\quad - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}_k\|^2 \end{aligned} \tag{20}$$

Furthermore, if F is ξ -monotone on $\mathcal{C}_=$, we have

$$\begin{aligned} & c \|\mathbf{x}_{k+1} - \mathbf{x}^\mu\|_2^\xi + f(\mathbf{x}_{k+1}) + g(\mathbf{y}_{k+1}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \\ &\leq \frac{1}{2\beta} \|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 \\ &\quad + \frac{\beta}{2} \|\mathbf{y}_k - \mathbf{y}^\mu\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}^\mu\|^2 \\ &\quad - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}_k\|^2. \end{aligned} \tag{21}$$

Proof. From the convexity of $f_k(\mathbf{x})$ and $g(\mathbf{y})$ and using Proposition 2 and (13), we have

$$\begin{aligned}
& f(\mathbf{x}_{k+1}) + g(\mathbf{y}_{k+1}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \\
& \leq f_k(\mathbf{x}_{k+1}) + g(\mathbf{y}_{k+1}) - f_k(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \\
& \leq \langle \hat{\nabla} f_k(\mathbf{x}_{k+1}), \mathbf{x}_{k+1} - \mathbf{x}^\mu \rangle + \langle \hat{\nabla} g(\mathbf{y}_{k+1}), \mathbf{y}_{k+1} - \mathbf{y}^\mu \rangle \\
& \quad + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \\
& \leq \frac{1}{2\beta} \|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 \\
& \quad + \frac{\beta}{2} \|\mathbf{y}_k - \mathbf{y}^\mu\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}^\mu\|^2 \\
& \quad - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}_k\|^2.
\end{aligned} \tag{22}$$

If F is ξ -monotone on $\mathcal{C}_=$, again by Proposition 2, we can add the term $c\|\mathbf{x}_{k+1} - \mathbf{x}^\mu\|_2^\xi$ to the first line and the inequality still holds. \square

Theorem 6. For Algorithm 1, we have

$$\begin{aligned}
f(\mathbf{x}_{k+1}) - f(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) &\rightarrow 0, \\
f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) &\rightarrow 0, \\
\mathbf{x}_{k+1} - \mathbf{y}_{k+1} &\rightarrow \mathbf{0},
\end{aligned}$$

as $k \rightarrow \infty$. Furthermore, if F is ξ -monotone on $\mathcal{C}_=$, we have

$$\mathbf{x}_{k+1} \rightarrow \mathbf{x}^\mu, \quad k \rightarrow \infty$$

Proof of Theorem 6. Proof From Lemma 2 and (20), we have

$$\begin{aligned}
& \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 + \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}_k\|^2 \\
& \leq \frac{1}{2\beta} \|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 \\
& \quad + \frac{\beta}{2} \|\mathbf{y}_k - \mathbf{y}^\mu\|^2 - \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}^\mu\|^2.
\end{aligned} \tag{23}$$

Summing over $k = 0, \dots, \infty$, we have

$$\begin{aligned}
& \sum_{k=0}^{\infty} \left(\frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 + \frac{\beta}{2} \|\mathbf{y}_{k+1} - \mathbf{y}_k\|^2 \right) \\
& \leq \frac{1}{2\beta} \|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^\mu\|^2 + \frac{\beta}{2} \|\mathbf{y}_0 - \mathbf{y}^\mu\|^2.
\end{aligned}$$

from which we deduce that $\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k \rightarrow \mathbf{0}$ and $\mathbf{y}_{k+1} - \mathbf{y}_k \rightarrow \mathbf{0}$. Moreover, $\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2$ and $\|\mathbf{y}_k - \mathbf{y}^\mu\|^2$ are bounded for all k , as well as $\|\boldsymbol{\lambda}_k\|$. Since

$$\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k = \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1}) = \beta(\mathbf{x}_{k+1} - \mathbf{x}^\mu) + \beta(\mathbf{y}_{k+1} - \mathbf{y}^\mu)$$

we deduce that $\mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rightarrow \mathbf{0}$ and $\mathbf{x}_{k+1} - \mathbf{x}^\mu$ is also bounded.

From (15) and the convexity of f and g , and using Proposition 2, we have:

$$\begin{aligned}
& f(\mathbf{x}_{k+1}) - f(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) \\
& \leq f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) \\
& \leq -\langle \boldsymbol{\lambda}_{k+1}, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle + \beta \langle \mathbf{y}_{k+1} - \mathbf{y}_k, \mathbf{x}_{k+1} - \mathbf{x}^\mu \rangle \rightarrow 0.
\end{aligned}$$

On the other hand, from (8), (9), and (10), we have:

$$\begin{aligned}
& f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) \\
& \geq f(\mathbf{x}_{k+1}) - f(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) \\
& \geq \langle -\boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{x}^\mu \rangle + \langle \boldsymbol{\lambda}^\mu, \mathbf{y}_{k+1} - \mathbf{y}^\mu \rangle \\
& = -\langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} \rangle \rightarrow 0.
\end{aligned}$$

Thus, we have $f(\mathbf{x}_{k+1}) - f(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) \rightarrow 0$ and $f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}^\mu) + g(\mathbf{y}_{k+1}) - g(\mathbf{y}^\mu) \rightarrow 0, k \rightarrow \infty$.

If F is ξ -monotone on $\mathcal{C}_=$, from Lemma 2 and (21) we have

$$\begin{aligned} & c\|\mathbf{x}_{k+1} - \mathbf{x}^\mu\|_2^\xi + \frac{1}{2\beta}\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 + \frac{\beta}{2}\|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2 \\ & \leq \frac{1}{2\beta}\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 - \frac{1}{2\beta}\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 \\ & + \frac{\beta}{2}\|-\mathbf{y}_k + \mathbf{y}^\mu\|^2 - \frac{\beta}{2}\|-\mathbf{y}_{k+1} + \mathbf{y}^\mu\|^2 \end{aligned} \quad (24)$$

From this we deduce that:

$$\begin{aligned} & c\|\mathbf{x}_{k+1} - \mathbf{x}^\mu\|_2^\xi + \sum_{k=0}^{\infty} \left(\frac{1}{2\beta}\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 + \frac{\beta}{2}\|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2 \right) \\ & \leq \frac{1}{2\beta}\|\boldsymbol{\lambda}^0 - \boldsymbol{\lambda}^\mu\|^2 + \frac{\beta}{2}\|-\mathbf{y}^0 + \mathbf{y}^\mu\|^2. \end{aligned}$$

Therefore, $\|\mathbf{x}_{k+1} - \mathbf{x}^\mu\|_2 \rightarrow 0, k \rightarrow \infty$. □

Lemma 8. For Algorithm 1, we have

$$\begin{aligned} & \frac{1}{2\beta}\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 + \frac{\beta}{2}\|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2 \\ & \leq \frac{1}{2\beta}\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}\|^2 + \frac{\beta}{2}\|-\mathbf{y}_k + \mathbf{y}_{k-1}\|^2. \end{aligned} \quad (25)$$

Furthermore, if F is ξ -monotone on $\mathcal{C}_=$, we have

$$\begin{aligned} & c\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 + \frac{1}{2\beta}\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 + \frac{\beta}{2}\|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2 \\ & \leq \frac{1}{2\beta}\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}\|^2 + \frac{\beta}{2}\|-\mathbf{y}_k + \mathbf{y}_{k-1}\|^2. \end{aligned} \quad (26)$$

Proof of Lemma 8. (15) gives:

$$\begin{aligned} & \langle \hat{\nabla} f_{k-1}(\mathbf{x}_k), \mathbf{x}_k - \mathbf{x} \rangle + \langle \hat{\nabla} g(\mathbf{y}_k), \mathbf{y}_k - \mathbf{y} \rangle \\ & = -\langle \boldsymbol{\lambda}_k, \mathbf{x}_k - \mathbf{y}_k - \mathbf{x} + \mathbf{y} \rangle + \beta \langle -\mathbf{y}_k + \mathbf{y}_{k-1}, \mathbf{x}_k - \mathbf{x} \rangle. \end{aligned} \quad (27)$$

Letting $(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) = (\mathbf{x}_k, \mathbf{y}_k, \boldsymbol{\lambda}_k)$ in (15) and $(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) = (\mathbf{x}_{k+1}, \mathbf{y}_{k+1}, \boldsymbol{\lambda}_{k+1})$ in (27), and adding them together, and using (7), we have

$$\begin{aligned}
& \langle \hat{\nabla} f_k(\mathbf{x}_{k+1}) - \hat{\nabla} f_{k-1}(\mathbf{x}_k), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle + \langle \hat{\nabla} g(\mathbf{y}_{k+1}) - \hat{\nabla} g(\mathbf{y}_k), \mathbf{y}_{k+1} - \mathbf{y}_k \rangle \\
&= -\langle \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k, \mathbf{x}_{k+1} - \mathbf{y}_{k+1} - \mathbf{x}_k + \mathbf{y}_k \rangle + \beta \langle -\mathbf{y}_{k+1} + \mathbf{y}_k - (-\mathbf{y}_k + \mathbf{y}_{k-1}), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \\
&= -\frac{1}{\beta} \langle \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k, \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k - (\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}) \rangle \\
&+ \langle -\mathbf{y}_{k+1} + \mathbf{y}_k + (\mathbf{y}_k - \mathbf{y}_{k-1}), \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k + \beta \mathbf{y}_{k+1} - (\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1} + \beta \mathbf{y}_k) \rangle \\
&= \frac{1}{2\beta} [\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}\|^2 - \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 - \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k - (\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1})\|^2] \\
&+ \frac{\beta}{2} [\|-\mathbf{y}_k + \mathbf{y}_{k-1}\|^2 - \|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2 - \|-\mathbf{y}_{k+1} + \mathbf{y}_k - (-\mathbf{y}_k + \mathbf{y}_{k-1})\|^2] \\
&+ \langle -\mathbf{y}_{k+1} + \mathbf{y}_k - (-\mathbf{y}_k + \mathbf{y}_{k-1}), \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k - (\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}) \rangle \\
&= \frac{1}{2\beta} (\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}\|^2 - \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2) + \frac{\beta}{2} (\|-\mathbf{y}_k + \mathbf{y}_{k-1}\|^2 - \|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2) \\
&- \frac{1}{2\beta} \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k - (\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1})\|^2 - \frac{\beta}{2} \|-\mathbf{y}_{k+1} + \mathbf{y}_k - (-\mathbf{y}_k + \mathbf{y}_{k-1})\|^2 \\
&+ \langle -\mathbf{y}_{k+1} + \mathbf{y}_k - (-\mathbf{y}_k + \mathbf{y}_{k-1}), \boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k - (\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}) \rangle \\
&\leq \frac{1}{2\beta} (\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_{k-1}\|^2 - \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2) + \frac{\beta}{2} (\|-\mathbf{y}_k + \mathbf{y}_{k-1}\|^2 - \|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2)
\end{aligned}$$

By the convexity of f_k and f_{k-1} , we have

$$\begin{aligned}
& \langle \hat{\nabla} f_k(\mathbf{x}_{k+1}), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \geq f_k(\mathbf{x}_{k+1}) - f_k(\mathbf{x}_k), \\
& -\langle \hat{\nabla} f_{k-1}(\mathbf{x}_k), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \geq f_{k-1}(\mathbf{x}_k) - f_{k-1}(\mathbf{x}_{k+1}).
\end{aligned}$$

Adding them together gives that:

$$\begin{aligned}
& \langle \hat{\nabla} f_k(\mathbf{x}_{k+1}) - \hat{\nabla} f_{k-1}(\mathbf{x}_k), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \geq f_k(\mathbf{x}_{k+1}) - f_{k-1}(\mathbf{x}_{k+1}) - f_k(\mathbf{x}_k) + f_{k-1}(\mathbf{x}_k) \\
& = \langle F(\mathbf{x}_{k+1}) - F(\mathbf{x}_k), \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \geq 0.
\end{aligned}$$

Thus by the monotonicity of F and $\hat{\nabla} g$, (25) follows. \square

Theorem 7. *If F is monotone on $\mathcal{C}_=$, then for Algorithm 1, we have*

$$\begin{aligned}
-\|\boldsymbol{\lambda}^\mu\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}} &\leq f(\mathbf{x}_{K+1}) + g(\mathbf{y}_{K+1}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) \\
&\leq f_K(\mathbf{x}_{K+1}) + g(\mathbf{y}_{K+1}) - f_K(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) \\
&\leq \frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^\mu\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}},
\end{aligned} \tag{28}$$

$$\|\mathbf{x}_{K+1} - \mathbf{y}_{K+1}\| \leq \sqrt{\frac{\Delta^\mu}{\beta(K+1)}}, \tag{29}$$

where $\Delta^\mu \triangleq \frac{1}{\beta} \|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^\mu\|^2 + \beta \|\mathbf{y}_0 - \mathbf{y}^\mu\|^2$.

Furthermore, if F is ξ -monotone on $\mathcal{C}_=$, we have

$$c \|\hat{\mathbf{x}}_{K+1} - \mathbf{x}^\mu\|^\xi \leq \frac{\Delta^\mu}{K+1}, \tag{30}$$

$$c \|\mathbf{x}_{K+1} - \mathbf{x}^\mu\|_2^\xi \leq \frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}}. \tag{31}$$

where $\hat{\mathbf{x}}_{K+1} = \frac{1}{K+1} \sum_{k=1}^{K+1} \mathbf{x}_k$.

Proof of Theorem 7. Summing (23) over $k = 0, 1, \dots, K$ and using the monotonicity of $\frac{1}{2\beta}\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|^2 + \frac{\beta}{2}\|-\mathbf{y}_{k+1} + \mathbf{y}_k\|^2$ from Lemma 8, we have:

$$\frac{1}{\beta}\|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}_K\|^2 + \beta\|-\mathbf{y}_{K+1} + \mathbf{y}_K\|^2 \leq \frac{1}{K+1} \left(\frac{1}{\beta}\|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^\mu\|^2 + \beta\|-\mathbf{y}_0 + \mathbf{y}^\mu\|^2 \right) \quad (32)$$

From the above we deduce that

$$\begin{aligned} \beta\|\mathbf{x}_{K+1} - \mathbf{y}_{K+1}\| &= \|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}_K\| \leq \sqrt{\frac{\beta\Delta^\mu}{(K+1)}}, \\ \|-\mathbf{y}_{K+1} + \mathbf{y}_K\| &\leq \sqrt{\frac{\Delta^\mu}{\beta(K+1)}}. \end{aligned}$$

On the other hand, (23) gives:

$$\begin{aligned} &\frac{1}{2\beta}\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}^\mu\|^2 + \frac{\beta}{2}\|-\mathbf{y}_{k+1} + \mathbf{y}^\mu\|^2 \\ &\leq \frac{1}{2\beta}\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}^\mu\|^2 + \frac{\beta}{2}\|-\mathbf{y}_k + \mathbf{y}^\mu\|^2 \\ &\leq \frac{1}{2\beta}\|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^\mu\|^2 + \frac{\beta}{2}\|-\mathbf{y}_0 + \mathbf{y}^\mu\|^2 = \frac{1}{2}\Delta^\mu. \end{aligned}$$

Hence, we have:

$$\|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}^{\mu_t}\| \leq \sqrt{\beta\Delta^{\mu_t}}, \quad (33)$$

$$\|-\mathbf{y}_{K+1} + \mathbf{y}^{\mu_t}\| \leq \sqrt{\frac{\Delta^{\mu_t}}{\beta}}. \quad (34)$$

Then from (16) and the convexity of f and g , we have:

$$\begin{aligned} &f(\mathbf{x}_{K+1}) - f(\mathbf{x}^\mu) + g(\mathbf{y}_{K+1}) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{K+1} - \mathbf{y}_{K+1} \rangle \\ &\leq f_K(\mathbf{x}_{K+1}) - f_K(\mathbf{x}^\mu) + g(\mathbf{y}_{K+1}) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \mathbf{x}_{K+1} - \mathbf{y}_{K+1} \rangle \\ &\leq \frac{1}{\beta}\|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}^\mu\|\|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}_K\| + \|-\mathbf{y}_{K+1} + \mathbf{y}_K\|\|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}_K\| \\ &\quad + \beta\|-\mathbf{y}_{K+1} + \mathbf{y}_K\|\|-\mathbf{y}_{K+1} + \mathbf{y}^\mu\| \\ &\leq \frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}}. \end{aligned} \quad (35)$$

From Lemma 3, we have (28).

If in addition F is ξ -monotone on $\mathcal{C}_=$, using (24), we can obtain the following inequality similar to (32):

$$\begin{aligned} &c \frac{\sum_{k=0}^K \|\mathbf{x}_{k+1} - \mathbf{x}^\mu\|_2^\xi}{K+1} + \frac{1}{\beta}\|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}_K\|^2 + \beta\|-\mathbf{y}_{K+1} + \mathbf{y}_K\|^2 \\ &\leq \frac{1}{K+1} \left(\frac{1}{\beta}\|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^\mu\|^2 + \beta\|-\mathbf{y}_0 + \mathbf{y}^\mu\|^2 \right) \end{aligned}$$

By the convexity of $\|\cdot\|_2^\xi$ we have $c\|\hat{\mathbf{x}}_{K+1} - \mathbf{x}^\mu\|_2^\xi \leq \frac{\Delta^\mu}{K+1}$. And from (35) we can see that $c\|\mathbf{x}_{K+1} - \mathbf{x}^\mu\|_2^\xi \leq \frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}}$. \square

Theorem 8. *If F is monotone on $\mathcal{C}_=$, then for Algorithm 1, we have*

$$\begin{aligned} |f(\hat{\mathbf{x}}_{K+1}) + g(\hat{\mathbf{y}}_{K+1}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu)| &\leq \frac{\Delta^\mu}{2(K+1)} + \frac{2\sqrt{\beta\Delta^\mu}\|\boldsymbol{\lambda}^\mu\|}{\beta(K+1)}, \\ \|\hat{\mathbf{x}}_{K+1} - \hat{\mathbf{y}}_{K+1}\| &\leq \frac{2\sqrt{\beta\Delta^\mu}}{\beta(K+1)} \end{aligned} \quad (36)$$

where $\hat{\mathbf{x}}_{K+1} = \frac{1}{K+1} \sum_{k=1}^{K+1} \mathbf{x}_k$, $\hat{\mathbf{y}}_{K+1} = \frac{1}{K+1} \sum_{k=1}^{K+1} \mathbf{y}_k$, and $\Delta^\mu \triangleq \frac{1}{\beta}\|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^\mu\|^2 + \beta\|\mathbf{y}_0 - \mathbf{y}^\mu\|^2$.

Proof of Theorem 8. Summing (20) over $k = 0, 1, \dots, K$, dividing both sides with $K + 1$, and using the definitions of $\hat{\mathbf{x}}_{K+1}$ and $\hat{\mathbf{y}}_{K+1}$ and the convexity of f and g , we have

$$f(\hat{\mathbf{x}}_{K+1}) + g(\hat{\mathbf{y}}_{K+1}) - f(\mathbf{x}^\mu) - g(\mathbf{y}^\mu) + \langle \boldsymbol{\lambda}^\mu, \hat{\mathbf{x}}_{K+1} - \hat{\mathbf{y}}_{K+1} \rangle \leq \frac{\Delta^\mu}{2(K+1)}.$$

From (7) and (33), we have:

$$\begin{aligned} \|\hat{\mathbf{x}}_{K+1} - \hat{\mathbf{y}}_{K+1}\| &= \frac{1}{\beta(K+1)} \left\| \sum_{k=0}^K (\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k) \right\| \\ &= \frac{1}{\beta(K+1)} \|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}_0\| \\ &\leq \frac{1}{\beta(K+1)} (\|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^\mu\| + \|\boldsymbol{\lambda}_{K+1} - \boldsymbol{\lambda}^\mu\|) \\ &\leq \frac{2\sqrt{\beta\Delta^\mu}}{\beta(K+1)} \end{aligned}$$

Finally, from Lemma 3, the conclusion follows. \square

From Proposition 1 and the fact that \mathcal{C} is compact we can see that $\lim_{\mu \rightarrow 0} \mathbf{x}^\mu$ and $\lim_{\mu \rightarrow 0} \boldsymbol{\lambda}^\mu$ exist and are finite. Let $\mathbf{x}^* = \lim_{\mu \rightarrow 0} \mathbf{x}^\mu$ and $\boldsymbol{\lambda}^* = \lim_{\mu \rightarrow 0} \boldsymbol{\lambda}^\mu$, then $\mathbf{x}^* \in \mathcal{S}_{\mathcal{C}, F}^*$.

Theorem 9. $\exists \tilde{\mu} > 0$, s.t. if $\mu_t < \tilde{\mu}$, then $\left| F(\mathbf{x}_{K+1}^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| \leq 2\left(\frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}}\right)$, $\left| F(\mathbf{x}^*)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| \leq 2\left(\frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}}\right)$ and $\left| F(\mathbf{x}^*)^\top (\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^*) \right| \leq 2\left(\frac{\Delta^\mu}{2(K+1)} + \frac{2\sqrt{\beta\Delta^\mu}\|\boldsymbol{\lambda}^*\|}{\beta(K+1)}\right)$, $\forall K \geq 0$.

Proof of Theorem 9. For simplicity we let $B(\mathbf{x})$ denote the log-barrier term $-\sum_{i=1}^m \log(-\varphi_i(\mathbf{x}))$. From Theorem 7 and 8 we have

$$\begin{aligned} &\left| F(\mathbf{x}^\mu)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^\mu) + \mu(B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^\mu)) \right| \\ &= \left| f(\mathbf{x}_{K+1}^{(t)}) - f(\mathbf{x}^\mu) + g(\mathbf{y}_{K+1}^{(t)}) - g(\mathbf{y}^\mu) \right| \\ &\leq \frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^\mu\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}} \end{aligned} \quad (37)$$

$$\begin{aligned} &\left| F(\mathbf{x}^{K+1})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^\mu) + \mu(B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^\mu)) \right| \\ &= \left| f_K(\mathbf{x}_{K+1}^{(t)}) - f_K(\mathbf{x}^\mu) + g(\mathbf{y}_{K+1}^{(t)}) - g(\mathbf{y}^\mu) \right| \\ &\leq \frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^\mu\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}} \end{aligned} \quad (38)$$

$$\begin{aligned} &\left| F(\mathbf{x}^\mu)^\top (\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^\mu) + \mu(B(\hat{\mathbf{y}}_{K+1}^{(t)}) - B(\mathbf{y}^\mu)) \right| \\ &= \left| f(\hat{\mathbf{x}}_{K+1}^{(t)}) - f(\mathbf{x}^\mu) + g(\hat{\mathbf{y}}_{K+1}^{(t)}) - g(\mathbf{y}^\mu) \right| \\ &\leq \frac{\Delta^\mu}{2(K+1)} + \frac{2\sqrt{\beta\Delta^\mu}\|\boldsymbol{\lambda}^\mu\|}{\beta(K+1)}. \end{aligned} \quad (39)$$

From Proposition 1 in App. A we know that when $\mu \rightarrow 0$, $\mathbf{x}^\mu \rightarrow \mathbf{x}^*$, $\boldsymbol{\lambda}^\mu \rightarrow \boldsymbol{\lambda}^*$ and $\mu(B(\mathbf{y}_{K+1}^{(t)} - B(\mathbf{y}^\mu))) \rightarrow 0$, so $\exists \tilde{\mu} > 0$, s.t. if $\mu_t < \tilde{\mu}$, then we have

$$\begin{aligned} \left| F(\mathbf{x}^*)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| &\leq 2 \left(\frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}} \right) \\ \left| F(\mathbf{x}_{K+1}^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| &\leq 2 \left(\frac{\Delta^\mu}{K+1} + \frac{2\Delta^\mu}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta^\mu}{\beta(K+1)}} \right) \\ \left| F(\mathbf{x}^*)^\top (\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^*) \right| &\leq 2 \left(\frac{\Delta^\mu}{2(K+1)} + \frac{2\sqrt{\beta\Delta^\mu}\|\boldsymbol{\lambda}^*\|}{\beta(K+1)} \right). \end{aligned}$$

□

To make the dependencies of the constants clear, here we restate Theorem 2 and Theorem 3 and provide their proofs.

Theorem 10 (restatement of Theorem 2). *Given an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$ monotone on $\mathcal{C}_=$ (Def. 1), and either F is strictly monotone on \mathcal{C} or one of φ_i is strictly convex. Assume there exists $r > 0$ or $s > 0$ such that F is star- ξ -monotone—as per Def. 1—on either $\hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$, resp. Let $\Delta \triangleq \frac{1}{\beta} \|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^*\|^2 + \beta \|\mathbf{y}_0 - \mathbf{y}^*\|^2$.*

Let $\mathbf{x}_K^{(t)}$ and $\hat{\mathbf{x}}_K^{(t)} \triangleq \frac{1}{K} \sum_{k=1}^K \mathbf{x}_k^{(t)}$ denote the last and average iterate of Algorithm 1, respectively, run with sufficiently small μ_{-1} . Then there exists $K_0 \in \mathbb{N}$, K_0 depends on r or s , s.t. $\forall K > K_0$, for all $t \in [T]$, we have that:

1. $\|\mathbf{x}_K^{(t)} - \mathbf{x}^*\| \leq \left(\frac{4}{c} \left(\frac{\Delta}{K} + \frac{2\Delta}{\sqrt{K}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta K}} \right) \right)^{1/\xi}$.
2. If in addition F is ξ -monotone on $\mathcal{C}_=$, we have $\|\hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^*\| \leq \left(\frac{2\Delta}{cK} \right)^{1/\xi}$.
3. Moreover, if F is L -Lipschitz on $\mathcal{C}_=$ —as per Assumption 1—then $\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) \leq M_0 \left(\frac{4}{c} \left(\frac{\Delta}{K} + \frac{2\Delta}{\sqrt{K}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta K}} \right) \right)^{1/\xi}$ and $\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) \leq M_0 \left(\frac{2\Delta}{c(K+1)} \right)^{1/\xi}$,

where $M_0 = DL + M$ is a linear function of L , see the proof of Lemma 1 in App. B.

Proof of Theorem 10. Without loss of generality, we suppose that F is star- ξ -monotone on $\hat{\mathcal{C}}_r$. Since you $\phi(\mathbf{y}_k) \leq \mathbf{0}$, from (29) we can see that $\exists K_0 \in \mathbb{N}$, K_0 depends on r , s.t. $\forall K \geq K_0$, $\mathbf{x}_k \in \hat{\mathcal{C}}_r$, so if $\mu_t < \tilde{\mu}$ ($\tilde{\mu}$ is defined in Theorem 9), by Theorem 9 we have

$$\begin{aligned} c \|\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*\|^\xi &\leq \left| (F(\mathbf{x}_{K+1}^{(t)})^\top - F(\mathbf{x}^*)^\top) (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| \\ &\leq \left| F(\mathbf{x}^*)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| + \left| F(\mathbf{x}_{K+1}^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| \\ &\leq 4 \left(\frac{\Delta}{K+1} + \frac{2\Delta}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta(K+1)}} \right). \end{aligned}$$

So we have

$$\|\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*\| \leq \left(\frac{4}{c} \left(\frac{\Delta}{K+1} + \frac{2\Delta}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta(K+1)}} \right) \right)^{1/\xi}.$$

If in addition F is ξ -monotone on $\mathcal{C}_=$, then from (30) we know that when $\hat{\mu}$ is small enough, we have

$$\|\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^*\| \leq \left(\frac{2\Delta}{c(K+1)} \right)^{1/\xi}.$$

If F is L -Lipschitz on $\mathcal{C}_=$, then from Lemma 1 we can see that

$$\begin{aligned} \mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) &\leq M_0 \left(\frac{4}{c} \left(\frac{\Delta}{K+1} + \frac{2\Delta}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta(K+1)}} \right) \right)^{1/\xi}, \\ \mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) &\leq M_0 \left(\frac{2\Delta}{c(K+1)} \right)^{1/\xi}. \end{aligned}$$

□

Theorem 11 (restatement of Theorem 3). *Given an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$, assume: (i) F is monotone on $\mathcal{C}_=$, as per Def. 1; (ii) either F is strictly monotone on \mathcal{C} or one of φ_i is strictly convex; and (iii)*

$\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} = a > 0$, where $S \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$. Let $\Delta \triangleq \frac{1}{\beta} \|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^*\|^2 + \beta \|\mathbf{y}_0 - \mathbf{y}^*\|^2$.*

Let $\mathbf{x}_K^{(t)}$ and $\hat{\mathbf{x}}_K^{(t)} \triangleq \frac{1}{K} \sum_{k=1}^K \mathbf{x}_k^{(t)}$ denote the last and average iterate of Algorithm 1, respectively, run with sufficiently small μ_{-1} . Then there exists $K_0 \in \mathbb{N}$, K_0 depends on r or s , s.t. $\forall t \in [T]$, $\forall K > K_0$, we have that:

1. $\|\mathbf{x}_K^{(t)} - \mathbf{x}^*\| \leq \frac{2}{a} \left(\frac{\Delta}{K} + \frac{2\Delta}{\sqrt{K}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta K}} \right)$.
2. If in addition $\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x}^*)^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} = b > 0$ (with $S \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$), then $\|\hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^*\| \leq \frac{2}{b} \left(\frac{\Delta}{2K} + \frac{2\sqrt{\beta\Delta}\|\boldsymbol{\lambda}^*\|}{\beta K} \right)$.
3. Moreover, if F is L -Lipschitz on $\mathcal{C}_=$ —as per Assumption 1—then $\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) \leq \frac{2M_0}{a} \left(\frac{\Delta}{K} + \frac{2\Delta}{\sqrt{K}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta K}} \right)$ and $\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) \leq \frac{2M_0}{b} \left(\frac{\Delta}{2K} + \frac{2\sqrt{\beta\Delta}\|\boldsymbol{\lambda}^*\|}{\beta K} \right)$,

where $M_0 = DL + M$ is a linear function of L , see the proof of Lemma 1 in App. B.

Proof of Theorem 11. Without loss of generality, we suppose $\inf_{\mathbf{x} \in \hat{\mathcal{C}}_r \setminus \{\mathbf{x}^*\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} = a > 0$.

When $K \geq K_0$ (K_0 and $\tilde{\mu}$ are defined in the proof of Theorem 10 and in Theorem 9, resp.), by Theorem 9 we have that

$$\begin{aligned} \|\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*\| &\leq \frac{1}{a} \left| F(\mathbf{x}_{K+1}^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| \\ &\leq \frac{2}{a} \left(\frac{\Delta}{K+1} + \frac{2\Delta}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta(K+1)}} \right). \end{aligned}$$

Similarly, if $\inf_{\mathbf{x} \in \tilde{\mathcal{C}}_s \setminus \{\mathbf{x}^*\}} F(\mathbf{x}^*)^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} = b > 0$, we have that

$$\|\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^*\| \leq \frac{2}{b} \left(\frac{\Delta}{2(K+1)} + \frac{2\sqrt{\beta\Delta}\|\boldsymbol{\lambda}^*\|}{\beta(K+1)} \right).$$

If F is L -Lipschitz on $\mathcal{C}_=$, then from Lemma 1 we can see that

$$\begin{aligned} \mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) &\leq \frac{2M_0}{a} \left(\frac{\Delta}{K+1} + \frac{2\Delta}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^*\| \sqrt{\frac{\Delta}{\beta(K+1)}} \right), \\ \mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) &\leq \frac{2M_0}{b} \left(\frac{\Delta}{2(K+1)} + \frac{2\sqrt{\beta\Delta}\|\boldsymbol{\lambda}^*\|}{\beta(K+1)} \right). \end{aligned}$$

□

From the above proofs, we can see that by setting μ_{-1} small enough, Theorem 2 and 3 hold true. But since we do not know exactly how small μ_{-1} should be, in practice, we can set a small μ_{-1} , and

μ_t could be smaller than any fixed positive number after a very small number of outer loops, as μ_t decays exponentially. Thus, Algorithm 1 is actually parameter-free.

Note that the assumption (iii) in Theorem 3 is the weakening of the assumption

$\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x}^*)^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0$, where $S = \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$. In fact, by the monotonicity of F , we

have $F(\mathbf{x})^\top (\mathbf{x} - \mathbf{x}^*) \geq F(\mathbf{x}^*)^\top (\mathbf{x} - \mathbf{x}^*)$, so if $\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x}^*)^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0$, there must be

$$\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0.$$

B.4 PARAMETER-FREE CONVERGENCE RATE

In this section, we give a convergence rate taking into account both the inner and the outer loop when the operator F is L -Lipschitz.

First, we bound $\|\mathbf{x}^* - \mathbf{x}^\mu\|$ under the assumptions in Theorem 2 and Theorem 3, resp, and give the following lemma:

Lemma 9. *For monotone F , then for any $\mu > 0$, we have:*

(i) *If F is star- ξ -monotone (Def. 1), then $\|\mathbf{x}^* - \mathbf{x}^\mu\| \leq \left(\frac{m}{c}\mu\right)^{\frac{1}{\xi}}$;*

(ii) *If $a \triangleq \inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0$, where $S \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$, then $\|\mathbf{x}^* - \mathbf{x}^\mu\| \leq \frac{m}{a}\mu$.*

Proof. Consider convex problem

$$\begin{cases} \min_{\mathbf{x}} F(\mathbf{x}^\mu)^\top \mathbf{x} \\ \text{s.t. } \varphi(\mathbf{x}) \leq \mathbf{0} \\ \mathbf{C}\mathbf{x} = \mathbf{d} \end{cases} \quad (\mathcal{P})$$

The Lagrangian of (\mathcal{P}) is

$$L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = F(\mathbf{x}^\mu)^\top \mathbf{x} + \boldsymbol{\lambda}^\top \varphi(\mathbf{x}) + \boldsymbol{\nu}^\top (\mathbf{C}\mathbf{x} - \mathbf{d}).$$

There exists $\bar{\boldsymbol{\lambda}}^\mu > \mathbf{0}$ and $\boldsymbol{\nu}^\mu$, s.t. $(\mathbf{x}^\mu, \bar{\boldsymbol{\lambda}}^\mu, \boldsymbol{\nu}^\mu)$ is a KKT point of (KKT-2). By the stationarity condition in (KKT-2), we have that

$$g(\bar{\boldsymbol{\lambda}}^\mu, \boldsymbol{\nu}^\mu) = \inf_{\mathbf{x}} L(\mathbf{x}^\mu, \bar{\boldsymbol{\lambda}}^\mu, \boldsymbol{\nu}^\mu) = F(\mathbf{x}^\mu)^\top \mathbf{x}^\mu + \bar{\boldsymbol{\lambda}}^\top \varphi(\mathbf{x}^\mu) + \boldsymbol{\nu}^\top (\underbrace{\mathbf{C}\mathbf{x}^\mu - \mathbf{d}}_{=0}).$$

Note that by the complementarity slackness condition in (KKT-2), we have $\bar{\boldsymbol{\lambda}}^\top \varphi(\mathbf{x}^\mu) = -m\mu$.

Since \mathbf{x}^* is a feasible point of (\mathcal{P}) , $(\bar{\boldsymbol{\lambda}}^\mu, \boldsymbol{\nu}^\mu)$ dual feasible, thus by the duality theory, we have:

$$F(\mathbf{x}^\mu)^\top \mathbf{x}^* \geq g(\bar{\boldsymbol{\lambda}}^\mu, \boldsymbol{\nu}^\mu) = F(\mathbf{x}^\mu)^\top \mathbf{x}^\mu - m\mu,$$

from where we deduce that

$$F(\mathbf{x}^\mu)^\top (\mathbf{x}^\mu - \mathbf{x}^*) \leq m\mu. \quad (40)$$

Therefore,

(i) If F is star- ξ -monotone:

Since $\mathbf{x}^* \in \mathcal{S}_{\hat{\mathcal{C}}, F}^*$, we have

$$F(\mathbf{x}^*)^\top (\mathbf{x}^\mu - \mathbf{x}^*) \geq 0. \quad (41)$$

Subtract (40) by (41) and using the star- ξ -monotonicity of F , we obtain

$$c \|\mathbf{x}^\mu - \mathbf{x}^*\|^\xi \leq (F(\mathbf{x}^\mu) - F(\mathbf{x}^*))^\top (\mathbf{x}^\mu - \mathbf{x}^*) \leq m\mu.$$

(ii) If $a \triangleq \inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} > 0$, where $S \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$, since $\mathbf{x}^\mu \in S$, by (40) we have that

$$a \|\mathbf{x}^\mu - \mathbf{x}^*\| \leq m\mu.$$

□

We give another lemma that would be used in the proofs of our main results in this section.

Lemma 10. *Assume F is monotone and L -smooth (Assumption 1), then we have*

$$\begin{aligned} & \left| F(\mathbf{x}^*)^\top(\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t(B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| \\ & \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + m\mu_t \\ & + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| + \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right), \end{aligned} \quad (42)$$

$$\begin{aligned} & \left| F(\mathbf{x}_K^{(t)})^\top(\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t(B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| \\ & \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + M \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\|, \end{aligned} \quad (43)$$

and

$$\begin{aligned} & \left| F(\mathbf{x}^*)^\top(\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t(B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| \\ & \leq \frac{\Delta^{\mu_t}}{2(K+1)} + \frac{2\sqrt{\beta\Delta^{\mu_t}}\|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K+1)} + m\mu_t \\ & + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| + \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right), \end{aligned} \quad (44)$$

where $M = \sup_{\mathbf{x} \in \mathcal{C}} \|F(\mathbf{x})\|$, and $B(\mathbf{x}) = -\sum_{i=1}^m \log(-\varphi_i(\mathbf{x}))$.

Proof. Note that

$$\left\| \mathbf{x}_{K+1}^{(t)} - \mathbf{x}^{\mu_t} \right\| = \left\| \mathbf{x}_{K+1}^{(t)} - \mathbf{y}^{\mu_t} \right\| \leq \left\| \mathbf{x}_{K+1}^{(t)} - \mathbf{y}_{K+1}^{(t)} \right\| + \left\| \mathbf{y}_{K+1}^{(t)} - \mathbf{y}^{\mu_t} \right\|. \quad (45)$$

Recall that (29) gives

$$\left\| \mathbf{x}_{K+1}^{(t)} - \mathbf{y}_{K+1}^{(t)} \right\| \leq \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}},$$

and (34) gives

$$\left\| \mathbf{y}_{K+1}^{(t)} - \mathbf{y}^{\mu_t} \right\| = \sqrt{\frac{\Delta^{\mu_t}}{\beta}}.$$

Plugging (29) and (34) into (45), we have

$$\left\| \mathbf{x}_{K+1}^{(t)} - \mathbf{x}^{\mu_t} \right\| \leq \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}}. \quad (46)$$

Note that

$$\begin{aligned} & F(\mathbf{x}^*)^\top(\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t(B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \\ & = F(\mathbf{x}^{\mu_t})^\top(\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}) + \mu_t(B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) + F(\mathbf{x}^{\mu_t})^\top(\mathbf{x}^{\mu_t} - \mathbf{x}^*) \\ & \quad + (F(\mathbf{x}^*) - F(\mathbf{x}^{\mu_t}))^\top(\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*). \end{aligned}$$

Thus by the L -Lipschitzness of F , using (40), (46) and (37) in the proof of Thm. 9, we have

$$\begin{aligned}
& \left| F(\mathbf{x}^*)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| \\
& \leq \left| F(\mathbf{x}^{\mu_t})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| + |F(\mathbf{x}_t^\mu)^\top (\mathbf{x}_t^\mu - \mathbf{x}^*)| \\
& \quad + \left| (F(\mathbf{x}^*) - F(\mathbf{x}^{\mu_t}))^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \right| \\
& \leq \left| F(\mathbf{x}^{\mu_t})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| + |F(\mathbf{x}^{\mu_t})^\top (\mathbf{x}^{\mu_t} - \mathbf{x}^*)| \\
& \quad + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\mathbf{x}^{\mu_t} - \mathbf{x}_{K+1}^{(t)}\| + \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \right) \\
& \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + m\mu_t \\
& \quad + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| + \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right).
\end{aligned}$$

Using (38), we have

$$\begin{aligned}
& \left| F(\mathbf{x}_{K+1}^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| \\
& = \left| F(\mathbf{x}^{\mu_t})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) + F(\mathbf{x}_{K+1}^{(t)})^\top (\mathbf{x}^{\mu_t} - \mathbf{x}^*) \right| \\
& \leq \left| F(\mathbf{x}^{\mu_t})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| + M \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \\
& \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + M \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\|.
\end{aligned}$$

Similarly, recall that (36) gives

$$\|\hat{\mathbf{x}}_{K+1} - \hat{\mathbf{y}}_{K+1}\| \leq \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)}.$$

By (34) and the convexity of $\|\cdot\|$, we have

$$\|\hat{\mathbf{y}}_{K+1}^{(t)} - \mathbf{y}^{\mu_t}\| = \left\| \frac{1}{K+1} \sum_{k=1}^{K+1} (\mathbf{y}_{k+1}^{(t)} - \mathbf{y}^{\mu_t}) \right\| \leq \sqrt{\frac{\Delta^{\mu_t}}{\beta}}. \quad (47)$$

Note that Thus we have

$$\|\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}\| = \|\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{y}^{\mu_t}\| \leq \|\hat{\mathbf{x}}_{K+1}^{(t)} - \hat{\mathbf{y}}_{K+1}^{(t)}\| + \|\hat{\mathbf{y}}_{K+1}^{(t)} - \mathbf{y}^{\mu_t}\|. \quad (48)$$

Plugging (36) and (47) into (48), we have

$$\|\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}\| \leq \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}}. \quad (49)$$

Using the L -Lipschitzness of F , (39) and (49), we have

$$\begin{aligned}
& \left| F(\mathbf{x}^*)^\top (\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\hat{\mathbf{y}}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| \\
& \leq \left| F(\mathbf{x}^{\mu_t})^\top (\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}) + \mu_t (B(\hat{\mathbf{y}}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| + |F(\mathbf{x}^{\mu_t})^\top (\mathbf{x}^{\mu_t} - \mathbf{x}^*)| \\
& \quad + \left| (F(\mathbf{x}^*) - F(\mathbf{x}^{\mu_t}))^\top (\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^*) \right| \\
& \leq \left| F(\mathbf{x}^{\mu_t})^\top (\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}) + \mu_t (B(\hat{\mathbf{y}}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| + |F(\mathbf{x}^{\mu_t})^\top (\mathbf{x}^{\mu_t} - \mathbf{x}^*)| \\
& \quad + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}\| + \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \right) \\
& \leq \frac{\Delta^{\mu_t}}{2(K+1)} + \frac{2\sqrt{\beta\Delta^{\mu_t}}\|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K+1)} + m\mu_t \\
& \quad + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| + \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right).
\end{aligned}$$

□

Now we are ready to give our main theorems in this section.

Theorem 12 (Complete convergence rate for star- ξ -monotone operator). *Given an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$ monotone and L -Lipschitz on $\mathcal{C}_=$ (Def. 1, 1), and either F is strictly monotone on \mathcal{C} or one of φ_i is strictly convex. Assume there exists $r > 0$ or $s > 0$ such that F is star- ξ -monotone—as per Def. 1—on either $\hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$, resp. Let $\Delta^{\mu_t} \triangleq \frac{1}{\beta} \|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^{\mu_t}\|^2 + \beta \|\mathbf{y}_0 - \mathbf{y}^{\mu_t}\|^2$,*

Let $\mathbf{x}_K^{(t)}$ and $\hat{\mathbf{x}}_K^{(t)} \triangleq \frac{1}{K} \sum_{k=1}^K \mathbf{x}_k^{(t)}$ denote the last and average iterate of Algorithm 1, respectively, run with sufficiently small μ_{t-1} . Then there exists $K_0 \in \mathbb{N}$, K_0 depends on r or s , s.t. $\forall K > K_0$, for all $t \in [T]$, we have that:

$$\begin{aligned}
1. \quad \|\mathbf{x}_K^{(t)} - \mathbf{x}^*\| & \leq \left(\frac{2}{c} \left(\frac{\Delta^{\mu_t}}{K} + \frac{2\Delta^{\mu_t}}{\sqrt{K}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta K}} \right) + \frac{LM_1}{c} \left(\frac{m}{c} \mu_{-1} \right)^{\frac{1}{\xi}} \delta^{\frac{t+1}{\xi}} \right)^{\frac{1}{\xi}}, \quad \text{and} \\
\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) & \leq M_0 \left(\frac{2}{c} \left(\frac{\Delta^{\mu_t}}{K} + \frac{2\Delta^{\mu_t}}{\sqrt{K}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta K}} \right) + \frac{LM_1}{c} \left(\frac{m}{c} \mu_{-1} \right)^{\frac{1}{\xi}} \delta^{\frac{t+1}{\xi}} \right)^{\frac{1}{\xi}}.
\end{aligned}$$

$$\begin{aligned}
2. \quad \text{If in addition } F \text{ is } \xi\text{-monotone on } \mathcal{C}_=, \text{ we have } \|\hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^*\| & \leq \left(\frac{\Delta^{\mu_t}}{cK} \right)^{\frac{1}{\xi}} + \left(\frac{m}{c} \mu_t \right)^{\frac{1}{\xi}}, \text{ and} \\
\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) & \leq M_0 \left(\left(\frac{\Delta^{\mu_t}}{cK} \right)^{1/\xi} + \left(\frac{m}{c} \mu_t \right)^{\frac{1}{\xi}} \right)^{\frac{1}{\xi}},
\end{aligned}$$

where $M_0 = DL + M$ is a linear function of L , see the proof of Lemma 1 in App. B, and $M_1 \triangleq \left(\frac{m}{c} \mu_{-1} \right)^{\frac{1}{\xi}} + 2\sqrt{\frac{\Delta^{\mu_t}}{\beta}} + \frac{M}{L}$.

Proof. By the star- ξ -monotonicity of course F , using (42) and (43) in Lemma 10, we have

$$\begin{aligned}
& c \left\| \mathbf{x}_{K+1}^{(t)} - \mathbf{x}^* \right\|^\xi \\
& \leq (F(\mathbf{x}_K^{(t)}) - F(\mathbf{x}^*))^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \\
& = F(\mathbf{x}_K^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \\
& \quad - (F(\mathbf{x}^*)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t}))) \\
& \leq \|F(\mathbf{x}_K^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t}))\| \\
& \quad + \|F(\mathbf{x}^*)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t}))\| \\
& \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + M \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \\
& \quad + \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + m\mu_t \\
& \quad + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| + \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right) \\
& \leq 2 \left(\frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} \right) \\
& \quad + L \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| \left(\|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| + \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} + \frac{M}{L} \right) + m\mu_t \\
& \leq 2 \left(\frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} \right) \\
& \quad + L \left(\frac{m}{c} \mu_t \right)^{\frac{1}{\xi}} \left(\left(\frac{m}{c} \mu_t \right)^{\frac{1}{\xi}} + \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} + \frac{M}{L} \right)
\end{aligned}$$

where in the last inequality we use Lemma 9 (i).

We let $M_1 \triangleq \left(\frac{m}{c} \mu_{-1} \right)^{\frac{1}{\xi}} + 2\sqrt{\frac{\Delta^{\mu_t}}{\beta}} + \frac{M}{L}$, then we have

$$\left\| \mathbf{x}_{K+1}^{(t)} - \mathbf{x}^* \right\| \leq \left(\frac{2}{c} \left(\frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} \right) + \frac{LM_1}{c} \left(\frac{m}{c} \mu_t \right)^{\frac{1}{\xi}} \right)^{\frac{1}{\xi}}.$$

By Lemma 1, we obtain

$$\mathcal{G}(\mathbf{x}_{K+1}^{(t)}, \mathcal{C}) \leq M_0 \left(\frac{2}{c} \left(\frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} \right) + \frac{LM_1}{c} \left(\frac{m}{c} \mu_{-1} \right)^{\frac{1}{\xi}} \delta^{\frac{t+1}{\xi}} \right)^{\frac{1}{\xi}}.$$

If in addition F is ξ -monotone on $\mathcal{C}_=$, then from (30) we have:

$$\left\| \hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^* \right\| \leq \left(\frac{\Delta^{\mu_t}}{c(K+1)} \right)^{\frac{1}{\xi}} + \left(\frac{m}{c} \mu_t \right)^{\frac{1}{\xi}}.$$

Again by Lemma 1, we have $\mathcal{G}(\hat{\mathbf{x}}_{K+1}^{(t)}, \mathcal{C}) \leq M_0 \left(\left(\frac{\Delta^{\mu_t}}{c(K+1)} \right)^{\frac{1}{\xi}} + \left(\frac{m}{c} \mu_t \right)^{\frac{1}{\xi}} \right)$. \square

Theorem 13 (Complete convergence rate for star- ξ -monotone operator). *Given an operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$, assume: (i) F is monotone and L -smooth on $\mathcal{C}_=$, as per Def. 1, 1; (ii) either F is strictly*

monotone on \mathcal{C} or one of φ_i is strictly convex; and (iii) $\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x})^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} = a > 0$, where $S \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$. Let $\Delta^{\mu_t} \triangleq \frac{1}{\beta} \|\boldsymbol{\lambda}_0 - \boldsymbol{\lambda}^{\mu_t}\|^2 + \beta \|\mathbf{y}_0 - \mathbf{y}^{\mu_t}\|^2$.

Let $\mathbf{x}_K^{(t)}$ and $\hat{\mathbf{x}}_K^{(t)} \triangleq \frac{1}{K} \sum_{k=1}^K \mathbf{x}_k^{(t)}$ denote the last and average iterate of Algorithm 1, respectively. Then there exists $K_0 \in \mathbb{N}$, K_0 depends on r or s , s.t. $\forall t \in [T], \forall K > K_0$, we have that:

1. $\|\mathbf{x}_K^{(t)} - \mathbf{x}^*\| = \frac{1}{a} \left(\frac{\Delta^{\mu_t}}{K} + \frac{2\Delta^{\mu_t}}{\sqrt{K}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta K}} \right) + \left(\frac{mM}{a^2} + \frac{B(\mathbf{y}^{\mu_t}) - B^*}{a} \right) \mu_{-1} \delta^{t+1}$, and $\mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) \leq M_0 \left(\frac{1}{a} \left(\frac{\Delta^{\mu_t}}{K} + \frac{2\Delta^{\mu_t}}{\sqrt{K}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta K}} \right) + \left(\frac{mM}{a^2} + \frac{B(\mathbf{y}^{\mu_t}) - B^*}{a} \right) \mu_{-1} \delta^{t+1} \right)$.
2. If in addition $\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x}^*)^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} = b > 0$ (with $S \equiv \hat{\mathcal{C}}_r$ or $\tilde{\mathcal{C}}_s$), then $\|\hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^*\| \leq \frac{1}{b} \left(\frac{\Delta^{\mu_t}}{2K} + \frac{2\sqrt{\beta\Delta^{\mu_t}}\|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K)} \right) + \frac{M_3}{b} \mu_{-1} \delta^{t+1}$, and $\mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) \leq M_0 \left(\frac{1}{b} \left(\frac{\Delta^{\mu_t}}{2K} + \frac{2\sqrt{\beta\Delta^{\mu_t}}\|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K)} \right) + \frac{M_3}{b} \mu_{-1} \delta^{t+1} \right)$,

where $M = \sup_{\mathbf{x} \in \mathcal{C}} \|F(\mathbf{x})\|$, $M_0 = DL + M$ is a linear function of L , see the proof of Lemma 1 in App. B, and $M_3 \triangleq m + \frac{Lm}{a} \left(\frac{m}{a} \mu_{-1} + \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right) + B(\mathbf{y}^{\mu_t}) - B^*$.

Proof. If the barrier term $B(\mathbf{y}_{K+1}^{(t)}) \geq B(\mathbf{y}^{\mu_t})$, using (43) in Lemma 10, we have

$$\begin{aligned} & F(\mathbf{x}_K^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \\ & \leq F(\mathbf{x}_K^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \\ & \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + M \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\|. \end{aligned}$$

If $B(\mathbf{y}_{K+1}^{(t)}) \leq B(\mathbf{y}^{\mu_t})$, since B is lower bounded in any compact set, and by (46) in Lemma 9 we have $\|\hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^{\mu_t}\| \leq \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \leq 3\sqrt{\frac{\Delta^{\mu_t}}{\beta}} \triangleq M_2$, we let $B^* = \inf_{\|\mathbf{x} - \mathbf{x}^*\| \leq M_2} B(\mathbf{x})$. Then we have

$$\begin{aligned} & F(\mathbf{x}_K^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \\ & \leq \left| F(\mathbf{x}_K^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| + \mu_t (B(\mathbf{y}^{\mu_t}) - B^*) \\ & \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + M \|\mathbf{x}^{\mu_t} - \mathbf{x}^*\| + \mu_t (B(\mathbf{y}^{\mu_t}) - B^*). \end{aligned}$$

Thus in both case we have

$$\begin{aligned} & a \|\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*\| \\ & \leq F(\mathbf{x}_K^{(t)})^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) \\ & \leq \frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} + \frac{mM}{a} \mu_t + \mu_t (B(\mathbf{y}^{\mu_t}) - B^*). \end{aligned}$$

where in the first inequality assumption (iii) is used and in the last inequality Lemma 9 is used. Thus we have:

$$\|\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*\| = \frac{1}{a} \left(\frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} \right) + \left(\frac{mM}{a^2} + \frac{B(\mathbf{y}^{\mu_t}) - B^*}{a} \right) \mu_t.$$

And by Lemma 1, we obtain

$$\mathcal{G}(\mathbf{x}_{K+1}^{(t)}, \mathcal{C}) \leq M_0 \left(\frac{1}{a} \left(\frac{\Delta^{\mu_t}}{K+1} + \frac{2\Delta^{\mu_t}}{\sqrt{K+1}} + \|\boldsymbol{\lambda}^{\mu_t}\| \sqrt{\frac{\Delta^{\mu_t}}{\beta(K+1)}} \right) + \left(\frac{mM}{a^2} + \frac{B(\mathbf{y}^{\mu_t}) - B^*}{a} \right) \mu_t \right).$$

If in addition $\inf_{\mathbf{x} \in S \setminus \{\mathbf{x}^*\}} F(\mathbf{x}^*)^\top \frac{\mathbf{x} - \mathbf{x}^*}{\|\mathbf{x} - \mathbf{x}^*\|} = b > 0$, then similarly, from (44) in Lemma 10, we have

$$\begin{aligned}
& b \left\| \hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^* \right\| \\
& \leq \left| F(\mathbf{x}^*)^\top (\mathbf{x}_{K+1}^{(t)} - \mathbf{x}^*) + \mu_t (B(\mathbf{y}_{K+1}^{(t)}) - B(\mathbf{y}^{\mu_t})) \right| + \mu_t (B(\mathbf{y}^{\mu_t}) - B^*) \\
& \leq \frac{\Delta^{\mu_t}}{2(K+1)} + \frac{2\sqrt{\beta\Delta^{\mu_t}} \|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K+1)} + m\mu_t \\
& \quad + L \frac{m}{a} \mu_t \left(\frac{m}{a} \mu_t + \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right) + \mu_t (B(\mathbf{y}^{\mu_t}) - B^*) \\
& = \frac{\Delta^{\mu_t}}{2(K+1)} + \frac{2\sqrt{\beta\Delta^{\mu_t}} \|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K+1)} \\
& \quad + \left(m + \frac{Lm}{a} \left(\frac{m}{a} \mu_{-1} + \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right) + B(\mathbf{y}^{\mu_t}) - B^* \right) \mu_t,
\end{aligned}$$

Let $M_3 \triangleq m + \frac{Lm}{a} \left(\frac{m}{a} \mu_{-1} + \frac{2\sqrt{\beta\Delta^{\mu_t}}}{\beta(K+1)} + \sqrt{\frac{\Delta^{\mu_t}}{\beta}} \right) + B(\mathbf{y}^{\mu_t}) - B^*$, then we have

$$\left\| \hat{\mathbf{x}}_{K+1}^{(t)} - \mathbf{x}^* \right\| \leq \frac{1}{b} \left(\frac{\Delta^{\mu_t}}{2(K+1)} + \frac{2\sqrt{\beta\Delta^{\mu_t}} \|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K+1)} \right) + \frac{M_3}{b} \mu_t.$$

By Lemma 1, we have

$$\mathcal{G}(\mathbf{x}_{K+1}^{(t)}, \mathcal{C}) \leq \frac{1}{b} \left(\frac{\Delta^{\mu_t}}{2(K+1)} + \frac{2\sqrt{\beta\Delta^{\mu_t}} \|\boldsymbol{\lambda}^{\mu_t}\|}{\beta(K+1)} \right) + \frac{M_3}{b} \mu_t.$$

□

Remark 5. In Theorem 12, for any $t \geq \mathcal{O}(\ln K)$, we have $\left\| \mathbf{x}_K^{(t)} - \mathbf{x}^* \right\|, \mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}\left(\frac{1}{K^{1/(2\epsilon)}}\right)$, and $\left\| \hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^* \right\|, \mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}\left(\frac{1}{K^{1/\epsilon}}\right)$. Similarly, in Theorem 13, for any $t \geq \mathcal{O}(\ln K)$, we have $\left\| \mathbf{x}_K^{(t)} - \mathbf{x}^* \right\|, \mathcal{G}(\mathbf{x}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}\left(\frac{1}{\sqrt{K}}\right)$, and $\left\| \hat{\mathbf{x}}_K^{(t)} - \mathbf{x}^* \right\|, \mathcal{G}(\hat{\mathbf{x}}_K^{(t)}, \mathcal{C}) \leq \mathcal{O}\left(\frac{1}{K}\right)$.

B.5 DISCUSSION ON ALGORITHM 1

Advantages and disadvantages of Algorithm 1. In Algorithm 1, the update of \mathbf{x} (step 8) requires solving (W-EQ). Our method is especially suitable for problems where (W-EQ) is easy to solve analytically. This includes the class of affine variational inequalities, low-dimensional problems, and when optimization variables represent probabilities, for example.

For problems where (W-EQ) is hard to solve—for example, min-max optimization problems in GANs—one could use other *unconstrained* methods like GDA or EG methods (*without* projection) so as to solve $VI(\mathbb{R}^n, G)$, where G is defined in (3). Algorithm 4 describes ACVI when using an unconstrained solver for the inner problems. We observe that Algorithm 4 outperforms the projection-based baseline algorithms, see for example Fig. 4.

Note that when there are constraints, the projection-based methods such as GDA, EG, OGD etc. require solving a quadratic programming problem in each iteration (or twice per iteration for EG). This problem is often nontrivial and solving it may require using an interior point method or variants (such as the Frank-Wolfe algorithm). Because of this, Algorithm 1 can be seen as an orthogonal approach to projection-based methods, or in other words, as a complementary tool to solve (cVI) and particularly relevant when the constraints are non-trivial.

ACVI with only equality or inequality constraints. If there are no equality constraints, then $\mathcal{C}_=$ becomes \mathbb{R}^n . In this case, we have that $\mathbf{x}_{k+1}^{(t)}$ is the solution of:

$$\mathbf{x} + \frac{1}{\beta} (F(\mathbf{x}) + \boldsymbol{\lambda}_k^{(t)}) - \mathbf{y}_k^{(t)} = \mathbf{0}.$$

When there are no inequality constraints, we let $\mathbf{y}_k = \mathbf{x}_k$ and $\boldsymbol{\lambda}_k = \mathbf{0}$ for every k , and we can remove the outer loop. Thus, Algorithm 1 can be simplified to update only one variable \mathbf{x} each iteration with the following updating rule:

$$\mathbf{x}_{k+1} \text{ is the unique solution of } \mathbf{x} + \frac{1}{\beta} \mathbf{P}_c F(\mathbf{x}) - \mathbf{P}_c \mathbf{x}_k - \mathbf{d}_c = \mathbf{0}.$$

C VARIANT OF THE ACVI ALGORITHM (v-ACVI)

The presented approach of combining interior point methods with ADMM can be used as a framework to derive additional algorithms that may be more suitable for some specific problems. More precisely, observe from Eq. 1 that we could also consider a different splitting than that in § 4. Following this approach, we present a variant of Algorithm 1 and discuss its advantages and disadvantages relative to Algorithm 1.

C.1 INTRODUCTION OF THE VARIANT ACVI

Deriving the v-ACVI algorithm. By considering a different splitting in (1) we get:

$$\begin{cases} \min_{\mathbf{x}, \mathbf{y}} F(\mathbf{w})^\top \mathbf{x} - \mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{x})) + \mathbb{1}[\mathbf{C}\mathbf{y} = \mathbf{d}] \\ \text{s.t.} \quad \mathbf{x} = \mathbf{y} \end{cases}, \quad (50)$$

where: $\mathbb{1}[\mathbf{C}\mathbf{y} = \mathbf{d}] = \begin{cases} 0, & \text{if } \mathbf{C}\mathbf{y} = \mathbf{d} \\ +\infty, & \text{if } \mathbf{C}\mathbf{y} \neq \mathbf{d} \end{cases}$.

The augmented Lagrangian of (50) is thus:

$$\mathcal{L}_\beta(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) = F(\mathbf{w})^\top \mathbf{x} - \mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{x})) + \mathbb{1}[\mathbf{C}\mathbf{y} = \mathbf{d}] + \langle \boldsymbol{\lambda}, \mathbf{x} - \mathbf{y} \rangle + \frac{\beta}{2} \|\mathbf{x} - \mathbf{y}\|^2 \quad (\text{AL-CVI})$$

where $\beta > 0$ is the penalty parameter. We have that for \mathbf{x} at step $k + 1$:

$$\begin{aligned} \mathbf{x}_{k+1} &= \arg \min_{\mathbf{x}} \mathcal{L}_\beta(\mathbf{x}, \mathbf{y}_k, \boldsymbol{\lambda}_k) \\ &= \arg \min_{\mathbf{x}} \frac{1}{2} \left\| \mathbf{x} - \mathbf{y}_k + \frac{1}{\beta} (F(\mathbf{w}) + \boldsymbol{\lambda}_k) \right\|^2 - \frac{\mu}{\beta} \sum_{i=1}^m \log(-\varphi_i(\mathbf{x})) \end{aligned} \quad (51)$$

The following proposition ensures the existence and uniqueness of \mathbf{x}_{k+1} in $\mathcal{C}_<$. i.e. We show that \mathbf{x}_{k+1} is the unique solution in $\mathcal{C}_<$ of the following closed-form equation (see App. C.2 for its proof):

$$\mathbf{x} - \mathbf{y}_k + \frac{1}{\beta} (F(\mathbf{w}) + \boldsymbol{\lambda}_k) - \frac{\mu}{\beta} \sum_{i=1}^m \frac{\nabla \varphi_i(\mathbf{x})}{\varphi_i(\mathbf{x})} = \mathbf{0}. \quad (\text{X-CF})$$

Proposition 3 (unique solution). *The problem (X-CF) has a solution in $\mathcal{C}_<$ and the solution is unique.*

For \mathbf{y} , the updating rule is

$$\begin{aligned} \mathbf{y}_{k+1} &= \operatorname{argmin}_{\mathbf{y}} \mathcal{L}_\beta(\mathbf{x}_{k+1}, \mathbf{y}, \boldsymbol{\lambda}_k) \\ &= \operatorname{argmin}_{\mathbf{y} \in \mathcal{C}_=} -\frac{1}{\beta} (\boldsymbol{\lambda}_k)^\top \mathbf{y} + \frac{1}{2} \|\mathbf{y} - \mathbf{x}_{k+1}\|_2^2 \\ &= \operatorname{argmin}_{\mathbf{y} \in \mathcal{C}_=} \frac{1}{2} \left\| \mathbf{y} - \mathbf{x}_{k+1} - \frac{1}{\beta} \boldsymbol{\lambda}_k \right\|_2^2 \\ &= \mathbf{P}_c(\mathbf{x}_{k+1} + \frac{\boldsymbol{\lambda}_k}{\beta}) + \mathbf{d}_c \end{aligned} \quad (\mathbf{y})$$

And the updating rule for dual variable $\boldsymbol{\lambda}$ is

$$\boldsymbol{\lambda}_{k+1} = \boldsymbol{\lambda}_k + \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1}) \quad (\boldsymbol{\lambda})$$

As in § 4, we would like to choose \mathbf{w}_k so that $\mathbf{w}_k = \mathbf{x}_{k+1}$. To this end, we need the following assumption:

Assumption 2. $\forall \mathbf{b} \in \mathbb{R}^n$ and $\mu > 0$, $\mathbf{x} + \frac{1}{\beta} F(\mathbf{x}) - \frac{\mu}{\beta} \sum_{i=1}^m \frac{\nabla \varphi_i(\mathbf{x})}{\varphi_i(\mathbf{x})} + \mathbf{b} = \mathbf{0}$ has a solution in $\mathcal{C}_<$.

If Assumption 2 holds true, we can let \mathbf{w}_k be a solution of

$$\mathbf{w} - \mathbf{y}_{k+1} + \frac{1}{\beta}(F(\mathbf{w}) + \boldsymbol{\lambda}_{k+1}) - \frac{\mu}{\beta} \sum_{i=1}^m \frac{\nabla \varphi_i(\mathbf{w})}{\varphi_i(\mathbf{w})} = \mathbf{0} \quad (52)$$

in $\mathcal{C}_<$. And by Proposition 3, we can let \mathbf{x}_{k+1} be the unique solution of

$$\mathbf{x} - \mathbf{y}_k + \frac{1}{\beta}(F(\mathbf{w}_k) + \boldsymbol{\lambda}_k) - \frac{\mu}{\beta} \sum_{i=1}^m \frac{\nabla \varphi_i(\mathbf{x})}{\varphi_i(\mathbf{x})} = \mathbf{0} \quad (\mathbf{x})$$

in $\mathcal{C}_<$.

Note that \mathbf{w}_k is also a solution of (\mathbf{x}) . By the uniqueness of the solution of (\mathbf{x}) shown in Prop. 3, we know that $\mathbf{w}_k = \mathbf{x}_{k+1}$.

So in the $(k+1)$ -th step, we can compute \mathbf{x} , \mathbf{y} , $\boldsymbol{\lambda}$ and \mathbf{w} by (\mathbf{x}) , (\mathbf{y}) , $(\boldsymbol{\lambda})$ and (52), respectively. Since $\mathbf{w}_k = \mathbf{x}_{k+1}$, we can simplify our algorithm by removing variable \mathbf{w} and only keep \mathbf{x} , \mathbf{y} and $\boldsymbol{\lambda}$, see Algorithm 3.

Algorithm 3 v-ACVI pseudocode.

- 1: **Input:** operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$, equality $\mathbf{C}\mathbf{x} = \mathbf{d}$ and inequality constraints $\varphi_i(\mathbf{x}) \leq 0, i = [m]$, hyperparameters $\mu_{-1}, \beta > 0, \delta \in (0, 1)$, number of outer and inner loop iterations T and K , resp.
 - 2: **Initialize:** $\mathbf{y}_0^{(0)} \in \mathbb{R}^n, \boldsymbol{\lambda}_0^{(0)} \in \mathbb{R}^n$
 - 3: $\mathbf{P}_c \triangleq \mathbf{I} - \mathbf{C}^\top(\mathbf{C}\mathbf{C}^\top)^{-1}\mathbf{C}$ where $\mathbf{P}_c \in \mathbb{R}^{n \times n}$
 - 4: $\mathbf{d}_c \triangleq \mathbf{C}^\top(\mathbf{C}\mathbf{C}^\top)^{-1}\mathbf{d}$ where $\mathbf{d}_c \in \mathbb{R}^n$
 - 5: **for** $t = 0, \dots, T - 1$ **do**
 - 6: $\mu_t = \delta \mu_{t-1}$
 - 7: Denote $\hat{\varphi}(\boldsymbol{\lambda}, \mathbf{y})$ as the solution of $\frac{1}{\beta}(\mu_t \sum_{i=1}^m \frac{\nabla \varphi_i(\mathbf{x})}{\varphi_i(\mathbf{x})} - F(\mathbf{x}) - \boldsymbol{\lambda}) + \mathbf{y} - \mathbf{x} = 0$ with respect to \mathbf{x} , where $\mathbf{y}, \boldsymbol{\lambda}$ are variables
 - 8: **for** $k = 0, \dots, K - 1$ **do**
 - 9: $\mathbf{x}_{k+1}^{(t)} = \hat{\varphi}(\boldsymbol{\lambda}_k^{(t)}, \mathbf{y}_k^{(t)})$ Ensures $\mathbf{x}_{k+1} \in \mathcal{X}_\leq$
 - 10: $\mathbf{y}_{k+1}^{(t)} = \mathbf{P}_c(\mathbf{x}_{k+1}^{(t)} + \frac{\boldsymbol{\lambda}_k^{(t)}}{\beta}) + \mathbf{d}_c$
 - 11: $\boldsymbol{\lambda}_{k+1}^{(t)} = \boldsymbol{\lambda}_k^{(t)} + \beta(\mathbf{x}_{k+1}^{(t)} - \mathbf{y}_{k+1}^{(t)})$
 - 12: **end for**
 - 13: $(\mathbf{y}_0^{(t+1)}, \boldsymbol{\lambda}_0^{(t+1)}) \triangleq (\mathbf{y}_K^{(t)}, \boldsymbol{\lambda}_K^{(t)})$
 - 14: **end for**
-

Discussion: ACVI Vs. v-ACVI. Relative to Algorithm 1, the subproblem for solving \mathbf{x} in line 7 in Algorithm 3 becomes more complex, whereas the subproblem for \mathbf{y} becomes simpler. Hence, in cases when the inequality constraints are simpler, or there are no inequality constraints Alg. 3 may be more suitable, as that simplifies the \mathbf{x} subproblem. However, Algorithm 1 balances better the complexities of the subproblems, hence it may be simpler to use for general problems.

Convergence of v-ACVI. By analogous proofs to those in App. B, we can get similar convergence results as for Algorithm 1, that is Theorems 2 and 3. Specifically, Theorem 2 and 3 hold for Algorithm 3, provided that we replace the assumption “ F is monotone on $\mathcal{C}_=$ ” with “ F is monotone on \mathcal{C}_\leq ” in Theorems 2 and 3.

C.2 PROOF OF PROPOSITION 3

To prove proposition 3, we will use the following lemma.

Lemma 11. $\forall \mathbf{b} \in \mathbb{R}^n, \forall \mathbf{x} \in \mathcal{C}_<, \frac{1}{2}\|\mathbf{x} - \mathbf{b}\|_2^2 - \frac{\mu}{\beta} \sum_{i=1}^m \log(-\varphi_i(\mathbf{x})) \rightarrow +\infty, \|\mathbf{x}\|_2 \rightarrow +\infty$

Proof of Lemma 11. We denote $\phi : \mathcal{C}_< \rightarrow \mathbb{R}$ by

$$\phi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x} - \mathbf{b}\|_2^2 - \frac{\mu}{\beta} \sum_{i=1}^m \log(-\varphi_i(\mathbf{x}))$$

let $B(\mathbf{x}) = -\frac{\mu}{\beta} \sum_{i=1}^m \log(-\varphi_i(\mathbf{x}))$. We choose an arbitrary $\mathbf{x}_0 \in \mathcal{C}_<$. Then by the convexity of $B(\mathbf{x})$ we deduce that

$$\forall \mathbf{x} \in \mathcal{C}_<, \phi(\mathbf{x}) \geq \frac{1}{2} \|\mathbf{x} - \mathbf{b}\|_2^2 + B(\mathbf{x}_0) + \nabla B(\mathbf{x}_0)^\top (\mathbf{x} - \mathbf{x}_0) \rightarrow +\infty, \|\mathbf{x}\|_2 \rightarrow +\infty$$

□

In the remaining, we prove Proposition 3 which guarantees that (X-CF) has a unique solution.

Proof of Proposition 3: uniqueness of the solution of (X-CF). Let $\phi : \mathcal{C}_< \rightarrow \mathbb{R}$ denote:

$$\phi(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}_k + \frac{1}{\beta} (F(\mathbf{w}) + \lambda_k)\|_2^2 - \frac{\mu}{\beta} \sum_{i=1}^m \log(-\varphi_i(\mathbf{x}))$$

We choose $\mathbf{x}_0 \in \mathcal{C}_<$. By Lemma 11, $\forall \mathbf{x} \in \mathcal{C}_<, \phi(\mathbf{x}) \rightarrow +\infty, \|\mathbf{x}\|_2 \rightarrow +\infty$.

So there exists $M > 0$ such that $\mathbf{x}_0 \in B(0, M)$ and $\forall \mathbf{x} \in S, \phi(\mathbf{x}) \leq \phi(\mathbf{x}_0), \mathbf{x}$ must belong to $B(0, M)$, where

$$B(0, M) = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| \leq M\}$$

It's clear that there exists $t > 0$ such that for every $\mathbf{x} \in \mathcal{C}_<$ that satisfies $\phi(\mathbf{x}) \leq \phi(\mathbf{x}_0), \mathbf{x}$ must belong to \mathcal{C}_t , where

$$\mathcal{C}_t = \{\mathbf{x} \in B(0, M) \mid \varphi_i(\mathbf{x}) \leq -t\} \subset \mathcal{C}_< \quad (53)$$

And we can make t small enough so that $\mathbf{x}_0 \in \mathcal{C}_t$. \mathcal{C}_t is a nonempty compact set and ϕ is continuous, so there exists $\mathbf{x}^* \in \mathcal{C}_t$ such that $\phi(\mathbf{x}^*) \leq \phi(\mathbf{x}), \forall \mathbf{x} \in \mathcal{C}_<$.

Note that $\forall \mathbf{x} \in \mathcal{C}_< \setminus \mathcal{C}_t, \phi(\mathbf{x}) \geq \phi(\mathbf{x}_0) \geq \phi(\mathbf{x}^*)$. Therefore, \mathbf{x}^* is a global minimizer of ϕ . ϕ is strongly-convex thus $\mathbf{x}_{k+1} = \mathbf{x}^*$ is its unique minimizer. So $\mathbf{x} = \mathbf{x}_{k+1}$ if and only if $\nabla \phi(\mathbf{x}) = \mathbf{0}$. Therefore, \mathbf{x}_{k+1} is the unique solution of $\mathbf{x} - \mathbf{y}_k + \frac{1}{\beta} (F(\mathbf{w}) + \lambda_k) - \frac{\mu}{\beta} \sum_{i=1}^m \frac{\nabla \varphi_i(\mathbf{x})}{\varphi_i(\mathbf{x})} = \mathbf{0}$. □

D DETAILS ON THE IMPLEMENTATION

In this section, we provide the details on the implementation of the experiments shown in the main part in 2D and higher dimension bilinear game, see § D.1 and § D.2, respectively. We also provide here in § D.3 the details of the MNIST experiments presented later in App. E. In addition, we provide the source code through the following link: <https://github.com/Chavdarova/ACVI>.

D.1 EXPERIMENTS IN 2D

We first state the considered problem fully, then describe the setup what includes the hyperparameters.

Problems. We consider the following constrained bilinear game (for the same experiment shown in Fig. 1 and 5):

$$\min_{x_1 \in \mathbb{R}_+} \max_{x_2 \in \mathbb{R}_+} 0.05x_1^2 + x_1x_2 - 0.05x_2^2. \quad (\text{cBG})$$

The *Von Neumann’s ratio game* (Von Neumann, 1971) (results in Fig. (a)) is as follows:

$$\min_{\mathbf{x} \in \Delta^2} \max_{\mathbf{y} \in \Delta^2} \frac{\langle \mathbf{x}, \mathbf{R}\mathbf{y} \rangle}{\langle \mathbf{x}, \mathbf{S}\mathbf{y} \rangle}, \quad (\text{RG})$$

where $\Delta^2 = \{z \in \mathbb{R}^2 | z \geq 0, e^\top z = 1\}$, $\mathbf{R} = \begin{pmatrix} -0.6 & -0.3 \\ 0.6 & -0.3 \end{pmatrix}$ and $\mathbf{S} = \begin{pmatrix} 0.9 & 0.5 \\ 0.8 & 0.4 \end{pmatrix}$.

The so called *Forsaken* (Hsieh et al., 2021) game—used in Fig. 2(b) and in App. E—is as follows:

$$\min_{x_1 \in \mathcal{X}_1} \max_{x_2 \in \mathcal{X}_2} x_1(x_2 - 0.45) + h(x_1) - h(x_2), \quad (\text{Forsaken})$$

where $h(z) = \frac{1}{4}z^2 - \frac{1}{2}z^4 + \frac{1}{6}z^6$. The original version is unconstrained $\mathcal{X} \equiv \mathbb{R}^2$. In Fig. 2(b) we use the constraint $x_1^2 + x_2^2 \leq 4$, and in App. E we use two other constraints: $x_1 \geq 0.08$ and $x_2 \geq 0.4$.

For the toy GAN experiments, shown in Fig. 2(c), the problem is as follows:

$$\begin{aligned} \min_{\theta} \max_{\varphi} \mathbb{E}_{x \sim \mathcal{N}(0,2)} (\varphi x^2) - \mathbb{E}_{z \sim \mathcal{N}(0,1)} (\varphi \theta^2 z^2) \\ \text{s.t. } \varphi^2 + \theta^2 \leq 4 \end{aligned} \quad (\text{toy-GAN})$$

In the experiment, we look at a finite sum (sample average) approximation, which we then solve deterministically in a full batch fashion.

Setup. For all the 2D problems, we set the step size of GDA, EG and OGDA to 0.1, we use $k = 5$ and $\alpha = 0.5$ for LA-GDA, we set $\beta = 0.08$, $\mu_{-1} = 10^{-5}$, $\delta = 0.5$ and $\lambda_0 = \mathbf{0}$ for ACVI; and run for 50 iterations. For ACVI, we set the number of outer loop iterations to $T = 20$. In the first 19 outer loop iterations, we only run one inner loop iteration, and in the last outer loop iteration, we run 30 inner loop iterations (for a total of 50 updates). In Fig. 1 and Fig. (c), we set the starting point for all algorithms. In Fig. 2(a) and (b), we set the starting point to be $(0.5, 0.5)^\top$ for all algorithms.

D.2 HIGH-DIMENSION BILINEAR GAME

We set the step size of GDA, EG, and OGDA to 0.1, using $k = 4$ and $\alpha = 0.5$ for LA-GDA. For ACVI, we set $\beta = 0.5$, $\mu_{-1} = 10^{-6}$, $\delta = 0.5$ and $\lambda_0 = \mathbf{0}$ for ACVI.

The solution of (HBG) is $\mathbf{x}^* = \frac{1}{500}\mathbf{e}$, where $\mathbf{e} \in \mathbb{R}^{1000}$. As a metric of the experiments on this problem, we use the relative error: $\varepsilon_r(\mathbf{x}_k) = \frac{\|\mathbf{x}_k - \mathbf{x}^*\|}{\|\mathbf{x}^*\|}$. In Fig.3(b), we set $\varepsilon = 0.02$ and compute the number of iterations of ACVI needed to reach the relative error given different rotation “strength” $1 - \eta$, $\eta \in (0, 1)$. Here we set the maximum number of iterations to be 50 for all algorithms. In Fig. 3(a), we set $\eta = 0.05$ and compute CPU times needed for ACVI, EG, OGDA, and LA4-GDA to reach different relative errors. Here we set the maximum run time to 1500 seconds for all algorithms. In Fig. 7 in App. E on the other hand, we fix $\eta = 0.05$, and for varying *CPU time* limits, we compute the relative error of the last iterates of ACVI, GDA, EG, OGDA, and LA4-GDA.

More general HD-BG game (g-HBG). Since (HBG) has perfect conditioning (that is, the interactive term $\mathbf{x}_1^\top \mathbf{B}\mathbf{x}_2$, is with $\mathbf{B} \equiv \mathbf{I}$), in App. E.2 we present results on the following more general high

dimensional bilinear game:

$$\begin{aligned} \min_{\mathbf{x}_1 \in \Delta} \max_{\mathbf{x}_2 \in \Delta} \quad & \frac{\eta}{2} \cdot \mathbf{x}_1^\top \mathbf{A} \mathbf{x}_1 + (1 - \eta) \mathbf{x}_1^\top \mathbf{B} \mathbf{x}_2 - \frac{\eta}{2} \mathbf{x}_2^\top \mathbf{C} \mathbf{x}_2, \\ \Delta = & \{\mathbf{x}_i \in \mathbb{R}^{500} \mid \mathbf{x}_i \geq -\mathbf{e}, \text{ and } \mathbf{e}^\top \mathbf{x}_i = 0\}. \end{aligned} \quad (\text{g-HBG})$$

Where \mathbf{A} , \mathbf{B} and \mathbf{C} are randomly generated 500×500 matrices, and \mathbf{A} , \mathbf{C} are positive semi-definite.

The solution of (g-HBG) is $\mathbf{x}^* = \mathbf{0}$, where $\mathbf{0} \in \mathbb{R}^{1000}$. As a metric of the experiments on this problem, we use the error $\varepsilon(\mathbf{x}_k) = \|\mathbf{x}_k\|$. The remaining settings are identical to those of (HBG), explained above.

Comparison with the Frank-Wolfe algorithm on general HD-BG (g-HBG) and (gg-HBG). In App. E, we compare with the FW method (see A.4) on two problems: (i) (g-HBG), and (ii) (gg-HBG), where the objective is the same as (g-HBG) but the constraint set becomes more general, in which \mathbf{C}_i is a randomly generated 10×500 matrix, $i = \{1, 2\}$. In both experiments, we implement FW as in Algorithm 2, where we choose γ to be $2/(2+t)$ at the t -th iterate, $t = 0, \dots, T$. For (i), the constraint set of (g-HBG) is a “shifted simplex”, hence its vertices are easy to compute. This allows us to solve the linear minimization problem in Algorithm 2 of (Gidel et al., 2017a) much faster, and we refer to this implementation as *fast-FW*. In contrast, for the (gg-HBG) problem, we cannot apply this, and in that case, we use the standard linear programming routine *covopt.solvers.lp* in Python—referred herein as *FW*.

$$\begin{aligned} \min_{\mathbf{x}_1 \in \mathcal{U}_1} \max_{\mathbf{x}_2 \in \mathcal{U}_2} \quad & \frac{\eta}{2} \cdot \mathbf{x}_1^\top \mathbf{A} \mathbf{x}_1 + (1 - \eta) \mathbf{x}_1^\top \mathbf{B} \mathbf{x}_2 - \frac{\eta}{2} \mathbf{x}_2^\top \mathbf{C} \mathbf{x}_2, \\ \mathcal{U}_i = & \{\mathbf{x}_i \in \mathbb{R}^{500} \mid -100\mathbf{e} \leq \mathbf{x}_i \leq 100\mathbf{e}, \text{ and } \mathbf{C}_i \mathbf{x}_i = \mathbf{0}\}, i = 1, 2. \end{aligned} \quad (\text{gg-HBG})$$

The solution of both (g-HBG) and (gg-HBG) is $\mathbf{0}$. As a metric of the experiments on this problem, we use the error $\varepsilon(\mathbf{x}_k) = \|\mathbf{x}_k\|$. The remaining settings are identical to those of (HBG), explained above.

D.3 MNIST AND FASHION-MNIST EXPERIMENTS

For the experiments on the MNIST dataset¹, we use the source code of Chavdarova et al. (2021b) for the baselines and we build on it to implement ACVI. For completeness, we provide an overview of the implementation.

Models. We used the DCGAN architectures (Radford et al., 2016), listed in Table 1, and the parameters of the models are initialized using PyTorch default initialization. For experiments on this dataset, we used the *non-saturating* GAN loss as proposed in (Goodfellow et al., 2014):

$$\mathcal{L}_D = \mathbb{E}_{\tilde{\mathbf{x}}_d \sim p_d} \log(D(\tilde{\mathbf{x}}_d)) + \mathbb{E}_{\tilde{\mathbf{z}} \sim p_z} \log(1 - D(G(\tilde{\mathbf{z}}))) \quad (\text{L-D})$$

$$\mathcal{L}_G = \mathbb{E}_{\tilde{\mathbf{z}} \sim p_z} \log(D(G(\tilde{\mathbf{z}}))), \quad (\text{L-G})$$

where $G(\cdot)$, $D(\cdot)$ denote the generator and discriminator, resp., and p_d and p_z denote the data and the latent distributions (the latter predefined as normal distribution).

Details on the ACVI implementation. When implementing ACVI on MNIST, we “remove” the outer loop of Algorithm 1 (that is we set $T = 1$), and fix μ to be a small number, in particular, we select $\mu = 10^{-9}$. We randomly initialize \mathbf{x} and \mathbf{y} and initialize λ to zero. For lines 8 and 9 of Algorithm 1, we run l steps of (unconstrained) EG and GD, respectively. For the update of \mathbf{x} (using EG), we use step-size $\eta_x = 0.001$, whereas for \mathbf{y} (using GD), we use step-size $\eta_y = 0.2$. We present results when $l \in \{1, 10\}$. At every iteration, we update λ using the expression in line 11 of Algorithm 1, with $\beta = 0.5$.

Because the problem in step 8 of Algorithm 1 does not change a lot over the iterations (as well as when computing \mathbf{y}), when we implement Algorithm 1 we do *not* reinitialize the variable \mathbf{x} . We

¹Provided under *Creative Commons Attribution-Share Alike 3.0*.

Generator	Discriminator
Input: $\mathbf{z} \in \mathbb{R}^{128} \sim \mathcal{N}(0, I)$	Input: $\mathbf{x} \in \mathbb{R}^{1 \times 28 \times 28}$
transposed conv. (ker: 3×3 , $128 \rightarrow 512$; stride: 1)	conv. (ker: 4×4 , $1 \rightarrow 64$; stride: 2; pad: 1)
Batch Normalization	LeakyReLU (negative slope: 0.2)
ReLU	conv. (ker: 4×4 , $64 \rightarrow 128$; stride: 2; pad: 1)
transposed conv. (ker: 4×4 , $512 \rightarrow 256$, stride: 2)	Batch Normalization
Batch Normalization	LeakyReLU (negative slope: 0.2)
ReLU	conv. (ker: 4×4 , $128 \rightarrow 256$; stride: 2; pad: 1)
transposed conv. (ker: 4×4 , $256 \rightarrow 128$, stride: 2)	Batch Normalization
Batch Normalization	LeakyReLU (negative slope: 0.2)
ReLU	conv. (ker: 3×3 , $256 \rightarrow 1$; stride: 1)
transposed conv. (ker: 4×4 , $128 \rightarrow 1$, stride: 2, pad: 1)	$Sigmoid(\cdot)$
$Tanh(\cdot)$	

Table 1: DCGAN architectures (Radford et al., 2016) used for experiments on MNIST. With “conv.” we denote a convolutional layer and “transposed conv” a transposed convolution layer (Radford et al., 2016). We use *ker* and *pad* to denote *kernel* and *padding* for the (transposed) convolution layers, respectively. With $h \times w$ we denote the kernel size. With $c_{in} \rightarrow c_{out}$ we denote the number of channels of the input and output, for (transposed) convolution layers. The models use Batch Normalization (Ioffe & Szegedy, 2015) layers.

instead use the one from the previous iteration as initialization and update it l times. The full details of the training when using an inner optimizer for step 8 of Algorithm 1 are provided in Algorithm 4, where we recall that $G(\mathbf{x})$ is defined as:

$$G(\mathbf{x}) \triangleq \mathbf{x} + \frac{1}{\beta} \mathbf{P}_c F(\mathbf{x}) - \mathbf{P}_c \mathbf{y}_k + \frac{1}{\beta} \mathbf{P}_c \boldsymbol{\lambda}_k - \mathbf{d}_c \quad (\text{G-EQ})$$

Note that in the case of MNIST, we consider only inequality constraints (and there are no equality constraints), therefore, the matrices \mathbf{P}_c and \mathbf{d}_c are identity and zero, respectively. Thus, there is no need for lines 3 and 4 in Algorithm 4.

Algorithm 4 Pseudocode for ACVI when using an inner optimizer (MNIST experiments).

- 1: **Input:** operator $F : \mathcal{X} \rightarrow \mathbb{R}^n$, equality $\mathbf{C}\mathbf{x} = \mathbf{d}$ and inequality constraints $\varphi_i(\mathbf{x}) \leq 0, i = [m]$, hyperparameters $\mu, \beta > 0, \delta \in (0, 1)$, inner optimizer \mathcal{A} (e.g. EG, GDA, OGDA), l number of steps for the inner-optimizer, number of iterations K .
 - 2: **Initialize:** $\mathbf{x}_0 \in \mathbb{R}^n, \mathbf{y}_0 \in \mathbb{R}^n, \boldsymbol{\lambda}_0 \in \mathbb{R}^n$
 - 3: $\mathbf{P}_c \triangleq \mathbf{I} - \mathbf{C}^\top (\mathbf{C}\mathbf{C}^\top)^{-1} \mathbf{C}$ where $\mathbf{P}_c \in \mathbb{R}^{n \times n}$
 - 4: $\mathbf{d}_c \triangleq \mathbf{C}^\top (\mathbf{C}\mathbf{C}^\top)^{-1} \mathbf{d}$ where $\mathbf{d}_c \in \mathbb{R}^n$
 - 5: **for** $k = 0, \dots, K - 1$ **do**
 - 6: To obtain \mathbf{x}_{k+1} : run l steps of \mathcal{A} solving $\text{VI}(\mathbb{R}^n, G)$, where G is defined in (G-EQ)
 - 7: To obtain \mathbf{y}_{k+1} : run l steps of GD to find \mathbf{y}_{k+1} that minimizes:

$$-\mu \sum_{i=1}^m \log(-\varphi_i(\mathbf{y})) + \frac{\beta}{2} \left\| \mathbf{y} - \mathbf{x}_{k+1} - \frac{1}{\beta} \boldsymbol{\lambda}_k \right\|^2$$
 - 8: $\boldsymbol{\lambda}_{k+1} = \boldsymbol{\lambda}_k + \beta(\mathbf{x}_{k+1} - \mathbf{y}_{k+1})$
 - 9: **end for**
 - 10: **Return:** \mathbf{x}_K
-

The implementation details for the Fashion-MNIST experiment are identical to those of the MNIST experiment.

Setup 1: MNIST. The MNIST experiment in Fig. 4 in the main part (and Fig. 11, 12) has 100 randomly generated linear inequality constraints for the Generator and 100 for the Discriminator.

Setup 1: projection details. Suppose the linear inequality constraints for the Generator are $\mathbf{A}\boldsymbol{\theta} \leq \mathbf{b}$, where $\boldsymbol{\theta} \in \mathbb{R}^n$ is the vector of all parameters of the Generator, $\mathbf{A} = (\mathbf{a}_1^\top, \dots, \mathbf{a}_{100}^\top)^\top \in \mathbb{R}^{100 \times n}$, $\mathbf{b} = (b_1, \dots, b_{100}) \in \mathbb{R}^{100}$. We use the *greedy projection algorithm* described in (Beck, 2017). A greedy projection algorithm is essentially a projected gradient method, it is easy to implement in

high-dimension problems, and it has a convergence rate of $O(1/\sqrt{K})$. See Chapter 8.2.3 in (Beck, 2017) for more details. Since the dimension n is very large, at each step of the projection, one could only project θ to one hyperplane $\mathbf{a}_i^\top \mathbf{x} = b_i$ for some $i \in \mathcal{I}(\theta)$, where

$$\mathcal{I}(\theta) \triangleq \{j | \mathbf{a}_j^\top \theta > b_j\}.$$

For every $j \in \{1, 2, \dots, 100\}$, let

$$\mathcal{S}_j \triangleq \{\mathbf{x} | \mathbf{a}_j^\top \mathbf{x} \leq b_j\}.$$

The greedy projection method chooses i so that $i \in \arg \max \{dist(\theta, \mathcal{S}_i)\}$. Note that as long as θ is not in the constraint set $C_{\leq} = \{\mathbf{x} | \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$, i would be in $\mathcal{I}(\theta)$. Algorithm 5 gives the details of the greedy projection method we use for the baseline, written for the Generator only for simplicity; the same projection method is used for the Discriminator as well.

Algorithm 5 Greedy projection method for the baseline.

```

1: Input:  $\theta \in \mathbb{R}^n$ ,  $\mathbf{A} = (\mathbf{a}_1^\top, \dots, \mathbf{a}_{100}^\top)^\top \in \mathbb{R}^{100 \times n}$ ,  $\mathbf{b} = (b_1, \dots, b_{100}) \in \mathbb{R}^{100}$ ,  $\varepsilon > 0$ 
2: while True do
3:    $\mathcal{I}(\theta) \triangleq \{j | \mathbf{a}_j^\top \theta > b_j\}$ 
4:   if  $\mathcal{I}(\theta) = \emptyset$  or  $\max_{j \in \mathcal{I}(\theta)} \frac{|\mathbf{a}_j^\top \theta - b_j|}{\|\mathbf{a}_j\|} < \varepsilon$  then
5:     break
6:   end if
7:   choose  $i \in \arg \max_{j \in \mathcal{I}(\theta)} \frac{|\mathbf{a}_j^\top \theta - b_j|}{\|\mathbf{a}_j\|}$ 
8:    $\theta \leftarrow \theta - \frac{|\mathbf{a}_i^\top \theta - b_i|}{\|\mathbf{a}_i\|^2} \mathbf{a}_i$ 
9: end while
10: Return:  $\theta$ 

```

Setup 2: MNIST & Fashion-MNIST. We add two constraints for the MNIST experiment: we set the squared sum of all parameters of the Generator and that of the Discriminator (separately) to be less than or equal to a hyperparameter M . We select a large number for M ; in particular, we set $M = 50$.

Metrics. We describe the metrics for the MNIST experiments shown later in App. E. We use the two standard GAN metrics, Inception Score (IS, Salimans et al., 2016) and Fréchet Inception Distance (FID, Heusel et al., 2017). Both FID and IS rely on a pre-trained classifier and take a finite set of \tilde{m} samples from the generator to compute these. Since MNIST has greyscale images, we used a classifier trained on this dataset and used $\tilde{m} = 5000$.

Metrics: IS. Given a sample from the generator $\tilde{\mathbf{x}}_g \sim p_g$ —where p_g denotes the data distribution of the generator—IS uses the softmax output of the pre-trained network $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}_g)$ which represents the probability that $\tilde{\mathbf{x}}_g$ is of class $c_i, i \in 1 \dots C$, i.e., $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}_g) \in [0, 1]^C$. It then computes the marginal class distribution $p(\tilde{\mathbf{y}}) = \int_{\tilde{\mathbf{x}}} p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}_g) p_g(\tilde{\mathbf{x}}_g)$. IS measures the Kullback–Leibler divergence \mathbb{D}_{KL} between the predicted conditional label distribution $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}_g)$ and the marginal class distribution $p(\tilde{\mathbf{y}})$. More precisely, it is computed as follows:

$$IS(G) = \exp \left(\mathbb{E}_{\tilde{\mathbf{x}}_g \sim p_g} \left[\mathbb{D}_{KL}(p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}_g) || p(\tilde{\mathbf{y}})) \right] \right) = \exp \left(\frac{1}{\tilde{m}} \sum_{i=1}^{\tilde{m}} \sum_{c=1}^C p(y_c|\tilde{\mathbf{x}}_i) \log \frac{p(y_c|\tilde{\mathbf{x}}_i)}{p(y_c)} \right). \quad (\text{IS})$$

It aims at estimating (i) if the samples look realistic i.e., $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}_g)$ should have low entropy, and (ii) if the samples are diverse (from different ImageNet classes), i.e., $p(\tilde{\mathbf{y}})$ should have high entropy. As these are combined using the Kullback–Leibler divergence, the higher the score is, the better the performance.

Metrics: FID. Contrary to IS, FID compares the synthetic samples $\tilde{\mathbf{x}}_g \sim p_g$ with those of the training dataset $\tilde{\mathbf{x}}_d \sim p_d$ in a feature space. The samples are embedded using the first several layers of a pre-trained classifier. It assumes p_g and p_d are multivariate normal distributions and estimates the means \mathbf{m}_g and \mathbf{m}_d and covariances \mathbf{C}_g and \mathbf{C}_d , respectively, for p_g and p_d in that feature space.

Finally, FID is computed as:

$$\mathbb{D}_{\text{FID}}(p_d, p_g) \approx \mathcal{D}_2((\mathbf{m}_d, \mathbf{C}_d), (\mathbf{m}_g, \mathbf{C}_g)) = \|\mathbf{m}_d - \mathbf{m}_g\|_2^2 + \text{Tr}(\mathbf{C}_d + \mathbf{C}_g - 2(\mathbf{C}_d \mathbf{C}_g)^{\frac{1}{2}}),$$

(FID)

where \mathcal{D}_2 denotes the Fréchet Distance. Note that as this metric is a distance, the lower it is, the better the performance.

Hardware. We used the Colab platform (<https://colab.research.google.com/>) and *Tesla P100* GPUs. The running times are reported in App. E.

E ADDITIONAL EMPIRICAL ANALYSIS

In this section, we provide some omitted plots/analyses of the results in the main paper as well as additional experiments. In particular, (i) App. E.1 lists results in 2D, (ii) App. E.2 on (HBG) and (g-HBG), whereas (iii) App. E.3 provides more detailed plots of those experiments summarized in Fig. 4 and presents additional results on other constraints on MNIST where we compare computationally-wise with *unconstrained* baselines.

E.1 ADDITIONAL EXPERIMENTS IN 2D, ON HBG AND ON G-HBG

Depicting the omitted baselines of Fig. 1. While Fig. 1 lists solely EG and ACVI for clarity, Fig. 5 depicts all the considered baselines in this paper on the cBG problem for completeness.




Figure 5: In addition to Fig. 1, here we depict all the considered baselines on the cBG problem. The solution at $(0, 0)$ is denoted with a red star (\star), and the vector field of this problem with gray arrows. See App. D.1 for details on the implementation.

Additional experiments: varying constraints on the Forsaken problem. The *Forsaken* game was first pointed out in (Hsieh et al., 2021) and is particularly relevant because it has *limit cycles*, despite that it is in 2D. Since we are missing a tool to detect if we are in a limit cycle when in higher dimensions, this example is a popular benchmark in many recent works. Interestingly, in Fig. 2(b) we observe that ACVI is the only method that escapes the limit cycle. However, since in those simulations, given the initial point the constraints are not active throughout the training, in this section, we run experiments with additional constraints. Fig. 6 depicts the baseline methods and ACVI on the *Forsaken* problem with two different constraints than that considered in Fig. 2 (that $x_1^2 + x_2^2 \leq 4$). Since this game is non-monotone, we observe that for some constraints the baseline methods—GDA, EG, OGDA, LA4-GDA—stay near the constraint (and do not converge). This may indicate that ACVI may have better chances of converging for *broader* problem classes than monotone VIs, relative to baseline methods whose convergence may depend on the constraints, and when hitting a constraint may be significantly slower (as Fig. 1 illustrates).

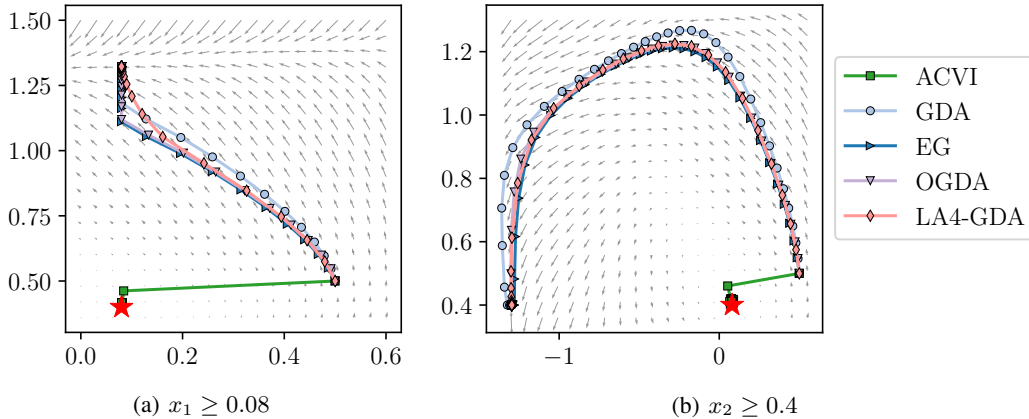


Figure 6: *Forsaken game with different constraints*: we consider two additional (to that in Fig. 2) constraints: (a) that $x_1 \geq 0.08$, and (b) that $x_2 \geq 0.4$. See App. D.1 for details on the implementation, and App. E.1 for a discussion.

E.2 ADDITIONAL EXPERIMENTS HBG AND ON G-HBG

Complementary analysis to those in Fig. 3. Similar to Fig. 3, in Fig. 7 we run experiments on the HBG problem. However, here for a given fixed CPU time, we depict the relative error of the considered baselines and ACVI.




Figure 7: Given varying CPU time (in seconds), depicting the relative error (see App. D.2) of GDA, EG, OGDA, LA4-GDA, and ACVI (Algorithm 1) on the HBG problem where η is fixed to $\eta = .05$ (hence the vector field is highly rotational). This experiment complements those in Fig. 3 in the main paper. For the details on the implementation, see App. D.2.

Additional experiments on (g-HBG). In Fig. 8 we run experiments on the generalized HBG problem (g-HBG). In figure 8(a), we compute the number of iterations needed to reach ϵ -distance to solution for varying intensity of the rotational component $(1 - \eta)$; in figure 8(b), we compute the error of the last iterate given fixed CPU time. We observe that despite the highly rotational monotone vector field, ACVI converges significantly faster in terms of wall clock time in higher dimensions as well.

Comparison with Frank-Wolf algorithm on (g-HBG) and (gg-HBG). Similar to Fig. 8, in Fig. 9 we also run experiments on (g-HBG), but here we compare ACVI with FW. We observe that ACVI outperforms FW even when we make use of the special structure of the simple constraint set when solving the linear minimization problem in FW (the *fast FW* method). Similar to Fig. 9(b), in Fig. 10

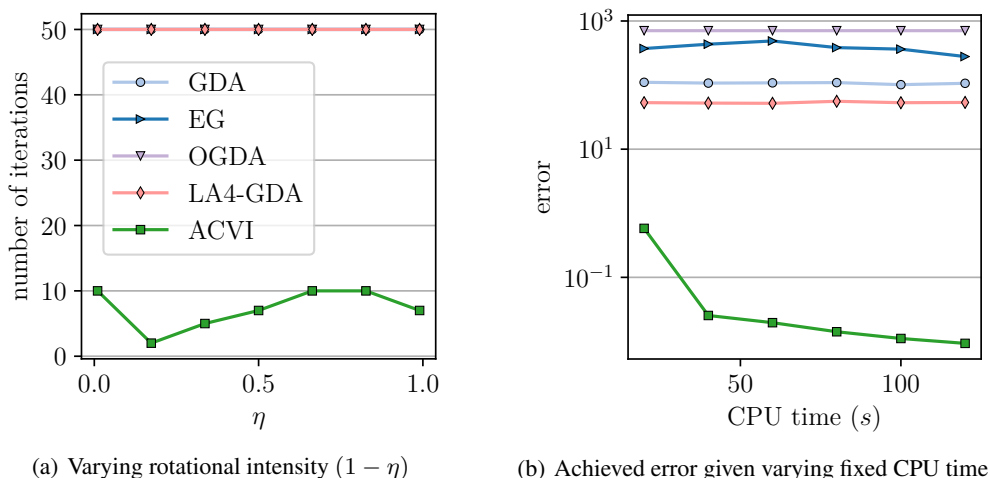


Figure 8: *General high-dimensional bilinear game (g-HBG)*: comparison of ACVI with the GDA, EG, OGDA, and LA4-GDA baselines (described in App. A.4). **Left:** number of iterations (y -axis) needed to reach an ϵ -distance to the solution, for varying intensity of the rotational component $1 - \eta$ (η is the x -axis) of the vector field (the smaller the η the higher the rotational component). We fix a threshold of the maximum number of iterations, and we stop the experiment. **Right:** distance to the solution (see App. D.2) of the last iterate (y -axis) for a varying wall-clock CPU time allowed to run each experiment (x -axis); in this experiment η is fixed to $\eta = 0.05$. See App. E.2 and D.2 for discussion and details on the implementation, respectively.

we run experiments on the (gg-HBG) problem, where we fix CPU time and depict the relative error of ACVI and FW.

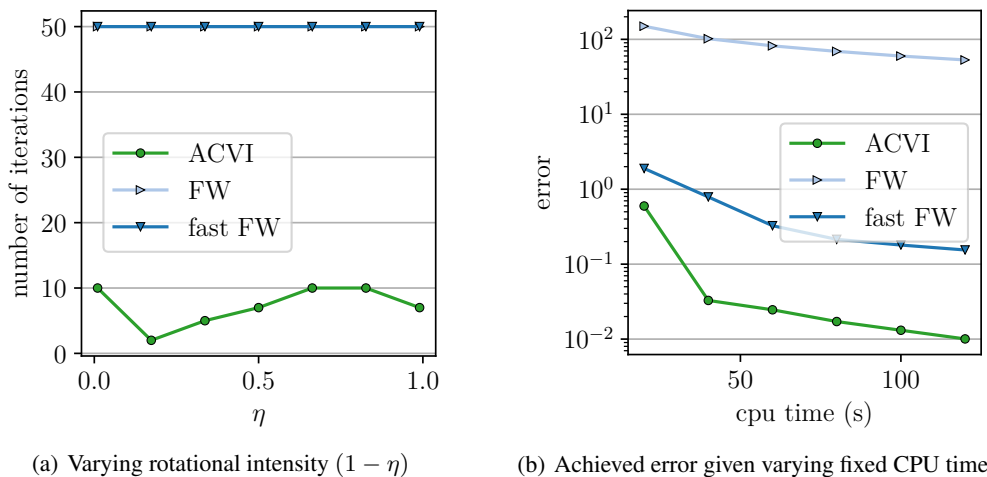


Figure 9: *General high-dimensional bilinear game (g-HBG)*: comparison of ACVI with FW baseline (Algorithm 2). **Left:** number of iterations (y -axis) needed to reach an ϵ -distance to the solution, for varying intensity of the rotational component $1 - \eta$ (η is the x -axis) of the vector field (the smaller the η the higher the rotational component). We fix a threshold of the maximum number of iterations, and we stop the experiment. **Right:** distance to the solution (see App. D.2) of the last iterate (y -axis) for a varying wall-clock CPU time allowed to run each experiment (x -axis); in this experiment η is fixed to $\eta = 0.05$. See App. E.2 and D.2 for discussion and details on the implementation, respectively.

Since FW (and variants, such as approximate and accelerated) rely on a specific structure of the constraints, FW can be extremely slow when those assumptions are not met—see discussion by Jaggi

(2013) in §3 as well as examples in §4 therein. In contrast, the herein-presented ACVI Algorithm focuses on constraints of a general form, and further variants can be derived out of it to also exploit the structure of the constraints. We leave exploiting such constraint structure—including extending FW to VIs and deriving variants of ACVI—for future work.




Figure 10: Given varying CPU time (in seconds), depicting the relative error (see App. D.2) of FW and ACVI (Algorithm 1) on the HBG problem where η is fixed to $\eta = .05$ (hence the vector field is highly rotational). For the details on the implementation, see App. D.2.

E.3 EXPERIMENTS ON MNIST AND FASHION-MNIST

E.3.1 SETUP 1: EXPERIMENTS ON MNIST WITH LINEAR INEQUALITIES

In this section, we present more detailed results of the summarizing plot in Fig. 4 of the main paper. For this experiment, we used linear inequalities as described in § D.3. Unlike in subsection E.3.2, here all the baselines are projected methods (that is, the same problem setting applies to ACVI and the baselines).

Fig. 11 and 12 depict the comparisons with projected GDA and projected EG, respectively. We observe that ACVI converges fast relative to the corresponding baseline. When choosing a larger number of steps for the inner problem $l = 10$ (see Algorithm 4) the wall-clock time per iteration increases, and interestingly the ACVI steps compensate for that and overall converge as fast as when $l = 1$.

E.3.2 SETUP 2: EXPERIMENTS ON (FASHION-)MNIST WITH QUADRATIC INEQUALITIES

In this section, we consider the MNIST and Fashion-MNIST datasets, which are unconstrained problems so as to make use of the well-established performance metrics (which are otherwise unclear in the non-monotone settings, where we do not know the optimal solution a priori). We augment the problem with a mild constraint which requires that the norm of the per-player parameters does not exceed a certain value (see App. D.3). We compare ACVI with *unconstrained* baselines, which sets ACVI at a disadvantage as the projection requires additional computation. However, the primary purpose of these experiments is to observe if Algorithm 1 is competitive computationally-wise when lines 8 and 9 are non-trivial and require an (unconstrained) solver. However note that since MNIST is a relatively easy problem, it may not answer the natural question if ACVI has advantages on problems augmented with constraints over standard unconstrained methods. We leave such analyses for future work. The implementation and the used metrics are described in App. D.3.

Fig. 13 summarizes the experiments in terms of the obtained FID score over time. We observe that ACVI (although it uses two solvers at each iteration) is yet performing *competitively* to *unconstrained* GDA and EG. Figures 14–19 provide in addition samples of the Generator and IS scores, separately for each method. Figures 20 and 21 depict samples generated by the different methods,

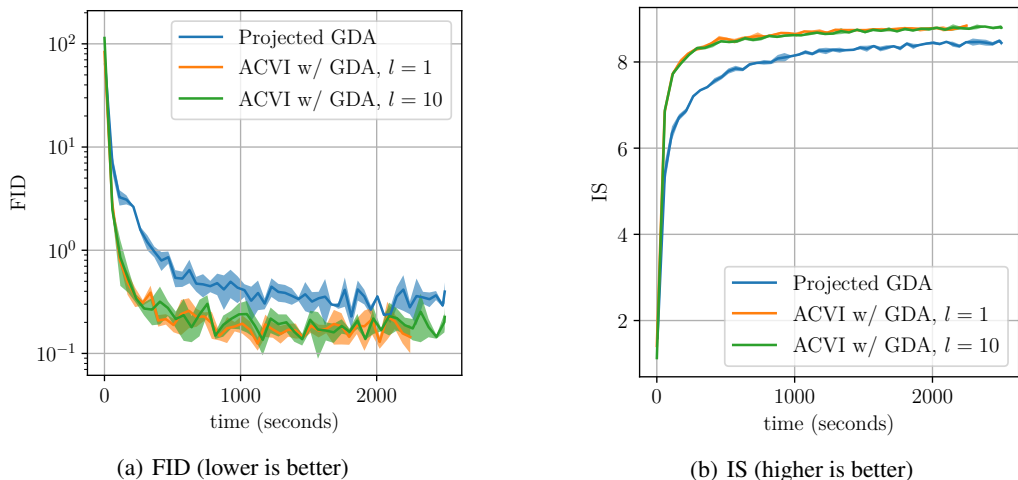


Figure 11: *Setup 1: Comparison between ACVI and GDA, and the **projected GDA** on MNIST with linear inequalities (described in § D.3). l denotes the number of steps for the inner problems, see Algorithm 4. The depicted results are over multiple seeds. The FID and IS metrics as well as the implementation details are described in App. D.3. See App. E.3.1 for discussion.*

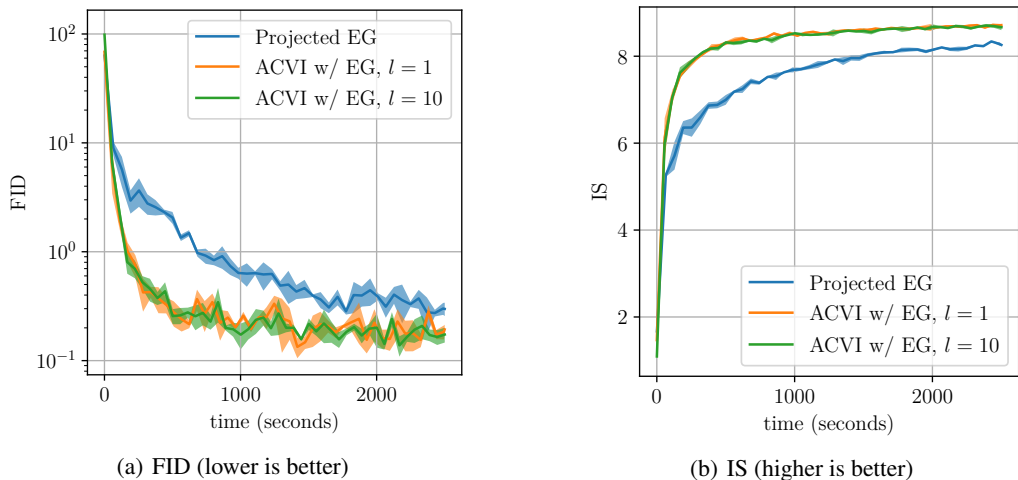


Figure 12: *Setup 1: Comparison between ACVI and EG, and the **projected EG** on MNIST with linear inequalities (described in § D.3). l denotes the number of steps for the inner problems, see Algorithm 4. The depicted results are over multiple seeds. The FID and IS metrics as well as the implementation details are described in App. D.3. See App. E.3.1 for discussion.*

when trained on the Fashion-MNIST dataset. We believe that further exploring the type of constraints to be added, or the implementation options (e.g., l , step-size) may be proven fruitful even for problems that are originally unconstrained—as such an approach may reduce the rotational component of the original vector field, what in turn causes faster convergence or may help in escaping limit cycles for problems beyond monotone ones.




Figure 13: Summary of the experiments on MNIST, using FID (lower is better). 13(a): GDA and ACVI with GDA, and 13(b): EG and ACVI with EG, using $l = \{1, 10\}$ for ACVI. Using step size of 0.001. The depicted results are over multiple seeds. See App. D.3 and E.3 for details on the implementation and discussion, resp.

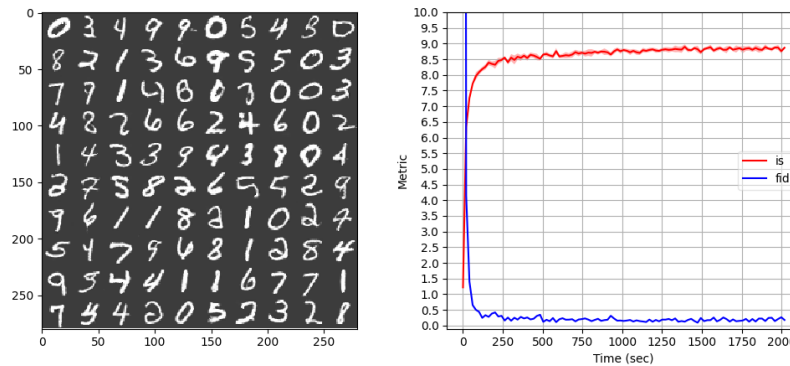


Figure 14: GDA on MNIST. Left: samples $\tilde{x}_g \sim p_g$ of the last iterate of the Generator. Right: FID and IS of GDA, depicted in blue and red, respectively. Using step size of 0.001.




Figure 15: ACVI (Algorithm 1) with 10 GDA steps on MNIST. Left: samples $\tilde{x}_g \sim p_g$ of the last iterate of the Generator. Right: FID and IS of GDA, depicted in blue and red, respectively. Using step size of 0.001 for x and 0.2 for y , and $l = 10$ both for x and y .

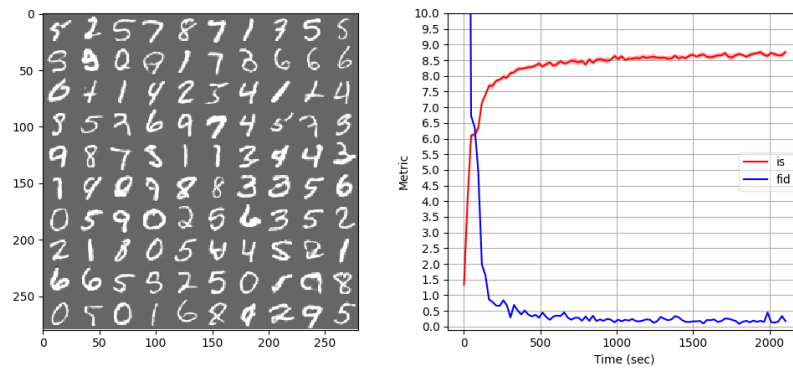


Figure 16: ACVI (Algorithm 1) with 1 GDA step on MNIST. Left: samples $\tilde{x}_g \sim p_g$ of the last iterate of the Generator. Right: FID and IS of GDA, depicted in blue and red, respectively. Using step size of 0.001 for x and 0.2 for y , and $l = 1$ both for x and y .

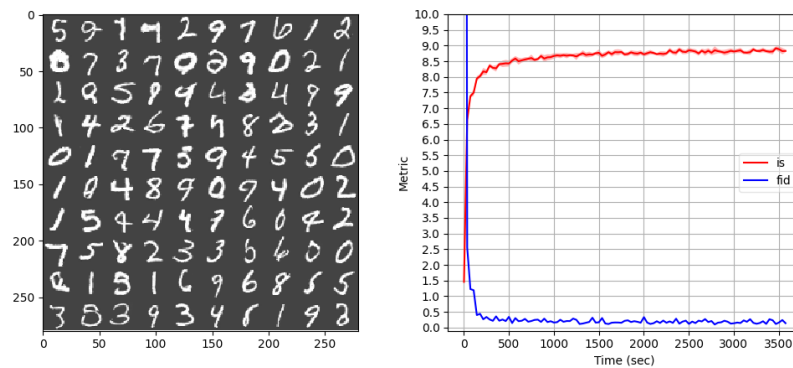


Figure 17: EG on MNIST. Left: samples $\tilde{x}_g \sim p_g$ of the last iterate of the Generator. Right: FID and IS of EG, depicted in blue and red, respectively. Using step size of 0.001.




Figure 18: ACVI (Algorithm 1) with 10 EG steps on MNIST. Left: samples $\tilde{x}_g \sim p_g$ of the last iterate of the Generator. Right: FID and IS of **GDA**, depicted in blue and red, respectively. Using step size of 0.001 for x and 0.2 for y , and $l = 10$ both for x and y .

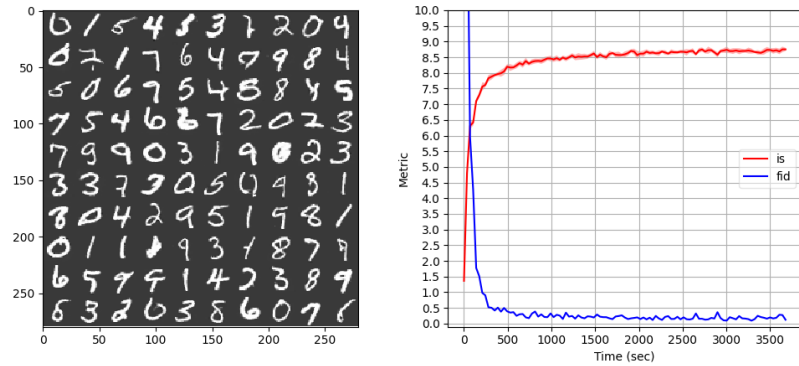


Figure 19: ACVI (Algorithm 1) with 1 EG step on MNIST. Left: samples $\tilde{x}_g \sim p_g$ of the last iterate of the Generator. Right: FID and IS of **GDA**, depicted in blue and red, respectively. Using step size of 0.001 for x and 0.2 for y , and $l = 1$ both for x and y .




Figure 20: Generated images at fixed wall-clock computation time (3000s) by: the baseline **EG**, and by ACVI with $l \in \{1, 10\}$ on the Fashion-MNIST (Xiao et al., 2017) dataset. See App. D.3 and E.3 for details on the implementation and discussion, resp.




Figure 21: Generated images at fixed wall-clock computation time (3000s) by: the baseline GDA, and by ACVI with $l \in \{1, 10\}$ on the Fashion-MNIST (Xiao et al., 2017) dataset. See App. D.3 and E.3 for details on the implementation and discussion, resp.