

---

# Coordinate Linear Variance Reduction for Generalized Linear Programming

---

**Chaobing Song\***

University of Wisconsin-Madison  
chaobing.song@wisc.edu

**Cheuk Yin Lin\***

University of Wisconsin-Madison  
cylin@cs.wisc.edu

**Stephen J. Wright**

University of Wisconsin-Madison  
swright@cs.wisc.edu

**Jelena Diakonikolas**

University of Wisconsin-Madison  
jelena@cs.wisc.edu

## Abstract

We study a class of generalized linear programs (GLP) in a large-scale setting, which includes a simple, possibly nonsmooth convex regularizer and simple convex set constraints. By reformulating (GLP) as an equivalent convex-concave min-max problem, we show that the linear structure in the problem can be used to design an efficient, scalable first-order algorithm, to which we give the name *Coordinate Linear Variance Reduction* (CLVR; pronounced “clever”). CLVR yields improved complexity results for (GLP) that depend on the max row norm of the linear constraint matrix in (GLP) rather than the spectral norm. When the regularization terms and constraints are separable, CLVR admits an efficient lazy update strategy that makes its complexity bounds scale with the number of nonzero elements of the linear constraint matrix in (GLP) rather than the matrix dimensions. Further, for the special case of linear programs and by exploiting sharpness, we propose a restart scheme for CLVR to obtain empirical linear convergence. Finally, we show that Distributionally Robust Optimization (DRO) problems with ambiguity sets based on both  $f$ -divergence and Wasserstein metrics can be reformulated as (GLPs) by introducing sparsely connected auxiliary variables. We complement our theoretical guarantees with numerical experiments that verify our algorithm’s practical effectiveness in terms of wall-clock time and number of data passes.

## 1 Introduction

We study the following generalized linear program (GLP):

$$\min_{\mathbf{x}} \{ \mathbf{c}^T \mathbf{x} + r(\mathbf{x}) : \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \in \mathcal{X} \}, \quad (\text{GLP})$$

where  $\mathbf{x}, \mathbf{c} \in \mathbb{R}^d$ ,  $\mathbf{A} \in \mathbb{R}^{n \times d}$ ,  $\mathbf{b} \in \mathbb{R}^n$ ,  $r : \mathbb{R}^d \rightarrow \mathbb{R}$  is a convex regularizer, and  $\mathcal{X} \subseteq \mathbb{R}^d$  is a closed convex set, such that a proximal/projection operator involving  $r$  and  $\mathcal{X}$  can be computed efficiently. When  $\mathcal{X}$  is the nonnegative orthant  $\{\mathbf{x} : x_i \geq 0, i \in [d]\}$  and  $r \equiv 0$ , (GLP) reduces to the standard form of a linear program (LP). When  $\mathcal{X}$  is a convex cone and  $r \equiv 0$ , (GLP) reduces to a conic linear program. (GLP) is an important paradigm in traditional engineering disciplines such as transportation, energy, telecommunications, and manufacturing. In modern data science, we note the renaissance of (GLP) due to its modeling power in such areas as reinforcement learning [19], optimal transport [57], and neural network verification [39]. For traditional engineering disciplines with moderate scale or exploitable sparsity, off-the-shelf interior point methods that form and factorize matrices in each

---

\*Equal contribution

iteration are often good choices as practical solvers [26]. In data science applications, however, where the data are often dense or of extreme scale, the amount of computation and/or memory required by matrix factorization is prohibitive. Thus, first-order methods that avoid matrix factorizations are potentially appealing options. In this context, because the presence of the linear equality constraint in (GLP) may complicate projection operations onto the feasible set, we consider an equivalent reformulation of (GLP) as a min-max problem involving the Lagrangian:

$$\min_{\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d} \max_{\mathbf{y} \in \mathbb{R}^n} \left\{ \mathcal{L}(\mathbf{x}, \mathbf{y}) := \mathbf{c}^T \mathbf{x} + r(\mathbf{x}) + \mathbf{y}^T \mathbf{A} \mathbf{x} - \mathbf{y}^T \mathbf{b} \right\}. \quad (\text{PD-GLP})$$

In data science applications, both  $n$  and  $d$  can be very large. (PD-GLP) can be viewed as a structured bilinearly coupled min-max problem, where the linearity of  $\mathcal{L}(\mathbf{x}, \mathbf{y})$  in the dual variable vector  $\mathbf{y}$  is vital to our algorithmic development.

## 1.1 Background

While there have been few papers that directly address (PD-GLP) — some special cases have been considered in [14, 25, 41–43, 60, 62, 63] — there has been significant recent work on first-order methods for general bilinearly coupled convex-concave min-max problems. Deterministic first-order methods include the proximal point method (PPM) [51], the extragradient/mirror-prox method (EGM) [34, 45], the primal-dual hybrid gradient (PDHG) method [15], and the alternating direction method of multipliers (ADMM) [20]. All these methods have per-iteration cost  $\Theta(\text{nnz}(\mathbf{A}))$  and convergence rate  $1/k$ , where  $\text{nnz}(\mathbf{A})$  denotes the number of nonzero elements of  $\mathbf{A}$  and  $k$  is the number of iterations.

For better scalability, stochastic counterparts of these methods have been proposed. [11, 33, 47, 49] have used “vanilla” stochastic gradients to replace the full gradients of their deterministic counterparts. [2, 13, 27] have exploited the finite-sum structure of the interaction term  $\langle \mathbf{y}, \mathbf{A} \mathbf{x} \rangle$  involving both primal and dual variables to perform variance reduction. With a separability assumption for the dual variables, [3] and [16] have combined incremental coordinate approaches on the dual variables with an implicit variance reduction strategy on the primal variables. Recently, under a separability assumption for dual variables, [55] proposed a new incremental coordinate method with an initialization step that requires a single access to the full data. This approach, known as *variance reduction via primal-dual accelerated dual averaging* (VRPDA<sup>2</sup>), obtains the first theoretical bounds that are better than their deterministic counterparts in the class of incremental coordinate approaches. The VRPDA<sup>2</sup> algorithm serves as the main motivation for our approach.

It is of particular interest to design algorithms that scale with the number of nonzero elements in  $\mathbf{A}$  for at least two reasons: (i) the data matrix can be sparse; and (ii) when we consider simplified reformulations of certain complicated models, we often need to introduce sparsely connected auxiliary variables. Nevertheless, the randomized coordinate algorithms of [3, 16, 55] have  $O(d)$  per-iteration cost regardless of the sparsity of  $\mathbf{A}$ . To address this issue, [24, 35] have proposed incremental primal-dual coordinate methods with per-iteration cost that scales with the number of nonzero elements in the row of  $\mathbf{A}$  used in each iteration, at the price of needing to take a smaller step than for dense  $\mathbf{A}$ . Moreover, [5] has proposed a random extrapolation approach that admits both low per-iteration cost and larger step size. Despite these developments, all these algorithms produce less accurate iterates than the methods with  $O(d)$  per-iteration cost, thus degrading their worst-case complexity.<sup>2</sup>

Finally, for the special case of LP, based on the positive Hoffman constant [30], [10] proved that the primal-dual formulation of LP exhibits a sharpness property that lower-bounds the growth of a normalized primal-dual gap from the same work. Leveraging this sharpness property, [10] proposed a restart scheme for the deterministic first-order methods discussed above to obtain linear convergence. [9] further extended this restart strategy using various heuristics to improve practical performance.

## 1.2 Motivation

We sharpen the focus from general bilinearly coupled convex-concave min-max problems to (GLP) and its primal-dual formulation (PD-GLP), because many complicated models can be reformulated as (GLP) and because this formulation possesses additional structure that can be exploited in algorithm

<sup>2</sup>Subsequent to this paper, a version of the PURE-CD algorithm of [5] that exploits sparsity in  $\mathbf{A}$  was developed and analyzed in [6].

design. Our motivation for focusing on (GLP) is to bridge the large gap between the well-studied stochastic variance reduced first-order methods [7, 32, 54, 55] and the increasingly popular and complicated, yet highly structured large-scale problems arising in distributionally robust optimization (DRO) [21–23, 31, 38, 44, 52, 58, 61]; see also a recent survey by [50] and references therein.

For DRO problems with ambiguity sets defined by  $f$ -divergence [31, 37, 44], the original formulation is a nonbilinearly coupled convex-concave min-max problem. Even the well constructed reformulation in [37] does not admit unbiased stochastic gradients, leading to complicated algorithms and analysis. For DRO problems with ambiguity sets defined by Wasserstein metric [23, 29, 38, 52, 61], the original formulation is in general infinite-dimensional. (Finite-dimensional reformulations [23, 52] exist for special cases of logistic regression and smooth convex losses.) Solvers that have been proposed for DRO with Wasserstein metric are either multiple-loop deterministic ADMM [38] or are designed for general convex-concave problems [61].

By introducing auxiliary variables with sparse connections,<sup>3</sup> we show that DRO with ambiguity sets based on both  $f$ -divergence and the Wasserstein metric can be reformulated as (GLP). Thus, complicated DRO problems can be addressed by a simple, efficient, and scalable algorithm for (GLP). Our algorithm for solving (GLP) and the proposed reformulations of DRO are our main contributions.

### 1.3 Contributions

**Algorithm.** Motivated by VRPDA<sup>2</sup> [55], we propose a simple, efficient, and scalable algorithm for (PD-GLP). Our algorithm combines an incremental *coordinate* method with exploitation of the *linear* structure for the dual variables in (PD-GLP) and the implicit *variance reduction* effect in the algorithm, so we name it *coordinate linear variance reduction* (CLVR, pronounced “clever”). CLVR is inspired by VRPDA<sup>2</sup> but customized to the particular structure of (PD-GLP). In particular, by exploiting the fact that the max problem is linear and unconstrained in the dual variable vector  $\mathbf{y} \in \mathbb{R}^n$ , we find that the expensive initialization step used in VRPDA<sup>2</sup> is not needed and we can take simpler and larger steps. Further, in the structured case in which  $\mathbf{A}$  is sparse and the convex constraint set  $\mathcal{X}$  and the regularizer  $r(\mathbf{x})$  are fully separable<sup>4</sup>, we show that the dual averaging update in CLVR enables us to design an efficient lazy update strategy for which the per-iteration cost of CLVR scales with the number of nonzero elements of the selected row from  $\mathbf{A}$  in each iteration, which is potentially much lower than the order- $d$  cost in VRPDA<sup>2</sup>. Finally, CLVR uses extrapolation on dual variables rather than on primal variables considered in VRPDA<sup>2</sup>, which significantly reduces implementation complexity of our lazy update strategy for structured variants of (PD-GLP). On the technical side, although both CLVR and VRPDA<sup>2</sup> are randomized algorithms that bound the primal-dual gap in expectation, the guarantee provided by CLVR is stronger as it allows bounding the expectation of the supremum gap as opposed to the supremum of expected gap in VRPDA<sup>2</sup>.

To state our complexity results, we make the following scaling assumption.

**Assumption 1.**  $L := \|\mathbf{A}\|$  and each row of  $\mathbf{A}$  in (GLP) is normalized with Euclidean norm  $R$ .

Preprocessing in modern LP solvers [26] often ensures normalized rows/columns for the data matrix. Observe that  $R \leq L \leq \sqrt{n}R$ , the upper bound being achieved when all elements of  $\mathbf{A}$  have identical value. Although the latter case is extreme, there exist ill-conditioned practical datasets where we can expect significant performance gains if the complexity can be reduced from  $O(L)$  to  $O(R)$ . (We provide empirical comparison between the values of  $L$  and  $R$  in practical problems in Section 5.)

In Table 1, we give the overall complexity bounds (total number of arithmetic operations) and the per-iteration cost of a representative set of existing algorithms, including our CLVR algorithm, for solving a structured form of (PD-GLP) in which the set  $\mathcal{X}$  and the function  $r$  have separable structure:  $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_d$  with  $\mathcal{X}_i \in \mathbb{R}$  ( $i \in [d]$ ) and  $r(\mathbf{x}) := \sum_{i=1}^d r(x^i)$ . To make the complexity results comparable, we assume further that for the stochastic algorithms [2, 16, 55] and our CLVR algorithm, we draw one row of  $\mathbf{A}$  per iteration uniformly at random. The general convex setting corresponds to  $r(\mathbf{x})$  being general convex ( $\sigma = 0$ ), while the strongly convex setting corresponds to  $r(\mathbf{x})$  being  $\sigma$ -strongly convex ( $\sigma > 0$ ).

<sup>3</sup>“Sparse connections” here means that even though the newly introduced variables may substantially increase the problem dimensions, the number of nonzero entries in the constraint matrix remains of the same order.

<sup>4</sup>We state the results here for the fully separable setting for convenience of comparison; however, our results are also applicable to the block separable setting.

Table 1: Overall complexity and per-iteration cost for solving structured (PD-GLP). (“—” indicates that the corresponding result does not exist or is unknown.)

Algorithm	General Convex (Primal-Dual Gap)	Strongly Convex (Distance to Solution)	Per-Iteration Cost
PDHG CP(2011)	$O\left(\frac{\text{nnz}(\mathbf{A})L}{\epsilon}\right)$	$O\left(\frac{(\text{nnz}(\mathbf{A})+n+d)L}{\sigma\sqrt{\epsilon}}\right)$	$O(\text{nnz}(\mathbf{A}))$
SPDHG CERS(2018)	$O\left(\frac{ndL}{\epsilon}\right)$	$O\left(\frac{ndL}{\sigma\sqrt{\epsilon}}\right)$	$O(d)$
EVR AM (2022)	$O\left(\text{nnz}(\mathbf{A}) + \frac{\sqrt{\text{nnz}(\mathbf{A})(n+d)nR}}{\epsilon}\right)$	—	$O(n+d)$
VRPDA <sup>2</sup> SWD(2021)	$O(nd \log \min\{\frac{1}{\epsilon}, n\} + \frac{ndR}{\epsilon})$	$O(nd \log \min\{\frac{1}{\epsilon}, n\} + \frac{ndR}{\sigma\sqrt{\epsilon}})$	$O(d)$
CLVR (This Paper)	$O\left(\frac{\text{nnz}(\mathbf{A})R}{\epsilon}\right)$	$O\left(\frac{\text{nnz}(\mathbf{A})R}{\sigma\sqrt{\epsilon}}\right)$	$O(\text{nnz}(\text{row}(\mathbf{A})))$

As shown in Table 1, all the algorithms have optimal dependence on  $\epsilon$  [48], while the dependence on the ambient dimensions  $n, d$ , the number of nonzero elements of  $\mathbf{A}$  ( $\text{nnz}(\mathbf{A})$ ), and the constants  $L$  and  $R$  are quite different. For both the general convex and strongly convex settings and among coordinate-type methods, CLVR is the first algorithm that reduces the runtime dependence on the input matrix size from  $nd$  to  $\text{nnz}(\mathbf{A})$ . Moreover, the complexity of CLVR depends on the max row norm  $R$  rather than the spectral norm  $L$ , and the per-iteration cost of CLVR depends only on the nonzero elements of the selected row from  $\mathbf{A}$  in each iteration, which can be far less than  $d$ .

By exploiting the linear structure again, we provide explicit guarantees for both the objective value and the constraint satisfaction of (GLP). Further, the analysis of CLVR applies to the more general *block-coordinate* update setting, which is better suited to modern parallel computing platforms. Finally, following the restart strategy based on the *normalized duality gap* for LP introduced in [10], we propose a more straightforward strategy to restart our CLVR algorithm (as well as other iterative algorithms for (PD-GLP)): Restart the algorithm every time a widely known metric for LP optimality [8] halves. Compared with the normalized duality gap, the LPMetric can be computed more efficiently and in a more straightforward fashion.

**DRO reformulations.** When the loss function is convex, DRO problems with ambiguity sets based on  $f$ -divergence [44] or Wasserstein metric [23] are convex. However, because both problems either have complicated constraints or are infinite-dimensional, vanilla first-order methods are inapplicable.

For DRO with  $f$ -divergence, we show that by using convex conjugates and introducing auxiliary variables, the problem can be reformulated as a (GLP). As a result, the issue of biased stochastic gradients encountered in [37] does not arise, and CLVR can be applied. Even though the resulting problem has larger dimensions, due to the sparseness of the introduced auxiliary variables and the lazy update strategy of CLVR, it can be solved with complexity scaling only with the number of nonzero elements of the data matrix. Due to being cast as a (GLP), the DRO problem can be solved with  $O(1/\epsilon)$  iteration complexity with CLVR, while existing methods such as [37] have  $O(1/\epsilon^2)$  iteration complexity, with higher iteration cost because of the batch of samples needed to reduce bias. This improvement is enabled in part by considering the primal-dual gap (rather than the primal gap considered in [37]) and by allowing the constraints to be approximately satisfied (see Corollary 1).

For DRO with Wasserstein metric, following the reformulation of [52, Theorem 1], we show further that the problem can be cast in the form of (GLP). Compared with the existing reformulations [23, 38, 52, 61], our reformulation can handle both smooth and nonsmooth convex loss functions. In fact, our reformulation can provide a more compact form for nonsmooth piecewise-linear convex loss functions (such as hinge loss). Moreover, compared with algorithms customized to this problem [38] and extragradient methods [34, 45, 61] for general convex-concave min-max problems, our CLVR method attains the best-known iteration complexity and per-iteration cost, as shown in Table 1.

## 2 Notation and preliminaries

For any positive integer  $p$ , we use  $[p]$  to denote  $\{1, 2, \dots, p\}$ . We assume that there is a given partition of the set  $[n]$  into sets  $S^j, j \in [m]$ , where  $|S^j| = n^j > 0$  and  $\sum_{j=1}^m n^j = n$ . For  $j \in [m]$ , we use  $\mathbf{A}^{S^j}$

to denote the submatrix of  $\mathbf{A}$  with rows indexed by  $S^j$  and  $\mathbf{y}^{S^j}$  to denote the subvector of  $\mathbf{y}$  indexed by  $S^j$ . We use  $\mathbf{0}_d$  and  $\mathbf{1}_d$  to denote the vectors with all ones and all zeros in  $d$  dimensions, respectively. Unless otherwise specified, we use  $\|\cdot\|$  to denote the Euclidean norm for vectors and the spectral norm for matrices. For a given proper convex lower semi-continuous function  $f : \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$ , we define the convex conjugate in the standard way as  $f^*(y) = \sup_{x \in \mathbb{R}} \{yx - f(x)\}$  (so that  $f^{**} = f$ ). For a vector  $\mathbf{u}$ , the inequality  $\mathbf{u} \geq \mathbf{0}$  is applied entry-wise. For a convex function  $r(\mathbf{x})$ , we use  $r'(\mathbf{x})$  to denote an element of the subdifferential set  $\partial r(\mathbf{x})$ . The proximal operator of  $r(\mathbf{x})$  over  $\mathcal{X}$  is

$$\text{prox}_r(\hat{\mathbf{x}}) = \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}\|^2 + r(\mathbf{x}) \right\}. \quad (1)$$

Further, we make the following assumptions, which apply throughout the convergence analysis.

**Assumption 2.** (PD-GLP) *attains at least one primal-dual solution  $(\mathbf{x}^*, \mathbf{y}^*)$ .  $\mathcal{W}^*$  denotes the set of all primal-dual solutions.*

Due to the convex-concave property of (PD-GLP),  $\mathcal{W}^*$  is a convex set in  $\mathcal{X} \times \mathbb{R}^n$ .

**Assumption 3.**  $\hat{L} = \max_{j \in [m]} \|\mathbf{A}^{S^j}\|$  is given at the input, where  $\|\mathbf{A}^{S^j}\| = \max_{\|\mathbf{x}\| \leq 1} \|\mathbf{A}^{S^j} \mathbf{x}\|$ .

Note that  $\hat{L}$  can be obtained either via preprocessing of the data or by parameter tuning. By combining Assumptions 1 and 3, it follows that  $R \leq \hat{L} \leq \sqrt{\max_{j \in [m]} |S^j|} R$ .

**Assumption 4.**  $r(\mathbf{x})$  is  $\sigma$ -strongly convex ( $\sigma \geq 0$ ); that is, for all  $\mathbf{x}_1$  and  $\mathbf{x}_2$  in  $\mathcal{X}$  and all  $r'(\mathbf{x}_2) \in \partial r(\mathbf{x}_2)$ , we have  $r(\mathbf{x}_1) \geq r(\mathbf{x}_2) + \langle r'(\mathbf{x}_2), \mathbf{x}_1 - \mathbf{x}_2 \rangle + \frac{\sigma}{2} \|\mathbf{x}_1 - \mathbf{x}_2\|^2$ .

For convex-concave min-max problems, a common metric for measuring solution quality is the primal-dual gap, which, for a feasible solution  $(\mathbf{x}, \mathbf{y})$  of (PD-GLP), is defined by

$$\sup_{(\mathbf{u}, \mathbf{v}) \in \mathcal{X} \times \mathbb{R}^n} \{ \mathcal{L}(\mathbf{x}, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \mathbf{y}) \}. \quad (2)$$

However, as the domain of  $\mathbf{v}$  is unbounded, the primal-dual gap can be infinite, which makes it a poor metric for measuring the progress of algorithms. As a result, for measuring the progress of our algorithm, we consider the following restricted primal-dual gap instead:

$$\sup_{(\mathbf{u}, \mathbf{v}) \in \mathcal{W}} \{ \mathcal{L}(\mathbf{x}, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \mathbf{y}) \}, \quad (3)$$

where  $\mathcal{W} \subset \mathcal{X} \times \mathbb{R}^n$  is a compact (i.e., closed and bounded) convex set. The use of a restricted version of primal-dual gap is standard in the existing literature; see, e.g., [15, 46].

### 3 The CLVR algorithm

#### 3.1 Algorithm and analysis for general formulation

Algorithm 1 specifies CLVR for (PD-GLP) in the general setting. The algorithm alternates the full update for  $\mathbf{x}_k$  in Step 4 ( $O(d)$  cost) with an incremental block coordinate update for  $\mathbf{y}_k$  in Steps 5 and 6 (with  $O(|S^{j_k}|d)$  cost for dense  $\mathbf{A}$ ). The auxiliary variables  $\mathbf{z}_k$  and  $\mathbf{q}_k$  accumulate the cancellation terms in the estimation sequence and give a pathway to a straightforward development of the lazified CLVR, which appears as Algorithm 2 in the appendix. The cost of updating auxiliary vectors  $\mathbf{z}_k$  and  $\mathbf{q}_k$  is  $O(|S^{j_k}|d)$  and  $O(d)$ , respectively. In essence, CLVR is a primal-dual coordinate method that uses a *dual averaging* update for  $\mathbf{x}_k$ , then updates the state variables  $\{\mathbf{q}_k\}$  by a *linear recursion*, and computes  $\mathbf{x}_k$  from  $\mathbf{q}_{k-1}$  via a *proximal step* without direct dependence on  $\mathbf{x}_{k-1}$ . The output  $\tilde{\mathbf{x}}_K$  is a convex combination of the iterates  $\{\mathbf{x}_k\}_{k=1}^K$ , as is standard for primal-dual methods. However,  $\tilde{\mathbf{y}}_K$  is only an *affine* (not convex) combination of  $\{\mathbf{y}_k\}_{k=0}^K$ , as it involves the term  $-(m-1)\mathbf{y}_0$  (whose coefficient is negative) and some of the coefficients  $ma_k - (m-1)a_{k+1}$  multiplying  $\mathbf{y}_k$  for  $k \in \{1, \dots, K-1\}$  may also be negative. An affine combination still provides valid bounds because the dual variable vector  $\mathbf{y}$  appears linearly in (PD-GLP). Moreover, in Step 9, the term  $ma_k(\mathbf{z}_k - \mathbf{z}_{k-1})$  serves to cancel certain errors from the randomization of the update w.r.t.  $\mathbf{y}_k$ , thus playing a key role in implicit variance reduction.

Theorem 1 provides the convergence results for Algorithm 1. The proof is provided in Appendix B. In the theorem (as in the algorithm),  $\gamma$  is a positive parameter that can be tuned.

---

**Algorithm 1** Coordinate Linear Variance Reduction (CLVR)
 

---

- 1: **Input:**  $\mathbf{x}_0 \in \mathcal{X}, \mathbf{y}_0 \in \mathbb{R}^n, \mathbf{z}_0 = \mathbf{A}^T \mathbf{y}_0, \gamma > 0, \hat{L} > 0, \sigma \geq 0, K, m, \{S^1, S^2, \dots, S^m\}$ .
  - 2:  $a_1 = A_1 = \frac{1}{2\hat{L}m}, \mathbf{q}_0 = a_1(\mathbf{z}_0 + \mathbf{c})$ .
  - 3: **for**  $k = 1, 2, \dots, K$  **do**
  - 4:  $\mathbf{x}_k = \text{prox}_{\frac{1}{\gamma} A_k r}(\mathbf{x}_0 - \frac{1}{\gamma} \mathbf{q}_{k-1})$ .
  - 5: Pick  $j_k$  uniformly at random in  $[m]$ .
  - 6:  $\mathbf{y}_k^{S^i} = \begin{cases} \mathbf{y}_{k-1}^{S^i}, & i \neq j_k \\ \mathbf{y}_{k-1}^{S^i} + \gamma m a_k (\mathbf{A}^{S^i} \mathbf{x}_k - \mathbf{b}^{S^i}), & i = j_k \end{cases}$ .
  - 7:  $a_{k+1} = \frac{\sqrt{1+\sigma A_k/\gamma}}{2\hat{L}m}, A_{k+1} = A_k + a_{k+1}$ .
  - 8:  $\mathbf{z}_k = \mathbf{z}_{k-1} + \mathbf{A}^{S^{j_k}, T}(\mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}})$ .
  - 9:  $\mathbf{q}_k = \mathbf{q}_{k-1} + a_{k+1}(\mathbf{z}_k + \mathbf{c}) + m a_k(\mathbf{z}_k - \mathbf{z}_{k-1})$ .
  - 10: **end for**
  - 11: **return**  $\tilde{\mathbf{x}}_K = \frac{1}{A_K} \sum_{k=1}^K a_k \mathbf{x}_k, \tilde{\mathbf{y}}_K = \frac{1}{A_K} \sum_{k=1}^K (a_k \mathbf{y}_k + (m-1)a_k(\mathbf{y}_k - \mathbf{y}_{k-1}))$ .
- 

**Theorem 1.** Let  $\mathbf{x}_k, \mathbf{y}_k, k \in [K]$ , be the iterates of Algorithm 1 and let  $\tilde{\mathbf{x}}_k, \tilde{\mathbf{y}}_k$  be defined by

$$\tilde{\mathbf{x}}_k = \frac{1}{A_k} \sum_{i=1}^k a_i \mathbf{x}_i, \quad \tilde{\mathbf{y}}_k = \frac{1}{A_k} \sum_{i=1}^k (a_i \mathbf{y}_i + (m-1)a_i(\mathbf{y}_i - \mathbf{y}_{i-1})), \quad (4)$$

for  $k \in [K]$ . Let  $\mathcal{W}_k \subset \mathcal{X} \times \mathbb{R}^n, k \in [K]$ , be a sequence of compact convex sets such that  $(\tilde{\mathbf{x}}_k, \tilde{\mathbf{y}}_k) \in \mathcal{W}_k \subset \mathcal{W} \subset \mathcal{X} \times \mathbb{R}^n$ , where  $\mathcal{W}$  is also convex and compact. Then:

$$\begin{aligned} & \mathbb{E} \left[ \sup_{(\mathbf{u}, \mathbf{v}) \in \mathcal{W}_k} \{ \mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k) \} \right] \\ & \leq \frac{1}{A_k} \left( \mathbb{E} \left[ \frac{\gamma}{2} \|\hat{\mathbf{u}} - \mathbf{x}_0\|^2 + \frac{1}{\gamma} \|\hat{\mathbf{v}} - \mathbf{y}_0\|^2 \right] + \frac{\gamma}{2} \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2 \right), \end{aligned} \quad (5)$$

where  $(\hat{\mathbf{u}}, \hat{\mathbf{v}}) = \arg \sup_{(\mathbf{u}, \mathbf{v}) \in \mathcal{W}_k} \{ \mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k) \}$ . Furthermore,

$$\mathbb{E} \left[ \frac{\gamma + \sigma A_k}{4} \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}^*\|^2 \right] \leq \frac{\gamma}{2} \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2. \quad (6)$$

Define  $K_0 = \lceil \frac{\sigma}{18\hat{L}m\gamma} \rceil$ . Then in the bounds above:

$$A_k \geq \max \left\{ \frac{k}{2\hat{L}m}, \frac{\sigma}{(6\hat{L}m)^2\gamma} \left( k - K_0 + \max \left\{ 3\sqrt{2\hat{L}m\gamma/\sigma}, 1 \right\} \right)^2 \right\}.$$

Observe that  $(\hat{\mathbf{u}}, \hat{\mathbf{v}})$  in the theorem statement exists because of compactness of  $\mathcal{W}_k$  and our assumptions on  $r(\cdot)$ . The parameter  $\gamma$  can be tuned to balance the relative weights of primal and dual initial quantities  $\|\mathbf{x}^* - \mathbf{x}_0\|$  and  $\|\mathbf{y}^* - \mathbf{y}_0\|$  (or estimates of these quantities), which can significantly influence practical performance of the method.

In addition to the guarantee on the variational form, due to the linear structure, we also provide explicit guarantees for both the objective and the constraints in (GLP), stated in the following corollary.

**Corollary 1.** In Algorithm 1, for all  $k \geq 1$ ,  $\tilde{\mathbf{x}}_k$  satisfies

$$\begin{aligned} \mathbb{E}[\|\mathbf{y}^*\| \cdot \|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|] & \leq \frac{\gamma \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2 + \frac{1}{\gamma} \mathbb{E}[\|\mathbf{v} - \mathbf{y}_0\|^2]}{A_k}, \\ \mathbb{E}[(\mathbf{c}^T \tilde{\mathbf{x}}_k + r(\tilde{\mathbf{x}}_k)) - (\mathbf{c}^T \mathbf{x}^* + r(\mathbf{x}^*))] & \leq \frac{\gamma \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2 + \frac{1}{\gamma} \mathbb{E}[\|\mathbf{v} - \mathbf{y}_0\|^2]}{A_k}, \end{aligned}$$

where  $\mathbf{v} = 2 \frac{\|\mathbf{y}^*\|}{\|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|} (\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b})$ .

In CLVR, we allow for arbitrary  $(\mathbf{x}_0, \mathbf{y}_0) \in \mathcal{X} \times \mathbb{R}^n$ . Nevertheless, by setting  $\mathbf{y} = \mathbf{0}_n$ , we can obtain  $\mathbf{z}_0 = \mathbf{0}_d$  at no cost — a useful strategy for large-scale problems since it avoids the (potentially expensive) single matrix-vector multiplication w.r.t.  $\mathbf{A}$ .

### 3.2 Lazy update for sparse and structured (PD-GLP)

In Algorithm 1, direct computation of the iterates  $(\mathbf{x}_k, \mathbf{y}_k)$  and the output points  $(\tilde{\mathbf{x}}_k, \tilde{\mathbf{y}}_k)$  can be expensive. However, [18] showed that it is possible to only update the averaged vector in the coordinate block chosen for that iteration. This strategy requires us to record the most recent update for each coordinate block and update it only when it is selected again, which is tricky and needs to be implemented carefully. For sparse and block coordinate-separable instances of (PD-GLP), we show that by introducing auxiliary variables that are sparsely connected, we can significantly simplify CLVR and make its complexity scale independently of the ambient dimension  $n \cdot d$ , instead scaling with  $\text{nnz}(\mathbf{A})$ . Due to space constraints, we defer technical details, including the lazy version of CLVR and associated proofs, to Appendix A.

### 3.3 Restart scheme

We now propose a fixed restart strategy with a fixed number of iterations per each restart epoch and discuss an adaptive restart strategy for the special case of standard-form LP, which corresponds to (GLP) with  $r(\mathbf{x}) \equiv 0$  and  $\mathcal{X} = \{\mathbf{x} : x_i \geq 0, i \in [d]\}$ . We write

$$\min_{\mathbf{x}} \mathbf{c}^T \mathbf{x} \text{ s. t. } \mathbf{A} \mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}_d, \quad (\text{LP})$$

and the primal-dual form

$$\min_{\mathbf{x} \geq \mathbf{0}_d} \max_{\mathbf{y} \in \mathbb{R}^n} \left\{ \mathcal{L}(\mathbf{x}, \mathbf{y}) = \mathbf{c}^T \mathbf{x} + \mathbf{y}^T \mathbf{A} \mathbf{x} - \mathbf{y}^T \mathbf{b} \right\}. \quad (\text{PD-LP})$$

This problem has a sharpness property that can be used to obtain linear convergence in first-order methods [10]. For convenience, in the following, we define  $\mathbf{w} = (\mathbf{x}, \mathbf{y})$ ,  $\hat{\mathbf{w}} = (\hat{\mathbf{x}}, \hat{\mathbf{y}})$ ,  $\tilde{\mathbf{w}} = (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$  and  $\mathbf{w}^* = (\mathbf{x}^*, \mathbf{y}^*)$ . Meanwhile, for  $\gamma > 0$ , we denote the weighted norm  $\|\mathbf{w}\|_{(\gamma)} := \sqrt{\gamma \|\mathbf{x} - \mathbf{x}^*\|_2^2 + \frac{1}{\gamma} \|\mathbf{y} - \mathbf{y}^*\|_2^2}$ . Further, we use  $\mathcal{W}^*$  to denote the optimal solution set of the LP and define the distance to  $\mathcal{W}^*$  by  $\text{dist}(\mathbf{w}, \mathcal{W}^*)_{(\gamma)} = \min_{\mathbf{w}^* \in \mathcal{W}^*} \|\mathbf{w} - \mathbf{w}^*\|_{(\gamma)}$ . When  $\gamma = 1$ ,  $\|\cdot\|_{(\gamma)}$  is the standard Euclidean norm. Then based on (PD-LP), we can use the following classical LPMetric<sup>5</sup> to measure the progress of iterative algorithms for LP:

$$\begin{aligned} & \text{LPMetric}(\mathbf{x}, \mathbf{y}) \\ &= \sqrt{\|\max\{-\mathbf{x}, \mathbf{0}\}\|_2^2 + \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2^2 + \|\max\{-\mathbf{A}^T \mathbf{y} - \mathbf{c}, \mathbf{0}\}\|_2^2 + |\max\{\mathbf{c}^T \mathbf{x} + \mathbf{b}^T \mathbf{y}, 0\}|^2}, \quad (7) \end{aligned}$$

which can be explicitly and directly computed. For the Euclidean case ( $\gamma = 1$ ), it is well-known [30] that there exists a Hoffman constant  $H_1$  such that

$$\text{LPMetric}(\mathbf{w}) \geq H_1 \text{dist}(\mathbf{w}, \mathcal{W}^*)_{(1)}. \quad (8)$$

Using the equivalence of norms in finite dimensions, for general  $\gamma > 0$ , we can conclude that there exists another constant  $H_\gamma$  (to which we refer as the generalized Hoffman's constant) such that

$$\text{LPMetric}(\mathbf{w}) \geq H_\gamma \text{dist}(\mathbf{w}, \mathcal{W}^*)_{(\gamma)}. \quad (9)$$

Using Eq. (9) and Theorem 1, we then obtain the following bounds for distance and LPMetric.

**Theorem 2.** Consider the CLVR algorithm applied to the standard-form LP problem (PD-LP), with input  $\mathbf{w}_0$  and output  $\tilde{\mathbf{w}}_k$ . Given  $\gamma > 0$ , define  $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}^*} \|\mathbf{w}_0 - \mathbf{w}\|_{(\gamma)}$ , and define  $C_0 = \gamma + 1/\gamma + (\sqrt{2} + 1)\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} + \|\mathbf{w}^*\|_{(\gamma)}$ . Then for  $H_\gamma$  defined as in (9), we have

$$\begin{aligned} \mathbb{E} \left[ \sqrt{\text{dist}(\tilde{\mathbf{w}}_k, \mathcal{W}^*)_{(\gamma)}} \right] &\leq 5 \sqrt{\frac{\hat{L} m C_0}{H_\gamma k}} \sqrt{\text{dist}(\mathbf{w}_0, \mathcal{W}^*)_{(\gamma)}}, \\ \mathbb{E} \left[ \sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)} \right] &\leq 5 \sqrt{\frac{\hat{L} m C_0}{H_\gamma k}} \sqrt{\text{LPMetric}(\mathbf{w}_0)}. \end{aligned}$$

<sup>5</sup>In (PD-LP), we dualize the constraint  $\mathbf{A} \mathbf{x} = \mathbf{b}$  by  $\mathbf{y}^T (\mathbf{A} \mathbf{x} - \mathbf{b})$  instead of  $\mathbf{y}^T (\mathbf{b} - \mathbf{A} \mathbf{x})$ , so in our LPMetric, there exist a sign difference for  $\mathbf{y}$  from the more common representation such as the one in [10].

As a result, by Theorem 2, if we know the values of  $\hat{L}$ ,  $\|\mathbf{w}^*\|_{(\gamma)}$  and  $H_\gamma$ , then by setting  $k = \frac{100\hat{L}mC_0}{H_\gamma k}$ , we can halve the square root of the distance and the LPMetric in expectation. Thus we can obtain linear convergence if we restart the CLVR algorithm after a fixed number of iterations. However, the values of  $\|\mathbf{w}^*\|_{(\gamma)}$  and  $H_\gamma$  are often unknown and thus make this strategy unrealistic in practice.

Compared with the above fixed restart strategy, a natural strategy is to restart whenever the LPMetric halves (summarized in Algorithm 4 in the appendix). Since LPMetric is easy to monitor and update, implementation of this strategy is straightforward. However, bounding the number of iterations required to halve the metric (in expectation or with high probability) seems nontrivial. What can be said (based on Theorem 2 and denoting by  $K$  the number of iterations on CLVR between restarts) is that  $\mathbb{P}[K > \frac{50\hat{L}mC_0}{\delta^2 H_\gamma}] \leq \delta$ . This follows by Markov inequality, as  $\mathbb{P}[K > k] = \mathbb{P}[\sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)} > \sqrt{\frac{\text{LPMetric}(\mathbf{w}_0)}{2}}] \leq 5\sqrt{2\frac{\hat{L}mC_0}{H_\gamma k}}$ . We provide a comparison between the adaptive restart scheme proposed in [10] and our proposed adaptive restart scheme in Section D.1 to demonstrate its practical competitiveness. Although we use adaptive restart in our experiments, we defer its convergence analysis to future work. Finally, as an independent and parallel work to ours, [40] proposed a high probability guarantee for scheduled restart for stochastic extragradient-type methods.

## 4 Application: DRO

Consider sample vectors  $\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$  with labels  $\{b_1, b_2, \dots, b_n\}$ , where  $b_i \in \{1, -1\}$  ( $i \in [n]$ ). The DRO problem with  $f$ -divergence based ambiguity set is

$$\min_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbf{p} \in \mathcal{P}_{\rho, n}} \sum_{i=1}^n p_i g(b_i \mathbf{a}_i^T \mathbf{x}), \quad (10)$$

where  $\mathcal{P}_{\rho, n} = \{\mathbf{p} \in \mathbb{R}^n : \sum_{i=1}^n p_i = 1, p_i \geq 0 (i \in [n]), D_f(\mathbf{p} \parallel \mathbf{1}/n) \leq \frac{\rho}{n}\}$  is the ambiguity set,  $g$  is a convex loss function and  $D_f$  is an  $f$ -divergence defined by  $D_f(\mathbf{p} \parallel \mathbf{q}) = \sum_{i=1}^n q_i f(p_i/q_i)$  with  $\mathbf{p}, \mathbf{q} \in \{\mathbf{p} \in \mathbb{R}^n : \sum_{i=1}^n p_i = 1, p_i \geq 0\}$  and  $f$  being a convex function [44]. The formulation (10) is a nonbilinearly coupled convex-concave min-max problem with constraint set  $\mathcal{P}_{\rho, n}$  for which efficient projections are not available in general. When  $g$  is a nonsmooth loss (e.g., the hinge loss), many well-known methods such as the extragradient [34, 45] cannot be used even if we could project onto  $\mathcal{P}_{\rho, n}$  efficiently. However, by introducing auxiliary variables, additional linear constraints, and simple convex constraints, we can make the interacting term between primal and dual variables bilinear, as shown next. (See Appendix C for a proof.)

**Theorem 3.** *Let  $\mathcal{X}$  be a compact convex set. Then the DRO problem in Eq. (10) is equivalent to*

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}, \mathbf{v}, \mathbf{w}, \boldsymbol{\mu}, \mathbf{q}, \gamma} \quad & \left\{ \gamma + \frac{\rho\mu_1}{n} + \frac{1}{n} \sum_{i=1}^n \mu_i f^*\left(\frac{q_i}{\mu_i}\right) \right\} \\ \text{s. t.} \quad & \mathbf{w} + \mathbf{v} - \frac{\mathbf{q}}{n} - \gamma \mathbf{1}_n = \mathbf{0}_n, \\ & u_i = b_i \mathbf{a}_i^T \mathbf{x}, \quad i \in [n] \\ & \mu_1 = \mu_2 = \dots = \mu_n, \\ & g(u_i) \leq w_i, \quad i \in [n] \\ & q_i \in \mu_i \text{ dom}(f^*), \quad i \in [n] \\ & v_i \geq 0, \mu_i \geq 0, \quad i \in [n] \\ & \mathbf{x} \in \mathcal{X}. \end{aligned}$$

In Theorem 3, the domain of the one-dimensional convex function  $f^*(\cdot)$  is an interval such as  $[a, b]$ , so that  $q_i \in \mu_i \text{ dom}(f^*)$  denotes the inequality  $\mu_i a \leq q_i \leq \mu_i b$ . Since the perspective function  $\mu f^*(\frac{q}{\mu})$  is a simple convex function of two variables, we can assume that the proximal operator for this function on the domain  $\{(\mu, q) : q \in \mu \text{ dom}(f^*), \mu > 0\}$  can be computed efficiently [12]. Similarly, we can assume that the constraint  $g(u) \leq w$  admits an efficiently computable projection operator. As a result, the formulation (10) can be solved by CLVR. When expressing (10) in the form of (PD-GLP), the primal and dual variable vectors have dimensions  $d + 1 + 4n$  and  $3n - 1$ ,

respectively. However, according to Table 1, provided that  $\mathcal{X}$  is coordinate separable, the overall complexity of CLVR will only be  $O\left(\frac{(\text{nnz}(\mathbf{A})+n)(R+1)}{\epsilon}\right)$ .

The original DRO problem with Wasserstein metric based ambiguity set is an *infinite*-dimensional *nonbilinearly* coupled convex-concave min-max problem defined by

$$\min_{\mathbf{w} \in \mathbb{R}^d} \sup_{\mathbb{P} \in \mathcal{P}_{\rho, \kappa}} \mathbb{E}^{\mathbb{P}}[g(\mathbf{b}\mathbf{a}^T \mathbf{w})], \quad (11)$$

where  $\mathbf{a} \in \mathbb{R}^d$ ,  $\mathbf{b} \in \{1, -1\}$ ,  $\mathbb{P}$  is a distribution on  $\mathbb{R}^d \times \{1, -1\}$ ,  $g$  is a convex loss function and  $\mathcal{P}_{\rho, \kappa}$  is the Wasserstein metric-based ambiguity set [52]. Our reformulation for Eq. (11) is in Appendix C.2.

## 5 Numerical experiments

We provide experimental evaluations of our algorithm for the reformulation of the DRO with Wasserstein metric based on the  $\ell_1$ -norm (with  $\kappa = 0.1$  and  $\rho = 10$ ) and hinge loss. For its LP formulation (see Theorem 4 in the Appendix), we compare our CLVR method with three representative methods: PDHG [15], SPDHG [16] and PURE-CD [5]. For all algorithms we use LPMetric (7) as the performance measure and use a restart strategy based on successive halving of LPMetric (Section 3.3) to obtain linear convergence. We implemented CLVR and other algorithms in Julia, optimizing all implementations to the extent possible. Full details of the experimental setup can be found in Appendix D. Our code is available at <https://github.com/ericclinc/Efficient-GLP>.

**Comparison between values of  $L$  and  $R$ .** As described in Section 1, a major advantage of CLVR is that the complexity of CLVR depends on the max row norm  $R$  instead of the spectral norm  $L$ , which in the worst case for ill-conditioned problems can lead to a factor of  $\sqrt{n}$  improvement. In practical problems where the problem instances are highly structured (e.g., reformulated DRO problems),  $R$  can be much smaller than  $L$ . Table 2 provides empirical evidence for this claim. In all our experiments, we normalize each rows of  $\mathbf{A}$  to  $R = 1$  as stated in Assumption 1, so the values of  $L$  demonstrate the theoretical improvements for the experiments described in Section 5.

Table 2: Values of the spectral norm  $L$  in the reformulated DRO problems with Wasserstein metric after each row is normalized to  $R = 1$ .

Reformulated a9a $d = 130738, n = 97929$	Reformulated gisette $d = 44002, n = 28000$	Reformulated rcv1 $d = 269914, n = 155198$	Reformulated news20 $d = 5500750, n = 2770370$
117.3	65.9	196.4	1041.6

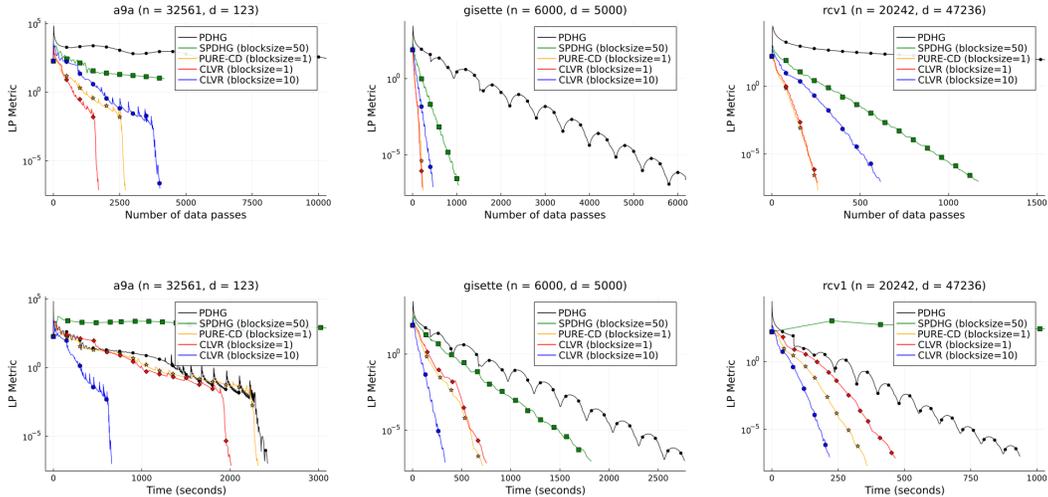


Figure 1: Comparison of numerical results in terms of number of data passes and wall-clock time.

**Comparison with primal-dual algorithms.** Figure 1 provides a comparison between algorithms in terms of the number of data passes and wall-clock time. The spikes in all the plots are due to restarts: At the beginning of each restart cycle, the value of LPMetric increases significantly, then decreases rapidly. For the number of data passes (top row), CLVR with block size 1 and PURE-CD perform best on all three datasets, CLVR with block size 10 and SPDHG with block size 50 have second-tier performance, and PDHG is worst. For the CLVR algorithm, smaller block size corresponds to smaller  $\hat{L}$  in Assumption 3, which corresponds to better complexity in terms of data passes by Theorem 1. Nevertheless, the gap between empirical performance and theoretical guarantee for SPDHG and PURE-CD deserves further research because, to date, they have only been shown to have the same iteration complexity as PDHG.<sup>6</sup> Empirically, on a9a, CLVR with block size 1 performs better than PURE-CD in terms of data passes.

In terms of wall-clock time (bottom row of Figure 1), because of different per-iteration costs of each algorithm and instruction-level parallelism in modern processors [28], the plots differ significantly from the plots for number of data passes. Even with block size 50, SPDHG spends the most wall-clock time for one data pass and is the slowest on sparse datasets a9a and rcv1, but is faster than PDHG on the dense dataset gisette. Meanwhile, while CLVR with block size 10 is not best in terms of data passes, it remains fastest in terms of wall-clock time on all datasets due to cheaper per-iteration cost and instruction-level parallelism. On rcv1, the per-iteration cost of PURE-CD is about 60% of that of CLVR with block size 1. Hence, despite having similar performance in terms of data passes, PURE-CD is faster than CLVR with block size 1, but is still slower than CLVR with block size 10.

**Comparison with production linear programming solvers.** Table 3 shows that CLVR is competitive against production-quality linear programming solvers such as GLPK [1] and Gurobi [26]. We observe that CLVR reached accurate solutions significantly faster than GLPK and Gurobi in the reformulated problems with gisette and rcv1 datasets. Although CLVR is much slower than Gurobi(barrier) on a9a dataset, we believe that much of the performance gap in this case is due to the redundancy in the problem formulation with the a9a dataset, much of which is removed by Gurobi presolver<sup>7</sup>. We leave presolving and other heuristic speedups of CLVR for future work.

Table 3: Comparison of numerical results between CLVR and three production solvers for linear programming, showing time required (in seconds) for each solver to reach accuracy  $10^{-8}$ .

Time (seconds)	Reformulated a9a $d = 130738, n = 97929$	Reformulated gisette $d = 44002, n = 28000$	Reformulated rcv1 $d = 269914, n = 155198$
JuMP+GLPK	899	$> 4 \times 10^4$	$> 4 \times 10^4$
JuMP+Gurobi(simplex)	893	2482	7008
JuMP+Gurobi(barrier)	<b>26</b>	1039.7	1039.5
CLVR	962	<b>697</b>	<b>582</b>

**Conclusion.** Our preliminary numerical experiments show that CLVR is fastest in both the number of data passes and wall-clock time on considered datasets, among all primal-dual algorithms that we implemented. It is also competitive with production-quality linear programming solvers. Since it has a theoretical guarantee that matches or improves the state of the art among primal-dual methods, we believe that CLVR could be a method of choice.

## Acknowledgments and Disclosure of Funding

CS was supported in part by the NSF grant 2023239. JD and CYL acknowledge support from the NSF award 2007757. JD was also supported by the Office of Naval Research under contract number N00014-22-1-2348 and the Wisconsin Alumni Research Foundation. SW was supported by NSF grants 2023239 and 2224213, the DOE under subcontract 8F-30039 from Argonne National Laboratory, and the AFOSR under subcontract UTA20-001224 from UT-Austin. Part of this work was done while JD, CS, and SW were visiting the Simons Institute for the Theory of Computing.

<sup>6</sup>The later paper [6] describes complexity results for a newly developed version of PURE-CD that exploits sparsity in  $A$ .

<sup>7</sup>In our DRO instance with a9a dataset, Gurobi presolver removed 25% of the columns and 58% of the nonzeros.

## References

- [1] GLPK (GNU linear programming kit). <https://www.gnu.org/software/glpk/>, 2022.
- [2] A. Alacaoglu and Y. Malitsky. Stochastic variance reduction for variational inequality methods. In *Conference on Learning Theory*, pages 778–816, 2022.
- [3] A. Alacaoglu, Q. T. Dinh, O. Fercoq, and V. Cevher. Smooth primal-dual coordinate descent algorithms for nonsmooth convex optimization. In *Proc. NIPS’17*, 2017.
- [4] A. Alacaoglu, O. Fercoq, and V. Cevher. On the convergence of stochastic primal-dual hybrid gradient. *arXiv preprint arXiv:1911.00799*, 2019.
- [5] A. Alacaoglu, O. Fercoq, and V. Cevher. Random extrapolation for primal-dual coordinate descent. In *Proc. ICML’20*, 2020.
- [6] A. Alacaoglu, V. Cevher, and S. J. Wright. On the complexity of a practical primal-dual coordinate method, 2022. URL <https://arxiv.org/abs/2201.07684>.
- [7] Z. Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. In *Proc. ACM STOC’17*, 2017.
- [8] E. D. Andersen and K. D. Andersen. The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In *High Performance Optimization*, pages 197–232. Springer, 2000.
- [9] D. Applegate, M. Díaz, O. Hinder, H. Lu, M. Lubin, B. O’Donoghue, and W. Schudy. Practical large-scale linear programming using primal-dual hybrid gradient. *Proc. NeurIPS’21*, 2021.
- [10] D. Applegate, O. Hinder, H. Lu, and M. Lubin. Faster first-order primal-dual methods for linear programming using restarts and sharpness. *arXiv preprint arXiv:2105.12715*, 2021.
- [11] P. Bianchi. Ergodic convergence of a stochastic proximal point algorithm. *SIAM Journal on Optimization*, 26(4):2235–2260, 2016.
- [12] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge university press, 2004.
- [13] Y. Carmon, Y. Jin, A. Sidford, and K. Tian. Variance reduction for matrix games. In *Proc. NeurIPS’19*, 2019.
- [14] Y. Carmon, Y. Jin, A. Sidford, and K. Tian. Coordinate methods for matrix games. In *Proceedings - 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS 2020)*, pages 283–293, November 2020.
- [15] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- [16] A. Chambolle, M. J. Ehrhardt, P. Richtárik, and C.-B. Schonlieb. Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications. *SIAM Journal on Optimization*, 28(4):2783–2808, 2018.
- [17] C.-C. Chang and C.-J. Lin. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):1–27, 2011.
- [18] C. D. Dang and G. Lan. Stochastic block mirror descent methods for nonsmooth and stochastic optimization. *SIAM Journal on Optimization*, 25(2):856–881, 2015.
- [19] D. P. De Farias and B. Van Roy. The linear programming approach to approximate dynamic programming. *Operations Research*, 51(6):850–865, 2003.
- [20] J. Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82(2):421–439, 1956.
- [21] J. Duchi and H. Namkoong. Learning models with uniform performance via distributionally robust optimization. *Annals of Statistics*, 49(3):1378–1406, 2021.

- [22] J. C. Duchi, P. W. Glynn, and H. Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021.
- [23] P. M. Esfahani and D. Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.
- [24] O. Fercoq and P. Bianchi. A coordinate-descent primal-dual algorithm with large step size and possibly nonseparable functions. *SIAM Journal on Optimization*, 29(1):100–134, 2019.
- [25] X. Gao and S.-Z. Zhang. First-order algorithms for convex optimization with nonseparable objective and coupled constraints. *Journal of the Operations Research Society of China*, 5(2): 131–159, 2017.
- [26] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2022. URL <https://www.gurobi.com>.
- [27] E. Y. Hamedani and A. Jalilzadeh. A stochastic variance-reduced accelerated primal-dual method for finite-sum saddle-point problems. *arXiv preprint arXiv:2012.13456*, 2020.
- [28] J. L. Hennessy and D. A. Patterson. *Computer architecture: a quantitative approach*. Elsevier, 2011.
- [29] N. Ho-Nguyen and S. J. Wright. Adversarial classification via distributional robustness with Wasserstein ambiguity. *Mathematical Programming, Series B*, 2022. URL <https://doi.org/10.1007/s10107-022-01796-6>.
- [30] A. J. Hoffman. On approximate solutions of systems of linear inequalities. In *Selected Papers Of Alan J Hoffman: With Commentary*, pages 174–176. World Scientific, 2003.
- [31] W. Hu, G. Niu, I. Sato, and M. Sugiyama. Does distributionally robust supervised learning give robust classifiers? In *Proc. ICML’18*, 2018.
- [32] R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Proc. NIPS’13*, 2013.
- [33] A. Juditsky, A. Nemirovski, and C. Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- [34] G. Korpelevich. The extragradient method for finding saddle points and other problems. *Ekonomika i Matematicheskie Metody*, 12(5):747–756 (in Russian; English translation in *Matekon*), 1976.
- [35] P. Latafat, N. M. Freris, and P. Patrinos. A new randomized block-coordinate primal-dual proximal algorithm for distributed optimization. *IEEE Transactions on Automatic Control*, 64(10):4050–4065, 2019.
- [36] S. Lee and S. J. Wright. Manifold identification in dual averaging for regularized stochastic online learning. *Journal of Machine Learning Research*, 13(6), 2012.
- [37] D. Levy, Y. Carmon, J. C. Duchi, and A. Sidford. Large-scale methods for distributionally robust optimization. In *Proc. NeurIPS’20*, 2020.
- [38] J. Li, S. Huang, and A. M.-C. So. A first-order algorithmic framework for Wasserstein distributionally robust logistic regression. *Proc. NeurIPS’19*, 2019.
- [39] C. Liu, T. Arnon, C. Lazarus, C. Strong, C. Barrett, M. J. Kochenderfer, et al. Algorithms for verifying deep neural networks. *Foundations and Trends® in Optimization*, 4, 2020.
- [40] H. Lu and J. Yang. Nearly optimal linear convergence of stochastic primal-dual methods for linear programming. *arXiv preprint arXiv:2111.05530*, 2021.
- [41] O. Mangasarian. A Newton method for linear programming. *Journal of Optimization Theory and Applications*, 121(1):1–18, 2004.

- [42] O. L. Mangasarian. Normal solutions of linear programs. In *Mathematical Programming at Oberwolfach II*, pages 206–216. Springer, 1984.
- [43] O. L. Mangasarian and R. Meyer. Nonlinear perturbation of linear programs. *SIAM Journal on Control and Optimization*, 17(6):745–752, 1979.
- [44] H. Namkoong and J. C. Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *Proc. NIPS’16*, 2016.
- [45] A. Nemirovski. Prox-method with rate of convergence  $O(1/t)$  for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- [46] Y. Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2-3):319–344, 2007.
- [47] H. Ouyang, N. He, L. Tran, and A. Gray. Stochastic alternating direction method of multipliers. In *Proc. ICML’13*, 2013.
- [48] Y. Ouyang and Y. Xu. Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems. *Mathematical Programming*, pages 1–35, 2019.
- [49] A. Patrascu and I. Necoara. Nonasymptotic convergence of stochastic proximal point methods for constrained convex optimization. *Journal of Machine Learning Research*, 18(1):7204–7245, 2017.
- [50] H. Rahimian and S. Mehrotra. Distributionally robust optimization: A Review. *arXiv preprint arXiv:1908.05659*, 2019.
- [51] R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.
- [52] S. Shafieezadeh Abadeh, P. Mohajerin Esfahani, and D. Kuhn. Distributionally robust logistic regression. In *Proc. NIPS’15*, 2015.
- [53] A. Shapiro. On duality theory of conic linear problems. In *Semi-Infinite Programming*, pages 135–165. Springer, 2001.
- [54] C. Song, Y. Jiang, and Y. Ma. Variance reduction via accelerated dual averaging for finite-sum optimization. In *Proc. NeurIPS’20*, 2020.
- [55] C. Song, S. J. Wright, and J. Diakonikolas. Variance reduction via primal-dual accelerated dual averaging for nonsmooth convex finite-sums. In *Proc. ICML’21*, 2021.
- [56] J. Stoer. Duality in nonlinear programming and the minimax theorem. *Numerische Mathematik*, 5(1):371–379, 1963.
- [57] C. Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.
- [58] W. Wiesemann, D. Kuhn, and M. Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.
- [59] L. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11(Oct):2543–2596, 2010.
- [60] Y. Xu. Primal-dual stochastic gradient method for convex programs with many functional constraints. *SIAM Journal on Optimization*, 30(2):1664–1692, 2020.
- [61] Y. Yu, T. Lin, E. Mazumdar, and M. I. Jordan. Fast distributionally robust learning with variance reduced min-max optimization. *arXiv preprint arXiv:2104.13326*, 2021.
- [62] L. Zhao and D.-L. Zhu. On iteration complexity of a first-order primal-dual method for nonlinear convex cone programming. *Journal of the Operations Research Society of China*, 10(1):53–87, 2022.

- [63] D. Zhu and L. Zhao. Linear convergence of randomized primal-dual coordinate method for large-scale linear constrained convex programming. In *International Conference on Machine Learning*, pages 11619–11628. PMLR, 2020.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] All theoretical claims are supported by proofs, which, due to space constraints, are provided in the appendix. Numerical results are provided in Section 5, with full technical details of the experiments provided in Appendix D.
  - (b) Did you describe the limitations of your work? [Yes] The three main examples are: (i) the discussion of lazy update strategy and guarantees for  $\hat{x}_k$  versus  $\tilde{x}_k$  in Appendix A; (ii) the discussion of linear convergence under the restart strategy and theoretical guarantees for scheduled restart versus empirical guarantees for adaptive restart in Section 3.3; and (iii) comparison to off-the-shelf solvers in Section 5 and use of preconditioning.
  - (c) Did you discuss any potential negative societal impacts of your work? [N/A] This work concerns a generic solver for generalized linear programs and as such is not tied to any specific application that could have a potential negative societal impact.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See assumptions in Section 2.
  - (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix B and C for complete proofs.
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We include our code in the supplementary material.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In Section 5 and Appendix D.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Error bars are not reported due to almost identical behaviors and results from the large number of stochastic iterations.
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In Appendix D.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes] We used datasets from LibSVM, which is an open source machine learning library and cited the creators [17] in Appendix D which provides full experimental details.
  - (b) Did you mention the license of the assets? [N/A] LibSVM datasets are distributed under the BSD license which is “a low restriction type of license for open source software that does not put requirements on redistribution.”
  - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Our code is provided in the supplementary material.
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A] As mentioned before, the datasets we use are open source.
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

# Coordinate Linear Variance Reduction for Generalized Linear Programming

## Appendix

**Outline.** The appendix of this paper is organized as follows:

- Section A contains details and the pseudo code for CLVR with lazy update.
- Section B provides the proofs for Section 3.
- Section C provides the proofs for Section 4.
- Section D provides the complementary details for Section 5.

### A Lazy update for sparse and structured instances of (PD-GLP)

In Algorithm 1, for dense  $\mathbf{A}$ , the  $O(|S^{jk}|d)$  cost of Steps 6 and 8 dominates the  $O(d)$  cost of Steps 4 and 9. However, when  $\mathbf{A}$  is sparse and  $|S^{jk}|$  is small, the cost of Steps 6 and 8 will be  $O(\text{nnz}(\mathbf{A}^{S^{jk}}))$ , which may be less than the  $O(d)$  cost of Steps 4 and 9. Using this observation, we show that the nature of the dual averaging update enables us to propose an efficient implementation whose complexity depends on  $\text{nnz}(\mathbf{A})$  rather than  $n \cdot d$ .

Recall that we partition  $[n]$  into subsets  $\{S^1, S^2, \dots, S^m\}$  and use  $\mathbf{A}^{S^j}$  ( $j \in [m]$ ) to denote the  $j^{\text{th}}$  row block of  $\mathbf{A}$ . For each block  $\mathbf{A}^{S^j}$ , we use  $C^j \subset [d]$  to denote the indices of those columns of  $\mathbf{A}^{S^j}$  that contain at least one nonzero element. (Of course,  $\{C^1, \dots, C^m\}$  is not in general a partition of  $[d]$  as different subsets  $S^j$  may have nonzeros in the same columns.) We assume further that  $\mathcal{X}$  and  $r$  are coordinate separable, that is,  $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_d$  with  $\mathcal{X}_i \subset \mathbb{R}$  for all  $i$  and  $r(\mathbf{x}) := \sum_{i=1}^d r(\mathbf{x}^i)$  with  $\mathbf{x}^i \in \mathcal{X}_i$ .

In Step 4 of Algorithm 1, the separability of  $\mathcal{X}$  and  $r$  means that an update to one coordinate block of  $\mathbf{x}$  — say the  $C^{jk}$  block in the update from  $\mathbf{x}_{k-1}$  to  $\mathbf{x}_k$ , which requires a projection and an application of the proximal operator — does not influence other coordinates  $\mathbf{x}_k^i$  for  $i \notin C^{jk}$ . Moreover, we can efficiently maintain an implicit representation of  $\mathbf{q}_k$ , in terms of a newly introduced auxiliary vector  $\mathbf{r}_k$ ; see Lemma 8 of Appendix B.2. Similarly, to update  $\tilde{\mathbf{y}}_k$  efficiently, we maintain an implicit representation of  $\tilde{\mathbf{y}}_k$  via an auxiliary vector  $\mathbf{s}_k$ , as shown in Lemma 9 of Appendix B.2.

It is generally not possible to maintain  $\tilde{\mathbf{x}}_k$  efficiently since, in principle, all components of  $\mathbf{x}_k$  can change on every iteration, and we wish to avoid the  $O(d)$  cost of evaluating every full  $\mathbf{x}_k$ . We seek instead to output a proxy  $\hat{\mathbf{x}}_K$  for  $\tilde{\mathbf{x}}_K$  from Algorithm 1 such that  $\mathbb{E}[\hat{\mathbf{x}}_K] = \tilde{\mathbf{x}}_K$ . One possible choice is to pick an index  $k' \in [K]$  from the weighted discrete distribution  $(\frac{a_1}{A_K}, \frac{a_2}{A_K}, \dots, \frac{a_K}{A_K})$  (computing the scalar quantities  $a_k$  and  $A_k$  for  $k \in [K]$  in advance), then setting  $\hat{\mathbf{x}}_K = \mathbf{x}_{k'}$ . A slightly more sophisticated strategy is to sample a predetermined number  $\hat{K}$  of vectors  $\mathbf{x}_k$ ,  $k \in [K]$ , and define  $\hat{\mathbf{x}}_K$  to be the simple average of these vectors. Once again, the indices are chosen from the weighted discrete distribution  $(\frac{a_1}{A_K}, \frac{a_2}{A_K}, \dots, \frac{a_K}{A_K})$ . Note that the total expected cost of the  $K$  iterations of the algorithm (excluding the full-vector updates) is  $O(K \text{nnz}(\mathbf{A})/m)$ , while the total cost of evaluating the  $\hat{K}$  full vectors  $\mathbf{x}_k$  and accumulating them into  $\hat{\mathbf{x}}_K$  is  $O(\hat{K}d)$ . Thus, for the full-step iterations not to dominate the total cost, we can choose  $\hat{K}$  to be  $O(K \text{nnz}(\mathbf{A})/(md))$ . Finally, we note that Theorem 1 is for  $\tilde{\mathbf{x}}_k$ , not  $\hat{\mathbf{x}}_K$ . However, since  $\mathbb{E}[\hat{\mathbf{x}}_k] = \tilde{\mathbf{x}}_k$ , we expect the convergence rates from the theorem to hold for  $\hat{\mathbf{x}}_K$  in practice.

An implementation of Algorithm 1 that exploits the form of  $\mathbf{q}_k$  in Lemma 8 and  $\tilde{\mathbf{y}}_k$  in Lemma 9, and evaluates explicitly only those components of  $\mathbf{x}_k$  needed to perform the rest of the iteration is given as Algorithm 2. This version also incorporates the strategy for obtaining  $\hat{\mathbf{x}}_K$  by sampling  $\hat{K}$  iterates on which to evaluate the full vector  $\mathbf{x}_k$ .

Due to the efficient implementation in Algorithm 2, to attain an  $\epsilon$ -accurate solution in terms of the primal-dual gap in Theorem 1, we need  $O(\frac{\text{nnz}(\mathbf{A})\hat{L}}{\epsilon})$  FLOPS, which corresponds to  $O(\frac{\hat{L}}{\epsilon})$  data passes. As a result, because a smaller batch size leads to a smaller  $\hat{L}$ , we attain the best performance in

terms of number of data passes when the batch size is set to one. However, as modern computer architecture has particular design for vectorized operations, lower runtime is obtained for a small batch size strictly larger than one (see Section 5).

**Remark 1.** While we consider the case of fully coordinate separable  $r$  and  $\mathcal{X}$  for simplicity, our lazy update approach is also applicable to the coordinate block partitioning case in which  $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_m$  with  $\mathcal{X}_i \in \mathbb{R}^{d_i}$  ( $i \in [m]$ ),  $\sum_{i=1}^m d_i = d$  and  $r(\mathbf{x}) := \sum_{i=1}^m r(\mathbf{x}^i)$  with  $\mathbf{x}^i \in \mathcal{X}_i$ . The difference is that for each coordinate block  $\mathbf{A}^{S^j}$  ( $j \in [m]$ ), we overload  $C^j \subset [m]$  to denote the set of blocks in  $\mathbf{A}^{S^j}$  where each coordinate block contains at least one nonzero element.

**Remark 2.** Dual averaging has been shown to have significant advantage in producing sparser iterates than mirror descent in the context of online learning [36, 59]. It further leads to better bounds in well-conditioned finite-sum optimization [54]. In this work, we show that dual averaging offers better flexibility with sparse matrices than does mirror descent.

---

**Algorithm 2** Coordinate Linear Variance Reduction with Lazy Update (Lazy CLVR)

---

- 1: **Input:**  $\tilde{\mathbf{x}}_0 = \mathbf{x}_0 = \mathbf{x}_{-1} \in \mathcal{X}$ ,  $\mathbf{y}_0 \in \mathbb{R}^n$ ,  $\mathbf{z}_0 = \mathbf{A}^T \mathbf{y}_0$ ,  $\gamma > 0$ ,  $\hat{L} > 0$ ,  $K$ ,  $\hat{K}$ ,  $m$ ,  $\{S^1, S^2, \dots, S^m\}$ ,  $\{C^1, C^2, \dots, C^m\}$ .
  - 2:  $a_0 = A_0 = 0$ ,  $a_1 = A_1 = \frac{1}{2\hat{L}m}$ ,  $\mathbf{q}_0 = a_1(\mathbf{z}_0 + \mathbf{c})$ ,  $\mathbf{r}_0 = \mathbf{0}_d$ ,  $\mathbf{s}_0 = \mathbf{0}_n$ .
  - 3: **for**  $k = 1, 2, \dots, K - 1$  **do**
  - 4:  $a_{k+1} = \frac{\sqrt{1 + \sigma A_k / \gamma}}{2\hat{L}m}$ ,  $A_{k+1} = A_k + a_{k+1}$ .
  - 5: **end for**
  - 6: Choose indices  $\{\ell_1, \dots, \ell_{\hat{K}}\}$  i.i.d. from  $[K]$  according to the distribution  $\left\{ \frac{a_1}{A_K}, \frac{a_2}{A_K}, \dots, \frac{a_K}{A_K} \right\}$ .
  - 7: **for**  $k = 1, 2, \dots, K$  **do**
  - 8: Pick  $j_k$  uniformly at random in  $[m]$ .
  - 9: **if**  $k = \ell_i$  for some  $i = 1, 2, \dots, \hat{K}$  **then**
  - 10:  $\mathbf{q}_{k-1} = A_k(\mathbf{c} + \mathbf{z}_{k-1}) + \mathbf{r}_{k-1}$ .
  - 11:  $\mathbf{x}_k = \text{prox}_{\frac{1}{\gamma} A_k r}(\mathbf{x}_0 - \frac{1}{\gamma} \mathbf{q}_{k-1})$ .
  - 12: **else**
  - 13:  $\mathbf{q}_{k-1}^{C^{j_k}} = A_k(\mathbf{c}^{C^{j_k}} + \mathbf{z}_{k-1}^{C^{j_k}}) + \mathbf{r}_{k-1}^{C^{j_k}}$ .
  - 14:  $\mathbf{x}_k^{C^{j_k}} = \text{prox}_{\frac{1}{\gamma} A_k r}(\mathbf{x}_0^{C^{j_k}} - \frac{1}{\gamma} \mathbf{q}_{k-1}^{C^{j_k}})$ .
  - 15: **end if**
  - 16:  $\mathbf{y}_k^{S^{j_k}} = \mathbf{y}_{k-1}^{S^{j_k}} + \gamma m a_k (\mathbf{A}^{S^{j_k}, C^{j_k}} \mathbf{x}_k^{C^{j_k}} - \mathbf{b}^{S^{j_k}})$ ,  $\mathbf{y}_{k-1}^{S^i}$  for all  $i \neq j_k$ .
  - 17:  $\mathbf{z}_k^{C^{j_k}} = \mathbf{z}_{k-1}^{C^{j_k}} + (\mathbf{A}^{S^{j_k}, C^{j_k}})^T (\mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}})$ ,  $\mathbf{z}_k^i = \mathbf{z}_{k-1}^i$  for all  $i \notin C^{j_k}$ ;
  - 18:  $\mathbf{r}_k^{C^{j_k}} = \mathbf{r}_{k-1}^{C^{j_k}} + (m a_k - A_k)(\mathbf{z}_k^{C^{j_k}} - \mathbf{z}_{k-1}^{C^{j_k}})$ ,  $\mathbf{r}_k^i = \mathbf{r}_{k-1}^i$  for all  $i \notin C^{j_k}$ ;
  - 19:  $\mathbf{s}_k^{S^{j_k}} = \mathbf{s}_{k-1}^{S^{j_k}} + ((m-1)a_k - A_{k-1})(\mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}})$ ,  $\mathbf{s}_k^i = \mathbf{s}_{k-1}^i$  for all  $i \notin S^{j_k}$ ;
  - 20: **end for**
  - 21:  $\hat{\mathbf{x}}_K = \frac{1}{K} \sum_{i=1}^{\hat{K}} \mathbf{x}_{\ell_i}$ .
  - 22:  $\tilde{\mathbf{y}}_K = \mathbf{y}_K + \frac{1}{A_K} \mathbf{s}_K$ .
  - 23: **return**  $\hat{\mathbf{x}}_K$  and  $\tilde{\mathbf{y}}_K$ .
- 

## B Omitted proofs from Section 3

### B.1 Omitted proofs from Section 3.1

We state a version of CLVR in Algorithm 3 that is convenient for the analysis, and is equivalent to Algorithm 1 in the main body of the paper. Before analyzing the convergence of CLVR, we justify our claim of equivalence in Proposition 1.

**Proposition 1.** *The iterates of Algorithm 1 and 3 are equivalent.*

*Proof.* To argue equivalence, we show that the iterates of Algorithm 1 and 3 solve the same optimization problems. To avoid ambiguity, here we will use  $\hat{\mathbf{x}}_k, \hat{\mathbf{y}}_k$  to denote the iterates  $\mathbf{x}_k, \mathbf{y}_k$  in Algorithm 3, while we retain the notation  $\mathbf{x}_k, \mathbf{y}_k$  for the iterates of Algorithm 1.

Let us first start by writing an equivalent definition of  $\mathbf{x}_k$  in Algorithm 1. To do so, we first unroll the recursive definitions of  $\mathbf{z}_k$  and  $\mathbf{q}_k$ . We can observe that, since by definition  $\mathbf{y}_k$  and  $\mathbf{y}_{k-1}$  only differ over the coordinate block  $S^{j_k}$ , we have

$$\mathbf{z}_k = \mathbf{z}_0 + \sum_{i=1}^k \mathbf{A}^T (\mathbf{y}_i - \mathbf{y}_{i-1}) = \mathbf{A}^T \mathbf{y}_k. \quad (12)$$

On the other hand, using Eq. (12), the definition of  $\mathbf{q}_k$  implies

$$\mathbf{q}_k = A_{k+1} \mathbf{c} + a_1 \mathbf{A}^T \mathbf{y}_0 + \sum_{i=1}^k \mathbf{A}^T [a_{i+1} \mathbf{y}_i + m a_i (\mathbf{y}_i - \mathbf{y}_{i-1})] \quad (13)$$

Using the definition of the proximal operator (see Eq. (1)) and the definition of  $\mathbf{x}_k$  in Step 4 of Algorithm 1, we have

$$\begin{aligned} \mathbf{x}_k &= \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{A_k}{\gamma} r(\mathbf{x}) + \frac{1}{2} \left\| \mathbf{x} - \mathbf{x}_0 + \frac{1}{\gamma} \mathbf{q}_{k-1} \right\|^2 \right\} \\ &= \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ A_k r(\mathbf{x}) + \frac{\gamma}{2} \left\| \mathbf{x} - \mathbf{x}_0 + \frac{1}{\gamma} \mathbf{q}_{k-1} \right\|^2 \right\}. \end{aligned} \quad (14)$$

Now let us consider the optimization problem that defines  $\hat{\mathbf{x}}_k$ . Assume for now that the definitions of  $\mathbf{y}_k$  and  $\hat{\mathbf{y}}_k$  agree (we justify this claim below). Observe first that the minimization problem defining  $\hat{\mathbf{x}}_k$  is independent of  $\mathbf{u}$ , so by unrolling the recursion for  $\phi_k$ , we have

$$\begin{aligned} \hat{\mathbf{x}}_k &= \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{\gamma}{2} \|\mathbf{x} - \mathbf{x}_0\|^2 + A_k r(\mathbf{x}) + a_1 \langle \mathbf{x}, \mathbf{c} + \mathbf{A}^T \bar{\mathbf{y}}_0 \rangle \right. \\ &\quad \left. + \sum_{i=2}^k a_i \left\langle \mathbf{x}, \mathbf{c} + \mathbf{A}^T \left( \mathbf{y}_{i-1} + \frac{m a_{i-1}}{a_i} (\mathbf{y}_{i-1} - \mathbf{y}_{i-2}) \right) \right\rangle \right\} \\ &= \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{\gamma}{2} \|\mathbf{x} - \mathbf{x}_0\|^2 + A_k r(\mathbf{x}) \right. \\ &\quad \left. + \left\langle \mathbf{x}, A_k \mathbf{c} + a_1 \mathbf{A}^T \mathbf{y}_0 + \mathbf{A}^T \left( \sum_{i=2}^k [a_i \mathbf{y}_{i-1} + m a_{i-1} (\mathbf{y}_{i-1} - \mathbf{y}_{i-2})] \right) \right\rangle \right\} \\ &= \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{\gamma}{2} \|\mathbf{x} - \mathbf{x}_0\|^2 + A_k r(\mathbf{x}) + \langle \mathbf{x}, \mathbf{q}_{k-1} \rangle \right\} \\ &= \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ A_k r(\mathbf{x}) + \frac{\gamma}{2} \left\| \mathbf{x} - \mathbf{x}_0 + \frac{1}{\gamma} \mathbf{q}_{k-1} \right\|^2 \right\} \\ &= \mathbf{x}_k. \end{aligned}$$

It remains to argue that the definitions of  $\mathbf{y}_k$  and  $\hat{\mathbf{y}}_k$  agree. First, observe that since the definitions of  $\psi_k$  and  $\psi_{k-1}$  differ only over block  $S^{j_k}$ , we have  $\hat{\mathbf{y}}_k^{S^j} = \hat{\mathbf{y}}_{k-1}^{S^j}$  for all  $j \neq j_k$ . For  $j = j_k$ , we have by unrolling the recursive definition of  $\psi_k$  that

$$\begin{aligned} \hat{\mathbf{y}}_k^{S^{j_k}} &= \arg \min_{\mathbf{y}^{S^{j_k}} \in \mathbb{R}^{n_{j_k}}} \left\{ \frac{1}{2\gamma} \|\mathbf{y}^{S^{j_k}} - \mathbf{y}_0^{S^{j_k}}\|^2 - \sum_{i=1}^k \mathbb{1}_{\{S^{j_i} = S^{j_k}\}} m a_i \left\langle \mathbf{y}^{S^{j_i}}, \mathbf{A}^{S^{j_i}} \mathbf{x}_i - \mathbf{b}^{S^{j_i}} \right\rangle \right\} \\ &= \mathbf{y}_0^{S^{j_k}} + \gamma \sum_{i=1}^k \mathbb{1}_{\{S^{j_i} = S^{j_k}\}} m a_i (\mathbf{A}^{S^{j_i}} \mathbf{x}_i - \mathbf{b}^{S^{j_i}}) \\ &= \hat{\mathbf{y}}_{k-1}^{S^{j_k}} + \gamma m a_k (\mathbf{A}^{S^{j_k}} \mathbf{x}_k - \mathbf{b}^{S^{j_k}}) \\ &= \mathbf{y}_k^{S^{j_k}}, \end{aligned}$$

as claimed.  $\square$

In the following three lemmas, we bound  $\phi_k(\mathbf{x}_k)$  and  $\psi_k(\mathbf{y}_k)$  below and above, which is then subsequently used to bound the primal-dual gap in Theorem 1.

---

**Algorithm 3** Coordinate Linear Variance Reduction (Analysis Version)

---

- 1: **Input:**  $\mathbf{x}_0 = \mathbf{x}_{-1} \in \mathcal{X}$ ,  $\mathbf{y}_0 = \bar{\mathbf{y}}_0 \in \mathbb{R}^n$ ,  $m, \{S^1, S^2, \dots, S^m\}, K, \gamma > 0, \hat{L} > 0$ .
  - 2:  $\phi_0(\cdot) = \frac{\gamma}{2} \|\cdot - \mathbf{x}_0\|^2$ ,  $\psi_0(\cdot) = \frac{\gamma}{2\gamma} \|\cdot - \mathbf{y}_0\|^2$ .
  - 3:  $a_1 = A_1 = \frac{1}{2\hat{L}m}$ .
  - 4: **for**  $k = 1, 2, 3, \dots, K$  **do**
  - 5:    $\mathbf{x}_k = \arg \min_{\mathbf{x} \in \mathcal{X}} \{\phi_k(\mathbf{x}) = \phi_{k-1}(\mathbf{x}) + a_k(\langle \mathbf{x} - \mathbf{u}, \mathbf{A}^T \bar{\mathbf{y}}_{k-1} + \mathbf{c} \rangle + r(\mathbf{x}) - r(\mathbf{u}))\}$ .
  - 6:   Pick  $j_k$  uniformly at random in  $[m]$ .
  - 7:    $\mathbf{y}_k = \arg \min_{\mathbf{y} \in \mathbb{R}^n} \{\psi_k(\mathbf{y}) = \psi_{k-1}(\mathbf{y}) + ma_k(-\langle \mathbf{y}^{S^{j_k}} - \mathbf{v}^{S^{j_k}}, \mathbf{A}^{S^{j_k}} \mathbf{x}_k - \mathbf{b}^{S^{j_k}} \rangle)\}$ .
  - 8:    $a_{k+1} = \frac{\sqrt{1 + \sigma A_k / \gamma}}{2\hat{L}m}$ ,  $A_{k+1} = A_k + a_{k+1}$ .
  - 9:    $\bar{\mathbf{y}}_k = \mathbf{y}_k + \frac{ma_k}{a_{k+1}}(\mathbf{y}_k - \mathbf{y}_{k-1})$ .
  - 10: **end for**
  - 11: **return**  $\tilde{\mathbf{x}}_K := \frac{1}{A_K} \sum_{k=1}^K a_k \mathbf{x}_k$ ,  $\tilde{\mathbf{y}}_K = \frac{1}{A_K} \sum_{k=1}^K (a_k \mathbf{y}_k + (m-1)a_k(\mathbf{y}_k - \mathbf{y}_{k-1}))$ .
- 

**Lemma 1.** For all steps of Algorithm 3 with  $k \geq 1$ , we have,  $\forall(\mathbf{u}, \mathbf{v}) \in \mathcal{X} \times \mathbb{R}^n$ ,

$$\begin{aligned} \phi_k(\mathbf{x}_k) &\leq \frac{\gamma}{2} \|\mathbf{u} - \mathbf{x}_0\|^2 - \frac{\gamma + \sigma A_k}{2} \|\mathbf{u} - \mathbf{x}_k\|^2, \\ \psi_k(\mathbf{y}_k) &\leq \frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_0\|^2 - \frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_k\|^2. \end{aligned}$$

*Proof.* By the definitions of  $\psi_k(\mathbf{y})$  and  $\phi_k(\mathbf{x})$  in Algorithm 3, it follows that,  $\forall k \geq 1$ ,

$$\phi_k(\mathbf{x}) = \sum_{i=1}^k a_i (\langle \mathbf{x} - \mathbf{u}, \mathbf{A}^T \bar{\mathbf{y}}_{i-1} + \mathbf{c} \rangle + r(\mathbf{x}) - r(\mathbf{u})) + \frac{\gamma}{2} \|\mathbf{x} - \mathbf{x}_0\|^2, \quad (15)$$

and

$$\psi_k(\mathbf{y}) = \sum_{i=1}^k ma_i \left( -\langle \mathbf{y}^{S^{j_i}} - \mathbf{v}^{S^{j_i}}, \mathbf{A}^{S^{j_i}} \mathbf{x}_i - \mathbf{b}^{S^{j_i}} \rangle \right) + \frac{1}{2\gamma} \|\mathbf{y} - \mathbf{y}_0\|^2. \quad (16)$$

Observe that, as function of  $\mathbf{x}$ ,  $\phi_k(\mathbf{x})$  is  $(\gamma + \sigma A_k)$ -strongly convex. As, by definition,  $\mathbf{x}_k = \arg \min_{\mathbf{x} \in \mathcal{X}} \phi_k(\mathbf{x})$ , it follows that

$$\phi_k(\mathbf{u}) \geq \phi_k(\mathbf{x}_k) + \frac{\gamma + \sigma A_k}{2} \|\mathbf{u} - \mathbf{x}_k\|^2. \quad (17)$$

Now, writing  $\phi_k(\mathbf{u})$  explicitly and rearranging the last inequality, the stated bound on  $\phi_k(\mathbf{x}_k)$  follows.

As a function of  $\mathbf{y}$ ,  $\psi_k(\mathbf{y})$  is  $1/\gamma$ -strongly convex. The proof for the bound on  $\psi_k$  uses the same argument and is omitted.  $\square$

In the following proof, for  $k \geq 1$ , let  $\mathcal{F}_k$  denote the natural filtration, containing all the randomness in the algorithm up to and including iteration  $k$ . Recall that  $\mathbf{A} = \begin{pmatrix} \mathbf{A}^{S^1} \\ \vdots \\ \mathbf{A}^{S^m} \end{pmatrix}$ , and let  $\mathbf{A}^{S^j}$  ( $j \in [m]$ )

denote the matrix  $\mathbf{A}$  with its  $S^j$  block of rows replaced by a zero block.

For convenience, for  $k = 1, 2, \dots$ , we define

$$\hat{\mathbf{y}}_k = \mathbf{y}_{k-1} + \gamma ma_k (\mathbf{A} \mathbf{x}_k - \mathbf{b}). \quad (18)$$

Then from the definition of  $\mathbf{y}_k$  in Algorithm 1, we have  $\mathbb{E}[\mathbf{y}_k - \mathbf{y}_{k-1} | \mathcal{F}_{k-1}] = \frac{1}{m} (\hat{\mathbf{y}}_k - \mathbf{y}_{k-1})$ .

Motivated by [4], we have the following result.

**Lemma 2.** Given the sequences  $\{\mathbf{y}_k\}$  in Algorithm 3 and  $\{\hat{\mathbf{y}}_k\}$  in Eq. (18), we define the sequence  $\{\check{\mathbf{v}}_k\}$  by

$$\check{\mathbf{v}}_k = (\mathbf{y}_k - \mathbf{y}_{k-1}) - \frac{1}{m} (\hat{\mathbf{y}}_k - \mathbf{y}_{k-1}).$$

Then for any  $\mathbf{v} \in \mathbb{R}^d$  that may be correlated with the randomness in  $\{\check{\mathbf{v}}_i\}_{i=1}^k$ , we have

$$\mathbb{E}\left[-\sum_{i=1}^k \langle \mathbf{v}, \check{\mathbf{v}}_i \rangle\right] \leq \mathbb{E}\left[\frac{1}{2}\|\mathbf{y}_0 - \mathbf{v}\|^2 + \frac{1}{2}\sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2\right], \quad (19)$$

where the expectation is taken over all the randomness in the history.

*Proof.* First, we prove a bound for  $\mathbb{E}\left[\sum_{i=1}^k \|\check{\mathbf{v}}_i\|^2\right]$ . By the fact that  $\mathbb{E}[\mathbf{y}_i - \mathbf{y}_{i-1} | \mathcal{F}_{i-1}] = \frac{1}{m}(\hat{\mathbf{y}}_i - \mathbf{y}_{i-1})$ , for each  $\check{\mathbf{v}}_i$ , we have

$$\mathbb{E}[\|\check{\mathbf{v}}_i\|^2 | \mathcal{F}_{i-1}] = \mathbb{E}[\|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 | \mathcal{F}_{i-1}] - \mathbb{E}\left[\left\|\frac{1}{m}(\hat{\mathbf{y}}_i - \mathbf{y}_{i-1})\right\|^2 | \mathcal{F}_{i-1}\right] \leq \mathbb{E}[\|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 | \mathcal{F}_{i-1}].$$

Taking expectation on all the randomness in the history, for  $\mathbb{E}\left[\sum_{i=1}^k \|\check{\mathbf{v}}_i\|^2\right]$ , we have

$$\mathbb{E}\left[\sum_{i=1}^k \|\check{\mathbf{v}}_i\|^2\right] = \mathbb{E}\left[\sum_{i=1}^k \mathbb{E}[\|\check{\mathbf{v}}_i\|^2 | \mathcal{F}_{i-1}]\right] \leq \mathbb{E}\left[\sum_{i=1}^k \mathbb{E}[\|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 | \mathcal{F}_{i-1}]\right] = \mathbb{E}\left[\sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2\right]. \quad (20)$$

Then to prove (19), we define the sequence  $\{\check{\mathbf{y}}_k\}_{k=0}^\infty$  by  $\check{\mathbf{y}}_0 = \mathbf{y}_0$  and  $\check{\mathbf{y}}_k = \check{\mathbf{y}}_{k-1} - \check{\mathbf{v}}_k$  for  $k = 1, 2, \dots$ . By expanding  $\frac{1}{2}\|\check{\mathbf{y}}_k - \mathbf{v}\|^2$ ,

$$\begin{aligned} \frac{1}{2}\|\check{\mathbf{y}}_k - \mathbf{v}\|^2 &= \frac{1}{2}\|\check{\mathbf{y}}_{k-1} - \mathbf{v}\|^2 - \langle \check{\mathbf{v}}_k, \check{\mathbf{y}}_{k-1} - \mathbf{v} \rangle + \frac{1}{2}\|\check{\mathbf{v}}_k\|^2 \\ \implies \langle \check{\mathbf{v}}_k, \check{\mathbf{y}}_{k-1} - \mathbf{v} \rangle &= \frac{1}{2}\|\check{\mathbf{y}}_{k-1} - \mathbf{v}\|^2 - \frac{1}{2}\|\check{\mathbf{y}}_k - \mathbf{v}\|^2 + \frac{1}{2}\|\check{\mathbf{v}}_k\|^2. \end{aligned}$$

By summing this expression and telescoping, we obtain

$$\begin{aligned} \sum_{i=1}^k \langle \check{\mathbf{y}}_{i-1} - \mathbf{v}, \check{\mathbf{v}}_i \rangle &= \frac{1}{2}\|\check{\mathbf{y}}_0 - \mathbf{v}\|^2 - \frac{1}{2}\|\check{\mathbf{y}}_k - \mathbf{v}\|^2 + \sum_{i=1}^k \frac{1}{2}\|\check{\mathbf{v}}_i\|^2 \\ &\leq \frac{1}{2}\|\check{\mathbf{y}}_0 - \mathbf{v}\|^2 + \sum_{i=1}^k \frac{1}{2}\|\check{\mathbf{v}}_i\|^2 \\ &= \frac{1}{2}\|\mathbf{y}_0 - \mathbf{v}\|^2 + \sum_{i=1}^k \frac{1}{2}\|\check{\mathbf{v}}_i\|^2, \end{aligned} \quad (21)$$

It follows that

$$\begin{aligned} \mathbb{E}\left[-\sum_{i=1}^k \langle \mathbf{v}, \check{\mathbf{v}}_i \rangle\right] &= \mathbb{E}\left[\sum_{i=1}^k \langle \check{\mathbf{y}}_{i-1} - \mathbf{v}, \check{\mathbf{v}}_i \rangle - \sum_{i=1}^k \langle \check{\mathbf{y}}_{i-1}, \check{\mathbf{v}}_i \rangle\right] \\ &\leq \mathbb{E}\left[\frac{1}{2}\|\mathbf{y}_0 - \mathbf{v}\|^2 + \sum_{i=1}^k \frac{1}{2}\|\check{\mathbf{v}}_i\|^2 - \sum_{i=1}^k \langle \check{\mathbf{y}}_{i-1}, \check{\mathbf{v}}_i \rangle\right] \\ &\leq \mathbb{E}\left[\frac{1}{2}\|\mathbf{y}_0 - \mathbf{v}\|^2 + \frac{1}{2}\sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 - \sum_{i=1}^k \langle \check{\mathbf{y}}_{i-1}, \check{\mathbf{v}}_i \rangle\right]. \end{aligned} \quad (22)$$

where the first inequality is by Eq. (21) and the second inequality is by Eq. (20). To obtain the result (19), we use the facts that  $\mathbb{E}[\check{\mathbf{v}}_i | \mathcal{F}_{i-1}] = 0$  and that  $\check{\mathbf{y}}_{i-1} \in \mathcal{F}_{i-1}$  to obtain  $\mathbb{E}\left[\sum_{i=1}^k \langle \check{\mathbf{y}}_{i-1}, \check{\mathbf{v}}_i \rangle\right] = 0$ .  $\square$

We emphasize that Lemma 2 holds for any  $\mathbf{v}$  that may be even correlated with the randomness in the algorithm.

**Lemma 3.** For all steps of Algorithm 3 with  $k \geq 1$ , we have for all  $(\mathbf{u}, \mathbf{v}) \in \mathcal{X} \times \mathbb{R}^n$  that

$$\begin{aligned}\phi_k(\mathbf{x}_k) &\geq \phi_{k-1}(\mathbf{x}_{k-1}) + \frac{\gamma + \sigma A_{k-1}}{2} \|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 + a_k(r(\mathbf{x}_k) - r(\mathbf{u})) \\ &\quad - a_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle + a_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T \mathbf{y}_k + \mathbf{c} \rangle \\ &\quad + ma_{k-1} \left( \langle \mathbf{x}_{k-1} - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle + \langle \mathbf{x}_k - \mathbf{x}_{k-1}, \mathbf{A}^T(\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle \right), \\ \psi_k(\mathbf{y}_k) &\geq \psi_{k-1}(\mathbf{y}_{k-1}) + \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 - a_k \langle \mathbf{A} \mathbf{x}_k - \mathbf{b}, \mathbf{y}_k - \mathbf{v} \rangle \\ &\quad - (m-1)a_k \langle \mathbf{A} \mathbf{x}_k - \mathbf{b}, \mathbf{y}_k - \mathbf{y}_{k-1} \rangle + \frac{1}{\gamma} \left\langle \mathbf{y}_{k-1} - \mathbf{v}, -(\mathbf{y}_k - \mathbf{y}_{k-1}) + \frac{1}{m}(\hat{\mathbf{y}}_k - \mathbf{y}_{k-1}) \right\rangle.\end{aligned}$$

*Proof.* For the first claim, we have from the definition of  $\phi_k(\mathbf{x}_k)$ , using that  $\phi_{k-1}(\mathbf{x}_{k-1})$  is  $(\gamma + \sigma A_{k-1})$ -strongly convex and minimized at  $\mathbf{x}_{k-1}$ , that

$$\begin{aligned}\phi_k(\mathbf{x}_k) &\geq \phi_{k-1}(\mathbf{x}_{k-1}) + \frac{\gamma + \sigma A_{k-1}}{2} \|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 \\ &\quad + a_k \left( \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T \bar{\mathbf{y}}_{k-1} + \mathbf{c} \rangle + r(\mathbf{x}_k) - r(\mathbf{u}) \right).\end{aligned}\tag{23}$$

Meanwhile, by the definition of  $\{\bar{\mathbf{y}}_k\}$  (using  $\mathbf{y}_{-1} = \mathbf{y}_0$  for the case of  $k = 1$ ), we have that

$$\begin{aligned}&a_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T \bar{\mathbf{y}}_{k-1} \rangle \\ &= a_k \left\langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T \left( \mathbf{y}_{k-1} + \frac{ma_{k-1}}{a_k} (\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \right) \right\rangle \\ &= -a_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle + a_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T \mathbf{y}_k \rangle \\ &\quad + ma_{k-1} \left( \langle \mathbf{x}_{k-1} - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle + \langle \mathbf{x}_k - \mathbf{x}_{k-1}, \mathbf{A}^T(\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle \right).\end{aligned}\tag{24}$$

The claimed lower bound on  $\phi_k(\mathbf{x}_k)$  follows when we combine (23) and (24).

For the second claim, we have by the definition of  $\psi_k$  that

$$\psi_k(\mathbf{y}_k) - \psi_{k-1}(\mathbf{y}_{k-1}) = \psi_{k-1}(\mathbf{y}_k) - \psi_{k-1}(\mathbf{y}_{k-1}) - ma_k \left\langle \mathbf{y}_k^{S^{j_k}} - \mathbf{v}^{S^{j_k}}, \mathbf{A}^{S^{j_k}} \mathbf{x}_k - \mathbf{b}^{S^{j_k}} \right\rangle.\tag{25}$$

To obtain the claimed lower bound on  $\psi_k$ , we proceed to bound the terms on the right-hand side in (25). First, since  $\psi_{k-1}$  is  $(1/\gamma)$ -strongly convex and minimized at  $\mathbf{y}_{k-1}$ , we have

$$\psi_{k-1}(\mathbf{y}_k) - \psi_{k-1}(\mathbf{y}_{k-1}) \geq \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2.\tag{26}$$

Second, by using the definition of  $\mathbf{A}^{S^{j_k}}$  and  $\mathbf{A}^{\bar{S}^{j_k}}$ , and by using several times that  $\mathbf{y}_{k-1}$  and  $\mathbf{y}_k$  differ only in their  $S^{j_k}$  components, we have

$$\begin{aligned}
& - \left\langle \mathbf{y}_k^{S^{j_k}} - \mathbf{v}^{S^{j_k}}, \mathbf{A}^{S^{j_k}} \mathbf{x}_k - \mathbf{b}^{S^{j_k}} \right\rangle \\
& = - \langle \mathbf{y}_k - \mathbf{v}, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle + \left\langle \mathbf{y}_k - \mathbf{v}, \mathbf{A}^{\bar{S}^{j_k}} \mathbf{x}_k - \mathbf{b}^{\bar{S}^{j_k}} \right\rangle \\
& = - \langle \mathbf{y}_k - \mathbf{v}, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle + \left\langle \mathbf{y}_{k-1} - \mathbf{v}, \mathbf{A}^{\bar{S}^{j_k}} \mathbf{x}_k - \mathbf{b}^{\bar{S}^{j_k}} \right\rangle \\
& = - \langle \mathbf{y}_k - \mathbf{v}, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle + \frac{m-1}{m} \langle \mathbf{y}_{k-1} - \mathbf{v}, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle \\
& \quad + \left\langle \mathbf{y}_{k-1} - \mathbf{v}, \mathbf{A}^{\bar{S}^{j_k}} \mathbf{x}_k - \mathbf{b}^{\bar{S}^{j_k}} - \frac{m-1}{m} (\mathbf{A} \mathbf{x}_k - \mathbf{b}) \right\rangle \\
& = - \frac{1}{m} \langle \mathbf{y}_k - \mathbf{v}, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle + \frac{m-1}{m} \langle \mathbf{y}_{k-1} - \mathbf{y}_k, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle \\
& \quad + \left\langle \mathbf{y}_{k-1} - \mathbf{v}^{S^{j_k}}, -(\mathbf{A}^{S^{j_k}} \mathbf{x}_k - \mathbf{b}^{S^{j_k}}) \right\rangle + \frac{1}{m} \langle \mathbf{y}_{k-1} - \mathbf{v}, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle \\
& = - \frac{1}{m} \langle \mathbf{y}_k - \mathbf{v}, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle + \frac{m-1}{m} \langle \mathbf{y}_{k-1} - \mathbf{y}_k, \mathbf{A} \mathbf{x}_k - \mathbf{b} \rangle \\
& \quad + \left\langle \mathbf{y}_{k-1} - \mathbf{v}, -\frac{1}{\gamma m a_k} (\mathbf{y}_k - \mathbf{y}_{k-1}) + \frac{1}{\gamma m^2 a_k} (\hat{\mathbf{y}}_k - \mathbf{y}_{k-1}) \right\rangle, \tag{27}
\end{aligned}$$

where in the last equality we used  $\mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}} = \gamma m a_k (\mathbf{A}^{S^{j_k}} \mathbf{x}_k - \mathbf{b}^{S^{j_k}})$  and  $\hat{\mathbf{y}}_k - \mathbf{y}_{k-1} = \gamma m a_k (\mathbf{A} \mathbf{x}_k - \mathbf{b})$ , which both hold by definitions of  $\hat{\mathbf{y}}_k$  and  $\mathbf{y}_k$ .

Finally, combining (25)–(27), we have the bound on  $\psi_k(\mathbf{y}_k)$  from the statement of the lemma.  $\square$

By combining the two lower bounds in Lemma 3, we obtain the following result.

**Lemma 4.** *For any  $(\mathbf{u}, \mathbf{v}) \in \mathcal{X} \times \mathbb{R}^n$ , the sum of  $\psi_k(\mathbf{y}_k) + \phi_k(\mathbf{x}_k)$  can be bounded as follows: for all  $k \geq 1$ ,*

$$\begin{aligned}
& \phi_k(\mathbf{x}_k) + \psi_k(\mathbf{y}_k) \\
& \geq \phi_{k-1}(\mathbf{x}_{k-1}) + \psi_{k-1}(\mathbf{y}_{k-1}) \\
& \quad - m a_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T (\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle + m a_{k-1} \langle \mathbf{x}_{k-1} - \mathbf{u}, \mathbf{A}^T (\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle \\
& \quad + \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 - \frac{1}{4\gamma} \|\mathbf{y}_{k-1} - \mathbf{y}_{k-2}\|^2 + Q_k, \tag{28}
\end{aligned}$$

where

$$\begin{aligned}
Q_k := & a_k \left( r(\mathbf{x}_k) - r(\mathbf{u}) - \langle \mathbf{A} \mathbf{u}, \mathbf{y}_k \rangle + \langle \mathbf{x}_k, \mathbf{A}^T \mathbf{v} \rangle + \langle \mathbf{x}_k - \mathbf{u}, \mathbf{c} \rangle + \langle \mathbf{b}, \mathbf{y}_k - \mathbf{v} \rangle \right. \\
& \left. - (m-1) \langle \mathbf{A} \mathbf{u} - \mathbf{b}, \mathbf{y}_k - \mathbf{y}_{k-1} \rangle \right) + \frac{1}{\gamma} \left\langle \mathbf{y}_{k-1} - \mathbf{v}, -(\mathbf{y}_k - \mathbf{y}_{k-1}) + \frac{1}{m} (\hat{\mathbf{y}}_k - \mathbf{y}_{k-1}) \right\rangle. \tag{29}
\end{aligned}$$

*Proof.* Before our proof, as  $A_{-1}$  and  $A_0$  are not used in Algorithm 3, without loss of generality, we set  $A_{-1} = A_0 = 0$ . Fix any  $(\mathbf{u}, \mathbf{v}) \in \mathcal{X} \times \mathbb{R}^n$ . By combining the bounds on  $\phi_k(\mathbf{x}_k)$  and  $\psi_k(\mathbf{y}_k)$  from Lemma 3, we have  $\forall k \geq 1$  that

$$\begin{aligned}
& \phi_k(\mathbf{x}_k) + \psi_k(\mathbf{y}_k) \\
& \geq \phi_{k-1}(\mathbf{x}_{k-1}) + \psi_{k-1}(\mathbf{y}_{k-1}) \\
& \quad - m a_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T (\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle + m a_{k-1} \langle \mathbf{x}_{k-1} - \mathbf{u}, \mathbf{A}^T (\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle \\
& \quad + P_k + Q_k, \tag{30}
\end{aligned}$$

where

$$P_k = \frac{\gamma + \sigma A_{k-1}}{2} \|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 + \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 + m a_{k-1} \langle \mathbf{x}_k - \mathbf{x}_{k-1}, \mathbf{A}^T (\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle \tag{31}$$

and  $Q_k$  is defined in Eq. (29).

To find a lower bound on  $P_k$  we start by bounding the magnitude of the inner product term in (31). Recall that  $\mathbf{y}_{k-2}$  and  $\mathbf{y}_{k-1}$  differ only on the coordinate block  $S^{j_{k-1}}$ . We thus have:

$$\begin{aligned}
& |ma_{k-1} \langle \mathbf{x}_k - \mathbf{x}_{k-1}, \mathbf{A}^T(\mathbf{y}_{k-1} - \mathbf{y}_{k-2}) \rangle| \\
&= |ma_{k-1} \langle \mathbf{A}^{S^{j_{k-1}}}(\mathbf{x}_k - \mathbf{x}_{k-1}), \mathbf{y}_{k-1}^{S^{j_{k-1}}} - \mathbf{y}_{k-2}^{S^{j_{k-1}}} \rangle| \\
&\leq ma_{k-1} \|\mathbf{A}^{S^{j_{k-1}}}(\mathbf{x}_k - \mathbf{x}_{k-1})\| \|\mathbf{y}_{k-1}^{S^{j_{k-1}}} - \mathbf{y}_{k-2}^{S^{j_{k-1}}}\| \\
&\leq (ma_{k-1})^2 \gamma \|\mathbf{A}^{S^{j_{k-1}}}(\mathbf{x}_k - \mathbf{x}_{k-1})\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_{k-1}^{S^{j_{k-1}}} - \mathbf{y}_{k-2}^{S^{j_{k-1}}}\|^2 \\
&\leq (m\hat{L}a_{k-1})^2 \gamma \|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_{k-1}^{S^{j_{k-1}}} - \mathbf{y}_{k-2}^{S^{j_{k-1}}}\|^2 \\
&= (m\hat{L}a_{k-1})^2 \gamma \|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_{k-1} - \mathbf{y}_{k-2}\|^2
\end{aligned} \tag{32}$$

where the second inequality is by Young's inequality, and the third inequality is by Assumption 3 where  $\|\mathbf{A}^{S^{j_{k-1}}}(\mathbf{x}_k - \mathbf{x}_{k-1})\| \leq \hat{L} \|\mathbf{x}_k - \mathbf{x}_{k-1}\|$ . With our setting  $a_{k-1} = \frac{\sqrt{1+\sigma A_{k-2}/\gamma}}{2m\hat{L}} \leq \frac{\sqrt{1+\sigma A_{k-1}/\gamma}}{2m\hat{L}}$ , for all  $k \geq 1$ , we have

$$(m\hat{L}a_{k-1})^2 \gamma = \frac{\gamma + \sigma A_{k-1}}{4}.$$

By substituting this equality into (32) and then combining with (31), we obtain

$$\begin{aligned}
P_k &\geq \frac{\gamma + \sigma A_{k-1}}{4} \|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 + \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 - \frac{1}{4\gamma} \|\mathbf{y}_{k-1} - \mathbf{y}_{k-2}\|^2 \\
&\geq \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 - \frac{1}{4\gamma} \|\mathbf{y}_{k-1} - \mathbf{y}_{k-2}\|^2.
\end{aligned} \tag{33}$$

We complete the proof by combining Eqs. (30), (31) and the bound Eq. (33) for  $P_k$ .  $\square$

By telescoping the inequality in Lemma 4, and using Lemmas 1 and 2, we obtain the next result.

**Lemma 5.** For all  $(\mathbf{u}, \mathbf{v}) \in \mathcal{X} \times \mathbb{R}^n$ , we have

$$\begin{aligned}
& A_k(\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k)) + \frac{\gamma + \sigma A_k}{4} \|\mathbf{u} - \mathbf{x}_k\|^2 + \frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_k\|^2 \\
&\leq \frac{\gamma}{2} \|\mathbf{u} - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_0\|^2 - \frac{1}{4\gamma} \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 + \frac{1}{\gamma} \sum_{i=1}^k \langle \mathbf{y}_{i-1}, \check{\mathbf{v}}_i \rangle - \frac{1}{\gamma} \sum_{i=1}^k \langle \mathbf{v}, \check{\mathbf{v}}_i \rangle,
\end{aligned} \tag{34}$$

where  $\check{\mathbf{v}}_i$  is defined in Lemma 2.

*Proof.* Telescoping the inequality in Lemma 4, we have

$$\begin{aligned}
\phi_k(\mathbf{x}_k) + \psi_k(\mathbf{y}_k) &\geq \phi_0(\mathbf{x}_0) + \psi_0(\mathbf{y}_0) \\
&\quad - ma_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle + ma_0 \langle \mathbf{x}_0 - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_0 - \mathbf{y}_{-1}) \rangle \\
&\quad + \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 - \frac{1}{4\gamma} \|\mathbf{y}_0 - \mathbf{y}_{-1}\|^2 + \frac{1}{4\gamma} \sum_{i=1}^{k-1} \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 + \sum_{i=1}^k Q_i \\
&= -ma_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle + \frac{1}{4\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 \\
&\quad + \frac{1}{4\gamma} \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 + \sum_{i=1}^k Q_i,
\end{aligned} \tag{35}$$

where the last equality is by the fact that  $\psi_0(\mathbf{y}_0) = \phi_0(\mathbf{x}_0) = 0$ ,  $a_0 = 0$ , and our convention that  $\mathbf{y}_{-1} := \mathbf{y}_0$ .

Then by the convexity of  $r(\cdot)$  and the definition of  $\{Q_i\}$ , we have

$$\begin{aligned}
\sum_{i=1}^k Q_i &\geq A_k(r(\tilde{\mathbf{x}}_k) - r(\mathbf{u}) - \langle \mathbf{A}\mathbf{u}, \tilde{\mathbf{y}}_k \rangle + \langle \tilde{\mathbf{x}}_k, \mathbf{A}^T \mathbf{v} \rangle + \langle \tilde{\mathbf{x}}_k - \mathbf{u}, \mathbf{c} \rangle + \langle \mathbf{b}, \tilde{\mathbf{y}}_k - \mathbf{v} \rangle) \\
&\quad + \frac{1}{\gamma} \sum_{i=1}^k \left\langle \mathbf{y}_{i-1} - \mathbf{v}, -(\mathbf{y}_i - \mathbf{y}_{i-1}) + \frac{1}{m}(\hat{\mathbf{y}}_i - \mathbf{y}_{i-1}) \right\rangle \\
&= A_k(\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k)) + \frac{1}{\gamma} \sum_{i=1}^k \langle \mathbf{y}_{i-1} - \mathbf{v}, -\tilde{\mathbf{v}}_i \rangle. \tag{36}
\end{aligned}$$

where  $\tilde{\mathbf{x}}_k = \frac{1}{A_k} \sum_{i=1}^k a_i \mathbf{x}_i$ ,  $\tilde{\mathbf{y}}_k = \frac{1}{A_k} \sum_{i=1}^k (a_i \mathbf{y}_i + (m-1)a_i(\mathbf{y}_i - \mathbf{y}_{i-1}))$  (as defined in (4)) and the last equality is by the definition of the Lagrangian  $\mathcal{L}(\cdot, \cdot)$  and  $\{\tilde{\mathbf{v}}_i\}$  in Lemma 2.

Then by combining Eqs. (35)-(36) and Lemma 1, we have

$$\begin{aligned}
&A_k(\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k)) \\
&\leq \left( \psi_k(\mathbf{y}_k) + \phi_k(\mathbf{x}_k) + ma_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle \right) - \frac{1}{4\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 \\
&\quad - \frac{1}{4\gamma} \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 + \frac{1}{\gamma} \sum_{i=1}^k \langle \mathbf{y}_{i-1}, \tilde{\mathbf{v}}_i \rangle - \frac{1}{\gamma} \sum_{i=1}^k \langle \mathbf{v}, \tilde{\mathbf{v}}_i \rangle \\
&\leq \left( \frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_0\|^2 - \frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_k\|^2 + \frac{\gamma}{2} \|\mathbf{u} - \mathbf{x}_0\|^2 - \frac{\gamma + \sigma A_k}{2} \|\mathbf{u} - \mathbf{x}_k\|^2 \right) \\
&\quad + ma_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle - \frac{1}{4\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 - \frac{1}{4\gamma} \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \\
&\quad + \frac{1}{\gamma} \sum_{i=1}^k \langle \mathbf{y}_{i-1}, \tilde{\mathbf{v}}_i \rangle - \frac{1}{\gamma} \sum_{i=1}^k \langle \mathbf{v}, \tilde{\mathbf{v}}_i \rangle.
\end{aligned}$$

Finally, we have from the fact that  $\mathbf{y}_k$  and  $\mathbf{y}_{k-1}$  differ in only the  $S^{j_k}$  components, Young's inequality, and the definition of  $a_k$  that

$$\begin{aligned}
&|ma_k \langle \mathbf{x}_k - \mathbf{u}, \mathbf{A}^T(\mathbf{y}_k - \mathbf{y}_{k-1}) \rangle| \\
&= |ma_k \langle \mathbf{A}^{S^{j_k}}(\mathbf{x}_k - \mathbf{u}), \mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}} \rangle| \\
&\leq ma_k \|\mathbf{A}^{S^{j_k}}(\mathbf{x}_k - \mathbf{u})\| \|\mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}}\| \\
&\leq (ma_k)^2 \gamma \|\mathbf{A}^{S^{j_k}}(\mathbf{x}_k - \mathbf{u})\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}}\|^2 \\
&\leq (m\hat{L}a_k)^2 \gamma \|\mathbf{x}_k - \mathbf{u}\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_k^{S^{j_k}} - \mathbf{y}_{k-1}^{S^{j_k}}\|^2 \\
&= (m\hat{L}a_k)^2 \gamma \|\mathbf{x}_k - \mathbf{u}\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 \\
&= \frac{\gamma + \sigma A_{k-1}}{4} \|\mathbf{x}_k - \mathbf{u}\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2 \\
&\leq \frac{\gamma + \sigma A_k}{4} \|\mathbf{x}_k - \mathbf{u}\|^2 + \frac{1}{4\gamma} \|\mathbf{y}_k - \mathbf{y}_{k-1}\|^2
\end{aligned}$$

leading to

$$\begin{aligned}
& A_k(\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k)) + \frac{1}{2\gamma}\|\mathbf{v} - \mathbf{y}_k\|^2 + \frac{\gamma + \sigma A_k}{4}\|\mathbf{u} - \mathbf{x}_k\|^2 \\
& \leq \frac{\gamma}{2}\|\mathbf{u} - \mathbf{x}_0\|^2 + \frac{1}{2\gamma}\|\mathbf{v} - \mathbf{y}_0\|^2 - \frac{1}{4\gamma}\sum_{i=1}^k\|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \\
& \quad + \frac{1}{\gamma}\sum_{i=1}^k\langle \mathbf{y}_{i-1}, \check{\mathbf{v}}_i \rangle - \frac{1}{\gamma}\sum_{i=1}^k\langle \mathbf{v}, \check{\mathbf{v}}_i \rangle
\end{aligned}$$

and completing the proof.  $\square$

**Lemma 6.** *Suppose that  $(\mathbf{x}^*, \mathbf{y}^*)$  is a Nash point for (PD-GLP). Then the iterates  $\mathbf{x}_k, \mathbf{y}_k$  from Algorithm 1 satisfy*

$$\begin{aligned}
& \mathbb{E}\left[\frac{\gamma + \sigma A_k}{4}\|\mathbf{x}^* - \mathbf{x}_k\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_k\|^2 + \frac{1}{4\gamma}\sum_{i=1}^k\|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2\right] \\
& \leq \frac{\gamma}{2}\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2,
\end{aligned} \tag{37}$$

where the expectation is w.r.t. all the randomness in the algorithm.

*Proof.* Note that the existence of a Nash point is assumed in Assumption 2. With  $(\mathbf{u}, \mathbf{v}) = (\mathbf{x}^*, \mathbf{y}^*)$ , by the definition of Nash equilibrium, we have

$$\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{y}^*) - \mathcal{L}(\mathbf{x}^*, \tilde{\mathbf{y}}_k) = (\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{y}^*) - \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*)) - (\mathcal{L}(\mathbf{x}^*, \tilde{\mathbf{y}}_k) - \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*)) \geq 0. \tag{38}$$

By setting  $(\mathbf{u}, \mathbf{v}) = (\mathbf{x}^*, \mathbf{y}^*)$  in the result of Lemma 5, using (38) to eliminate the first term on the left-hand side of the inequality, and rearranging, we obtain

$$\begin{aligned}
& \frac{\gamma + \sigma A_k}{4}\|\mathbf{x}^* - \mathbf{x}_k\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_k\|^2 + \frac{1}{4\gamma}\sum_{i=1}^k\|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \\
& \leq \frac{\gamma}{2}\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2 + \frac{1}{\gamma}\sum_{i=1}^k\langle \mathbf{y}_{i-1}, \check{\mathbf{v}}_i \rangle - \frac{1}{\gamma}\sum_{i=1}^k\langle \mathbf{v}, \check{\mathbf{v}}_i \rangle.
\end{aligned}$$

The result will follow when we show that the expectation (with respect to all the randomness) of the last two terms on the right-hand side is zero. Since  $\mathbf{y}_{i-1} \in \mathcal{F}_{i-1}$ , we have

$$\mathbb{E}\left[\sum_{i=1}^k\langle \mathbf{y}_{i-1}, \check{\mathbf{v}}_i \rangle\right] = \mathbb{E}\left[\mathbb{E}\left[\sum_{i=1}^k\langle \mathbf{y}_{i-1}, \check{\mathbf{v}}_i \rangle \middle| \mathcal{F}_{i-1}\right]\right] = \mathbb{E}\left[\sum_{i=1}^k\langle \mathbf{y}_{i-1}, \mathbb{E}[\check{\mathbf{v}}_i | \mathcal{F}_{i-1}]\rangle\right] = 0, \tag{39}$$

which takes care of the second-last term. For any fixed  $\mathbf{v} \in \mathbb{R}^n$ , we also have

$$\mathbb{E}\left[\frac{1}{\gamma}\sum_{i=1}^k\langle \mathbf{v}, \check{\mathbf{v}}_i \rangle\right] = 0, \tag{40}$$

which takes care of the last term, and completes the proof.  $\square$

Next we state a technical lemma, proved in an earlier paper, which provides the final ingredient for the proof of Theorem 1.

**Lemma 7.** *[[55]] Let  $\{A_k\}_{k \geq 0}$  be a sequence of nonnegative real numbers such that  $A_0 = 0$  and  $A_k$  is defined recursively via  $A_k = A_{k-1} + \sqrt{c_1^2 + c_2 A_{k-1}}$ , where  $c_1 > 0$ , and  $c_2 \geq 0$ . Define  $K_0 = \lceil \frac{c_2}{9c_1} \rceil$ . Then*

$$A_k \geq \begin{cases} \frac{c_2}{9}\left(k - K_0 + \max\left\{3\sqrt{\frac{c_1}{c_2}}, 1\right\}\right)^2, & \text{if } c_2 > 0 \text{ and } k > K_0, \\ c_1 k, & \text{otherwise.} \end{cases}$$

We are now ready to prove our main result, which we restate here.

**Theorem 1.** Let  $\mathbf{x}_k, \mathbf{y}_k, k \in [K]$ , be the iterates of Algorithm 1 and let  $\tilde{\mathbf{x}}_k, \tilde{\mathbf{y}}_k$  be defined by

$$\tilde{\mathbf{x}}_k = \frac{1}{A_k} \sum_{i=1}^k a_i \mathbf{x}_i, \quad \tilde{\mathbf{y}}_k = \frac{1}{A_k} \sum_{i=1}^k (a_i \mathbf{y}_i + (m-1)a_i(\mathbf{y}_i - \mathbf{y}_{i-1})), \quad (4)$$

for  $k \in [K]$ . Let  $\mathcal{W}_k \subset \mathcal{X} \times \mathbb{R}^n, k \in [K]$ , be a sequence of compact convex sets such that  $(\tilde{\mathbf{x}}_k, \tilde{\mathbf{y}}_k) \in \mathcal{W}_k \subset \mathcal{W} \subset \mathcal{X} \times \mathbb{R}^n$ , where  $\mathcal{W}$  is also convex and compact. Then:

$$\begin{aligned} & \mathbb{E} \left[ \sup_{(\mathbf{u}, \mathbf{v}) \in \mathcal{W}_k} \{ \mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k) \} \right] \\ & \leq \frac{1}{A_k} \left( \mathbb{E} \left[ \frac{\gamma}{2} \|\hat{\mathbf{u}} - \mathbf{x}_0\|^2 + \frac{1}{\gamma} \|\hat{\mathbf{v}} - \mathbf{y}_0\|^2 \right] + \frac{\gamma}{2} \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2 \right), \end{aligned} \quad (5)$$

where  $(\hat{\mathbf{u}}, \hat{\mathbf{v}}) = \arg \sup_{(\mathbf{u}, \mathbf{v}) \in \mathcal{W}_k} \{ \mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k) \}$ . Furthermore,

$$\mathbb{E} \left[ \frac{\gamma + \sigma A_k}{4} \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{2\gamma} \|\mathbf{y}_k - \mathbf{y}^*\|^2 \right] \leq \frac{\gamma}{2} \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2. \quad (6)$$

Define  $K_0 = \lceil \frac{\sigma}{18\hat{L}m\gamma} \rceil$ . Then in the bounds above:

$$A_k \geq \max \left\{ \frac{k}{2\hat{L}m}, \frac{\sigma}{(6\hat{L}m)^2\gamma} \left( k - K_0 + \max \left\{ 3\sqrt{2\hat{L}m\gamma/\sigma}, 1 \right\} \right)^2 \right\}.$$

*Proof.* To provide a guarantee in terms of primal-dual gap, we need to take the supremum on  $\mathbf{u}, \mathbf{v}$  of  $\{ \mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k) \}$  over  $\mathcal{W}_k$ ; we denote the arg sup by  $(\hat{\mathbf{u}}, \hat{\mathbf{v}})$ . When we subsequently take the expectation, we have to account for the fact that  $\hat{\mathbf{v}}$  will be correlated with the randomness in the iteration history. We can however, use Lemmas 2 and 6 to give the upper bound on  $\mathbb{E} \left[ -\sum_{i=1}^k \langle \hat{\mathbf{v}}, \tilde{\mathbf{v}}_i \rangle \right]$ .

From (34), using the fact that  $\frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_k\|^2 + \frac{\gamma + \sigma A_k}{4} \|\mathbf{u} - \mathbf{x}_k\|^2 \geq 0$ , we have

$$\begin{aligned} & \mathbb{E} \left[ A_k \sup_{(\mathbf{u}, \mathbf{v}) \in \mathcal{W}_k} \{ \mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k) \} \right] \\ & \leq \mathbb{E} \left[ \frac{\gamma}{2} \|\hat{\mathbf{u}} - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}_0 - \hat{\mathbf{v}}\|^2 \right] - \frac{1}{4\gamma} \mathbb{E} \left[ \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] + \frac{1}{\gamma} \mathbb{E} \left[ \sum_{i=1}^k \langle \mathbf{y}_{i-1}, \tilde{\mathbf{v}}_i \rangle \right] \\ & \quad + \frac{1}{\gamma} \mathbb{E} \left[ -\sum_{i=1}^k \langle \hat{\mathbf{v}}, \tilde{\mathbf{v}}_i \rangle \right] \\ & \leq \mathbb{E} \left[ \frac{\gamma}{2} \|\hat{\mathbf{u}} - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}_0 - \hat{\mathbf{v}}\|^2 \right] - \frac{1}{4\gamma} \mathbb{E} \left[ \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] \\ & \quad + \frac{1}{\gamma} \mathbb{E} \left[ \frac{1}{2} \|\mathbf{y}_0 - \hat{\mathbf{v}}\|^2 + \frac{1}{2} \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] \\ & = \mathbb{E} \left[ \frac{\gamma}{2} \|\hat{\mathbf{u}} - \mathbf{x}_0\|^2 + \frac{1}{\gamma} \|\mathbf{y}_0 - \hat{\mathbf{v}}\|^2 \right] + \frac{1}{4\gamma} \mathbb{E} \left[ \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] \\ & \leq \mathbb{E} \left[ \frac{\gamma}{2} \|\hat{\mathbf{u}} - \mathbf{x}_0\|^2 + \frac{1}{\gamma} \|\mathbf{y}_0 - \hat{\mathbf{v}}\|^2 \right] + \frac{\gamma}{2} \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2. \end{aligned} \quad (41)$$

where the first inequality is by Eq. (34), the second inequality is by Lemma 2 and (39), and the last inequality is by Lemma 6. This proves the first claim (5). The second claim (6) follows from Lemma 6 with the fact that  $\frac{1}{4\gamma} \mathbb{E} \left[ \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] \geq 0$ . The final claim concerning  $A_k$  follows from Lemma 7 when we set

$$c_1 = \frac{1}{2\hat{L}m}, \quad c_2 = \frac{\sigma}{(2\hat{L}m)^2\gamma}.$$

□

**Corollary 1.** In Algorithm 1, for all  $k \geq 1$ ,  $\tilde{\mathbf{x}}_k$  satisfies

$$\begin{aligned} \mathbb{E}[\|\mathbf{y}^*\| \cdot \|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|] &\leq \frac{\gamma\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2 + \frac{1}{\gamma}\mathbb{E}[\|\mathbf{v} - \mathbf{y}_0\|^2]}{A_k}, \\ |\mathbb{E}[(\mathbf{c}^T \tilde{\mathbf{x}}_k + r(\tilde{\mathbf{x}}_k)) - (\mathbf{c}^T \mathbf{x}^* + r(\mathbf{x}^*))]| &\leq \frac{\gamma\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2 + \frac{1}{\gamma}\mathbb{E}[\|\mathbf{v} - \mathbf{y}_0\|^2]}{A_k}, \end{aligned}$$

where  $\mathbf{v} = 2\frac{\|\mathbf{y}^*\|}{\|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|}(\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b})$ .

*Proof.* Assume that  $\|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\| \neq 0$ , as otherwise the first bound follows trivially. Let  $\mathbf{u} = \mathbf{x}^*$  and  $\mathbf{v} = \frac{2\|\mathbf{y}^*\|(\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b})}{\|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|}$ . Then we have

$$\begin{aligned} &\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k) \\ &= (\mathbf{c}^T \tilde{\mathbf{x}}_k + r(\tilde{\mathbf{x}}_k) + 2\|\mathbf{y}^*\| \|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|) - (\mathbf{c}^T \mathbf{x}^* + r(\mathbf{x}^*) + \tilde{\mathbf{y}}_k^T (\mathbf{A}\mathbf{x}^* - \mathbf{b})) \\ &= (\mathbf{c}^T (\tilde{\mathbf{x}}_k - \mathbf{x}^*) + r(\tilde{\mathbf{x}}_k) - r(\mathbf{x}^*)) + 2\|\mathbf{y}^*\| \|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|. \end{aligned} \quad (42)$$

For any fixed  $\mathbf{u}$ , and any  $\mathbf{v} \in \mathbb{R}^n$  possibly depending on the randomness in the algorithm, we have from Lemma 5, taking expectation over all the randomness, that

$$\begin{aligned} &A_k \mathbb{E}[\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k)] \\ &\leq \mathbb{E} \left[ A_k (\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{v}) - \mathcal{L}(\mathbf{u}, \tilde{\mathbf{y}}_k)) + \frac{1}{2\gamma} \|\mathbf{v} - \mathbf{y}_k\|^2 + \frac{\gamma + \sigma A_k}{4} \|\mathbf{u} - \mathbf{x}_k\|^2 \right] \\ &\leq \frac{\gamma}{2} \|\mathbf{u} - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \mathbb{E}[\|\mathbf{v} - \mathbf{y}_0\|^2] - \mathbb{E} \left[ \frac{1}{4\gamma} \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] \\ &\quad + \frac{1}{\gamma} \mathbb{E} \left[ \sum_{i=1}^k \langle \mathbf{y}_{i-1}, \tilde{\mathbf{v}}_i \rangle \right] - \frac{1}{\gamma} \mathbb{E} \left[ \sum_{i=1}^k \langle \mathbf{v}, \tilde{\mathbf{v}}_i \rangle \right]. \end{aligned} \quad (43)$$

Meanwhile, we have

$$\begin{aligned} &-\frac{1}{4\gamma} \mathbb{E} \left[ \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] - \frac{1}{\gamma} \mathbb{E} \left[ \sum_{i=1}^k \langle \mathbf{v}, \tilde{\mathbf{v}}_i \rangle \right] \\ &\leq \frac{1}{2\gamma} \mathbb{E}[\|\mathbf{y}_0 - \mathbf{v}\|^2] + \frac{1}{4\gamma} \mathbb{E} \left[ \sum_{i=1}^k \|\mathbf{y}_i - \mathbf{y}_{i-1}\|^2 \right] \\ &\leq \frac{1}{2\gamma} \mathbb{E}[\|\mathbf{y}_0 - \mathbf{v}\|^2] + \frac{\gamma}{2} \|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma} \|\mathbf{y}^* - \mathbf{y}_0\|^2, \end{aligned}$$

where the first inequality is by Lemma 2 and the second inequality is by Lemma 6. Since  $\mathbf{y}_{i-1} \in \mathcal{F}_{i-1}$ , we have as in (39) that

$$\frac{1}{\gamma} \mathbb{E} \left[ \sum_{i=1}^k \langle \mathbf{y}_{i-1}, \tilde{\mathbf{v}}_i \rangle \right] = 0. \quad (44)$$

By combining Eq. (42)-(44) with  $\mathbf{u} = \mathbf{x}^*$ , we have

$$\begin{aligned} &\mathbb{E}[(\mathbf{c}^T (\tilde{\mathbf{x}}_k - \mathbf{x}^*) + r(\tilde{\mathbf{x}}_k) - r(\mathbf{x}^*)) + 2\|\mathbf{y}^*\| \|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|] \\ &\leq \frac{\gamma\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{\gamma}\mathbb{E}[\|\mathbf{v} - \mathbf{y}_0\|^2] + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2}{A_k}. \end{aligned} \quad (45)$$

By the KKT condition of (PD-GLP) and the optimality of  $(\mathbf{x}^*, \mathbf{y}^*)$ , we have for all  $\mathbf{x} \in \mathcal{X}$  that

$$(\mathbf{c}^T \mathbf{x} + r(\mathbf{x})) - (\mathbf{c}^T \mathbf{x}^* + r(\mathbf{x}^*)) - \langle \mathbf{y}^*, \mathbf{A}\mathbf{x} - \mathbf{b} \rangle \geq 0, \quad (46)$$

and thus

$$(\mathbf{c}^T \mathbf{x} + r(\mathbf{x})) - (\mathbf{c}^T \mathbf{x}^* + r(\mathbf{x}^*)) \geq -\|\mathbf{y}^*\| \|\mathbf{A}\mathbf{x} - \mathbf{b}\|. \quad (47)$$

By combining Eq. (45) and Eq. (47), we have

$$\mathbb{E}[\|\mathbf{y}^*\| \cdot \|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|] \leq \frac{\gamma\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{\gamma}\mathbb{E}[\|\mathbf{v} - \mathbf{y}_0\|^2] + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2}{A_k}, \quad (48)$$

proving our first claim. The second claim is obtained by combining Eqs. (45), (47), and (48).  $\square$

## B.2 Omitted proofs from Section 3.2

We need to compute explicitly only those components of  $\mathbf{q}_{k-1}$  and  $\mathbf{x}_k$  that are needed to update  $\mathbf{y}_k$ .

**Lemma 8.** For  $\{\mathbf{q}_k\}$  defined in Algorithm 1, we have

$$\mathbf{q}_k = A_{k+1}(\mathbf{c} + \mathbf{z}_k) + \mathbf{r}_k, \quad k = 0, 1, 2, \dots,$$

where  $\mathbf{r}_0 = \mathbf{0}$  and for all  $k = 1, 2, \dots$

$$\mathbf{r}_k = \mathbf{r}_{k-1} + (ma_k - A_k)(\mathbf{z}_k - \mathbf{z}_{k-1}).$$

*Proof.* The proof is by induction. For  $k = 0$ , we have  $\mathbf{q}_0 = A_1(\mathbf{c} + \mathbf{z}_0)$  which is true by definition. Assuming that the result holds for some index  $k$ , we show that it continues to hold for  $k + 1$ . We have from Step 9 of Algorithm 1, then using the inductive assumption, that

$$\begin{aligned} \mathbf{q}_{k+1} &= a_{k+2}(\mathbf{z}_{k+1} + \mathbf{c}) + ma_{k+1}(\mathbf{z}_{k+1} - \mathbf{z}_k) + \mathbf{q}_k \\ &= a_{k+2}(\mathbf{z}_{k+1} + \mathbf{c}) + ma_{k+1}(\mathbf{z}_{k+1} - \mathbf{z}_k) + A_{k+1}(\mathbf{c} + \mathbf{z}_k) + \mathbf{r}_k \\ &= A_{k+2}\mathbf{c} + (A_{k+2} - A_{k+1})\mathbf{z}_{k+1} + ma_{k+1}(\mathbf{z}_{k+1} - \mathbf{z}_k) + A_{k+1}\mathbf{z}_k + \mathbf{r}_k \\ &= A_{k+2}(\mathbf{c} + \mathbf{z}_{k+1}) + (ma_{k+1} - A_{k+1})(\mathbf{z}_{k+1} - \mathbf{z}_k) + \mathbf{r}_k \\ &= A_{k+2}(\mathbf{c} + \mathbf{z}_{k+1}) + \mathbf{r}_{k+1}, \end{aligned}$$

as required.  $\square$

This lemma indicates that we can reconstruct  $\mathbf{q}_k$  (or any subvector of  $\mathbf{q}_k$  that we need, on demand), provided we maintain  $\mathbf{z}_k$  and  $\mathbf{r}_k$ . Note that the update from  $\mathbf{z}_k$  to  $\mathbf{z}_{k+1}$  is sparse; these two vectors differ only in the components corresponding to the block  $C^{j_k}$ . To obtain  $\mathbf{r}_{k+1}$  from  $\mathbf{r}_k$ , we need to add a scalar times the sparse vector  $\mathbf{z}_{k+1} - \mathbf{z}_k$ , so this update is also efficient.

We can also maintain an implicit representation of the averaged vector  $\tilde{\mathbf{y}}_k$  efficiently, as shown in the following lemma. We defer the proofs of Lemmas 8 and 9 to Appendix B.

**Lemma 9.** For  $\{\tilde{\mathbf{y}}_k\}$  defined in (4), we have

$$\tilde{\mathbf{y}}_k = \mathbf{y}_k + \frac{1}{A_k}\mathbf{s}_k, \quad k = 1, 2, \dots,$$

where  $\mathbf{s}_0 = \mathbf{0}$  and for all  $k = 1, 2, \dots$ ,

$$\mathbf{s}_k = \mathbf{s}_{k-1} + ((m-1)a_k - A_{k-1})(\mathbf{y}_k - \mathbf{y}_{k-1}).$$

*Proof.* Recall that  $\tilde{\mathbf{y}}_K$  ( $K \geq 1$ ) is defined in Step 11 of Algorithm 1. The proof is by induction. For  $k = 1$ ,  $\tilde{\mathbf{y}}_1 = \mathbf{y}_1 + \frac{1}{A_1}(m-1)a_1(\mathbf{y}_1 - \mathbf{y}_0) = \mathbf{y}_1 + \frac{1}{A_1}\mathbf{s}_1$  which is true by the definition of  $\mathbf{s}_1$  and  $A_0$ . Next, we assume that the result holds for some  $k \geq 2$  and show that it continues to hold for  $k + 1$ . We have

$$\begin{aligned} A_{k+1}\tilde{\mathbf{y}}_{k+1} &= \sum_{i=1}^{k+1} (a_i\mathbf{y}_i + (m-1)a_i(\mathbf{y}_i - \mathbf{y}_{i-1})) \\ &= A_k\tilde{\mathbf{y}}_k + a_{k+1}\mathbf{y}_{k+1} + (m-1)a_{k+1}(\mathbf{y}_{k+1} - \mathbf{y}_k) \\ &= A_k\mathbf{y}_k + \mathbf{s}_k + a_{k+1}\mathbf{y}_{k+1} + (m-1)a_{k+1}(\mathbf{y}_{k+1} - \mathbf{y}_k) \\ &= A_k(\mathbf{y}_k - \mathbf{y}_{k+1} + \mathbf{y}_{k+1}) + \mathbf{s}_k + a_{k+1}\mathbf{y}_{k+1} + (m-1)a_{k+1}(\mathbf{y}_{k+1} - \mathbf{y}_k) \\ &= A_{k+1}\mathbf{y}_{k+1} + \mathbf{s}_k + ((m-1)a_{k+1} - A_k)(\mathbf{y}_{k+1} - \mathbf{y}_k) \\ &= A_{k+1}\mathbf{y}_{k+1} + \mathbf{s}_{k+1}. \end{aligned}$$

as required.  $\square$

## B.3 Omitted proofs from Section 3.3

The standard-form LP (PD-LP) is derived by setting  $r(\mathbf{x}) \equiv 0$  and  $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^d : x_i \geq 0, i \in [d]\}$  in (PD-GLP). Given any  $\mathbf{w} \in \mathcal{X} \times \mathbb{R}^n$ , we define a compact convex subset  $\mathcal{W}_{\zeta, \gamma}(\mathbf{w})$  of  $\mathcal{X} \times \mathbb{R}^n$  as follows:

$$\mathcal{W}_{\zeta, \gamma}(\mathbf{w}) = \{\hat{\mathbf{w}} \in \mathcal{X} \times \mathbb{R}^n : \|\hat{\mathbf{w}} - \mathbf{w}\|_{(\gamma)} \leq \zeta, \zeta > 0, \gamma > 0\}, \quad (49)$$

where we have defined  $\|\cdot\|_{(\gamma)}$  by  $\|\mathbf{w}\|_{(\gamma)} = \sqrt{\gamma\|\mathbf{x}\|^2 + \frac{1}{\gamma}\|\mathbf{y}\|^2}$  in Section 3.3.

**Lemma 10.** Consider the standard-form LP (LP). Given  $\tau > 0$  and  $\mathbf{w} \in \mathcal{X} \times \mathbb{R}^n$ , if  $\|\mathbf{w}\|_{(\gamma)} \leq \tau$  and  $\zeta \leq \tau$ , then

$$\sup_{\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\mathbf{w})} \{\mathcal{L}(\mathbf{x}, \hat{\mathbf{y}}) - \mathcal{L}(\hat{\mathbf{x}}, \mathbf{y})\} \geq \frac{\zeta}{\gamma + 1/\gamma + \tau} \text{LPMetric}(\mathbf{w}). \quad (50)$$

*Proof.* Let  $\mathbf{F}(\mathbf{w}) = \begin{bmatrix} \mathbf{A}^T \mathbf{y} + \mathbf{c} \\ \mathbf{b} - \mathbf{A} \mathbf{x} \end{bmatrix}$ . Then we have

$$\rho_{\zeta, \gamma}(\mathbf{w}) := \sup_{\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\mathbf{w})} \{\mathcal{L}(\mathbf{x}, \hat{\mathbf{y}}) - \mathcal{L}(\hat{\mathbf{x}}, \mathbf{y})\} = \sup_{\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\mathbf{w})} \mathbf{F}(\mathbf{w})^T (\mathbf{w} - \hat{\mathbf{w}}) \geq 0, \quad (51)$$

where the inequality follows from  $\mathbf{w} \in \mathcal{W}_{\zeta, \gamma}(\mathbf{w})$ .

First, we prove

$$\rho_{\zeta, \gamma}(\mathbf{w}) \geq \frac{\zeta \|\max\{-\mathbf{A}^T \mathbf{y} - \mathbf{c}, \mathbf{0}\}\|_2}{\gamma}. \quad (52)$$

If  $\|\max\{-\mathbf{A}^T \mathbf{y} - \mathbf{c}, \mathbf{0}\}\|_2 > 0$ , for  $\hat{\mathbf{w}}_1 = (\hat{\mathbf{x}}_1, \hat{\mathbf{y}}_1)$ , let  $\hat{\mathbf{x}}_1 = \mathbf{x} + \frac{\zeta}{\gamma \|\max\{-\mathbf{A}^T \mathbf{y} - \mathbf{c}, \mathbf{0}\}\|_2} \max\{-\mathbf{A}^T \mathbf{y} - \mathbf{c}, \mathbf{0}\} \in \mathcal{X}$  and  $\hat{\mathbf{y}}_1 = \mathbf{y}$ . Then we have  $\|\hat{\mathbf{w}}_1 - \mathbf{w}\|_{(\gamma)} = \zeta$  and thus  $\hat{\mathbf{w}}_1 \in \mathcal{W}_{\zeta, \gamma}(\mathbf{w})$ . So, as  $\rho_{\zeta, \gamma}(\mathbf{w}) \geq \mathbf{F}(\mathbf{w})^T (\mathbf{w} - \hat{\mathbf{w}}_1)$ , with the definition of  $\hat{\mathbf{w}}_1$ , Eq. (52) holds. If  $\|\max\{-\mathbf{A}^T \mathbf{y} - \mathbf{c}, \mathbf{0}\}\|_2 = 0$ , then Eq. (52) holds trivially.

Second, we prove

$$\rho_{\zeta, \gamma}(\mathbf{w}) \geq \zeta \gamma \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2. \quad (53)$$

If  $\|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2 > 0$ , for  $\hat{\mathbf{w}}_2 = (\hat{\mathbf{x}}_2, \hat{\mathbf{y}}_2)$ , let  $\hat{\mathbf{x}}_2 = \mathbf{x}$  and  $\hat{\mathbf{y}}_2 = \mathbf{y} + \frac{\zeta \gamma}{\|\mathbf{A} \mathbf{x} - \mathbf{b}\|} (\mathbf{A} \mathbf{x} - \mathbf{b})$ . Then we have  $\|\hat{\mathbf{w}}_2 - \mathbf{w}\|_{(\gamma)} = \zeta$  and thus  $\hat{\mathbf{w}}_2 \in \mathcal{W}_{\zeta, \gamma}(\mathbf{w})$ . Then as  $\rho_{\zeta, \gamma}(\mathbf{w}) \geq \mathbf{F}(\mathbf{w})^T (\mathbf{w} - \hat{\mathbf{w}}_2)$ , Eq. (53) holds. If  $\|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2 = 0$ , then Eq. (53) holds trivially.

Third, we prove that

$$\rho_{\zeta, \gamma}(\mathbf{w}) \geq \frac{\zeta}{\tau} \max\{\mathbf{c}^T \mathbf{x} + \mathbf{b}^T \mathbf{y}, \mathbf{0}\}. \quad (54)$$

Note that the inequality holds trivially if  $\mathbf{c}^T \mathbf{x} + \mathbf{b}^T \mathbf{y} \leq 0$ , so we assume that  $\mathbf{c}^T \mathbf{x} + \mathbf{b}^T \mathbf{y} > 0$ , and note that  $\mathbf{F}(\mathbf{w})^T \mathbf{w} = \mathbf{c}^T \mathbf{x} + \mathbf{b}^T \mathbf{y} > 0$  in this case. This condition implies that  $\mathbf{w} \neq \mathbf{0}$ , so we can define  $\hat{\mathbf{w}}_3 = \mathbf{w} - \min\{\frac{\zeta}{\|\mathbf{w}\|_{(\gamma)}}, 1\} \mathbf{w}$ . Then we have  $\|\hat{\mathbf{w}}_3 - \mathbf{w}\|_{(\gamma)} \leq \frac{\zeta}{\|\mathbf{w}\|_{(\gamma)}} \|\mathbf{w}\|_{(\gamma)} \leq \zeta$ . Meanwhile, we also have  $\hat{\mathbf{x}}_3 \geq \mathbf{x}_3 - \mathbf{x}_3 = \mathbf{0}$ , so that  $\hat{\mathbf{w}}_3 \in \mathcal{X} \times \mathbb{R}^n$ . Thus, we have  $\hat{\mathbf{w}}_3 \in \mathcal{W}_{\zeta, \gamma}(\mathbf{w})$ . Then, with the assumptions  $\zeta \leq \tau$  and  $\|\mathbf{w}\|_{(\gamma)} \leq \tau$ , together with  $\mathbf{F}(\mathbf{w})^T \mathbf{w} > 0$ , we have

$$\rho_{\zeta, \gamma}(\mathbf{w}) \geq \min\left\{\frac{\zeta}{\|\mathbf{w}\|_{(\gamma)}}, 1\right\} \mathbf{F}(\mathbf{w})^T \mathbf{w} \geq \min\left\{\frac{\zeta}{\tau}, 1\right\} \mathbf{F}(\mathbf{w})^T \mathbf{w} = \frac{\zeta}{\tau} \mathbf{F}(\mathbf{w})^T \mathbf{w} = \frac{\zeta}{\tau} (\mathbf{c}^T \mathbf{x} + \mathbf{b}^T \mathbf{y}), \quad (55)$$

completing the proof of the claim (54).

By combining Eqs. (52), (53) and (54), we obtain

$$\begin{aligned} (\gamma + 1/\gamma + \tau)^2 \rho_{\zeta, \gamma}^2(\mathbf{w}) &\geq \zeta^2 (\|\max\{-\mathbf{A}^T \mathbf{y} - \mathbf{c}, \mathbf{0}\}\|_2^2 + \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2^2 + \max\{\mathbf{c}^T \mathbf{x} + \mathbf{b}^T \mathbf{y}, \mathbf{0}\}^2) \\ &\geq \zeta^2 (\text{LPMetric}(\mathbf{w}))^2, \end{aligned} \quad (56)$$

from which the result follows.  $\square$

**Lemma 11.** Consider the standard-form LP (LP), and let  $\mathbf{w}^* = (\mathbf{x}^*, \mathbf{y}^*)$  be a Nash point (that is, a solution of the primal-dual form (PD-LP)). For a starting point  $\mathbf{w}_0$ , define  $\zeta = \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}$  and choose  $\gamma > 0$ . Then for all  $k = 1, 2, \dots$ , we have

$$\mathbb{E}[\|\tilde{\mathbf{w}}_k - \mathbf{w}^*\|_{(\gamma)}] \leq \sqrt{2} \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)},$$

where the expectation is taken w.r.t. all the randomness up to iteration  $k$ . Further, for  $\mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)$  defined as in (49), we have

$$\mathbb{E}\left[\sup_{\mathbf{w} \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)} \{\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{y}) - \mathcal{L}(\mathbf{y}, \tilde{\mathbf{y}}_k)\}\right] \leq \frac{25 \hat{L} m}{k} \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2. \quad (57)$$

*Proof.* By Theorem 1, we have

$$\mathbb{E}\left[\frac{\gamma}{4}\|\mathbf{x}^* - \mathbf{x}_k\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_k\|^2\right] \leq \frac{\gamma}{2}\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2 = \frac{1}{2}\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2,$$

by the definition of  $\|\cdot\|_{(\gamma)}$ . Using this definition again, we have that the left-hand side in this expression is bounded below by  $\frac{1}{4}\|\mathbf{w}_k - \mathbf{w}^*\|_{(\gamma)}^2$ , so that

$$\mathbb{E}\left[\|\mathbf{w}_k - \mathbf{w}^*\|_{(\gamma)}^2\right] \leq 2\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2, \quad \forall k \geq 1. \quad (58)$$

By convexity of  $\|\cdot\|^2$ , we obtain

$$\mathbb{E}[\|\tilde{\mathbf{w}}_k - \mathbf{w}^*\|_{(\gamma)}^2] = \mathbb{E}\left[\left\|\frac{1}{k}\sum_{i=1}^k(\mathbf{w}_i - \mathbf{w}^*)\right\|_{(\gamma)}^2\right] \leq \frac{1}{k}\mathbb{E}\left[\sum_{i=1}^k\|\mathbf{w}_i - \mathbf{w}^*\|_{(\gamma)}^2\right] \leq 2\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2. \quad (59)$$

The first claim of the lemma now follows by applying Jensen's inequality to this bound.

For  $\hat{\mathbf{w}} = (\hat{\mathbf{x}}, \hat{\mathbf{y}}) \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)$ , we have the following bound on  $\mathbb{E}\left[\gamma\|\hat{\mathbf{x}} - \mathbf{x}_0\|^2 + \frac{1}{\gamma}\|\hat{\mathbf{y}} - \mathbf{y}_0\|^2\right]$ :

$$\begin{aligned} & \mathbb{E}\left[\gamma\|\hat{\mathbf{x}} - \mathbf{x}_0\|^2 + \frac{1}{\gamma}\|\hat{\mathbf{y}} - \mathbf{y}_0\|^2\right] \\ &= \mathbb{E}[\|\mathbf{w}_0 - \hat{\mathbf{w}}\|_{(\gamma)}^2] \\ &= \mathbb{E}[\|\mathbf{w}_0 - \mathbf{w}^* + \mathbf{w}^* - \tilde{\mathbf{w}}_k + \tilde{\mathbf{w}}_k - \hat{\mathbf{w}}\|_{(\gamma)}^2] \\ &\leq \mathbb{E}[(\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} + \|\mathbf{w}^* - \tilde{\mathbf{w}}_k\|_{(\gamma)} + \|\tilde{\mathbf{w}}_k - \hat{\mathbf{w}}\|_{(\gamma)})^2] \\ &\leq \mathbb{E}[3(\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2 + \|\mathbf{w}^* - \tilde{\mathbf{w}}_k\|_{(\gamma)}^2 + \|\tilde{\mathbf{w}}_k - \hat{\mathbf{w}}\|_{(\gamma)}^2)] \\ &\leq \mathbb{E}[3(\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2 + 2\|\mathbf{w}^* - \mathbf{w}_0\|_{(\gamma)}^2 + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2)] \\ &= 12\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2, \end{aligned} \quad (60)$$

where the first inequality is by the triangle inequality, the second inequality is by the arithmetic inequality, the third inequality is by Eq. (59) and our assumption  $\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)$  with  $\zeta = \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}$ .

For the case of linear programming, the strong convexity constant is  $\sigma = 0$ , so that  $A_k = \frac{k}{2\hat{L}m}$  in CLVR. Thus, by applying Theorem 1 with  $\mathcal{W}_k = \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)$  and  $\hat{\mathbf{w}} = (\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \arg \sup_{\mathbf{w}=(\mathbf{x}, \mathbf{y}) \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)} \{\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{y}) - \mathcal{L}(\mathbf{x}, \tilde{\mathbf{y}}_k)\}$ , and using the definition of  $\|\cdot\|_{(\gamma)}$  and Eq. (60), we have

$$\begin{aligned} & \mathbb{E}\left[\sup_{\mathbf{w} \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)} \{\mathcal{L}(\tilde{\mathbf{x}}_k, \mathbf{y}) - \mathcal{L}(\mathbf{x}, \tilde{\mathbf{y}}_k)\}\right] \\ &\leq \frac{2\hat{L}m}{k} \left( \mathbb{E}\left[\frac{\gamma}{2}\|\hat{\mathbf{x}} - \mathbf{x}_0\|^2 + \frac{1}{\gamma}\|\hat{\mathbf{y}} - \mathbf{y}_0\|^2\right] + \frac{\gamma}{2}\|\mathbf{x}^* - \mathbf{x}_0\|^2 + \frac{1}{2\gamma}\|\mathbf{y}^* - \mathbf{y}_0\|^2 \right) \\ &\leq \frac{2\hat{L}m}{k} \left( \mathbb{E}[\|\mathbf{w}_0 - \hat{\mathbf{w}}\|_{(\gamma)}^2] + \frac{1}{2}\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2 \right) \\ &\leq \frac{2\hat{L}m}{k} \left( 12\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2 + \frac{1}{2}\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2 \right) \\ &\leq \frac{25\hat{L}m}{k} \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2. \end{aligned} \quad (61)$$

□

Then with Theorem 1, Lemmas 10 and 11, we give our theorem for the fixed restart strategy.

**Theorem 2.** Consider the CLVR algorithm applied to the standard-form LP problem (PD-LP), with input  $\mathbf{w}_0$  and output  $\tilde{\mathbf{w}}_k$ . Given  $\gamma > 0$ , define  $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}^*} \|\mathbf{w}_0 - \mathbf{w}\|_{(\gamma)}$ , and define

$C_0 = \gamma + 1/\gamma + (\sqrt{2} + 1)\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} + \|\mathbf{w}^*\|_{(\gamma)}$ . Then for  $H_\gamma$  defined as in (9), we have

$$\begin{aligned}\mathbb{E}\left[\sqrt{\text{dist}(\tilde{\mathbf{w}}_k, \mathcal{W}^*)_{(\gamma)}}\right] &\leq 5\sqrt{\frac{\hat{L}mC_0}{H_\gamma k}}\sqrt{\text{dist}(\mathbf{w}_0, \mathcal{W}^*)_{(\gamma)}}, \\ \mathbb{E}\left[\sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)}\right] &\leq 5\sqrt{\frac{\hat{L}mC_0}{H_\gamma k}}\sqrt{\text{LPMetric}(\mathbf{w}_0)}.\end{aligned}$$

*Proof.* Applying Lemma 10 with  $\mathbf{w} = \tilde{\mathbf{w}}_k$ ,  $\tau = \|\tilde{\mathbf{w}}_k\|_{(\gamma)} + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}$  and  $\zeta = \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} \leq \tau$ , we have

$$\sup_{\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)} \{\mathcal{L}(\tilde{\mathbf{x}}_k, \hat{\mathbf{y}}) - \mathcal{L}(\hat{\mathbf{x}}, \tilde{\mathbf{y}}_k)\} \geq \frac{\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} \text{LPMetric}(\tilde{\mathbf{w}}_k)}{\gamma + 1/\gamma + \|\tilde{\mathbf{w}}_k\|_{(\gamma)} + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}}. \quad (62)$$

As  $\frac{x^2}{y}$  is jointly convex in  $x, y$  on the domain  $\{(x, y) : x \in \mathbb{R}, y > 0\}$  [12, Example 3.18] using Jensen's inequality, we have that  $\mathbb{E}\left[\frac{x^2}{y}\right] \geq \frac{(\mathbb{E}[x])^2}{\mathbb{E}[y]}$ . (In our case, this simply follows as  $x, y$  depend on the same source of randomness.) Applying this inequality to (62) with  $x = \sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)}$  and  $y = \gamma + 1/\gamma + \|\tilde{\mathbf{w}}_k\|_{(\gamma)} + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}$ , we obtain

$$\begin{aligned}\mathbb{E}\left[\sup_{\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)} \{\mathcal{L}(\tilde{\mathbf{x}}_k, \hat{\mathbf{y}}) - \mathcal{L}(\hat{\mathbf{x}}, \tilde{\mathbf{y}}_k)\}\right] &\geq \mathbb{E}\left[\frac{\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} (\sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)})^2}{\gamma + 1/\gamma + \|\tilde{\mathbf{w}}_k\|_{(\gamma)} + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}}\right] \\ &\geq \frac{\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} \left(\mathbb{E}[\sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)}]\right)^2}{\mathbb{E}[\gamma + 1/\gamma + \|\tilde{\mathbf{w}}_k\|_{(\gamma)} + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}]}.\end{aligned} \quad (63)$$

Using Lemma 11, we have

$$\begin{aligned}&\mathbb{E}[\gamma + 1/\gamma + \|\tilde{\mathbf{w}}_k\|_{(\gamma)} + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}] \\ &= \gamma + 1/\gamma + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} + \mathbb{E}[\|\tilde{\mathbf{w}}_k\|_{(\gamma)}] \\ &\leq \gamma + 1/\gamma + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} + \mathbb{E}[\|\mathbf{w}^*\|_{(\gamma)} + \|\tilde{\mathbf{w}}_k - \mathbf{w}^*\|_{(\gamma)}] \\ &\leq \gamma + 1/\gamma + \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} + \|\mathbf{w}^*\|_{(\gamma)} + \sqrt{2}\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} \\ &= \gamma + 1/\gamma + (\sqrt{2} + 1)\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} + \|\mathbf{w}^*\|_{(\gamma)}.\end{aligned} \quad (64)$$

By combining Eqs. (63) and (64) and using the definition of  $C_0$ , we have

$$\begin{aligned}\mathbb{E}\left[\sup_{\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)} \{\mathcal{L}(\tilde{\mathbf{x}}_k, \hat{\mathbf{y}}) - \mathcal{L}(\hat{\mathbf{x}}, \tilde{\mathbf{y}}_k)\}\right] &\geq \frac{\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}}{C_0} \left(\mathbb{E}[\sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)}]\right)^2 \\ &\geq \frac{\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}}{C_0} \left(\mathbb{E}[\sqrt{H_\gamma \text{dist}(\tilde{\mathbf{w}}_k, \mathcal{W}^*)_{(\gamma)}}]\right)^2,\end{aligned} \quad (65)$$

where the last inequality is by the definition of Hoffman constant in Eq. (9). Meanwhile, by Lemma 11, we have:

$$\mathbb{E}\left[\sup_{\hat{\mathbf{w}} \in \mathcal{W}_{\zeta, \gamma}(\tilde{\mathbf{w}}_k)} \{\mathcal{L}(\tilde{\mathbf{x}}_k, \hat{\mathbf{y}}) - \mathcal{L}(\hat{\mathbf{x}}, \tilde{\mathbf{y}}_k)\}\right] \leq \frac{25\hat{L}m}{k} \|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)}^2. \quad (66)$$

Now, recalling that  $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}^*} \|\mathbf{w}_0 - \mathbf{w}\|_{(\gamma)}$  and using Eq. (9), we have

$$\|\mathbf{w}_0 - \mathbf{w}^*\|_{(\gamma)} = \text{dist}(\mathbf{w}_0, \mathcal{W}^*)_{(\gamma)} \leq \frac{1}{H_\gamma} \text{LPMetric}(\mathbf{w}_0). \quad (67)$$

By combining Eqs. (65)-(67), we have

$$\begin{aligned}\frac{1}{C_0} \left(\mathbb{E}[\sqrt{H_\gamma \text{dist}(\tilde{\mathbf{w}}_k, \mathcal{W}^*)_{(\gamma)}}]\right)^2 &\leq \frac{1}{C_0} \left(\mathbb{E}[\sqrt{\text{LPMetric}(\tilde{\mathbf{w}}_k)}]\right)^2 \\ &\leq \frac{25\hat{L}m}{k} \text{dist}(\mathbf{w}_0, \mathcal{W}^*)_{(\gamma)} \\ &\leq \frac{25\hat{L}m}{H_\gamma k} \text{LPMetric}(\mathbf{w}_0).\end{aligned} \quad (68)$$

Both bounds follow from this chain of inequalities.  $\square$

Finally, we summarize the adaptive restart strategy in Algorithm 4.

---

**Algorithm 4** Lazy CLVR with Adaptive Restarts

---

- 1: **Input:**  $\epsilon > 0$ ,  $\mathbf{x}_0 \in \mathcal{X}$ ,  $\mathbf{y}_0 \in \mathbb{R}^n$ ,  $\gamma > 0$ ,  $\hat{L} > 0$ ,  $K, \hat{K}, m, \{S^1, S^2, \dots, S^m\}, \{C^1, C^2, \dots, C^m\}$
  - 2:  $t = 0$ ,  $\mathbf{x}_0^{(0)} = \mathbf{x}_0$ ,  $\mathbf{y}_0^{(0)} = \mathbf{y}_0$ ,  $\mathbf{w}^{(0)} = (\mathbf{x}_0^{(0)}, \mathbf{y}_0^{(0)})$
  - 3: **repeat**
  - 4:   Run Lazy CLVR (Algorithm 2) until  $\text{LPMetric}(\mathbf{w}^{(t+1)}) \leq \frac{1}{2} \text{LPMetric}(\mathbf{w}^{(t)})$  where,  $\mathbf{w}^{(t+1)} = \tilde{\mathbf{w}}_K^{(t)} = (\tilde{\mathbf{x}}_K^{(t)}, \tilde{\mathbf{y}}_K^{(t)})$  and  $\tilde{\mathbf{x}}_K^{(t)}, \tilde{\mathbf{y}}_K^{(t)}$  are the output points of Lazy CLVR
  - 5:    $t = t + 1$
  - 6: **until**  $\text{LPMetric}(\mathbf{w}^{(t)}) \leq \epsilon$
  - 7: **Return:**  $\mathbf{w}^{(t)}$
- 

## C Omitted proofs from Section 4

### C.1 DRO with $f$ -divergence-based ambiguity set

**Theorem 3.** *Let  $\mathcal{X}$  be a compact convex set. Then the DRO problem in Eq. (10) is equivalent to*

$$\begin{aligned}
& \min_{\mathbf{x}, \mathbf{u}, \mathbf{v}, \mathbf{w}, \boldsymbol{\mu}, \mathbf{q}, \gamma} \left\{ \gamma + \frac{\rho \mu_1}{n} + \frac{1}{n} \sum_{i=1}^n \mu_i f^* \left( \frac{q_i}{\mu_i} \right) \right\} \\
& \text{s. t. } \mathbf{w} + \mathbf{v} - \frac{\mathbf{q}}{n} - \gamma \mathbf{1}_n = \mathbf{0}_n, \\
& \quad u_i = b_i \mathbf{a}_i^T \mathbf{x}, \quad i \in [n] \\
& \quad \mu_1 = \mu_2 = \dots = \mu_n, \\
& \quad g(u_i) \leq w_i, \quad i \in [n] \\
& \quad q_i \in \mu_i \text{ dom}(f^*), \quad i \in [n] \\
& \quad v_i \geq 0, \mu_i \geq 0, \quad i \in [n] \\
& \quad \mathbf{x} \in \mathcal{X}.
\end{aligned}$$

*Proof.* Since the context is clear, to simplify the notation, in the following we use  $\mathcal{P}$  to denote  $\mathcal{P}_{\rho, n}$ . First, using Sion's minimax theorem, we have that

$$\min_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbf{p} \in \mathcal{P}} \sum_{i=1}^n p_i g(b_i \mathbf{a}_i^T \mathbf{x}) = \sup_{\mathbf{p} \in \mathcal{P}} \min_{\mathbf{x} \in \mathcal{X}} \sum_{i=1}^n p_i g(b_i \mathbf{a}_i^T \mathbf{x}).$$

Introducing auxiliary variables  $\mathbf{w}$  and  $\mathbf{u}$ , the problem is further equivalent to

$$\sup_{\mathbf{p} \in \mathcal{P}} \min_{\substack{\mathbf{x} \in \mathcal{X}, \mathbf{w}, \mathbf{u}: \\ u_i = b_i \mathbf{a}_i^T \mathbf{x}, i \in [n], \\ g(u_i) \leq w_i, i \in [n]}} \mathbf{p}^T \mathbf{w} \equiv \min_{\substack{\mathbf{x} \in \mathcal{X}, \mathbf{w}, \mathbf{u}: \\ u_i = b_i \mathbf{a}_i^T \mathbf{x}, i \in [n], \\ g(u_i) \leq w_i, i \in [n]}} \sup_{\mathbf{p} \in \mathcal{P}} \mathbf{p}^T \mathbf{w},$$

where the last equivalence is by applying minimax equality, which holds due to compactness of  $\mathcal{P}$  [56]. Hence, we can conclude that

$$\min_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbf{p} \in \mathcal{P}} \sum_{i=1}^n p_i g(b_i \mathbf{a}_i^T \mathbf{x}) = \min_{\substack{\mathbf{x} \in \mathcal{X}, \mathbf{w}, \mathbf{u}: \\ u_i = b_i \mathbf{a}_i^T \mathbf{x}, i \in [n], \\ g(u_i) \leq w_i, i \in [n]}} \sup_{\mathbf{p} \in \mathcal{P}} \mathbf{p}^T \mathbf{w}. \quad (69)$$

For a fixed tuple  $(\mathbf{x}, \mathbf{w}, \mathbf{u})$ , using Lagrange multipliers to enforce the constraints from  $\mathcal{P}$ , we can further write

$$\begin{aligned} \sup_{\mathbf{p} \in \mathcal{P}} \mathbf{p}^T \mathbf{w} &= - \inf_{\mathbf{p} \in \mathcal{P}} -\mathbf{p}^T \mathbf{w} \\ &= - \inf_{\mathbf{p} \in \mathbb{R}^n} \left( -\mathbf{p}^T \mathbf{w} + \sup_{\mathbf{v} \geq 0, \gamma \in \mathbb{R}, \mu \geq 0} \left( -\mathbf{p}^T \mathbf{v} + \gamma \left( \sum_{i=1}^n p_i - 1 \right) + \mu \left( D_f(\mathbf{p} \parallel \mathbf{1}/n) - \frac{\rho}{n} \right) \right) \right) \\ &= \sup_{\mathbf{p} \in \mathbb{R}^n} \inf_{\mathbf{v} \geq 0, \gamma \in \mathbb{R}, \mu \geq 0} \left( \mathbf{p}^T \mathbf{w} + \mathbf{p}^T \mathbf{v} - \gamma \left( \sum_{i=1}^n p_i - 1 \right) - \mu \left( D_f(\mathbf{p} \parallel \mathbf{1}/n) - \frac{\rho}{n} \right) \right). \end{aligned}$$

Now, using the definitions of  $D_f$  and the convex conjugate of  $f$ , we have

$$D_f(\mathbf{p} \parallel \mathbf{1}/n) = \sum_{i=1}^n \frac{1}{n} f(np_i) = \sum_{i=1}^n \frac{1}{n} \sup_{\nu_i \in \text{dom}(f^*)} (np_i \nu_i - f^*(\nu_i)). \quad (70)$$

As a result, we have

$$\begin{aligned} &\inf_{\mu \geq 0} -\mu \left( D_f(\mathbf{p} \parallel \mathbf{1}/n) - \frac{\rho}{n} \right) \\ &= \inf_{\mu \geq 0} -\frac{\mu}{n} \left( \sum_{i=1}^n \sup_{\nu_i \in \text{dom}(f^*), i \in [n]} (np_i \nu_i - f^*(\nu_i)) - \rho \right) \\ &= \inf_{\mu \geq 0, \nu_i \in \text{dom}(f^*), i \in [n]} -\frac{\mu}{n} \left( \sum_{i=1}^n (np_i \nu_i - f^*(\nu_i)) - \rho \right) \\ &= \inf_{\mu \geq 0, q_i \in \mu \text{ dom}(f^*), i \in [n]} \frac{1}{n} \sum_{i=1}^n \left( -p_i q_i + \mu f^*\left(\frac{q_i}{n\mu}\right) \right) + \frac{\mu\rho}{n} \\ &= \inf_{\substack{\mu_1 = \mu_2 = \dots = \mu_n \geq 0, \\ q_i \in \mu_i \text{ dom}(f^*), i \in [n]}} \frac{1}{n} \sum_{i=1}^n \left( -p_i q_i + \mu_i f^*\left(\frac{q_i}{n\mu_i}\right) \right) + \frac{\mu_1 \rho}{n}, \end{aligned}$$

where the first equality is by Eq. (70), the third one is by the variable substitution  $q_i = n\mu\nu_i$ , and the last one is by introducing  $\mu_1, \mu_2, \dots, \mu_n$  to replace  $\mu$ . Then each  $n\mu_i f^*\left(\frac{q_i}{n\mu_i}\right) (i \in [n])$  is a perspective function of  $f^*$  [12], which is jointly convex w.r.t.  $(\mu_i, q_i)$ . Hence, we can conclude that

$$\begin{aligned} &\sup_{\mathbf{p} \in \mathcal{P}} \mathbf{p}^T \mathbf{w} \\ &= \sup_{\mathbf{p} \in \mathbb{R}^n} \inf_{\substack{\gamma \in \mathbb{R}, \mathbf{v} \geq 0, \\ \mu_1 = \mu_2 = \dots = \mu_n \geq 0, \\ q_i \in \mu_i \text{ dom}(f^*), i \in [n]}} \left( \mathbf{p}^T \mathbf{w} + \mathbf{p}^T \mathbf{v} - \gamma \left( \sum_{i=1}^n p_i - 1 \right) + \frac{1}{n} \sum_{i=1}^n \left( -p_i q_i + \mu_i f^*\left(\frac{q_i}{n\mu_i}\right) \right) + \frac{\mu_1 \rho}{n} \right) \\ &= \inf_{\substack{\gamma \in \mathbb{R}, \mathbf{v} \geq 0, \\ \mu_1 = \mu_2 = \dots = \mu_n \geq 0, \\ q_i \in \mu_i \text{ dom}(f^*), i \in [n]}} \sup_{\mathbf{p} \in \mathbb{R}^n} \left( \mathbf{p}^T \mathbf{w} + \mathbf{p}^T \mathbf{v} - \gamma \left( \sum_{i=1}^n p_i - 1 \right) + \frac{1}{n} \sum_{i=1}^n \left( -p_i q_i + \mu_i f^*\left(\frac{q_i}{n\mu_i}\right) \right) + \frac{\mu_1 \rho}{n} \right), \end{aligned}$$

where the last line is by strong duality. Thus, combining with Eq. (69), we conclude that the original DRO problem with  $f$ -divergence based ambiguity set is equivalent to the following problem:

$$\begin{aligned} &\min_{\mathbf{x}, \mathbf{u}, \mathbf{v}, \mathbf{w}, \boldsymbol{\mu}, \mathbf{q}, \gamma} \max_{\mathbf{p} \in \mathbb{R}^n} \left\{ \gamma + \frac{\rho\mu_1}{n} + \frac{1}{n} \sum_{i=1}^n \mu_i f^*\left(\frac{q_i}{\mu_i}\right) + \mathbf{p}^T \left( \mathbf{w} + \mathbf{v} - \frac{\mathbf{q}}{n} - \gamma \mathbf{1}_n \right) \right\} \\ &\text{s. t. } u_i = \mathbf{b}_i \mathbf{a}_i^T \mathbf{x}, \quad i \in [n] \\ &\quad \mu_1 = \mu_2 = \dots = \mu_n, \\ &\quad g(u_i) \leq w_i, \quad i \in [n] \\ &\quad q_i \in \mu_i \text{ dom}(f^*), \quad i \in [n] \\ &\quad v_i \geq 0, \mu_i \geq 0, \quad i \in [n] \\ &\quad \mathbf{x} \in \mathcal{X}. \end{aligned}$$

Finally, noticing that the maximization problem over  $\mathbf{p} \in \mathbb{R}^n$  enforces the equality constraint  $\mathbf{w} + \mathbf{v} - \frac{\mathbf{q}}{n} - \gamma \mathbf{1}_n = \mathbf{0}_n$ , we obtain

$$\begin{aligned} & \min_{\mathbf{x}, \mathbf{u}, \mathbf{v}, \boldsymbol{\mu}, \mathbf{q}, \gamma} \left\{ \gamma + \frac{\rho \mu_1}{n} + \frac{1}{n} \sum_{i=1}^n \mu_i f^* \left( \frac{q_i}{\mu_i} \right) \right\} \\ \text{s. t. } & \mathbf{w} + \mathbf{v} - \frac{\mathbf{q}}{n} - \gamma \mathbf{1}_n = \mathbf{0}_n, \\ & u_i = b_i \mathbf{a}_i^T \mathbf{x}, \quad i \in [n] \\ & \mu_1 = \mu_2 = \dots = \mu_n, \\ & g(u_i) \leq w_i, \quad i \in [n] \\ & q_i \in \mu_i \text{ dom}(f^*), \quad i \in [n] \\ & v_i \geq 0, \mu_i \geq 0, \quad i \in [n] \\ & \mathbf{x} \in \mathcal{X}, \end{aligned}$$

as claimed.  $\square$

**Example: Conditional Value at Risk (CVaR) with hinge loss.** As a specific example of an application of Theorem 3, we consider CVaR at level  $\alpha \in (0, 1)$ , which leads to the optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \sup_{\substack{\mathbf{p} \geq 0, \mathbf{1}_n^T \mathbf{p} = 1 \\ p_i \leq \frac{1}{\alpha n} (i \in [d])}} \sum_{i=1}^n p_i g(b_i \mathbf{a}_i^T \mathbf{x}), \quad (71)$$

where  $g(u_i) = \max\{0, 1 - u_i\}$  is the hinge loss. Here the ambiguity set constraint reduces to simple bounds  $p_i \leq \frac{1}{\alpha n}$  for  $i \in [n]$ , so the reformulation based on the convex perspective function can be avoided altogether and replaced by simple Lagrange multipliers for this linear constraint. In particular, in the proof of Theorem 3, we can write

$$\begin{aligned} \sup_{\mathbf{p} \in \mathcal{P}} \mathbf{p}^T \mathbf{w} &= \sup_{\mathbf{p} \in \mathbb{R}^n} \inf_{\mathbf{v} \geq 0, \gamma \in \mathbb{R}, \boldsymbol{\mu} \geq 0} \left( \mathbf{p}^T \mathbf{w} + \mathbf{p}^T \mathbf{v} - \gamma \left( \sum_{i=1}^n p_i - 1 \right) - \sum_{i=1}^n \mu_i \left( p_i - \frac{1}{\alpha n} \right) \right) \\ &= \sup_{\mathbf{p} \in \mathbb{R}^n} \inf_{\mathbf{v} \geq 0, \gamma \in \mathbb{R}, \boldsymbol{\mu} \geq 0} \left( \gamma + \mathbf{p}^T (\mathbf{w} + \mathbf{v} - \gamma \mathbf{1}_n - \boldsymbol{\mu}) + \frac{1}{\alpha n} \boldsymbol{\mu}^T \mathbf{1}_n \right) \\ &= \inf_{\mathbf{v} \geq 0, \gamma \in \mathbb{R}, \boldsymbol{\mu} \geq 0} \sup_{\mathbf{p} \in \mathbb{R}^n} \left( \gamma + \mathbf{p}^T (\mathbf{w} + \mathbf{v} - \gamma \mathbf{1}_n - \boldsymbol{\mu}) + \frac{1}{\alpha n} \boldsymbol{\mu}^T \mathbf{1}_n \right). \end{aligned}$$

Following the argument from the proof of Theorem 3, and expressing  $w_i \geq g(u_i)$  equivalently as the pair of constraints  $w_i \geq 0$  and  $w_i \geq 1 - u_i$ , the problem reduces to

$$\begin{aligned} & \min_{\mathbf{x}, \mathbf{u}, \mathbf{v}, \mathbf{w}, \boldsymbol{\mu}, \gamma} \left\{ \gamma + \frac{1}{\alpha n} \boldsymbol{\mu}^T \mathbf{1}_n \right\} \\ \text{s. t. } & \mathbf{w} + \mathbf{v} - \gamma \mathbf{1}_n - \boldsymbol{\mu} = \mathbf{0}_n, \\ & u_i = b_i \mathbf{a}_i^T \mathbf{x}, \quad i \in [n], \\ & w_i \geq 0, w_i \geq 1 - u_i, \quad i \in [n], \\ & v_i \geq 0, \mu_i \geq 0, \quad i \in [n], \end{aligned} \quad (72)$$

which is a linear program. To write it in the standard form, we further introduce slack variables  $\mathbf{s} \in \mathbb{R}^n$ ,  $\mathbf{s} \geq \mathbf{0}$ , to replace inequality constraints  $w_i \geq 1 - u_i$  by  $s_i - u_i - w_i = -1$ . For implementation purposes, we define  $\mathcal{X}$  to be the set of simple non-negativity constraints ( $w_i \geq 0$ ,  $s_i \geq 0$ ,  $v_i \geq 0$ ,  $\mu_i \geq 0$ ,  $\forall i \in [n]$ ) from Eq. (72). The problem then becomes

$$\begin{aligned} & \min_{\mathbf{x}, \mathbf{u}, \mathbf{v}, \mathbf{w}, \boldsymbol{\mu}, \mathbf{s}, \gamma} \left\{ \gamma + \frac{1}{\alpha n} \boldsymbol{\mu}^T \mathbf{1}_n \right\} \\ \text{s. t. } & \mathbf{w} + \mathbf{v} - \gamma \mathbf{1}_n - \boldsymbol{\mu} = \mathbf{0}_n, \\ & u_i - b_i \mathbf{a}_i^T \mathbf{x} = 0, \quad i \in [n], \\ & \mathbf{s} - \mathbf{u} - \mathbf{w} = -\mathbf{1}_n, \\ & \mathbf{w} \geq \mathbf{0}_n, \mathbf{v} \geq \mathbf{0}_n, \boldsymbol{\mu} \geq \mathbf{0}_n, \mathbf{s} \geq \mathbf{0}_n. \end{aligned}$$

## C.2 DRO with Wasserstein metric-based ambiguity set

**Definition 1.** Let  $\mu$  and  $\nu$  be two probability distributions supported on  $\Theta = \mathbb{R}^d \times \{1, -1\}$  and let  $\Pi(\mu, \nu)$  denote the set of all joint distributions between  $\mu$  and  $\nu$ . Then the Wasserstein metric between  $\mu$  and  $\nu$  is defined by

$$W(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{\Theta \times \Theta} \zeta(\xi, \xi') \pi(d\xi, d\xi'), \quad (73)$$

where  $\xi, \xi' \in \Theta$  and  $\zeta(\cdot, \cdot) : \Theta \times \Theta \rightarrow \mathbb{R}_+$  is a convex cost function defined by

$$\zeta((\mathbf{a}, b), (\mathbf{a}', b')) = \|\mathbf{a} - \mathbf{a}'\| + \kappa|b - b'|,$$

where  $\|\cdot\|$  denotes a general norm and  $\kappa > 0$  is used to balance the feature mismatch and label uncertainty. Let  $\mathbb{Q} = \frac{1}{n} \sum_{i=1}^n \delta_{(\mathbf{a}_i, b_i)}$ , where  $\delta_{(\mathbf{a}_i, b_i)}$  is the Dirac Delta function (or a point mass) at point  $(\mathbf{a}_i, b_i)$ . Then, the Wasserstein metric based ambiguity set is defined as

$$\mathcal{P}_{\rho, \kappa} = \left\{ \mathbb{P} : W(\mathbb{P}, \mathbb{Q}) \leq \rho \right\}. \quad (74)$$

**Assumption 5.**  $\mathbb{R} \ni M = \sup_{\theta \in \text{dom}(g^*)} |\theta|$ .

We first show the following auxiliary lemma which is used in the proof of Theorem 4. A similar result for the special case of a logistic loss function can be found in [52].

**Lemma 12.** Let  $(\mathbf{a}', b') \in \mathbb{R}^d \times \{1, -1\}$  be a given pair of data sample and label. Then, for every  $\lambda > 0$ , we have

$$\sup_{\mathbf{a} \in \mathbb{R}^d} g(b' \mathbf{a}^T \mathbf{w}) - \lambda \|\mathbf{a} - \mathbf{a}'\| = \begin{cases} g(b' \mathbf{a}'^T \mathbf{w}), & \text{if } \|\mathbf{w}\|_* \leq \lambda/M \\ +\infty, & \text{otherwise} \end{cases}. \quad (75)$$

*Proof.* Since  $g$  is assumed to be proper, convex, and lower semicontinuous, by the Fenchel-Moreau theorem, it is equal to its biconjugate. Applying this property, we have

$$\begin{aligned} & \sup_{\mathbf{a} \in \mathbb{R}^d} \{g(b' \mathbf{a}^T \mathbf{w}) - \lambda \|\mathbf{a} - \mathbf{a}'\|\} \\ &= \sup_{\substack{\mathbf{a} \in \mathbb{R}^d, \\ \theta \in \text{dom}(g^*)}} \{\theta b' \mathbf{a}^T \mathbf{w} - g^*(\theta) - \lambda \|\mathbf{a} - \mathbf{a}'\|\} \\ &= \sup_{\substack{\mathbf{a} \in \mathbb{R}^d, \\ \theta \in \text{dom}(g^*)}} \{\theta b' (\mathbf{a} - \mathbf{a}')^T \mathbf{w} + \theta b' (\mathbf{a}')^T \mathbf{w} - g^*(\theta) - \lambda \|\mathbf{a} - \mathbf{a}'\|\}. \end{aligned}$$

Applying the change of variable  $\mathbf{v} := \mathbf{a} - \mathbf{a}'$ , we further have

$$\begin{aligned} & \sup_{\mathbf{a} \in \mathbb{R}^d} \{g(b' \mathbf{a}^T \mathbf{w}) - \lambda \|\mathbf{a} - \mathbf{a}'\|\} \\ &= \sup_{\substack{\mathbf{v} \in \mathbb{R}^d, \\ \theta \in \text{dom}(g^*)}} \{\theta b' \mathbf{v}^T \mathbf{w} + \theta b' (\mathbf{a}')^T \mathbf{w} - g^*(\theta) - \lambda \|\mathbf{v}\|\} \\ &= \sup_{\theta \in \text{dom}(g^*)} \{\mathbb{1}_{\{\|\theta b' \mathbf{w}\|_* \leq \lambda\}} + \theta b' (\mathbf{a}')^T \mathbf{w} - g^*(\theta)\} \\ &= \begin{cases} g(b' (\mathbf{a}')^T \mathbf{w}), & \text{if } \sup_{\theta \in \text{dom}(g^*)} \|\theta b' \mathbf{w}\|_* \leq \lambda \\ +\infty, & \text{otherwise} \end{cases}, \end{aligned}$$

where the second equality is by the convex conjugate of a norm  $\|\cdot\|$  being equal to the convex indicator of the unit ball w.r.t. the dual norm  $\|\cdot\|_*$ . Finally, it remains to use that, by Assumption 5,  $\sup_{\theta \in \text{dom}(g^*)} |\theta| = M$ .  $\square$

Then following [52, Theorem 1], we provide reformulations of the problem from Eq. (11) that can be addressed by computationally efficient solvers.

**Theorem 4.** *The optimization problem from Eq. (11) is equivalent to:*

$$\begin{aligned}
\min_{\mathbf{w}, \lambda, \mathbf{u}, \mathbf{v}, \mathbf{s}, \mathbf{t}} \quad & \rho\lambda + \frac{1}{n} \sum_{i=1}^n s_i \\
\text{s. t.} \quad & u_i = b_i \mathbf{a}_i^T \mathbf{w}, \quad i \in [n], \\
& v_i = -u_i, \quad i \in [n], \\
& t_i = 2\kappa\lambda + s_i, \quad i \in [n], \\
& g(u_i) \leq s_i, \quad i \in [n], \\
& g(v_i) \leq t_i, \quad i \in [n], \\
& \|\mathbf{w}\|_* \leq \lambda/M.
\end{aligned} \tag{76}$$

*Proof.* Let  $\mathbf{z} = (\mathbf{a}, b) \in \Theta := \mathbb{R}^d \times \{1, -1\}$  and let  $h_{\mathbf{w}}(\mathbf{z}) := g(\mathbf{b}\mathbf{a}^T \mathbf{w})$ . Then by the definition of the Wasserstein metric,

$$\sup_{\mathbb{P} \in \mathcal{P}_{\rho, \kappa}} \mathbb{E}^{\mathbb{P}}[g(\mathbf{b}\mathbf{a}^T \mathbf{w})] = \begin{cases} \sup_{\pi \in \Pi(\mathbb{P}, \hat{\mathbb{P}}_n)} \int_{\Theta} h_{\mathbf{w}}(\mathbf{z}) \pi(d\mathbf{z}, \Theta) \\ \text{s. t.} \quad \int_{\Theta \times \Theta} \zeta(\mathbf{z}, \mathbf{z}') \pi(d\mathbf{z}, d\mathbf{z}') \leq \rho. \end{cases} \tag{77}$$

Assume that the conditional distribution of  $\mathbf{z}$  given  $\mathbf{z}' = (\mathbf{a}_i, b_i)$  is  $\mathbb{P}^i$ , for all  $i \in [n]$ . Then, based on the definition of  $\hat{\mathbb{P}}_n = \frac{1}{n} \sum_{i=1}^n \delta_{(\mathbf{a}_i, b_i)}$ , we have

$$\pi(d\mathbf{z}, d\mathbf{z}') = \frac{1}{n} \sum_{i=1}^n \delta_{(\mathbf{a}_i, b_i)} \mathbb{P}^i(d\mathbf{z}). \tag{78}$$

As a result, the problem from Eq. (77) is equivalent to

$$\sup_{\mathbb{P} \in \mathcal{P}_{\rho, \kappa}} \mathbb{E}^{\mathbb{P}}[g(\mathbf{b}\mathbf{a}^T \mathbf{w})] = \begin{cases} \sup_{\mathbb{P}^i} \frac{1}{n} \sum_{i=1}^n \int_{\Theta} h_{\mathbf{w}}(\mathbf{z}) \mathbb{P}^i(d\mathbf{z}) \\ \text{s. t.} \quad \frac{1}{n} \sum_{i=1}^n \int_{\Theta} \zeta(\mathbf{z}, \mathbf{z}') \mathbb{P}^i(d\mathbf{z}) \leq \rho \\ \int_{\Theta} \mathbb{P}^i(d\mathbf{z}) = 1. \end{cases} \tag{79}$$

Then substituting in  $\mathbf{z} = (\mathbf{a}, b)$ , using that the domain of  $y$  is  $\{1, -1\}$ , and decomposing  $\mathbb{P}^i$  into unnormalized measures  $\mathbb{P}_{\pm 1}^i(d\mathbf{a}) = \mathbb{P}^i(d\mathbf{a}, \{b = \pm 1\})$  supported on  $\mathbb{R}^d$ , the RHS of Eq. (79) can be simplified to

$$\begin{aligned}
& \sup_{\mathbb{P}_{\pm}^i} \frac{1}{n} \sum_{i=1}^n \int_{\mathbb{R}^d} (h_{\mathbf{w}}(\mathbf{a}, 1) \mathbb{P}_1^i(d\mathbf{a}) + h_{\mathbf{w}}(\mathbf{a}, -1) \mathbb{P}_{-1}^i(d\mathbf{a})) \\
& \text{s. t.} \quad \frac{1}{n} \sum_{i=1}^n \int_{\mathbb{R}^d} (\zeta((\mathbf{a}, 1), (\mathbf{a}_i, \mathbf{b}_i)) \mathbb{P}_1^i(d\mathbf{a}) + \zeta((\mathbf{a}, -1), (\mathbf{a}_i, \mathbf{b}_i)) \mathbb{P}_{-1}^i(d\mathbf{a})) \leq \rho \\
& \int_{\mathbb{R}^d} (\mathbb{P}_1^i(d\mathbf{a}) + \mathbb{P}_{-1}^i(d\mathbf{a})) = 1.
\end{aligned}$$

With the definition of the cost function  $\zeta((\mathbf{a}, b), (\mathbf{a}', b')) = \|\mathbf{a} - \mathbf{a}'\| + \kappa|b - b'|$ , it follows that

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n \int_{\mathbb{R}^d} \zeta((\mathbf{a}, 1), (\mathbf{a}_i, \mathbf{b}_i)) \mathbb{P}_1^i(d\mathbf{a}) + \zeta((\mathbf{a}, -1), (\mathbf{a}_i, \mathbf{b}_i)) \mathbb{P}_{-1}^i(d\mathbf{a}) \\
& = \frac{1}{n} \int_{\mathbb{R}^d} \sum_{b_i=1} [ \|\mathbf{a} - \mathbf{a}_i\| \mathbb{P}_1^i(d\mathbf{a}) + \|\mathbf{a} - \mathbf{a}_i\| \mathbb{P}_{-1}^i(d\mathbf{a}) + 2\kappa \mathbb{P}_{-1}^i(d\mathbf{a}) ] \\
& \quad + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{b_i=-1} [ \|\mathbf{a} - \mathbf{a}_i\| \mathbb{P}_{-1}^i(d\mathbf{a}) + \|\mathbf{a} - \mathbf{a}_i\| \mathbb{P}_1^i(d\mathbf{a}) + 2\kappa \mathbb{P}_1^i(d\mathbf{a}) ] \\
& = \frac{2\kappa}{n} \int_{\mathbb{R}^d} \left( \sum_{b_i=1} \mathbb{P}_{-1}^i(d\mathbf{a}) + \sum_{b_i=-1} \mathbb{P}_1^i(d\mathbf{a}) \right) + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i=1}^n \|\mathbf{a} - \mathbf{a}_i\| (\mathbb{P}_{-1}^i(d\mathbf{a}) + \mathbb{P}_1^i(d\mathbf{a})). \tag{80}
\end{aligned}$$

Thus, we have

$$\sup_{\mathbb{P} \in \mathcal{P}_{\rho, \kappa}} \mathbb{E}^{\mathbb{P}}[g(\mathbf{b}\mathbf{a}^T \mathbf{w})] = \begin{cases} \sup_{\mathbb{P}_{\pm 1}^i} \frac{1}{n} \sum_{i=1}^n \int_{\mathbb{R}^d} (h_{\mathbf{w}}(\mathbf{a}, 1) \mathbb{P}_1^i(d\mathbf{a}) + h_{\mathbf{w}}(\mathbf{a}, -1) \mathbb{P}_{-1}^i(d\mathbf{a})) \\ \text{s. t. } \frac{2\kappa}{n} \int_{\mathbb{R}^d} \left( \sum_{b_i=1} \mathbb{P}_{-1}^i(d\mathbf{a}) + \sum_{b_i=-1} \mathbb{P}_1^i(d\mathbf{a}) \right) \\ \quad + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i=1}^n \|\mathbf{a} - \mathbf{a}_i\| (\mathbb{P}_{-1}^i(d\mathbf{a}) + \mathbb{P}_1^i(d\mathbf{a})) \leq \rho \\ \int_{\mathbb{R}^d} (\mathbb{P}_1^i(d\mathbf{a}) + \mathbb{P}_{-1}^i(d\mathbf{a})) = 1. \end{cases}$$

The above problem w.r.t.  $\mathbb{P}_{\pm 1}^i$  is an infinite-dimensional linear program with a finite number of constraints. By [53, Proposition 3.4], we get the following equivalent dual formulation:

$$\sup_{\mathbb{P} \in \mathcal{P}_{\rho, \kappa}} \mathbb{E}^{\mathbb{P}}[g(\mathbf{b}\mathbf{a}^T \mathbf{w})] = \begin{cases} \min_{\lambda, s_i} \rho\lambda + \frac{1}{n} \sum_{i=1}^n s_i \\ \text{s. t. } \sup_{\mathbf{a} \in \mathbb{R}^d} h_{\mathbf{w}}(\mathbf{a}, 1) - \lambda \|\mathbf{a} - \mathbf{a}_i\| - \lambda\kappa(1 - b_i) \leq s_i, \quad i \in [n] \\ \sup_{\mathbf{a} \in \mathbb{R}^d} h_{\mathbf{w}}(\mathbf{a}, -1) - \lambda \|\mathbf{a} - \mathbf{a}_i\| - \lambda\kappa(1 + b_i) \leq s_i, \quad i \in [n] \\ \lambda \geq 0. \end{cases}$$

Then, recalling that  $h_{\mathbf{w}}(\mathbf{a}, \pm 1)$  is short for  $g(\pm \mathbf{a}^T \mathbf{w})$ , by Lemma 12, we have

$$\sup_{\mathbf{a} \in \mathbb{R}^d} h_{\mathbf{w}}(\mathbf{a}, \pm 1) - \lambda \|\mathbf{a} - \mathbf{a}_i\| = \begin{cases} g(\pm \mathbf{a}_i^T \mathbf{w}), & \text{if } \sup_{\theta \in \text{dom}(g^*)} \|\theta \mathbf{w}\|_* \leq \lambda \\ +\infty, & \text{otherwise.} \end{cases}$$

Finally, the resulting reformulation is

$$\begin{aligned} \min_{\mathbf{w}, \lambda, \mathbf{s}} \quad & \rho\lambda + \frac{1}{n} \sum_{i=1}^n s_i \\ \text{s. t.} \quad & g(b_i \mathbf{a}_i^T \mathbf{w}) \leq s_i, \quad i \in [n], \\ & g(-b_i \mathbf{a}_i^T \mathbf{w}) - 2\lambda\kappa \leq s_i, \quad i \in [n], \\ & \sup_{\theta \in \text{dom}(g^*)} \|\theta \mathbf{w}\|_* \leq \lambda. \end{aligned}$$

Finally, recalling that, by assumption,  $\sup_{\theta \in \text{dom}(g^*)} |\theta| = M$ , it follows that the constraint  $\sup_{\theta \in \text{dom}(g^*)} \|\theta \mathbf{w}\|_* \leq \lambda$  is equivalent to  $\|\mathbf{w}\|_* \leq \lambda/M$ . Meanwhile, by introducing  $u_i = b_i \mathbf{a}_i^T \mathbf{w}$ ,  $v_i = -u_i$  ( $i \in [n]$ ) and  $s_i, t_i$  ( $i \in [n]$ ), we obtain Theorem 4.  $\square$

In Theorem 4, when we assume that the conic constraints  $g(u) \leq s$  and  $\|\mathbf{w}\|_* \leq \lambda/M$  in Eq. (76) admit efficient proximal operators, we can formulate this problem as (PD-GLP) and apply CLVR. The resulting complexity bounds are similar to those discussed above for the  $f$ -divergence formulation.

## D Experiment details

### D.1 Comparison of adaptive restart schemes

We provide a brief empirical comparison between our adaptive restart scheme that uses LPMetric and the adaptive restart scheme using the normalized duality gap proposed in [10]. We compared the performance of PDHG on benchmark problem sets qap10, qap15, nug08, and nug20 used in [10], using the two adaptive restart criteria. We ran PDHG until reaching accuracy as described in [10] (that is, until normalized duality gap is at most  $10^{-6}$  and primal and dual infeasibility is at most  $10^{-8}$ ).

Table 4: Number of iterations required for the normalized duality gap and primal and dual infeasibility to fall below  $10^{-6}$  and  $10^{-8}$ , respectively.

Problem Name	Adaptive Normalized Duality Gap	Adaptive LPMetric
qap10	<b>13041</b>	14521
qap15	12561	<b>961</b>
nug08	<b>841</b>	1481
nug20	22001	<b>16281</b>

Table 4 shows that the two restart criteria give similar performance in terms of iteration complexity. Normalized gap is better on qap10 and nug08, while LPMetric is better on qap15 and nug20. For further details, Figure 2 plots the normalized duality gap vs iteration count. The two adaptive restart schemes lead to similar performance of PDHG over iterations. Comparisons based on wall-clock time are shown in Figure 3; the behavior is similar. We conclude that our restart criterion based on LPMetric seems comparable with normalized duality gap, in terms of iteration complexity.

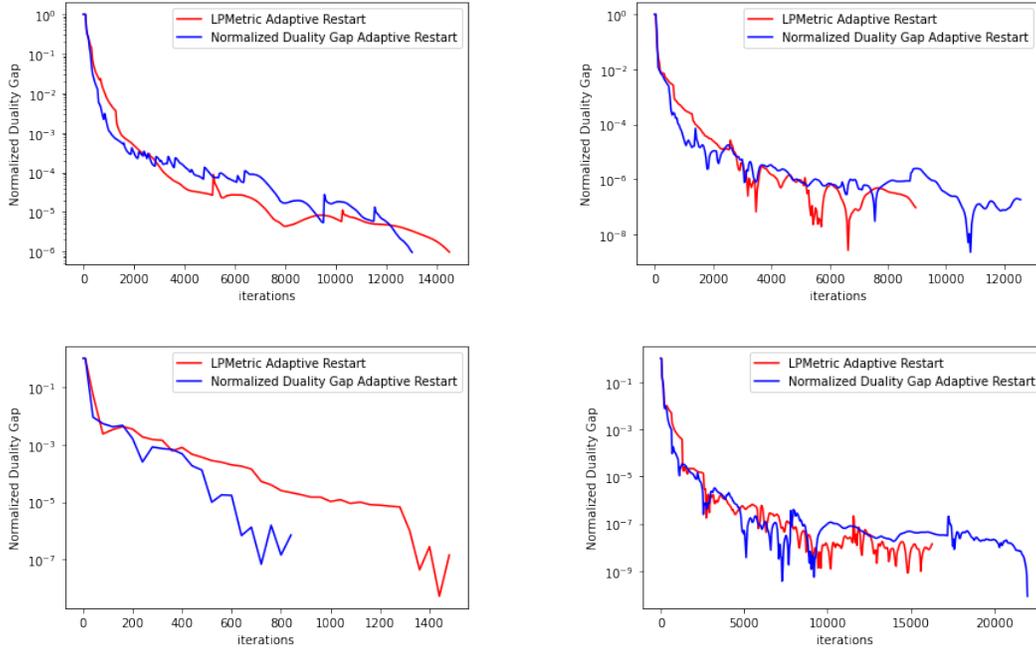


Figure 2: Comparisons of restart schemes that use LPMetric and that use the normalized duality gap against number of iterations. The plots from left to right and then top to bottom are for qap10, qap15, nug08, and nug20.

## D.2 Details of experiments in Section 5

When we consider the DRO problem with Wasserstein metric of  $\ell_1$ -norm and with hinge loss, we observe that the reformulation described in Theorem 4 can be further reformulated into an ordinary LP, as the dual norm of  $\ell_1$  norm is  $\ell_\infty$  norm and hinge loss can be decomposed linearly with additional auxiliary variables. Thus in the following instances we consider, we apply our adaptive restart scheme with respect to LPMetric as illustrated in Section 3.3 to achieve heuristic linear convergence rate in terms of the number of data passes. We compare our CLVR method with three representative methods: PDHG [15], SPDHG [16] and PURE-CD [5]. We implemented CLVR and other algorithms in Julia, optimizing all implementations to the best of our ability.<sup>8</sup> For SPDHG, whose per-iteration cost is at least  $O(d)$ , we consider a large batch size of 50 to balance the effect of the  $O(d)$  cost and improve

<sup>8</sup>Julia is particularly designed for high performance numerical computation.

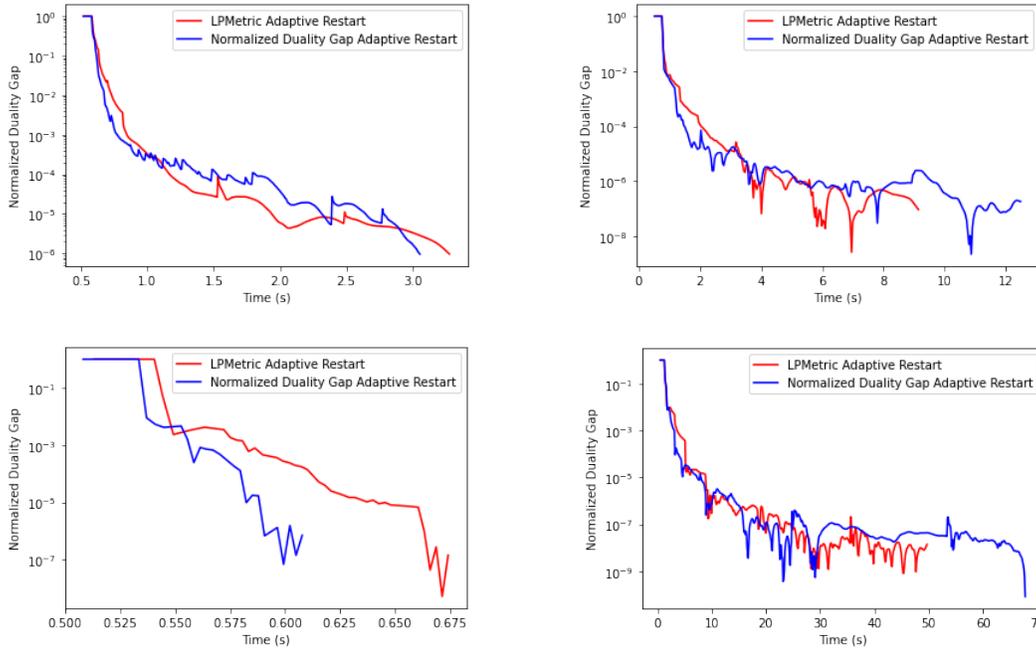


Figure 3: Comparisons of restart schemes that use LPMetric and that use the normalized duality gap against wall clock time in terms of seconds. The plots from left to right and then top to bottom are for qap10, qap15, nug08, and nug20.

the overall efficiency. Meanwhile, PURE-CD with block size 1 is already well suited to sparsity. For CLVR, we experiment with block sizes 1 and 10.

We conducted our experiments on LibSVM [17] datasets `a9a`, `gisette`, and `rcv1.binary`, each with different sparsity levels. We run each algorithm using one CPU core, on a Linux machine with a second generation Intel Xeon Scalable Processor (Cascade Lake-SP) with 128 GB of RAM. Because the weight parameter  $\gamma$  between primal and dual variables (see Theorem 1) strongly influences empirical performance, we tune it for all datasets by trying the values  $\{10^{-i}\}$  for  $i \in \mathbb{Z}$ , for each of the methods. We set the Lipschitz constant of PDHG to be the largest singular value of the constraint matrix in the LP formulation. For PURE-CD and CLVR with block size 1, because the rows of the matrix are normalized, we set the Lipschitz constant to 1. For CLVR with block size 10 and SPDHG with block size 50, the Lipschitz constants are tuned to 3 and 9, respectively.

**Remark 3** (Comparisons of using different block sizes). *We conducted experiments to compare the practical performance of CLVR against different choices of block sizes, with results shown in Figure 4. We ran the DRO with Wasserstein metric using the same setup as described in Section 5, on the `rcv1` dataset and using an early stopping criterion of LPMetric at  $10^{-1}$ . In the plot, we can see that CLVR converges to an approximate solution fastest when the block size is set to 10, providing support for our choice of 10 in Section 5. As illustrated in Figure 1, CLVR is most efficient in terms of the number of data passes when the block size is 1, but in terms of the execution time, running CLVR with larger block size yields better performance. We attribute this phenomenon to the instruction-level parallelism [28] in modern processors, allowing more computations to be completed in the same number of clock cycles.*

In Table 5, we list information about the three datasets and the corresponding matrices in the reformulations. As we see, due to the sparse connectivity of auxiliary variables, all the matrices in reformulations are quite sparse. As a commonly adopted preprocessing step for LP, we normalize the matrix in the standard-form LP so that each row has Euclidean norm 1.

**Remark 4** (Performance comparison using multiple cores). *We conducted further experiments to examine the effects of allowing the algorithms to run on more computing cores. However, we did not observe any meaningful difference in performances in terms of wall-clock time when we repeated the*

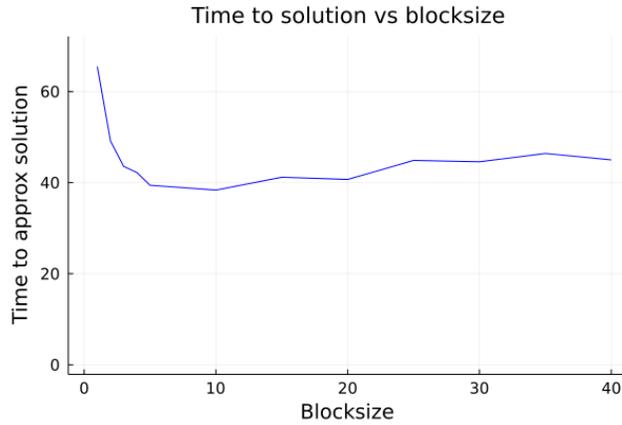


Figure 4: Comparison of the performance of CLVR with various choices of blocksizes.

Table 5: The dimension and sparsity of the original datasets and the corresponding matrices in reformulations.

Dataset	Original $(d, n)$	#nonzeros / $(d \times n)$	Reformulated $(d, n)$	#nonzeros / $(d \times n)$
<b>a9a</b>	(123, 32561)	0.11	(130738, 97929)	$9.6 \times 10^{-5}$
<b>gisette</b>	(5000, 6000)	0.99	(44002, 28000)	$4.9 \times 10^{-2}$
<b>rcv1</b>	(47236, 20242)	$1.5 \times 10^{-3}$	(269914, 155198)	$8.8 \times 10^{-5}$

*experiments described above and in Section 5 using 2 CPU cores per algorithm. Our interpretation of this observation is that because most of the steps within CLVR and other algorithms we are comparing against are simple and cheap, involving very few large matrix-vector multiplications and no matrix factorization, the practical performance of algorithms becomes memory-bound, hence additional cores do not make much difference.*