

Speeding up Linear Programming using Randomized Linear Algebra

Agniva Chowdhury* Palma London† Haim Avron‡ Petros Drineas§

Abstract

Linear programming (LP) is an extremely useful tool and has been successfully applied to solve various problems in a wide range of areas, including operations research, engineering, economics, or even more abstract mathematical areas such as combinatorics. It is also used in many machine learning applications, such as ℓ_1 -regularized SVMs, basis pursuit, nonnegative matrix factorization, etc. Interior Point Methods (IPMs) are one of the most popular methods to solve LPs both in theory and in practice. Their underlying complexity is dominated by the cost of solving a system of linear equations at each iteration. In this paper, we consider *infeasible* IPMs for the special case where the number of variables is much larger than the number of constraints. Using tools from Randomized Linear Algebra, we present a preconditioning technique that, when combined with the Conjugate Gradient iterative solver, provably guarantees that infeasible IPM algorithms (suitably modified to account for the error incurred by the approximate solver), converge to a feasible, approximately optimal solution, without increasing their iteration complexity. Our empirical evaluations verify our theoretical results on both real-world and synthetic data.

1 Introduction

Linear programming (LP) is one of the most useful tools available to theoreticians and practitioners throughout science and engineering. It has been extensively used to solve various problems in a wide range of areas, including operations research, engineering, economics, or even in more abstract mathematical areas such as combinatorics. Also in machine learning and numerical optimization, LP appears in numerous settings, including ℓ_1 -regularized SVMs [54], basis pursuit (BP) [51], sparse inverse covariance matrix estimation (SICE) [52], the nonnegative matrix factorization (NMF) [42], MAP inference [34], etc. Not surprisingly, designing and analyzing LP algorithms is a topic of paramount importance in computer science and applied mathematics.

One of the most successful paradigms for solving LPs is the family of Interior Point Methods (IPMs), pioneered by Karmarkar in the mid 1980s [23]. Path-following IPMs and, in particular, long-step path following IPMs, are among the most practical approaches for solving linear programs. Consider the standard form of the primal LP problem:

$$\min \mathbf{c}^\top \mathbf{x}, \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, and $\mathbf{c} \in \mathbb{R}^n$ are the inputs, and $\mathbf{x} \in \mathbb{R}^n$ is the vector of the primal variables. The associated dual problem is

$$\max \mathbf{b}^\top \mathbf{y}, \text{ subject to } \mathbf{A}^\top \mathbf{y} + \mathbf{s} = \mathbf{c}, \mathbf{s} \geq \mathbf{0}, \quad (2)$$

*Department of Statistics, Purdue University, West Lafayette, IN, USA, chowdhu5@purdue.edu.

†Department of Computer Science, California Institute of Technology, Pasadena, CA, USA, plondon@caltech.edu.

‡School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel, haimav@tauex.tau.ac.il.

§Department of Computer Science, Purdue University, West Lafayette, IN, USA, pdrineas@purdue.edu.

where $\mathbf{y} \in \mathbb{R}^m$ and $\mathbf{s} \in \mathbb{R}^n$ are the vectors of the dual and slack variables respectively. Triplets $(\mathbf{x}, \mathbf{y}, \mathbf{s})$ that uphold both (1) and (2) are called *primal-dual solutions*. Path-following IPMs typically converge towards a primal-dual solution by operating as follows: given the current iterate $(\mathbf{x}^k, \mathbf{y}^k, \mathbf{s}^k)$, they compute the Newton search direction $(\Delta\mathbf{x}, \Delta\mathbf{y}, \Delta\mathbf{s})$ and update the current iterate by following a step towards the search direction. To compute the search direction, one standard approach [38] involves solving the *normal equations*¹:

$$\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\Delta\mathbf{y} = \mathbf{p}. \quad (3)$$

Here, $\mathbf{D} = \mathbf{X}^{1/2}\mathbf{S}^{-1/2}$ is a diagonal matrix, $\mathbf{X}, \mathbf{S} \in \mathbb{R}^{n \times n}$ are diagonal matrices whose i -th diagonal entries are equal to \mathbf{x}_i and \mathbf{s}_i , respectively, and $\mathbf{p} \in \mathbb{R}^m$ is a vector whose exact definition is given in eqn. (22)². Given $\Delta\mathbf{y}$, computing $\Delta\mathbf{s}$ and $\Delta\mathbf{x}$ only involves matrix-vector products.

The core computational bottleneck in IPMs is the need to solve the linear system of eqn. (3) at each iteration. This leads to two key challenges: first, for high-dimensional matrices \mathbf{A} , solving the linear system is computationally prohibitive. Most implementations of IPMs use a *direct solver*; see Chapter 6 of [38]. However, if $\mathbf{A}\mathbf{D}^2\mathbf{A}^\top$ is large and dense, direct solvers are computationally impractical. If $\mathbf{A}\mathbf{D}^2\mathbf{A}^\top$ is sparse, specialized direct solvers have been developed, but these do not apply to many LP problems, especially those arising in machine learning applications, due to irregular sparsity patterns. Second, an alternative to direct solvers is the use of iterative solvers, but the situation is further complicated since $\mathbf{A}\mathbf{D}^2\mathbf{A}^\top$ is typically ill-conditioned. Indeed, as IPM algorithms approach the optimal primal-dual solution, the diagonal matrix \mathbf{D} becomes ill-conditioned, which also results in the matrix $\mathbf{A}\mathbf{D}^2\mathbf{A}^\top$ becoming ill-conditioned. Additionally, using approximate solutions for the linear system of eqn. (3) causes certain invariants, which are crucial for guaranteeing the convergence of IPMs, to be violated; see Section 1.1 for details.

In this paper, we address the aforementioned challenges, for the special case where $m \ll n$, i.e., the number of constraints is much smaller than the number of variables; see Section 5 for a generalization. This is a common setting in many applications of LP solvers. For example, in machine learning, ℓ_1 -SVMs and basis pursuit problems often exhibit such structure when the number of available features (n) is larger than the number of objects (m). Indeed, this setting has been of interest in recent work on LPs [16, 3, 29]. For simplicity of exposition, we also assume that the constraint matrix \mathbf{A} has full rank, equal to m . First, we propose and analyze a preconditioned Conjugate Gradient (CG) iterative solver for the normal equations of eqn. (3), using matrix sketching constructions from the Randomized Linear Algebra (RLA) literature. We develop a preconditioner for $\mathbf{A}\mathbf{D}^2\mathbf{A}^\top$ using matrix sketching which allows us to prove strong convergence guarantees for the *residual* of CG solvers. Second, building upon the work of [35], we propose and analyze a provably accurate long-step *infeasible* IPM algorithm. The proposed IPM solves the normal equations using iterative solvers. In this paper, for brevity and clarity, we primarily focus our description and analysis on the CG iterative solver. We note that a non-trivial concern is that the use of iterative solvers and matrix sketching tools implies that the normal equations at each iteration will be solved only approximately. In our proposed IPM, we develop a novel way to *correct* for the error induced by the approximate solution in order to guarantee convergence. Importantly, this correction step is relatively computationally light, unlike a similar step proposed in [35]. Third, we empirically show that our algorithm performs well in practice. We consider solving LPs that arise from ℓ_1 -regularized SVMs and test them on a variety of synthetic and real-world data sets. Several extensions of our work are discussed in Section 5.

¹Another widely used approach is to solve the augmented system [38]. This approach is less relevant for this paper.

²The superscript k in eqn. (22) simply indicates iteration count and is omitted here for notational simplicity.

1.1 Our contributions

Our point of departure in this work is the introduction of preconditioned, iterative solvers for solving eqn. (3). Preconditioning is used to address the ill-conditioning of the matrix $\mathbf{A}\mathbf{D}^2\mathbf{A}^\top$. Iterative solvers allow the computation of approximate solutions using only matrix-vector products while avoiding matrix inversion, Cholesky or LU factorizations, etc. A preconditioned formulation of eqn. (3) is:

$$\mathbf{Q}^{-1}\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\Delta\mathbf{y} = \mathbf{Q}^{-1}\mathbf{p}, \quad (4)$$

where $\mathbf{Q} \in \mathbb{R}^{m \times m}$ is the preconditioning matrix; \mathbf{Q} should be easily invertible (see [2, 20] for background). An alternative yet equivalent formulation of eqn. (4), which is more amenable to theoretical analysis, is

$$\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\mathbf{z} = \mathbf{Q}^{-1/2}\mathbf{p}, \quad (5)$$

where $\mathbf{z} \in \mathbb{R}^m$ is a vector such that $\Delta\mathbf{y} = \mathbf{Q}^{-1/2}\mathbf{z}$. Note that the matrix in the left-hand side of the above equation is always symmetric, which is not necessarily the case for eqn. (4). We do emphasize that one can use eqn. (4) in the actual implementation of the preconditioned solver; eqn. (5) is much more useful in theoretical analyses.

Recall that we focus on the special case where $\mathbf{A} \in \mathbb{R}^{m \times n}$ has $m \ll n$, i.e., it is a short-and-fat matrix. Our first contribution starts with the design and analysis of a preconditioner for the Conjugate Gradient solver that satisfies, with high probability,

$$\frac{2}{2+\zeta} \leq \sigma_{\min}^2(\mathbf{Q}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}) \leq \sigma_{\max}^2(\mathbf{Q}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}) \leq \frac{2}{2-\zeta}, \quad (6)$$

for some error parameter $\zeta \in [0, 1]$. In the above, $\sigma_{\min}(\cdot)$ and $\sigma_{\max}(\cdot)$ correspond to the smallest and largest singular value of the matrix in parentheses. The above condition says that the preconditioner effectively reduces the condition number of $\mathbf{A}\mathbf{D}$ to a constant. We note that the particular form of the lower and upper bounds in eqn. (6) was chosen to simplify our derivations. RLA matrix-sketching techniques allow us to construct preconditioners for all short-and-fat matrices that satisfy the above inequality *and* can be inverted efficiently. Such constructions go back to the work of [1]; see Section 3 for details on the construction of \mathbf{Q} and its inverse. Importantly, given such a preconditioner, we then prove that the resulting CG iterative solver satisfies

$$\|\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^t - \mathbf{Q}^{-1/2}\mathbf{p}\|_2 \leq \zeta^t \|\mathbf{Q}^{-1/2}\mathbf{p}\|_2. \quad (7)$$

Here $\tilde{\mathbf{z}}^t$ is the approximate solution returned by the CG iterative solver after t iterations. In words, the above inequality states that the *residual* achieved after t iterations of the CG iterative solver drops exponentially fast. To the best of our knowledge, this result is not known in the CG literature: indeed, it is actually well-known that the residual error of CG may oscillate, even in cases where the energy norm of the solution error decreases monotonically. However, we prove that if the preconditioner is sufficiently good, i.e., it satisfies the constraint of eqn. (6), then the residual error decreases as well.

Our second contribution is the analysis of a novel variant of a long-step *infeasible* IPM algorithm proposed by [35]. Recall that such algorithms can, in general, start with an initial point that is not necessarily feasible, but does need to satisfy some, more relaxed, constraints. Following the lines of [53, 35], let \mathcal{S} be the set of feasible and optimal solutions of the form $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{s}^*)$ for the primal and dual problems of eqns. (1) and (2) and assume that \mathcal{S} is not empty. Then, long-step infeasible IPMs can start with any initial point $(\mathbf{x}^0, \mathbf{y}^0, \mathbf{s}^0)$ that satisfies $(\mathbf{x}^0, \mathbf{s}^0) > 0$ *and* $(\mathbf{x}^0, \mathbf{s}^0) \geq (\mathbf{x}^*, \mathbf{s}^*)$, for some feasible and optimal solution $(\mathbf{x}^*, \mathbf{s}^*) \in \mathcal{S}$. In words, the starting primal and slack variables must be strictly positive *and* larger (element-wise) when compared to some feasible, optimal primal-dual

solution. See Chapter 6 of [49] for a discussion regarding why such choices of starting points are also relevant to computational practice.

The flexibility of infeasible IPMs comes at a cost: long-step *feasible* IPMs converge in $\mathcal{O}(n \log 1/\epsilon)$ iterations, while long-step *infeasible* IPMs need $\mathcal{O}(n^2 \log 1/\epsilon)$ iterations to converge [53, 35]. Here ϵ is the accuracy of the approximate LP solution returned by the IPM; see Algorithm 2 for the exact definition. Let

$$\mathbf{A}\mathbf{x}^0 - \mathbf{b} = \mathbf{r}_p^0, \tag{8}$$

$$\mathbf{A}^\top \mathbf{y}^0 + \mathbf{s}^0 - \mathbf{c} = \mathbf{r}_d^0, \tag{9}$$

where $\mathbf{r}_p^0 \in \mathbb{R}^n$ and $\mathbf{r}_d^0 \in \mathbb{R}^m$ are the *primal* and *dual* residuals, respectively, and characterize how far the initial point is from being feasible. As long-step infeasible IPM algorithms iterate and update the primal and dual solutions, the residuals are updated as well. Let $\mathbf{r}^k = (\mathbf{r}_p^k, \mathbf{r}_d^k) \in \mathbb{R}^{n+m}$ be the primal and dual residual at the k -th iteration: it is well-known that the convergence analysis of infeasible long-step IPMs critically depends on \mathbf{r}^k lying on the line segment between 0 and \mathbf{r}^0 . Unfortunately, using approximate solvers (such as the CG solver proposed above) for the normal equations violates this invariant. A simple solution to fix this problem by adding a perturbation vector \mathbf{v} to the current primal-dual solution that guarantees that the invariant is satisfied is proposed in [35]. Again, we use RLA matrix sketching principles to propose an efficient construction for \mathbf{v} that provably satisfies the invariant. Next, we combine the above two primitives to prove that Algorithm 2 in Section 4 satisfies the following theorem.

Theorem 1. *Let $0 \leq \epsilon \leq 1$ be an accuracy parameter. Consider the long-step infeasible IPM Algorithm 2 (Section 4) that solves eqn. (5) using the CG solver of Algorithm 1 (Section 3). Assume that the CG iterative solver runs with accuracy parameter $\zeta = 1/2$ and iteration count $t = \mathcal{O}(\log n)$. Then, with probability at least 0.9, the long-step infeasible IPM converges after $\mathcal{O}(n^2 \log 1/\epsilon)$ iterations.*

We note that the 0.9 success probability above is for simplicity of exposition and can be easily amplified using standard techniques. Also, at each iteration of our infeasible long-step IPM algorithm, the running time is $\mathcal{O}((\text{nnz}(A) + m^3) \log n)$. See Section 4 for a detailed discussion of the overall running time.

Our empirical evaluation demonstrates that our algorithm requires an order of magnitude much fewer inner CG iterations than a standard IPM using CG, while producing a comparably accurate solution (see Section 6). In practice, our empirical evaluation also indicates that using a CG solver with our sketching-based preconditioner does not increase the number of (outer) iterations of the infeasible IPM, compared to unpreconditioned CG or a direct linear solver. In particular, there are instances where our solver performs much better than unpreconditioned CG in terms of (outer) iteration count.

1.2 Comparison with Related Work

There is a large body of literature on solving LPs using IPMs. We only review literature that is immediately relevant to our work. Recall that we solve the normal equations inexactly at each iteration, and develop a way to *correct* for the error incurred. We also focus on IPMs that can use an sufficiently positive, infeasible initial point (see Section 1.1). We discuss below two papers that present related ideas.

The use of an approximate iterative solver for eqn. (3), followed by a correction step to “fix” the approximate solution was proposed in [35] (see our discussion in Section 1.1). We propose efficient, RLA-based approaches to precondition and solve eqn. (3), as well as a novel approach to correct for the approximation error in order to guarantee the convergence of the IPM algorithm. Specifically, [35] propose to solve eqn. (3) using the so-called *maximum weight basis* preconditioner [43]. However,

computing such a preconditioner needs access to a maximal linearly independent set of columns of \mathbf{AD} in each iteration, which is costly, taking $\mathcal{O}(m^2n)$ time in the worst-case. More importantly, while [36] was able to provide a bound on the condition number of the preconditioned matrix, that depends only on properties of \mathbf{A} , and is independent of \mathbf{D} , this bound might, in general, be very large. In contrast, our bound is a constant and it does not depend on properties of \mathbf{A} or its dimension. In addition, [35] assumed a bound on the two-norm of the residual of the preconditioned system, but it is unclear how their preconditioner guarantees such a bound. Similar concerns exist for the construction of the correction vector \mathbf{v} proposed by [35], which our work alleviates.

The line of research in the Theoretical Computer Science literature that is closest to our work is [14], who presented an IPM that uses an approximate solver in each iteration. However, their accuracy guarantee is in terms of the final objective value which is different from ours. More importantly, [14] focuses on *short-step*, feasible IPMs, whereas ours is *long-step* and does not require a feasible starting point. Finally, the approximate solver proposed by [14] works only for the special case of input matrices that correspond to graph Laplacians, following the lines of [44, 45].

We also note that in the Theoretical Computer Science literature, [24, 25, 26, 27, 28, 11] proposed and analyzed theoretically ground-breaking algorithms for LPs based on novel tools such as the so-called *inverse maintenance* for accelerating the linear system solvers in IPMs. However, all these endeavors are primarily focused on the theoretically fast but practically inefficient short-step feasible IPMs. In contrast, our work is focused on infeasible *long-step* IPMs, known to work efficiently in practice. Very recently, [6] proposed another fast, short-step, feasible IPM for solving tall and dense LPs. The output of their algorithm does not satisfy the linear constraints exactly (similar to [14]) and the final convergence guarantee is somewhat different from our work.

Another relevant line of research is the work of [13], which proposed solving eqn. (3) using preconditioned Krylov subspace methods, including variants of *generalized minimum residual* (GMRES) or CG methods. Indeed, [13] conducted extensive numerical experiments on LP problems taken from standard benchmark libraries, but did not provide any theoretical guarantees.

From a matrix-sketching perspective, our work was partially motivated by [7], which presented an iterative, sketching-based algorithm to solve under-constrained ridge regression problems, but did not address how to make use of such approaches in an IPM-based framework, as we do here. Recent papers proposed the so-called *Newton sketch* [40, 50] to construct an approximate Hessian matrix for more general convex objective functions of which LP is a special case. Nevertheless, these randomized second-order methods are significantly faster than the conventional approach only when the data matrix is over-constrained, *i.e.* $m \gg n$. It is unclear whether the approach of [40, 50] is faster than IPMs when the optimization problem to be solved is linear. A probabilistic algorithm to solve LP approximately in a random projection-based reduced feature-space was proposed in [46]. A possible drawback of this paper is that the approximate solution is infeasible with respect to the original region.

Finally, we refer the interested reader to the surveys [48, 18, 22, 31, 17] for more background on Randomized Linear Algebra.

2 Notation and Background

$\mathbf{A}, \mathbf{B}, \dots$ denote matrices and $\mathbf{a}, \mathbf{b}, \dots$ denote vectors. For vector \mathbf{a} , $\|\mathbf{a}\|_2$ denotes its Euclidean norm; for a matrix \mathbf{A} , $\|\mathbf{A}\|_2$ denotes its spectral norm and $\|\mathbf{A}\|_F$ denotes its Frobenius norm. We use $\mathbf{0}$ to denote a null vector or null matrix, dependent upon context, and $\mathbf{1}$ to denote the all-ones vector. For any matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ with $m \leq n$ of rank m a thin Singular Value Decomposition (SVD) is a product $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, with $\mathbf{U} \in \mathbb{R}^{m \times m}$ (the matrix of the left singular vectors), $\mathbf{V} \in \mathbb{R}^{n \times m}$ (the matrix of the top- m right singular vectors), and $\mathbf{\Sigma} \in \mathbb{R}^{m \times m}$ a diagonal matrix whose entries are equal to the singular values of \mathbf{X} . We use $\sigma_i(\cdot)$ to denote the i -th singular value of the matrix in parentheses.

For any two symmetric positive semidefinite (resp. positive definite) matrices \mathbf{A}_1 and \mathbf{A}_2 of appropriate dimensions, $\mathbf{A}_1 \preceq \mathbf{A}_2$ ($\mathbf{A}_1 \prec \mathbf{A}_2$) denotes that $\mathbf{A}_2 - \mathbf{A}_1$ is positive semidefinite (resp. positive definite).

We now briefly discuss a result on matrix sketching [12, 10] that is particularly useful in our theoretical analyses. In our parlance, [12] proved that, for any matrix $\mathbf{Z} \in \mathbb{R}^{m \times n}$, there exists a sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$ such that

$$\|\mathbf{Z}\mathbf{W}\mathbf{W}^\top\mathbf{Z}^\top - \mathbf{Z}\mathbf{Z}^\top\|_2 \leq \frac{\zeta}{4} \left(\|\mathbf{Z}\|_2^2 + \frac{\|\mathbf{Z}\|_F^2}{r} \right) \quad (10)$$

holds with probability at least $1 - \delta$ for any $r \geq 1$. Here $\zeta \in [0, 1]$ is a (constant) accuracy parameter. Ignoring constant terms, $w = \mathcal{O}(r \log(r/\delta))$; \mathbf{W} has $\mathcal{O}(\log(r/\delta))$ non-zero entries per row; and the product $\mathbf{Z}\mathbf{W}$ can be computed in time $\mathcal{O}(\log(r/\delta) \cdot \text{nnz}(\mathbf{Z}))$.

3 Conjugate Gradient Solver

In this section, we discuss the computation of the preconditioner \mathbf{Q} (and its inverse), followed by a discussion on how such a preconditioner can be used to satisfy eqns. (6) and (7).

Algorithm 1 Solving eqn. (5) via CG

Input: $\mathbf{AD} \in \mathbb{R}^{m \times n}$, $\mathbf{p} \in \mathbb{R}^m$, sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$, iteration count t ;

1. Compute \mathbf{ADW} and its SVD: let $\mathbf{U}_\mathbf{Q} \in \mathbb{R}^{m \times m}$ be the matrix of its left singular vectors and let $\Sigma_\mathbf{Q}^{1/2} \in \mathbb{R}^{m \times m}$ be the matrix of its singular values;
2. Compute $\mathbf{Q}^{-1/2} = \mathbf{U}_\mathbf{Q}\Sigma_\mathbf{Q}^{-1/2}\mathbf{U}_\mathbf{Q}^\top$;
3. Initialize $\tilde{\mathbf{z}}^0 \leftarrow \mathbf{0}_m$ and run standard CG on the preconditioned system of eqn. (5) for t iterations;

Output: return $\tilde{\mathbf{z}}^t$;

Algorithm 1 takes as input the sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$, which we construct as discussed in Section 2. Our preconditioner \mathbf{Q} is equal to

$$\mathbf{Q} = \mathbf{AD}\mathbf{W}\mathbf{W}^\top\mathbf{D}\mathbf{A}^\top. \quad (11)$$

Notice that we only need to compute $\mathbf{Q}^{-1/2}$ in order to use it to solve eqn. (5). Towards that end, we first compute the sketched matrix $\mathbf{ADW} \in \mathbb{R}^{m \times w}$. Then, we compute the SVD of the matrix \mathbf{ADW} : let $\mathbf{U}_\mathbf{Q}$ be the matrix of its left singular vectors and let $\Sigma_\mathbf{Q}^{1/2}$ be the matrix of its singular values. Notice that the left (and right) singular vectors of $\mathbf{Q}^{-1/2}$ are equal to $\mathbf{U}_\mathbf{Q}$ and its singular values are equal to $\Sigma_\mathbf{Q}^{-1/2}$. Therefore, $\mathbf{Q}^{-1/2} = \mathbf{U}_\mathbf{Q}\Sigma_\mathbf{Q}^{-1/2}\mathbf{U}_\mathbf{Q}^\top$.

Let $\mathbf{AD} = \mathbf{U}\Sigma\mathbf{V}^\top$ be the thin SVD representation of \mathbf{AD} . We apply the results of [12] (see Section 2) to the matrix $\mathbf{Z} = \mathbf{V}^\top \in \mathbb{R}^{m \times n}$ with $r = m$ to get that, with probability at least $1 - \delta$,

$$\|\mathbf{V}^\top\mathbf{W}\mathbf{W}^\top\mathbf{V} - \mathbf{I}_m\|_2 \leq \frac{\zeta}{4} \left(\|\mathbf{V}\|_2^2 + \frac{\|\mathbf{V}\|_F^2}{m} \right) \leq \frac{\zeta}{2}. \quad (12)$$

In the above we used $\|\mathbf{V}\|_2 = 1$ and $\|\mathbf{V}\|_F^2 = m$. The running time needed to compute the sketch \mathbf{ADW} is equal to (ignoring constant factors) $\mathcal{O}(\text{nnz}(\mathbf{A}) \cdot \log(m/\delta))$. Note that $\text{nnz}(\mathbf{AD}) = \text{nnz}(\mathbf{A})$. The cost of computing the SVD of \mathbf{ADW} (and therefore $\mathbf{Q}^{-1/2}$) is $\mathcal{O}(m^3 \log(m/\delta))$. Overall, computing $\mathbf{Q}^{-1/2}$ can be done in time

$$\mathcal{O}(\text{nnz}(\mathbf{A}) \cdot \log(m/\delta) + m^3 \log(m/\delta)). \quad (13)$$

Given these results, we now discuss how to satisfy eqns. (6) and (7) using the sketching matrix \mathbf{W} . We start with the following bound, which is relatively straight-forward given prior RLA work.

Lemma 2. *If the sketching matrix \mathbf{W} satisfies eqn. (12), then, for all $i = 1 \dots m$,*

$$(1 + \zeta/2)^{-1} \leq \sigma_i^2(\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}) \leq (1 - \zeta/2)^{-1}.$$

Proof. Consider the condition of eqn. (12):

$$\|\mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V} - \mathbf{I}_m\|_2 \leq \frac{\zeta}{2} \Leftrightarrow -\frac{\zeta}{2} \mathbf{I}_m \preceq \mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V} - \mathbf{I}_m \preceq \frac{\zeta}{2} \mathbf{I}_m \quad (14)$$

$$\Leftrightarrow -\frac{\zeta}{2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \preceq \mathbf{A} \mathbf{D} \mathbf{W} \mathbf{W}^\top \mathbf{D} \mathbf{A}^\top - \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \preceq \frac{\zeta}{2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \quad (15)$$

$$\Leftrightarrow \left(1 - \frac{\zeta}{2}\right) \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \preceq \underbrace{\mathbf{A} \mathbf{D} \mathbf{W} \mathbf{W}^\top \mathbf{D} \mathbf{A}^\top}_{\mathbf{Q}} \preceq \left(1 + \frac{\zeta}{2}\right) \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top. \quad (16)$$

We obtain eqn. (15) by pre- and post-multiplying the previous inequality by $\mathbf{U} \Sigma$ and $\Sigma \mathbf{U}^\top$ respectively and using the facts that $\mathbf{A} \mathbf{D} = \mathbf{U} \Sigma \mathbf{V}^\top$ and $\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top = \mathbf{U} \Sigma^2 \mathbf{U}^\top$. Also, from eqn. (14), note that all the eigenvalues of $\mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V}$ lie between $(1 - \frac{\zeta}{2})$ and $(1 + \frac{\zeta}{2})$ and thus $\text{rank}(\mathbf{V}^\top \mathbf{W}) = m$. Therefore, $\text{rank}(\mathbf{A} \mathbf{D} \mathbf{W}) = \text{rank}(\mathbf{U} \Sigma \mathbf{V}^\top \mathbf{W}) = m$, as $\mathbf{U} \Sigma$ is non-singular and we know that the rank of a matrix remains unaltered by pre- or post-multiplying it by a non-singular matrix. So, we have $\text{rank}(\mathbf{Q}) = m$; in words \mathbf{Q} has full rank. Therefore, all the diagonal entries of $\Sigma_{\mathbf{Q}}$ are positive and $\mathbf{Q}^{-1/2} \mathbf{Q} \mathbf{Q}^{-1/2} = \mathbf{I}_m$.

Using the above arguments, pre- and post- multiplying eqn. (16) by $\mathbf{Q}^{-1/2}$, we get

$$\begin{aligned} & \left(1 - \frac{\zeta}{2}\right) \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \preceq \mathbf{I}_m \preceq \left(1 + \frac{\zeta}{2}\right) \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \\ \Rightarrow & \left(1 + \frac{\zeta}{2}\right)^{-1} \mathbf{I}_m \preceq \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \preceq \left(1 - \frac{\zeta}{2}\right)^{-1} \mathbf{I}_m. \end{aligned} \quad (17)$$

Eqn. (17) implies that all the eigenvalues of $\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2}$ are bounded between $\left(1 + \frac{\zeta}{2}\right)^{-1}$ and $\left(1 - \frac{\zeta}{2}\right)^{-1}$, which concludes the proof of the lemma. \square

The above lemma directly implies eqn. (6). We now proceed to show that the above construction for $\mathbf{Q}^{-1/2}$, when combined with the conjugate gradient solver to solve eqn. (5), indeed satisfies eqn. (7)³. We do note that in prior work most of the convergence guarantees for CG focus on the error of the approximate solution. However, in our work, we are interested in the convergence of the *residuals* and it is known that even if the energy norm of the error of the approximate solution decreases monotonically, the norms of the CG residuals may oscillate. Interestingly, we can combine a result on the residuals of CG from [5] with Lemma 2 to prove that in our setting the norms of the CG residuals also decrease monotonically.

Let $\tilde{\mathbf{f}}^{(j)}$ be the residual at the j -th iteration of the CG algorithm:

$$\tilde{\mathbf{f}}^{(j)} = \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{z}}^j - \mathbf{Q}^{-1/2} \mathbf{p}.$$

Recall from Algorithm 1 that $\tilde{\mathbf{z}}^0 = \mathbf{0}$ and thus $\tilde{\mathbf{f}}^{(0)} = -\mathbf{Q}^{-1/2} \mathbf{p}$. In our parlance, Theorem 8 of [5] proved the following bound.

³See Chapter 9 of [30] for a detailed overview of CG.

Lemma 3 (Theorem 8 of [5]). Let $\tilde{\mathbf{f}}^{(j-1)}$ and $\tilde{\mathbf{f}}^{(j)}$ be the residuals obtained by the CG solver at steps $j-1$ and j . Then,

$$\|\tilde{\mathbf{f}}^{(j)}\|_2 \leq \frac{\kappa^2(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}) - 1}{2} \|\tilde{\mathbf{f}}^{(j-1)}\|_2,$$

where $\kappa(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D})$ is the condition number of $\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}$.

From Lemma 2, we get

$$\kappa^2(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}) = \frac{\sigma_{\max}^2(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D})}{\sigma_{\min}^2(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D})} \leq \frac{1 + \zeta/2}{1 - \zeta/2}. \quad (18)$$

Combining eqn. (18) with Lemma 3,

$$\|\tilde{\mathbf{f}}^{(j)}\|_2 \leq \frac{\frac{1+\zeta/2}{1-\zeta/2} - 1}{2} \|\tilde{\mathbf{f}}^{(j-1)}\|_2 = \frac{\zeta}{2-\zeta} \|\tilde{\mathbf{f}}^{(j-1)}\|_2 \leq \zeta \|\tilde{\mathbf{f}}^{(j-1)}\|_2, \quad (19)$$

where the last inequality follows from $\zeta \leq 1$. Applying eqn. (19) recursively, we get

$$\|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \zeta \|\tilde{\mathbf{f}}^{(t-1)}\|_2 \leq \dots \leq \zeta^t \|\tilde{\mathbf{f}}^{(0)}\|_2 = \zeta^t \|\mathbf{Q}^{-1/2}\mathbf{p}\|_2,$$

which proves the condition of eqn. (7).

We remark that one can consider using MINRES [39] instead of CG. Our results hinges on bounding the two-norm of the residual. MINRES finds, at each iteration, the optimal vector with respect the two-norm of the residual inside the same Krylov subspace of CG for the corresponding iteration. Thus, the bound we prove for CG applies to MINRES as well.

4 The Infeasible IPM algorithm

In order to avoid spurious solutions, primal-dual path-following IPMs bias the search direction towards the *central path* and restrict the iterates to a neighborhood of the central path. This search is controlled by the *centering parameter* $\sigma \in [0, 1]$. At each iteration, given the current solution $(\mathbf{x}^k, \mathbf{y}^k, \mathbf{s}^k)$, a standard infeasible IPM obtains the search direction $(\Delta\mathbf{x}^k, \Delta\mathbf{y}^k, \Delta\mathbf{s}^k)$ by solving the following system of linear equations:

$$\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\Delta\mathbf{y}^k = \mathbf{p}^k, \quad (20a)$$

$$\Delta\mathbf{s}^k = -\mathbf{r}_d^k - \mathbf{A}^\top\Delta\mathbf{y}^k, \quad (20b)$$

$$\Delta\mathbf{x}^k = -\mathbf{x}^k + \sigma\mu_k\mathbf{S}^{-1}\mathbf{1}_n - \mathbf{D}^2\Delta\mathbf{s}^k. \quad (20c)$$

Here \mathbf{D} and \mathbf{S} are computed given the current iterate $(\mathbf{x}^k$ and $\mathbf{s}^k)$; we skip the indices on \mathbf{D} and \mathbf{S} for notational simplicity. After solving the above system, the infeasible IPM Algorithm 2 proceeds by computing a step-size $\bar{\alpha}$ to return:

$$(\mathbf{x}^{k+1}, \mathbf{y}^{k+1}, \mathbf{s}^{k+1}) = (\mathbf{x}^k, \mathbf{y}^k, \mathbf{s}^k) + \bar{\alpha}(\Delta\mathbf{x}^k, \Delta\mathbf{y}^k, \Delta\mathbf{s}^k). \quad (21)$$

Recall that $\mathbf{r}^k = (\mathbf{r}_p^k, \mathbf{r}_d^k)$ is a vector with $\mathbf{r}_p^k = \mathbf{A}\mathbf{x}^k - \mathbf{b}$ and $\mathbf{r}_d^k = \mathbf{A}^\top\mathbf{y}^k + \mathbf{s}^k - \mathbf{c}$ (the primal and dual residuals). We also use the *duality measure* $\mu_k = \mathbf{x}^{k\top}\mathbf{s}^k/n$ and the vector

$$\mathbf{p}^k = -\mathbf{r}_p^k - \sigma\mu_k\mathbf{A}\mathbf{S}^{-1}\mathbf{1}_n + \mathbf{A}\mathbf{x}^k - \mathbf{A}\mathbf{D}^2\mathbf{r}_d^k. \quad (22)$$

Given $\Delta \mathbf{y}^k$ from eqn. (20a), $\Delta \mathbf{s}^k$ and $\Delta \mathbf{x}^k$ are easy to compute from eqns. (20b) and (20c), as they only involve matrix-vector products. However, since we use Algorithm 1 to solve eqn. (20a) approximately using the sketching-based preconditioned CG solver, the primal and dual residuals *do not* lie on the line segment between $\mathbf{0}$ and \mathbf{r}^0 . This invalidates known proofs of convergence for infeasible IPMs.

For notational simplicity, we now drop the dependency of vectors and scalars on the iteration counter k . Let $\hat{\Delta} \mathbf{y} = \mathbf{Q}^{-1/2} \hat{\mathbf{z}}^t$ be the approximate solution to eqn. (20a). In order to account for the loss of accuracy due to the approximate solver, we compute $\hat{\Delta} \mathbf{x}$ as follows:

$$\hat{\Delta} \mathbf{x} = -\mathbf{x} + \sigma \mu \mathbf{S}^{-1} \mathbf{1}_n - \mathbf{D}^2 \hat{\Delta} \mathbf{s} - \mathbf{S}^{-1} \mathbf{v}. \quad (23)$$

Here $\mathbf{v} \in \mathbb{R}^n$ is a perturbation vector that needs to exactly satisfy the following invariant at each iteration of the infeasible IPM:

$$\mathbf{A} \mathbf{S}^{-1} \mathbf{v} = \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p}. \quad (24)$$

We note that the computation of $\hat{\Delta} \mathbf{s}$ is still done using, essentially, eqn. (20b), namely

$$\Delta \hat{\mathbf{s}}^k = -\mathbf{r}_d^k - \mathbf{A}^\top \hat{\Delta} \mathbf{y}^k. \quad (25)$$

In [35] it is argued that if \mathbf{v} satisfies eqn. (24), the primal and dual residuals lie in the correct line segment.

Construction of \mathbf{v} . There are many choices for \mathbf{v} satisfying eqn. (24). To prove convergence, it is desirable for \mathbf{v} to have a small norm, hence a general choice is

$$\mathbf{v} = (\mathbf{A} \mathbf{S}^{-1})^\dagger (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p}),$$

which involves the computation of the pseudoinverse $(\mathbf{A} \mathbf{S}^{-1})^\dagger$, which is expensive, taking time $\mathcal{O}(m^2 n)$. Instead, we propose to construct \mathbf{v} using the sketching matrix \mathbf{W} of Section 2. More precisely, we construct the perturbation vector

$$\mathbf{v} = (\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p}). \quad (26)$$

The following lemma proves that the proposed \mathbf{v} satisfies eqn. (24).

Lemma 4. *Let $\mathbf{W} \in \mathbb{R}^{n \times w}$ be the sketching matrix of Section 2 and \mathbf{v} be the perturbation vector of eqn. (26). Then, with probability at least $1 - \delta$, $\text{rank}(\mathbf{A} \mathbf{D} \mathbf{W}) = m$ and \mathbf{v} satisfies eqn. (24).*

Proof. Let $\mathbf{A} \mathbf{D} = \mathbf{U} \Sigma \mathbf{V}^\top$ be the thin SVD representation of $\mathbf{A} \mathbf{D}$. We use the exact same \mathbf{W} as discussed in Section 3. Therefore, eqn. (12) holds with probability $1 - \delta$ and it directly follows from the proof of Lemma 2 that $\text{rank}(\mathbf{A} \mathbf{D} \mathbf{W}) = m$. Recall that $\mathbf{A} \mathbf{D} \mathbf{W}$ has full row-rank and thus $\mathbf{A} \mathbf{D} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger = \mathbf{I}_m$. Therefore, taking $\mathbf{v} = (\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p})$, we get

$$\begin{aligned} \mathbf{A} \mathbf{S}^{-1} \mathbf{v} &= \mathbf{A} \mathbf{S}^{-1} (\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p}) \\ &= \mathbf{A} \mathbf{D} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p}) \\ &= \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p}, \end{aligned}$$

where the second equality follows from $\mathbf{D} = \mathbf{X}^{1/2} \mathbf{S}^{-1/2}$. \square

We emphasize here that we use the same exact sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$ to form the preconditioner used in the CG algorithm of Section 3 *as well as* the vector \mathbf{v} in eqn. (26). This allows us to sketch $\mathbf{A} \mathbf{D}$ only once, thus saving time in practice. Next, we present a bound for the two-norm of the perturbation vector \mathbf{v} of eqn. (26).

Lemma 5. *With probability at least $1 - \delta$, our perturbation vector \mathbf{v} in Lemma 4 satisfies*

$$\|\mathbf{v}\|_2 \leq \sqrt{3n\mu} \|\tilde{\mathbf{f}}^{(t)}\|_2, \quad (27)$$

with $\tilde{\mathbf{f}}^{(t)} = \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{z}}^t - \mathbf{Q}^{-1/2} \mathbf{p}$.

Proof. Recall that $\mathbf{Q} = \mathbf{A} \mathbf{D} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\top = \mathbf{U}_\mathbf{Q} \Sigma_\mathbf{Q} \mathbf{U}_\mathbf{Q}^\top$. Also, $\mathbf{U}_\mathbf{Q}$ and $\Sigma_\mathbf{Q}^{1/2}$ are (respectively) the matrices of the left singular vectors and the singular values of $\mathbf{A} \mathbf{D} \mathbf{W}$. Now, let $\hat{\mathbf{V}}$ be the right singular vector of $\mathbf{A} \mathbf{D} \mathbf{W}$. Therefore, $\mathbf{A} \mathbf{D} \mathbf{W} = \mathbf{U}_\mathbf{Q} \Sigma_\mathbf{Q}^{1/2} \hat{\mathbf{V}}^\top$ is the thin SVD representation of $\mathbf{A} \mathbf{D} \mathbf{W}$. Also, from Lemma 2, we know that \mathbf{Q} has full rank. Therefore, $\mathbf{Q}^{1/2} \mathbf{Q}^{-1/2} = \mathbf{I}_m$. Next, we bound $\|\mathbf{v}\|_2$:

$$\begin{aligned} \|\mathbf{v}\|_2 &= \|(\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p})\|_2 \\ &= \|(\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger \mathbf{Q}^{1/2} \mathbf{Q}^{-1/2} (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p})\|_2 \\ &\leq \|(\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} (\mathbf{A} \mathbf{D} \mathbf{W})^\dagger \mathbf{Q}^{1/2}\|_2 \|\tilde{\mathbf{f}}^{(t)}\|_2. \end{aligned} \quad (28)$$

In the above we used $\mathbf{Q}^{-1/2} (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \hat{\Delta} \mathbf{y} - \mathbf{p}) = \tilde{\mathbf{f}}^{(t)}$. Using the SVD of $\mathbf{A} \mathbf{D} \mathbf{W}$ and \mathbf{Q} , we get $(\mathbf{A} \mathbf{D} \mathbf{W})^\dagger \mathbf{Q}^{1/2} = \hat{\mathbf{V}} \Sigma_\mathbf{Q}^{-1/2} \mathbf{U}_\mathbf{Q}^\top \mathbf{U}_\mathbf{Q} \Sigma_\mathbf{Q}^{1/2} \mathbf{U}_\mathbf{Q}^\top = \hat{\mathbf{V}} \mathbf{U}_\mathbf{Q}^\top$. Now, note that $\mathbf{U}_\mathbf{Q} \in \mathbb{R}^{m \times m}$ is an orthogonal matrix and $\|\hat{\mathbf{V}}\|_2 = 1$. Therefore, combining with eqn. (28) yields

$$\begin{aligned} \|\mathbf{v}\|_2 &\leq \|(\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} \hat{\mathbf{V}} \mathbf{U}_\mathbf{Q}^\top\|_2 \|\tilde{\mathbf{f}}^{(t)}\|_2 = \|(\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} \hat{\mathbf{V}}\|_2 \|\tilde{\mathbf{f}}^{(t)}\|_2 \\ &\leq \|(\mathbf{X} \mathbf{S})^{1/2} \mathbf{W}\|_2 \|\tilde{\mathbf{f}}^{(t)}\|_2. \end{aligned} \quad (29)$$

The first equality follows from the unitary invariance property of the spectral norm and the second inequality follows from the sub-multiplicativity of the spectral norm and $\|\hat{\mathbf{V}}\|_2 = 1$. Our construction for \mathbf{W} implies that eqn. (10) holds for any matrix \mathbf{Z} and, in particular, for $\mathbf{Z} = (\mathbf{X} \mathbf{S})^{1/2}$. Eqn. (10) implies that

$$\left\| (\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} \mathbf{W}^\top (\mathbf{X} \mathbf{S})^{1/2} - (\mathbf{X} \mathbf{S}) \right\|_2 \leq \frac{\zeta}{4} \left(\|(\mathbf{X} \mathbf{S})^{1/2}\|_2^2 + \frac{\|(\mathbf{X} \mathbf{S})^{1/2}\|_F^2}{m} \right) \quad (30)$$

holds with probability at least $1 - \delta$. Applying Weyl's inequality on the left hand side of the eqn. (30), we get

$$\left| \left\| (\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} \right\|_2^2 - \left\| (\mathbf{X} \mathbf{S})^{1/2} \right\|_2^2 \right| \leq \frac{\zeta}{4} \left(\|(\mathbf{X} \mathbf{S})^{1/2}\|_2^2 + \frac{\|(\mathbf{X} \mathbf{S})^{1/2}\|_F^2}{m} \right). \quad (31)$$

Using $\zeta \leq 1$ and $\|(\mathbf{X} \mathbf{S})^{1/2}\|_2 \leq \|(\mathbf{X} \mathbf{S})^{1/2}\|_F \leq \mathbf{x}^\top \mathbf{s} = n\mu$, we get⁴

$$\left\| (\mathbf{X} \mathbf{S})^{1/2} \mathbf{W} \right\|_2^2 \leq 3 \|(\mathbf{X} \mathbf{S})^{1/2}\|_F^2 = 3n\mu. \quad (32)$$

Finally, combining eqns. (29) and (32), we conclude

$$\|\mathbf{v}\|_2 \leq \sqrt{3n\mu} \|\tilde{\mathbf{f}}^{(t)}\|_2.$$

□

⁴The constant 3 in eqn. (32) could be slightly improved to $3/2$; we chose to keep the suboptimal constant to make our results compatible with the work in [35] and avoid reproving theorems from [35]. The better constant does not result in any significant improvements in the number of iterations of the Algorithm 2.

The above result is particularly useful in proving the convergence of Algorithm 2. More precisely, combining a result from [35] with our preconditioner $\mathbf{Q}^{-1/2}$, we can prove that $\|\mathbf{Q}^{-1/2}\mathbf{p}\|_2 \leq \mathcal{O}(n)\sqrt{\mu}$. This bound allows us to prove that if we run Algorithm 1 for $\mathcal{O}(\log n)$ iterations, then

$$\|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \frac{\gamma\sigma}{4\sqrt{n}}\sqrt{\mu} \text{ and } \|\mathbf{v}\|_2 \leq \frac{\gamma\sigma}{4}\mu.$$

The last two inequalities are critical in the convergence analysis of Algorithm 2; see Appendix C for details.

We are now ready to present the infeasible IPM algorithm. We will need the following definition for the neighborhood

$$\mathcal{N}(\gamma) = \left\{ (\mathbf{x}^k, \mathbf{y}^k, \mathbf{s}^k) : x_i^k s_i^k \geq (1 - \gamma)\mu \text{ and } \frac{\|\mathbf{r}^k\|_2}{\|\mathbf{r}^0\|_2} \leq \frac{\mu_k}{\mu_0} \right\}.$$

Here $\gamma \in (0, 1)$ and we note that the duality measure μ_k steadily reduces at each iteration.

Algorithm 2 Infeasible IPM

Input: $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, $\mathbf{c} \in \mathbb{R}^n$, $\gamma \in (0, 1)$, tolerance $\epsilon > 0$, centering parameter $\sigma \in (0, 4/5)$;

Initialize: $k \leftarrow 0$; initial point $(\mathbf{x}^0, \mathbf{y}^0, \mathbf{s}^0)$;

while $\mu_k > \epsilon$ **do**

(a) Compute sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$ (Section 2) with $\zeta = 1/2$ and $\delta = \mathcal{O}(n^{-2})$;

(b) Compute $\mathbf{r}_p^k = \mathbf{A}\mathbf{x}^k - \mathbf{b}$; $\mathbf{r}_d^k = \mathbf{A}^\top \mathbf{y}^k + \mathbf{s}^k - \mathbf{c}$; and \mathbf{p}^k from eqn. (22);

(c) Solve the linear system of eqn. (5) for \mathbf{z} using Algorithm 1 with \mathbf{W} from step (a) and $t = \mathcal{O}(\log n)$. Compute $\hat{\Delta}\mathbf{y} = \mathbf{Q}^{-1/2}\mathbf{z}$;

(d) Compute \mathbf{v} using eqn. (26) with \mathbf{W} from step (a); $\hat{\Delta}\mathbf{s}$ using eqn. (20b); $\hat{\Delta}\mathbf{x}$ using eqn. (23);

(e) Compute $\tilde{\alpha} = \operatorname{argmax}\{\alpha \in [0, 1] : (\mathbf{x}^k, \mathbf{y}^k, \mathbf{s}^k) + \alpha(\hat{\Delta}\mathbf{x}^k, \hat{\Delta}\mathbf{y}^k, \hat{\Delta}\mathbf{s}^k) \in \mathcal{N}(\gamma)\}$.

(f) Compute $\bar{\alpha} = \operatorname{argmin}\{\alpha \in [0, \tilde{\alpha}] : (\mathbf{x}^k + \alpha\hat{\Delta}\mathbf{x}^k)^\top (\mathbf{s}^k + \alpha\hat{\Delta}\mathbf{s}^k)\}$.

(g) Compute $(\mathbf{x}^{k+1}, \mathbf{y}^{k+1}, \mathbf{s}^{k+1}) = (\mathbf{x}^k, \mathbf{y}^k, \mathbf{s}^k) + \bar{\alpha}(\hat{\Delta}\mathbf{x}^k, \hat{\Delta}\mathbf{y}^k, \hat{\Delta}\mathbf{s}^k)$; set $k \leftarrow k + 1$;

end while

Running time. We start by discussing the running time to compute \mathbf{v} . As discussed in Section 3, $(\mathbf{ADW})^\dagger$ can be computed in $\mathcal{O}(\operatorname{nnz}(\mathbf{A}) \cdot \log(m/\delta) + m^3 \log(m/\delta))$ time. Now, as \mathbf{W} has $\mathcal{O}(\log(m/\delta))$ non-zero entries per row, pre-multiplying by \mathbf{W} takes $\mathcal{O}(\operatorname{nnz}(\mathbf{A}) \log(m/\delta))$ time (assuming $\operatorname{nnz}(\mathbf{A}) \geq n$). Since \mathbf{X} and \mathbf{S} are diagonal matrices, computing \mathbf{v} takes $\mathcal{O}(\operatorname{nnz}(\mathbf{A}) \cdot \log(m/\delta) + m^3 \log(m/\delta))$ time, which is asymptotically the same as computing $\mathbf{Q}^{-1/2}$ (see eqn. (13)).

We now discuss the overall running time of Algorithm 2. At each iteration, with failure probability δ , the preconditioner $\mathbf{Q}^{-1/2}$ and the vector \mathbf{v} can be computed in $\mathcal{O}(\operatorname{nnz}(\mathbf{A}) \cdot \log(m/\delta) + m^3 \log(m/\delta))$ time. In addition, for $t = \mathcal{O}(\log n)$ iterations of Algorithm 1, all the matrix-vector products in the CG solver can be computed in $\mathcal{O}(\operatorname{nnz}(\mathbf{A}) \cdot \log n)$ time. Therefore, the computational time for steps (a)-(d) is given by $\mathcal{O}(\operatorname{nnz}(\mathbf{A}) \cdot (\log n + \log(m/\delta)) + m^3 \log(m/\delta))$. Finally, taking a union bound over all iterations with $\delta = \mathcal{O}(n^{-2})$ (ignoring constant factors), Algorithm 2 converges with probability at least 0.9. The running time at each iteration is given by $\mathcal{O}((\operatorname{nnz}(\mathbf{A}) + m^3) \log n)$.

5 Extensions

We briefly discuss extensions of our work. First, there is nothing special about using a CG solver for solving eqn. (5). We analyze two more solvers that could replace the proposed CG solver without any

loss in accuracy or any increase in the number of iterations for the long-step infeasible IPM Algorithm 2 of Section 4. In Appendix A, we analyze the performance of the preconditioned Richardson Iteration and in Appendix B, we analyze the performance of the preconditioned Steepest Descent. In both cases, if the respective preconditioned solver (with the preconditioner of Section 3) runs for $t = \mathcal{O}(\log n)$ steps, Theorem 1 still holds, with small differences in the constant terms. While preconditioned Richardson iteration and preconditioned Steepest Descent are interesting from a theoretical perspective, they are not particularly practical. In future work, we will also consider the preconditioned Chebyshev semi-iterative method, which offers practical advantages compared to PCG in parallel settings.

Second, recall that our approach focused on full rank input matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m \ll n$. Our overall approach still works if \mathbf{A} in any $m \times n$ matrix that is low-rank, e.g., $\text{rank}(\mathbf{A}) = k \ll \min\{m, n\}$. In that case, using the thin SVD of \mathbf{A} , we can rewrite the linear constraints as follows $\mathbf{U}_\mathbf{A} \boldsymbol{\Sigma}_\mathbf{A} \mathbf{V}_\mathbf{A}^\top \mathbf{x} = \mathbf{b}$, where $\mathbf{U}_\mathbf{A} \in \mathbb{R}^{m \times k}$ and $\mathbf{V}_\mathbf{A} \in \mathbb{R}^{n \times k}$ are the matrices of left and right singular vectors of \mathbf{A} respectively; $\boldsymbol{\Sigma}_\mathbf{A} \in \mathbb{R}^{k \times k}$ is the diagonal matrix with the k non-zero singular values of \mathbf{A} as its diagonal elements. The LP of eqn. (1) can be restated as

$$\min \mathbf{c}^\top \mathbf{x}, \text{ subject to } \mathbf{V}_\mathbf{A}^\top \mathbf{x} = \tilde{\mathbf{b}}, \mathbf{x} \geq \mathbf{0}, \quad (33)$$

where $\tilde{\mathbf{b}} = \boldsymbol{\Sigma}_\mathbf{A}^{-1} \mathbf{U}_\mathbf{A}^\top \mathbf{b}$. Note that, $\text{rank}(\mathbf{V}_\mathbf{A}) = k \ll n$ and therefore eqn. (33) can be solved using our framework. The matrices $\mathbf{U}_\mathbf{A}$, $\mathbf{V}_\mathbf{A}$, and $\boldsymbol{\Sigma}_\mathbf{A}$ can be approximately recovered using the fast SVD algorithms of [22, 4, 9]. However, the accuracy of the final solution will depend on the accuracy of the approximate SVD and we defer this analysis to future work.

Third, even though we chose to use the Count-Min sketch and its analysis from [12] (Section 2), there are many other alternative sketching matrix constructions that would lead to similar results. A particularly simple one is the Gaussian sketching matrix $\mathbf{W}_G \in \mathbb{R}^{n \times w}$, where every entry is a $\mathcal{N}(0, 1)$ random variable. Setting $w = \mathcal{O}(m + \log(1/\delta)/\zeta^2)$ would result in the same accuracy guarantees as the sketching matrix of Section 2. However, the (theoretical) running time needed to compute $\mathbf{A}\mathbf{D}\mathbf{W}$ increases to $\mathcal{O}(m \cdot \text{nnz}(\mathbf{A}))$. In practice, at least for relatively small matrices, using Gaussian sketching matrices is a reasonable alternative; see the discussion in [33] which argued that the Gaussian matrix sketching-based solvers are considerably better than direct solvers. We also opted to use Gaussian matrices in our empirical evaluation, since we primarily interested in measuring the accuracy of the final solution as a function of the number of iterations of the solver and the IPM algorithm. Other known constructions of sketching matrices that are also applicable in our setting include (any) sub-gaussian sketching matrix; the Subsampled Randomized Hadamard transform (SRHT); and any of the Sparse Subspace Embeddings of [8, 37, 32, 10].

6 Experiments

Here we demonstrate the empirical performance of our algorithm on a variety of real-world data sets from the UCI ML Repository [19], such as ARCENE, DEXTER [21], DrivFace [15], and a gene expression cancer RNA-Sequencing dataset that is part of the PANCAN dataset [47]. See Table 1 for a description of the datasets. The experiments were implemented in Python and we observed that the results for both synthetic data (generated as described in Appendix D.2) and real-world data were qualitatively similar. Thus, we highlight results on several representative real datasets.

As an application, we consider ℓ_1 -regularized SVMs. All of the data sets are concerned with binary classification with $m \ll n$, where n is the number of features. In Appendix D.1, we describe the ℓ_1 -SVM problem and how it can be formulated as an LP. Here, m is the number of training points, n is the feature dimension, and the size of the constraint matrix in the LP becomes $m \times (2n + 1)$.

Comparisons and Metrics. We compare our Algorithm 2 with a standard IPM (see Chapter 10, [41]) using CG, and a standard IPM using a direct solver. We also use CVXPY as a benchmark to compare





Figure 1: *ARCENE* (top row) and *DEXTER* (bottom row) data sets: Our algorithm (Sk. IPM) requires an order of magnitude fewer inner iterations than the Standard IPM with CG at each outer iteration, as demonstrated in (a). This is possibly due to the improved conditioning of $\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^T\mathbf{Q}^{-1/2}$ compared to $\mathbf{A}\mathbf{D}^2\mathbf{A}^T$, as shown in (b). For all experiments $tolCG = 10^{-5}$ and $\tau = 10^{-9}$.

the accuracy of the solutions; we define the *relative error* $\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2 / \|\mathbf{x}^*\|_2$, where $\hat{\mathbf{x}}$ is our solution and \mathbf{x}^* is the solution generated by CVXPY. We also consider the number of *outer iterations*, namely the number of iterations of the IIPM algorithm, as well as the number of *inner iterations*, namely the number of iterations of the CG solver. We denote the relative stopping tolerance for CG by $tolCG$ and we denote the outer iteration residual error by τ . If not specified: $\tau = 10^{-9}$, $tolCG = 10^{-5}$, and $\sigma = 0.5$. We evaluated a Gaussian sketching matrix, and the initial triplet $(\mathbf{x}, \mathbf{y}, \mathbf{s})$ for all IPM algorithms was set to be all ones.

Experimental Results. Figure 1(a) shows that our Algorithm 2 uses an order of magnitude fewer *inner* iterations than the un-preconditioned standard solver. This is due to the improved conditioning of the respective matrices in the normal equations, as demonstrated in Figure 1(b). Across various real-world and synthetic data sets, the results were qualitatively similar to those shown in Figure 1. Results for several real-world data sets are summarized in Table 1.

In general, our preconditioned CG solver used in Algorithm 2 does not increase the total number of *outer* iterations as compared to the standard IPM with CG, and the standard IPM with a direct linear solver (denoted IPM w/Dir), as seen in Table 1. Actually, for unpreconditioned CG there is clearly more outer iterations, especially for Gene RNA, which has x5 outer iterations. Figure 1 also demonstrates the relative insensitivity to the choice of w (the sketching dimension, i.e., the number of columns of the sketching matrix \mathbf{W} of Section 2). For smaller values of w , our algorithm requires



(a) Max. Inner CG Iterations.

(b) Max. Condition Number.

Figure 2: *ARCENE* data set: for various $(w, tolCG)$ settings, (a) the maximum number of inner iterations used by our algorithm and (b) the maximum condition number of $\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^T\mathbf{Q}^{-1/2}$, across outer iterations. The standard IPM, across all settings, needed on the order of 1,000 iterations and $\kappa(\mathbf{A}\mathbf{D}^2\mathbf{A}^T)$ was on the order of 10^8 . The relative error was fixed to 0.04%.

Table 1: Comparison of (our) sketched IPM with CG, standard IPM with CG, and Standard IPM with a direct solver, for the ℓ_1 -SVM problem on UCI Machine Learning Repository [19] data sets. Across all, $\tau = 10^{-9}$ and a relative error of 10^{-3} or less was achieved. We define $\kappa_{\text{Sk}} = \kappa(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^T\mathbf{Q}^{-1/2})$ and $\kappa_{\text{Stan}} = \kappa(\mathbf{A}\mathbf{D}^2\mathbf{A}^T)$.

Problem	Size ($m \times N$)	Sketch IPM w/ Precond. CG				Stand. IPM w/ Unprec. CG			IPM w/ Dir.
		w	In. It.	Out. It.	κ_{Sk}	In. It.	Out. It.	κ_{Stan}	Out. It.
ARCENE	(100 × 10K)	200	30	50	38.09	1.1K	59	4.4×10^8	50
DEXTER	(300 × 20K)	500	39	39	75.42	4.6K	39	7.6×10^9	39
DrivFace	(606 × 6400)	1000	50	42	68.87	139K	43	17×10^{12}	42
Gene RNA	(801 × 20531)	2000	27	44	20.03	101K	208	4.7×10^{12}	44

more inner iterations. However, across various choices of w , the number of inner iterations is always an order of magnitude smaller than the number required by the standard solver.

Figure 2 shows the performance of our algorithm for a range of $(w, tolCG)$ pairs. Figure 2(a) demonstrates that the number of the inner iterations is robust to the choice of $tolCG$ and w . The number of inner iterations varies between 15 and 35 for the ARCENE data set, while the standard IPM took on the order of 1,000 iterations across all parameter settings. Across all settings, the relative error was fixed at 0.04%. In general, our sketched IPM is able to produce an extremely high accuracy solution across parameter settings. Thus we do not report additional numerical results for the relative error, which was consistently 10^{-3} or less. Figure 2(b) demonstrates a trade-off of our approach: as both $tolCG$ and w are increased, the condition number $\kappa(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^T\mathbf{Q}^{-1/2})$ decreases, corresponding to better conditioned systems. As a result, fewer inner iterations are required. In this context, Figure 3 shows that how the number of inner CG iterations (Figure 3(a)) or the condition number of $\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^T\mathbf{Q}^{-1/2}$ (Figure 3(b)) decreases with the increase in sketching dimension w for various $tolCG$.




Figure 3: *ARCENE* data set: As w increases, (a) the number of inner iterations decreases, and is relatively robust to tol_{CG} and (b) the condition number decreases as well.

7 Conclusions

We proposed and analyzed an infeasible IPM algorithm using a preconditioned conjugate gradient solver for the normal equations and a novel perturbation vector to correct for the error due to the approximate solver. Thus, we speed up each iteration of the IPM algorithm, without increasing the overall number of iterations. We demonstrate empirically that our IPM requires an order of magnitude fewer inner iterations within each linear solve than standard IPMs. It would be interesting to extend our work to analyze feasible IPMs. More precisely, we would like to apply Algorithm 2 of Section 4 starting with a strictly feasible point. In that case, the analysis should be simpler and the iteration complexity of the IPM algorithm should reduce to $\mathcal{O}(n \log(1/\epsilon))$, which is the best known for feasible long-step path following IPM algorithms. We chose to present the more technically challenging approach in this paper and delegate the feasible case to future work.

Acknowledgements

AC and PD were partially supported by NSF FRG 1760353 and NSF CCF BSF 1814041. HA was partially supported by BSF grant 2017698. PL was supported by an Amazon Graduate Fellowship in Artificial Intelligence.

References

- [1] Haim Avron, Petar Maymounkov, and Sivan Toledo. Blendenpik: Supercharging LAPACK’s least-squares solver. *SIAM Journal on Scientific Computing*, 32(3):1217–1236, 2010.
- [2] Owe Axelsson and Vincent A. Barker. *Finite element solution of boundary value problems: Theory and computation*, volume 35. Society for Industrial and Applied Mathematics, 1984.
- [3] Daniel Bienstock and Garud Iyengar. Approximating fractional packings and coverings in $\mathcal{O}(1/\epsilon)$ iterations. *SIAM Journal on Computing*, 35(4):825–854, 2006.
- [4] Christos Boutsidis, Petros Drineas, and Malik Magdon-Ismail. Near-optimal column-based matrix reconstruction. *SIAM Journal on Computing*, 43(2):687–717, 2014.

- [5] R. Bouyouli, Gérard Meurant, Laurent Smoch, and Hassane Sadok. New results on the convergence of the conjugate gradient method. *Numerical Linear Algebra with Applications*, 16(3):223–236, 2009.
- [6] Jan van den Brand, Yin Tat Lee, Aaron Sidford, and Zhao Song. Solving tall dense linear programs in nearly linear time. *arXiv preprint arXiv:2002.02304*, 2020.
- [7] Agniva Chowdhury, Jiasen Yang, and Petros Drineas. An iterative, sketching-based framework for ridge regression. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 988–997, 2018.
- [8] Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the 45th Annual ACM symposium on Theory of Computing*, pages 81–90, 2013.
- [9] Kenneth L Clarkson and David P Woodruff. Low-rank approximation and regression in input sparsity time. *Journal of the ACM (JACM)*, 63(6):54, 2017.
- [10] Michael B. Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 278–287, 2016.
- [11] Michael B. Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing*, pages 938–942, 2019.
- [12] Michael B. Cohen, Jelani Nelson, and David P. Woodruff. Optimal approximate matrix product in terms of stable rank. In *43rd International Colloquium on Automata, Languages, and Programming*, pages 11:1–11:14, 2016.
- [13] Yiran Cui, Keiichi Morikuni, Takashi Tsuchiya, and Ken Hayami. Implementation of interior-point methods for LP based on krylov subspace iterative solvers with inner-iteration preconditioning. *Computational Optimization and Applications*, 74(1):143–176, 2019.
- [14] Samuel I. Daitch and Daniel A. Spielman. Faster approximate lossy generalized flow via interior point algorithms. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 451–460, 2008.
- [15] Katerine Diaz-Chito, Aura Hernández-Sabaté, and Antonio M. López. A reduced feature set for driver head pose estimation. *Applied Soft Computing*, 45:98–107, 2016.
- [16] David L. Donoho and Jared Tanner. Sparse nonnegative solution of underdetermined linear equations by linear programming. In *Proceedings of the National Academy of Sciences of the United States of America*, pages 9446–9451, 2005.
- [17] Petros Drineas and Michael W. Mahoney. RandNLA: Randomized numerical linear algebra. *Communications of the ACM*, 59(6):80–90, 2016.
- [18] Petros Drineas and Michael W. Mahoney. *Lectures on randomized numerical linear algebra*, volume 25 of *The Mathematics of Data, IAS/Park City Mathematics Series*. American Mathematical Society, 2018.
- [19] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [20] Gene H. Golub and Charles F. Van Loan. *Matrix computations*, volume 3. Johns Hopkins University Press, 2012.
- [21] Isabelle Guyon, Steve Gunn, Asa Ben-Hur, and Gideon Dror. Result analysis of the NIPS 2003 feature selection challenge. In *Advances in Neural Information Processing Systems*, pages 545–552,

2005.

- [22] Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [23] Narendra Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing*, pages 302–311, 1984.
- [24] Yin Tat Lee and Aaron Sidford. Path finding I: Solving linear programs with $\tilde{O}(\sqrt{rank})$ linear system solves. *arXiv preprint arXiv:1312.6677*, 2013.
- [25] Yin Tat Lee and Aaron Sidford. Path finding II: An $\tilde{O}(m\sqrt{n})$ algorithm for the minimum cost flow problem. *arXiv preprint arXiv:1312.6713*, 2013.
- [26] Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\sqrt{rank})$ iterations and faster algorithms for maximum flow. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science*, pages 424–433, 2014.
- [27] Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science*, pages 230–249, 2015.
- [28] Yin Tat Lee and Aaron Sidford. Solving linear programs with $\tilde{O}(\sqrt{rank})$ linear system solves. *arXiv preprint arXiv:1910.08033*, 2019.
- [29] Palma London, Shai Vardi, Adam Wierman, and Hanling Yi. A parallelizable acceleration framework for packing linear programs. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 3706 – 3713, 2018.
- [30] David G. Luenberger and Yinyu Ye. *Linear and Nonlinear Programming*. Springer Publishing Company, Incorporated, 3rd edition, 2008.
- [31] Michael W. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2):123–224, 2011.
- [32] Xiangrui Meng and Michael W. Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 91–100, 2013.
- [33] Xiangrui Meng, Michael A. Saunders, and Michael W. Mahoney. LSRN: A parallel iterative solver for strongly over- or underdetermined systems. *SIAM Journal on Scientific Computing*, 36(2):95–118, 2014.
- [34] Ofer Meshi and Amir Globerson. An alternating direction method for dual MAP LP relaxation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 470–483. Springer, 2011.
- [35] Renato D. C. Monteiro and Jerome W. O’Neal. Convergence analysis of a long-step primal-dual infeasible interior-point LP algorithm based on iterative linear solvers. *Georgia Institute of Technology*, 2003.
- [36] Renato D. C. Monteiro, Jerome W. O’Neal, and Takashi Tsuchiya. Uniform boundedness of a preconditioned normal matrix used in interior-point methods. *SIAM Journal on Optimization*, 15(1):96–100, 2004.
- [37] Jelani Nelson and Huy L. Nguyễn. OSNAP: Faster numerical linear algebra algorithms via sparser subspace embeddings. In *Proceedings of the 54th IEEE Symposium on Foundations of Computer*

Science, pages 117–126, 2013.

- [38] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [39] Christopher C. Paige and Michael A. Saunders. Solution of sparse indefinite systems of linear equations. *SIAM Journal on Numerical Analysis*, 12(4):617–629, 1975.
- [40] Mert Pilanci and Martin J. Wainwright. Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence. *SIAM Journal on Optimization*, 27(1):205–245, 2017.
- [41] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. Numerical recipes 3rd edition: The art of scientific computing. In *The Oxford Handbook of Innovation*, chapter 10. Cambridge University Press, 2007.
- [42] Ben Recht, Christopher Re, Joel Tropp, and Victor Bittorf. Factoring nonnegative matrices with linear programs. In *Advances in Neural Information Processing Systems*, pages 1214–1222, 2012.
- [43] Mauricio G. C. Resende and Geraldo Veiga. An implementation of the dual affine scaling algorithm for minimum-cost flow on bipartite uncapacitated networks. *SIAM Journal on Optimization*, 3(3):516–537, 1993.
- [44] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, volume 4, pages 81–90, 2004.
- [45] Daniel A. Spielman and Shang-Hua Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014.
- [46] Ky Vu, Pierre-Louis Poirion, and Leo Liberti. Random projections for linear programming. *Mathematics of Operations Research*, 43(4):1051–1071, 2018.
- [47] John N. Weinstein, Eric A. Collisson, Gordon B. Mills, Kenna R. Mills Shaw, Brad A. Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, Joshua M. Stuart, et al. The cancer genome atlas pan-cancer analysis project. *Nature Genetics*, 45(10):1113–1120, 2013.
- [48] David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2), 2014.
- [49] Stephen J. Wright. *Primal-dual interior-point methods*, volume 54. Society for Industrial and Applied Mathematics, 1997.
- [50] Peng Xu, Jiyan Yang, Farbod Roosta-Khorasani, Christopher Ré, and Michael W. Mahoney. Sub-sampled Newton methods with non-uniform sampling. In *Advances in Neural Information Processing Systems*, pages 3000–3008, 2016.
- [51] Junfeng Yang and Yin Zhang. Alternating direction algorithms for ℓ_1 -problems in compressive sensing. *SIAM Journal on Scientific Computing*, 33(1):250–278, 2011.
- [52] Ming Yuan. High dimensional inverse covariance matrix estimation via linear programming. *Journal of Machine Learning Research*, 11(Aug):2261–2286, 2010.
- [53] Yin Zhang. On the convergence of a class of infeasible interior-point methods for the horizontal linear complementarity problem. *SIAM Journal on Optimization*, 4(1):208–227, 1994.
- [54] Ji Zhu, Saharon Rosset, Robert Tibshirani, and Trevor J. Hastie. 1-norm support vector machines. In *Advances in Neural Information Processing Systems*, pages 49–56, 2004.

Appendix A Richardson Iteration

Here, we show that all our analyses still hold if we replace Step 4 of Algorithm 1 (CG solver) with Richardson's iteration. Basically, all we need to show is that the condition of eqn. (7) holds. Note that the condition of eqn. (6) already holds from Lemma 2, as we use the sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$ discussed in Section 3.

Algorithm 3 Richardson Iteration Solver

Input: $\mathbf{AD} \in \mathbb{R}^{m \times n}$, $\mathbf{p} \in \mathbb{R}^m$; number of iterations $t > 0$; sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$;
Initialize: $\tilde{\mathbf{z}}^0 \leftarrow \mathbf{0}_m$;
for $j = 1$ **to** t **do**
 $\tilde{\mathbf{z}}^j \leftarrow \tilde{\mathbf{z}}^{j-1} + \mathbf{Q}^{-1/2}(\mathbf{p} - \mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^{j-1})$;
end for
Output: return $\tilde{\mathbf{z}}^t$;

Our first result expresses the residual vector $\tilde{\mathbf{f}}^{(j)}$ in terms of $\tilde{\mathbf{f}}^{(j-1)}$ for $j = 1 \dots t$.

Lemma 6. *Let $\tilde{\mathbf{f}}^{(j)}$, $j = 1 \dots t$ be the residual vectors at each iteration. Then,*

$$\tilde{\mathbf{f}}^{(j)} = \left(\mathbf{I}_n - \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} \right) \tilde{\mathbf{f}}^{(j-1)}. \quad (34)$$

Recall that $\mathbf{Q} = \mathbf{AD}\mathbf{W}\mathbf{W}^\top\mathbf{D}\mathbf{A}^\top$ and $\tilde{\mathbf{f}}^{(j)} = \mathbf{Q}^{-1/2}(\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^j - \mathbf{p})$.

Proof. Using Algorithm 3, we express $\tilde{\mathbf{f}}^{(j)}$ as

$$\begin{aligned} \tilde{\mathbf{f}}^{(j)} &= \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^j - \mathbf{Q}^{-1/2}\mathbf{p} \\ &= \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} \left(\tilde{\mathbf{z}}^{j-1} + \mathbf{Q}^{-1/2}(\mathbf{p} - \mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^{j-1}) \right) - \mathbf{Q}^{-1/2}\mathbf{p} \\ &= \left(\mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^{j-1} - \mathbf{Q}^{-1/2}\mathbf{p} \right) \\ &\quad - \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} \left(\mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^{j-1} - \mathbf{Q}^{-1/2}\mathbf{p} \right) \\ &= \left(\mathbf{I}_m - \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} \right) \left(\mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^{j-1} - \mathbf{Q}^{-1/2}\mathbf{p} \right) \\ &= \left(\mathbf{I}_m - \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} \right) \tilde{\mathbf{f}}^{(j-1)}, \end{aligned}$$

which concludes the proof. \square

In the next result, we show that the spectral norm of $\mathbf{I}_m - \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}$ is upper bounded by ζ .

Lemma 7. *Let the condition of eqn. (6) holds for the sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$. Then*

$$\|\mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} - \mathbf{I}_m\|_2 \leq \zeta.$$

Proof. As the condition in eqn. (6) holds, we can go backwards in the proof of Lemma 2 and note that eqn. (17) holds. So, we subtract \mathbf{I}_m from each side of eqn. (17) to get

$$\begin{aligned} \left(\frac{2}{2+\zeta} - 1 \right) \mathbf{I}_m &\preceq \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \left(\frac{2}{2-\zeta} - 1 \right) \mathbf{I}_m \\ \Leftrightarrow -\frac{\zeta}{2+\zeta}\mathbf{I}_m &\preceq \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \frac{\zeta}{2-\zeta}\mathbf{I}_m \\ \Rightarrow -\frac{\zeta}{2-\zeta}\mathbf{I}_m &\preceq \mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \frac{\zeta}{2-\zeta}\mathbf{I}_m \end{aligned} \quad (35)$$

$$\Leftrightarrow \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} - \mathbf{I}_m\|_2 \leq \frac{\zeta}{2-\zeta} \leq \zeta. \quad (36)$$

Eqn. (35) holds as $\frac{\zeta}{2+\zeta} \leq \frac{\zeta}{2-\zeta}$ and the last inequality of eqn. (36) follows from $\zeta < 1$. \square

Satisfying eqn. (6). Note that the condition in eqn. (6) already holds from Lemma 2, as we use the exact same sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$.

Satisfying eqn. (7). Using Lemma 7 and applying Lemma 6 recursively, we get

$$\|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \zeta \|\tilde{\mathbf{f}}^{(t-1)}\|_2 \leq \dots \leq \zeta^t \|\tilde{\mathbf{f}}^{(0)}\|_2 = \zeta^t \|\mathbf{Q}^{-1/2} \mathbf{p}\|_2.$$

Appendix B Steepest Descent

We now replace Step 4 of Algorithm 1 (our proposed PCG solver) by preconditioned steepest descent. We again prove that our analysis of the proposed infeasible long-step IPM remains essentially the same.

First, we construct the sketching matrix \mathbf{W} as discussed in Section 2, with a slightly more stringent accuracy guarantee. More specifically, we necessitate that

$$\|\mathbf{v}^\top \mathbf{W} \mathbf{W}^\top \mathbf{v} - \mathbf{I}_m\|_2 \leq \frac{\zeta(1-\zeta)}{2} \quad (37)$$

holds with probability at least $1 - \delta$ for a constant $\zeta \in [0, 1]$. Notice that the sketching dimension $w = \mathcal{O}(m \log(m/\delta))$ and the running time needed to compute $\mathbf{Q}^{-1/2}$ (which is $\mathcal{O}(\text{nnz}(\mathbf{A}) \cdot \log(m/\delta) + m^3 \log(m/\delta))$) remain, asymptotically, the same. In the case of steepest descent, it turns out that at each iteration the search direction is the negative of the gradient, which is equal to the residual $\tilde{\mathbf{f}}^{(j)}$. Moreover, the step size α_j is determined by an exact *line search* that minimizes the underlying quadratic function:

$$\alpha_j = \frac{\tilde{\mathbf{f}}^{(j)\top} \tilde{\mathbf{f}}^{(j)}}{\tilde{\mathbf{f}}^{(j)\top} \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{f}}^{(j)}}.$$

For this choice of α_j , it is easy to verify that the current gradient is orthogonal to the previous one.

Algorithm 4 Steepest Descent Solver

Input: $\mathbf{A} \mathbf{D} \in \mathbb{R}^{m \times n}$, $\mathbf{p} \in \mathbb{R}^m$; number of iterations $t > 0$; sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$;

Initialize: $\tilde{\mathbf{z}}^0 \leftarrow \mathbf{0}_m$;

for $j = 0$ **to** $t - 1$ **do**

$$\alpha_j = \frac{\tilde{\mathbf{f}}^{(j)\top} \tilde{\mathbf{f}}^{(j)}}{\tilde{\mathbf{f}}^{(j)\top} \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{f}}^{(j)}};$$

$$\tilde{\mathbf{z}}^{j+1} \leftarrow \tilde{\mathbf{z}}^j - \alpha_j \tilde{\mathbf{f}}^{(j)};$$

end for

Output: return $\tilde{\mathbf{z}}^t$;

Similar to Lemma 6, our next result reveals a recursive relation between the search directions which will be instrumental in bounding $\tilde{\mathbf{f}}^{(t)}$.

Lemma 8. *Let $\tilde{\mathbf{f}}^{(j)}$, $j = 1 \dots t$ be the residual vectors at each iteration and let α_j be as in Algorithm 4. Then,*

$$\tilde{\mathbf{f}}^{(j+1)} = \left(\mathbf{I}_n - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \right) \tilde{\mathbf{f}}^{(j)}. \quad (38)$$

Recall that $\mathbf{Q} = \mathbf{A} \mathbf{D} \mathbf{W} \mathbf{W}^\top \mathbf{D} \mathbf{A}^\top$ and $\tilde{\mathbf{f}}^{(j)} = \mathbf{Q}^{-1/2} (\mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{z}}^j - \mathbf{p})$.

Proof. In Algorithm 4, we pre-multiply $\tilde{\mathbf{z}}^{j+1}$ by $\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2}$ and then subtract $\mathbf{Q}^{-1/2} \mathbf{p}$ to get

$$\begin{aligned} \tilde{\mathbf{f}}^{(j+1)} &= \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{z}}^{j+1} - \mathbf{Q}^{-1/2} \mathbf{p} \\ &= \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{z}}^j - \mathbf{Q}^{-1/2} \mathbf{p} - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{f}}^{(j)} \\ &= \tilde{\mathbf{f}}^{(j)} - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{f}}^{(j)} = \left(\mathbf{I}_m - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \right) \tilde{\mathbf{f}}^{(j)}, \end{aligned}$$

which concludes the proof. \square

Next, using this new condition in eqn. (37), we bound $\left\| \mathbf{I}_m - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \right\|_2$.

Lemma 9. *If eqn. (37) is satisfied, then $|\alpha_j - 1| \leq \frac{\zeta(1-\zeta)}{2}$.*

Proof. First, we rewrite eqn. (37) as follows,

$$-\frac{\zeta(1-\zeta)}{2} \mathbf{I}_m \preceq \mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V} - \mathbf{I}_m \preceq \frac{\zeta(1-\zeta)}{2} \mathbf{I}_m.$$

Next, we pre- and post-multiply the above expression by $\mathbf{U} \Sigma$ and $\Sigma \mathbf{U}^\top$ to get

$$-\frac{\zeta(1-\zeta)}{2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \preceq \underbrace{\mathbf{A} \mathbf{D} \mathbf{W} \mathbf{W}^\top \mathbf{D} \mathbf{A}^\top}_{\mathbf{Q}} - \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \preceq \frac{\zeta(1-\zeta)}{2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top. \quad (39)$$

Now, pre and post-multiplying eqn. (39) again by $\mathbf{Q}^{-1/2}$, we get

$$\begin{aligned} &\left(1 - \frac{\zeta(1-\zeta)}{2} \right) \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \preceq \mathbf{I}_m \preceq \left(1 + \frac{\zeta(1-\zeta)}{2} \right) \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \\ \Rightarrow &\left(1 - \frac{\zeta(1-\zeta)}{2} \right) \tilde{\mathbf{f}}^{(j)\top} \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{f}}^{(j)} \leq \tilde{\mathbf{f}}^{(j)\top} \tilde{\mathbf{f}}^{(j)} \leq \left(1 + \frac{\zeta(1-\zeta)}{2} \right) \tilde{\mathbf{f}}^{(j)\top} \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{f}}^{(j)} \\ \Rightarrow &\left(1 - \frac{\zeta(1-\zeta)}{2} \right) \leq \frac{\tilde{\mathbf{f}}^{(j)\top} \tilde{\mathbf{f}}^{(j)}}{\tilde{\mathbf{f}}^{(j)\top} \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{f}}^{(j)}} \leq \left(1 + \frac{\zeta(1-\zeta)}{2} \right) \\ \Rightarrow &|\alpha_j - 1| \leq \frac{\zeta(1-\zeta)}{2}, \text{ for } j = 1 \dots t. \end{aligned} \quad (40)$$

\square

Our next result shows that if eqn. (37) holds, then $\left\| \mathbf{I}_m - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \right\|_2$ is upper bounded by a small quantity for $j = 1 \dots t$.

Lemma 10. *If eqn. (37) is satisfied, then $\left\| \mathbf{I}_m - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \right\|_2 \leq \zeta$, for $j = 1 \dots t$.*

Proof. We note that eqn. (37) directly implies

$$\left\| \mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V} - \mathbf{I}_m \right\|_2 \leq \frac{\zeta}{2}. \quad (41)$$

Now, as eqn. (41) holds, from eqn. (17) in the proof of Lemma 2, we have

$$\begin{aligned} &\left(1 + \frac{\zeta}{2} \right)^{-1} \mathbf{I}_m \preceq \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \preceq \left(1 - \frac{\zeta}{2} \right)^{-1} \mathbf{I}_m \\ \Leftrightarrow &\left(\frac{2\alpha_j}{2+\zeta} - 1 \right) \mathbf{I}_m \preceq \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \left(\frac{2\alpha_j}{2-\zeta} - 1 \right) \mathbf{I}_m \end{aligned}$$

$$\Leftrightarrow \frac{2(\alpha_j - 1) - \zeta}{2 + \zeta} \mathbf{I}_m \preceq \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \frac{2(\alpha_j - 1) + \zeta}{2 - \zeta} \mathbf{I}_m. \quad (42)$$

The above expression follows by multiplying eqn. (17) by α_j and then subtracting \mathbf{I}_m . Now, from Lemma 9, we get $-\zeta(1 - \zeta) \leq 2(\alpha_j - 1) \leq \zeta(1 - \zeta)$ for $j = 1 \dots t$. Using this in eqn. (42), we further have

$$\begin{aligned} & -\frac{\zeta(1 - \zeta) + \zeta}{2 + \zeta} \mathbf{I}_m \preceq \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \frac{\zeta(1 - \zeta) + \zeta}{2 - \zeta} \mathbf{I}_m \\ \Leftrightarrow & -\frac{\zeta(2 - \zeta)}{2 + \zeta} \mathbf{I}_m \preceq \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \zeta \mathbf{I}_m \\ \Rightarrow & -\zeta \mathbf{I}_m \preceq \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} - \mathbf{I}_m \preceq \zeta \mathbf{I}_m \\ \Rightarrow & \left\| \mathbf{I}_m - \alpha_j \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \right\|_2 \leq \zeta, \end{aligned} \quad (43)$$

where eqn. (43) is due to the fact that $\frac{2-\zeta}{2+\zeta} \leq 1$. \square

Satisfying eqn. (6). As eqn. (41) holds, eqn. (6) directly follows from Lemma 2.

Satisfying eqn. (7). Using Lemma 10 and applying Lemma 8 recursively, we get

$$\|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \zeta \|\tilde{\mathbf{f}}^{(t-1)}\|_2 \leq \dots \leq \zeta^t \|\tilde{\mathbf{f}}^{(0)}\|_2 = \zeta^t \|\mathbf{Q}^{-1/2} \mathbf{p}\|_2.$$

Appendix C Convergence analysis of Algorithm 2

C.1 Additional notation

For any two vectors $\mathbf{a} = (a_1, \dots, a_\ell)^\top$ and $\mathbf{b} = (b_1, \dots, b_\ell)^\top$ let $\mathbf{a} \circ \mathbf{b} = (a_1 b_1, \dots, a_\ell b_\ell)^\top$. For any vector $\mathbf{a} \in \mathbb{R}^n$ its ℓ_∞ norm is defined as $\|\mathbf{a}\|_\infty = \max_i |a_i|$. We heavily use the following standard inequality to prove results in this section:

$$\left| \frac{\mathbf{a}^\top \mathbf{1}_n}{n} \right| \leq \|\mathbf{a}\|_\infty \leq \|\mathbf{a}\|_2. \quad (44)$$

C.2 Number of iterations for the CG solver

In this section, most of the proofs follow [35] except for the fact that we used our sketching based preconditioner $\mathbf{Q}^{-1/2}$. Recall that \mathcal{S} is the set of optimal and feasible solutions for the proposed LP.

Lemma 11. *Let $(\mathbf{x}^0, \mathbf{y}^0, \mathbf{s}^0)$ be the initial point with $(\mathbf{x}^0, \mathbf{s}^0) > \mathbf{0}$ and $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{s}^*) \in \mathcal{S}$ such that $(\mathbf{x}^*, \mathbf{s}^*) \leq (\mathbf{x}^0, \mathbf{s}^0)$ with $\mathbf{s}^0 \geq |\mathbf{A}^\top \mathbf{y}^0 - \mathbf{c}|$. Then, for any point $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$ such that $\mathbf{r} = \eta \mathbf{r}^0$ and $0 \leq \eta \leq \min \left\{ 1, \frac{\mathbf{s}^0 \mathbf{x}}{\mathbf{s}^0 \mathbf{x}^0} \right\}$, we get*

$$(i) \quad \eta (\mathbf{x}^\top \mathbf{s}^0 + \mathbf{s}^\top \mathbf{x}^0) \leq 3n\mu, \quad (45a)$$

$$(ii) \quad \eta \|\mathbf{S}(\mathbf{x}^* - \mathbf{x}^0)\|_2 \leq \eta \|\mathbf{S}\mathbf{x}^0\|_2 \leq \eta \mathbf{s}^\top \mathbf{x}^0 \leq 3n\mu, \quad (45b)$$

$$(iii) \quad \eta \|\mathbf{X}(\mathbf{s}^0 + \mathbf{A}^\top \mathbf{y}^0 - \mathbf{c})\|_2 \leq 2\eta \|\mathbf{X}\mathbf{s}^0\|_2 \leq 2\eta \mathbf{x}^\top \mathbf{s}^0 \leq 6n\mu. \quad (45c)$$

Proof. We prove eqns. (45a)–(45c) below.

Proof of eqn. (45a). For completeness, we provide a proof of eqn. (45a) following [35]. Since $(\mathbf{x}^*, \mathbf{s}^*, \mathbf{y}^*) \in \mathcal{S}$, the following equalities hold:

$$\mathbf{A}\mathbf{x}^* = \mathbf{b} \quad (46a)$$

$$\mathbf{A}^\top \mathbf{y}^* + \mathbf{s}^* = \mathbf{c}. \quad (46b)$$

Furthermore, $\mathbf{r} = \eta \mathbf{r}^0$ implies

$$\mathbf{A}\mathbf{x} - \mathbf{b} = \eta(\mathbf{A}\mathbf{x}^0 - \mathbf{b}) \quad (47a)$$

$$\mathbf{A}^\top \mathbf{y} + \mathbf{s} - \mathbf{c} = \eta(\mathbf{A}^\top \mathbf{y}^0 + \mathbf{s}^0 - \mathbf{c}). \quad (47b)$$

Combining eqn. (46a) with eqn. (47a) and eqn. (46b) with eqn. (47b), we get

$$\mathbf{A}(\mathbf{x} - \eta \mathbf{x}^0 - (1 - \eta)\mathbf{x}^*) = \mathbf{0} \quad (48a)$$

$$\mathbf{A}^\top (\mathbf{y} - \eta \mathbf{y}^0 - (1 - \eta)\mathbf{y}^*) + (\mathbf{s} - \eta \mathbf{s}^0 - (1 - \eta)\mathbf{s}^*) = \mathbf{0}. \quad (48b)$$

Multiplying eqn. (48b) by $(\mathbf{x} - \eta \mathbf{x}^0 - (1 - \eta)\mathbf{x}^*)^\top$ on the left and using eqn. (48a), we get

$$(\mathbf{x} - \eta \mathbf{x}^0 - (1 - \eta)\mathbf{x}^*)^\top (\mathbf{s} - \eta \mathbf{s}^0 - (1 - \eta)\mathbf{s}^*) = 0.$$

Expanding we get

$$\begin{aligned} \eta (\mathbf{x}^{0\top} \mathbf{s} + \mathbf{x}^\top \mathbf{s}^0) &= \eta^2 \mathbf{x}^{0\top} \mathbf{s}^0 + (1 - \eta)^2 (\mathbf{x}^*)^\top \mathbf{s}^* + \mathbf{x}^\top \mathbf{s} \\ &+ \eta(1 - \eta) (\mathbf{x}^{0\top} \mathbf{s}^* + (\mathbf{x}^*)^\top \mathbf{s}^0) - (1 - \eta) ((\mathbf{x}^*)^\top \mathbf{s} + \mathbf{x}^\top \mathbf{s}^*). \end{aligned} \quad (49)$$

Next, we use the given conditions and rewrite eqn. (49) as

$$\begin{aligned} \eta (\mathbf{x}^{0\top} \mathbf{s} + \mathbf{s}^{0\top} \mathbf{x}) &\leq \eta^2 \mathbf{x}^{0\top} \mathbf{s}^0 + \mathbf{x}^\top \mathbf{s} + \eta(1 - \eta) (\mathbf{x}^{0\top} \mathbf{s}^* + \mathbf{s}^{0\top} \mathbf{x}^*) \\ &\leq \eta^2 \mathbf{x}^{0\top} \mathbf{s}^0 + \mathbf{x}^\top \mathbf{s} + 2\eta(1 - \eta) \mathbf{x}^{0\top} \mathbf{s}^0 \\ &\leq 2\eta \mathbf{x}^{0\top} \mathbf{s}^0 + \mathbf{x}^\top \mathbf{s} \leq 3\mathbf{x}^\top \mathbf{s} = 3n\mu. \end{aligned} \quad (50)$$

The first inequality in eqn. (50) follows from the following facts. First, $(1 - \eta)((\mathbf{x}^*)^\top \mathbf{s} + \mathbf{x}^\top \mathbf{s}^*) \geq 0$ as $(\mathbf{x}^*, \mathbf{s}^*) \geq \mathbf{0}$ and $(\mathbf{x}^0, \mathbf{s}^0) \geq \mathbf{0}$. Second, as $(\mathbf{x}^*, \mathbf{s}^*, \mathbf{y}^*) \in \mathcal{S}$ (which implies $\mathbf{x}^* \circ \mathbf{s}^* = \mathbf{0}$), we have $(\mathbf{x}^*)^\top \mathbf{s}^* = 0$. The second inequality in eqn. (50) holds as $\mathbf{x}^* \leq \mathbf{x}^0$, $\mathbf{s}^* \leq \mathbf{s}^0$, $(\mathbf{x}^*, \mathbf{s}^*) \geq \mathbf{0}$, and $(\mathbf{x}^0, \mathbf{s}^0) \geq \mathbf{0}$; combining them we get $(\mathbf{x}^{0\top} \mathbf{s}^* + \mathbf{s}^{0\top} \mathbf{x}^*) \leq 2\mathbf{x}^{0\top} \mathbf{s}^0$. Third inequality in eqn. (50) is true as we have $\eta^2 \mathbf{x}^{0\top} \mathbf{s}^0 + 2\eta(1 - \eta) \mathbf{x}^{0\top} \mathbf{s}^0 = 2\eta \mathbf{x}^{0\top} \mathbf{s}^0 - \eta^2 \mathbf{x}^{0\top} \mathbf{s}^0 \leq 2\eta \mathbf{x}^{0\top} \mathbf{s}^0$. The final inequality holds as $\eta \leq \frac{\mathbf{x}^\top \mathbf{s}}{\mathbf{x}^{0\top} \mathbf{s}^0}$. \square

Proof of eqn. (45b). The last inequality follows from eqn. (45a). The second to last inequality is also easy to prove as

$$\|\mathbf{S}\mathbf{x}^0\|_2 = \sqrt{\sum_{i=1}^s (s_i x_i^0)^2} \leq \sqrt{\left(\sum_{i=1}^s s_i x_i^0\right)^2} = \mathbf{s}^\top \mathbf{x}^0. \quad (51)$$

To prove the first inequality in eqn. (45b), we use the fact $\mathbf{x}^0 \geq \mathbf{x}^*$ as follows:

$$\begin{aligned} \|\mathbf{S}\mathbf{x}^0\|_2^2 - \|\mathbf{S}(\mathbf{x}^* - \mathbf{x}^0)\|_2^2 &= \sum_{i=1}^n (s_i x_i^0)^2 - \sum_{i=1}^n s_i^2 \left((x_i^*)^2 + (x_i^0)^2 - 2x_i^* x_i^0 \right) \\ &= \sum_{i=1}^n s_i^2 \left(2x_i^* x_i^0 - (x_i^*)^2 \right) \geq 0. \end{aligned}$$

\square

Proof of eqn. (45c). To prove this we use a similar approach as in eqn. (45b). The last inequality directly follows from eqn. (45a); the second to last inequality is also easy to prove as

$$\|\mathbf{X}\mathbf{s}^0\|_2 = \sqrt{\sum_{i=1}^n (x_i s_i^0)^2} \leq \sqrt{\left(\sum_{i=1}^n x_i s_i^0\right)^2} = \mathbf{x}^\top \mathbf{s}^0. \quad (52)$$

For the first inequality, we proceed as follows:

$$\begin{aligned} \|\mathbf{X}(\mathbf{s}^0 + \mathbf{A}^\top \mathbf{y}^0 - \mathbf{c})\|_2^2 &= \|\mathbf{X}\mathbf{s}^0\|_2^2 + \|\mathbf{X}(\mathbf{A}^\top \mathbf{y}^0 - \mathbf{c})\|_2^2 + 2\mathbf{s}^{0\top} \mathbf{X}^\top \mathbf{X}(\mathbf{A}^\top \mathbf{y}^0 - \mathbf{c}) \\ &= \|\mathbf{X}\mathbf{s}^0\|_2^2 + \sum_{i=1}^n x_i^2 (\mathbf{A}^\top \mathbf{y}^0 - \mathbf{c})_i^2 + 2 \sum_{i=1}^n x_i^2 s_i^0 (\mathbf{A}^\top \mathbf{y}^0 - \mathbf{c})_i \\ &\leq \|\mathbf{X}\mathbf{s}^0\|_2^2 + \sum_{i=1}^n (x_i s_i^0)^2 + 2 \sum_{i=1}^n (x_i s_i^0)^2 \\ &= \|\mathbf{X}\mathbf{s}^0\|_2^2 + \|\mathbf{X}\mathbf{s}^0\|_2^2 + 2\|\mathbf{X}\mathbf{s}^0\|_2^2 = 4\|\mathbf{X}\mathbf{s}^0\|_2^2. \end{aligned} \quad (53)$$

The inequality in eqn. (53) follows from $x_i \geq 0$, $s_i^0 \geq 0$ and $|(\mathbf{A}^\top \mathbf{y}^0 - \mathbf{c})_i| \leq s_i^0$ for all $i = 1 \dots n$. \square

Our next result bounds $\|\mathbf{Q}^{-1/2} \mathbf{p}\|_2$ which is instrumental in proving the final bound.

Lemma 12. *Let $(\mathbf{x}^0, \mathbf{y}^0, \mathbf{s}^0)$ be the initial point with $(\mathbf{x}^0, \mathbf{s}^0) > \mathbf{0}$ such that $\mathbf{x}^0 \geq \mathbf{x}^*$ and $\mathbf{s}^0 \geq \max\{\mathbf{s}^*, |\mathbf{c} - \mathbf{A}^\top \mathbf{y}^0|\}$ for some $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{s}^*) \in \mathcal{S}$. Furthermore, let $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$ with $\mathbf{r} = \eta \mathbf{r}^0$ for some $0 \leq \eta \leq 1$. If the sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$ satisfies the condition in eqn. (6), then*

$$\|\mathbf{Q}^{-1/2} \mathbf{p}\|_2 \leq \sqrt{2} \left(\frac{9n}{\sqrt{1-\gamma}} + \sigma \sqrt{\frac{n}{1-\gamma}} + \sqrt{n} \right) \sqrt{\mu}.$$

Recall that $\mathbf{r} = (\mathbf{r}_p, \mathbf{r}_d) = (\mathbf{A}\mathbf{x} - \mathbf{b}, \mathbf{A}^\top \mathbf{y} + \mathbf{s} - \mathbf{c})$ and $\mathbf{r}^0 = (\mathbf{r}_p^0, \mathbf{r}_d^0) = (\mathbf{A}\mathbf{x}^0 - \mathbf{b}, \mathbf{A}^\top \mathbf{y}^0 + \mathbf{s}^0 - \mathbf{c})$.

Proof. Note that after correcting the approximation error of the CG solver using \mathbf{v} , the primal and dual residuals $\mathbf{r} = (\mathbf{r}_p, \mathbf{r}_d)$ corresponding to an iterate $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$ always lie on the line segment between zero and $\mathbf{r}^{(0)}$. In other words, $\mathbf{r} = \eta \mathbf{r}^{(0)}$ always holds for some $\eta \in [0, 1]$. This was formally proven in Lemma 3.3 of [35]. In order to bound $\|\mathbf{Q}^{-1/2} \mathbf{p}\|_2$, first we express \mathbf{p} as in eqn. (3) and rewrite

$$\mathbf{Q}^{-1/2} \mathbf{p} = \mathbf{Q}^{-1/2} \left(-\mathbf{r}_p - \sigma \mu \mathbf{A} \mathbf{S}^{-1} \mathbf{1}_n + \mathbf{A} \mathbf{x} - \mathbf{A} \mathbf{D}^2 \mathbf{r}_d \right). \quad (54)$$

Then, applying the triangle inequality to $\|\mathbf{Q}^{-1/2} \mathbf{p}\|_2$ in eqn. (54), we get

$$\|\mathbf{Q}^{-1/2} \mathbf{p}\|_2 \leq \Delta_1 + \Delta_2 + \Delta_3 + \Delta_4, \quad (55)$$

where

$$\begin{aligned} \Delta_1 &= \|\mathbf{Q}^{-1/2} \mathbf{r}_p\|_2, \\ \Delta_2 &= \sigma \mu \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D} (\mathbf{X} \mathbf{S})^{-1/2} \mathbf{1}_n\|_2, \\ \Delta_3 &= \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D} \mathbf{D}^{-1} \mathbf{x}\|_2, \\ \Delta_4 &= \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{r}_d\|_2. \end{aligned}$$

To bound Δ_1 , Δ_2 , Δ_3 and Δ_4 we heavily use the condition of eqn. (6).

Bounding Δ_1 . Using $\mathbf{r}_p = \eta \mathbf{r}_p^0$, $\mathbf{r}_p^0 = \mathbf{A}\mathbf{x}^0 - \mathbf{b}$ and $\mathbf{b} = \mathbf{A}\mathbf{x}^*$, we rewrite Δ_1 as

$$\begin{aligned}
\Delta_1 &= \eta \|\mathbf{Q}^{-1/2} \mathbf{A}(\mathbf{x}^0 - \mathbf{x}^*)\|_2 \\
&= \eta \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D} \mathbf{D}^{-1}(\mathbf{x}^0 - \mathbf{x}^*)\|_2 \\
&\leq \eta \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}\|_2 \|\mathbf{D}^{-1}(\mathbf{x}^0 - \mathbf{x}^*)\|_2 \\
&\leq \sqrt{2} \eta \|\mathbf{D}^{-1}(\mathbf{x}^0 - \mathbf{x}^*)\|_2 \\
&= \sqrt{2} \eta \|(\mathbf{X} \mathbf{S})^{-1/2} \mathbf{S}(\mathbf{x}^0 - \mathbf{x}^*)\|_2 \\
&\leq \sqrt{2} \eta \|(\mathbf{X} \mathbf{S})^{-1/2}\|_2 \|\mathbf{S}(\mathbf{x}^0 - \mathbf{x}^*)\|_2,
\end{aligned} \tag{56}$$

where the above steps follow from submultiplicativity and eqn. (6). From eqn. (6), note that we have $\|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}\|_2 \leq \sqrt{2}$ as $\zeta \leq 1$. Now, applying eqn. (45b) and $\|(\mathbf{X} \mathbf{S})^{-1/2}\|_2 = \max_{1 \leq i \leq n} \frac{1}{\sqrt{x_i s_i}}$, we further have

$$\begin{aligned}
\Delta_1 &\leq \sqrt{2} \max_{1 \leq i \leq n} \frac{1}{\sqrt{x_i s_i}} \cdot 3n\mu \\
&\leq 3\sqrt{2} n \sqrt{\frac{\mu}{1-\gamma}},
\end{aligned} \tag{57}$$

where the last inequality follows from $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$.

Bounding Δ_2 . Applying submultiplicativity, we get

$$\begin{aligned}
\Delta_2 &= \sigma \mu \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D} (\mathbf{X} \mathbf{S})^{-1/2} \mathbf{1}_n\|_2 \\
&\leq \sigma \mu \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}\|_2 \|(\mathbf{X} \mathbf{S})^{-1/2} \mathbf{1}_n\|_2 \\
&\leq \sqrt{2} \sigma \mu \|(\mathbf{X} \mathbf{S})^{-1/2} \mathbf{1}_n\|_2 \\
&= \sqrt{2} \sigma \mu \sqrt{\sum_{i=1}^n \frac{1}{x_i s_i}} \leq \sqrt{2} \sigma \mu \sqrt{\sum_{i=1}^n \frac{1}{(1-\gamma)\mu}} \\
&= \sqrt{2} \sigma \sqrt{\frac{n\mu}{(1-\gamma)}},
\end{aligned} \tag{58}$$

where the second to last inequality follows from eqn. (6) and the last inequality holds as $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$.

Bounding Δ_3 . Using $\mathbf{D} = \mathbf{S}^{-1/2} \mathbf{X}^{1/2}$ and $\mathbf{x} = \mathbf{X} \mathbf{1}_n$ we get

$$\begin{aligned}
\Delta_3 &= \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D} (\mathbf{S}^{1/2} \mathbf{X}^{-1/2}) \mathbf{X} \mathbf{1}_n\|_2 \\
&= \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D} (\mathbf{S} \mathbf{X})^{1/2} \mathbf{1}_n\|_2 \\
&\leq \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}\|_2 \|(\mathbf{S} \mathbf{X})^{1/2} \mathbf{1}_n\|_2 \\
&\leq \sqrt{2} \sqrt{\sum_{i=1}^n x_i s_i} = \sqrt{2n\mu},
\end{aligned} \tag{59}$$

where the inequalities follow from submultiplicativity and eqn. (6).

Bounding Δ_4 . Using $\mathbf{r}_d = \eta \mathbf{r}_d^0$, we have

$$\begin{aligned}\Delta_4 &= \eta \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{r}_d^0\|_2 \\ &\leq \eta \|\mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}\|_2 \|(\mathbf{X} \mathbf{S})^{-1/2} \mathbf{X} \mathbf{r}_d^0\|_2 \\ &\leq \sqrt{2} \eta \|(\mathbf{X} \mathbf{S})^{-1/2} \mathbf{X} (\mathbf{A}^\top \mathbf{y}^0 + \mathbf{s}^0 - \mathbf{c})\|_2 \\ &\leq \sqrt{2} \eta \|(\mathbf{X} \mathbf{S})^{-1/2}\|_2 \|\mathbf{X} (\mathbf{A}^\top \mathbf{y}^0 + \mathbf{s}^0 - \mathbf{c})\|_2,\end{aligned}$$

where the above inequalities follow from submultiplicativity and eqn. (6). Now, applying eqn. (45c) and $\|(\mathbf{X} \mathbf{S})^{-1/2}\|_2 \leq \frac{1}{\sqrt{(1-\gamma)\mu}}$, we further have

$$\Delta_4 \leq 6\sqrt{2}n \sqrt{\frac{\mu}{1-\gamma}}. \quad (60)$$

Final bound. Combining eqns. (55), (57), (58), (59) and (60), we get

$$\|\mathbf{Q}^{-1/2} \mathbf{p}\|_2 \leq \sqrt{2} \left(\frac{9n}{\sqrt{1-\gamma}} + \sigma \sqrt{\frac{n}{1-\gamma}} + \sqrt{n} \right) \sqrt{\mu}. \quad (61)$$

This concludes the proof of Lemma 12. \square

Lemma 13. *Let the sketching matrix \mathbf{W} satisfy the conditions of eqns. (6) and (7). Then, after $t \geq \frac{\log(4\sqrt{6n}\psi/\gamma\sigma)}{\log(1/\zeta)}$ iterations of the CG solver in Algorithm 1,*

$$\|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \frac{\gamma\sigma}{4\sqrt{n}} \sqrt{\mu} \quad \text{and} \quad \|\mathbf{v}\|_2 \leq \frac{\gamma\sigma}{4} \mu.$$

Here $\psi = \left(\frac{9n}{\sqrt{1-\gamma}} + \sigma \sqrt{\frac{n}{1-\gamma}} + \sqrt{n} \right)$ and $\tilde{\mathbf{f}}^{(t)} = \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{D}^2 \mathbf{A}^\top \mathbf{Q}^{-1/2} \tilde{\mathbf{z}}^t - \mathbf{Q}^{-1/2} \mathbf{p}$ is the residual of the CG solver.

Proof. Combining Lemma 12 and the condition in eqn. (7), we get

$$\|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \zeta^t \psi \sqrt{2\mu}. \quad (62)$$

Now, $\|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \frac{\gamma\sigma}{4\sqrt{n}} \sqrt{\mu}$ holds if $\sqrt{2}\psi \zeta^t \sqrt{\mu} \leq \frac{\gamma\sigma}{4\sqrt{n}} \sqrt{\mu}$, which holds if $\left(\frac{1}{\zeta}\right)^t \geq \frac{4\sqrt{2n}\psi}{\gamma\sigma}$. The last inequality holds for our choice of t . Next, combining Lemma 5 and eqn. (62) we get

$$\|\mathbf{v}\|_2 \leq \sqrt{3n\mu} \|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \sqrt{6n} \zeta^t \psi \mu$$

Therefore, $\|\mathbf{v}\|_2 \leq \frac{\gamma\sigma\mu}{4}$ holds if $\sqrt{6n}\psi \zeta^t \mu \leq \frac{\gamma\sigma\mu}{4}$, which holds for our choice of t . Now, fixing γ , σ , and ζ , after $t = \mathcal{O}(\log n)$ iterations of Algorithm 1 the conclusions of the lemma hold. \square

C.3 Determining step-size, bounding the number of iterations, and proof of Theorem 1

Assume that the triplet $(\hat{\Delta}\mathbf{x}, \hat{\Delta}\mathbf{y}, \hat{\Delta}\mathbf{s})$ satisfies eqns. (23), (24) and (25). We rewrite this system in the following alternative form:

$$\mathbf{A} \hat{\Delta}\mathbf{x} = -\mathbf{r}_p, \quad (63a)$$

$$\mathbf{A}^\top \hat{\Delta}\mathbf{y} + \hat{\Delta}\mathbf{s} = -\mathbf{r}_d, \quad (63b)$$

$$\mathbf{X}\hat{\Delta}\mathbf{s} + \mathbf{S}\hat{\Delta}\mathbf{x} = -\mathbf{X}\mathbf{S}\mathbf{1}_n + \sigma\mu\mathbf{1}_n - \mathbf{v}. \quad (63c)$$

Indeed, we now show how to derive eqns. (23), (24) and (25) from eqn. (63). Pre-multiplying both sides of eqn. (63c) by $\mathbf{A}\mathbf{S}^{-1}$ and noting that $\mathbf{D}^2 = \mathbf{X}\mathbf{S}^{-1}$, we get

$$\begin{aligned} \mathbf{A}\mathbf{D}^2\hat{\Delta}\mathbf{s} + \mathbf{A}\hat{\Delta}\mathbf{x} &= -\mathbf{A}\mathbf{X}\mathbf{1}_n + \sigma\mu\mathbf{A}\mathbf{S}^{-1}\mathbf{1}_n - \mathbf{A}\mathbf{S}^{-1}\mathbf{v} \\ \Rightarrow \mathbf{A}\mathbf{D}^2\hat{\Delta}\mathbf{s} &= -\mathbf{A}\mathbf{x} + \mathbf{r}_p + \sigma\mu\mathbf{A}\mathbf{S}^{-1}\mathbf{1}_n - \mathbf{A}\mathbf{S}^{-1}\mathbf{v}. \end{aligned} \quad (64)$$

Eqn. (64) holds as $\mathbf{A}\mathbf{X}\mathbf{1}_n = \mathbf{A}\mathbf{x}$ and, from eqn. (63a), $\mathbf{A}\hat{\Delta}\mathbf{x} = -\mathbf{r}_p$. Next, pre-multiplying eqn. (63b) by $\mathbf{A}\mathbf{D}^2$, we get

$$\begin{aligned} \mathbf{A}\mathbf{D}^2\mathbf{A}^\top\hat{\Delta}\mathbf{y} + \mathbf{A}\mathbf{D}^2\hat{\Delta}\mathbf{s} &= -\mathbf{A}\mathbf{D}^2\mathbf{r}_d \\ \Rightarrow \mathbf{A}\mathbf{D}^2\mathbf{A}^\top\hat{\Delta}\mathbf{y} &= -\mathbf{r}_p - \sigma\mu\mathbf{A}\mathbf{S}^{-1}\mathbf{1}_n + \mathbf{A}\mathbf{x} - \mathbf{A}\mathbf{D}^2\mathbf{r}_d + \mathbf{A}\mathbf{S}^{-1}\mathbf{v} = \mathbf{p} + \mathbf{A}\mathbf{S}^{-1}\mathbf{v}. \end{aligned} \quad (65)$$

The first equality in eqn. (65) follows from eqn. (64) and the definition of \mathbf{p} . This establishes eqn. (24). Eqn. (25) directly follows from eqn. (63b). Finally, we get eqn. (23) by pre-multiplying eqn. (63c) by \mathbf{S}^{-1} .

Next, we define each new point traversed by the algorithm as $(\mathbf{x}(\alpha), \mathbf{y}(\alpha), \mathbf{s}(\alpha))$, where

$$(\mathbf{x}(\alpha), \mathbf{y}(\alpha), \mathbf{s}(\alpha)) = (\mathbf{x}, \mathbf{y}, \mathbf{s}) + \alpha(\hat{\Delta}\mathbf{x}, \hat{\Delta}\mathbf{y}, \hat{\Delta}\mathbf{s}), \quad (66)$$

$$\mu(\alpha) = \mathbf{x}(\alpha)^\top \mathbf{s}(\alpha) / n, \quad (67)$$

$$\mathbf{r}(\alpha) = \mathbf{r}(\mathbf{x}(\alpha), \mathbf{s}(\alpha), \mathbf{y}(\alpha)). \quad (68)$$

The goal in this section is to bound the number of iterations required by Algorithm 2. Towards that end, we bound the magnitude of the step size α . First, we provide an upper bound on α , which allows us to show that each new point $(\mathbf{x}(\alpha), \mathbf{s}(\alpha), \mathbf{y}(\alpha))$ traversed by the algorithm stays within the neighborhood $\mathcal{N}(\gamma)$. Second, we provide a lower bound on α , which allows us to bound the number of iterations required. We use multiple lemmas from [35], which we reproduce here, without their proofs.

First, we provide an upper bound on α , ensuring that each new point $(\mathbf{x}(\alpha), \mathbf{y}(\alpha), \mathbf{s}(\alpha))$ traversed by the algorithm stays within the neighborhood $\mathcal{N}(\gamma)$.

Lemma 14 (Lemma 3.5 of [35]). *Assume $(\hat{\Delta}\mathbf{x}, \hat{\Delta}\mathbf{y}, \hat{\Delta}\mathbf{s})$ satisfies eqns. (63) for some $\sigma > 0$, $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$ (for $\gamma \in (0, 1)$), and $\|\mathbf{v}\|_2 \leq \frac{\gamma\sigma\mu}{4}$. Then, $(\mathbf{x}(\alpha), \mathbf{y}(\alpha), \mathbf{s}(\alpha)) \in \mathcal{N}(\gamma)$ for every scalar α such that*

$$0 \leq \alpha \leq \min \left\{ 1, \frac{\gamma\sigma\mu}{4\|\hat{\Delta}\mathbf{x} \circ \hat{\Delta}\mathbf{s}\|_\infty} \right\}. \quad (69)$$

We now provide a lower bound on the values of $\bar{\alpha}$ and the corresponding $\mu(\bar{\alpha})$; see Algorithm 2.

Lemma 15 (Lemma 3.6 of [35]). *In each iteration of Algorithm 2, if $\|\mathbf{v}\|_2 \leq \frac{\gamma\sigma\mu}{4}$, then the step size $\bar{\alpha}$ satisfies*

$$\bar{\alpha} \geq \min \left\{ 1, \frac{\min\{\gamma\sigma, (1 - \frac{5}{4}\sigma)\}\mu}{4\|\hat{\Delta}\mathbf{x} \circ \hat{\Delta}\mathbf{s}\|_\infty} \right\} \quad (70)$$

and

$$\mu(\bar{\alpha}) = \left[1 - \frac{\bar{\alpha}}{2} \left(1 - \frac{5}{4}\sigma \right) \right] \mu. \quad (71)$$

At this point, we have provided a lower bound (eqn. (70)) for the allowed values of the step size $\bar{\alpha}$. Next, we show that this lower bound is bounded away from zero. From eqn. (70) this is equivalent to showing that $\|\hat{\Delta}\mathbf{x} \circ \hat{\Delta}\mathbf{s}\|_\infty$ is bounded.

Lemma 16 (Lemma 3.7 of [35] (slightly modified)). *Let $(\mathbf{x}^0, \mathbf{y}^0, \mathbf{s}^0)$ be the initial point with $(\mathbf{x}^0, \mathbf{s}^0) > 0$ and $(\mathbf{x}^0, \mathbf{s}^0) \geq (\mathbf{x}^*, \mathbf{s}^*)$ for some $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{s}^*) \in \mathcal{S}$. Let $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$ be such that $\mathbf{r} = \eta \mathbf{r}^0$ for some $\eta \in [0, 1]$ and $\|\mathbf{v}\|_2 \leq \frac{\gamma\sigma\mu}{4}$. Then, the search direction $(\hat{\Delta}\mathbf{x}, \hat{\Delta}\mathbf{y}, \hat{\Delta}\mathbf{s})$ produced by Algorithm 2 at each iteration satisfies*

$$\max\{\|\mathbf{D}^{-1}\hat{\Delta}\mathbf{x}\|_2, \|\mathbf{D}\hat{\Delta}\mathbf{s}\|_2\} \leq \left(1 + \frac{\sigma^2}{1-\gamma} - 2\sigma\right)^{1/2} \sqrt{n\mu} + \frac{6n}{\sqrt{(1-\gamma)}}\sqrt{\mu} + \frac{\gamma\sigma}{4\sqrt{1-\gamma}}\sqrt{\mu}. \quad (72)$$

We should note here that the above lemma is slightly different than Lemma 3.7 of [35]. Indeed, Lemma 3.7 of [35] actually proves the following bound:

$$\max\{\|\mathbf{D}^{-1}\hat{\Delta}\mathbf{x}\|_2, \|\mathbf{D}\hat{\Delta}\mathbf{s}\|_2\} \leq \left(1 + \frac{\sigma^2}{1-\gamma} - 2\sigma\right)^{1/2} \sqrt{n\mu} + \frac{6n}{\sqrt{(1-\gamma)}}\sqrt{\mu} + \frac{\gamma\sigma}{4\sqrt{n}}\sqrt{\mu}. \quad (73)$$

Notice that there is slight difference in the last term in the right-hand side, which does not asymptotically change the bound. The underlying reason for this difference is the fact that [35] constructed the vector \mathbf{v} differently. In our case, we need to bound $\|(\mathbf{X}\mathbf{S})^{-1/2}\mathbf{v}\|_2$, which we do as follows:

$$\|(\mathbf{X}\mathbf{S})^{-1/2}\mathbf{v}\|_2 \leq \|(\mathbf{X}\mathbf{S})^{-1/2}\|_2 \|\mathbf{v}\|_2 \leq \frac{1}{\min_i \sqrt{x_i s_i}} \frac{\gamma\sigma\mu}{4}, \quad (74)$$

where in the above expression we use the fact that $\|(\mathbf{X}\mathbf{S})^{-1/2}\|_2 = \frac{1}{\min_i \sqrt{x_i s_i}}$. Now as $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$, we further have $x_i s_i \geq (1-\gamma)\mu$ for all $i = 1 \dots n$. Combining this with eqn. (74), we get

$$\|(\mathbf{X}\mathbf{S})^{-1/2}\mathbf{v}\|_2 \leq \frac{\gamma\sigma\mu}{4\sqrt{(1-\gamma)\mu}} = \frac{\gamma\sigma}{4\sqrt{1-\gamma}}\sqrt{\mu}. \quad (75)$$

On the other hand, [35] had a different construction of \mathbf{v} for which $\|(\mathbf{X}\mathbf{S})^{-1/2}\mathbf{v}\|_2 = \|\tilde{\mathbf{f}}^{(t)}\|_2$ holds. Therefore they had the following bound:

$$\|(\mathbf{X}\mathbf{S})^{-1/2}\mathbf{v}\|_2 = \|\tilde{\mathbf{f}}^{(t)}\|_2 \leq \frac{\gamma\sigma}{4\sqrt{n}}\sqrt{\mu}.$$

The next lemma bounds the number of iterations that Algorithm 2 needs when started with an infeasible point that is sufficiently positive.

Lemma 17 (Theorem 2.6 of [35]). *Assume that the constants γ and σ are such that $\max\{\gamma^{-1}, (1-\gamma)^{-1}, \sigma^{-1}, (1-\frac{5}{4}\sigma)^{-1}\} = \mathcal{O}(1)$. Let the initial point $(\mathbf{x}^0, \mathbf{s}^0, \mathbf{y}^0)$ satisfy $(\mathbf{x}^0, \mathbf{s}^0) \geq (\mathbf{x}^*, \mathbf{s}^*)$ for some $(\mathbf{x}^*, \mathbf{s}^*, \mathbf{y}^*) \in \mathcal{S}$ and $\|\mathbf{v}\|_2 \leq \frac{\gamma\sigma\mu}{4}$. Algorithm 2 generates an iterate $(\mathbf{x}^k, \mathbf{s}^k, \mathbf{y}^k)$ satisfying $\mu_k \leq \epsilon\mu_0$ and $\|\mathbf{r}^k\|_2 \leq \epsilon\|\mathbf{r}^0\|_2$ after $\mathcal{O}(n^2 \log 1/\epsilon)$ iterations.*

Finally, Theorem 1 follows from Lemmas 13 and 17.

Appendix D Additional notes on experiments

D.1 Support Vector Machines (SVMs)

The classical ℓ_1 -SVM problem is as follows. We consider the task of fitting an SVM to data pairs $S = \{(x_i, y_i)\}_{i=1}^m$, where $x_i \in \mathbb{R}^n$ and $y_i \in \{+1, -1\}$. Here, m is the number of training points, and n is the feature dimension. The SVM problem with an ℓ_1 regularizer has the following form:

$$\begin{aligned} & \underset{w}{\text{minimize}} && \|w\|_1 \\ & \text{subject to} && y_i(w^T x_i + b') \geq 1, \quad i = 1 \dots m. \end{aligned} \quad (76)$$

This problem can be written as an LP by introducing the variables w^+ and w^- , where $w = w^+ - w^-$. The objective becomes $\sum_{j=1}^n w_j^+ + w_j^-$, and we constrain $w_i^+ \geq 0$ and $w_i^- \geq 0$. Note that the size of the constraint matrix in the LP becomes $m \times (2n + 1)$.

D.2 Random data

We generate random synthetic instances of linear programs as follows. To generate $\mathbf{A} \in \mathbb{R}^{m \times n}$, we set $a_{ij} \sim_{i.i.d.} U(0, 1)$ with probability p and $a_{ij} = 0$ otherwise. We then add $\min\{m, n\}$ i.i.d. draws from $U(0, 1)$ to the main diagonal, to ensure each row of \mathbf{A} has at least one nonzero entry. We set $\mathbf{b} = \mathbf{A}\mathbf{x} + 0.1\mathbf{z}$, where \mathbf{x} and \mathbf{z} are random vectors drawn from $N(0, 1)$. Finally, we set $c \sim N(0, 1)$.

D.3 Real-world data

We used a gene expression cancer RNA-Sequencing dataset, taken from the UCI Machine Learning repository. It is part of the RNA-Seq (HiSeq) PANCAN data set [47] and is a random extraction of gene expressions from patients who have different types of tumors: BRCA, KIRC, COAD, LUAD, and PRAD. We considered the binary classification task of identifying BRCA versus other types.

We also used the DrivFace dataset taken from the UCI Machine Learning repository. In the DrivFace dataset, each sample corresponds to an image of a human subject, taken while driving in real scenarios. Each image is labeled as corresponding to one of three possible gaze directions: left, straight, or right. We considered the binary classification task of identifying two different gaze directions: straight, or to either side (left or right).