

# Quantum algorithms for Second-Order Cone Programming and Support Vector Machines

Iordanis Kerenidis<sup>1,2</sup>, Anupam Prakash<sup>1,2</sup>, and Dániel Szilágyi<sup>2</sup>

<sup>1</sup>QCWare, Palo Alto, California

<sup>2</sup>Université de Paris, CNRS, IRIF, F-75006, Paris, France

We present a quantum interior-point method (IPM) for second-order cone programming (SOCP) that runs in time  $\tilde{O}\left(n\sqrt{r}\frac{\zeta\kappa}{\delta^2}\log(1/\epsilon)\right)$  where  $r$  is the rank and  $n$  the dimension of the SOCP,  $\delta$  bounds the distance of intermediate solutions from the cone boundary,  $\zeta$  is a parameter upper bounded by  $\sqrt{n}$ , and  $\kappa$  is an upper bound on the condition number of matrices arising in the classical IPM for SOCP. The algorithm takes as its input a suitable quantum description of an arbitrary SOCP and outputs a classical description of a  $\delta$ -approximate  $\epsilon$ -optimal solution of the given problem.

Furthermore, we perform numerical simulations to determine the values of the aforementioned parameters when solving the SOCP up to a fixed precision  $\epsilon$ . We present experimental evidence that in this case our quantum algorithm exhibits a polynomial speedup over the best classical algorithms for solving general SOCPs that run in time  $O(n^{\omega+0.5})$  (here,  $\omega$  is the matrix multiplication exponent, with a value of roughly 2.37 in theory, and up to 3 in practice). For the case of random SVM (support vector machine) instances of size  $O(n)$ , the quantum algorithm scales as  $O(n^k)$ , where the exponent  $k$  is estimated to be 2.59 using a least-squares power law. On the same family random instances, the estimated scaling exponent for an external SOCP solver is 3.31 while that for a state-of-the-art SVM solver is 3.11.

## 1 Introduction

It is well known that many interesting and relevant optimization problems in the domain of Machine Learning can be expressed in the framework of convex optimization [6, 10]. The landmark result in this area was the discovery of interior-point methods (IPM) by [23], and their subsequent generalization to all “self-scaled” (i.e. symmetric) cones by [29, 30]. Very recently, [13] have shown that it is possible to solve linear programs (LP) in  $\tilde{O}(n^\omega)$ , the time it takes to multiply two matrices (as long as  $\omega \geq 2 + 1/6$ , which is currently the case). This result has been further extended in [27] to a slightly more general class of cones, however, their techniques did not yield improved complexities for second-order (SOCP) and semidefinite programming (SDP). An efficient algorithm for SOCP would also yield efficient algorithms for many interesting problems, such as (standard and quadratically-constrained) convex quadratic programming, portfolio optimization, and many others [1].

Iordanis Kerenidis: [jkeren@irif.fr](mailto:jkeren@irif.fr)

Anupam Prakash: [anupamprakash1@gmail.com](mailto:anupamprakash1@gmail.com)

Dániel Szilágyi: [dszilagy@irif.fr](mailto:dszilagy@irif.fr)

Starting with the landmark results of [20, 35], and, more recently, [21], it has been demonstrated that quantum computers offer significant (sometimes even exponential) asymptotic speedups for a number of important problems. More recently, there has been substantial work in the area of convex optimization. Quantum speedups for gradient descent were investigated by [17], whereas [7, 9, 40, 41] presented quantum algorithms for SDP based on the the multiplicative weights framework of [4]. However, it has been difficult to establish asymptotic speedups for this family of quantum SDP solvers as their the running time depends on problem-specific parameters, including a 5th-power dependence on the width of the SDP. Interestingly, the recent result of [8] suggests that such a speedup might be obtained when applying an SDP algorithm of this type to some low-precision instances of quadratic binary optimization.

In an orthogonal approach, [26] proposed a quantum algorithm for LPs and SDPs by quantizing a variant of the classical interior point method and using the state of the art quantum linear algebra tools [11, 18] – in particular, the matrix multiplication and inversion algorithms whose running time is sub-linear in the input size. However, the complexity of this algorithm depends on the condition number of  $O(n^2)$ -sized matrices that is difficult to bound theoretically. It therefore remains an open question to find an end-to-end optimization problems for which quantum SDP solvers achieve an asymptotic speedup over state of the art classical algorithms. In this paper, we propose support vector machines (SVM) as a candidate for such an end-to-end quantum speedup using a quantum interior point method based algorithm.

## 1.1 Our results and techniques

In this section, we provide a high level sketch of our results and the techniques used for the quantum interior point method for SOCPs, we begin by discussing the differences between classical and quantum interior point methods.

A classical interior point method solves an optimization problem over symmetric cones by starting with a feasible solution and iteratively finding solutions with a smaller duality gap while maintaining feasibility. A single iterative step consists of solving a system of linear equations called the Newton linear system and updating the current iterate using the solutions of the Newton system. The analysis of the classical IPM shows that in each iteration, the updated solutions remain feasible and the duality gap is decreased by a factor of  $(1 - \alpha/\sqrt{n})$  where  $n$  is the dimension of the optimization problem and  $\alpha > 0$  is a constant. The algorithm therefore converges to a feasible solution with duality gap  $\epsilon$  in  $O(\sqrt{n} \log(1/\epsilon))$  iterations.

A quantum interior point method [26] uses a quantum linear system solver instead of classical one in each iteration of the IPM. However, there is an important difference between classical and quantum linear algebra procedures for solving the linear system  $A\mathbf{x} = \mathbf{b}$ . Unlike classical linear system solvers which return an exact description of  $\mathbf{x}$ , quantum tomography procedures can return an  $\epsilon$ -accurate solution  $\tilde{\mathbf{x}}$  such that  $\|\tilde{\mathbf{x}} - \mathbf{x}\| \leq \epsilon\|\mathbf{x}\|$  with  $O(n/\epsilon^2)$  runs of the quantum linear system solver. Additionally, these linear system solvers require  $A$  and  $\mathbf{b}$  to be given as block-encodings [11], so this input model is used by our algorithm as well. One of the main technical challenges in developing a quantum interior point method (QIPM) is to establish convergence of the classical IPM which uses  $\epsilon$ -approximate solutions of the Newton linear system (in the  $\ell_2$  norm) instead of the exact solutions in the classical IPM.

The quantum interior point method for second-order cone programs requires additional ideas going beyond those used for the quantum interior point methods for SDP [26]. Second order cone programs are optimization problems over the product of second-order or

Lorentz cones (see section 2.1 for definitions), interior point methods for SOCP can be described using the Euclidean Jordan algebra framework [28]. The Euclidean Jordan algebra framework provides analogs of concepts like eigenvalues, spectral and Frobenius norms for matrices and positive semidefinite constraints for the case of SOCPs. Using these conceptual ideas from the Euclidean Jordan algebra framework [28] and the analysis of the approximate SDP interior point method [26] we provide an approximate IPM for SOCP that converges in  $O(\sqrt{n} \log(1/\epsilon))$  iterations. Approximate IPMs for SOCP have not been previously investigated in the classical or the quantum optimization literature, this analysis is one of the main technical contributions of this paper.

From an algorithmic perspective, SOCPs are much closer to LPs (Linear Programs) than to SDPs, since for cones of dimension  $n$ , the Newton linear systems arising in LP and SOCP IPMs are of size  $O(n)$ , whereas in the SDP case they are of size  $O(n^2)$ . Namely, a second-order conic constraint of dimension  $n$  can be expressed as a single PSD constraint on a (sparse)  $n \times n$  matrix [1] – this allows us to embed an SOCP with  $n$  variables and  $m$  constraints in an SDP with an  $n \times n$  matrix and  $m$  constraints. The cost of solving that SDP would have a worse dependence on the error [40] or the input size [26]. On the other hand, the quantum representations (block encodings) of the Newton linear systems for SOCP are also much simpler to construct than those for SDP. The smaller size of the SOCP linear system also makes it feasible to empirically estimate the condition number for these linear systems in a reasonable amount of time allowing us to carry out extensive numerical experiments to validate the running time of the quantum algorithm.

The theoretical analysis of the quantum algorithm for SOCP shows that its worst-case running time is

$$\tilde{O}\left(n\sqrt{r}\frac{\zeta\kappa}{\delta^2}\log\left(\frac{1}{\epsilon}\right)\right), \quad (1)$$

where  $r$  is the rank and  $n$  the dimension of the SOCP,  $\delta$  bounds the distance of intermediate solutions from the cone boundary,  $\zeta$  is a parameter bounded by  $\sqrt{n}$ ,  $\kappa$  is an upper bound on the condition number of matrices arising in the interior-point method for SOCPs, and  $\epsilon$  is the target duality gap. The running time of the algorithm depends on problem dependent parameters like  $\kappa$  and  $\delta$  that are difficult to bound in terms of the problem dimension  $n$  – this is also the case with previous quantum SDP solvers [26, 40] and makes it important to validate the quantum optimization speedups empirically. Interestingly, since we require a classical solution of the Newton system, the linear system solver could also be replaced by a classical iterative solver [34] which would yield a complexity of  $O(n^2\sqrt{r}\kappa\log(n/\epsilon))$ .

Let us make a remark about the complexity: as it is the case with all approximation algorithms, (1) depends on the inverse of the target duality gap  $\epsilon$ , thus the running time of our algorithm grows to infinity as  $\epsilon$  approaches zero, as in the case of classical IPM. Our running time also depends on  $\kappa$ , which in turn is empirically observed to grow inversely with the duality gap (in particular as  $O(1/\epsilon)$ ) which again makes the running time go to infinity as  $\epsilon$  approaches zero. The quantum IPM is a low precision method, unlike the classical IPM, and it can offer speedups for settings where the desired precision  $\epsilon$  is moderate or low. Thus although at first glance it seems that the  $\epsilon$ -dependence in (1) is logarithmic, experimental evidence suggests that the factor  $\frac{\zeta\kappa}{\delta^2}$  depends polynomially on  $1/\epsilon$ .

Support Vector Machines (SVM) is an important application in machine learning, where even a modest value of  $\epsilon = 0.1$  yields an almost optimal classifier. Since the SVM training problem can be reduced to SOCP, the quantum IPM for SOCP can be used to obtain an efficient quantum SVM algorithm. We perform extensive numerical experiments to evaluate our algorithm on random SVM instances and compare it against state of the art

classical SOCP and SVM solvers.

The numerical experiments on random SVM instances indicate that the running time of the quantum algorithm scales as roughly  $O(n^{2.591})$ , where all the parameters in the running time are taken into account and the exponent is estimated using a least squares fit. We also benchmarked the exponent for classical SVM algorithms on the same instances and for a comparable accuracy, the scaling exponent was found to be 3.31 for general SOCP solvers and 3.11 for state-of-the-art SVM solvers. We note that this does not prove a worst-case asymptotic speedup, but the experiments on unstructured SVM instances provide strong evidence for a significant polynomial speedup of the quantum SVM algorithm over state-of-the-art classical SVM algorithms. We can therefore view SVMs as a candidate problem for which quantum optimization algorithms can achieve a polynomial speedup over state of the art classical algorithms for an end-to-end application.

## 1.2 Related work

Our main result is the first specialized quantum algorithm for training support-vector machines (SVM). While several quantum SVM algorithms have been proposed, they are unlikely to offer general speedups for the most widely used formulation of SVM – the soft-margin ( $\ell_1$ -)SVM (defined in eq. (10)). On one hand, papers such as [3] formulate the SVM as a SDP, and solve that using existing quantum SDP solvers such as [7, 9, 40, 41] – with the conclusion being that a speedup is observed only for very specific sparse instances. On the other hand, [33] solves an easier related problem – the least-squares SVM ( $\ell_2$ -SVM or LS-SVM, see eq. (11) for its formulation), thus losing the desirable sparsity properties of  $\ell_1$ -SVM [38]. It turns out that applying our algorithm to this problem also yields the same complexity as in [33] for the  $\ell_2$ -SVM. Very recently, a quantum algorithm for yet another variant of SVM (SVM-perf, [22]) has been presented in [2].

## 2 Preliminaries

### 2.1 Second-order cone programming

For the sake of completeness, in this section we present the most important results about classical SOCP IPMs, from [1, 28]. We start by defining SOCP as the optimization problem over the product of second-order (or Lorentz) cones  $\mathcal{L} = \mathcal{L}^{n_1} \times \dots \times \mathcal{L}^{n_r}$ , where  $\mathcal{L}_k \subseteq \mathbb{R}^k$  is defined as  $\mathcal{L}^k = \{ \mathbf{x} = (x_0; \tilde{\mathbf{x}}) \in \mathbb{R}^k \mid \|\tilde{\mathbf{x}}\| \leq x_0 \}$ . In this paper we consider the problem (2) and its dual (3):

$$\begin{aligned} \min \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & A\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \in \mathcal{L}, \end{aligned} \quad (2) \qquad \begin{aligned} \max \quad & \mathbf{b}^\top \mathbf{y} \\ \text{s.t.} \quad & A^\top \mathbf{y} + \mathbf{s} = \mathbf{c} \\ & \mathbf{s} \in \mathcal{L}, \mathbf{y} \in \mathbb{R}^m. \end{aligned} \quad (3)$$

We call  $n := \sum_{i=1}^r n_i$  the *size* of the SOCP (2), and  $r$  is its *rank*.

A solution  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$  satisfying the constraints of both (2) and (3) is *feasible*, and if in addition it satisfies  $\mathbf{x} \in \text{int } \mathcal{L}$  and  $\mathbf{s} \in \text{int } \mathcal{L}$ , it is *strictly feasible*. If at least one constraint of (2) or (3), is violated, the solution is *infeasible*. The duality gap of a feasible solution  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$  is defined as  $\mu := \mu(\mathbf{x}, \mathbf{s}) := \frac{1}{r} \mathbf{x}^\top \mathbf{s}$ . As opposed to LP, and similarly to SDP, strict feasibility is required for *strong duality* to hold [1]. From now on, we assume that our SOCP has a strictly feasible primal-dual solution (this assumption is valid, since the homogeneous self-dual embedding technique from [42] allows us to embed (2) and (3) in a slightly larger SOCP where this condition is satisfied).

## 2.2 Euclidean Jordan algebras

The cone  $\mathcal{L}^n$  has an algebraic structure similar to that of symmetric matrices under the matrix product. Here, we consider the Jordan product of  $(x_0, \tilde{\mathbf{x}}) \in \mathbb{R}^n$  and  $(y_0, \tilde{\mathbf{y}}) \in \mathbb{R}^n$ , defined as

$$\mathbf{x} \circ \mathbf{y} := \begin{bmatrix} \mathbf{x}^T \mathbf{y} \\ x_0 \tilde{\mathbf{y}} + y_0 \tilde{\mathbf{x}} \end{bmatrix}, \text{ and its identity element } \mathbf{e} := \begin{bmatrix} 1 \\ 0^{n-1} \end{bmatrix}.$$

This product is closely related to the (linear) matrix representation  $\text{Arw}(\mathbf{x}) := \begin{bmatrix} x_0 & \tilde{\mathbf{x}}^T \\ \tilde{\mathbf{x}} & x_0 I_{n-1} \end{bmatrix}$ , which in turn satisfies the following equality:

$$\mathbf{x} \circ \mathbf{y} = \text{Arw}(\mathbf{x})\mathbf{y} = \text{Arw}(\mathbf{x}) \text{Arw}(\mathbf{y})\mathbf{e}.$$

The key observation is that this product induces a spectral decomposition of any vector  $\mathbf{x}$ , that has similar properties as its matrix relative. Namely, for any vector  $\mathbf{x}$  we define

$$\begin{aligned} \lambda_1(\mathbf{x}) &:= x_0 + \|\tilde{\mathbf{x}}\|, & \mathbf{c}_1(\mathbf{x}) &:= \frac{1}{2} \begin{bmatrix} 1 \\ \frac{\tilde{\mathbf{x}}}{\|\tilde{\mathbf{x}}\|} \end{bmatrix}, \\ \lambda_2(\mathbf{x}) &:= x_0 - \|\tilde{\mathbf{x}}\|, & \mathbf{c}_2(\mathbf{x}) &:= \frac{1}{2} \begin{bmatrix} 1 \\ -\frac{\tilde{\mathbf{x}}}{\|\tilde{\mathbf{x}}\|} \end{bmatrix}. \end{aligned} \quad (4)$$

We use the shorthands  $\lambda_1 := \lambda_1(\mathbf{x})$ ,  $\lambda_2 := \lambda_2(\mathbf{x})$ ,  $\mathbf{c}_1 := \mathbf{c}_1(\mathbf{x})$  and  $\mathbf{c}_2 := \mathbf{c}_2(\mathbf{x})$  whenever  $\mathbf{x}$  is clear from the context, so we observe that  $\mathbf{x} = \lambda_1 \mathbf{c}_1 + \lambda_2 \mathbf{c}_2$ . The set of eigenvectors  $\{\mathbf{c}_1, \mathbf{c}_2\}$  is called the *Jordan frame* of  $\mathbf{x}$ , and satisfies several properties:

**Proposition 1** (Properties of Jordan frames). *Let  $\mathbf{x} \in \mathbb{R}^n$  and let  $\{\mathbf{c}_1, \mathbf{c}_2\}$  be its Jordan frame. Then, the following holds:*

1.  $\mathbf{c}_1 \circ \mathbf{c}_2 = 0$  (the eigenvectors are “orthogonal”)
2.  $\mathbf{c}_1 \circ \mathbf{c}_1 = \mathbf{c}_1$  and  $\mathbf{c}_2 \circ \mathbf{c}_2 = \mathbf{c}_2$
3.  $\mathbf{c}_1, \mathbf{c}_2$  are of the form  $(\frac{1}{2}; \pm \tilde{\mathbf{c}})$  with  $\|\tilde{\mathbf{c}}\| = \frac{1}{2}$

On the other hand, just like a given matrix is positive (semi)definite if and only if all of its eigenvalues are positive (nonnegative), a similar result holds for  $\mathcal{L}^n$  and  $\text{int } \mathcal{L}^n$  (the Lorentz cone and its interior):

**Proposition 2.** *Let  $\mathbf{x} \in \mathbb{R}^n$  have eigenvalues  $\lambda_1, \lambda_2$ . Then, the following holds:*

1.  $\mathbf{x} \in \mathcal{L}^n$  if and only if  $\lambda_1 \geq 0$  and  $\lambda_2 \geq 0$ .
2.  $\mathbf{x} \in \text{int } \mathcal{L}^n$  if and only if  $\lambda_1 > 0$  and  $\lambda_2 > 0$ .

Now, using this decomposition, we can define arbitrary real powers  $\mathbf{x}^p$  for  $p \in \mathbb{R}$  as  $\mathbf{x}^p := \lambda_1^p \mathbf{c}_1 + \lambda_2^p \mathbf{c}_2$ , and in particular the “inverse” and the “square root”

$$\begin{aligned} \mathbf{x}^{-1} &= \frac{1}{\lambda_1} \mathbf{c}_1 + \frac{1}{\lambda_2} \mathbf{c}_2, \text{ if } \lambda_1 \lambda_2 \neq 0, \\ \mathbf{x}^{1/2} &= \sqrt{\lambda_1} \mathbf{c}_1 + \sqrt{\lambda_2} \mathbf{c}_2, \text{ if } \mathbf{x} \in \mathcal{L}^n. \end{aligned}$$

Moreover, we can also define some operator norms, namely the Frobenius and the spectral one:

$$\begin{aligned}\|\mathbf{x}\|_F &= \sqrt{\lambda_1^2 + \lambda_2^2} = \sqrt{2}\|\mathbf{x}\|, \\ \|\mathbf{x}\|_2 &= \max\{|\lambda_1|, |\lambda_2|\} = |x_0| + \|\tilde{\mathbf{x}}\|.\end{aligned}$$

Finally, we define an analogue to the operation  $Y \mapsto XYX$ . It turns out that for this we need another matrix representation (*quadratic representation*)  $Q_{\mathbf{x}}$ , defined as

$$Q_{\mathbf{x}} := 2 \operatorname{Arw}^2(\mathbf{x}) - \operatorname{Arw}(\mathbf{x}^2) = \begin{bmatrix} \|\mathbf{x}\|^2 & 2x_0\tilde{\mathbf{x}}^T \\ 2x_0\tilde{\mathbf{x}} & \lambda_1\lambda_2 I_n + 2\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T \end{bmatrix}. \quad (5)$$

Now, the matrix-vector product  $Q_{\mathbf{x}}\mathbf{y}$  will behave as the quantity  $XYX$ . To simplify the notation, we also define the matrix  $T_{\mathbf{x}} := Q_{\mathbf{x}^{1/2}}$ .

The definitions that we introduced so far are suitable for dealing with a single constraint  $\mathbf{x} \in \mathcal{L}^n$ . For dealing with multiple constraints  $\mathbf{x}_1 \in \mathcal{L}^{n_1}, \dots, \mathbf{x}_r \in \mathcal{L}^{n_r}$ , we need to deal with block-vectors  $\mathbf{x} = (\mathbf{x}_1; \mathbf{x}_2; \dots; \mathbf{x}_r)$  and  $\mathbf{y} = (\mathbf{y}_1; \mathbf{y}_2; \dots; \mathbf{y}_r)$ . We call the number of blocks  $r$  the *rank* of the vector (thus, up to now, we were only considering rank-1 vectors). Now, we extend all our definitions to rank- $r$  vectors.

1.  $\mathbf{x} \circ \mathbf{y} := (\mathbf{x}_1 \circ \mathbf{y}_1; \dots; \mathbf{x}_r \circ \mathbf{y}_r)$
2. The matrix representations  $\operatorname{Arw}(\mathbf{x})$  and  $Q_{\mathbf{x}}$  are the block-diagonal matrices containing the representations of the blocks:

$$\operatorname{Arw}(\mathbf{x}) := \operatorname{Arw}(\mathbf{x}_1) \oplus \dots \oplus \operatorname{Arw}(\mathbf{x}_r) \text{ and } Q_{\mathbf{x}} := Q_{\mathbf{x}_1} \oplus \dots \oplus Q_{\mathbf{x}_r}$$

3.  $\mathbf{x}$  has  $2r$  eigenvalues (with multiplicities) – the union of the eigenvalues of the blocks  $\mathbf{x}_i$ . The eigenvectors of  $\mathbf{x}$  corresponding to block  $i$  contain the eigenvectors of  $\mathbf{x}_i$  as block  $i$ , and are zero everywhere else.
4. The identity element is  $\mathbf{e} = (\mathbf{e}_1; \dots; \mathbf{e}_r)$ , where  $\mathbf{e}_i$ 's are the identity elements for the corresponding blocks.

Thus, all things defined using eigenvalues can also be defined for rank- $r$  vectors:

1. The norms are extended as  $\|\mathbf{x}\|_F^2 := \sum_{i=1}^r \|\mathbf{x}_i\|_F^2$  and  $\|\mathbf{x}\|_2 := \max_i \|\mathbf{x}_i\|_2$ , and
2. Powers are computed blockwise as  $\mathbf{x}^p := (\mathbf{x}_1^p; \dots; \mathbf{x}_r^p)$  whenever the corresponding blocks are defined.

Some further matrix-algebra inspired properties of block vectors are stated in the following two claims:

**Claim 1** (Algebraic properties). *Let  $\mathbf{x}, \mathbf{y}$  be two arbitrary block-vectors. Then, the following holds:*

1. *The spectral norm is subadditive:  $\|\mathbf{x} + \mathbf{y}\|_2 \leq \|\mathbf{x}\|_2 + \|\mathbf{y}\|_2$ .*
2. *The spectral norm is less than the Frobenius norm:  $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_F$ .*
3. *If  $A$  is a matrix with minimum and maximum singular values  $\sigma_{\min}$  and  $\sigma_{\max}$  respectively, then the norm  $\|A\mathbf{x}\|$  is bounded as  $\sigma_{\min}\|\mathbf{x}\| \leq \|A\mathbf{x}\| \leq \sigma_{\max}\|\mathbf{x}\|$ .*
4. *The minimum eigenvalue of  $\mathbf{x} + \mathbf{y}$  is bounded as  $\lambda_{\min}(\mathbf{x} + \mathbf{y}) \geq \lambda_{\min}(\mathbf{x}) - \|\mathbf{y}\|_2$ .*

5. The following submultiplicativity property holds:  $\|\mathbf{x} \circ \mathbf{y}\|_F \leq \|\mathbf{x}\|_2 \cdot \|\mathbf{y}\|_F$ .

In general, the proofs of these statements are analogous to the matrix case, with a few notable differences: First, the vector spectral norm  $\|\cdot\|_2$  is not actually a norm, since there exist nonzero vectors outside  $\mathcal{L}$  which have zero norm. It is, however, still bounded by the Frobenius norm (just like in the matrix case), which is in fact a proper norm. Secondly, the minimum eigenvalue bound also holds for matrix spectral norms, with the exact same statement. Finally, the last property is reminiscent of the matrix submultiplicativity property  $\|A \cdot B\|_F \leq \|A\|_2 \|B\|_F$ .

We finish with several well-known properties of the quadratic representation  $Q_{\mathbf{x}}$  and  $T_{\mathbf{x}}$ .

**Proposition 3** (Properties of  $Q_{\mathbf{x}}$ , from [1]). *Let  $\mathbf{x} \in \text{int } \mathcal{L}$ . Then, the following holds:*

1.  $Q_{\mathbf{x}}\mathbf{e} = \mathbf{x}^2$ , and thus  $T_{\mathbf{x}}\mathbf{e} = \mathbf{x}$ .
2.  $Q_{\mathbf{x}^{-1}} = Q_{\mathbf{x}}^{-1}$ , and more generally  $Q_{\mathbf{x}^p} = Q_{\mathbf{x}}^p$  for all  $p \in \mathbb{R}$ .
3.  $\|Q_{\mathbf{x}}\|_2 = \|\mathbf{x}\|_2^2$ , and thus  $\|T_{\mathbf{x}}\|_2 = \|\mathbf{x}\|_2$ .
4.  $Q_{\mathbf{x}}$  preserves  $\mathcal{L}$ , i.e.  $Q_{\mathbf{x}}(\mathcal{L}) = \mathcal{L}$  and  $Q_{\mathbf{x}}(\text{int } \mathcal{L}) = \text{int } \mathcal{L}$ .

### 2.3 Interior-point methods

Our algorithm follows the general IPM structure, i.e. it uses Newton's method to solve a sequence of increasingly strict relaxations of the Karush-Kuhn-Tucker (KKT) optimality conditions:

$$\begin{aligned} A\mathbf{x} &= \mathbf{b}, \quad A^\top \mathbf{y} + \mathbf{s} = \mathbf{c} \\ \mathbf{x} \circ \mathbf{s} &= \nu \mathbf{e}, \quad \mathbf{x} \in \mathcal{L}, \mathbf{s} \in \mathcal{L}, \end{aligned} \tag{6}$$

where the parameter  $\nu > 0$  is decreased by a factor  $\sigma < 1$  in each iteration. Since  $\mathbf{x} \circ \mathbf{s} = \nu \mathbf{e}$  implies that the duality gap is  $\mu = \nu$ , by letting  $\nu \rightarrow 0$  the IPM converges towards the optimal solution. The curve traced by (feasible) solutions  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$  of (6) for  $\nu > 0$  is called the *central path*, and we note that all points on it are strictly feasible.

More precisely, in each iteration we need to find  $\Delta \mathbf{x}, \Delta \mathbf{y}, \Delta \mathbf{s}$  such that  $\mathbf{x}_{\text{next}} := \mathbf{x} + \Delta \mathbf{x}$ ,  $\mathbf{y}_{\text{next}} := \mathbf{y} + \Delta \mathbf{y}$  and  $\mathbf{s}_{\text{next}} := \mathbf{s} + \Delta \mathbf{s}$  satisfy (6) for  $\nu = \sigma \mu$ . After linearizing the product  $\mathbf{x}_{\text{next}} \circ \mathbf{s}_{\text{next}}$ , we obtain the following linear system – the *Newton system*:

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^\top & I \\ \text{Arw}(\mathbf{s}) & 0 & \text{Arw}(\mathbf{x}) \end{bmatrix} \begin{bmatrix} \Delta \mathbf{x} \\ \Delta \mathbf{y} \\ \Delta \mathbf{s} \end{bmatrix} = \begin{bmatrix} \mathbf{b} - A\mathbf{x} \\ \mathbf{c} - \mathbf{s} - A^\top \mathbf{y} \\ \sigma \mu \mathbf{e} - \mathbf{x} \circ \mathbf{s} \end{bmatrix}. \tag{7}$$

As a final remark, it is not guaranteed that  $(\mathbf{x}_{\text{next}}, \mathbf{y}_{\text{next}}, \mathbf{s}_{\text{next}})$  is on the central path (6), or even that it is still strictly feasible. Luckily, it can be shown that long as  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$  starts out in a neighborhood  $\mathcal{N}$  of the central path,  $(\mathbf{x}_{\text{next}}, \mathbf{y}_{\text{next}}, \mathbf{s}_{\text{next}})$  will remain both strictly feasible and in  $\mathcal{N}$ . It can be shown that this algorithm halves the duality gap every  $O(\sqrt{r})$  iterations, so indeed, after  $O(\sqrt{r} \log(\mu_0/\epsilon))$  it will converge to a (feasible) solution with duality gap at most  $\epsilon$  (given that the initial duality gap was  $\mu_0$ ).



## 2.4 Quantum linear algebra

As it was touched upon in the paper, the main speedup in our algorithms comes from the fact that we use the quantum linear algebra algorithms from [11, 18] whose complexity is sublinear in the dimension. Of course, we need to change our computational model for this sentence to make any sense: namely, we encode  $n$ -dimensional unit vectors as quantum states of a  $\lceil \log_2(n) \rceil$ -qubit system. In other words, for a vector  $\mathbf{z} \in \mathbb{R}^{2^k}$  with  $\|\mathbf{z}\| = 1$ , we use the notation  $|\mathbf{z}\rangle$  to refer to the  $k$ -qubit state  $|\mathbf{z}\rangle := \sum_{i=0}^{2^k-1} z_i |i\rangle$ , where  $|i\rangle$  are the standard basis vectors of  $\mathbb{C}^{2^k}$ , the state space of a  $k$ -qubit system [32].

Given a quantum state  $|\mathbf{z}\rangle$ , we have very limited ways of interacting with it: we can either apply a unitary transformation  $U : \mathbb{C}^{2^k} \rightarrow \mathbb{C}^{2^k}$ , or we can *measure* it, which means that we discard the state and obtain a single random integer  $0 \leq i \leq 2^k - 1$ , with the probability of measuring  $i$  being  $z_i^2$ . In particular this means that we can neither observe the *amplitudes*  $z_i$  directly, nor can we create a copy of  $|\mathbf{z}\rangle$  for an arbitrary  $|\mathbf{z}\rangle$ . In addition to this, it is *a priori* not clear how (and whether it is even possible) to implement the state  $|\mathbf{z}\rangle$  or an arbitrary unitary  $U$  using the gates of a quantum computer. Luckily, there exists a quantum-classical framework using *QRAM* data structures described in [24] that provides a positive answer to both of these questions.

The QRAM can be thought of as the quantum analogue to RAM, i.e. an array  $[\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(m)}]$  of  $w$ -bit bitstrings, whose elements we can access in poly-logarithmic time given their address (position in the array). More precisely, QRAM is just an efficient implementation of the unitary transformation

$$|i\rangle |0\rangle^{\otimes w} \mapsto |i\rangle |b_1^{(i)} \dots b_w^{(i)}\rangle, \text{ for } i \in [m].$$

The usefulness of QRAM data structures becomes clear when we consider the *block encoding* framework:

**Definition.** Let  $A \in \mathbb{R}^{n \times n}$  be a symmetric matrix. Then, the  $\ell$ -qubit unitary matrix  $U \in \mathbb{C}^{2^\ell \times 2^\ell}$  is a  $(\zeta, \ell)$  block encoding of  $A$  if  $U = \begin{bmatrix} A/\zeta & \cdot \\ \cdot & \cdot \end{bmatrix}$ . For an arbitrary matrix  $B \in \mathbb{R}^{n \times m}$ , a block encoding of  $B$  is any block encoding of its symmetrized version  $\text{sym}(B) := \begin{bmatrix} 0 & B \\ B^\top & 0 \end{bmatrix}$ .

We want  $U$  to be implemented efficiently, i.e. using an  $\ell$ -qubit quantum circuit of depth (poly-)logarithmic in  $n$ . Such a circuit would allow us to efficiently create states  $|A_i\rangle$  corresponding to rows (or columns) of  $A$ . Moreover, we need to be able to construct such a data structure efficiently from the classical description of  $A$ . It turns out that we are able to fulfill both of these requirements using a data structure built on top of QRAM.

**Theorem 1** (Block encodings using QRAM [24, 25]). *There exist QRAM data structures for storing vectors  $\mathbf{v}_i \in \mathbb{R}^n$ ,  $i \in [m]$  and matrices  $A \in \mathbb{R}^{n \times n}$  such that with access to these data structures one can do the following:*

1. Given  $i \in [m]$ , prepare the state  $|\mathbf{v}_i\rangle$  in time  $\tilde{O}(1)$ . In other words, the unitary  $|i\rangle |0\rangle \mapsto |i\rangle |\mathbf{v}_i\rangle$  can be implemented efficiently.
2. A  $(\zeta(A), 2 \log n)$  unitary block encoding for  $A$  with  $\zeta(A) = \|A\|_2^{-1} \min(\|A\|_F, s_1(A))$ , where  $s_1(A) = \max_i \sum_j |A_{i,j}|$  can be implemented in time  $O(\log n)$ . Moreover, this block encoding can be constructed in a single pass over the matrix  $A$ , and it can be updated in  $O(\log^2 n)$  time per entry.



From now on, we will also refer to storing vectors and matrices in QRAM, meaning that we use the data structure from Theorem 1. Note that this is the same quantum oracle model that has been used to solve SDPs in [26] and [40].

Once we have these block encodings, we may use them to perform linear algebra. In particular, we want to construct the quantum states  $|A\mathbf{b}\rangle$  and  $|A^{-1}\mathbf{b}\rangle$ , corresponding to the matrix-vector product  $A\mathbf{b}$  and the solution of the linear system  $A\mathbf{x} = \mathbf{b}$ :

**Theorem 2** (Quantum linear algebra with block encodings [11, 18]). *Let  $A \in \mathbb{R}^{n \times n}$  be a matrix with non-zero eigenvalues in the interval  $[-1, -1/\kappa] \cup [1/\kappa, 1]$ , and let  $\epsilon > 0$ . Given an implementation of an  $(\zeta, O(\log n))$  block encoding for  $A$  in time  $T_U$  and a procedure for preparing state  $|b\rangle$  in time  $T_b$ ,*

1. *A state  $\epsilon$ -close to  $|A^{-1}\mathbf{b}\rangle$  can be generated in time  $O((T_U\kappa\zeta + T_b\kappa) \text{polylog}(\kappa\zeta/\epsilon))$ .*
2. *A state  $\epsilon$ -close to  $|A\mathbf{b}\rangle$  can be generated in time  $O((T_U\kappa\zeta + T_b\kappa) \text{polylog}(\kappa\zeta/\epsilon))$ .*
3. *For  $\mathcal{A} \in \{A, A^{-1}\}$ , an estimate  $\Lambda$  such that  $\Lambda \in (1 \pm \epsilon)\|\mathcal{A}\mathbf{b}\|$  can be generated in time  $O((T_U + T_b) \frac{\kappa\zeta}{\epsilon} \text{polylog}(\kappa\zeta/\epsilon))$ .*

Finally, in order to recover classical information from the outputs of a linear system solver, we require an efficient procedure for *quantum state tomography*. The tomography procedure is linear in the dimension of the quantum state.

**Theorem 3** (Efficient vector state tomography, [26]). *There exists an algorithm that given a procedure for constructing  $|\mathbf{x}\rangle$  (i.e. a unitary mapping  $U : |0\rangle \mapsto |\mathbf{x}\rangle$  and its controlled version in time  $T_U$ ) and precision  $\delta > 0$  produces an estimate  $\bar{\mathbf{x}} \in \mathbb{R}^n$  with  $\|\bar{\mathbf{x}}\| = 1$  such that  $\|\mathbf{x} - \bar{\mathbf{x}}\| \leq \sqrt{7}\delta$  with probability at least  $(1 - 1/n^{0.83})$ . The algorithm runs in time  $O\left(T_U \frac{n \log n}{\delta^2}\right)$ .*

Of course, repeating this algorithm  $\tilde{O}(1)$  times allows us to increase the success probability to at least  $1 - 1/\text{poly}(n)$ . Putting Theorems 1, 2 and 3 together, assuming that  $A$  and  $\mathbf{b}$  are already in QRAM, we obtain that the complexity of a completely self-contained algorithm for solving the system  $A\mathbf{x} = \mathbf{b}$  with error  $\delta$  is  $\tilde{O}\left(n \cdot \frac{\kappa\zeta}{\delta^2}\right)$ . For well-conditioned matrices, this presents a significant improvement over  $O(n^\omega)$  (or, in practice,  $O(n^3)$ ) needed for solving linear systems classically, especially when  $n$  is large and the desired precision is not too high. This can be compared with classical iterative linear algebra algorithms [34] that have  $O(\text{nnz}(A))$  complexity per iteration, as well as the new quantum-inspired solvers [19] that have high-degree polynomial dependence on the rank, error, and the condition number.

### 3 A quantum interior-point method

Having introduced the classical IPM for SOCP, we can finally introduce our quantum IPM (Algorithm 1). The main idea is to use quantum linear algebra as much as possible (including solving the Newton system – the most expensive part of each iteration), and falling back to classical computation when needed. Since quantum linear algebra introduces inexactness, we need to deal with it in the analysis. The inspiration for the “quantum part” of the analysis is [26], whereas the “classical part” is based on [28]. Nevertheless, the SOCP analysis is unique in many aspects and a number of hurdles had to be overcome to make the analysis go through. In the rest of the section, we give a brief introduction of

**Require:** Matrix  $A$  and vectors  $\mathbf{b}, \mathbf{c}$  in QRAM, precision parameter  $\epsilon$

1. Find feasible initial point  $(\mathbf{x}, \mathbf{y}, \mathbf{s}, \mu)$  and store it in QRAM.
2. Repeat the following steps for  $O(\sqrt{r} \log(\mu_0/\epsilon))$  iterations:
  - (a) Compute the vector  $\sigma\mu\mathbf{e} - \mathbf{x} \circ \mathbf{s}$  classically and store it in QRAM.
  - (b) Prepare and update the block encodings of the LHS and the RHS of the Newton system (7)
  - (c) Solve the Newton system to obtain  $|(\Delta\mathbf{x}; \Delta\mathbf{y}; \Delta\mathbf{s})\rangle$ , and obtain a classical approximate solution  $(\overline{\Delta\mathbf{x}}; \overline{\Delta\mathbf{y}}; \overline{\Delta\mathbf{s}})$  using tomography.
  - (d) Update  $\mathbf{x} \leftarrow \mathbf{x} + \overline{\Delta\mathbf{x}}$ ,  $\mathbf{s} \leftarrow \mathbf{s} + \overline{\Delta\mathbf{s}}$  and store in QRAM.
  - (e) Update  $\mu \leftarrow \frac{1}{r} \mathbf{x}^\top \mathbf{s}$ .
3. Output  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$ .

**Algorithm 1:** A quantum IPM for SOCP

the most important quantum building blocks we use, as well as present a sketch of the analysis.

First, we note that the algorithms from the previous section allow us to “forget” that Algorithm 1 is quantum, and treat it as a small modification of the classical IPM, where the system (7) is solved up to an  $\ell_2$ -error  $\delta$ . Since Algorithm 1 is iterative, the main part of the analysis is proving that a single iteration preserves closeness to the central path, strict feasibility, and improves the duality gap. In the remainder of this section, we state our main results informally, while the exact statements and proofs of all claims can be found in the supplementary material.

**Theorem 4** (Per-iteration correctness, informal). *Let  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$  be a strictly feasible primal-dual solution that is close to the central path, with duality gap  $\mu$ , and at distance at least  $\delta$  from the boundary of  $\mathcal{L}$ . Then, the Newton system (7) has a unique solution  $(\Delta\mathbf{x}, \Delta\mathbf{y}, \Delta\mathbf{s})$ . There exist positive constants  $\xi, \alpha$  such that the following holds: If we let  $\overline{\Delta\mathbf{x}}, \overline{\Delta\mathbf{s}}$  be approximate solutions of (7) that satisfy*

$$\|\Delta\mathbf{x} - \overline{\Delta\mathbf{x}}\|_F \leq \xi\delta \text{ and } \|\Delta\mathbf{s} - \overline{\Delta\mathbf{s}}\|_F \leq \xi\delta,$$

and let  $\mathbf{x}_{next} := \mathbf{x} + \overline{\Delta\mathbf{x}}$  and  $\mathbf{s}_{next} := \mathbf{s} + \overline{\Delta\mathbf{s}}$  be the updated solution, then:

1. The updated solution is strictly feasible, i.e.  $\mathbf{x}_{next} \in \text{int } \mathcal{L}$  and  $\mathbf{s}_{next} \in \text{int } \mathcal{L}$ .
2. The updated solution is close to the central path, and the new duality gap is less than  $(1 - \alpha/\sqrt{r})\mu$ .

The proof of this theorem consists of 3 main parts:

1. Rescaling  $\mathbf{x}$  and  $\mathbf{s}$  so that they commute in the Jordan-algebraic sense [1]. This part can be reused from the classical analysis [28].
2. Bounding the norms of  $\overline{\Delta\mathbf{x}}$  and  $\overline{\Delta\mathbf{s}}$ , and proving that  $\mathbf{x} + \overline{\Delta\mathbf{x}}$  and  $\mathbf{s} + \overline{\Delta\mathbf{s}}$  are still strictly feasible (in the sense of belonging to  $\text{int } \mathcal{L}$ ). This part of the analysis is also inspired by the classical analysis, but it has to take into account the inexactness of the Newton system solution.

3. Proving that the new solution  $(\mathbf{x} + \overline{\Delta\mathbf{x}}, \mathbf{y} + \overline{\Delta\mathbf{y}}, \mathbf{s} + \overline{\Delta\mathbf{s}})$  is in the neighborhood of the central path, and the duality gap/central path parameter have decreased by a factor of  $1 - \alpha/\sqrt{r}$ , where  $\alpha$  is constant. This part is the most technical, and while it is inspired by [26], it required using many of the Jordan-algebraic tools from [1, 28].

Theorem 4 formalizes the fact that the Algorithm 1 has the same iteration invariant as the classical IPM. Since the duality gap is reduced by the same factor in both algorithms, their iteration complexity is the same, and a simple calculation shows that they need  $O(\sqrt{r})$  iterations to halve the duality gap. On the other hand, the cost of each iteration varies, since the complexity of the quantum linear system solver depends on the precision  $\xi\delta$ , the condition number  $\kappa$  of the Newton matrix, as well as its  $\zeta$ -parameter. While exactly bounding these quantities is the subject of future research, it is worth noting that for  $\zeta$ , we have the trivial bound  $\zeta \leq \sqrt{n}$ , and research on iterative linear algebra methods [14] suggests that  $\kappa = O(1/\mu) = O(1/\epsilon)$ . The final complexity is summarized in the following theorem:

**Theorem 5.** *Let (2) be a SOCP with  $A \in \mathbb{R}^{m \times n}$ ,  $m \leq n$ , and  $\mathcal{L} = \mathcal{L}^{n_1} \times \dots \times \mathcal{L}^{n_r}$ . Then, Algorithm 1 achieves duality gap  $\epsilon$  in time*

$$T = \tilde{O} \left( \sqrt{r} \log(\mu_0/\epsilon) \cdot \frac{n\kappa\zeta}{\delta^2} \log \left( \frac{\kappa\zeta}{\delta} \right) \right),$$

where the  $\tilde{O}(\cdot)$  notation hides the factors that are poly-logarithmic in  $n$  and  $m$ .

Finally, the quality of the resulting (classical) solution is characterized by the following theorem:

**Theorem 6.** *Let (2) be a SOCP as in Theorem 5. Then, after  $T$  iterations, the (linear) infeasibility of the final iterate  $\mathbf{x}, \mathbf{y}, \mathbf{s}$  is bounded as*

$$\begin{aligned} \|A\mathbf{x} - \mathbf{b}\| &\leq \delta\|A\|, \\ \|A^\top \mathbf{y} + \mathbf{s} - \mathbf{c}\| &\leq \delta(\|A\| + 1). \end{aligned}$$

## 4 Technical results

In this section, we present our main technical results – the proofs of Theorems 4, 5 and 6.

### 4.1 Central path

In addition to the central path defined in (6), we define the distance from the central path as  $d(\mathbf{x}, \mathbf{s}, \nu) = \|T_{\mathbf{x}}\mathbf{s} - \nu\mathbf{e}\|_F$ , so the corresponding  $\eta$ -neighborhood is given by

$$\mathcal{N}_\eta(\nu) := \{(\mathbf{x}, \mathbf{y}, \mathbf{s}) \mid (\mathbf{x}, \mathbf{y}, \mathbf{s}) \text{ strictly feasible and } d(\mathbf{x}, \mathbf{s}, \nu) \leq \eta\nu\}.$$

Using this neighborhood definition, we can specify what exactly do we mean when we claim that the important properties of the central path are valid in its neighborhood as well.

**Lemma 1** (Properties of the central path). *Let  $\nu > 0$  be arbitrary and let  $\mathbf{x}, \mathbf{s} \in \text{int } \mathcal{L}$ . Then,  $\mathbf{x}$  and  $\mathbf{s}$  satisfy the following properties:*

1. For all  $\nu > 0$ , the duality gap and distance from the central path are related as

$$|\mathbf{x}^\top \mathbf{s} - r\nu| \leq \sqrt{\frac{r}{2}} \cdot d(\mathbf{x}, \mathbf{s}, \nu).$$

2. The distance from the central path is symmetric in its arguments i.e.  $d(\mathbf{x}, \mathbf{s}, \nu) = d(\mathbf{s}, \mathbf{x}, \nu)$ .

3. Let  $\mu = \frac{1}{r} \mathbf{x}^\top \mathbf{s}$ . If  $d(\mathbf{x}, \mathbf{s}, \mu) \leq \eta \mu$ , then  $(1 + \eta) \|\mathbf{s}^{-1}\|_2 \geq \|\mu^{-1} \mathbf{x}\|_2$ .

*Proof.* For part 1, let  $\{\lambda_i\}_{i=1}^{2r}$  be the eigenvalues of  $T_{\mathbf{x}} \mathbf{s}$ , note that  $T_{\mathbf{x}}$  is invertible as  $x \in \mathcal{L}$ . Then using the properties of  $T_{\mathbf{x}}$ , we have

$$\mathbf{x}^\top \mathbf{s} = \mathbf{x}^\top T_{\mathbf{x}}^{-1} T_{\mathbf{x}} \mathbf{s} = (T_{\mathbf{x}^{-1}} \mathbf{x})^\top T_{\mathbf{x}} \mathbf{s} = \mathbf{e}^\top T_{\mathbf{x}} \mathbf{s} = \frac{1}{2} \sum_{i=1}^{2r} \lambda_i.$$

We can therefore bound the duality gap  $\mathbf{x}^\top \mathbf{s}$  as follows,

$$\mathbf{x}^\top \mathbf{s} = \frac{1}{2} \sum_{i=1}^{2r} \lambda_i \leq r\nu + \frac{1}{2} \sum_{i=1}^{2r} |\lambda_i - \nu| \leq r\nu + \sqrt{\frac{r}{2}} \sqrt{\sum_{i=1}^{2r} (\lambda_i - \nu)^2} = r\nu + \sqrt{\frac{r}{2}} \cdot d(\mathbf{x}, \mathbf{s}, \nu).$$

The second step used the Cauchy-Schwarz inequality while the third follows from the definition  $d(\mathbf{x}, \mathbf{s}, \nu)^2 = \sum_{i=1}^{2r} (\lambda_i - \nu)^2$ . The proof of the lower bound is similar, but starts instead with the inequality

$$\frac{1}{2} \sum_{i=1}^{2r} \lambda_i \geq r\nu - \frac{1}{2} \sum_{i=1}^{2r} |\nu - \lambda_i|.$$

For part 2, it suffices to prove that  $T_{\mathbf{x}} \mathbf{s}$  and  $T_{\mathbf{s}} \mathbf{x}$  have the same eigenvalues. This follows from part 2 of Theorem 10 in [1]. Finally for part 3, as  $d(\mathbf{s}, \mathbf{x}, \mu) \leq \eta \mu$  we have,

$$\begin{aligned} \eta \mu &\geq \|T_{\mathbf{s}} \mathbf{x} - \mu \mathbf{e}\|_F \\ &= \|T_{\mathbf{s}} \mathbf{x} - \mu (T_{\mathbf{s}} T_{\mathbf{s}^{-1}}) \mathbf{e}\|_F \\ &= \|T_{\mathbf{s}} (\mathbf{x} - \mu T_{\mathbf{s}^{-1}} \mathbf{e})\|_F \\ &\geq \lambda_{\min}(T_{\mathbf{s}}) \|\mathbf{x} - \mu T_{\mathbf{s}^{-1}} \mathbf{e}\|_F \\ &\geq \lambda_{\min}(T_{\mathbf{s}}) \|\mathbf{x} - \mu \mathbf{s}^{-1}\|_2 \\ &= \frac{1}{\|\mathbf{s}^{-1}\|_2} \cdot \mu \cdot \|\mu^{-1} \mathbf{x} - \mathbf{s}^{-1}\|_2 \end{aligned}$$

Therefore,  $\eta \|\mathbf{s}^{-1}\|_2 \geq \|\mu^{-1} \mathbf{x} - \mathbf{s}^{-1}\|_2$ . Finally, by the triangle inequality for the spectral norm,

$$\eta \|\mathbf{s}^{-1}\|_2 \geq \|\mu^{-1} \mathbf{x}\|_2 - \|\mathbf{s}^{-1}\|_2,$$

so we can conclude that  $\|\mathbf{s}^{-1}\|_2 \geq \frac{1}{1+\eta} \|\mu^{-1} \mathbf{x}\|_2$ .  $\square$

## 4.2 A single quantum IPM iteration

Recall that the essence of our quantum algorithm is repeated solution of the Newton system (7) using quantum linear algebra. As such, our goal is to prove the following theorem:

**Theorem 4** (Per-iteration correctness, formal). *Let  $\chi = \eta = 0.01$  and  $\xi = 0.001$  be positive constants and let  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$  be a feasible solution of (2) and (3) with  $\mu = \frac{1}{r} \mathbf{x}^\top \mathbf{s}$  and  $d(\mathbf{x}, \mathbf{s}, \mu) \leq \eta \mu$ . Then, for  $\sigma = 1 - \chi/\sqrt{n}$ , the Newton system (7) has a unique solution  $(\Delta \mathbf{x}, \Delta \mathbf{y}, \Delta \mathbf{s})$ . Let  $\overline{\Delta \mathbf{x}}, \overline{\Delta \mathbf{s}}$  be approximate solutions of (7) that satisfy*

$$\|\Delta \mathbf{x} - \overline{\Delta \mathbf{x}}\|_F \leq \frac{\xi}{\|T_{\mathbf{x}^{-1}}\|}, \quad \|\Delta \mathbf{s} - \overline{\Delta \mathbf{s}}\|_F \leq \frac{\xi}{2\|T_{\mathbf{s}^{-1}}\|},$$

where  $T_{\mathbf{x}}$  and  $T_{\mathbf{x}}$  are the square roots of the quadratic representation matrices in equation (5). If we let  $\mathbf{x}_{\text{next}} := \mathbf{x} + \overline{\Delta \mathbf{x}}$  and  $\mathbf{s}_{\text{next}} := \mathbf{s} + \overline{\Delta \mathbf{s}}$ , the following holds:

1. The updated solution is strictly feasible, i.e.  $\mathbf{x}_{next} \in \text{int } \mathcal{L}$  and  $\mathbf{s}_{next} \in \text{int } \mathcal{L}$ .
2. The updated solution satisfies  $d(\mathbf{x}_{next}, \mathbf{s}_{next}, \bar{\mu}) \leq \eta \bar{\mu}$  and  $\frac{1}{r} \mathbf{x}_{next}^\top \mathbf{s}_{next} = \bar{\mu}$  for  $\bar{\mu} = \bar{\sigma} \mu$ ,  $\bar{\sigma} = 1 - \frac{\alpha}{\sqrt{r}}$  and a constant  $0 < \alpha \leq \chi$ .

Since the Newton system (7) is the same as in the classical case, we can reuse Theorem 1 from [28] for the uniqueness part of Theorem 4. Therefore, we just need to prove the two parts about strict feasibility and improving the duality gap. Our analysis is inspired by the general case analysis from [5], the derived SDP analysis from [26], and uses some technical results from the SOCP analysis in [28]. The proof of Theorem 4 consists of three main steps:

1. Rescaling  $\mathbf{x}$  and  $\mathbf{s}$  so that they share the same Jordan frame.
2. Bounding the norms of  $\Delta \mathbf{x}$  and  $\Delta \mathbf{s}$ , and proving that  $\mathbf{x} + \Delta \mathbf{x}$  and  $\mathbf{s} + \Delta \mathbf{s}$  are still strictly feasible (in the sense of belonging to  $\text{int } \mathcal{L}$ ).
3. Proving that the new solution  $(\mathbf{x} + \Delta \mathbf{x}, \mathbf{y} + \Delta \mathbf{y}, \mathbf{s} + \Delta \mathbf{s})$  is in the  $\eta$ -neighborhood of the central path, and the duality gap/central path parameter have decreased by a factor of  $1 - \alpha/\sqrt{n}$ , where  $\alpha$  is constant.

### 4.3 Rescaling $\mathbf{x}$ and $\mathbf{s}$

As in the case of SDPs, the first step of the proof uses the symmetries of the Lorentz cone to perform a commutative scaling, that is to reduce the analysis to the case when  $\mathbf{x}$  and  $\mathbf{s}$  share the same Jordan frame. Although  $\circ$  is commutative by definition, two vectors sharing a Jordan frame are akin to two matrices sharing a system of eigenvectors, and thus commuting (some authors [1] say that the vectors *operator commute* in this case). The easiest way to achieve this is to scale by  $T_{\mathbf{x}} = Q_{\mathbf{x}^{1/2}}$  and  $\mu^{-1}$ , i.e. to change our variables as

$$\mathbf{x} \mapsto \mathbf{x}' := T_{\mathbf{x}}^{-1} \mathbf{x} = \mathbf{e} \text{ and } \mathbf{s} \mapsto \mathbf{s}' := \mu^{-1} T_{\mathbf{x}} \mathbf{s}.$$

Note that for convenience, we have also rescaled the duality gap to 1. Recall also that in the matrix case, the equivalent of this scaling was  $X \mapsto X^{-1/2} X X^{-1/2} = I$  and  $S \mapsto \mu^{-1} X^{1/2} S X^{1/2}$ . We use the notation  $\mathbf{z}'$  to denote the appropriately-scaled vector  $\mathbf{z}$ , so that we have

$$\Delta \mathbf{x}' := T_{\mathbf{x}}^{-1} \Delta \mathbf{x}, \quad \Delta \mathbf{s}' := \mu^{-1} T_{\mathbf{x}} \Delta \mathbf{s}$$

For approximate quantities (e.g. the ones obtained using tomography, or any other approximate linear system solver), we use the notation  $\overline{\cdot}$ , so that the increments become  $\overline{\Delta \mathbf{x}}$  and  $\overline{\Delta \mathbf{s}}$ , and their scaled counterparts are  $\overline{\Delta \mathbf{x}'} := T_{\mathbf{x}}^{-1} \overline{\Delta \mathbf{x}}$  and  $\overline{\Delta \mathbf{s}'} := \mu^{-1} T_{\mathbf{x}} \overline{\Delta \mathbf{s}}$ . Finally, we denote the scaled version of the next iterate as  $\mathbf{x}'_{next} := \mathbf{e} + \overline{\Delta \mathbf{x}'}$  and  $\mathbf{s}'_{next} := \mathbf{e} + \overline{\Delta \mathbf{s}'}$ . Now, we see that the statement of Theorem 4 implies the following bounds on  $\|\Delta \mathbf{x}' - \overline{\Delta \mathbf{x}'}\|_F$  and  $\|\Delta \mathbf{s}' - \overline{\Delta \mathbf{s}'}\|_F$ :

$$\begin{aligned} \|\Delta \mathbf{x}' - \overline{\Delta \mathbf{x}'}\|_F &= \|T_{\mathbf{x}^{-1}} \Delta \mathbf{x} - T_{\mathbf{x}^{-1}} \overline{\Delta \mathbf{x}}\|_F \\ &\leq \|T_{\mathbf{x}^{-1}}\| \cdot \|\Delta \mathbf{x} - \overline{\Delta \mathbf{x}}\|_F \leq \xi, \text{ and} \\ \|\Delta \mathbf{s}' - \overline{\Delta \mathbf{s}'}\|_F &= \mu^{-1} \|T_{\mathbf{x}} \Delta \mathbf{s} - T_{\mathbf{x}} \overline{\Delta \mathbf{s}}\|_F \\ &\leq \mu^{-1} \|T_{\mathbf{x}}\| \|\Delta \mathbf{s} - \overline{\Delta \mathbf{s}}\|_F \\ &= \mu^{-1} \|\mathbf{x}\|_2 \|\Delta \mathbf{s} - \overline{\Delta \mathbf{s}}\|_F \\ &\leq (1 + \eta) \|\mathbf{s}^{-1}\|_2 \|\Delta \mathbf{s} - \overline{\Delta \mathbf{s}}\|_F \text{ by Lemma 1} \\ &\leq 2 \|T_{\mathbf{s}^{-1}}\| \|\Delta \mathbf{s} - \overline{\Delta \mathbf{s}}\|_F \leq \xi. \end{aligned}$$

Throughout the analysis, we will make use of several constants:  $\eta > 0$  is the distance from the central path, i.e. we ensure that our iterates stay in the  $\eta$ -neighborhood  $\mathcal{N}_\eta$  of the central path. The constant  $\sigma = 1 - \chi/\sqrt{r}$  is the factor by which we aim to decrease our duality gap, for some constant  $\chi > 0$ . Finally constant  $\xi > 0$  is the approximation error for the scaled increments  $\overline{\Delta \mathbf{x}'}, \overline{\Delta \mathbf{s}'}$ . Having this notation in mind, we can state several facts about the relation between the duality gap and the central path distance for the original and scaled vectors.

**Claim 2.** *The following holds for the scaled vectors  $\mathbf{x}'$  and  $\mathbf{s}'$ :*

1. *The scaled duality gap is  $\frac{1}{r} \mathbf{x}'^\top \mathbf{s}' = 1$ .*
2.  *$d(\mathbf{x}, \mathbf{s}, \mu) \leq \eta \mu$  is equivalent to  $\|\mathbf{s}' - \mathbf{e}\| \leq \eta$ .*
3.  *$d(\mathbf{x}, \mathbf{s}, \mu \sigma) = \mu \cdot d(\mathbf{x}', \mathbf{s}', \sigma)$ , for all  $\sigma > 0$ .*

At this point, we claim that it suffices to prove the two parts of Theorem 4 in the scaled case. Namely, assuming that  $\mathbf{x}'_{\text{next}} \in \text{int } \mathcal{L}$  and  $\mathbf{s}'_{\text{next}} \in \text{int } \mathcal{L}$ , by construction we get

$$\mathbf{x}_{\text{next}} = T_{\mathbf{x}} \mathbf{x}'_{\text{next}} \text{ and } \mathbf{s}_{\text{next}} = T_{\mathbf{x}^{-1}} \mathbf{s}'_{\text{next}}$$

and thus  $\mathbf{x}_{\text{next}}, \mathbf{s}_{\text{next}} \in \text{int } \mathcal{L}$ .

On the other hand, if  $\mu d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \bar{\sigma}) \leq \eta \bar{\mu}$ , then  $d(\mathbf{x}_{\text{next}}, \mathbf{s}_{\text{next}}, \bar{\mu}) \leq \eta \bar{\mu}$  follows by Claim 2. Similarly, from  $\frac{1}{r} \mathbf{x}'_{\text{next}} \mathbf{s}'_{\text{next}} = \bar{\sigma}$ , we also get  $\frac{1}{r} \mathbf{x}_{\text{next}} \mathbf{s}_{\text{next}} = \bar{\mu}$ . We conclude this part with two technical results from [28], that use the auxiliary matrix  $R_{xs}$  defined as  $R_{xs} := T_{\mathbf{x}} \text{Arw}(\mathbf{x})^{-1} \text{Arw}(\mathbf{s}) T_{\mathbf{x}}$ . These results are useful for the later parts of the proof of Theorem 4.

**Claim 3** ([28], Lemma 3). *Let  $\eta$  be the distance from the central path, and let  $\nu > 0$  be arbitrary. Then,  $R_{xs}$  is bounded as*

$$\|R_{xs} - \nu I\| \leq 3\eta\nu.$$

**Claim 4** ([28], Lemma 5, proof). *Let  $\mu$  be the duality gap. Then, the scaled increment  $\Delta \mathbf{s}'$  is*

$$\Delta \mathbf{s}' = \sigma \mathbf{e} - \mathbf{s}' - \mu^{-1} R_{xs} \Delta \mathbf{x}'.$$

#### 4.4 Maintaining strict feasibility

The main tool for showing that strict feasibility is conserved is the following bound on the increments  $\Delta \mathbf{x}'$  and  $\Delta \mathbf{s}'$ :

**Lemma 2** ([28], Lemma 6). *Let  $\eta$  be the distance from the central path and let  $\mu$  be the duality gap. Then, we have the following bounds for the scaled direction:*

$$\begin{aligned} \|\Delta \mathbf{x}'\|_F &\leq \frac{\Theta}{\sqrt{2}} \\ \|\Delta \mathbf{s}'\|_F &\leq \Theta \sqrt{2}, \quad \text{where } \Theta = \frac{2\sqrt{\eta^2/2 + (1-\sigma)^2 r}}{1-3\eta} \end{aligned}$$

Moreover, if we substitute  $\sigma$  with its actual value  $1 - \chi/\sqrt{r}$ , we get  $\Theta = \frac{\sqrt{2\eta^2 + 4\chi^2}}{1-3\eta}$ , which we can make arbitrarily small by tuning the constants. Now, we can immediately use this result to prove  $\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}} \in \text{int } \mathcal{L}$ .

**Lemma 3.** *Let  $\eta = \chi = 0.01$  and  $\xi = 0.001$ . Then,  $\mathbf{x}'_{\text{next}}$  and  $\mathbf{s}'_{\text{next}}$  are strictly feasible, i.e.  $\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}} \in \text{int } \mathcal{L}$ .*



*Proof.* By Lemma 2,  $\lambda_{\min}(\bar{\mathbf{x}}) \geq 1 - \|\overline{\Delta \mathbf{x}'}\|_F \geq 1 - \frac{\Theta}{\sqrt{2}} - \xi$ . On the other hand, since  $d(\mathbf{x}, \mathbf{s}, \mu) \leq \eta\mu$ , we have  $d(\mathbf{x}', \cdot, 1) \leq \eta$ , and thus

$$\eta^2 \geq \|\cdot, -e\|_F^2 = \sum_{i=1}^{2r} (\lambda_i(\cdot) - 1)^2$$

The above equation implies that  $\lambda_i(\cdot) \in [1 - \eta, 1 + \eta], \forall i \in [2r]$ . Now, since  $\|\mathbf{z}\|_2 \leq \|\mathbf{z}\|_F$ ,

$$\begin{aligned} \lambda_{\min}(\bar{\mathbf{s}}) &\geq \lambda_{\min}(\cdot, + \Delta \mathbf{s}') - \|\overline{\Delta \mathbf{s}'} - \Delta \mathbf{s}'\|_F \\ &\geq \lambda_{\min}(\cdot) - \|\Delta \mathbf{s}'\|_F - \|\overline{\Delta \mathbf{s}'} - \Delta \mathbf{s}'\|_F \\ &\geq 1 - \eta - \Theta\sqrt{2} - \xi, \end{aligned}$$

where we used Lemma 2 for the last inequality. Substituting  $\eta = \chi = 0.01$  and  $\xi = 0.001$ , we get that  $\lambda_{\min}(\bar{\mathbf{x}}) \geq 0.8$  and  $\lambda_{\min}(\bar{\mathbf{s}}) \geq 0.8$ .  $\square$

#### 4.5 Maintaining closeness to central path

Finally, we move on to the most technical part of the proof of Theorem 4, where we prove that  $\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}$  is still close to the central path, and the duality gap has decreased by a constant factor. We split this into two lemmas.

**Lemma 4.** *Let  $\eta = \chi = 0.01$ ,  $\xi = 0.001$ , and let  $\alpha$  be any value satisfying  $0 < \alpha \leq \chi$ . Then, for  $\bar{\sigma} = 1 - \alpha/\sqrt{r}$ , the distance to the central path is maintained, that is,  $d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \bar{\sigma}) < \eta\bar{\sigma}$ .*

*Proof.* By Claim 2, the distance of the next iterate from the central path is

$$d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \bar{\sigma}) = \|T_{\mathbf{x}'_{\text{next}}} \mathbf{s}'_{\text{next}} - \bar{\sigma} \mathbf{e}\|_F,$$

and we can bound it from above as

$$\begin{aligned} d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \bar{\sigma}) &= \|T_{\mathbf{x}'_{\text{next}}} \mathbf{s}'_{\text{next}} - \bar{\sigma} \mathbf{e}\|_F \\ &= \|T_{\mathbf{x}'_{\text{next}}} \mathbf{s}'_{\text{next}} - \bar{\sigma} T_{\mathbf{x}'_{\text{next}}} T_{\mathbf{x}'_{\text{next}}}^{-1} \mathbf{e}\|_F \\ &\leq \|T_{\mathbf{x}'_{\text{next}}}\| \cdot \|\mathbf{s}'_{\text{next}} - \bar{\sigma} \cdot (\mathbf{x}'_{\text{next}})^{-1}\|. \end{aligned}$$

So, it is enough to bound  $\|\mathbf{z}\|_F := \|\mathbf{s}'_{\text{next}} - \bar{\sigma} \cdot (\mathbf{x}'_{\text{next}})^{-1}\|_F$  from above, since

$$\|T_{\mathbf{x}'_{\text{next}}}\| = \|\mathbf{x}'_{\text{next}}\|_2 \leq 1 + \|\overline{\Delta \mathbf{x}'}\|_2 \leq 1 + \|\Delta \mathbf{x}'\|_2 + \xi \leq 1 + \frac{\Theta}{\sqrt{2}} + \xi.$$

We split  $\mathbf{z}$  as

$$\mathbf{z} = \underbrace{\left(\cdot, + \overline{\Delta \mathbf{s}'} - \bar{\sigma} \mathbf{e} + \overline{\Delta \mathbf{x}'}\right)}_{\mathbf{z}_1} + \underbrace{(\bar{\sigma} - 1)\overline{\Delta \mathbf{x}'}}_{\mathbf{z}_2} + \underbrace{\bar{\sigma} \left(e - \overline{\Delta \mathbf{x}'} - (e + \overline{\Delta \mathbf{x}'})^{-1}\right)}_{\mathbf{z}_3},$$

and we bound  $\|\mathbf{z}_1\|_F$ ,  $\|\mathbf{z}_2\|_F$ , and  $\|\mathbf{z}_3\|_F$  separately.

1. By the triangle inequality,  $\|\mathbf{z}_1\|_F \leq \|\cdot, + \Delta \mathbf{s}' - \bar{\sigma} \mathbf{e} + \Delta \mathbf{x}'\|_F + 2\xi$ . Furthermore, after substituting  $\Delta \mathbf{s}'$  from Claim 4, we get

$$\begin{aligned} \cdot, + \Delta \mathbf{s}' - \bar{\sigma} \mathbf{e} + \Delta \mathbf{x}' &= \sigma \mathbf{e} - \mu^{-1} R_{xs} \Delta \mathbf{x}' - \bar{\sigma} \mathbf{e} + \Delta \mathbf{x}' \\ &= \frac{\alpha - \chi}{\sqrt{r}} \mathbf{e} + \mu^{-1} (\mu I - R_{xs}) \Delta \mathbf{x}'. \end{aligned}$$

Using the bound for  $\|\mu I - R_{xs}\|$  from Claim 3 as well as the bound for  $\|\Delta \mathbf{x}'\|_F$  from Lemma 2, we obtain

$$\|\mathbf{z}_1\|_F \leq 2\xi + \frac{\chi}{\sqrt{r}} + \frac{3}{\sqrt{2}}\eta\Theta.$$

2.  $\|\mathbf{z}_2\|_F \leq \frac{\chi}{\sqrt{r}} \left( \frac{\Theta}{\sqrt{2}} + \xi \right)$ , where we used the bound from Lemma 2 again.
3. Here, we first need to bound  $\|(e + \Delta \mathbf{x}')^{-1} - (e + \overline{\Delta \mathbf{x}'})^{-1}\|_F$ . For this, we use the submultiplicativity of  $\|\cdot\|_F$ :

$$\begin{aligned} \|(e + \Delta \mathbf{x}')^{-1} - (e + \overline{\Delta \mathbf{x}'})^{-1}\|_F &= \|(e + \Delta \mathbf{x}')^{-1} \circ (e - (e + \Delta \mathbf{x}') \circ (e + \overline{\Delta \mathbf{x}'})^{-1})\|_F \\ &\leq \|(e + \Delta \mathbf{x}')^{-1}\|_2 \cdot \|e - (e + \overline{\Delta \mathbf{x}'})^{-1} \circ (e + \Delta \mathbf{x}')\|_F \\ &= \|(e + \Delta \mathbf{x}')^{-1}\|_2 \cdot \|(\Delta \mathbf{x}' - \overline{\Delta \mathbf{x}'}) \circ (e + \overline{\Delta \mathbf{x}'})^{-1}\|_F \\ &\leq \|(e + \Delta \mathbf{x}')^{-1}\|_2 \cdot \|\Delta \mathbf{x}' - \overline{\Delta \mathbf{x}'}\|_F \cdot \|(e + \overline{\Delta \mathbf{x}'})^{-1}\|_2 \\ &\leq \xi \cdot \|(e + \Delta \mathbf{x}')^{-1}\|_2 \cdot \|(e + \overline{\Delta \mathbf{x}'})^{-1}\|_2. \end{aligned}$$

Now, we have the bound  $\|(e + \Delta \mathbf{x}')^{-1}\|_2 \leq \frac{1}{1 - \|\Delta \mathbf{x}'\|_F}$  and similarly  $\|(e + \overline{\Delta \mathbf{x}'})^{-1}\|_2 \leq \frac{1}{1 - \|\Delta \mathbf{x}'\|_F - \xi}$ , so we get

$$\|(e + \Delta \mathbf{x}')^{-1} - (e + \overline{\Delta \mathbf{x}'})^{-1}\|_F \leq \frac{\xi}{(1 - \|\Delta \mathbf{x}'\|_F - \xi)^2}.$$

Using this, we can bound  $\|\mathbf{z}_3\|_F$ :

$$\|\mathbf{z}_3\|_F \leq \bar{\sigma} \left( \|e - \Delta \mathbf{x}' - (e + \Delta \mathbf{x}')^{-1}\|_F + \xi + \frac{\xi}{(1 - \|\Delta \mathbf{x}'\|_F - \xi)^2} \right).$$

If we let  $\lambda_i$  be the eigenvalues of  $\Delta \mathbf{x}'$ , then by Lemma 2, we have

$$\begin{aligned} \|e - \Delta \mathbf{x}' - (e + \Delta \mathbf{x}')^{-1}\|_F &= \sqrt{\sum_{i=1}^{2r} \left( (1 - \lambda_i) - \frac{1}{1 + \lambda_i} \right)^2} \\ &= \sqrt{\sum_{i=1}^{2r} \frac{\lambda_i^4}{(1 + \lambda_i)^2}} \leq \frac{\Theta}{\sqrt{2} - \Theta} \sqrt{\sum_{i=1}^{2r} \lambda_i^2} \\ &\leq \frac{\Theta^2}{2 - \sqrt{2}\Theta}. \end{aligned}$$

Combining all bound from above, we obtain

$$\begin{aligned} d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \bar{\sigma}) &\leq \left( 1 + \frac{\Theta}{\sqrt{2}} + \xi \right) \cdot \left( 2\xi + \frac{\chi}{\sqrt{r}} + \frac{3}{\sqrt{2}}\eta\Theta \right. \\ &\quad \left. + \frac{\chi}{\sqrt{r}} \left( \frac{\Theta}{\sqrt{2}} + \xi \right) \right. \\ &\quad \left. + \bar{\sigma} \left( \frac{\Theta^2}{2 - \sqrt{2}\Theta} + \xi + \frac{\xi}{(1 - \Theta/\sqrt{2} - \xi)^2} \right) \right). \end{aligned}$$

Finally, if we plug in  $\chi = 0.01$ ,  $\eta = 0.01$ ,  $\xi = 0.001$ , we get  $d(\bar{\mathbf{x}}, \bar{\mathbf{s}}, \bar{\sigma}) \leq 0.005\bar{\sigma} \leq \eta\bar{\sigma}$ .  $\square$

Now, we prove that the duality gap decreases.

**Lemma 5.** *For the same constants, the updated solution satisfies  $\frac{1}{r}\mathbf{x}'_{\text{next}}\mathbf{s}'_{\text{next}} = \left(1 - \frac{\alpha}{\sqrt{r}}\right)$  for  $\alpha = 0.005$ .*

*Proof.* Since  $\mathbf{x}'_{\text{next}}$  and  $\mathbf{s}'_{\text{next}}$  are scaled quantities, the duality gap between their unscaled counterparts is  $\frac{\mu}{r}\mathbf{x}'_{\text{next}}\mathbf{s}'_{\text{next}}$ . Applying Lemma 1 (and Claim 2) with  $\nu = \sigma\mu$  and  $\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}$ , we obtain the upper bound

$$\mu\mathbf{x}'_{\text{next}}\mathbf{s}'_{\text{next}} \leq r\sigma\mu + \sqrt{\frac{r}{2}}\mu d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \sigma),$$

which in turn implies

$$\frac{1}{r}\mathbf{x}'_{\text{next}}\mathbf{s}'_{\text{next}} \leq \left(1 - \frac{0.01}{\sqrt{r}}\right) \left(1 + \frac{d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \sigma)}{\sigma\sqrt{2r}}\right).$$

By instantiating Lemma 1 for  $\alpha = \chi$ , from its proof, we obtain  $d(\mathbf{x}'_{\text{next}}, \mathbf{s}'_{\text{next}}, \sigma) \leq 0.005\sigma$ , and thus

$$\frac{1}{r}\mathbf{x}'_{\text{next}}\mathbf{s}'_{\text{next}} \leq 1 - \frac{0.005}{\sqrt{r}}$$

Therefore, the final  $\alpha$  for this Lemma is 0.005.  $\square$

#### 4.6 Final complexity and feasibility

In every iteration we need to solve the Newton system to a precision dependent on the norms of  $T_{\mathbf{x}-1}$  and  $T_{\mathbf{s}-1}$ . Thus, to bound the running time of the algorithm (since the complexity of the vector state tomography procedure depends on the desired precision), we need to bound  $\|T_{\mathbf{x}-1}\|$  and  $\|T_{\mathbf{s}-1}\|$ . Indeed, by the properties of the quadratic representation, we get

$$\begin{aligned} \|T_{\mathbf{x}-1}\| &= \|\mathbf{x}^{-1}\| = \lambda_{\min}(\mathbf{x})^{-1} \text{ and} \\ \|T_{\mathbf{s}-1}\| &= \|\mathbf{s}^{-1}\| = \lambda_{\min}(\mathbf{s})^{-1}. \end{aligned}$$

If the tomography precision for iteration  $i$  is chosen to be at least (i.e. smaller than)

$$\delta_i := \frac{\xi}{4} \min \{\lambda_{\min}(\mathbf{x}_i), \lambda_{\min}(\mathbf{s}_i)\}, \quad (8)$$

then the premises of Theorem 4 are satisfied. The tomography precision for the entire algorithm can therefore be chosen to be  $\delta := \min_i \delta_i$ . Note that these minimum eigenvalues are related to how close the current iterate is to the boundary of  $\mathcal{L}$  – as long as  $\mathbf{x}_i, \mathbf{s}_i$  are not “too close” to the boundary of  $\mathcal{L}$ , their minimal eigenvalues should not be “too small”.

There are two more parameters that impact the complexity of the quantum linear system solver: the condition number of the Newton matrix  $\kappa_i$  and the matrix parameter  $\zeta_i$  of the QRAM encoding in iteration  $i$ . For both of these quantities we define their global versions as  $\kappa = \max_i \kappa_i$  and  $\zeta = \max_i \zeta_i$ . Therefore, we arrive to the following statement about the complexity of Algorithm 1.

**Theorem 5** (restatement). *Let (2) be a SOCP with  $A \in \mathbb{R}^{m \times n}$ ,  $m \leq n$ , and  $\mathcal{L} = \mathcal{L}^{n_1} \times \dots \times \mathcal{L}^{n_r}$ . Then, our algorithm achieves duality gap  $\epsilon$  in time*

$$T = \tilde{O} \left( \sqrt{r} \log(\mu_0/\epsilon) \cdot \frac{n\kappa\zeta}{\delta^2} \log \left( \frac{\kappa\zeta}{\delta} \right) \right).$$

This complexity can be easily interpreted as product of the number of iterations and the cost of  $n$ -dimensional vector tomography with error  $\delta$ . So, improving the complexity of the tomography algorithm would improve the running time of our algorithm as well.

Note that up to now, we cared mostly about strict (conic) feasibility of  $\mathbf{x}$  and  $\mathbf{s}$ . Now, we address the fact that the linear constraints  $A\mathbf{x} = \mathbf{b}$  and  $A^\top \mathbf{y} + \mathbf{s} = \mathbf{c}$  are not exactly satisfied during the execution of the algorithm. Luckily, it turns out that this error is not accumulated, but is instead determined just by the final tomography precision:

**Theorem 6** (restatement). *Let (2) be a SOCP as in Theorem 5. Then, after  $T$  iterations, the (linear) infeasibility of the final iterate  $\mathbf{x}, \mathbf{y}, \mathbf{s}$  is bounded as*

$$\begin{aligned} \|A\mathbf{x}_T - \mathbf{b}\| &\leq \delta \|A\|, \\ \|A^\top \mathbf{y}_T + \mathbf{s}_T - \mathbf{c}\| &\leq \delta (\|A\| + 1). \end{aligned}$$

*Proof.* Let  $(\mathbf{x}_T, \mathbf{y}_T, \mathbf{s}_T)$  be the  $T$ -th iterate. Then, the following holds for  $A\mathbf{x}_T - \mathbf{b}$ :

$$A\mathbf{x}_T - \mathbf{b} = A\mathbf{x}_0 + A \sum_{t=1}^T \overline{\Delta \mathbf{x}_t} - \mathbf{b} = A \sum_{t=1}^T \overline{\Delta \mathbf{x}_t}. \quad (9)$$

On the other hand, the Newton system at iteration  $T$  has the constraint  $A\Delta \mathbf{x}_T = \mathbf{b} - A\mathbf{x}_{T-1}$ , which we can further recursively transform as,

$$\begin{aligned} A\Delta \mathbf{x}_T &= \mathbf{b} - A\mathbf{x}_{T-1} = \mathbf{b} - A(\mathbf{x}_{T-2} + \overline{\Delta \mathbf{x}_{T-1}}) \\ &= \mathbf{b} - A\mathbf{x}_0 - \sum_{t=1}^{T-1} \overline{\Delta \mathbf{x}_t} = - \sum_{t=1}^{T-1} \overline{\Delta \mathbf{x}_t}. \end{aligned}$$

Substituting this into equation (9), we get

$$A\mathbf{x}_T - \mathbf{b} = A(\overline{\Delta \mathbf{x}_T} - \Delta \mathbf{x}_T).$$

Similarly, using the constraint  $A^\top \Delta \mathbf{y}_T + \Delta \mathbf{s}_T = \mathbf{c} - \mathbf{s}_{T-1} - A^\top \mathbf{y}_{T-1}$  we obtain that

$$A^\top \mathbf{y}_T + \mathbf{s}_T - \mathbf{c} = A^\top (\overline{\Delta \mathbf{y}_T} - \Delta \mathbf{y}_T) + (\overline{\Delta \mathbf{s}_T} - \Delta \mathbf{s}_T).$$

Finally, we can bound the norms of these two quantities,

$$\begin{aligned} \|A\mathbf{x}_T - \mathbf{b}\| &\leq \delta \|A\|, \\ \|A^\top \mathbf{y}_T + \mathbf{s}_T - \mathbf{c}\| &\leq \delta (\|A\| + 1). \end{aligned}$$

□

## 5 Quantum Support-Vector Machines

In this section we present our quantum support vector machine (SVM) algorithm as an application of our SOCP solver. Given a set of vectors  $\mathcal{X} = \{\mathbf{x}^{(i)} \in \mathbb{R}^n \mid i \in [m]\}$  (*training examples*) and their *labels*  $y^{(i)} \in \{-1, 1\}$ , the objective of the SVM training process is to find the “best” hyperplane that separates training examples with label 1 from those with label  $-1$ . In this paper we focus on the (traditional) *soft-margin* ( $\ell_1$ -)SVM, which can be expressed as the following optimization problem:

$$\begin{aligned} \min_{\mathbf{w}, b, \boldsymbol{\xi}} \quad & \|\mathbf{w}\|^2 + C\|\boldsymbol{\xi}\|_1 \\ \text{s.t.} \quad & y^{(i)}(\mathbf{w}^\top \mathbf{x}^{(i)} + b) \geq 1 - \xi_i, \quad \forall i \in [m] \\ & \boldsymbol{\xi} \geq 0. \end{aligned} \quad (10)$$

Here, the variables  $\mathbf{w} \in \mathbb{R}^n$  and  $b \in \mathbb{R}$  correspond to the hyperplane,  $\boldsymbol{\xi} \in \mathbb{R}^m$  corresponds to the “linear inseparability” of each point, and the constant  $C > 0$  is a hyperparameter that quantifies the tradeoff between maximizing the margin and minimizing the constraint violations.

As a slightly less traditional alternative, one might also consider the  $\ell_2$ -SVM (or least-squares SVM, LS-SVM) [37], where the  $\|\boldsymbol{\xi}\|_1$  regularization term is replaced by  $\|\boldsymbol{\xi}\|^2$ . This formulation arises from considering the least-squares regression problem with the constraints  $y^{(i)}(\mathbf{w}^\top \mathbf{x}^{(i)} + b) = 1$ , which we solve by minimizing the squared 2-norm of the residuals:

$$\begin{aligned} \min_{\mathbf{w}, b, \boldsymbol{\xi}} \quad & \|\mathbf{w}\|^2 + C\|\boldsymbol{\xi}\|^2 \\ \text{s.t.} \quad & y^{(i)}(\mathbf{w}^\top \mathbf{x}^{(i)} + b) = 1 - \xi_i, \forall i \in [m] \end{aligned} \quad (11)$$

Since this is a least-squares problem, the optimal  $\mathbf{w}, b$  and  $\boldsymbol{\xi}$  can be obtained by solving a linear system. In [33], a quantum algorithm for LS-SVM is presented, which uses a single quantum linear system solver. Unfortunately, replacing the  $\ell_1$ -norm with  $\ell_2$  in the objective of (11) leads to the loss of a key property of ( $\ell_1$ -)SVM – weight sparsity [38].

## 5.1 Reducing SVM to SOCP

Finally, we are going to reduce the SVM problem (10) to SOCP. In order to do that, we define an auxiliary vector  $\mathbf{t} = (t + 1; t; \mathbf{w})$ , where  $t \in \mathbb{R}$  – this allows us to “compute”  $\|\mathbf{w}\|^2$  using the constraint  $\mathbf{t} \in \mathcal{L}^{n+2}$  since

$$\mathbf{t} \in \mathcal{L}^{n+2} \Leftrightarrow (t + 1)^2 \geq t^2 + \|\mathbf{w}\|^2 \Leftrightarrow 2t + 1 \geq \|\mathbf{w}\|^2.$$

Thus, minimizing  $\|\mathbf{w}\|^2$  is equivalent to minimizing  $t$ . Note we can restrict our bias  $b$  to be nonnegative without any loss in generality, since the case  $b < 0$  can be equivalently described by a bias  $-b > 0$  and weights  $-\mathbf{w}$ . Using these transformations, we can restate (10) as the following SOCP:

$$\begin{aligned} \min_{t, b, \boldsymbol{\xi}} \quad & \begin{bmatrix} 0 & 1 & 0^n & 0 & C^m \end{bmatrix} \begin{bmatrix} t & b & \boldsymbol{\xi} \end{bmatrix}^\top \\ \text{s.t.} \quad & \begin{bmatrix} 0 & 0 & & 1 \\ \vdots & \vdots & X^\top & \vdots \\ 0 & 0 & & 1 \\ 1 & -1 & 0^n & 0 & 0^m \end{bmatrix} \begin{bmatrix} t \\ b \\ \boldsymbol{\xi} \end{bmatrix} = \begin{bmatrix} \mathbf{y} \\ 1 \end{bmatrix} \\ & \mathbf{t} \in \mathcal{L}^{n+2}, b \in \mathcal{L}^1, \xi_i \in \mathcal{L}^1 \quad \forall i \in [m] \end{aligned} \quad (12)$$

Here, we use the notation  $X \in \mathbb{R}^{n \times m}$  for the matrix whose columns are the training examples  $\mathbf{x}^{(i)}$ , and  $\mathbf{y} \in \mathbb{R}^m$  for the vector of labels. This problem has  $O(n + m)$  variables, and  $O(m)$  conic constraints (i.e. its rank is  $r = O(m)$ ). Therefore, in the interesting case of  $m = \Theta(n)$ , it can be solved in  $\tilde{O}(\sqrt{n})$  iterations. More precisely, if we consider both the primal and the dual, in total they have  $3m + 2n + 7$  scalar variables and  $2m + 4$  conic constraints.

In practice (as evidenced by the LIBSVM and LIBLINEAR libraries [12, 16]), a small modification is made to the formulations (10) and (11): instead of treating the bias separately, all data points are extended with a constant unit coordinate. In this case, the SOCP formulation remains almost identical, with the only difference being that the constraints  $\mathbf{t} \in \mathcal{L}^{n+2}$  and  $b \in \mathcal{L}^1$  are replaced by a single conic constraint  $(\mathbf{t}; b) \in \mathcal{L}^{n+3}$ . This change allows us to come up with a simple feasible initial solution in our numerical experiments, without going through the homogeneous self-dual formalism of [42].




Figure 1: Observed complexity of Algorithm 1, its power law fit, and 95% confidence interval.




Figure 2: Empirical CDF of the difference in accuracy between SVMs trained in different ways.

Note also that we can solve the LS-SVM problem (11), by reducing it to a SOCP in a similar manner. In fact, this would have resulted in just  $O(1)$  conic constraints, so an IPM would converge to a solution in  $\tilde{O}(1)$  iterations, which is comparable with the result from [33].

## 5.2 Experimental results

We next present some experimental results to assess the running time parameters and the performance of our algorithm for random instances of SVM. If an algorithm demonstrates a speedup on unstructured instances like these, it is reasonable to extrapolate that the speedup is generic, as it could not have used any special properties of the instance to derive an advantage. For a given dimension  $n$  and number of training points  $m$ , we denote our distribution of random SVMs with  $\mathcal{SVM}(n, m, p)$ , where  $p$  denotes the probability that a datapoint is misclassified by the optimal separating hyperplane. Additionally, for every training set sampled from  $\mathcal{SVM}(n, m, p)$ , a corresponding test set of size  $\lfloor m/3 \rfloor$  was also sampled from the same distribution. These test sets are used to evaluate the generalization error of SVMs trained in various ways.

Our experiments consist of generating roughly 16000 instances of  $\mathcal{SVM}(n, 2n, p)$ , where  $n$  is chosen to be uniform between  $2^2$  and  $2^9$  and  $p$  is chosen uniformly from the discrete set  $\{0, 0.1, \dots, 0.9, 1\}$ . The instances are then solved using a simulation of Algorithm 1 (with the target duality gap of  $\epsilon = 0.1$ ) as well as using the ECOS SOCP solver [15] (with the default target duality gap). We simulate the execution of Algorithm 1 by implementing the classical IPM and adding noise to the solution of the Newton system (7). The noise added to each coordinate is uniform, from an interval selected so that the noisy increment  $(\Delta \mathbf{x}, \Delta \mathbf{y}, \Delta \mathbf{s})$  simulates the outputs of the tomography algorithm with precision determined by Theorem 4. The SVM parameter  $C$  is set to be equal to 1 in all experiments. Additionally, a separate, smaller experiment with roughly 1000 instances following the same distribution is performed for comparing Algorithm 1 with LIBSVM [12] using a linear kernel.

The experiments are performed on a Dell Precision 7820T workstation with two Intel Xeon Silver 4110 CPUs and 64GB of RAM, and experiment logs are available at [39]. By finding the least-squares fit of the power law  $y = ax^b$  through the observed values of the quantity  $\frac{n^{1.5} \kappa \zeta}{\delta^2}$ , we obtain the exponent  $b = 2.591$ , and its 95% confidence interval [2.564, 2.619] (this interval is computed in the standard way using Student's  $t$ -distribution,



as described in [31]). These observations, the power law, and its confidence interval are shown on Figure 1. Thus, we can say that for random  $\mathcal{SVM}(n, 2n, p)$ -instances, and fixed  $\epsilon = 0.1$ , the complexity of Algorithm 1 scales as  $O(n^{2.591})$ . This represents a polynomial improvement over general dense SOCP solvers with complexity  $O(n^{\omega+0.5})$ . In practice, the polynomial speedup is conserved when compared to ECOS [15], that has a measured running time scaling of  $O(n^{3.314})$ , with a 95% confidence interval for the exponent of [3.297, 3.330] (this is consistent with the internal usage of a [36]-like matrix multiplication algorithm, with complexity  $O(n^{2.807})$ ). Neglecting constant factors, this gives us a speedup of  $10^4$  for  $n = 10^6$ . The results from the LIBSVM solver indicate that the training time with a linear kernel has a complexity of  $O(n^{3.112})$ , with a 95% confidence interval for the exponent of [2.799, 3.425]. These results suggest Algorithm 1 retains its advantage even when compared to state-of-the-art specialized classical algorithms.

Additionally, we use the gathered data to verify that the accuracy of our quantum (or approximate) SVM is close to the optimum: Figure 2 shows that both at train- and at test-time the accuracies of all three classifiers are most of the time within a few percent of each other, with Algorithm 1 often outperforming the exact SOCP SVM classifier.

In conclusion, the performed numerical experiments indicate that Algorithm 1 provides a polynomial speedup for solving SOCPs with low- and medium precision requirements. In particular, for SVM, we achieve a polynomial speedup with no detriment to the quality of the trained classifier.

## Acknowledgments

This work was partly supported by IdEx Université de Paris ANR-18-IDEX-0001, as well as the French National Research Agency (ANR) projects QuBIC and QuDATA.

## References

- [1] F. Alizadeh and D. Goldfarb. Second-order cone programming. *Math. Program.*, 95 (1, Ser. B):3–51, 2003. ISSN 0025-5610. DOI: [10.1007/s10107-002-0339-5](https://doi.org/10.1007/s10107-002-0339-5). ISMP 2000, Part 3 (Atlanta, GA).
- [2] Jonathan Allcock and Chang-Yu Hsieh. A quantum extension of SVM-perf for training nonlinear SVMs in almost linear time. *Quantum*, 4:342, October 2020. ISSN 2521-327X. DOI: [10.22331/q-2020-10-15-342](https://doi.org/10.22331/q-2020-10-15-342).
- [3] Tomasz Arodz and Seyran Saeedi. Quantum sparse support vector machines, 2019. URL <https://arxiv.org/abs/1902.01879>.
- [4] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory Comput.*, 8:121–164, 2012. DOI: [10.4086/toc.2012.v008a006](https://doi.org/10.4086/toc.2012.v008a006).
- [5] Aharon Ben-Tal and Arkadi Nemirovski. *Lectures on modern convex optimization*. MPS/SIAM Series on Optimization. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA; Mathematical Programming Society (MPS), Philadelphia, PA, 2001. ISBN 0-89871-491-5. DOI: [10.1137/1.9780898718829](https://doi.org/10.1137/1.9780898718829). Analysis, algorithms, and engineering applications.
- [6] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004. ISBN 0-521-83378-7. DOI: [10.1017/CBO9780511804441](https://doi.org/10.1017/CBO9780511804441).
- [7] Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP solvers: large speed-ups, optimality, and applications to quantum learning. In *46th International Colloquium on Automata,*

- Languages, and Programming*, volume 132 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 27, 14. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. DOI: [10.4230/LIPICs.ICALP.2019.27](https://doi.org/10.4230/LIPICs.ICALP.2019.27).
- [8] Fernando G. S L. Brandão, Richard Kueng, and Daniel Stilck França. Faster quantum and classical sdp approximations for quadratic binary optimization, 2020. URL <https://arxiv.org/abs/1909.04613>.
- [9] Fernando G.S.L. Brandão and Krysta M. Svore. Quantum speed-ups for solving semidefinite programs. In *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*, pages 415–426. IEEE Computer Soc., Los Alamitos, CA, 2017. DOI: [10.1109/focs.2017.45](https://doi.org/10.1109/focs.2017.45).
- [10] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015. DOI: [10.1561/22000000050](https://doi.org/10.1561/22000000050).
- [11] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. The power of block-encoded matrix powers: improved regression techniques via faster Hamiltonian simulation. In *46th International Colloquium on Automata, Languages, and Programming*, volume 132 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 33, 14. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. DOI: [10.4230/LIPICs.ICALP.2019.33](https://doi.org/10.4230/LIPICs.ICALP.2019.33).
- [12] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):1–27, April 2011. DOI: [10.1145/1961189.1961199](https://doi.org/10.1145/1961189.1961199).
- [13] Michael B. Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. In *STOC’19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 938–942. ACM, New York, 2019. DOI: [10.1145/3313276.3316303](https://doi.org/10.1145/3313276.3316303).
- [14] Hilary Dollar. *Iterative linear algebra for constrained optimization*. PhD thesis, University of Oxford, 2005. URL <https://www.numerical.rl.ac.uk/people/hsd/thesismain.pdf>.
- [15] Alexander Domahidi, Eric Chu, and Stephen Boyd. ECOS: An SOCP solver for embedded systems. In *2013 European Control Conference (ECC)*. IEEE, July 2013. DOI: [10.23919/ecc.2013.6669541](https://doi.org/10.23919/ecc.2013.6669541).
- [16] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. LIBLINEAR: A library for large linear classification. *Journal of machine learning research*, 9(Aug):1871–1874, 2008. URL <https://www.jmlr.org/papers/volume9/fan08a/fan08a>.
- [17] András Gilyén, Srinivasan Arunachalam, and Nathan Wiebe. Optimizing quantum optimization algorithms via faster quantum gradient computation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1425–1444. SIAM, Philadelphia, PA, 2019. DOI: [10.1137/1.9781611975482.87](https://doi.org/10.1137/1.9781611975482.87).
- [18] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *STOC’19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204. ACM, New York, 2019. DOI: [10.1145/3313276.3316366](https://doi.org/10.1145/3313276.3316366).
- [19] András Gilyén, Seth Lloyd, and Ewin Tang. Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension, 2018. URL <https://arxiv.org/abs/1811.04909>.
- [20] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Com-*

- puting (Philadelphia, PA, 1996), pages 212–219. ACM, New York, 1996. DOI: 10.1145/237814.237866.
- [21] Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15), October 2009. DOI: 10.1103/physrevlett.103.150502.
- [22] Thorsten Joachims. Training linear SVMs in linear time. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '06, page 217–226, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595933395. DOI: 10.1145/1150402.1150429.
- [23] N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing - STOC '84*, pages 302–311. ACM Press, 1984. DOI: 10.1145/800057.808695.
- [24] Iordanis Kerenidis and Anupam Prakash. Quantum recommendation systems. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49:1–49:21, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-029-3. DOI: 10.4230/LIPIcs.ITCS.2017.49.
- [25] Iordanis Kerenidis and Anupam Prakash. Quantum gradient descent for linear systems and least squares. *Phys. Rev. A*, 101:022316, Feb 2020. DOI: 10.1103/PhysRevA.101.022316.
- [26] Iordanis Kerenidis and Anupam Prakash. A quantum interior point method for LPs and SDPs. *ACM Transactions on Quantum Computing*, 1(1), October 2020. ISSN 2643-6809. DOI: 10.1145/3406306.
- [27] Yin Tat Lee, Zhao Song, and Qiuyi Zhang. Solving empirical risk minimization in the current matrix multiplication time. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2140–2157, Phoenix, USA, 25–28 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v99/lee19a.html>.
- [28] Renato D. C. Monteiro and Takashi Tsuchiya. Polynomial convergence of primal-dual algorithms for the second-order cone program based on the MZ-family of directions. *Math. Program.*, 88(1, Ser. A):61–83, 2000. ISSN 0025-5610. DOI: 10.1007/PL00011378.
- [29] Yu. E. Nesterov and M. J. Todd. Self-scaled barriers and interior-point methods for convex programming. *Math. Oper. Res.*, 22(1):1–42, 1997. ISSN 0364-765X. DOI: 10.1287/moor.22.1.1.
- [30] Yu. E. Nesterov and M. J. Todd. Primal-dual interior-point methods for self-scaled cones. *SIAM J. Optim.*, 8(2):324–364, 1998. ISSN 1052-6234. DOI: 10.1137/S1052623495290209.
- [31] John Neter, Michael H Kutner, Christopher J Nachtsheim, and William Wasserman. *Applied linear statistical models*, volume 4. Irwin Chicago, 1996.
- [32] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2009. DOI: 10.1017/cbo9780511976667.
- [33] Patrick Reberntrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13), September 2014. DOI: 10.1103/physrevlett.113.130503.
- [34] Yousef Saad. *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, second edition, 2003. ISBN 0-89871-534-2. DOI: 10.1137/1.9780898718003.

- [35] Peter W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *35th Annual Symposium on Foundations of Computer Science (Santa Fe, NM, 1994)*, pages 124–134. IEEE Comput. Soc. Press, Los Alamitos, CA, 1994. DOI: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700).
- [36] Volker Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13(4): 354–356, August 1969. DOI: [10.1007/bf02165411](https://doi.org/10.1007/bf02165411).
- [37] J.A.K. Suykens and J. Vandewalle. Least squares support vector machine classifiers. *Neural Processing Letters*, 9(3):293–300, 1999. DOI: [10.1023/a:1018628609742](https://doi.org/10.1023/a:1018628609742).
- [38] J.A.K. Suykens, J. De Brabanter, L. Lukas, and J. Vandewalle. Weighted least squares support vector machines: robustness and sparse approximation. *Neurocomputing*, 48 (1-4):85–105, October 2002. DOI: [10.1016/s0925-2312\(01\)00644-0](https://doi.org/10.1016/s0925-2312(01)00644-0).
- [39] Daniel Szilagyi, Iordanis Kerenidis, and Anupam Prakash. Quantum SVM via SOCP experiment logs. Mar 2021. DOI: [10.6084/m9.figshare.11778189.v1](https://doi.org/10.6084/m9.figshare.11778189.v1).
- [40] Joran van Apeldoorn and András Gilyén. Improvements in quantum SDP-solving with applications. In *46th International Colloquium on Automata, Languages, and Programming*, volume 132 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 99, 15. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. DOI: [10.4230/LIPICs.ICALP.2019.99](https://doi.org/10.4230/LIPICs.ICALP.2019.99).
- [41] Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: Better upper and lower bounds. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, October 2017. DOI: [10.1109/focs.2017.44](https://doi.org/10.1109/focs.2017.44).
- [42] Yinyu Ye, Michael J. Todd, and Shinji Mizuno. An  $O(\sqrt{n}L)$ -iteration homogeneous and self-dual linear programming algorithm. *Math. Oper. Res.*, 19(1):53–67, 1994. ISSN 0364-765X. DOI: [10.1287/moor.19.1.53](https://doi.org/10.1287/moor.19.1.53).