

# A Catalyst Framework for the Quantum Linear System Problem via the Proximal Point Algorithm

Junhyung Lyle Kim<sup>1</sup>, Nai-Hui Chia<sup>1</sup>, Anastasios Kyriallidis<sup>1</sup>

<sup>1</sup>Department of Computer Science, Rice University  
{jlylekim, nc67, anastasios}@rice.edu

## Abstract

Solving systems of linear equations is a fundamental problem, but it can be computationally intensive for classical algorithms in high dimensions. Existing quantum algorithms can achieve exponential speedups for the quantum linear system problem (QLSP) in terms of the problem dimension, but the advantage is bottlenecked by condition number of the coefficient matrix. In this work, we propose a new quantum algorithm for QLSP inspired by the classical proximal point algorithm (PPA). Our proposed method can be viewed as a meta-algorithm that allows inverting a modified matrix via an existing `QLSP_solver`, thereby directly approximating the solution vector instead of approximating the inverse of the coefficient matrix. By carefully choosing the step size  $\eta$ , the proposed algorithm can effectively precondition the linear system to mitigate the dependence on condition numbers that hindered the applicability of previous approaches. Importantly, this is the first iterative framework for QLSP where a tunable parameter  $\eta$  and initialization  $x_0$  allows controlling the trade-off between the runtime and approximation error.

Extended version — <https://arxiv.org/pdf/2406.13879>

## 1 Introduction

**Background.** Solving systems of linear equations is a fundamental problem with many applications spanning science and engineering. Mathematically, for a given (Hermitian) matrix  $A \in \mathbb{C}^{N \times N}$  and vector  $b \in \mathbb{C}^N$ , the goal is to find the  $N$ -dimensional vector  $x^*$  that satisfies  $Ax^* = b$ . While classical algorithms, such as Gaussian elimination (Gauss 1877; Higham 2011), conjugate gradient method (Hestenes and Stiefel 1952), and  $LU$  decomposition (Schwarzenberg-Czerny 1995; Shabat et al. 2018) can solve this problem, their complexity scales (at worst) cubically with  $N$ , the dimension of  $A$ , motivating the development of quantum algorithms that could potentially achieve speedups.

Indeed, for the quantum linear system problem (QLSP) —a BQP-complete problem<sup>1</sup>— in Definition 2.1, Harrow, Hassidim, and Lloyd (2009) showed that dependence on the problem dimension *exponentially* reduces to

$\mathcal{O}(\text{poly} \log(N))$ , with query complexity of  $\mathcal{O}(\kappa^2/\varepsilon)$ , under some (quantum) access model for  $A$  and  $b$  (c.f., Definitions 2.2 and 2.3). Here,  $\kappa$  is the condition number of  $A$ , defined as the ratio of the largest to the smallest singular value of  $A$ . Subsequent works, such as Ambainis (2012) and Childs, Kothari, and Somma (2017) (a.k.a. the CKS algorithm), improve the dependence on  $\kappa$  and  $\varepsilon$ ; see Table 1. The best quantum algorithm for QLSP is based on the discrete adiabatic theorem (Costa et al. 2022), achieving the query complexity of  $\mathcal{O}(\kappa \cdot \log(1/\varepsilon))$ , (asymptotically) matching the lower bound (Orsucci and Dunjko 2021).<sup>2</sup>

Solving QLSP is a fundamental subroutine in many quantum algorithms. For instance, it is used in quantum recommendation systems (Kerenidis and Prakash 2016), quantum SVM (Rebentrost, Mohseni, and Lloyd 2014), unsupervised learning (Wiebe, Kapoor, and Svore 2014), and solving differential equations (Liu et al. 2021), to name a few. Hence, improving the overall runtime of a generic `QLSP_solver` is crucial in developing more sophisticated and efficient quantum algorithms.

**Challenges in existing methodologies.** A common limitation of existing quantum algorithms is that the dependence on the condition number  $\kappa$  must be small to achieve the (exponential) quantum advantage. To put more context, for any quantum algorithm in Table 1 to achieve an exponential advantage over classical algorithms,  $\kappa$  must be of the order  $\text{poly} \log(N)$ , where  $N$  is the dimension of  $A$ . For instance, for  $N = 10^3$ ,  $\kappa$  needs to be around 4 to exhibit the exponential advantage. However, condition numbers are often large in real-world problems (Pappayan 2020).

Moreover, it was proven (Orsucci and Dunjko 2021, Proposition 6) that for QLSP, even when  $A$  is *positive-definite*, the dependence on condition number *cannot* be improved from  $\mathcal{O}(\kappa)$ . This is in contrast to the classical algorithms, such as the conjugate gradient method, which achieves  $\mathcal{O}(\sqrt{\kappa})$  reduction in complexity for the linear systems of equations with  $A \succ 0$ . Such observation reinforces the importance of alleviating the dependence on the condition number  $\kappa$  for *quantum algorithms*, which is our aim.

**Our contributions.** We present a novel meta-algorithm for solving the QLSP based on the proximal point algorithm

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>A proof can found for instance in Prasad and Zhuang (2022).

<sup>2</sup>This also matches the iteration complexity of the (classical) conjugate gradient method (Hestenes and Stiefel 1952).




Figure 1: (Left) Query complexity scaling with respect to the condition number  $\kappa$ . Here, the SOTA quantum algorithm (Costa et al. 2022), which enjoys (asymptotically) optimal query complexity  $\Omega(\kappa \log \frac{1}{\varepsilon})$ , is used as `QLSP_solver` for the subroutine in Algorithm 1. Simply by “wrapping” the `QLSP_solver`, one can achieve much better scaling with respect to the condition number  $\kappa$ . (Right) Query complexity improvement with warm start. Baseline (blue) is again the (asymptotically) optimal query complexity by Costa et al. (2022). The other three lines are the improved query complexity in (17) using Algorithm 1 where  $x_0$  is initialized with the output of  $\{200, 500, 1000\}$  steps of GD.

(PPA) (Rockafellar 1976; Güler 1991); see Algorithm 1. Notably, in contrast to the existing methods that approximate  $A^{-1}$  (Harrow, Hassidim, and Lloyd 2009; Ambainis 2012; Childs, Kothari, and Somma 2017; Gribling, Kerenidis, and Szilágyi 2021; Orsucci and Dunjko 2021), our algorithm directly approximates  $x^* = A^{-1}b$  through an iterative process based on PPA. Classically, PPA is known to improve the “conditioning” of the problem at hand, compared to the gradient descent (Toulis, Airoidi, and Rennie 2014; Toulis and Airoidi 2017; Ahn and Sra 2022); it can also be accelerated (Güler 1992; Kim, Toulis, and Kyrillidis 2022).

An approximate proximal point algorithm for classical convex optimization has been proposed under the name of “catalyst” (Lin, Mairal, and Harchaoui 2015) in the machine learning community. Our proposed method operates similarly and can be viewed as a generic acceleration scheme for QLSP where one can plug in different `QLSP_solver`—e.g., HHL—to achieve generic (constant-level) acceleration.

In Figure 1 (left), we illustrate the case where the (asymptotically) optimal quantum algorithm for QLSP by Costa et al. (2022), based on the discrete adiabatic theorem, is utilized as the `QLSP_solver` subroutine for Algorithm 1. Simply by “wrapping” it with our meta-algorithm, one can achieve significant constant-level improvement in the query complexity to achieve a fixed accuracy. Importantly, the improvement gets more pronounced as  $\kappa$  increases.

Intuitively, by the definition of PPA detailed in Section 3, our method allows one to invert a modified matrix  $I + \eta A$ , with initialization  $x_0$ , and arrive at the same solution:

$$\hat{x} = (I + \eta A)^{-1}(x_0 + \eta b) \underset{\eta \rightarrow \infty}{=} A^{-1}b.$$

Yet, a key feature and the main distinction from existing `QLSP_solver` is the introduction of a *tunable (step size) parameter*  $\eta$  that allows preconditioning the linear system. By appropriately choosing  $\eta$ , we can invert the (normalized) modified matrix  $I + \eta A$  that is better conditioned than  $A$ , thereby mitigating the dependence on  $\kappa$  of the existing quantum algorithms; see also Remark 3.4.

Furthermore, one can also exploit the freedom of choosing the *initialization*  $x_0$ , another feature not possible in existing quantum algorithms. For example,  $x_0$  can be initialized with the output of gradient descent, as illustrated in Figure 1 (right). The baseline (blue) is the (asymptotically) optimal query complexity  $\Omega(\kappa \log \frac{1}{\varepsilon})$  (Costa et al. 2022), which can effectively be halved (red) via Algorithm 1 with warm start; see Section 4 for details.

Our framework complements and provides advantages over prior works on QLSP. Most importantly, it provides knobs to maneuver existing `QLSP_solvers`, via the step size  $\eta$  and the initialization  $x_0$ , enabling quantum speedups for a broader class of problems where  $\kappa$  may be large. Our contributions can be summarized as follows:

- We propose a meta-algorithmic framework for the quantum linear system problem (QLSP) based on the proximal point algorithm. Unlike existing quantum algorithms for QLSP, which rely on different unitary approximations of  $A^{-1}$ , our proposed method allows to invert a modified matrix with a smaller condition number (c.f., Remark 3.4 and Lemma 3.3), when  $A$  is positive-definite.
- To our knowledge, this is the first framework for QLSP where a *tunable parameter*  $\eta$  and *initialization*  $x_0$  allow users to control the trade-off between the runtime (query complexity) and the approximation error (solution quality). Importantly, there exist choices of  $\eta > 0$  that allow to decrease the runtime while maintaining the same error level (c.f., Theorem 4.5).
- Our proposed method allows to achieve significant constant-level improvements in the query complexity, even compared to the (asymptotically) optimal quantum algorithm for QLSP (Costa et al. 2022), simply by using it as a subroutine of Algorithm 1. This is possible as the improvement “grows faster” than the overhead (c.f., Figure 2 and Theorem 4.8), which can further be improved via warm start (c.f., Section 4) in practice.

## 2 Problem Setup and Related Work

**Notation.** Matrices are represented with uppercase letters as in  $A \in \mathbb{C}^{N \times N}$ ; vectors are represented with lowercase letters as in  $b \in \mathbb{C}^N$ , and are distinguished from scalars based on the context. The condition number of a matrix  $A$ , denoted as  $\kappa$ , is the ratio of the largest to the smallest singular value of  $A$ . We denote  $\|\cdot\|$  as the Euclidean norm. A *Qubit* is the fundamental unit in quantum computing, analogous to a *bit* in classical computing. The state of a qubit is represented using the bra-ket notation, where a single qubit state  $|\psi\rangle \in \mathbb{C}^2$  can be expressed as a linear combination of the basis states  $|0\rangle = [0 \ 1]^\top \in \mathbb{C}^2$  and  $|1\rangle = [1 \ 0]^\top \in \mathbb{C}^2$ , as in  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ ; here,  $\alpha, \beta \in \mathbb{C}$  are called *amplitudes* and encode the probability of the qubit collapsing to either, such that  $|\alpha|^2 + |\beta|^2 = 1$ .  $|\cdot\rangle$  is a column vector (called *bra*), and its conjugate transpose (called *ket*), denoted by  $\langle \cdot |$ , is defined as  $\langle \cdot | = |\cdot\rangle^*$ . Generalizing this, an  $n$ -qubit state is a unit vector in  $n$ -qubit Hilbert space, defined as the Kronecker product of  $n$  single qubit states, i.e.,  $\mathcal{H} = \otimes_{i=1}^n \mathbb{C}^2 \cong \mathbb{C}^{2^n}$ . It is customary to write  $2^n = N$ . Quantum states can be manipulated using quantum gates, represented by unitary matrices that act on the state vectors. For example, a single-qubit gate  $U \in \mathbb{C}^{2 \times 2}$  acting on a qubit state  $|\psi\rangle$  transforms it to  $U|\psi\rangle$ , altering the state's probability amplitudes.

### The Quantum Linear System Problem

In the quantum setting, the goal of quantum linear system problem (QLSP) is to prepare a *quantum state* proportional to the vector  $x^*$ . That is, we want to output  $|x^*\rangle := \frac{\sum_i x_i^* |i\rangle}{\|\sum_i x_i^* |i\rangle\|}$  where the vector  $x^* = [x_1^*, \dots, x_N^*]^\top$  satisfies  $Ax^* = b$ . We formally define QLSP below, based on Childs, Kothari, and Somma (2017).

**Definition 2.1** (Quantum Linear System Problem). Let  $A$  be an  $N \times N$  Hermitian matrix satisfying  $\|A\| = 1$  with condition number  $\kappa$  and at most  $s$  nonzero entries in any row or column. Let  $b$  be an  $N$ -dimensional vector, and let  $x^* := A^{-1}b$ . We define the quantum states  $|b\rangle$  and  $|x^*\rangle$  as

$$|b\rangle := \frac{\sum_{i=0}^{N-1} b_i |i\rangle}{\|\sum_{i=0}^{N-1} b_i |i\rangle\|} \text{ and } |x^*\rangle := \frac{\sum_{i=0}^{N-1} x_i^* |i\rangle}{\|\sum_{i=0}^{N-1} x_i^* |i\rangle\|}. \quad (1)$$

Given access to  $A$  via  $\mathcal{P}_A$  in Definition 2.3 or  $\mathcal{U}_A$  in Definition 2.4, and access to the state  $|b\rangle$  via  $\mathcal{P}_B$  in Definition 2.2, the goal of QLSP is to output a state  $|\hat{x}\rangle$  such that  $\|\hat{x} - |x^*\rangle\| \leq \epsilon$ .

Following previous works (Harrow, Hassidim, and Lloyd 2009; Childs, Kothari, and Somma 2017; Ambainis 2012; Gribling, Kerenidis, and Szilágyi 2021; Orsucci and Dunjko 2021; Costa et al. 2022), we assume that access to  $A$  and  $b$  is provided by black-box subroutines that we detail below. We start with the state preparation oracle for the vector  $b \in \mathbb{C}^N$ .

**Definition 2.2** (State preparation oracle (Harrow, Hassidim, and Lloyd 2009)). Given a vector  $b \in \mathbb{C}^N$ , there exists a procedure  $\mathcal{P}_B$  that prepares the state  $|b\rangle := \frac{\sum_i b_i |i\rangle}{\|\sum_i b_i |i\rangle\|}$  in time  $\mathcal{O}(\text{poly} \log(N))$ .

We assume two encoding models for  $A$ : the sparse-matrix-access in Definition 2.3 denoted by  $\mathcal{P}_A$ , and the matrix-block-encoding model (Gilyén et al. 2019; Low and Chuang 2019) in Definition 2.4, denoted by  $\mathcal{U}_A$ .

**Definition 2.3** (Sparse matrix access (Childs, Kothari, and Somma 2017)). Given a  $N \times N$  Hermitian matrix  $A$  with operator norm  $\|A\| \leq 1$  and at most  $s$  nonzero entries in any row or column,  $\mathcal{P}_A$  allows the following mapping:

$$|j, \ell\rangle \mapsto |j, \nu(j, \ell)\rangle \quad \forall j \in [N] \text{ and } \ell \in [s], \quad (2)$$

$$|j, k, z\rangle|0\rangle \mapsto |j, k, z \oplus A_{jk}\rangle \quad \forall j, k \in [N], \quad (3)$$

where  $\nu : [N] \times [s] \rightarrow [N]$  in (2) computes the row index of the  $\ell$ th nonzero entry of the  $j$ th column, and the third register of (3) holds a bit string representing the entry  $A_{jk}$ .

**Definition 2.4** (Matrix block-encoding (Gilyén et al. 2019)). A unitary operator  $\mathcal{U}_A$  acting on  $n + c$  qubits is called an  $(\alpha, c, \epsilon)$ -matrix-block-encoding of a  $n$ -qubit operator  $A$  if

$$\|A - \alpha(\langle 0^c | \otimes I) \mathcal{U}_A (|0^c\rangle \otimes I)\| \leq \epsilon.$$

The above can also be expressed as follows:

$$\mathcal{U}_A = \begin{bmatrix} \tilde{A} & * \\ \alpha & * \\ * & * \end{bmatrix} \quad \text{with} \quad \|\tilde{A} - A\| \leq \epsilon,$$

where  $*$ 's denote arbitrary matrix blocks with appropriate dimensions.

In Childs, Kothari, and Somma (2017), it was shown that a  $(s, 1, 0)$ -matrix-block-encoding of  $A$  is possible using a constant number of calls to  $\mathcal{P}_A$  in Definition 2.3 (and  $\mathcal{O}(\text{poly}(n))$  extra elementary gates). In short,  $\mathcal{P}_A$  in Definition 2.3 implies efficient implementation of  $\mathcal{U}_A$  in Definition 2.4 (c.f., (Gilyén et al. 2019, Lemma 48)). We present both for completeness as different works rely on different access models; however, our proposed meta-algorithm can provide generic acceleration for any QLSP solver, regardless of the encoding method.

### Related Work

**Quantum algorithms.** We summarize the related quantum algorithms for QLSP and their query complexities in Table 1; all QLSP solvers share the exponential improvement on the input dimension,  $\mathcal{O}(\text{poly} \log(N))$ . The HHL algorithm (Harrow, Hassidim, and Lloyd 2009) utilizes quantum subroutines including (i) Hamiltonian simulation (Feynman 1982; Lloyd 1996; Childs et al. 2018) that applies the unitary operator  $e^{iAt}$  to  $|b\rangle$  for a superposition of different times  $t$ , (ii) phase estimation (Kitaev 1995) that allows to decompose  $|b\rangle$  into the eigenbasis of  $A$  and to find its corresponding eigenvalues, and (iii) amplitude amplification (Brassard and Hoyer 1997; Grover 1998; Brassard et al. 2002) that allows to implement the final state with amplitudes the same with the elements of  $x^*$ . Subsequently, Ambainis (2012) achieved a quadratic improvement on the condition number at the cost of worse error dependence; the main technical contribution was to improve the amplitude amplification, the previous bottleneck. CKS (Childs, Kothari, and Somma 2017) significantly improved the suboptimality by the linear combination of unitaries. We review these subroutines in the supplementary material.

QLSP_solver	Query Complexity	Key Technique/Result
HHL (Harrow, Hassidim, and Lloyd 2009)	$\mathcal{O}(\kappa^2/\varepsilon)$	First quantum algorithm for QLSP
Ambainis (Ambainis 2012)	$\mathcal{O}(\kappa \log^3(\kappa)/\varepsilon^3)$	Variable Time Amplitude Amplification
CKS (Childs, Kothari, and Somma 2017)	$\mathcal{O}(\kappa \cdot \text{poly} \log(\kappa/\varepsilon))$	(Truncated) Chebyshev bases via LCU
Subaşı et al. (Subaşı, Somma, and Orsucci 2019)	$\mathcal{O}((\kappa \log \kappa)/\varepsilon)$	Adiabatic Randomization Method
An & Lin (An and Lin 2022)	$\mathcal{O}(\kappa \cdot \text{poly} \log(\kappa/\varepsilon))$	Time-Optimal Adiabatic Method
Lin & Tong (Lin and Tong 2020)	$\mathcal{O}(\kappa \cdot \log(\kappa/\varepsilon))$	Zeno Eigenstate Filtering
Costa, et al. (Costa et al. 2022)	$\mathcal{O}(\kappa \cdot \log(1/\varepsilon))$	Discrete Adiabatic Theorem
<b>Ours</b>	$\kappa \rightarrow \frac{\kappa(1+\eta)}{\kappa+\eta}$ for all above	Proximal Point Algorithm

Table 1: *Query complexities and key results used in related works on QLSP.* Our proposed framework in Algorithm 1 allows to improve the dependence on the condition number  $\kappa$  for any `QLSP_solver`, so long as the input matrix  $A$  is positive-definite.

Costa et al. (2022) is the state-of-the-art QLS algorithm based on the adiabatic framework, which was spearheaded by Subaşı, Somma, and Orsucci (2019) and improved in An and Lin (2022); Lin and Tong (2020). These are significantly different from the aforementioned HHL-based approaches. Importantly, Algorithm 1 is oblivious to such differences and provides generic acceleration without assuming additional structure.

**Lower bounds.** Along with the first quantum algorithm for QLSP, Harrow, Hassidim, and Lloyd (2009) also proved the lower bound of  $\Omega(\kappa)$  queries to the entries of the matrix is needed for general linear systems. In Orsucci and Dunjko (2021, Proposition 6), this lower bound was surprisingly extended to the case of *positive-definite* systems. This is in contrast to the classical optimization literature, where methods such as the conjugate gradient method (Hestenes and Stiefel 1952) achieve  $\sqrt{\kappa}$ -acceleration for positive-definite systems. Therefore, a constant-level improvement we achieve in this work is the most one can hope for.

### 3 The Proximal Point Algorithm for QLSP

We now introduce our proposed method, summarized in Algorithm 1. Our method can be viewed as a meta-algorithm, where one can plug in any existing `QLSP_solver` as a subroutine to achieve generic acceleration. We first review the proximal point algorithm (PPA), a classical optimization method on which Algorithm 1 is based.

#### The Proximal Point Algorithm

We take a step back from the QLSP in Definition 2.1 and introduce the proximal point algorithm (PPA), a fundamental optimization method in convex optimization (Rockafellar 1976; Güler 1991; Parikh, Boyd et al. 2014; Bauschke and Combettes 2019). PPA is an iterative algorithm that proceeds by minimizing the original function plus an additional quadratic term, as in:

$$x_{t+1} = \arg \min_x \left\{ f(x) + \frac{1}{2\eta} \|x - x_t\|_2^2 \right\}. \quad (4)$$

As a result, it changes the “conditioning” of the problem; if  $f(\cdot)$  is convex, the optimization problem in (4) can be strongly convex (Ahn and Sra 2022). By the first-order optimality condition (Boyd and Vandenberghe 2004, Eq. (4.22)),

(4) can be written in the following form, also known as the implicit gradient descent (IGD):

$$x_{t+1} = x_t - \eta \nabla f(x_{t+1}). \quad (5)$$

(5) is an implicit method and generally cannot be implemented. However, the case we are interested in is the quadratic minimization problem:

$$\min_x f(x) = \frac{1}{2} x^\top A x - b^\top x. \quad (6)$$

Then, we have the closed-form update for (5) as follows:

$$x_{t+1} = x_t - \eta(Ax_{t+1} - b) = x_t - \eta A(x_{t+1} - x^*).$$

In particular, by rearranging and unfolding, we have

$$\begin{aligned} x_{t+1} &= (I + \eta A)^{-1}(x_t + \eta b) = \dots \\ &= (I + \eta A)^{-(t+1)}x_0 + \eta b \sum_{k=1}^{t+1} (I + \eta A)^{-k}. \end{aligned} \quad (7)$$

The above expression sheds some light on how applying PPA can differ from simple inversion:  $A^{-1}b$ . In particular, PPA enables inverting a modified matrix  $I + \eta A$  based on  $\eta$ ; see also Remark 3.4 below.

Further, since  $b = Ax^*$ , we can equivalently express the series of operations in (7) as follows:

$$\begin{aligned} x_{t+1} - x^* &= (I + \eta A)^{-1}(x_t - x^*) = \dots \\ &= (I + \eta A)^{-t}(x_0 - x^*). \end{aligned} \quad (8)$$

The above expression helps computing the number of iterations required for PPA, given  $\eta > 0$ , for finding  $\varepsilon$ -approximate solution, as we detail in Section 4.

#### Meta-Algorithm for QLSP via PPA

We present our proposed method, which is extremely simple as summarized in Algorithm 1.

Line 2 is the cornerstone of the algorithm where one can employ any `QLSP_solver`-like HHL (Harrow, Hassidim, and Lloyd 2009), CKS (Childs, Kothari, and Somma 2017) or the recent work based on discrete adiabatic approach (Costa et al. 2022)– to the (normalized) matrix  $(I + \eta A)/\|I + \eta A\|$ , enabled by the PPA approach. In other words, line 3 can be seen as the output of applying `QLSP_solver` $(\frac{I+\eta A}{\|I+\eta A\|}, |x_0 + \eta b\rangle)$ . We make some remarks on the input.

---

**Algorithm 1: Proximal Point Algorithm for the Quantum Linear System Problem**


---

- 1: **Input:** An oracle  $\mathcal{P}_{\eta,A}$  that prepares sparse-access or block-encoding of  $\frac{I+\eta A}{\|I+\eta A\|}$  (c.f., Definition 2.3 and Definition 2.4); a state preparation oracle  $\mathcal{P}_{b,x_0,\eta}$  that prepares  $|x_0 + \eta b\rangle$  (c.f., Definition 2.2); and a tunable step size  $\eta > 0$ .
  - 2: **Subroutine:** Invoke any `QLSP_solver` such that  $|\psi_{I+\eta A, x_0+\eta b}\rangle \approx \left| \left( \frac{I+\eta A}{\|I+\eta A\|} \right)^{-1} (x_0 + \eta b) \right\rangle$ .
  - 3: **Output:** Normalized quantum state  $|\psi_{I+\eta A, x_0+\eta b}\rangle$ .
  - 4: **Benefit:** Improved dependence on condition number as in Remark 3.4.
- 

*Remark 3.1* (Access to  $(I + \eta A)/\|I + \eta A\|$ ). Our algorithm necessitates an oracle  $\mathcal{P}_{\eta,A}$ , which can provide either sparse access to  $(I + \eta A)/\|I + \eta A\|$  (as in Definition 2.3) or its block encoding (as in Definition 2.4). For sparse access to  $(I + \eta A)/\|I + \eta A\|$ , direct access is feasible from sparse access to  $A$ , when all diagonal entries of  $A$  are non-zero. Specifically, the access described in (3) is identical to that of  $A$ , while the access in (2) modifies  $A_{jj}$  to  $(1 + \eta A_{jj})/\|I + \eta A\|$ . If  $A$  has zero diagonal entries, sparse access to  $(I + \eta A)/\|I + \eta A\|$  requires only two additional uses of (3).

For block encoding,  $(I + \eta A)/\|I + \eta A\|$  can be achieved through a linear combination of block-encoded matrices, as demonstrated in (Gilyén et al. 2019, Lemma 52). This method allows straightforward adaptation of existing data structures that facilitate sparse-access or block-encoding of  $A$  to also support  $(I + \eta A)/\|I + \eta A\|$ .

In addition, original data structures that provide sparse access or block-encoding of  $A$  can be easily modified for access to  $(I + \eta A)/\|I + \eta A\|$ . For instance, suppose that we are given a matrix as a sum of multiple small matrices as in the local Hamiltonian problem or Hamiltonian simulation. Then the sparse access and block encoding of both  $A$  and  $(I + \eta A)/\|I + \eta A\|$  can be efficiently derived from the sum of small matrices.

*Remark 3.2* (Access to  $|x_0 + \eta b\rangle$ ). We prepare the initial state  $|x_0 + \eta b\rangle$ , instead of  $|b\rangle$ , reflecting the PPA update in (7). For this step, we assume there exists an oracle  $\mathcal{P}_{b,x_0,\eta}$  such that  $|x_0 + \eta b\rangle$  can be efficiently prepared similarly to Definition 2.2.

Since we are inverting the (normalized) modified matrix  $\frac{I+\eta A}{\|I+\eta A\|}$ , the spectrum changes as follows.

**Lemma 3.3.** *Let  $A$  be an  $N \times N$  Hermitian positive-definite matrix satisfying  $\|A\| = 1$  with condition number  $\kappa$ . Then, the condition number of the modified matrix in Algorithm 1,  $\frac{I+\eta A}{\|I+\eta A\|}$ , is given by  $\hat{\kappa} = \frac{\kappa(1+\eta)}{\kappa+\eta}$ .*

Notice that the modified condition number  $\hat{\kappa}$  depends both on  $\kappa$  and the step size parameter  $\eta$  for PPA. As a result,  $\eta$  plays a crucial role in the overall performance of Algorithm 1. The introduction of the tunable parameter  $\eta$  in the context of QLSP is one of the main properties that differentiates Algorithm 1 from other quantum algorithms. We summarize the trade-off of  $\eta$  in the following remark.

*Remark 3.4* ( $\eta$  trade-off). The modified condition number,  $\hat{\kappa} = \frac{\kappa(1+\eta)}{\kappa+\eta}$ , in conjunction with the PPA convergence in (8) introduces a trade-off based on the tunable parameter  $\eta$ .

- Large  $\eta$  regime: PPA to converge fast, as can be seen in (15). On the other hand, the benefit of the modified condition number diminishes and recovers the original  $\kappa$ :

$$\hat{\kappa} = \frac{\kappa(1+\eta)}{\kappa+\eta} \xrightarrow{\eta \rightarrow \infty} \kappa,$$

- Small  $\eta$  regime: PPA requires more number of iterations, as can be seen in (15). However, the modified condition number becomes increasingly better conditioned, as in:

$$\hat{\kappa} = \frac{\kappa(1+\eta)}{\kappa+\eta} \xrightarrow{\eta \rightarrow 0} 1.$$

## 4 Theoretical Analysis

As shown in Algorithm 1 as well as the PPA iteration explained in (7), the main distinction of our proposed method from the existing QLSP algorithms is that we invert the modified matrix  $(I + \eta A)/\|I + \eta A\|$  instead of the original matrix  $A$ . Precisely, our goal is to bound the following:

$$\left\| \underbrace{|\psi_{I+\eta A, x_0+\eta b}\rangle - \left| \left( \frac{I+\eta A}{\|I+\eta A\|} \right)^{-1} (x_0 + \eta b) \right\rangle}_{\text{QLSP\_solver error (c.f., Proposition 4.1)}} + \underbrace{\left| \left( \frac{I+\eta A}{\|I+\eta A\|} \right)^{-1} (x_0 + \eta b) \right\rangle - |A^{-1}b\rangle}_{\text{PPA error (c.f., Proposition 4.4)}} \right\| \leq \varepsilon, \quad (9)$$

where  $|\psi_{I+\eta A, x_0+\eta b}\rangle$  is the output of Algorithm 1, and  $|A^{-1}b\rangle = \frac{A^{-1}b}{\|A^{-1}b\|}$  is the target quantum state of QLSP based on Definition 2.1. In between, the term  $\left| \left( \frac{I+\eta A}{\|I+\eta A\|} \right)^{-1} (x_0 + \eta b) \right\rangle$  is added and subtracted, reflecting the modified inversion due to PPA. Specifically, the first pair of terms quantifies the error coming from the inexactness of any `QLSP_solver` used as a subroutine of Algorithm 1. The second pair of terms quantifies the error coming from PPA in estimating  $A^{-1}b$ , as can be seen in (8). In the following subsections, we analyze each term carefully. All proofs can be found in the supplementary material.

### Inverting the Modified Matrix $(I + \eta A)/\|I + \eta A\|$

In the first part, we apply any `QLSP_solver` to invert the modified matrix  $(I + \eta A)/\|I + \eta A\|$ . Importantly, we first need to encode it into a quantum computer, and this requires normalization so that the resulting encoded matrix is unitary, which is necessary for any quantum computer operation (c.f., Remarks 3.1 and 3.2). Since we are invoking existing `QLSP_solver`, we have to control the error  $\varepsilon_1$  and its contribution to the final accuracy  $\varepsilon$ , as summarized below.

**Proposition 4.1** (`QLSP_solver` error). *Assume access to the oracle  $\mathcal{P}_{\eta,A}$  in Algorithm 1, similarly to the Definition 2.3 or 2.4 (c.f., Remark 3.1). Further, assume access to the oracle  $\mathcal{P}_{b,x_0,\eta}$  similarly to the Definition 2.2 (c.f.,*

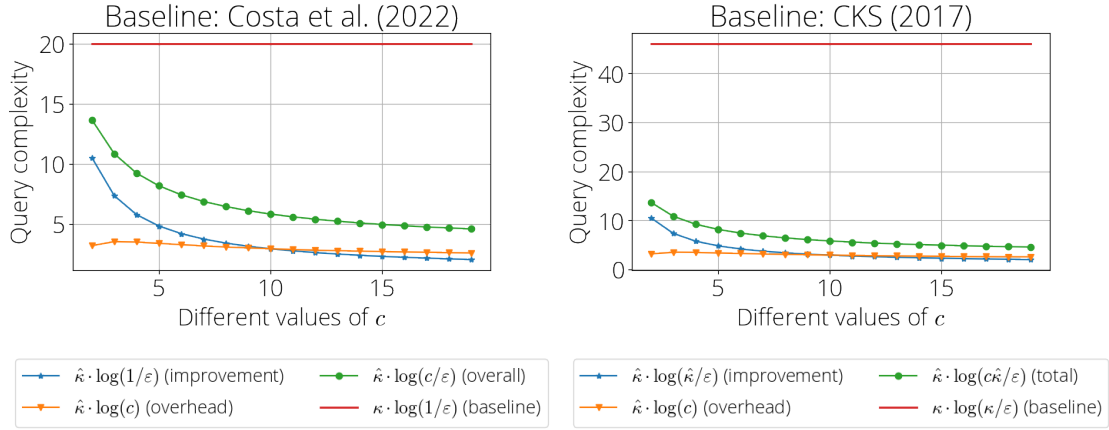


Figure 2: How the “improvement” term, the “overhead” term, and their sum “total” allows improvement compared to the baselines from (Childs, Kothari, and Somma 2017) and (Costa et al. 2022) (c.f., Theorem 4.8). (Right): using CKS (Childs, Kothari, and Somma 2017) as baseline (c.f., Theorem 4.6); (Left): using (Costa et al. 2022) as baseline (c.f., Theorem 4.7). The key insight is that the rate of “improvement” is faster than the rate of “overhead,” which plateaus quickly thanks to the logarithm as can be seen in Theorem 4.8.

Remark 3.2), respectively. Then, there exists quantum algorithms, such as the ones in Table 1, satisfying

$$\left\| \left| \psi_{I+\eta A, x_0+\eta b} \right\rangle - \left| \left( \frac{I+\eta A}{\|I+\eta A\|} \right)^{-1} (x_0 + \eta b) \right\rangle \right\| \leq \varepsilon_1. \quad (10)$$

**Example 4.2** (CKS polynomial (Childs, Kothari, and Somma 2017)). CKS performed the following polynomial approximation of  $A^{-1}$ :

$$A^{-1} \approx P(A) = \sum_i \alpha_i \mathcal{T}_i(A), \quad (11)$$

where  $\mathcal{T}_i$  denotes the Chevshev polynomials (of the first kind). Then, given that  $\|P(A) - A^{-1}\| \leq \varepsilon_1 < \frac{1}{2}$ , (Childs, Kothari, and Somma 2017, Proposition 9) proves:

$$\left\| \frac{P(A)|\phi\rangle}{\|P(A)|\phi\rangle\|} - \frac{A^{-1}|\phi\rangle}{\|A^{-1}|\phi\rangle\|} \right\| \leq 4\varepsilon_1. \quad (12)$$

**Adapting Example 4.2 to Algorithm 1.** The approximation in (11) is only possible when the norm of the matrix to be inverted is upper-bounded by 1. Hence, we can use:

$$\sum_i \alpha_i \mathcal{T}_i \left( \frac{I + \eta A}{\|I + \eta A\|} \right) = P \left( \frac{I + \eta A}{\|I + \eta A\|} \right), \quad (13)$$

which allows similar step as (12) to hold, and further allows the dependence on  $\kappa$  for the CKS algorithm to be alleviated as summarized in Lemma 3.3; see also the illustration in Figure 2 (right).

### Approximating $x^* = A^{-1}b$ via PPA

For the PPA error (second pair of terms in (9)), the step size  $\eta$  needs to be properly set up so that a single step of PPA, i.e.,  $x_1 = (I + \eta A)^{-1}(x_0 + \eta b)$  is close enough to  $A^{-1}b = x^*$ . To achieve that, observe from (8):

$$\begin{aligned} \|x_{t+1} - x^*\| &\leq \|(I + \eta A)^{-t}\| \cdot \|x_0 - x^*\| \\ &\leq \frac{1}{(1 + \eta \sigma_{\min})^t} \cdot \|x_0 - x^*\|, \end{aligned} \quad (14)$$

where  $\sigma_{\min}$  is the smallest singular value of  $A$ . We desire the RHS to be less than  $\varepsilon_2$ . Denoting  $\|x_0 - x^*\| := d$  and using  $\sigma_{\min} = 1/\kappa$  (c.f., Definition 2.1), we can compute the lower bound on the number of iterations  $t$  as:

$$\frac{d}{(1 + \eta/\kappa)^t} \leq \varepsilon_2 \implies \frac{\log(\frac{d}{\varepsilon_2})}{\log(1 + \frac{\eta}{\kappa})} \leq t. \quad (15)$$

Based on the above analysis, we can compute the number of iterations  $t$  required to have  $\varepsilon$ -optimal solution of the quadratic problem in (6). Here, (15) is well defined in the sense that the lower bound on  $t$  is positive, as long as  $\eta > 0$ . In other words, if the lower bound of  $t$  in (15) is less than 1, that means PPA converges to  $\varepsilon_2$ -approximate solution in one step, with a proper step size.

Again, our goal is to achieve (9). Hence, we have to characterize how (14) results in the proximity in corresponding normalized quantum states. We utilize Lemma 4.3 below.

**Lemma 4.3.** *Let  $x$  and  $y$  be two vectors. Suppose  $\|x - y\| \leq \epsilon$  for some small positive scalar  $\epsilon$ . Then, the distance between the normalized vectors satisfies the following:*

$$\|x/\|x\| - y/\|y\|\| \leq \epsilon / \sqrt{\|x\| \cdot \|y\|}.$$

Now, we can characterize the error of the normalized quantum state of as an output of PPA.

**Proposition 4.4.** *Running the PPA in (7) for a single iteration with  $\eta = \kappa \left( \frac{d}{\varepsilon_2} - 1 \right)$ , where  $d := \|x_0 - A^{-1}b\|$ , results in the normalized quantum state satisfying:*

$$\left\| \left| \left( \frac{I + \eta A}{\|I + \eta A\|} \right)^{-1} (x_0 + \eta b) \right\rangle - |A^{-1}b\rangle \right\| \leq \frac{\varepsilon_2}{\Psi}, \quad (16)$$

where  $\Psi := \sqrt{\|(I + \eta A)^{-1}(x_0 + \eta b)\| \cdot \|A^{-1}b\|}$ .

### Overall Complexity and Improvement

Equipped with Propositions 4.1 and 4.4, we obtain the following main result.

**Theorem 4.5 (Main result).** Consider solving QLSP in Definition 2.1 with Algorithm 1, of which the approximation error  $\varepsilon$  can be decomposed as (9), recalled below:

$$\left\| \underbrace{|\psi_{I+\eta A, x_0+\eta b}\rangle - |(\frac{I+\eta A}{\|I+\eta A\|})^{-1}(x_0+\eta b)\rangle}_{\text{QLSP\_solver error} \leq \varepsilon_1 \text{ (c.f., Proposition 4.1)}} + \underbrace{|(\frac{I+\eta A}{\|I+\eta A\|})^{-1}(x_0+\eta b)\rangle - |A^{-1}b\rangle}_{\text{PPA error} \leq \varepsilon_2/\Psi \text{ (c.f., Proposition 4.4)}} \right\| \leq \varepsilon.$$

Suppose the existence of a QLSP\_solver satisfying the assumptions of Proposition 4.1 such that (10) is satisfied with  $\varepsilon_1 = \varepsilon/c$ , for  $c > 1$ . Further, suppose the assumptions of Proposition 4.4 hold, i.e., single-run PPA with  $\eta = \kappa \left(\frac{d}{\varepsilon_2} - 1\right)$  is implemented with accuracy  $\varepsilon_2 = \left(1 - \frac{1}{c}\right) \varepsilon \cdot \Psi$ . Then, the output of Algorithm 1 satisfies:

$$\left\| |\psi_{I+\eta A, x_0+\eta b}\rangle - |A^{-1}b\rangle \right\| \leq \varepsilon.$$

Moreover, the dependence on the condition number of QLSP\_solver changes from  $\kappa$  to  $\hat{\kappa} := \frac{\kappa(1+\eta)}{\kappa+\eta}$ .

We now further interpret the analysis from the previous subsection and compare it with other QLSP algorithms (e.g., CKS (Childs, Kothari, and Somma 2017)). We first recall the CKS query complexity below.

**Theorem 4.6** (CKS complexity (Childs, Kothari, and Somma 2017, Theorem 5)). The QLSP in Definition 2.1 can be solved to  $\varepsilon$ -accuracy by a quantum algorithm that makes  $\mathcal{O}(\kappa \cdot \text{poly} \log(\frac{\kappa}{\varepsilon}))$  queries to the oracles  $\mathcal{P}_A$  in Definition 2.3 and  $\mathcal{P}_B$  in Definition 2.2.

It was an open question whether there exists a quantum algorithm that matches the lower bound  $\Omega(\kappa \cdot \log(\frac{1}{\varepsilon}))$  (Harrow, Hassidim, and Lloyd 2009; Orsucci and Dunjko 2021). A recent quantum algorithm based on the discrete adiabatic theorem (Costa et al. 2022) was shown to (asymptotically) match this lower bound. We recall their result.

**Theorem 4.7** (Optimal complexity (Costa et al. 2022, Theorem 19)). The QLSP in Definition 2.1 can be solved to  $\varepsilon$ -accuracy by a quantum algorithm that makes  $\mathcal{O}(\kappa \cdot \log(1/\varepsilon))$  queries to the oracles  $\mathcal{U}_A$  in Definition 2.4 and  $\mathcal{P}_B$  in Definition 2.2.

A natural direction to utilize Algorithm 1 is to use the best QLSP\_solver (Costa et al. 2022), which has the query complexity  $\mathcal{O}(\kappa \cdot \log(1/\varepsilon))$  from Theorem 4.7. We summarize this in the next theorem.

**Theorem 4.8** (Improving the optimal complexity). Consider running Algorithm 1 with (Costa et al. 2022) as the candidate for QLSP\_solver, which has the original complexity of  $\mathcal{O}(\kappa \cdot \log(1/\varepsilon))$  (c.f., Theorem 4.7). The modified complexity of Algorithm 1 via Theorem 4.5 can be written and decomposed to:

$$\hat{\kappa} \cdot \log\left(\frac{c}{\varepsilon}\right) = \underbrace{\frac{\kappa(1+\eta)}{\kappa+\eta} \cdot \log\left(\frac{1}{\varepsilon}\right)}_{\text{Improvement}} + \underbrace{\frac{\kappa(1+\eta)}{\kappa+\eta} \cdot \log(c)}_{\text{Overhead}}, \quad (17)$$

where the ‘‘improvement’’ comes from  $\hat{\kappa} \leq \kappa$ , and the ‘‘overhead’’ is due to the weight of  $\varepsilon_1 = \varepsilon/c$  (and subsequently  $\varepsilon_2 = \left(1 - \frac{1}{c}\right) \varepsilon \cdot \Psi$ ) in Theorem 4.5. Further, with  $\eta = \kappa \left(\frac{d}{\varepsilon_2} - 1\right)$ , it follows

$$\hat{\kappa} = \frac{\kappa(1+\eta)}{\kappa+\eta} = \kappa - \frac{(c-1)(\kappa-1)\Psi\varepsilon}{c \cdot d} \leq \kappa. \quad (18)$$

We illustrate Theorem 4.8 in Figure 2 (left); thanks to the flexibility of Algorithm 1 and Theorem 4.5, similar analysis can be done with CKS (Childs, Kothari, and Somma 2017) as the baseline, as illustrated in Figure 2 (right).

Intuitively, based on the decomposition in (17), we can see that the constant  $c$ —which controls the weight of  $\varepsilon_1$  and  $\varepsilon_2$  in Theorem 4.5—enters a logarithmic term in the ‘‘overhead.’’ On the contrary, the (additive) improvement in  $\kappa$  is proportional to the term  $\frac{(c-1)}{c}$ , as can be seen in (18).

### Practical Improvement via Warm Start

The modified condition number  $\hat{\kappa}$  in (18) depends on the initial point of PPA,  $x_0$ , via  $d := \|x_0 - x^*\|$ . Therefore, a better initialization  $x_0$  via warm start can result in a bigger improvement in the overall complexity. We note again that such approach is *not possible* with any existing QLSP\_solver.

Let us provide a simple example. Suppose we initialize  $x_0$  such that  $d := \|x_0 - x^*\| = 2 \cdot \frac{\kappa-1}{\kappa} \cdot \varepsilon_2$ . This is possible since  $\varepsilon_2$  is the level of error achieved by the (classical) PPA with the specified step size  $\eta$  (c.f., Proposition 4.4). As  $\frac{\kappa-1}{\kappa} \approx 1$  for large  $\kappa$ , one can simply run a few iterations of classical optimization (e.g., gradient descent), and initialize  $x_0$  with the output that satisfies roughly twice the desired  $\varepsilon_2$ . Then, based on (18), we have

$$\hat{\kappa} = \kappa - \frac{(c-1)(\kappa-1)\Psi\varepsilon}{c \cdot d} \stackrel{d \leftarrow \frac{2(\kappa-1)\varepsilon_2}{\kappa}}{=} \frac{\kappa}{2}.$$

That is, simply by running a few steps of gradient descent classically such that  $\|x_{\text{GD}} - x^*\| \approx 2 \cdot \varepsilon_2$ , and initializing  $x_0 \leftarrow x_{\text{GD}}$  in Algorithm 1, the (asymptotically) optimal query complexity from Costa et al. (2022) can be halved.

We illustrate the practicality of warm start in Figure 1 (right), where we generate (normalized)  $A$  and  $b$  from  $\mathcal{N}(0, 1)$ . We vary the condition number of  $A$  to be  $\kappa \in \{100, 200, 300, 400, 500\}$ . We plot the overall query complexity of Algorithm 1 with warm start, where  $x_0$  is initialized with the last iterate of  $\{200, 500, 1000\}$  steps of gradient descent. The baseline (blue) is the optimal query complexity  $\Omega(\kappa \log \frac{1}{\varepsilon})$  (Costa et al. 2022), which can effectively be halved (red) via Algorithm 1 with warm start.

## 5 Conclusion

In this work, we proposed a novel quantum algorithm for solving the quantum linear systems problem (QLSP), based on the proximal point algorithm (PPA). Specifically, we showed that implementing a single-step PPA is possible by utilizing existing QLSP solvers. We designed a meta-algorithm where any QLSP solver can be utilized as a subroutine to improve the dependence on the condition number. Even the (asymptotically) optimal quantum algorithm (Costa et al. 2022) can be significantly accelerated via Algorithm 1, especially when the problem is ill-conditioned.

## References

- Ahn, K.; and Sra, S. 2022. Understanding Nesterov’s Acceleration via Proximal Point Method. In *Symposium on Simplicity in Algorithms (SOSA)*, 117–130. SIAM.
- Ambainis, A. 2012. Variable time amplitude amplification and quantum algorithms for linear algebra problems. In *STACS’12 (29th Symposium on Theoretical Aspects of Computer Science)*, volume 14, 636–647. LIPIcs.
- An, D.; and Lin, L. 2022. Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm. *ACM Transactions on Quantum Computing*, 3(2): 1–28.
- Bauschke, H.; and Combettes, P. 2019. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, corrected printing.
- Boyd, S. P.; and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Brassard, G.; and Hoyer, P. 1997. An exact quantum polynomial-time algorithm for Simon’s problem. In *Proceedings of the Fifth Israeli Symposium on Theory of Computing and Systems*, 12–23. IEEE.
- Brassard, G.; Hoyer, P.; Mosca, M.; and Tapp, A. 2002. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305: 53–74.
- Childs, A. M.; Kothari, R.; and Somma, R. D. 2017. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6): 1920–1950.
- Childs, A. M.; Maslov, D.; Nam, Y.; Ross, N. J.; and Su, Y. 2018. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38): 9456–9461.
- Costa, P. C.; An, D.; Sanders, Y. R.; Su, Y.; Babbush, R.; and Berry, D. W. 2022. Optimal scaling quantum linear-systems solver via discrete adiabatic theorem. *PRX Quantum*, 3(4): 040303.
- Feynman, R. P. 1982. Simulating Physics with Computers. *International Journal of Theoretical Physics*, 21(6/7).
- Gauss, C. F. 1877. *Theoria motus corporum coelestium in sectionibus conicis solem ambientium*, volume 7. FA Perthes.
- Gilyén, A.; Su, Y.; Low, G. H.; and Wiebe, N. 2019. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 193–204.
- Gribling, S.; Kerenidis, I.; and Szilágyi, D. 2021. Improving quantum linear system solvers via a gradient descent perspective. *arXiv preprint arXiv:2109.04248*.
- Grover, L. K. 1998. Quantum computers can search rapidly by using almost any transformation. *Physical Review Letters*, 80(19): 4329.
- Güler, O. 1991. On the convergence of the proximal point algorithm for convex minimization. *SIAM journal on control and optimization*, 29(2): 403–419.
- Güler, O. 1992. New proximal point algorithms for convex minimization. *SIAM Journal on Optimization*, 2(4): 649–664.
- Harrow, A. W.; Hassidim, A.; and Lloyd, S. 2009. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15): 150502.
- Hestenes, M. R.; and Stiefel, E. 1952. Methods of conjugate gradients for solving. *Journal of research of the National Bureau of Standards*, 49(6): 409.
- Higham, N. J. 2011. Gaussian elimination. *Wiley Interdisciplinary Reviews: Computational Statistics*, 3(3): 230–238.
- Kerenidis, I.; and Prakash, A. 2016. Quantum recommendation systems. *arXiv preprint arXiv:1603.08675*.
- Kim, J. L.; Toulis, P.; and Kyriallidis, A. 2022. Convergence and stability of the stochastic proximal point algorithm with momentum. In *Learning for Dynamics and Control Conference*, 1034–1047. PMLR.
- Kitaev, A. Y. 1995. Quantum measurements and the Abelian stabilizer problem. *arXiv preprint quant-ph/9511026*.
- Lin, H.; Mairal, J.; and Harchaoui, Z. 2015. A universal catalyst for first-order optimization. *Advances in neural information processing systems*, 28.
- Lin, L.; and Tong, Y. 2020. Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems. *Quantum*, 4: 361.
- Liu, J.-P.; Kolden, H. Ø.; Krovi, H. K.; Loureiro, N. F.; Trivisa, K.; and Childs, A. M. 2021. Efficient quantum algorithm for dissipative nonlinear differential equations. *Proceedings of the National Academy of Sciences*, 118(35): e2026805118.
- Lloyd, S. 1996. Universal quantum simulators. *Science*, 273(5278): 1073–1078.
- Low, G. H.; and Chuang, I. L. 2019. Hamiltonian simulation by qubitization. *Quantum*, 3: 163.
- Orsucci, D.; and Dunjko, V. 2021. On solving classes of positive-definite quantum linear systems with quadratically improved runtime in the condition number. *Quantum*, 5: 573.
- Papayan, V. 2020. Traces of class/cross-class structure pervade deep learning spectra. *The Journal of Machine Learning Research*, 21(1): 10197–10260.
- Parikh, N.; Boyd, S.; et al. 2014. Proximal algorithms. *Foundations and trends® in Optimization*, 1(3): 127–239.
- Prasad, T.; and Zhuang, A. 2022. Proving the BQP-Completeness of the Quantum Linear Systems Problem Using a Clock Construction. *Technical Report*.
- Rebentrost, P.; Mohseni, M.; and Lloyd, S. 2014. Quantum support vector machine for big data classification. *Physical review letters*, 113(13): 130503.
- Rockafellar, R. T. 1976. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5): 877–898.
- Schwarzenberg-Czerny, A. 1995. On matrix factorization and efficient least squares solution. *Astronomy and Astrophysics Supplement*, v. 110, p. 405, 110: 405.

Shabat, G.; Shmueli, Y.; Aizenbud, Y.; and Averbuch, A. 2018. Randomized LU decomposition. *Applied and Computational Harmonic Analysis*, 44(2): 246–272.

Subaşı, Y.; Somma, R. D.; and Orsucci, D. 2019. Quantum algorithms for systems of linear equations inspired by adiabatic quantum computing. *Physical review letters*, 122(6): 060504.

Toulis, P.; Airoldi, E.; and Rennie, J. 2014. Statistical analysis of stochastic gradient methods for generalized linear models. In *International Conference on Machine Learning*, 667–675. PMLR.

Toulis, P.; and Airoldi, E. M. 2017. Asymptotic and finite-sample properties of estimators based on stochastic gradients. *The Annals of Statistics*, 45(4): 1694–1727.

Wiebe, N.; Kapoor, A.; and Svore, K. 2014. Quantum algorithms for nearest-neighbor methods for supervised and unsupervised learning. *arXiv preprint arXiv:1401.2142*.