

An improved quantum-inspired algorithm for linear regression

András Gilyén¹, Zhao Song², and Ewin Tang³

¹Alfréd Rényi Institute of Mathematics

²Adobe Research

³University of Washington

We give a classical algorithm for linear regression analogous to the quantum matrix inversion algorithm [Harrow, Hassidim, and Lloyd, Physical Review Letters'09] for low-rank matrices [Wossnig, Zhao, and Prakash, Physical Review Letters'18], when the input matrix A is stored in a data structure applicable for QRAM-based state preparation.

Namely, suppose we are given an $A \in \mathbb{C}^{m \times n}$ with minimum non-zero singular value σ which supports certain efficient ℓ_2 -norm importance sampling queries, along with a $b \in \mathbb{C}^m$. Then, for some $x \in \mathbb{C}^n$ satisfying $\|x - A^+b\| \leq \varepsilon \|A^+b\|$, we can output a measurement of $|x\rangle$ in the computational basis and output an entry of x with classical algorithms that run in $\tilde{O}\left(\frac{\|A\|_F^6 \|A\|^6}{\sigma^{12} \varepsilon^4}\right)$ and $\tilde{O}\left(\frac{\|A\|_F^6 \|A\|^2}{\sigma^8 \varepsilon^4}\right)$ time, respectively. This improves on previous “quantum-inspired” algorithms in this line of research by at least a factor of $\frac{\|A\|_F^{16}}{\sigma^{16} \varepsilon^2}$ [Chia, Gilyén, Li, Lin, Tang, and Wang, STOC'20]. As a consequence, we show that quantum computers can achieve at most a factor-of-12 speedup for linear regression in this QRAM data structure setting and related settings. Our work applies techniques from sketching algorithms and optimization to the quantum-inspired literature. Unlike earlier works, this is a promising avenue that could lead to feasible implementations of classical regression in a quantum-inspired settings, for comparison against future quantum computers.

András Gilyén: gilyen@renyi.hu, Formerly at the Institute for Quantum Information and Matter, California Institute of Technology.

Zhao Song: zsong@adobe.com

Ewin Tang: ewint@cs.washington.edu

1 Introduction

An important question for the future of quantum computing is whether we can use quantum computers to speed up machine learning [1]. Answering this question is a topic of active research [2, 3, 4]. One potential avenue to an affirmative answer is through quantum linear algebra algorithms, which can perform linear regression [5] (and compute similar linear algebra expressions [6]) in time poly-logarithmic in input dimension, with some restrictive “quantum” assumptions [2, 3]. For machine learning tasks, a natural setting of these assumptions is that the input is low-rank, in which case the above quantum algorithms require that the input is given in a manner allowing efficient quantum state preparation. This state preparation assumption is typically satisfied via a data structure instantiated with quantum random access memory (QRAM)¹ [7, 8]. A recent paper shows that, by exploiting these assumptions, classical algorithms can also perform quantum linear algebra with only a polynomial slowdown [9], meaning that here quantum computers do not give the exponential speedup in dimension that one might hope for. However, these algorithms have running times with high exponents, leaving open the possibility that quantum computers admit large polynomial speedups for these problems. For example, given our current understanding, the quantum recommendation systems algorithm [10], with running time $\mathcal{O}^*(\frac{\|A\|_F}{\sigma})$, could still give a large polynomial quantum speedup, since the best-known classical algorithm has running time $\tilde{\mathcal{O}}^*(\frac{\|A\|_F^6 \|A\|^{16}}{\sigma^{24}})$ [9] (improving the original running time of $\tilde{\mathcal{O}}^*(\frac{\|A\|_F}{\sigma})^{24}$) [11].² It is an open question *whether any quantum linear algebra algorithm admits a large polynomial speedup*, compared to classical numerical linear algebra [1, 12, 13].

We focus on this open question for the problem of low-rank linear regression, where we are given a matrix $A \in \mathbb{C}^{m \times n}$ with minimum singular value σ and a vector $b \in \mathbb{C}^m$, and asked to approximate $x^* := \arg \min_x \frac{1}{2} \|Ax - b\|_2^2$. Linear regression has been an influential primitive for quantum machine learning (QML) [4] since the invention of Harrow, Hassidim, and Lloyd’s algorithm (HHL) originally introduced for sparse A [5, 14]. The low-rank variant that we consider commonly appears in the QML literature [15, 16]. These algorithms, like all the algorithms we consider, are called “low-rank” because their runtimes depend on $\frac{\|A\|_F^2}{\sigma^2} \geq \text{rank } A$, so in some sense, good performance requires the matrix A to be strictly low-rank. However, this restriction can be relaxed with regularization, which replaces σ^2 by $\sigma^2 + \lambda$, so that for reasonable choices of the regularization parameter λ , A only needs to be approximately low-rank. The current state-of-the-art quantum algorithm can produce a state ε -close to $|A^+b\rangle$ in ℓ_2 -norm in $\mathcal{O}^*(\frac{\|A\|_F}{\|A\|} \frac{\|A\|}{\sigma})$ time [17], given A and b in the aforementioned data structure. We think about this running time as depending polynomially on the (square root of) stable rank $\frac{\|A\|_F}{\|A\|}$ and the condition number $\frac{\|A\|}{\sigma}$. Note that being able to produce quantum states $|x\rangle := \frac{1}{\|x\|} \sum_i x_i |i\rangle$ corresponding to a desired vector x is akin to a classical sampling problem, and is different from outputting x itself. Though recent “quantum-inspired” classical algorithms [18] demonstrate that this quantum algorithm does not admit exponential speedups, the best previous classical algorithm of Chia, Gilyén, Li, Lin, Tang and Wang [9] runs in $\tilde{\mathcal{O}}^*(\frac{\|A\|_F}{\|A\|})^6 (\frac{\|A\|}{\sigma})^{28}$ time, leaving significant room for *polynomial* quantum speedups.³

¹In this paper, we always use QRAM in combination with a data structure used for efficiently preparing states $|0\rangle \rightarrow \sum_i \frac{x_i}{\|x\|} |i\rangle$ corresponding to vectors $x \in \mathbb{C}^n$.

²We define $\tilde{\mathcal{O}}(T)$ as $\mathcal{O}(T \cdot \text{polylog}(T))$, and define \mathcal{O}^* to be big \mathcal{O} notation, hiding polynomial dependence on ε and poly-logarithmic dependence on dimension.

³The ε dependence for the quantum algorithm is $\log \frac{1}{\varepsilon}$, compared to the classical algorithm which gets

Our main result tightens this gap, giving an algorithm running in $\mathcal{O}^*\left(\left(\frac{\|A\|_F}{\|A\|}\right)^6 \left(\frac{\|A\|}{\sigma}\right)^{12}\right)$ time. Roughly, this tightens the separation between quantum and classical from 1-to-28 to 1-to-12. As a bonus, this algorithm is a factor of $\frac{\|A\|_F^4}{\sigma^4}$ faster if we want to compute an entry of the output, rather than to sample an index. Our algorithm is a carefully-analyzed instance of stochastic gradient descent that exploits the sampling access provided in our setting, combined with a technical sketching step to enforce sparsity of the input vector b . Since it is an iterative algorithm, it is potentially more practical to implement and use as a benchmark against future scalable quantum computers [19]. Our result suggests that other tasks that can be solved through iterative methods, like primal-dual algorithms, may have quantum-classical gaps that are smaller than existing quantum-inspired results suggest.

1.1 Our model

We now introduce the input and output model that we and all prior quantum-inspired algorithms in this line of research use [20, 9]. The motivation for this model is to be a classical analogue to the input model of QML algorithms: many such algorithms assume that input is stored in a particular data structure in QRAM [15, 16, 8, 4, 1, 21], including current QML algorithms for low-rank linear regression (our focus) [17].⁴ So, our comparison classical algorithms must also assume that input is stored in this “quantum-inspired” data structure, which supports several fully-classical operations. We first define the quantum-inspired data structure for vectors (Definition 1.1), then for matrices (Definition 1.2).

Definition 1.1 (Vector-based data-structure, $\text{SQ}(v)$ and $\text{Q}(v)$). *For any vector $v \in \mathbb{C}^n$, let $\text{SQ}(v)$ denote a data-structure that supports the following operations:*

1. *SAMPLE()*, which outputs the entry i with probability $|v_i|^2/\|v\|^2$.
2. *QUERY(i)*, which takes $i \in [n]$ as input and outputs v_i .
3. *NORM()*, which outputs $\|v\|$.⁵

Let $\mathcal{T}(v)$ denote the max time it takes for the data structure to respond to any query. If we only allow the *QUERY* operation, the data-structure is called $\text{Q}(v)$.

Notice that the *SAMPLE* function is the classical analogue of quantum state preparation, since the distribution being sampled is identical to the one attained by measuring $|v\rangle$ in the computational basis. We now define the quantum-inspired data structure for matrices.

Definition 1.2 (Matrix-based data-structure, $\text{SQ}(A)$). *For any matrix $A \in \mathbb{C}^{m \times n}$, let $\text{SQ}(A)$ denote a data-structure that supports the following operations:*

1. *SAMPLE1()*, which outputs $i \in [m]$ with probability $\|A_{i,*}\|^2/\|A\|_F^2$.

¹ $\frac{1}{\varepsilon^6}$. While this suggests an exponential speedup in ε , it appears to only hold for sampling problems, such as measuring $|A^+b\rangle$ in the computational basis. Learning information from the output quantum state $|A^+b\rangle$ generally requires $\text{poly}(\frac{1}{\varepsilon})$ samples, preventing exponential separation for computational problems.

⁴Although current QRAM proposals suggest that quantum hardware implementing QRAM may be realizable with essentially only logarithmic overhead in the running time [7], an actual physical implementation would require substantial advances in quantum technology in order to maintain coherence for a long enough time [22].

⁵When we claim that we can call *NORM()* on output vectors, we will mean that we can output a constant approximation: a number in $[0.9\|v\|, 1.1\|v\|]$.

2. *SAMPLE2*(i), which takes a row index $i \in [m]$ as input and outputs the column index $j \in [n]$ with probability $|A_{i,j}|^2 / \|A_{i,*}\|^2$.
3. *QUERY*(i, j), which takes $i \in [m]$ and $j \in [n]$ as input and outputs $A_{i,j}$.
4. *NORM*(i), which takes $i \in [m]$ as input and outputs $\|A_{i,*}\|$.
5. *NORM*(\cdot), which outputs $\|A\|_F$.

Let $\mathcal{T}(A)$ denote the max time the data structure takes to respond to any query.

For the sake of simplicity, we will assume all input SQ data structures respond to queries in $\mathcal{O}(1)$ time. There are data structures that can do this in the word RAM model [9, Remark 2.15]. The dynamic data structures that commonly appear in the QML literature [10] respond to queries in $\mathcal{O}(\log(mn))$ time, so using such versions only increases our running time by a logarithmic factor.

The specific choice of data structure does not appear to be important to the quantum-classical separation, though. The QRAM and data structure assumptions of typical QML algorithms can be replaced with any *state preparation assumption*, which is an assumption implying that quantum states corresponding to input data can be prepared in time polylogarithmic in dimension. However, to the best of our knowledge, all models admitting efficient protocols that take v stored in the standard way as input and output the quantum state $|v\rangle$ also admit corresponding efficient classical sample and query operations that can replace the data structure described above [9, Remark 2.15]. In other words, classical algorithms in the sample and query access model can be run whenever the corresponding QML algorithm can, assuming that the input data is classical. Consequently, our results on the quantum speedup for low-rank matrix inversion appear robust to changing quantum input models.

Notations. For a vector $v \in \mathbb{C}^n$, $\|v\|$ denotes ℓ_2 norm. For a matrix $A \in \mathbb{C}^{m \times n}$, A^\dagger , A^+ , $\|A\|$, and $\|A\|_F$ denote the conjugate transpose, pseudoinverse, operator norm, and Frobenius norm of A , respectively. We use $A_{i,j}$ to denote the entry of A at the i -th row and j -th column. We use $A_{i,*}$ and $A_{*,j}$ to denote the i -th row and j -th column of A .

We use $\mathbb{E}[\cdot]$ and $\mathbb{V}[\cdot]$ to denote expectation and variance of a random variable. Abusing notation, for a random vector v , we denote $\mathbb{V}[v] = \mathbb{E}[\|v - \mathbb{E}[v]\|^2]$. For a differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, ∇f denotes the gradient of f . We say f is convex if, for all $x, y \in \mathbb{R}^n$ and $t \in [0, 1]$, $f(tx + (1-t)y) \leq tf(x) + (1-t)f(y)$.

We use the shorthand $g \lesssim h$ for $g = \mathcal{O}(h)$, and respectively $g \gtrsim h$ and $g \approx h$ for $g = \Omega(h)$ and $g = \Theta(h)$. $\tilde{\mathcal{O}}(T)$ denotes $\mathcal{O}(T \cdot \text{polylog}(T))$ and \mathcal{O}^* denotes big \mathcal{O} notation that hides polynomial dependence on ε and poly-logarithmic dependence on dimension. Note that the precise exponent on $\log(mn)$ depends on the choice of word size in our RAM models. Quantum and classical algorithms have different conventions for this choice, hence we use notation that elides this detail for simplicity.

1.2 Our results

In this paper, we focus on solving the following problem. Prior work in this line of research has studied this problem in the $\lambda = 0$ case [9].

Problem (Regression with regularization). Given $A \in \mathbb{C}^{m \times n}$, $b \in \mathbb{C}^m$, and a regularization parameter $\lambda \geq 0$, we define the function $f : \mathbb{C}^n \rightarrow \mathbb{R}$ as

$$f(x) := \frac{1}{2}(\|Ax - b\|^2 + \lambda\|x\|^2).$$

Let⁶ $x^* := \arg \min_{x \in \mathbb{C}^n} f(x) = (A^\dagger A + \lambda I)^+ A^\dagger b$.

We will manipulate vectors by manipulating their *sparse descriptions*, defined as follows:

Definition 1.3 (sparse description). *We say we have an s -sparse description of $x \in \mathbb{C}^d$ if we have an s -sparse $v \in \mathbb{C}^n$ such that $x = A^\dagger v$. We use the convention that any s -sparse description is also a t -sparse description for all $t \geq s$.*

Our main result is that we can solve regression efficiently, assuming $\text{SQ}(A)$. Algorithm 1 and the corresponding Theorem 1.4 directly improves the previous best result [9] due to Chia, Gilyén, Li, Lin, Tang and Wang by a factor of $\frac{\|A\|_{\text{F}}^{20}}{\sigma^{20}\varepsilon^2}$.

Algorithm 1 Quantum-inspired regression via stochastic gradient descent

Input: $\text{SQ}(A)$, $\text{Q}(b)$, ε , and $\lambda = \mathcal{O}(\|A\|_{\text{F}}\|A\|)$.

- 1: **Init:** Sparsify b as described in Lemma 2.7 to obtain an $s = 800 \frac{\|A\|_{\text{F}}^2 \|b\|^2}{(\sigma^2 + \lambda)^2 \varepsilon^2 \|x^*\|^2}$ -sparse \hat{b} .
- 2: Set $v^{(0)} := 0$ (and implicitly define $x^{(t)} = A^\dagger v^{(t)}$).
- 3: Set $\eta := \frac{\varepsilon^2(\sigma^2 + \lambda)}{32\|A\|_{\text{F}}^2\|A\|^2 + 16\lambda^2}$ and $T := \frac{\ln(8/\varepsilon^2)}{\eta(\sigma^2 + \lambda)} = 32 \ln\left(\frac{\sqrt{8}}{\varepsilon}\right) \frac{2\|A\|_{\text{F}}^2\|A\|^2 + \lambda^2}{\varepsilon^2(\sigma^2 + \lambda)^2}$.
- 4: **for** $t = 0, 1, \dots, T - 1$ **do**
- 5: Sample a row index r according to the row norms $\frac{\|A_{r,*}\|^2}{\|A\|_{\text{F}}^2}$.
- 6: Sample $C := \frac{\|A\|_{\text{F}}^2}{\|A\|^2}$ column indices $\{c_i\}_{i \in [C]}$ i.i.d. according to $\frac{|A_{r,c}|^2}{\|A_{r,*}\|^2}$.
- 7: Define $v^{(t+1)}$ as follows:

$$v^{(t+1)} := (1 - \eta\lambda)v^{(t)} + \eta\hat{b} - \eta \frac{\|A\|_{\text{F}}^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v^{(t)} \right) e_r. \quad (1)$$

8: **end for**

- 9: **Output:** $v^{(T)}$ which is $32 \frac{2\|A\|_{\text{F}}^2\|A\|^2 + \lambda^2}{(\sigma^2 + \lambda)^2 \varepsilon^2} \left(\frac{25\|b\|^2}{2\|A\|^2\|x^*\|^2} + \ln\left(\frac{\sqrt{8}}{\varepsilon}\right) \right)$ -sparse; set $x := A^\dagger v^{(T)}$.

Queries to x : $x_j = \sum A_{j,i}^\dagger v_i^{(T)}$, so query $A_{i,j}$ for all non-zero $v_i^{(T)}$ and compute the sum.

Sampling from $|x_j|^2/\|x\|^2$: Perform rejection sampling according to Lemma 2.5.

Theorem 1.4 (Main result). *Suppose we are given $\text{SQ}(A) \in \mathbb{C}^{m \times n}$, $\text{Q}(b) \in \mathbb{C}^n$. Denote $\sigma := \|A^+\|^{-1}$ and consider $f(x)$ for $\lambda \lesssim \|A\|_{\text{F}}\|A\|$. Algorithm 1 runs in*

$$\mathcal{O}\left(\frac{\|A\|_{\text{F}}^6\|A\|^2}{(\sigma^2 + \lambda)^4 \varepsilon^4} \left(\frac{\|b\|^2}{\|A\|^2\|x^*\|^2} + \log \frac{1}{\varepsilon} \right) \log \frac{1}{\varepsilon}\right)$$

time and outputs an $\mathcal{O}\left(\frac{\|A\|_{\text{F}}^2\|A\|^2}{(\sigma^2 + \lambda)^2 \varepsilon^2} \left(\frac{\|b\|^2}{\|A\|^2\|x^\|^2} + \log \frac{1}{\varepsilon} \right)\right)$ -sparse description of an x such that $\|x - x^*\| \leq \varepsilon\|x^*\|$ with probability ≥ 0.9 . This description admits $\text{SQ}(x)$ for*

$$\mathcal{T}(x) = \mathcal{O}\left(\frac{\|A\|_{\text{F}}^6\|A\|^6}{(\sigma^2 + \lambda)^6 \varepsilon^4} \log^2 \frac{1}{\varepsilon} \left(\frac{\|b\|^4}{\|A\|^4\|x^*\|^4} + \log^2 \frac{1}{\varepsilon} \right) \left(\frac{\|b\|^2}{\|A\|_{\text{F}}^2\|x^*\|^2} + 1 \right)\right).$$

⁶Note that for $\lambda = 0$ the function f might not be strictly convex and so the minimizer is not necessarily unique. Nevertheless, on the image of A^\dagger , f is strictly convex. Since we will search a solution within the image of A^\dagger , we implicitly restrict f to this subspace, thus we can assume without loss of generality that f is strictly convex.

Note that the running time stated for $\mathcal{T}(x)$ corresponds to the running time of SAMPLE and NORM; the running time of QUERY is $\mathcal{O}\left(\frac{\|A\|_{\mathbb{F}}^2 \|A\|^2}{(\sigma^2 + \lambda)^2 \varepsilon^2} \left(\frac{\|b\|^2}{\|A\|^2 \|x^*\|^2} + \log \frac{1}{\varepsilon}\right)\right)$, the sparsity of the description. So, unlike previous quantum-inspired algorithms, it may be the case that the running time of the SQ query dominates, though it's conceivable that this running time bound is an artifact of our analysis.

Factors of $\frac{\|b\|^2}{\|A\|^2 \|x^*\|^2}$ (which is always at least one) arise because sampling error is additive with respect to $\|b\|$, and need to be rescaled relative to $\|x^*\|$; the quantum algorithm must also pay such a factor. The expression $\frac{\|A\|^2 \|x^*\|^2}{\|b\|^2}$ can be thought of as the fraction of b 's "mass" that is in the well-conditioned row space of A . To make this rigorous, we define the following *thresholded projector*.

Definition 1.5. For an $A \in \mathbb{C}^{m \times n}$ with singular value decomposition $A = \sum \sigma_i u_i v_i^\dagger$ and $\lambda \geq 0$, we define $\Pi_{A,\lambda} = \sum p_{A,\lambda}(\sigma_i) u_i u_i^\dagger$ where

$$p_{A,\lambda}(\sigma) = \begin{cases} 0 & \sigma = 0 \\ \frac{2\sigma\sqrt{\lambda}}{\sigma^2 + \lambda} & 0 < \sigma \leq \sqrt{\lambda} \\ 1 & \sqrt{\lambda} < \sigma. \end{cases}$$

For intuition, $\Pi_{A,0}$ projects onto the row space of A , and for nonzero λ smoothly projects away the u_i 's with singular values σ_i smaller than $\sqrt{\lambda}$ roughly linearly, since $\frac{\sigma}{\sqrt{\lambda}} \leq p_{A,\lambda}(\sigma) \leq \frac{2}{\sqrt{\lambda}}$ for $\sigma \in [0, \sqrt{\lambda}]$. We use the definition presented here because $\|\Pi_{A,\lambda} b\|$ gives a natural lower bound for $\|x^*\|$: by Fact 2.8, $\frac{\|A\|^2 \|x^*\|^2}{\|b\|^2} \geq \frac{\|\Pi_{A,\lambda} b\|^2}{2\|b\|^2}$ when $\lambda \leq \|A\|^2$. This is the "fraction of mass" type quantity that one would expect.⁷

As mentioned previously, we use stochastic gradient descent to solve this problem, for $T := \mathcal{O}\left(\frac{\|A\|_{\mathbb{F}}^2 \|A\|^2}{(\sigma^2 + \lambda)^2 \varepsilon^2} \log \frac{1}{\varepsilon}\right)$ iterations. Such optimization algorithms are standard for solving regression in this setting [23, 24]: the idea is that, instead of explicitly computing the closed form of the minimizer to f , which requires computing a matrix pseudoinverse, one can start at $x^{(0)} = \vec{0}$ and iteratively find $x^{(t+1)}$ from $x^{(t)}$ such that the sequence $\{x^{(t)}\}_{t \in \mathbb{N}}$ converges to a local minimum of f . Updating $x^{(t)}$ by nudging it in the direction of f 's gradient produces such a sequence, even when we only use decent (stochastic) estimators of the gradient [25]. Because f is convex, it only has one minimizer, so $x^{(t)}$ converges to the global minimum.⁶

Challenges. Though the idea of SGD for regression is simple and well-understood, we note that *some standard analyses fail in our setting, so some care is needed in analysing SGD correctly.*

- First, our stochastic gradient (Definition 2.1) has variance that depends on the norm of the current iterate, and projection onto the ball of vectors of bounded norm is too costly, so our analyses cannot assume a uniform bound on the second moment of our stochastic gradient.
- Second, we want a bound for the final iterate $x^{(T)}$, not merely a bound on averaged iterates as is common for SGD analyses, since using an averaging scheme would complicate the analysis in Section 2.2.

⁷Note that this also means $\frac{\|A\|^2 \|x^*\|^2}{\|b\|^2} \geq \frac{\|Mb\|^2}{2\|b\|^2}$ for any matrix M satisfying $M \preceq \Pi_{A,\lambda}$. In particular, it holds when M projects onto only those u_i 's such that $\sigma_i \geq \sqrt{\lambda}$, i.e., when M is a true projector onto the well-conditioned subspace of the row space of A .

- Third, we want a bound on the output x of our algorithm of the form $\|x - x^*\| \leq \varepsilon \|x^*\|$. We could have chosen to aim for the typical bound in the optimization literature of $f(x) - f(x^*) \leq \varepsilon (f(x^{(0)}) - f(x^*))$, but our choice is common in the QML literature, and it only makes sense to have $\text{SQ}(x)$ when x is indeed close to x^* up to multiplicative error.
- Finally, note that acceleration via methods like the accelerated proximal point algorithm (APPA) [26] framework cannot be used in this setting, since we cannot pay the linear time in dimension necessary to recenter gradients.

The SGD analysis of Bach and Moulines [27] meets all of our criteria, which inspired the error analysis of this paper, while we use quantum-inspired ideas for the computation of the stochastic gradients (Definition 2.1).

The running time of SGD is only good for our choice of gradient when b is sparse, but through a sketching matrix we can reduce to the case where b has $\mathcal{O}\left(\frac{\|A\|_{\mathbb{F}}^2 \|A\|^2}{\sigma^4 \varepsilon^2} \frac{\|b\|^2}{\|AA^\dagger + b\|^2}\right)$ non-zero entries. Additional analysis is necessary to argue that we have efficient sample and query access to the output, since we need to rule out the possibility that our output x is given as a description $A^\dagger \tilde{x}$ such that \tilde{x} is much larger than x , in which case computing, say, $\|x\|$ via rejection sampling would be intractable.

Our analysis of the stochastic gradient descent is completely self-contained and the analysis of the stochastic gradients only relies on a few standard results about quantum-inspired sampling techniques, so the paper shall be accessible even without having a background on convex optimization.

1.3 Discussion

Comparisons to prior work. First, we describe the prior work on quantum-inspired classical low-rank linear regression algorithms in more detail. There are three prior papers here: [28] achieves a running time of $\tilde{\mathcal{O}}\left(\frac{\|A\|_{\mathbb{F}}^6 k^6 \|A\|^{16}}{\sigma^{22} \varepsilon^6}\right)$ (where $k \leq \frac{\|A\|_{\mathbb{F}}^2}{\sigma^2}$ is the rank of A), [29] achieves an incomparable running time of $\mathcal{O}\left(\frac{\|A\|_{\mathbb{F}}^6 \|A\|^{22}}{\sigma^{28} \varepsilon^6}\right)$, and [9] achieves the same running time as [29] for a more general problem (where σ can be chosen to be a threshold, and isn't necessarily the minimum singular value of A). For simplicity of exposition, we only compared our algorithm to the latter running time. However, our algorithm improves on all previous running times: an improvement of $k^6 \frac{\|A\|^{10}}{\sigma^{10} \varepsilon^2}$ on the former, and $\frac{\|A\|^{16}}{\sigma^{16} \varepsilon^2}$ on the latter.

Comparison to [23] In some sense, this work gives a version of the regression algorithm of Gupta and Sidford [23] (which, alongside [24], is the current state-of-the-art iterative algorithm for regression in many numerically sparse settings) that trades running time for weaker, more quantum-like input assumptions. Their algorithm takes A and b , with only query access to the input's nonzero entries, and explicitly outputs a vector x satisfying $\|x - x^*\| \leq \varepsilon \|x^*\|$ in $\tilde{\mathcal{O}}^*(\text{nnz}(A) + \frac{\|A\|_{\mathbb{F}}}{\sigma} \text{nnz}(A)^{\frac{2}{3}} n^{\frac{1}{6}})$ time (shown here after naively bounding numerical sparsity by n). They use a stochastic first-order iterative method, that we adapt by mildly weakening the stochastic gradient to be efficient assuming the quantum-inspired data structure (instead of the bespoke data structure they design for their algorithm).

To elaborate, we make a direct comparison. Suppose we wish to apply our algorithm from Theorem 1.4 to Gupta and Sidford's setting, meaning that we are given A and b as lists of non-zero entries, and wish to output the vector x in full satisfying $\|x - x^*\| \leq \varepsilon \|x^*\|$. Because of the difference in setting, we need additional $\mathcal{O}(\text{nnz}(A))$ time to set

up the data structures to get $\text{SQ}(A)$, and we also need $\mathcal{O}^*(\frac{\|A\|_{\mathbb{F}}^2\|A\|^2}{\sigma^4}n)$ time to query for all the entries of x , given the $\text{SQ}(x)$ from the output of Algorithm 1. In total, this takes $\mathcal{O}(\text{nnz}(A)) + \mathcal{O}^*(\frac{\|A\|_{\mathbb{F}}^6\|A\|^2}{\sigma^8} + \frac{\|A\|_{\mathbb{F}}^2\|A\|^2}{\sigma^4}n)$ time. This is worse than Gupta and Sidford’s time, even after accounting for the fact that, unlike in the setting we consider, we can apply the acceleration framework in [26]. Acceleration speeds up our algorithm and reduces our error dependence to poly-logarithmic in $\frac{1}{\varepsilon}$, at the cost of $\text{nnz}(A)$ time per “phase” throughout the logarithmically many phases. The quantum algorithm performs comparably, since one needs to run it $\mathcal{O}^*(n)$ times so that state tomography can convert copies of $|x\rangle$ to an explicit list of x ’s entries, resulting in a total running time of $\mathcal{O}(\text{nnz}(A)) + \tilde{\mathcal{O}}^*(\frac{\|A\|_{\mathbb{F}}}{\sigma}n)$.

However, our algorithm performs better than Gupta and Sidford’s algorithm in a weaker setting. QML algorithms often don’t take the form specified above, where the goal is to output a full vector and we assume nothing about the input. For example, the quantum recommendation system [10] only outputs a measurement from the output vector, and works in a setting (providing online recommendations, in essence a dynamic data structure problem) where the $\text{nnz}(A)$ cost for building a data structure can be amortized over many runs of their algorithm. This is the setting we design our algorithm for, where we are given $\text{SQ}(A)$ directly and only wish to know a measurement from the output vector via $\text{SQ}(x)$, and we achieve a runtime independent of dimension via Algorithm 1. Gupta and Sidford’s algorithm (even without acceleration) does not behave well in this setting, and would have a running time polynomial in input size, thus being exponentially slower than the quantum-inspired (and quantum) algorithm. Our goal with this result is to advance the complexity theoretic understanding of quantum speedups by *limiting the possibility of quantum speedup in the most general setting possible*. To this end, we assume fairly little about our input, so our classical algorithms can work for many different applications where one might hope for quantum speedups. This comes at the cost of increased runtime, including a polynomial dependence on ε^{-1} which, being exponentially larger than in other classical algorithms, potentially reduces the applicability of our result as an algorithm. On the other hand, a similar drawback is also present in the related quantum algorithms, so the resulting runtimes can be at least more directly compared.

Finally, we note that sketching algorithms could also have been used to improve upon previous quantum-inspired results. In particular, one can view the ℓ_2 -type samples that our input data structure supports as $\frac{\|A\|_{\mathbb{F}}^2}{k\sigma^2}$ -oversampled leverage score samples, enabling the application of typical sketching results [30] associated with approximating matrix products and approximate subspace embeddings. However, we were not quite able to find an algorithm with running time as good as Theorem 1.4 through the sketching literature, though something similar might be possible (say, the same running time but with the alternate guarantee that $\|Ax - b\| \leq (1 + \varepsilon)\|Ax^* - b\|$). Our guarantee that $\|x - x^*\| \leq \varepsilon\|x^*\|$ is perhaps more natural in this quantum-inspired setting, since we need x to be bounded away from zero to enable efficient (rejection) sampling. Sketching may be more beneficial in the slightly less restrictive setting where one wishes to find *any* dynamic data structure that can store A, b and output measurements from $|\approx A^+b\rangle$. In this setting, one wouldn’t be required to use the QRAM data structure, and could build alternative data structures that could exploit, for example, oblivious sketches.⁸

Comparisons to concurrent and subsequent work. Around the same time as this work, Chepurko, Clarkson, Horesh, and Woodruff described algorithms for linear regression

⁸To our knowledge, prior versions of [31] produced algorithms in this setting, but the current version only uses the standard QRAM data structure.

and “quantum-inspired” recommendation systems, heavily using technical tools from the sketching literature [31]. Their algorithm gives a roughly $\frac{\|A\|_{\mathbb{F}}^2}{\|A\|^2}$ improvement over our algorithm, ignoring minor details involving problem-specific parameters.⁹ This is done using an algorithm similar to that of [28], but with an improved analysis and an application of conjugate gradient to solve a system of linear equations (instead of inverting the matrix directly).

Subsequent work by Shao and Montanaro [32] gives an algorithm improving on this result by a factor of $\frac{\|A\|^4}{\sigma^4 \varepsilon^2}$ when $\lambda = 0$ and $Ax^* = b$ exactly. Their work uses the randomized Kaczmarz method [33], which is a version of stochastic gradient descent [34]. If one changes our stochastic gradient ∇g to the randomized Kaczmarz update used in [32], replacing $A^\dagger b$ with the sketch $\frac{\|A\|_{\mathbb{F}}^2}{\|A_{r,*}\|^2} A_{r,*}^\dagger b_r$, (i.e., replacing \hat{b} in Equation (1) by $\frac{\|A\|_{\mathbb{F}}^2}{\|A_{r,*}\|^2} b_r e_r$) then our analysis recovers a similar result for the special case when $\lambda = 0$ and $Ax^* = b$.¹⁰

Outlook. Our result and the aforementioned follow-up work show that simple iterative algorithms like stochastic gradient descent can be nicely combined with quantum-inspired linear algebra techniques, resulting in quantum-inspired algorithms that compare favorably to direct approaches based on finding approximate singular value decompositions [9]. However, one needs to be careful, as more involved iterative steps require more involved quantum-inspired subroutines that can rapidly blow up the overall complexity, as occurs in the quantum-inspired SDP solver of [9], for example. We leave it to future work to study which other iterative approaches can be similarly advantageously combined with quantum-inspired linear algebra techniques.

2 Proofs

For simplicity we assume knowledge of $\|A\|$, σ , and later, $\|x^*\|$, exactly, despite that our $\text{SQ}(A)$ only gives us access to $\|A\|_{\mathbb{F}}$. Just an upper bound on $\|A\|$ and lower bounds on σ and $\|x^*\|$ respectively suffice, giving a running time bound by replacing these quantities in our complexity bounds by any respective upper or lower bounds. This holds because these quantities are only used to choose the internal parameters s, η^{-1}, T , and C , and replacing these quantities by their respective bounds only increase these internal parameters, only improving the resulting error bounds of Algorithm 1.

The algorithm we use to solve regression is stochastic gradient descent (SGD). Suppose we wish to minimize a convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Consider the following recursion, starting from $x^{(0)} \in \mathbb{R}^n$, with a random function $\nabla g : \mathbb{R}^n \rightarrow \mathbb{R}$ and a deterministic sequence of positive scalars $(\eta_t)_{k \geq 1}$.

$$x^{(t)} = x^{(t-1)} - \eta_t \nabla g(x^{(t-1)}) \quad (2)$$

⁹This comparison is between [31, Theorem 18], which, upon taking $\lambda = 0$ and $d' = 1$, gives a runtime of $\tilde{O}(\varepsilon^{-4} \|A\|_{\mathbb{F}}^4 \|A\|^4 \sigma^{-8} \log(d))$ to return their output description, and our Theorem 1.4, which gives a runtime of $\tilde{O}(\varepsilon^{-4} \|A\|_{\mathbb{F}}^6 \|A\|^2 \sigma^{-8})$ to return our description of x .

¹⁰Specifically, the bottleneck of our algorithm is the large bound on the residual $\mathbb{E}[\|\nabla g(x^*)\|^2]$, since this is the piece of the error that does not decrease exponentially. If $\lambda = 0$ and $Ax^* = b$, then making the described change leads to Ax being replaced by $Ax - b$ in the variance bound of Lemma 2.2, which is then zero for $x = x^*$. So, unlike in the original analysis, increasing C beyond $\frac{\|A\|_{\mathbb{F}}^2}{\|A\|^2}$ actually helps to reduce this residual. For example, increasing C by a factor of $\frac{1}{\varepsilon^2}$ reduces the number of iterations by a factor of $\frac{1}{\varepsilon^2}$, decreasing the runtime (which is linear in C and quadratic in iteration count) by a factor of $\frac{1}{\varepsilon^2}$. We point the reader to [34] for further exploration of improving the residual in SGD and randomized Kaczmarz.

The idea behind SGD is that if ∇g is a good stochastic approximation to ∇f , then $x^{(t)}$ will converge to the minimizer of f .

In our case, as defined in the introduction, $f(x) = \frac{1}{2}(\|Ax - b\|^2 + \lambda\|x\|^2)$, so we need a stochastic approximation to $\nabla f(x) = A^\dagger Ax - A^\dagger b + \lambda x$. Our choice of stochastic gradient comes from observing that $A^\dagger Ax = \sum_{r=1}^m \sum_{c=1}^n (A_{r,*})^\dagger A_{r,c} x_c$, so we can estimate this by sampling some of the summands.

Definition 2.1. We define $\nabla g(x)$ to be the random function resulting from the following process. Draw $r \in [m]$ from the distribution that is r with probability $\frac{\|A_{r,*}\|_{\mathbb{F}}^2}{\|A\|_{\mathbb{F}}^2}$, and then draw c_1, \dots, c_C i.i.d. from the distribution that is c with probability $\frac{|A_{r,c}|^2}{\|A_{r,*}\|_{\mathbb{F}}^2}$. For this choice of r and c_1, \dots, c_C , take

$$\nabla g(x) = \frac{\|A\|_{\mathbb{F}}^2}{\|A_{r,*}\|_{\mathbb{F}}^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|_{\mathbb{F}}^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j} \right) (A_{r,*})^\dagger - A^\dagger b + \lambda x. \quad (3)$$

Notice that the first term of this expression is the average of copies of the random vector

$$\frac{\|A\|_{\mathbb{F}}^2}{|A_{r,c}|^2} A_{r,c} x_c (A_{r,*})^\dagger \quad \text{with probability } \frac{|A_{r,c}|^2}{\|A\|_{\mathbb{F}}^2}.$$

Algorithm 1 is simply running the iteration Eq. (2) with $x^{(0)} = \vec{0}$, ∇g as defined in Definition 2.1, and $\eta_t := \frac{\varepsilon^2(\sigma^2 + \lambda)}{32\|A\|_{\mathbb{F}}^2\|A\|^2 + 16\lambda^2}$ (as we will see, this comes from applying Proposition 2.3 with $\varepsilon' = \varepsilon/2$). In Section 2.1, we prove that $x^{(T)}$ satisfies the desired error bound for $T := \Theta\left(\frac{\|A\|_{\mathbb{F}}^2\|A\|^2 + \lambda^2}{(\sigma^2 + \lambda)^2\varepsilon^2} \log \frac{1}{\varepsilon}\right)$. In Section 2.2, we prove that computing $x^{(T)}$ and simulating $\text{SQ}(x^{(T)})$ can be done in the desired running time, assuming that b is sparse. Finally, in Section 2.3, we show how to generalize our result to non-sparse b .

2.1 Error analysis of stochastic gradient descent

We now list the properties of the stochastic gradient from Definition 2.1: the proofs are straightforward computation and so are deferred to the appendix.

Lemma 2.2. For fixed $x, y \in \mathbb{C}^n$ and the random function $\nabla g(\cdot)$ defined in Definition 2.1, the following properties hold ($\nabla g(x)$ and $\nabla g(y)$ use the same instance of the random function):

- Part 1.** $\mathbb{E}[\nabla g(x)] = \nabla f(x) = A^\dagger Ax - A^\dagger b + \lambda x$
- Part 2.** $\mathbb{V}[\nabla g(x)] = \frac{1}{C} \|A\|_{\mathbb{F}}^4 \|x\|^2 + \left(1 - \frac{1}{C}\right) \|A\|_{\mathbb{F}}^2 \|Ax\|^2 - \|A^\dagger Ax\|^2$
- Part 3.** $\mathbb{E}[\|\nabla g(x) - \nabla g(y)\|_2^2] = \|(A^\dagger A + \lambda I)(x - y)\|^2 + \mathbb{V}[\nabla g(x - y)]$

We take $C = \|A\|_{\mathbb{F}}^2 / \|A\|^2$, which is the largest value we can set C to before it stops having an effect on the variance of ∇g . In Section 2.2, we will see that it's good for SGD's running time to take C to be as large as possible. In this setting, and more generally if $C \leq \frac{\|A\|_{\mathbb{F}}^2 \|x\|^2}{\|Ax\|^2}$ the variance in Part 2 can be bounded as $\frac{2}{C} \|A\|_{\mathbb{F}}^4 \|x\|^2$.

Now, we show that performing SGD for T iterations gives an x sufficiently close to the optimal vector x^* .

Proposition 2.3. Consider a matrix $A \in \mathbb{C}^{m \times n}$, a vector $b \in \mathbb{C}^m$, a regularization parameter $\lambda \geq 0$, and an error parameter $\varepsilon \in (0, 1]$. Denote $\sigma := \|A^+\|^{-1}$. Let $x^{(T)}$ be defined

as Eq. (2), with $x^{(0)} = \vec{0}$, $\eta_t := \eta := \frac{\varepsilon^2(\sigma^2 + \lambda)}{8\|A\|_{\mathbb{F}}^2\|A\|^2 + 4\lambda^2}$, and ∇g as defined in Definition 2.1 with $C := \frac{\|A\|_{\mathbb{F}}^2}{\|A\|^2}$. Then for $T := \frac{\ln(2/\varepsilon^2)}{\eta(\sigma^2 + \lambda)} = 8 \ln\left(\frac{\sqrt{2}}{\varepsilon}\right) \frac{2\|A\|_{\mathbb{F}}^2\|A\|^2 + \lambda^2}{\varepsilon^2(\sigma^2 + \lambda)^2}$,

$$\mathbb{E}[\|x^{(T)} - x^*\|^2] \leq \varepsilon^2 \|x^*\|^2.$$

In particular, by Chebyshev's inequality, with probability ≥ 0.96 , $\|x^{(T)} - x^*\| \leq 5\varepsilon \|x^*\|$.

Proof. Stochastic gradient descent is known to require a number of iterations linear in (something like the) second moment of the stochastic gradient. To analyze SGD, we loosely follow a strategy used by Bach and Moulines [27]. First, note that

$$\begin{aligned} & \mathbb{E}[\|\nabla g(x)\|^2] \\ & \leq \mathbb{E}[2\|\nabla g(x) - \nabla g(x^*)\|^2 + 2\|\nabla g(x^*)\|^2] \quad \text{since } \|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2 \\ & = 2\left(\|(A^\dagger A + \lambda I)(x - x^*)\|^2 + \mathbb{V}[\nabla g(x - x^*)] + \mathbb{V}[\nabla g(x^*)]\right) \quad \text{since } \nabla f(x^*) = \vec{0} \\ & \leq 2\left(\|(A^\dagger A + \lambda I)(x - x^*)\|^2 + 2\|A\|_{\mathbb{F}}^2\|A\|^2(\|x - x^*\|^2 + \|x^*\|^2)\right) \quad \text{by Lemma 2.2} \\ & \leq 2(\|A^\dagger A + \lambda I\|^2 + 2\|A\|_{\mathbb{F}}^2\|A\|^2)\|x - x^*\|^2 + 4\|A\|_{\mathbb{F}}^2\|A\|^2\|x^*\|^2. \end{aligned}$$

Next, we use the well-known fact that $f(x)$ is $(\sigma^2 + \lambda)$ -strongly convex,¹¹ implying that $\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq (\sigma^2 + \lambda)\|x - y\|^2$ [25, Lemma 3.11], which can also be directly verified using the formula $\nabla f(x) = A^\dagger Ax - A^\dagger b + \lambda x$ for the gradient.

$$\begin{aligned} & \mathbb{E}[\|x^{(t)} - x^*\|^2 \mid x^{(t-1)}] \\ & = \mathbb{E}[\|x^{(t-1)} - x^*\|^2 + 2\langle x^{(t-1)} - x^*, x^{(t)} - x^{(t-1)} \rangle + \|x^{(t)} - x^{(t-1)}\|^2 \mid x^{(t-1)}] \\ & = \mathbb{E}[\|x^{(t-1)} - x^*\|^2 - 2\eta_t \langle x^{(t-1)} - x^*, \nabla g(x^{(t-1)}) \rangle + \eta_t^2 \|\nabla g(x^{(t-1)})\|^2 \mid x^{(t-1)}] \\ & = \|x^{(t-1)} - x^*\|^2 - 2\eta_t \langle x^{(t-1)} - x^*, \nabla f(x^{(t-1)}) \rangle + \eta_t^2 \mathbb{E}[\|\nabla g(x^{(t-1)})\|^2 \mid x^{(t-1)}] \\ & \leq (1 - 2\eta_t(\sigma^2 + \lambda)) + 2\eta_t^2(\|A^\dagger A + \lambda I\|^2 + 2\|A\|_{\mathbb{F}}^2\|A\|^2)\|x^{(t-1)} - x^*\|^2 + 4\eta_t^2\|A\|_{\mathbb{F}}^2\|A\|^2\|x^*\|^2. \end{aligned}$$

We conditioned on $x^{(t-1)}$ in the above computation, so the randomness in ∇g came only from one iteration. Let $\delta_t := \mathbb{E}[\|x^{(t)} - x^*\|^2]$. Then by taking expectation over $x^{(t-1)}$ for both sides of the above computation, we get

$$\begin{aligned} \delta_t & \leq (1 - 2\eta_t(\sigma^2 + \lambda)) + 2\eta_t^2(\|A^\dagger A + \lambda I\|^2 + 2\|A\|_{\mathbb{F}}^2\|A\|^2)\delta_{t-1} + 4\eta_t^2\|A\|_{\mathbb{F}}^2\|A\|^2\|x^*\|^2 \\ & \leq (1 - 2\eta_t(\sigma^2 + \lambda)) + 2\eta_t^2(4\|A\|_{\mathbb{F}}^2\|A\|^2 + 2\lambda^2)\delta_{t-1} + 4\eta_t^2\|A\|_{\mathbb{F}}^2\|A\|^2\|x^*\|^2, \end{aligned}$$

where the last step follows from the inequality $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$. Now, by substituting $\eta_t^2 = \eta^2 = \frac{\varepsilon^2(\sigma^2 + \lambda)}{8\|A\|_{\mathbb{F}}^2\|A\|^2 + 4\lambda^2}$ and using that $\varepsilon \leq 1$ we get

$$\delta_t \leq (1 - \eta(\sigma^2 + \lambda))\delta_{t-1} + \frac{\varepsilon^2}{2}\eta(\sigma^2 + \lambda)\|x^*\|^2.$$

This is a purely deterministic recursion on δ_t , which we can bound by

$$\delta_t \leq \left(\exp(-t\eta(\sigma^2 + \lambda)) + \frac{\varepsilon^2}{2} \right) \|x^*\|^2.$$

¹¹Since we search for a sparse description of a solution vector we are effectively restricting the function to the image of A^\dagger , and on this subspace the function is indeed $(\sigma^2 + \lambda)$ -strongly convex.

Since $x^{(0)} = \vec{0}$ we have $\delta_0 = \|x^*\|^2$ so the bound holds for $t = 0$, and by induction we have

$$\begin{aligned}\delta_t &\leq (1 - \eta(\sigma^2 + \lambda)) \left(\exp(-(t-1)\eta(\sigma^2 + \lambda)) + \frac{\varepsilon^2}{2} \right) \|x^*\|^2 + \frac{\varepsilon^2}{2} \eta(\sigma^2 + \lambda) \|x^*\|^2 \\ &= \left((1 - \eta(\sigma^2 + \lambda)) \exp(-(t-1)\eta(\sigma^2 + \lambda)) + \frac{\varepsilon^2}{2} \right) \|x^*\|^2 \\ &\leq \left(\exp(-t\eta(\sigma^2 + \lambda)) + \frac{\varepsilon^2}{2} \right) \|x^*\|^2,\end{aligned}\tag{4}$$

which is at most $\varepsilon^2 \|x^*\|^2$ for $t \geq T = \frac{\ln(2/\varepsilon^2)}{\eta(\sigma^2 + \lambda)}$. \square

2.2 Time complexity analysis of SGD for sparse b

Now, we show that, assuming $\text{SQ}(A)$, it's possible to perform the gradient steps when b is sparse (that is, the number of non-zero entries of b , $\|b\|_0$, is small). Recall from Definition 1.3 that we say we have an s -sparse description of $x \in \mathbb{C}^d$ if we have an s -sparse $v \in \mathbb{C}^n$ such that $x = A^\dagger v$.

Lemma 2.4. *Given $\text{SQ}(A)$, we can output $x^{(t)}$ as a $(t + \|b\|_0)$ -sparse description in time*

$$\mathcal{O}(Ct(t + \|b\|_0)).$$

Proof. First, suppose we are given $x^{(t)}$ as an s -sparse description, and wish to output a sparse description for the next iterate $x^{(t+1)}$, which from Definition 2.1, satisfies

$$\begin{aligned}x^{(t+1)} &= x^{(t)} - \eta_{t+1} \nabla g(x^{(t)}) \\ &= x^{(t)} - \eta_{t+1} \cdot \left(\frac{\|A\|_{\text{F}}^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j}^{(t)} \right) (A_{r,*})^\dagger - A^\dagger b + \lambda x^{(t)} \right).\end{aligned}$$

The r and c_j 's are drawn from distributions, and $\text{SQ}(A)$ can produce such samples with its `SAMPLE` queries, taking $\mathcal{O}(C)$ time.

From inspection of the above equation, if we have $x^{(t)}$ in terms of its description as $A^\dagger v^{(t)}$, then we can write $x^{(t+1)}$ as a description $A^\dagger v^{(t+1)}$ where $v^{(t+1)}$ satisfies (as in Equation (1))

$$v^{(t+1)} = v^{(t)} - \eta_{t+1} \cdot \left(\frac{\|A\|_{\text{F}}^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v^{(t)} \right) e_r - b + \lambda v^{(t)} \right).\tag{5}$$

Here, e_r is the vector that is one in the r th entry and zero otherwise. So, if $v^{(t)}$ is s -sparse, and has a support that includes the support of b , then $v^{(t+1)}$ is $(s+1)$ -sparse. Furthermore, by exploiting the sparsity of $v^{(t)}$, computing $v^{(t+1)}$ takes $\mathcal{O}(Cs)$ time (including the time taken to use $\text{SQ}(A)$ to query A for all of the relevant norms and entries).

So, if we wish to compute $x^{(t)}$, we begin with $x^{(0)}$, which we have trivially as an $\|b\|_0$ -sparse description ($v^{(0)} = \vec{0}$). It is sparser, but if we consider $x^{(0)}$ as having the same support as b , by the argument described above, we can then compute $x^{(1)}$ as an $(\|b\|_0 + 1)$ -sparse description in $\mathcal{O}(C\|b\|_0)$ time.

By iteratively computing $v^{(i+1)}$ from $v^{(i)}$ in $\mathcal{O}(C(\|b\|_0 + i))$ time, we can output $x^{(t)}$ as a $(t + \|b\|_0)$ -sparse description in

$$\mathcal{O}\left(\sum_{i=0}^t C(\|b\|_0 + i)\right) = \mathcal{O}(Ct(t + \|b\|_0))$$

time as desired. \square

If $\lambda = \mathcal{O}(\|A\|_F \|A\|)$ by Proposition 2.3, to get a good enough iterate $x^{(T)}$ we can take $T = \mathcal{O}\left(\frac{\|A\|^2 \|A\|_F^2}{(\sigma^4 + \lambda^2)\varepsilon^2} \log \frac{1}{\varepsilon}\right)$ and $C = \frac{\|A\|_F^2}{\|A\|^2}$, giving a running time of

$$\mathcal{O}\left(\|b\|_0 \frac{\|A\|_F^4}{(\sigma^4 + \lambda^2)\varepsilon^2} \log \frac{1}{\varepsilon} + \frac{\|A\|^2 \|A\|_F^6}{(\sigma^8 + \lambda^4)} \log^2 \frac{1}{\varepsilon}\right).$$

Notice that it's good to scale C up to be as large as possible, since it means a corresponding linear decrease in the number of iterations, which our algorithm's running time depends on quadratically.

After performing SGD, we have an $x^{(T)}$ that is close to x^* as desired, and it is given as a sparse description $A^\dagger v^{(T)}$. We want to say that we have $\text{SQ}(x^{(T)})$ from its description. For this we invoke a result from [11], describing how to length-square sample from a vector that is a linear combination of length-square accessible vectors—that is, getting $\text{SQ}(c_1 v_1 + \dots + c_d v_d)$ given $\text{SQ}(v_1), \dots, \text{SQ}(v_d)$.

Lemma 2.5 ([9, Lemmas 2.9 and 2.10]). *Suppose we have $\text{SQ}(M^\dagger)$ for $M \in \mathbb{C}^{n \times d}$ and $Q(x) \in \mathbb{C}^d$. Denote $y := Mx$ and $\Delta := \sum_{i=1}^d \|M_{*,i}\|^2 x_i^2 / \|y\|_2^2$. Then we can implement $\text{SQ}(y) \in \mathbb{C}^n$ with $\mathcal{T}(y) = \mathcal{O}(d^2 \Delta \log(1/\delta) \cdot \mathcal{T}(M))$, where queries succeed with probability $\geq 1 - \delta$. Namely, we can:*

- (a) query for entries with complexity $\mathcal{O}(d \cdot \mathcal{T}(M))$;
- (b) sample from y with running time $\mathcal{T}_{\text{SAMPLE}}(y)$ satisfying

$$\mathbb{E}[\mathcal{T}_{\text{SAMPLE}}(y)] = \mathcal{O}(d^2 \Delta \cdot \mathcal{T}(M))$$

and $\Pr[\mathcal{T}_{\text{SAMPLE}}(y) = \mathcal{O}(d^2 \Delta \cdot \log(1/\delta))] \geq 1 - \delta$.

- (c) estimate $\|y\|$ to $(1 \pm \varepsilon)$ multiplicative error with success probability at least $1 - \delta$ in complexity

$$\mathcal{O}\left(\frac{d^2 \Delta}{\varepsilon^2} \mathcal{T}(M) \cdot \log(1/\delta)\right).$$

So we care about the quantity $\Delta = \frac{1}{\|A^\dagger v^{(t)}\|_2^2} \sum_{i=1}^m \|A_{i,*}\|^2 |v_i^{(t)}|^2$, with $t = T$, where $v^{(t)}$ follows the recurrence according to Eq. (5) (recalling from before that r and c_1, \dots, c_C are sampled randomly and independently each iteration):

$$v^{(t+1)} = v^{(t)} - \eta_{t+1} \cdot \underbrace{\left(\frac{\|A\|_F^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{\|A_{r,c_j}\|^2} A_{r,c_j} (A_{*,c_j})^\dagger v^{(t)}\right) e_r - b + \lambda v^{(t)}\right)}_{\nabla \tilde{g}(v^{(t)})}.$$

Roughly speaking, Δ encodes the amount of cancellation that could occur in the product $A^\dagger v^{(t)}$. We will consider it as a norm: $\Delta = \frac{\|v^{(t)}\|_D^2}{\|x^{(t)}\|^2}$, for D a diagonal matrix with $D_{ii} = \|A_{i,*}\|$, where $\|v\|_D := \sqrt{v^\dagger D^\dagger D v}$. The rest of this section will be devoted to bounding Δ where $x^{(t)}$ comes from SGD as in Proposition 2.3.

First, notice that we can show similar moment bounds for $\nabla \tilde{g}(v^{(t)})$ in the D norm as those for $\nabla g(x^{(t)})$ in Lemma 2.2; we defer the proof of this to the appendix.

Lemma 2.6. *Let $\nabla \tilde{g}(v)$ denote*

$$\frac{\|A\|_F^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v \right) e_r - b + \lambda v,$$

where r, c_1, \dots, c_C are sampled according to the same distribution as ∇g is (Definition 2.1). For a fixed $v \in \mathbb{C}^m$ and $C = \frac{\|A\|_F^2}{\|A\|^2}$, we have:

$$\begin{aligned} \text{Part 1.} \quad & \mathbb{E}[\nabla \tilde{g}(v)] = AA^\dagger v - b + \lambda v \\ \text{Part 2.} \quad & \mathbb{E}[\|\nabla \tilde{g}(v) - \mathbb{E}[\nabla \tilde{g}(v)]\|_D^2] \leq 2\|A\|_F^2 \|A\|^2 \|A^\dagger v\|^2. \end{aligned}$$

To bound Δ , we will show a recurrence via Lemma 2.6. Note that

$$\begin{aligned} & \mathbb{E}[\|v^{(t+1)}\|_D \mid v^{(t)}] \\ & \leq \|v^{(t)}\|_D (1 - \eta_{t+1} \lambda) + \eta_{t+1} \mathbb{E}[\|\nabla \tilde{g}(v^{(t)}) - \lambda v^{(t)}\|_D \mid v^{(t)}] && \text{by triangle inequality} \\ & \leq \|v^{(t)}\|_D + \eta_{t+1} \sqrt{\mathbb{E}[\|\nabla \tilde{g}(v^{(t)}) - \lambda v^{(t)}\|_D^2 \mid v^{(t)}]} && \text{by } \mathbb{E}[Z^2] \leq \mathbb{E}[Z^2] \\ & \leq \|v^{(t)}\|_D + \eta_{t+1} \sqrt{\|AA^\dagger v^{(t)} - b\|_D^2 + 2\|A\|_F^2 \|A\|^2 \|A^\dagger v^{(t)}\|^2} && \text{by } \mathbb{E}[Z^2] = \mathbb{E}[Z]^2 + \mathbb{V}[Z] \\ & \leq \|v^{(t)}\|_D + \eta_{t+1} (\|AA^\dagger v^{(t)} - b\|_D + \sqrt{2}\|A\|_F \|A\| \|A^\dagger v^{(t)}\|) && \text{by } \sqrt{a+b} \leq \sqrt{a} + \sqrt{b} \\ & \leq \|v^{(t)}\|_D + \eta_{t+1} (\|b\|_D + (1 + \sqrt{2})\|A\|_F \|A\| \|x^{(t)}\|). && \text{by } x^{(t)} = A^\dagger v^{(t)} \text{ and } \|u\|_D \leq \|A\| \|u\| \end{aligned}$$

Taking expectation over $v^{(t)}$ we get

$$\mathbb{E}[\|v^{(t+1)}\|_D] \leq \mathbb{E}[\|v^{(t)}\|_D] + \eta_{t+1} (\|b\|_D + 3\|A\|_F \|A\| \mathbb{E}[\|x^{(t)}\|]),$$

and since $v^{(0)} = \vec{0}$ we can trivially solve this recurrence resulting in

$$\mathbb{E}[\|v^{(t)}\|_D] \leq \sum_{t'=1}^t \eta_{t'} (\|b\|_D + 3\|A\|_F \|A\| \mathbb{E}[\|x^{(t'-1)}\|]). \quad (6)$$

Further, using the parameter choices of Proposition 2.3 due to Eq. (4) we have

$$\mathbb{E}[\|x^{(t')}\|] \leq \|x^*\| + \sqrt{\mathbb{E}[\|x^{(t')} - x^*\|^2]} = \|x^*\| + \sqrt{\delta_{t'}} \leq 3\|x^*\|,$$

so

$$\mathbb{E}[\|v^{(t)}\|_D] \leq t\eta (\|b\|_D + 9\|A\|_F \|A\| \|x^*\|). \quad (7)$$

Finally, we can bound our cancellation constant. Using Markov's inequality for Eq. (7) and combining it with Proposition 2.3 by a union bound we get that with probability ≥ 0.9 , both $\|x^{(T)} - x^*\| \leq 5\epsilon \|x^*\|$ and

$$\|v^{(T)}\|_D \leq 20T\eta (\|b\|_D + 9\|A\|_F \|A\| \|x^*\|) = 20 \frac{\ln(2/\epsilon^2)}{(\sigma^2 + \lambda)} (\|b\|_D + 9\|A\|_F \|A\| \|x^*\|)$$

When both those bounds hold and $\varepsilon \leq 1/10$, we have

$$\begin{aligned} \Delta &= \frac{\|v^{(T)}\|_D^2}{\|x^{(T)}\|^2} \leq \frac{\|v^{(T)}\|_D^2}{\|x^*\|^2(1-5\varepsilon)^2} \lesssim \frac{\log^2(1/\varepsilon)}{(\sigma^2 + \lambda)^2} \left(\frac{\|b\|_D^2}{\|x^*\|^2} + \|A\|_F^2 \|A\|^2 \right) \\ &\lesssim \frac{\log^2(1/\varepsilon)}{(\sigma^2 + \lambda)^2} \left(\frac{\|A\|^2 \|b\|^2}{\|x^*\|^2} + \|A\|_F^2 \|A\|^2 \right). \end{aligned} \quad (8)$$

The last inequality follows from using that $\|b\|_D^2 = \|Db\|^2 \leq \|D\|^2 \|b\|^2 \leq \|A\|^2 \|b\|^2$.

2.3 Extending to non-sparse b : Proof of Theorem 1.4

In previous sections, we have shown how to solve our regularized regression problem for sparse b : from Proposition 2.3, performing SGD for $T \approx \frac{\|A\|_F^2 \|A\|^2 + \lambda^2}{(\sigma^2 + \lambda)^2 \varepsilon^2} \log \frac{1}{\varepsilon}$ iterations outputs an x with the desired error bound; from Lemma 2.4, it takes $\mathcal{O}\left(\frac{\|A\|_F^2}{\|A\|^2} T(T + \|b\|_0)\right)$ time to output x as a sparse description; and from Section 2.2, we have sample and query access to that output x given its sparse description. Let $\lambda = \mathcal{O}(\|A\|_F \|A\|)$ in the rest of the section.

Now, all that remains is to extend this work to the case that b is non-sparse. In this case, we will simply replace b with a sparse \hat{b} that behaves similarly, and show that running SGD with this value of \hat{b} gives all the same results. The sparsity of \hat{b} will be $\mathcal{O}\left(\frac{\|A\|_F^2 \|b\|^2}{(\sigma^2 + \lambda)^2 \varepsilon^2 \|x^*\|^2}\right)$, giving a total running time of

$$\mathcal{O}\left(\frac{\|A\|_F^6 \|A\|^2}{(\sigma^2 + \lambda)^4 \varepsilon^4} \left(\frac{\|b\|^2}{\|A\|^2 \|x^*\|^2} + \log \frac{1}{\varepsilon} \right) \log \frac{1}{\varepsilon}\right).$$

We show below that the bound in Eq. (8) also holds in this case. Using Lemma 2.5, the time it takes to respond to a query to $\text{SQ}(x^{(T)})$ with probability 0.99 is $(T + \|\hat{b}\|_0)^2 \Delta$, which gives the running time in Theorem 1.4,

$$\mathcal{O}\left((T + \|\hat{b}\|_0)^2 \Delta\right) = \mathcal{O}\left(\frac{\|A\|_F^6 \|A\|^6}{(\sigma^2 + \lambda)^6 \varepsilon^4} \log^2 \frac{1}{\varepsilon} \left(\frac{\|b\|^4}{\|A\|^4 \|x^*\|^4} + \log^2 \frac{1}{\varepsilon} \right) \left(\frac{\|b\|^2}{\|A\|_F^2 \|x^*\|^2} + 1 \right)\right).$$

The crucial observation for sparsifying b is that we can use importance sampling to approximate the matrix product $A^\dagger b$, which suffices to approximate the solution x^* .

Lemma 2.7 (Matrix multiplication to Frobenius norm error, [35, Lemma 4]). *Consider $X \in \mathbb{C}^{m \times n}$, $Y \in \mathbb{C}^{m \times p}$, and let $S \in \mathbb{R}^{s \times m}$ be an importance sampling matrix for X . That is, let each $S_{i,*}$ be independently sampled to be $\frac{\|X\|_F}{\sqrt{s \|X_{i,*}\|}} e_i$ with probability $\frac{\|X_{i,*}\|^2}{\|X\|_F^2}$. Then*

$$\mathbb{E}[\|X^\dagger S^\dagger S Y - X^\dagger Y\|_F^2] \leq \frac{1}{s} \|X\|_F^2 \|Y\|_F^2.$$

We have $\text{SQ}(A)$, so we can use Lemma 2.7 with $s \leftarrow 200 \frac{\|A\|_F^2 \|b\|^2}{(\sigma^2 + \lambda)^2 \varepsilon^2 \|x^*\|^2}$, $X \leftarrow A$, and $Y \leftarrow b$, to find an S in $\mathcal{O}(s)$ time that satisfies the guarantee

$$\|A^\dagger S^\dagger S b - A^\dagger b\| \leq \sqrt{\frac{200}{s}} \|A\|_F \|b\| = \varepsilon(\sigma^2 + \lambda) \|x^*\| \quad (9)$$

with probability ≥ 0.995 , using Markov's inequality. Recall that D is defined to be the diagonal matrix with $D_{i,i} = \|A_{i,*}\|$, so $D^\dagger = D$, $\|D\|_F = \|A\|_F$, and importance sampling matrices for A are also importance sampling matrices for D . Consequently, we can apply Lemma 2.7 with the same argument to conclude that, with probability ≥ 0.995 ,

$$\|DS^\dagger S b - D b\| \leq \sqrt{\frac{200}{s}} \|D\|_F \|b\| = \varepsilon(\sigma^2 + \lambda) \|x^*\| \quad (10)$$

By union bound, both Eq. (9) and Eq. (10) hold *for the same* S with probability ≥ 0.99 . Assuming Eqs. (9) and (10) hold, we can perform SGD as in Proposition 2.3 on $\hat{b} := S^\dagger S b$ (which is s -sparse) to find an x such that $\|x - \hat{x}^*\| \leq \varepsilon \|\hat{x}^*\|$, where \hat{x}^* is the optimum $(A^\dagger A + \lambda I)^{-1} A^\dagger \hat{b}$. This implies that $\|x - x^*\| \leq 2\varepsilon \|x^*\|$, since by Eq. (9),¹²

$$\|\hat{x}^* - x^*\| = \|(A^\dagger A + \lambda I)^+ A^\dagger (\hat{b} - b)\| \leq \frac{1}{\sigma^2 + \lambda} \|A^\dagger (\hat{b} - b)\| \leq \varepsilon \|x^*\|. \quad (11)$$

To bound the running times of $\text{SQ}(x)$, we modify the analysis of Δ from Section 2.2: recalling from Eq. (8), we have that, with probability ≥ 0.9 ,

$$\Delta \lesssim \frac{\log^2(1/\varepsilon)}{(\sigma^2 + \lambda)^2} \left(\frac{\|\hat{b}\|_D^2}{\|\hat{x}^*\|^2} + \|A\|_F^2 \|A\|^2 \right) \lesssim \frac{\log^2(1/\varepsilon)}{(\sigma^2 + \lambda)^2} \left(\frac{\|b\|_D^2}{\|x^*\|^2} + \|A\|_F^2 \|A\|^2 \right). \quad (12)$$

In other words, the same bound on Δ from Eq. (8) holds after approximating b by \hat{b} . The last step follows from the following upper bound on $\frac{\|\hat{b}\|_D^2}{\|\hat{x}^*\|^2}$:

$$\begin{aligned} \frac{\|\hat{b}\|_D^2}{\|\hat{x}^*\|^2} &\leq \frac{\|\hat{b}\|_D^2}{\|x^*\|^2 (1 - \varepsilon)^2} = \frac{\|DS^\dagger S b\|^2}{\|x^*\|^2 (1 - \varepsilon)^2} \leq \frac{(\|D b\| + \|D b - DS^\dagger S b\|)^2}{\|x^*\|^2 (1 - \varepsilon)^2} \\ &\lesssim \frac{\|D b\|^2 + \varepsilon^2 (\sigma^2 + \lambda)^2 \|x^*\|^2}{\|x^*\|^2 (1 - \varepsilon)^2} \lesssim \frac{\|b\|_D^2}{\|x^*\|^2} + \varepsilon^2 (\sigma^2 + \lambda)^2 \end{aligned}$$

where the first inequality follows from Eq. (11), the second inequality is triangle inequality, the third inequality follows from Eq. (10), and the last inequality holds for $\varepsilon \leq \frac{1}{2}$ which we can assume without loss of generality. The added term of $\varepsilon^2 (\sigma^2 + \lambda)^2$ is dominated by the $\|A\|_F^2 \|A\|^2$ term in Eq. (12).

Finally, we observe the following fact about $\|x^*\|$ which can be used to convert the runtime of Theorem 1.4 to other parameters.

Fact 2.8. $\|x^*\| \geq \frac{\|A\|}{\|A\|^2 + \lambda} \|\Pi_{A,\lambda} b\|$ when $\|A\| \geq \sqrt{\lambda}$, and $\|x^*\| = \frac{1}{2\sqrt{\lambda}} \|\Pi_{A,\lambda} b\|$ otherwise.

$$\begin{aligned} \|x^*\| &= \|(A^\dagger A + \lambda I)^+ A^\dagger b\| \\ &= \|(A^\dagger A + \lambda I)^+ A^\dagger \Pi_{A,\lambda}^+ \Pi_{A,\lambda} b\| && \text{since } A^\dagger = A^\dagger \Pi_{A,\lambda}^+ \Pi_{A,\lambda} \\ &\geq \|\Pi_{A,\lambda} b\| \min_{\substack{v \in \text{span}(A^\dagger) \\ \|v\|=1}} \|(A^\dagger A + \lambda I)^+ A^\dagger \Pi_{A,\lambda}^+ v\| \end{aligned} \quad (13)$$

The min term is equal to the minimum non-zero singular value of $(A^\dagger A + \lambda I)^+ A^\dagger \Pi_{A,\lambda}^+$. The non-zero singular values of this expression are $g(\sigma_i)$, where σ_i is a non-zero singular value of A and

$$g(\sigma) = \frac{\sigma}{(\sigma^2 + \lambda)(p_{A,\lambda}(\sigma))} = \begin{cases} \frac{1}{2\sqrt{\lambda}} & 0 < \sigma \leq \sqrt{\lambda} \\ \frac{\sigma}{\sigma^2 + \lambda} & \sqrt{\lambda} < \sigma. \end{cases}$$

¹²We obtain the constants in Algorithm 1 by using the error parameter $\hat{\varepsilon} := \varepsilon/2$ to satisfy $\|x - x^*\| \leq 2\hat{\varepsilon} \|x^*\| = \varepsilon \|x^*\|$.

This function is non-increasing, so the minimum singular value of $(A^\dagger A + \lambda I)^+ A^\dagger \Pi_{A,\lambda}^+$ is $g(\|A\|)$, proving the above fact by observing that for $\|A\| \leq \sqrt{\lambda}$ the operator $(A^\dagger A + \lambda I)^+ A^\dagger \Pi_{A,\lambda}^+$ is proportional to a projector, thus Eq. (13) becomes an equality.

Acknowledgments

E.T. thanks Kevin Tian immensely for discussions integral to these results. Z.S. thanks Runzhou Tao and Ruizhe Zhang for giving helpful comments on the draft.

A.G. acknowledges funding provided by Samsung Electronics Co., Ltd., for the project “The Computational Power of Sampling on Quantum Computers”, and additional support by the Institute for Quantum Information and Matter, an NSF Physics Frontiers Center (NSF Grant PHY-1733907), as well as support by ERC Consolidator Grant QPROGRESS, by QuantERA project QuantAlgo 680-91-034 and also by the EU’s Horizon 2020 Marie Skłodowska-Curie program 891889-QuantOrder. Z.S. was partially supported by Schmidt Foundation and Simons Foundation. E.T. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1762114.

References

- [1] John Preskill. “Quantum Computing in the NISQ era and beyond”. *Quantum* **2**, 79 (2018). [arXiv:1801.00862](#).
- [2] Andrew M Childs. “Equation solving by simulation”. *Nature Physics* **5**, 861–861 (2009).
- [3] Scott Aaronson. “Read the fine print”. *Nature Physics* **11**, 291–293 (2015).
- [4] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. “Quantum machine learning”. *Nature* **549**, 195–202 (2017). [arXiv:1611.09347](#).
- [5] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. “Quantum algorithm for linear systems of equations”. *Physical Review Letters* **103**, 150502 (2009). [arXiv:0811.3171](#).
- [6] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. “Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics”. In Proceedings of the 51st ACM Symposium on the Theory of Computing (STOC). Pages 193–204. (2019). [arXiv:1806.01838](#).
- [7] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. “Quantum random access memory”. *Physical Review Letters* **100**, 160501 (2008). [arXiv:0708.1879](#).
- [8] Anupam Prakash. “Quantum algorithms for linear algebra and machine learning”. PhD thesis. University of California at Berkeley. (2014). url: www2.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-211.pdf.
- [9] Nai-Hui Chia, András Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. “Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning”. In Proceedings of the 52nd ACM Symposium on the Theory of Computing (STOC). Page 387–400. (2020). [arXiv:1910.06151](#).
- [10] Iordanis Kerenidis and Anupam Prakash. “Quantum recommendation systems”. In Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS). Pages 49:1–49:21. (2017). [arXiv:1603.08675](#).
- [11] Ewin Tang. “A quantum-inspired classical algorithm for recommendation systems”. In Proceedings of the 51st ACM Symposium on the Theory of Computing (STOC). Pages 217–228. (2019). [arXiv:1807.04271](#).

- [12] Iordanis Kerenidis, Jonas Landman, Alessandro Luongo, and Anupam Prakash. “q-means: A quantum algorithm for unsupervised machine learning”. In *Advances in Neural Information Processing Systems*. Volume 32. (2019). [arXiv:1812.03584](#).
- [13] Iordanis Kerenidis and Anupam Prakash. “Quantum gradient descent for linear systems and least squares”. *Physical Review A* **101**, 022316 (2020). [arXiv:1704.04992](#).
- [14] Danial Dervovic, Mark Herbster, Peter Mountney, Simone Severini, Nairi Usher, and Leonard Wossnig. “Quantum linear systems algorithms: a primer” (2018). [arXiv:1802.08227](#).
- [15] Leonard Wossnig, Zhikuan Zhao, and Anupam Prakash. “Quantum linear system algorithm for dense matrices”. *Physical Review Letters* **120**, 050502 (2018). [arXiv:1704.06174](#).
- [16] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. “Quantum support vector machine for big data classification”. *Physical Review Letters* **113**, 130503 (2014). [arXiv:1307.0471](#).
- [17] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. “The power of block-encoded matrix powers: Improved regression techniques via faster Hamiltonian simulation”. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*. Pages 33:1–33:14. (2019). [arXiv:1804.01973](#).
- [18] Nai-Hui Chia, András Gilyén, Han-Hsuan Lin, Seth Lloyd, Ewin Tang, and Chunhao Wang. “Quantum-inspired algorithms for solving low-rank linear equation systems with logarithmic dependence on the dimension”. In *Proceedings of the 31st International Symposium on Algorithms and Computation (ISAAC)*. Pages 47:1–47:17. (2020). [arXiv:1811.04852](#) and [1811.04909 \(merged\)](#).
- [19] Juan Miguel Arrazola, Alain Delgado, Bhaskar Roy Bardhan, and Seth Lloyd. “Quantum-inspired algorithms in practice”. *Quantum* **4**, 307 (2020). [arXiv:1905.10415](#).
- [20] Ewin Tang. “Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions”. *Physical Review Letters* **127**, 060503 (2021). [arXiv:1811.00414](#).
- [21] Carlo Ciliberto, Mark Herbster, Alessandro Davide Ialongo, Massimiliano Pontil, Andrea Rocchetto, Simone Severini, and Leonard Wossnig. “Quantum machine learning: a classical perspective”. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **474**, 20170551 (2018). [arXiv:1707.08561](#).
- [22] Srinivasan Arunachalam, Vlad Gheorghiu, Tomas Jochym-O’Connor, Michele Mosca, and Priyaa Varshinee Srinivasan. “On the robustness of bucket brigade quantum RAM”. *New Journal of Physics* **17**, 123010 (2015). [arXiv:1502.03450](#).
- [23] Neha Gupta and Aaron Sidford. “Exploiting numerical sparsity for efficient learning: Faster eigenvector computation and regression”. In *Advances in Neural Information Processing Systems*. Pages 5269–5278. (2018). [arXiv:1811.10866](#).
- [24] Yair Carmon, Yujia Jin, Aaron Sidford, and Kevin Tian. “Coordinate methods for matrix games”. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. Pages 283–293. IEEE (2020). [arXiv:2009.08447](#).
- [25] Sébastien Bubeck. “Convex optimization: Algorithms and complexity”. *Foundations and Trends in Machine Learning* **8**, 231–357 (2015). [arXiv:1405.4980](#).
- [26] Roy Frostig, Rong Ge, Sham Kakade, and Aaron Sidford. “Un-regularizing: Approximate proximal point and faster stochastic algorithms for empirical risk minimization”. In *International Conference on Machine Learning*. Pages 2540–2548. (2015). [arXiv:1506.07512](#).
- [27] Francis Bach and Eric Moulines. “Non-asymptotic analysis of stochastic approximation algorithms for machine learning”. In *Advances in Neural Information Processing Systems*.

- tens. Pages 451–459. (2011). url: <http://papers.nips.cc/paper/4316-non-asymptotic-analysis-of-stochastic-approximation-algorithms-for-machine-learning.pdf>.
- [28] András Gilyén, Seth Lloyd, and Ewin Tang. “Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension” (2018). [arXiv:1811.04909](https://arxiv.org/abs/1811.04909).
- [29] Nai-Hui Chia, Han-Hsuan Lin, and Chunhao Wang. “Quantum-inspired sublinear classical algorithms for solving low-rank linear systems” (2018). [arXiv:1811.04852](https://arxiv.org/abs/1811.04852).
- [30] David P. Woodruff. “Sketching as a tool for numerical linear algebra”. *Foundations and Trends in Theoretical Computer Science* **10**, 1–157 (2014).
- [31] Nadiia Chepurko, Kenneth L. Clarkson, Lior Horesh, Honghao Lin, and David P. Woodruff. “Quantum-inspired algorithms from randomized numerical linear algebra” (2020). [arXiv:2011.04125](https://arxiv.org/abs/2011.04125).
- [32] Changpeng Shao and Ashley Montanaro. “Faster quantum-inspired algorithms for solving linear systems” (2021). [arXiv:2103.10309](https://arxiv.org/abs/2103.10309).
- [33] Thomas Strohmer and Roman Vershynin. “A randomized Kaczmarz algorithm with exponential convergence”. *Journal of Fourier Analysis and Applications* **15**, 262–278 (2008). [arXiv:math/0702226](https://arxiv.org/abs/math/0702226).
- [34] Deanna Needell, Nathan Srebro, and Rachel Ward. “Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm”. *Mathematical Programming* **155**, 549–573 (2015). [arXiv:1310.5715](https://arxiv.org/abs/1310.5715).
- [35] Petros Drineas, Ravi Kannan, and Michael W Mahoney. “Fast Monte Carlo algorithms for matrices II: Computing a low-rank approximation to a matrix”. *SIAM Journal on Computing* **36**, 158–183 (2006).

A Stochastic gradient bounds

Definition 2.1. We define $\nabla g(x)$ to be the random function resulting from the following process. Draw $r \in [m]$ from the distribution that is r with probability $\frac{\|A_{r,*}\|_F^2}{\|A\|_F^2}$, and then draw c_1, \dots, c_C i.i.d. from the distribution that is c with probability $\frac{|A_{r,c}|^2}{\|A_{r,*}\|_F^2}$. For this choice of r and c_1, \dots, c_C , take

$$\nabla g(x) = \frac{\|A\|_F^2}{\|A_{r,*}\|_F^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|_F^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j} \right) (A_{r,*})^\dagger - A^\dagger b + \lambda x. \quad (3)$$

Lemma 2.2. For fixed $x, y \in \mathbb{C}^n$ and the random function $\nabla g(\cdot)$ defined in Definition 2.1, the following properties hold ($\nabla g(x)$ and $\nabla g(y)$ use the same instance of the random function):

- Part 1.** $\mathbb{E}[\nabla g(x)] = \nabla f(x) = A^\dagger A x - A^\dagger b + \lambda x$
- Part 2.** $\mathbb{V}[\nabla g(x)] = \frac{1}{C} \|A\|_F^4 \|x\|^2 + \left(1 - \frac{1}{C}\right) \|A\|_F^2 \|Ax\|^2 - \|A^\dagger A x\|^2$
- Part 3.** $\mathbb{E}[\|\nabla g(x) - \nabla g(y)\|_2^2] = \|(A^\dagger A + \lambda I)(x - y)\|^2 + \mathbb{V}[\nabla g(x - y)]$

Proof.

Part 1. First, we show $\nabla g(x)$ is unbiased.

$$\mathbb{E}[\nabla g(x)] = \mathbb{E} \left[\frac{1}{C} \sum_{j=1}^C \frac{\|A\|_F^2}{\|A_{r,*}\|_F^2} \frac{\|A_{r,*}\|_F^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j} (A_{r,*})^\dagger \right] - A^\dagger b + \lambda x \quad \text{by definition of } \nabla g$$

$$\begin{aligned}
&= \mathbb{E}_{r,c_1} \left[\frac{\|A\|_{\mathbb{F}}^2}{|A_{r,c_1}|^2} A_{r,c_1} x_{c_1} (A_{r,*})^\dagger \right] - A^\dagger b + \lambda x && \text{since the } c_i \text{'s are i.i.d.} \\
&= \left(\sum_r \sum_{c_1} \frac{\|A_{r,c_1}\|^2}{\|A\|_{\mathbb{F}}^2} \frac{\|A\|_{\mathbb{F}}^2}{\|A_{r,c_1}\|^2} A_{r,c_1} x_{c_1} (A_{r,*})^\dagger \right) - A^\dagger b + \lambda x \\
&= A^\dagger A x - A^\dagger b + \lambda x = \nabla f(x).
\end{aligned}$$

Part 2. Now variance. We use that, for a random vector v , $\mathbb{V}[v] = \mathbb{E}[\|v - \mathbb{E}[v]\|^2] = \mathbb{E}[\|v\|^2] - \|\mathbb{E}[v]\|^2$.

$$\begin{aligned}
\mathbb{V}[\nabla g(x)] &= \mathbb{V}[\nabla g(x) + A^\dagger b - \lambda x] \\
&= \mathbb{E}[\|\nabla g(x) + A^\dagger b - \lambda x - \mathbb{E}[\nabla g(x) + A^\dagger b - \lambda x]\|^2] && \text{by definition of } \mathbb{V}[\cdot] \\
&= \mathbb{V} \left[\frac{\|A\|_{\mathbb{F}}^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j} \right) (A_{r,*})^\dagger \right] && \text{by definition of } \nabla g \\
&= \mathbb{V} \left[\left(\frac{1}{C} \sum_{j=1}^C \frac{\|A\|_{\mathbb{F}}^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j} \right) (A_{r,*})^\dagger \right] \\
&= \mathbb{E} \left[\left\| \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A\|_{\mathbb{F}}^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j} \right) (A_{r,*})^\dagger \right\|^2 \right] - \|A^\dagger A x\|^2. \\
&= \mathbb{E} \left[\underbrace{\left\| \frac{1}{C} \sum_{j=1}^C \frac{\|A\|_{\mathbb{F}}^2}{|A_{r,c_j}|^2} A_{r,c_j} x_{c_j} \right\|^2}_{V} \|A_{r,*}\|^2 \right] - \|A^\dagger A x\|^2. \tag{14}
\end{aligned}$$

We expand the first term, using that it is an average of i.i.d. random variables:

$$\begin{aligned}
V &= \sum_{i=1}^m \frac{\|A_{i,*}\|^2}{\|A\|_{\mathbb{F}}^2} \mathbb{E}_{c_1, \dots, c_C} \left[\left\| \frac{1}{C} \sum_{j=1}^C \frac{\|A\|_{\mathbb{F}}^2}{|A_{i,c_j}|^2} A_{i,c_j} x_{c_j} \right\|^2 \right] \|A_{i,*}\|^2 \\
&= \|A\|_{\mathbb{F}}^2 \sum_{i=1}^m \mathbb{E}_{c_1, \dots, c_C} \left[\left\| \frac{1}{C} \sum_{j=1}^C \frac{\|A_{i,*}\|^2}{|A_{i,c_j}|^2} A_{i,c_j} x_{c_j} \right\|^2 \right] \\
&= \|A\|_{\mathbb{F}}^2 \sum_{i=1}^m \left(\frac{1}{C} \mathbb{V}_{c_1} \left[\frac{\|A_{i,*}\|^2}{|A_{i,c_1}|^2} A_{i,c_1} x_{c_1} \right] + |A_{i,*} x|^2 \right) \\
&= \|A\|_{\mathbb{F}}^2 \sum_{i=1}^m \left(\frac{1}{C} \sum_{j=1}^n \left(\frac{|A_{i,j}|^2}{\|A_{i,*}\|^2} \frac{\|A_{i,*}\|^4}{|A_{i,j}|^4} |A_{i,j} x_j|^2 \right) - \frac{1}{C} |A_{i,*} x|^2 + |A_{i,*} x|^2 \right) \\
&= \|A\|_{\mathbb{F}}^2 \sum_{i=1}^m \left(\frac{1}{C} \|A_{i,*}\|^2 \|x\|^2 + \left(1 - \frac{1}{C}\right) (A_{i,*} x)^2 \right) \\
&= \frac{1}{C} \|A\|_{\mathbb{F}}^4 \|x\|^2 + \left(1 - \frac{1}{C}\right) \|A\|_{\mathbb{F}}^2 \|A x\|^2
\end{aligned}$$

Part 3. Finally, the expression for $\mathbb{E}[\|\nabla g(x) - \nabla g(y)\|^2]$ follows from the observation that $\nabla g(x) - \nabla g(y)$ is simply the expression for $\nabla g(x - y)$, taking b to be the zero vector. \square

Lemma 2.6. Let $\nabla \tilde{g}(v)$ denote

$$\frac{\|A\|_{\mathbb{F}}^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v \right) e_r - b + \lambda v,$$

where r, c_1, \dots, c_C are sampled according to the same distribution as ∇g is (Definition 2.1). For a fixed $v \in \mathbb{C}^m$ and $C = \frac{\|A\|_F^2}{\|A\|^2}$, we have:

$$\begin{aligned} \text{Part 1.} \quad & \mathbb{E}[\nabla \tilde{g}(v)] = AA^\dagger v - b + \lambda v \\ \text{Part 2.} \quad & \mathbb{E}[\|\nabla \tilde{g}(v) - \mathbb{E}[\nabla \tilde{g}(v)]\|_D^2] \leq 2\|A\|_F^2 \|A\|^2 \|A^\dagger v\|^2. \end{aligned}$$

Proof.

Part 1. We first show the expectation,

$$\begin{aligned} \mathbb{E}[\nabla \tilde{g}(v)] &= \mathbb{E} \left[\frac{\|A\|_F^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v \right) e_r - b + \lambda v \right] \\ &= \sum_{r=1}^m \mathbb{E}_{c_1, \dots, c_C} \left[\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v \right] e_r - b + \lambda v \\ &= \sum_{r=1}^m \sum_{c=1}^n A_{r,c} (A_{*,c})^\dagger v e_r - b + \lambda v \\ &= AA^\dagger v - b + \lambda v. \end{aligned}$$

Part 2. For the variance in the D norm, we reuse a computation from the proof of Lemma 2.2:

$$\begin{aligned} \mathbb{E}[\|\nabla \tilde{g}(v) - \mathbb{E}[\nabla \tilde{g}(v)]\|_D^2] &= \mathbb{V} \left[D \frac{\|A\|_F^2}{\|A_{r,*}\|^2} \left(\frac{1}{C} \sum_{j=1}^C \frac{\|A_{r,*}\|^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v \right) e_r \right] \\ &= \mathbb{E} \left[\underbrace{\|A_{r,*}\|^2 \left| \frac{1}{C} \sum_{j=1}^C \frac{\|A\|_F^2}{|A_{r,c_j}|^2} A_{r,c_j} (A_{*,c_j})^\dagger v \right|^2}_{V \text{ from Eq. (14), where } x = A^\dagger v} \right] - \|AA^\dagger v\|_D^2 \\ &= \frac{1}{C} \|A\|_F^4 \|A^\dagger v\|^2 + \left(1 - \frac{1}{C}\right) \|A\|_F^2 \|AA^\dagger v\|^2 - \|AA^\dagger v\|_D^2 \\ &\leq 2\|A\|_F^2 \|A\|^2 \|A^\dagger v\|^2. \quad \square \end{aligned}$$