

ℓ_1 Regression using Lewis Weights Preconditioning and Stochastic Gradient Descent

David Durfee

Kevin A. Lai

Saurabh Sawlani

Georgia Institute of Technology

DDURFEE@GATECH.EDU

KEVINLAI@GATECH.EDU

SAWLANI@GATECH.EDU

Editors: Sebastien Bubeck, Vianney Perchet and Philippe Rigollet

Abstract

We present preconditioned stochastic gradient descent (SGD) algorithms for the ℓ_1 minimization problem $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$ in the overdetermined case, where there are far more constraints than variables. Specifically, we have $\mathbf{A} \in \mathbb{R}^{n \times d}$ for $n \gg d$. Commonly known as the Least Absolute Deviations problem, ℓ_1 regression can be used to solve many important combinatorial problems, such as minimum cut and shortest path. SGD-based algorithms are appealing for their simplicity and practical efficiency. Our primary insight is that careful preprocessing can yield preconditioned matrices $\tilde{\mathbf{A}}$ with strong properties (besides good condition number and low-dimension) that allow for faster convergence of gradient descent. In particular, we precondition using Lewis weights to obtain an isotropic matrix with fewer rows and strong upper bounds on all row norms. We leverage these conditions to find a good initialization, which we use along with recent smoothing reductions and accelerated stochastic gradient descent algorithms to achieve ϵ relative error in $\tilde{O}(nnz(\mathbf{A}) + d^{2.5}\epsilon^{-2})$ time with high probability, where $nnz(\mathbf{A})$ is the number of non-zeros in \mathbf{A} . This improves over the previous best result using gradient descent for ℓ_1 regression. We also match the best known running times for interior point methods in several settings.

Finally, we also show that if our original matrix \mathbf{A} is approximately isotropic and the row norms are approximately equal, we can give an algorithm that avoids using fast matrix multiplication and obtains a running time of $\tilde{O}(nnz(\mathbf{A}) + sd^{1.5}\epsilon^{-2} + d^2\epsilon^{-2})$, where s is the maximum number of non-zeros in a row of \mathbf{A} . In this setting, we beat the best interior point methods for certain parameter regimes.

Keywords: ℓ_1 regression, stochastic gradient descent, Lewis weights

1. Introduction

Stochastic gradient descent (SGD) is one of the most widely-used practical algorithms for optimization problems due to its simplicity and practical efficiency (Nesterov and Vial, 2008; Nemirovski et al., 2009). We consider SGD methods to solve the unconstrained overdetermined ℓ_1 regression problem, commonly known as the Least Absolute Deviations problem, which is defined as follows:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$ and $n \gg d$. Compared to Least Squares (ℓ_2) regression, the ℓ_1 regression problem is more robust and is thus useful when outliers are present in the data. Moreover, many

important combinatorial problems, such as minimum cut or shortest path, can be formulated as ℓ_1 regression problems (Chin et al., 2013), and high accuracy ℓ_1 regression can be used to solve general linear programs.¹ Since (1) can be formulated as a linear program (Portnoy and Koenker, 1997; Chen et al., 2001), generic methods for solving linear programs, such as the interior-point method (IPM), can be used to solve it (Portnoy and Koenker, 1997; Portnoy, 1997; Meng and Mahoney, 2013b; Lee and Sidford, 2015).

SGD algorithms are popular in practice for ℓ_1 and other regression problems because they are simple, scalable, and efficient. State-of-the-art algorithms for solving (1) utilize sketching techniques from randomized linear algebra to achieve $\text{poly}(d, \epsilon^{-1})$ running times, whereas a naive extension of Nesterov acceleration (Nesterov, 1983) to certain classes of non-smooth functions (Nesterov, 2005b,a, 2007; Allen-Zhu and Hazan, 2016) takes $\text{poly}(n, \epsilon^{-1})$ time. This difference is significant because $n \gg d$ in the overdetermined setting. Ideally, the only dependence on n in the running time will be implicitly in an additive $nnz(\mathbf{A})$ term.

Sketching techniques from randomized numerical linear algebra look to find a low-distortion embedding of \mathbf{A} into a smaller subspace, after which popular techniques for ℓ_1 regression can be applied on the reduced matrix. Efforts to build these sampled matrices or “coresets” have been made using random sampling (Clarkson, 2005), fast matrix multiplication (Sohler and Woodruff, 2011), and ellipsoidal rounding (Dasgupta et al., 2009; Clarkson et al., 2013). All of these methods produce coresets of size $\text{poly}(d, \epsilon^{-1}) \times d$ in time $O(n \cdot \text{poly}(d))$. Meng and Mahoney (2013a) and Woodruff and Zhang (2013) improve these techniques to produce similar coresets in $O(nnz(\mathbf{A}) + \text{poly}(d))$ time.

In addition to using sketching as a preprocessing step, one can also apply *preconditioners*. Preconditioners transform the input matrix into one with additional desirable properties, such as good condition number, that speed up subsequent computation. For our setting, we will use the term “preconditioning” to refer to dimensionality reduction followed by additional processing of the matrix. Of particular interest for our setting is the preconditioning technique of Cohen and Peng (2015), which utilizes a Lewis change of density (Lewis, 1978) to sample rows of \mathbf{A} with probability proportional to their Lewis weights such that the sampled matrix $\tilde{\mathbf{A}}$ approximately preserves ℓ_1 distances, which is to say that $\|\mathbf{A}\mathbf{x}\|_1 \approx \|\tilde{\mathbf{A}}\mathbf{x}\|_1$ for any vector \mathbf{x} . Lewis weights are in essence the “correct” sampling weights for dimensionality reduction in ℓ_1 regression, and they are used by the previous best SGD-based methods for solving (1) (Yang et al., 2016). As it turns out, Lewis weights also lead to nice ℓ_2 conditions for the sampled matrix. One of the key insights of this paper is to leverage these additional guarantees to obtain significantly faster running times for SGD after Lewis weight preconditioning.

Techniques Our techniques for solving the ℓ_1 regression problem follow the popular paradigm of preconditioning and then using gradient descent methods on the resulting problem. Typically, the preconditioner is a black-box combination of a sketching method with a matrix rotation, which yields a well-conditioned low-dimensional matrix. The crucial idea in this paper is that the sketch can give us some strong properties in addition to the low-dimensional embedding. By more tightly integrating the components of the preconditioner, we obtain faster running times for ℓ_1 regression.

In particular, preconditioning the given matrix-vector pair $[\mathbf{A} \ \mathbf{b}]$ using Lewis weights (Cohen and Peng, 2015) achieves a low-dimensional $[\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]$ such that $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1 \approx \|\tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}}\|_1$ for every

1. For instance, one can determine if $\{\mathbf{x} | \mathbf{A}\mathbf{x} - \mathbf{b}, \mathbf{x} \geq 0\}$ is feasible by writing an objective of the form $\alpha(\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1 + \|\mathbf{x}\|_1 + \|\mathbf{x} - \beta\mathbf{1}\|_1)$ where α and β are sufficiently large polynomials in the input size.

$\mathbf{x} \in \mathbb{R}^d$. It is then possible to apply the low-dimensional embedding properties of Lewis weights in a black-box manner to ℓ_1 regression algorithms, using the fact that this embedding reduces the row-dimension from n to $O(d\epsilon^{-2} \log n)$ and then plugging these new parameters into the runtime guarantees. Our critical observation will be that sampling by Lewis weights also has the important property that all leverage scores of the new matrix are approximately equal. Since rotations of a matrix do not change its leverage scores, we are free to rotate $\tilde{\mathbf{A}}$ to place it into isotropic position, in which case the leverage score condition implies that the row norms are tightly bounded.

The isotropic and row norm conditions yield two essential phenomena. First, a careful choice of initial vector can be shown to be close to optimal. Second, we get strong bounds on the gradient of any row. Using these properties, it is almost immediately implied that standard SGD only requires $O(d^2\epsilon^{-2})$ iterations to arrive at a solution $\hat{\mathbf{x}}$ with relative error² ϵ , leading to a total running time of $\tilde{O}(nnz(\mathbf{A}) + d^3\epsilon^{-2})$. These properties can be further applied to smoothing reductions by [Allen-Zhu and Hazan \(2016\)](#) and accelerated SGD algorithms by [Allen Zhu \(2017\)](#) to improve the runtime to $\tilde{O}(nnz(\mathbf{A}) + nd^{\omega-1} + \sqrt{nd}^{1.5}\epsilon^{-1})$. As previously mentioned, sampling by Lewis weights guarantees that $n \leq O(d\epsilon^{-2} \log n)$, so we also obtain a running time of $\tilde{O}(nnz(\mathbf{A}) + d^{2.5}\epsilon^{-2})$. Algorithm 1 gives the basic framework of our ℓ_1 solver.

Algorithm 1: General structure of our algorithm

Input: Matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, and vector $\mathbf{b} \in \mathbb{R}^n$, along with error parameter $\epsilon > 0$.

1. Precondition $[\mathbf{A} \ \mathbf{b}]$ by Lewis weight sampling as in [Cohen and Peng \(2015\)](#), along with a matrix rotation.
 2. Initialize \mathbf{x}_0 as the exact or approximate ℓ_2 minimizer of the preconditioned problem.
 3. Run a stochastic gradient descent algorithm on the preconditioned problem with starting point \mathbf{x}_0 .
-

Our Results Our first main theorem uses smoothing reductions from [Allen-Zhu and Hazan \(2016\)](#) with the accelerated SGD algorithm of [Allen Zhu \(2017\)](#):

Theorem 1 *Given $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$, assume $\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2$ is either 0 or bounded below by n^{-c} for some constant $c > 0$. Then for any $\epsilon > 0$, there is a routine that outputs $\tilde{\mathbf{x}}$ such that with high probability,³*

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon) \min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_1$$

with a runtime of $O(nnz(\mathbf{A}) \log^2 n + d^{2.5}\epsilon^{-2} \log^{1.5} n)$ whenever $n \geq d\epsilon^{-2} \log n$, and a runtime of $O(\sqrt{nd}^2\epsilon^{-1} \log n + nd^{\omega-1} \log n)$ when $n \leq d\epsilon^{-2} \log n$.

Achieving the bounds when $n \leq d\epsilon^{-2} \log n$ requires some additional technical work. Theorem 1 is proved in Section 3, where we also show our $\tilde{O}(nnz(\mathbf{A}) + d^3\epsilon^{-2})$ running time for standard SGD.

Our second main theorem is motivated by the fact that the theoretical running time bounds of fast matrix multiplication are difficult to achieve in practice, and most implementations of algorithms

2. Relative error ϵ means that $f(\hat{\mathbf{x}}) - f(\mathbf{x}^*) \leq \epsilon f(\mathbf{x}^*)$, where $f(\mathbf{x}) = \mathbf{Ax} - \mathbf{b}$ and $\mathbf{x}^* = \arg \min_{\mathbf{x}} f(\mathbf{x})$.
 3. Throughout this paper, we let “with high probability” mean that ξ is the failure probability and our runtime has dependence on $\log(1/\xi)$, which we ignore for ease of notation.

actually use naive matrix multiplication. Thus, it would be ideal for an algorithm’s running time to be independent of fast matrix multiplication. It turns out that our only dependence on fast matrix multiplication is during the preconditioning stage. Accordingly, if we are given a matrix which is already approximately isotropic with all row norms approximately equal, then we can eliminate the usage of fast matrix multiplication and still prove the same time bound. Moreover, this method preserves the row-sparsity of \mathbf{A} . The primary difficulty of this approach is in computing an appropriate initialization when \mathbf{A} is only approximately isotropic. To resolve this issue, we use efficient ℓ_2 regression solvers that do not rely on fast matrix multiplication.

Theorem 2 *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{b} \in \mathbb{R}^n$ be such that the matrix-vector pair $[\mathbf{A} \ \mathbf{b}]$ satisfies $[\mathbf{A} \ \mathbf{b}]^T [\mathbf{A} \ \mathbf{b}] \approx_{O(1)} \mathbf{I}^4$ and for each row i of $[\mathbf{A} \ \mathbf{b}]$, $\|[\mathbf{A} \ \mathbf{b}]_{i,:}\|_2^2 \approx_{O(1)} d/n$. Assume $\|\mathbf{b}\|_2 \leq n^c$ and $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ is either 0 or bounded below by n^{-c} for some constant $c > 0$. Then for any $\epsilon > 0$, there is a routine that computes $\tilde{\mathbf{x}}$ such that with high probability,*

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon) \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$$

with a runtime of $O(\text{nnz}(\mathbf{A}) \log^2 n + s \cdot d^{1.5} \epsilon^{-2} \log^{1.5} n + d^2 \epsilon^{-2} \log^2 n)$, where s is the maximum number of entries in any row of \mathbf{A} .

The added assumption on $\|\mathbf{b}\|_2$ in Theorem 2 comes from the bounds of the ℓ_2 solvers. We prove Theorem 2 in Section 4.

Comparison of our results with previous work Our algorithms are significantly faster than the previous best SGD-based results for ℓ_1 regression, which took $O(\text{nnz}(\mathbf{A}) \log n + d^{4.5} \epsilon^{-2} \sqrt{\log d})$ time (Yang et al., 2016). Furthermore, our standard SGD bounds are especially likely to be achievable in practice. Table 1 compares the running time of our algorithm to the fastest gradient descent methods (Clarkson, 2005; Nesterov, 2009; Yang et al., 2016), interior point methods (Meng and Mahoney, 2013b; Lee and Sidford, 2015), and multiplicative weights update methods (Chin et al., 2013). Since we can apply sampling by Lewis weights prior to any algorithm⁵, we can replace n with $O(d\epsilon^{-2} \log n)$ and $\text{nnz}(\mathbf{A})$ with $O(d^2 \epsilon^{-2} \log n)$ at the cost of an additive $\tilde{O}(\text{nnz}(\mathbf{A}) + d^\omega)$ overhead for any running time bound in Table 1, where d^ω is time to multiply two $d \times d$ matrices.⁶

Note that we match the theoretical performance of the current best IPM (Lee and Sidford, 2015) in several regimes. In low to medium-precision ranges, for example $\epsilon \geq 10^{-3}$, both the best IPM and our algorithm have a running time of $\tilde{O}(\text{nnz}(\mathbf{A}) + d^{2.5})$. If all of the algorithms are implemented with naive matrix multiplication, Lee and Sidford (2015) takes $\tilde{O}(\text{nnz}(\mathbf{A}) + d^3)$ time, while we prove an identical running time for standard non-accelerated SGD with Lewis weights preconditioning. For general ϵ , if $\text{nnz}(\mathbf{A}) \geq d^2 \epsilon^{-2} \log n$, then Lee and Sidford (2015) will use Lewis weights sampling and both their algorithm and our algorithm will achieve a running time of $\tilde{O}(\text{nnz}(\mathbf{A}) + d^{2.5} \epsilon^{-2})$. This is significant because our SGD-based algorithms are likely to be far more practical⁷. Finally, we also note that in the setting where \mathbf{A} is approximately isotropic with approximately equal row norms and \mathbf{A} is row-sparse, with at most s non-zeros per row, our algorithm in Theorem 2 has the best dimensional-dependence out of any existing algorithm for

4. As defined in the notation section, we say that $\mathbf{A} \approx_\kappa \mathbf{B}$ if and only if $\frac{1}{\kappa} \mathbf{B} \preceq \mathbf{A} \preceq \kappa \mathbf{B}$.

5. Sampling \mathbf{A} by Lewis weights creates a dense $(d\epsilon^{-2} \log n) \times d$ matrix $\tilde{\mathbf{A}}$.

6. The current best value for ω is approximately 2.373 (Williams, 2012; Davie and Stothers, 2013; Le Gall, 2014).

7. Lewis weights preconditioning is also fast in practice

$s < d$. In particular, we beat Lee and Sidford (2015) whenever $s < d\epsilon^2$ or whenever $s < d$ and $\text{nnz}(\mathbf{A}) \geq d^2\epsilon^{-2}\log n$.

Solver	Running time ⁸
Subgradient descent Clarkson (2005)	$\tilde{O}(nd^5\epsilon^{-2}\log(1/\epsilon))$
Smoothing and gradient descent Nesterov (2009)	$\tilde{O}(n^{1.5}d\epsilon^{-1})$
Randomized IPCPM ⁹ Meng and Mahoney (2013b)	$\tilde{O}(nd^{\omega-1} + (\text{nnz}(\mathbf{A}) + \text{poly}(d))d\log(1/\epsilon))$
Multiplicative weights Chin et al. (2013)	$\tilde{O}(n^{1/3}(\text{nnz}(\mathbf{A}) + d^2)\epsilon^{-8/3})$
IPM and fast inverse maintenance Lee and Sidford (2015)	$\tilde{O}((\text{nnz}(\mathbf{A}) + d^2)\sqrt{d}\log(1/\epsilon))$
Weighted SGD Yang et al. (2016)	$\tilde{O}(\text{nnz}(\mathbf{A}) + d^{4.5}\epsilon^{-2})$
Lewis weights and SGD (this work)	$\tilde{O}(\text{nnz}(\mathbf{A}) + d^3\epsilon^{-2})$
Lewis weights and accelerated SGD (this work) ¹⁰	$\tilde{O}(\text{nnz}(\mathbf{A}) + d^{2.5}\epsilon^{-2})$

Table 1: Running time of several iterative ℓ_1 regression algorithms. All running times are to find a solution with ϵ relative error, with constant failure probability. Note that the first three algorithms in the table could also be sped up using the preconditioning method we use, i.e., Lewis weights row sampling (Cohen and Peng, 2015). Doing this would need a preconditioning time of $\tilde{O}(\text{nnz}(\mathbf{A}) + d^\omega)$, and enable us to use the fact that $n \leq O(d\epsilon^{-2}\log n)$ after preconditioning. However, this still does not make them faster than later algorithms.

Another related work by Bubeck et al. (2017) gives algorithms for ℓ_p regression for $p \in (1, \infty)$ that run in time $\tilde{O}(\text{nnz}(\mathbf{A}) \cdot n^{|1/2-1/p|} \log(1/\epsilon))$ and $\tilde{O}(\text{nnz}(\mathbf{A})n^{|1/2-1/p|-1/2}d^{1/2} + n^{|1/2-1/p|}d^2 + d^\omega) \log(1/\epsilon)$. They also use preconditioning and accelerated stochastic gradient descent techniques as a subroutine. However, they don't give bounds for ℓ_1 regression. Also, for p close to 1, these bounds are worse than ours. In contrast to Bubeck et al. (2017), our algorithms are specific to the ℓ_1 setting. We use the fact that the gradients are bounded for ℓ_1 regression, which doesn't hold for general ℓ_p regression. Moreover, our initialization using an ℓ_2 minimizer doesn't give good bounds for general ℓ_p regression. In this sense, our algorithms really leverage the special structure of the ℓ_1 problem.

1.1. Organization

The paper is organized as follows. Section 2 contains definitions and basic lemmas which we will use throughout the paper. Section 3 contains our main contribution, i.e., once we are given a suitably preconditioned matrix, it shows how we arrive at an approximate ℓ_1 minimizer within the claimed time bounds, for both non-accelerated and accelerated versions of stochastic gradient descent. Section 4 shows that if we restrict our input to slightly weaker preconditions, then we can eliminate the need for fast matrix multiplication to achieve the same time bounds. Appendix A contains the primary proof details for our main result in Section 3. In Appendix B, we show that row sampling using Lewis weights Cohen and Peng (2015), along with matrix rotation, suffices to

8. \tilde{O} hides terms polylogarithmic in d and n .

9. Interior Point Cutting Plane Methods

10. This running time only assumes that $\omega \leq 2.5$

give us a matrix satisfying our precondition requirements. Appendix C contains proof details from Section 4. Finally, Appendix D will give secondary and straightforward technical details for our main results that we include for completeness.

2. Preliminaries

In this section, we describe some of the notation and important definitions we use in the paper. We represent matrices and vectors using bold variables. We let $\mathbf{A}_{i,:}$ denote the i^{th} row of a matrix \mathbf{A} , and we use $\text{nnz}(\mathbf{A})$ to denote the number of non-zero elements in \mathbf{A} . \mathbf{A}^\dagger refers to the Moore-Penrose pseudoinverse of \mathbf{A} . When \mathbf{A} has linearly-independent columns, $\mathbf{A}^\dagger = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T$. Also, we assume that the input \mathbf{A} has full rank.

Definition 3 (ℓ_p -norm) The ℓ_p norm of a vector $\mathbf{v} \in \mathbb{R}^n$ is defined as

$$\|\mathbf{v}\|_p \stackrel{\text{def}}{=} (\sum_{i=1}^n \mathbf{v}_i^p)^{1/p}.$$

Accordingly, the ℓ_p norm of a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ is defined as

$$\|\mathbf{A}\|_p \stackrel{\text{def}}{=} \sup_{\mathbf{x} \in \mathbb{R}^d, \mathbf{x} \neq 0} \frac{\|\mathbf{A}\mathbf{x}\|_p}{\|\mathbf{x}\|_p}.$$

Definition 4 (Matrix approximation) We say that $\mathbf{A} \approx_\kappa \mathbf{B}$ if and only if

$$\frac{1}{\kappa} \mathbf{B} \preceq \mathbf{A} \preceq \kappa \mathbf{B}.$$

Here, \preceq refers to the Löwner partial ordering of matrices, where we say that $\mathbf{A} \preceq \mathbf{B}$ if $\mathbf{B} - \mathbf{A}$ is positive semi-definite.

Note that we also use \approx similarly in the case of scalars, as is commonplace.

Definition 5 (IRB) A matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with $n \geq d$ is said to be isotropic row-bounded (IRB) if the following hold:

1. $\mathbf{A}^T \mathbf{A} = \mathbf{I}$,
2. For all rows of \mathbf{A} , $\|\mathbf{A}_{i,:}\|_2^2 \leq O(d/n)$.

Definition 6 Given a matrix \mathbf{A} , we define the statistical leverage score of row $\mathbf{A}_{i,:}$ to be

$$\tau_i(\mathbf{A}) \stackrel{\text{def}}{=} \mathbf{A}_{i,:} (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}_{i,:}^T = \left\| (\mathbf{A}^T \mathbf{A})^{-1/2} \mathbf{A}_{i,:}^T \right\|_2^2.$$

Definition 7 For a matrix \mathbf{A} , the ℓ_1 Lewis weights $\bar{\mathbf{w}}$ are the unique weights such that for each row i we have

$$\bar{\mathbf{w}}_i = \tau_i \left(\bar{\mathbf{W}}^{-1/2} \mathbf{A} \right)$$

or equivalently

$$\bar{\mathbf{w}}^2 = \mathbf{A}_{i,:} \left(\mathbf{A}^T \bar{\mathbf{W}}^{-1} \mathbf{A} \right)^{-1} \mathbf{A}_{i,:}^T$$

where $\bar{\mathbf{W}}$ is the diagonal matrix formed by putting the elements of $\bar{\mathbf{w}}$ on the diagonal.

Definition 8 A function f is L -smooth if for any $x, y \in \mathbb{R}^d$, $\|\nabla f(x) - \nabla f(y)\|_2 \leq L\|x - y\|_2$.

Definition 9 A function f is σ -strongly convex if for any $x, y \in \mathbb{R}^d$,

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\sigma}{2}\|x - y\|_2^2.$$

Definition 10 A function f is G -Lipschitz continuous if for any $x, y \in \mathbb{R}^d$,

$$\|f(x) - f(y)\|_2 \leq G\|x - y\|_2.$$

3. Stochastic Gradient Descent for ℓ_1 Regression

In this section, we describe how we achieve the bounds in Theorem 1. We first introduce the preconditioning technique by [Cohen and Peng \(2015\)](#), which, along with rotating the matrix, will reduce our problem to ℓ_1 minimization where the input matrix \mathbf{A} is isotropic and the norms of all its rows have strong upper bounds, i.e. it is IRB by Definition 5. We relegate the details and proof of this preconditioning procedure to Appendix B.

In Section 3.2, we prove that known stochastic gradient descent algorithms will run provably faster if we assume that \mathbf{A} is IRB. In particular, if \mathbf{A} is IRB, we can find an initialization \mathbf{x}_0 that is close to the optimum \mathbf{x}^* , which in addition to bounding the gradient of our objective, will then allow us to plug these bounds into standard stochastic gradient descent algorithms and achieve a runtime of $\tilde{O}(nnz(\mathbf{A}) + d^3\epsilon^{-2})$. Finally, in Section 3.3 we take known smoothing techniques by [Allen-Zhu and Hazan \(2016\)](#) along with the Katyusha accelerated stochastic gradient descent by [Allen Zhu \(2017\)](#) to achieve a runtime of $\tilde{O}(nnz(\mathbf{A}) + d^{2.5}\epsilon^{-2})$.

3.1. Preconditioning with Lewis weights

The primary tool in our preconditioning routine will be a sampling scheme by Lewis weights, introduced in [Cohen and Peng \(2015\)](#), that was shown to approximately preserve the ℓ_1 norm. Specifically, we will use the combination of two primary theorems from [Cohen and Peng \(2015\)](#) that approximately compute the Lewis weights of a matrix quickly and then sample accordingly while still approximately preserving ℓ_1 norm distances with high probability.

Theorem 11 (Theorem 2.3 and 6.1 from [Cohen and Peng \(2015\)](#)) Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with ℓ_1 Lewis weights $\bar{\mathbf{w}}$ and an error parameter $\epsilon > 0$, then for any function $h(n, \epsilon) \geq O(\epsilon^{-2} \log n)$, we can find sampling values

$$\mathbf{p}_i \approx_{O(1)} \bar{\mathbf{w}}_i h(n, \epsilon)$$

for each $i \in \{1, 2, \dots, n\}$, and generate a matrix \mathbf{S} with $N = \sum_i \mathbf{p}_i$ rows, each chosen independently as the i^{th} standard basis vector of dimension n , times $\frac{1}{\mathbf{p}_i}$ with probability proportional to \mathbf{p}_i , such that with high probability we have

$$\|\tilde{\mathbf{A}}\mathbf{x}\|_1 \approx_{1+\epsilon} \|\mathbf{A}\mathbf{x}\|_1$$

for all $\mathbf{x} \in \mathbb{R}^d$, where $\tilde{\mathbf{A}} = \mathbf{S}\mathbf{A}$. Computing these sampling values requires $O(nnz(\mathbf{A}) \log n + d^\omega)$ time.

In Appendix B, we will show that this sampling scheme also ensures that each row of $\tilde{\mathbf{A}}$ has approximately the same leverage score. This proof will involve applying known facts about leverage scores and their connections with Lewis weights, along with standard matrix concentration bounds. Furthermore, we will obtain additional nice properties by rotating \mathbf{A} and showing that a solution to our reduced problem gives an approximate solution to the original problem, culminating in the following lemma:

Lemma 12 *There is a routine that takes a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, a vector $\mathbf{b} \in \mathbb{R}^n$ and $\epsilon > 0$, then produces a matrix $[\tilde{\mathbf{A}}, \tilde{\mathbf{b}}] \in \mathbb{R}^{N \times (d+1)}$ with $N = O(d\epsilon^{-2} \log n)$ and an invertible matrix $\mathbf{U} \in \mathbb{R}^{d \times d}$ such that matrix $\tilde{\mathbf{A}}\mathbf{U}$ is IRB and if $\tilde{\mathbf{x}}_{\mathbf{U}}^*$ minimizes $\|\tilde{\mathbf{A}}\mathbf{U}\mathbf{x} - \tilde{\mathbf{b}}\|_1$, then for any $\tilde{\mathbf{x}}$ such that*

$$\|\tilde{\mathbf{A}}\mathbf{U}\tilde{\mathbf{x}} - \tilde{\mathbf{b}}\|_1 \leq (1 + \delta)\|\tilde{\mathbf{A}}\mathbf{U}\tilde{\mathbf{x}}_{\mathbf{U}}^* - \tilde{\mathbf{b}}\|_1,$$

holds for some $\delta > 0$, we must have

$$\|\mathbf{A}(\mathbf{U}\tilde{\mathbf{x}}) - \mathbf{b}\|_1 \leq (1 + \epsilon)^2(1 + \delta)\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1$$

with high probability.

Furthermore, the full running time is $O(nnz(\mathbf{A}) \log n + d^{\omega-1} \min\{d\epsilon^{-2} \log n, n\} + \Upsilon)$ where $\Upsilon = \min\{d\epsilon^{-2} \log n, (d\epsilon^{-2} \log n)^{1/2+o(1)} + n \log^2 n\}$.

As a result, we will assume that all of our matrices \mathbf{A} are already in the same form as $\tilde{\mathbf{A}}\mathbf{U}$, i.e. we assume \mathbf{A} is IRB, since relative error guarantees for the preconditioned system apply to the original system.

3.2. Isotropic and Row-Bounded \mathbf{A} for Stochastic Gradient Descent

To demonstrate the usefulness of the properties of our preconditioned \mathbf{A} , we consider standard stochastic gradient descent and the bounds on its running time. We let $\mathbf{x}^* = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$. We will use the following theorem to prove runtime bounds for standard SGD:

Theorem 13 (Ruszczynski and Syski (1986)) *Given a function f and \mathbf{x}_0 such that $\|\mathbf{x}_0 - \mathbf{x}^*\|_2 \leq R$ and L is an upper bound on the ℓ_2 norm of the stochastic subgradients of f , then projected subgradient descent ensures that after t steps:*

$$f(\mathbf{x}_t) - f(\mathbf{x}^*) \leq O\left(\frac{RL}{\sqrt{t}}\right).$$

where $\mathbf{x}^* = \arg \min_{\mathbf{x}} f(\mathbf{x})$.

To use this theorem, we must prove bounds on the initialization distance $\|\mathbf{x}_0 - \mathbf{x}^*\|_2$ and the norm of the stochastic subgradients we use, i.e. $\nabla(n \cdot |\mathbf{A}_{i,:}\mathbf{x} - \mathbf{b}_i|)$ for each i . We show that both of these bounds come from our assumptions on \mathbf{A} .

Lemma 14 *If \mathbf{A} is IRB, then by setting $\mathbf{x}_0 = \mathbf{A}^T \mathbf{b}$ we have*

$$\|\mathbf{x}_0 - \mathbf{x}^*\|_2^2 \leq O(d/n) \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1^2.$$

Proof

$$\begin{aligned}
 \|\mathbf{x}^{(0)} - \mathbf{x}^{(*)}\|_2^2 &= \|\mathbf{A}^T \mathbf{b} - \mathbf{x}^{(*)}\|_2^2 \\
 &= \|\mathbf{A}^T (\mathbf{b} - \mathbf{A} \mathbf{x}^{(*)})\|_2^2 && \text{by assumption } \mathbf{A}^T \mathbf{A} = \mathbf{I} \\
 &= \|\sum_i \mathbf{A}_{i,:} (\mathbf{b} - \mathbf{A} \mathbf{x}^{(*)})\|_2^2 \\
 &\leq \|\mathbf{A} \mathbf{x}^* - \mathbf{b}\|_1^2 \cdot \max_i \|\mathbf{A}_{i,:}\|_2^2 && \text{by convexity of } \|\cdot\|_2, \text{ also shown in Lemma 40} \\
 &= O\left(\frac{d}{n}\right) \|\mathbf{A} \mathbf{x}^* - \mathbf{b}\|_1^2.
 \end{aligned}$$

■

Lemma 15 *If \mathbf{A} is IRB, then $\|\nabla(n \cdot |\mathbf{A}_{i,:} \mathbf{x} - \mathbf{b}_i|)\|_2^2 \leq O(nd)$ for all i .*

Proof We see that $\nabla(n \cdot |\mathbf{A}_{i,:} \mathbf{x} - \mathbf{b}_i|) = n \cdot \mathbf{A}_{i,:}^T \text{sgn}(\mathbf{A}_{i,:} \mathbf{x} - \mathbf{b}_i)$, and $\|\mathbf{A}_{i,:}\|_2^2 \leq O(d/n)$ for all i by our assumption that \mathbf{A} is IRB. This then implies our desired inequality. ■

These bounds, particularly the initialization distance, are stronger than the bounds for general \mathbf{A} , and together will give our first result that improves upon the runtime given by [Yang et al. \(2016\)](#) by using our preconditioning.

Theorem 16 *Given $\mathbf{A} \in \mathbb{R}^{n \times d}$, we can find $\tilde{\mathbf{x}} \in \mathbb{R}^d$ using preconditioning and stochastic gradient descent such that*

$$\|\mathbf{A} \tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon) \|\mathbf{A} \mathbf{x}^* - \mathbf{b}\|_1$$

in time $O(nnz(\mathbf{A}) \log^2 n + d^3 \epsilon^{-2} \log n)$.

Proof By preconditioning with Lemma 12 and error $O(\epsilon)$ we obtain an $N \times d$ matrix $\tilde{\mathbf{A}} \mathbf{U}$ in time $O(nnz(\mathbf{A}) \log n + d^\omega \epsilon^{-2})$.

By Theorem 13, we then need to run $O(d^2 \epsilon^{-2})$ iterations of standard stochastic gradient descent to achieve absolute error of $O(\epsilon \cdot f(\mathbf{x}^*))$ which is equivalent to relative error of $O(\epsilon)$. The required runtime is then $O(d^3 \epsilon^{-2})$. Technically, the input to stochastic gradient descent will require the value R , i.e. the upper bound on initialization distance, which requires access to a constant factor approximation of $f(\mathbf{x}^*)$. We will show in Appendix D.2 that we can assume that we have such an approximation at the cost of a factor of $\log n$ in the running time.

Combining the preconditioning and stochastic gradient descent will produce $\tilde{\mathbf{x}}$ with $O(\epsilon)$ relative error to the optimal objective function value in time $O(nnz(\mathbf{A}) \log n + d^3 \epsilon^{-2})$. Adding the factor $\log n$ overhead from estimating $f(\mathbf{x}^*)$ gives the desired runtime. ■

3.3. Smoothing Reductions and Katyusha Accelerated SGD

In order to improve the running time guarantees, we consider whether our strong initialization distance bound will allow us to apply black-box accelerated stochastic gradient descent methods. These methods generally require smoothness and strong convexity of the objective function, neither of which are necessarily true for our objective function. Previous results ([Nesterov, 2005b, 2007](#);

Duchi et al., 2012; Ouyang and Gray, 2012; Allen-Zhu and Hazan, 2016) have addressed this general issue and given reductions from certain classes of objective functions to similar functions with smoothness and strong convexity, while still maintaining specific error and runtime guarantees. Accordingly, we will first show how our initialization distance fits into the reduction of Allen-Zhu and Hazan (2016), then apply Katyusha’s accelerated gradient descent algorithm by Allen Zhu (2017) to their framework. We state the theorem below, and relegate the details of its proof to Appendix A.

Theorem 1 *Given $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$, assume $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ is either 0 or bounded below by n^{-c} for some constant $c > 0$. Then for any $\epsilon > 0$, there is a routine that outputs $\tilde{\mathbf{x}}$ such that with high probability,¹¹*

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon) \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$$

with a runtime of $O(\text{nnz}(\mathbf{A}) \log^2 n + d^{2.5} \epsilon^{-2} \log^{1.5} n)$ whenever $n \geq d\epsilon^{-2} \log n$, and a runtime of $O(\sqrt{nd} \epsilon^{-1} \log n + nd^{\omega-1} \log n)$ when $n \leq d\epsilon^{-2} \log n$.

The last bound in Theorem 1 may seem odd, as Lewis weights sampling produces a matrix with row-dimension $d\epsilon^{-2} \log n$, which will in fact hurt our running time if $n \ll d\epsilon^{-2} \log n$. Moreover, we cannot simply avoid running the Lewis weights sampling because our algorithm critically relies on the resulting leverage score properties. Instead, if $n \ll d\epsilon^{-2} \log n$, we can simulate the sampling procedure in $O(n)$ time and keep a count for each of the n unique rows. Since the simulated sample matrix will look like the original but with duplicated rows, we will be able to carry out the rest of our linear algebraic manipulations in time dependent on n rather than $d\epsilon^{-2} \log n$. We will address this in Appendix D.1.

4. Row-Sparsity Bounds for ℓ_1 Regression

In this section, we explain how to avoid using matrix multiplication, which we use in Lemma 12, to get an IRB matrix, and in Lemma 14, to get a good initialization. To avoid both procedures, we assume that our given matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and vector $\mathbf{b} \in \mathbb{R}^n$ are such that $[\mathbf{A} \ \mathbf{b}]^T [\mathbf{A} \ \mathbf{b}] \approx_{O(1)} \mathbf{I}$ and for each row i of $[\mathbf{A} \ \mathbf{b}]$ we have $\|[\mathbf{A} \ \mathbf{b}]_{i,:}\|_2^2 \approx_{O(1)} d/n$. Notice that these conditions imply $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$ and $\|\mathbf{A}_{i,:}\|_2^2 \leq O(d/n)$, which are nearly the properties of the preconditioned matrix $\tilde{\mathbf{A}}\mathbf{U}$ generated in Lemma 12. However, we still face two new complications: (1) the row count of matrix \mathbf{A} has not been reduced from n to $O(d\epsilon^{-2} \log n)$, and (2) \mathbf{A} is only approximately isotropic.

We will account for the first complication by showing that, under these conditions, the Lewis weights are approximately equal, which implies that uniform row sampling is nearly equivalent to that in Theorem 11. We then describe how to find a good initialization point using conjugate gradient methods when $\tilde{\mathbf{A}}$ is only approximately isotropic. Finally, we use the reduction in Allen-Zhu and Hazan (2016) and the Katyusha stochastic gradient descent algorithm from Allen Zhu (2017) to achieve a total running time of $\tilde{O}(\text{nnz}(\mathbf{A}) + sd^{1.5} \epsilon^{-2} + d^2 \epsilon^{-2})$, where s is the maximum number of non-zeros in a row of \mathbf{A} . Note that as a byproduct of this algorithm, we achieve row-sparsity-preserving ℓ_1 minimization.

11. Throughout this paper, we let “with high probability” mean that ξ is the failure probability and our runtime has dependence on $\log(1/\xi)$, which we ignore for ease of notation.

UNIFORM SAMPLING OF \mathbf{A}

We deal with the first complication of avoiding Lemma 12 by sampling \mathbf{A} uniformly. In particular, if we uniformly sample $N = O(d\epsilon^{-2} \log d)$ rows of $[\mathbf{A} \ \mathbf{b}]$, this yields a smaller matrix $[\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]$ such that $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \approx_{O(1)} \mathbf{I}$ and $\|\tilde{\mathbf{A}}_{i,:}\|_2^2 \leq O(d/N)$ for each row i , which is to say $\tilde{\mathbf{A}}$ is almost IRB. This then culminates in the following lemma whose proof we relegate to Appendix C.1.

Lemma 17 *Suppose we are given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ such that $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$ and $\|\mathbf{A}_{i,:}\|_2^2 \approx_{O(1)} d/n$. If we uniformly sample $N = O(d\epsilon^{-2} \log n)$ rows independently and rescale each row by n/N to obtain matrix $\tilde{\mathbf{A}}$, then with high probability the following properties hold:*

1. $\|\mathbf{A}\mathbf{x}\|_1 \approx_{1+\epsilon} \|\tilde{\mathbf{A}}\mathbf{x}\|_1$ for all $\mathbf{x} \in \mathbb{R}^d$.
2. $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \approx_{O(1)} \left(\frac{n}{N}\right) \mathbf{I}$.
3. $\|\tilde{\mathbf{A}}_{i,:}\|_2^2 \approx_{O(1)} dn/N^2$ for all rows $i \in \{1, 2, \dots, N\}$.

INITIALIZATION USING APPROXIMATE ℓ_2 MINIMIZER

It now remains to show that we can still find a good initialization \mathbf{x}_0 for gradient descent even with our relaxed assumptions on \mathbf{A} . Previously, when we had $\mathbf{A}^T \mathbf{A} = \mathbf{I}$, we used $\mathbf{x}_0 = \mathbf{A}^T \mathbf{b} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$. It turns out that for $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$, the ℓ_2 minimizer $\mathbf{x}_0 = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ is still a good initialization point. But finding an exact ℓ_2 minimizer would take a prohibitive amount of time or would require matrix multiplication. However, an approximate ℓ_2 minimizer suffices, and we can find such a point quickly using the conjugate gradient method.

For this section, we define $\mathbf{x}^* \stackrel{\text{def}}{=} \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$ and $\mathbf{x}_0 \stackrel{\text{def}}{=} \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$. Our main result is the following:

Lemma 18 (Approximate ℓ_2 minimizer is close to \mathbf{x}^*) *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be such that $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$ and for each row i of \mathbf{A} , $\|\mathbf{A}_{i,:}\|_2^2 \leq O(d/n)$. Assume that $\|\mathbf{b}\|_2 \leq n^c$ and $\|\mathbf{A}\mathbf{x}_0 - \mathbf{b}\|_2 \geq n^{-c}$ for some constant $c > 0$.*

$$\|\tilde{\mathbf{x}}_0 - \mathbf{x}^*\|_2 \leq O(\sqrt{d/n}) \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1.$$

Moreover, $\tilde{\mathbf{x}}_0$ can be computed in $O((t_{\mathbf{A}^T \mathbf{A}} + d) \log(n/\epsilon))$ time, where $t_{\mathbf{A}^T \mathbf{A}}$ denotes the time to multiply a vector by $\mathbf{A}^T \mathbf{A}$.

We prove Lemma 18 using the following two lemmas whose proofs are deferred until Appendix C.2. Lemma 19 shows that the ℓ_2 minimizer is close to the ℓ_1 minimizer even when \mathbf{A} is only approximately isotropic, and the proof is similar to that of Lemma 14. Lemma 20 shows that we can find a good estimate for the ℓ_2 minimizer.

Lemma 19 (Exact ℓ_2 minimizer is close to \mathbf{x}^*) *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be such that $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$ and for each row i of \mathbf{A} , $\|\mathbf{A}_{i,:}\|_2^2 \leq O(d/n)$. Then*

$$\|\mathbf{x}_0 - \mathbf{x}^*\|_2 \leq O(\sqrt{d/n}) \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1.$$

Lemma 20 (Conjugate gradient finds an approximate ℓ_2 minimizer) *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be such that $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$ and for each row i of \mathbf{A} , $\|\mathbf{A}_{i,:}\|_2^2 \leq O(d/n)$. Assume that $\|\mathbf{b}\|_2 \leq n^c$ and $\|\mathbf{A}\mathbf{x}_0 - \mathbf{b}\|_2 \geq n^{-c}$. Then for any $\epsilon > 0$, the conjugate gradient method can find an $\tilde{\mathbf{x}}_0$ such that*

$$\|\tilde{\mathbf{x}}_0 - \mathbf{x}_0\|_2 \leq \epsilon \|\mathbf{A}\mathbf{x}_0 - \mathbf{b}\|_2.$$

Moreover, $\tilde{\mathbf{x}}_0$ can be found in time $O((t_{\mathbf{A}^T \mathbf{A}} + d) \log(n/\epsilon))$, where $t_{\mathbf{A}^T \mathbf{A}}$ is the time to multiply a vector by $\mathbf{A}^T \mathbf{A}$.

Proof [Proof of Lemma 18] We use conjugate gradient with $\epsilon = \sqrt{d/n}$ to find an $\tilde{\mathbf{x}}_0$ in $O((t_{\mathbf{A}^T \mathbf{A}} + d) \log(n/\epsilon))$ time by Lemma 20 to achieve our desired initialization time bounds. Note that by definition of \mathbf{x}_0 and by a standard norm inequality, we have:

$$\|\mathbf{A}\mathbf{x}_0 - \mathbf{b}\|_2 \leq \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_2 \leq \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1$$

Then by the triangle inequality and Lemma 19 we have:

$$\begin{aligned} \|\tilde{\mathbf{x}}_0 - \mathbf{x}^*\|_2 &\leq \|\tilde{\mathbf{x}}_0 - \mathbf{x}_0\|_2 + \|\mathbf{x}_0 - \mathbf{x}^*\|_2 \\ &\leq \sqrt{d/n} \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1 + O(\sqrt{d/n}) \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1 \end{aligned}$$

which gives our desired initialization accuracy. ■

FAST ROW-SPARSITY-PRESERVING ℓ_1 MINIMIZATION

Finally, we combine the matrix achieved by uniform sampling in Lemma 17 and the initialization from Lemma 18 to achieve fast row-sparsity-preserving ℓ_1 minimization. This then gives the following main theorem that we prove in Appendix C.3.

Theorem 2 *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{b} \in \mathbb{R}^n$ be such that the matrix-vector pair $[\mathbf{A} \ \mathbf{b}]$ satisfies $[\mathbf{A} \ \mathbf{b}]^T [\mathbf{A} \ \mathbf{b}] \approx_{O(1)} \mathbf{I}$ ¹² and for each row i of $[\mathbf{A} \ \mathbf{b}]$, $\|[\mathbf{A} \ \mathbf{b}]_{i,:}\|_2^2 \approx_{O(1)} d/n$. Assume $\|\mathbf{b}\|_2 \leq n^c$ and $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ is either 0 or bounded below by n^{-c} for some constant $c > 0$. Then for any $\epsilon > 0$, there is a routine that computes $\tilde{\mathbf{x}}$ such that with high probability,*

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon) \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$$

with a runtime of $O(\text{nnz}(\mathbf{A}) \log^2 n + s \cdot d^{1.5} \epsilon^{-2} \log^{1.5} n + d^2 \epsilon^{-2} \log^2 n)$, where s is the maximum number of entries in any row of \mathbf{A} .

Acknowledgments

We are grateful to Richard Peng for many insightful discussions. We would also like to thank Michael Cohen for his helpful comments and Junxing Wang for his feedback. We thank the anonymous reviews of an earlier version of this manuscript for their very valuable comments, especially with regards to the performances of the many existing algorithms in the ℓ_1 case.

12. As defined in the notation section, we say that $\mathbf{A} \approx_{\kappa} \mathbf{B}$ if and only if $\frac{1}{\kappa} \mathbf{B} \preceq \mathbf{A} \preceq \kappa \mathbf{B}$.

References

- Zeyuan Allen Zhu. Katyusha: the first direct acceleration of stochastic gradient methods. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 1200–1205, 2017. doi: 10.1145/3055399.3055448. URL <http://doi.acm.org/10.1145/3055399.3055448>.
- Zeyuan Allen-Zhu and Elad Hazan. Optimal black-box reductions between optimization objectives. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pages 1614–1622, USA, 2016. Curran Associates Inc. ISBN 978-1-5108-3881-9. URL <http://dl.acm.org/citation.cfm?id=3157096.3157277>.
- Sébastien Bubeck, Michael B Cohen, Yin Tat Lee, and Yuanzhi Li. An homotopy method for ℓ_p regression provably beyond self-concordance and in input-sparsity time. *arXiv preprint arXiv:1711.01328*, 2017. Available at: <https://arxiv.org/abs/1711.01328>.
- Scott Shaobing Chen, David L. Donoho, and Michael A. Saunders. Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159, January 2001. ISSN 0036-1445. doi: 10.1137/S003614450037906X. URL <http://dx.doi.org/10.1137/S003614450037906X>.
- Hui Han Chin, Aleksander Madry, Gary L. Miller, and Richard Peng. Runtime guarantees for regression problems. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science, ITCS '13*, pages 269–282, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1859-4. doi: 10.1145/2422436.2422469. URL <http://doi.acm.org/10.1145/2422436.2422469>.
- Kenneth L. Clarkson. Subgradient and sampling algorithms for ℓ_1 regression. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05*, pages 257–266, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics. ISBN 0-89871-585-7. URL <http://dl.acm.org/citation.cfm?id=1070432.1070469>.
- Kenneth L. Clarkson, Petros Drineas, Malik Magdon-Ismael, Michael W. Mahoney, Xiangrui Meng, and David P. Woodruff. The fast cauchy transform and faster robust linear regression. In *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '13*, pages 466–477, Philadelphia, PA, USA, 2013. Society for Industrial and Applied Mathematics. ISBN 978-1-611972-51-1. URL <http://dl.acm.org/citation.cfm?id=2627817.2627851>.
- Michael B. Cohen and Richard Peng. ℓ_p row sampling by lewis weights. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing, STOC '15*, pages 183–192, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3536-2. doi: 10.1145/2746539.2746567. URL <http://doi.acm.org/10.1145/2746539.2746567>.
- Michael B. Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15*, pages 181–190, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3333-7. doi: 10.1145/2688073.2688113. URL <http://doi.acm.org/10.1145/2688073.2688113>.

- Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W. Mahoney. Sampling algorithms and coresets for ℓ_p regression. *SIAM J. Comput.*, 38(5):2060–2078, February 2009. ISSN 0097-5397. doi: 10.1137/070696507. URL <http://dx.doi.org/10.1137/070696507>.
- A. M. Davie and A. J. Stothers. Improved bound for complexity of matrix multiplication. *Proceedings of the Royal Society of Edinburgh: Section A Mathematics*, 143(2):351–369, 2013. doi: 10.1017/S0308210511001648.
- John C Duchi, Peter L Bartlett, and Martin J Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- Martín Farach-Colton and Meng-Tsung Tsai. Exact sublinear binomial sampling. *Algorithmica*, 73(4):637–651, December 2015. ISSN 0178-4617. doi: 10.1007/s00453-015-0077-8. URL <http://dx.doi.org/10.1007/s00453-015-0077-8>.
- F. G. Foster. On the stochastic matrices associated with certain queuing processes. *The Annals of Mathematical Statistics*, 24(3):355–360, 1953. ISSN 00034851. URL <http://www.jstor.org/stable/2236286>.
- Nick Harvey. Matrix concentration and sparsification. Workshop on “Randomized Numerical Linear Algebra (RandNLA): Theory and Practice”, 2012. Available at: http://www.drineas.org/RandNLA/slides/Harvey_RandNLA@FOCS_2012.pdf.
- François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, ISSAC ’14, pages 296–303, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2501-1. doi: 10.1145/2608628.2608664. URL <http://doi.acm.org/10.1145/2608628.2608664>.
- Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, FOCS ’15, pages 230–249, Washington, DC, USA, 2015. IEEE Computer Society. ISBN 978-1-4673-8191-8. doi: 10.1109/FOCS.2015.23. URL <http://dx.doi.org/10.1109/FOCS.2015.23>.
- D. Lewis. Finite dimensional subspaces of ℓ_p . *Studia Mathematica*, 63(2):207–212, 1978. URL <http://eudml.org/doc/218208>.
- Xiangrui Meng and Michael W. Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC ’13, pages 91–100, New York, NY, USA, 2013a. ACM. ISBN 978-1-4503-2029-0. doi: 10.1145/2488608.2488621. URL <http://doi.acm.org/10.1145/2488608.2488621>.
- Xiangrui Meng and Michael W. Mahoney. Robust regression on mapreduce. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML’13, pages III–888–III–896. JMLR.org, 2013b. URL <http://dl.acm.org/citation.cfm?id=3042817.3043036>.

- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM J. on Optimization*, 19(4):1574–1609, January 2009. ISSN 1052-6234. doi: 10.1137/070704277. URL <http://dx.doi.org/10.1137/070704277>.
- Yu. Nesterov and J. Ph. Vial. Confidence level solutions for stochastic programming. *Automatica*, 44(6):1559–1568, June 2008. ISSN 0005-1098. doi: 10.1016/j.automatica.2008.01.017. URL <http://dx.doi.org/10.1016/j.automatica.2008.01.017>.
- Yurii Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27, pages 372–376, 1983.
- Yurii Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM Journal on Optimization*, 16(1):235–249, May 2005a. ISSN 1052-6234. doi: 10.1137/S1052623403422285. URL <http://dx.doi.org/10.1137/S1052623403422285>.
- Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, May 2005b. ISSN 0025-5610. doi: 10.1007/s10107-004-0552-5. URL <http://dx.doi.org/10.1007/s10107-004-0552-5>.
- Yurii Nesterov. Smoothing technique and its applications in semidefinite optimization. *Mathematical Programming*, 110(2):245–259, March 2007. ISSN 0025-5610. doi: 10.1007/s10107-006-0001-8. URL <http://dx.doi.org/10.1007/s10107-006-0001-8>.
- Yurii Nesterov. Unconstrained convex minimization in relative scale. *Mathematics of Operations Research*, 34(1):180–193, February 2009. ISSN 0364-765X. doi: 10.1287/moor.1080.0348. URL <http://dx.doi.org/10.1287/moor.1080.0348>.
- Hua Ouyang and Alexander Gray. Stochastic smoothing for nonsmooth minimizations: Accelerating SGD by exploiting structure. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, ICML’12, pages 1523–1530, USA, 2012. Omnipress. ISBN 978-1-4503-1285-1. URL <http://dl.acm.org/citation.cfm?id=3042573.3042768>.
- Stephen Portnoy. On computation of regression quantiles: Making the laplacian tortoise faster. *Lecture Notes-Monograph Series*, 31:187–200, 1997. ISSN 07492170. URL <http://www.jstor.org/stable/4355977>.
- Stephen Portnoy and Roger Koenker. The gaussian hare and the laplacian tortoise: computability of squared-error versus absolute-error estimators. *Statistical Science*, 12(4):279–300, 11 1997. doi: 10.1214/ss/1030037960. URL <http://dx.doi.org/10.1214/ss/1030037960>.
- Andrzej Ruszczyński and Wojciech Syski. On convergence of the stochastic subgradient method with on-line stepsize rules. *Journal of Mathematical Analysis and Applications*, 114(2):512 – 527, 1986. ISSN 0022-247X. doi: [http://dx.doi.org/10.1016/0022-247X\(86\)90104-6](http://dx.doi.org/10.1016/0022-247X(86)90104-6). URL <http://www.sciencedirect.com/science/article/pii/0022247X86901046>.
- Sushant Sachdeva and Nisheeth K. Vishnoi. Faster algorithms via approximation theory. *Foundations and Trends® in Theoretical Computer Science*, 9(2):125–210, March 2014. ISSN 1551-305X. doi: 10.1561/04000000065. URL <http://dx.doi.org/10.1561/04000000065>.

Christian Sohler and David P. Woodruff. Subspace embeddings for the ℓ_1 -norm with applications. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing, STOC '11*, pages 755–764, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0691-1. doi: 10.1145/1993636.1993736. URL <http://doi.acm.org/10.1145/1993636.1993736>.

Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, August 2012. ISSN 1615-3375. doi: 10.1007/s10208-011-9099-z. URL <http://dx.doi.org/10.1007/s10208-011-9099-z>. Available at <http://arxiv.org/abs/1004.4389>.

Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing, STOC '12*, pages 887–898, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1245-5. doi: 10.1145/2213977.2214056. URL <http://doi.acm.org/10.1145/2213977.2214056>.

David P. Woodruff and Qin Zhang. Subspace embeddings and ℓ_p -regression using exponential random variables. In *COLT 2013 - The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA*, pages 546–567, 2013. URL <http://jmlr.org/proceedings/papers/v30/Woodruff13.html>.

Jiyan Yang, Yin-Lam Chow, Christopher Ré, and Michael W. Mahoney. Weighted SGD for ℓ_p regression with randomized preconditioning. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '16*, pages 558–569, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics. ISBN 978-1-611974-33-1. URL <http://dl.acm.org/citation.cfm?id=2884435.2884476>.

Appendix A. Proof of Theorem 1

In this section, we prove our primary result, stated in Theorem 1. Specifically, we further examine whether our strong initialization distance bound will allow us to improve the running time with black-box accelerated stochastic gradient descent methods. The first step towards this is to apply smoothing reductions to our objective function.

A.1. Smoothing the Objective Function and Adding Strong Convexity

As before, we let $\mathbf{x}^* = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$. For clarity, we will borrow some of the notation from Allen-Zhu and Hazan (2016) to more clearly convey their black-box reductions.

Definition 21 A function $f(x)$ is (L, σ) -smooth-sc if it is both L -smooth and σ -strongly-convex.

Definition 22 An algorithm $\mathcal{A}(f(x), x_0)$ is a $\text{TIME}_{\mathcal{A}}(L, \sigma)$ -minimizer if $f(x)$ is (L, σ) -smooth-sc and $\text{TIME}_{\mathcal{A}}(L, \sigma)$ is the time it takes \mathcal{A} to produce x' such that $f(x') - f(x^*) \leq \frac{f(x_0) - f(x^*)}{4}$ for any starting x_0 .

Allen-Zhu and Hazan assume access to efficient $\text{TIME}_{\mathcal{A}}(L, \sigma)$ -minimizer algorithms, and show how a certain class of objective functions can be slightly altered to meet the smoothness and strong convexity conditions to apply these algorithms without losing too much in terms of error and run-time.

Theorem 23 (Theorem C.2 from Allen-Zhu and Hazan (2016)) Consider the problem of minimizing an objective function

$$f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$$

such that each $f_i(\cdot)$ is G -Lipschitz continuous. Let \mathbf{x}_0 be a starting vector such that $f(\mathbf{x}_0) - f(\mathbf{x}^*) \leq \Delta$ and $\|\mathbf{x}_0 - \mathbf{x}^*\|_2^2 \leq \Theta$. Then there is a routine that takes as input a $\text{TIME}_{\mathcal{A}}(L, \sigma)$ -minimizer, \mathcal{A} , alongside $f(x)$ and x_0 , with $\beta_0 = \Delta/G^2$, $\sigma_0 = \Delta/\Theta$ and $T = \log_2(\Delta/\epsilon)$, and produces x_T satisfying $f(x_T) - f(x^*) \leq O(\epsilon)$ in total running time

$$\sum_{t=0}^{T-1} \text{TIME}_{\mathcal{A}}(2^t/\beta_0, \sigma_0 \cdot 2^{-t}).$$

It is then straightforward to show that our objective function fits the necessary conditions to utilize Theorem 23.

Lemma 24 If \mathbf{A} is IRB, then the function $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$ can be written as $\frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x})$ such that each $f_i(\cdot)$ is $O(\sqrt{nd})$ -Lipschitz continuous.

Proof By the definition of 1-norm,

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1 = \sum_{i=1}^n |\mathbf{A}_{i,:}\mathbf{x} - \mathbf{b}_i|.$$

We then set $f_i(\mathbf{x}) = n \cdot |\mathbf{A}_{i,:}\mathbf{x} - \mathbf{b}_i|$ and the result follows from Lemma 15. ■

We can then incorporate our objective into the routine from Theorem 23, along with our initialization of \mathbf{x}_0 .

Lemma 25 Let \mathcal{A} be a $\text{TIME}_{\mathcal{A}}(L, \sigma)$ -minimizer, along with objective $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$ such that \mathbf{A} is IRB and $\mathbf{x}_0 = \mathbf{A}^T \mathbf{b}$, then the routine from Theorem 23 produces \mathbf{x}_T satisfying $f(\mathbf{x}_T) - f(\mathbf{x}^*) \leq O(\epsilon)$ in total running time

$$\sum_{t=0}^{T-1} \text{TIME}_{\mathcal{A}} \left(O \left(\frac{nd2^{2t}}{\Delta} \right), O \left(\frac{n\Delta}{d \cdot f(\mathbf{x}^*)^2 2^{2t}} \right) \right).$$

Proof Lemma 24 implies that we can apply Theorem 23 where $G = O(\sqrt{nd})$, and Lemma 14 gives $\Theta = O(\frac{d}{n} f(\mathbf{x}^*)^2)$. We then obtain $\beta_0 = O(\frac{\Delta}{nd})$ and $\sigma_0 = O(\frac{n\Delta}{df(\mathbf{x}^*)^2})$ and substitute these values in the running time of Theorem 23. ■

A.2. Applying Katyusha Accelerated SGD

Now that we have shown how our initialization can be plugged into the smoothing construction of Allen-Zhu and Hazan (2016), we simply need an efficient $\text{TIME}_{\mathcal{A}}(L, \sigma)$ -minimizer to obtain all the necessary pieces to prove our primary result.

Theorem 26 (Corollary 3.8 in Allen Zhu (2017)) *There is a routine that is a $\text{TIME}_{\mathcal{A}}(L, \sigma)$ -minimizer where $\text{TIME}_{\mathcal{A}}(L, \sigma) = d \cdot O(n + \sqrt{nL/\sigma})$.*

We can then precondition the matrix to give our strong bounds on the initialization distance of \mathbf{x}_0 from the optimal \mathbf{x}^* , which allows us to apply the smoothing reduction and Katyusha accelerated gradient descent more efficiently.

Theorem 1 *Given $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$, assume $\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2$ is either 0 or bounded below by n^{-c} for some constant $c > 0$. Then for any $\epsilon > 0$, there is a routine that outputs $\tilde{\mathbf{x}}$ such that with high probability,¹³*

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon) \min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_1$$

with a runtime of $O(nnz(\mathbf{A}) \log^2 n + d^{2.5} \epsilon^{-2} \log^{1.5} n)$ whenever $n \geq d\epsilon^{-2} \log n$, and a runtime of $O(\sqrt{nd^2} \epsilon^{-1} \log n + nd^{\omega-1} \log n)$ when $n \leq d\epsilon^{-2} \log n$.

Proof Once again, by preconditioning with Lemma 12 and error $O(\epsilon)$ we obtain a matrix $\tilde{\mathbf{A}} \mathbf{U} \in \mathbb{R}^{N \times d}$ and a vector $\tilde{\mathbf{b}} \in \mathbb{R}^n$ in time $O(nnz(\mathbf{A}) \log n + d^{\omega-1} \min\{d\epsilon^{-2} \log n, n\})$. We utilize the routine in Theorem 26 as the $\text{TIME}_{\mathcal{A}}(L, \sigma)$ -minimizer for Lemma 25, and plug the time bounds in to achieve an absolute error of $O(\delta)$ in the preconditioned objective function with the following running time:

$$\begin{aligned} d \cdot O\left(\sum_{t=0}^{T-1} N + \sqrt{N \cdot \left(\frac{Nd2^t}{\Delta}\right) \left(\frac{d \cdot f(x^*)^{22t}}{N\Delta}\right)}\right) &= d \cdot O\left(N \log \frac{\Delta}{\delta} + \frac{d\sqrt{N} \cdot f(x^*)}{\Delta} \sum_{t=0}^{T-1} 2^t\right) \\ &= d \cdot O\left(N \log \frac{\Delta}{\delta} + \frac{d\sqrt{N} \cdot f(x^*)}{\delta}\right). \end{aligned}$$

To achieve our desired relative error of ϵ we need to set $\delta = O(\epsilon f(x^*))$. Technically, this means that the input to gradient descent will require at least a constant factor approximation to $f(x^*)$. We will show in Appendix D.2 that we can assume that we have such an approximation at the cost of a factor of $\log n$ in the running time. We assume that $f(x^*) \geq n^{-c}$ for some fixed constant¹⁴ c in order to upper bound $\log \frac{\Delta}{\delta}$, which gives a runtime of $O\left(dN \log(n\epsilon^{-1}) + \frac{d\sqrt{N}}{\epsilon}\right)$.

Here, we used the fact that $N = O(d\epsilon^{-2} \log n)$, but can also assume that computationally, $N \leq n$, as will be addressed in Appendix D.1. This gives a runtime of

$$O\left(\min\{d^{2.5} \epsilon^{-2} \sqrt{\log n}, nd \log(n/\epsilon) + \sqrt{nd^2} \epsilon^{-1}\}\right),$$

which, combined with our preconditioning runtime (where Υ is a lower order term if we assume the ϵ is at most polynomially small in n) and the factor $\log n$ overhead from estimating $f(x^*)$ which we address in Appendix D.2, gives the desired runtime. Furthermore, since the error in our preconditioning was $O(\epsilon)$, by Lemma 12 we have achieved a solution with $O(\epsilon)$ relative error in the original problem. ■

13. Throughout this paper, we let “with high probability” mean that ξ is the failure probability and our runtime has dependence on $\log(1/\xi)$, which we ignore for ease of notation.

14. Note that if $f(x^*) = 0$, then our initialization $\mathbf{x}_0 = \mathbf{A}^T \mathbf{b}$ will be equal to \mathbf{x}^* .

Appendix B. Preconditioning with Lewis Weights and Rotation

In this section, we show how to precondition a given matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ into a “good” matrix, primarily using techniques by [Cohen and Peng \(2015\)](#), and will ultimately prove Lemma 12. Recall that our overall goal was to efficiently transform \mathbf{A} into a matrix $\tilde{\mathbf{A}}$ such that the ℓ_1 norm is approximately maintained for all \mathbf{x} , along with $\tilde{\mathbf{A}}$ being isotropic and having all row norms approximately equal.

Accordingly, our preconditioning will be done in the following two primary steps:

1. We sample $N = O(d\epsilon^{-2} \log d)$ rows from \mathbf{A} according to Lewis weights ([Cohen and Peng, 2015](#)) to construct a matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times d}$. The guarantees of [Cohen and Peng \(2015\)](#) ensure that for all $\mathbf{x} \in \mathbb{R}^d$, $\|\tilde{\mathbf{A}}\mathbf{x}\|_1 \approx_{1+\epsilon} \|\mathbf{A}\mathbf{x}\|_1$ with high probability. We then further show that this sampling scheme gives $\tau_i(\tilde{\mathbf{A}}) = O(d/N)$ for all $1 \leq i \leq N$ with high probability.
2. We then find an invertible matrix \mathbf{U} such that $\tilde{\mathbf{A}}\mathbf{U}$ still has the two necessary properties from Lewis weight sampling and is also isotropic.

The matrix $\tilde{\mathbf{A}}\mathbf{U}$ then has all the prerequisite properties to run our ℓ_1 minimization algorithms, and it only becomes necessary to show that running an ℓ_1 -minimization routine on $\tilde{\mathbf{A}}\mathbf{U}$ will help us find an approximate solution to the original problem.

In Appendix B.1, we show that Lewis weight sampling gives a matrix with approximately equal leverage scores. In Appendix B.2, we find the invertible matrix \mathbf{U} that makes $\tilde{\mathbf{A}}\mathbf{U}$ isotropic while preserving other properties. In Appendix B.3, we show that an approximate solution with respect to the preconditioned matrix will give an approximate solution with respect to the original matrix. Finally, we prove our primary preconditioning result, Lemma 12, in Appendix B.4.

Before we do this, the following facts are useful.

Fact 27 (Foster’s theorem ([Foster, 1953](#))) For a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$,

$$\sum_{i=1}^n \tau_i(\mathbf{A}) = d.$$

Fact 28 (Lemma 2 in [Cohen et al. \(2015\)](#)) Given a matrix \mathbf{A} , for all rows i ,

$$\tau_i(\mathbf{A}) = \min_{\mathbf{A}^T \mathbf{x} = \mathbf{A}_{i,:}^T} \|\mathbf{x}\|_2^2.$$

B.1. Lewis Weight Sampling gives Approximately Equal Leverage Scores

In this section, we prove that sampling according to Lewis weights gives a matrix with approximately equal leverage scores. This proof will largely rely on showing that, up to row rescaling, the sampled matrix $\tilde{\mathbf{A}}$ is such that $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}$ is spectrally close to $\mathbf{A}^T \mathbf{A}$. This proof will boil down to a standard application of matrix concentration bounds for sampling according to leverage scores. Our primary lemma in this section will then mostly follow from the following lemma, which will be proven at the end of this section.

Lemma 29 Given a matrix \mathbf{A} that is sampled according to Theorem 11 with error ϵ and gives matrix $\tilde{\mathbf{A}}$, then

$$\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \approx_{O(1)} \frac{1}{h(n, \epsilon)} \mathbf{A}^T \overline{\mathbf{W}}^{-1} \mathbf{A}$$

with high probability.

Using this, we can prove our key lemma.

Lemma 30 *Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ that is sampled according to Theorem 11 and gives matrix $\tilde{\mathbf{A}}$, then for all rows i of $\tilde{\mathbf{A}}$,*

$$\tau_i(\tilde{\mathbf{A}}) \approx_{O(1)} \frac{d}{N}$$

with high probability.

Proof Lemma 29 implies that

$$\tau_i(\tilde{\mathbf{A}}) = \tilde{\mathbf{A}}_{i,:} \left(\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \right)^{-1} \tilde{\mathbf{A}}_{i,:}^T \approx_{O(1)} h(n, \epsilon) \cdot \tilde{\mathbf{A}}_{i,:} \left(\mathbf{A}^T \overline{\mathbf{W}}^{-1} \mathbf{A} \right)^{-1} \tilde{\mathbf{A}}_{i,:}^T$$

with high probability. Theorem 11 implies that every row i of $\tilde{\mathbf{A}}$ is simply some row j of \mathbf{A} , scaled by $\frac{1}{p_j}$. Therefore, for any row i of $\tilde{\mathbf{A}}$ we must have

$$\tau_i(\tilde{\mathbf{A}}) \approx_{O(1)} h(n, \epsilon) \cdot \tilde{\mathbf{A}}_{i,:} \left(\mathbf{A}^T \overline{\mathbf{W}}^{-1} \mathbf{A} \right)^{-1} \tilde{\mathbf{A}}_{i,:}^T = h(n, \epsilon) \cdot \frac{\mathbf{A}_{j,:}}{p_j} \left(\mathbf{A}^T \overline{\mathbf{W}}^{-1} \mathbf{A} \right)^{-1} \frac{\mathbf{A}_{j,:}^T}{p_j}.$$

From Definition 7 we have

$$\overline{w}_j^2 = \mathbf{A}_{j,:} \left(\mathbf{A}^T \overline{\mathbf{W}}^{-1} \mathbf{A} \right)^{-1} \mathbf{A}_{j,:}^T$$

which along with the fact that $p_j \approx_{O(1)} \overline{w}_j \cdot h(n, \epsilon)$ reduces the leverage score to

$$\tau_i(\tilde{\mathbf{A}}) \approx_{O(1)} h(n, \epsilon) \cdot \frac{\overline{w}_j^2}{p_j^2} \approx_{O(1)} \frac{1}{h(n, \epsilon)}.$$

Finally Fact 27 gives us that the sum of Lewis weights must be d because they are leverage scores of $\overline{\mathbf{W}}^{-1/2} \mathbf{A}$, which implies $\frac{1}{h(n, \epsilon)} \approx_{O(1)} \frac{d}{N}$ by our definition of $N = \sum_i p_i$. \blacksquare

It now remains to prove Lemma 29. The proof follows similarly to the proof of Lemma 4 in Cohen et al. (2015), except that their leverage score sampling scheme draws each row without replacement, and we need a fixed number of sampled rows with replacement. Accordingly, we will also use the following matrix concentration result from Harvey (2012), which is a variant of Corollary 5.2 in Tropp (2012):

Lemma 31 (Harvey (2012)) *Let $\mathbf{Y}_1 \dots \mathbf{Y}_k$ be independent random positive semidefinite matrices of size $d \times d$. Let $\mathbf{Y} = \sum_{i=1}^k \mathbf{Y}_i$, and let $\mathbf{Z} = \mathbb{E}[\mathbf{Y}]$. If $\mathbf{Y}_i \preceq R \cdot \mathbf{Z}$ then*

$$\Pr \left[\sum_{i=1}^k \mathbf{Y}_i \preceq (1 - \epsilon) \mathbf{Z} \right] \leq de^{-\frac{\epsilon^2}{2R}}$$

and

$$\Pr \left[\sum_{i=1}^k \mathbf{Y}_i \succeq (1 + \epsilon) \mathbf{Z} \right] \leq de^{-\frac{\epsilon^2}{3R}}.$$

Proof [Proof of Lemma 29] First, we define $\bar{\mathbf{A}} \stackrel{\text{def}}{=} \bar{\mathbf{W}}^{-1/2} \mathbf{A}$. Then, by Definition 7, $\bar{w}_i = \tau_i(\bar{\mathbf{A}})$. Since $\bar{\mathbf{W}}$ is the diagonal matrix of Lewis weights \bar{w} , each row of $\bar{\mathbf{A}}$ is simply $\bar{\mathbf{A}}_{i,:} = \bar{w}_i^{-1/2} \mathbf{A}_{i,:}$.

By construction of our random $\tilde{\mathbf{A}}$ in Theorem 11 we choose a row j of \mathbf{A} with probability $\frac{p_j}{N}$ and scale by $\frac{1}{p_j}$. Therefore, if we let \mathbf{Y}_i be the random variable

$$\mathbf{Y}_i = \begin{cases} \frac{\mathbf{A}_{j,:} \mathbf{A}_{j,:}^T}{p_j^2}, & \text{with probability } \frac{p_j}{N} \text{ for each } j \end{cases}$$

then,

$$\mathbf{Y} = \sum_{i=1}^N \mathbf{Y}_i = \sum_{i=1}^N \tilde{\mathbf{A}}_{i,:} \tilde{\mathbf{A}}_{i,:}^T = \tilde{\mathbf{A}}^T \tilde{\mathbf{A}}.$$

Furthermore, we can substitute $\bar{\mathbf{A}}_{j,:} \sqrt{\bar{w}_i}$ for $\mathbf{A}_{j,:}$ and use the fact that $p_j \approx_{O(1)} \bar{w}_j \cdot h(n, \epsilon)$ to obtain

$$\frac{\mathbf{A}_{j,:} \mathbf{A}_{j,:}^T}{p_j^2} \approx_{O(1)} \frac{\bar{\mathbf{A}}_{j,:} \bar{\mathbf{A}}_{j,:}^T}{p_j \cdot h(n, \epsilon)}.$$

As a result, we have

$$\begin{aligned} \mathbf{Z} &= \mathbb{E} \left[\sum_{i=1}^N \mathbf{Y}_i \right] = \sum_{i=1}^N \mathbb{E} [\mathbf{Y}_i] \\ &\approx_{O(1)} \sum_{i=1}^N \sum_{j=1}^n \frac{\bar{\mathbf{A}}_{j,:} \bar{\mathbf{A}}_{j,:}^T}{N \cdot h(n, \epsilon)} \\ &= \frac{1}{h(n, \epsilon)} \sum_{j=1}^n \bar{\mathbf{A}}_{j,:} \bar{\mathbf{A}}_{j,:}^T = \frac{1}{h(n, \epsilon)} \mathbf{A}^T \bar{\mathbf{W}}^{-1} \mathbf{A}. \end{aligned}$$

In order to apply Lemma 31 we need to find R such that $\mathbf{Y}_i \preceq R \cdot \mathbf{Z}$, which by our construction of \mathbf{Y}_i requires

$$\frac{\mathbf{A}_{j,:} \mathbf{A}_{j,:}^T}{p_j^2} \preceq R \cdot \mathbf{Z}$$

for all j . We use our constant factor approximations of \mathbf{Z} and $\frac{\mathbf{A}_{j,:} \mathbf{A}_{j,:}^T}{p_j^2}$ to see that it also suffices to show

$$\frac{\bar{\mathbf{A}}_{j,:} \bar{\mathbf{A}}_{j,:}^T}{p_j \cdot h(n, \epsilon)} \preceq \frac{R}{O(1)} \cdot \frac{1}{h(n, \epsilon)} \mathbf{A}^T \bar{\mathbf{A}}.$$

Given that $\tau_j(\bar{\mathbf{A}}) = \bar{w}_j$ and $p_j \approx_{O(1)} \bar{w}_j \cdot h(n, \epsilon)$, we have

$$\frac{\bar{\mathbf{A}}_{j,:} \bar{\mathbf{A}}_{j,:}^T}{p_j \cdot h(n, \epsilon)} \preceq \frac{O(1) \bar{\mathbf{A}}_{j,:} \bar{\mathbf{A}}_{j,:}^T}{\tau_j(\bar{\mathbf{A}}) \cdot h(n, \epsilon)^2}$$

which along with the fact (Equation 12 in the proof of Lemma 4 from Cohen et al. (2015)) that

$$\frac{\bar{\mathbf{A}}_{j,:} \bar{\mathbf{A}}_{j,:}^T}{\tau_j(\bar{\mathbf{A}})} \preceq \bar{\mathbf{A}}^T \bar{\mathbf{A}}$$

implies that

$$\mathbf{Y}_i \preceq \frac{O(1)}{h(n, \epsilon)} \mathbf{Z}.$$

By Theorem 11 we know that $h(n, \epsilon) \geq c\epsilon^{-2} \log n$ for some constant c . Plugging this in for R in Lemma 31 gives that

$$\mathbf{Y} \approx_{1+\epsilon} \mathbf{Z}$$

or, substituting our values of \mathbf{Y} and \mathbf{Z} ,

$$\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \approx_{O(1)} \frac{1}{h(n, \epsilon)} \mathbf{A}^T \overline{\mathbf{W}}^{-1} \mathbf{A}$$

with probability at least $1 - 2de^{-\frac{\epsilon^{-2}}{3R}} \geq 1 - 2de^{-\frac{c \log n}{O(1)}} \geq 1 - 2dn^{-c/O(1)}$. This implies that the statement in the lemma is true with high probability for c bigger than $O(1)$ (where the $O(1)$ comes from our p_i approximation of $\overline{w}_i \cdot h(n, \epsilon)$) and our assumption on $n \geq d$. \blacksquare

B.2. Rotating the Matrix to Achieve Isotropic Position

Now that we have sampled by Lewis weights and achieved all leverage scores to be approximately equal, we will show that we can efficiently rotate the matrix into isotropic position while still preserving the fact that all leverage scores are approximately equal.

Lemma 32 *If $\mathbf{U} \in \mathbb{R}^{d \times d}$ is an invertible matrix and $\mathbf{U}^T \mathbf{U} = (\mathbf{A}^T \mathbf{A})^{-1}$ then*

1. $(\mathbf{A}\mathbf{U})^T (\mathbf{A}\mathbf{U}) = \mathbf{I}$.
2. For all rows i , $\tau_i(\mathbf{A}) = \tau_i(\mathbf{A}\mathbf{U})$.

Proof For the first condition, we see that

$$\mathbf{U}^T \mathbf{A}^T \mathbf{A} \mathbf{U} = \mathbf{I} \iff \mathbf{A}^T \mathbf{A} = (\mathbf{U}^T)^{-1} \mathbf{U}^{-1} \iff (\mathbf{A}^T \mathbf{A})^{-1} = \mathbf{U}^T \mathbf{U}.$$

For the second condition, the i th row of $\mathbf{A}\mathbf{U}$ will be $\mathbf{A}_{i,:} \mathbf{U}$, which by the definition of leverage scores then gives,

$$\begin{aligned} \tau_i(\mathbf{A}\mathbf{U}) &= \mathbf{A}_{i,:} \mathbf{U} ((\mathbf{A}\mathbf{U})^T (\mathbf{A}\mathbf{U}))^{-1} (\mathbf{A}_{i,:} \mathbf{U})^T \\ &= \mathbf{A}_{i,:} \mathbf{U} \mathbf{U}^{-1} (\mathbf{A}^T \mathbf{A})^{-1} (\mathbf{U}^T)^{-1} \mathbf{U}^T \mathbf{A}_{i,:}^T \\ &= \mathbf{A}_{i,:} (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}_{i,:}^T \\ &= \tau_i(\mathbf{A}). \end{aligned}$$

\blacksquare

It is clear then that we want to rotate our matrix by \mathbf{U} as above, so it only remains to efficiently compute such a \mathbf{U} .

Lemma 33 *Given a full rank matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times d}$, there is a routine ROTATE that can find an invertible \mathbf{U} such that $\mathbf{U}^T \mathbf{U} = (\tilde{\mathbf{A}}^T \tilde{\mathbf{A}})^{-1}$ in time $O(Nd^{\omega-1} + d^\omega)$.*

Proof Computing $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}$ can be done in $O(Nd^{\omega-1})$ time using fast matrix multiplication. Inverting $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}$, a $d \times d$ matrix that must have an inverse because $\tilde{\mathbf{A}}$ is full rank, will require $O(d^\omega)$ time. Finally, we perform a QR-decomposition of $(\tilde{\mathbf{A}}^T \tilde{\mathbf{A}})^{-1}$ in $O(d^\omega)$ time to obtain our square invertible matrix \mathbf{U} .¹⁵ ■

Lastly, we want to ensure that by rotating our matrix, we can still use an approximate solution to the rotated matrix to obtain an approximate solution of the original matrix.

Lemma 34 *Given a matrix-vector pair $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$, another matrix-vector pair $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times d}$, $\tilde{\mathbf{b}} \in \mathbb{R}^N$, and an invertible matrix $\mathbf{U} \in \mathbb{R}^{d \times d}$,*

$$\|[\mathbf{A}, \mathbf{b}]\mathbf{x}\|_1 \approx_{1+\epsilon} \|[\tilde{\mathbf{A}}, \tilde{\mathbf{b}}]\mathbf{x}\|_1 \forall \mathbf{x} \in \mathbb{R}^{d+1} \iff \|[\mathbf{A}\mathbf{U}, \mathbf{b}]\mathbf{y}\|_1 \approx_{1+\epsilon} \|[\tilde{\mathbf{A}}\mathbf{U}, \tilde{\mathbf{b}}]\mathbf{y}\|_1 \forall \mathbf{y} \in \mathbb{R}^{d+1}.$$

Proof This follows immediately from the fact that for any \mathbf{x} satisfying the LHS, there exists a \mathbf{y} satisfying the RHS, and vice versa. Specifically $\mathbf{y}_{[1,d]} = \mathbf{U}^{-1}\mathbf{x}_{[1,d]}$ and $\mathbf{y}_{d+1} = \mathbf{x}_{d+1}$, and equivalently $\mathbf{U}\mathbf{y}_{[1,d]} = \mathbf{x}_{[1,d]}$ and $\mathbf{y}_{d+1} = \mathbf{x}_{d+1}$. ■

B.3. Translating between Preconditioned and Original Matrix Solutions

Our preconditioning combination of Lewis weights and rotating the matrix gives our desired conditions, specifically an IRB matrix, but it remains to be seen that we can take a solution to this preconditioned matrix and translate it back into an approximate solution of the original matrix. In the following lemma we will show that this is in fact true.

Lemma 35 *Given a matrix-vector pair $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$, another matrix-vector pair $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times d}$, $\tilde{\mathbf{b}} \in \mathbb{R}^N$, and an invertible matrix $\mathbf{U} \in \mathbb{R}^{d \times d}$; if*

$$\|[\tilde{\mathbf{A}} \tilde{\mathbf{b}}]\mathbf{y}\|_1 \approx_{1+\epsilon} \|[\mathbf{A} \mathbf{b}]\mathbf{y}\|_1$$

for all $\mathbf{y} \in \mathbb{R}^{d+1}$, and if $\tilde{\mathbf{x}}_{\mathbf{U}}^*$ minimizes $\|\tilde{\mathbf{A}}\mathbf{U}\mathbf{x} - \tilde{\mathbf{b}}\|_1$, then for any $\tilde{\mathbf{x}}$ such that

$$\|\tilde{\mathbf{A}}\mathbf{U}\tilde{\mathbf{x}} - \tilde{\mathbf{b}}\|_1 \leq (1 + \delta)\|\tilde{\mathbf{A}}\mathbf{U}\tilde{\mathbf{x}}_{\mathbf{U}}^* - \tilde{\mathbf{b}}\|_1$$

we must have

$$\|\mathbf{A}(\mathbf{U}\tilde{\mathbf{x}}) - \mathbf{b}\|_1 \leq (1 + \epsilon)^2(1 + \delta)\|\mathbf{A}\tilde{\mathbf{x}}_{\mathbf{U}}^* - \mathbf{b}\|_1$$

with high probability.

Proof By assumption we have

$$\|[\tilde{\mathbf{A}} \tilde{\mathbf{b}}]\mathbf{y}\|_1 \approx_{1+\epsilon} \|[\mathbf{A} \mathbf{b}]\mathbf{y}\|_1$$

for all $\mathbf{y} \in \mathbb{R}^{d+1}$, and we can then use Lemma 34 to obtain

$$\|[\tilde{\mathbf{A}}\mathbf{U} \tilde{\mathbf{b}}]\mathbf{y}\|_1 \approx_{1+\epsilon} \|[\mathbf{A}\mathbf{U} \mathbf{b}]\mathbf{y}\|_1$$

15. For an invertible matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$, it is easy to see that $\mathbf{M}(\mathbf{M}^T \mathbf{M})^{-1/2}$ is an orthonormal basis for \mathbf{M} . We can compute $(\mathbf{M}^T \mathbf{M})^{-1}$ using Schur decomposition in $O(d^\omega)$ time, and by careful analysis of that algorithm, we can also compute $(\mathbf{M}^T \mathbf{M})^{-1/2}$ in the same amount of time.

for all $\mathbf{y} \in \mathbb{R}^{d+1}$. By fixing \mathbf{y} to be $\begin{pmatrix} \mathbf{x} \\ -1 \end{pmatrix}$, we get

$$\|\tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}}\|_1 \approx_{1+\epsilon} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1 \quad \forall \mathbf{x} \in \mathbb{R}^d, \quad (2)$$

$$\|\tilde{\mathbf{A}}\mathbf{U}\mathbf{x} - \tilde{\mathbf{b}}\|_1 \approx_{1+\epsilon} \|\mathbf{A}\mathbf{U}\mathbf{x} - \mathbf{b}\|_1 \quad \forall \mathbf{x} \in \mathbb{R}^d. \quad (3)$$

(2) gives

$$\|\mathbf{A}\mathbf{U}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon)\|\tilde{\mathbf{A}}\mathbf{U}\tilde{\mathbf{x}} - \tilde{\mathbf{b}}\|_1.$$

Using our initial assumption and defining $\tilde{\mathbf{x}}^* \stackrel{\text{def}}{=} \mathbf{U}\tilde{\mathbf{x}}_{\mathbf{U}}^*$ then gives us

$$\|\mathbf{A}\mathbf{U}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon)(1 + \delta)\|\tilde{\mathbf{A}}\tilde{\mathbf{x}}^* - \tilde{\mathbf{b}}\|_1.$$

Notice that if $\tilde{\mathbf{x}}_{\mathbf{U}}^*$ minimizes $\|\tilde{\mathbf{A}}\mathbf{U}\mathbf{x} - \tilde{\mathbf{b}}\|_1$, then $\tilde{\mathbf{x}}^*$ must minimize $\|\tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}}\|_1$ because \mathbf{U} is invertible. Therefore, $\|\tilde{\mathbf{A}}\tilde{\mathbf{x}}^* - \tilde{\mathbf{b}}\|_1 \leq \|\tilde{\mathbf{A}}\mathbf{x}^* - \tilde{\mathbf{b}}\|_1$ and we have

$$\|\mathbf{A}\mathbf{U}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon)(1 + \delta)\|\tilde{\mathbf{A}}\mathbf{x}^* - \tilde{\mathbf{b}}\|_1.$$

Finally, applying (3) gives

$$\|\mathbf{A}\mathbf{U}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon)^2(1 + \delta)\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1.$$

■

B.4. Proof of Lemma 12

We now have all the necessary pieces to prove our primary preconditioning lemma, which we will now restate and prove.

Lemma 36 *There is a routine that takes a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, a vector $\mathbf{b} \in \mathbb{R}^n$ and $\epsilon > 0$, then produces a matrix $[\tilde{\mathbf{A}}, \tilde{\mathbf{b}}] \in \mathbb{R}^{N \times (d+1)}$ with $N = O(d\epsilon^{-2} \log n)$ and an invertible matrix $\mathbf{U} \in \mathbb{R}^{d \times d}$ such that matrix $\tilde{\mathbf{A}}\mathbf{U}$ is IRB and if $\tilde{\mathbf{x}}_{\mathbf{U}}^*$ minimizes $\|\tilde{\mathbf{A}}\mathbf{U}\mathbf{x} - \tilde{\mathbf{b}}\|_1$, then for any $\tilde{\mathbf{x}}$ such that*

$$\|\tilde{\mathbf{A}}\mathbf{U}\tilde{\mathbf{x}} - \tilde{\mathbf{b}}\|_1 \leq (1 + \delta)\|\tilde{\mathbf{A}}\mathbf{U}\tilde{\mathbf{x}}_{\mathbf{U}}^* - \tilde{\mathbf{b}}\|_1,$$

holds for some $\delta > 0$, we must have

$$\|\mathbf{A}(\mathbf{U}\tilde{\mathbf{x}}) - \mathbf{b}\|_1 \leq (1 + \epsilon)^2(1 + \delta)\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1$$

with high probability.

Furthermore, the full running time is $O(nnz(\mathbf{A}) \log n + d^{\omega-1} \min\{d\epsilon^{-2} \log n, n\} + \Upsilon)$ where $\Upsilon = \min\{d\epsilon^{-2} \log n, (d\epsilon^{-2} \log n)^{1/2+o(1)} + n \log^2 n\}$.

Proof From Theorem 11 we have that

$$\|[\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]\mathbf{y}\|_1 \approx_{1+\epsilon} \|[\mathbf{A} \ \mathbf{b}]\mathbf{y}\|_1$$

for all $\mathbf{y} \in \mathbb{R}^{d+1}$ with high probability. Lemma 35 then gives

$$\|\mathbf{A}(\mathbf{U}\tilde{\mathbf{x}}) - \mathbf{b}\|_1 \leq (1 + \epsilon)^2(1 + \delta) \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1$$

by our assumption on $\tilde{\mathbf{x}}$.

Lemma 29 and the assumption that \mathbf{A} is full rank imply that $\tilde{\mathbf{A}}$ is full rank with high probability. Our use of ROTATE to generate \mathbf{U} , such that $\mathbf{U}^T \mathbf{U} = (\tilde{\mathbf{A}}^T \tilde{\mathbf{A}})^{-1}$, along with Lemma 32, gives $(\tilde{\mathbf{A}}\mathbf{U})^T \tilde{\mathbf{A}}\mathbf{U} = \mathbf{I}$ and also that $\tau_i(\tilde{\mathbf{A}}\mathbf{U}) = \tau_i(\tilde{\mathbf{A}})$ for all i . Fact 28 gives $\tau_i(\tilde{\mathbf{A}}) \leq \tau_i([\tilde{\mathbf{A}} \tilde{\mathbf{b}}])$, which along with Lemma 30, implies $\tau_i(\tilde{\mathbf{A}}\mathbf{U}) \leq O(d/N)$ for all i . Finally, by Definition 6 and the fact that $(\tilde{\mathbf{A}}\mathbf{U})^T \tilde{\mathbf{A}}\mathbf{U} = \mathbf{I}$, we then have

$$\tau_i(\tilde{\mathbf{A}}\mathbf{U}) = \left\| \left(\tilde{\mathbf{A}}\mathbf{U} \right)_{i,:} \right\|_2.$$

The sampling of \mathbf{A} is done according to the technique by Cohen and Peng (2015), which requires $O(nnz(\mathbf{A}) \log n + d^\omega)$ time to obtain the sampling probabilities. Then the actual sampling requires $O(\min\{d\epsilon^{-2} \log n, (d\epsilon^{-2} \log n)^{1/2+o(1)} + n \log^2 n\})$ -time according to Corollary 46 shown in Appendix D.1. Computing the invertible matrix \mathbf{U} for input $\tilde{\mathbf{A}}$ takes $O(Nd^{\omega-1} + d^\omega)$ time from Lemma 33, and the number of rows of $\tilde{\mathbf{A}}$ is $N = O(d\epsilon^{-2} \log n)$. Finally, Lemma 42 and Corollary 43 in Appendix D.1 show that this computation time can also be bounded with $N \leq n$, which then gives our desired runtime. \blacksquare

Appendix C. Proofs from Section 4

In this section we provide the omitted proofs from Section 4.

C.1. Proof of Lemma 17

In this section we reduce the number of rows in \mathbf{A} by uniform sampling while still preserving certain guarantees. Note that we will ultimately sample from $[\mathbf{A} \ \mathbf{b}]$, but for simplicity in notation, we will just use \mathbf{A} here.

To prove Lemma 17, we need the following lemma, which states the key fact that the conditions on \mathbf{A} ensure approximately uniform Lewis weights.

Lemma 37 (Almost-uniform leverage scores imply almost-uniform Lewis weights) *Consider a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ such that $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$ and $\|\mathbf{A}_{i,:}\|_2^2 \approx_{O(1)} d/n$. Let $\bar{\mathbf{w}}$ denote the ℓ_1 Lewis weights for \mathbf{A} . Then for each row i , we have $\bar{\mathbf{w}}_i \approx_{O(1)} d/n$.*

Proof [Proof of Lemma 17] Note that by Lemma 37 we have

$$\mathbf{p}_i = \frac{N}{n} \approx_{O(1)} \frac{d \cdot O(\epsilon^{-2} \log n)}{n} \approx_{O(1)} \bar{\mathbf{w}}_i \cdot O(\epsilon^{-2} \log n).$$

Thus, if we use $\mathbf{p}_i = N/n$ for each i in Theorem 11, we get the first property while avoiding the cost of computing \mathbf{p}_i 's stated in Theorem 11.

The second property follows from Lemma 29 in Appendix B. Specifically, we have

$$\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \approx_{O(1)} \frac{1}{O(\epsilon^{-2} \log n)} \mathbf{A}^T \bar{\mathbf{W}}^{-1} \mathbf{A},$$

which then implies that

$$\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \approx_{O(1)} \frac{n}{d \cdot O(\epsilon^{-2} \log n)} \mathbf{A}^T \mathbf{A} \approx_{O(1)} \frac{n}{N} \mathbf{I}.$$

Let τ denote the leverage scores for \mathbf{A} . Now, for the third property, it follows from the definition of leverage scores and the second property that

$$\tau_i(\tilde{\mathbf{A}}) = \left\| \left(\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \right)^{-1/2} \tilde{\mathbf{A}}_{i,:}^T \right\|_2^2 \approx_{O(1)} \left\| \sqrt{N/n} \tilde{\mathbf{A}}_{i,:}^T \right\|_2^2.$$

Furthermore, Lemma 30 in Appendix B shows that $\tau_i(\tilde{\mathbf{A}}) \approx_{O(1)} d/N$. Factoring this into the equation gives us

$$\left\| \tilde{\mathbf{A}}_{i,:}^T \right\|_2^2 \approx_{O(1)} dn/N^2. \quad \blacksquare$$

Now, to prove Lemma 37, we need the following definition and lemma.

Definition 38 (Definition 5.1 of α -almost Lewis weights for ℓ_1 from Cohen and Peng (2015)) For a matrix \mathbf{A} , an assignment of weights \mathbf{w} is α -almost Lewis if

$$\mathbf{A}_{i,:} (\mathbf{A}^T \mathbf{W}^{-1} \mathbf{A})^{-1} \mathbf{A}_{i,:}^T \approx_{\alpha} \mathbf{w}_i^2,$$

where \mathbf{W} is the diagonal matrix form of \mathbf{w} .

Lemma 39 (Definition 5.2 and Lemma 5.3 from Cohen and Peng (2015)) Any set of α -almost Lewis weights satisfy

$$\bar{\mathbf{w}}_i \approx_{\alpha} \mathbf{w}_i.$$

Proof [Proof of Lemma 37] We know that $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$ and for each row i , $\|\mathbf{A}_{i,:}\|_2^2 \approx_{O(1)} d/n$. Then,

$$\begin{aligned} \tau_i(\mathbf{A}) &= \mathbf{A}_{i,:} (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}_{i,:}^T \approx_{O(1)} \mathbf{A}_{i,:} \mathbf{A}_{i,:}^T \\ \implies \tau_i(\mathbf{A}) &\approx_{O(1)} d/n. \end{aligned}$$

That is, all of the leverage scores are approximately equal. Then we can show that $\mathbf{w} = (d/n) \mathbb{1}$, where $\mathbb{1}$ is the all ones vector. Then,

$$\mathbf{A}_{i,:} (\mathbf{A}^T \mathbf{W}^{-1} \mathbf{A})^{-1} \mathbf{A}_{i,:}^T = (d/n) \mathbf{A}_{i,:} (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}_{i,:}^T \approx_{O(1)} d^2/n^2 = \mathbf{w}_i^2.$$

Thus, \mathbf{w} is $O(1)$ -almost Lewis. The result then follows by Lemma 39. \blacksquare

C.2. Proof of Lemma 19 and Lemma 20

To prove Lemma 19, we use the following lemma:

Lemma 40 *Let $\mathbf{v} \in \mathbb{R}^n$ be a vector with $\|\mathbf{v}\|_1 = 1$. Then, for a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$,*

$$\|\mathbf{A}^T \mathbf{v}\|_2 \leq \max_i \|\mathbf{A}_{i,:}\|_2.$$

Proof

$$\|\mathbf{A}^T \mathbf{v}\|_2 = \left\| \sum_i \mathbf{A}_{i,:} v_i \right\|_2 \leq \max_i \|\mathbf{A}_{i,:}\|_2.$$

where the inequality follows by the convexity of $\|\cdot\|_2$ and since $\sum_i |v_i| = 1$, ■

Proof [Proof of Lemma 19] By our assumptions on \mathbf{A} , we have $\mathbf{A}^T \mathbf{A} + \mathbf{B} = \mathbf{I}$ for some symmetric \mathbf{B} where $\|\mathbf{B}\|_2 \leq O(1)$. Since $\mathbf{x}_0 = \arg \min_x \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$, we have $\mathbf{x}_0 = \mathbf{A}^\dagger \mathbf{b} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$.

$$\begin{aligned} \|\mathbf{x}_0 - \mathbf{x}^*\|_2 &= \|(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b} - \mathbf{x}^*\|_2 \\ &= \|(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b} - (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{A} \mathbf{x}^*\|_2 \\ &= \|(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T (\mathbf{b} - \mathbf{A} \mathbf{x}^*)\|_2. \end{aligned}$$

Let $\mathbf{v} = (\mathbf{A} \mathbf{x}^* - \mathbf{b}) / \|\mathbf{A} \mathbf{x}^* - \mathbf{b}\|_1$.

$$\begin{aligned} \|\mathbf{x}_0 - \mathbf{x}^*\|_2 &= \|(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{v}\|_2 \|\mathbf{b} - \mathbf{A} \mathbf{x}^*\|_1 \\ &= \|(\mathbf{A}^T \mathbf{A} + \mathbf{B})(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{v}\|_2 \|\mathbf{b} - \mathbf{A} \mathbf{x}^*\|_1 \\ &= \|(\mathbf{I} + \mathbf{B}(\mathbf{A}^T \mathbf{A})^{-1}) \mathbf{A}^T \mathbf{v}\|_2 \|\mathbf{b} - \mathbf{A} \mathbf{x}^*\|_1 \\ &\leq \|\mathbf{I} + \mathbf{B}(\mathbf{A}^T \mathbf{A})^{-1}\|_2 \|\mathbf{A}^T \mathbf{v}\|_2 \|\mathbf{b} - \mathbf{A} \mathbf{x}^*\|_1. \end{aligned}$$

Now note:

$$\begin{aligned} \|\mathbf{I} + \mathbf{B}(\mathbf{A}^T \mathbf{A})^{-1}\|_2 &\leq \|\mathbf{I}\|_2 + \|\mathbf{B}\|_2 \|(\mathbf{A}^T \mathbf{A})^{-1}\|_2 \\ &\leq O(1). \end{aligned}$$

Also, by Lemma 40 and the assumptions on \mathbf{A} ,

$$\|\mathbf{A}^T \mathbf{v}\|_2 \leq O(\sqrt{d/n}).$$

Thus, we have:

$$\|\mathbf{x}_0 - \mathbf{x}^*\|_2 \leq O(\sqrt{d/n}) \|\mathbf{b} - \mathbf{A} \mathbf{x}^*\|_1. \quad \blacksquare$$

To prove Lemma 20, we use the following theorem from [Sachdeva and Vishnoi \(2014\)](#):

Theorem 41 (Theorem 9.1 from Sachdeva and Vishnoi (2014)) *Given an symmetric positive definite matrix $M \in \mathbb{R}^{n \times n}$ and a vector $\mathbf{y} \in \mathbb{R}^n$, the Conjugate Gradient method can find a vector \mathbf{x} such that $\|\mathbf{x} - M^{-1}\mathbf{y}\|_M \leq \delta \|M^{-1}\mathbf{y}\|_M$ in time $O((t_M + n) \cdot \sqrt{\kappa(M)} \log(1/\delta))$, where t_M is the time required to multiply M with a given vector and $\kappa(M)$ is the condition number of M .*

Proof [Proof of Lemma 20] Let $M = \mathbf{A}^T \mathbf{A}$ and $\mathbf{y} = \mathbf{A}^T \mathbf{b}$. Then by Theorem 41, the conjugate gradient method finds a vector $\tilde{\mathbf{x}}_0$ such that $\|\tilde{\mathbf{x}}_0\|_{\mathbf{A}^T \mathbf{A}} \leq \delta \|(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}\|_{\mathbf{A}^T \mathbf{A}}$ in time $O((t_{\mathbf{A}^T \mathbf{A}} + d) \log(1/\delta))$. Noting that $\mathbf{A}^T \mathbf{A} \approx_{O(1)} \mathbf{I}$, we get

$$\|\tilde{\mathbf{x}}_0 - \mathbf{x}_0\|_2 \leq O(\delta) \|\mathbf{x}_0\|_2.$$

Next, we note:

$$\begin{aligned} \|\mathbf{b}\|_2 &\geq \|\mathbf{A}\mathbf{x}_0 - \mathbf{b}\|_2 \geq \|\mathbf{A}\mathbf{x}_0\|_2 - \|\mathbf{b}\|_2 \\ \implies \|\mathbf{x}_0\|_2 &\leq O(\|\mathbf{b}\|_2). \end{aligned}$$

Now, since we assume that $\|\mathbf{b}\|_2 \leq n^c$ and $\|\mathbf{A}\mathbf{x}_0 - \mathbf{b}\|_2 \geq 1/n^c$ for some c , we can set $\delta = O(\epsilon/(n^c))$ to get:

$$\|\tilde{\mathbf{x}}_0 - \mathbf{x}_0\|_2 \leq \epsilon/n^c \leq \epsilon \|\mathbf{A}\mathbf{x}_0 - \mathbf{b}\|_2. \quad \blacksquare$$

C.3. Proof of Theorem 2

Theorem 2 *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{b} \in \mathbb{R}^n$ be such that the matrix-vector pair $[\mathbf{A} \ \mathbf{b}]$ satisfies $[\mathbf{A} \ \mathbf{b}]^T [\mathbf{A} \ \mathbf{b}] \approx_{O(1)} \mathbf{I}$ ¹⁶ and for each row i of $[\mathbf{A} \ \mathbf{b}]$, $\|[\mathbf{A} \ \mathbf{b}]_{i,:}\|_2^2 \approx_{O(1)} d/n$. Assume $\|\mathbf{b}\|_2 \leq n^c$ and $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ is either 0 or bounded below by n^{-c} for some constant $c > 0$. Then for any $\epsilon > 0$, there is a routine that computes $\tilde{\mathbf{x}}$ such that with high probability,*

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_1 \leq (1 + \epsilon) \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$$

with a runtime of $O(\text{nnz}(\mathbf{A}) \log^2 n + s \cdot d^{1.5} \epsilon^{-2} \log^{1.5} n + d^2 \epsilon^{-2} \log^2 n)$, where s is the maximum number of entries in any row of \mathbf{A} .

Proof We first prove correctness. We apply Lemma 17 to $[\mathbf{A} \ \mathbf{b}]$ and rescale $[\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]$ by $\sqrt{N/n}$. Note that rescaling does not change the relative error of our output $\tilde{\mathbf{x}}$. From this rescaling, we have $[\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]^T [\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}] \approx_{O(1)} \mathbf{I}$ and thus $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} \approx_{O(1)} \mathbf{I}$. This implies that $\tau_i([\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]) \approx_{O(1)} \|[\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]_{i,:}\|_2^2$ and $\tau_i(\tilde{\mathbf{A}}) \approx_{O(1)} \|\tilde{\mathbf{A}}_{i,:}\|_2^2$. By Fact 28, we have $\tau_i(\tilde{\mathbf{A}}) \leq \tau_i([\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}])$, so we have $\|\tilde{\mathbf{A}}_{i,:}\|_2^2 \leq O(d/N)$ for all rows i . As a result, we can find $\tilde{\mathbf{x}}_0$ according to Lemma 18. The rest of the correctness follows exactly as in the proof of Theorem 1.

We now examine the running time and note that uniform sampling will take $O(\text{nnz}(\mathbf{A}) + d^2 \epsilon^{-2} \log n)$ time to produce $[\tilde{\mathbf{A}} \ \tilde{\mathbf{b}}]$. By Lemma 18, we can then find $\tilde{\mathbf{x}}_0$ in time $O(d^2 \epsilon^{-2} \log n)$ because $\tilde{\mathbf{A}}$ is a $d \epsilon^{-2} \log n \times d$ matrix, so $t_{\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}} = O(d^2 \epsilon^{-2} \log n)$. Finally from the analysis of

16. As defined in the notation section, we say that $\mathbf{A} \approx_{\kappa} \mathbf{B}$ if and only if $\frac{1}{\kappa} \mathbf{B} \preceq \mathbf{A} \preceq \kappa \mathbf{B}$.

Theorem 1 we know that accelerated stochastic gradient descent requires $O(d^{2.5} \log^{1/2} n \cdot \epsilon^{-2})$ time. However, we note that the extra factor of d came from Theorem 26 where we substituted d for the time per iteration of stochastic gradient descent. This value can actually be upper bounded by the maximum number of entries in any row of $\tilde{\mathbf{A}}$, which because of our uniform sampling is upper bounded by the maximum number of entries in any row of \mathbf{A} . Adding a runtime overhead of $\log n$ for computing an approximation of the optimal objective, as in Appendix D.2, gives the desired runtime. ■

Appendix D. Secondary technical details for our main results

In this section, we address a couple assumptions that were made in the proofs of our main results. These assumptions were minor details, but we now include proofs for completeness. First, we always assumed that our row dimension after preprocessing was $O(\min(n, d\epsilon^{-2} \log n))$, and we will address this in Appendix D.1. Second, we required a constant factor approximation of the optimal objective value for which we give the procedure in Appendix D.2.

D.1. Simulated Sampling of \mathbf{A}

In Lemma 12, our primary preconditioning lemma, we set N to be the minimum of n and $O(d\epsilon^{-2} \log n)$. However, all of our sampling above assumed that $O(d\epsilon^{-2} \log n)$ rows were sampled to achieve certain matrix concentration results. Accordingly, we will still assume that $O(d\epsilon^{-2} \log n)$ rows are sampled, but show that we can reduce the computational cost of any duplicate rows to $O(1)$, and hence the computation factor of N can be assumed to be $\min\{n, O(d\epsilon^{-2} \log n)\}$. The sampling procedure itself can be done in about $O(n)$ time. At the end of this section, we explain how the running time of Katyusha can be made to depend on n , rather than $d\epsilon^{-2} \log n$.

Ultimately, our proof of Lemma 12 will critically use the fact that $\tilde{\mathbf{A}}$ has $O(d\epsilon^{-2} \log n)$ rows in several places. The following lemmas will then show how we can reduce this computation for duplicate rows, allowing us to substitute n for $O(d\epsilon^{-2} \log n)$ in the running time when $n \ll d\epsilon^{-2} \log n$.

Lemma 42 *Let $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times d}$ be a matrix with at most n unique rows, and for each unique row, we are given the number of copies in $\tilde{\mathbf{A}}$. Then computing $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}$ takes at most $O(nd^{\omega-1})$ time.*

Proof By definition

$$\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} = \sum_i \tilde{\mathbf{A}}_{i,:}^T \tilde{\mathbf{A}}_{i,:}.$$

Therefore, if we have k copies of row $\tilde{\mathbf{A}}_{i,:}$, we know that they contribute $k \tilde{\mathbf{A}}_{i,:}^T \tilde{\mathbf{A}}_{i,:}$ to the summation. Accordingly, if we replaced all of them with one row $\sqrt{k} \tilde{\mathbf{A}}_{i,:}$, then this row would contribute an equivalent amount to the summation. As a result, we can combine all copies of unique rows to achieve an $n \times d$ matrix $\tilde{\mathbf{A}}'$ and compute $\tilde{\mathbf{A}}'^T \tilde{\mathbf{A}}'$ which will be equivalent to $\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}$. ■

Corollary 43 *Let $[\tilde{\mathbf{A}} \tilde{\mathbf{b}}] \in \mathbb{R}^{N \times (d+1)}$ be a matrix with at most n unique rows, and for each unique row, we are given the number of copies in $[\tilde{\mathbf{A}} \tilde{\mathbf{b}}]$. Then computing $\tilde{\mathbf{A}} \mathbf{U}$ where $\mathbf{U} \in \mathbb{R}^{d \times d}$, and computing $\tilde{\mathbf{A}}^T \tilde{\mathbf{b}}$ takes $O(nd^{\omega-1})$ and $O(nd)$ time, respectively.*

Proof We can similarly use the fact that $\tilde{\mathbf{A}}_{i,:} \mathbf{U}$ is equivalent for all copies of $\tilde{\mathbf{A}}_{i,:}$ and combine k copies into the row $k\tilde{\mathbf{A}}_{i,:}$.

Analogously, we have $\tilde{\mathbf{A}}^T \tilde{\mathbf{b}} = \sum_i \tilde{\mathbf{A}}_{i,:}^T \tilde{\mathbf{b}}_i$, so we can combine duplicate rows. \blacksquare

Furthermore, we need to show that we can efficiently sample $O(d\epsilon^{-2} \log n)$ rows (ideally in $O(n)$ -time) even when $O(d\epsilon^{-2} \log n) \gg n$. We will achieve this through known results on fast binomial distribution sampling.

Theorem 44 (Theorem 1.1 in Farach-Colton and Tsai (2015)) *Given a binomial distribution $B(n, p)$ for $n \in \mathbb{N}$, $p \in \mathbb{Q}$, drawing a sample from it takes $O(\log^2 n)$ time using $O(n^{1/2+\epsilon})$ space w.h.p., after $O(n^{1/2+\epsilon})$ -time preprocessing for small $\epsilon > 0$. The preprocessing does not depend on p and can be used for any $p' \in \mathbb{Q}$ and for any $n' \leq n$.*

This result implies that sampling m items independently can be done more efficiently if $m \gg n$, where we are only concerned with the number of times each item in the state space is sampled.

Corollary 45 *Given a probability distribution $\mathcal{P} = (p_1, \dots, p_n)$ over a state space of size n , sampling m items independently from \mathcal{P} takes $O(m^{1/2+\epsilon} + n \log^2 n)$ -time.*

Proof Note that sampling independently m times is equivalent to determining how many of each item is sampled by using the binomial distribution and updating after each item. More specifically, we can iterate over all $i \in [n]$ and draw $k_i \sim B(m, p_i)$, then update m to be $m - k_i$ and scale up each p_j (where $j > i$) by $(1 - p_i)^{-1}$. It is straightforward to make the scaling up of each p_j efficient, and according to Theorem 44 we can obtain the binomial sample in $O(\log^2 n)$ -time.

Furthermore, because m is decreasing at each iteration, we can use the original preprocessing in Theorem 44 for each step to achieve our desired running time. \blacksquare

Corollary 46 *Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, with a probability distribution over each row, we can produce a matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{O(d\epsilon^{-2} \log n) \times d}$ according to the given distribution in time at most $O(\min(d\epsilon^{-2} \log n, (d\epsilon^{-2} \log n)^{1/2+o(1)} + n \log^2 n))$.*

Finally, our application of Theorem 26 assumes that it is given an $N \times d$ matrix, but we assumed that the computational cost could assume $N = \min\{n, O(d\epsilon^{-2} \log n)\}$. A closer examination of Algorithm 2 in Allen Zhu (2017), which is the routine for Theorem 26, shows that the factor of N comes from a full gradient calculation, which can be done more quickly by combining rows in an equivalent manner to the lemma and corollary above.

D.2. Approximating the Optimal Objective Value

For ease of notation, we let $f^* = \min_x \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1$ and $f_2^* = \min_x \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ in this section. In our proof of both Theorem 16 and 1, we assumed access to a constant approximation of f^* with a runtime overhead of $\log n$. We will obtain access to this value by giving polynomially approximate upper and lower bounds on f^* and using our primary algorithm on $\log n$ guesses for f^* within this range. We start with the following lemma that gives upper and lower bounds on f^* :

Lemma 47 *Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a vector $\mathbf{b} \in \mathbb{R}^n$, if \mathbf{x}_2^* minimizes $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ then*

$$\|\mathbf{A}\mathbf{x}_2^* - \mathbf{b}\|_2 \leq \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1 \leq \sqrt{n} \|\mathbf{A}\mathbf{x}_2^* - \mathbf{b}\|_2.$$

Proof By known properties of ℓ_1 and ℓ_2 , for any $\mathbf{x} \in \mathbb{R}^n$, we have $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1 \leq \sqrt{n} \|\mathbf{x}\|_2$. Accordingly, we must have

$$\|\mathbf{A}\mathbf{x}_2^* - \mathbf{b}\|_2 \leq \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_2 \leq \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1,$$

where the first inequality follows from \mathbf{x}_2^* being the ℓ_2 -minimizer. Similarly, we also have

$$\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_1 \leq \|\mathbf{A}\mathbf{x}_2^* - \mathbf{b}\|_1 \leq \sqrt{n} \|\mathbf{A}\mathbf{x}_2^* - \mathbf{b}\|_2,$$

where the first inequality follows from \mathbf{x}^* being the ℓ_1 -minimizer. ■

Since $\mathbf{A}^T \mathbf{A} = \mathbf{I}$, our initialization of $\mathbf{A}^T \mathbf{b}$ is equal to \mathbf{x}_2^* . Then we can compute $\|\mathbf{A}\mathbf{x}_2^* - \mathbf{b}\|_2$ in $O(Nd)$ time.¹⁷ Consequently, if we let f_2^* be the minimized objective function $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$, we can compute polynomially close upper and lower bounds, f_2^* and $\sqrt{n}f_2^*$ respectively, for f^* .

Lemma 48 *In both variants of our primary algorithm for Theorem 16 and 1, we can run the respective algorithms with a constant approximation of f^* by running them $\log n$ times using different approximations of f^* , which we will denote by \tilde{f}^* . Furthermore, the runtime of each is independent of the choice of \tilde{f}^* .*

Proof We first examine the latter claim and note that the gradient descent portion of both algorithms take upper bounds on $\|\mathbf{x}_0 - \mathbf{x}^*\|_2$ as inputs. Therefore, given a certain \tilde{f}^* we can input the upper bound $O(\sqrt{d/n})\tilde{f}^*$ and following the analysis of the proofs in Theorems 16 and 1, in runtime $O(d^3\epsilon^{-2})$ and $O(d^{2.5}\epsilon^{-2}\sqrt{\log n})$ respectively, we are guaranteed that we achieve $\tilde{\mathbf{x}}$ such that $f(\tilde{\mathbf{x}}) - f^* \leq \epsilon\tilde{f}^*$ with high probability. However, note that this is only true if $f^* \leq \tilde{f}^*$. Otherwise, we are given no guarantee on the closeness of $f(\tilde{\mathbf{x}})$ to f^* .

The runtime of each algorithm is then not affected by our approximation of \tilde{f}^* , however, the closeness guarantees are affected. Accordingly, we will run the gradient descent procedure in each respective algorithm $\log n$ times with $\tilde{f}^* = f_2^* \cdot 2^i$ for $i = 0$ to $\log n$, and whichever iteration produces $\tilde{\mathbf{x}}$ that minimizes $f(\cdot)$ will be output. Lemma 47 implies that there must exist some i such that $f_2^* \cdot 2^i \leq f^* \leq f_2^* \cdot 2^{i+1}$. Therefore, when we run our algorithm with $\tilde{f}^* = f_2^* \cdot 2^{i+1}$, the algorithm will succeed with high probability. Thus, the overall success probability is at least as high as any individual run of the algorithm. Moreover, the output $\tilde{\mathbf{x}}$ is guaranteed to have $f(\tilde{\mathbf{x}}) - f^* \leq 2\epsilon f^*$. ■

17. Note that our sampled and rotated $\tilde{\mathbf{A}}\mathbf{U}$ from Lemma 12 loses any sparsity guarantees that \mathbf{A} may have had.