

Query Complexity of Least Absolute Deviation Regression via Robust Uniform Convergence

Xue Chen

George Mason University

XUECHEN@GMU.EDU

Michał Dereziński

University of California, Berkeley

MDEREZIN@BERKELEY.EDU

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

Consider a regression problem where the learner is given a large collection of d -dimensional data points, but can only query a small subset of the real-valued labels. How many queries are needed to obtain a $1 + \epsilon$ relative error approximation of the optimum? While this problem has been extensively studied for least squares regression, little is known for other losses. An important example is least absolute deviation regression (ℓ_1 regression) which enjoys superior robustness to outliers compared to least squares. We develop a new framework for analyzing importance sampling methods in regression problems, which enables us to show that the query complexity of least absolute deviation regression is $\Theta(d/\epsilon^2)$ up to logarithmic factors. We further extend our techniques to show the first bounds on the query complexity for any ℓ_p loss with $p \in (1, 2)$. As a key novelty in our analysis, we introduce the notion of *robust uniform convergence*, which is a new approximation guarantee for the empirical loss. While it is inspired by uniform convergence in statistical learning, our approach additionally incorporates a correction term to avoid unnecessary variance due to outliers. This can be viewed as a new connection between statistical learning theory and variance reduction techniques in stochastic optimization, which should be of independent interest.

Keywords: robust regression, query complexity, uniform convergence, Lewis weight sampling

1. Introduction

Consider a linear regression problem defined by an $n \times d$ data matrix A and a vector $y \in \mathbb{R}^n$ of labels (or responses). Our goal is to approximately find a vector $\beta \in \mathbb{R}^d$ that minimizes the total loss over the entire dataset, given by $L(\beta) = \sum_{i=1}^n l(a_i^\top \beta - y_i)$, where a_i^\top is the i th row of A . Suppose that we are only given the data matrix A and we can choose to query some number of individual labels y_i (with the remaining labels hidden). When the cost of obtaining those labels dominates other computational costs (for example, because it requires performing a complex and resource-consuming measurement), it is natural to ask how many queries are required to obtain a good approximation of the optimal solution over the entire dataset. This so-called query complexity arises in a number of statistical learning tasks such as active learning and experimental design.

Definition 1 (Query complexity) *For a given loss function $l : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, the query complexity $\mathcal{M} := \mathcal{M}_l(d, \epsilon, \delta)$ is the smallest number such that there is a randomized algorithm that, given any $n \times d$ matrix A and any hidden vector $y \in \mathbb{R}^n$, queries \mathcal{M} entries of y and returns $\tilde{\beta} \in \mathbb{R}^d$, so that:*

$$L(\tilde{\beta}) \leq (1 + \epsilon) \cdot \min_{\beta} L(\beta) \quad \text{with probability } 1 - \delta, \quad \text{where } L(\beta) = \sum_{i=1}^n l(a_i^\top \beta - y_i).$$

Note that since the algorithm has access to the entire dataset on which it will be evaluated, it can use that information to select the queries better than, say uniformly at random, which differentiates this problem from the traditional sample complexity in statistical learning theory. Also, note that we require a relative error approximation, as opposed to an additive one. This is because we are not restricting the range of the labels, so a relative error approximation provides a more useful scale-invariant guarantee (this is a common practice when approximating regression problems).

Significant work has been dedicated to various notions of query complexity in classification (for an overview, see [Hanneke, 2014](#)). On the other hand, in the context of regression, prior literature has primarily focused on the special case of least squares regression, $l(a) = a^2$, where the optimum solution has a closed form expression, which considerably simplifies the analysis. In this setting, the query complexity is known to be $\mathcal{M} = O(d/\epsilon)$ ([Chen and Price, 2019](#)). An important drawback of the square loss in linear regression is its sensitivity to outliers, and to that end, a number of other loss functions are commonly used in practice to ensure robustness to outliers. Those losses no longer yield a closed form solution, so the techniques from least squares do not apply. The primary example is *least absolute deviation* regression, also known as ℓ_1 regression, which uses the loss function $l(a) = |a|$. A natural way of interpolating between the robustness of the ℓ_1 loss and the smoothness of least squares, is to use an ℓ_p loss, i.e., $l(a) = |a|^p$ for $p \in (1, 2)$. Other popular choices include the Huber loss and the Tukey loss; for more discussion on robust regression see [Clarkson and Woodruff \(2015\)](#). The basic question that motivates this paper is: What is the query complexity of robust regression? This question has remained largely open even for the special case of least absolute deviation regression.

1.1. Main Results

Our first main result gives nearly-matching (up to logarithmic factors) upper and lower bounds for the query complexity of least absolute deviation regression. While the randomized algorithm that achieves the upper bound is in fact based on an existing importance sampling method, our key contribution is to provide a new analysis of this method that overcomes a significant limitation of the prior work, resulting in the first non-trivial guarantee for query complexity.

Theorem 2 *The query complexity of least absolute deviation regression, $l(a) = |a|$, satisfies:*

$$\Omega((d + \log(1/\delta))/\epsilon^2) \leq \mathcal{M}_l(d, \epsilon, \delta) \leq O(d \log(d/\epsilon\delta)/\epsilon^2).$$

To obtain the upper bound, we use the so-called Lewis weight sampling originally due to [Lewis \(1978\)](#) (see [Section 2](#)), which can be implemented in time that is nearly-linear in the support size of the matrix A ([Cohen and Peng, 2015](#)). Lewis weight sampling is known to provide a strong guarantee called the ℓ_1 subspace embedding property ([Talagrand, 1990](#); [Cohen and Peng, 2015](#)), which is an important tool in randomized linear algebra for approximately solving least absolute deviation regression. However, this strategy only works if the algorithm has unrestricted access to the vector y when computing the weights (e.g., when the motivation is computational efficiency rather than query complexity), or if we require only a constant factor approximation (i.e., $\epsilon > 1$).¹ So, prior work on Lewis weight sampling provides no guarantees for query complexity with $\epsilon \leq 1$.

1. A constant factor approximation with $\epsilon = 7$ and $\delta = 0.4$ is folklore (e.g., Theorem 6 in [Dasgupta et al. \(2009\)](#)). To obtain a more accurate approximation, one can use the ℓ_1 subspace embedding property with respect to the matrix $[A; y]$ (which has $d + 1$ columns) instead of A , however this requires unrestricted access to the vector y .

In this paper, we develop a new analysis of Lewis weight sampling, which relies on what we call *robust uniform convergence*: an approximation guarantee for the empirical loss that is similar to uniform convergence in statistical learning theory, except it incorporates a correction to reduce the variance due to outliers. This leads to an intriguing connection between uniform convergence in statistical learning and variance reduction techniques in stochastic optimization (see Section 1.2). Our new approach not only allows us to compute the weights without accessing the label vector, thereby obtaining guarantees for query complexity, but it also leads to simplifications of existing methods for approximately solving the regression problem, both in terms of the algorithms and the analysis (see Section 1.3). Our analysis framework is not specific to the ℓ_1 loss, and is likely of interest beyond robust regression. To demonstrate this, our second main result uses robust uniform convergence to provide an upper bound on the query complexity of regression with any ℓ_p loss where $p \in (1, 2)$, which is again the first non-trivial guarantee of this kind.

Theorem 3 *For any $p \in (1, 2)$, the query complexity of ℓ_p regression, $l(a) = |a|^p$, satisfies:*

$$\mathcal{M}_l(d, \epsilon, \delta) \leq O(d^2 \log(d/\epsilon\delta)/\epsilon^2\delta).$$

While this result also relies on Lewis weight sampling, the analysis is considerably more challenging, since it requires interpolating between the techniques for ℓ_1 and ℓ_2 losses. We expect that the upper bound can be improved to match the one from Theorem 2 in terms of the dependence on d , and we leave this as a new direction for future work.

1.2. Key Technique: Robust Uniform Convergence

Next, we give a brief overview of the proof techniques used to obtain the main results, focusing on the notion of robust uniform convergence which is central to our analysis.

A key building block in the algorithms for approximately solving regression problems is randomized sampling, which is used to construct an estimate of the loss $L(\beta) = \sum_{i=1}^n l(a_i^\top \beta - y_i)$. This involves constructing a sparse sequence of non-negative random variables s_1, \dots, s_n , such that $\tilde{L}(\beta) = \sum_{i=1}^n s_i l(a_i^\top \beta - y_i)$ is an unbiased estimate of $L(\beta)$, i.e., $\mathbb{E}[\tilde{L}(\beta)] = L(\beta)$. The algorithm then solves the regression problem defined by the loss estimate, $\tilde{\beta} = \arg \min_{\beta} \tilde{L}(\beta)$, and returns $\tilde{\beta}$ as the approximate solution to the original problem. To minimize the number of queries, we have to make sure that the number of non-zero s_i 's is small (if $s_i > 0$, then we must query y_i), while at the same time preserving the quality of the approximate solution.

A standard approach for establishing guarantees in statistical learning is via the framework of uniform convergence: Showing that $\sup_{\beta} |L(\beta) - \tilde{L}(\beta)|$ vanishes at some rate with the sample size (for an overview, see Vapnik, 1999, 2013). This approach can be adapted to the setting of relative error approximation, requiring that for sufficiently large sample size:

$$\text{(uniform convergence of relative error)} \quad \sup_{\beta} \frac{|L(\beta) - \tilde{L}(\beta)|}{L(\beta)} \leq \epsilon \quad \text{with probability } 1 - \delta.$$

This has proven successful for importance sampling in regression problems where the access to label vector y is unrestricted and the primary aim is computational efficiency. However, in the query model, it may happen that one entry y_i (unknown to the algorithm) is an outlier which significantly contributes to the loss $L(\beta)$, and without sampling that entry we will not obtain a good estimate $\tilde{L}(\beta)$. In the worst-case, a randomized algorithm that is oblivious to y may need to sample almost

all of the entries before discovering the outlier. Does this mean that it is impossible to obtain a relative error approximation of the optimum without catching the outlier? In fact it does not, and this is particularly intuitive in the case of least absolute deviation regression, where our goal is specifically to ignore the outliers. For this reason, we use a modified version of the uniform convergence property, where the contribution of the outliers is subtracted from the loss, using a correction term denoted by Δ in the following definition.

Definition 4 (Robust uniform convergence) *A randomized algorithm satisfies the robust uniform convergence property with query complexity $m(d, \epsilon, \delta)$ if, given any $n \times d$ matrix A and hidden vector $y \in \mathbb{R}^n$, it queries $m(d, \epsilon, \delta)$ entries of y and then with probability $1 - \delta$ returns $\tilde{L}(\cdot)$ s.t.:*

$$\sup_{\beta} \frac{|L(\beta) - \tilde{L}(\beta) - \Delta|}{L(\beta)} \leq \epsilon, \quad \text{where } \Delta = L(\beta^*) - \tilde{L}(\beta^*), \quad \beta^* = \arg \min_{\beta} L(\beta). \quad (1)$$

Note that as long as $\tilde{L}(\cdot)$ is an unbiased estimate of $L(\cdot)$, then Δ is a mean zero correction of the error quantity in uniform convergence. This bares much similarity to variance reduction techniques which have gained considerable attention in stochastic optimization (for an overview, see [Gower et al., 2020](#)). The key difference is that we are using the correction purely for the analysis, and not for the algorithm. Thus, whereas in optimization algorithms such as Stochastic Variance Reduced Gradient (SVRG, [Johnson and Zhang, 2013](#)), one must explicitly compute the correction based on an estimate of β^* , in our setting one never has to compute Δ , so we can simply use β^* itself to define the correction. Nevertheless, the analogy is apt in that the corrected error quantity will in fact have reduced variance, particularly in the presence of outliers, which is what enables our query complexity analysis.

The guarantee from (1) immediately implies that the regression estimate $\tilde{\beta} = \arg \min_{\beta} \tilde{L}(\beta)$ is a relative error approximation of β^* , as long as $\epsilon < 1$:

$$L(\tilde{\beta}) - L(\beta^*) \leq \tilde{L}(\tilde{\beta}) - \tilde{L}(\beta^*) + \epsilon \cdot L(\tilde{\beta}) \leq \epsilon \cdot L(\tilde{\beta}),$$

where we used that $\tilde{L}(\tilde{\beta}) \leq \tilde{L}(\beta^*)$, and after some manipulations, we get $L(\tilde{\beta}) \leq (1 + \frac{\epsilon}{1-\epsilon}) \cdot L(\beta^*)$. Thus to bound the query complexity in Theorems 2 and 3, it suffices to bound the complexity of ensuring robust uniform convergence for a given loss function.

Finally, to obtain our results we must still use carefully chosen importance sampling to construct the loss estimate, where the importance weights depend on the data matrix A . Here, we use the Lewis weights, which is an extension of statistical leverage scores (see Section 2 for details) that is known to be effective in approximating ℓ_p losses. Existing guarantees for Lewis weight sampling (namely, the ℓ_p subspace embedding property) prove insufficient for establishing robust uniform convergence, so we develop new techniques, which are presented for $p = 1$ in Section 3 and for $p \in (1, 2)$ in Section 4. Moreover, in Section 2, we discuss Lewis weights and their connection to natural importance weights defined as $\max_{\beta} \frac{|a_i^\top \beta|^p}{\|A\beta\|_p^p}$ for each row a_i . Also, note that our results only require a constant factor approximation of Lewis weights, where the constant enters into the query complexity. As a consequence, we can obtain new relative error guarantees for uniform sampling, where the sample size depends on a notion of matrix coherence based on the degree of non-uniformity of Lewis weights.

1.3. Related Work

There is significant prior work related to query complexity and relative error approximations for linear regression, which we summarize below.

Much of the work on importance sampling for linear regression has been done in the context of Randomized Numerical Linear Algebra (RandNLA; see, e.g., [Drineas and Mahoney, 2016](#); [Dereziński and Mahoney, 2021](#)), where the primary goals are computational efficiency or reducing the size of the problem. This line of work was initiated by [Drineas et al. \(2006\)](#), using importance sampling via statistical leverage scores to obtain relative error approximations of ℓ_2 regression. Leverage score sampling is known to require $\Theta(d \log d + d/\epsilon)$ queries (see, e.g., [Dereziński et al., 2019](#)), where we let the failure probability δ be a small constant for simplicity. Other techniques, such as random projections (e.g., [Sarlos, 2006](#); [Clarkson and Woodruff, 2013](#)), have been used for efficiently solving regression problems, however those methods generally require unrestricted access to the label vector y . More recently, [Dereziński and Warmuth \(2017\)](#) considered the query complexity of ℓ_2 regression, showing that d queries are sufficient to obtain a relative error approximation with $\epsilon = d$, by using a non-i.i.d. importance sampling technique called Volume Sampling. Note that at least d queries are necessary for regression with any non-trivial loss ([Dereziński and Warmuth, 2018](#)). Efficient algorithms for Volume Sampling were given by [Dereziński et al. \(2018, 2019\)](#). A different non-i.i.d. sampling approach was used by [Chen and Price \(2019\)](#) to reduce the query complexity of least squares to $O(d/\epsilon)$, with a matching lower bound.

Sampling algorithms for regression with a general ℓ_p loss were studied by [Dasgupta et al. \(2009\)](#). The importance weights they used are closely related to Lewis weights, and most of our analysis can be adapted to work with those weights (we use Lewis weights because they yield a better polynomial dependence on d). They use a two-stage sampling scheme, where the importance weights in the first stage are computed without accessing the label vector. Their analysis of the first stage leads to $O(d^{2.5})$ query complexity for obtaining a relative error approximation with $\epsilon = O(1)$. However, to further improve the approximation, their second stage constructs refined importance weights using the entire label vector. Remarkably, using robust uniform convergence one can show that, at least for $p \in [1, 2]$, their second stage weights are not necessary (it suffices to use the weights from the first stage with an adjusted sample size), simplifying both the algorithm and the analysis. Other randomized methods have been proposed for robust regression (e.g., see [Clarkson, 2005](#); [Li et al., 2014](#); [Meng and Mahoney, 2013](#); [Duffee et al., 2018](#)). In particular, [Clarkson and Woodruff \(2015\)](#) provide a more general framework that includes the Huber loss. However, all of these approaches require unrestricted access to the label vector, and so they do not imply any bounds for query complexity. Finally, [Cohen and Peng \(2015\)](#) proposed to use Lewis weights as an improvement to the importance weights of [Dasgupta et al. \(2009\)](#), however their results again focus on computational efficiency. We discuss this in more detail in Section 2.

Query complexity has been studied extensively in the context of active learning (see, e.g., [Cohn et al., 1994](#); [Balcan et al., 2009](#); [Hanneke, 2014](#)). These works focus mostly on classification problems, where one can take advantage of adaptivity: selecting next query based on the previously obtained labels. Our framework (Definition 1) does allow for such adaptivity, however, in the context of regression it appears to be of limited use beyond the initial selection of sampling weights. Approaches without adaptivity are sometimes referred to as pool-based active learning ([Sugiyama and Nakajima, 2009](#)) or experimental design ([Chaloner, 1984](#); [Dereziński et al., 2019](#)).

Finally, in classification problems where each $y_i \in \{0, 1\}$, prior work considered a variance reducing correction term to obtain faster rates of uniform convergence (for an overview, see [Boucheron et al., 2005](#)). Despite some similarities, our notion of robust uniform convergence serves a different purpose in that it is used to remove the effect of outliers in a regression problem where y_i could be unbounded.

2. Preliminaries

Throughout this work, let A denote the data matrix of dimension $n \times d$ and a_1, \dots, a_n be its n rows such that $A^\top = [a_1^\top, \dots, a_n^\top]$. Then let $\beta \in \mathbb{R}^d$ denote the coefficient vector and $y \in \mathbb{R}^n$ be the hidden vector of labels. When p is clear, β^* denotes the minimizer of ℓ_p regression $L(\beta) = \|A\beta - y\|_p^p$, where $\|\cdot\|_p$ denotes the ℓ_p vector norm.

We always use \tilde{L} to denote the empirical loss $\tilde{L}(\beta) = \sum_{i=1}^n s_i \cdot |a_i^\top \beta - y_i|^p$ with weights (s_1, \dots, s_n) . When we write down the weights as a diagonal matrix $S = \text{diag}(s_1^{1/p}, \dots, s_n^{1/p})$, the loss becomes $\tilde{L}(\beta) = \|SA\beta - Sy\|_p^p$. Since it is more convenient to study the concentration of the random variable $\tilde{L}(\beta) - \tilde{L}(\beta^*)$ around its mean $L(\beta) - L(\beta^*)$, we will rewrite (1) as

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm \epsilon \cdot L(\beta) \quad \forall \beta \in \mathbb{R}^d,$$

where $a = b \pm \epsilon$ means that a and b are ϵ -close, i.e., $a \in [b - \epsilon, b + \epsilon]$. Furthermore, for two variables (or numbers) a and b , we use $a \approx_\alpha b$ to denote that a is an α -approximation of b for $\alpha \geq 1$, i.e., that $b/\alpha \leq a \leq \alpha \cdot b$. Also, let $\text{Poisson}(\lambda)$ denote the Poisson random variable with mean λ and $\text{sign}(x)$ be the sign function which is 0 for $x = 0$ and $x/|x|$ for $x \neq 0$.

Importance weights and Lewis weights. Given A and the norm ℓ_p , we define the importance weight of a row a_i to be

$$\sup_{\beta} \frac{|a_i^\top \beta|^p}{\|A\beta\|_p^p}.$$

Note that this definition is rotation free, i.e., for any rotation matrix $R \in \mathbb{R}^{d \times d}$, vector $R^\top a_i$ has the same importance weight in AR since we could replace β in the above definition by $R^{-1}\beta$.

While most of our results hold for importance weights, it will be more convenient to work with Lewis weights ([Lewis, 1978](#)), which have efficient approximation algorithms and provide additional useful properties. The Lewis weights (w_1, \dots, w_n) of $A \in \mathbb{R}^{n \times d}$ are defined as the unique solution ([Lewis, 1978](#); [Cohen and Peng, 2015](#)) of the following set of equations:

$$\forall i \in [n], a_i^\top \cdot (A^\top \cdot \text{diag}(w_1, \dots, w_n)^{1-2/p} \cdot A)^{-1} a_i = w_i^{2/p}. \quad (2)$$

Note that Lewis weights always satisfy $\sum_i w_i = d$. We will use extensively the following relation between the importance weights and Lewis weights, which is potentially of independent interest.

Theorem 5 *Given any $p \in [1, 2]$ and matrix A , let (w_1, \dots, w_n) be the Lewis weights of A defined in (2). Then the importance weight of every row i in A is bounded by*

$$\max_{\beta \in \mathbb{R}^d} \frac{|a_i^\top \beta|^p}{\|A\beta\|_p^p} \in \left[d^{-(1-p/2)} \cdot w_i, w_i \right].$$

We remark that when $p = 2$, the Lewis weights are equal to the statistical leverage scores of A (the i th leverage score is defined as $a_i^\top (A^\top A)^\dagger a_i$, where \dagger denotes the Moore-Penrose pseudoinverse), which are known to be equal to the ℓ_2 importance weights (Spielman and Srivastava, 2011; Chen and Price, 2019). The proof of Theorem 5 is via a study of the relationships between importance weights, statistical leverages scores, and Lewis weights. In particular, our analysis (Lemma 17 and Claim 20) also gives the sample complexity of uniform sampling for ℓ_p regression by comparing Lewis weights with uniform sampling in terms of the non-uniformity of the statistical leverage scores. We defer the detailed discussions and proofs to Appendix A.

Computing Lewis weights. Given two sequences of weights (w'_1, \dots, w'_n) and (w_1, \dots, w_n) , we say w' is a γ -approximation of w if $w'_i \approx_\gamma w_i$ for every $i \in [n]$. Now we invoke the contraction algorithm by Cohen and Peng (2015) to approximate w_i . While we state it for a $(1+\epsilon)$ -approximation, our results only need a constant approximation say $\gamma = 2$.

Lemma 6 [Theorem 1.1 in Cohen and Peng (2015)] *Given any matrix $A \in \mathbb{R}^{n \times d}$ and $p \in [1, 2]$, there is an algorithm that runs in time $\log(\log \frac{n}{\epsilon}) \cdot O(\mathbf{nnz}(A) \log n + d^{\omega+o(1)})$ to output a $(1+\epsilon)$ -approximation of the Lewis weights of A , where $\mathbf{nnz}(A)$ denotes the number of nonzero entries in A and ω is the matrix-multiplication exponent.*

Subspace embedding. Since the diagonal matrix $S = \text{diag}(s_1^{1/p}, \dots, s_n^{1/p})$ will be generated from the Lewis weights with support size $O(\frac{d \log \frac{d}{\epsilon \delta}}{\epsilon^2})$ for $p = 1$ and $O(\frac{d^2 \log d / \epsilon \delta}{\epsilon^2 \delta})$ for $p \in (1, 2]$, by the main results in Cohen and Peng (2015), we can assume that SA satisfies the ℓ_p subspace embedding property, given below.

Fact 7 *Given any $p \in [1, 2]$ and A , the sketch matrix S generated in Theorem 8 for $p = 1$ or Theorem 10 for $p \in (1, 2]$ has the following property: With probability $1 - \delta$,*

$$\|SA\beta\|_p^p \approx_{1+\epsilon} \|A\beta\|_p^p \quad \text{for any } \beta.$$

3. Near-optimal Analysis for ℓ_1 Regression

We consider the ℓ_1 loss function $L(\beta) = \|A\beta - y\|_1$ in this section. The main result is to prove the upper bound in Theorem 2, namely that we can use Lewis weight sampling with support size $\tilde{O}(d/\epsilon^2)$ to construct $\tilde{L}(\beta)$ that satisfies the robust uniform convergence property.

Theorem 8 *Given an $n \times d$ matrix A and $\epsilon, \delta \in (0, 1)$, let (w'_1, \dots, w'_n) be a γ -approximation of the Lewis weights (w_1, \dots, w_n) of A , i.e., $w'_i \approx_\gamma w_i$ for all $i \in [n]$. Let $u = \Theta(\frac{\epsilon^2}{\log \gamma d / \delta \epsilon})$ and $p_i = \min\{\gamma w'_i / u, 1\}$. For each $i \in [n]$, with probability p_i , set $s_i = 1/p_i$ (otherwise, let $s_i = 0$).*

Let $L(\beta) = \|A\beta - y\|_1$ with an unknown label vector $y \in \mathbb{R}^n$. With probability $1 - \delta$, the support size of s is $O(\frac{\gamma^2 d \cdot \log \gamma d / \epsilon \delta}{\epsilon^2})$ and $\tilde{L}(\beta) := \sum_{i=1}^n s_i \cdot |a_i^\top \beta - y_i|$ satisfies robust uniform convergence:

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm \epsilon \cdot L(\beta) \quad \text{for all } \beta.$$

First of all, we bound the sample size (i.e., the support size of s). Since $\sum_i w_i = d$ and $w'_i \approx_\gamma w_i$, let $m := \sum_i p_i$ denote the expected sample size, which is at most $\gamma/u \cdot \sum_i w'_i \leq \gamma^2 d / u$. Since the variance of $|\text{supp}(s)|$ is at most m , with probability $1 - \delta/2$, the support size of

s is $O(m + \sqrt{m \cdot \log 1/\delta} + \log 1/\delta) = O(\frac{\gamma^2 d \cdot \log \gamma d / \varepsilon \delta}{\varepsilon^2})$ by the Bernstein inequalities. Moreover, Lemma 6 provides an efficient algorithm to generate s with $\gamma = 2$.

In the rest of this section, we prove the robust uniform convergence property in Theorem 8. Before showing the formal proof in Section 3.1, we discuss the main technical tool — a concentration bound for all β with bounded $\|A(\beta - \beta^*)\|_1$.

Lemma 9 *Given $A \in \mathbb{R}^{n \times d}$, ε and a failure probability δ , let w_1, \dots, w_n be the Lewis weight of each row A_i . Let p_1, \dots, p_n be a sequence of numbers upper bounding w , i.e.,*

$$p_i \geq \min\{w_i/u, 1\} \text{ for } u = \Theta\left(\frac{\varepsilon^2}{\log(m/\delta + d/\varepsilon\delta)}\right) \text{ where } m = \sum_i p_i.$$

Suppose that the coefficients (s_1, \dots, s_n) in $\tilde{L}(\beta) := \sum_{i=1}^n s_i \cdot |a_i^\top \beta - y_i|$ are generated as follows: For each $i \in [n]$, with probability p_i , we set $s_i = 1/p_i$. Given any subset $B \subset \mathbb{R}^d$ of β , with probability at least $1 - \delta$,

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm \varepsilon \cdot \sup_{\beta \in B} \|A(\beta - \beta^*)\|_1 \quad \text{for all } \beta \in B.$$

We remark that Lemma 9 holds for any choice of B while the error depends on its radius via the term $\varepsilon \cdot \sup_{\beta \in B} \|A(\beta - \beta^*)\|_1$. For example, we could apply this lemma to $B := \{\beta \mid \|A(\beta - \beta^*)\|_1 \leq 3L(\beta^*)\}$ to conclude the error is at most $\varepsilon \cdot \|A(\beta - \beta^*)\|_1 \leq 3\varepsilon \cdot L(\beta^*)$, which is at most $3\varepsilon \cdot L(\beta)$ by the definition of β^* . However, this only gives guarantee for bounded β . To bound the error in terms of $\varepsilon L(\beta)$ for all $\beta \in \mathbb{R}^d$ with high probability, we partition $\beta \in \mathbb{R}^d$ into several subsets and apply Lemma 9 to those subsets separately. We defer the detailed proof of Theorem 8 to Section 3.1.

The proof of Lemma 9 is based on the contraction principle from Ledoux and Talagrand (1991) and chaining arguments introduced in Talagrand (1990); Cohen and Peng (2015). The main observation in its proof is that after rewriting $\tilde{L}(\beta) - \tilde{L}(\beta^*) = \sum_{i=1}^n s_i \cdot (|a_i^\top \beta - y_i| - |a_i^\top \beta^* - y_i|)$, all vectors in this summation, $(|a_i^\top \beta - y_i| - |a_i^\top \beta^* - y_i|)_{i \in [n]}$, contract from the vector $(|a_i^\top (\beta - \beta^*)|)_{i \in [n]}$ in the ℓ_1 summation. So, using the contraction principle (Ledoux and Talagrand, 1991), $\tilde{L}(\beta) - \tilde{L}(\beta^*)$ will concentrate at least as well as $\|SA(\beta - \beta^*)\|_1$, which can be analyzed using the ℓ_1 subspace embedding property (see Fact 7). We defer the formal proof to Appendix B.

3.1. Proof of Theorem 8

Recall that $m = \sum_i p_i \leq \gamma^2 d/u$ and $u = \Theta\left(\frac{\varepsilon^2}{d/\varepsilon\delta}\right)$. Since $\log m/\varepsilon\delta = O(\log d\gamma/\varepsilon\delta)$, we choose a small constant for u such that for some large C , $u \leq \frac{\varepsilon^2}{C \log(m/\varepsilon\delta + d/\varepsilon\delta)}$ and $p_i \geq \min\{1, w_i/u\}$.

For convenience, and without loss of generality, we assume $\beta^* = 0$ such that $L(\beta^*) = \|y\|_1$ (this is because we can always shift y to $y - A\beta^*$ and β^* to 0). Thus we can write

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = \|SA\beta - Sy\|_1 - \|Sy\|_1 \quad \text{and} \quad L(\beta) - L(\beta^*) = \|A\beta - y\|_1 - \|y\|_1$$

so that we can use the triangle inequality of the ℓ_1 norm. Next, notice that $\mathbb{E}_S[\|Sy\|_1] = \|y\|_1$. By Markov's inequality, $\|Sy\|_1 \leq \frac{1}{\delta} \cdot \|y\|_1$ holds with probability $1 - \delta$. We now apply Lemma 9 multiple times with different choices of B_β . We choose $t = O(1/\varepsilon^2\delta)$ and $B_i = \{\beta \mid \|A\beta\|_1 \leq$

$(3 + \varepsilon \cdot t)\|y\|_1\}$ for $i = 0, 1, \dots, t$. Then we apply Lemma 9 with failure probability $\frac{\delta}{t+1} = O(\delta^2 \varepsilon^2)$ to guarantee that with probability $1 - \delta$, we have for *all* $i = 0, 1, \dots, t$:

$$\|SA\beta - Sy\|_1 - \|Sy\|_1 = \|A\beta\|_1 - \|y\|_1 \pm \varepsilon \cdot \sup_{\beta \in B_i} \|A\beta\|_1 \quad \text{for all } \beta \in B_i.$$

Furthermore, with probability $1 - \delta$ over S , we have the ℓ_1 subspace embedding property:

$$\forall \beta, \|SA\beta\|_1 = (1 \pm \varepsilon) \cdot \|A\beta\|_1.$$

We can now argue that the concentration holds for all β , by considering the following three cases:

1. β has $\|A\beta\|_1 < 3\|y\|_1$: From the concentration of B_0 , the error is $\leq \varepsilon \cdot 3\|y\|_1 \leq 3\varepsilon \cdot L(\beta)$.
2. β has $\|A\beta\|_1 \in [(3 + \varepsilon \cdot i) \cdot \|y\|_1, (3 + \varepsilon \cdot (i + 1)) \cdot \|y\|_1)$ for $i < t$: Let β' be the rescaling of β with $\|A\beta'\|_1 = (3 + \varepsilon \cdot i)\|y\|_1$ such that $\|A(\beta - \beta')\|_1 \leq \varepsilon\|y\|_1$. Then we rewrite $\|A\beta - y\|_1 - \|y\|_1$ as

$$\|A(\beta - \beta' + \beta') - y\|_1 - \|y\|_1 = \|A\beta' - y\|_1 - \|y\|_1 \pm \|A(\beta - \beta')\|_1 = \|A\beta' - y\|_1 - \|y\|_1 \pm \varepsilon\|y\|_1.$$

Similarly, we rewrite $\tilde{L}(\beta) - \tilde{L}(\beta^*)$ as $\|SA\beta - Sy\|_1 - \|Sy\|_1$ and bound it by

$$\begin{aligned} \|SA\beta' - Sy\|_1 - \|Sy\|_1 \pm \|SA(\beta - \beta')\|_1 &= \|SA\beta' - Sy\|_1 - \|Sy\|_1 \pm (1 + \varepsilon) \cdot \|A(\beta - \beta')\|_1 \\ &= \|SA\beta' - Sy\|_1 - \|Sy\|_1 \pm (1 + \varepsilon) \cdot \varepsilon\|y\|_1, \end{aligned}$$

where we use the ℓ_1 subspace embedding in the middle step. Next we use the guarantee of $\beta' \in B_i$ to bound the error between $\|A\beta' - y\|_1 - \|y\|_1$ and $\|SA\beta' - Sy\|_1$ by $\varepsilon \cdot \|A\beta'\|_1$. So the total error is $O(\varepsilon) \cdot (\|A\beta'\|_1 + \|y\|_1) = O(\varepsilon) \cdot L(\beta)$ since $L(\beta) \geq \|A\beta'\|_1/2$ by the triangle inequality.

3. β has $\|A\beta\|_1 \geq \frac{25}{\varepsilon\delta} \cdot \|y\|_1$: We always have $\|A\beta - y\|_1 - \|y\|_1 = \|A\beta\|_1 \pm 2\|y\|_1$ by the triangle inequality. On the other hand, the ℓ_1 subspace embedding and the bound on $\|Sy\|_1$ imply:

$$\|SA\beta - Sy\|_1 - \|Sy\|_1 = \|SA\beta\|_1 \pm 2\|Sy\|_1 = (1 \pm \varepsilon)\|A\beta\|_1 \pm \frac{2}{\delta}\|y\|_1.$$

Since $\|A\beta\|_1 \geq \frac{25}{\varepsilon\delta} \cdot \|y\|_1$, this becomes $(1 \pm \varepsilon)\|A\beta\|_1 \pm \frac{2}{\delta} \cdot \frac{\varepsilon\delta}{25}\|A\beta\|_1 = (1 \pm \varepsilon \pm \frac{2\varepsilon}{25})\|A\beta\|_1$. Again, the error is $2\varepsilon\|A\beta\|_1 = O(\varepsilon) \cdot L(\beta)/2$ by the triangle inequality.

4. General Analysis for ℓ_p Regression

In this section, we consider the ℓ_p loss function $L(\beta) := \sum_{i=1}^n |a_i^\top \beta - y_i|^p$ for a given $p \in (1, 2]$. The main result is to show the upper bound in Theorem 3, i.e., that $\tilde{L}(\beta)$, when generated properly according to the Lewis weights with sample size $\tilde{O}(\frac{d^2}{\varepsilon^2\delta})$, satisfies robust uniform convergence.

Theorem 10 *Given $p \in (1, 2]$ and a matrix $A \in \mathbb{R}^{n \times d}$, let (w'_1, \dots, w'_n) be a γ -approximation of the Lewis weights (w_1, \dots, w_n) of A , i.e., $w_i \approx_\gamma w'_i$ for all $i \in [n]$. For $m = O(\frac{\gamma \cdot d^2 \log d / \varepsilon \delta}{\varepsilon^2} + \frac{\gamma \cdot d^{2/p}}{\varepsilon^2 \delta})$ with a sufficiently large constant, we sample $s_i \sim \frac{d}{m \cdot w'_i} \cdot \text{Poisson}(\frac{m \cdot w'_i}{d})$ for each $i \in [n]$.*

Then with probability $1 - \delta$, $|\text{supp}(s)| = O(m)$ and $\tilde{L}(\beta) := \sum_{i=1}^n s_i \cdot |a_i^\top \beta - y_i|^p$ satisfies robust uniform convergence. Namely,

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm \varepsilon \cdot L(\beta) \quad \text{for any } \beta \in \mathbb{R}^d.$$

First of all, we bound the support size of s . Note that $\Pr[s_i > 0] = 1 - e^{-\frac{m \cdot w'_i}{d}} \leq \frac{m \cdot w'_i}{d}$. So $\mathbb{E}[|\text{supp}(s)|] \leq \sum_i \frac{m \cdot w'_i}{d} \leq \gamma \cdot m$. Since the variance of $|\text{supp}(s)|$ is bounded by $\gamma \cdot m$, with probability $1 - \delta/2$, the sample size is $O(\gamma m + \sqrt{\gamma m \cdot \log 1/\delta} + \log 1/\delta) = O(\gamma m)$ by the Bernstein inequalities. Also, Lemma 6 provides an efficient algorithm to generate s with $\gamma = 2$.

In the rest of this section, we outline the proof of the robust uniform convergence property in Theorem 10, while the formal proof is deferred to Appendix C. The proof has two steps.

The 1st step is a reduction to the case of uniform Lewis weights. Specifically, we create an equivalent problem with matrix $A' \in \mathbb{R}^{N \times d}$ and $y' \in \mathbb{R}^N$ such that $L(\beta) = \|A'\beta - y'\|_p^p$ and the Lewis weights of A' are almost uniform. Moreover, we show an equivalent way to generate (s_1, \dots, s_n) so that $\tilde{L}(\beta) = \|S'A'\beta - S'y'\|_p^p$. Thus, showing that the new empirical loss $\|S'A'\beta - S'y'\|_p^p$ satisfies robust uniform convergence with respect to $\|A'\beta - y'\|_p^p$ will imply the same property for SA .

In the 2nd step, we interpolate techniques for ℓ_1 and ℓ_2 losses to prove the property for the special case of almost uniform Lewis weights (namely A' in the above paragraph). Let us discuss the key ideas in this interpolation. For ease of exposition, we assume $\beta^* = 0$ and $A^\top A = I$. Furthermore let α denote the approximation parameter of the uniform sampling compared to the Lewis weights of A after the 1st step reduction, i.e., $w_i \approx_\alpha d/n$.

We will again consider the robust uniform convergence as a concentration of $\tilde{L}(\beta) - \tilde{L}(\beta^*)$, which can be written as $\sum_i s_i \cdot (|a_i^\top \beta - y_i|^p - |y_i|^p)$ from the assumption $\beta^* = 0$. Our proof heavily relies on the following two properties of A when its Lewis weights are almost uniform.

Fact 11 *Let A be an $n \times d$ matrix such that $A^\top A = I$. If $w_i \approx_\alpha d/n$ for every $i \in [n]$, then:*

1. *The leverage scores are almost uniform (recall that when $A^\top A = I$, the i th leverage score becomes $\|a_i\|_2^2$):*

$$\|a_i\|_2^2 \approx_{\alpha C_p} d/n \quad \text{for} \quad C_p = 4/p - 1. \quad (3)$$

2. *The importance weights of $\|A\beta\|_p^p$ are almost uniform:*

$$\forall \beta, \quad \forall i \in [n], \quad \frac{\|a_i^\top \beta\|_p^p}{\|A\beta\|_p^p} \leq w_i \leq \alpha \cdot d/n. \quad (4)$$

The first property is implied by Claim 20 in Appendix A and the 2nd one is from Theorem 5.

The major difference between the analyses of $p > 1$ and $p = 1$ is that we cannot use the contraction principle to remove the effect of outliers in $\sum_i s_i \cdot |a_i^\top \beta - y_i|^p$. For example, when $p = 2$, we always have the cross term $2 \sum_i s_i \cdot (a_i^\top \beta) y_i$ besides $\sum_i s_i ((a_i^\top \beta)^2 + y_i^2)$. For general $p \in (1, 2)$, the cross term is replaced by the 1st order Taylor expansion of $\sum_i s_i \cdot |a_i^\top \beta - y_i|^p$. While the assumption that $\beta^* = 0$ implies that its expectation is $\sum_i (a_i^\top \beta) y_i = 0$, we can only use Chebyshev's inequality to conclude that the cross term is small for a fixed β . Since the coefficient s_i is independent of y_i , we cannot rely on a stronger Chernoff/Bernstein type concentration. This turns out to be insufficient to use the union bound for all β .

So our key technical contribution in the 2nd step is to bound the cross term for all β simultaneously. To do this, we rewrite it as an inner product between β and $(\langle A[*], j \rangle, S^p \cdot y)_{j \in [d]}$ (say $p = 2$), where $A[*], j$ denotes the j th column of A . To save the union bound for β , we bound the ℓ_2 norms of these two vectors separately to apply the Cauchy-Schwartz inequality. While prior work (Chen and Price, 2019) for ℓ_2 loss provides a way to bound the 2nd vector $(\langle A[*], j \rangle, S^p \cdot y)_{j \in [d]}$

given bounded leverage scores in (3), the new ingredient is to bound $\|\beta\|_2$ in terms of $\|A\beta\|_p$ for $p < 2$ when the Lewis weights are uniform. This is summarized in the following claim, stated for general $p \in (1, 2)$, in which we bound the 1st order Taylor expansion of $\sum_i s_i \cdot |a_i^\top \beta - y_i|^p$.

Claim 12 *If $w_i \approx_\alpha d/n$ for every $i \in [n]$, then with probability at least $1 - \delta$, we have the following bound for all $\beta \in \mathbb{R}^d$:*

$$\sum_{i=1}^n s_i \cdot p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta = O\left(\sqrt{\frac{\alpha^{\frac{2+p}{p}} \cdot \gamma \cdot d^{2/p}}{\delta \cdot m}} \cdot \|A\beta\|_p \cdot \|y\|_p^{p-1}\right). \quad (5)$$

The rest of the proof (given in Appendix C) is centered on obtaining a strong concentration bound for the 2nd order Taylor expansion using (4). We remark that in the special case of $p = 2$, this analysis simplifies considerably, and it can be easily adapted to show an $O(\frac{d \log d}{\delta \epsilon^2})$ bound on the query complexity of robust uniform convergence for least squares regression.

5. Lower Bound for ℓ_1 Regression

In this section, we prove the information-theoretic lower bound in Theorem 2 for the query complexity of ℓ_1 regression.

Theorem 13 *Given any $\epsilon, \delta < 0.01$, and d , consider the $n \times d$ matrix A defined as follows:*

$$A^\top = \underbrace{[e_1^\top, \dots, e_1^\top]}_{n/d}, e_2, \dots, e_{d-1}, \underbrace{[e_d^\top, \dots, e_d^\top]}_{n/d}$$

for the canonical basis $e_1, \dots, e_d \in \mathbb{R}^d$ with a sufficiently large n . It takes $m = \Omega(\frac{d + \log 1/\delta}{\epsilon^2})$ queries to y to produce $\tilde{\beta}$ satisfying $\|A\tilde{\beta} - y\|_1 \leq (1 + \epsilon) \cdot \|A\beta^* - y\|_1$ with probability $1 - \delta$.

Our proof will make a reduction to the classical biased coin testing problem (Kleinberg, 2020), by relying on Yao's minmax principle so that we can consider a deterministic querying algorithm and a randomized label vector y .

We rely on the fact that the ℓ_1 regression problem $\|A\beta - y\|_1$ for A specified in Theorem 13 can be separated into d independent subproblems of dimension $n' \times 1$ for $n' := n/d$. First of all, let us consider the 1-dimensional subproblem, which is a reformulation of the biased coin test. In this subproblem, we will use the following two distributions with different biases to generate $y' \in \{\pm 1\}^{n'}$:

1. D_1 : Each label $y'_i = 1$ with probability $1/2 + \epsilon$ and -1 with probability $1/2 - \epsilon$.
2. D_{-1} : Each label $y'_i = 1$ with probability $1/2 - \epsilon$ and -1 with probability $1/2 + \epsilon$.

The starting point of our reduction is that a $(1 + \epsilon)$ -approximate solution of β^* could distinguish between the two biased coins of D_1 and D_{-1} .

Claim 14 *Let $d = 1$, $n' \geq \frac{100 \log 1/\delta}{\epsilon^2}$, and $A' = [1, \dots, 1]^\top$ whose dimension is $n' \times 1$. Let $L(\beta)$ for $\beta \in \mathbb{R}$ denote $\|A'\beta - y'\|_1$ for $y' \in \mathbb{R}^{n'}$ and β^* be the minimizer of $L(\beta)$.*

If y' is generated from D_1 , then with probability $1 - \frac{\delta}{100}$, we will have $\beta^ = 1$ and any $(1 + \epsilon)$ -approximation $\tilde{\beta}$ will be positive. Similarly, if $y' \sim D_{-1}$, then with probability $1 - \frac{\delta}{100}$ we will have $\beta^* = -1$ and any $(1 + \epsilon)$ -approximation $\tilde{\beta}$ will be negative.*

Proof Let us consider y' generated from D_1 of dimension $n' \geq \frac{100 \log 1/\delta}{\epsilon^2}$. Let n_1 and n_{-1} denote the number of 1s and -1 s in y' . From the standard concentration bound, with probability $1 - \frac{\delta}{100}$, we have $n_1 \geq (1/2 + \epsilon/2)n'$ for a random string y' from D_1 . So for this 1-dimensional problem, β^* minimizing $\|A'\beta - y'\|_1$ will equal 1 with loss $L(\beta^*) := \|A'\beta^* - y'\|_1 = 2 \cdot n_{-1} \leq (1 - \epsilon)n$. Let $\tilde{\beta}$ be any $(1 + \epsilon)$ -approximation s.t. $\|A'\tilde{\beta} - y'\|_1 \leq (1 + \epsilon) \cdot \|A'\beta^* - y'\|_1$.

We show that $\tilde{\beta}$ must be positive given $n_1 \geq (1/2 + \epsilon/2)n$. If $\tilde{\beta} \in (-1, 0]$, then the loss of $\tilde{\beta}$ is given by

$$n_1 \cdot (1 - \tilde{\beta}) + n_{-1} \cdot (1 + \tilde{\beta}) = n - \tilde{\beta}(n_1 - n_{-1}) \geq n \geq \frac{L(\beta^*)}{1 - \epsilon}.$$

Otherwise if $\tilde{\beta} \leq -1$, the ℓ_1 loss of $\tilde{\beta}$ becomes

$$n_1 \cdot (1 - \tilde{\beta}) + n_{-1} \cdot (-\tilde{\beta} - 1) = -\tilde{\beta} \cdot n + n_1 - n_{-1} \geq n \geq \frac{L(\beta^*)}{1 - \epsilon}.$$

In both cases, $L(\tilde{\beta})$ does not give a $(1 + \epsilon)$ -approximation of $L(\beta^*)$. Similarly when y is generated from D_{-1} and $n_{-1} \geq (1/2 + \epsilon/2)n$, then any $\tilde{\beta}$ that is a $(1 + \epsilon)$ -approximation of β^* , must be negative. \blacksquare

Given the connection between the 1-dimensional subproblem and the biased coin test in Claim 14, the $\Omega(\frac{\log 1/\delta}{\epsilon^2})$ lower bound in the main theorem comes from the classical lower bound for distinguishing between biased coins.

Lemma 15 *Let us consider the following game between Alice and Bob:*

1. *Let Alice generate $\alpha \sim \{\pm 1\}$ randomly.*
2. *Alice generates $y \in \{\pm 1\}^{n'}$ from D_α defined above and allows Bob to query m' entries in y .*
3. *After all queries, Bob wins the game only if he predicts α correctly.*

If Bob wants to win this game with probability $1 - \delta$, he needs to make $m' = \Omega(\frac{\log 1/\delta}{\epsilon^2})$ queries.

The proof of Lemma 15 follows from KL divergence or the variation distance (a.k.a. statistical distance). We refer to Theorem 1.5 in Kleinberg (2020) for a proof with minor modifications such as replacing the fair coin by one of bias $1/2 - \epsilon$.

Next, we discuss how to obtain the $\Omega(d/\epsilon^2)$ part. Due to the space constraint, we give a high-level overview and defer the formal proof of Theorem 13 to Appendix D. First of all, for each string in $b \in \{\pm 1\}^d$, we consider the n -dimensional distribution of y generated as $D_b = D_{b_1} \circ D_{b_2} \circ \dots \circ D_{b_d}$. Namely, each chunk of n/d bits of y is generated from D_1 or D_{-1} mentioned above.

The second observation is that when we generate $y \sim D_b$ for a random $b \in \{\pm 1\}^n$, then a $(1 + \epsilon)$ -approximation algorithm could decode almost all entries of b except a tiny fraction (w.h.p. based on Claim 14). Then we show the existence of $b \in \{\pm 1\}^d$ and its flip $b^{(i)}$ on the i th coordinate such that: (1) the algorithm could decode entry i in b and $b^{(i)}$ (w.h.p. separately) when $y \sim D_b$ or $D_{b^{(i)}}$; (2) the algorithm makes $O(m/d)$ queries to decode entry b_i . Such a pair of b and $b^{(i)}$ provides a good strategy for Bob to win the game in Lemma 15 with a constant high probability using $O(m/d)$ queries, which gives a lower bound of m .

6. Conclusions and Future Directions

We provided nearly-matching upper and lower bounds for the query complexity of least absolute deviation regression. In this setting, the learner is allowed to use a sampling distribution that depends on the unlabeled data. In the process, we proposed *robust uniform convergence*, a framework for showing sample complexity guarantees in regression tasks, which includes a correction term that minimizes the effect of outliers. Further, we extended our results to robust regression with any ℓ_p loss for $p \in (1, 2)$. Note that the guarantees in this paper also apply to data-oblivious sampling, in which case we pay an additional data-dependent factor in the sample complexity.

Many new directions for future work arise from our results. First of all, we expect that the query complexity of ℓ_p regression should match the guarantee for ℓ_1 (our current bounds exhibit quadratic dependence on the input dimension d for $p > 1$, and linear dependence for $p = 1$). Moreover, the query complexity question remains open for a number of other robust losses, such as the Huber loss and the Tukey loss. More broadly, we ask whether the robust uniform convergence framework can be used to obtain sample complexity bounds that are robust to the presence of outliers for other settings in statistical learning, e.g., for classification and non-linear regression problems, and with additive as well as multiplicative error bounds.

Acknowledgments

We would like to acknowledge DARPA, IARPA, NSF, and ONR via its BRC on RandNLA for providing partial support of this work. Our conclusions do not necessarily reflect the position or the policy of our sponsors, and no official endorsement should be inferred.

References

- Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Agnostic active learning. *Journal of Computer and System Sciences*, 75(1):78–89, 2009.
- Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: A survey of some recent advances. *ESAIM: probability and statistics*, 9:323–375, 2005.
- Kathryn Chaloner. Optimal bayesian experimental design for linear models. *The Annals of Statistics*, 12(1):283–300, 1984.
- Xue Chen and Eric Price. Active regression via linear-sample sparsification. In Alina Beygelzimer and Daniel Hsu, editors, *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pages 663–695. PMLR, 2019.
- Kenneth L Clarkson. Subgradient and sampling algorithms for ℓ_1 regression. In *Symposium on Discrete Algorithms: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, volume 23, pages 257–266. Citeseer, 2005.
- Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pages 81–90, New York, NY, USA, 2013. ACM.

- Kenneth L. Clarkson and David P. Woodruff. Sketching for M -estimators: A unified approach to robust regression. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 921–939. SIAM, 2015.
- Michael B. Cohen and Richard Peng. L_p row sampling by lewis weights. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 183–192, 2015.
- David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. *Machine learning*, 15(2):201–221, 1994.
- 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, 2009.
- Michał Dereziński and Michael W Mahoney. Determinantal point processes in randomized numerical linear algebra. *Notices of the American Mathematical Society*, 68(1):34–45, 2021.
- Michał Dereziński and Manfred K. Warmuth. Unbiased estimates for linear regression via volume sampling. In *Advances in Neural Information Processing Systems 30*, pages 3087–3096, Long Beach, CA, USA, 2017.
- Michał Dereziński and Manfred K. Warmuth. Reverse iterative volume sampling for linear regression. *Journal of Machine Learning Research*, 19(23):1–39, 2018.
- Michał Dereziński, Manfred K. Warmuth, and Daniel Hsu. Leveraged volume sampling for linear regression. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 2510–2519. Curran Associates, Inc., 2018.
- Michał Dereziński, Kenneth L. Clarkson, Michael W. Mahoney, and Manfred K. Warmuth. Minimax experimental design: Bridging the gap between statistical and worst-case approaches to least squares regression. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1050–1069, Phoenix, USA, 25–28 Jun 2019.
- Michał Dereziński, Manfred K. Warmuth, and Daniel Hsu. Correcting the bias in least squares regression with volume-rescaled sampling. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 944–953. PMLR, 16–18 Apr 2019.
- Michał Dereziński, Manfred K. Warmuth, and Daniel Hsu. Unbiased estimators for random design regression. *arXiv e-prints*, art. arXiv:1907.03411, Jul 2019.
- Petros Drineas and Michael W. Mahoney. RandNLA: Randomized numerical linear algebra. *Communications of the ACM*, 59:80–90, 2016.
- Petros Drineas, Michael W Mahoney, and S. Muthukrishnan. Sampling algorithms for ℓ_2 regression and applications. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1127–1136. Society for Industrial and Applied Mathematics, 2006.

- David Durfee, Kevin A. Lai, and Saurabh Sawlani. ℓ_1 regression using lewis weights preconditioning and stochastic gradient descent. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1626–1656. PMLR, 06–09 Jul 2018. URL <http://proceedings.mlr.press/v75/durfee18a.html>.
- Robert M Gower, Mark Schmidt, Francis Bach, and Peter Richtarik. Variance-reduced methods for machine learning. *Proceedings of the IEEE*, 108(11):1968–1983, 2020.
- Steve Hanneke. Theory of disagreement-based active learning. *Foundations and Trends® in Machine Learning*, 7(2-3):131–309, 2014.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26:315–323, 2013.
- Robert Kleinberg. *Lecture notes on Theoretical Foundations of Machine Learning*. Available at https://www.cs.cornell.edu/courses/cs6781/2020sp/lectures/15_sequential-notes.pdf, 2020.
- M. Ledoux and M. Talagrand. *Probability in Banach spaces*. Springer, 1991.
- D. Lewis. Finite dimensional subspaces of L_p . *Studia Mathematica*, 63(2):207–212, 1978.
- Mu Li, Gary L. Miller, and Richard Peng. Iterative row sampling. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 00:127–136, 2014. ISSN 0272-5428. doi: [doi:10.1109/FOCS.2013.22](https://doi.org/10.1109/FOCS.2013.22).
- 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, 2013.
- Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, FOCS ’06, pages 143–152, Washington, DC, USA, 2006. IEEE Computer Society.
- Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- Masashi Sugiyama and Shinichi Nakajima. Pool-based active learning in approximate linear regression. *Mach. Learn.*, 75(3):249–274, June 2009.
- Michel Talagrand. Embedding subspaces of ℓ_1 into ℓ_2 . *Proceedings of the American Mathematical Society*, 108(2):363–369, 1990. ISSN 00029939, 10886826.
- Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999.
- Vladimir N Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- Roman Vershynin. *High-Dimensional Probability*. Cambridge University Press, 2018.

Appendix A. Lewis Weights Bound Importance Weights

We finish the proof of Theorem 5 in this section. Specifically, given any matrix A , we show that Lewis weights (w_1, \dots, w_n) defined in Equation (2) bound the importance weights for each row:

$$\max_{x \in \mathbb{R}^d} \frac{|a_i^\top x|^p}{\|Ax\|_p^p} \in \left[d^{-(1-p/2)} \cdot w_i, w_i \right].$$

Remark 16 *There are matrices that obtain both upper and lower bounds in the above inequality (up to constants). For example, $A = I$ gives the upper bound.*

For the lower bound, consider $A \in \{\pm 1\}^{2^d \times d}$ constituted by $n = 2^d$ distinct vectors in $\{\pm 1/\sqrt{d}\}^d$. Then for each a_i , we set $x = a_i$ such that $a_i^\top x = 1$ but $w_i = d/n$ and $\|Ax\|_p^p \approx n \cdot \mathbb{E}_{Z \sim N(0, 1/d)} [|Z|^p] = \Theta(n \cdot d^{-p/2})$.

To prove Theorem 5, we will show the following lemma for the special case where the Lewis weights are almost uniform and make a reduction from the general case to the almost uniform case. Recall that $a \approx_\alpha b$ means $a \in [b/\alpha, b \cdot \alpha]$ and $W = \text{diag}(w_1, \dots, w_n)$ denotes the diagonal Lewis-weight matrix when the matrix A and parameter p are fixed.

Lemma 17 *Suppose the ℓ_p Lewis weights of A are almost uniform: for a parameter $\alpha \geq 1$, $w_i \approx_\alpha d/n$. Then for every row i , its importance weight satisfies:*

$$\max_{x \in \mathbb{R}^d} \frac{|a_i^\top x|^p}{\|Ax\|_p^p} \in \left[\alpha^{-O(1)} \cdot d^{p/2}/n, \alpha^{O(1)} \cdot d/n \right].$$

In particular, if $\alpha = 1$ (the Lewis weights are uniform), this indicates the importance weights are uniformly upper bounded. We defer the proof to Appendix A.2. We remark that Claim 20 in Appendix A.2 shows that when the statistical leverage scores are almost uniform — $a_i^\top (A^\top A)^{-1} a_i \approx_\alpha d/n$ for some α , the ℓ_p Lewis weights have $w_i \approx_{\alpha^{O_p(1)}} d/n$ for all $i \in [n]$ and any $p \in [1, 2]$. This bounds the sample complexity of uniform sampling in terms of α by plugging $\gamma = \alpha^{O(1)}$ in Theorem 8 and Theorem 10.

While this could give a bound on the importance weights in terms of the Lewis weights for the non-uniform case, we will use the following properties to obtain the tight bound in Theorem 5 via a reduction.

Claim 18 *Given $A \in \mathbb{R}^{n \times d}$ whose Lewis weights are (w_1, \dots, w_n) , let $A' \in \mathbb{R}^{(n+k-1) \times d}$ be the matrix of splitting one row, say the last row a_n , into k copies: $a'_i = a_i$ for $i < n$ and $a'_i = a_n/k^{1/p}$ for $i \geq n$. Then the Lewis weights of A' are $(w_1, \dots, w_{n-1}, w_n/k, \dots, w_n/k)$.*

And the same property holds for the importance weight.

Claim 19 *Given $A \in \mathbb{R}^{n \times d}$ whose importance weights are (u_1, \dots, u_n) , let $A' \in \mathbb{R}^{(n+k-1) \times d}$ be the matrix of splitting the last row a_n into k copies: $a'_i = a_i$ for $i < n$ and $a'_i = a_n/k^{1/p}$ for $i \geq n$. Then the importance weights of A' are $(u_1, \dots, u_{n-1}, u_n/k, \dots, u_n/k)$.*

We defer the proofs of these two claims to Section A.1 and finish the proof of Theorem 5.

Proof of Theorem 5. Suppose the Lewis weights of A are (w_1, \dots, w_n) and the importance weights are (u_1, \dots, u_n) . Let ε be a tiny constant and $N_1 = \lceil w_1/\varepsilon \rceil, \dots, N_n = \lceil w_n/\varepsilon \rceil$. We define $A' \in \mathbb{R}^{N \times d}$ with $N = \sum_i N_i$ as

$$\begin{bmatrix} a_1/N_1^{1/p} \\ \vdots \\ a_1/N_1^{1/p} \\ a_2/N_2^{1/p} \\ \vdots \\ a_n/N_n^{1/p} \end{bmatrix}$$

where there are N_1 rows of $a_1/N_1^{1/p}$, N_2 rows of $a_2/N_2^{1/p}$, and so on. From Claim 18, we have the Lewis weights of A' are

$$\left(\underbrace{w_1/N_1, \dots, w_1/N_1}_{N_1}, \underbrace{w_2/N_2, \dots, w_2/N_2}_{N_2}, \dots, \underbrace{w_n/N_n, \dots, w_n/N_n}_{N_n} \right).$$

Since w_1, \dots, w_n are fixed, we know $w_i/N_i = w_i/\lceil w_i/\varepsilon \rceil \in [\varepsilon/(1 + \varepsilon/w_1), \varepsilon]$. So let α be the parameter satisfying $w_i/N_i \approx_\alpha d/N$. If w_1, \dots, w_n are multiples of ε , $d/N = \varepsilon$ and $\alpha = 1$. Since $\lim_{\varepsilon \rightarrow 0} d/N = \varepsilon$ (by the property $\sum_i w_i = d$ and the definition of N), we know $\lim_{\varepsilon \rightarrow 0} \alpha = 1$.

From Claim 19, we have the importance weights of A' are

$$\left(\underbrace{u_1/N_1, \dots, u_1/N_1}_{N_1}, \underbrace{u_2/N_2, \dots, u_2/N_2}_{N_2}, \dots, \underbrace{u_n/N_n, \dots, u_n/N_n}_{N_n} \right).$$

By Lemma 17, we have $w_i/N_i \in [\alpha^{-O(1)} \cdot d^{-(1-p/2)} \cdot u_1/N_1, \alpha^{O(1)} \cdot u_1/N_1]$ for each i . By taking $\varepsilon \rightarrow 0$ and $\alpha \rightarrow 1$, this shows both upper and lower bounds. \blacksquare

A.1. Proofs of Claim 18 and 19

To prove Claim 18, we use the property that Lewis weights constitute the unique diagonal matrix satisfying (2) (Cohen and Peng, 2015). So we only need to verify $(w_1, \dots, w_{n-1}, w_n/k, \dots, w_n/k)$ satisfying (2) for A' . First of all, we show the inverse in the equation is the same:

$$\begin{aligned} A'^\top (W')^{1-2/p} A' &= \sum_{i=1}^{n+k-1} (w'_i)^{1-2/p} \cdot (a'_i) \cdot (a'_i)^\top \\ &= \sum_{i=1}^{n-1} (w_i)^{1-2/p} a_i \cdot a_i^\top + \sum_{i=1}^k (w_n/k)^{1-2/p} \cdot (a_n/k^{1/p}) \cdot (a_n/k^{1/p})^\top \\ &= \sum_{i=1}^{n-1} (w_i)^{1-2/p} a_i \cdot a_i^\top + k^{2/p} \cdot w_n^{1-2/p} \cdot (a_n/k^{1/p}) \cdot (a_n/k^{1/p})^\top = \sum_{i=1}^n (w_i)^{1-2/p} a_i \cdot a_i^\top. \end{aligned}$$

Then it is straightforward to verify $(a_n/k^{1/p})^\top (A'^\top W'^{-1} A')^{1-2/p} \cdot (a_n/k^{1/p}) = (w_n/k)^{2/p}$.

Next We prove Claim 19. We note that for any x , we always have $\|Ax\|_p^p = \|A'x\|_p^p$ and $|(a'_i)^\top x|^p = |a_i^\top x|^p$ for any $i < n$. These two indicate $u_i = u'_i$ for $i < n$.

Now we prove $u_n = u'_n/k$. For any x , we still have $\|Ax\|_p^p = \|A'x\|_p^p$ but $|(a'_n)^\top x|^p = |a_n^\top x|^p/k$. These two indicate $u'_n = u_n/k$.

A.2. Almost Uniform Lewis Weights

Without loss of generality, we assume $A^\top A = I$. We recall the definition of leverage scores that will be used in this proof. Given $p = 2$ and A with rows a_1, \dots, a_n , the leverage score of a_i is $a_i^\top (A^\top A)^{-1} a_i$. Since we assume $A^\top A = I$, this simplifies the score to $\|a_i\|_2^2$.

Let us start with the case of the uniform Lewis weights for ease of exposition. Because $w_1 = w_2 = \dots = w_n = d/n$, we simplify the equation of the Lewis weights

$$a_i^\top \cdot (A^\top W^{-1} A)^{1-2/p} a_i = w_i^{2/p}$$

to $a_i^\top \cdot w_i^{2/p-1} \cdot I \cdot a_i = w_i^{2/p}$. This indicates that the leverage score $\|a_i\|_2^2 = w_i = d/n$ is also uniform.

Then we show the upper bound $\frac{|a_i^\top x|^p}{\|Ax\|_p^p} \leq d/n$. By Cauchy-Schwartz,

$$|a_i^\top x| \leq \|a_i\|_2 \cdot \|x\|_2 = \sqrt{d/n} \cdot \|x\|_2.$$

Next we lower bound $\|Ax\|_p^p$ by

$$\|Ax\|_p^p \cdot \|Ax\|_\infty^{2-p} \geq \|Ax\|_2^2. \quad (6)$$

Because $A^\top A = I$, the right hand side is $\|x\|_2^2$. So

$$\|Ax\|_p^p \geq \frac{\|x\|_2^2}{(\|x\|_2 \cdot \sqrt{d/n})^{2-p}} = \frac{\|x\|_2^p}{(d/n)^{1-p/2}}.$$

We combine the upper bound and lower bound to obtain

$$\frac{|a_i^\top x|^p}{\|Ax\|_p^p} \leq \frac{(\sqrt{d/n} \cdot \|x\|_2)^p}{\|x\|_2^p / (d/n)^{1-p/2}} = d/n.$$

The lower bound $\frac{|a_i^\top x|^p}{\|Ax\|_p^p} \geq \frac{d^{p/2}}{n}$ follows from choosing $x = a_i$ and replacing (6) by the Holder's inequality $\|Ax\|_p^p \leq (n)^{\frac{2-p}{2}} \cdot \|Ax\|_2^p$.

In the non-uniform case, we will use the stability of Lewis weights (Definition 5.1 and 5.2 in [Cohen and Peng \(2015\)](#)) to finish the proof. Consider any $\bar{W} = \text{diag}[\bar{w}_1, \dots, \bar{w}_n]$ satisfies

$$\forall i \in [n], a_i^\top \left(A^\top \bar{W}^{1-2/p} A \right)^{-1} a_i \approx_\alpha \bar{w}_i^{2/p}.$$

Lemma 5.3 in [Cohen and Peng \(2015\)](#) shows that for each $i \in [n]$, $\bar{w}_i \approx_{\alpha c_p} w_i$ for the Lewis weights (w_1, \dots, w_n) where the constant $c_p = \frac{p/2}{1-|p/2-1|}$.

Back to our problem, if the ℓ_p Lewis weights of A are almost uniform, we show the uniformity for its ℓ_q Lewis weights.

Claim 20 Given p and q less than 4, let $A \in \mathbb{R}^{n \times d}$ be a matrix whose ℓ_p Lewis weight (w_1, \dots, w_n) satisfies $w_i \approx_\alpha d/n$ for each $i \in [n]$. Then the ℓ_q Lewis weight of A satisfies $w'_i \approx_{\alpha^C} d/n$ for constant $C = (4/p - 1) \cdot c_q$. In particular, when $p = 2$ and $q = 1$, $C = 1$.

Proof Since $w_i \approx_\alpha d/n$, $\alpha^{-1}d/n \cdot I \preceq W \preceq \alpha d/n \cdot I$. We plug this sandwich-bound into $A^\top W^{1-2/p} A$, which becomes between $(\alpha d/n)^{1-2/p} \cdot A^\top A$ and $(\alpha^{-1}d/n)^{1-2/p} \cdot A^\top A$. Similarly, we bound its inverse $(A^\top W^{1-2/p} A)^{-1}$ by

$$(\alpha^{-1}d/n)^{2/p-1} \cdot (A^\top A)^{-1} \preceq (A^\top W^{1-2/p} A)^{-1} \preceq (\alpha d/n)^{2/p-1} \cdot (A^\top A)^{-1}.$$

From the definition of ℓ_p Lewis weights $a_i^\top (A^\top W^{1-2/p} A)^{-1} a_i = w_i^{2/p}$, we have

$$(\alpha^{-1}d/n)^{2/p-1} \cdot a_i^\top (A^\top A)^{-1} a_i \leq w_i^{2/p} \leq (\alpha d/n)^{2/p-1} \cdot a_i^\top (A^\top A)^{-1} a_i.$$

We combine the lower bound above with the property $w_i \approx_\alpha d/n$ to upper bound the leverage score $a_i^\top (A^\top A)^{-1} a_i$ by

$$(\alpha \cdot d/n)^{2/p} \cdot (\alpha)^{2/p-1} \cdot (d/n)^{1-2/p} = \alpha^{4/p-1} \cdot d/n.$$

Similarly, we lower bound the leverage score $a_i^\top (A^\top A)^{-1} a_i$ by $\alpha^{1-4/p} \cdot d/n$. These two imply the leverage scores are almost uniform: $a_i^\top (A^\top A)^{-1} a_i \approx_{\alpha^{4/p-1}} d/n$. So for the ℓ_q Lewis weights, we approximate it by $\bar{W} = d/n \cdot I$ and get

$$a_i^\top \left(A^\top \bar{W}^{1-2/q} A \right)^{-1} a_i = (d/n)^{2/q-1} \cdot a_i^\top (A^\top A)^{-1} a_i \approx_{\alpha^{4/p-1}} (d/n)^{2/q}.$$

From the stability of the Lewis weight, we know the actual ℓ_q Lewis weights (w'_1, \dots, w'_n) satisfy $w'_i \approx_{\alpha^{(4/p-1) \cdot c_q}} d/n$ ■

Now we restate Lemma 17 here.

Lemma 21 Given p , let the matrix A satisfy (1) $A^\top A = I$ and (2) its Lewis weights $w_i \approx_\alpha d/n$ for each row i . We have the ℓ_p importance weights of A satisfy

$$\max_x \frac{|a_i^\top x|^p}{\|Ax\|_p^p} \in \left[\alpha^{-Cp/2} \cdot d^{p/2}/n, \alpha^C \cdot d/n \right]$$

for $C = 4/p - 1$.

Proof For the upper bound, we have $|a_i^\top x| \leq \|a_i\|_2 \cdot \|x\|_2$ for all i . In particular, this implies $\|Ax\|_\infty \leq \sqrt{\alpha^C \cdot d/n} \cdot \|x\|_2$ for $C = 4/p - 1$ after plugging $q = 2$ into Claim 20 to bound $\|a_i\|_2^2$ by the leverage score. Then the rest of the proof is the same as the uniform case shown in the beginning of this section: We lower bound

$$\|Ax\|_p^p \geq \|Ax\|_2^2 / \|Ax\|_\infty^{2-p} = \|x\|_2^2 / (\sqrt{\alpha^C \cdot d/n} \cdot \|x\|_2)^{2-p}.$$

Now we upper bound $\max_x \frac{|a_i^\top x|^p}{\|Ax\|_p^p}$ by

$$\frac{(\sqrt{\alpha^C \cdot d/n} \cdot \|x\|_2)^p}{\|x\|_2^2 / (\sqrt{\alpha^C \cdot d/n} \cdot \|x\|_2)^{2-p}} \leq \alpha^C \cdot d/n.$$

For the lower bound, we choose $x = a_i$ such that $|a_i^\top x| = \|a_i\|_2 \cdot \|x\|_2$. Then by Holder's inequality, $\|Ax\|_p^p \leq (n)^{\frac{2-p}{2}} \cdot \|Ax\|_2^p$. These two show

$$\frac{|a_i^\top x|^p}{\|Ax\|_p^p} \geq \frac{(\sqrt{\alpha^{-C} \cdot d/n} \cdot \|x\|_2)^p}{(n)^{\frac{2-p}{2}} \cdot \|x\|_2^p} \geq \alpha^{-C \cdot p/2} \cdot d^{p/2}/n.$$

■

Appendix B. Additional Proofs from Section 3

Since any row with $p_i \geq 1$ will always be selected in S , those rows will not affect the random variable $\tilde{L}(\beta)$ considered in this section. We restrict our attention to rows with $p_i < 1$ in this proof. As a warm up, we will first prove the following lemma, which is a simplified version of Lemma 9.

Lemma 22 *Given $A \in \mathbb{R}^{n \times d}$ and ϵ , let w_1, \dots, w_n be the Lewis weight of each row A_i . Let p_1, \dots, p_n be a sequence of numbers upper bounding w , i.e.,*

$$p_i \geq w_i/u \text{ for } u = \Theta\left(\frac{\epsilon^2}{\log(m + d/\epsilon)}\right) \text{ where } m = \sum_i p_i.$$

Suppose the coefficients (s_1, \dots, s_n) in $\tilde{L}(\beta) := \sum_{i=1}^n s_i \cdot |a_i^\top \beta - y_i|$ is generated as follows: For each $i \in [n]$, with probability p_i , we set $s_i = 1/p_i$ (when $p_i \geq 1$, s_i is always 1). Given any subset $B_\beta \subset \mathbb{R}^d$ of β ,

$$\mathbb{E} \left[\sup_{\beta \in B_\beta} \left\{ \tilde{L}(\beta) - \tilde{L}(\beta^*) - (L(\beta) - L(\beta^*)) \right\} \right] \leq \epsilon \cdot \sup_{\beta \in B_\beta} \|A(\beta - \beta^*)\|_1.$$

We remark that while Lemma 22 implies that rescaling of u to $u = u/\delta^2$ gives

$$\text{w.p. } 1 - \delta, \tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm \epsilon \cdot \sup_{\beta \in B_\beta} \|A(\beta - \beta^*)\|_1 \text{ for all } \beta \in B_\beta,$$

Lemma 9 has a better dependency on δ .

We need a few ingredients about random Gaussian processes to finish the proofs of Lemma 22 and Lemma 9. Except the chaining arguments in Cohen and Peng (2015), the proof of Lemma 22 will use the following Gaussian comparison theorem (e.g. Theorem 7.2.11 and Exercise 8.6.4 in Vershynin (2018)). In the rest of this section, we use g to denote an i.i.d. Gaussian vector $N(0, 1)^d$ when the dimension is clear.

Theorem 23 (Slepian-Fernique) *Let v_0, \dots, v_n and u_0, \dots, u_n be two sets of vectors in \mathbb{R}^d where $v_0 = u_0 = \mathbf{0}$. Suppose that*

$$\|v_i - v_j\|_2 \geq \|u_i - u_j\|_2 \text{ for all } i, j = 0, \dots, n.$$

Then $\mathbb{E}_g [\max_i |\langle v_i, g \rangle|] \geq C_0 \cdot \mathbb{E}_g [\max_i |\langle u_i, g \rangle|]$ for some constant C_0 .

The proof of Lemma 9 follows the same outline except that we will use the following higher moments version for a better dependency on δ .

Corollary 24 (Corollary 3.17 of Ledoux and Talagrand (1991)) *Let v_0, \dots, v_n and u_0, \dots, u_n be two sets of vectors satisfying the conditions in the above Theorem. Then for any $\ell > 0$,*

$$4^\ell \cdot \mathbb{E}_g \left[\max_i |\langle v_i, g \rangle|^\ell \right] \geq \mathbb{E}_g \left[\max_i |\langle u_i, g \rangle|^\ell \right].$$

We will use the following concentration bound for ℓ_1 ball when the Lewis weights are uniformly small for all rows. It is an extension of Lemma 8.2 in [Cohen and Peng \(2015\)](#); but for completeness, we provide a proof in Section B.2.

Lemma 25 *Let A be a matrix with Lewis weight upper bounded by u . For any set $S \subseteq \mathbb{R}^d$,*

$$\mathbb{E}_g \left[\max_{\beta \in S} |\langle g, A\beta \rangle| \right] \lesssim \sqrt{u \cdot \log n} \cdot \max_{\beta \in S} \|A\beta\|_1.$$

We finish the proof of Lemma 22 here and defer the proof of Lemma 9 to Section B.1.

Proof of Lemma 22. For convenience, we define

$$\delta_i(\beta) := |a_i^\top \beta - y_i| \text{ such that } \tilde{L}(\beta) - \tilde{L}(\beta^*) = \sum_i s_i (\delta_i(\beta) - |y_i|)$$

given the assumption $\beta^* = 0$ and rewrite $L(\beta) - L(\beta^*) = \sum_i (\delta_i(\beta) - |y_i|)$ similarly.

So, we have

$$\mathbb{E}_S \left[\sup_{\beta \in B_\beta} \left\{ \tilde{L}(\beta) - \tilde{L}(\beta^*) - (L(\beta) - L(\beta^*)) \right\} \right] = \mathbb{E}_S \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \sum_{i=1}^n (\delta_i(\beta) - |y_i|) \right| \right].$$

Since each $(\delta_i(\beta) - |y_i|)$ is the expectation of $s_i \cdot (\delta_i(\beta) - |y_i|)$, from the standard symmetrization and Gaussianization in [Ledoux and Talagrand \(1991\)](#) (which is shown in Section B.3 for completeness), this is upper bounded by

$$\sqrt{2\pi} \mathbb{E}_S \left[\mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n g_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| \right] \right].$$

Then we fix S and plan to apply the Gaussian comparison Theorem 23 to the following process.

$$\mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \sum_j g_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| \right] = \mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \left\langle g, (s_j \cdot (\delta_j(\beta) - |y_j|))_{j \in [n]} \right\rangle \right| \right].$$

Let us verify the condition of the Gaussian comparison Theorem 23 to upper bound this by

$$C_2 \cdot \mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \left\langle g, SA\beta \right\rangle \right| \right].$$

Note that for any β and β' , the j th term of β and β' in the above Gaussian process is upper bounded by

$$|s_j \cdot (\delta_j(\beta) - |y_j|) - s_j \cdot (\delta_j(\beta') - |y_j|)| \leq s_j \cdot |\delta_j(\beta) - \delta_j(\beta')| \leq s_j \cdot |a_j^\top \beta - a_j^\top \beta'|.$$

At the same time, for the $\mathbf{0}$ vector, we always have

$$|s_j \cdot (\delta_j(\beta) - |y_j|) - 0| \leq s_j \cdot \left| |a_j^\top \beta - y_j| - |y_j| \right| \leq s_j \cdot |a_j^\top \beta|.$$

Hence, the Gaussian comparison Theorem 23 implies that the Gaussian process is upper bounded by

$$\mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \sum_j g_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| \right] \leq \mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \langle g, SA\beta \rangle \right| \right].$$

Next we plan to use Lemma 25 to bound this Gaussian process on $SA\beta$.

However, Lemma 25 requires that SA has bounded Lewis weight. To show this, the starting point is that if we replace a_i in (2) by $s_i \cdot a_i = \frac{1}{p_i} a_i$, its Lewis weight becomes $w_i/p_i \leq u$ by our setup. So we only need to handle the inverse in the middle of (2). The rest of the proof is very similar to the proof of Lemma 7.4 in Cohen and Peng (2015): We add a matrix A' into the Gaussian process to have this property of bounded Lewis weight. Recall that $u \leq \frac{\varepsilon^2}{C \log(m+d/\varepsilon)}$ for a large constant C . By sampling the rows in A with the Lewis weight and scaling every sample by a factor of u , we have the existence of A' with the following 3 properties (see Lemma B.1 in Cohen and Peng (2015) for the whole argument using the Lewis weight),

1. A' has $\tilde{O}(d/u)$ rows and each row has Lewis weight at most u .
2. The Lewis weight W' of A' satisfies $A'^\top \overline{W}'^{1-2/p} A' \succeq A^\top \overline{W}^{1-2/p} A$.
3. $\|A'x\|_1 = O(\|Ax\|_1)$ for all x .

Let A'' be the union of SA and A' (so the j th row of A'' is $s_j \cdot a_j$ for $j \leq [n]$ and $A'[j, *]$ for $j \geq n+1$). Because the first n entries of $A'' \cdot \beta$ are the same as $SA\beta$ for any β and the rest entries will only increase the energy, we have

$$\mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \langle g, SA\beta \rangle \right| \right] \leq \mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \langle g, A'' \cdot \beta \rangle \right| \right].$$

To apply Lemma 25, let us verify the Lewis weight of A'' is bounded. First of all, $A''^\top \overline{W}''^{1-2/p} A'' \succeq A^\top \overline{W}^{1-2/p} A$ from Lemma 5.6 in Cohen and Peng (2015). Now for each row $s_j a_j$ in A'' , its Lewis weight is upper bounded by

$$\left(\frac{1}{p_j} a_j^\top (A^\top \overline{W}^{1-2/p} A)^{-1} \cdot \frac{1}{p_j} a_j \right)^{1/2} \leq w_j/p_j \leq u$$

by the assumption $p_j \geq w_j/u$ in the lemma. At the same time, the Lewis weight of A' is already bounded by u from the definition, which indicates the rows in A'' added from A' are also bounded by u . So Lemma 25 implies

$$\begin{aligned} \mathbb{E}_g \left[\sup_{\beta \in B_\beta} \left| \langle g, A''\beta \rangle \right| \right] &\leq \sqrt{u \log(m+d/u)} \cdot \sup_{\beta \in B_\beta} \|A''\beta\|_1 \\ &\leq \sqrt{u \cdot \log(m+d/u)} \cdot \left(\sup_{\beta \in B_\beta} \|SA\beta\|_1 + O\left(\sup_{\beta \in B_\beta} \|A\beta\|_1 \right) \right) \end{aligned}$$

by the definition of A'' and the last property of A' . Finally, we bring the expectation of S back to bound $\mathbb{E}_S[\sup_{\beta \in B_\beta} \|SA\beta\|_1] \leq (1 + \varepsilon) \sup_{\beta \in B_\beta} \|A\beta\|_1$ by Lemma 7.4 in [Cohen and Peng \(2015\)](#) about the ℓ_1 subspace embedding.

From all discussion above, the deviation $\mathbb{E}_g \left[\sup_{\beta \in B_\beta} |\langle g, A''\beta \rangle| \right]$ is upper bounded by $O(\varepsilon) \cdot \sup_{\beta \in B_\beta} \|A\beta\|_1$ given our choice of u . \blacksquare

B.1. Proof of Lemma 9

The proof follows the same outline of the proof of Lemma 22 except using Corollary 24 to replace Theorem 23 with a higher moment. We choose the moment $\ell = O(\log n/\delta)$ and plan to bound $\mathbb{E}_S \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \sum_{i=1}^n (\delta_i(\beta) - |y_i|) \right|^\ell \right]$ for a better dependence on the failure probability. We apply symmetrization and Gaussianization again with different constants:

$$\begin{aligned}
 & \mathbb{E}_S \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \sum_{i=1}^n (\delta_i(\beta) - |y_i|) \right|^\ell \right] \\
 &= \mathbb{E}_S \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \mathbb{E}_{S'} \sum_{j=1}^n s'_j \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right] \\
 & \quad \text{(Use the convexity of } |\cdot|^\ell \text{ to move } \mathbb{E}_{S'} \text{ out)} \\
 &\leq \mathbb{E}_{S, S'} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \sum_{j=1}^n s'_j \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right] \\
 & \quad \text{(Use the symmetry of } S \text{ and } S') \\
 &\leq \mathbb{E}_{S, S', \varepsilon \in \{\pm 1\}^n} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \varepsilon_j \cdot (s_j - s'_j) \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right] \\
 & \quad \text{(Split } S \text{ and } S') \\
 &\leq \mathbb{E}_{S, S', \varepsilon \in \{\pm 1\}^n} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \varepsilon_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) - \sum_{j=1}^n \varepsilon_j \cdot s'_j \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right] \\
 & \quad \text{(Pay an extra } 2^\ell \text{ factor to bound the cross terms)} \\
 &\leq 2^\ell \cdot \mathbb{E}_{S, \varepsilon \in \{\pm 1\}^n} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \varepsilon_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right] + 2^\ell \cdot \mathbb{E}_{S', \varepsilon \in \{\pm 1\}^n} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \varepsilon_j \cdot s'_j \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right] \\
 &\leq 2^{\ell+1} \cdot \mathbb{E}_{S, \varepsilon \in \{\pm 1\}^n} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \varepsilon_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right] \\
 & \quad \text{(Gaussianize it)} \\
 &\leq 2^{\ell+1} \cdot \sqrt{\pi/2}^\ell \cdot \mathbb{E}_{S, g} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n g_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right|^\ell \right].
 \end{aligned}$$

Then we replace the vectors $(s_j \cdot (\delta_j(\beta) - |y_j|))_j$ by $SA\beta$ using the Gaussian comparison Corollary 24 (we omit the verification here because it is the same as the verification in the proof of Lemma 22):

$$C_1^\ell \cdot \mathbb{E}_S \mathbb{E}_g \left[\sup_{\beta} \left| \langle g, SA\beta \rangle \right|^\ell \right].$$

The proofs of Lemma 8.4 and Theorem 2.3 in Section 8 of [Cohen and Peng \(2015\)](#) show that

$$\mathbb{E}_S \mathbb{E}_g \left[\sup_{\beta} \left| \langle g, SA\beta \rangle \right|^\ell \right] \leq C_2^\ell \cdot \varepsilon^\ell \cdot \delta \cdot \sup_{\beta \in B_\beta} \|A\beta\|_1^\ell.$$

From all discussion above, we have

$$\mathbb{E}_S \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \sum_{i=1}^n (\delta_i(\beta) - |y_i|) \right|^\ell \right] \leq (C_1 C_2)^\ell \cdot \varepsilon^\ell \delta \cdot \sup_{\beta \in B_\beta} \|A\beta\|_1^\ell.$$

This implies with probability $1 - \delta$, the R.H.S. is at most $(C_1 C_2)^\ell \cdot \varepsilon^\ell \cdot \sup_{\beta \in B_\beta} \|A\beta\|_1^\ell$. So $\sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) = \sum_{i=1}^n (\delta_i(\beta) - |y_i|) \pm (C_1 C_2) \cdot \varepsilon \sup_{\beta \in B_\beta} \|A\beta\|_1$ for all $\beta \in B_\beta$. Finally, we finish the proof by rescaling ε by a factor of $C_1 \cdot C_2$.

B.2. Proof of ℓ_1 Concentration

We finish the proof of Lemma 25 in this section.

Proof of Lemma 25.

We consider the natural inner product induced by the Lewis weights (w_1, \dots, w_n) and define the weighted projection onto the column space of A with $W = \text{diag}(w_1, \dots, w_n)$ as $\Pi := A \cdot (A^\top W^{-1} A)^{-1} \cdot A^\top W^{-1}$. It is straightforward to verify $\Pi A = A$. Thus we write $\langle g, A\beta \rangle$ as $\langle g, \Pi A\beta \rangle$. In the rest of this proof, let Π_1, \dots, Π_n denote its n columns.

Then we upper bound the inner product in the Gaussian process by $\|\cdot\|_1 \cdot \|\cdot\|_\infty$:

$$\mathbb{E}_g \left[\max_{\beta \in S} \left| \langle \Pi^\top g, A\beta \rangle \right| \right] \leq \mathbb{E}_g \left[\max_{\beta \in S} \{ \|\Pi^\top g\|_\infty \cdot \|A\beta\|_1 \} \right].$$

So we further simplify the Gaussian process as

$$\mathbb{E}_g \left[\|\Pi^\top g\|_\infty \right] \cdot \max_{\beta \in S} \|A\beta\|_1 \leq \sqrt{2 \log n} \cdot \max_{i \in [n]} \|\Pi_i\|_2 \cdot \max_{\beta \in S} \|A\beta\|_1,$$

where we observe the entry i of $\Pi^\top g$ is a Gaussian variable with variance $\|\Pi_i\|_2^2$ and apply a union bound over n Gaussian variables. In the rest of this proof, we bound $\|\Pi_i\|_2^2$.

$$\begin{aligned}
 & \sum_j (A \cdot (A^\top W^{-1} A)^{-1} \cdot A^\top W^{-1})_{j,i}^2 \\
 &= \sum_j (w_i^{-1} \cdot a_i^\top \cdot (A^\top W^{-1} A)^{-1} \cdot a_j)^2 \\
 &\leq u \sum_j (w_i^{-1} \cdot a_i^\top \cdot (A^\top W^{-1} A)^{-1} \cdot a_j w_j^{-1/2})^2 \quad (w_j \leq u \text{ from the assumption}) \\
 &\leq u w_i^{-2} \cdot \sum_j a_i^\top \cdot (A^\top W^{-1} A)^{-1} \cdot a_j w_j^{-1} a_j^\top \cdot (A^\top W^{-1} A)^{-1} \cdot a_i \\
 &\leq u w_i^{-2} \cdot a_i^\top (A^\top W^{-1} A)^{-1} \cdot \left(\sum_j a_j w_j^{-1} a_j^\top \right) \cdot (A^\top W^{-1} A)^{-1} \cdot a_i \\
 &\leq u w_i^{-2} a_i^\top (A^\top W^{-1} A)^{-1} \cdot a_i \quad (\text{by the definition of Lewis weights}) \\
 &\leq u w_i^{-2} \cdot w_i^2 = u.
 \end{aligned}$$

■

B.3. Symmetrization and Gaussianization

We start with a standard symmetrization by replacing $\sum_{i=1}^n (\delta_i(\beta) - |y_i|)$ by its expectation $\mathbb{E}_{S'} \left[\sum_{j=1}^n s'_j \cdot (\delta_j(\beta) - |y_j|) \right]$:

$$\mathbb{E}_S \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \mathbb{E}_{S'} \left[\sum_{j=1}^n s'_j \cdot (\delta_j(\beta) - |y_j|) \right] \right| \right].$$

Using the convexity of the absolute function, we move out the expectation over S' and upper bound this by

$$\begin{aligned}
 & \mathbb{E}_{S, S'} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - \sum_{j=1}^n s'_j \cdot (\delta_j(\beta) - |y_j|) \right| \right] \\
 &= \mathbb{E}_{S, S'} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n s_j \cdot (\delta_j(\beta) - |y_j|) - s'_j \cdot (\delta_j(\beta) - |y_j|) \right| \right].
 \end{aligned}$$

Because S and S' are symmetric and each coordinate is independent, this expectation is equivalent to

$$\begin{aligned} & \mathbb{E}_{S, S', \sigma} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \sigma_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) - \sigma_j \cdot s'_j \cdot (\delta_j(\beta) - |y_j|) \right| \right] \\ & \leq \mathbb{E}_{S, S', \sigma} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \sigma_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| + \left| \sum_{j=1}^n \sigma_j \cdot s'_j \cdot (\delta_j(\beta) - |y_j|) \right| \right] \\ & \leq 2 \mathbb{E}_{S, \sigma} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \sigma_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| \right]. \end{aligned}$$

Then we apply Gaussianization: Since $\mathbb{E}[|g_j|] = \sqrt{2/\pi}$, the expectation is upper bounded by

$$\sqrt{2\pi} \mathbb{E}_{S, \sigma} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n \mathbb{E}[|g_j|] \sigma_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| \right].$$

Using the convexity of the absolute function again, we move out the expectation over g_j and upper bound this by

$$\sqrt{2\pi} \mathbb{E}_{S, \sigma, g} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n |g_j| \sigma_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| \right].$$

Now $|g_j| \sigma_j$ is a standard Gaussian random variable, so we simplify it to

$$\sqrt{2\pi} \mathbb{E}_{S, g} \left[\sup_{\beta \in B_\beta} \left| \sum_{j=1}^n g_j \cdot s_j \cdot (\delta_j(\beta) - |y_j|) \right| \right].$$

Appendix C. Additional Proofs from Section 4

We finish the proof of Theorem 10 in this section. As mentioned in Section 4, we state the following lemma for the 2nd step about matrices with almost uniform Lewis weights. It may be convenient to assume $\gamma = O(1)$, $\alpha = O(1)$, $w'_i \approx d/n$ such that the probability $\frac{m \cdot w'_i}{d} < 1$ in this statement.

Lemma 26 *Given any matrix A , let (w'_1, \dots, w'_n) be a γ -approximation of the Lewis weights (w_1, \dots, w_n) of A , i.e., $w'_i \approx_\gamma w_i$ for all $i \in [n]$. Further, suppose the Lewis weights are almost uniform: $w_i \approx_\alpha d/n$ for a given parameter α .*

Let $m = O\left(\frac{\alpha^{O(1)} \gamma \cdot d^2 \log d / \varepsilon \delta}{\varepsilon^2} + \frac{\alpha^{O(1)} \gamma \cdot d^{2/p}}{\varepsilon^2 \delta}\right)$. For each $i \in [n]$, we randomly generate $s_i = \frac{d}{m \cdot w'_i}$ with probability $\frac{m \cdot w'_i}{d}$ and 0 otherwise. Then with probability $1 - \delta$, for $\tilde{L}(\beta) := \sum_i s_i \cdot |a_i^\top \beta - y_i|^p$, we have

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm \varepsilon \cdot L(\beta) \text{ for any } \beta \in \mathbb{R}^d.$$

Now we are ready to finish the proof of Theorem 10.

Proof of Theorem 10. Because we could always adjust w'_i and m by a constant factor, let $\varepsilon \rightarrow 0$ be a tiny constant such that $N_1 = w'_1/\varepsilon, N_2 = w'_2/\varepsilon, \dots, N_n = w'_n/\varepsilon$ are integers. We define $A' \in \mathbb{R}^{N \times d}$ with $N = \sum_i N_i$ rows as

$$\begin{bmatrix} a_1/N_1^{1/p} \\ \vdots \\ a_1/N_1^{1/p} \\ a_2/N_2^{1/p} \\ \vdots \\ a_n/N_n^{1/p} \end{bmatrix}$$

where the first N_1 rows are of the form $a_1/N_1^{1/p}$, the next N_2 rows are of the form $a_2/N_2^{1/p}$, and so on. Similarly, we define $y' \in \mathbb{R}^N$ as

$$\left(\underbrace{y_1/N_1^{1/p}, \dots, y_1/N_1^{1/p}}_{N_1}, \underbrace{y_2/N_2^{1/p}, \dots, y_2/N_2^{1/p}}_{N_2}, \dots, \underbrace{y_n/N_n^{1/p}, \dots, y_n/N_n^{1/p}}_{N_n} \right).$$

By the definition, $L(\beta) = \|A\beta - y\|_p^p = \|A'\beta - y'\|_p^p$.

Then we consider $\tilde{L}(\beta)$. First of all, the Lewis weights of A' are

$$\left(\underbrace{w_1/N_1, \dots, w_1/N_1}_{N_1}, \underbrace{w_2/N_2, \dots, w_2/N_2}_{N_2}, \dots, \underbrace{w_n/N_n, \dots, w_n/N_n}_{N_n} \right).$$

Second, (w'_1, \dots, w'_n) on A induces weights

$$\bar{w} = \left(\underbrace{w'_1/N_1, \dots, w'_1/N_1}_{N_1}, \underbrace{w'_2/N_2, \dots, w'_2/N_2}_{N_2}, \dots, \underbrace{w'_n/N_n, \dots, w'_n/N_n}_{N_n} \right) \text{ on } A'.$$

Then we consider (s'_1, \dots, s'_N) generated in the way described in Lemma 26: For $j \in [N]$, $s'_j = \frac{d}{m \cdot \bar{w}_j}$ with probability $\frac{m \cdot \bar{w}_j}{d}$ and 0 otherwise. Let $i \in [n]$ be the row $a_i/N_i^{1/p}$ corresponding to a'_j , i.e., $j \in (N_1 + \dots + N_{i-1}, N_1 + \dots + N_i]$. Since $\bar{w}_j = w'_i/N_i$, $s'_j = \frac{d \cdot N_i}{m \cdot w'_i}$ with probability $\frac{m \cdot w'_i}{N_i \cdot d}$. So the contribution of all j 's corresponding to i is

$$\begin{aligned} & \sum_{j=N_1+\dots+N_{i-1}+1}^{N_1+\dots+N_i} s'_j \cdot |(a'_j)^\top \cdot \beta - y'_j|^p \\ &= \sum_{j=N_1+\dots+N_{i-1}+1}^{N_1+\dots+N_i} 1(s'_j > 0) \frac{d \cdot N_i}{m \cdot w'_i} \cdot |a_i^\top / N_i^{1/p} \cdot \beta - y_i / N_i^{1/p}|^p \\ &= \left(\sum_j 1(s'_j > 0) \right) \cdot \frac{d}{m \cdot w'_i} \cdot |a_i^\top \beta - y_i|^p. \end{aligned}$$

Next the random variable $\left(\sum_j 1(s'_j > 0)\right)$ generated by s'_j converges to a Poisson random variable with mean $N_i \cdot \frac{m \cdot w'_i}{N_i \cdot d} = \frac{m \cdot w'_i}{d}$ when $\varepsilon \rightarrow 0$ and $N_i \rightarrow +\infty$. So $s_i \sim \frac{d}{m \cdot w'_i} \cdot \text{Poisson}\left(\frac{m \cdot w'_i}{d}\right)$ generated in this theorem will converge to the coefficient $\left(\sum_j 1(s'_j > 0)\right) \cdot \frac{d}{m \cdot w'_i}$ in the above calculation. This implies that $\sum_{i=1}^n s_i \cdot |a_i^\top \beta - y_i|^p$ generated in Lemma 26 for A' with weights \bar{w} converges to $\tilde{L}(\beta) = \sum_j s'_j \cdot |(a'_j)^\top \cdot \beta - y'_j|^p$ generated in this theorem when $\varepsilon \rightarrow 0$.

From all discussion above, we have the equivalences of $L(\beta) = \|A\beta - y\|_p^p = \|A'\beta - y'\|_p^p$ and $\tilde{L}(\beta) = \sum_{i=1}^n s_i \cdot |a_i^\top \beta - y_i|^p = \sum_j s'_j \cdot |(a'_j)^\top \cdot \beta - y'_j|^p$. Then Lemma 26 (with the same γ and $\alpha = 1$) provides the guarantee of \tilde{L} for all β . \blacksquare

Next we prove Lemma 26. We discuss a few properties and ingredients that will be used in this proof besides Fact 11 and Claim 12 mentioned earlier.

The key property of β^* that will be used in this proof is that the partial derivative in every direction is 0, i.e., $\partial L(\beta^*) = \mathbf{0}$. Since β^* is assumed to be $\mathbf{0}$, this is equivalent to

$$\forall j \in [d], \sum_{i=1}^n p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} = 0. \quad (7)$$

It will be more convenient to write $L(\beta) - L(\beta^*)$ in Lemma 26 as $\sum_{i=1}^n (|a_i^\top \beta - y_i|^p - |y_i|^p)$ (given $\beta^* = \mathbf{0}$) and similarly for $\tilde{L}(\beta) - \tilde{L}(\beta^*)$. Then we will use the following claim to approximately bound $|a_i^\top \beta - y_i|^p - |y_i|^p$ in the comparison of L and \tilde{L} .

Claim 27 *For any real numbers a and b ,*

$$|a - b|^p - |a|^p + p \cdot |a|^{p-1} \cdot \text{sign}(a) \cdot b = O(|b|^p).$$

In particular, this implies that for any β , a_i , and y_i , we always have

$$|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta = O(|a_i^\top \beta|^p).$$

We will use the following claim about random events associated with s_1, \dots, s_n to finish the proof. Recall that s_i is generated according to the weight w'_i . So one property about w'_i that will be used in this proof is $w'_i \approx_\gamma w_i$ and $w_i \approx_\alpha d/n$, implying that $w'_i \approx_{\gamma\alpha} d/n$. The claim considers the concentration for the approximation in terms of Claim 27.

Claim 28 *Given any β , with probability at least $1 - \exp\left(-\Omega\left(\frac{\varepsilon^2 m}{\alpha^{O(1)} \cdot \gamma \cdot d}\right)\right)$ (over s_1, \dots, s_m),*

$$\begin{aligned} & \sum_{i=1}^n s_i \cdot \left(|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta \right) \\ &= \sum_{i=1}^n \left(|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta \right) \pm \varepsilon \|A\beta\|_p^p. \end{aligned}$$

Note that the extra term on the R.H.S., $\sum_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta$, is always zero for any β from (7). So the last piece in this proof is to bound its counterpart on the left hand side by Claim 12 mentioned earlier.

We defer the proof of Claim 27 to Section C.1, the proof of Claim 28 to Section C.2, and the proof of Claim 12 to Section C.3 separately.

Proof of Lemma 26. Recall that we assume $\beta^* = 0$ in this proof. Let $C_0 = \Theta(1/\epsilon\delta)$ with a sufficiently large constant, $\epsilon' = \epsilon \cdot \delta/50$ and $B_{\epsilon'}$ be an ϵ' -net (in ℓ_p norm) of the ℓ_p ball with radius $C_0\|y\|_p$ such that for any β with $\|A\beta\|_p \leq C_0\|y\|_p$, $\exists\beta' \in B_{\epsilon'}$ satisfies $\|A(\beta - \beta')\|_p \leq \epsilon'\|y\|_p$. We set $m = O\left(\frac{\alpha^{O(1)}\gamma \cdot d^2 \log C_0/\epsilon'}{\epsilon^2} + \frac{\alpha^{O(1)}\gamma \cdot d^{2/p}}{\epsilon^2\delta}\right)$ such that we have the following properties of s_1, \dots, s_n :

1. Claim 28 holds for all $\beta \in B_{\epsilon'}$.
2. For all β , we always have $\sum_i |a_i^\top \beta|^p \approx_{1+\epsilon} \sum_i s_i \cdot |a_i^\top \beta|^p$. In other words, the subspace embedding $\|SA\beta\|_p^p \approx_{1+\epsilon} \|A\beta\|_p^p$ for all β .
3. By Claim 12, with probability $1 - \delta$, the R.H.S. of Equation (5) is $\leq \epsilon \cdot \|A\beta\|_p \cdot \|y\|_p^{p-1}$, which is upper bounded $\epsilon(\|A\beta\|_p^p + \|y\|_p^p)$.
4. Moreover, since $\mathbb{E}_s[\tilde{L}(\beta^*)] = L(\beta^*)$, we assume $\tilde{L}(\beta^*) := \|Sy\|_p^p = \sum_{i=1}^n s_i |y_i|^p$ is $\leq \frac{1}{\delta}\|y\|_p^p$ w.p. $1 - \delta$ by the Markov's inequality.

First, we argue the conclusion $\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm O(\epsilon) \cdot (\|A\beta\|_p^p + \|y\|_p^p)$ holds for all $\beta \in B_{\epsilon'}$. As mentioned earlier, we rewrite $L(\beta) - L(\beta^*)$ as

$$\sum_{i=1}^n \left(|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta \right) \quad (8)$$

since the cross term is zero from (7). Then we rewrite $\tilde{L}(\beta) - \tilde{L}(\beta^*)$ as

$$\underbrace{\sum_{i=1}^n s_i \cdot \left(|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta \right)}_{T_1} - \underbrace{\sum_{i=1}^n s_i \cdot p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta}_{T_2}.$$

By Claim 28, the first term T_1 is $\epsilon \cdot \|A\beta\|_p^p$ -close to (8). Moreover, by Property 3 from Claim 12 mentioned above, the second term T_2 is always upper bounded by $\epsilon(\|y\|_p^p + \|A\beta\|_p^p)$.

Second, we argue $\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm O(\epsilon + \epsilon') \cdot (\|A\beta\|_p^p + \|y\|_p^p)$ holds for all β with $\|A\beta\|_p \leq C_0\|y\|_p$. Let us fix such a β and define β' to be the closest vector to it in $B_{\epsilon'}$ with $\|A\beta - A\beta'\|_p \leq \epsilon'\|y\|_p$. We rewrite

$$L(\beta) - L(\beta^*) = \|A\beta' + A(\beta - \beta') - y\|_p^p - L(\beta^*). \quad (9)$$

By triangle inequality, we bound

$$\|A\beta' + A(\beta - \beta') - y\|_p \in [\|A\beta' - y\|_p - \|A(\beta - \beta')\|_p, \|A\beta' - y\|_p + \|A(\beta - \beta')\|_p]$$

and use the approximation in Claim 27 with $a = \|A\beta' - y\|_p$ and $b = \|A(\beta - \beta')\|_p$ to approximate its p th power in (9) as

$$\begin{aligned}
 & L(\beta) - L(\beta^*) \\
 &= \|A\beta' - y\|_p^p - L(\beta^*) \pm O(p \cdot \|A\beta' - y\|_p^{p-1} \cdot \|A(\beta - \beta')\|_p + \|A(\beta - \beta')\|_p^p) \\
 &= \|A\beta' - y\|_p^p - L(\beta^*) \pm O(p \cdot \|A\beta' - y\|_p^{p-1} \cdot \varepsilon' \|y\|_p + \varepsilon'^p \cdot \|y\|_p^p). \\
 &= \|A\beta' - y\|_p^p - L(\beta^*) \pm \varepsilon' \cdot O(\|A\beta'\|_p^{p-1} \cdot \|y\|_p + \|y\|_p^p + \|y\|_p^p) \quad (\text{by triangle inequality}) \\
 &= L(\beta') - L(\beta^*) \pm \varepsilon' \cdot O(\|A\beta'\|_p^{p-1} \cdot \|y\|_p + \|y\|_p^p).
 \end{aligned}$$

Now we approximate $\tilde{L}(\beta) - \tilde{L}(\beta^*)$ by defining a weighted L_p norm $\|x\|_{s,p}$ for a vector $x \in \mathbb{R}^n$ as $(\sum_i s_i |x_i|^p)^{1/p}$. Note that $\|\cdot\|_{s,p}$ satisfies the triangle inequality. At the same time, $\tilde{L}(\beta) = \|A\beta - y\|_{s,p}$ by the definition and $\|A\beta\|_p^p \approx_{1+\varepsilon} \|A\beta\|_{s,p}^p$ by the 2nd property mentioned in this proof. By the same argument as above, we have

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = \tilde{L}(\beta') - \tilde{L}(\beta^*) \pm \varepsilon' \cdot O(\|A\beta'\|_{s,p}^{p-1} \cdot \|y\|_{s,p} + \|y\|_{s,p}^p).$$

Since β' has $\tilde{L}(\beta') - \tilde{L}(\beta^*) = L(\beta') - L(\beta^*) \pm O(\varepsilon) \cdot (\|A\beta'\|_p^p + \|y\|_p^p)$, from all the discussion above, the error between $\tilde{L}(\beta) - \tilde{L}(\beta^*)$ and $L(\beta) - L(\beta^*)$ is at most

$$\begin{aligned}
 & O(\varepsilon) \cdot (\|A\beta'\|_p^p + \|y\|_p^p) + \varepsilon' \cdot O(\|A\beta'\|_{s,p}^{p-1} \cdot \|y\|_{s,p} + \|y\|_{s,p}^p + \|A\beta'\|_p^{p-1} \cdot \|y\|_p + \|y\|_p^p) \\
 &= O(\varepsilon) \cdot (\|A\beta'\|_p^p + \|y\|_p^p) + \varepsilon'/\delta \cdot O(\|A\beta'\|_p^{p-1} \cdot \|y\|_p + \|y\|_p^p)
 \end{aligned}$$

where we replace $\|y\|_{s,p}^p = \|Sy\|_p^p$ by $\frac{1}{\delta}\|y\|_p^p$ and apply the subspace embedding Property 2 to $\|A\beta\|_{s,p}$. Given $\varepsilon' = O(\varepsilon\delta)$, the error is at most $\varepsilon \cdot O(\|A\beta\|_p^p + \|y\|_p^p)$

The last case of β is $\|A\beta\|_p > C_0\|y\|_p$, where we bound $\tilde{L}(\beta) - \tilde{L}(\beta^*) = L(\beta) - L(\beta^*) \pm O(\varepsilon) \cdot \|A\beta\|_p^p$. Note that $L(\beta) := \|A\beta - y\|_p^p$ is in $(\|A\beta\|_p \pm \|y\|_p)^p$ by the triangle inequality. From the approximation in Claim 27, this is about

$$\|A\beta\|_p^p \pm p\|A\beta\|_p^{p-1} \cdot \|y\|_p \pm O(\|y\|_p^p),$$

which bounds

$$L(\beta) - L(\beta^*) = \|A\beta\|_p^p \pm p\|A\beta\|_p^{p-1} \cdot \|y\|_p \pm O(\|y\|_p^p) = \|A\beta\|_p^p \pm O(1/C_0)\|A\beta\|_p^p.$$

Similarly, we have $\tilde{L}(\beta) - \tilde{L}(\beta^*) = \|A\beta\|_{s,p}^p \pm p\|A\beta\|_{s,p}^{p-1} \cdot \|y\|_{s,p} \pm O(\|y\|_{s,p}^p)$ for the weighted ℓ_p norm defined as above. Given $\|y\|_{s,p}^p \leq \frac{1}{\delta}\|y\|_p^p$ and $\|A\beta\|_{s,p}^p = (1 \pm \varepsilon)\|A\beta\|_p^p$, so

$$\tilde{L}(\beta) - \tilde{L}(\beta^*) = \|A\beta\|_p^p \pm \varepsilon\|A\beta\|_p^p \pm O\left(\frac{1}{\delta C_0}\right)\|A\beta\|_p^p.$$

When $C_0 = \Theta(1/\varepsilon\delta)$ for a sufficiently large constant, the error is $O(\varepsilon) \cdot \|A\beta\|_p^p$.

Finally we conclude that the error is always bounded by $\varepsilon \cdot L(\beta) := \varepsilon \cdot \|A\beta - y\|_p^p$.

1. When $\|A\beta\|_p \leq 2\|y\|_p$, $L(\beta) \geq \|y\|_p^p$ by the definition of β^* . So $\varepsilon \cdot O(\|y\|_p^p + \|A\beta\|_p^p) = O(\varepsilon) \cdot L(\beta)$.
2. Otherwise $\|A\beta\|_p > 2\|y\|_p$. Then $L(\beta) \geq (\|A\beta\|_p - \|y\|_p)^p \geq \|A\beta\|_p^p/4$, we also have $\varepsilon \cdot (\|y\|_p^p + \|A\beta\|_p^p) = O(\varepsilon) \cdot L(\beta)$.

■

C.1. Proof of Claim 27

Without loss of generality, we assume $a > 0$. We consider two cases:

1. $|b| < a/2$: We rewrite $|a-b|^p - |a|^p + p \cdot |a|^{p-1} \cdot b$ as $\int_{a-b}^a p \cdot a^{p-1} - px^{p-1} dx$. Now we bound $p \cdot a^{p-1} - px^{p-1}$ using Taylor's theorem, whose reminder is $f'(\zeta) \cdot (a-x)$ for $f(t) = p \cdot t^{p-1}$ and some $\zeta \in (x, a)$. So we upper bound the integration by

$$\int_{a-b}^a p \cdot a^{p-1} - px^{p-1} dx \leq \int_{a-b}^a p(p-1) \cdot |\zeta_x|^{p-2} \cdot |a-x| dx.$$

Since $p \in (1, 2]$, $x \in (a-b, a)$ and $\zeta_x \in (x, a)$, we always upper bound $|\zeta_x|^{p-2}$ by $\frac{1}{|a/2|^{2-p}}$.

So this integration is upper bounded by $\frac{p(p-1) \cdot b^2}{2|a/2|^{2-p}} = O(|b|^p)$ given $|b| < a/2$.

2. $|b| \geq a/2$: We upper bound

$$\begin{aligned} |a-b|^p - |a|^p + p \cdot |a|^{p-1} \cdot b &\leq (|a| + |b|)^p - |a|^p + p \cdot |a|^{p-1} \cdot |b| \\ &\leq (3|b|)^p + (2|b|)^p + p \cdot 2^{p-1} \cdot |b|^p = O(|b|^p). \end{aligned}$$

C.2. Proof of Claim 28

Since $s_i = \frac{d}{m \cdot w'_i}$ with probability $\frac{m \cdot w'_i}{d}$ (otherwise 0), $\mathbb{E}[s_i] = 1$ and

$$\mathbb{E} \sum_{i=1}^n s_i \left(|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \text{sign}(y_i) \cdot a_i^\top \beta \right) = \sum_{i=1}^n \left(|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \text{sign}(y_i) a_i^\top \beta \right).$$

To bound the deviation, we plan to use the following version of Bernstein's inequality for $X_i = (s_i - 1) \cdot (|a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta)$:

$$\Pr \left(\sum_i X_i \geq t \right) \leq \exp \left(- \frac{\frac{1}{2} t^2}{\sum_i \mathbb{E}[X_i^2] + \frac{1}{3} M t} \right).$$

First of all, we bound $M = \sup |X_i|$ as

$$\begin{aligned} &\frac{d}{m \cdot w'_i} \left| |a_i^\top \beta - y_i|^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta \right| \\ &\leq \frac{d}{m \cdot w'_i} \cdot C |a_i^\top \beta|^p && \text{(by Claim 27)} \\ &\leq \frac{d}{m} \cdot C \cdot \frac{|a_i^\top \beta|^p}{w'_i} \\ &\leq \frac{d}{m} \cdot C \gamma \cdot \frac{|a_i^\top \beta|^p}{w_i} && \text{(recall } w'_i \approx_\gamma w_i) \\ &\leq \frac{d}{m} \cdot C \gamma \cdot \|A\beta\|_p^p && \text{(by Property (4))} \end{aligned}$$

Next we bound $\sum_i \mathbb{E}[X_i^2]$ as

$$\begin{aligned}
 & \sum_{i=1}^n \mathbb{E} \left[(s_i - 1)^2 \cdot (a_i^\top \beta - y_i)^p - |y_i|^p + p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot a_i^\top \beta \right]^2 \\
 & \leq \sum_{i=1}^n \mathbb{E}[(s_i - 1)^2] \cdot C^2 \cdot |a_i^\top \beta|^{2p} && \text{(by Claim 27)} \\
 & \leq \sum_{i=1}^n \left(1 + \frac{d}{m \cdot w'_i}\right) \cdot C^2 \cdot \left(\frac{\alpha \cdot d}{n} \|A\beta\|_p^p\right) \cdot |a_i^\top \beta|^p && \text{(by Property (4))} \\
 & \leq \frac{2d \cdot \alpha \gamma \cdot n}{m \cdot d} \cdot C^2 \left(\frac{\alpha \cdot d}{n} \|A\beta\|_p^p\right) \cdot \sum_{i=1}^n |a_i^\top \beta|^p && \text{(by } w'_i \approx_{\alpha\gamma} d/n) \\
 & = \frac{2C^2 \cdot \alpha^2 \gamma \cdot d}{m} \|A\beta\|_p^{2p}.
 \end{aligned}$$

So for $m = \Omega(\alpha^2 \gamma d / \varepsilon^2)$ and $t = \varepsilon \|A\beta\|_p^p$, we have the deviation is at most t with probability $1 - \exp(-\Omega(\frac{\varepsilon^2 m}{\alpha^2 \gamma \cdot d}))$.

C.3. Proof of Claim 12

By Cauchy-Schwartz, we upper bound

$$\begin{aligned}
 & \sum_{i=1}^n s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot \sum_j A_{i,j} \beta_j \\
 & = \sum_j \beta_j \cdot \left(\sum_i s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right) \\
 & \leq \left\| \left(\sum_i s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right)_{j \in [d]} \right\|_2 \cdot \|\beta\|_2.
 \end{aligned}$$

Now we bound the ℓ_2 norm of those two vectors separately.

Fact 29 $\mathbb{E} \left\| \left(\sum_i s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right)_{j \in [d]} \right\|_2^2 \leq \frac{2\alpha d}{m} \cdot \sum_i |y_i|^{2p-2}.$

Proof We rewrite the ℓ_2 norm square as

$$\begin{aligned}
 & \mathbb{E} \left\| \left(\sum_i s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right)_{j \in [d]} \right\|_2^2 - \sum_j \mathbb{E} \left[\sum_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right]^2 \\
 & \hspace{15em} \text{(the make up term is always zero by Property (7))} \\
 &= \sum_j \mathbb{E} \left[\left(\sum_i (s_i - 1) p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right)^2 \right] \\
 &= \sum_j \sum_i \mathbb{E} \left((s_i - 1) p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right)^2 \\
 & \hspace{15em} \text{(use } \mathbb{E}[s_i] = 1 \text{ and the independence between different } s_i \text{'s)} \\
 &\leq \sum_j \sum_i \left(\frac{d}{m \cdot w'_i} + 1 \right) \cdot |y_i|^{2p-2} \cdot A_{i,j}^2 \\
 &\leq \sum_i |y_i|^{2p-2} \cdot \left(\frac{d}{m \cdot w'_i} + 1 \right) \cdot \sum_j A_{i,j}^2 \\
 & \hspace{15em} \text{(use the definition of leverage score in (3) and } w'_i \approx_{\alpha\gamma} d/n) \\
 &\leq \sum_i |y_i|^{2p-2} \cdot \frac{2\alpha\gamma n}{m} \cdot \frac{\alpha^{C_p} d}{n} \\
 &= \frac{2\alpha^{1+C_p}\gamma d}{m} \cdot \sum_i |y_i|^{2p-2}.
 \end{aligned}$$

■

Fact 30 Suppose $A^\top A = I_d$ and its leverage score is almost uniform: $\|a_i\|_2^2 \approx_\alpha d/n$. Then $\|\beta\|_2 \leq \|A\beta\|_p \cdot \left(\frac{\alpha^{C_p} \cdot d}{n} \right)^{\frac{2-p}{2p}}$.

Proof Since $A^\top A = I_d$, $\|\beta\|_2 = \|A\beta\|_2$. Then we upper bound $\|A\beta\|_2^2$:

$$\begin{aligned}
 \|A\beta\|_2^2 &= \sum_{i=1}^n |a_i^\top \beta|^2 \\
 &\leq \left(\sum_i |a_i^\top \beta|^p \right) \cdot \max_i |a_i^\top \beta|^{2-p} \\
 &\leq \|A\beta\|_p^p \cdot \left(\frac{\alpha \cdot d}{n} \cdot \|A\beta\|_p \right)^{\frac{2-p}{p}} \hspace{10em} \text{(by Property (4))} \\
 &\leq \|A\beta\|_p^2 \cdot \left(\frac{\alpha \cdot d}{n} \right)^{\frac{2-p}{p}}.
 \end{aligned}$$

■

Now we are ready to finish the proof. From the 1st fact, with probability $1 - \delta$,

$$\left\| \left(\sum_i s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right)_{j \in [d]} \right\|_2 \leq 10 \sqrt{\frac{2\alpha^{O(1)}\gamma d}{\delta m}} \cdot \|y\|_{2p-2}.$$

From the 2nd fact, we always have

$$\|\beta\|_2 \leq \|A\beta\|_p \cdot \left(\frac{\alpha \cdot d}{n}\right)^{\frac{2-p}{2p}}.$$

So with probability $1 - \delta$,

$$\begin{aligned} \sum_{i=1}^n s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot \sum_j A_{i,j} \beta_j &\leq \left\| \left(\sum_i s_i p \cdot |y_i|^{p-1} \cdot \text{sign}(y_i) \cdot A_{i,j} \right)_{j \in [d]} \right\|_2 \cdot \|\beta\|_2 \\ &\leq 20 \sqrt{\frac{\alpha^{O(1)} \gamma d}{\delta m}} \cdot \|y\|_{2p-2}^{p-1} \cdot \|A\beta\|_p \cdot \left(\frac{\alpha \cdot d}{n}\right)^{\frac{2-p}{2p}} \\ &\leq C \sqrt{\frac{\alpha^{O(1)} \gamma d}{\delta m}} \cdot \|A\beta\|_p \cdot (\alpha \cdot d)^{\frac{2-p}{2p}} \cdot \|y\|_{2p-2}^{p-1} \cdot (1/n)^{\frac{2-p}{2p}}. \end{aligned}$$

Finally we use Holder's inequality to bound the last two terms $\|y\|_{2p-2}^{p-1} \cdot (1/n)^{\frac{2-p}{2p}}$ in the above calculation. We set $q_1 = \frac{p}{2-p}$ and $q_2 = \frac{p}{2p-2}$ (such that $1 = 1/q_1 + 1/q_2$) to obtain

$$\left(\sum_i |y_i|^{2p-2} \right) \leq \left(\sum_i 1 \right)^{1/q_1} \cdot \left(\sum_i |y_i|^{(2p-2) \cdot q_2} \right)^{1/q_2} = n^{\frac{2-p}{p}} \cdot \|y\|_p^{2p-2}.$$

So we further simplify the above calculation

$$C \sqrt{\frac{\alpha^{O(1)} \gamma d}{\delta m}} \cdot \|A\beta\|_p \cdot (\alpha \cdot d)^{\frac{2-p}{2p}} \cdot \|y\|_{2p-2}^{p-1} \cdot (1/n)^{\frac{2-p}{2p}} \leq C \sqrt{\frac{\alpha^{O(1)} \cdot \gamma \cdot d^{2/p}}{\delta m}} \cdot \|A\beta\|_p \cdot \|y\|_p^{p-1}.$$

Appendix D. Additional Proofs from Section 5

We finish the proof of Theorem 13 in this section.

Proof of Theorem 13. For contradiction, we assume there is an algorithm P that outputs an $(1 + \varepsilon/200)$ -approximation with probability $1 - \delta$ using m queries on y . We first demonstrate $m = \Omega(\frac{\log 1/\delta}{\varepsilon^2})$. We pick $\alpha \in \{\pm 1\}$ and generate y as follows:

1. $(y_1, \dots, y_{n/d})$ are generated from D_α defined in Section 5.
2. The remaining entries are 0, i.e. $y_i = 0$ for $i > n/d$.

So $\beta^* = (\alpha, 0, \dots, 0)$ with probability $1 - \delta/d$. For any $(1 + \varepsilon/200)$ -approximation $\tilde{\beta}$, its first entry has $\text{sign}(\tilde{\beta}_1) = \alpha$ from Claim 14. By the lower bound in Lemma 15, the algorithm must make $\Omega(\frac{\log 1/\delta}{\varepsilon^2})$ queries to $y_1, \dots, y_{n/d}$.

Then we show $m = \Omega(d/\varepsilon^2)$. For convenience, the rest of the proof considers a fixed $\delta = 0.01$. By Yao's minmax principle, we consider a deterministic algorithm P in this part. Now let us define the distribution D_b over $\{\pm 1\}^n$ for $b \sim \{\pm 1\}^d$ as follows:

1. We sample $b \sim \{\pm 1\}^d$.
2. We generate $y \sim D_b$ where $D_b = D_{b_1} \circ D_{b_2} \circ \dots \circ D_{b_d}$.

For any b , when $n > 100d \log d / \varepsilon^2$ and $n' = n/d$, with probability 0.99, for each $i \in [d]$, D_{b_i} will generate n' bits where the number of bits equaling b_i is in the range of $[(1/2 + \varepsilon/2)n', (1/2 + 3\varepsilon/2)n']$. We assume this condition in the rest of this proof. From Claim 14, β^* minimizing $\|A\beta - y\|_1$ will have $b_i = \text{sign}(\beta_i^*)$ for every $i \in [d]$.

Next, given A and y , let $\tilde{\beta}$ be the output of P . We define b' as $b'_i = \text{sign}(\tilde{\beta}_i)$ for each coordinate i . We show that b' will agree with b (the string used to generate y) on 0.99 fraction of bits when P outputs an $(1 + \varepsilon/200)$ -approximation. As discussed before, $\|A\tilde{\beta} - y\|_1$ is the summation of d subproblems for d coordinates separately, i.e., $L_1(\tilde{\beta}_1), \dots, L_d(\tilde{\beta}_d)$. In particular, $\|A\tilde{\beta} - y\|_1 \leq (1 + \varepsilon/200)\|A\beta^* - y\|_1$ implies

$$\sum_{i=1}^d L_i(\tilde{\beta}_i) \leq (1 + \varepsilon/200) \sum_{i=1}^d L_i(\beta_i^*).$$

At the same time, we know $L_i(\tilde{\beta}_i) \geq L_i(\beta_i^*)$ and $L_i(\beta_i^*) \in [(1 - \varepsilon)n', (1 - 3\varepsilon)n']$ for any i . This implies that for at least 0.99 fraction of $i \in [d]$, $L_i(\tilde{\beta}_i) \leq (1 + \varepsilon)L_i(\beta_i^*)$: Otherwise the approximation ratio is not $(1 + \varepsilon/200)$ given

$$0.99 \cdot (1 - \varepsilon) + 0.01 \cdot (1 - 3\varepsilon) \cdot (1 + \varepsilon) > (1 + \varepsilon/200) \cdot (0.99 \cdot (1 - \varepsilon) + 0.01(1 - 3\varepsilon)). \quad (10)$$

From Claim 14, for such an i , $\tilde{\beta}_i$ will have the same sign with β_i^* . So b' agree with b on at least 0.99 fraction of coordinates.

For each b , let $m_i(b)$ denote the expected queries of P on $y_{(i-1)(n/d)+1}, \dots, y_{i(n/d)}$ (over the randomness of $y \sim D_b$). Since P makes at most m queries, we have $\mathbb{E}_y[\sum_i m_i(b)] \leq m$.

Now for each $b \in \{\pm 1\}^d$ and coordinate $i \in [d]$, we say that the coordinate i is good in b if (1) $\mathbb{E}_y[m_i(b)] \leq 60m/d$; and (2) $b'_i = b_i$ with probability 0.8 over $y \sim D_b$ when $b' = \text{sign}(\tilde{\beta})$ defined from the output $\tilde{\beta}$ of P .

Let $b^{(i)}$ denote the flip of b on coordinate i . We plan to show the existence of b and a coordinate $i \in [d]$ such that i is good in both b and $b^{(i)}$: Let us consider the graph G on the Boolean cube that corresponds to $b \in \{\pm 1\}^d$ and has edges of $(b, b^{(i)})$ for every b and i . To show the existence, for each edge $(b, b^{(i)})$ in the Boolean cube, we remove it if i is not good in b or $b^{(i)}$. We prove that after removing all bad events by a union bound, G still has edges inside.

For the 1st condition (1) $\mathbb{E}[m_i(b)] \leq 60m/d$, we will remove at most $d/60$ edges from each vertex b . So the fraction of edges removed by this condition is at most $1/30$. For the 2nd condition, from the guarantee of P , we know

$$\Pr_{b,y}[\tilde{\beta} \text{ is an } (1 + \varepsilon/200) \text{ approximation}] \geq 0.99.$$

So for 0.8 fraction of points b , we have $\Pr_{y \sim D_b}[\tilde{\beta} \text{ is an } (1 + \varepsilon/200) \text{ approximation}] \geq 0.95$ (Otherwise we get a contradiction since $0.8 + 0.2 \cdot 0.95 = 0.99$). Thus, when $y \sim D_b$ for such a string b , w.p. 0.95, b' will agree with b on at least 0.99 fraction of coordinates from the above discussion. In another word, for such a string b ,

$$\mathbb{E}_{y \sim D_b} \left[\sum_i 1(b_i = b'_i) \right] \geq 0.95 \cdot 0.99n.$$

This indicates that there are at least 0.65 fraction of coordinates in $[d]$ that satisfy $\Pr_y[b_i = b'_i] \geq 0.8$ (otherwise we get a contradiction since $0.65 + 0.35 \cdot 0.8 = 0.93$). Hence, such a string b will have 0.65 fraction of *good* coordinates.

Back to the counting of bad edges removed by the 2nd condition, we will remove at most $2 \cdot (0.2 + 0.8 \cdot 0.35) = 0.96$ fraction of edges. Because $1/30 + 0.96 < 1$, we know the existence of b and $b^{(i)}$ such that i is a good coordinate in b and $b^{(i)}$.

Now we use $b, b^{(i)}$ and P to construct an algorithm for Bob to win the game in Lemma 15. From the definition, P makes $60m/d$ queries in expectation on entries $y_{(i-1)(n/d)+1}, \dots, y_{i(n/d)}$ and outputs $b'_i = b_i$ with probability 0.8. Since halting P at $20 \cdot 60m/d$ queries on entries $y_{(i-1)(n/d)+1}, \dots, y_{i(n/d)}$ will reduce the success probability to $0.8 - 0.05 = 0.75$. We assume P makes at most $1200m/d$ queries with success probability 0.75. Now we describe Bob's strategy to win the game:

1. Randomly sample $b_j \sim \{\pm 1\}$ for $j \neq i$.
2. Simulate the algorithm P : each time when P asks the label y_ℓ for $\ell \in [(j-1)n/d+1, jn/d]$.
 - (a) If $j = i$, Bob queries the corresponding label from Alice.
 - (b) Otherwise, Bob generates the label using D_{b_j} by himself.
3. Finally, we use P to produce $b' \in \{\pm 1\}^n$ and outputs b'_i .

Bob wins this game with probability at least

$$\Pr_a[a = -1] \cdot \Pr_y[b'_i = -1] + \Pr_a[a = 1] \cdot \Pr_y[b'_i = 1] \geq 0.75$$

from the properties of i in b and $b^{(i)}$. So we know $1200m/d = \Omega(1/\varepsilon^2)$, which lower bounds $m = \Omega(d/\varepsilon^2)$. ■