
Efficient Symmetric Norm Regression via Linear Sketching*

Zhao Song

University of Washington
magic.linuxkde@gmail.com

Ruosong Wang

Carnegie Mellon University
ruosongw@andrew.cmu.edu

Lin F. Yang

University of California, Los Angeles
linyang@ee.ucla.edu

Hongyang Zhang

Toyota Technological Institute at Chicago
hongyanz@ttic.edu

Peilin Zhong

Columbia University
pz2225@columbia.edu

Abstract

We provide efficient algorithms for overconstrained linear regression problems with size $n \times d$ when the loss function is a symmetric norm (a norm invariant under sign-flips and coordinate-permutations). An important class of symmetric norms are Orlicz norms, where for a function G and a vector $y \in \mathbb{R}^n$, the corresponding Orlicz norm $\|y\|_G$ is defined as the unique value α such that $\sum_{i=1}^n G(|y_i|/\alpha) = 1$. When the loss function is an Orlicz norm, our algorithm produces a $(1 + \varepsilon)$ -approximate solution for an arbitrarily small constant $\varepsilon > 0$ in input-sparsity time, improving over the previously best-known algorithm which produces a $d \cdot \text{polylog } n$ -approximate solution. When the loss function is a general symmetric norm, our algorithm produces a $\sqrt{d} \cdot \text{polylog } n \cdot \text{mmc}(\ell)$ -approximate solution in input-sparsity time, where $\text{mmc}(\ell)$ is a quantity related to the symmetric norm under consideration. To the best of our knowledge, this is the first input-sparsity time algorithm with provable guarantees for the general class of symmetric norm regression problem. Our results shed light on resolving the universal sketching problem for linear regression, and the techniques might be of independent interest to numerical linear algebra problems more broadly.

1 Introduction

Linear regression is a fundamental problem in machine learning. For a data matrix $A \in \mathbb{R}^{n \times d}$ and a response vector $b \in \mathbb{R}^n$ with $n \gg d$, the overconstrained linear regression problem can be formulated as solving the following optimization problem:

$$\min_{x \in \mathbb{R}^d} \mathcal{L}(Ax - b), \quad (1)$$

where $\mathcal{L} : \mathbb{R}^n \rightarrow \mathbb{R}$ is a loss function. Via the technique of linear sketching, we have witnessed many remarkable speedups for linear regression for a wide range of loss functions. Such technique involves designing a sketching matrix $S \in \mathbb{R}^{r \times n}$, and showing that by solving a linear regression instance on the data matrix SA and the response vector Sb , which is usually much smaller in size, one can obtain

*All authors contribute equally.

Table 1: M -estimators

HUBER	$\begin{cases} x^2/2 & x \leq c \\ c(x - c/2) & x > c \end{cases}$
$\ell_1 - \ell_2$	$2(\sqrt{1 + x^2/2} - 1)$
“FAIR”	$c^2 (x /c - \log(1 + x /c))$

an approximate solution to the original linear regression instance in (1). Sarlós showed in [29] that by taking S as a Fast Johnson-Lindenstrauss Transform matrix [1], one can obtain $(1 + \varepsilon)$ -approximate solutions to the least square regression problem ($\mathcal{L}(y) = \|y\|_2^2$) in $O(nd \log n + \text{poly}(d/\varepsilon))$ time. The running time was later improved to $O(\text{nnz}(A) + \text{poly}(d/\varepsilon))$ [12, 26, 28, 23, 15]. Here $\text{nnz}(A)$ is the number of non-zero entries in the data matrix A , which could be much smaller than nd for sparse data matrices. This technique was later generalized to other loss functions. By now, we have $\tilde{O}(\text{nnz}(A) + \text{poly}(d/\varepsilon))$ time algorithms for ℓ_p norms ($\mathcal{L}(y) = \|y\|_p^p$) [18, 26, 35, 16, 32], the quantile loss function [36], M -estimators [14, 13] and the Tukey loss function [11].

Despite we have successfully applied the technique of linear sketching to many different loss functions, ideally, it would be more desirable to design algorithms that work for a wide range of loss functions, instead of designing a new sketching algorithm for every specific loss function. Naturally, this leads to the following problem, which is the linear regression version of the universal sketching problem² studied in streaming algorithms [10, 9]. We note that similar problems are also asked and studied for various algorithmic tasks, including principal component analysis [31], sampling [21], approximate nearest neighbor search [4, 3], discrepancy [17, 8], sparse recovery [27] and mean estimation with statistical queries [19, 22].

Problem 1. *Is it possible to design sketching algorithms for linear regression, that work for a wide range of loss functions?*

Prior to our work, [14, 13] studied this problem in terms of M -estimators, where the loss function employs the form $\mathcal{L}(y) = \sum_{i=1}^n G(y_i)$ for some function G . See Table 1 for a list of M -estimators. However, much less is known for the case where the loss function $\mathcal{L}(\cdot)$ is a norm, except for ℓ_p norms. Recently, Andoni et al. [2] tackle Problem 1 for Orlicz norms, which can be seen as a scale-invariant version of M -estimators. For a function G and a vector $y \in \mathbb{R}^n$ with $y \neq 0$, the corresponding Orlicz norm $\|y\|_G$ is defined as the unique value α such that

$$\sum_{i=1}^n G(|y_i|/\alpha) = 1. \quad (2)$$

When $y = 0$, we define $\|y\|_G$ to be 0. Note that Orlicz norms include ℓ_p norms as special cases, by taking $G(z) = |z|^p$ for some $p \geq 1$. Under certain assumptions on the function G , [2] obtains the first input-sparsity time algorithm for solving Orlicz norm regression. More precisely, in $\tilde{O}(\text{nnz}(A) + \text{poly}(d \log n))$ time, their algorithm obtains a solution $\hat{x} \in \mathbb{R}^d$ such that $\|A\hat{x} - b\|_G \leq d \cdot \text{polylog } n \cdot \min_{x \in \mathbb{R}^d} \|Ax - b\|_G$.

There are two natural problems left open by the work of [2]. First, the algorithm in [2] has approximation ratio as large as $d \cdot \text{polylog } n$. Although this result is interesting from a theoretical point of view, such a large approximation ratio is prohibitive for machine learning applications in practice. Is it possible to obtain an algorithm that runs in $\tilde{O}(\text{nnz}(A) + \text{poly}(d/\varepsilon))$ time, with approximation ratio $1 + \varepsilon$, for arbitrarily small ε , similar to the case of ℓ_p norms? Moreover, although Orlicz norm includes a wide range of norms, many other important norms, e.g., top- k norms (the sum of absolute values of the leading k coordinates of a vector), max-mix of ℓ_p norms (e.g. $\max\{\|x\|_2, c\|x\|_1\}$ for some $c > 0$), and sum-mix of ℓ_p norms (e.g. $\|x\|_2 + c\|x\|_1$ for some $c > 0$), are not Orlicz norms. More complicated examples include the k -support norm [5] and the box-norm [25], which have found applications in sparse recovery. In light of Problem 1, it is natural to ask whether it is possible to apply the technique of linear sketching to a broader class of norms. In this paper, we obtain affirmative answers to both problems, and make progress towards finally resolving Problem 1.

Notations. We use $\tilde{O}(f)$ to denote $f \text{ polylog } f$. For a matrix $A \in \mathbb{R}^{n \times d}$, we use $A_i \in \mathbb{R}^d$ to denote its i -th row, viewed as a column vector. For n real numbers x_1, x_2, \dots, x_n , we define

²https://sublinear.info/index.php?title=Open_Problems:30.

$\text{diag}(x_1, x_2, \dots, x_n) \in \mathbb{R}^{n \times n}$ to be the diagonal matrix where the i -th diagonal entry is x_i . For a vector $x \in \mathbb{R}^n$ and $p \geq 1$, we use $\|x\|_p$ to denote its ℓ_p norm, and $\|x\|_0$ to denote its ℓ_0 norm, i.e., the number of non-zero entries in x . For two vectors $x, y \in \mathbb{R}^n$, we use $\langle x, y \rangle$ to denote their inner product. For any $n > 0$, we use $[n]$ to denote the set $\{1, 2, \dots, n\}$. For $0 \leq p \leq 1$, we define $\text{Ber}(p)$ to be the Bernoulli distribution with parameter p . We use \mathbb{S}^{n-1} to denote the unit ℓ_2 sphere in \mathbb{R}^n , i.e., $\mathbb{S}^{n-1} = \{x \in \mathbb{R}^n \mid \|x\|_2 = 1\}$. We use $\mathbb{R}_{\geq 0}$ to denote the set of all non-negative real numbers, i.e., $\mathbb{R}_{\geq 0} = \{x \in \mathbb{R} \mid x \geq 0\}$.

1.1 Our Contributions

Algorithm for Orlicz Norms. Our first contribution is a unified algorithm which produces $(1 + \varepsilon)$ -approximate solutions to the linear regression problem in (1), when the loss function $\mathcal{L}(\cdot)$ is an Orlicz norm. Before introducing our results, we first give our assumptions on the function G which appeared in (2).

Assumption 1. We assume the function $G : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ satisfies the following properties:

1. G is a strictly increasing convex function on $[0, \infty)$;
2. $G(0) = 0$, and for all $x \in \mathbb{R}$, $G(x) = G(-x)$;
3. There exists some $C_G > 0$, such that for all $0 < x < y$, $G(y)/G(x) \leq C_G(y/x)^2$.

The first two conditions in Assumption 1 are necessary to make sure the corresponding Orlicz norm $\|\cdot\|_G$ is indeed a norm, and the third condition requires the function G to have at most quadratic growth, which can be satisfied by all M -estimators in Table 1 and is also required by prior work [2]. Notice that our assumptions are weaker than those in [2]. In [2], it is further required that $G(x)$ is a linear function when $x > 1$, and G is twice differentiable on an interval $(0, \delta_G)$ for some $\delta_G > 0$. Given our assumptions on G , our main theorem is summarized as follows.

Theorem 1. For a function G that satisfies Assumption 1, there exists an algorithm that, on any input $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^n$, finds a vector x^* in time $\tilde{O}(\text{nnz}(A) + \text{poly}(d/\varepsilon))$, such that with probability at least 0.9, $\|Ax^* - b\|_G \leq (1 + \varepsilon) \min_{x \in \mathbb{R}^d} \|Ax - b\|_G$.

To the best of our knowledge, this is the first input-sparsity time algorithm with $(1 + \varepsilon)$ -approximation guarantee, that goes beyond ℓ_p norms, the quantile loss function, and M -estimators. See Table 2 for a more comprehensive comparison with previous results.

Algorithm for Symmetric Norms. We further study the case when the loss function $\mathcal{L}(\cdot)$ is a symmetric norm. Symmetric norm is a more general class of norms, which includes all norms that are invariant under sign-flips and coordinate-permutations. Formally, we define symmetric norms as follow.

Definition 1. A norm $\|\cdot\|_\ell$ is called a *symmetric norm*, if $\|(y_1, y_2, \dots, y_n)\|_\ell = \|(s_1 y_{\sigma_1}, s_2 y_{\sigma_2}, \dots, s_n y_{\sigma_n})\|_\ell$ for any permutation σ and any assignment of $s_i \in \{-1, 1\}$.

Symmetric norm includes ℓ_p norms and Orlicz norms as special cases. It also includes all examples provided in the introduction, i.e., top- k norms, max-mix of ℓ_p norms, sum-mix of ℓ_p norms, the k -support norm [5] and the box-norm [25], as special cases. Understanding this general set of loss functions can be seen as a preliminary step to resolve Problem 1. Our main result for symmetric norm regression is summarized in the following theorem.

Theorem 2. Given a symmetric norm $\|\cdot\|_\ell$, there exists an algorithm that, on any input $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^n$, finds a vector x^* in time $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$, such that with probability at least 0.9, $\|Ax^* - b\|_\ell \leq \sqrt{d} \cdot \text{polylog } n \cdot \text{mmc}(\ell) \cdot \min_{x \in \mathbb{R}^d} \|Ax - b\|_\ell$.

In the above theorem, $\text{mmc}(\ell)$ is a characteristic of the symmetric norm $\|\cdot\|_\ell$, which has been proven to be essential in streaming algorithms for symmetric norms [7]. See Definition 7 for the formal definition of $\text{mmc}(\ell)$, and Section 3 for more details about $\text{mmc}(\ell)$. In particular, for ℓ_p norms with $p \leq 2$, top- k norms with $k \geq n / \text{polylog } n$, max-mix of ℓ_2 norm and ℓ_1 norm ($\max\{\|x\|_2, c\|x\|_1\}$ for some $c > 0$), sum-mix of ℓ_2 norm and ℓ_1 norm ($\|x\|_2 + c\|x\|_1$ for some $c > 0$), the k -support norm, and the box-norm, $\text{mmc}(\ell)$ can all be upper bounded by $\text{polylog } n$, which implies our algorithm has approximation ratio $\sqrt{d} \cdot \text{polylog } n$ for all these norms. This clearly demonstrates the generality of our algorithm.

Table 2: Comparison among input-sparsity time linear regression algorithms

Reference	Loss Function	Approximation Ratio
[18, 26, 35, 16, 32]	ℓ_p norms	$1 + \varepsilon$
[36]	Quantile loss function	$1 + \varepsilon$
[14, 13]	M -estimators	$1 + \varepsilon$
[2]	Orlicz norms	$d \cdot \text{polylog } n$
Theorem 1	Orlicz norms	$1 + \varepsilon$
Theorem 2	Symmetric norms	$\sqrt{d} \cdot \text{polylog } n \cdot \text{mmc}(\ell)$

Empirical Evaluation. In Section E of the supplementary material, we test our algorithms on real datasets. Our empirical results quite clearly demonstrate the practicality of our methods.

1.2 Technical Overview

Similar to previous works on using linear sketching to speed up solving linear regression, our core technique is to provide efficient dimensionality reduction methods for Orlicz norms and general symmetric norms. In this section, we discuss the techniques behind our results.

Row Sampling Algorithm for Orlicz Norms. Compared to prior work on Orlicz norm regression [2] which is based on random projection³, our new algorithm is based on row sampling. For a given matrix $A \in \mathbb{R}^{n \times d}$, our goal is to output a *sparse* weight vector $w \in \mathbb{R}^n$ with at most $\text{poly}(d \log n / \varepsilon)$ non-zero entries, such that with high probability, for all $x \in \mathbb{R}^d$,

$$(1 - \varepsilon) \|Ax - b\|_G \leq \|Ax - b\|_{G,w} \leq (1 + \varepsilon) \|Ax - b\|_G. \quad (3)$$

Here, for a weight vector $w \in \mathbb{R}^n$ and a vector $y \in \mathbb{R}^n$, the *weighted Orlicz norm* $\|y\|_{G,w}$ is defined as the unique value α such that $\sum_{i=1}^n w_i G(|y_i|/\alpha) = 1$. See Definition 4 for the formal definition of weighted Orlicz norm. To obtain a $(1 + \varepsilon)$ -approximate solution to Orlicz norm regression, by (3), it suffices to solve

$$\min_{x \in \mathbb{R}^d} \|Ax - b\|_{G,w}. \quad (4)$$

Since the vector $w \in \mathbb{R}^n$ has at most $\text{poly}(d \log n / \varepsilon)$ non-zero entries, and we can ignore all rows of A with zero weights, there are at most $\text{poly}(d \log n / \varepsilon)$ remaining rows in A in the optimization problem in (4). Furthermore, as we show in Lemma 3, $\|\cdot\|_{G,w}$ is a seminorm, which implies we can solve the optimization problem in (4) in $\text{poly}(d \log n / \varepsilon)$ time, by simply solving a convex program with size $\text{poly}(d \log n / \varepsilon)$. Thus, we focus on how to obtain the weight vector $w \in \mathbb{R}^n$ in the remaining part. Furthermore, by taking \bar{A} to be a matrix whose first d columns are A and last column is b , to satisfy (3), it suffices to find a weight vector w such that for all $x \in \mathbb{R}^{d+1}$,

$$(1 - \varepsilon) \|\bar{A}x\|_G \leq \|\bar{A}x\|_{G,w} \leq (1 + \varepsilon) \|\bar{A}x\|_G. \quad (5)$$

Hence, we ignore the response vector b in the remaining part of the discussion.

We obtain the weight vector w via importance sampling. We compute a set of sampling probabilities $\{p_i\}_{i=1}^n$ for each row of the data matrix A , and sample the rows of A according to these probabilities. The i -th entry of the weight vector w is then set to be $w_i = 1/p_i$ with probability p_i and $w_i = 0$ with probability $1 - p_i$. However, unlike ℓ_p norms, Orlicz norms are not “entry-wise” norms, and it is not even clear that such a sampling process gives an unbiased estimation. Our key insight here is that for a vector Ax with unit Orlicz norm, if for all $x \in \mathbb{R}^d$,

$$(1 - \varepsilon) \sum_{i=1}^n G((Ax)_i) \leq \sum_{i=1}^n w_i G((Ax)_i) \leq (1 + \varepsilon) \sum_{i=1}^n G((Ax)_i), \quad (6)$$

then (5) holds, which follows from the convexity of the function G . See Lemma 7 and its proof for more details. Therefore, it remains to develop a way to define and calculate $\{p_i\}_{i=1}^n$, such that the total number of sampled rows is small.

³Even for ℓ_p norms with $p < 2$, embeddings based on random projections will necessarily induce a distortion factor polynomial in d , as shown in [32].

Our method for defining and computing sampling probabilities p_i is inspired by row sampling algorithms for ℓ_p norms [18]. Here, the key is to obtain an upper bound on the contribution of each entry to the summation $\sum_{i=1}^n G((Ax)_i)$. Indeed, suppose for some vector $u \in \mathbb{R}^n$ such that $G(Ax)_i \leq u_i$ for all $x \in \mathbb{R}^d$ with $\|Ax\|_G = 1$, we can then sample each row of A with sampling probability proportional to u_i . Now, by standard concentration inequalities and a net argument, (6) holds with high probability. It remains to upper bound the total number of sampled rows, which is proportional to $\sum_{i=1}^n u_i$.

We use the case of ℓ_2 norm, i.e., $G(x) = x^2$, as an example to illustrate our main ideas for choosing the vector $u \in \mathbb{R}^n$. Suppose $U \in \mathbb{R}^{n \times d}$ is an orthonormal basis matrix of the column space of A , then the *leverage score*⁴ is defined to be the squared ℓ_2 norm of each row of U . Indeed, leverage score gives an upper bound on the contribution of each row to $\|Ux\|_2^2$, since by Cauchy-Schwarz inequality, for each row U_i of U , we have $\langle U_i, x \rangle^2 \leq \|U_i\|_2^2 \|x\|_2^2 = \|U_i\|_2^2 \|Ux\|_2^2$, and thus we can set $u_i = \|U_i\|_2^2$. It is also clear that $\sum_{i=1}^n u_i = d$.

For general Orlicz norms, leverage scores are no longer upper bounds on $G((Ux)_i)$. Inspired by the role of orthonormal bases in the case of ℓ_2 norm, we first define well-conditioned basis for general Orlicz norms as follow.

Definition 2. Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. We say $U \in \mathbb{R}^{n \times d}$ is a well-conditioned basis with condition number $\kappa_G = \kappa_G(U)$ if for all $x \in \mathbb{R}^d$, $\|x\|_2 \leq \|Ux\|_G \leq \kappa_G \|x\|_2$.

Given this definition, when $\|Ux\|_G = 1$, by Cauchy-Schwarz inequality and monotonicity of G , we can show that $G((Ux)_i) \leq G(\|U_i\|_2 \|x\|_2) \leq G(\|U_i\|_2 \|Ux\|_G) \leq G(\|U_i\|_2)$. This also leads to our definition of Orlicz norm leverage scores.

Definition 3. Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. For a given matrix $A \in \mathbb{R}^{n \times d}$ and a well-conditioned basis U of the column space of A , the *Orlicz norm leverage score* of the i -th row of A is defined to be $G(\|U_i\|_2)$.

It remains to give an upper bound on the summation of Orlicz norm leverage scores of all rows. Unlike the ℓ_2 norm, it is not immediately clear how to use the definition of well-conditioned basis to obtain such an upper bound for general Orlicz norms. To achieve this goal, we use a novel probabilistic argument. Suppose one takes x to be a vector with i.i.d. Gaussian random variables. Then each entry of Ux has the same distribution as $\|U_i\|_2 \cdot g_i$, where $\{g_i\}_{i=1}^n$ is a set of standard Gaussian random variables. Thus, with constant probability, $\sum_{i=1}^n G((Ux)_i)$ is an upper bound on the summation of Orlicz norm leverage scores. Furthermore, by the growth condition of the function G , we have $\sum_{i=1}^n G((Ux)_i) \leq C_G \|Ux\|_G^2$. Now by Definition 2, $\|Ux\|_G \leq \kappa_G \|x\|_2$, and $\|x\|_2 \leq O(\sqrt{d})$ with constant probability by tail inequalities of Gaussian random variables. This implies an upper bound on the summation of Orlicz norm leverage scores. See Lemma 4 and its proof for more details.

Our approach for constructing well-conditioned bases is inspired by [30]. In Lemma 5, we show that given a subspace embedding Π which embeds the column space of A with Orlicz norm $\|\cdot\|_G$ into the ℓ_2 space with distortion κ , then one can construct a well-conditioned basis with condition number $\kappa_G \leq \kappa$. The running time is dominated by calculating ΠA and doing a QR-decomposition on ΠA . To this end, we can use the oblivious subspace embedding for Orlicz norms in Corollary 12⁵ to construct well-conditioned bases. The embedding in Corollary 12 has $O(d)$ rows and $\kappa = \text{poly}(d \log n)$, and calculating ΠA can be done in $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time. Using such an embedding to construct the well-conditioned basis, our row sampling algorithm produces a vector w that satisfies (6) with $\|w\|_0 \leq \text{poly}(d \log n / \varepsilon)$ in time $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$.

We would like to remark that our sampling algorithm still works if the third condition in Assumption 1 does not hold. In general, suppose the function $G : \mathbb{R} \rightarrow \mathbb{R}$ satisfies that for all $0 < x < y$, $G(y)/G(x) \leq C_G (y/x)^p$, for the Orlicz norm induced by G , given a well-conditioned basis with condition number κ_G , our sampling algorithm returns a matrix with roughly $O((\sqrt{d} \kappa_G)^p \cdot d / \varepsilon^2)$ rows such that Theorem 1 holds. One may use the Löwner–John ellipsoid as the well-conditioned

⁴See, e.g., [24], for a survey on leverage scores.

⁵Alternatively, we can use the oblivious subspace embedding in [2] for this step. However, as we have discussed, the oblivious subspace embedding in [2] requires stronger assumptions on the function $G : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ than those in Assumption 1, which restricts the class of Orlicz norms to which our algorithm can be applied.

basis (as in [18]) which has condition number $\kappa_G = \sqrt{d}$ for any norm. However, calculating the Löwner–John ellipsoid requires at least $O(nd^5)$ time. Moreover, our method described above fails when $p > 2$ since it requires an oblivious subspace embedding with $\text{poly}(d)$ distortion, and it is known that such embedding does not exist when $p > 2$ [10]. Since we focus on input-sparsity time algorithms in this paper, we only consider the case that $p \leq 2$.

Finally, we would like to compare our sampling algorithm with that in [13]. First, the algorithm in [13] works for M -estimators, while we focus on Orlicz norms. Second, our definitions for Orlicz norm leverage score and well-conditioned basis, as given in Definition 2 and 3, are different from all previous works and are closely related to the Orlicz norm under consideration. The algorithm in [13], on the other hand, simply uses ℓ_p leverage scores. Under our definition, we can prove that the sum of leverage scores is bounded by $O(C_G d \kappa_G^2)$ (Lemma 4), whose proof requires a novel probabilistic argument. In contrast, the upper bound on sum of leverage scores in [13] is $O(\sqrt{nd})$ (Lemma 38 in [11]). Thus, the algorithm in [13] runs in an iterative manner since in each round the algorithm can merely reduce the dimension from n to $O(\sqrt{nd})$, while our algorithm is one-shot.

Oblivious Subspace Embeddings for Symmetric Norms. To obtain a faster algorithm for linear regression when the loss function is a general symmetric norm, we show that there exists a distribution over embedding matrices, such that if S is a random matrix drawn from that distribution, then for any $n \times d$ matrix A , with constant probability, for all $x \in \mathbb{R}^d$, $\|Ax\|_\ell \leq \|SAx\|_2 \leq \text{poly}(d \log n) \cdot \text{mmc}(\ell) \cdot \|Ax\|_\ell$. Moreover, the embedding matrix S is *sparse*, and calculating SA requires only $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time. Another favorable property of S is that it is an *oblivious subspace embedding*, meaning the distribution of S does not depend on A . To achieve this goal, it is sufficient to construct a random diagonal matrix D such that for any fixed vector $x \in \mathbb{R}^n$,

$$\Pr[\|Dx\|_2 \geq \Omega(1/\text{poly}(d \log n)) \cdot \|x\|_\ell] \geq 1 - \exp(-\Omega(d \log n)), \quad (7)$$

and

$$\Pr[\|Dx\|_2 \leq \text{poly}(d \log n) \cdot \text{mmc}(\ell) \cdot \|x\|_\ell] \geq 1 - O(1/d). \quad (8)$$

Our construction is inspired by the sub-sampling technique in [20], which was used for sketching symmetric norms in data streams [7]. Throughout the discussion, we use $\xi^{(q)} \in \mathbb{R}^n$ to denote a vector with q non-zero entries and each entry is $1/\sqrt{q}$. Let us start with a special case where the vector $x \in \mathbb{R}^n$ has s non-zero entries and each non-zero entry is 1. It is easy to see $\|x\|_\ell = \sqrt{s} \|\xi^{(s)}\|_\ell$. Now consider a random diagonal matrix D which corresponds to a sampling process, i.e., each diagonal entry is set to be 1 with probability p and 0 with probability $1 - p$. Our goal is to show that $\sqrt{1/p} \|\xi^{(1/p)}\|_\ell \cdot \|Dx\|_2$ is a good estimator of $\|x\|_\ell$. If $p = \Theta(d \log n/s)$, then with probability at least $1 - \exp(-\Omega(d \log n))$, Dx will contain at least one non-zero entry from x , in which case (7) is satisfied. However, we do not know s in advance. Thus, we use $t = O(\log n)$ different matrices D_1, D_2, \dots, D_t , where D_i has sampling probability $1/2^i$. Clearly at least one such D_j can establish (7). For the upper bound part, if p is much smaller than $1/s$, then Dx will never contain a non-zero entry from x . Otherwise, in expectation Dx will contain ps non-zero entries, in which case our estimation will be roughly $\sqrt{s} \|\xi^{(1/p)}\|_\ell$, which can be upper bounded by $O(\log n \cdot \text{mmc}(\ell) \cdot \sqrt{s} \|\xi^{(s)}\|_\ell)$. At this point, (8) follows from Markov’s inequality. See Section C.5 for the formal argument, and Section 3 for a detailed discussion on $\text{mmc}(\ell)$.

To generalize the above argument to general vectors, for a vector $x \in \mathbb{R}^n$, we conceptually partition its entries into $\Theta(\log n)$ groups, where the i -th group contains entries with magnitude in $[2^i, 2^{i+1})$. By averaging, at least one group of entries contributes at least $\Omega(1/\log n)$ fraction to the value of $\|x\|_\ell$. To establish (7), we apply the lower bound part of the argument in the previous paragraph to this “contributing” group. To establish (8), we apply the upper bound part of the argument to all groups, which will only induce an additional $O(\log n)$ factor in the approximation ratio, by triangle inequality.

Since our oblivious subspace embedding embeds a given symmetric norm into the ℓ_2 space, in order to obtain an approximate solution to symmetric norm regression, we only need to solve a least squares regression instance with much smaller size. This is another advantage of our subspace embedding, since the least square regression problem is a well-studied problem in optimization and numerical linear algebra, for which many efficient algorithms are known, both in theory and in practice.

2 Linear Regression for Orlicz Norms

In this section, we introduce our results for Orlicz norm regression. We first give the definition of weighted Orlicz norm.

Definition 4. For a function G that satisfies Assumption 1 and a weight vector $w \in \mathbb{R}^n$ such that $w_i \geq 0$ for all $i \in [n]$, for a vector $x \in \mathbb{R}^n$, if $\sum_{i=1}^n w_i \cdot |x_i| = 0$, then the weighted Orlicz norm $\|x\|_{G,w}$ is defined to be 0. Otherwise, the weighted Orlicz norm $\|x\|_{G,w}$ is defined as the unique value $\alpha > 0$ such that $\sum_{i=1}^n w_i G(|x_i|/\alpha) = 1$.

When $w_i = 1$ for all $i \in [n]$, we have $\|x\|_{G,w} = \|x\|_G$ where $\|x\|_G$ is the (unweighted) Orlicz norm. It is well known that $\|\cdot\|_G$ is a norm. We show in the following lemma that $\|\cdot\|_{G,w}$ is a seminorm.

Lemma 3. For a function G that satisfies Assumption 1 and a weight vector $w \in \mathbb{R}^n$ such that $w_i \geq 0$ for all $i \in [n]$, for all $x, y \in \mathbb{R}^n$, we have (i) $\|x\|_{G,w} \geq 0$, (ii) $\|x + y\|_{G,w} \leq \|x\|_{G,w} + \|y\|_{G,w}$, and (iii) $\|ax\|_{G,w} = |a| \cdot \|x\|_{G,w}$ for all $a \in \mathbb{R}$.

Leverage Scores and Well-Conditioned Bases for Orlicz Norms. The following lemma establishes an upper bound on the summation of Orlicz norm leverage scores defined in Definition 3.

Lemma 4. Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. Let $U \in \mathbb{R}^{n \times d}$ be a well-conditioned basis with condition number κ_G as in Definition 2. Then we have $\sum_{i=1}^n G(\|U_i\|_2) \leq O(C_G d \kappa_G^2)$,

Now we show that given a subspace embedding which embeds the column space of A with Orlicz norm $\|\cdot\|_G$ into the ℓ_2 space with distortion κ , then one can construct a well-conditioned basis with condition number $\kappa_G \leq \kappa$.

Lemma 5. Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. For a given matrix $A \in \mathbb{R}^{n \times d}$ and an embedding matrix $\Pi \in \mathbb{R}^{s \times n}$, suppose for all $x \in \mathbb{R}^d$, $\|Ax\|_G \leq \|\Pi Ax\|_2 \leq \kappa \|Ax\|_G$. Let $Q \cdot R = \frac{1}{\kappa} \Pi A$ be a QR-decomposition of $\frac{1}{\kappa} \Pi A$. Then AR^{-1} is a well-conditioned basis (see Definition 2) with $\kappa_G(AR^{-1}) \leq \kappa$.

The following lemma shows how to estimate Orlicz norm leverage scores given a change of basis matrix $R \in \mathbb{R}^{d \times d}$, in $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time.

Lemma 6. Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. For a given matrix $A \in \mathbb{R}^{n \times d}$ and $R \in \mathbb{R}^{d \times d}$, there exists an algorithm that outputs $\{u_i\}_{i=1}^n$ such that with probability at least 0.99, $u_i = \Theta(G(\|(AR^{-1})_i\|_2))$ for all $1 \leq i \leq n$. The algorithm runs in $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time.

The Row Sampling Algorithm. Based on the notion of Orlicz norm leverage scores and well-conditioned bases, we design a row sampling algorithm for Orlicz norms.

Lemma 7. Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. Let $U \in \mathbb{R}^{n \times d}$ be a well-conditioned basis with condition number $\kappa_G = \kappa_G(U)$ as in Definition 2. For sufficiently small ε and δ , and sufficiently large constant C , let $\{p_i\}_{i=1}^n$ be a set of sampling probabilities satisfying $p_i \geq \min\{1, C(\log(1/\delta) + d \log(1/\varepsilon))\varepsilon^{-2} G(\|U_i\|_2)\}$. Let w be a vector whose i -th entry is set to be $w_i = 1/p_i$ with probability p_i and $w_i = 0$ with probability $1 - p_i$, then with probability at least $1 - \delta$, for all $x \in \mathbb{R}^d$, we have $(1 - \varepsilon)\|Ux\|_G \leq \|Ux\|_{G,w} \leq (1 + \varepsilon)\|Ux\|_G$.

Solving Linear Regression for Orlicz Norms. Now we combine all ingredients to give an algorithm for Orlicz norm regression. We use $\bar{A} \in \mathbb{R}^{n \times (d+1)}$ to denote a matrix whose first d columns are A and the last column is b . The algorithm is described in Figure 1, and we prove its running time and correctness in Theorem 8. We assume we are given an embedding matrix Π , such that for all $x \in \mathbb{R}^{d+1}$, $\|\bar{A}x\|_G \leq \|\Pi \bar{A}x\|_2 \leq \kappa \|\bar{A}x\|_G$. The construction of Π and the value κ will be given in Corollary 12. In Section D.1 of the supplementary material, we use Theorem 8 and Corollary 12 to formally prove Theorem 1.

1. For the given embedding matrix Π , calculate $\Pi\bar{A}$ and invoke QR-decomposition on $\Pi\bar{A}/\kappa$ to obtain $Q \cdot R = \Pi\bar{A}/\kappa$.
2. Invoke Lemma 6 to obtain $\{u_i\}_{i=1}^n$ such that $u_i = \Theta(G(\|(\bar{A}R^{-1})_i\|_2))$.
3. For a sufficiently large constant C , let $\{p_i\}_{i=1}^n$ be a set of sampling probabilities with $p_i \geq \min\{1, C \cdot d \cdot \varepsilon^{-2} \log(1/\varepsilon) \cdot G(\|(\bar{A}R^{-1})_i\|_2)\}$, and w be a vector whose i -th entry $w_i = 1/p_i$ with probability p_i and $w_i = 0$ with probability $1 - p_i$.
4. Calculate $x^* = \operatorname{argmin}_{x \in \mathbb{R}^d} \|Ax - b\|_{G,w}$. Return x^* .

Figure 1: Algorithm for Orlicz norm regression

Theorem 8. *Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. Given an embedding matrix Π , such that for all $x \in \mathbb{R}^d$, $\|\bar{A}x\|_G \leq \|\Pi\bar{A}x\|_2 \leq \kappa\|\bar{A}x\|_G$, with probability at least 0.9, the algorithm in Figure 1 outputs $x^* \in \mathbb{R}^d$ in time $\operatorname{poly}(d\kappa/\varepsilon) + \mathcal{T}_{\text{QR}}(\Pi\bar{A})$, such that $\|Ax^* - b\|_G \leq (1 + \varepsilon) \min_{x \in \mathbb{R}^d} \|Ax - b\|_G$. Here, $\mathcal{T}_{\text{QR}}(\Pi\bar{A})$ is the running time for calculating $\Pi\bar{A}$ and invoking QR-decomposition on $\Pi\bar{A}$.*

3 Linear Regression for Symmetric Norms

In this section, we introduce **SymSketch**, a subspace embedding for symmetric norms.

Definition of SymSketch. We first formally define **SymSketch**. Due to space limitation, we give the definition of Gaussian embeddings, **CountSketch** embeddings and their compositions in Section C.1.1 of the supplementary material.

Definition 5 (Symmetric Norm Sketch (**SymSketch**)). Let $t = \Theta(\log n)$. Let $\tilde{D} \in \mathbb{R}^{n(t+1) \times n}$ be a matrix defined as $\tilde{D} = [(w_0 D_0)^\top \quad (w_1 D_1)^\top \quad \dots \quad (w_t D_t)^\top]^\top$, where for each $i \in \{0, 1, \dots, t\}$, $D_i = \operatorname{diag}(z_{i,1}, z_{i,2}, \dots, z_{i,n}) \in \mathbb{R}^{n \times n}$ and $z_{i,j} \sim \operatorname{Ber}(1/2^i)$ for each $j \in [n]$. Moreover, $w_i = \|(1, 1, \dots, 1, 0, \dots, 0)\|_\ell$ (there are 2^i 1s). Let $\Pi \in \mathbb{R}^{O(d) \times n(t+1)}$ be a composition of Gaussian embedding and **CountSketch** embedding (Definition 12) with $\varepsilon = 0.1$, and $S = \Pi\tilde{D}$. We say $S \in \mathbb{R}^{O(d) \times n}$ is a **SymSketch**.

Modulus of Concentration. Now we give the definition of $\operatorname{mmc}(\ell)$ for a symmetric norm.

Definition 6 ([7]). Let \mathcal{X} denote the uniform distribution over \mathbb{S}^{n-1} . The *median* of a symmetric norm $\|\cdot\|_\ell$ is the unique value M_ℓ such that $\Pr_{x \sim \mathcal{X}}[\|x\|_\ell \geq M_\ell] \geq 1/2$ and $\Pr_{x \sim \mathcal{X}}[\|x\|_\ell \leq M_\ell] \geq 1/2$.

Definition 7 ([7]). For a given symmetric norm $\|\cdot\|_\ell$, we define the *modulus of concentration* to be $\operatorname{mc}(\ell) = \max_{x \in \mathbb{S}^{n-1}} \|x\|_\ell / M_\ell$, and define the *maximum modulus of concentration* to be $\operatorname{mmc}(\ell) = \max_{k \in [n]} \operatorname{mc}(\ell^{(k)})$, where $\|\cdot\|_{\ell^{(k)}}$ is a norm on \mathbb{R}^k which is defined to be $\|(x_1, x_2, \dots, x_k)\|_{\ell^{(k)}} = \|(x_1, x_2, \dots, x_k, 0, \dots, 0)\|_\ell$.

It has been shown in [7] that $\operatorname{mmc}(\ell) = \Theta(n^{1/2-1/p})$ for ℓ_p norms when $p > 2$, $\operatorname{mmc}(\ell) = \Theta(1)$ for ℓ_p norms when $p \leq 2$, $\operatorname{mmc}(\ell) = \tilde{\Theta}(\sqrt{n/k})$ for top- k norms, and $\operatorname{mmc}(\ell) = O(\log n)$ for the k -support norm [5] and the box-norm [25]. We show that $\operatorname{mmc}(\ell)$ is upper bounded by $O(1)$ for max-mix of ℓ_2 norm and ℓ_1 norm and sum-mix of ℓ_2 norm and ℓ_1 norm.

Lemma 9. *For a real number $c > 0$, let $\|x\|_{\ell_a} = \|x\|_2 + c\|x\|_1$ and $\|x\|_{\ell_b} = \max\{\|x\|_2, c\|x\|_1\}$. We have $\operatorname{mmc}(\ell_a) = O(1)$ and $\operatorname{mmc}(\ell_b) = O(1)$.*

Moreover, we show that for an Orlicz norm $\|\cdot\|_G$ induced by a function G which satisfies Assumption 1, $\operatorname{mmc}(\ell)$ is upper bounded by $O(\sqrt{C_G} \log n)$.

Lemma 10. *For an Orlicz norm $\|\cdot\|_G$ on \mathbb{R}^n induced by a function G which satisfies Assumption 1, $\operatorname{mmc}(\ell)$ is upper bounded by $O(\sqrt{C_G} \log n)$.*

Subspace Embedding. The following theorem shows that **SymSketch** is a subspace embedding.

Theorem 11. *Let $S \in \mathbb{R}^{O(d) \times n}$ be a **SymSketch** as defined in Definition 5. For a given matrix $A \in \mathbb{R}^{n \times d}$, with probability at least 0.9, for all $x \in \mathbb{R}^d$,*

$$\Omega\left(1/(\sqrt{d} \cdot \log^3 n)\right) \cdot \|Ax\|_\ell \leq \|SAx\|_2 \leq O\left(\text{mmc}(\ell) \cdot d^2 \cdot \log^{5/2} n\right) \cdot \|Ax\|_\ell.$$

Furthermore, the running time of computing SA is $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$.

Combine Theorem 11 with Lemma 10, we have the following corollary.

Corollary 12. *Let $\|\cdot\|_G$ be an Orlicz norm induced by a function G which satisfies Assumption 1. Let $S \in \mathbb{R}^{O(d) \times n}$ be a **SymSketch** as defined in Definition 5. For a given matrix $A \in \mathbb{R}^{n \times d}$, with probability at least 0.9, for all $x \in \mathbb{R}^d$,*

$$\Omega\left(1/(\sqrt{d} \cdot \log^3 n)\right) \cdot \|Ax\|_\ell \leq \|SAx\|_2 \leq O\left(\sqrt{C_G} \cdot d^2 \cdot \log^{7/2} n\right) \cdot \|Ax\|_\ell.$$

Furthermore, the running time of computing SA is $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$.

4 Conclusion

In this paper, we give efficient algorithms for solving the overconstrained linear regression problem, when the loss function is a symmetric norm. For the special case when the loss function is an Orlicz norm, our algorithm produces a $(1 + \varepsilon)$ -approximate solution in $\tilde{O}(\text{nnz}(A) + \text{poly}(d/\varepsilon))$ time. When the loss function is a general symmetric norm, our algorithm produces a $\sqrt{d} \cdot \text{polylog } n \cdot \text{mmc}(\ell)$ -approximate solution in $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time.

In light of Problem 1, there are a few interesting problems that remain open. Is that possible to design an algorithm that produces $(1 + \varepsilon)$ -approximate solutions to the linear regression problem, when the loss function is a general symmetric norm? Furthermore, is that possible to use the technique of linear sketching to speed up the overconstrained linear regression problem, when the loss function is a general norm? Answering these problems could lead to a better understanding of Problem 1.

Acknowledgements

P. Zhong is supported in part by NSF grants (CCF-1703925, CCF-1421161, CCF-1714818, CCF-1617955 and CCF-1740833), Simons Foundation (#491119), Google Research Award and a Google Ph.D. fellowship. R. Wang is supported in part by NSF grant IIS-1763562, Office of Naval Research (ONR) grants (N00014-18-1-2562, N00014-18-1-2861), and Nvidia NVAIL award. Part of this work was done while Z. Song, L. F. Yang, H. Zhang and P. Zhong were interns at IBM Research - Almaden and while Z. Song, R. Wang and H. Zhang were visiting the Simons Institute for the Theory of Computing. Z. Song and P. Zhong would like to thank Alexandr Andoni, Kenneth L. Clarkson, Yin Tat Lee, Eric Price, Clifford Stein and David P. Woodruff for insight discussions.

References

- [1] N. Ailon and B. Chazelle. Approximate nearest neighbors and the fast johnson-lindenstrauss transform. In *STOC*, pages 557–563, 2006.
- [2] A. Andoni, C. Lin, Y. Sheng, P. Zhong, and R. Zhong. Subspace embedding and linear regression with Orlicz norm. In *ICML*, pages 224–233, 2018.
- [3] A. Andoni, A. Naor, A. Nikolov, I. Razenshteyn, and E. Waingarten. Hölder homeomorphisms and approximate nearest neighbors. In *FOCS*, pages 159–169, 2018.
- [4] A. Andoni, H. L. Nguyen, A. Nikolov, I. Razenshteyn, and E. Waingarten. Approximate near neighbors for general symmetric norms. In *STOC*, pages 902–913, 2017.
- [5] A. Argyriou, R. Foygel, and N. Srebro. Sparse prediction with the k -support norm. In *NIPS*, pages 1457–1465, 2012.
- [6] H. Auerbach. *On the area of convex curves with conjugate diameters*. PhD thesis, University of Lwów, 1930.
- [7] J. Błasiok, V. Braverman, S. R. Chestnut, R. Krauthgamer, and L. F. Yang. Streaming symmetric norms via measure concentration. In *STOC*, pages 716–729, 2017.
- [8] P. Brändén. Hyperbolic polynomials and the kadison-singer problem. *arXiv preprint arXiv:1809.03255*, 2018.
- [9] V. Braverman, S. R. Chestnut, D. P. Woodruff, and L. F. Yang. Streaming space complexity of nearly all functions of one variable on frequency vectors. In *PODS*, pages 261–276, 2016.
- [10] V. Braverman and R. Ostrovsky. Zero-one frequency laws. In *STOC*, pages 281–290, 2010.
- [11] K. L. Clarkson, R. Wang, and D. P. Woodruff. Dimensionality reduction for tukey regression. In *ICML*, pages 1262–1271, 2019.
- [12] K. L. Clarkson and D. P. Woodruff. Low rank approximation and regression in input sparsity time. In *STOC*, pages 81–90, 2013.
- [13] K. L. Clarkson and D. P. Woodruff. Input sparsity and hardness for robust subspace approximation. In *FOCS*, pages 310–329, 2015.
- [14] K. L. Clarkson and D. P. Woodruff. Sketching for M-estimators: A unified approach to robust regression. In *SODA*, pages 921–939, 2015.
- [15] M. B. Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In *SODA*, pages 278–287, 2016.
- [16] M. B. Cohen and R. Peng. ℓ_p row sampling by lewis weights. In *STOC*, pages 183–192, 2015.
- [17] D. Dadush, A. Nikolov, K. Talwar, and N. Tomczak-Jaegermann. Balancing vectors in any norm. In *FOCS*, 2018.
- [18] A. Dasgupta, P. Drineas, B. Harb, R. Kumar, and M. W. Mahoney. Sampling algorithms and coresets for ℓ_p regression. *SIAM Journal on Computing*, 38(5):2060–2078, 2009.
- [19] V. Feldman, C. Guzmán, and S. S. Vempala. Statistical query algorithms for mean vector estimation and stochastic convex optimization. In *SODA*, pages 1265–1277, 2017.
- [20] P. Indyk and D. P. Woodruff. Optimal approximations of the frequency moments of data streams. In *STOC*, pages 202–208, 2005.
- [21] Y. T. Lee, Z. Song, and S. S. Vempala. Algorithmic theory of odes and sampling from well-conditioned logconcave densities. *arXiv preprint arXiv:1812.06243*, 2018.
- [22] J. Li, A. Nikolov, I. Razenshteyn, and E. Waingarten. On mean estimation for general norms with statistical queries. In *COLT*, 2019.

- [23] M. Li, G. L. Miller, and R. Peng. Iterative row sampling. In *FOCS*, pages 127–136, 2013.
- [24] M. W. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends® in Machine Learning*, 3(2):123–224, 2011.
- [25] A. M. McDonald, M. Pontil, and D. Stamos. Spectral k -support norm regularization. In *NIPS*, pages 3644–3652, 2014.
- [26] X. Meng and M. W. Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *STOC*, pages 91–100, 2013.
- [27] V. Nakos, Z. Song, and Z. Wang. Robust sparse recovery via m-estimators. *Manuscript*, 2019.
- [28] J. Nelson and H. L. Nguyễn. Sparsity lower bounds for dimensionality reducing maps. In *STOC*, pages 101–110, 2013.
- [29] T. Sarlócs. Improved approximation algorithms for large matrices via random projections. In *FOCS*, pages 143–152, 2006.
- [30] C. Sohler and D. P. Woodruff. Subspace embeddings for the l_1 -norm with applications. In *STOC*, pages 755–764, 2011.
- [31] Z. Song, D. P. Woodruff, and P. Zhong. Towards a zero-one law for column subset selection. In *NeurIPS*, 2019.
- [32] R. Wang and D. P. Woodruff. Tight bounds for ℓ_p oblivious subspace embeddings. In *SODA*, pages 1825–1843, 2019.
- [33] P. Wojtaszczyk. *Banach spaces for analysts*, volume 25. Cambridge University Press, 1996.
- [34] D. P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- [35] D. P. Woodruff and Q. Zhang. Subspace embeddings and ℓ_p -regression using exponential random variables. In *COLT*, pages 546–567, 2013.
- [36] J. Yang, X. Meng, and M. Mahoney. Quantile regression for large-scale applications. In *ICML*, pages 881–887, 2013.

Supplementary Material for “Efficient Symmetric Norm Regression via Linear Sketching”

A Preliminaries

Notations. For a matrix $A \in \mathbb{R}^{n \times d}$, we use A_i to denote its i -th row, A^i to denote its i -th column, $\|A\|_F$ to denote the Frobenius norm of A , and $\|A\|_2$ to denote the spectral norm of A . For any $n' \leq n$, we define $\xi^{(n')} \in \mathbb{R}^n$ to be a vector $\xi^{(n')} = \frac{1}{\sqrt{n'}}(1, 1, \dots, 1, 0, 0, \dots, 0)$.

ε -nets. We use the standard upper bound on size of ε -nets.

Definition 8. For a given set \mathcal{S} and a norm $\|\cdot\|$, we say $\mathcal{N} \subseteq \mathcal{S}$ is a ε -net of \mathcal{S} if for any $s \in \mathcal{S}$, there exists some $\bar{s} \in \mathcal{N}$ such that $\|s - \bar{s}\| \leq \varepsilon$.

Lemma 13 ([33, II.E, 10]). *Given a matrix $A \in \mathbb{R}^{n \times d}$ and a norm $\|\cdot\|$, let \mathcal{S} be the unit $\|\cdot\|$ -norm ball in the column space of A , i.e., $\mathcal{S} = \{Ax \mid \|Ax\| = 1\}$. For $\varepsilon \in (0, 1)$, there exists an ε -net \mathcal{N} of \mathcal{S} with size $|\mathcal{N}| \leq (1 + 1/\varepsilon)^d$.*

B Missing Proofs in Section 2

In this section, we give missing proofs in Section 2.

We first show that if a function G satisfies Assumption 1, then G has at least linear growth. We will use this fact in later proofs.

Lemma 14. *Given a function G that satisfies property \mathcal{P} , then for any $0 < x \leq y$, $y/x \leq G(y)/G(x)$.*

Proof. Due to the convexity of G and $G(0) = 0$, for any $y > x > 0$, we have

$$G(x) \leq G(y)x/y + G(0)(1 - x/y) = G(y)x/y.$$

□

B.1 Proof of Lemma 3

Proof. The first condition is clear from the definition of $\|x\|_{G,w}$.

Now we prove the second condition. When $\|x + y\|_{G,w} = 0$, the triangle inequality clearly holds since $\|x\|_{G,w} \geq 0$ and $\|y\|_{G,w} \geq 0$. When $\|x\|_{G,w} = 0$ and $\|x + y\|_{G,w} \neq 0$, for any $\alpha > 0$, we have

$$\sum_{i=1}^n w_i G(|x_i + y_i|/\alpha) = \sum_{i|w_i > 0} w_i G(|x_i + y_i|/\alpha) = \sum_{i|w_i > 0} w_i G(|y_i|/\alpha) = \sum_{i=1}^n w_i G(|y_i|/\alpha),$$

which implies $\|x + y\|_{G,w} = \|y\|_{G,w}$. Similarly, the second condition also holds if $\|y\|_{G,w} = 0$ and $\|x + y\|_{G,w} \neq 0$. If $\|x + y\|_{G,w} \neq 0$, $\|x\|_{G,w} \neq 0$ and $\|y\|_{G,w} \neq 0$, by definition of $\|\cdot\|_{G,w}$, we have

$$\sum_{i=1}^n w_i G(x_i/\|x\|_{G,w}) = 1$$

and

$$\sum_{i=1}^n w_i G(y_i/\|y\|_{G,w}) = 1.$$

Thus,

$$\begin{aligned}
& \sum_{i=1}^n w_i G\left(\frac{x_i + y_i}{\|x\|_{G,w} + \|y\|_{G,w}}\right) \\
& \leq \sum_{i=1}^n w_i G\left(\frac{|x_i| + |y_i|}{\|x\|_{G,w} + \|y\|_{G,w}}\right) \quad (G \text{ is increasing}) \\
& \leq \sum_{i=1}^n w_i \left(\frac{\|x\|_{G,w}}{\|x\|_{G,w} + \|y\|_{G,w}} \cdot G(|x_i|/\|x\|_{G,w}) + \frac{\|y\|_{G,w}}{\|x\|_{G,w} + \|y\|_{G,w}} \cdot G(|y_i|/\|y\|_{G,w}) \right) \\
& \quad (G \text{ is convex}) \\
& = 1,
\end{aligned}$$

which implies $\|x + y\|_{G,w} \leq \|x\|_{G,w} + \|y\|_{G,w}$.

For the third condition, for any $a \in \mathbb{R}$ and $x \in \mathbb{R}^n$, if $\|x\|_{G,w} = 0$ then $\|ax\|_{G,w} = 0$. If $a = 0$, we have $\|ax\|_{G,w} = 0$. Otherwise, we have

$$\sum_{i=1}^n w_i G(x_i/\|x\|_{G,w}) = 1,$$

which implies

$$\sum_{i=1}^n w_i G\left(\frac{a \cdot x_i}{|a|\|x\|_{G,w}}\right) = 1,$$

and thus $\|ax\|_{G,w} = |a|\|x\|_{G,w}$. □

B.2 Proof of Lemma 4

Proof. Let $g \in \mathbb{R}^d$ be a vector whose entries are i.i.d. Gaussian random variables with zero mean and standard deviation 10^2 . We show that with probability at least 0.8,

$$\sum_{i=1}^n G(\|U_i\|_2) \leq O\left(\sum_{i=1}^n G(\langle U_i, g \rangle)\right) \leq O(\max\{1, C_G \|Ug\|_G^2\}) \leq O(C_G d \kappa_G^2).$$

We divide our proofs into three parts.

Part I. We will show that with probability at least 0.9,

$$\sum_{i=1}^n G(\|U_i\|_2) \leq O\left(\sum_{i=1}^n G(\langle U_i, g \rangle)\right).$$

For each $i \in [n]$, $\langle U_i, g \rangle$ has the same distribution as $10^2 \cdot \|U_i\|_2 \cdot \mathcal{N}(0, 1)$. For each $i \in [n]$, we let B_i be the random variable such that

$$B_i = \begin{cases} 1 & |\langle U_i, g \rangle| \leq \|U_i\|_2 \\ 0 & \text{otherwise} \end{cases}.$$

By tail inequalities of standard Gaussian random variables, $\Pr[B_i = 1] \leq 0.01$. Thus,

$$\mathbf{E}[B_i \cdot G(\|U_i\|_2)] \leq 0.01 \cdot G(\|U_i\|_2),$$

which implies

$$\mathbf{E}\left[\sum_{i=1}^n B_i \cdot G(\|U_i\|_2)\right] \leq 0.01 \cdot \sum_{i=1}^n G(\|U_i\|_2),$$

By the monotonicity of G , since

$$G(\langle U_i, g \rangle) \geq (1 - B_i)G(\|U_i\|_2),$$

we have

$$\sum_{i=1}^n G(\langle U_i, g \rangle) \geq \sum_{i=1}^n (1 - B_i) G(\|U_i\|_2).$$

By Markov's inequality, with probability at least 0.9, we have

$$\sum_{i=1}^n B_i \cdot G(\|U_i\|_2) \leq 0.1 \sum_{i=1}^n G(\|U_i\|_2),$$

which implies

$$\sum_{i=1}^n G(\langle U_i, g \rangle) \geq 0.9 \sum_{i=1}^n G(\|U_i\|_2).$$

Part II. We will show that

$$\sum_{i=1}^n G(\langle U_i, g \rangle) \leq \max\{1, C_G \cdot \|Ug\|_G^2\}.$$

When $\|Ug\|_G \leq 1$, by monotonicity of G , we must have

$$\sum_{i=1}^n G(\langle U_i, g \rangle) \leq 1.$$

When $\|Ug\|_G \geq 1$, we have

$$\sum_{i=1}^n G(\langle U_i, g \rangle / \|Ug\|_G) = 1.$$

Since

$$G(\langle U_i, g \rangle) \leq G(\langle U_i, g \rangle / \|Ug\|_G) \cdot C_G \|Ug\|_G^2$$

and

$$\sum_{i=1}^n G(\langle U_i, g \rangle / \|Ug\|_G) = 1,$$

we must have

$$\sum_{i=1}^n G(\langle U_i, g \rangle) \leq \sum_{i=1}^n G(\langle U_i, g \rangle / \|Ug\|_G) \cdot C_G \cdot \|Ug\|_G^2 \leq C_G \cdot \|Ug\|_G^2.$$

Part III. We will show that $\|Ug\|_G^2 \leq O(C_G d \kappa_G^2)$. By definition of a well-conditioned basis and tail inequalities of Gaussian random variables, with probability at least 0.9, we have

$$\|Ug\|_G \leq \kappa_G \|g\|_2 \leq O(\kappa_G \sqrt{d}).$$

Thus, applying a union bound over three parts of the proof, we have with probability at least 0.8,

$$\sum_{i=1}^n G(\|U_i\|_2) \leq O(C_G d \kappa_G^2). \quad (9)$$

However, the condition in (9) is deterministic. Thus, the condition in (9) always holds. \square

B.3 Proof of Lemma 5

Proof. Notice that for any $x \in \mathbb{R}^d$,

$$\|AR^{-1}x\|_G \leq \|\Pi AR^{-1}x\|_2 = \kappa \|Qx\|_2 = \kappa \|x\|_2$$

and

$$\|AR^{-1}x\|_G \geq \frac{1}{\kappa} \|\Pi AR^{-1}x\|_2 = \|Qx\|_2 = \|x\|_2.$$

\square

B.4 Proof of Lemma 6

Proof. In Theorem 2.13 of [34], it has been shown how to calculate $\{l_i\}_{i=1}^n$ such that $l_i = \Theta(\|(AR^{-1})_i\|_2)$ in $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time with probability at least 0.99. We simply take $u_i = G(l_i)$. By Lemma 14 and the growth condition of G , we must have $u_i = \Theta(G(\|(AR^{-1})_i\|_2))$. \square

B.5 Proof of Lemma 7

Proof. By homogeneity, we only need to prove that with probability $1 - \delta$, for all x which satisfies $\|Ux\|_G = 1$,

$$(1 - \varepsilon)\|Ux\|_G \leq \|Ux\|_{G,w} \leq (1 + \varepsilon)\|Ux\|_G.$$

We first prove that for any fixed $x \in \mathbb{R}^d$ such that $\|Ux\|_G = 1$, with probability $1 - \delta(1 + 4/\varepsilon)^{-d}$,

$$(1 - \varepsilon/4)\|Ux\|_G \leq \|Ux\|_{G,w} \leq (1 + \varepsilon/4)\|Ux\|_G.$$

Let $x \in \mathbb{R}^d$ that satisfies $\|Ux\|_G = 1$ and $y = Ux$. Let Z_i be a random variable which denotes the value of $w_i G(y_i)$ and $Z = \sum_{i=1}^n Z_i$.

We will first show that if $Z \in [1 - \varepsilon/4, 1 + \varepsilon/4]$, then $\|y\|_{G,w} \in [1 - \varepsilon/4, 1 + \varepsilon/4]$. There are three cases:

1. If $\|y\|_{G,w} = 1$, then $\|y\|_{G,w}$ is already in $[1 - \varepsilon/4, 1 + \varepsilon/4]$.
2. If $\|y\|_{G,w} > 1$, then by Lemma 14, we have

$$\sum_{i=1}^n w_i G(y_i) \geq \sum_{i=1}^n w_i \|y\|_{G,w} \cdot G(y_i / \|y\|_{G,w}).$$

Since

$$\sum_{i=1}^n w_i \cdot G(y_i / \|y\|_{G,w}) = 1,$$

we must have

$$\|y\|_{G,w} \leq \sum_{i=1}^n w_i G(y_i) = Z \leq 1 + \varepsilon/4.$$

3. If $\|y\|_{G,w} < 1$, then by Lemma 14, we have

$$1 = \sum_{i=1}^n w_i G(y_i / \|y\|_{G,w}) \geq 1 / \|y\|_{G,w} \cdot \sum_{i=1}^n w_i G(y_i),$$

which implies

$$\|y\|_{G,w} \geq \sum_{i=1}^n w_i G(y_i) = Z \geq 1 - \varepsilon/4.$$

Thus, it suffices to prove that

$$\Pr[Z \in [1 - \varepsilon/4, 1 + \varepsilon/4]] \geq 1 - \delta(1 + 4/\varepsilon)^{-d}.$$

Consider the expectation of Z , we have

$$\mathbf{E}[Z] = \sum_{i=1}^n \mathbf{E}[Z_i] = \sum_{i=1}^n \mathbf{E}[w_i] \cdot G((Ux)_i) = \sum_{i=1}^n G((Ux)_i) = 1,$$

where the last equality follows since $\|Ux\|_G = 1$.

Notice that $|Z_i - \mathbf{E}[Z_i]|$ is always upper bounded by

$$\begin{aligned} w_i G(y_i) - \mathbf{E}[Z_i] &= w_i G((Ux)_i) - \mathbf{E}[w_i G((Ux)_i)] \leq w_i G(\|U_i\|_2 \cdot \|x\|_2) - \mathbf{E}[w_i G(\|U_i\|_2)] \\ &\leq G(\|U_i\|_2) / p_i \leq \frac{\varepsilon^2}{C(\log(1/\delta) + d \log(1/\varepsilon))}, \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz inequality, the second inequality follows from the definition of well-conditioned basis in Definition 2 and monotonicity of G , the third inequality follows from definition of w_i and the last inequality follows from the choice of p_i .

Consider the variance of Z , we have:

$$\begin{aligned} \text{Var}(Z) &= \sum_{i|p_i < 1} \text{Var}(Z_i) \leq \sum_{i|p_i < 1} \mathbf{E}(Z_i^2) = \sum_{i|p_i < 1} (G((Ux)_i))^2 / p_i \\ &\leq \left(\sum_{i|p_i < 1} G((Ux)_i) \right) \cdot \max_{i|p_i < 1} G((Ux)_i) / p_i \leq \frac{\varepsilon^2}{C(\log(1/\delta) + d \log(1/\varepsilon))}, \end{aligned}$$

where the second inequality follows from Hölder's inequality and the last inequality follows from the upper bound of $G((Ux)_i)/p_i$ and $\|Ux\|_G = 1$.

Thus, by Bernstein inequality, we have:

$$\Pr(|Z - 1| > \varepsilon/4) \leq (1 + 4/\varepsilon)^{-d} \delta.$$

Thus, for a fixed x , with probability at least $1 - (1 + 4/\varepsilon)^{-d} \delta$, we have

$$(1 - \varepsilon/4)\|Ux\|_G \leq \|Ux\|_{G,w} \leq (1 + \varepsilon/4)\|Ux\|_G.$$

Let \mathcal{S} be the unit $\|\cdot\|_G$ -norm ball in the column space of U , i.e., $\mathcal{S} = \{Ux \mid \|Ux\|_G = 1\}$. According to Lemma 13, there exists an $\varepsilon/4$ -net \mathcal{N} of \mathcal{S} with $|\mathcal{N}| \leq (1 + 4/\varepsilon)^d$. We use \mathcal{E} to denote the event that for all $y \in \mathcal{N}$, $\|y\|_{G,w} \in [1 - \varepsilon/4, 1 + \varepsilon/4]$. By taking union bound over all vectors in \mathcal{N} , we have $\Pr[\mathcal{E}] \geq 1 - \delta$.

Conditioned on \mathcal{E} , now we show that for all $y \in \mathcal{S}$, $\|y\|_{G,w} \in [1 - \varepsilon, 1 + \varepsilon]$. Consider a fixed vector $y \in \mathcal{S}$, since \mathcal{N} is an $\varepsilon/4$ -net of \mathcal{S} , we can choose a vector $u^{(1)} \in \mathcal{N}$ such that

$$\|y - u^{(1)}\|_G \leq \varepsilon/4.$$

Thus, we have that

$$\|y\|_{G,w} \leq \|u^{(1)}\|_{G,w} + \|y - u^{(1)}\|_{G,w} \leq (1 + \varepsilon/4) + \|y - u^{(1)}\|_{G,w}.$$

Let $\alpha^{(1)} = 1/\|y - u^{(1)}\|_G$. Then we have $\alpha^{(1)}(y - u^{(1)}) \in \mathcal{S}$. Thus, there exist $u^{(2)} \in \mathcal{N}$ such that

$$\|u^{(2)} - \alpha^{(1)}(y - u^{(1)})\|_G \leq \varepsilon/4.$$

It implies that

$$\|(y - u^{(1)}) - u^{(2)}/\alpha^{(1)}\|_G \leq \varepsilon/(4\alpha^{(1)}) \leq (\varepsilon/4)^2.$$

Thus,

$$\|y - u^{(1)}\|_{G,w} \leq \|u^{(2)}\|_{G,w}/\alpha^{(1)} + \|y - u^{(1)} - u^{(2)}/\alpha^{(1)}\|_{G,w} \leq (1 + \varepsilon/4)\varepsilon/4 + \|y - u^{(1)} - u^{(2)}/\alpha^{(1)}\|_{G,w}.$$

Let $\alpha^{(2)} = 1/\|y - u^{(1)} - u^{(2)}/\alpha^{(1)}\|_G$. Then we can repeat the above argument and get

$$\begin{aligned} \|y\|_{G,w} &\leq (1 + \varepsilon/4) + (1 + \varepsilon/4)\varepsilon/4 + (1 + \varepsilon/4)(\varepsilon/4)^2 + \dots \\ &= (1 + \varepsilon/4)/(1 - \varepsilon/4) \leq 1 + \varepsilon. \end{aligned}$$

By applying the above upper bound on $\|\alpha^{(1)}(u^{(1)} - y)\|_{G,w}$, we can get

$$\begin{aligned} \|y\|_{G,w} &\geq \|u^{(1)}\|_{G,w} - \|u^{(1)} - y\|_{G,w} \\ &\geq (1 - \varepsilon/4) - \|u^{(1)} - y\|_{G,w} \\ &\geq (1 - \varepsilon/4) - \frac{1 + \varepsilon}{\alpha^{(1)}} \\ &\geq 1 - \varepsilon/2 - \varepsilon^2/4 \\ &\geq 1 - \varepsilon. \end{aligned}$$

Thus, conditioned on \mathcal{E} , which holds with probability $1 - \delta$, we have $\|y\|_{G,w} \in [1 - \varepsilon, 1 + \varepsilon]$ for all $y = Ux$ with $\|y\|_G = 1$. \square

B.6 Proof of Theorem 8

Proof. We first analyze the running time of the algorithm. In Step 1, we calculate $\Pi\bar{A}$ and invoke QR-decomposition on $\Pi\bar{A}$. In Step 2, we apply the algorithm in Lemma 6, which runs in $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time. Obtaining the weight vector $w \in \mathbb{R}^n$ in Step 3 requires $O(n)$ time.

Since for all $x \in \mathbb{R}^d$,

$$\|\bar{A}x\|_G \leq \|\Pi\bar{A}x\|_2 \leq \kappa\|\bar{A}x\|_G.$$

we have

$$\begin{aligned} \mathbf{E}[\|w\|_0] &= \sum_{i=1}^n p_i = \sum_{i=1}^n O(d \log(1/\varepsilon)/\varepsilon^2 \cdot G(\|(\bar{A}R^{-1})_i\|_2)) \\ &\leq O\left(d \log(1/\varepsilon)/\varepsilon^2 \cdot C_G d (\kappa_G(\bar{A}R^{-1}))^2\right) && \text{(Lemma 4)} \\ &\leq O(C_G d^2 \kappa^2 \log(1/\varepsilon)/\varepsilon^2). && \text{(Lemma 5)} \end{aligned}$$

By Markov's inequality, with constant probability we have $\|w\|_0 \leq O(C_G d^2 \kappa^2 \log(1/\varepsilon)/\varepsilon^2)$. Moreover, in order to solve $\min_x \|Ax - b\|_{G,w}$, we can ignore all rows of A with zero weights, and thus there are at most $O(C_G d^2 \kappa^2 \log(1/\varepsilon)/\varepsilon^2)$ remaining rows in A . Furthermore, as we show in Lemma 3, $\|\cdot\|_{G,w}$ is a seminorm, which implies we can solve $\min_x \|Ax - b\|_{G,w}$ in $\text{poly}(C_G d \kappa/\varepsilon)$ time, by simply solve a convex program with size $O(C_G d^2 \kappa^2 \log(1/\varepsilon)/\varepsilon^2)$.

Now we prove the correctness of the algorithm. The algorithm in Lemma 6 succeeds with constant probability. By Lemma 7, with constant probability, simultaneously for all $x \in \mathbb{R}^{d+1}$,

$$(1 - \varepsilon/3)\|\bar{A}R^{-1}x\|_G \leq \|\bar{A}R^{-1}x\|_{G,w} \leq (1 + \varepsilon/3)\|\bar{A}R^{-1}x\|_G.$$

Equivalently, with constant probability, simultaneously for all $x \in \mathbb{R}^{d+1}$,

$$(1 - \varepsilon/3)\|\bar{A}x\|_G \leq \|\bar{A}x\|_{G,w} \leq (1 + \varepsilon/3)\|\bar{A}x\|_G.$$

Since $x^* = \text{argmin}_x \|Ax - b\|_{G,w}$, for all $x \in \mathbb{R}^d$, we have

$$\begin{aligned} \|Ax^* - b\|_G &\leq 1/(1 - \varepsilon/3)\|Ax^* - b\|_{G,w} \leq 1/(1 - \varepsilon/3)\|Ax - b\|_{G,w} \\ &\leq (1 + \varepsilon/3)/(1 - \varepsilon/3)\|Ax - b\|_G \leq (1 + \varepsilon)\|Ax - b\|_G \end{aligned}$$

for sufficiently small ε . Thus, x^* is a $(1 + \varepsilon)$ -approximate solution to $\min_x \|Ax - b\|_G$.

Note that the failure probability of the algorithm can be reduced to an arbitrarily small constant by independent repetitions and taking the best solution found among all repetitions. \square

C Missing Proofs in Section 3

In this section, we give missing proofs in Section 3.

Without loss of generality, throughout this section, for the symmetric norm $\|\cdot\|_\ell$ under consideration, we assume $\|\xi^{(1)}\|_\ell = 1$.

C.1 Background

C.1.1 Known ℓ_2 Oblivious Subspace Embeddings

In this section, we recall some known ℓ_2 subspace embeddings.

Definition 9. We say $S \in \mathbb{R}^{t \times n}$ is an ℓ_2 subspace embedding for the column space of $A \in \mathbb{R}^{n \times d}$, if for all $x \in \mathbb{R}^d$,

$$(1 - \varepsilon)\|Ax\|_2 \leq \|SAx\|_2 \leq (1 + \varepsilon)\|Ax\|_2.$$

Definition 10. A CountSketch embedding is defined to be $\Pi = \Phi D \in \mathbb{R}^{m \times n}$ with $m = \Theta(d^2/\varepsilon^2)$, where D is an $n \times n$ random diagonal matrix with each diagonal entry independently chosen to be $+1$ or -1 with equal probability, and $\Phi \in \{0, 1\}^{m \times n}$ is an $m \times n$ binary matrix with $\Phi_{h(i),i} = 1$ and all remaining entries being 0, where $h: [n] \rightarrow [m]$ is a random map such that for each $i \in [n]$, $h(i) = j$ with probability $1/m$ for each $j \in [m]$.

Theorem 15 ([12]). For a given matrix $A \in \mathbb{R}^{n \times d}$ and $\varepsilon \in (0, 1/2)$. Let $\Pi \in \mathbb{R}^{\Theta(d^2/\varepsilon^2) \times n}$ be a *CountSketch* embedding. With probability at least 0.9999, Π is an ℓ_2 subspace embedding for the column space of A . Furthermore, ΠA can be computed in $O(\text{nnz}(A))$ time.

Definition 11. A Gaussian embedding S is defined to be $\frac{1}{\sqrt{m}} \cdot G \in \mathbb{R}^{m \times n}$ with $m = \Theta(d/\varepsilon^2)$, where each entry of $G \in \mathbb{R}^{m \times n}$ is chosen independently from the standard Gaussian distribution.

Theorem 16 ([34]). For a given matrix $A \in \mathbb{R}^{n \times d}$ and $\varepsilon \in (0, 1/2)$. Let $S \in \mathbb{R}^{\Theta(d/\varepsilon^2) \times n}$ be a Gaussian embedding. With probability at least 0.9999, S is an ℓ_2 subspace embedding for the column space of A .

Definition 12. A composition of Gaussian embedding and *CountSketch* embedding is defined to be $S' = S\Pi$, where $\Pi \in \mathbb{R}^{\Theta(d^2/\varepsilon^2) \times n}$ is a *CountSketch* embedding and $S \in \mathbb{R}^{\Theta(d/\varepsilon^2) \times \Theta(d^2/\varepsilon^2)}$ is a Gaussian embedding.

The following corollary directly follows from the above two theorems.

Corollary 17. For a given matrix $A \in \mathbb{R}^{n \times d}$ and $\varepsilon \in (0, 1/2)$. Let $S' \in \mathbb{R}^{\Theta(d/\varepsilon^2) \times n}$ be a composition of Gaussian embedding and *CountSketch* embedding. With probability at least 0.9998, S' is an ℓ_2 subspace embedding for the column space of A . Furthermore, $S' A$ can be computed in $O(\text{nnz}(A) + d^4/\varepsilon^4)$ time.

We remark that all ℓ_2 subspace embeddings introduced in this section are *oblivious*, meaning that the distribution of the embedding matrix does not depend on the matrix A .

C.1.2 Properties of Symmetric Norms

General Properties. We first introduce several general properties of symmetric norms.

Lemma 18 (Lemma 2.1 in [7]). For any symmetric norm $\|\cdot\|_\ell$ and $x, y \in \mathbb{R}^n$ such that for all $i \in [n]$ we have $|x_i| \leq |y_i|$, then $\|x\|_\ell \leq \|y\|_\ell$.

Lemma 19 (Fact 2.2 in [7]). Suppose $\|\xi^{(1)}\|_\ell = 1$, for any vector $x \in \mathbb{R}^n$,

$$\|x\|_\infty \leq \|x\|_\ell \leq \|x\|_1.$$

Lemma 20 (Lemma 3.12 in [7]). Let $\|\cdot\|_\ell$ be a symmetric norm. Then

$$\Omega(M_\ell/\sqrt{\log n}) \leq \|\xi^{(n)}\|_\ell \leq O(M_\ell),$$

where M_ℓ is as defined in Definition 6.

Modulus of Approximation. We need the following quantity of a symmetric norm.

Definition 13. The *maximum modulus of approximation* of a symmetric norm $\|\cdot\|_\ell$ is defined as

$$\text{mma}(\ell, r) = \max_{1 \leq a \leq b \leq ar \leq n} \frac{M_{\ell^{(ar)}}}{M_{\ell^{(b)}}},$$

where $\|\cdot\|_{\ell^{(k)}}$ is a norm on \mathbb{R}^k which is defined to be

$$\|(x_1, x_2, \dots, x_k)\|_{\ell^{(k)}} = \|(x_1, x_2, \dots, x_k, 0, \dots, 0)\|_\ell,$$

and $M_{\ell^{(k)}}$ is as defined in Definition 6.

Intuitively, $\text{mma}(\ell, r)$ characterizes how well the original symmetric norm can be approximated by a lower dimensional induced norm. We show in the following lemma that $\text{mma}(\ell, r) \leq O(\sqrt{r \log n})$ for any symmetric norm.

Lemma 21. For any symmetric norm $\|\cdot\|_\ell$ and $r \in [n]$, $\text{mma}(\ell, r) \leq O(\sqrt{r \log n})$.

Proof. By Lemma 20, for any $i \in [n]$, $\Omega(M_{\ell^{(i)}}/\sqrt{\log n}) \leq \|\xi^{(i)}\|_\ell \leq O(M_{\ell^{(i)}})$. Let $ar = c_1 b + c_2$, where c_1, c_2 are non-negative integers with $c_1 \leq ar/b$ and $c_2 \leq b$. Observe that we can rewrite $\xi^{(ar)}$ as

$$\xi^{(ar)} = \left[\frac{\sqrt{b}}{\sqrt{ar}} \cdot \left(\underbrace{\xi^{(b)}, \xi^{(b)}, \xi^{(b)}, \dots, \xi^{(b)}}_{c_1 \text{ times}} \right), \frac{\sqrt{c_2}}{\sqrt{ar}} \xi^{(c_2)}, 0, \dots, 0 \right].$$

Therefore, by triangle inequality, we have

$$\begin{aligned}
\|\xi^{(ar)}\|_\ell &\leq \frac{\sqrt{b}}{\sqrt{ar}} \cdot c_1 \cdot \|\xi^{(b)}\|_\ell + \frac{\sqrt{c_2}}{\sqrt{ar}} \cdot \|\xi^{(c_2)}\|_\ell \\
&\leq \sqrt{\frac{b}{ar}} \cdot \frac{ar}{b} \cdot \|\xi^{(b)}\|_\ell + \frac{\sqrt{b}}{\sqrt{ar}} \cdot \|\xi^{(b)}\|_\ell && (c_1 \leq ar/b \text{ and } c_2 \leq b) \\
&\leq \sqrt{r} \cdot \|\xi^{(b)}\|_\ell + \|\xi^{(b)}\|_\ell && (a \leq b \leq ar) \\
&\leq 2\sqrt{r} \cdot \|\xi^{(b)}\|_\ell. && (r \geq 1)
\end{aligned}$$

Now we apply Lemma 20 on both sides, which implies

$$\frac{M_{\ell(ar)}}{\sqrt{\log n}} \leq O(\sqrt{r} \cdot M_{\ell(b)})$$

as desired. \square

Properties of SymSketch. Now we introduce several properties of SymSketch.

The following lemma shows that for a data matrix $A \in \mathbb{R}^d$, calculating SA requires $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time for a SymSketch S .

Lemma 22. *For a given matrix $A \in \mathbb{R}^{n \times d}$, let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as in Definition 5. SA can be computed in $O(\text{nnz}(A)) + \text{poly}(d)$ time in expectation, and in $O(\text{nnz}(A) \log n) + \text{poly}(d)$ time in the worst case.*

Proof. Since S is a SymSketch, $S = \Pi \tilde{D} = \Pi \cdot \begin{bmatrix} w_0 D_0 \\ w_1 D_1 \\ \vdots \\ w_t D_t \end{bmatrix}$, where $\Pi \in \mathbb{R}^{O(d) \times O(n \log n)}$.

Since D_i is a diagonal matrix, $\text{nnz}(D_i A) \leq \text{nnz}(A)$, and thus $\text{nnz}(\tilde{D}A) \leq (t+1) \cdot \text{nnz}(A) = O(\text{nnz}(A) \log n)$, which implies $\tilde{D}A$ can be computed in $(t+1) \cdot \text{nnz}(A) = O(\text{nnz}(A) \log n)$ time.

On the other hand, the expected number of nonzero entries of $D_i A$ is $2^{-i} \text{nnz}(A)$. Thus, $\tilde{D}A$ has $O(\text{nnz}(A))$ nonzero entries in expectation, which implies $\tilde{D}A$ can be computed in $O(\text{nnz}(A))$ time.

Finally, notice that Π is a composition of Gaussian embedding and CountSketch embedding, which implies $\Pi \tilde{D}A$ can be computed in $\text{nnz}(\tilde{D}A) + \text{poly}(d)$ time. \square

The following lemma shows that with constant probability, for all $x \in \mathbb{R}^n$, $\|Sx\|_2 \leq \text{poly}(n)\|x\|_2$.

Lemma 23. *Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 5, then with probability at least 0.9999 $\|S\|_2 \leq \text{poly}(n)$.*

Proof. Notice that $S = \Pi \tilde{D}$, since $\|S\|_2 \leq \|\Pi\|_2 \cdot \|\tilde{D}\|_2$, it suffices to bound $\|\Pi\|_2$ and $\|\tilde{D}\|_2$. Since Π is a composition of Gaussian embedding and CountSketch embedding (Definition 12),

with probability at least 0.9999, $\|\Pi\|_2 \leq \|\Pi\|_F \leq \text{poly}(n)$. Now consider $\tilde{D} = \begin{bmatrix} w_0 D_0 \\ w_1 D_1 \\ \vdots \\ w_t D_t \end{bmatrix}$. By

Lemma 19, for all $j \in [t]$, $w_j \leq \text{poly}(n)$. Furthermore, $\|D_j\|_2 \leq 1$ and $t = \Theta(\log n)$, which implies $\|\tilde{D}\|_2 \leq \text{poly}(n)$. \square

Throughout this whole section we assume that for any non-zero vector $x \in \mathbb{R}^n$, we have $1 \leq |x_j| \leq \text{poly}(n)$ for all $j \in [n]$. Notice that this assumption is without loss of generality, as shown in the following lemma.

Lemma 24. For any non-zero vector $x \in \mathbb{R}^n$, let $\bar{x} \in \mathbb{R}^n$ be a vector with $\bar{x} = \frac{\text{poly}(n) \cdot x}{\|x\|_\infty}$, and $x' \in \mathbb{R}^n$ where

$$x'_i = \begin{cases} \bar{x}_i & \text{if } \bar{x}_i \geq 1 \\ 0 & \text{otherwise} \end{cases}.$$

For a symmetric norm $\|\cdot\|_\ell$, suppose $\|S\|_2 \leq \text{poly}(n)$ and

$$\alpha \|x'\|_\ell \leq \|Sx'\|_2 \leq \beta \|x'\|_\ell$$

for some $\alpha, \beta \in [1/\text{poly}(n), \text{poly}(n)]$, then

$$\Omega(\alpha) \|x\|_\ell \leq \|Sx\|_2 \leq O(\beta) \|x\|_\ell.$$

Proof. By triangle inequality and Lemma 18, we have

$$\|\bar{x}\|_\ell - n \leq \|x'\|_\ell \leq \|\bar{x}\|_\ell.$$

By Lemma 19,

$$\|x'\|_\ell \geq \|x'\|_\infty = \|\bar{x}\|_\infty = \text{poly}(n),$$

we have

$$(1 - 1/\text{poly}(n)) \|\bar{x}\|_\ell \leq \|x'\|_\ell \leq \|\bar{x}\|_\ell.$$

Notice that $\|S\bar{x}\|_2 = \|Sx' + S(\bar{x} - x')\|_2$. By triangle inequality we have

$$\|Sx'\|_2 - \|S\|_2 \|\bar{x} - x'\|_2 \leq \|S\bar{x}\|_2 \leq \|Sx'\|_2 + \|S\|_2 \|\bar{x} - x'\|_2.$$

By the given conditions, we have

$$(1 - 1/\text{poly}(n)) \|Sx'\|_2 \leq \|S\bar{x}\|_2 \leq (1 + 1/\text{poly}(n)) \|Sx'\|_2,$$

which implies

$$\Omega(\alpha) \|\bar{x}\|_\ell \leq \|S\bar{x}\|_2 \leq O(\beta) \|\bar{x}\|_\ell.$$

Since $\bar{x} = \frac{\text{poly}(n) \cdot x}{\|x\|_\infty}$, we have

$$\Omega(\alpha) \|x\|_\ell \leq \|Sx\|_2 \leq O(\beta) \|x\|_\ell.$$

□

By Lemma 24, we can focus on those non-zero vectors $x \in \mathbb{R}^n$ such that $1 \leq |x_j| \leq \text{poly}(n)$ for all $j \in [n]$.

Definition 14. For a given vector $x \in \mathbb{R}^n$, suppose for all $j \in [n]$,

$$1 \leq |x_j| \leq \text{poly}(n).$$

Let $g = \Theta(\log n)$. For each $i \in \{0, 1, \dots, g\}$, we define

$$L_i(x) = \{j \mid 2^i \leq |x_j| < 2^{i+1}\}.$$

For each $i \in \{0, 1, \dots, g\}$, we define $V_i(x) \in \mathbb{R}^n$ to be the vector

$$V_i(x) = \underbrace{(2^i, 2^i, \dots, 2^i)}_{|L_i(x)|}, 0, \dots, 0).$$

For each $i \in \{0, 1, \dots, g\}$, we say a level i to be *contributing* if

$$\|V_i(x)\|_\ell \geq \Omega(1/g) \cdot \|x\|_\ell.$$

Lemma 25. Let $g = \Theta(\log n)$. For a given vector $x \in \mathbb{R}^n$ such that for all $j \in [n]$, $1 \leq |x_j| \leq 2^g$, there exists at least one level $i \in \{0, 1, \dots, g\}$ which is contributing.

Proof. If none of $i \in \{0, 1, \dots, g\}$ is contributing, then $\|x\|_\ell \leq \sum_{i=0}^g \|V_i(x)\|_\ell \leq 1/(2g) \cdot \sum_{i=0}^g \|x\|_\ell \leq \frac{1}{2} \|x\|_\ell$, which leads to a contradiction. □

C.2 Proof of Lemma 9

Proof. Consider a fixed $n' \in [n]$. By Lemma 20, we have

$$M_{\ell_a^{(n')}} = \Omega \left(\|\xi^{(n')}\|_{\ell_a^{(n')}} \right) = \Omega \left(1 + c\sqrt{n'} \right)$$

and

$$M_{\ell_b^{(n')}} = \Omega \left(\|\xi^{(n')}\|_{\ell_b^{(n')}} \right) = \Omega \left(\max \left(1, c\sqrt{n'} \right) \right).$$

It is also straightforward to verify that

$$\max_{x \in \mathbb{S}^{n'-1}} \|x\|_{\ell_a} = 1 + c\sqrt{n'}$$

and

$$\max_{x \in \mathbb{S}^{n'-1}} \|x\|_{\ell_b} = \max \left(1, c\sqrt{n'} \right).$$

Taking the ratio between $\max_{x \in \mathbb{S}^{n'-1}} \|x\|_{\ell}$ and $M_{\ell^{(n')}}$ for $\ell \in \{\ell_a, \ell_b\}$, we complete the proof. \square

C.3 Proof of Lemma 10

Proof. Let $\bar{G}(x) = \sqrt{x} \cdot G^{-1}(1/x)$, where $G^{-1}(1/x)$ is the unique value in $[0, \infty)$ such that $G(G^{-1}(1/x)) = 1/x$. We first show that $\bar{G}(x)$ is an approximately decreasing function for $x \in (0, \infty)$. Let m, n be two real numbers with $0 < m \leq n$. We have $1/n \leq 1/m$, which implies $0 < G^{-1}(1/n) \leq G^{-1}(1/m)$ by monotonicity of G . By the third condition in Assumption 1, we have

$$\frac{G(G^{-1}(1/m))}{G(G^{-1}(1/n))} \leq C_G \cdot \left(\frac{G^{-1}(1/m)}{G^{-1}(1/n)} \right)^2,$$

which implies

$$\sqrt{n} \cdot G^{-1}(1/n) \leq \sqrt{C_G} \cdot \sqrt{m} \cdot G^{-1}(1/m).$$

Hence $\bar{G}(n) \leq \sqrt{C_G} \cdot \bar{G}(m)$.

We are now ready to prove the lemma. Recall that for the Orlicz norm $\|\cdot\|_{\ell} = \|\cdot\|_G$, we have

$$\text{mmc}(\ell) = \max_{n' \in [n]} \text{mc}(\ell^{(n')}) = \max_{n' \in [n]} \frac{\max_{x \in \mathbb{S}^{n'-1}} \|x\|_{\ell^{(n')}}}{M_{\ell^{(n')}}}.$$

By Lemma 20, we have,

$$\Omega(1) \cdot \|\xi^{(n')}\|_{\ell^{(n')}} \leq M_{\ell^{(n')}} \leq O(\sqrt{\log n}) \cdot \|\xi^{(n')}\|_{\ell^{(n')}}.$$

Thus $\|\xi^{(n')}\|_{\ell^{(n')}}$ provides an approximation to $M_{\ell^{(n')}}$. By definition of $\|\cdot\|_G$, we have

$$\|\xi^{(n')}\|_{\ell^{(n')}} = \frac{1}{\sqrt{n'} \cdot G^{-1}(1/n')} = \frac{1}{\bar{G}(n')}.$$

Hence

$$\Omega(1) \leq M_{\ell^{(n')}} \cdot \bar{G}(n') \leq O(\sqrt{\log n}).$$

Next, we compute $\max_{x \in \mathbb{S}^{n'-1}} \|x\|_{\ell^{(n')}}$. For an arbitrary unit vector $x \in \mathbb{S}^{n'-1}$, we denote

$$B_j = \{i \in [n] : |x_i| \in [1/2^j, 1/2^{j-1}]\}$$

and $b_j = |B_j|$. For each j , let $x^{B_j} \in \mathbb{R}^n$ be the vector such that

$$x_i^{B_j} = \begin{cases} x_i & \text{if } i \in B_j \\ 0 & \text{otherwise} \end{cases}.$$

Note that non-zero coordinates in x^{B_j} have magnitude close to each other (within a factor of 2), we thus have

$$\|x^{B_j}\|_{\ell} = \|x^{B_j}\|_2 \cdot \left\| \frac{x^{B_j}}{\|x^{B_j}\|_2} \right\|_{\ell} \leq \frac{\sqrt{b_j}}{2^{j-1}} \cdot \left\| \frac{x^{B_j}}{\|x^{B_j}\|_2} \right\|_{\ell} \leq \frac{\sqrt{b_j}}{2^{j-2}} \cdot \|\xi^{(b_j)}\|_{\ell^{(b_j)}} = \frac{\sqrt{b_j}}{2^{j-2}} \cdot \frac{1}{\bar{G}(b_j)}.$$

Similarly,

$$\|x^{B_j}\|_\ell \geq \frac{\sqrt{b_j}}{2^{j+2}} \cdot \|\xi^{(b_j)}\|_{\ell^{(b_j)}} = \frac{\sqrt{b_j}}{2^{j+2}} \cdot \frac{1}{\overline{G}(b_j)} \geq \frac{\sqrt{b_j}}{2^{j+2}} \cdot \frac{1}{\sqrt{C_G} \cdot \overline{G}(1)}.$$

We claim there exists an constant $c > 0$ such that

$$\sum_{\substack{j > c \log n \\ b_j > 0}} \|x^{B_j}\|_\ell \leq \sum_{j' \leq c \log n} \|x^{B_{j'}}\|_\ell.$$

To show this, by Lemma 14, for any $b \geq 1$,

$$b = \frac{G(G^{-1}(1))}{G(G^{-1}(1/b))} \geq \frac{G^{-1}(1)}{G^{-1}(1/b)},$$

which implies

$$G^{-1}(1/b) \geq \frac{G^{-1}(1)}{b}.$$

Next, since $\|x\|_2 = 1$, there exists an $0 \leq \tilde{j} \leq 4 \log n$ such that $b_{\tilde{j}} \geq 1$. Therefore,

$$\|x^{B_{\tilde{j}}}\|_\ell \geq \frac{1}{2^{\tilde{j}+2}} \cdot \frac{1}{\sqrt{C_G} \cdot \overline{G}(1)}.$$

Thus, we have

$$\begin{aligned} \sum_{\substack{j > c \log n \\ b_j > 0}} \|x^{B_j}\|_\ell &= \sum_{\substack{j > c \log n \\ b_j > 0}} \|x^{B_j}\|_2 \cdot \left\| \frac{x^{B_j}}{\|x^{B_j}\|_2} \right\|_\ell \leq \sum_{\substack{j > c \log n \\ b_j > 0}} \frac{\sqrt{n}}{2^{j-2}} \cdot \|\xi^{(b_j)}\|_{\ell^{(b_j)}} \\ &\leq \sum_{\substack{j > c \log n \\ b_j > 0}} \frac{\sqrt{n}}{2^{j-2}} \cdot \frac{1}{\sqrt{b_j} G^{-1}(1/b_j)} \leq n \cdot \frac{\sqrt{n}}{2^{c \log n - 2}} \cdot \frac{b_j}{\sqrt{b_j} G^{-1}(1)} \\ &\leq n \cdot 2^{\tilde{j}+2} \cdot \frac{\sqrt{n}}{2^{c \log n - 2}} \cdot \sqrt{C_G n} \cdot \|x^{B_{\tilde{j}}}\|_\ell \leq \|x^{B_{\tilde{j}}}\|_\ell \leq \sum_{j' \leq c \log n} \|x^{B_{j'}}\|_\ell \end{aligned}$$

for some sufficiently large constant c .

Let

$$j^* = \operatorname{argmax}_{j \leq c \log n} \|x^{B_j}\|_\ell,$$

we have

$$\|x^{B_{j^*}}\|_\ell \leq \|x\|_\ell \leq \sum_{j \leq c \log n} \|x^{B_j}\|_\ell + \sum_{\substack{j > c \log n \\ b_j > 0}} \|x^{B_j}\|_\ell \leq 2 \sum_{j \leq c \log n} \|x^{B_j}\|_\ell \leq O(\log n) \cdot \|x^{B_{j^*}}\|_\ell.$$

Thus,

$$\begin{aligned} \max_{x \in \mathbb{S}^{n'-1}} \|x\|_{\ell^{(n')}} &\leq O(\log n') \max_{x \in \mathbb{S}^{n'-1}} \|x^{B_{j^*}}\|_\ell \leq O(\log n') \max_{x \in \mathbb{S}^{n'-1}} \|x^{B_{j^*}}\|_2 \cdot \left\| \frac{x^{B_{j^*}}}{\|x^{B_{j^*}}\|_2} \right\|_\ell \\ &\leq O(\log n') \max_{x \in \mathbb{S}^{n'-1}} \left\| \frac{x^{B_{j^*}}}{\|x^{B_{j^*}}\|_2} \right\|_\ell \leq O(\log n') \max_{b_{j^*} \leq n'} \|\xi^{(b_{j^*})}\|_{\ell^{(b_{j^*})}} \leq O(\log n') \max_{b_{j^*} \leq n'} \frac{1}{\overline{G}(b_{j^*})} \leq \frac{O(\sqrt{C_G} \log n')}{\overline{G}(n')}. \end{aligned}$$

Thus, we have

$$\operatorname{mmc}(\ell) = \max_{n' \in [n]} \frac{\max_{x \in \mathbb{S}^{n'-1}} \|x\|_{\ell^{(n')}}}{M_{\ell^{(n')}}} \leq O(\sqrt{C_G} \log n).$$

□

C.4 Contraction Bound of SymSketch

In this section we give the contraction bound of SymSketch. We first show that for a fixed vector $x \in \mathbb{R}^n$, $\|\tilde{D}x\|_2 \geq 1/\text{poly}(d \log n) \cdot \|x\|_\ell$ with probability $1 - 2^{-\Theta(d \log n)}$.

Lemma 26. *Let \tilde{D} be the matrix defined in Definition 5. For any fixed $x \in \mathbb{R}^n$, with probability $1 - 2^{-\Theta(d \log n)}$, $\|\tilde{D}x\|_2 \geq 1/\alpha_0 \cdot \|x\|_\ell$, where $\alpha_0 = O(\text{mma}(\ell, d) \cdot \log^{5/2} n) = O(\sqrt{d} \log^3 n)$.*

Proof. The lemma follows from the following two claims. Recall that $t = \Theta(\log n)$.

Claim 1. *For any fixed $x \in \mathbb{R}^n$. If there is a contributing level $i^* \in \{0, 1, 2, \dots, g\}$ such that $|L_{i^*}(x)| = \Theta(d \log n) \cdot 2^j$ for some $j \in [t]$, then with probability at least $1 - 2^{-\Theta(d \log n)}$,*

$$\|w_j D_j x\|_2 \geq \Omega\left(\frac{1}{\text{mma}(\ell, d) \cdot \log^{5/2} n}\right) \|x\|_\ell.$$

Proof. Let y_h be a random variable such that

$$y_h = \begin{cases} 1 & \text{if the } h\text{-th diagonal entry of } D_j \text{ is 1} \\ 0 & \text{otherwise} \end{cases}.$$

Let $Y = \sum_{h \in L_{i^*}(x)} y_h$. By Chernoff bound, we have

$$\Pr[Y \geq \Omega(d \log n)] \geq 1 - 2^{-\Theta(d \log n)}.$$

Conditioned on $Y \geq \Omega(d \log n)$, we have

$$\begin{aligned} \frac{\|w_j D_j x\|_2}{\|x\|_\ell} &= \frac{w_j \|D_j x\|_2}{\|x\|_\ell} \\ &\geq \frac{2^{i^*} w_j \sqrt{d \log n}}{\|x\|_\ell} \\ &\geq \frac{2^{i^*} \sqrt{2^j} M_{\ell(2^j)} \sqrt{d \log n}}{2^{i^*+1} g M_{\ell(|L_{i^*}(x)|)} \sqrt{\log n} \sqrt{|L_{i^*}(x)|}} \\ &\geq \Omega(1/\log^{3/2} n) \cdot \frac{M_{\ell(2^j)}}{M_{\ell(|L_{i^*}(x)|)}} \\ &= \Omega(1/\log^{3/2} n) \cdot \frac{M_{\ell(2^j)}}{M_{\ell(|L_{i^*}(x)|/\log n)}} \cdot \frac{M_{\ell(|L_{i^*}(x)|/\log n)}}{M_{\ell(|L_{i^*}(x)|)}} \\ &\geq \Omega\left(\frac{1}{\text{mma}(\ell, d) \cdot \text{mma}(\ell, \log n) \cdot \log^{3/2} n}\right) \\ &\geq \Omega\left(\frac{1}{\text{mma}(\ell, d) \cdot \log^{5/2} n}\right). \end{aligned}$$

Here the first inequality follows from the fact that there are at least $\Omega(d \log n)$ coordinates sampled from $L_{i^*}(x)$. The second inequality follows from Lemma 20 and the fact that level i^* is a contributing level. The third inequality follows from $|L_{i^*}(x)| = \Theta(d \log n) \cdot 2^j$ and $g = \Theta(\log n)$. The fourth inequality follows from Definition 13. The last inequality follows from Lemma 21. \square

Claim 2. *For any fixed $x \in \mathbb{R}^n$. If there is a contributing level $i^* \in \{0, 1, 2, \dots, g\}$ such that $|L_{i^*}(x)| = O(d \log n)$, then we have*

$$\|w_0 D_0 x\|_2 \geq \Omega\left(\frac{1}{\text{mma}(\ell, d) \cdot \log^{5/2} n}\right) \|x\|_\ell.$$

Proof.

$$\begin{aligned}
\frac{\|w_0 D_0 x\|_2}{\|x\|_\ell} &= \frac{w_0 \|x\|_2}{\|x\|_\ell} \\
&\geq \frac{2^{i^*} w_0 \sqrt{|L_{i^*}(x)|}}{\|x\|_\ell} \\
&\geq \frac{2^{i^*} M_{\ell(1)} \sqrt{|L_{i^*}(x)|}}{2^{i^*+1} g M_{\ell(|L_{i^*}(x)|)} \sqrt{\log n} \sqrt{|L_{i^*}(x)|}} \\
&\geq \Omega(1/\log^{3/2} n) \cdot \frac{M_{\ell(1)}}{M_{\ell(|L_{i^*}(x)|/\log n)}} \cdot \frac{M_{\ell(|L_{i^*}(x)|/\log n)}}{M_{\ell(|L_{i^*}(x)|)}} \\
&\geq \Omega\left(\frac{1}{\text{mma}(\ell, d) \cdot \log^{5/2} n}\right)
\end{aligned}$$

The first inequality follows from the fact that we only consider the contribution of the coordinates in $L_{i^*}(x)$. The second inequality follows from Lemma 20 and the fact that level i^* is a contributing level. The third inequality follows from $g = \Theta(\log n)$. The last inequality follows from Definition 13 and Lemma 21. \square

By Claim 1 and Claim 2, since any vector $x \in \mathbb{R}^n$ contains at least one contributing level, with probability at least $1 - 2^{-\Theta(d \log n)}$ we have $\|\tilde{D}x\|_2 \geq \Omega(1/(\text{mma}(\ell, d) \cdot \log^{5/2} n)) \cdot \|x\|_\ell$. We complete the proof by combining this with Lemma 21. \square

Now we show how to combine the contraction bound in Lemma 26 with a net argument to give a contraction bound for all vectors in a subspace.

Lemma 27. *Let $S \in \mathbb{R}^{O(d) \times n}$ be a random matrix. For any $\alpha_0 = \text{poly}(n)$ and $A \in \mathbb{R}^{n \times d}$, if*

1. $\|S\|_2 \leq \text{poly}(n)$ holds with probability at least 0.999;
2. for any fixed $x \in \mathbb{R}^n$, $\|Sx\|_2 \geq 1/\alpha_0 \cdot \|x\|_\ell$ holds with probability $1 - e^{-Cd \log n}$ for a sufficiently large constant C ,

then with probability at least 0.998, for all $y \in \mathbb{R}^n$ in the column space of A ,

$$\|Sy\|_2 \geq \Omega(1/\alpha_0) \|y\|_\ell.$$

Proof. For the matrix $A \in \mathbb{R}^{n \times d}$, we define the set $\mathcal{B} = \{y \mid y = Ax, \|y\|_2 = 1\}$. We define $\mathcal{N} \subset \mathbb{R}^n$ to be an ε -net of \mathcal{B} as in Definition 8. By Lemma 13, we have $|\mathcal{N}| \leq (1 + 1/\varepsilon)^d$, and for all $y \in \mathcal{B}$, there exists $z \in \mathcal{N}$ such that $\|y - z\|_2 \leq \varepsilon$. We take $\varepsilon = 1/\text{poly}(n)$ here.

Due to the second condition, since $|\mathcal{N}| \leq e^{O(d \log n)}$, by taking union bound over all vectors in \mathcal{N} , we know that with probability $1 - e^{-\Theta(d \log n)}$, for all $z \in \mathcal{N}$, $\|Sz\|_2 \geq 1/\alpha_0 \cdot \|z\|_\ell$.

Now, for any vector $y \in \mathcal{B}$, there exists $z \in \mathcal{N}$ such that $\|y - z\|_2 \leq 1/\text{poly}(n)$, and we define $w = y - z$.

$$\begin{aligned}
\|Sy\|_2 &= \|S(z + w)\|_2 \\
&\geq \|Sz\|_2 - \|Sw\|_2 && \text{(triangle inequality)} \\
&\geq 1/\alpha_0 \cdot \|z\|_\ell - \|Sw\|_2 && \text{(by the second condition)} \\
&\geq 1/\alpha_0 \cdot \|z\|_\ell - \|S\|_2 \|w\|_2 && (\|Sw\|_2 \leq \|S\|_2 \cdot \|w\|_2) \\
&\geq 1/\alpha_0 \cdot \|z\|_\ell - \text{poly}(n) \cdot \|w\|_2 && \text{(by the first condition)} \\
&\geq 1/\alpha_0 \cdot \|y - w\|_\ell - \text{poly}(n) \cdot \|w\|_2 && (y = z + w) \\
&\geq 1/\alpha_0 \cdot \|y\|_\ell - 1/\alpha_0 \cdot \|w\|_\ell - \text{poly}(n) \cdot \|w\|_2 && \text{(triangle inequality)} \\
&\geq 1/\alpha_0 \cdot \|y\|_\ell - 1/\alpha_0 \cdot \sqrt{n} \|w\|_2 - \text{poly}(n) \cdot \|w\|_2 && \text{(Lemma 19)} \\
&\geq 1/\alpha_0 \cdot \|y\|_\ell - (1/\alpha_0 \cdot \sqrt{n} + \text{poly}(n)) \varepsilon && (\|w\|_2 \leq \varepsilon) \\
&\geq 0.5/\alpha_0 \cdot \|y\|_\ell.
\end{aligned}$$

\square

Lemma 28. For a given matrix $A \in \mathbb{R}^{n \times d}$. Let $S \in \mathbb{R}^{O(d) \times n}$ be a **SymSketch** as defined in Definition 5. With probability at least 0.995, for all $x \in \mathbb{R}^d$, $\|SAx\|_2 \geq 1/\alpha_0 \cdot \|Ax\|_\ell$ where $\alpha_0 = O(\sqrt{d} \log^3 n)$.

Proof. By Lemma 26 and Lemma 23, the two conditions in Lemma 27 are satisfied. By Lemma 27, with probability at least 0.998, for all $x \in \mathbb{R}^d$, $\|\tilde{D}Ax\|_2 \geq \Omega(1/\alpha_0)\|Ax\|_\ell$. Since $\Pi \in \mathbb{R}^{O(d) \times n(t+1)}$ is a composition of Gaussian embedding and CountSketch embedding with $\varepsilon = 0.1$, by Corollary 17, with probability at least 0.999, for all $x \in \mathbb{R}^d$, $\|\Pi\tilde{D}Ax\|_2 \geq \Omega(\|\tilde{D}Ax\|_2)$. By a union bound, we know that with probability at least 0.995, for all $x \in \mathbb{R}^d$, $\|SAx\|_2 \geq \Omega(1/\alpha_0)\|Ax\|_\ell$. \square

C.5 Dilation Bound of SymSketch

In this section we give the dilation bound of **SymSketch**. We first show that for any fixed $x \in \mathbb{R}^n$, with high probability, $\|\tilde{D}x\|_2 \leq \text{poly}(d \log n) \cdot \text{mmc}(\ell) \cdot \|x\|_\ell$.

Lemma 29. Let \tilde{D} be the matrix defined in Definition 5. For any fixed vector $x \in \mathbb{R}^n$, with probability $1 - \delta$, $\|\tilde{D}x\|_2 \leq \alpha_1/\delta \cdot \|x\|_\ell$, where $\alpha_1 = O(\text{mmc}(\ell) \log^{5/2} n)$.

Proof. Consider a fixed vector $x \in \mathbb{R}^n$. Recall that $t = \Theta(\log n)$. Let $c > 0$ be a fixed constant. We define the j -heavy level set H_j as

$$H_j = \left\{ i \mid |L_i(x)| \geq c \frac{\delta 2^j}{\log^2 n}, 0 \leq i \leq g \right\}.$$

Let \bar{H}_j be the j -light level set, i.e., $\bar{H}_j = \{0, 1, \dots, g\} \setminus H_j$. Notice that

$$\sum_{i \in \bar{H}_j} |L_i(x)| \cdot 2^{-j} \leq g \cdot c \delta 2^j / \log^2 n \cdot 2^{-j} \leq O(\delta / \log n).$$

By Markov's inequality, with probability at least $1 - \delta/(2t)$, no element from a j -light level is sampled by D_j , i.e., for all $i \in \bar{H}_j$, $k \in L_i(x)$, the k -th diagonal entry of D_j is 0. By taking union bound over all $j \in [t]$, with probability at least $1 - \delta/2$, for all $j \in [t]$, no element from a j -light level is sampled by D_j . Let ζ denote this event. We condition on this event in the remaining part of the proof. In the following analysis, we show an upper bound of $\|w_j D_j x\|_2^2$ for each $j \in [t]$. Let H_j be the set of j -heavy levels.

Consider a fixed $j \in [t]$, we have

$$\begin{aligned} \mathbf{E}_{D_j} \left[\|w_j D_j x\|_2^2 \mid \zeta \right] &= w_j^2 \mathbf{E}_{D_j} \left[\|D_j x\|_2^2 \mid \zeta \right] \\ &= w_j^2 \mathbf{E}_{D_j} \left[\sum_{h=1}^n (D_j(h, h))^2 x_h^2 \mid \zeta \right] \\ &= w_j^2 \mathbf{E}_{D_j} \left[\sum_{i=0}^g \sum_{h \in L_i(x)} (D_j(h, h))^2 x_h^2 \mid \zeta \right] \\ &= w_j^2 \frac{1}{2^j} \sum_{i \in H_j} \sum_{h \in L_i(x)} x_h^2 \\ &\leq w_j^2 \frac{1}{2^j} \sum_{i \in H_j} |L_i(x)| \cdot (2^{i+1})^2. \end{aligned}$$

Claim 3. $w_j^2 2^{-j} \leq O((M_{\ell(2^j)})^2)$.

Proof.

$$\begin{aligned}
w_j^2 2^{-j} &= (\|(1, 1, \dots, 1, 0, \dots, 0)\|_\ell)^2 \cdot 2^{-j} \\
&= \left(\left\| \frac{1}{\sqrt{2^j}} (1, 1, \dots, 1, 0, \dots, 0) \right\|_\ell \sqrt{2^j} \right)^2 \cdot 2^{-j} \\
&= (\|\xi^{(2^j)}\|_\ell)^2 \cdot 2^j \cdot 2^{-j} \\
&= (\|\xi^{(2^j)}\|_\ell)^2 \\
&\leq O(M_{\ell(2^j)}),
\end{aligned}$$

where the third step follows from the definition of $\xi^{(2^j)}$, and the last step follows from Lemma 20. \square

Using the above claim, we have

$$\begin{aligned}
&w_j^2 \sum_{i \in H_j} |L_i(x)| \cdot 2^{2i-j} \\
&= \sum_{i \in H_j} \frac{w_j^2 2^{-j}}{(M_{\ell(|L_i(x)|)})^2} \cdot |L_i(x)| \cdot (M_{\ell(|L_i(x)|)})^2 \cdot 2^{2i} \\
&\leq O \left(\sum_{i \in H_j} \left(\frac{M_{\ell(2^j)}}{M_{\ell(|L_i(x)|)}} \right)^2 \cdot |L_i(x)| \cdot (M_{\ell(|L_i(x)|)})^2 \cdot 2^{2i} \right) \\
&\leq O \left(\sum_{i \in H_j} \left(\frac{M_{\ell(2^j)}}{M_{\ell(|L_i(x)|)}} \right)^2 w_{\log |L_i(x)|}^2 \log n \cdot 2^{2i} \right) \\
&= \log n \cdot \underbrace{\sum_{i \in H_j, |L_i(x)| \leq 2^j} \left(\frac{M_{\ell(2^j)}}{M_{\ell(|L_i(x)|)}} \right)^2 \cdot w_{\log |L_i(x)|}^2 \cdot 2^{2i}}_{\diamond} \\
&\quad + \log n \cdot \underbrace{\sum_{i \in H_j, |L_i(x)| > 2^j} \left(\frac{M_{\ell(2^j)}}{M_{\ell(|L_i(x)|)}} \right)^2 \cdot w_{\log |L_i(x)|}^2 \cdot 2^{2i}}_{\heartsuit},
\end{aligned}$$

where the second step follows from $w_j^2 2^{-j} \leq O((M_{\ell(2^j)})^2)$ (Claim 3), and the third step follows from $|L_i(x)| \cdot (M_{\ell(|L_i(x)|)})^2 \leq O(w_{\log |L_i(x)|}^2 \log n)$ (Lemma 20). It remains to upper bound \diamond and \heartsuit .

To given an upper bound for \diamond , we have

$$\begin{aligned}
\diamond &\leq O \left(\sum_{i \in H_j, |L_i(x)| \leq 2^j} \text{mma}^2(\ell, \log^2 n/\delta) \cdot w_{\log |L_i(x)|}^2 \cdot 2^{2i} \right) \\
&\leq O \left(\text{mma}^2(\ell, \log^2 n/\delta) \left(\sum_{i=0}^g w_{\log |L_i(x)|} \cdot 2^i \right)^2 \right) \\
&\leq O(\text{mma}^2(\ell, \log^2 n/\delta) \|x\|_\ell^2) \\
&\leq O(\log^3 n/\delta) \|x\|_\ell^2,
\end{aligned}$$

where the first step follows from the definition of mma , the second step follows from Minkowski inequality, the third step follows from the definition of $L_i(x)$, $w_{\log |L_i(x)|}$ and triangle inequality, the last step follows from Lemma 21.

To give an upper bound for \heartsuit , we have

$$\begin{aligned}\heartsuit &\leq O\left(\log n \cdot \sum_{i \in H_j, |L_i(x)| > 2^j} \text{mmc}^2(\ell) w_{\log |L_i(x)|}^2 \cdot 2^{2i}\right) \\ &\leq O\left(\log n \cdot \text{mmc}^2(\ell) \cdot \left(\sum_{i=0}^g w_{\log |L_i(x)|} \cdot 2^i\right)^2\right) \\ &\leq O\left(\log n \cdot \text{mmc}^2(\ell) \cdot \|x\|_\ell^2\right),\end{aligned}$$

where the first step follows from $(M_{\ell(2^j)}/M_{\ell(|L_i(x)|)})^2 \leq O(\log n \cdot \text{mmc}^2(\ell))$ (Lemma 3.14 in [7]).

Putting it all together, we have

$$\mathbf{E}_{D_j}[\|w_j D_j x\|_2^2 | \zeta] \leq \log n \cdot (\diamond + \heartsuit) \leq O(\log^4 n / \delta + \log^2 n \cdot \text{mmc}^2(\ell)) \|x\|_\ell^2.$$

Thus,

$$\mathbf{E}_{\tilde{D}}[\|\tilde{D}x\|_2^2 | \zeta] \leq \sum_{j=0}^t \mathbf{E}_{D_j}[\|w_j D_j x\|_2^2 | \zeta] \leq O(\log^5 n / \delta + \log^3 n \cdot \text{mmc}^2(\ell)) \|x\|_\ell^2.$$

By Markov's inequality, conditioned on ζ , with probability at least $1 - \delta/2$,

$$\|\tilde{D}x\|_2^2 \leq O(\log^5 n / \delta + \log^3 n \cdot \text{mmc}^2(\ell)) \|x\|_\ell^2 / \delta.$$

Since ζ holds with probability at least $1 - \delta/2$, with probability at least $1 - \delta$, we have

$$\|\tilde{D}x\|_2 \leq O(\log^{5/2} n / \delta \cdot \text{mmc}(\ell)) \cdot \|x\|_\ell.$$

□

Now we show how to use the dilation bound for a fixed vector in Lemma 29 to prove a dilation bound for all vectors in a subspace. We need the following existential result in our proof.

Lemma 30 ([6]). *Given a matrix $A \in \mathbb{R}^{n \times m}$ and a norm $\|\cdot\|$, there exists a basis matrix $U \in \mathbb{R}^{n \times d}$ of the column space of A , such that*

$$\sum_{i=1}^d \|U^i\| \leq d,$$

and for all $x \in \mathbb{R}^d$,

$$\|x\|_\infty \leq \|Ux\|.$$

Lemma 31. *Given a matrix $A \in \mathbb{R}^{n \times d}$. Let $S \in \mathbb{R}^{O(d) \times n}$ be a **SymSketch** as defined in Definition 5. With probability at least 0.99, for all $x \in \mathbb{R}^d$,*

$$\|SAx\|_2 \leq O(\alpha_1 d^2) \|Ax\|_\ell,$$

where $\alpha_1 = O(\text{mmc}(\ell) \cdot \log^{5/2} n)$.

Proof. Recall that $S = \Pi \tilde{D}$. Let U be a basis matrix of the column space of A as in Lemma 30. By Lemma 29, for a fixed $i \in [d]$, with probability at least $1 - 1/(100d)$, $\|\tilde{D}U^i\|_2 \leq O(\alpha_1 d) \|U^i\|_\ell$. By taking a union bound over $i \in [d]$, with probability at least 0.999, for all $i \in [d]$, $\|\tilde{D}U^i\|_2 \leq$

$\alpha_1 d \|U^i\|_\ell$. Thus, for any $x \in \mathbb{R}^d$,

$$\begin{aligned}
\|\tilde{D}Ux\|_2 &\leq \sum_{i=1}^d |x_i| \cdot \|\tilde{D}U^i\|_2 \\
&\leq \|x\|_\infty \cdot \sum_{i=1}^d \|\tilde{D}U^i\|_2 \\
&\leq \|Ux\|_\ell \cdot \sum_{i=1}^d \|\tilde{D}U^i\|_2 \\
&\leq O(\alpha_1 d) \cdot \|Ux\|_\ell \cdot \sum_{i=1}^d \|U^i\|_\ell \\
&\leq O(\alpha_1 d^2) \cdot \|Ux\|_\ell,
\end{aligned}$$

where the first step follows from triangle inequality, the second step follows from $|x_i| \leq \|x\|_\infty$ for all $i \in [d]$, the third step follows from $\|x\|_\infty \leq \|Ux\|_\ell$, the fourth step follows from $\|\tilde{D}U^i\|_2 \leq O(\alpha_1 d) \|U^i\|_\ell$, the last step follows from $\sum_{i=1}^d \|U^i\|_\ell \leq d$.

By Corollary 17, with probability at least 0.999, Π is an ℓ_2 subspace embedding with $\varepsilon = 0.1$ for the column space of $\tilde{D}U$. Thus, with probability at least 0.99, for all $x \in \mathbb{R}^d$, $\|SAx\|_2 \leq O(\alpha_1 d^2) \|Ax\|_\ell$. \square

C.6 Proof of Theorem 11

Proof. It directly follows from Lemma 28, Lemma 31 and Lemma 22. \square

D Missing Proofs of Main Theorems

D.1 Proof of Theorem 1

Let $S \in \mathbb{R}^{O(d) \times n}$ be a **SymSketch** as defined in Definition 5, and $\Pi = O(\sqrt{d} \log^3 n) \cdot S$. By Corollary 12, for a given matrix $A \in \mathbb{R}^{n \times d}$, with probability at least 0.9, for all $x \in \mathbb{R}^d$,

$$\|Ax\|_G \leq \|\Pi Ax\|_2 \leq \kappa \|Ax\|_G,$$

where $\kappa = O(\sqrt{C_G} d^{5/2} \log^{13/2} n)$. We prove Theorem 1 by combining Theorem 8 with the embedding matrix Π constructed above.

D.2 Proof of Theorem 2

Let $S \in \mathbb{R}^{O(d) \times n}$ be a **SymSketch** as defined in Definition 5. For a given data matrix $A \in \mathbb{R}^{n \times d}$ and response vector $b \in \mathbb{R}^n$, we calculate $x^* = \operatorname{argmin}_x \|SAx - Sb\|_2$ and return x^* . The algorithm runs in $O(\operatorname{nnz}(A) + \operatorname{poly}(d))$ time, since by Lemma 22, the expected running time for calculating SA is $O(\operatorname{nnz}(A) + \operatorname{poly}(d))$, and $x^* = (SA)^+ Sb$ can be calculated in $\operatorname{poly}(d)$ time.

To see the correctness, let $\bar{x} = \operatorname{argmin}_x \|Ax - b\|_\ell$. With probability at least 0.99, we have

$$\begin{aligned}
\|Ax^* - b\|_\ell &\leq O(\sqrt{d} \log^3 n) \|SAx^* - Sb\|_2 \\
&\leq O(\sqrt{d} \log^3 n) \|SA\bar{x} - Sb\|_2 \\
&\leq O(\sqrt{d} \log^3 n) \|\tilde{D}A\bar{x} - \tilde{D}b\|_\ell \\
&= O(\sqrt{d} \log^{11/2} n) \cdot \operatorname{mmc}(\ell) \cdot \|A\bar{x} - b\|_\ell.
\end{aligned}$$

The first step follows by applying Lemma 28 on \bar{A} , where we use $\bar{A} \in \mathbb{R}^{n \times (d+1)}$ to denote a matrix whose first d columns are A and the last column is b . The second step follows from the fact that $x^* = \operatorname{argmin}_x \|SAx - Sb\|_2$. The third step follows by Definition 5 and Corollary 17. The last step follows by applying Lemma 29 on $A\bar{x} - b$.

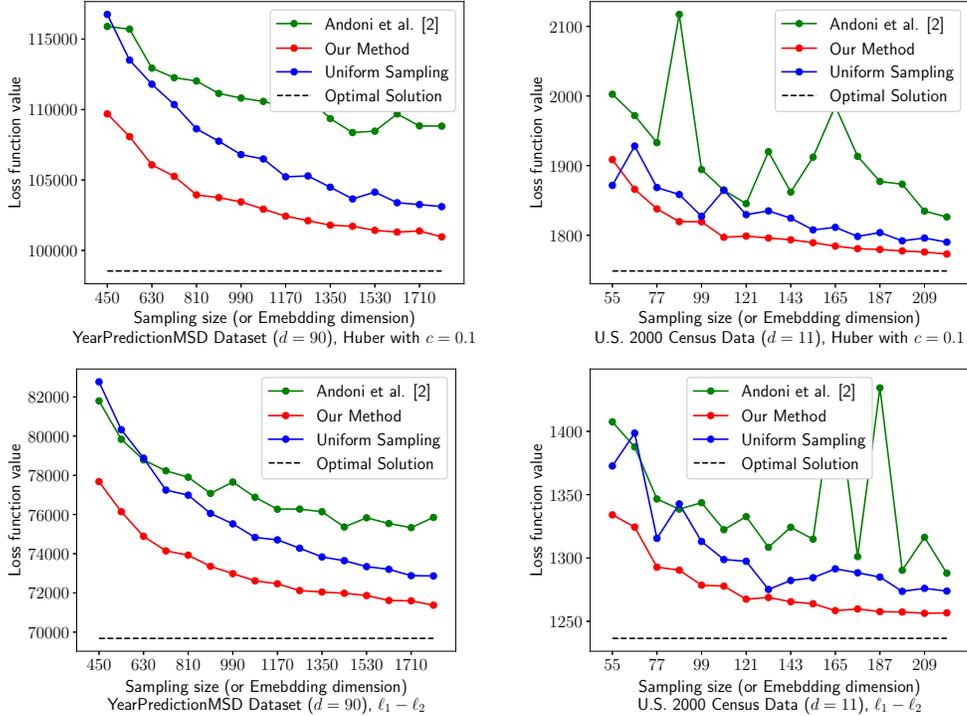


Figure 2: Experiments on Orlicz norm.

E Experiments

In this section, we perform experiments to validate the practicality of our methods.

Experiment Setup. We compare the proposed algorithms with baseline algorithms on the U.S. 2000 Census Data containing $n = 5 \times 10^6$ rows and $d = 11$ columns and UCI YearPredictionMSD dataset which has $n = 515,345$ rows and $d = 90$ columns. All algorithms are implemented in Python 3.7. To solve the optimization problems induced by the regression problems and their sketched versions, we invoke the `minimize` function in `scipy.optimize`. Each experiment is repeated for 25 times, and the mean of the loss function value is reported. In all experiments, we vary the sampling size or embedding dimension from $5d$ to $20d$, and observe their effects on the quality of approximation.

Experiments on Orlicz Norm. We compare our algorithm in Section 2 with uniform sampling and the embedding in [2]. We also calculate the optimal solution to verify the approximation ratio. We try Orlicz norms induced by two different G functions: Huber with $c = 0.1$ and “ $\ell_1 - \ell_2$ ”. See Table 1 for definitions. Our experimental results in Figure 2 clearly demonstrate the practicality of our algorithm. In both datasets, our algorithm outperforms both baseline algorithms by a significant margin, and achieves the best accuracy in almost all settings.

Experiments on Symmetric Norm. We compare our algorithm in Section 3 (SymSketch) with the optimal solution to verify the approximation ratio. We try two different symmetric norms: top- k norm with $k = n/5$ and sum-mix of ℓ_1 and ℓ_2 norm ($\|x\|_1 + \|x\|_2$). As shown in Figure 3, SymSketch achieves reasonable approximation ratios with moderate embedding dimension. In particular, the algorithm achieves an approximation ratio of 1.25 when the embedding dimension is only $5d$.

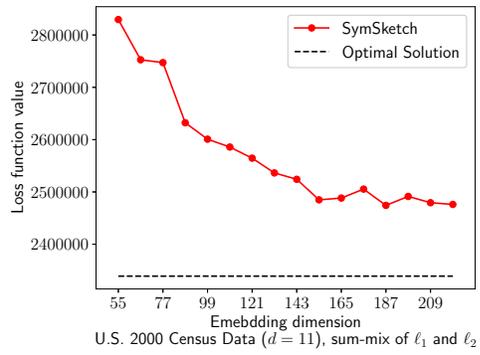
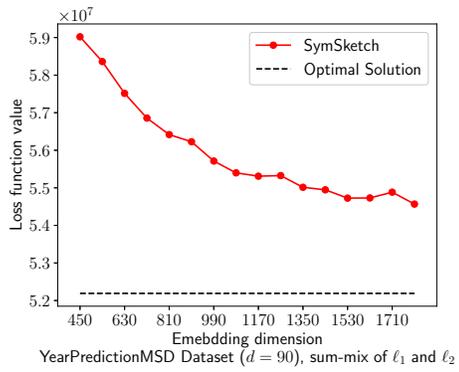
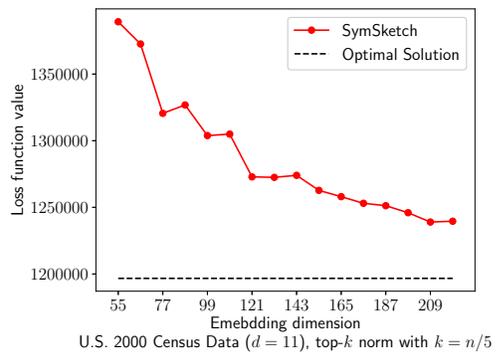
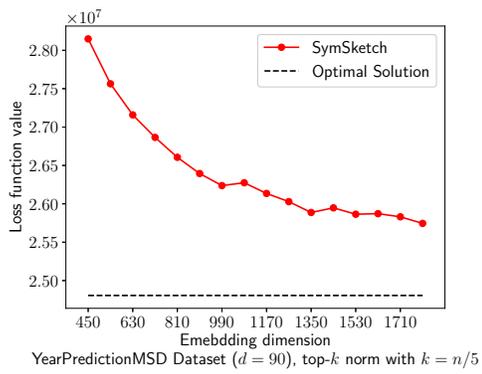


Figure 3: Experiments on symmetric norm.