The ℓ_p-Subspace Sketch Problem in Small Dimensions with Applications to Support Vector Machines

Yi Li* Honghao Lin† David P. Woodruff‡

Abstract

In the ℓ_p-subspace sketch problem, we are given an $n \times d$ matrix A with $n > d$, and asked to build a small memory data structure $Q(A, \varepsilon)$ so that, for any query vector $x \in \mathbb{R}^d$, we can output a number in $(1 \pm \varepsilon)\|Ax\|_p^p$ given only $Q(A, \varepsilon)$. This problem is known to require $\widetilde{\Omega}(d\varepsilon^{-2})$ bits of memory for $d = \Omega(\log(1/\varepsilon))$. However, for $d = o(\log(1/\varepsilon))$, no data structure lower bounds were known. Small constant values of d are particularly important for estimating point queries for support vector machines (SVMs) in a stream (Andoni et al. 2020), where only tight bounds for $d = 1$ were known.

We resolve the memory required to solve the ℓ_p-subspace sketch problem for any constant d and integer p, showing that it is $\Omega(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ bits and $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ words, where the $\widetilde{O}(\cdot)$ notation hides poly$(\log(1/\varepsilon))$ factors. This shows that one can *beat* the $\Omega(\varepsilon^{-2})$ lower bound, which holds for $d = \Omega(\log(1/\varepsilon))$, for any constant d. Further, we show how to implement the upper bound in a single pass stream, with an additional multiplicative poly$(\log \log n)$ factor and an additive poly$(\log n)$ cost in the memory. Our bounds extend to loss functions other than the ℓ_p-norm, and notably they apply to point queries for SVMs with additive error, where we show an optimal bound of $\widetilde{\Theta}(\varepsilon^{-\frac{2d}{d+3}})$ for every constant d. This is a near-quadratic improvement over the $\Omega(\varepsilon^{-\frac{d+1}{d+3}})$ lower bound of Andoni et al. Further, previous upper bounds for SVM point query were noticeably lacking: for $d = 1$ the bound was $\widetilde{O}(\varepsilon^{-1/2})$ and for $d = 2$ the bound was $\widetilde{O}(\varepsilon^{-4/5})$, but all existing techniques failed to give any upper bound better than $\widetilde{O}(\varepsilon^{-2})$ for any other value of d. Our techniques, which rely on a novel connection to low dimensional techniques from geometric functional analysis, completely close this gap.

1 Introduction

We consider the subspace sketch problem, which is the problem of designing a low memory data structure to compress a given $n \times d$ matrix A, so that later given only the compressed version of A, one can query norms of vectors of the form Ax for $x \in \mathbb{R}^d$. Formally,

Definition 1.1. *In the subspace sketch problem, we are given an $n \times d$ matrix A with entries specified by $O(\log(nd))$ bits, an accuracy parameter $\varepsilon > 0$, and a function $\Phi : \mathbb{R}^n \to \mathbb{R}^{\geq 0}$, and the goal is to design a data structure Q_Φ so that, with constant probability, simultaneously for all $x \in \mathbb{R}^d$, $Q_\Phi(x) = (1 \pm \varepsilon)\Phi(Ax)$.*

An important case of the above is when the functions correspond to the classical ℓ_p-norms, i.e., $\Phi(x) = \sum_{i=1}^n |x_i|^p$ for some $p \geq 1$. A space bound of $d^{O(p)}$ words is known for even integers p, independent of ε. For p that is not an even integer, it was shown in [LWW21] that there is an $\widetilde{\Omega}(d\varepsilon^{-2})$ lower bound on the memory required to solve the subspace sketch problem for $d = \Omega(\log(1/\varepsilon))$. Here and throughout, the $\widetilde{O}(\cdot)$ notation hides poly-logarithmic factors in its arguments. The fact that $d = \Omega(\log(1/\varepsilon))$ was crucial for the arguments in [LWW21], and a natural question is if the same $\widetilde{\Omega}(d\varepsilon^{-2})$ lower bound holds for smaller d, in particular for constant d.

The interest in constant d is particularly motivated given the recent work of [ABL+20], which studied the support vector machines (SVM) problem in constant dimensions in the streaming setting. Here x can be thought

*Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University. Email: yili@ntu.edu.sg. Supported in part by Singapore Ministry of Education Tier 1 grant RG75/21.
†Computer Science Department, Carnegie Mellon University. Email: honghaol@andrew.cmu.edu. Supported by National Science Foundation (NSF) Grant CCF-1815840.
‡Computer Science Department, Carnegie Mellon University. Email: dwoodruf@andrew.cmu.edu. Supported in part by National Science Foundation (NSF) Grant CCF-1815840.

Parameters	Lower Bound		Upper Bound	
$p = 1$	$\Omega\big(\varepsilon^{-\frac{2(d-1)}{d+2p}}\big)$	Theorem 3.1	$\widetilde{O}\big(\varepsilon^{-\frac{2(d-1)}{d+2}}\big)$	[Mat96]
$p \in \mathbb{Z} \setminus 2\mathbb{Z}$			$\widetilde{O}\big(\varepsilon^{-\frac{2(d-1)}{d+2p}}\big)$	Theorem 4.2
$p \in [1,\infty) \setminus \mathbb{Z}$			$\widetilde{O}\big(\varepsilon^{-\frac{2(d^q-1)}{d^q+2}}\big)$	$q = \lceil \frac{p}{2} \rceil$, Section A.2
$p \in [1,\infty) \setminus \mathbb{Z}$, $d \geq 5$			$\widetilde{O}\big(\varepsilon^{-\frac{2(d-1)}{d+2}}\big)$	Theorem A.1
$p \in (d-1,\infty) \setminus \mathbb{Z}$			$\widetilde{O}\big(\varepsilon^{-\frac{2d}{2p-d+2}}\big)$	Section A.3
$p \in 2\mathbb{Z}$	no dependence on ε		$O(1)$	[Sch11, LWW21]
SVM	$\Omega\big(\varepsilon^{-\frac{2d}{d+3}}\big)$	Theorem 7.2	$\widetilde{O}\big(\varepsilon^{-\frac{2d}{d+3}}\big)$	Theorem 7.1

Table 1: A summary of existing results and our results for the ℓ_p-subspace sketch problem. The lower bounds are in terms of the number of bits and the upper bounds are in the number of words.

of as a pair $(\theta, b) \in \mathbb{R}^d \times \mathbb{R}$ and each of the n rows of A can be thought of as a pair $(x_i, y_i) \in \mathbb{R}^d \times \{-1, 1\}$, and for a parameter $\lambda > 0$ we have:

$$\Phi((\theta, b)) = \frac{\lambda}{2}\|(\theta, b)\|_2^2 + \frac{1}{n}\sum_{i=1}^{n} \max\{0, 1 - y_i(\theta^T x_i + b)\}. \tag{1.1}$$

The authors refer to the above as the "point query" version of SVM and show an $\Omega\big(\varepsilon^{-\frac{d+1}{d+3}}\big)$ bit lower bound for any single-pass streaming algorithm for solving this problem. In fact, their lower bound applies to the memory required of any data structure for solving the subspace sketch problem with Φ as in (1.1). In terms of upper bounds, [ABL+20] show an $\widetilde{O}(\varepsilon^{-1/2})$ bound for $d = 1$ and an $\widetilde{O}(\varepsilon^{-4/5})$ bound for $d = 2$. For any $d > 2$, the best known upper bound is a trivial $\widetilde{O}(\varepsilon^{-2})$ bound obtained by uniform sampling. In fact, these upper bounds are also all one-pass streaming algorithms. One of the major open questions of [ABL+20] was to close this nearly quadratic gap for large constant d.

The ℓ_p-subspace sketch problem has also been studied in functional analysis for constant values of d for the special case of $p = 1$, and for the special case of requiring an embedding, i.e., a low dimension m and a matrix B of m rows so that $\|Bx\|_1 = (1 \pm \varepsilon)\|Ax\|_1$ for all x. In particular, a dimension of $\widetilde{O}\big(\varepsilon^{-\frac{2(d-1)}{d+2}}\big)$ was established in a sequence of work [BM83, Gor85, Sch87, Lin89, BL88, Mat96], while a matching lower bound for this particular type of subspace sketch (and for $p = 1$ and constant d) was shown in [BLM89]. There are many natural questions left open by the functional analysis work: (1) can the upper bound be made a streaming upper bound with a small amount of memory? (2) does the lower bound hold for arbitrary data structures, (3) can the arguments extend to $p > 1$, etc.?

Throughout the remainder of this section, we assume that $d \geq 2$ and $p \geq 1$ are constants.

1.1 Our Results A summary of our results is provided in Table 1.

In this paper, we show that the $\Omega\big(\varepsilon^{-\frac{2(d-1)}{d+2}}\big)$ lower bound actually holds for any type of data structure for $p = 1$. Furthermore, for every $p \in [1, \infty) \setminus 2\mathbb{Z}$, we obtain a lower bound of $\Omega\big(\varepsilon^{-\frac{2(d-1)}{d+2p}}\big)$.

Theorem 1.1. *Suppose that $p \in [1, \infty) \setminus 2\mathbb{Z}$. Any data structure that solves the ℓ_p-subspace sketch problem for dimension d and accuracy parameter ε requires $\Omega(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ bits of space.*

For every integer p, we obtain an $\widetilde{O}\big(\varepsilon^{-\frac{2(d-1)}{d+2p}}\big)$ upper bound, matching the lower bound up to logarithmic factors.

Theorem 1.2. (Informal) *Suppose that $A \in \mathbb{R}^{n \times d}$ and p is a positive integer. There is a polynomial-time algorithm that maintains a data structure using $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ words of space, which solves the ℓ_p-subspace sketch problem.*

Moreover, we show the upper bound above can be implemented in a single pass row-arrival stream, with an additional multiplicative poly$(\log\log n)$ factor and an additive poly$(\log n)$ cost in the memory.

THEOREM 1.3. (INFORMAL) *Let $A = a_1 \circ \cdots \circ a_n$ be a stream of n rows, where $a_i \in \mathbb{R}^{1\times d}$. There is an algorithm that maintains a data structure in $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ words of space, which solves the ℓ_p-subspace sketch problem. Moreover, the algorithm can be implemented in $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}} + \log^{\frac{3d+2p-2}{d+2p}} n)$ words of space.*

We also obtain an $O(1)$-update time algorithm with a slightly worse $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ words of space bound for general $d = O(1)$ and a tight space bound of $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ for $d \le 2p + 2$.

THEOREM 1.4. (INFORMAL) *Let $A = a_1 \circ \cdots \circ a_n$ be a stream of n rows, where $a_i \in \mathbb{R}^{1\times d}$. There is an algorithm which maintains a data structure Q of $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ words of space which solves the ℓ_p-subspace sketch problem. Moreover, the algorithm updates Q in $O(1)$ time and can be implemented using $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ words of space.*

When $d \le 2p + 2$, the size of Q can be improved to $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ words of space and the whole algorithm can be implemented in $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ words of space.

To obtain a tight bound for the SVM point query problem mentioned above, we also study the following version of the affine ℓ_p-subspace sketch problem, where

$$\Phi(x, b) = \sum_{i=1}^{n} |b - \langle A_i, x\rangle| .$$

for every $x \in \mathbb{R}^d$ and $b \in \mathbb{R}$. We show a tight space complexity of $\widetilde{\Theta}(\varepsilon^{-\frac{2d}{d+2p+1}})$.

THEOREM 1.5. (INFORMAL) *Suppose that $p \in [1, \infty) \setminus 2\mathbb{Z}$. Any data structure that solves the affine ℓ_p-subspace sketch problem for dimension d and accuracy parameter ε requires $\Omega(\varepsilon^{-\frac{2d}{d+2p+1}})$ bits of space.*

When p is an integer, there is a polynomial-time algorithm that returns a data structure with $\widetilde{O}(\varepsilon^{-\frac{2d}{d+2p+1}})$ words of space, which solves the affine ℓ_p subspace sketch problem.

Based on these results, we show that a tight space bound for the point query problem can be derived via a black box reduction with $p = 1$.

THEOREM 1.6. (INFORMAL) *Any data structure which solves the d-dimensional point estimation problem for SVM requires $\Omega(\varepsilon^{-\frac{2d}{d+3}})$ bits of space.*

Furthermore, there is an algorithm that maintains a data structure using $\widetilde{O}(\varepsilon^{-\frac{2d}{d+3}})$ words of space, which solves the d-dimensional point query problem for SVM. The algorithm can be implemented using $\widetilde{O}(\varepsilon^{-\frac{2d}{d+3}})$ words of space.

Lastly, we obtain results for non-integer p, giving algorithms with $o(\varepsilon^{-2})$ words of space even for such p. The details are postponed to Appendix A.

1.2 Our Techniques One of our key technical contributions is to connect the SVM point query problem to the ℓ_1-subspace sketch problem and to use techniques for the latter from geometric functional analysis, which previously had not been considered in the context of the SVM problem [ABL+20]. Throughout this section, we assume that $p \ge 1$ is not an even integer.

1.2.1 Lower Bound The idea behind our lower bound for the subspace sketch problem is to give Alice one of $m = 2^{\Omega(n)}$ possible subsets $S_1, \ldots, S_m$ of $n/2$ points of a fixed set $S = \{p_1, \ldots, p_n\}$ of n points on the unit sphere $\mathbb{S}^{d-1}$, where $\|p_i - p_j\|_2 \ge \eta$ for all $i \ne j$. Here we have $|S_i \cap S_j| \le n/4$ for every pair $i \ne j$. If Alice has a specific subset S_i, she can send the subspace sketch of her set to Bob. Bob then pretends he has an S_j and enumerates over all possible queries x, and by construction of our sets S_i, we will (abusing notation and writing a set as the matrix whose rows correspond to the entries in the set) have that $|\|S_i x\|_p - \|S_j x\|_p|$ is larger than the tolerable subspace sketch error, and Bob will be able to determine that $i \ne j$.

Our main novelty is Lemma 3.2, which says that for two sets A and B of points on the sphere, each symmetric around the origin and such that no point in A is close to any point in B, there is some direction x on the sphere for which $|\|Ax\|_p^p - \|Bx\|_p^p|$ is large. The proof of this lemma is inspired by ideas from geometric functional analysis [BLM89], which give lower bounds for the specific case when the data structure is an embedding for $p = 1$ and for constant d. Indeed, as in their proof, we make use of spherical harmonics. However, we require a significant strengthening of the arguments in [BLM89]. In particular, the lemma in [BLM89] can be seen as lower bounding $\max_x \|Ax\|_p^p - \min_x \|Ax\|_p^p$, which, after expanding $\|Ax\|_p^p$ in a spherical harmonic series, boils down to lower bounding

$$\sum_{i,j} f(\langle A_i, A_j\rangle), \quad \text{where} \quad f(t) = \frac{1 - r^4}{(1 + r^4 - 2r^2 t)^{d/2}},$$

where r is a parameter to be determined. The lower bound of [BLM89] is obtained simply by considering only the terms for which "$i = j$". In our case, we need to lower bound $\max_x(\|Ax\|_p^p - \|Bx\|_p^p) - \min_x(\|Ax\|_p^p - \|Bx\|_p^p)$, which reduces to lower bounding a more complicated quantity:

$$\sum_{i,j} f(\langle A_i, A_j\rangle) + \sum_{i,j} f(\langle B_i, B_j\rangle) - \sum_{i,j} f(\langle A_i, B_j\rangle). \tag{1.2}$$

The first two terms can be lower bounded similarly by taking only the "$i = j$" terms as before. However, there is a third term which causes additional complications since it requires a good upper bound. To do this, for each point A_i, we partition the points B_j into level sets of geometrically increasing distances from A_i. The critical observation is that the number of points in each level set grows at a slower rate than the function f decays. Hence, the contribution from each level set is geometrically decreasing and the total contribution for each fixed A_i can thus be controlled. In the end, we are able to show that the third term in (1.2) is at most a constant fraction of the first two terms, which leads to our desired lower bound.

1.2.2 Upper Bound Our upper bound is inspired from an argument of Matousek [Mat96] for $p = 1$. We give a high-level description of this idea, assuming first that each point $A_i \in \mathbb{S}^{d-1}$. The points A_i are partitioned into a number of groups $P_1, \ldots, P_s$ each of diameter at most $\eta = \varepsilon^{2/(d+2)}$ (Lemma 2.3), then $\Theta(d)$ random points are chosen from each group, such that the barycenter of the randomly selected points is the same as the barycenter of the group (Lemma 4.2). The data structure stores the selected points in each group as a surrogate for the group.

For a fixed query point x, each group P_j belongs to one of three types, based on its relative position to the the equator $\{y : \langle x, y\rangle = 0\}$: positive type, if $\langle A_i, x\rangle > C\eta$ for all $A_i \in P_j$; negative type, if $\langle A_i, x\rangle < -C\eta$ for all $A_i \in P_j$; zero type, if $|\langle A_i, x\rangle| \le C\eta$ for all $A_i \in P_j$, where $C > 0$ is an absolute constant. When P_j is of positive type, we have $\sum_{A_i \in P_j} |\langle A_i, x\rangle| = \sum_{A_i \in P_j} \langle A_i, x\rangle = \langle \sum_{A_i \in P_j} A_i, x\rangle$, which can be calculated exactly without error from the sampled points, as guaranteed by the barycenter property. When P_j is of negative type, the contribution can be calculated exactly in a similar manner. Next, consider the groups P_j of zero type. Since all summands $|\langle A_i, x\rangle|$ are small, a Bernstein bound shows that using randomly selected points can approximate the *total* contribution from *all* zero-type groups up to a small additive error with high probability. Taking a union bound over a net for query points x, the overall sum $\sum_i |\langle A_i, x\rangle|$ can be estimated with a small additive error ε with high probability simultaneously for all $x \in \mathbb{S}^{d-1}$.

Now, suppose that p is an odd integer. The key observation is a "tensor trick"

$$\langle x, y\rangle^p = \langle x^{\otimes p}, y^{\otimes p}\rangle,$$

where $x^{\otimes p}$ and $y^{\otimes p}$ are d^p-dimensional vectors. Lemma 4.2 is then applied to a group of d^p-dimensional points, so $\Theta(d^p)$ points are selected randomly in each group and stored in the data structure. Another crucial observation is that in this way, all the error comes from the terms $|\langle A_i, x\rangle|^p$ such that $|\langle A_i, x\rangle| \le \eta$, which implies $|\langle A_i, x\rangle|^p \le \eta^p \ll \eta$. This allows for tighter concentration than what is possible for $p = 1$. We can therefore estimate $\sum_i |\langle A_i, x\rangle|^p$ up to a small additive error ε for all $x \in \mathbb{S}^{d-1}$, assuming that $\|A_i\|_2 = O(1)$ for all i.

The procedure above can be generalized to estimate $\sum_i w_i |\langle A_i, x\rangle|^p$ up to an additive error of $\varepsilon \sum_i w_i$, where $w_i \ge 0$ is the weight associated with the point A_i. This requires that the random points selected from each group carry (new) weights such that the weighted barycenters are the same. This was already attained in Matousek's work (Lemma 4.2).

In order to obtain a multiplicative error for a general matrix A, we perform a preprocessing step, which transforms the John ellipsoid of $\{x \in \mathbb{R}^d : \|Ax\|_p \le 1\}$ to the unit ℓ_2 ball in $\mathbb{R}^d$ via a linear transformation, giving a matrix A' for which we can show that $\|A'_i\|_2 = O(1)$ and $\|A'x\|_p^p = \Omega(\sum_i \|A'_i\|_2^p)$. This is sufficient to deduce that the additive error to the normalized version of A' is in fact a multiplicative error for A.

We remark that having a p-th power in $|\langle A_i, x\rangle|^p$ is only useful when $\langle A_i, x\rangle$ is small, and improvements exploiting this benefit may not occur in other algorithms. For example, Matousek proposed another algorithm [Mat96] which removes the $\log(1/\varepsilon)$ factors for $d \geq 5$, obtaining a clean $O(\varepsilon^{-\frac{2(d-1)}{d+2}})$ bound. The analysis of this algorithm relies on the fact that the function $x \mapsto |\langle u, x\rangle| - |\langle v, x\rangle|$ is Lipschitz on the region where $\langle u, x\rangle$ and $\langle v, x\rangle$ have the same sign and the Lipschitz constant is proportional to $\|u - v\|_2$. However, we cannot expect a substantially smaller Lipschitz constant for the function $x \mapsto |\langle u, x\rangle|^p - |\langle v, x\rangle|^p$. Interestingly, as we shall show in Appendix A, despite it failing to give a tight bound for integer $p > 1$, this algorithm actually implies an $O\big(\varepsilon^{-\frac{2(d-1)}{d+2}}\big)$ upper bound for all constant $p > 1$ once $d \geq 5$.

1.2.3 Streaming Upper Bound The preceding approach for the upper bound does not give a streaming algorithm since computing the groups $P_1, \ldots, P_s$ to perform the partition requires storing all of the points. Nevertheless, it can be viewed as a coreset procedure which, given a set of weighted points (A, w) (where w is a vector in which w_i is the weight for A_i), outputs a small subset $B \subseteq A$ of size $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ together with new weights w' such that $\sum_j w'_j \,|\langle B_j, x\rangle|^p = (1 \pm \varepsilon) \sum_i w_i \,|\langle A_i, x\rangle|^p$. The standard merge-and-reduce framework (see, e.g., [BDM+20] in the context of numerical linear algebra) can then be applied, leading to a streaming algorithm using $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}} \operatorname{polylog}(n/\varepsilon))$ words of space. However, this memory bound depends on the product of a term involving ε and a term involving $\log n$. By instead running the algorithm with $\varepsilon = \Theta(1)$ and using it to estimate the so-called ℓ_p-sensitivities, according to which we can sample the points A_i, we can replace n with $\operatorname{poly}(d/\varepsilon) \log n$, and then run the merge-and-reduce framework on this new value of n, resulting in only a $\operatorname{poly}(\log\log n)$ factor multiplying the term depending on ε, plus an additive $\operatorname{poly}(\log n)$ term.

1.2.4 Connection to SVM As mentioned, one of our key contributions is showing that the SVM point query problem can be related the ℓ_1-subspace sketch problem via a black-box reduction. As shown in [ABL+20], the function Φ for SVM can be modified to

$$\Phi(\theta, b) = \frac{1}{n} \sum_{i=1}^{n} \max\{0, b - \theta^T x_i\} \ ,$$

without loss of generality. Consider the special case when $b = 0$. Suppose that $X = \{x_i\}$ is the point set given by the data stream. Let $-X = \{-x : x \in X\}$ and observe that

$$\underset{X}{\Phi}(\theta, 0) + \underset{-X}{\Phi}(\theta, 0) = \frac{1}{n} \sum_i \big(\max\{0, \theta^\top x_i\} + \max\{0, -\theta^\top x_i\}\big) = \frac{1}{n} \sum_i \left|\theta^\top x_i\right| \ ,$$

which means that a lower bound for the d-dimensional ℓ_1-subspace actually yields a lower bound for the d-dimensional point query problem for SVM. To obtain a tight lower bound, we instead study a special affine version of the ℓ_p-subspace sketch problem, where

$$\Phi(x, b) = \sum_{i=1}^{n} |b - \langle A_i, x\rangle|$$

for a given $x \in \mathbb{R}^d$ and $b \in \mathbb{R}$.

The key observation is the following. Given a matrix $A \in \mathbb{R}^{n\times d}$, let $B \in \mathbb{R}^{n\times(d+1)}$ be the matrix in which the i-th row $B_i = \begin{pmatrix} A_i & -1 \end{pmatrix}$. Suppose that $x \in \mathbb{R}^d$ and $b \in \mathbb{R}$ are the query vector and value, and let $y = \begin{pmatrix} x^\top & b \end{pmatrix}^\top \in \mathbb{R}^{d+1}$. Then

$$\sum_i |\langle A_i, x\rangle - b|^p = \sum_i |\langle B_i, y\rangle|^p = \|By\|_p^p \, .$$

Hence, the d-dimensional affine ℓ_p-subspace sketch is related to the $(d+1)$-dimensional ℓ_p-subspace sketch problem where for all points A_i, the last coordinate is 1. A closer examination of the lower bound for the ℓ_p subspace sketch problem reveals that the lower bound for $\max_x(\|Ax\|_p^p - \|Bx\|_p^p) - \min_x(\|Ax\|_p^p - \|Bx\|_p^p)$ continues to hold, up to a constant factor, even when A and B do not necessarily lie on $\mathbb{S}^{d-1}$ but rather in a thin spherical shell, i.e., $\|A_i\|_2, \|B_i\|_2 = \Theta(1)$, provided that their projections on $\mathbb{S}^{d-1}$, $A_i/\|A_i\|_2$ and $B_i/\|B_i\|_2$, are sufficiently separated from each other. Hence, the point set S in the communication problem can now be chosen from the spherical cap $\{x \in \mathbb{S}^d : x_{d+1} = \Omega(1)\}$ so that normalizing the last coordinate x_{d+1} to 1 yields a vector $x' = x/x_{d+1}$ lying inside the spherical shell. The same lower bound (up to a constant factor) for the $(d+1)$-dimensional ℓ_p sketch problem then follows.

2 Preliminaries

Notation Let $\mathbb{S}^{d-1}$ denote the unit sphere in $\mathbb{R}^d$, i.e., $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$, and let Δ_{d-1} denote the standard $(d-1)$-simplex, i.e., $\Delta_{d-1} = \{x \in \mathbb{R}^d : x_1 + \cdots + x_d = 1 \text{ and } x_i \geq 0 \text{ for all } i\}$.

For two matrices A and B of the same number of columns, we denote their vertical concatenation by $A \circ B$.

Spherical Harmonics The spherical harmonics $\{Y_{k,j}\}_{k,j}$ form an orthonormal basis in $L^2(\mathbb{S}^{d-1}, \sigma_{d-1})$, where $\sigma_{d-1}(x)$ denotes the normalized rotationally-invariant measure on $\mathbb{S}^{d-1}$. Here $k \geq 0$ is an integer and $j = 1, \ldots, M(d,k)$ for each k, where $M(d,k) = \binom{k+d-2}{d-2} + \binom{k+d-3}{d-2} = O(k^{d-2})$. The following are some useful properties of spherical harmonics (see, e.g. [AH12, DX13]).

- (Addition Theorem) For all $x, y \in \mathbb{S}^{d-1}$,

$$\sum_j Y_{k,j}(x)Y_{k,j}(y) = M(d,k)P_{k,d}(\langle x, y\rangle), \tag{2.3}$$

  where $P_{k,d}(t)$ is the Legendre polynomial of degree k in dimension d.

- (Funk-Hecke Formula) Suppose that $f : [-1,1] \to \mathbb{R}$ is a function and $y \in \mathbb{S}^{d-1}$. For $h : \mathbb{S}^{d-1} \to \mathbb{R}$ defined as $h(x) = f(\langle x, y\rangle)$, it holds that

$$\langle h, Y_{k,j}\rangle = \int_{\mathbb{S}^{d-1}} h(x)Y_{k,j}(x)\, d\sigma_{d-1}(x) = \lambda_k Y_{k,j}(y), \tag{2.4}$$

  where

$$\lambda_k = \frac{\Gamma(\frac{d}{2})}{\sqrt{\pi}\Gamma(\frac{d-1}{2})} \int_{-1}^{1} f(t)(1-t^2)^{\frac{d-3}{2}} P_{k,d}(t)dt. \tag{2.5}$$

- (Poisson Identity) It holds for all $0 \leq r < 1$ and all $t \in [-1,1]$ that

$$\frac{1-r^2}{(1+r^2-2rt)^{d/2}} = \sum_{k=0}^{\infty} M(d,k)r^k P_{k,d}(t). \tag{2.6}$$

LEMMA 2.1. *Suppose that $d \geq 2$ and $p \geq 1$ are constants. Let $f(t) = |t|^p$ and λ_k be as defined in* (2.5). *It holds that (i) $\lambda_k = 0$ for odd k; (ii) when p is not an even integer, $\lambda_k \neq 0$ for all even k and*

$$|\lambda_k| \sim_{d,p} \frac{1}{k^{d/2+p}} \sin\frac{\pi p}{2}, \quad \text{even } k \to \infty;$$

(iii) when p is an even integer, $\lambda_k = 0$ for all $k > p$.

Proof. Since $P_{k,d}(t)$ is an odd function when k is odd, it is clear that $\lambda_k = 0$ when k is odd. We shall assume that k is even in the rest of the proof. In this case $P_{k,d}$ is an even function.

Note that $P_{n,k}$ is the normalized Gegenbauer polynomial, $P_{k,d}(t) = C_k^{\alpha}(t)/C_k^{\alpha}(1)$ with $\alpha = d/2-1$. It is known that $C_k^{\alpha}(1) = (2\alpha)_k/(k!)$. Invoking the identity (2.21.1.1) from [PBM88], we have that

$$\lambda_k = c(d)(-1)^{k/2}\Gamma\left(\alpha + \frac{1}{2}\right) \frac{\Gamma(\frac{p+1}{2})}{\Gamma(1+\frac{p}{2}+\alpha+\frac{k}{2})} \left(\frac{-p}{2}\right)_{k/2}.$$

Now it is clear that $\lambda_k = 0$ when p is an even integer and $k > p$, and $\lambda_k \neq 0$ for all even k when p is not an even integer. In the latter case, when $k > 0$,

$$\lambda_k = c(d,p)(-1)^{k/2+1} \sin\left(\frac{\pi p}{2}\right) \frac{\Gamma(\frac{k-p}{2})}{\Gamma(\frac{k+d+p}{2})} \sim_{d,p} (-1)^{k/2+1} \sin\left(\frac{\pi p}{2}\right) \frac{1}{k^{d/2+p}}.$$

□

Volume of Spherical Caps For a point $x \in \mathbb{S}^{d-1}$ and $r > 0$, consider the spherical cap $C(x,r) = \{y \in \mathbb{S}^{d-1} : \|x-y\|_2 \leq r\}$. It is clear that $\sigma_{d-1}(C(x,r))$ is independent of x and so we write it as $\sigma_{d-1}(\mathrm{Cap}(r))$. We cite a result on the volume of spherical caps from [BW03] as follows.

LEMMA 2.2. ([BW03, LEMMA 3.1]) *When $r^2 \leq 2(1-1/\sqrt{d+1})$, it holds that*

$$\kappa_d \frac{r^{d-1}}{1-r^2/2}\left(1-\frac{r^2}{4}\right)^{\frac{d-1}{2}} \leq \sigma_{d-1}(\mathrm{Cap}(r)) \leq \kappa_d \frac{r^{d-1}}{1-r^2/2}\left(1-\frac{r^2}{4}\right)^{\frac{d-1}{2}},$$

where $\kappa_d = \Gamma(\frac{d}{2})/(2\sqrt{\pi}\Gamma(\frac{d+1}{2}))$.

Partition of Sphere We shall need the following partition lemma in [Mat96].

LEMMA 2.3. ([MAT96, LEMMA 5]) *Let P be an N-point set in $\mathbb{S}^{d-1}$, and let $s \geq 2$ be a constant. There is an $O(N^{d-\frac{1}{d-1}})$-time deterministic algorithm which computes a collection of disjoint s-point subsets $P_1, \ldots, P_t \subset P$, which together contain at least half the points of P, and with the following properties:*

(i) For every $x \in \mathbb{S}^{d-1}$, the hyperplane $\{y : \langle x, y\rangle = 0\}$ only intersects at most $O(N^{1-\frac{1}{d-1}})$ sets among P_i. Here "intersect" means that there exist x and y in P_i such that x, y are on different sides of the hyperplane.

(ii) Each P_i has diameter at most $O(N^{-\frac{1}{d-1}})$.

3 ℓ_p-Subspace Sketch Lower Bound

We first state an auxiliary lemma.

LEMMA 3.1. *Suppose that $c(d) \leq r < 1$. Define a function $f : [-1,1] \to \mathbb{R}$ as*

$$f(t) = \frac{1-r^4}{(1+r^4-2r^2t)^{d/2}} + \frac{1-r^4}{(1+r^4+2r^2t)^{d/2}}.$$

Then for $0 \leq t \leq 1$, it holds that

$$f(t) \leq \frac{2(1-r^2)}{(2r^2)^{d/2}}\left(\frac{1}{(1-t)^{d/2}}+1\right).$$

Proof. Let $q = 1 - r^2$. Then the first term of $f(t)$ is

$$\frac{1-r^4}{(1+r^4-2r^2t)^{d/2}} \leq \frac{2q-q^2}{(2-2q+q^2-2(1-q)t)^{d/2}} \leq \frac{2q}{(2(1-q)(1-t))^{d/2}}$$

and the second term of $f(t)$ is

$$\frac{1-r^4}{(1+r^4+2r^2t)^{d/2}} \leq \frac{2q-q^2}{(2-2q+q^2)^{d/2}} \leq \frac{2q}{(2(1-q))^{d/2}}.$$

The claimed result follows. □

Our result is mainly based on the following lemma.

LEMMA 3.2. *Suppose that $p \in [1,\infty) \setminus 2\mathbb{Z}$ is a constant. Let A and B be sets of $n \le N$ points on $\mathbb{S}^{d-1}$. Suppose that A and B are symmetric around the origin and $\|A_i - B_j\|_2 \ge \eta = C_1 N^{-\frac{1}{d-1}}$ for all i, j. Then we have*

$$\delta \equiv \sup_{x\in\mathbb{S}^{d-1}} \frac{1}{n} \left| \|Ax\|_p^p - \|Bx\|_p^p \right| \ge c_2 N^{-\frac{d+2p}{2(d-1)}}.$$

In the statement above, $C_1 > 0$ is a constant that depends only on d and $c_2 > 0$ is a constant that depends on d and p only.

Proof. The proof is inspired by the proof of [BLM89, Proposition 6.6], which uses spherical harmonics. Let $h_A(x) = \frac{1}{n}\|Ax\|_p^p = \frac{1}{n}\sum_i |\langle A_i, x\rangle|^p$ and let $h_B(x)$ be defined similarly. We expand $h_A - h_B$ into

$$h_A - h_B = \sum_{k=0}^{\infty}\sum_j \langle h_A - h_B, Y_{k,j}\rangle Y_{k,j}.$$

where $\langle h, Y_{k,j}\rangle$ denotes the inner product in $L_2(\sigma_d)$ for which

$$\langle h, Y_{k,j}\rangle = \int_{\mathbb{S}^{d-1}} h(x) Y_{k,j}(x)\, d\sigma_{d-1}(x).$$

It follows from Parseval's identity that

$$\delta^2 \ge \|h_A - h_B\|_{L_2(\sigma_{d-1})}^2 = \sum_{k=1}^{\infty}\sum_j \langle h_A - h_B, Y_{k,j}\rangle^2.$$

Now, from (2.3) and (2.6) we have for all $u, v \in \mathbb{S}^{d-1}$ that

$$\frac{1-r^2}{(1+r^2-2r\langle u,v\rangle)^{d/2}} = \sum_{k=0}^{\infty} r^k \sum_j Y_{k,j}(u) Y_{k,j}(v). \tag{3.7}$$

Combining with (2.4) and using the fact that A and B are symmetrically distributed on $\mathbb{S}^{d-1}$, we obtain that

$$\frac{1}{n}\sum_{i=1}^{n} \frac{(1-r^2)}{(1+r^2-2r\langle u, A_i\rangle)^{d/2}} = 1 + \sum_{\text{even } k\ge 2} r^k \lambda_k^{-1} \sum_j \langle h_A, Y_{k,j}\rangle Y_{k,j}(u),$$

and

$$\frac{1}{n}\sum_{i=1}^{n} \frac{(1-r^2)}{(1+r^2-2r\langle u, B_i\rangle)^{d/2}} = 1 + \sum_{\text{even } k\ge 2} r^k \lambda_k^{-1} \sum_j \langle h_B, Y_{k,j}\rangle Y_{k,j}(u).$$

Hence,

$$\frac{1}{n}\sum_{i=1}^{n} \frac{(1-r^2)}{(1+r^2-2r\langle u, A_i\rangle)^{d/2}} - \frac{1}{n}\sum_{i=1}^{n} \frac{(1-r^2)}{(1+r^2-2r\langle u, B_i\rangle)^{d/2}} = \sum_{\text{even } k\ge 2} r^k \lambda_k^{-1} \sum_j \langle h_A - h_B, Y_{k,j}\rangle Y_{k,j}(u),$$

Integrating with respect to u on $\mathbb{S}^{d-1}$, we have that

$$\begin{aligned} &\left\| \frac{1}{n}\sum_{i=1}^{n} \frac{(1-r^2)}{(1+r^2-2r\langle u, A_i\rangle)^{d/2}} - \frac{1}{n}\sum_{i=1}^{n} \frac{(1-r^2)}{(1+r^2-2r\langle u, B_i\rangle)^{d/2}} \right\|_{L_2(\sigma_{d-1})}^2 \\ &= \sum_{\text{even } k\ge 2} r^{2k}\lambda_k^{-2} \sum_j \langle h_A - h_B, Y_{k,j}\rangle^2 \\ &\le \delta^2 \max_{\text{even } k\ge 2} (r^{2k}\lambda_k^{-2}), \end{aligned} \tag{3.8}$$

The leftmost side of (3.8) equals (using (3.7))

$$\frac{1}{n^2}\left(\sum_{i,j}\frac{1-r^4}{(1+r^4-2r^2\langle A_i,A_j\rangle)^{d/2}}+\sum_{i,j}\frac{1-r^4}{(1+r^4-2r^2\langle B_i,B_j\rangle)^{d/2}}-2\sum_{i,j}\frac{1-r^4}{(1+r^4-2r^2\langle A_i,B_j\rangle)^{d/2}}\right).$$

Let I,J be the index sets of A and B such that the points in A_I,B_J are on the same hemisphere. We can rewrite the expansion above as

$$\frac{2}{n^2}\left(\sum_{i\in I,j\in I}f(\langle A_i,A_j\rangle)+\sum_{i\in J,j\in J}f(\langle B_i,B_j\rangle)-2\sum_{i\in I,j\in J}f(\langle A_i,B_j\rangle)\right), \tag{3.9}$$

where $f(t)$ is as defined in Lemma 3.1.

We now choose $(1-r^2)^{-d+1}=N$. Suppose that N is sufficiently large such that $1-r^2\le c_4(d)<1$, where $c_4(d)$ is a constant depending only on d. Next we turn to lower bounding the expression in (3.9).

Considering the summands with $i=j$. We see that the first two terms inside the bracket of (3.9) are at least

$$\sum_{i\in I}f(\langle A_i,A_i\rangle)+\sum_{j\in J}f(\langle B_j,B_j\rangle)\ge n\frac{1-r^4}{(1-r^2)^d}\ge n(1-r^2)^{-d+1}=nN.$$

Next we bound the cross terms. We first show that, for every A_i,

$$\sum_{j\in J}f(\langle A_i,B_j\rangle)\le\frac{1}{4}N. \tag{3.10}$$

Without loss of generality, we assume $\langle A_i,B_j\rangle\ge 0$ for all j, as otherwise we can replace B_j with $-B_j$. To achieve this, consider a fixed $i\in I$. Let $I_k=\{j\in J:\ 2^{k-1}\eta\le\|A_i-B_j\|_2<2^k\eta\}$ for $k=1,2,\dots,K$, where K is the smallest positive integer such that $2^K\eta\ge 1/2$. We also define $I_{K+1}=\{j\in J:\|A_i-B_j\|_2\ge 2^K\eta\}$. Then, we have that

$$\sum_{j\in J}f(\langle A_i,B_j\rangle)=\sum_{k=1}^{K+1}S_k=\sum_{k=1}^{K+1}\left(\sum_{j\in I_k}f(\langle A_i,B_j\rangle)\right).$$

Next, we will bound S_k individually. From Lemma 2.2 we have that

$$|I_k|\le\frac{\sigma_{d-1}(\mathrm{Cap}(2^k\eta))}{\sigma_{d-1}(\mathrm{Cap}(\eta/2))}\le C_3(d)\cdot(2^k)^{d-1}.$$

By Lemma 3.1, when $0\le t\le 1-2^{2k-1}\eta^2$ and $k\le K$, we have that

$$\begin{aligned}f(t)&\le\frac{2(1-r^2)}{(2r^2)^{d/2}}\left(\frac{1}{(1-t)^{d/2}}+1\right)\\&\le\frac{2(1-r^2)}{(2(1-c_4))^{d/2}}\left(\frac{1}{2^{(k-1/2)d}\eta^d}+1\right)\\&\le\frac{C_4(d)}{2^{kd}}\frac{1-r^2}{\eta^d}.\end{aligned}$$

Also, when $0\le t<7/8$, we have that

$$\begin{aligned}f(t)&\le\frac{2(1-r^2)}{(2r^2)^{d/2}}\left(\frac{1}{(1-t)^{d/2}}+1\right)\\&\le\frac{2(1-r^2)}{(2(1-c_4))^{d/2}}\left(\frac{1}{(1/8)^{d/2}}+1\right)\\&\le C_4(d)(1-r^2).\end{aligned}$$

Consequently (noting that $\langle A_i, B_j\rangle = 1 - \frac{\|A_i - B_j\|_2^2}{2}$),

$$\sum_{j\in J} f(\langle A_i, B_j\rangle) \le \left(\sum_{k\le K} C_3(d)(2^k)^{d-1}\cdot C_4(d)\frac{1-r^2}{\eta^d}\frac{1}{2^{kd}}\right) + \frac{n}{2}\cdot C_4(d)(1-r^2)$$
$$\le C_5(d)\left(\frac{1-r^2}{\eta^d} + \frac{N}{2}(1-r^2)\right) \le \frac{1}{4}N$$

provided that

$$\frac{1}{4}N \ge C_5(d)\left(\frac{1-r^2}{\eta^d} + \frac{N}{2}(1-r^2)\right) = C_5(d)\left(N^{-\frac{1}{d-1}}\eta^{-d} + \frac{1}{2}N^{1-\frac{1}{d-1}}\right),$$

which holds whenever

$$\eta \ge C_1(d)N^{-\frac{1}{d-1}}$$

and N is sufficiently large. It follows from (3.10) that

$$2\sum_{i\in I, j\in J} f(\langle A_i, B_j\rangle) \le \frac{1}{2}nN,$$

which implies that the expression in (3.9) is at least

$$\frac{2}{n^2}\left(nN - \frac{1}{2}nN\right) \ge 1.$$

Plugging this back into (3.8) yields

$$\delta^2 \max_{\text{even } k}(r^{2k}\lambda_k^{-2}) \ge 1.$$

By Lemma 2.1, there exists a constant $c(d)$ such that

$$c(d,p)\delta^2 \max_{\text{even } k}(r^{2k}k^{d+2p}) \ge \Omega(1).$$

The maximum is attained when $k \approx (d+2)/\ln(1/r) \sim \frac{d}{1-r^2}$, from which we obtain that

$$\delta \ge c(d,p)N^{-\frac{d+2p}{2(d-1)}}.$$

□

Equipped with the lemma above, we are now ready to prove our lower bound. We need the following lemma on the size of intersecting families.

Lemma 3.3. ([BBD15, p14]) *Suppose that $0 < \beta < \alpha < 1$ and n is sufficiently large. There exists a family $\mathcal{S}$ of subsets of $[n]$ such that (i) $|S| = \alpha n$ for each $S \in \mathcal{S}$, (ii) $|S\cap T| \le \beta n$ for every pair of distinct $S, T \in \mathcal{S}$ and (iii) $|\mathcal{S}| = \Omega((1/\alpha)^{\beta n})$.*

Let $\eta = C_1(d)N^{-\frac{1}{d-1}}$, where $C_1(d)$ is as in Lemma 3.2. It follows from Lemma 2.2 that we can take $n = c(d)N \le N/2$ points $p_1, \dots, p_n$ on some spherical cap $C(x, \sqrt{2-\eta})$ such that $\|p_i - p_j\|_2 \ge \eta$ for every $1 \le i < j \le n$. Since the p_i are chosen from a spherical cap of radius $\sqrt{2-\eta}$, we also have that $\|p_i + p_j\|_2 \ge \eta$ for all pairs $i \ne j$. Let $S = \{p_1, ..., p_n\}$. Applying Lemma 3.3, we can find $m = 2^{\Omega(n)}$ subsets $S_1, \dots, S_m$ of S such that $|S_i| = n/2$ for each i and $|S_i \cap S_j| \le n/4$ for every pair $i \ne j$.

Given the approximation parameter ε, let $N = c(d)\varepsilon^{-\frac{2(d-1)}{d+2p}}$ be such that $c_2 N^{-\frac{d+2p}{2(d-1)}} = 12\varepsilon$, where c_2 is as in Lemma 3.2.

Now we consider the following communication problem: suppose that Alice has one of the subsets $S_1, S_2, \dots, S_m$ and she wants to send a message to Bob, who needs to determine which subset Alice has. We shall show that we can solve this problem if we have an ℓ_p-subspace sketch data structure. Suppose that the subset

Alice has is S_t and Q is a data structure such that $Q(x) = (1 \pm \varepsilon)\|S_t x\|_p$ for all $x \in \mathbb{S}^{d-1}$. Then, Alice sends such a data structure Q to Bob.

On the one hand, we have for the subset $S_i = S_t$ that

$$|Q(x) - \|S_i x\|_p| \leq n\varepsilon, \quad \forall x \in \mathbb{S}^{d-1}. \tag{3.11}$$

On the other hand, if $S_t \neq S_i$, we know from the construction of the subsets S_i that $|S_i \setminus S_t| = |S_t \setminus S_i| \geq n/4$. Applying Lemma 3.2 to $S_i \cup (-S_i)$ and $S_t \cup (-S_t)$, we see that there exists an $x \in \mathbb{S}^{d-1}$ such that

$$|\|S_i x\|_p - \|S_t x\|_p| \geq \frac{1}{4} n \cdot c_2 N^{-\frac{d+2p}{2(d-1)}} \geq 3n\varepsilon,$$

whence it follows that

$$|Q(x) - \|S_i x\|_p| \geq |\|S_i x\|_p - \|S_t x\|_p| - |\|S_t x\|_p - Q(x)| \geq 3n\varepsilon - n\varepsilon = 2n\varepsilon. \tag{3.12}$$

Combining (3.11) and (3.12), we immediately see that Bob can decide which of the two cases he is in by querying a sufficient number of points on $\mathbb{S}^{d-1}$, which implies a lower bound of $\Omega(\log m) = \Omega(n) = \Omega(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ bits of space.

THEOREM 3.1. *Suppose that $p \in [1, \infty) \setminus 2\mathbb{Z}$ and d are constants. Any data structure that solves the ℓ_p-subspace sketch problem for dimension d and accuracy parameter ε requires $\Omega(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ bits of space.*

REMARK 3.1. *Our proof does not work when p is an even integer, which is as expected. In this case, $\|Ax\|_p^p$ is a polynomial and the spherical harmonic series will be finite. The term $\max_k(r^{2k}\lambda_k^{-2})$ in (3.8) will be a constant instead of a quantity depending on N.*

4 ℓ_p-Subspace Sketch Upper Bound

Our approach follows [Mat96], which deals with the case of $p = 1$. The following is our key lemma, which is an analogue of [Mat96, Proposition 7] for a general p.

LEMMA 4.1. *Let $P \in \mathbb{R}^d$ be an N-point set with each point having ℓ_2 norm $O(1)$, and let $w \in \Delta_{N-1}$ be an associated weight vector. There is an $O(d^{3/2}N^d)$-time deterministic algorithm which computes a subset $P' \subset P$ of at most $\frac{7}{8}N$ points with a weight vector $w' \in \Delta_{|P'|-1}$ such that with probability $1 - 1/N$, for every $x \in \mathbb{S}^{d-1}$*

$$\left| \sum_i w_i |\langle P_i, x\rangle|^p - \sum_i w_i' |\langle P_i', x\rangle|^p \right| = O(N^{-\frac{d+2p}{2(d-1)}} \sqrt{\log N}) .$$

Given matrix A and error parameter ε, we apply Lemma 4.1 repeatedly as in [Mat96], obtaining a sequence of subsets with $\frac{7}{8}N$, $(\frac{7}{8})^2 N$, ... points, until a subset of size $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ is obtained. Note that the errors of the successive approximations form a geometric progression, and hence the final error will be $O(\varepsilon)$. The following theorem follows immediately.

THEOREM 4.1. *Suppose that $A \in \mathbb{R}^{n\times d}$ satisfies that $\|A_i\|_2 = O(1)$ for all i and p is a positive integer constant. There is an $O(d^{3/2}n^d)$-time deterministic algorithm which computes a matrix B consisting of $m = O(\varepsilon^{-\frac{2(d-1)}{d+2p}} \log^{\frac{d-1}{d+2p}}(\varepsilon^{-1}))$ rows of A, and an associated weight vector $w \in \Delta_{m-1}$, such that with high probability for every $x \in \mathbb{S}^{d-1}$,*

$$\left| \sum_{i=1}^m w_i |\langle B_i, x\rangle|^p - \frac{1}{n} \|Ax\|_p^p \right| = O(\varepsilon) .$$

To prove Lemma 4.1, we first give intuition, which is inspired by the proof of [Mat96, Proposition 7]. Given a point set P, from Lemma 2.3 we know that for at least half of the points of P, we can divide them into groups $P_1, P_2, \ldots, P_t$ which satisfy the corresponding conditions (for the points that are not on the sphere, we scale them to unit vectors when applying the partition lemma). That is, (i) each P_i is within a small area and (ii) for every $x \in \mathbb{S}^{d-1}$, the hyperplane $\{y : \langle x, y\rangle = 0\}$ only intersects a small number of sets among P_i.

Given a query point $x \in \mathbb{S}^{d-1}$, let H be the hyperplane $\{y : \langle x, y\rangle = 0\}$. For the subset P_i, we first consider the case that H does not intersect P_i. In this case we have

$$\sum_{y\in P_i} |\langle x, y\rangle|^p = \sum_{y\in P_i} \langle x, y\rangle^p \ .$$

The tensor product trick $\langle x, y\rangle^p = \langle x^{\otimes p}, y^{\otimes p}\rangle$ implies that

$$\sum_{y\in P_i} |\langle x, y\rangle|^p = \sum_{y\in P_i} \langle x^{\otimes p}, y^{\otimes p}\rangle = \langle x^{\otimes p}, \sum_{y\in P_i} y^{\otimes p}\rangle \ .$$

Hence, what we need to store is $\sum_{y\in P_i} y^{\otimes p}$, which can be seen as a d^p-dimensional vector.

For the other case where H intersects P_i, the method above will not work. However, note that in this case $|\langle x, y\rangle|$ is small for every $y \in P_i$ as each P_i lies in a small region. We also know that H intersects only a small number of groups P_i; therefore, sampling points from the intersection is enough to achieve an ε-additive error.

However, one issue here is that it is not easy to determine whether H intersects P_i if we only store a limited number of points. To address this, for each subset P_i, we choose a random subset $T \subseteq P_i$ with associated weights $\{w_y\}_{y\in T}$ such that (i) $\sum_{y\in T} w_y \cdot y^{\otimes p} = \sum_{y\in P_i} y^{\otimes p}$ (which is helpful for the first case) and (ii) $\mathbb{E} \sum_{y\in T} w_y \cdot |\langle x, y\rangle|^p = \sum_{y\in P_i} \langle x, y\rangle^p$ (which we will show is enough for the second case). The following lemma will be useful.

Lemma 4.2. ([Mat96], Lemma 8) *Let $P = \{P_1, \ldots, P_s\} \subset \mathbb{R}^d$ be a set of $s \geq d + 1$ points with an associated weight vector $u \in \Delta_{s-1}$. There is a deterministic algorithm which computes a group of the subsets $T_1, \ldots, T_{s'}$, each associated with a probability p_i and a weight vector $w_i \in \Delta_{t_i-1}$, where $s' \leq s - d$ and $t_i = |T_i|$, such that*

- *(i)* $\sum_{i=1}^{s'} p_i = 1$;
- *(ii)* $t_i \leq d + 1$ *for each* $i \in [s']$;
- *(iii)* $\sum_{j=1}^{t_i} w_{i,j} T_{i,j} = \sum_{i=1}^{s} u_i P_i$ *for each* $T_i = \{T_{i,j}\}_{j=1,\ldots,t_i}$;
- *(iv)* $\sum_{i\in I_k} p_i w_{i,j(i,k)} = u_k$ *for each* $k \in [s]$, *where* $I_k = \{i \in [s'] : P_k \in T_i\}$ *and* $j(i, k)$ *is the index* j *such that* $T_{i,j} = P_k$ *for* $i \in I_k$.

Furthermore, the running time of the algorithm is dominated by that of finding a feasible solution to a linear program with s variables and $d + 1$ constraints.

We now are ready to prove our Lemma 4.1.

Proof. [Proof of Lemma 4.1] Since $\sum_i u_i = 1$, there are at least $\frac{1}{2}N$ points with weight $w_i \leq \frac{2}{N}$. Applying Lemma 2.3 to these points with $s = 2(d^p + 1)$ (recall that d and p are both constants), we obtain a collection of disjoint s-points subsets $P_1, ..., P_t$, each of which has diameter $O(N^{\frac{1}{d-1}})$. Furthermore, for every $x \in \mathbb{S}^{d-1}$, the corresponding hyperplane $\{y : \langle x, y\rangle = 0\}$ intersects at most $O(N^{1-\frac{1}{d-1}})$ subsets. For the remaining points, we keep them with the same weights. For each P_i, we sample at most half of the points for each group as below.

For each $x \in \mathbb{R}^d$, let $T(x) = x^{\otimes p}$ denote its p-th tensor product, viewed as a d^p-dimensional vector. For each subset P_i, let $T(P_i) = \{T(x) : x \in P_i\}$ and $w(P_i) = \sum_{y\in P_i} w_y$. We then apply Lemma 4.2 to $T(P)$ with the normalized weights $w'_y = w_y/w(P_i)$, obtaining a group of the subsets $T_1, \ldots, T_{s'}$ with weights $v_1, \ldots, v_{s'}$ and probabilities $p_1 \ldots, p'_s$. We choose a random index $j \in \{1, 2, \ldots, s'\}$ according to probabilities $p_1, \ldots, p'_s$ and take the subset T_j to be our sample set of the points. The guarantee of Lemma 4.2 implies that T_j contains at most $d^p + 1$ points, which is at most half of the size of P_i. Repeating this procedure for each P_i, we finally obtain a subset of P containing at most $\frac{7}{8}N$ points.

Now we analyze correctness of our algorithm. Fix $x \in \mathbb{S}^{d-1}$. For each P_i, let Q_i denote the subsets of P_i with the points we sampled in the procedure above with the associated weight v_i. It suffices to show that

$$\left| \sum_i \sum_{y\in Q_i} v_y |\langle x, y\rangle|^p - \sum_i \sum_{y\in P_i} w_y |\langle x, y\rangle|^p \right| = O(N^{-\frac{d+2p}{2(d-1)}} \sqrt{\log N})$$

holds with high probability.

Let I denote the set of indices in $\{1, 2, \ldots, t\}$ for which the hyperplane $H = \{y : \langle x, y\rangle = 0\}$ intersects P_i. Then for each $i \notin I$, from the condition

$$\sum_{y \in Q_i} v_y y^{\otimes p} = \sum_{y \in P_i} w_y y^{\otimes p}$$

we have that

$$\langle x^{\otimes p}, \sum_{y \in Q_i} v_y y^{\otimes p}\rangle = \sum_{y \in Q_i} v_y \langle x, y\rangle^p = \sum_{y \in P_i} w_y \langle x, y\rangle^p = \langle x^{\otimes p}, \sum_{y \in P_i} w_y y^{\otimes p}\rangle \,.$$

Note that in this case $\langle x, y\rangle$ has the same sign for all $y \in P_i$. Hence

$$\sum_{y \in Q_i} v_y |\langle x, y\rangle|^p = \sum_{y \in P_i} w_y |\langle x, y\rangle|^p$$

Now we focus on the case $i \in I$. Recall that $|I| = O(N^{1-\frac{1}{d-1}})$. Define a random variable X_i as

$$X_i = \sum_{y \in Q_i} v_y |\langle x, y\rangle|^p - \sum_{y \in P_i} w_y |\langle x, y\rangle|^p \,,$$

where the randomness is taken over the choice of the subsets Q_i. Lemma 4.2(iii) implies that

$$\mathbb{E}\left[\sum_{y \in Q_i} v_y |\langle x, y\rangle|^p\right] = \sum_{y \in P_i} w_y |\langle x, y\rangle|^p \,.$$

Hence $\mathbb{E}\, X_i = 0$. Since H intersects P_i and the diameter of P_i is $O(N^{-\frac{1}{d-1}})$, it holds that $|\langle x, y\rangle|^p = O(N^{-\frac{p}{d-1}})$ for all $y \in Q_i$. Recalling our definitions of v_y and s, we have $\sum_{y \in Q_y} v_y = w(P_i) = O(s/N)$ and thus $|X_i| = O(N^{-1-\frac{p}{d-1}})$. Let U be this upper bound for $|X_i|$. It then follows from Bernstein's inequality that

$$\Pr\left[\left|\sum_{i \in I} X_i\right| > \lambda U \sqrt{|I|}\right] < 2e^{-\lambda^2/2},$$

where $U\sqrt{|I|} = O(N^{-\frac{d+2p}{2(d-1)}})$. Taking $\lambda \sim \sqrt{\log N}$, we obtain that for a fixed x,

$$\left|\sum_i \sum_{y \in Q_i} v_y |\langle x, y\rangle|^p - \sum_i \sum_{y \in P_i} w_y |\langle x, y\rangle|^p\right| = O(N^{-\frac{d+2p}{2(d-1)}} \sqrt{\log N}) \tag{4.13}$$

with probability at least $1 - 1/\operatorname{poly}(N)$.

Next we do a net argument. Choose a constant $D > \frac{d+2p}{2(d-1)}$ and let $\gamma = N^{-D}$. We can take a γ-net $\mathcal{N}$ on $\mathbb{S}^{d-1}$ with $O(\gamma^{-(d-1)})$ points. By a union bound, we have that with probability at least $1 - 1/N$, the bound (4.13) holds for all $x \in \mathcal{N}$ simultaneously. Note that for two $x, x' \in \mathbb{S}^{d-1}$, if $\|x - x'\|_2 \le \gamma$, then $||\langle x, y\rangle|^p - |\langle x', y\rangle|^p| = O(\gamma)$. It follows that with probability at least $1 - 1/N$, the target bound (4.13) holds for all $x \in \mathbb{S}^{d-1}$ simultaneously, which is what we need.

We now analyze the time complexity of our algorithm. The first step is to compute the subsets $P_1, \ldots, P_t$, which can be done in $O(N^{d-\frac{1}{d-1}})$ time by Lemma 2.3. Then, for each subset P_i, by Lemma 4.2, the subset T and the associated weights w can be computed in $\widetilde{O}(d^{\frac{3}{2}} s) = \widetilde{O}(d^{\frac{3}{2}+p})$ time (for instance, using the LP algorithm in [LS15]). Therefore, the total runtime of the algorithm is $O(d^{3/2} N^d)$, as claimed. ∎

Achieving $(1 \pm \varepsilon)$-Multiplicative Error. Consider a general matrix $A \in \mathbb{R}^{n \times d}$. Without loss of generality, assume that A has full column rank.

Consider the John ellipsoid of $Z(A) := \{x \in \mathbb{R}^d : \|Ax\|_p \le 1\}$. There exists an invertible linear transformation $T : \mathbb{R}^d \to \mathbb{R}^d$ such that $B_{\ell_2}^d \subseteq Z(AT) \subseteq \sqrt{d} B_{\ell_2}^d$. Let $A' = AT$. Then $1/\sqrt{d} \le \|A'x\|_p \le 1$ when $x \in \mathbb{S}^{d-1}$. Let $A_i'' = A_i' / \|A_i'\|_2$. Then $A_i'' \in \mathbb{S}^{d-1}$. Note that

$$1 \ge \|A'x\|_p^p = \sum_i \|A_i'\|_2^p \left|\langle A_i'', x\rangle\right|^p, \quad \forall x \in \mathbb{S}^{d-1}.$$

Integrate both sides over $\mathbb{S}^{d-1}$ w.r.t. x. Observe that $\int_{\mathbb{S}^{d-1}} |\langle u, x\rangle|^p \, d\sigma_{d-1}(x)$ is a constant, independent of u, whenever $u \in \mathbb{S}^{d-1}$. Denote this constant by $\beta_{d,p}$. It follows that

$$1 \ge \sum_i \|A_i'\|_2^p \cdot \beta_{d,p},$$

and thus $\|A'x\|_p^p \ge 1/d^{p/2} \ge (\beta_{d,p}/d^{p/2}) \sum_i \|A_i'\|_2^p$ for all $x \in \mathbb{S}^{d-1}$.

Next, define weights $w_i' = w_i \|A_i'\|_2^p$. Then

$$\sum_i w_i \left|\langle A_i', x\rangle\right|^p = \sum_i w_i' \left|\langle A_i'', x\rangle\right|^p.$$

Suppose that $w_i = 1/n$ for all i. We apply Lemma 4.1 to A'' with weights $w_i'' = w_i' / \sum_j w_j'$ and obtain a subset of points W with weights v_i, where $|W| = O(\varepsilon^{-\frac{2(d-1)}{d+2p}})$, such that with high probability,

$$\left|\sum_i v_i |\langle W_i, x\rangle|^p - \frac{1}{n\sum_j w_j'} \|A'x\|_p^p\right| = O(\varepsilon), \quad \forall x \in \mathbb{S}^{d-1}. \tag{4.14}$$

Recall that we showed above that $\|A'x\|_p^p \ge c(d,p) \cdot n \sum_j w_j' = c(d,p) \sum_i \|A_i'\|_2^p$ for all $x \in \mathbb{S}^{d-1}$, where $c(d,p) > 0$ is some constant depending only on d and p.

Let $v_i' = n v_i \sum_j w_j'$. It follows from (4.14) that

$$\left|\sum_i v_i' |\langle W_i, x\rangle|^p - \|A'x\|_p^p\right| = O(\varepsilon) \|A'x\|_p^p, \quad \forall x \in \mathbb{S}^{d-1}.$$

Therefore,

$$\left|\sum_i v_i' |\langle (T^{-1})^\top W_i, Tx\rangle|^p - \|ATx\|_p^p\right| = O(\varepsilon) \cdot \|ATx\|_p^p, \quad \forall x \in \mathbb{R}^d.$$

This implies that the rows of WT^{-1} (which form a row-subsampled matrix of A) with weights v_i' are exactly what we need. Our final theorem is now immediate.

THEOREM 4.2. *Suppose that $A \in \mathbb{R}^{n \times d}$ and p is a positive integer constant. There is a polynomial-time algorithm which computes a subset B of $m = \widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ rows of A and an associated weight vector $w \in \mathbb{R}^m$, such that with high probability for every $x \in \mathbb{S}^{d-1}$,*

$$\left|\sum_{i=1}^m w_i |\langle B_i, x\rangle|^p - \|Ax\|_p^p\right| = O(\varepsilon) \cdot \|Ax\|_p^p.$$

Proof. First, we show that a constant-factor approximation of the John ellipsoid of $Z(A)$ can be found in polynomial time. We can compute in polynomial time a decomposition $A = UT$, where $U \in \mathbb{R}^{n \times d}$ has orthonormal columns and $T \in \mathbb{R}^{d \times d}$ is an invertible matrix. It then suffices to find the John ellipsoid of $Z(U)$. It is clear that $B(0,r) \subseteq Z(U) \subseteq B(0,R)$ for $r = 1/n$ and $R = n^{\max\{1/2-1/p,0\}}$ and that a separation oracle for $Z(U)$ can be implemented in polynomial time. Then an ellipsoid E satisfying $E \subseteq Z(U) \subseteq \sqrt{d(d+1)}E$ can be found in polynomial time via the shallow-cut ellipsoid method [GLS88, Theorem 4.6.3]. Therefore, $T^{-1}E \subseteq Z(A) \subseteq \sqrt{d(d+1)}T^{-1}E$, that is, $T^{-1}E$ is a constant-factor approximation of the John ellipsoid of $Z(A)$.

Correctness then follows from the discussion preceding the theorem statement (which goes through with a constant-factor approximation of John's ellipsoid) and Theorem 4.1, which also implies that the remaining procedure can be completed in polynomial time. ∎

5 ℓ_p-Subspace Sketch in a One-Pass Stream

In this section, we implement our algorithm in the previous section in a one-pass stream, where each row arrives one at a time. We assume that each entry of matrix A can be saved in $\log(n)$ bits of space. We first show that a coreset can be constructed in linear space when the number of rows is not large, which will be used as a subroutine in our full algorithm. Then we present the full algorithm, which is based on the standard merge-and-reduce paradigm and uses an additional factor of $\mathrm{polylog}(n)$ space. Finally, we show that the $\mathrm{polylog}(n)$ factor can be reduced to $\mathrm{poly}(\log\log n)$. At the end of this section, we shall give another $O(\varepsilon)$-additive error streaming algorithm with $O(1)$ update time, with a slightly worse space complexity $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ for general $d = O(1)$ and the near-tight bound $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ for $d \le 2p+2$.

5.1 Basic Step: Coreset Our algorithm is based on the following lemma.

LEMMA 5.1. *Suppose that p is a positive integer constant, $A \in \mathbb{R}^{n\times d}$ and $w \in \Delta_{n-1}$ is the associated weight vector. There is an algorithm* CORESET(A, w, ε) *which computes a subset B of $m = \widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ rows of A and a weight vector $v \in \Delta_{m-1}$, such that with high probability, it holds for every $x \in \mathbb{S}^{d-1}$,*

$$\left|\sum_i v_i |\langle B_i, x\rangle|^p - \sum_i w_i |\langle A_i, x\rangle|^p\right| = O(\varepsilon)\cdot\left(\sum_i w_i |\langle A_i, x\rangle|^p\right).$$

Furthermore, the algorithm can be implemented in $O(n)$ space.

Proof. The ellipsoid method employed in the proof of Theorem 4.2 operates in $\mathbb{R}^d$ with $\mathrm{poly}(d)$ space, except for the separation oracle which can be clearly implemented in $O(n)$ space. Therefore, the linear transformation that normalizes the John ellipsoid of $Z(A)$ can be computed in $O(n)$ space. Recall that we only need a constant-factor approximation. Hence, we can assume that after the linear transformation, each entry can still be stored in $O(\log n)$ bits of space.

Since we only need to invoke Lemma 4.1 a total of $\log n$ times, it suffices to show that $O(N)$ space is enough in each invocation, where N is the number of rows of the input. From the proof of Lemma 4.1, there exist desired groups $P_1, \ldots, P_t$, which can be found in $O(N)$ space by enumeration. Each subset P_i corresponds to only $O(s)$ rows, and thus the weights can be computed in $\mathrm{poly}(s)$ space. Thus, each run of Lemma 4.1 can be implemented in $O(N)$ space. This finishes the proof. ◻

5.2 Merge and Reduce Given the coreset procedure (Lemma 5.1), the general coreset framework in [BDM+20] can be readily applied, leading to an algorithm similar to Algorithm 6 therein. For completeness, a full description of our algorithm is presented in Algorithm 1. We maintain a number of blocks $\mathbf{B}_0, \mathbf{B}_1, \ldots$, each of size $\mathsf{m}_{\mathsf{space}}$ (which is determined by the coreset size). The most recent rows are stored in $\mathbf{B}_0$; whenever $\mathbf{B}_0$ is full, the successive non-empty blocks $\mathbf{B}_0, \ldots, \mathbf{B}_i$ are merged and reduced to a new coreset, which will be stored in $\mathbf{B}_{i+1}$. Since there are n data points in the stream, the next lemma guarantees that maintaining $(\log n + 1)$ blocks $\mathbf{B}_0, \ldots, \mathbf{B}_{\log n}$ suffices.

LEMMA 5.2. ([BDM+20]) *Suppose that $\mathbf{B}_0, \ldots, \mathbf{B}_{i-1}$ are all empty while $\mathbf{B}_i$ is non-empty. Then $\mathbf{B}_i$ with the associated weight vector w_i is a $(1 + \frac{\varepsilon}{\log n})^i$-coreset for the last $2^{i-1}\mathsf{m}_{\mathsf{space}}$ rows.*

THEOREM 5.1. *Let $A = a_1 \circ \cdots \circ a_n$ be a stream of n rows, where $a_i \in \mathbb{R}^{1\times d}$. There is an algorithm which computes a subset B of $m = \widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ rows of A and an associated weight vector $w \in \mathbb{R}^m$, such that with high probability, for every $x \in \mathbb{S}^{d-1}$,*

$$\left|\sum_i w_i |\langle B_i, x\rangle|^p - \|Ax\|_p^p\right| = O(\varepsilon)\cdot\|Ax\|_p^p.$$

Moreover, the algorithm can be implemented in $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}}\cdot \mathrm{polylog}(n))$ words of space.

Algorithm 1: Merge-and-reduce framework for the ℓ_p subspace sketch for constant d.

Input: A stream of rows $a_1, a_2, \ldots, a_n \in \mathbb{R}^{1\times d}$ with weights $u_1, \ldots, u_n$, and approximation factor ε;
Output: A coreset B with associated weights w ;
Initialize blocks $\mathbf{B}_0, \mathbf{B}_1, \ldots, \mathbf{B}_{\log n} \leftarrow \emptyset$, $\mathsf{m}_{\mathsf{space}} \leftarrow \widetilde{O}(\gamma^{-\frac{2(d-1)}{d+2p}})$ where $\gamma = \frac{\varepsilon}{\log n}$;
foreach *row* a_t *and weight* u_t **do**
  if $\mathbf{B}_0$ *does not contain* $\mathsf{m}_{\mathsf{space}}$ *rows* **then**
    $\mathbf{B}_0 \leftarrow a_t \circ \mathbf{B}_0$;
  else
    Let $i > 0$ be the minimal index such that $\mathbf{B}_i = \emptyset$;
    $\mathbf{B}_i, w_i \leftarrow \textsc{Coreset}\left(\mathbf{M}, w, \frac{\varepsilon}{\log n}\right)$, where $\mathbf{M} = \mathbf{B}_0 \circ \cdots \circ \mathbf{B}_{i-1}$ and $w = w_0 \circ \cdots \circ w_{i-1}$;
    for $j = 0$ **to** $j = i-1$ **do**
      $\mathbf{B}_j \leftarrow \emptyset$;
    end
    $\mathbf{B}_0 \leftarrow a_t$, $w_0 \leftarrow u_t$;
  end
end
$\mathbf{B}^*, w^* \leftarrow \textsc{Coreset}(\mathbf{B}_{\log n} \circ \cdots \circ \mathbf{B}_0, w_{\log n} \circ \cdots \circ w_0, \varepsilon)$;
return $\mathbf{B}^*$ and w^*

Proof. We assume that each call to the subroutine $\textsc{Coreset}(\mathbf{M}, w, \frac{\varepsilon}{\log n})$ is successful, which holds with high probability after taking a union bound. It then follows from Lemma 5.2 that $\mathbf{B}_{\log n} \circ \cdots \circ \mathbf{B}_0$ is a $(1 + \frac{\varepsilon}{\log n})^{\log n} = (1 + O(\varepsilon))$-coreset of the rows of A and, after a further coreset operation, $\mathbf{B}^*$ is a $(1 + O(\varepsilon))(1 + \varepsilon) = (1 + O(\varepsilon))$-coreset of A, as desired.

Now we analyze the space complexity of our algorithm. Let $\gamma = \frac{\varepsilon}{\log n}$. The algorithm stores $O(\log n)$ blocks $\mathbf{B}_i$, each taking at most $O(\mathsf{m}_{\mathsf{space}} \cdot d) = \widetilde{O}(\gamma^{-\frac{2(d-1)}{d+2p}})$ words of space. Lemma 5.1 implies that each call to the subroutine $\textsc{Coreset}(\mathbf{M}, w, \gamma)$ takes $O(\mathsf{m}_{\mathsf{space}} \cdot \log n)$ words of space, since the number of total rows in the input does not exceed $O(\mathsf{m}_{\mathsf{space}} \cdot \log n)$. Therefore, the total space of our algorithm is $O(\mathsf{m}_{\mathsf{space}} \cdot \log n) = O(\varepsilon^{-\frac{2(d-1)}{d+2p}} \cdot \mathrm{polylog}(n/\varepsilon))$. ∎

5.3 Reducing Space Complexity In this section, we show that the multiplicative $\mathrm{poly}(\log n)$ factor in the space complexity can be reduced to $\mathrm{poly}(\log \log n)$ at the cost of an extra additive $\mathrm{poly}(\log n)$ term. The basic idea is that if we can sample a few of the rows which form a good approximation to the matrix A, then we can then treat the sampled rows as a new stream and the new input to our Algorithm 1. To this end, we consider the following ℓ_p sensitivities of the rows of A.

DEFINITION 5.1. *For a matrix $A = a_1 \circ \cdots \circ a_n \in \mathbb{R}^{n\times d}$, the ℓ_p-sensitivity of a_i, denoted by $s_i(A)$, is defined to be $s_i(A) = \sup_{x\in\mathbb{R}^d\setminus\{0\}} \frac{|\langle a_i, x\rangle|}{\|Ax\|_p^p}$.*

The following lemma shows that sampling according to the ℓ_p-sensitivities of A gives a subspace embedding of A. This is a generalization of the ℓ_1-sensitivity sampling [BDM+20, Lemma 4.4] and the proof follows the same approach.

LEMMA 5.3. *Let $A \in \mathbb{R}^{n\times d}$ and $1 \le p < \infty$. The matrix B is a submatrix of A such that the rescaled i-th row $p_i^{-1/p} a_i$ is included in B with probability $p_i \ge \min(\beta s_i(A), 1)$. Then, there is a constant c such that when $\beta \ge c\varepsilon^{-2} d \log(1/\varepsilon)$, the matrix B is a $(1 \pm \varepsilon)$-subspace embedding of A with probability at least 9/10.*

Proof. Since the row space of B is contained in that of A, in order to show that B is a $(1\pm\varepsilon)$-subspace embedding of A, it suffices to show that $\|Bx\|_p = 1 \pm \varepsilon$ for all x such that $\|Ax\|_p = 1$.

Fix an x and let $y = Ax$. Define the random variable Z_i to be $|y_i|^p/p_i$ with probability p_i and 0 otherwise, so $\mathbb{E}\, Z_i = |y_i|^p$. Let $Z = \|Bx\|_p^p$. Then $Z = \sum_i Z_i$ and

$$\mathbb{E}[Z] = \sum_i \mathbb{E}[Z_i] = \|y\|_p^p \ .$$

We also have

$$\mathbf{Var}[Z_i] \le \mathbb{E}[Z_i^2] = \frac{|y_i|^{2p}}{p_i^2} \cdot p_i = \frac{|y_i|^{2p}}{p_i}$$

and

$$Z_i \le \frac{|y_i|^p}{p_i} \le \frac{1}{\beta} \frac{|y_i|^p \|y\|_p^p}{|y_i|^p} = \frac{\|y\|_p^p}{\beta} \ ,$$

so

$$\mathbf{Var}[Z] = \sum_i \mathbf{Var}[Z_i] \le \sum_{i=1}^n \frac{|y_i|^{2p}}{p_i} \le \frac{1}{\beta} \sum_{i=1}^n |y_i|^p \frac{|y_i|^p \|y\|_p^p}{|y_i|^p} = \frac{\|y\|_p^{2p}}{\beta} \ .$$

It follows from Bernstein's inequality that

$$\Pr\left[|Z - \mathbb{E}[Z]| \ge \varepsilon \|y\|_p^p\right] \le 2 \exp\left(-\frac{\beta}{2} \frac{\varepsilon^2 \|y\|_p^{2p}}{\|y\|_p^{2p} + \varepsilon \|y\|_p^{2p}/3}\right) = \varepsilon^{-\Omega(d)} \ .$$

Rescaling ε by a constant factor (depending on p), we have that $\|Bx\|_p = (1 \pm \varepsilon)\|Ax\|_p$ with probability at least $1 - \varepsilon^{\Omega(d)}$ for each fixed x.

We next need a net argument. Let $\mathcal{S} = \{Ax : x \in \mathbb{R}^d, \|Ax\|_p = 1\}$ be the unit ℓ_p ball and $\mathcal{N}$ be a net of size $(3/\varepsilon)^d$ of $\mathcal{S}$ under the ℓ_p distance. By a union bound, we have that $\|Bx\|_p = (1 \pm \varepsilon)\|Ax\|_p$ for every $Ax \in \mathcal{N}$ simultaneously with probability at least $9/10$. Conditioned on this event, for each $y = Ax \in \mathcal{S}$, we choose a sequence of points $y_0, y_1, \dots \in \mathcal{S}$ as follows.

- Choose $y_0 \in \mathcal{S}$ such that $\|y - y_0\|_p \le \varepsilon$ and let $\alpha_0 = 1$;
- After choosing $y_0, y_1, \dots, y_i$, we choose y_{i+1} such that

$$\left\|\frac{y - \alpha_0 y_0 - \alpha_1 y_1 - \cdots - \alpha_i y_i}{\alpha_{i+1}} - y_{i+1}\right\|_p \le \varepsilon,$$

  where $\alpha_{i+1} = \|y - \alpha_0 y_0 - \alpha_1 y_1 - \cdots - \alpha_i y_i\|_p$.

The choice of y_{i+1} means that

$$\alpha_{i+2} = \|y - \alpha_0 y_0 - \alpha_1 y_1 - \cdots - \alpha_i y_i - \alpha_{i+1} y_{i+1}\|_p \le \alpha_{i+1}\varepsilon.$$

A simple induction yields that $\alpha_i \le \varepsilon^i$. Hence

$$y = y_0 + \sum_{i\ge1} \alpha_i y_i, \quad |\alpha_i| \le \varepsilon^i \ .$$

Suppose that $y_i = Ax_i$, we have

$$\|Bx\|_p \le \|Bx_0\|_p + \sum_{i\ge1} \epsilon^i \|Bx_i\|_p \le (1+\epsilon) + \sum_{i\ge1} \epsilon^i (1+\epsilon) = 1 + O(\epsilon),$$

and

$$\|Bx\|_p \ge \|Bx_0\|_p - \sum_{i\ge1} \epsilon^i \|Bx_i\|_p \ge (1-\epsilon) - \sum_{i\ge1} \epsilon^i (1-\epsilon) = 1 - O(\epsilon).$$

Rescaling ε again gives the result. □

Since the rows of A are given in the streaming model, we shall sample according to the online ℓ_p sensitivities of the rows, which are defined below.

Definition 5.2. (Online ℓ_p Sensitivities [BDM+20, WY22]) *Let $A \in \mathbb{R}^{n\times d}$ and let $1 \le p < \infty$. Then, for each $i \in [n]$, the i-th online ℓ_p sensitivity is defined as*

$$s_i^{\mathsf{OL}}(A) := \begin{cases} \min\left\{\sup_{x\in\mathrm{rowspace}(A)\setminus\{0\}} \frac{|\langle a_i, x\rangle|^p}{\|A_{i-1}x\|_p^p}, 1\right\} & a_i \in \mathrm{rowspace}(A_{i-1}) \\ 1 & \textit{otherwise}, \end{cases}$$

where $A_j \in \mathbb{R}^{j\times d}$ denotes the submatrix of A formed by the first j rows.

It is clear from the definition that the online ℓ_p sensitivity is at least as large as the ℓ_p-sensitivity of the same row and we can thus use the online ℓ_p-sensitivities in an online algorithm to achieve the ℓ_p-subspace embedding property. Our full algorithm is given in Algorithm 2.

Algorithm 2: Algorithm for ℓ_p subspace sketch for constant d.

Input: A streaming of rows $a_1, a_2, \ldots, a_n \in \mathbb{R}^{1\times d}$ and approximation factor γ;
ALG_1 is an instance of Algorithm 1 with $\varepsilon = O(1)$;
ALG_2 is an instance of Algorithm 1 with $\varepsilon = \gamma$;
$\beta \leftarrow \mathrm{poly}(d)\log(1/\varepsilon)/\varepsilon^2$;
foreach *row* a_t **do**
  $M_t \leftarrow$ a coreset of A_{t-1} from ALG_1;
  $\tau_t \leftarrow$ a $\mathrm{poly}(d)$-approximation of $s_t^{\mathsf{OL}}(A)$ from a_t and M_t;
  $p_t \leftarrow \min(1, \beta\tau_t)$;
  With probability p_t, feed $a_t/p_t^{1/p}$ to ALG_2;
  Feed a_t to ALG_1;
end
return a coreset M from ALG_2.

It follows from Theorem 5.1 and Lemma 5.3 that the output M is a $(1\pm\varepsilon)$-coreset of A with high probability. To bound the space complexity of our algorithm, we need to bound the sum of the online ℓ_p sensitivities of the rows a_i, which is given below in Lemma 5.4. It follows that the expected number of sampled rows is $O(\varepsilon^{-2}\,\mathrm{poly}(d)\log(1/\varepsilon)\log n) = O(\varepsilon^{-2}\log(1/\varepsilon)\log n)$.

Lemma 5.4. ([BDM+20, Lemma 4.7], [WY22, Theorem 3.10]) *Suppose that $A \in \mathbb{R}^{n\times d}$ and $p \in \{1\}\cup[2,\infty)$. Let $q = 1$ when $p = 1$ or $q = p/2$ when $p \ge 2$, and κ be the condition number of A. Then, we have*

$$\sum_{i=1}^n s_i^{\mathsf{OL}}(A) = O(d^q(\log n)\log^q \kappa).$$

Moreover, if $A \in \mathbb{Z}^{n\times d}$ is the integer matrix with entries bounded by $\mathrm{poly}(n)$*, we have*

$$\sum_{i=1}^n s_i^{\mathsf{OL}}(A) = O(d^q \log^{q+1} n).$$

The only thing remaining is to compute a $\mathrm{poly}(d)$-approximation of s_i from a_i and M_i. Since $\|A_{i-1}x\|_p = (1\pm\varepsilon)\|M_i x\|_p$ for all x, it must hold that $(\mathrm{rowspace}(A_{i-1}))^\perp = (\mathrm{rowspace}(M_i))^\perp$ and so $\mathrm{rowspace}(A_{i-1}) = \mathrm{rowspace}(M_i)$. Hence, we can use M_i to determine whether $a_i \in \mathrm{rowspace}(A_{i-1})$. When $a_i \in \mathrm{rowspace}(A_{i-1})$, we need an efficient algorithm to find τ_t, for which we consider a well-conditioned basis of A_{i-1}, defined below.

Definition 5.3. (Well-conditioned basis, [DDH+09]) *Suppose that $A \in \mathbb{R}^{n\times d}$ has rank r and $p \ge 1$. An $n\times r$ matrix U is an (α,β,p)-well-conditioned basis for A if (i)* $\mathrm{colspace}(U) = \mathrm{colspace}(A)$*, (ii) $\|U\|_p \le \alpha$ and (iii) $\|z\|_q \le \beta\|Uz\|_p$ for all $z \in \mathbb{R}^d$, where $q = p/(p-1)$ is the conjugate index of p.*

THEOREM 5.2. ([DDH+09]) *Let A be an $n \times d$ matrix of rank r, $p \in [1, \infty)$ and q be the conjugate index of p. There exists an (α, β, p)-well-conditioned basis U for the column space of A such that:*

(1) if $p < 2$ then $\alpha = r^{\frac{1}{p}+\frac{1}{2}}$ and $\beta = 1$,

(2) if $p = 2$ then $\alpha = \sqrt{d}$ and $\beta = 1$, and

(3) if $p > 2$ then $\alpha = r^{\frac{1}{p}+\frac{1}{2}}$ and $\beta = r^{\frac{1}{p}-\frac{1}{2}}$.

Moreover, there is a deterministic procedure that computes a decomposition $A = UT$, where $U \in \mathbb{R}^{n\times r}$ is a well-conditioned basis as described above and $T \in \mathbb{R}^{r\times d}$ is of full row rank, in time $O(ndr + nd^5 \log n)$ for $p \neq 2$ and $O(ndr)$ if $p = 2$.

Suppose that A_{i-1} has a decomposition $A_{i-1} = UT$ as in Theorem 5.2 and assume that $a_i \in \operatorname{rowspace}(A_{i-1}) = \operatorname{rowspace}(T)$. Since T has full row rank, it has a right inverse $T^\dagger \in \mathbb{R}^{d\times r}$ such that $TT^\dagger = I$ and $\operatorname{colspace}(T^\dagger) = \operatorname{rowspace}(T)$. Now, the online ℓ_p sensitivity of a_i becomes

$$\sup_{x\in\mathbb{R}^r\setminus\{0\}} \frac{|\langle a_i, T^\dagger x\rangle|^p}{\|A_{i-1}T^\dagger x\|_p^p} = \sup_{x\in\mathbb{R}^r\setminus\{0\}} \frac{|\langle b, x\rangle|^p}{\|Ux\|_p^p} = \sup_{x\in\mathbb{S}^{r-1}} \frac{|\langle b, x\rangle|^p}{\|Ux\|_p^p} ,$$

where $b = (T^\dagger)^\top a_i$. The definition of the well-conditioned basis indicates that $1/\operatorname{poly}(d) \le \|Ux\|_p^p \le \operatorname{poly}(d)$. Suppose that x_0 is the vector that attains the supremum. Then $|\langle b, x_0\rangle|^p \le |x_0|_2^p \cdot \|b\|_2^p = \|b\|_2^p$, implying that $s_i^{\mathsf{OL}}(A) \le \operatorname{poly}(r) \cdot \|b\|_2^p$. On the other hand, taking $x = b/\|b\|_2$ leads to $s_i^{\mathsf{OL}}(A) \ge \|b\|_2^p / \operatorname{poly}(r)$. Therefore, $\|b\|_2^p$ is a $\operatorname{poly}(d)$-approximation to $s_i^{\mathsf{OL}}(A)$. Putting everything together, we obtain the following theorem.

THEOREM 5.3. *Let $A = a_1 \circ \cdots \circ a_n$ be a stream of n rows, where $a_i \in \mathbb{R}^{1\times d}$. There is an algorithm which computes a subset B of $m = \widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ rows of A and an associated weight vector $w \in \mathbb{R}^m$ such that with high probability for every $x \in \mathbb{S}^{d-1}$,*

$$\left|\sum_i w_i |\langle B_i, x\rangle|^p - \|Ax\|_p^p\right| = O(\varepsilon) \cdot \|Ax\|_p^p .$$

Moreover, the algorithm can be implemented in $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}} + \log^{\frac{3d+2p-2}{d+2p}} n)$ words of space.

5.4 Faster Update Time In this section, we give a different one-pass streaming algorithm which achieves $O(\varepsilon)$-additive error with $O(1)$ update time, assuming that $\|A_i\|_2 = O(1)$. The basic idea is to partition the sphere, rather than the points, into a number of regions and maintain a sketch for each region individually. This idea of sphere partitioning was previously used by Bourgain and Lindenstrauss [BL88] to obtain suboptimal results for ℓ_1-subspace embeddings for $d \ge 5$. We remark that our new algorithm does not give a subspace embedding, unlike the previous ones.

To begin with, we state a partition lemma for spheres.

LEMMA 5.5. *Suppose that $\eta \in (0, 1/2)$. There exists a partition of $\mathbb{S}^{d-1}$ with $c_1(1/\eta)^{d-1}$ regions such that (i) the diameter of each region is at most 2η and (ii) for every $x \in \mathbb{S}^{d-1}$, the hyperplane $\{y : \langle x, y\rangle = 0\}$ only intersects at most $c_2(1/\eta)^{d-2}$ regions. Here c_1 and c_2 are constants that only depend on d.*

Proof. Take a maximal set $\mathcal{N} \subset \mathbb{S}^{d-1}$ such that $d(x, y) > \eta$ for all distinct $x, y \in \mathcal{N}$. It is a standard fact that $m := |\mathcal{N}| \le c_1(d)/\eta^{d-1}$. Suppose that $\mathcal{N} = \{v_1, \dots, v_m\}$. For each i, define $Q_i = \{x \in \mathbb{S}^{d-1} : d(x, v_i) = d(x, \mathcal{N})\}$ and $R_i = Q_i \setminus \bigcup_{j<i} R_j$. We can see the regions $R_1, \dots, R_m$ form a partition of $\mathbb{S}^{d-1}$. By the construction of $\mathcal{N}$, we can see that $\operatorname{int}(B(v_i, \eta/2)) \subseteq R_i \subseteq B(v_i, \eta)$, where $\operatorname{int}(\cdot)$ denotes the relative interior of a set on $\mathbb{S}^{d-1}$.

Now, given any $x \in \mathbb{S}^{d-1}$, we bound the number of the regions that intersect the equator $E_x = \{y \in \mathbb{S}^{d-1} : \langle x, y\rangle = 0\}$. On the one hand, if a region R_i intersects the equator E_x, it is covered by the band $\{y \in \mathbb{S}^{d-1} : d(y, E_x) \le 2\eta\}$, which has area at most $c_2(d)\eta$. On the other hand, each region R_i contains an open neighbourhood $B(v_i, \eta/2)$, thus having area at least $c_3(d)\eta^{d-1}$. Since the regions are disjoint, the number of regions that intersect E_x must be at most $c_2(d)\eta/c_3(d)/\eta^{d-1} = c_4(d)/\eta^{d-2}$. ∎

Algorithm Description. We now describe our algorithm. Given approximation parameter ε, invoking Lemma 5.5 with $\eta = \varepsilon^{\frac{2}{d+2p-1}}$, we obtain a partition of the sphere $R_1, \ldots, R_t$ where $t = O(\eta^{d-1}) = O(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$. For each region we maintain the following two things throughout the stream (we rescale the points which are not on the sphere to unit vectors when determining the region for them): (i) the sum of the p-th tensor products over the points in the region $q_i = \sum_{y \in R_i} y^{\otimes p}$. (ii) a sample point z_i in the region, for which we employ reservoir sampling so that the sample point is uniformly chosen from all the points in the region. If the number of points in one region exceeds $\eta^{d-1}n$ in the middle of the stream, we will treat this region as a new region and create a separate sketch for the new incoming points in it. Note that such an operation will create at most $O((1/\eta)^{d-1})$ new regions.

Query Algorithm. Given a query point $y \in \mathbb{S}^{d-1}$, we perform the following procedure: for each region R_i, similar to the analysis in Section 4, if the hyperplane $H = \{y : \langle x, y\rangle = 0\}$ does not intersect R_i, we have that

$$\langle x^{\otimes p}, \sum_{y\in R_i} y^{\otimes p}\rangle = \sum_{y \in R_i} \langle x, y\rangle^p = \langle x^{\otimes p}, q_i\rangle \ .$$

Hence we obtain $\sum_{y\in R_i} |\langle x, y\rangle|^p$ with zero error. Now we focus on the case that H intersects R_i. In this case, since the value of $|\langle x, y\rangle|^p$ is small for all $y \in R_i$, we can use the sample point z_i to estimate $\sum_{y\in R_i} |\langle x, y\rangle|^p$. Specifically, suppose that R_i contains c_i points. We define an estimator $Z_i = |\langle x, z\rangle|^p \cdot c_i$ and let $Z = \sum_{i\in I} Z_i$ be our final estimator, where I is the set of indices of the regions which H intersects.

It is easy to see that Z is an unbiased estimator. To analyze the concentration, note that for each such area R_i, $\mathbf{Var}[Z_i] \le \sum_{v\in R_i} \frac{1}{c_i}\cdot(c_i\cdot \ell)^2 = c_i^2\ell^2$, where $\ell = \max_{v\in R_i} |\langle v, x\rangle|^p \le \eta^p$. Recall that $c_i \le \eta^{d-1}n$ by our construction, and there are $O(\eta^{-(d-1)})$ regions. We have that $\mathbf{Var}[Z_i] \le \eta^{2(d-1)}n^2\eta^{2p}$ and thus $\mathbf{Var}[Z] \le \eta^{d-1+2p}n^2 = \varepsilon^2 n^2$. It then follows from Chebyshev's inequality that the error $|Z - \mathbb{E}\, Z| = O(\varepsilon n)$ with probability at least $9/10$. Rescaling by a normalization factor of $1/n$ yields the following theorem.

THEOREM 5.4. (FOR-EACH VERSION) *Let $A = a_1 \circ \cdots \circ a_n$ be a stream of n rows, where $a_i \in \mathbb{R}^{1\times d}$. There is an algorithm which maintains a data structure Q of $O(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ words of space such that for each $x \in \mathbb{S}^{d-1}$, with probability at least $9/10$,*

$$\left|Q(x) - \frac{1}{n}\|Ax\|_p^p\right| = O(\varepsilon) \ . \tag{5.15}$$

Moreover, the algorithm updates Q in $O(1)$ time and can be implemented in $O(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ words of space.

Applying the median trick and a net argument, we obtain the following for-all version of Theorem 5.4.

COROLLARY 5.1. (FOR-ALL VERSION) *Let $A = a_1 \circ \cdots \circ a_n$ be a stream of n rows, where $a_i \in \mathbb{R}^{1\times d}$. There is an algorithm which maintains a data structure Q of $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ words of space such that (5.15) holds for all $x \in \mathbb{S}^{d-1}$ simultaneously with high probability. Moreover, the algorithm updates Q in $O(1)$ time and can be implemented in $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p-1}})$ words of space.*

Tight Bound for $d \le 2p+2$. Below we show that when $d \le 2p+2$, the algorithm above can achieve a tight space complexity with a small modification. Let $\eta = \varepsilon^{\frac{2}{d+2p}}$. When a region contains more than $\eta^{d-1}n$ points, we split it into a number of new sub-regions with diameter at least half that of the larger region and we stop splitting when a region has diameter less than $O(\mathrm{poly}(\varepsilon))$ because the error is negligible at this point. Specifically, for a region with diameter $O(\eta/2^i)$, we take a new $O(\eta/2^{i+1})$-partition in Lemma 5.5 and partition the region using the new partitions.

Now we turn to bound the variance. We split Z as $Z = \sum_{i\ge 0}\sum_{j\in I_i} Z_j$, where I_i is the set of indices of regions which H intersects and has diameter in $(\eta/2^{i-1}, \eta/2^i]$. From Lemma 5.5 we have that $|I_i| = O(\eta^{-(d-2)}\cdot 2^{i(d-2)})$ and so $\sum_{j\in I_i}\mathbf{Var}[Z_j] \le (\eta^{d-1}n)^2\cdot(\eta/2^i)^{2p}\cdot|I_i| = O(\eta^{d+2p}n) = O(\varepsilon^2n^2)$. Summing over $i = 0, 1, \ldots, O(\log(1/\varepsilon))$, it follows that $\mathbf{Var}[Z] = O(\varepsilon^2 n^2\log(1/\varepsilon))$. Using the same argument as before and after a rescaling of ε, we conclude with the following theorem.

THEOREM 5.5. *Suppose that* $d \le 2p+2$. *Let* $A = a_1 \circ \cdots \circ a_n$ *be a stream of* n *rows, where* $a_i \in \mathbb{R}^{1\times d}$. *There is an algorithm which maintains a data structure* Q *of* $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ *words of space such that* (5.15) *holds for all* $x \in \mathbb{S}^{d-1}$ *simultaneously with high probability. Moreover, the algorithm updates* Q *in* $O(1)$ *time and can be implemented using* $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2p}})$ *words of space.*

6 Affine ℓ_p-Subspace Sketch

In this section, we consider the following special version of the affine ℓ_p-subspace sketch problem, which can be seen as a generalization of the ℓ_p subspace sketch problem. Given a matrix $A \in \mathbb{R}^{n\times d}$ and for any given $x \in \mathbb{R}^d$ and $b \in \mathbb{R}$, we want to determine a $(1 \pm \varepsilon)$-approximation to $\sum_i |\langle A_i, x\rangle - b|^p$.

Upper Bound. We first consider the upper bound. The crucial observation is that, given matrix $A \in \mathbb{R}^{n\times d}$, let $B \in \mathbb{R}^{n\times(d+1)}$ be the matrix for which the i-th row $B_i = \begin{pmatrix} A_i & -1 \end{pmatrix}$. Suppose that $x \in \mathbb{R}^d$ and $b \in \mathbb{R}$ are the query vector and value, and let $y = \begin{pmatrix} x^\top & b \end{pmatrix}^\top \in \mathbb{R}^{d+1}$. Then

$$\sum_i |\langle A_i, x\rangle - b|^p = \sum_i |\langle B_i, y\rangle|^p = \|By\|_p^p .$$

Hence, any data structure that solves the ℓ_p subspace sketch problem in $(d+1)$ dimensions solves the affine ℓ_p subspace sketch problem for d dimensions. The following theorem follows immediately from Theorem 4.2.

THEOREM 6.1. *Suppose that* $A \in \mathbb{R}^{n\times d}$ *and* p *is a positive integer constant. Then there is a polynomial-time algorithm that returns a data structure with* $\widetilde{O}(\varepsilon^{-\frac{2d}{d+2p+1}})$ *bits of space, which solves the affine* ℓ_p *subspace sketch problem.*

Lower Bound. We now turn to the lower bound. We will show that interestingly, we obtain a lower bound with the same space complexity. To achieve this, we need the following lemma, which can be seen as a stronger version of Lemma 3.2.

LEMMA 6.1. *Suppose that* $p \in [1,\infty) \setminus 2\mathbb{Z}$ *is a constant. Let* A *and* B *be sets of* $n \le N$ *points in a spherical shell* $\{x \in \mathbb{R}^d : \alpha \le \|x\|_2 \le \beta\}$, *where* $\alpha, \beta > 0$ *are constants such that* $\alpha < \beta < (\frac{1+\sqrt{3}}{2})^{\frac{1}{p}}\alpha$. *Suppose that* A *and* B *are symmetric around the origin and* $\left\| \frac{A_i}{\|A_i\|_2} - \frac{B_j}{\|B_j\|_2} \right\|_2 \ge \eta = C_1 N^{-\frac{1}{d-1}}$ *for all* $i \ne j$. *Then we have*

$$\delta \equiv \sup_{x\in\mathbb{S}^{d-1}} \frac{1}{n} \left| \|Ax\|_p^p - \|Bx\|_p^p \right| \ge c_2 N^{-\frac{d+2p}{2(d-1)}} .$$

In the statement above, $C_1 > 0$ *is a constant that depends only on* d *and* $c_2 > 0$ *is a constant that depends on* d, p, α, β *only.*

Proof. Normalizing the points in A and B, we consider the equivalent form of this problem: A and B are still subsets of $\mathbb{S}^{d-1}$, while the target error δ becomes

$$\delta \equiv \sup_{x\in\mathbb{S}^{d-1}} \left(\sum_i w_{A_i} |\langle A_i, x\rangle|^p - \sum_i w_{B_i} |\langle B_i, x\rangle|^p \right),$$

where the w_i are weights such that $w_i \in [\frac{1}{\beta^p n}, \frac{1}{\alpha^p n}]$. Define the function $h_A = \sum_i w_{A_i} |\langle A_i, x\rangle|^p$ and $h_B = \sum_i w_{B_i} |\langle B_i, x\rangle|^p$. Following similar steps as in the proof of Lemma 3.2, we obtain that

$$\sum_{i=1}^n \frac{w_{A_i}(1-r^2)}{(1+r^2-2r\langle u, A_i\rangle)^{d/2}} = c_A + \sum_{\text{even } k\ge 2} r^k \lambda_k^{-1} \sum_j \langle h_A, Y_{k,j}\rangle Y_{k,j}(u)$$

and

$$\sum_{i=1}^n \frac{w_{B_i}(1-r^2)}{(1+r^2-2r\langle u, B_i\rangle)^{d/2}} = c_B + \sum_{\text{even } k\ge 2} r^k \lambda_k^{-1} \sum_j \langle h_B, Y_{k,j}\rangle Y_{k,j}(u),$$

where $c_A = \sum_i w_{A_i}$ and $c_B = \sum_i w_{B_i}$ are both contained in $[\frac{1}{\beta}, \frac{1}{\alpha}]$. Hence

$$\sum_{i=1}^{n} \frac{w_{A_i}(1-r^2)}{(1+r^2-2r\langle u, A_i\rangle)^{d/2}} - \sum_{i=1}^{n} \frac{w_{B_i}(1-r^2)}{(1+r^2-2r\langle u, B_i\rangle)^{d/2}} = c_A - c_B + \sum_{\text{even } k\geq 2} r^k \lambda_k^{-1} \sum_j \langle h_A - h_B, Y_{k,j}\rangle Y_{k,j}(u).$$

Integrating with respect to u on $\mathbb{S}^{d-1}$, we have that

$$\begin{aligned} &\left\| \sum_{i=1}^{n} \frac{w_{A_i}(1-r^2)}{(1+r^2-2r\langle u, A_i\rangle)^{d/2}} - \sum_{i=1}^{n} \frac{w_{B_i}(1-r^2)}{(1+r^2-2r\langle u, B_i\rangle)^{d/2}} \right\|_{L_2(\sigma_{d-1})}^2 \\ &= \sum_{\text{even } k\geq 2} r^{2k}\lambda_k^{-2} \sum_j \langle h_A - h_B, Y_{k,j}\rangle^2 + (c_A - c_B)^2 \\ &\leq \delta^2 \max_{\text{even } k\geq 2} (r^{2k}\lambda_k^{-2}) + (c_A - c_B)^2, \end{aligned} \tag{6.16}$$

Following the same steps as in Lemma 3.2, we obtain that the left hand side of (6.16) is at least $\frac{2}{\beta^{2p}} - \frac{1}{\alpha^{2p}} > 0$. Note that we have c_A and c_B are both in $[\frac{1}{\beta^p}, \frac{1}{\alpha^p}]$ so $(c_A - c_B)^2 \geq (\frac{1}{\alpha^p} - \frac{1}{\beta^p})^2$. Hence

$$\delta^2 \max_{\text{even } k} (r^{2k}\lambda_k^{-2}) \geq \frac{2}{\beta^{2p}} - \frac{1}{\alpha^{2p}} - (\frac{1}{\alpha^p} - \frac{1}{\beta^p})^2 = \frac{1}{\beta^{2p}} + \frac{2}{\alpha^p\beta^p} - \frac{2}{\alpha^{2p}} > 0$$

by our assumptions on α and β. It follows that

$$\delta \geq c(d,p,\alpha,\beta) N^{-\frac{d+2p}{2(d-1)}}.$$

□

Consider the same $S = \{p_1, \ldots, p_n\}$ as in Section 3 for $(d+1)$ dimensions. Consider the spherical cap

$$C = \left\{ x \in \mathbb{S}^d : x_{d+1} \geq \left(\frac{3}{4}\right)^{\frac{1}{p}} \right\}.$$

and let $T = S \cap C$. From Lemma 2.2 we have that the area of C is at least $c(d,p)$, a constant depending only on d and p. Hence, $|T| = c(d,p) \cdot \Omega(\varepsilon^{-\frac{2d}{d+2p+1}})$. Let $T' \subset \mathbb{R}^{d+1}$ be the set of points obtained by scaling the $(d+1)$-st coordinate of each point in T to 1. Note that every point in T' has the same last coordinate and ℓ_2 norm in $[1, (\frac{4}{3})^{1/p}]$. Following the same steps in Section 3 and combining with Lemma 6.1, we have that any affine ℓ_p subspace sketch data structure solves the same subset identification problem. Our theorem follows immediately.

THEOREM 6.2. *Suppose that $p \in [1,\infty) \setminus 2\mathbb{Z}$ and d are constants. Any data structure that solves the affine ℓ_p-subspace sketch problem for dimension d and accurary parameter ε requires $\Omega(\varepsilon^{-\frac{2d}{d+2p+1}})$ bits of space.*

7 Point Estimation for SVMs

As an application of our results, we obtain tight space bounds for point estimation for the streaming SVM problem. In this section, we consider the following (regularized) SVM objective function:

$$F_\lambda(\theta, b) := \frac{\lambda}{2}\|(\theta, b)\|_2^2 + \frac{1}{n}\sum_{i=1}^{n} \max\{0, 1 - y_i(\theta^\top x_i + b)\}, \tag{7.17}$$

where n data points $(x_i, y_i) \in \mathbb{R}^d \times \{-1,+1\}$, with $\|x_i\|_2 = O(1)$ and $(\theta, b) \in \mathbb{R}^d \times \mathbb{R}$ are the unknown model parameters, and λ is the regularization parameter. In this section, we consider the point estimation problem, that is, given (θ, b), we want to output an approximation to $F_\lambda(\theta, b)$. As mentioned in [ABL+20], when $d \geq 2$, it is impossible to obtain a $(1 \pm \varepsilon)$-approximation in $o(n)$ space and so we consider $O(\varepsilon)$-additive error. Throughout this section, we assume that $\lambda = O(1)$, which is a common setting for this problem.

DEFINITION 7.1. (SVM POINT ESTIMATION) *Given* (θ, b) *with* $\|(\theta, b)\|_2 = O(1)$*, compute a value* Z *such that* $|Z - F_\lambda(\theta, b)| \le \varepsilon$*, where* $F_\lambda(\theta, b)$ *is as defined in* (7.17) *with* $\|x_i\|_2 = O(1)$ *for all* i*,* $\lambda = O(1)$ *and* $y_i \in \{-1, +1\}$ *for all* i*.*

First, we demonstrate that it suffices to consider a simplified SVM objective, as mentioned in [ABL+20]. We can assume that $\lambda = 0$ because we can compute the regularization term exactly. Next, recall that $y_i = \pm 1$, and so we can estimate the contribution from the positive labels and negative labels separately and so we can assume, without loss of generality, that $y_i = 1$ for all i. With a further replacement of b with $1-b$, the objective is changed to

$$F(\theta, b) := \frac{1}{n}\sum_{i=1}^{n} \max\{0, b - \theta^\top x_i\} .$$

Upper Bound. The observation is similar to that in Section 6. Given points $x_i \in \mathbb{R}^d$ with positive labels, let $y_i = \begin{pmatrix} -x_i^\top & 1 \end{pmatrix}^\top \in \mathbb{R}^{d+1}$. Suppose that $\theta \in \mathbb{R}^d$ and $b \in \mathbb{R}$ are the query vector and value, and let $\theta' = \begin{pmatrix} \theta^\top & b \end{pmatrix}^\top \in \mathbb{R}^{d+1}$. Then

$$F(\theta, b) = \frac{1}{n}\sum_i \max\{0, b - \theta^\top x_i\} = \frac{1}{n}\sum_i \max\{0, \theta'^\top y_i\}.$$

It is now clear that the problem can be seen as a variant of the ℓ_1-subspace sketch problem in dimension $d+1$, where $|\theta^\top x_i|$ is replaced with $\max\{0, \theta^\top x_i\}$. Our earlier algorithms also work for the $\max\{x, 0\}$-loss for the subspace sketch problem. Recall the proof of Lemma 4.1: we partition the data points into groups $P_1, \ldots, P_t$. Given a query point θ, for a group P_i for which the hyperplane $H = \{x : \langle x, \theta\rangle = 0\}$ does not intersect, we can still obtain the exact value of

$$\sum_{x\in P_i} \max\{0, \langle \theta, x\rangle\} = \sum_{x\in P_i} \langle \theta, x\rangle \text{ or } \sum_{x\in P_i} \max\{0, \langle \theta, x\rangle\} = 0 .$$

For the groups which H intersects, since $\max\{0, \langle\theta, x\rangle\} \le |\langle\theta, x\rangle|$, our analysis of the concentration for the sampling method still holds. Therefore, an analogous version of Lemma 5.1 holds, that is, we can find a coreset of size $\widetilde{O}(\varepsilon^{-\frac{2d}{d+3}})$, which approximates $F(\theta, b)$ up to an additive error of ε in $O(n)$ space. To reduce the space usage, the key observation is that uniformly sampling $O(1/\varepsilon^2)$ points, by Hoeffding's inequality, is sufficient for an $O(\varepsilon)$ additive error. Therefore, we have effectively reduced n to $O(1/\varepsilon^2)$ as it suffices to find a coreset for $O(1/\varepsilon^2)$ uniformly sampled points, for which we use Algorithm 1. Thus, we arrive at the following theorem.

THEOREM 7.1. *Let* $x_1, \ldots, x_n$ *be a stream of* n *points, where* $x_i \in \mathbb{R}^d$*. There is an algorithm which computes a subset* y_i*'s of* $m = \widetilde{O}(\varepsilon^{-\frac{2d}{d+3}})$ *points of* a_i*'s and weights* $w \in \mathbb{R}^m$ *such that with high probability, for every* $(\theta, b) \in \mathbb{R}^d \times \mathbb{R}$ *with* $\|(\theta, b)\|_2 = O(1)$*,*

$$\left|\frac{1}{n}\sum_i w_i \cdot \max\{0, b - \theta^\top y_i\} - \frac{1}{n}\sum_i \max\{0, b - \theta^\top x_i\}\right| = O(\varepsilon) .$$

Moreover, the algorithm can be implemented in $\widetilde{O}(\varepsilon^{-\frac{2d}{d+3}})$ *words of space.*

Lower Bound. We now turn to lower bounds for the point estimation problem. Suppose that $X = \{x_i\}$ is the point set given by the data stream. Let $-X = \{-x : x \in X\}$ and observe that

$$\underset{X}{F}(\theta, 0) + \underset{-X}{F}(\theta, 0) = \frac{1}{n}\sum_i \left(\max\{0, \theta^\top x_i\} + \max\{0, -\theta^\top x_i\}\right) = \frac{1}{n}\sum_i \left|\theta^\top x_i\right| .$$

Thus, if we can solve the d-dimensional point estimation problem for SVM, we can solve the d-dimensional ℓ_1-subspace sketch problem, whence a lower bound of $\Omega(\varepsilon^{-\frac{2(d-1)}{d+2}})$ bits follows from Theorem 3.1. To obtain a tight bound, consider the affine ℓ_1-subspace sketch problem in Section 6. Specifically, we have

$$\underset{X}{F}(\theta, b) + \underset{-X}{F}(\theta, -b) = \frac{1}{n}\sum_i \left(\max\{0, b - \theta^\top x_i\} + \max\{0, -b + \theta^\top x_i\}\right) = \frac{1}{n}\sum_i \left|b - \theta^\top x_i\right| ,$$

which implies that if we can solve the d-dimensional point estimation problem for SVM, we would be able to solve the d-dimensional affine ℓ_1-subspace sketch problem. Our theorem follows immediately from Theorem 6.2.

THEOREM 7.2. *Suppose that d is constant. Any data structure that solves the d-dimensional point estimation problem for SVM requires $\Omega(\varepsilon^{-\frac{2d}{d+3}})$ bits of space.*

We remark that our analysis shows tight space complexity via a black box reduction to the ℓ_1 subspace sketch problem, which is much simpler than the analysis in previous work [ABL+20].

References

[ABL+20] Alexandr Andoni, Collin Burns, Yi Li, Sepideh Mahabadi, and David P. Woodruff. Streaming complexity of svms. In Jaroslaw Byrka and Raghu Meka, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2020, August 17-19, 2020, Virtual Conference*, volume 176 of *LIPIcs*, pages 50:1–50:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[AH12] Kendall Atkinson and Weimin Han. *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*, volume 2044 of *Lecture Notes in Mathematics*. Springer-Verlag Berlin Heidelberg 2012, 2012.

[BBD15] Jeremiah Blocki, Manuel Blum, and Anupam Datta. Naturally rehearsing passwords. *IACR Cryptol. ePrint Arch.*, page 166, 2015.

[BDM+20] Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P. Woodruff, and Samson Zhou. Near optimal linear algebra in the online and sliding window models. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 517–528. IEEE, 2020.

[BL88] J. Bourgain and J. Lindenstrauss. Distribution of points on spheres and approximation by zonotopes. *Israel Journal of Mathematics*, 64:25–31, 1988.

[BLM89] J. Bourgain, J. Lindenstrauss, and V. Milman. Approximation of zonoids by zonotopes. *Acta Mathematica*, 162(none):73 – 141, 1989.

[BM83] U. Betke and P. McMullen. Estimating the sizes of convex bodies from projections. *Journal of the London Mathematical Society*, s2-27:525–538, June 1983.

[BW03] Károly Böröczky and Gergely" Wintsche. Covering the sphere by equal spherical balls. In Boris Aronov, Saugata Basu, János Pach, and Micha Sharir, editors, *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, pages 235–251. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.

[DDH+09] Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W. Mahoney. Sampling algorithms and coresets for ℓ_p regression. *SIAM J. Comput.*, 38(5):2060–2078, 2009.

[Dor76] L. E. Dor. Potentials and isometric embeddings in L_1. *Israel Journal of Mathematics*, 24(3–4):260–268, 1976.

[DX13] Feng Dai and Yuan Xu. *Approximation theory and harmonic analysis on spheres and balls*. Springer Monogr. Math. New York, NY: Springer, 2013.

[FG11] Omer Friedland and Olivier Guédon. Random embedding of ℓ_p^n into ℓ_r^n. *Mathematische Annalen*, 350(4):953–972, 2011.

[GLS88] Martin Grötschel, Lászlo Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag Berlin Heidelberg, 1988.

[Gor85] Yehoram Gordon. Some inequalities for gaussian processes and applications. *Israel Journal of Mathematics*, 50:265–289, 1985.

[Lin89] J. Linhart. Approximation of a ball by zonotopes using uniform distribution on the sphere. *Archiv der Mathematik*, 53:82–86, 1989.

[LS15] Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In Venkatesan Guruswami, editor, *56th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 230–249. IEEE Computer Society, 2015.

[LWW21] Yi Li, Ruosong Wang, and David P. Woodruff. Tight bounds for the subspace sketch problem with applications. *SIAM J. Comput.*, 50(4):1287–1335, 2021.

[Mat96] Jiří Matoušek. Improved upper bounds for approximation by zonotopes. *Acta Mathematica*, 177(1):55–73, 1996.

[PBM88] A. P. Prudnikov, Yu. A. Brychkov, and O. I. Marichev. *Integrals and series, Vol 2, Special functions*. Gordon & Breach Science Publishers, New York, 1988. Translated from the Russian by N. M. Queen.

[Rag71a] David L. Ragozin. Constructive polynomial approximation on spheres and projective spaces. *Transactions of the American Mathematical Society*, 162:157–170, December 1971.

[Rag71b] David L. Ragozin. Uniform convergence of spherical harmonic expansions. *Mathematische Annalen*, 195:87–94, 1971.

[Sch87] Gideon Schechtman. More on embedding subspaces of L_p in l_r^n. *Compositio Mathematica*, 61:159–169, 1987.
[Sch11] Gideon Schechtman. Tight embedding of subspaces of L_p in ℓ_p^n for even p. *Proceedings of the American Mathematical Society*, 139(12):4419–4421, December 2011.
[WY22] David P. Woodruff and Taisuke Yasuda. High-dimensional geometric streaming in polynomial space. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, Oct 31–Nov 3, 2022*. IEEE, 2022. To appear.

A ℓ_p-Subspace Sketch Upper Bound for Non-integer p

In this section, we consider the ℓ_p-subspace sketch problem for non-integers p. We shall show that for every constant d and non-integer constant p, it is still possible to obtain a sketch of size $o(\varepsilon^{-2})$.

A.1 Subspace Embedding for $d \geq 5$. In Section 4, we adapted a proof of Matousek [Mat96], giving a bound that contains logarithmic factors. Matousek's work contains a second result, which shows a tight $O(\varepsilon^{-2(d-1)/(d+2)})$ bound without logarithmic factors for $d \geq 5$ and $p = 1$. In fact, the same bound also holds for $p > 1$, which we shall demonstrate below.

The starting point and the main change is the following generalization of a proposition in [Mat96] to $p > 1$.

Lemma A.1. (Generalization of [Mat96, Proposition 9]) *Let $d \geq 3$ and $P \subseteq \mathbb{S}^{d-1}$ be a point set of size N, where $N \geq N_0$ for some large constant N_0. There exist a subset $P^* \subseteq P$ of $N^* \geq N/8$ points with N^* even, and a subset $Q \subseteq P^*$ of size $N^*/2$ such that for all $x \in \mathbb{S}^{d-1}$*

$$\left| \sum_{y \in P^*} |\langle x, y\rangle|^p - 2 \sum_{y \in P^*} |\langle x, y\rangle|^p \right| = O(N^{\frac{1}{2} - \frac{3}{2(d-1)}})$$

Proof. We only highlight the changes from the original proof in [Mat96]. Take k such that $2^k \sim N^2$ and for each $i = 1, \dots, k$, let $\mathcal{N}_i$ be a 2^{-i}-net on $\mathbb{S}^{d-1}$ and $\pi_i(x) : \mathbb{S}^{d-1} \to \mathcal{N}_i$ be the induced projection map. Then define $\phi_{i,q} \in \mathbb{S}^{d-1} \to \mathbb{R}$ for $q \in \mathcal{N}_i$ as

$$\phi_{i,q}(y) = \begin{cases} |\langle q, y\rangle|^p, & \text{if } i = 1; \\ |\langle q, y\rangle|^p - |\langle \pi_{i-1}(q), y\rangle|^p, & \text{if } i > 1. \end{cases}$$

The claim is that the functions $\phi_{i,q}$ satisfy the following three properties:

(i) $|\phi_{i,q}(y)| = O(2^{-i})$ for all $y \in \mathbb{S}^{d-1}$;

(ii) Define

$$L_{i,q}^+ = \{y \in \mathbb{S}^{d-1} : \langle q, y\rangle \geq 0 \text{ and } \langle \pi_{i-1}(q), y\rangle \geq 0\}$$
$$L_{i,q}^- = \{y \in \mathbb{S}^{d-1} : \langle q, y\rangle \leq 0 \text{ and } \langle \pi_{i-1}(q), y\rangle \leq 0\}$$

Then $\phi_{i,q}$ is $O(2^{-i})$-Lipschitz on $L_{i,q}^+$ and $L_{i,q}^-$.

(iii) The expansion

$$|\langle x, y\rangle|^p = \sum_{i=1}^{k} \phi_{i,q_i}(y) + r_x(y),$$

where $q_k = \pi_k(x)$ and $q_{i-1} = \pi_{i-1}(q_i)$, has the remainder term $|r_x(y)| = O(N^{-2})$ for all $y \in \mathbb{S}^{d-1}$.

The three properties are easy to verify when $p = 1$. Now we shall verify them for a general $p > 1$. For notational convenience, let $u = \pi_{i-1}(q) - q$. Then $\|u\|_2 \leq 2^{-i}$. Property (i) is easy to verify as $|\phi_{1,q}(y)| \leq 1$ and

$$|\phi_{i,q}(y)| = ||\langle q, y\rangle|^p - |\langle q, y\rangle + \langle u, y\rangle|^p| \leq p\,|\langle u, y\rangle| \leq p2^{-i},$$

where we used the fact that $1 - x^p \leq p(1 - x)$ when $x \in (0, 1)$ and thus $||a|^p - |b|^p| \leq p\,|a - b|$ when $|a|, |b| \leq 1$. Property (iii) is also easy to verify, as

$$|r_x(y)| = ||\langle x, y\rangle|^p - |\langle \pi_k(x), y\rangle|^p| \leq p\,|\langle x - \pi_k(x), y\rangle| \leq p\,\|x - \pi_k(x)\|_2 \leq p2^{-k} = O(N^{-2}).$$

Next, we verify Property (ii). Suppose that $y, z \in L_{i,q}^+$. We first consider

$$\sup_{\substack{y,z\in L_{i,q}^+ \\ \langle u, y-z\rangle\neq 0}} \left|\frac{\phi_{i,q}(y)-\phi_{i,q}(z)}{\langle u,y\rangle-\langle u,z\rangle}\right| = \sup_{\substack{y,z\in L_{i,q}^+ \\ \langle u, y-z\rangle\neq 0}} \left|\frac{\langle q,y\rangle^p-(\langle q,y\rangle+\langle u,y\rangle)^p-\langle q,z\rangle^p+(\langle q,z\rangle+\langle u,z\rangle)^p}{\langle u,y\rangle-\langle u,z\rangle}\right|$$

By relating the expression to the definition of derivatives, it is easy to see that this supremum is upper bounded by a constant L (which depends on p only), thus

$$|\phi_{i,q}(y)-\phi_{i,q}(z)| \le L\,|\langle u, y-z\rangle| \le 2L\,\|u\|_2 = O(2^{-i}).$$

By continuity, this bound also holds for $\langle u,y\rangle = \langle u,z\rangle$. Therefore, we have verified that $\phi_{i,q}$ is $O(2^{-i})$-Lipschitz on $L_{i,q}^+$. A similar argument works for $L_{i,q}^-$ and thus we have verified (ii).

The rest of Matousek's original proof goes through, establishing the lemma. □

Next, we repeat Matousek's argument. Repeatedly applying the preceding lemma yields P_1^* and Q_1 from P, P_2^* and Q_2 from $P\setminus P_1^*$, P_3^* and Q_3 from $P\setminus(P_1^*\cup P_2^*)$, and so forth. Let Q be the union of these Q_i's. This shows that for any set P of N points, one can find a subset $P'\subset P$ of size $N/2$ such that

$$\left|\frac{1}{N}\sum_{y\in P}|\langle x,y\rangle|^p - \frac{2}{N}\sum_{y\in Q}|\langle x,y\rangle|^p\right| = O(N^{-\frac{d+2}{2(d-1)}}), \quad \forall x\in\mathbb{S}^{d-1}.$$

An iterative application of the step above leads to the final bound of $O(\varepsilon^{-2(d-1)/(d+2)})$, where we can repeat a point several times to accommodate different weights. Formally, we have

THEOREM A.1. *Suppose that $p\ge 1$ and $d\ge 5$ are constants, $A\in\mathbb{R}^{n\times d}$ and $w\in\Delta_{n-1}$ is the associated weight vector. There exists a polynomial time algorithm which outputs a subset B of $m = O(\varepsilon^{-\frac{2(d-1)}{d+2}})$ rows of A and a weight vector $v\in\Delta_{m-1}$ such that with high probability it holds for every $x\in\mathbb{S}^{d-1}$,*

$$\left|\sum_i v_i|\langle B_i,x\rangle|^p - \sum_i w_i|\langle A_i,x\rangle|^p\right| = O(\varepsilon)\cdot\left(\sum_i w_i|\langle A_i,x\rangle|^p\right).$$

A.2 Algorithm for all $d\ge 2$ and Non-Integer p. First, consider the case of $1<p<2$. It is known that ℓ_p^n $(1+\varepsilon)$-embeds into ℓ_1^N for some $N = C(p)n/\varepsilon$ (see, e.g., [FG11]), which means for every matrix $A\in\mathbb{R}^{n\times d}$, there is a matrix $T\in\mathbb{R}^{N\times n}$ such that $\|TAx\|_1 = (1\pm\varepsilon)\|Ax\|_p$ for all $x\in\mathbb{R}^d$. Thus, we can apply our existing upper bound to TA under the ℓ_1 norm, yielding an $\widetilde{O}(\varepsilon^{-\frac{2(d-1)}{d+2}})$ upper bound. We remark that the matrix T constructed in [FG11] is oblivious and has independent columns, so we can write $(TA)_{\cdot,j} = \sum_{i=1}^n T_{\cdot,i}A_{i,j}$ and compute $T_{\cdot,i}A_{i,j}$ for all $j=1,\dots,d$ sequentially for each $i=1,\dots,n$. In this manner, the entire algorithm takes $O(N) = O(n/\varepsilon)$ words of space. We remark that such an embedding-based approach does not seem amenable to the streaming setting because it would require storing the entire A to compute $(TA)_{i,\cdot} = T_{i,\cdot}A$.

Now consider $p>2$. Let $q = \lceil p/2\rceil$, then for every $x\in\mathbb{R}^d$,

$$|\langle A_i,x\rangle|^p = (|\langle A_i,x\rangle|^q)^{p/q} = \left|\langle A_i^{\otimes q}, x^{\otimes q}\rangle\right|^{p/q}.$$

Hence, the problem can be reduced to the d^q-dimensional ℓ_r-subspace sketch problem with $1<r=p/q<2$, which leads to an $\widetilde{O}(\varepsilon^{-\frac{2(d^q-1)}{d^q+2}})$ upper bound. Again, the algorithm uses $O(n/\varepsilon)$ words of space and is not amenable to the streaming setting. We remark that we cannot expect a better upper bound by reducing an ℓ_p-subspace sketch problem to an ℓ_r-subspace sketch problem for an integer $r>2$ because ℓ_p^2 does not $(1+\varepsilon)$-embed into ℓ_r^N when $p>r>2$ [Dor76]. The following is a short proof we include for completeness. Suppose that ℓ_p^2 does $(1+\varepsilon)$-embed into ℓ_r^N. Then there exist $u,v\in\ell_r^N$ such that $\|au+bv\|_r \le (1+\varepsilon)(|a|^p+|b|^p)^{1/p}$ for all $u,v\in\mathbb{R}$. This means that $(1+\varepsilon)2^{r/p} \ge \frac{1}{2}(\|u+v\|_r^r + \|u-v\|_r^r) \ge \|u\|_r^r + \|v\|_r^r = 2$, which is a contradiction for all small ε.

A.3 Streaming Algorithm for $p > d-1$. Next, we present a streaming algorithm when $p > d-1$, using approximation theory on the unit sphere. Given a function $h : \mathbb{S}^{d-1} \to \mathbb{R}$, consider its series expansion in spherical harmonics

$$h(x) = \sum_{k=0}^{\infty} \sum_{j=1}^{M(d,k)} \langle h, Y_{k,j} \rangle Y_{k,j}(x)$$

and the truncation of this series of order at most K

$$(Q_K h)(x) = \sum_{k=0}^{K} \sum_{j=1}^{M(d,k)} \langle h, Y_{k,j} \rangle Y_{k,j}(x).$$

It is a well-studied problem in approximation theory on the unit sphere that (see, e.g. [AH12, Theorem 2.35])

$$\|h - Q_K h\|_\infty \le (1 + \|Q_K\|_{C\to C}) E_{K,\infty}(h),$$

where $\|Q_K\|_{C\to C}$ is the operator norm of Q_K when viewed as an operator from $C(\mathbb{S}^{d-1})$ to $C(\mathbb{S}^{d-1})$, and $E_{n,\infty}(f)$ is the minimum error in the ℓ_∞ norm of approximating f by polynomials of total degree at most K on $\mathbb{S}^{d-1}$. It was shown in [Rag71b] that $\|Q_K\|_{C\to C} \simeq K^{d/2-1}$ when $d \ge 3$ and it is a classical result in Fourier analysis that $\|Q_K\|_{C\to C} \simeq \ln K$ when $d = 2$.

In our problem, without loss of generality, consider $h(x) = \sum_i w_i |\langle a_i, x\rangle|^p$ with $a_i \in \mathbb{S}^{d-1}$ and $\sum_i w_i = O(1)$. Let us first consider the approximation to $f(x) = |\langle a_i, x\rangle|^p$. It follows from the result of Ragozin [Rag71a] (see also [AH12, Eq. (4.49)]) that

$$E_{K,\infty}(f) \le \frac{C_p}{n^p}, \tag{A.1}$$

where $C_p > 0$ is a constant depending only on p. Thus,

$$E_{K,\infty}(h) \le \sum_i w_i E_{K,\infty}(f) \le \frac{C_p'}{n^p}$$

and

$$\|h - Q_K h\|_\infty \lesssim_p \begin{cases} \ln K / K^p & d = 2 \\ 1/K^{p-d/2+1} & d \ge 3. \end{cases}$$

Therefore, we can take

$$K \simeq_p \begin{cases} (1/\varepsilon)^{1/p} \log^{1/p}(1/\varepsilon) & d = 2 \\ (1/\varepsilon)^{1/(p-d/2+1)} & d \ge 3 \end{cases}$$

such that $\|h - Q_K h\|_\infty \le \varepsilon$. For our $h(x)$, we have from Funk-Hecke Theorem that

$$(Q_K h)(x) = \sum_{k=0}^{K} \lambda_k \sum_{j=1}^{M(d,k)} Y_{k,j}(x) \left(\sum_i w_i Y_{k,j}(a_i) \right).$$

Therefore the streaming algorithm needs only to maintain $\sum_i w_i Y_{k,j}(a_i)$ for each $k = 0, \ldots, K$ and $j = 0, \ldots, M(d,k)$ to output $(Q_K h)(x)$. This is clearly feasible in the data stream setting, as we can calculate for each new point a_i the value of $Y_{k,j}(a_i)$ for all pairs (k,j) with $k \le K$. The number of such values to maintain is

$$O\left(\sum_{k=0}^{K} M(d,k) \right) = O\left(\sum_{k=0}^{K} k^{d-2} \right) = O(K^{d-1}).$$

However, this approach suffers from a precision problem when $d \ge 3$. Using the explicit expression of $Y_{k,j}$ in terms of the Gegenbauer polynomials (see, e.g. [DX13, Theorem 1.5.1]) would cause the intermediate results to be as small as $1/K^{O(Kd)}$ or as large as $K^{O(Kd)}$, thus requiring $\widetilde{O}(K)$ words of space to calculate and store the value of each $Y_{k,j}(x)$. This leads to an overall space of

$$\widetilde{O}(K^d) = \widetilde{O}(\varepsilon^{-2d/(2p-d+2)})$$

words when $d \geq 3$, which is $o(\varepsilon^{-2})$ when $p > d-1$. This precision problem, however, does not exist when $d = 2$, since the spherical harmonics degenerate to exactly sines and cosines and the intermediate values can fit in a word. Thus, the streaming algorithm uses $O(K) = \widetilde{O}(\varepsilon^{-1/p})$ words of space when $d = 2$, which is close to the lower bound of $\Omega(\varepsilon^{-1/(p+1)})$ bits.