

Nearly-optimal bounds for sparse recovery in generic norms, with applications to k -median sketching

Arturs Backurs^{*†} Piotr Indyk[‡] Ilya Razenshteyn[§] David P. Woodruff[¶]

Abstract

We initiate the study of trade-offs between sparsity and the number of measurements in sparse recovery schemes for *generic* norms. Specifically, for a norm $\|\cdot\|$, sparsity parameter k , approximation factor $K > 0$, and probability of failure $P > 0$, we ask: what is the minimal value of m so that there is a distribution over $m \times n$ matrices A with the property that for any x , given Ax , we can recover a k -sparse approximation to x in the given norm with probability at least $1 - P$? We give a partial answer to this problem, by showing that for norms that admit efficient linear sketches, the optimal number of measurements m is closely related to the *doubling dimension* of the metric induced by the norm $\|\cdot\|$ on the set of all k -sparse vectors. By applying our result to specific norms, we cast known measurement bounds in our general framework (for the ℓ_p norms, $p \in [1, 2]$) as well as provide new, measurement-efficient schemes (for the Earth-Mover Distance norm). The latter result directly implies more succinct linear sketches for the well-studied planar k -median clustering problem. Finally, our lower bound for the doubling dimension of the EMD norm enables us to resolve the open question of [Frahling-Sohler, STOC'05] about the space complexity of clustering problems in the dynamic streaming model.

1 Introduction

The field of sparse recovery studies the following question: for a signal x , when is it possible to compute an approximation \hat{x} to x that is parameterized by only a small number coefficients, given only a small number of linear measurements of x ? The answers to this basic question, i.e., the *sparse recovery schemes*, have found a

surprising number of applications in a broad spectrum of fields, including *compressive sensing* [CRT06, Don06], data stream computing [Mut05] (see also the resources at `sublinear.info`) and Fourier sampling [GIIS14].

A particularly useful and well-studied formalization of this question is that of *stable sparse recovery*. A general formulation of the problem is as follows. For a norm $\|\cdot\|$, sparsity parameter k , probability of failure P and an approximation factor $K > 0$, design a distribution over $m \times n$ matrices A which has the following property:

There is an algorithm \mathcal{A} that, for any x , given Ax , recovers a vector $\hat{x} = \mathcal{A}(Ax)$ such that

$$(1.1) \quad \|x - \hat{x}\| \leq K \cdot \min_{k\text{-sparse } x'} \|x - x'\|$$

with probability at least $1 - P$.

Here we say that x' is k -sparse if it has at most k non-zero coordinates¹. The typical choices of the norm $\|\cdot\|$ are either ℓ_1 or ℓ_2 . However, several other variants have been studied as well: [BGI⁺08, AZGR15] studied sparse recovery under general ℓ_p norms, [GIP10, IP11, MD13] considered the Earth-Mover-Distance (EMD) norm, while [KW14] considered rearrangement-invariant block norms.

It is easy to observe that the number of measurements m must depend on the sparsity parameter k : the more information about the signal we want to acquire, the more measurements must be taken. For ℓ_1 and ℓ_2 norms, the tradeoff between m and k is well-understood: it is known that $m = O(k \log(n/k))$ measurements suffice [CRT06], and this bound is tight [Don06, DIPW10]. For other norms, however, our understanding of the tradeoffs is much more limited.

1.1 Our results In this paper we initiate the study of sparsity-measurements trade-offs for *generic* norms².

^{*}The first three authors were in part supported by the NSF and the Simons Foundation.

[†]MIT

[‡]MIT

[§]MIT

[¶]IBM Almaden. D. Woodruff was supported in part by the XDATA program of the Defense Advanced Research Projects Agency (DARPA), administered through Air Force Research Laboratory contract FA8750-12-C-0323.

¹Further generalizations of the problem can be obtained by allowing the sparsity in arbitrary basis, or by allowing different norms on the LHS and RHS of Equation 1.1. Although important, we will not consider these generalizations in this paper.

²In fact, our results hold even for *quasi-norms*, e.g., ℓ_p norms for $p < 1$ (see Preliminaries for more details). However, for the

Our results generalize the previously known tradeoffs, and provide improved bounds for specific norms, notably EMD and ℓ_p for $p \in (0, 1)$. Further, our results for EMD immediately yield new sketching algorithms and new lower bounds for the low-dimensional k -median clustering problem.

Our first result shows that, for norms that *admit efficient linear sketches* the number of measurements sufficient for sparse recovery is closely related to the doubling dimension of k -sparse vectors under that norm. Formally, we prove the following theorem.

THEOREM 1.1. *Suppose that $X = (\mathbb{R}^n, \|\cdot\|)$ is an n -dimensional normed space and $1 \leq k \leq n$ be the sparsity parameter. Assume that, for some (distortion) parameter $D \geq 1$ there is a distribution over $s \times n$ random (sketch) matrices S and an (estimator) function $E : \mathbb{R}^s \rightarrow \mathbb{R}$ such that for any x and a k -sparse y we have*

$$\Pr[\|x - y\| \leq E(Sx, Sy) \leq D\|x - y\|] \geq 2/3.$$

Furthermore, let d be the doubling dimension of the set of k -sparse vectors from \mathbb{R}^n with respect to the metric induced by $\|\cdot\|$. Then, for every $0 < \epsilon, \tau < 1/3$ there exists a distribution over random matrices $A \in \mathbb{R}^{m \times n}$ with

$$m = O\left(s \cdot \left(d \cdot \log(D/\epsilon) + \log \log(1/\tau)\right)\right)$$

such that for every $x \in \mathbb{R}^n$ given Ax we can recover with probability at least $2/3$ a vector $\hat{x} \in \mathbb{R}^n$ such that

$$(1.2) \quad \|x - \hat{x}\| \leq (1 + \epsilon)D \min_{k\text{-sparse } x^*} \|x - x^*\| + \tau\|x\|.$$

To explain the theorem, we first observe that the guarantee given by (1.2) is analogous to the one given by (1.1), with the exception of the extra additive term $\tau\|x\|$. The “precision parameter” τ can be made arbitrarily small, at a price of increasing the number of measurements by an extra $\log \log(1/\tau)$ term. Similar tradeoffs between the precision and the number of measurements are quite common in compressive sensing schemes³, although we do not know whether this extra term is necessary in our setting. Apart from the precision dependence, the number of measurements m is linear in the doubling dimension d , linear in the sketch length s and logarithmic in the distortion D .

³For sake of simplicity, in the rest of the paper we will mostly focus on norms.

³E.g., in most of the existing sparse Fourier transform algorithms the sample complexity depends logarithmically on the precision parameter [GIIS14].

Our theorem requires that the normed space of interest admits efficient linear sketches. We believe that some variant of this assumption is necessary for sparse recovery, as such sketches are needed if one wants to estimate the approximation error, i.e., the RHS of Equation 1.2. However, this intuition does not lend itself to a formal argument, as e.g., for the ℓ_1 norm there exist sparse recovery schemes [CRT06, Don06] that satisfy Equation 1.2 *without* explicitly estimating the approximation error. Still, the ℓ_1 norm supports efficient sketches, which suggests that some form of sketchability of the norm could be a necessary condition.

To illustrate Theorem 1.1, consider the case of the ℓ_p (quasi)-norms for $p \in [0, 2]$. It is known [Ind06] that these norms allow sketches with distortion $D = 1 + \epsilon$ and dimension $s = O(1/\epsilon^2)$ for any $\epsilon > 0$, and it is also immediate that the doubling dimension d is $O(k \log(n/k))$. Therefore, for $p \in [1, 2]$ our theorem reproduces the known optimal $O(k \log(n/k))$ measurement bound, up to the dependence on the precision parameter τ . The same bound is obtained for $p \in (0, 1)$. The latter result is, to the best of our knowledge, new.

We note that Theorem 1.1 is not efficient: it does not provide a *polynomial time* algorithm for recovering \hat{x} from Ax . Given the generality of the setting, in particular, the fact that it allows a general (sketchable) norm $\|\cdot\|$, we believe that a general polynomial time recovery algorithm is unlikely to exist. However, it is possible that efficient algorithms exist for specific norms which have good computational properties. For example, we show that for the case of the Earth-Mover Distance norm discussed in more detail below, the recovery algorithm runs in time polynomial in n and $\log^k n$. In particular, the running time is polynomial for any constant k .

Lower bound The ℓ_p norm example shows that the bound of Theorem 1.1 is tight for *some* norms. In fact, one can show that the linear dependence on the doubling dimension d is necessary for *all* norms whose “aspect ratio” is bounded by a polynomial in n . In particular, we show the following theorem.

THEOREM 1.2. *Consider any norm $\|\cdot\|$ over \mathbb{R}^n for which $\frac{1}{n^c} \leq \frac{\|x\|}{\|x\|_2} \leq n^c$ for some constant c . Let $T_k \subset \mathbb{R}^n$ denote the set of k -sparse vectors and $d > 1$ denote the doubling dimension of $[0, \infty)^n \cap T_k$ with $\|\cdot\|$. Then any sparse recovery scheme for $[0, \infty)^n$ with approximation factor K requires $m = \Omega(d/\log K)$ measurements.*

Note that the theorem holds even for vectors $x \geq 0$, which will be useful in the context of the Earth-Mover Distance.

Rand./Deterministic	Sketch length m	Approx.
Deterministic	$k \log n \log(n/k)$	$O(1)$
Deterministic	$k \log(n/k)$	$\sqrt{\log(n/k)}$
Randomized	$k \log(n/k)$	$O(1)$

Figure 1: Performance of sparse recovery schemes for the EMD from [IP11]. The schemes assume that the input vector x is non-negative. Each result implies a sketching scheme for the k -median problem with the same parameters.

Earth-Mover Distance Our results have direct implications for sparse recovery over the Earth-Mover Distance (EMD) norm. This norm is defined over n -dimensional vectors with $n = \Delta^2$, where such vectors can be interpreted as functions $[\Delta]^2 \rightarrow \mathbb{R}$. Informally, for vectors $x, y : [\Delta]^2 \rightarrow \mathbb{R}_+$ which have the same ℓ_1 norm, the EMD is defined as the cost of the min-cost flow that transforms x into y , where the cost of transporting a “unit” of mass from a point $p \in [\Delta]^2$ of x to a point $q \in [\Delta]^2$ of y is equal to the ℓ_1 distance⁴ between p and q . See Preliminaries for a formal definition.

Earth-Mover Distance and its variants are popular metrics for estimating similarity between images and feature sets [RTG00, GD05]. Furthermore, the k -sparse approximation of non-negative vectors under the EMD norm has the following natural interpretation. Let \hat{x} be the k -sparse vector closest to x under this norm. Then one can observe that the non-zero entries of \hat{x} correspond to the cluster centers in the best k -median⁵ clustering of x . Thus, sparse recovery schemes for the EMD norm provide methods for recovering near-optimal solutions to the planar k -median problem from few linear measurements of the input point-sets, a problem that has attracted a considerable attention in streaming and sketching literature [Ind04, FS05, IP11].

The state of the art schemes for this problem are listed in Figure 1. In particular, the best known bound for the number of measurements is $O(k \log(n/k))$, which mimics the best possible bound achievable for sparse recovery in the ℓ_1 norm.

We show that Theorem 1.1 provides new results for this problem. Specifically, we show that the doubling dimension of the EMD norm over k -sparse vectors is

⁴One can also use the ℓ_2 distance. Note that the two distances differ by at most a factor of $\sqrt{2}$ for two-dimensional images.

⁵For completeness, in our context the k -median clustering problem is defined as follows. First, each pixel $p \in [\Delta]^2$ is interpreted as a point with weight x_p . Then the goal is to find a set $C \subset [\Delta]^2$ of k “medians” that minimizes the objective function $\sum_{p \in [\Delta]^2} x_p \cdot \min_{c \in C} \|p - c\|_2$.

only $O(k \log \log n)$. Combined with the known fact that the EMD norm can be embedded into ℓ_1 with distortion $O(\log n)$ [Cha02, IT03] (and therefore its sketching complexity s is constant), this implies that there exist a sparse recovery scheme for EMD with approximation factor $O(\log n)$ that uses only $O(k(\log \log n)^2)$ measurements (ignoring the dependence on the precision). The running time of recovery procedure is polynomial in Δ and $\log^k \Delta$ (again ignoring the dependence on the precision), which is polynomial in Δ for any k up to $\log \Delta / \log \log \Delta$. We further show that the result can be strengthened in three ways:

- By performing a more careful analysis of the embedding procedure of [IT03], we show that it in fact incurs a distortion of $O(\log k + \log \log n)$ with constant probability, which is sufficient for our purposes.
- By using a variant of the embedding (given in [Ind07]) and combining it with a sketch of [VZ12], we show the distortion can be reduced further to $O(\log k)$ while increasing the sketch length by a factor of $O(\log^\delta n)$ for any constant $\delta > 0$. Note that in the case of *constant* k , the approximation we obtain is constant as well.
- Finally, we consider vectors x with the property that, for some integer N , all entries x_p are multiples of $1/N$ (in this case we say that x has *granularity* $1/N$). Such vectors correspond to characteristic vectors of multisets of size N , and naturally occur in the *unweighted* k -median problem over point sets of size N . In this case we show that, in the bounds for the doubling dimension and distortion, we can replace $\log \log n$ by $\log \log N$. Notably, the bounds we obtain in this case are *independent of the ambient dimension* n .

By combining these bounds with Theorem 1.1 we obtain sparse recovery schemes for EMD with the guarantees as in Figure 2 (see also Section 7 for the formal statement of the results).

The aforementioned bounds are quite surprising, as they are provably impossible to achieve for the ℓ_1 or ℓ_2 norms. In particular—for ℓ_1 and ℓ_2 —one needs $\Omega(\log n)$ measurements to achieve constant approximation factor even for $k = 1$, and $\Omega(\log n / \log \log n)$ measurements to achieve $O(\log n)$ distortion [DIPW10]. This means that the EMD norm is actually *easier* than ℓ_p norms from the sparse recovery perspective, at least in a range of parameters.

We also show that at least one $\log \log n$ factor in the measurement bound is necessary as long as $k \geq 2$, by proving a lower bound for the doubling dimension of

Randomized/Deterministic	Sketch length m	Approximation factor
Randomized	$k(\log \log n)(\log(\log k + \log \log n)) + \log \log(1/\tau)$	$O(\log k + \log \log n)$
Randomized	$k \log^\delta n + \log \log(1/\tau)$	$O(\log k)$
Randomized, lower bound	$\Omega(k(\log \log(n/k))/\log K)$	$K \geq 2$

Figure 2: Performance of our sparse recovery schemes for the EMD. The schemes assume that the input vector x is non-negative. The first two results imply a sketching scheme for the k -median problem with the same parameters.

k -sparse vectors under EMD and using Theorem 1.2. In fact, our lower bound argument applies almost verbatim to the space complexity of the following *data stream* problem: design a data structure that maintains a vector x under increments and decrements of its coordinates which, when queried, reports a k -sparse approximation to x with approximation factor K with a constant probability. As discussed earlier, in the context of the EMD norm this task corresponds to the problem of maintaining a k -median clustering of a dynamic point set where points can be inserted and deleted (i.e., the coordinates of x can be incremented and decremented). As we show in Theorem 4.1, the space *bit* complexity of this problem is $\Omega(\frac{d}{\log K} \log n)$ for general norms, thus, in particular, $\Omega(\frac{k}{\log K} \log(\log(\frac{\Delta^2}{k})) \cdot \log \Delta)$ for the EMD norm. The last bound answers the open question of [FS05, Soh] (Section 7) who asked whether it is possible to maintain a constant size (for fixed k and K) “core-set”⁶ for the k -median and k -means problem in dynamic data streams using a constant number of memory words. In this paper we show that any algorithm that solves k -median and k -means⁷ in the dynamic data stream model must use a *super-constant* number of words of size $\log \Delta$, even for constant k and K .

Finally, we show that for the case of $k = 1$, a sparse recovery scheme exists with $O(1)$ measurements for constant d and ϵ , independent of n . This is again in sharp contrast to ℓ_1 or ℓ_2 norms, as well as the aforementioned case of $k \geq 2$.

1.2 Our techniques and related work Our upper bound for the number of measurements relies on the connection between sparse recovery and the approximate nearest neighbor search. Specifically, our goal can be phrased as finding the nearest neighbor of x in a set of bounded doubling dimension. The latter prob-

⁶Informally, a core-set for the k -median problem over a set of points P is a weighted subset $C \subset P$ such that a solution to C provides an approximate solution to P . Core-sets provide a tool for solving streaming problem for k -means and k -median problems in data streams. See [FS05] for more details.

⁷The lower bound for the k -means problem is presented in Section 9.

lem can be solved using the *navigating nets* data structure [KL04], and indeed we are using a similar top-down search approach in our algorithm. However, we need to deal with complications that arise due to the fact that in our setting we can only estimate distances approximately and with a certain probability. Specifically, to obtain the desired bound, we need to ensure that the total number of distances that our sketch needs to preserve is only *linear* in the depth of the tree. This allows us to bound the probability of failure of the algorithm by taking the union bound over a small number of events. It is easy to observe, however, that the path in the tree taken by the search algorithm is *adaptive*, i.e., the approximation errors incurred by the sketch at one level affect the points considered by the algorithm at the next level. Nevertheless we show that the path cannot be *too adaptive*, and that one can identify a set of points of size linear in the tree depth so that preserving all the distances from those points to x ensures the correctness of the algorithm. The details are in Section 3

Our lower bound builds on the argument from [DIPW10], where the number of measurements was lower bounded by encoding long bit sequences into the signal x , such that those bits could be unambiguously decoded by the sparse recovery algorithm. The encoding proceeded on several distance scales. At each scale, the encoding used a large set of almost equidistant k -sparse vectors as the “dictionary”. Since the maximum size of such sets is directly related to the doubling dimension of the space, the lower bound argument goes through in the setting of a *general* norm. The details are in Section 4.

The doubling dimension of the set of k -sparse vectors under EMD was previously studied by [GKK10], who showed that it is at most $O(k \log k)$ for the special case of measures induced by k -sets, i.e., measures of granularity $1/k$. For this case it is in fact not difficult to improve the bound to $O(k)$ and we give an outline of the improved argument in Section 5.1.1. However, for our applications we need a bound that holds for general measures. This makes the argument more complex, since we need to deal with general flows. In both cases the idea of the proof is to explicitly construct a covering of a ball of radius R using a small number

of balls of radius $R/2$, by using the geometric and combinatorial properties of planar flows. The details are in Section 5.1.1 (for general non-negative vectors) and Section 5.1.2 (for vectors of bounded granularity).

Our improved analysis of the embedding of [IT03], as well as the analysis of the embedding from [Ind07], utilize the fact that our application allows us to relax the standard embedding definition in two ways. First, we only need to preserve the distance between a k -sparse vector and a general vector, as opposed to between any pair of vectors (see the statement of Theorem 1.1 for the precise guarantee that we are after). Second, we only need to ensure that the distances are preserved with constant probability, not in expectation, which means that we can tolerate events that incur high distortion as long as they occur with low enough probability. Combining the two relaxations⁸ with a more careful analysis allows us to achieve the improved bound, surprisingly almost without any modifications to the embeddings themselves. The details are in Section 6.

We note that if one wants to preserve EMD between two vectors that are *both* k -sparse, then one can embed those vectors into $\ell_{1-\epsilon}$ with distortion $O(\log k)$ [BI14], which yields a sketch with the same distortion and constant size [Ind06]. Also, for the case when one of the vectors is k -sparse, a recent work [YO14] shows a sketch with distortion $O(\min(k^3, \log n))$ and size roughly $O(\log^4 n)$. The sketch in this paper substantially improves over the latter bound.

For the 1-median problem we solve an ℓ_1 -regression problem. We give oblivious sketches that provide subspace embeddings for the ℓ_1 -norm for d -dimensional subspaces with a “disjoint basis” property that arises in this setting. Our embedding works when the basis is expressible as the union of a small number of sets of vectors, where in each set the vectors have disjoint support. Unlike existing oblivious embeddings for ℓ_1 [CDM⁺13, MM13, SW11, WZ13], we obtain $(1 + \epsilon)$ instead of $\text{poly}(d)$ distortion, and low δ instead of constant probability of failure (to simultaneously preserve norms of all vectors in the space). Our embedding maps n -dimensional vectors to $O(d/\epsilon^2 \log(d/(\delta\epsilon)))$ dimensions. We overcome non-embeddability results for ℓ_1 [BC05, CS02] by using a non-convex estimator. This is reminiscent of estimators for data streams [Ind06], but complicated here by the fact that we require the stronger notion of a subspace embedding. It is known (see, e.g., [ABS10]) that for constant d and ϵ one can

⁸It can be seen that both relaxations are needed in order to achieve the better bound. In particular, the *expected* distortion of the embedding is $\Theta(\log n)$, even for a pair of 1-sparse vectors. Similarly, if $k = n$, the distortion of the embedding is $\Omega(\log n)$ with probability $1 - o(1)$.

solve the 1-median by taking $O(1)$ samples and solving the problem on the samples, but this cannot be expressed as a linear sketch with fewer than $\Omega(\log n)$ measurements (the sampling lower bound follows from Theorem 8 of [JST11]), whereas we achieve $O(1)$ measurements. The details are in Section 8.

2 Preliminaries

EMD. We start by defining EMD. Consider any two non-negative vectors $x, y : [\Delta]^d \rightarrow \mathbb{R}_+$ such that $\|x\|_1 = \|y\|_1$. Let $\Gamma(x, y)$ be a set of functions $\gamma : [\Delta]^d \times [\Delta]^d \rightarrow \mathbb{R}_+$, such that for any $i, j \in [\Delta]^d$ we have $\sum_l \gamma(i, l) = x_i$ and $\sum_l \gamma(l, j) = y_j$. Then we define

$$\text{EMD}^*(x, y) = \inf_{\gamma} \sum_{i, j \in [\Delta]^d} \gamma(i, j) \|i - j\|_1$$

Note that if x and y are characteristic vectors of some sets $A, B \subset [\Delta]^d$, then $\text{EMD}^*(x, y)$ is equal to the value of the minimum cost matching between A and B .

For the case of general vectors x, y , we define

$$\text{EMD}(x, y) = \inf_{\substack{x' \leq x, y' \leq y \\ \|x'\|_1 = \|y'\|_1}} \text{EMD}^*(x', y') + D[\|x - x'\|_1 + \|y - y'\|_1]$$

where $D = d\Delta$ is the diameter of the set $[\Delta]^d$.

Metric spaces. For a metric space (X, d_X) , we define $B_X(u, r)$ or, equivalently, $\text{Ball}_X(u, r)$ to be the ball centered at u or radius r containing all points from X within r from u : $B_X(u, r) := \{x \in X : d_X(u, x) \leq r\}$. Further, for a metric space (X, d_X) , the *doubling dimension* is the smallest number d such that, for every $r > 0$ and any $x \in X$, we can choose $x_1, x_2, \dots, x_{2^d} \in X$ with

$$B_X(x, r) \subseteq B_X(x_1, r/2) \cup B_X(x_2, r/2) \cup \dots \cup B_X(x_{2^d}, r/2).$$

Finally, for $K \geq 1$ we define a K -quasi-metric space as a variant of a metric space, where we have the following *relaxed* triangle inequality: $d(x, y) \leq K \cdot (d(x, z) + d(z, y))$. Thus, every metric space is a 1-quasi-metric space. We define K -quasi-norms in an analogous way.

3 General upper Bound on Measurement Complexity

Suppose we have a K -quasi-metric space $\mathcal{M} = (X, \rho)$ and a closed subset $Y \subseteq X$ with doubling dimension d . Let us assume we can sketch distances between points from X and Y with distortion D , sketch size s and success probability at least $2/3$ (see Theorem 1.1 for the formal definition).

The following Lemma builds on a result from [KL04] on approximate nearest neighbor search in doubling spaces.

Search Procedure:

$y_0 \leftarrow$ the only element of N_0
for $i \leftarrow 1 \dots L$ **do**
 $S_i \leftarrow N_i \cap B_X(y_{i-1}, \beta r_i)$
 $y_i \leftarrow \operatorname{argmin}_{y \in S_i} q(y)$
 if $q(y_i) > \gamma r_i$ **then**
 return y_{i-1}
return y_L

LEMMA 3.1. *For every $0 < \epsilon < 1/2$, $0 < \lambda < \Lambda$ and $y_0 \in Y$ one can sketch points of X with sketch size*

$$O\left(s \cdot \left(d \log(DK/\epsilon) + \log \log(\Lambda/\lambda)\right)\right)$$

so that from this sketch for $x \in X$ with $\rho(x, y_0) \leq \Lambda$ we can recover with probability at least $2/3$ a point $\hat{y} \in Y$ such that

$$(3.3) \quad \rho(x, \hat{y}) \leq \max((1 + \epsilon)DK \cdot \rho(x, Y), \lambda).$$

Proof. First, we describe the recovery procedure and then show how to sketch points. For now, we assume that for the point of interest $x \in X$ and for every $y \in Y$ we know a number $q(y)$ such that

$$(3.4) \quad \rho(x, y) \leq q(y) \leq D \cdot \rho(x, y).$$

The recovery procedure we describe has several parameters: a positive integer L , a real $0 < \alpha < 1$ and real $\beta, \gamma > D$. For the reasons that will be clear later we require that

$$(3.5) \quad K^2 \cdot (\alpha + 2\gamma) \leq \alpha\beta.$$

The recovery procedure is as follows. First, for every $0 \leq i \leq L$ we build a r_i -net N_i of $Y \cap B_X(y_0, \Lambda)$, where $r_i = 2\alpha^i \Lambda$ such that all pairs of points from N_i have pairwise distances larger than r_i . In particular, $|N_0| = 1$, and for every i the size of N_i is finite since the doubling dimension d of Y is finite. Such a net can be found using a straightforward greedy algorithm. Second, given a point $x \in X$ with $\rho(x, y_0) \leq \Lambda$ we recover an approximate nearest neighbor from Y as follows:

Now let us analyze this procedure. Denote $y^* = \operatorname{argmin}_{y \in Y} \rho(x, y)$ one of the nearest neighbors for x from Y (note that y^* exists, since Y is assumed to be closed). The proof follows from the following three claims (the proofs are in Section 3.1).

CLAIM 3.1. *If (3.5) holds and for some $1 \leq i \leq L$ one has $q(y_{i-1}) \leq \gamma r_{i-1}$, then $\rho(y^*, S_i) \leq r_i$.*

Now let us analyze the case when the algorithm returns y_{i-1} for some $1 \leq i \leq L$.

CLAIM 3.2. *If (3.5) holds and the algorithm returns y_{i-1} for some $1 \leq i \leq L$, then*

$$\frac{\rho(x, y_{i-1})}{\rho(x, y^*)} < \frac{DK\gamma}{\alpha(\gamma - DK)}.$$

Next, suppose that our algorithm returns y_L .

CLAIM 3.3. *If (3.5) holds and the algorithm returns y_L , then $\rho(x, y_L) \leq 2\gamma\alpha^L \Lambda$.*

Let us now show how to set L , α , β and γ . Claims 3.2 and 3.3 imply that in order to satisfy (3.3) we need to satisfy together with (3.5) the following conditions:

$$(3.6) \quad \frac{DK\gamma}{\alpha(\gamma - DK)} \leq (1 + \epsilon)DK,$$

$$(3.7) \quad 2\gamma\alpha^L \Lambda \leq \lambda.$$

It is immediate to see that we can satisfy (3.5), (3.6) and (3.7) simultaneously by setting $\alpha = 1 - \Theta(\epsilon)$, $\beta = \Theta(DK^3/\epsilon)$, $\gamma = \Theta(DK/\epsilon)$ and $L = \Theta\left(\frac{1}{\epsilon} \cdot \log \frac{DK\Lambda}{\epsilon\lambda}\right)$.

So far we assumed that we have access to a function $q(\cdot)$ that satisfies (3.4). In reality we build such a function from sketches of distances between points from X and Y . Suppose we can build a subset⁹ $Q \subseteq Y$ with $|Q| \leq N$ such that for a given $x \in X$ the recovery procedure can query $q(y)$ only for $y \in Q$. Then, we can use the standard amplification argument for the median estimator, and sketch x in size $O(s \log N)$ to get a randomized function $q'(\cdot)$ such that for every $y \in Q$ one has $\Pr[\rho(x, y) \leq q'(y) \leq D \cdot \rho(x, y)] \geq 1 - \frac{1}{3N}$. Now we use $q'(\cdot)$ for the recovery and by the union bound the recovery algorithm succeeds with probability at least $2/3$. It is only left to upper bound N for an appropriately chosen set Q .

It is clear that we query $q(\cdot)$ for points only from

$$\bigcup_{\substack{i \in [L] \\ q(x, y_{i-1}) \leq \gamma \cdot r_{i-1}}} S_i \subseteq \bigcup_{\substack{i \in [L] \\ \rho(x, y_{i-1}) \leq \gamma \cdot r_{i-1}}} S_i$$

(the inclusion is by (3.4)). By Claim 3.1 the right-hand side is included in

$$Q = \bigcup_{1 \leq i \leq L} (N_i \cap B_X(y^*, K^2 \cdot (1 + 2\beta) \cdot r_i)).$$

Since Y has doubling dimension d and points from N_i are r_i -separated, we get $N = |Q| \leq L \cdot (K^2 \cdot (1 + 2\beta))^{O(d)}$.

⁹Note that Q is more than just a single path from the “root” to the solution, as the behavior of the algorithm is not deterministic and depends on the random bits chosen by the sketching procedure.

Now using the values of L , α , β and γ , we get that the final sketch size is $O(\log N) \leq O(s \cdot (d \log(K \cdot (1 + \beta)) + \log L)) \leq O\left(s \cdot \left(d \log(DK/\epsilon) + \log \log(\Lambda/\lambda)\right)\right)$.

COROLLARY 3.1. *Suppose that X is induced by a norm of dimension n , and that there is an algorithm that computes the sets S_i defined by the search procedure in time $|S_i|^{O(1)}$. Then the search procedure runs in time polynomial in $N = L \cdot (K^2 \cdot (1 + 2\beta))^{O(d)}$ and n .*

3.1 Proofs

Proof. [Proof of Claim 3.1] Let $y' \in N_i$ be a point such that $\rho(y^*, y') \leq r_i$ (recall that N_i is an r_i -net of Y). Clearly, it is sufficient to prove that $y' \in S_i$. This is equivalent to the condition $\rho(y', y_{i-1}) \leq \beta r_i$. Let us verify the latter:

$$\begin{aligned} \rho(y', y_{i-1}) &\leq K \cdot (\rho(y^*, y') + \rho(y^*, y_{i-1})) \\ &\leq K \cdot (r_i + \rho(y^*, y_{i-1})) \leq K^2 \cdot (r_i + \rho(x, y^*) + \rho(x, y_{i-1})) \leq \\ &\leq K^2 \cdot (r_i + 2\rho(x, y_{i-1})) \leq K^2 \cdot (r_i + 2q(y_{i-1})) \\ &\leq K^2 \cdot (r_i + 2\gamma r_{i-1}) = K^2 \cdot (\alpha + 2\gamma) \cdot r_{i-1} \leq \alpha \beta r_{i-1} = \beta r_i, \end{aligned}$$

where the third inequality follows from the definition of y' , the fourth inequality follows from the definition of y^* , the fifth inequality follows from (3.4), the sixth step follows from the statement of the Claim, and the penultimate step follows from (3.5).

Proof. [Proof of Claim 3.2] First, observe that by (3.4) and the fact that the algorithm returns y_{i-1} we have $\rho(x, y_{i-1}) \leq q(y_{i-1}) \leq \gamma r_{i-1}$. Second, by (3.4) and Claim 3.1,

$$\begin{aligned} \frac{\gamma r_i}{D} &< \frac{q(y_i)}{D} \leq \rho(x, y_i) = \rho(x, S_i) \\ &\leq K \cdot (\rho(x, y^*) + \rho(y^*, S_i)) \leq K \cdot (\rho(x, y^*) + r_i). \end{aligned}$$

Thus,

$$\rho(x, y^*) > \left(\frac{\gamma}{DK} - 1\right) \cdot r_i.$$

Overall,

$$\frac{\rho(x, y_{i-1})}{\rho(x, y^*)} < \frac{\gamma r_{i-1}}{\left(\frac{\gamma}{DK} - 1\right) \cdot r_i} = \frac{DK\gamma}{\alpha(\gamma - DK)}.$$

4 General lower Bound on Measurement Complexity

We use $a \lesssim b$ to denote that there exists a universal constant C such that $a \leq Cb$. We use $a \gtrsim b$ to denote $b \lesssim a$ and $a \approx b$ to denote $a \lesssim b \lesssim a$.

We work with the linear sparse recovery scheme as in the Introduction (Equation (1.1)). We set the probability of error to be $P = \frac{1}{4}$.

The following lemma generalizes the result in [DIPW10] to general norms and nonnegative inputs.

LEMMA 4.1. *Consider any norm $\|\cdot\|$ over \mathbb{R}^n for which $\frac{1}{n^c} \leq \frac{\|x\|}{\|x\|_2} \leq n^c$ for some constant c . Further suppose that there exists a set $X \subset [0, \infty)^n$ of k -sparse vectors such that $\|x\| \approx 1$ for all $x \in X$ and $\|x - x'\| \gtrsim 1$ for all $x \neq x' \in X$. Then any linear sparse recovery scheme with approximation factor K over $[0, \infty)^n$ must use $m \gtrsim \frac{\log|X|}{\log K}$ linear measurements.*

Proof. We first show a set of assumptions we can make without loss of generality, then give an algorithm to solve augmented indexing using sparse recovery, then analyze the algorithm.

WLOG assumptions and setup. First, we show that we can assume that $x \in X$ have coordinates that are multiples of $1/n^{c+1}$. Let x' be x rounded to the nearest multiple of $1/n^{c+1}$ in each coordinate, so $\|x - x'\|_\infty \leq 1/n^{c+1}$. Therefore $\|x - x'\|_2 \leq \sqrt{n}/n^{c+1}$ or $\|x - x'\| \leq 1/\sqrt{n}$. This means that replacing x with x' would also satisfy the conditions with negligibly worse constants and have coordinates that are multiples of $1/n^{c+1}$.

We would like to give a lower bound for all randomized sparse recovery schemes that work for each input with $3/4$ probability. By Yao's minimax principle, it suffices to give an explicit distribution on inputs for which no deterministic sparse recovery scheme (A, \mathcal{A}) can work with $3/4$ probability. Furthermore, we may assume that $A \in \mathbb{R}^{m \times n}$ has orthonormal rows (otherwise, if $A = U\Sigma V^T$ is its singular value decomposition, $\Sigma^+ U^T A$ has this property and the transformation can be inverted before applying the algorithm).

We use the following lemma from [DIPW10]:

LEMMA 4.2. *Consider any $m \times n$ matrix A with orthonormal rows. Let A' be the result of rounding A to b bits per entry. Then for any $v \in \mathbb{R}^n$ there exists an $s \in \mathbb{R}^n$ with $A'v = A(v - s)$ and $\|s\|_1 < n^2 2^{-b} \|v\|_1$.*

Proof. Let $A'' = A - A'$ be the roundoff error, so each entry of A'' is less than 2^{-b} . Then for any v and $s = A^T A'' v$, we have $As = A'' v$ and

$$\begin{aligned} \|s\|_1 &= \|A^T A'' v\|_1 \leq \sqrt{n} \|A'' v\|_1 \leq m \sqrt{n} 2^{-b} \|v\|_1 \\ &\leq n^2 2^{-b} \|v\|_1. \end{aligned}$$

Now, let A' be A rounded to $c' \log n$ bits per entry for c' to be chosen later. By Lemma 4.2, for any v we have $A'v = A(v - s)$ for some s with $\|s\|_1 \leq n^2 2^{-c' \log n} \|v\|_1$, so

$$\|s\| \leq n^{2c+2-c'} \|v\|.$$

We are now ready to construct the lower bound of m via a reduction from the one-way augmented indexing problem in communication complexity. In this problem, Alice has a bit string b of length $r \log |X|$ for $r = \log n$, and Bob has an index $i^* \in [r \log |X|]$ as well as b_1, \dots, b_{i^*-1} . Alice must send a message to Bob, who must output b_{i^*} with $2/3$ probability. It is known that the message must contain $\Omega(r \log |X|) = \Omega(\log n \log |X|)$ bits. We will show a way to use the sparse recovery algorithm to solve augmented indexing with $O(m \cdot \log n \cdot \log K)$ bits, giving the lower bound of $m \gtrsim \frac{\log |X|}{\log K}$.

Algorithm to solve augmented indexing.

Alice turns her $r \log |X|$ bits into a list $x_1, \dots, x_r \in X$. She then defines

$$z = \sum_{i=1}^r x_i / (KC)^i$$

for a sufficiently large constant integer C to be specified later, and

$$y = A'z.$$

Since $\|z\| \leq \sum_{i=1}^r \|x_i\| / (KC)^i \lesssim 1$, we have that $y = A(z - s)$ for some s with $\|s\| \leq n^{2c+2-c'}$. Alice then sends y to Bob.

Transmitting y takes $O(m \cdot \log n \cdot \log K)$ bits. To see this, note that each coordinate of z is a multiple of $\frac{1}{n^{c+1}(KC)^r}$ that is at most n^c , and each coordinate of A' is a multiple of $1/n^{c'}$ that is at most 1. Hence each coordinate of $y = A'z$ is a multiple of $\frac{1}{n^{c'+c+1}(KC)^r}$ that is at most n^{c+1} , which can be represented in $\log(n^{c'+2c+2}(KC)^r) \lesssim \log n \cdot (c' + \log K)$ bits. There are m coordinates, so transmitting y takes $O(m \cdot \log n \cdot (c' + \log K))$ bits.

Now, based on his inputs b_1, \dots, b_{i^*-1} and i^* , Bob can figure out x_1, \dots, x_{i^*-1} and wants to figure out x_{i^*} for $i^* = 1 + \lfloor i^* / \log |X| \rfloor$. Once he learns $y = A'z = A(z - s)$, Bob chooses $u \in [0, \frac{1}{KCn^{c+1}}]^n$ uniformly at random, and computes

$$y' = (KC)^{i'} (y - A \sum_{i=1}^{i'-1} x_i / (KC)^i) + Au.$$

Bob then performs sparse recovery using \mathcal{A} on y' getting a result \hat{x} . He rounds \hat{x} to the $x \in X$ minimizing $\|x - \hat{x}\|$. We will show that $x = x_{i^*}$ with at least $2/3$ probability; if this happens, Bob can recover b_{i^*} from the associated vector x_{i^*} .

Analysis of algorithm. We have that $y' = A(z' - s + u)$ for z', s with $\|s\| \lesssim n^{2c+2-c'}$ and

$$z' = x_{i'} + \sum_{j=1}^{r-i'} x_{i'+j} / (KC)^j =: x_{i'} + w$$

for $w = \sum_{j=1}^{r-i'} x_{i'+j} / (KC)^j$ having $\|w\| \lesssim 1/(KC)$. Then $y' = A(x_{i'} + w + u - s)$.

For now, pretend that Bob performed sparse recovery on $A(x_{i'} + w + u)$ instead of $A(x_{i'} + w + u - s)$. The distribution of $x_{i'} + w + u$ depends on the distribution of inputs to the augmented indexing problem, but it is independent of the choice of A and is over $[0, \infty)^n$. Therefore we can choose our A to be a matrix that lets us perform sparse recovery with $3/4$ probability over this distribution. Then the result \hat{x} of sparse recovery satisfies

$$(4.8) \quad \begin{aligned} \|\hat{x} - (x_{i'} + w + u)\| &\lesssim K \min_{k\text{-sparse } \bar{x}} \|x_{i'} + w + u - \bar{x}\| \\ &\leq K \|w + u\| \end{aligned}$$

with $3/4$ probability, or

$$(4.9) \quad \|\hat{x} - x_{i'}\| \lesssim K(\|w\| + \|u\|) \lesssim 1/C$$

If C is a sufficiently large constant, this is less than

$$(4.10) \quad \min_{x \neq x' \in X} \|x - x'\| / 2 \gtrsim 1.$$

Therefore, when Bob rounds \hat{x} to X , he gets $x_{i'}$ whenever sparse recovery succeeds, as happens with $3/4$ probability.

In fact, Bob performs sparse recovery on $A(x_{i'} + w + u - s)$ not $A(x_{i'} + w + u)$. However, the latter is statistically close to the former. In particular, $\|s\|_\infty \lesssim n^{3c+2-c'}$ so that the total variation distance

$$\text{TV}(u, u - s) \lesssim n \cdot \frac{n^{3c+2-c'}}{1/(Kn^{c+1})} \leq Kn^{4c+4-c'}$$

Setting $c' = 4c + 5 + \frac{\log K}{\log n}$, we get that

$$\text{TV}(A(x_{i'} + w + u), A(x_{i'} + w + u - s)) \leq \text{TV}(u, u - s) \lesssim 1/n.$$

Therefore Bob's rounding of \hat{x} to X will equal $x_{i'}$ with probability at least $3/4 - O(1/n) > 2/3$. This solves the augmented indexing problem with only $O(m \log n \cdot (c' + \log K)) = O(m \log n \cdot \log K)$ bits of communication. Since augmented indexing requires $\Omega(r \log |X|) = \Omega(\log n \log |X|)$ bits of communication in this setting, we have $m \gtrsim \frac{\log |X|}{\log K}$.

Proof. [Proof of Theorem 1.2.] Define $S = [0, \infty)^n \cap T_k \cap \{x \in \mathbb{R}^n : \|x\| \leq 1\}$.

Because the space and the norm are homogeneous, we have by the definition of doubling dimension that covering S requires 2^d balls of radius $1/2$. Therefore we can find a packing $X \subset S$ of 2^d points such that $\min_{x \neq x' \in X} \|x - x'\| \geq 1/2$. This also means at most one $x \in X$ has $\|x\| < 1/4$. Throwing this possible element out, we get a set of size $2^d - 1$ satisfying the constraints of Lemma 4.1, giving that $m \gtrsim \frac{\log(2^d - 1)}{\log K} \gtrsim \frac{d}{\log K}$.

4.1 Lower bound for streaming algorithms In this section we show a lower bound on the space bit complexity of any streaming algorithm that maintains an approximately best k -sparse approximation of a vector with respect to any norm $\|\cdot\|$ on \mathbb{R}^n such that

$$(4.11) \quad n^{-O(1)} \leq \frac{\|x\|}{\|x\|_2} \leq n^{O(1)}$$

for every $x \in \mathbb{R}^n$.

THEOREM 4.1. *Suppose that there is an algorithm that can maintain a vector $x \in \mathbb{R}^n$ under updates of the form $x_i := x_i + \delta_i$, where $\delta_i \in \mathbb{Z}$, and, moreover, suppose that we are promised that all entries of x at any moment of time are integers between 0 and $n^{O(1)}$. In the end, the algorithm is required to output a vector y such that*

$$\|x - y\| \leq K \cdot \min_{k\text{-sparse } x^*} \|x - x^*\|,$$

where $K > 2$ is some approximation factor. Then, the space bit complexity of the algorithm must be at least

$$\Omega\left(d \cdot \frac{\log n}{\log K}\right),$$

where d is the doubling dimension of the non-negative k -sparse vectors under $\|\cdot\|$.

The rest of this section is devoted to proving this Theorem. We roughly follow the above argument for proving the lower bound for sparse recovery. However, in this case, the argument is even simpler since we do not need to handle issues related to the sketching matrix.

First, we take $r = 2^{\Omega(d)}$ non-negative k -sparse vectors v_1, \dots, v_r whose $\|\cdot\|$ -norm is $\Theta(1)$ and that are $\Omega(1)$ -pairwise separated wrt $\|\cdot\|$. We will show how Alice and Bob can solve Augmented Indexing on $b = \Omega\left(d \cdot \frac{\log n}{\log K}\right)$ -bit strings using the assumed algorithm. Alice partitions her b -bit sequence into blocks of length $\log r = \Omega(d)$, encodes each block in one of the v_i 's (denote it by u_j for $1 \leq j \leq b/\log r$), and then feeds the (properly rescaled and discretized) vector

$$U = \sum_{j=1}^{b/\log r} \frac{u_j}{(CK)^j},$$

where $C > 0$ is a sufficiently large constant, to the algorithm. Bob takes over, starting from this moment, subtracts the part of U that corresponds to his prefix and then uses the algorithm to recover the next u_j .

Overall, we have that the required space is at least $\Omega(b) = \Omega\left(d \cdot \frac{\log n}{\log K}\right)$. The only remaining fact we need

to argue about is why the accuracy $[0, n^{O(1)}]$ per entry is sufficient. First, we use (4.11) to claim that the polynomial in n accuracy is enough to represent v_i 's with the required conditions. Second, since $b/\log r \approx \log n/\log K$, we get that U can be represented with accuracy polynomial in n .

5 Doubling dimension of EMD

5.1 Upper bounds

5.1.1 General measures We will prove that the doubling dimension of k -sparse probability measures over $[\Delta]^2$ equipped with EMD is $O(k \log \log \Delta)$. For a weaker and simple bound $O(k \log k)$ on the doubling dimension in the case of k -sparse subsets, see [GKK10]. In fact, it is not hard to prove upper bound $O(k)$ on the doubling dimension for k -sparse subsets. Notice that in this case the upper bound on the doubling dimension does not depend on the size of the grid. We will now provide an intuition why the upper bound $O(k)$ holds.

We have an EMD ball $\text{Ball}_{\text{EMD}}(\mu, R)$ or radius R centered at k -sparse measure μ such that $\mu(x, y) = 1$ for all $(x, y) \in \text{supp}(\mu)$. (We can think of μ as a k -sparse set.) And we would like to cover all k -sparse subsets within $\text{Ball}_{\text{EMD}}(\mu, R)$ with $2^{O(k)}$ EMD balls of radius $R/2$ centered at k -sparse subsets.

First, let's show how to cover all subsets $o \in \text{Ball}_{\text{EMD}}(\mu, R)$ with o satisfying $\|o_i - \mu_i\|_1 = \Theta(R/k)$ for all $i \in [k]$. o_i and μ_i denote points in $\text{supp}(o)$ and $\text{supp}(\mu_i)$ and they get matched together in the optimal transportation between o and μ . For this, we take $R/(100k)$ -net of $\text{Ball}_{\ell_1}(\mu_i, 10R/k)$ for every $i \in [k]$. Every such net is of size $O(1)$. To cover all the $o \in \text{Ball}_{\text{EMD}}(\mu, R)$, we need to take a representative from a net from $\text{Ball}_{\ell_1}(\mu_i, 10R/k)$ for all $i \in [k]$ and combine representatives in k -sparse subsets. There are $2^{O(k)}$ possible ways to construct subsets by taking representatives.

In the case when we do not have the mentioned guarantee at the beginning of the previous paragraph, we can guess values $\|o_i - \mu_i\|_1$ up to a constant factor and construct covers for all guesses. We need to show that it is enough to take at most $2^{O(k)}$ guesses. And it can indeed be shown by noticing that we do not need to cover ℓ_1 balls of very small radius (when $\|o_i - \mu_i\|_1$ is small).

We proceed by showing upper bound on the doubling dimension when we consider arbitrary measures with support of size at most k living in a square of side length Δ .

LEMMA 5.1. *The doubling dimension of the set of k -sparse probability measures over $[\Delta]^2$ under EMD met-*

ric is $O(k \log \log \Delta)$.

Proof. Let μ be a k -sparse probability measure over $[\Delta]^2$ and let $R > 0$ be some real number. Our goal is to cover $B_{\text{EMD}}(\mu, R)$ with $\log^{O(k)} \Delta$ EMD-balls centered in k -sparse measures and of radius $R/2$. In order to achieve this it is sufficient to cover $B_{\text{EMD}}(\mu, R)$ with $\log^{O(k)} \Delta$ EMD-balls centered in *arbitrary* measures and of radius $R/4$.

The pseudocode in Figure 3 builds a set of measures \mathcal{M} that serve as centers of balls with radius $R/4$ that together cover $B_{\text{EMD}}(\mu, R)$. Roughly speaking, we first guess the topology of the optimal flow. Then we guess the lengths of the corresponding edges. Then we guess the support. And finally we guess the masses transported over the edges.

We assume that $\text{BUILDNET}(p, r)$ returns an $(r/100)$ -net of $B_{\ell_1^2}(p, r) \cap [\Delta]^2$. It is immediate that $|\mathcal{M}| \leq \log^{O(k)} \Delta$, and that the running time of the above procedure is also $\log^{O(k)} \Delta$. It is left to show that for every k -sparse μ' such that $\text{EMD}(\mu, \mu') \leq R$, there exists $\mu'' \in \mathcal{M}$ with $\text{EMD}(\mu'', \mu') \leq R/4$.

CLAIM 5.1. *There exists an optimal flow between μ and μ' that is supported on at most $2k$ pairs of points.*

Proof. Consider an optimal flow from μ to μ' . Consider an undirected graph $G = (V, E)$ with $V = \text{supp}(\mu) \cup \text{supp}(\mu')$. We connect two vertices $(x, y) \in \text{supp}(\mu)$ and $(x', y') \in \text{supp}(\mu')$ iff there non-zero amount flowing from (x, y) to (x', y') in the flow.

If $|E| \geq 2k + 1$, then there is a cycle $e_1, e_2, \dots, e_{2m} \in E$ of even length $2m$ in the graph G . W.l.o.g. assume that the total length of e_i with even i is at most the total length of e_i with odd i . Let us increase all flows over e_i with even i and decrease all flows over e_i with odd i by the same amount such that at least one edge carries zero flow. Clearly, the total cost can only decrease.

Repeating the above process several times, we arrive at a flow supported on at most $2k$ edges.

Thus, in our enumeration algorithm at least one $c(x, y)$ corresponds to the number of outgoing flow edges from $(x, y) \in \text{supp} \mu$ (Line 6). When we enumerate l there is at least one choice that guesses all the lengths of the corresponding edges within a multiplicative factor of 1.01 (Line 8). Thus, there exists a measure $\tilde{\mu}$ (not necessarily k -sparse) such that

- $\text{supp} \tilde{\mu} \subseteq \text{supp} \mu \cup \{p(x, y, i)\}_{(x, y, i) \in \mathcal{I}}$;
- $\text{EMD}(\mu', \tilde{\mu}) \leq R/50$;
- there exists a flow between μ and $\tilde{\mu}$ of cost at most $1.02 \cdot R$ that transports mass from a point

(x, y) to $\{(x, y)\} \cup \{p(x, y, i)\}_{i: (x, y, i) \in \mathcal{I}}$ for every $(x, y) \in \text{supp} \mu$.

We have that $\tilde{\mu}$ covers μ' , i.e., that $\text{EMD}(\mu', \tilde{\mu}) \leq R/100$, and that the procedure in the pseudocode will guess $\text{supp}(\tilde{\mu})$ but not necessarily $\tilde{\mu}$. After guessing the support (Line 9), the pseudocode proceeds by trying to guess the measure at the support (Line 10). We will show that there will be guess μ'' made by the pseudocode with $\text{supp}(\mu'') \subseteq \text{supp}(\tilde{\mu})$ that satisfy $\text{EMD}(\tilde{\mu}, \mu'') \leq 2R/50$.

Fix $(x, y) \in \text{supp} \mu$. We show to deal with the multi-set $\{(x, y)\} \cup \{p(x, y, i) : (x, y, i) \in \mathcal{I}\}$. We round down the mass in $\tilde{\mu}$ at the coordinates $\{p(x, y, i)\}$ to the closest element of $\{0, m_0, 1.01 \cdot m_0, 1.01^2 \cdot m_0^2, \dots, \min(1, R)\}$ (Line 11). Let μ'' be the resulting measure. We also set $\mu''((x, y)) := \sum_{i: (x, y, i) \in \mathcal{I}} (\tilde{\mu}(p(x, y, i)) - \mu''(p(x, y, i)))$ (Line 19). One can observe that μ'' is included in the set measures enumerated by our algorithm.

We now show that $\text{EMD}(\mu'', \tilde{\mu}) \leq 2R/50$. The cost of $\text{EMD}(\mu'', \tilde{\mu})$ comes from two sources:

1. Contribution from $(x, y, i) \in \mathcal{I}$ for which $\tilde{\mu}(p(x, y, i)) < m_0$. Then $\mu''(p(x, y, i)) = 0$. There are at most $2k$ such $(x, y, i) \in \mathcal{I}$. But we can reroute these small masses with cost at most $2\Delta k m_0 \leq 0.02R$.
2. Contribution from $(x, y, i) \in \mathcal{I}$ for which $\tilde{\mu}(p(x, y, i)) \geq m_0$. This implies that the value of $\tilde{\mu}(p(x, y, i))$ is within 1% of $\mu''(p(x, y, i))$. Therefore, the total contribution of such coordinates $(x, y, i) \in \mathcal{I}$ is at most $0.01 \cdot \text{EMD}(\mu, \tilde{\mu}) \leq 0.02R$.

Thus, overall we have $\|\mu' - \mu''\|_{\text{EMD}} \leq \|\mu' - \tilde{\mu}\|_{\text{EMD}} + \|\tilde{\mu} - \mu''\|_{\text{EMD}} \leq (0.02 + 0.02 + 0.02) \cdot R < R/4$.

5.1.2 Measures with bounded granularity

LEMMA 5.2. *Consider set S of all k -sparse measures μ such that, for all coordinates (x, y) , $\mu(x, y)$ is equal to i/N for some non-negative integer i and the total measure of μ is 1. The set S , under EMD, has doubling dimension $O(k \log \log N)$.*

Proof. Let $\text{Ball}_{\text{EMD}}(p, r)$ be EMD ball of radius r containing k -sparse probability measures of granularity $1/N$ over plane. $p \in S$ is the center of the ball. Further down we will denote $\text{Ball}_{\text{EMD}}(p, 1)$ by B .

Case 1. $|\text{supp}(p)| = 1$. WLOG, the entire probability mass of p is at point $(0, 0)$. We can verify that $\mu \in B$ implies $\text{supp}(\mu) \subseteq [-100N, 100N]^2$.

Let B' be the set of all probability measures μ with properties that μ has granularity $1/n$, $\text{supp}(\mu) \subseteq$

```

1:  $m_0 \leftarrow R/(100\Delta k)$ 
2:  $\mathcal{M} \leftarrow \emptyset$ 
3: for  $c: \text{supp } \mu \rightarrow \mathbb{Z}_{>0}$  such that  $\sum_{(x,y) \in \text{supp } \mu} c(x,y) \leq 2k$  do
4:    $\mathcal{I} \leftarrow \{(x,y,i) \mid (x,y) \in \text{supp } \mu, 1 \leq i \leq c(x,y)\}$ 
5:   for  $l: \mathcal{I} \rightarrow \{1, 1.01, 1.01^2, \dots, 2\Delta\}$  do
6:     for  $(x,y,i) \in \mathcal{I}$  and for all  $p(x,y,i) \in \text{BUILDNET}((x,y), l(x,y,i))$  do
7:       for  $m: \mathcal{I} \rightarrow \{0, m_0, 1.01 \cdot m_0, 1.01^2 \cdot m_0, \dots, \min(1, R)\}$  do
8:         if for every  $(x,y) \in \text{supp } \mu$  we have  $\sum_{i: (x,y,i) \in \mathcal{I}} m(x,y,i) \leq \mu(x,y)$  then
9:           let  $\mu'$  be a measure over  $[\Delta]^2$  that is identically zero
10:          for  $(x,y) \in \text{supp } \mu$  do
11:             $s \leftarrow 0$ 
12:            for  $i: (x,y,i) \in \mathcal{I}$  do
13:               $s \leftarrow s + m(x,y,i)$ 
14:               $\mu'(p(x,y,i)) \leftarrow \mu'(p(x,y,i)) + m(x,y,i)$ 
15:               $\mu'(x,y) \leftarrow \mu'(x,y) + \mu(x,y) - s$ 
16:             $\mathcal{M} \leftarrow \mathcal{M} \cup \{\mu'\}$ 

```

Figure 3: Pseudocode for the net construction

$[-100N, 100N]^2$, all coordinates of points from $\text{supp}(\mu)$ are of the form $\frac{i}{1000k}$ for an integer i . We can verify that, for every $\mu \in B$, there exists $\mu' \in B'$ with $\|\mu - \mu'\| \leq \frac{|\text{supp}(\mu)|}{1000k} \leq 1/1000$.

Therefore, if we construct $\frac{1}{60}$ -cover X of B' as per Lemma 5.1 of size $|X| = (\log N)^{O(k)}$, X is also $\frac{1}{100}$ -cover of B and we get the required upper bound.

There is one issue, though. It might be that the measures from cover X does not have granularity $1/N$. To deal with this, we first build $\frac{1}{200}$ -cover X of B' according to Lemma 5.1. Then, for every measure μ from the cover, if μ is not of granularity $1/N$ and EMD ball of radius $\frac{1}{200}$ around μ does not contain any measure of granularity $1/N$, we discard this measure because it does not cover any measure of interest. If, on the other hand, the EMD ball of radius $1/200$ does contain measure μ' of granularity $1/N$, we replace μ with μ' in the cover. Clearly, increasing the radius by a factor of 2, still covers all the previous points, i.e., $\text{Ball}_{\text{EMD}}(\mu, 1/200) \subseteq \text{Ball}_{\text{EMD}}(\mu', 1/100)$.

Case 2. $|\text{supp}(p)| > 1$. We denote elements of $\text{supp}(p)$ by (x,y) . Because of granularity of measures, if $\mu \in B$, then $\text{supp}(\mu) \subseteq \bigcup_{(x,y) \in \text{supp}(p)} S_{(x,y)}$, where

$$S_{(x,y)} = [x - 100N, x + 100N] \times [y - 100N, y + 100N]$$

denotes a square in plane with side length $200N$.

We construct a graph with vertices $S_{(x,y)}$, $(x,y) \in \text{supp}(p)$. We connect two vertices if the corresponding squares have non-empty intersection. We consider connected components of the resulting graph. We want to move the connected components so that, in the end, all of them live inside a square of side length $10^9 N^2$

and distance between any two connected components is $\geq 10^5 N$. We can verify that we can do that.

Let p' denote the resulting measure and $(x', y') \in \text{supp}(p')$ be the resulting elements of the support. We round the coordinates of the elements of $\text{supp}(p')$ so that all x' and y' are of the form $\frac{i}{1000k}$ for some integer i . Let p'' be the measure after the rounding.

We can check that $\text{EMD}(p', p'') \leq \frac{\text{supp}(p')}{1000k} \leq 1/1000$. Therefore, if we construct $\frac{1}{100}$ -cover of $\text{Ball}_{\text{EMD}}(p'', 1.1)$, we get $\frac{1}{100}$ -cover of $\text{Ball}_{\text{EMD}}(p', 1)$.

Consider all probability measures from $\text{Ball}_{\text{EMD}}(p'', 1.2)$ with the property that all coordinates of elements of supports of measures have form $\frac{i}{1000k}$ for some integer i . We denote this set by $\text{Ball}'_{\text{EMD}}(p'', 1.2)$. $\frac{1}{200}$ -cover for $\text{Ball}'_{\text{EMD}}(p'', 1.2)$ gives $\frac{1}{200} + \frac{1}{1000} < \frac{1}{100}$ -cover for $\text{Ball}_{\text{EMD}}(p'', 1.1)$.

To construct $\frac{1}{200}$ -cover for $\text{Ball}'_{\text{EMD}}(p'', 1.2)$, we start by constructing $\frac{1}{400}$ -cover of $\text{Ball}'_{\text{EMD}}(p'', 1.2)$ by measures not necessarily having granularity $\frac{1}{N}$. To get measures with granularity $\frac{1}{N}$, we proceed in the same way as in Case 1, i.e., we consider 2 cases. If a measure in the $\frac{1}{400}$ -cover does not have a measure of granularity $\frac{1}{N}$ within EMD distance $\frac{1}{400}$, then discard this measure from the cover. Otherwise, replace the measure with the measure that has granularity $\frac{1}{N}$. We can see that the set of measures that these operations produce, is $2 \cdot \frac{1}{400} = \frac{1}{200}$ -cover $\text{Ball}'_{\text{EMD}}(p'', 1.2)$ and has granularity $\frac{1}{N}$. From Lemma 5.1, the size of the cover is $(\log N)^{O(k)}$. As a result, we have $\frac{1}{100}$ -cover of $\text{Ball}_{\text{EMD}}(p', 1)$. All measures in the cover have granularity $\frac{1}{N}$.

Given $\frac{1}{100}$ -cover of $\text{Ball}_{\text{EMD}}(p', 1)$, we would like to construct $\frac{1}{100}$ -cover of B . Given that p and all measures

from the cover have granularity $\frac{1}{N}$, we can make the following assumption. The optimal transportation of probability measure from p to every measure from the cover has probability mass on every edge of amount $\frac{i}{N}$ for some non-negative integer i . As a result, if μ is a measure from $\frac{1}{100}$ -cover of $\text{Ball}_{\text{EMD}}(p', 1)$, in the optimal transportation of p' into μ (that achieves cost $\text{EMD}(p', \mu)$), μ has non-zero amount on edges to elements of $\text{supp}(p')$ that corresponds to at most one component of the graph. (This follows because the connected components are highly separated in p' .) This gives that we can move components independently. We move the components to their original positions (according in p) and accordingly transform measures in the cover. This gives $\frac{1}{100}$ -cover for B .

5.2 Lower bounds

LEMMA 5.3. *Weighted point-sets over $[\Delta]$ of cardinality k under EMD has doubling dimension $\Omega(k) \cdot \log(\Omega(\log \frac{\Delta}{k}))$ for $k > 1$.*

Proof. WLOG, we assume that Δ is an integer power of 2. By (x, w) we denote a point with coordinate x and weight w . Let set A be a weighted point-set of size k . For $i = 1, 2, 3, \dots, k/2$, we set the i -th point of A to be $A_i = (2i\Delta/k, 2)$. The remaining $k/2$ points has weight 0 and arbitrary coordinates on the line.

Let $I = i_1, i_2, \dots, i_{k/2}$ for $0 \leq i_j \leq \log U$ (we will later set $U = \Delta/k$), and let B_I be a point-set defined as

$$B_I = \cup_{j=1}^{k/2} \{(2j\Delta/k, 2 - 2^{-i_j}), (2j\Delta/k + 2^{i_j}, 2^{-i_j})\},$$

We constructed B_I such that $\text{EMD}(A, B_I) = k/2$ for all I .

Consider an EMD ball of radius $k/2$ around A , i.e., $\text{Ball}_{\text{EMD}}(A, k/2)$. We will show that the number of EMD balls of radius $k/200$ needed to cover it is $(\log \frac{\Delta}{k})^{\frac{k}{2} \cdot \frac{9}{10}} / 2^k$, which yields the result.

Consider a $\frac{k}{200}$ -cover of $\text{Ball}_{\text{EMD}}(A, k/2)$. We will show that the size of the cover must be large. Consider a weighted pointset B_I for some $I = i_1, i_2, \dots, i_{k/2}$. Given that B_I is covered by an element from the cover, there must be an element C from the cover with the property that at least 9/10 fraction of intervals

$$[2j\Delta/k + 2^{i_j} - 2^{i_j}/10, 2j\Delta/k + 2^{i_j} + 2^{i_j}/10]$$

(for $j = 1 \dots k/2$) contains an element from $\text{supp}(C)$. We call that C hits B_I and the set of elements of $\text{supp}(C)$ that is contained in some interval we call the hitting set of B_I . Otherwise, for any C that does not satisfy this property, we have

$$\text{EMD}(B_I, C) > \left(1 - \frac{9}{10}\right) \cdot |\text{supp}(B_I)| \cdot \frac{1}{10} = \frac{k}{200}.$$

There are

$$\begin{aligned} & |\{B_I | I = i_1, i_2, \dots, i_{k/2} \text{ and } 0 \leq i_j \leq \log U \text{ for } j \in [k/2]\}| \\ & = \left(\log \frac{\Delta}{k}\right)^{\frac{k}{2}} \end{aligned}$$

pointsets B_I that are covered.

Consider an element C from the cover. We will show that C can hit at most $2^k \cdot (\log \frac{\Delta}{k})^{\frac{k}{2} \cdot (1 - \frac{9}{10})}$ sets B_I . This will finish the proof.

There are at most 2^k subsets D of $\text{supp}(C)$ that can be a hitting set for some B_I . Every D can be a hitting set for at most $(\log \frac{\Delta}{k})^{\frac{k}{2} \cdot (1 - \frac{9}{10})}$ sets B_I because $|D| \geq \frac{9}{10} \cdot \frac{k}{2}$. This finishes the proof.

COROLLARY 5.1. *Weighted point-sets over $[\Delta]^2$ of cardinality k under EMD has doubling dimension $\Omega(k) \cdot \log(\Omega(\log \frac{\Delta^2}{k}))$ for $k > 1$.*

Proof. We want to choose a point-set A and a lot of “highly separated” point-sets B_I similarly as in 5.3 with $\text{EMD}(A, B_I) = k/4$.

For that, we place $k/4$ points with non-zero weight on a line of length $\Delta \cdot \frac{\sqrt{k}}{2}$ to construct A and B_I s analogously as in Lemma 5.3. The difference is that, instead of placing $k/2$ points, we place $k/4$ points and that, instead of having an interval of length Δ , we have interval of length $\Delta \cdot \frac{\sqrt{k}}{2}$. Then we split points of A with their counterparts of B_I into $\sqrt{k}/2$ consecutive sequences of points each containing $\sqrt{k}/2$ points. We put i -th sequence in $i \cdot \frac{2\Delta}{\sqrt{k}}$ -row of the grid.

We can verify that the resulting point-sets satisfy the necessary properties.

6 Better Sketches for EMD

6.1 Refined analysis of the grid embedding In this section we recall the embedding of $\text{EMD}_{[\Delta]^2}$ into ℓ_1 from [IT03] (building on [Cha02]) and provide the refined analysis of a variant of it, under the assumption that we are embedding a measure that can be represented as a difference of two non-negative measures that both sum to one such that one of them is k -sparse.

We state the following simple lemma without a proof.

LEMMA 6.1. *For any vector $y : [\Delta] \rightarrow \mathbb{R}_+$, define $\text{CDF}(y) : [\Delta] \rightarrow \mathbb{R}_+$ by*

$$\text{CDF}(y)_i = \sum_{j \leq i} y_j.$$

Then for any $y, y' : [\Delta] \rightarrow \mathbb{R}_+$ with $\|y\| = \|y'\| = 1$ we have

$$\text{EMD}(y, y') = \|\text{CDF}(y) - \text{CDF}(y')\|_1.$$

The Cauchy distribution is continuous probability distribution with the probability density function $\frac{1}{\pi(\gamma + \frac{x^2}{\gamma})}$, where γ is the scale parameter. If not otherwise specified, we will refer to a Cauchy variable as one which is drawn from distribution with $\gamma = 1$.

First, we need the following folklore claim that will be useful for us later.

CLAIM 6.1. *Let X_1, X_2, \dots, X_n are (not necessarily independent) non-negative random variables such that for every i and $t > 0$ we have*

$$\Pr[X_i \geq t] \leq \frac{C}{t},$$

where $C > 0$ is some constant. Suppose that $S = \sum_i \alpha_i X_i$, where $\alpha_i \geq 0$, $\sum_i \alpha_i = 1$. Then, for every $\delta > 0$ we have

$$\Pr[S \leq O_{C,\delta}(H(\alpha))] \geq 1 - \delta,$$

where $H(\alpha)$ is the entropy of the distribution over $[n]$ defined by α . In particular, $H(\alpha) \leq \log_2 n$.

Proof. Let T_1, T_2, \dots, T_n be non-negative parameters to be chosen later. Denote \mathcal{E} the event ‘‘for every i one has $X_i \leq T_i$ ’’. Then, by the union bound,

$$\Pr[\neg \mathcal{E}] \leq \sum_{i=1}^n \frac{C}{T_i},$$

and for every i one has $E[X_i | \mathcal{E}] \leq O_C(\log T_i)$. Thus, by Markov inequality,

$$\Pr\left[S \leq O_{C,\delta}\left(\sum_i \alpha_i \log T_i\right) \mid \mathcal{E}\right] \geq 1 - \delta/2.$$

Thus, we are looking for T_i 's such that $\sum_{i=1}^n \frac{C}{T_i} \leq \delta/2$ and $\sum_i \alpha_i \log T_i$ is minimized. Via simple calculus, we obtain the desired inequality.

Let us remind, how the embedding from [IT03] of $\text{EMD}_{[\Delta]^2}$ into ℓ_1 works. For the sake of exposition, let us assume that $\Delta = 2^l$ for a non-negative integer l .

For $s = (s_1, s_2) \in \mathbb{Z}^2$ and $0 \leq t \leq l$ we define a linear map $G_{s,t} : \mathbb{R}^{[\Delta]^2} \rightarrow \ell_1$ as follows. We first impose a grid $\mathcal{G}_{s,t}$ over \mathbb{Z}^2 with side length 2^t so that one of the corners is located in $s = (s_1, s_2)$. Then, for a measure $\mu \in \mathbb{R}^{[\Delta]^2}$ we define $G_{s,t}\mu \in \ell_1$ as follows: for every square of the grid we count the total mass of μ that

is located there. Then, we define the following (linear) embedding G_s of $\mathbb{R}^{[\Delta]^2}$ into ℓ_1 parametrized by a shift $s = (s_1, s_2) \in \mathbb{Z}^2$: $G_s\mu := \bigoplus_{t=0}^l 2^t \cdot G_{s,t}\mu$.

In [IT03] the following properties of G_s have been proved.

THEOREM 6.1. ([IT03]) *For every $\mu \in \text{EMD}_{[\Delta]^2}$:*

- for every $s = (s_1, s_2) \in \mathbb{Z}^2$, one has $\|\mu\|_{\text{EMD}} \leq O(1) \cdot \|G_s\mu\|_1$;
- $E_{s \in [\Delta]^2} [\|G_s\mu\|_1] \leq O(\log \Delta) \cdot \|\mu\|_{\text{EMD}}$.

Then, concatenating G_s for all $s \in [\Delta]^2$ one obtains a deterministic embedding of $\text{EMD}_{[\Delta]^2}$ into ℓ_1 with distortion $O(\log \Delta)$.

Now we turn to the refined analysis of the above embedding.

DEFINITION 6.1. *For $x \in \mathbb{R}^2$ and $0 < R \leq 2\Delta$ consider an ℓ_1 -ball in the plane $B_{\ell_1^2}(x, R)$. Suppose that we sample a shift $s = (s_1, s_2) \in [\Delta]^2$ uniformly at random. Consider the following random variable $\mathcal{A}_{x,R}(s)$:*

$$\mathcal{A}_{x,R}(s) := \frac{\min\left(\left\{2^t : \mathcal{G}_{s,t} \text{ doesn't cut } B_{\ell_1^2}(x, R)\right\}\right)}{R}.$$

In words, we are looking for the side length of the finest out of $l+1$ grids that does not cut the ball of interest, or 2^{l+1} , if it does not exist.

Implicit in [IT03] are the following two Lemmas.

LEMMA 6.2. ([IT03]) *There exists $C > 0$ such that for every $x \in \mathbb{R}^2$, $0 < R \leq 2\Delta$ and $t > 0$ one has*

$$\Pr_{s \in [\Delta]^2} [\mathcal{A}_{x,R}(s) \geq t] \leq \frac{C}{t}.$$

Proof. One has $\mathcal{A}_{x,R}(s) \geq t$ iff the coarsest grid with side length less than $R \cdot t$ (which is $\Theta(R \cdot t)$) cuts the ball $B_{\ell_1}(x, R)$. It can be easily verified that this probability is $O(1/t)$.

LEMMA 6.3. ([IT03]) *For every two points $x, y \in [\Delta]^2$ and uniformly random $s = (s_1, s_2) \in [\Delta]^2$ the quantity $\|G_s(e_x - e_y)\|_1$, where e_x and e_y are basis vectors that correspond to points x and y , respectively, is upper bounded by $O(1) \cdot R \cdot \mathcal{A}_{u,R}(s)$ for every $u \in \mathbb{R}^2$ and $0 < R \leq 2\Delta$ such that the ball $B_{\ell_1}(u, R)$ contains both x and y .*

Proof. All grids that are of side length at least $R \cdot \mathcal{A}_{u,R}(s)$ do not contribute to $\|G_s(e_x - e_y)\|_1$ by the definition of $\mathcal{A}_{u,R}$. All finer grids contribute towards $\|G_s(e_x - e_y)\|_1$ the geometric series, whose total sum can be upper bounded by $O(1) \cdot R \cdot \mathcal{A}_{u,R}(s)$.

Combining Lemma 6.2, Lemma 6.3 and the triangle inequality, we obtain the following Claim, which later will be very useful for our refined analysis of the embedding from [IT03]. Basically, we show that we can upper bound $\|G_s \mu\|_1$ for $\mu \in \text{EMD}_{[\Delta]^2}$ using Claim 6.1.

CLAIM 6.2. *Suppose that μ and ν are two non-negative measures over $[\Delta]^2$ that both sum to one. Assume that the optimal transportation of μ to ν consists of moving mass w_i from the point $x_i \in [\Delta]^2$ to the point $y_i \in [\Delta]^2$ for $1 \leq i \leq p$. Let $\{B_j = B_{\ell_1}(u_j, R_j)\}_{j=1}^q$ be a collection of ℓ_1 -balls in the plane such that for every $1 \leq i \leq p$ there exists $1 \leq j^*(i) \leq q$ such that both x_i and y_i belong to B_{j^*} . For every $1 \leq j \leq q$ define*

$$\tilde{w}_j = \sum_{i: j^*(i)=j} w_i.$$

Suppose we sample a shift $s = (s_1, s_2) \in [\Delta]^2$ uniformly at random. Then, the random variable

$$\|G_s(\mu - \nu)\|_1 \leq \sum_{i=1}^p w_i \|G_s(e_{x_i} - e_{y_i})\|_1$$

is dominated by $S = \sum_{j=1}^q \tilde{w}_j R_j \cdot X_i$ for some non-negative (not necessarily independent) random variables X_1, X_2, \dots, X_q such that for every i and $t > 0$ one has

$$\Pr[X_i \geq t] \leq \frac{C}{t}$$

for some absolute constant $C > 0$.

Now applying Claim 6.1 we conclude the following.

CLAIM 6.3. *Assuming the notation and conditions from Claim 6.2, we have*

$$\Pr_s [\|G_s(\mu - \nu)\|_1 \leq O(1) \cdot H(\alpha) \cdot T] \geq 0.99,$$

where $T = \sum_{j=1}^q \tilde{w}_j R_j = \sum_{i=1}^p w_i R_{j^(i)}$ and α is the following distribution over $[q]$:*

$$\alpha_j = \frac{\tilde{w}_j R_j}{T}.$$

Now we state two applications of this claim that are our main goal.

LEMMA 6.4. *Suppose that μ and ν are two non-negative measures over $[\Delta]^2$ that both sum to one and, in addition, μ has support of size at most k for some $1 \leq k \leq \Delta^2$. Then,*

$$\Pr_s [\|G_s(\mu - \nu)\|_1 \leq O(\log k + \log \log \Delta) \cdot \|\mu - \nu\|_{\text{EMD}}] \geq 0.99.$$

Proof. Suppose that $\{x_1, x_2, \dots, x_k\} \subseteq [\Delta]^2$ is the support of μ . Consider the following family of $O(k \log \Delta)$ balls: $\{B(x_i, 2^j)\}_{1 \leq i \leq k, 0 \leq j \leq \log \Delta + 1}$. Next, consider the optimal transportation of μ to ν . Every edge of length l participating in this transportation can be enclosed in one of the balls of radius $O(l)$. Thus, we can apply Claim 6.3 with $T \leq O(1) \cdot \|\mu - \nu\|_{\text{EMD}}$. It is left to upper bound $H(\alpha)$. In this Lemma we use a crude bound: namely, that $H(\alpha) \leq \log O(k \log \Delta) \leq O(\log k + \log \log \Delta)$, since the support of α is of size at most $O(k \log \Delta)$.

LEMMA 6.5. *Suppose that μ and ν are two non-negative measures over $[\Delta]^2$ that both sum to one, and all the weights of μ and ν are multiples of $1/N$, where $N \geq 1$ is some integer. Moreover, assume that μ is k -sparse for some $1 \leq k \leq N$. Then,*

$$\Pr_s [\|G_s(\mu - \nu)\|_1 \leq O(\log k + \log \log N) \cdot \|\mu - \nu\|_{\text{EMD}}] \geq 0.99.$$

Proof. The proof is the same as in Lemma 6.4, but we need to upper bound $H(\alpha)$ in a slightly fancier way. Let us recall the definition of α . For each of the $O(k \log \Delta)$ balls we compute the total mass transported over edges that are allocated to this ball and multiply it by the radius of the ball. Since all the masses are multiples of $1/N$ and for every $j \leq \log \Delta$ we have k balls of radius 2^j , we can reformulate the question of upper bounding $H(\alpha)$ as follows. Suppose that we have a bin for every $i \in [k]$ and $j \geq 0$. Then, we put N balls into these bins (adversarially). Then, for each bin indexed by (i, j) we multiply the number of balls there by 2^j and then normalize the resulting numbers so that they sum to 1. What is the upper bound of the entropy of this distribution? We prove that it is $O(\log k + \log \log N)$ as follows. Denote j^* the largest j such that there is $i \in [k]$ such that the bin (i, j) is non-empty. Then, the bins with $j \leq j^* - 100 \log N$ contribute to the entropy negligibly, since we multiply the number of balls in these bins by $2^j \leq 2^{j^*} / N^{100}$. But the entropy for bins with $j \geq j^* - 100 \log N$ is $\log O(k \log N) = O(\log k + \log \log N)$, since the total number of these ‘‘important’’ bins is $O(k \log N)$.

Remark: The terms $O(\log \log \Delta)$ and $O(\log \log N)$ in Lemma 6.4 and Lemma 6.5, might appear to be unfortunate artifacts of our analyses. However, one can show that in both cases the bounds for the embedding from [IT03] are *in fact tight*. Nevertheless, in the next section we show how to achieve approximation $O(\log k)$, if we allow embeddings into more complex spaces (that still allow reasonably good sketches).

6.2 Embedding of EMD into the ℓ_1 -sum of the small EMD instances In this section we provide a refined analysis of the embedding of $\text{EMD}_{[\Delta]^2}$ from [Ind07].

Suppose our goal is to sketch $\text{EMD}_{[\Delta]^2}$, where $\Delta = 2^l$ for some integer $l \geq 0$. Let $0 \leq t \leq l$ be a parameter to be chosen later. Let us impose a randomly shifted hierarchy of nested grids with side lengths $\Delta/2^t, \Delta/2^{2t}, \dots, 1$ ($O(\frac{\log \Delta}{t})$ grids in total). By “randomly shifted” we mean that the coarsest grid has a corner in a point $s = (s_1, s_2) \in [\Delta]^2$ chosen uniformly at random, and all the finer grids are imposed by subdividing the cruder ones. Now let us define the sketching procedure. First, we sketch $\text{EMD}_{[\Delta/2^t]^2}$ instances induced by the crudest grid recursively (we have $O(2^{2t})$ of these). Second, for each of these instances we remember the total mass. Now, to estimate EMD, we estimate EMD for the smaller instances, add these estimates, then compute EMD for the instance induced by the total masses we remembered, multiply it by $\Delta/2^t$ (the side length of the crudest grid), and add it to the result. This can be seen as a randomized embedding $f_s: \text{EMD}_{[\Delta]^2} \rightarrow \ell_1(\text{EMD}_{[O(2^t)]^2})$. In [Ind07] the following properties of f_s are shown:

THEOREM 6.2. [Ind07] *For every $\mu \in \text{EMD}_{[\Delta]^2}$:*

- for every s , one has $\|\mu\|_{\text{EMD}_{[\Delta]^2}} \leq O(1) \cdot \|f_s \mu\|_{\ell_1(\text{EMD}_{[O(2^t)]^2})}$;
- $\mathbb{E}_s \left[\|f_s \mu\|_{\ell_1(\text{EMD}_{[O(2^t)]^2})} \right] \leq O\left(\frac{\log \Delta}{t}\right) \cdot \|\mu\|_{\text{EMD}_{[\Delta]^2}}$.

In what follows we improve upon the second item in the above theorem under the following additional assumptions on μ . Namely, suppose we apply f_s for random s to a difference $\nu - \tau$, where ν and τ are non-negative measures over $[\Delta]^2$ that sum to one and ν is k -sparse.

LEMMA 6.6. *If ν and τ as above, then*

$$\Pr_s \left[\|f_s(\nu - \tau)\|_{\ell_1(\text{EMD}_{[O(2^t)]^2})} \leq O\left(1 + \frac{\log k + \log \log \Delta}{t}\right) \cdot \|\nu - \tau\|_{\text{EMD}_{[\Delta]^2}} \right] \geq 0.99.$$

Proof. As in the proof of Lemma 6.4, we cover the edges of optimal transportation of ν to τ with $O(k \log \Delta)$ balls $\{B_j\}$ such that every edge of length r lies within a ball of radius $O(r)$. Define the event \mathcal{E} as follows: “every ball B_j is not cut by a grid with side length at least radius of B_j times $Ck \log \Delta$ ”. We can choose C such that $\Pr_s[\mathcal{E}] \geq 0.999$ (we can take the union bound over

the balls B_j and for every fixed ball we proceed as in Claim 6.2).

Now let us consider a fixed edge of length r from the optimal transportation. The goal is to argue that, conditioned on \mathcal{E} , the expected contribution of the edge to $\|f_s(\nu - \tau)\|_{\ell_1(\text{EMD}_{[O(2^t)]^2})}$ is

$$O\left(r \cdot \left(\frac{\log k + \log \log \Delta}{t} + 1\right)\right).$$

Then we will be done by the triangle inequality, Markov’s inequality and the fact that $\Pr_s[\mathcal{E}] \geq 0.999$.

Let us argue about the contribution of the edge for every grid separately. First, all grids with side length less than $r/10$ contribute at most $O(r)$ in total, because the endpoints end up in different subproblems, and thus the contribution is proportional to the side length. The side lengths accumulate as geometric series, so we have that the sum is $O(r)$ in total.

Grids with side length at least $C' \cdot r \cdot k \log \Delta$ (with C' being large enough) do not contribute anything, conditioned on \mathcal{E} .

Grids with side lengths between $r/10$ and $C' \cdot r \cdot k \log \Delta$ contribute in expectation $O(r)$ each (see Lemma 3.3 in [Ind07]). Conditioning on \mathcal{E} can change the expectation by at most a constant factor, since $\Pr_s[\mathcal{E}] \geq 0.999$. Since we have $O((\log k + \log \log \Delta)/t)$ such grids, the required bound follows.

6.3 Implications for sketching of EMD

THEOREM 6.3. *One can sketch linearly $\text{EMD}_{[\Delta]^2}$ for measures that are differences of two non-negative measures that sum to 1, one of which is k -sparse as follows:*

- with sketch size $O(1)$ and approximation $O(\log k + \log \log \Delta)$;
- with sketch size $O(\log^\delta \Delta)$ and approximation $O(\log k)$ for every constant $0 < \delta < 1$.
- Moreover, if both measures have all the weights being multiples of $1/n$, where N is a positive integer, then the first of the results can be improved to having approximation $O(\log k + \log \log N)$.

Proof. The first result follows from composing the first item of Theorem 6.1 and Lemma 6.4 with a sketch for ℓ_1 from [Ind06]. The third result is similar, except we use Lemma 6.5.

As for the second result, the starting point is the first item of Theorem 6.2 together with our Lemma 6.6. Let us set $t = \delta \log \log \Delta$. This way, we get a randomized embedding of $\text{EMD}_{[\Delta]^2}$ into $\ell_1(\text{EMD}_{[O(\log^\delta \Delta)]^2})$ with distortion $O(\log k)$. Then, we apply the result of Verbin

and Zhang [VZ12] to perform dimension reduction. Namely, we need to apply their randomized map twice to reduce the dimension to $O(\log \log \Delta)$. As a result, we get a sketch of size $O(\log^{O(\delta)} \Delta)$ and distortion $O(\log k)$, if δ is a (small) positive constant.

THEOREM 6.4. *One can sketch linearly $\text{EMD}_{[\Delta]}$ over interval $[\Delta]$ of measures that are differences of two non-negative measures that sum to 1, one of which is k -sparse. We can achieve sketch size $O(1/\epsilon^2)$ and approximation $1 + \epsilon$.*

Proof. Using Lemma 6.1, we can isometrically embed EMD over the interval $[\Delta]$ into ℓ_1 . Now we can sketch ℓ_1 using the sketch from [Ind06]. This give sketch size $O(1/\epsilon^2)$ and approximation $1 + \epsilon$.

7 Sparse recovery for EMD : putting it all together

The following three theorems follow from Lemma 3.1 and Theorem 6.3.

THEOREM 7.1. *There is a linear sketching scheme of probability distributions over $[\Delta]^2$ with the following guarantees. The size of the sketch is*

$$O(k(\log \log \Delta) \log(\log k + \log \log \Delta) + \log \log(\Delta/\lambda))$$

and, given a sketch of x , we can recover x^* such that

$$\text{EMD}(x, x^*) \leq \max(O(\log k + \log \log \Delta) \min_{k\text{-sparse } x'} \text{EMD}(x, x'), \lambda).$$

in time polynomial in Δ and $\log^{O(k)} \Delta$.

Proof. Lemma 5.1 gives that the doubling dimension of k -sparse probability measures over $[\Delta]^2$ is $O(k \log \log \Delta)$. Combining this with Lemma 3.1 and the first result from Theorem 6.3, we get the stated guarantees.

THEOREM 7.2. *There is a linear sketching scheme of probability distributions over $[\Delta]^2$ with the following guarantees. The size of the sketch is*

$$O(1)(\log^\delta \Delta)(k(\log \log \Delta) \log \log k + \log \log(\Delta/\lambda))$$

for some constant $\delta > 0$. Given a sketch of x , we can recover x^* such that

$$\text{EMD}(x, x^*) \leq \max(O(\log k) \min_{k\text{-sparse } x'} \text{EMD}(x, x'), \lambda).$$

in time polynomial in Δ and $\log^{O(k)} \Delta$.

Proof. Lemma 5.1 gives that the doubling dimension of k -sparse probability measures over $[\Delta]^2$ is $O(k \log \log \Delta)$. Combining this with Lemma 3.1 and the second result from Theorem 6.3, we get the stated guarantees.

THEOREM 7.3. *Let N be a positive integer. There is a linear sketching scheme of probability measures that have granularity $1/N$. The size of the sketch is*

$$O(k(\log \log N) \log(\log k + \log \log N) + \log \log(\Lambda/\lambda))$$

and, given a sketch of x , we can recover x^* such that

$$\text{EMD}(x, x^*) \leq \max(O(\log k + \log \log N) \min_{k\text{-sparse } x'} \text{EMD}(x, x'), \lambda).$$

in time polynomial in Δ and $\log^{O(k)} \Delta$. Λ is the upper bound on $\text{EMD}(x, y)$ for the starting k -sparse approximation y of x .

Proof. Lemma 5.2 gives that the doubling dimension of k -sparse probability measures with granularity $1/n$ is $O(k \log \log N)$. Combining this with Lemma 3.1 and the third result from Theorem 6.3, we get the stated guarantees.

THEOREM 7.4. *There is a linear sketching scheme of probability distributions over interval $[\Delta]$ with the following guarantees. The size of the sketch is*

$$O(1/\epsilon^2)(k(\log \log \Delta) \log \frac{1}{\epsilon} + \log \log(\Delta/\lambda))$$

and, given a sketch of x , we can recover x^* such that

$$\text{EMD}(x, x^*) \leq \max((1 + \epsilon) \min_{k\text{-sparse } x'} \text{EMD}(x, x'), \lambda).$$

in time polynomial in Δ and $\log^{O(k)} \Delta$.

Proof. Lemma 5.1 also gives that the doubling dimension of k -sparse probability measures over interval $[\Delta]$ is $O(k \log \log \Delta)$. Combining this with Lemma 3.1 and Theorem 6.4, we get the stated guarantees.

7.1 Lower Bounds Lemma 5.3 and Lemma 4.1 gives the following two theorems 7.5 and 7.6.

THEOREM 7.5. *Any linear sparse recovery scheme with approximation factor K with respect to EMD over interval $[\Delta]$ requires*

$$m \geq \frac{\Omega(k) \log(\Omega(\log \frac{\Delta}{k}))}{\log K}$$

measurements for sparsity $k > 1$.

We want to compare guarantees of Theorem 7.4 with the lower bound that we achieve in Theorem 7.5.

Theorem 7.4 and assumptions that ϵ is a constant and $\lambda \geq 2^{-(\log \Delta)^{O(k)}}$ gives approximation guarantee

$$(7.12) \quad \text{EMD}(x, x^*) \leq \max(O(1) \min_{k\text{-sparse } x'} \text{EMD}(x, x'), \lambda)$$

with $O(k \log \log \Delta)$ number of measurements.

Theorem 7.5 and assumption that $k < \Delta^{1-c}$ for some constant $c > 0$ give lower bound $\Omega(k \log \log \Delta)$ on the number of measurements for constant approximation factor. However, this lower bound holds for the case when λ is equal to 0 in guarantee 7.12.

From the proof of Lemma 4.1 (equations (4.8), (4.9) and (4.10)) and Lemma 5.3 (we construct $\frac{k}{200}$ -cover for EMD ball of radius $k/2$) we see that we are actually good as long as λ is sufficiently small. As long as $\lambda \leq \frac{k}{C}$ for some large constant C . Therefore, our lower bound holds if $\frac{k}{C} \geq \lambda \geq 2^{-(\log \Delta)^{O(k)}}$.

We see that the upper bound and the lower bound match for the described range of parameters.

THEOREM 7.6. *Any linear sparse recovery scheme with approximation factor K with respect to EMD over square $[\Delta]^2$ requires*

$$m \geq \frac{\Omega(k) \log \left(\Omega \left(\log \frac{\Delta^2}{k} \right) \right)}{\log K}$$

measurements for sparsity $k > 1$.

We want to compare guarantees of Theorem 7.1 with the lower bound that we achieve in Theorem 7.6.

Theorem 7.1 and assumptions that ϵ is a constant and $\lambda \geq 2^{-(\log \Delta)^{O(k)}}$ gives approximation guarantee

$$(7.13) \quad \text{EMD}(x, x^*) \leq \max(O(\log k + \log \log \Delta) \min_{k\text{-sparse } x'} \text{EMD}(x, x'), \lambda)$$

with $O(k(\log \log \Delta) \log(\log k + \log \log \Delta))$ number of measurements.

Theorem 7.6 and assumption that $k < \Delta^{2-c}$ for some constant $c > 0$ give lower bound $\Omega(1) \frac{k \log \log \Delta}{\log(\log k + \log \log \Delta)}$ on the number of measurements for approximation factor $O(\log k + \log \log \Delta)$. However, this lower bound holds for the case when λ is equal to 0 in guarantee 7.13.

From the proof of Lemma 4.1 (equations (4.8), (4.9) and (4.10)) and Corollary 5.1 (we construct $\frac{k}{200}$ -cover for EMD ball of radius $k/2$) we see that we are actually

good as long as λ is sufficiently small. As long as $\lambda \leq \frac{k}{C}$ for some large constant C . Therefore, our lower bound holds if $\frac{k}{C} \geq \lambda \geq 2^{-(\log \Delta)^{O(k)}}$.

We see that the upper bound and the lower bound match up to a factor of $\log^2(\log k + \log \log \Delta)$ for the described range of parameters.

8 Sketching of 1-Median

For a vector $x \in \mathbb{R}^n$, we use $\|x\|_{\text{med}}$ to denote the median over $i \in [n]$ of $|x_i|$.

8.1 Subspace embeddings

LEMMA 8.1. *Let L be a d -dimensional subspace of \mathbb{R}^n . Let $A \in \mathbb{R}^{m \times n}$ be a matrix with $m = O(\frac{1}{\epsilon^2} d \log \frac{d}{\epsilon \delta})$ and i.i.d. Cauchy entries with scale parameter $\gamma = 1$. With $1 - \delta$ probability, for all $x \in L$ we have*

$$(1 - \epsilon)\|x\|_1 \leq \|Ax\|_{\text{med}} \leq (1 + \epsilon)\|x\|_1.$$

Proof. In an abuse of notation, let L be an orthonormal basis for the subspace L . For any threshold $\tau = \text{poly}(\frac{d}{\epsilon \delta})$, the probability that any entry of AL has absolute value larger than τ is $O(\sqrt{d}/\tau)$, using that the ℓ_1 norms of the columns of L is \sqrt{d} . Setting $\tau = O(d^{2.5}/\delta)$, we have that every entry of AL is at most τ with probability $1 - \delta/2$. Suppose this happens.

Then for all $x \in \mathbb{R}^d$, we have that $\|Lx\|_1 \geq \|Lx\|_2 = \|x\|_2 \geq \|x\|_1/\sqrt{d}$ and $\|ALx\|_\infty \leq \tau\|x\|_1 \leq \tau\sqrt{d}\|Lx\|_1$. Thus for all $y \in L$ we have

$$\|Ay\|_\infty \leq (d^3/\delta)\|y\|_1.$$

Let $\tau' = d^3/\delta$.

We construct an $\frac{\epsilon}{\tau'}$ -net T in the ℓ_1 norm for the unit ℓ_1 ball intersect L , which has size at most $(1 + \tau'/\epsilon)^d = e^{O(d \log \frac{d}{\epsilon \delta})}$ by the standard volume argument.

For any $x \in \mathbb{R}^n$, we say Ax is “good” if only a $\frac{1}{2} - C_2\epsilon$ fraction of coordinates are too large or too small, i.e.

$$\begin{aligned} |\{i : |(Ax)_i| < (1 - \epsilon)\|x\|_1\}| &\leq \left(\frac{1}{2} - C_2\epsilon\right)m \\ |\{i : |(Ax)_i| > (1 + \epsilon)\|x\|_1\}| &\leq \left(\frac{1}{2} - C_2\epsilon\right)m \end{aligned}$$

for some small constant C_2 . If Ax is “good”, then for any y with at most $C_2\epsilon m$ coordinates larger than $\epsilon\|x\|_1$, we have

$$(8.14) \quad (1 - 2\epsilon)\|x\|_1 \leq \|Ax + y\|_{\text{med}} \leq (1 + 2\epsilon)\|x\|_1.$$

Because $(Ax)_i$ is a Cauchy variable with scale $\|x\|_1$, we have that

$$\begin{aligned} \Pr[|(Ax)_i| < (1 - \epsilon)\|x\|_1] &< 1/2 - \Omega(\epsilon) \\ \Pr[|(Ax)_i| > (1 + \epsilon)\|x\|_1] &< 1/2 - \Omega(\epsilon). \end{aligned}$$

By a Chernoff bound, for sufficiently small C_2 we have that Ax is “good” with all but $e^{-\Omega(\epsilon^2 m)}$ probability. For our choice of m , we can union bound to have that Ax is “good” for all $x \in T$ with all but $e^{-\Omega(\epsilon^2 m)} \leq \delta^{\Omega(d)}$ probability.

Every $y \in L$ with $\|y\|_1 = 1$ can be expressed as $x+z$ for $x \in T$ and $\|z\|_1 \leq \epsilon/\tau'$. We have that Ax is “good” and that $\|\infty\|Az \leq \tau'\|z\|_1 \leq \epsilon$. Hence by (8.14),

$$(1 - 2\epsilon)\|x\|_1 \leq \|Ay\|_{med} \leq (1 + 2\epsilon)\|x\|_1.$$

which implies

$$(1 - 3\epsilon)\|y\|_1 \leq \|Ay\|_{med} \leq (1 + 3\epsilon)\|y\|_1.$$

Since A is linear, the restriction to $\|y\|_1 = 1$ is unnecessary; rescaling ϵ then gives the result.

COROLLARY 8.1. *Let A have $O(d \log(d/(\epsilon\delta))/\epsilon^2)$ rows and Cauchy entries with scale $\gamma = 1$. For any subspace L of dimension d and subset $S \subset L$, with $1-\delta$ probability we have that*

$$\hat{x} := \operatorname{argmin}_{x \in S} \|Ax\|_{med}$$

satisfies

$$\|\hat{x}\|_1 \leq (1 + \epsilon) \min_{x \in S} \|x\|_1.$$

8.2 1-median in d dimensions

8.2.1 1-median in 1 dimension

THEOREM 8.1. *We can find a $1 + \epsilon$ -approximation to the 1-median in 1 dimensions using $O(\log(1/\epsilon)/\epsilon^2)$ linear measurements and $\exp(\text{poly}(1/\epsilon))$ time.*

Proof. Define $B \in \mathbb{R}^{n \times 2}$ by $B_{i,1} = i$ and $B_{i,2} = 1$ for all $i \in [n]$. For any $x \in \mathbb{R}^n$, define $D_x \in \mathbb{R}^{n \times n}$ to be the diagonal matrix with $D_{i,i} = x_i$. Then for any $j \in [n]$ and $z = (j, -1)$, we have that

$$\|D_x Bz\|_1 = \sum_i |x_i| |i - j|$$

is the cost of using j as the median for x .

Let $A \in \mathbb{R}^{m \times n}$ for $m = O(\log(1/\epsilon)/\epsilon^2)$ have i.i.d. Cauchy entries. Then $AD_x B \in \mathbb{R}^{m \times 2}$ consists of $2m$ linear measurements of x .

Furthermore, the set of $S = \{D_x Bz \mid z_2 = -1\}$ is a subset of a 2-dimensional subspace. Hence, by Corollary 8.1,

$$\hat{z} = \operatorname{argmin}_{\substack{z \in \mathbb{R}^2 \\ z_2 = -1}} \|AD_x Bz\|_{med}$$

satisfies

$$\|D_x B\hat{z}\|_1 \leq (1 + \epsilon) \min_{\substack{z \in \mathbb{R}^2 \\ z_2 = -1}} \|D_x Bz\|_1 = (1 + \epsilon) \operatorname{cost}(x).$$

Given $AD_x B$ we can compute \hat{z} , from which we recover z_1 as a $(1 + \epsilon)$ approximation to the 1-median.

8.2.2 1-median in d dimensions

CLAIM 8.1. (DVORETSKY'S THEOREM [DVO60]) *Let $G \in \mathbb{R}^{m \times d}$ have suitably scaled i.i.d. Gaussian entries, for $m = O(d/\epsilon^2)$. Then with all but $e^{-\Omega(d)}$ probability, for all $x \in \mathbb{R}^d$ we have*

$$\|Gx\|_1 \leq \|x\|_2 \leq (1 + \epsilon)\|Gx\|_1.$$

THEOREM 8.2. *We can find a $1 + \epsilon$ -approximation to the Euclidean 1-median in d dimensions using $O(d^2 \log(d/\epsilon)/\epsilon^2)$ linear measurements and $\exp(\text{poly}(d/\epsilon))$ time.*

Proof. Let $G \in \mathbb{R}^{t \times d}$ for $t = O(d/\epsilon^2)$ satisfy Claim 8.1, so

$$\|Gp\|_1 \leq \|p\|_2 \leq (1 + \epsilon)\|Gp\|_1$$

for all $p \in [n]^d$. For each point $p \in [n]^d$, define the matrix $B^{(p)} \in \mathbb{R}^{t \times (t+1)}$ by the first t columns being the identity matrix and column $t+1$ being Gp .

Define $G' \in \mathbb{R}^{(t+1) \times (d+1)}$ to equal G over the first $t \times d$ submatrix, $G'_{t+1,d+1} = 1$, and zero elsewhere. For any point $p \in [n]^d$ define $z^{(p)} \in \mathbb{R}^{d+1}$ by $z_i = p_i$ for $i \leq d$ and $z_{d+1} = -1$. For any $p, q \in [n]^d$, we have

$$B^{(q)} G' z_p = \begin{pmatrix} I & Gq \end{pmatrix} \begin{pmatrix} G & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} p \\ -1 \end{pmatrix} = Gq - Gp$$

Hence

$$\|B^{(q)} G' z_p\|_1 = \|Gq - Gp\|_1 \leq \|p - q\|_2 \leq (1 + \epsilon) \|B^{(q)} G' z_p\|_1.$$

For $x \in \mathbb{R}^d$, define $C_x \in \mathbb{R}^{tn \times (t+1)}$ to be the concatenation of the matrices $x_p B^{(p)}$ for all $p \in [n]^d$. Then for all $x \in \mathbb{R}^d$ and $p \in [n]^d$, therefore,

$$(8.15) \quad \|C_x G' z^{(p)}\|_1 \leq \operatorname{cost}(x, p) \leq (1 + \epsilon) \|C_x G' z^{(p)}\|_1.$$

Let $A \in \mathbb{R}^{m \times tn}$ for $m = O(d \log(d/\epsilon)/\epsilon^2)$ have i.i.d. Cauchy entries. Our method observes

$$AC_x G' \in \mathbb{R}^{m \times (d+1)}$$

which is a set of $m(d+1) = O(d^2 \log(d/\epsilon)/\epsilon^2)$ linear measurements of x .

By Corollary 8.1, with good probability we have that

$$\hat{z} = \operatorname{argmin}_{\substack{z \in \mathbb{R}^{d+1} \\ z_{d+1} = -1}} \|AC_x G' z\|_{med}$$

satisfies

$$\|C_x G' \hat{z}\|_1 \leq (1 + \epsilon) \min_{\substack{z \in \mathbb{R}^{d+1} \\ z_{d+1} = -1}} \|C_x G' z\|_1 = (1 + \epsilon) \operatorname{cost}(x).$$

Hence by (8.15), for $\hat{p} = (\hat{z}_1, \dots, \hat{z}_d)$,

$$\operatorname{cost}(x, \hat{p}) \leq (1 + \epsilon)^2 \operatorname{cost}(x).$$

Given $AC_x G'$ we can compute \hat{z} , from which we get \hat{p} as a $(1 + \epsilon)$ approximation to the 1-median.

9 Lower bounds for k -means

In this section we prove lower bounds for sketching and streaming k -means.

First, one can extend the definition of EMD to the sum of squares of distances. Let us denote the corresponding “distance” EMD^2 . It is immediate to see that $\mathbb{R}^{[\Delta]^2}$ equipped with EMD^2 is a 2-quasi-metric space. Sparse recovery with respect to EMD^2 is equivalent to the k -means clustering.

Second, observe that the construction from Section 5.2 can be translated verbatim to EMD^2 to show that the doubling dimension of the latter is $\Omega(k \cdot \log \log \frac{\Delta^2}{k})$ as well.

Finally, observe that the results of Section 4 can be applied to EMD^2 as well. Indeed, EMD^2 enjoys the polynomial aspect ratio and relaxed triangle inequality, and these two happen to be enough for the argument to go through.

As a result, we get the lower bound $\Omega(k \cdot \log \log \frac{\Delta^2}{k} / \log K)$ on the number of measurements necessary for the linear sketching of k -means with approximation K .

Alternatively, we can consider the streaming model and reuse the proof from Section 4 to show that *streaming* k -means with approximation K requires

$$\Omega\left(\frac{k}{\log K} \cdot \log \log \frac{\Delta^2}{k} \cdot \log \Delta\right)$$

bits.

10 Acknowledgments

We would like to thank Eric Price for his significant contributions to this paper. We thank Alexandr Andoni for helpful discussions.

References

- [ABS10] Marcel R. Ackermann, Johannes Blömer, and Christian Sohler. Clustering for metric and nonmetric distance measures. *ACM Transactions on Algorithms*, 6(4), 2010.
- [AZGR15] Zeyuan Allen-Zhu, Rati Gelashvili, and Ilya Razenshteyn. Restricted isometry property for general p -norms. In *Proceedings of 31st International Symposium on Computational Geometry (SoCG 2015)*, 2015.
- [BC05] Bo Brinkman and Moses Charikar. On the impossibility of dimension reduction in l_1 . *J. ACM*, 52(5):766–788, 2005.
- [BGI⁺08] Radu Berinde, Anna C. Gilbert, Piotr Indyk, Howard J. Karloff, and Martin J. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. *Allerton*, 2008.
- [BI14] Artūrs Bačkurs and Piotr Indyk. Better embeddings for planar earth-mover distance over sparse sets. In *Annual Symposium on Computational Geometry*, page 280. ACM, 2014.
- [CDM⁺13] Kenneth L. Clarkson, Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, Xiangrui Meng, and David P. Woodruff. The fast cauchy transform and faster robust linear regression. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 466–477, 2013.
- [Cha02] Moses Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 380–388, 2002.
- [CRT06] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1208–1223, 2006.
- [CS02] Moses Charikar and Amit Sahai. Dimension reduction in the ℓ_1 norm. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 551–560, 2002.
- [DIPW10] Khanh Do Ba, Piotr Indyk, Eric Price, and David P Woodruff. Lower bounds for sparse recovery. In *SODA*, volume 10, pages 1190–1197. SIAM, 2010.
- [Don06] D. L. Donoho. Compressed Sensing. *IEEE Trans. Info. Theory*, 52(4):1289–1306, Apr. 2006.
- [Dvo60] A. Dvoretzky. Some results on convex bodies and Banach spaces. 1960.
- [FS05] G. Frahling and C. Sohler. Coresets in dynamic geometric data streams. *STOC*, 2005.
- [GD05] K. Grauman and T. Darrell. The pyramid match kernel: Discriminative classification with sets of image features. *ICCV*, 2005.
- [GHS14] A Gilbert, P Indyk, M Iwen, and L Schmidt. Recent developments in the sparse fourier transform: A compressed fourier transform for big data. *Signal Processing Magazine, IEEE*, 31(5):91–100, 2014.
- [GIP10] Rishi Gupta, Piotr Indyk, and Eric Price. Sparse recovery for earth mover distance. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 1742–1744. IEEE, 2010.
- [GKK10] Lee-Ad Gottlieb, Leonid Kontorovich, and Robert Krauthgamer. Efficient classification for metric data. In *Proceedings of the 23rd Conference on Learning Theory (COLT 2010)*, pages 433–440, 2010.
- [Ind04] P. Indyk. Algorithms for dynamic geometric problems over data streams. *STOC*, 2004.
- [Ind06] P. Indyk. Stable distributions, pseudorandom generators, embeddings and data stream computation. *J. ACM*, 53(3):307–323, 2006.
- [Ind07] Piotr Indyk. A near linear time constant factor approximation for Euclidean bichromatic matching (cost). In *Proceedings of the 18th ACM-SIAM Sym-*

- posium on Discrete Algorithms (SODA 2007)*, pages 39–42, 2007.
- [IP11] Piotr Indyk and Eric Price. K-median clustering, model-based compressive sensing, and sparse recovery for earth mover distance. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 627–636. ACM, 2011.
- [IT03] Piotr Indyk and Nitin Thaper. Fast image retrieval via embeddings. In *the 3rd International Workshop on Statistical and Computational Theories of Vision (at ICCV 2003)*, 2003.
- [JST11] Hossein Jowhari, Mert Saglam, and Gábor Tardos. Tight bounds for lp samplers, finding duplicates in streams, and related problems. In *Proceedings of the 30th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2011, June 12-16, 2011, Athens, Greece*, pages 49–58, 2011.
- [KL04] Robert Krauthgamer and James R. Lee. Navigating nets: simple algorithms for proximity search. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA 2004)*, pages 798–807, 2004.
- [KW14] Felix Krahmer and Rachel Ward. A unified framework for linear dimensionality reduction in l_1 . *arXiv preprint arXiv:1405.1332*, 2014.
- [MD13] Dian Mo and Marco F Duarte. Compressive parameter estimation with earth mover’s distance via k-median clustering. In *SPIE Optical Engineering+ Applications*, pages 88581P–88581P. International Society for Optics and Photonics, 2013.
- [MM13] Xiangrui Meng and Michael W. Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 91–100, 2013.
- [Mut05] S Muthukrishnan. *Data streams: Algorithms and applications*. Now Publishers Inc, 2005.
- [RTG00] Y. Rubner, C. Tomassi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000.
- [Soh] Christian Sohler. Personal Communication.
- [SW11] Christian Sohler and David P. Woodruff. Subspace embeddings for the l_1 -norm with applications. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 755–764, 2011.
- [VZ12] Elad Verbin and Qin Zhang. Rademacher-sketch: A dimensionality-reducing embedding for sum-product norms, with an application to Earth-Mover Distance. In *Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP 2012)*, pages 834–845, 2012.
- [WZ13] David P. Woodruff and Qin Zhang. Subspace embeddings and l_1 -regression using exponential random variables. In *COLT 2013 - The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA*, pages 546–567, 2013.
- [YO14] Arman Yousefi and Rafail Ostrovsky. Improved approximation algorithms for earth-mover distance in data streams. *arXiv preprint arXiv:1404.6287*, 2014.