Optimal compression of approximate Euclidean distances

Noga Alon[1] Bo'az Klartag[2]

Abstract

Let X be a set of n points of norm at most 1 in the Euclidean space R^k, and suppose $\varepsilon > 0$. An ε-distance sketch for X is a data structure that, given any two points of X enables one to recover the square of the (Euclidean) distance between them up to an *additive* error of ε. Let $f(n, k, \varepsilon)$ denote the minimum possible number of bits of such a sketch. Here we determine $f(n, k, \varepsilon)$ up to a constant factor for all $n \geq k \geq 1$ and all $\varepsilon \geq \frac{1}{n^{0.49}}$. Our proof is algorithmic, and provides an efficient algorithm for computing a sketch of size $O(f(n, k, \varepsilon)/n)$ for each point, so that the square of the distance between any two points can be computed from their sketches up to an additive error of ε in time linear in the length of the sketches.

1 The problem and main results

Let X be a set of n points of norm at most 1 in the Euclidean space R^k, and suppose $\varepsilon > 0$. An ε-distance sketch for X is a data structure that, given any two points of X enables one to recover the square of the (Euclidean) distance between them up to an *additive* error of ε. What is the minimum possible number of bits of such a sketch ? Denote this minimum by $f(n, k, \varepsilon)$. Here $1 \leq k \leq n$ and we assume that $\frac{1}{n^{0.49}} \leq \varepsilon \leq 0.1$.

The most basic case is when $k = n$, that is, there is no restriction on the dimension. In this case one can apply the Johnson-Lindenstrauss Lemma [5] to project the points into R^m where $m = O(\log n/\varepsilon^2)$ with distortion at most $\varepsilon/2$, and then round each point to the closest one in an $\varepsilon/2$-net in the ball of radius $1 + \varepsilon/2$ in R^m. As the size of the net is $[O(1/\varepsilon)]^m$, this enables us to represent each point by $O(m \log(1/\varepsilon))$ bits showing that

$$f(n, n, \varepsilon) \leq O(\frac{n \log n}{\varepsilon^2} \log(1/\varepsilon)).$$

[1]Sackler School of Mathematics and Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. Email: nogaa@tau.ac.il. Research supported in part by a USA-Israeli BSF grant 2012/107, by an ISF grant 620/13 and by the Israeli I-Core program.

[2]Sackler School of Mathematics, Tel Aviv University, Tel Aviv 69978, Israel and Department of Mathematics, Weizmann Institute of Science, Rehovot 7610001, Israel. Email: klartagb@tau.ac.il. Research supported in part by an ERC grant.

On the other hand it is not difficult to deduce from the recent construction in [7] that

$$f(n,n,\varepsilon) \geq \Omega(\frac{n \log n}{\varepsilon^2}).$$

A better upper bound follows from the results of Kushilevitz, Ostrovsky and Rabani in [6], where the authors show that all inner products between the pairs of n points on the unit sphere in R^n can be approximated up to a relative error of ε by storing only $O(\frac{\log n}{\varepsilon^2})$ bits per point. This easily implies that $f(n,n,\varepsilon) = \Theta(\frac{n \log n}{\varepsilon^2})$ and in view of the discussion above

$$f(n,k,\varepsilon) = \Theta(\frac{n \log n}{\varepsilon^2})$$

for all $\frac{\log n}{\varepsilon^2} \leq k \leq n$.

What happens for smaller k ? In this short paper we determine $f(n,k,\varepsilon)$ up to a constant factor for all admissible n, k and ε. This is stated in the following Theorem.

Theorem 1.1 *For all n and $\frac{1}{n^{0.49}} \leq \varepsilon \leq 0.1$ the function $f(n,k,\varepsilon)$ satisfies the following*

- *For $\frac{\log n}{\varepsilon^2} \leq k \leq n$,*
$$f(n,k,\varepsilon) = \Theta(\frac{n \log n}{\varepsilon^2}).$$
- *For $\log n \leq k \leq \frac{\log n}{\varepsilon^2}$,*
$$f(n,k,\varepsilon) = \Theta(nk \log(2 + \frac{\log n}{\varepsilon^2 k})).$$
- *For $1 \leq k \leq \log n$,*
$$f(n,k,\varepsilon) = \Theta(nk \log(1/\varepsilon)).$$

As mentioned above, the first part of the theorem is known, by the results of [6], [7]. For completeness, and since our proof is different, we include here a proof of this part as well. We present two proofs of the upper bound in the theorem. The first, described in Section 2, is based on a short probabilistic (or volume) argument. Its disadvantage is that it is not constructive and provides neither an efficient algorithm for producing the sketch for a given set of points X, nor an efficient algorithm for recovering the approximate square distance between two desired points of X, given the sketch. The second proof, presented in Section 3, is algorithmic. It provides an efficient randomized algorithm for computing a sketch consisting of $O(f(n,k,\varepsilon)/n)$ bits for each point of X, so that the square of the distance between any two points can be recovered, up to an additive error of ε, from their sketches, in time linear in the length of the sketches.

The proofs of the upper bound here and in [6] are different. In particular, our proof(s) yield sharp results for all values of k while the argument in [6] is suboptimal for $k = o(\frac{\log n}{\varepsilon^2})$. We describe the lower bound in Section 4, and the final Section 5 contains some concluding remarks and open problems.

Theorem 1.1 supplies an alternative proof of the main result of [7] about dimension reduction. For $n \geq k \geq \ell$ and $\varepsilon \geq \frac{1}{n^{0.49}}$ we say that there is an $(n, k, \ell, \varepsilon)$-Euclidean dimension reduction if for any points $x_1, \ldots, x_n \in R^k$ of norm at most one, there exist points $y_1, \ldots, y_n \in R^\ell$ satisfying

$$|x_i - x_j|^2 - \varepsilon \leq |y_i - y_j|^2 \leq |x_i - x_j|^2 + \varepsilon \qquad (i, j = 1, \ldots, n). \tag{1}$$

Corollary 1.2 *There exists an absolute positive constant $c > 0$ so that for any $n \geq k > ck \geq \ell$ and for $1/n^{0.49} \leq \varepsilon \leq 0.1$, there is an $(n, k, \ell, \varepsilon)$-Euclidean dimension reduction if and only if $\ell = \Omega(\log n/\varepsilon^2)$.*

Moreover, the same holds if we replace additive distortion by multiplicative distortion, i.e., if we replace condition (1) by the following condition

$$(1 - \varepsilon) \cdot |x_i - x_j|^2 \leq |y_i - y_j|^2 \leq (1 + \varepsilon) \cdot |x_i - x_j|^2 \qquad (i, j = 1, \ldots, n). \tag{2}$$

Corollary 1.2 means that if $k \geq c_1 \log n/\varepsilon^2$, then there is an $(n, k, \log n/\varepsilon^2, \varepsilon)$-Euclidean dimension reduction (by the Johnson-Lindesntrauss Lemma), and that if there is an $(n, k, \ell, \varepsilon)$-Euclidean dimension reduction with $\ell = o(k)$ then necessarily $k \geq \ell \geq c_2 \log n/\varepsilon^2$, for some absolute constants $c_1, c_2 > 0$. This statement for $k \geq \Omega(\log n/\varepsilon^2)$ is proved in [7], and the result for smaller k is an easy consequence.

Throughout the proofs we make no serious attempt to optimize the absolute constants involved. For convenience we sometimes bound $f(n, k, 2\varepsilon)$ or $f(n, k, 5\varepsilon)$ instead of $f(n, k, \varepsilon)$, the corresponding bounds for $f(n, k, \varepsilon)$ follow, of course, by replacing ε by $\varepsilon/2$ or $\varepsilon/5$ in the expressions we get, changing the estimates only by a constant factor.

2 The upper bound

It is convenient to split the proof into three lemmas, dealing with the different ranges of k.

Lemma 2.1 *For $\frac{\log n}{\varepsilon^2} \leq k \leq n$,*

$$f(n, k, 5\varepsilon) \leq O(\frac{n \log n}{\varepsilon^2}).$$

Proof: Since $f(n,k,5\varepsilon)$ is clearly a monotone increasing function of k, it suffices to prove the upper bound for $k=n$. By [5] we can replace the points of $X \subset B^k$, where B^k is the unit ball in R^k, by points in R^m where $m = 40\frac{\log n}{\varepsilon^2}$ so that all distances and norms of the points change by at most ε. Hence we may and will assume that our set of points X lies in R^m. Note that given the squares of the norms of two vectors up to an additive error of ε and given their inner product up to an additive error of ε we get an approximation of the square of their distance up to an additive error of 4ε. It thus suffices to show the existence of a sketch that can provide the approximate norm of each of our vectors and the approximate inner products between pairs. The approximate norms can be stored trivially by $O(\log(1/\varepsilon))$ bits per vector. (Note that here the cost for storing even a much better approximation for the norms is negligible, so if the constants are important we can ensure that the norms are known with almost no error). It remains to prepare a sketch for the inner products.

The Gram matrix $G(w_1, w_2, \ldots, w_n)$ of n vectors $w_1, \ldots, w_n$ is the n by n matrix G given by $G(i,j) = w_i^t w_j$. We say that two Gram matrices G_1, G_2 are ε-separated if there are two indices $i \neq j$ so that $|G_1(i,j) - G_2(i,j)| > \varepsilon$. Let $\mathcal{G}$ be a maximal (with respect to containment) set of Gram matrices of ordered sequences of n vectors $w_1, \ldots, w_n$ in R^m, where the norm of each vector w_i is at most 2, so that every two distinct members of $\mathcal{G}$ are ε-separated. Note that by the maximality of $\mathcal{G}$, for every Gram matrix M of n vectors of norms at most 2 in R^m there is a member of $\mathcal{G}$ in which all inner products of pairs of distinct points are within ε of the corresponding inner products in M, meaning that as a sketch for M it suffices to store (besides the approximate norms of the vectors), the index of an appropriate member of $\mathcal{G}$. This requires $\log |\mathcal{G}|$ bits. It remains to prove an upper bound for the cardinality of $\mathcal{G}$. We proceed with that.

Let $V_1, V_2, \ldots, V_n$ be n vectors, each chosen randomly, independently and uniformly in the ball of radius 3 in R^m. Let $T = G(V_1, V_2, \ldots, V_n)$ be the Gram matrix of the vectors V_i. For each $G \in \mathcal{G}$ let A_G denote the event that for every $1 \leq i \neq j \leq n$, $|T(i,j) - G(i,j)| < \varepsilon/2$. Note that since the members of $\mathcal{G}$ are ε-separated, all the events A_G for $G \in \mathcal{G}$ are pairwise disjoint. We claim that the probability of each event A_G is at least $0.5(1/3)^{mn}$. Indeed, fix a Gram matrix $G = G(w_1, \ldots, w_n) \in \mathcal{G}$ for some $w_1, \ldots, w_n \in R^m$ of norm at most 2. For each fixed i the probability that V_i lies in the unit ball centered at w_i is exactly $(1/3)^m$. Therefore the probability that this happens for all i is exactly $(1/3)^{nm}$. The crucial observation is that conditioning on that, each vector V_i is uniformly distributed in the unit ball centered at w_i. Therefore, after the conditioning, for each $i \neq j$ the probability that the inner product $(V_i - w_i)^t w_j$ has absolute value at

least $\varepsilon/2$ is at most $2e^{-\varepsilon^2 m/8} < 1/(2n^2)$. (Here we used the fact that the norm of w_j is at most 2). Similarly, since the norm of V_i is at most 3, the probability that the inner product $V_i^t(V_j - w_j)$ has absolute value at least $\varepsilon/2$ is at most $2e^{-\varepsilon^2 m/12} < 1/2n^2$. It follows that with probability bigger than $0.5(1/3)^{nm}$ all these inner products are smaller than $\varepsilon/2$, implying that

$$|V_i^t V_j - w_i^t w_j| \leq |(V_i - w_i)^t w_j| + |V_i^t(V_j - w_j)| < \varepsilon.$$

This proves that the probability of each event A_G is at least $0.5(1/3)^{nm}$, and as these are pairwise disjoint their number is at most $2 \cdot 3^{nm}$, completing the proof of the lemma. □

Lemma 2.2

For $\log n \leq k \leq \frac{\log n}{\varepsilon^2}$,

$$f(n, k, 4\varepsilon) \leq O(nk \log(2 + \frac{\log n}{\varepsilon^2 k})).$$

Proof: The proof is nearly identical to the second part of the proof above. Note, first, that by monotonicity and the fact that the expressions above change only by a constant factor when ε changes by a constant factor, it suffices to prove the required bound for $k = \frac{\delta^2}{\varepsilon^2} \log n$ where $2\varepsilon \leq \delta \leq 1/2$. Let $\mathcal{G}$ be a maximal set of ε-separated Gram matrices of n vectors of norm at most 1 in R^k. (Here it suffices to deal with norm 1 as we do not need to start with the Johnson-Lindenstrauss Lemma which may slightly increase norms). In order to prove an upper bound for $\mathcal{G}$ consider, as before, a fixed Gram matrix $G = G(w_1, \ldots, w_n)$ of n vectors of norm at most 1 in R^k. Let $V_1, V_2, \ldots, V_n$ be random vectors distributed uniformly and independently in the ball of radius 2 in R^k, let T denote their Gram matrix, and let A_G be, as before, the event that $T(i,j)$ and $G(i,j)$ differ by less than $\varepsilon/2$ in each non-diagonal entry. The probability that each V_i lies in the ball of radius, say, $\delta/5$ centered at w_i is exactly $(\delta/10)^{kn}$. Conditioning on that, the probability that the inner product $(V_i - w_i)^t w_j$ has absolute value at least $\varepsilon/2$ is at most

$$2e^{-\varepsilon^2 25k/4\delta^2} < 1/(2n^2).$$

Similarly, the probability that the inner product $V_i^t(V_j - w_j)$ has absolute value at least $\varepsilon/2$ is at most

$$2e^{-\varepsilon^2 25k/8\delta^2} < 1/2n^2.$$

As before, this implies that $|\mathcal{G}| \leq 2(10/\delta)^{kn}$, establishing the assertion of the lemma. □

Lemma 2.3

For $k \leq \log n$,

$$f(n,k,\varepsilon) \leq O(nk\log(1/\varepsilon)).$$

Proof: Fix an $\varepsilon/2$-net of size $(1/\varepsilon)^{O(k)}$ in the unit ball in R^k. The sketch here is simply obtained by representing each point by the index of its closest neighbor in the net. □

3 An algorithmic proof

In this section we present an algorithmic proof of the upper bound of Theorem 1.1. We first reformulate the theorem in its algorithmic version. Note that the first part also follows from the results in [6].

Theorem 3.1 *For all* n *and* $\frac{1}{n^{0.49}} \leq \varepsilon \leq 0.1$ *there is a randomized algorithm that given a set of* n *points in* B^k *computes, for each point, a sketch of* $g(n,k,\varepsilon)$ *bits. Given two sketches, the square of the distance between the points can be recovered up to an additive error of* ε *in time* $O(\frac{\log n}{\varepsilon^2})$ *for* $\frac{\log n}{\varepsilon^2} \leq k \leq n$ *and in time* $O(k)$ *for all smaller* k. *The function* $g(n,k,\varepsilon)$ *satisfies the following*

- *For* $\frac{\log n}{\varepsilon^2} \leq k \leq n$,

  $$g(n,k,\varepsilon) = \Theta(\frac{\log n}{\varepsilon^2})$$

  and the sketch for a given point can be computed in time $O(k\log k + \log^3 n/\varepsilon^2)$.

- *For* $\log n \leq k \leq \frac{\log n}{\varepsilon^2}$,

  $$g(n,k,\varepsilon) = \Theta(k\log(2+\frac{\log n}{\varepsilon^2 k})).$$

  and the sketch for a given point can be computed in time $O(k)$.

- *For* $1 \leq k \leq \log n$,

  $$g(n,k,\varepsilon) = \Theta(k\log(1/\varepsilon))$$

  and the sketch for a given point can be computed in time $O(k)$.

In all cases the length of the sketch is optimal up to a constant factor.

As before, it is convenient to deal with the different possible ranges for k separately. Note first that the proof given in Section 2 for the range $k \leq \log n$ is essentially constructive, since it is well known (see, for example [2] or the argument below) that there are explicit constructions of ε nets of size $(1/\varepsilon)^{O(k)}$ in B^k, and it is enough to round each vector to a point of the net which is ε-close to it (and not necessarily to its nearest neighbor).

For completeness we include a short description of a δ-net which will also be used later. For $0 < \delta < 1/4$ and for $k \geq 1$ let $N = N(k, \delta)$ denote the set of all vectors of Euclidean norm at most 1 in which every coordinate is an integral multiple of $\frac{\delta}{\sqrt{k}}$. Note that each member of N can be represented by k signs and k non-negative integers n_i whose sum of squares is at most k/δ^2. Representing each number by its binary representation (or by two bits, say, if it is 0 or 1) requires at most $2k + \sum_i \log n_i$ bits, where the summation is over all $n_i \geq 2$. Note that $\sum_i \log n_i = 0.5 \log(\Pi_i n_i^2)$ which is maximized when all numbers are equal and gives an upper bound of $k \log(1/\delta) + 2k$ bits per member of the net. Given a vector in B^k we can round it to a vector of the net that lies within distance $\delta/2$ from it by simply rounding each coordinate to the closest integral multiple of $\delta/\sqrt{k}$. The computation of the distance between two points of the net takes time $O(k)$. The size of the net is $(1/\delta)^k 2^{O(k)}$, as each point is represented by $k \log(1/\delta) + 2k$ bits and k signs.

The above description of the net suffices to prove Theorem 3.1 for $k \leq \log n$. We proceed with the proof for larger k.

For $k \geq \frac{40 \log n}{\varepsilon^2}$ we first apply the Johnson-Lindenstrauss Lemma (with the fast version described in [1]) to project the points to R^m for $m = 40 \log n/\varepsilon^2$ without changing any square distance or norm by more than ε. It is convenient to now shrink all vectors by a factor of $1 - \varepsilon$ ensuring they all lie in the unit ball B^m while the square distances, norms and inner products are still within 3ε of their original values. We thus may assume from now on that all vectors lie in B^m.

As done in Section 2, we handle norms separately, namely, the sketch of each vector contains some $O(\log(1/\varepsilon))$ bits representing a good approximation for its norms. The rest of the sketch, which is its main part, will be used for recovering approximate inner products between vectors. This is done by replacing each of our vectors w_i by a randomized rounding of it chosen as follows. Each coordinate of the vector, randomly and independently, is rounded to one of the two closest integral multiples of $1/\sqrt{m}$, where the probabilities are chosen so that its expectation is the original value of the coordinate. Thus, if the value of a coordinate is $(i + p)/\sqrt{m}$ with $0 \leq p \leq 1$ it is rounded to $i/\sqrt{m}$ with probability $(1 - p)$ and to $(i + 1)/\sqrt{m}$ with probability p. Let V_i be the random vector obtained from w_i in this way. Then the expectation of each coordinate of $V_i - w_i$ is zero. For each $j \neq i$ the random variable $(V_i - w_i)^t w_j$ is a sum of m independent random variables where the expectation of each of them is 0 and the sum of squares of the difference between the maximum value of each random variable and its minimum value is the square of the norm of w_j divided by m. Therefore this sum is at most $1/m$, and by Hoeffding's Inequality (see [3], Theorem 2) the probability that this inner product is in absolute value at least

$\varepsilon/2$ is at most $2e^{-\varepsilon^2 m/8}$ which is smaller than $1/n^5$. Similar reasoning shows that the probability that $V_i^t(V_j - w_j)$ is of absolute value at least $\varepsilon/2$ is smaller than $1/n^5$. As in the proof in Section 2, it follows that with probability at least $1 - 2/n^3$ all inner products of distinct vectors in our rounded set lie within ε of their original values, as needed. The claims about the running time follow from [1] and the description above. This completes the proof of the first part of Theorem 3.1.

The proof of the second part is essentially identical (without the projection step using the Johnson-Lindenstrauss Lemma). The only difference is in the parameters. If $k = \frac{40\delta^2 \log n}{\varepsilon^2}$ with $\varepsilon \leq \delta \leq 1/2$ we round each coordinate randomly to one of the two closest integral multiples of $\delta/\sqrt{k}$, ensuring the expectation will be the original value of the coordinate. The desired result follows as before, from the Hoeffding Inequality. This completes the proof of Theorem 3.1. □

4 The lower bound

Lemma 4.1 *If*

$$k = \delta^2 \log n/(200\varepsilon^2)$$

where $2\varepsilon \leq \delta \leq 1/2$, *then* $f(n, k, \varepsilon/2) \geq \Omega(kn \log(1/\delta)$

Proof: Fix a maximal set of points N in the unit ball B^k of R^k so that the Euclidean distance between any two of them is at least δ. It is easy and well known that the size of N is $(1/\delta)^{(1+o(1))k}$ (where the $o(1)$-term tends to 0 as δ tends to 0). For the lower bound we construct a large number of ε-separated Gram matrices of n vectors in B^k. Each collection of n vectors consists of a fixed set R of $n/2$ vectors, whose existence is proved below, together with $n/2$ points of the set N. The set R of fixed points will ensure that all the corresponding Gram matrices are ε-separated.

We claim that there is a choice of a set R of $n/2$ points in B^k so that the inner products of any two distinct points from N with some point of R differ by more than ε. Indeed, for any two fixed points of N, the difference between them has norm at least δ, hence the probability that the product of a random point of B^k with this difference is bigger than ε is at least $e^{-1.5\varepsilon^2 k/\delta^2}$ (with room to spare). It thus suffices to have

$$(1 - e^{-1.5\varepsilon^2 k/\delta^2})^{n/2} < 1/|N|^2$$

hence the following will do:

$$(n/2)e^{-2\varepsilon^2 k/\delta^2} > (2 + o(1))k \log(1/\delta).$$

Thus it suffices to have

$$2\varepsilon^2 k/\delta^2 < \log(n/5k\log(1/\delta))$$

and as the left hand side is equal to $(\log n)/100$ this indeed holds. Thus a set R with the desired properties exists.

Fix a set R as above. Note that every two distinct choices of ordered sets of $n/2$ members of N provide ε-separated Gram matrices. This implies that

$$f(n,k,\varepsilon/2) \geq \log |N|^{n/2} = \Omega(n \log |N|) = \Omega(nk\log(1/\delta)),$$

completing the proof of the lemma. □

By monotonicity and the case $\delta = 1/2$ in the above Lemma the desired lower bound in Theorem 1.1 for all $k \geq \log n$ follows.

It remains to deal with smaller k. Here we fix a set N of size $(1/2\varepsilon)^{(1+o(1))k}$ in B^k so that the distance between any two points is at least 2ε. As before, the inner products with all members of a random set R of $n/2$ points distinguishes, with high probability, between any two members of N by more than ε. Fixing R and adding to it in all possible ways an ordered set of $n/2$ members of N we conclude that in this range

$$f(n,k,\varepsilon/2) \geq \log(|N|^{n/2}) = \Omega(nk\log(1/\varepsilon))$$

completing the proof of the lower bound and hence that of Theorem 1.1. □

We conclude this section by observing that the proof of the lower bound implies that the size of the sketch per point given by Theorem 3.1 is tight, up to a constant factor, for all admissible values of the parameters. Indeed, in the lower bounds we always have a fixed set R of $n/2$ points and a large net N, so that if our set contains all the points of R then no two distinct points of N can have the same sketch, as for any two distinct $u, v \in N$ there is a member of R whose inner products with u and with v differ by more than ε. The lower bound for the length of the sketch is thus $\log N$, by the pigeonhole principle.

5 Concluding remarks

- By the first two parts of Theorem 1.1, $f(n,n,2\varepsilon)$ is much bigger than $f(n,k,\varepsilon)$ for any $k < c\frac{\log n}{\varepsilon^2}$ for some absolute constant $c > 0$, implying that, as proved recently by Larsen and Nelson [7], the $\frac{\log n}{\varepsilon^2}$ bound in the Johnson-Lindenstrauss Lemma [5] is tight. The first part of Corollary 1.2 follows by a similar reasoning. It can also be derived directly from the result for $k = \log n/\varepsilon^2$. As for the "Moreover" part,

it follows by combining the Johnson-Lindenstrauss Lemma with the lower bound of Theorem 1.1.

- It is worth noting that in the proof of Theorem 3.1 the inner product of each rounded vector with itself is typically not close to the square of its original value and hence it is crucial to keep the approximate norms separately. An alternative, less natural possibility is to store two independent rounded copies of each vector and use their inner product as an approximation for its norm. This, of course, doubles the length of the sketch and there is no reason to do it. For the same reason in the proof of Theorem 1.1 in Section 2 we had to handle norms separately and consider only inner products between distinct vectors. Indeed, in this proof after the conditioning V_i is likely to have much bigger norm than w_i, and yet the inner products of distinct V_i, V_j is typically very close to that of distinct w_j, w_j.

- The problem of maintaining all square distances between the points up to a relative error of ϵ is more difficult than the one considered here. Our lower bounds, of course, hold, see [4] for the best known upper bounds. For this problem there is still a logarithmic gap between the upper and lower bounds.

Acknowledgment We thank Jaroslaw Blasiok, Kasper Green Larsen and especially Jelani Nelson for helpful comments, and for noting the relation to the paper [6].

References

[1] N. Ailon and B. Chazelle, The fast Johnson-Lindenstrauss transform and approximate nearest neighbors, SIAM J. Comput. 39 (2009), 302–322.

[2] N. Alon, T. Lee, A. Shraibman and S. Vempala, The approximate rank of a matrix and its algorithmic applications, Proc. STOC 2013, 675–684.

[3] W. Hoeffding, Probability inequalities for sums of bounded random variables, Journal of the American Statistical Association. 58 (301) (1963), 13–30.

[4] P. Indyk and T. Wagner, Near-Optimal (Euclidean) Metric Compression, arXiv 1609.06295, 2016.

[5] W. B. Johnson and J. Lindenstrauss, Extensions of Lipschitz maps into a Hilbert space, Contemp Math 26 (1984), 189–206.

[6] E. Kushilevitz, R. Ostrovsky and Y. Rabani, Efficient search for approximate nearest neighbor in high-dimensional spaces, Proc. STOC 1998, 614–623.

[7] K. G. Larsen and J. Nelson, Optimality of the Johnson-Lindenstrauss Lemma, arXiv:1609.02094, 2016.