# Reporting Neighbors in High-Dimensional Euclidean Space * 

Dror Aiger ${ }^{\dagger} \quad$ Haim Kaplan ${ }^{\ddagger} \quad$ Micha Sharir ${ }^{\text {§ }}$


#### Abstract

We consider the following problem, which arises in many database and web-based applications: Given a set $P$ of $n$ points in a high-dimensional space $\mathbb{R}^{d}$ and a distance $r$, we want to report all pairs of points of $P$ at Euclidean distance at most $r$. We present two randomized algorithms, one based on randomly shifted grids, and the other on randomly shifted and rotated grids. The running time of both algorithms is of the form $C(d)(n+k) \log n$, where $k$ is the output size and $C(d)$ is a constant that depends on the dimension $d$. The $\log n$ factor is needed to guarantee, with high probability, that all neighbor pairs are reported, and can be dropped if it suffices to report, in expectation, an arbitrarily large fraction of the pairs. When only translations are used, $C(d)$ is of the form $(a \sqrt{d})^{d}$, for some (small) absolute constant $a \approx 0.484$; this bound is worst-case tight, up to an exponential factor of about $2^{d}$. When both rotations and translations are used, $C(d)$ can be improved to roughly $6.74^{d}$, getting rid of the super-exponential factor $\sqrt{d}^{d}$. When the input set (lies in a subset of $d$-space that) has low doubling dimension $\delta$, the performance of the first algorithm improves to $C(d, \delta)(n+k) \log n$ (or to $C(d, \delta)(n+k))$, where $C(d, \delta)=O\left((e d / \delta)^{\delta}\right)$, for $\delta \leq \sqrt{d}$. Otherwise, $C(d, \delta)=$ $O\left(e^{\sqrt{d} \sqrt{d}}{ }^{\delta}\right)$.

We also present experimental results on several large datasets, demonstrating that our algorithms run significantly faster than all the leading existing algorithms for reporting neighbors.


## 1 Introduction

Problems involving distances between points in high-dimensional Euclidean space are of major importance in many areas, including pattern recognition [30], searching in multimedia data [30], vector compression [24], computational statistics [20], and data mining [41]. Most of these applications involve huge data sets (tens of billions of images or documents) and the dimensionality of the points (that is, number of real attributes representing or encoding

[^0]a point) is usually large as well (typically in the hundreds). It is therefore important to have algorithms that scale well with the input size and with the dimension.

The class of proximity problems of this kind includes nearest neighbor searching, near neighbor reporting, closest pair, diameter, minimum spanning tree and a variety of clustering problems; see [25]. In this paper we focus on the all near-neighbors reporting problem. That is, given an input set $P$ of $n$ points in $\mathbb{R}^{d}$, we wish to report all pairs of points of $P$ at (Euclidean) distance at most some given value $r$.

We continue with a review of some (rather small) subset of the vast work on such proximity problems which is most relevant to our work.

Background. For arbitrary fixed dimension $d$ and for the $L_{\infty}$-distance, the grid technique used by Lenhof and Smid [34] can be modified to solve the all near-neighbors problem in optimal $O(n+k)$ time, where $k$ is the output size. The algorithm simply inspects all pairs of points that lie in adjacent grid cells (or in the same cell) of some fixed uniform grid of cell size $r$. The efficiency of the algorithm follows from a packing lemma, established in [34] and similar to an earlier variant by Salowe [38]. It asserts that, for any $r>0$, the number of pairs at $L_{\infty}$-distance at most $2 r$ is at most a constant multiple of $n$ plus the number of pairs at $L_{\infty}$-distance at most $r$; see also Chan [15] for a simplified version.

A major drawback of the algorithms in $[15,34]$ is that their analysis only caters to $L_{\infty}$-distances. Although the algorithms also work for any $L_{p}$-distance, no analysis of the corresponding variant of the packing lemma is provided; consequently, no sharp bounds on the performance of the algorithms have been established. Our paper considers the case of Euclidean distance and provides a rather intricate analysis of this case, which replaces the simpler $L_{\infty}$-based analysis of $[15,34]$.

The use of the $L_{2}$-metric as a proximity measure is common, in particular in computer vision applications, for identifying similar images based on distances between feature vectors extracted from the images [11]. In high dimensional spaces the sparsity of the data makes the estimation of similarity by distance problematic. Specifically, consider the $L_{p}$ norm for some $p \geq 1$. It has been shown [13, 29] (see also [1]) that for $n$ points drawn independently at random from any distribution in $d$ dimensions ${ }^{1}$ (think of $n$ as fixed and of $d$ as going to infinity) the difference between the distances of the farthest point and of the nearest point from the origin increases at a rate of $d^{\frac{1}{p}-\frac{1}{2}}$. The expression $d^{\frac{1}{p}-\frac{1}{2}}$ goes to 0 for $p>2$, equals 1 for $p=2$, and goes to $\infty$ for $1 \leq p<2$. This suggests $L_{1}$ and $L_{2}$ as the best computationally simple norms for similarity estimation in high dimensional spaces, with an advantage for $L_{1}$. In this paper we consider only the $L_{2}$-metric. Extending our ideas to deal with the $L_{1}$-metric is an interesting open problem.

We note that the constants of proportionality in the bounds in [15,34] are exponential in $d$, which might make the algorithms rather inefficient when $d$ is large (as is generally the case in practice). In fact, this "curse of dimensionality" is omnipresent in most of the algorithms for proximity problems in higher dimensions, and the bounds in our paper are no exception, although we make every effort to reduce the dependence on $d$.

The results mentioned so far are for "one-shot" problems, where we are given an input set and wish to compute certain proximity attributes thereof. A considerably larger effort has been invested in preprocessing-and-query problems, where the goal is to construct some

[^1]data structure on the input point set which supports proximity queries of a certain kind. The most ubiquitous of this kind of problems is nearest neighbor searching, where we are given a set $P$ of $n$ points in a high-dimensional space $\mathbb{R}^{d}$, and wish to construct a data structure that, given a query point $q$, finds the point in $P$ closest to $q$. Extensive research on this problem has led to a variety of interesting solutions, both exact and approximate. Here too the dependence on $d$ of the performance of the resulting algorithms is at least exponential.

Many of the known exact and approximate nearest neighbor searching data structures can be modified to report all (or most) points of $P$ that are within a certain given distance $r$ from a query point $q$. Consequently, one can use such a structure to solve the all nearneighbors reporting problem, by first building a nearest-neighbor data structure for the points of $P$, and then querying this structure with each $q \in P$, thereby obtaining the points in $P$ at distance at most $r$ from $q$. The running time of such a solution is the time it takes to build the nearest neighbor data structure plus the time for $n$ queries. The space required is the total space used by the data structure plus the output size.

We give a brief review of the state-of-the-art approximate nearest neighbor data structures. For more information and related references see Har-Peled's recent book [26]. These approximate nearest neighbor data structures return, for any query point $q$, a point $p$ whose distance from $q$ is at most $(1+\varepsilon)$ times the distance between $q$ and its nearest neighbor.

A data structure based on Box-Decomposition Trees, partitioning space into axis-aligned boxes, was given by Arya et al. [9]. This structure takes $O(n)$ space, can be constructed in $O(n \log n)$ time, and answers a query in $O\left(\frac{1}{\varepsilon^{d}} \log n\right)$ time. Several improvements to this data structure have been given $[12,16,17,21]$; in particular, the query time was reduced to $O\left(\frac{1}{\varepsilon^{d}}+\log n\right) .^{2}$ A nice feature of this data structure is that $\varepsilon$ does not have to be specified at the preprocessing stage but only at query time (and can vary with the query). In subsequent work it was shown that the query time can be reduced further, at the expense of larger storage. One of the main tools used to obtain these results is the approximate Voronoi diagram of Har-Peled, with a quadtree-like search structure on top of it [27]. To date, many trade-offs between query time and space have been achieved, and in all of the more efficient ones the product of the term depending on $\varepsilon$ in the storage and the square of the term depending on $\varepsilon$ in the query time is roughly $\frac{1}{\varepsilon^{d}}$ [8]. The best known trade-off has been achieved by Arya et al. [7].

To overcome the exponential dependence on $d$ of the performance of all these data structures, Indyk and Motwani introduced a different technique called Locality Sensitive Hashing (LSH) [32]. The first component in this method is a reduction from the problem of approximate nearest neighbor search to the problem of finding a neighbor at distance $\leq(1+\varepsilon) r$ if there exists a neighbor at distance $r$, for some prespecified $r>0$. Then the latter problem is solved using a family of hash functions that tend to map close points to the same bin (an LSH family in short). If we neglect polylogarithmic factors (but this time do not hide factors exponential in the dimension) then the solution of Indyk and Motwani answers a query in $O\left(n^{\frac{1}{1+\varepsilon}}\right)$ time and takes $O\left(n^{1+\frac{1}{1+\varepsilon}}\right)$ preprocessing time and storage. This was later improved, using more complex hash functions, to $O\left(n^{\frac{1}{(1+\varepsilon)^{2}}}\right)$ query time and $O\left(n^{1+\frac{1}{(1+\varepsilon)^{2}}}\right)$ preprocessing and storage [4, 19]. For a survey on the LSH technique see [6].

[^2]Our results. In this paper we present two simple randomized algorithms for reporting all near neighbors in a set $P$ of $n$ points in a high-dimensional Euclidean space $\mathbb{R}^{d}$. That is, given a threshold distance $r$, we want to report all (or most) pairs of input points at Euclidean distance at most $r$.

Our algorithms are straightforward to implement, and indeed this work has started with the implementation of these algorithms, and with the realization that they work very well in practice. We have then aimed at a rigorous analysis of the worst-case performance of the algorithms, and its dependence on $n$, on the output size $k$, and, equally importantly, on the dimension $d$. This analysis has turned out to be rather intricate; it requires the exploitation of several interesting properties of balls, Minkowski sums, and other structures in high-dimensional spaces.

Our first (and simpler) algorithm lays down a randomly shifted uniform grid of certain cell size $c$ (that is, each cell is an axis-parallel cube with side-length $c$ ), collects the points of $P$ in each grid cell, and inspects all pairs of points within each cell (this latter feature is similar to the approach of [34]), reporting those that are at Euclidean distance at most $r$ from each other. This is repeated a suitable number of times to guarantee that, with high probability, all, or most pairs of points at Euclidean distance at most $r$ will be detected, that is, will both fall into the same grid cell. See later for full details concerning the issue of reporting all vs. most pairs.

There are two factors that determine the efficiency of the algorithm. The first factor is the number of repetitions of the grid layout procedure just described. This number depends (reciprocally) on the probability that, for a specific pair $a, b$ of points at Euclidean distance at most $r$, both $a$ and $b$ will fall into the same cell of a randomly shifted grid of cell size $c$.

For a very naive lower bound on this probability, put $x=a-b$, so $\|x\| \leq r$. The probability that $a$ and $b$ fall into different cells is at most $\left(\sum_{i=1}^{d} x_{i}\right) / c$, which, by the CauchySchwarz inequality, is at most $\|x\| \sqrt{d} / c \leq r \sqrt{d} / c$. This means that two points $a$ and $b$ at distance at most $r$ end up in the same grid cell with probability at least $1-r \sqrt{d} / c$. Notice however that this bound is meaningful only for $c \geq r \sqrt{d}$. Our experiments however showed that one can get very good results with much smaller cells, even of size close to $r$. This has motivated us to prove a better lower bound on this probability which is meaningful also for small values of $c$. The analysis of this probability estimation is presented in Section 2.

We note that a related approach is to cover the point set $P$ with random balls, such that each point is assigned to one particular ball. As shown in [18, Section 3] it is rather easy to construct such a cover with balls of diameter $c$ such that the probability that two points $a$ and $b$ are in different balls is at most $\|x\| \sqrt{d} / c$ where $x=a-b$ as before. This is similar to the randomly shifted grid technique but harder to implement as it requires to draw a random center in a union of balls. We also recall that the hash functions underlying the most efficient variants of LSH map points into balls such that close points tend to map to the same ball. Whether such a scheme based on balls can be made practical and if so how well does it perform in practice are open questions.

The other factor determining the performance of such partitioning schemes is the ratio between the number of pairs that the algorithm inspects (which, for a fixed grid, is $\sum_{\tau}\binom{|P \cap \tau|}{2}$, where the sum is over all cells $\tau$ of the grid), and the output size $k$, namely, the number of pairs at Euclidean distance at most $r$. The intuition is that if a cell $\tau$ contains many points, and if the cell size $c$ is relatively small, many pairs of these points should be
close to each other in the Euclidean metric. While this intuition is correct, sharp calibration of the number of such pairs (especially its dependence on the dimension $d$ ) turns out to be rather involved. This analysis is presented in Section 3.

There is a trade-off in the choice of the cell size $c$. On one hand we would like to make it large, to increase the probability of capturing near neighbors in the same grid cell (and thereby decrease the number of repetitions of the basic grid layout procedure), but on the other hand a large value of $c$ will increase the ratio between the number of pairs inspected by the algorithm and the number of pairs at (Euclidean) distance at most $r$ (the output size), which will increase the running time of each step. One interesting finding of our analysis is that in most cases the (theoretically) best choice for $c$ is a value very close to $r$.

Specifically, we show (see Corollary 3.4) that, for a given set $P$ of $n$ points in $\mathbb{R}^{d}$ and for a threshold distance $r$, the running time of the algorithm is $C(d)(n+k) \log n$ (for reporting, with high probability, all pairs at distance at most $r$ ), or $C(d)(n+k)$ (for reporting, in expectation, an arbitrarily large fraction of these pairs), where $C(d)=\tilde{O}\left(e^{\sqrt{d} \frac{(\sqrt{d})^{d} e^{\frac{3}{2}(\pi / 2)^{1 / 3} d^{2 / 3}}}{(\sqrt{\pi e / 2})^{d}}}\right)$. Here the " $\tilde{O}$ " notation only hides factors depending polynomially on $d, k$ is the output size, and the success probability or the expected fraction of reported pairs can be made arbitrarily close to 1 by increasing $C(d)$ by an appropriate absolute constant factor.

A major drawback of the first algorithm is that the cell size $c$ must clearly be greater than $r$ (in particular, if we want to report all pairs). In this case the ratio between the number of pairs of points in a cell and the number of these pairs at Euclidean distance at most $r$ is large in the worst case, and its dependence on $d$ is super-exponential, about $\sqrt{d}^{d}$. Informally, this is a consequence of the fact that the volume of the $d$-dimensional Euclidean ball is much smaller than the volume of a cube of the same size (the ratio between the volumes involves the super-exponential factor just mentioned).

To overcome this super-exponential dependence on $d$, we turn to our second algorithm, which is similar to the first one, except that at each step it first applies a random rotation of the coordinate frame, and then lays down a randomly shifted grid of an appropriate cell size $c$. As might seem surprising at first glance, here we can choose $c$ to be much smaller than $r$; specifically, we choose $c$ to be proportional to $r / \sqrt{d}$. An appropriate such choice of $c$ still ensures non-negligible positive probability of capturing a pair of points at Euclidean distance at most $r$ in the same grid cell. This probability is much smaller than the corresponding probability in the first algorithm (where $c$ was chosen to be much larger), and consequently the number of repetitions of the basic step is much larger. However, because the cell size is now so small, we more than compensate for this degradation by drastically reducing the ratio between the number of inspected pairs and the output size. Overall, the dependence on $d$ of the performance of the algorithm now becomes only exponential in $d$ (where the best base that this method seems to yield is about 6.74 ). Specifically, the running time of the algorithm is $O\left(6.74^{d}(n+k) \log n\right)$ (for reporting, with high probability, all pairs at distance at most $r$ ), or $O\left(6.74^{d}(n+k)\right.$ ) (for reporting, in expectation, an arbitrarily large fraction of these pairs), where the constant of proportionality depends polynomially on $d, k$ is the output size, and the success probability or the expected fraction of reported pairs can be made arbitrarily close to 1 by increasing the bound by an appropriate absolute constant factor. This algorithm and its analysis are presented in Section 5.

We note that our algorithms also lead to an algorithm for finding the closest pair in the
dataset: First we find a 2-approximation to $r$, the distance of the closest pair, and then we run either of the algorithms for this value of $r$.

In spite of all this progress, an exponential dependence on $d$ is far from satisfactory, especially when $d$ is really large. In retrospect, though, it appears that a major factor affecting the observed efficiency of the algorithms in practice is the fact that the input point sets tend to be of low doubling dimension $\delta \ll d$ with respect to the Euclidean metric [10]. That is, for every Euclidean ball $B$, the set $B \cap P$ can be covered by at most $2^{\delta}$ balls of half the radius. As it turns out, our first algorithm is suitable for handling such input sets, and its analysis can be modified, to yield a much improved dependence of its performance on $d$ and $\delta$. Specifically, the dependence is now sub-exponential in $d$ and exponential only in $\delta$; see Theorem 4.1 for the precise bound. This provides theoretical substantiation for the efficiency of the algorithm in practice. Low doubling dimension occurs naturally in practical applications and was exploited before for fast algorithms for approximate nearest neighbor searching $[28,33]$.

As already mentioned, our work has started with the implementation of our algorithms and with testing them on large data sets, observing that they work very well in practice. After providing the theoretical analysis that supports these findings, we end the paper by presenting some experimental results that manifest the efficiency of our implementations. In these results we compare our first algorithm (based on randomly shifted grids) with several of the leading existing software packages for reporting neighbors. As the results show, our algorithm runs significantly faster than any of its competitors. These results are presented in Section 6.

We note that the paper takes great care in choosing the constants that are the bases of the exponential factors that appear in our bounds, even though the practical performance of the algorithms is much better than what is predicted by the theoretical upper bounds. The main reasons for investing so much in an attempt to optimize the constants are the following. (i) It has guided us in obtaining very good practical results by tweaking the parameters, such as the cell size. (ii) We believe that the analysis itself will be useful for other related problems. (iii) The analysis indicates that our choice of constants is close to the optimal for the proposed schemes.

The techniques used by our algorithms are not new. Randomly shifted grids and random rotations have been used in the past in algorithms for various proximity problems and metric embeddings. Examples can be found in Har-Peled's book [26] and in the survey on metric embeddings by Indyk and Matoušek [31]. In particular, randomly shifted grids have been suggested and analyzed as an LSH family for the $L_{1}$-distance in [5]. A similar LSH family for $L_{2}$-distance based on projecting into quantized random lines was analyzed in [19]. Nevertheless, our precise and intricate theoretical and experimental analysis of algorithms based on randomly shifted grids for reporting all near neighbors appears to be new.

## 2 Randomly shifted grids

In this section we present and analyze our first algorithm, which is based on randomly shifted grids. To simplify the notation, and without loss of generality, we consider the problem of reporting all (or most) pairs of points at Euclidean distance at most 1 in a set $P$ of $n$ points in $\mathbb{R}^{d}$.

The simple algorithm consists of the following steps.
(a) Fix a parameter $c>1$ and fairly close to 1 ; the exact choice of $c$ will be discussed in detail below.
(b) Place a randomly shifted uniform axis-parallel grid $\mathcal{G}$ of cell size $c$ in $\mathbb{R}^{d}$, and compute, for each cell $\tau$ of $\mathcal{G}$, the subset $P_{\tau}=P \cap \tau$.
(c) For each cell $\tau$ with $\left|P_{\tau}\right| \geq 2$, go over all pairs of points of $P_{\tau}$ and report those pairs at Euclidean distance at most 1.
(d) Repeat steps (b) and (c) $M$ times, for an appropriate parameter $M$ (that depends on $c, d$, and optionally on $n$; see below for its precise choice). In doing so we filter out pairs that have already been reported.

The algorithm is clearly very simple to implement. Step (b) can be implemented in $O(n)$ time if we assume availability of a constant-cost implementation of the floor function and of hashing. Similarly, each filtering operation in step (d) can be implemented in $O(1)$ time using constant-cost hashing. All the other steps are trivial to implement.

As the analysis below will show, for appropriate choices of $c$ and $M$, the algorithm has the following properties: (i) With high probability, every pair at distance at most 1 will be reported; alternatively, with an appropriate choice of a smaller value of $M$, an arbitrarily large fraction of these pairs will be reported in expectation. (ii) The algorithm is outputsensitive, in the sense that the number of pairs that it examines is at most $C(d)(n+k) \log n$ (for reporting all pairs), or $C(d)(n+k)$ (for reporting an expected arbitrarily large fraction of the pairs), where $C(d)$ is a constant that depends on $d$, and $k$ is the output size (the number of pairs at Euclidean distance at most 1).

The constant $C(d)$ that we obtain grows super-exponentially in $d$; specifically, it is of the form $(a \sqrt{d})^{d}$, for some (small) absolute constant $a$. We show that an approach like ours must incur a constant that in the worst case is at least roughly $(\sqrt{d} / \sqrt{2 \pi e}) \approx(0.242 \sqrt{d})^{d}$. In spite of the super-exponential factor $\sqrt{d}^{d}$, it is still a challenge to make the base $a$ in our bound as small as possible. The best parameter that we can obtain is $a \approx \sqrt{\frac{2}{\pi e}} \approx 0.484$, when $d$ is sufficiently large, using some non-trivial analysis. This gets us reasonably close to the worst-case lower bound (within a factor of roughly $2^{d}$ ). Later, in Section 5 , we will be able to reduce the constant to roughly $6.74^{d}$, thereby getting rid of the super-exponential factor $\sqrt{d}^{d}$. This is achieved by using in addition random rotations of the coordinate frame.

Capturing a unit vector in a grid cell. One of the main steps in the analysis is to establish property (i) of the algorithm, namely that every pair of points at Euclidean distance at most 1 will be captured in the same cell of a randomly shifted grid with at least some probability $\zeta>0$. Then, repeating the grid construction $M=\frac{b}{\zeta} \ln n$ times, the probability that a fixed pair will not be captured will be at most $(1-\zeta)^{M} \approx e^{-b \log n}=$ $1 / n^{b}$, so, by the probability union bound, the probability that all appropriate pairs will be captured will be at least $1-\binom{n}{2} / n^{b}>1-1 /\left(2 n^{b-2}\right)$, which we can make arbitrarily small by choosing $b>2$ sufficiently large. Alternatively, repeating the process only $M^{\prime}=\frac{b}{\zeta}$ times, the success probability of capturing any fixed pair is $\approx 1-e^{-b}$. Hence, choosing $b$ arbitrarily large, we can ensure that an arbitrarily large fraction of the desired pairs will be
reported. As noted, this is sufficient in many applications.
We now proceed to estimate $\zeta$. The precise bound is stated in the summary Theorem 2.1, whose statement is based on certain choices of parameters in the forthcoming analysis.

As is intuitively clear, the worst lower bound for $\zeta$ occurs when the distance between the points in the pair is 1 ; this will be briefly justified more rigorously at the end of the analysis.

So let $c>1$ be fixed, and let $u v$ be a unit vector in $\mathbb{R}^{d}$. Let $\mathcal{G}$ be a randomly shifted uniform axis-parallel grid in $\mathbb{R}^{d}$ of cell size $c$; that is, we draw a point $o$ uniformly at random from $[0, c]^{d}$ and shift the grid so that $o$ is one of its vertices. Our goal is to derive a lower bound for the probability $\zeta$ that both $u$ and $v$ lie in the same grid cell.

Write $\overrightarrow{u v}=\left(x_{1}, x_{2}, \ldots, x_{d}\right)$, with $x_{1}^{2}+\cdots+x_{d}^{2}=1$. For each $i=1, \ldots, d$, the probability that the $x_{i}$-projections of $u$ and $v$ both lie in the same edge of a grid cell is $\left(c-\left|x_{i}\right|\right) / c$, so the probability of capturing both $u$ and $v$ in the same grid cell is

$$
\frac{\left(c-\left|x_{1}\right|\right)\left(c-\left|x_{2}\right|\right) \cdots\left(c-\left|x_{d}\right|\right)}{c^{d}} .
$$

In other words, the problem can be stated as follows. Let $\sigma$ denote the unit $(d-1)$-sphere in $\mathbb{R}^{d}$, given by $x_{1}^{2}+x_{2}^{2}+\cdots+x_{d}^{2}=1$, and let $c>1$ be a parameter. We want to find the minimum value of the function

$$
F_{c}(\mathbf{x})=\left(1-\frac{\left|x_{1}\right|}{c}\right)\left(1-\frac{\left|x_{2}\right|}{c}\right) \cdots\left(1-\frac{\left|x_{d}\right|}{c}\right)
$$

for $\mathbf{x}=\left(x_{1}, \ldots, x_{d}\right) \in \sigma$.
To simplify the analysis, let $\sigma^{+}$denote the portion of $\sigma$ in the positive orthant. It suffices to find the minimum of $F_{c}$ on $\sigma^{+}$, in which case we can rewrite $F_{c}(\mathbf{x})$ as

$$
F_{c}(\mathbf{x})=\left(1-\frac{x_{1}}{c}\right)\left(1-\frac{x_{2}}{c}\right) \cdots\left(1-\frac{x_{d}}{c}\right) .
$$

The case $c>\sqrt{2}$. Assume first that $c>\sqrt{2}$. In this case, using Lagrange multipliers, we claim that the unique extremum of $F_{c}$ in the relative interior of $\sigma^{+}$(namely, at points whose coordinates are all positive) is attained at

$$
x_{1}=x_{2}=\cdots=x_{d}=\frac{1}{\sqrt{d}},
$$

and the value of $F_{c}$ at that point is

$$
F_{c}^{(d)}=\left(1-\frac{1}{c \sqrt{d}}\right)^{d} .
$$

To see this, we use the Lagrange multipliers technique for finding extrema of $F_{c}(\mathbf{x})$ within the relative interior of the sphere $\sum_{i} x_{i}^{2}=1$ inside the positive orthant. This yields the equations

$$
\begin{equation*}
2 \lambda x_{i}=-\frac{1}{c-x_{i}} F_{c}(\mathbf{x}), \quad \text { for } i=1, \ldots, d, \tag{1}
\end{equation*}
$$

or, equivalently, all the products $x_{i}\left(c-x_{i}\right)$ are equal. Thus $x_{i}\left(c-x_{i}\right)=x_{j}\left(c-x_{j}\right)$ for any $i \neq j$, which is equivalent to $\left(x_{i}-x_{j}\right)\left(x_{i}+x_{j}-c\right)=0$. However, since $x_{i}^{2}+x_{j}^{2} \leq 1$ we have
$x_{i}+x_{j} \leq \sqrt{2}<c$, so we must have $x_{i}=x_{j}$, implying that all coordinates at the extremum are equal.

The other local extrema of $F_{c}$ on $\sigma^{+}$must occur on the boundary of $\sigma^{+}$, at points where some coordinates are 0 . By symmetry, it suffices to examine $d-1$ additional subproblems, where in the $j$-th subproblem, $j=1, \ldots, d-1$, we seek the extrema of $F_{c}$ on

$$
\sigma_{j}^{+}=\sigma^{+} \cap\left\{\mathbf{x} \mid x_{j+1}=x_{j+2}=\cdots=x_{d}=0\right\}
$$

Arguing as in the case of the whole $\sigma^{+}$, the unique extremum of $F_{c}$ in the relative interior of $\sigma_{j}^{+}$is attained at

$$
x_{1}=x_{2}=\cdots=x_{j}=\frac{1}{\sqrt{j}}, \quad x_{j+1}=x_{j+2}=\cdots=x_{d}=0
$$

and the value of $F_{c}$ at that point is

$$
F_{c}^{(j)}=\left(1-\frac{1}{c \sqrt{j}}\right)^{j}
$$

In other words, the minimum of $F_{c}$ over $\sigma$ is the minimum of the sequence

$$
s_{j}(c)=\left(1-\frac{1}{c \sqrt{j}}\right)^{j}, \quad j=1, \ldots, d
$$

To simplify matters further, consider instead the minimum of the sequence

$$
t_{j}=\ln s_{j}(c)=j \ln \left(1-\frac{1}{c \sqrt{j}}\right), \quad j=1, \ldots, d
$$

or rather its continuous version

$$
f(x)=x \ln \left(1-\frac{1}{c \sqrt{x}}\right)
$$

We have

$$
\begin{aligned}
f^{\prime}(x)= & \ln \left(1-\frac{1}{c \sqrt{x}}\right)+\frac{x}{1-\frac{1}{c \sqrt{x}}} \cdot \frac{1}{2 c x \sqrt{x}}= \\
& \ln \left(1-\frac{1}{c \sqrt{x}}\right)+\frac{1}{2(c \sqrt{x}-1)} .
\end{aligned}
$$

Put $y=c \sqrt{x}$. Since we are interested in the range $x \geq 1$, we have $y>\sqrt{2}$. However, to accommodate also the case $1<c \leq \sqrt{2}$, discussed next, we consider the extended range $y>1$. We need to solve

$$
2 \ln \left(1-\frac{1}{y}\right)+\frac{1}{y-1}=0
$$

Note that this equation is independent of $c$. Using the WolframAlpha software and some numerical calculations, we find that there is a unique zero at $y_{0} \approx 1.39795$, where $f$ attains its maximum ( $f^{\prime}$ is positive for $y<y_{0}$ and negative for $y>y_{0}$ ). That means that the maximum of $f$ is attained at $x_{0}=y_{0}^{2} / c^{2}<0.977<1($ for $c>\sqrt{2})$.

In particular, $f$ decreases for $x>x_{0}$, and therefore the minimum value of $s_{j}(c)$, for $j \geq 1>x_{0}$, is attained at $j=d$. That is, we have shown that

$$
\begin{equation*}
F_{c}(\mathbf{x}) \geq\left(1-\frac{1}{c \sqrt{d}}\right)^{d} \tag{2}
\end{equation*}
$$

for all $\mathbf{x} \in \sigma$, provided that $c>\sqrt{2}$.

The case $1<c \leq \sqrt{2}$. We first note that the analysis given for the preceding case continues to apply here too, in the sense that all the values $s_{j}(c)$, for $1 \leq j \leq d$, continue to be local extrema of $F_{c}$ (more precisely, $s_{j}(c)$ continues to be a local extremum in the relative interior of a $(j-1)$-dimensional sub-sphere of $\sigma)$. The maximum of the corresponding continuous function $f(x)$ is still attained at $x=y_{0}^{2} / c^{2}$ which is smaller than 2 for any $c>1$, implying that the suffix $s_{2}(c), s_{3}(c), \ldots, s_{d}(c)$ is decreasing, so its minimum is still $s_{d}(c)$, as given in (2). Here however we have a new contender for the minimum, namely

$$
s_{1}(c)=1-\frac{1}{c}
$$

which, when $c$ is very close to 1 , can indeed be smaller than $s_{d}(c)$. To avoid this case, we will require that

$$
1-\frac{1}{c} \geq e^{-\frac{\sqrt{d}}{c}} \geq\left(1-\frac{1}{c \sqrt{d}}\right)^{d}
$$

The right inequality always holds, being an instance of the inequality $e^{-x} \geq 1-x$ for $x \geq 0$. The left inequality will hold, for $c \leq \sqrt{2}$, if we choose

$$
c \geq \frac{1}{1-e^{-\sqrt{d / 2}}}
$$

which in turn holds if

$$
\begin{equation*}
c \geq 1+2 e^{-\sqrt{d / 2}} \tag{3}
\end{equation*}
$$

The latter implication follows since $1+2 x \geq \frac{1}{1-x}$ for $x \leq 1 / 2$, and $e^{-\sqrt{d / 2}} \leq 1 / 2$ for $d \geq 2$.

In the remainder of this section we will only consider the case where $c$ satisfies (3).
Equation (3) guarantees that $s_{1}(c) \geq s_{d}(c)$ but this still does not make $s_{d}(c)$ the minimum of $F_{c}(x)$ over $\sigma^{+}$since for $1<c \leq \sqrt{2}, F_{c}$ has some additional local extrema. Specifically, using the Lagrange multipliers technique, and the implied analysis following Equation (1), we conclude that, at an extremum $\mathbf{x}$ of $F_{c}$ on $\sigma^{+}$, some coordinates $x_{i}$ may be zero, and any pair $x_{i}, x_{j}$ of nonzero coordinates must satisfy either $x_{i}=x_{j}$ or $x_{i}+x_{j}=c$. This implies right away that (on $\sigma^{+}$) the nonzero coordinates of $\mathbf{x}$ have only two distinct values, which we denote as $t$ and $c-t$, and assume that $t \leq c-t$, so $c-1 \leq t \leq \frac{c}{2}$.

Suppose that we have $a$ coordinates equal to $t$ and $b$ coordinates equal to $c-t$, with $a+b \leq d$. We then must have

$$
\begin{equation*}
a t^{2}+b(c-t)^{2}=1 \tag{4}
\end{equation*}
$$

Since $t \leq c-t$ we get that $c-t>1 / 2$, which is easily seen to imply that $b \leq 3$.

The case $t \geq 1-\frac{1}{\sqrt{2}}$. In this case we have

$$
a=\frac{1-b(c-t)^{2}}{t^{2}} \leq \frac{1-b / 4}{t^{2}} \leq \frac{3 / 4}{3 / 2-\sqrt{2}}=3(3 / 2+\sqrt{2})<9
$$

so $a \leq 8$. In this case we have

$$
F_{c}(\mathbf{x})=\left(\frac{t}{c}\right)^{b}\left(1-\frac{t}{c}\right)^{a} \geq\left(\frac{t}{c}\right)^{3}\left(1-\frac{t}{c}\right)^{8}
$$

Since $t / c \geq t / \sqrt{2} \geq(\sqrt{2}-1) / 2$ and $1-t / c=(c-t) / c \geq(c-t) / \sqrt{2} \geq 1 /(2 \sqrt{2}), F_{c}(\mathbf{x})$ is at least some absolute constant

$$
\mu \geq\left(\frac{\sqrt{2}-1}{2}\right)^{3}\left(\frac{1}{2 \sqrt{2}}\right)^{8}
$$

Since $s_{d}(c)$ decreases to 0 as $d$ grows (as argued above, it is upper bounded by $e^{-\sqrt{d} / c} \leq$ $\left.e^{-\sqrt{d / 2}}\right), F_{c}(\mathbf{x})$, for any new local extremum $\mathbf{x}$, will be larger than $s_{d}(c)$ when $d$ is at least some sufficiently large constant $d_{0}$ (independent of $c$ ).

The case $t \leq 1-\frac{1}{\sqrt{2}}$. In this case $1-t \geq \frac{1}{\sqrt{2}}$, or $(c-t)^{2}>1 / 2$, and then (4) implies that $b=1$, and also that

$$
a=\frac{1-(c-t)^{2}}{t^{2}}<\frac{2 t-t^{2}}{t^{2}}=\frac{2}{t}-1<\frac{2}{t} .
$$

Substituting $a$ and $b$ into the expression for $F_{c}(\mathbf{x})$, we get that

$$
F_{c}(\mathbf{x}) \geq \frac{t}{c}\left(1-\frac{t}{c}\right)^{2 / t} \geq \frac{t}{c} e^{-4 / c}
$$

Here we use the inequality $1-x \geq e^{-2 x}$, which holds for $x=\frac{t}{c}<1 / 2$. If

$$
\begin{equation*}
t e^{-4 / c} \geq c e^{-\sqrt{d} / c} \tag{5}
\end{equation*}
$$

then $F_{c}(\mathbf{x})=\frac{t}{c} e^{-4 / c}$ will be larger than $s_{d}(c)$. Inequality (5) is equivalent to

$$
t \geq \frac{c}{e^{\frac{\sqrt{d}-4}{c}}}
$$

If $c-1 \geq \frac{c}{e^{\frac{\sqrt{d}-4}{c}}}$ we are done since we always have $t \geq c-1$, so the new extremal values of $F_{c}$ will all be greater than $s_{d}(c)$. (Note that for this inequality to hold, $d$ must be at least 17.) So assume that $c-1<\frac{c}{e^{\frac{\sqrt{d}-4}{c}}}$, or

$$
\begin{equation*}
c<1+\frac{c}{e^{\frac{\sqrt{d}-4}{c}}} \leq 1+\frac{\sqrt{2}}{e^{\frac{\sqrt{d}-4}{\sqrt{2}}}} \leq 1+\sqrt{2} e^{2 \sqrt{2}} e^{-\sqrt{d / 2}}<1+24 e^{-\sqrt{d / 2}} \tag{6}
\end{equation*}
$$

It remains to consider the case where $c$ satisfies this inequality and $c-1<t \leq \frac{c}{e^{\frac{\sqrt{d}-4}{c}}}$. Here we use the fact that $a \leq d-1$. So we also have the bound

$$
F_{c}(\mathbf{x}) \geq(t / c)(1-t / c)^{d-1}
$$

The function $g(t)=(t / c)(1-t / c)^{d-1}$ on the right-hand side has a maximum at $t=c / d$, as is easily verified. Assume that $d$ is sufficiently large so that

$$
\frac{c}{e^{\frac{\sqrt{d}-4}{c}}} \leq \frac{c}{d}
$$

Then, for $c-1<t \leq \frac{c}{e^{\frac{\sqrt{d}-4}{c}}} \leq \frac{c}{d}$, we have $F_{c}(\mathbf{x})>g(c-1)=(c-1) / c^{d}$.

However, for $d$ sufficiently large, we have

$$
\frac{c-1}{c^{d}} \geq\left(1-\frac{1}{c \sqrt{d}}\right)^{d}=s_{d}(c)
$$

Indeed, recall that $c$ is assumed to satisfy (3) and (6); that is,

$$
1+2 e^{-\sqrt{d / 2}} \leq c<1+24 e^{-\sqrt{d / 2}}
$$

We then have (using the inequality $1+x<e^{x}$ for $x>0$ )

$$
\frac{c-1}{c^{d}}>\frac{2 e^{-\sqrt{d / 2}}}{\left(1+24 e^{-\sqrt{d / 2}}\right)^{d}}>\frac{2 e^{-\sqrt{d / 2}}}{e^{24 d e^{-\sqrt{d / 2}}}}
$$

which, for $d$ sufficiently large, will be larger than

$$
e^{-\sqrt{d / 2}} \geq e^{-\sqrt{d} / c} \geq\left(1-\frac{1}{c \sqrt{d}}\right)^{d}
$$

as claimed.
In summary, we have obtained the following result.

Theorem 2.1 There exists some threshold dimension $d_{0}$ so that, for every $d \geq d_{0}$ and for every $c$ satisfying $c \geq 1+2 e^{-\sqrt{d / 2}}$, we have $F_{c}(\mathbf{x}) \geq s_{d}(c)$ for every $\mathbf{x} \in \sigma$. As a consequence, the probability of capturing both endpoints of a unit vector in $\mathbb{R}^{d}$ inside the same cell of a randomly shifted grid of cell size $c$ is at least $s_{d}(c)$.

Remarks. (1) The analysis so far has only considered unit vectors, but it can easily be adapted to apply to vectors of smaller length. The simplest way of doing this is perhaps as follows. If the length of the vector $u v$ is $r<1$ then, by scaling space by the factor $1 / r$, we turn $u v$ into a unit vector, and $c$ is replaced by the larger parameter $c / r$. The probability of capturing $u$ and $v$ in the same cell is now at least $s_{d}(c / r)$, which is larger than $s_{d}(c)$.
(2) It is perhaps tempting to use a larger value of $c$ so as to increase the probability $s_{d}(c)$ of capturing $u$ and $v$ in the same cell. The value of $s_{d}(c)$, for $c$ close to 1 , e.g., for $c$ meeting the lower bound in (3), is rather small; namely it is approximately $e^{-\sqrt{d} / c} \approx e^{-\sqrt{d}}$. However, to make this probability really large, we need to choose $c$ quite large, about $\sqrt{d}$. The penalty for using such a large grid size will be incurred in having to inspect too many pairs of points in $P$. This trade-off, between (i) the success probability of capturing a unit vector in a grid cell, and (ii) the ratio between the number of pairs at Euclidean distance at most 1 and the overall number of pairs in a grid cell, is a major theme in our analysis, here and in the following sections.
(3) It would be interesting to obtain a sharp calibration of the value of the threshold dimension $d_{0}$, and to analyze the situation when $d<d_{0}$. We have not made an attempt to sharply calibrate $d_{0}$. The rather weak inequalities that we have employed lead to a value close to 100 , but the actual threshold is probably much smaller.

As already noted at the beginning of the section, if we repeat the randomly shifted grid construction $\frac{b}{s_{d}(c)} \log n$ (or just $\left.\frac{b}{s_{d}(c)}\right)$ times, for some $b>2$ sufficiently large, then, with
the first choice, we ensure, with high probability, that every pair of points at Euclidean distance at most 1 will be captured in the same cell of one of the grids; with the second choice, we ensure an arbitrarily large expected number of such pairs. This however is only one side of the story, because, as noted in Remark (2) above, the algorithm examines all pairs of points in the same grid cell, and we wish to ensure that the latter number is not significantly larger than the number of pairs, within the same cell, at Euclidean distance at most 1. This analysis will be undertaken in the next section.

## 3 The number of Euclidean neighbors in a grid cell

Let $\mathcal{G}$ be a randomly shifted grid, of cell size $c$, with $c$ satisfying the condition in (3). Let $\tau$ be a cell of $\mathcal{G}$ and put $P_{\tau}=P \cap \tau$. The algorithm of Section 2 examines every pair of points of $P_{\tau}$, and reports those pairs at Euclidean distance at most 1 . In this section we argue that the number of pairs examined by the algorithm, namely $\binom{\left|P_{\tau}\right|}{2}$, is comparable with the number $N_{2}\left(P_{\tau}\right)$ of pairs of points of $P_{\tau}$ at Euclidean distance at most 1. Specifically, we show that $N_{2}\left(P_{\tau}\right) \geq \gamma(d)\binom{\left|P_{\tau}\right|}{2}-\frac{1}{2}\left|P_{\tau}\right|$, where $\gamma(d)$ is some constant fraction that depends on $d$.

To establish this inequality, we fix a grid cell $\tau$ of size $c>0$ (e.g., the value chosen in Section 2; however, we will later on consider values of $c$ smaller than 1, and the analysis in this section applies to these choices too). For simplicity, we regard $P_{\tau}$ as the entire set $P$, and denote its size by $n$.

An upper bound on $\gamma(d)$. To understand what values of $\gamma(d)$ one might expect to get, we start with a simple upper bound on that quantity. Let $P$ be a set of $n$ points drawn independently and uniformly at random in $\tau$. A ball of radius 1 around one such point $p$ contains, in expectation, at most $\frac{V_{d} \cdot(n-1)}{c^{d}}$ points of $P \backslash\{p\}$, where

$$
V_{d}=\frac{\pi^{d / 2}}{\Gamma(1+d / 2)}
$$

is the volume of a unit ball in $\mathbb{R}^{d} ; \Gamma(\cdot)$ is the Gamma function; see, e.g., [40, p. 136]. This bound on the expectation follows from the fact that the points of $P$ are drawn independently, so the probability of another point $q \in P$ to fall into the ball is the volume of its portion within $\tau$ divided by $\operatorname{Vol}(\tau)=c^{d}$. The point $p$ forms with each point $q \in P$ in this ball a pair at Euclidean distance at most 1, and every such pair $(p, q)$ arises exactly twice in this manner, once for the ball centered at $p$ and once for the ball centered at $q$.

Summing up over all points $p \in P$, we get that the expected number of pairs at Euclidean distance at most 1 is at most

$$
\begin{equation*}
\frac{n}{2} \cdot \frac{V_{d} \cdot(n-1)}{c^{d}}=\binom{n}{2} \frac{V_{d}}{c^{d}} \tag{7}
\end{equation*}
$$

This shows that we cannot hope for a worst-case lower bound for $\gamma(d)$ larger than $V_{d} / c^{d}$. Using the above expression for $V_{d}$ and applying Stirling's approximation, we get that, up to factors polynomial in $d$,

$$
\begin{equation*}
\gamma(d) \leq \frac{V_{d}}{c^{d}}=\frac{\pi^{d / 2}}{\Gamma(1+d / 2) c^{d}} \approx \frac{(\sqrt{2 \pi e})^{d}}{\sqrt{d}^{d} c^{d}} \tag{8}
\end{equation*}
$$

In other words, in this case (assuming that $c$ is sufficiently close to 1 ) the number of pairs that the algorithm examines within $\tau$ is at least $\sqrt{d}^{d} c^{d} /(\sqrt{2 \pi e})^{d} \approx \sqrt{d}^{d} / 4.14^{d}$ times the number of pairs at Euclidean distance 1. This shows that to get rid of the super-exponential factor $\sqrt{d}^{d}$ we have to use a very small value of $c$, proportional to $1 / \sqrt{d}$. We clearly cannot do that when only random translations are used. However, as will be shown in Section 5, one can choose such small values for $c$ when random rotations are also employed.

A lower bound for $\gamma(d)$. We partition the cell $\tau$ of size $c$ into a grid of $(c \sqrt{d})^{d}$ smaller cells of size $\frac{1}{\sqrt{d}}$ each. Let $\xi$ be one of the small cells, and let $n_{\xi}$ be the number of points of $P$ in $\xi$. Every pair of points in $\xi$ are at distance at most 1 , so we have at least $\sum_{\xi}\binom{n_{\xi}}{2}$ pairs at distance at most 1 in $\tau$. Using the Cauchy-Schwarz inequality we get that

$$
\begin{equation*}
N_{2}\left(P_{\tau}\right) \geq \sum_{\xi}\binom{n_{\xi}}{2}=\sum_{\xi} \frac{n_{\xi}\left(n_{\xi}-1\right)}{2} \geq \frac{n^{2}}{2(c \sqrt{d})^{d}}-\frac{n}{2} \tag{9}
\end{equation*}
$$

Another lower bound for $\gamma(d)$. This lower bound is not the best that can be achieved (as will be demonstrated later), but we include it to introduce and motivate the analysis that follows in the next subsection. We enlarge $\tau$ by shifting each of its facets $f$ by $1 / 2$ away from $\tau$ in the direction orthogonal to $f$; let $\tau^{*}$ denote the cube formed by the hyperplanes that support the shifted facets. We then consider a ball of radius $1 / 2$ whose center is picked uniformly at random in $\tau^{*}$. The probability that this ball contains a point $p \in \tau$ is at least $V_{d}\left(\frac{1}{2}\right)^{d} /(c+1)^{d}$, which is the volume of such a ball divided by the volume of $\tau^{*}$. (The random ball contains $p$ iff we pick its center in a ball of radius $1 / 2$ around $p$, and this latter ball is contained in its entirety in $\tau^{*}$.) It follows, in particular, that there is a ball of radius $1 / 2$ that contains at least $\frac{V_{d}\left(\frac{1}{2}\right)^{d}}{(c+1)^{d}} n$ points of $P$. We take one such ball, and remove from $P$ all points in that ball. As long as we have not removed more than $n / 2$ points, a probabilistic argument similar to the one just given implies that there is a ball of radius $1 / 2$ containing at least $\frac{V_{d}\left(\frac{1}{2}\right)^{d}}{(c+1)^{d}} \frac{n}{2}$ remaining points of $P$. So we keep collecting such balls, each containing a set of at least $\frac{V_{d}\left(\frac{1}{2}\right)^{d}}{(c+1)^{d}} \frac{n}{2}$ points of $P$, so that these sets are pairwise disjoint, until we are left with fewer than $n / 2$ points. The number of balls that we collect is therefore at most $\frac{(c+1)^{d}}{V_{d}\left(\frac{1}{2}\right)^{d}}$. Let $n_{i}$ be the number of points associated with the $i$ th ball. By construction we have $\sum_{i} n_{i} \geq n / 2$. Since each pair of points associated with a particular ball are at Euclidean distance at most 1, the Cauchy-Schwarz inequality implies, as above, that at least

$$
\begin{equation*}
\sum_{i}\binom{n_{i}}{2}=\sum_{i} \frac{n_{i}\left(n_{i}-1\right)}{2} \geq \frac{n^{2}}{8} \frac{V_{d}\left(\frac{1}{2}\right)^{d}}{(c+1)^{d}}-\frac{n}{2} \tag{10}
\end{equation*}
$$

pairs of points are at Euclidean distance at most 1. Substituting the explicit expression for $V_{d}$, and using Stirling's approximation, as in Equation (8), we get that this lower bound is approximately

$$
\frac{n^{2}}{8} \frac{(\sqrt{\pi e / 2})^{d}}{\sqrt{d}^{d}(c+1)^{d}}-\frac{n}{2}
$$

The main drawback in this second lower bound is the presence of the factor $(c+1)^{d}$ in the denominator. (If we could have replaced $(c+1)^{d}$ by $c^{d}$ then since $\sqrt{\pi e / 2} \approx 2.07$ the
second lower bound in Equation (10) would have clearly dominated the first lower bound in Equation (9).) The factor $(c+1)^{d}$ essentially makes the lower bound in Equation (10) useless when $c$ is very small, such as $O\left(\frac{1}{\sqrt{d}}\right)$; such small values of $c$ will indeed be used in Section 5. In the next subsection we overcome this drawback by drawing the center of the random ball of radius $1 / 2$ in the Minkowski sum of $\tau$ and a ball of radius $1 / 2$, rather than in the much larger cell $\tau^{*}$ (and we will show that the Minkowski sum is indeed much smaller than $\left.\tau^{*}\right)$.

### 3.1 The Minkowski sum of a cube and a ball

Let $K$ denote the unit cube $[0,1]^{d}$ in $\mathbb{R}^{d}$, and let $r>0$ be a parameter. Let $B_{r}$ denote the ball of radius $r$ centered at the origin, and let $K_{r}=K \oplus B_{r}$ denote their Minkowski sum. Our next task is to estimate the volume of $K_{r}$, which is accomplished in the following lemma. ${ }^{3}$

## Lemma 3.1

$$
\operatorname{Vol}\left(K_{r}\right) \leq \begin{cases}\beta e^{\frac{3}{2}(2 \pi)^{1 / 3} d^{2 / 3} r^{2 / 3}} & \text { for } r \leq \sqrt{\frac{d}{2 \pi}}  \tag{11}\\ \left.\frac{(2 \pi)^{1 / 2} e^{3 / 2} r}{d^{1 / 2}}\right)^{d} & \text { for } r>\sqrt{\frac{d}{2 \pi}},\end{cases}
$$

where $\beta$ is a constant that depends polynomially on $d$.

Proof. We note that $\operatorname{Vol}\left(K_{r}\right)$ can be expressed as a special case of Steiner's formula, which asserts that, for any convex body $K$ in $\mathbb{R}^{d}$,

$$
\operatorname{Vol}\left(K \oplus B_{r}\right)=\sum_{j=0}^{d}\binom{d}{j} W_{j}(K) r^{j},
$$

where $W_{j}(K)$, called the $j$-th quermassintegral of $K$, is the mixed volume

$$
V(\underbrace{K, K, \ldots, K}_{d-j \text { times }}, \underbrace{B, B, \ldots, B}_{j \text { times }})
$$

of $d-j$ copies of $K$ and $j$ copies of the unit ball $B$; see, e.g., [22, 39]. As follows from the forthcoming analysis, the unit cube $K$ has the property that $W_{j}(K)$ is the volume $V_{j}$ of the $j$-dimensional unit ball, for $j=0, \ldots, d$. We did not manage to find in the literature an explicit statement of this property, although we are fairly certain that it has been made. It can be established using properties of Minkowski sums involving segments (the cube, as well as any of its faces, is such a sum); see, e.g., formula (A.43), p. 407, in [22]. We establish this property in the following explicit manner.

We first obtain the following representation of $K_{r}$. For each $j=0, \ldots, d$ there are $2^{j}\binom{d}{j}$ $(d-j)$-faces of $K$, so that each such face $f$ is defined in terms of a subset $\sigma$ of $\{1, \ldots, d\}$ of size $j$ and a mapping $\delta: \sigma \mapsto\{0,1\}$, so that

$$
f=\left\{\mathbf{x} \in K \mid x_{i}=\delta(i), \text { for } i \in \sigma\right\} .
$$

[^3]That is, $f$ is defined by fixing $j$ coordinates, each to 0 or 1 . Given $\sigma$ and $\delta$, we denote by $f_{\sigma, \delta}$ the corresponding face of $K$.

For each such face $f_{\sigma, \delta}$, we "push it outwards" by forming its Minkowski sum with an appropriate portion of a $j$-dimensional ball of radius $r$, in the directions of the coordinates that are fixed on $f_{\sigma, \delta}$. Technically, for $i \in \sigma$ and for $\varepsilon \in\{0,1\}$, define the hyperplane

$$
H_{i, \varepsilon}= \begin{cases}\left\{x_{i} \geq 0\right\} & \text { if } \varepsilon=1 \\ \left\{x_{i} \leq 0\right\} & \text { if } \varepsilon=0\end{cases}
$$

Now put

$$
B_{r}^{(j)}(\delta)=B_{r}^{(j)} \cap \bigcap_{i \in \sigma} H_{i, \delta(i)},
$$

where $B_{r}^{(j)}$ is the $j$-dimensional ball of radius $r$ centered at the origin, within the subspace spanned by the fixed coordinates of $\sigma$, and put

$$
K_{r}\left(f_{\sigma, \delta}\right)=f_{\sigma, \delta} \oplus B_{r}^{(j)}(\delta)
$$

With all these preparations, we can finally write $K_{r}$ as the disjoint union

$$
K_{r}=\bigcup_{j=0}^{d} \bigcup_{\sigma, \delta} K_{r}\left(f_{\sigma, \delta}\right)
$$

where the inner union is over all $\binom{d}{j}$ possible choices $\sigma$ of $j$ coordinates and $2^{j} 0 / 1$ "sign patterns" $\delta$ on these coordinates. Note that the special case $j=0$ yields $K$ itself (there are no coordinates to be "pushed"). (The proofs of the above equality and of the disjointness of the constituents of the union are straightforward albeit a bit tedious, and we omit them.)


Figure 1: The Minkowski sum of a square and a disk.
The reader is invited to check Figure 1 for an illustration of the case $d=2$, and to work out the concrete form of the above expression for $d=3$, where the "fringe" of $K_{r}$ consists of six boxes (obtained by pushing out the facets), twelve quarters of cylinders (obtained by pushing out the edges), and eight eighths of balls (obtained by pushing out the vertices). ${ }^{4}$

The volume of $K_{r}$. Note that each of the sums $K_{r}\left(f_{\sigma, \delta}\right)=f_{\sigma, \delta} \oplus B_{r}^{(j)}(\delta)$ is a $\left(1 / 2^{j}\right)$ portion of the cylinder $f_{\sigma, \delta} \oplus B_{r}^{(j)}$ whose axis is $f_{\sigma, \delta}$ and whose radius is $r$. Since the volume

[^4]of any face $f$ of $K$ (within its affine hull) is always 1 , the volume of $K_{r}\left(f_{\sigma, \delta}\right)$ is $1 / 2^{j}$ of the volume of $B_{r}^{(j)}$, which we can write as $V_{j} r^{j}$, where, as before,
$$
V_{j}=\frac{\pi^{j / 2}}{\Gamma(1+j / 2)}
$$
is the volume of a unit $j$-dimensional ball. Altogether we thus obtain
\[

$$
\begin{equation*}
\operatorname{Vol}\left(K_{r}\right)=\sum_{j=0}^{d}\binom{d}{j} \cdot 2^{j} \cdot \frac{1}{2^{j}} \cdot V_{j} r^{j}=\sum_{j=0}^{d}\binom{d}{j} \frac{\pi^{j / 2} r^{j}}{\Gamma(1+j / 2)} \tag{12}
\end{equation*}
$$

\]

Note that this establishes our claim that the $j$-th quermassintegral $W_{j}(K)$ of the unit cube $K$ is $V_{j}$, for each $j=0, \ldots, d$.

We approximate the sum, up to factors that depend polynomially on $d$, as follows. We use the inequalities

$$
\begin{gathered}
\binom{d}{j} \leq \frac{d^{j}}{j!} \quad \text { and } \\
\Gamma(1+j / 2) \approx \frac{j(j-2)(j-4) \cdots}{2^{j / 2}} \geq \frac{(j!)^{1 / 2}}{2^{j / 2}}
\end{gathered}
$$

where $\approx$ denotes equality up to a factor that depends polynomially in $d$. Hence

$$
\operatorname{Vol}\left(K_{r}\right) \leq \beta_{0} \sum_{j=0}^{d} \frac{d^{j}}{j!} \cdot \frac{2^{j / 2}}{(j!)^{1 / 2}} \cdot \pi^{j / 2} r^{j}
$$

for a suitable constant $\beta_{0}$ depending (polynomially) on $d$. Using Stirling's formula we further obtain that

$$
\begin{equation*}
\operatorname{Vol}\left(K_{r}\right) \leq \beta_{1} \sum_{j=1}^{d}\left(\frac{2^{1 / 2} \pi^{1 / 2} e^{3 / 2} d r}{j^{3 / 2}}\right)^{j} \tag{13}
\end{equation*}
$$

for another similar constant $\beta_{1}$.
To obtain an upper bound for this sum, we replace its terms by the corresponding function

$$
f(x)=\left(\frac{A}{x^{3 / 2}}\right)^{x}, \quad \text { where } \quad A=2^{1 / 2} \pi^{1 / 2} e^{3 / 2} d r
$$

over $x \geq 1$. It is simpler to analyze

$$
g(x)=\ln f(x)=x \ln A-\frac{3}{2} x \ln x .
$$

We have

$$
g^{\prime}(x)=\ln A-\left(\frac{3}{2} \ln x+\frac{3}{2}\right)
$$

so $g^{\prime}(x)=0$ at $x=x_{0}=A^{2 / 3} / e=(2 \pi)^{1 / 3} d^{2 / 3} r^{2 / 3}$. It can easily be checked that $g(x)$ attains its maximum at this value, and so does $f(x)$. The maximum value of $f$ is therefore

$$
f\left(x_{0}\right)=f\left(A^{2 / 3} / e\right)=e^{\frac{3}{2} x_{0}}=e^{\frac{3}{2}(2 \pi)^{1 / 3} d^{2 / 3} r^{2 / 3}} .
$$

Two cases can arise:
(i) Assume first that $A^{2 / 3} / e<d$; that is, $r<\sqrt{\frac{d}{2 \pi}}$. In this case $\operatorname{Vol}\left(K_{r}\right)$ is at most $(d+1) \beta_{1} f\left(x_{0}\right)$; that is,

$$
\operatorname{Vol}\left(K_{r}\right) \leq \beta f\left(x_{0}\right)=\beta e^{\frac{3}{2}(2 \pi)^{1 / 3} d^{2 / 3} r^{2 / 3}}
$$

for another suitable constant $\beta=(d+1) \beta_{1}$ that depends polynomially on $d$.
(ii) Assume next that $A^{2 / 3} / e \geq d, r \geq \sqrt{\frac{d}{2 \pi}}$. In this case the maximum value of $f(x)$ in our range is attained at $x=d$, and is then equal to $\left((2 \pi)^{1 / 2} e^{3 / 2} r / d^{1 / 2}\right)^{d}$. Hence

$$
\operatorname{Vol}\left(K_{r}\right) \leq \beta\left(\frac{(2 \pi)^{1 / 2} e^{3 / 2} r}{d^{1 / 2}}\right)^{d}
$$

for the same constant $\beta=(d+1) \beta_{1}$ as above.
This completes the proof of Lemma 3.1.
Hence, as long as $r=o\left(d^{1 / 2}\right), \operatorname{Vol}\left(K_{r}\right)=e^{o(d)}$.

Remark. To appreciate the significance of this bound, note that if $r=\sqrt{d}$ (choosing $r$ to be any constant multiple of $\sqrt{d}$ would lead to a similar phenomenon) then $K_{r}$ fully contains the cube $[-1,2]^{d}$, namely the cube obtained by "pushing" $K$ by 1 in all directions. Indeed the points of the expanded cube that are farthest from $K$ are its vertices, and each of them lies at distance $\sqrt{d}$ from a corresponding vertex of $K$. In other words, for $r=\sqrt{d}$ the volume of $K_{r}$ is at least $3^{d}$, but for sufficiently smaller values of $r \ll \sqrt{d}$, its volume becomes, relatively speaking, negligible. More precisely, it is then larger than $\operatorname{Vol}(K)=1$ by only a sub-exponential factor.

The case of a grid cell of size $c$. We can apply the preceding analysis to the case where $K$ is a grid cell of arbitrary size $c \neq 1$. To do so, we simply scale $d$-space by the factor $1 / c$, so $K$ becomes a unit cube and $B_{r}$ becomes the ball $B_{r / c}$. We then scale space back, and obtain the bound

$$
\operatorname{Vol}\left(K_{r}\right) \leq \begin{cases}\beta c^{d} e^{\frac{3}{2}(2 \pi)^{1 / 3}(d / r c)^{2 / 3}} & \text { for } r \leq \sqrt{\frac{c d}{2 \pi}}  \tag{14}\\ \left.c^{d} \frac{(2 \pi)^{1 / 2} e^{3 / 2} r}{c d^{1 / 2}}\right)^{d} & \text { for } r>\sqrt{\frac{c d}{2 \pi}}\end{cases}
$$

The preceding remark now observes that $\operatorname{Vol}\left(K_{r}\right) / \operatorname{Vol}(K)$ is sub-exponential as long as $r \ll c \sqrt{d}$. This will become significant in the next section, where we choose $c=\Theta(1 / \sqrt{d})$.

An improved lower bound for $\gamma(d)$. We next improve our second lower bound, using the bound on $\operatorname{Vol}\left(K_{r}\right)$ just derived (in (14)). We now draw a ball of radius $1 / 2$ whose center is picked uniformly at random in $K_{1 / 2}$, rather than in the enlarged cube. A probabilistic argument similar to the one given above shows that there exists such a ball that contains at least $\frac{V_{d}\left(\frac{1}{2}\right)^{d}}{\operatorname{Vol}\left(K_{1 / 2}\right)} n$ points of $P$. We remove these points and repeat the drawing process. The
same argument as before implies that at least

$$
\frac{n^{2}}{8} \cdot \frac{V_{d}\left(\frac{1}{2}\right)^{d}}{\operatorname{Vol}\left(K_{1 / 2}\right)}-\frac{n}{2} \geq \beta^{\prime} n^{2} \frac{(\sqrt{\pi e / 2})^{d}}{\sqrt{d}^{d} c^{d} e^{\frac{3}{2}(\pi / 2)^{1 / 3}(d / c)^{2 / 3}}}-\frac{n}{2}
$$

pairs of points are at Euclidean distance at most 1 , where $\beta^{\prime}$ is yet another constant that depends polynomially on $d$, and where we have estimated $\operatorname{Vol}\left(K_{1 / 2}\right)$ using (14) (here $r<$ $c \sqrt{\frac{d}{2 \pi}}$ unless $d$ is very small), and have approximated $V_{d}$ using Stirling's formula as before. We summarize this finding in the following lemma.

Lemma 3.2 Let $P$ be a set of $n$ points in a cube of size $c>0$ in $\mathbb{R}^{d}$. Then at least

$$
\beta^{\prime} n^{2} \frac{(\sqrt{\pi e / 2})^{d}}{\sqrt{d}^{d} c^{d} e^{\frac{3}{2}(\pi / 2)^{1 / 3}(d / c)^{2 / 3}}}-\frac{n}{2}
$$

pairs of points of $P$ are at Euclidean distance at most 1, where $\beta^{\prime}$ is a constant that depends polynomially on $d$.

In other words, the constant fraction multiplying $n^{2}$ that we get here is larger than the one we got before roughly by the factor

$$
\left(\frac{c+1}{c}\right)^{d} \cdot \frac{1}{e^{\frac{3}{2}(\pi / 2)^{1 / 3}(d / c)^{2 / 3}}} \approx e^{d / c-\frac{3}{2}(\pi / 2)^{1 / 3}(d / c)^{2 / 3}}
$$

which is a significant improvement as long as $c \ll d$.
Combining the bound in Lemma 3.2 with the one in Section 2, in which we approximate $s_{d}(c)$ by $e^{-\sqrt{d} / c}$, we obtain the following summary result.

Theorem 3.3 Given a set $P$ of $n$ points in $\mathbb{R}^{d}$ and a parameter $c \geq 1+2 e^{-\sqrt{d / 2}}$, we can report, with high probability, all pairs of points of $P$ at Euclidean distance at most 1 in time $C(d)(n+k) \log n$, where

$$
C(d)=\tilde{O}\left(e^{\sqrt{d} / c} \frac{(\sqrt{d})^{d} c^{d} e^{\frac{3}{2}(\pi / 2)^{1 / 3}(d / c)^{2 / 3}}}{(\sqrt{\pi e / 2})^{d}}\right)
$$

Here $k$ is the output size. Alternatively, we can report, in expectation, an arbitrarily large fraction of the pairs of points of $P$ at Euclidean distance at most 1 in time $C(d)(n+k)$, with the same asymptotic bound on $C(d)$. The probability of reporting all pairs in the first variant, and the expected fraction of reported pairs in the second variant, can be made arbitrarily close to 1 by increasing $C(d)$ by a sufficiently large absolute constant factor.

Remark. One might attempt to minimize the value of $C(d)$ as a function of $c$. Straightforward calculation shows that $C(d)$ is minimized at $c=c_{0}=z^{3} / \sqrt{d} \approx 2.6 / \sqrt{d}$, where $z \approx 1.375$ is the positive root of

$$
z^{3}-\left(\frac{\pi}{2}\right)^{1 / 3} z-1=0
$$

Of course, we cannot choose this value of $c$, which is much smaller than 1 (for $d \geq 7$ ). Since $C(d)$ increases for $c>c_{0}$, we conclude that our best choice for $c$ is the smallest possible value, namely, $c=1+2 e^{-\sqrt{d / 2}}$. This implies, as is easily checked, the following corollary.

Corollary 3.4 With the above choice of c, the running time of the algorithm is

$$
\tilde{O}\left(e^{\sqrt{d}} \frac{(\sqrt{d})^{d} e^{\frac{3}{2}(\pi / 2)^{1 / 3} d^{2 / 3}}}{(\sqrt{\pi e / 2})^{d}}\right)(n+k) \log n
$$

for reporting, with high probability, all pairs at Euclidean distance at most 1, and it is

$$
\tilde{O}\left(e^{\sqrt{d}} \frac{(\sqrt{d})^{d} e^{\frac{3}{2}(\pi / 2)^{1 / 3} d^{2 / 3}}}{(\sqrt{\pi e / 2})^{d}}\right)(n+k)
$$

for reporting, in expectation, an arbitrarily large fraction of the pairs; $k$ is the output size.
Remark. Ignoring subexponential factors, the preceding bound is roughly $\sqrt{d}^{d} / \sqrt{\pi e / 2}^{d}$, which is roughly $2^{d}$ times the lower bound.

### 3.2 Packing lemma

As another interesting application of Lemma 3.1, we derive in this subsection the following so-called packing lemma. It is reminiscent of a similar packing lemma used in Lenhof and Smid [34], and it can also be regarded as a more natural extension of Lemma 3.2, in which the points were confined to lie in a fixed cube.

Lemma 3.5 (Packing Lemma) Let $P$ be a set of $n$ points in $\mathbb{R}^{d}$, and let $r>0$ be a parameter. Let $N_{\infty}(P)$ (resp., $N_{2}(P)$ ) be the number of pairs of points of $P$ at $L_{\infty}$-distance (resp., $L_{2}$-distance) at most $r$. Then we have

$$
N_{\infty}(P) \leq B_{d}\left(N_{2}(P)+n\right)
$$

where

$$
B_{d}=\tilde{O}\left((4 / \sqrt{2 \pi e})^{d} d^{d / 2} e^{\frac{3}{4}(9 \pi)^{1 / 3} d^{2 / 3}}\right)
$$

Proof. As done so far in this paper, we may assume, without loss of generality, that $r=1$. Cover $\mathbb{R}^{d}$ by balls of diameter 1. An old result of Rogers (see, e.g., [14, Section 8.2]) asserts that this can be done (by a lattice covering) so that each point lies in at most $\lambda=O(d \log d)$ balls. Let $\mathcal{B}$ denote the (necessarily finite) sub-collection of the cover consisting of those balls containing points of $P$. For each ball $B \in \mathcal{B}$, let $n_{B}$ denote the number of points of $P \cap B$. We clearly have $\sum_{B \in \mathcal{B}} n_{B} \leq \lambda n$ and $\sum_{B \in \mathcal{B}}\binom{n_{B}}{2} \leq \lambda N_{2}(P)$.

Let $K$ denote the cube $[-1,1]^{d}$, and let $B(r)$ denote the ball of radius $r$ centered at the origin, for any $r>0$. Let $p$ and $q$ be a pair of points of $P$ with $\|p-q\|_{\infty} \leq 1$. Let $B_{p}$ and $B_{q}$ be balls of $\mathcal{B}$ containing $p$ and $q$, respectively, and assume, without loss of generality, that $n_{B_{q}} \leq n_{B_{p}}$. Since $q \in p+K$, it follows that $B_{q} \subset B_{p} \oplus K \oplus B(1)$. Since (i) the Minkowski sum is associative, (ii) $B_{p}$ is congruent to $B(1 / 2)$, and (iii) $B(1 / 2) \oplus B(1)=B(3 / 2)$, an
application of Lemma 3.1, or rather of the bound in (14), with $c=2$ and $r=3 / 2$ (again, unless $r$ is very small, we are in the first case of (14)), yields

$$
\operatorname{Vol}\left(B_{p} \oplus K \oplus B(1)\right) \leq \beta 2^{d} e^{\frac{3}{4}(9 \pi)^{1 / 3} d^{2 / 3}}
$$

The number $M$ of balls of $\mathcal{B}$ that are fully contained in $B_{p} \oplus K \oplus B(1)$ satisfies

$$
\begin{aligned}
M & \leq \frac{\lambda \operatorname{Vol}\left(B_{p} \oplus K \oplus B(1)\right)}{V_{d}(1 / 2)^{d}} \\
& \leq \frac{\lambda \beta 4^{d} e^{\frac{3}{4}(9 \pi)^{1 / 3} d^{2 / 3}} \Gamma(1+d / 2)}{\pi^{d / 2}} \\
& =\tilde{O}\left((4 / \sqrt{\pi})^{d} e^{\frac{3}{4}(9 \pi)^{1 / 3} d^{2 / 3}}(d /(2 e))^{d / 2}\right) \\
& =\tilde{O}\left((4 / \sqrt{2 \pi e})^{d} d^{d / 2} e^{\frac{3}{4}(9 \pi)^{1 / 3} d^{2 / 3}}\right),
\end{aligned}
$$

where the final constant of proportionality depends polynomially on $d$.
We charge the pair $(p, q)$ to $B_{p}$, and note that every ball $B \in \mathcal{B}$ is charged by at most

$$
M n_{B}^{2}=2 M\binom{n_{B}}{2}+M n_{B}
$$

such pairs (recall our assumption that $n_{B_{q}} \leq n_{B_{p}}$ ). Summing over all balls of $\mathcal{B}$, and using the inequalities observed at the beginning of the proof, we obtain

$$
N_{\infty}(P) \leq M \sum_{B \in \mathcal{B}} n_{B}^{2}=2 M \sum_{B \in \mathcal{B}}\binom{n_{B}}{2}+M \sum_{B \in \mathcal{B}} n_{B} \leq 2 M \lambda N_{2}(P)+M \lambda n
$$

and the lemma follows.
Remark. It is perhaps simpler to absorb the last sub-exponential factor of $M$ in its first exponential factor, by slightly increasing its base. Also, this base is smaller than 1 , so a simpler, slightly weaker form of the lemma is that $N_{\infty}(P) \leq \sqrt{d}^{d}\left(N_{2}(P)+n\right)$.

## 4 Low doubling dimension

As noted in the introduction, in many applications the input set $P$ has a restricted structure, in the sense that its points have much fewer "degrees of freedom" than the ambient dimension $d$. For example, all the points of $P$ might lie on, or near, some manifold of much smaller dimension. In these cases one hopes for a better performance of the algorithm, in the sense that the constant $C(d)$ in the bounds in Theorem 3.3 can be replaced by a much smaller one. We show in this section that this is indeed the case.

One formal way of capturing such a special structure of $P$ is through the notion of doubling dimension; see $[28,33]$. Specifically, a metric space $(X, \rho)$ has doubling dimension $\delta$ if every ball in $X$ of any radius $r$ can be covered by at most $2^{\delta}$ balls of radius $r / 2$. In what follows we assume that $P$, under the Euclidean metric in $d$-space, has doubling dimension $\delta \ll d$.

The mechanism that has led to Theorem 3.3 can easily be adapted to this case. Specifically, we have the following result.

Theorem 4.1 Given a set $P$ of $n$ points in $\mathbb{R}^{d}$ of doubling dimension $\delta$ (with respect to the Euclidean distance), and a parameter $c \geq 1+2 e^{-\sqrt{d / 2}}$, we can report, with high probability, all pairs of points of $P$ at Euclidean distance at most 1 in time $C(d, \delta)(n+k) \log n$, where

$$
C(d, \delta)=\tilde{O}\left(e^{\sqrt{d} / c}(2 c \sqrt{d})^{\delta}\right)
$$

and where $k$ is the output size. Alternatively, we can report, in expectation, an arbitrarily large fraction of the pairs of points of $P$ at Euclidean distance at most 1 in time $C(d, \delta)(n+$ $k$ ), with the same asymptotic bound on $C(d)$. The probability of reporting all pairs in the first variant, and the expected fraction of reported pairs in the second variant, can be made arbitrarily close to 1 by increasing $C(d, \delta)$ by a sufficiently large absolute constant factor.

Remark. Simple calculation shows that the bound on $C(d, \delta)$ is minimized at $c=\sqrt{d} / \delta$, and then $C(d, \delta)=O\left((2 e d / \delta)^{\delta}\right)$. Of course, this choice of $c$ makes sense only when $\delta$ is smaller than $\sqrt{d}$; otherwise we choose $c=1+2 e^{-\sqrt{d / 2}}$, as in Corollary 3.4, and get $C(d, \delta)=\tilde{O}\left(e^{\sqrt{d}}(2 \sqrt{d})^{\delta}\right)$.
Proof. We use a randomly shifted grid of cell size $c$, as in Section 2, and conclude that any fixed unit vector $\mathbf{x}$ is captured in a grid cell with probability at least $s_{d}(c) \approx e^{-\sqrt{d} / c}$. (This part of the analysis does not seem to benefit from the fact that $P$ has low doubling dimension, but at any rate the dependence of $s_{d}(c)$ on $d$ is "tame", that is, sub-exponential.)

What can be improved is the analysis that leads to Lemma 3.2. Specifically, fix a grid cell $\tau$, containing $n_{\tau}$ points of $P$. Observe that $\tau$ is fully contained in a ball of radius $c \sqrt{d} / 2$, and in a ball of radius $c \sqrt{d}$ centered at any point of $P_{\tau}=P \cap \tau$. It easily follows from the definition of doubling dimension that we can cover $P_{\tau}$ by $(c \sqrt{d})^{\delta}$ balls of radius $1 / 2$. Within each such ball, every pair of points is at Euclidean distance at most 1.

Let $\mathcal{B}$ be a minimum-size cover of $P_{\tau}$ by balls of radius $1 / 2$. As just shown, we have $M=|\mathcal{B}| \leq(c \sqrt{d})^{\delta}$. We claim that no point of $P_{\tau}$ is contained in more than $2^{\delta}$ balls of $\mathcal{B}$. Indeed, suppose to the contrary that there exists a point $p \in P_{\tau}$ that is contained in more than $2^{\delta}$ balls of $\mathcal{B}$. Then all these balls are fully contained in the ball $B_{1}(p)$ of radius 1 centered at $p$. Cover $P \cap B_{1}(p)$ by at most $2^{\delta}$ balls of radius $1 / 2$, and replace the balls of $\mathcal{B}$ that are contained in $B_{1}(p)$ by these $2^{\delta}$ balls. This gives us a cover of $P_{\tau}$ by fewer balls, a contradiction that establishes the claim.

It follows that no pair of points of $P_{\tau}$ can lie together in more than $2^{\delta}$ balls of $\mathcal{B}$, so we have

$$
\sum_{B \in \mathcal{B}}\binom{\left|P_{\tau} \cap B\right|}{2} \leq 2^{\delta} N_{2}\left(P_{\tau}\right)
$$

where $N_{2}\left(P_{\tau}\right)$ is the number of pairs of points of $P_{\tau}$ at Euclidean distance at most 1. On the other hand, using the Cauchy-Schwarz inequality, as in the derivation of Equation (9), we have

$$
\sum_{B \in \mathcal{B}}\binom{\left|P_{\tau} \cap B\right|}{2} \geq \frac{1}{2 M}\left(\sum_{B}\left|P_{\tau} \cap B\right|\right)^{2}-\frac{1}{2}\left(\sum_{B}\left|P_{\tau} \cap B\right|\right) \geq \frac{\left|P_{\tau}\right|^{2}}{2 M}-2^{\delta-1}\left|P_{\tau}\right|
$$

Combining the two inequalities, we obtain

Lemma 4.2 Let $P$ be a set of $n$ points in a cube of size $c>1 / \sqrt{d}$ in $\mathbb{R}^{d}$, of doubling dimension $\delta$ (with respect to the Euclidean distance). Then at least

$$
\frac{n^{2}}{2(2 c \sqrt{d})^{\delta}}-\frac{n}{2}
$$

pairs of points of $P$ are at Euclidean distance $\leq 1$.

As is easily checked, this completes the proof of Theorem 4.1.
Remark. We note that Datar et al. [19, Appendix A] present an algorithm (also based on gridding, though with non-orthogonal directions) for near neighbor search that also achieves improved performance when the dataset has "low dimension" (although they use a slightly stronger notion of dimension, called the Karger-Ruhl dimension).

## 5 Randomly rotated and shifted grids

Clearly, in the preceding analysis we cannot choose $c \leq 1$ because some unit vectors, such as axis-parallel unit vectors, will have zero probability of being captured in a grid cell. However, as we show in this section, if we also apply a random rotation, we will be able to capture any vector of length (at most) 1 within a grid cell of size $c \leq 1$ (of a randomly shifted grid), with some fixed positive probability (depending, as usual, on $c$ and $d$ ), provided that $c$ is not too small. Although this probability will be significantly smaller than the corresponding probability obtained in the preceding sections, the small value of $c$ will lead to much fewer repetitions of the procedure and consequently to a significant improvement in the overall running time of the algorithm.

Concretely, the main result of this section is the following.

Theorem 5.1 Given a set $P$ of $n$ points in $\mathbb{R}^{d}$, we can report, with high probability, all pairs of points of $P$ at Euclidean distance at most 1, in time $O\left(6.74^{d}(n+k) \log n\right)$, where $k$ is the output size. Alternatively, we can report, in expectation, an arbitrarily large fraction of these pairs in time $O\left(6.74^{d}(n+k)\right)$; the success probability and the expected fraction of reported pairs can be made arbitrarily close to 1 by increasing the constant factor hidden by the "big- $O$ ".

Proof. Fix some vector $\mathbf{x}_{0}$ of length 1 (as before, this is clearly the worst-case assumption; vectors of smaller length will be captured with higher probability). The probabilistic model assumed in this section is that we choose a random rotation $\rho$ of $d$-space (there are several ways of doing this; see, e.g., Genz [23] for an $O\left(d^{2} \log d\right)$ procedure) followed by placing a randomly shifted grid of cell size $c$ in the rotated space. We want to derive a lower bound on the probability of the event $A\left(\mathbf{x}_{0}\right)$ that both endpoints of the rotated image $\rho\left(\mathbf{x}_{0}\right)$ of $\mathbf{x}_{0}$ fall into the same grid cell.

The random rotation $\rho$ maps $\mathbf{x}_{0}$ to a random vector $\mathbf{x}=\rho\left(\mathbf{x}_{0}\right)$ (uniformly) on $\sigma$. The probability of capturing $\mathbf{x}=\left(x_{1}, \ldots, x_{d}\right)$ in a cell of a randomly shifted grid of cell size $c \leq 1$ is

$$
F_{c}(\mathbf{x})=\left(1-\frac{\left|x_{1}\right|}{c}\right)^{+} \cdots\left(1-\frac{\left|x_{d}\right|}{c}\right)^{+}
$$

where $w^{+}=\max \{w, 0\}$, for $w \in \mathbb{R}$. This is the probability of the event $A\left(\mathbf{x}_{0}\right)$ conditioned on the choice of the rotation $\rho$, so the actual probability of $A\left(\mathbf{x}_{0}\right)$ is the expectation (average) of $F_{c}\left(\rho\left(\mathbf{x}_{0}\right)\right)$ over the space of rotations.

The symmetry of the space of rotations is easily seen to imply that the expected value $\mu$ of $F_{c}\left(\rho\left(\mathbf{x}_{0}\right)\right)$, over the space of rotations, is the same as the expected value of $F_{c}(\mathbf{x})$ over $\sigma$ (because a random rotation takes $\mathbf{x}_{0}$ to a random point on $\sigma$, and each point has the same (differential) probability of being reached). That is,

$$
\mu=\frac{1}{|\sigma|} \int_{\sigma} F_{c}(\mathbf{x}) d \mathbf{x}
$$

where $|\sigma|$ is the surface volume of $\sigma$, which is (see [40])

$$
|\sigma|=\frac{2 \pi^{d / 2}}{\Gamma(d / 2)}
$$

To estimate $\mu$, we use the easily established fact, observed by Muller [37], that if we draw $d$ independent normal variables $x_{i}, 1 \leq i \leq d$, from the standard normal distribution $\mathbf{N}(0 ; 1)$ and put $\mathbf{x}=\left(x_{1}, \ldots, x_{d}\right)$, then $\mathbf{x} /\|\mathbf{x}\|$ is a uniform random direction (a random point on $\sigma)$. It follows from this observation that

$$
\begin{equation*}
\mu=\frac{1}{(2 \pi)^{d / 2}} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty}\left(1-\frac{\left|x_{1}\right|}{c\|\mathbf{x}\|}\right)^{+} \cdots\left(1-\frac{\left|x_{d}\right|}{c\|\mathbf{x}\|}\right)^{+} e^{-\|\mathbf{x}\|^{2} / 2} d x_{1} \cdots d x_{d} \tag{15}
\end{equation*}
$$

To obtain a lower bound on $\mu$ we fix some threshold radius $r_{0}$, and write $\mu=\mu^{-}+\mu^{+}$, where $\mu^{-}$(resp., $\mu^{+}$) is the portion of $\mu$ obtained by integrating over the multivariate normal distribution within the region $\|\mathbf{x}\| \leq r_{0}$ (resp., $\|\mathbf{x}\| \geq r_{0}$ ). Clearly $\mu \geq \mu^{+}$so we focus on lower-bounding $\mu^{+}$.

For $\|\mathbf{x}\| \geq r_{0}$ we have

$$
\left(1-\frac{\left|x_{1}\right|}{c\|\mathbf{x}\|}\right)^{+} \cdots\left(1-\frac{\left|x_{d}\right|}{c\|\mathbf{x}\|}\right)^{+} \geq\left(1-\frac{\left|x_{1}\right|}{c r_{0}}\right)^{+} \cdots\left(1-\frac{\left|x_{d}\right|}{c r_{0}}\right)^{+}
$$

so

$$
\mu^{+} \geq \frac{1}{(2 \pi)^{d / 2}} \int_{\|\mathbf{x}\| \geq r_{0}}\left(1-\frac{\left|x_{1}\right|}{c r_{0}}\right)^{+} \cdots\left(1-\frac{\left|x_{d}\right|}{c r_{0}}\right)^{+} e^{-\|\mathbf{x}\|^{2} / 2} d x_{1} \cdots d x_{d}
$$

Write this right-hand integral as $I-I_{0}$, where $I$ (resp., $I_{0}$ ) is obtained by integrating over the entire $\mathbb{R}^{d}$ (resp., over $\|\mathbf{x}\| \leq r_{0}$ ).
$I$ is easy to compute (here we use the independence of the choice of coordinates $x_{i}$ ), and we get

$$
\begin{aligned}
I & =\frac{1}{(2 \pi)^{d / 2}}\left(\int_{-\infty}^{\infty}\left(1-\frac{|x|}{c r_{0}}\right)^{+} e^{-x^{2} / 2} d x\right)^{d} \\
& =\frac{1}{(2 \pi)^{d / 2}}\left(2 \int_{0}^{c r_{0}}\left(1-\frac{x}{c r_{0}}\right) e^{-x^{2} / 2} d x\right)^{d} \\
& =\left(\frac{2}{\pi}\right)^{d / 2}\left(\int_{0}^{c r_{0}} e^{-x^{2} / 2} d x-\frac{1}{c r_{0}} \int_{0}^{c r_{0}} x e^{-x^{2} / 2} d x\right)^{d} \\
& =\left(\frac{2}{\pi}\right)^{d / 2}\left(\int_{0}^{c r_{0}} e^{-x^{2} / 2} d x-\frac{1-e^{-\left(c r_{0}\right)^{2} / 2}}{c r_{0}}\right)^{d}
\end{aligned}
$$

Using integration by parts, it is easily checked that

$$
\frac{e^{-\left(c r_{0}\right)^{2} / 2}}{c r_{0}}=\int_{c r_{0}}^{\infty} e^{-x^{2} / 2} d x+\int_{c r_{0}}^{\infty} \frac{1}{x^{2}} e^{-x^{2} / 2} d x
$$

which implies that

$$
\begin{aligned}
& \int_{0}^{c r_{0}} e^{-x^{2} / 2} d x+\frac{e^{-\left(c r_{0}\right)^{2} / 2}}{c r_{0}}= \\
& \int_{0}^{c r_{0}} e^{-x^{2} / 2} d x+\int_{c r_{0}}^{\infty} e^{-x^{2} / 2} d x+\int_{c r_{0}}^{\infty} \frac{1}{x^{2}} e^{-x^{2} / 2} d x= \\
& \int_{0}^{\infty} e^{-x^{2} / 2} d x+\int_{c r_{0}}^{\infty} \frac{1}{x^{2}} e^{-x^{2} / 2} d x \geq \sqrt{\pi / 2}
\end{aligned}
$$

where in the final inequality we neglect the second integral. Hence

$$
I \geq\left(\frac{2}{\pi}\right)^{d / 2}\left(\sqrt{\pi / 2}-\frac{1}{c r_{0}}\right)^{d}=\left(1-\frac{\sqrt{2}}{\sqrt{\pi} c r_{0}}\right)^{d}
$$

We obtain an upper bound for $I_{0}$ by upper bounding $\left(1-\frac{\left|x_{1}\right|}{c r_{0}}\right)^{+} \cdots\left(1-\frac{\left|x_{d}\right|}{c r_{0}}\right)^{+}$by 1 , so we have

$$
I_{0} \leq \frac{1}{(2 \pi)^{d / 2}} \int_{\|\mathbf{x}\| \leq r_{0}} e^{-\|\mathbf{x}\|^{2} / 2} d x_{1} \cdots d x_{d} .
$$

We express this integral in spherical coordinates $(r, \theta)$, where $r$ is the radius (distance from the origin) and $\theta \in \sigma$ is the orientation. The spherical symmetry of the multivariate normal distribution is easily seen to yield

$$
\begin{equation*}
I_{0} \leq \frac{1}{(2 \pi)^{d / 2}} \int_{\sigma} 1 d \theta \int_{0}^{r_{0}} r^{d-1} e^{-r^{2} / 2} d r \tag{16}
\end{equation*}
$$

where in the first integral $\theta$ denotes a point on $\sigma$ and $d \theta$ denotes a surface volume element on $\sigma$. Informally, the identity follows since a volume element at a point $\mathbf{x}=(r, \theta)$ can be written as $r^{d-1} d \theta$ (the surface volume element on the sphere through $\mathbf{x}$, of radius $r$ ) times $d r$.

We note that $\int_{\sigma} 1 d \theta$ is simply the surface volume of $\sigma$, which is $|\sigma|=\frac{2 \pi^{d / 2}}{\Gamma(d / 2)}$, and we thus get

$$
I_{0} \leq \frac{1}{(2 \pi)^{d / 2}} \cdot \frac{2 \pi^{d / 2}}{\Gamma(d / 2)} \int_{0}^{r_{0}} r^{d-1} e^{-r^{2} / 2} d r=\frac{1}{2^{d / 2-1} \Gamma(d / 2)} \int_{0}^{r_{0}} r^{d-1} e^{-r^{2} / 2} d r
$$

Replacing $r$ by $\sqrt{2 z}$, putting $r_{0}=\sqrt{t d}$, for some $0<t$, and writing $d / 2=k$, we get

$$
I_{0} \leq \frac{1}{\Gamma(k)} \int_{0}^{t k} z^{k-1} e^{-z} d z
$$

The function $z^{k-1} e^{-z}$ is maximized at $z=k-1$. If we assume that $t \leq 1-2 / d$ then $k-1 \geq t k$, so the integrand is an increasing function in the range $0 \leq z \leq t k$, and we get

$$
\int_{0}^{t k} z^{k-1} e^{-z} d z \leq t k \cdot(t k)^{k-1} e^{-t k}=(t k)^{k} e^{-t k}
$$

We can also lower bound $\Gamma(k)$, using Stirling's formula, by $\sqrt{2 \pi / k}(k / e)^{k}$. Hence we get

$$
\begin{aligned}
I_{0} & \leq(k /(2 \pi))^{1 / 2}(e / k)^{k}(t k)^{k} e^{-t k} \\
& =(k /(2 \pi))^{1 / 2}(e t)^{k} e^{-t k}=(k /(2 \pi))^{1 / 2}\left(t e^{1-t}\right)^{k} .
\end{aligned}
$$

Putting

$$
\beta(t):=t^{1 / 2} e^{(1-t) / 2}
$$

we conclude that

$$
I_{0} \leq(k /(2 \pi))^{1 / 2} \beta(t)^{d} .=(d /(4 \pi))^{1 / 2} \beta(t)^{d} .
$$

Note that $\beta(t)<1$ for any $t<1$.
Putting the lower bound on $I$ and the upper bound on $I_{0}$ together, we get

$$
\mu \geq \mu^{+} \geq I-I_{0} \geq\left(1-\frac{\sqrt{2}}{\sqrt{\pi} c r_{0}}\right)^{d}-(d /(4 \pi))^{1 / 2} \beta(t)^{d}
$$

Choosing $c$. Note that we can, and indeed intend to, take $c$ to be much smaller than 1. The only constraint that we need to enforce is that

$$
1-\frac{\sqrt{2}}{\sqrt{\pi} c r_{0}}>\left(\frac{d}{4 \pi}\right)^{1 /(2 d)} \beta(t)
$$

The maximum value of $\left(\frac{d}{4 \pi}\right)^{1 /(2 d)}$ is smaller than 1.015 , as is easily checked, so we replace this constraint by

$$
1-\frac{\sqrt{2}}{\sqrt{\pi} c r_{0}}>1.015 \cdot \beta(t)
$$

Recalling the choice of $r_{0}=\sqrt{d t}$ and the definition of $\beta(t)$, this becomes

$$
c>\frac{a(t)}{\sqrt{d}}
$$

where

$$
\begin{equation*}
a(t)=\frac{\sqrt{2}}{\sqrt{\pi t}\left(1-1.015 t^{1 / 2} e^{(1-t) / 2}\right)}>\frac{\sqrt{2}}{\sqrt{\pi t}} . \tag{17}
\end{equation*}
$$

That is, we can take $t$ to be some constant fraction (the optimal choice of $t$ will be discussed later in the section) and then take $c$ to be slightly larger than $\frac{a(t)}{\sqrt{d}}$; that is, we take $c$ to be some constant multiple of $1 / \sqrt{d}$. In this case the (lower bound on the) expected success probability $\mu$ will decrease exponentially with $d$. In contrast, choosing $c$ to be some fixed absolute fraction, similar to the choice in the purely translational case treated in Sections 2 and 3 , yields a lower bound for $\mu$ of the form $e^{-\alpha \sqrt{d} / c}$, for $\alpha \approx \sqrt{\frac{2}{\pi t}}$ (in this case any choice of $t$ not too close to 1 will do).

The rationale for taking $c$ so small is that this (significant) degradation in the success probability will be more than offset by the ratio (arising in the analysis of our algorithm) between the number of pairs of points in a grid cell and the number of such pairs at Euclidean
distance at most 1 , making the combination of these two factors only exponential in $d$, rather than super-exponential, as was the case in Section 3.

Let $a(t)$ be as defined in (17), and let us choose $c=a / \sqrt{d}$, for some slightly larger $a>a(t)$. As can be easily checked, the smallest value for $a(t)$, attained at $t \approx 0.11$, is about 5.07. (Observe that we cannot choose $c$ to be smaller than $1 / \sqrt{d}$, because then no unit vector will fit into a cell of size $c$.) The best choice of the constant $a$ is different though, and is discussed next.

For $c=a / \sqrt{d}, \mu$ is at least roughly $\left(1-\frac{\sqrt{2}}{a \sqrt{\pi t}}\right)^{d}$. That is, for any fixed unit vector $\mathbf{x}_{0}$, a random rotation followed by a placement of a randomly shifted grid of cell size $c$, for $c=a / \sqrt{d}$ as above, captures both endpoints of $\mathbf{x}_{0}$ in the same grid cell with probability at least $\left(1-\frac{\sqrt{2}}{a \sqrt{\pi t}}\right)^{d}$, so the expected number of trials until a successful one (for the fixed vectors $\left.\mathbf{x}_{0}\right)$ is at most $M_{0}=1 /\left(1-\frac{\sqrt{2}}{a \sqrt{\pi t}}\right)^{d}$. More precisely, repeating this step $b M_{0} \log n$ times ensures with high probability that every pair at Euclidean distance at most 1 is captured in a grid cell. Repeating this step only $b M_{0}$ times ensures that an arbitrarily large expected fraction of these pairs (a fraction approaching 1 when $b$ increases) will be captured.

On the other hand, assume that a grid cell $\tau$ of size $c=a / \sqrt{d}$ contains $n_{\tau}$ points of $P$. We use Lemma 3.2, and note that, since $a$ is not too close to 1 , we are again at the first case of (14). We conclude that the number of pairs of points of $P \cap \tau$ at Euclidean distance at most 1 is at least

$$
\beta^{\prime} n_{\tau}^{2} \frac{(\sqrt{\pi e / 2})^{d}}{\sqrt{d}^{d} c^{d} e^{\frac{3}{2}(\pi / 2)^{1 / 3}(d / c)^{2 / 3}}}-\frac{n_{\tau}}{2}
$$

Substituting $c=a / \sqrt{d}$, this becomes

$$
\beta^{\prime}\left(\frac{\sqrt{\pi e / 2}}{a e^{\frac{3}{2}\left(\pi /\left(2 a^{2}\right)\right)^{1 / 3}}}\right)^{d} n_{\tau}^{2}-\frac{n_{\tau}}{2}
$$

It follows that, with high probability, we report each of the $k$ pairs of points at Euclidean distance at most 1 (resp., in expectation, an arbitrarily large fraction of these pairs) by inspecting no more than $C(d, a) \cdot(n+k) \log n$ (resp., $C(d, a) \cdot(n+k))$ pairs of points of $P$, where $C(d, a)$ is

$$
\begin{aligned}
C(d, a) & =\tilde{O}\left(\frac{1}{\left(1-\frac{\sqrt{2}}{a \sqrt{\pi t}}\right)^{d}} \cdot\left(\frac{a e^{\frac{3}{2}\left(\pi /\left(2 a^{2}\right)\right)^{1 / 3}}}{\sqrt{\pi e / 2}}\right)^{d}\right) \\
& =\tilde{O}\left(\left(\frac{a^{2} e^{\frac{3}{2}\left(\pi /\left(2 a^{2}\right)\right)^{1 / 3}}}{a \sqrt{\pi e / 2}-\sqrt{e / t}}\right)^{d}\right)
\end{aligned}
$$

Increasing $C(d, a)$ by a constant factor increases the overall success probability that we indeed report all pairs at distance 1 (or the constant fraction of the pairs reported).

We thus get rid of the super-exponential factor $d^{d / 2}$ that we encounter when $c$ is a fixed fraction or when no rotation is applied, and obtain an overall exponential factor.

Recall that $a=a(t)$ is a function of $t$ given by Equation (17). So the best lower bound on the base in the "overhead" constant factor $C(d, a)$ that this method yields is about

$$
\min _{t}\left\{\frac{a(t)^{2} e^{\frac{3}{2}\left(\pi /\left(2 a(t)^{2}\right)\right)^{1 / 3}}}{a(t) \sqrt{\pi e / 2}-\sqrt{e / t}}\right\}
$$

and numerical calculations show that this minimum value is about 6.74 , obtained for $t \approx$ 0.25 .

This completes the proof of Theorem 5.1.
Remark. The case where $P$ has low doubling dimension $\delta \ll d$ does not seem to fit the framework of this section because, as in the preceding sections, we do not know how to make the success probability of capturing a unit vector within a grid cell depend also on $\delta$. While this probability was only sub-exponentially small (in $d$ ) in the case of pure translation, here, with the very small value of $c$ that we choose, this probability becomes exponentially small. This in turn makes the constant in the bound depend (at least) exponentially on $d$, a much larger value than that obtained in Section 4.

## 6 Experimental results

We have implemented both algorithms in C++ and have compared them to several leading software packages for reporting neighbors. These do not include algorithms that lay down a (fixed) grid of cell size $r$ and test each point in each nonempty cell with the points in its neighboring nonempty cells, such as the one in [34]. These algorithms are very slow due to the large number of neighboring cells, each of which must be tested to determine whether it is empty or not.

A natural approach to solve our problem is by using certain "off the shelf" data structures for (approximate) nearest neighbor queries. We first construct such a data structure on our input set and then query it with each input point. Evidently in order to find all pairs of points at distance at most $r$ we need each query with a point $q$ to return all points at distance at most $r$ from $q$, rather than just the nearest neighbor of $q$. Fortunately, the available packages do provide functions that can be adapted for this task. Since these packages approximate the nearest neighbors they in fact only find "most" neighbors of distance at most $r$. By fine-tuning the parameters of the appropriate package, we can control the expected fraction of the pairs at distance at most $r$ that this package reports.

We have implemented such a scheme using three publicly available software packages: the E2LSH package for Locality Sensitive Hashing [3, 4], the ANN package by Arya et al. [9], and the FLANN package [35, 36]. ANN and FLANN use various hierarchical space decompositions (such as kd-trees and box decomposition trees).

The E2LSH package provides a function to learn the configuration parameters needed for a particular success probability. This function takes a data set to index, a set of queries, a target probability $p$, and a distance $r$. It then computes the parameters required to ensure that for each query $q$ we report every point $w$ at distance at most $r$ from $q$, with probability $p$. We applied this function with target probability of $p=0.9$ to a moderate-size data set (arising from 900 images, see below) which we used both as the data set to index, and as the
set of queries. Then in all our experiments we used the configuration parameters produced by this training phase.

The ANN and FLANN packages support a search for the $k$ nearest neighbors and allow to stop the search when the distance to the $j$ th neighbor (for some $j<k$ ) exceeds some prespecified value $r$. We used this search with a large $k$ and our target distance $r$. As an outcome for a query point $q$ and a parameter $\varepsilon$, we got all points $p$ at distance at most $r /(1+\varepsilon)$ from $q$, and some points $p$ at distance at most $r$ from $q$ (we never get points lying further away from $q$ ). We have set $\varepsilon=1.4$ since for this value the number of points reported by ANN and FLANN was comparable to the number of points reported by E2LSH (with $p=0.9$ ) on the training data set mentioned above. In our first algorithm, which uses a randomly shifted grid, we set the size of a grid cell to be $c=3.6 r$ and we used 5 different randomly shifted grids. For these parameters the total number of pairs that we report on the training data set was comparable to the number of pairs reported by the other packages with the parameters mentioned above.

Our points lie in Euclidean space of dimension $d=64$, but $c / r=3.6<\sqrt{d}=8$. This is exactly the situation mentioned in the introduction, where the easy lower bound of $1-r \sqrt{d} / c$ on the success probability is meaningless. Nevertheless our choice of $c$ is not as small as our analysis suggests. The reason for that is the small intrinsic ("doubling") dimension of our data sets (which is probably closer to 10 than to 64 ). This small dimension makes the number of points that we see in a grid cell relatively small and as a result moderate values of $c$ work better in practice. As already remarked in Section 2 , the assumption that $d>d_{0}$, as required by the theoretical analysis, is likely to hold here. Nevertheless, this is not a crucial issue in analyzing the practical results.

Note that by running the training phase of E2LSH, it is guaranteed to report $90 \%$ of the pairs of distance at most $r$ of each other. ANN and FLANN do not provide such a guarantee directly but they do so indirectly since we tuned their parameters to report about the same number of pairs as E2LSH. ANN and FLANN are guaranteed, however, to report all pairs at distance smaller than $r /(1+\varepsilon)$. Our theoretical analysis in Section 2 guarantees that a pair of distance $r$ falls into the same grid cell of one random grid with probability at least $s_{d}(c / r)=s_{d}(3.6) \approx 0.1$. So the probability that we discover a close pair using 5 grids is $1-0.9^{5}=0.41$. The fact that with 5 grids we get an output size comparable to that of E2LSH indicates that our lower bound, $s_{d}(c / r)$, on the probability of locating each pair is pessimistic, and that this probability is larger in practice.

We have constructed our input set by extracting representative descriptors in $\mathbb{R}^{64}$ from images, using the Speeded Up Robust Feature (SURF) method [11], which is used for computer vision problems such as object recognition and 3D reconstruction. SURF first identifies interesting points at significant corner locations of the image. Then it summarizes the region around each interesting point, using the Haar wavelet transform, into a descriptor of 64 numbers in $[0,1]$. The vectors of these descriptors form the points in $\mathbb{R}^{64}$ that we used as our input data set.

Finding all pairs of descriptors at distance at most $r$ from each other is a common method for identifying similar images. Images are declared to be similar if there are sufficiently many pairs of descriptors, one from each image, at distance at most $r$ from each other. We used $r=0.08$ which is appropriate for the underlying image matching application. We used a collection of images containing pairs that were taken at the same geographic location or of
the same object. Such pairs of images have many close descriptors. These pairs of close descriptors are the pairs of points in $\mathbb{R}^{64}$ that our algorithm tries to identify. (For training LSH we used a set of descriptors extracted from 900 images from the same set.)

The table in Figure 2 shows the results for four data sets obtained respectively from 50, 100,200 , and 400 images. The second column gives the total number of points (in $\mathbb{R}^{64}$ ) extracted from these images. This is about 260 points per image. The third column gives the exact number of pairs of points at distance at most $r$. The last four columns give the results for the four algorithms that we compared. We show two numbers in each entry of these columns: the number of reported pairs at distance at most $r$, and in parenthesis the running time in seconds. We can see from this table that the fraction of reported pairs in all instances is about $96 \%$ of the total number of pairs at distance at most $r$, except for FLANN, which failed to report a sufficiently large number of pairs for 400 images. Automatic parameter selection for FLANN gave bad results and we set manually the parameters to get the best performance. Specifically, we set the index parameter to 4 and search parameter to 64 . We can see that FLANN is considerably slower than all the other algorithms. Our algorithm is the fastest.

We note that FLANN is highly sensitive to the required accuracy. As can be seen in the table, it reports fewer pairs than the other algorithms, and this is more noticeable where the number of points increases. In a recent paper by the first author [2], for the slightly different problem of range searching with arbitrary query points, it was observed that when FLANN is allowed to report even fewer points it becomes faster than ANN.

| \#images | \#pts | \#r-close <br> pairs | ANN | FLANN | LSH | Ours |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 50 | 12981 | 1024 | 998 <br> $(0.31)$ | 1012 <br> $(1.19)$ | 1013 <br> $(0.46)$ | 1001 <br> $(0.07)$ |
| 100 | 25532 | 1784 | 1734 <br> $(0.64)$ | 1755 <br> $(2.3)$ | 1755 <br> $(0.89)$ | 1705 <br> $(0.16)$ |
| 200 | 53106 | 4628 | 4347 <br> $(1.61)$ | 4154 <br> $(5.12)$ | 4532 <br> $(1.95)$ | 4405 <br> $(0.46)$ |
| 400 | 106336 | 8538 | 8040 <br> $(4.46)$ | 7295 <br> $(10.4)$ | 8319 <br> $(4.16)$ | 8416 <br> $(1.17)$ |

Figure 2: Comparison of the numbers of reported points and the running times (in seconds) for the different algorithms.

To appreciate the asymptotic improvement of our algorithm, we also compared it to ANN and E2LSH on much larger data sets. FLANN is much slower and was not able to compete with the other algorithms on these instances. Figure 3 shows the running times of the three algorithms as a function of the size of the data set. The asymptotic supremacy of our method can clearly be observed. The reported number of pairs is shown in Figure 4.

### 6.1 Using rotations

The number of pairs of points in a grid cell of size $c$ typically increases rapidly with $c$, so, as discussed earlier, we want to use very small values of $c$ which, combined with a random


Figure 3: Running times on inputs of increasing sizes.


Figure 4: Number of reported neighboring pairs.
rotation of the coordinate frame, yields a procedure for which we were able to obtain a much improved theoretical upper pound on its expected performance.

On our real data sets, however, as it turns out, rotations do not yield an advantage. This is probably due to the fact that the data sets have low doubling dimension, and that the algorithm that uses rotation does not take advantage of this property, as does the purely translational algorithm. Decreasing the value of $c$ indeed decreases the number of pairs which we inspect. ${ }^{5}$ But the gain in running time due to this decrease is subsumed by the loss incurred by the larger number of experiments which we need to perform. Figure 5 shows the running times on our data set produced from 100 images, of the algorithms with rotations and without rotations for various values of $c$. In all experiments we reported about 1700 pairs. The best running time is obtained for $c=3.6$ without rotation. A peculiar phenomenon about this data set is that for higher values of $c$ the algorithm with rotation becomes faster. This phenomenon is somewhat surprising and we do not fully understand it. It may be related to a special initial alignment of this particular point set. (Points in this data set seem to tend to be near one of the main diagonals of the unit cube.) For $c=7$ without rotation we inspect 8528906 pairs out of which 1782 are at distance at most $r$. When we rotate the points we inspect only 1626949 pairs out of which 1713 are at distance at most $r$.


Figure 5: Running times with and without rotation on real data as a function of the cell size.

For other data sets containing points which are distributed differently, random rotations can give a substantial advantage. To demonstrate this we experimented with a random data set of 100000 points in the unit cube of dimension 16 (which do not have low doubling dimension). We set $r=0.5$ for this data set. There are 3600 pairs at distance at most $r$ in the data set of which we reported about 3500 in all our experiments. Here the number of

[^5]

Figure 6: Running times with and without rotation on random data as a function of the cell size.
inspected pairs increases more drastically with $c .^{6}$. As a result it pays off to use rotations with a smaller value of $c$ and a larger number of experiments. Figure 6 shows the running times for the algorithms with rotations and without rotations for different values of $c$. We see in this plot that the best running time is obtained for $c=1.3$ with rotations. Recall that this is obtained despite the overhead of $O\left(d^{2} n\right)$ time which we invest in each experiment to randomly rotate the points.

Acknowledgements. The authors would like to thank Shiri Artstein-Avidan for helpful discussions concerning Minkowski sums, and Sariel Har-Peled and Amir Rothschild for helpful discussions about the problem and its existing literature. Thanks are also due to two anonymous referees for their very helpful comments.

## References

[1] C. C. Aggarwal, A. Hinneburg, and D. A. Keim, On the surprising behavior of distance metrics in high dimensional spaces, in Proc. 8th Internat. Conf. Database Theory, 2001, 420-434.
[2] D. Aiger, E. Kokiopoulou, and E. Rivlin, Random grids: Fast approximate nearest neighbors and range searching for image search, in Proc. IEEE Annu. Internat. Conf. Computer Vision (ICCV) (2013), 3471-3478.
[3] A. Andoni, Implementation of LSH: E2LSH, http://www.mit.edu/~andoni/LSH/

[^6][4] A. Andoni and P. Indyk, Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions, in Proc. 47 th Annu. IEEE Sympos. Found. Comput. Sci., 2006, 459-468.
[5] A. Andoni and P. Indyk, Efficient algorithms for substring near neighbor problem, in Proc. $1^{77}$ th Annu. ACM-SIAM Sympos. Discrete Algorithms, 2006, 1203-1212.
[6] A. Andoni and P. Indyk, Near-optimal hashing algorithms for approximate nearest neighbors in high dimensions, Commun. ACM 51(1) (2008), 117-122.
[7] S. Arya, G. D. Fonseca, D. M. Mount, Approximate polytope membership queries, Proc. 43rd Annu. ACM Sympos. Theory of Computing (2011), 579-586.
[8] S. Arya, T. Malamatos, and D. M. Mount, Space-time tradeoffs for approximate nearest neighbor searching, J. ACM, 57(1) (2009), 1:1-1:54.
[9] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu, An optimal algorithm for approximate nearest neighbor searching in fixed dimensions, J. ACM 45(6) (1998), 891-923.
[10] P. Assouad, Plongements Lipschitziens dans R, Bull. Soc. Math. France, 111(4) (1983), 429-448.
[11] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, SURF: Speeded Up Robust Features, Computer Vision and Image Understanding 110(3) (2008), 346-359.
[12] S. N. Bespamyatnikh, Dynamic algorithms for approximate neighbor searching, in Proc. 8th Canad. Conf. Comput. Geom., 1996, 252-257.
[13] K. Beyer, J. Goldstein, R. Ramakrishnan, and U. Shaft, When is "Nearest Neighbor" meaningful?, in Proc. Internat. Conf. Database Theory, 1999, 217-235.
[14] K. Böröczky, Finite Packing and Covering, Cambridge Tracts in Mathematics, Vol. 154, Cambridge University Press, Cambridge 2004.
[15] T. M. Chan, On enumerating and selecting distances, Internat. J. Comput. Geom. Appls. 11(3) (2001), 291-304.
[16] T. M. Chan, Closest-point problems simplified on the RAM, in Proc. 13th Annu. ACMSIAM Sympos. Discrete Algorithms, 2002, 472-473.
[17] T. M. Chan, A minimalist's implementation of an approximate nearest neighbor algorithm in fixed dimensions. Manuscript. www.cs.vwaterloo.ca/tmchan/sss.ps.
[18] M. Charikar, C. Chekuri, A. Goel, S. Guha, and S. A. Plotkin, Approximating a finite metric by a small number of tree metrics, in Proc. 39th Annu. Sympos. Found. Comput. Sci., 1998, 379-388.
[19] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni, Locality-sensitive hashing scheme based on $p$-stable distributions, in Proc. 20th Annu. Sympos. Comput. Geom., 2004, 253-262. Also Chapter 3 in [30], pp. 61-72.
[20] L. Devroye and T. J. Wagner, Nearest neighbor methods in discrimination, in Handbook of Statistics, volume 2, P. R. Krishnaiah and L. N. Kanal, editors, North-Holland, Amsterdam 1982, 193-197.
[21] C. A. Duncan, M. T. Goodrich, and S. Kobourov, Balanced aspect ratio trees: Combining the advantages of k-d trees and octrees, J. Algorithms 38 (2001), 303-333.
[22] R. J. Gardner, Geometric Tomography, 2nd edition, Encyclopedia of Mathematics and its Applications, Vol 58, Cambridge University Press, New York, 2006.
[23] A. Genz, Methods for generating random orthogonal matrices, in Monte Carlo and Quasi Monte Carlo Methods 1998, H. Niederreite and J. Spanier (Eds.), SpringerVerlag, Berlin (1999), pp. 199-213.
[24] A. Gersho and R. M. Gray, Vector Quantization and Data Compression, Kluwer Academic Press, Boston, 1991.
[25] J. E. Goodman and J. O'Rourke, editors, Handbook of Discrete and Computational Geometry, Second Edition, CRC Press LLC, Boca Raton, FL, 2004.
[26] S. Har-Peled, Geometric Approximation Algorithms, Mathematical Surveys and Monographs, volume 173, AMS Press, Providence, RI, 2011.
[27] S. Har-Peled, A replacement for Voronoi diagrams of near linear size, in Proc. 42nd Annu. IEEE Sympos. Found. Comput. Sci., 2001, 94-103.
[28] S. Har-Peled and M. Mendel, Fast construction of nets in low-dimensional metrics and their applications, SIAM J. Comput. 35(5) (2006), 1148-1184.
[29] A. Hinneburg, C. C. Aggarwal, and D. A. Keim, What is the nearest neighbor in high dimensional spaces?, in Proc. 26th Internat. Conf. Very Large Data Bases, 2000, 506-515.
[30] G. Shakhnarovish, T. Darrell, and P. Indyk, Eds., Nearest-Neighbor Methods in Learning and Vision, MIT Press, Cambridge MA, 2006.
[31] P. Indyk and J. Matoušek, Low-distortion embeddings of finite metric spaces, in [25, pp. 177-196].
[32] P. Indyk and R. Motwani, Approximate nearest neighbors: Towards removing the curse of dimensionality, in Proc. 30th Annu. ACM Sypos. Theory Comput., 1998, 604-613.
[33] R. Krauthgamer and J. R. Lee, Navigating nets: Simple algorithms for proximity search, in Proc. 15th Annu. ACM-SIAM Sympos. Discrete Algorithms, 2004, 798-807.
[34] H.-P. Lenhof and M. Smid, Sequential and parallel algorithms for the $k$ closest pairs problem, Internat. J. Comput. Geom. Appls. 5 (1995), 273-288.
[35] M. Muja and D. G. Lowe, Fast approximate nearest neighbors with automatic algorithm configuration, in Internat. Conf. Computer Vision Theory Appls. (VISAPP'09), 2009, 331-340
[36] M. Muja and D. G. Lowe, The FLANN implementation, http://people.cs.ubc.ca/ ~mariusm/index.php/FLANN/FLANN
[37] M. E. Muller, A note on a method for generating points uniformly on $n$-dimensional spheres, Commun. ACM 2(4) (1959), 19-20.
[38] J. S. Salowe, Enumerating interdistances in space, Internat. J. Comput. Geom. Appls. 2 (1992), 49-59.
[39] R. Schneider, Convex Bodies: The Brunn-Minkowski Theory, Cambridge University Press, Cambridge, 1993.
[40] D. M. Y. Sommerville, An Introduction to the Geometry of $n$ Dimensions, Dover, New York, 1958.
[41] G. T. Toussaint, Geometric proximity graphs for improving nearest neighbor methods in instance-based learning and data mining, Internat. J. Comput. Geom. Appls. 15 (2) (2005), 101-150.


[^0]:    *Work by Haim Kaplan and Micha Sharir has been supported by the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11). Work by Haim Kaplan has also been supported by grant 822/10 from the Israel Science Fund and by Grant 2006/204 from the U.S.-Israel Binational Science Foundation. Work by Micha Sharir has also been supported by Grant 338/09 from the Israel Science Fund, and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University. A preliminary version of the paper has appeared in Proc. 24th ACM-SIAM Sympos. Discrete Algorithm, 2013.
    ${ }^{\dagger}$ Google Inc.; email: aigerd@google.com
    ${ }^{\ddagger}$ School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; email: haimk@tau.ac.il
    ${ }^{\text {§ School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of }}$ Mathematical Sciences, New York University, New York, NY 10012, USA; email: michas@tau.ac.il

[^1]:    ${ }^{1}$ Technically it is a different arbitrary distribution for each dimension $d$.

[^2]:    ${ }^{2}$ We still hide factors exponential in $d$ which are independent of $\epsilon$.

[^3]:    ${ }^{3}$ Note that we simplify the analysis by considering a unit cube $K$ instead of the grid cell $\tau$. We will later scale the bound that we obtain to fit it to the actual setup with an arbitrary cell size $c$.

[^4]:    ${ }^{4}$ The 3-dimensional case is also presented in detail in [22].

[^5]:    ${ }^{5}$ For our data set produced from 100 images we inspect about $40,240,4700$, and 90000 pairs for $c$ equals to $0.8,1,2$, and 3.6 , respectively.

[^6]:    ${ }^{6}$ We inspect about 80000 pairs for $c=0.8,500000$ pairs for $c=1.0$, and 120000000 pairs for $c=2.0$. For $c=3.6$ the number of pairs is too large for the computation to finish in a reasonable amount of time.

