Maximal and Convex Layers of Random Point Sets

Meng He, Cuong P. Nguyen, and Norbert Zeh

Faculty of Computer Science, Dalhousie University, Halifax, Nova Scotia, Canada
mhe@cs.dal.ca, cn536386@dal.ca, nzeh@cs.dal.ca

Abstract. We study two problems concerning the maximal and convex layers of a point set in d dimensions. The first is the average-case complexity of computing the first k layers of a point set drawn from a uniform or component-independent (CI) distribution. We show that, for $d \in \{2, 3\}$, the first $n^{1/d-\epsilon}$ maximal layers can be computed using $dn + o(n)$ scalar comparisons with high probability. For $d \geq 4$, the first $n^{1/2d-\epsilon}$ maximal layers can be computed within this bound with high probability. The first $n^{1/d-\epsilon}$ convex layers in 2D, the first $n^{1/2d-\epsilon}$ convex layers in 3D, and the first $n^{1/(d^2+2)}$ convex layers in $d \geq 4$ dimensions can be computed using $2dn + o(n)$ scalar comparisons with high probability. Since the expected number of maximal layers in 2D is $2\sqrt{n}$, our result for 2D maximal layers shows that it takes $dn + o(n)$ scalar comparisons to compute a $1/n^{\epsilon}$-fraction of all layers in the average case. The second problem is bounding the expected size of the kth maximal and convex layer. We show that the kth maximal and convex layer of a point set drawn from a continuous CI distribution in d dimensions has expected size $O(k^d \log^{d-1}(n/k^d))$.

Keywords: Maximal layers, skyline, convex layers, average-case analysis

1 Introduction

Maximal and convex layers are fundamental geometric structures with applications for example in data mining [5], pattern recognition and statistics [8,15]. A point p *dominates* another point q if p is no less than q in any dimension and p is greater than q in at least one dimension. The *skyline* (*first maximal layer*) of a d-dimensional point set S is the set of all points in S not dominated by any other point in S. A point $p \in S$ belongs to the *convex hull* (*first convex layer*) of S if there exists a $(d-1)$-dimensional hyperplane through p that has all points of S on the same side. For $k > 1$, the *kth maximal or convex layer* is the skyline or convex hull of the subset of S obtained by removing the first $k-1$ maximal or convex layers, respectively.

Computing maximal and convex layers are problems that have been studied extensively. A classical result of Kung et al. [12] shows that the skyline of a point set in 2D or 3D can be found in $O(n \log n)$ time; for any constant $d \geq 4$, the cost

in d dimensions is $O(n \log^{d-2} n)$. The convex hull of a 2D or 3D point set can also be found in $O(n \log n)$ time [3], while the cost of finding the convex hull in $d \geq 4$ dimensions is $\Theta(n^{\lfloor d/2 \rfloor})$ in the worst case [7]. A simple adversary argument shows that, in the worst case, $\Omega(n \log n)$ comparisons between scalars are necessary to compute the skyline or convex hull in d dimensions for any $d \geq 2$. For component-independent (CI) point distributions, on the other hand, expected linear-time algorithms exist, where a point distribution is component-independent if it is continuous and the coordinates of each point are chosen independently. The algorithm of Bentley et al. [1] uses $dn + o(n)$ scalar comparisons in expectation to find the skyline of a point set in d dimensions. For the convex hull, they presented an algorithm that uses $2dn + o(n)$ expected scalar comparisons for $d \in \{2, 3\}$. For $d \geq 4$, they presented an algorithm that finds a superset of the convex hull of expected size $O(\log^{d-1} n)$ using $2dn + o(n)$ expected scalar comparisons. They also proved that dn scalar comparisons is a lower bound for computing either the skyline or convex hull.

All maximal layers of a point set can be computed in $O(n \log n)$ time in 2D [4] and 3D [6]. For $d \geq 4$, no optimal algorithm for computing multiple maximal layers is known. The convex layers of a point set in 2D can be computed in $O(n \log n)$ time [8]. For $d \geq 3$, no optimal algorithm for computing multiple convex layers is known. Nielsen [15] presented an output-sensitive algorithm for finding the first k convex or maximal layers of a 2D point set in $O(n \log h_k)$ time, where h_k is the number of points in these layers.

One of the key ingredients of Bentley et al.'s skyline and convex hull algorithms [1] is the ability to quickly identify a small subset of points that is likely to contain all skyline or convex hull points. The skyline or convex hull can then be computed by applying one of the algorithms above to this subset of points. Such a small subset can exist only if the skyline or convex hull is small. Bentley et al. [2] proved that the expected size of the skyline or the expected number of vertices of the convex hull over an arbitrary CI distribution is $O(\log^{d-1} n)$. Note that the work on the expected complexity of geometric structures, including that of Bentley et al. [2], is of independent interest. Many other problems have also been defined and studied under similar assumptions. For instance, Dalal [9] shows that the expected number of convex layers is $\Theta(n^{2/(d+1)})$ for a set of n points independently chosen from a uniform distribution inside any bounded, nonempty region in $\mathbb{R}^d$. We refer to Okabe et al. [16] for a review of many problems in this area.

Our results. We extend Bentley et al.'s results [1] to multiple layers and strengthen the algorithm analysis by proving high-probability bounds on the number of scalar comparisons. Our first main result is a reduction that allows us to obtain an algorithm that computes the first k maximal or convex layers using $dn + o(n)$ or $2dn + o(n)$ expected scalar comparisons, respectively, given an algorithm that computes these layers using $O(k^c n^{1+\epsilon})$ scalar comparisons in the worst case. The exact bound on k is given in the following theorem.

Theorem 1. *Let S be a set of n points drawn from an arbitrary CI distribution in d dimensions. Suppose there is an algorithm M that can compute the first k maximal (or convex) layers of S using $O(k^c n^{1+\epsilon})$ scalar comparisons in the worst case, where c and ϵ are constants with $c \geq 0$ and $0 < \epsilon < \frac{1}{(c+1)d}$. Then the first $\kappa = n^{\frac{1}{(c+1)d}-\epsilon}$ maximal (or convex) layers of S can be computed using $dn + o(n)$ (or $2dn + o(n)$) expected scalar comparisons, and the actual number of comparisons is within the same bounds with probability $1 - o(n^{-n^{\epsilon'}})$ for any $\epsilon' \in (0, (c\epsilon + \frac{\epsilon^2}{2(\epsilon+1)})d)$.*

To achieve this result, our main strategy is to generalize the algorithms of Bentley et al. [1] to compute more than one maximal or convex layer. While it is not difficult to generalize the algorithms themselves, it is more challenging to analyze their running times. To perform the analysis, our key strategy is to further conceptually subdivide some objects defined by these algorithms into even smaller objects, such that a layer must contain a point inside a particular smaller object with high probability. These constructs may be of general interest, as they may be useful to the tasks of performing some other similar analysis over multiple layers of the given point set.

The existing algorithms discussed previously allow us to find the first k maximal layers using $O(n^{1+\epsilon})$ comparisons for $d \in \{2, 3\}$ and using $O(kn^{1+\epsilon})$ comparisons for $d \geq 4$. The first k convex layers can be computed using $O(n^{1+\epsilon})$ comparisons in 2D and using $O(kn^{1+\epsilon})$ comparisons in 3D. Thus, we obtain the following corollary of Theorem 1:

Corollary 1. *Let S be a set of n points drawn from an arbitrary CI distribution in d dimensions. If $d \in \{2, 3\}$, the first $n^{\frac{1}{d}-\epsilon}$ maximal layers of S can be computed using $dn + o(n)$ expected scalar comparisons. If $d \geq 4$, the first $n^{\frac{1}{2d}-\epsilon}$ maximal layers can be computed using this expected number of scalar comparisons. If $d = 2$, the first $n^{\frac{1}{d}-\epsilon}$ convex layers of S can be computed using $2dn+o(n)$ expected scalar comparisons. If $d = 3$, the first $n^{\frac{1}{2d}-\epsilon}$ convex layers can be computed using this number of expected scalar comparisons. In all these cases, the actual number of comparisons is within the same upper bounds on the expected number of comparisons with probability $1 - o(n^{-n^{\epsilon'}})$.*

Our results are the first that show that more than one maximal or convex layer can be computed using the optimal number of scalar comparisons on random point sets up to lower order terms and, in the case of convex hull, up to a constant factor of 2. With the exception of a high-probability analysis of an alternative skyline algorithm by Bentley et al. [1] provided by Golin [11], only expected bounds on the number of scalar comparisons were known even for computing only the first convex or maximal layer.

The number of maximal layers of a point set S is the length of a longest monotonically increasing subsequence (LMIS) of the sequence of y-coordinates of the points in S sorted by their x-coordinates. If S is drawn from a CI distribution, this sequence of y-coordinates is a uniform random permutation of

the y-coordinates. Thus, by a result of [10], the expected length of an LMIS of this sequence, and thus the number of maximal layers of S approaches $2\sqrt{n}$ as n approaches infinity. Therefore, for $d = 2$, our algorithm finds a $1/n^{\epsilon}$-fraction of all maximal layers in the average case using the optimal number of scalar comparisons up to lower-order terms.

For $d \geq 4$ dimensions, no convex hull algorithm using expected $2dn + o(n)$ comparisons on random point sets was known, and we cannot satisfy the condition of Theorem 1 even for $k = 1$ since computing the convex hull takes $\Theta(n^{\lfloor d/2 \rfloor})$ time in the worst case. However, the construction that proves Theorem 1 can be combined with the $\Theta(n^{\lfloor d/2 \rfloor})$-time convex hull algorithm to obtain the following theorem:

Theorem 2. *Let S be a set of n points in $d \geq 4$ dimensions drawn from an arbitrary CI distribution. For any $k \leq n^{1/(d^2+2)}$, the first k convex layers of S can be found using $2dn + o(n)$ scalar comparisons with probability $1 - O(\frac{1}{n^{1/d-\epsilon}})$, for any $\epsilon > 0$.*

This result is the first that computes multiple convex layers in four or higher dimension in linear time with high probability.

Our second main result bounds the size of the kth maximal or convex layer of a d-dimensional point set. Previously, only bounds on the expected size of the *first* maximal or convex layer were known.

Theorem 3. *For any point set S drawn from a continuous CI distribution in d dimensions, the kth maximal or convex layer has expected size $O(k^d \log^{d-1}(n/k^d))$.*

2 Algorithm Overview

Bentley et al.'s algorithms [1] for computing the skyline or convex hull of a point set S using expected $dn+o(n)$ or $2dn+o(n)$ comparisons uses the following simple idea: Find a rectangular *inner region* I that is expected to contain almost all points in S and is likely to be completely below the skyline of S or inside the convex hull of S. See Figure 1. In particular, with high probability, the points in S that belong to the skyline or convex hull are all contained in the *outer region* $O = \mathbb{R}^d \setminus I$. The algorithm partitions the point set S into two subsets $S_I = S \cap I$ and $S_O = S \cap O$ and computes the skyline or convex hull L_O of S_O using some standard skyline or convex hull algorithm, which takes $o(n)$ time in expectation because S_O is small. Finally, the algorithm checks whether certain subregions of O (C and $C_1, \ldots, C_4$, respectively, in Figure 1) each contain at least one point of S_O. If so, I is completely below or inside L_O, which implies that L_O is also the skyline or convex hull of S because no point in $S \setminus S_O = S_I \subseteq I$ can be on the maximal layer or convex hull of S. Thus, the algorithm terminates in this case. Otherwise, the algorithm runs a standard skyline or convex hull algorithm on S to compute the skyline or convex hull of S. While this is costly, this happens infrequently because I is likely to be below the skyline or inside the convex hull of S, so the expected cost of this final step is again $o(n)$.

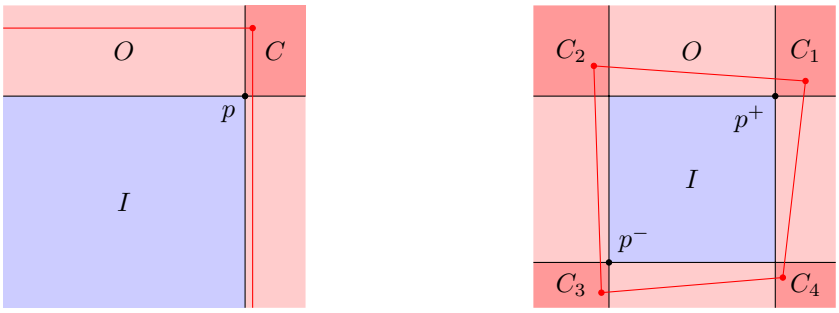

(a) Maximal layers (b) Convex layers

Fig. 1: The inner and outer regions used in Bentley et al.'s [1] and our algorithm illustrated for the 2D case. I is shaded blue. O is shaded pink, including the darker regions, which are the corners that are tested by the algorithm whether they contain a point not on the first k maximal or convex layers. As illustrated in red, any point in C dominates I in the case of maximal layers; in the case of convex layers, the convex hull of any four points in $C_1, \ldots, C_4$ encloses I.

For the skyline algorithm, $I = (-\infty, x_1(p)] \times (-\infty, x_2(p)] \times \cdots \times (-\infty, x_d(p)]$, where p is an appropriate point and $x_i(p)$ denotes the ith coordinate of p, so the partition into S_I and S_O can be obtained using dn scalar comparisons. For the convex hull algorithm, $I = [x_1(p^-), x_1(p^+)] \times [x_2(p^-), x_2(p^+)] \times \cdots \times [x_d(p^-), x_d(p^+)]$ for an appropriate pair of corner points (p^-, p^+), so the partition can be performed using $2dn$ scalar comparisons. The corner points of I can be found without comparisons by setting $x_i(p^-) = \epsilon$ and $x_i(p) = x_i(p^+) = 1 - \epsilon$ for all $1 \leq i \leq d$ and some appropriate value $\epsilon > 0$. At least this is the case for points distributed uniformly at random in the unit hypercube. For an arbitrary CI distribution, p, p^-, and p^+ can each be found using $dn + o(n)$ scalar comparisons using randomized linear-time selection. The partitioning of S into S_I and S_O can be done as part of the selection process without incurring any additional comparisons. We discuss this in more detail as part of our high-probability analysis and extension to multiple layers (Lemmas 1 and 4). Overall, the expected cost of the algorithm is $dn + o(n)$ or $2dn + o(n)$ comparisons for finding p or p^- and p^+ and computing the partition of S into S_I and S_O plus $o(n)$ expected comparisons for computing the maximal layer or convex hull.

To extend Bentley et al.'s result [1] to multiple maximal or convex layers, we need to show that there exists a point p or a pair of points (p^-, p^+) that defines inner and outer regions I and O as above such that, again, almost all points in S are inside I and the first k layers are unlikely to intersect I. To achieve a running time of $dn + o(n)$ or $2dn + o(n)$ *with high probability*, we also need to strengthen the analysis of Bentley et al. [1] to (a) show that these points can be found using $dn + o(n)$ or $2dn + o(n)$ scalar comparisons with high probability and (b) with high probability, I does not intersect the first k layers.

Since the proofs are slightly simpler, we present our result for maximal layers first. Then, in Section 4, we argue that the same approach, with minor modifications, can also be used to compute convex layers.

3 Maximal Layers

Throughout this section, we use $p \nearrow q$ to indicate that q dominates p. Given a point set S drawn from a CI distribution $\mathcal{D}$ and some value $\tau \in [0,1]$, we call a point p a *τ-pivot* of S if, for all $1 \leq i \leq d$ and any point p' chosen uniformly at random from S, $P[x_i(p') \geq x_i(p)] = \tau$; recall that $x_i(p)$ denotes the ith coordinate of point p. Point p is not necessarily in S. We first prove the following lemma on locating p.

Lemma 1. *Let S be a point set drawn from a CI distribution. For any value $t > 0$, any value $\tau \in (0, n^{-t}] \cup [1-n^{-t}, 1)$, and any constant $\epsilon' \in (0,1)$, a τ-pivot p and a partition of S into two subsets $S_I = S \cap I$ and $S_O = S \cap (\mathbb{R}^d \setminus I)$ can be computed using $dn + o(n)$ scalar comparisons in expectation and with probability at least $1 - o\left(n^{-n^{\epsilon'}}\right)$, where I is the region dominated by p.*

Proof. If S is drawn uniformly at random from the unit hypercube, then $p = (1-\tau, \ldots, 1-\tau)$ is a τ-pivot and can be found without any comparisons. The partition of S into S_I and S_O can be computed by deciding for each point whether it is dominated by p (and thus belongs to S_I) or not (and thus belongs to S_O). This takes d comparisons per point in S, incurring dn comparisons in total, that is, the lemma holds in the worst case for a uniform random distribution. For an arbitrary CI distribution, set $x_i(p)$ to be the (τn)th largest coordinate in dimension i among the points in S. Then p is a τ-pivot. Each value $x_i(p)$ can be found using $n + o(n)$ scalar comparisons in expectation and with probability at least $1 - o\left(n^{-n^{\epsilon'}}\right)$ using a simplified version of LazySelect [14]; we omit the details due to page constraints. In the process, every point in S is tagged as having ith coordinate less than or equal to, or greater than $x_i(p)$. Doing this for all d dimensions produces p and takes $dn+o(n)$ scalar comparisons in expectation and with probability at least $1 - o\left(n^{-n^{\epsilon'}}\right)$. The partition of S into S_I and S_O is then obtained without additional scalar comparisons by collecting all points tagged as greater than p in at least one dimension into S_O, and the remaining points into S_I. □

The following observation and lemmas are needed for our proof of Theorem 1.

Observation 4 *Let p be a τ-pivot of S and consider the corresponding partition of S into subsets S_I and S_O as in Lemma 1. If there exist $k+1$ points $p_1, p_2, \ldots, p_{k+1}$ in S_O such that $p \nearrow p_{k+1} \nearrow \cdots \nearrow p_1$, then the first k maximal layers of S and S_O are identical and p_{k+1} is not part of these layers.*

Lemma 2. *Let S be a point set drawn from a CI distribution, let $0 < \epsilon_1 < \epsilon_2 < 1$ be constants, let $\frac{\tau}{k+1} \geq n^{(\epsilon_2-1)/d}$, and let $h_0, h_1, \ldots, h_{k+1}$ be $k+2$ points such that h_j is a $\left(\frac{j}{k+1}\tau\right)$-pivot of S for all $0 \leq j \leq k+1$. Then with probability at least $1 - o\left(n^{-n^{\epsilon_1}}\right)$, each hyperrectangle H_j defined by points h_{j-1} and h_j, for $1 \leq j \leq k+1$, contains a point $p_j \in S$. These points satisfy $p_{k+1} \nearrow p_k \nearrow \cdots \nearrow p_1$.*

Proof. Consider an arbitrary hyperrectangle H_j. Since h_{j-1} is a $\left(\frac{j-1}{k+1}\tau\right)$-pivot and h_j is a $\left(\frac{j}{k+1}\tau\right)$-pivot, each point $p \in S$ satisfies $x_i(h_{j-1}) \leq x_i(p) \leq x_i(h_j)$ with probability $\frac{\tau}{k+1}$ for each $1 \leq i \leq d$. Since the coordinates are chosen independently, $p \in H_j$ with probability $\left(\frac{\tau}{k+1}\right)^d$. Thus, $E(|H_j \cap S|) = \left(\frac{\tau}{k+1}\right)^d n$. Since $|H_j \cap S|$ is the sum of independent Bernoulli random variables, the Chernoff bound states that $P(H_j \cap S = \emptyset) < e^{-(\tau/(k+1))^d n/4}$ and the probability that there exists an index $1 \leq j \leq k+1$ such that $H_j \cap S = \emptyset$ is less than $(k+1)e^{-(\tau/(k+1))^d n/4}$. For $\frac{\tau}{k+1} \geq n^{(\epsilon_2-1)/d}$, this is bounded by $(k+1)e^{-n^{\epsilon_2}/4} \leq n^{1-n^{\epsilon_2}/(4\ln n)} = o\left(n^{-n^{\epsilon_1}}\right)$ for any $\epsilon_1 < \epsilon_2$ because $k+1 \leq n$. □

Lemma 3. *Let S be a point set drawn from a CI distribution, let $0 < \epsilon_1 < \epsilon_2 < 1$ be constants, let $\tau \geq n^{\epsilon_2-1}$, let p be a τ-pivot of S, let $S_I \subseteq S$ be the set of points dominated by p, and let $S_O = S \setminus S_I$. Then $E(|S_O|) \leq d\tau n$ and $P(|S_O| > 2d\tau n) = o\left(n^{-n^{\epsilon_1}}\right)$.*

Proof. We can cover the outer region O with d halfspaces $B_1, B_2, \ldots, B_d$, where $B_i = \{p' \in \mathbb{R}^d \mid x_i(p') \geq x_i(p)\}$. Since a point $p' \in S$ satisfies $x_i(p') \geq x_i(p)$ with probability τ, we have $E(|B_i \cap S|) = \tau n$ and $E(|S_O|) \leq \sum_{i=1}^d E(|B_i \cap S|) = d\tau n$. Since $|B_i \cap S|$ is the sum of independent Bernoulli random variables, the Chernoff bound states that $P(|B_i \cap S| > 2\tau n) < e^{-\tau n/3} \leq n^{-n^{\epsilon_2}/(3\ln n)} = o\left(n^{-n^{\epsilon_1}}\right)$. Thus, $P(|S_O| > 2d\tau n) \leq \sum_{i=1}^d P(|B_i| > 2\tau n) = o\left(dn^{-n^{\epsilon_1}}\right) = o\left(n^{-n^{\epsilon_1}}\right)$. □

Proof (Proof of Theorem 1 (Maximal Layers).). Our algorithm finds a τ-pivot p of S, partitions S into S_I and S_O, computes the first k maximal layers of S_O using M, and checks whether there exists a point in S_O that is not on the computed maximal layers but dominates p. If this test succeeds, then the maximal layers of S and S_O are the same, so the algorithm reports the computed maximal layers. Otherwise, it runs M on S to compute the first k maximal layers of S.

We prove that this algorithm uses $dn + o(n)$ scalar comparisons with high probability. The analysis of the expected number of comparisons is analogous. The number of comparisons the algorithm performs is $dn + o(n)$ if (a) computing p and partitioning S into S_I and S_O takes $dn + o(n)$ comparisons, (b) running algorithm M on S_O incurs $o(n)$ comparisons, and (c) there exists a point in S_O that is not on the first k maximal layers and dominates p, that is, the fallback option of running M on the entire point set S is not invoked. Thus, it suffices to bound the probability that any of these three conditions fails.

By Lemma 1, (a) fails with probability $o\left(n^{-n^{\epsilon'}}\right)$, for any $\epsilon' \in (0,1)$, as long as $\tau = n^{-t}$ for some $t > 0$. Running algorithm M on S_O incurs $o(n)$

scalar comparisons if $|S_O| = o\big(n^{1/(1+\epsilon)}\big)/k^c$. By Lemma 3, $|S_O| \leq 2d\tau n$ with probability $1 - o\Big(n^{-n^{\epsilon'}}\Big)$ as long as $\tau \geq n^{\epsilon_2 - 1}$ for some $\epsilon_2 > \epsilon'$. Therefore, (b) fails with probability $o\Big(n^{-n^{\epsilon'}}\Big)$ as long as $\tau n = o\big(n^{1/(1+\epsilon)}\big)/k^c$ and $\tau \geq n^{\epsilon_2 - 1}$. By Observation 4 and Lemma 2, (c) fails with probability $o\Big(n^{-n^{\epsilon'}}\Big)$ as long as $\frac{\tau}{k+1} \geq n^{(\epsilon_2-1)/d}$ for some $\epsilon_2 > \epsilon'$. Thus, the probability that any of these three conditions fails is $o(n^{-n^{\epsilon'}})$, provided we can choose τ so that the above constraints are satisfied.

First observe that $\epsilon_2 - 1 < 0$. Thus, $\tau \geq n^{\epsilon_2 - 1}$ if $\frac{\tau}{k+1} \geq n^{(\epsilon_2-1)/d}$, so we have to choose a value of $\tau = n^{-t}$, for some $t > 0$, such that $\frac{\tau}{k+1} \geq n^{(\epsilon_2-1)/d}$ and $\tau n = o\big(n^{1/(1+\epsilon)}\big)/k^c$. The last two constraints imply that $k^{c+1} = o(n^{-\epsilon/(\epsilon+1)+(1-\epsilon_2)/d})$ or $k = o(n^{\frac{1}{(c+1)d} - \epsilon + \delta})$ where $\delta = \epsilon - \frac{\epsilon_2}{(c+1)d} - \frac{\epsilon}{(\epsilon+1)(c+1)}$. For any $\epsilon_2 < \big(c\epsilon + \frac{\epsilon^2}{2(\epsilon+1)}\big)d$, we have $\delta > 0$, that is, we can compute up to $n^{\frac{1}{(c+1)d} - \epsilon}$ maximal layers and, since $(c\epsilon + \frac{\epsilon^2}{2(\epsilon+1)}) > 0$, we can choose values ϵ' and ϵ_2 such that $0 < \epsilon' < \epsilon_2 < (c\epsilon + \frac{\epsilon^2}{2(\epsilon+1)})d$. It remains to choose τ. We have $\tau n = o\big(n^{1/(1+\epsilon)}\big)/k^c$ if $t > \frac{\epsilon}{1+\epsilon} + \frac{c}{(c+1)d} - \epsilon c$. To satisfy $\frac{\tau}{k+1} \geq n^{(\epsilon_2-1)/d}$, we need $t = -\log_n \tau \leq -\log_n(k+1) - \frac{\epsilon_2 - 1}{d}$. To compute the first $n^{\frac{1}{(c+1)d} - \epsilon}$ maximal layers, we replace k by $n^{\frac{1}{(c+1)d} - \epsilon}$ in this inequality, and it holds for large enough n if t is a constant and $t < \epsilon - \frac{1}{(c+1)d} - \frac{\epsilon_2 - 1}{d}$, which is true as long as t is a constant satisfying $t < \epsilon - \frac{1}{(c+1)d} - \epsilon c - \frac{\epsilon^2}{2(1+\epsilon)} + \frac{1}{d}$ because $\epsilon_2 < (c\epsilon + \frac{\epsilon^2}{2(\epsilon+1)})d$. It is easy to verify that $\frac{\epsilon}{1+\epsilon} + \frac{c}{(c+1)d} - \epsilon c < \epsilon - \frac{1}{(c+1)d} - \epsilon c - \frac{\epsilon^2}{2(1+\epsilon)} + \frac{1}{d}$. Thus, we can choose a value of t that satisfies both constraints and set $\tau = n^{-t}$. In addition, since $\epsilon < \frac{1}{(c+1)d}$, we have $\frac{\epsilon}{1+\epsilon} + \frac{c}{(c+1)d} - \epsilon c > 0$, that is, $t > 0$. □

4 Convex Layers

Convex Layers in Two and Three Dimensions: To apply the framework from Section 3 to compute convex layers, we need to extend the notion of dominance to the 2^d possible quadrants of a point in $\mathbb{R}^d$. We identify each quadrant using a *sign vector* $\sigma \in \{+1, -1\}^d$. We say a point $q \in \mathbb{R}^d$ *σ-dominates* another point $p \in \mathbb{R}^d$, written as $p \nearrow_\sigma q$ if $\sigma \circ q$ dominates $\sigma \circ p$, where $p \circ q$ is the Hadamard product: $p \circ q = (x_1(p), x_2(p), \ldots, x_d(p)) \circ (x_1(q), x_2(q), \ldots, x_d(q)) = (x_1(p)x_1(q), x_2(p)x_2(q), \ldots, x_d(p)x_d(q))$. We call a point p a *(τ, σ)-pivot* of S if, for all $1 \leq i \leq d$ and any point p' chosen uniformly at random from S, $P(x_i(\sigma)x_i(p') \geq x_i(\sigma)x_p(p)) = \tau$. Note that $\mathbf{1}$-dominance is the same as normal dominance, a $(\tau, \mathbf{1})$-pivot is just a τ-pivot, and a $(\tau, -\mathbf{1})$-pivot is a $(1-\tau)$-pivot, where $\mathbf{1} = (1, \ldots, 1)$ and $-\mathbf{1} = (-1, \ldots, -1)$. A pair of points (p^-, p^+), where $0 < \tau < 1/2$, p^{-1} is a $(\tau, -\mathbf{1})$-pivot, and p^+ is a $(\tau, \mathbf{1})$-pivot, divides $\mathbb{R}^d$ into an *inner region* I containing all points in $\mathbb{R}^d$ that dominate p^- and are dominated by p^+, and an *outer region* $O = \mathbb{R}^d \setminus I$; see Figure 1. Similar to maximal

layers, we define $S_I = S \cap I$ and $S_O = S \cap O$. The corners of I are the points $\{p^\sigma \mid \sigma \in \{+1,-1\}^d\}$, where $p^\sigma = \frac{1}{2}((\mathbf{1}+\sigma) \circ p^+ + (\mathbf{1}-\sigma) \circ p^-)$. Since S is drawn from a CI distribution, each such corner p^σ is a (τ,σ)-pivot of S.

Our algorithm finds (p^-, p^+), partitions S into S_I and S_O, computes the first k convex layers of S_O using M, and checks whether, for every $\sigma \in \{+1,-1\}^d$, there exists a point in S_O that is not on the computed convex layers but σ-dominates p^σ. If this test succeeds, then the convex layers of S and S_O are the same, so the algorithm reports the computed convex layers. Otherwise, it runs M on S to compute the first k convex layers of S. To analyze this algorithm, we first prove the following lemmas and observation.

Lemma 4. *Let S be a point set drawn from a CI distribution. For any value $t > 0$, any value $\tau \in (0, n^{-t}]$, and any constant $\epsilon' \in (0,1)$, a pair of points (p^-, p^+) such that p^- is a $(\tau, -\mathbf{1})$-pivot of S and p^+ is a $(\tau, \mathbf{1})$-pivot of S and a partition of S into two subsets $S_I = S \cap I$ and $S_O = S \cap O$ can be computed using $2dn + o(n)$ scalar comparisons in expectation and with probability at least $1 - o\Big(n^{-n^{\epsilon'}}\Big)$.*

Proof. Since p^+ is a τ-pivot and p^- is a $(1-\tau)$-pivot of S, we can find these two points using the claimed number of scalar comparisons by applying Lemma 1 twice. In the case of an arbitrary CI distribution, the selection of the coordinates of p^- and p^+ also tags each point as having ith coordinate less than $x_i(p^-)$, between $x_i(p^-)$ and $x_i(p^+)$ or greater than $x_i(p^+)$, for each $1 \le i \le d$. Thus, S_I and S_O can be produced without any additional scalar comparisons by placing each point that has at least one coordinate less than p^- or at least one coordinate greater than p^+ into S_O and all remaining points into S_I. □

Observation 5 *Let $h_0^-, h_1^-, \ldots, h_{k+1}^-$ and $h_0^+, h_1^+, \ldots, h_{k+1}^+$ be points such that $h_0^- \nearrow h_1^- \nearrow \ldots \nearrow h_{k+1}^- \nearrow h_{k+1}^+ \nearrow h_k^+ \nearrow \cdots \nearrow h_0^+$ and consider the regions I and O defined by the pair of points $p^- = h_{k+1}^-$ and $p^+ = h_{k+1}^+$. Each pair of points (h_j^-, h_j^+) defines a hyperrectangle with corner set $\{h_i^\sigma \mid \sigma \in \{+1,-1\}^d\}$ similar to the corner set $\{p^\sigma \mid \sigma \in \{+1,-1\}^d\}$ of I defined by (p^-, p^+). If, for every sign vector $\sigma \in \{+1,-1\}^d$, there exist $k+1$ points $p_1^\sigma, p_2^\sigma, \ldots, p_{k+1}^\sigma$ in S_O such that $h_j^\sigma \nearrow^\sigma p_j^\sigma \nearrow^\sigma h_{j-1}^\sigma$ for all $1 \le j \le k+1$, then the first k convex layers of S and S_O are identical and p_{k+1}^σ is not part of these layers for any $\sigma \in \{+1,-1\}^d$.*

Lemma 5. *Let S be a point set drawn from a CI distribution, let $0 < \epsilon_1 < \epsilon_2 < 1$ be constants, let $\frac{\tau}{k+1} \ge n^{(\epsilon_2-1)/d}$, and let $h_0^-, h_1^-, \ldots, h_{k+1}^-, h_0^+, h_1^+, \ldots, h_{k+1}^+$ be points such that h_j^- is a $\big(\frac{j}{k+1}\tau, -\mathbf{1}\big)$-pivot and h_j^+ is a $\big(\frac{j}{k+1}\tau, \mathbf{1}\big)$-pivot for all $0 \le j \le k+1$. Then with probability at least $1 - o\big(n^{-n^{\epsilon_1}}\big)$, every hyperrectangle H_j^σ defined by the points h_{j-1}^σ and h_j^σ, for $1 \le j \le k+1$ and every sign vector $\sigma \in \{+1,-1\}^d$, contains a point $p_j^\sigma \in S$.*

Proof. Analogous to the proof of Lemma 2, $P(H_j^\sigma \cap S = \emptyset) < e^{-(\tau/(k+1))^d n/4}$, so the probability that there exists a pair (j,σ) such that $H_j^\sigma \cap S = \emptyset$ is less than $(k+$

$1)2^d e^{-(\tau/(k+1))^d n/4}$. As shown in the proof of Lemma 2, $(k+1)e^{-(\tau/(k+1))^d n/4} = o\big(n^{-n^{\epsilon_1}}\big)$. Since d is a constant, this implies that $(k+1)2^d e^{-(\tau/(k+1))^d n/4} = o\big(n^{-n^{\epsilon_1}}\big)$. □

Lemma 6. *Let S be a point set drawn from a CI distribution, let $0 < \epsilon_1 < \epsilon_2 < 1$ be constants, let $\tau \geq n^{\epsilon_2 - 1}$, let p^- be a $(\tau, -\mathbf{1})$-pivot of S, let p^+ be a $(\tau, \mathbf{1})$-pivot of S, let I be the hyperrectangle defined by (p^-, p^+), let $S_I = S \cap I$, and let $S_O = S \setminus S_I$. Then $E(|S_O|) \leq 2d\tau n$ and $P(|S_O| > 4d\tau n) = o\big(n^{-n^{\epsilon_1}}\big)$.*

Proof. The proof is identical to the proof of Lemma 3 after observing that S_O can be covered with $2d$ halfspaces $B_1^-, B_2^-, \ldots, B_d^-, B_1^+, B_2^+, \ldots, B_d^+$, where $B_i^- = \{p' \in \mathbb{R}^d \mid x_i(p') \leq x_i(p^-)\}$ and $B_i^+ = \{p' \in \mathbb{R}^d \mid x_i(p') \geq x_i(p^+)\}$ for all $1 \leq i \leq d$. □

With these lemmas and observation, we claim that the analysis of the algorithm in Section 4 that computes the first k convex layers in two or three dimensions is identical to the proof of Theorem 1 for maximal layers, using Lemmas 4, 5, and 6 and Observation 5 in place of Lemmas 1, 2, and 3 and Observation 4. This completes the proof of Theorem 1.

Convex Layers in Four or Higher Dimensions: We now consider the problem of computing the first k convex layers of a point set S drawn from an arbitrary CI distribution in four or higher dimensions, for $k \leq n^{1/(d^2+2)}$. The framework of Theorem 1 cannot be applied directly to this problem because the best known algorithm for computing even the convex *hull* in $d \geq 4$ dimensions [7] takes $O(n^{\lfloor d/2 \rfloor})$ comparisons.

Bentley et al. [2] showed how to use 2^d skyline computations to produce a superset Q' of the convex hull Q of S of small expected size. Matoušek [13] later called this structure the *quadrant hull* of S. We opt for *orthant hull* here because an orthant is the generalization of quadrants to higher dimensions. We will show later in this section that, with high probability, the size of Q' is small enough so that applying Chazelle's convex hull algorithm to Q' takes $O(n)$ comparisons. Combined with Kung et al.'s algorithm for computing the skyline in d dimensions in $O(n \log^{d-2} n)$ time [12], this gives an algorithm M that computes the convex hull of S using $O(n \log^{d-2} n)$ comparisons with high probability. To compute $k > 1$ convex layers, M repeats this process k times: the ith iteration computes the convex hull of the point set left after removing the first $i-1$ convex layers. With high probability, this will take $O(kn \log^{d-2} n)$ time because, as we show below, the size not only of the orthant hull but in fact of the first k *orthant layers* is small enough to apply Chazelle's algorithm k times.

To prove Theorem 2, we use M in conjunction with Theorem 1 where $c = 1$. Theorem 1 requires M to achieve a running time of $O(kn^{1+\epsilon})$ in the worst case. However, it is easily verified that the proof of Theorem 1 continues to hold if M achieves this running time with some probability $p > 0$, in which case Theorem 1 produces a convex hull algorithm that uses $2dn + o(n)$ scalar comparisons with probability $\Theta\Big(\min\Big(p, 1 - o\Big(n^{-n^{\epsilon'}}\Big)\Big)\Big)$. Since we prove below that M achieves a

running time of $O(n \log^{d-2} n)$ with probability $1 - O\left(\frac{1}{n^{1/d-\epsilon}}\right)$, Theorem 2 follows by setting $\epsilon = 1/(2d) - 1/(d^2+2)$ in Theorem 1.

Let $\sigma \in \{+1, -1\}^d$ be a sign vector, let $\sigma \circ S = \{\sigma \circ p \mid p \in S\}$, and let L^σ be the set of points $p \in S$ such that $\sigma \circ p$ belongs to the skyline of $\sigma \circ S$. We call L^σ the *σ-skyline* of S. The *orthant hull* of S is $Q' = \bigcup_{\sigma \in \{+1,-1\}^d} L^\sigma$, and Bentley et al. proved that $Q \subseteq Q'$. To define the *orthant layers* of S, let Q' be the first orthant layer of S and, for $i > 1$, let the ith orthant layer of S be the orthant hull of the subset of S obtained after removing the first $i-1$ orthant layers from S.

Let $Q_1, Q_2, \ldots, Q_k$ be the first k convex layers of S, let $S_i = S \setminus \bigcup_{j=1}^{i-1} Q_j$, and let Q'_i be the orthant hull of S_i for all $1 \leq i \leq k$. Since Q_i is the convex hull of S_i, Bentley et al.'s result shows that $Q_i \subseteq Q'_i$ and it is not hard to see that $\bigcup_{i=1}^k Q'_i$ is a subset of the points on the first k orthant layers[1] of S. Computing Q'_i in the ith iteration takes $O(2^d n \log^{d-2} n) = O(n \log^{d-2} n)$ time by applying Kung et al.'s algorithm once for each sign vector σ. Summing over all k iterations, we get an upper bound $O(kn \log^{d-2} n)$. To compute Q_i, we apply Chazelle's algorithm to Q'_i, which takes $O(|Q'_i|^{\lfloor d/2 \rfloor})$ time. Summing over all k layers, we obtain that computing the first k convex layers using M takes $O(kn \log^{d-2} n + \sum_{i=1}^k |Q'_i|^{\lfloor d/2 \rfloor}) = O(kn \log^{d-2} n + k|Q''|^{\lfloor d/2 \rfloor})$ time, where Q'' is the set of points on the first k orthant layers of S. As we show in Section 5.2, $E(k^{2/d}|Q''|) = O(k^{2/d} k^d \log^{d-1} n) = O(n^{1/d} \log^{d-1} n)$ because $k \leq n^{1/(d^2+2)}$. Thus, by Markov's inequality, $P(k^{2/d}|Q''| > n^{2/d}) \leq n^{-1/d+\epsilon}$, that is, M takes $O(kn \log^{d-2} n)$ time with probability at least $1 - n^{-1/d+\epsilon}$, as claimed.

5 Expected Size of the First k Layers

Bentley et al. [2] proved that the expected size of the skyline of a point set drawn from a CI distribution in d dimensions is $O(\log^{d-1} n)$. They also used this result to give a bound of $O(\log^{d-1} n)$ on the expected number of vertices on the convex hull. It seems difficult to extend their technique to subsequent layers. In Section 5.1, we show that, for continuous CI distributions, the kth maximal layer has expected size $O(k^d \log^{d-1}(n/k^d))$. The proof is based on a proof sketch for 2D suggested by an anonymous reviewer of an earlier draft of this paper. In Section 5.2, we show how to extend the argument to obtain the same bound (up to a factor of 4^d) for convex and orthant layers. This proves Theorem 3.

5.1 Maximal Layers

First consider a point set S drawn uniformly at random from the unit hypercube. To simplify the discussion, we bound the size of the kth *minimal* layer of S, which is equivalent to the kth maximal layer via the transformation $(x_1, x_2, \ldots, x_d) \mapsto (1 - x_1, 1 - x_2, \ldots, 1 - x_d)$. For every point $p \in \mathbb{R}^d$, let D_p be the set of points

[1] Note that $\bigcup_{i=1}^k Q'_i$ is not a subset of the first k *σ-skyline* of S. A counterexample will be given in the full version of this paper.

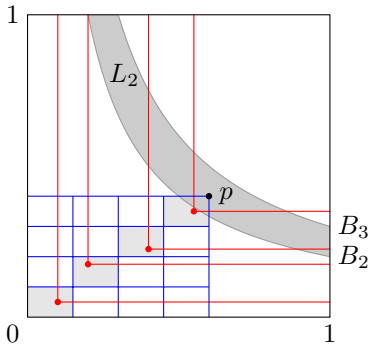

Fig. 2: Bound on the probability that a point in L_t belongs to the first k minimal layers in 2D. Here, $k = 2$ and $n = 160$. The region L_2 is shaded. Each of the grey grid cells contains two points of S in expectation. Any point in such a grid cell (red) is dominated by p and all points in grid cells to its top right. Thus, unless more than two of these grid cells are empty, p does not belong to the first two minimal layers.

dominated by p, and let $|D_p|$ be the volume of D_p. For every integer $t \geq 0$, let B_t be the set of all points $p \in [0,1]^d$ with $|D_p| = \frac{(2k)^d t}{n}$, that is, a point $(x_1, x_2, \ldots, x_d)$ belongs to B_t if and only if $x_1 x_2 \cdots x_d = \frac{(2k)^d t}{n}$ and $0 \leq x_i \leq 1$ for all $1 \leq i \leq d$. B_t splits the unit hypercube into two regions: L_t^- includes $(0, \ldots, 0)$ and L_t^+ includes $(1, \ldots, 1)$. For $t \geq 0$, let $L_t = L_t^+ \cap L_{t+1}^-$ be the region between by B_t and B_{t+1}. See Figure 2 for an example. The volume of L_t^- is bounded by

$$\int_{(2k)^d t/n}^{1} \int_{(2k)^d t/n}^{1} \cdots \int_{(2k)^d t/n}^{1} \frac{(2k)^d t}{n x_1 x_2 \cdots x_{d-1}} dx_1 dx_2 \cdots dx_{d-1} + \frac{d(2k)^d t}{n} = O\left(\frac{k^d t \log^{d-1}(n/k^d)}{n}\right).$$

Since $L_t \subseteq L_{t+1}^-$, this implies that $|L_t| = O\left(\frac{k^d t \log^{d-1}(n/k^d)}{n}\right)$. Next consider a point $p \in L_t$ and divide D_p into a uniform grid with $(2k)^d$ cells by dividing each side of D_p into $2k$ equal intervals. Each cell of this grid has volume at least t/n and thus contains at least t points in expectation. Thus, using the Chernoff bound, any of these grid cells is empty with probability less than $e^{-t/4}$. For p to be on one of the first k minimal layers, at least k of the $2k$ cells on the diagonal of the grid must be empty, which happens with probability less than $2e^{-t/4}$, by Markov's inequality. Thus, any point in S belongs to L_t and to one of the first k layers with probability $O(\frac{k^d t \log^{d-1}(n/k^d)}{n e^{t/4}})$. The expected number of points in S that belong to L_t and to one of the first k layers is thus $O(\frac{k^d t \log^{d-1}(n/k^d)}{e^{t/4}})$.

Since $\sum_{t=0}^{\infty} \frac{t}{e^{t/4}} = O(1)$, the expected number of points on the first k minimal layers is thus $O(k^d \log^{d-1}(n/k^d))$.

For an arbitrary continuous CI distribution $\mathcal{D}$, let P_i be the cumulative distribution function of the probability distribution of the ith coordinate. Then, for any point set S drawn from $\mathcal{D}$, the mapping $\phi : (x_1, x_2, \ldots, x_d) \mapsto (P_1(x_1), P_2(x_2), \ldots, P_d(x_d))$ produces a point set S' drawn uniformly at random from the unit hypercube and $p \in S$ dominates $q \in S$ if and only if $\phi(p)$ dominates $\phi(q)$. Thus, the total expected size of the first k maximal layers, and thus the expected size of the kth maximal layer, of S is $O(k^d \log^{d-1}(n/k^d))$.

5.2 Convex Layers

By the argument in the previous paragraph, it suffices to prove that the first k convex layers of a point set S drawn uniformly at random from the unit hypercube have expected size $O(k^d \log^{d-1}(n/k^d))$. Let o be the center point of the unit hypercube. Point o splits the unit hypercube into 2^d orthants $\mathcal{O}^\sigma$ with side length $1/2$ for $\sigma \in \{+1, -1\}^d$. More precisely, $\mathcal{O}^\sigma$ is the set of all points in the unit hypercube that σ-dominate o. We prove that the expected number of points in each orthant $\mathcal{O}^\sigma$ that belong to the first k convex layers is $O((2k)^d \log^{d-1}(n/k^d))$. Summing over all 2^d orthants, we obtain a bound of $O((4k)^d \log^{d-1}(n/k^d)) = O(k^d \log^{d-1}(n/k^d))$ on the expected size of the first k convex layers of S. W.l.o.g. consider the orthant $\mathcal{O}^{\mathbf{-1}}$; the argument for any other orthant $\mathcal{O}^\sigma$ is analogous after negating the point coordinates of S in all dimensions where σ and $-\mathbf{1}$ differ.

We define sets $L_1, L_2, \ldots$ as in Section 5.1. The key of our proof is to show that any point in $L_t \cap \mathcal{O}^{\mathbf{-1}}$ belongs to one of the first k convex layers with probability at most $2^{d+1} e^{-t/4}$. Since $L_t \cap \mathcal{O}^{\mathbf{-1}} \subseteq L_t$, the exact same calculation as in Section 5.1 then shows that the expected number of points in $L_t \cap \mathcal{O}^{\mathbf{-1}}$ that belong to the first k convex layers is $O\Big(\frac{(2k)^d t \log^{d-1}(n/k^d)}{e^{t/4}}\Big)$ and again, since $\sum_{t=0}^{\infty} \frac{t}{e^{t/4}} = O(1)$, the expected number of points in $\mathcal{O}^{\mathbf{-1}}$ that belong to the first k convex layers is $O((2k)^d \log^{d-1}(n/k^d))$, as claimed.

So consider a point $p \in L_t \cap \mathcal{O}^{\mathbf{-1}}$. Let D_p^σ be the part of the unit hypercube σ-dominated by p. We divide each hyperrectangle D_p^σ into a uniform grid with $(2k)^d$ cells by dividing each side of D_p^σ into $2k$ equal intervals. Consider the diagonal of D_p^σ connecting p with the opposite corner of D_p^σ and let $H_1^\sigma, H_2^\sigma, \ldots, H_{2k}^\sigma$ be the grid cells intersected by this diagonal, ordered by increasing distance from p. The argument in Section 5.1 shows that $P(S \cap H_i^{\mathbf{1}} = \emptyset) < e^{-t/4}$ for all $1 \le i \le 2k$, since $D_p = D_p^{\mathbf{1}}$. Since $p \in \mathcal{O}^{\mathbf{-1}}$, we have $|H_i^\sigma| \ge |H_i^{\mathbf{1}}|$ for all i and all σ. Thus, $P(S \cap H_i^\sigma = \emptyset) \le P(S \cap H_i^{\mathbf{1}}) < e^{-t/4}$. Now, if there exist k indices $1 \le i_1 < i_2 < \cdots < i_k \le 2k$ such that, for all $1 \le j \le k$ and all $\sigma \in \{-1, +1\}^d$, $H_{i_j}^\sigma \cap S \neq \emptyset$, then p is not on the kth convex layer. Thus, for p to be on the kth convex layer, there have to be at least k indices $1 \le i'_1 < i'_2 < \cdots < i'_k \le 2k$ and sign vectors $\sigma_1, \sigma_2, \ldots, \sigma_k$ such that $S \cap H_{i'_j}^{\sigma_j} = \emptyset$ for all $1 \le j \le k$. For any fixed index $1 \le i \le 2k$, the probability that there exists a sign vector σ_i such that

$S \cap H_i^{\sigma_i} = \emptyset$ is less than $2^d e^{-t/4}$, since $P(S \cap H_i^{\sigma} = \emptyset) < e^{-t/4}$ for any fixed i and σ. Thus, the expected number of indices i such that $S \cap H_i^{\sigma_i} = \emptyset$ for some sign vector σ_i is less than $2^{d+1} k e^{-t/4}$. By Markov's inequality, the probability that there are at least k such indices is thus less than $2^{d+1} e^{-t/4}$. Since this is an upper bound on the probability that p belongs to the first k convex layers, this finishes the proof of Theorem 3 for convex layers.

To obtain the same bound for the expected size of the first k orthant layers, observe that point p does not belong to the first k orthant layers if there exist k indices $1 \leq i_1 < i_2 < \cdots < i_k \leq 2k$ such that, for all $1 \leq j \leq k$ and all $\sigma \in \{-1, +1\}^d$, $H_{i_j}^{\sigma} \cap S \neq \emptyset$. Since this is the same condition we used to bound the size of the first k convex layers, the above argument also shows that the first k orthant layers of S have expected size $O(k^d \log^{d-1}(n/k^d))$.

References

1. J. L. Bentley, K. L. Clarkson, and D. B. Levine. Fast linear expected-time algorithms for computing maxima and convex hulls. *Algorithmica*, 9(2):168–183, 1993.
2. J. L. Bentley, H. T. Kung, M. Schkolnick, and C. D. Thompson. On the average number of maxima in a set of vectors and applications. *J. ACM*, 25(4):536–543, 1978.
3. M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 3rd edition, 2008.
4. H. Blunck and J. Vahrenhold. In-place algorithms for computing (layers of) maxima. In *Proc. 10th Scand. Workshop on Algorithm Theory*, pages 363–374, 2006.
5. S. Börzsönyi, D. Kossmann, and K. Stocker. The skyline operator. In *Proc. 17th International Conference on Data Engineering*, pages 421–430, 2001.
6. A. L. Buchsbaum and M. T. Goodrich. Three-dimensional layers of maxima. *Algorithmica*, 39(4):275–286, 2004.
7. B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10(4):377–409, 1993.
8. B. Chazelle. On the convex layers of a planar set. *IEEE Trans. Inf. Theor.*, 31(4):509–517, 2006.
9. K. Dalal. Counting the onion. *Random Struct. Algorithms*, 24(2):155–165, 2004.
10. A. Frieze. On the length of the longest monotone subsequence in a random permutation. *The Annals of Applied Probability*, 1(2):301–305, 1991.
11. M. J. Golin. A provably fast linear-expected-time maxima-finding algorithm. *Algorithmica*, 11(6):501–524, 1994.
12. H. T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, 1975.
13. J. Matoušek and P. Plecháč. On functional separately convex hulls. *Discrete & Computational Geometry*, 19(1):105–130, Jan 1998.
14. R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1995.
15. F. Nielsen. Output-sensitive peeling of convex and maximal layers. *Information Processing Letters*, 59(5):255–259, 1996.
16. A. Okabe, B. Boots, K. Sugihara, S. Chiu, and D. G. Kendall. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams, 2nd Edition*. John Wiley & Sons, Inc., 2008.