

MINIMUM CONVEX PARTITIONS AND MAXIMUM EMPTY POLYTOPES*

Adrian Dumitrescu,[†] Sarel Har-Peled,[‡] and Csaba D. Tóth[§]

ABSTRACT. Let S be a set of n points in \mathbb{R}^d . A Steiner convex partition is a tiling of $\text{conv}(S)$ with empty convex bodies. For every integer d , we show that S admits a Steiner convex partition with at most $\lceil (n-1)/d \rceil$ tiles. This bound is the best possible for points in general position in the plane, and it is best possible apart from constant factors in every fixed dimension $d \geq 3$. We also give the first constant-factor approximation algorithm for computing a minimum Steiner convex partition of a planar point set in general position.

Establishing a tight lower bound for the maximum volume of a tile in a Steiner convex partition of any n points in the unit cube is equivalent to a famous problem of Danzer and Rogers. It is conjectured that the volume of the largest tile is $\omega(1/n)$. Here we give a $(1 - \varepsilon)$ -approximation algorithm for computing the maximum volume of an empty convex body amidst n given points in the d -dimensional unit box $[0, 1]^d$.

Keywords: Steiner convex partition, Horton set, epsilon-net, lattice polytope, approximation algorithm.

1 Introduction

Let S be a set of $n \geq d + 1$ points in \mathbb{R}^d , $d \geq 2$. A convex body C is *empty* if its interior is disjoint from S . A *convex partition* of S is a partition of the convex hull $\text{conv}(S)$ into empty convex bodies (called *tiles*) such that the vertices of the tiles are in S . In a *Steiner convex partition* of S the vertices of the tiles are arbitrary: they can be points in S or *Steiner points*. For instance, any triangulation of S is a convex partitions of S , where the convex bodies are simplices, and so $\text{conv}(S)$ can be always partitioned into $O(n^{\lfloor d/2 \rfloor})$ empty convex tiles [13].

In this paper, we study the minimum number of tiles that a Steiner convex partition of every n points in \mathbb{R}^d admits, and the maximum volume of a single tile for a given point set. The research is motivated by a longstanding open problem by Danzer and Rogers [2,

*A preliminary version of this paper appeared in the *Proceedings of the 13th Scandinavian Symposium and Workshops on Algorithm Theory, Helsinki, Finland, July, 2012*. The first author was partially supported by NSF grant DMS-1001667, the second author was partially supported by NSF AF award CCF-0915984, and the third author was partially supported by NSERC grant RGPIN 35586.

[†]Department of Computer Science, University of Wisconsin–Milwaukee, WI 53201-0784, USA. dumitres@uwm.edu

[‡]Department of Computer Science, University of Illinois at Urbana–Champaign, Urbana, IL 61801-2302, USA. sariel@cs.uiuc.edu

[§]Department of Mathematics, California State University, Northridge, Los Angeles, CA 91330-8313; and Department of Computer Science, Tufts University, Medford, MA 02155, USA. cdtoth@acm.org

6, 10, 19, 35]: What is the maximum volume of an empty convex body $C \subset [0, 1]^d$ that can be found amidst any set $S \subset [0, 1]^d$ of n points in a unit cube? The current best bounds are $\Omega(1/n)$ and $O(\log n/n)$, respectively (for a fixed d). The lower bound, for instance, can be deduced by decomposing the unit cube by n parallel hyperplanes, each containing at least one point, into at most $n + 1$ empty convex bodies. The upper bound is tight apart from constant factors for n randomly and uniformly distributed points in the unit cube. It is suspected that the largest volume is $\omega(1/n)$ in any dimension $d \geq 2$, i.e., the ratio between this volume and $1/n$ tends to ∞ .

For a convex body C in \mathbb{R}^d , denote by $\text{vol}(C)$ the Lebesgue measure of C , i.e., its area when $d = 2$, or its volume when $d \geq 3$.

Minimum number of tiles in a convex partition. A *minimum convex partition* of S is a convex partition of S with a minimum number of tiles. Denote this number by $f_d(S)$. Further define (by slightly abusing notation)

$$f_d(n) = \max\{f_d(S) : S \subset \mathbb{R}^d, |S| = n\}.$$

Similarly define a *minimum Steiner convex partition* of S as one with a minimum number of tiles and let $g_d(S)$ denote this number. We also define

$$g_d(n) = \max\{g_d(S) : S \subset \mathbb{R}^d, |S| = n\}.$$

There has been substantial work on estimating $f_2(n)$, and computing $f_2(S)$ for a given set S in the plane. It has been shown successively that $f_2(n) \leq \frac{10n-18}{7}$ by Neumann-Lara et al. [34], $f_2(n) \leq \frac{15n-24}{11}$ by Knauer and Spillner [29], and $f_2(n) \leq \frac{4n-6}{3}$ for $n \geq 6$ by Sakai and Urrutia [36]. From the other direction, García-López and Nicolás [20] proved that $f_2(n) \geq \frac{12n-22}{11}$, for $n \geq 4$, thereby improving an earlier lower bound $f_2(n) \geq n + 2$ by Aichholzer and Krasser [1]. Knauer and Spillner [29] have also obtained a $\frac{30}{11}$ -factor approximation algorithm for computing a minimum convex partition for a given set $S \subset \mathbb{R}^2$, no three of which are collinear. There are also a few exact algorithms, including three fixed-parameter algorithms [17, 21, 38].

The state of affairs is much different in regard to Steiner convex partitions. As pointed out in [15], no corresponding results are known for the variant with Steiner points. Here we take the first steps in this direction, and obtain the following results.

Theorem 1. *For $n \geq d + 1$, we have $g_d(n) \leq \lceil \frac{n-1}{d} \rceil$. For $d = 2$, this bound is the best possible, that is, $g_2(n) = \lceil (n-1)/2 \rceil$; and for every fixed $d \geq 2$, we have $g_d(n) = \Omega(n)$.*

We say that a set of points in \mathbb{R}^d is in *general position* if every k -dimensional affine subspace contains at most $k + 1$ points for $0 \leq k < d$. We show that in the plane every Steiner convex partition for n points in general position, i of which lie in the interior of the convex hull, has $\Omega(i)$ tiles. This leads to a simple constant-factor approximation algorithm.

Theorem 2. *Given a set S of n points in general position in the plane, a ratio 3 approximation of a minimum Steiner convex partition of S can be computed in $O(n \log n)$ time.*

The *average* volume of a tile in a Steiner convex partition of n points in the unit cube $[0, 1]^d$ is an obvious lower bound for the maximum possible volume of a tile, and for the maximum volume of any empty convex body $C \subset [0, 1]^d$. The lower bound $g_d(n) = \Omega(n)$ in Theorem 1 shows that the average volume of a tile is $O(1/n)$ in some instances, where the constant of proportionality depends only on the dimension. This implies that a simple “averaging” argument is not a viable avenue for finding a solution to the problem of Danzer and Rogers.

Maximum empty polytope among n points in a unit cube. In the second part of the paper, we consider the following problem: Given a set of n points in a rectangular box B in \mathbb{R}^d , find a maximum-volume empty convex body $C \subset B$. Since the ratio between volumes is invariant under affine transformations, we may assume without loss of generality that $B = [0, 1]^d$. We therefore have the problem of computing a maximum volume empty convex body $C \subset [0, 1]^d$ for a set of n points in $[0, 1]^d$. It can be argued that the maximum volume empty convex body is a polytope, however, the number and location of its vertices is unknown and this represents the main difficulty. For $d = 2$ there is a polynomial-time exact algorithm (see Section 5) while for $d \geq 3$ we are not aware of any exact algorithm. Thus the problem of finding faster approximations naturally suggests itself.

There exist exact algorithms for some related problems. Eppstein et al. [16] find the maximum area empty convex k -gon with vertices among n points in $O(kn^3)$ time, if it exists. As a byproduct, a maximum area empty convex polygon with vertices among n given points can be computed exactly in $O(n^4)$ time with their dynamic programming algorithm. The running time was subsequently improved to $O(n^3 \log n)$ by Fischer [18] and then to $O(n^3)$ by Bautista-Santiago et al. [9].

By John’s ellipsoid theorem [32], the maximum volume empty ellipsoid in $[0, 1]^d$ gives a $1/d^d$ -approximation. Here we present a $(1 - \varepsilon)$ -approximation for a maximum volume empty convex body C_{opt} by first guessing a good approximation of the bounding hyperrectangle of C_{opt} of minimum volume, and then finding a sufficiently close approximation of C_{opt} inside it. We obtain the following two approximation algorithms. The planar algorithm runs in quadratic time in n , however, the running time degrades with the dimension.

Theorem 3. *Given a set S of n points in $[0, 1]^2$ and parameter $\varepsilon > 0$, one can compute an empty convex body $C \subseteq [0, 1]^2$, such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. The running time of the algorithm is $O(\varepsilon^{-6}n^2)$.*

Theorem 4. *Given a set S of n points in $[0, 1]^d$, $d \geq 3$, and a parameter $\varepsilon > 0$, one can compute an empty convex body $C \subseteq [0, 1]^d$, such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. The running time of the algorithm is*

$$\exp\left(O\left(\varepsilon^{-d(d-1)/(d+1)} \log \varepsilon^{-1}\right)\right) n^{1+d(d-1)/2} \log^d n.$$

As far as the problem of Danzer and Rogers is concerned, one need not consider convex sets—it suffices to consider simplices—and for simplices the problems considered are much simpler. Specifically, every convex body C in \mathbb{R}^d , $d \geq 2$, contains a simplex T

of volume $\text{vol}(T) \geq \text{vol}(C) / (d+2)^d$ [30]. That is, for fixed d , the largest empty simplex amidst n points in the unit cube $[0, 1]^d$ yields a constant-factor approximation of the largest volume convex body (polytope) amidst the same n points. Consequently, the asymptotic dependencies on n of the volumes of the largest empty simplex and convex body are the same. For $d = 2$ there is a polynomial-time exact algorithm for computing the largest empty triangle amidst n points in $[0, 1]^2$ (see Section 5) while for $d \geq 3$ we are not aware of any exact algorithm for computing the largest empty simplex amidst n points in $[0, 1]^d$.

Related work. Decomposing polygonal domains into convex sub-polygons has been also studied extensively. We refer to the article by Keil [26] for a survey of results up to the year 2000. For instance, when the polygon may contain holes, obtaining a minimum convex partition is NP-hard, regardless of whether Steiner points are allowed. For polygons without holes, Chazelle and Dobkin [12] obtained an $O(n + r^3)$ time algorithm for the problem of decomposing a polygon with n vertices, r of which are reflex, into convex parts, with Steiner points permitted. Keil [26] notes that although there are an infinite number of possible locations for the Steiner points, a dynamic programming approach is amenable to obtain an exact (optimal) solution; see also [27, 37].

Fevens et al. [17] designed a polynomial time algorithm for computing a minimum convex partition for a given set of n points in the plane if the points are arranged on a constant number of convex layers. The problem of minimizing the total Euclidean length of the edges of a convex partition has been also considered. Grantson and Levcopoulos [20], and Spillner [38] proved that the shortest convex partition and Steiner convex partition problems are fixed parameter tractable, where the parameter is the number of points of P lying in the interior of $\text{conv}(P)$. Dumitrescu and Tóth [15] proved that every set of n points in \mathbb{R}^2 admits a Steiner convex partition which is at most $O(\log n / \log \log n)$ times longer than the minimum spanning tree, and this bound cannot be improved. Without Steiner points, the best upper bound for the ratio of the minimum length of a convex partition and the length of a minimum spanning tree (MST) is $O(n)$ [28].

A largest area convex polygon contained in a given (non-convex) polygon with n vertices can be found by the algorithm of Chang and Yap [11] in $O(n^7)$ time. The problem is known as the *potato-peeling problem*. On the other hand, a largest area triangle contained in a simple polygon with n vertices, can be found by the algorithm of Melissaratos and Souvaine [33] in $O(n^4)$ time. Hall-Holt et al. [22] compute a constant approximation in time $O(n \log n)$. The same authors show how to compute a $(1 - \varepsilon)$ -approximation of the largest *fat* triangle inside a simple polygon (if it exists) in time $O(n)$. Given a triangulated polygon (with possible holes) with n vertices, Aronov et al. [4] compute the largest area convex polygon respecting the triangulation edges in $O(n^2)$ time.

For finding a maximum volume empty axis-parallel box amidst n points in $[0, 1]^d$, Backer and Keil [5] reported an algorithm with worst-case running time of $O(n^d \log^{d-2} n)$. An empty axis-aligned box whose volume is at least $(1 - \varepsilon)$ of the maximum can be computed in $O\left(\left(\frac{8ed}{\varepsilon^2}\right)^d n \log^d n\right)$ time by the algorithm of Dumitrescu and Jiang [14].

Lawrence and Morris [31] studied the minimum integer $k_d(n)$ such that the complement $\mathbb{R}^d \setminus S$ of any n -element set $S \subset \mathbb{R}^d$, not all in a hyperplane, can be covered by

$k_d(n)$ convex sets. They prove $k_d(n) = \Omega(\log n / d \log \log n)$. It is known that covering the complement of n uniformly distributed points in $[0, 1]^d$ requires $\Omega(n / d \log n)$ convex sets, which follows from the upper bound in the problem of Danzer and Rogers.

2 Combinatorial bounds

In this section we prove Theorem 1. We start with the upper bound. The following simple algorithm returns a Steiner convex partition with at most $\lceil (n-1)/d \rceil$ tiles for any n points in \mathbb{R}^d .

Algorithm **A1**:

STEP 1. Compute the convex hull $R \leftarrow \text{conv}(S)$ of S . Let $A \subseteq S$ be the set of hull vertices, and let $B = S \setminus A$ denote the remaining points.

STEP 2. Compute $\text{conv}(B)$, and let H be the supporting hyperplane of an arbitrary $(d-1)$ -dimensional face of $\text{conv}(B)$. Denote by H^+ the halfspace that contains B , and $H^- = \mathbb{R}^d \setminus H^+$. The hyperplane H contains d points of B , and it decomposes R into two convex bodies: $R \cap H^-$ is empty and $R \cap H^+$ contains all points in $B \setminus H$. Update $B \leftarrow B \setminus H$ and $R \leftarrow R \cap H^+$.

STEP 3. Repeat STEP 2 with the new values of R and B until B is the empty set. (If $|B| < d$, then any supporting hyperplane of B completes the partition.)




Figure 1: Steiner convex partitions with Steiner points drawn as hollow circles. Left: A Steiner convex partition of a set of 13 points. Middle: A Steiner partition of a set of 12 points into three tiles. Right: A Steiner partition of the same set of 12 points into 4 tiles, generated by Algorithm **A1** (the labels reflect the order of execution).

It is obvious that the algorithm generates a Steiner convex partition of S . An illustration of Algorithm **A1** on a small planar example appears in Figure 1 (right). Let h and i denote the number of hull and interior points of S , respectively, so that $n = h + i$. Each hyperplane used by the algorithm removes d interior points of S (with the possible exception of the last round if i is not a multiple of d). Hence the number of convex tiles is $1 + \lceil i/d \rceil$, and we have $1 + \lceil i/d \rceil = \lceil (i+d)/d \rceil \leq \lceil (n-1)/d \rceil$, as required.

Lower bound in the plane. A matching lower bound in the plane is given by the following construction. For $n \geq 3$, let $S = A \cup B$, where A is a set of 3 non-collinear points in the plane, and B is a set of $n-3$ points that form a regular $(n-3)$ -gon in the interior

of $\text{conv}(A)$, so that $\text{conv}(S) = \text{conv}(A)$ is a triangle. If $n = 3$, then $\text{conv}(S)$ is an empty triangle, and $g_2(S) = 1 = \lceil (n-1)/2 \rceil$. If $4 \leq n \leq 5$, S is not in convex position, and so $g_2(S) \geq 2 = \lceil (n-1)/2 \rceil$. Suppose now that $n \geq 6$.

Consider an arbitrary convex partition of S . Let o be a point in the interior of $\text{conv}(B)$ such that the lines os , $s \in S$, do not contain any edges of the tiles. Refer to Figure 2. For each point $s \in B$, choose a *reference point* $r(s) \in \mathbb{R}^2$ on the ray \vec{os} in $\text{conv}(A) \setminus \text{conv}(B)$ sufficiently close to point s , and lying in the interior of a tile. Note that the convex tile containing o cannot contain any reference points. We claim that any tile contains at most 2 reference points. This immediately implies $g_2(S) \geq 1 + \lceil (n-3)/2 \rceil = \lceil (n-1)/2 \rceil$.

Suppose, to the contrary, that a tile τ contains 3 reference points r_1, r_2, r_3 , corresponding to the points s_1, s_2, s_3 . Refer to Figure 2. Note that o cannot be in the interior of




Figure 2: Lower bound construction in \mathbb{R}^2 .

τ , otherwise τ would contain all points s_1, s_2, s_3 in its interior. Hence $\text{conv}(\{o, s_1, s_2, s_3\})$ is a quadrilateral, and $\text{conv}(\{o, r_1, r_2, r_3\})$ is also a quadrilateral, since the reference points are sufficiently close to the corresponding points in B . We may assume w.l.o.g. that vertices of $\text{conv}(\{o, s_1, s_2, s_3\})$ are o, s_1, s_2, s_3 in counterclockwise order. Then s_2 lies in the interior of $\text{conv}(\{r_1, r_2, r_3\})$. Hence the tile containing r_1, r_2 , and r_3 , must contain point s_2 in its interior, a contradiction. We conclude that every tile τ contains at most 2 reference points, as required.

Lower bounds for $d \geq 3$. A similar construction works in for any $d \geq 2$, but the lower bound no longer matches the upper bound $g_d(n) \leq \lceil (n-1)/d \rceil$ for $d \geq 3$.

Recall that a *Horton set* [25] is a set S of n points in the plane such that the convex hull of any 7 points is non-empty. Valtr [39] generalized Horton sets to \mathbb{R}^d . For every $d \in \mathbb{N}$, there exists a minimal integer $h(d)$ with the property that for every $n \in \mathbb{N}$ there is a set S of n points in general position in \mathbb{R}^d such that the convex hull of any $h(d) + 1$ points in S is non-empty. It is known that $h(2) = 6$, and Valtr proved that $h(3) \leq 22$, and in general that $h(d) \leq 2^{d-1}(N(d-1) + 1)$, where $N(k)$ is the product of the first k primes.

We construct a set S of $n \geq d + 1$ points in \mathbb{R}^d as follows. Let $S = A \cup B$, where A is a set of $d + 1$ points in general position in \mathbb{R}^d , and B is a generalized Horton set of

$n - (d + 1)$ points in the interior of $\text{conv}(A)$, such that the interior of any $h(d) + 1$ points from B contains some point in B .

Consider an arbitrary Steiner convex partition of S . Every point $b \in B$ is in the interior of $\text{conv}(S)$, and so it lies on the boundary of at least 2 convex tiles. For each $b \in B$, place two *reference points* in the interiors of 2 distinct tiles incident to b . Every tile contains at most $h(d)$ reference points. Indeed, if a tile contains $h(d) + 1$ reference points, then it is incident to $h(d) + 1$ points in B , and some point of B lies in the interior of the convex hull of these points, a contradiction.

There are $2(n - d - 1)$ reference points, and every tile contains at most $h(d)$ of them. So the number of tiles is at least $\lceil 2(n - d - 1)/h(d) \rceil$. Consequently, for every fixed $d \geq 2$, we have $g_d(n) = \Omega(n)$.

3 Approximating the minimum Steiner convex partition in \mathbb{R}^2

In this section we prove Theorem 2 by showing that our simple-minded algorithm **A1** from Section 2 achieves a constant-factor approximation in the plane if the points in S are in general position.

Approximation ratio. Recall that algorithm **A1** computes a Steiner convex partition of $\text{conv}(S)$ into at most $1 + \lceil i/2 \rceil$ parts, where i stands for the number of interior points of S .

If $i = 0$, the algorithm computes an optimal partition, i.e., $\text{ALG} = \text{OPT} = 1$. Assume now that $i \geq 1$. Consider an optimal Steiner convex partition Π of S with OPT tiles. We construct a planar multigraph $G = (V, E)$ as follows. The *faces* of G are the convex tiles and the exterior of $\text{conv}(S)$ (the outer face). The *vertices* V are the points in the plane incident to at least 3 faces (counting the outer face as well). Since $i \geq 1$, G is non-empty and we have $|V| \geq 2$. Each *edge* in E is a Jordan arc on the common boundary of two faces. An edge between two bounded faces is a straight line segment, and so it contains at most two interior points of S . An edge between the outer face and a bounded face is a convex arc, containing hull points from S . Double edges are possible if two vertices of the outer face are connected by a straight line edge and a curve edge along the boundary—in this case these two parallel edges bound a convex face. No loops are possible in G . Since Π is a convex partition, G is connected.

Let v , e , and f , respectively, denote the number of vertices, edges, and bounded (convex) faces of G ; in particular, $f = \text{OPT}$. By Euler's formula for planar multigraphs, we have $v - e + f = 1$, that is, $f = e - v + 1$. By construction, each vertex of G is incident to at least 3 edges, and every edge is incident to two vertices. Therefore, $3v \leq 2e$, or $v \leq 2e/3$. Consequently, $f = e - v + 1 \geq e - 2e/3 + 1 = e/3 + 1$. Since S is in general position, each straight-line edge of G contains at most 2 interior points from S . Curve edges along the boundary do not contain interior points. Hence each edge in E is incident to at most two interior points in S , thus $i \leq 2e$. Substituting this into the previous inequality on f yields $\text{OPT} = f \geq e/3 + 1 \geq i/6 + 1$. Comparing this lower bound with the upper bound $\text{ALG} \leq \lceil i/2 \rceil + 1$, we conclude that $\frac{\text{ALG}}{\text{OPT}} \leq \frac{\lceil i/2 \rceil + 1}{i/6 + 1} \leq 3 \frac{i + 3}{i + 6} < 3$, and the approximation ratio of 3 follows.

Tightness of the approximation ratio. We first show that the ratio 3 established above is tight for Algorithm **A1**. We construct a planar point set S as follows. Refer to Figure 3 (left). Consider a large (say, hexagonal) section of a hexagonal lattice. Place Steiner vertices at the lattice points, and place two points in S on each lattice edge. Slightly perturb the lattice, and add a few more points in S near the boundary, and a few more Steiner points, so as to obtain a Steiner convex partition of S with no three points collinear. Denote by v , e , and f , the elements of the planar multigraph G as before. Since we consider a large lattice section, we have $v, e, f \rightarrow \infty$. We write $a \sim b$, whenever $a/b \rightarrow 1$. As before, we have $f + v = e + 1$, and since each non-boundary edge is shared by two convex faces, we have $e \sim 6f/2 = 3f$. By construction, $i \sim 2e \sim 6f$, hence $f \sim i/6$. Therefore the convex partition constructed above has $f \sim i/6$, while Algorithm **A1** constructs one with about $i/2$ faces. Letting $e \rightarrow \infty$, then $i \rightarrow \infty$, and the ratio ALG/OPT approaches 3 in the limit: $\text{ALG}/\text{OPT} \sim (i/2)/(i/6) = 3$.




Figure 3: Left: two points on each edge of a section of a perturbed hexagonal lattice in \mathbb{R}^2 , and four extra vertices of a bounding box. Right: Points in general position on a saddle surface in \mathbb{R}^3 .

Time analysis. Algorithm **A1** can be implemented to run in $O(n \log n)$ time for a set S of n points in the plane. We employ the semi-dynamic (delete only) convex hull data structure of Hershberger and Suri [24]. This data structure supports point deletion in $O(\log n)$ time, and uses $O(n)$ space and $O(n \log n)$ preprocessing time. We maintain the boundary of a convex polygon R in a binary search tree, a set $B \subset S$ of points lying in the interior of R , and the convex hull $\text{conv}(B)$ with the above semi-dynamic data structure [24]. Initially, $R = \text{conv}(S)$, which can be computed in $O(n \log n)$ time; and $B \subset S$ is the set of interior points. In each round of the algorithm, consider the supporting line H of an arbitrary edge e of $\text{conv}(B)$ such that B lies in the halfplane H^+ . The two intersection points of H with the boundary of R can be computed in $O(\log n)$ time. At the end of the round, we can update $B \leftarrow B \setminus H$ and $\text{conv}(B)$ in $O(k \log n)$ time, where k is the number of points removed from B ; and we can update $R \leftarrow R \cap H^+$ in $O(\log n)$ time. Every point is removed from B exactly once, and the number of rounds is at most $\lceil (n-3)/2 \rceil$, so the total update time is $O(n \log n)$ throughout the algorithm.

Remark. Interestingly enough, in dimensions 3 and higher, Algorithm **A1** does not give a constant-factor approximation. For every integer n , one can construct a set S of n points in general position in \mathbb{R}^3 such that $i = n - 4$ of them lie in the interior of $\text{conv}(S)$, but the minimum Steiner convex partition has only $O(\sqrt{n})$ tiles. In contrast, Algorithm **A1** computes a Steiner partition with $i/3 = (n - 4)/3$ convex tiles.

We first construct the convex tiles, and then describe the point set S . Specifically, S consists of 4 points of a large tetrahedron, and 3 points in general position on the common boundary of certain pairs of adjacent tiles.

Let $k = \lceil \sqrt{(n - 4)/3} \rceil$. Place $(k + 1)^2$ Steiner points $(a, b, a^2 - b^2)$ on the saddle surface $z = x^2 - y^2$ for pairs of integers $(a, b) \in \mathbb{Z}^2$, $- \lfloor k/2 \rfloor \leq a, b \leq \lceil k/2 \rceil$. The four points $\{(x, y, x^2 - y^2) : x \in \{a, a + 1\}, y \in \{b, b + 1\}\}$ form a parallelogram for every $(a, b) \in \mathbb{Z}^2$, $- \lfloor k/2 \rfloor \leq a, b \leq \lceil k/2 \rceil - 1$. Refer to Figure 3 (right). These parallelograms form a terrain over the region $\{(x, y) : - \lfloor k/2 \rfloor \leq x, y \leq \lceil k/2 \rceil\}$. Note that no two parallelograms are coplanar. Subdivide the space *below* this terrain by vertical planes $x = a$, $- \lfloor k/2 \rfloor \leq a \leq \lceil k/2 \rceil$. Similarly, subdivide the space *above* this terrain by planes $y = b$, $- \lfloor k/2 \rfloor \leq b \leq \lceil k/2 \rceil$. We obtain $2k$ interior-disjoint convex regions, k above and k below the terrain, such that the common boundary of a region above and a region below is a parallelogram of the terrain. The points in \mathbb{R}^3 that do not lie above or below the terrain can be covered by 4 convex wedges.

Enclose the terrain in a sufficiently large tetrahedron T . Clip the $2k$ convex regions and the 4 wedges into the interior of T . These $2k + 4$ convex bodies tile T . Choose 3 noncollinear points of S in each of the k^2 parallelograms, such that no 4 points are coplanar and no 2 are collinear with vertices of T . Let the point set S be the set of 4 vertices of the large tetrahedron T and the $3k^2$ points selected from the parallelograms.

4 Approximating the maximum empty convex body

Let S be a set of points in the unit cube $[0, 1]^d \subseteq \mathbb{R}^d$. Our task is to approximate the largest convex body $C \subseteq [0, 1]^d$ that contains no points of S in its interior. Let $C_{\text{opt}} = C_{\text{opt}}(S)$ denote this body.

4.1 Approximation by the discrete hull

In the following, assume that $m > 0$ is some integer, and consider the grid point set

$$\mathcal{G}(m) = \left\{ (i_1, \dots, i_d)/m \mid i_1, \dots, i_d \in \{0, 1, \dots, m\} \right\}.$$

Let $S \subseteq [0, 1]^d$ be a point set, and let C_{opt} be the corresponding largest empty convex body in $[0, 1]^d$. Given a grid $\mathcal{G}(m)$, we call $\text{conv}(C_{\text{opt}} \cap \mathcal{G}(m))$ the *discrete hull* of C_{opt} [23]. We need the following easy lemma.

Lemma 1. *Let $C \subseteq [0, 1]^d$ be a convex body and $D = \text{conv}(C \cap \mathcal{G}(m))$. Then we have $\text{vol}(C) - \text{vol}(D) = O(1/m)$, where the constant of proportionality depends only on d .*

Proof. Consider a point $p \in C$, and the cube $p + [-2, 2]^d/m$ centered at p with side length $4/m$. If this cube is contained in C , then all grid points of the grid cell of $\mathcal{G}(m)$ containing p are in C , and p lies in D . Therefore, for every point $p \in C \setminus D$, the cube $p + [-2, 2]^d/m$ is not contained in C . By convexity, at least one of the vertices of the cube $p + [-2, 2]^d/m$ lies outside of C . Therefore, the distance from p to the boundary of C is at most the distance from p to a corner of this cube, which is $2\sqrt{d}/m$.

It follows that all the points in the corridor $C \setminus D$ are at distance at most $2\sqrt{d}/m$ from the boundary of C . The volume of the boundary of C is bounded from above by the volume of the boundary of the unit cube, namely $2d$. As such, the volume of this corridor is $\text{vol}(\partial C) 2\sqrt{d}/m \leq (2d)(2\sqrt{d}/m) = O(d^{3/2}/m)$. For a fixed d , this is $O(1/m)$, as claimed. \square

Lemma 1 implies that if $\text{vol}(C_{\text{opt}}) \geq \rho$, in order to obtain a $(1 - \varepsilon)$ -approximation, we can concentrate our search on convex polytopes that have their vertices at grid points in $\mathcal{G}(m)$, where $m = O(1/(\varepsilon\rho))$. If ρ is a constant, then the maximum volume empty lattice polytope in $\mathcal{G}(m)$ with $m = O(1/\varepsilon)$ is an $(1 - \varepsilon)$ -approximation for C_{opt} . However, for arbitrary $\text{vol}(C_{\text{opt}}) = \Omega(1/n)$, a much finer grid would be necessary to achieve this approximation.

4.2 An initial brute force approach

In this section we present approximation algorithms (for all d) relying on Lemma 1 alone, approximating the maximum volume empty polytope by a lattice polytope in a sufficiently fine lattice (grid). We shall refine our technique in Subsections 4.3 and 4.4.

For the plane, we take advantage of the existence of an efficient solution for a related search problem. Refining a natural dynamic programming approach by Eppstein et al. [16] and Fischer [18], Bautista-Santiago et al. [9] obtained the following result.

Lemma 2 ([9]). *Given a set S of m points and a set Q of $O(m)$ points in the plane, one can compute a convex polygon with the largest area with vertices in S that does not contain any point of Q in its interior in $O(m^3)$ time.*

Remark. The algorithm has the same running time if Q is a set of $O(m)$ forbidden rectangles.

The combination of Lemmas 1 and 2 readily yields an approximation algorithm for the plane, whose running time depends on $\text{vol}(C_{\text{opt}})$.

Lemma 3. *Given a set $S \subseteq [0, 1]^2$ of n points, such that $\text{vol}(C_{\text{opt}}) \geq \rho$, and a parameter $\varepsilon > 0$, one can compute an empty convex body $C \subseteq [0, 1]^2$ such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. The running time of the algorithm is $O(n + 1/(\varepsilon\rho)^6)$.*

Proof. Consider the grid $\mathcal{G}(m)$ with $m = O(1/(\varepsilon\rho))$. By Lemma 1 we can restrict our search to a grid polygon. Going a step further, we mark all the grid cells containing points of S as forbidden. Arguing as in Lemma 1, one can show that the area of the largest convex grid polygon avoiding the forbidden cells is at least $\text{vol}(C_{\text{opt}}) - c/m$, where c is a constant.

We now restrict our attention to the task of finding a largest polygon. We have a set Q of $O(m^2)$ grid points that might be used as vertices of the grid polygon, and a set of $O(m^2)$ grid cells that cannot intersect the interior of the computed polygon. By Lemma 2, a largest empty polygon can be found in $O(m^6)$ time. Setting $m = O(1/(\varepsilon\rho))$, we get an algorithm with overall running time $O(n + 1/(\varepsilon\rho)^6)$. \square

For dimensions $d \geq 3$, we are not aware of any analogue of the dynamic programming algorithm in Lemma 2. Instead, we use a brute force approach that enumerates all feasible subsets of a sufficiently fine grid.

Lemma 4. *Given a set $S \subseteq [0, 1]^d$ of n points such that $\text{vol}(C_{\text{opt}}) \geq \rho$, and a parameter $\varepsilon > 0$, one can compute an empty convex body $C \subseteq [0, 1]^d$, such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. The running time of the algorithm is $O(n) + \exp(O(m^{d(d-1)/(d+1)} \log m))$, where $m = O(1/(\varepsilon\rho))$ and d is fixed.*

Proof. Consider the grid $\mathcal{G}(m)$ with $m = O(1/(\varepsilon\rho))$. Let X be the set of vertices of all grid cells of $\mathcal{G}(m)$ that contain some point from S (i.e., 2^d vertices per cell). Note that $|X| = O(m^d)$. Andrews [3] proved that a convex lattice polytope of volume V has $O(V^{(d-1)/(d+1)})$ vertices. Hence a convex lattice polytope in $\mathcal{G}(m)$ has $O(m^{d(d-1)/(d+1)})$ vertices. By the well-known inequality $\sum_{i=0}^k \binom{n}{i} \leq (\frac{en}{k})^k$, the number of subsets of size $O(m^{d(d-1)/(d+1)})$ from $\mathcal{G}(m)$ is

$$\sum_{i=0}^{O(m^{d(d-1)/(d+1)})} \binom{m^d}{i} \leq \left(m^{2d/(d+1)}\right)^{O(m^{d(d-1)/(d+1)})} \leq \exp\left(O(m^{d(d-1)/(d+1)} \log m)\right).$$

For each such candidate subset G of size $O(m^{d(d-1)/(d+1)})$, test whether $\text{conv}(G)$ is empty of points from X . For each point in X , the containment test reduces to a linear program that can be solved in time polynomial in m . Returning the subset with the largest hull volume found yields the desired approximation. The runtime of the algorithm is $\exp(O(m^{d(d-1)/(d+1)} \log m))$. \square

Bárány and Vershik [8] proved that there are $\exp(O(m^{d(d-1)/(d+1)}))$ convex lattice polytopes in $\mathcal{G}(m)$. If the polytopes can also be enumerated in this time (as in the planar case [7]), then the runtime in Lemma 4 reduces accordingly.

4.3 A faster approximation in the plane

If C_{opt} is long and skinny (e.g., ρ is close to $1/n$), then the uniform grid $\mathcal{G}(m)$ we used in Lemmas 3 and 4 is unsuitable for finding a $(1 - \varepsilon)$ -approximation efficiently. Instead, we employ a rotated and stretched grid (an affine copy of $\mathcal{G}(m)$) that has similar orientation and aspect ratio as C_{opt} . This overcomes one of the main difficulties in obtaining a good approximation. Since we do not know the shape and orientation of C_{opt} , we guess these parameters via the minimum area rectangle containing C_{opt} .

Lemma 5. *Given a set $S \subseteq [0, 1]^2$ of n points such that $\text{vol}(C_{\text{opt}}) \geq \rho$, and a parameter $\varepsilon > 0$, one can compute an empty convex body $C \subseteq [0, 1]^2$ such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. The running time of the algorithm is $O(\rho^{-1}(n + \rho^{-1}\varepsilon^{-6}))$.*

Proof. The idea is to first guess a rectangle R that contains C_{opt} such that $\text{vol}(C_{\text{opt}})$ is at least a constant fraction of the area of $\text{vol}(R)$, and then to apply Lemma 3 to the rectangle R (as the unit square) to get the desired approximation.

Let B_0 be the minimum area rectangle (of arbitrary orientation) that contains C_{opt} ; see Figure 4 (left). We guess an approximate copy of B_0 . In particular, we guess the lengths of the two sides of B_0 (up to a factor of 2) and the orientation of B_0 (up to an angle of $O(1/n)$), and then try to position a scaled copy of the guessed rectangle so that that it fully contains C_{opt} .




Figure 4: Left: C_{opt} , a minimum area rectangle B_0 , $C_{\text{opt}} \subseteq B_0$, and a minimum area rectangle B_1 , $B_0 \subseteq B_1$, with canonical side lengths and the same orientation as B_0 . Right: Rectangle B_1 , a rotated copy B_2 with the closest canonical orientation, and a minimum area scaled copy kB_2 such that $B_1 \subseteq kB_2$.

Assume for convenience that $n \geq 10$. We now show that $\text{vol}(C_{\text{opt}}) \geq \sqrt{2}/n$, using Theorem 1. Augment the point set S with the four corners of the unit square $[0, 1]^2$ into a set of $n + 4$ points. By Theorem 1, the augmented point set has a Steiner convex partition into at most $g_2(n + 4) = \lceil \frac{n+4-1}{2} \rceil$ tiles. The area of the largest tile is at least that of the average tile in this partition, that is, $\text{vol}(C_{\text{opt}}) \geq 1/\lceil \frac{n+4-1}{2} \rceil \geq \frac{2}{n+4} \geq \frac{2}{\sqrt{2}n} = \frac{\sqrt{2}}{n}$, for $n \geq 10$. Therefore, we may assume that $\text{vol}(C_{\text{opt}}) \geq \rho \geq \sqrt{2}/n$.

Denote by a and b the lengths of the two sides of B_0 , where $a \leq b$. It is clear that $b \leq \sqrt{2}$, the diameter of the unit square. We also have $a = \text{vol}(B_0)/b \geq \text{vol}(C_{\text{opt}})/b \geq \sqrt{2}/(bn) \geq 1/n$, hence the aspect ratio of B_0 is $b/a \leq \sqrt{2}/a \leq \sqrt{2}n$.

Assume now that $2^{i-1}\rho \leq \text{vol}(C_{\text{opt}}) < 2^i\rho$ for some $i = 1, 2, \dots, \lceil \log_2(\sqrt{2}n) \rceil$. If we want to guess the aspect ratio of B_0 up to a factor of two, we need to consider only $O(\log \rho^{-1})$ possibilities. Indeed, we consider the *canonical aspect ratios* 2^j for $j = 0, \dots, \lceil \log_2(\sqrt{2}/\rho) \rceil - i$, and *canonical side lengths* $2^{(i+j)/2}\sqrt{\rho}$ and $2^{(i-j)/2}\sqrt{\rho}$. Let B_1 be a minimum area rectangle

with canonical side lengths and the same orientation as B_0 , so that $B_0 \subseteq B_1$.

The *orientation* of a rectangle is given by the angle between one side and the x -axis. We approximate the orientation of B_0 by *canonical orientations* $\alpha = r\pi/(5 \cdot 2^j)$, for $r = 0, 1, \dots, 5 \cdot 2^j - 1$. Let B_2 be a congruent copy of B_1 rotated clockwise to the nearest canonical orientation about the center of B_1 . We show that $B_1 \subset 2B_2$, i.e., a scaled copy of B_2 contains B_1 . Let $k \geq 1$ be the minimum scale factor such that $B_1 \subseteq kB_2$. Refer to Figure 4 (right). Denote by o the common center of B_1 and B_2 , let x be a vertex of B_1 on the boundary of kB_2 , and let y be the corresponding vertex of kB_2 . Clearly, $\sin(\angle xoy) \leq \pi/(5 \cdot 2^j)$ since we rotate by at most $\pi/(5 \cdot 2^j)$. The aspect ratio of the rectangle kB_2 is $\cot(\angle oxy) = 2^j$. Since $\angle oxy < \pi/4$, we have $\sin(\angle oxy) = \tan(\angle oxy) \cos(\angle oxy) \geq 2^{-j} \cos \frac{\pi}{4} = 2^{-j-1/2} > \pi/(5 \cdot 2^j)$. The law of sines yields $|ox| > |xy|$; and we have $|ox| + |xy| > |oy|$ by the triangle inequality. It follows that $|oy| < 2|ox|$, and so $k \leq 2$ suffices. Summing over all possible areas, canonical aspect ratios, and orientations, the number of possibilities is

$$\begin{aligned} \sum_{i=0}^{\lceil \log_2(\sqrt{2}/\rho) \rceil} \sum_{j=0}^{\lceil \log_2(\sqrt{2}/\rho) \rceil - i} 5 \cdot 2^j &\leq \sum_{i=0}^{\lceil \log_2(\sqrt{2}/\rho) \rceil} 10 \cdot 2^{\lceil \log_2(\sqrt{2}/\rho) \rceil - i} \\ &\leq 20 \cdot 2^{\lceil \log_2(\sqrt{2}/\rho) \rceil} = O(\rho^{-1}). \end{aligned}$$

So far we have guessed the canonical side lengths and orientation of B_2 , however, we do not know its location in the plane. If a translated copy $B_2 + v$ of B_2 intersects C_{opt} , then $3B_2 + v$ contains it, since $C_{\text{opt}} \subseteq B_0 \subseteq B_1 \subseteq 2B_2$. Consider an arbitrary tiling of the plane with translates of B_2 . By a packing argument, only $O(1/\rho)$ translates intersect the unit square $[0, 1]^2$. One of these translates, say $B_2 + v$, intersects C_{opt} , and hence the rectangle $R = 3B_2 + v$ contains C_{opt} .

We can apply Lemma 3 to the rectangle R (as the unit square) to get the desired approximation. Specifically, let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be an affine transformation that maps R into the unit square $[0, 1]^2$, and apply Lemma 3 for the point set $T(S \cap R)$ and $T(R \cap [0, 1]^2)$. The grid $\mathcal{G}(m)$ clipped in $T(R \cap [0, 1]^2)$ corresponds to a stretched and rotated grid in R ; each grid cell of $\mathcal{G}(m)$ is stretched to a rectangle with the same aspect ratio as R . The convex polygon C_{opt} occupies a constant fraction of the area of R , and so the resulting running time is $O(n_1 + 1/\varepsilon^6)$, where n_1 is the number of points in R . Note that the algorithm of Lemma 3 partitions R into a grid with $O(1/\varepsilon^2)$ cells. The approximation algorithm only cares about which cells are empty and which are not.

Since the algorithm of Lemma 3 is repeated for all possible positions of R , the overall running time is $O(\rho^{-1}(n + \rho^{-1}\varepsilon^{-6}))$, where the first factor of ρ^{-1} counts possible areas, canonical aspect ratios, and orientations, and the second factor of ρ^{-1} inside the parenthesis counts possible positions of the rectangle R . \square

Remark. If $\rho = \Omega(1)$ the running time of this planar algorithm is linear in n .

Since $\rho = \Omega(1/n)$, the running time of the algorithm in Lemma 5 is bounded by $O(\varepsilon^{-6}n^2)$. We summarize our result for the plane in the following.

Theorem 3. *Given a set S of n points in $[0, 1]^2$ and a parameter $\varepsilon > 0$, one can compute an empty convex body $C \subseteq [0, 1]^2$, such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. The running time of the algorithm is $O(\varepsilon^{-6}n^2)$.*

4.4 A faster approximation in higher dimensions

Given a set $S \subseteq [0, 1]^d$ of n points and a parameter $\varepsilon > 0$, we compute an empty convex body $C \subseteq [0, 1]^d$ such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. Similarly to the algorithm in Subsection 4.3, we guess a hyperrectangle R that contains C_{opt} such that $\text{vol}(C_{\text{opt}})$ is at least a constant fraction of $\text{vol}(R)$; and then apply Lemma 4 to R (as the hypercube) to obtain the desired approximation.

Consider a hyperrectangle B_0 of minimum volume (and arbitrary orientation) that contains C_{opt} . The d edges incident to a vertex of a hyperrectangle B are pairwise orthogonal. We call these d directions the *axes* of B ; and the *orientation* of B is the set of its axes.

We next enumerate all possible discretized hyperrectangles of volume $\Omega(1/n)$, guessing the lengths of their axes, their orientations, and their locations as follows:

Guess the length of every axis up to a factor of 2. Since the minimum length of an axis in our case is $\Omega(1/n)$ and the maximum is \sqrt{d} , the number of possible lengths to be considered is $O(\log^d n)$. Let B_1 be a hyperrectangle of minimum volume with canonical side lengths and the same orientation as B_0 such that $B_0 \subseteq B_1$.

We can discretize the orientation of a hyperrectangle as follows. We spread a dense set of points on the sphere of directions, with angular distance $O(1/n)$ between any point on the sphere and its closest point in the chosen set. $O(n^{d-1})$ points suffice for this purpose. We try each point as the direction of the first axis of the hyperrectangle, and then generate the directions of the remaining axes analogously in the orthogonal hyperplane for the chosen direction. Overall, this generates $O(n^{\sum_{i=1}^{d-1} i}) = O(n^{d(d-1)/2})$ possibilities.

Successively replace each axis of B_1 by an approximate axis that makes an angle at most $\alpha = 1/(cn)$ with its corresponding axis, where $c = c(d)$ is a constant depending on d . Let B_2 be a congruent copy of B_1 obtained in this way. If $c = c(d)$ is sufficiently small, then $B_1 \subseteq 2B_2$.

Consider a tiling of \mathbb{R}^d with translates of B_2 . Note that only $O(1/\text{vol}(C_{\text{opt}})) = O(n)$ translates intersect the unit cube $[0, 1]^d$. One of these translates $B_2 + v$ intersects C_{opt} , and then the hyperrectangle $R = 3B_2 + v$ contains C_{opt} . Since $C_{\text{opt}}(S)$ takes a constant fraction of the volume of R , we can deploy Lemma 4 in this case, and get the desired $(1 - \varepsilon)$ -approximation in $\exp(O(\varepsilon^{-d(d-1)/(d+1)} \log \varepsilon^{-1}))$ time. Putting everything together, we obtain the following.

Theorem 4 *Given a set S of n points in $[0, 1]^d$, $d \geq 3$, and a parameter $\varepsilon > 0$, one can compute an empty convex body $C \subseteq [0, 1]^d$, such that $\text{vol}(C) \geq (1 - \varepsilon)\text{vol}(C_{\text{opt}})$. The running time of the algorithm is*

$$\exp\left(O\left(\varepsilon^{-d(d-1)/(d+1)} \log \varepsilon^{-1}\right)\right) n^{1+d(d-1)/2} \log^d n.$$

Remark. Consider a set S of n points in \mathbb{R}^d . The approximation algorithm we have presented can be modified to approximate the largest empty tile, i.e., the largest empty convex body contained in $\text{conv}(S)$, rather than $[0, 1]^d$. The running time is slightly worse, since we need to take the boundary of $\text{conv}(S)$ into account. We omit the details.

5 Conclusions

In this section we briefly outline two exact algorithms for finding the largest area empty convex polygon and the largest area empty triangle amidst n points in the unit square. At the end we list a few open problems.

Largest area convex polygon. Let $S \subset U = [0, 1]^2$, where $|S| = n$. Let T be the set of four vertices of U . Observe that the boundary of an optimal convex body, C_{opt} , contains at least two points from $S \cup T$. By convexity, the midpoint of one of these $O(n^2)$ segments lies in C_{opt} . For each such midpoint m , create a weakly simple polygon P_m by connecting each point $p \in S$ to the boundary of the square along the ray mp . The polygon P_m has $O(n)$ vertices and is empty of points from S in its interior. Then apply the algorithm of Chang and Yap [11] for the potato-peeling problem (mentioned in Section 1) in these $O(n^2)$ weakly simple polygons. The algorithm computes a largest area empty convex polygon contained in a given (non-convex) polygon with n vertices in $O(n^7)$ time. Finally, return the largest convex polygon obtained in this way. The overall running time is $O(n^9)$.

The running time can be reduced to $O(n^8 \log n)$ as follows. Instead of considering the $O(n^2)$ midpoints, compute a set P of $O(n \log n)$ points so that every convex set of area at least $2/(n+4)$ contains at least one of these points. In particular, C_{opt} contains a point from P . The set P can be computed by starting with a $O(n) \times O(n)$ grid, and then computing an ε -net for it, where $\varepsilon = O(1/n)$, using discrepancy [32]. The running time of this deterministic procedure is roughly $O(n^2)$, and the running time of the overall algorithm improves to $O(n^7 \cdot n \log n) = O(n^8 \log n)$.

Largest area empty triangle. The same reduction can be used for finding largest area empty triangle contained in U , resulting in $O(n^2)$ weakly simple polygons P_m . Then the algorithm of Melissaratos and Souvaine [33] for finding a largest area triangle contained in a polygon is applied to each of these $O(n^2)$ polygons. The algorithm finds such a triangle in $O(n^4)$ time, given a polygon with n vertices. Finally, return the largest triangle obtained in this way. The overall running time is $O(n^6)$. Via the ε -net approach (from the previous paragraph) the running time of the algorithm improves to $O(n^4 \cdot n \log n) = O(n^5 \log n)$.

Open questions. Interesting questions remain open regarding the structure of optimal Steiner convex partitions and the computational complexity of computing such partitions.

Other questions relate to the problem of finding the largest empty convex body in the presence of points.

- (1) Is there a polynomial-time algorithm for computing a minimum Steiner convex partition of a given set of n points in \mathbb{R}^d ? Is there one for points in the plane?
- (2) Is there a constant-factor approximation algorithm for the minimum Steiner convex partition of an arbitrary point set in \mathbb{R}^d (without the general position restriction)? Is there one for points in the plane?
- (3) For $d > 2$, the running time of our approximation algorithm for the maximum empty polytope has a factor of the form $n^{O(d^2)}$. It seems natural to conjecture that this term can be reduced to $n^{O(d)}$. Another issue of interest is extending Lemma 2 to higher dimensions for a faster overall algorithm.
- (4) Given n points in $[0, 1]^d$, the problem of finding the largest convex body in $[0, 1]^d$ that contains up to k (outlier) points naturally suggests itself and appears to be also quite challenging.

Acknowledgement. The authors thank Joe Mitchell for helpful discussions regarding the exact algorithms in Section 5, in particular for suggesting the reduction of the maximum-area-empty-convex-body problem to the potato-peeling problem. Many thanks also go to Sergio Cabello and Maria Saumell for pointing us to the recent results of Bautista-Santiago et al. [9] and for suggesting logarithmic factor improvements in the running time of the approximation algorithm in Section 4.3.

References

- [1] O. Aichholzer and H. Krasser, The point set order type data base: A collection of applications and results, in *Proc. 13th Canadian Conf. on Comput. Geom.*, Waterloo, ON, Canada, 2001, pp. 17–20.
- [2] N. Alon, I. Bárány, Z. Füredi, and D. Kleitman, Point selections and weak ε -nets for convex hulls, *Combinatorics, Probability & Computing* **1** (1992), 189–200.
- [3] G. E. Andrews, A lower bound for the volumes of strictly convex bodies with many boundary points, *Transactions of the AMS* **106** (1963), 270–279.
- [4] B. Aronov, M. van Kreveld, M. Löffler, and R. I. Silveira, Peeling meshed potatoes, *Algorithmica* **60**(2) (2011), 349–367.
- [5] J. Backer and M. Keil, The mono- and bichromatic empty rectangle and square problems in all dimensions, in *Proc. 9th Latin American Sympos. on Theoretical Informatics*, vol. 6034 of LNCS, Springer, 2010, pp. 14–25.
- [6] R. P. Bambah and A. C. Woods, On a problem of Danzer, *Pacific J. Math.* **37**(2) (1971), 295–301.

- [7] I. Bárány and J. Pach, On the number of convex lattice polygons, *Combinatorics, Probability & Computing* **1** (1992), 295–302.
- [8] I. Bárány and A. M. Vershik, On the number of convex lattice polytopes, *Geometric and Functional Analysis* **2** (1992), 381–393.
- [9] C. Bautista-Santiago, J. M. Díaz-Báñez, D. Lara, P. Pérez-Lantero, J. Urrutia, and I. Ventura, Computing optimal islands, *Oper. Res. Lett.* **39** (4) (2011), 246–251.
- [10] J. Beck and W. Chen, *Irregularities of Distributions*, vol. 89 of Cambridge Tracts in Mathematics, Cambridge University Press, Cambridge, 1987.
- [11] J.-S. Chang and C.-K. Yap, A polynomial solution for the potato-peeling problem, *Discrete Comput. Geom.* **1** (1986), 155–182.
- [12] B. Chazelle and D. P. Dobkin, Decomposing a polygon into its convex parts, in *Proc. 11th Sympos. Theory of Computing*, ACM Press, 1979, pp. 38–48.
- [13] J. A. De Loera, J. Rambau, and F. Santos, Triangulations: Structures for Algorithms and Applications, vol. 25 of *Algorithms and Computation in Mathematics*, Springer, 2010.
- [14] A. Dumitrescu and M. Jiang, On the largest empty axis-parallel box amidst n points, *Algorithmica* **66**(2) (2013), 225–248.
- [15] A. Dumitrescu and Cs. D. Tóth, Minimum weight convex Steiner partitions, *Algorithmica* **60**(3) (2011), 627–652.
- [16] D. Eppstein, M. Overmars, G. Rote, and G. Woeginger, Finding minimum area k -gons, *Discrete Comput. Geom.* **7**(1) (1992), 45–58.
- [17] T. Fevens, H. Meijer, and D. Rappaport, Minimum convex partition of a constrained point set, *Discrete Appl. Math.* **109**(1-2) (2001), 95–107.
- [18] P. Fischer, Sequential and parallel algorithms for finding a maximum convex polygon, *Comput. Geom. Theory Appl.* **7** (1997), 187–200.
- [19] Z. Füredi and J. Pach, Traces of finite sets: extremal problems and geometric applications, in *Extremal Problems for Finite Sets (P. Frankl, Z. Füredi, G. Katona, D. Miklós, editors)*, vol. 3 of Bolyai Society Mathematical Studies, Budapest, 1994, pp. 251–282.
- [20] J. García-López and C. Nicolás, Planar point sets with large minimum convex partitions, in *Proc. 22nd European Workshop on Comput. Geom.*, Delphi, Greece, 2006, pp. 51–54.
- [21] M. Grantson and C. Levcopoulos, A fixed-parameter algorithm for the minimum number convex partition problem, in *Proc. Japanese Conf. on Discrete Comput. Geom.*, vol. 3742 of LNCS, Springer, 2005, pp. 83–94.

- [22] O. A. Hall-Holt, M. J. Katz, P. Kumar, J. S. B. Mitchell, and A. Sityon, Finding large sticks and potatoes in polygons, in *Proc. 17th ACM-SIAM Sympos. Discrete Algorithms*, 2006, pp. 474–483.
- [23] S. Har-Peled, An output sensitive algorithm for discrete convex hulls, *Comput. Geom. Theory Appl.* **10** (1998), 125–138.
- [24] J. Hershberger and S. Suri, Applications of a semi-dynamic convex hull algorithm, *BIT* **32(2)** (1992), 249–267.
- [25] J. Horton, Sets with no empty convex 7-gons, *Canadian Math. Bulletin* **26** (1983), 482–484.
- [26] J. M. Keil, Polygon partition, in *Handbook of Computational Geometry (J.-R. Sack, J. Urrutia, editors)*, Elsevier, 2000, pp. 491–518.
- [27] J. M. Keil and J. Snoeyink, On the time bound for convex partition of simple polygons, *Internat. J. Comput. Geom. Appl.* **12** (2002), 181–192.
- [28] D. G. Kirkpatrick, A note on Delaunay and optimal triangulations, *Inform. Proc. Lett.* **10(3)** (1980), 127–128.
- [29] C. Knauer and A. Spillner, Approximation algorithms for the minimum convex partition problem, in *Proc. SWAT*, vol. 4059 of LNCS, Springer, 2006, pp. 232–241.
- [30] M. Lassak, Approximation of convex bodies by inscribed simplices of maximum volume, *Contributions to Algebra and Geometry* **52(2)** (2011), 389–394.
- [31] J. Lawrence and W. Morris, Finite sets as complements of finite unions of convex sets, *Discrete Comput. Geom.* **42(2)** (2009), 206–218.
- [32] J. Matoušek, *Lectures on Discrete Geometry*, Springer, 2002.
- [33] E. A. Melissaratos and D. L. Souvaine, Shortest paths help solve geometric optimization problems in planar regions, *SIAM J. Comput.* **21(4)** (1992), 601–638.
- [34] V. Neumann-Lara, E. Rivera-Campo, and J. Urrutia, A note on convex partitions of a set of points in the plane, *Graphs and Combinatorics* **20(2)** (2004), 223–231.
- [35] J. Pach and G. Tardos, Piercing quasi-rectangles—on a problem of Danzer and Rogers, *J. of Combinatorial Theory, Series A* **119** (2012), 1391–1397.
- [36] T. Sakai and J. Urrutia, Convex partitions of point sets in the plane, in *Proc. 7th Japan Conf. on Comput. Geom. and Graphs (Kanazawa, 2009)*, JAIST.
- [37] T. Shermer, Recent results in art galleries, *Proceedings of the IEEE* **80(9)** (1992), 1384–1399.
- [38] A. Spillner, A fixed parameter algorithm for optimal convex partitions, *J. Discrete Algorithms* **6(4)** (2008), 561–569.
- [39] P. Valtr, Sets in \mathbb{R}^d with no large empty convex subsets, *Discrete Mathematics* **108(1-3)** (1992), 115–124.