

Halving by a Thousand Cuts or Punctures

Sariel Har-Peled*

Da Wei Zheng†

August 23, 2022

Abstract

For point sets P_1, \dots, P_k , a set of lines L is *halving* if any face of the arrangement $\mathcal{A}(L)$ contains at most $|P_i|/2$ points of P_i , for all i . We study the problem of computing a halving set of lines of minimal size. Surprisingly, we show a polynomial time algorithm that outputs a halving set of size $O(\sigma^{3/2})$, where σ is the size of the optimal solution. Our solution relies on solving a new variant of the weak ε -net problem for corridors, which we believe to be of independent interest.

We also study other variants of this problem, including an alternative settings, where one needs to introduce a set of guards (i.e., points), such that no convex set avoiding the guards contains more than half the points of each point set.

1. Introduction

A basic problem in algorithms is to partition data effectively in order to apply divide and conquer algorithms, or just store the data or manipulate it efficiently in distributed or parallel fashion. In the context of Computational Geometry, such tasks are usually achieved using cuttings [CF90], partitions [Mat92], or even hashing [AI08]. More recently, there was significant progress [She22] on using polynomials to perform such partitions (e.g., polynomial ham-sandwich theorem) to derive better combinatorial bounds (and in some cases, algorithms).

Partitioning a point set $P \subset \mathbb{R}^2$ via polynomials is quite powerful, as such partitions can have many desirable properties not achievable by the other techniques. However, while computing the partitioning polynomial can be done efficiently [She22], using such partitions algorithmically is challenging. As a concrete example, consider a two dimensional polynomial in the plane $p(x, y)$ used to partition a set of points P . It partitions the plane into cells via its zero set $Z = \{(x, y) \mid p(x, y) = 0\}$ – that is, every connected component C of $\mathbb{R}^2 \setminus Z$ induces a cluster in the partition of P (i.e., $C \cap P$). However, computing these clusters is not algorithmically easy (or convenient) as dealing with roots of high degree polynomials is cumbersome and computationally slow. If one remembers how the polynomial p was computed, in some cases, such tasks become easier – however, other tasks like adjusting the partition when the underlying point set changes remains a challenge, as multi-variable high-degree polynomials are unwieldy.

1.1. Problem I: Separating multiple point sets by lines/planes

For a set P of n points in \mathbb{R}^d , a set L of (hyper)planes *separates* P , if for any pair of points of $p, q \in P$, there is a plane in L that intersects the interior of the segment pq (which also does not contain p or q). In the plane L is a set of lines. The *separability* of P , denoted by $S_n = \text{sep}(P)$, is the size of the smallest set of lines that separates P . The separability of a point set captures how grid-like the point set is. In particular, the separability of the $\sqrt{n} \times \sqrt{n}$ grid is $2\sqrt{n} - 2$, while for n points in convex position the separability is $\lceil n/2 \rceil$ (and this is the worst case assuming general position).

*Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu; <http://sarielhp.org/>. Work on this paper was partially supported by a NSF AF award CCF-1907400.

†Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; dwzheng2@illinois.edu; <https://davidzheng.web.illinois.edu/>.

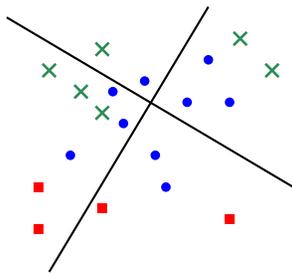


Figure 1.1: Given three point sets, suppose the goal is to break the green (cross) point set into sets with at most three points, the blue (dot) point set into sets with most four points, and the red (square) point set into sets with at most two points. This can be achieved using two separating lines.

This problem can be stated as a hitting set problem (i.e., pick a minimal size set of planes that hits all the segments formed by pairs of points). The standard greedy algorithm yields a $O(\log n)$ approximation, and at least in the plane, it can be sped up by using data-structures [HJ20]. Somewhat surprisingly, the separability of random points (picked inside a unit square) is (roughly) $\Theta(n^{2/3})$, in contrast to grids where it is $\Theta(\sqrt{n})$ [HJ20]. Since the separability of n points requires $\Omega(\sqrt{n})$ lines, an approximation quality of $O(\log n)$ is somewhat more acceptable (although, whether this approximation ratio can be improved in this case remains open).

Separating point sets by lines. One can define murkier partition problems, such as partitioning several point sets in a balanced way simultaneously.

Problem 1.1. An instance \mathcal{I} of the *reduction* problem is specified by k point sets $P_1, \dots, P_k \subset \mathbb{R}^2$, not necessarily disjoint, and corresponding fractions $\xi_1, \dots, \xi_k \in (0, 1]$. The *size* of \mathcal{I} is $m = \sum_i |P_i|$. The goal is to compute the smallest set of lines L , such that for every cell C in the arrangement $\mathcal{A}(L)$ of lines, $|P_i \cap C| \leq \xi_i |P_i|$, for all i . See Figure 1.1 for an example. In the *halving problem*, we have that $\xi_i = 1/2$. Let σ denote the size of the optimal set L .

Observe that σ might be a small constant even if k and n are large.

Current solutions to the reduction problem. This problem can be reduced to several “parallel” instances of partial set cover [HJ18], and this problem can be stated as a submodular optimization problem, which provides an $O(\log n)$ approximation by the greedy algorithm. The basic idea is to define a potential function which captures for every point how far it is from being happily separated enough from the remaining points. Then, the greedy algorithm chooses the line that its addition to the partitioning set reduces this potential function the most, see [HJ18] for details.

Using other techniques, Chekuri *et al.* [CIQ+22] improved the approximation to $O(\log k)$ (assuming that $\xi \geq 1/2$).

The challenge. In light of the above, the interesting case of the reduction problem is when the number of sets k is polynomially large (e.g., $k = \sqrt{m}$), and the optimal solution σ is small (say, a constant). Can one get a constant approximation in this case?

1.2. Problem II: Guarding multiple point sets against convex regions

The second “dual” problem is computing simultaneous weak nets for several point sets.

Problem 1.2. An instance \mathcal{I} of the *guarding* problem is defined by k point sets $P_1, \dots, P_k \in \mathbb{R}^d$ (not necessarily disjoint), and corresponding fractions $\xi_1, \dots, \xi_k \in (0, 1]$. The *size* of \mathcal{I} is $m = \sum_i |P_i|$. The goal is to find a minimum set of points \mathcal{N} , such that every convex region D where $|D \cap P_i| \geq \xi_i |P_i|$ has nonempty intersection with \mathcal{N} . Let σ denote the size of the optimal set \mathcal{N} .

This problem can also be viewed as a hitting set problem, where we wish to hit all convex polygons that contain too many points of a given point set P_i by at least one point.

1.3. Background

Weak ε -nets and guarding points from convex regions. For a set P of n points in \mathbb{R}^d , a set $S \subset \mathbb{R}^d$ is a *weak ε -net* if for every convex region D where $|D \cap P| \geq \varepsilon n$ has nonempty intersection with S . We can view S as a set of guards in the plane that protects the points against any convex set which contains many points. The goal is to pick a discrete point set where $|S|$ is as small as possible. This problem is well studied, see [BFL90, ABFK92, CEG+95, MW04]. The state of the art is the recent results by Rubinfeld [Rub18, Rub21] showing the existence of weak ε -nets of size $O_d(\varepsilon^{-(d-0.5+\alpha)})$ for arbitrarily small $\alpha > 0$. For more detailed history of the problem, see the introduction of Rubinfeld [Rub18, Rub21]. As for a lower bound, Bukh *et al.* [BMN09] gave constructions of point sets for which any weak ε -net must have size $\Omega_d(\varepsilon^{-1} \log^{d-1} \varepsilon^{-1})$. Closing this gap remains a major open problem. See [MV17] for a recent survey of ε -nets and related concepts.

Round-and-cut.¹ Many approximation algorithm works by rounding a fractional optimal solution to an associated LP. Sometime the LP is implicit, and it can be solved using the ellipsoid algorithm via a separation oracle. In particular, it is well known that LP can be solved in (weakly) polynomial time by such an algorithm. At every step, the LP solver asks the separation oracle about the status of a specific solution/point. The oracle either finds a violated constraint and returns it, or alternatively returns that the given query point is feasible. Once a solution to the LP is found, the approximation algorithm rounds the LP to get (hopefully) a good approximation.

In the *round-and-cut* approach [CFLP00], one combines the two steps. Given a query point (i.e., a fractional assignment), the oracle either returns a violated constraint (if one such constraint is easy to find), or tries to round this fractional value. If the rounding is successful – a good approximation was found. Otherwise, the failure of the rounding provides a violated constraint which is returned by the separation oracle. This is especially useful where we do not know how to implement the standard separation oracle, or the standard separation oracle requires exponential time.

A variant of the round-and-cut technique was used (implicitly) in computational geometry. The multiplicative weight update (MWU) algorithms can be viewed as solving an LP. In particular, Clarkson’s algorithm [Cla93] for set cover/hitting set (see also [BG95]), work by assigning weights to points (i.e., think about these as the LP values assigned to the points), and each stage computing an ε -net (the value of ε is guessed in advance). Either a small ε -net is found, or alternatively a multiplicative weight update is applied (i.e., the values of the LP are adjusted). This connection between these MWU algorithms and LP is discussed by Har-Peled [Har11, Chapter 6].

1.4. Our results

We provide polynomial time approximation algorithms for both problems. The reduction problem is solved by studying a fractional version of the line separation problem. We solve the later problem using the round-and-cut framework, where the rounding procedure (essentially) requires a solution to a new problem, which is “dual” to the weak ε -net problem. Specifically, given a set of lines L , one need to find a minimum number of lines that intersect all convex regions intersecting more than ε -fraction of the lines of L . We refer to this problem as the *weak ε -cutting* problem. This problem has similar flavor to the weak ε -net problem.

Surprisingly, unlike for the weak ε -net problem, this problem has a direct $O(1/\varepsilon^2)$ solution. Indeed, setting $r = 1/\varepsilon$, one can compute a $1/r$ -cutting of L . This decomposes \mathbb{R}^d into $O(r^d)$ simplices, so that each one intersects at most n/r lines (i.e., planes or hyperplanes in higher dimensions), where $n = |L|$. In 2d, adding the lines supporting the edges of the triangles readily yields a weak ε -cutting with $O(r^2)$ lines. This solution also works in higher dimensions, yielding a weak ε -cutting of size $O(1/\varepsilon^d)$ in d dimensions.

¹Not to be confused with *cut-and-run* or probably the more correct name for this technique “*cut as long as you can not round*”.

Smaller weak cuttings in 2d. In two dimensions we show how to efficiently construct weak ε -cutting of size $\tilde{O}(1/\varepsilon^{3/2})$. The construction requires over-sampling together with a refinement of larger faces into “large” polygons, and using known combinatorial bounds on the complexity of many faces. See [Theorem 3.3](#) for details.

Weak ε -net for corridors. In the dual, the above problem becomes the following – given a set P of m points in the plane, compute a set \mathcal{N} of points (not in P), such that any corridor containing εm points of P must contain a point of \mathcal{N} , where a *corridor* is the region bound between the upper and lower envelopes of any set of lines. That is, this is the problem of computing *weak ε -net for corridors*. The above constructions readily implies a weak ε -net of size $\tilde{O}(1/\varepsilon^{3/2})$.

These two problems were not studied before, and we consider this result (and its primal) to be quite surprising.

Approximation to the reduction problem. We transform the reduction problem (using lines for separation) to an implicit hitting set problem. We solve the LP relaxation of the later problem, by repeatedly using the weak cutting construction algorithm above to perform rounding, and find a violated constraint if such a constraint exists. This replaces the original constraints involving multiple sets, into “monochromatic” constraints. This yields a set of $O(\phi^{3/2} \log^{3/2} \phi)$ lines that performs the desired separation, where ϕ is the size of the optimal solution, see [Theorem 4.4](#) for details. Thus, when ϕ is a constant, our algorithm is the first constant approximation algorithm for this problem.

Interestingly, to get a fast algorithm, we show that one can reduce the number of lines under consideration. In particular, we show that instead of the $O(m^2)$ lines, one can quickly generate a candidate set of lines of size $\tilde{O}((k/\xi)^2)$, such that it contains a constant approximation to the optimal solution, see [Lemma 4.2](#) for details (here ξ is the minimum fraction of separation required of any set in the original instance). Using this as a preprocessing stage, yields a near linear time approximation algorithm for the reduction problem.

Approximation algorithm for the convex guarding problem. The same approach works for the convex guarding problem, except that the rounding now is done via the “standard” weak ε -net construction. Furthermore, the separation oracle requires finding a bad convex polygon given a suggested net, which is done via dynamic programming, which might be of independent interest (see [Lemma 5.2](#)). In this case, we do not have a way to reduce the candidate set of points being used as part of the net, and thus the running time is worse (i.e., polynomial). We get a $\tilde{O}(\sqrt{\phi})$ approximation in polynomial time, see [Theorem 5.3](#) for details.

Paper organization. We start at [Section 2](#) with some standard background. In [Section 3](#) we present the construction of weak ε -cutting for lines, which in the dual is weak ε -net for corridors. We present the approximation algorithm for the reduction problem in [Section 4](#). In [Section 5](#) we present the approximation algorithm for the convex guarding problem.

2. Preliminaries

2.1. Notations

For an integer n , let $\llbracket n \rrbracket = \{1, \dots, n\}$. In the following, *plane* denotes a flat of dimension $d - 1$ contained in \mathbb{R}^d .

Duality. A plane $g \equiv x_d = b_1 x_1 + \dots + b_{d-1} x_{d-1} + b_d$ in \mathbb{R}^d can be interpreted as a function from \mathbb{R}^{d-1} to \mathbb{R} . Given a point $p = (p_1, \dots, p_d)$, let $g(p) = b_1 p_1 + \dots + b_{d-1} p_{d-1} + b_d$. Thus, a point p lies *above* the plane g if $p_d > g(p)$. As such, a point lies *on* the plane g if $g(p) = p_d$. The *duality* between points and planes is

defined as

$$\begin{aligned} p = (p_1, \dots, p_d) &\implies p^* \equiv x_d = p_1x_1 + \dots + p_{d-1}x_{d-1} - p_d \\ \mathbf{g} \equiv x_d = a_1x_1 + \dots + a_{d-1}x_{d-1} + a_d &\implies \mathbf{g}^* = (a_1, \dots, a_{d-1}, -a_d). \end{aligned}$$

The following is well known [Har11].

Lemma 2.1. *For a point $p = (b_1, \dots, b_d)$, we have the following:*

- (A) $p^{**} = p$.
- (B) A point p lies above/below/on the plane $\mathbf{g} \iff$ the point \mathbf{g}^* lies above/below/on the plane p^* .
- (C) The vertical distance between p and \mathbf{g} is the same as that between p^* and \mathbf{g}^* .
- (D) The distance between two parallel planes \mathbf{g} and \mathbf{h} is the length of the vertical segment $\mathbf{g}^*\mathbf{h}^*$.

2.2. The reduction problem as a hitting set problem

Problem 1.1 specifies the given instance. A convex region D , is **bad** for a set of lines L if D does not intersect any line of L , and there is an index i , such that $|P_i \cap D| > \xi_i |P_i|$. In particular, **Problem 1.1** can be interpreted as computing a minimal set of lines that intersects the interior of all the bad regions. That is, this problem can be stated as a hitting set problem. Since the family of all convex regions induced by a set of all allowable lines L' has an exponential size in $|L'|$, it is not possible to compute this family of regions explicitly in polynomial time or explicitly state the LP associated with this problem (since it has exponential size).

2.3. Corridors

Given a plane (i.e., a line in two dimensions) $\mathbf{g} \equiv x_d = \sum_{i=1}^{d-1} \lambda_i x_i + \lambda_d$, and a non-zero real number $\beta \in \mathbb{R}$, let $\alpha \otimes \mathbf{g}$ be the **scaled** plane $\alpha \otimes \mathbf{g} \equiv x_d = \sum_{i=1}^{d-1} \alpha \lambda_i x_i + \alpha \lambda_d$. Similarly, given a second plane $\mathbf{h} \equiv x_d = \sum_{i=1}^{d-1} \lambda'_i x_i + \lambda'_d$, let their **sum** be the plane $\mathbf{g} \oplus \mathbf{h} \equiv x_d = \sum_{i=1}^{d-1} (\lambda_i + \lambda'_i) x_i + (\lambda_d + \lambda'_d)$. Thus, planes form a vector space with scalar multiplication \otimes , and vector addition \oplus (this is an immediate consequence of ‘‘importing’’ the corresponding operations from the dual space).

Thus, given two planes $\mathbf{g}_1, \mathbf{g}_2$, their **convex combination**, for $t \in [0, 1]$, is the plane

$$\mathbf{g}(t) = ((1-t) \otimes \mathbf{g}_1) \oplus (t \otimes \mathbf{g}_2).$$

For two lines in the plane, the set $\{\mathbf{g}(t) \mid t \in [0, 1]\}$ is the set of all lines in the double wedge between \mathbf{g}_1 and \mathbf{g}_2 passing through the intersection point $\mathbf{g}_1 \cap \mathbf{g}_2$. More generally, given $\alpha_1, \dots, \alpha_m \in [0, 1]$, with $\sum_i \alpha_i = 1$, and planes $\mathbf{g}_1, \dots, \mathbf{g}_m$, they define the **convex combination** $(\alpha_1 \otimes \mathbf{g}_1) \oplus \dots \oplus (\alpha_m \otimes \mathbf{g}_m)$. Given a set of planes L , their **hull**, denoted by $\mathcal{H}(L)$, is the set of all their convex-combinations.

A set of planes L is **convex** if $L = \mathcal{H}(L)$. For a set of planes L , its **corridor** $\text{corr}(L) = \cup \mathcal{H}(L) = \cup_{\mathbf{g} \in \mathcal{H}(L)} \mathbf{g}$ is the union of planes in $\mathcal{H}(L)$ – geometrically, it is the region bounded by the upper and lower envelopes of L . A set of planes L is **bounded**, if all the coefficients used by planes in L are bounded. The closure of a bounded convex set of planes does not contain vertical planes. A region $R \subseteq \mathbb{R}^d$ is a **corridor** if there is a set of planes L , such that $R = \text{corr}(L)$.

Lemma 2.2. *Let L be a finite (bounded) set of planes in \mathbb{R}^d . We have the following:*

- (A) The dual of $\mathcal{C} = \text{corr}(L)$ is a convex polytope $\mathcal{C}^* = (\text{corr}(L))^* = \mathcal{CH}(L^*)$.
- (B) A point $p \in \text{corr}(L) \iff$ the plane p^* intersects $\mathcal{CH}(L^*)$.
- (C) A plane $\mathbf{g} \subseteq \text{corr}(L) \iff$ the point $\mathbf{g}^* \in \mathcal{CH}(L^*)$.
- (D) A plane $\mathbf{g} \subseteq \text{corr}(L) \implies \exists H = \{\mathbf{h}_1, \dots, \mathbf{h}_{d+1}\} \subseteq L$, such that $\mathbf{g} \subseteq \text{corr}(H)$. Furthermore, the set H can be computed in $O(n)$ time.

Proof: (A), (B) and (C) are immediate implications of duality.

(D) Indeed, we have $\mathbf{g} \subseteq \text{corr}(L) \iff \mathbf{g}^* \in \mathcal{CH}(L^*)$. By Carathéodory’s theorem, there is a set of $d+1$ points $H^* = \{\mathbf{h}_1^*, \dots, \mathbf{h}_{d+1}^*\} \subseteq L^*$ such that $\mathbf{g}^* \in \mathcal{CH}(H^*)$. This set can be computed in linear time using low dimensional linear programming [Har11]. As such, for $H = H^{**}$, we have that $\mathbf{g} = \mathbf{g}^{**} \subseteq \text{corr}(H)$. \blacksquare

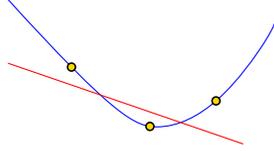


Figure 2.1: Including a point in the corridor by choosing a line separating it from remaining points.

Definition 2.3. For a set P of m points in the plane, and a parameter $\varepsilon \in (0, 1)$, a **weak ε -net for corridors** is a set of points $\mathcal{N} \subseteq \mathbb{R}^2$, such that for any corridor \mathcal{C} in the plane, that contains at least εm points of P , must contain at least one point of \mathcal{N} .

The following implies that “strong” versions of weak ε -nets for corridors, where \mathcal{N} is restricted to be a subset of P , must contain a large number of points.

Lemma 2.4. *Let $\varepsilon \in (0, 1)$, and let m be any positive integer number. There is a set P of m points in the plane, such that any weak ε -net for corridors that is restricted to be a subset of P must be of size $\geq m - \lceil \varepsilon m \rceil + 1$.*

Proof: Let $f(x) = x^2$ (or any other convex smooth function). Let $P = \{(i, f(i)) \mid i \in \llbracket m \rrbracket\}$, and let ℓ be any line lying below P .

Consider any set $X \subseteq P$. For any point $p \in X$, introduce a non-vertical line lying above p that separates p from the remaining points of P . Clearly, for the resulting set of lines L_X , we have $\text{corr}(L_X \cup \{\ell\}) \cap P = X$. Thus, any ε -net $Y \subseteq P$, such that $|Y| < m - \lceil \varepsilon m \rceil + 1$ does not stab the complement set $X = P \setminus Y$, which is both ε -heavy and realizable as a corridor, as $\text{corr}(L_X \cup \{\ell\}) \cap P = X$, a contradiction. ■

2.4. ε -net theorem

Definition 2.5. A set $\mathcal{N} \subseteq X$ is an **ε -net** for X if for any range $\mathbf{r} \in \mathcal{R}$, if $|\mathbf{r} \cap X| \geq \varepsilon |X|$, then \mathbf{r} contains at least one point of \mathcal{N} (i.e., $\mathbf{r} \cap \mathcal{N} \neq \emptyset$).

Theorem 2.6 (ε -net theorem, [HW87]). *Let (X, \mathcal{R}) be a range space of VC dimension δ , and suppose that $0 < \varepsilon \leq 1$ and $\varphi < 1$. Let \mathcal{N} be a set obtained by $m = \Omega(\varepsilon^{-1}(\log \varphi^{-1} + \delta \log \varepsilon^{-1}))$ random independent draws from X . Then \mathcal{N} is an ε -net for X with probability at least $1 - \varphi$.*

3. Weak ε -net for corridors

The input is a set L of n lines in the plane, and a parameter $\varepsilon \in (0, 1)$. Our purpose here is to compute a set \mathcal{N} of lines in the plane, such that any convex body B in the plane such that $|B \cap L| \geq \varepsilon n$, we have that a line of \mathcal{N} intersects B , where

$$B \cap L = \{\ell \in L \mid \ell \cap B \neq \emptyset\}.$$

In the dual, this corresponds to the property that any corridor \mathcal{C} containing more than εn points of L^* , contains a point of \mathcal{N}^* . The set \mathcal{N} is a **weak ε -cutting**, as every face of the arrangement $\mathcal{A}(\mathcal{N})$ intersects at most εn lines of L . This definition extends naturally to higher dimensions.

We start with an easy construction.

Lemma 3.1. *For a set L of n planes in \mathbb{R}^d , one can compute a weak ε -cuttings size $O(1/\varepsilon^d)$.*

Proof: Compute regular cuttings of size $O(r^d)$ of L , for $r = \lceil 1/\varepsilon \rceil$. Furthermore, such cuttings decompose \mathbb{R}^d into $O(r^d)$ simplices, where each simplex intersects at most n/r planes of L . We replace each $(d-1)$ -dimensional face of a simplex in the cutting by the plane that supports it. Clearly, the resulting set \mathcal{N} of planes is of size $O(r^d)$, and fulfils the requirement of being a weak ε -cutting.

Indeed, consider any convex region $B \subseteq \mathbb{R}^d$. If B is not fully contained in a simplex of the cuttings, then it must intersect one of the planes of \mathcal{N} . Otherwise, it is contained in a simple simplex of the cutting, say ∇ . But then, we have $|B \cap L| \leq |\nabla \cap L| \leq \varepsilon n$. ■

3.1. A better construction in two dimensions

A better construction, in two dimensions, arises by oversampling coupled with breaking down the large faces into polygons with fewer edges.

Complexity of the i th largest face. Let L be a set of lines in the plane, and consider the arrangement $\mathcal{A} = \mathcal{A}(L)$. The *complexity* of a face F of \mathcal{A} , denoted by $|F|$, is the number of edges/rays on the boundary of F .

Lemma 3.2. *For a set L of ν lines in the plane, let c_i be the complexity of the i th face of $\mathcal{A} = \mathcal{A}(L)$ in decreasing order of the complexity of the faces. Then $c_i = O(\nu^{2/3}/i^{1/3} + \nu/i + 1)$.*

Proof: Let c_i be the complexity of the largest i th face in \mathcal{A} . The complexity of i faces in the arrangement of \mathcal{A} is $M_i = O(\nu^{2/3}i^{2/3} + \nu + i)$ [SA95]. Namely, we have $ic_i \leq \sum_{j=1}^i c_j \leq M_i$, which implies $c_i \leq M_i/i$. ■

Theorem 3.3. *Let L be a set of n lines in \mathbb{R}^2 , and let $\varepsilon \in (0, 1)$ be a parameter. One can compute a set R of lines of size $O(\varepsilon^{-3/2} \log^{3/2} \varepsilon^{-1})$, such that R is a weak ε -cuttings of L . That is, any open convex region D that avoids the lines of R , intersects at most εn lines of L .*

In the dual, R^ is a weak ε -net for corridors for the point set L^* . That is, any corridor \mathcal{C} that avoids the points of R^* contains at most εn points of L^* .*

Proof: Let $r = \lceil 10/\varepsilon \rceil$. Let R_1 be a random sample from L of size $\nu = car \log r$, where c is a sufficiently large constant, and $\alpha \in \llbracket r^3 \rrbracket$ is a parameter. The sample R_1 is an δ -net for L for vertical trapezoids, where $\delta = 1/(2c\alpha r)$, with probability close to one. In particular, consider a face F of $\mathcal{A}(R_1)$. If F has at most α edges, then it can be decomposed into α vertical trapezoids, each one intersecting at most δn lines of L . As such, $|F \cap L| \leq \alpha \delta n \leq n/r$,

Thus, we need to fix only large faces with strictly more than α edges. Let F be such a face, and sweep it from left to right by a vertical line, whenever the sweep line encounters the α th vertex of F , introduce a vertical line to break it into smaller faces. Let R_2 be the resulting set of new lines introduced. If the total number of edges of faces with more α edges is I , then the overall number of lines introduced is I/α . Let $R = R_1 \cup R_2$.

Consider a face F of $\mathcal{A}(R)$ that is contained in a face F' of $\mathcal{A}(R_1)$. There are several possibilities:

- (i) F has at most α edges. Then, F can be decomposed into α vertical trapezoids that avoids R_1 , and $|F \cap L| \leq n/r$, following the argument above.
- (ii) The face F' has at most α edges – the same argument implies $|F \cap L| \leq |F' \cap L| \leq n/r$.
- (iii) Otherwise, the above process introduced vertical lines into R_2 that break F' into polygons with at most α edges. In particular, one of these polygons, say, $D \subseteq F'$ contains F . Arguing as above, we have $|F \cap L| \leq |D \cap L| \leq n/r$.

Now, consider any convex region B that avoids the lines of R , and observe that it is contained in a single face of $\mathcal{A}(R)$, which intersects at most εn lines of L , by the above. Thus R is the desired weak ε -cutting.

Recall that $\nu = |R_1| = car \log r$. By Lemma 3.2, if we require that $\nu^{2/3}/m^{1/3} + \nu/m + 1 \leq c'\alpha$, for c' a sufficiently small constant. Thus, the complexity of the m th face in $\mathcal{A}(R_1)$ is at most α . This holds if $\nu^{2/3}/m^{1/3} \leq c'\alpha/3$ and $\nu/m \leq c'\alpha/3$. This in turns holds if

$$m \geq \left(\frac{\nu^{2/3}}{c'\alpha/3} \right)^3 = \Omega\left(\frac{\nu^2}{\alpha^3} \right) = \Omega\left(\frac{r^2 \log^2 r}{\alpha} \right).$$

The total complexity of these m large faces is $I = O(m^{2/3}\nu^{2/3} + m + \nu) = O\left(\left(\frac{r^2 \log^2 r}{\alpha} \right)^{2/3} (\alpha r \log r)^{2/3} \right) = O(r^2 \log^2 r)$. As such, the set R has size $\nu + I/\alpha = O\left(\alpha r \log r + \frac{r^2 \log^2 r}{\alpha} \right)$, which is minimized for $\alpha = O(\sqrt{r \log r})$. ■

4. The reduction problem: Approximation algorithm

The input instance \mathcal{I} is made of k point sets P_1, \dots, P_k in \mathbb{R}^2 , not necessarily disjoint, and fractions $\xi_1, \dots, \xi_k \in (0, 1]$. Let $\xi = \min_i \xi_i$. Furthermore, let $m_i = |P_i|$, for $i \in [k]$, and $m = \sum_i |P_i|$. As a reminder, the goal is to compute the smallest set of lines L , such that for every cell C in the arrangement $\mathcal{A}(L)$ of lines, $|P_i \cap C| \leq \xi_i |P_i|$, for all i .

4.1. Reducing the number of candidate cutting lines

Observation 4.1. *Consider the range space where the ground set is \mathbb{R}^2 , and ranges are corridors formed by 3 lines. The VC dimension of this range space is $O(1)$.*

Given a set Q of points in general position. Let $\text{lines}(Q)$ be the set of all lines passing through pairs of points of Q . The set $\text{lines}(Q)$ has size $O(|Q|^2)$, and it can be computed in this time.

Lemma 4.2. *Given an instance \mathcal{I} to the reduction problem, in the plane, with k different sets and $\xi = \min_i \xi_i$. One can compute a set of lines L , of size $O((k/\xi)^2 \log^2(k/\xi))$, such that there is a solution for \mathcal{I} of size $\leq 3\sigma$, made out of lines from L , where σ is the size of the optimal solution (where any line in the plane can be used). The running time of the algorithm is bounded by the output size.*

Proof: Let \mathcal{N}_i be a $\xi_i/2$ -net of P_i for corridors formed by three lines. By the above, a sample of size $O(\xi_i^{-1} \log(k/\xi_i))$ is such a net with probability $\geq 1 - 1/k^{O(1)}$. Let $\mathcal{N} = \cup_i \mathcal{N}_i$, and let $L = \text{lines}(\mathcal{N})$.

We claim that L is the desired set of lines. Assume there is an optimal solution $\mathcal{O} = \{\psi_1, \dots, \psi_\sigma\}$. Each line ψ_i , separates \mathcal{N} into two sets Q_i^+, Q_i^- . Consider the polygons $\mathcal{CH}(Q_i^-)$ and $\mathcal{CH}(Q_i^+)$. Let $L_i \subseteq L$ be the lines the edges of these convex-hulls induces, as well as the two lines realizing the cross tangents.

The corridor $\text{corr}(L_i)$ does not contain any point of \mathcal{N} in its interior, and $\psi_i \subseteq \text{corr}(L_i)$. By dual Carathéodory theorem, [Lemma 2.2](#) (D), there is a set $L'_i \subseteq L_i$ of three lines, such that $\psi_i \subseteq \text{corr}(L'_i)$. We claim that $H = \cup_i L'_i$ is a valid solution to the reduction set, and $|H| \leq 3\sigma$.

Consider any face F of $\mathcal{A}(H)$. If it is contained in a face F' of $\mathcal{A}(\mathcal{O})$, then, for any i , we have $|F \cap P_i| \leq |F' \cap P_i| \leq \xi_i |P_i|$. Otherwise, F must be crossed by an optimal line, say, ψ_i . But then, $F \subseteq \text{corr}(L'_i)$. This implies that, for any j , we have

$$|F \cap P_j| \leq |\text{corr}(L'_i) \cap P_j| \leq (\xi_j/2) |P_j|,$$

since $\text{corr}(L'_i)$ does not contain any point of \mathcal{N} (and thus of \mathcal{N}_j) in its interior, and \mathcal{N}_j is a $\xi_j/2$ -net for corridors induced by three lines for P_j . In the above, we treated both F and $\text{corr}(L'_i)$ as open sets. One need to repeat the above argument also for edges of the arrangement $\mathcal{A}(H)$, but this case is easier, as can be easily verified. We have $|\mathcal{N}| = \sum_i O(\xi_i^{-1} \log(k/\xi_i)) = O((k/\xi) \log(k/\xi))$, and thus $|L| = O(|\mathcal{N}|^2) = O((k/\xi)^2 \log^2(k/\xi))$. ■

4.2. Implicit LP, separation oracle and fractional solution

Given an instance \mathcal{I} of the reduction problem, a set L of n lines, consider the set of all “bad” regions. To this end, let \mathcal{D} be the set of all convex sets in the plane. We assume that no two vertices of $\mathcal{A}(L)$ have the same x value.

A convex region $\sigma \in \mathcal{D}$ is **bad** if there is an index j , such that $|P_j \cap \sigma| > \xi_j |P_j|$, and let \mathcal{D}_{bad} be the set of all bad regions in \mathcal{D} . The associated LP for computing a hitting set, of at most t lines, for all bad polygons is:

$$v(L) = \sum_{\ell \in L} x_\ell \leq t \tag{4.1}$$

$$1 \geq x_\ell \geq 0 \quad \forall \ell \in L$$

$$v(\sigma) = \sum_{\ell \in L \cap \sigma} x_\ell \geq 1 \quad \forall \sigma \in \mathcal{D}_{\text{bad}}. \tag{*}$$

A *separation oracle* is a procedure that gets assignment of fractional values for the variables of the LP, and returns a violated constraint if such a constraint exists. For our purposes it is enough to find an approximate violation. The first two conditions in the above LP can be checked directly, and if they are violated, they are returned as violated to the LP solver. Otherwise, the algorithm tries to round the solution as described next.

4.2.1. The rounding attempt

A fractional solution to the above LP can be efficiently rounded. One need to adapt the algorithm of [Section 3](#) to work for the case that the lines have weights. Conceptually, we treat $\alpha = v(L)$ as the number of lines we have, and $\varepsilon = 1/(2\alpha)$ as the desired threshold. The ε -net theorem applies verbatim in the weighted settings (the sampling has to be adapted to the weights, but this is standard), and the algorithm of [Section 3](#) applies verbatim. We get a weak ε -cutting of size $O(\alpha^{3/2} \log^{3/2} \alpha)$, realized by a set \mathcal{W} of lines. For every face F of $\mathcal{A}(\mathcal{W})$ the algorithm computes

$$v(F) = \sum_{\ell \in L \cap F} v(\ell), \quad \text{and} \quad |F \cap P_j|, \quad j = 1, \dots, k.$$

If $v(F) \leq 1/2$ then the weak cutting computed failed (which happens with low probability), and the algorithm recomputes the weak cutting. If there is any j such that $|F \cap P_j| > \xi_j |P_j|$, then the rounding failed. Namely, we found a bad region $\sigma \in \mathcal{D}_{\text{bad}}$ (i.e., a constraint of the LP that is violated). The algorithm returns the corresponding constraint of (*) as being violated

$$v(F) = \sum_{\ell \in L \cap F} x_\ell \geq 1,$$

as $v(F) < 1/2$. If all the faces of $\mathcal{A}(\mathcal{W})$ are good, then \mathcal{W} is a valid solution, with $O(\alpha^{3/2} \log^{3/2} \alpha)$ lines.

Remark 4.3. The ellipsoid algorithm (with a separation oracle) solves an LP with n variables, using a number of iterations that is polynomial in n and \log of the largest number in the LP, which is n (in our case). See [\[GLS93\]](#).

Theorem 4.4. *Given an instance $\mathcal{I} = (P_1, \xi_1, \dots, P_k, \xi_k)$ of the *reduction problem* in the plane of size m , with k sets and $\xi = \min_i \xi_i$, one can compute a set L of $O(\alpha^{3/2} \log^{3/2} \alpha)$ lines, such that for any cell C of $\mathcal{A}(L)$, and any $j \in \llbracket k \rrbracket$, we have that $|P_j \cap C| \leq \xi_j |P_j|$, where α is the minimum size of any set of lines with this property.*

The expected running time of this algorithm is $O(m \log(k/\xi) + (k/\xi)^{O(1)})$, and the algorithm succeeds with probability $\geq 1 - (\xi/k)^{O(1)}$.

Proof: A naive upper bound on α is $O(k/\xi)$, as each set P_i can be partitioned using $\lceil 1/\xi_i \rceil$ lines.

We first generate a small set of lines L as candidates for cutting the set, using [Lemma 4.2](#). This stage succeeds with probability $\geq 1 - (\xi/k)^{O(1)}$. The resulting set of lines L has size $n = O((k/\xi)^2 \log^2(k/\xi))$.

We now run the above algorithm with exponential search on the parameter t used in the LP from 1 up to $O(k/\xi)$, see [Eq. \(4.1\)](#), stopping as soon as the algorithm succeeds. The number of separation oracle calls performed in each attempt to solve the LP is $n^{O(1)}$. Each such attempt involves computing the weak cuttings, which can be done in $O(t^3)$ time. Verifying that no face contains too many points of any set P_j , can be done by preprocessing the arrangement for point-location queries. This takes $O(m \log t)$ time for point locations per rounding attempt. Overall, this results in running time $(m + k/\xi)^{O(1)}$.

To get an improved running time, we have to avoid performing the point-location stage at each rounding attempt. To this end, observe that the weak cuttings only adds vertical lines that passes through vertices of $\mathcal{A}(L)$. Let L' be this set of vertical lines, and compute the arrangement $\mathcal{A}(L \cup L')$. This arrangement has $(k/\xi)^{O(1)}$ faces, and it is enough to compute for each face in this arrangement how many points of P_j falls into this face, as all polygons/faces considered by the above algorithm are disjoint union of such basic faces. Reducing P_1, \dots, P_k in this way takes $O((k/\xi)^{O(1)} + m \log(k/\xi))$ time. From this point on, we have that the “reduced” point sets have total size $(k/\xi)^{O(1)}$, and thus each rounding attempt can be performed in $(k/\xi)^{O(1)}$ time. ■

5. Approximation algorithm for the guarding problem

The input instance \mathcal{I} is made of k point sets P_1, \dots, P_k in \mathbb{R}^2 not necessarily distinct, and fractions $\xi_1, \dots, \xi_k \in (0, 1]$. Let $\xi = \min_i \xi_i$ and $m_i = |P_i|$, for $i \in \llbracket k \rrbracket$, and let $m = \sum_i m_i$. As a reminder, the goal is to compute the smallest set of points \mathcal{N} that is simultaneously a weak ξ_i -net for every point set P_i , that is for any convex set D in \mathbb{R}^2 with $|P_i \cap D| \geq \xi_i m_i$, we have that $\mathcal{N} \cap D \neq \emptyset$. Let $P = \cup_{i=1}^k P_i$.

5.1. Reducing the number of candidate points

Given an instance \mathcal{I} of the guarding problem, we again consider the set of all convex sets \mathcal{D} , and focus on the set of bad polygons \mathcal{D}_{bad} . As a reminder, $\sigma \in \mathcal{D}$ is *bad*, if there is an index $j \in \llbracket k \rrbracket$, such that $|P_j \cap \sigma| \geq \xi_j |P_j|$. We make a few observations about bad polygons:

- (I) Since we care only with how a convex set interact with the set P , it suffices to restrict \mathcal{D}_{bad} to the set of bad convex polygons that are the convex hull of some subset of points of P . We denote this set by $\mathcal{B} = \mathcal{D}_{\text{bad}}(P) = \{\mathcal{CH}(\sigma \cap P) \mid \sigma \in \mathcal{D}_{\text{bad}}\}$.
- (II) Any solution \mathcal{N} stabs all the polygons in \mathcal{B} .
- (III) Thus, consider the arrangement $\mathcal{A}(\mathcal{B})$. Any point $p \in \mathcal{N}$, can be moved to a vertex of the face of this arrangement that contains it, and it would still stab the same polygons of \mathcal{B} . Since the vertices are defined by edges of convex hulls of points in P , we can restrict \mathcal{N} to be a subset of the vertices of this arrangement.
- (IV) Specifically, it suffices to restrict to point guards that lie on the intersection of two line segments joining pairs of points of P . There are $O(|P|^2)$ such line segments and thus $O(|P|^4)$ points defined this way. Let Q be the set of all such points.

5.2. The implicit LP

The associated LP for hitting set of the polygons \mathcal{B} with at most t points is:

$$\begin{aligned} v(Q) &= \sum_{p \in Q} x_p \leq t \\ 1 &\geq x_p \geq 0 && \forall p \in Q \\ v(\sigma) &= \sum_{p \in Q \cap \sigma} x_p \geq 1 && \forall \sigma \in \mathcal{B}. \end{aligned} \tag{**}$$

5.3. The rounding scheme

We will use following result of Rubin [Rub18] about weak ε -nets.

Theorem 5.1. *Let P be a set of m points in \mathbb{R}^2 , and let $\varepsilon \in (0, 1)$ be a parameter. For any $\alpha > 0$, one can compute a weak ε -net S of size $O(\varepsilon^{-3/2-\alpha})$ in time $\tilde{O}(m^2/\sqrt{\varepsilon})$, where \tilde{O} hides polylogarithmic factors in m .*

Let $\alpha = v(Q)$. The above construction applies for discrete point sets, but we can apply the construction to the points of Q by including point $p \in Q$ with multiplicity $v'(p) = \lfloor 4|Q|v(p) \rfloor$ and compute a weak ε -net $\mathcal{W} \subseteq Q$, with $\varepsilon = 1/(4\alpha)$, of size $O(\alpha^{3/2+\alpha})$ for some constant $\alpha > 0$. This set can be constructed in $\tilde{O}(|Q|\alpha) = \tilde{O}(m^8\alpha^{3/2})$ time. Observe that for any multiset of the points S , $4|Q|v(S) \geq v'(S) \geq 4|Q|v(S) - |S|$. Consider a convex set σ with $v(\sigma) = v(\sigma \cap Q) \geq 1/2 \geq 2\varepsilon\alpha$, and observe that

$$v'(\sigma \cap Q) \geq 4|Q|v(\sigma \cap Q) - |\sigma \cap Q| \geq 8\varepsilon\alpha|Q| - |\sigma \cap Q| \geq |Q| = \varepsilon \cdot 4|Q|v(Q) \geq \varepsilon v'(Q).$$

This implies that any convex set $\sigma \in \mathcal{B}$ that contains no points of \mathcal{W} must have $v(\sigma) < 1/2$. The idea is now to test if there exists a convex polygon $\sigma \in \mathcal{B}$ that contains no point of \mathcal{W} . If no such polygon exists, then we have found the desired set of guards (i.e., we successfully rounded the given LP solution). Otherwise, we found a polygon $\sigma \cap \mathcal{W} = \emptyset$ (i.e., $v(\sigma) < 1/2$) – namely, we found a constraint belonging to (**) that is being violated.

5.3.1. Searching for a bad polygon

This step was easier for [Problem 1.1](#) as we only needed to check every face of the arrangement of the computed lines. However, here there are exponentially many (canonical) convex sets that avoid the set \mathcal{W} of guards that need to be checked.

Lemma 5.2. *Given a set P of m points, and another set \mathcal{W} of at most m points, one can decide, in $O(m^4 \log m)$ time, if there exists a closed convex set σ that satisfies $|P \cap \sigma| \geq \xi m$ and $|\mathcal{W} \cap \sigma| = 0$.*

Proof: For the simplicity of exposition, we assume the x -values of all the points under consideration are all distinct – this can be ensured by slightly perturbing the points. Arguing as above, it suffices to consider only polygons σ that are formed by the convex hull of some subset of points in P . Consider two segments with endpoints in P and their vertical decomposition – there might be at most one vertical trapezoid has these two segments as a floor and ceiling segments. Let \mathcal{U} be the set of all such trapezoids. Clearly, the set \mathcal{U} can be computed in $O(m^4)$ time. Furthermore, using simplex range searching, one can count for each such trapezoid τ how many points of P it contains (ignoring say points that lie on its right wall), denoted by $w(\tau)$, and how many points of \mathcal{W} it contains. With $O(m^{2+o(1)})$ preprocessing, such queries can be answered in $O(\log m)$ time [CSW92]. Let \mathcal{T} be the set \mathcal{U} after we remove from it all the trapezoids that contains any point of \mathcal{W} . The set \mathcal{T} can be computed in $O(m^4 \log m)$ time, and has size $O(m^4)$.

Two trapezoids $\tau_1, \tau_2 \in \mathcal{T}$ that share a vertical wall are *compatible*, if there is a polygon $\sigma \in \mathcal{D}$ such that $\sigma \cap \text{slab}(\tau_1) = \tau_1$ and $\sigma \cap \text{slab}(\tau_2) = \tau_2$, where $\text{slab}(\tau_i)$ is the minimal vertical strip containing τ_i . Note that this is a local condition and can be checked in constant time.

We create a DAG G over \mathcal{T} , where an edge $\tau_1 \rightarrow \tau_2$ is in G if τ_1 and τ_2 are compatible, and τ_1 is to the left of τ_2 . By the assumption of unique x -coordinates, τ_1 and τ_2 must share either the bottom or top supporting lines. The DAG G has $O(m^4)$ vertices, and potentially there are $O(m)$ out going edges from each vertex – as by assumption two adjacent compatible trapezoids changes only either the floor or ceiling supporting segment. This results in a graph G with $O(m^5)$ edges. However, by adding special entrance vertices to each trapezoid, and chaining them by slope of the changing segment, one can reduce the number of edges to $O(m^4)$. The graph G can be computed in $O(m^4 \log m)$ time.

Note, that any convex polygon in \mathcal{B} corresponds to a maximal path in the DAG G . A trapezoid $\tau \in \mathcal{T}$ is a *start* (resp. *final*) trapezoid if its left (resp., right) wall is a vertex. Note, that the leftmost (resp., rightmost) trapezoid in any polygon of \mathcal{D} must be a start (resp., final) trapezoid.

The problem thus reduces to computing the longest path in the DAG G with vertex weights $w(\tau)$ for each $\tau \in \mathcal{T}$. This can be done with a standard dynamic program in time linear in the size of the DAG, which takes $O(m^4)$ time. ■

5.4. The result

The above algorithm provides us with a procedure for computing a bad polygon if it exists for the currently suggested solution – namely, we can use it in the round-and-cut framework.

Theorem 5.3. *Given an instance $\mathcal{I} = (P_1, \xi_1, \dots, P_k, \xi_k)$ of the *guarding problem* in the plane of size $m = \sum_i |P_i|$, with k sets and $\xi = \min_i \xi_i$, one can compute a set \mathcal{W} of $O(\sigma^{3/2+\alpha})$ points, for any fixed $\alpha > 0$, such that \mathcal{W} is a weak ξ_i -net for P_i for all $i \in \llbracket k \rrbracket$. That is, for any convex polygon B , and any $j \in \llbracket k \rrbracket$, if $|B \cap P_j| > \xi_j |P_j|$ then $B \cap \mathcal{W} \neq \emptyset$. The algorithm has running time polynomial in m , assuming $\xi_1, \dots, \xi_k \geq 1/m$.*

Proof: Observe that m is a naive upper bound on the number of guards needed, as we can simply guard all the points. We can get better upper bounds by taking weak ε -nets of each class of points, but this is not needed.

We restrict our attention to our candidate set Q with $n = O(m^4)$ points by (IV). We now run the round-and-cut algorithm with exponential search on t from 1 up to m , stopping as soon as the algorithm succeeds. The number of separation oracle calls performed in each attempt to solve the LP is $n^{O(1)} = m^{O(1)}$. Each rounding step takes $\tilde{O}(m^8 \sigma^{3/2})$ where σ is at most m . The total running time is $m^{O(1)}$. ■

6. Conclusions

We revisited the natural geometric divide-and-conquer **reduction problem**. In the process we introduced a new kind of weak ε -nets for corridors (i.e., weak ε -cutting). We presented a non-trivial construction that provides such nets of size $\tilde{O}(1/\varepsilon^{3/2})$. Using this construction of nets as a rounding scheme, used within the round-and-cut framework, we were able to get a $\tilde{O}(\sqrt{\phi})$ -approximation to the optimal solution, where ϕ was the size of the optimal solution. While this approximation quality is somewhat “underwhelming” it is still a significant improvement when ϕ is small (say a constant), where previously only logarithmic approximation was known, and previous approaches seems unlikely to lead to a sublogarithmic approximation in the general case.

We then solved the dual problem of guarding a point set against convex regions by inserting guards, except that in this case in addition to the rounding provided by the “standard” weak ε -net, we had to use dynamic programming to find bad convex regions if they exists.

There are numerous open problems for further research raised by our work. The first one is improving the approximation quality even further. Secondly, further improving the size of the construction of weak ε -nets for corridors (i.e., weak ε -cutting for lines). The bound we get is mysteriously very similar to the best known bound for the weak ε -net for points (for convex regions). It is natural to further investigate this connection. Ultimately, improving and simplifying Rubin’s construction in 2d seems like a worthy problem for further research.

Beyond that, this work emphasize the “rounding is approximation” approach. It is natural to wonder if there are other natural geometric problems where better rounding is possible because of the geometry, which would lead to better approximation algorithms.

References

- [ABFK92] N. Alon, I. Bárány, Z. Füredi, and D. J. Kleitman. Point selections and weak ε -nets for convex hulls. *Comb. Probab. Comput.*, 1: 189–200, 1992.
- [AI08] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM*, 51(1): 117–122, 2008.
- [BFL90] I. Bárány, Z. Füredi, and L. Lovász. On the number of halving planes. *Combinatorica*, 10: 175–183, 1990.
- [BG95] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite vc-dimension. *Discrete Comput. Geom.*, 14(4): 463–479, 1995.
- [BMN09] B. Bukh, J. Matoušek, and G. Nivasch. Lower bounds for weak epsilon-nets and stair-convexity. *Proc. 25th Annu. Sympos. Comput. Geom. (SoCG)*, 1–10, 2009.
- [CEG+95] B. Chazelle, H. Edelsbrunner, M. Grigni, L. J. Guibas, M. Sharir, and E. Welzl. Improved bounds on weak epsilon-nets for convex sets. *Discret. Comput. Geom.*, 13: 1–15, 1995.
- [CF90] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3): 229–249, 1990.
- [CFLP00] R. D. Carr, L. Fleischer, V. J. Leung, and C. A. Phillips. Strengthening integrality gaps for capacitated network design and covering problems. *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA*, 106–115, 2000.
- [CIQ+22] C. Chekuri, T. Inamdar, K. Quanrud, K. Varadarajan, and Z. Zhang. Algorithms for covering multiple submodular constraints and applications. *Journal of Combinatorial Optimization*, 2022.
- [Cla93] K. L. Clarkson. Algorithms for polytope covering and approximation. *Proc. 3th Workshop Algorithms Data Struct. (WADS)*, vol. 709. 246–252, 1993.
- [CSW92] B. Chazelle, M. Sharir, and E. Welzl. Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Algorithmica*, 8(5&6): 407–429, 1992.

- [GLS93] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. 2nd. Vol. 2. Algorithms and Combinatorics. Berlin Heidelberg: Springer-Verlag, 1993.
- [Har11] S. Har-Peled. *Geometric Approximation Algorithms*. Vol. 173. Math. Surveys & Monographs. Boston, MA, USA: Amer. Math. Soc., 2011.
- [HJ18] S. Har-Peled and M. Jones. Few cuts meet many point sets. *CoRR*, abs/1808.03260, 2018. arXiv: 1808.03260.
- [HJ20] S. Har-Peled and M. Jones. On separating points by lines. *Discret. Comput. Geom.*, 63(3): 705–730, 2020.
- [HW87] D. Haussler and E. Welzl. ε -nets and simplex range queries. *Discrete Comput. Geom.*, 2: 127–151, 1987.
- [Mat92] J. Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8: 315–334, 1992.
- [MV17] N. H. Mustafa and K. Varadarajan. *Epsilon-approximations and epsilon-nets*. 2017. arXiv: 1702.03676.
- [MW04] J. Matoušek and U. Wagner. New constructions of weak ε -nets. *Discrete Comput. Geom.*, 32(2): 195–206, 2004.
- [Rub18] N. Rubin. An improved bound for weak epsilon-nets in the plane. *Proc. 59th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, 224–235, 2018.
- [Rub21] N. Rubin. Stronger bounds for weak epsilon-nets in higher dimensions. *Proc. 53rd Annu. ACM Sympos. Theory Comput. (STOC)*, to appear.
- [SA95] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. New York: Cambridge University Press, 1995.
- [She22] A. Sheffer. *Polynomial Methods and Incidence Theory*. Cambridge University Press, Mar. 2022.