Generalized Ham-Sandwich Cuts

William Steiger · Jihui Zhao

Received: 17 March 2009 / Revised: 4 August 2009 / Accepted: 14 August 2009 /

Published online: 11 September 2009

© Springer Science+Business Media, LLC 2009

Abstract Bárány, Hubard, and Jerónimo recently showed that for given well-separated convex bodies $S_1, \dots, S_d$ in R^d and constants $\beta_i \in [0, 1]$, there exists a unique hyperplane h with the property that $\text{Vol}(h^+ \cap S_i) = \beta_i \cdot \text{Vol}(S_i)$; h^+ is the closed positive transversal halfspace of h , and h is a “generalized ham-sandwich cut.” We give a discrete analogue for a set S of n points in R^d which are partitioned into a family $S = P_1 \cup \dots \cup P_d$ of *well-separated* sets and are in *weak general position*. The combinatorial proof inspires an $O(n(\log n)^{d-3})$ algorithm which, given positive integers $a_i \leq |P_i|$, finds the unique hyperplane h incident with a point in each P_i and having $|h^+ \cap P_i| = a_i$. Finally we show two other consequences of the direct combinatorial proof: the first is a stronger result, namely that in the discrete case, the conditions assuring existence and uniqueness of generalized cuts are also necessary; the second is an alternative and simpler proof of the theorem in Bárány et al., and in addition, we strengthen the result via a partial converse.

Keywords Partitions of measures · Ham-sandwich cuts · Algorithms

1 Introduction

Given d sets $S_1, S_2, \dots, S_d \in R^d$, a ham-sandwich cut is a hyperplane h that simultaneously bisects each S_i . “Bisect” means that $\mu(S_i \cap h^+) = \mu(S_i \cap h^-) < \infty$, h^+ , h^- the closed halfspaces defined by h , and μ a suitable, “nice” measure on Borel sets

W. Steiger (✉) · J. Zhao

Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, New Jersey 08854-8004, USA

e-mail: steiger@cs.rutgers.edu

J. Zhao

e-mail: zhaojih@cs.rutgers.edu

in R^d , e.g., the volume. The well-known ham-sandwich theorem guarantees the existence of such a cut. As with other consequences of the Borsuk-Ulam theorem [10], there is a discrete version that applies to sets $P_1, \dots, P_d$ of points in general position in R^d . For example, Lo et al. [9] gave a direct proof of a discrete version of the ham-sandwich theorem which inspired an efficient algorithm to compute a cut. More recently Bereg [4] studied a discrete version of a result of Bárány and Matoušek [2] that showed the existence of wedges that simultaneously equipartition three measures on R^2 (they are called equitable two-fans). By seeking a direct, combinatorial proof of a discrete version (for counting measure on points sets in R^2) he was able to strengthen the original result and also obtained a beautiful, nearly optimal algorithm to construct an equitable two-fan. Finally, Roy, and Steiger [13] followed a similar path to obtain complexity results for several other combinatorial consequences of the Borsuk-Ulam theorem.

The present paper is in the same spirit. The starting point is a recent, interesting result about generalized ham-sandwich cuts.

Definition 1 (See [8]) Given $k \leq d + 1$, a family $S_1, \dots, S_k$ of connected sets in R^d is well separated if, for every choice of $x_i \in S_i$, the affine hull of $x_1, \dots, x_k$ is a $(k - 1)$ -dimensional flat in R^d .

Bárány et al. [1] proved the following:

Proposition 1 Let $K_1, \dots, K_d$ be well-separated convex bodies in $\mathbb{R}^d$, and $\beta_1, \dots, \beta_d$ given constants with $0 \leq \beta_i \leq 1$. Then there is a unique hyperplane $h \subset R^d$ with the property that $\text{Vol}(K_i \cap h^+) = \beta_i \cdot \text{Vol}(K_i)$, $i = 1, \dots, d$.

Here h^+ denotes the closed, positive transversal halfspace defined by h : that is, the halfspace where, if Q is an interior point of h^+ and $z_i \in K_i \cap h$, the d -simplex $\Delta(z_1, \dots, z_d, Q)$ is negatively oriented [1]. Specifying this choice of halfspaces is what forces h to be uniquely determined. Bárány et al. give analogous results for such *generalized ham-sandwich cuts* for other kinds of well-separated sets that support suitable measures.

We are interested in a version of Proposition 1 for n points partitioned into d sets in R^d ; i.e., points in $S = P_1 \cup \dots \cup P_d$, $P_i \cap P_j = \emptyset$, $i \neq j$, $|S| = n$. For this context, we use the following:

Definition 2 Point sets $P_1, \dots, P_d$ in R^d are well separated if their convex hulls, $\text{Conv}(P_1), \dots, \text{Conv}(P_d)$, are well separated.

We need some kind of general position and will assume the following weaker form.

Definition 3 Points in $S = P_1 \cup \dots \cup P_d$ have weak general position if, for each $(x_1, \dots, x_d)$, $x_i \in P_i$, the affine hull of $x_1, \dots, x_d$ is a $(d - 1)$ -dimensional flat that contains no other point of S .

This does not prohibit more than d data points from being in a hyperplane, e.g., if they are all in the same P_i . For the discrete analogue of a generalized cut, we use the following:

Definition 4 Given positive integers $a_i \leq |P_i|$, an $(a_1, \dots, a_d)$ -cut is a hyperplane h for which $h \cap P_i \neq \emptyset$ and $|h^+ \cap P_i| = a_i$, $1 \leq i \leq d$.

As in Proposition 3, a cut is a transversal hyperplane (here incident with at least one data point in each P_i), and h^+ its positive closed halfspace. The discrete version of Proposition 3 is the following:

Theorem 1 If $P_1, \dots, P_d$ are well-separated point sets in R^d , and $a_1, \dots, a_d$ are positive integers, $a_i \leq |P_i|$, then

- (i) if an $(a_1, \dots, a_d)$ -cut exists, it is unique. Also,
- (ii) if the points have weak general position, then a cut exists for every $(a_1, \dots, a_d)$, $1 \leq a_i \leq |P_i|$.

It might be possible to prove this using the results of [1] along with a standard argument that takes the average of n symmetric d -dimensional normal distributions, one centered at each data point. The variance of the distributions is decreased to zero, and one argues about the limit of the cuts (see [7]). Instead we give a direct combinatorial proof in Sect. 2. We did this because of the interest in the algorithmic problem where, given n points distributed among d well-separated sets in R^d , and in weak general position, the object is to find the cut for given $a_1, \dots, a_d$. Our combinatorial proof of Theorem 1 leads directly to the formulation of an efficient, $O(n(\log n)^{d-3})$ algorithm to compute generalized cuts, $d \geq 3$. This algorithm is one of our main results, and is described in Sect. 3.¹

There are two other useful consequences of the combinatorial proof. The ideas in our proof of Theorem 1 can be applied in the continuous case, and we get an alternative, simpler proof for the original theorem. These results appear in Sect. 4. Also, as a corollary to Theorem 1, we observe that in the discrete case, the conditions for the existence and uniqueness of all cuts are also necessary. This enabled us to strengthen the original theorems by showing that something similar holds in the continuous context.

2 Proof of the Discrete Version

There are several equivalent forms of the well-separation property for connected sets [3], in particular the fact that such a family is well separated if and only if the convex hulls are well separated. Others include:

1. Sets $S_1, \dots, S_k$, $k \leq d + 1$, are well separated if and only if, when I and J are disjoint subsets of $1, \dots, d + 1$, there is a hyperplane separating the sets S_i , $i \in I$, from the sets S_j , $j \in J$.

¹ An earlier version of some of these results appeared in [15].

2. $S_1, \dots, S_d$ are well separated in R^d if and only if there is no $(d-2)$ -dimensional flat that has nonempty intersection with all $\text{Conv}(S_i)$, $i = 1, \dots, d$.

In view of Definition 2, they hold for the discrete context as well.

Given points $p_i \in \text{Conv}(P_i)$, $i = 1, \dots, d$ (not necessarily data points in S), the hyperplane h determined by $\{p_1, \dots, p_d\}$ is a transversal hyperplane of dimension $d-1$. As in Bárány et al. [1], if a unit vector c satisfies $\langle c, p_i \rangle = t$ for some fixed constant t and for all i , the unit normal vector v of h can be chosen as either c or $-c$. The *positive transversal hyperplane* arises when v is chosen so that

$$\det \begin{vmatrix} p_1 & p_2 & \dots & p_d & v \\ 1 & 1 & \dots & 1 & 0 \end{vmatrix} > 0.$$

We can write h as $\{p \in R^d : \langle p, v \rangle = t\}$, and h^+ , the *positive transversal halfspace*, as

$$h^+ = \{p \in R^d : \langle p, v \rangle \leq t\}.$$

The relation $p \in h^+$ is invariant under translation and rotation.

Proof of Theorem 1 The proof is by induction. The base case $d = 2$ is probably folklore (but see [12]). “Well-separated” implies that points in P_1 may be dualized to (red) lines having positive slopes and those in P_2 , to (blue) lines having negative slope. If a red/blue intersection q has a_1 red lines and a_2 blue lines above it, vertex q is the dual of an (a_1, a_2) -cut. It must be the unique one because the red levels have positive slope and blue ones have negative slope, proving (i).

If P_1 and P_2 also have weak general position, every red/blue intersection in the dual is a distinct vertex, $|P_1| \cdot |P_2|$ of them in all, and each is incident with just those two lines. This implies that each level in the first arrangement has a unique intersection with every level of the second, proving (ii). In fact the unique intersection can be found in linear time by adapting the prune-and-search algorithm given in [12] for intersection of median levels.

Next, suppose that the claim holds in every dimension $j < d$; we show that it also holds in R^d . Let π be a hyperplane that separates P_1 from $\bigcup_{i=2}^d P_i$. Fix a point $x \in \text{Conv}(P_1)$ and project each data point $z \in \bigcup_{i=2}^d P_i$ onto π via the mapping $P_x : z \rightarrow \overline{xz} \cap \pi$. Write P'_i for the multiset of images in π of the points $z \in P_i$.

Fact 1 $P'_2, \dots, P'_d$ are $d-1$ well-separated sets in π .

If not, there is a $(d-3)$ -flat $\rho \subset \pi$ that meets all $\text{Conv}(P'_i)$, $i \geq 2$. But the span of x and ρ is a $(d-2)$ -flat that meets all $\text{Conv}(P_1), \dots, \text{Conv}(P_d)$, a contradiction.

Fact 2 If $P_1, \dots, P_d$ have weak general position and if we project from a point $x \in P_1$, then $P'_2, \dots, P'_d$ have weak general position in π .

Each $z' \in P'_i$ is the image of a distinct point $z \in P_i$, $i > 1$. A transversal flat $\rho_x \subset \pi$ has dimension $d-2$ by Fact 1. If it contains one point x'_i from each P'_i , $i > 1$, and

any other $z' \in \bigcup_{i=2}^d P'_i$, then x and ρ_x span a hyperplane that violates weak general position for $P_1, \dots, P_d$.

These facts show that the induction hypotheses apply to the images $P'_2, \dots, P'_d$ in π .

Given a point $x \in \text{Conv}(P_1)$ and $(a_2, \dots, a_d)$, a hyperplane h_x containing x is an $(a_2, \dots, a_d)$ semi-cut (or just a semi-cut) if, for each $i > 1$, it is incident with a point $p_i \in P_i$ and $|h_x^+ \cap P_i| = a_i$. The following useful fact is straightforward:

Lemma 1 *Given $x \in \text{Conv}(P_1)$ and $(a_2, \dots, a_d)$, if there is an $(a_2, \dots, a_d)$ semi-cut h_x , then it is unique.*

Proof Suppose that h_1 and h_2 are distinct $(a_2, \dots, a_d)$ semi-cuts incident with $x \in \text{Conv}(P_1)$. Then there are points $q_i = h_1 \cap P_i$ and $q'_i = h_2 \cap P_i$, $i = 2, \dots, n$, and the images of these points in π would be distinct $(a_2, \dots, a_d)$ cuts, in violation of the induction hypothesis. $\square$

To advance the induction, fix $(a_1, \dots, a_d)$ and suppose that h_x is a cut with these values, $x \in P_1$. By Lemma 1, it is the unique semi-cut containing x , so suppose that there is an $(a_2, \dots, a_d)$ semi-cut h_y through $y \in P_1$, $y \notin h_x$. Hyperplanes h_x and h_y cannot meet in $\text{Conv}(P_1)$ since any such point would be in two different $(a_2, \dots, a_d)$ semi-cuts, violating Lemma 1. But this implies that $a_1 \neq |P_1 \cap h_y^+|$: if $y \in h_x^+$, so is every $z \in P_1 \cap h_x^+$; if $y \in h_x^-$, so is every $z \in P_1 \cap h_x^-$. Therefore h_x is unique, which proves statement (i) in Theorem 2.

Now suppose that $P_1, \dots, P_d$ have weak general position and fix $x \in P_1$ and $(a_2, \dots, a_d)$. Projecting from x , there is a unique $(a_2, \dots, a_d)$ -cut $\rho_x \subset \pi$ by the induction hypothesis and the fact that each $z' \in \bigcup_{i=2}^d P'_i$ is the image of a distinct $z \in \bigcup_{i=2}^d P_i$. The hyperplane h_x determined by x and ρ_x is an $(m_x, a_2, \dots, a_d)$ -cut, m_x denoting $|P_1 \cap h_x^+|$. Lemma 1 implies that there is no other $(m_x, a_2, \dots, a_d)$ -cut. Also, repeating this procedure for every $x \in P_1$, existence and uniqueness imply that the integers m_x , $x \in P_1$ form a permutation of $1, \dots, |P_1|$. So for some $z \in P_1$, we have the unique $(a_1, \dots, a_d)$ -cut, and this proves statement (ii). $\square$

In fact the conditions of the Theorem are also necessary.

Corollary 1 *Well separation and weak general position are necessary if all $(a_1, \dots, a_d)$ -cuts exist and are unique.*

Weak general position is necessary for the existence and uniqueness of all $(a_1, \dots, a_d)$ -cuts by simple counting. There are $|P_1| \cdot |P_2| \cdots |P_d|$ different d -tuples $(a_1, \dots, a_d)$, and there are this many different transversal hyperplanes through data points only if we have weak general position.

Now suppose that $P_1, \dots, P_d$ are not well separated. By Property 1 at the beginning of this section, there is a partition $I \cup J$ of $\{1, \dots, d\}$ such that $A = \text{Conv}(\bigcup_{i \in I} P_i) \cap \text{Conv}(\bigcup_{j \in J} P_j) \neq \emptyset$. For points in A on the boundaries of the convex hulls, weak general position is violated. For points of A interior to both convex hulls, any halfspace containing $\bigcup_{i \in I} P_i$ also contains at least one point of $\bigcup_{j \in J} P_j$

in its interior. If we set $a_i = 1$ for $i \in I$ and $a_i = |P_i|$ for $i \in J$, no $(a_1, \dots, a_d)$ -cut can exist.

3 An Algorithm for Generalized Cuts

From now on we assume weak general position and well separation. Theorem 2 implies that for every $1 \leq a_i \leq |P_i|$, $i = 1, \dots, d$, there is a unique set of data points $p_1, \dots, p_d$, $p_i \in P_i$, for which the affine hull of $(p_1, \dots, p_d)$ is an $(a_1, \dots, a_d)$ -cut. So we could use a brute force enumeration and find it in $O(n^{d+1})$, $O(n)$ being the cost to test each d -tuple.

A small improvement can be obtained by resorting to the following algorithmic result of [9] (slightly restated to reflect new upper bounds on k -sets [6, 11, 14]).

Proposition 2 *Given n points in R^d which are partitioned into d sets $P_1, \dots, P_d$ in R^d , a ham-sandwich cut can be computed in time proportional to the (worst-case) time needed to construct a given level in the arrangement of n given hyperplanes in R^{d-1} . The latter problem can be solved within the following bounds:*

$$\begin{aligned} O(n^{4/3} \log^2 n / \log^* n) & \text{ for } d = 3, \\ O(n^{5/2} \log^{1+\delta} n) & \text{ for } d = 4, \\ O(n^{4-\frac{2}{45}} \log^{1+\delta} n) & \text{ for } d = 5, \\ O(n^{d-1-a(d)}) & \text{ for } d \geq 6. \end{aligned}$$

$\delta > 0$ is an appropriate constant, and $a(d) > 0$ a small constant; also $a(d) \rightarrow 0$ as $d \rightarrow \infty$.

It is not difficult to verify that the ham-sandwich algorithms given in [9] may be extended to find generalized cuts for well-separated points sets having weak general position—given that they exist—and in this way, the complexity of finding generalized cuts may be reduced to $O(n^{d-1-a(d)})$.

Here, we will describe a much more practical algorithm, applying ideas from the proof in Sect. 2. We showed there that for each data point $x \in P_1$ and $(a_2, \dots, a_d)$, there is a unique $(m_x, a_2, \dots, a_d)$ -cut h_x that contains x . Furthermore, for each j , $1 \leq j \leq |P_1|$, there is a unique $x \in P_1$ for which $m_x = j$, where $m_x = |P_1 \cap h_x^+|$. Thus we could consider in turn all $x \in P_1$. For each, we project onto π , find the unique $(a_2, \dots, a_d)$ -cut $\rho_x \subset \pi$, and compute $m_x = |P_1 \cap h_x^+|$ for h_x , the hyperplane spanned by x and ρ_x . At some stage we will discover the unique $z \in P_1$ for which $m_z = a_1$ and h_z is the $(a_1, \dots, a_d)$ -cut. The cost would be bounded by the cost to solve n problems in R^{d-1} .

In fact we will find the desired $z \in P_1$ by solving at most $O(\log n)$ problems in R^{d-1} . The key is the ability to prune a fixed fraction of remaining points in P_1 after a search step with $x \in P_1$ by using the fact that if $m_x < a_1$, no point $y \in h_x^+ \cap P_1$ has $m_y = a_1$.

ALGORITHM GEN-CUT

1. choose $c > 0$, a small, fixed integer (say 10)
2. Find a hyperplane π that separates P_1 from $P_2 \cup \dots \cup P_d$
3. $C \leftarrow P_1$
4. $a \leftarrow a_1$
5. WHILE $|C| > c$ DO
 - (a) Construct A , an ϵ -approximation to C with respect to halfspaces
 - (b) FOR each $x \in A$ DO
 - i. Project each $y \in P_2 \cup \dots \cup P_d$ onto π ; let P'_i denote the projections of the points in P_i
 - ii. Find the $(a_2, \dots, a_d)$ -cut $\rho_x \subset \pi$ for the projections $P'_2, \dots, P'_d$ by solving a $(d-1)$ -dimensional problem
 - iii. Get h_x , the hyperplane that spans x and ρ_x .
 - iv. Compute the number of points of C in the positive transversal halfspace h_x^+
 - v. END FOR
 - (c) Prune from C points $x \in P_1$ whose n_x is too small or too large, and adjust C and a
 - (d) END WHILE
6. For each remaining data point in $x \in C$, project, find the $(a_2, \dots, a_d)$ -cut ρ_x in π for the projections by solving a $(d-1)$ -dimensional problem, get h_x , and compute $n_x = |P_1 \cap h_x^+|$, stopping when $n_x = a_1$.

In Step 2, finding a separating hyperplane π can be formulated as a linear programming problem and can be solved in time $O(n)$, for fixed dimension d . In Step 3, C is the set of candidates for the sought point $z \in P_1$; initially $C = P_1$. The number of undeleted points in the positive transversal halfspace of z 's semicut is denoted by a ; initially $a = a_1$.

In the WHILE loop we construct an ϵ -approximation to C . The range space $(C, \mathcal{A})$ has VC dimension $d+1$, where $\mathcal{A}$ denotes the set of all halfspaces in $\mathbb{R}^d$ that contain some points in C . By [5], in $O(|C|)$ time [i.e., linear; in fact, it is $O((d+1)^{3(d+1)}(\frac{d+1}{\epsilon^2} \log \frac{d+1}{\epsilon})^{d+1}|C|)$] we can construct an ϵ -approximation $A \subset C$, having constant size [in fact, $|A| = k = O(\frac{d+1}{\epsilon^2} \log \frac{d+1}{\epsilon})$].

The FOR loop in 5b is traversed $k \equiv |A|$ times. The cost of each traversal is dominated by $O(B_{d-1})$, the cost of the $(d-1)$ -dimensional problem in (ii); the cost of (i) is $O(n)$, and that of (iv) is $O(|C|)$.

At the end of the FOR we have for each $x \in A$, the value of $n_x = |h_x^+ \cap C|$. These distinct values order the elements $x \in A$, and our target value, a , is (1) less than the smallest n_x , (2) greater than the largest n_x , or (3) between a successive pair in the ordering. In the first case we delete all $y \in C$, $y \notin h_u^+$, where $n_u = \min(n_x, x \in A)$. In the second case we delete all $y \in C$, $y \in h_v^+$, where $n_v = \max(n_x, x \in A)$; here we also reduce a by $a \leftarrow a - n_v$. The middle case is similar. Since A is an ϵ -approximation, only a constant fraction ($< 1/(k+1) + 2\epsilon$) of the points in C remains after pruning.

The geometric decrease in $|C|$ implies that the number of iterations of the WHILE loop is bounded by $O(\log |P_1|) = O(\log n)$. Therefore Step 5b contributes $O(B_{d-1} \log n)$ to the total cost of the loop, where B_k denotes the complexity of the present algorithm in dimension k . This dominates the total cost of the loop because all other steps have cost either $O(n)$ or $O(|C|)$ and contribute a total of $O(n \log n)$ to the loop.

When the loop terminates, each remaining point in C is treated in time $O(B_{d-1})$ by executing Steps (i)–(iv) in 5b. Then, instead of Step 5c, we test whether $|h^+ \cap P_1| = a_1$; exactly one point will have this property. Since the base case for dimension $d = 2$ has linear running time, the present algorithm will find a generalized cut in $O(n(\log n)^{d-2})$.

Finally, for $d = 3$, Lo et al. [9] showed how to find a ham-sandwich cut for well-separated point sets in linear time. That algorithm is easily adapted to generalized cuts. Using this as the base case for $d > 2$, the algorithm just described will now have running time $O(n(\log n)^{d-3})$ for dimensions $d \geq 3$, and we have shown the following:

Theorem 2 *Given n points partitioned into well-separated sets $P_1, \dots, P_d$ and having weak general position, and $(a_1, \dots, a_d) \in [0, 1]^d$, an $(a_1, \dots, a_d)$ -cut can be found in time $O(n(\log n)^{d-3})$, $d \geq 3$, and in linear time if $d = 2$.*

4 A Simple Proof for the Continuous Case

In this section, we apply the inductive approach of the proof of Theorem 1 to give a new proof for the continuous case. We need to extend the approach to *nice measures* and, to be self-contained, we repeat notation and terminology from Bárány et al. [1].

Writing $v \in S^{d-1}$ for the unit outer normal vector of a halfspace H , we denote the halfspace $\{x \in R^d : \langle x, v \rangle \leq t\}$ by $H(v \leq t)$. Analogously we write $H(v = t) = \{x \in R^d : \langle x, v \rangle = t\}$. Given a set $K \subset R^d$, a unit vector v , and a scalar t , we denote the set $H(v = t) \cap K$ by $K(v = t)$; analogously $K(v \leq t) = H(v \leq t) \cap K$.

Let μ be a finite measure on the Borel subsets of R^d , and let $v \in S^{d-1}$ be a unit vector. Define

$$t_0 = t_0(v) = \inf\{t \in R : \mu(H(v \leq t)) > 0\},$$

$$t_1 = t_1(v) = \sup\{t \in R : \mu(H(v \leq t)) < \mu(R^d)\}.$$

We write $H(s_0 \leq v \leq s_1)$ for the closed slab between the hyperplanes $H(v = s_0)$ and $H(v = s_1)$ and define the set K by

$$K = \bigcap_{v \in S^{d-1}} H(t_0(v) \leq v \leq t_1(v)).$$

K is called the support of μ . It is convex, and $\mu(R^d \setminus K) = 0$.

Barany et al. [1] used the following:

Definition 5 A measure μ on R^d is nice if:

- (1) $t_0(v)$ and $t_1(v)$ are finite for every $v \in S^{d-1}$.
- (2) $\mu(H(v=t)) = 0$ for all $v \in S^{d-1}$ and $t \in R$.
- (3) $\mu(H(s_0 \leq v \leq s_1)) > 0$ for every $v \in S^{d-1}$ and for all s_0, s_1 satisfying $t_0(v) \leq s_0 < s_1 \leq t_1(v)$.

We observe that

Fact 3 Condition (3) in the definition is equivalent to

- (3') For any two hyperplanes h_1, h_2 with $h_1 \cap K \neq \emptyset$ and $h_2 \cap K \neq \emptyset$ but $h_1 \cap h_2 \cap K = \emptyset$, the measure of the closed slab of K between h_1 and h_2 is positive.

Proof Condition (3) is a special case of (3'). On the other hand, given $h_1 \cap h_2 \cap K = \emptyset$, let $x \in h_1 \cap K$ be a point with the smallest distance from h_2 . Then the hyperplane h' incident with x and parallel to h_2 , together with h_2 , forms a slab with positive measure, by (3), and this is a subset of the slab defined by h_1 and h_2 since K is convex. $\square$

One of the main results (Theorem 3) in [1] is the following:

Proposition 3 Suppose that μ_i is a nice measure on R^d with support K_i , $i \in \{1, \dots, d\}$. Assume that the family $\mathcal{F} = \{K_1, \dots, K_d\}$ is well separated and let $\alpha = (\alpha_1, \dots, \alpha_d) \in [0, 1]^d$. Then there is a unique positive transversal halfspace H such that $\mu_i(K_i \cap H) = \alpha_i \cdot \mu_i(K_i)$, $i = 1, \dots, d$.

Here is a simple proof along the lines we used for Theorem 1.

Proof We will use induction on dimension d and normalize each measure so that $\mu_i(R^d) = 1$. As before, a hyperplane h is an $(\alpha_1, \dots, \alpha_d)$ -cut if its corresponding halfspace H has measure $\mu_i(K_i \cap H) = \alpha_i \cdot \mu_i(K_i)$, $i = 1, \dots, d$.

For the base case, take $d = 1$ under the nice measure μ_1 . From (1), the support K_1 is a finite line segment $[l, u]$. From Definition 5 and Fact 3, the function $f : x \mapsto \mu_1(v \leq x)$ is easily seen to satisfy (i) $f(x) = 0$, $x \leq l$, and $f(x) = 1$, $x \geq u$, and (ii) $f(x)$ is strictly increasing and continuous on $[l, u]$, properties that guarantee the existence and uniqueness of an α_1 -cut for every $\alpha_1 \in [0, 1]$.

Now suppose that the claim holds for every dimension $j < d$. Let π be a hyperplane that separates K_1 from $\bigcup_{i=2}^d K_i$. For a point $x \in K_1$ and $y \in \bigcup_{i=2}^d K_i$, define $P_x(y) = \overline{xy} \cap \pi$, a mapping that projects each K_i onto π , $i > 1$. We write

$$PK_i := P_x(K_i)$$

for the image of K_i in π and define the measure μ'_i on π by

$$\mu'_i(S) := \mu_i\{v \mid v \in R^d, P_x(v) \in S\}$$

for all measurable $S \subseteq \pi$.

From the definition of nice measure it easily follows that for each $i = 2, \dots, d$, μ'_i is also a nice measure on π . In addition,

Fact $PK_2, \dots, PK_d$ are $d - 1$ well-separated sets in π .

Therefore the induction hypotheses apply to $PK_2, \dots, PK_d$ and an $(\alpha_2, \dots, \alpha_d)$ -cut, $\rho_x \subset \pi$ exists for $PK_2, \dots, PK_d$ under measures $\mu'_2, \dots, \mu'_d$, and it is unique. Here we still call the hyperplane h_x determined by x and ρ_x an $(\alpha_2, \dots, \alpha_d)$ -semicut (or just a semi-cut). By the definition of μ'_i we have $\mu_i(K_i \cap H_x) = \alpha_i$, $i = 2, \dots, d$, where H_x is the positive halfspace of h_x . As in the discrete case, we have the following:

Lemma 2 Fix $x \in K_1$ and $(\alpha_2, \dots, \alpha_d)$. An $(\alpha_2, \dots, \alpha_d)$ semi-cut h_x exists and is unique.

This in turn implies that, for any $x \neq y \in K_1$, the semicuts h_x, h_y either are the same or do not meet in K_1 . Finally, fix $(\alpha_2, \dots, \alpha_d)$ and define the function $f : x \in K_1 \mapsto \mu_1(K_1 \cap H_x)$, where H_x the positive halfspace of semicut h_x . Because μ_1 is a nice measure, and in view of (3'), $f(y) < f(x)$ if $y \in H_x$, and f is continuous. Therefore the existence and uniqueness of an $(\alpha_1, \dots, \alpha_d)$ -cut follows, and the induction advances.

Now consider a partition $I \cup J$ of $\{1, \dots, d\}$. A cut h with $\alpha_i = 0$ for $i \in I$ and $\alpha_j = 1$ for $j \in J$ exists only if $\bigcup_{i \in I} K_i$ and $\bigcup_{j \in J} K_j$ are in different closed half-spaces of h . Letting I and J range over all 2^d partitions of $\{1, \dots, d\}$, we either have well-separation or, for some I, J , and corresponding h , there are points $P_i \in K_i \cap h$, and these d points do not span a $(d - 1)$ -flat. We call this *degenerate nonseparation*, and we have the following partial converse to Proposition 2. $\square$

Corollary 2 If all possible cuts exist, either the K_i are well separated or they are degenerately nonseparated.

Remark We tried to find a way to do the inductive step in constant time, similar to the way Lo et al. [9] did for separated ham-sandwich cuts in R^3 , but we did not succeed. A main open question is whether there is an $O(n)$ algorithm for this problem.

References

1. Bárány, I., Hubard, A., Jerónimo, J.: Slicing convex sets and measures by a hyperplane. *Discrete Comput. Geom.* **39**(1–3), 67–75 (2008)
2. Bárány, I., Matoušek, J.: Simultaneous partitions of measures by k -fans. *Discrete Comput. Geom.* **25**(3), 317–334 (2001)
3. Bárány, I., Valtr, P.: A positive fraction Erdős–Szekeres theorem. *Discrete Comput. Geom.* **19**(3), 335–342 (1998)
4. Bereg, S.: Equipartitions of measures by 2-fans. *Discrete Comput. Geom.* **34**(1), 87–96 (2005)
5. Chazelle, B.: *The Discrepancy Method: Randomness and Complexity*. Cambridge University Press, Cambridge (2000)
6. Dey, T.: Improved bounds for planar k -sets and related problems. *Discrete Comput. Geom.* **19**(3), 373–382 (1998)
7. Edelsbrunner, H.: *Algorithms in Combinatorial Geometry*. Springer, New York (1987)
8. Kramer, H., Nemeth, A.: Supporting spheres for families of independent convex sets. *Arch. Math.* **24**, 91–96 (1973)

9. Lo, C.-Y., Matoušek, J., Steiger, W.: Algorithms for ham-sandwich cuts. *Discrete Comput. Geom.* **11**, 433–452 (1994)
10. Matoušek, J.: *Using the Borsuk–Ulam Theorem*. Springer, Berlin (2003)
11. Matoušek, J., Sharir, M., Smorodinsky, S., Wagner, U.: On k -sets in four dimensions. *Discrete Comput. Geom.* **35**(2), 177–191 (2006)
12. Megiddo, N.: Partitioning with two lines in the plane. *J. Algorithms* **6**, 430–433 (1985)
13. Roy, S., Steiger, W.: Some combinatorial and algorithmic aspects of the Borsuk–Ulam theorem. *Graphs Comb.* **23**, 331–341 (2007)
14. Sharir, M., Smorodinsky, S., Tardos, G.: An improved bound for k -sets in three dimensions. *Discrete Comput. Geom.* **26**(2), 195–204 (2001)
15. Steiger, W., Zhao, J.: Generalized ham-sandwich cuts for well-separated point sets. In: *Proceedings of the 20th Canadian Conference on Computational Geometry*, Montreal, August 2008