Generalized Ham-Sandwich Cuts for Well Separated Point Sets

William Steiger* Jihui Zhao‡

Abstract

Bárány, Hubard, and Jerónimo recently showed that for given well separated convex bodies $S_1, \dots, S_d$ in $\mathbb{R}^d$ and constants $\beta_i \in [0,1]$, there exists a unique hyperplane h with the property that $\mathrm{Vol}(h^+ \cap S_i) = \beta_i \cdot \mathrm{Vol}(S_i)$; 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 $\mathbb{R}^d$ which is 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 $a_i \le |P_i|$, finds the unique hyperplane h incident with a point in each P_i and having $|h^+ \cap P_i| = a_i$. Finally we show that the conditions assuring existence and uniqueness of generalized cuts are also necessary.

1 Introduction.

Given d sets $S_1, S_2, ..., S_d \in \mathbb{R}^d$, 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 in $\mathbb{R}^d$, 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 $\mathbb{R}^d$. 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 $\mathbb{R}^2$ (they are called equitable two-fans). By seeking a direct, combinatorial proof of a discrete version (for counting measure on points sets in $\mathbb{R}^2$) 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 [12] followed a similar path to obtain complexity results for several other combinatorial consequences of the Borsuk-Ulam theorem.

*Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, New Jersey 08854-8004; ({steiger,zhaojih}@cs.rutgers.edu.)

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 $k \le d+1$, a family $S_1, \dots, S_k$ of connected sets in $\mathbb{R}^d$ 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 $\mathbb{R}^d$.*

Bárány et.al. [1] proved

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

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 $\Delta(z_1, ..., z_d, Q)$ 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 $\mathbb{R}^d$; i.e., points in $S = P_1 \cup \dots \cup P_d$, $P_i \cap P_j = \phi, i \ne j$, $|S| = n$. For this context we use

Definition 2: *Point sets $P_1, ..., P_d$ are well separated if their convex hulls,* $\mathrm{Conv}(P_1), \dots, \mathrm{Conv}(P_n)$*, are well separated.*

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

Definition 3: *Points in $S = P_1 \cup \ldots \cup P_d$ have weak general position if, for each $(x_1, \dots, x_d)$, $x_i \in P_i$, aff$(x_1, \dots, x_d)$ is a $(d-1)$-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

Definition 4: *Given positive integers $a_i \le |P_i|$, an $(a_1, ..., a_d)$-cut is a hyperplane h for which $h \cap P_i \ne \phi$ and $|h^+ \cap P_i| = a_i$, $1 \le i \le d$.*

As in Proposition 1, a cut is a transversal hyperplane (here incident with at least one data point in each P_i)

and h^+ its positive halfspace. The discrete version of Proposition 1 is

Theorem 2 *If $P_1, \ldots, P_d$ are well separated point sets in $\mathbb{R}^d$, then (i) if an $(a_1, ..., a_d)$-cut exists, it is unique. Also (ii) if the points have weak general position, cuts exist for every $(a_1, ..., a_d)$, $1 \leq a_i \leq |P_i|$.*

This might be proved using some results of [1] via a standard argument that takes the average of n probability measures, one centered at each data point. The variance of the measures is decreased to zero, and one argues about the limit (see [7]). Instead we give a direct combinatorial proof in Section 2. In addition, and of independent interest, we show that both well-separated and weak general position are also necessary for every possible $(a_1, ..., a_d)$-cut to exist and be unique. An analogous converse is likely to hold for Proposition 1.

There is also interest in the algorithmic problem where, given n points distributed among d well separated sets in $\mathbb{R}^d$, and in weak general position, the object is to find the cut for given $a_1, \ldots, a_d$. Our proof of Theorem 2 leads to the formulation of an efficient, $O(n(\log n)^{d-3})$ algorithm for generalized cuts. This appears in Section 3. Throughout, because of space limitations, some details are omitted.

2 Proof of the Discrete Version.

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

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

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

Given points $p_i \in \text{Conv}(P_i), i = 1, \ldots, d$ (not necessarily data points in S), the hyperplane $h \equiv \text{aff}\{p_1, \ldots, p_d\}$ 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,

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

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

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

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

Proof of Theorem 2: The proof is by induction. The base case $d = 2$ is probably folklore (but see [11]). 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 [11] for intersection of median levels.

Next, suppose the claim holds for dimension $j < d$. Let π be a hyperplane that separates P_1 from $\bigcup_{i=2}^d P_i$. Fix a point $x \in \text{Conv}(P_1)$, project each data point $z \in \bigcup_{i=2}^d P_i$ onto π, and write P'_i for the set of images in π of the points $z \in P_i$.

Fact 1: $P'_2, \ldots, P'_d$ 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 $P_1, \ldots, P_d$, a contradiction. ■

Fact 2: If $P_1, \ldots, P_d$ have weak general position and if $x \in P_1$ then $P'_2, \ldots, P'_d$ have weak general position in π.

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 $P_1, \ldots, P_d$. ■

Therefore the induction hypotheses apply to $P'_2, \ldots, P'_d$.

Given a point $x \in \text{Conv}(P_1)$ and $(a_2, \ldots, a_d)$, a hyperplane h_x containing x is *an* $(a_2, \ldots, a_d)$ *semi-cut (or just a semi-cut)* if, for each $i > 1$, it is incident with a point $p_i \in P_i$ and $|h_x^+ \cap P_i| = a_i$. It's not hard to prove the following useful fact.

Lemma 3 *Given $x \in \text{Conv}(P_1)$ and $(a_2, \ldots, a_d)$. If there is an $(a_2, \ldots, a_d)$ semi-cut h_x then it is unique.*

To advance the induction, fix $(a_1, \ldots, a_d)$ and suppose h_x is a cut with these values, $x \in P_1$. By Lemma 3, it is the unique semi-cut cut containing x, so suppose there is an $(a_2, \dots, a_d)$ semi-cut h_y through $y \in P_1$, $y \notin h_x$. h_x and h_y cannot meet in $\mathrm{Conv}(P_1)$ since any such point would be in two different $(a_2, \dots, a_d)$ semi-cuts, violating Lemma 3. But this implies that $a_1 \neq |P_1 \cap h_y^+|$. Therefore h_x is unique, which proves (i).

Now suppose $P_1, \dots, P_d$ 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 $z \in \bigcup_{i=2}^{d} P_i$. x and ρ_x span a hyperplane h_x that is an $(m_x, a_2, \dots, a_d)$-cut, m_x denoting $|P_1 \cap h_x^+|$. Lemma 3 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 (ii). ■

In fact the conditions of the Theorem are also necessary.

Corollary 4 *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 $P_1, \dots, P_d$ 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 $A = \mathrm{Conv}(\bigcup_{i \in I} P_i) \cap \mathrm{Conv}(\bigcup_{j \in J} P_j) \neq \phi$. 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 half space containing $\bigcup_{i \in I} P_i$ also contains at least one point in $\bigcup_{j \in J} P_j$ in its interior. If we set $a_i = 1$ for $i \in I$, $a_i = |P_i|$ 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 there is a unique set of data points $p_1, \dots, p_d$, $p_i \in P_i$, for which $\mathrm{aff}(p_1, \dots, p_d)$ 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], [13]).

Proposition 2. *Given n points in $\mathbb{R}^d$ which are partitioned into d sets $P_1, \dots, P_d$ in $\mathbb{R}^d$, 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 $\mathbb{R}^{d-1}$. The latter problem can be solved within the following bounds:*

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

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

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 $O(n^{d-1-a(d)})$.

Finally, we will describe a much more practical algorithm, applying ideas from the proof in Section 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 $m_x = j$. Thus we could loop through all $x \in P_1$, project onto π, find the unique $(a_2, \dots, a_d)$ cut $\rho_x \subset \pi$, and count $m_x = |P_1 \cap h_x^+|$ for h_x, the hyperplane spanned by x and ρ_x. At some stage we will find the $z \in P_1$ 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 $(d-1)$-dimensional problems.

In fact we will find the desired $z \in P_1$ by solving at most $O(\log n)$ $(d-1)$-dimensional problems. The key is to be able to prune a fixed fraction of remaining points in P_1 after each search step, and uses the fact that if $n_x < a_1$, no point $y \in h_x^+ \cap P_1$ has $n_y = a_1$.

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**
   - **Construct A, an ϵ-approximation to C**
   - **FOR each $x \in A$ DO**
     (a) **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**
     (b) **Find the $(a_2, ..., a_d)$-cut $\rho_x \subset \pi$ for the projections $P_2', \dots, P_d'$ by solving a $(d-1)$-dimensional problem**

(c) **Get h_x, the hyperplane that spans x and ρ_x.**

(d) **Compute the number of points of C in the positive transversal halfspace h_x^+**

(e) **END FOR**

- **Prune from C points $x \in P_1$ whose m_x is too small or too large, and adjust C and a**
- **END WHILE**

6. **For each remaining data point in $x \in C$, project, find the $(a_2, ..., a_d)$-cut ρ_x in π for the projections by solving a $(d-1)$-dimensional problem, get h_x and compute $m_x = |P_1 \cap h_x^+|$, stopping when $m_x = a_1$.**

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. C is the set of candidates for the sought point $z \in P_1$; initially $C = P_1$. a denotes the number of undeleted points in the positive transversal halfspace of z's semicut; 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 its $O((d+1)^{3(d+1)}(\frac{d+1}{\epsilon^2} \log \frac{d+1}{\epsilon})^{d+1}|C|)$] 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 is traversed $k = |A|$ times. The cost of each traversal is dominated by $O(B_{d-1})$, the cost of the $(d-1)$-dimensional problem in (b); the cost of (a) is $O(n)$ and (d) 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 less than the smallest n_x, greater than the largest n_x, or 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 third 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 5a through 5d. Then, instead of Step 5e, 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, the algorithm just described now has running time $O(n(\log n)^{d-3})$ for dimensions $d \geq 3$.

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

Acknowledgement: We thank the reviewers for insightful comments.

References

[1] Bárány, I., Hubard, A., Jerónimo, J. " Slicing Convex Sets and Measures by a Hyperplane". *Discrete and Computational Geometry 39*, 67-75 (2008).

[2] I. Bárány and J. Matoušek. "Separating Measures by Two-Fans". *Discrete and Computational Geometry*

[3] I. Bárány and P. Valtr. "A positive fraction Erdős-Szekeres theorem". *Discrete and Computational Geometry 19*, 335-342 (1998).

[4] S. Bereg. "Equipartitions of Measures by 2-Fans". *Discrete and Computational Geometry 34 (1)*, 87-96 (2005).

[5] B. Chazelle: The Discrepancy Method: Randomness and Complexity, Cambridge University Press (2000).

[6] T. Dey. "Improved Bounds for Planar k-Sets and Related Problems". *Discrete and Computational Geometry 19 (3)* 373- 382 (1998).

[7] H. Edelsbrunner. Algorithms in Combinatorial Geometry. Springer, New York (1987)

[8] H. Kramer and A. Nemeth. "Supporting spheres for families of independent convex sets". *Arch. Math. 24*, 91-96 (1973).

[9] C.-Y. Lo, J. Matoušek, and W. Steiger. "Algorithms for Ham-Sandwich Cuts". *Discrete and Computational Geom. 11*, 433-452 (1994).

[10] J. Matoušek. *Using the Borsuk-Ulam Theorem.* Springer-Verlag (2003).

[11] N. Megiddo. "Partitioning with two lines in the plane". *J. Algorithms 6*, 430-433 (1985)

[12] S. Roy and W. Steiger. "Some Combinatorial and Algorithmic Aspects of the Borsuk-Ulam Theorem". *Graphs and Combinatorics 23*, 331-341, (2007).

[13] M. Sharir, S. Smorodinsky, and G. Tardos. "An Improved Bound for k-Sets in Three Dimensions". *Discrete and Computational Geometry 26*, 195-204 (2001).