Information Processing Letters 112 (2012) 532–534

Contents lists available at SciVerse ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Computing generalized ham-sandwich cuts ☆

Sergey Bereg

Department of Computer Science, University of Texas at Dallas, Box 830688, Richardson, TX 75083, USA

ARTICLE INFO

Article history:
Received 24 August 2011
Received in revised form 11 March 2012
Accepted 15 March 2012
Available online 19 March 2012
Communicated by F.Y.L. Chin

Keywords:
Computational geometry
Ham-sandwich theorem

ABSTRACT

Bárány et al. (2008) [1] proved that, for any $\beta \in [0,1]^d$ and any d well-separated convex bodies $S_1, S_2, \ldots, S_d$ in $\mathbb{R}^d$, there exists a hyperplane (a generalized ham-sandwich cut) splitting the volume of S_i as $(\beta_i, 1-\beta_i)$ for all i. Steiger and Zhao (2010) [4] proved a discrete analogue for n points in weak general position. The (elegant!) proof inspired an algorithm for computing this hyperplane (Steiger and Zhao, 2010 [4]) with running time $O(n \log^{d-3} n)$ for $d \geqslant 3$ and $O(n)$ for $d = 2$. In this note we show that their algorithm can be modified to compute a generalized ham-sandwich cut in linear time for any fixed d.

1. Introduction

Given d convex bodies in $\mathbb{R}^d$, a ham-sandwich cut is a hyperplane that cuts off half of the volume of each convex body. What if we wish to slice different portions of the bodies? Bárány et al. [1] found a sufficient condition for the existence of such a cut.

Definition 1. (See [3].) Given $k \leqslant d+1$, a family $S_1, \ldots, S_k$ of connected sets in $\mathbb{R}^d$ is *well separated* if, for every choice of $x_i \in S_i$, the affine hull of $x_1, \ldots, x_k$ is a $(k-1)$-dimensional flat in $\mathbb{R}^d$.

Theorem 2 *(Bárány et al.). (See [1].) Let $K_1, \ldots, K_d$ be well-separated convex bodies in $\mathbb{R}^d$, and let $\beta_1, \ldots, \beta_d$ given constants with $0 \leqslant \beta_i \leqslant 1$. Then there is a unique hyperplane h such that $\mathrm{Vol}(K_i \cap h^+) = \beta_i \cdot \mathrm{Vol}(K_i)$ for all i. By h^+ we assume the positive closed halfspace determined by h.*

Steiger and Zhao [4] studied a discrete version of the problem where convex bodies are replaced by points.

Definition 3. (See [4].) Points in $S = P_1 \cup \cdots \cup P_d$ are *well separated* if their convex hulls $\mathrm{Conv}(P_1), \ldots, \mathrm{Conv}(P_d)$ are well separated.

Definition 4. (See [4].) Points in $S = P_1 \cup \cdots \cup P_d$ are in *weak general position* if, for each $(x_1, \ldots, x_d)$, $x_i \in P_i$, the affine hull of $x_1, \ldots, x_d$ is a $(d-1)$-dimensional flat that contains no other point of S.

Definition 5. (See [4].) Given positive integers $a_i \leqslant |P_i|$, an $(a_1, \ldots, a_d)$*-cut* is a hyperplane h such that $h \cap P_i \neq \emptyset$ and $|h^+ \cap P_i| = a_i$, $1 \leqslant i \leqslant d$.

Theorem 6 *(Steiger and Zhao). (See [4].) If $P_1, \ldots, P_d$ are well-separated point sets in $\mathbb{R}^d$, and $a_1, \ldots, a_d$ are positive integers, $a_i \leqslant |P_i|$, then*

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

Steiger and Zhao [4] also found an efficient algorithm for computing generalized ham-sandwich cuts for points. Let $n = \sum |P_i|$ be the total number of points. An $(a_1, \ldots, a_d)$-cut can be found in time $O(n \log^{d-3} n)$, $d \geqslant 3$, and in linear time if $d = 2$. The running time exceeds linear bound if $d \geqslant 4$. Steiger and Zhao made a remark: "We

☆ This research has been partially supported by project FEDER MEC MTM2009-08652.

E-mail address: besp@utdallas.edu.

0020-0190/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ipl.2012.03.013

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 $\mathbb{R}^3$, but we did not succeed. A main open question is whether there is an on $O(n)$ algorithm for this problem." In this note we answer their question in the affirmative.

Theorem 7. *Let P be any set of n points in weak general position in $\mathbb{R}^d$. For any partition of P into well-separated finite point sets $P_1, \ldots, P_d$ in $\mathbb{R}^d$ and any positive integers $a_i \leqslant |P_i|$, an $(a_1, \ldots, a_d)$-cut exists and can be found in linear time.*

Our algorithm is a modification of the algorithm by Steiger and Zhao [4] which uses a prune-and-search approach. The main difference is that the pruning step may alternate dimensions.

2. Linear time algorithm

We briefly describe the algorithm by Steiger and Zhao [4] for computing a generalized ham-sandwich cut. The algorithm reduces the d-dimensional problem to a constant number of $(d-1)$-dimensional problems by repeating $O(\log n)$ pruning steps. The pruning step removes a portion of points in the first set P_1. Only a constant number of points remain in P_1 after pruning. Suppose that we want to compute an $(a_1', a_2, a_3, \ldots, a_d)$-cut after pruning P_1. For each remaining point p in P_1, the algorithm

(i) projects sets $P_2, P_3, \ldots, P_d$ onto a hyperplane separating P_1 from them (using the projection from p),
(ii) computes an $(a_2, a_3, \ldots)$-cut for the projected sets in the $(d-1)$-dimensional space, and
(iii) computes the hyperplane spanning the cut from (ii) and p.

Then the running time $T_d(n)$ of the d-dimensional problem satisfies the recurrence

$$T_d(n) = O(n) + O(\log n) \cdot T_{d-1}(n).$$

It results in $T_d(n) = O(n \log^{d-3} n)$ for $d \geqslant 3$ and $O(n)$ for $d = 2$ since linear time algorithms for computing ham-sandwich cuts in two and three dimensions can be adapted to generalized cuts.

2.1. Pruning

The pruning step is based on an ε-approximation of P_1. For a finite set of points S in $\mathbb{R}^d$ and the set of all halfspaces in $\mathbb{R}^d$, the range space (S, H) has VC dimension $d+1$ [2]. An ε-approximation of (S, H) is a set A of points in $\mathbb{R}^d$ such that, for any halfspace $h' \in H$,

$$\left| \frac{|h' \cap S|}{|S|} - \frac{|h' \cap A|}{|A|} \right| \leqslant \varepsilon. \tag{1}$$

By [2], an ε-approximation $A \subset P_1$ of (P_1, H) of constant size can be constructed in $O(|P_1|)$ time. To be more precise, $|A| = O(\frac{d+1}{\varepsilon^2} \log \frac{d+1}{\varepsilon})$ and the running time for computing A is $O((d+1)^{3(d+1)} (\frac{d+1}{\varepsilon^2} \log \frac{d+1}{\varepsilon})^{d+1} |P_1|)$.

The pruning step copies P_1 to a set C and is repeated until C contains at most c points, where c is a small constant. The number of steps is $O(\log n)$.

We modify the algorithm so that the set for pruning is not fixed but is selected as the largest set in $\{P_1, P_2, \ldots, P_d\}$. In order to prune this way, we first copy sets into new sets $Q_1, Q_2, \ldots, Q_d$ and copy the parameters a_i of the cut into an array $b_1, b_2, \ldots, b_d$. Coefficient b_i is adjusted after the pruning step where Q_i is pruned. After all pruning steps, each of the sets $Q_1, Q_2, \ldots, Q_d$ is small, i.e. $|Q_1|, |Q_2|, \ldots, |Q_d| \leqslant c$ for some constant c, and we compute a $(b_1, b_2, \ldots, b_d)$-cut using a brute-force search. The correctness of the modified algorithm follows from the correctness of the pruning in the algorithm by Steiger and Zhao [4]. We describe the pseudo-code for completeness.

Algorithm 1: GenCut($P_1, \ldots, P_d, a_1, \ldots, a_d$).

step 1 Let c be a small integer constant.
  for $i \leftarrow 1$ **to** d **do**
    $Q_i \leftarrow P_i$
    $b_i \leftarrow a_i$
    Find a hyperplane π_i that separates P_i from $P_1 \cup \cdots \cup P_{i-1} \cup P_{i+1} \cup \cdots \cup P_d$.
  while $|Q_i| > c$ *for some* i **do**
step 2 Find i such that $|Q_i| = \max_j |Q_j|$.
    Construct A, an ε-approximation of Q_i.
step 3 **foreach** $x \in A$ **do**
(a) **for** $j \leftarrow 1$ **to** d **do**
      if $j \neq i$ **then**
        Compute the projection of Q_j onto π_i.
(b) Find the $(b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_d)$-cut ρ_x in π_i for sets $Q_1, \ldots, Q_{i+1}, Q_{i+1}, \ldots, Q_d$ by solving a $(d-1)$-dimensional problem.
(c) Compute h_x, the hyperplane that spans x and ρ_x.
(d) Compute n_x, the number of points of Q_i in the positive transversal halfspace h_x^+.
    Prune points from Q_i and adjust b_i.
step 4 Find the $(b_1, \ldots, b_d)$-cut π for sets $Q_1, \ldots, Q_d$ by enumerating all d-tuples $(q_1, \ldots, q_d)$, $q_i \in Q_i$.
  return π

2.2. Running time analysis

We assume that dimension d is a constant. Let $T_d(n)$ be the running time for computing a cut for n points in d dimensions. Step 1, Steps 3a, c, d and the pruning after the loop (Step 3) take $O(n)$ time. Set A of constant size is constructed in Step 2 in $O(|Q_i|)$ time. Step 3b takes $O(T_{d-1}(n - |Q_i|))$ time. Step 4 takes $O(c^d) = O(1)$ time. Let Q_i' be the pruned set of Q_i in the iteration of the while loop, i.e. Q_i' is set Q_i in the next iteration. Then

$$T_d(n) = \begin{cases} |A| \cdot T_{d-1}(n - |Q_i|) + T_d(n - |Q_i| + |Q_i'|) \\ \quad + O(n) \quad \text{if } |Q_i| = \max_j |Q_j| > c, \\ O(1) \quad \text{otherwise.} \end{cases}$$

A linear bound for $T_d(n)$ follows from the following lemma.

Lemma 8. *For any $d \geqslant 2$, there exists a constant c_d such that $T_d(n) \leqslant c_d n$ for all $n \geqslant 1$.*

Proof. We prove the lemma by induction on d. We assume that $T_d(n) \leqslant C_d n$ for some constant C_d and $n \leqslant cd$. This covers the base case in Step 4.

Consider the inductive step. Suppose that $d \geqslant 3$ and $\varepsilon = 1/3$. We have

$$|Q_i| \geqslant \frac{n}{d} \quad \text{and} \quad n - |Q_i| \leqslant \frac{d-1}{d} n.$$

Next, we bound $|Q_i'|$ from below. $|A|$ hyperplanes constructed in Step 3b split Q_i into $|A|+1$ sets. Only one of these sets, Q_i', will survive. Suppose that $Q_i' = h_x^+ - h_y^+$ where h_x^+ and h_y^+ are closed halfspaces determined by hyperplanes h_x and h_y, respectively. Note that $h_x^+ \cap A = h_y^+ \cap A$. Let $a = |h_x^+ \cap A|/|A|$. Since A is an ε-approximation of (Q_i, H) [4], by inequality (1)

$$\left|\frac{|h_x' \cap Q_i|}{|Q_i|} - a\right| \leqslant \varepsilon \quad \text{and} \quad \left|\frac{|h_y' \cap Q_i|}{|Q_i|} - a\right| \leqslant \varepsilon.$$

Then

$$\frac{|Q_i'|}{|Q_i|} = \left|\frac{|h_x^+ \cap Q_i|}{|Q_i|} - \frac{|h_y^+ \cap Q_i|}{|Q_i|}\right| \leqslant 2\varepsilon = \frac{2}{3}.$$

Then

$$|Q_i| - |Q_i'| \geqslant |Q_i| - \frac{2}{3}|Q_i| = |Q_i|/3 \geqslant \frac{n}{3d}.$$

The last term in the recurrence for $T_d(n)$ can be bounded from above by $c_d' n$ for some constant c_d'. Then

$$\begin{aligned} T_d(n) &\leqslant |A| \cdot T_{d-1}\left(\frac{d-1}{d} n\right) + T_d\left(n - \frac{n}{3d}\right) + c_d' n \\ &\leqslant |A| \cdot c_{d-1} \frac{d-1}{d} n + c_d\left(n - \frac{n}{3d}\right) + c_d' n \\ &\leqslant c_d n + \left(|A| \cdot c_{d-1} \frac{d-1}{d} + c_d' - \frac{c_d}{3d}\right) n \leqslant c_d n \end{aligned}$$

if $c_d \geqslant 3(d-1)|A| \cdot c_{d-1} + 3dc_d'$. The lemma follows if we choose $c_d = \max(C_d, 4(d-1) \cdot |A| \cdot c_{d-1} + 3dc_d')$. □

This completes the proof of Theorem 7.

References

[1] I. Bárány, A. Hubard, J. Jerónimo, Slicing convex sets and measures by a hyperplane, Discrete Comput. Geom. 39 (1–3) (2008) 65–75.
[2] B. Chazelle, The Discrepancy Method: Randomness and Complexity, Cambridge University Press, 2000.
[3] H. Kramer, A. Nemeth, Supporting spheres for families of independent convex sets, Arch. Math. 24 (1973) 91–96.
[4] W. Steiger, J. Zhao, Generalized ham-sandwich cuts, Discrete Comput. Geom. 44 (3) (2010) 535–545.