Discrete Comput Geom 6:263–275 (1991)

Discrete & Computational
Geometry

Bisectors of Linearly Separable Sets

Lee R. Nackman and Vijay Srinivasan

Manufacturing Research Department, IBM Research Division,
Thomas J. Watson Research Center, P.O. Box 218,
Yorktown Heights, NY 10598, USA

Abstract. A bisector of two sets is the set of points equidistant from them. Bisectors arise naturally in several areas of computational geometry. We show that bisectors of weakly linearly separable sets in E^d have many properties of interest. Among these, the bisector of a restricted class of linearly separated sets is a homeomorphic image of the linear separator. We also give necessary and sufficient conditions for the existence of a particular continuous map from (a portion of) any linear separator to the bisector.

1. Introduction

Bisectors, which are defined as the set of points equidistant from two given sets, arise naturally in computing the symmetric axis transform [1] and in computing Voronoi diagrams [5–9]. For example, a common means of computing the Voronoi diagram is the following recursive approach: compute separately the Voronoi diagrams of two subproblems and then merge them together; during this merge, the bisector of the two subproblems is used to "trim" the Voronoi diagrams of the subproblems. Certain properties of the bisector are the key to an efficient merge. For example, one algorithm for computing the Voronoi diagram of point sites in E^2 using a divide-and-conquer strategy [8] partitions the points into two almost equal-sized sets separated by a line. The bisector between these two sets is connected and is a single-valued map of the dividing line. In the case of computing the Voronoi diagram of multiply connected polygonal domains [9], certain bisectors are simple, closed curves. In each of these two examples, properties of the bisector allow a linear-time merge. Moreover, knowing the topology of the bisector helps in choosing an appropriate data structure for the bisector. For example, if the bisector of two sets in E^3 is known to be a 2-manifold, then the QuadEdge data structure [2] can be used to represent and manipulate it.

In this paper we show that bisectors of linearly separable sets have many properties of interest. The results presented here are for general sets in E^d. There are several reasons for this. First, many of the results and proofs are simpler when the details of a particular class of sets do not intrude. More importantly, we want to broaden the study of bisectors beyond their use in algorithms for computing the Voronoi diagram of point sites in E^2 because we believe that proximity properties of more general geometric elements in higher dimensions have important applications.

2. Linearly Separable Sets

We denote the closure of a set S by cl S, the interior by int S, the boundary by ∂S, and the closure of the convex hull of S by CH(S). Boldface lowercase characters denote points in E^d and p_i denotes the ith coordinate of a point $\mathbf{p}$.

The (*Euclidean*) *distance* between two points $\mathbf{p}$ and $\mathbf{q}$ is denoted by $d(\mathbf{p}, \mathbf{q})$. The *distance* between a point $\mathbf{p}$ and a nonempty set S is $d(\mathbf{p}, S) = \text{glb}\{d(\mathbf{p}, \mathbf{q}): q \in S\}$. The *nearness* of S_1 and S_2 is $n(S_1, S_2) = \text{glb}\{d(\mathbf{p}, \mathbf{q}): \mathbf{p} \in S_1,\ \mathbf{q} \in S_2\}$. Let $\text{Sl}(\pi_1, \pi_2)$ denote the *open slab* between the two distinct parallel hyperplanes π_1 and π_2. S_1 and S_2 are *separated by a slab* $\text{Sl}(\pi_1, \pi_2)$ if S_1 and S_2 lie in different components of $E^d - \text{Sl}(\pi_1, \pi_2)$. S_1 and S_2 are *strongly linearly separable* if there exists an open slab that separates cl S_1 and cl S_2. A hyperplane contained in such an open slab is called a *strong linear separator*. S_1 and S_2 are *linearly separable* if there exists a hyperplane π, called a *linear separator*, such that cl S_1 and cl S_2 are in different components of $E^d - \pi$. Similarly, S_1 and S_2 are *weakly linearly separable* if there exists a

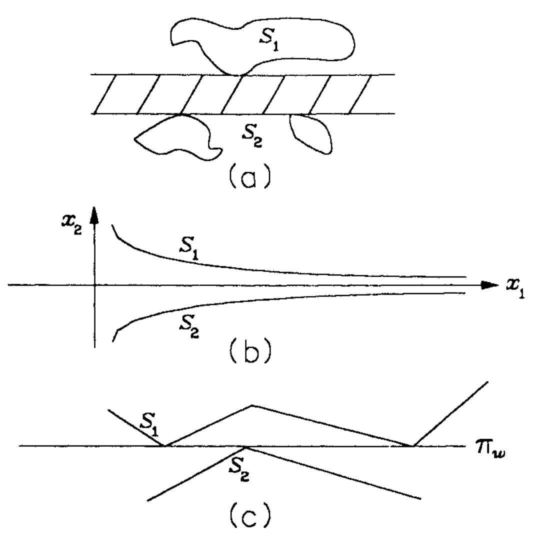

Fig. 1. Three types of linear separability. (a) S_1 and S_2 are strongly linearly separable since they are separated by the cross-hatched slab. (b) $S_1: x_2 = x_1^{-1}$, $S_2: x_2 = -x_1^{-1}$, $x_1 > 0$, are linearly separable (but not strongly linearly separable) and $x_2 = 0$ is the linear separator. (c) S_1 and S_2 are weakly linearly separable (but not linearly separable) and π_w is the weak linear separator.

hyperplane π_w called a *weak linear separator*, such that cl S_1 and cl S_2 lie in the closures of different components of $E^d - \pi$. Figure 1 shows some examples in E^2 of separable sets and separators.

Note that strongly linearly separable sets are also linearly separable and linearly separable sets are also weakly linearly separable. Likewise, a strong linear separator is also a linear separator and a linear separator is also a weak linear separator. Most results presented in this paper are formulated in terms of weakly linearly separable sets and weak linear separators.

A practically useful characterization of strong linear separability is given by:

Theorem 1. *If at least one of S_1 and S_2 is bounded, then S_1 and S_2 are strongly linearly separable if and only if* $\mathrm{CH}(S_1) \cap \mathrm{CH}(S_2) = \varnothing$.

Proof. Since at least one of S_1 and S_2 is bounded, $\mathrm{CH}(S_1) \cap \mathrm{CH}(S_2) = \varnothing$ implies that $n(\mathrm{CH}(S_1), \mathrm{CH}(S_2)) > 0$ (Corollary 7.2, Section 1.8, of [4]). Since the nearness is positive, there exists a slab that separates $\mathrm{CH}(S_1)$ and $\mathrm{CH}(S_2)$ (Theorem 4, Section 6.2, of [4]). S_1 and S_2 are therefore strongly linearly separable. The converse is trivial. □

3. Bisectors

In this section we investigate properties of the bisector of two weakly linearly separable sets S_1 and S_2. These are of use in Section 4 in studying the topology of the bisector. Hereafter, we assume that S_1 and S_2 are nonempty sets in E^d such that cl $S_1 \cap$ cl $S_2 = \varnothing$.

The *bisector* $B(S_1, S_2)$ of two sets S_1 and S_2 is the set of points equidistant from S_1 and S_2, i.e., $B(S_1, S_2) = \{\mathbf{p} \in E^d: d(\mathbf{p}, S_1) = d(\mathbf{p}, S_2)\}$. Let $\mathbf{q}$ be a point on $B(S_1, S_2)$ and define the *maximal ball* $\mathscr{B}$ to be the open ball centered at $\mathbf{q}$ with radius $r = d(\mathbf{q}, S_1) = d(\mathbf{q}, S_2)$. Also define the *maximal sphere* $\mathscr{S} = \partial\mathscr{B}$. Observe that, by definition, $\mathscr{B}$ does not contain points of S_1 or S_2 and that $\mathscr{S}$ contains at least one point $\mathbf{p}_1$ from cl S_1 and at least one point $\mathbf{p}_2$ from cl S_2. The points $\mathbf{p}_1$ and $\mathbf{p}_2$ are called *touching points* of $\mathscr{S}$ on S_1 and S_2, respectively, and $\mathscr{S}$ is said to *touch* S_1 and S_2. Notice that $\mathbf{p}_1 \neq \mathbf{p}_2$ as cl $S_1 \cap$ cl $S_2 = \varnothing$. One maximal ball cannot contain another because the included maximal sphere would lack touching points.

Without loss of generality, we assume that a weak linear separator π_w is the hyperplane $x_d = 0$, that S_1 is contained in the closure of the open half-space π_w^1: $x_d > 0$, and that S_2 is contained in the closure of the open half-space π_w^2: $x_d < 0$. The mapping we consider takes a point on a weak linear separator vertically up or down to a point in the bisector. We first show in Section 3.1 that if a line perpendicular to π_w intersects $B(S_1, S_2)$, then the intersection is connected. Then, in Section 3.2, we establish the necessary and sufficient conditions under which all lines perpendicular to π_w intersect $B(S_1, S_2)$. Finally, in Section 3.3, we give the necessary and sufficient conditions for a specific line perpendicular to π_w to intersect $B(S_1, S_2)$.

3.1. Connectedness of Intersection with a Perpendicular

Theorem 2. *If a line l perpendicular to a weak linear separator π_w of S_1 and S_2 intersects $B(S_1, S_2)$, then it does so in a connected subset of l. Moreover, if a point $\mathbf{q}$ is in the relative interior of $l \cap B(S_1, S_2)$, then all the touching points of the maximal sphere centered at $\mathbf{q}$ are in π_w.*

Proof. To show that the intersection is connected, assume that there are two distinct points $\mathbf{q}_a$ and $\mathbf{q}_b$ in $l \cap B(S_1, S_2)$. Let $\mathscr{B}_a$, $\mathscr{S}_a$, $\mathscr{B}_b$, and $\mathscr{S}_b$ be the corresponding maximal balls and spheres. Since $\mathbf{q}_a$ and $\mathbf{q}_b$ are distinct, $\mathscr{B}_a \neq \mathscr{B}_b$. Since one maximal ball cannot be contained in another, there are three remaining cases:

Case 1: *$\mathscr{S}_a$ and $\mathscr{S}_b$ are disjoint.* Since $\mathbf{q}_a$ and $\mathbf{q}_b$ are in l, there exists a hyperplane λ perpendicular to l that is a strong linear separator of $\mathscr{S}_a$ and $\mathscr{S}_b$. See Fig. 2(a). Since λ is parallel to π_w, $\mathscr{S}_a$ cannot touch S_2 and/or $\mathscr{S}_b$ cannot touch S_1—a contradiction.

Case 2: *$\mathscr{S}_a$ and $\mathscr{S}_b$ intersect at one point* $\mathbf{p}$. Let λ be the hyperplane perpendicular to l through $\mathbf{p}$. See Fig. 2(b). $\mathscr{S}_a$ and $\mathscr{S}_b$ can touch both S_1 and S_2 only if $\mathbf{p} \in \text{cl } S_1 \cap \text{cl } S_2$, which violates the assumption that $\text{cl } S_1 \cap \text{cl } S_2 = \emptyset$.

Case 3: *$\mathscr{S}_a$ and $\mathscr{S}_b$ intersect in a nondegenerate lower-dimensional sphere.* Let λ be the hyperplane through the sphere of intersection. Note that λ is parallel to π_w. If $\lambda \neq \pi_w$, the sets of touching points of $\mathscr{S}_a$ and of $\mathscr{S}_b$ lie in opposite closed half-spaces bounded by λ. See Fig. 2(c). Therefore, since λ is parallel to π_w $\mathscr{S}_a$ cannot touch S_2 and/or $\mathscr{S}_b$ cannot touch S_1, which is a contradiction. Thus $\lambda = \pi_w$. $\mathscr{S}_a$ and $\mathscr{S}_b$ can each touch points in both S_1 and S_2 only if one of the touching points on S_1 and one of the touching points on S_2 are both on π_w. These touching points are contained in $\mathscr{S}_a \cap \mathscr{S}_b$. Let $\mathbf{q}$ be strictly between $\mathbf{q}_a$ and $\mathbf{q}_b$, and let $\mathscr{S}_q$ be the sphere centered at $\mathbf{q}$ and passing through $\mathscr{S}_a \cap \mathscr{S}_b$. See Fig. 2(d). Clearly the open ball $\mathscr{B}_q$ defined by $\mathscr{S}_q$ is contained in $\mathscr{B}_a \cup \mathscr{B}_b$ and hence is free of points of $S_1 \cup S_2$. Thus $\mathscr{S}_q$ is a maximal sphere touching both S_1 and S_2 in π_w, and $\mathbf{q} \in B(S_1, S_2)$.

The theorem follows directly. □

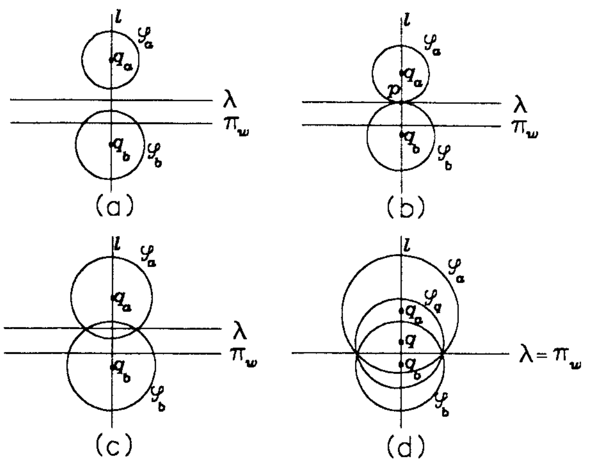

Fig. 2. Illustration of various cases for proof of Theorem 2.

3.2. *Intersection with Every Perpendicular*

In Section 3.1 we proved that if the intersection between the bisector and a line perpendicular to a weak linear separator exists, then it is connected; in this section we give necessary and sufficient conditions for every line perpendicular to the weak linear separator to intersect the bisector.

Let $I = \{\mathbf{x} \in E^{d-1}; \alpha_i \le x_i \le \beta_i,\ i = 1, 2, \ldots, d-1\}$, for real numbers $\alpha_i < \beta_i$, $i = 1, 2, \ldots, d-1$, be a closed $(d-1)$-cell in $\pi_{\mathbf{w}}$.

Theorem 3. *Every line perpendicular to $\pi_{\mathbf{w}}$ intersects $B(S_1, S_2)$ if and only if*

(1) $S_1, S_2 \not\subset \pi_{\mathbf{w}}$, *or*
(2) $S_2 \subset \pi_{\mathbf{w}}$ *and* $\mathrm{cl}\, S_1 \cap \pi_{\mathbf{w}} = \varnothing$, *or vice versa.*

Moreover, if every line perpendicular to $\pi_{\mathbf{w}}$ intersects $B(S_1, S_2)$, then, for any $(d-1)$-cell I in $\pi_{\mathbf{w}}$, $B(S_1, S_2) \cap (I \times \mathbb{R})$ is bounded. □

We prove Theorem 3 by showing the sufficient conditions in Lemmas 1 and 2 and the necessary condition in Lemma 3.

Lemma 1. *If $S_1, S_2 \not\subset \pi_{\mathbf{w}}$, then every line perpendicular to $\pi_{\mathbf{w}}$ intersects $B(S_1, S_2)$. Moreover, $B(S_1, S_2) \cap (I \times \mathbb{R})$ is bounded.*

Proof. Let $\mathbf{q}$ be a point on the line $l = \{\mathbf{x}: x_i = q_i, i = 1, \ldots, d-1\}$ perpendicular to $\pi_{\mathbf{w}}$. Define $f(\mathbf{q}) = d^2(\mathbf{q}, S_1) - d^2(\mathbf{q}, S_2)$, and observe that $\mathbf{q} \in B(S_1, S_2)$ if and only if $f(\mathbf{q}) = 0$. First assume that $q_d > 0$. Let $\mathbf{p} \in S_1, p_d > 0$; such a point must exist since $S_1 \not\subset \pi_{\mathbf{w}}$. See Fig. 3. Since $d^2(\mathbf{q}, S_1) \le d^2(\mathbf{q}, \mathbf{p}) = \sum_{i=1}^{d} (p_i - q_i)^2$ and $d^2(\mathbf{q}, S_2) \ge d^2(\mathbf{q}, \pi_{\mathbf{w}}) = q_d^2$, $f(\mathbf{q}) \le \sum_{i=1}^{d} (p_i - q_i)^2 - q_d^2 = \sum_{i=1}^{d-1} (p_i - q_i)^2 + p_d^2 - 2p_d q_d$. For sufficiently large q_d, $f(\mathbf{q}) < 0$. By a symmetric argument, for sufficiently small $q_d < 0$, $f(\mathbf{q}) > 0$. But $d(\mathbf{p}, S_i)$ is a continuous function of $\mathbf{p}$ (Theorem 3,

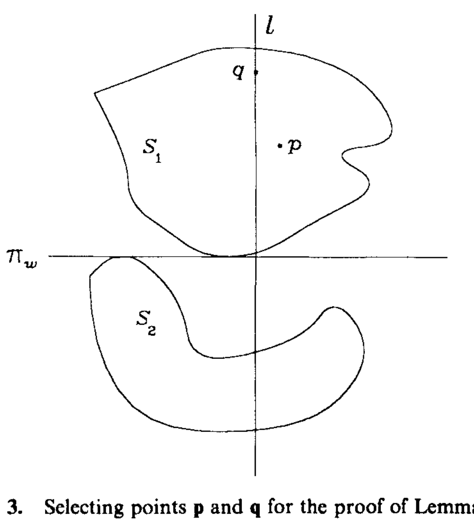

Fig. 3. Selecting points $\mathbf{p}$ and $\mathbf{q}$ for the proof of Lemma 1.

Section 1.8, of [4]). Therefore, since f changes sign, it must have at least one zero, which implies that the line l intersects $B(S_1, S_2)$.

It remains to show that $B(S_1, S_2) \cap (I \times \mathbb{R})$ is bounded. Observe that for $q_d > \left(\sum_{i=1}^{d-1} (p_i - q_i)^2 + p_d^2\right)/2p_d$, $f(\mathbf{q}) < 0$. Thus the intersection of $B(S_1, S_2)$ with l is bounded in the positive x_d-direction by a continuous function of $q_1, \ldots, q_{d-1}$. It is similarly bounded in the negative x_d-direction. In particular, as I is compact, $B(S_1, S_2) \cap (I \times \mathbb{R})$ is bounded. □

Lemma 2. *If $S_2 \subset \pi_w$ and $\text{cl}\, S_1 \cap \pi_w = \emptyset$, then every line perpendicular to π_w intersects $B(S_1, S_2)$. Moreover, $B(S_1, S_2) \cap (I \times \mathbb{R})$ is bounded.*

Proof. Again, let $\mathbf{q}$ be a point on a line l perpendicular to π_w and vary q_d so that $\mathbf{q}$ moves along l. By the arguments in the proof of Lemma 1, $f(\mathbf{q}) = d^2(\mathbf{q}, S_1) - d^2(\mathbf{q}, S_2) < 0$ for sufficiently large $q_d > 0$.

We now show that there exists a $q_d \leq 0$ such that $f(\mathbf{q}) \geq 0$. Let $\mathbf{u} = l \cap \pi_w$ and $\mathbf{p}$ be a point of $\text{cl}\, S_2$ closest to $\mathbf{u}$. If $\mathbf{p} = \mathbf{u}$, then $d^2(\mathbf{q}, S_2) = q_d^2$. Therefore, since $d^2(\mathbf{q}, S_1) > d^2(\mathbf{q}, \pi_w) = q_d^2$ for all $q_d \leq 0$, $f(\mathbf{q}) > 0$ for all $q_d \leq 0$.

If $\mathbf{p} \neq \mathbf{u}$, then consider an open ball $\mathbf{B_u}$ of radius $d(\mathbf{u}, \mathbf{p})$ centered at $\mathbf{u}$. If $\text{cl}\, S_1 \cap \text{cl}\, \mathbf{B_u} = \emptyset$, then $d^2(\mathbf{u}, S_1) > d^2(\mathbf{u}, S_2)$, which leads to $f(\mathbf{u}) > 0$. Otherwise, since $\text{cl}\, S_1 \cap \text{cl}\, \mathbf{B_u}$ is compact, there exists a point $\mathbf{t} \in \text{cl}\, S_1 \cap \text{cl}\, \mathbf{B_u}$ with smallest x_d-coordinate. Moreover, $t_d > 0$ because $\text{cl}\, S_1 \cap \pi_w = \emptyset$. See Fig. 4. For $q_d \leq 0$, $d^2(\mathbf{q}, S_1) \geq (t_d - q_d)^2$. Therefore, since $d^2(\mathbf{q}, S_2) = \sum_{i=1}^{d} (p_i - q_i)^2$, $f(\mathbf{q}) \geq t_d^2 - 2t_d q_d - \sum_{i=1}^{d-1} (p_i - q_i)^2$ if $q_d \leq 0$. For sufficiently small $q_d \leq 0$, $f(\mathbf{q}) \geq 0$. Thus $f(\mathbf{q})$ has a zero and the intersection result follows. Boundedness is guaranteed, since $\mathbf{t}$ is confined to a bounded set as long as $\mathbf{u} \in I$. □

To complete the proof of Theorem 3, it remains to prove the necessary condition.

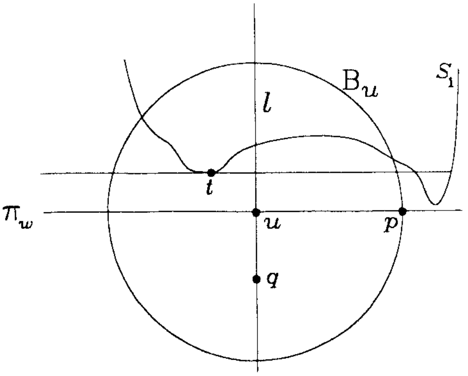

Fig. 4. Construction for proof of Lemma 2.

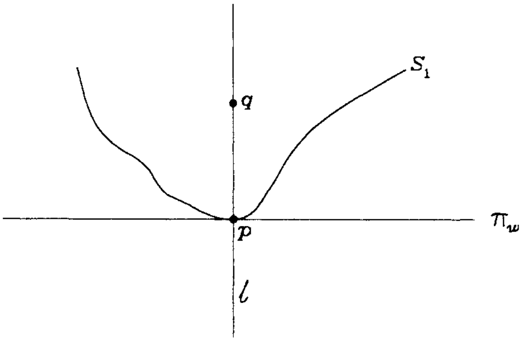

Fig. 5. Construction for proof of Lemma 3.

Lemma 3. *If $S_2 \subset \pi_w$ and $\text{cl}\, S_1 \cap \pi_w \neq \emptyset$, then there exist lines perpendicular to π_w that do not intersect $B(S_1, S_2)$.*

Proof. Consider the line l perpendicular to π_w through a point $\mathbf{p} \in \text{cl}\, S_1 \cap \pi_w$. See Fig. 5. For any $\mathbf{q} \in l$, $d^2(\mathbf{q}, S_1) \leq d^2(\mathbf{q}, \mathbf{p}) = q_d^2$. Since $\mathbf{p} \notin \text{cl}\, S_2$ and $S_2 \subset \pi_w$, $d^2(\mathbf{q}, S_2) > q_d^2$. Therefore $d^2(\mathbf{q}, S_1) < d^2(\mathbf{q}, S_2)$, $\forall \mathbf{q} \in l$, so l cannot intersect $B(S_1, S_2)$. □

Figure 6 shows two examples as an illustration of Lemma 3. From Theorems 2 and 3 we have:

Corollary 1. *If S_1 and S_2 are weakly linearly separated by π_w and $\text{cl}\, S_1 \subset \pi_w^1$, then every line perpendicular to π_w intersects $B(S_1, S_2)$ in a single point.* □

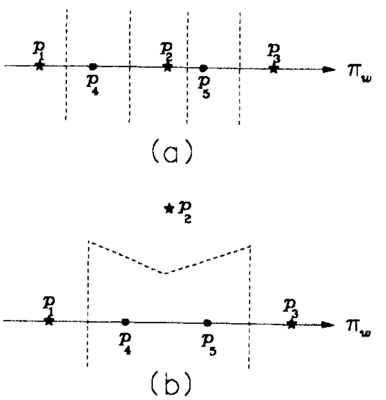

Fig. 6. Sets illustrating the condition of Lemma 3. Let $S_1 = \{\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3\}$ and $S_2 = \{\mathbf{p}_4, \mathbf{p}_5\}$. (a) Both S_1 and S_2 are completely contained in their weak linear separator π_w. $B(S_1, S_2)$ consists of just four lines (shown dashed), all of which are perpendicular to π_w. (b) S_1 is only partially contained in π_w. In both cases, there exist lines perpendicular to π_w that do not intersect $B(S_1, S_2)$.

As a special case we also have:

Corollary 2. *If a hyperplane π is a linear separator of S_1 and S_2, then every line perpendicular to π intersects $B(S_1, S_2)$ in a single point.* □

3.3. *Intersection with a Specific Perpendicular*

Theorem 3 gave the necessary and sufficient conditions under which *every* line perpendicular to π_w intersects $B(S_1, S_2)$. Even if such global conditions do not hold, it is still possible to obtain local results. This section gives necessary and sufficient conditions under which any *specific* line perpendicular to π_w intersects $B(S_1, S_2)$.

Lemma 4. *Let $S_1, S_2 \subset \pi_w$ and let $\mathbf{u}$ be a point in π_w. A line l perpendicular to π_w and passing through $\mathbf{u}$ intersects $B(S_1, S_2)$ if and only if $d(\mathbf{u}, S_1) = d(\mathbf{u}, S_2)$. Moreover, if l intersects $B(S_1, S_2)$, then $l \subset B(S_1, S_2)$.*

Proof. Let $\mathbf{q} \in l$. Then $d^2(\mathbf{q}, S_i) = d^2(\mathbf{u}, S_i) + q_d^2$, for $i = 1, 2$. Therefore, if $d(\mathbf{u}, S_1) = d(\mathbf{u}, S_2)$, then $d(\mathbf{q}, S_1) = d(\mathbf{q}, S_2)$, which implies that $l \subset B(S_1, S_2)$. Conversely, if $\mathbf{q} \in B(S_1, S_2)$, then $d^2(\mathbf{q}, S_1) = d^2(\mathbf{q}, S_2)$ and hence $d(\mathbf{u}, S_1) = d(\mathbf{u}, S_2)$. □

Lemma 5. *Let $S_2 \subset \pi_w$, $S_1 \not\subset \pi_w$, $\mathrm{cl}\, S_1 \cap \pi_w \neq \varnothing$, and $\mathbf{u} \in \pi_w$ such that $d(\mathbf{u}, S_2) \neq d(\mathbf{u}, \mathrm{cl}\, S_1 \cap \pi_w)$. A line l perpendicular to π_w and passing through $\mathbf{u}$ intersects $B(S_1, S_2)$ if and only if $d(\mathbf{u}, S_2) < d(\mathbf{u}, \mathrm{cl}\, S_1 \cap \pi_w)$. Moreover, if l intersects $B(S_1, S_2)$, then it does so in a single point.*

Proof. Let $\mathbf{q}$ be a point in l. See Fig. 7. Sufficiency follows from arguments similar to those of the proof of Lemma 2. To show necessity, assume that $\mathbf{q} \in B(S_1, S_2)$. Since $d^2(\mathbf{q}, S_1) = d^2(\mathbf{q}, S_2)$, $d^2(\mathbf{q}, S_2) = d^2(\mathbf{u}, S_2) + q_d^2$, $d^2(\mathbf{q}, \mathrm{cl}\, S_1 \cap \pi_w) = d^2(\mathbf{u}, \mathrm{cl}\, S_1 \cap \pi_w) + q_d^2$, and $d^2(\mathbf{q}, S_1) \le d^2(\mathbf{q}, \mathrm{cl}\, S_1 \cap \pi_w)$, we have $d^2(\mathbf{u}, S_2) + q_d^2 \le d^2(\mathbf{u}, \mathrm{cl}\, S_1 \cap \pi_w) + q_d^2$. Necessity follows since, by hypothesis, $d(\mathbf{u}, S_2) \neq d(\mathbf{u}, \mathrm{cl}\, S_1 \cap \pi_w)$.

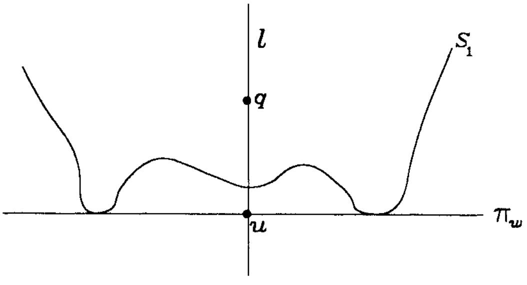

Fig. 7. Construction for proof of Lemma 5.

To see that the intersection is a single point, assume the contrary. Then, by Theorem 2, there exists $\mathbf{q} \in \text{int}[l \cap B(S_1, S_2)]$ such that the touching points of the maximal sphere centered at $\mathbf{q}$ are contained in π_w. This implies that the touching points in S_1 are in $\text{cl } S_1 \cap \pi_w$ and that $d^2(\mathbf{q}, S_2) = d^2(\mathbf{q}, \text{cl } S_1 \cap \pi_w)$ — a contradiction. □

Let $\mathbf{B}$ be an open $(d-1)$-dimensional ball and let $\mathbf{v}$ be a point on a line perpendicular to the hyperplane that contains $\mathbf{B}$, and passing through the center of $\mathbf{B}$. The *truncated semicone* $C(\mathbf{v}, \mathbf{B})$ is defined to be $\text{int CH}(\mathbf{v} \cup \mathbf{B})$.

Lemma 6. *Let* $S_2 \subset \pi_w$, $S_1 \not\subset \pi_w$, $\text{cl } S_1 \cap \pi_w \neq \emptyset$, *and* $\mathbf{u} \in \pi_w$ *such that* $d(\mathbf{u}, S_2) = d(\mathbf{u}, \text{cl } S_1 \cap \pi_w)$. *Let* $\mathbf{B} \subset \pi_w$ *be the* $(d-1)$*-dimensional ball of radius* $d(\mathbf{u}, S_2)$ *centered at* $\mathbf{u}$. *A line* l *perpendicular to* π_w *and passing through* $\mathbf{u}$ *intersects* $B(S_1, S_2)$ *if and only if there exists a point* $\mathbf{q} \in l \cap \pi_w^1$ *such that* $C(\mathbf{q}, \mathbf{B}) \cap S_1 = \emptyset$. *Furthermore, if* l *intersects* $B(S_1, S_2)$, *then a half-line of* l *is contained in* $B(S_1, S_2)$.

Proof. To show sufficiency, let there be a point $\mathbf{q}_1 \in l \cap \pi_w^1$ such that $C(\mathbf{q}_1, \mathbf{B}) \cap S_1 = \emptyset$. See Fig. 8. Now, consider the one-parameter family of d-dimensional balls (and the associated boundary spheres) the intersect π_w in $\mathbf{B}$ and whose centers lie on l. Some of these boundary spheres must intersect π_w^1 within $C(\mathbf{q}_1, \mathbf{B})$. Since $\mathbf{B}$ intersects neither $\text{cl } S_1$ nor $\text{cl } S_2$, but the boundary of $\mathbf{B}$ touches both $\text{cl } S_1$ and $\text{cl } S_2$, such spheres must be maximal spheres. Furthermore, any member of the family whose center has smaller x_d-coordinate must also be maximal. Thus, a half-line of l is contained in $B(S_1, S_2)$.

To show necessity, consider a maximal sphere $\mathscr{S}$ (and the associated ball $\mathscr{B}$) centered at $\mathbf{q} \in l \cap B(S_1, S_2)$. Note that $\mathbf{u} \notin \text{cl } (S_1 \cup S_2)$ because $d(\mathbf{u}, S_2) = d(\mathbf{u}, \text{cl } S_1 \cap \pi_w)$ and $\text{cl } S_1 \cap \text{cl } S_2 = \emptyset$. Therefore, $\exists \mathbf{q}_2 \in \pi_w^1 \cap l \cap \mathscr{S}$. See Fig. 8. Since $C(\mathbf{q}_2, \mathbf{B}) \subset \mathscr{B}$, $C(\mathbf{q}_2, \mathbf{B}) \cap S_1 = \emptyset$. □

An example where the conditions of Lemma 6 do not hold is shown in Fig. 9.

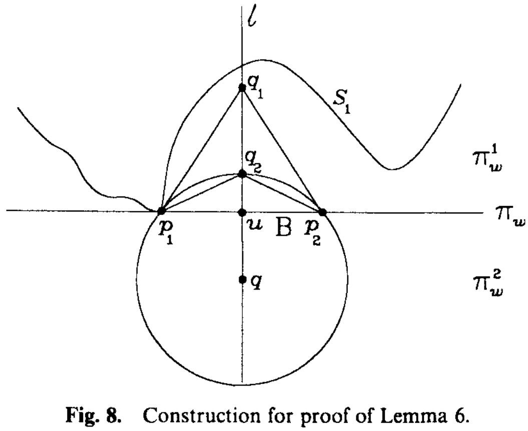

Fig. 8. Construction for proof of Lemma 6.

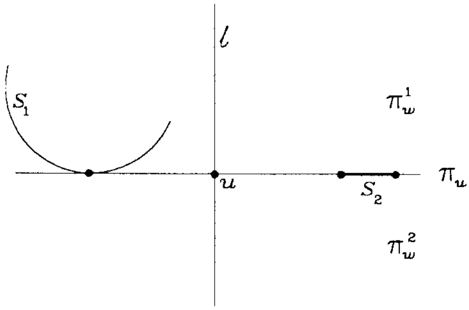

Fig. 9. An example where the conditions of Lemma 6 do not hold. S_1 is an arc of a circle, $S_2 \subset \pi_w$ and $d(\mathbf{u}, S_2) = d(\mathbf{u}, \text{cl } S_1 \cap \pi_w)$. The line l will not intersect $B(S_1, S_2)$.

Necessary and sufficient conditions for any specific line perpendicular to π_w to intersect $B(S_1, S_2)$ follow directly from Theorem 3 and Lemmas 4–6.

Theorem 4. *Let $\mathbf{u} \in \pi_w$ and let l be the line perpendicular to π_w passing through $\mathbf{u}$. The line l intersects $B(S_1, S_2)$ if and only if (up to a switch of S_1 and S_2)*

1. $S_1 \not\subset \pi_w$ and $S_2 \not\subset \pi_w$, *or*
2. $S_2 \subset \pi_w$, *and* $\text{cl } S_1 \cap \pi_w = \emptyset$, *or*
3. $S_1 \subset \pi_w$, $S_2 \subset \pi_w$ *and* $d(\mathbf{u}, S_1) = d(\mathbf{u}, S_2)$, *or*
4. $S_2 \subset \pi_w$, $S_1 \not\subset \pi_w$, $\text{cl } S_1 \cap \pi_w \neq \emptyset$, *and* $d(\mathbf{u}, S_2) < d(\mathbf{u}, \text{cl } S_1 \cap \pi_w)$, *or*
5. $S_2 \subset \pi_w$, $S_1 \not\subset \pi_w$, $\text{cl } S_1 \cap \pi_w \neq \emptyset$, $d(\mathbf{u}, S_2) = d(\mathbf{u}, \text{cl } S_1 \cap \pi_w)$, *and there exists a point* $\mathbf{q} \in l \cap \pi_w^1$ *such that* $C(\mathbf{q}, \mathbf{B}) \cap S_1 = \emptyset$, *where* $\mathbf{B} \subset \pi_w$ *is the open* $(d-1)$*-dimension ball of radius* $d(\mathbf{u}, S_2)$ *centered at* $\mathbf{u}$.

4. Continuous Mapping from Linear Separator to Bisector

We have thus far shown exactly when a line perpendicular to π_w intersects $B(S_1, S_2)$ at a single point. This defines a mapping which lifts points of the separator up to the bisector. We now show that wherever such a mapping exists, it is continuous. Notice that this map will automatically be a homeomorphism, as its inverse is the orthogonal projection onto the separator — certainly a well-defined continuous map.

Theorem 5. *Let the hyperplane π_w: $x_d = 0$ be a weak linear separator of S_1 and S_2 and let M be a relatively open subset of π_w such that $\forall \mathbf{p} \in M$, the line through $\mathbf{p}$ perpendicular to π_w intersects $B(S_1, S_2)$ in a single point. Then the mapping b: $M \to \mathbb{R}$ such that $(x_1, \ldots, x_{d-1}, b(x_1, \ldots, x_{d-1})) \in B(S_1, S_2)$ is continuous.*

To prove this we use the following lemma. Let E and F be topological spaces. A function g: $E \to F$ is said to have a *closed graph* if its graph $\{(x, y): y = g(x), x \in E\}$ in the product space $E \times F$ is a closed set.

Lemma 7 (Theorem 3, Chapter VII of [3]). *Let E, F be topological spaces and let $g: E \to F$ be a function. If g has a closed graph and F is compact, then g is continuous.*

Proof of Theorem 5. The mapping b is a function by hypothesis. We show that b is continuous at an arbitrary $\mathbf{p} \in M$. Let $I = \{\mathbf{x} \in E^{d-1}: \alpha_i \leq x_i \leq \beta_i, i = 1, 2, \ldots, d - 1\}$, for real numbers $\alpha_i < \beta_i, i = 1, 2, \ldots, d - 1$, be a closed $(d - 1)$-cell in M containing $\mathbf{p}$ in its interior and let b_I denote the restriction of b to I. We claim that b_I has a compact graph and thus, by Lemma 7, b is continuous at $\mathbf{p}$.

The graph of b_I is the intersection of $B(S_1, S_2)$ and $I \times \mathbb{R}$. Referring to the proof of Lemma 1, $B(S_1, S_2) = f^{-1}(0)$. Since the inverse image of a closed set under a continuous map is closed (p. 35 of [4]), $B(S_1, S_2)$ is closed. Therefore, the graph of b_I is closed.

It remains to show that the range of b is also bounded and, hence, compact. Since the line perpendicular to each point in M intersects $B(S_1, S_2)$, one of the five conditions of Theorem 4 must hold for each point in M. Furthermore, since by hypothesis, each such line intersects $B(S_1, S_2)$ in a single point, Lemmas 4 and 6 imply that, up to a switch of S_1 and S_2, one of the following must hold:

1. $S_1 \not\subset \pi_w$ and $S_2 \not\subset \pi_w$, or
2. $S_2 \subset \pi_w$, and $\mathrm{cl}\, S_1 \cap \pi_w = \emptyset$, or
3. $S_2 \subset \pi_w$, $S_1 \not\subset \pi_w$, $\mathrm{cl}\, S_1 \cap \pi_w \neq \emptyset$, and $d(\mathbf{u}, S_2) < d(\mathbf{u}, \mathrm{cl}\, S_1 \cap \pi_w)$, $\forall \mathbf{u} \in M$.

In the first two cases, Theorem 3 establishes that the graph is bounded. In the third case, Lemma 5 and the fact that I is compact establish that the graph is bounded. Therefore, the graph of b_I is compact, which implies that b is continuous at every point of M. □

When π_w is a linear separator, we have:

Corollary 3. *Let the hyperplane $\pi: x_d = 0$ be a linear separator of S_1 and S_2. If b is the mapping $b: \pi \to \mathbb{R}$ such that $(x_1, \ldots, x_{d-1}, b(x_1, \ldots, x_{d-1})) \in B(S_1, S_2)$, then b is a continuous function. In fact, the perpendicular projection of $B(S_1, S_2)$ onto π is a homeomorphism.*

Corollary 3 generalizes the notion described in [8] that the bisector of linearly separated point sites in E^2 is a monotone chain. More importantly, it shows that $B(S_1, S_2)$ is a $(d - 1)$-manifold in E^d.

5. Summary

In this paper we have presented some general properties of bisectors of sets in E^d that are separated by hyperplanes. We have given necessary and sufficient conditions for the perpendicular projection of the bisector of two weakly linearly separated sets onto a separator to be a homeomorphism. This study needs to be expanded in two major directions.

Throughout this paper we have required that $\mathrm{cl}\, S_1 \cap \mathrm{cl}\, S_2 = \emptyset$. When the closures of the sets are not disjoint, the bisector need not be a manifold, as

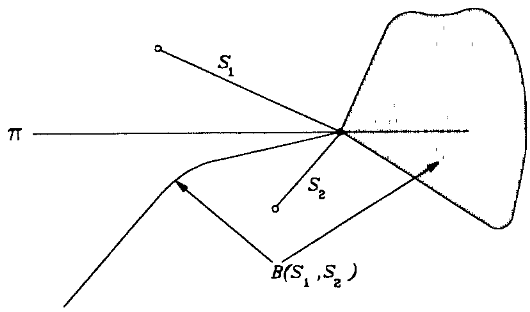

Fig. 10. A nonhomogeneously two-dimensional bisector. $B(S_1, S_2)$ is the bisector between two open line-segments that do not intersect. However, the closures of the line-segments intersect at an endpoint.

illustrated in Fig. 10. This issue has been addressed in the literature on Voronoi diagrams of point sites and open line segments in E^2 by defining the bisectors between individual elements *ab initio* so that they are always homogeneously one-dimensional [5]-[7], [9]. We think that this needs further investigation.

Another direction in which the theory could be generalized is to investigate the general conditions under which the bisector of two sets is a $(d - 1)$-manifold that partitions E^d into two disjoint regions. In this paper we have shown that if the two sets are linearly separable, then their bisector has this property. Also, some sufficient conditions for the bisector to be a simple closed curve were given in [9] for sets in E^2 that are not even weakly linearly separable. We are not aware of any other result related to this problem.

Acknowledgments

We thank Michael O'Connor for helping us to formalize the proof of Lemma 1 and for suggesting the current proof of Theorem 5, which is much simpler than our original proof. We also thank him for encouraging us to generalize our results in E^2 to E^d. We thank Jarek Rossignac for suggesting an earlier proof of Lemmas 4 and 5, and the referee for suggesting several important improvements.

References

1. F. L. Bookstein, The Line-Skeleton, *Comput. Graphics and Image Process.*, vol. 11, pp. 123-137, October 1979.
2. L. Guibas and J. Stolfi, Primitives for the Manipulation of General Subdivisions and the Computation of Voronoi Diagrams, *ACM Trans. Graphics*, vol. 4, no. 2, pp. 74-123, April 1985.
3. T. Husain, *Topology and Maps*, New York: Plenum, 1977.
4. P. Kelly and M. Weiss, *Geometry and Convexity: A Study in Mathematical Methods*, New York: Wiley, 1979.
5. D. T. Lee, Medial Axis Transformation of a Planar Shape, *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 4, no. 4, pp. 363-369, July 1982.

6. D. T. Lee and R. L. Drysdale, III, Generalization of Voronoi Diagrams in the Plane, *SIAM J. Comput.*, vol. 10, no. 1, pp. 73-87, February 1981.
7. D. Leven and M. Sharir, Intersection and Proximity Problems and Voronoi Diagrams, in J. T. Schwartz and C. Yap, editors, *Algorithmic and Geometric Aspects of Robotics*, pp. 187-228, Hillsdale, New Jersey: Erlbaum, 1987.
8. F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*, New York: Springer-Verlag, 1985.
9. V. Srinivasan and L. R. Nackman, Voronoi Diagram for Multiply-Connected Polygonal Domains I: Algorithm, *IBM J. Res. Develop.*, vol. 31, no. 3, pp. 361-372, May 1987.

Received November 3, 1988, and in revised form March 22, 1989, and January 24, 1990.