Computing a Center-Transversal Line[*]

Pankaj K. Agarwal [1]   Sergio Cabello [2]   J. Antoni Sellarès [3]   Micha Sharir [4]

[1] Department of Computer Science, Duke University, P.O. Box 90129, Durham NC 27708-0129, USA, `pankaj@cs.duke.edu`

[2] Department of Mathematics, IMFM, and Department of Mathematics, FMF, University of Ljubljana, Slovenia, `sergio.cabello@fmf.uni-lj.si`

[3] Institut d'Informàtica i Aplicacions, Universitat de Girona, Spain, `sellares@ima.udg.es`

[4] School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA, `michas@post.tau.ac.il`

Abstract

A center-transversal line for two finite point sets in $\mathbb{R}^3$ is a line with the property that any closed halfspace that contains it also contains at least one third of each point set. It is known that a center-transversal line always exists [14, 29], but the best known algorithm for finding such a line takes roughly n^{12} time. We propose an algorithm that finds a center-transversal line in $O(n^{1+\varepsilon}\kappa^2(n))$ worst-case time, for any $\varepsilon > 0$, where $\kappa(n)$ is the maximum complexity of a single level in an arrangement of n planes in $\mathbb{R}^3$. With the current best upper bound $\kappa(n) = O(n^{5/2})$ of [26], the running time is $O(n^{6+\varepsilon})$, for any $\varepsilon > 0$. We also show that the problem of deciding whether there is a center-transversal line parallel to a given direction can be solved in $O(n \log n)$ expected time. Finally, we extend the concept of center-transversal line to that of bichromatic depth of lines in space, and give an algorithm that computes a deepest line exactly in time $O(n^{1+\varepsilon}\kappa^2(n))$, and a linear-time approximation algorithm that computes, for any specified $\delta > 0$, a line whose depth is at least $1 - \delta$ times the maximum depth.

1 Introduction

Two classical notions in discrete geometry are the notions of center points and ham-sandwich cuts. Given a set P of points in $\mathbb{R}^d$, a point q, not necessarily in P, is a *center point* with respect to P if any closed halfspace that contains q also contains at least $|P|/(d+1)$ points of P. The existence of center points is a consequence of Helly's theorem [19]. Given d finite point sets $P_0, \ldots, P_{d-1}$ in $\mathbb{R}^d$ with n points in total, a *ham-sandwich cut* is a hyperplane h such that each of the open halfspaces bounded by h contains at most $|P_i|/2$ points of P_i, for every $i = 0, 1, \ldots, d-1$. Dol'nikov [14], and Živaljević and Vrećica [29] proved the following theorem, called *center-transversal theorem*, which yields a generalization of center points and ham-sandwich cuts.

Theorem 1.1 (Center-Transversal Theorem) *Given $k+1$ finite point sets $P_0, P_1, \ldots, P_k$ in $\mathbb{R}^d$, for any $0 \leq k \leq d-1$, there exists a k-flat f such that any closed halfspace that contains f also contains at least $\frac{1}{d-k+1}|P_i|$ points of P_i, for each $i = 0, 1, \ldots, k$.*

Observe that when $k = 0$, f is a center point, and when $k = d - 1$, f is a ham-sandwich cut. Therefore, the center-transversal theorem can be seen as an "interpolation" between these two theorems. A weaker result with $|P_i|/(d+1)$ instead of $|P_i|/(d-k+1)$ can easily be obtained by considering the k-flat passing through a center point of each of the P_i, $i = 0, 1, \ldots, k$.

In this paper we consider in detail the case $d = 3$, $k = 1$. Given two finite point sets P_0, P_1 in $\mathbb{R}^3$, we say that a line ℓ is a *center-transversal line* for P_0, P_1 if any closed half-space that contains ℓ also contains at least $|P_i|/3$ points of P_i, for $i = 0, 1$. The center-transversal theorem asserts that, for any finite point sets P_0, P_1 in $\mathbb{R}^3$, there exists a center-transversal line. However, the original proofs [14, 29] of this result are non-constructive and do not lead to an algorithm for finding a center-transversal line. The running time of the best known algorithm for this problem [5] is rather large (about n^{12}—see below). We present a considerably more efficient algorithm for finding such a line, and consider several other related problems.

Related work. A more detailed review of center points, ham sandwich cuts, and related problems can be found in Matoušek [19]. Efficient algorithms are known for computing a center point in $\mathbb{R}^2$ and $\mathbb{R}^3$ [12, 17, 20]. A center point in $\mathbb{R}^d$ can be found using linear programming with $\Theta(n^d)$ linear inequalities, and there exists a faster algorithm, due to Clarkson et al. [11], for computing an *approximate* center point in arbitrary dimensions; that is, a point q such that any closed halfspace containing q contains at least $\Omega(n/d^2)$ points of P. Efficient algorithms have also been developed for constructing the *center region*, namely, the set of all center points, in $\mathbb{R}^2$ and $\mathbb{R}^3$ [4, 7, 18]. The concept of center point leads to generalizations that have been useful in robust statistics. The *halfspace depth* (also called location depth, data depth) of a point q relative to a data set P in $\mathbb{R}^d$, is the smallest number of data points in any closed halfspace whose boundary passes through q. A center point is a point with depth at least $|P|/(d+1)$, and a halfspace median, or a *Tukey point*, is a point with maximum halfspace depth. Chan [7], improving upon previous results, has obtained a randomized $O(n \log n + n^{d-1})$ expected-time algorithm for computing a Tukey point in $\mathbb{R}^d$.

The problem that we consider can be related to *multivariate regression depth*, a generalization, introduced by Bern and Eppstein [5], of *regression depth*, a quality measure for robust linear re-

gression defined by Rousseeuw and Hubert [16, 24, 25]. In particular, Bern and Eppstein [5] give a general-purpose algorithm, which can be easily modified to yield an algorithm that constructs a center-transversal line in $\mathbb{R}^3$ in $O(n^{12+\varepsilon})$ time, for any $\varepsilon > 0$.

Our contributions. Let P_0, P_1 be two finite point sets in $\mathbb{R}^3$ with a total of n points.

- We present an algorithm that constructs a center-transversal line for P_0 and P_1 in $O(n^{1+\varepsilon}\kappa^2(n))$ worst-case time, for any $\varepsilon > 0$, where $\kappa(n)$ is the maximum complexity of a single level in an arrangement of n planes in $\mathbb{R}^3$. With the current best upper bound $\kappa(n) = O(n^{5/2})$ of [26], the running time is $O(n^{6+\varepsilon})$, for any $\varepsilon > 0$. This is a considerable improvement over the algorithm by Bern and Eppstein [5].[1] This improvement is attained by analyzing the problem structure carefully, by conducting the search for candidate center-transversal lines in a controlled recursive manner, and by using (standard) range-searching data structures for interacting lines with polyhedral terrains. See Section 2.

- Using a simple relation between center-transversal lines and center points in two dimensions, we show how to decide in $O(n \log n)$ time, for a given direction, whether there exists a center-transversal line of P_0 and P_1 with that direction. See Section 3.

- We introduce the notion of the *bichromatic depth* of a line ℓ, with respect to P_0 and P_1, extending similar earlier concepts. Specifically, it is the minimum fraction size ρ of the points in either set that lie in a halfspace that contains ℓ; that is, each halfspace containing ℓ contains at least $\rho|P_0|$ points of P_0 and $\rho|P_1|$ points of P_1. This concept generalizes that of center-transversal line (which has bichromatic depth at least $1/3$). We show how to compute a deepest line in $O(n^{1+\varepsilon}\kappa^2(n))$ time, for any $\varepsilon > 0$, and give a linear-time approximation algorithm that computes, for any $\delta > 0$, a line whose depth is at least $1-\delta$ times the maximum depth. See Section 4.

2 Finding a Center-Transversal Line

We consider the problem of computing a center-transversal line in dual space, where the problem is reformulated in terms of levels in arrangements of planes. We generate a set of candidate lines that is guaranteed to contain a center-transversal line and use a data structure to determine which of these candidate lines is a center transversal line. For simplicity, we assume that $P_0 \cup P_1$ are in general position in the sense that no four of them are coplanar.

Center-transversal lines in the dual. The widely used *duality* transform maps a point p in $\mathbb{R}^d$ to a hyperplane p^* in $\mathbb{R}^d$ and vice-versa, so that the incidence and above/below relationships are preserved. There are many variants of duality [19]; we use the following one: A point $a = (a_1, \ldots, a_d) \in \mathbb{R}^d$ is mapped to the nonvertical hyperplane $a^* : x_d = a_1 x_1 + \cdots + a_{d-1} x_{d-1} - a_d$, and a hyperplane $h : x_d = \alpha_1 x_1 + \cdots + \alpha_{d-1} x_{d-1} + \alpha_d$ is mapped to the point $h^* =$

---

[1]We note though that an algorithm with running time near n^8 is not hard to obtain.

2

$(\alpha_1, \ldots, \alpha_{d-1}, -\alpha_d)$, so $(a^*)^* = a$. A point p lies below (resp., above, on) a hyperplane h if the dual point h^* lies below (resp., above, on) the dual hyperplane p^*. The *pencil* of hyperplanes passing through a line ℓ in $\mathbb{R}^d$, for $d \geq 3$, maps to the set of points in $\mathbb{R}^d$ lying on a line ℓ^*; we refer to ℓ^* as the dual of ℓ. For a set A of objects, set $A^* = \{a^* \mid a \in A\}$.

Let P be a set of n points in $\mathbb{R}^3$, and let $H = P^*$ be the set of n non-vertical planes in $\mathbb{R}^3$ dual to the points in P. The *level* of a point $p \in \mathbb{R}^3$, with respect to H, is the number of planes in H that lie *below* p. For $0 \leq k < n$, the *k-level* of H, denoted $\mathcal{L}_k(H)$ (or simply $\mathcal{L}_k$ if the set H is understood), is the closure of the set of all points on any of the planes of H that are at level k. The k-level $\mathcal{L}_k$ is a *polyhedral terrain*, that is, an xy-monotone piecewise-linear continuous surface formed by a subset of the faces of the arrangement $\mathcal{A}(H)$. The combinatorial complexity of $\mathcal{L}_k$ is the number of faces of all dimensions in $\mathcal{L}_k$. Let $\kappa(n)$ denote the maximum complexity of a level in any arrangement of n planes in $\mathbb{R}^3$. The best known upper bound for $\kappa(n)$ is $O(n^{5/2})$ [26], which differs substantially from the best known lower bound $n^2 e^{\Omega(\sqrt{\log n})}$ [28]. See [3] for more details on arrangements and levels.

If h is a plane in $\mathbb{R}^3$ so that each of the two halfspaces bounded by h contains at least k points of P, then h^* lies between $\mathcal{L}_k(H)$ and $\mathcal{L}_{n-k}(H)$. If ℓ is a line in $\mathbb{R}^3$ so that any halfspace containing ℓ contains at least k points of P, then the entire dual line ℓ^* lies between $\mathcal{L}_k(H)$ and $\mathcal{L}_{n-k}(H)$. Hence, the problem of computing a center-transversal line for P_0 and P_1 reduces to computing a line in the dual space that lies above $\Sigma_0 = \mathcal{L}_{k_0}(H_0), \Sigma_1 = \mathcal{L}_{k_1}(H_1)$ and below $\Sigma_2 = \mathcal{L}_{n_0-k_0}(H_0), \Sigma_3 = \mathcal{L}_{n_1-k_1}(H_1)$, where $H_i = P_i^*$, $n_i = |P_i|$, and $k_i = \lceil n_i/3 \rceil$ for $i = 0, 1$. We note that each of these four terrains can be computed in $O(n^\varepsilon \kappa(n))$ time, for any $\varepsilon > 0$ [2].

We thus have four terrains $\Sigma_0, \Sigma_1, \Sigma_2, \Sigma_3$, and we wish to compute a line that lies above Σ_0, Σ_1 and below Σ_2, Σ_3. Note that such a line cannot be *z-vertical*, i.e., parallel to the z-axis. Let E_i be the set of edges in Σ_i, for $i = 0, 1, 2, 3$, and $E = \bigcup_{i=0}^3 E_i$. Set $m := |E| \leq 4\kappa(n)$, and assume that $m \geq n$ (or else the problem can be solved much faster than the time bound of our algorithm). Let $H = H_0 \cup H_1$. Each edge in E_i lies in the intersection line of a pair of planes in H. We define a "sidedness function" $\chi : E \to \{+1, -1\}$, where $\chi(e) = +1$ if $e \in E_0 \cup E_1$ and $\chi(e) = -1$ if $e \in E_2 \cup E_3$. Let V be the set of endpoints of edges in E. By the general-position assumption, each point of V is incident upon at most three edges of E. For an object (point, line, segment) Δ in $\mathbb{R}^3$, let $\tilde{\Delta}$ denote its xy-projection in $\mathbb{R}^2$.

Definition 2.1 Let ℓ be a nonvertical line in $\mathbb{R}^3$, and let e be a nonvertical segment in $\mathbb{R}^3$ so that $\tilde{\ell}$ intersects $\tilde{e}$. We say that ℓ lies *above* (resp., *below*) e if the oriented line in the $(+z)$-direction that passes through $\tilde{\ell} \cap \tilde{e}$ meets e before (resp., after) ℓ. The line ℓ is in *compliance* with an edge $e \in E$ if (i) $\tilde{\ell}$ does not intersect $\tilde{e}$, or (ii) ℓ does not lie below (resp., above) e if $\chi(e) = +1$ (resp., $\chi(e) = -1$). We say that ℓ is in compliance with a subset $R \subseteq E$ if it is in compliance with every edge in R. In particular, we have:

Lemma 2.2 *A nonvertical line ℓ in $\mathbb{R}^3$ lies above Σ_0, Σ_1 and below Σ_2, Σ_3 if and only if ℓ is in compliance with E.*

The problem of computing a center-transversal line now reduces to finding a line that is in compliance with E. Let $\mathbb{L}$ be the set of all lines in $\mathbb{R}^3$ that are not parallel to the yz-plane. We

restrict the search for a line that is in compliance with E to lines in $\mathbb{L}$. This involves no loss of generality: The lines in $\mathbb{R}^3$ parallel to the yz-plane have three degrees of freedom and a center-transversal line among them, if there exists one, can be found using a much simpler (and more efficient) algorithm. Alternatively, we can run our algorithm twice, exchanging the roles of the x- and y-axes in the second run.

Overview of the algorithm. We show that, for each line $\ell \in \mathbb{L}$, there exists a "witness set" of $O(n)$ edges of E, so that ℓ is in compliance with E if and only if it is in compliance with its witness set. We then group the lines in $\mathbb{L}$ into equivalence classes so that all lines in the same class have the same witness set. Using this reduction, we present an algorithm that works in three stages. The first stage, called the *filtering stage*, splits the problem into $O(m^2/n^2)$ subproblems, each aiming to compute a line that is in compliance with some set of $O(n)$ edges. The second stage, a *recursive candidate generation stage*, computes, for each subproblem, a set of $O(n^{3+\varepsilon})$ candidate lines, for any $\varepsilon > 0$, which is guaranteed to contain a line in compliance with the corresponding subset if there exists one. The final stage, the *verification stage*, checks which of the candidate lines generated by the previous step is in compliance with E, and report the first such line that it encounters (which is guaranteed to exist). We now describe each of these steps in detail.

Witness sets and equivalence classes. For a line $\ell \in \mathbb{L}$ and a subset $R \subseteq E$ of edges, we define the *witness set* of ℓ for R, denoted by $W(\ell, R)$, as follows. For $i = 0, 1, 2, 3$, let $R_i \subseteq R$ be the sequence of edges in $R \cap E_i$ whose xy-projections intersect $\tilde{\ell}$, sorted by the order of the intersection points along $\tilde{\ell}$. For a plane $h \in H_0 \cup H_1$, let $e^-_{h,i}, e^+_{h,i} \in R_i$ be, respectively, the first and the last edges in the i-th sequence that lie on h, where only planes in H_0 (resp., H_1) are considered for $i = 0, 2$ (resp., $i = 1, 3$). We set

$$W(\ell, R) = \{e^-_{h,i}, e^+_{h,i} \mid h \in H, \ 0 \leq i \leq 3\}.$$

By definition, $\tilde{\ell}$ intersects the xy-projection of every edge in $W(\ell, R)$. Note that $|W(\ell, R)| = O(n)$.

Lemma 2.3 *For a subset $R \subseteq E$, a line $\ell \in \mathbb{L}$ is in compliance with R if and only if ℓ is in compliance with $W(\ell, R)$.*

The proof of the lemma follows from the simple observation that if ℓ lies above (resp., below) both $e^-_{h,i}, e^+_{h,i}$ then it lies above (resp., below) all edges in R_i that lie in h.

We define, for a subset $R \subseteq E$, an equivalence relation on $\mathbb{L}$ so that for any two lines ℓ_1, ℓ_2 in the same equivalence class, $W(\ell_1, R) = W(\ell_2, R)$. This will discretize the search for a center-transversal line. For this we need a few notations. For a point or a line ξ in $\mathbb{R}^3$, let $\varphi(\xi)$ denote the dual (in $\mathbb{R}^2$) of $\tilde{\xi}$, i.e., $\varphi(\xi) = \tilde{\xi}^*$. [2] For an edge $e = uv$ in E, let $\varphi(e) \subseteq \mathbb{R}^2$ be the double wedge that is formed by the lines $\varphi(u)$ and $\varphi(v)$ and does not contain the line in $\mathbb{R}^2$ passing through their intersection point and parallel to the y-axis. By standard properties of the duality transform in $\mathbb{R}^2$, a line γ in $\mathbb{R}^2$ intersects $\tilde{e}$ if and only if $\gamma^* \in \varphi(e)$. Moreover if the points $\gamma_1^*, \gamma_2^* \in \mathbb{R}^2$ lie in the same (left or right) wedge of $\varphi(e)$, then γ_1, γ_2 intersect $\tilde{e}$ from the *same side*, in the sense that the

---

[2] Note that $\varphi(\ell)$ is not defined if ℓ is parallel to the yz-plane. That is why we exclude these lines from $\mathbb{L}$.

same endpoint of $\tilde{e}$ lies in each of the positive halfplanes bounded by γ_1 and γ_2, respectively (that is, the halfplanes above these lines).

Let $R \subseteq E$ be a fixed subset of edges, let $V_R \subseteq V$ be the set of endpoints of the edges in R, and let $\Lambda(R) = \{\varphi(v) \mid v \in V_R\}$ be the corresponding set of lines in $\mathbb{R}^2$. For each face f in the arrangement $\mathcal{A}(\Lambda(R))$ of $\Lambda(R)$, let $R(f)$ denote the set of those edges $e \in R$ for which $\varphi(e)$ contains f. For a line $\ell \in \mathbb{L}$, if f is the face containing $\varphi(\ell)$ then, by construction, $R(f)$ is the set of edges of R whose xy-projections intersect $\tilde{\ell}$. By definition, $W(\ell, R) \subseteq R(f)$.

Definition 2.4 We call two lines $\ell_1, \ell_2 \in \mathbb{L}$ *equivalent* (with respect to R), denoted by $\ell_1 \equiv_R \ell_2$, if $\varphi(\ell_1)$ and $\varphi(\ell_2)$ lie in the same face of $\mathcal{A}(\Lambda(R))$.

Lemma 2.5 *Let $R \subseteq E$ be a set of edges, and let $\ell_1, \ell_2 \in \mathbb{L}$ be two lines so that $\ell_1 \equiv_R \ell_2$. Then $W(\ell_1, R) = W(\ell_2, R)$.*

Proof: Let f be the face of $\mathcal{A}(\Lambda(R))$ that contains $\varphi(\ell_1)$ and $\varphi(\ell_2)$. Set $R_i(f) := R(f) \cap E_i$ and $L_i := \Lambda(R_i(f)) \subseteq \Lambda(R)$, for $i = 0, 1, 2, 3$. Clearly, $\varphi(\ell_1), \varphi(\ell_2)$ lie in the same face of $\mathcal{A}(L_i)$. Since the edges of E_i all belong to the same terrain, their xy-projections are pairwise disjoint. An easy observation (due to [1]) shows that $\tilde{\ell}_1, \tilde{\ell}_2$ intersect the xy-projections of the edges in $R_i(f)$ in the same order. This immediately implies that $W(\ell_1, R) \cap E_i = W(\ell_2, R) \cap E_i$, from which the lemma follows. $\square$

In view of the preceding lemma, we define, for each face f of $\mathcal{A}(R)$, $W_f(R) \subseteq R$ to be the common witness set for any line in the equivalence class corresponding to f.

The filtering stage. Given a set L of lines in $\mathbb{R}^2$, a triangle Δ_0, and a parameter $1 \leq r \leq |L|$, a $(1/r)$-*cutting* of (L, Δ_0) is a triangulation Ξ of Δ_0 so that each triangle of Ξ is crossed by at most $|L|/r$ lines of L. It is known that a $(1/r)$-cutting consisting of $O(r^2)$ triangles, along with the set of lines crossing each of its triangles, can be computed in $O(|L|r)$ time [8].

Let $\Lambda = \Lambda(E)$. We set $\Delta_0 = \mathbb{R}^2$ and $r = m/n$, and compute a $(1/r)$-cutting Ξ of (Λ, Δ_0). For each triangle $\Delta \in \Xi$, let Λ_Δ be the set of lines of Λ that cross Δ; since Ξ is a $(1/r)$-cutting, we have $|\Lambda_\Delta| \leq m/r = n$. Let $E_\Delta \subseteq E$ be the set of edges $e = uv$ so that either $\varphi(u) \in \Lambda_\Delta$ or $\varphi(v) \in \Lambda_\Delta$. Since each vertex of V is an endpoint of at most three edges of E, we have $|E_\Delta| \leq 3|\Lambda_\Delta| \leq 3n$. For each $\Delta \in \Xi$, let $F_\Delta = \{e \in E \setminus E_\Delta \mid \Delta \subseteq \varphi(e)\}$. We refer to the edges in E_Δ as *short* and to the edges in F_Δ as *long*. Finally, let $\mathbb{L}_\Delta = \{\ell \in \mathbb{L} \mid \varphi(\ell) \in \Delta\}$.

Since Δ is contained in a face of $\mathcal{A}(\Lambda(F_\Delta))$ (the arrangement of lines dual to the xy-projections of the endpoints of the edges in F_Δ), Lemma 2.5 implies that $W(\ell, F_\Delta)$ is the same for all lines $\ell \in \mathbb{L}_\Delta$; let W_Δ denote this common witness set. Observe that $|W_\Delta| = O(n)$.

If two triangles Δ and Δ' in Ξ share an edge, then $F_\Delta \oplus F_{\Delta'} \subseteq E_\Delta \cup E_{\Delta'}$. Therefore W_Δ can be computed from $W_{\Delta'}$ in $O(|E_\Delta| + |E_{\Delta'}|) = O(n)$ time. Hence, by performing a traversal of Ξ, we can compute W_Δ for all triangles $\Delta \in \Xi$, in overall time $O(m^2/n)$.

The next lemma follows from Lemmas 2.3 and 2.5.

Lemma 2.6 *For any $\Delta \in \Xi$, a line $\ell \in \mathbb{L}_\Delta$ is in compliance with E if and only if ℓ is in compliance with $E_\Delta \cup W_\Delta$.*

Hence, for each $\Delta \in \Xi$, we have a subproblem $(\Delta, E_\Delta, W_\Delta)$, in which we want to determine whether there is a line in $\mathbb{L}_\Delta$ that is in compliance with $E_\Delta \cup W_\Delta$ (and thus with E). Since $\bigcup_\Delta \mathbb{L}_\Delta = \mathbb{L}$, these subproblems together exhaust the overall problem of computing a line in $\mathbb{L}$ that is in compliance with E. There are $O(m^2/n^2)$ such subproblems, and the total time spent in generating them is $O(m^2/n)$.

The recursive candidate generation stage. Let $(\Delta, E_\Delta, W_\Delta)$ be one of the subproblems generated in the previous stage. We generate a set of "candidate" lines that contains a line in compliance with $E_\Delta \cup W_\Delta$ if there exists one. Let $\ell \in \mathbb{L}_\Delta$ be such a line. We move it around while keeping it in the set $\mathbb{L}_\Delta$ and in compliance with $E_\Delta \cup W_\Delta$, until we reach a critical position of ℓ at which one of the following events occurs (for the following enumeration, recall that passing above, below, or through an endpoint of an edge in W_Δ can occur only when $\varphi(\ell)$ reaches the boundary of Δ):

(E0) $\varphi(\ell)$ is a vertex of Δ;

(E1) ℓ passes through a pair of endpoints of edges in E_Δ;

(E2) ℓ passes through an endpoint of an edge in E_Δ, $\varphi(\ell)$ lies on an edge of Δ, and ℓ touches the relative interior of an edge of $E_\Delta \cup W_\Delta$;

(E3) ℓ passes through an endpoint of an edge in E_Δ and touches the relative interior of two edges of $E_\Delta \cup W_\Delta$;

(E4) $\varphi(\ell)$ lies on an edge of Δ, and ℓ touches the relative interior of three edges of $E_\Delta \cup W_\Delta$;

(E5) ℓ touches the relative interior of four edges of $E_\Delta \cup W_\Delta$.

Since (E0)–(E4) are defined by at most three edges of $E_\Delta \cup W_\Delta$ and there are $O(1)$ lines for each such event (assuming general position), we generate all critical lines of these types (the $O(n^3)$ cost of producing these lines is subsumed by the cost of generating the lines of type (E5)—see below). We add all the resulting lines that belong to $\mathbb{L}_\Delta$ to the candidate set. Hence, it suffices to describe an algorithm for computing the set of candidate lines that satisfy (E5). Let $\mathcal{C}(\Delta, E_\Delta, W_\Delta)$ denote this set. We compute a superset of $\mathcal{C}(\Delta, E_\Delta, W_\Delta)$ with a divide-and-conquer algorithm that employs *Plücker coordinates* [22]. Our approach for generating candidate lines is very similar to that used by Pellegrini [21] (see also [23]).

Before describing the algorithm, we briefly review the representation of lines in Plücker space. An oriented line ℓ in $\mathbb{R}^3$ can be mapped to a point $\pi(\ell) \in \mathbb{R}^5$, called the *Plücker point* of ℓ, that lies on the so-called 4-dimensional *Plücker hypersurface* Π, or to a hyperplane $\varpi(\ell)$ in $\mathbb{R}^5$, called the *Plücker hyperplane* of ℓ. (The actual Plücker space is the *real projective* 5-space, but since we exclude lines parallel to the yz-plane, it is easy (though some care is needed) to embed the Plücker structure into the real 5-dimensional space.) Abusing the notation a little, we use $\pi(e)$ and $\varpi(e)$ to denote the Plücker point and hyperplane, respectively, of the line supporting an oriented segment e in $\mathbb{R}^3$.

We orient every line of $\mathbb{L}$ and every edge of E in the $(+x)$-direction (this is well defined for lines in $\mathbb{L}$, by definition, and for edges of E, by the general position assumption). For two oriented

lines ℓ_1, ℓ_2 in $\mathbb{R}^3$, $\pi(\ell_1)$ lies above $\varpi(\ell_2)$ (which is the same as $\pi(\ell_2)$ lying above $\varpi(\ell_1)$) if and only if the simplex spanned by a vector $\vec{u}_1$ lying on ℓ_1 with the same orientation, and by a vector $\vec{u}_2$ lying on ℓ_2 with the same orientation, is positively oriented. This is easily seen to imply that, when ℓ_1 and ℓ_2 are non-vertical, ℓ_1 passes above ℓ_2 if and only if either (i) $\pi(\ell_1)$ lies above $\varpi(\ell_2)$ and $\tilde{\ell}_1$ lies counterclockwise to $\tilde{\ell}_2$, or (ii) $\pi(\ell_1)$ lies below $\varpi(\ell_2)$ and $\tilde{\ell}_1$ lies clockwise to $\tilde{\ell}_2$; see [22] for more details.

We now proceed to describe the construction of the set of lines $\mathcal{C}(\Delta, E_\Delta, W_\Delta)$. We choose a constant r and construct a $(1/6r)$-cutting T of $(\Lambda(E_\Delta), \Delta)$. As in the filtering stage, we define, for each $\tau \in T$, $E_\tau \subseteq E_\Delta$ to be the set of short edges in τ, and $F_\tau \subseteq E_\Delta$ to be the set of long edges in τ. We have $|E_\tau| \leq 3|\Lambda(E_\Delta)|/6r \leq |E_\Delta|/r$. Set $W_\tau := F_\tau \cup W_\Delta$. Define $\mathbb{L}_\tau = \{\ell \in \mathbb{L}_\Delta \mid \varphi(\ell) \in \tau\}$, and note that $\bigcup_{\tau \in T} \mathbb{L}_\tau = \mathbb{L}_\Delta$. For each $\tau \in T$, we compute a set of *candidate* lines $\mathcal{C}_\tau \subset \mathbb{L}_\tau$, with the property that $\mathcal{C}(\Delta, E_\Delta, W_\Delta) \subseteq \bigcup_{\tau \in T} \mathcal{C}_\tau$.

Consider a triangle $\tau \in T$. We want to construct a set of candidate lines $\mathcal{C}_\tau$ that includes the lines in $\mathbb{L}_\tau$ of type (E5). Hence, it suffices to consider only the edges $E_\tau \cup W_\tau$ in its construction. The line ℓ is in compliance with an edge $e \in W_\tau$ if $\pi(\ell)$ lies in one specific halfspace Γ_e bounded by $\varpi(e)$. Γ_e depends on the function $\chi(e)$ and on the clockwise order of $\tilde{\ell}$ and $\tilde{e}$ (when oriented in the positive x-direction). Since τ is a subset of a fixed wedge of $\varphi(e)$, this clockwise order is the same for all lines $\ell \in \mathbb{L}_\tau$; hence Γ_e is the same halfspace for all lines in $\mathbb{L}_\tau$. Set $\mathcal{K} := \bigcap_{e \in W_\tau} \Gamma_e$; $\mathcal{K}$ is a convex polyhedron in $\mathbb{R}^5$ with $O(n)$ facets, so its overall combinatorial complexity is $O(n^2)$.

Let $\ell \in \mathbb{L}_\tau$ be a line that touches the relative interior of four edges of $E_\tau \cup W_\tau$, and let $B(\ell)$ denote the set of these four edges. There are four cases, depending on how many edges of W_τ the line ℓ touches.

$B(\ell) \subseteq W_\tau$. If all edges of $B(\ell)$ belong to W_τ and ℓ is in compliance with $E_\tau \cup W_\tau$, then $\pi(\ell) \in \mathcal{K}$. Since ℓ touches four edges of W_τ, it lies on an edge of $\mathcal{K}$. Therefore, we find lines of this type by intersecting each edge of $\mathcal{K}$ with the (quadratic) Plücker hypersurface Π, and by adding the (at most) two lines corresponding to the two intersection points to the candidate set $\mathcal{C}_\tau$, if they belong to $\mathbb{L}_\tau$. The total time spent is $O(n^2)$.

$|B(\ell) \cap W_\tau| = 3$. For any line ℓ with this kind of contacts, $\pi(\ell)$ lies on the intersection edge of some 2-face of $\mathcal{K}$ and the Plücker hyperplane $\varpi(e)$ for some $e \in E_\tau$. For each pair $e \in E_\tau$ and 2-face ϕ of $\mathcal{K}$, we compute the at most two intersection points of $\phi \cap \varpi(e) \cap \Pi$, and add the corresponding lines to the candidate set $\mathcal{C}_\tau$, if they belong to $\mathbb{L}_\tau$. Since the polyhedron $\mathcal{K}$ has $O(n^2)$ 2-faces, the total number of lines generated in this case is $O(n^2|E_\tau|) = O(n^3/r)$, and their construction takes $O(n^3/r)$ time.

$|B(\ell) \cap W_\tau| = 2$. Let $e_1, e_2 \in E_\tau$ be the two edges that belong to $B(\ell)$. The Plücker subspace F of lines (in $\mathbb{L}$) that touch e_1 and e_2 is a 3-dimensional flat in $\mathbb{R}^5$, and $\pi(\ell) \in F \cap \mathcal{K}$. Since $F \cap \mathcal{K}$ is a convex 3-polyhedron with $O(n)$ facets, it only has $O(n)$ edges. We form, as above, the intersections of each edge of $F \cap \mathcal{K}$ with the Plücker surface Π, and add the (at most two) resulting lines to our candidate set $\mathcal{C}_\tau$, if they belong to $\mathbb{L}_\tau$. The total number of lines generated in this case is $O(|E_\tau|^2 n) = O(n^3/r^2)$, and their computation takes $O(|E_\tau|^2 n \log n) = O((n^3/r^2) \log n)$ time, where the costliest step is the construction, repeated $O(|E_\tau|^2)$ times, of convex 3-polyhedra, each defined by at most n inequalities.

$|B(\ell) \cap W_\tau| \leq 1$. We partition W_τ (arbitrarily) into $u = O(r)$ subsets $W_\tau^{(1)}, \ldots, W_\tau^{(u)}$ so that $|W_\tau^{(i)}| \leq n/r$ for each i. We recursively compute the set of candidate lines $\mathcal{C}(\tau, E_\tau, W_\tau^{(i)})$, for $1 \leq i \leq u$ and for $\tau \in T$. We thus recursively solve $O(r)$ subproblems, all of whose outputs are added to our candidate set $\mathcal{C}_\tau$. Clearly, all lines of this type (and perhaps more) are found by this recursive procedure.

The correctness of the procedure is fairly straightforward. Let $T(n)$ denote the maximum time needed to compute $\bigcup_{\tau \in T} \mathcal{C}_\tau$, which is a superset of $\mathcal{C}(\Delta, E_\Delta, W_\Delta)$, when $|E_\Delta|, |W_\Delta| \leq n$. For each $\tau \in T$, we spend $O(n^2 + n^3/r + (n^3/r^2)\log n)$ time plus the time needed to solve $O(r)$ recursive calls where the size of each of the two sets of edges is at most n/r. Since the cutting T consists of $O(r^2)$ triangles, we obtain the following recurrence.

$$T(n) = O(r^3)T(n/r) + O(n^2 r^2 + n^3 r + n^3 \log n).$$

The solution of this recurrence is $T(n) = O(n^{3+\varepsilon})$, for any $\varepsilon > 0$ (for which we need to choose r sufficiently large, as a function of ε). The size of $\mathcal{C}(\Delta, E_\Delta, W_\Delta)$ is also bounded by this quantity.

Repeating this procedure for the $O(m^2/n^2)$ subproblems generated by the filtering stage, we construct, in $O(m^2 n^{1+\varepsilon})$ overall time, a candidate set $\mathcal{C}$ of $O(m^2 n^{1+\varepsilon})$ lines.

The verification stage. To complete the algorithm, we test which of the lines in $\mathcal{C}$ is in compliance with E. Using the data structure described in [10], we can preprocess, in $O(m^{2+\varepsilon})$ time, each E_i into a data structure of size $O(m^{2+\varepsilon})$ so that we can determine in $O(\log n)$ time whether a line $\ell \in \mathbb{L}$ passes above or below the terrain Σ_i, or, equivalently, whether ℓ is in compliance with E_i. Querying each line in $\mathcal{C}$ with this data structure for every E_i, we can determine, in $O(m^{2+\varepsilon} + m^2 n^{1+\varepsilon}) = O(m^2 n^{1+\varepsilon})$ time, which of the lines in $\mathcal{C}$ are in compliance with E. Since a center-transversal line always exists, it belongs to $\mathcal{C}$, by construction, and will be found by this procedure. Putting everything together, and recalling that $m \leq 4\kappa(n)$, where $\kappa(n)$ is the maximum complexity of a level in an arrangement of n planes in $\mathbb{R}^3$, we obtain the following main result of the paper. For the concrete time bound, we use the currently best known upper bound $\kappa(n) = O(n^{5/2})$ of [26].

Theorem 2.7 *A center-transversal line for two sets P_0, P_1 with a total of n points in $\mathbb{R}^3$ can be constructed in $O(n^{1+\varepsilon}\kappa^2(n))$ time, for any $\varepsilon > 0$. This time bound is $O(n^{6+\varepsilon})$, for any $\varepsilon > 0$.*

Remarks. (1) It is strongly believed that $\kappa(n) = O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, in which case our algorithm takes $O(n^{5+\varepsilon})$, for any $\varepsilon > 0$.
(2) The only place where we use the general-position assumption on P_0 and P_1 is in bounding the size of E_Δ (or E_τ). If we define the weight of a line $\varphi(v) \in \Lambda(E)$ to be the number of edges of E incident upon v and compute *weighted* $(1/r)$-cuttings [9], the same bound on the size of E_Δ (or E_τ) can be obtained.

Terrains with many coplanar faces. Pellegrini [21] and Halperin and Sharir [15] have shown that the complexity of the envelope of lines above a terrain of complexity k is $O(k^{3+\varepsilon})$, for any $\varepsilon > 0$. The complexity of this envelope corresponds to the number of lines that are tangent to the

terrain while lying above it. In our scenario, we have taken advantage of the fact that the faces of our terrains are contained in few planes. It is not clear how to plug this hypothesis into the techniques used in [15, 21]. However, using the ideas of witness sets and the filtering stage as we have done, we directly obtain the following result, which may be of independent interest.

Theorem 2.8 *Let Σ be a terrain of complexity k in $\mathbb{R}^3$, all of whose facets lie on n different planes. Then the complexity of the envelope of lines that pass above Σ is $O(n^{1+\varepsilon}k^2)$, for any $\varepsilon > 0$.*

3   Center-Transversal Line in a Given Direction

In this section we present a randomized algorithm for deciding whether there exists a center-transversal line of P_0 and P_1 in a given direction, say, the z-direction. Let $\tilde{P}_0$ (resp., $\tilde{P}_1$) be the xy-projection of P_0 (resp., P_1). A center-transversal line of P_0 and P_1 exists in the z-direction if and only if the intersection of the center regions of $\tilde{P}_0$ and $\tilde{P}_1$ is nonempty. Since each of these center regions can be computed in $O(n \log^2 n)$ randomized expected time [7] and their intersection can be computed in linear time, we can compute a center-transversal line in the z-direction, if it exists, in $O(n \log^2 n)$ expected time. Here we improve the expected running time to $O(n \log n)$.

For $i = 0, 1$, let H_i be the set of lines dual to $\tilde{P}_i$, and, for an integer k, let $\mathcal{L}_k(H_i)$ (resp., $\mathcal{U}_k(H_i)$) be the set of points whose level in $\mathcal{A}(H_i)$ is at most k (resp., at least $|H_i| - k$). In the dual setting, the problem of computing a center-transversal line in the z-direction reduces to determining whether there exists a line in the dual plane that lies above $\mathcal{L}_{k_0}(H_0) \cup \mathcal{L}_{k_1}(H_1)$ and below $\mathcal{U}_{k_0}(H_0) \cup \mathcal{U}_{k_1}(H_1)$, where $k_i = \lceil |P_i|/3 \rceil$ for $i = 0, 1$.

Let $\mathbb{L}$ be the set of all lines in $\mathbb{R}^2$. Suppose we have a (possibly infinite) set S of points in $\mathbb{R}^2$, in which each point is colored red or blue. We wish to compute

$$
\begin{aligned}
\omega(S) \quad &:= \quad \min_{\ell \in \mathbb{L}} \operatorname{slope}(\ell) \\
&\text{s.t. } \ell \text{ lies above the red points of } S \\
&\text{and } \ell \text{ lies below the blue points of } S \qquad (1)
\end{aligned}
$$

As argued by Chan [7], this is an instance of linear programming. By setting $S = \mathcal{L}_k(H_0) \cup \mathcal{L}_k(H_1) \cup \mathcal{U}_k(H_0) \cup \mathcal{U}_k(H_1)$, where the points of $\mathcal{L}_k(H_0) \cup \mathcal{L}_k(H_1)$ are colored red and the points of $\mathcal{U}_k(H_0) \cup \mathcal{U}_k(H_1)$ are colored blue, we can reduce our problem to an instance of (1). Although the set S is infinite in our case, it suffices to consider the vertices of $\mathcal{L}_k(H_i)$ and $\mathcal{U}_k(H_i)$, for $i = 0, 1$. However we cannot afford to compute the vertices of the levels explicitly if we are aiming for an $O(n \log n)$-time algorithm, as the best known upper bound on the complexity of a level in $\mathcal{A}(H_i)$ is $O(n^{4/3})$ [13], and a lower bound of $n \cdot 2^{\Omega(\sqrt{\log n})}$ exists [28]. We use Chan's randomized technique for solving LP-type problems in which the constraints are defined implicitly by a set of input objects, and which satisfy certain properties (see Lemma 3.1 below).

Given a set $\mathbb{H}$ of constraints and a totally ordered set W, a weight function $\omega : 2^{\mathbb{H}} \to W$ is called *LP-type* of dimension at most d if the following three conditions are satisfied for every subset $H \subseteq \mathbb{H}$ and each constraint $h \in \mathbb{H}$:

- There exists a subset B of size at most d, called a *basis* of H, so that $\omega(H) = \omega(B)$.

- $\omega(H \cup \{h\}) \geq \omega(H)$.

- Let $F \subseteq H$ such that $\omega(F) = \omega(H)$. Then $\omega(H \cup \{h\}) > \omega(H) \Leftrightarrow \omega(F \cup \{h\}) > \omega(F)$.

Since linear programming, with ω being the corresponding linear objective function, is an LP-type problem of dimension $d + 1$, (1) is an LP-type problem. See [27] for more details. The following lemma is the main result behind Chan's technique.

Lemma 3.1 (Chan [7]) *Let $\omega : 2^{\mathbb{H}} \to W$ be an LP-type function of constant dimension d, and let $\alpha < 1$ and s be fixed constants. Suppose $f : \mathbb{P} \to 2^{\mathbb{H}}$ is a function that maps inputs from some set $\mathbb{P}$ to sets of constraints with the following properties:*

(C1) *For inputs $P_1, \ldots, P_d \in \mathbb{P}$ of constant size, a basis for $f(P_1) \cup \cdots \cup f(P_d)$ can be computed in constant time.*

(C2) *For any input $P \in \mathbb{P}$ and any basis $B \subseteq f(P)$, we can decide in $O(D(n))$ time whether B satisfies $f(P)$, i.e., $\omega(f(P)) = \omega(B)$.*

(C3) *For any input $P \in \mathbb{P}$, we can construct, in $O(D(n))$ time, inputs $P_1, \ldots, P_s \in \mathbb{P}$ each of size at most $\lceil \alpha n \rceil$, so that $f(P) = f(P_1) \cup \cdots \cup f(P_s)$.*

Then we can compute a basis for $f(P)$ in $O(D(n))$ expected time, assuming that $D(n)/n^\varepsilon$ is monotonically increasing.

This lemma is a multidimensional version of an earlier technique that Chan proposed in [6]; he used this technique to compute the Tukey depth of a point set. A very slight (straightforward) variant of this algorithm can be used to solve our problem. For the sake of completeness, we sketch the algorithm here.

We formulate the problem in a slightly more general framework. Given $G_i \subseteq H_i$, for $i = 0, 1$, a triangle τ, and integers a_0, a_1, b_0, b_1, let

$$f(G_0, G_1, \tau, a_0, a_1, b_0, b_1) = \tau \cap (\mathcal{L}_{a_0}(G_0) \cup \mathcal{L}_{a_1}(G_1) \cup \mathcal{U}_{b_0}(G_0) \cup \mathcal{U}_{b_1}(G_1)). \qquad (2)$$

The points of $L = (\mathcal{L}_{a_0}(G_0) \cup \mathcal{L}_{a_1}(G_1)) \cap \tau$ are colored red, and the points of $U = (\mathcal{U}_{b_0}(G_0) \cup \mathcal{U}_{b_1}(G_1)) \cap \tau$ are colored blue. We wish to compute $\omega(f(G_0, G_1, \tau, a_0, a_1, b_0, b_1))$, as defined in (1).

We show that (2) satisfies (C1)–(C3). Condition (C1) is trivial because we can solve the problem explicitly in $O(1)$ time for constant-size inputs, by constructing the full arrangement of the input lines.

As for (C2), let ℓ be the line defined by a basis B. We need to determine whether ℓ lies above L and below U. We describe how to determine whether ℓ lies above L. Let τ^+ be the portion of τ lying above ℓ, then ℓ lies above L if and only if $\tau^+ \cap L = \emptyset$; the latter holds if and only if none of the edges of τ^+ intersects L. Let e be an edge of τ^+. We compute the intersection points of e with the lines in $G_0 \cup G_1$ and sort them along e. By computing the level of an endpoint of e with respect to G_0 and G_1 and then traversing the list of the intersection points, we can determine in linear time whether L intersects e. Hence (C2) holds with $D(n) = O(n \log n)$.

10

As for (C3), we choose a constant r and compute in linear-time a $(1/r)$-cutting Ξ of $G_0 \cup G_1$ of size $O(r^2)$ within τ [9]. For a triangle $\Delta \in \Xi$, let $G_i^\Delta \subseteq G_i$ be the set of lines that intersect Δ. Let a_i^Δ (resp., b_i^Δ) be the number of lines of G_i that lie below (resp., above) Δ. Then

$$
\begin{aligned}
f(G_0, G_1, \tau, a_0, a_1, b_0, b_1) &= \bigcup_{\Delta \in \Xi} f(G_0, G_1, \Delta, a_0, a_1, b_0, b_1) \\
&= \bigcup_{\Delta \in \Xi} f(G_0^\Delta, G_1^\Delta, \Delta, a_0 - a_0^\Delta, a_1 - a_1^\Delta, b_0 - b_0^\Delta, b_1 - b_1^\Delta).
\end{aligned}
$$

Since $|G_0^\Delta \cup G_1^\Delta| \leq |G_0 \cup G_1|/r$, condition (C3) is satisfied.

Hence, we can compute a basis for $\omega(H_0, H_1, \mathbb{R}^2, k_0, k_1, k_0, k_1)$ in randomized expected $O(n \log n)$ time. Putting everything together, we conclude the following.

Theorem 3.2 *Given two finite point sets P_0, P_1 in $\mathbb{R}^3$ with a total of n points and a direction u, we can compute a center-transversal line for P_0, P_1 in direction u, or decide that no such line exists, in $O(n \log n)$ expected time.*

4 Variations

Bichromatically deepest line. The algorithm that we have presented in Section 2 can be extended so that, for any given number $\alpha \in [0, 1]$, it finds a line ℓ with the property that any closed halfspace containing ℓ also contains at least $\lceil \alpha |P_i| \rceil$ points of P_i, for $i = 0, 1$, or determines that no such line exists. The running time remains $O(n^{1+\varepsilon} \kappa^2(n))$, for any $\varepsilon > 0$.

We define the *bichromatic depth* of a line ℓ with respect to P_0, P_1 as follows:

$$
\text{DEPTH}(\ell; P_0, P_1) = \min_h \left\{ \frac{|P_0 \cap h|}{|P_0|}, \frac{|P_1 \cap h|}{|P_1|} \right\} \in [0, 1],
$$

where the minimum is taken over all closed halfspaces h containing ℓ. Equivalently, $\text{DEPTH}(\ell; P_0, P_1) \geq \alpha$ means that any closed halfspace containing ℓ also contains at least $\lceil \alpha |P_i| \rceil$ points of P_i, for $i = 0, 1$. A line ℓ_0 is a *bichromatically deepest line* if it has maximum bichromatic depth. The center-transversal theorem (Theorem 1.1) implies that there always exists a line of depth at least $1/3$. By conducting a binary search and using the extended version of the algorithm of Section 2, we can easily find a line with maximum depth. We thus obtain the following.

Theorem 4.1 *Given two finite point sets P_0, P_1 in $\mathbb{R}^3$ with a total of n points, we can compute a bichromatically deepest line for P_0, P_1 in $O(n^{1+\varepsilon} \kappa^2(n))$ time, for any $\varepsilon > 0$.*

Computing an almost-deepest line. We next observe that, for any fixed $\delta > 0$, we can compute in linear time a line ℓ whose bichromatic depth with respect to P_0, P_1 is at least $1 - \delta$ times the maximum depth of a line. An *ε-approximation* of a point set P (with respect to closed halfspace ranges) is a subset $A \subseteq P$ such that, for any closed halfspace h we have

$$
\left| \frac{|A \cap h|}{|A|} - \frac{|P \cap h|}{|P|} \right| \leq \varepsilon.
$$

11

As is well known [9], for any fixed ε, an ε-approximation of size $O\left(\frac{1}{\varepsilon^2}\log\frac{1}{\varepsilon}\right)$ can be computed deterministically in $O(n)$ time.

We fix $\varepsilon = \frac{\delta}{6}$, and compute for each P_i an ε-approximation subset $A_i \subset P_i$ as above. We then compute a bichromatic deepest line ℓ_A for A_0 and A_1 in $O(1)$ time and return ℓ_A. We now argue that ℓ_A is an almost-deepest line. Observe that for any line ℓ we have (where h ranges over all closed halfspaces containing ℓ)

$$\begin{aligned}
\text{DEPTH}(\ell; P_0, P_1) &= \min_h \min_{i=0,1}\{|P_i \cap h|/|P_i|\} \geq \min_h \min_{i=0,1}\{|A_i \cap h|/|A_i|\} - \varepsilon \\
&= \text{DEPTH}(\ell; A_0, A_1) - \varepsilon,
\end{aligned}$$

and similarly

$$\text{DEPTH}(\ell; P_0, P_1) \leq \text{DEPTH}(\ell; A_0, A_1) + \varepsilon.$$

Let ℓ_{opt} be a bichromatically deepest line for P_0, P_1. Since $\text{DEPTH}(\ell_{opt}; P_0, P_1) \geq \frac{1}{3}$, we have

$$\begin{aligned}
\text{DEPTH}(\ell_A; P_0, P_1) &\geq \text{DEPTH}(\ell_A; A_0, A_1) - \varepsilon \geq \text{DEPTH}(\ell_{opt}; A_0, A_1) - \varepsilon \\
&\geq \text{DEPTH}(\ell_{opt}; P_0, P_1) - \frac{\delta}{3} \geq (1-\delta)\text{DEPTH}(\ell_{opt}; P_0, P_1).
\end{aligned}$$

We thus conclude the following.

Theorem 4.2 *For a fixed parameter $\delta > 0$, and two finite point sets $P_0, P_1 \subset \mathbb{R}^3$ with a total of n points, we can compute in $O(n)$ time a line ℓ whose bichromatic depth is at least $1 - \delta$ times the maximum bichromatic depth.*

5   Conclusions

The efficiency of our algorithm in Section 2 depends on the worst-case complexity $\kappa(n)$ of a k-level in an arrangement of n planes in three dimensions. The currently best known bound $\kappa(n) = O(n^{5/2})$ of [26] is probably not tight, and reducing it would have direct impact on the running time bound of our algorithm. Also, it is not clear that our approach best exploits the geometric structure of the problem in 3-space. For example, the analysis of Section 3 gives a simple reduction of the problem to the problem of finding a direction in which the projections of P_0 and P_1 have a common center point. Can we find an efficient characterization of "critical" directions of this kind, and then test each of them efficiently? Finally, can our algorithm be turned into a constructive proof of the existence of a center-transversal line?

Acknowledgements

References

[1] P. K. Agarwal. Ray shooting and other applications of spanning trees with low stabbing number. *SIAM J. Comput.*, 21:540–570, 1992.

[2] P. K. Agarwal and J. Matoušek. Dynamic half-space range reporting and its applications. *Algorithmica*, 13:325–345, 1995.

[3] P. K. Agarwal and M. Sharir. Arrangements and their applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 49–119. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[4] P. K. Agarwal, M. Sharir, and E. Welzl. Algorithms for center and Tverberg points. In *Proc. 20th Annu. ACM Sympos. Comput. Geom.*, pages 61–67, 2004. Also to appear in *ACM Trans. Algorithms*.

[5] M. Bern and D. Eppstein. Multivariate regression depth. *Discrete Comput. Geom.*, 28(1):1–17, 2002.

[6] T. M. Chan. Geometric applications of a randomized optimization technique. *Discrete Comput. Geom.*, 22:547–567, 1999.

[7] T. M. Chan. An optimal randomized algorithm for maximum Tukey depth. In *Proc. 15th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 430–436, 2004.

[8] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.

[9] B. Chazelle. *The Discrepancy Method: Randomness and Complexity*. Cambridge University Press, New York, 2001.

[10] B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir. Algorithms for bichromatic line segment problems and polyhedral terrains. *Algorithmica*, 11:116–132, 1994.

[11] K. L. Clarkson, D. Eppstein, G. L. Miller, C. Sturtivant, and S.-H. Teng. Approximating center points with iterated Radon points. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 91–98, 1993.

[12] R. Cole, M. Sharir, and C. K. Yap. On k-hulls and related problems. *SIAM J. Comput.*, 16:61–77, 1987.

[13] T. K. Dey. Improved bounds on planar k-sets and related problems. *Discrete Comput. Geom.*, 19:373–382, 1998.

[14] V. Dol'nikov. A generalization of the sandwich theorem. *Mathematical Notes*, 52:771–779, 1992.

[15] D. Halperin and M. Sharir. New bounds for lower envelopes in three dimensions, with applications to visibility in terrains. *Discrete Comput. Geom.*, 12:313–326, 1994.

[16] M. Hubert and P. Rousseeuw. The catline for deep regression. *J. Multivariate Analysis*, 66:270–296, 1998.

[17] S. Jadhav and A. Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. *Discrete Comput. Geom.*, 12:291–312, 1994.

[18] J. Matoušek. Computing the center of planar point sets. In J. E. Goodman, R. Pollack, and W. Steiger, editors, *Computational Geometry: Papers from the DIMACS Special Year*, pages 221–230. American Mathematical Society, Providence, 1991.

[19] J. Matoušek. *Lectures on Discrete Geometry*. Springer Verlag, Berlin, 2002.

[20] N. Naor and M. Sharir. Computing a point in the center of a point set in three dimensions. In *Proc. 2nd Canad. Conf. Comput. Geom.*, pages 10–13, 1990.

[21] M. Pellegrini. On lines missing polyhedral sets in 3-space. *Discrete Comput. Geom.*, 12:203–221, 1994.

[22] M. Pellegrini. Ray shooting and lines in space. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 37, pages 839–856. CRC Press LLC, Boca Raton, FL, 2nd edition, 2004.

[23] M. Pellegrini and P. Shor. Finding stabbing lines in 3-space. *Discrete Comput. Geom.*, 8:191–208, 1992.

[24] P. Rousseeuw and M. Hubert. Depth in an arrangement of hyperplanes. *Discrete Comput. Geom.*, 22:167–176, 1999.

[25] P. Rousseeuw and M. Hubert. Regression depth. *J. Amer. Stat. Assoc.*, 94:388–402, 1999.

[26] M. Sharir, S. Smorodinsky, and G. Tardos. An improved bound for k-sets in three dimensions. *Discrete Comput. Geom.*, 26:195–204, 2001.

[27] M. Sharir and E. Welzl. A combinatorial bound for linear programming and related problems. In *Proc. 9th Annu. Sympos. Theoretical Aspects of Computer Science*, pages 569–579, 1992.

[28] G. Toth. Point sets with many k-sets. *Discrete Comput. Geom.*, 26(2):187–194, 2001.

[29] R. T. Živaljević and S. T. Vrećica. An extension of the ham sandwich theorem. *Bull. London Math. Soc.*, 22:183–186, 1990.