

Generalized Halfspaces in Restricted-Orientation Convexity[‡]

Eugene Fink* and Derick Wood[†]

Technical Report HKUST-CS95-45
September 1995

*Department of Computer Science
Carnegie–Mellon University
Pittsburgh, PA 15213
U. S. A.

[†]Department of Computer Science
Hong Kong University of Science & Technology
Clear Water Bay, Kowloon
Hong Kong

Abstract

Restricted-orientation convexity (*\mathcal{O} -convexity*) is the study of geometric objects whose intersections with lines from some fixed set \mathcal{O} of orientations are empty or connected. The notion of \mathcal{O} -convexity generalizes standard convexity, as well as several other types of non-traditional convexity.

We introduce and study \mathcal{O} -halfspaces, which are analogs of standard halfspaces in the theory of \mathcal{O} -convexity, and directed \mathcal{O} -halfspaces, which are restricted \mathcal{O} -halfspaces. We explore some of the basic properties of them and outline their relationships to \mathcal{O} -convex sets and \mathcal{O} -connected sets, which are restricted \mathcal{O} -convex sets, in two and more dimensions.

For \mathcal{O} -halfspaces, we prove that: Every \mathcal{O} -halfspace is \mathcal{O} -convex; if \mathcal{O} has the point-intersection property, then the number of connected components of an \mathcal{O} -halfspace in d dimensions is at most 2^{d-1} , and this bound is attainable; and the closed complement of an \mathcal{O} -halfspace is an \mathcal{O} -halfspace if and only if the boundary of the \mathcal{O} -halfspace is \mathcal{O} -convex. In addition, for directed \mathcal{O} -halfspaces, we prove that: Every \mathcal{O} -halfspace is \mathcal{O} -connected, and the closed complement of a directed \mathcal{O} -halfspace is a directed \mathcal{O} -halfspace.

[‡]This research was supported by grants from the Natural Sciences and Engineering Research Council of Canada and from the Information Technology Research Centre of Ontario.



Generalized Halfspaces in Restricted-Orientation Convexity¹

Eugene Fink²

Derick Wood³

Abstract

Restricted-orientation convexity (*\mathcal{O} -convexity*) is the study of geometric objects whose intersections with lines from some fixed set \mathcal{O} of orientations are empty or connected. The notion of \mathcal{O} -convexity generalizes standard convexity, as well as several other types of nontraditional convexity.

We introduce and study \mathcal{O} -halfspaces, which are analogs of standard halfspaces in the theory of \mathcal{O} -convexity, and directed \mathcal{O} -halfspaces, which are restricted \mathcal{O} -halfspaces. We explore some of the basic properties of them and outline their relationships to \mathcal{O} -convex sets and \mathcal{O} -connected sets, which are restricted \mathcal{O} -convex sets, in two and more dimensions.

For \mathcal{O} -halfspaces, we prove that: Every \mathcal{O} -halfspace is \mathcal{O} -convex; if \mathcal{O} has the point-intersection property, then the number of connected components of an \mathcal{O} -halfspace in d dimensions is at most 2^{d-1} , and this bound is attainable; an interior-closed set with an \mathcal{O} -convex boundary is an \mathcal{O} -halfspace; and the closed complement of an \mathcal{O} -halfspace is an \mathcal{O} -halfspace if and only if the boundary of the \mathcal{O} -halfspace is \mathcal{O} -convex.

In addition, for directed \mathcal{O} -halfspaces, we prove that: Every \mathcal{O} -halfspace is \mathcal{O} -connected; if \mathcal{O} has the point-intersection property, then the intersection of a directed \mathcal{O} -halfspace with every \mathcal{O} -flat is connected; the boundary of a directed \mathcal{O} -halfspace is \mathcal{O} -convex; if \mathcal{O} has the point-intersection property, then the boundary of a directed \mathcal{O} -halfspace is \mathcal{O} -connected; and the closed complement of a directed \mathcal{O} -halfspace is a directed \mathcal{O} -halfspace.

1 Introduction

The study of convex sets is a branch of geometry, analysis, and linear algebra [5; 7] that has applications in many practical areas of computer science, including VLSI design, com-

¹This work was supported under grants from the Natural Sciences and Engineering Research Council of Canada and the Information Technology Research Centre of Ontario.

²Department of Computer Science, Carnegie–Mellon University, Pittsburgh, PA, U. S. A.

³Department of Computer Science, Hong Kong University of Science & Technology, Clear Water Bay, Kowloon, Hong Kong.

puter graphics, architectural databases, and geographic databases [12]. Many notions of nontraditional convexity have been studied, such as orthogonal convexity [9; 10; 11], finitely oriented convexity [6; 20; 14], strong convexity [13; 4], NESW convexity [8; 18; 20], and link convexity [1; 19; 17].

Rawlins introduced the notion of **restricted-orientation convexity**, also called **\mathcal{O} -convexity**, in his doctoral dissertation [13] and demonstrated that this notion generalized standard convexity and orthogonal convexity. Rawlins and Wood have studied \mathcal{O} -convex sets and demonstrated that the properties of these sets are similar to the properties of standard convex sets [15; 16]. Schuierer has continued their exploration and presented an extensive study of geometrical and computational properties of \mathcal{O} -convex sets in his doctoral thesis [17].

Our goal is to investigate nontraditional convexities in multidimensional space. We generalized **strong \mathcal{O} -convexity** [13] to higher dimensions and explored the properties of strongly \mathcal{O} -convex sets [4; 3]. We then demonstrated that \mathcal{O} -convexity can also be extended to multidimensional space [2]. We introduced **\mathcal{O} -convex sets** and **\mathcal{O} -connected sets**, proved some of the major properties of them, and showed that their properties are much richer than the properties of planar \mathcal{O} -convex sets. The notion of multidimensional \mathcal{O} -convexity generalizes not only planar \mathcal{O} -convexity, but also standard multidimensional convexity.

We now further develop the theory of \mathcal{O} -convexity by introducing and investigating \mathcal{O} -halfspaces and directed \mathcal{O} -halfspaces, and describing their relationship to \mathcal{O} -convex and \mathcal{O} -connected sets. We restrict our attention to the exploration of closed sets. We conjecture that most of the results hold for nonclosed sets; however, some of our proofs work only for closed sets.

Closed-Set Assumption: We consider only closed geometric objects. An object is **closed** if, for every convergent sequence of its points, the limit of the sequence belongs to the object.

The paper is organized as follows. In Section 2, we briefly describe the notion of \mathcal{O} -convexity in two dimensions and introduce \mathcal{O} -halfplanes. In Section 3, we extend the definition of \mathcal{O} -convexity to higher dimensions and define **\mathcal{O} -halfspaces**. In particular, we prove that in d dimensions, an \mathcal{O} -halfspace has at most 2^{d-1} connected components and that this bound is attainable for all d . In Section 4, we introduce **directed \mathcal{O} -halfspaces** that are \mathcal{O} -halfspaces with several special properties. In Section 5, we study the boundaries of \mathcal{O} -halfspaces and present an analog of the boundary-convexity property for standard convex sets, and, in Section 6, we generalize the complementation property for standard halfspaces to \mathcal{O} -halfspaces. We conclude, in Section 7, with some open problems and suggestions for further work.




Figure 1: Ortho-convex sets (a–c) and ortho-halfplanes (d–h).

2 \mathcal{O} -convexity and \mathcal{O} -halfplanes in two dimensions

We begin by reviewing the notions of orthogonal convexity [11] and \mathcal{O} -convexity [13] in two dimensions and defining analogs of halfplanes for these two types of convexity. We then explore basic properties of these \mathcal{O} -convexity halfplanes.

We can describe convex sets through their intersections with lines: a set of points is convex if its intersection with every line is connected. Note that *we consider the empty set to be connected*; this convention simplifies our definitions and results.

We define **orthogonal convexity** by considering the intersections of sets of points with vertical and horizontal lines. A closed set is **ortho-convex** if its intersection with every vertical and horizontal line is connected. We give three examples of ortho-convex sets in Figures 1(a)–(c). Note that, unlike standard convex sets, ortho-convex sets may be disconnected (see Figure 1c).

Standard halfplanes can also be characterized in terms of their intersections with lines.

Proposition 2.1 *A set is a halfplane if and only if its intersection with every line is empty, a ray, or a line.*

We define ortho-halfplanes in terms of their intersection with vertical and horizontal lines. A closed set is an **ortho-halfplane** if its intersection with every vertical and horizontal line is empty, a ray, or a line. We give five examples of ortho-halfplanes in Figures 1(d)–(h). The last of these examples demonstrates that ortho-halfplanes may be disconnected. Note that the *empty set* and the *whole plane* are considered to be ortho-halfplanes. This convention also simplifies some of the results.

We now present some basic properties of ortho-halfplanes.

Lemma 2.2

1. *Every translation of an ortho-halfplane is an ortho-halfplane.*
2. *Every standard halfplane is an ortho-halfplane.*




Figure 2: Proof of Lemma 2.2.

3. *Every ortho-halfplane is ortho-convex.*
4. *An ortho-halfplane is either connected or consists of two connected components.*
5. *A disconnected set is an ortho-halfplane if and only if each of its components is an ortho-halfplane and no vertical or horizontal line intersects two components.*

Proof.

(1) If the intersection of a set with every vertical and horizontal line is empty, a ray, or a line, then the same holds for every translation of the set.

(2) The intersection of a standard halfplane with every vertical and horizontal line is empty, a ray, or a line. Therefore, every halfplane is an ortho-halfplane.

(3) The intersection of an ortho-halfplane with every vertical and horizontal line is connected. Therefore, every ortho-halfplane is ortho-convex.

(4) We prove that every ortho-halfplane P has at most two connected components by demonstrating that, for every three points $p, q, x \in P$, two of them are in the same component. Since P is an ortho-halfplane, one of the two horizontal rays with endpoint p is contained in P (we show this ray in Figure 2). Similarly, we may choose a horizontal ray with endpoint q and a horizontal ray with endpoint x contained in P .

We select two of these three rays that point in the same direction. Without loss of generality, we assume that the two selected rays are the rays with endpoints p and q . We choose a vertical line that intersects the two selected rays and denote the intersection points by y and z , respectively (see Figure 2). The polygonal line (p, y, z, q) is wholly in P ; therefore, p and q are in the same connected component.

(5) If P is the union of disjoint ortho-halfplanes and no vertical or horizontal line intersects two of these ortho-halfplanes, then the intersection of P with every vertical and horizontal line is empty, a ray, or a line and, hence, P is an ortho-halfplane.

If one of P 's components is not an ortho-halfplane, the intersection of this component with some vertical or horizontal line is not empty, not a ray, and not a line. Then, the intersection of P with this line is not empty, not a ray, and not a line; therefore, P is not an ortho-halfplane. Finally, if some vertical or horizontal line intersect two components, then the intersection of P with this line is disconnected; therefore, P is not an ortho-halfplane. \square

We next describe \mathcal{O} -convexity [13], which is a generalization of ortho-convexity and standard convexity [15], and a corresponding analog of halfplanes.




Figure 3: Planar \mathcal{O} -convex sets (b–d) and \mathcal{O} -halfplanes (e–g).

To obtain this generalization, we first introduce the notion of an orientation set. We define an **orientation set** \mathcal{O} as a (finite or infinite) set of lines through some fixed point o . An example of a finite orientation set is shown in Figure 3(a). A line parallel to one of the lines of \mathcal{O} is called an \mathcal{O} -line. For example, the dashed lines in Figure 3 are \mathcal{O} -lines. We define \mathcal{O} -convex sets in terms of their intersections with \mathcal{O} -lines.

Definition 2.1 (\mathcal{O} -Convexity) *A closed set is \mathcal{O} -convex if its intersection with every \mathcal{O} -line is connected.*

For the orientation set given in Figure 3(a), the sets in Figures 3(b)–(d) are \mathcal{O} -convex (some \mathcal{O} -lines intersecting these sets are shown by dashed lines).

We define an \mathcal{O} -halfplane, in the same way as ortho-halfplanes: A closed set is an **\mathcal{O} -halfplane** if its intersection with every \mathcal{O} -line is empty, a ray, or a line. We show examples of \mathcal{O} -halfplanes in Figures 3(e)–(g).

Let us recall the properties of ortho-halfplanes given in Lemma 2.2. We can readily generalize Properties 1–3 and 5: These properties hold for \mathcal{O} -halfplanes and their proofs are the same as the proofs for ortho-halfplanes. If the orientation set \mathcal{O} contains at least two distinct lines, then Property 4 also holds (again, the proof is the same as the one for ortho-halfplanes).

We next present the intersection property of \mathcal{O} -convex sets [16], which generalizes the observation that every standard convex set is the intersection of halfplanes.

Lemma 2.3 *A connected set is \mathcal{O} -convex if and only if it is the intersection of \mathcal{O} -halfplanes.*

Proof. Suppose that a set P is the intersection of \mathcal{O} -halfplanes. For every \mathcal{O} -line l , the intersection of each \mathcal{O} -halfplane with l is connected and, hence, the intersection of P with l is also connected. Therefore, P is \mathcal{O} -convex.

Suppose, conversely, that P is \mathcal{O} -convex. We show that P is the intersection of \mathcal{O} -halfplanes by demonstrating that, for every point p not in P , some \mathcal{O} -halfplane contains P and does not contain p .




Figure 4: Proof of Lemma 2.3.




Figure 5: Proof of Lemma 2.4.

We draw the two lines through p that support P (see Figure 4a). If the marked angle between these lines is less than π , then there is a standard halfplane that contains P and does not contain p (Figure 4b). This halfplane is the desired \mathcal{O} -halfplane.

If the marked angle is greater than or equal to π , we consider the set Q shown by shading in Figure 4(c). The boundary of Q consists of the segment of P 's boundary between the supporting lines and the parts of the supporting lines that extend this segment. We show that Q is an \mathcal{O} -halfplane.

If the intersection of Q with some \mathcal{O} -line l is disconnected, then there is an \mathcal{O} -line parallel to l whose intersection with P is disconnected (see Figure 4d), contradicting the assumption that P is \mathcal{O} -convex. On the other hand, there is no line whose intersection with Q is a point or segment. Therefore, the intersection of Q with every \mathcal{O} -line is empty, a ray, or a line. \square

If \mathcal{O} contains at least three lines, \mathcal{O} -halfplanes have some additional basic properties.

Lemma 2.4 *Suppose that the orientation set \mathcal{O} contains at least three distinct lines. Then, if the intersection of an \mathcal{O} -halfplane with two parallel \mathcal{O} -lines forms two rays, these rays point in the same direction (rather than in opposite directions).*

Proof. Suppose that the intersection of an \mathcal{O} -halfplane P with parallel \mathcal{O} -lines l_1 and l_2 forms rays that point in opposite directions (in Figure 5(b), we show these two rays by solid lines). We assume, for convenience, that l_1 is below l_2 and the lower ray points to the left.

Let l be the element of \mathcal{O} parallel to l_1 and l_2 , and let m and n be two other elements of \mathcal{O} (see Figure 5a). We assume that the marked angle between l and m is smaller than the marked angle between l and n .

We choose a point $p \in l_1$ and draw a line n_1 through p parallel to n . We select this point p in such a way that p is *not* in P and n_1 intersects the upper ray (see Figure 5b).



Figure 6: (a) \mathcal{O} -halfplane is directed. (b) \mathcal{O} -halfplane is *not* directed.




Figure 7: Proof of Lemma 2.5.

Since n_1 is an \mathcal{O} -line, its intersection with P must be empty, a ray, or a line. Therefore, the part of n_1 above l_2 (shown by a solid line) is in P .

We now choose a point $q \in l_2$ and draw a line m_1 through q parallel to m . We select q is such a way that q is *not* in P and m_1 intersects the lower ray. Note that m_1 intersects n_1 above l_2 ; that is, m_1 intersects the part of n_1 contained in P . Since m_1 is an \mathcal{O} -line, its intersection with P is connected; therefore, the segment of m_1 between l_1 and n_1 is in P , contradicting the assumption that q is *not* in P . \square

We illustrate the directed-ray property of \mathcal{O} -halfplanes in Figure 6(a), where the intersection of an \mathcal{O} -halfplane with several parallel \mathcal{O} -lines is shown by dashed rays. The \mathcal{O} -halfplanes that satisfy this directed-ray property are called **directed \mathcal{O} -halfplanes**.

If \mathcal{O} contains two lines, an \mathcal{O} -halfplane may not be directed. For example, the \mathcal{O} -halfplane in Figure 6(b) is not directed, since the dashed \mathcal{O} -rays point in opposite directions.

In the following result, we provide two basic properties of directed \mathcal{O} -halfplanes.

Lemma 2.5 *Suppose that \mathcal{O} contains at least two distinct lines. Then:*

1. *Every directed \mathcal{O} -halfplane is connected.*
2. *The boundary of every directed \mathcal{O} -halfplane is connected and \mathcal{O} -convex.*

Proof.

(1) We show that every two points p and q of a directed \mathcal{O} -halfplane P can be connected by a polygonal line in P . We choose two parallel \mathcal{O} -rays, with endpoints p and q , that are contained in P and point in the same direction (see Figure 7a). We next choose an \mathcal{O} -line that intersects these two rays and denote the intersection points by x and y , respectively. The polygonal line (p, x, y, q) is wholly in P .

(2) Suppose that the boundary of a directed \mathcal{O} -halfplane P is *not* connected. Since P is connected, the complement of P is disconnected and we may choose points p and q in

different connected components of P 's complement. We next choose two parallel \mathcal{O} -rays, with endpoints p and q , that do *not* intersect P and point in the same direction (we reuse Figure 7(a) to illustrate this construction). Finally, we choose an \mathcal{O} -line l that intersects these two rays and denote the intersection points by x and y , respectively.

The segment of l between x and y does *not* intersect P : if some point z of this segment were in P , then one of the two contained-in- l rays with endpoint z would be in P , contradicting the assumption that x and y are not in P . Therefore, the polygonal line (p, x, y, q) is wholly in P 's complement, contradicting the assumption that p and q are in different components of P 's complement.

Now suppose that the boundary of P is *not* \mathcal{O} -convex. Then, the intersection of some \mathcal{O} -line l with P 's boundary is disconnected and we may select points $p, q \in l$ that are in the boundary and a point $x \in l$ between them that is *not* in the boundary (see Figure 7b). We assume, for convenience, that p is to the left of x .

Since the intersection of P with l is connected, x is in the interior of P and we can choose a circle $B_x \subseteq P$ centered at x . Either all left-pointed or all right-pointed rays with endpoints in B_x are contained in P . We assume that the left-pointed rays (shown by dotted lines) are in P . Then, some circle B_p centered at p is wholly contained in P ; therefore, p is in P 's interior, which yields a contradiction. \square

3 \mathcal{O} -convexity and \mathcal{O} -halfspaces in higher dimensions

We extend the notion of \mathcal{O} -convexity and \mathcal{O} -halfspaces to d -dimensional space \mathcal{R}^d . We assume that the space \mathcal{R}^d is fixed; however, the results are independent of the particular value of d . We introduce a set \mathcal{O} of hyperplanes through a fixed point o , show how this set gives rise to \mathcal{O} -lines, and define \mathcal{O} -halfspaces in terms of their intersections with \mathcal{O} -lines. We then explore basic properties of \mathcal{O} -halfspaces.

A **hyperplane** in d dimensions is a subset of \mathcal{R}^d that is a $(d-1)$ -dimensional space. For example, hyperplanes in three dimensions are the usual planes. Analytically, we can define a hyperplane as a set of points satisfying a linear equation, $a_0 + a_1x_1 + a_2x_2 + \dots + a_dx_d = 0$, in Cartesian coordinates. Two hyperplanes are **parallel** if they are translations of each other. Analytically, two hyperplanes are parallel if their equations differ only by the value of a_0 .

Definition 3.1 (Orientation set and \mathcal{O} -hyperplanes) *An orientation set \mathcal{O} is a set of hyperplanes through a fixed point o . A hyperplane parallel to one of the elements of \mathcal{O} is an \mathcal{O} -hyperplane.*

Note that every translation of an \mathcal{O} -hyperplane is an \mathcal{O} -hyperplane and a particular choice of the point o is not important. In Figure 8, we show examples of finite orientation sets in three dimensions. The set in Figure 8(a) contains three mutually orthogonal planes; we call it an **orthogonal-orientation set**.




Figure 8: Examples of finite orientation sets.

\mathcal{O} -lines in \mathcal{R}^d are formed by the intersections of \mathcal{O} -hyperplanes. In other words, a line is an \mathcal{O} -line if it is the intersection of several \mathcal{O} -hyperplanes. Note that every translation of an \mathcal{O} -line is an \mathcal{O} -line.

Since every \mathcal{O} -hyperplane is parallel to one of the hyperplanes of the orientation set \mathcal{O} , every \mathcal{O} -line is parallel to some line formed by the intersection of several elements of \mathcal{O} . For example, the intersections of the four planes of the orientation set given in Figure 8(b) form six different lines through o and every \mathcal{O} -line is parallel to one of these six lines.

Lemma 3.1 *If the orientation set \mathcal{O} is nonempty and the intersection of the elements of \mathcal{O} is the point o (rather than a superset of o), then:*

1. *There is at least one \mathcal{O} -line.*
2. *For every line, there is an \mathcal{O} -hyperplane that intersects it and does not contain it.*

Proof.

(1) We consider a minimal set of \mathcal{O} -hyperplanes whose intersection forms o . We denote these hyperplanes by $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_n$. Then, $\mathcal{H}_2 \cap \dots \cap \mathcal{H}_n$ is a flat (affine variety) different from the point o and the intersection of this flat with \mathcal{H}_1 is o , which may happen only if $\mathcal{H}_2 \cap \dots \cap \mathcal{H}_n$ is a line. This line is an \mathcal{O} -line.

(2) We consider a line l and assume, without loss of generality, that this line is through o . The intersection of all elements of \mathcal{O} is o ; therefore, for some \mathcal{O} -hyperplane $\mathcal{H} \in \mathcal{O}$, the line l is not contained in \mathcal{H} . \square

When \mathcal{O} is nonempty and the intersection of the elements of \mathcal{O} is the point o , we say that \mathcal{O} has the **point-intersection property**. For example, the orientation sets in Figures 8(a) and 8(b) have this property, whereas the set in Figure 8(c) does not. Some of our results hold only for orientation sets with the point-intersection property and we use Lemma 3.1 to prove these results.

If the orientation set \mathcal{O} does not have the point-intersection property, then the intersection of the elements of \mathcal{O} is either a line or a higher-dimensional flat (affine variety). If this intersection is a line (see Figure 8c), then there is exactly one \mathcal{O} -line through o and all other

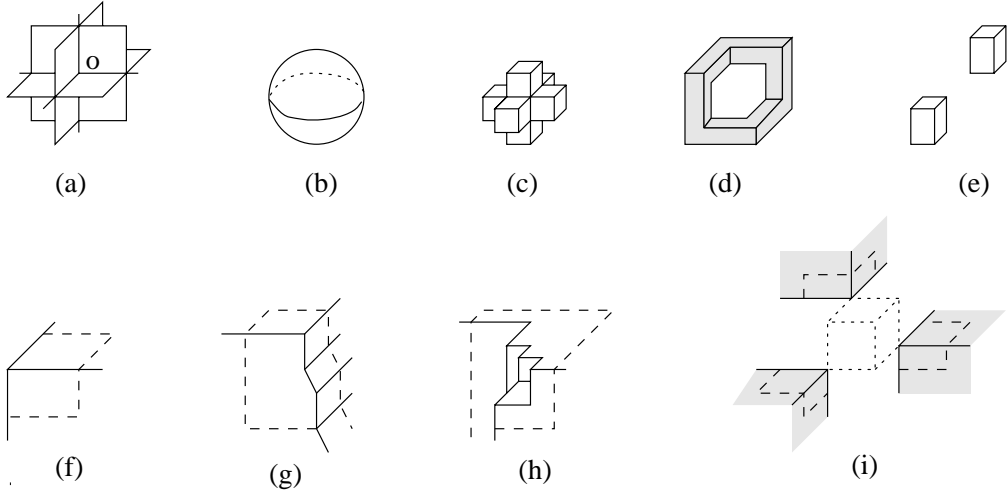


Figure 9: \mathcal{O} -convex sets (b–e) and \mathcal{O} -halfspaces (f–i) in three dimensions.

\mathcal{O} -lines are parallel to this \mathcal{O} -line. If the intersection is neither a point nor a line, then there are no \mathcal{O} -lines at all.

We define \mathcal{O} -convex sets in higher dimensions in the same way as in two dimensions: a closed set in \mathcal{R}^d is \mathcal{O} -convex if its intersection with every \mathcal{O} -line is connected. For example, the sets in Figures 9(b)–(e) are \mathcal{O} -convex for the orthogonal-orientation set shown in Figure 9(a). The notion of \mathcal{O} -halfspaces is the higher-dimensional analog of \mathcal{O} -halfplanes.

Definition 3.2 (\mathcal{O} -halfspaces) *A closed set in \mathcal{R}^d is an \mathcal{O} -halfspace if its intersection with every \mathcal{O} -line is empty, a ray, or a line.*

Figures 3(f)–(i) provide examples of \mathcal{O} -halfspaces for the orthogonal-orientation set. We use dashed lines to show infinite planar regions in the boundaries of these \mathcal{O} -halfspaces. Note that disconnected \mathcal{O} -halfspaces in higher dimensions may have more than two components. For example, the \mathcal{O} -halfspace in Figure 3(i) consists of three components, located around the dotted cube.

We next present some simple properties of \mathcal{O} -halfspaces and compare them with properties of \mathcal{O} -halfplanes. We begin by recalling the properties of ortho-halfplanes given in Lemma 2.2. Properties 1–3 hold for \mathcal{O} -halfspaces and their proofs are the same as the proofs for ortho-halfplanes. The most important of them is Property 3: every \mathcal{O} -halfspace is \mathcal{O} -convex, which is a generalization of the convexity property of standard halfspaces. In the next result, we give necessary and sufficient conditions under which an \mathcal{O} -convex set is an \mathcal{O} -halfspace.

Lemma 3.2 *A set P is an \mathcal{O} -halfspace if and only if the following two conditions hold:*

1. P is an \mathcal{O} -convex set.

2. For every point p in P and every \mathcal{O} -line l , one of the two parallel-to- l rays with endpoint p is wholly contained in P .

Proof. Clearly, Conditions 1 and 2 hold for every \mathcal{O} -halfspace. Suppose, conversely, that P satisfies the two conditions and let us show that the intersection of P with every \mathcal{O} -line is empty, a ray, or a line. Since P is \mathcal{O} -convex (Condition 1), the intersection of P with an \mathcal{O} -line is empty, a point, a segment, a ray, or a line. By Condition 2, the intersection is neither a point nor a segment; therefore, it is empty, a ray, or a line. \square

We next generalize Properties 4 and 5 of \mathcal{O} -halfplanes to \mathcal{O} -halfspaces. We show that every component of a disconnected \mathcal{O} -halfspace is an \mathcal{O} -halfspace and that, if the orientation set has the point-intersection property, then the number of components of an \mathcal{O} -halfspace is bounded.

Theorem 3.3

1. A disconnected set is an \mathcal{O} -halfspace if and only if every connected component of the set is an \mathcal{O} -halfspace and no \mathcal{O} -line intersects two components.
2. In d dimensions, if the orientation set \mathcal{O} has the point-intersection property, then the number of connected components of an \mathcal{O} -halfspace is at most 2^{d-1} .

Proof.

(1) If P is the union of \mathcal{O} -halfspaces and no \mathcal{O} -line intersects two of them, then, clearly, P is an \mathcal{O} -halfspace. If some connected component of P is not an \mathcal{O} -halfspace, then the intersection of this component with some \mathcal{O} -line is not empty, not a ray, and not a line; therefore, the intersection of P with this \mathcal{O} -line is not empty, not a ray, and not a line. Finally, if some \mathcal{O} -line intersects two components, the intersection of P with this line is disconnected.

(2) We show that, for every $(2^{d-1} + 1)$ points of an \mathcal{O} -halfspace P , two of these points are in the same connected component. We use induction on the dimension d . We have already proved this result for \mathcal{O} -halfplanes, which provides us with an induction base.

The induction step is based on Lemma 4.5, which we will present in Section 4: the intersection of an \mathcal{O} -halfspace P with an \mathcal{O} -hyperplane \mathcal{H} is an \mathcal{O} -halfspace in the $(d - 1)$ -dimensional space \mathcal{H} . That is, the intersection of $P \cap \mathcal{H}$ with every \mathcal{O} -line contained in \mathcal{H} is empty, a ray, or a line.

We denote the $(2^{d-1} + 1)$ points in P by $p_0, p_1, \dots, p_{2^{d-1}-1}$ (in Figure 10, we illustrate the proof for $d = 3$). By Lemma 3.1, since \mathcal{O} has the point-intersection property, we may choose some \mathcal{O} -line l and an \mathcal{O} -hyperplane \mathcal{H} that intersects l and does not contain it. For every point p_k , one of the two parallel-to- l rays with endpoint p_k is contained in P . Thus, we get $(2^{d-1} + 1)$ parallel rays in P , and at least $(2^{d-2} + 1)$ of these rays point in the same direction. We assume that these $(2^{d-1} + 1)$ same-direction rays correspond to the points $p_0, p_1, \dots, p_{2^{d-2}}$.




Figure 10: Proof of Lemma 3.3.




Figure 11: \mathcal{O} -halfspace with 2^{d-1} connected components in (a) two and (b) three dimensions.

We select an \mathcal{O} -plane \mathcal{H}' , parallel to \mathcal{H} , that intersects these $(2^{d-2} + 1)$ rays and denote the intersection points by $q_0, q_1, \dots, q_{2^{d-2}}$, respectively.

$P \cap \mathcal{H}'$ (shaded in Figure 10) is an \mathcal{O} -halfspace in the $(d - 1)$ -dimensional space \mathcal{H} . Therefore, by the induction hypothesis, some points q_k and q_m belong to the same connected component of $P \cap \mathcal{H}'$, which implies that p_k and p_m belong to the same component of P . \square

If the orientation set \mathcal{O} does not have the point-intersection property, then an \mathcal{O} -halfspace may have infinitely many components. For example, if there is only one \mathcal{O} -line through o , then any collection of lines parallel to this \mathcal{O} -line forms an \mathcal{O} -halfspace. Such an \mathcal{O} -halfspace, formed by the union of parallel lines, may have infinitely many components.

We next demonstrate that the bound on the number of components given in Theorem 3.3 is tight.

Example: An \mathcal{O} -halfspace that has 2^{d-1} connected components.

We consider the orthogonal-orientation set in d dimensions and construct an \mathcal{O} -halfspace P whose components are rectangular polyhedral angles (quadrants). Note the orthogonal-orientation set has the point-intersection property. In Figure 11, we illustrate the construction for two and three dimensions.

We choose a cube (shown by dotted lines) whose facets are parallel to elements of \mathcal{O} . The cube has 2^d vertices; we select 2^{d-1} vertices no two of which are adjacent. We define P as the union of the 2^{d-1} polyhedral angles vertical to the selected angles of the cube.

We may describe P analytically, using \mathcal{O} -lines through the cube's center (shown by dashed




Figure 12: Every \mathcal{O} -halfspace containing the \mathcal{O} -convex set P also contains the point $p \notin P$.

lines in Figure 11a) as coordinate axes and taking one of the cube's vertices for $(1, 1, \dots, 1)$. The equation for P is as follows:

$$|x_1|, |x_2|, \dots, |x_d| \geq 1 \text{ and } x_1 \times x_2 \times \dots \times x_d > 0.$$

To show that P is indeed an \mathcal{O} -halfspace, we observe that each \mathcal{O} -line either intersects one component, in which case the intersection is a ray, or does not intersect any component. \square

We have shown, in Lemma 2.3, that, in two dimensions, every connected \mathcal{O} -convex set is formed by the intersection of \mathcal{O} -halfplanes. In higher dimensions, the analog of this result does not hold.

Example: A connected \mathcal{O} -convex set that is not the intersection of \mathcal{O} -halfspaces. We consider the orthogonal-orientation set \mathcal{O} and the set P shown in Figure 12(b). This set is a horizontal disk on eight vertical, equally spaced “pillars.” The pillars are rays located in such a way that no \mathcal{O} -line intersects two of them. Clearly, P is \mathcal{O} -convex. We readily see that every \mathcal{O} -halfspace containing P also contains the point $p \notin P$, located under the center of the disk. Thus, the intersection of all \mathcal{O} -halfspaces containing P is *not* P . \square

4 Directed \mathcal{O} -halfspaces

We now consider the \mathcal{O} -halfspaces that have the directed-ray property as described in Section 2 and establish their special properties.

Definition 4.1 (Directed \mathcal{O} -halfspaces) *An \mathcal{O} -halfspace is a directed \mathcal{O} -halfspace if, for every two parallel \mathcal{O} -lines whose intersection with the \mathcal{O} -halfspace forms rays, these rays point in the same direction.*

For example, the \mathcal{O} -halfspaces in Figures 13(b) and 13(c) are directed for the orthogonal-orientation set shown in Figure 13(a). On the other hand, the \mathcal{O} -halfspace in Figure 13(d) is not directed, because the two dotted rays, formed by the intersection of this \mathcal{O} -halfspace with vertical \mathcal{O} -lines, point in opposite directions.

The next result immediately follows from the definition of directed \mathcal{O} -halfspaces.




Figure 13: \mathcal{O} -halfspaces (b) and (c) are directed. \mathcal{O} -halfspace (d) is *not* directed.

Lemma 4.1

1. Every translation of a directed \mathcal{O} -halfspace is a directed \mathcal{O} -halfspace.
2. Every standard halfspace is a directed \mathcal{O} -halfspace.

We next introduce the notion of **\mathcal{O} -flats**, which are formed by the intersections of \mathcal{O} -hyperplanes, and demonstrate that, for the orientation set with the point-intersection property, the intersection of a directed \mathcal{O} -halfspace with every \mathcal{O} -flat is connected. This property of directed \mathcal{O} -halfspaces is one of the main tools in our exploration.

A **flat**, also known as an **affine variety**, is a subset of \mathcal{R}^d that is itself a lower-dimensional space. For example, points, lines, two-dimensional planes, and hyperplanes are flats. The whole space \mathcal{R}^d is also a flat. Analytically, a k -dimensional flat is represented in Cartesian coordinates as a system of $(d - k)$ independent linear equations. We will use the following properties of flats.

Proposition 4.2 (Properties of flats)

1. The intersection of a collection of flats is either empty or a flat.
2. The intersection of a k -dimensional flat η and a hyperplane is empty, η , or a $(k - 1)$ -dimensional flat.

We now define \mathcal{O} -flats.

Definition 4.2 (\mathcal{O} -flats) *A flat formed by the intersection of several \mathcal{O} -hyperplanes in an \mathcal{O} -flat. \mathcal{O} -hyperplanes themselves and the whole space \mathcal{R}^d are also \mathcal{O} -flats.*

Since every \mathcal{O} -hyperplane is a translation of some element of \mathcal{O} , every \mathcal{O} -flat is a translation of some flat formed by the intersection of several elements of \mathcal{O} . For example, the orthogonal-orientation set in three dimensions (Figure 8a) gives rise to the following \mathcal{O} -flats through o : the whole space \mathcal{R}^3 , the three mutually orthogonal \mathcal{O} -planes, the three \mathcal{O} -lines formed by the intersections of these planes, and the point o .

The next two properties of \mathcal{O} -flats readily follow from the definition.




Figure 14: Lower-dimensional orientation sets.

Lemma 4.3

1. Every translation of an \mathcal{O} -flat is an \mathcal{O} -flat.
2. The intersection of a collection of \mathcal{O} -flats is either empty or an \mathcal{O} -flat.

We now describe lower-dimensional orientation sets contained in \mathcal{O} -flats. We consider an \mathcal{O} -flat η and denote the dimension of η by k . We treat η as an independent k -dimensional space and define the orientation set \mathcal{O}_η and \mathcal{O}_η -flats in this lower-dimensional space.

An \mathcal{O}_η -flat is an \mathcal{O} -flat contained in η . The $(k-1)$ -dimensional \mathcal{O}_η -flats (“ \mathcal{O}_η -hyperplanes”) through some fixed point o_η form the lower-dimensional orientation set \mathcal{O}_η .

For example, consider the orthogonal-orientation set in Figure 14(a). The \mathcal{O} -plane η contains vertical and horizontal \mathcal{O} -lines. Therefore, the lower-dimensional orientation set \mathcal{O}_η contains the vertical and horizontal line through o_η . In Figure 14(b), we show another example of a lower-dimensional orientation set.

We demonstrate that \mathcal{O}_η -flats have all necessary properties of \mathcal{O} -flats and that, if \mathcal{O} has the point-intersection property, then \mathcal{O}_η also has this property.

Lemma 4.4 *Let η be an \mathcal{O} -flat and k be the dimension of η .*

1. Every translation of an \mathcal{O}_η -flat within the space η is an \mathcal{O}_η -flat.
2. A set $H \subseteq \eta$ is an \mathcal{O}_η -flat if and only if it is either η itself or the intersection of several $(k-1)$ -dimensional \mathcal{O}_η -flats.
3. If the orientation set \mathcal{O} has the point-intersection property, then \mathcal{O}_η also has the point-intersection property; that is, the orientation set \mathcal{O}_η is nonempty and $\bigcap \mathcal{O}_\eta = o_\eta$.

Proof.

(1) Every \mathcal{O}_η -flat is an \mathcal{O} -flat; therefore, a translation of an \mathcal{O}_η -flat is also an \mathcal{O} -flat (Lemma 4.3). If the translation is in η , then it is an \mathcal{O}_η -flat.

(2) Since \mathcal{O}_η -flats are \mathcal{O} -flats, the intersection of $(k-1)$ -dimensional \mathcal{O}_η -flats is an \mathcal{O} -flat. This \mathcal{O} -flat is in η and, hence, it is an \mathcal{O}_η -flat.




Figure 15: For (a) any orientation set, the intersection of (b) an \mathcal{O} -halfspace with (c) an \mathcal{O} -flat η is (d) an \mathcal{O}_η -halfspace.

We next show, conversely, that every \mathcal{O}_η -flat H distinct from η is the intersection of $(k-1)$ -dimensional \mathcal{O}_η -flats. Since H is an \mathcal{O} -flat, it is formed by the intersection of several \mathcal{O} -hyperplanes, say $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_n$. Let $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$ be those hyperplanes among them that do *not* contain η . Then, $\mathcal{H}_1 \cap \eta, \mathcal{H}_2 \cap \eta, \dots, \mathcal{H}_k \cap \eta$ are $(k-1)$ -dimensional flats (see Proposition 4.2). These flats are \mathcal{O}_η -flats and their intersection forms H .

(3) We assume, without loss of generality, that $o_\eta = o$. If the intersection of the elements of \mathcal{O} is the point o , then, for some \mathcal{O} -hyperplane \mathcal{H} through o , we have $\eta \not\subseteq \mathcal{H}$. Then, $\mathcal{H} \cap \eta$ is a $(k-1)$ -dimensional \mathcal{O}_η -flat, which implies that \mathcal{O}_η is nonempty.

We next show that $\bigcap \mathcal{O}_\eta = o$. Let \mathcal{O}' be the set of \mathcal{O} -hyperplanes through o that do *not* contain η and $\{\mathcal{H} \cap \eta \mid \mathcal{H} \in \mathcal{O}'\}$ be the set of their intersections with η . Then, all elements of the latter set are $(k-1)$ -dimensional \mathcal{O}_η -planes through o and their intersection is o . We conclude that the intersection of all $(k-1)$ -dimensional \mathcal{O}_η -planes through o is exactly o ; that is, $\bigcap \mathcal{O}_\eta = o$. \square

The orientation set \mathcal{O}_η gives rise to lower-dimensional \mathcal{O}_η -halfspaces in the space η . These lower-dimensional \mathcal{O}_η -halfspaces are defined in the same way as \mathcal{O} -halfspaces, in terms of their intersection with \mathcal{O}_η -lines.

The following result, illustrated in Figure 15, will enable us to use induction on the dimension d in proofs about \mathcal{O} -halfspaces.

Lemma 4.5 *The intersection of a (directed) \mathcal{O} -halfspace with an \mathcal{O} -flat η is a (directed) \mathcal{O}_η -halfspace.*

Proof. Suppose that P is an \mathcal{O} -halfspace. For every \mathcal{O} -line $l \subseteq \eta$, we have $l \cap (P \cap \eta) = l \cap P$, which implies that the intersection of l with $P \cap \eta$ is empty, a ray, or a line. Thus, $P \cap \eta$ is an \mathcal{O}_η -halfspace. If P is a directed \mathcal{O} -halfspace, then the rays formed by the intersection of $P \cap \eta$ with parallel \mathcal{O} -lines in η always point in the same direction and, hence, $P \cap \eta$ is a directed \mathcal{O}_η -halfspace. \square

We use Lemma 4.5 to prove that, for orientation sets with the point-intersection property, the intersection of a directed \mathcal{O} -halfspace with an \mathcal{O} -flat is always connected.




Figure 16: Proof of Theorem 4.6.

Theorem 4.6 *If the orientation set \mathcal{O} has the point-intersection property, then the intersection of a directed \mathcal{O} -halfspace with every \mathcal{O} -flat is connected.*

Proof. We first prove, by induction on the dimension d , that all directed \mathcal{O} -halfspaces are connected. By Lemma 2.5, directed \mathcal{O} -halfplanes are connected, which provides us with an induction base.

We consider a directed \mathcal{O} -halfspace P in d dimensions and show that every two points $p, q \in P$ are in the same connected component. By Lemma 3.1, since \mathcal{O} has the point-intersection property, we may select some \mathcal{O} -line l and an \mathcal{O} -hyperplane \mathcal{H} that intersects l and does not contain it (see Figure 16). Since P is a directed \mathcal{O} -halfspace, there are two parallel-to- l rays, with endpoints p and q , that are contained in P and point in the same direction.

We select an \mathcal{O} -plane \mathcal{H}' , parallel to \mathcal{H} , that intersects the two rays with endpoints p and q and denote the intersection points by x and y , respectively. By Lemma 4.5, $P \cap \mathcal{H}'$ is a directed $\mathcal{O}_{\mathcal{H}'}$ -halfspace; that is, $P \cap \mathcal{H}'$ is a $(d - 1)$ -dimensional directed \mathcal{O} -halfspace. By Lemma 4.4, the lower-dimensional orientation set $\mathcal{O}_{\mathcal{H}'}$ has the point-intersection property. Therefore, by the induction hypothesis, $P \cap \mathcal{H}'$ is connected. We conclude that x and y are in the same component of P , which implies that p and q are also in the same component.

Finally, we note that the intersection of a directed \mathcal{O} -halfspace with every \mathcal{O} -flat η is a directed \mathcal{O}_η -halfspace (Lemma 4.5) and the lower-dimensional orientation set \mathcal{O}_η has the point-intersection property (Lemma 4.4). Therefore, the intersection of P with η is connected. \square

Theorem 4.6 cannot be extended to orientation sets that do not have the point-intersection property. For example, if there is only one \mathcal{O} -line through o , then the union of several lines parallel to this \mathcal{O} -line is a disconnected directed \mathcal{O} -halfspace.

The property of directed \mathcal{O} -halfspaces stated in Theorem 4.6 is called \mathcal{O} -connectedness.

Definition 4.3 (\mathcal{O} -Connectedness) *A closed set is \mathcal{O} -connected if its intersection with every \mathcal{O} -flat is connected.*

Note that \mathcal{O} -connected sets are connected, since the whole space \mathcal{R}^d is an \mathcal{O} -flat.

Directed \mathcal{O} -halfspaces are not the only \mathcal{O} -connected sets. For example, the set in Figure 17(b) is \mathcal{O} -connected, even though it is not an \mathcal{O} -halfspace. In fact, every standard convex set is \mathcal{O} -connected. As another example, the set in Figure 17(c) is \mathcal{O} -connected for




Figure 17: For the orthogonal-orientation set (a), the sets (b) and (c) are \mathcal{O} -connected, whereas the sets (d–f) are *not* \mathcal{O} -connected.

the orthogonal-orientation set in Figure 17(a). On the other hand, the set in Figure 17(d) is not \mathcal{O} -connected because it is disconnected, the set in Figure 17(e) is not \mathcal{O} -connected because its intersection with the dashed \mathcal{O} -line is disconnected, and the set in Figure 17(f) is not \mathcal{O} -connected because its intersection with the dashed \mathcal{O} -plane is disconnected. We have studied \mathcal{O} -connected sets as a part of our exploration of \mathcal{O} -convexity [2].

5 Boundary convexity

We explore the properties of the boundaries of \mathcal{O} -halfspaces and present analogs of the boundary-convexity property for standard halfspaces (see Section 1).

We begin by showing that, if an orientation set has the point-intersection property, then all points in the boundary of an \mathcal{O} -halfspace are “infinitely close” to its interior; that is, an \mathcal{O} -halfspace is the closure of its interior.

Lemma 5.1 *Let P be an \mathcal{O} -halfspace and P_{int} be the interior of P . If the orientation set \mathcal{O} has the point-intersection property, then $\text{Closure}(P_{\text{int}}) = P$.*

Proof. We show that $\text{Closure}(P_{\text{int}}) = P$ by demonstrating that, for every point p in the boundary of P and every distance ϵ , there is an interior point within distance ϵ of p . We use induction on the dimension d .

By Lemma 3.1, we may select some \mathcal{O} -line l and an \mathcal{O} -hyperplane \mathcal{H} that intersects l and does not contain it. Let \mathcal{H}' be the \mathcal{O} -hyperplane through p parallel to \mathcal{H} , let Q be the intersection of P and \mathcal{H}' , and let Q_{int} be the interior of Q in the $(d-1)$ -dimensional space \mathcal{H}' (rather than in \mathcal{R}^d). If $d > 2$, the set Q is a lower-dimensional \mathcal{O} -halfspace (Lemma 4.5) and, by the induction hypothesis, $\text{Closure}(Q_{\text{int}}) = Q$. On the other hand, if $d = 2$, then Q is empty, a ray, or a line; therefore, $\text{Closure}(Q_{\text{int}}) = Q$, which establishes an induction base.

We select an interior point of Q within distance $\epsilon/2$ of p and a $(d-1)$ -dimensional ball $B \subseteq Q$ centered at this point such that the radius of B is at most $\epsilon/2$ (see Figure 18a).

We assume that the line l is vertical and we divide the rays parallel to l into “upward” and “downward” rays. For every point $q \in B$, the upward or downward ray with endpoint q is contained in P . Let U be the set of all points of B such that the upward rays from them




Figure 18: Proof of Lemma 5.1.

are in P and D be the set of all points of B such that the downward rays from them are in P .

We consider two cases. First, suppose that the interior of U in the $(d - 1)$ -dimensional space \mathcal{H}' is nonempty, which means that there is a $(d - 1)$ -dimensional ball $B' \subseteq U$ (see Figure 18a). The union of the upward rays from all points of B' forms a “cylinder,” wholly contained in P . Some of the interior points of this cylinder are within distance ϵ of p . Clearly, the interior of the cylinder is in the interior of P ; therefore, some of P 's interior points are within ϵ of p .

Now suppose that the interior of U is empty. Then, every $(d - 1)$ -dimensional ball $B' \subseteq B$ contains a point of the set D , which implies that $\text{Closure}(D) = B$. Let \mathcal{H}'' be a plane parallel to \mathcal{H}' and located below \mathcal{H}' (see Figure 18b). The intersection of the downward rays from all points of D with \mathcal{H}'' forms a translation of D . This translation of D is in P and its closure is a $(d - 1)$ -dimensional ball, say B'' , which is a translation of B .

By the closed-set assumption (see Section 1), P is closed; therefore, B'' is in P . The union of the vertical segments joining B and B'' forms a cylinder, wholly contained in P . The interior of this cylinder is in P 's interior and some of the cylinder's interior points are within distance ϵ of p . \square

We call sets satisfying the property stated in Lemma 5.1 **interior-closed sets**: A set P is interior-closed if $\text{Closure}(P_{\text{int}}) = P$.

Our purpose in the rest of the section is to present \mathcal{O} -convexity analogs of the following boundary-convexity characterization of standard halfspaces.

Proposition 5.2 *An interior-closed set is a halfspace if and only if its boundary is a nonempty convex set.*

We first generalize the “if” part of this boundary-convexity characterization.

Lemma 5.3 *An interior-closed set with an \mathcal{O} -convex boundary is an \mathcal{O} -halfspace.*

Proof. Let P be an interior-closed set with an \mathcal{O} -convex boundary. We show that the intersection of P with every \mathcal{O} -line l is empty, a ray, or a line. The intersection of l with P 's boundary may be a line, a ray, a segment, a point, or empty. If this intersection is a line,




Figure 19: The eight cases considered in the proof of Lemma 5.3.



Figure 20: An \mathcal{O} -halfplane whose boundary is not \mathcal{O} -convex.

then $P \cap l$ is also a line (see Figure 19a). If the intersection of l with P 's boundary is a ray, then $P \cap l$ is a ray or a line (Figures 19b and 19c). If l does not intersect P 's boundary, then $P \cap l$ is empty or a line (Figures 19d and 19e).

Finally, suppose that the intersection of l with P 's boundary is a segment or point. We show, by contradiction, that, in this case, $P \cap l$ is a line or ray (Figures 19f and 19g). If not, then $P \cap l$ is a segment or point (see Figure 19h). We select points $p, q \in l$ that are outside of P , on different sides of the intersection. We consider equal-sized balls, B_p and B_q , that are centered at these points and do *not* intersect P . We next select a point x in the intersection of l and P 's boundary and consider the ball B_x , centered at x , of the same size as B_p and B_q . Since P is interior-closed, we may select a point $y \in B_x$ that is in the interior of P . We consider the line l_y through y parallel to l . This line intersects the balls B_p and B_q , which are outside of P . Since x is in the interior of P , we conclude that the intersection of the \mathcal{O} -line l_y with the boundary of P is not connected, contradicting the \mathcal{O} -convexity of P 's boundary. \square

The converse of Lemma 5.3 does not hold: the boundary of an \mathcal{O} -halfspace may not be \mathcal{O} -convex. In Figure 20, we show an \mathcal{O} -halfplane whose boundary is not \mathcal{O} -convex: the intersection of the boundary with the dashed \mathcal{O} -line is disconnected. We now present a necessary and sufficient characterization of \mathcal{O} -halfspaces in terms of their boundary.

Theorem 5.4 (Boundary of \mathcal{O} -halfspaces) *An interior-closed set P is an \mathcal{O} -halfspace if and only if, for every \mathcal{O} -line l , one of the following two conditions holds:*

1. *The intersection of l with the boundary of P is connected.*
2. *The intersection of l with the boundary of P consists of two disconnected rays and the segment of l between these rays is in P .*




Figure 21: Proof of Theorem 5.4.

Proof. Suppose that the two conditions hold. We demonstrated in the proof of Lemma 5.3 that, if the intersection of l with the boundary of P is connected (Condition 1), then $P \cap l$ is empty, a ray, or a line. On the other hand, if the intersection of l with P 's boundary satisfies Condition 2, then $P \cap l$ is a line. We conclude that the intersection of P with every \mathcal{O} -line is empty, a ray, or a line; thus, P is an \mathcal{O} -halfspace.

To prove the converse, suppose that P is an \mathcal{O} -halfspace and the intersection of an \mathcal{O} -line l with the boundary of P is *not* connected. We show that this intersection satisfies Condition 2.

Since the boundary is closed, we can select points $p, q \in l$ in P 's boundary such that all points of l between p and q are *not* in the boundary. Since the intersection of l with P is connected, the segment of l between p and q is in P .

We next show, by contradiction, that all points of l outside of this segment are in P 's boundary. Suppose that some point x is *not* in P 's boundary. Without loss of generality, we assume that x is to the left of p and that q is to the right of p (see Figure 21a).

If x is an interior point of P , then there is a ball $B_x \subseteq P$ centered at x (Figure 21a). Let $y \in l$ be some point between p and q , $B_y \subseteq P$ be a ball centered at y , and B_p be a ball, centered at p , that is no larger than B_x and B_y . Since p is in the boundary of P , we may choose a point $z \in B_p$ that is *not* in P . We consider the \mathcal{O} -line l_z through z parallel to l . This line intersects the balls B_x and B_y , which are in P . Since z is not in P , the intersection of l_z with P is disconnected, contradicting the assumption that P is an \mathcal{O} -halfspace.

If x is an exterior point of P , we may select a ball B_x centered at x that does not intersect P (see Figure 21b). Let $y \in l$ be some point between p and q , $B_y \subseteq P$ be a ball centered at y , and B_q be a ball, centered at q , that is no larger than B_x and B_y . Since q is in the boundary of P , we may select a point $z \in B_q$ that is *not* in P . We again consider the \mathcal{O} -line l_z through z parallel to l . This line intersects the balls B_x and B_y , which implies that the intersection of l_z with P is not empty, not a ray, and not a line, contradicting the assumption that P is an \mathcal{O} -halfspace. \square

Observe that, if the intersection of P 's boundary with a line l consists of two rays and the segment of the line between these rays is in P , then l is wholly contained in P . We use this observation to simplify Condition 2 in Theorem 5.4.

Corollary 5.5 *An interior-closed set P is an \mathcal{O} -halfspace if and only if, for every \mathcal{O} -line l , one of the following two conditions holds:*




Figure 22: An \mathcal{O} -halfspace that is *not* directed, even though its boundary is \mathcal{O} -connected.

1. *The intersection of l with the boundary of P is connected.*
2. *l is wholly contained in P .*

Proof. We immediately conclude from Theorem 5.4 that, if P is an \mathcal{O} -halfspace, then the two conditions hold. Suppose, conversely, that these conditions hold. If the intersection of l with P 's boundary is connected (Condition 1), then $P \cap l$ is empty, a ray, or a line (see the proof of Lemma 5.3). If l is contained in P (Condition 2), then $P \cap l$ is a line. Therefore, P is an \mathcal{O} -halfspace. \square

We have shown in Lemma 2.5 that the boundary of a directed \mathcal{O} -halfplane is always \mathcal{O} -convex. We now observe that the proof of Lemma 2.5 carries over to higher dimensions, which gives us the following property of directed \mathcal{O} -halfspaces.

Lemma 5.6 *The boundary of a directed \mathcal{O} -halfspace is \mathcal{O} -convex.*

If the orientation set has the point-intersection property, then the *boundary of a directed \mathcal{O} -halfspace is \mathcal{O} -connected*, which is a stronger property. We will prove this result in Section 6, using properties of the complement of a directed \mathcal{O} -halfspace. The converse of this result does not hold, as we demonstrate in the following example.

Example: An \mathcal{O} -halfspace, with an \mathcal{O} -connected boundary, that is not directed. We show such an \mathcal{O} -halfspace in Figure 22(b). This \mathcal{O} -halfspace consists of two rectangular polyhedral angles (quadrants), which touch each other along one of their faces. The boundary of the \mathcal{O} -halfspace is \mathcal{O} -connected for the orthogonal-orientation set in Figure 22(a). The \mathcal{O} -halfspace, however, is *not* directed, since the dotted rays (Figure 22c), formed by the intersection of this \mathcal{O} -halfspace with vertical \mathcal{O} -lines, point in opposite directions. \square

6 Complementation

We have observed in Section 1 that the closure of the complement of a standard halfspace is a halfspace. We now generalize this result to \mathcal{O} -halfspaces.

We call the closure of the complement of a set the **closed complement**. In general, the closed complement of an \mathcal{O} -halfspace may *not* be an \mathcal{O} -halfspace. For example, the closed




Figure 23: For (a) the orthogonal orientation set, the set (b) is an \mathcal{O} -halfspace, whereas (c) its closed complement is *not* an \mathcal{O} -halfplane.




Figure 24: Proofs of (a) Theorem 6.1 and (b) Theorem 6.2.

complement of the \mathcal{O} -halfplane in Figure 23(b) is not an \mathcal{O} -halfplane, since its intersection with the dashed \mathcal{O} -line (Figure 23c) is disconnected. We state a necessary and sufficient condition under which the closed complement of an \mathcal{O} -halfspace is an \mathcal{O} -halfspace.

Theorem 6.1 (Complements of \mathcal{O} -halfspaces) *The closed complement of an \mathcal{O} -halfspace P is an \mathcal{O} -halfspace if and only if the boundary of P is \mathcal{O} -convex.*

Proof. We denote the closed complement of P by Q . Note that Q is interior-closed and the boundary of Q is the same as the boundary of P . If the boundary of P is \mathcal{O} -convex, then Q is an interior-closed set with an \mathcal{O} -convex boundary, which implies that Q is an \mathcal{O} -halfspace (Lemma 5.3).

Suppose, conversely, that Q is an \mathcal{O} -halfspace. We show, by contradiction, that the boundary of P is \mathcal{O} -convex. If the boundary is not \mathcal{O} -convex, then there are points p , x , and q on some \mathcal{O} -line l , such that p and q are in the boundary, whereas x , located between p and q , is *not* in the boundary (see Figure 24a). Note that p and q belong both to P and to Q , whereas x is either in P 's interior or in Q 's interior. If x is in P , then the intersection of P with the \mathcal{O} -line l is disconnected and, hence, P is not an \mathcal{O} -halfspace. Similarly, if x is in Q , then Q is not an \mathcal{O} -halfspace. \square

We next generalize the complementation property of halfspaces to directed \mathcal{O} -halfspaces.

Theorem 6.2 (Complements of directed \mathcal{O} -halfspaces) *The closed complement of a directed \mathcal{O} -halfspace is a directed \mathcal{O} -halfspace.*

Proof. The boundary of a directed \mathcal{O} -halfspace P is \mathcal{O} -convex (Lemma 5.6). Therefore, by Theorem 6.1, the closed complement of P is an \mathcal{O} -halfspace.

We show, by contradiction, that the closed complement of P is directed. Suppose that the intersection of the closed complement with two parallel \mathcal{O} -lines, l_1 and l_2 , forms rays that point in opposite directions. We denote the endpoints of these rays by p and q , respectively. In Figure 24(b), we show the rays by solid lines. Note that p and q are in the boundary of P and the dashed parts of the lines are in the interior of P .

We select a point in the dashed part of l_1 and a ball $B \subseteq P$ centered at this point. Let B_p be the same-sized ball centered at p . Since p is in the boundary of P , we may select a point $x \in B_p$ that is *not* in P . We consider the \mathcal{O} -line l_x through x parallel to l_1 . The intersection of l_x with P is empty, a ray, or a line. Since x is not in P and, on the other hand, l_x intersects the ball $B \subseteq P$, we conclude that the intersection of l_x with P is a ray pointing “to the right.”

Using a similar construction with l_2 , we get an \mathcal{O} -line whose intersection with P is a ray pointing “to the left” (see Figure 24b), which means that the \mathcal{O} -halfspace P is *not* directed, giving a contradiction. \square

We use the complementation property of directed \mathcal{O} -halfspaces to demonstrate that, for orientation sets with the point-intersection property, the boundary of a directed \mathcal{O} -halfspace is \mathcal{O} -connected.

Theorem 6.3 (Boundaries of directed \mathcal{O} -halfspaces) *If the orientation set \mathcal{O} has the point-intersection property, then the boundary of every directed \mathcal{O} -halfspace is \mathcal{O} -connected.*

Proof. To demonstrate that the boundary of a directed \mathcal{O} -halfspace P is \mathcal{O} -connected, we first show that the boundary is connected and then use this result to demonstrate that the intersection of the boundary with every \mathcal{O} -flat η is also connected.

We establish the connectedness of P 's boundary by contradiction. Since the orientation set \mathcal{O} has the point-intersection property, P is connected (Theorem 4.6). Therefore, if the boundary of P is disconnected, then the closed complement of P is also disconnected. On the other hand, the closed complement of P is a directed \mathcal{O} -halfspace (Theorem 6.2), contradicting the connectedness of directed \mathcal{O} -halfspaces.

We next show that the intersection of P 's boundary, denoted by $\text{Bdry}(P)$, with every \mathcal{O} -flat η is connected. Let Q be the intersection of P with η and $\text{Bdry}(Q)$ be the boundary of Q in the lower-dimensional space η (rather than in the whole space \mathcal{R}^d). Clearly, $\text{Bdry}(Q) \subseteq \text{Bdry}(P) \cap \eta$. The set Q is a directed \mathcal{O}_η -halfspace (Lemma 4.5) and the orientation set \mathcal{O}_η has the point-intersection property (Lemma 4.4). Therefore, by the first part of the proof, $\text{Bdry}(Q)$ is connected.

Suppose that $\text{Bdry}(P) \cap \eta$ is *not* connected. Then, $\text{Bdry}(P) \cap \eta$ contains a component disconnected from the boundary of Q (see Figure 25, where this component is shown by the shaded region). This component is contained in Q ; therefore, it is surrounded in the flat η by interior points of P (the interior points are shown by the dashed region in Figure 25).

We now consider the intersection of the closed complement of P with η . This intersection contains $\text{Bdry}(Q)$ and the component of the $\text{Bdry}(P) \cap \eta$ disconnected from $\text{Bdry}(Q)$. The




Figure 25: Proof of Theorem 6.3.

intersection does *not* contain any interior points of P . Therefore, the intersection of the closed complement of P with η is disconnected. On the other hand, the closed complement of P is a directed \mathcal{O} -halfspace (Theorem 6.2), which implies that the intersection of the closed complement of P with the \mathcal{O} -flat η is connected (Theorem 4.6), yielding a contradiction. \square

7 Concluding remarks

We investigated \mathcal{O} -halfspaces and demonstrated that the properties of \mathcal{O} -halfspaces are somewhat similar to those of standard halfspaces. In defining \mathcal{O} -halfspaces, we lose connectedness; therefore, we introduced directed \mathcal{O} -halfspaces, which are always connected.

We established many basic properties of \mathcal{O} -halfspaces and directed \mathcal{O} -halfspaces; however, we are unable to prove the following two appealing, and tantalizing, conjectures, which we leave as open problems. An \mathcal{O} -connected set is the intersection of the directed \mathcal{O} -halfspaces that contain it. A connected \mathcal{O} -halfspace is the intersection of the directed \mathcal{O} -halfspaces that contain it.

The work presented here extends our previous studies of strong \mathcal{O} -convex sets, \mathcal{O} -convex sets, and \mathcal{O} -connected sets [3; 2]. We plan to further develop the theory of \mathcal{O} -convexity and also study the computational properties of \mathcal{O} -convex sets and \mathcal{O} -halfspaces.

References

- [1] C. K. Bruckner and J. B. Bruckner. On L_n -sets, the Hausdorff metric, and connectedness. *Proceedings of the American Mathematical Society*, 13:765–767, 1962.
- [2] Eugene Fink and Derick Wood. Fundamentals of restricted-orientation convexity. Unpublished manuscript, 1995.
- [3] Eugene Fink and Derick Wood. Strong restricted-orientation convexity. Unpublished manuscript, 1995.
- [4] Eugene Fink and Derick Wood. Three-dimensional strong convexity and visibility. In *Proceedings of the Vision Geometry IV Conference*, 1995.
- [5] Branko Grünbaum, Victor Klee, M. A. Perles, and G. C. Shephard. *Convex Polytopes*. John Wiley & Sons, New York, NY, 1967.

- [6] Ralf Hartmut Güting. Stabbing C -oriented polygons. *Information Processing Letters*, 16:35–40, 1983.
- [7] Victor Klee. What is a convex set? *American Mathematical Monthly*, 78:616–631, 1971.
- [8] Witold Lipski and Christos H. Papadimitriou. A fast algorithm for testing for safety and detecting deadlock in locked transaction systems. *Journal of Algorithms*, 2:211–226, 1981.
- [9] D. Y. Montuno and Alain Fournier. Finding the x - y convex hull of a set of x - y polygons. Technical Report 148, University of Toronto, Toronto, Ontario, 1982.
- [10] Tina M. Nicholl, D. T. Lee, Y. Z. Liao, and C. K. Wong. Constructing the X-Y convex hull of a set of X-Y polygons. *BIT*, 23:456–471, 1983.
- [11] Thomas Ottmann, Eljas Soisalon-Soininen, and Derick Wood. On the definition and computation of rectilinear convex hulls. *Information Sciences*, 33:157–171, 1984.
- [12] Franco P. Preparata and Michael I. Shamos. *Computational Geometry*. Springer-Verlag, New York, NY, 1985.
- [13] Gregory J. E. Rawlins. *Explorations in Restricted-Orientation Geometry*. PhD thesis, University of Waterloo, 1987. Technical Report CS-89-48.
- [14] Gregory J. E. Rawlins and Derick Wood. Optimal computation of finitely oriented convex hulls. *Information and Computation*, 72:150–166, 1987.
- [15] Gregory J. E. Rawlins and Derick Wood. Ortho-convexity and its generalizations. In Godfried T. Toussaint, editor, *Computational Morphology*, pages 137–152. Elsevier Science Publishers B. V., North-Holland, 1988.
- [16] Gregory J. E. Rawlins and Derick Wood. Restricted-orientation convex sets. *Information Sciences*, 54:263–281, 1991.
- [17] Sven Schuierer. *On Generalized Visibility*. PhD thesis, Universität Freiburg, Germany, 1991.
- [18] Eljas Soisalon-Soininen and Derick Wood. Optimal algorithms to compute the closure of a set of iso-rectangles. *Journal of Algorithms*, 5:199–214, 1984.
- [19] F. A. Valentine. Local convexity and L_n -sets. *Proceedings of the American Mathematical Society*, 16:1305–1310, 1965.
- [20] Peter Widmayer, Y. F. Wu, and C. K. Wong. On some distance problems in fixed orientations. *SIAM Journal on Computing*, 16:728–746, 1987.