

Fundamentals of Restricted-Orientation Convexity*

Eugene Fink[†]

Derick Wood[‡]

Abstract

A *restricted-orientation convex set*, also called an \mathcal{O} -convex set, is a set of points whose intersection with lines from some fixed set is empty or connected. The notion of \mathcal{O} -convexity generalizes standard convexity and orthogonal convexity.

We explore some of the basic properties of \mathcal{O} -convex sets in two and higher dimensions. We also study *\mathcal{O} -connected sets*, which are a subclass of \mathcal{O} -convex sets, with several special properties. We introduce and investigate restricted-orientation analogs of lines, flats, and hyperplanes, and characterize \mathcal{O} -convex and \mathcal{O} -connected sets in terms of their intersections with hyperplanes. We then explore properties of \mathcal{O} -connected curves; in particular, we show when replacing a segment of an \mathcal{O} -connected curve with a new curvilinear segment yields an \mathcal{O} -connected curve and when the concatenation of several curvilinear segments forms an \mathcal{O} -connected segment. We use these results to characterize an \mathcal{O} -connected set in terms of \mathcal{O} -connected segments, joining pairs of its points, that are wholly contained in the set.

We also identify some of the major properties of standard convex sets that hold for \mathcal{O} -convexity. In particular, we demonstrate that the intersection of a collection of \mathcal{O} -convex sets is an \mathcal{O} -convex set, every \mathcal{O} -connected curvilinear segment is a segment of some \mathcal{O} -connected curve, and, for every two points of an \mathcal{O} -convex set, there is an \mathcal{O} -convex segment joining them that is wholly contained in the set.

1 Introduction

The study of convex sets is a branch of geometry that has numerous connections with other areas of mathematics, including analysis, linear algebra, statistics, number theory, and combinatorics [5, 10, 9]. The importance of convexity theory stems from the fact that convex sets arise frequently in many areas of mathematics and are often amenable to rather elementary reasoning. The concept of convexity serves to unify a wide range of mathematical phenomena.

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

[†]School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, U. S. A. E-mail: eugene@cs.cmu.edu. WWW: <http://www.cs.cmu.edu/~eugene>.

[‡]Department of Computer Science, Hong Kong University of Science & Technology, Clear Water Bay, Kowloon, Hong Kong. E-mail: dwood@cs.ust.hk. WWW: <http://www.cs.ust.hk/faculty/dwood>.

The application of convexity theory to practical problems [15] led to the exploration of nontraditional notions of convexity, such as orthogonal convexity [12, 13, 14], finitely oriented convexity [7, 24, 17], restricted-orientation convexity [16, 20, 21], NESW convexity [11, 22], and link convexity [1, 23, 21]. These nontraditional convexities have been used in pixel graphics, locked transaction systems, VLSI design, motion planning, and other areas.

Güting introduced the notion of *restricted orientations* in two dimensions in his doctoral thesis [7]. He investigated computational properties of polygons whose edges are parallel to the elements of some fixed finite set of lines, called an *orientation set* [6, 8].

Rawlins applied restricted orientations in his definition of two new types of generalized convexity, which he called *\mathcal{O} -convexity* and *strong \mathcal{O} -convexity* [16]. *\mathcal{O} -convexity*, also called *restricted-orientation convexity*, was defined in terms of the intersection of a geometric object with lines parallel to the elements of a fixed orientation set \mathcal{O} (see Section 2). Rawlins and Wood showed that the notion of *\mathcal{O} -convexity* generalizes standard convexity and orthogonal convexity, and that the properties of *\mathcal{O} -convex sets* are similar to the properties of standard convex sets [18, 20]. Schuierer continued their exploration and presented an extensive study of geometric and computational properties of *\mathcal{O} -convex sets* in his doctoral thesis [21].

In this paper, the second in a series [3, 2], we demonstrate that the notion of *\mathcal{O} -convexity* can be extended to higher dimensions. This extension yields a generalization of planar *\mathcal{O} -convexity* and standard higher-dimensional convexity. In the first paper of the series [3], we described a generalization of strong *\mathcal{O} -convexity* [16] to higher dimensions and explored the properties of strongly *\mathcal{O} -convex sets* [4]. In the third paper [2], we describe a restricted-orientation generalization of halfspaces and investigate properties of these generalized halfspaces.

We restrict our attention to the exploration of closed sets. We conjecture that most of the results hold for nonclosed sets as well; 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 present basic properties of planar *\mathcal{O} -convex sets*. In Section 3, we generalize *\mathcal{O} -convexity* and its basic properties to higher dimensions. In Section 4, we describe *\mathcal{O} -connected sets*, which are a subclass of *\mathcal{O} -convex sets*, with several special properties. In Section 5, we explore properties of *\mathcal{O} -connected curves*. In Section 6, we present visibility results for *\mathcal{O} -convex* and *\mathcal{O} -connected sets*. Finally, we conclude, in Section 7, with a discussion of open problems and of future work.

2 *\mathcal{O} -convexity* in two dimensions

We begin by reviewing the notion of *\mathcal{O} -convexity* in two dimensions [16] and presenting some of the basic properties of planar *\mathcal{O} -convex sets*. Rawlins introduced this notion as a generalization of orthogonal convexity [14] and standard convexity. He defined *\mathcal{O} -convex*

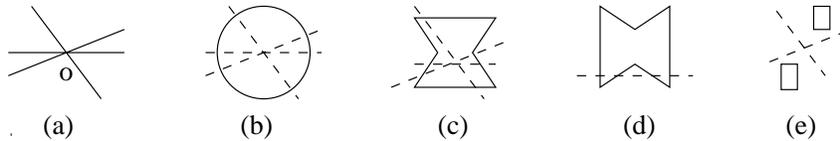


Figure 1: Planar \mathcal{O} -convexity.

sets in terms of their intersection with straight lines, by analogy with one of the definitions of standard convexity. Rawlins, Wood, and Schuierer explored properties of \mathcal{O} -convex sets in two dimensions and showed that many properties of these sets are similar to the properties of standard convex sets [18, 20, 21].

We can describe convex sets through their intersections with straight lines: a set of points is convex if its intersection with every line is empty or connected. We define \mathcal{O} -convexity by considering the intersection of a set of points with lines from a *certain set* (rather than all lines). In other words, we select some collection of lines and say that a set is \mathcal{O} -convex if its intersection with every line from this collection is empty or connected.

To define this restricted collection of lines, we first introduce the notion of an orientation set. An *orientation set* \mathcal{O} is a (finite or infinite) closed set of lines through some fixed point o . An example of a finite orientation set is shown in Figure 1(a). A straight line parallel to one of the lines of \mathcal{O} is called an \mathcal{O} -line. For example, the dashed lines in Figures 1(b)–(e) are \mathcal{O} -lines. We use the collection of all \mathcal{O} -lines in defining \mathcal{O} -convexity.

Definition 1 (\mathcal{O} -Convexity) *A closed set is \mathcal{O} -convex if its intersection with every \mathcal{O} -line is empty or connected.*

For the orientation set in Figure 1(a), the sets shown in Figures 1(b) and 1(c) are \mathcal{O} -convex (some \mathcal{O} -lines intersecting these sets are shown by dashed lines). On the other hand, the set in Figure 1(d) is not \mathcal{O} -convex, since its intersection with the dashed \mathcal{O} -line is disconnected. Note that the set in Figure 1(d) is a rotation of the set in Figure 1(c). This observation shows that rotations may not preserve \mathcal{O} -convexity.

Unlike standard convex sets, \mathcal{O} -convex sets may be disconnected. In Figure 1(e), we show a disconnected \mathcal{O} -convex set, which consists of two rectangles.

We now prove some basic properties of planar \mathcal{O} -convex sets [20].

Lemma 1

1. *Every translation of an \mathcal{O} -convex set is \mathcal{O} -convex.*
2. *Every standard convex set is \mathcal{O} -convex.*
3. **(Intersection)** *If C is a collection of \mathcal{O} -convex sets, then the intersection $\bigcap C$ of these sets is also an \mathcal{O} -convex set.*
4. *A disconnected set is \mathcal{O} -convex if and only if every connected component of the set is \mathcal{O} -convex and no \mathcal{O} -line intersects two components.*
5. *If \mathcal{O} is not empty, then every connected \mathcal{O} -convex set is simply connected.*

Proof.

(1) By definition, every translation of an \mathcal{O} -line is an \mathcal{O} -line. Therefore, if the intersection of a set with every \mathcal{O} -line is empty or connected, then the same holds for every translation of the set.

(2) The intersection of a convex set with every straight line is empty or connected. In particular, the intersection of a convex set with every \mathcal{O} -line is empty or connected. Therefore, a convex set is \mathcal{O} -convex.

(3) If C is a collection of \mathcal{O} -convex sets, then, for each \mathcal{O} -line l , the intersection of every element of C with l is empty or connected; therefore, the intersection of $\bigcap C$ with l is also empty or connected. We conclude that the intersection of $\bigcap C$ with every \mathcal{O} -line is empty or connected, which means that $\bigcap C$ is \mathcal{O} -convex.

(4) If a set P is the union of disjoint \mathcal{O} -convex components and no \mathcal{O} -line intersects two components, then the intersection of P with every \mathcal{O} -line is empty or connected; therefore, P is \mathcal{O} -convex.

If one of the components of the set P is not \mathcal{O} -convex, then the intersection of this component with some \mathcal{O} -line is disconnected. The intersection of P with this \mathcal{O} -line is also disconnected; therefore, P is not \mathcal{O} -convex. Finally, if some \mathcal{O} -line intersects two (or more) components, then the intersection of P with this line is disconnected; therefore, we again conclude that P is not \mathcal{O} -convex.

(5) If a set P is connected and not simply connected, then P has a hole and there is some \mathcal{O} -line that cuts the hole. The intersection of P with this line is disconnected; thus, P is not \mathcal{O} -convex. \square

3 \mathcal{O} -convexity in higher dimensions

We now extend the notion of \mathcal{O} -convexity to d -dimensional space \mathcal{R}^d . We assume that space \mathcal{R}^d is fixed; however, all our 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 d -dimensional \mathcal{O} -convex sets in terms of their intersections with \mathcal{O} -lines. We then explore basic properties of \mathcal{O} -convexity in \mathcal{R}^d .

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, a hyperplane is 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 2 (Orientation set and \mathcal{O} -hyperplanes) *An orientation set \mathcal{O} in d dimensions is a closed set of hyperplanes through a fixed point o . A hyperplane parallel to one of the elements of \mathcal{O} is called 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 2, we give two examples of finite orientation sets in three dimensions. The first set contains three mutually orthogonal planes; we call it an *orthogonal-orientation set*. The second orientation set consists of four planes.

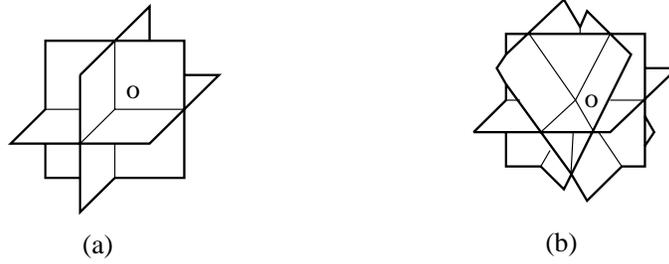


Figure 2: Orientation sets.

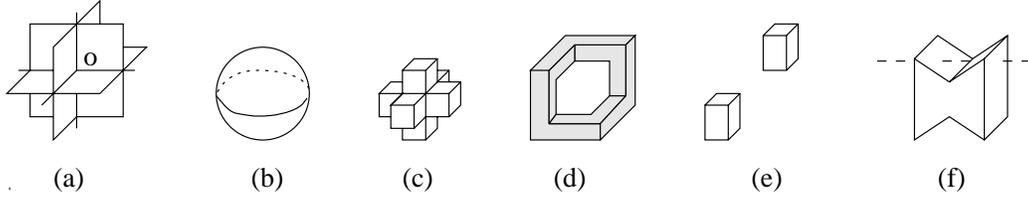


Figure 3: \mathcal{O} -convexity in three dimensions.

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

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 2(b) form six different lines through o and every \mathcal{O} -line for this orientation set is parallel to one of these six lines.

We define \mathcal{O} -convexity in higher dimensions in the same way as in two dimensions.

\mathcal{O} -Convexity A closed set in \mathcal{R}^d is \mathcal{O} -convex if its intersection with every \mathcal{O} -line is empty or connected.

For example, the sets in Figures 3(b)–(e) are \mathcal{O} -convex for the orthogonal-orientation set shown in Figure 3(a). On the other hand, the set in Figure 3(f) is not \mathcal{O} -convex, because its intersection with the dashed \mathcal{O} -line is disconnected.

Let us now recall the properties of planar \mathcal{O} -convexity given in Section 2 (see Lemma 1). We can readily generalize Properties 1–4 to higher dimensions: these properties hold in \mathcal{R}^d and their proofs are the same as the proofs in \mathcal{R}^2 . The most important of them is Property 3, which is a generalization of the intersection property of standard convex sets: the intersection of a collection of \mathcal{O} -convex sets is \mathcal{O} -convex.

Property 5 does not hold in \mathcal{R}^d : a connected \mathcal{O} -convex set may not be simply connected. For example, the set in Figure 3(d) is \mathcal{O} -convex with respect to the orthogonal-orientation set and it is not simply connected.

We next characterize higher-dimensional \mathcal{O} -convex sets in terms of their intersections with \mathcal{O} -hyperplanes.

Lemma 2 *A set is \mathcal{O} -convex if and only if its intersection with every \mathcal{O} -hyperplane is \mathcal{O} -convex.*

Proof. All \mathcal{O} -hyperplanes are convex and, hence, they are \mathcal{O} -convex. Since the intersection of two \mathcal{O} -convex sets is \mathcal{O} -convex, we conclude that the intersection of an \mathcal{O} -convex set with every \mathcal{O} -hyperplane is \mathcal{O} -convex.

Suppose, conversely, that the intersection of a set P with every \mathcal{O} -hyperplane is \mathcal{O} -convex. To demonstrate that the intersection of P with every \mathcal{O} -line l is empty or connected, we choose some \mathcal{O} -hyperplane \mathcal{H} that contains l . Since $P \cap \mathcal{H}$ is \mathcal{O} -convex, the intersection of $P \cap \mathcal{H}$ with l is empty or connected. We now note that $P \cap \mathcal{H} \cap l = P \cap l$, which implies that the intersection of P with l is also empty or connected. \square

4 \mathcal{O} -Connectedness

We have seen that \mathcal{O} -convex sets may not be connected (see Figure 3e), whereas standard convex sets are always connected. We now describe a subclass of \mathcal{O} -convex sets that has the connectedness property: all sets of this subclass are connected, just like standard convex sets. We define these sets in terms of their intersection with flats formed by the intersections of \mathcal{O} -hyperplanes.

A *flat*, also known as an *affine variety*, in d dimensions is a subset of \mathcal{R}^d that is itself a lower-dimensional space. For example, points, straight 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. Two flats are *parallel* if they are translations of each other (note that parallel flats are of the same dimension).

We now define \mathcal{O} -flats.

Definition 3 (\mathcal{O} -flats) *A flat formed by the intersection of several \mathcal{O} -hyperplanes is 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 parallel to one of the hyperplanes of the orientation set \mathcal{O} , every \mathcal{O} -flat is parallel to some flat formed by the intersection of several elements of \mathcal{O} .

For example, the orthogonal-orientation set in three dimensions (Figure 2a) 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 result readily follows from the definition of \mathcal{O} -flats.

Lemma 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 define \mathcal{O} -connected sets in terms of path-connectedness of their intersection with \mathcal{O} -flats. A set is *path-connected* if every two points of the set can be connected by a path that is wholly contained in the set. (This property is stronger than usual connectedness.)

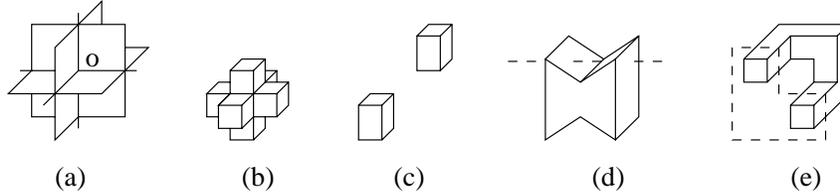


Figure 4: \mathcal{O} -connectedness in three dimensions.

Definition 4 (\mathcal{O} -Connectedness) *A closed set is \mathcal{O} -connected if its intersection with every \mathcal{O} -flat is empty or path-connected; that is, every two points of the intersection can be connected by a path that is wholly contained in the intersection.*

For example, the set in Figure 4(b) is \mathcal{O} -connected for the orthogonal-orientation set shown in Figure 4(a). On the other hand, the set in Figure 4(c) is not \mathcal{O} -connected because it is disconnected, the set in Figure 4(d) is not \mathcal{O} -connected because its intersection with the dashed \mathcal{O} -line is disconnected, and the set in Figure 4(e) is not \mathcal{O} -connected because its intersection with the dashed \mathcal{O} -plane is disconnected.

In two dimensions, every path-connected \mathcal{O} -convex set is \mathcal{O} -connected, since the intersections of a path-connected \mathcal{O} -convex set with \mathcal{O} -lines and with the whole plane \mathcal{R}^2 are path-connected. In higher dimensions, a path-connected \mathcal{O} -convex set may not be \mathcal{O} -connected. For example, the set in Figure 4(e) is a path-connected \mathcal{O} -convex set that is not \mathcal{O} -connected.

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

Lemma 4

1. *Every translation of an \mathcal{O} -connected set is \mathcal{O} -connected.*
2. *Every convex set is \mathcal{O} -connected and every \mathcal{O} -connected set is \mathcal{O} -convex.*

The intersection of \mathcal{O} -connected sets may *not* be \mathcal{O} -connected. For example, the intersection of the set in Figure 4(b) with some straight lines is disconnected, even though this set and all straight lines are \mathcal{O} -connected. Because of this “drawback” of \mathcal{O} -connectedness, we do not consider it a “true” generalization of convexity.

We can characterize \mathcal{O} -connected sets in terms of their intersections with \mathcal{O} -hyperplanes, much in the same way as we characterized \mathcal{O} -convex sets (see Lemma 2).

Theorem 5 *A set is \mathcal{O} -connected if and only if it is path-connected and its intersection with every \mathcal{O} -hyperplane is empty or \mathcal{O} -connected.*

Proof. Suppose that P is an \mathcal{O} -connected set. We show that P ’s intersection with every \mathcal{O} -hyperplane \mathcal{H} is \mathcal{O} -connected by demonstrating that, for every \mathcal{O} -flat η , the intersection of $P \cap \mathcal{H}$ with η is empty or path-connected. Since $\mathcal{H} \cap \eta$ is empty or an \mathcal{O} -flat (Lemma 3) and P is \mathcal{O} -connected, the intersection of $\mathcal{H} \cap \eta$ with P is empty or path-connected. We now note that this intersection is identical to the intersection of $P \cap \mathcal{H}$ with η ; therefore, the intersection of $P \cap \mathcal{H}$ with η is empty or path-connected.

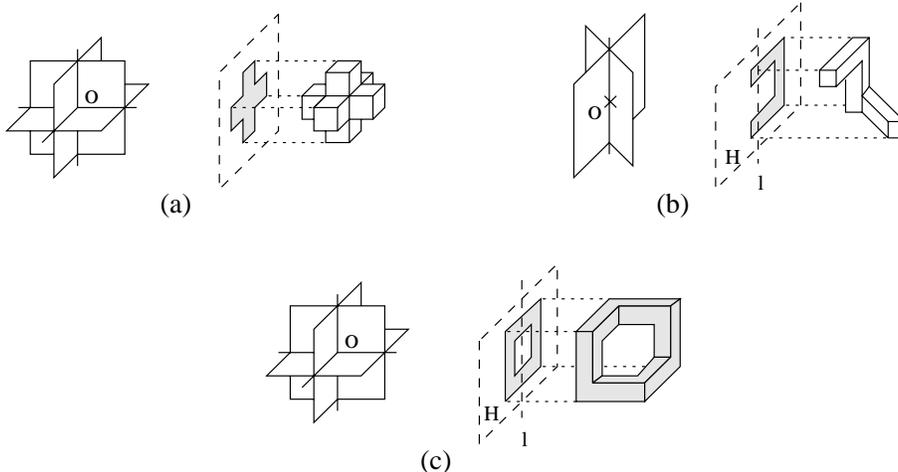


Figure 5: Orthogonal projections of \mathcal{O} -connected and \mathcal{O} -convex sets onto \mathcal{O} -planes.

Suppose, conversely, that P is a path-connected set and its intersection with every \mathcal{O} -hyperplane is \mathcal{O} -connected. To demonstrate that the intersection of P with every \mathcal{O} -flat η is path-connected, we choose some \mathcal{O} -hyperplane \mathcal{H} that contains η . Since $P \cap \mathcal{H}$ is \mathcal{O} -connected, the intersection of $P \cap \mathcal{H}$ with η is empty or path-connected. We next note that $P \cap \mathcal{H} \cap \eta = P \cap \eta$; therefore, the intersection of P with η is empty or path-connected. \square

If the orientation set \mathcal{O} is composed of mutually orthogonal hyperplanes, \mathcal{O} -connected sets have one more interesting property: an orthogonal projection of an \mathcal{O} -connected set onto an \mathcal{O} -flat is \mathcal{O} -connected. We illustrate this property in Figure 5(a), where a three-dimensional \mathcal{O} -connected cross is projected into a planar \mathcal{O} -connected cross.

Theorem 6 *If an orientation set \mathcal{O} is composed of mutually orthogonal hyperplanes, then the orthogonal projection of every \mathcal{O} -connected set onto every \mathcal{O} -flat is \mathcal{O} -connected.*

Proof. Let P be an \mathcal{O} -connected set, η be an \mathcal{O} -flat, and P_η be the orthogonal projection of P onto η . Note that P_η is path-connected, because it is a projection of a path-connected set. We now show that, for every \mathcal{O} -flat $\eta_1 \neq \mathcal{R}^d$, the intersection of P_η and η_1 is empty or path-connected.

The \mathcal{O} -flat η_1 is the intersection of several \mathcal{O} -hyperplanes, $\eta_1 = \mathcal{H}_1 \cap \mathcal{H}_2 \cap \dots \cap \mathcal{H}_n$. If one of these hyperplanes does not intersect η , then the intersection of P_η and η_1 is empty. If all these hyperplanes contain η , then $\eta \subseteq \eta_1$. In this case, the intersection of P_η and η_1 is the set P_η itself; therefore, this intersection is path-connected.

Finally, we consider the case when the hyperplanes $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_n$ all intersect η and some of them, say $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$, do not contain η . Since the orientation set \mathcal{O} is composed of mutually orthogonal hyperplanes, the \mathcal{O} -hyperplanes $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$ are all orthogonal to η (see Figure 6). We consider the \mathcal{O} -flat $\eta_2 = \mathcal{H}_1 \cap \mathcal{H}_2 \cap \dots \cap \mathcal{H}_k$. We note that $P_\eta \cap \eta_1 = P_\eta \cap \eta_2 \cap \eta$; therefore, $P_\eta \cap \eta_1$ is the projection of $P \cap \eta_2$ onto η (Figure 6). Since P is \mathcal{O} -connected, the intersection of P with the \mathcal{O} -flat η_2 is empty or path-connected. Since $P_\eta \cap \eta_1$ is the projection of $P \cap \eta_2$, we conclude that the intersection of P_η and η_1 is also empty or path-

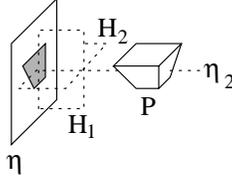


Figure 6: Proof of Theorem 6.

connected. □

If elements of the orientation set \mathcal{O} are not mutually orthogonal, the projection of an \mathcal{O} -connected set onto an \mathcal{O} -flat may not be \mathcal{O} -connected. We provide an example of such a situation in Figure 5(b), where the orientation set contains two planes. The set shown in Figure 5(b) is \mathcal{O} -connected; however, its projection onto the \mathcal{O} -plane H is not \mathcal{O} -connected, since the intersection of this projection with the \mathcal{O} -line l is disconnected.

The orthogonal projection of an \mathcal{O} -convex set onto an \mathcal{O} -flat may not be \mathcal{O} -convex, even for the orthogonal-orientation set. For example, the projection of the \mathcal{O} -convex set in Figure 5(c) onto the \mathcal{O} -plane H is not \mathcal{O} -convex, since the intersection of this projection with the \mathcal{O} -line l is disconnected.

5 \mathcal{O} -connected curves

We now study properties of \mathcal{O} -connected curves and curvilinear segments, and generalize some properties of planar \mathcal{O} -connected curves [20, 21] to higher dimensions.

Definition 5 (Curves and their segments) *A curve in d dimensions is the image of a continuous mapping from a straight line into \mathcal{R}^d . A segment of the curve is the image of a segment of the line.*

We restrict our attention to the exploration of simple \mathcal{O} -connected curves, where a simple curve is defined as follows.

Definition 6 (Simple curves) *A curve c is simple if, for every two points p and q of c , the shortest path from p to q that is wholly contained in c is a segment of c .*

Informally, this definition says that the shortest way to reach p from q while remaining in c is to follow c . Self-intersecting curves are *not* simple: if p and q are points on the different sides of a loop, the shortest path from p to q does not traverse the loop. Some unusual curves are not simple even though they are not self-intersecting. For example, a Peano curve that covers all points of a unit square is not simple even though it may not be self-intersecting.

Note that every straight line is an \mathcal{O} -connected curve. In Figure 7, we show two examples of more complex \mathcal{O} -connected curves, which run along the edges of the dotted cubes.

We now explore basic properties of simple \mathcal{O} -connected curves. We begin by characterizing \mathcal{O} -connected curves and curvilinear segments in terms of their intersection with \mathcal{O} -hyperplanes.

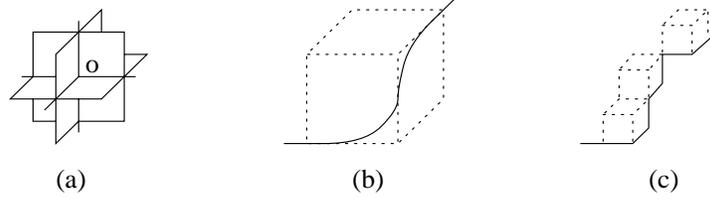


Figure 7: \mathcal{O} -connected curves.

Lemma 7 *A simple curve (curvilinear segment) is \mathcal{O} -connected if and only if its intersection with every \mathcal{O} -hyperplane is empty or connected.*

Proof. The “only if” part directly follows from the definition of \mathcal{O} -connected sets. Suppose, conversely, that the intersection of a simple curve (curvilinear segment) c with every \mathcal{O} -hyperplane is connected. We show that c is \mathcal{O} -connected by demonstrating that, for every \mathcal{O} -flat η and every two points $p, q \in c \cap \eta$, the segment $c[p, q]$ of the curve c is wholly in η , which implies that the intersection of c with every \mathcal{O} -flat η is path-connected.

If η is an \mathcal{O} -flat, then η is the intersection of several \mathcal{O} -hyperplanes. For each of these hyperplanes, its intersection with the curve c is connected. Since c is simple, this observation implies that $c[p, q]$ is wholly contained in all these hyperplanes; thus, $c[p, q]$ is wholly in η . \square

Observe that Lemma 7 holds only for simple curves. If a curve is not simple, it may not be \mathcal{O} -connected even if its intersection with every \mathcal{O} -hyperplane is connected. For example, consider a Peano curve in \mathcal{R}^3 that covers all points of a ball’s boundary. The intersection of this curve with every plane is empty or connected (a circle); however, the curve is not \mathcal{O} -connected, since its intersections with \mathcal{O} -lines through the center of the ball are disconnected.

We next show that a segment of an \mathcal{O} -connected curve is always \mathcal{O} -connected and, conversely, every \mathcal{O} -connected curvilinear segment can be extended to an \mathcal{O} -connected curve.

Lemma 8 (Segment extension)

1. *For every simple \mathcal{O} -connected curve c and every two points p and q of c , the segment $c[p, q]$ of the curve c is \mathcal{O} -connected.*
2. *For every simple \mathcal{O} -connected segment $c[p, q]$, there is a simple \mathcal{O} -connected curve c such that $c[p, q]$ is a segment of c .*

Proof.

(1) Let c be an \mathcal{O} -connected curve. We have shown in the proof of Lemma 7 that, for every \mathcal{O} -flat η and every two points $u, v \in c \cap \eta$, the segment $c[u, v]$ of the curve c is wholly contained in η . In particular, this observation holds for every two points $u, v \in c[p, q] \cap \eta$. Therefore, the intersection of $c[p, q]$ with every \mathcal{O} -flat η is path-connected; thus, $c[p, q]$ is \mathcal{O} -connected.

(2) Let $c[p, q]$ be an \mathcal{O} -connected curvilinear segment and l be the straight line through its endpoints p and q (see Figure 8a). We consider the curve c obtained from l by replacing the straight segment between p and q with the curvilinear segment $c[p, q]$ (this curve is shown by solid lines in Figure 8).

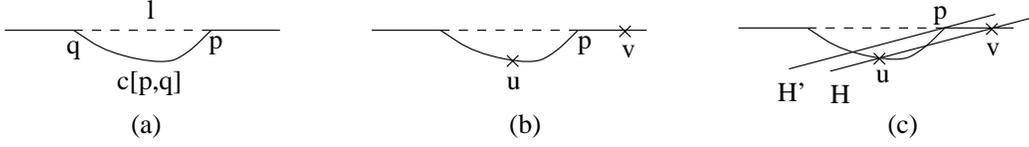


Figure 8: Proof of Lemma 8.

We prove that the curve c is \mathcal{O} -connected by demonstrating that, for every \mathcal{O} -hyperplane \mathcal{H} and every two points $u, v \in c \cap \mathcal{H}$, the segment $c[u, v]$ of the curve c is wholly in \mathcal{H} . We then conclude that the intersection of c with every \mathcal{O} -hyperplane \mathcal{H} is connected, which implies that c is \mathcal{O} -connected (Lemma 7).

If u and v are in $c[p, q]$, then the segment of $c[p, q]$ between u and v (that is, the segment $c[u, v]$) is wholly in \mathcal{H} , because $c[p, q] \cap \mathcal{H}$ is connected and $c[p, q]$ is simple.

If u and v are in l , then $l \subseteq \mathcal{H}$ and, hence, p and q are in \mathcal{H} . Since the intersection of $c[p, q]$ and \mathcal{H} is connected and $c[p, q]$ is simple, we conclude that $c[p, q]$ is wholly in \mathcal{H} . Thus, in this case, the curve c is wholly contained in \mathcal{H} and, hence, $c[u, v]$ is wholly in \mathcal{H} .

Finally, we consider the case when u is in $c[p, q]$ and v is in l (that is, v is in one of the two rays extending $c[p, q]$, as shown in Figure 8b). Without loss of generality, we assume that p (rather than q) is between u and v on the curve c (see Figure 8b). If $p \notin \mathcal{H}$, then the intersection of the segment $c[p, q]$ and the \mathcal{O} -hyperplane \mathcal{H}' through p parallel to \mathcal{H} is not connected (see Figure 8c), contradicting the \mathcal{O} -connectedness of $c[p, q]$. We conclude that $p \in \mathcal{H}$; therefore, the straight segment joining v and p and the segment of $c[p, q]$ between p and u are in \mathcal{H} . Thus, the segment of c between u and v (that is, the segment $c[u, v]$) is wholly in \mathcal{H} . \square

We next show that, if we cut a segment from an \mathcal{O} -connected curve and replace it with another \mathcal{O} -connected segment, then the resulting curve is also \mathcal{O} -connected (see Figure 9a).

Theorem 9 (Cutting and pasting) *Let p and q be two points of a simple \mathcal{O} -connected curve c . If we replace the part of the curve c between p and q with some other simple \mathcal{O} -connected segment, then the resulting curve c' is also \mathcal{O} -connected.*

Proof. We show that the curve c' is \mathcal{O} -connected by demonstrating that, for every \mathcal{O} -hyperplane \mathcal{H} and every two points $u, v \in c' \cap \mathcal{H}$, the segment $c'[u, v]$ of the curve c' is wholly in \mathcal{H} . Then, the intersection of c' with every \mathcal{O} -hyperplane \mathcal{H} is connected, which implies that c' is \mathcal{O} -connected (Lemma 7).

We denote the three parts of c' by c_1 , $c'[p, q]$, and c_2 , as shown in Figure 9(b). Note that c_1 and c_2 are \mathcal{O} -connected, because they are segments of the \mathcal{O} -connected curve c (Lemma 8). We use a case analysis, similar to the proof of Part 2 of Lemma 8. The cases are determined by the positions of the points u and v .

If u and v are in $c'[p, q]$, then the segment of $c'[p, q]$ between u and v is wholly in \mathcal{H} , because $c'[p, q] \cap \mathcal{H}$ is connected and $c'[p, q]$ is simple.

Similarly, if $u, v \in c_1$ (or $u, v \in c_2$), then the segment of the curve between u and v is wholly in \mathcal{H} , because c_1 is \mathcal{O} -connected and, hence, $c_1 \cap \mathcal{H}$ is connected.

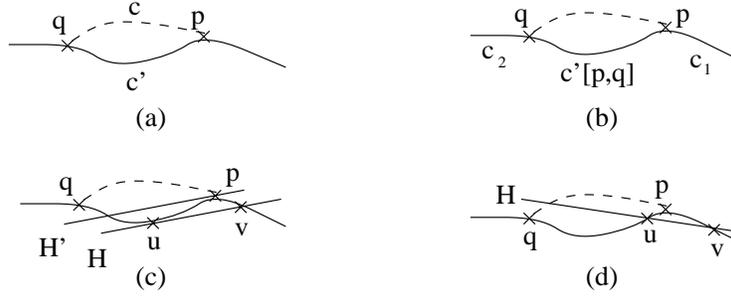


Figure 9: Proof of Theorem 9.

If $u \in c_1$ and $v \in c_2$, then the segment of the \mathcal{O} -connected curve c between u and v is wholly in \mathcal{H} ; therefore, p and q are in \mathcal{H} . Since the segment $c'[p, q]$ is \mathcal{O} -connected and its endpoints, p and q , are in \mathcal{H} , we conclude that $c'[p, q]$ is in \mathcal{H} . Thus, the segment of c' between u and v is wholly in \mathcal{H} .

Finally, we consider the case when $u \in c'[p, q]$ and $v \in c_1$ (or $v \in c_2$), as shown in Figure 9(c). We first demonstrate, by contradiction, that $p \in \mathcal{H}$. Suppose that $p \notin \mathcal{H}$. If p and q are “on the same side” of \mathcal{H} (see Figure 8c), then the intersection of $c'[p, q]$ and the \mathcal{O} -hyperplane \mathcal{H}' through p (or through q) parallel to \mathcal{H} is not connected (Figure 8c), contradicting the \mathcal{O} -connectedness of $c'[p, q]$. On the other hand, if p and q are “on different sides” of \mathcal{H} , or if $q \in \mathcal{H}$ (see Figure 8d), then the intersection of the curve c and the \mathcal{O} -hyperplane \mathcal{H} is not connected, since c contains the points v, p , and q (in this order), which again contradicts the \mathcal{O} -connectedness of c . We conclude that $p \in \mathcal{H}$; therefore, the segment of $c'[p, q]$ between u and p and the segment of c_1 between p and v are both in \mathcal{H} . Thus, the segment of c' between u and v is wholly in \mathcal{H} . \square

Finally, we state a condition under which the catenation of several curvilinear segments is an \mathcal{O} -connected segment (see Figure 10a).

Lemma 10 (Catenation) *Let p_0, p_1, \dots, p_n be a sequence of points connected by simple curvilinear segments $c[p_0, p_1], c[p_1, p_2], \dots, c[p_{n-1}, p_n]$. The union of these segments is an \mathcal{O} -connected segment if and only if the following two conditions hold:*

1. *Each of the segments is \mathcal{O} -connected.*
2. *For every \mathcal{O} -hyperplane \mathcal{H} , if \mathcal{H} intersects two segments, $c[p_{k-1}, p_k]$ and $c[p_m, p_{m+1}]$ (where $k \leq m$), then the points p_k, p_{k+1}, \dots, p_m are contained in \mathcal{H} .*

Proof. The union of the segments $c[p_0, p_1], c[p_1, p_2], \dots, c[p_{n-1}, p_n]$ is a curvilinear segment. We denote this segment by $c[p_0, p_n]$.

Suppose that the segment $c[p_0, p_n]$ is \mathcal{O} -connected. Then, by Lemma 8, every subsegment $c[p_k, p_{k+1}]$ of the segment $c[p_0, p_n]$ is \mathcal{O} -connected. The intersection of $c[p_0, p_n]$ with every \mathcal{O} -hyperplane \mathcal{H} is connected; therefore, if \mathcal{H} intersects subsegments $c[p_{k-1}, p_k]$ and $c[p_m, p_{m+1}]$, then the points p_k, \dots, p_m are in \mathcal{H} (see Figure 10a).

To prove the converse, suppose that the segments $c[p_0, p_1], c[p_1, p_2], \dots, c[p_{n-1}, p_n]$ are all *straight* segments (see Figure 10b). Then, Condition 2 of the lemma immediately implies

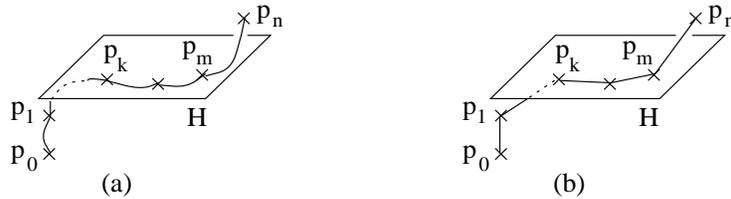


Figure 10: Proof of Lemma 10.

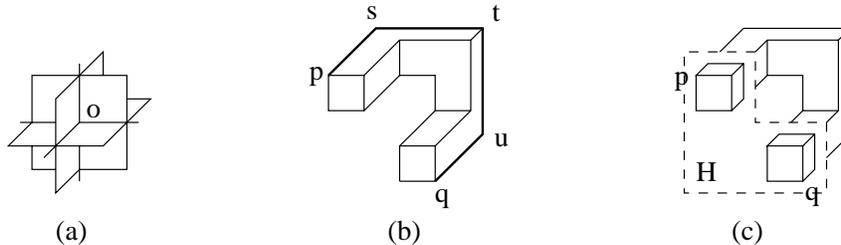


Figure 11: Generalized visibility.

that the intersection of the polygonal line $c[p_0, p_n]$ with every \mathcal{O} -hyperplane is connected. Therefore, by Lemma 7, $c[p_0, p_n]$ is \mathcal{O} -connected. If we now replace every straight segment $c[p_k, p_{k+1}]$ with an arbitrary \mathcal{O} -connected segment, the resulting new segment $c[p_0, p_n]$ is also \mathcal{O} -connected by Lemma 9. \square

6 Visibility

We present two notions of generalized visibility and characterize \mathcal{O} -convex and \mathcal{O} -connected sets in terms of this generalized visibility.

In standard convexity, two points of a set are *visible* to each other if the straight segment joining them is wholly contained in the set. For example, the points p and s of the set in Figure 11(b) are visible to each other, whereas p and q are not. We can characterize convex sets in terms of visibility: a set is convex if and only if every two points of the set are visible to each other.

Since \mathcal{O} -convexity is weaker than standard convexity, some points of an \mathcal{O} -convex set may not be visible to each other. For example, the set in Figure 11(b) is \mathcal{O} -convex for the orthogonal-orientation set shown in Figure 11(a), and the points p and q in that set are not visible to each other.

We define weaker visibility, which enables us to characterize \mathcal{O} -convex sets: we say that two points are visible to each other if there is an \mathcal{O} -convex curvilinear segment joining them that is wholly in the set. For example, we can join the points p and q in Figure 11(b) by the \mathcal{O} -convex polygonal line (p, s, t, u, q) , which is contained in the set.

We can describe path-connected \mathcal{O} -convex sets in terms of this generalized visibility much in the same way as we describe standard convex sets through standard visibility.

Theorem 11 (Visibility for \mathcal{O} -convex sets) *A path-connected set is \mathcal{O} -convex if and only if every two points of the set can be joined by an \mathcal{O} -convex curvilinear segment that is*



Figure 12: Proof of Theorem 11.

wholly contained in the set.

Proof. Suppose that every two points of a set P can be joined by an \mathcal{O} -convex segment that is wholly contained in P . We observe that, if a line through two points p and q of P is an \mathcal{O} -line, then the only \mathcal{O} -convex curvilinear segment joining p and q is the straight segment; therefore, this straight segment joining p and q is wholly in P . This observation implies that the intersection of every \mathcal{O} -line with P is connected; thus, P is \mathcal{O} -convex.

Suppose, conversely, that a set P is path-connected and \mathcal{O} -convex. To demonstrate that every two points p and q of P can be joined by an \mathcal{O} -convex segment that is wholly contained in P , we consider a shortest curvilinear segment $c[p, q]$ joining p and q in P . (Since all sets in our exploration are assumed to be closed, there is such a shortest segment.) We prove, by contradiction, that $c[p, q]$ is \mathcal{O} -convex.

Suppose that the segment $c[p, q]$ is *not* \mathcal{O} -convex. Then, the intersection of $c[p, q]$ with some \mathcal{O} -line l is disconnected (see Figure 12). Therefore, for some points $u, v \in c[p, q] \cap l$, the segment of the line l between p and q is not in $c[p, q]$. Since P is \mathcal{O} -convex, this segment is wholly in P . Replacing the subsegment of $c[p, q]$ between u and v with the straight segment, we obtain a shorter path from p to q in P , contradicting the assumption that $c[p, q]$ is a shortest path. \square

We can characterize \mathcal{O} -connected sets in a similar way, if we define visibility in terms of \mathcal{O} -connected segments joining points of a set. This type of visibility is stronger than \mathcal{O} -convex visibility: two points sometimes cannot be joined by an \mathcal{O} -connected segment even when they can be joined by an \mathcal{O} -convex segment. For example, the points p and q in Figure 11(b) cannot be joined by an \mathcal{O} -connected segment contained in the set, because the intersection of the \mathcal{O} -plane H (Figure 11c) with every segment joining these two points is disconnected.

Theorem 12 (Visibility for \mathcal{O} -connected sets) *A set is \mathcal{O} -connected if and only if every two points of the set can be joined by an \mathcal{O} -connected curvilinear segment that is wholly contained in the set.*

Proof. Suppose that every two points p and q of a set P can be joined by an \mathcal{O} -connected segment. If p and q are in some \mathcal{O} -flat, then the \mathcal{O} -connected segment joining them is wholly in this flat (see the proof of Lemma 7). We conclude that the intersection of P with every \mathcal{O} -flat is path-connected and, hence, P is \mathcal{O} -connected.

We use induction on the dimension d to prove that, conversely, every two points of an \mathcal{O} -connected set P can be joined by an \mathcal{O} -connected path in P . In two dimensions, every two points of an \mathcal{O} -connected set can be joined by an \mathcal{O} -convex path that is wholly in the set (Theorem 11) and every \mathcal{O} -convex path is \mathcal{O} -connected, which establishes an induction base.

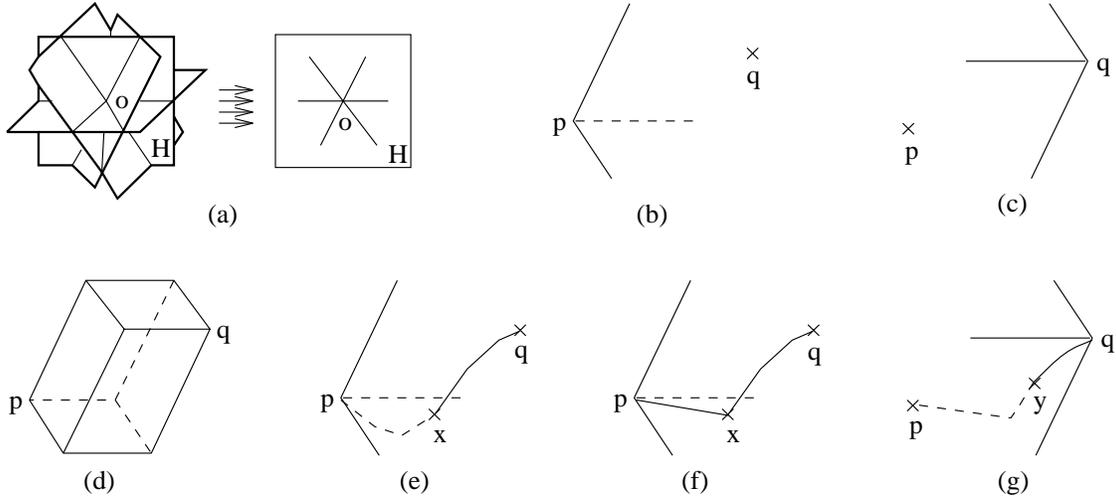


Figure 13: Constructions in the proof of Theorem 12.

The proof of the induction step consists of three parts. First, we show that, if two points p and q of an \mathcal{O} -connected set P in d dimensions are contained in some \mathcal{O} -hyperplane, then they can be joined by an \mathcal{O} -connected path in P . We then consider the case when there is no \mathcal{O} -hyperplane through p and q , introduce the notion of the \mathcal{O} -block of p and q , and show that there is a path from p to q that is wholly contained in $P \cap \mathcal{O}\text{-block}(p, q)$. Finally, we use this result to construct an \mathcal{O} -connected path in P from p to q .

Suppose that points p and q of P belong to some \mathcal{O} -hyperplane \mathcal{H} . We may view \mathcal{H} as an independent $(d - 1)$ -dimensional space and define an orientation set $\mathcal{O}_{\mathcal{H}}$ in this space. We define an $\mathcal{O}_{\mathcal{H}}$ -flat as an \mathcal{O} -flat contained in \mathcal{H} ; the $(d - 2)$ -dimensional $\mathcal{O}_{\mathcal{H}}$ -flats through some fixed point form the orientation set $\mathcal{O}_{\mathcal{H}}$. We illustrate this definition in Figure 13(a): the \mathcal{O} -plane H contains three \mathcal{O} -lines through o , which form the three-element orientation set \mathcal{O}_H .

The $\mathcal{O}_{\mathcal{H}}$ -flats have all necessary properties of \mathcal{O} -flats [2]:

1. Every translation of an $\mathcal{O}_{\mathcal{H}}$ -flat within the space \mathcal{H} is an $\mathcal{O}_{\mathcal{H}}$ -flat.
2. A set $\eta \in \mathcal{H}$ is an $\mathcal{O}_{\mathcal{H}}$ -flat if and only if it is either \mathcal{H} itself or the intersection of $(d - 2)$ -dimensional $\mathcal{O}_{\mathcal{H}}$ -flats ($\mathcal{O}_{\mathcal{H}}$ -hyperplanes).

The orientation set $\mathcal{O}_{\mathcal{H}}$ gives rise to $\mathcal{O}_{\mathcal{H}}$ -connected sets in the space \mathcal{H} , which are defined in terms of the path-connectedness of their intersections with $\mathcal{O}_{\mathcal{H}}$ -flats. Since $\mathcal{O}_{\mathcal{H}}$ -flats are \mathcal{O} -flats, a set in \mathcal{H} is $\mathcal{O}_{\mathcal{H}}$ -connected if and only if it is \mathcal{O} -connected.

The intersection of P with \mathcal{H} is \mathcal{O} -connected (Theorem 5) and, hence, it is $\mathcal{O}_{\mathcal{H}}$ -connected. By the induction hypothesis, the points p and q can be connected by an $\mathcal{O}_{\mathcal{H}}$ -connected segment in $P \cap \mathcal{H}$; therefore, they can be connected by an \mathcal{O} -connected segment in P .

Now suppose that there is no \mathcal{O} -hyperplane through points p and q of P . We define the \mathcal{O} -block of p and q [3] and show that there is a path from p to q in $P \cap \mathcal{O}\text{-block}(p, q)$.

Let S_p be the intersection of all the halfspaces, containing q , whose boundaries are \mathcal{O} -hyperplanes through p . The set S_p is a polyhedral angle with vertex p (see Figure 13b).

Similarly, let S_q be the intersection of all the halfspaces, containing p , whose boundaries are \mathcal{O} -hyperplanes through q (see Figure 13c). The \mathcal{O} -block of p and q is the intersection of S_p and S_q (see Figure 13d). Observe that every \mathcal{O} -connected path from p to q is wholly in $\mathcal{O}\text{-block}(p, q)$, since the intersection of such a path with every \mathcal{O} -hyperplane is connected.

To show that there is a path, from p to q , that is wholly in $P \cap \mathcal{O}\text{-block}(p, q)$, we consider some path $c[p, q] \subseteq P$ from p to q . If $c[p, q]$ is *not* wholly in S_p , we select a point x in the intersection of $c[p, q]$ with the boundary of S_p such that the segment of $c[p, q]$ between x and q is wholly in S_p (see Figure 13e). We note that there is some \mathcal{O} -hyperplane \mathcal{H} through p and x ; therefore, we can join p and x by an \mathcal{O} -connected path in $P \cap \mathcal{H}$. We replace the segment of $c[p, q]$ between p and x with this \mathcal{O} -connected path, thus obtaining a new path, say $c_1[p, q]$, from p to q , which is wholly contained in $P \cap S_p$ (see Figure 13f).

If $c_1[p, q]$ is *not* wholly in S_q , we select a point y in the intersection of $c_1[p, q]$ with the boundary of S_q such that the segment of $c_1[p, q]$ between p and y is wholly in S_q (see Figure 13g). We replace the segment of $c_1[p, q]$ between y and q with an \mathcal{O} -connected path in P , thus obtaining a new path, from p to q , which is wholly in $P \cap \mathcal{O}\text{-block}(p, q)$.

We now construct an \mathcal{O} -connected path in P from p to q . We consider some path from p to q in $P \cap \mathcal{O}\text{-block}(p, q)$ and choose a point z in this path such that the distance between p and z is equal to the distance between z and q . We now consider some path from p to z in $P \cap \mathcal{O}\text{-block}(p, z)$ and some path from z to q in $P \cap \mathcal{O}\text{-block}(z, q)$, and choose a point in each of these two paths, in the same way as we chose the point z in the path from p to q .

We recursively repeat this point-selection operation; on the n -th level of recursion, we obtain $2^n - 1$ intermediate points, connected by 2^n segments that form a path from p to q . The distances between consecutive points converge to zero as n tends to infinity. Observe that, by construction, the intersection of this path with every \mathcal{O} -hyperplane through any of the $2^n - 1$ intermediate points is connected.

The closure of the set of points selected in infinitely many recursive steps is a path, from p to q , contained in P . The intersection of this path with every \mathcal{O} -hyperplane is connected and, hence, the path is \mathcal{O} -connected. \square

7 Concluding remarks

We generalized \mathcal{O} -convexity to three and higher dimensions and demonstrated that the properties of \mathcal{O} -convex sets are similar to the properties of standard convex sets. The main property of convex sets that we lose in \mathcal{O} -convexity is connectedness: a convex set is always connected, whereas an \mathcal{O} -convex set may be disconnected. To bridge this difference, we introduced \mathcal{O} -connected sets, which are always connected, and demonstrated that their properties are also similar to the properties of standard convex sets.

The work presented in this paper is just beginning; it leaves many open research problems, which we are currently trying to solve.

For example, we have not established the *contractability* of \mathcal{O} -connected sets. Intuitively, a set is contractable if it is connected and does not have holes. For example, flats and balls are contractable. On the other hand, a hollow sphere is not contractable, because it has a cavity inside. A doughnut (torus) is not contractable either, because it has a hole through it. To put it more formally, a set is contractable if it can be continuously transformed

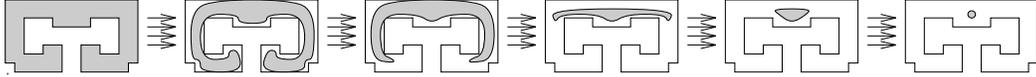


Figure 14: The contraction of a set to a point.

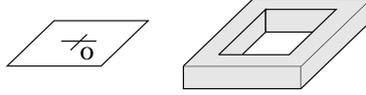


Figure 15: An \mathcal{O} -connected set that is not contractible.

(contracted) to a point in such a way that all intermediate stages of the transformation are contained in the original set (see Figure 14).

In two dimensions, contractability is equivalent to simple connectedness. In higher dimensions, however, contractability is a stronger property than simple connectedness. For example, a hollow sphere is simply connected, even though it is not contractible.

Convex sets are always contractible. Connected \mathcal{O} -convex sets in *two dimensions* are also contractible, if the orientation set \mathcal{O} contains at least one line (see Lemma 1). This property of \mathcal{O} -convex sets, however, does not hold in higher dimensions. In Figure 3(d), we provide an example of a connected \mathcal{O} -convex set that is not contractible.

If the orientation set \mathcal{O} contains too few hyperplanes, then even \mathcal{O} -connected sets may not be contractible. Here, by too few we mean that the intersections of \mathcal{O} -hyperplanes do not form \mathcal{O} -lines. For example, consider the orientation set in \mathcal{R}^3 that contains only one plane (see Figure 15). Then, there are no \mathcal{O} -lines and some \mathcal{O} -connected sets are not contractible. We show one such set in Figure 15: it is \mathcal{O} -connected, since its intersection with every \mathcal{O} -plane is path-connected; however, it is not contractible.

We conjecture that, *if the set of \mathcal{O} -lines is nonempty, then every \mathcal{O} -connected set is contractible*. This result would generalize Property 5 of planar \mathcal{O} -convex sets (see Lemma 1).

As another example of an open problem, we have not characterized the boundaries of \mathcal{O} -convex polytopes. In two dimensions, if the orientation set contains n lines, the boundary of every \mathcal{O} -convex polygon can be partitioned into at most n \mathcal{O} -convex polygonal lines [16]. We conjecture that a similar result holds in higher dimensions: for every finite orientation set \mathcal{O} , there is some fixed number n such that the boundary of every \mathcal{O} -convex polytope can be partitioned into at most n connected \mathcal{O} -convex regions.

We also have not studied the computational aspects of \mathcal{O} -convexity, such as verifying \mathcal{O} -convexity and \mathcal{O} -connectedness of a polytope and computing the \mathcal{O} -convex hull. Finally, we may explore other generalizations of convexity, such as NESW convexity [16, 19], in higher dimensions.

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. Generalized halfspaces in 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. *Conquering Contours: Efficient Algorithms for Computational Geometry*. PhD thesis, Universität Dortmund, Germany, 1983.
- [7] Ralf Hartmut Güting. Stabbing C -oriented polygons. *Information Processing Letters*, 16:35–40, 1983.
- [8] Ralf Hartmut Güting. Dynamic C -oriented polygonal intersection searching. *Information and Control*, 63:143–163, 1984.
- [9] Paul J. Kelly and Max L. Weiss. *Geometry and Convexity: A Study in Mathematical Methods*. John Wiley & Sons, New York, NY, 1978.
- [10] Victor Klee. What is a convex set? *American Mathematical Monthly*, 78:616–631, 1971.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] Thomas Ottmann, Eljas Soisalon-Soininen, and Derick Wood. On the definition and computation of rectilinear convex hulls. *Information Sciences*, 33:157–171, 1984.
- [15] Franco P. Preparata and Michael I. Shamos. *Computational Geometry*. Springer-Verlag, New York, NY, 1985.
- [16] Gregory J. E. Rawlins. *Explorations in Restricted-Orientation Geometry*. PhD thesis, University of Waterloo, 1987. Technical Report CS-89-48.
- [17] Gregory J. E. Rawlins and Derick Wood. Optimal computation of finitely oriented convex hulls. *Information and Computation*, 72:150–166, 1987.

- [18] 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.
- [19] Gregory J. E. Rawlins and Derick Wood. A decomposition theorem for convexity spaces. *Journal of Geometry*, 36:263–281, 1989.
- [20] Gregory J. E. Rawlins and Derick Wood. Restricted-orientation convex sets. *Information Sciences*, 54:263–281, 1991.
- [21] Sven Schuierer. *On Generalized Visibility*. PhD thesis, Universität Freiburg, Germany, 1991.
- [22] 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.
- [23] F. A. Valentine. Local convexity and L_n -sets. *Proceedings of the American Mathematical Society*, 16:1305–1310, 1965.
- [24] Peter Widmayer, Y. F. Wu, and C. K. Wong. On some distance problems in fixed orientations. *SIAM Journal on Computing*, 16:728–746, 1987.