Restricted-Orientation Convexity in Higher-Dimensional Spaces

Eugene Fink   Derick Wood

University of Waterloo, Waterloo, Ont., Canada N2L3G1
{efink, dwood}@violet.waterloo.edu

ABSTRACT

A *restricted-oriented convex set* is a set whose intersection with any line from a fixed set of orientations is either empty or connected. This notion generalizes both orthogonal convexity and normal convexity. The aim of this paper is to establish a mathematical foundation for the theory of restricted-oriented convex sets in higher-dimensional spaces.

1 Introduction

Restricted-orientation geometry is the study of the properties of geometric objects, whose facet orientations are restricted, and interaction of such restricted objects with unrestricted geometrical objects. The study of restricted-oriented objects was initiated by Guting [2, 3] and further developed by Widmayer et al. [9]. The restricted-orientation convexity was introduced by Rawlins as a subarea of restricted-orientation geometry [4, 5]. The research in this area was continued by Schuierer [8]. In all cases, only planar sets have been studied.

In this paper we explore properties of restricted-oriented convex sets in a higher-dimensional space $\mathcal{R}^d$. We show that some of planar-case results may be generalized to a higher-dimensional case. It turns out that planar-case theorems are considerably harder to prove for higher-dimensional spaces, and many of them do not hold, or hold only partially. Also, we show a connection between properties of restricted-oriented convex objects in the space $\mathcal{R}^d$ and in the lower-dimensional subspaces of $\mathcal{R}^d$.

All results are presented without proofs, but the proofs may be found in [1].

2 Basic definitions and notation

We are going to work with a higher-dimensional Euclidean space $\mathcal{R}^d$. We use the letter d to denote the dimension of $\mathcal{R}^d$. All geometrical objects discussed in the paper are assumed to be closed, unless otherwise specified.

We denote subsets of $\mathcal{R}^d$ by capital letters, usually P or Q, and points by lower case letters, usually p or q. We denote a straight line by the letter l, or, if it passes through points p and q, by (p, q), a curve by the letter c, and a curvilinear segment with endpoints p and q by $c[p, q]$.

A *k-flat* in $\mathcal{R}^d$ is a subset of $\mathcal{R}^d$ which is itself a k-dimensional space. For example, a 1-flat is a straight line, a 2-flat is a usual 2-dimensional plane, and a 0-flat is a point. We denote flats by the Greek letters η and μ. $(d-1)$-flats are called *hyperplanes*. A hyperplane is denoted by $\mathcal{H}$, and a halfspace bounded by $\mathcal{H}$ is denoted by $[\mathcal{H}, p)$, where p is some point in the interior of the halfspace.

A *solid angle* is a union of rays in $\mathcal{R}^d$ with a common endpoint o, such that the intersection of this union with any hyperplane is connected (see Fig. 1). The point o is called the *vertex* of the solid angle.

3 Sets of orientations and ranges

3.1 Orientations

We say that two k-flats have the same *orientation* if they are parallel. Thus, the orientation of a k-flat η may be viewed as the set of all k-flats parallel to η. The orientation of a k-flat is called a k-orientation, and the orientation of a hyperplane is called a *hyperorientation*. We denote the orientation of η by $\bar{\eta}$ (the same letter with the bar above it).

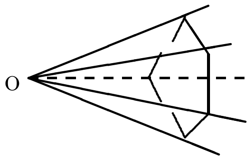

Fig. 1. Solid angle

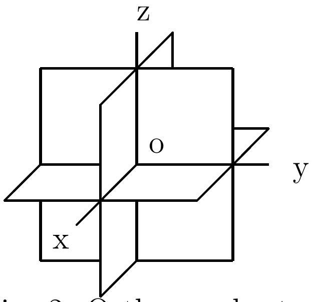

Fig. 2. Orthogonal set of orientations

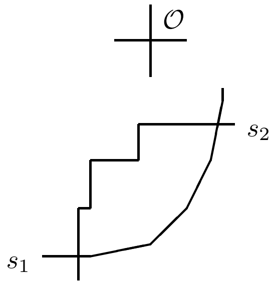

Fig. 3. Intersection of two $\mathcal{O}$-connected curves

Let us fix some point o in $\mathcal{R}^d$. It's easy to see that for each k-orientation, there is exactly one k-flat of this orientation that contains o. Thus, there is a one-to-one correspondence between orientations and flats through o. We use this correspondence as one of our basic tools in presenting properties of orientations.

Let l be a straight line through o, and η be a flat containing o. If $l \subseteq \eta$, we say that $\bar{l} \in \bar{\eta}$. Thus, we represent every k-orientation as a set of line orientations. We consider the empty set to be a 0-orientation, which is the orientation of a point. We define the inclusion relationship between orientations as the usual set inclusion and the intersection of orientations as the set intersection:

$$\text{Inclusion: } \bar{\eta} \subseteq \bar{\mu} \iff (\forall \bar{l} \in \bar{\eta})\ \bar{l} \in \bar{\mu} \qquad \text{Intersection: } \bigcap_{i \in I} \bar{\eta}_i = \{\bar{l} \mid (\forall i \in I)\ l \in \bar{\eta}_i\}$$

The following properties of orientations show the correspondence between orientations and their representation as flats through a fixed point o.

(1) Let η and μ be two flats whose intersection is not empty. Then $\eta \subseteq \mu$ if and only $\bar{\eta} \subseteq \bar{\mu}$.
(2) Let $\{\eta_i\}_{i \in I}$ be a family of flats, and $\{\bar{\eta}_i\}_{i \in I}$ be the corresponding family of orientations. Then $\bar{\eta} = \bigcap_{i \in I} \bar{\eta}_i$ is an orientation. Moreover, if $\bigcap_{i \in I} \eta_i$ is not empty, then it is a flat, whose orientation is $\bar{\eta}$.

3.2 Sets of orientations

We use the symbol $\mathcal{O}$ to refer to a *closed* set of hyperorientations. (To define a *closed* set we need some notion of the *distance* between two orientations. We use the *angle* between flats as a measure of the distance between them.) The elements of $\mathcal{O}$ are called *$\mathcal{O}$-hyperorientations.* We will slightly abuse notation by writing $\mathcal{H} \in \mathcal{O}$ in the case when $\overline{\mathcal{H}} \in \mathcal{O}$. If $\mathcal{H} \in \mathcal{O}$, we call $\mathcal{H}$ an *$\mathcal{O}$-hyperplane.* An *orthogonal* set of orientations $\mathcal{O}_\perp$ is a set of d hyperplane orientations in $\mathcal{R}^d$ such that any two orientations are perpendicular to each other.

We denote by $\mathcal{O}^k$ the set of all k-orientations formed by the intersections of elements of $\mathcal{O}$:

$$\mathcal{O}^k = \{\bar{\eta} \mid \dim(\bar{\eta}) = k \text{ and } \bar{\eta} = \bigcap_{i \in I} \overline{\mathcal{H}}_i, \text{ where } \{\overline{\mathcal{H}}_i\}_{i \in I} \subseteq \mathcal{O}\}$$

The elements of $\mathcal{O}^k$ are called k-dimensional *$\mathcal{O}$-orientations* or, shortly, $\mathcal{O}^k$-orientations (see Fig. 2). If $\bar{\eta} \in \mathcal{O}^k$, we write $\eta \in \mathcal{O}^k$ and call η an *$\mathcal{O}$-oriented* flat or, shortly, *$\mathcal{O}$-flat.* The elements of $\mathcal{O}^1$ are called *$\mathcal{O}$-lines.* We define the d-dimensional $\mathcal{O}$-orientation as the orientation of the whole space $\mathcal{R}^d$, that is the set of all line orientations. The following are basic properties of $\mathcal{O}$-orientations:

(1) For all $k \in [1..d]$, the set $\mathcal{O}^k$ is closed.
(2) The intersection of $\mathcal{O}$-orientations is always an $\mathcal{O}$-orientation.
(3) The intersection of $\mathcal{O}$-flats is either empty or an $\mathcal{O}$-flat.
(4) For any natural numbers $k \le m \le d$, every $\mathcal{O}^k$-orientation is a subset of some $\mathcal{O}^m$-orientation.
(5) Let n be the number of $\mathcal{O}$-hyperorientations, $n = |\mathcal{O}|$. Then the number of distinct $\mathcal{O}^k$-orientations is no greater than $C_n^{d-k} = \frac{n!}{(d-k)! \cdot (n-d+k)!}$.

Now let us consider some $\mathcal{O}$-hyperplane $\mathcal{H}$. We define $\mathcal{O}(\mathcal{H})$ as the set of the intersections of $\overline{\mathcal{H}}$ with all other $\mathcal{O}$-orientations: $\mathcal{O}(\mathcal{H}) = \{\overline{\mathcal{H}} \cap \overline{\mathcal{H}}' \mid \overline{\mathcal{H}}' \in (\mathcal{O} - \{\overline{\mathcal{H}}\})\}$. Clearly, $\mathcal{O}(\mathcal{H})$ is a set of $(d-2)$-orientations, and all elements of $\mathcal{O}(\mathcal{H})$ are subsets of the orientation $\overline{\mathcal{H}}$: $(\forall \bar{\eta} \in \mathcal{O}(\mathcal{H}))\ \bar{\eta} \subseteq \overline{\mathcal{H}}$. A hyperplane $\mathcal{H}$ is a

$(d-1)$-dimensional space, and $\mathcal{O}(\mathcal{H})$ may be viewed as a set of hyperorientations in this space. It may be shown that $\mathcal{O}(\mathcal{H})$ *is closed*. The set $\mathcal{O}(\mathcal{H})$ in the space $\mathcal{H}$ has properties similar to the properties of $\mathcal{O}$ in $\mathcal{R}^d$. In particular, the intersections of elements of $\mathcal{O}(\mathcal{H})$ produce the lower-dimensional orientations. We define $\mathcal{O}^k(\mathcal{H})$ as the set of k-orientations produced by the intersections of elements of $\mathcal{O}(\mathcal{H})$. It may be shown that a k-flat is an $\mathcal{O}(\mathcal{H})$-oriented if and only if it is $\mathcal{O}$-oriented and contained in $\mathcal{H}$:

$$(\forall k \in [1..d-2])\ \mathcal{O}^k(\mathcal{H}) = \{\overline{\eta} \in \mathcal{O}^k \mid \overline{\eta} \subseteq \overline{\mathcal{H}}\}$$

3.3 Hyperranges

Let us represent a set of orientations as a set of flats through a fixed point o. Consider a solid angle Θ with the vertex o, such that the interior of Θ is *non-empty* and *connected*, and Θ is the closure of its own interior. Informally, Θ is a closed solid angle, and each point of its boundary is "infinitely close" to the interior. We define $\overline{\Theta}$ as the set of orientations of all lines through o that intersect Θ *not only* in o: $\overline{\Theta} = \{\overline{(o,p)} \mid p \in (\Theta - \{o\})\}$. $\overline{\Theta}$ is called a *hyperrange* of orientations. The interior of a hyperrange $\overline{\Theta}$ is defined as the set of orientations of all lines through o that intersect the interior of Θ: $\overline{\Theta}_{int} = \{\overline{(o,p)} \mid p \in \Theta_{int}\}$.

A hyperrange $\overline{\Theta}$ is called $\mathcal{O}$-*free* if the intersection of its interior with any $\mathcal{O}$-hyperrange is empty: $(\forall \overline{\mathcal{H}} \in \mathcal{O})\ \overline{\Theta}_{int} \cap \overline{\mathcal{H}} = \emptyset$. Equivalently, we can say that no $\mathcal{O}$-hyperplane through o intersects the interior of the corresponding solid angle Θ. An $\mathcal{O}$-free range which is not a proper subset of any other $\mathcal{O}$-free range is called an $\mathcal{O}$-*hyperrange*. The solid angles corresponding to $\mathcal{O}$-hyperranges are called $\mathcal{O}$-*hyperangles*.

Lemma 3.1 *Let o and p be two points in $\mathcal{R}^d$ not lying on the same $\mathcal{O}$-hyperplane, i.e. $(\nexists \mathcal{H} \in \mathcal{O})\ o, p \in \mathcal{H}$, and Θ be an $\mathcal{O}$-hyperangle with the vertex o that contains p. Then $\Theta = \bigcap\{[\mathcal{H}, p) \mid \mathcal{H} \in \mathcal{O} \text{ and } o \in \mathcal{H}\}$.*

Observe that according to this lemma, if line orientation is not lying within an $\mathcal{O}$-hyperorientation, it is contained in the unique $\mathcal{O}$-hyperrange. The lemma allows us to describe the set of all $\mathcal{O}$-hyperranges in terms of $\mathcal{O}$-hyperorientations. To present this description, we fix some point o and draw all possible $\mathcal{O}$-hyperplanes through o. The hyperplanes partition the space into solid angles. These solid angles, bounded by $\mathcal{O}$-hyperranges, are the $\mathcal{O}$-hyperangles with the vertex o. Fig. 2 presents the described picture for the orthogonal set of orientations in $\mathcal{R}^3$, where three $\mathcal{O}$-hyperplanes partition the space into eight $\mathcal{O}$-hyperangles.

3.4 Ranges

Unfortunately, $\mathcal{O}$-hyperranges do not provide a disjoint partition of the set of line orientations. A line orientation on the boundary of $\mathcal{O}$-hyperranges belongs to two different $\mathcal{O}$-hyperranges. On the other hand, a line orientation may not belong to any $\mathcal{O}$-hyperrange. For example, if $\mathcal{O}$ contains *all* hyperorientations, then there are no $\mathcal{O}$-hyperranges at all. In this subsection we show that the picture may be improved by introducing the notion of lower-dimensional $\mathcal{O}$-*ranges*.

Recall that every $\mathcal{O}$-hyperplane $\mathcal{H}$ may be viewed as a $(d-1)$-dimensional space with the set of orientations $\mathcal{O}(\mathcal{H})$. An $\mathcal{O}(\mathcal{H})$-hyperrange $\overline{\theta}$ in this space is called an $\mathcal{O}^{d-1}$-*range*. The interior of $\overline{\theta}$ *in the space* $\mathcal{H}$ is denoted by $\overline{\theta}_{int}$. (Note that we define the interior of $\overline{\theta}$ in the $(d-1)$-dimensional space $\mathcal{H}$, not in the whole space $\mathcal{R}^d$, where the interior of $\overline{\theta}$ is trivially empty.) Similarly, we may define $\mathcal{O}^k$-ranges and their interiors for all $k \in [1..d]$. The solid angles corresponding to $\mathcal{O}$-ranges are called $\mathcal{O}$-*angles*.

$\mathcal{O}$-ranges may be described informally by representing them as $\mathcal{O}$-angles with a common vertex (see Fig. 2). The facets of hyperangles are $\mathcal{O}^{d-1}$-angles (e.g. the planar angle $\angle xoy$ on Fig. 2), the subfacets are $\mathcal{O}^{d-2}$-angles, and so on. Generally, a solid angle is an $\mathcal{O}^k$-angle if and only if it is a k-dimensional face of some $\mathcal{O}$-hyperangle. Observe that $\mathcal{O}^1$-angles are $\mathcal{O}$-lines, and the interiors of such angles (in 1-dimensional space) are the whole $\mathcal{O}$-lines themselves. One can easily see that every point except the common vertex o belongs to the interior of exactly one $\mathcal{O}$-angle. This gives rise to the following theorem.

Theorem 3.1 (The Partition Theorem) *The set of interiors of all $\mathcal{O}$-ranges is a disjoint partition of the set of all line orientations.*

The last theorem in this section generalizes Lemma 3.1. It presents a formal description of $\mathcal{O}$-angles in terms of hyperplanes.

Theorem 3.2 (The Range Theorem) *Let o and p be two points in $\mathcal{R}^d$, and o be the $\mathcal{O}$-angle with the vertex o whose interior contains p. Then $\theta = \bigcap\{[\mathcal{H}, p) \mid \mathcal{H} \in \mathcal{O} \text{ and } o \in \mathcal{H}\}$.*

4 Restricted orientation convexity

4.1 Abstract convexity

Before discussing restricted orientation convexity, we review basic definitions from abstract convexity theory. Given a set S and a family C of subsets of S, the structure (S, C) is said to be a *convexity space* if $\emptyset, S \in C$ and C is closed under intersection: $(\forall C' \subseteq C)\ \bigcap C' \in C$. For example, the convex sets in $\mathcal{R}^d$ form a convexity space. In this convexity space $S = \mathcal{R}^d$ and C is the set of the usual convex sets. Another example of a convexity space is the set of $\mathcal{O}$-orientations, where S is the set of all line orientations, and $C = \mathcal{O}^0 \cup \mathcal{O}^1 \cup \ldots \cup \mathcal{O}^d$. (Recall that the intersection of $\mathcal{O}$-orientations is always an $\mathcal{O}$-orientation.)

Given a convexity space (S, C), we define the *hull* of a subset P of S as the intersection of all convex sets containing P: $hull(P) = \bigcap\{Q \mid P \subseteq Q \text{ and } Q \in C\}$.

A *nested chain* is a (possibly infinite) sequence of sets $P_0, P_1, P_2, \ldots$ such that $P_0 \subseteq P_1 \subseteq P_2 \subseteq \ldots$ A convexity space is called *aligned* if, for any nested chain $\mathcal{N}$ of convex sets, the union of all elements of $\mathcal{N}$ is also a convex set: $(\forall \text{ chain } \mathcal{N} \subseteq C)\ \bigcup \mathcal{N} \in C$. Alignment is an important property that allows to us prove several basic facts about convexity spaces [5].

4.2 $\mathcal{O}$-convexity

The following definition of higher-dimensional $\mathcal{O}$-convexity, which we use in this article, was stated in [4].

Definition 4.1 ($\mathcal{O}$-convexity)
*(1) Let P be a planar set, $P \subseteq \mathcal{R}^2$, and $\mathcal{O}$ be a set of orientations in $\mathcal{R}^2$ (in this case $\mathcal{O}$ is a set of line orientations). P is $\mathcal{O}$-*convex *if for every $l \in \mathcal{O}$, $P \cap l$ is connected. (We consider the empty set to be connected.)*

*(2) Let P be a d-dimensional set, $P \subseteq \mathcal{R}^d$, and $\mathcal{O}$ be a set of orientations in $\mathcal{R}^d$. P is $\mathcal{O}$-*convex *if for every $\mathcal{O}$-hyperplane $\mathcal{H}$, $P \cap \mathcal{H}$ is $\mathcal{O}(\mathcal{H})$-convex, that is $P \cap \mathcal{H}$ is $\mathcal{O}$-convex in $(d-1)$-dimensional space $\mathcal{H}$, w.r.t. the set of orientations $\mathcal{O}(\mathcal{H})$.*

The next theorem shows that this recursive definition may be considerably simplified.

Theorem 4.1 *Let $P \subseteq \mathcal{R}^d$, and $\mathcal{O}$ be a set of orientations in $\mathcal{R}^d$. Then P is $\mathcal{O}$-convex if and only if for every $\mathcal{O}$-line $l \in \mathcal{O}^1$, $P \cap l$ is connected.*

This theorem allows us to prove basic properties of $\mathcal{O}$-convex sets, presented below. (For the planar case, the first three of these properties were presented in [7].)

(1) For any set of orientations $\mathcal{O}$, if P is convex then P is $\mathcal{O}$-convex.
(2) The intersection of $\mathcal{O}$-convex sets is an $\mathcal{O}$-convex set.
(3) A set is $\mathcal{O}$-convex if and only if it is the union of disjoint connected components such that each component is $\mathcal{O}$-convex and no $\mathcal{O}$-line intersects any pair of components.
(4) For any $\mathcal{O}$, the set of all connected $\mathcal{O}$-convex sets is aligned.
(5) Let $\mathcal{H}$ be an $\mathcal{O}$-hyperplane, and $P \subseteq \mathcal{H}$. Then P is $\mathcal{O}$-convex if and only if it is $\mathcal{O}(\mathcal{H})$-convex.

Since the empty set and the whole space $\mathcal{R}^d$ are clearly $\mathcal{O}$-convex, and the intersection of $\mathcal{O}$-convex sets is always $\mathcal{O}$-convex, we conclude that *$\mathcal{O}$-convex sets form a convexity space.* We denote the hull of a set P in this space by $\mathcal{O}$-hull(P). Below we present two properties of $\mathcal{O}$-hulls. (For the planar case, these properties were stated in [4].)

(1) The $\mathcal{O}$-hull of a connected set is connected.
(2) Let $\{\mathcal{O}_i\}_{i \in I}$ be a family of sets of orientations, and $\mathcal{O}$ be such a set of orientations that $\mathcal{O}^1 = \bigcup\{\mathcal{O}_i^1\}$. Then for any set P, $\bigcup_{i \in I}(\mathcal{O}_i\text{-hull}(P)) \subseteq \mathcal{O}\text{-hull}(P)$.

4.3 $\mathcal{O}$-connectedness

Definition 4.2 ($\mathcal{O}$-connectedness) *A planar set is $\mathcal{O}$-*connected *if it is $\mathcal{O}$-convex and connected. A d-dimensional set P is $\mathcal{O}$-*connected *if it is connected, and for every $\mathcal{O}$-hyperplane $\mathcal{H}$, $P \cap \mathcal{H}$ is $\mathcal{O}(\mathcal{H})$-connected.*

$\mathcal{O}$-connected sets are similar to convex sets in many respects. However, they do *not* form a convexity space, because the intersection of two $\mathcal{O}$-connected sets may not be $\mathcal{O}$-connected. Fig. 3 shows an example of $\mathcal{O}$-connected polygonal curves s_1 and s_2 whose intersection is disconnected. Below we list basic properties of $\mathcal{O}$-connected sets.

(1) Every $\mathcal{O}$-connected set is $\mathcal{O}$-convex and connected. (The reverse does not hold.)
(2) Let $\mathcal{H}$ be an $\mathcal{O}$-hyperplane, and $P \subseteq \mathcal{H}$. Then P is $\mathcal{O}$-connected if and only if it is $\mathcal{O}(\mathcal{H})$-connected.
(3) The intersection of an $\mathcal{O}$-connected set with any $\mathcal{O}$-flat is $\mathcal{O}$-connected.
(4) Let P be a connected subset of $\mathcal{R}^d$. P is $\mathcal{O}$-connected if and only if for any $\mathcal{O}$-flat η, $P \cap \eta$ is connected.
(5) The set of all $\mathcal{O}$-connected sets in $\mathcal{R}^d$ is aligned.
(6) This property holds only for orthogonal convexity. Let P be an $\mathcal{O}_\perp$-connected set, and $\mathcal{H}$ be an $\mathcal{O}_\perp$-hyperplane. Then the perpendicular projection of P onto $\mathcal{H}$ is $\mathcal{O}_\perp(\mathcal{H})$-connected.

The important property of $\mathcal{O}$-connected sets that makes them similar to usual convex sets is simple connectedness. Unfortunately, we have proved this property only for the planar and three-dimensional cases.

Conjecture 4.1 (The Connectedness Conjecture) *If the set of $\mathcal{O}$-line orientations $\mathcal{O}^1$ is not empty, then every $\mathcal{O}$-connected set is simply connected.*

4.4 $\mathcal{O}$-stairlines

In this section we introduce the notion of $\mathcal{O}$-stairlines, which play the same role for restricted-orientation convexity as lines do for usual convexity.

Definition 4.3 ($\mathcal{O}$-stairlines) *An $\mathcal{O}$-*stairline *is an $\mathcal{O}$-connected curve. An $\mathcal{O}$-*stairsegment *is an $\mathcal{O}$-connected curvilinear segment* (see Fig. 3).

We denote an $\mathcal{O}$-stairline by the letter s, and a stairsegment with endpoints p and q by $s[p,q]$. To describe $\mathcal{O}$-stairlines in terms of $\mathcal{O}$-ranges, we introduce the notion of the *span* of a curve. The span of c is the set of orientations of all lines that intersect c in at least two points: $\mathit{span}(c) = \{\overline{(p,q)} \mid p,q \in c\}$.

The next theorem gives us a convenient description of $\mathcal{O}$-stairlines and $\mathcal{O}$-stairsegments.

Theorem 4.2 *A curve (curvilinear segment) is an $\mathcal{O}$-stairline ($\mathcal{O}$-stairsegment) if and only if its intersection with every $\mathcal{O}$-hyperplane is connected.*

The following two properties of $\mathcal{O}$-stairlines are corollaries of the above theorem.

(1) Let $s[p,q]$ be an $\mathcal{O}$-stairsegment and $\overline{\theta}$ be an $\mathcal{O}$-range such that $\overline{(p,q)} \in \overline{\theta}$. Then $\mathit{span}(s[p,q]) \subseteq \overline{\theta}$.
(2) Let $s[p,q]$ be an $\mathcal{O}$-stairsegment and η be an $\mathcal{O}$-flat such that $p,q \in \eta$. Then $s[p,q] \subseteq \eta$.

Now we present the main theorem of this subsection, which describes sufficient and necessary conditions for a curve to be a stairline in terms of the span of the curve.

Theorem 4.3 (The Span Theorem) *A curve (curvilinear segment) is an $\mathcal{O}$-stairline ($\mathcal{O}$-stairsegment) if and only if its span is completely contained within some $\mathcal{O}$-range.*

Below we state three corollaries of the Span Theorem that allow us to combine a long $\mathcal{O}$-stairsegment from several small pieces. (For the planar case, these properties were stated in [8].)

(1) Let p, q, and r be points in $\mathcal{R}^d$, and there exists an $\mathcal{O}$-range $\overline{\theta}$ such that $\overline{(p,q)} \in \overline{\theta}$ and $\overline{(q,r)} \in \overline{\theta}$. If $s[p,q]$ is an $\mathcal{O}$-stairsegment from p to q and $s[q,r]$ is an $\mathcal{O}$-stairsegment from q to r, then $s[p,q] \cup s[q,r]$ is an $\mathcal{O}$-stairsegment from p to r.
(2) Let v and w be points on the $\mathcal{O}$-stairsegment $s[p,q]$. If we replace the part of $s[p,q]$ between v and w by some other $\mathcal{O}$-stairsegment, the resulting curvilinear segment $s'[p,q]$ is still an $\mathcal{O}$-stairsegment.
(3) Let s be a polygonal line consisting of the edges $e_1, e_2, \ldots, e_n$. Then s is an $\mathcal{O}$-stairsegment if and only if there exists an $\mathcal{O}$-hyperrange $\overline{\Theta}$ such that $\bar{e}_i \in \overline{\Theta}$ for all $i \in [1..n]$. (Here $\bar{e}_i$ is the orientation of e_i.)

The next conjecture describes $\mathcal{O}$-connected sets via the notion of $\mathcal{O}$-stairsegments. The conjecture was proved for the planar case in [4]. In the higher-dimensional case, we found its proof only for the orthogonal convexity.

Conjecture 4.2 *A set P is $\mathcal{O}$-connected if and only if any two points p and q in P may be connected by an $\mathcal{O}$-stairsegment $s[p,q]$ such that $s[p,q] \subseteq P$.*

We have proved that a similar result holds for $\mathcal{O}$-convex sets and $\mathcal{O}$-convex curvilinear segments.

Theorem 4.4 *Let P be a connected set. P is $\mathcal{O}$-convex if and only if for all $p, q \in P$, there exists an $\mathcal{O}$-convex curvilinear segment $c[p,q]$ such that $c[p,q] \subseteq P$.*

4.5 $\mathcal{O}$-stairsurfaces and $\mathcal{O}$-halfspaces

An *$\mathcal{O}$-stairsurface* in a d-dimensional space $\mathcal{R}^d$ is an $\mathcal{O}$-convex surface. An *$\mathcal{O}$-halfspace candidate* is a set P whose intersection with any $\mathcal{O}$-line is either empty, or a ray, or a line. Now let Q be the closure of the complement of P: $Q = \text{closure}(\mathcal{R}^d - P)$. If both P and Q are $\mathcal{O}$-halfspace candidates, then P is called an $\mathcal{O}$-halfspace. The notion of $\mathcal{O}$-stairsurfaces corresponds to the notion of hyperplanes in usual geometry, and $\mathcal{O}$-halfspaces correspond to usual halfspaces. Unlike usual halfspaces, $\mathcal{O}$-halfspaces may not be connected. For example, the union of the right angles P_1 and P_2 on Fig. 4 is a disconnected $\mathcal{O}$-halfsapce.

Now we present elementary properties of $\mathcal{O}$-halfspaces.

(1) For any set of orientations $\mathcal{O}$, a usual halfspace is an $\mathcal{O}$-halfspace.
(2) Let $Q = \text{closure}(\mathcal{R}^d - P)$. If P is an $\mathcal{O}$-halfspace, then Q is also an $\mathcal{O}$-halfspace.
(3) A set P is an $\mathcal{O}$-halfspace ($\mathcal{O}$-halfspace candidate) if and only if for any $\mathcal{O}$-hyperplane $\mathcal{H}$, $P \cap \mathcal{H}$ is an $\mathcal{O}(\mathcal{H})$-halfspace ($\mathcal{O}(\mathcal{H})$-halfspace candidate).
(4) A set is an $\mathcal{O}$-halfspace ($\mathcal{O}$-halfspace candidate) if and only if it is the union of disjoint connected components such that each component is an $\mathcal{O}$-halfspace ($\mathcal{O}$-halfspace candidate), and no $\mathcal{O}$-line intersects any pair of components.

The next theorem shows the connection between $\mathcal{O}$-stairsurfaces and $\mathcal{O}$-halfspaces.

Theorem 4.5 (The Boundary Theorem for $\mathcal{O}$-halfspaces) *A set P is an $\mathcal{O}$-halfspace if and only if its boundary consists of $\mathcal{O}$-stairsurfaces, and no $\mathcal{O}$-line intersects any pair of these $\mathcal{O}$-stairsurfaces.*

The last result stated in this section, the Separation Conjecture, was proved in [4] for the planar case. It probably works in higher-dimensional spaces too, but we found a proof only for the case of the orthogonal $\mathcal{O}$-convexity in the three-dimensional space.

Conjecture 4.3 (The Separation Conjecture) *Let P be an $\mathcal{O}$-connected set, and p be a point outside P. Then there exists an $\mathcal{O}$-halfspace that completely contains P and does not contain p.*

A corollary of this result would be:

Corollary 4.1 (of the Separation Conjecture) *An $\mathcal{O}$-connected set is the intersection of all $\mathcal{O}$-halfspaces containing it.*

4.6 Characterizing $\mathcal{O}$-convex sets

Consider some arbitrary surface in $\mathcal{R}^d$. Intuitively, we wish to divide this surface into $\mathcal{O}$-convex regions. We call such regions *$\mathcal{O}$-stairfacets*. Formally, a *curvilinear facet* is an open connected subset of a surface together with its boundary, and an *$\mathcal{O}$-stairfacet* is an $\mathcal{O}$-convex curvilinear facet.

Definition 4.4 ($\mathcal{O}$-extremal points) *A point p is said to be an $\mathcal{O}$-extremal point of a surface if p is a point of support of this surface with respect to some $\mathcal{O}$-line.*

Theorem 4.6 *A curvilinear facet is $\mathcal{O}$-convex if and only if none of its interior points are $\mathcal{O}$-extremal.*

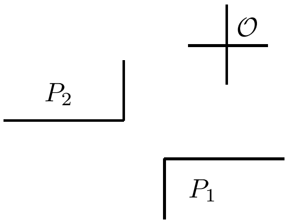

Fig. 4. $(P_1 \cup P_2)$ is an $\mathcal{O}$-halfspace

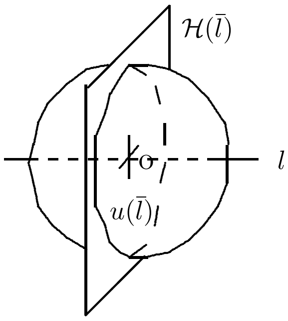

Fig. 5.

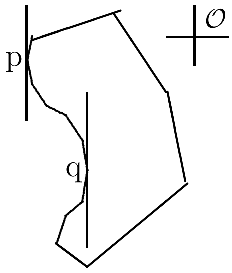

Fig. 6. $\mathcal{O}$-convexity point p and $\mathcal{O}$-concavity point q

Thus, we may view a surface as a collection of $\mathcal{O}$-stairfacets, whose boundaries consist of $\mathcal{O}$-extremal points. Pictorially, this view may be presented as follows. Consider a unit sphere U and some $\mathcal{O}$-line orientation $\bar{l}$. Let $\mathcal{H}(\bar{l})$ be the hyperplane through the center of U perpendicular to $\bar{l}$, and $u(\bar{l}) = U \cup \mathcal{H}(\bar{l})$ (see Fig. 5). It may be shown that $u(\bar{l})$ *is the set of points of support of U w.r.t. $\bar{l}$-oriented lines.* Thus if we draw $u(\bar{l})$ for every $\bar{l} \in \mathcal{O}$, we find all $\mathcal{O}$-extremal points of the sphere, and regions on the sphere bounded by $u(\bar{l})$'s are $\mathcal{O}$-stairfacets. Now consider some surface S. We define the *direction* of each point p of S as the direction of the normal vector to S at p (we assume that all normal vectors point to the "same side" of the surface), and consider the mapping from S to the sphere U such that each point of S is mapped onto the point of U with the same direction[1]. This mapping has two important properties:

(1) The image of every $\mathcal{O}$-extremal point of S is an $\mathcal{O}$-extremal point of U. (The reverse does not hold.)
(2) A curvilinear facet on the surface S is $\mathcal{O}$-convex if and only if the image of its interior is completely contained in some $\mathcal{O}$-stairfacet of U.

The second property means that intuitively we may envision S as the collection of $\mathcal{O}$-stairfacets that are preimages of the stairfacets of U.

Next consider a simply connected set P, and let S be its boundary. We may divide the $\mathcal{O}$-extremal points of the surface S into two groups. The first kind is the points where the supporting $\mathcal{O}$-lines lie in the exterior of P, and thus these $\mathcal{O}$-lines are also lines of support of the set P. These points are called *points of $\mathcal{O}$-convexity.* The second kind is the points where the supporting lines "touch" the surface from the interior of P. More formally, in some neighborhood of the extremal point the supporting $\mathcal{O}$-line lies within P. Such points are called *points of $\mathcal{O}$-concavity* (see Fig. 6).

Theorem 4.7 (The Boundary Theorem for $\mathcal{O}$-convex Sets) *A simply connected set P is $\mathcal{O}$-convex if and only if its boundary does not contain any points of $\mathcal{O}$-concavity.*

This theorem means that we may envision the boundary of a simply connected $\mathcal{O}$-convex set as a collection of $\mathcal{O}$-stairfacets, whose boundaries consist of points of $\mathcal{O}$-convexity. If $\mathcal{O}$ contains *all* orientations, that is if we consider the usual convexity, then $\mathcal{O}$-stairfacets become usual facets of a polytope, and the Boundary Theorem states that a polytope is convex if and only if all angles between adjacent facets are convex.

4.7 Strong $\mathcal{O}$-convexity

A set P is said to be *strongly $\mathcal{O}$-convex* if, for every two points $p, q \in P$, every $\mathcal{O}$-stairsegment joining p and q is contained in P. The union of all stairsegments joining p and q is called an *$\mathcal{O}$-parallelotope* and denoted by $\mathcal{O}$-$\|[p,q]$. The next two propositions give us a convenient description of $\mathcal{O}$-$\|[p,q]$.

(1) Let $\mathcal{O} \neq \emptyset$, and $\overline{\mathcal{H}} \in \mathcal{O}$. Consider $\overline{\mathcal{H}}$-oriented hyperplanes $\mathcal{H}_p$ and $\mathcal{H}_q$ such that $p \in \mathcal{H}_p$ and $q \in \mathcal{H}_q$. We define $\overline{\mathcal{H}}$-layer(p,q) as the set of all points between $\mathcal{H}_p$ and $\mathcal{H}_q$, including $\mathcal{H}_p$ and $\mathcal{H}_q$ themselves. Then $\mathcal{O}$-$\|[p,q] = \bigcap\{\overline{\mathcal{H}}\text{-layer}(p,q) \mid \overline{\mathcal{H}} \in \mathcal{O}\}$.
(2) Let $\overline{\theta}$ be the $\mathcal{O}$-range such that $(p,q) \in \overline{\theta}_{int}$. Let θ_p be the $\overline{\theta}$-oriented angle with the vertex p that contains q, and θ_q be the $\overline{\theta}$-oriented angle with the vertex q that contains p. Then $\mathcal{O}$-$\|[p,q] = \theta_p \cap \theta_q$.

[1]This mapping may not be a function, since if S is not smooth, some points of S have more than one direction.

Using the above description of $\mathcal{O}$-parallelotopes, we may prove the following properties of strongly $\mathcal{O}$-convex sets. (For the planar case, the first two properties were presented in [4].)

(1) P is strongly $\mathcal{O}$-convex if and only if the intersection of P and every $\mathcal{O}$-stairline is connected.
(2) P is strongly $\mathcal{O}$-convex if and only if P is a convex polytope all facets of which are $\mathcal{O}$-oriented.
(3) A flat is strongly $\mathcal{O}$-convex if and only if it is $\mathcal{O}$-oriented.
(4) Let $\mathcal{H}$ be an $\mathcal{O}$-hyperplane, and $P \subseteq \mathcal{H}$. Then P is strongly $\mathcal{O}$-convex if and only if it is strongly $\mathcal{O}(\mathcal{H})$-convex.
(5) The intersection of a strongly $\mathcal{O}$-connected set with any $\mathcal{O}$-flat is strongly $\mathcal{O}$-convex.
(6) If a set P is strongly $\mathcal{O}$-convex, it is also $\mathcal{O}$-connected.
(7) The intersection of strongly $\mathcal{O}$-convex sets is strongly $\mathcal{O}$-convex.
(8) For any $\mathcal{O}$, the set of all $\mathcal{O}$-convex sets is aligned.

Observe that both the empty set and the whole space is strongly $\mathcal{O}$-convex, and therefore it follows from the seventh property that strongly $\mathcal{O}$-convex sets form a convexity space.

5 Conclusion

We have presented four types of objects: $\mathcal{O}$-convex sets, $\mathcal{O}$-connected sets, $\mathcal{O}$-halfspaces, and strongly $\mathcal{O}$-convex sets. In the case when $\mathcal{O}^1$ is the set of all line orientations, $\mathcal{O}$-halfspaces become usual halfspaces, and the three other types become usual convex sets. We have shown that these four kinds of objects have a lot of similar properties. (However the proofs of the same properties are quite different for different objects.)

$\mathcal{O}$-convex and strongly $\mathcal{O}$-convex sets have many properties similar to the properties of usual convex objects and form abstract convexity spaces, and $\mathcal{O}$-halfspaces are similar to usual halfspaces. The main characteristic of convex sets and halfspaces that we have lost in the generalization to $\mathcal{O}$-convexity is *connectedness*: convex sets are always connected, while $\mathcal{O}$-convex sets may be disconnected.

References

[1] E. Fink and D. Wood. Restricted-orientation geometry in higher-dimensional spaces. Forthcoming.

[2] R. H. Guting. Stabbing C-oriented polygons. *Information processing letters*, 16, 1983.

[3] R. H. Guting. Dynamic C-oriented polygonal intersection searching. *Information and control*, 63, 1984.

[4] G. J. E. Rawlins. *Explorations in restricted-orientation geometry.* PhD thesis, Department of Computer Science, University of Waterloo, Waterloo, 1987. Research report CS-87-57.

[5] G. J. E. Rawlins, S. Schuierer, and D. Wood. *Visibility, skulls, and kernels in convexity spaces.* Department of Computer Science, University of Waterloo, Waterloo, 1989. Research report CS-89-48.

[6] G. J. E. Rawlins, S. Schuierer, and D. Wood. Convexity, visibility, and orthogonal polygons. *Contemporary mathematics*, 119, 1991.

[7] G. J. E. Rawlins and D. Wood. Restricted-orientation convex sets. *Information sciences*, 54, 1991.

[8] S. Schuierer. *On generalized visibility.* PhD thesis, Universitat Freiburg, Germany, 1991.

[9] P. Widmayer, Y. F. Wu, M. D. F. Schlag, and C. K. Wong. On some union and intersection problems for polygons with fixed orientations, *Computing*, 36, 1986.