International Journal of Computer Vision 47(1/2/3), 161–170, 2002

Stereo with Oblique Cameras

TOMÁŠ PAJDLA*
Center for Machine Perception, Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, Karlovo nám. 13, 121-35 Praha 2, Czech Republic
pajdla@cmp.felk.cvut.cz

Abstract. Mosaics acquired by pushbroom cameras, stereo panoramas, omnivergent mosaics, and spherical mosaics can be viewed as images taken by non-central cameras, i.e. cameras that project along rays that do not all intersect at one point. It has been shown that in order to reduce the correspondence search in mosaics to a one-parametric search along curves, the rays of the non-central cameras have to lie in double ruled epipolar surfaces. In this work, we introduce the oblique stereo geometry, which has non-intersecting double ruled epipolar surfaces. We analyze the configurations of mutually oblique rays that see every point in space. These configurations, called oblique cameras, are the most non-central cameras among all cameras. We formulate the assumption under which two oblique cameras posses oblique stereo geometry and show that the epipolar surfaces are non-intersecting double ruled hyperboloids and two lines. We show that oblique cameras, and the corresponding oblique stereo geometry, exist and give an example of a physically realizable oblique stereo geometry. We introduce linear oblique cameras as those which can be generated by a linear mapping from points in space to camera rays and characterize those collineations which generate them. We show that all linear oblique cameras are obtained by a collineation from one example of an oblique camera. Finally, we relate oblique cameras to spreads known from incidence geometries.

Keywords: non-central camera, stereo panorama, epipolar geometry, spread

1. Introduction

A complete theory as well as computational techniques have been elaborated in order to reconstruct three-dimensional scenes from images acquired by central cameras (Hartley and Zisserman, 2000). It is human nature to ask whether any useful multiview theory is restricted only to central cameras or if it can be extended by relaxing the requirement that all rays have to intersect at one point. Besides the pure intellectual curiosity, yet another motivation stems from the literature describing mosaics, which can often be viewed as images taken along rays that do not all intersect at one point.

For instance, Rademacher and Bishop (1998) introduced multiple center of projection images to facilitate image-based rendering. They required that the images were taken by a smoothly moving camera. They mentioned epipolar geometry between such images and pointed out that, in general, epipolar lines are replaced by epipolar curves. Gupta and Hartley (1997) analyzed linear pushbroom cameras, which are formed by a pencil of rays swept at a constant speed along a line that is perpendicular to the pencil. They proposed a linear pushbroom camera model and studied relative configurations of two such cameras. Epipolar geometry of two linear pushbroom cameras in a general position was defined and it was shown that for two rays to correspond, a cubic constraint has to be satisfied. Concentric mosaics, concentric symmetric panoramas, or circular panoramas (Peleg et al., 2001) are formed by rotating a linear camera—i.e. a pencil of rays—along a circle that

*This research was supported by the following projects: GACR 102/01/0971, MSMT Kontakt 2001/09, and MSM 212300013.

is tangent to the pencil. Peleg et al. (2001) proposed a realization of a concentric panorama using a specially shaped mirror in order to capture stereo panoramic images of moving scenes. Shum et al. (1999) proposed a non-central camera called an omnivergent sensor in order to reconstruct scenes with minimal error. We have shown recently that some non-central cameras provide a generalization of epipolar geometries (Pajdla, 2001a). The same idea was independently introduced by Seitz (2001) and demonstrated on existing mosaicing techniques and some new mosaicing techniques were proposed.

In this paper, we show an interesting generalization of the epipolar geometry. Our generalization leads to non-central cameras, which have pairwise oblique rays. We show that such cameras can form a stereo geometry with double ruled epipolar surfaces such that the sets of mutually oblique rays are partitioned into disjoint subsets of rays. Each subset is either a line or a set of lines lying in a double ruled quadric. All rays in one of the subsets form two one-parametric families of rays. Similarly, each classical epipolar plane is spanned by two one-parametric pencils of rays, one pencil from each camera. It is important that each family is independent and one-parametric. Only then can the correspondence problem be solved independently in each subset by finding a 1D mapping between the two families of rays. The ordering along the parameter can be used in the same way as the ordering along epipolar lines.

The main contribution of the paper is theoretical. The geometry of two non-central cameras with non-intersecting epipolar surfaces—*oblique stereo geometry*—is introduced and studied. However, it is also shown that the oblique stereo geometry can be realized in practice with the use of a catadioptric camera system, which was already proposed and used for mosaicing (Nayar and Karmarkar, 2000).

The structure of the paper is the following. Notations and concepts are given in Section 2. Section 3 introduces the notion of visibility closure to interpret the classical epipolar geometry of two central cameras in a new way, which allows for a generalization. Section 4 describes the oblique camera. In Section 5, the geometry of two oblique cameras satisfying natural requirements is derived. The structure of visibility closures and the existence of entities corresponding to epipolar planes and epipoles are discussed. Section 6 gives an example of an oblique stereo-geometry and describes how to realize it in practice. Linear oblique cameras are introduced and characterized in Section 7. Section 8 summarizes the work.

The concept of visibility closure was, for the first time, introduced in our conference paper (Pajdla, 2001a). The structure of visibility closures of oblique cameras was first described in another conference paper (Pajdla, 2001b). Here we integrate our preliminary results into a complete work. We have also improved some concepts and simplified the proofs so that it would be easier to generalize the concept of visibility closures to other non-central cameras. All material regarding linear oblique cameras is entirely new.

2. Notations and Concepts

Let $\exp A$ denote the set of all subsets of a set A. $A \subseteq B$ means "A is a subset of B" and $A \subset B$ means "A is a proper subset of B". The three-dimensional real projective space will be denoted by $\mathbb{P}^3$ througout the paper. Space $\mathbb{P}^3$ consists of a set of points, a set of lines, and an incidence relation "$\circ$" satisfying the axioms of three-dimensional projective space (Mihalek, 1972). Points are denoted by upper case letters, for instance X. Their representatives in $\mathbb{R}^4$ are denoted as upper case bold letters $\mathbf{X}$. Lines in $\mathbb{P}^3$ are denoted by lower case letters, e.g. l. When we refer to a line generated by points X and Y we write $X \vee Y$ or, in the linear representation, $span(\mathbf{XY})$. By $k \wedge l$ we mean the set of all points that are incident with lines k and l. We understand the lines as sets of points and therefore "$\circ$" also means "$\in$". We say that two lines intersect if there is exactly one point incident with both. We say that a point lies on a line if it is incident with the line.

By a camera we understand a subset of the set of lines in $\mathbb{P}^3$, Fig. 1(b). We often refer to the lines of a camera as to rays. In this paper, the term ray will mean a non-oriented line and it will be used whenever we want to stress that a line is a line of a camera.

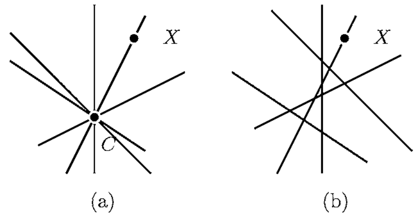

Figure 1. (a) A camera is a set of rays in $\mathbb{P}^3$. (b) A central camera is a set of all rays incident with one point—the projection center C.

Our notion of a general camera does not impose any constraint on the rays of the camera. By central camera, Fig. 1(a), we understand a set of rays in $\mathbb{P}^3$ that are all incident with one point, the center of projection. By *imaging* we mean a mapping that assigns rays from a camera to points in $\mathbb{P}^3$. Thus, for each point X in $\mathbb{P}^3$, other than the projection center C, a unique ray is assigned to X by choosing the line $X \vee C$. No unique ray can be assigned this way to the center of projection but all other points are imaged exactly once.

We say that point X is seen by camera $\mathcal{C}$ if there is a ray $r \in \mathcal{C}$ such that $X \circ r$. We say that a point is seen once by a camera if there exists just one ray from the camera that is incident with the point. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be two cameras, i.e. two sets of rays in $\mathbb{P}^3$. We say that a ray r from camera $\mathcal{C}_1$ is visible from camera $\mathcal{C}_2$ if each point of r is incident with a ray from $\mathcal{C}_2$.

We say that a quadric is proper if it is not degenerated, i.e. it is represented in $\mathbb{R}^4$ by a full rank matrix.

3. Visibility Closures

Figure 2 shows a diagram of an epipolar plane of two central cameras. The first resp. the second camera, denoted $\mathcal{C}_1$ resp. $\mathcal{C}_2$, is formed by the set of all lines in $\mathbb{P}^3$ passing through point C_1 resp. C_2. Let us study the visibility of rays from one camera by the rays of the other camera.

Let us choose arbitrary ray $r_1 \in \mathcal{C}_1$, $r_1 \neq e$. The set U_2 of all rays from $\mathcal{C}_2$ intersecting r_1 spans the epipolar plane. Symmetrically, the set $U_1 \subseteq \mathcal{C}_1$ of rays intersecting some $r_2 \in U_2$, $r_2 \neq e$, spans the same epipolar plane because both C_1 and C_2 lie in the epipolar plane. Any epipolar plane can be thus viewed as the set of points that are double-covered by the set of lines $U = U_1 \cup U_2$, $U_1 \subseteq \mathcal{C}_1$ and $U_2 \subseteq \mathcal{C}_2$, such that each line from U_1 is visible from U_2 and vice versa. We define *visibility closure* to formalize the above concept.

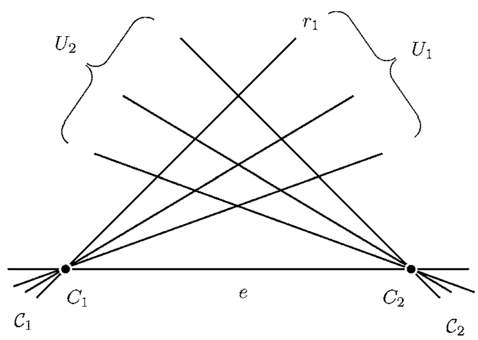

Figure 2. Epipolar geometry of two central cameras, see text.

Definition 1 (Visibility closure). We say that set $U = U_1 \cup U_2$, $U_1 \subseteq \mathcal{C}_1$, $U_2 \subseteq \mathcal{C}_2$ is a visibility closure of rays in cameras $\mathcal{C}_1$, $\mathcal{C}_2$ iff it holds that

$$\begin{aligned} &\forall k \in U_1, \forall X \in \mathbb{P}^3, X \circ k : \exists l \in U_2 \quad \text{such that } X \circ l \\ &\forall k \in U_2, \forall X \in \mathbb{P}^3, X \circ k : \exists l \in U_1 \quad \text{such that } X \circ l \end{aligned} \tag{1}$$

The empty set is a visibility closure.

Let us again look at the structure of visibility closures of two central cameras. First, there is one one-line closure consisting of the line e joining camera centers C_1, C_2 in Fig. 2. The line e contains exactly those points in space that project into epipoles and which cannot be reconstructed by intersection of rays from cameras $\mathcal{C}_1$, $\mathcal{C}_2$. Secondly, there is a fan of planar closures—epipolar planes—all intersecting in line e. Finally, the largest closure is formed by $\mathcal{C}_1 \cup \mathcal{C}_2$. The closures are partially ordered by the inclusion: line $\{e\} \subseteq$ rays spanning epipolar planes $\subseteq \mathcal{C}_1 \cup \mathcal{C}_2$.

We will show that the notion of visibility closure is well-defined even for other arrangement of rays than the one provided by two central cameras. We can show that our visibility closure is the set theoretical closure (Hazewinkel, 1995) under the following assumption.

Assumption 1 (about unique visibility). Let it hold that for each of the two cameras all points in $\mathbb{P}^3$ that are incident with a ray of one of the cameras are incident with exactly one ray from both of them.

Let us define the operation that creates visibility closures from sets of rays and show that it is the set-theoretical closure.

Definition 2. Let $\mathcal{C}_1$, $\mathcal{C}_2$ be two cameras satisfying Assumption 1. Let us define the operation $\overline{}$: $\exp(\mathcal{C}_1 \cup \mathcal{C}_2) \rightarrow \exp(\mathcal{C}_1 \cup \mathcal{C}_2)$ that assigns to a set of lines $L \subseteq \mathcal{C}_1 \cup \mathcal{C}_2$ the smallest (ordered by the set-theoretical inclusion) visibility closure $U \in \exp(\mathcal{C}_1 \cup \mathcal{C}_2)$ containing L. Let furthermore $\bar{\emptyset} = \emptyset$.

To show that the above definition is correct, we need to prove the following lemma.

Lemma 1. *Let A, B be two visibility closures in cameras $\mathcal{C}_1$, $\mathcal{C}_2$ satisfying Assumption* 1. *Then, $A \cap B$ and $A \cup B$ are also visibility closures in $\mathcal{C}_1$, $\mathcal{C}_2$.*

Proof: Empty $A \cap B$ is a visibility closure by the definition. A nonempty $A \cap B$ contains at least one line. Let WLOG $m \in A \cap B \cap \mathcal{C}_1$. Since m is in A, which is a visibility closure, all points of m are incident with a line from $A \cap \mathcal{C}_2$. Let us call the set of those lines L. Since m is also in B, which is a visibility closure, all points of m are incident with a line from $B \cap \mathcal{C}_2$. Let us call the set of those lines M. It follows from Assumption 3 that there is exactly one line in $\mathcal{C}_2$ incident with each point of m. Therefore $L = M$ and since $L \subset A \cap \mathcal{C}_2$ and $L = M \subset B \cap \mathcal{C}_2$, then $L \subset A \cap B \cap \mathcal{C}_2$. Thus, every point on every line from $A \cap B \cap \mathcal{C}_1$ is incident with a line from $A \cap B \cap \mathcal{C}_2$ and vice versa. The set $A \cap B$ is a visibility closure in $\mathcal{C}_1$, $\mathcal{C}_2$.

The set $A \cup B$ satisfies (1) because both A and B satisfy (1). □

It follows from Definition 1 that $\emptyset$ is a visibility closure. From Assumption 1 it follows that $\mathcal{C}_1 \cup \mathcal{C}_2$ is also a visibility closure. Thus, for each $L \subseteq \mathcal{C}_1 \cup \mathcal{C}_2$ there is a visibility closure—$\mathcal{C}_1 \cup \mathcal{C}_2$—containing L. Let us take $\mathcal{V}_L$, the set of all visibility closures that contain L. The set $\bigcap_{\forall V \in \mathcal{V}_L} V$ contains L, is smaller or equal to all elements in $\mathcal{V}_L$, and is a visibility closure according to Lemma 1. Thus $\overline{}$ is well defined on $\mathcal{C}_1 \cup \mathcal{C}_2$.

Seeing that operation $\overline{}$ is correct, we are at the position to formulate and prove the following theorem.

Theorem 1. *The operation $\overline{}$ is the set theoretical closure in $\mathcal{C}_1$, $\mathcal{C}_2$.*

Proof: We have to show (Hazewinkel, 1995) that (1) $\overline{A \cup B} = \bar{A} \cup \bar{B}$, (2) $A \subseteq \bar{A}$, (3) $\bar{\emptyset} = \emptyset$, (4) $\bar{\bar{A}} = \bar{A}$, for all $A, B \subseteq \mathcal{C}_1 \cup \mathcal{C}_2$. Properties (2), (3), and (4) follow trivially from Definition 2. Let us show that (1) holds.

First of all, $D \subseteq E \Rightarrow \bar{D} \subseteq \bar{E}$ for each $D, C \subseteq \mathcal{C}_1 \cup \mathcal{C}_2$ because $\mathcal{V}_E \subseteq \mathcal{V}_D$ and all $V \in \mathcal{V}_E$ contain D, where $\mathcal{V}_D$ resp. $\mathcal{V}_E$ is the set of all visibility closures containing D resp. E. It is clear that the smallest set from $\mathcal{V}_E$ is greater or equal to the smallest set from $\mathcal{V}_D$.

It holds that $A \subseteq A \cup B$ and $B \subseteq A \cup B$ and therefore it follows from the previous paragraph that $\bar{A} \subseteq \overline{A \cup B}$ and $\bar{B} \subseteq \overline{A \cup B}$ thus yielding $\bar{A} \cup \bar{B} \subseteq \overline{A \cup B}$.

The set $\bar{A} \cup \bar{B}$ is a visibility closure by Lemma 1. It contains both A and B and therefore it is a visibility closure containing $A \cup B$. The set $\overline{A \cup B}$ is, according to Definition 2, the smallest visibility closure containing $A \cup B$ and therefore $\overline{A \cup B} \subseteq \bar{A} \cup \bar{B}$. □

The above theorem justifies the choice of the name "visibility closure". The visibility closure of any A can be obtained as the set-theoretical closure $\overline{}$ of some set of lines in $\mathcal{C}_1 \cup \mathcal{C}_2$, e.g. as $\bar{A}$.

Visibility closures are partially ordered by the set-theoretical inclusion. We say that a nonempty visibility closure is *atomic* if it does not contain any smaller nonempty closure. The following lemma about atomic closures will be necessary for proving an important theorem later.

Lemma 2. *Let l be a ray in a pair of cameras satisfying Assumption* 1. *Then $\bar{l}$ is atomic.*

Proof: The closure $\bar{l}$ is not empty because it contains l. If $\bar{l}$ is not atomic, then there is a visibility closure A, $\emptyset \subset A \subset \bar{l}$. Either $l \in A$ or $l \notin A$. If $l \in A$, then $\bar{l} \subseteq A$ since $\bar{l}$ is the smallest visibility closure containing l. However, $A \subset \bar{l}$ and thus $l \notin A$. Therefore $l \in \bar{l} \setminus A \subset \bar{l}$. Let us denote the set $\bar{l} \setminus A$ by B.

We ought to show that B is a visibility closure. First, let us show that no ray from A has a common point with a ray from B. Let WLOG $k \in A \cap \mathcal{C}_1$. Since A is a visibility closure, there is at least one $m \in \mathcal{C}_2$ for all $X \in \mathbb{P}^3$, $X \circ k$, such that $X \circ m$. By Assumption 1, through each point of a ray in A passes exactly one ray from $\mathcal{C}_2$. Therefore, all rays from $\mathcal{C}_2$ incident with X are in A and none of them are in B because $A \cap B = \emptyset$. Secondly, let WLOG $p \in B \cap \mathcal{C}_1$. The set $A \cup B = \bar{l}$ is a visibility closure and thus every point on p is on some $q \in A \cup B \cap \mathcal{C}_2$. Ray q is not in A since $p \in B$ has a point in common with q. Therefore, $q \in B \cap \mathcal{C}_2$. Thus all points on every ray from $B \cap \mathcal{C}_1$ are incident with a ray from $B \cap \mathcal{C}_2$ and vice versa, which makes B a visibility closure.

The fact that $B \subset \bar{l}$ is a visibility closure and contains l is in contradiction with $\bar{l}$ being the smallest visibility closure containing l. Therefore, there is no nonempty visibility closure $A \subset \bar{l}$. □

4. Oblique Cameras

The geometry of central cameras is characterized by the requirement that all camera rays intersect at one point. When generalizing to a non-central camera, there is a plethora possible constraints to impose on the rays of the camera. We may require all the rays to intersect a circle like it is in case of a circular panoramas (Peleg et al., 2001) or a line as it is for a pushbroom camera (Gupta and Hartley, 1997). Let us concentrate on cameras that *image all points in space exactly once*.

Definition 3 (Oblique camera). We say that set $\mathcal{C}$ is an *oblique camera* if

$$\forall X \in \mathbb{P}^3 \; \exists! l \in \mathcal{C} \quad \text{such that} \quad X \circ l \, . \tag{2}$$

The above requirement is sufficient to provide our camera with a geometrical structure that justifies the name "oblique camera".

Observation 1. Two rays of an oblique camera are either identical or oblique.

Proof: Let k, l be two rays of an oblique camera. Line k does not intersect line l because if k intersected l there would be a point in $\mathbb{P}^3$ imaged by two rays from one camera what contradicts (2). □

The oblique cameras are exactly on the opposite side of the spectrum of cameras to central cameras. While all rays of a central camera intersect at a projection center, there are no intersecting rays in any oblique camera.

The structure of an oblique camera is not completely fixed by the above requirement and it is also not clear whether there is any such camera. In what follows, we will first assume to have two different oblique cameras and will study their possible visibility closures. Then, we will show that oblique cameras with oblique stereo geometries exist.

5. Oblique Stereo Geometry

We will show that by adopting the following constraint on the relationship between the rays of two oblique cameras, interesting visibility closures are obtained.

Assumption 2 (about two oblique cameras). Let us assume that the rays of oblique cameras $\mathcal{C}_1 \neq \mathcal{C}_2$ are in such a configuration that for each three mutually distinct rays l_1, l_2, l_3 from $\mathcal{C}_1$ resp. $\mathcal{C}_2$ the following holds. If there is one ray in $\mathcal{C}_2$ resp. $\mathcal{C}_1$ that intersects l_1, l_2, l_3 simultaneously then all lines from $\mathbb{P}^3$, which intersect l_1, l_2, l_3 simultaneously, are in $\mathcal{C}_2$ resp. $\mathcal{C}_1$.

Let us show the structure of visibility closures for oblique cameras satisfying Assumption 2.

5.1. Structure of Visibility Closures

Visibility closures of oblique cameras satisfying Assumption 2 are subsets of the set of all lines in $\mathbb{P}^3$. Therefore, it will be useful to restate the following classical geometrical theorem (Hilbert and Cohn-vossen, 1999).

Theorem 2 (*Double ruled surfaces in* $\mathbb{P}^3$)**.** *Let k_1, k_2, k_3 be three mutually oblique lines in $\mathbb{P}^3$. Then, the set of all points that lie on all lines intersecting k_1, k_2, k_3 form a proper double ruled quadric. Lines k_1, k_2, k_3 lie in one ruling while the lines intersecting them lie in the other. Both rulings are generated by any three mutually distinct lines from the other ruling. Every line from one ruling intersects all the lines from the other ruling. Proper double ruled quadrics and planes are the only double ruled surfaces in three-dimensional real projective space. The proper double ruled quadrics are either a hyperboloid of one sheet or a hyperbolic paraboloid in Euclidean three-dimensional space.*

Proof: See Hilbert and Cohn-Vossen (1999) and Knarr (1995). □

Applying Theorem 2 to oblique cameras in the configuration satisfying Assumption 2 yields the following lemma.

Lemma 3. *An atomic visibility closure in two oblique cameras satisfying Assumption 2 is either a ray or a proper double ruled quadric.*

Proof: Let $U = U_1 \cup U_2$, where $U_1 \subseteq \mathcal{C}_1$, $U_2 \subseteq \mathcal{C}_2$ be an atomic visibility closure of a ray belonging to an oblique camera satisfying Assumption 2. Closure U is nonempty. Let WLOG $l \in U_1$. Then, either $l \in U_2$ or $l \notin U_2$.

Let $l \in U_2$. Then $\bar{l} = l$ according to Definition 1.

Let $l \notin U_2$, Fig. 3(a). Then, according to Definition 3, through all points of l passes a line from U_2. Let us take three different lines from U_2 that intersect l, Fig. 3(b). It follows from Observation 1 that they are mutually

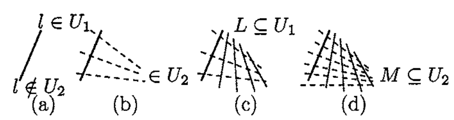

Figure 3. The illustration for the proof of Lemma 3, see text.

oblique. They are intersected by l and therefore due to Assumption 2, the set L of all lines that are simultaneously intersecting the three chosen lines is a subset of U_1, Fig. 3(c). The points incident with all lines L form a proper double ruled quadric as follows from Theorem 2, Fig. 3(d). Again thanks to Assumption 2, the set M of all lines that simultaneously intersect any three distinct lines from L is a subset of U_2 and therefore all points incident with lines in M form the same proper double ruled quadric according to Theorem 2. Therefore $\bar{l} = L \cup M$ since each point of l is intersected by a line from the proper double ruled quadric. □

Putting Theorem 1, Lemma 2, and Lemma 3 together yields the following theorem, which characterizes the structure of intersection closures of oblique cameras.

Theorem 3. *A visibility closure of rays in a pair of oblique cameras satisfying Assumption* 2 *is a union of mutually disjoint lines and proper double ruled quadrics.*

Proof: The visibility closure of a nonempty set L of rays is not empty. It follows from Theorem 1 that the visibility closure $\bar{L}$ is given by $\bigcup_{\forall l \in L} \bar{l}$. It follows from Lemma 2 that the closure $\bar{l}$ is atomic. Atomic closures are either lines of proper double ruled quadrics as follows from Lemma 3. Two distinct atomic closures have no line in common. □

5.2. Epipoles, Epipolar Lines, Epipolar Surfaces

Let $U = U_1 \cup U_2$, where $U_1 \in \mathcal{C}_1$ and $U_2 \in \mathcal{C}_2$ be a visibility closure in a pair of oblique cameras satisfying Assumption 2. We see from Theorem 3 that there are two kinds of atomic closures in U.

The atomic closures of the first kind are lines that are themselves visibility closures. No point on such lines can be reconstructed by intersecting a line from U_1 with a line from U_2. The line l, which is a visibility closure, is in U_1 as well as in U_2. Since there is only one line in each U_i that passes through a point in space, there is no line in U_1, other than l, that would intersect a line from U_2 at a point on l.

The atomic closures of the second kind are proper double ruled quadrics. Each of their points can be reconstructed as an intersection of a line from U_1 with a line from U_2.

We see that a pair of oblique cameras can be constructed as follows. There are two rulings on each proper double ruled quadric. Thus, on each quadric we assign lines from one ruling to one camera and the lines of the other ruling to the other camera. An oblique camera is fixed once we make thi assignment on each proper double ruled quadric. Many different cameras, which are not projectively equivalent, can be constructed by exchanging the rulings between the cameras on each of the quardrics. All lines in each ruling can be parameterized along a line from the other ruling since every line from one ruling intersects all the lines in the other.

Analogically to the classical epipolar plane, the set of rays emanating from one oblique camera and spanning a proper double ruled quadric is in a one-to-one correspondence with a line. Thus solving for a correspondence inside one quadric visibility closure of two oblique cameras, i.e. finding correspondences between the ray sets U_1 and U_2 in one U, amounts to finding a correspondence between points of two lines. It can be done exactly the same way as it is done for epipolar lines in a pair of central images.

While the notion of epipolar lines in central images has a symmetrical notion in oblique images, there is no equivalent notion for epipoles. An epipole in a central camera C_1 is the image of the center of a second camera C_2, i.e. the image of a point in space that is on more than on ray of C_2. There are no epipoles in oblique stereo geometry since every point is on exactly one ray of an oblique camera.

6. Example of Oblique Stereo Geometry

Theorem 3 characterizes nonempty visibility closures. However, is there any arrangement of lines in $\mathbb{P}^3$ such that they might form two oblique cameras satisfying Assumption 2? Figure 4 shows one such example. The figure shows two lines, l and k_∞, and a set of rotational hyperboloids of one sheet with axis l. The hyperboloids fill the space between l and k_∞. Line k_∞ lies in the plane at infinity and therefore it is, in a Euclidean space of the figure, depicted as a circle. One oblique camera is for example formed by the set of all rulings depicted by black lines while the other is depicted by gray lines. Both cameras contain the lines l and k_∞, which are the only one-line visibility closures of the arrangement.

All hyperboloids as well as the two lines form the set of quadrics $Q = \{Q(s) \mid s \in [0, 1] \subset \mathbb{R}\}$ in $\mathbb{P}^3$ with

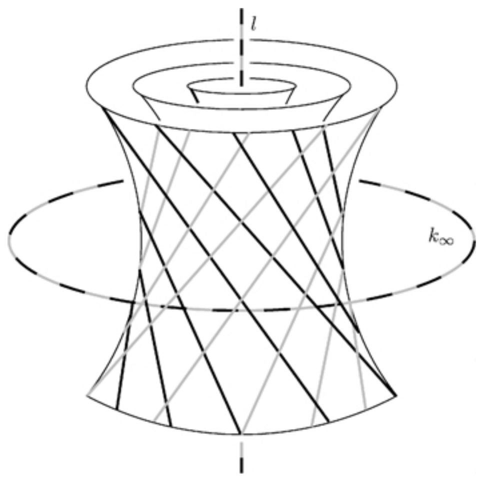

Figure 4. Two sets of lines—black and gray—that form oblique cameras satisfying Assumption 2.

$Q(s)$ defined as

$$\mathbf{X}^T \begin{pmatrix} s & & & \\ & s & & \\ & & s-1 & \\ & & & s-1 \end{pmatrix} \mathbf{X} = 0, \qquad (3)$$

where $\mathbf{X} \in \mathbb{R}^4 \backslash \mathbf{0}$ stands for a vector from the linear representation of $\mathbb{P}^3$. Quadrics $Q(s)$ defined by (3) are lines for $s = 0$ and $s = 1$, and double ruled rotational hyperboloids for $s \in (0, 1)$. For each point $X \in \mathbb{P}^3$ there is exactly one $s \in \mathbb{R}$ such that $Q(s)$ contains X.

The lines that form oblique cameras that coincide with the lines of rulings of the hyperboloids in Q can be written as

$$l_1 = span \begin{pmatrix} x & -y \\ y & x \\ z & w \\ w & -z \end{pmatrix}, \ l_2 = span \begin{pmatrix} x & -y \\ y & x \\ z & -w \\ w & -z \end{pmatrix} \qquad (4)$$

for all points $\mathbf{X} = (x, y, z, w)^T \in \mathbb{R}^4 \backslash \mathbf{0}$ (Pajdla, 2001b). The lines l_1, l_2 form the reguli of quadrics Q. The lines l_1 form camera $\mathcal{C}_1$, lines l_2 form camera $\mathcal{C}_2$.

All other arrangements of lines that are obtained from the above example by a collineation also form two oblique cameras satisfying Assumption 2 since collineations preserve incidence.

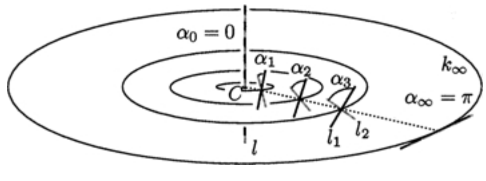

Figure 5. The arrangement of lines from Fig. 4 can be generated by rotating two lines along a set of concentric circles.

6.1. *Realization*

The arrangement of lines depicted on Fig. 4 can be realized by rotating two intersecting lines l_1, l_2 along concentric circles that lie in a plane perpendicular to the line l, see Fig. 5. The center of circles is at the point C where l intersects the plane. The angle α between l_1, l_2 is zero for the zero radius, and therefore l_1, l_2 merge together to one line l. It equals π for the infinite radius, and therefore the lines merge together to one line k_∞. If the angle is neither zero nor π, the set of lines obtained by rotating l_1, l_2 forms a rotational hyperboloid of one sheet with one ruling generated by l_1 and the other by l_2.

It is impossible to realize an oblique camera that would see whole $\mathbb{P}^3$ as in the above example. However, it is possible to realize a set of rays containing just the rays of a subset of the set of visibility closures. The following camera has been proposed by Nayar and Karmarkar (2000) in order to obtain complete spherical mosaics.

Let us have a conical mirror observed by a telecentric lens (Watanabe and Nayar, 1996) in the direction of the mirror rotation axis o as shown in Figs. 6 and 7. Let furthermore the mirror be such that all the parallel rays going through the telecentric lens are reflected to the rays that are perpendicular to the mirror axis thus

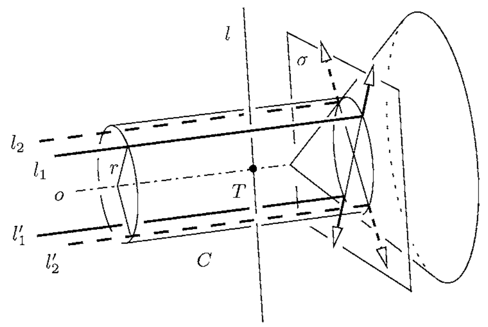

Figure 6. The two lines from Fig. 5 can be generated by selecting four pixels from an image of the scene reflected by a conical mirror.

Figure 7. Two sets of rays forming hyperboloidal closures can be realized by rotating catadioptric camera consisting from a telecentric optics and a conical mirror.

spanning plane σ perpendicular to o. As the radius r of the cylinder grows, the plane moves away from the tip of the conical mirror. If one had an infinitely large mirror as well as the telecentric lens, the moving plane would fill the whole half-space.

On each cylinder, denoted by C in Fig. 6, of parallel rays we can select four rays, e.g. l_1, l_1' and l_2, l_2', such that l_1, l_1' as well as l_2, l_2' reflect into the same line in σ. The angle between the lines in σ can be made arbitrary by the choice of the rays on C. Thus, by selecting four proper rays from each cylinder, one can obtain couples of rays lying in planes perpendicular to line o.

The subset of the set of visibility closures is then obtained by rotating the mirror around the line l, which intersects the axis o at the point T. There are only those closures that intersect the volume swept by the rotating mirror. Only the points, which are not contained in the swept volume, are seen.

7. Linear Oblique Cameras

The example given above is not the only example of an oblique camera. In mathematical literature, e.g. in Buekenhout (1995), oblique cameras are called spreads. A *spread* is a set of lines (i.e. sets of points) that partitions the points of the space. Spreads were studied in finite as well as infinite projective spaces in connection with translation planes (Knarr, 1995; Hirschfeld, 1998). Spreads in $\mathbb{P}^3$ are characterized by certain mappings over affine planes, which are called transversal (Knarr, 1995). There is a spread for each transversal mapping and there are many different spreads, which are not projectively equivalent.

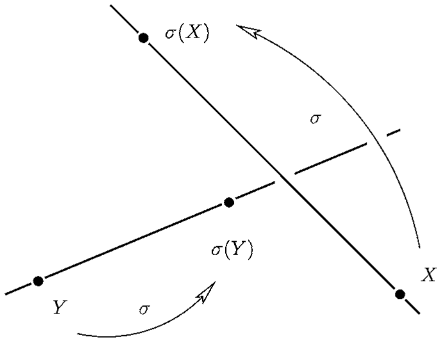

Figure 8. A spread can be generated by a mapping σ over set $\mathcal{S} \subseteq \mathbb{P}^3$ of points as the set of lines $\{X \vee \sigma(X) | \forall X \in \mathcal{S}\}$.

Central cameras assign lines to points in space by connecting the points with a projection center C. Imaging can be thus viewed as a mapping π from points of the space to the lines of the space $X \xrightarrow{\pi} X \vee C$. The mapping π is, in a linear representation of the projective space, a linear singular mapping that can be represented by a 4×4 real matrix P of rank three (Hartley and Zisserman, 2000). The right null-space of P corresponds to the projection center C. In other words, central cameras assign rays to points by linear mappings in the linear representation of the projective space.

Spreads can also be constructed such that to each point X in the space a unique line l is assigned by a linear mapping. We shall say that a mapping $\sigma : \mathcal{S} \subseteq \mathbb{P}^3 \rightarrow \mathbb{P}^3$, generates a spread over the set of points $\mathcal{S}$ if the set of lines $\{X \vee \sigma(X) \mid \forall X \in \mathcal{S}\}$ is a spread, see Fig. 8. It is interesting to look at those spreads, which are generated by collineations over $\mathbb{P}^3$.

Not all collineations in $\mathbb{P}^3$ generate spreads over $\mathbb{P}^3$. For instance, no collineation that has a fixed point, let say X, generates a spread because $X \vee \sigma(X) = X$, which is not a line. Thus no σ in $\mathbb{P}^3$, which generates a spread, has one-dimensional invariant subspace. On the other hand, σ must have enough two-dimensional invariant subspaces such that each point of the space is exactly in one such subspace. The intersection of any two two-dimensional invariant subspaces has to be the zero vector because otherwise there would be a one-dimensional invariant subspace. Therefore, if a line l is generated as $l = X \vee \sigma(X)$ for some $X \in \mathbb{P}^3$, then it holds true that $Y \circ l \Rightarrow \sigma(Y) \circ l$ for all Y incident with l, see Fig. 9.

The following theorem characterizes collineations that generate spreads over $\mathbb{P}^3$ in terms of the matrices of the corresponding linear mappings over $\mathbb{R}^4$.

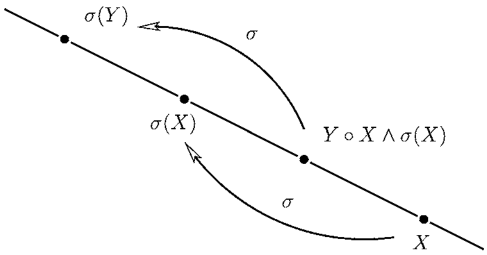

Figure 9. Collineation σ maps every point Y from the line $X \vee \sigma(X)$ back to the line.

Theorem 4. *A collineation* $\sigma : \mathbb{P}^3 \rightarrow \mathbb{P}^3$ *represented by the matrix* $\mathsf{S} \in \mathbb{R}^{4\times 4}$ *generates a spread over* $\mathbb{P}^3$ *if and only if*

$$\mathsf{S} \approx \begin{pmatrix} \alpha & 1 & & \\ -1 & \alpha & & \\ & & \alpha & 1 \\ & & -1 & \alpha \end{pmatrix} \qquad (5)$$

for some $\alpha \in \mathbb{R}$, *where* $\approx$ *stands for the matrix similarity* [2].

Proof: A line has to be assigned to all points of space and therefore S must be regular. Each spread contains at least two non-intersecting lines. The lines are mapped to themselves by σ. Therefore, a basis can be chosen from subspaces corresponding to the two lines in which S is block diagonal. The basis consists of four vectors chosen such that two vectors are in either line. By requiring that for each $l = X \wedge \sigma(X)$ holds true that $Y \circ l \Rightarrow \sigma(Y) \circ l$ for all Y incident with l, see Fig. 9, the collineation σ is obtained in the form (5). See Pajdla (2001c) for more details. □

As pointed out by Kohout (2001), the spread generated by (5) does not depend on the parameter α. The parameter only affects the action of the collineation inside the lines of the spread. The spread itself, i.e. the set of lines, remains the same for all values of α. It is therefore natural to formulate the following theorem.

Theorem 5. *A spread generated over* $\mathbb{P}^3$ *by a collineation with the matrix similar to* (5) *for some* $\alpha \in \mathbb{R}$ *can be generated over* $\mathbb{P}^3$ *by the collineation with the matrix similar to* (5) *for* $\alpha = 0$.

Proof: Denote the matrix of the collineation that generates the spread by S. Then, by Theorem 4, there exists a basis $\mathcal{B}$, in which the generating matrix can be written as (5) for some real α. Denote the generating matrix as $\mathsf{S}_\mathcal{B}(\alpha)$. It is enough to show that for each $X \in \mathbb{P}^3$, and any $\alpha \in \mathbb{R}$ it holds true that $\mathbf{X}$, $\mathsf{S}_\mathcal{B}(0)\mathbf{X}$, $\mathsf{S}_\mathcal{B}(\alpha)\mathbf{X}$ are of rank two. For any given real α, we can find coefficients α, 1, -1, which are not all zero, such that $\alpha\mathbf{X} + \mathsf{S}_\mathcal{B}(0)\mathbf{X} - \mathsf{S}_\mathcal{B}(\alpha)\mathbf{X} = 0$. □

Note that when transforming $X \in \mathbb{P}^3$ into $X' \in \mathbb{P}^3$ by a collineation P as $\mathbf{X}' = \mathsf{P}\mathbf{X}$, the line $span(\mathbf{X}\ \mathsf{S}\mathbf{X})$ is transformed into $span(\mathbf{X}', \mathsf{PSP}^{-1}\mathbf{X}')$ because

$$\begin{aligned} span(\mathbf{X}, \mathsf{S}\mathbf{X}) &= span(\mathsf{P}^{-1}\mathbf{X}', \mathsf{SP}^{-1}\mathbf{X}') \\ &= \mathsf{P}^{-1} span(\mathbf{X}', \mathsf{PSP}^{-1}\mathbf{X}'), \end{aligned} \qquad (6)$$

which is equivalent to transforming the points of $\mathbb{P}^3$ by a collineation P, by which S goes to PSP^{-1}.

We can see that Theorem 5 is important because it says that all spreads generated over $\mathbb{P}^3$ by collineations can be obtained by a collineation from the spread generated over $\mathbb{P}^3$ by (5) for $\alpha = 0$. By comparing the action of the matrix (5) for $\alpha = 0$ with the lines (4), we clearly see that all linear oblique cameras can be obtained by a collineation from the example described in Section 6.

8. Summary and Conclusions

We have generalized the notion of a camera by replacing the requirement that all the rays of a camera intersect at one point. Instead, we have introduced oblique cameras, for which it holds that no two rays from one camera intersect.

We have introduced the notion of visibility closures to interpret epipolar planes as one-parametric visibility closures of rays of two central cameras. When adopting Assumption 2 about the rays of two oblique cameras under scrutiny we arrived at generalized one-parametric closure spaces. We have shown that in such a case the generalization of epipolar planes leads to proper double ruled quadrics—in Euclidean space either a hyperboloid of one sheet or a hyperbolic paraboloid. Such surfaces can be parameterized by one parameter along a line, the same way as the rays in epipolar planes can be parameterized along epipolar lines. It was therefore natural to introduce the notion of oblique stereo geometry for the arrangement of non-intersecting rays allowing for one-parametric closures.

From the practical point of view, it is important that there is an arrangement of rays in $\mathbb{P}^3$ that allows for two oblique cameras with nontrivial one-parametric intersection closures. We have also proposed a technical

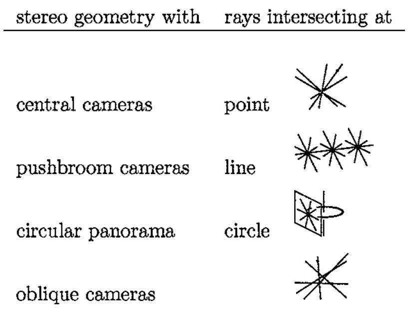

Figure 10. All stereo-geometries are "between" the central and oblique stereo geometry.

realization of an oblique camera pair having a subset of such visibility closures.

We mentioned that our oblique cameras are known in mathematical literature as spreads. Linear oblique cameras were defined to be analogical to central linear cameras. The characterization of collineations generating spreads was given.

We believe that oblique stereo geometry is an interesting theoretical concept because all non-central cameras lie between the central camera and oblique camerasy, Fig. 10. Similarly, all stereo geometries, namely those produced by various mosaicing techniques, lie between the classical epipolar geometry and oblique stereo geometry.

References

Buekenhout, E. (Ed.). 1995. *Handbook of Incidence Geometry: Buildings and Foundations*. Elsevier: Amsterdam, The Netherlands.

Friedberg, S.H., Insel, A.J., and Spence, L.E. 1997. *Linear Algebra* 3rd. ed. Prentice-Hall: Englewood Cliffs, NJ, USA.

Gupta, R. and Hartley, R.I. 1997. Linear pushbroom cameras. *PAMI*, 19(9):963–975.

Hartley, R. and Zisserman, A. 2000. *Multiple View Geometry in Computer Vision*. Cambridge University Press: Cambridge, UK.

Hazewinkel, M. (Ed.). 1995. *Encyclopaedia of Mathematics*, vol. 1. Kluwer Academic Publishers: Boston, MA.

Hilbert, D. and Cohn-Vossen, S. 1999. *Geometry and the Imagination*, 2nd ed. AMS Chelsea Publishing: Providence, RI.

Hirschfeld, J.W.P. 1998. *Projective Geometries over Finite Fields*, 2nd ed. Oxford Mathematical Monographs. Clarendon Press, Oxford, UK.

Knarr, N. 1995. *Translation Planes: Foundations and Construction Principles*, vol. 1611 of Lecture Notes in Mathematics. Springer: Berlin, Germany.

Kohout, V. 2001. Spreads generated by collineations over P_3—personal communication. Czech Technical University in Prague.

Mihalek, R.J. 1972. *Projective Geometry and Algebraic Structures*. Accademic Press: New York.

Nayar, S.K. and Karmarkar, A. 2000. 360 × 360 mosaics. In *CVPR '2000*, vol. 2, pp. 388–395.

Pajdla, T. 2001a. Epipolar geometry of some non-classical cameras. In *Computer Vision Winter Workshop*, Ljubljana, Slovenia, pp. 223–233. Slovenian Pattern Recorgnition Society.

Pajdla, T. 2001b. An example of oblique cameras with double ruled visibility closures. Czech Technical University, Prague, RR CTU–CMP–2001–10, CMP K333 FEE. Available at ftp://cmp.felk.cvut.cz/pub/cmp/articles/pajdla/Pajdla-TR-2001-10.pdf.

Pajdla, T. 2001c. Oblique cameras generated by collineations. Czech Technical University, Prague, RR CTU–CMP–2001–14, CMP K333 FEE. Available at ftp://cmp.felk.cvut.cz/pub/cmp/articles/pajdla/ Pajdla-TR-2001-14.ps.gz.

Pajdla, T. 2001. Stereo with oblique cameras. In *IEEE Workshop on Stereo and Multi-Baseline Vision*. Kauai, Hawaii, USA, pp. 85–91.

Peleg, S., Ben-Ezra, M., and Pritch, Y. 2001. Omnistereo: Panoramic stereo imaging. *PAMI*, 23(3):279–290.

Rademacher, P. and Bishop, G. 1998. Multiple-center-of-projection images. In *SIGGRAPH'98* Orlando, Fl. pp. 199–206, in *ACM SIGGRAPH*.

Seitz, S. 2001. The space of all stereo images. In *ICCV '2001*, Vancouver, Canada, pp. 26–33.

Shum, H., Kalai, A., and Seitz, S. 1999. Omnivergent stereo. In *ICCV '99*, Corfu, Greece, pp. 22–29.

Watanabe, M. and Nayar, S.K. 1996. Telecentric optics for computational vision. In *ECCV96*, vol. II, pp. 439–451.