Stereo with Oblique Cameras

Tomás Pajdla
Center for Machine Perception
Czech Technical University in Prague
Faculty of Electrical Engineering
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. We call such configurations oblique cameras. We argue that oblique cameras are important because they are the most non-central cameras among all cameras. We show that oblique cameras, and the corresponding oblique stereo geometry, exist and give an example of a physically realizable oblique stereo geometry.

1. Introduction

A complete theory[1] as well as computational techniques have been elaborated in order to reconstruct a three-dimensional scene form images acquired by central cameras [2]. 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 Bischop [9] 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 be replaced by epipolar curves. Gupta and Hartley [1] analyzed linear pushbroom cameras, which are formed by a pencil of rays swept in a constant speed along a line that is perpendicular to the pencil. They proposed a linear pushbroom camera model and studied relative configurations of such two 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 [8] are formed by rotating a linear camera – i.e. a pencil of rays – along a circle that is tangent to the pencil. Peleg, Ben-Ezra, and Pritch [8] proposed a realization of a concentric panorama using a specially shaped mirror in order to capture stereo panoramic images of moving scenes. Shum, Kalai, and Seitz [11] proposed a non-central camera called omnivergent sensor in order to reconstruct scenes with minimal error. It has been recently shown that some non-central cameras provide a generalization of epipolar geometries [7]. The idea was demonstrated on existing mosaicing techniques as well as some new mosaicing techniques were proposed [10].

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 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 the correspondence problem can 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 to enforce (piecewise) continuity of the mapping the same way as the ordering along epipolar lines is used.

[1]This research was supported by the EU FP5 project OMNIVIEWS No. 1999-29017, GACR 102/01/0971, MSMT projects KONTAKT 2001/09 and MSM 212300013.

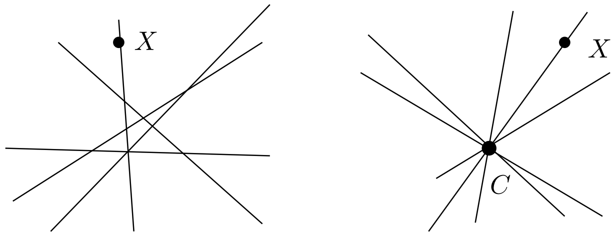

(a) (b)

Figure 1: (a) A camera is a set of rays in P_3. (b) A central camera is a set of all rays incident with one point – the projection center C

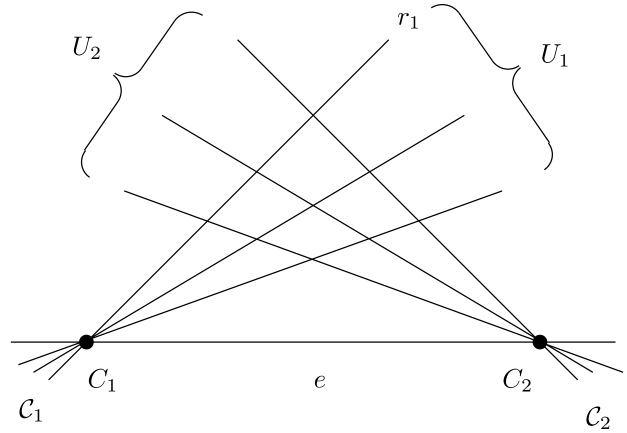

Figure 2: Epipolar geometry of two central cameras, see text

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 [6].

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 5.2 describes the oblique camera. In Section 5, the geometry of two oblique cameras satisfying natural requirements is derived. Both geometrical as well as algebraic characterizations of epipolar surfaces are given. Section 7 summarizes the work.

2. Notation and concepts

Let $expA$ denote the set of all subsets of a set A. A three-dimensional real projective space will be called *space* and denoted by P_3 through the paper. Space 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 [5].

By a camera we understand a subset of the set of lines in P_3, Figure 1(a). We often refer to the lines of a camera as to the 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.

Our notion of a general camera does not impose any constraint on the rays of the camera. By the central camera, Figure 1(b), we understand a set of rays in 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 P_3. Thus, for each point X in P_3, other than the projection center, a unique ray is assigned to X by choosing the line that is incident with the center and with X. 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 $X \neq C$ is seen by camera $\mathcal{C}$ iff 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 P_3. We say that a ray r from camera $\mathcal{C}_1$ is visible from camera $\mathcal{C}_2$ iff each point of r is incident with a ray from $\mathcal{C}_2$.

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 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 $U_1 \subset \mathcal{C}_1$ 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 $U = U_1 \cup U_2$, $U_1 \subset \mathcal{C}_1$ and $U_2 \subset \mathcal{C}_2$, such that each line from U_1 is visible from U_2 and vice versa. We define *visibility closures* to formalize the above concept.

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*

$$\begin{aligned} &\forall k \in U_1, \forall X \in P_3, X \circ k : \ \exists l \in U_2 \text{ such that } X \circ l \\ &\forall k \in U_2, \forall X \in P_3, X \circ k : \ \exists l \in U_1 \text{ such that } X \circ l \end{aligned} \quad (1)$$

The empty set is a visibility closure.

Two pencils of rays passing through C_1 resp. C_2 and spanning the epipolar plane containing ray r form the visibility closure of r.

We will show that the notion of visibility closure is well-defined even for more general arrangements of rays than the one provided by two classical central cameras. We can show that our visibility closure is the set theoretical closure [3] under the following assumption.

Assumption 1 (Unique visibility) *Let us assume that it holds for each two cameras that all points in space that are seen by one of the cameras are seen exactly once by 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 $\overline{\quad}$: $exp(\mathcal{C}_1 \cup \mathcal{C}_2) \to exp(\mathcal{C}_1 \cup \mathcal{C}_2)$ 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 $\overline{\emptyset} = \emptyset$.*

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.*

Proof. For $A \cap B$ being a visibility closure, it is enough to show that for any ray p from $A \cap B$, every point $X \circ p$ is seen by $k \in \mathcal{C}_1$ and $l \in \mathcal{C}_2$ such that $k, l \in A \cap B$. Since A is a visibility closure and Assumption 1 holds, $X \circ p \in A$ is seen by exactly one ray $k \in A \cap \mathcal{C}_1$ and by exactly one ray $l \in A \cap \mathcal{C}_2$. Since B is a visibility closure and Assumption 1 holds, $X \circ p \in B$ is seen by exactly one ray $m \in B \cap \mathcal{C}_1$ and by exactly one ray $n \in B \cap \mathcal{C}_2$. It follows from Assumption 1 that $k = m$ and $l = n$. Therefore $k \in A \cap B \cap \mathcal{C}_1$ and $l \in A \cap B \cap \mathcal{C}_2$.

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 $\overline{\quad}$ is well defined on $\mathcal{C}_1 \cup \mathcal{C}_2$.

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

Theorem 1 *The visibility closure is the set theoretical closure.*

Proof. We have to show [3] that (1) $\overline{A \cup B} = \overline{A} \cup \overline{B}$, (2) $A \subseteq \overline{A}$, (3) $\overline{\emptyset} = \emptyset$, (4) $\overline{\overline{A}} = \overline{A}$, 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, $D \subseteq E \Rightarrow \overline{D} \subseteq \overline{E}$ 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 intersection 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 $\overline{A} \subseteq \overline{A \cup B}$ and $\overline{B} \subseteq \overline{A \cup B}$ thus yielding $\overline{A} \cup \overline{B} \subseteq \overline{A \cup B}$.

The set $\overline{A} \cup \overline{B}$ 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 $\overline{A \cup B} \subseteq \overline{A} \cup \overline{B}$. □

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 $\overline{\quad}$ of some set of lines in $\mathcal{C}_1 \cup \mathcal{C}_2$, e.g. as $\overline{A}$. 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 nonempty closure.

Let us 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 Figure 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 maximal closure is formed by $\mathcal{C}_1 \cup \mathcal{C}_2$. The closures are partially ordered by the inclusion: line $\{e\} \subset$ rays spanning epipolar planes $\subset \mathcal{C}_1 \cup \mathcal{C}_2$.

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 we have 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 [8], or a line as it is for a pushbroom camera [1]. 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*

$$\forall X \in P_3 \ \exists! l \in \mathcal{C} \text{ such that } X \circ l \, . \tag{2}$$

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 P_3 imaged by two rays from one camera what contradicts (2). □

The oblique cameras are exactly on the opposite pole of the spectrum of cameras than the central cameras. While all rays of a central camera intersect at a projection center, there are no intersecting rays in an 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. Later 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 (Configuration of 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 rays from 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 P_3. Therefore, it will be useful to restate the following classical geometrical theorem [4].

Theorem 2 (Double ruled surfaces in P_3) *Let k_1, k_2, k_3 be three mutually oblique lines in P_3. Then, the set of all points that lie on all lines intersecting k_1, k_2, k_3 form a 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.*

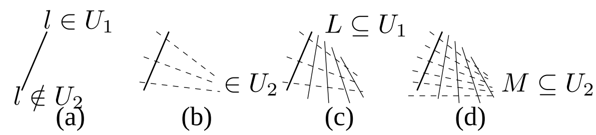

Figure 3: The illustration for the proof of Lemma 2, see text

Proof. See [4]. □

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

Lemma 2 *The smallest visibility closure of a ray 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 a visibility closure of a ray belonging to an oblique camera satisfying Assumption 2. 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 $l \notin U_2$, Figure 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, Figure 3(b). It follows from Observation 1 that they are mutually 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, Figure 3(c). The points incident with all lines L form a double ruled surface as follows from Theorem 2, Figure 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 double ruled surface according to Theorem 2. Therefore $\bar{l} = L \cup M$ since each point of l is intersected by a line from the double ruled surface. □

The following important observation follows from the fact that any line from one ruling intersects all the lines from the other ruling on a double ruled quadric.

Observation 2 *Let $U = U_1 \cup U_2$, $U_1 \in \mathcal{C}_1$, $U_2 \in \mathcal{C}_2$ be the smallest visibility closure of a ray in in a pair of oblique cameras $\mathcal{C}_1, \mathcal{C}_2$ satisfying Assumption 2. Then, the visibility closure $\bar{l}$ of any ray $l \in U$ equals U.*

Proof. It follows from Theorem 2 that every ray l in U_1 resp. U_2 can be used to generate U_2 resp. U_1 as the set of all rays that intersect l and two other rays from U_1 resp. U_2. □

The following consequence of the previous observation will be useful.

Observation 3 *The smallest visibility closures of two rays in a pair of oblique cameras satisfying Assumption 2 are either identical or disjoint.*

Proof. Two visibility closures U, V are either disjoint or have at least one ray, let say l, in common. Let them have l in common. Since $l \in U$ and $l \in V$ then by Observation 2 $U = \bar{l}$ and $V = \bar{l}$ and therefore $U = V$. □

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

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

Proof. A nonempty closure contains at least one line. It follows from the set theoretical closure properties of $\overline{\quad}$ that the smallest intersection closure of a set of rays is given by the union of the the smallest closures of each ray in the set. The smallest closure of one ray is atomic. Atomic closures are either lines of double ruled quadrics as follows from Lemma 2. It follows from Observation 3 that distinct atomic closures are disjoint. □

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. Their each point can be reconstructed as an intersection of a line form U_1 with a line from U_2.

Wee see that oblique cameras can be constructed as sets of lines from non-intersecting double ruled quadrics such that one camera contains all lines form one of the rulings on each double ruled quadric. Thus, an oblique camera is fixed once we choose one of the two rulings on each double ruled

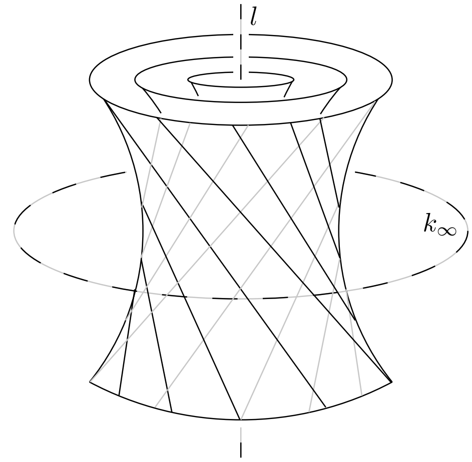

Figure 4: Two sets of lines — black and gray — that form oblique cameras satisfying Assumption 2

quadric. 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 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 line in central images has a symmetrical notion in oblique images, there is no equivalent notion for epipoles. An epipole in the camera $\mathcal{C}_1$ is the image of the center of the camera $\mathcal{C}_2$, i.e. the image of a point in space that is not uniquely imaged by $\mathcal{C}_2$. There are no epipoles in oblique stereo geometry since all points are imaged uniquely by any oblique camera,

6. Example of oblique stereo geometry

Theorem 3 characterizes nonempty visibility closures. However, is there any arrangement of lines in 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 k_∞ and l. 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 e.g. 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

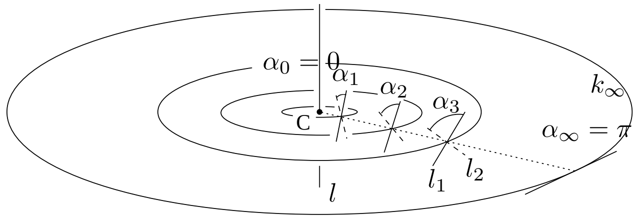

Figure 5: The arrangement of lines from Figure 4 can be generated by rotating two lines along a set of concentric circles

arrangement.

Besides the above example, all other arrangements of lines that are obtained from the example by a collineation also form two oblique cameras satisfying Assumption 2 since the incidence is preserved by collineations.

The arrangement of lines depicted on Figure 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 Figure 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, thus forming one line l. It equals π for the infinite radius, thus forming line k_∞. If the angle is not zero or π, the set of all lines obtained by rotating l_1, l_2 forms a rotational hyperboloid of one sheet with one ruling generated by rotating l_1 and the other by rotating l_2.

It is impossible to realize the oblique camera that would see whole P_3 from 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 in [6] in order to obtain a complete spherical mosaics.

Let us have a conical mirror observed by a telecentric lens [12] in the direction of the mirror rotation axis o as shown in Figure 6. 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 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, e.g. C, of parallel rays we can select four rays, e.g. l_1, l_1' and l_2, l_2' in Figure 6, such that l_1, l_1' as well as l_2, l_2' reflect into two rays lying on 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 certain four 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. 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 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. We have adopted a further assumption about the rays of two oblique cameras under scrutiny in order to arrive at generalized one-parametric closure spaces. We have shown that, under the accepted Assumption 2, the generalization of epipolar planes leads to

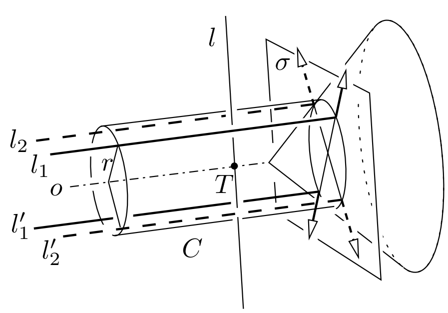

Figure 6: The two lines from Figure 5 can be generated by selecting four pixels from an image of the scene reflected by a conical mirror

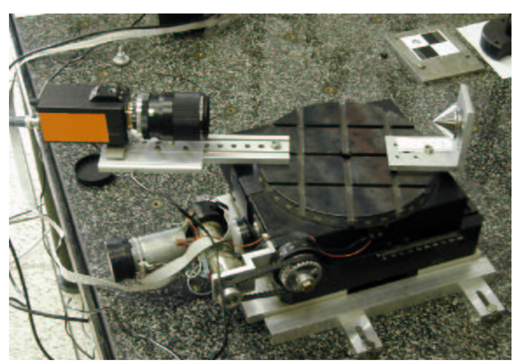

Figure 7: Two sets of rays forming hyperboloidal closures can be realized by rotating catadioptric camera consisting from a telecentric optics and a conical mirror

stereo geometry with	rays intersect at
central cameras	point
pushbroom cameras	line
circular panorama	circle
oblique cameras	

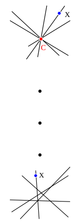

Figure 8: All stereo-geometries are "between" the central and oblique stereo geometry.

double ruled quadrics – in Euclidean space either a hyperboloid of one sheet or a hyperbolic paraboloid. We point out again that such visibility closures can be parameterized by one parameter along a line, the same way as as the rays in epipolar plane can be parameterized along epipolar lines. We have introduced the notion of oblique stereo geometry for the arrangement of non-intersecting rays allowing for one-parametric closures.

We have shown that there is an arrangement of rays in P_3 that allows for two oblique cameras with nontrivial one-parametric intersection closures and proposed a technical realization of an oblique camera pair having a subset of such visibility closures.

Besides the fact that oblique stereo geometry can be practically realized, it is an interesting theoretical concept because all other stereo geometries, namely those produced by various mosaicing techniques, lie between the central and oblique geometry, see Figure 8.

References

[1] R. Gupta and R. I. Hartley. Linear pushbroom cameras. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19(9):963–975, September 1997.

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

[3] M. Hazewinkel, editor. *Encyclopaedia of Mathematics*, volume 1. Kluwer Academic Publishers, 1995.

[4] D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination*. AMS Chelsea Publishing, Providence, Rhode Island, 2nd edition, 1999.

[5] R.J. Mihalek. *Projective Geometry and Algebraic Structures*. Accademic Press, Inc., New York, 1972.

[6] S. K. Nayar and A. Karmarkar. 360 x 360 mosaics. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 388–395, 2000.

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

[8] S. Peleg, M. Ben-Ezra, and Y. Pritch. Omnistereo: Panoramic stereo imaging. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(3):279–290, March 2001.

[9] P. Rademacher and G. Bishop. Multiple-center-of-projection images. In *SIGGRAPH'98 Orlando, FL*. ACM SIGGRAPH, July 1998.

[10] S. Seitz. The space of all stereo images. In *International Conference on Computer Vision, Vancouver, Canada*, pages 26–33, July 2001.

[11] H. Shum, A. Kalai, and S. Seitz. Omnivergent stereo. In *International Conference on Computer Vision, Corfu, Greece*, pages 22–29, September 1999.

[12] M. Watanabe and S.K. Nayar. Telecentric optics for computational vision. In *European Conference on Computer Vision*, pages II:439–451, 1996.