Combinatorial Constraints on Multiple Projections of a Set of Points

Tomáš Werner

Center for Machine Perception, Czech Technical University, Prague, Czech Republic

werner@cmp.felk.cvut.cz

Abstract

Multiple projections of a scene cannot be arbitrary, the allowed configurations being given by matching constraints. This paper presents new matching constraints on multiple projections of a rigid point set by uncalibrated cameras, obtained by formulation in the oriented projective rather than projective geometry. They follow from consistency of orientations of camera rays and from the fact that the scene is the affine rather that projective space. For their non-parametric nature, we call them combinatorial. The constraints are derived in a unified theoretical framework using the theory of oriented matroids. For example, we present constraints on 4 point correspondences for 2D camera resectioning, on 3 correspondences in two 1D cameras, and on 4 correspondences in two 2D cameras.

1. Introduction

It[1] has been recognized [12, 8, 17] that the *oriented projective geometry* (also called two-sided projective) is better to model multiple projections of a rigid scene than the projective geometry, because it can represent ray orientations and the fact that the scene is affine rather than projective space (i.e., a scene plane at infinity exists). Closely related to this are also papers on ch(e)irality and quasi-affine reconstruction [6] and on orientation-based combinatorial geometry [3], which used oriented projective geometry implicitly.

In the theory of multiple projection, considering orientations has brought an important result called quasi-affine reconstruction [6]. Using ray orientations (called *ch(e)irality* there), image point correspondences were shown not only to determine scene points and cameras (up to projectivity) but also the set of allowed scene planes at infinity.

This paper presents another important result for multiple projections, namely *new matching constraints*.

Constraints satisfied by uncalibrated projections of a set of rigid points have been almost all discovered. They are equalities of linear or multilinear functions in homogeneous coordinates of the points, the coefficients of these functions being tensors such as homography matrix, camera projection matrix, fundamental matrix, or trifocal tensor [7]. There is a certain number N of correspondences for each

---

constraint which determines its parameters uniquely (or up to a finite number of solutions). More correspondences prescribed arbitrarily yield an empty solution for parameters, less of them do not determine parameters uniquely. For example, N is 4 for a 2D homography, $5\frac{1}{2}$ for camera resection, 7 for a fundamental matrix, 6 for trifocal tensor, 7 for three 1D cameras [10], and no number of correspondences determines parameters of the system of two 1D cameras.

We show that orientation yields constraints on N or less correspondences. Because of non-parametric nature of these new constraints, we have chosen to call them *combinatorial constraints*. There are quite many of them for various number of images, points, dimension of the world, and particular types of constraints imposed. The presented new constraints are on, e.g.: 3 point pairs in two 2D images related by a homography, 5 point pairs for resectioning a conventional camera, 4 point pairs for resectioning a panoramic camera, 3 point pairs in two 1D panoramic cameras, 4 point pairs in two panoramic 2D cameras, 5 point triplets in three conventional cameras (three-view component only).

The new constraints are derived in *unified theoretical framework* using oriented matroids. *Oriented matroid* [1, 3] is a combinatorial object that captures much of the qualitative structure of a point set. Reconstructibility of a set of image correspondences can be expressed in oriented matroid language as existence of an acyclic extension of the partial chirotope defined by the correspondences.

As oriented matroids are known since as late as 1970's, this unified approach to handling orientation in multiple projections was most likely unavailable to classical photogrammetry and projective geometry, and seems to be new.

2. Notation and concepts

$\{a, b\}$ denotes a set and (a, b) an ordered tuple. $\{x \,|\, \phi(x)\}$ denotes the set defined by a property ϕ of its elements.

$\mathbb{R}$ denotes the field of reals numbers. Elements of a *vector space* (e.g., $\mathbb{R}^n$) are typeset in boldface. The symbol $\sim$ denotes equivalence of two vectors up to scale, $\mathbf{x} \sim \lambda\mathbf{x}$ for all $\lambda \neq 0$. The symbol $\overset{+}{\sim}$ denotes equivalence of two vectors up to a positive scale, $\mathbf{x} \overset{+}{\sim} \lambda\mathbf{x}$ for all $\lambda > 0$.

$|\mathbf{A}|$ denotes determinant. Determinant of the matrix with columns $\mathbf{x}_1, \ldots, \mathbf{x}_n$ is abbreviated by $[\mathbf{x}_1, \ldots, \mathbf{x}_n]$. The symbol $\mathbf{A}^\wedge$ denotes the column n-vector being the wedge product of the rows of $(n-1) \times n$ matrix $\mathbf{A}$; the n-th entry

of $\mathbf{A}^\wedge$ is the minor of the first $n-1$ columns of $\mathbf{A}$ and the remaining entries are obtained by cyclically permuting the indices. In the particular case $n=4$, it is $[\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4]^\wedge = [|\mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4|, |\mathbf{a}_3, \mathbf{a}_4, \mathbf{a}_1|, |\mathbf{a}_4, \mathbf{a}_1, \mathbf{a}_2|, |\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3|]^\top$.

The *projective d-space* is the set of equivalence classes $\mathbb{P}^d = (\mathbb{R}^{d+1}\backslash\{\mathbf{0}\})/\sim$, where $/$ denotes factorization by equivalence. Elements of $\mathbb{P}^d$ are called points and denoted in italics, e.g., x. A homogeneous vector $\mathbf{x}$ which belongs to an equivalence class x is said to *represent* that point.

Analogically, the *oriented* (also called two-sided) *projective d-space* is $\mathbb{T}^d = (\mathbb{R}^{d+1}\backslash\{\mathbf{0}\})/\stackrel{+}{\sim}$. It is isomorphic to the d-sphere $\mathbb{S}^d = \{\mathbf{x} \in \mathbb{R}^{d+1}| \|\mathbf{x}\|_2 = 1\}$. The *oriented projective geometry* is a structure on $\mathbb{T}^d$, discussed in [12].

The *d-hemisphere* is an open half of $\mathbb{T}^d$ and is denoted by $(\mathbb{T}^d, \pi)$. It consists of points lying on the positive side of a hyperplane π, i.e., of points $\mathbf{x}$ satisfying $\boldsymbol{\pi}\mathbf{x} > 0$ (π is represented by *row* 4-vector $\boldsymbol{\pi}$). $(\mathbb{T}^d, \pi)$ is isomorphic to the affine d-space $\mathbb{A}^d$; then π is interpreted as the (oriented) hyperplane at infinity. We make identification $\mathbb{A}^d = (\mathbb{T}^d, \pi)$.

3. Oriented matroids

This section surveys the facts from the theory of *oriented matroids* we will need. It is inevitably very brief; for details, we recommend the book [1] and the introductions [11, 2].

Let $E = \{1, \ldots, N\}$. *Oriented matroid of rank r* is a pair (E, χ), where its *chirotope* $\chi: E^r \to \{-1, 0, +1\}$ is an alternating[2] map satisfying the *chirotope axioms*[3]:

(Chi1) Set $B = \{\{n_1, \ldots, n_r\} \in \binom{E}{r}| \chi(n_1, \ldots, n_r) \neq 0\}$ is the set of bases of a *matroid*[4], i.e., B satisfies the *exchange axiom*: for all $b_1, b_2 \in B$ and $e \in b_1\backslash b_2$, there exists $f \in b_2\backslash b_1$ such that $(b_1\backslash\{e\}) \cup \{f\} \in B$.

(Chi2) For $\mathbf{n} = (n_1, \ldots, n_{r-2}) \in E^{r-2}$ and $a, b, c, d \in E$, the set $\{\chi(\mathbf{n}, a, b)\chi(\mathbf{n}, c, d), -\chi(\mathbf{n}, a, c)\chi(\mathbf{n}, b, d), \chi(\mathbf{n}, a, d)\chi(\mathbf{n}, b, c)\}$ equals either $\{-1, +1\}$ or $\{0\}$.

Vector configuration is a matrix $\mathbf{V} = (\mathbf{x}_1, \ldots, \mathbf{x}_N)$ of full rank, where $\mathbf{x}_n \in \mathbb{R}^r$. *Chirotope of vector configuration* $\mathbf{V}$ is the map $\chi_{\mathbf{V}}: E^r \to \{-1, 0, +1\}$ defined by

$$\chi_{\mathbf{V}}(n_1, \ldots, n_r) = \text{sign}[\mathbf{x}_{n_1}, \ldots, \mathbf{x}_{n_r}]. \qquad (1)$$

Chirotope $\chi_{\mathbf{V}}$ of any vector configuration $\mathbf{V}$ satisfies the chirotope axioms.

A *realization* of an oriented matroid (E, χ) is a vector configuration $\mathbf{V}$ such that $\chi = \chi_{\mathbf{V}}$. An oriented matroid is *realizable* if it possesses a realization. The chirotope axioms are necessary but not sufficient for realizability. Unlike in projective geometry, realizability cannot be ensured

by adding another axiom; in fact, testing for realizability is NP-complete. However, all matroids for $r = 2$, for $r = 3$ and $N \leq 8$, and for $r \geq 4$ and $N \leq r + 3$, are realizable; otherwise, the ratio of non-realizable and realizable oriented matroids is 'small' for 'small' r and N [1, 5].

An oriented matroid is *uniform* if $\chi(n_1, \ldots, n_r) \neq 0$ for all $\{n_1, \ldots, n_r\} \in \binom{E}{r}$. An oriented matroid is *acyclic*[5] if there exists $n = \{n_1, \ldots, n_{r-1}\} \in \binom{E}{r-1}$ such that $\chi(n_1, \ldots, n_{r-1}, a)$ has the same sign for all $a \in E\backslash n$.

Where do the axioms come from? (Chi1) is a natural requirement for B to form a 'set of bases', known from linear algebra. *(Chi1) is always true for uniform oriented matroids.* As this paper restricts to uniform matroids, (Chi2) is more important for it. To motivate (Chi2), consider identities $[\mathbf{n}ab][\mathbf{n}cd] - [\mathbf{n}ac][\mathbf{n}bd] + [\mathbf{n}ad][\mathbf{n}bc] = 0$, relating six-tuples of the $r \times r$ submatrices of $\mathbf{V}$, where $[\mathbf{n}ab]$ stands for $[\mathbf{x}_{n_1}, \ldots, \mathbf{x}_{n_{r-2}}, \mathbf{x}_a, \mathbf{x}_b]$, etc. For the LHS of the identity to vanish, the signs of its three summands cannot be arbitrary, which is what (Chi2) requires.

In this paper, the vector configuration $\mathbf{V}$ is interpreted as the ordered N-tuple of points from $\mathbb{T}^{r-1}$ represented by $\mathbf{x}_n$. Then, uniformity of an oriented matroid means that no r points of any realization (if it exists) lie in a single hyperplane. Acyclicity means that there exists a hyperplane $\boldsymbol{\pi}$ such that all points lie on its positive side, $\forall n(\boldsymbol{\pi}\mathbf{x}_n > 0)$.

3.1. Extending a partial chirotope

An alternating map defined on E^r is fully determined by its values $\chi(\mathbf{n})$ where $\mathbf{n} \in \Lambda(N, r) = \{(n_1, \ldots, n_r) \in E^r| 1 \leq n_1 < \cdots < n_r \leq N\}$, the rest is given by antisymmetry. Note that $|\Lambda(N, r)| = \binom{N}{r}$. *Partial chirotope* is an alternating map $\hat{\chi}: F \to \{-1, 0, +1\}$ where $F \not\supset \Lambda(N, r)$, i.e., some of its values are unknown. An *extension of partial chirotope* $\hat{\chi}$ is a chirotope χ, defined on the whole of E^r and satisfying the chirotope axioms, such that $\chi(\mathbf{n}) = \hat{\chi}(\mathbf{n})$ for $\mathbf{n} \in F$. A partial chirotope is *extendable* if it has an extension. Testing partial chirotope for extendability is NP-complete [15].

We use the following algorithm[6] to compute all uniform extensions of a partial uniform chirotope. It starts with partial chirotope $\hat{\chi}_0$ defined on $F_0 \not\supset \Lambda(N, r)$, and tries to define $\hat{\chi}_0$ on a new argument $\mathbf{n}_1 \notin F_0$. If it succeeds, the result is a chirotope $\hat{\chi}_1$, defined on the larger set $F_1 = F_0 \cup \mathbf{n}_1$. This step is repeated until reaching a full chirotope χ, defined on the whole of E^r (if $\hat{\chi}_0$ is extendable).

Two values, -1 and $+1$ (0 is omitted because we assume uniform $\hat{\chi}_i$), can be assigned to $\hat{\chi}_i(\mathbf{n}_i)$ in each step. This results in a tree search, yielding all extensions of $\hat{\chi}_0$. The tree search causes the (inevitable) exponential time complexity; the algorithm is practical only for 'small' N and r.

---

[2]*Alternating* means $\chi(\pi(\mathbf{n})) = \text{sign}(\pi)\chi(\mathbf{n})$ for any permutation π.

[3]The chirotope axioms can be formulated in several different forms. We used the form from [11] which suits best our purpose. See [1, 2] for more about the justification of these axioms.

[4]*Matroid* is a different object than *oriented matroid*. While the oriented matroid of $\mathbf{V}$ is the pair $(E, \chi_{\mathbf{V}})$, the matroid of $\mathbf{V}$ is fully given by $(E, |\chi_{\mathbf{V}}|)$, i.e., it captures only incidence.

[5]In the literature, acyclicity is defined using the oriented graph interpretation of oriented matroids, as absence of a positive circuit (cycle) [1].

[6]The algorithm has been designed *ad hoc*, for playing with the combinatorial matching constraints. Perhaps, a more efficient one can be found in the literature (e.g., see [15]).

The basic block of the algorithm is the *single value extension test*, testing whether $\hat{\chi}_{i-1}$, known to satisfy (Chi2), can attain a value $\hat{\chi}_i(\mathbf{n}_i)$ in the new argument $\mathbf{n}_i$, such that $\hat{\chi}_i$ satisfies (Chi2). (Chi1) need not be tested thanks to uniformity. The advantage of doing a single extension at a time is that only 6-tuples of arguments (which are subject to the condition in (Chi2)) containing $\mathbf{n}_i$ can be considered.

The algorithm can be used to extend a partial chirotope not only to a full chirotope, but also to a larger partial one.

4. Representing scene, image, and camera

In computer vision and photogrammetry, scene is traditionally modeled by $\mathbb{P}^3$ and images by $\mathbb{P}^2$. Using oriented projective geometry allows for a better model, since ray orientations and the fact that the scene is affine are expressible.

The *scene* is represented by the affine 3-space $\mathbb{A}^3 = (\mathbb{T}^3, \pi_\infty)$, i.e., it consists of points $\mathbf{X}$ being in front of the *plane at infinity*, $\pi_\infty \mathbf{X} > 0$.

We distinguish two types of images. *Directional image* contains rays from a single hemisphere in scene; it is obtained with conventional (e.g., TV or photographic) cameras. *Spherical image* contains rays that do not fit in a single hemisphere; it is obtained by omnidirectional cameras[7]. Spherical image is represented by $\mathbb{T}^2$. Directional image is represented by $\mathbb{A}^2 = (\mathbb{T}^2, l_\infty)$, where l_∞ is the *image line at infinity*; it consists of points $\mathbf{x}$ satisfying $l_\infty \mathbf{x} > 0$.

Camera is a linear mapping from the scene to the image, $\mathbf{x} = \mathbf{PX}$ where $\mathbf{P} \in \mathbb{R}^{3 \times 4}$. The *directional camera* is besides characterized by its line at infinity, hence denoted by $(\mathbf{P}, l_\infty)$. The condition $l_\infty \mathbf{x} > 0$ is equivalent to saying that the camera sees only scene points $\mathbf{X}$ that lie in front of the scene plane $l_\infty \mathbf{P}$, passing through the projection center.

Vector $\mathbf{P}^\wedge$ represents the *oriented projection center* [14, 6, 7, 17, 4] of camera $\mathbf{P}$; note that $\mathbf{PP}^\wedge = \mathbf{0}$ and $|(\mathbf{P}^\top | \mathbf{P}^\wedge)| > 0$. Camera $\mathbf{P}$ in the affine scene $(\mathbb{T}^3, \pi_\infty)$ is *right-handed* if $\pi_\infty \mathbf{P}^\wedge > 0$ and *left-handed* if $\pi_\infty \mathbf{P}^\wedge < 0$. Since a left-handed camera can be always transformed to a right-handed one by changing the orientation of its image basis, we assume that all cameras are right-handed[8].

5. Realizability of point correspondences

Consider multiple points and cameras. We denote the n-th scene point by $\mathbf{X}_n$, the k-th camera by $\mathbf{P}^k$, and the projection of $\mathbf{X}_n$ in camera $\mathbf{P}^k$ by $\mathbf{x}_n^k$. It follows from the above adopted representations that the set of image point correspondences $\{\mathbf{x}_n^k\}$ (where $n = 1, \ldots, N$ and $k = 1, \ldots, K$), and the set of scene points $\{\mathbf{X}_n\}$ and right-handed cameras $\{\mathbf{P}^k\}$ in the affine scene $(\mathbb{T}^3, \pi_\infty)$ satisfy

$$\forall k, n(\mathbf{x}_n^k \overset{\centerdot}{\sim} \mathbf{P}^k \mathbf{X}_n, \pi_\infty \mathbf{X}_n > 0, \pi_\infty \mathbf{P}^{k\wedge} > 0) . \quad (2)$$

We call the set $\{\mathbf{x}_n^k\}$ *realizable* if there exist $\{\mathbf{X}_n\}$, $\{\mathbf{P}^k\}$ and π_∞ satisfying (2).

Using identity $[\mathbf{PX}_1, \mathbf{PX}_2, \mathbf{PX}_3] = [\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{P}^\wedge]$, the first condition $\forall n(\mathbf{x}_n \overset{\centerdot}{\sim} \mathbf{PX}_n)$ can be shown equivalent to two conditions holding simultaneously: $\forall n(\mathbf{x}_n \sim \mathbf{PX}_n)$ and $\forall a, b, c(\text{sign}[\mathbf{x}_a, \mathbf{x}_b, \mathbf{x}_c] = \text{sign}[\mathbf{X}_a, \mathbf{X}_b, \mathbf{X}_c, \mathbf{P}^\wedge])$. Thus, the set $\{\mathbf{x}_n^k\}$ is realizable if and only if there exist $\{\mathbf{X}_n\}$ and $\{\mathbf{P}^k\}$ satisfying the following three conditions:

$$\forall k, n(\mathbf{x}_n^k \sim \mathbf{P}^k \mathbf{X}_n) \quad (3)$$
$$\forall k, a, b, c(\text{sign}[\mathbf{x}_a^k, \mathbf{x}_b^k, \mathbf{x}_c^k] = \text{sign}[\mathbf{X}_a, \mathbf{X}_b, \mathbf{X}_c, \mathbf{P}^{k\wedge}]) \quad (4)$$
$$\exists \pi_\infty[\forall n(\pi_\infty \mathbf{X}_n > 0), \forall k(\pi_\infty \mathbf{P}^{k\wedge} > 0)] \quad (5)$$

We will refer to (3), (4), and (5) as *projective, orientation* and *affine constraint* on correspondences[9], respectively.

6. Combinatorial matching constraints

Unlike the projective constraint, the orientation and affine constraints are combinatorial in nature. We will show that they can be expressed in oriented matroid terms.

Let $\mathbf{V} = (\mathbf{X}_1, \ldots, \mathbf{X}_N, \mathbf{P}^{1\wedge}, \ldots, \mathbf{P}^{K\wedge})$ be the vector configuration consisting of scene points and camera centers. The *scene chirotope* $\chi_\mathbf{V}$ consists of signs of all 4×4 minors of $\mathbf{V}$. If $K \geq 4$, there are five types of these determinants, [xxxx], [xxxc], [xxcc], [xccc], and [cccc], where 'x' denotes a scene point and 'c' a camera center. The orientation constraint (4) says that the signs of determinants of type [xxxc] are all known; in other words, image correspondences define a *partial scene chirotope*, $\hat{\chi}_\mathbf{V}$. The affine constraint (5) further requires $\chi_\mathbf{V}$ to be acyclic. Put together,

> *the chirotope of the vector configuration formed by scene points and camera centers is an acyclic extension of the partial chirotope defined by the image point correspondences.*

In this paper, *we assume that chirotopes of $\mathbf{V}$ and chirotopes of image point lists are uniform.* Thus, all extensions of $\hat{\chi}_\mathbf{V}$ can be computed by the algorithm in Section 3.1. Discarding cyclic ones[10] yields all possible chirotopes $\chi_\mathbf{V}$.

Existence of an acyclic extension is necessary for the set of correspondences to be realizable. Further in this section, we discuss *minimal non-realizable configurations* of correspondences for setups of various number of cameras and points and dimensions of the scene.

For each setup, the orientation constraint alone is discussed first, and then the affine one is added (i.e., the simultaneous orientation+affine constraint is discussed).

---

[7]Omnidirectional cameras can be built, e.g., by combining an ordinary camera and a curved mirror [9]. Another example is the fish-eye lens.

[8]E.g., the usual choice used in computer vision [7], $\pi_\infty = (0, 0, 0, 1)$ and $\mathbf{P} = \mathbf{KR}(\mathbf{I} | -\mathbf{t})$ where $|\mathbf{KR}| > 0$, yields a right-handed camera.

[9]The affine constraint is called *ch(e)iral inequalities* in [6, 7]; these inequalities assume that the overall scale signs of $\mathbf{X}_n$ and $\mathbf{P}^k$ already satisfy $\mathbf{x}_n^k \overset{\centerdot}{\sim} \mathbf{P}^k \mathbf{X}_n$. Scene and cameras satisfying $\mathbf{x}_n^k \overset{\centerdot}{\sim} \mathbf{P}^k \mathbf{X}_n$ are called *oriented projective reconstruction* in [17]. The terms *projective* and *orientation constraint* are used in [16] in the same meaning as here.

[10]Cyclicity can be tested in time $O[\binom{N}{r}]$ using the definition in Sec. 3.

6.1. Camera resectioning

Camera resectioning [7] is the simplest[11] example. We want to compute $\mathbf{P}$ from set $\{(\mathbf{X}_n, \mathbf{x}_n)\}$ of scene-image correspondences in a general position, $n = 1, \ldots, N$.

We will consider first the 1D camera ($\mathbf{X}_n \in \mathbb{R}^3$ and $\mathbf{x}_n \in \mathbb{R}^2$). Let $\mathbf{V} = (\mathbf{X}_1, \ldots, \mathbf{X}_N, \mathbf{P}^\wedge)$. The scene chirotope $\chi_{\mathbf{V}}$ contains signs of minors of two types, [xxx] and [xxc]. All are known from the correspondences.

Orientation constraint requires that $\chi_{\mathbf{V}}$ satisfies (Chi2).

For $N = 3$, there are $2^{|\Lambda(|\mathbf{V}|,3)|} = 16$ different alternating maps $\chi_{\mathbf{V}} : E^{|\mathbf{V}|} \to \{-1, +1\}$. All satisfy (Chi2), that is, all are chirotopes. Thus (4) imposes no constraint for $N = 3$.

For $N = 4$, only 384 of $2^{10} = 1024$ alternating maps are chirotopes. Therefore *there are configurations of four correspondences violating the orientation constraint.* They exist for both spherical and directional images. For the latter, $\frac{2}{3}$ of all configuration classes are allowed and $\frac{1}{3}$ forbidden.

Affine constraint requires that $\{\mathbf{X}_n\} \cup \{\mathbf{P}^\wedge\}$ fits in a 2-hemisphere[12], i.e., $\chi_{\mathbf{V}}$ is acyclic. Note, the set $\{\mathbf{X}_n\}$ alone already fits in a 2-hemisphere because the scene is affine.

For $N = 3$, 2 of the 2^4 chirotopes are acyclic, corresponding to a single (up to mirroring) configuration of 4 points on the 2-sphere that does not fit in a 2-hemisphere. It results in *the single (up to non-mirroring homographies) forbidden configuration of three correspondences,* $\{(\mathbf{X}_n, \mathbf{x}_n)\} = \{((1,0,0),(0,1)), ((0,1,0),(1,0)), ((0,0,1),(-1,-1))\}$. Note that the set $\{\mathbf{x}_n\}$ in this configuration is spherical, there is no constraint for $N = 3$ and a directional image.

2D camera. For a 2D camera ($\mathbf{X}_n \in \mathbb{R}^4$, $\mathbf{x}_n \in \mathbb{R}^3$), the situation is entirely analogous. Configurations violating the orientation constraint[13] exist for $N \geq 5$. For $N = 4$, there is a single (up to non-mirroring homographies) configuration violating the affine constraint.

6.2. Two 1D cameras

Let $\{(\mathbf{x}_n, \mathbf{x}'_n)\}$ be N correspondences from two 1D cameras in 2D scene, $\mathbf{x}_n, \mathbf{x}'_n \in \mathbb{R}^2$, $\mathbf{X}_n \in \mathbb{R}^3$. Unlike for 3D scene, the projective constraint alone is not violated by any configuration for any N, since any two lines in $\mathbb{P}^2$ intersect.

---

[11]Even simpler (but trivial) example is a *homography-related image pair*. Let $\{(\mathbf{x}_n, \mathbf{x}'_n)\}$ be the set of correspondences in two cameras with equal orientations of image coordinate bases. It is $\forall n(\mathbf{x}'_n \stackrel{+}{\simeq} \mathbf{H}\mathbf{x}_n)$. If the homography is induced by a scene plane such that both cameras lie on its single side, it is $|\mathbf{H}| > 0$ and thus $\chi = \chi'$; if the cameras lie on its opposite sides (the plane must be transparent then), it is $|\mathbf{H}| < 0$ and $\chi = -\chi'$. If the camera rotates around its center or zooms, $|\mathbf{H}| > 0$.

[12]This is not exactly true. If the scene points and camera have been transformed by a *mirroring homography*, there should be $\{\mathbf{X}_n\} \cup \{-\mathbf{P}^\wedge\}$ instead. It is because the oriented camera center $\mathbf{P}^\wedge$ does not transform co-variantly (i.e., as an ordinary scene point). It can be seen from the formula $(\mathbf{P}\mathbf{T}^{-1})^\wedge = |\mathbf{T}|^{-1}\mathbf{T}\mathbf{P}^\wedge$; if $|\mathbf{T}| < 0$, then $\mathbf{P}^\wedge$ swaps sign [6, 7, 17].

[13]Projectively, $5\frac{1}{2}$ correspondences uniquely determine $\mathbf{P}$ [7]. 5 correspondences constrain the camera center $\mathbf{P}^\wedge$ on a twisted cubic, passing through $\mathbf{P}^\wedge$. The orientation constraint further restricts $\mathbf{P}^\wedge$ either on a *segment* of the cubic, or (for forbidden configurations) on the empty set.

The scene chirotope $\chi_{\mathbf{V}}$, $\mathbf{V} = \{\mathbf{X}_1, \ldots, \mathbf{X}_N, \mathbf{P}^\wedge, \mathbf{P}'^\wedge\}$, contains signs of minors of type [xxx], [xxc], and [xcc]. Of them, [xxc] are known, defining partial chirotope $\hat{\chi}_{\mathbf{V}}$.

To discover forbidden configurations, we proceeded by enumerating and testing representants of all different classes of configurations. The classes are defined by equality of both image chirotopes.

Orientation constraint requires $\hat{\chi}_{\mathbf{V}}$ to be extendable. There is no forbidden configuration for $N = 3$. However, up to image homographies, relabelling, and replacing any pair $(\mathbf{x}, \mathbf{x}')$ with $(-\mathbf{x}, -\mathbf{x}')$, *there is a single forbidden configuration of four correspondences violating the orientation constraint* (see Figure 1; points are depicted in $\mathbb{T}^1$).

In this minimal configuration, one image is directional and one spherical, and this property is invariant to replacements $(\mathbf{x}, \mathbf{x}') \mapsto (-\mathbf{x}, -\mathbf{x}')$. Presence of a spherical image is crucial here: *if both images are directional, any configuration of any number of correspondences satisfy (4).*
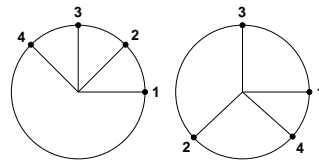

Figure 1. The forbidden configuration of 4 points in two 1D cameras violating the orientation constraint.
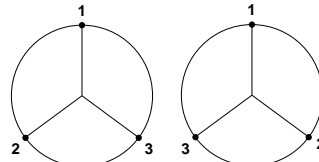

Figure 2. The forbidden configuration of 3 points in two 1D cameras violating the affine constraint.

Affine constraint requires that at least one extension of $\hat{\chi}_{\mathbf{V}}$ is acyclic. Figure 2 shows *the single forbidden configuration of three points violating the orientation+affine constraint*. It is given as follows: (i) both images are spherical and (ii) $(\mathbf{x}_n)$ and $(\mathbf{x}'_n)$ are related by a mirroring homography, i.e., $\forall n(\mathbf{x}'_n \simeq \mathbf{H}\mathbf{x}_n)$ for some $\mathbf{H}$ such that $|\mathbf{H}| < 0$.

6.3. Two 2D cameras

Consider realizability of the set $\{(\mathbf{x}_n, \mathbf{x}'_n)\}$ of N correspondences from two 2D cameras in the 3D scene, $\mathbf{x}_n, \mathbf{x}'_n \in \mathbb{R}^3$, $\mathbf{X}_n \in \mathbb{R}^4$. This case is very similar to two 1D cameras.

Orientation constraint does not forbid any configurations of $N \leq 4$ correspondences. However, *it does constitute a constraint on $N = 5$ correspondences*. An example forbidden configuration is in Figure 3 (points depicted in $\mathbb{A}^2$).

Affine constraint provides a *single forbidden configuration of four correspondences*, defined as follows: the images are spherical and related by a mirroring 2D homography.
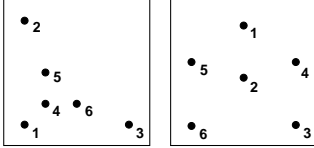
Figure 3. A forbidden configuration of 5 points in two 2D cameras violating the orientation constraint.
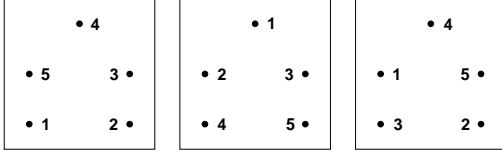

Figure 4. A configuration of 5 points in three 2D cameras that does not violate any of two-view orientation constraints but violates the three-view one.
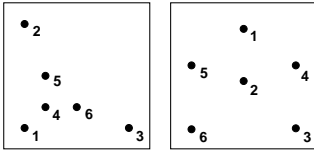

Figure 5. A configuration of 6 points in two 2D cameras that does not violate the orientation constraint but violates combined projective+orientation one.

6.4. Three and more cameras

For three views, a natural question arises, if the three-view constraint is stronger than three pair-wise constraints imposed simultaneously on view pairs $(1, 2)$, $(1, 3)$ and $(2, 3)$.

For the orientation constraint, the answer is yes for both four points in 1D cameras and five points in 2D cameras. An example of a non-decomposable forbidden configuration for 2D cameras is in Figure 4 (shown in $\mathbb{A}^2$).

However, the simultaneous orientation+affine constraint on three points in three 1D cameras is equivalent to three orientation+affine constraints applied pair-wise. This holds analogically for four points in three 2D cameras.

For 1D cameras, the orientation constraint does not forbid any configuration of 3 points across any number of views. This holds analogically for 2D cameras and 4 points.

6.5. Oriented matroids vs. constraints on epipoles

Consider the same setup as in Section 6.3. It is known [17, 16] that epipoles $\mathbf{e}$ and $\mathbf{e}'$ and any two corresponding pairs $(\mathbf{x}_a, \mathbf{x}'_a)$ and $(\mathbf{x}_b, \mathbf{x}'_b)$ in two 2D images satisfy

$$\text{sign}[\mathbf{x}_a, \mathbf{x}_b, \mathbf{e}] = -\text{sign}[\mathbf{x}'_a, \mathbf{x}'_b, \mathbf{e}'] . \qquad (6)$$

How is this related to the oriented matroid approach?

Condition (6) follows from (4) and antisymmetry of determinants. It is obtained from two equations stated in (4), written down for the first and then for the second camera.

Chirotope of the configuration $\mathbf{v} = (\mathbf{x}_1, \ldots, \mathbf{x}_N, \mathbf{e})$ contains signs of determinants of type [xxx] and [xxe], where

'x' denotes image point and 'e' epipole. A positions of $\mathbf{e}$ in the first image is described by the signs of type [xxe]. Possible positions of $\mathbf{e}$ are given by all extensions of the partial chirotope containing only type [xxx] signs.

Testing the partial scene chirotope $\hat{\chi}_{\mathbf{V}}$ for extendability is equivalent to the following test: (i) for each epipole, find the set of partial chirotopes describing its possible position in each image; (ii) check if any pair of these partial chirotopes satisfies (6) for all pairs (a, b).

The same applies also to 1D cameras. For them, (6) is replaced by condition $\text{sign}[\mathbf{x}_a, \mathbf{e}] = -\text{sign}[\mathbf{x}'_a, \mathbf{e}']$.

7. Knowing less/more of the partial chirotope

The partial scene chirotope $\hat{\chi}_{\mathbf{V}}$ specified by the images contains determinant signs of type [xxxc]. It is possible to consider a slightly different partial scene chirotope, by putting less or more knowledge in it.

Thus, *scene constraints* can be imposed by specifying some signs of type [xxxx], [xxcc], [xccc] or [cccc]. E.g., we can test if the image correspondences are consistent with the scene plane 123 separating points 4 and 5, in other words, [1234][1235]< 0. Or, with two camera centers separated by the convex hull of scene points, etc.

Also, we can leave some signs of type [xxxc] undefined, i.e., specify signs of only some triplets of image points. An obvious example are *occlusions*, when some points were not measured in some images. As another example, consider an image distorted by a non-linear continuous transformation. While signs of the triplets whose convex hulls are large are likely to be affected by this distortion, signs of the triplets with small convex hull are likely to be unaffected, and provide some knowledge about the scene.

8. Adding the projective constraint?

Existence of scene points and cameras satisfying the orientation and affine constraints is necessary but *sometimes not sufficient* for realizability of image correspondences. The sufficient condition is obtained by additionally imposing the projective constraint (3).

For sufficiently large N, (3) might induce *non-linear* dependencies among scene points and camera centers. E.g., fixing one camera center might restrict the other center(s) to a variety in the scene. This can exclude some scene chirotopes, otherwise consistent with (4) and (5).

Unfortunately, these dependencies are difficult to incorporate in the unified framework described in Section 6.

E.g., consider two 2D cameras. If five image point correspondences are given, (3) relates the two epipoles by a 5th degree Cremona mapping [13, 16]. Requiring simultaneously (6) resulted in the constraint on five points in two images presented in [16]. By exhaustive enumeration of cases, we have proven that this constraint is *equivalent to (6)*, i.e., to the orientation constraint from Section 6.3 alone[14].

---

[14]Currently, I do not see any simple evidence for this.

8.1. Example: six points in two 2D cameras

However, for six points in two 2D images there exist configuration with extendable partial chirotope $\hat{\chi}_{\mathbf{V}}$ violating the simultaneous projective+orientation constraint. An example is in Figure 5. This section presents an algorithm to test whether six correspondences in two images satisfy simultaneously (3) and (4). It is closely related to the paper [16].

We start with the constraint imposed on the epipoles e and e' by (3) [13]. Let $\{x_n, x'_n\}$ be correspondences in two images, $n = 1, \ldots, 6$ and $x_n, x'_n \in \mathbb{P}^2$. Let h_{ij} denote the homography relating all except the i- and j-th correspondence, $h_{ij}(x_n) = x'_n$ for all $n \in \{1, \ldots, 6\} \backslash \{i, j\}$. Let Q_i denote the conic passing through points $\{x_1, \ldots, x_6\} \backslash \{x_i\}$. Let $Q'_{ij} = h_{ij}(Q_i)$ denote Q_i mapped to the second image by homography h_{ij}. Let 12 new points in the second image be defined by $y'_n = \cap_{i \neq n} Q'_{ni}$ and $z'_n = \cap_{i \neq n} Q'_{in}$. It turns out that the 18 points x'_n, y'_n, z'_n lie on a common cubic, C'.

Symmetrically by interchanging the images, we obtain points y_n and z_n and cubic C in the first image.

The first epipole e can lie only on C and the second one on C'. The second epipole is uniquely given by the first one, $e' = f(e) \in C'$. Moreover[15], $f(x_n) = z'_n$ for each n.

The test is done as follows. If, for any n, one of the pair of epipoles $(e, e') = (x_n, z'_n)$ and $(e', e) = (z_n, x'_n)$ can be equipped with orientation (i.e., the overall scales of their representing vectors $\mathbf{e}$ and $\mathbf{e}'$ chosen) such that (6) holds for all pairs $a, b \in \{1, \ldots, 6\} \backslash \{n\}$, then there are scene points and cameras such that (3) and (4) simultaneously hold.

9. Concluding remarks

This paper is an application of oriented matroids to geometry of multiple uncalibrated projections. Using the fact that signs of determinants made of all 4-tuples of scene points and camera centers must obey the chirotope axioms, we have shown that some configurations of correspondences, allowed in the (unoriented) projective geometry, are impossible in reality. The oriented matroid approach allows to derive combinatorial constraints for any scene dimension.

Practically, the combinatorial constraints are quite weak compared to the projective ones. Yet they can be useful in applications, e.g., for wide-baseline stereo. Moreover, we believe they are important and interesting theoretically.

There is a number of open questions. The unoriented (projective) matching constraints stop with $K = 4$ images and the quadrifocal tensor. Do the combinatorial constraints, e.g. for 1D cameras, stop for any K? How is it with applying the projective constraint? E.g., can the constraint on 5 points in 3 images be made stronger by imposing it?

We considered only forbidden configurations for minimal number N of points. If N is such that the scene points

and cameras are uniquely given by (3) (e.g., $N = 7$ for two 2D cameras), the test is easy: compute scene points and cameras and test for (4) and (5). How is it with N in between, e.g., $4 < N < 7$ for three 1D cameras?

In many cases, extendability of the partial scene chirotope $\hat{\chi}_{\mathbf{V}}$ to a full scene chirotope $\chi_{\mathbf{V}}$ is equivalent to extendability of $\hat{\chi}_{\mathbf{V}}$ to a scene chirotope containing only signs of type [xxxc] and [xxxx], or even [xxxc] and [xxcc], etc. How is it in general?

References

[1] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler. *Oriented Matroids, Encyclopaedia of Mathematics 46*. Cambridge University Press, 1993.

[2] J. Bokowski. Oriented matroids. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, chapter 2.5. Elsevier Science Publishers, 1993.

[3] S. Carlsson. Combinatorial geometry for shape representation and indexing. In *Object Representation in Computer Vision*, pages 53–78, 1996.

[4] O. Chum, T. Werner, and T. Pajdla. Joint orientation of epipoles. In *Brittish Machine Vision Conf. BMVC'03*, 2003.

[5] L. Finschi. *A Graph Theoretical Approach for Reconstruction and Generation of Oriented Matroids*. PhD thesis, Swiss Federal Institute of Technology, Zurich, 2001.

[6] R. I. Hartley. Chirality. *Int. Jour. Computer Vision IJCV*, 26(1):41–61, 1998.

[7] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.

[8] S. Laveau and O. Faugeras. Oriented projective geometry for computer vision. In *ECCV96*, pages I:147–156, 1996.

[9] T. Pajdla, T. Svoboda, and V. Hlaváč. Epipolar geometry of central panoramic cameras. In R. Benosman and S. B. Kang, editors, *Panoramic Vision : Sensors, Theory, and Applications*. Springer Verlag, Berlin, Germany, 2000.

[10] L. Quan. Two-way ambiguity in 2D projective reconstruction from three uncalibrated 1D images. *PAMI*, 23(2):212–216, February 2001.

[11] J. Richter-Gebert and G. J. Ziegler. Oriented matroids. A preprint from Algorithmic and Discrete Mathematics group, Technical University Berlin. 23 pages. Also on www.

[12] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, Inc., 1250 Sixth Avenue, San Diego, CA 92101, 1991.

[13] R. Sturm. *Die Lehre von den geometrischen Verwandschaften*, volume 1 of *B. G. Teubners Sammlung von Lehrbüchern auf dem Gebiete der mathematischen Wissenschaften*. Druck und Verlag von B. G. Teubner, Leipzig und Berlin, 1908.

[14] B. Triggs. Matching constraints and the joint image. In *ICCV95*, pages 338–343, 1995. See also the full version on www, submitted to IJCV.

[15] F. Tschirschnitz. Testing extendability for partial chirotopes is NP-complete. On www of Institute for Theoretical Computer Science, ETH Zürich, tschirsc@inf.ethz.ch, 1999.

[16] T. Werner. Constraint on five points in two images. In *Comp. Vision and Pattern Recognition Conf. CVPR'03*, 2003.

[17] T. Werner and T. Pajdla. Oriented matching constraints. In T. Cootes and C. Taylor, editors, *British Machine Vision Conference 2001*, pages 441–450, London, UK, September 2001. British Machine Vision Association.

---

[15]In [13], only y'_n are defined and shown to lie on a common cubic C' with x'_n. Points z'_n are not mentioned there. However, the correspondence $f(x_n) = z'_n$ is crucial for the orientation constraint test.