Cheirality Invariants

Richard I. Hartley

GE - Corporate Research and Development,
P.O. Box 8, Schenectady, NY, 12301.
Ph : (518)-387-7333
Fax : (518)-387-6845
email : hartley@crd.ge.com

Abstract

It is known that a set of points in 3 dimensions is determined up to projectivity from two views with uncalibrated cameras. It is shown in this paper that this result may be improved by distinguishing between points in front of and behind the camera. Any point that lies in an image must lie in front of the camera producing that image. Using this idea, it is shown that the scene is determined from two views up to a more restricted class of mappings known as good projectivities, *which are precisely those projectivities that preserve the convex hull of an object of interest. An invariant of good projectivity known as the* cheirality invariant *of a set of points is defined and it is shown how the cheirality invariant may be computed using two uncalibrated views. As demonstrated theoretically and by experiment the cheirality invariant may distinguish between sets of points that are projectively equivalent (but not via a good projectivity). These results lead to necessary and sufficient conditions for a set of corresponding pixels in two images to be realizable as the images of a set of points in 3 dimensions.*

Using similar methods, a necessary and sufficient condition is given for the the orientation *of a set of points to be determined by two views. If the perspective centres are not separated from the point set by a plane, then the orientation of the set of points is determined from two views.*

Good projectivities and the cheirality invariant are also defined for point sets in a plane, which allows these new methods to be applied to images of planar objects.

1 Introduction

Consider a set of points $\{\mathbf{x}_i\}$ lying in a plane in space and let $\{\mathbf{u}_i\}$ and $\{\mathbf{u}'_i\}$ be two images of these points taken with arbitrary uncalibrated perspective (pinhole) cameras. It is well known that the points $\mathbf{u}_i$ and $\mathbf{u}'_i$ are related by a planar projectivity, that is, there exists h a projectivity of the plane such that $h\mathbf{u}_i = \mathbf{u}'_i$ for all i. This fact has been used for the recognition of planar objects. For instance in [Rothwell-92] planar projective invariants were used to define indexing functions allowing look-up of the objects in an object data-base. Since the indexing functions are invariant under plane projectivities, they provide the same value independent of the view of the object.

In a similar way, it has been shown in [Faugeras-92] and [Hartley-Gupta-92] that a set of points in 3-dimensions is determined up to a 3-dimensional projectivity by two images taken with uncalibrated cameras. Both these papers give a constructive method for determining the point configuration (up to projectivity). This permits the computation of projective invariants of sets of points seen in two views. An experimental verification of these results has been reported in [Hartley-92] and is summarized in this paper.

The papers just cited make no distinction between points that lie in front of the camera and those that lie behind. The specification of the front of a camera will be termed the *cheirality* of the camera (from Greek : $\chi\epsilon\iota\rho =$ *hand* or *side*). It is well know that camera cheirality is valuable in determining scene geometry for calibrated cameras. Longuet-Higgins [Higgins-81] uses it to distinguish between four different scene reconstructions. No systematic treatment of cheirality of uncalibrated cameras has previously appeared, however. Investigation of this phenomenon turns out to be quite fruitful, as is seen in the present paper. Cheirality is valuable in distinguishing different point sets in space, especially in allowing projectively equivalent point sets to be distinguished.

The major results of this paper are summarized now. In Definition 4.4 a class of projectivities called *good projectivities* is defined, consisting of those ones that preserve the convex hull of a set of points of interest. In section 5 an invariant of good projectivity is defined – the cheirality invariant. Theorem 6.13 strengthens the result of [Faugeras-92, Hartley-Gupta-92] by showing that a 3-dimensional point set is determined up to good projectivity by its image in two uncalibrated views. This sharpening of the theorem of [Faugeras-92, Hartley-Gupta-92] results from a consideration of the cheirality of the cameras. In section 9 an example of computation of the cheirality invariant for 3D point sets shows that it is useful in distinguishing between non-equivalent point sets. In section 7 the concept of good-projectivity is applied to orientation of point sets, explaining why some point sets allow two differently oriented reconstructions from two views, whereas some do not. The relationship of this result to human visual perception of 3D scenes is briefly mentioned, suggesting that the brain accepts various interpretations of a scene differing by good projectivities, but not by arbitrary projectivities.

2 Notation

We will consider object space to be the 3-dimensional Euclidean space R^3 and represent points in R^3 as 3-vectors. Similarly, image space is the 2-dimensional Euclidean space R^2 and points are represented as 2-vectors. Euclidean space, R^3 is embedded in a natural way in projective 3-space $\mathcal{P}^3$by the addition of a plane at infinity. Similarly, R^2 may be embedded in the projective 2-space $\mathcal{P}^2$by the addition of a line at infinity. The simplicity of considering projections between $\mathcal{P}^3$and $\mathcal{P}^2$has led many authors to identify $\mathcal{P}^3$and $\mathcal{P}^2$as the object and images space. This point of view will not be followed here however, although when necessary we will consider points in R^2 and R^3 to as lying in $\mathcal{P}^2$and $\mathcal{P}^3$respectively, via the natural embedding.

Vectors will be represented as bold-face lower case letters, such as $\mathbf{x}$. Such a notation represents a column vector. The corresponding row vector will be denoted by $\mathbf{x}^\top$. The notation $\mathbf{x}$ usually denotes a vector in R^3, whereas $\mathbf{u}$ represents a vector in R^2. Elements in projective spaces $\mathcal{P}^3$and $\mathcal{P}^2$will be denoted with a tilde accent. For instance, $\tilde{\mathbf{x}}$ is a homogeneous 4-vector representing an element in $\mathcal{P}^3$, and $\tilde{\mathbf{u}}$ is a homogeneous 3-vector representing an element of $\mathcal{P}^2$.

The notation $\approx$ represents equality of matrices or homogeneous vectors up to an arbitrary non-zero factor. If $\mathbf{x} = (x, y, z)^\top$ is a 3-vector representing a point in R^3, then $\hat{\mathbf{x}}$ is the vector $(x, y, z, 1)^\top$. Similarly, if $\mathbf{u} = (u, v)^\top$, then $\hat{\mathbf{u}}$ represents the vector $(u, v, 1)^\top$.

The notation $a \doteq b$ means that a and b have the same sign. For instance $a \doteq 1$ has the same meaning as $a > 0$.

3 Projections in $\mathcal{P}^3$

A projection from $\mathcal{P}^3$into $\mathcal{P}^2$is represented by a 3×4 matrix P, whereby a point $\tilde{\mathbf{x}}$ maps to the point $\tilde{\mathbf{u}} \approx P\tilde{\mathbf{x}}$. It will be assumed that P has rank 3. Since P has 4 columns but rank 3, there is a unique point $\tilde{\mathbf{t}}$ such that $P\tilde{\mathbf{t}} = (0, 0, 0)^\top$. In other words, the projective transformation is undefined at the point $\tilde{\mathbf{t}}$, since $(0, 0, 0)^\top$ is not a valid homogeneous 3-vector. The point $\tilde{\mathbf{t}}$ will be called the *perspective centre* of the camera. We will assume that the perspective centre is not a point at infinity so we may write $\tilde{\mathbf{t}} \approx \hat{\mathbf{t}} = \begin{pmatrix} \mathbf{t} \\ 1 \end{pmatrix}$ where $\mathbf{t}$ is the perspective center as a point in R^3.

Now, the camera matrix P may be written in block form as $P = (M \mid \mathbf{c})$ where M is a 3×3 block and $\mathbf{c}$ is a column vector. Now

$$P\hat{\mathbf{t}} = (M \mid \mathbf{c})\begin{pmatrix} \mathbf{t} \\ 1 \end{pmatrix} = M\mathbf{t} + \mathbf{c} = 0 \ ,$$

and so $\mathbf{c} = -M\mathbf{t}$. In future, we will write $P = (M \mid -M\mathbf{t})$. Now since P has rank 3 and $-M\mathbf{t}$ is a linear combination of the columns of M, it follows that M must have rank 3. In other words, M is non-singular. Summarizing this discussion we have

Proposition 3.1. *If P is a camera transform matrix for a camera with perspective centre not at infinity, then P can be written as $P = (M \mid -M\mathbf{t})$ where M is a non-singular 3×3 matrix and $\mathbf{t}$ represents the perspective centre in R^3.*

There exist points in $\mathcal{P}^3$that are mapped to points at infinity in the image. To find what they are, we suppose that $\tilde{\mathbf{u}} = (u, v, 0)^\top = P\tilde{\mathbf{x}}$. Letting $\mathbf{p}_1{}^\top$, $\mathbf{p}_2{}^\top$

and $\mathbf{p}_3{}^\top$ be the rows of P, we see that $\mathbf{p}_3{}^\top\tilde{\mathbf{x}} = 0$. In other words, a point $\tilde{\mathbf{x}}$ in $\mathcal{P}^3$that maps to a point at infinity in the image must satisfy the equation $\tilde{\mathbf{x}}^\top\mathbf{p}_3 = 0$. Looked at another way, if $\mathbf{p}_3$ is taken as representing a plane in $\mathcal{P}^3$, then a point $\tilde{\mathbf{x}}$ lies on the plane $\mathbf{p}_3$ if and only if $\tilde{\mathbf{x}}^\top\mathbf{p}_3 = 0$. In other words, the condition for $\tilde{\mathbf{x}}$ to map to a point at infinity is the same as the condition for $\tilde{\mathbf{x}}$ to lie on the plane $\mathbf{p}_3$. Since $P\hat{\mathbf{t}} = 0$, we see in particular that $\mathbf{p}_3{}^\top\hat{\mathbf{t}} = 0$, and so $\hat{\mathbf{t}}$ lies on the plane $\mathbf{p}_3$. To summarize this paragraph, the set of points in $\mathcal{P}^3$mapping to a point at infinity in the image is a plane passing through the perspective centre and represented by $\mathbf{p}_3$, where $\mathbf{p}_3{}^\top$ is the last row of P. This plane will be called the *meridian plane* of the camera.

Restricting now to R^3, consider a point $\mathbf{x}$ in space, not lying on the meridian plane. It is projected by the camera with matrix P onto a point $\mathbf{u}$ where $w\hat{\mathbf{u}} = P\hat{\mathbf{x}}$ for some scale factor w. The value of w will vary continuously with $\mathbf{x}$ and the set of points where it vanishes is precisely the meridian plane. It follows that on one side of the meridian plane $w > 0$ and on the other side, $w < 0$. It can be shown, but is not used in this paper, that w is in fact proportional to the distance of $\mathbf{x}$ from the meridian plane.

Any real camera can only view points on one side of the meridian plane, those points that are "in front of" the camera. Points on the other side will not be visible. In order to distinguish the front of the camera from the back, a convention is necessary.

Definition 3.2. A camera matrix $P = (M \mid -M\mathbf{t})$ is said to be *normalized* if $\det(M) > 0$. If P is a normalized camera matrix, a point $\mathbf{x}$ is said to lie *in front of* the camera if $P\hat{\mathbf{x}} = w\hat{\mathbf{u}}$ with $w > 0$. Points $\mathbf{x}$ for which $w < 0$ are said to be *behind* the camera.

Any camera matrix may be normalized by multiplying it by -1 if necessary. It will **always** be assumed that camera matrices are normalized. The selection of which side of the camera is the front is simply a convention, consistent with the assumption that for a camera with matrix $(I \mid 0)$, points with positive z-coordinate lie in front of the camera. This is the usual convention in computer vision literature, used for instance in [Higgins-81].

The following statement expresses the fact that a camera sees only those points that lie in front of it.

Proposition 3.3. *A point $\mathbf{x}$ in R^3 is mapped to a point $\mathbf{u}$ in R^2 by a camera with normalized matrix P if and only if $w\hat{\mathbf{u}} = P\hat{\mathbf{x}}$ for some constant $w > 0$.*

4 Good Projectivities

A subset B of R^n is called convex if the line segment joining any two points in B also lies entirely within B. The convex hull of B, denoted $\bar{B}$ is the smallest convex set containing B.

Definition 4.4. Let B be a subset of R^n and let h be a projectivity of $\mathcal{P}^n$. The projectivity h is said to be a "good projectivity" with respect to the set B if $h^{-1}(L_\infty)$ does not meet $\bar{B}$, where L_∞ is the plane (or line) at infinity.

A good projectivity with respect to B is precisely one that preserves the convex hull of B. It may be verified that if h is a good projectivity with respect to B, then h^{-1} is a good projectivity with respect to $h(B)$. Details are omitted for the sake of brevity. We will be considering sets of points $\{\mathbf{u}_i\}$ and $\{\mathbf{u}'_i\}$ that correspond via a projectivity. When we speak of the projectivity being *good*, we will mean good with respect to the set $\{\mathbf{u}_i\}$.

An alternative characterization of good projectivities is given in the following theorem.

Theorem 4.5. *A projectivity $h : \mathcal{P}^n \rightarrow \mathcal{P}^n$ represented by a matrix H is good with respect to a set $B = \{\mathbf{u}_i\} \subset R^n - h^{-1}(L_\infty)$ if an only if there exist constants w_i, all of the same sign, such that $H\hat{\mathbf{u}}_i = w_i\hat{\mathbf{u}}'_i$*

Proof. To prove the forward implication, we assume that h is a good projectivity. By definition, constants w_i exist such that $H\hat{\mathbf{u}}_i = w_i\hat{\mathbf{u}}'_i$. What needs proof is that they all have the same sign. The value of w in the mapping $w\hat{\mathbf{u}}'_i = H\hat{\mathbf{u}}_i$ is a continuous function of the point $\mathbf{u}$. If $w_i > 0$ for some point $\mathbf{u}_i$, and $w_j < 0$ for another point $\mathbf{u}_j$, then there must some point $\mathbf{u}_\infty$ on the line segment joining $\mathbf{u}_i$ to $\mathbf{u}_j$ for which $w = 0$. This means that $h(\mathbf{u}_\infty)$ lies on the line at infinity, contrary to hypothesis.

Next, to prove the converse, we assume that there exist such constants w_i all of the same sign. Let S be the subset of R^n consisting of all points $\mathbf{u}$ satisfying the condition $H\hat{\mathbf{u}} = w\hat{\mathbf{u}}'$ such that w has the same sign as all w_i. The set S contains B and it is clear that $S \cap h^{-1}(L_\infty) = \emptyset$. All that remains to show is that S is convex, for then S must contain the convex hull of B. If $\mathbf{u}_i$ and $\mathbf{u}_j$ are points in S with corresponding constants w_i and w_j, then any intermediate point $\mathbf{u}$ between $\mathbf{u}_i$ and $\mathbf{u}_j$ must have w value intermediate between w_i and w_j. Consequently, the value of w must have the same sign as w_i and w_j, and so $\mathbf{u}$ lies in S also. This completes the proof. □

As just noted, if a projectivity is not good, then there are points in the convex hull for which w equals nought (0). For this reason, a projectivity that is not good will be called "naughty"[1].

This theorem gives an effective method of identifying good projectivities. The question remains whether good projectivities form a useful class. This question will be answered by the following theorem.

Theorem 4.6. *If B is a point set in a plane (the "object plane") in R^3 lying entirely in front of a projective camera, then the mapping from the object plane to the image plane defined by the camera is a good projectivity with respect to B.*

Proof. That there is a projectivity h mapping the object plane to the image plane is well known. What is to be proven is that the projectivity is good with respect to B. Let L be the line in which the meridian plane of the camera meets the object plane. Since B lies entirely in front of the camera, L does

[1]This terminology was suggested to me by David Forsyth

not meet the convex hull of B. However, by definition of the meridian plane $L = h^{-1}(L_\infty)$, where L_∞ is the line at infinity in the image plane. Therefore, h is a good projectivity with respect to B. $\square$

As an example, Fig. 1 shows an image of a comb and the image resampled according to a naughty projectivity. Most people will agree that the resampled image is unlike any view of a comb seen by camera or human eye. Nevertheless, the two images are projectively equivalent and will have the same projective invariants.

Note that if points $\mathbf{u}_i$ are visible in an image, then the corresponding object points must lie in front of the camera. Applying Theorem 4.6 to a sequence of imaging operations (for instance, a picture of a picture of a picture, etc), it follows that the original and final images in the sequence are connected by a planar projectivity which is good with respect to any point set in the object plane visible in the final image.

Similarly, if two images are taken of a set of point $\{\mathbf{x}_i\}$ in a plane, $\mathbf{u}_i$ and $\mathbf{u}'_i$ being correponding points in the two images, then there is a good projectivity (with respect to the $\mathbf{u}_i$) mapping each $\mathbf{u}_i$ to $\mathbf{u}'_i$, and so Theorem 4.5 applies, yielding the following proposition.

Proposition 4.7. *If $\{\mathbf{u}_i\}$ and $\{\mathbf{u}'_i\}$ are corresponding points in two views of a set of object points $\{\mathbf{x}_i\}$ lying in a plane, then there is a matrix H representing a planar projectivity such that $H\hat{\mathbf{u}}_i = w_i\hat{\mathbf{u}}'_i$ and all w_i have the same sign.*

This fact was pointed out to me by Charles Rothwell (private communication) and served as a starting point for the current investigation. Rothwell derived this result using the methods of [Sparr-92].

5 An integer valued invariant

Given a set of $N \geq n+2$ points $\{\mathbf{u}_i\}$, $i = 1, \ldots, N$ in R^n, it is possible to define an invariant of good projectivity as follows. Let $\mathbf{e}_1, \ldots \mathbf{e}_{n+2}$ be points in R^n such that $\{\hat{\mathbf{e}}_i\}$ form a canonical projective basis for $\mathcal{P}^n$. For $n = 2$, the points $(0,0)^\top$, $(1,0)^\top$, $(0,1)^\top$ and $(1,1)^\top$ will do. Assume that the points $\mathbf{u}_i$ are numbered in such a way that the first $n+2$ of them are in general position (meaning that no $n+1$ of them lie in a codimension 1 hyperplane). In this case, there exists a projectivity g (not necessarily good) such that $g(\mathbf{u}_i) = \mathbf{e}_i$ for $i = 1, \ldots, n+2$. Now, for each $i = 1, \ldots, N$ we define a value η_i as follows. If $g(\mathbf{u}_i)$ lies on the plane at infinity, we set $\eta_i = 0$. Otherwise, there exists a further $\mathbf{e}_i$ such that $g(\mathbf{u}_i) = \mathbf{e}_i$. If g is represented by a matrix G, then η_i is defined by the equation $G\hat{\mathbf{u}}_i = \eta_i\hat{\mathbf{e}}_i$. We show that, except for possible simultaneous negation, the values $\mathrm{sign}(\eta_i)$ are an invariant of good projectivity. Here $\mathrm{sign}(\eta_i)$ is defined to equal 1, -1 or 0 depending on whether η_i is positive, negative or zero respectively. The invariant value is of course dependent on the choice of canonical basis $\{\mathbf{e}_i\}$.

To prove the invariance, suppose that h is a good projectivity with respect to points $\{\mathbf{u}_i\}$ and let $h(\mathbf{u}_i) = \mathbf{u}'_i$. Consider the projectivity g' defined by

$g'(\mathbf{u}'_i) = \mathbf{e}_i$ for $i = 1, \ldots, n+2$. Values η'_i may be defined as before in terms of the projectivity g'. On the other hand, values w_i may be defined in terms of the projectivity h mapping each $\mathbf{u}_i$ to $\mathbf{u}'_i$ as in Theorem 4.5.

Since h and $g'^{-1} \circ g$ agree on a set of basis points, it follows that $h = g'^{-1} \circ g$. Consequently, $w_i = \eta_i / \eta'_i$. However, under the assumption that h is a good projectivity, all the w_i have the same sign, and so, for all i, we have $\eta_i \doteq \epsilon\, \eta'_i$, where $\epsilon = \pm 1$. In other words, the set of values $\mathrm{sign}(\eta_i)$ are an invariant under good projectivity, except for possible simultaneous negation.

It is possible to code the values η_i into a single number according to the formula

$$\chi(\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_N) = \left| \sum_{i=1}^{N} \mathrm{sign}(\eta_i) 3^{i-1} \right| \tag{1}$$

The value $\chi(\mathbf{u}_i)$ is an invariant under good projectivity of the ordered set of points $\mathbf{u}_i$. It will be called the *cheirality invariant* of the points.

6 Three dimensional point sets

We now consider three-dimensional point sets. The question that will be addressed is : "Under what conditions can points $\mathbf{u}_i$ and $\mathbf{u}'_i$ in two views be the images of a three dimensional point set $\mathbf{x}_i$ corresponding to two arbitrary uncalibrated cameras ?". One well-known necessary condition ([Higgins-81]) is the epipolar constraint, $\hat{\mathbf{u}}_i'^{\top} Q \hat{\mathbf{u}}_i = 0$ for all i and some rank-two matrix Q. We will ignore the effects of noise, so that the epipolar constraint equation will be assumed to hold exactly. The question is whether this is also a sufficient condition. The answer is no.

It will be assumed that there are sufficient points for the matrix Q to be determined unambiguously, that is at least 7 ([Hartley-92]) or 8 ([Higgins-81]) points. Under these conditions as shown in [Hartley-Gupta-92] and [Faugeras-92] it is possible to determine the location of points $\tilde{\mathbf{x}}_i$ and cameras P and P' such that $\hat{\mathbf{u}}_i \approx P\tilde{\mathbf{x}}_i$ and $\hat{\mathbf{u}}'_i \approx P'\tilde{\mathbf{x}}_i$, and furthermore, the choice is unique up to projectivity of $\mathcal{P}^3$. Assuming that none of the reconstructed points $\mathbf{x}_i$ is at infinity, we can write

$$\begin{aligned} w_i \hat{\mathbf{u}}_i &= P\hat{\mathbf{x}}_i \\ w'_i \hat{\mathbf{u}}'_i &= P'\hat{\mathbf{x}}_i \end{aligned} \tag{2}$$

If all the w_i and w'_i are positive, then according to Proposition 4.7 the points $\mathbf{x}_i$ map to points $\mathbf{u}_i$ and $\mathbf{u}'_i$ in the two images. Normally, this will not be the case. It is possible, however, that another choice of P, P' and $\mathbf{x}_i$ exists with the desired property.

We introduce some new terminology. A triplet $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ is called an *epipolar configuration* if Q is a rank 2 matrix satisfying the epipolar constraint equation $\hat{\mathbf{u}}_i'^{\top} Q \hat{\mathbf{u}}_i = 0$ for all i. A *weak realization* of $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ is a triplet $(P, P', \{\mathbf{x}_i\})$, where P and P' are a choice of normalized camera matrices corresponding to the essential matrix Q and the points $\{\mathbf{x}_i\}$ are object points satisfying the equations (2) for each i. A *strong realization* is such a triplet satisfying the additional condition that all the w_i and w'_i are positive. The triplet $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ is called a *feasible configuration* if a strong realization exists.

The following lemma sets notation and derives a basic technical result.

Lemma 6.8. *Let $(P, P', \{\mathbf{x}_i\})$ and $(\bar{P}, \bar{P}', \{\bar{\mathbf{x}}_i\})$ be two weak realization of a feasible configuration $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$. There exists a 4×4 matrix H such that $P \approx \bar{P}H$, $P' \approx \bar{P}'H$ and $\hat{\mathbf{x}}_i \approx H^{-1}\hat{\bar{\mathbf{x}}}_i$. Assume that P, P', $\bar{P}$ and $\bar{P}'$ are normalized and let constants ϵ, η_i, w_i and $\bar{w}_i$ be defined by the equations*

$$\begin{aligned} P &= \epsilon \bar{P} H \\ \hat{\mathbf{x}}_i &= \eta_i H^{-1} \hat{\bar{\mathbf{x}}}_i \\ w_i \hat{\mathbf{u}}_i &= P \hat{\mathbf{x}}_i \\ \bar{w}_i \hat{\bar{\mathbf{u}}}_i &= \bar{P} \hat{\bar{\mathbf{x}}}_i \end{aligned} \tag{3}$$

Then $w_i \bar{w}_i \epsilon \eta_i \doteq 1$.

If constants w'_i, $\bar{w}'_i$, and ϵ' are defined in a similar way then $w'_i \bar{w}'_i \epsilon' \eta_i \doteq 1$.

Proof. The existence of the matrix H is proven in [Hartley-Gupta-92]. Now,

$$\begin{aligned} w_i \hat{\mathbf{u}}_i &= P \hat{\mathbf{x}}_i \\ &= \epsilon \eta_i \bar{P} H H^{-1} \hat{\bar{\mathbf{x}}}_i \\ &= \epsilon \eta_i \bar{w}_i \hat{\mathbf{u}}_i \end{aligned}$$

whence $w_i = \epsilon \eta_i \bar{w}_i$. Multiplying each side of this equation by w_i gives the required result. The proof for the primed quantities is of course the same. $\square$

A further useful technical result follows.

Lemma 6.9. *Let H be the matrix*

$$H = \begin{pmatrix} I & 0 \\ k\mathbf{v}^\top & k \end{pmatrix} .$$

Then with the notation used in Lemma 6.8, $\epsilon \doteq \hat{\mathbf{v}}^\top \hat{\mathbf{t}}$, $\epsilon' \doteq \hat{\mathbf{v}}^\top \hat{\mathbf{t}}'$ and for each i, $\eta_i \doteq k \hat{\mathbf{v}}^\top \hat{\mathbf{x}}_i$, where $\mathbf{t}$ and $\mathbf{t}'$ are the perspective centres of P and P'. (Remember that $\doteq$ denotes equality of sign.)

Proof. One verifies that

$$H^{-1} = \begin{pmatrix} I & 0 \\ -\mathbf{v}^\top & k^{-1} \end{pmatrix} .$$

Let $P = (M \mid -M\mathbf{t})$ with $\det(M) > 0$ and $\bar{P} = (\bar{M} \mid -\bar{M}\bar{\mathbf{t}})$ with $\det(\bar{M}) > 0$. The from the $\epsilon \bar{P} = P H^{-1}$ it follows that $\epsilon \bar{M} = M(I + \mathbf{t}\mathbf{v}^\top)$. Taking determinants and signs gives

$$\epsilon \doteq \det(I + \mathbf{t}\mathbf{v}^\top) = 1 + \mathbf{v}^\top \mathbf{t} = \hat{\mathbf{v}}^\top \hat{\mathbf{t}}$$

as required. The same proof holds for ϵ'.

From (3) we have $H\hat{\mathbf{x}}_i = \eta_i \hat{\bar{\mathbf{x}}}_i$. Multiplying this out and considering only the last component yields $\eta_i = k(\mathbf{v}^\top \mathbf{x}_i + 1) = k\hat{\mathbf{v}}^\top \hat{\mathbf{x}}_i$ as required. $\square$

Applying Lemma 6.8 to the case where one of the realizations is a strong realization leads to a necessary and sufficient condition for an epipolar configuration to be feasible.

Theorem 6.10. *Let* $(P, P', \{\mathbf{x}_i\})$ *be any weak realization of an epipolar configuration* $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ *and let* w_i *and* w'_i *be defined as in (2). There exists a strong realization* $(\bar{P}, \bar{P}', \bar{\mathbf{x}}_i)$ *of* $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ *if and only if* $w_i w'_i$ *has the same sign for all* i.

Proof. We begin by proving the *if* part of this theorem, and apply Lemma 6.8 to the case where $(\bar{P}, \bar{P}', \bar{\mathbf{x}}_i)$ is a strong realization. In this case, $\bar{w}_i \doteq 1$ and so $w_i \eta_i \epsilon \doteq 1$. Similarly, $w'_i \eta_i \epsilon' \doteq 1$. Therefore $w_i w'_i \eta_i^2 \epsilon \epsilon' \doteq 1$. from which it follows that $w_i w'_i \doteq \epsilon \epsilon'$ which is constant for all i.

Now, we turn to prove the converse. Let X^+ be the set of points $\mathbf{x}_i$ such that $w_i > 0$ and let X^- be the set of points such that $w_i < 0$. The sets X^+ and X^- are separated by the meridian planes of each of the cameras. Now, we seek a plane that separates X^- from X^+ and satisfies the additional condition that the perspective centres of the two camera lie on the same side of the plane if $w_i w'_i > 0$ for all i, or on opposite sides of the plane if $w_i w'_i < 0$ for all i. Such a plane can easily be found by slightly displacing the meridian plane of one of the cameras[2].

Let this separating plane be represented by a 4-vector $\hat{\mathbf{v}}$. The condition that both perspective centres $\mathbf{t}$ and $\mathbf{t}'$ lie on the same or opposite sides of the plane may be written as $\hat{\mathbf{v}}^\top \hat{\mathbf{t}} \doteq \kappa$ and $\hat{\mathbf{v}}^\top \hat{\mathbf{t}}' \doteq \kappa w_i w'_i$ where κ is some non-zero value and $\mathrm{sign}(w_i w'_i)$ is a constant for all i by hypothesis. The condition that the plane $\hat{\mathbf{v}}$ separates X^- from X^+ may be written as $\hat{\mathbf{v}}^\top \hat{\mathbf{x}}_i \doteq \xi w_i$ for some constant ξ. Now, let H be the matrix

$$H = \begin{pmatrix} I & 0 \\ \kappa \xi \mathbf{v}^\top & \kappa \xi \end{pmatrix} .$$

Then according to Lemma 6.9, $\epsilon \doteq \hat{\mathbf{v}}^\top \hat{\mathbf{t}} \doteq \kappa$, $\epsilon' \doteq \hat{\mathbf{v}}^\top \hat{\mathbf{t}}' \doteq \kappa w_i w'_i$ and $\eta_i \doteq \kappa \xi \hat{\mathbf{v}}^\top \hat{\mathbf{x}}_i \doteq \kappa \xi^2 w_i$. Now substituting into the equation $w_i \bar{w}_i \epsilon \eta_i \doteq 1$ from Lemma 6.8 yields $w_i \bar{w}_i \kappa^2 \xi^2 w_i \doteq 1$ from which it follows that $\bar{w}_i \doteq 1$ as required. Similarly, from the equation $w'_i \bar{w}'_i \epsilon' \eta_i \doteq 1$ we derive $w'_i \bar{w}'_i \kappa w_i w'_i \kappa \xi^2 w_i \doteq 1$, from which it follows that $\bar{w}_i \doteq 1$. This shows that $(\bar{P}, \bar{P}', \{\bar{x}_i\})$ is a strong realization as required. $\square$

Since the epipolar configuration derived from two images of a real scene must have a strong realization, this theorem gives a necessary and sufficient condition for a set of image correspondences to be realizable as a three dimensional scene. Theorem 6.10 is illustrated in Fig 2.

For planar object sets, Theorem 4.6 established the existence of a good projectivity between the object plane and the image plane. For *non-planar* objects seen in two views, strong realizations of the epipolar configuration take the rôle played by sets of image points in the two dimensional case.

[2] For this construction to work, it seems necessary to make the additional assumption that the point set $\{\mathbf{u}_i\}$ is bounded in the image plane. This assumption will be true for any reasonable pinhole camera, which can not have an image of infinite extent.

Theorem 6.11. *Let* $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ *be an epipolar configuration and let* $(P, P', \{\mathbf{x}_i\})$ *and* $(\bar{P}, \bar{P}', \{\bar{\mathbf{x}}_i\})$ *be two separate strong realizations of the configuration. Then the projectivity mapping each point* $\mathbf{x}_i$ *to* $\bar{\mathbf{x}}_i$ *is good.*

Proof. With notation as in (3), $w_i \doteq \bar{w}_i \doteq 1$, and hence from Lemma 6.8, $\eta_i \epsilon \doteq 1$, which means that all η_i have the same sign. Therefore, by Theorem 4.5, H is a good projectivity. □

The particular case where one of the two realizations is the "correct" realization is of interest. It is the analogue in three dimensions of Proposition 4.6.

Corollary 6.12. *If* $\{\mathbf{x}_i\}$ *are points in* R^3*, image coordinates* $\{\mathbf{u}_i\}$ *and* $\{\mathbf{u}'_i\}$ *are corresponding image points in two uncalibrated views,* Q *is the essential matrix derived from the image correspondences* $\mathbf{u}_i \leftrightarrow \mathbf{u}'_i$ *and* $(P, P', \{\bar{\mathbf{x}}_i\})$ *is a strong realization of the triple* $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$*, then there is a good projectivity taking each* $\mathbf{x}_i$ *to* $\bar{\mathbf{x}}_i$*.*

From this corollary, we can deduce one of the main results of this paper.

Theorem 6.13. *Let* $(P, P', \{\mathbf{x}_i\})$ *and* $(\bar{P}, \bar{P}', \{\bar{\mathbf{x}}_i\})$ *be two different reconstructions of 3D scene geometry derived as strong realizations of possibly different epipolar configurations corresponding to possibly different pairs of images of a 3D point set. Then there is a good projectivity mapping each point* $\mathbf{x}_i$ *to* $\bar{\mathbf{x}}_i$*.*

What this theorem is saying is that if a point set in R^3 is reconstructed as a strong realization from two separate pairs of views, then the two results are the same up to a good projectivity.

Proof. By corollary 6.12 there exist good projectivities mapping each of the sets of reconstructed points $\{\mathbf{x}_i\}$ and $\{\bar{\mathbf{x}}_i\}$ to the actual 3D locations of the points. The result follows by composing one of these projectivities with the inverse of the other. □

7 Orientation

We now consider the question of image orientation. A mapping h from R^n to itself is called orientation-preserving at a point $\mathbf{x}$ if the Jacobian of h has positive determinant at $\mathbf{x}$. Otherwise h is called orientation reversing. Reflection of points of R^n with respect to a hyperplane (that is mirror imaging) is an example of an orientation reversing mapping. A projectivity h from $\mathcal{P}^n$ to itself restricts to a mapping from $R^n - h^{-1}(L_\infty)$ to R^n, where L_∞ is the hyperplane (line, plane) at infinity. Consider the case $n = 3$ and let H be a 4×4 matrix representing the projectivity h. We wish to determine at which points $\mathbf{x}$ in $R - h^{-1}(L_\infty)$ the map h is orientation preserving. It may be verified (quite easily using Mathematica [Wolfram-88]) that if $H\hat{\mathbf{x}} = w\hat{\mathbf{x}}'$ and J is the Jacobian of h evaluated at $\mathbf{x}$, then $\det(J) = \det(H)/w^4$. This gives the following result.

Proposition 7.14. *A projectivity h of $\mathcal{P}^3$ represented by a matrix H is orientation preserving at any point in $R^3 - h^{-1}(L_\infty)$ if and only if* $\det(H) > 0$.

Of course, the concept of orientability may be extended to the whole of $\mathcal{P}^3$, and it may be shown that h is orientation-preserving on the whole of $\mathcal{P}^3$if and only if $\det(H) > 0$. The essential feature here is that as a topological manifold, $\mathcal{P}^3$is orientable. The situation is somewhat different for $\mathcal{P}^2$, which is not orientable as a topological space. In this case, with notation similar to that used above, it may be verified that $\det(J) = \det(H)/w^3$. Therefore, the following proposition is true.

Proposition 7.15. *A projectivity h of $\mathcal{P}^2$is orientation preserving at a point* $\mathbf{u}$ *in $R^2 - h^{-1}(L_\infty)$ if and only if $w\det(H) > 0$, where $H\hat{\mathbf{u}} = w\hat{\mathbf{u}}'$.*

This theorem allows us to strengthen the statement of Theorem 4.5 somewhat.

Corollary 7.16. *If h is a good projectivity of $\mathcal{P}^2$with respect to a set of points $\{\mathbf{u}_i\}$ in R^2, then h is either orientation-preserving or orientation-reversing at all points $\mathbf{u}_i$. Suppose the matrix H corresponding to h is normalized to have positive determinant (by possible multiplication by -1) and let $H\hat{\mathbf{u}}_i = w_i\hat{\mathbf{u}}'_i$. Then h is orientation-preserving if and only if $w_i > 0$ for all i.*

An example where Corollary 7.16 applies is in the case where two images of a planar object are taken from the same side of the object plane. In this case, an orientation-preserving good projectivity will exist between the two images. Consequently, all the w_i defined with respect to a matrix H will be positive, provided that H is normalized to have positive determinant.

The situation in 3-dimensions is rather more involved and more interesting. Two sets of points $\{\mathbf{x}_i\}$ and $\{\bar{\mathbf{x}}_i\}$ that correspond via a good projectivity are said to be *oppositely oriented* if the projectivity is orientation-reversing. This definition extends also to two strong realizations $(P, P', \{\mathbf{x}_i\})$ and $(\bar{P}, \bar{P}', \{\bar{\mathbf{x}}_i\})$ of a common epipolar configuration $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$, since in view of Theorem 6.11 the point sets are related via a good projectivity. Whether or not oppositely oriented strong realizations exist depends on the imaging geometry. Common experience provides some clues here. In particular a stereo pair may be viewed by presenting one image to one eye and the other image to the other eye. If this is done correctly, then the brain perceives a 3-D reconstruction of the scene (a strong realization of the image pair). If, however, the two images are swapped and presented to the opposite eyes, then the perspective will be reversed – hills become valleys and vice versa. In effect, the brain is able to compute two oppositely oriented reconstructions of the image pair. It seems, therefore, that in certain circumstances, two oppositely oriented realizations of an image pair exist. It may be surprising to discover that this is not always the case, as is shown in the following theorem.

Theorem 7.17. *Let $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ be an epipolar configuration and let $(P, P', \{\mathbf{x}_i\})$ be a strong realization of $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$. There exists a different oppositely oriented strong realization $(\bar{P}, \bar{P}', \{\bar{\mathbf{x}}_i\})$ if and only if there exists a plane in R^3 such that the perspective centres of both cameras P and P' lie on one side of the plane, and the points $\mathbf{x}_i$ lie on the other side.*

Before proving this theorem, we need a lemma.

Lemma 7.18. *Let* $(P, P', \{\mathbf{x}_i\})$ *be a strong realization of an epipolar configuration* $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$*. Then there exists a similarly oriented strong realization* $(\bar{P}, \bar{P}', \{\bar{\mathbf{x}}_i\})$ *for which* $\bar{P} = (I \mid 0)$*.*

Proof. Suppose $P = (M \mid -M\mathbf{t})$, with $\det(M) > 0$. Then multiplication by the matrix

$$H = \begin{pmatrix} M^{-1} & \mathbf{t} \\ 0 & 1 \end{pmatrix}$$

transforms P to the required form. Furthermore, H^{-1} defines an orientation-preserving good projectivity on the points $\mathbf{x}_i$. □

Now, we may prove the theorem.

Proof. (Theorem 7.17) In light of Lemma 7.18 it may be assumed that P' and $\bar{P}'$ are both of the form $(I \mid 0)$, because an oppositely oriented pair of realizations exist if and only if an oppositely oriented pair exist satisfying this additional condition.

Let us assume that such an oppositely oriented pair of strong realizations exists and H represents the orientation-reversing good projectivity relating them. We define ϵ, ϵ' and η_i as in (3). If necessary, H may be multiplied by a constant so that $\epsilon' \doteq 1$. Since $w_i \doteq w'_i \doteq \bar{w}_i \doteq \bar{w}'_i \doteq 1$, it follows from Lemma 6.8 that $\eta_i \doteq 1$ for all i and $\epsilon \doteq 1$. From the equation $(I \mid 0)H = (I \mid 0)$ the form of H may be deduced :

$$H = \begin{pmatrix} I & 0 \\ k\mathbf{v}^\top & k \end{pmatrix}$$

for some 3-vector $\mathbf{v}$ and, since H is orientation reversing, $k \doteq -1$.

Now, according to Lemma 6.9, $\eta_i \doteq k\hat{\mathbf{v}}^\top\hat{\mathbf{x}}_i$, and since $\eta_i \doteq 1$ and $k \doteq -1$ it follows that $\hat{\mathbf{v}}^\top\hat{\mathbf{x}}_i \doteq -1$. This condition may be interpreted as meaning that all the $\mathbf{x}_i$ lie on one side of the plane defined by $\hat{\mathbf{v}}$.

On the other hand, by applying Lemma 6.9, we get $\hat{\mathbf{v}}^\top\hat{\mathbf{t}} \doteq \epsilon \doteq 1$ and $\hat{\mathbf{v}}^\top\hat{\mathbf{t}}' \doteq \epsilon' \doteq 1$. These equations mean that $\mathbf{t}$ and $\mathbf{t}'$ lie on the opposite side of the plane $\hat{\mathbf{v}}$ from all the points $\mathbf{x}_i$. This completes the *only if* part of the proof.

The converse may be proven by working backwards through this proof. Assuming the existence of a separating plane $\hat{\mathbf{v}}$ one constructs the orientation reversing matrix H as above and verifies that the resulting $(\bar{P}, \bar{P}', \{\bar{\mathbf{x}}_i\})$ is a strong realization. □

Note that the existence of such a separating plane as described in Theorem 7.17 may be checked using any strong realization.

8 3D cheirality invariants

The cheirality invariant of a set of points may be computed from two views by constructing a strong realization of the epipolar configuration and then invoking

Theorem 6.13. If in addition each pair of views is discovered to satisfy the condition of Theorem 7.17 then the orientation of the set of points with respect to a canonical basis gives a further invariant.

In general, finding a strong realization involves substantial computation. It is therefore convenient to be able to compute the cheirality invariant of a set of points from a weak realization. This may be done using the following theorem

Theorem 8.19. *Suppose $(P, P', \{\mathbf{x}_i\})$ is a weak realization of an epipolar configuration $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$ and let constants $\breve{\eta}_i$ be defined for each $\mathbf{x}_i$ as in the definition of the cheiral invariant. Suppose that $P\hat{\mathbf{x}}_i = w_i \hat{\mathbf{u}}_i$ and define $\eta_i = \breve{\eta}_i w_i$, then $\left|\sum_{i=1}^{N} \mathrm{sign}(\eta_i) 3^{i-1}\right|$ is the cheiral invariant of a strong realization of $(Q, \{\mathbf{u}_i\}, \{\mathbf{u}'_i\})$.*

Details of the proof will not be given. It is simply a matter of considering the composition of two projectivities : from the strong realization to the weak realization and from the weak realization to the canonical frame.

9 Experimental results

In considering real images of 3-D configurations it is necessary to take into account the effects of noise. In particular, because of measurement inaccuracies, it will (virtually) never be the case that a point $\mathbf{x}_i$ in a strong realization will map by chance exactly onto the plane at infinity under the mapping to the canonical basis. For this reason, in practical experiments I have preferred to define the cheiral invariant by interpreting the values η_i as bits of a binary integer : $\eta_i > 0$ corresponds to a 1 bit and $\eta_i < 0$ to a 0 bit. In some cases, a value of η_i will lie so close to 0 variations due to noise can swap its sign. For robust evaluation of a cheiral invariant value, it is necessary to select a noise model and determine how errors in the input data affect the sign of each η_i. In the following discussion, noise effects are ignored, however.

In [Hartley-93] projective invariants of 3D point sets were discussed. As an experiment in that paper, a set of images of some model houses were acquired. Figures 3, 4 and 5 show the three images. Corresponding vertices were selected by hand from among those detected automatically. The 13 vertices used are shown in Fig 6.

Six sets of six points were chosen as in the following table which shows the indices of the points as given in Fig 6.

$$\begin{array}{lcrrrrrrl} S_1 & = & \{1, & 2, & 3, & 6, & 9, & 10\} & , \\ S_2 & = & \{2, & 4, & 6, & 8, & 10, & 12\} & , \\ S_3 & = & \{1, & 3, & 5, & 7, & 9, & 11\} & , \\ S_4 & = & \{1, & 2, & 3, & 6, & 7, & 8\} & , \\ S_5 & = & \{1, & 4, & 7, & 10, & 13, & 12\} & , \\ S_6 & = & \{2, & 5, & 8, & 11, & 12, & 13\} & \end{array}$$

From image correspondences in two views (Figs 3 and 4) the essential matrix Q was found and a weak realization $(P, P', \{\mathbf{x}_i\})$ was computed. For each of the six sets of indices i shown above a complete projective invariant of the points $\{x_i\}$ was computed by mapping the first five points onto a canonical basis. The coordinates of the mapped sixth point constitute a projective invariant of the set of six points.

This computation was repeated with a different pair of views (Figs 4 and 5). Theory predicts that the invariants should have the same value when computed from different views, and should distinguish between non-equivalent point sets.

Table (4) shows the comparison of the computed invariant values.

$$\begin{array}{cccccc}
\mathbf{0.026} & 0.970 & 0.975 & 0.619 & 0.847 & 0.823 \\
0.995 & \mathbf{0.015} & 0.064 & 0.841 & 0.252 & 0.548 \\
0.967 & 0.066 & \mathbf{0.013} & 0.863 & 0.276 & 0.516 \\
0.617 & 0.830 & 0.873 & \mathbf{0.016} & 0.704 & 0.752 \\
0.861 & 0.238 & 0.289 & 0.708 & \mathbf{0.005} & 0.590 \\
0.828 & 0.544 & 0.519 & 0.719 & 0.574 & \mathbf{0.026}
\end{array} \tag{4}$$

The (i, j)-th entry of the table shows the distance according to an appropriate metric between the invariant of set S_i as computed from the first image pair with that of set S_j as computed from the second image pair. The diagonal entries of the matrix (in bold) should be close to 0.0, which indicates that the invariants had the same value when computed from different pairs of views.

Although the projective invariants computed here are quite effective at discriminating between different point sets, indicated by the fact that most off-diagonal entries are not close to zero, entries $(2, 3)$ and $(3, 2)$ are small indicating that the point sets numbered 2 and 3 are close to being equivalent up to projectivity.

Next, the cheirality invariants for each of the point sets were computed from the weak realization using the method described here. The computed values for each of the six point sets were as follows : $\chi(S_1) = 28$, $\chi(S_2) = 3$, $\chi(S_3) = 59$, $\chi(S_4) = 60$, $\chi(S_5) = 21$, $\chi(S_6) = 27$. As expected these invariant values were the same whether computed using the first pair of views or the second pair. Note that the cheirality invariant clearly distinguishes point sets 2 and 3. In fact, all six point sets are distinguished.

Reordering : Although there are no invariants of projectivity for 5 points in $\mathcal{P}^3$, the cheirality invariant is defined. In order to estimate its effectiveness for distinguishing different configurations the following experiment was carried out. Five points in $\mathcal{P}^3$were selected and the cheirality invariant computed for all permutations of the five points. The result was that 10 different invariant values were found (out of 16 possible), each one occuring 12 times. It may be seen that this will be true whichever 5 points are selected (though the invariant values will be different). In short, there is about one chance in 10 that two sets of five arbitrarily selected points will have the same cheirality.

When this experiment was carried out with 6 points arbitrarily chosen the results were seen to vary according to the particular configuration of the points. For various choices of points it was seen that the probability of getting a chance match for arbitrary permutations of the point set is about one chance in 20 or 30.

Conclusions : These results show that the cheirality invariant is quite effective at distinguishing between arbitrary sets of points. Given the relative ease with which the cheirality invariant may be computed, it may be extremely useful in grouping points. In addition, it may conveniently be used as an indexing function in an object recognition system. It has been demonstrated that the cheirality invariant gives supplementary information that is not available in projective invariants. As a theoretical tool, the cheirality invariants provide conditions under which image point matches may be realized by real point configurations.

References

[Faugeras-92] Faugeras, O., "What can be seen in three dimensions with an uncalibrated stereo rig?", Proc. of ECCV-92, G. Sandini Ed., LNCS-Series Vol. 588, Springer- Verlag, 1992, pp. 563 – 578.

[Higgins-81] H.C. Longuet-Higgins, *"A computer algorithm for reconstructing a scene from two projections,"* Nature, Vol. 293, 10 Sept. 1981.

[Hartley-92] R. Hartley, "Invariants of 3D point sets", submitted for publication.

[Hartley-Gupta-92] R. Hartley, R. Gupta and T. Chang, "Stereo from Uncalibrated Cameras", Proceedings Computer Vision and Pattern Recognition conference (CVPR-92), 1992.

[Hartley-93] R. Hartley, "Invariants of Lines in Space", In these proceedings.

[Wolfram-88] Wolfram, S., Mathematica, "A System for Doing Mathematics by Computer," Addison-Wesley, Redwood City, California, 1988.

[Rothwell-92] Charles A. Rothwell, Andrew Zisserman, David A. Forsyth and Joseph L. Mundy, "Canonical Frames for Planar Object Recognition", Proc. of ECCV-92, G. Sandini Ed., LNCS-Series Vol. 588, Springer- Verlag, 1992, pp. 757-772.

[Sparr-92] Gunnar Sparr, "Depth Computations from Polyhedral Images", Proc. of ECCV-92, G. Sandini Ed., LNCS-Series Vol. 588, Springer- Verlag, 1992, pp. 378 – 386.

Figure 1. At the left a comb. At the right a naughty projection of the comb.

Figure 2. Each camera is shown symbolically as a line representing the meridian plane and an arrow indicating the direction of the front of the camera. Each diagram represents a weak realization of an epipolar configuration. The two top configurations of points and cameras satisfy the condition of Theorem 6.10 and may be converted to strong realizations. The two lower configurations do not, and hence can not be weak realizations of a real scene.