The Non-existence of General-case View-invariants[1]

J. Brian Burns, Richard S. Weiss and Edward M. Riseman

Computer Science Dept.
University of Massachusetts
Amherst, MA 01003

Abstract

The recognition of 3D objects from a single, unknown view can be facilitated by the use of image features that are invariant with respect to view position. This paper concerns invariants that are functions of projected 3D point sets. View invariants exist for special classes of objects (3D point sets), such as those constrained to planes; however, there does not exist a function that is view-invariant for all possible 3D point sets of size n, for any n. The non-existence of such a *general-case* view invariant is shown for the true perspective, weak perspective, and orthographic projection models.

Though there are no general-case weak-perspective invariants, there are special-case invariants of practical importance. In the remaining portion of the paper, the special-case weak-perspective invariants cited in the literature are expressed in terms of linear dependences.

[1]Published in *Geometric Invariance in Machine Vision*, J. Mundy and A. Zisserman (Eds.), MIT Press, 1992. This research was supported by the Defense Advanced Research Projects Agency under contract F30602-87-C-0140 and (via TACOM) under contract DAAE07-91-C-R035, and by Army ETL contract DACA76-89-C-0017.

1 Introduction

Familiar three-dimensional objects are typically sensed visually in the form of a two-dimensional projection. While model-independent understanding of 3D surface structure is possible from motion, stereo, shading and texture, quite often these cues will be unavailable, unreliable or provide only a rough indication of the object structure. Thus, one of the outstanding problems in visual object recognition is the development of systems capable of recognizing 3D objects from a single, unknown view using 2D features in the image. An image feature is useful in discriminating some subset of the modelled objects to the degree that its distribution of values with respect to each object is narrow and the distributions with respect to the different objects are well separated. Since the camera viewpoint may not be constrained (or only partly constrained), clearly the usefulness of a feature of the object projection is a function of how it varies with view.

Ideally, a feature would be *view-invariant*, that is, unaffected by change in view. However, the contribution of this paper is a theoretical result that establishes that there are no *general-case* view invariant features for any number of 3D points, given a single view using true perspective, weak perspective or orthographic projection models. That is, there does not exist a function that is view-invariant for every possible 3D point set of a given size n, for any n [2]. If one assumes constraints on the 3D point set being viewed, then there may be useful invariants; however, without any assumed constraints, there are no useful invariants. The result is of fundamental importance because if such a view invariant existed, it would be a recognition feature with a value distribution of zero variance for every rigid 3D object. Only added noise in the image, together with small separation in the values for the different objects, would hinder its usefulness in discriminating arbitrary collections of objects. This point-set result trivially extends to sets of line segments, polyhedral surfaces or any other structure that can be completely characterized by an *arbitrary* finite set of points. When one can determine constraints on the structures being viewed, interesting invariants may be derivable: in the case of the polyhedral surfaces, useful invariants can be assumed if one

[2] It can be argued that even the point count is not invariant since there are some views where multiple 3D points project to the same image point, thus changing the count. Technically these views are not degenerate in our sense of the term since the feature is still defined from these views (Section 2).

can determine which points lie on the same face and these coplanar points number greater than three.

The absence of general-case view-invariant features has led some researchers to use *special-case* invariants, features that are only view-invariant for special configurations of 3D points or sets of views. For example, the viewpoint was restricted in [Tuc88] so that all projections were related by 2D rigid motions. This allows very effective use of features that are invariant for those motions, such as angle and distance. Features called *affine coordinates* are invariant to view if one assumes, as in [Lam88a], that the object points are planar and the projection can be approximated by the weak perspective model. In Section 3, the special-case weak-perspective invariants cited in the literature are expressed in terms of linear dependences. Forsyth *et al* use a true perspective model and assume a planar surface. They have shown that by fitting two conics to two image curves projected from this surface, one can also obtain an invariant that depends on the coefficients of the conics [For90]. Important invariants may also be associated with special-case configurations of points other than planar; for example, Forsyth *et al* are currently investigating the cases of cubics and surfaces of revolution [For91]. The property of view-invariance can also be useful in other areas of vision, such as the use of the cross-ratio in scene reconstruction [Moh90]. This invariant requires that four points lie on a straight line, or dually that four lines intersect at a point. However, as is shown in this paper, invariants that do not require any restrictions are impossible, and recognition systems for the unrestricted problem cannot be expected to be as effective as the above special-case schemes.

Prior to our work establishing this result [Bur90], the non-existence of general-case invariants does not seem to have been determined [Ahu68, Low85, Bin87, Ali87, Lam88b, Wei88]. In a chapter surveying projective geometry, Duda et al., explicitly mention the current lack of understanding on the subject of invariants of 2D projections of 3D points [Dud73]. In a study subsequent to ours, Clemens and Jacobs established a special-case of the result for the weak perspective projection model [Cle91]. Recently we have seen increased efforts toward in understanding how projective invariants can be used in vision [For90, For91]. This provides a relevant context for the results presented here.

2 General-Case invariants

This section of the paper is devoted to a proof that there are no view invariant features for the general case of n three-dimensional points under perspective projection. This result is then extended to orthographic and weak perspective projection models in the section that follows.

More specifically, we establish here that there is no function of the perspective projection of n 3D points, for any integer n, that is both a general case view invariant and non-trivial. We consider an invariant function to satisfy the general case if for each 3D point set, it maps the set into the same value under all possible view transformations. We consider the invariant function to be non-trivial if it is not a constant; i.e., it does not map all point sets into the same value.

2.1 Definitions

The perspective object-to-image transformation is

$$\pi_{R,\vec{T}}(\vec{P}) = (F/z) \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} (R\vec{P} + \vec{T})$$

for 3D point[3] $\vec{P}$ at depth z, 3D rotation R, 3D translation $\vec{T}$ and focal length F. Each $(R, \vec{T})$ represents a distinct *view* of the point.

We are considering features that are functions of the projections of 3D point sets and are defined almost everywhere. More precisely, a feature is some function $f : \mathcal{R}^{2n} \rightarrow \mathcal{R}$, where n is some fixed number of projected points, and the function is restricted in the following sense. A view is *degenerate* for a given feature and 3D point set if the feature is undefined for the projection of the point set from this view. A *3D point set* is degenerate for a feature if the set of degenerate views for this feature and point set is not measure zero. We are considering features that have at most a measure zero collection of degenerate point sets, a weak restriction that should not limit the result in any significant way.

The projection of the 3D point set $S \in \mathcal{R}^{3n}$ will be compactly represented as $\pi_{R,\vec{T}}(S)$, and the value for a given feature f, view $(R, \vec{T})$ and 3D point

[3]Henceforth, lower case signifies positions in the image coordinate frame, and upper case, positions in the object coordinate frame.

set S is $f(\pi_{R,\vec{T}}(S))$. Since the point sets are ordered, whenever we consider multiple sets that are themselves indexed, we will refer to the jth point of the ith set as $\vec{P}_{i,j}$.

Def: A feature $f : \mathcal{R}^{2n} \rightarrow \mathcal{R}$ is *view-invariant* for a given 3D point set $S \in \mathcal{R}^{3n}$, if for all non-degenerate pairs of views $(R_1, \vec{T}_1, R_2, \vec{T}_2)$,

$$f(\pi_{R_1,\vec{T}_1}(S)) = f(\pi_{R_2,\vec{T}_2}(S)).$$

Def: A feature $f : \mathcal{R}^{2n} \rightarrow \mathcal{R}$ is a *general-case* view-invariant if it is view-invariant for all 3D point sets $S \in \mathcal{R}^{3n}$. In other words, in contrast to the special-cases discussed in the introduction, there are no constraints assumed for the 3D point sets.

There is always a trivial f that can satisfy the above property: any constant feature, where a *constant* feature has the same value for all 3D point sets and views. Thus, the following is another important property for a feature.

Def: A feature $f : \mathcal{R}^{2n} \rightarrow \mathcal{R}$ is *non-trivial* if there exist two different, non-degenerate point sets S_1 and S_2, and views of these point sets $(R_1, \vec{T}_1)$ and $(R_2, \vec{T}_2)$, such that $f(\pi_{R_1,\vec{T}_1}(S_1)) \neq f(\pi_{R_2,\vec{T}_2}(S_2))$.

Much of classical projective geometry studies transformations from one space to another of the same dimension; for example, from one plane to another. However, we are considering the 2D projection of point sets that are not constrained to be coplanar, and we must therefore generalize some standard definitions. These concepts are required to establish the main result; their use will be explained in the section that follows.

Def: *View correspondence.* Two point sets of the same size are in view correspondence if there exist a camera position and a pencil of rays through the focal point such that there is a one-one correspondence between the sets, with corresponding points on the same ray (for example, sets S_1 and S_2 in Figure 1). In addition, all rays must be constrained to the same half-space to create an image by sectioning the pencil with a plane. Clearly, two such 3D point sets project to the same image if the camera focal point is the ray intersection point and the rays are sectioned by the image plane; in other words, $\pi_{R_1,\vec{T}_1}(S_1) = \pi_{R_2,\vec{T}_2}(S_2)$, for some $(R_1, \vec{T}_1, R_2, \vec{T}_2)$.

Def: *Transitive view correspondence.* Two point sets are in transitive view correspondence if they can be connected by a chain of view correspondences. For example, in Figure 1, S_1 and S_3 are in transitive view correspondence through view correspondences with S_2.

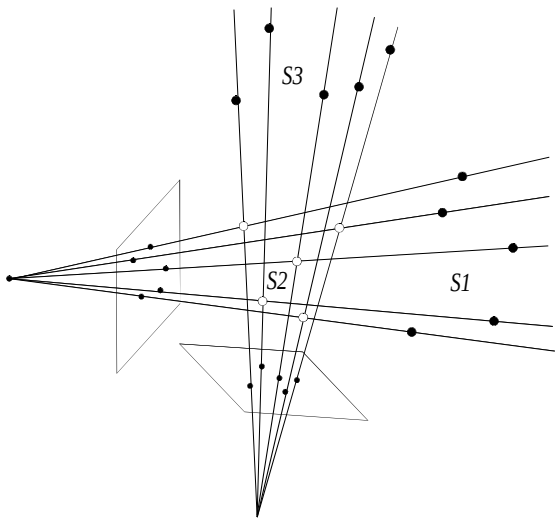

Figure 1: View correspondence. Point sets S_1 and S_2, and S_2 and S_3, are in view correspondence; point sets S_1 and S_3 are in transitive view correspondence through S_2.

View correspondence is a generalization of the concept of *perspective correspondence* found in traditional projective geometry [Dud73]. In the traditional concept, the point sets are constrained to planes.

2.2 Overview of Proof

In order to guide the reader, let us give an overview of the *proof by contradiction* that there are no general-case invariant functions. Suppose that such a *non-trivial* invariant function f existed. Then, for some n, there would be two sets S and S' of n three-dimensional points such that

$$f(\pi_{R,\vec{T}}(S)) \neq f(\pi_{R',\vec{T}'}(S'))$$

for a pair of views $(R, \vec{T})$ and $(R', \vec{T}')$ non-degenerate for each point set respectively.

The proof involves constructing a chain of intermediate sets starting with $S_1 = S$ and terminating in $S_n = S'$, such that any two consecutive sets must have the same value for the invariant function. This contradicts the assumption, proving that the only general-case invariants are constant functions.

Lemma 1 proves that point sets in transitive view correspondence maintain the same value of an invariant function and is a direct consequence of the constancy under view correspondence and change of viewpoint. It is fairly obvious that 3D point sets in view correspondence will project to the same set of 2D points, and therefore the value of any function of the projection must be the same. Since we are assuming an invariant function, the value will not change as a function of viewpoint. Thus, any chain of view correspondences with a change of viewpoint between pairs will leave the value of an invariant unchanged.

Figure 1 is an example of point sets S_1 and S_3 in transitive view correspondence, but note that in this example the two sets are not in general 3D position; otherwise, the corresponding rays would not cross. Lemmas 2 and 3 establish that two sets of the same size S_1 and S_n are in transitive view correspondence in the general case. Lemma 2 shows that pairs of sets identical up to the position of one point are in transitive view correspondence, and in Lemma 3, a sequence of sets $S_2, \cdots S_{n-1}$ are constructed where each adjacent pair of sets $(S_i, S_{i+1}), 1 \leq i < n$, are identical up to the position of one point, thus establishing the correspondence.

The proof of the theorem is now straightforward because the sets $S = S_1$ and $S' = S_n$ of size n must be in transitive view correspondence, and therefore have the same value for an invariant. This contradicts the original assumption that $f(\pi_{R,T}(S)) \neq f(\pi_{R',T'}(S'))$. Clearly views of relevant point sets from which the feature is undefined cannot be used to establish this contradiction; therefore, much of the proof detail deals with avoiding such views.

2.3 Theorem and proof

Lemma 1 *A general-case view-invariant feature has the same value for point sets in transitive view correspondence if this correspondence is established through views that are not degenerate for the feature and point sets involved.*

Proof. For point sets in view correspondence S_1 and S_2, this is immediate. Being a general-case invariant, f has a single value f_1 for all projections of S_1, and similarly f_2 for S_2. Since $\pi_{R_1,\vec{T}_1}(S_1) = \pi_{R_2,\vec{T}_2}(S_2)$, for some pair of views $(R_1, \vec{T}_1, R_2, \vec{T}_2)$ non-degenerate for their respective point sets, we have $f_1 = f(\pi_{R_1,\vec{T}_1}(S_1)) = f(\pi_{R_2,\vec{T}_2}(S_2)) = f_2$.

By definition, for point sets in transitive view correspondence S_1 and S_n, there exists a sequence of point sets $S_2, \ldots, S_{n-1}$ such that every pair (S_i, S_{i+1}) is in view correspondence, for $1 \leq i < n$. This means that there exist views $(R_i, \vec{T}_i, R_{i+1}, \vec{T}_{i+1})$ such that $\pi_{R_i,T_i}(S_i) = \pi_{R_{i+1},T_{i+1}}(S_{i+1})$. Given that these correspondences are established through non-degenerate views of the point sets involved, by transitivity of equality, the feature values associated with S_1 and S_n must be equal. □

Lemma 2 *If two 3D point sets of the same size are identical for all but one point, then they are in transitive view correspondence through views that are not degenerate for a given feature and the point sets involved, assuming the two point sets are themselves not degenerate for this feature.*

Proof. Consider two such sets S_1 and S_2, and say that they differ in the position of the jth point; i.e., $\vec{P}_{1,j} \neq \vec{P}_{2,j}$ (Figure 2a). In order to show *view* correspondence between the sets, it is sufficient to show that there exists a pencil of rays that contains both sets. This can be done by placing the focal point (point of intersection for the rays) in the line that passes through $\vec{P}_{1,j}$ and $\vec{P}_{2,j}$. Clearly these two points will then have a ray in common and the other point pairs will share rays by virtue of the fact that they occupy the same positions in space.

Unfortunately, this view correspondence is established using a focal point selected from a restricted set, the line through $\vec{P}_{1,j}$ and $\vec{P}_{2,j}$, and it is possible that the views associated with this set are all degenerate for the points and a given feature. For example, in Figure 2(b) this line contains another point $\vec{P}_{1,i} = \vec{P}_{2,i}$, and thus the projection of these points will be coincident for all views associated with the line, a degenerate condition for many features (Section 4 and 5).

Nevertheless, a *transitive* view correspondence can still be established through views non-degenerate for a given feature. This can be done by considering a third, intermediate set of points S_3 which is identical to both S_1 and S_2 except for point $\vec{P}_{3,j}$ (Figure 2c). For a given feature, a point $\vec{P}_{3,j}$ can always be found so that S_3 is in view correspondence with both S_1 and S_2 via non-degenerate views of S_1, S_2 and S_3. From the construction (Figure 2c), two focal points $\vec{F}_1$ and $\vec{F}_2$ can be used to establish the non-degenerate view correspondences of (S_1, S_3) and (S_3, S_2), respectively, if $\vec{P}_{3,j}$ is taken to be the intersection of the lines through $(\vec{F}_1, \vec{P}_{1,j})$ and $(\vec{F}_2, \vec{P}_{2,j})$, and the

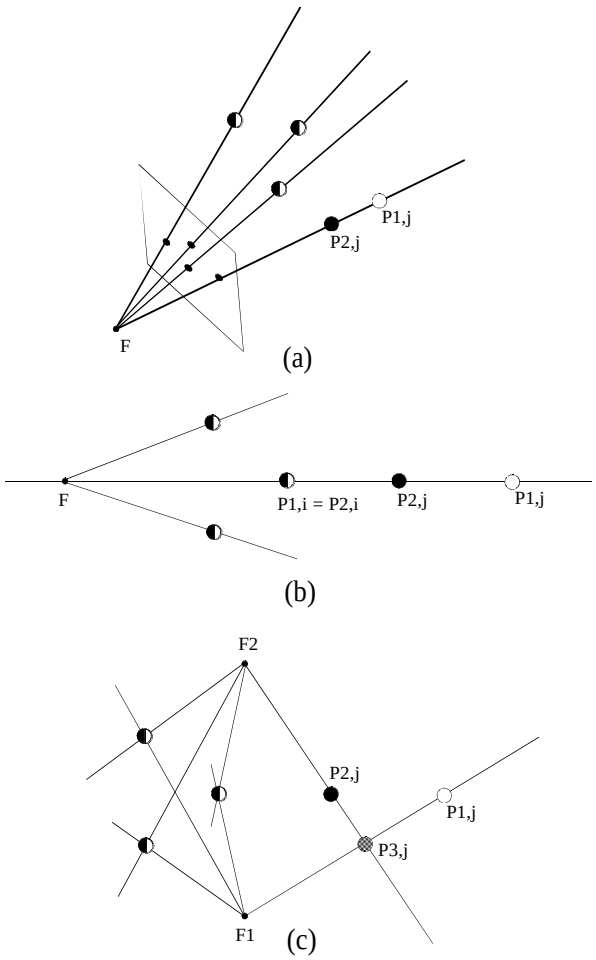

Figure 2: 3D point sets S_1 (white) and S_2 (black) are identical up to $\vec{P}_{1,j}$ and $\vec{P}_{2,j}$. (a) To show that such sets are in view correspondence select a focal point in the line through $\vec{P}_{1,j}$ and $\vec{P}_{2,j}$; in this way each pair of corresponding points, one from each set, share a ray. This correspondence cannot always be established using non-degenerate views, for example in (b), the projections of $\vec{P}_{1,j}$ and $\vec{P}_{1,i}$ will be coincident. (c) *Transitive* view correspondence through views non-degenerate for a given feature can be established using a set S_3, which is identical to S_1 and S_2 for all but $\vec{P}_{3,j}$.

pair of focal points $\vec{F}_1$ and $\vec{F}_2$ satisfy the following three properties: (1) $\vec{F}_1$, $\vec{F}_2$, $\vec{P}_{1,j}$ and $\vec{P}_{2,j}$ are coplanar, (2) the lines through $(\vec{F}_1, \vec{P}_{1,j})$ and $(\vec{F}_2, \vec{P}_{2,j})$ are not parallel, and (3) $\vec{F}_1$ and $\vec{F}_2$ produce non-degenerate views for S_1 and S_2 respectively [4]. The first two properties are sufficient conditions for the existence of the intersection point $\vec{P}_{3,j}$.

Given that there is only a measure zero set of degenerate views for any (non-degenerate) 3D point set and feature, we show that it is always possible to find a pair of points $\vec{F}_1$ and $\vec{F}_2$ that satisfy these conditions. Consider the pencil of planes that contain $\vec{P}_{2,j}$ and $\vec{P}_{1,j}$. At most a measure zero set of these planes can contain more than a measure zero set of degenerate $\vec{F}_1$, otherwise the total set of degenerate $\vec{F}_1$ in R^3 would not be of measure zero. Since the same situation holds for $\vec{F}_2$, and since the intersection of two measure zero sets is itself measure zero, at least one plane in the pencil contains at most a measure zero set of degenerate $\vec{F}_1$ and at most a measure zero set of degenerate $\vec{F}_2$. Consider this plane and a non-degenerate point $\vec{F}_1$ in it. Since there are non-degenerate $\vec{F}_2$ almost everywhere in the plane, there must be at least one not in the line through $\vec{P}_{2,j}$ that is parallel to the line $(\vec{F}_1, \vec{P}_{1,j})$. Hence, there is always a pair of points $(\vec{F}_1, \vec{F}_2)$ that satisfy the above three conditions, establishing the non-degenerate transitive view correspondence. □

Lemma 3 *Every pair of 3D point sets of size n are in transitive view correspondence through views that are not degenerate for the point sets involved and a given feature, assuming the two point sets are themselves not degenerate for this feature.*

Proof. Consider two sets of size n, S_1 and S_n. For these two sets, construct a sequence of sets $(S_2, \ldots, S_{n-1})$ such that adjacent pairs S_j and S_{j+1} are identical, except for their jth points, for $1 \leq j < n$. This sequence can be constructed by defining each S_j to be the union of the subset of $j-1$ points $\{\vec{P}_{1,i} | 1 \leq i < j\}$ of S_n and the subset of $n-j+1$ points $\{\vec{P}_{n,i} | j \leq i \leq n\}$ of S_1.

Every adjacent pair of point sets in this sequence differs by a single point, and therefore by Lemma2 is in transitive view correspondence through views

[4] Given this property, the focal points also produce non-degenerate views for S_3, since this point set is in view correspondence with S_1 and S_2 via these focal points.

non-degenerate for any given feature, provided that these point sets are not themselves degenerate. By transitivity, the original point sets S_1 and S_n are also in transitive view correspondence.

It is possible for some of the intermediate sets S_j to be degenerate for the given feature, even though S_1 and S_n are not; e.g. if S_1 and S_n had points in common but were labeled differently. Fortunately, in such cases, it is always possible to change both of the sets using view correspondence so that the degeneracy no longer occurs (i.e. by moving a point along a ray). This is always possible since there can only be a measure-zero set of degenerate cases for the feature and any given set S_1 is always in transitive view correspondence, through non-degenerate views, with a non-measure-zero set of S's [5] $\square$

Theorem 1: There is no feature of n projected points, for any n, that is both a general-case view-invariant and non-trivial in the senses defined above.

Proof. Consider a general-case view-invariant f. For f to be non-trivial, there must exist at least one pair of point sets S_1 and S_2, such that $f(\pi_{R_1,\vec{T}_1}(S_1)) \neq f(\pi_{R_2,\vec{T}_2}(S_2))$ for some $(R_1, \vec{T}_1, R_2, \vec{T}_2)$.

Since all non-degenerate 3D point sets of the same size are in transitive view correspondence through non-degenerate views (Lemma 3) and every pair in transitive view correspondence through non-degenerate views must have the same value for any general-case view-invariant (Lemma 1), the necessary condition for a non-trivial feature cannot be established for any such general-case view-invariant. $\square$

2.4 Extension of result to other projection models

Theorem 1 also holds for two other commonly used projection models: orthographic and weak perspective. Orthographic projection is the same as perspective, except that the projection rays must always intersect at a point at infinity. Only Lemma 2 and Lemma 3 consider the projection geometry. In Lemma 2, since the construction step allows the ray intersection (focal) points to lie anywhere in given lines, they can always be placed at infin-

[5] This can be pictured by noting that, for a given ray bundle, each point in S_1 can move independently along its ray to produce a family of new sets S'. By repeatedly moving the points in this way, using non-degenerate and non-collinear ray intersection points, the set of S's that can be reached has positive measure.

ity. In addition, the existence of appropriate *non-degenerate* views can be demonstrated along the same lines as in the proof of Lemma 2. In Lemma 3, degenerate intermediate point sets can be avoided by suitable perturbation of the initial point sets in the same way as in the perspective case.

Weak perspective is identical to orthographic except that the projection is scaled by the inverse distance from the camera to the centroid of the 3D point set. Given this, from the viewpoint for which two point sets are in orthographic view correspondence, their weak perspective projections will always be identical with the original orthographic projection up to some scale factor (possibly different for each set). By suitable translation of viewpoints, each of the two sets can always produce the original orthographic projection exactly, and if the orthographic view is non-degenerate for a given feature, the aligned weak-perspective views must also be. By defining *weak perspective* view correspondence to be this alignment of projections, the extension of the theorem to orthography must also be possible for weak perspective.

3 Special-case view-invariants under weak perspective

The *weak perspective* projection model is an approximation to perspective projection that is applied extensively in object recognition research; it can simplify the analysis and computation for 3D object recognition, with reasonable results when the camera is far enough from the object relative to the depth variation of the object [Bro81,Tho87,Hut87,Lam88a]. Weak perspective is the same as the perspective projection defined in the last section, except that all of the points on a 3D object are treated as being at the same distance from the camera:

$$\pi_{R,\vec{T}}(\vec{P}) = (F/z_0) \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} (R\vec{P} + \vec{T})$$

for average distance z_0. Given this, the object-to-image transformation becomes much simpler: there is a single scale factor F/z_0 instead of a different variable z_i for each point, each with a non-linear effect. Another reason for working with a weak perspective model is that there are some useful special-case view-invariants that are approximately invariant for true perspective

at an appropriately large camera distance. For example, parallelism can be effectively used as a line grouping criterion in grouping-based recognition [Low85].

Given the weak perspective projection model, the object-to-image transformation is a singular affine transformation that can be represented by a 2×3 rank 2 matrix and 2D translation vector [Lam88b]. By considering only point position *differences*, the translation component can be subtracted out and the vectors representing the 3D point position differences can be related to their projections through a linear transformation. This is important since, for non-degenerate views, the property of linear dependence for a set of vectors, the coefficients expressing the linear dependence, and the subspace dimension associated with the dependence are all invariant to linear transformation[6].

All of the weak perspective invariants discussed in the recognition literature [Bro81, Low85, Lam88b] follow from this observation, and other, related special-case invariants can be deduced that may also be useful for recognition algorithms. For each type of special-case invariant, the required 3D point-set conditions can be specified as the linear dependence of some set of their position differences, and the invariant feature itself can be defined as a coefficient or subspace dimension associated with this linear dependence.

The different special-case conditions can be distinguished by two properties: the dimension of the subspace containing the set of 3D difference vectors, and the pattern of the point differences. The subspace dimensions given 3D point sets are two, one, and zero; the point difference patterns specify which pairs of points are being subtracted. Two types of difference patterns can be found in the invariants surveyed: differences from a single reference point $\vec{p}_1$ selected from the set, or $\{(\vec{p}_i - \vec{p}_1) | i \geq 2\}$, and differences between points that have been paired off, $\{(\vec{p}_{2i} - \vec{p}_{2i-1}) | i \geq 1\}$.

Table 1 shows the classification of special-case weak-perspective invariants based on the above two properties. It is interesting to note that special-case invariants that have not been found in the recognition literature can be deduced from this framework; see for example table entry (3), the invariant classified under (*dimension two, paired-off point differences*). It is also interesting to note that a 3D point set need not be planar for this invariant

[6]This follows from the fact that the null vector $\vec{0}$ is a fixed point for all linear transformations.

Subspace dimension	Pattern of 3D point differences	
	Differences from single reference point $\{(\vec{P}_i - \vec{P}_1) \mid i \geq 2\}$	Paired-off point differences $\{(\vec{P}_{2i} - \vec{P}_{2i-1}) \mid i \geq 1\}$
two	**Special-case constraint:** Four 3D points such that $\sum_{i=2}^{4} c_i(\vec{P}_i - \vec{P}_1) = \vec{0}$ for non-zero c_i. (Four coplanar points.) **View invariants:** (1) $(c_3/c_2, c_4/c_2)$. Example of use: the *affine coordinate* features in [Lam88a] (2) Subspace dimension of dependence is 2, which is not significant for 2D images.	**Special-case constraint:** Six 3D points such that $\sum_{i=1}^{3} c_i(\vec{P}_{2i} - \vec{P}_{2i-1}) = \vec{0}$ for non-zero c_i. (Note: the *points* need not be coplanar.) **View invariants:** (3) $(c_2/c_1, c_3/c_1)$. No known example of use in recognition literature. (4) Subspace dimension of dependence is 2, which not significant for 2D images.
one	**Special-case constraint:** Three 3D points such that $\sum_{i=2}^{3} c_i(\vec{P}_i - \vec{P}_1) = \vec{0}$ for non-zero c_i. (Three collinear points.) **View invariants:** (5) c_3/c_2. Ratio of distances along line, example of use: the *approach ratio* in [Bro81]. (6) Subspace dimension of dependence is 1, *collinearity*.	**Special-case constraint:** Four 3D points such that $\sum_{i=1}^{2} c_i(\vec{P}_{2i} - \vec{P}_{2i-1}) = \vec{0}$ for non-zero c_i. (Endpoints of two parallel line segments.) **View invariants:** (7) c_2/c_1. Ratio of distances measured in parallel directions. (8) Subspace dimension of dependence is 1, *parallelism*. Example of use: grouping-based recognition in [Low85].
zero	**Special-case constraint:** Two 3D points such that $c_2(\vec{P}_2 - \vec{P}_1) = \vec{0}$ for non-zero c_2. (Two coincident points.) **View invariants:** (9) c_2. Uninteresting (can always be anything but zero). (10) Subspace dimension of dependence is 0, *coincidence*. Example of use: coincidence of line segment endpoints for grouping-based recognition in [Low85].	Same case as left panel.

Table 1: Classification of special-case view-invariants under weak perspective. All invariants discussed in the literature are functions of point position differences, and they can be distinguished by two properties: the dimension of the subspace containing the set of 3D difference vectors (row), and the pattern of the point differences (column).

to be valid. Many of the special-case invariants represented in the table can be equivalently expressed as special cases for one of the four *line-segment* features analyzed in the next section; this will be made clear in that section.

Since the weak perspective object-to-image transformation is a strict subset of the singular affine transformations just discussed, it is possible that there are additional invariants, valid for the weak perspective case only. This is an open issue that would be useful to resolve.

There are also special-case invariants under perspective projection [Dud73, Lam88b] analogous to those under the column *Differences from single reference point* in Table 1. The perspective invariant called the *cross-ratio* is analogous to the ratio of coefficients (table entry 5), but requires four points instead of three, and *projective coordinates* correspond to the affine coordinates (table entry 1), but require five points instead of four. There are also projective invariants of planar curves that can be used in recognition [For90]. In addition, important projective invariants may be associated with special-case configurations of points other than planar; for example, Forsyth *et al* are currently investigating the cases of cubics and surfaces of revolution [For91]. There do not seem to be special-case invariants under perspective projection analogous to the ones in the second column (table entries 3, 7, 8).

4 Summary and conclusions

One of the outstanding problems in visual object recognition is the development of systems capable of recognizing 3D objects from unknown view by matching them to a single 2D image. A fundamental issue involves reliable recognition in spite of the variation in the 2D projection as a function of viewpoint. Thus, an important part of our study of 3D object recognition in 2D images has been an analysis of the variation with respect to view of features that are functions of projected sets of points [Bur91].

This paper establishes that general-case view invariants do not exist for any number of points, given true perspective, weak perspective or orthographic projection models. The absence of view-invariant features has led some researchers to use special-case invariants, features that are only view-invariant for special configurations of 3D points or sets of views. In this paper, the special-case weak-perspective invariants cited in the literature are derived from linear dependence relations and the invariance of this type of

relation to linear transformation. However, as is also shown in this paper, invariants that do not require any restrictions are impossible, and recognition systems for the unrestricted problem cannot be expected to be as effective as the special-case schemes. Since view invariants cannot be guaranteed to exist or be useful for any given recognition problem (i.e., collection of 3D objects), it is imperative that systems be developed that are capable of utilizing view-varying features as well. For example, features such as relative orientation, size and length of projected line segments in arbitrary positions can be useful for the discrimination of many objects, from most of their views [Ben90, Burns91]. Because features crucial for recognition may vary with view, their distributions may be non-trivial. Systems developed to use these features must effectively model the distributions and apply probabilistic means of generating and evaluating matches [Bin87, Ben90, Burns91].

ACKNOWLEDGMENTS: The authors wish to thank Harpreet Sawhney, Mark Snyder and Poornima Balasubramanyam for useful discussions and comments.

References

[1] Ahuja, D.V. and S.A. Coons, "Geometry for Construction and Display", IBM Systems Journal, 7(3-4), pp. 188-205, 1968.

[2] Aliomonos, Y. and M. Swain, "Paraperspective Projection: Between Orthography and Perspective", CAR-TR-320, U. Maryland, 1987.

[3] Ben-Aire J., "The Probabilistic Peaking Effect of Viewed Angles and Distances with Application to 3D Object Recognition", IEEE PAMI-12, no. 8, pp. 760-774, 1990.

[4] Binford, T.O., T.S. Levitt, W.B. Mann, "Baysian Inference in Model-Based Machine Vision", Proc. AAAI Uncert. Work., 1987.

[5] Burns, J.B., "View-variation of point-set and line-segment features", *DARPA Image Understaning Workshop*, 1990.

[6] Burns, J.B., "Matching 2D Images to Multiple 3D Objects Using View Description Networks", PhD. Thesis, UMass Amherst, 1991.

[7] Clemens, D.T. and D.W. Jacobs, "Model group indexing for recognition", *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition*, June 1991.

[8] Duda, R. and P. Hart, *Pattern Classification and Scene Analysis*, Wiley-Interscience, 1973.

[9] Forsyth, D.A., J.L. Mundy, A. Zisserman, and C.M. Brown, "Projectively Invariant Representations Using Implicit Algebraic Curves", *Proc. First European Conf. Comput. Vision*, Springer LNCS, 1990.

[10] Forsyth, D.A., J.L. Mundy, A. Zisserman, and C.M. Brown, work in preparation, 1991.

[11] Lamdan, Y., J.T. Schwartz and H.J. Wolfson, "Object Recognition by Affine Invariant Matching", Proc. CVPR, pp. 335-344, 1988.

[12] Lamdan, Y., H.J. Wolfson, "Geometric Hashing: A General and Efficient Model-Based Recognition Scheme", Proc. ICCV, pp. 238-249, 1988.

[13] Lowe, D.G., *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, 1985.

[14] Mohr, R., L. Morin, C. Inglebert and L. Quan, "Geometric Solutions to 3D Vision Problems", in preparation, 1990.

[15] Tucker, L.W., C.R. Feynman, D.M. Fritzsche, "Object Recognition using the Connection Machine", Proc. CVPR, 1988.

[16] Weiss, I., "Projective Invariants of Shapes", DARPA Image Understanding Workshop, pp. 1125-1134, 1988.