On Invariants of Sets of Points or Line Segments Under Projection *

Daniel P. Huttenlocher and Jon M. Kleinberg
Computer Science Department
Cornell University
Ithaca NY 14853

Abstract

We consider the problem of computing invariant functions of the image of a set of points or line segments in $\Re^3$ under projection. Such functions are in principle useful for machine vision systems, because they allow different images of a given geometric object to be described by an invariant 'key'. We show that if a geometric object consists of an arbitrary set of points or line segments in $\Re^3$, and the object can undergo a general rotation, then there are no invariants of its image under projection. For certain constrained rotations, however, there are invariants (e.g., rotation about the viewing direction). Thus we precisely delimit the conditions for the existence or nonexistence of invariants of arbitrary sets of points or line segments under projection.

*This work was supported in part by National Science Foundation PYI grant IRI-9057928 and matching funds from General Electric, Kodak and Xerox, and by Air Force contract AFOSR-91-0328.

1. Introduction

For many problems in computer vision, it is desirable to be able to describe an object independent of the direction from which it is viewed. In other words, it is useful to be able to form a 'key' corresponding to a given object, where that key remains unchanged regardless of the relative position and orientation of the object and the viewer. In order to compute such viewpoint-independent descriptions, a number of recent investigations have applied techniques from classical invariant theory (e.g., [3, 6, 7, 8]). Invariant theory is a branch of algebra concerned with functions for which the value of the function remains unchanged after it is composed with members of some distinguished set of functions $\mathcal{H}$ (cf. [4]). For example, a function f on the elements of some set S is invariant with respect to $\mathcal{H}$, if $f(x) = f(h(x))$ for all $x \in S$, $h \in \mathcal{H}$. A function f is said to be trivial if it has the same value for all $x \in S$ — clearly such constant functions are invariant but they are uninteresting because they cannot distinguish different elements of S.

If the set of functions $\mathcal{H}$ represents the change in the shape of an object as the viewing direction varies, then any nontrivial invariants with respect to $\mathcal{H}$ provide viewpoint-independent 'keys' for describing an object. Such invariants do exist in certain cases, and have been successfully applied to a number of model-based recognition problems (in which geometric models of known objects are compared against an image containing unknown objects) [3, 7]. The classical theory, however, studies invariants with respect to sets of functions that form a group. In particular, this means that each function in the set $\mathcal{H}$ must have an inverse also in the set. For many machine vision applications this is an unfortunate constraint, because the projection of the three-dimensional world onto a two-dimensional image plane is not invertible. Thus applications of invariant theory to computer vision have been primarily restricted to the case where the object being viewed is flat (planar), so that the viewing transformation forms a group. One notable exception is the work of [7] who develops invariants of orthographic projections of three orthogonal vectors in space.

In this paper we investigate the problem of computing invariants of the two-dimensional images of geometric objects in three-space. We focus on the case of finite point sets in $\Re^3$, and then show how to extend the results for point sets to sets of line segments in $\Re^3$. We consider the image of a set of points or line segments in $\Re^3$ under some projection operation (e.g., orthographic or central projection), where the set may be transformed by the action of some group G acting on $\Re^3$. The invariant functions, if they exist, operate on the projections of the point set. Thus these functions have no direct access to the action of the group G, but only to the images of those actions under the projection operation (in other words the set $\mathcal{H}$ here is the composition of G with the projection operation). We delineate the classes of groups for which there exist (or do not exist) invariants of sets of points or line segments under projection. In order to do this, we define a more general notion of projection that encompasses both central and orthographic projection (among others).

A number of papers in the computer vision literature have recently considered the question of the existence of nontrivial invariants of three-dimensional objects under projection into the plane. One of these works [1] concludes after somewhat lengthy arguments that there are no (nontrivial) invariants of finite point sets in $\Re^3$ under Euclidean motion and central projection. Another paper [2] makes an informal argument that there cannot be invariants under orthographic projection from $\Re^3$ into $\Re^2$. A third paper makes a more rigorous argument that there cannot be invariants under orthographic projection from $\Re^3$ into $\Re^2$, and investigates certain restricted object classes where there are invariants (e.g., symmetric sets) [5]. None of these investigations precisely characterizes the conditions for the existence of invariants for point sets under projection, nor do they consider sets of line segments.

Our basic result is that if the group G acting on a three-dimensional set (points or line segments) contains a rotation about some fixed but arbitrary axis, then there are no invariants under projection. For certain rotation axes (such as the viewing direction under orthographic projection) there are invariants. We thus provide a precise formulation of which groups have (or do not have) nontrivial invariants under projection.

2. Projections and Invariants

There are at least two standard projections we could consider: orthographic (where lines of projection are perpendicular to the image plane) and central (where lines of projection all pass through a single point). Rather than prove all of our results twice, we use a more general definition of projection, which we introduce here. The set of lines of projection (or lines of sight) is the key structure underlying a projection. In general, for each point $p \in \Re^3$ there should be some unique corresponding line of projection on which it lies. Clearly in orthographic projection this is just the line through p and normal to the viewing plane. For central projection it is the line through p and the center of projection. Note that in this latter case the center of projection itself is a *degenerate point* – it does not have a unique projection line, as all the lines go through it.

We define the set of lines of projection, Λ, to be a set of lines that cover all of $\Re^3$, such that each line $\ell \in \Lambda$ has at most one degenerate point. We add an additional constraint that the set of all degenerate points be small (see condition 3 below). Denote this set of degenerate points by ϕ_Λ, and let $D_\Lambda = \Re^3 - \phi_\Lambda$. A given set of lines of projection, Λ, thus defines a naturally associated map $\lambda : D_\Lambda \Rightarrow \Lambda$. That is, each non-degenerate point of $\Re^3$ (the points of D_Λ) lies on a unique line of Λ, and the function λ maps each such point to its corresponding line. Note that this definition of a projection does not map points of $\Re^3$ to points of $\Re^2$; rather it maps points of $\Re^3$ to lines of Λ. These lines could then be cut by some (fixed) viewing plane, or manifold, yielding a unique point of $\Re^2$ for each line of Λ.

A little more precisely, a projection λ is the natural function from D_Λ to Λ associated with a given set of projection lines Λ, where Λ has the following three properties:

1. $\bigcup_{\ell \in \Lambda} \ell = \Re^3$
2. With $\phi_\Lambda = \{\ell \cap \ell' : \ell, \ell' \in \Lambda\}$, each set $\ell \cap \phi_\Lambda$ ($\ell \in \Lambda$) is empty or contains a single point.
3. ϕ_Λ contains no ray or circle. (Among other things, this guarantees that it is a "one-dimensional" subset of $\Re^3$.)

We wish to note that our results still hold for a projection in which the degenerate set is a line (e.g., as would occur with a prism); the proofs, however, become considerably more involved.

In the following, we consider the case of invariants of point sets under projection, and defer discussion of sets of line segments (which we call frameworks) until Section 6. For a finite point set $P \subset D_\Lambda$, we define $L = \lambda P$ in the natural way as the set of lines $\{\lambda p : p \in P\}$. While the map λ is not invertible, it is often desirable to speak of a particular set of points that could project to $L = \lambda P$. We thus define a *realization* of $L \subset \Lambda$ to be any set of points that could project to L; i.e. a point set $P \subset D_\Lambda$ such that $\lambda P = L$. In other words, a realization of $L \subset \Lambda$ is a point set that results from simply picking one non-degenerate point on each of the lines $\ell \in L$.

Next, we would like to define more precisely what is meant by an invariant function under projection. These functions will operate over finite sets of lines of Λ (viewing lines) of some fixed cardinality. We fix the cardinality because it is a nontrivial invariant, though not a very interesting (nor geometric) one. Moreover, so as to be able to prove the strongest non-existence results we can, we will assume that our invariant functions are only defined on sets of n distinct lines of Λ — this means that if one point of P occludes another in the projection, f is simply not defined. If this assumption is not made, the proofs below become much easier, but carry less weight as they rely on bizarre degeneracies based on the viewing direction. We use the notation Λ_n to denote the set of all subsets $L \subset \Lambda$ of cardinality n. Let $\mathcal{F}_\lambda$ be the set of all functions on sets $L \in \Lambda_n$ for some fixed n. We take all elements of $\mathcal{F}_\lambda$ to have some common range set, but the only assumption we make about this set is that it is at least as large as the reals. Let G be a transformation group that acts on points in $\Re^3$, and let λ be a projection induced by the covering set Λ (as defined above).

We will say that $f \in \mathcal{F}_\lambda$ is invariant with respect to G and λ if $f(\lambda P) = f(\lambda g P)$ for any set of points P such that $\lambda P, \lambda g P \in \Lambda_n$. (Note that $P \in \Lambda_n$ exactly when each point of P lies on some unique line of Λ, i.e., $P \subset D_\lambda$ and $\forall p, p' \in P : \lambda p = \lambda p' \rightarrow p = p'$). Let $I_\lambda(G) \subset \mathcal{F}_\lambda$ denote the set of all functions in $\mathcal{F}_\lambda$ that are invariant with respect to G and λ. Clearly the constant functions are in $I_\lambda(G)$ for all G, λ; these are said to be the *trivial* invariants, and if $I_\lambda(G)$ contains only constant functions, we will say that $I_\lambda(G)$ is trivial.

In other words, $I_\lambda(G) \subset \mathcal{F}_\lambda$ is the set of functions on Λ_n that remain unchanged under the action of G on a point set $P \subset D_\Lambda$ (where λP and $\lambda g P$ are both sets of n distinct lines of Λ). Thus the structure is that of a group acting on a space (D_Λ) that is not directly 'observable' by the functions in $\mathcal{F}_\lambda$ (all that the functions see is Λ_n). Classical invariant theory was also developed in part to study the images of geometric objects under projection (e.g., the planar affine and projective groups). In the classic case, however, the composition of G and the projection operation λ itself forms a group. Thus the object of study is still a group action. In the general case that we consider here the composition of G and λ is no longer invertible.

Let us make note of a subtlety in all this. Our invariant functions are based on *a priori* assumptions about the structure of our sets in $\Re^3$ and Λ. That is, we have the following conditions on P and λP:

- The three-dimensional condition: P is a set of n points, none of which lie in ϕ_Λ.
- The two-dimensional condition: No two points of P are mapped to the same line of Λ by λ.

As noted above, these assumptions make the following non-existence proofs more, not less, general. In what follows, we shall be careful the realizations and projections we construct always satisfy this pair of conditions.

There is a powerful inductive technique that we use in our proofs, which relies on a property of set functions that we will call **Property A**. Consider two sets S and S' of n elements each that differ in at most one element (i.e., $S \cap S'$ has cardinality at least $n-1$). We say that f has Property A when $f(S) = f(S')$ for any such S, S'.

The following result now holds for functions on sets of cardinality n.

Lemma 1 *If f has Property A, then f is constant.*

Proof. Let $T = \{t_1, \ldots, t_n\}$, $T' = \{t'_1, \ldots, t'_n\}$. Set $T_0 = T$, $T_n = T'$, and for $i = 1, \ldots, n-1$, define $T_i = \{t_1, \ldots, t_i, t_{i+1}', \ldots, t'_n\}$. Then for $i = 0, \ldots, n-1$, $f(T_i) = f(T_{i+1})$, by Property A, so $f(T) = f(T')$. Since T and T' were arbitrary, f is constant. ∎

Using Property A, we will devise a characterization of groups G that have only trivial invariants under the projection function λ. We wish to note that there are analogous ways of phrasing Property A and Lemma 1 that correspond to labeled point sets. Our non-existence proofs then carry over to the case of labeled point sets with almost no modification.

Before developing the last few tools needed to show some general results on existence and non-existence of invariants, we would like to demonstrate the simplicity of our main technique by proving one result immediately.

Theorem 2 *Let E_3 denote the group of rigid (Euclidean) motions in $\Re^3$ (i.e., the composition of the special orthogonal and translation groups). For any projection λ, $I_\lambda(E_3)$ is trivial.*

Proof. Let f be invariant with respect to λ and E_3, and let L and L' be subsets of Λ of size n that differ in at most one element. Let P and P' be *realizations* of L and L' such that $p_j = p'_j$ for all $j \neq i$. (Recall P is a realization of L when $\lambda P = L$.) Let η be the line through p_i and p'_i. For any non-degenerate line $\ell \in \Lambda$, there is a rigid motion g such that $g\eta = \ell$ and neither p_i nor p'_i lands on $\ell \cap \phi_\Lambda$. Moreover, we can choose P and P' such that no two points of gP or gP' lie on the same line of Λ or in the degenerate set. Then $\lambda gP = \lambda gP'$, so $f(L) = f(\lambda gP) = f(\lambda gP') = f(L')$; thus f has Property A and is constant. ∎

The key step of this proof is showing that given two arbitrary points in $\Re^3$, we can apply an element of E_3 that positions them on the same line of projection; in this way, two point sets differing by a single element cannot be distinguished by a function invariant with respect to $G = E_3$ and λ. Then the result follows immediately from Lemma 1. This result is a generalization of the claims in [1] and [2], on the nonexistence of invariants under central and orthographic projection. Our proof of this fact is very simple and does not rely on any degeneracies, such as multiple collinear or coplanar points. Moreover, we apply the same general technique to delineate precisely the conditions for the existence or nonexistence of invariants under projection.

We also note that

Fact 3 *If H is a subgroup of G, then $I_\lambda(H) \supseteq I_\lambda(G)$.*

Thus there are no invariants under projection for any group that contains E_3 as a subgroup.

The proof of Theorem 2 can be generalized to show that other (smaller) groups do not have nontrivial invariants under projection. By Fact 3, these groups in turn rule out invariants for any groups containing them. We now turn to developing some additional tools needed to establish these results. The main difference in the technique is that for more restricted groups we are no longer able to use the group action to place two points on some line of Λ. Instead we form a finite sequence of group actions and motions along a line of Λ that places two points on the same line of Λ. (Note that motions along a line of Λ are not detectable by functions in $\mathcal{F}_\lambda$.) In the end we derive a precise characterization of those groups with trivial invariants under projection.

3. Groups With(out) Invariants Under Projection

In the proof of Theorem 2 we used the fact that for any two points $p, q \in \Re^3$, there exists a Euclidean motion $g \in E_3$ that positions p and q on any line of $\Re^3$. Thus it was possible to establish Property A. We now consider groups for which this is no longer the case (the only assumption that we make about the group G is that it maps lines to lines). Consider the rotation group (i.e., the special orthogonal group SO_3); clearly it is not possible to rotate the line containing two arbitrary points p and q in order to make it coincident with any line of $\Re^3$. Thus we make use of the additional fact that motion of a point along a line $\ell \in \Lambda$ is not detectable (all points along a given viewing line are indistinguishable). We then build a finite sequence of group actions and motions along ℓ that brings p and q onto a given viewing line.

We express this notion as follows. Let $\ell^0, \ell^1 \in \Lambda$ and write $\ell^0 \sim \ell^1$ if there exist $g_1, \ldots, g_k \in G$, $\ell_1^0, \ldots, \ell_k^0, \ell_1^1, \ldots, \ell_k^1 \in \Lambda$ such that with $\ell^0 = \ell_0^0$, $\ell^1 = \ell_0^1$ we have

$$g_i(\ell_{i-1}^j - \phi_\Lambda) \cap (\ell_i^j - \phi_\Lambda) \neq \phi$$

$$(i = 1, \ldots, k, j = 0, 1)$$

$$\ell_k^0 = \ell_k^1 \tag{1}$$

We will sometimes refer to these sequences as a chain from ℓ^0 to ℓ^1; the intermediate lines ℓ_i^j are the two "sides" of the chain. If $\ell^0 \sim \ell^1$ for every pair of lines of Λ, then we say that (G, λ) has **Property B**; that is, there is a chain of the form in (1) for every pair of lines of Λ. We first prove that $\sim$ is in fact an equivalence relation

Lemma 4 *$\sim$ is an equivalence relation on the elements of Λ.*

Proof. Reflexivity and symmetry are easy to show. To establish transitivity, assume that $\ell \sim \ell'$ and $\ell' \sim \ell''$. Thus we have $g_1, \ldots, g_k \in G$, $\ell_1^0, \ldots, \ell_k^0, \ell_1^1, \ldots, \ell_k^1 \in \Lambda$ forming a chain between ℓ and ℓ'; we have $h_1, \ldots, h_p \in G$, $\eta_1^0, \ldots, \eta_p^0, \eta_1^1, \ldots, \eta_p^1 \in \Lambda$ forming a chain between ℓ' and ℓ''.

For $i = 1, \ldots, p$, inductively define $\nu_i \in \Lambda$ such that $\nu_0 = \ell$ and $h_i(\nu_{i-1} - \phi_\Lambda)$ intersects $\nu_i - \phi_\Lambda$ (this is always possible by Property 3 of the covering set Λ). Then it is easily verified that we have a chain between ℓ and ℓ'', with

$$h_1, \ldots, h_p, h_{p-1}^{-1}, \ldots, h_1^{-1}, g_1, \ldots, g_k \in G$$

and the lines

$$\nu_1, \ldots, \nu_p, \nu_{p-1}, \ldots, \nu_1, \ell_1^0, \ldots, \ell_k^0 \in \Lambda$$
$$\eta_1^1, \ldots, \eta_p^1, \eta_{p-1}^0, \ldots, \eta_1^0, \ell_1^1, \ldots, \ell_k^1 \in \Lambda$$

forming the two sides of the chain. ∎

Let Γ be the set of equivalence classes induced by $\sim$, and let Γ_n denote the set of all cardinality-n multisets of Γ. Define the function $\gamma : \Lambda \to \Gamma$ to map a line of Λ to its equivalence class, and define $\gamma_n : \Lambda_n \to \Gamma_n$ analogously for multisets. We are now in a position to consider the relationship between $\sim$ and the structure of $I_\lambda(G)$.

Lemma 5 *Let $L^0, L^1 \in \Lambda_n$, if $\gamma_n(L^0) = \gamma_n(L^1)$ and $f \in I_\lambda(G)$, then $f(L^0) = f(L^1)$.*

Proof. As usual, it is enough to assume that L^0 and L^1 differ by only one element; that is, there is some $\ell^0 \in L^0$ and $\ell^1 \in L^1$ with $\ell^0 \neq \ell^1$. The relation $\sim$ says that we may find $g_1, \ldots, g_k \in G$, $\ell_1^0, \ldots, \ell_k^0, \ell_1^1, \ldots, \ell_k^1 \in \Lambda$ satisfying the condition of (1). Using this chain, we are going to construct a sequence of cardinality-n subsets of Λ, beginning with L^0 and ending with L^1, such that the property of invariance ensures f takes the same value on each consecutive pair.

Let $P_0^{0\prime}$ be a realization of L^0 and $P_0^{1\prime}$ a realization of L^1, with $p_0^{0\prime}$ and $p_0^{1\prime}$ the realizations of ℓ^0 and ℓ^1. These are chosen so that $g_1 p_0^{0\prime} \in \ell_1^0 - \phi_\Lambda$ and $g_1 p_0^{1\prime} \in \ell_1^1 - \phi_\Lambda$, and neither $g_1 P_0^{0\prime}$ nor $g_1 P_0^{1\prime}$ contains two points on the same line of Λ. By the properties of the chain from (1) and since ϕ_Λ contains no ray, it is easy to verify that these constraints can be satisfied. Set $P_1^0 = g_1 P_0^{0\prime}$ and $P_1^1 = g_1 P_0^{1\prime}$. These two sets differ only by points lying on ℓ_1^0 and ℓ_1^1.

Proceeding inductively, assume we have point sets P_i^0 and P_i^1 differing only by points that lie on ℓ_i^0 and ℓ_i^1. Call these points p_i^0 and p_i^1. We construct $P_i^{j\prime}$ $(j = 0, 1)$ by replacing p_i^j with $p_i^{j\prime} \in \lambda p_i^j$ such that $g_{i+1} p_i^{j\prime} \in \ell_{i+1}^j - \phi_\Lambda$. This is possible by the properties of the chain (1). If necessary, we also adjust the other points (keeping their projections the same) so that no two are mapped to the same line of Λ by g_{i+1}. Now setting $P_{i+1}^j = g_{i+1} P_i^{j\prime}$, we have two point sets differing only by points that lie on ℓ_{i+1}^0 and ℓ_{i+1}^1.

We observe the following facts:

- $\lambda P_i^j = \lambda P_i^{j\prime}$, so $f(\lambda P_i^j) = f(\lambda P_i^{j\prime})$.
- $P_{i+1}^j = g_{i+1} P_i^{j\prime}$, so $f(\lambda P_{i+1}^j) = f(\lambda P_i^{j\prime})$.

Thus by transitivity, $f(L^0) = f(\lambda P_k^0)$ and $f(L^1) = f(\lambda P_k^1)$. But P_k^0 and P_k^1 differ only by points on ℓ_k^0 and ℓ_k^1, *which are the same line of* Λ. Thus, $\lambda P_k^0 = \lambda P_k^1$, so $f(\lambda P_k^0) = f(\lambda P_k^1)$ and $f(L^0) = f(L^1)$. The result now follows from Lemma 1. ∎

Lemma 6 *Let $g \in G$, $\ell_0, \ell_1, \ell_0', \ell_1' \in \Lambda$. Assume that $\ell_0 \sim \ell_1$, $g(\ell_0 - \phi_\Lambda)$ intersects $\ell_0' - \phi_\Lambda$, and $g(\ell_1 - \phi_\Lambda)$ intersects $\ell_1' - \phi_\Lambda$. Then $\ell_0' \sim \ell_1'$.*

Proof. Since $\ell_0 \sim \ell_1$, there is a chain between them of the form in (1). Applying g^{-1} to ℓ_0' and ℓ_1' yields a chain from ℓ_0' to ℓ_1'. ∎

Thus, we can view G as acting on Γ: if $x \in \Gamma$, and $\ell \in x$ such that $g(\ell - \phi_\Lambda)$ intersects $\ell' - \phi_\Lambda$, then gx is equal to $\gamma(\ell')$. Lemma 6 shows that this is well-defined. Moreover, it is easily verified that for $g, h \in G$ and $x \in \Gamma$, $g(hx) = (gh)x$. We can extend this group action to Γ_n, thereby obtaining the main "structure theorem" of this section.

We will say that a function g refines a function f if for all sets on which g is constant, f is also constant. We will say that g strictly refines f if g refines f and for some x, y, $f(x) = f(y)$ but $g(x) \neq g(y)$. Call $f \in I_\lambda(G)$ *maximal* if there is no $g \in I_\lambda(G)$ that strictly refines it. Then we have

Theorem 7 *Let f be a maximal invariant in $I_\lambda(G)$. Then the cardinality of the range of f is equal to the cardinality of the set of orbits of Γ_n with respect to G.*

Proof. Any function in $I_\lambda(G)$ must be constant on all the sets in a single element of Γ_n, by Lemma 5. Thus, it must be constant on all the sets in a single orbit of Γ_n.

Now, consider a function f that takes a different value on each orbit of Γ_n. Clearly this is possible, since f maps into a set that has cardinality at least as large as that of the reals. We claim that f is invariant with respect to λ and G. If not, there is some $P \subset D_\Lambda$ of cardinality n and $g \in G$ such that $f(\lambda P) \neq f(\lambda g P)$. But applying g to the lines of λP, we get a set of lines that intersect $\lambda g P$ at the points of $gP \subset D_\Lambda$; by the definition of an orbit, this means that $\gamma_n(\lambda P)$ and $\gamma_n(\lambda g P)$ belong to the same orbit of Γ_n. Since f is constant on any given orbit, this is a contradiction.

f is a maximal invariant, since there is clearly no $g \in I_\lambda(G)$ that refines it. Conversely, every maximal invariant must take a different value on each orbit of Γ_n, since otherwise there would exist an invariant strictly refining it. Thus, we have the stated result. ∎

We can phrase implications of this result in a number of ways.

Corollary 8 *Let G be a group acting on $\Re^3$ and λ a projection.*

(i) If (G, λ) has Property B, then $I_\lambda(G)$ is trivial.

(ii) If $\sim$ partitions Λ into k equivalence classes, then any function in $I_\lambda(G)$ takes on at most $\binom{n+k-1}{k-1}$ values.

(iii) If Γ_n has infinitely many orbits with respect to G, then $I_\lambda(G)$ contains an infinite-valued invariant.

Now that we have a clearer picture of what it means for $I_\lambda(G)$ to be trivial, we can try to identify those subgroups of E_3 that have only trivial invariants. There is one more statement we can make about projections in general; after that, we will turn to the special cases of orthographic and central projection.

Theorem 9 *Let λ be any projection and G the group of rotations about a point $z \in D_\Lambda$. Then $I_\lambda(G)$ is trivial.*

Proof. We will show that (G, λ) has Property B. Let $\ell, \ell' \in \Lambda$; we consider only the case in which neither line is λz, as all other cases are easier.

Choose $s \in \ell - \phi_\Lambda$ and $t \in \ell' - \phi_\Lambda$. There exists a rotation g_0 such that $g_0 s \in \lambda z - \phi_\Lambda$, since there are two points on λz to which s can be rotated and at most one lies in ϕ_Λ. Consider the set $H = \{g \in G : gs = g_0 s\}$. The set Ht is a circle; by the definition of a covering set, ϕ_Λ cannot contain this entire circle. Thus there exists a rotation g such that $gs \in \lambda z - \phi_\Lambda$ and $gt \in D_\Lambda$.

Now choose g' such that $g'gt \in \lambda z - \phi_\Lambda$. Thus using $g, g' \in G$, we have a chain of the form in (1), with $\ell, \lambda z, \lambda z$ and $\ell', \lambda gt, \lambda z$ constituting the two sides of the chain. ∎

Consider the example in which λ is central projection and G is the group of all rotations about the center of projection. This shows that the requirement $z \in D_\Lambda$ is necessary, since in this example, it is easy to construct non-trivial functions that are invariant with respect to G and λ.

4. Orthographic and Central Projection

We now consider groups consisting of a single rotation about an axis; if u is the axis, we will denote the corresponding rotation group by R_u. First we consider the case of orthographic projection, where all of the lines of λ are at the same orientation. Call this orientation $\hat{v}$.

Lemma 10 *The group R_u, where u is parallel or perpendicular to $\hat{v}$, has non-trivial, infinite-valued, invariants.*

Proof. If u is parallel to $\hat{v}$, then there are clearly non-trivial invariants (because the rotation is about the direction of projection). Let us consider the case in which u is perpendicular to $\hat{v}$. If S is a 3D point set, then for all $g \in R_u$, $\lambda g S$ always has the same projection onto λu. Any non-trivial function based on this projection will be a non-trivial invariant. ∎

Lemma 11 *If u is not parallel or perpendicular to $\hat{v}$, then R_u has only trivial invariants.*

Proof. As usual, we show that (G, λ) has Property B. Let $\ell, \ell' \in \Lambda$, and choose any rotation g such that $g\ell$ and $g\ell'$ are not vertical. Let ℓ_1 and ℓ_1' be lines of Λ that pass through u and intersect $g\ell$ and $g\ell'$ respectively. Consider the orbit of ℓ_1 with respect to G — it is a cone $\mathcal{C}$ with axis u and apex $\ell_1 \cap u$, and there is a point at which ℓ_1' intersects $\mathcal{C}$. Thus for some $g' \in G$, $g'\ell_1'$ intersects ℓ_1. Since $g'\ell_1$ intersects ℓ_1 (at $\ell_1 \cap u$), we have a chain from ℓ to ℓ', with ℓ, ℓ_1, ℓ_1 and $\ell', \ell_1', g'\ell_1'$ as the two sides. ∎

These two results can be elegantly summed up in the obvious way:

Theorem 12 *R_u has non-trivial invariants if and only if u is parallel or perpendicular to $\hat{v}$.*

Now we turn to central projection. First of all, we observe that if u is a line through the center of projection, then it is easy to construct non-trivial, infinite-valued functions invariant with respect to R_u and λ. As it turns out, rotations about any other line result in invariants that are only finite-valued.

Theorem 13 *If u does not pass through the center of projection, then the greatest number of values taken on by a function in $I_\lambda(R_u)$ is $n+1$.*

Proof. Let $\mathcal{P}$ be the plane through the center of projection and normal to u. We note the following three facts.

1. If $\ell \subset \mathcal{P}$, $\ell' \not\subset \mathcal{P}$, then ℓ and ℓ' are not related by $\sim$.
2. If $\ell, \ell' \subset \mathcal{P}$, then $\ell \sim \ell'$.
3. If $\ell, \ell' \not\subset \mathcal{P}$, then $\ell \sim \ell'$.

These three facts can be verified by arguments very similar to those above. Thus, $\sim$ partitions D_Λ into 2 equivalence classes, so by Corollary 8, any non-trivial invariant of R_u takes on at most $n+1$ values. Such an invariant exists – consider the function that counts the number of lines lying in $\mathcal{P}$. ∎

We now turn to two applications of the model we have developed here.

5. The Overhead Viewing Problem

Consider the following practical problem. We have an overhead camera viewing a tabletop below, and each object sitting on the table has some small (fixed) number of stable resting positions. We represent each object as a collection of models, one for each such stable position. Thus, if we picture the z-axis normal to the workspace, a given model is essentially free to undergo translation in x and y, and rotation about z, the viewing direction. We would like to know whether invariants are useful for recognizing what is sitting on the table.

Let us consider the two main models of projection — orthographic and central — in turn. For the case of orthographic projection, and a point set $P \subset \Re^3$ under these conditions, G and λ have a nice commutativity property: if G' is the group of all rigid motions in x and y, then there is an obvious isomorphism φ between G and G', and we observe that $\lambda g P = \varphi(g)\lambda P$. That is, to the overhead camera, P might as well be a two-dimensional point set that is free to undergo full rigid motion. Hence any classical invariant f of two-dimensional point sets with respect to rigid motion will be a non-trivial invariant with respect to G and λ.

Thus in this application, we would simply need to create a separate two-dimensional point set for each stable position of P, and label P with the corresponding finite set of values of f. Thus, we conclude that despite the depressing overtones of the preceding sections, invariants are well-suited to certain practical problems.

The results for central projection, however, are much bleaker. We observed in Section 4 that there are non-trivial invariants here with respect to rotation about a line of Λ; but the addition of translation changes everything, in view of the following result.

Lemma 14 *If λ is central projection and G is translation in x and y (no rotation!), then any function invariant with respect to G and λ takes on at most $n+1$ values.*

Proof. Let $\mathcal{P}$ be the plane through the center of projection and perpendicular to $\hat{v}$ (the viewing direction). Very much as in the proof of Theorem 13, we have three important facts.

1. If $\ell \subset \mathcal{P}$, $\ell' \not\subset \mathcal{P}$, then ℓ and ℓ' are not related by $\sim$.
2. If $\ell, \ell' \subset \mathcal{P}$, then $\ell \sim \ell'$.
3. If $\ell, \ell' \not\subset \mathcal{P}$, then $\ell \sim \ell'$.

As before, an (n+1)-valued invariant would be one that counted the number of lines in $\mathcal{P}$. ∎

A final note is that in the practical problem, a camera that is reasonably high above the workspace is taking pictures essentially by orthographic projection, and so it appears that invariants are very much of use in such a situation.

6. Segments and Frameworks

It is somewhat surprising that the model developed so far can be extended, with almost no effort, to prove that there are no non-trivial invariants of a much broader class of three-dimensional objects. Specifically, define a *framework* F to be the geometric realization of a graph in $\Re^3$: it consists of a set of n points $V \subset \Re^3$, and a set of segments E with endpoints in V. Let us fix some graph $\mathcal{T}$ and denote by $R(\mathcal{T})$ the set of all frameworks that are isomorphic to $\mathcal{T}$ as graphs (this implicitly fixes the cardinality of the frameworks).

The projection of a framework into two dimensions is simply a graph on points in $\Re^2$; however, there are several ways in which one could choose to represent the projection of a framework. In this section, we project frameworks "symbolically" — if $F \in R(\mathcal{T})$, then λF is a graph (L, E_L), where L is a set of n lines of Λ (the vertices) and E_L is a set of unordered pairs of these lines (the edges). This will be the "input" to any function defined on projections of frameworks. We denote the set $\{\lambda F : F \in R(\mathcal{T})\}$ by $\lambda R(\mathcal{T})$. In this context, the definition of a function invariant with respect to G and λ remains the same as in the previous sections. Observe the new two- and three-dimensional conditions on our objects.

- The three-dimensional condition: F is a framework isomorphic to $\mathcal{T}$, and none of its points lie in ϕ_Λ.
- The two-dimensional condition: No two points of F are mapped to the same line by λ.

We will say that (L, E_L) and (L', E'_L) in $\lambda R(\mathcal{T})$ differ by a single vertex if the deletion of one line and all its incident edges from each set makes them identical. We say that a function f has Property A' if $f(L, E_L) = f(L', E'_L)$ for all $(L, E_L), (L', E'_L)$ differing by a single vertex.

Lemma 15 *If f has Property A' then f is constant.*

Proof. Let (L, E_L) and (L', E'_L) be arbitrary members of $\lambda R(\mathcal{T})$. They are isomorphic as graphs, so there is some function φ such that $\varphi(\ell_i) = \ell'_i$ and $(\varphi(\ell_i), \varphi(\ell_j))$ is an edge in E'_L if and only if (ℓ_i, ℓ_j) is an edge in E_L. Set $(L^0, E^0_L) = (L, E_L)$, $(L^n, E^n_L) = (L', E'_L)$, and for $i = 1, \ldots, n-1$, define (L^i, E^i_L) such that $L^i = \{\ell_1, \ldots, \ell_i, \ell'_{i+1}, \ldots, \ell'_n\}$ and E^i_L is the set of edges that makes (L^i, E^i_L) isomorphic to (L, E_L) as a graph. Then for $i = 0, \ldots, n-1$, (L^i, E^i_L) differs from (L^{i+1}, E^{i+1}_L) by a single vertex (i.e. $\ell'_{i+1} \in L^i$ and $\ell_{i+1} \in L^{i+1}$), so $f(L^i, E^i_L) = f(L^{i+1}, E^{i+1}_L)$. By transitivity, we have $f(L, E_L) = f(L', E'_L)$; since these two graphs were arbitrary, f is constant. ∎

Now let us consider a function f on projections of frameworks that is invariant with respect to some G and λ, such that (G, λ) has Property B. Then the proof of Lemma 5 carries over with no modification to show that if (L^0, E_L^0) and (L^1, E_L^1) differ by a single vertex, then for $j = 0, 1$ we have a sequence of frameworks $F_0^{j\prime}, F_1^j, F_1^{j\prime}, \ldots, F_{k-1}^{j\ \prime}, F_k^j$ such that

- $\lambda F_i^j = \lambda F_i^{j\prime}$.
- For some $g \in G$, $gF_i^{j\prime} = F_{i+1}^j$.
- $\lambda F_k^0 = \lambda F_k^1$.

Since f is invariant with respect to G and λ, f takes the same value on each consecutive pair in these sequences, so by transitivity, $f(L^0, E_L^0) = f(L^1, E_L^1)$. Thus f has Property A' and is trivial. We therefore conclude

Theorem 16 *If (G, λ) has Property B and f is a function on $R(\mathcal{T})$ that is invariant with respect to G and λ, then f is trivial.*

This can be restated in a somewhat looser but more descriptive way.

Corollary 17 *If (G, λ) has Property B, then the only invariant of a framework with respect to G and λ is its topology as a graph.*

This very general theorem includes our results on point sets as a special case, and it also shows that if (G, λ) has Property B, there are no invariants on sets of n segments, n triangles, n tetrahedra, n bent segments with a fixed number of bends per segment, and so on.

7. Summary and Conclusions

We have provided an algebraic characterization of groups and projections for which there are non-trivial invariants, and discovered that for most commonly considered cases, there are in fact no non-trivial invariants. Our proofs do not rely on degeneracies based on the viewing direction, and thus hold even under the assumption that functions on images can recognize multiple or occluded points. Moreover, the proof technique generalizes to show that the same algebraic characterization holds for invariants of wire-frame models. Nevertheless, there are still cases involving projection in which invariants are useful in object recognition. The most notable is the overhead viewing problem, considered in Section 5.

The challenge now appears to be coming up with reasonable sets of conditions under which there are non-trivial invariants under projection. For example, what if we were guaranteed that all the three-dimensional models we were viewing were drawn from some restricted class? In effect, we would be strengthening the "three-dimensional conditions" mentioned above, so that the techniques used in our non-existence proofs — construction of arbitrary realizations — would no longer apply.

We know in fact that there are invariants under projection of sets of points known to be coplanar: these are precisely the affine invariants. Other restricted object classes for which there are invariants under projection are considered in [5]. It seems that a promising idea would be to define classes of three-dimensional objects more general than coplanar point sets for which invariant-based techniques are still of use.

Acknowledgements

The authors would like to thank Paul Chew, John Dalbec, Robert Kleinberg, Paul Pedersen, Bernd Sturmfels, and Peter Wayner for helpful discussions and comments on the topic of this paper.

References

[1] J. Burns, R. Weiss, and E. Riseman (1991) "View variations of point set and line segment features," DARPA-ESPRIT Workshop on Invariants in Computer Vision, Iceland.

[2] D.T. Clemens, and D.W. Jacobs (1991) "Indexing and Object Recognition", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(10), pp. 1007–1017.

[3] D. Forsyth, J.L. Mundy, A. Zisserman, C. Coelho, A. Heller and C. Rothwell (1991) "Invariant descriptors for 3-D object recognition and pose determination", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(10), pp. 971-991.

[4] Klein, F. *Elementary Mathematics from an Advanced Standpoint: Geometry*, MacMillan, New York, 1939.

[5] Y. Moses and S. Ullman (1992) "Limitations of Non-Model-Based Recognition Schemes," *Proc. 2nd European Computer Vision Conference*, LNCS 488, pp. 820–828, Springer-Verlag, Berlin.

[6] L. VanGool, P. Kempenaers, and A. Oosterlinck (1991) "Recognition and semi-differential invariants", *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 454–460.

[7] P. Wayner (1991) "Efficiently Using Invariant Theory for Model-based Matching", *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.*

[8] I. Weiss (1989) "Projective Invariants of Shapes" *Proceedings of the DARPA Image Understanding Workshop*, pp. 1125–1134.