

GEOMETRIC CONSTRAINTS FROM PLANAR SURFACE PATCH MATCHING

Robert B. Fisher

Dept. of Artificial Intelligence, University of Edinburgh
5 Forrest Hill, Edinburgh EH1 2QL, Scotland, United Kingdom

Some three dimensional scene analysis programs pair data surface patches (i.e. as extracted from range data) to model surface patch features. From the correspondences, a variety of rotation and translation constraints of the model can be deduced, depending on the assumptions behind the correspondences. This paper elaborates some of the constraints available when single planar patches are related. The constraints are expressed algebraically for patches that are: nearby, touch, contained, or registered. The constraints are used to estimate object position, given a single patch correspondence. Examples of their use are given.

Fisher and Orr [1] have been investigating algebraic constraints as a representation of partial position constraints, much as in the ACRONYM system [2]. This research is being undertaken as part of the IMAGINE II scene analysis system [3, 4], which is currently concerned with surface-based object recognition.

The data primitives in the matching are surface patches, which can be obtained by a variety of techniques for segmenting range data (e. g. [5, 6]). These processes result in data patches with consistent curvature signs. Objects are defined using the SMS models [7], with surface patches as one of the primitive features.

The model matching process is not described here, but results in making correspondences between model and data surface patches. With these correspondences, it is possible to estimate the position of the object containing the patches (assuming the correct model has been chosen and the correspondences were made correctly) or reject the correspondences as inconsistent.

A key element of the geometric analysis is formulating algebraic constraints between a rotated and translated model and the observed surface patches. The constraints attempt to define or refine the set of allowable positions for the object, consistent with the observed data patches. For example, pairing a planar data patch with a planar model patch immediately constrains the rotations allowable, because the transformed model normal \vec{n}_m must be aligned with the data normal \vec{n}_d , within an angular observation error ϵ . Hence,

$$T\vec{n}_m \cdot \vec{n}_d \geq \cos(\epsilon)$$

is a constraint on the rotation component of the position T (here expressed as a prefix transformation on the model normal \vec{n}_m .)

There are a variety of constraints on both the rotation and translation, depending on the assumptions made about the relationship between the model and data patches, such as whether the patches overlap or one is contained in the other. The generation of the constraints is the responsibility of the matching program, which depends on the constraint analysis program for solving them. Each model-to-data correspondence may create several constraints. Moreover, there may be several correspondences that each add constraints. Altogether, an analysis may generate many constraints that have to be solved to generate a single position estimate.

We have been using a value passing network [1] based on interval arithmetic [8] to solve the constraints, though other representation and solution techniques are possible, such as symbolic algebra [2] or statistical [9] techniques. The networks are created from an underlying algebraic description of three-dimensional geometric relations and transformations. These networks implement the five key visual geometric reasoning functions: locate, predict, transform, inverse and merge [10]. The networks also produce tighter bounds than simply using symbolic algebra and have greatly improved efficiency (over, e.g. ACRONYM [2]).

The forms of the algebraic constraints tend to be few and hence standard subnetwork modules can be developed, with instances allocated as geometric constraints were identified during model matching. For example, network modules exist for constraints like: "a model vector is transformed to a data vector" and "a data boundary must lie inside a transformed model boundary".

As there are a variety of model and data surface patches, a variety of constraints can be generated. Elsewhere [11], we consider the constraints generated from pairing instances of planar, cylindrical, conical and generic surface patches. This paper presents the algebraic position constraints that we have identified for simply pairing planar model and data patches, under a variety of assumptions about the correspondence.

There is some similarity between these results and our earlier work [12], except that that work used two-dimensional image measurements, whereas the new results use three-dimensional measurements, and the previous work also used a set of heuristics to infer the positions, whereas here the position constraints are expressed explicitly and solved in a network computation.




Figure 1: *Planar Patch Descriptions*




Figure 2: *Width Bounding of An Irregular Figure at One Orientation*

1 Data Patch Descriptions

In this paper, we assume that the features extracted from the image data are planar surface patches. A planar patch (see Figure 1) is flat and has bounded extent. Irregular shapes cause problems, because it is possible that a data patch might only be paired with a model patch in one way, and we would like to exploit as much information as possible. However, we would also like to use a simple scheme that can be used with all patches, irrespective of shape. Hence, we represent planar patches by the bounding width ($\omega_d(\phi)$) at each orientation (ϕ) (see Figure 2) for our position constraints. This width is a measure of the maximum extent of the surface in each direction. The orientation $\phi = 0$ is chosen arbitrarily, and each width measure has an error range of $\pm\delta_w$ associated with it. Additionally, we use the position (\vec{p}_d) of the center-of-mass (with the isotropic error range $\pm\vec{\gamma}$), the direction \vec{n}_d of the surface normal at that point (with angular error $\pm\epsilon$) and an arbitrary direction \vec{a}_d in the patch (observed with angular error $\pm\beta$).

2 Model Patch Descriptions

The constraints given in the following sections are designed for use with the planar surface primitives of the SMS object representation system [7] used in the IMAGINE II system [3] being developed.

SMS is a geometric modeling system, based on multi-

ple, alternative representations (curve, surface or first and second-order volumetric primitives). The models need not have all features specified, and the features are usually represented in a subcomponent hierarchy. The features may have parameterized sizes and properties (e.g. surface curvatures) and there may be unconstrained degrees-of-freedom in the reference frame transformations. The goal of SMS is to represent the visual aspects of an object that characterize its identity, rather than describe its shape. Hence, the modeling approach aims to provide representations that closely correspond to reliably extractable image features.

The model surface patches are characterized similarly to the data patches except that they are represented with a m subscript instead of a d subscript. Not all of the shape information is modeled explicitly in the SMS scheme, but all are straightforward and trivial extensions.

3 Constraints From The Planar Patch Pairings

3.1 Specifying Position Constraints

Our key geometric reasoning data type for high-level computer vision is the position. It represents the relative spatial relationship between two visual features (e.g. world-to-camera, camera-to-model, or model-to-subcomponent). Positions are represented using a 3-vector representing for relative translation and a unit quaternion (4-vector) for relative orientation (of the form $(\cos(\theta/2), \sin(\theta/2)\vec{w})$ for a rotation of θ about the axis \vec{w}).

The key geometric relationships concern relative position and have two forms – exact and partially constrained. An example of an exact form is: let object A be at global position (\vec{r}_A, \vec{t}_A) , (rotation \vec{r}_A and translation \vec{t}_A) and object B be at $(\vec{r}_{AB}, \vec{t}_{AB})$ relative to A. Then, the global position of B is:

$$(\vec{r}_B, \vec{t}_B) = (\vec{r}_{AB} * \vec{r}_A, \vec{r}_{AB} * \vec{t}_A * \vec{r}'_{AB} + \vec{t}_{AB})$$

where $*$ is the quaternion multiplication operator and $'$ is the quaternion inverse operator.

A partially constrained position is given by an inequality constraint, such as:

$$t_{Az} \geq 50$$

This means that the z component of A's global position is at least 50.

Other relationships concern vectors or points linked by a common transformation, as in $T\vec{v}_1 = \vec{v}_2$, or the proximity of points or vectors:

$$\|\vec{p}_1 - \vec{p}_2\| \leq \epsilon$$

A complication is that each data measurement may have some error or uncertainty, and hence the estimated values may also have these. Or, a variable may be only partially constrained in the model or by a *a priori* scene

information. Hence, each numerical quantity is represented by an interval [8] of possible values.

Decisions about patch pairings is the responsibility of the scene analysis program, that then determines which constraints are generated. The main consideration is the extent to which we can constrain their relative position, based on *a priori* knowledge. For example, if we are observing rigid man-made identical objects, then we can assume that the observed data patch will be a subset of the corresponding model patch. This is our assumption A3, and there are both stronger and weaker assumptions considered. The assumptions (from weak-to-strong) are:

A1 Both model and data patches are part of a larger surface of bounded size R , but the patches need not overlap.

A2 The model and data patches must touch or overlap.

A3 The data patch is a subset of the model patch.

A4 The centroids of both the model and data patches correspond to the same point.

Note that each assumption is increasingly strict and so any constraint that applies when A1 holds also applies when (e.g.) A3 holds.

Position constraints are expressed algebraically, using the prefix transformation T applied to some model feature (e.g. surface normal), and showing the relation the transformed model feature must have with observed data features. Where possible, both rotation and translation constraints are specified. Also, some constraints are not perfect, to simplify their formulation.

Translation constraints are usually of the form:

$$(Tp_m - p_d) \in I(\vec{r})$$

meaning the vector distance between the transformed model point and a given data point lies within the interval vector [8] $I(\vec{r})$ defined by:

If $\vec{r} = (r_x, r_y, r_z)$ then

$$I(\vec{r}) = \{(x, y, z) : |x| \leq |r_x|, |y| \leq |r_y|, |z| \leq |r_z|\}$$

The interval vector expression $I(\vec{a}) + I(\vec{b})$ has the value:

$$\{\vec{c} + \vec{d} : \vec{c} \in I(\vec{a}), \vec{d} \in I(\vec{b})\}$$

The translation constraints are based on displacements in three orthogonal directions determined by the surface normal (\vec{n}_d), the elongation axis (\vec{a}_d) and the cross product (\vec{c}_d) of the two:

$$\vec{c}_d = \frac{\vec{n}_d \times \vec{a}_d}{\|\vec{n}_d \times \vec{a}_d\|}$$

3.2 Planar Data-to-Model Patch Constraints

We now look at the constraints that can be generated by pairing planar model and data patches. The main

rotation constraints are based on aligning normals and elongation axes. The translation constraints are largely based on proximity and distance between the designated points on the surface.

The first rotation constraint relates the surface normals, within the observed angular error of ϵ , assuming A1:

$$Tn_m \cdot \vec{n}_d \geq \cos(\epsilon)$$

If A3 holds, we can relate the surface orientations, because the bounding widths may limit the rotation of the model surface in the plane of the data surface, such that it fits around the observed data patch. Let:

$\omega_m(\phi)$ = bounding width (see Figure 2) of model surface at orientation ϕ

$\omega_d(\phi)$ = bounding width of the data surface at orientation ϕ

$\Omega = \{\theta : \forall \phi (\omega_m(\phi - \theta) \geq \omega_d(\phi) - \delta_\omega)\}$ is the set of possible rotations of the model that can fit around the data

\vec{a}_m = the nominal unit vector in the plane of the model surface (at orientation ϕ_m)

\vec{a}_d = average orientation of \vec{a}_m over Ω (at orientation ϕ_d)

If $size(\Omega)/2 + \beta < \pi$ then:

$$Ta_m \cdot \vec{a}_d \geq \cos(size(\Omega)/2 + \beta)$$

If more than one interval of angles exists in Ω , then each interval creates an independent constraint.

If assumption A4 holds, then a similar constraint can be developed, except where $\omega(\phi)$ refers to the cross-section width about the patch centers.

All translation constraints have the form:

$$(Tp_m - p_d) \in I(A\vec{n}_d) + I(B\vec{a}_d) + I(C\vec{c}_d) + I(\vec{r})$$

where the constants A , B and C depend on the specific assumptions:

Asmp.	A	B
A1	$\epsilon\sqrt{B^2 + C^2}$	$R - \frac{1}{2}(M + \omega_d(\phi_d)) - \delta_\omega$
A2	$\epsilon\sqrt{B^2 + C^2}$	$\frac{1}{2}(N + \omega_d(\phi_d) + \delta_\omega)$
A3	$\epsilon\sqrt{B^2 + C^2}$	$\frac{1}{2}(U - \omega_d(\phi_d) + \delta_\omega)$

Asmp.	C
A1	$R - \frac{1}{2}(M + \omega_d(\phi_d + \pi/2)) - \delta_\omega$
A2	$\frac{1}{2}(N + \omega_d(\phi_d + \pi/2) + \delta_\omega)$
A3	$\frac{1}{2}(V - \omega_d(\phi_d + \pi/2) + \delta_\omega)$

Where:

$$M = \min_\phi(\omega_m(\phi))$$

$$N = \max_\phi(\omega_m(\phi))$$

$$U = \max_\phi(\omega_m(\phi)) \text{ over } \phi - \phi_m + \phi_d \in \Omega$$

$$V = \max_\phi(\omega_m(\phi + \pi/2)) \text{ over } \phi - \phi_m + \phi_d \in \Omega$$

If assumption A1 or A2 holds, then the orientation of \vec{a}_d is arbitrary, and is chosen to correspond with angle

$\phi_d = 0$. If assumption A4 holds, then all translation constraints are of the form:

$$(T\vec{p}_m - \vec{p}_d) \in I(\vec{\gamma})$$

and are not listed further below. ($\vec{\gamma}$ is the isotropic error on the observed point \vec{p}_d).

4 Examples

There are a great many constraints possible from the results presented in this paper, and it is not useful to present experiments for all. Instead, this section will present results from the planar model patch to planar data patches pairing relationships for both assumptions A2 (overlap) and A3 (data subset of model).

The experiments will involve the estimation of the reference frame of a single data patch, given its observation, and applying the constraints given in Section 3.

The experiments start with a known model surface normal ($\vec{n}_m = (0, 0, -1)$), axis ($\vec{a}_m = (1, 0, 0)$) and central point ($\vec{p}_m = (0, 0, 0)$).

We then randomly pick a rotation quaternion (distributed uniformly over the 4-dimensional unit sphere) and a translation (each component distributed uniformly over the range ± 100). Using this position (P_{true}), we transform the input features to derive some scene features: \vec{n}_s , \vec{a}_s and \vec{p}_s .

The scene features are then corrupted with randomly generated noise to generate the observed features: \vec{n}_d , \vec{a}_d and \vec{p}_d . The corruptions were uniformly distributed subject to:

$$\begin{aligned} \vec{n}_s \cdot \vec{n}_d &\geq \cos(\epsilon) \\ \vec{a}_s \cdot \vec{a}_d &\geq \cos(\beta) \\ (\vec{p}_s - \vec{p}_d) &\in I(\vec{\gamma}) \end{aligned}$$

The observed features were then paired with the modeled features, to derive the geometric position constraints, based on the results of Section 3.2.

The constraints were input into a geometric reasoning network [1] set up for evaluating these constraints. Figure 3 shows the network for the experiments. This network implements a SUP/INF interval bounding of all components of the position, as a function of the intervals bounding the inputs. The inputs for this network are:

- the model vectors \vec{n}_m , \vec{a}_m and \vec{p}_m extended to form intervals with width 0.0001,
- the data vectors \vec{n}_d and \vec{a}_d extended to form intervals with width 0.0001,
- the estimated position \vec{p}_d extended to form a larger interval given by the translation constraint and
- the rotation dot product constraints dot_n and dot_a given by the rotation constraints

$$T\vec{n}_m \cdot \vec{n}_d \geq dot_n$$

$$T\vec{a}_m \cdot \vec{a}_d \geq dot_a$$

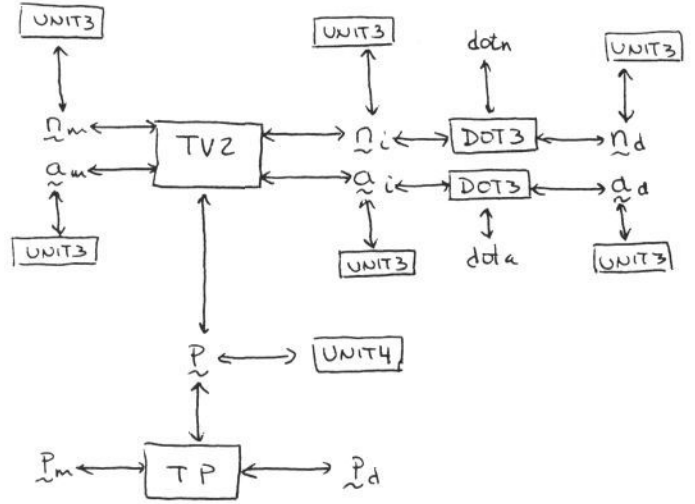


Figure 3: Geometric Reasoning Network For Estimating Position

The main component of the network is the “TV2” module near the center. This module relates the input and output vectors according to the position \vec{P} (near the bottom left) and is defined by the algebraic relationships between the input and output vectors. Similarly, the translation component of \vec{P} is estimated by the “TP” module that relates the model and observed feature points. There are also two “DOT3” modules that implement the rotation dot product constraints. The vectors \vec{n}_i and \vec{a}_i are intermediate results. Finally, there are several “UNIT3” and “UNIT4” modules that enforce unit 3-vector and 4-vector constraints.

This sets up the network for estimating the position \vec{P} . We then evaluate the network to estimate the position.

The relationship between the true and estimated position provides the statistical output described below. All the reported results below are based on 100 randomly generated trials of the particular experiment. All experiments used the same 100 trial data sets.

The following parameters were used (where p is a percentage parameter described below):

Model Value	Data Value	Error
$\omega_m(\phi)$	$\omega_d(\phi)$	$p\omega_m(\phi)$
\vec{n}_m	\vec{n}_d	transformed
\vec{a}_m	\vec{a}_d	transformed
\vec{p}_m	\vec{p}_d	transformed

* - the test figure was a rectangle of length 100 and width 30

The experiment was run using assumptions A2 and A3 and the parameter p was allowed to vary from 1.0 to 0.1 in steps of 0.1. This allowed the exploration of the constraint obtainable because the patch was limited in its ability to rotate relative to the model patch.

In all experiments, the true object position was contained within the estimated interval position, so the constraints

appear to be correct. There is then the question of how tight the constraints are, and this is the main point of the experiments.

Five statistics were extracted from the experiments:

1. the number of times that the rotation had at least two independent constraints and hence was fully constrained (though subject to error).
2. the mean angle between the input and output quaternion vectors (a measure of how well the rotation was estimated, and ideally should be zero)
3. the average width of the interval of possible values for a component of the rotation quaternion (a measure of how tight the estimates are)
4. the mean distance between the input and estimated translation (a measure of how well the translation was estimated, and ideally should be zero)
5. the average width of the interval of possible values for a component of the translation (a measure of how tight the estimates are)

More precisely, the error measures are defined by:

$$\begin{aligned} \text{Angle} &= \frac{1}{N} \sum_{i=1}^N \cos^{-1}(\hat{d}_{0i}q_{0i} + \hat{d}_{1i}q_{1i} + \hat{d}_{2i}q_{2i} \\ &\quad + \hat{d}_{3i}q_{3i}) \\ \text{Q-Width} &= \left(\frac{1}{N} \sum_{i=1}^N \text{width}(D_{0i}) * \text{width}(D_{1i}) \right. \\ &\quad \left. * \text{width}(D_{2i}) * \text{width}(D_{3i})\right)^{\frac{1}{4}} \\ \text{Trans} &= \frac{1}{N} \sum_{i=1}^N \sqrt{\text{dist}} \\ \text{dist} &= (t_{xi} - v_{xi})^2 + (t_{yi} - v_{yi})^2 + (t_{zi} - v_{zi})^2 \\ \text{T-Width} &= \left(\frac{1}{N} \sum_{i=1}^N \text{width}(T_{xi}) * \text{width}(T_{yi}) \right. \\ &\quad \left. * \text{width}(T_{zi})\right)^{\frac{1}{3}} \end{aligned}$$

where:

- q_{ji} is the j^{th} component of the i^{th} true scalar rotation quaternion \bar{q}_i
- d_{ji} is the mean value of the j^{th} component (D_{ji}) of the i^{th} estimated rotation interval \bar{D}_i
- \hat{d}_{ji} is the unit-quaternion normalized value of d_{ji}
- t_{xi} is the mean value of the x-component (T_{xi}) of the i^{th} estimated translation interval vector \bar{T}_i
- v_{xi} is the x-component of the i^{th} true translation \bar{v}_i
- $\text{width}(X)$ is the width of scalar interval X

The results of the experiment for assumption A3 are shown in Table 1. As the shrinkage percentage was increased, the rotation and translation errors increased, as expected, because the model patch was freer to move about the data patch. At 20% there was a dramatic increase in the errors, because the patch size was small enough that it could rotate freely within the model patch, and hence only one direction vector was effectively constrained (the surface normal). The mean translation error was constant and mainly resulted from the isotropic error on the observed central point.

Table 1: *Plane-Plane Parameter Estimation Statistics for Assumption A3*

p (%)	Con	Angle	Q_Width	Trans	T_Width
100	100	0.23	0.69	1.01	3.82
90	100	0.27	0.74	1.01	12.74
80	100	0.31	0.77	1.01	22.75
70	100	0.41	0.89	1.01	32.80
60	100	0.46	0.95	1.01	44.52
50	100	0.53	1.04	1.01	56.83
40	100	0.77	1.25	1.01	72.26
30	100	1.07	1.48	1.01	92.86
20	0	1.42	1.52	1.01	100.28
10	0	1.42	1.52	1.01	107.20

Table 2: *Plane-Plane Parameter Estimation Statistics for Assumption A2*

p (%)	Con	Angle	Q_Width	Trans	T_Width
100	0	1.42	1.52	1.01	182.89
90	0	1.42	1.52	1.01	176.03
80	0	1.42	1.52	1.01	169.17
70	0	1.42	1.52	1.01	162.31
60	0	1.42	1.52	1.01	155.44
50	0	1.42	1.52	1.01	148.57
40	0	1.42	1.52	1.01	141.69
30	0	1.42	1.52	1.01	134.81
20	0	1.42	1.52	1.01	127.92
10	0	1.42	1.52	1.01	121.02

When assumption A2 was used, the results in Table 2 were obtained. Here, only one rotation constraint was obtained (that of the aligned normals). This explained the constant large rotation error. As before, the translation error was constant and small. The translation error range started larger, because of the larger area over which the model surface could move and still have overlap with the data surface. However, unlike the assumption A3 case, as the percentage of the observed surface decreased, the translation error also decreased, because the model surface had to overlap with the smaller data surface.

In general, we can see below that as the size of the data patch tends to zero, the two assumptions give nearly the same results, as is expected (because the effective translation constraints become the same).

5 Discussion

The paper identified several geometric constraints for estimating surface positions, given information about the shapes of planar patches and the relationships between the model and data patches. The experiments of the last section showed both that the constraints are valid, and how well they perform.

We would like to now discuss some of the implications of the constraints and experiments.

For simplicity and uniformity, we mainly used properties from a width bounding of the patches. Clearly, there are some boundary shapes (such as with a long, narrow protrusion) where this uniform bounding loses information. Other types of bounding could be developed, but the use of the width did allow a precise statement of the range of allowable rotations.

It is possible that tighter constraints could be developed for the cases. For example, the position constraints generally treated the three axis directions independently, and then transformed the bounds relative to these to bounds in the three coordinate directions. There is likely to be some dependence between the factors, allowing tighter position estimates.

The estimation of the patch orientation is not always accurate. This stems from three problems:

1. the axis constraints are not particularly strong, so the overall orientation is dramatically affected,
2. the TV2 rotation estimation module is not particularly effective when estimating rotations from inputs with substantial variation and
3. each stage of the TV2 rotation estimation process increases the size of the output intervals, so that modest errors are amplified in the output.

The result of these problems is that there are many rotation estimates where the precise direction of the rotation cannot be determined (i.e. determined to be \vec{q} or $-\vec{q}$) and hence the output rotation intervals are large.

While some features are strong and reliable, such as the normals on planar surfaces or curvature axes on strongly curved surfaces, many other direction vectors are weak, and do not greatly constrain object position. Work in progress shows that using constraints from pairs of surfaces (e.g. two normals) gives better position estimates. For example, the vector between two distant features gives much more reliable data than the axis directions on either.

Though our interval-based position estimation process may not be ideal, the position constraints presented in this paper hold, irrespective of the process used to estimate the position parameters from the constraints.

Acknowledgements

This work was funded by the University of Edinburgh. I'd like to thank J. Hallam, H. Hughes, and M. Cameron-Jones for suggestions with the text and work.

Bibliography

1. Fisher, R. B., Orr, M. J. L., "Solving Geometric Constraints in a Parallel Network", *Image and Vision Computing*, Vol 6, No 2, 1988.
2. Brooks, R. A., "Symbolic reasoning among 3-D models and 2-D images", *Artificial Intelligence*, Vol 17, p285, 1981.
3. Fisher, R. B., "The Design of the IMAGINE II Scene Analysis Program", Univ. of Edinburgh, Dept. of Artificial Intelligence, Working Paper 211, December 1987.
4. Fisher, R. B., *From Surfaces to Objects: Computer Vision and Three Dimensional Scene Analysis*, John Wiley and Sons Ltd, Chichester, 1989.
5. Besl, P. J., "Surfaces in Early Range Image Understanding", PhD Dissertation, Electrical Engineering and Computer Science Department (RSD-TR-10-86), University of Michigan, 1986.
6. Cai, L. D., Forthcoming PhD Thesis, Department of Artificial Intelligence, University of Edinburgh, 1989.
7. Fisher, R. B., "SMS: A Suggestive Modeling System For Object Recognition", *Image and Vision Computing*, Vol 5, pp 98 - 104, 1987.
8. Alefeld, G. and Herzberger, J., *Introduction to Interval Computations*, Academic Press, 1983.
9. Durrant-Whyte, H. F., "Uncertain geometry in robotics", *Proceedings of the IEEE Conference on Robotics and Automation*, Vol.2, p851, 1987.
10. Orr, M. J. L. and Fisher, R. B., "Geometric Reasoning for Computer Vision", *Image and Vision Computing*, Vol 5, p233, 1987.
11. Fisher, R. B., 3D Geometric Constraints From General Surface Patch Matching, submitted to *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
12. Fisher, R. B., "Using Surfaces and Object Models to Recognize Partially Obscured Objects", *Proc. 8th Int. Joint Conf. on Artificial Intelligence*, pp 989-995, 1983.