Probabilistic 3D Object Recognition *

Ilan Shimshoni and Jean Ponce
Department of Computer Science and Beckman Institute
University of Illinois, Urbana, IL 61801

Abstract: *A probabilistic 3D object recognition algorithm is presented. In order to guide the recognition process the probability that match hypotheses between image features and model features are correct is computed. A model is developed which uses the probabilistic peaking effect of measured angles and ratios of lengths by tracing iso-angle and iso-ratio curves on the viewing sphere. The model also accounts for various types of uncertainty in the input such as incomplete and inexact edge detection. For each match hypothesis the pose of the object and the pose uncertainty which is due to the uncertainty in vertex position are recovered. This is used to find sets of hypotheses which reinforce each other by matching features of the same object with compatible uncertainty subsets. A probabilistic expression is used to rank these hypothesis sets. The hypothesis sets with the highest rank are output. The algorithm has been fully implemented, and tested on real images.*

1 Introduction

One of the major problems in computer vision is to recognize 3D objects appearing in a scene as instances from a database of models. Several recognition systems extract features such as points [9, 13] or lines [3] from the image and match them to corresponding features in the database. They then verify each candidate hypothesis using other features in the image and finally rank the verified hypotheses.

When constructing such a system, the major problems to be addressed are: how to generate match hypotheses, in which order to process these hypotheses, how to verify them using additional image features, and how to rank the verified hypotheses. In our approach, we define a match hypothesis as the matching of a pair of edges or a trihedral corner (Figure 1) to corresponding feature sets in the model database. Combining these two types of features into a single recognition algorithm is one of the contributions of this paper. The hypotheses are ranked by computing the probability that each one is correct using the probabilistic peaking effect [4, 5, 6].

*This work was supported in part by the Beckman Institute and the Center for Advanced Study of the University of Illinois at Urbana-Champaign, by NSF under grant IRI-9224815, and by NASA under grant NAG 1-613.

In order to verify a hypothesis, we compute the pose of the object assuming it is correct using a simple approach which deals with both types of features. Hypotheses which match other features in the image to features in the same model should yield a compatible pose if they belong to the same instance of the model [2, 7]. However, by taking uncertainty in the values measured in the image into account, the pose is transformed from a point in the pose space to a subset of that space, and for hypotheses to possibly reinforce each other we check if their pose uncertainty subsets intersect. Sets of hypotheses whose pose subsets have a non-empty intersection reinforce each other; they are ranked by the probability that the match of the whole set is correct.

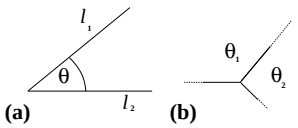

Figure 1: Image feature sets: (a)two adjacent lines; (b) a trihedral corner.

The rest of the paper is organized as follows. In Section 2 we develop a method for computing probability functions using equations we develop for angles and ratios. We use these probabilities for ranking match hypotheses in Section 3. In Section 4 we discuss pose estimation and estimate the effects of uncertainty on the pose. In Section 5 we present our probabilistic expression to rank matching hypotheses. Finally, in Section 6 we show experimental recognition results.

2 Hypothesis Probability Computation

In this section we compute the probability that a given match hypothesis is correct. We compute the probability for one value measured in the image (e.g., the ratio between lengths or angles) to be within a certain range when the actual value (measured in the model) is given. We then extend this technique to joint probability functions for two such values, and use the results to rank match hypotheses.

When measuring the ratio of line lengths in an image or the angle between lines, we want to know from

what viewing directions this image could be scanned given a match between the lines and certain edges in the model. Consider two segments l_1 and l_2 in the image (Figure 1(a)) which are projections of edges $\mathbf{u}_1$ and $\mathbf{u}_2$ respectively in the model, such that the ratio between the length of l_1 and l_2 is ρ and the angle between them is θ.

The ratio between the lengths of the projections of $\mathbf{u}_1$ and $\mathbf{u}_2$ is ρ for viewing directions which satisfy:

$$|\mathbf{u}_1 \times \mathbf{v}| - \rho|\mathbf{u}_2 \times \mathbf{v}| = 0.$$

Viewing directions which satisfy:

$$(\mathbf{u}_1 \times \mathbf{v}) \cdot (\mathbf{u}_2 \times \mathbf{v}) - \cos\theta|\mathbf{u}_1 \times \mathbf{v}||\mathbf{u}_2 \times \mathbf{v}| = 0,$$

yield an angle θ between the projections of $\mathbf{u}_1$ and $\mathbf{u}_2$.

For tracing the curves we use an algorithm for tracing algebraic curves [11] which relies on homotopy continuation [12] to find all curve singularities and construct a discrete representation of the smooth branch curves. In order to insure that the viewing directions lie on the viewing sphere we add the equation $|\mathbf{v}|^2 = 1$ to the equations derived earlier.

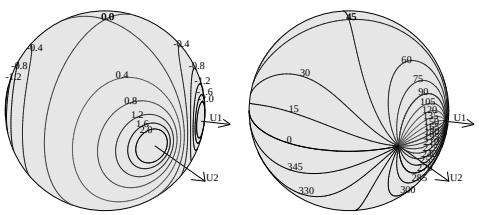

(a) (b)

Figure 2: For two lines of equal length with 45° between them: (a) the viewing sphere with iso-ratio curves with $\log_2(\rho)$ in $[-2, 2]$; (b) the viewing sphere with iso-angle curves.

For reasons of symmetry, ratios are plotted using a $\log_2$ scale. Therefore we traced curves for values of ρ with equal intervals on the $\log_2 \rho$ scale. Figure 2(a) shows curves for values $\log_2 \rho$ between -2 and 2 where the the ratio in the model is 1. At viewing directions parallel to $\pm\mathbf{u}_1$ ($\pm\mathbf{u}_2$) $\mathbf{u}_1$ ($\mathbf{u}_2$) is totally foreshortened and ρ is zero (infinity). When the viewing direction is orthogonal to the plane of $\mathbf{u}_1$ and $\mathbf{u}_2$ the ratio is the "real" ratio.

In Figure 2(b) we show curves for angles between 0 and 360° and the "real" angle is 45°. All the curves go between viewing directions parallel to $\pm\mathbf{u}_1$ and viewing directions parallel to $\pm\mathbf{u}_2$. This occurs because at those directions one of the edges has been foreshortened to zero length and therefore the angle between the two edges is not defined. However, by slightly changing the viewing direction any angle can be obtained. At the viewing direction orthogonal to the $\mathbf{u}_1$,$\mathbf{u}_2$ plane, the angle is the "real" angle.

In order to rank match hypotheses we must be able to compute how likely the values measured in the image are when the match hypothesis is correct. The likelihood of a match hypothesis is measured by the value of the joint p.d.f. for the values measured in the image. This can be done during the recognition process for all the matching hypotheses. However in order to speed up the recognition process we build look-up tables off-line. The $(\log_2 \rho, \theta)$ and (θ_1, θ_2) spaces are divided into rectangles and for each rectangle the average value of the joint p.d.f is computed and stored in the table.

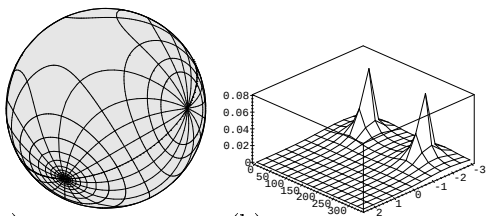

(a) (b)

Figure 3: (a) Tessellation of the viewing sphere into ratio/angle regions; (b) a graph of the joint probability function for ratios and angles as estimated by the area of the regions in (a).

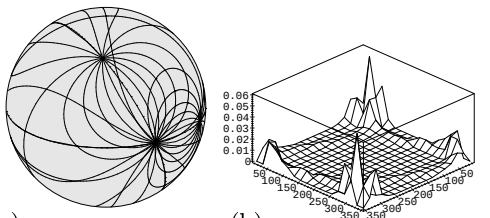

(a) (b)

Figure 4: (a) Tessellation of the viewing sphere into angle/angle regions; (b) a graph of the joint probability function for pairs of angles as estimated by the area of the regions in (a).

This requires computing for two adjacent edges e_1 and e_2 and for values of ρ and θ measured in the image, the probability that $P(\rho_1 < \rho < \rho_2, \theta_1 < \theta < \theta_2)$. We will denote this region of the (ρ, θ) space by $R(\rho_1, \rho_2, \theta_1, \theta_2)$. Using the technique described above we trace the curves of ρ_1, ρ_2, θ_1 and θ_2 and find the curve bounding the region. The discrete points on the curve bound a polygon whose area is computed. To obtain a probability the result is divided by the area of the sphere 4π. Dividing this value by the area of the rectangle in the $(\log_2 \rho, \theta)$ space yields the average joint p.d.f. for ratios and angles within that rectangle.

We have traced the corresponding regions on the viewing sphere of each rectangle, producing a tessellation of the viewing sphere shown in Fig 3(a). The areas of the regions which represent $P((\rho, \theta) \in R(\rho_1, \rho_2, \theta_1, \theta_2))$ are plotted in Figure 3(b). For trihedral corner feature sets curves for the two angles are plotted in Figure 4(a). The areas of the regions which represent $P((\theta_1, \theta_2) \in R(\theta_{11}, \theta_{12}, \theta_{21}, \theta_{22}))$ were plotted in Figure 4(b).

We account for occlusion of the feature sets by other features in the model using the aspect graph

of the object. The regions computed earlier are intersected with the non-critical regions on the viewing sphere where the corners are visible. For regions in which edges in the edge pair feature set are partially occluded, we use a different equation for computing ratios of lengths of the visible parts of the edges, and use it to compute the probabilities.

3 Ranking Match Hypotheses

In this section we use the techniques presented in the previous section for ranking match hypotheses. We build look-up tables of lists of hypotheses sorted by probability off-line and test the sorted hypotheses during the recognition stage.

In the preprocessing stage, look-up tables $T_1(\rho, \theta)$ and $T_2(\theta_1, \theta_2)$ are built for the ratio/angle and angle/angle pairs respectively. The tables are built by computing joint probability distributions which were described in the previous section for all feature sets in the model database.

The probabilities computed are used to determine the induced probabilities on the identity of the object. We denote by E a feature set measured in the image with values within a certain region of the (ρ, θ) or (θ_1, θ_2) space, by M_i the event in which the i^{th} model from the database appears in the image, and by $H_j^{(i)}$ the hypothesis that the j^{th} model feature set of the i^{th} model matches E. In order to compute $P(H_j^{(i)}, M_i|E)$ we use Bayes' law, which yields:

$$P(H_j^{(i)}, M_i|E) = \frac{P(E, H_j^{(i)}|M_i)P(M_i)}{P(E)}.$$

$P(E, H_j^{(i)}|M_i)$ is the probability that was computed in the previous section, and $P(E)$ is the sum of all the probabilities of all the hypotheses in which E can be measured. Therefore,

$$P(H_j^{(i)}, M_i|E) = \frac{P(E, H_j^{(i)}|M_i)P(M_i)}{\sum_{m,j} P(E, H_j^{(m)}|M_m)P(M_m)}. \quad (1)$$

$P(M_i)$ can be determined using prior knowledge we have about the likelihood of a model to appear in the image. $P(H_j^{(i)}, M_i|E)$ is computed for the hypotheses in the list for each entry in the tables and the lists are sorted by probability.

It is important to note that the regions in the (ρ, θ) or (θ_1, θ_2) spaces do not have to be of equal size. Moreover, the recognition algorithm will perform better if regions with multiple probabilistic peaks are divided into smaller regions with one peak in each.

Up until now we have not accounted for the uncertainty in the input to the recognition process. We identify two types of uncertainty which have to be dealt with: certain feature sets which appear in the image will not be recovered by the edge and corner detectors, and the image features might not match any model features.

We denote by $P_d(E, H)$ the probability that a feature set which appears in the image, for which we measure values E, and for which the hypothesis H is correct, will be detected by the edge and corner detectors. For each hypothesis evidence pair $P_d(E, H)$ can be in principle obtained empirically by testing the performance of edge and corner detectors on many typical scenes. In general, feature sets appearing on the silhouette of the object will have a higher $P_d(E, H)$ than internal features. Incorporating $P_d(E, H)$ into (1) yields:

$$P(H_j^{(i)}, M_i|E) = \frac{P(E, H_j^{(i)}|M_i)P(M_i)P_d(E, H_j^{(i)})}{\sum_m \sum_l P(E, H_l^{(m)}, M_m)P_d(E, H_l^{(m)})}.$$

The second type of uncertainty deals with feature sets which are not "legal" feature sets (e.g., features not correctly recovered, features belonging to several objects or features belonging to the background). Computing the average number of "illegal" feature sets P_{il} in a scene depends on the types of scenes which the recognition process analyzes and the quality of the edge and corner detectors. However, P_{il} can be estimated empirically by testing the system on images of typical scenes. In order to incorporate this information into the look-up tables we compute the probability that we measured values within a region in the value space due to random features and multiply that by P_{il}. In each entry in the tables a new hypothesis is added to the list, which represents the "illegal" feature set.

4 Pose Estimation

For all correct hypotheses which match features from an instance of an object to features in the model, the pose recovered for all of these match hypotheses should be the same. Therefore we use pose estimation to find sets of hypotheses which produce the same pose to reinforce the recognition hypothesis. Uncertainty in the values measured in the image induces uncertainty in the pose and are accounted for when testing the compatibility of match hypotheses.

The problem of estimating the pose from three points [1, 8, 9] or a trihedral corner [14] has been extensively studied. We present here a simple approach which deals with both types of feature sets. We regard the pose of an object in weak perspective projection as a combination of the following four components: a viewing direction $\mathbf{v}$ which is a point on the viewing sphere, a rotation of the image by δ degrees about the viewing direction, a scale s, and a translation $\mathbf{t}$ in the image. The projection p_i of a point p of the object in the image is:

$$p_i = sR(\delta)(p \cdot \mathbf{v}_2, p \cdot \mathbf{v}_3) + \mathbf{t},$$

where $\mathbf{v}, \mathbf{v}_2, \mathbf{v}_3$ is an orthonormal basis of $\mathbb{R}^3$ and $R(\delta)$ is a rotation by δ degrees.

Each measured angle or ratio imposes a one-dimensional constraint on possible viewing directions. In order to determine $\mathbf{v}$ two such constraints are needed. Two pairs of types of curves were considered: the ratio/angle pair for pairs of adjacent edges (Figure 1(a)) and the angle/angle pair for trihedral corners (Figure 1(b)). Using the two equations for the two measured values and the constraint $|\mathbf{v}^2| = 1$ we recover the viewing direction component using homotopy continuation. Once the viewing direction has been recovered the other components of the pose are recovered using standard 2D pose estimation. The scale and translation can not be obtained for the angle/angle case. Therefore only in the case of the ratio/angle pair they are recovered.

For each match hypothesis we are able to recover certain components of the pose of the object as described above. However to reinforce these hypotheses we need to find sets of hypotheses matching different feature sets in the image to features of the same object. In that case the values of the components of the pose recovered for each match hypothesis should be the same. When uncertainty in the values measured in the image is accounted for, instead of points in the pose space we get pose uncertainty subsets, and the subsets associated with compatible hypotheses should have a non-empty intersection.

We measure the position of three vertices in the ratio/angle case and the position of four vertices in the angle/angle case. Assuming that each of the measured coordinates has a normal distribution with the measured value as its mean and with a certain variance σ^2 (which can be found empirically), this induces a probability density function on the pose of the object due to these features.

In order to account for uncertainty we have to calculate the probability that the two poses are close enough. This will be discussed in Sect. 5. Here we describe how to eliminate most pairs of poses whose matching probability is negligible by estimating the pose uncertainty subsets and checking if they intersect. We bound the vertex position uncertainty by a circle of radius $k\sigma$. Under that assumption we compute the size of the subset of the pose space which could generate the feature set. If the subsets for both hypotheses have an empty intersection the hypothesis pair is discarded.

We divide the computation into two stages. In the first stage we compute the uncertainty in the values measured in the image such as ratios, angles and lengths due to the uncertainty in vertex position. This stage does not depend on the model features. In the second stage we combine the results of the first stage with the uncertainty due to interrelations between the various components of the pose to produce the estimated pose uncertainty subset.

In the first stage we find the maximum and minimum values for the various values measured in the image assuming the uncertainty is bounded by $\epsilon = k\sigma$. For the ratio/angle case (Figure 5(a)) given vertices

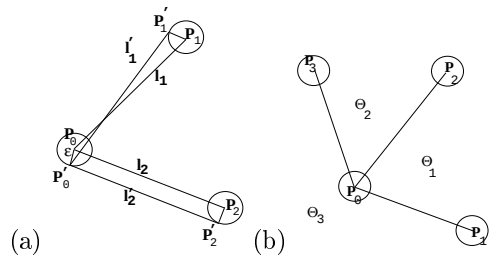

Figure 5: Given vertices P_i in the image, find P_i' within their respective uncertainty regions which yield extreme values for ratios, angles or lengths: (a) the two-edge case; (b) the trihedral corner case.

P_0,P_1 and P_2 we find P_0',P_1' and P_2' within their respective uncertainty regions which yield extreme values for ratios angles or lengths. We assume that the extreme values happen on the boundary of the uncertainty region, and therefore we only need to find three angles ϕ_i such that $P_i' = P_i + (\epsilon\cos\phi_i, \epsilon\sin\phi_i)$.

To demonstrate the general approach we shall deal with the case of finding extreme ratios between the length of the lines l_1' and l_2'. We define $F(\phi_0, \phi_1, \phi_2) = (|l_1'|/|l_2'|)$. The maximum (minimum) is obtained when $\nabla F(\phi_0, \phi_1, \phi_2) = 0$. At first it seems that a three-parameter optimization must be performed to obtain the maximum (minimum) ratio. However when given a value for ϕ_0, the values for ϕ_1 and ϕ_2 can be computed directly. In this case ϕ_i is chosen such that $P_i - P_0'$ is parallel to $P_i' - P_i$ which extends (contracts) l_1 by ϵ and contracts (extends) l_2 by ϵ achieving the maximum (minimum) ratio. Thus we are left with a one-dimensional optimization which can be performed using standard numerical techniques[10]. In Table 1 we summarize how to find all extreme values. For the case of a corner (Figure 5(b)) we measure the minimum and maximum values for $\theta_1, \theta_2, \theta_3 = 2\pi - (\theta_1 + \theta_2)$ and the rotation.

Value	Function	Computing ϕ_1,ϕ_2
Ratio	$\lvert l_1'\rvert/\lvert l_2'\rvert$	$(P_i - P_0') \parallel (P_i' - P_i)$
Angle	$cos^{-1}(\frac{(l_1' \cdot l_2')}{(\lvert l_1'\rvert \lvert l_2'\rvert)})$	$(P_i' - P_i) \perp (P_i' - P_0')$
Scale	$(\frac{\lvert l_1'\rvert}{\lvert l_1\rvert} + \frac{\lvert l_2'\rvert}{\lvert l_2\rvert})/2$	$(P_i - P_0') \parallel (P_i' - P_i)$
Rotation	$Bis^{\perp} \cdot Bis'$ where Bis is the normalized bisector	$(P_i' - P_i) \perp (P_i' - P_0')$
Translation	Maximal value is ϵ	$\phi_i = \phi_0$

Table 1: A table describing the different values which we would like to maximize (minimize), the function which is maximized (minimized) and how to compute ϕ_1 and ϕ_2 when ϕ_0 is given.

In the second stage of the computation, components of the pose subset are computed. When computing the projection of a model point, transformations due to $\mathbf{v}, \alpha, s$ and $\mathbf{t}$ are applied to the point in that order. Therefore, uncertainty in the initial components of the projection can increase the uncertainty of the other components of the pose. The uncertainty of α and s

is affected by the uncertainty in $\mathbf{v}$, and the uncertainty in $\mathbf{t}$ is affected by the uncertainty in $\mathbf{v}, \alpha$ and s.

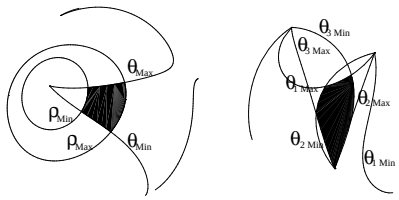

(a) (b)

Figure 6: The uncertainty region on the viewing sphere bounded by curves for the possible minimal and maximal values measured in the image. (a) the two-edge case; (b) the trihedral corner case.

For the two-edge case the region of the viewing sphere which belongs to the subset is the region bounded by the curves for $\rho_{min}, \rho_{max}, \theta_{min}$ and θ_{max} which were computed in the previous step (Figure 6(a)). We find the minimal and maximal values for the position of one of the vertices, the rotation, and the scale of the projected features in that region. Using the extreme values for the rotation α and scale s from both stages, we compute a worst case estimate for them, and compute the translation component by taking into account the effects of all other components of the pose on it.

For the case of the trihedral corner we compute the boundary of the region on the viewing sphere bounded by curves for the minimum and maximum values of $\theta_1, \theta_2, \theta_3$ (Figure 6 (b)) and using that region recover the rotation component of the uncertainty subset.

Although the technique described above gave us insight into the nature of these subsets, a more efficient technique is needed to estimate these subsets at the time of recognition and find the intersection between the subsets. Assuming a match hypothesis is correct the pose p is a function of the vertex positions $\mathbf{a}$ measured in the image. Using a Taylor expansion, the effect of a small uncertainty ξ in $\mathbf{a}$ on the pose can be estimated by:

$$p(\mathbf{a} + \xi) \approx p(\mathbf{a}) + \nabla p(\mathbf{a})\xi.$$

Using this approximation the uncertainty of each component of the pose is estimated as the sum of the contributions due to each component of $\mathbf{a}$. For most components of the pose computing the uncertainty is straightforward, however for the viewing direction component $\mathbf{v}$ of the pose we compute the uncertainty in two directions orthogonal to the unperturbed viewing direction on the viewing sphere.

In order to check if two pose uncertainty subsets have a non-empty intersection, all components of the pose are compared. When one or both of them are angle/angle hypotheses, we estimate the scale and translation components of the pose subset using the vertex position and the uncertainty of the corner of the other hypothesis and then compare the resulting pose subspaces.

5 Ranking Recognition Results

In the final stage of the algorithm, pairs of hypotheses whose pose uncertainty subsets have a non-empty intersection are ranked by probability. Given a set of image features which participate in a match hypothesis (ratio/angle or angle/angle), the p.d.f.'s of the measured values in the image induce a p.d.f. on the pose space.

Given two feature sets in the image e_1 and e_2 and two respective hypotheses h_1 and h_2, we define H as the hypothesis that h_1 and h_2 are true and both match image feature sets to the same instance of a certain model M. We compute $P(h_1, h_2, H|e_1, e_2)$ using Bayes' rule yielding:

$$P(h_1, h_2, H|e_1, e_2) = \frac{P(e_1, e_2|h_1, h_2, H)P(h_1, h_2, H)}{P(e_1, e_2)}. \quad (2)$$

We can write $P(e_1, e_2|h_1, h_2, H)$ as a marginal distribution with respect to the pose p. Therefore

$$P(e_1, e_2|h_1, h_2, H) = \int_p P(e_1, e_2|h_1, h_2, H, p) f_p(p) dp,$$

leading to

$$P(e_1, e_2|h_1, h_2, H) = \int_p f_1(p) f_2(p) f_p(p) dp, \quad (3)$$

where f_1 and f_2 are the multi-normal distribution functions applied to the difference between the back-projected features assuming pose p and the features measured in the image, and $f_p(p)$ is the p.d.f. of poses in the pose space which we assume is uniformly distributed. $P(h_1, h_2, H)$ is the probability that the object M was in the scene, both feature sets were visible and were detected by the edge detector. Therefore,

$$P(h_1, h_2, H) = |S(h_1, h_2)|P(M)P_d(e_1, h_1)P_d(e_2, h_2),$$

where $|S(h_1, h_2)|$ is the area of the viewing sphere in which both model feature sets are visible.

In computing $P(e_1, e_2)$ we sum over every pair of hypotheses h_i, h_j which could generate e_1 and e_2 respectively and whether e_1 and e_2 belong to the same object ($H^{(i,j)}$). When the two feature sets belong to the same object a term like the numerator of (2) is computed, and when they belong to two different objects the product of the terms for the two feature sets is computed.

When computing the various terms in (2), we evaluate (3) and the term for the single feature set, using a Monte-Carlo technique. As the value of the integrand is negligible outside the pose uncertainty subsets, we only have to evaluate it for poses in the pose uncertainty subset for a single feature set, and in the intersection of the subsets for (3). This reduces the number of evaluations which have to be made, speeding up the computation. For single feature sets we estimate its value off-line for typical hypotheses in regions of the (ρ, θ) and (θ_1, θ_2) spaces and store the values in the

look-up tables. For (3) we estimate its value on-line, however that is only performed for the small number of hypothesis pairs whose pose uncertainty subset have a non-empty intersection.

When evaluating (2), terms similar to (3) appear in the numerator and the denominator of the expression. For hypothesis pairs whose pose uncertainty subset do not intersect, this term will be negligible. Therefore, when computing the rank of a hypothesis pair we can assume at first that all the terms of that type except the one for the pair (h_1, h_2) are zero. When we compute the rank of another hypothesis pair h_1' and h_2' which matches the same image feature sets e_1 and e_2, we will add the value we computed for the numerator of (2) to the denominator of the probability of h_1 and h_2. As the hypotheses are traversed in decreasing probability order there is a high probability that the best recognition results will have the highest rank even before all hypotheses have been examined. Therefore, if the recognition process has to be interrupted due to lack of time we can still assume that the correct recognition results will be ranked high on the list.

Another important characteristic of our ranking process is the way "popular" sets of feature sets such as rectangular faces are dealt with. In that case the same feature sets will participate in many hypothesis sets. Therefore their numerators will contribute to the denominators of their probabilities. Thus reducing the ranks of them all. This is reasonable as by using this set of features we can not discriminate between the different hypotheses. Where as a less "popular" but probably correct set of feature sets will have a higher rank, as not many competing hypotheses will exist for those two feature sets.

6 Experimental Recognition Results

In this section we present preliminary recognition results performed on a real image. Figure 7(a) shows the image. Figure 7(b) shows the feature sets extracted from the image. Note that not all features that appear in the image were detected and that some of the feature sets extracted are due to the background and shadows. An important feature of our algorithm is that it works well even with suboptimal results from the initial stages. In Figure 7(c) the objects recognized by the algorithm are shown. Many false recognition candidates with rectangular faces were recovered using features belonging to the rectangular face of the prism. However, these results were discarded because all these results contributed to the denominator of (2), reducing their rank.

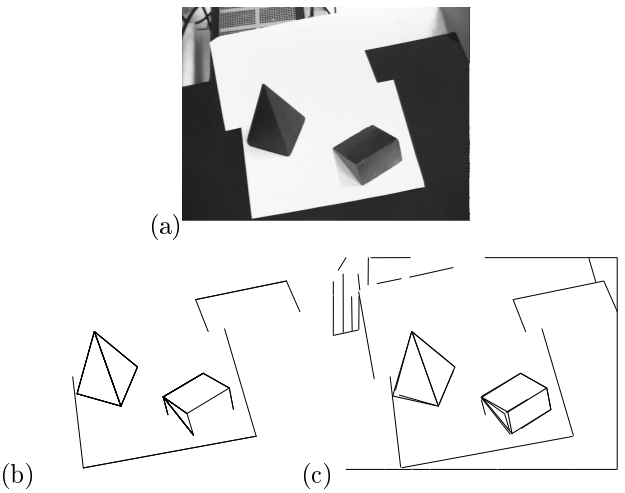

Figure 7:(a) An image of a number of objects; (b) feature sets recovered from the object; (c) recognized objects.

References

[1] T. D. Alter. 3D pose from 3 points using weak-perspective. *PAMI*, 16(8):802–808, August 1994.

[2] T. D. Alter and D. W. Jacobs. Error propagation in full 3D-from-2D object recognition. In *Proc. CVPR*, pages 892–898, Seattle, Washington, 1994.

[3] N. Ayache and O. D. Faugeras. HYPER: a new approach for the recognition and positioning of 2D objects. *PAMI*, 8(1):44–54, January 1986.

[4] J. Ben-Arie. The probabilistic peaking effect of viewed angles and distances with application to 3-D object recognition. *PAMI*, 12(8):760–774, August 1990.

[5] T.O. Binford, T. Levitt, and W. Mann. Bayesian inference in model-based machine vision. In *Workshop on Uncertainty in Artificial Intelligence*, 1987.

[6] J. B. Burns, R. S. Weiss, and E. M. Riseman. View variation of point-set and line-segment features. *PAMI*, 15(1):51–68, January 1993.

[7] W. E. L. Grimson, D. P. Huttenlocher, and T. D. Alter. Recognizing 3D objects from 2D images; an error analysis. In *Proc. CVPR*, pages 316–321, Champaign, Illinois, 1992.

[8] R. M. Haralick, C. Lee, K. Ottenberg, and M. Nölle. Surface reflection: Physical and geometrical perspectives. *IJCV*, 13(3):331–356, December 1994.

[9] D. Huttenlocher and S. Ullman. Recognizing 3D solid objects by alignment with an image. *IJCV*, 5(2):195–212, 1990.

[10] D. Kahaner, C. Moler, and S. Nash. *Numerical Methods and Software*. Prentice Hall, Englewood Cliffs, NJ, 1989.

[11] D.J. Kriegman and J. Ponce. A new curve tracing algorithm and some applications. In P.J. Laurent, A. Le Méhauté, and L.L. Schumaker, editors, *Curves and Surfaces*, pages 267–270. Academic Press, New York, 1991.

[12] A.P. Morgan. *Solving Polynomial Systems using Continuation for Engineering and Scientific Problems*. Prentice Hall, Englewood Cliffs, NJ, 1987.

[13] C. F. Olsen. Fast alignment using probabilistic indexing. In *Proc. CVPR*, pages 387–392, New York, New York, 1993.

[14] Y. Wu, S. S. Iyenger, R. Jain, and S. Bose. A new generalized computational framework for finding object orientation using perspective trihedral angle constraint. *PAMI*, 16(10):961–975, October 1994.