Using symmetry, ellipses and perceptual groups for detecting generic surfaces of revolution in 2D images

Geoff A.W. West and Paul L. Rosin

Cognitive Systems Group, School of Computing,
Curtin University of Technology, Perth,
Western Australia, 6001
email: {geoff,rosin}@cs.curtin.edu.au

ABSTRACT

A number of objects in the real world can be described as surfaces of revolution. These are a particular type of generalised cylinder with a straight axis whose 3D shape is formed by rotating a 2D plane about the axis. Examples of such objects are vases, many chess pieces, light bulbs, table lamps etc. This paper describes a number of techniques that can be used to recognise this class of object in a typical cluttered scene under perspective projection. Use is made of the symmetry of the occluding boundary, perceptual grouping of ellipses, 3D models and the hypothesis that an ellipse is a circle in the real world.

1. INTRODUCTION

A large number of objects in the real world can be described as surfaces of revolution (SORs) e.g. vases, many chess pieces, light bulbs, table lamps, etc. These are a particular form of generalised cylinder, i.e. right, circular, straight, homogeneous generalised cylinders (RCSHGC's)[15]. SORs are an interesting type of object to investigate for recognition as they contain straight and curved edges and surfaces and exhibit symmetry. In addition they are relatively easy to describe as 3D models. Figure 1 shows views of the models used.

A number of suggestions have been made for the recognition of SORs. All use some cue derived from the image to hypothesise a pose, and can be separated into model based and bottom-up approaches.

Dhome[3] proposed two techniques. 1) Detect one ellipse, hypothesise that this matches a circle on the model, and then sequentially check the four possible poses by iteratively aligning the model of the surface of revolution with the occluding boundary in the image using Lowe's technique[8] 2) The detection of zero-curvature points on the boundary define a cone of revolution that again can be used to generate possible object poses.

Ponce *et al*[10] and Kriegman *et al*[7] have developed techniques based on elimination theory to match an analytic representation of the occluding boundary to image contours. This is similar to the least mean squares technique proposed by Lowe[8] for aligning a model with image features in that the result is an accurate measure of the pose of the model in the real world. One possible problem with this method is the requirement for an analytic or implicit description of the SOR.

Several related bottom-up techniques are available for recovering the axis of a SOR. These involve finding the intersection of pairwise combinations of either 1) tangents at all edge points[10], 2) tangents at zeros of curvature along the boundary[10], or 3) bitangents[4]. Since the intersections lie on the symmetry axis the axis can then be found by accumulating the points, and searching for straight lines (e.g. using the Hough Transform).

Perceptual grouping has been proposed for hypothesising the poses of SORs under orthographic projection[14]. Ellipses are hypothesised as being the projections of circles, and their poses are determined. Groupings of ellipses with similar pose parameters suggest potential axes of symmetry. These hypotheses are verified by detecting the occluding boundary, which is symmetric about the axis of the SOR, by matching edge pixels or straight lines on either side of the axis[14].

0-8194-1200-7/93/$4.00

Techniques based on the analysis of the occluding contour of an SOR work for well segmented scenes containing a small number of objects. The majority of the techniques described above use simple scenes containing a single isolated object and no background clutter. The object has no surface markings or texture, and fills the image. However, for more realistic scenes, accurately and robustly extracting the location of zeros of curvature is difficult to achieve. The calculation of curvature is extremely susceptible to noise, requiring substantial smoothing. Smoothing the contours requires large filters that requires in turn long connected boundaries for the objects. Any breakup of the contours (e.g. by a textured background) will reduce the detectability of the zeros of curvature. An important consideration is the requirement for a parameter for the degree of smoothing which determines how many zeros are detected. Detection of bitangents relies on the two tangents being present and on the same contour. If one is missing then the analysis will not work. If the contours are fragmented then an expensive search has to be made for all possible bitangents over all pairs of edge lists.

For any method based on detecting ellipses, a reliable ellipse detection algorithm is required. In real images the ellipse data is often incomplete and noisy, and the parameters of the fitted ellipses are often inaccurate. This would cause Dhome's technique to fail since the hypothesised boundaries would be incorrect. This is because a choice has to be made of which image features (straight segments) match the occluding boundary after which the pose is refined. Choosing the wrong image features will cause the pose to deteriorate. We propose to use a more robust method of determining the best match based on 1) the distance transform and 2) model perturbations based on an estimate of the goodness of the ellipse fit.

All the above techniques can be used to generate tentative hypotheses for the possible presence of SORs in a scene. In some cases an approximate pose is determined. It is proposed to use all forms of bottom-up cues to generate tentative hypotheses and reduce the search space for determining the presence of particular SORs and their poses. This general purpose scheme should work with arbitrary SORs with discontinuities in radius, and cope with occlusion and cluttered scenes. Most of the techniques described in this paper rely on the detection of ellipses in the image. Use is made of the combined ellipse and line detection algorithm[13] which segments edges into perceptually significant combinations of lines and ellipses. Figure 2 shows the ellipses and lines detected for various images. Figure 3 just shows the ellipses.

This paper is divided into two main parts. The first part is concerned with the generation of tentative hypotheses using three methods: symmetry analysis, ellipse grouping and single ellipses. The second part is concerned with verification of the presence of a particular SOR by matching the occluding boundary determined from the model with edge data or the distance transform. A process of refinement can be used to determine the pose accurately.

2. HYPOTHESIS GENERATION

Originally[14] we assumed orthographic projection - this enabled ellipses to be grouped as belonging to a SOR if they had similar tilt and orientation, and their centres lay on a straight line. In many cases orthographic projection is a reasonable approximation to perspective projection, and the advantage of this technique is its simplicity. To allow for more extreme cases where the approximation breaks down we have extended our technique to recover the 3D parameters of circles under perspective projection assuming a known camera geometry. Use is made of Dhome's formulation for the 3D recovery[3]. The ellipses (i.e. 3D circles) can then be grouped as before to detect SORs.

A number of bottom-up cues are proposed to generate initial hypotheses for the presence of SORs in a scene.

- group ellipses and determine axis of symmetry and pose
- for each ellipse determine axis of symmetry and pose
- use a general symmetry finding algorithm

Dhome's formulation for the 3D recovery of the pose of a circle in the real world from an ellipse in the image requires knowledge of the radius of the circle. If there is no knowledge about the SORs that may be present then there is no indication of the 3D pose of each circle because the radius of the circles are unknown. However the tilt

and rotation angles will be correct whereas the coordinates of the centres will be unknown. The angles can be used as a starting point for grouping. Without knowledge of circle radii, the 3D centres of the circles will not be guaranteed to lie on a straight line (the axis of the SOR). However the 2D projection of the centres into the image plane will lie on a straight line in the image as straight lines in 3D map onto straight lines in 2D under perspective projection. The goodness of fit of a straight line to the centres in the image can be used to reinforce the hypothesis. The advantage of the ellipse grouping approach is that even if the parameters of each individual ellipse in the group are inaccurate due to the problems of image noise, clutter, and occlusion, the complete grouping should still provide a reasonably accurate hypothesis.

If insufficient ellipses have been detected to allow grouping then a single ellipse can be used to hypothesise the axis of symmetry. Then possible groupings of boundary fragments are used to more accurately estimate the orientation of the axis of symmetry. Symmetry of the occluding boundary would be a useful cue but under perspective projection this is not truly symmetric[6]. However a coarse analysis for symmetry is useful and can be applied to ellipse groupings as well as single ellipses.

If no ellipses can be found then the third approach is used. Here the axis of symmetry is found by detecting reflectional symmetry in the image. A symmetry finding algorithm[12] is used to hypothesise bilateral symmetry after which the possible occluding contours are determined from the axis. Figure 4 shows the result of analysing an image of the can for bilateral symmetry.

3. HYPOTHESIS VERIFICATION

3.1 Model matching

The above three techniques are all bottom-up in operation and can give tentative evidence for surfaces of revolution. To recognise a particular SOR a model of the SOR is required. A simple model is used consisting of "real" circles, generator circles and curved surface generators which can generate all surfaces of revolution with circular cross section. Real circles are those that may be visible in a scene and represent positions of surface discontinuity parallel to the axis in the SOR. Generator circles and curved surface generators are used to describe the shape of the SOR. They will not be visible in the scene but are necessary to determine the occluding boundary. The model can be transformed and projected to form an image of the model under various projection models. The main features extracted from the model for recognition are the positions of the real circles (transformed to match the ellipses in the image) and the pixels in the occluding boundary which are determined from the generators. The occluding boundary has to be determined after perspective projection of the model because it is not symmetric about the axis of the SOR and so cannot be determined by simple techniques. A number of models of different SORs are used eg. one for each of a number of table lamps. The model description will be extended to include a parametric form that consists of a number of circles that can be of different diameter and separation. The parameterised model will allow a generic type of model to be recognised.

3.2 Initial pose generation

The initial pose generation is obtained from an analysis of the bottom-up hypothesis generation techniques. Each of the techniques produces a range of possible views and reduces the search space for the matching process.

The initial pose(s) are generated by determining the possible 3D transformation parameters (translation: t_x, t_y, t_z) and rotation (τ_y, θ_z). Only two rotation parameters are necessary as the third represents rotation about the axis of the SOR which has no effect. Dhome's formulation[3] is used (see West & Rosin[17]) which requires the radius of the circle to match, the camera parameters (focal length) and the ellipse parameters.

For each grouping of ellipses detected, the image ellipses are matched with each of the model circles. Each match scales the model and the other image ellipses and model circles can now be compared for correct size and position. If a match is found then the occluding boundary of the model for that particular pose is compared with edges from the image. The metric used for the matching is to use a distance transform on the model boundary and to determine the overall match by summing the distances for each pixel hypothesised to be on the occluding

boundary. The fact that the projection of a number of model circles matches a number of image ellipses correctly in 2D space indicates that a good match and good localisation has been obtained.

In our experience, for typical images, the poor quality of the feature extraction because of clutter etc. limits the usefulness of this approach. In addition some ellipses are not visible such as when viewing a lamp from almost directly above. Therefore the approach based on matching a single ellipse has to be considered. For each model, the ellipse could match with each circle (unless other constraints are available such as size of the object) at one of four poses. In addition, the estimation of the axis of symmetry is less accurate because of the reliance on the estimation of the parameters of a single ellipse.

Figure 5 shows the results for the model `lamp.sor` and image `lamp` for the correct circle to ellipse matches and the three other wrong hypotheses. The perspective distortion is apparent in these results. The correct matches are approximately correct and may be refined using Lowe's technique[3] if good segmentation of the features in the image is possible. Good segmentation means the relevant features are extracted from the image. However, this is not possible in many of the images used for this research. Surface markings (`sheen1`) and background clutter (`lamp`) confuse the Marr-Hildreth edge detector into combining edge pixels of the features with other features. This is well known problem caused by the proximity of edges compared with the size of the convolution mask[5].

Having grouped the ellipses and hypothesised the pose of a SOR the hypothesis is verified by detecting the occluding boundary which should be symmetric about the axis of the SOR. We have previously shown how edge pixels and line segments can be grouped about the axis under the assumption of orthographic projection (edges and lines equidistant from the central axis) (Rosin & West[14]). This is only approximately true under perspective projection. In many cases there are a number of possible positions where the boundary can occur due to accidental alignment of edges and lines in the image. In the absence of a model to disambiguate them and determine the true occluding boundary, they are all considered as possible boundaries. If a model is available, then the occluding boundary of the model can be used to search in the image for matching edges. This requires some form of match metric.

3.3 Match metrics

For each hypothesised pose a measure of match is determined by considering other features in the image. Initially the projected occluding boundary of the model is matched against the edge pixels in the image. This can be achieved by the following algorithm:

```
match_metric = 0
set current_minimum = diagonal length of the image
for each pixel in the occluding boundary
      for each edge pixel in the image
            determine the euclidean_distance
            if euclidean_distance < current_minimum then
                  set current_minimum  = euclidean_distance
      match_metric += current_minimum
```

Although feasible, the algorithm is computationally expensive considering the typical values for the number of occluding boundary pixels in the model (e.g. 1000) and edge pixels in the image (e.g. 25,000). Brisdon[2] defined a metric based on using a Gaussian distance measure between model and image features. The search was made perpendicular to the axis of the feature (straight lines in this case). This has advantages over a simple metric although it was decided to try a much simpler approach to reduce the computation time.

Currently a distance transform is used which is computed at the same time as the feature extraction occurs. This inputs a binary edge map and outputs an image in which each pixel is set to the approximate distance to the nearest edge point. The true Euclidean distance could be calculated, but the 3-4 chamfer[1] is accurate enough and is much faster. A similar distance transform incorporating orientation has been used by Mundy *et al*[9] for confirming hypotheses for aircraft on the ground. Figure 6 shows the distance transform for `sheen1` and `lamp`

for 20 iterations - i.e. only distances up to 20 are calculated. A much larger number of iterations could be used but again this slows the computation. Additionally, because all pixels not processed are assigned the value 255 then this can be used as a non-linear penalty on grossly misaligned hypotheses. A problem with the distance transform approach is that even small amounts of noise or clutter will distort the measure, producing small distances to the noise rather than large distances to more significant features. Currently a distance map is generated from a edge image obtained using the Marr-Hildreth for a sigma of 5.0 (feature extraction uses a sigma of 1.0) which removes some noise and clutter while not affecting the main features.

The match metric is determined by summing the distances for each pixel in the occluding boundary. For the correct hypothesis this measure should be the minimum encountered. However a number of factors prevent this. First the measure favours the hypotheses with the lowest number of occluding pixels. So hypotheses that result in small far away objects will be favoured. This can be overcome by determining the average distance per pixel. A second factor is occlusion. Only part of the object may be visible because of occlusion by other objects or because only part of the object is in the image, the rest being outside the image boundary. This can be partly accommodated by further dividing the match metric by the number of pixels visible in the image. Tables 1 and 2 show the results for a number of images: `sheen1` and `lamp`, and models: `sheen.sor` and `lamp.sor`. These are the best ranked hypotheses after matching all the real circles in the models with all the ellipses in the images. Generally the correct hypothesis is one of the top two in the tables. The best hypotheses are shown in fig. 7. These show the original image with the model overlayed (real circles and occluding boundary).

image ellipse	model circle	pose interpr-etation	total boundary (t)	visible pixels (v)	total distance (d)	average distance (d/v)	match (d/vt)
3	1	3	846	846	5345.0	6.32	0.0075
3	3	1	846	846	5448.0	6.44	0.0076
2	1	1	1228	131	1588.0	12.12	0.0099
2	1	2	1162	132	1600.0	12.12	0.0104
2	3	4	1162	131	1612.0	12.31	0.0106
2	3	3	1228	133	1804.0	13.56	0.0110
5	3	1	992	277	3725.0	13.45	0.0136
5	1	3	992	278	3792.0	13.64	0.0138
5	3	2	896	280	3771.0	13.47	0.0150
5	1	4	896	278	3881.0	13.96	0.0156

Table 1. Best ten hypotheses for `sheen.sor` and `sheen1`.

image ellipse	model circle	pose interpr-etation	total boundary (t)	visible pixels (v)	total distance (d)	average distance (d/v)	match (d/vt)
9	11	1	1422	1422	9864.0	6.94	0.0049
1	10	1	1372	1372	10382.0	7.57	0.0055
11	10	3	266	266	1023.0	3.85	0.0145
26	3	3	638	638	5934.0	9.30	0.0146
10	1	3	1378	1211	25065.0	20.70	0.0150
11	10	1	280	280	1432.0	5.11	0.0183
26	1	3	466	466	4062.0	8.72	0.0187
22	11	1	506	506	4987.0	9.86	0.0195
23	11	3	526	432	4445.0	10.29	0.0196
15	10	3	194	194	740.0	3.81	0.0197

Table 2. Best ten hypotheses for `lamp.sor` and `lamp`.

More complex images have been considered. `Lamp32` contains a simple table lamp in a more cluttered scene. Other objects are present and the background consists of a tartan table cloth which is heavily patterned. This has

a similar effect to the surface markings of the sheen images in that features describing the edges of the SOR are not detected. An additional problem with these images is the distance transform for which a reduced descriminatory performance is obtained because there are many edge pixels and for many wrong matches a good match metric is still obtained. For lamp32, the best ranked hypotheses are tabulated in table 3.

rank	image ellipse	model circle	pose interpr-etation	total boundary (t)	visible pixels (v)	total distance (d)	average distance (d/v)	match (d/vt)
1	6	11	1	1871	159	926.0	5.82	0.0031
2	10	12	1	1496	80	466.0	5.82	0.0039
3	6	11	2	1824	158	1125.0	7.12	0.0039
4	10	12	2	1302	80	458.0	5.72	0.0044
5	14	11	1	1286	190	1097.0	5.77	0.0045
6	6	12	2	1823	138	1153.0	8.36	0.0046
7	6	12	1	1858	136	1249.0	9.18	0.0049
...								
...								
88	9	1	3	1012	1012	12987.0	12.83	0.0127

Table 3. Best eight and the correct hypotheses for lamp3.sor and lamp32.

12669	12308	11523	9724	9381	9541	10808	11667	11969	12320	12777
12624	12289	11240	9763	9519	10183	11329	11934	12182	12524	12737
12711	12362	11409	10025	9938	10694	11671	12127	12360	12693	12772
12750	12460	11843	10924	10883	11532	12359	12552	12588	12704	12754
12713	12717	12003	11468	11808	12706	13064	13051	12868	12884	12980
12938	13121	12314	11554	12099	12987	13605	13548	13321	13052	13205
13250	13234	12591	11890	12230	13409	13892	13807	13719	13101	13150
13617	13509	12948	12231	12363	13543	14287	14217	13936	13269	13172
13753	13737	13206	12539	12488	13598	14449	14493	14238	13616	13282
13993	14124	13258	12543	12778	14077	14769	14780	14427	13787	13293
14301	14220	13550	12532	12778	13882	14784	14864	14537	13897	13379

Table 4. Match metric for variation in rotation (x axis) and tilt (y axis). See fig. 8 for graphical representation

The best two matches and the correct interpretation are shown in fig. 8. The two best matches are matching the bottom of the lamp with spurious ellipses in the tartan pattern of the table cloth. The correct match is ranked 88 out of the possible 540 matches. One reason for the poor result is the poor estimation of the model pose from the single circle to ellipse match. Perturbating the pose by small changes in tilt and rotation results in a well defined valley in match metric space indicating good correlation, see table 4 and figure 9. By contrast perturbating the poses for the first two ranked matches in the table results in no significant change in the match metric indicating that good correlation is not being obtained. Further results for recognition can be found in West[16].

4. FURTHER WORK

It is proposed to investigate better match metrics. A multi-scale approach is favoured. Ideally, the distance measure at a point should be calculated for all features in the image and an average of the distances used, weighting each distance by the scale lifetime of the generating feature. This would be extremely computationally expensive, and instead we use an approximation. The distance transform is applied to a series of edge maps generated at different scales, and a simple average is taken.

A number of other distance transforms have been considered. Non-linear transforms such as that obtained by convolving a gaussian with the edge image, and logarithmic distance can be used to increasingly weight close matches. Poor matches would be penalised more. A distance transform that is non-linear is a binary distance transform that has a value of 0 for an image pixel within a distance *d* of an edge and 255 otherwise. This would allow a certain amount of misalignment before a penalty is incurred. The performance of these match metrics is the subject of a future paper (West & Rosin[18]).

With real scenes instances of the model may be hypothesised in cluttered areas. Even if the hypothesis is false most of the occluding boundary will be close to some edge point. This can be discriminated from correct hypotheses by taking into account the orientation of the edges. When the edge map is generated an orientation map is also calculated. We do this by applying a Sobel operator to the intensity image, and then smoothing the orientation map at the same scale as the edge detection. When the distance transform is calculated and minimum distances are propagated the corresponding orientations are also propagated in a separate image.

The need to perturb the model to get a better match because of pose estimation errors is a problem with many object recognition paradigms. However, in the absence of high level features to match it appears that its use is inevitable. It is hoped that better match metrics may remove the need for perturbation. The amount of perturbation needed depends on the accuracy of ellipse detection. Using an ellipse detector[11] that gives an estimate of the accuracy of the ellipse fit will further reduce the amount of perturbation required.

5. SUMMARY

This paper has discussed techniques for the recognition of surfaces of revolution (SORs) in single grey scale images obtained using perspective projection. A number of bottom-up techniques have been proposed for the tentative generation of hypotheses of the pose of SORs in 3D scenes. These would be used in combination with other techniques to produce a comprehensive hypothesis generation process that would work on images containing numerous shapes of SOR.

For each hypothesis a model-based matching approach similar to Dhome's[3] is proposed. The occluding boundary of the projected model is matched with the image edge data using a distance transform. A distance transform is used so that no reliance is made on detecting higher order features in the image.

The technique has been shown to work for a number of scenes containing SORs and other objects and clutter. However it fails on numerous images. Perturbating the model to account for errors in initial pose estimation and multi-scale distance transforms are being investigated to improve the performance.

6. REFERENCES

1. Barrow H.G., Tenenbaum J.M., Bolles R.C., Wolfe H.C., "Parametric Correspondence and Chamfer: Two New Techniques for Image Matching", *Proc. IJCAI*, pp. 650-663, 1977.

2. Brisdon, K., "Alvey MMI-007 Vehicle Exemplar: Evaluation and Verification of Model Instances", *Proc. 3rd Alvey Vision Conference*, Cambridge, UK, September, pp. 33-37, 1987.

3. Dhome, M., Lapreste, J.T., Rives, G. and Richetin, M., "Spatial localisation of modelled objects of revolution in monocular perspective vision", *Proc. 1st ECCV'*, Antibes, France, pp. 475-485, 1990.

4. Forsyth D., Mundy, J.L., Zisserman, A. and Rothwell, C., "Recognising rotationally symmetric surfaces from their outlines", *Proc. 2nd. ECCV*, Santa Margherita Ligure, Italy, pp. 639-647, 1992.

5. Heurtas, A. andMedioni, G., "Detection of intensity changes with subpixel accuracy using laplacian-gaussian masks", *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 8, No. 5, pp. 651-663, 1986.

6. Koenderink J.J., "What does the occluding contour tell us about solid shape?", *Perception*, Vol. 13, pp 321-330, 1984.

7. Kriegman, D.J. and Ponce, J., "On recognisiong and positioning curved 3D objects from image contours", *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 12, No. 12, pp. 1127- 1137, 1990.

8. Lowe, D.G., *Perceptual organisation and visual recognition*, Kluwer Academic Publishers, Mass. USA, 1985.

9. Mundy, J.L. and Heller, A.J., "The evolution and testing of model-based object recognition system", *Proc. IEEE International Conference on Computer Vision*, Osaka, Japan, December, pp. 268-282, 1990.

9. Ponce, J., Chelberg, D. and Mann, W., "Invariant properties of straight homogeneous generalised cylinders and their contours", *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 11, No. 9, pp. 951-966, 1989.

10. Ponce,J., Hoogs, A. and Kriegman, D.J., "On using CAD models to compute the pose of curved 3D objects", *Proc. IEEE Workshop on Directions in Automated CAD-based Vision*, pp. 136-145, Maui, USA, 1991.

11. Porrill, J., "Fitting ellipses and predicting confidence envelopes using a bias corrected kalman filter", *Proc. 5th. Alvey Vision Conference*, Reading, UK, September, pp. 175-180, 1989.

12. Reisfeld, D., Wolfson, H. and Yeshurun, Y. "Detection of interest points using symmetry", *Proc. 3rd IEEE ICCV*, Osaka, Japan, pp. 62-65, 1990.

13. Rosin, P.L. and West, G.A W. "Segmenting curves into elliptic arcs and straight lines", *Proc. IEEE International Conference on Computer Vision*, Osaka, Japan, December, pp. 75-78, 1990.

14. Rosin, P.L., and West, G.A.W., "Detection and verification of surfaces of revolution by perceptual grouping", *Pattern Recognition Letters*, vol. 13, pp. 53-461, 1992.

15. Shafer, S.A. , *Shadows and silhouettes in computer vision*, Kluwer, Boston, USA, 1985.

16. West, G.A.W., "Model-based recognition of surfaces of revolution using elliptical cues for initial pose estimation", Technical Report 1993/2, School of Computing, Curtin University, Perth, WA, 1993.

17. West, G.A.W. and Rosin, P.L., "Generating object hypotheses by perceptual grouping of ellipses", Technical Report 1992/10, School of Computing, Curtin University, Perth, WA, 1992.

18. West, G.A.W. and Rosin, P.L., "An analysis of the performance of a number of match metrics for 3D object recognition", Technical Report 1993/3, School of Computing, Curtin University, Perth, WA, 1993.

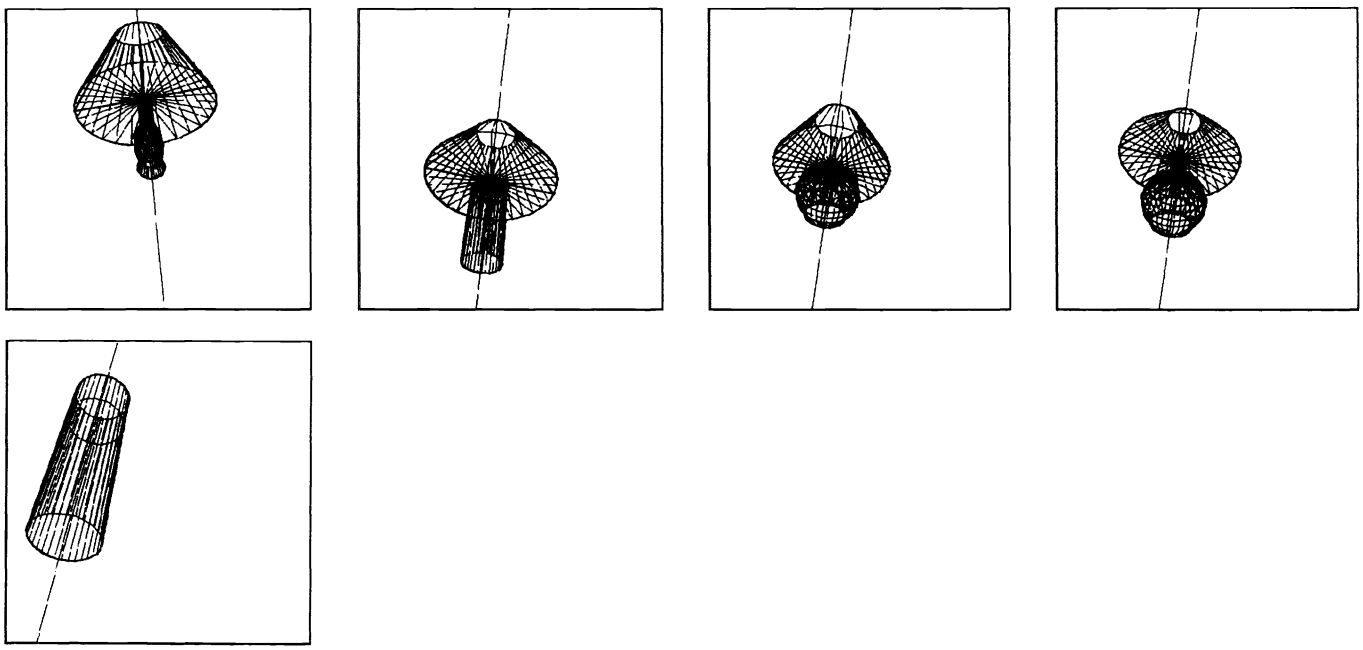

Figure 1. The five different SORs used: `lamp.sor`, `lamp2.sor`, `lamp3.sor`, `lamp4.sor` and `sheen.sor`, showing generators and plotted with perspective projection.

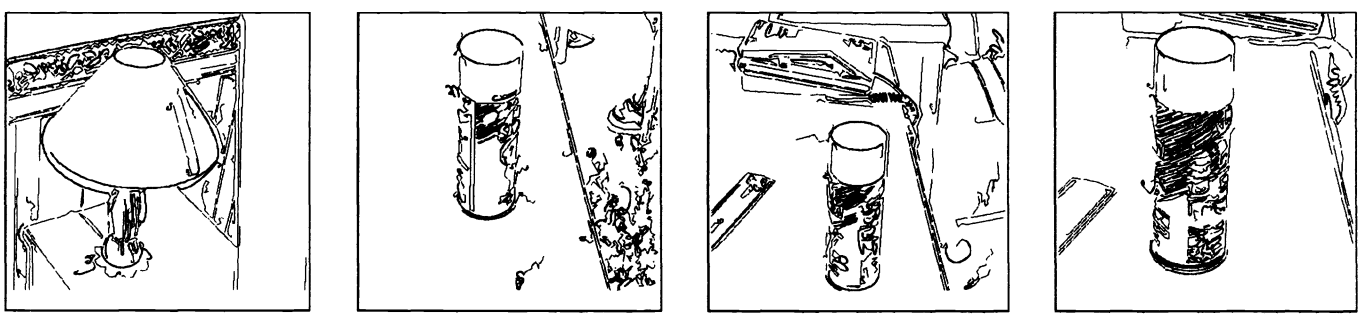

Figure 2. Ellipses and lines detected for the objects `lamp.sor` and `sheen.sor` for one view of the lamp and three views of the can.

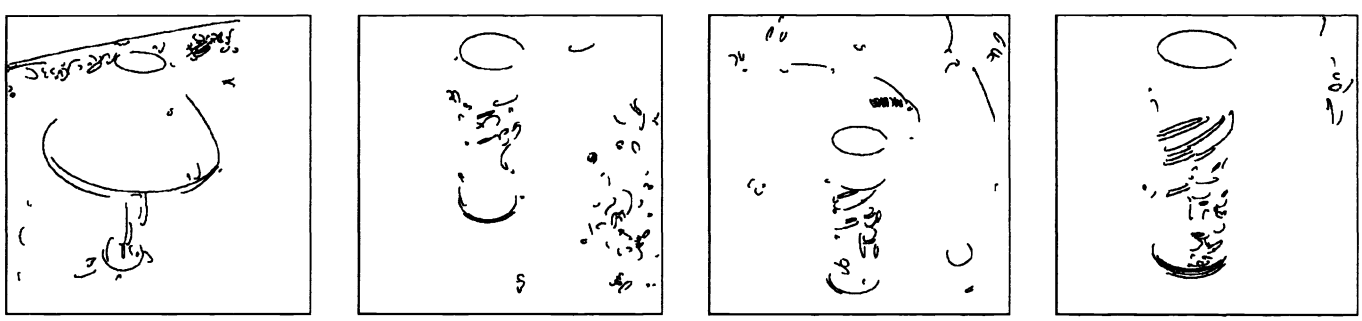

Figure 3. Ellipses detected for the objects `lamp.sor` and `sheen.sor` for one view of the lamp and three views of the can.

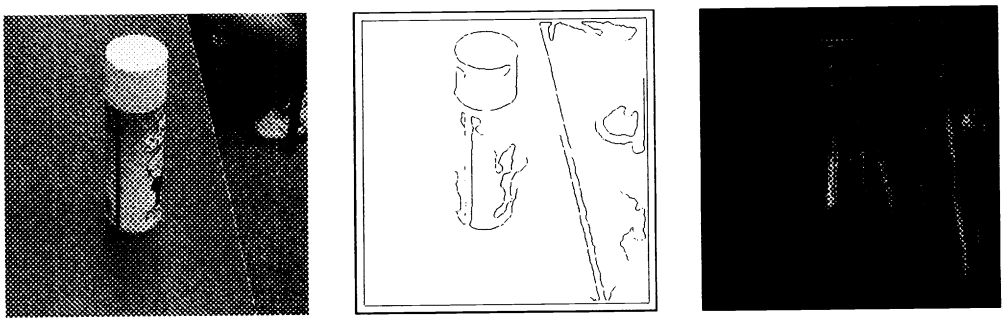

Figure 4. The can image `sheen1` , significant edges and the map of symmetric pixels indicating possible bilateral symmetry.

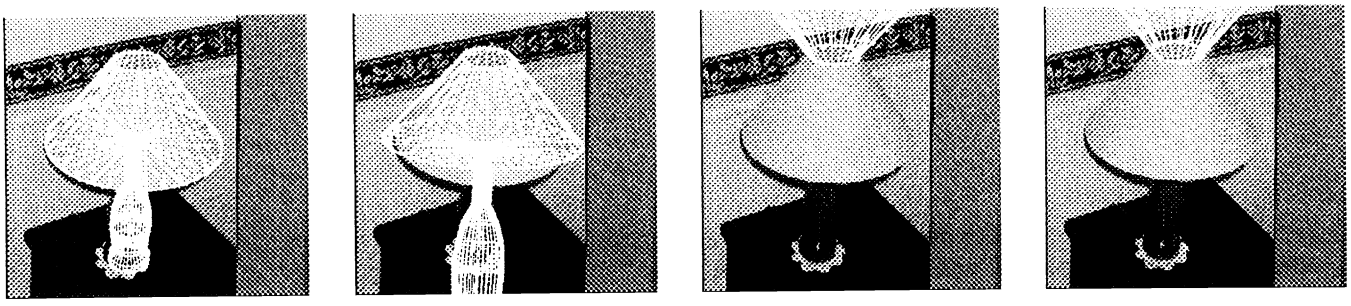

Figure 5. The four possible poses for a particular model circle to image ellipse match for `lamp.sor` and `lamp`

Figure 6. The distance transforms for the can and lamp. A sigma of 5.0 was used for the can and 1.0 for the lamp.

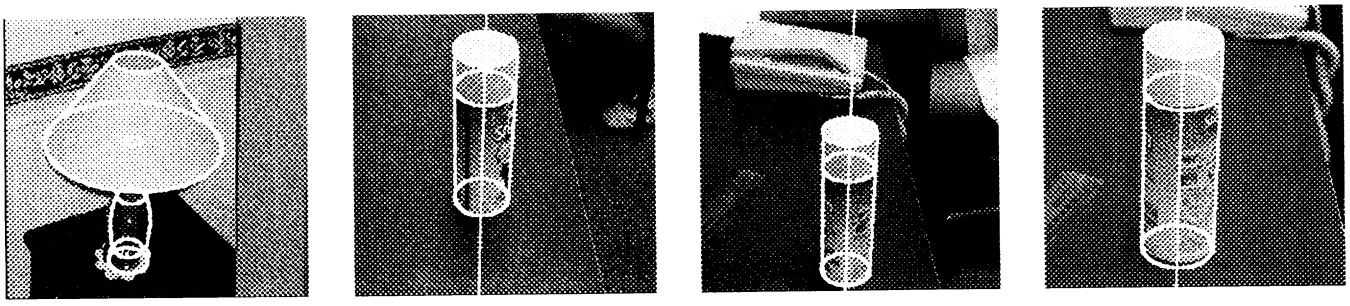

Figure 7. Correct matches for the `lamp` and three views of the can (`sheen1` - 3) showing the models superimposed on the original images.

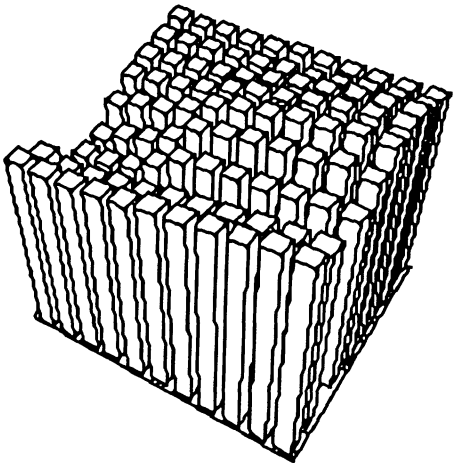

Figure 9. Variation of the match metric for the correct interpretation of lamp3.sor and image lamp32 for variation in pose by small values of tilt and rotation angle (see table 4).

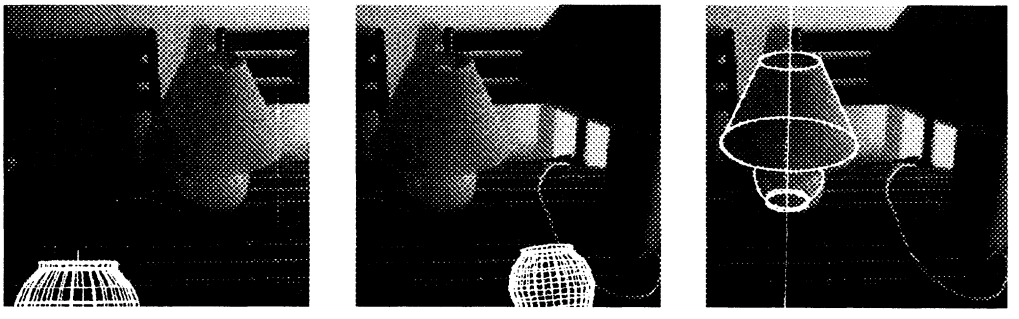

Figure 8. The two best matches and the correct match for `lamp3.sor` and image `lamp32`. Note the slight error in pose for the correct match.