3D Symmetry Detection Using The Extended Gaussian Image

Changming Sun and Jamie Sherrah

Abstract—**Symmetry detection is important in the area of computer vision. A 3-D symmetry detection algorithm is presented in this correspondence. The symmetry detection problem is converted to the correlation of the Gaussian image. Once the Gaussian image of the object has been obtained, the algorithm is independent of the input format. The algorithm can handle different kinds of images or objects. Simulated and real images have been tested in a variety of formats, and the results show that the symmetry can be determined using the Gaussian image.**

Keywords— **Symmetry detection, reflectional and rotational symmetry, extended Gaussian image, sphere tessellation, orientation histogram, principal axis.**

I. Introduction

Many objects around us exhibit some form of symmetry. Symmetry is a powerful concept which facilitates object detection and recognition in many situations. For instance, symmetry information is useful in robotics for recognition, inspection, grasping, and reasoning. Symmetry may be defined in terms of three transformations in n-dimensional Euclidean space E^n: reflection, rotation and translation. Formally, a subset S of E^n is symmetric with respect to a transformation T if $T(S) = S$. Reflectional symmetry has a reflection plane, for which the left half-space is a mirror image of the right half-space. Rotational symmetry has an symmetry axis and an *order* of symmetry f ($f \geq 2$). After an object with order f symmetry has been rotated about its symmetry axis by $(m \times 360/f)^\circ$ ($1 \leq m < f$), it is indistinguishable from its original orientation.

Most of the work carried out on 3-D symmetry detection has been based on edge, contour or point set information. Alt *et al* [1] presented algorithms for computing exact or approximate congruences and symmetries of point sets. Wolter *et al* [2] described exact algorithms for detecting all rotational and involutional symmetries in point sets, polygons and polyhedra. Zabrodsky *et al* [3] defined a Continuous Symmetry Measure to quantify the symmetry of objects. They also presented a multiresolution scheme that hierarchically detects symmetric and almost symmetric patterns [4]. Jiang and Bunke presented an algorithm for determining rotational symmetries of polyhedral objects [5]. Parry-Barwick and Bowyer [6] developed methods that can detect both hierarchical and partial symmetry of two-dimensional set-theoretic models with components constructed from a few straight edges or polynomials. This method has the disadvantage of being computationally intensive.

Ishikawa *et al* [7] presented an interactive technique for measuring the symmetry degree of an arbitrary three-dimensional object. This technique works on 3-D dense CT or MRI data after the octree representation has been obtained. Minovic *et al* [8] used an octree representation to find the symmetry of a 3-D object, and showed the results for synthetic objects. These octree methods are similar to the work presented in this correspondence. The advantage of the approach taken here is that it works for a range of image types, including real images rather than just synthetic objects, and is therefore quite flexible.

C. Sun is with CSIRO Division of Mathematics and Statistics, Locked Bag 17, North Ryde, NSW 2113, Australia; e-mail: changming.sun@cmis.csiro.au.

J. Sherrah is with Department of Electrical and Electronic Engineering, The University of Adelaide, SA 5005, Australia; e-mail: jsherrah@eleceng.adelaide.edu.au.

The extended Gaussian image (EGI) is useful for such tasks as recognition [9] and determining the orientation of an object in space [10], [11]. Brown [12] used the 3-D histogram to study the orientation of dendritic fields. Kang and Ikeuchi [13] used the complex EGI for 3-D pose determination. But none has used EGI for symmetry detection yet. In this article, we investigate the use of the extended Gaussian image for detecting symmetry in 3-D objects. The method of symmetry detection used is based on our observation that in many cases, if an object is symmetric, then so is its EGI. The EGI is sampled at regular surface patches to give the object's orientation histogram, a discrete form of the extended Gaussian image. After the orientation histogram has been obtained, it is tested for symmetry without further consideration of the original image. The object can be represented in different formats: CT or MRI images; Range images; 2-D images; 3-D wire-frame models; and point data sets.

The aim of this correspondence is to determine: (1) the position and orientation of the plane of reflectional symmetry; and (2) the axis and order of rotational symmetry for images in a variety of formats. Note that it is not an objective to determine whether or not an image is symmetric. It is assumed that the images under consideration possess some degree of symmetry, although it need not be perfect. Not all symmetries are required to be found, but only the symmetry which is strongest in the object. For rotational symmetry, the highest order is sought.

Section II describes a method for tessellating a sphere and discusses how these tessellations are used to obtain the orientation histogram. Section III details the method for obtaining the principal axes, used to reduce the symmetry search space. Section IV describes the algorithms for finding the reflectional symmetry and the rotational symmetry parameters. Section V shows the results of the algorithm on simulated and real images. The conclusions drawn from the work are found in Section VI.

II. Sphere Tessellation and the Orientation Histogram

In order to construct and view the 3-D orientation histogram used for symmetry detection, it is necessary to tessellate the surface of a unit sphere. The tessellation is constructed from an icosahedron. The steps followed to construct the tessellation are: (a) Subdivide each triangular face of the icosahedron into four smaller triangles by joining the mid-points of the three sides; (b) Normalise the new triangle vertices onto the unit sphere; (c) Repeat these steps until the desired resolution is reached; and (d) Convert the triangular facets into hexagonal facets to make the cells rounded. Each vertex of the triangular tessellation has six triangles connected to it, except for the twelve original vertices of the icosahedron, which have five. The hexagonal facets are formed by joining the centres of the connected triangles. The tessellation chosen is not optimal, but struck a satisfactory compromise between the criteria mentioned in [10].

To perform useful analysis using the tessellated sphere, a unique labelling system is required for the triangular facets. The method used for numbering the triangular facets is explained in detail in [14]. The task of finding which facet (or *bin*) a given vector lies in has to be performed many times during symmetry detection, so an efficient algorithm is required. Rather than searching each bin direction and choosing the direction with the maximum dot product, a hierarchical approach is taken.

The Extended Gaussian image is a function defined on the face of the unit sphere, with the value in the direction of the unit normal $\mathbf{n}$ given by $K(\mathbf{n})$. The method for transforming an

object to its EGI depends on the object's image format. Range images can be differentiated to find the gradient vector at each pixel, and this direction is used as the normal for the pixel. Note that in an intensity image, the grey level does not have the same dimensionality as the x-y grid. We will treat an intensity image the same as a range image to get the normals. Point data can be represented, with each point direction providing the corresponding orientation. For 3-D data sets from CT or MRI, the grey-level gradients can be used to obtain the EGI. The extended Gaussian image can also be obtained from wire-frame objects, such as those in CAD models. Each wire-frame facet has a surface orientation which can be weighted by the area of the face. Some information is discarded in the formation of the extended Gaussian image, since positional information is not taken into account. Nevertheless, for convex objects the representation is unique [11].

The orientation histogram is a discrete version of the extended Gaussian image. Each point on the Gaussian sphere lies in a particular facet of the histogram. The orientation histogram is obtained by adding each non-zero value of $K(\mathbf{n})$ to the histogram bin in which the vector $\mathbf{n}$ lies. The bins used are the facets of the hexagonal tessellation. The resulting image is like a spiky ball, as shown in Figure 1.

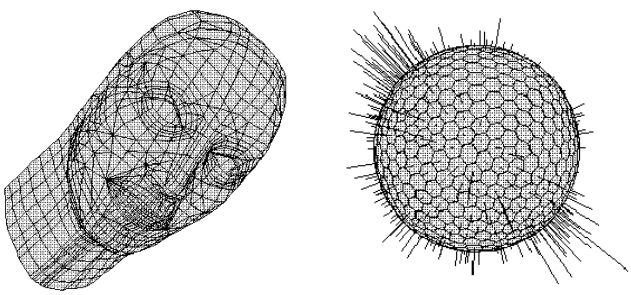

(a) (b)

Fig. 1. Example of an orientation histogram for a wire-frame object: (a) wire-frame model of a human head; (b) the corresponding orientation histogram.

III. Reduce Search Space

In order to determine the vector associated with the strongest degree of symmetry, the obvious solution is to test every bin direction on the histogram. The number of bins on the histogram can be quite large, however, so this approach is computationally prohibitive. The search time can be halved by noting that only one hemisphere of the histogram needs to be searched. For high resolution tessellations, though, this can still be very slow. The method used to narrow down the search for the symmetry axis is to check only the directions of the principal axes of the object, and their 5 or 6 neighbouring bins. The motivation for this approach is based on two theorems found in [15], which state that:

- *Any plane of symmetry of a body is perpendicular to a principal axis.*
- *Any axis of symmetry of a body is a principal axis.*

The inertia matrix J of an object can be formed according to the following equation:

$$\mathsf{J} = \mathrm{tr}(\mathsf{C})\mathsf{I} - \mathsf{C} = \begin{bmatrix} \mu_{020}+\mu_{002} & -\mu_{110} & -\mu_{101} \\ -\mu_{110} & \mu_{200}+\mu_{002} & -\mu_{011} \\ -\mu_{101} & -\mu_{011} & \mu_{200}+\mu_{020} \end{bmatrix}$$

where $\mathrm{tr}(\cdot)$ denotes the trace of a matrix and I is the 3×3 identity matrix. C is the covariance matrix of an object which is derived from the second-order central moments. The three eigenvectors of J are the principal axes, and the eigenvalues vary inversely with the variance of the object along the corresponding axes.

IV. Symmetry Detection

Detection of symmetry in an object is based on our observation that, in most practical cases, an object's orientation histogram exhibits the same symmetry as the object. The degree of symmetry in the histogram is measured by the correlation of the histogram with itself after some transformation. The histogram is searched for the bin direction which is referential to the largest degree of symmetry. For this correspondence, two programs were written to perform symmetry detection: the first program finds the position of and normal to the plane of reflectional symmetry of the input image, and the second program determines the direction and position of the object's axis of rotational symmetry, along with the detected order of symmetry. The steps performed during both forms of symmetry detection are:

1. Read image and calculate the principal axes and centre of mass;
2. Build the orientation histogram from the image, and smooth the histogram;
3. Search near the principal axis directions for the strongest symmetry by performing correlation operations on the histogram; and
4. Present the symmetry data in a useful format.

Steps 3 and 4 require different processing and presentation methods for reflectional and rotational symmetry detection, which are discussed separately below.

A. Reflectional Symmetry

The degree of reflectional symmetry in a plane is measured by the correlation of the histogram with itself after reflection in the plane. The reflected (or mirrored) vector $\mathbf{v}_{\mathsf{mir}}$ of $\mathbf{v}$ w.r.t plane $\mathbf{n}$ can be described as:

$$\mathbf{v}_{\mathsf{mir}} = \mathsf{mir}(\mathbf{n})\mathbf{v} \quad \text{and} \quad \mathsf{mir}(\mathbf{n}) = \mathsf{I} - 2(\mathbf{n}\otimes\mathbf{n}) \tag{1}$$

where $\otimes$ denotes the tensor product of two vectors.

The correlation is formed by visiting each histogram bin and multiplying its value by the value in the bin which mirrors it in the plane. Each resulting product is added to the correlation. Only the three principal axis directions and their 5 or 6 neighbours are considered as candidates for the symmetry plane normal. The plane of strongest reflectional symmetry is that which yields the highest histogram correlation.

B. Rotational Symmetry

We denote the rotation around vector $\mathbf{n}$ through an angle θ by $\mathsf{rot}(\mathbf{n},\theta)$. The rotated vector $\mathbf{v}_{\mathsf{rot}}$ of $\mathbf{v}$ can be obtained by:

$$\mathbf{v}_{\mathsf{rot}} = \mathsf{rot}(\mathbf{n},\theta)\mathbf{v} \quad \text{and} \quad \mathsf{rot}(\mathbf{n},\theta) = e^{\mathsf{H}} = \mathsf{I} + \frac{sin\theta}{\theta}\mathsf{H} + \frac{1-cos\theta}{\theta^2}\mathsf{H}^2 \tag{2}$$

where H is a skew-symmetric matrix composed of the scaled elements of the rotation vector. The scale factor is equal to the rotation angle.

The *order* of rotational symmetry is the number of unique rotations of the object about the symmetry axis which bring it into coincidence with itself. The corresponding *symmetry angle* is the minimum angle of rotation required to bring the object into coincidence with itself. For each order of symmetry under consideration, the histogram is rotated by the corresponding symmetry angle and correlated with the original. The correlation can be performed for each unique rotation, and averaged to

reduce sensitivity to noise. The order which results in a peak in the correlation function is the strongest order of symmetry for the given axis. This is performed for all the axes under consideration, and the largest correlation from these axes indicates the axis of rotational symmetry. Since most rotationally-symmetric objects have a modestly-sized order of symmetry, only orders 2, 3, $\cdots$, 12 were considered (order 12 corresponds to symmetry angle of 30°).

V. Experimental Results

A. Reflectional Symmetry Results

Reflectional symmetry detection was performed on a number of images in a variety of formats using a Sun Sparc10. Figures 2 and 3 show the results of the symmetry detection algorithm on simulated objects. Figure 2(a) shows a wire-frame mushroom, annotated with the detected normal and plane of symmetry. Figure 2(b) is the mushroom's orientation histogram. The mushroom model is not perfectly symmetric, yet the algorithm found the most likely orientation of the normal to the plane of symmetry. The symmetry planes and normals found for other wire-frame objects are displayed in Figure 3.

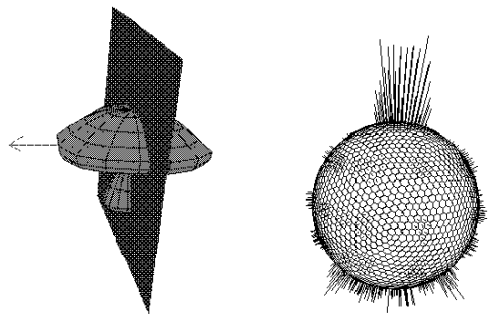

(a) (b)

Fig. 2. Symmetry detection results for simulated objects. (a) mushroom model consisting of quadrilateral and triangular patches; (b) the corresponding orientation histogram.

A number of 2-D grey-scale images were used to test the program, and the results are shown in Figure 4. Each image has been annotated with a line drawn from the centre of mass in the direction of the normal to the plane of symmetry. For more complex images, the intensity was not a meaningful descriptor of shape, and the principle axes found were intuitively incorrect.

Figure 5 shows several range images with the normal to the plane of symmetry overlaid. Figure 5(d) is a noise added image of Figure 5(c). For these two images, the detected normals of the symmetry planes are the same.

Samples of independent observations of unit vectors, taken from [16], were used as point data. The data set consists of 155 measurements of facing directions of conically folded bedding planes, originally provided in (plunge, plunge-azimuth) format. The program took 49.91 secs to complete, and the results are displayed in Figure 6(a).

The CT data set used is a 256×256×31 image of human ribs. The plane of symmetry was found to divide the rib-cage, passing through the spine. The program took 18.86 mins to complete on a Sun Sparc10. The long computation time is due to the Sobel operation on the large data set. The rendering package **Explorer** was used on a Silicon Graphics machine to display the CT image and its plane of reflectional symmetry, as shown in Figure 6(b).

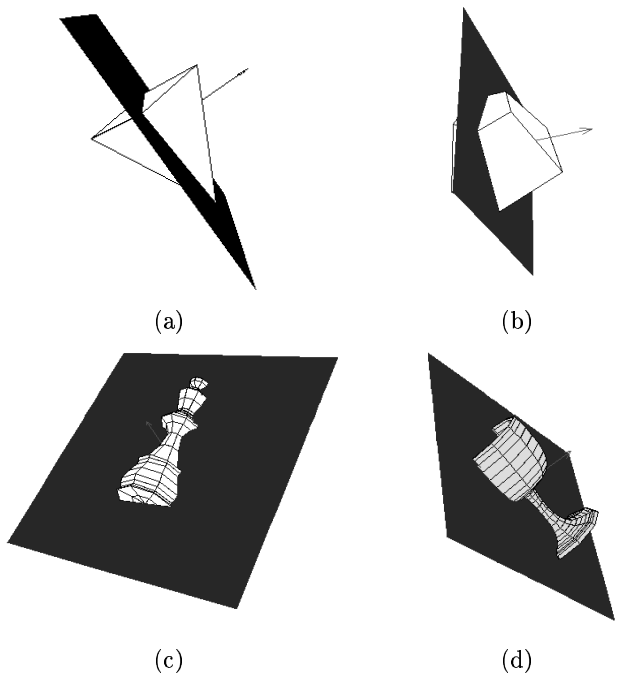

(a) (b)

(c) (d)

Fig. 3. Display of reflectional symmetry detection results for wire-frame objects.

B. Rotational Symmetry Results

The rotational symmetry detection program was tested using a number of images on a Sun Sparc10. The symmetry axes and rays detected for wire-frame objects are displayed in Figure 7. Figure 7(c) shows a king chess piece with order 8 symmetry, denoted by the rays drawn emanating from the object's centre of mass. The long vector is the axis of rotational symmetry. Figure 7(d) shows a goblet with order 10 symmetry, so that if the object is rotated by 36° about the axis, it is indistinguishable from its original orientation.

A number of 2-D images were used to test the program, and the results are shown in Figure 8. Figure 8(a) is a triangle with order 3 symmetry about the axis directed out of the page, and Figure 8(c) is an order 5 symmetric pentagon whose axis of symmetry also points out of the page. The images are annotated with rays drawn from the centre of mass which are equal in number to the order of symmetry and separated by the symmetry angle. Figure 9 shows the rotational symmetry detection result for a range image.

C. Discussion

Principal component analysis alone cannot obtain the order of 3-D rotational symmetries; it needs to be followed by another search or analysis to obtain the order of symmetry. Because of the effect of digitisation on the orientation histogram, and the possibly imperfect symmetry of the input object, a search in the neighbouring directions of the principal axes is necessary.

Burel and Henocq described a method for estimation of the orientation of 3-D objects using spherical harmonics [17]. We choose the principal component analysis procedure for its simplicity to reduce our searching space. The octree approach only works for dense 3-D data such as CT and MRI, while our method can deal with different image formats.

The symmetry property of the orientation histogram is a necessary-but-not-sufficient condition for 3-D symmetry detection. That is, for certain non-symmetric objects, the orientation

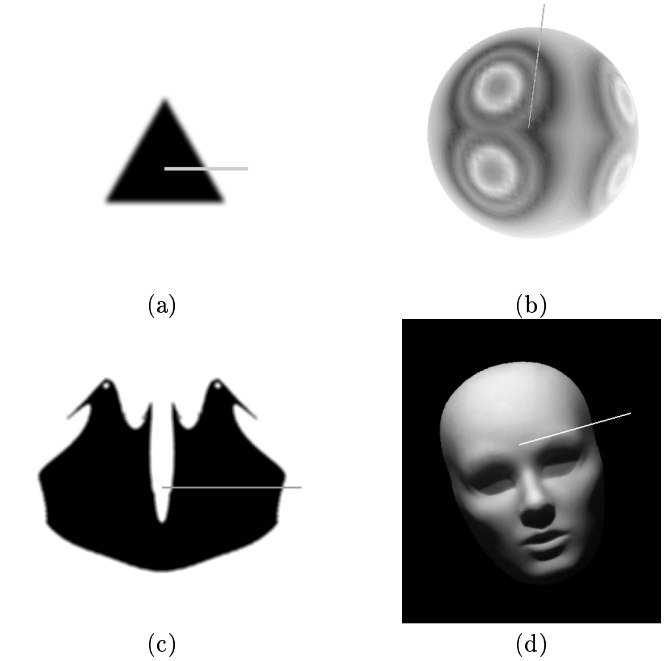

Fig. 4. Display of reflectional symmetry detection results for 2-D image

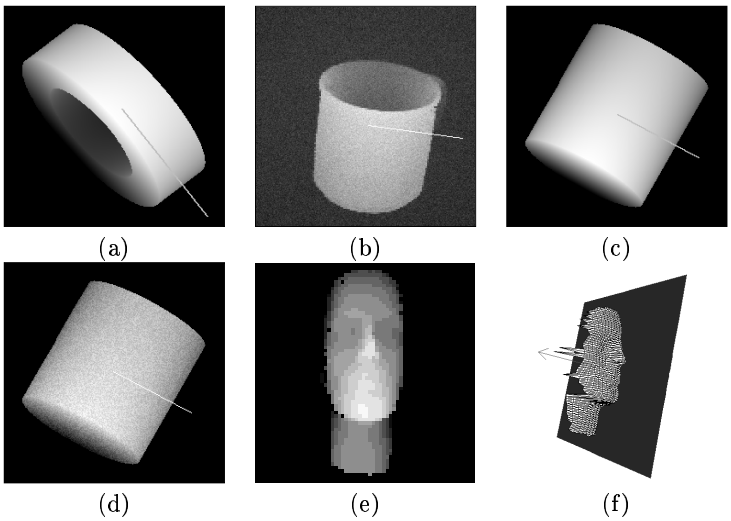

Fig. 5. Reflectional symmetry detection results for 2-D range images. (a) A tape roll. (b) A range cup image. (c) A Harris cup. (d) Noise added to (c), and the symmetry was also detected. (e) A range image of a human head. (f) The reflectional symmetry detected for the human head image. The arrow is the normal of the reflectional plane.

histogram might still be symmetric. As long as we know that the object in the image is reflectionally or rotationally symmetric, we can apply our algorithm to detect symmetry. When the orientation histogram shows reflectional or rotational symmetry but the object is not symmetric, a further step might be necessary to check whether the object does exhibit same symmetry. The symmetry hypothesis can be tested using a symmetry measure (see [3]). If this symmetry measure gives a low value, we say that the object is not symmetric. Note that if an object is convex, then the histogram symmetry is a sufficient condition for 3-D symmetry in the object.

The orientation histogram for a spherical object (with a uniform distribution of EGI intensities) and for a flat object whose facets are randomly oriented will be approximately the same.

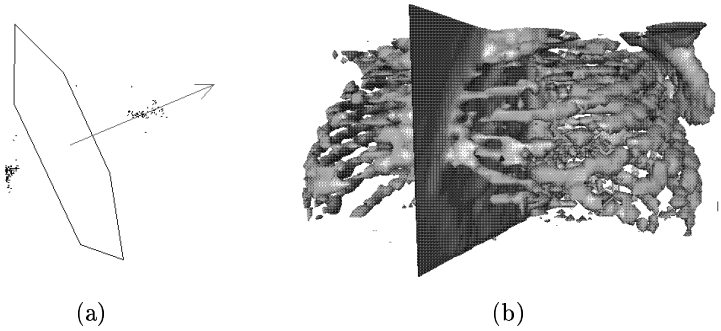

Fig. 6. (a) Results of reflectional symmetry detection for a point data set; (b) Results of reflectional symmetry detection for a 3-D CT image.

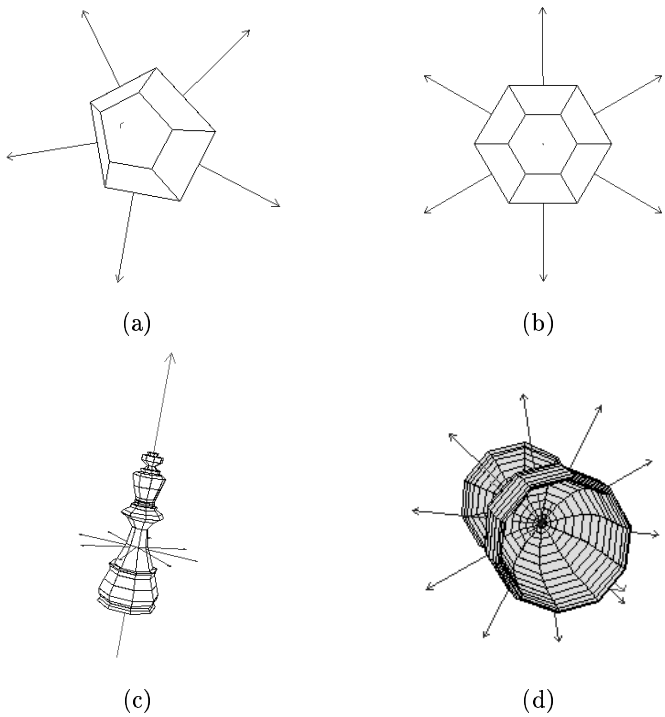

Fig. 7. Rotational symmetry detection results for wire-frame objects.

This ambiguity will be resolved by the principal component analysis step. For objects that are not rotationally symmetric, there will be a trivial symmetry: if an object is rotated 360 degrees, it will come back to its original position. The order number in this case will be one.

For reflectional symmetry, if i is the resolution of the sphere tessellation, the algorithm complexity will be $O(i4^i)$. For rotational symmetry, the complexity will be $O(i4^i f_{max}^2)$, where f_{max} is the maximum order of symmetry sought.

VI. Conclusions

Using our approach, the symmetry detection problem has been converted to the correlation of the Gaussian image. The algorithm accommodates different kinds of images or objects. After the orientation histogram of the image has been obtained, the symmetry detection algorithm is independent of the image format. From the results shown in this correspondence, rotational and reflectional symmetry directions can be determined using the statistics of the orientation histogram. The symmetry information can then be employed to represent the object economically, perform recognition, determine a suitable grasping

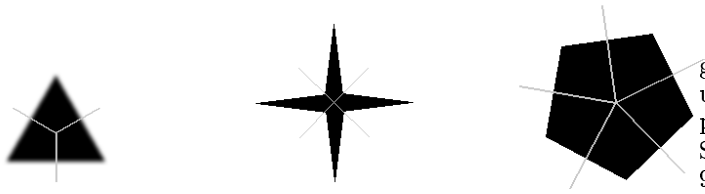

(a) (b) (c)

Fig. 8. Rotational symmetry detection results for 2-D images.

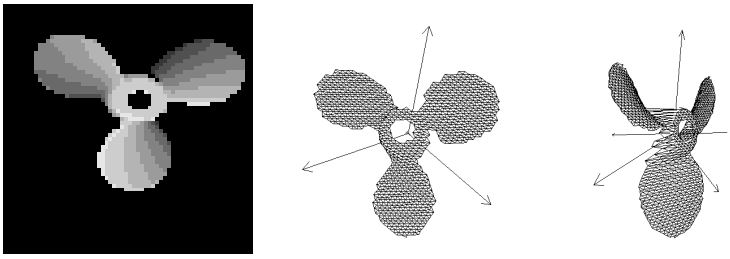

(a) (b) (c)

Fig. 9. Rotational symmetry detection results for a 2-D range image. (a) Range image of a propeller; (b) Detected symmetry information overlaid on the wire-frame representation of the range data; (c) A second view of (b) (the arrow in (c) near horizontal direction is the symmetry axis).

orientation, and so on.

The symmetry found in the objects is not an exhaustive list — only the symmetry which resulted in the maximum correlation is chosen. The speed of the algorithm varied with its application. Rotational symmetry took longer to find than reflectional symmetry. Increasing the resolution of the histogram invariably increased execution time. In most cases the algorithm took about 1 min for reflectional symmetry, and 5-10 mins for rotational symmetry on Sun Sparc10. The execution time compares favourably with the octree method presented in [18], where symmetries are detected for simple modelled objects in 1-5 mins, but for complex medical images in 100 mins on an Apollo Domain 4000.

As for accuracy, the algorithm can theoretically be as accurate as the sphere resolution. However, some factors of the implementation prevent this increase in accuracy. First, the principal components have limited accuracy: in the wire-frame case, the accuracy is set by the integration step-size. In the 2-D and 3-D image cases, the integration involved visiting each point so that the resolution could not be increased further. Second, the method of searching the adjacent facets of the principal component axes is restrictive. As the resolution of the sphere is increased, the neighbours become closer and closer to the principal axis bin directions, but the accuracy of the principal axes remains the same. Therefore in some cases increasing the resolution can give worse results, as the correct symmetry direction is never searched. This problem could be obviated in future implementations by searching a larger neighbourhood of the principal axes.

The execution time could be reduced by optimising the code, and a further increase in efficiency could be introduced when finding the symmetry order. Rather than checking all orders up to some maximum, only prime numbers need to be checked, and the search space for the order of symmetry can be reduced through deduction.

Acknowledgements

We thank the reviewers for their valuable comments. Thanks go to Nick Fisher, who supplied the data set shown in Figure 6(a), and Steve Davies for his help with the **Explorer** package. We are grateful to C. S. Chua, M. Takatsuka and S. Geddes for providing the range images in Figures 5(e) and 9(a). Hugues Talbot made the observation about checking only prime orders of symmetry and proof-read the paper. The code for the hierarchical representation of the triangular tessellations was written by Dr E. Trucco and Dr M. Umasuthan from Heriot-Watt University, Scotland.

References

[1] H. Alt, K. Mehlhorn, H. Wagener, and E. Welzl, "Congruence, similarity and symmetries of geometric objects," *Discrete Computational Geometry*, vol. 3, pp. 237–256, 1988.

[2] J. D. Wolter, T. C. Woo, and R. A. Volz, "Optimal algorithms for symmetry detection in two and three dimensions," *The Visual Computer*, vol. 1, pp. 37–48, 1985.

[3] H. Zabrodsky, S. Peleg, and D. Avnir, "Symmetry as a continuous feature," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 12, pp. 1154–1165, December 1995.

[4] H. Zabrodsky, S. Peleg, and D. Avnir, "Hierarchical symmetry," in *Proceedings of International Conference on Pattern Recognition*, August 30–September 3, The Hague, The Netherlands, 1992, vol. III, pp. 9–11.

[5] X. Y. Jiang and H. Bunke, "A simple and efficient algorithm for determining the symmetries of polyhedra," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, vol. 54, no. 1, pp. 91–95, January 1992.

[6] S. Parry-Barwick and A. Bowyer, "Symmetry analysis and geometric modelling," in *Digital Image Computing: Techniques and Applications*, K. K. Fung and A. Ginige, Eds. Dec. 1993, vol. I, pp. 39–46, Australian Pattern Recgnition Society.

[7] S. Ishikawa, K. Sato, P. Minovic, and K. Kato, "An interactive 3D symmetry analysis system," in *IAPR Workshop on Machine Vision Applications*, Dec. 7-9 1992, pp. 375–378.

[8] P. Minovic, S. Ishikawa, and K. Kato, "Symmetry identification of a 3D object represented by octree," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 5, pp. 507–154, May 1993.

[9] K. Ikeuchi, "Recognition of 3-D objects using the extended Gaussian image," in *IJCAI*, Aug.24-28 1981, pp. 595–600.

[10] B. K. P. Horn, *Robot Vision*, The MIT Press, 1986.

[11] P. Brou, "Using the Gaussian image to find the orientation of objects," *International Journal of Robotics Research*, vol. 3, no. 4, pp. 89–125, 1984.

[12] C. M. Brown, "Representing the orientation of dendritic fields with geodesic tessellations," Tech. Rep. TR-13, Computer Science Dept., Univ. of Rochester, 1977.

[13] S. B. Kang and K. Ikeuchi, "The complex EGI: A new representation for 3D pose determination," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 7, pp. 707–721, July 1993.

[14] M. R. Korn and C. R. Dyer, "3-D multiview object representations for model-based object recognition," *Pattern Recognition*, vol. 20, no. 1, pp. 91–103, 1987.

[15] P. Minovic, S. Ishikawa, and K. Kato, "Three-dimensional symmetry identification. part I: Theory," Tech. Rep. 21, Kyushu Institute of Technology, Japan, 1992.

[16] N. I. Fisher, T. Lewis, and B. J. J. Embleton, *Statistical Analysis of Spherical Data*, Cambridge University Press, 1987.

[17] G. Burel and H. Hennocq, "Determination of the orientation of 3D objects using spherical harmonics," *Graphical Models and Image Processing*, vol. 57, no. 5, pp. 400–408, September 1995.

[18] P. Minovic, S. Ishikawa, and K. Kato, "Three-dimensional symmetry identification. part III: Experimental results," Tech. Rep. 21, Kyushu Institute of Technology, Japan, 1992.