

Finding Axes of Symmetry From Potential Fields

V. Shiv Naga Prasad and B. Yegnanarayana

Abstract—This paper addresses the problem of detecting axes of bilateral symmetry in images. In order to achieve robustness to variation in illumination, only edge-gradient information is used. To overcome the problem of edge breaks, a potential field is developed from the edge map which spreads the information in the image plane. Pairs of points in the image plane are made to vote for their axes of symmetry with some confidence values. To make the method robust to overlapping objects, only local features in the form of Taylor coefficients are used for quantifying symmetry. We define an axis of symmetry histogram, which is used to accumulate the weighted votes for all possible axes of symmetry. To reduce the computational complexity of voting, a hashing scheme is proposed, wherein pairs of points, whose potential fields are too asymmetric, are pruned by not being counted for the vote. Experimental results indicate that the proposed method is fairly robust to edge breaks and is able to detect symmetries even when only 0.05% of the possible pairs are used for voting.

Index Terms—Axis of symmetry histogram, bilateral symmetry, gradient vector flow field, potential field, voting.

I. INTRODUCTION

SYMMETRY is an important consideration in Gestalt laws for perceptual grouping [1]. The reason is that it is very difficult for two completely unrelated objects to be placed in such a manner as to produce spatial symmetry. It is much more likely that if a group of objects exhibit symmetry, then they are related and are perceived in relation to one another (i.e., they are grouped together perceptually). This has motivated research in the detection of symmetries in images and shapes.

There are two types of symmetries [2], as follows.

- 1) **Bilateral Symmetry** (also called mirror symmetry): A figure is said to possess bilateral symmetry if it is invariant under a reflection about a line (called the axis of symmetry) passing through the centroid of the figure [3]–[6].
- 2) **Rotational Symmetry**: A figure is said to possess rotational symmetry (C_n) of order n if it is invariant under rotations of $(2\pi/n)$ radians about its center of mass. Central symmetry is a special case of rotational symmetry with $n = 2$ (see [6]–[8] for work on rotational symmetry and [9], [10] for central symmetry).

Manuscript received January 9, 2003; revised February 13, 2004. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ivan W. Selesnick.

V. S. N. Prasad was with the Department of Computer Science and Engineering, Indian Institute of Technology–Madras, Chennai 600036, India. He is now with the Institute of Advanced Computer Studies (UMIACS), University of Maryland, College Park, MD 20742 USA (e-mail: shivnaga@cs.umd.edu).

B. Yegnanarayana is with the Department of Computer Science and Engineering, Indian Institute of Technology–Madras, Chennai 600036, India (e-mail: yegna@cs.iitm.ernet.in).

Digital Object Identifier 10.1109/TIP.2004.837564




Fig. 1. Example of different types of symmetries. (a) Rotational symmetry of order 4 (C_4). (b) Bilateral symmetry.

Fig. 1 illustrates the two types of symmetries. When symmetric objects undergo skew transformation, they produce skew symmetries [11]–[13].

Symmetry of an object is a binary characteristic, i.e., an object is either symmetric, or it is not. However, perfect symmetry is rarely observed in real images. Objects, though inherently symmetrical, usually exhibit some deviations from the ideal. Thus, in the real world there is a continuum in symmetry. This issue has been discussed in [14], [15], where the amount of deviation of a given figure from perfect symmetry is computed.

Some of the desired properties of any method used for detecting symmetry are the following:

- a) avoid a brute-force search in the image;
- b) robustness to disjoint contours, including contours formed out of dots;
- c) robustness to minor deviations from perfect symmetry;
- d) ability to detect multiple symmetries simultaneously;
- e) robust to overlap and occlusion;
- f) should not assume simple curves (e.g., nonself-intersection, convexity, etc.).

Several techniques have been proposed for detecting symmetry based on the computation of the medial-axis transform (MAT) of the image, e.g., hierarchical Voronoi skeletons [16], Hamilton–Jacobi transform [17], magnetic field-based methods [18], etc. However, it is a known problem that the medial axis of a figure can undergo radical deformation in the case of overlapping figures. The other generic approach cited in literature is region based [6]. Although region-based methods are robust to deformation, it is difficult to employ them in the case of broken edges and overlapping figures.

In view of the above discussion, we adopt a middle course between MAT-based and region-based approaches. We characterize the space near/around the contour/figures and use local features to quantify symmetry. The proposed method does not address the problem of skew symmetries. The methodology adopted is as follows. A potential field is generated taking into account the location, orientation, and magnitude of the edge gradient. Local features in the form of Taylor coefficients are computed from the field at all points in the image plane. A hashing algorithm is then used to detect pairs of points which might possibly have symmetric fields (this avoids a brute-force search). Pairs of points which pass this test are then made to vote for their axes of symmetry, their vote being weighted by a

confidence measure. All the votes are accumulated in a global axis of symmetry histogram, which shows the net confidence values for all possible axes. Potential fields have the property of smoothing or “filling up” spaces (indicated by the Laplace equation). As a result, they are robust to edge breaks and dotted lines. Local features are used for hashing and voting, so that the method is robust to overlap and occlusion. To accommodate deviations from symmetry, the votes are weighted with confidence scores. A global axis of symmetry histogram is used so that multiple symmetries can be detected simultaneously, and deviations from ideal symmetry can be accommodated.

Section II describes the development of potential fields from the edge gradient. Section III describes the criteria for any two points in the potential field to have locally symmetric fields. Section IV describes the hashing scheme, the voting method, and the axis of symmetry histogram. Section V describes ways for further enhancing the axis histograms. Section VI presents experimental results, and Section VII summarizes the work and discusses possible extensions to the present work.

II. DERIVING POTENTIAL FIELD FROM EDGE GRADIENT

The objective of developing a potential field from the edge gradient is to bring about interaction between the gradients at different points in the image plane. The motivation for bringing about interaction is that it makes the system robust to noise and deformation. The potential field employed in the present work is the gradient vector flow (GVF) field [19]. Let the image be $I(x, y)$ and let $f(x, y)$ be its edge map, where $f(x, y)$ is larger near the edges and ∇f is the edge gradient. The GVF for the image is denoted by $\mathbf{v} = [u(x, y), v(x, y)]$, where $u(x, y)$ and $v(x, y)$ are two fields defined on the image plane. Let u_x, u_y, v_x and v_y be the partial derivatives of $u(\cdot)$ and $v(\cdot)$. $\mathbf{v}(x, y)$ is obtained by minimizing the energy functional \mathcal{E} , defined as

$$\mathcal{E} = \int \int [\mu (u_x^2 + u_y^2 + v_x^2 + v_y^2) + |\nabla f|^2 |\mathbf{v} - \nabla f|^2] dx dy. \quad (1)$$

The first term in the integrand in (1) is the smoothing factor, similar to the formulation of optical flow used by Horn and Schunk [20]. This term imposes a penalty on the divergence and curl of \mathbf{v} [21]. At points where $|\nabla f|$ is small (i.e., edge gradient is low), the smoothing term dominates.

The second term is responsible for introducing image information into \mathbf{v} . Wherever $|\nabla f|$ is significant (high edge gradient is present), \mathbf{v} is forced toward ∇f . The scaling factor μ decides the relative importance of the two terms in the integrand in (1).

Using a calculus of variations [22], it can be shown that for \mathcal{E} to be minimum, the following are necessary conditions:

$$\mu \nabla^2 u - (u - f_x) |\nabla f|^2 = 0 \quad (2)$$

$$\mu \nabla^2 v - (v - f_y) |\nabla f|^2 = 0. \quad (3)$$

It can be seen that when $\nabla f = 0$ the Euler equations (2) and (3) reduce to Laplace equations on u and v . In other words, u and v can be considered as potential fields with their values forced toward f_x and f_y , respectively, at points where the edge




Fig. 2. GVF fields of images depicting an ideal square and squares with distortions. Each column shows a 32×32 pixel image and the corresponding normalized field. The images are (a) ideal square, (b) square formed of irregular dots, and (c) square with an overlapping curve.

gradient is nonzero. An iterative method for the computation of $\mathbf{v}(x, y)$ using (2) and (3) is given in [19].

Potential fields have the property of additivity and the ability to “fill up” spaces. For example, the potential field produced by discrete charges placed along a curve is similar to the one produced by a continuous charge along the curve, provided that the measurement is not very close to the curve. Fig. 2 illustrates the GVF obtained for images of an ideal square and squares with distortions. Each column shows the 32×32 pixel image, and the corresponding GVF. Fig. 2(a) and (b) show the normalized¹ GVF for continuous and discontinuous curves. The fields produced by images in Fig. 2(a) and (b) are similar when observed at some distance from the edges. The potential field at points situated away from the edges will show gross or averaged features. This makes potential fields robust to edge breaks and slight deformations.

If there is no interference, like overlap or occlusion, then the GVF for a figure exhibits valleys corresponding to the medial axis. For a symmetric curve, the axis of symmetry will coincide with one of the valleys. However, in the case of overlap and interference from noise, this property is lost. For example, in Fig. 2(c), the medial axis of the square is lost due to overlap. One way of overcoming this problem is to compute local features of the GVF at various points in the image plane, and then get the axes of symmetry based on these local features. The GVF field at a point depends more on points nearby than those at some distance. Hence, interference near a particular curve from other curves will be less. This is reflected in the local features, which predominantly characterize the curve to which they are the closest.

The local features used in the present work are the coefficients of the Taylor series expansions of $u(x, y)$ and $v(x, y)$ fields. We limit ourselves to coefficients up to the first order. The truncated Taylor expansion represents the actual field reasonably well in a small neighborhood around the center.

III. CONDITIONS FOR SYMMETRY

Let (p_i, q_i) and (p_j, q_j) be two points in the image plane, and let the GVF be $\mathbf{v} = [u(x, y), v(x, y)]$. Let $p_j = p_i + r \cos \phi$ and $q_j = q_i + r \sin \phi$, where $r > 0$. Let $u_i = u(p_i, q_i)$, $u_{ix} =$

¹We normalize the GVFs before plotting them, so as to make the field clearly visible. Let $\mathbf{v}_n(x, y)$ be the normalized GVF $\mathbf{v}(x, y)$, then $\mathbf{v}_n(x, y) = (\mathbf{v}(x, y)) / (\sqrt{u^2(x, y) + v^2(x, y)})$.

$(\partial u/\partial x)|_{p_i, q_i}$ and $u_{iy} = (\partial u/\partial y)|_{p_i, q_i}$. Similarly, we define $v_i, u_j, v_j, v_{ix}, v_{iy}, u_{jx}, u_{jy}, v_{jx}$, and v_{jy} . If the fields at the two points are symmetric, then the field at (p_j, q_j) can be obtained by “flipping” the field at (p_i, q_i) about the y axis, and then rotating it anticlockwise by 2ϕ radians.

By comparing the coefficients of the Taylor series expansions of the fields $u(\cdot)$ and $v(\cdot)$ at the points (p_i, q_i) and (p_j, q_j) , it can be shown that, for the fields at the two points to be bilaterally symmetrical, the following conditions are necessary:

$$-u_i = u_j \cos 2\phi + v_j \sin 2\phi + d_1 \quad (4)$$

$$u_{iy} = -u_{jy} \cos^2 2\phi - (v_{jy} - u_{jx}) \sin 2\phi \cos 2\phi + v_{jx} \sin^2 2\phi + d_2 \quad (5)$$

$$u_{ix} = u_{jx} \cos^2 2\phi + (v_{jx} + u_{jy}) \sin 2\phi \cos 2\phi + v_{jy} \sin^2 2\phi + d_3 \quad (6)$$

$$v_i = -u_j \sin 2\phi + v_j \cos 2\phi + d_4 \quad (7)$$

$$v_{iy} = u_{jx} \sin^2 2\phi + (-v_{jx} - u_{jy}) \sin 2\phi \cos 2\phi + v_{jy} \cos^2 2\phi + d_5 \quad (8)$$

$$v_{ix} = u_{jy} \sin^2 2\phi + (u_{jx} - v_{jy}) \sin 2\phi \cos 2\phi - v_{jx} \cos^2 2\phi + d_6 \quad (9)$$

where $d_1 \dots d_6$ are the error terms. [see Appendix for details of (4)–(9)]. For perfectly symmetric fields, the error terms will be zero. Let $V(x, y) = u^2(x, y) + v^2(x, y)$, $h_1(x, y) = u_x(x, y) + v_y(x, y)$, and $h_2 = v_x(x, y) - u_y(x, y)$, then from (4)–(9), the necessary, but not sufficient, conditions for perfect symmetry (i.e., $d_1 = d_2 = \dots = d_6 = 0$) are

$$C_1 : V(p_i, q_i) = V(p_j, q_j)$$

$$C_2 : h_1(p_i, q_i) = h_1(p_j, q_j)$$

$$C_3 : h_2(p_i, q_i) = -h_2(p_j, q_j).$$

Unlike (4)–(9), conditions C_1, C_2 and C_3 can be checked *independent* of ϕ . This fact will be used later in the hashing stage.

Another local feature extracted from GVF is the curvature of GVF ($C(x, y)$). The curvature captures the bends in the edges and is invariant under translation and rotation. It is computed as follows:

$$C(x, y) = \frac{1}{|v|^3} [(v_x + u_y)uv - u_x v^2 - v_y u^2]. \quad (10)$$

For the two points (p_i, q_i) and (p_j, q_j) to have locally symmetric fields, the curvature of GVF at the two points must be similar, i.e., condition $C_4 : C(p_i, q_i) = C(p_j, q_j)$ must hold true. One of the advantages of $C(x, y)$ is that it can be computed independently for each point, irrespective of the pairing (that is independent of ϕ).

In order to observe the behavior of $V(x, y), h_1(x, y), h_2(x, y)$, and $C(x, y)$ in the presence of interference and noise, the functions were computed for the three images in Fig. 2. The results are presented in Fig. 3. The plots for the ideal square [column (a)] are used as reference. We can make the following observations from Fig. 3.

- 1) In case of edge breaks and discontinuous edges [Fig. 3(b)], the four functions show the desired symmetry as the distance from the edges increases. This is because the potential field, through superimposability,




Fig. 3. Behavior of $V(x, y), h_1(x, y), h_2(x, y)$, and $C(x, y)$ for images in Fig. 2. Plots are in the same order.

is able to reduce the disruption caused by discontinuous edges.

- 2) In the case of overlap [Fig. 3(c)], the functions exhibit symmetry *near* the edges of the squares. The reason is that the edge gradients of the square are able to dominate in these regions and produce the desired symmetry.
- 3) The deviation in the conditions $C_1 \dots C_4$ depends upon the level of interference in the image; hence, the tolerance values for the conditions need to be adjusted according to the case at hand.

IV. VOTING FOR AXES OF SYMMETRY

In order to ensure that the detection of axes of symmetry is robust to overlap, occlusion, etc., only local GVF features are checked for symmetry. As the extracted features give only local knowledge, we need to employ a voting scheme to compute globally consistent information about the axes of symmetry. Voting schemes have been extensively studied and employed in image processing [23]–[25]. For the present problem, during the voting process, points are taken in pairs and their mutual symmetry is quantified. This is then used as a weightage factor for the vote each pair casts for its probable axis of symmetry. But, the number of votes for an $m \times n$ image will be $O(m^2 n^2)$, which is highly inefficient, as most of the pairs are asymmetric and hence their votes are redundant.

One way of increasing the efficiency of voting is to employ a hashing technique to eliminate pairs of points which are too asymmetric. In the present case, we map the points in the image plane into a hash space such that points mapped to the same bin in the hash space might be symmetric and, hence, are allowed to vote. Those falling into different bins are too asymmetric for their vote to count and, hence, are not allowed to vote (pruned). For the hashing operation, only those features

are useful which can be independently computed for each point in the image plane (i.e., independent of ϕ). The functions $V(x, y)$, $h_1(x, y)$, $h_2(x, y)$ and $C(x, y)$ meet this criterion. Let $\langle i, j \rangle$ denote the pair of points (p_i, q_i) and (p_j, q_j) . The pair $\langle i, j \rangle$ will be allowed to vote for its axis of symmetry only if the condition \mathcal{D} is true, where \mathcal{D} is defined as

$$\begin{aligned} \mathcal{D} : & (V(p_i, q_i) - V(p_j, q_j))^2 < \Delta_1 \\ & (h_1(p_i, q_i) - h_1(p_j, q_j))^2 < \Delta_2 \\ & (h_2(p_i, q_i) + h_2(p_j, q_j))^2 < \Delta_3 \quad \text{and} \\ & (C(p_i, q_i) - C(p_j, q_j))^2 < \Delta_4. \end{aligned}$$

Here, $\Delta_1 \dots \Delta_4$ set the granularity of the bins (tolerance values for the deviations) and, thus, control the degree of pruning to be carried out.

The fields $u(x, y)$ and $v(x, y)$ are predominantly smoothly varying, except near the edges. Consequently, $h_1(x, y)$ and $h_2(x, y)$, which are linear combinations of the partial derivatives of $u(\cdot)$ and $v(\cdot)$, have very low variation compared to that of $V(x, y)$ and $C(x, y)$. Therefore, $V(x, y)$ and $C(x, y)$ perform the major portion of the hashing operation, $h_1(x, y)$ and $h_2(x, y)$ being used for fine tuning the hashing.

Once the hashing operation has been carried out, pairs of points which have a good chance of having symmetric fields are obtained. The next step is to make each of these pairs vote for its axis of symmetry with some confidence value. Let $\langle i, j \rangle$ be one of the pairs of points obtained after hashing. It will vote for a line denoted by (θ, ρ) , where the line is defined as

$$\begin{aligned} y(t) &= -\rho \cos \theta + t \sin \theta \\ x(t) &= \rho \sin \theta + t \cos \theta \end{aligned} \quad (11)$$

where t is an independent parameter

$$\rho = \frac{p_i + p_j}{2} \sin \theta - \frac{q_i + q_j}{2} \cos \theta$$

and

$$\theta = \begin{cases} \phi + \frac{\pi}{2}, & \text{if } \phi \in [-\frac{\pi}{2}, \frac{\pi}{2}) \\ \phi - \frac{\pi}{2}, & \text{if } \phi \in [\frac{\pi}{2}, \frac{3\pi}{2}) \end{cases}.$$

For calculating the weight factor for the vote, the field at (p_i, q_i) is predicted using the field at (p_j, q_j) and the angle ϕ . The degree of asymmetry is given by the deviation of the actual field from the predicted field. We use relations (4)–(9) to predict the field and obtain the errors $d_1 \dots d_6$. For a pair $\langle i, j \rangle$, let $e_1(\langle i, j \rangle) = (d_1^2 + d_4^2)/(2)$ and $e_2(\langle i, j \rangle) = (d_2^2 + d_3^2 + d_5^2 + d_6^2)/(4)$, where $e_1(\langle i, j \rangle)$ gives the average error in predicting $u(\cdot)$ and $v(\cdot)$, and $e_2(\langle i, j \rangle)$ gives the average error in predicting the partial derivatives. The weight factor for the pair $\langle i, j \rangle$'s vote is given by

$$w(\langle i, j \rangle) = e^{-(\frac{e_1(\langle i, j \rangle)}{\sigma_1} + \frac{e_2(\langle i, j \rangle)}{\sigma_2})} \min(V(p_i, q_i), V(p_j, q_j)) \quad (12)$$

where σ_1 and σ_2 are positive. The first factor is an inverse exponential of the errors, and its sensitivity to deviations from symmetry is adjusted using the σ_k 's. The second factor gives more




Fig. 4. Axis histogram obtained from the ideal square in Fig. 2(a). θ varies from 0° to 180° , ρ varies from -45 to 45 . The four peaks are numbered corresponding to the four axes of symmetry present in the square.

weight to votes with higher gradient flow magnitude. Images with sparse edge maps have large regions with very low gradient flow magnitudes. In the absence of the second factor, these regions would dominate the voting process and obscure the actual axes of symmetry.

The votes cast by the pairs obtained after hashing are accumulated in an axis of symmetry histogram (called axis histogram in short), which has θ and ρ as its basis. The histogram is an $l_\rho \times l_\theta$ matrix of bins which quantize the $\theta - \rho$ space. The range of values taken by ρ for an $m \times n$ image is $[-\sqrt{m^2 + n^2}, \sqrt{m^2 + n^2})$ units and by θ is $[0, \pi)$ radians. In our implementation, the bins are of uniform size, with their width along ρ being 1 unit and along θ being .05 radians ($\approx 3^\circ$). Thus, for an $m \times n$ image, we obtain an $l_\rho \times l_\theta$ histogram, where $l_\rho = 2\sqrt{m^2 + n^2}$ and $l_\theta = 64$. Fig. 4 shows the axis histogram obtained from the image in Fig. 2(a) (ideal square), after 0% pruning, i.e., all possible pairs are allowed to vote.

Usually, when distortions are introduced into an image depicting a symmetric figure, the symmetry of the figure is reduced. As a result, the axis histogram for a distorted image will have less distinct peaks. Ideally, the degradation in the axis histogram due to distortions in the image should be graceful. When the amount of pruning is increased (i.e., Δ_k 's are made more stringent), then the amount of information available to compute the histogram is reduced. Ideally, the method should be robust to pruning, and the degradation in the histograms with increasing levels of pruning should be gradual and smooth.

As part of the first experiment, we observed the behavior of the proposed method in the presence of different distortions in images for different levels of pruning. Fig. 5 shows axis histograms computed for each of the three images in Fig. 2 for two levels of pruning: 0% and 99% (percentage of pairs of points which were not allowed to vote). We can make the following observations from the results.

- 1) The peaks in the histograms for the square formed from dots [row (b)] are least distinct. This is because the field becomes symmetrical only after some distance from the dots, but, by then, the magnitude of the GVF becomes very low. As a result, the weight factors for the votes are low.
- 2) The method is able to handle overlap [row (c)], as well.
- 3) For all the three images, as the pruning is increased from 0% to 99%, the degradation in the histograms is smooth, i.e., except for an overall decrease in magnitude, the histograms retain their character. Thus, the method is not very sensitive to pruning.

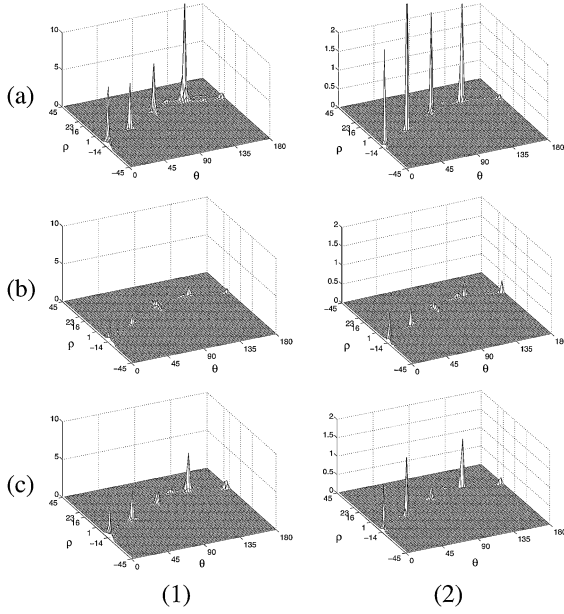


Fig. 5. Axis of symmetry histograms for the images in Fig. 2 for two prune levels. Each row shows the two axis histograms for each of the three images. Order of images is same as that in Fig. 2. The first column is for 0% pruning and the second column is for 99% pruning. Note that the z -axis scale for histograms at 0% pruning is from 0 to 10, and for those at 99% pruning, the scale is from 0 to 2.

V. ENHANCING THE AXIS HISTOGRAMS

In order to minimize the influence of an image's boundaries on its axis histogram, we need to set appropriate boundary conditions for the fields $u(\cdot)$ and $v(\cdot)$. In our implementation, this was done by forcing a whirlpool-like field along the border of the images.

As only local features are considered for quantifying symmetry, the axis histogram might show false peaks, corresponding to symmetries that are not visually significant. A few (even two or three) perfectly symmetrical voters are enough to cause a significant peak, even though the other voters may not be making any significant contribution. Such peaks are undesirable and should be suppressed. This can be achieved to some extent by preventing perfectly symmetrical pairs from voting. For visually dominant symmetries, a large number of voters will be present, and there will be a continuum in the voter's confidence values. Therefore, peaks corresponding to visually dominant symmetries are not much affected by the elimination of perfectly symmetrical voters. We define τ_α such that all voting pairs $\langle i, j \rangle$ with $e^{-((e_1(\langle i, j \rangle)/\sigma_1) + (e_2(\langle i, j \rangle)/\sigma_2))} > \tau_\alpha$ are not allowed to vote. Too high a value of τ_α will be ineffective, whereas too low a value will lead to excessive degradation of the histograms. Its value was experimentally fixed at 0.7, which is the value used in this paper. Fig. 6 shows a test image, its edge map, original axis histogram, and the axis histogram obtained after eliminating perfectly symmetrical voters. The level of pruning is 99.92%. It can be seen that the peak corresponding to the visually dominant symmetry becomes more prominent in the second histogram [Fig. 6(d)].

False peaks in the axis histogram can be further suppressed by verifying whether the underlying edge map is symmetrical.

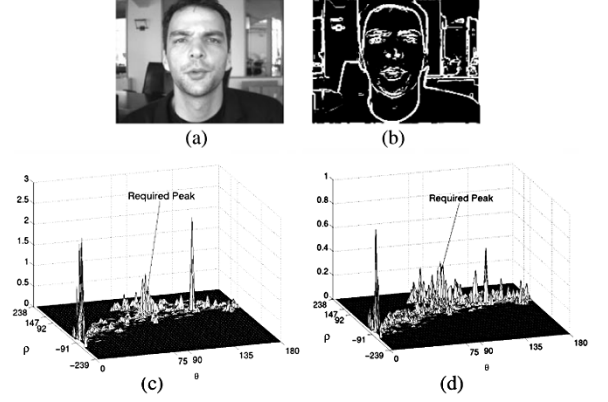


Fig. 6. Example of axis histogram obtained after elimination of perfectly symmetrical voting pairs. (a) Image. (b) Edge map. (c) Original histogram. (d) Histogram obtained after the elimination.

This verification is done for only those peaks whose magnitude is greater than quarter of the histogram's maximum peak value. Let K be the original histogram² and let K' be the thresholded histogram, then

$$K'(s, t) = \begin{cases} K(s, t), & K(s, t) > \frac{1}{4} \max_{a, b} K(a, b) \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

We can summarize the post-processing as follows. For each nonzero bin in K' , do the following.

- 1) Get the voting pairs in its 3×3 neighborhood.
- 2) Let $L(\langle i, j \rangle)$ denote the line segment joining the points of the pair $\langle i, j \rangle$. Arrange the collected voting pairs in increasing order of the perpendicular distances of $L(\cdot)$ s from the origin.
- 3) Generate two sequences S_1 and S_2 of $\sqrt{V(\cdot)}$ values by going in straight lines from one voting point to another along the two sides of the axis of symmetry.
- 4) If the underlying edge map is symmetrical, then the two sequences would be identical or nearly so. Let D be the averaged difference in the sequences, T be their length, and p be the number of voters

$$D = \frac{1}{Tp} \sum_t |S_1(t) - S_2(t)|. \quad (14)$$

- 5) Scale the peak of the bin under consideration using the error D . Let K'' be the histogram thus obtained. Then

$$K''(s, t) = K'(s, t) \frac{\max_{a, b} D(a, b)}{D(s, t)}. \quad (15)$$

Fig. 7 shows the enhanced histogram (K'') obtained from the histogram in Fig. 6(d). It can be seen that a number of false peaks are suppressed, making the peak for the visually dominant symmetry relatively more prominent.

VI. EXPERIMENTAL RESULTS

Henceforth, by axis histograms, we mean the enhanced axis histograms obtained after the above-discussed processing. The

²From here onward, the histograms are the ones with the perfectly symmetrical voters eliminated.




Fig. 7. Enhanced histogram obtained for the image in Fig. 6 after taking into account the underlying edge map.

proposed method was tested with several images from CalTech Image database, BioID database, and SAMPL Image database. In order to analyze its behavior for a given figure, a prominent peak in the obtained histogram was chosen, and pairs of points voting for the corresponding bins were plotted with the GVF overlayed. This plot is called the Voting-Pair plot. Each pair of points in the Voting-Pair plot is designated by a line segment whose ends are located at the corresponding points in the image plane. In our implementation the pairs plotted in the Voting-Pair plot are chosen as follows. Upon selecting a particular peak in the histogram, bins in the 3×3 neighborhood of the peak's maxima are taken, and all pairs of points voting for one of these bins are plotted in the Voting-Pair plot.

The Voting-Pair plot helps in observing the following:

- 1) how the pairs are chosen for voting, i.e., which pairs are rejected, and which are allowed to vote;
- 2) how the axis of symmetry histogram aids in clustering together pairs of points which have the same/similar axis of symmetry.

Fig. 8 shows one of the test images used in the second experiment, its edge map, the enhanced axis of symmetry histogram, and the Voting-Pair plot for one of the peaks. The parameters used were: $\Delta_1 = 10^{-6}$, $\Delta_2 = \Delta_3 = \Delta_4 = .01$ and $\sigma_1 = \sigma_2 = 10^{-4}$. An additional constraint imposed upon on the voters was that only points with $V(x, y) > 0.01$ were allowed to vote. This was done to reduce the computation. The level of pruning was 99.96% (i.e., 13 958 pairs out of the possible 39 139 128 were allowed to vote).

From the histogram in Fig. 8, we can see that the peak corresponding to the axis of symmetry of the male's face (peak 2) is not prominent. This is because of the pronounced skew in the posture. On the other hand, the peak corresponding to the axis of symmetry of the female's face (peak 1) is prominent, as the skew in this case is very small. The most prominent peak corresponds to the symmetry of the right eye of the female. Since the method was designed to accommodate slight deviations from ideal symmetry, it is robust to small skews exhibited by three-dimensional (3-D) objects. However, it does not handle cases where the skew is significant. Fig. 9(a) and (b) show the axis histograms for two other test cases: (a) when there is a large amount of overlap and interference and (b) highly broken and irregular edges.

We can see that, for a given image, although, perceptually, we may see only one axis of symmetry, the axis histogram shows up several symmetries (in the form of peaks). This is because of two reasons: 1) the method uses only local features and 2) humans

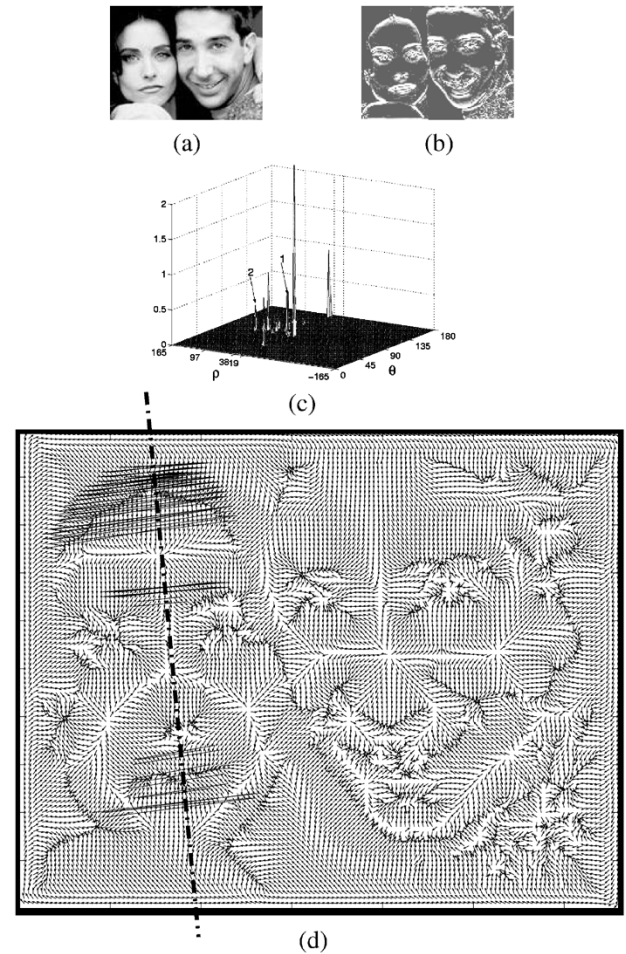


Fig. 8. Axis of symmetry histogram for a real-world image. (a) Cropped image taken from "Friends" photograph. (b) Corresponding edge map. (c) Axis histogram. (d) Voting-Pair plot for peak 1.

eliminate a number of possible symmetries by using semantic knowledge. Another possible reason for spurious peaks is that the proposed method uses only edge gradient information. The edge map is almost an outline image. The use of color/grayscale values might help in eliminating some of the possible symmetries highlighted by the axis histogram, but this would be sensitive to changes in illumination. If the illumination is asymmetric, then the grayscale information can no longer be directly used for quantifying symmetry.

The axis histogram can be used to detect clusters of pairs of points with similar axes of bilateral symmetry, which can, in turn, be used to aid perceptual grouping. The presence of multiple peaks in the axis histograms makes a completely objective evaluation of the method difficult. We tested the method with a number of images (more than 50) and found that, in almost all cases, for pruning levels of the order of 99.9%, the required symmetry was present as a prominent peak in the axis histograms. When the fraction of point pairs disqualified from voting is very high (above 99.99%), then the axis histograms are so degraded that the expected peaks are no longer prominent in the histograms. Fig. 10 shows a plot illustrating this trend for the axis histograms obtained for the image shown in Fig. 6(a), for different levels of pruning.

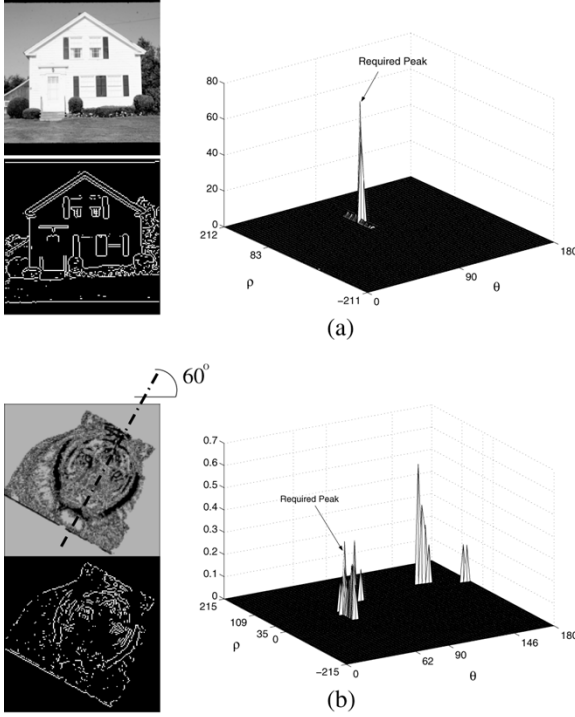


Fig. 9. (a) Example of overlapping figures and clutter: The vertical symmetry of the house suffers interference from the door and windows. (b) Example of highly broken edges: Frontal photograph of a tiger.




Fig. 10. Plot showing typical peak height variation with increasing levels of pruning. The y axis is the ratio of the height of the peak corresponding to the desired symmetry to that of the other tallest peak present in the axis histogram. Note that the x axis has been drawn with nonuniform scaling.

VII. CONCLUSION

A method for detecting axes of bilateral symmetry in images was proposed. The method does not address the problem of skew symmetry. The results indicate that the proposed method is fairly robust to distortions in the image and is able to detect symmetries, even when the pruning at the hashing stage is high. For real-world images, the method is able to detect symmetry, even when only 0.05% of the pairs of points are allowed to vote.

The method can be used in applications requiring fast detection of symmetries (e.g., aiding face detection), for perceptual grouping and object tracking.

There are three major directions for future work.

- 1) Make the method iterative, i.e., start off with a very high level of pruning (99.9%) and get the approximate axes of symmetry. Use this as *a priori* knowledge in a more detailed search.
- 2) Incorporate allowances for skew symmetries. This will enable the method to handle cases where 3-D objects in the image have undergone skew transformations.
- 3) Use the symmetry information obtained from the axis histogram to modify the GVF, and then recompute the axis histogram. By repeating this a number of times, we can highlight features in the GVF which show symmetry along a particular axis, which, in turn, can be used to enhance the corresponding regions in the image.

APPENDIX

Let (p_i, q_i) and (p_j, q_j) be the two points under consideration for symmetry. First, consider the case where $\phi = 0$, where ϕ is as defined in Section III. Let $(p_i + x, q_i + y)$ be a point in the neighborhood of (p_i, q_i) . For the field in the neighborhood of the two points (p_i, q_i) and (p_j, q_j) to be symmetric, we require that

$$\begin{bmatrix} -u(p_i + x, q_i + y) \\ v(p_i + x, q_i + y) \end{bmatrix} = \begin{bmatrix} u(p_j - x, q_j + y) \\ v(p_j - x, q_j + y) \end{bmatrix}. \quad (16)$$

Now, if we extend the above case by including the possibility for ϕ to be nonzero, then we require

$$\begin{bmatrix} -u(p_i + x_i, q_i + y_i) \\ v(p_i + x_i, q_i + y_i) \end{bmatrix} = \begin{bmatrix} \cos 2\phi & \sin 2\phi \\ -\sin 2\phi & \cos 2\phi \end{bmatrix} \begin{bmatrix} u(p_j + x_j, q_j + y_j) \\ v(p_j + x_j, q_j + y_j) \end{bmatrix} \quad (17)$$

where

$$\begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (18)$$

$$\begin{bmatrix} x_j \\ y_j \end{bmatrix} = \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix} \begin{bmatrix} -x \\ y \end{bmatrix} \quad (19)$$

and x and y are independent parameters. We compare the coefficients of the Taylor series expansions of fields $u(\cdot)$ and $v(\cdot)$ up to the first order. By eliminating sine and cosine terms on the L.H.S., and appending the error terms, we get (4)–(9).

REFERENCES

- [1] M. A. Arbib, Ed., *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press, 1995.
- [2] H. Weyl, *Symmetry*. Princeton, NJ: Princeton Univ. Press, 1952.
- [3] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, "Comparing images using the Hausdorff distance," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, pp. 850–863, Aug. 1993.
- [4] D. Chetverikov and A. Lerch, "A multiresolution algorithm for rotation-invariant matching of planar shapes," *Pattern Recognit. Lett.*, vol. 13, pp. 669–676, 1992.
- [5] V. S. Nalwa, "Line-drawing interpretation: Bilateral symmetry," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, pp. 1117–1120, Oct. 1989.
- [6] H. J. A. M. Heijmans and A. V. Tuzikov, "Similarity and symmetry measures for convex shapes using Minkowski addition," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 20, pp. 980–993, Sept. 1998.
- [7] F. Mokhtarian and A. Mackworth, "Scale-based description and recognition of planar curves and two-dimensional shapes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 8, pp. 34–43, Jan. 1986.

- [8] F. P. Kuhl and C. R. Giardina, "Elliptic Fourier features of a closed contour," *Comput. Vis., Graph. Image Processing*, vol. 18, pp. 236–258, 1982.
- [9] E. M. Arkin, L. P. Chew, D. P. Huttenlocher, K. Kedem, and J. S. B. Michell, "An efficiently computable metric for comparing polygonal shapes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, pp. 209–216, Feb. 1991.
- [10] D. Ballard and C. Brown, *Computer Vision*. Englewood Cliffs, NJ: Prentice-Hall, 1982.
- [11] C. C. Lin and R. Chellappa, "Classification of partial 2D shapes using Fourier descriptors," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 9, pp. 686–690, May 1987.
- [12] E. Persoon and K. S. Fu, "Shape discrimination using Fourier descriptors," *IEEE Trans. Syst., Man, Cybern. A*, vol. 7, pp. 170–179, Jan. 1977.
- [13] P. S. Mark, H. Rom, and G. Medioni, "B-spline contour representation and symmetry detection," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, pp. 1191–1197, Nov. 1993.
- [14] H. Zabrodsky, S. Peleg, and D. Avnir, "Symmetry as a continuous feature," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 17, pp. 1154–1166, Dec. 1995.
- [15] K. Kamatani, "Comments on 'Symmetry as a continuous feature'," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 19, pp. 246–247, Feb. 1995.
- [16] R. L. Ogniewicz and O. Kubler, "Hierarchical Voronoi skeletons," *Pattern Recognit.*, vol. 28, pp. 1839–1844, Dec. 1995.
- [17] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. W. Zucker, "The Hamilton-Jacobi transform," in *Proc. Int. Conf. Computer Vision*, vol. 2, Corfu, Greece, 1999, pp. 828–834.
- [18] H. Shroff and J. Ben-Arie, "Finding shape axes using magnetic fields," *IEEE Trans. Image Processing*, vol. 8, pp. 1388–1394, Oct. 1999.
- [19] C. Xu and J. L. Prince, "Snakes, shapes, and gradient vector flow," *IEEE Trans. Image Processing*, vol. 7, pp. 359–369, Mar. 1998.
- [20] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.*, vol. 17, pp. 185–203, 1981.
- [21] S. N. Gupta and J. L. Prince, "Stochastic models for DIV-CURL optical flow methods," *IEEE Signal Processing Lett.*, vol. 3, pp. 32–35, Jan. 1996.
- [22] R. Courant and D. Hilbert, *Methods of Mathematical Physics*. New York: Interscience, 1953, vol. 1.
- [23] R. C. Veltkamp and M. Hagedoorn, "State of the art in shape matching," Utrecht Univ., Utrecht, The Netherlands, Tech. Rep. UU-CS-1999-27, 1999.
- [24] D. C. W. Pao, H. F. Li, and R. Jayakumar, "Shapes recognition using the straight line Hough transform: Theory and generalization," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, pp. 1076–1089, Nov. 1992.
- [25] G. Sapiro, "Color and illuminant voting," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 21, pp. 1210–1215, Nov. 1999.



V. Shiv Naga Prasad received the B.Tech. and M.Tech. degrees from the Department of Computer Science and Engineering, Indian Institute of Technology (IIT)—Madras, Chennai, in 2003. He is currently pursuing the Ph.D. degree at the Institute of Advanced Computer Studies, University of Maryland, College Park.

His current research interests include gesture recognition and perceptual organization.



B. Yegnanarayana was born in India in 1944. He received the B.E., M.E., and the Ph.D. degrees in electrical communication engineering from the Indian Institute of Science (IIS), Bangalore, in 1964, 1966, and 1974, respectively.

He was a member of the faculty in the Department of Electrical Communication Engineering, IIS, from 1966 to 1978. From 1978 to 1980, he was a Visiting Associate Professor of computer science at Carnegie-Mellon University, Pittsburgh, PA. Since 1980, he has been a Professor in the Department of Computer Science and Engineering, Indian Institute of Technology—Madras, Chennai. He is the author of the book *Artificial Neural Networks* (New Delhi, India: Prentice-Hall, 1999). His current research interests are in signal processing, speech, vision, and neural networks, and he has published papers in reviewed journals in these areas.

Dr. Yegnanarayana is a Fellow of the Indian National Science Academy, the Indian National Academy of Engineering, and the Indian Academy of Sciences. He is also an Associate Editor of the IEEE TRANSACTIONS ON SPEECH AND AUDIO PROCESSING.