A Fast Radial Symmetry Transform for Detecting Points of Interest

Gareth Loy and Alexander Zelinsky

Australian National University, Canberra 0200, Australia,
{gareth, alex}@syseng.anu.edu.au
http://www.syseng.anu.edu.au/~gareth

Abstract. A new feature detection technique is presented that utilises local radial symmetry to identify regions of interest within a scene. This transform is significantly faster than existing techniques using radial symmetry and offers the possibility of real-time implementation on a standard processor. The new transform is shown to perform well on a wide variety of images and its performance is tested against leading techniques from the literature. Both as a facial feature detector and as a generic region of interest detector the new transform is seen to offer equal or superior performance to contemporary techniques whilst requiring drastically less computational effort.

1 Introduction

Automatic detection of points of interest in images is an important topic in computer vision. Point of interest detectors can be used to selectively process images by concentrating effort at key locations in the image, they can identify salient features and compare the prominence of such features, and real-time interest detectors can provide attentional mechanisms for active vision systems [11].

In this paper a new point of interest operator is presented. It is a simple and fast gradient-based interest operator which detects points of high radial symmetry. The approach was inspired by the results of the generalised symmetry transform [8, 4, 9], although the method bares more similarity to the work of Sela and Levine [10] and the circular Hough transform [5, 7]. The approach presented herein determines the contribution each pixel makes to the symmetry of pixels around it, rather than considering the contribution of a local neighbourhood to a central pixel. Unlike previous techniques that have used this approach [5, 7, 10] it does not require the gradient to be quantised into angular bins, the contribution of *every* orientation is computed in a single pass over the image. The new method works well with a general fixed parameter set, however, it can also be tuned to exclusively detect particular kinds of features. Computationally the algorithm is very efficient, being of order $O(KN)$ when considering local radial symmetry in $N \times N$ neighbourhoods across an image of K pixels.

Section 2 of this paper defines the new radial symmetry transform. Section 3 discusses the application of the transform, including selection of parameters and

some additional refinements. Section 4 shows the performance of the new transform on a variety of images, and compares it to existing techniques, and Section 5 presents the conclusions.

2 Definition of the Transform

The new transform is calculated over a set of one or more ranges N depending on the scale of the features one is trying to detect. The value of the transform at range $n \in N$ indicates the contribution to radial symmetry of the gradients a distance n away from each point. Whilst the transform can be calculated for a continuous set of ranges this is generally unnecessary as a small subset of ranges is normally sufficient to obtain a representative result.

At each range n an *orientation projection image* O_n and a *magnitude projection image* M_n are formed. These images are generated by examining the gradient $\mathbf{g}$ at each point $\mathbf{p}$ from which a corresponding *positively-affected pixel* $\mathbf{p}_{+ve}(\mathbf{p})$ and *negatively-affected pixel* $\mathbf{p}_{-ve}(\mathbf{p})$ are determined, as shown in Figure 1. The *positively-affected pixel* is defined as the pixel that the gradient vector $\mathbf{g}(\mathbf{p})$ is pointing to, a distance n away from $\mathbf{p}$, and the *negatively-affected pixel* is the pixel a distance n away that the gradient is pointing directly away from.

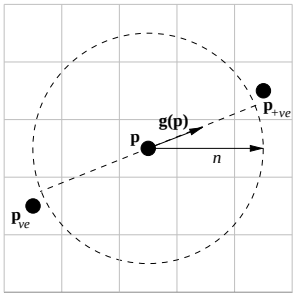

Fig. 1. The locations of pixels $\mathbf{p}_{+ve}(\mathbf{p})$ and $\mathbf{p}_{-ve}(\mathbf{p})$ affected by the gradient element $\mathbf{g}(\mathbf{p})$ for a range of $n = 2$. The dotted circle shows all the pixels which can be affected by the gradient at $\mathbf{p}$ for a range n.

The coordinates of the positively-affected pixel are given by

$$\mathbf{p}_{+ve}(\mathbf{p}) = \mathbf{p} + \text{round}\left(\frac{\mathbf{g}(\mathbf{p})}{\|\mathbf{g}(\mathbf{p})\|}n\right)$$

while those of the negatively-affected pixel are

$$\mathbf{p}_{-ve}(\mathbf{p}) = \mathbf{p} - \text{round}\left(\frac{\mathbf{g}(\mathbf{p})}{\|\mathbf{g}(\mathbf{p})\|}n\right)$$

where 'round' rounds each vector element to the nearest integer.

The orientation and projection images are initially zero. For each pair of affected pixels the corresponding point $\mathbf{p}_{+ve}$ in the orientation projection image O_n and magnitude projection image M_n is incremented by 1 and $\|\mathbf{g}(\mathbf{p})\|$ respectively, while the point corresponding to $\mathbf{p}_{-ve}$ is decremented by these same quantities in each image. That is

$$O_n(\mathbf{p}_{+ve}(\mathbf{p})) = O_n(\mathbf{p}_{+ve}(\mathbf{p})) + 1$$

$$O_n(\mathbf{p}_{-ve}(\mathbf{p})) = O_n(\mathbf{p}_{-ve}(\mathbf{p})) - 1$$

$$M_n(\mathbf{p}_{+ve}(\mathbf{p})) = M_n(\mathbf{p}_{+ve}(\mathbf{p})) + ||\mathbf{g}(\mathbf{p})||$$

$$M_n(\mathbf{p}_{-ve}(\mathbf{p})) = M_n(\mathbf{p}_{-ve}(\mathbf{p})) - ||\mathbf{g}(\mathbf{p})||$$

The radial symmetry contribution at a range n is defined as the convolution

$$S_n = F_n * A_n \tag{1}$$

where

$$F_n(\mathbf{p}) = ||\tilde{O}_n(\mathbf{p})||^{(\alpha)}\tilde{M}_n(\mathbf{p}), \tag{2}$$

$$\tilde{O}_n(\mathbf{p}) = \frac{O_n}{\max_{\mathbf{p}}\{||O_n(\mathbf{p})||\}},$$

$$\tilde{M}_n(\mathbf{p}) = \frac{M_n}{\max_{\mathbf{p}}\{||M_n(\mathbf{p})||\}},$$

α is the radial strictness parameter, and A_n is a two-dimensional Gaussian. These parameters are discussed in more detail in Section 3.

The full transform is defined as the sum of the symmetry contributions over all the ranges considered,

$$S = \sum_{n \in N} S_n \tag{3}$$

If the gradient is calculated so it points from dark to light then the output image S will have positive values corresponding to bright radially symmetric regions and negative values indicating dark symmetric regions (see Figure 2 for example).

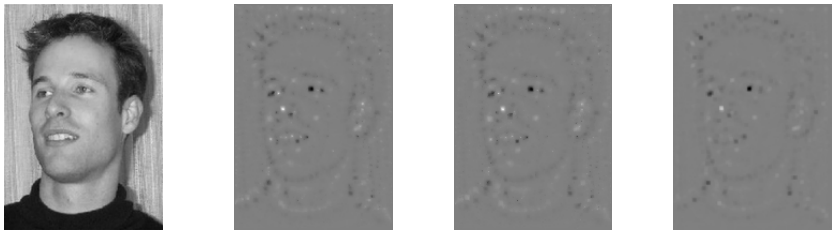

(a) original image (b) $\sum_{n=1}^{5} S_n$ (c) $\sum_{n=1,3,5} S_n$ (d) S_5

Fig. 2. Varying the set of ranges N.

3 Applying the Transform

In order to apply the transform there are a number of parameters that must first be defined, namely, the a set of ranges $N = \{n_1, n_2, ...\}$ at which to calculate S_n, the Gaussian kernels A_n, and the radial strictness parameter α. Some additional refinements are also considered, including ignoring small gradient elements, and only searching for dark or light radially symmetric regions.

The traditional approach to local symmetry detection [3, 8, 10] is to calculate the symmetry apparent in a local neighbourhood about each point. This can be achieved by calculating S_n for a continuous set of ranges $N = \{1, 2, ..., n_{max}\}$ and combining using equation 3. However, since the symmetry contribution is calculated independently for each range n it is simple to determine the result at a single range, or an arbitrary selection of ranges that need not be continuous. Furthermore, the results obtained by examining a representative subset of ranges give a good approximation of the output obtained by examining a continuous selection of ranges, while saving on computation.

Figure 2 shows the combined output S calculated for a continuous range of n from 1 to 5 (b) is closely approximated by combining only $n = 1$, 3 and 5 (c). Also, if the scale of a radially symmetric feature is know *a priori* then the feature can be efficiently detected by only determining the transform at the appropriate range, this is demonstrated by the effective highlighting of the eyes (that have radius 5 pixels) by S_5 in Figure 2 (d).

The purpose of the Gaussian kernel A_n is to spread the influence of the positively- and negatively-affected pixels as a function of the range n. A two-dimensional Gaussian is chosen because it is radially symmetric so it will have a consistent effect over all gradient orientations, and it is separable so its convolution can be efficiently determined. Figure 3 shows the contribution for a single gradient element $\mathbf{g}(\mathbf{p})$. By scaling the standard deviation linearly with the range n, we define an arc of influence that applies to all affected pixels. The width of the arc is defined by scaling the standard deviation of A_n with respect to n.

The parameter α determines how strictly radial the radial symmetry must be for the transform to return a high interest value. Figure 4 shows the effect of choosing α to be 1, 2 and 3 on S_1 for an image exhibiting strong radial values

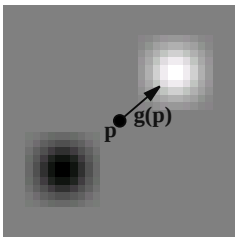

Fig. 3. The contribution of a single gradient element, with A_n chosen to be a 2D Gaussian of size $n \times n$ and standard deviation $\sigma = 0.25n$, and $n = 10$.

around the eyes. Note how a higher α eliminates non-radially symmetric features such as lines. A choice of $\alpha = 2$ is suitable for most applications. Choosing a higher α starts attenuating points of interest, whilst a lower α gives too much emphasis to non-radially symmetric features, however, choosing α as 1 minimises the computation when determining F_n in Equation 2.

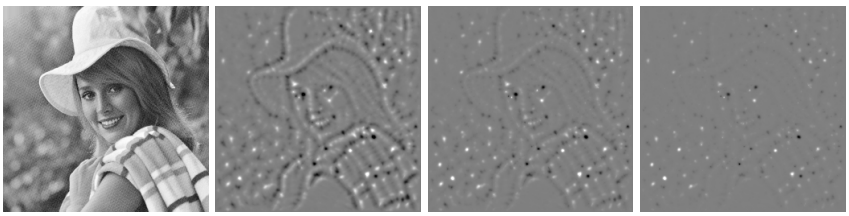

Input image $\alpha = 1$ $\alpha = 2$ $\alpha = 3$

Fig. 4. Effect of varying α. Original image from USC-SIPI Image Database [1]

Gradient elements with small magnitudes have less reliable orientations, are more easily corrupted by noise, and tend to correspond to features that are not immediately apparent to the human eye. Since the purpose of the transform is to pick out points of interest in the image it is logical to ignore such elements in our calculation. A gradient threshold parameter β is introduced for this purpose, and when calculating images O_n and M_n all gradient elements whose magnitudes are below β are ignored. The effect of a non-zero β is shown in Figure 5. The main advantage of a non-zero β is an increase in the speed of the algorithm, since there are less gradient elements considered, and hence less affected pixels to be calculated.

The transform can be tuned to look only for dark or bright regions of symmetry. To look exclusively for dark regions, only the negatively-affected pixels need be considered when determining M_n and O_n (see Section 2). Likewise, to detect

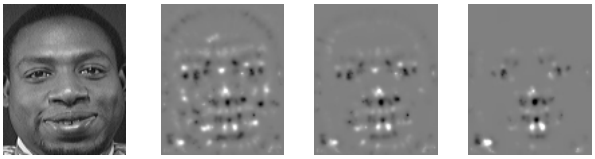

Input image $\beta = 0$ $\beta = 20\%$ $\beta = 40\%$

Fig. 5. The effect of different values of β on S. Here β is measured as a percentage of the maximum possible gradient magnitude. Original image from Database of Faces, AT&T Laboratories Cambridge [2]

bright symmetry only positive affected pixels need be considered. Examples of dark symmetry are shown in Section 4.

4 Performance of the Transform

The performance of the new transform was demonstrated on a range of images and compared with several prominent transforms from the literature.

Figure 6 demonstrates the performance of the new transform on faces and other images. These figures were generated using the parameter settings presented in Table 1, and show how the transform can provide a useful cue for the location of facial features – especially eyes – in face images, as well as highlighting generic points of interest that are characterised by high contrast and radial symmetry.

Table 1. Parameter Settings used for Experimentation

Parameter	Setting		
	Full	Fast	Fast Dark
Set of ranges N	$\{n : n = 1, 2, ..., 6\}$	$\{n : n = 1, 3, 5\}$	$\{n : n = 1, 3, 5\}$
Gaussian kernel			
Volume under kernel	n^2	n^2	n^2
Size	n	n	n
Standard deviation	$0.5n$	$0.5n$	$0.5n$
Radial strictness α	2	2	2
Small gradients ignored	0	20% ignored	20% ignored
Dark symmetry	Yes	Yes	Yes
Bright symmetry	Yes	Yes	No

Figure 7 compares the performance of the transform against existing techniques from the literature. Each transform is applied to the image in the centre of the figure (the standard 256×256 lena image) for which the intuitive points of interest are the eyes.

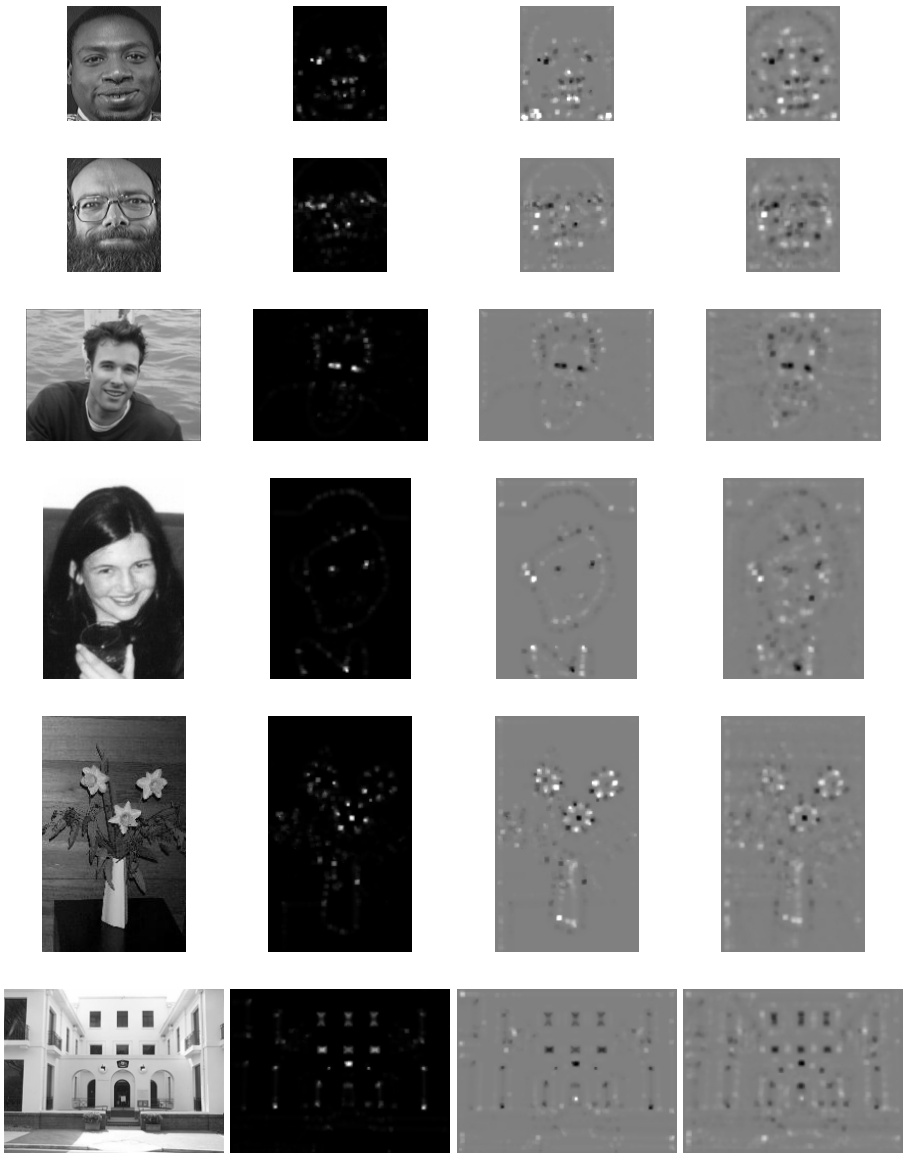

Original image | Radial symmetry (fast dark) | Radial symmetry (fast) | Radial symmetry) (full)

Fig. 6. The new transform applied to a variety of images. The parameter settings are indicated beneath each column and refer to the values detailed in Table 1. The two top most images are from the Database of Faces [2].

All methods were implemented with a local neighbourhood radius of 6 pixels, and where necessary the gradient orientation was quantised into 8 bins.

Each of the transforms was implemented in Matlab. For the majority of the transforms an estimate of the approximate number of floating point operations involved was obtained from Matlab, however, for Di Gesù *et al.*'s discrete symmetry transform and Sela and Levine's real-time attention mechanism this was not feasible. These transforms involve optimised low-level processes that were not practical to implement in Matlab, so the number of operations required is not reported here. It suffices to say that the *non-optimised* implementations used to generate the visual results shown required computation well in excess of the other methods. The estimated computations obtained are presented in Table 2.

The new transform effectively highlights the points of interest (eyes) in Figure 7. Of the existing transforms Reisfeld's generalised (dark and radial) symmetry provide the next best result, and while the other transforms do highlight the eye regions they tend to highlight many other points as well reducing their overall effectiveness.

Table 2. Estimated Computation Required for Different Transforms to compute the results in Figure 7

Transform	Computations (Mflop)
New Transform	
Full	19.7
Fast	7.93
Fast Dark	7.02
Existing Transforms	
Generalised Symmetry	
Radial [8]	259
Dark [9]	179
Circular Hough [7]	33.9

Table 3 lists the theoretical order of computation required to compute the transforms on an image of K pixels, where local symmetry is considered in an $N \times N$ neighbourhood, and for those methods that require gradient quantisation the gradient is quantized into B bins. The complexity $O(KN)$ of the new transform is lower than all other transforms considered, with the possible exception of Di Gesu *et al.*'s Discrete Symmetry Transform that has complexity $O(KB)$. However, as was discussed in Section 3, it is not necessary to calculate the new transform at all ranges $1..N$, so the computational order can be further reduced, whereas it is essential to calculate Di Gesu *et al.*'s Discrete Symmetry Transform across four or more angular bins. Furthermore the results from the Discrete Symmetry Transform do not appear as effective for locating points of interest (see Figure 7).

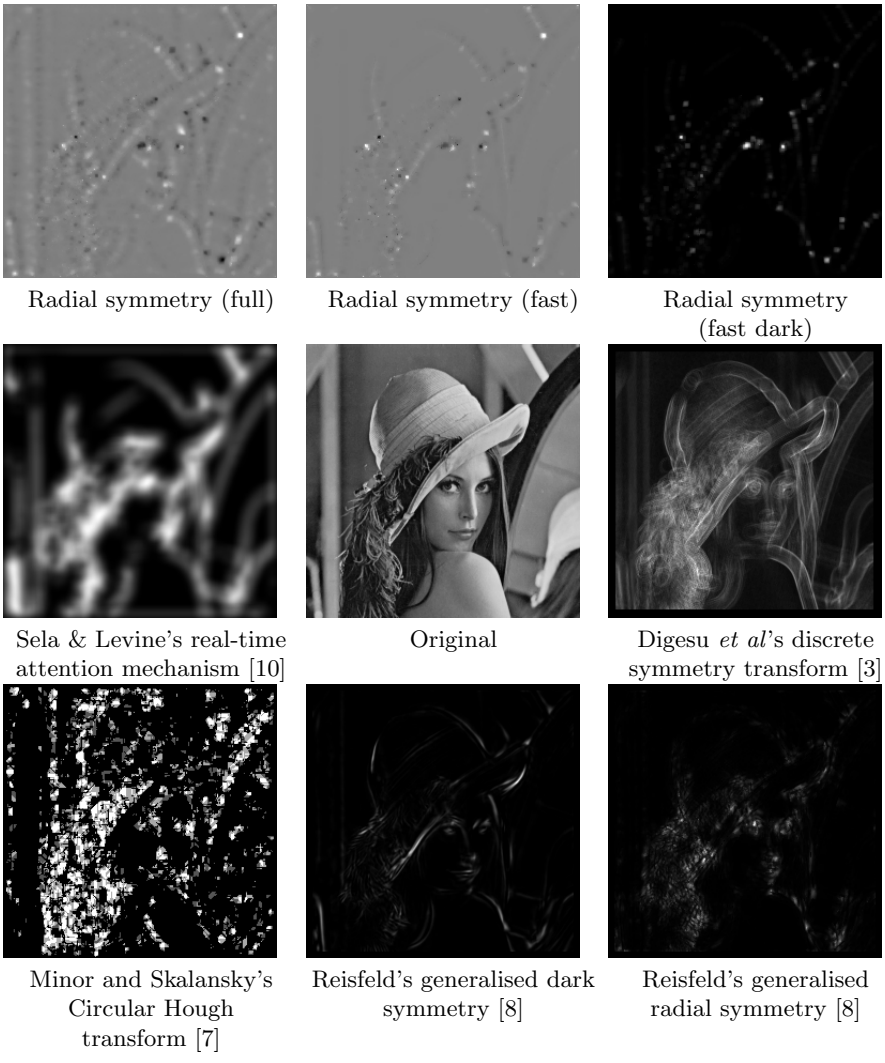

Fig. 7. Comparison of new transform (top row) with other available transforms. In order to compare the output of Sela & Levine's real-time attention mechanism with the other transforms, the final step, which involved identifying local maximums in the output as points of interest, has been omitted.

Table 3. Computational Order of Different Transforms

Transform	Order
New Radial Symmetry Transform	KN
Reisfeld's Generalised Symmetry Transform [8]	KN^2
Lin and Lin's Gradient-based Inhibitory Mechanism [6]	KN^2
Di Gesu *et al.*'s Discrete Symmetry Transform [3]	KB
Sela and Levine's Real-Time Attentional Mechanism [10]	KBN
Circular Hough Transform [7]	KBN

Figure 7 and Tables 2 and 3 demonstrate that the new transform can provide comparable or superior results to existing techniques whilst requiring significantly less computation and complexity.

The key to the speed of the new transform lies in the use of *affected pixels* to project the effect of gradient elements. This allows an approximation of the effect of each gradient element on the radial symmetry of the pixels around it, without specifically considering neighbourhoods about each point, as did [6, 8], or requiring multiple calculations for different gradient orientations, as did [3, 7, 10].

Unlike other transforms the fast symmetry transform differentiates between dark and bright regions of radial symmetry, while allowing both to be computed simultaneously. Alternatively just dark (or bright) points of symmetry can be considered exclusively with an associated reduction in computation.

5 Conclusion

A novel point of interest detector has been presented that uses the gradient of an image to locate points of high radial symmetry. The method has been demonstrated on a series of face images and other scenes, and compared against a number of contemporary techniques from the literature. As a point of interest operator the new transform provides equal or superior performance on the images tested while offering significant savings in both the computation required and the complexity of the implementation. The efficiency of this transform makes it well suited to real-time vision applications.

References

1. The USC-SIPI Image Database, University of Southern California Signal & Image Processing Institute. *http:// sipi.usc.edu/ services/ database/ Database.html*.
2. Database of Faces. AT&T Laboratories Cambridge, *http:// www.cam-orl.co.uk/ facedatabase.html*.
3. V. Di Gesù and C. Valenti. The discrete symmetry transform in computer vision. Technical Report DMA 011 95, Palermo University, 1995.
4. N. Intrator, D. Reisfeld, and Y. Yeshurun. Extraction of facial features for recognition using neural networks, 1995.

5. C. Kimme, D. Ballard, and J. Sklansky. Finding circles by an array of accumulators. *Communications of the Association for Computing Machinery*, 18(2):120–122, February 1975.
6. Cheng-Chung Lin and Wei-Chung Lin. Extracting facial features by an inhibitory mechanism based on gradient distributions. *Pattern Recognition*, 29(12):2079–2101, 1996.
7. L. G. Minor and J. Sklansky. Detection and segmetaion of blobs in infrared images. *IEEE Tranactions on Systems Man and Cyberneteics*, SMC-11(3):194–201, March 1981.
8. D. Reisfeld, H. Wolfson, and Y. Yeshurun. Context free attentional operators: the generalized symmetry transform. *International Journal of Computer Vision, Special Issue on Qualitative Vision*, 14:119–130, 1995.
9. D. Reisfeld and Y. Yeshurun. Preprocessing of face images: Detection of features and pose normalisation. *Computer Vision and Image Understanding*, 71(3):413–430, September 1998.
10. Gal Sela and Martin D. Levine. Real-time attention for robotic vision. *Real-Time Imaging*, 3:173–194, 1997.
11. Orson Sutherland, Harley Truong, Sebastien Rougeaux, and Alexander Zelinsky. Advancing active vision systems by improved design and control. In *Proceedings of International Symposium on Experimental Robotics (ISER2000)*, December 2000.