Non-linear voting in the space variant Hough transform

Nick Barnes

National ICT Australia, Locked Bag 8001, Canberra, ACT 2601, AUSTRALIA.
Nick.Barnes@nicta.com.au

Abstract. For space variant sensors the volume of world space that projects onto a single sensing element varies significantly with its spatial position in the sensing array. This includes the log-polar sensor, and central catadioptric cameras. In this paper, we consider the Hough transform for space variant sensors. We show that for real discrete images the voting equation must be modified to allow for the fact that a larger sub-space of parallel world planes can intersect with a larger sensing element than with a smaller one. We derive a voting algorithm for the space variant Hough Transform, and apply this algorithm to the log-polar sensor. We compare the standard voting algorithm with the proposed non-linear voting scheme for the space variant log-polar Hough transform. We demonstrate experimentally that for the standard Hough transform, lines that cover the majority of the sensor do not register a significantly larger vote count than shorter lines, while they are gain full support in the space variant transform. Further, in real images we show that this can lead to poor definition of peaks in the regular Hough transform, while similar peaks are well defined in the space variant Hough transform. Experiments are conducted on simulated images, and images from a true log-polar sensor.

1 Introduction

The Hough Transform originated with a paper [4], and a patent in 1962 [5]. Since this time it has been the topic of a great number of research papers and is one of the most widely used algorithms in image processing. The standard Hough transform is an algorithm for finding straight lines in an image. It sums the number of sensing elements that are in support of a particular line, defined in a uniformly discretised parametric Hough space. The usual form of this is defined in polar coordinates, as described by Duda and Hart [1]. The vast majority of research on the Hough transform has treated sensors that use a uniform distribution of sensing elements. To be more specific, in the absence of effects such as radial distortion, each sensing element is the projection of a volume of world space, and that volume is assumed to be the same for all sensing elements.[1] However, there is another class of sensors where the volume of world space projected onto a single

[1] With effects such as optical vignetting [11] this is not quite true, but is a reasonable approximation in contrast to true space variant sensors.

sensing element varies significantly with its spatial position in the sensing array. Space-variant sensors defined in this manner, include sensors where the sensing elements are placed in a space-variant distribution, such as the log-polar sensor [8]. In the log-polar sensor, the resolution varies logarithmically with distance from the fovea. Another group of sensors that can be included in this class are cameras that use distorted lenses to achieve a similar effect, (e.g., [7]), and some members of the class of catadioptric cameras that combine the use of a mirror and a lens, typically to achieve an omni-directional view. Of interest here are the members that give a space variant projection, including for example, central catadioptric cameras which can be modelled as having a spherical projection [3]. These include parabolic, hyperbolic and elliptical mirror-based cameras. In this paper, we address sensors that have significant variation in the size of sensing elements. By significant we mean that there are sensing elements that are greater than four times the area of some others in the sensor.

The importance of line-finding in computer vision is clear from the number of papers on the topic, and the Hough transform remains one of the most commonly used algorithms for this function. The Hough transform is particularly attractive due to its low computational cost facilitating real-time operation. Lines are fundamental feature in images, and can be useful for a vast array of manufactured objects, but can also be the basis of calibration algorithms [3]. These topics are just as relevant to space variant sensors as to space invariant sensors.

For space-variant sensors, the topic of the Hough transform has certainly been addressed. The log-hough transform was proposed for the log-polar camera [12], and authors have dealt with the projection of straight-lines in central catadioptric cameras [3]. A Hough transform has also been proposed for some non-single view point cameras [2]. However, these papers all deal with the projection of straight-lines in a continuous space geometric sense, and ignore the issues of sampling in the sensors. In real sensors a pixel does not represent an infinitely small impulse. Also, when an edge is detected as passing though a pixel, this does not imply that the edge has passed through the centre of that sensing element. The family of world planes that intersect a sensing element is not a pencil, but a larger set that pass through the area of space covered by the sensing element. This means that when the size of sensing elements varies, the set of planes passing through the element also varies. In this paper, we address the consequences of space varying sensor elements on the Hough transform. We show that the vote space for a sensing element of a space variant sensor not always one dimensional, but in places has finite width. Finally, we give an efficient algorithm for finding the space variant Hough Transform, and demonstrate an improvement in results for detecting lines on a space-variant sensor.

2 The Hough Transform

The Hough transform is a map from points in an image to sinusoidal curves in Hough space. In the ideal case, this curve represents the family of straight-lines that could pass through the point, which arise from a pencil of world planes

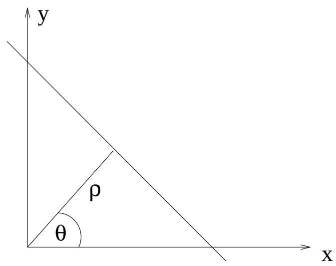

Fig. 1. ρ is the closest distance to the origin, and θ the slope of the normal to the line.

that intersect that point. The standard form of the transform is the equation of a straight-line, parameterised in polar coordinates:

$$\rho = x_i \cos\theta + y_i \sin\theta, \tag{1}$$

where (x_i, y_i) are the pixel coordinates, and ρ and θ are defined as per Figure 1.

Using this transform, any edge point that appears in the image votes for all lines that could possibly pass through that point. In this way, if there is a real line in the image, by transforming all of its points to Hough space, we will accumulate a large number of votes for the actual line, and only one for any other line (assuming no noise) in the image.

3 Edges in finite size sensing element arrays

However, in a real image, the image space is not continuous, but made up of a discrete set of sensing elements, each of which samples light over a finite area. Further, Hough space is represented as a discrete set of samples over (ρ, θ) space.

Following the definition by [6], in the typical discrete case, standard Hough space is an accumulator array, where θ is discretised between $\frac{-\pi}{2}$ and π. Suppose that it is discretised in intervals of $\frac{1}{sf_\theta}$, then:

$$\theta \epsilon [\frac{-\pi}{2}, \frac{-\pi}{2} + \frac{1}{sf_\theta}, ..., \pi] \tag{2}$$

Further ρ is discretised between ρ_{min} and ρ_{max} in intervals of $\frac{1}{sf_\rho}$:

$$\rho \epsilon [\rho_{min}, \rho_{min} + \frac{1}{sf_\rho}, ..., \rho_{max}] \tag{3}$$

In the Hough Transform algorithm, we step through the discrete elements of θ for each edge pixel (x_i, y_i), and calculate the closest ρ from Equation (1). This means there may be multiple votes for a single ρ value, but there will be at most one vote per θ. As there is a finite spacing between the pixels, the region of space over which a line will be classified as passing through the ρ dimension of a quanta that is centred at, ρ_{cm}, can be defined as:

$$\rho_{cm} - 0.5 \leq \rho < \rho_{cm} + 0.5 \tag{4}$$

4 The log-polar sensor

Schwartz [9, 10] derived an analytical formulation of biological vision systems based on experimental measures of the mapping from the retina to the visual cortex of monkeys. Visual data is transformed from the retinal plane in polar coordinates (ρ, θ) to log-polar Cartesian coordinates (ξ, γ) in the cortical plane. The relation can be expressed as:

$$\xi = log_a \frac{\rho}{\rho_0}, \gamma = q\eta, \tag{5}$$

where (ρ, η) are the polar coordinates of a point on the retinal plane and ρ_0, q, and a are constants determined by the physical layout of the sensor. Thus, sensing elements appear in a non-uniform distribution, with a high density at the central fovea, and continuously decreasing density toward the periphery.

5 The log-polar Hough transform

The parametric equation of a line, with slope θ, and where its closest distance to the origin is r is:

$$\rho = \frac{r}{sin(\gamma - \theta)}, 0 < \gamma - \theta < \pi, \tag{6}$$

where r and θ are the polar parameters of the line, for polar image points (γ, ρ). Substituting the equations for the log-polar sensor from Equation (5), we obtain:

$$\xi = log_a(\frac{r}{\rho_0}) - log_a(\sin q\eta - \theta) \tag{7}$$

This is essentially the same as the log-Hough transform [12], with the introduction of real sensor constants a, q, and ρ_0. The Hough space here is the standard, non-space variant Hough space, with uniform quantisation of ρ and γ. In moving to a digital image, we must consider how the space should be sampled. Consider this intuitively, for a line that only passes through coarsely quantised sensing elements, a single quanta of Hough space may be quite large, as we cannot determine its parameters to a very fine resolution. Whereas for a line that passes through finely quantised sensing elements we will be able to place a tighter constraint on its parameters.

Consider the family of space varying sensors where the spatial resolution is higher at the centre of the image, and varies symmetrically about the image centre, decreasing monotonically. We will call such sensors **centre foveated**

sensors. Centre foveated sensors include central catadioptric cameras. For a centre foveated space variant sensor, the natural Hough transform coordinate system is from the image centre (as opposed to the fact that a corner is typically used for Cartesian sensors). By placing the origin of space-variant Hough space at the image centre for such images, we can take advantage of symmetry, and consider the transform as a one dimensional problem. For such a Hough transform, a sensing element in the camera at (ρ_1, θ_1) can vote only for lines for which $\rho \leq \rho_1$. Thus, for any point in Hough space (ρ_a, θ_a), the smallest sensing element that can vote for it is the one nearest to (ρ_a, θ_a). As Hough quanta towards the fovea are supported by sensing elements that are closer together, we can obtain a tighter bound on the parameters of lines that pass closer to the centre of the sensor. Thus, for any central foveated sensor the resolution of quantisation of Hough space must be maximal near the fovea, but can be less towards the periphery.

It can be demonstrated that an efficient invertible Hough transform (in the sense of [6]) is formed by using a space varying r, specifically that a linear mapping that is sufficiently dense to be invertible for small values of r will be more dense than necessary for large values of r, and that r should vary with ξ. We define the space varying log-polar Hough space as having $r^{'} = log_a \frac{r}{\rho_0}$. Substituting into Equation (7) we form the space varying log-polar Hough transform:

$$\xi = r^{'} - log_a(\sin \frac{\eta}{q} - \theta) \tag{8}$$

The voting equation for this is:

$$r^{'} = \xi + log_a(\sin \frac{\eta}{q} - \theta) \tag{9}$$

We also define that the quantisation of the space variant log-polar Hough transform quantisation precisely matches the placement of sensing elements. We do not present any proof of this transform or sampling as they are not necessary for the argument of the paper, and do not change the validity of subsequent derivations, but we will use Equation (8) as the Hough transform for this sensor.

6 Space variation and the Hough transform

Figure 2 represents Hough space across a polar sensor, with even spacing in the ρ dimension, but a uniform quantisation of θ for every radii. In such a sensor, elements that are further from the centre will cover a larger spatial area than elements that are closer to the centre. Two sensing elements have been highlighted. It can be seen that there are more distinct ρ, θ curves passing through the outer element than the inner. Given a uniform quantisation of (ρ, θ), it is intuitively clear that more curves will pass through a larger area of space than a smaller one. Thus, all space variant sensors must consider the possibility of a family of parallel curves from smaller sensing elements passing through a single larger sensing element. For central foveated sensors, this implies that an edge through an outer sensing element votes on many parallel curves in the inner part of Hough space.

Consider two parallel world planes that project onto sensing elements that are at the outer part of the image as well as the inner. Also, that these planes are close, and project onto a single pixel at an outer part of the image. Consider the dimension of world sensing element projection that is orthogonal to the planes. The angle subtended by a sensing element at the inner part of the image is smaller than at the outer part of the image. Thus, depending on the size difference between the pixels, although the planes fall within a single pixel at the lower resolution part of the sensor, it may be that they fall on different pixels at the inner part. For a non-space variant sensor, it may be that the two planes project to neighbouring pixels due to aliasing effects in sampling. However, for a space-variant sensor, they may fall many pixels apart, depending on the size difference of the two-dimensional angle projected onto the the pixel. We now derive this result formally for the space variant log-polar Hough transform.

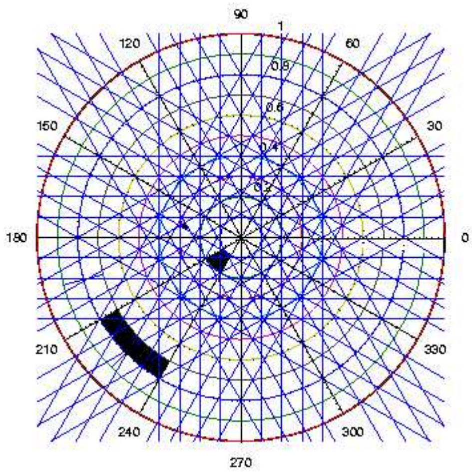

Fig. 2. More Hough lines pass through outer than inner sensing elements.

Consider Figure 3, let the closest point on the inner line pass through (r_1, θ_0), this being its Hough parameterisation. Similarly, let the second, parallel, line pass through (r_2, θ_0). Let the line (r_1, θ_0) also pass through the centre of an outer sensing element (ξ_k, γ). Hence, for the log-polar sensor, we have:

$$\xi_k = r_1 - log_a(\sin(\gamma - \theta_0)) \qquad (10)$$

At γ, the curve (r_2, θ_0) passes through the point ξ_p defined:

$$\xi_p = r_2 - log_a(\sin(\gamma - \theta_0)) \tag{11}$$

For the curve defined by (r_2, θ_0) to pass through the sensing element (ξ_k, γ), it is sufficient that:

$$\xi_k < \xi_p < \xi_k + (\xi_{k+1} - \xi_k)/2, \tag{12}$$

where ξ_{k+1} is the ξ value for the neighbouring outer sensing element with the same value of γ. $(\xi_{k+1} - \xi_k)/2$ is the width from the centre to the outer edge of the sensing element in the ξ dimension.

As we have $\rho_1 < \rho_2$, trivially, $\xi_k < \xi_p$. Subtracting Equation (10) from (11), we obtain:

$$\xi_p = \xi_k - (\rho_1 - \rho_2) \tag{13}$$

Substituting into Equation (12):

$$\xi_k - (\rho_1 - \rho_2) < (\xi_k + (\xi_{k+1} - \xi_k)/2 \tag{14}$$

Thus,

$$2(\rho_2 - \rho_2) < \xi_{k+1} - \xi_k \tag{15}$$

Thus, if the retinal spacing between any pair of Hough quantisations at the same value of θ is less than double the spacing between a pair of sensing elements of the same value of γ at an outer part of the sensor, and the inner Hough quantisation passes through the centre (or to the inside of it) of the sensing element (ξ_k, γ) then both curves will pass through sensing element (ξ_k, γ). This means for typical space variant sensors where the above condition is satisfied, that a single sensing element in the outer part of a sensor may vote on multiple Hough quantisations from the inner part of the sensor. For the log-polar sensor, the distance between the outer intersections in retinal space (as opposed to the logarithmic mapping) is: $\rho_0(a^{\xi_{k+1}} - a^{\xi_k})$, which will be substantially larger than the inner spacing for any sufficiently large k, where $a > 1$. In the Giotto sensor, more than 20 parallel lines passing through separate inner sensing elements pass through a single outer sensing element.

We may define a space variant Hough transform by Equation (15) in conjunction with Equation (8). However, a more elegant method is to define the inverse of the space variant Hough transform, and use this to compute a lookup table for space variant Hough space instead of a function. In practice, a lookup table is a more computationally efficient implementation for any Hough transform. For log-polar space, the sensor is radially symmetric about the fovea, so all columns of pixels are identical. Also, translation in ξ yields magnification [13], thus, it can be shown that translating a straight-line radially shifts all the sensing elements that it projects onto in log-polar space by a uniform amount. Thus, we may

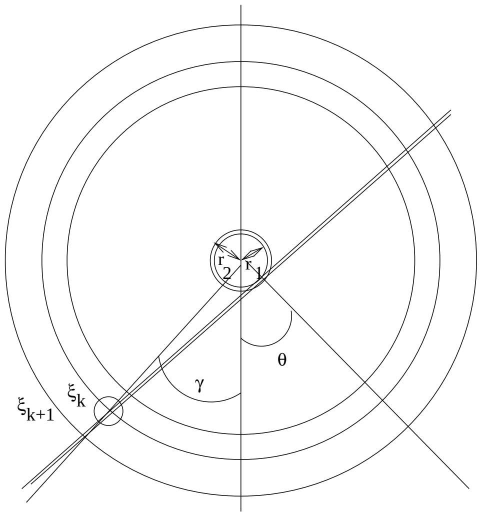

Fig. 3. Two parallel lines passing through inner sensing elements pass through a single outer sensing element.

represent a Hough transform as a single lookup entry for all sensing elements, where the Hough space curve is simply translated by offsets for (ξ, γ).

From Equation (8), the log-polar space-variant Hough transform, we can form the inverse voting equation:

$$\xi = r^{'} - log_a(\sin(\gamma - \theta)) \tag{16}$$

To form the lookup table, take Equation (16) and for each value of θ, $r^{'}$, hold γ to a constant value, and add the $(\theta,\, r^{'})$ parameters to the lookup table for the computed value of ξ. The log-polar space-variant Hough transform for other values of γ can be simply computed by offsetting the lookup value of θ by a corresponding offset to that of γ compared to the constant value.

The same procedure can be used to derive the space-variant Hough transform lookup table for any sensor where pixel resolution monotonically decreases towards the periphery. The above derivation is only valid for the log-polar sensor.

7 Results

We generated a simulated log-polar edge image, based on a model of the Giotto log polar sensor. The angular space of the sensor contains 128 evenly spaced sensing elements, and there are 76 separate circles of elements. The constants used for the simulation were $\rho_0 = 2.72195$, and $a = 1.0528432$. These are the approximate values for the Giotto sensor. In the Giotto sensor, the number of sensing elements is reduced for the inner 20 circles, this was not modelled.

Using this model, we formed edge images where the ρ value for the edge precisely passed through the centre of each sensing element, for a single value of θ. Note that by symmetry, all values of θ produce the same result. These 76 edge images were then transformed using the standard non-space variant Log Hough transform for the log-polar sensor defined in Equation (9), as well as through the space-variant Hough transform log-polar. There was no noise introduced in these images, only a single edge appeared, and it passed through all intermediate pixels from one side of the sensor to the other.

Figure 4 shows the highest peak (verified to always be the correct line) and the second highest peak for both the Space Variant and standard Log Hough Transform. The x-axis is the value of ρ that was the closest passing point of the line - its Hough space ρ parameter. For all perfectly placed, noise free edge images, every edge pixels always voted precisely once for the correct line for the space variant Hough transform (this follows from the definition). It can be seen that the correct line has significantly more votes than any other peak, until outermost edges, where it degrades gracefully. It can also be seen that the standard transform does not take advantage of all the pixels of the edge as it only samples the line at regular intervals of γ, and thus does not consider the possible parallel lines that could have passed through the outer sensing elements. This means that in the standard Hough transform, lines that cover the majority of the sensor do not get a significantly larger vote, as can be clearly seen in Figure 4. This is problematic if the vote is used to select lines that cover the most pixels, as is fundamental to the definition of the Hough transform.

The space variant Hough transform was also tested on a real log-polar image taken by the Giotto camera in which many edges appear. Figure 5 shows the original and edge images as well as the space variant Hough transform. The edge were extracted using the Sobel detector, with thresholding. In this case no ground truth data is available to verify the precise location the edges. For the edges that appear to have the most pixels, the point at the apex of their edge curve in log-polar space was hand determined to be the true location of the edge. This can only be accurate to within a few pixels. There were four lines identified, being a thin line near the top, and three near parallel curves at the bottom. The locations were estimated to be (11, 32), (8, 97), (18,95), and (33,95).

The peaks for each local area for the space variant Hough transform were (7, 32), (6, 96), (16,97), and (33, 96). The values at these peaks were 92, 109, 100, 78 respectively. For the standard Hough Transform the peaks were (7,32), (7,97), (15, 96), and (33,97), with values 30, 52, 55, and 47 respectively. Thus, both the space variant and standard log-polar Hough transform have correctly identified the edges with the most pixels, to within the accuracy of the experiment (although they have slightly different results it is not possible in this case to identify which is more correct). It is clear that the space variant Hough transform has higher peaks than the standard Hough transform. This is because the space variant Hough transform gives votes to the outer pixels that support the

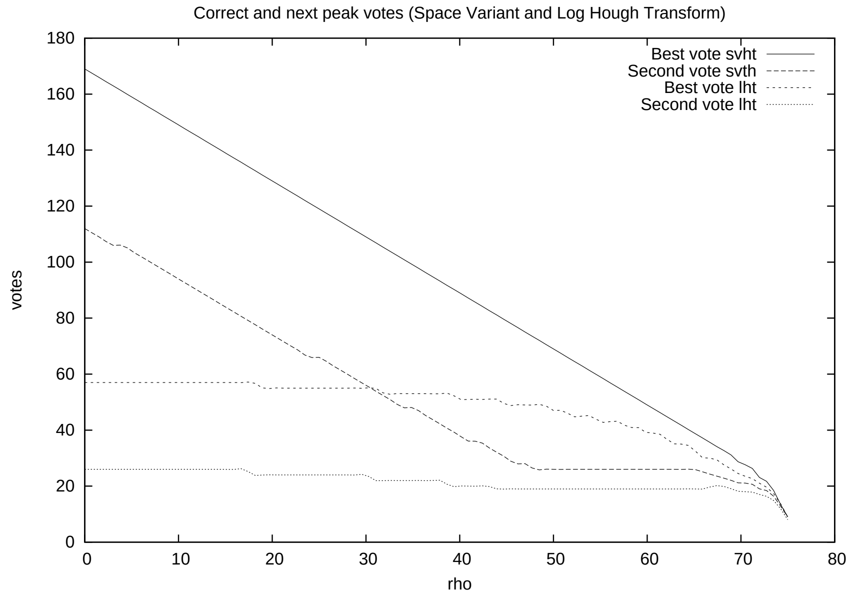

Fig. 4. The highest peak and the next highest for the Space Variant and Log Hough Transforms based on simulated noise-free images. Rho is the log of the radial coordinate of the nearest point on the line to the centre of the sensor.

line, but these are ignored in a Hough transform that does not take space variance into account. Note that we can expect some aliasing in this image as the Sobel edge detector does not thin edges, and as can be seen from the image some of the edges are quite broad. The image is a real indoor scene, and shadowing means some of the edges themselves are quite broad.

Consider Figure 6 which shows the histograms for the two Hough transforms of Figure 5. The vertical lines show the position of the lowest peaks that correspond to clear lines in the transforms, 78 in the space variant, and 30 in the regular. The lowest peak for the regular Hough transform corresponds to the large line at the top of the image. The vote count is low because the line is thin, and a little patchy close to the fovea. The line has much stronger support towards the periphery. For the regular Hough transform it has fallen in a region of the histogram that has a lot of similar votes. Whereas the same line in the space variant Hough transform actually has 90 votes, and is very strongly supported by its outer sensor pixels, well clear of the bulk of other peaks. The 78 votes for an edge in the space variant Hough transform corresponds to the outermost of the three near parallel lines. This has fewer votes because it is smaller, as would be expected. However, 78 is well towards the upper end of the histogram.

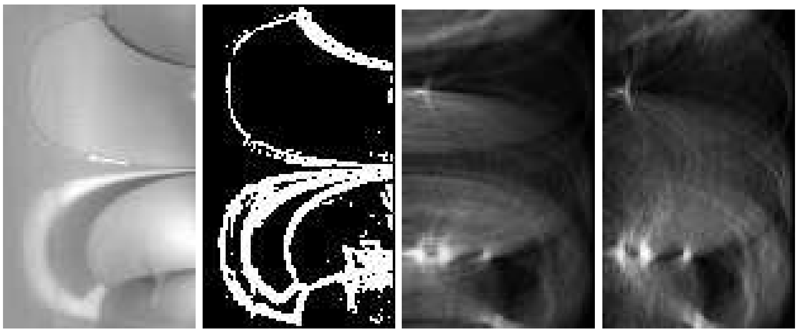

(a) (b) (c) (d)

Fig. 5. Edges and the Hough transforms for a real image from the Giotto sensor. (a) The original log-polar image, (b) edges found by thresholding the Sobel image, (c) the space variant Hough transform of this image, and (d) the standard log-hough transform.

8 Conclusion

We have argued that the voting in a space variant Hough transform should be non-linear. Variation in size of the world space represented by sensing elements means that it is not just a one dimensional family of world planes that can intersect with a larger sensing element. If Hough space is partitioned finely enough to represent all possibly distinguishable lines in the high resolution part of the sensor, a single low-resolution sensing element may vote on a family of parallel lines. We proposed an efficient algorithm for the space variant Hough transform that uses a lookup table generated using the inverse of the voting equation. We derived this specifically on the log-polar sensor, and presented an implementation. We showed that applying the regular Hough transform to log-polar images may ignore support from outer sensing elements for lines close to the fovea. It was demonstrated that this can lead to poor peak definition in real images. This problem does not occur in the space variant log-polar Hough transform.

Acknowledgments

National ICT Australia is funded by the Australian Department of Communications, Information Technology and the Arts and the Australian Research Council through *Backing Australia's ability* and the ICT Centre of Excellence Program.

References

1. R O Duda and P E Hart. *Pattern Classification and Scene Analysis.* Wiley, New York, 1973.

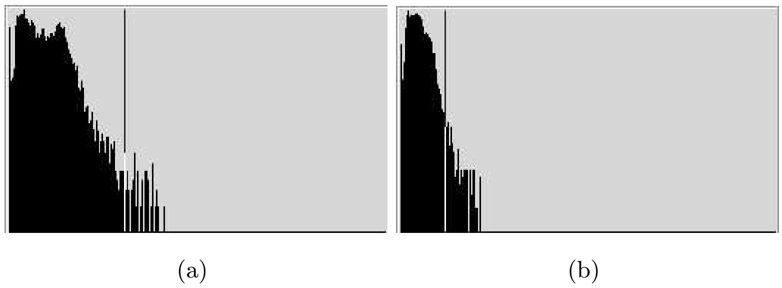

Fig. 6. Histograms for the (a) space variant Hough transform, and (b) the standard log Hough transform from the real image shown in Figure 5. The full length vertical lines in (a) and (b) are at 78 and 30 respectively, the value of the lowest peaks for each transform. It can be seen that in (a) this value is higher than the bulk of the histogram, while in (b) a large number of values fall at a similar level.

2. M Fiala and A Basu. Hough transform for feature detection in panoramic images. *Pattern Recognition Letters*, 23:1863–1874, 2002.
3. C Geyer and K Daniilidis. Catadioptric projective geometry. *International Journal of Computer Vision*, 43:223–243, 2001.
4. P V C Hough. Machine analysis of bubble chamber pictures. In *International Conference on High Energy Accelerators and Instrumentation*. CERN, 1959.
5. P V C Hough. Method and means for recognizing complex patterns, December 1962. U.S. Patent, 3,069,654.
6. A L Kesidis and N Papamarkos. On the inverse hough transform. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(12):1329–1343, Dec. 1999.
7. Y Kuniyoshi, N Kita, K Sugimoto, S Nakamura, and T Suehiro. A foveated wide angle lens for active vision. In *IEEE International Conference on Robotics and Automation*, pages 2982–2985, Nagoya, Japan, 1995.
8. P Questa and G Sandini. Time to contact computation with a space-variant retina-line c-mos sensor. In *Proceedings of the International Conference on Intelligent Robots and Systems*, Osaka, Japan, 1996.
9. E L Schwartz. Spatial mapping in the primate sensory projection: Analytical structure and the relevance to perception. *Biological Cybernetics*, 25:181–194, 1977.
10. E L Schwartz. A quantitive model of the functional architecture of human striate cortex with application to visual illustration and cortical texture analysis. *Biological Cybernetics*, 37:63–76, 1980.
11. J Strong. *Concepts of Classical Optics*. W H Freeman and Company, San Francisco: USA, 1958.
12. C F R Weiman. Polar exponential sensor arrays unify iconic and hough space representation. In *Proceedings SPIE:Intelligent Robots and Computer Vision VIII:Algorithms and Techniques*, pages 832–842, 1989.
13. C F R Weiman and G Chaikin. Logarithmic spiral grids for image processing and display. *Computer Graphics and Image Processing*, 11:197–226, 1979.