

Straight lines and circles in the log-polar image

David Young

School of Cognitive and Computing Sciences

University of Sussex

Brighton, BN1 9QH, UK

email: davidy@cogs.susx.ac.uk

Cognitive Sciences Research Paper CSRP 522

July 2000

To appear in *Proceedings of the Eleventh British Machine Vision Conference*,
Bristol, September 2000

Abstract

Foveal or spatially-variant image representations are important components of active vision systems. Log-polar sampling is a particularly powerful example as a result of the simplicity with which expansion and rotation can be handled. These properties are exploited here for the detection of general straight lines, line

form sampling, and partly because without effective active cameras non-uniform sampling sacrifices too much potential information. Now that active vision systems are becoming more common, non-uniform sampling is likely to increase in importance.

Despite its difficulties, non-uniform sampling, and in particular log-polar sampling, has received a certain amount of attention. Funt [3] demonstrated some of the fundamental advantages of an active foveated system for representing solid motion in 2-D, whilst Weiman and Chaikin [11] laid some mathematical groundwork. Wilson [12] emphasised the approximate log-polar mapping of the optic array onto the visual cortex in primates. A number of researchers, notably Tistarelli and Sandini [7, 8] have used the scheme in the context of motion detection; Tunley and Young [9] investigated the advantages of log-polar representations in estimating first-order optic flow. Also using log-polar sampling, Lim, West and Venkatesh [4] have developed mechanisms for precise foveation of features, Peters and Bishay [5] have described foveation on vanishing points, and Bederson, Wallace and Schwartz [1] have described an active vision system incorporating log-polar sampling.

The present paper builds on the theoretical work of Weiman and Chaikin [11] to explore the representation and detection of straight lines and circles in log-polar sampled images. An efficient new algorithm for finding these structures is described and its performance on real images investigated. The algorithm is intended to be applied in the context of a system like that of Brunnström, Eklundh and Uhlin [2], where a representation of a scene is built up using directed foveations.

2 The log-polar sampled image

In log-polar sampling, pixels are indexed by ring number R and wedge number W , related to ordinary x, y image coordinates by the mapping

$$\begin{aligned} r &= [(x - x_c)^2 + (y - y_c)^2]^{1/2}, & \theta &= \tan^{-1}\left(\frac{y - y_c}{x - x_c}\right) \\ R &= \frac{(n_r - 1)\log(r/r_{min})}{\log(r_{max}/r_{min})}, & W &= \frac{n_w \theta}{2\pi} \end{aligned} \quad (1)$$

where (r, θ) are polar coordinates, (x_c, y_c) is the position of the centre of the log-polar sampling pattern, n_r and n_w are the numbers of rings and wedges respectively, and r_{min} and r_{max} are the radii of the smallest and largest rings of samples. We also define $\rho = \log r$.

A log-polar sampled image is one whose samples are centred on points mapping to integral R and W , $R \in \{0, \dots, n_r - 1\}$, $W \in \{0, \dots, n_w - 1\}$. The separation between sample points is proportional to distance from the sam-

pling centre, as shown in Fig. 1a. This arrangement appears to be approximated by the ganglion cells of the primate retina and the visual cortex [6]. In this representation, image expansions and rotations about (x_c, y_c) become shifts in R and W , but image translation has a more complex effect.

In order to keep a pixel's nearest neighbours in orthogonal directions at approximately equal distances from it, the following constraint is needed

$$(2)$$

Log-polar sampled images are often displayed on orthogonal (R, W) axes, as in Fig. 1b, but this is misleading since it leads them to be regarded as “distorted” representations. In fact, the distortion only arises when they are displayed on the page or screen: as a mapping from coordinate values to position on a plane, the log-polar representation is no more distorted than the conventional one. When displayed with the correct mapping to position, as in Fig. 1c, the significant observable feature is the loss of resolution towards the periphery, as the samples become further apart.

These images should ideally be generated using special-purpose cameras, such as those described in [8]. However, a reasonable approximation for research is obtained by resampling a conventionally digitised image, and this method is used in the present work.

3 The straight line in the log-polar image

3.1 The log-polar straight line and its Fourier transform

Any straight line, not passing through (x_c, y_c) , can be mapped into any other straight line by a rotation (to make the lines parallel) followed by a uniform expansion with (x_c, y_c) fixed. This property can be exploited to allow easy detection of straight lines in the log-polar image.

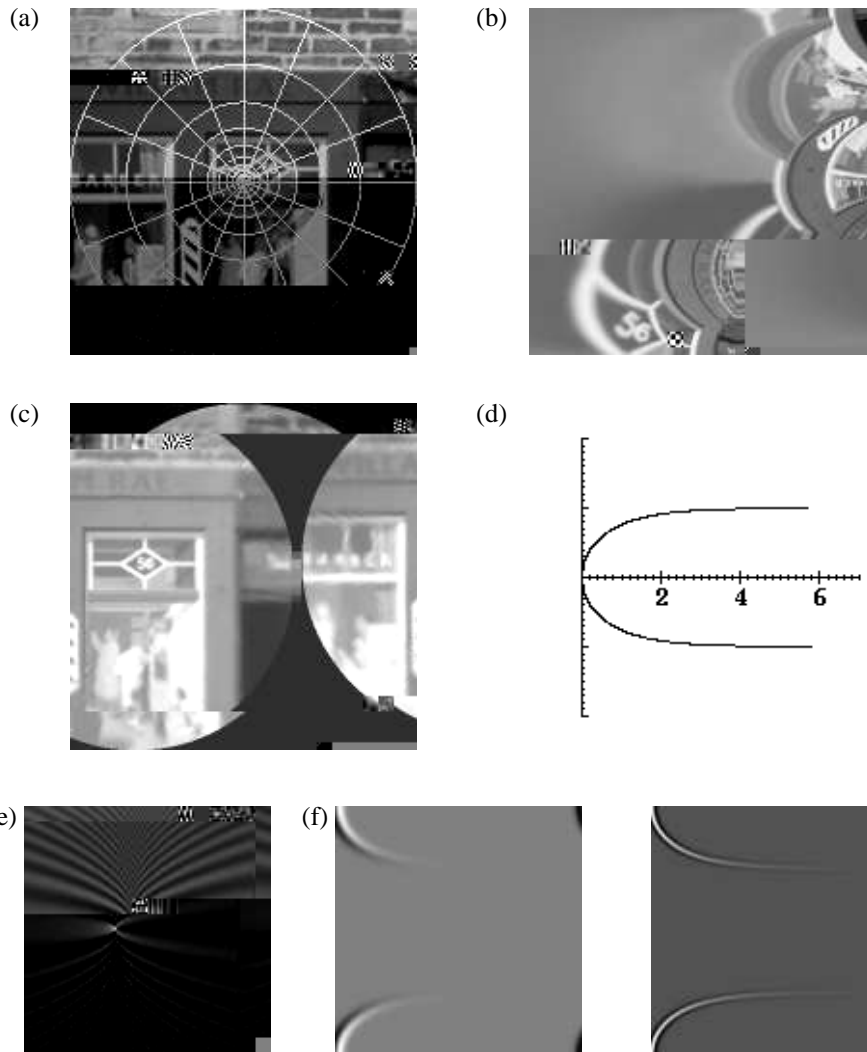


Figure 1: (a)

polar space. Although in practice it might be adequate to synthesise the straight line in a log-polar array and apply the discrete Fourier transform, computing its transform directly avoids noise caused by starting from a discrete representation of the line. The formula for the transform also opens up the possibility of further analysis of the properties of the process in the frequency domain, though this is not exploited here.

To find the transform, we take a path integral along the line in log-polar space; if S is the standard line $\rho = -\log \cos \theta$ with element ds in (ρ, θ) space, the integral is

$$F(k_\rho, k_\theta) = \int_S e^{-ik_\rho \rho} e^{-ik_\theta \theta} w(\rho, \theta) ds \quad (3)$$

where $w(\rho, \theta)$ is a weighting factor to allow convergence. This must be smooth and tend to zero for large ρ . A suitable choice is

$$w(\rho, \theta) = (\cos \theta)^{1-\alpha}, \quad 0 < \alpha < 1 \quad (4)$$

where a larger α makes the template more localised round the minimum of ρ . In all the examples in this paper, $\alpha = 0.2$. Since $d\theta/ds = \cos \theta$ the integral becomes

$$F(k_\rho, k_\theta) = \int_{-\pi/2}^{\pi/2} (\cos \theta)^{ik_\rho - \alpha} e^{-ik_\theta \theta} d\theta \quad (5)$$

Rearranging and using standard tables, this evaluates to

$$F(k_\rho, k_\theta) = \frac{\Gamma(\alpha - ik_\rho)}{\Gamma(\alpha)^2} \frac{\Gamma(\alpha - ik_\rho - ik_\theta)}{\Gamma(\alpha - ik_\rho - ik_\theta)} e^{-ik_\theta \theta} \quad (6)$$

where Γ is the complex gamma function. The line is at an arbitrary position in the log-polar grid; to make a useful mask we choose as the template . The discrete Fourier transform of this is obtained by evaluating $F(k_\rho, k_\theta)$ at unit intervals of k_R and k_W from 0 to and respectively, with and .

3.2 Implementation of straight line detection

Transformed straight line templates were generated using Eq. 6. Because of symmetry, the sum of the real and imaginary parts is sufficient; an example of a computed template in the frequency domain is shown in Fig. 1e. Templates were multiplied by the frequency domain representations of a variety

of other operators, allowing Gaussian smoothing, differentiation with respect to R or W , and difference of Gaussians convolution to be combined with line detection in a single step. This allows matching to various kinds of boundaries. Examples of the masks generated are shown in Figs 1f and 1g.

Any linear operation can be thus combined with line detection, but the images used in this study were first subjected to a non-linear process to reduce sensitivity to the grey-level range. This involved subtracting a local average of the grey level, obtained by Gaussian smoothing, from each pixel in the original conventional images, and then applying the logistic function to each pixel of the result.

This compression was carried out on the conventional image input prior to

peak, if the convolution had been carried out in the spatial domain. The size of these values indicates how much each pixel of the image contributed to the peak.

We then simply project this array onto the W axis, by summing over all R from 0 to $W-1$ for each W . Starting from the maximum in the resulting one-dimensional array, we search outwards in each direction, wrapping round if necessary, to find a value less than some constant times the peak value. This gives the limits in θ of the line segment which contributed most to the detected line. For a line with parameters ρ

4 The circle in the log-polar image

4.1 The log-polar circle and its Fourier transform

It is reasonable to ask whether the simplicity of straight line detection in log-

5 ‘Eye movements’: a resampling strategy

To get some sense of the potential value of log-polar line detection, it is important to simulate the way it might be exploited in an active vision system. To this end, a simple recentring strategy was used to move the log-polar pattern around in conventional images, in rough simulation of saccadic eye movements in the optic array. The strategy adopted is designed to demonstrate the possibilities of the approach rather than to be optimum in any respect.

By analogy with eye movements, a foveation is taken to mean extraction of information for a single sampling centre (x_c, y_c) , and a saccade to mean a movement of the sampling centre. Various kinds of saccade were programmed, including:

- (i) a small random step from a Gaussian distribution round the current centre;
- (ii) a random movement to anywhere in the original image with a probability depending on the density of previous sampling;

here, straight line detection is demonstrably effective, but circle detection is hampered by the inaccuracy of the representation. A significant advantage of the approach is that no edge or feature detection precedes the line detection; the main computational cost is an FFT of complexity $O(N \log N)$ for each foveation, and this could be carried out in suitable hardware.

Processes on log-polar images are not easy to compare with processes on conventional images, since they are designed to be embedded in a system with foveal sampling and an active camera at the hardware level. In particular, the results of log-polar line and circle detection are not comparable with those of,

say, edge detectors or the Hough transform operating on conventional images. Weiman [10] has described Hough transform detection of straight lines in log-polar space, but results on real images are not reported.

A full evaluation of the system described here requires further work. Tests on synthetic images show that the method can readily locate the boundaries of polygons where the standard deviation of the noise exceeds the grey-level difference between the interior and exterior, but this is not surprising given the incorporation of smoothing and the integration of evidence from a substantial area of the image. More to the point, Figs 4d and 4f give some indication of the extent to which a combination of log-polar line detection and camera movements might work together to build up a structural representation. An appropriate benchmark would involve a higher-level task which demanded that image structure be extracted.

These processes could play a valuable role if integrated into an active vision system, with spatially variant sampling at the hardware level. However, the most significant challenges are not at the level of feature detection, but lie in developing a strategy for foveation and saccade so as to integrate information effectively. This will require a more task-directed, purposive approach.

References

- [1] Bederson, B.B., Wallace, R.S. and Schwartz, E. (1995) A miniaturized space-variant active vision system: Cortex-1. *Machine Vision and Applications*, **8**, 101-109.
- [2] Brunnström, K., Eklundh, J-O., and Uhlin, T. (1996) Active fixation for scene exploration. *International Journal of Computer Vision*, **17**, 137-162.
- [3] Funt, B.V. (1980) Problem-solving with diagrammatic representations. *Artificial Intelligence*, **13**, 201-230.
- [4] Lim, F.L., West, G.A.W. and Venkatesh, S. (1997) Use of log polar space for foveation and feature recognition. *IEE Proc. Vis. Image Signal Process.*, **144**, 323-331.
- [5] Peters, R.A. II and Bishay, M. (1996) Centering peripheral features in an indoor environment using a binocular log-polar 4DOF camera head. *Robotics and Autonomous Systems*, **18**, 271-281.
- [6] Schwartz, E.L. Spatial mapping in the primate sensory projection: analytic structure and relevance to perception. *Biological Cybernetics*, **25**, 181-194.
- [7] Tistarelli, M. and Sandini, G. (1992) Dynamic aspects in Active Vision. *CVGIP: Image Understanding*, **56**, 108-129.

[8] Tistarelli, M. and Sandini, G. (1993) On the advantages of polar and log-polar mapping for direct estimation of time-to-impact from optical flow. *IEEE Trans. Pattern Analysis and Machine Intelligence*, **15**, 401-410.

[9] Tunley, H. and Young, D. (1994) Dynamic fixation of a moving surface using log polar sampling. In Hancock, E. (ed) *Proc. of the 5th British Machine Vision Conf.*,