

Contrast Marginalised Gradient Template Matching

Saleh Basalamah¹, Anil Bharath¹, and Donald McRobbie²

¹ Dept. of Bioengineering, Imperial College London, London SW7 2AZ, UK
{s.basalamah,a.bharath}@imperial.ac.uk

² Imaging Sciences Dept., Charing Cross Hospital, Imperial College London, UK

Abstract. This paper addresses a key problem in the detection of shapes via template matching: the variation of accumulator-space response with object-background contrast. By formulating a probabilistic model for planar shape location within an image or video frame, a vector-field filtering operation may be derived which, in the limiting case of vanishing noise, leads to the Hough-transform filters reported by Kerbyson & Atherton [5]. By further incorporating a model for contrast uncertainty, a contrast invariant accumulator space is constructed, in which local maxima provide an indication of the most probable locations of a sought planar shape. Comparisons with correlation matching, and Hough transforms employing gradient magnitude, binary and vector templates are presented. A key result is that a posterior density function for locating a shape marginalised for contrast uncertainty is obtained by summing the functions of the outputs of a series of spatially invariant filters, thus providing a route to fast parallel implementations.

1 Introduction

Primate vision employs various strategies in visual processing that appear to be highly effective. The apparent separation of visual input into various streams, or channels, of representation is one obvious example. These interpretations, collectively referred to as “channel models” of vision have been suggested for decades, and is backed up by much physiological [12] and psychophysical evidence. From a pattern recognition viewpoint, such ideas might not be surprising: the construction of multiple channels is closely analogous to the extraction of the components of a feature space constructed for pattern recognition purposes. In this paper, the issue of constructing a posterior distribution over spatial variables from a generic gradient field is discussed as a means of obtaining shape-specific responses.

Gradient field estimates are thought to play a significant role in biological vision, and particularly for distinguishing form (shape). At the same time, gradient field estimates are used as the basis of the Compact Hough Transform for detecting closed shapes [1].

In the rest of the paper, a generic probabilistic framework is suggested for locating planar shapes from gradient field observations, and this leads to (i) a

new interpretation of the Hough Transform filters suggested by Kerbyson and Atherton [5], and related techniques for medial axis detection [9] suggested by Bharath [2] (ii) an improvement in the principle of contrast invariant shape measures based on non-linear combinations of the outputs of linear spatial filters.

2 A Probabilistic Formulation

This paper addresses the task of locating compact planar shapes from gradient-fields of digital images, expressed as follows:

Q1: Given an estimate of a gradient field of a digital image, what is the most likely location of a known planar shape, S_0 , in the image ?

In answering this question, one may derive techniques that are related to existing methods of planar shape localisation and are in the spirit of Hough-transform approaches. For compact shapes, *Q1* requires the definition of a unique point of origin on the shape. The problem of finding the position of a known shape is reduced to one of finding its point of origin within the image plane.

Remark 1. The reason for restricting the observations to a gradient field, rather than the original image itself, is to focus on the nature of *shape*. There are, however, some other intriguing reasons for investigating this route: in sparse image (transform coding) it is quite easy to extract a gradient field estimate from the transform domain representation.

A statistical formulation for the question posed above exists in the form of conditional probability, and the most probable shape location, \mathbf{t}_{opt} , given the gradient field data may be found from

$$\mathbf{t}_{opt} = \arg \max_{\mathbf{t} \in \mathcal{R}_I} \{f_t(\mathbf{t}|\mathcal{B}, \mathcal{X}; \Theta)\} \quad (1)$$

where $(\mathcal{B}, \mathcal{X})$ represents N pairs of vector tuples, $\{\mathbf{b}^{(i)}, \mathbf{x}^{(i)}\}_{i=1..N}$, with $\mathbf{b}^{(i)}$ denoting an estimate (observation) of the gradient vector field at position $\mathbf{x}^{(i)}$. The possible location of the shape is denoted by the position vector, \mathbf{t} . More generally, one may seek local maxima of the posterior distribution on the right hand side of Equation (1) as an indication of the most probable locations of the shape in question, or of multiple instances of the same shape. The information about the particular shape being sought, and its observation conditions (such as rotation and scale), are contained in the parameter vector Θ .

Equation (1) represents a concise formulation of the shape detection problem, given observed data and known parameters. The first difficulty is that of estimating this probability density function (or, more precisely, finding its local maxima). Using Bayes' Theorem,

$$f_t(\mathbf{t}|\mathcal{B}, \mathcal{X}; \Theta) = \frac{f(\mathcal{B}, \mathcal{X}|\mathbf{t}; \Theta)f_t(\mathbf{t}|\Theta)}{f(\mathcal{B}, \mathcal{X}|\Theta)} \quad (2)$$

For a suitable choice of likelihood function and prior, it may be shown that Equation (2) leads to several algorithms for efficient construction of an accumulator space through convolution. Some of these have been identified previously, but new results may be obtained by a careful consideration of the likelihood model, and by marginalising over variations in contrast.

3 Plausible Likelihood Functions

A plausible form for the likelihood function, $f(\mathcal{B}, \mathcal{X}|\mathbf{t}; \Theta)$ is required. For a scalar intensity image, the value of the gradient vector field at location $\mathbf{x}^{(i)}$, $\mathbf{b}(\mathbf{x}^{(i)})$, is typically obtained by convolving the image with gradient masks¹ Such masks are designed to have zero average value, and are well localised in space. Thus, these masks approximate the properties of discrete wavelets, and the resulting gradient fields, despite their non-completeness, have univariate statistics that are well understood from the wavelet literature [10,11].

To characterise the statistical behaviour of a gradient field estimate consisting of two components, the *joint* statistics of the components of the vector $\mathbf{b}(\mathbf{x}^{(i)})$ should be determined, as should the joint statistics between two different field locations. Such high-dimensional density functions are difficult to model and approximate. To simplify the specification of a probabilistic model of gradient field behaviour, independence of these components, across all scenery, might be suggested. This is unrealistic, even for quite simple visual scenes.

A solution is to assume that the statistics are conditionally independent; that is, *given* the presence of some structure in the image, and certain restricted observation conditions, the vector components at any pixel deviate from a mean *vector* in an independent fashion. Conditional independence between observations is precisely what is needed for the construction of a useful likelihood function. Conditional independence between the components represents a simplified model, but it is a far weaker assumption than that of unconditional independence. Furthermore, although the independence assumption may fail under certain conditions, it is likely that there exist some variables such that the independence assumption of the joint statistics *conditioned on those variables* does hold. Finally, it should also be pointed out that marginalising over the variables under which the density function is defined yields unconditional joint statistics which are *no longer* independent.

3.1 Models for Gradient Field Statistics

It is widely recognized that the univariate statistics of wavelet coefficient space are well approximated by a Laplacian probability density function (PDF). For

¹ The dual notation that is used serves to (a) formulate the detection problem in terms recognizable as Bayesian inference, (b) illustrate the important link with spatial correlation/convolution, and (c) lay the ground for future work in which possible spatial distortion (uncertainty in the location $\mathbf{x}^{(i)}$) is corrected for.

example, Sendur and Selesnick [10] developed a bivariate Laplacian model for a wavelet coefficient space which they applied to denoising.

It turns out that the Laplacian density function suggested above is not critical for expressing the distribution of values in a *conditional* gradient field in a region near to and enclosing a shape of interest. Indeed, adopting the view of wavelet statistics, the Laplacian distribution is merely a reflection of the sparseness of the “detail” coefficient space. A Laplacian-like PDF can arise from marginalising a conditional bivariate Gaussian model over the location variable, \mathbf{t} .

3.2 Importance of Conditional Statistics

To develop the formulation in a concise way, it will be assumed that the observation conditions, Θ , are presumed fixed over all experimental space; the effect of contrast variations will be developed in Section (3.3). Then, *given* that a shape S_0 is present at location \mathbf{t} in an image, the statistics of the transform space can be modelled by a distribution in a simple, but elegant way: the *deviation* from the conditional mean vector is approximated by an appropriate distribution, where the conditional mean vector field is specified as a function of position in the observation field. To express this more clearly, the auxiliary random vector field conditional variable $\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0)$ is introduced:

$$\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0) = \mathbf{b}(\mathbf{x}^{(i)}) - \boldsymbol{\mu}(\mathbf{x}^{(i)}|\mathbf{t}, S_0) \quad (3)$$

where $\boldsymbol{\mu}(\mathbf{x}^{(i)}|\mathbf{t}, S_0)$ is the mean gradient vector field (over experimental space) of the gradient field *conditional* on shape S_0 being present at location \mathbf{t} .

A bivariate Gaussian version of the statistics of $\beta(\mathbf{x}^{(i)})$ is, therefore, both convenient and realistic, as will be demonstrated; it has the form

$$f_{\beta}(\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0), \boldsymbol{\mu}(\mathbf{x}^{(i)}|\mathbf{t}, S_0)) = \frac{\alpha_{\beta}}{\pi} e^{-\alpha_{\beta}|\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)}|\mathbf{t}, S_0))|^2} \quad (4)$$

which is a “standard” bivariate normal distribution, with α_{β} being half the inverse variance.

Remark 2 Caution is required in interpreting Equation(4); it is *not* a specification of the distribution of the *magnitude* of the gradient field deviation, but a bivariate distribution which decays (as a Gaussian) with the magnitude of the field deviation - this is a subtle, but crucial, distinction.

By considering well-defined shape localisation problems based on gradient estimates, it is possible to demonstrate that distributions of the form of Equation(4) can be representative of real gradient field observations, and, furthermore lead to the derivation of previous reported results and new methodologies for shape detection.

To validate the proposed model for gradient field statistics in the presence of objects, a simple experiment was conducted. An image of garden peas scattered across a rug, Figure (1a), was acquired using a standard consumer digital camera

(FujiFilm FinePix, 4700) on the “Normal Quality” setting at a resolution of 1280×960 . Normal artificial room lighting conditions were used, and the object-camera distance was of the order of 1.5 metres. The “V” channel of the HSV colourspace representation was extracted from the JPEG file of the camera. Using this channel, a thresholding operation was applied to binarise the image, then connected components labelling used to find individual peas. The centroids of the binary pea images were found, and a 33×33 area centred on several pea centroids was extracted from a gradient field estimate by using simple Prewitt operators. A mean gradient field was computed, and the bivariate probability density function of the deviation of actual gradient estimates from this mean was estimated by histogram binning. The result is shown in Fig. (1b). The total number of pixels used to estimate this histogram (on a 20×20 grid of bins) is approximately 20,000.




Fig. 1. (a) Peas-on-Rug image (b) Experimentally estimated bivariate PDF or region containing single peas from the “Pea-on-rug” image obtained by methods described in the text. (c) Sampling from $f_{\beta}(\beta_x, \beta_y | A)$ for $A = 1$. (d) Sampling from $f_{\beta}(\beta_x, \beta_y | A)$ and the prior on A , $f(A)$, and marginalising over A .

3.3 A Model for Contrast Uncertainty

Between any two gradient images, there will be a certain degree of contrast uncertainty. This contrast uncertainty will scale all values, including the x and y components of the gradient estimate around objects of interest, and might be considered one mechanism which leads to a violation of independence between components or pixels of a single image.

Although there will be uncertainty in this contrast value, some prior information about this uncertainty, and particularly about its role in providing “coupling” between the components of the gradient field, may be postulated. The following model for auxiliary conditional gradient field is proposed,

$$\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)}), A) = A\mathbf{b}(\mathbf{x}^{(i)}) - \boldsymbol{\mu}(\mathbf{x}^{(i)}|\mathbf{t}, S_0) \quad (5)$$

together with a prior for the contrast (amplitude scaling) parameter, A :

$$f_A(A) = \sqrt{\frac{\alpha_A}{\pi}} e^{-\alpha_A(A-\mu_A)^2} \quad (6)$$

and, for the auxiliary conditional random vector field deviation,

$$f_\beta(\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)}), A) = \frac{\alpha_\beta}{\pi} e^{-\alpha_\beta|\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)}), A)|^2} \quad (7)$$

In Fig. (2b), a sample is drawn from the gradient field model of small, bright circles of radius around 3 pixels against a darker background. The parameter μ_A was set to 1, α_β was set to 200, and $\alpha_A = 1$, i.e. corresponding standard deviations of 0.05 and $\sqrt{1/2}$ respectively. The conditional mean vector field, $\boldsymbol{\mu}(\mathbf{x}^{(i)}|\mathbf{t} = (0, 0), A = 1)$ is illustrated in Fig. (2a). The reader should note that the model includes a variation in contrast which introduces coupling between x and y gradient components and *independent* noise on the x and y components.

To show that this approach leads this to plausible PDF's that reflect those of real image data, one may use sampling principles [3] to marginalise the conditional density function, $f_\beta(\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)}), A)$ over A using a standard normally distributed random number generator², with $\mu_A = 1$, and $\alpha_A=1$. The result of this is shown in Fig. (1d). This should be compared with the estimates from the “pea on carpet” image of Fig. (1b).

3.4 Analytic Marginalisation over Contrast

Analytically, the marginalisation of Equation (7) over A using the prior of Equation (6) is also tractable, and leads to the main new useful results of this paper:

$$f_\beta(\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0, \boldsymbol{\mu}_t(\mathbf{x}^{(i)})) = \int_{-\infty}^{\infty} f_\beta(\beta(\mathbf{x}^{(i)}|\mathbf{t}, S_0, A, \boldsymbol{\mu}_t(\mathbf{x}^{(i)})) f_A(A) dA \quad (8)$$

² Matlab 6.5 running under Windows 2000




Fig. 2. (a)Left. Illustration of mean vector field conditional on a bright circle being present at location (0,0) on a dark background. (b) Right. Sample of vector field drawn from $f_{\beta}(\beta(\mathbf{x}^{(i)})|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)}), A)$ using vector field to the left. S_0 is a circle of radius 3 pixels

where, using a standard book of integral tables [4]

$$\int_{-\infty}^{\infty} \exp \{-p^2 x^2 \pm qx\} dx = \frac{\sqrt{\pi}}{p} \exp \left\{ \frac{q^2}{4p^2} \right\}, \quad p > 0 \quad (9)$$

and Equation (3) results in the following,

$$\begin{aligned} f_b(\mathbf{b}(\mathbf{x}^{(i)})|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)})) &= \frac{\alpha_{\beta}}{\pi} \sqrt{\frac{\alpha_A}{Z_0^{(i)}}} \exp \left\{ \frac{\alpha_{\beta}}{Z_0^{(i)}} \left[2\alpha_A \left(b_x(\mathbf{x}^{(i)})\mu_x(\mathbf{x}^{(i)}|\mathbf{t}) \right. \right. \right. \\ &\quad \left. \left. + b_y(\mathbf{x}^{(i)})\mu_y(\mathbf{x}^{(i)}|\mathbf{t}) \right) \right. \\ &\quad \left. + 2\alpha_{\beta} \left(b_x(\mathbf{x}^{(i)})b_y(\mathbf{x}^{(i)})\mu_x(\mathbf{x}^{(i)}|\mathbf{t})\mu_y(\mathbf{x}^{(i)}|\mathbf{t}) \right) \right. \\ &\quad \left. - \left(\alpha_A + \alpha_{\beta}(b_y^2(\mathbf{x}^{(i)})) \right) \mu_x^2(\mathbf{x}^{(i)}|\mathbf{t}) \right. \\ &\quad \left. - \left(\alpha_A + \alpha_{\beta}(b_x^2(\mathbf{x}^{(i)})) \right) \mu_y^2(\mathbf{x}^{(i)}|\mathbf{t}) \right. \\ &\quad \left. \left. - \alpha_A \left(b_x^2(\mathbf{x}^{(i)}) + b_y^2(\mathbf{x}^{(i)}) \right) \right] \right\} \quad (10) \end{aligned}$$

with

$$Z_0^{(i)} = \alpha_A + \alpha_{\beta} \left[b_x^2(\mathbf{x}^{(i)}) + b_y^2(\mathbf{x}^{(i)}) \right] \quad (11)$$

and where $\mu_x(\mathbf{x}^{(i)}|\mathbf{t})$ and $\mu_y(\mathbf{x}^{(i)}|\mathbf{t})$ are the x and y components of the mean gradient field, $\boldsymbol{\mu}(\mathbf{x}^{(i)}|\mathbf{t})$; $b_x(\mathbf{x}^{(i)})$ and $b_y(\mathbf{x}^{(i)})$ are the x and y components of the observed vector gradient field field at the i^{th} pixel, located at $\mathbf{x}^{(i)}$.

Assuming that all gradient observation pixels are drawn independently from this model, then, for a set of N observations,

$$\begin{aligned} f_b(\{\mathbf{b}(\mathbf{x}^{(i)})\}_{i=1..N}|\mathbf{t}, S_0, \boldsymbol{\mu}(\mathbf{x}^{(i)})) &= \prod_{i=1}^N f_b(\mathbf{b}(\mathbf{x}^{(i)})|\mathbf{t}, S_0, \boldsymbol{\mu}_t) \\ &= Z_N \exp\{\alpha_\beta(2\alpha_A C_1 + 2C_2 - C_3 - C_4 - \alpha_A C_5)\} \end{aligned} \quad (12)$$

where $Z_N = \prod_{i=1}^N \frac{\alpha_\beta}{\pi} \sqrt{\frac{\alpha_A}{Z_0^{(i)}}}$, and

$$C_1 = \sum_{i=1}^N \frac{b_x(\mathbf{x}^{(i)})\mu_x(\mathbf{x}^{(i)}|\mathbf{t}) + b_y(\mathbf{x}^{(i)})\mu_y(\mathbf{x}^{(i)}|\mathbf{t})}{Z_0^{(i)}} \quad (13)$$

$$C_2 = \sum_{i=1}^N \frac{b_x(\mathbf{x}^{(i)})b_y(\mathbf{x}^{(i)})\mu_x(\mathbf{x}^{(i)}|\mathbf{t})\mu_y(\mathbf{x}^{(i)}|\mathbf{t})}{Z_0^{(i)}} \quad (14)$$

$$C_3 = \sum_{i=1}^N \frac{\alpha_A + \alpha_\beta b_y^2(\mathbf{x}^{(i)})\mu_x^2(\mathbf{x}^{(i)}|\mathbf{t})}{Z_0^{(i)}} \quad (15)$$

$$C_4 = \sum_{i=1}^N \frac{\alpha_A + \alpha_\beta b_x^2(\mathbf{x}^{(i)})\mu_y^2(\mathbf{x}^{(i)}|\mathbf{t})}{Z_0^{(i)}} \quad (16)$$

and

$$C_5 = \sum_{i=1}^N \frac{b_x^2(\mathbf{x}^{(i)}) + b_y^2(\mathbf{x}^{(i)})}{Z_0^{(i)}} \quad (17)$$

Each of the terms C_1, C_2, C_3, C_4 involves spatial convolution (or cross correlation) between the appropriate mask and simple functions of the gradient field components. Term C_1 , for example approaches Kerbyson and Atherton [5] filtering technique if α_A is significantly greater than α_β . Returning to Equations (1) and (2), one may write

$$f_t(\mathbf{t}|\mathcal{B}, \mathcal{X}; \Theta_1) \propto Z_N \exp\{\alpha_\beta(2\alpha_A C_1 + 2C_2 - C_3 - C_4 - \alpha_A C_5)\} f_t(\mathbf{t}|\Theta_1) \quad (18)$$

The choice of the prior $f(\mathbf{t}|\Theta_1)$ may be completely non-informative, or may be chosen as coming from other data fields in a real image. For example, in searching for faces in a scene, the prior may be derived from some measure in HVS colour space; in searching for sub-components within an automotive inspection application, geometric considerations may yield a useful prior.

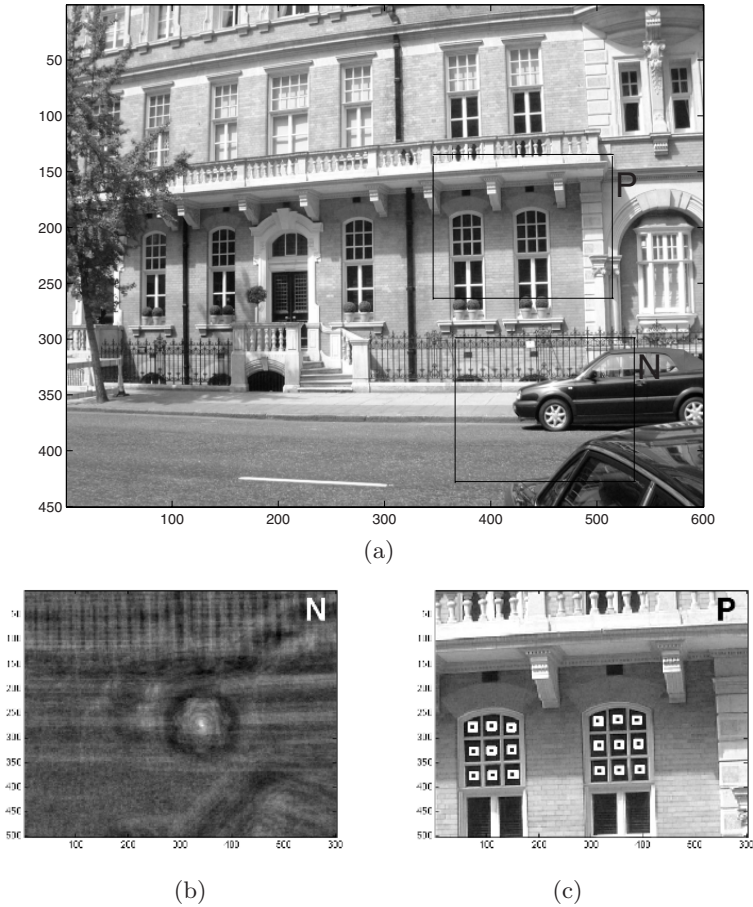


Fig. 3. (a) Image of building and car (downsampled by a factor of 4) (b) Accumulator space for circles, or $\log(f_t(\mathbf{t}|\mathcal{B}, \mathcal{X}; \Theta_1))$ (c) Detected locations of square centroids. For reasons of space and clarity, only the results after eliminating disconnected pixels in thresholded space are shown. An image with overlaid accumulator space is presented in the supplementary material.

4 Experimental Results

The performance of the technique has been tested on a range of images. Comparisons with other detection methods have been performed such as the correlation and gradient magnitude implementation of the Compact Hough transform. In all cases a flat prior on \mathbf{t} was chosen.

4.1 Natural Image Test

To represent the performance of the technique in natural images a test has been conducted on a 1800x2400 image of a building and car (Fig. (3a)), where the

aim is to detect the small rectangular windows in the image and the car wheel. The windows have a size of 34×48 pixels and the car wheel has a diameter of 80 pixels. Two different geometrical templates (a rectangle for detecting the windows and a circle for detecting the wheel) were created using generalized rectangular functions [6] that mimic optical blurring.




Fig. 4. (a) Average face template. (b) Image of people. (c) Marginalized gradient accumulator space. (d) Thresholded accumulator space showing face locations. (e) Correlation accumulator space. (f) Thresholded correlation accumulator space.

4.2 Face Detection

The technique has also been applied for detecting non-geometric objects. “A typical averaged face”, Fig.(4a), prepared by Adam D. Bradley, was downloaded from Dr. Robert Frischholz’s Face Detection Homepage³ and was used to detect faces in Fig. (4b). The technique correctly detects the faces as they are the highest local maxima in the accumulator space (Fig.(4c)). In comparison, the accumulator space of correlation gives high responses at most face locations, but also responds to other objects in the image as well (Fig.(4e)). The design of the gradient template, including the scale of the gradient field operators, remains an important aspect of this formulation and much work would need to be done in this area.

4.3 Circle and Square Test

To test the behaviour of the technique under different contrast conditions and noise levels, a 241x241 image containing circles and squares was synthesized. The background contrast of the image is varied linearly along the horizontal axis and also random noise is added to the image. There are squares and circles of the same size, aligned horizontally and vertically. Circles and squares with variation of both size and contrast are also placed along the vertical axis. Fig.(5a) represents the image with a random noise of $\sigma = 1$.

To demonstrate the discriminative performance of the technique, receiver operating characteristic (ROC) curves were plotted for each of 4 techniques: (a) Cross Correlation between image and template, (b) Template matching between gradient magnitude of image and template, (c) Template matching between thresholded gradient magnitude of image and binary template, (d) Kerbyson and Atherton’s method: template matching between x-y gradients of image and x-y gradients of template.

The ROC is a plot of true positive rate (TP) against the false positive rate (FP). The true positives represent the correct detection of squares with a side-length of 9 pixels, whereas the false positives represent the detection of all the other objects (circles with different diameters and squares with different side-lengths). 100 threshold levels are applied on the accumulator space (Fig. (5b)) where TP and FP are calculated for each level. The ROC plot in Fig. (4c) compares our technique with 4 other detection techniques in images with random noise of $\sigma = 2$.

Since the total number of objects in the image is only 25, the number of experiments is increased to include 10 images, each with different noise realisations thus raising the total number of objects at each threshold level to 250.

5 Conclusions and Further Work

Marginalisation over a contrast variation solves the problems of early threshold setting that plague many Hough transform techniques. Whilst exact sampling

³ <http://home.t-online.de/home/Robert.Frischholz/face.htm>

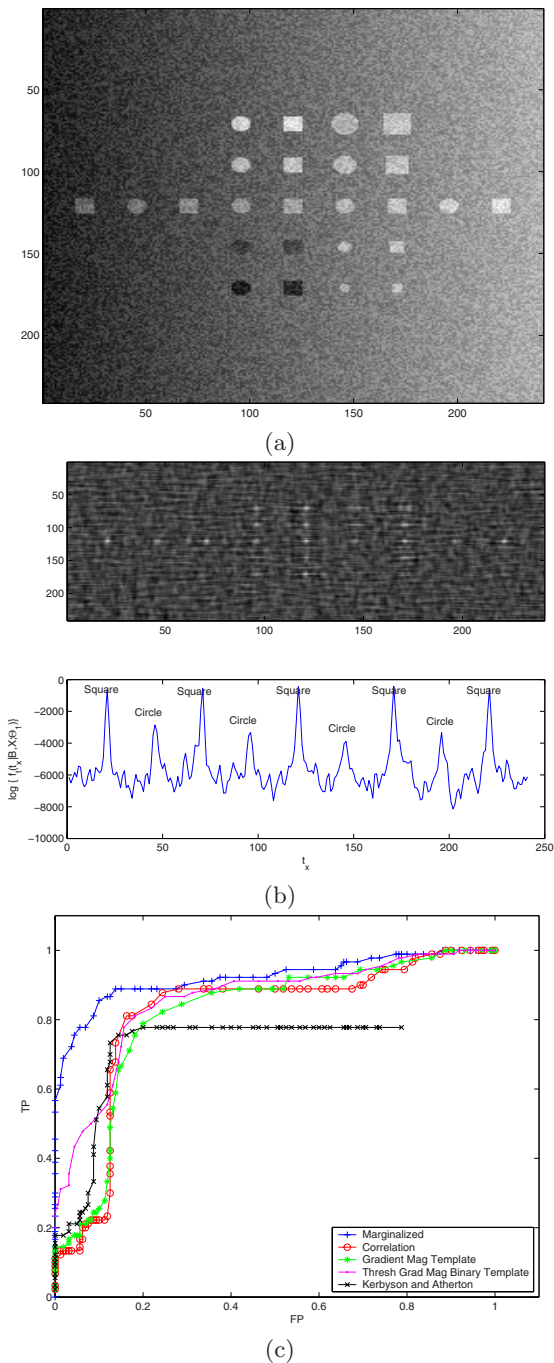


Fig. 5. (a) Synthesized image of circles and squares. (b) (Top) $\log(f_t(\mathbf{t}|\mathcal{B}, \mathcal{X}; \Theta_1))$ for 9 pixel squares detection, (Bottom) Profile through top response at row=121. (c) ROC comparison of different detection methods.

methods provide the possibility of more accurate estimates of posterior density functions of shape locations, the approach employed in this paper leads to a direct implementation through spatial convolution. This work has not addressed the crucial problem of handling within plane affine geometric transformation of the shape being sought [8]. Strategies for handling this problem using either sequential parameter estimation techniques [7] or marginalisation over, for example, rotation, are under investigation. Finally it should also be pointed out that the scale of gradient estimators is likely to play a critical role in distinguishing shapes that are close to each other and in the tolerance of marginalised filtering to noise.

Acknowledgement. This work is sponsored by the Saudi Ministry of Higher Education and is partially supported by the UK Research Council under the Basic Technology Research Programme "Reverse Engineering Human Visual Processes" GR/R87642/02.

References

1. D.H. Ballard. Generalising the Hough Transform to detect arbitrary shapes. *Pattern Recognition*, 13:111–122, 1981.
2. A. A. Bharath and C.J. Huberson. Obtaining medial responses from steerable filters. *IEE Proceedings on Vision, Image and Signal Processing*, 146(5):1087–1088, 1999.
3. Michael Evans and Tim Swartz. *Approximating Integrals via Monte Carlo and Deterministic Methods*. Oxford University Press, 2000.
4. I.S. Gradshteyn, I.M. Ryzhik, and Alan Jeffrey. *Tables of Integrals, Series and Products: 5th Edition*. Academic Press, 1994.
5. J.H. Kerbyson and T.J. Atherton. Circle detection using hough transform filters. In *Image Analysis and its Applications, 1995, IEE Conference Publication No. 410*, pages 370–374, 1995.
6. M.J. Lighthill. *Introduction to Fourier Analysis and Generalised Functions*. Cambridge monographs on mechanics and applied mathematics. Cambridge: Cambridge University Press, 1958.
7. David J.C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
8. Maria Petrou and Alexander Kadyrov. Affine invariant features from the trace transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:30–44, 2004.
9. Stephen M. Pizer, Christina A. Burbeck, James M. Coggins, Daniel S. Fritsch, and Bryan S. Morse. Object shape before boundary shape: Scale-space medial axes. *Journal of Mathematical Imaging and Vision*, 4:303–313, 1994.
10. Levent Sendur and Ivan W. Selesnick. Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency. *IEEE Transactions on Signal Processing*, 50(11), November 2002.
11. E.P. Simoncelli and E.H. Adelson. Noise removal via bayesian wavelet coring. In *Proceedings of the 1996 International Conference on Image Processing*, volume 1, pages 379–382, September 1996.
12. Martin J. Tovée. *An Introduction to the human visual system*. Cambridge Press, 1996.