STATE UNIVERSITY OF NEW YORK AT
STONY BROOK
ELECTRICAL ENGINEERING DEPARTMENT

CEAS TECHNICAL REPORT 615

Object Recognition by Quasi-Invariant Filters

John J. Murray and Yong Yu

January 2, 1992

Abstract

A modified approach for rotation-invariant pattern recognition in two dimensions is proposed. A class of filters is derived by optimizing a stochastic performance measure which tends to yield a spike at the location of an object regardless of its orientation, and small values elsewhere. These filters are quasi-invariant under rotation, handle additive noise and resolve closely-space objects. Simulation results are shown, in which the filters are implemented by using the FFT.

Index Term: Computer vision, Image processing, Object detection, Optimal filter, Pattern recognition.

1 Introduction

Object and feature recognition are important areas in computer vision. Human observers are generally capable of recognizing patterns independently of their orientation, position, and size within an image; a fundamental obstacle to doing this on a computer, however, is the dimension of the space to be searched. Techniques have, therefore, been developed that, to a certain degree, enable one to recognize the pattern no matter what its rotation and size; this paper presents improved methods for doing this in the case of rotation.

Using simple circular harmonic components, Hsu [1] proposed a method for rotation invariant digital pattern recognition. In this method, a reference pattern is expressed in polar coordinates by its circular harmonic components, of which only one is used to cross-correlate with the input image. Multiplication of the Fourier spectra provides linear shift-invariant correlation operations. Wu and Stark [2] modified this method by using a set of circular harmonic components instead of only one; since

several harmonic components are used to determine the pattern, pattern specificity is improved.

Schils and Sweeney [3] also proposed a rotation-and translation-invariant matching algorithm. This rotation and translation invariant filter uses the linear combination of circular harmonic components in such a way that the magnitude of the correlation between the input and filter is constant, independent of orientation. An iterative method is used to obtain a group of coefficients to achieve this result. Caelli and Liu [4] developed an adaptive matched filtering technique for the invariant recognition problem where, for a given recognition criterion, the number of templates needed to achieve invariant recognition varies with the pattern structure. With this set of templates, invariance with respect to translation, rotation, and dilation is achieved, and uniqueness is preserved up to the threshold chosen.

Moment invariants were used by Hu [5] as image recognition features which have the desirable property of being invariant under such variations of the image content as shifting, scaling, and rotation. Since then, they have been given considerable attention in the literature and satisfactory experimental results [6, 7, 8] have been reported.

Kummar and Pochapsky [9] proposed a modified matched spatial filter based on the use of a training set of images. The training set consists of N images $r_1(x), r_2(x), ..., r_n(x)$, which are obtained by deliberately distorting a given image $r(x)$. The ECP [10] filter used in this method is matched to a new image $h(x)$ chosen such that it produces the same cross correlation value at the origin with all training images $r_i(x)$, i=1,2,...,N. It is shown in the paper [9] that the price paid for detecting N independent images with a singe filter is a decrease in the SNR by a factor of N with respect to the matched spatial filter for detection of a known image in the presence of white noise.

In this paper, we propose another approach for pattern recognition. Instead of the

traditional matched filter method that maximizes the output at the object location, we derive a class of filters based on optimization of a stochastic criterion subject to ignorance of the orientation of the object in question. The effect of this is to give a peak in the output at the object location and small values elsewhere. From another point of view, the present approach closely resembles a Wiener filter, rather than a matched filter. This has the disadvantage that it is more sensitive to noise but has the advantage that it yields sharper peaks,and therefore greater specificity and location accuracy.

The paper first introduces the basic stochastic criterion to be optimized, and derives the optimal filters. In part 3, the relationship of these filters to rotation-invariant filters is discussed. Next, an algorithm which uses several of the filters simultaneously is presented. Experimental results are shown in section 5, and a summary and discussion of future directions are given in the last section. The proof of the main result, and a discussion of the relationship with harmonic components, are given in the appendices.

2 Optimal Filter

The problem to be solved is the derivation of a filter to detect and locate one or more instances of a given object at unknown orientations in a noisy image, possibly while rejecting other objects. It is assumed that the problem data, and therefore the solution, are translation invariant, and attention is restricted to linear filters.

The major problem is that a direct approach requires a search in three dimensions (two translation and one rotation). To reduce the dimensionality of the problem we will, instead of searching in three dimensions for the given deterministic object, search in two dimensions for a corresponding stochastic object, namely, the given object with a random orientation.

The precise formulation is as follows. We will assume that the two-dimensional object to be detected is given by the function $s(\boldsymbol{x})$ in some standard orientation and position, and that the background noise and clutter can be modelled by an additive two-dimensional random noise process $n(\boldsymbol{x})$.

The input image is then given by

$$r(\boldsymbol{x}) = s_\theta(\boldsymbol{x}) + n(\boldsymbol{x})$$

where θ is a random variable, uniformly distributed on the interval $[0, 2\pi)$, and s_θ is the object rotated through an angle of θ degrees.

For simplicity, we assume that $n(\boldsymbol{x})$ is a zero-mean, second-order stationary, isotropic random process; here isotropic means that the autocorrelation function, and hence the power spectral density, are invariant under rotation.

Since the filter is linear, it is sufficient to look for one instance of the object; since it is assumed shift-invariant, it is sufficient to search for an object located at $(0, 0)$.

We then look for a linear filter $f(\boldsymbol{x})$ to maximize the functional

$$J(f) = E_\theta \left\{ |(f * s_\theta)(0,0)|^2 \right\}$$

subject to

$$k_1 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E_\theta \{ |(f * s_\theta)(\boldsymbol{x})|^2 \} d^2\boldsymbol{x} + k_2 E\{ |f * n|^2 \} = 1$$

The effect of the maximization is that, averaged over θ, the largest possible peak in the output at $(0, 0)$ (the location of the object) is obtained, subject to the condition that the mean square over the entire output, (a weighted sum of contributions due to the background clutter and the output from the object at points other than $(0, 0)$), is held constant. This has the effect of giving a sharp peak at the object location, which is desireable to resolve multiple instances of the object. From another point of view, one can regard the above objective function as minimizing the effect of clutter,

where instances of the object at locations other than $(0,0)$ are regarded as part of the clutter.

Since the solution of the problem intrinsically involves complex filters, we will assume from here on that all of the filters involved may take complex values.

The solution to the above minimization problem is not straightforward, since both the objective and the constraint are quadratic; however, it is shown in appendix A that the solutions are given by the following theorem.

Theorem 1 *Let*

$$S(\boldsymbol{\omega}) = \mathcal{F}\{s(\boldsymbol{x})\}$$

and

$$Q(\boldsymbol{\omega}) = \frac{k_1}{4\pi^2} E_\theta\{|S_\theta(\boldsymbol{\omega})|^2\} + k_2 P_n(\boldsymbol{\omega})$$

where $S_\theta(\boldsymbol{\omega})$ denotes the Fourier transform of s_θ and $P_n(\boldsymbol{\omega})$ is the power spectral density of the noise n. Then, if the noise is isotropic, and the orientation θ is uniformly distributed on the interval $[-\pi, \pi]$, the function Q depends only on the magnitude of the vector $\boldsymbol{\omega}$. If further

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{|S(\boldsymbol{\omega})|^2}{Q(\boldsymbol{\omega})} d^2\boldsymbol{\omega} < \infty \tag{1}$$

then the stationary points of the functional

$$J(f) = E_\theta\left\{|(f * s_\theta)(0,0)|^2\right\}$$

subject to the constraint

$$k_1 \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} E_\theta\{|(f * s_\theta)(\boldsymbol{x})|^2\} d^2\boldsymbol{x} + k_2 E\{(f * n)^2\} = 1 \tag{2}$$

are given by the functions f_m whose Fourier transforms are:

$$F_m(\boldsymbol{\omega}) = K_m \frac{1}{Q(\boldsymbol{\omega})} \int_{-\pi}^{\pi} S_\theta^*(\boldsymbol{\omega}) e^{jm\theta} d\theta \tag{3}$$

where the constant K_m is chosen so that the constraint 2 is satisfied.

Finally, the value of the objective functional at the stationary points is given by

$$\lambda_m = J(f_m) = \frac{1}{32\pi^5} \int_0^\infty \left| \int_{-\pi}^{\pi} \tilde{S}(\rho, \phi) e^{jm\phi} d\phi \right|^2 \frac{\rho}{\tilde{Q}(\rho)} d\rho$$

where $\tilde{S}$ and $\tilde{Q}$ denote S and Q respectively in polar coordinates.

The proof of this theorem is given in the appendix; here some comments may be made. First, it follows that the solution(s) of the problem are found by picking the value(s) of m which maximize the integral 1, and using the corresponding f_m given by 3 as the optimal filter(s). Second, we note that the constant K_m in the expression for F_m is simply a scale factor, and does not affect the performance; what we are really concerned with is the ratio of the objective functional to the constraint functional. Third, the condition 1 only serves to ensure that the optimum is finite; it is therefore always satisfied under realistic conditions, e.g., if the signal energy is finite, and there is a component of white noise. Finally, we note that if the conditions of isotropy or uniform distribution of θ are not satisfied, the solutions are of the same general form as 3. However, the functions $e^{jm\theta}$ are replaced by functions related to the eigenvectors of a symmetric integral equation, and the values λ_m are the corresponding eigenvalues; in all but the simplest cases, these will have to be calculated numerically.

Examples of the performance of these filters are given in section 5; in the next section we indicate their relationship to rotation-invariant filters.

3 Rotation Quasi-Invariant Filters

As was mentioned previously, pattern recognition is usually a problem of searching in a high dimensional space, and the goal of most recognition algorithms is to reduce the size of the space to be searched. Among the many different ways of doing this, (e.g., heuristic, pruned tree searches), the present approach is hierarchical; it involves a two-dimensional search for points at which there is likely to be an instance of the

given object at any orientation, followed (if necessary) by a one-dimensional search at these points. A two-dimensional search followed by a small number of one-dimensional searches requires an order of magnitude fewer computations. The penalty paid for this is reduced specificity.

The usual approaches to reducing the search dimension in the context of unknown orientations is to perform matched filtering for one or more of the circular harmonic components of the object being sought, rather than for the object itself, or to use rotation-invariant filtering. The relationship between the present approach and circular harmonics is discussed in appendix B; here we concentrate on the relationship to rotation invariance.

By a (rotation) quasi-invariant filter we will mean a filter which has the property that, if its input is rotated through an angle θ_0, the output will be rotated through θ_0, and multiplied by a constant, k_{θ_0}, of magnitude 1. This is a natural generalization to complex-valued filters of the idea of rotation invariance.

Because of the basic group property of rotation, it is easily seen that we must have

$$k_{\theta_0} = e^{jm\theta_0}$$

for some integer m. Rotation invariance is simply the special case where $m = 0$.

It is also easily seen that quasi-invariance is equivalent to the filter's impulse response, $h(x, y)$, having the property:

$$h(x\cos\theta_0 + y\sin\theta_0, -x\sin\theta_0 + y\cos\theta_0) = e^{-jm\theta_0} h(x, y)$$

for some integer m; if h is expressed in polar coordinates (ρ, θ) as

$$h(x, y) = h(r\cos\theta, r\sin\theta) = \tilde{h}(r, \theta)$$

then one has

$$\tilde{h}(r, \theta) = \tilde{h}(r)e^{jm\theta}$$

In the frequency domain, if

$$H(\boldsymbol{\omega}) = \mathcal{F}\{h(\boldsymbol{x})\}$$

and H is expressed in polar coordinates (ρ, ϕ) as

$$\tilde{H}(\rho, \phi) = H(\rho\cos\phi, \rho\sin\phi) = H(\boldsymbol{\omega})$$

then quasi-invariance translates to

$$\tilde{H}(\rho, \phi) = \tilde{H}(\rho, 0)e^{jm\phi}$$

If we now recall that $S_\theta(\boldsymbol{\omega})$ can be expressed in polar coordinates (ρ, ϕ) as $\tilde{S}(\rho, \phi + \theta)$, and that $\tilde{Q}(\rho)$ is independent of ϕ, it is easily seen that the optimal filters given by equation 3 are quasi-invariant.

In fact, it can be shown that these filters could have been derived by the usual Schwartz inequality approach used in matched and noncausal Wiener filtering, subject to the additional constraint of quasi-invariance; that is to say, each F_m given by equation 3 is the function which maximizes the functional:

$$I(f) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(\boldsymbol{\omega})S(\boldsymbol{\omega})d^2\boldsymbol{\omega}$$

subject to:

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} |F(\boldsymbol{\omega})|^2 Q(\boldsymbol{\omega})d^2\boldsymbol{\omega} = constant$$

and

$$\tilde{F}(\rho, \phi) = \tilde{F}(\rho, 0)e^{jm\phi}$$

Here, as before, $\tilde{F}$ is F expressed in polar coordinates.

We feel, however, that the approach in the previous section is preferable, for two reasons: first, it derives the quasi-invariance of the optimizing solutions, rather than imposing it as an a priori constraint; and second, it gives more insight into the basic

statistical assumptions involved. As an illustration of the latter point, consider the case of object detection subject to unknown magnification, rather than orientation. If one imposes the quasi-invariance constraint a priori, it is natural to try to derive an analogous quasi-invariant filter for the case of magnification. It is well known that this yields very poor results. From the point of view of the previous section, however, this is to be expected; the crucial ingredient which yields a quasi-invariant filter is the assumption of a uniform a priori distribution for the random variable θ, and the analogous assumption in the magnification case does not make sense.

4 Use Of Multiple Components

As must be expected, and as is confirmed by the results of section 5, the use of the optimal filter derived in section 2 results in a lower peak value than the optimal filter when the orientation is known. In the present situation, however, we have an infinite set of independent stationary solutions; the question therefore arises as to whether it is possible to combine several of them in such a way as to increase the specificity of the filter.

Before turning to this, it should be noted that the loss of specificity can range from mild to severe, depending on how the energy of the object is distributed among its various circular harmonics. Objects for which one of the λ_k is much larger than the others will have only a small loss of specificity when the optimal filter is used, while those for which many λ_k are clustered around the maximum value will give poor optimal filters. In terms of shape, objects which are close to circular will work well, while long, narrow objects will work poorly. It is also worth noting that an algorithm which combines a large number of the components will be using the same information as, and will be computationally as demanding as, a three-dimensional search.

There is no linear way of combining the various quasi-invariant filters to get better

performance (since the filter we have derived is the optimal linear filter). It is therefore necessary to combine them in a nonlinear fashion. The method we have chosen is to apply a simple pattern recognition approach to the outputs of a moderate number of the filters corresponding to the largest λ_k. Rotation invariance is preserved, since a rotation does not affect the magnitude of λ_k, and accuracy is improved substantially.

In greater detail, the algorithm used may be described as follows:

1) For a given object, calculate λ_l ($0 \leq l \leq L$). Since the object is real, we need consider only integer $l \geq 0$.

2) Pick a number M ($M \leq L$); M is the number of components which we can afford to use.

3) Pick the M values of $l_1, ..., l_M$, which give the largest values of λ_l and pick tolerances ϵ_m. Apply the filters $F_{l_1}, ..., F_{l_M}$ to the input image; denote the output images by $V_{l_1}, ..., V_{l_M}$

4) Let the final output be given by

$V(x,y) =$ number of j, $1 \leq j \leq M$ such that $||V_{l_j}(x,y)| - \lambda_{l_j}| \leq \epsilon_j$

This assumes that the expected input intensity of the object to be detected is known; if not, the ratios of the λ_l can be used instead of the actual values.

This algorithm is essentially a robust matching algorithm on the magnitudes of the outputs. Other matching algorithms (e.g, correlation-coefficient based) could be used, but they tend to be insufficiently robust to account for the fact that, in a digitized image, rotations can give only approximate versions of the rotated object. It should also be noted that, totally apart from the fact that only a finite number of components is used, this method can not achieve ideal performance, since it would be possible to have two objects with the same λ_l values, but whose components have different relative orientations. Nonetheless, the performance is quite good, as is shown by the examples in the next section.

5 Experimental Results

The algorithms described above have been implemented on a VAX 11/780; convolutions were performed by means of the FFT, and interpolation was used for polar-to-rectangular and rectangular-to-polar conversions.

Initially, the optimal single filter given by equation 3 was implemented with the parameters $k_1 = 1, k_2 = 100$ and zero-mean white Gaussian noise with standard deviation $\sigma = 0.0001$. The object to be detected was a square, and in this case the optimal filter was given by $m = 0$. The input image in Fig. 1 is the original image (64 by 64 resolution) contaminated with zero-mean white Gaussian noise with $\sigma = 0.0001$. The magnitude of the output is shown in Fig. 2a; it can be clearly seen that there are two sharp peaks at the location of the squares, and substantially lower peaks in the vicinity of the triangle locations. Further, the size of the peaks at the square location is independent of the orientation of the square. Conversely, if we use the triangle as the reference object (with the same parameters as before), the output is as shown in Fig. 2b, with much the same results as before.

For comparison, the outputs obtained by using the matched filter corresponding to the zero-order harmonic components of the square and the triangle are shown in Fig. 2c and 2d, respectively. The following points may be noted: first, both filters achieve rotation-invariant responses; second, the output from the matched filter is much smoother, with none of the noise-like peaks which occur in the output from our filter; third, the penalty for this smoothness, however, is much broader, more ill-defined peak in the object locations; and fourth, although the matched filter for square detection correctly has its largest peaks at the locations of the squares, the matched filter for triangle detection incorrectly places the largest peaks at the location of the square also. The problem for the matched filter here is that each square may be thought of as two triangles in close proximity, and the matched filter has an output

contribution for each of the two triangles, without the resolving power to separate them into distinct peaks.

Although the single component filter gives good results for the square and triangle, this may not hold for more complex shapes. For example if the input image is Fig. 3 contaminated with zero-mean white Gaussian noise with $\sigma = 0.0001$ and all parameters are again as in the previous examples, the outputs due to using the zero-order quasi-invariant filter for detecting "E" and "A" are shown in Fig. 4a and Fig. 4b, respectively. In Fig. 4a, the peaks do occur at the correct location, but there is considerably more interference, in the shape of broader main peaks and noise like smaller peaks, than there was for the simpler shape; while in Fig. 4b, it is not even clear which character gives the higher peak. (For comparison, the outputs from the corresponding matched filters are shown in Fig. 4c and 4d and here the "E" peak is clearly higher in both cases.)

For this reason, we turn to examples in which multiple components are used as described in section 4. Since, in using multiple components, we want to find those with largest λ_l values, the λ_l ($l = 0$ to 10) values for the square and triangle objects were calculated. As expected, for the square, the only significant values occur when l is a multiple of 4. The triangle, on the other hand, gives significant values for all values of l up to 10, except for $l = 1$. In the following, we use a number of these components to implement the algorithm of section 4, with $\epsilon_l = \lambda_l/3$, $k_1 = 1$ and $k_2 = 100$.

Fig. 5a is the original image contaminated with zero-mean white Gaussian noise with $\sigma = 0.1$. The outputs are shown in Figs. 5b and 5c for detection of the square (using three f_m) and triangle (using four f_m), respectively. It is clear that a significant improvement in detection power has been achieved even with high noise interference. This is even more clearly shown in Figs. 6b and 6c which are the outputs from the

image in Fig. 6a; this is the image of Fig. 3 contaminated with zero-mean white Gaussian noise with $\sigma = 0.1$. All parameter are the same as in the previous example, except that six and eight f_m have been used to detect the "E" (Fig. 6b) and the "A" (Fig. 6c), respectively.

6 Conclusion

Two methods for the rotation-invariant detection of objects in images have been presented. Both are based on the optimization of a functional which achieves a peak at the object location, while balancing between noise suppression and peak sharpness. The first method uses a single linear quasi-invariant filter which optimizes the functional, while the second achieves greater specificity, at the cost of computational complexity, by combining a moderate number of quasi-invariant filters with a simple pattern-matching algorithms. Examples have been presented which compare the algorithms to each other, and to the classical rotation-invariant matched filter.

Among the directions for further research in this approach may be mentioned extending the method to the case of non-uniform prior rotational distribution; to the dilation case (where the prior distribution must be non-uniform); the case of feature, rather than object, detection; the problem of handling perturbation, rather than additive noise; and the problem of extending the present approach to non-linear filtering. These topics are currently under active investigation.

A Proof of Theorem 1

We want to find the maximum of the quadratic functional

$$J(f) = E_\theta \left\{ |(f * s_\theta)(0,0)|^2 \right\}$$

subject to the constraint

$$k_1 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E_\theta\{|(f * s_\theta)(\boldsymbol{x})|^2\} d^2\boldsymbol{x} + k_2 E[(f * n)^2] = 1$$

If we let P_n be the power spectral density of the noise process n, the constraint can be written, in frequency-domain terms, as

$$\frac{k_1}{4\pi^2} E_\theta\{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |F(\boldsymbol{\omega}) S_\theta(\boldsymbol{\omega})|^2\} d^2\boldsymbol{\omega} + k_2 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |F(\boldsymbol{\omega})|^2 P_n(\boldsymbol{\omega}) d^2\boldsymbol{\omega} = 1$$

and, with

$$Q(\boldsymbol{\omega}) = \frac{k_1}{4\pi^2} E_\theta\{|S_\theta(\boldsymbol{\omega})|^2\} + k_2 P_n(\boldsymbol{\omega})$$

and

$$G(\boldsymbol{\omega}) = F(\boldsymbol{\omega}) \sqrt{Q(\boldsymbol{\omega})}$$

the constraint becomes

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |G(\boldsymbol{\omega})|^2 d^2\boldsymbol{\omega} = 1$$

In frequency domain terms , the objective functional becomes:

$$\begin{aligned} J &= E_\theta \left\{ \left| \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(\boldsymbol{\omega}) S_\theta(\boldsymbol{\omega}) d^2\boldsymbol{\omega} \right|^2 \right\} \\ &= \frac{1}{16\pi^4} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(\boldsymbol{\omega}_1) F^*(\boldsymbol{\omega}_2) E_\theta\{S_\theta(\boldsymbol{\omega}_1) S_\theta{}^*(\boldsymbol{\omega}_2)\} d^2\boldsymbol{\omega}_1 d^2\boldsymbol{\omega}_2 \\ &= \frac{1}{16\pi^4} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G(\boldsymbol{\omega}_1) G^*(\boldsymbol{\omega}_2) \frac{E_\theta\{S_\theta(\boldsymbol{\omega}_1) S_\theta{}^*(\boldsymbol{\omega}_2)\}}{\sqrt{Q(\boldsymbol{\omega}_1) Q(\boldsymbol{\omega}_2)}} d^2\boldsymbol{\omega}_1 d^2\boldsymbol{\omega}_2 \end{aligned}$$

and so the problem becomes one of maximizing the positive quadratic form J over all G on the unit sphere.

It is well known that the stationary points of this problem are the eigenvectors of the Hermitian operator

$$R : L_2(\mathbf{R}^2) \longrightarrow L_2(\mathbf{R}^2)$$

defined by

$$R(U) = V$$

where

$$V(\omega) = \frac{1}{16\pi^4} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{E_\theta\{S_\theta{}^*(\omega) S_\theta(\omega_1)\}}{\sqrt{Q(\omega)Q(\omega_1)}} U(\omega_1) d^2\omega_1$$

The maximizing function is then the eigenvector belonging to the largest eigenvalue, and the optimal value of J is this eigenvalue. (It will be shown below that every point, with the possible exception of zero, in the spectrum of R is an eigenvalue.)

To find the non-zero spectrum of R, we first define the operator

$$T : L_2([-\pi, \pi]) \longrightarrow L_2(\mathbf{R}^2)$$

by

$$T(u) = V$$

where

$$V(\omega) = \frac{1}{4\pi^2 \sqrt{2\pi Q(\omega)}} \int_{-\pi}^{\pi} S_\theta{}^*(\omega) u(\theta) d\theta$$

It is easy to show that this defines a bounded linear operator provided

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{|S(\omega)|^2}{Q(\omega)} d^2\omega < \infty$$

and that its adjoint

$$T^* : L_2(\mathbf{R}^2) \longrightarrow L_2([-\pi, \pi])$$

is given by

$$(T^*U)(\theta) = \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{S_\theta(\omega)}{\sqrt{2\pi Q(\omega)}} U(\omega) d^2\omega$$

It then follows that

$$\begin{aligned}
(TT^*U)(\omega) &= \frac{1}{16\pi^4 2\pi \sqrt{Q(\omega)}} \int_{-\pi}^{\pi} S_\theta{}^*(\omega) \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{S_\theta(\omega_1)}{\sqrt{Q(\omega_1)}} U(\omega_1) d^2\omega_1 \\
&= \frac{1}{16\pi^4} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\frac{1}{2\pi} \int_{-\pi}^{\pi} S_\theta{}^*(\omega) S_\theta(\omega_1) d\theta}{\sqrt{Q(\omega)Q(\omega_1)}} U(\omega_1) d^2\omega_1 \\
&= (RU)(\omega)
\end{aligned}$$

since θ is assumed to be uniformly distributed on $[-\pi, \pi]$.

We now use the fact that, with the possible exception of 0, the spectrum of TT^* is the same as that of T^*T. The operator T^*T is given by

$$\begin{aligned}(T^*Tu)(\theta) &= \frac{1}{16\pi^4}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\frac{S_\theta(\boldsymbol{\omega})}{2\pi Q(\boldsymbol{\omega})}\int_{-\pi}^{\pi}S_\psi^*(\boldsymbol{\omega})u(\psi)d\psi d^2\boldsymbol{\omega} \\ &= \int_{-\pi}^{\pi}\frac{1}{32\pi^5}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\frac{S_\theta(\boldsymbol{\omega})S_\psi^*(\boldsymbol{\omega})}{Q(\boldsymbol{\omega})}d^2\boldsymbol{\omega}\, u(\psi)d\psi \qquad (4)\end{aligned}$$

Now, if we use tildes to denote functions in polar coordinates (ρ,ϕ), and use the fact that the random variable θ has a uniform distribution on $[-\pi,\pi]$, we get

$$\begin{aligned}\tilde{Q}(\rho,\phi) &= Q(\boldsymbol{\omega}) \\ &= \frac{k_1}{4\pi^2}\frac{1}{2\pi}\int_{-\pi}^{\pi}|\tilde{S}(\rho,\phi+\theta)|^2 d\theta + k_2 P_n(\rho,\phi)\end{aligned}$$

which is independent of ϕ since the noise is assumed to be isotropic; we will therefore write $\tilde{Q}(\rho,\phi)$ simply as $\tilde{Q}(\rho)$.

In polar coordinates, the integral with respect to $\boldsymbol{\omega}$ in equation 4 then becomes

$$\begin{aligned}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\frac{S_\theta(\boldsymbol{\omega})S_\psi^*(\boldsymbol{\omega})}{Q(\boldsymbol{\omega})}d^2\boldsymbol{\omega} &= \int_0^{\infty}\frac{\rho}{\tilde{Q}(\rho)}\int_{-\pi}^{\pi}\tilde{S}(\rho,\phi+\theta)\tilde{S}^*(\rho,\phi+\psi)d\phi\, d\rho \\ &= k(\psi-\theta)\end{aligned}$$

where

$$k(\psi) = \int_0^{\infty}\frac{\rho}{\tilde{Q}(\rho)}\int_{-\pi}^{\pi}\overline{\tilde{S}(\rho,\phi)}\tilde{S}(\rho,\phi+\psi)d\phi\, d\rho$$

Equation 4 then becomes

$$(T^*Tu)(\theta) = \frac{1}{32\pi^5}\int_{-\pi}^{\pi}k(\psi-\theta)u(\psi)d\psi$$

and, since this is a convolution equation on the circle, it follows immediately that the spectrum of T^*T consists of eigenvalues

$$\begin{aligned}\lambda_m &= \frac{1}{32\pi^5}\int_{-\pi}^{\pi}k(\psi)e^{jm\psi}d\psi \\ &= \int_{-\pi}^{\pi}\frac{1}{32\pi^5}\int_0^{\infty}\frac{\rho}{\tilde{Q}(\rho)}\int_{-\pi}^{\pi}\tilde{S}(\rho,\phi)\tilde{S}^*(\rho,\phi+\psi)d\phi\, d\rho e^{jm\psi}d\psi\end{aligned}$$

$$= \frac{1}{32\pi^5}\int_0^\infty \frac{\rho}{\tilde{Q}(\rho)}\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}\tilde{S}(\rho,\phi)\tilde{S}^*(\rho,\phi+\psi)e^{jm\psi}d\psi d\phi\, d\rho$$
$$= \frac{1}{32\pi^5}\int_0^\infty \frac{\rho}{\tilde{Q}(\rho)}\left|\int_{-\pi}^{\pi}\tilde{S}(\rho,\chi)e^{jm\chi}d\chi\right|^2 d\rho$$

with corresponding eigenvectors

$$u_m(\psi) = e^{jm\psi}$$

Finally, since

$$T^*Tu_m = \lambda_m u_m$$

applying the operator T to both sides gives

$$TT^*(Tu_m) = \lambda_m Tu_m$$

and so each λ_m is an eigenvalue of TT^*, with corresponding eigenvector G_m given by

$$G_m(\boldsymbol{\omega}) = (Tu_m)(\boldsymbol{\omega}) = \frac{1}{4\pi^2\sqrt{2\pi Q(\boldsymbol{\omega})}}\int_{-\pi}^{\pi} S_\theta{}^*(\boldsymbol{\omega})u_m(\theta)d\theta$$
$$= \frac{1}{4\pi^2\sqrt{2\pi Q(\boldsymbol{\omega})}}\int_{-\pi}^{\pi} S_\theta{}^*(\boldsymbol{\omega})e^{jm\theta}d\theta$$

and so the stationary points are given by

$$F_m(\boldsymbol{\omega}) = \frac{K'_m}{\sqrt{Q(\boldsymbol{\omega})}}G_m(\boldsymbol{\omega})$$
$$= \frac{K_m}{Q(\boldsymbol{\omega})}\int_{-\pi}^{\pi} S_\theta{}^*(\boldsymbol{\omega})e^{jm\theta}d\theta$$

This completes the proof; it may be worth remarking that in the case where the probability density function of the random variable θ is nonuniform or the noise is non-isotropic, the solution will have the same form, but the u_m and λ_m will usually have to be calculated numerically by finding the eigenvectors and eigenvalues of an integral equation.

B Circular harmonics and quasi-invariant filters

Throughout this section, polar coordinates will be used, and the tildes which have been used in the paper will be omitted.

The object $s(r,\theta)$ may be expanded in a series of circular harmonics

$$s(r,\theta) = \sum_{l=-\infty}^{\infty} s_l(r)e^{jl\theta}$$

It can then be shown that

$$\mathcal{F}\{s_m(r)e^{jm\theta}\} = S_m(\rho)e^{jm\phi}$$

where

$$S_m(\rho) = \frac{1}{2\pi}\int_0^{2\pi} S(\rho,\phi)e^{-jm\phi}d\phi$$

We therefore have

$$F_m(\rho,\phi) = 2\pi K_m \frac{(S_m(\rho)e^{jm\phi})}{Q(\rho)} \tag{5}$$

Even in the noise-free case, this is not an ideal deconvolution filter; however, if we take the weighted sum

$$\begin{aligned} F(\rho,\phi) &= \sum_l \frac{1}{2\pi K_l} F_l(\rho,\phi) && (6)\\ &= \frac{1}{Q(\rho)}(\sum_l S_l(\rho)e^{jl\phi})^* && (7)\\ &= \frac{S(\rho,\phi)^*}{Q(\rho)} && (8) \end{aligned}$$

Then $F(\rho,\phi)$ is the ideal deconvolution filter for the unrotated object in the noise-free case. We may therefore regard the quasi-invariant filters as the harmonic components of the optimal filter for the unrotated object.

From another point of view, one can imagine obtaining the optimal filter for each circular component directly, using the Schwartz inequality approach:

$$F_m(\rho,\phi) = K_m \frac{(S_m(\rho)e^{jm\phi})^*}{\frac{k_1}{4\pi^2}|S_m(\rho)|^2 + k_2 P_n(\rho)}$$

In the noise-free case, this would be an ideal deconvolution filter for the corresponding component. This component, however, never occurs its own, and the quasi-invariant filter may be regarded as adding the energy of the other components in the denominator to account for the fact that they always appear in conjunction with the given component. From this point of view, the quasi-invariant filter may be regarded as the optimal filter for a given circular component, treating all of the other components as unavoidable noise.

References

[1] Y. Hsu. Rotation-invariant digital pattern recognition using circular harmonic expansion. *Applied Optics*, vol. 21(no. 22):pp. 4012–4015, 1982.

[2] R. Wu and H. Stark. Rotation invariant pattern recognition using a vector reference. *Applied Optics*, vol. 21:pp. 838–840, 1984.

[3] G. Shils and D. Sweeney. Interactive technique for the synthesis of optical-correlation filters. *J. Opt. Soc. Amer.*, vol. A3:pp. 1433–1441, 1986.

[4] T. Caelli and Z. Liu. On the minimum number of templates required for shift, rotation and size invariant pattern recognition. *Pattern Recognition*, vol. 21(no. 3):pp. 205–216, 1988.

[5] M. K. Hu. Visual pattern recognition by moment invariant. *IEEE Trans. Inform. Theory*, vol. C-26:pp. 179–187, 1962.

[6] M. R. Teague. Image analysis via the general theory of moment. *J. Opt. Soc. Amer.*, vol. 70:pp. 920–930, 1980.

[7] Y. Abu-Mostafa and D. Psaltis. Recognition aspects of moment invariants. *IEEE Trans. PAMI*, vol. 6(no. 6):pp. 698–706, 1984.

[8] C-H. Teh and R. T. Chin. On image analysis by the method of moment. *IEEE Trans. PAMI*, vol. 10:pp. 496–513, 1988.

[9] B. Kumar and E. Pochapsky. Signal-to-noise ratio considerations in modified matched special filters. *J. Opt. Soc. Amer.*, vol. 3(no. 6):pp. 777–785, 1986.

[10] C. F. Hester and D. Casasent. Multivariant technique for multi-class pattern recognition. *Applied Optics*, vol. 19:pp. 1758–1761, 1980.

List of Captions

Figure 1: Input Image with Low Noise

Figure 2: a)Output from Optimal Square Filter; b)Output from Optimal Triangle Filter; c)Output from Matched Filter for Square; d)Output from Matched Filter for Triangle

Figure 3: More Complex Input Image

Figure 4: a)Output from Optimal "E" Filter; b)Output from Optimal "A" Filter; c)Output from Matched Filter for "E"; d)Output from Matched Filter for "A"

Figure 5: a)Input Image with High Noise; b)Output from Three-Component Square Filter; c)Output from Four-Component Triangle Filter

Figure 6: a)More Complex Input Image with High Noise; b)Output from Six-Component Filter for "E"; c)Output from Eight-Component Filter for "A"

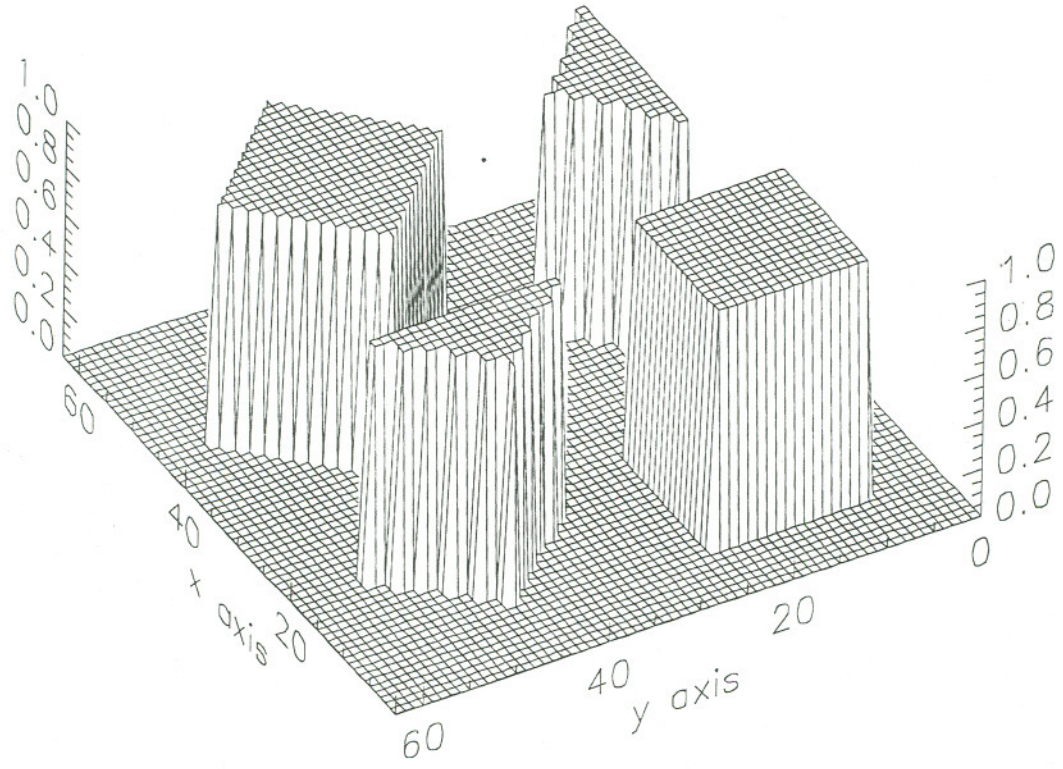

Figure 1.

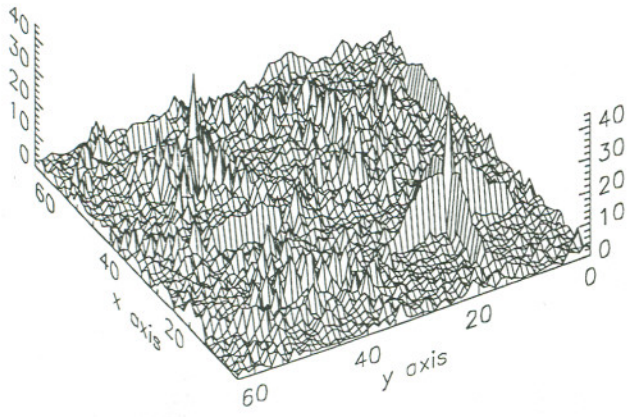

a)

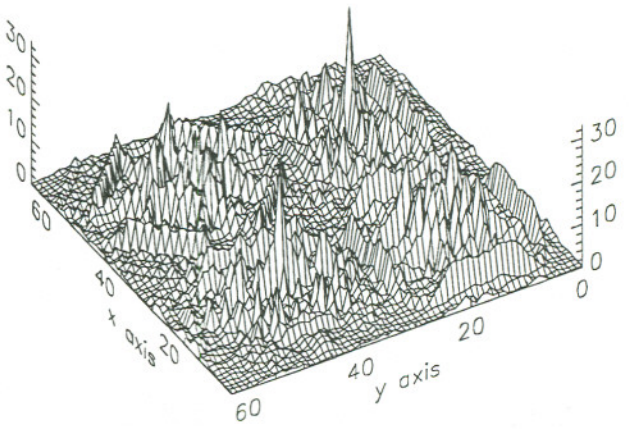

b)

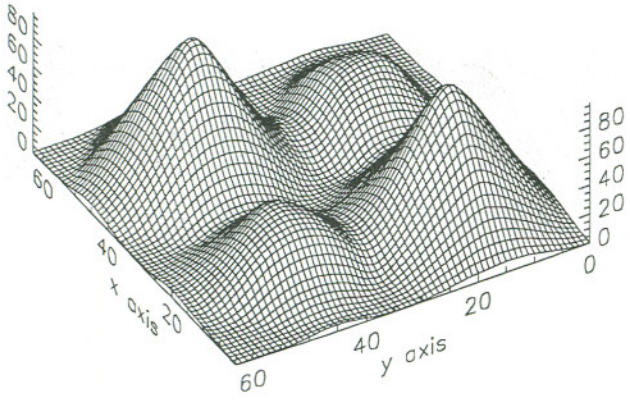

c)

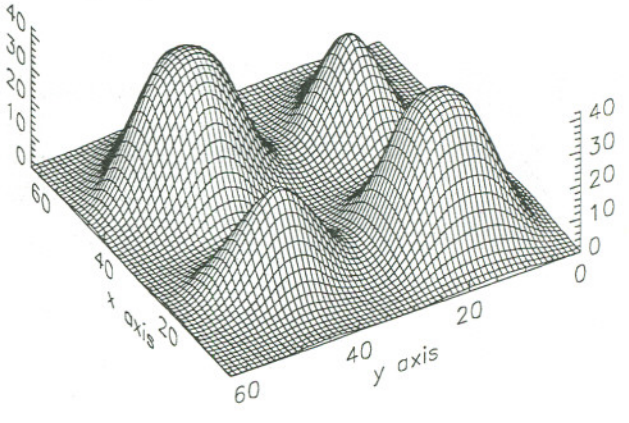

d)

Figure 2.

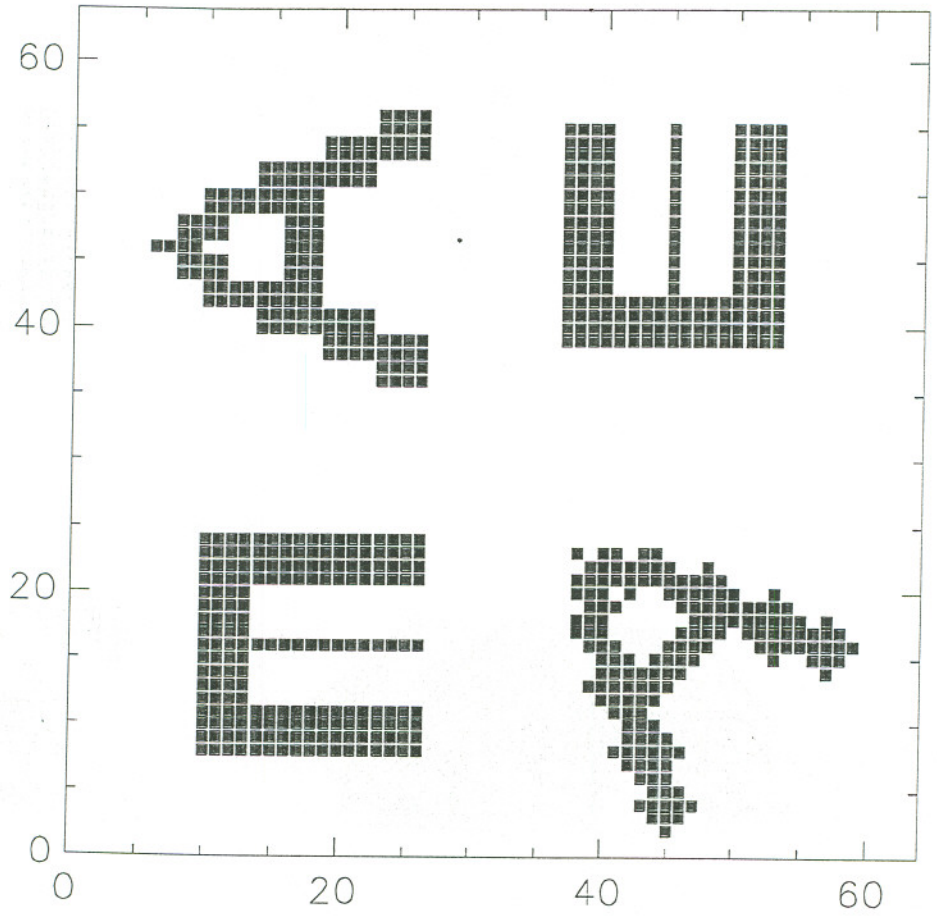

Figure 3.

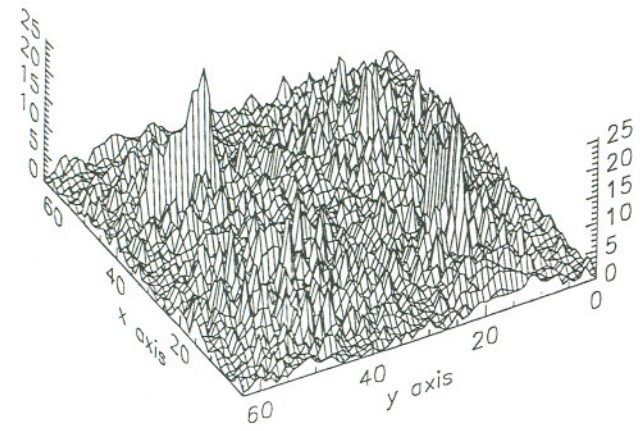

a)

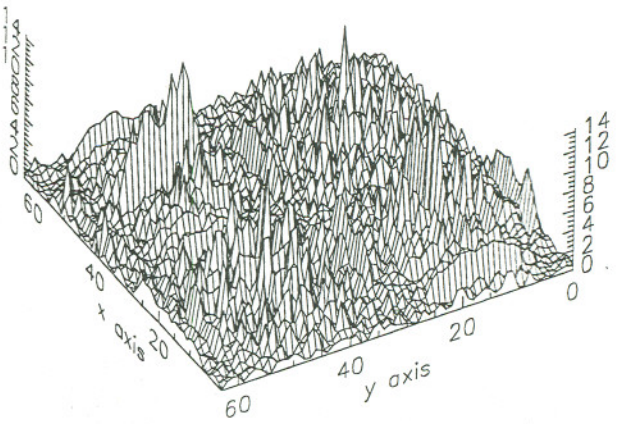

b)

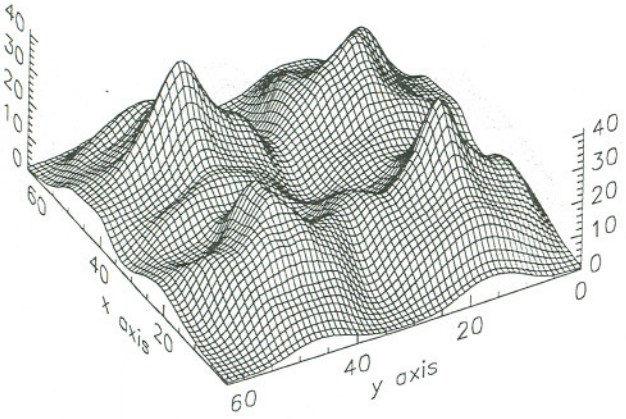

c)

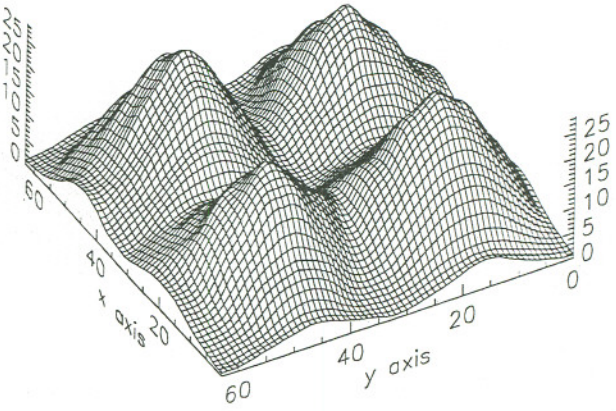

d)

Figure 4.

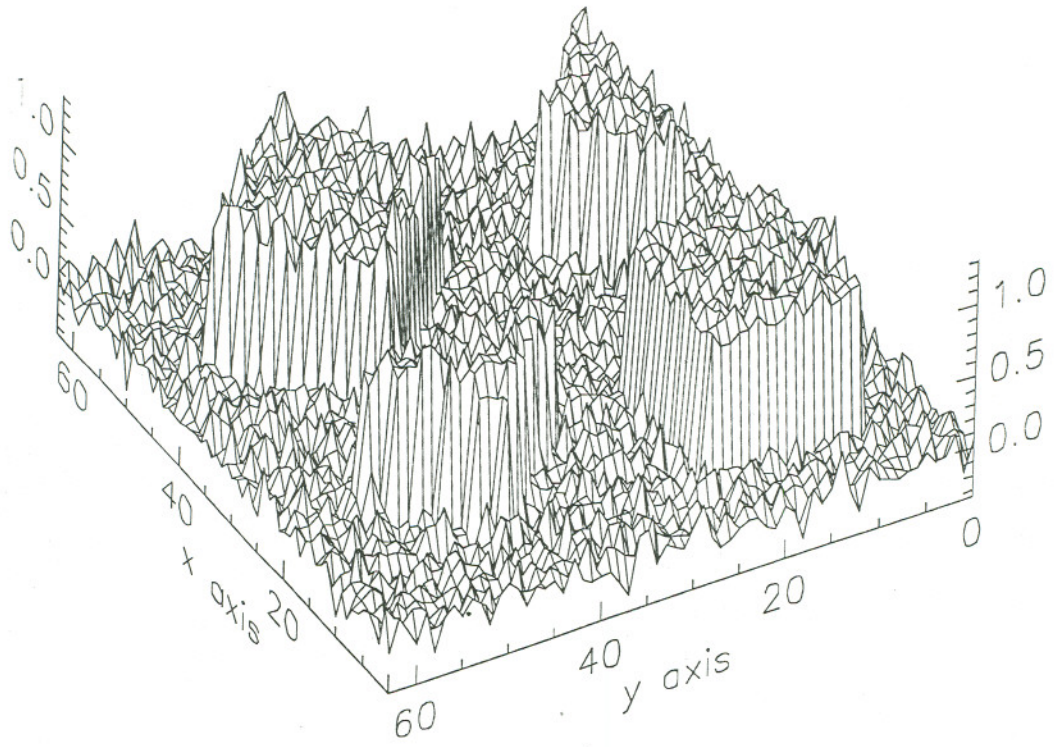
1.0
0.5
0.0
60
40
x axis
20
0
20
40
y axis
60

a)

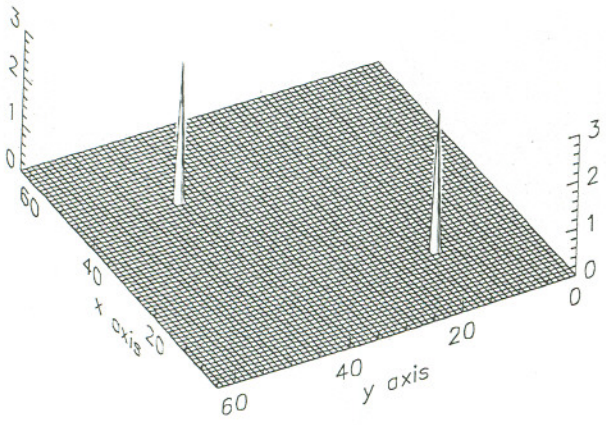
3
2
1
0
60
40
x axis
20
0
20
40
y axis
60

b)

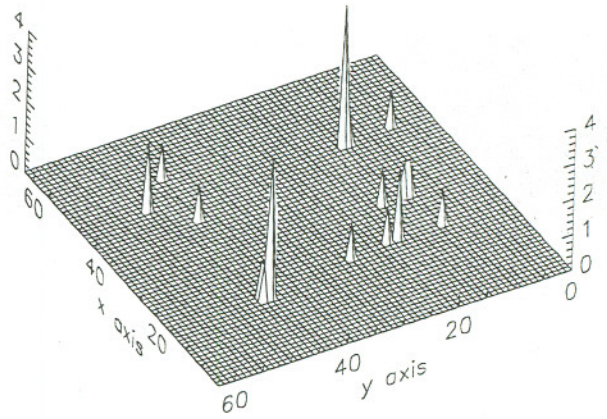
4
3
2
1
0
60
40
x axis
20
0
20
40
y axis
60

c)

Figure 5.

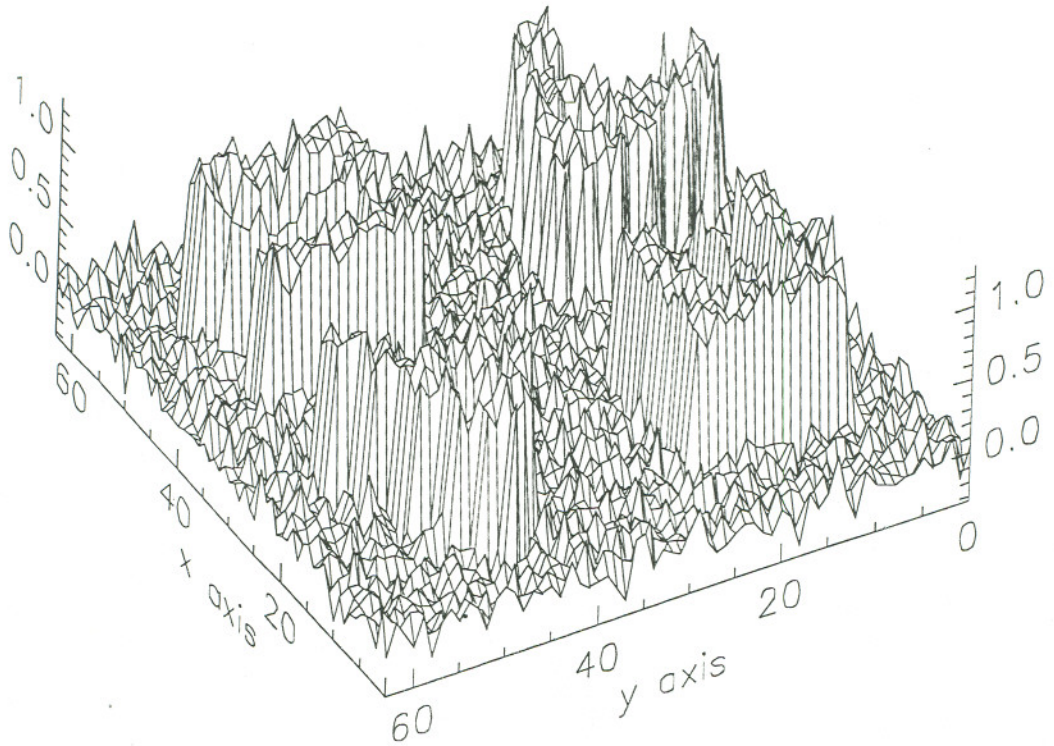
1.0
0.5
0.0
60
40
x axis
20
60
40
y axis
20
0
1.0
0.5
0.0

a)

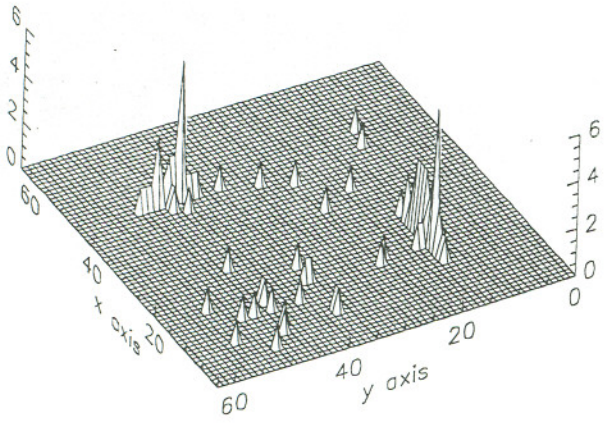
6
4
2
0
60
40
x axis
20
60
40
y axis
20
0
6
4
2
0

b)

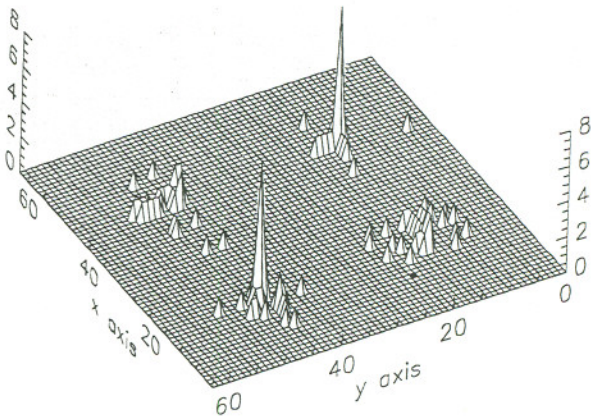
8
6
4
2
0
60
40
x axis
20
60
40
y axis
20
0
8
6
4
2
0

c)

Figure 6.