AD-A257 163

RL-TR-92-174
Final Technical Report
June 1992

OPTICAL FILTERING AND CORRELATOR EVALUATION

Florida Institute of Technology

Samuel P. Kozaitis
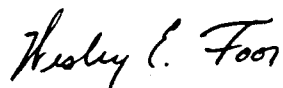
DTIC
S ELECTE
E NOV 0 5 1992
D

92-28917

**Rome Laboratory
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700**

92 11 04 063

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-92-174 has been reviewed and is approved for publication.

APPROVED:

Wesley E. Foor

WESLEY E. FOOR
Project Engineer

FOR THE COMMANDER

JAMES W. YOUNGBERG, Lt Col, USAF
Deputy Director
Surveillance and Photonics Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 1992	Final Jan 91 – Sep 91

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
OPTICAL FILTERING AND CORRELATOR EVALUATION	C – F30602-88-D-0028, Task P1-6013
6. AUTHOR(S)	PE – 62702F
Samuel P. Kozaitis	PR – 4600
	TA – P1
	WU – PC

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Florida Institute of Technology Department of Electrical and Computer Engineering Melbourne FL 32901-6988	N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
Rome Laboratory (OCPA) Griffiss AFB NY 13441-5700	RL-TR-92-174

11. SUPPLEMENTARY NOTES

Rome Laboratory Project Engineer: Wesley E. Foor/OCPA/(315) 330-2944

12a. DISTRIBUTION/AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for public release; distribution unlimited.	

13 ABSTRACT (Maximum 200 words)

In an optical correlator, binary phase-only filters (BPOFs) that recognize objects that vary in a nonrepeatable way are essential for recognizing objects from actual sensors. An approach is required that is as descriptive as a BPOF yet robust to object and background variations of an unknown or unrepeatable type. We investigated the use of BPOFs to recognize objects undergoing unknown variations. In our study, we developed a BPOF that was more robust than other designs. This was done by creating a filter that retained the invariant features of a training set. Our feature-based filter offered a range of performance by setting a parameter to different values. At one extreme, the filter offered similar performance to that of a synthetic discriminant function (SDF) filter. As the value of the parameter was changed, correlation peaks within the training set became more consistent and broader. In addition, the feature-based filter was potentially useful for recognizing objects outside the training set. Furthermore, the feature-based filter was more easily calculated and trained than an SDF filter.

14. SUBJECT TERMS		15 NUMBER OF PAGES
Pattern Recognition, Optical Correlation, Binary Phased-Only Filter		28
		16 PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

1 Introduction

The major difficulty encountered when using a binary phase-only filter (BPOF) in an optical correlator is its sensitivity to changes in the object's appearance. In imagery from actual sensors, the same object can vary significantly depending on aspect angle, lighting, atmospheric effects, and a host of other variables. In addition, object boundaries may be poorly defined and buried in the background. Identifying an object that has a nonrepeatable signature is one of the key technical challenges of automatic object recognition.[1]

An optical correlator that uses binary spatial light modulators (SLMs) requires the conversion of sensor imagery to binary imagery. The conversion process is highly vulnerable to noise and variations of the object and background. Therefore, an object can appear differently after the image is converted to a binary image due to environmental or other conditions. The reliability of the conversion process is critical for object recognition because the binary image contains shape features of the object. In real imagery, the global shape of an object is frequently too perturbed to generate a reliable, specific version of the object. The binary result is often a version of the object that changes in an unknown or nonrepeatable way.

BPOFs are useful for recognizing fixed objects in stationary backgrounds. BPOFs exhibit large and narrow correlation peaks and are effective for multi-class discrimination.[2,3] They can work well in the presence of background clutter or when an object is partially obscured.[4,5] BPOFs have provided suitable solutions when objects have a repeatable signature. If an object varies in a limited or known manner, more complex filters can be used.

Distortion-invariant filters, including synthetic discriminant function (SDF) filters, circular harmonic filters, and lock and tumble filters have been developed for optical correlators.[6-9] Methods for directly creating distortion-invariant BPOFs have also been developed.[10-13] These approaches have been primarily used to identify rotated or scaled versions of an object and to specify the response of in-class and out-of-class objects. For example, the SDF method uses training images that are distorted versions of the reference image. A combination of training images is used to create a filter that will produce the same cross-correlation result with all training images. Using these techniques, BPOFs can be created that can perform well when objects vary in a known manner.

In reality, imagery from sensors is often nonrepeatable; the same object can vary in appearance depending on weather conditions, lighting, and other variables. These objects, whose distortion cannot be repeated exactly, are said to vary in an unknown manner. The creation of BPOFs to recognize objects that vary in a nonrepeatable way is essential for recognizing objects from actual sensors. Recently, approaches based on SDF filters have been applied to areas such as handwriting verification, and aircraft identification.[14-15] In one study, many variations of an pattern was incorporated into the filter to obtain robustness.[14] Another approach used digital image processing techniques to confine images to vary in a limited but unknown manner, then filters with a small number of training images were used to identify objects.[15] In both cases, an approach is required that is as descriptive as a BPOF yet robust to object or background variations of an unknown or nonrepeatable type.

We investigated the use of BPOFs to recognize objects undergoing unknown variations. We also derived a filter that was calculated from objects' features. By attempting to recognize an object

based on its features, we made a BPOF more robust. To help evaluate the potential of our approach, we used imagery from actual sensors that were not from the original training set. The next section briefly discusses thresholding and is followed by SDF filter formation in the following section. Details of a filter made from features of a training set are in Section 4, and Section 5 presents the performance of the filter. Finally, discussion and conclusions are presented in Sections 6 and 7 respectively.

2 Preprocessing

Infrared (IR) imagery (8-14 μm) of ground scenes from actual sensors were used to evaluate filters. Images were digitized with 128 x 128 pixels with 8 bits/pixel. Because the application was to binary SLMs, the imagery ultimately had to be thresholded.

Thresholding was performed by choosing a single threshold value for the entire image. Threshold values can be chosen several different ways. They can be based on the noise statistics of an image, a histogram of the image, or a fixed value chosen near the middle of the available pixel values.[16] When the object and background are within an image are obvious, a threshold value can be easily chosen, and different methods usually give similar results. Other techniques can be used if the imagery is more complex.[17] In either case, a thresholding method should be automatic in that it should perform similarly with a variety of imagery under various lighting and atmospheric conditions.

In the imagery we examined, the background and object were easily separated; however, edges were not well-defined. We used digital image processing techniques to implement an isodata thresholding method.[18] This technique examines peak values in the histogram of an image. A threshold value was chosen between peaks that were associated with the object and the background so that the object could be segmented from the background. If noise or atmospheric distortion is present, the peaks of the histogram will change their position or shape, but the peaks are usually identified. Choosing a threshold value between peaks of a histogram often results in an image that is similar to the silhouette of the object. As variables such as lighting, noise, and atmospheric effects within an image change, the resulting thresholded images will remain similar but will often be different in an unknown or nonrepeatable way.

3 Filtering Methods
3.1 Binary Phase-Only Filters Made From Single Objects

When attempting to recognize objects that changed from the filter image, we found that various types of BPOFs made from a single object had similar performance. For example, the thresholding line angle (TLA) of a BPOF can be chosen to improve its performance. Sometimes choosing a specific TLA is equivalent to creating a specific filter such as that obtained with the Hartley transform;[19] however, choosing an optimal TLA seems to increase the SNR by only 1-2dB.[20] Furthermore, BPOFs can oe designed to yield a maximum SNR by limiting their bandwidth,[21] or by identifying an optimum support function (pixels that are not set to zero) in the Fourier plane.[22] By choosing the optimal support function, the SNR can be increased from 4-6 dB. All of these filters are optimized for a specific input. There is nothing to suggest that any of these filters will correlate well when an input undergoes a change with respect a filter image.

We chose a TLA of zero degrees because others have achieved satisfactory results using this value.[22] Our BPOFs were created in the following way. The Fourier transform of a desired image was computed and its phase was set to either 0 or π at each pixel. If the calculated phase angle was between 0 and π, then the phase at that pixel was set to 0. If the calculated phase angle was between π and 2π, then the phase at that pixel was set to π.

We considered an IR photograph like that described in the previous section that contained a total of ten aircraft. A portion of the photograph is shown in Fig. 1. We created ten 128 x 128 images from the photograph where each aircraft could fit in a square of about 90 x 90 pixels and labeled the images image1 - image10. A BPOF was made from each of the ten aircraft and each filter was cross-correlated with all ten images. In some cases, aircraft were rotated so that the orientation of the input and filter aircraft were the same. With all filters, the autocorrelation response produced the highest correlation peak with some of the other responses potentially being useful. The data for the ten filters were similar, so the results from only the case when image1 was used as a filter is shown in Fig. 2. The data was normalized so that the correlation heights were expressed as a percentage of the autocorrelation height. For grey-leveled inputs, the maximum response other than the autocorrelation response was 45.3% of the autocorrelation height. Because the application is to binary SLMs, we converted the input images to binary as discussed in the previous section and repeated the experiment. In the case of thresholded imagery, the maximum correlation height other that the autocorrelation decreased to 24.9%. A few of the responses due to the grey-leveled images could be useful, but the responses due to the thresholded imagery were generally not useful. This implies that more filters or more complex filters are needed to recognize these images as the same class. As previously reported, the conversion to binary is an important factor that can make images that appear to be similar before thresholding produce poor correlation results after thresholding.[15,23]

Figure 1 Portion of IR image showing some of the images used in our experiment. Thresholded versions of the five aircraft were used as training images for other filters.
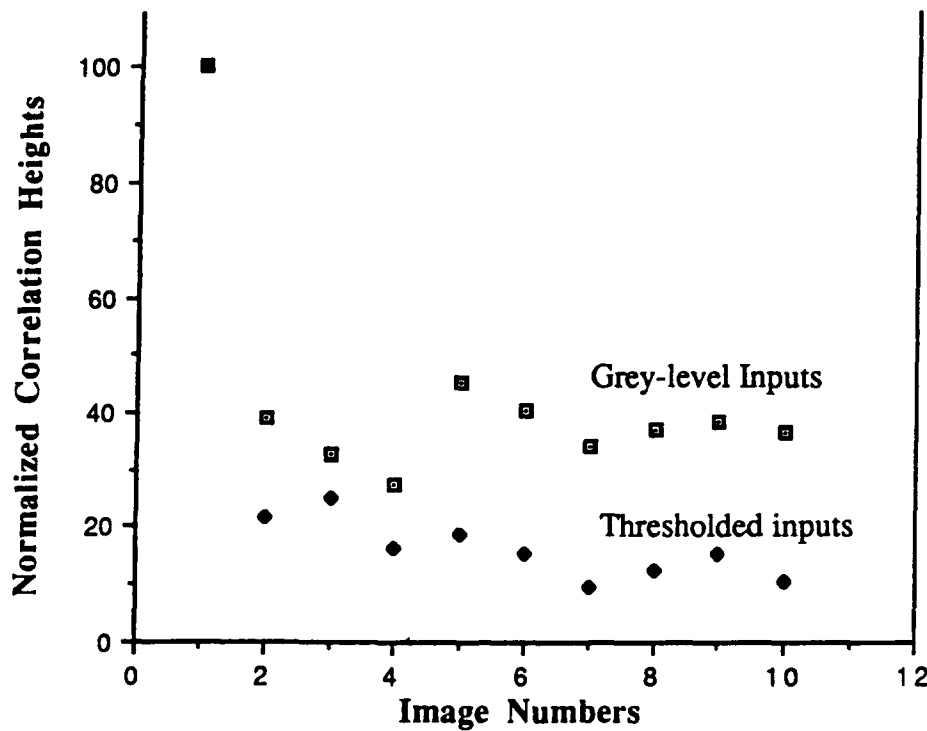
Figure 2 Graph showing correlation results using a BPOF for similar gray-level images before and after thresholding.

3.2 Synthetic Discriminant Function Filters

This section contains a brief discussion of how the SDF filters we used were made. A more detailed description of the approach we followed has been previously reported.[24] The main advantage of the SDF filter approach is that the need for displaying different filters is reduced or eliminated. A conventional SDF is a combination of images that can be described as

$$s(x,y) = \sum_n a_n t_n(x,y)$$

(1)

where t_n are centered training images and a_n are weight coefficients. SDF synthesis techniques may be used to determine the weight coefficients.[6,25] The complex conjugate of the Fourier transform of $s(x,y)$ is the matched filter

$$S(u,v) = F[s(x,y)]^*$$

(2)

where F is the Fourier transform operator.

Converting the SDF-matched filter to a BPOF may result in a severe information loss. Recently, an improved version of an SDF, called a filter SDF (fSDF), has been introduced that includes the function modulation characteristics of the device onto which the filter is mapped in the synthesis

equations.[24] The significance is that virtually any type of filter can be made into an SDF filter using this approach. For fSDF-BPOFs, the coefficients an can be iterated based on the formula

$$a_n^{i+1} = a_n^i + \beta \left[c_n - c_0 \left(\frac{m_n^i}{m_0^i} \right) \right]$$

(3)

where i is the iteration number, β is a damping constant, and m_n^i is the modulus of the peak correlation response of image $t_n(x, y)$ with a filter made with the coefficient vector a^i. In the experiments described in this paper, the initial solution vector was taken to be the desired correlation response vector, $a^0 = c = 1$. The initial fSDF, $s(x,y)$, was then found and cross-correlated with each training image. The values of the correlation heights were then placed into Eq. (3), and an updated version of a was found. A new $s(x,y)$ was found and the procedure is repeated. The modulation characteristics of an a BPOF SLM was included by calculating intermediate cross-correlation in Eq. (3) with BPOFs.

We used five thresholded images that were derived from Fig. 1 as training images to create an fSDF filter. We also obtained nine additional images (image11 - image19) from another photograph that contained the same type of aircraft. We then cross-correlated the filter with the training set and the remaining fourteen images. The data are shown in Fig. 3 and are normalized such that the largest value of the correlation height was set to 100. The results for imagery within the training set were useful for recognition with average and minimum values of 70.8 and 45.3% of the maximum respectively. The fSDF filter produced useful results for images within the training set but the results were not useful for data outside the training set. The training set would have to be expanded to obtain useful results for image11 - image19. The peak-to-root mean square ratio (PRMSR) can be used as a measure of sharpness of a correlation peak and is the ratio of the maximum intensity of the correlation peak divided by the rms value of all pixels below half the height of the correlation peak.[26] For objects in the training set, sharp peaks were produced with an average PRMSR of 15.3
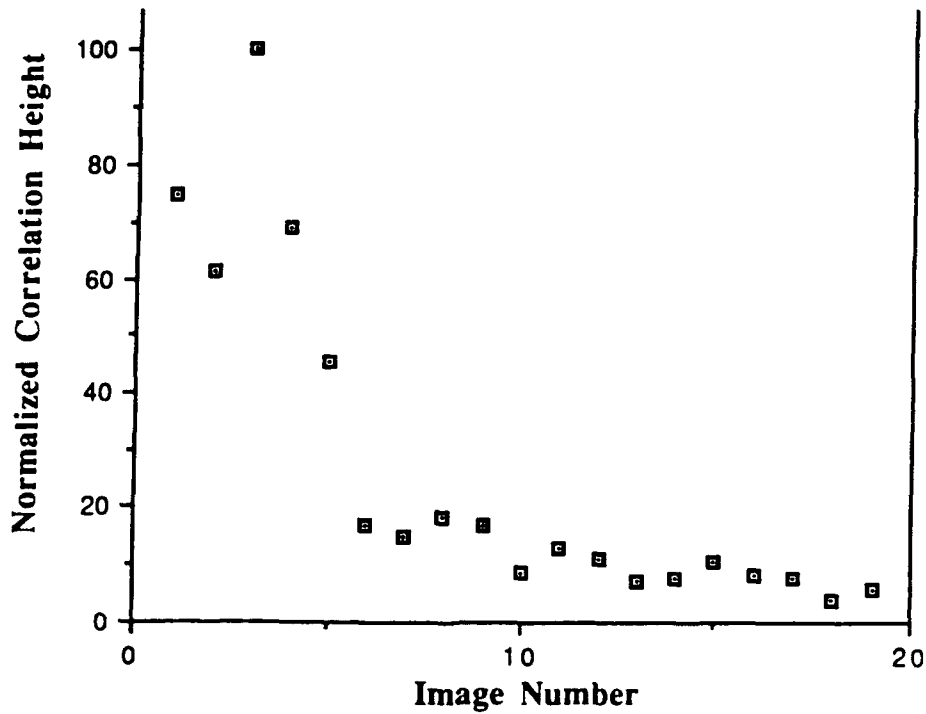
Figure 3 Graph of normalized correlation heights for various imagery using an fSDF filter where image1 - image5 belong to the training set.

4.0 Feature-based filtering

Generally, the cross-correlation between a input image and an image used to make a filter is maximized when the mean squared error (MSE) between the two images is minimized.[27] In the formulation of an SDF filter, the MSE between each training image and the filter image are minimized subject to the constraint that all the MSEs are equal. This is accomplished by multiplying each input image by a weight value. The idea behind the fSDF filter is similar to that of the SDF filter; however, the weights are chosen to produce the desired fil..er after the SDF has been quantized to form a BPOF.

In contrast to the SΓᴿ filter formulation, we developed a filter whose values were determined by features of the training set. We attempted to find a filter that represents the critical characteristics of an object so that objects outside the training set could be identified. Therefore, we examined features that were invariant with respect to the training set. Because correlation filters are derived from the Fourier transform of an object, we examined feature extraction in the Fourier domain. Optical feature extraction has been previously performed in the Fourier domain by using the power spectrum of an object as features to train a neural network for classification.[28,29] Because BPOFs have provided suitable solutions for recognizing objects, we viewed the BPOFs from several objects as a collection of features to create a new filter for an optical correlator. We retained those features that were invariant among a training set by examining all filters at the same pixel on a pixel by pixel basis. We rejected those features that were not invariant for a particular distortion. For example, if the same value occurred in the BPOFs at a particular pixel, we considered that value for that pixel

in a new BPOF. In this way we constructed a new BPOF that represented features of the objects in the training set.

The main difference between the feature-based approach and the fSDF approach to filter generation can be seen by examining Fig. 4. We considered the values of the discrete Fourier Transform (FT) at a pixel (i, j) for four training images. Three of the values are in the vicinity of the real axis which we assume to be our thresholding axis for the BPOF. The fourth value is well into the lower quadrant. Because the SDF approach attempts to minimize the MSE error among the four points subject to the constraint that the total error from all pixels for each training set be the same; generally, the value of a particular pixel in an SDF filter will be in the vicinity of the average of the values for that pixel. From Fig. 4, it can be seen that any value in the vicinity of the average of the four values will be in the lower quadrant and ultimately represented by a -1 value in a BPOF. Had the BPOFs of each of the three training images been considered separately, the values of that pixel would have been, +1, +1, +1 and -1. If we considered these values of the BPOF as features of the object, we could see that the +1 value occurs more often for this training set. Therefore, we would consider a +1 value for this pixel because more +1's occurred than -1's. In this way, we considered the BPOFs of a training set as a collection of features and created a filter based on these features.
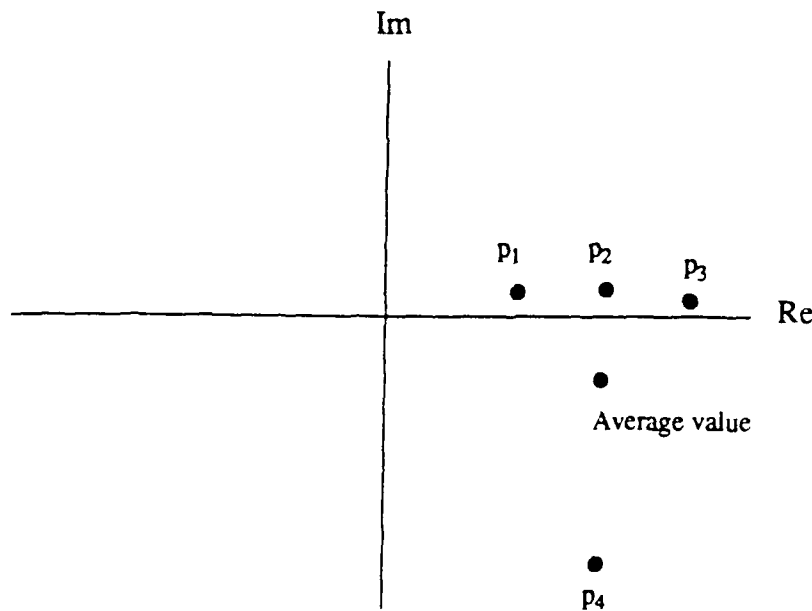
Im

Figure 4 Fourier plane representation at pixel (i, j) for four images.

Sometimes features may offer conflicting information.[1] For example, in Fig. 4, two points could have been above the real axis and two below. Other than arbitrarily choosing a value to represent these values or using an odd number of training images, we could avoid the use of this feature by setting the pixel value to zero. In this way, we would not use features that offered conflicting information.

Our filters were generated from n, N x N training images as follows. The BPOF for each training image was calculated. Then, the BPOFs of all the training images are added together pixel by pixel. The result was an N x N array of integer values. If the value at a pixel of the array was greater than a positive-valued threshold, then that pixel was set to +1. If the value at a pixel of the array was less than the negative of the threshold, then the pixel was set to -1, otherwise the pixel was set to zero. The filter was described as

$$
G(u, v) = \left\{
\begin{array}{c}
1 \ if \ p \leq \sum_{n} H_n(u, v) \\[2mm]
-1 \ if \ (-p) \geq \sum_{n} H_n(u, v) \\[2mm]
0 \ otherwise
\end{array}
\right\}
\tag{4}
$$

where $H_n(u, v)$ is the BPOF of a training image $t_n(x, y)$ and p is the threshold that determines the order of the filter.

5.0 Experiment

We tested our feature-based filter by computer simulation using the same imagery that was used with the fSDF filter. A comparison of the normalized correlation heights for the fSDF filter and the feature-based filter with $p = 1$ is shown in Fig. 5. The maximum value for each filter was set to 100; note that the maximum value of the feature-based filter was 90.8% of the maximum of the fSDF filter. The filters had similar performance. For images within the training set, the average correlation height produced by the fSDF and feature-based filters were 70.1 and 76.4% the maximum respectively. The PRMSR values for images within the training set were 11.6 for the feature-based filter compared to the 15.3 for the fSDF filter. Neither filter produced very useful correlation heights for images outside the training set. The feature-based filter produced more consistent results than the fSDF filter at the expense of slightly decreased and broader correlation peaks.
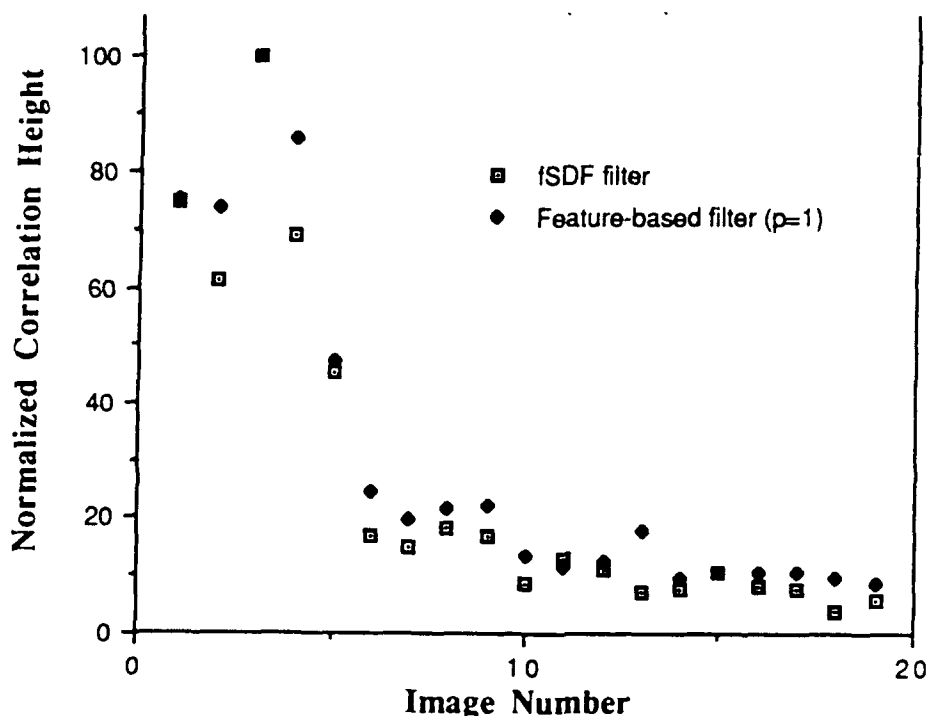
Figure 5 Comparison of correlation heights obtained with fSDF filter and a feature-based filter with $p = 1$, for a variety of test imagery. Note that only image1 - image5 are in the training set

A comparison of the normalized correlation heights between the fSDF filter and feature-based filter with $p = 3, 5$ is shown in Fig. 6. The maximum value for each filter was set to 100 as was in the previous experiment. For images within the training set, the average correlation height produced by the feature-based filters for $p = 3, 5$ were 78.3 and 94.5% the maximum respectively and the PRMSR values were 9.6 and 3.2 respectively. For $p = 3$, the correlation peaks were slightly broader than the $p = 1$ case; however, the $p = 5$ filter had significantly broader peaks than the $p = 1$ case. The maximum values of the correlation height for $p = 3, 5$ were 52.8 and 11.1% of the maximum of the fSDF filter respectively.

In contrast to the fSDF filter, some correlation responses for images outside the training set were useful using a feature-based filter with $p = 3, 5$. For the feature-based filter with $p = 3$, some correlation heights were above 30% the maximum. For $p = 5$, some values were above 60%. As in the previous experiment, the feature-based filter produced more consistent results than the fSDF filter at the expense of decreased and broader correlation peaks.
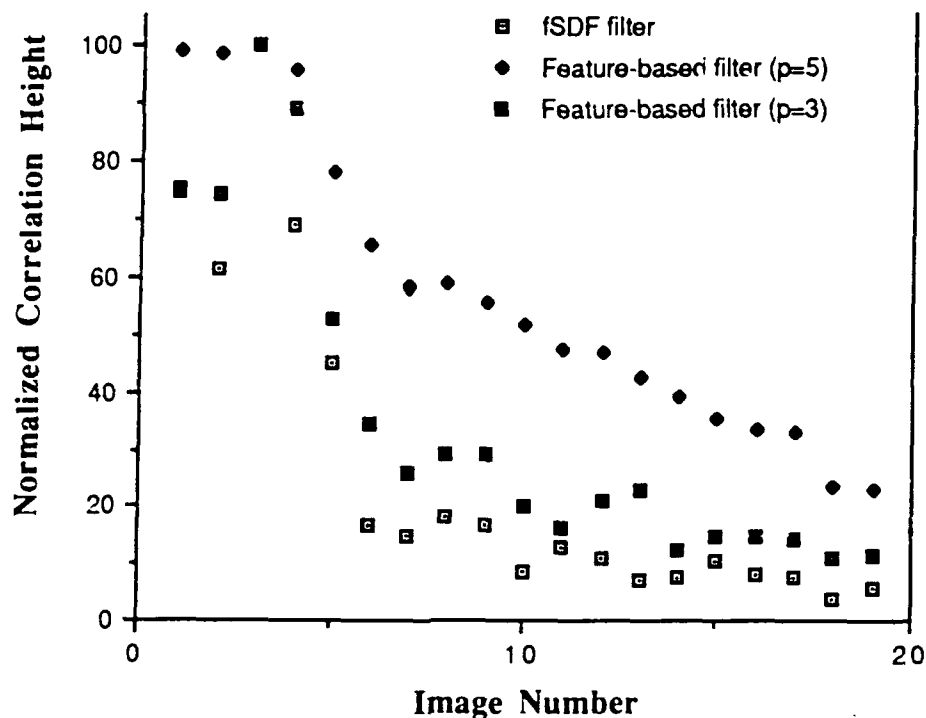
Figure 6 Comparison of correlation heights obtained with feature-based filters with p =3, 5, and an fSDF filter for various imagery. Note that only image1 - image5 are in the training set.

The correlation peaks for the filter when p = 5 were broad when compared to other filters with different values for p or the fSDF filter. In an effort to improve the PRMSR values, we considered blocking pixels of the filter near the DC component. The blocking of the DC is consistent with our approach because all objects will have a DC component; therefore, this feature will not be useful for discrimination and can be eliminated. We blocked the central pixel of the filter where p = 5 and compared its results to the case where the central pixel was not blocked. The results are shown in Fig. 7. The normalized correlation heights are almost the same in both cases but the PRMSR values had increased from 3.2 to 4.2. We also blocked a 3 x 3 pixel area of the filter around the DC and obtained the results shown in Fig. 7. The PRMSR further increased to 5.3 for the training set but the correlation heights generally decreased. A summary of the performance of the various filters are shown in Table 1. As fewer but more consistent features were retained, the correlation responses were more consistent but the correlation response became broader. By blocking the pixels around the DC component, the average correlation height of the training set was made sharper at the expense of less consistent correlation heights. As an example of the correlation results obtained with the different filters, the cross-correlation responses for image4 and the six filters are shown in Fig. 8 for comparison; the height of each response was set to the same value.
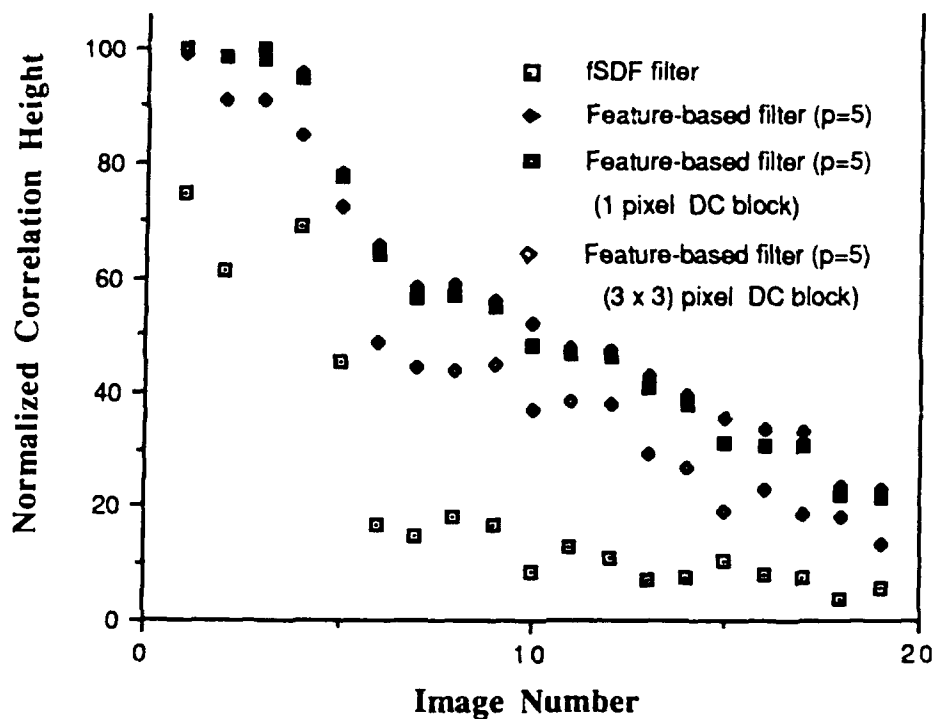
Figure 7 Comparison of correlation heights obtained with fSDF filter and feature-based filters with $p = 5$ for a variety of test imagery. Note that only image1 - image5 are in the training set.

Table 1: Performance of Various Filters for Images in Training Set

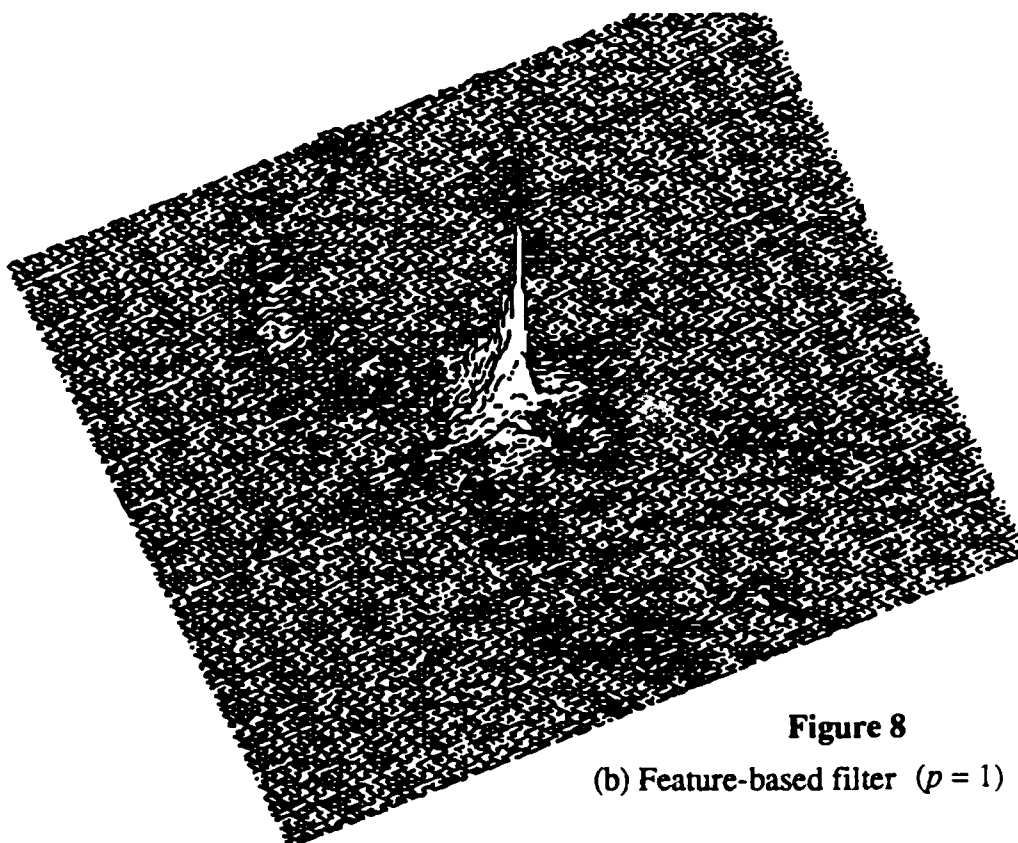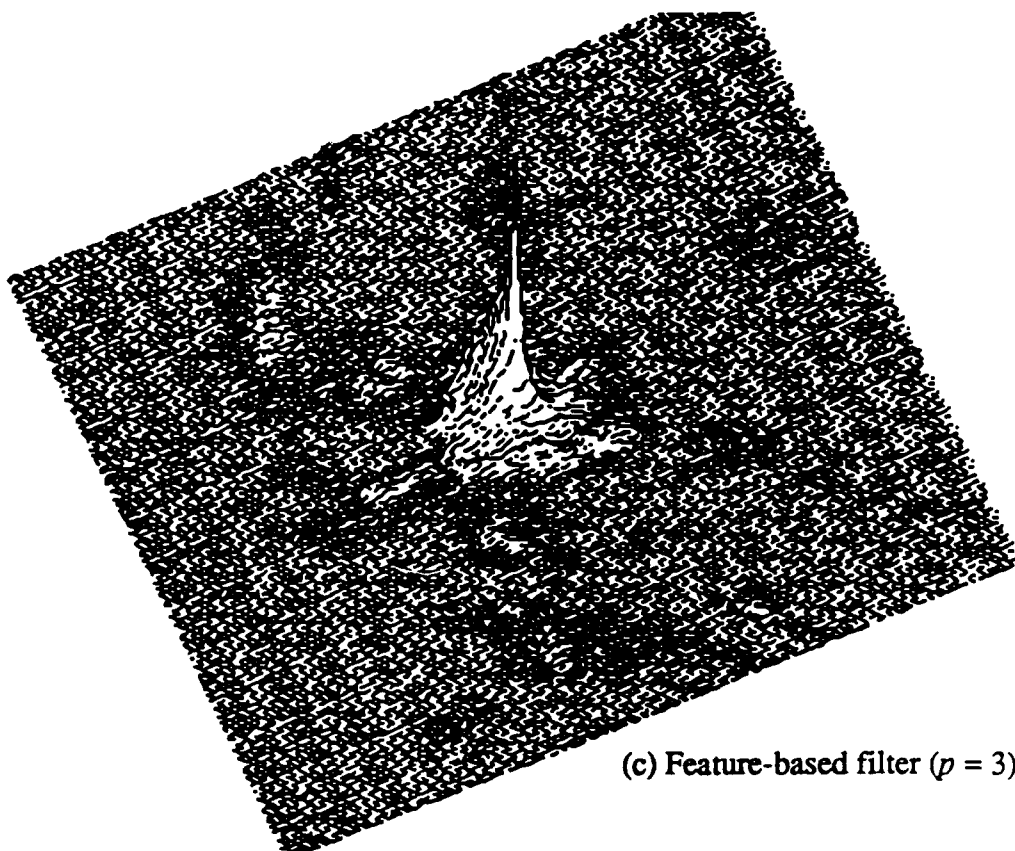| Filter	Average Normalized Correlation Height	Average PRMSR value
fSDF	70.8	15.3
Feature-based ($p=1$)	76.4	11.6
Feature-based ($p=3$)	78.3	9.6
Feature-based ($p=5$)	94.5	3.2
Feature-based ($p=5$) 1 pixel DC block	94.0	4.2
Feature-based ($p=5$) 3x3 pixel DC block	87.8	5.3
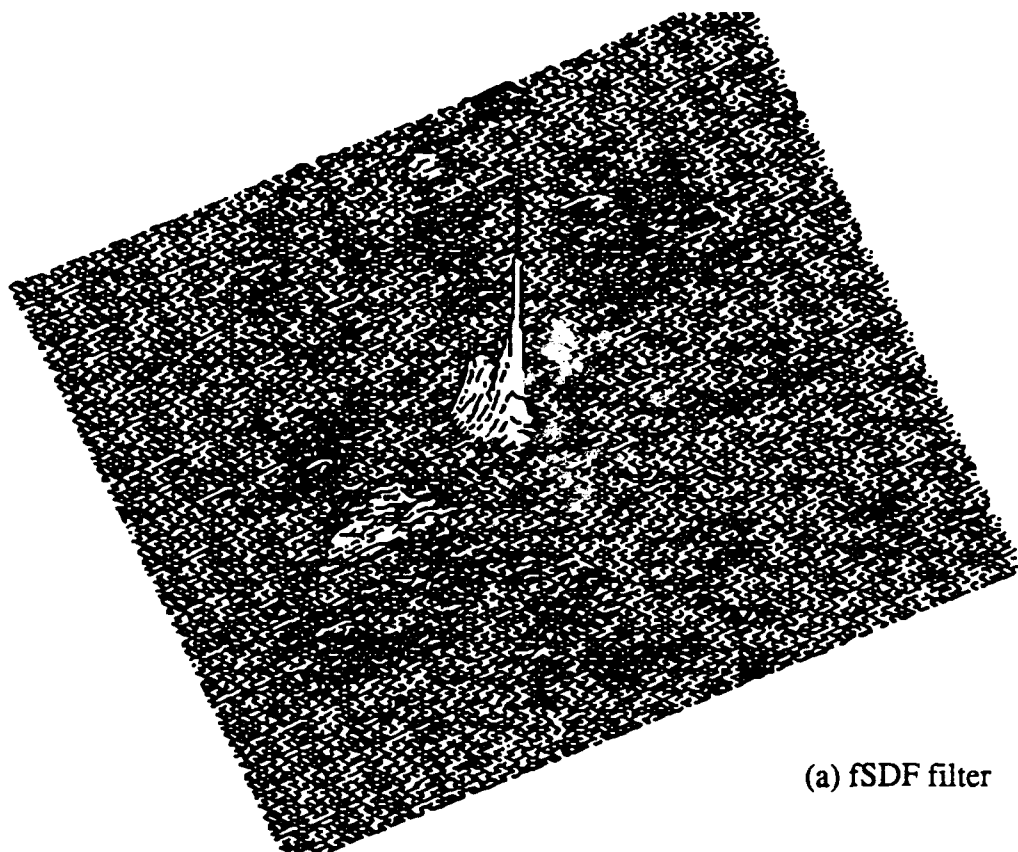
(a) fSDF filter

Figure 8

(b) Feature-based filter $(p = 1)$
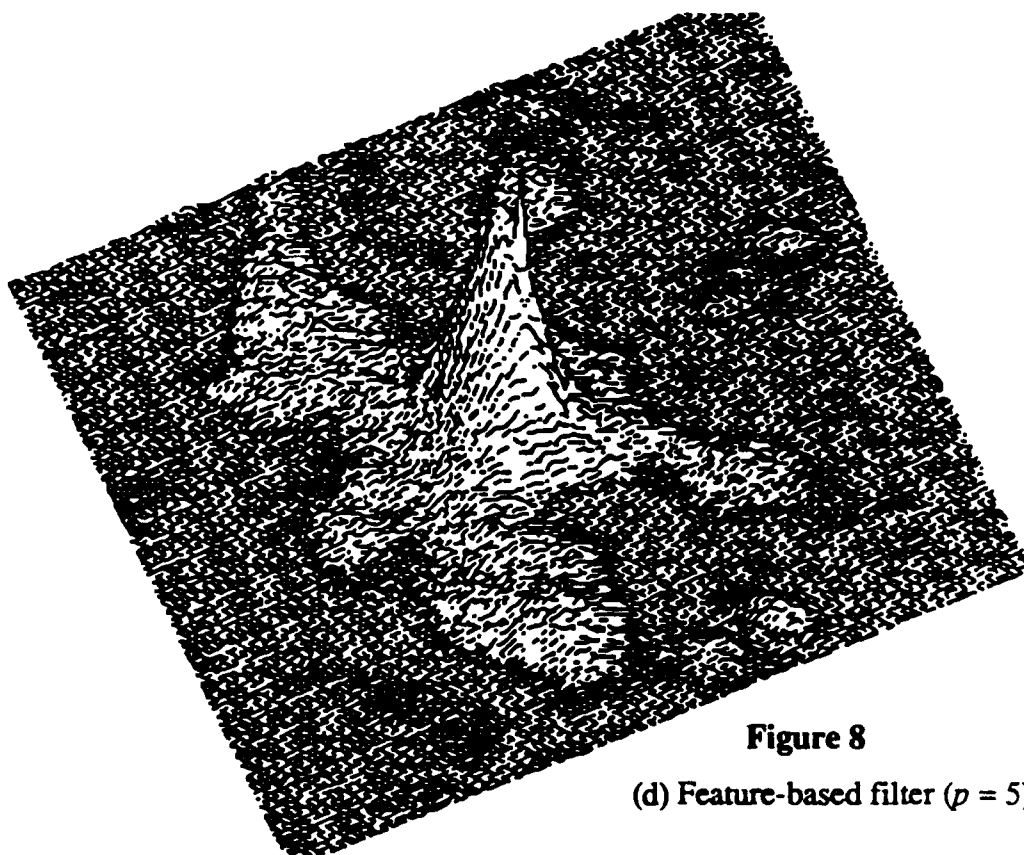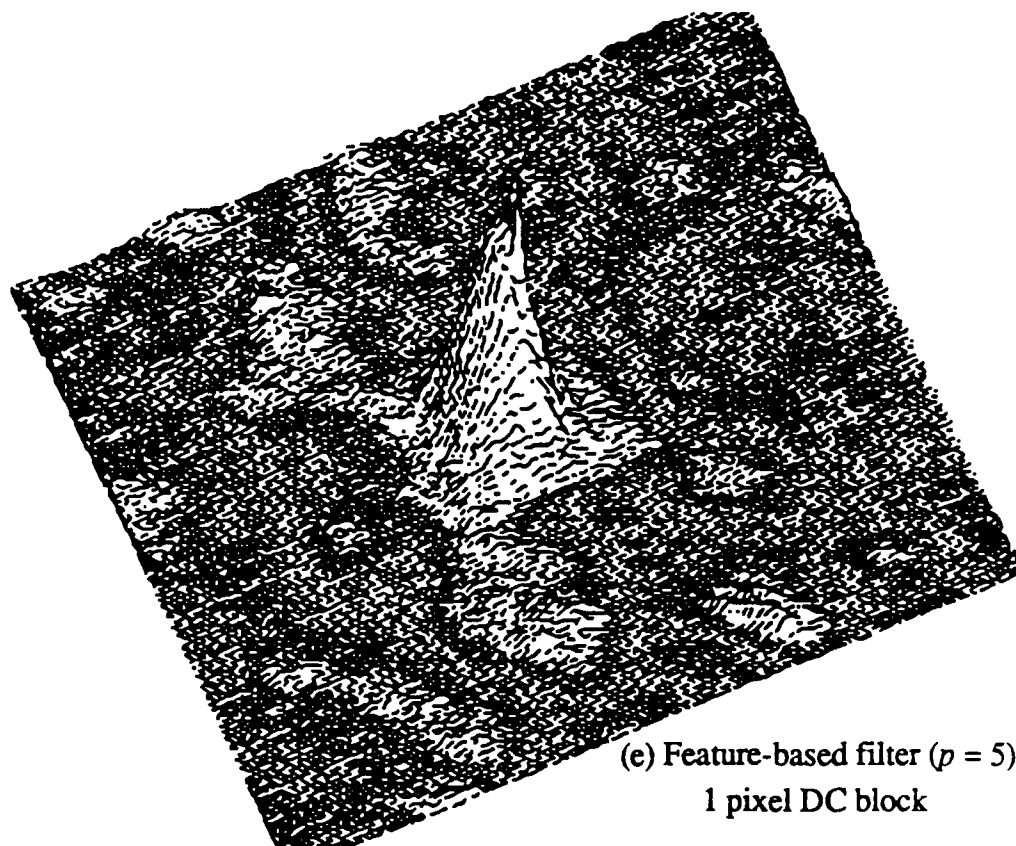
(c) Feature-based filter ($p = 3$)

Figure 8

(d) Feature-based filter ($p = 5$)

(e) Feature-based filter ($p = 5$)
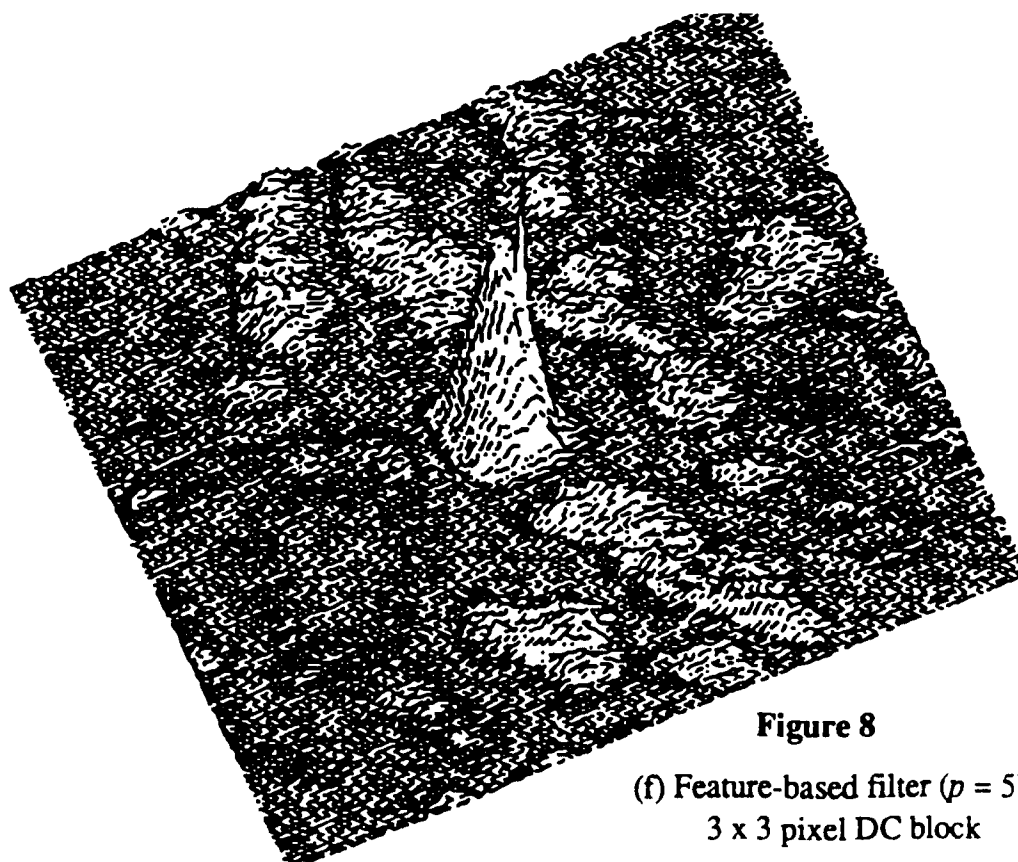1 pixel DC block

Figure 8

(f) Feature-based filter ($p = 5$)
3 x 3 pixel DC block

In contrast to the fSDF filter, the feature-based filters were potentially useful for images outside the training set. We arbitrarily set a threshold as a percentage of the maximum correlation height of the training set to examine the responses of the different filters. A correlation response greater than the threshold indicates recognition of an object. Although changing the threshold would change the results, similar conclusions can be drawn unless the threshold value is too low or too high. We examined the case where the threshold is 30% of the maximum and examined the data in Figs. 5, 6, and 7, for correlation results outside the training set. The data was summarized in Table 2 and showed that the performance of the fSDF filter and the filters with $p = 1, 3$, were similar. These filters were generally not useful for recognizing images outside the training set; a total of 1 of the 14 images outside the training set was recognized for these filters. When p was increased to 5, 12 of the 14 images were recognized. As pixels near the DC were blocked to improve the PRM-SR value obtained with the $p = 5$ filter, the average correlation heights of the responses decreased. Therefore, fewer images could be recognized. The performance of the filters with $p = 5$ with and without the central pixel blocked were similar.

Table 2: Performance of Various Filters for Images outside of Training Set

Filter	Number of images recognized (above 30% of maximum)	Average Normalized Correlation Height of recognized images
fSDF	0	-
Feature-based (p=1)	0	-
Feature-based (p=3)	1	34.6
Feature-based (p=5)	12	47.6
Feature-based (p=5) 1 pixel DC block	12	45.4
Feature-based (p=5) 3 x 3 pixel DC block	7	42.2

6.0 Discussion

The feature-based filter recognized a set of objects based on features. This was done by creating a filter that retained features that were invariant with respect to the training set. The feature domain contained the sign of the phase of the Fourier transform of training images at each pixel. By using features that were invariant to the training set, objects that were outside the training set could be recognized. When features offering conflicting information were used, objects outside the training set were not necessarily recognized because the features were not entirely invariant outside the training set. As more consistent features were used, the correlation heights of objects outside the training set increased.

The feature-based filter was calculated in a simpler manner than the fSDF filter. Both the time and number of operations required to calculate the feature-based filter was on the order of calculating n FFTs. In contrast, the fSDF approach requires a cross-correlation between every training image and the filter every iteration. This requires on the order of ni FFT calculations where i is the number of iterations and has been generally set to ten.[24] Therefore, the time and number of operations required to calculate the feature-based filter was about an order of magnitude less than for the fSDF filter.

A feature-based filter was more easily trained than an fSDF filter. We consider the case where an fSDF filter made from n training images is presented with a new image that belong to the same class as the training set but does not produce a useful correlation result. In this case, the fSDF filter must be completely recalculated in the usual way with $n + 1$ images in the training set. For k new images, $(n+k)i$ FFTs are required. In contrast, the feature-based filter required k FFTs when k new images were presented.

The proper features must have existed in the Fourier domain to provide recognition using the feature-based filter. It has been shown that a proper choice of support function can improve correlation results;[22] therefore, a better choice of features may provide improved correlation results.

7.0 Conclusion

The feature-based filter offered a range of performance. In the case where none of the pixels were set to zero in the filter, the fSDF and feature-based filter offered similar performance. The feature-based filter was slightly more consistent and had broader correlation peaks for objects within the training set than the fSDF filter. Neither filter appeared to be useful for recognizing objects outside the training set.

As pixels of the filter were set to zero in the feature-based filter, the correlation peaks within the training set became more consistent even though their average height decreased. As the number of pixels set to zero increased, the correlation heights became more consistent but broader. When images of the same class as the training set but not in the training set were used as inputs, the feature-based filter was potentially useful. Our experiments involved five training images. The use of more training images suggests that more possibilities are available in trading off between consistency and broadness of the correlation results. In this way, the feature-based filter can be made robust to recognized object outside the training set.

Finally, a feature-based filter was more easily calculated and trained than an fSDF filter. The feature-based filter required about an order of magnitude fewer calculation than the fSDF filter. In addition, new images added to the training set, required the fSDF filter to be completely recalculated which was not the case in the feature-based filter.

8.0 References

[1] M. W. Roth, "Survey of neural network technology for automatic target recognition," *IEEE Trans. Neur. Net.* 1(1) 28-43 (1990)

[2] M. A. Flavin and J. L. Horner, "Correlation experiments with a binary phase-only filter implemented on a quartz substrate," *Opt. Eng.* 28(5), 470-473 (1989)

[3] J. L. Horner and H. O. Bartlett, "Two-bit correlation," *Appl. Opt.* 24(18), 2889-2897 (1985)

[4] K. H. Fielding and J. L. Horner, "Clutter effects on the adaptive joint transform correlator," *Opt. Eng.* 31(3) 602-605 (1992)

[5] E. Ochoa, G. F. Schills, and D. W. Sweeny, "Detection of multiple views of an object in the presence of clutter," *Opt. Eng.* 27(4), 266-273 (1988)

[6] D. Casasent, "Unified synthetic discriminant function computational formulation," *Appl. Opt.* 23(10), 1620-1627 (1984)

[7] G. F. Schills and D. W. Sweeny, "Rotational invariant correlation filtering," *J. Opt. Soc. Amer. A* 2(9), 1411-1418 (1985)

[8] G. F. Schills and D. W. Sweeny, "Iterative technique for the synthesis of optical-correlation filters," *J. Opt. Soc. Amer. A* 3(9), 1433-1442 (1986)

[9] T. Szoplik and H. H. Arsenault, "Shift and scale-invariant anamorphic Fourier correlator using multiple circular harmonic filters," *Appl. Opt.* 24(19), 3179-3183 (1985)

[10] B. Javidi, S. F. Odeh, and Y. F. Chen, "Rotation and scale sensitivities of the binary phase-only filter," *Opt. Comm.* 65(4), 233-238 (1988)

[11] D. Casasent and W. A. Rozzi, "Computer-generated and phase-only synthetic discriminant function filters," *Appl. Opt.* 25(20), 3767-3772 (1986)

[12] D. L. Flannery, J. S. Loomis, M. E. Milkovich, and P. E. Keller, "Application of binary phase-only correlation to machine vision," *Opt. Eng.* 27(4), 309-320 (1988)

[13] R. R. Kallman, "Direct construction of phase-only filters," *Appl. Opt.* 26(24), 5200-5201 (1987)

[14] T. S. Wilkinson, D. A. Pender, and J. W. Goodman, "Use of synthetic discriminant functions for handwritten-signature verification," *Appl. Opt.* 30(23), 3345-3353 (1991)

[15] S. P. Kozaitis and W. E. Foor, "Performance of synthetic discriminant functions for binary phase-only filtering of thresholded imagery," *Opt. Eng.* 31(4) (1992)

[16] J. F. Haddon, "Generalised threshold selection for edge detection," *Pattern Recognition* 21(3),

195-203 (1988)

[17] M. M. Trivedi, C. A. Anderson, R. W. Conners, and S. Goh, "Object detection based on gray level cooccurrence," *Comp. Vision, Graphics & Im. Proc.*, 28(2), 199-219 (1984)

[18] R. Kohler, "A segmentation system based on thresholding," *Comp. Vision, Graphics & Im. Proc.*, 15(4), 319-338 (1981)

[19] D. M. Cottrell, R. A. Lilley, J. A. Davis, and T. Day, "Optical correlator performance of binary phase-only filters using Fourier and Hartley transforms," *Appl. Opt.* 26, 3755-3761 (1987)

[20] M. W. Farn and J. W. Goodman, "Optimal binary phase-only matched filters," *Appl. Opt.*, 27 4431-4437 (1988)

[21] B. V. K. Vijaya Kumar and Z. Bahri, "Phase-only filters with improved signal to noise ratio," *Appl. Opt.* 28(2) 250-257 (1989)

[22] B. V. K. Vijaya Kumar and Z. Bahri, "Efficient algorithm for designing a ternary valued filter yielding maximum signal to noise ratio," *Appl. Opt.* 28(10) 1919-1925 (1989)

[23] S. P. Kozaitis, S. Halby, and W. Foor, "Experimental performance of a binary phase-only optical correlator using visual and infrared imagery," *Optical Information Processing Systems and Architectures*, B. Javidi, Ed., Proc. SPIE 1296, 140-151 (1990)

[24] D. A. Jared and D. J. Ennis, "Inclusion of filter modulation in synthetic discriminant function construction," *Appl. Opt.* 28(2), 232-239 (1989)

[25] Z. Bahri and B. V. K. Vijaya Kumar, "Generalized synthetic discriminant functions," *J. Opt. Soc. Amer. A* 5(4), 562-571 (1988)

[26] B. V. K. Kumar, and L. Hasselbrook, "Performance measures for correlation filters," Appl. Opt. 29(20), 2997-3006 (1990)

[27] P. Maragos, "Morphological correlation and mean absolute error criteria," *Proc. of ICASSP* '89, IEEE, 1568-1571 (1989)

[28] D. E. Glover, An optical Fourier/electronic neurocomputer automated inspection system," *IEEE Int. Conf. Neural Networks*, vol.I, IEEE, 569-576 (1988)

[29] D. Clark and D. Casasent, "Practical optical Fourier analysis for high speed inspection," *Opt. Eng.* 27(5), 365-371 (1988)

MISSION

OF

ROME LABORATORY

Rome Laboratory plans and executes an interdisciplinary program in research, development, test, and technology transition in support of Air Force Command, Control, Communications and Intelligence (C^3I) activities for all Air Force platforms. It also executes selected acquisition programs in several areas of expertise. Technical and engineering support within areas of competence is provided to ESD Program Offices (POs) and other ESD elements to perform effective acquisition of C^3I systems. In addition, Rome Laboratory's technology supports other AFSC Product Divisions, the Air Force user community, and other DOD and non-DOD agencies. Rome Laboratory maintains technical competence and research programs in areas including, but not limited to, communications, command and control, battle management, intelligence information processing, computational sciences and software producibility, wide area surveillance/sensors, signal processing, solid state sciences, photonics, electromagnetic technology, superconductivity, and electronic reliability/maintainability and testability.