

Circular optimal trade-off and distance-classifier correlation filters

Samuel P. Kozaitis, MEMBER SPIE
Sila Thangwaritorn, MEMBER SPIE
Florida Institute of Technology
Division of Electrical, and Computer Science
& Engineering
150 W. University Blvd.
Melbourne, Florida 32901
E-mail: kozaitis@zach.fit.edu

Abstract. We use circular versions of advanced distortion-invariant filters such as optimal trade-off synthetic discriminant function and distance classifier correlation filters to obtain rotation invariance with an optical correlator. The filter noise performance is compared using a common measure of probability of error because the filters have different characteristics. The filters are real-valued so they can be implemented on a variety of SLMs. The circular symmetry of the filters significantly decreases their computational requirement. © 2000 Society of Photo-Optical Instrumentation Engineers. [S0091-3286(00)00605-X]

Subject terms: circular filter; correlator; distance-classifier correlation filter; optimal trade-off filter; rotation invariance.

Paper ATR-007 received Aug. 27, 1999; revised manuscript received Nov. 15, 1999; accepted for publication Dec. 8, 1999.

1 Introduction

There are several methods available to perform rotation-invariant object recognition using correlation filters.¹⁻¹¹ Recent advances have created correlation filters that can identify many rotated views of an object with one or a small number of filters. One popular approach has been the use of circular harmonic filters (CHF), which provide full rotation and shift invariance.^{1,2} This method is based on circular harmonic components of an input object, for which rotation is invariant. The first CHF used one circular harmonic component, while more advanced versions use linear combinations of circular harmonics of a training set of images, or minimize the average correlation energy. These later filters generally have reduced correlation plane side-lobes and improved discrimination when compared to earlier filters. Furthermore, phase-only and binary phase-only CHF filters have been used to yield sharp correlation peaks and improve discrimination.^{6,12,13} Other advances include a CHF filter designed to provide a specified response to rotations, implementation with a joint transform correlator, and a wavelet-based CHF filter.¹⁴⁻¹⁶ Using these approaches, rotation invariance can be achieved using circular harmonic expansions.

Rotation invariance can be obtained in a straightforward way by using rotated versions of an input image as training images to a generalized distortion-invariant filter. For example, using the original synthetic discriminant function (SDF) filter, rotated objects can be made to have the same correlation peak value.¹⁷ However, the SDF filter is limited because it constrains only one value in the correlation plane. The minimum average correlation energy (MACE) filter is a type of SDF filter that attempts to control the entire correlation plane.¹⁸ Generally, the MACE filter lacks noise tolerance and is sensitive to intraclass variations. The MACE filter is actually a special case of an optimal trade-off SDF (OTSDF) filter. In this filter, a performance measure is optimized while holding others constant and satisfying correlation peak constraints.^{19,20} To determine an

OTSDF filter for a particular problem, an energy function, which is the weighted sum of performance measures, is minimized. In this way, the OTSDF filter provides the best trade-off between the performance measures for the specified weights. In addition, some of these filters have been combined with CHFs to provide rotation invariance.^{4,9,11,16} To properly use OTSDF filters, prior knowledge of the recognized and rejected classes are required, which may not always be entirely available.

The use of distance classifier correlation filters (DCCFs) can also yield rotation invariance. DCCFs represent a different approach in that they optimally separate available target and reference image statistics while making them as compact as possible. The distortion tolerance of the approach requires only a similarity in features of objects rather than a precise match in pixel values. These filters are used with more than one class of objects and measure the similarity between shapes of the correlation plane and the ideal shape for that class.²¹⁻²⁵ In addition, they are capable of quadratic decision boundaries. The design of filters that contain many views of an object can be computationally intensive. Because pattern recognition problems could involve many classes, objects within classes, and full rotation invariance, the computation of a filter could be significant. A simpler method that could significantly decrease the computational load would enable more extensive development to occur.

Some disadvantages of SDF-based filters and DCCFs are related to their implementation. These filters are in general complex-valued, which limits the devices on which they can be implemented; however, modifications may be used to find filters that may be implemented on spatial light modulators (SLMs) with specific constraints. In addition, most development of advanced distortion-invariant filters have been in the frequency domain and are often implemented with a $4f$ correlator. The joint transform correlator (JTC) has some advantages when compared to the $4f$ correlator; however, the JTC cannot fully take advantage of

many filters because an amplitude or phase-only form description of a filter is often required.

As the angular spacing of the training images used in the SDF-based filters and DCCFs decreases to a small value, the resulting filter will appear circularly symmetric. In the limit, the filters will consist of concentric rings and can be referred to as circular filters. They can be thought of in polar coordinates with only the radius as a variable. By using circularly symmetric filters, rotation-invariant texture classification was achieved.^{26,27} In addition, a genetic algorithm was used to determine the radial magnitudes of a binary circular filter.²⁸ This approach used a performance criterion of a weighted sum of other performance criteria. The filters generated can be implemented on binary SLMs in the frequency domain as opposed to the generally complex values of the former filters.

We describe circular versions of OTSDF filters and DCCFs for implementation in the frequency or space domains, although they were designed in the frequency domain. Because the filters are circularly symmetric and the phase of the Fourier transform of a real function is an odd function, the filters are real-valued. In this way, they can be represented completely with SLMs with limited display characteristics. In addition, the spatial domain description is also real, which enables implementation on a JTC. The circularly symmetric nature of the filter simplifies its computation. Because the only variable is the value of the filter as a function of radius, the filter for an $n \times n$ image can be described by a vector of length p . In the next section, we briefly describe circular and OTSDF filters and DCCFs. Then, because some filters use different parameters in their design and sometimes different performance measures, we compare the filter performance in terms of an equivalent effect of probability of error. Because complete prior knowledge of the recognized and rejected classes may not always be entirely available, we test the filters with objects not in the training set, and in the presence of noise.

2 Circular Correlation Filters

When designing an SDF-based filter for rotation invariance, often a set of images at equally spaced angles are combined to form the filter. Sometimes, the number of training images N is kept small to reduce the computational requirement. The value of N is also significant because it affects the value of the output SNR (Refs. 29–31). As the value of N increases, usually better rotation invariance is achieved; the minimum output correlation peak of the training set is larger. However, with increasing N , the maximum SNR of the training set usually decreases and levels off. There may be many image classes, which increases the size of N and the computational demand of the filter.

Circular symmetric filter design attempts to eliminate variations in the output response of rotated input objects. In particular, only one image is required for an object to be rotation invariant, so the computational requirement is significantly reduced. The disadvantage of circular filters is that discrimination may suffer. Therefore, we designed circular symmetric filters using advanced filter designs specifically for detection and discrimination. In contrast to traditional approaches where multiple views of an object are used as a training set for the filter, we summed the Fourier

transform of an object across each angle for a particular radius as

$$\mathbf{x}_i = \sum_{\theta} X_i(r, \theta), \quad (1)$$

where $X_i(r, \theta)$ is the Fourier transform of the i 'th input object in polar coordinates, and \mathbf{x}_i is a vector that contains the sum of the Fourier transform as a function of radius r . Using the values of \mathbf{x}_i in a circular filter is equivalent to using a training set with a large value of N (close angular spacing). This approach requires only one image for each object, and the resulting filter values are contained in a vector as opposed to a matrix. Because the filter is circularly symmetric, the phase values cancel and the filter will be real-valued.

2.1 OTSDF Filters

We considered an example of an OTSDF filter that considers parameters of discrimination and noise tolerance as performance values. Additional parameters can be easily incorporated into the filter, but our example demonstrates the circular OTSDF method. In our example, the filter was designed to minimize the effect of both additive noise on the correlation output, and the average correlation energy. This is actually a trade-off between MACE and minimum variance SDF (MVSDf) filters.³² In the MVSDf filter the output correlations are generally broad, while the MACE filter usually has sharp correlation peaks. The circular OTSDF design can be described as

$$\mathbf{h} = \mathbf{T}^{-1} \mathbf{X} (\mathbf{X}^+ \mathbf{T}^{-1} \mathbf{X})^{-1} \mathbf{c}, \quad (2)$$

where \mathbf{X} is a matrix and contains the training set of circular filter values \mathbf{x}_i as columns, \mathbf{c} is a vector that contains the desired correlation peak of each training image, and the matrix \mathbf{T} contains the weighted sum of performance parameters. The vector \mathbf{h} contains the values of the circular OTSDF filter. In our example, $\mathbf{T} = \alpha \mathbf{P} + (1 - \alpha^2 \mathbf{D}_x)^{1/2}$, where \mathbf{D}_x is a diagonal matrix with that contains the average power spectrum of the circular filter values, and \mathbf{P} is the noise power spectral density which is assumed to be white. The parameter α varies from $0 \leq \alpha \leq 1$. If $\alpha = 0$ then the resulting filter is a MACE filter, if $\alpha = 1$ the result is a MVSDf filter.

2.2 DCCF Filters

The DCCF filter was developed to directly measure the similarity between the shape of the correlation plane and the ideal shape for that class.²¹ The DCCF filter was designed to optimally separate classes while making them as compact as possible. In other words, the filter increases the interclass distances while making each class more compact to simultaneously improve distortion tolerance and discrimination. For two classes, the circular DCCF filter can be described by

$$\mathbf{h} = \mathbf{S}^{-1} (\mathbf{m}_x - \mathbf{m}_y), \quad (3)$$

where \mathbf{m}_x and \mathbf{m}_y are the mean vectors of classes x and y , respectively, and \mathbf{S} is the diagonal matrix,

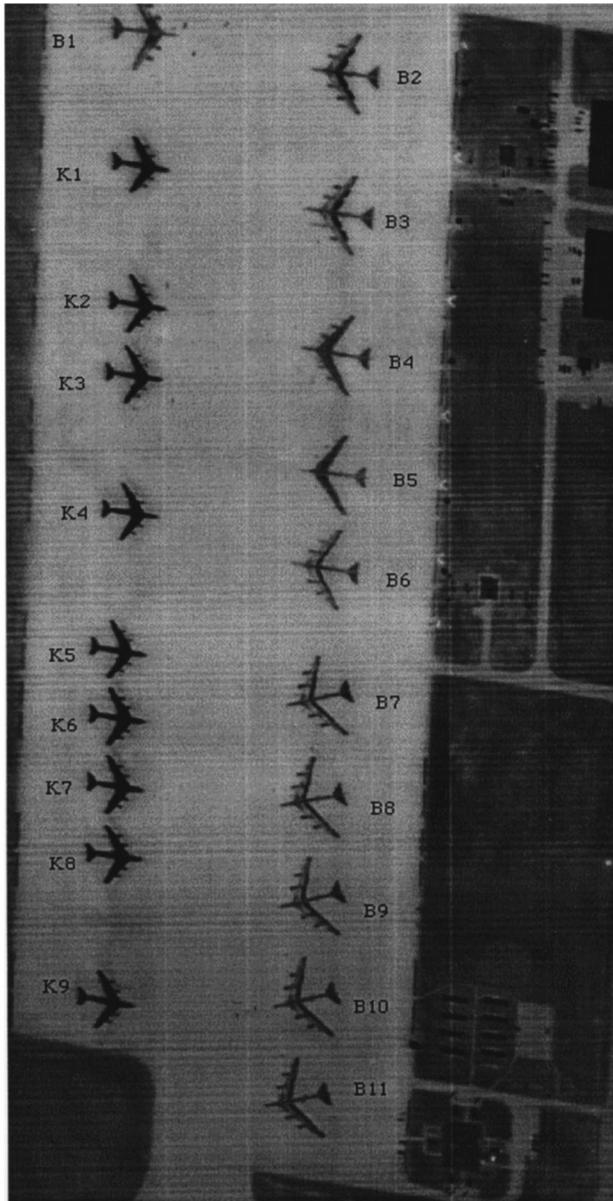


Fig. 1 IR image used in experiments.

$$\mathbf{S} = [1/N_x \sum (\mathbf{X}_i - \mathbf{M}_x)(\mathbf{X}_i - \mathbf{M}_x)^* + 1/N_y \sum (\mathbf{Y}_i - \mathbf{M}_y)(\mathbf{Y}_i - \mathbf{M}_y)^*], \quad (4)$$

where \mathbf{M}_x and \mathbf{M}_y are the diagonal matrix versions of \mathbf{m}_x and \mathbf{m}_y , and \mathbf{X}_i and \mathbf{Y}_i are the diagonal matrix versions of the i 'th image in classes x and y . The first summation in Eq. (4) is over class x , and the second summation is over class y , and N_x and N_y are the number of images used for each class.

3 Simulation Results

3.1 Recognition

We examined the performance of circular matched and OTSDF filters and DCCFs by comparing their results from the same sensor imagery. We used two classes, each with

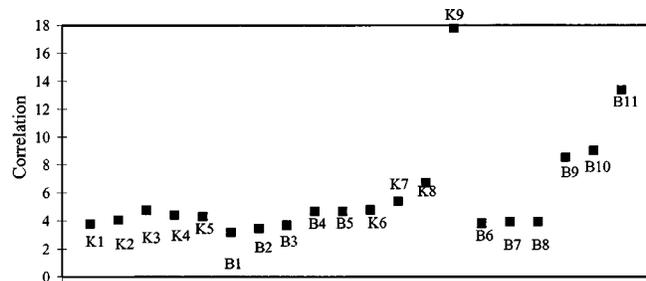


Fig. 2 Correlation peak values resulting from a matched filter with object K2 as the reference.

five training objects with each object taken from the IR image shown in Fig. 1. Each object image was 128×128 pixels and contained one object. The circular matched filter was derived from object K2, and the training sets for the OTSDF filter and the DCCF were from objects K1 to K5 and B1 to B5, for class x (in class) and class y (out of class), respectively. In the approach of Ref. 28, a genetic algorithm determined the radii of the circular filters while we considered circular filters with radii of 2 pixels.

We initially compared the correlation peak values using different filters as a measure of performance. With all filters we cross-correlated the input images at one rotation angle with the circular filter; input images at various rotation angles gave the same result due to the filter being circularly symmetric. The results for the matched filter are shown in Fig. 2. The data showed that most correlation peaks were similar to or greater than the peak associated with the object from which the filter was made. The results indicated that the matched filter provided no discrimination between the two classes of objects.

We show the peak correlation values of the OTSDF filter in Fig. 3 for $\alpha = 0.90$. The results show that the discrimination among the training set was satisfactory, indicated by the large separation between the two sets of correlation peak values. The discrimination of objects not in the training decreased even though the images visually appeared similar to the training set.

The correlation peak values for the DCCF filter are shown in Fig. 4. For the training set, the two classes were separated, but it did not appear to be as great as with the OTSDF filter. Although the discrimination of the remaining objects appeared to decrease, the discrimination appeared to be greater than that for the OTSDF filter. A useful performance measure associated with DCCF filters is the log-distance ratio (LDR). It is written as

$$\text{LDR} = \log(d_x/d_y), \quad (5)$$

where d_x and d_y are the transformed distances between the input object and the class centers. For in-class objects, the LDR should be negative, indicating that the distance to the center of class x is smaller than the distance to the center of class y . In addition, the LDR should be positive for out-of-class objects. The LDR results for the DCCF filter are shown in Fig. 5. The data show that class y was correctly classified in all cases; however, there were two classification errors with class x .

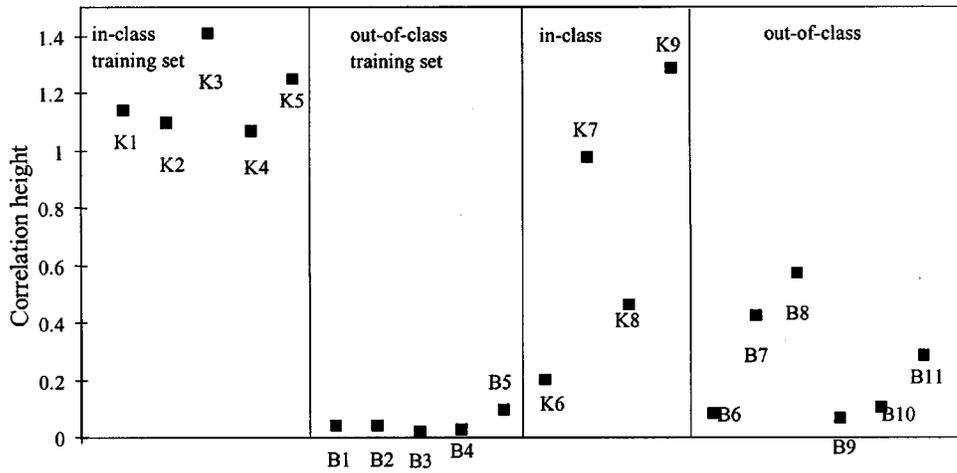


Fig. 3 Correlation peak values resulting from the OTSDF filter.

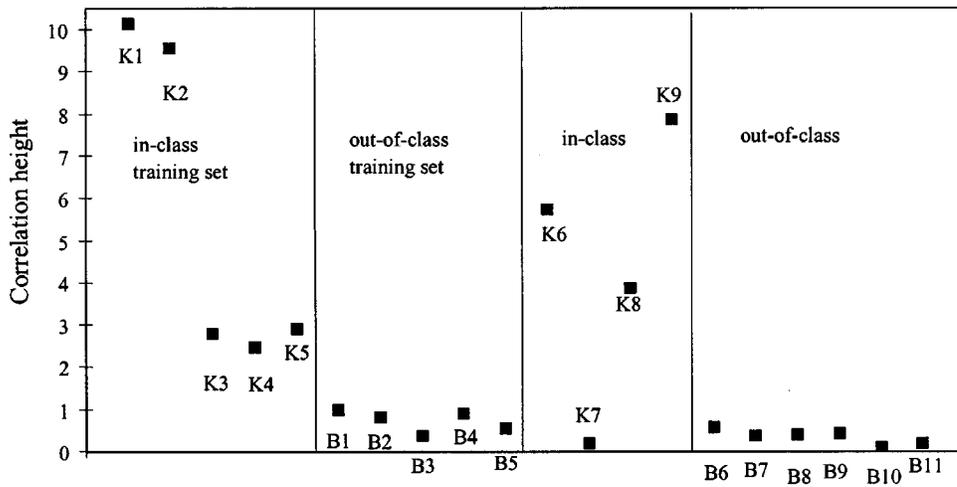


Fig. 4 Correlation peak values resulting from the DCCF.

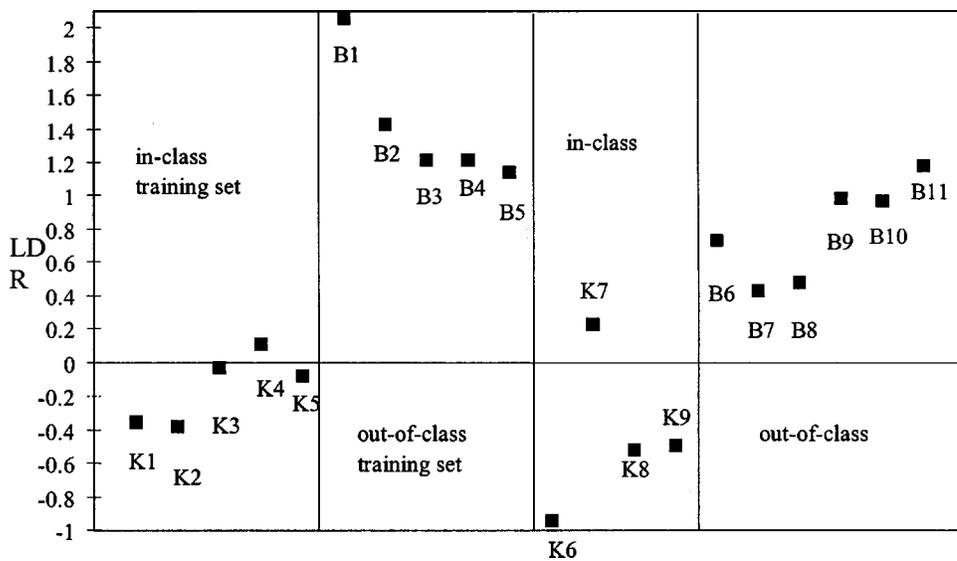


Fig. 5 LDR values resulting from the DCCF.

Table 1 PCE values for different filters.

Filter	PCE
Matched	0.89
OTSDF	1.84
DCCF	0.77

The peak correlation energy (PCE) values for the different filters were indicated in Table 1. The PCE for the matched filter was the value obtained with K2 as the input object. The PCE values for the OTSDF filter and the DCCF were obtained from the average of the in-class training sets, and were 1.84, and 0.77 for the OTSDF filter and the DCCF, respectively. The correlation plane responses for an input SNR of 10 when the object K2 was used as an input are shown in Figs. 6 and 7 for the OTSDF filter and the DCCF, respectively. Figure 6 shows narrow correlation peaks with large sidelobes, and Fig. 7 shows a much broader response without sidelobes.

3.2 Noise Performance

We compared the noise performance of the OTSDF filter and the DCCF as a function of the SNR of the input images using a methodology based on the human psychophysics literature.³³ In any detection task, there are some variables that affect the SNR in an image, which affects the performance of detection or discrimination. To measure noise performance we added noise to all input images. Then, two groups of correlation peak values were generated: target and nontarget values, as shown in Fig. 8. The two curves overlap each other and generate two areas when a threshold is chosen; one corresponds to error detection, where the input image is classified as not containing the target when it does. In contrast, there is a false alarm if the input image is classified as containing the target when in fact it does not. The probability of error detection $P(\text{error detection})$ is defined as the area under the curve indicated. Similarly for the probability of false alarm $P(\text{false alarm})$. For an equal probability of target and nontarget the probability of error $P(E)$, is defined as

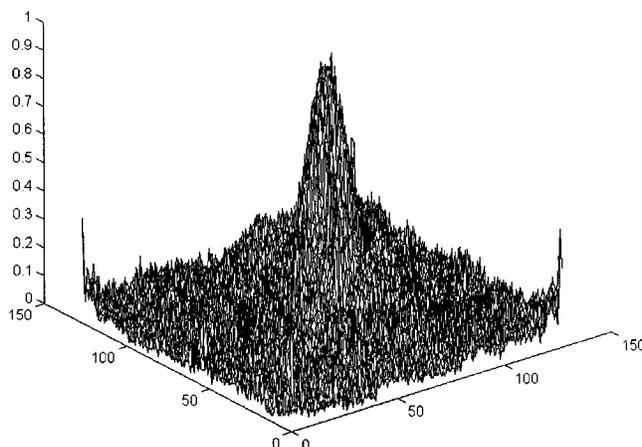


Fig. 7 Correlation response of object K2 and the DCCF.

$$P(E) = \frac{P(\text{error detection}) + P(\text{false alarm})}{2} \tag{6}$$

We compared the performance of both the OTSDF filter and the DCCF for two cases; one when the 10 training images were used as the test images and the other when the test images contained the training set and an additional 10 images from Fig. 1. We used SNRs from 1 to 10, and for each SNR value, we used 100 versions of each image containing independent noise samples. We then determined the histogram of the correlation peak values and found the value that produced the minimum value of the probability of error $P(E)$ for the two classes.

In the case of the OTSDF filter, Fig. 9 shows the minimum $P(E)$ for all sets of test images for a filter with $\alpha = 0.90$. When only the training images were used as the test images, there was a relatively large separation between the two sets of correlation values for all SNR values. The results indicate that a relatively large amount of noise was needed to produce errors.

In the case where all images were used as test images, the error due to the OTSDF filter increased when compared to when only the training set images were used. Even when

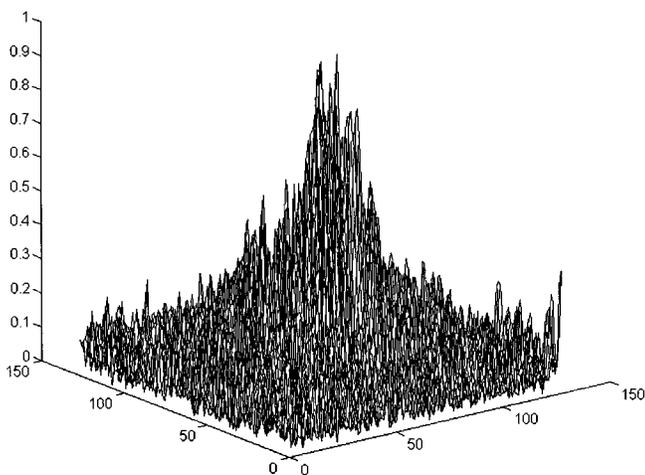


Fig. 6 Correlation response of object K2 and the OTSDF filter.

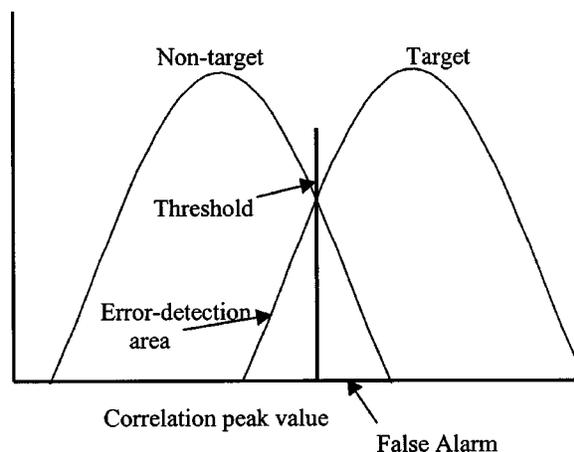


Fig. 8 Histogram of target (in-class) and nontarget (out-of-class) peak correlation values.

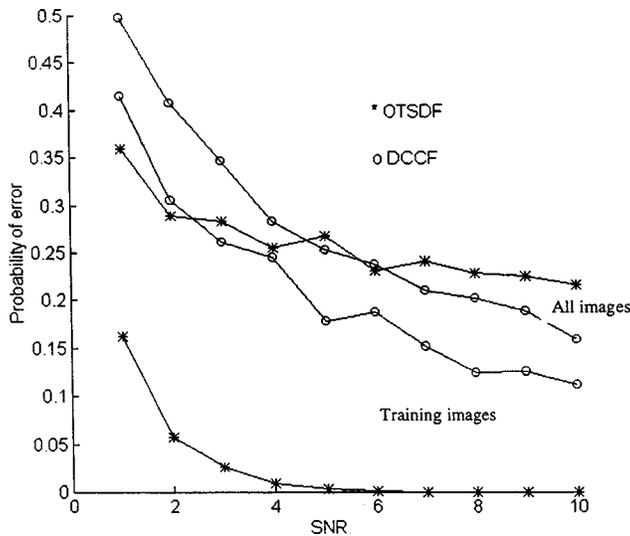


Fig. 9 Minimum probability of error using correlation peak values for both the DCCF and the OTSDF filter.

noise was not present, some error occurred. Although the threshold for the two cases were similar, the separation of the correlation peak values of the images not in the training set was smaller than for the training set images; therefore, the $P(E)$ increased.

We also showed the results of experiments using correlation peak values using the DCCF with the same two sets of input images. In the case where only the training set images were used to test the filter, the $P(E)$ was higher than for the OTSDF filter for all values of input SNR. When all images were used as test images the $P(E)$ of the DCCF was lower than that of the OTSDF filter at large values of SNR, and higher than that of the OTSDF filter for lower SNR values. In addition, the $P(E)$ did not increase as much as with the OTSDF filter when all images were used when compared to when only the training set was used to test the filter.

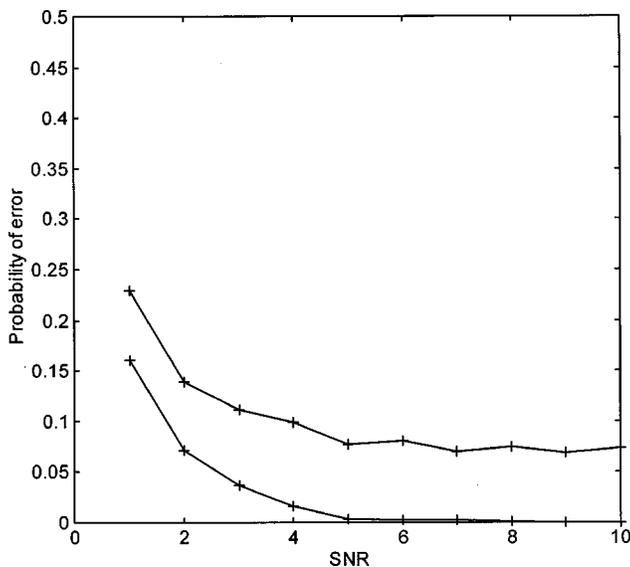


Fig. 10 Probability of error using LDR values for the DCCF.

We also used the LDR values of the DCCFs to determine the probability of error. We calculated the LDR values instead of correlation peak values and chose a threshold value of zero. Note that the zero threshold indicates to which class an object is most similar, but does not necessarily indicate the optimum $P(E)$. Using the LDR as the performance measure, the $P(E)$ generally decreased, as shown in Fig. 10. The results using the training set as the test set gave results similar to that of the OTSDF filter using the same training set and correlation peak values as a measure of performance. In addition, the $P(E)$ remained lower than previous results when the additional images were added to the test set.

4 Conclusion

Advanced distortion-invariant filters such as the OTSDF filter and the DCCF can be used as circular filters to obtain rotation invariance. The filters' performance was compared using a probability of error so that filters having different probability of error could be compared with a single measure. Using correlation peak values and a test set consisting of only the training set used to make the filter, the noise performance of the OTSDF filter was superior to that of the DCCF. This is not surprising as the OTSDF filter is supposed to specifically set the correlation peak value. When additional images were used to test the performance of both filters, the results were similar. Therefore the OTSDF filter was more sensitive to images not in the training set than the DCCF. Using the LDR value as a performance measure the DCCF gave results similar to those of the OTSDF filter using correlation peak values, when only the training set was considered for both filters. When images not in the training set were used to test the DCCF, its probability of error was lower than that of the OTSDF filter. Finally, the filters were real-valued so they can be implemented on a variety of SLMs. The circular symmetry of the filters significantly decreased the computational requirement of the filters.

References

1. Y.-N. Hsu, H. H. Arsenault, and G. April, "Rotation invariant digital pattern recognition using circular harmonic expansion," *Appl. Opt.* **21**, 4012–4015 (1982).
2. Y.-N. Hsu and H. H. Arsenault, "Optical pattern recognition using circular harmonic expansion," *Appl. Opt.* **23**, 841–844 (1984).
3. D. Asselin and H. H. Arsenault, "Rotation and scale invariance with polar and log-polar coordinate transforms," *Opt. Commun.* **104**, 391–404 (1994).
4. D. Casasent, A. Iyer, and G. Ravichandran, "Circular harmonic function MACE filters," *Appl. Opt.* **30**, 5169–5175 (1991).
5. G. Ravichandran and D. Casasent, "Advanced in-plane rotation-invariant correlation filters," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-16**, 415–420 (1994).
6. J. Rosen and J. Shamir, "Circular harmonic phase filters for efficient rotation-invariant pattern recognition," *Appl. Opt.* **27**, 2895–2899 (1988).
7. R. Wu and H. Stark, "Rotation-invariant pattern recognition using optimum feature extraction," *Appl. Opt.* **24**, 179–184 (1985).
8. Y. Sheng and H. H. Arsenault, "Object detection from a real scene using the correlation peak coordinates of multiple circular harmonic filters," *Appl. Opt.* **28**, 245–249 (1989).
9. J. Garcia, J. Campos, and C. Ferreira, "Circular-harmonic minimum average correlation energy filter for color pattern recognition," *Appl. Opt.* **33**, 2180–2187 (1994).
10. V. Zavala-Hamz and J. Alvarez-Borrego, "Circular harmonic filters for the recognition of marine microorganisms," *Appl. Opt.* **36**, 484–489 (1997).
11. D. Casasent and W. Cox, "RI-MINACE filters to augment segmentation of touching objects," *Pattern Recogn.* **31**, 1311–1317 (1998).
12. L. Leclerc, Y. Sheng, and H. H. Arsenault, "Rotation invariant

- phase-only and binary phase-only correlation," *Appl. Opt.* **28**, 1251–1256 (1989).
13. X. Chen and Z. Chen, "Amplitude-modulated circular-harmonic filter for pattern recognition," *Appl. Opt.* **34**, 879–885 (1995).
 14. B. V. K. Vijaya Kumar and T. K. Ng, "Multiple circular-harmonic-function correlation filter providing specified response to in-plane rotation," *Appl. Opt.* **35**, 1871–1878 (1996).
 15. S. Chang, S. Boothroyd, P. Palacharia, and S. Paparao, "Rotation-invariant pattern recognition using a joint transform correlator," *Opt. Commun.* **127**, 107–116 (1996).
 16. S. H. Lee, J. W. Kim, H. W. Lee, D. S. Noh, and S. J. Kim, "Wavelet CHF-SDF filter for distortion-invariant pattern recognition," *Proc. SPIE* **2898**, 238–248 (1996).
 17. C. F. Hester and D. Casasent, "Multivariate techniques for multiclass pattern recognition," *Appl. Opt.* **19**, 1758–1761 (1980).
 18. A. Mahalanobis, B. V. K. Vijaya Kumar, and D. Casasent, "Minimum average correlation energy filters," *Appl. Opt.* **26**, 3633–3640 (1987).
 19. P. Refregier, "Filter design for optical pattern recognition: multi-criteria optimization approach," *Opt. Lett.* **15**, 854–856 (1990).
 20. V. Laude and P. Refregier, "Multi-criteria characterization of coding domains with optimal Fourier spatial light modulator filters," *Appl. Opt.* **33**, 4465–4471 (1994).
 21. A. Mahalanobis, D. W. Carlson, B. V. K. Vijaya Kumar, and S. R. F. Sims, "Distance classifier correlation filters," *Proc. SPIE* **2238**, 2–13 (1994).
 22. B. V. K. Vijaya Kumar and A. Mahalanobis, "Recent advances in distortion-invariant correlation filter design," *Proc. SPIE* **2490**, 2–13 (1995).
 23. A. Mahalanobis, A. V. Forman, N. Day, M. Bower, and R. Cherry, "Multi-class SAR ATR using shift-invariant correlation filters," *Pattern Recogn.* **27**, 619–626 (1994).
 24. A. Mahalanobis, B. V. K. Vijaya Kumar, and S. R. F. Sims, "Distance-classifier correlation filters for multiclass target recognition," *Appl. Opt.* **35**, 3127–3133 (1996).
 25. A. Mahalanobis, L. A. Ortiz, and B. V. K. Vijaya Kumar, "Performance of the MACH filter and DCCF algorithms on the 10-class public release MSTAR data set," *Proc. SPIE* **3721**, 285–291 (1999).
 26. R. Porter and N. Canagarajah, "Gabor filters for rotation invariant texture classification," in *Proc. IEEE Int. Symp. Cir. Syst.* pp. 1193–1196 (1997).
 27. R. Porter and N. Canagarajah, "Robust rotation-invariant texture classification: wavelet, Gabor filter and GMRF based schemes," *Proc. IEEE: Vis. Image Signal Process.* **144**, 180–188 (1997).
 28. L. Singher, O. K. Ersoy, and G. E. Miles, "Optimization of binary circular filters," *Opt. Eng.* **36**(3), 922–934 (1997).
 29. H. Mostafavi and F. Smith, "Image correlation with geometric distortion—part I: acquisition performance," *IEEE Trans. Aerosp. Electron. Syst.* **AES-14**, 487–493 (1978).
 30. B. V. K. Vijaya Kumar and E. Pochapsky, "Signal-to-noise ratio considerations in modified matched spatial filters," *J. Opt. Soc. Am. A* **3**, 1579–1584 (1986).
 31. B. V. K. Vijaya Kumar, "Tutorial of composite filter designs for optical correlators," *Appl. Opt.* **31**, 4773–4801 (1992).
 32. B. V. K. Vijaya Kumar, "Minimum variance synthetic discriminant functions," *J. Opt. Soc. Am. A* **3**, 1579–1584 (1986).
 33. T. Kanungo, M. Y. Jaismha, J. Palmer, and R. M. Haralick, "A methodology for quantitative performance evaluation of detection algorithms," *IEEE Trans. Image Process.* **4**(12), 1667–1673 (1995).



Samuel P. Kozaitis received his PhD in electrical engineering in 1986 from Wayne State University, where he was an assistant research professor. He has also been with General Motors Research Laboratories, was a research fellow at the Photonics Center of the Rome Laboratory from 1988 to 1997, spent two summers at the National Aeronautics and Space Administration Kennedy Space Center, and has consulted for government and industry. He is currently an associate professor in the Division of Electrical, Computer Science, and Engineering at the Florida Institute of Technology. His major research interests include optical pattern recognition, wavelet transforms, image processing, and higher order statistics.

Sila Thangwaritorn received his BS in physics from Ramkhamhaeng University and his MS in electrical engineering from King-mongkut's Institute of Technology, Thailand, in 1985 and 1989. He also received MS and PhD degrees in electrical engineering from the Florida Institute of Technology in 1993 and 1998. He was a research assistant involved with fiber optic sensor, communications, and optical pattern recognition at Florida Institute of Technology from 1992 to 1998. He is currently a system engineer with Intelligent Machine Concepts. His research interests include image processing, pattern recognition, and fiber optics.