AD-A235 553

RADC-TR-90-108
In-House Report
April 1990

FOURIER OPTICAL SIGNAL PROCESSORS

Dr. Joseph L. Horner and David L. Flannery

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

DTIC
ELECTE
MAY 14 1991
S B D

Rome Air Development Center
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700
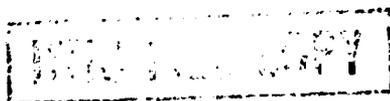
91 5 14 010

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

RADC-TR-90-108 has been reviewed and is approved for publication.

APPROVED:
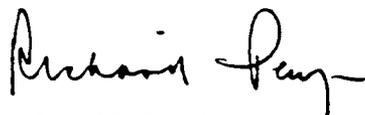
RICHARD PAYNE
Chief, Electro-Optic Device Tech Division
Directorate of Solid State Sciences

APPROVED:
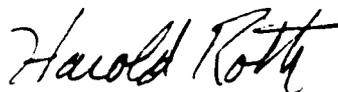
HAROLD ROTH
Director of Solid State Sciences

FOR THE COMMANDER:
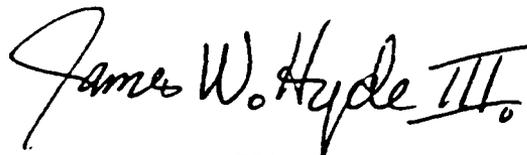
JAMES W. HYDE III
Directorate of Plans & Programs

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

| 1. AGENCY USE ONLY (Leave blank) | 2. REPORT DATE
1990 April | 3. REPORT TYPE AND DATES COVERED
In-House Oct 87 - Sep 89 |
|---|---|---|

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Fourier Optical Signal Processors	PE 61102F
PR   2305
TA   J7
WU   11

6. AUTHOR(S)

Horner, Joseph L.; Flannery, David L.*

| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Rome Air Development Center (ESOP)
Hanscom AFB
Massachusetts 01731-5000 | 8. PERFORMING ORGANIZATION
REPORT NUMBER

RADC-TR-90-108 |
|---|---|

| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | 10. SPONSORING/MONITORING
AGENCY REPORT NUMBER |
|---|---|
| | |

11. SUPPLEMENTARY NOTES

*University of Dayton Research Institute, Dayton, OH 45469

| 12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited	12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Progress in Fourier optical processing techniques is reviewed with particular emphasis on real-time pattern recognition, which recently has received increased interest as the result of new filter formulations that can be implemented with existing spatial light modulators. Architectures for coherent optical correlation, and the recently reported phase-only filters are reviewed. "Smart filters" which attack the inherent distortion sensitivity of correlation, are reviewed including those suitable for implementation with phase-only modulation. Correlation experiments implementing real-time variation of both input and reference patterns are reviewed. The potential for successful near-term application of these techniques is defined.

| 14. SUBJECT TERMS

Pattern recognition correlation, holography, spatial filters | 15. NUMBER OF PAGES
44 |
|---|---|
| | 16. PRICE CODE |

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

Contents

Illustrations

Fourier Optical Signal Processors

I. INTRODUCTION

Fourier optical signal processing is certainly the oldest form of optical signal processing with roots going back to the last century with the works of Abbe in 1893. By Fourier optical signal processing, we mean the intentional modification of the spectrum of an image to modify that image or some process based on that image (such as image restoration) or recognition of that image by an autonomous system. It is this latter application that we will deal with at some length in this article. There are two basic reasons for that. Pattern recognition has not been "solved" by any stretch of the imagination, either by digital or optical processing. However, with new algorithms and devices we think a hybrid, optical/digital system can make a significant advance in this area, in the form of a system that is ready for transitioning to the work place.

The basic premise of the Fourier optical processor is the ability of a lens to extract a two-dimensional Fourier transform of an input signal, as shown in Figure 1. Note that the input illumination must be highly coherent light, as from a laser, and that the computation is done fully in parallel, literally at the speed of light, about 1 ft/ns. This is equivalent to a throughput of about $4 \cdot 10^{15}$ mathematical operations/sec for a 256 x 256 sample data array in a Fourier optical system approximately 1 ft long, using a 20 mW laser diode light source. This estimate is based on the number of digital multiplications and additions required to perform the same transform using an FFT algorithm. It is this vast "number crunching" potential of optical processing that has stimulated and sustained research and development in this area since the advent of the laser as a practical source of such a system. Up until the past few years, the bottleneck in using this concept for anything practical has been getting the data into and out of the system. Fortunately, recent developments in hardware and software (filter algorithms), have mitigated this problem, and spurred a flurry of research and development activity in this field.

1
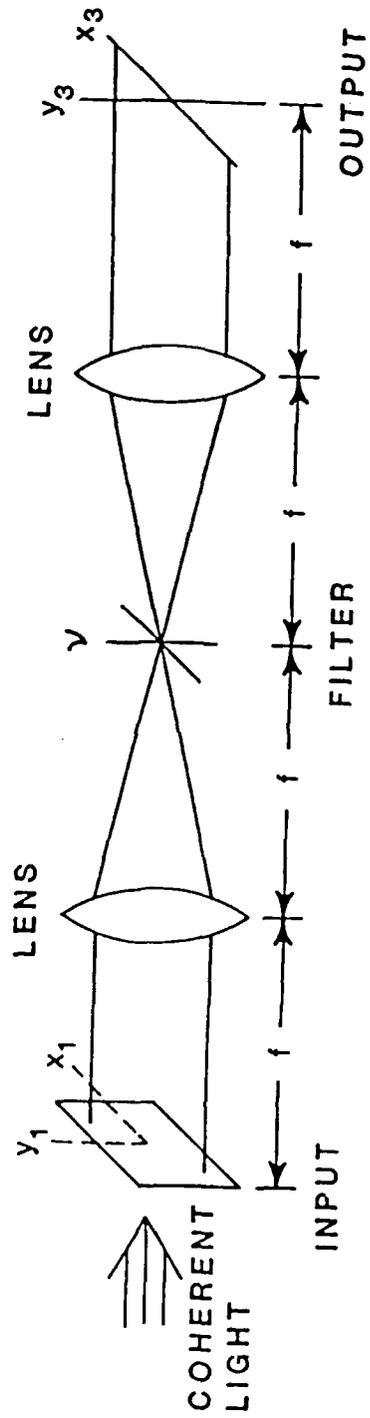
Figure 1. Frequency plane optical correlator. Reference function is inserted in filter plane. Correlation between input signal at plane (x1, y1) and reference is produced at plane (x3, y3).

2

2. BACKGROUND

In 1964, Vander Lugt proposed using a Fourier plane mask for pattern recognition,[1] and this architecture is usually referred to as a frequency plane correlator. His system performed a correlation between two functions, f1 and f2, and is based on the autocorrelation theorem and the Fourier transform property of a lens and coherent light:

$$C(x) = FT^{-1} \{ F_1(v) \cdot F_2^*(v) \}, \tag{1}$$

where $C(x)$ is the correlation function, FT stands for the Fourier transformation, $F1$ and $F2$ are the Fourier transforms of $f1$ and $f2$ respectively, and the $*$ denotes complex conjugation. We will use the convention of letting "x" represent $x, y,$ "v" for $v_x, v_y,$ etc. This is implemented by the Fourier transform system of Figure 1. The function F_2^* is equivalent to a matched filter and is a complex function expressible in terms of magnitude and phase:

$$F_2 = |F_2| \cdot e^{i\phi}. \tag{2}$$

Vander Lugt's insight was realizing a way to write this complex function on an energy-sensitive medium, photographic film. The method proposed was holography, which results in encoding the complex function on a spatial carrier, and involves an object beam modulated by the reference pattern. With the introduction of computer-generated holography,[2, 3] (CGH), it became possible to make matched filters of mathematically synthesized objects, but the CGH filter still was written on photographic film, and therefore could not be written in real time.

Another optical architecture, which has become known as the joint transform correlator (JTC), was proposed several years later by Weaver and Goodman,[4] and is shown in Figure 2. Here the reference function and unknown object are presented together in the input plane, and their combined or joint Fourier transform is produced in the focal plane behind the first lens. If the joint Fourier transform is recorded on a square-law or energy detector, originally photographic film, and the inverse Fourier transform is taken, we see that again we have implemented Eq. (1) and have produced a matched filter correlation between the two signals. One advantage of this system is that one does not need a matched filter; one simply introduces the reference function, usually a real function, beside the input pattern at the input plane. This contrasts with frequency plane correlation which requires producing a Fourier transform of the reference function, and encoding and recording this complex function as a physical matched filter to implement modulation in the Fourier plane. A disadvantage of the joint transform correlator is that the space-bandwidth product available in the input plane must be shared by the input and reference signals, as well as a guard band to ensure separation of correlation patterns from undesired terms at the output plane. This architecture recently has attracted much interest as researchers have replaced the photographic film in the Fourier plane with CCD detectors, and the input film with a miniature liquid crystal TV (LCTV) display,[5, 6] thus yielding real-time operation. In addition, an architecture has been proposed which, by using a time-sequencing scheme, uses only a single spatial light modulator (SLM) for both the input and Fourier planes.[7, 8] This greatly reduces the size and cost of this correlation system.

These then are the basic Fourier-optical pattern recognition architectures we have to work with — based on the simple principle that a lens focusing a beam of coherent (laser) light modulated with an image produces a 2D optical analog Fourier transform of that image. Recent developments in hardware (SLMs), and new filter algorithms particularly amenable to optical implementation (phase-only filters), have combined to create new

3
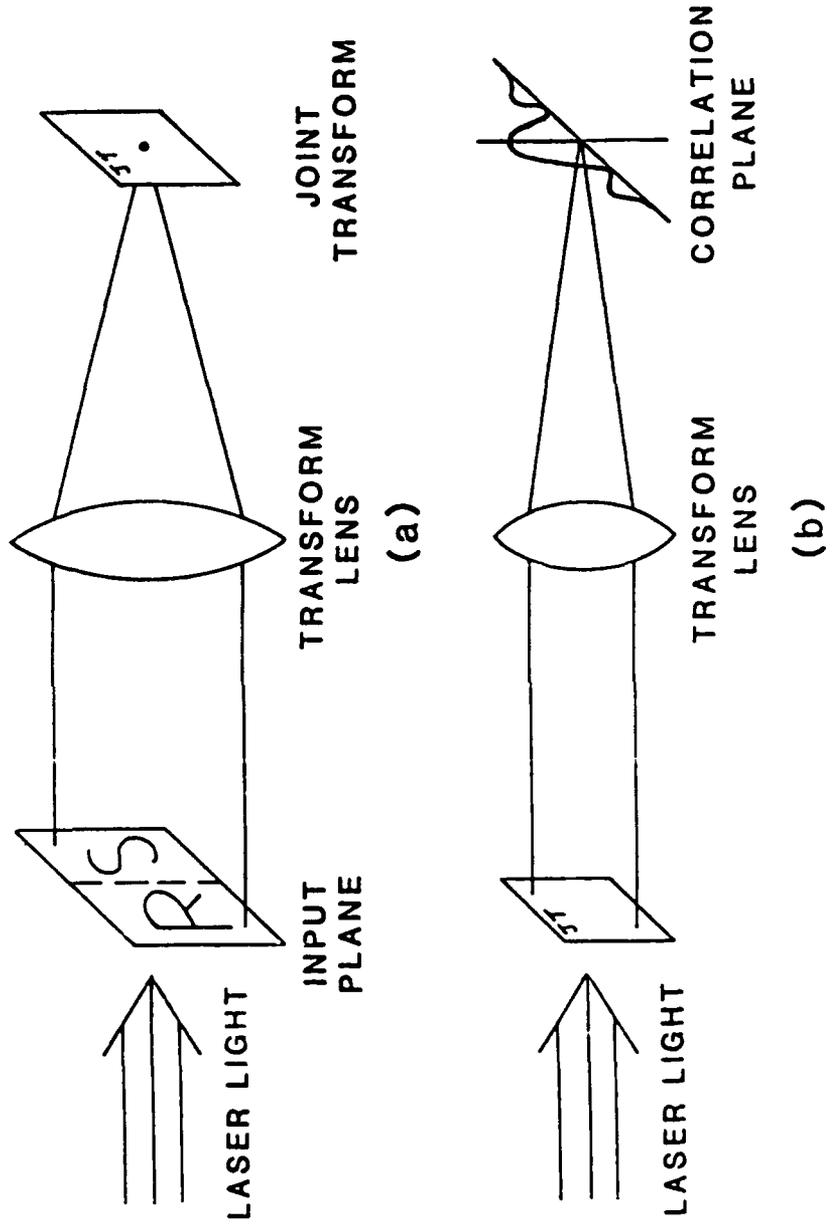
Figure 2. Joint transform correlator. In (a), the joint Fourier spectrum is recorded. In (b), the joint spectrum is re-illuminated and Fourier transformed to produce the cross correlation between a reference (R) and signal (S).

4

interest and potential in this technology.

3. REAL-TIME IMPLEMENTATION ISSUES AND NEW FILTER ALGORITHMS

Real-time implementations are critical to the practical application of optical correlation, and the key device technology is the spatial light modulator (SLM), which introduces signal or image patterns as either corresponding light amplitude or phase modulation patterns across the optical beam. Spatial light modulator technology will not be reviewed here, but it is fair to say that the performance of these devices limits the performance of correlators, particularly with regard to space-bandwidth product and speed. Two important categorizations of SLMs are addressing method (electrical or optical), and modulation type (phase and/or amplitude) Correlator systems need SLMs at both input and Fourier planes, except as noted above,[7, 8] and the requirements differ depending on the application.

About a decade ago, a device to replace photographic film as the input means in the correlator of Figure 1, named the Liquid Crystal Light Valve, was made commercially available by Hughes Aircraft Corp. It was an optically addressed incoherent-to-coherent image converter, and as such could take an incoherent image, as from a CRT or imaging lens, and use it to spatially modulate a beam of laser light. Now for the first time real-time inputs could be used with the optical correlator.

Probably the most challenging SLM requirement is that of real-time matched filter modulation at the Fourier plane of a Vander Lugt correlator. The matched filter generally is a complex function, requiring independently controlled phase and magnitude spatial modulation. Most SLMs, such as the LCLV in its normal mode of operation, provide only magnitude modulation, sometimes accompanied by incidental phase modulation which is not independently controllable. Both Vander Lugt and Goodman realized clever ways to record the complex filter on an intensity or amplitude-only medium, for example, photographic film. However, film is not a recording medium that allows for quick, easy, or simple writing of new or modified filters, as is necessary for searching through an array of reference functions. The result was that optical correlation did not find much application in the real world.

The resolution and space-bandwidth product of the LCLV and other candidate SLMs are not sufficient to support practical holographic matched filters. However, with Lohmann's simplified version of the true holographic (Vander Lugt) filter, [the computer-generated holographic (CGH) filter], it is now possible to write very low resolution reference images on current SLMs. The state-of-the-art space-bandwidth product in pixellated devices is roughly 256 x 256 individually addressable elements. One problem with this is the very low diffraction efficiency of the SLM when used as a holographic filter. Most of the light goes into the zeroth order, a residual on-axis DC spike. We will now discuss some recent developments that greatly facilitate implementation of filters.

3.1 Phase-Only Filters

As we have seen, practical application of the matched filter in the past was limited because the filter is a complex quantity, and no effective means exists to implement such a filter in real time. Previous work in image analysis and synthesis had shown that the phase information, in the Fourier plane, was significantly more important than the amplitude information in image reconstruction.[9, 10] Caulfield[10] and Horner[11] originally conceived the idea of using only the Fourier phase of the signal to make a matched-like filter, as reported by Caul-

field. Horner and Gianino[12] first showed the feasibility of this idea through computer simulation, naming it the "phase-only filter" (POF). A much earlier paper[13] treated the case in which the amplitude information was omitted by hard-clipping from a holographic matched filter. The authors concluded that in a filter with "only phase control" the SNR in the correlation plane actually decreased by 29 percent. To our knowledge, this is the only mention in the literature of a correlation filter using only phase until References 10 and 12. The computer simulations showed that a phase-only filter would indeed compare quite favorably to the matched filter in that the structure of the correlation peak was typically a single narrow spike with negligible sidelobes.[12] The computer simulations also showed that the correlation peak was anywhere from 50 to 500 times higher than that of the matched filter. A simulation for a matched filter is shown in Figure 3(a), and the phase-only filter in Figure 3(b). The input signal was the letter 'G'. For an optical system to have a pure phase element in the filter plane made good sense because all the light gets through without attenuation. Everyday examples of this are lenses and prisms. In contrast, the amplitude portion of a matched filter attenuates the light since it is written on a passive device such as film or an SLM, and the peak transmission cannot exceed 100 percent, which usually occurs at the origin of the frequency plane. Also the light budget of a holographically-recorded matched filter is further eroded by diffraction efficiency, particularly if one tries to implement it on a spatial light modulator. Thus, the 100 percent light efficiency of the POF is a distinct advantage in a low-powered optical correlator where system and detector noise could be a problem. There are several sources of noise in an optical correlator. One is the so-called artifact noise. Since we use coherent light, any slight imperfections in the lenses, beam splitters, SLM, etc. generate their own diffraction patterns and propagate through to the correlation plane. Another source of noise in systems with insufficient laser source power is detector noise. Figure 3 also shows the large increase in correlation peak intensity obtained with the phase-only filter — a factor of over 300 when compared to the matched filter for the particular case studied. This corresponds to the case where the input laser energy and pattern are kept constant and only the filter type is changed. The light through the phase-only filter has a much better chance of breaking through the residual noise background and being detected. The reason for the considerably higher peak is simple conservation of energy. First, with the POF, all the energy gets through the filter plane (Figure 1) since there's no amplitude function to attenuate it; second, the energy that does get through to the correlation is concentrated in a much narrower peak, approximating a delta function in most cases.

Another advantage of the POF is the reduction in data required to store a filter because only the phase is retained. The Vander Lugt matched filter uses a carrier frequency to encode the complex amplitude information. A phase-only filter does not require a carrier or encoding scheme to write or record it, as it is a pure phase entity. This results in a large savings in space-bandwidth product (SBP). For example, an eight-level Lohmann type computer-generated hologram pattern[3] encoding a 128 x 128-sample matched filter requires 131 kb (kilobytes of storage. The same filter as a POF requires only 16 kb. In the next section we will see that we can further reduce this by going to the binary POF, which will require only 2 kb.

3.2 Binary Phase-Only Correlation

Shortly after the phase-only filter was proposed, a simplification of this concept was proposed independently by two researchers. Psaltis et al[14] showed experimental results using a symmetric input test signal and writing its corresponding bipolar Fourier transform on a binary SLM. Through a simple but clever modification they adapted the SLM, which was designed to write binary amplitude modulation, to write binary phase functions, that is, $e^{i\pi}$ and e^{i0} or -1 and +1. Horner and Leger[15] showed through computer simulation that arbitrary sig-
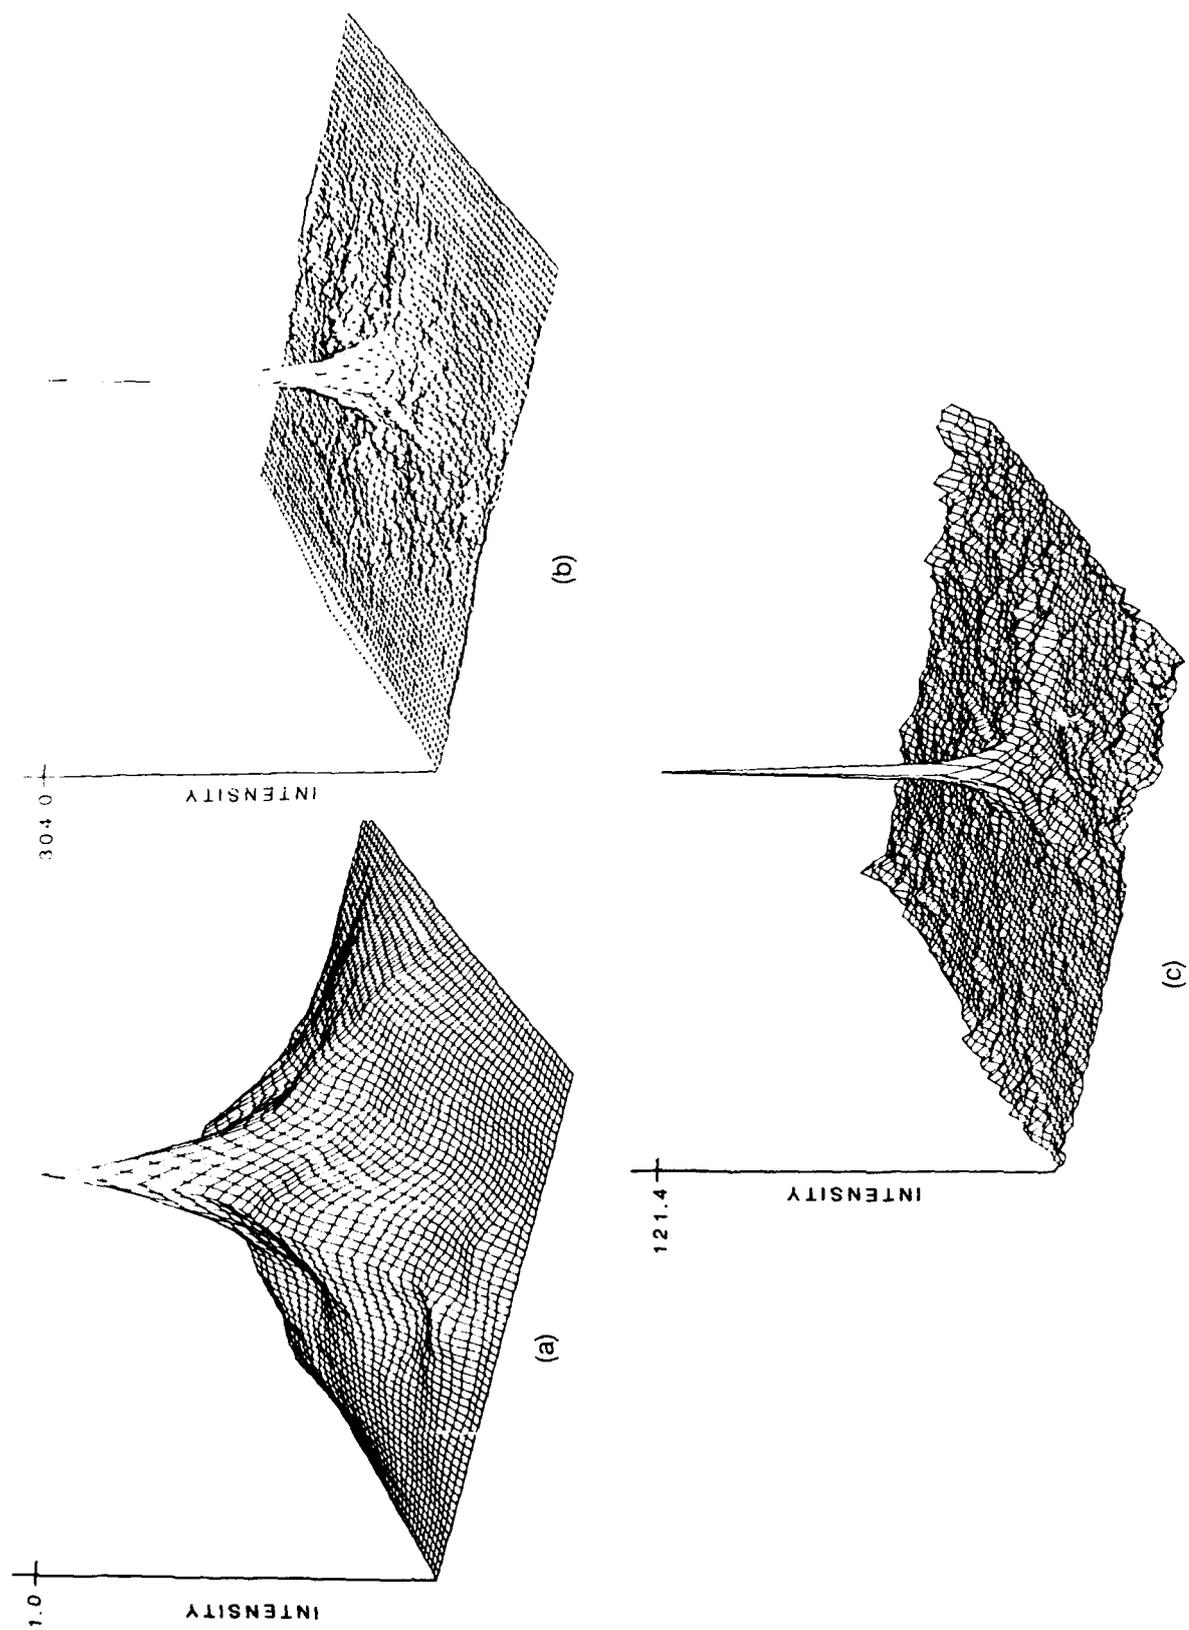
Figure 3. Simulation of correlation with (a) matched filter, (b) phase-only filter, (c) binary phase-only filter.

7

nals having no particular symmetry could be Fourier transformed, "brute force" binarized and the resulting binary phase-only filter (BPOF) used for pattern recognition. The basic idea is shown in Figure 4, where the continuous phase of the Fourier transform is shown in part (a). Curve (b) shows the binary signal derived from Figure 4 (a) by setting

$$\phi\text{BPOF} \begin{cases} = 0 \ degrees, \ F_r(v) \geq 0 \\ \\ = 180 \ degrees, \ F_r(v) < 0 \end{cases} , \tag{3}$$

where $F_r(v)$ stands for the real part of the Fourier transform.

It should be noted that the correlation plane output of an input correlated with a BPOF consists of both a correlation and a convolution signal. To see how this comes about, consider a sample point in the Fourier plane with a phase angle ϕ_i as shown in Figure 5. We can make this point binary by adding a phasor at a conjugate angle as shown by the dashed line. The sum of these two phasors, normalized to unity, is a phasor at zero angle. A similar argument applies to the left half-plane. Doing this at every point in the Fourier plane produces a function equivalent to a BPOF,

$$H(v) = k \cdot e^{i\phi} + k \cdot e^{-i\phi} , \tag{4}$$

where $k = 1/(2 \cdot cos(\phi))$. The impulse response of this function is the Fourier transform of Eq. (4). The Fourier transform of $e^{i\phi}$ gives an edge-enhanced version of the original signal, and the transform of $e^{-i\phi}$ gives an inverted edge-enhanced likeness of the original signal. Viewing correlation in terms of a shift, multiply, and integrate process, we can see that the inverted (180-degree rotated) component is responsible for the convolution while the normal image yields correlation. A photograph of the impulse response of a BPOF is shown in Figure 6(b). The filter was fabricated on a fused quartz substrate using VLSI and reactive-ion etching techniques.[16] The image from which it was made is shown in Figure 6(a). An interesting technique has been proposed for reducing the convolutional response by incorporating a random phase mask in the filter function and at the Fourier plane.[17]

The choice of angle of the Fourier-plane binarization axis used to generate a BPOF was studied by Kallman[18] who showed that if the axis was allowed to rotate and a BPOF made for each case, the output SNR would go through a maximum. This concept was further investigated and the name "threshold line angle" (TLA) was proposed for the binarization axis.[19, 20] The filter type investigated by Psaltis[14] involved thresholding on the real part of the transform, corresponding to threshold line angle = 0, and is called a cosine-BPOF because it can be derived from only the cosine transform. This BPOF type embodies information concerning only the even part of the reference pattern. The sine-BPOF is defined with threshold line angle = 90 degrees, corresponding to thresholding on the imaginary part of the Fourier transform, and was first used by Horner and Leger.[15] It encodes only odd-part information from the reference pattern. Hartley-BPOFs[19] result from thresholding on the Hartley transform or equivalently by thresholding on the Fourier transform with threshold line angle = 45 degrees.[20] This type BPOF encodes a mixture of equal components of even and odd reference pattern information. Flannery et al[20] studied BPOF performance using a tank embedded in a cluttered background and found a variation of about 2 db. ir signal-to-clutter ratio when the threshold line angle was varied over its full range.
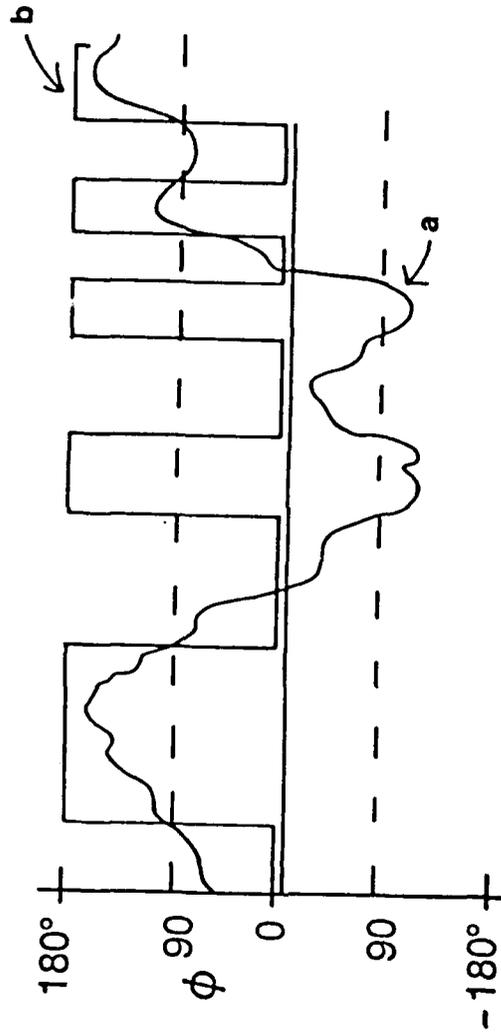
Figure 4. Binarization of a function by thresholding; curve (a) is the original function, and (b) is the binarization by setting $\phi = 0$ when REAL F (v) $\geq$ 0, and $\phi = \pi$ when REAL F (v) < 0.

9
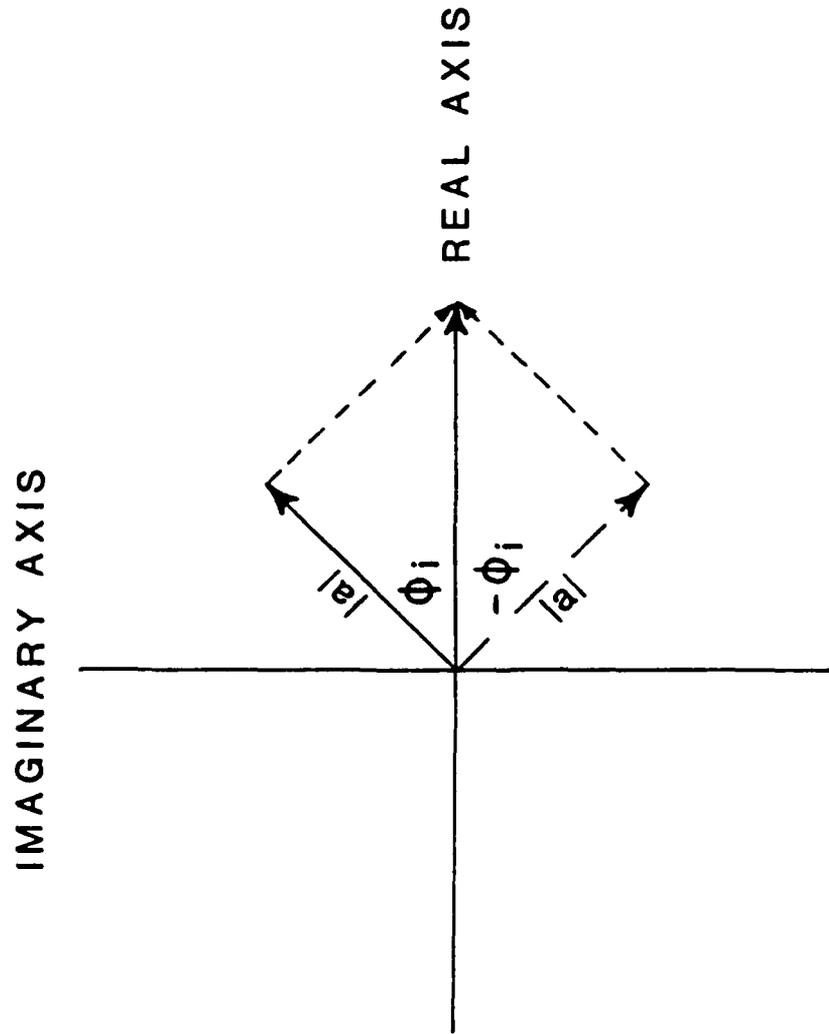
Figure 5. Binarization of continuous phase function by the addition of its complex conjugate.

(b)

(a)
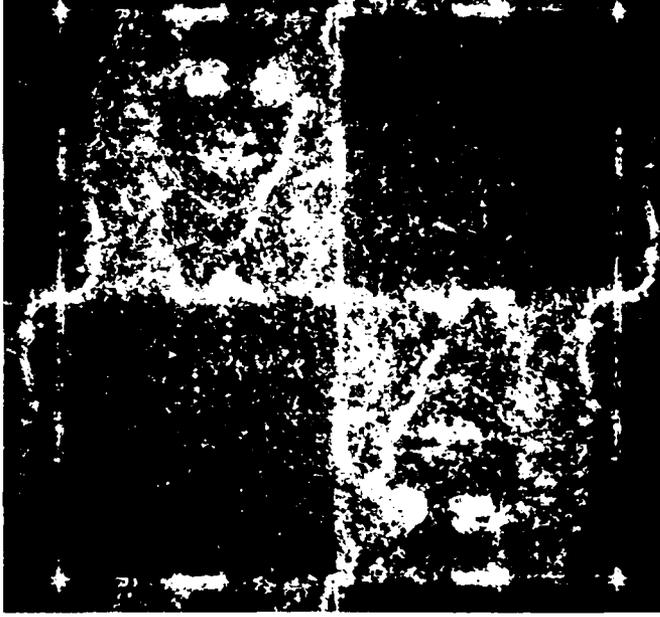
Figure 6. Impulse response of BPOF; (a) original image, (b) impulse response of BPOF made from (a).

3.3 Noise Performance of Phase-Only Filters

The theoretical noise performance of POFs and BPOFs relative to the classical matched filter has been investigated. The results depend on the definition of signal-to-noise ratio (SNR), and the type of noise used. The POF with an appropriate choice of bandpass has been shown to be the optimum phase-only filter function for classical noise (white, additive, zero mean, Gaussian) with typical SNR figures of 5 db less than the classical matched filter.[21, 22] However, when the sidelobes of the correlation pattern are included in the noise metric, as originally proposed by Horner et al,[24] the POF and BPOF inevitably have higher SNRs than the matched filter. Recently Dickey and Romero[23] showed that the POF, in the original form proposed,[12] is the optimum of all possible unity-amplitude filters (that is, generalized POFs), with regard to a metric they called peak-to-sidelobe ratio; it is nearly identical to the SNR defined by Horner et al,[24] for correlation.

Dickey and Romero[23] also established a dual optimality of the POF or unity-amplitude filter. The latter term implies a more general definition of the POF whereby any phase value can be assigned to a filter element by an arbitrary formulation. In discussion up to this point, the phase-only filter was defined in the sense that it was the phase derived specifically from the Fourier transform. The more general meaning of phase-only, that is, unity-amplitude, will also be used later in this report, and the choice of specific or general meaning of phase-only should be clear from context.

Dickey et al[25] have established a lower bound of -6 dB for the signal-to-noise performance of the Hartley-BPOF relative to an optimal POF, and an upper bound of 3 dB for the improvement over the Hartley case obtainable by the use of the optimum threshold line angle value, although the latter bound has not been approached in any reported cases.

Recent research shows that the POF and BPOF do a considerably better job than a classical matched filter when the noise is more akin to what is actually found in the real world,[26] instead of the white Gaussian noise associated with the matched filter. Figure 7(a) shows a tank in a cluttered desert background; Figure 7(b) shows the correlation response for a matched filter made from just the tank (without background), and 7(c) and (d) show the responses using phase-only and binary phase-only filters. The matched filter shows no clear correlation peak at the on-axis position where it should be located. The other two filters clearly show the correlation peak exactly where the object is located.

In comparing the various filters it must be remembered that the POF generally provides a narrower peak, higher optical efficiency, and greater discrimination against other shapes, all of which may be of great practical value. The tradeoff between these performance aspects is highly application specific.

3.4 Amplitude-Encoded Phase-Only Filter

It is possible to encode the BPOF modulation onto an amplitude device and preserve the essential features of phase-only filtering.[27] To see how this is possible, note that when the binarized phase values are 0 degrees and 180 degrees, the filter values are 1 and -1. If we add a bias term and renormalize, these become 1 and 0. The bias yields an extra term in the correlation result:

$$C(x_1) = FT^{-1} \left\{ F_1(v) \cdot (\tfrac{1}{2} + \tfrac{1}{2} \cdot e^{i\phi b}) \right\},$$

$$= \tfrac{1}{2} \cdot \{ f_1(x_3) + C_b(x_3) \}$$

(5)

where $C_b(x_3)$ is the correlation signal obtained with the true BPOF. Thus, an image of the input pattern

12

EFFECTS OF CLUTTER ON CORRELATOR PERFORMANCE
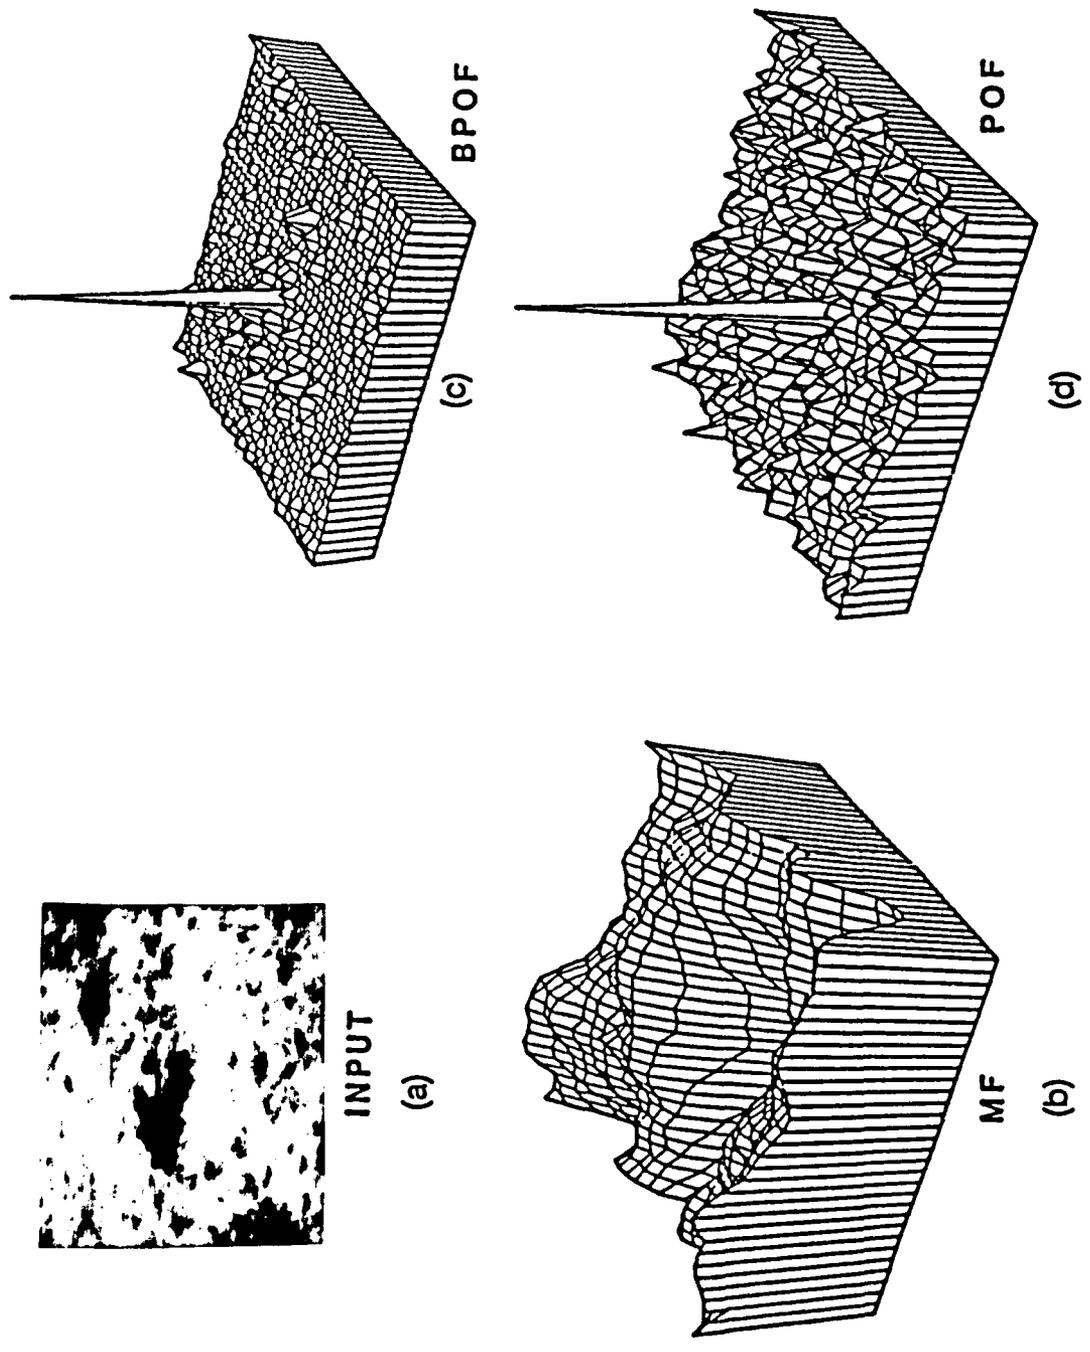
Figure 7. Comparison of matched filter (b), binary phase-only filter (c) and phase-only filter (d) when the target (a) is immersed in clutter noise.

13

appears superimposed on the bin: '_' phase-only correlation response. If the signal to the correlator is an electronic signal, and the signal in the correlation plane is detected using a CCD array, the two could be subtracted to improve the apparent SNR in the output plane, although this would require additional post-processing time. However, experiments and simulations[27] show that the correlation peak is typically ten times stronger (in intensity) than any secondary peaks. Amplitude-modulating SLMs are more common than phase-modulating ones, and this technique permits their use for binary phase-only filtering. Media-like photographic film can also be used. Thus, a researcher with only a small desktop computer can fabricate BPOF-like filters by photographing a plot of the amplitude-encoded binary pattern from his CRT display or laser printer.

4. SMART FILTERS

The matched filter has a well known sensitivity to distortion of the input pattern from the reference pattern.[28] For example, in-plane rotations of several degrees, or scale changes of several percent, typically result in a 50 percent reduction in the correlation peak intensity. The sensitivity strongly depends on specific parameters such as object shape and spatial frequency bandpass of the correlation filter. It has been clear that this distortion sensitivity severely limits the practical application of correlators if they incorporate only a fixed single filter, since most applications of interest involve many possible distortions of target object patterns. The application of a number of filters to a given input scene, that is, real-time filtering, seems an obvious approach and, as will be seen, such system concepts are under investigation. However, the motivation also exists to make each filter as powerful as possible to minimize the number of filters required and to simplify system operation. Such filter formulations, broadly labeled "smart filters," are a current topic of much research interest and will be reviewed here. Only space-invariant filters will be discussed. They provide location estimates intrinsic to their operation, do not require centering of the input object for proper operation as do space-variant types, and are the type most frequently used in Fourier optical processors.

4.1 Problem Definition

The problem of designing distortion-invariant correlation filters is best viewed in a discrimination context. In any practical application not only must expected distortions of targeted, or in-class, patterns (objects) be recognized (usually by virtue of a high correlation peak response), but one or more types of noise, or out-of-class, patterns must be rejected (usually by virtue of a relatively low correlation response). The classical matched filter is the optimal solution for the case of a single in-class pattern and regarding stochastic noise as the out-of-class pattern — a case having limited practical application.

A problem with complexity more typical of practical applications is recognizing certain types of vehicles, regardless of orientation, in scenes that include other (nontargeted) vehicles and natural or man-made clutter such as trees, bushes, buildings, roads, etc. Random noise also may be present in the input patterns as a result of the imaging system or viewing conditions. As will be discussed, the need to reject out-of-class patterns places serious constraints on designing widened in-class response (distortion invariance).

Many applications would benefit from the use of the space-invariant feature of Fourier-based correlation to furnish a location estimate of each recognized target, based on correlation peak location in the output plane. This involves additional trade-offs in filter design since, in general, optimization of peak localization (that is, achieving narrow peaks) is inconsistent with maximizing classification accuracy (for example, achieving highest

signal-to-noise).

Performance evaluation, and thus optimization, of correlation filters requires a systematic set of performance metrics. Unfortunately, universally-accepted, standardized metrics do not exist, and indeed it may be argued that a single set of such metrics is inappropriate due to differing requirements of various applications. Nevertheless, three broad areas of performance measurement usually important to most applications can be identified. These are discrimination, optical efficiency, and peak localization.

Discrimination includes the concepts normally associated with terms such as signal-to-noise (SNR) and peak-to-clutter), that is, the ratio of target, or in-class, correlation peaks to responses associated with non-target, or out-of-class, inputs. A common definition of SNR uses some measure of correlation peak energy divided by the RMS energy computed over the correlation plane. An alternate concept, frequently used in defining peak-to-clutter, is to divide by the intensity of the highest non-target output peak, yielding a ratio sometimes called "peak-to-secondary." In any case, issues involving the specification of the effective size of the desired correlation peak and the effective energy-integration area (for example, detector resolution) must be addressed.

Horner efficiency measures the fraction of processor beam power concentrated in useful correlation peaks at the output plane.[10, 11, 24]

Peak localization addresses those properties of correlation peak shape that impact the accuracy and degree of difficulty of forming location estimates, essentially peak width and sidelobes. This area of performance measurement is the poorest-defined of the three discussed here. Concepts used for the characterization of smoothing-window performance in spectrum estimation using the Fourier transform would seem to have applicability;[29] for example, the 6 dB (10 dB) width (including any sidelobes) of a correlation peak might be a reasonable metric. Another approach, illustrated by Kallman,[30] is to chose an output region size determined by application parameters (for example, required location accuracy and/or detector resolution) and count any response within the region as a valid correlation peak component and all others as nontarget responses. The latter approach has appeal because it corresponds well to the actual situation that arises when simple thresholding of detector array outputs is used to define correlation peaks.

The general problem of smart filter formulation may now be stated: Specify the modulation values to be encoded in each resolvable or addressable element of a correlation filter to optimize correlation performance for a given correlation (discrimination) problem. We note that the number of free parameters involved is determined by the space-bandwidth capacity of the filter medium and, in general, can be quite large, for example, 16,384 for a 128 x 128-element filter. The specification of the desired correlation performance also can be quite complex, involving simultaneous optimization of several performance metrics over a robust set of input scenes. Considerable judgment will probably be required just to formulate an overall metric that properly reflects the tradeoffs between individual performance metrics such as discrimination, efficiency, and peak localization, and incorporates a discrimination performance metric appropriately weighted or averaged over the input domain. An additional complication is the constraint of listed ranges or discrete values of modulation imposed by certain filter types such as POFs and BPOFs. The nonlinear relationships of these filter functions to the Fourier transform upon which they are based precludes the straightforward application of many smart filter formulations developed for continuous complex filter modulation. Yet formulations optimized for limited encoding ranges are of practical importance due to the limitations of spatial light modulator devices needed to implement filters in real time.

Thus, the general smart filter design problem is seen to be both formidably complex and very case-specific, that is, the optimum filter depends on the particular correlation problem and upon designer judgment regarding performance tradeoffs. It is not surprising that no definitive analytical theory of smart filter design exists.

However, effective smart filter formulations have been developed and notable instances of these will be reviewed here.

4.2 Smart Filters For Continuous-Valued Complex Modulation

First we discuss smart filter formulations designed with no restrictions imposed on their modulation values. The full range of complex numbers is assumed available in a SLM or hologram for physical realization of the filter function. Filter formulations designed for implementation with limited modulation ranges (for example, binary phase modulation) will be discussed in a following section.

4.2.1 COMPOSITE FILTERS AND LINEAR DISCRIMINANT THEORY

A concept underlying many space-invariant smart filter formulations is the composite filter[31] defined by an impulse response that is a weighted sum of image components, frequently a training set. (Whether the sum is performed in image or Fourier, or filter, space is irrelevant since transforms are linearly related to images.) The intuitive concept of a simple composite filter for distortion-invariant object recognition is to superimpose different views of the object in the filter so each view will have an equal correlation response — in its simplest form an average filter. The concept is valid although not a panacea, as will be seen. Equal correlation response to each view requires a more sophisticated method of weighting the superimposed components than the equal weighting implied above, because they are not orthogonal. This was first addressed by Caulfield and Maloney[31] and later in the synthetic discriminant function filter formulation, to be discussed.

Composite filter formulations can be used to illustrate a tradeoff between distortion-invariance (or widened in-class response) and discrimination for correlation filters which appears to be fundamental, although subject to optimization to varying degrees with smart filter techniques. Kumar and Popchapsky[32] analyzed and studied this tradeoff for composite filters designed to recognize in-plane rotations and scale changes of patterns in the presence of additive random noise. Figure 8, taken from that work, illustrates that SNR (that is, discrimination against the out-of-class pattern called noise) degraded by 7 dB as the number of superimposed filter components was increased from 1 to 72 to make the filter rotationally invariant (that is, widen the in-class response). Note that the SNR degradation is close to the number of components used, that is, $10\log(72) = 18.6\,dB$. This example clearly illustrates the basic challenge of smart filter design.

Another important concept of composite filters, interpolation between training set components, also is illustrated by Figure 8. Note that when the training set components are spaced by 30 degrees (12 training set components), there is a large droop in correlation response for input rotations between the training set angles. Training with five-degree increments appears necessary to achieve good interpolation. The filter designer must choose a training set that ensures good interpolation between components. An intuitive approach is to base the interval between components on the distortion sensitivity exhibited by a single-view filter, which seems to apply to the case shown in Figure 8. A broader perspective of the interpolation issue is that a training set must be statistically representative of its class. This is expected to be a significant problem in many cases because of the high dimensionality of the images being processed relative to that of typical composite filters (which is on the order of, or equal to, the number of training set components). The composite filter approach may be viewed as a method of parametrizing filter design which dramatically reduces the dimensionality of the problem (from the space-bandwidth of the filter) but involves attendant limitations including the interpolation problem.
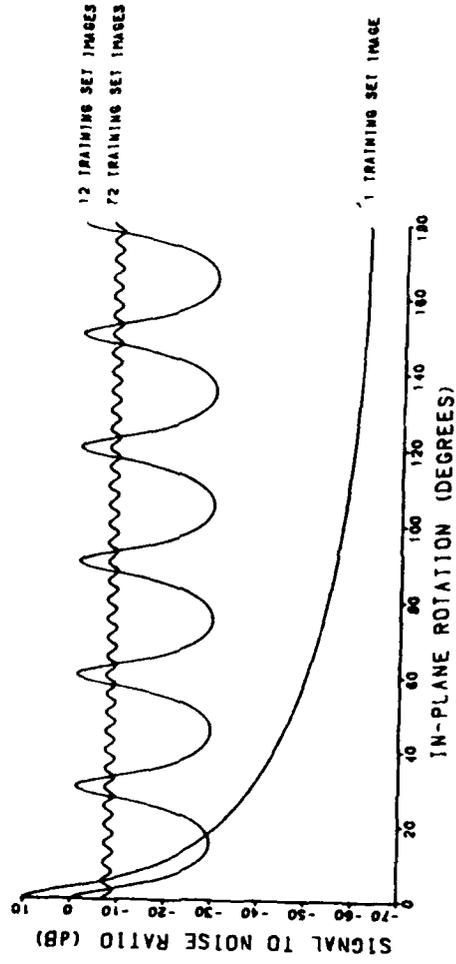
Figure 8. Analytical results indicating the interpolation-vs-discrimination tradeoff of composite filters. Correlation signal-to-noise vs rotation angle of input pattern for filters computed using various training-set intervals, taken from Kumar and Popchapsky [Ref. 32].

The on-axis correlation value, (the amplitude value found at the center of the correlation plane for a centered input pattern), is the functional inner product of the unshifted input and reference patterns, which may be viewed either as two-dimensional-or one-dimensional- (sequentially-scanned) functions or (when sampled) vectors. Thus, the correlation (with respect to the on-axis result only) implements a linear discriminant function, which has led to a large body of work[33-37] concerning optical implementation of linear discrimination concepts.[38] The performance of such classifiers can be limited in some applications by their reliance on only the on-axis correlation response. In-class distortions or out-of-class inputs not represented in the filter can yield other peaks greater than the on-axis value. These are confused with valid peaks since, in general, the postprocessor does not know the location of inputs and, in fact, is usually relying on the correlation process to furnish a location estimate, as mentioned.

4.2.2 SYNTHETIC DISCRIMINANT FUNCTION FILTERS

The first widely used, and probably best known, smart filter was the synthetic discriminant function (SDF) filter,[33]although several closely related approaches previously had been reported.[31] The SDF filter is a composite filter in which the weights are set, using linear discrimination techniques, to yield specified on-axis correlation responses over the in-class and out-of-class training sets. In its first form[33] called the equal-correlation-peak (ECP) SDF, the on-axis covariance matrix for the training set is formed and linear algebra techniques are applied to solve the composite filter weights which ensure equal-correlation-peak response over the in-class training set. Problems with (spurious) correlation responses larger than the (desired) on-axis response were addressed by the correlation-SDF (c-SDF) variant which incorporates control of sidelobe responses not included in the equal-correlation-peak formulation.[39]

Variants of the SDF have been developed, several of which are called optimal linear discriminant function (OLDF) filters.[35] Notable recent developments include the minimum average correlation energy (MACE) SDF[40] and the minimum-variance SDF (MVSDF).[41] A generalized SDF formulation, encompassing most of the previous types as special cases, has been reported recently.[42]

SDF filters have demonstrated good performance both in simulations[39-41, 43-47] and experiments.[40, 47] Except for the cSDF, performance of these analytical SDF formulations is limited by their control of only the on-axis correlation response and that no phase variation of that response is allowed, although the detected output of the correlator is intensity, which is not sensitive to phase. Also, a practical disadvantage is that straightforward application of minimum-variance SDF and generalized-SDF formulations requires linear algebra operations on huge matrixes (N-squared on a side for N X N images), since the eigensolution of the raw image space is part of these formulations. However, the SDF filter variants probably comprise the best developed and most powerful analytical smart filter formulations.

Generalized matched filters[37] also are based on linear discrimination concepts, applied in this case in the Fourier (filter) domain. The goal of the filter formulation is a (complex) linear discriminant function (the filter) that best separates the target and nontarget classes. Rigorous application of the formalism involves eigensolution of huge matrixes, but with a simplifying assumption (that different Fourier components are uncorrelated over a class) tractable solutions result. Simulations using the simplified formalism have yielded good results in a case in which nine distortions of an aircraft silhouette were to be discriminated from zero-mean Gaussian noise.[37] This formalism also has the limitation of controlling only the on-axis correlation value.

Kallman has reported SDFs formulated by iterative numerical techniques using min-max performance criteria in which complex weighting coefficients were allowed and response was controlled over the entire output

plane.[30, 48] In test simulations, these filters performed substantially better than equal-correlation-peak SDF filters for the same inputs, as illustrated in Figure 9. Part (a) is a typical input scene from these simulations,[48] consisting a target tank embedded in clutter composed of a "junkyard of tank parts." Figure 9 (b) and (c) are profiles of the correlation planes for the equal-correlation-peak and Kallman-recipe SDF filters, demonstrating substantial improvement in peak-to-clutter.

The iterative-search filter formulation algorithm used by Kallman does not require operations on huge matrixes, but is computationally intensive in its own right, in part because each point in the correlation plane is examined and controlled. The use of complex weighting coefficients doubles the dimensionality of the filter parametrization over that of the equal-correlation-peak SDF filter formulations using real weights. Although not expressible in closed form, and thus not amenable to general analytical performance assessment, the Kallman SDF approach uses a filter performance metric that is more appropriate for many applications, and it has the potential to yield superior filters as a result of the additional free parameters associated with complex weights. The Kallman recipe is an example of a general class of filter design approaches involving parametrization of the filter (to reduce dimensionality) and iterative search over the parameter space to find a (probably local) maximum of a performance metric. Such approaches presumably could benefit from global search concepts such as simulated thermal annealing[49] which are purported to find performance maxima near the global maximum. A simulation demonstrating the use of this approach to optimize discrimination between two characters in a BPOF filter has been reported.[50]

4.2.3 CIRCULAR HARMONIC FILTERS

Filters based on circular harmonic decompositions are naturally suggested when invariance to in-plane rotation is desired, and they have been the basis for a substantial body of research.[51-55] A fundamental limitation of circular harmonic composite filters is that only in-plane rotation invariant response is obtained, and only if a single circular harmonic number is incorporated in the filter.[51, 55] This comprises a serious limitation for many practical cases, since a single harmonic may contain only a small fraction of the shape information critical to recognition and discrimination. One attractive approach for pattern recognition using circular harmonics involves forming a feature vector based on the responses of a bank of filters, each based on a different circular harmonic number.[53-55] The resulting feature vector is rotationally invariant and can embody robust information related to object shape, supporting good recognition and discrimination.

Although a given circular harmonic always provides rotation-invariant correlation intensity response, the distribution of correlation responses over harmonic number depends on the center-of-expansion chosen for the decomposition. Excessive sidelobes can result if the proper center is not used.

Fourier-Mellin descriptors, comprising radial moments of circular harmonic filters, have been investigated for scale- and rotation-invariant pattern recognition.[55]

Circular harmonic smart filter techniques are particularly attractive for applications where in-plane rotation is the predominant form of in-class distortion.

4.2.4 MELLIN RADIAL HARMONIC FILTER

Mendlovic et al[56] have reported a filter based on the decomposition of a function in terms of azimuthally weighted radial functions (called Mellin Radial Harmonics) that exhibit radial scale invariance. This formalism is a dual of the circular harmonic filter and provides filters that are invariant to in-plane scale changes, provided

Figure 9a. Example of SDF filter performance achieved by Kallman [Ref. 48]. Input scene: tank in clutter described as "a junkyard of tank parts."
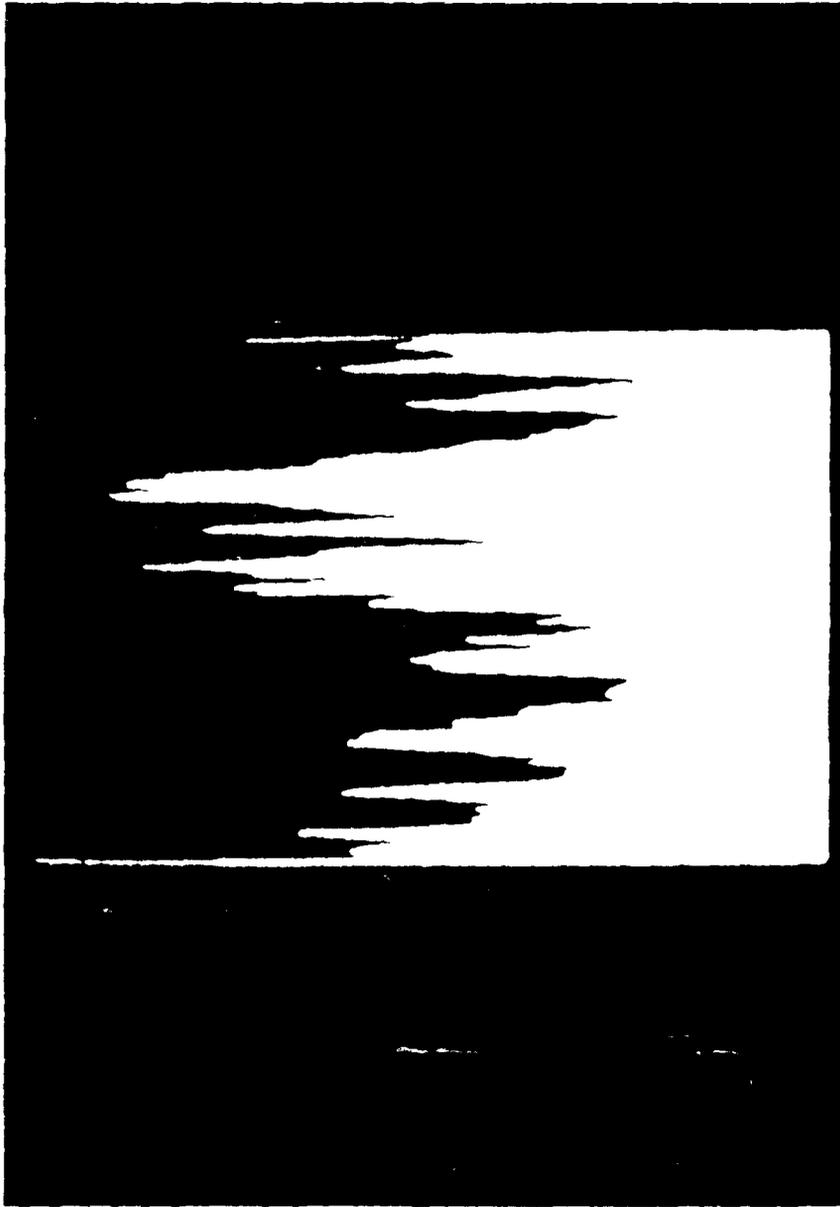
Figure 9b. Example of SDF filter performance achieved by Kallman [Ref. 48]. Correlation profile using usual SDF formulation.
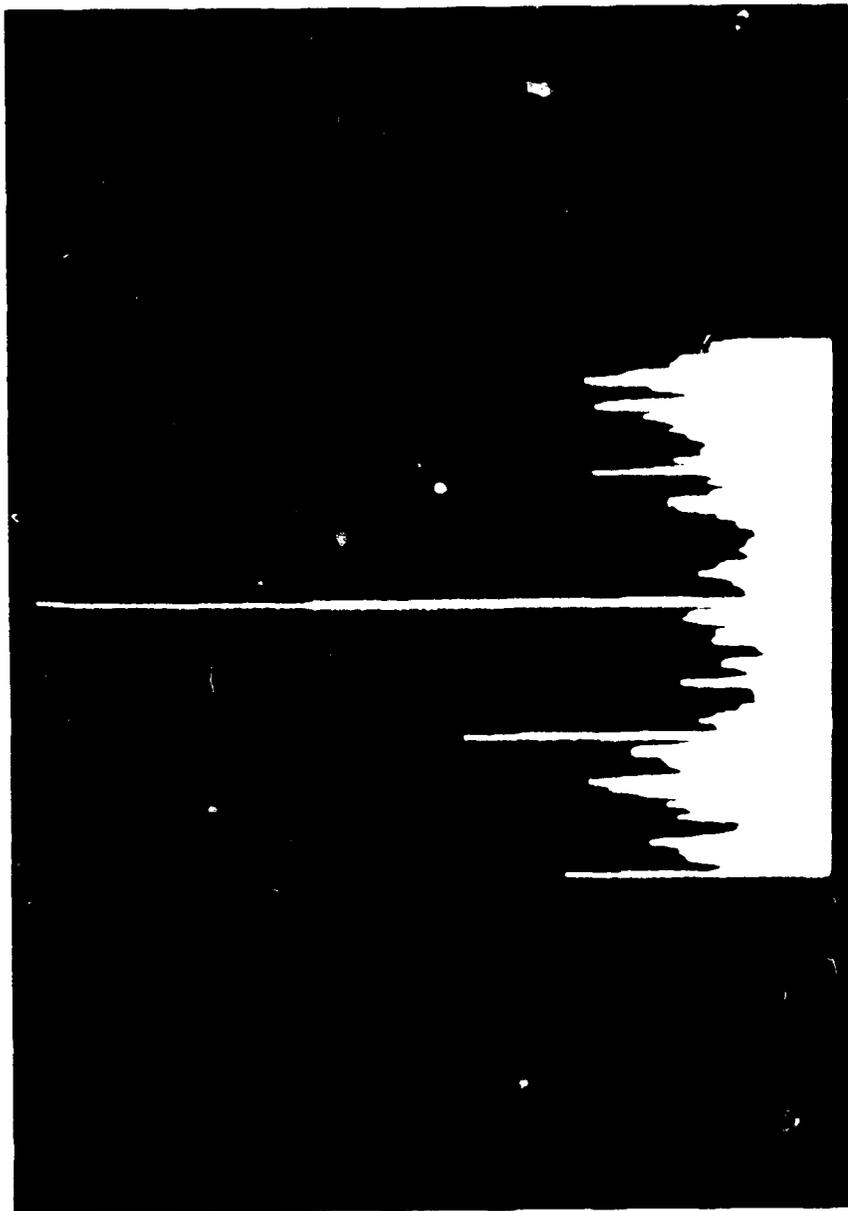
Figure 9c. Example of SDF filter performance achieved by Kallman [Ref. 48]. Correlation profile using Kallman formulation.

only a single Mellin harmonic is embodied in the filter. An analogous dependence on proper choice of center-of-expansion also exists. Experimental correlations demonstrating recognition and discrimination of the letter 'E' over a 4:1 range of scale were reported.[56]

4.2.5 SPATIAL-TEMPORAL CORRELATION FILTER

The (space-invariant) spatial-temporal correlation (STC) filter[57] formulation is based on the polar-log coordinate transform, which was the basis of a well-known spacevariant correlation technique possessing both in-plane scale and rotation invariance.[58] The method constructs a filter to recognize one object and defines a signature control function (SCF) that specifies the variation of correlation peak amplitude over 360 degrees of rotation and a range of scale change. If the singature control function is set constant over the distortion control region, an equal-correlation-peak response filter is produced, but in principle any variation of peak response as a function of distortion variables may be programmed.

A particular case demonstrated in reported examples involves programming for uniform response over the scale distortion range and a sharply peaked response in rotation angle. Thus, if the input is rotated, a temporal signature which peaks at a time determined by the rotation angle of the target object is produced at the correlation plane.[57] Simulations using computer-simulated aircraft outlines demonstrated the expected behavior for this type filter.[57] The programmed temporal signature was obtained only at an output location corresponding to the location of the target aircraft in the input, and was obtained for differently-scaled versions of the target. In other simulations,[59] spatial-temporal correlation filters for estimating the rotation angle of airplane outlines based on correlation peak value were constructed.

The spatial-temporal correlation filter controls only the on-axis correlation peak response for a single target pattern, and addresses only in-plane distortions. It uses no training set, but addresses distortion in a continuous fashion. The filter has a closed-form mathematical expression and its formulation is straightforward, involving log-polar coordinate transformations and discrete Fourier transforms.

4.2.6 THE LOCK-AND-TUMBLER FILTER

The lock-and-tumbler (LAT) filter[60-62] is a composite filter based on a superposition of eigenimages of a training set representing any range of any type of distortion. A spectral iteration technique is used to find the (complex) weights that solve the design constraint of the filter, which is that the on-axis intensity be constant over a bank of LAT filters for a given training-set input. Typically 100 or more eigenimage components are included in each filter. In the special cases where only in-plane rotation or scale distortions are addressed, the eigenimages are circular harmonics or Mellin components respectively.

The use of the LAT filter is fundamentally different from all other filters in that the correlation output pattern must be post-processed to find the point(s) exhibiting minimum variance taken over a number (typically 20) of input-filter correlation pairs. Recognition and location of a target is based on existence of a peak in the post-processed pattern of mean-to-variance ratio calculated for each pixel. Simulations[60, 62] and experiments[61] have verified the feasibility of the filter.

The bank of LAT filters is distinct only by virtue of the starting-seed of expansion coefficients used in the spectral iteration. A multiplicity of filters satisfying the identical design constraint is possible since the number of spectral weights greatly exceeds the number of constraints applied in the filter solution. Eigensolution involves linear algebra operations on M X M matrixes, where M is the number of training set images. To minimize

computations but retain sufficient interpolation, the training set must be carefully chosen, and this is done with regard to the cross-correlation coefficients of the set.[62] In its first reported form,[60] the LAT involved circular harmonic eigenimages and the equivalent of a bank of N LAT filters could be obtained by taking N different relative rotations of the input and a single filter.

The LAT has been applied successfully in simulations to provide perspective-invariant recognition of a model tank, using captured, median filtered, and edge-enhanced images, over 360 degrees of rotation and 60 degrees of zenith-angle variation.[62] In this case 380 training set images and an equal number of eigenimages were used. Interpolation and discrimination of other shapes was demonstrated using 20 LAT filters. Figure 10 is an example taken from these simulations. Figure 10(a) is an input scene containing three different perspective views of the target tank and one nontarget vehicle image. Figure 10(b) is a processed output-pattern plot indicating recognition, with good peak-to-clutter ratio, of the three targets and rejection of the nontarget.[62]

LAT filter formulation is complicated and computationally-intensive, but its use of eigenimages to ensure incorporation of all training set information is well founded and the performance exhibited in the perspective-invariant recognition problem just discussed is impressive. LAT issues requiring attention include the effects of clutter, nonzero backgrounds, and nonuniform illumination, and real-time implementations. (The filter is complex valued, hence no SLM can write it directly.)

4.2.7 SUMMARY OF STATUS OF SMART FILTERS FOR CONTINUOUS-VALUED ENCODING

Our discussion has indicated that a diverse selection of smart filter techniques is evolving. They differ in mathematical approach from simple ad hoc concepts to sophisticated treatments; in solution method from closed form to iterative searches; and in goals from controlling only peak response to controlling responses over the entire correlation output plane. Some use training sets; others do not. As stated, no general all-purpose formulation has emerged, and such a development would be surprising in view of the complex and sometimes conflicting demands that apply.

Correlation (discrimination) problems also span a diverse range. They differ markedly in complexity and goals, and some readily furnish training sets while others may not. Thus, it seems fitting that a wide spectrum of smart filter design tools is being developed. All of the filter formulations have shown good performance, at least in computer simulations, applied to specific problems. With the possible exception of the equal correlation peak SDF filter, none have been tested sufficiently, either experimentally or by simulation, to establish a basis for definitive assessment of their overall capability.

Experimental implementation of the filters discussed in this section has lagged theory and simulations, both because simulations are a logical and easier first step, and because spatial light modulator technology does not readily support the real time implementation of continuous complex-valued filter functions. Continued research and development is expected to attack both these limitations.

4.3 Smart Filters for Limited Modulation Range Encoding

In this section we discuss smart filter formulations for implementation with limited modulation ranges, in particular for POF, BPOF, and ternary phase-amplitude filter (TPAF) encoding. The TPAF is a more recent development and will be described. As mentioned, the motivation for such filters is practical, related primarily to the necessity of implementation with available SLMs that only modulate either phase or amplitude, with only
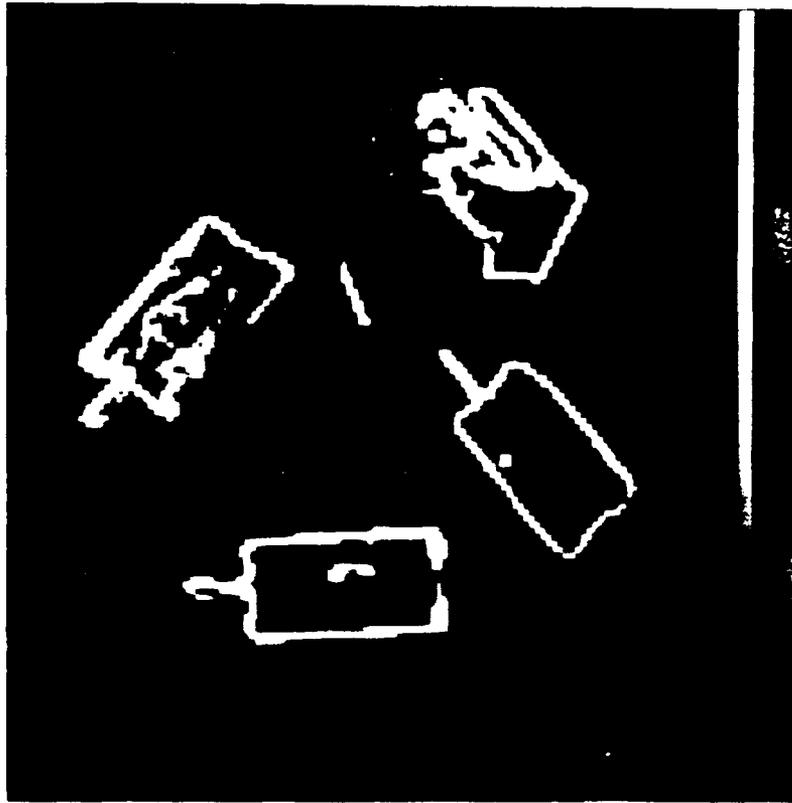
(a)
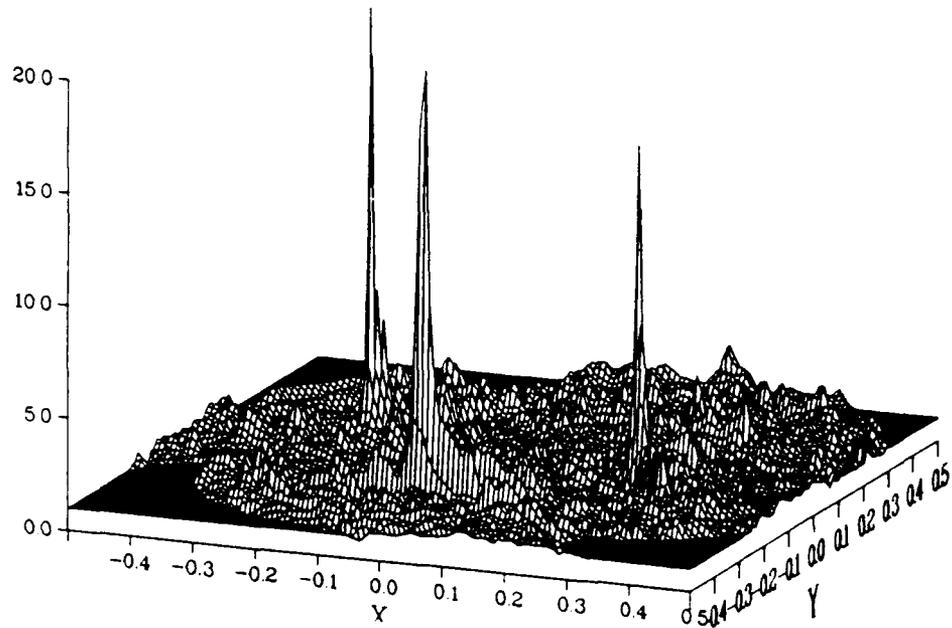
(b)

Figure 10. Example of Lock-and-tumbler filter performance [Ref. 62]. (a) Input scene containing three views of target tank (two of which were not in training set) and one nontarget (bulldozer), (b) Output plane showing peaks corresponding to three views of target tank.

binary or ternary (three-level) modulation available in some cases, and to reduce on-line storage requirements.

The fundamental smart filter design problem has been defined and is the same here except that the additional constraint of the limited encoding range must be treated. As mentioned, many filter design formalisms are profoundly impacted by this constraint. Processing a filter function to yield the required encoding, (for example, taking only the phase, or thresholding to define binary phase), is a nonlinear operation and produces a totally new filter function not necessarily satisfying the design goals of the original continuous-valued filter function. Many formalisms, including the original equal-correlation-peak SDF and LAT filters, cannot readily encompass the required nonlinear step, and thus have not been revised for limited-range encoding.

4.3.1. CONVERTED CONTINUOUS-VALUED COMPLEX FILTERS

The phase-only concept is a very general principle that can be applied to any complex filter function by merely retaining the phase (or binarizing it) and setting the modulus to unity, as first proposed by Horner and Gianino for the p-SDF filter.[24]

Researchers first created smart POFs and BPOFs for testing by simply converting continuous-valued complex smart filters of the equal-correlation-peak SDF, Kallman-SDF, or GMF types[18, 24, 45, 63] It was recognized that this was sub-optimal, for the reasons stated above, but the simplicity of the approach was appealing and, initially, no filters specifically optimized for these encoding methods existed. Results were mixed. In some cases reasonably good filters resulted, although they did not provide the performance (for example, equal correlation peaks over the in-class set) designed into the equal-correlation-peak SDF.

4.3.2 DIRECTLY OPTIMIZED POF AND BPOF FILTERS

Smart filters optimized for phase-only and binary phase-only encoding over in-class and out-of-class training sets were reported by Kallman[64] and by Jared and Ennis,[65, 66] using distinctly different approaches.

Kallman directly optimized the N^2 phase values of a $N \times N$ POF using iterative search techniques and his min-max performance criteria.[64] BPOF filters were created from the resulting POF patterns after searching for an optimum threshold line angle, and thus, were not directly optimized for that encoding. These POF and BPOF filters provided significantly better discrimination performance than did converted SDFs for out-of-plane distortions of military vehicles in random backgrounds.[64] The filters also exhibited interpolation between training set components. The numerical methods used in the formulation of these filters involve considerable computation but they address meaningful optimization criteria and provide good performance.

Jared and Ennis[65] developed an ad hoc iterative technique, called the filter-SDF (fSDF) which adjusts the weights of a composite transform, based on training set images, so as to control the on-axis correlation response of the converted POF or BPOF over the training set, that is, addressing the goals of the equal-correlation-peak SDF filter. The formulation process is numerically relatively simple, but does not converge for some cases. The use of only real weights and control of only the on-axis peaks are additional limitations of the approach as reported, just as for the original SDF filter. Impressive simulation results for a limited set of test images have been reported.[66] If the above limitations are acceptable, and convergence can be obtained for the problem at hand, this technique provides a relatively simple method to construct POF and BPOF filters of the SDF type.

4.3.3 CIRCULAR HARMONIC PHASE-ONLY COMPOSITE FILTERS

Circular harmonic phase-only filters have recently been reported.[67, 68] As with continuous-valued complex modulation, composite circular harmonic phase-only filters must be restricted to one harmonic to retain rotation invariance. Such filters have demonstrated performance superior to continuous-valued versions for test cases involving a small set of printed characters[68] and images of the space shuttle on a background.[67] Even better performance was obtained when some amplitude information derived from the nontarget letter was incorporated in the filter, but such a filter is no longer phase-only.

The behavior of phase-only circular harmonic filters relative to continuous-valued versions is as might be predicted based on the comparison of phase-only matched filters to their continuous complex counterparts. Narrower peaks, greater discrimination, and much higher Horner efficiencies are exhibited by the continuous complex filters. Basic characteristics of circular harmonic filters, including in-plane rotation invariance (for a single harmonic) and sensitivity to center-of-expansion are retained.

4.3.4 TPAF SMART FILTERS

Smart filter formulations for ternary phase-amplitude encoding (-1, 0, 1 modulation levels) have been reported.[69, 70] This work treated the filter as a product of binary amplitude (0, 1 modulation) and binary phase (-1, 1 modulation) patterns that were independently formulated, that is, a "two-bit" filter. This concept goes beyond simple high-pass and low-pass filter cutoff concepts in that each filter pixel can be turned off (set to zero modulation) in a pattern designed to optimize the filter's performance. Ternary filter patterns also were investigated incidentally in work previously discussed in the context of BPOF formulation. The Kallman BPOFs[64] included a low frequency box of zero-modulation (DC block) chosen by trial-and-error to optimize filter performance. Casasent and Rozzi[45] investigated SDF filters converted to TPAF modulation with the amplitude pattern determined by thresholding on the magnitude of an SDF filter function. The use of bandpass filtering has been shown to optimize SNR for simple POFs and BPOFs,[21] and in the latter case, a TPAF is created. Here, we are concerned with formulations making more sophisticated use of the amplitude modulation capability to optimize performance, though we note that choice of an optimum processor overall bandpass should always be considered.

The primary benefit expected of TPAF modulation, by virtue of the zero-modulation state added to the BPOF modulation values, is improved discrimination. This is believed to result from selective blocking of spatial frequencies associated with out-of-class patterns. This concept was the basis of the transform-ratio (TR) technique for setting the zero-modulation state (or equivalently, setting the binary amplitude portion).[70] The basic idea of the TR method is to form, in some appropriate manner, representative in-class and out-of-class power spectral density patterns. The two patterns are compared, point-by-point, and the binary-amplitude pattern is established by thresholding on the ratio of in-class to out-of-class values. It is of practical importance to treat the TR threshold as a design parameter of the filter, varying it to search for optimum performance. The phase-only portion of the filter can be set according to any method normally used for BPOF formulation, although it would seem desirable to address distortion invariance (widened in-class response) with the binary phase pattern, since the primary purpose of the binary amplitude pattern is to improve discrimination. Simulations have demonstrated the expected discrimination improvements (relative to a simple BPOF) against both noise and other objects.[70]

The filter-TPAF (fTPAF) formulation[69] combines the TR method of setting a binary-amplitude pattern

with the iterative method used in the fSDF[65] to form a binary-phase pattern, resulting in a smart TPAF which has controlled distortion-invariance and discrimination. Simulations addressing a multi-font character recognition problem have demonstrated the superior classification performance of this type filter compared with converted SDF filters and the fSDF-BPOF formulation,[71] as illustrated in Figure 11. Figure 11(a) shows the various in-class and out-of-class characters used in the study. Figure 11(b) and (c) are bar plots showing the improved discrimination performance of the fTPAF over the binarized projection SDF filter.

The fTPAF formulation evaluated filter performance over the entire correlation output plane after each iteration of filter weights, a significant departure from the reported fSDF iteration procedure. Thus, the filters controlled performance over the entire correlation plane rather than only the on-axis responses. This was critical to success for this application.[71] As for the fSDF, the ad hoc fTPAF iterative formulation procedure does not always converge.

TPAFs have exhibited low correlation efficiencies relative to POFs and BPOFs in reported simulations,[69, 70, 71] this is not surprising since they typically block up to 90 percent of the filter plane with zero-modulation states. In the worst cases, the degradation factor was about two orders of magnitude. For proper perspective, it must be remembered that POFs and BPOFs typically exhibit correlation efficiencies 20-50 times that of classical matched filters. Thus, even the least efficient TPAFs may have efficiencies within an order of magnitude of a classical matched filter. (This discussion has ignored practical efficiency factors of the optical implementation and filter medium, which are independent of filter formulation.)

4.4 Summary

Phase-only and discrete-valued filters are a relatively recent development in optical correlation, first reported in 1984. Thus, the development of smart filter techniques specifically optimized for such filter encodings is at an early stage. However, the experimental testing of such smart filters will be accelerated by the availability of SLMs suitable for their implementation, since this practical factor was one motivation for initially considering these filter types.

The current selection of smart filter formulations for limited-range encoding is not as diverse as for continuous filter functions, but research is proceeding and the filters already developed have shown good promise, as discussed.

Simple forms of BPOF smart filter already have been implemented experimentally, demonstrating good performance and agreement with computer simulations.[72, 73]

TPAFs having non-trivial zero-modulation patterns (that is, other than simple low- and high-pass filters) have not been implemented experimentally, but this situation is likely to change in the near future, since it recently has been shown that TPAF modulation can be achieved with relative ease in magneto-optic spatial light modulators.[73]

POF and discrete-valued smart filters are attractive due to their potential for near-term real-time implementation with available SLMs. A definitive assessment of the relative performance of these filter types is not available, but the reported results are promising and investigation of them is continuing.

5. REAL-TIME IMPLEMENTATION

This action will review experimental realizations of real-time Fourier optical correlation, limited by definition to those architectures providing rapid changes of *both* input and reference patterns. Spatial light modula-

E

NEW YORK (NY-E)

E

MANHATTAN (MN-3)

F

NY-F

E

PIERCE (PI-E)

E

SAN DIEGO (SD-E)

R

NY-R

E

WASHINGTON D.C. (WA-E)

E

RAVENNA (RV-E)

L

NY-L

E

BOSTON (BS-E)

E

SAIGON (SG-E)

P

NY-P

E

BROADWAY (BR-E)

E

COURIER (CD-E)

B

NY-B

E

MEMPHIS (MP-E)

E

DALLAS (DL-E)

8

NY-8

E

STARFLEET (SF-E)

ε

FLORENCE (FL-E)

I

NY-I

E

TIFFANY (TF-E)

T

NY-T

IN-CLASS

OUT-OF-CLASS

Figure 11a. Example of fTPAF character recognition simulation results [Ref. 71]. Characters used in study.

29

(b)

(c)
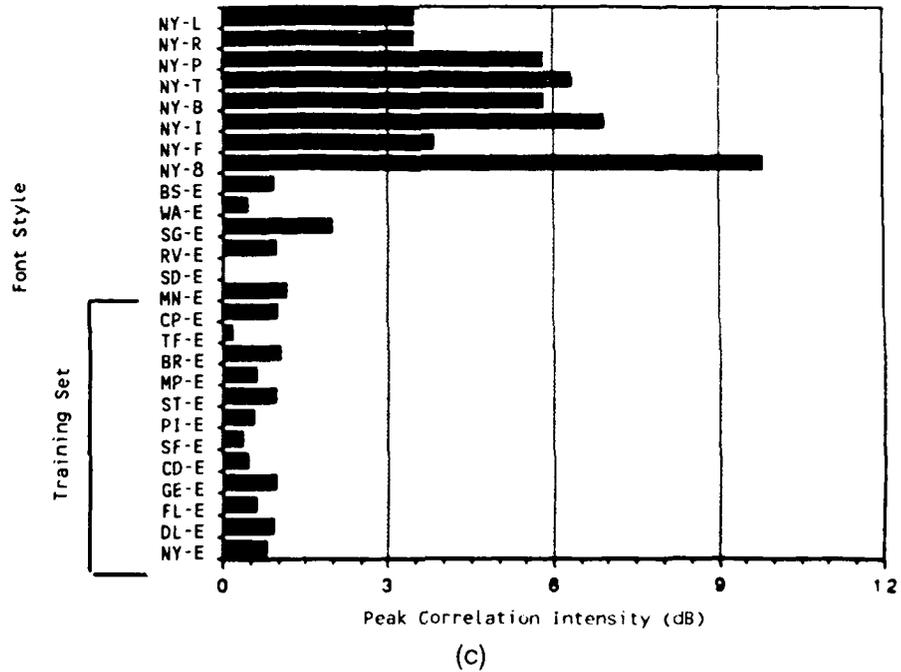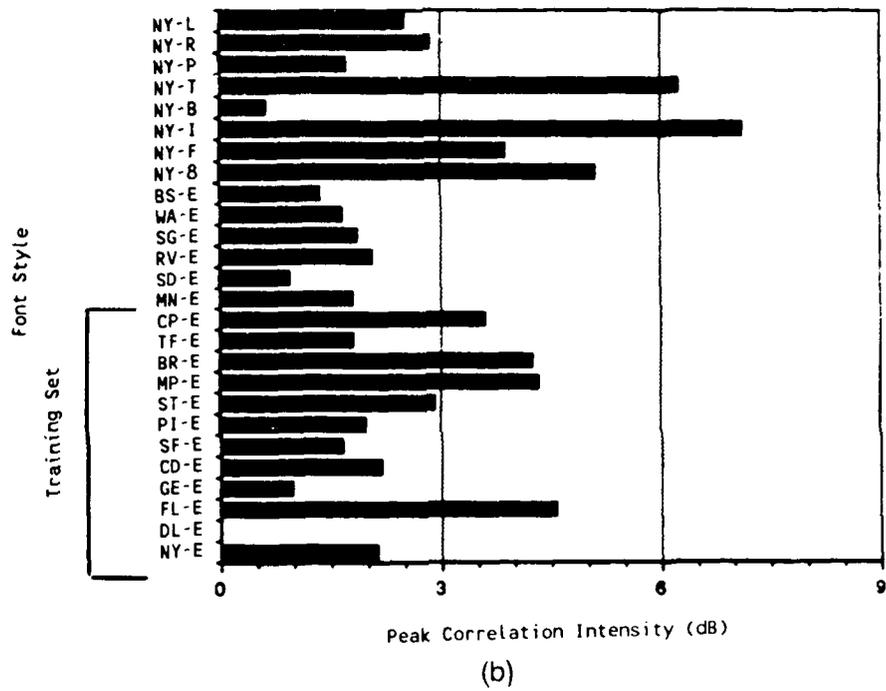
Figure 11. Example of fTPAF character recognition simulation results [Ref. 71]. (b) Discrimination achieved with pSDF filter, (c) Discrimination achieved with fTPAF.

tors (SLM) are the key element in such systems. They will be mentioned where applicable, but not reviewed systematically. The reader is referred to the literature[74-76] and the paper by J. Neff and R. Athale in the October 1989 *Proc. IEEE* [77 (No. 2)] for further information on the technology of these devices. Correlator performance is limited by SLM technology, particularly in the areas of resolution, frame rate, and fidelity of modulation. The correlation architectures now considered attractive for practical application have evolved through advances in both SLM technology and theoretical concepts which capitalize on available SLM performance.

Optical correlators have not achieved everyday use in either commercial or military contexts, although their potential for both areas has been recognized and has motivated much of the research performed. Thus, mainly we will be discussing laboratory experiments, although the work has been carried into the breadboard or brassboard stage in a few instances, which will be noted.

Most real-time correlation experiments fit into the two broad architectural categories already discussed, frequency plane, and joint-transform correlation. Variants involving phase-only, binary, or ternary encodings exist in both categories.

5.1 Real-Time Matched Filtering Experiment

Any type of SLM can be used for the real-time input function, with optically or electrically addressed types being selected as required by the application. Real-time implementation of complex continuous-valued filter functions (for example, matched filters) requires modulation capability (that is, independent magnitude and phase programming) not provided by any existing SLM, unless a holographic (carrier-frequency) encoding is used. The latter, although possible, is expensive in terms of required SLM space-bandwidth, typically exacting a penalty of a four-fold reduction in information capacity to establish the carrier spatial frequency, and an additional factor if magnitude is encoded in other than analog form, as in CGH schemes.

Despite the difficulties mentioned, experiments implementing continuous complex spatial filters in real time have been reported and will be reviewed here.

Four-wave mixing in photorefractive crystal, such as bismuth silicon oxide (BSO), can be used to form correlation filters in real time.[77] Good agreement between digitally calculated correlations and experimental results were reported, and correlation rise times of less than 200 ns were observed with a Q-switched frequency-doubled Nd:YAG source. The architecture is somewhat akin to joint-transform correlation and requires input of the reference pattern in the image (rather than Fourier) domain.

A portable brassboard correlator achieving rapid sequencing between four fixed holographic filters using an array of laser diode sources has been constructed and tested.[78] A Hughes liquid crystal light valve (optically addressed) SLM served as the input device.

5.2 Real-Time Binary Phase-Only and Ternary Phase-Amplitude Correlator

The first real-time correlator implementing BPOF filters, reported in 1986, used 48 x 48 element electrically-addressed magneto-optic SLMs (MOSLM) at both input and filter plane.[79] Nontrivial patterns could be correlated with this limited resolution because no filter encoding capacity was sacrificed to implement a holographic spatial carrier. The processor is an in-line configuration. It cycled new inputs and/or filters at the rate of four per second, limited by the controlling microcomputer (an Apple IIe) rather than the SLMs.

Further experiments with this correlator have demonstrated good correlation performance in character and machine part recognition applications, including successful implementation of simple smart filters.[72] Real-time

31

magneto-optic SLMS correlators using 128 x 128 element devices also have been reported.[73, 80] In one extensive set of field experiments, the correlator was operated at 20 frames per second while recognizing fly-by aircraft images.[80]

Recently, BPOF correlation using the Texas Instruments Deformable Mirror Device (DMD) SLM have been reported.[81] This electrically-addressed device has 128 x 128 element resolution, cycles readily at TV frame rates and higher, and is capable of continuous phase modulation, which was not used in these initial experiments.

Also recently, the magneto-optic SLM devices have been shown to provide easily-accessed ternary (1, 0, -1) modulation levels, and simple TPAF correlations showing enhanced signal-to-noise were demonstrated.[73]

5.3 Real-Time Joint Transform Correlation Experiment

Real-time joint transform correlation also requires real-time SLM devices, with slightly different performance requirements. The input and reference patterns are introduced side-by-side in the input SLM and a buffer zone between them usually must be allowed so the desired correlation component and other terms are separated at the output plane. Thus, the space-bandwidth product (that is, resolution) of the input SLM must be 2 to 3 times that of the images being processed if a single SLM is used. Alternatively, two adjacent SLMs may be used, if practicable. At the Fourier plane, an interference pattern must be square-law detected, modulated onto a coherent beam, and optically Fourier-transformed to produce the correlation pattern (and other components). This requires an optically-addressed SLM, or an imaging photodetector whose output signal drives an electrically addressed SLM. The resolution of this SLM must be 2 to 3 times that of the processed images, but there is no requirement for complex (phase and amplitude) modulation since the signal is square-law detected.

As mentioned, the use of four-wave mixing in nonlinear crystals to form a real-time filter[77] may be viewed as a form of joint transform correlation.

Real-time joint transform correlation has been implemented using the Hughes Liquid Crystal Light Valve SLM as the Fourier plane device and a magneto-optic SLM as the input device.[82]

Figure 12 diagrams an experiment using a liquid-crystal TV (LCTV, the familiar inexpensive consumer appliance) for both input and joint transform modulation, performed sequentially.[83] The adjacent input-reference pattern was first introduced and the resulting joint transform pattern detected by the vidicon and recorded on tape. This signal was replayed into the same liquid-crystal TV and transformed to generate the correlation pattern. Preliminary results indicated successful correlation. The ultimate utility of the liquid-crystal TV for coherent optical processing depends largely on the degree to which phase distortions normally exhibited can be reduced in improved devices or ameliorated by other techniques.[84]

5.4 Binary Joint Transform Correlator

Reports of the experimental demonstration of binary joint transform correlation are likely in the near future since two of the leading SLM types, the magneto-optic SLMs and the deformable mirror device, appear to be good candidates for implementing binary joint transform modulation.

5.5 Real-Time Fourier-Feature Processor

A Fourier-feature processor derives its output based on features consisting of spectral components of the spatial Fourier transform of the input pattern, and thus, does not strictly fall within the scope defined for this
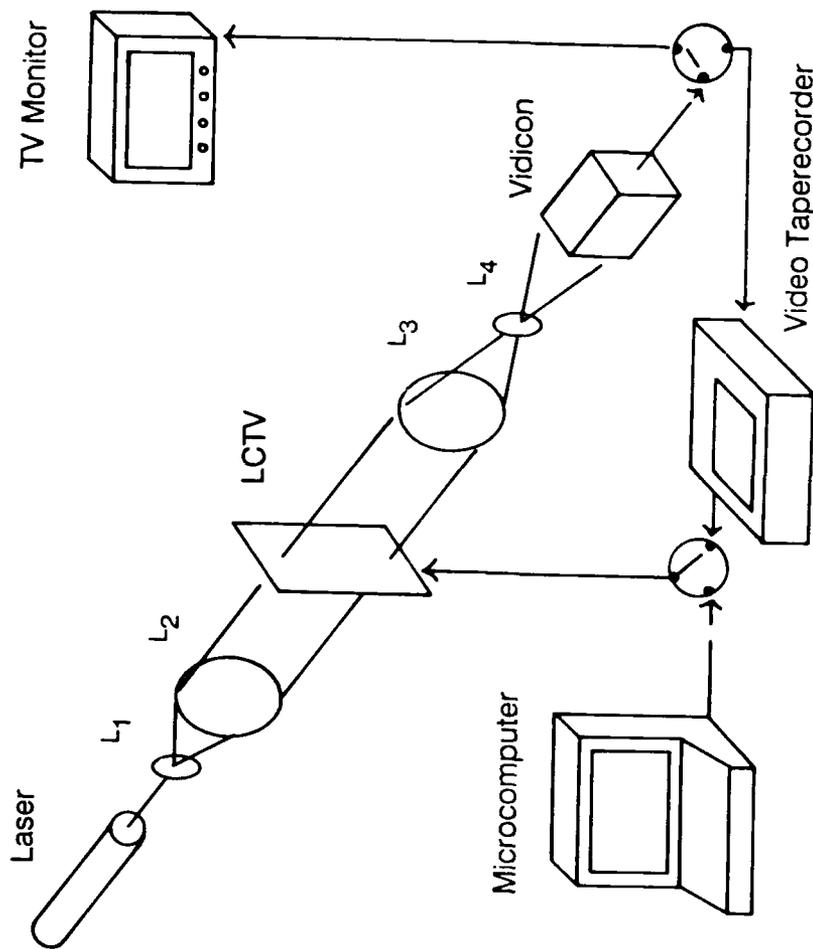
Figure 12. Diagram of real-time joint-transform correlator based on a liquid crystal TV SLM [Ref. 83].

section. Nevertheless, the process is closely related to correlation and a notable effort has been underway to develop real-time commercial systems for machine vision inspection.[85] Global Holonetics Corporation, Fairfield, Iowa, is developing such systems using ring-wedge detectors to read out radial and azimuthal Fourier features which are then processed (non-optically) using appropriate discrimination techniques to derive the desired classification between acceptable or reject items, for example, a misprinted or misaligned label on a container.

6. CONCLUSION

Hardware and analytical advances in the last five years have led to a rapid evolution of real-time correlation architectures having high potential for practical application to military and commercial pattern recognition problems. With the potential to perform correlation of reasonably high resolution (for example, 512 x 512 pixels) images at rates approaching 1,000 frames per second (for both input and reference patterns), in packages well under a cubic foot in volume, there is little doubt that this technology can provide performance-to-size ratios orders of magnitude above any other approach, now or in the foreseeable future.

The perception of this potential is leading to considerable research on both the fundamentals and application of correlation to problems such as automatic target recognition, signal processing, quality control and inspection, and robotic vision.

As reviewed here, a variety of real-time correlation architectures and filter formulations are being advanced to attack the various applications, and this is appropriate because the different approaches have complementary strengths so that no single technique is likely to prevail in all applications. A generic hybrid (optical/electronic) configuration in which optics does the computing and electronics performs filter storage, control, and post-processing, is emerging as a paradigm for the optimum configuration of real-time correlation systems realizable in the near future.

References

[1] Vander Lugt, A. (1964) Signal Detection By Complex Filtering, *IEEE Trans.* Information Theory, IT-10 (No. 2):139

[2] Brown, G.R., and Lohmann, A.W. (1966) Complex Spatial Filtering with Binary Masks, *Appl. Opt.*, 6:967-969.

[3] Lohmann, A.W., and Paris, D.P. (1967) Binary Fraunhofer Holograms Generated by Computer, *Appl. Opt.*, 6:1739-1748.

[4] Weaver, C.S., and Goodman, J.W. (1966) Technique for Optically Convolving Two Functions, *Appl. Opt.*, 5:1248-1249.

[5] Gregory, D.A. (1986) Real-Time Pattern Recognition Using a Modified LCTV in a Coherent Optical Correlator. *Appl. Opt.*, 25:467-468.

[6] Gregory, D. (1988) *Time Multiplexed Miniature Optical Correlator*, Report RD-RE-88-02, U.S. Army Missile Command.

[7] Javidi, B., and Horner, J.L. (1989) Single SLM Joint Transform Correlator, *Appl. Opt.*, 28:1027-1032.

[8] Javidi, B. Nonlinear Optical Correlation, submitted to *Applied Optics*.

[9] Oppenheim, A., and Lim, J. (1981) The Importance of Phase in Signals, *Proc. IEEE*, 69:529-541.

[10] Caulfield, H.J. (1982) Role of Horner Efficiency in the Optimization of Spatial Filters for Pattern Recognition, *Appl. Opt.*, 24:4391-4392.

[11] Horner, J.L. (1982) Light Utilization in Optical Correlators, *Appl. Opt.*, 21:4511-4514.

[12] Horner, J.L., and Gianino, P.D. (1984) Phase-Only Matched Filtering, *Appl. Opt.* 23:812-816.

[13] Kozma, A., and Kelly, D. (1965) Spatial Filtering for Detection of Signals Submerged in Noise, *Appl. Opt.*, 4:387-392.

[14] Psaltis, D., Pack, E., and Venkatesh, S. (1984) Optical Image Correlation with Binary Spatial Light Modulator, *Optical Eng.*, pp. 698-704.

[15] Horner, J.L., and Leger, J.R. (1985) Pattern Recognition with Binary Phase-Only Filters, *Appl. Opt.*, 24:609-611.

[16] Flavin, M., and Horner, J.L. (1989) Correlation Experiments with a Binary Phase-Only Filter, *Optical Eng.*, 28:470-473.

[17] Zhang, G., Su, X., and Guo, L. (1987) Binary Synthetic Phase-Only Filter with Random Mask, *S.P.I.E. Proc.*, 812:108-112.

[18] Kallman, R.R. (1986) Optimal Low Noise Phase-Only and Binary Phase-Only Optical Correlation Filters for Threshold Detectors, *Appl. Opt.*, 25.4216-4217.

[19] Cottrell, D., Lilly, R., and Davis, T. Day. (1987) Optical Correlator Performance of Binary Phase-Only Filters Using Fourier and Hartley Transforms, *Appl. Opt.*, 26:3755-3761.

[20] Flannery, D., Loomis, J., and Milkovich, M. (1988) Design elements of binary phase-only correlation filters, *Appl. Opt.*, 27:4231-4235.

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

[22] Kumar, B., and Bahri, Z. Efficient algorithm for designing optimal binary phase-only filters, *Appl. Opt.* review.

[23] Dickey, F.M., and Romero, L.A. (1989) Dual Optimality of Phase-Only Filter, *Opt. Lett.*, 14:4-5).

[24] Horner, J., and Gianino, P. (1985) Applying the phase-only filter concept to the synthetic discriminant function correlation filter, *Appl. Opt.*, 24:851-855.

[25] Dickey, F., Stalker, K, and Mason, J. (1988) Bandwidth considerations for binary phase-only filters, *Appl. Opt.*, 27:3811-3818.

[26] Horner, J.L., and Fielding, K. Performance of Phase-Only Filters with Cluttered Input Imagery, submitted to Applied Optics.

[27] Flavin, M.A., and Horner, J.L. Amplitude Encoded Phase-Only Filters, submitted to Applied Optics.

[28] Goodman, J.W. (1968) Introduction to Fourier Optics, New York: McGraw-Hill, ch. 7, pp. 183-184.

[29] Harris, F.J. (1978) On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform, *Proc. IEEE*, 66:51-83.

[30] Kallman, R.R. (1986) The construction of low noise optical correlation filters, *Appl. Opt.*, 25:1032-1033.

[31] Caulfield, H., and Maloney, W. (1969) Improved Discrimination in Optical Character Recognition, *Appl. Opt.*, 8:2354-2356.

[32] Kumar, B., and Pochapsky, E. (1986) Signal-to-noise ratio considerations in modified matched spatial filters, *J. Opt. Am. A*, 3:777-786.

[33] Hester, C.F., and Casasent, D. (1980) Multivariant technique for multi-class pattern recognition, *Appl. Opt.*, **19**:1758-1761.

[34] Casasent, D. (1984) Unified synthetic discriminant function computation formulation, *Appl. Opt.*, **23**1620-1627.

[35] Sharma, V., and Casasent, D. (1984) Optimal Linear Discriminant Functions, *Proc. SPIE*, **519**:50-55.

[36] Bahri, Z., and Kumar, B. (1988) Generalized synthetic discriminant functions, *J. Opt. Soc. Am. A*, **5**:562-571.

[37] Caulfield, H., and Weinberg, M. (1982) Computer recognition of 2-D patterns using generalized matched filters, *Appl. Opt.*, **21**:1699-1704.

[38] Duda, R., and Hart, P. (1973) Pattern Classification and Scene Analysis, New York: Wiley.

[39] Casasent, D., and Chang, W. (1986) Correlation synthetic discriminant functions, *App. Opt.*, **25**:2343-2350.

[40] Mahalanobis, A., Kumar, B., and Casasent, D. (1987) Minimum average correlation energy filters, *App. Opt.*, **25**:3633-3640.

[41] Kumar, B. (1986) Minimum-variance synthetic discriminant functions, *J. Opt. Soc. Am. A*, **3**:1579-1584.

[42] Bahri, Z., and Kumar, B. (1988) Generalized synthetic discriminant functions, *J. Opt. Soc. Am. A*, **5**:562-571.

[43] Casasent, D., and Sharma, V. (1983) Performance of Synthetic Discriminant Functions for Infrared Ship Classification, *Proc. SPIE*, **422**:193-196.

[44] Riggins, J., and Butler, S. (1984) Simulation of synthetic discriminant function optical implementation, *Optical Eng.*, **23**:721-626.

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

[46] Casasent, D., and Mahalamobis, A. (1987) Optical iconic filters for large class recognition, *Appl. Opt.*, **26**:2266-2273.

[47] Casasent, D., and Ye, R. (1988) Optical Laboratory Realization of Distortion Invariant Filters, *Proc. SPIE*, **939**:105-120.

[48] Kallman, R. (1986) Construction of low noise correlation filters, *Proc. SPIE*, **638**:214-217.

[49] Kirkpatrick, S., Gelatt, C., and Vecchi, M. (1983) Optimization by Simulated Annealing, *Science*, **220**:671-680.

[50] Kim, M., Feldman, M., and Guest, C. (1989) Optimization of Binary Phase Only Filter with Simulated Annealing Algorithm, Paper TUI-31 in *Optical Computing*, Optical Society of America 1989 Technical Digest Series, **9**:342-345.

[51] Hsu, Y., and Arsenault, H. (1982) Optical pattern recognition using circular harmonic expansion, *Appl. Opt.*, **21**:4016-4019.

[52] Hsu, Y., Arsenault, H., and April, G. (1982) Rotation-invariant digital pattern recognition using circular harmonic expansion, *Appl. Opt.*, **21**:4012-4015.

[53] Hsu Y., and Arsenault, H. (1984) Pattern discrimination by multiple circular harmonic components, *Appl. Opt.*, **23**:841-844.

[54] Wu, R., and Stark, H. (1984) Rotation-invariant pattern recognition using a vector reference, *Appl. Opt.*, **23**:838-840.

[55] Arsenault, H., and Sheng, Y. (1986) Properties of the circular harmonic expansion for rotation-invariant pattern recognition, *Appl. Opt.*, **24**:3225-3229.

[56] Mendlovic, A., Marom, E., and Konforti, N. (1988) Shift and Scale Invariant Pattern Recognition Using Mellin Radial Harmonics, *Optics Communications*, **67**:172-176.

[57] Mahalanobis, A., Kumar, B., and Casasent, D. (1986) Spatial-temporal correlation filter for in-plane distortion invariance, *Appl. Opt.*, **25**:4466-4472.

[58] Casasent, D., and Psaltis, D. (1976) Position, rotation, and scale invariant optical correlation, *Appl. Opt.*, **15**:1795-1799.

[59] Kumar, B., Lee, A., and Connelly, J. (1988) Correlation filters for orientation estimation, *Proc. SPIE*, **938**:190-197.

[60] Schils, G., and Sweeney, D. (1986) Iterative technique for the synthesis of optical-correlation filters, *J. Opt. Soc. Am. A*, **3**:1433-1442.

[61] Sweeney, D., Ochoa, E., and Schils, G. (1987) Experimental use of iteratively designed rotation invariant filters, *Appl. Opt.*, **26**:3458-3465.

[62] Schils, G., and Sweeney, D. (1988) Optical processor for recognition of three-dimensional targets viewed from any direction, *J. Opt. Soc. Am. A*, **5**:1309-1321.

[63] Gaynor, E., and Caulfield, H. (1984) Orientation Variability in Generalized Matched Filters, *Proc. SPIE*, **519**:56-61.

[64] Kallman, R. (1987) Direct construction of phase-only correlation filters, *Proc. SPIE*, **827**:184-190.

[65] Jared, D., and Ennis, D. (1989) Inclusion of filter modulation in synthetic-discriminant function filters, *Appl. Opt.*, **28**:232-239.

[66] Jared, D., Ennis, D., and Dreskin, S. (1988) Evaluation of binary phase-only filters for distortion-invariant pattern recognition, *Proc. SPIE*, **884**:139-145.

[67] Arsenault, H., Leclerc, L., and Sheng, Y. (1988) Similarity and invariance in pattern recognition, *Proc. SPIE*, **960**:2-17.

[68] Rosen, J., and Shamir, J. (1988) Circular harmonic phase filters for efficient rotation-invariant pattern recognition, *Appl. Opt.*, **27**:2895-2899.

[69] Flannery, D., Loomis, J., and Milkovich, M. (1988) New formulations for discrete-valued correlation filters, *Proc. SPIE*, **938**:206-211.

[70] Flannery, D., Loomis, J., and Milkovich, M. (1988) Transform-ratio ternary phase-amplitude filter formulation for improved correlation discrimination, *Appl. Opt.*, 27:4079-4083.

[71] Milkovich, M., Flannery, D., and Loomis, J. (1989) A Study of Transform-ratio Ternary Phase-Amplitude Filter Formulations for Character Recognition, *Optical Engineering*, 28:487-493.

[72] Flannery, D., Loomis, J., Milkovich, M., and Keller, P. (1988) Application of binary phase-only correlation to machine vision, *Optical Engineering*, 27:309-320.

[73] Kast, B., Giles, M., Lindell, S., and Flannery, D. (1989) Implementation of Ternary Phase-amplitude Filters Using a Magneto-optic Spatial Light Modulator, *Appl. Opt.*, 28:1044-1046.

[74] Spatial Light Modulators and Applications, (1988) Technical Digest Series Volume 8, Summaries of papers presented at the Spatial Light Modulators and Applications topical Meeting, June 15-17, 1988, South Lake Tahoe, Nevada, Optical Society of America, Washington, D.C.

[75] Warde, C., and Fisher, A. (1987) Spatial Light Modulators: Applications and Functional Capabilities, in OPTICAL SIGNAL PROCESSING, J.L. Horner, Ed., Sand Diego, CA, Academic Press, ch. 7.2, pp. 476-523.

[76] Latest Applied Optics on SLMs (U. Efron editor, spring 1989).

[77] Nicholson, M., Cooper, I., Gibbons, G., and Petts, C. (1987) Optimization of an updatable optical image correlator, *Optical Engineering*, 26:445-452.

[78] Upatnieks, J. (1983) Portable real-time coherent optical correlator, *Appl. Opt.*, 22:2798-2803.

[79] Flannery, D., Biernacki, A., Loomis, J., and Cartwright, S. (1986) Real-time Coherent Correlator Using Binary Magnetooptic Spatial Light Modulators at Input and Fourier Planes, *Appl. Opt.*, 25:466.

[80] Mills, S., and Ross, W. (1987) Dynamic magneto-optic correlator: real-time operation, *Proc. SPIE*, 753:54-63.

[81] Florence, J., and Gale, R. (1988) Coherent optical correlation using a deformable mirror device spatial light modulator in the Fourier plane, *Appl. Opt.*, 27:2091-2093.

[82] Yu, F., and Lu, X. (1984) A Real-Time Programmable Joint Transform Correlator, *Optics Communications*, 52:10.

[83] Yu, F., Jutamulia, S., Lin, T., and Gregory, D. (1987) Adaptive real-time pattern recognition using a liquid crystal TV based joint transform correlator, *Appl. Opt.*, 26:1370-1372.

[84] Horner, J. (1988) Is phase correction required in SLM-based optical correlators?, *Appl. Opt.*, 27:436-438.

[85] Clark, D. An Optical Feature Extractor for Machine Vision, *Proc. VISION '87 Conference*, pp. 7-23 - 7-49, society of Manufacturing Engineers, Dearborn, Michigan.

MISSION
of
Rome Air Development Center

RADC plans and executes research, development, test and selected acquisition programs in support of Command, Control, Communications and Intelligence (C^3I) activities. 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. The areas of technical competence include communications, command and control, battle management, information processing, surveillance sensors, intelligence data collection and handling, solid state sciences, electromagnetics, and propagation, and electronic, maintainability, and compatibility.