Robust Classification of Blurred Imagery

Deepa Kundur, *Member, IEEE*, Dimitrios Hatzinakos, *Senior Member, IEEE*, and Henry Leung, *Member, IEEE*

***Abstract*—In this paper, we present two novel approaches for the classification of blurry images. It is assumed that the blur is linear and space invariant, but that the exact blurring function is unknown. The proposed fusion-based approaches attempt to perform the simultaneous tasks of blind image restoration and classification. We call such a problem *blind image fusion*. The techniques are implemented using the nonnegativity and support constraints recursive inverse filtering (NAS-RIF) algorithm for blind image restoration and the Markov random field (MRF)-based fusion method for classification by Schistad-Solberg *et al.*. Simulation results on synthetic and real photographic data demonstrate the potential of the approaches. The algorithms are compared with one another and to situations in which blind blur removal is not attempted.**

***Index Terms*—Blind image restoration, classification, multispectral image fusion.**

I. Introduction

THE inaccuracy of many image classification strategies often results from attempting to fuse data that exhibits motion-induced blurring or defocusing effects. Compensation for such blurring is inherently sensor-dependent and is nontrivial as the exact blur is often time-varying and unknown [1]. In such a situation, restoration is frequently performed on the degraded data prior to classification. A major obstacle arises when the exact blurring function is unknown. This may be overcome by using blind image restoration algorithms. *Blind image restoration* refers to the dual process of blur identification and image restoration and is used only when partial information about the image degradation process is available. It is suited for situations in which the blurring and noise characteristics of the imaging system may be unknown due to their time-varying nature. A recent survey of such algorithms can be found in [2].

The weakness of most blind restoration techniques stems from numerical instability; there is often insufficient information to provide an accurate regularized solution to the problem and noise amplification may result [3]. In addition, it may be impractical for a human observer to supervise the restorations to prevent instability. This may result in an unreliable image estimate, which in turn can effect the accuracy of subsequent higher-level processing tasks such as classification.

In this paper, we consider the classification of blurred and noisy images using data fusion strategies. It is important that we distinguish between our use of the terms "fusion" and "classification."*Multisensor data fusion* refers to the acquisition, processing, and synergistic combination of information from various knowledge sources to provide a better understanding of the situation under consideration [4]. *Classification* is an information processing task in which specific entities are mapped to general categories. For the classification of multispectral images, the specific goal is to assign each vector-valued pixel of the combined multispectral image to its appropriate category using tonal and/or textural data; the result is a single color-coded image showing the several types of classes in the scene. In this paper, image classification is the specific goal we wish to achieve and data fusion is the process by which we accomplish the task.

The accuracy of image classification is often highly dependent on the quality of the restored images[1]. We argue that performing blind image restoration prior to and separately from classification results in suboptimal solutions. We attempt to combine both processes using multisensor data fusion strategies to produce a more regularized solution; we call this process *blind image fusion* for classification. The problem has applications to many areas of image analysis, such as robot vision, remote sensing for land-use classification and medical imaging, among others. The imagery may be multisensor/multispectral in nature or can consist of a single sensor image band.

The main contributions of this paper are as follows.

1) The integration of blind image restoration with traditional data fusion concepts to provide a more optimal image classification. Two novel approaches based on different philosophies are proposed to address the problem of blind image fusion.
   a) In our first technique we take a new perspective on the image estimates at each iteration of a recursive restoration algorithm. Each estimate is treated as a different reading of an image sensor and the estimates are simultaneously fused to produce a more regularized classified image.
   b) In the second approach we define an error metric which uses the results from intermediate-level image fusion and classification to produce an appropriate[2] termination point for the restoration algorithm.
2) The implementation of the two architectures to assess the practical feasibility of blind image fusion. The authors are unaware of other advanced research into combining

Manuscript received May 4, 1998; revised April 28, 1999. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Scott T. Acton.

D. Kundur and D. Hatzinakos are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, Ont., Canada, M5S 3G4 (e-mail: dimitris@comm.toronto.edu).

H. Leung is with the Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alta., Canada, T2N 1N4.

Publisher Item Identifier S 1057-7149(00)01158-1.

[1]This is especially true if the classification technique is trained on or uses models fof nonblurred imagery.

[2]By the term *appropriate* we mean that the image estimate at termination results in an accurate classification. This termination point is not always the best image estimate visually or in the mean square sense.

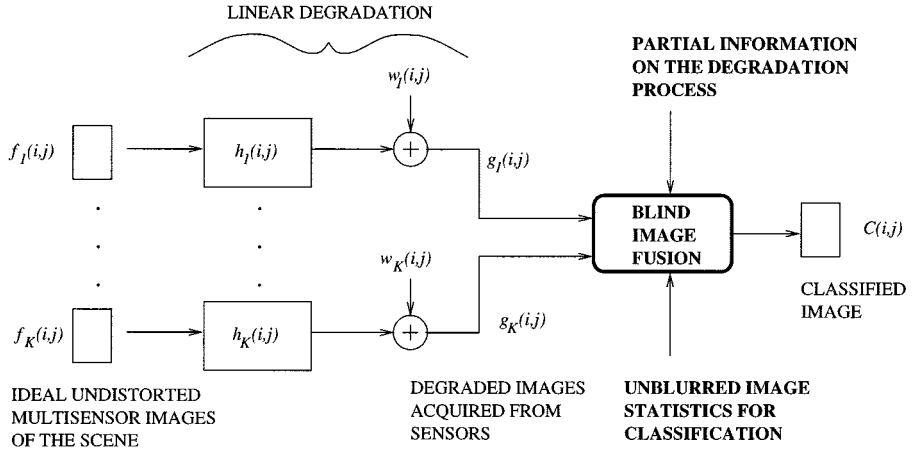

Fig. 1. The blind image restoration and classification problem. The sensor images are assumed to be degraded according to the linear degradation model. These images are used to produce a classification of the scene using partial information about the undistorted imaged scene, and the blurring process.

restoration and fusion and assess the potential of the technology.

3) A comparative study of the effects of blur removal for image classification.

In the next section we define the specific problem to address and discuss our assumptions. In Section III we propose the two novel approaches for blind image fusion. Implementation issues are discussed in Section IV. Simulation results and comparisons with existing techniques are given in Section V, and final remarks are conveyed in Section VI.

Assumption Set 1 (General Assumptions for the Blind Image Fusion Problem)

1) Multisensor images of the same scene are registered. If the images are, for example, the different color bands of a color image, then it is safe to assume that the images are registered. Otherwise, a registration algorithm must be applied to the images prior to processing.
2) The blurring process obeys the linear degradation model of (1). This model is successfully used in imaging applications in which the blurring process is *isoplantic* (i.e., the blurring operation is linear and shift invariant). The model makes the solution to the image restoration problem tractable [10].
3) Partial information about the imaged scene is available to perform blind image restoration. This information is specific to the algorithm employed. Section IV-A1 provides the specific assumptions for the NAS-RIF algorithm.
4) Statistical information about the undistorted images of the scene is available for classification. This information is also specific to the particular classification approach implemented. Section IV-A2 provides the specifics for the MRF classification algorithm.

II. Problem Formulation

We address the issue of the robust classification of blurred imagery. The difference between the problem we consider in this paper and that addressed in other literature on image classification [5]–[9] is that we explicitly account for the blurring degradation in the image while other approaches do not. We focus on the design of techniques for the simultaneous restoration and classification of blurred imagery using fusion-based strategies.

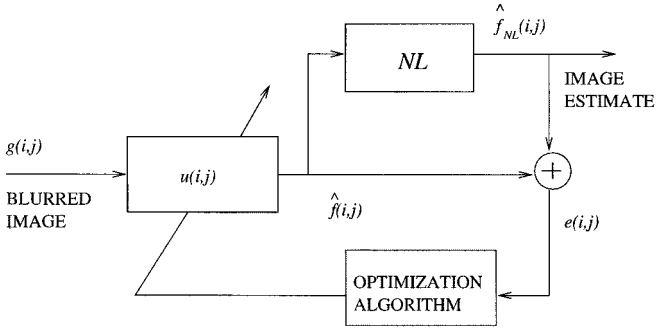

Fig. 2. NAS-RIF algorithm for blind image deconvolution.

We consider the scenario in which we have K different degraded registered sensor images $g_1, g_2, \cdots, g_k$, where $K \geq 1$, of a given scene. Each acquired image g_k is assumed to be degraded according to the following linear degradation model.

$$g_k(i,j) = f_k(i,j) * h_k(i,j) + w_k(i,j) \tag{1}$$

where $g_k(i,j)$ is the blurred noisy image from the kth sensor, $f_k(i,j)$ is the undistorted true image of the scene, $h_k(i,j)$ is the blurring function also known as the point-spread function (PSF), $w_k(i,j)$ is zero-mean additive white Gaussian noise, and $*$ represents the two-dimensional (2-D) linear convolution operator.[3] If we neglect the noise term in (1), the process of recovering $f_k(i,j)$ from $g_k(i,j)$ is called *blind image deconvolution*. Our goal is to effectively combine the information in these degraded images to provide an accurate classification of the scene C. Any training or reference data used for classification is assumed to apply to the undistorted images f_k. Fig. 1 gives an overview of the problem.

[3]The PSF models imperfection in the system such as motion blurring from the movement of satellites or from the use of out-of-focus cameras; the additive noise term models nonidealities in sensor equipment and electronic instrumentation [10].

TABLE I
STATISTICAL MRF FUSION METHOD USING THE ICM ALGORITHM

1. Initialization:

   For each pixel (i,j),

   Perform a search to choose the class $C(i,j)$ that minimizes the sensor-specific image statistic function $\sum_{s=1}^{K} \alpha_s U_{data}(\mathbf{g}_s, C(i,j))$ where U_{data} is given by Equation (9).

2. For $k = 1, 2, \ldots, N_{ICM}$

   For each pixel (i,j),

   Perform a search to choose the class $C(i,j)$ that minimizes the following equation:

$$U(\mathbf{g}_1, \mathbf{g}_2, \ldots, \mathbf{g}_K, C(i,j)) = \sum_{s=1}^{K} \alpha_s U_{data}(\mathbf{g}_s, C(i,j)) + \hat{U}_{sp}(C(i,j), \hat{\mathbf{C}}_{k-1})$$

   where U_{data} is given by Equation (9), and $\hat{U}_{sp}$ is an approximation of U_{sp} such that the classification of pixel (i,j) depends on the square neighbourhood $G_{(i,j)}$ centered about (i,j) of the previous classification estimate $\hat{\mathbf{C}}_{k-1}$. That is,

$$\hat{U}_{sp}(C(i,j), \hat{\mathbf{C}}_{k-1}) = \beta_{sp} \sum_{(k,l) \in G_{(i,j)}} I(C(i,j), \hat{C}_{k-1}(k,l)), \tag{19}$$

   where $G_{(i,j)}$ is the set of all pixels adjacent and diagonal to (i,j), $I(\cdot,\cdot)$ is given by Equation (11), and β_{sp} is a user-specified scalar parameter as specified in Equation (10). N_{ICM} is the number of iterations in the ICM algorithm [16]. The use of $\hat{U}_{sp}$ term makes estimating the global minimum of Equation (8) less computationally complex, but may result in a suboptimal solution. This is the trade-off proposed by the ICM algorithm.

We list the general assumptions made in our problem formulation below. More specific assumptions are made during implementation as specific algorithms are employed for blind image fusion. These assumptions are discussed in Sections IV-A1 and IV-A2.

III. BLIND IMAGE FUSION

The two major considerations in the design of our blind image fusion scenarios was the portability of the techniques to different applications and the ease of implementation. As a result, we have designed our methods such that they can incorporate existing iterative blind image restoration and classification algorithms. This allows the flexibility to select the algorithms most appropriate for a given application. In addition, the use of well-known and well-documented methods makes implementation and testing easier, and allows a user to predict and assess the behavior of the overall blind fusion scheme to a greater degree.

The selection of the data fusion approach is based on existing research demonstrating the success of fusion in effectively combining complementary and redundant information (see [1], [4] and references therein). In addition, the authors' previous research into data fusion suggests its promise for the processing of data exhibiting signal-dependent noise [11]. In the next section we show how a restored image can be considered to suffer from signal-dependent noise. In Section III-B we discuss the general scenarios.

A. Blurring as Signal-Dependent Noise

In this section, we demonstrate how a restored image estimate is considered to suffer from both signal-dependent and -independent additive noise processes. We assume that image restoration is performed by the process of inverse filtering (i.e., the filtering of $g(i,j)$ to produce an estimate of $f(i,j)$).[4] Assuming the linear degradation model of (1), the restored image $\hat{f}(i,j)$ is given by

$$\hat{f}(i,j) = g(i,j) * u(i,j) \tag{2}$$

$$= f(i,j) * h(i,j) * u(i,j) + w(i,j) * u(i,j) \tag{3}$$

$$= f(i,j) + [h(i,j) * u(i,j) - \delta(i,j)] * f(i,j) + w(i,j) * u(i,j) \tag{4}$$

$$= f(i,j) + \hat{s}(i,j) + \hat{w}(i,j) \tag{5}$$

where $\delta(i,j)$ is the Kronecker delta function, $\hat{s}(i,j) = [h(i,j) * u(i,j) - \delta(i,j)] * f(i,j)$ and $\hat{w}(i,j) = w(i,j) * u(i,j)$. As we can see, the restored image estimate suffers from two errors: $\hat{s}(i,j)$ which is a signal-dependent noise process (i.e., dependent on $f(i,j)$), and $\hat{w}(i,j)$ which is a signal-independent

[4]Inverse filtering is a common approach to perform blind deconvolution [12].

random process. The $\hat{s}(i,j)$ term represents the degree of partial blurring that remains in the restored result and $\hat{w}(i,j)$ is the filtered additive noise term.

The goal of blind image restoration is to find a filter $u(i,j)$ which minimizes the overall restoration error $\hat{s}(i,j)+\hat{w}(i,j)$. Ideally, if the blurred image does not contain additive noise (i.e., $w(i,j)=0$ for all (i,j)), then $u(i,j)=h^{-1}(i,j)$, where $h^{-1}(i,j)$ is the inverse of the PSF,[5] minimizes the restoration error such that $\hat{s}(i,j)=0$ and $\hat{f}(i,j)=f(i,j)$. In practice, however, $w(i,j)$ is not negligible and $h(i,j)$ is low pass which suggests that $h^{-1}(i,j)$ is high pass. As a result, inverse filtering will amplify the noise process so that $\hat{w}(i,j)$ becomes significant.

In fact, it is well-known that a tradeoff exists between the degree of deblurring and the extent of noise amplification [13]. Many classification methods perform poorly as a result of partial blurring or excessive noise amplification in the restored result. Regularization techniques can be employed to find an appropriate compromise. The question arises as to how to find the degree of regularization which will produce the best classified result.

B. The Fusion-Based Strategies

We investigate two general architectures for the robust classification of blurred images which exploit the concepts of image fusion. The two approaches are as follows.

1) Our first approach involves exploiting the complementarity of distinctly regularized restorations. We treat the restored images as the outputs of different bands of an image sensor; the bands are assumed to experience additive and correlated signal-dependent noise. By fusing the images in the classification stage we can attempt to extract the salient features from each of the different restorations to produce a more regularized overall result. The fusion process is applied to the selected image estimates of a given degraded image band as well as any other available imagery from other sensors.

   The approach provides the user with the flexibility to choose the restorations to fuse for classification. The major disadvantage is that the image selection process is ad hoc.

2) Based on the limitation of our first approach, we investigate a technique to automate the process of choosing a good restoration for classification. The selection technique is not straight-forward as the visually most appealing restoration does not necessarily produce the most accurate classification. We base our approach on the premise that feedback from the classification stage to enhance the iterative restoration process can result in a more optimal overall classification. We devise a technique (which uses a quantitative figure of merit) to assess the restoration that produces the most accurate classification. Image fusion is applied at the classification stage which incorporates information from an intermediate image estimate as well as other available sensor images of the scene. This information that is fed back to the restoration stage to enhance the image estimate.

   The advantage of this technique is that the process of selecting a good image estimate for classification is algorithmic. However, as the classification result makes use of only one image estimate, the solution is not as regularized as in our first approach.

[5]If the PSF is noninvertible then we may consider $h^{-1}(i,j)$ to be the *pseudo-inverse* [3].

IV. Implementation of the Architectures

A. Restoration and Classification Algorithms

In each of the approaches proposed the blind image restoration and classification stages are distinct, but interconnected. For the implementation of the techniques we make use of the NAS-RIF algorithm and the MRF classification technique by Schistad-Solberg *et al.*.

1) The NAS-RIF Algorithm for Blind Image Restoration: The nonnegativity and support constraints recursive inverse filtering (NAS-RIF) algorithm [14] is applicable to situations in which an object of finite support is imaged against a uniform background. Fig. 2 gives an overview of the algorithm. The blurred image pixels $g(i,j)$ are input to a 2-D variable coefficient finite impulse response (FIR) filter $u(i,j)$ whose output represents an estimate of the true image denoted $\hat{f}(i,j)$. This estimate is then mapped to the set of all nonnegative images of given finite support by replacing the negative pixels with zero and the pixels outside the region of support with the appropriate background pixel color value L_B to produce $\hat{f}_{NL}(i,j)$. The difference between $\hat{f}$ and $\hat{f}_{NL}$ is used to update the filter u. The algorithm involves the minimization of the following convex cost function with respect to $u(i,j)$:

$$\begin{aligned} J &= \sum_{\forall(i,j)} [\hat{f}_{NL}(i,j)-\hat{f}(i,j)]^2 \\ &= \sum_{(i,j)\in D_{sup}} \hat{f}^2(i,j)\left[\frac{1-\mathrm{sgn}(\hat{f}(i,j))}{2}\right] \\ &\quad + \sum_{(i,j)\in \overline{D}_{sup}} [\hat{f}(i,j)-L_B]^2 + \gamma\left[\sum_{\forall(i,j)} u(i,j)-1\right]^2 \end{aligned} \tag{6}$$

where $\hat{f}(i,j)=g(i,j)*u(i,j)$, and $\mathrm{sgn}(f)=-1$ if $f<0$ and $\mathrm{sgn}(f)=1$, if $f\geq 0$. D_{sup} is the set of all pixels inside the region of support, and $\overline{D}_{sup}$ is the set of all pixels outside the region of support. The support is defined as the smallest rectangle which can completely encompass the true unblurred object. The variable γ in third term of the equation is nonzero only when L_B is zero, i.e., the background color is black. The third term is used to constrain the parameters away from the trivial all-zero global minimum for this situation [15]. The user-specified parameters in the algorithm are the FIR filter dimensions, the unblurred object support and the parameter γ. The parameter γ is set in general to 1 when the background color is black. The initial setting for $u(i,j)$ is the discrete unit impulse (i.e., a value of one in the center of the FIR filter parameters and zero elsewhere). Assuming

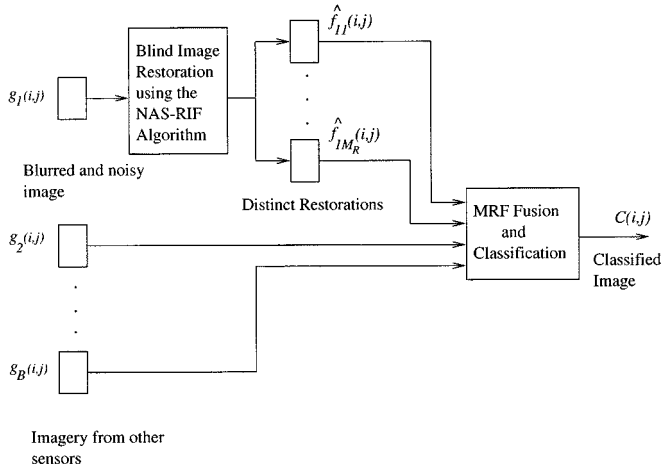

Fig. 3. Blind Image Fusion Method 1: The simultaneous merging of distinct restorations.

TABLE II
PERCENTAGE CLASSIFICATION ACCURACIES FOR THE SYNTHETIC DATA. IMAGE 2 IS AN ADDITIONAL REGISTERED IMAGE OF THE SAME ISLAND-LIKE SCENE WHICH IS DEGRADED SOLELY BY ADDITIVE WHITE GAUSSIAN NOISE WITH A SNR OF 20 dB. THE *boldface rows* REPRESENT THE RESULTS OF THE PROPOSED APPROACH

Image(s) Fused	Overall	Class 1	Class 2	Class 3	Class 4
Blurred Image (Fig. 5(b))	74.7	61.3	63.3	1.6	88.6
Rest. I (Fig. 5(c))	91.4	86.1	90.9	75.4	96.0
Rest. II (Fig. 5(d))	62.6	9.2	100.0	0.0	99.3
Rest. I & II	**93.8**	**88.1**	**95.2**	**75.0**	**98.7**
Image 2 (not shown)	92.2	83.7	98.5	49.5	100.0
Rest. I & Image 2	98.1	97.2	99.2	76.0	100.0
Rest. II & Image 2	96.2	98.3	100.0	0.0	100.0
Rest I & II, & Image 2	**98.3**	**98.8**	**98.5**	**65.7**	**100.0**

the steepest-descent minimization routine, the update law for $u(i,j)$ is given by

$$u_{k+1}(i,j) = u_k(i,j) - 2\eta \cdot \left[\sum_{(m,n)\in D_{sup}} g(m-i,n-j)\hat{f}_k(m,n)s \cdot (-\hat{f}_k(m,n)) + \sum_{(m,n)\in \overline{D}_{sup}} g(m-i,n-j)[\hat{f}_k(m,n) - L_B] + \gamma\left(\sum_{\forall(m,n)} u(m,n) - 1\right)\right] \quad (7)$$

where $\eta > 0$ is the update step-size, $s(\cdot)$ is the unit step function, $u_k(i,j)$ is the value of the filter u at the kth iteration, and $\hat{f}_k(i,j) = g(i,j) * u_k(i,j)$. Alternatively, the conjugate gradient routine may also be implemented to minimize J. Convergence using this approach is usually faster than the steepest-descent method. The specifics are provided in [15].

The assumptions of the NAS-RIF algorithm are as follows.

Assumption Set 2 (Assumptions Specific to the NAS-RIF Algorithm)

1) The imaged scene is comprised of a finite extent object against a uniform background. In general, a blind deconvolution algorithm requires a *reference* to perform deconvolution. The reference is often in the form of *a priori* information about the true undistorted image. In the case of the NAS-RIF algorithm, the information assumed about the true image is the nonnegativity of the pixels and the known finite extent of the original undistorted object.[6]
2) The support of the image is known or estimated *a priori*. This information can be obtained from undistorted versions of different sensor images of the same scene or by the use of support-finding algorithms [14].

The main advantage of the NAS-RIF method is its superior convergence properties to other techniques of its class [2]. In addition it is based on well-developed theory and is not computationally complex [15]. The major disadvantage is the excessive noise amplification that results due to the ill-posed nature of the blind deconvolution problem. One successful method of combating the numerical instability is to terminate the algorithm before excessive noise occurs in the restored output [15].

2) Statistical MRF Fusion for Classification: The MRF classification algorithm is attractive because it allows the incorporation of spatial, spectral and temporal information in a unified framework. A previous investigation by the authors [11] suggests that the MRF method is successful in classifying images degraded by non-Gaussian signal-dependent and nonadditive noise. This implies that MRF fusion may also perform well in the classification of restored images which we can consider to exhibit a type of signal-dependent noise as discussed in Section III-A.

Markov random field (MRF) models provide a methodological framework which allows the images from the different sensors to be merged statistically in a consistent way. The MRF classification method [5] uses the MRF models to form a likelihood function. The goal of the method is to find a classification that maximizes the posterior (MAP) distribution of the likelihood function. Determination of this MAP estimate requires the minimization of a sum of *energy functions* [5]. The overall function originally proposed in [5] is more complex than that implemented and discussed in this paper. The simplification arises because we assume that all images used for classification are taken at the same time instant and that there are no old ground truth maps available to aid classification.

Specifically, our MAP classification estimate requires the minimization of the following function

$$U(\mathbf{g}_1, \mathbf{g}_2, \cdots, \mathbf{g}_K, \mathbf{C}) = \sum_{s=1}^{K} \alpha_s U_{data}(\mathbf{g}_s, \mathbf{C}) + U_{sp}(\mathbf{C}) \quad (8)$$

where the subscript s denotes data from a particular sensor, K is the number of sensors to be fused, $\mathbf{g}_s, s = 1, 2, \cdots, K$ are the sensor images (in lexicographical order), α_s is the sensor-specific reliability factor (usually taken to be the individual sensor classification accuracy), and $\mathbf{C}$ is the classification of the scene and is the same dimensions as the sensor imagery. Each element of $\mathbf{C}$ is an integer representing a particular class. It is a value between 1 and N_{class} where N_{class} is the total number of classes in a given scene. U_{data} and U_{sp} are called the sensor-specific

[6]This information can be obtained in astronomy and satellite remote sensing applications.

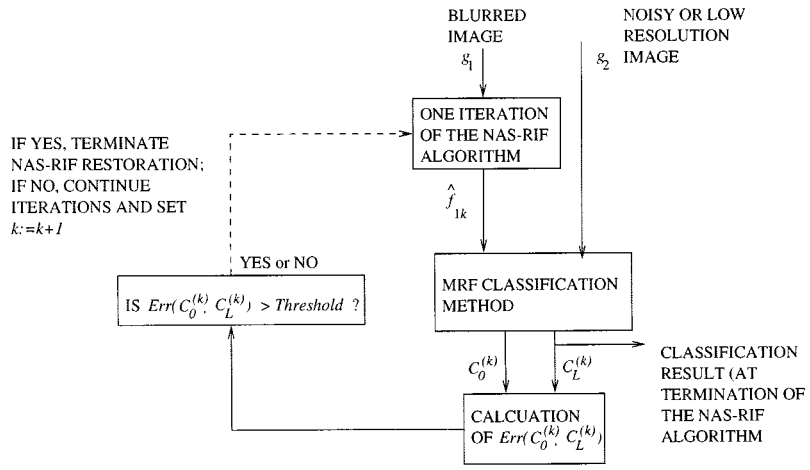

Fig. 4. Blind Image Fusion Method 2: The use of a fusion-based stopping criterion for blind image restoration.

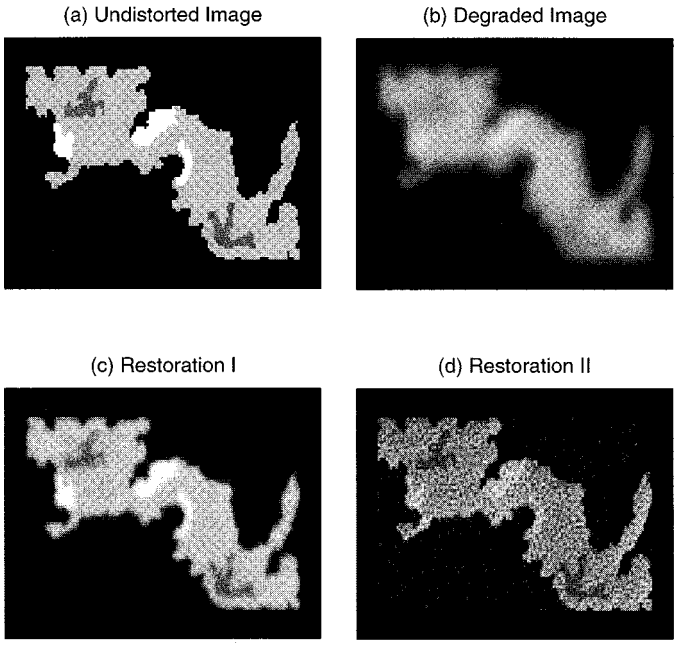

Fig. 5. Synthetic Image Data to Test Approach 1. The grey-levels represent the actual simulated radiance of the scene. (a) Original, (b) degraded image with BSNR of 40 dB, (c) restoration at sixth iteration, and (d) restoration at seventh iteration.

image statistic function and the spatial context energy function, respectively.

In the simulations performed we assumed that the image noise statistics can be modeled by a normal distribution and that the noise processes are independent from one another.[7]

As a result, U_{data} is given by

$$U_{data}(\mathbf{g}_s, C(i,j)) = \frac{B_s}{2} \lim |2\pi \boldsymbol{\Sigma}_c| + \frac{1}{2}(\mathbf{g}_s(i,j) - \mu_c)^T \boldsymbol{\Sigma}_c^{-1}(\mathbf{g}_s(i,j) - \boldsymbol{\mu}_c) \quad (9)$$

[7]This was successfully done by Schistad-Solberg *et al.* for non-Gaussian noise in [5].

where B_s is the number of bands for sensor s, $\boldsymbol{\Sigma}_c$ and $\boldsymbol{\mu}_c$ are the covariance matrix and the mean value vector of the radiance for class $C(i,j)$ for the bands in image sensor s, respectively, and $\mathbf{g}_s(i,j)$ is the image pixel vector corresponding to pixel (i,j).[8] The parameter sets $\boldsymbol{\Sigma}_c$ and $\boldsymbol{\mu}_c$ are assumed to be known prior to fusion. The spatial context function, which promotes regions of uniform classification, is given by [5]

$$U_{sp}(\mathbf{C}) = \beta_{sp} \sum_{(k,l)\in G_{(i,j)}} I(C(i,j), C(k,l)) \quad (10)$$

where β_{sp} is a user-specified nonnegative scalar parameter, $G_{(i,j)}$ is the set of all pixels adjacent and diagonal to (i,j), and $I(\cdot,\cdot)$ is given by

$$I(c_1, c_2) = \begin{cases} -1, & \text{if } c_1 = c_2 \\ 0, & \text{if } c_1 \neq c_2. \end{cases} \quad (11)$$

The overall algorithm we implemented is shown in Table I.

The value of β_{sp} is selected through experience. For the simulations in this paper the values were selected as suggested in [5]. The technique performs pixel-level fusion for $\beta_{sp} = 0$[9] and intermediate-level fusion for $\beta_{sp} > 0$. Equation (8) is not convex for $\beta_{sp} > 0$. The U_{sp} term implies that the classification of a pixel at location (i,j) is dependent on all other pixels in a square neighborhood $G_{(i,j)}$ centered about (i,j). This makes determining the global minimum of (8) computationally complex. To reduce the complexity, we implement the minimization routine using Besag's iterate conditional modes (ICM) algorithm [16] as suggested in [5]. The algorithm can become potentially trapped in local minima, but this was not a problem during the actual simulations performed.

It should be noted that if a single image is used for classification, then no fusion is performed by the MRF technique. The algorithm merely classifies the image.

[8]Because there are B_s bands for the image sensor, each pixel can be considered to have an associated vector whose elements consist of the radiance values from each of the bands.

[9]If $\beta_{sp} = 0$, then $U_{sp} = U_{data}$. Therefore, it can be seen from (9) that the minimum of U_{data} with respect to $C(i,j)$ only depends on the corresponding pixel value of each of the sensor images and, hence, the classification procedure involves pixel-level fusion.

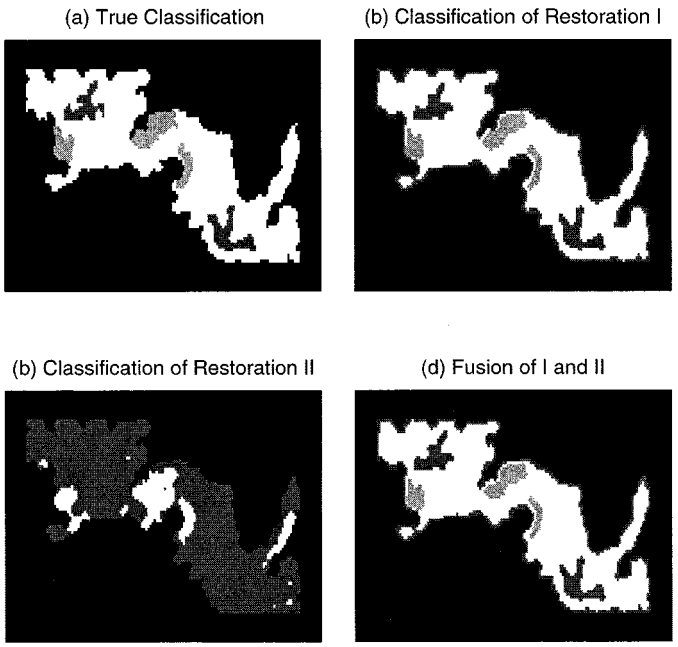

Fig. 6. Classification results for Approach 1 on synthetic data. The four different grey-levels represent each of the different classes in the image. The white, black, light grey and dark grey colors denote classes 1, 2, 3, and 4, respectively.

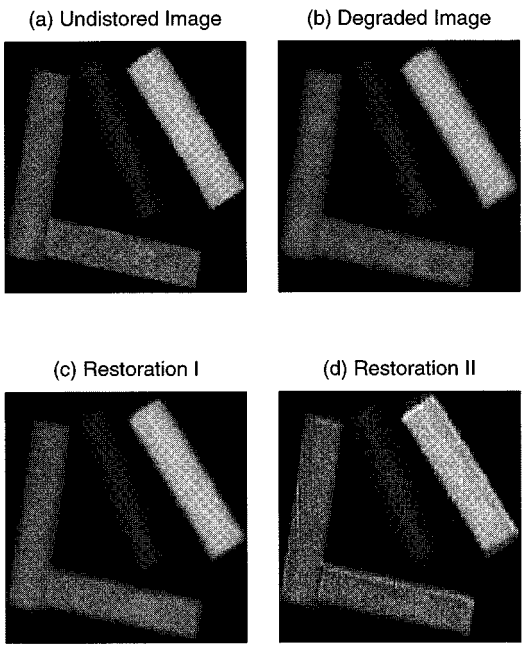

Fig. 7. Photographic color image data. The images represent the red band of a color photograph of chalk. (a) Original, (b) degraded image with BSNR of 40 dB, (c) restoration at first iteration, and (d) restoration at eleventh iteration.

The specific assumptions made by the algorithm are provided below.

Assumption Set 3 (Assumptions Specific to the MRF Classification Method)

1) The mean and (co-)variances of the radiance of the different classes in the scene are known or estimated from similar (nonblurred) data. The tonal and textural information about each class is used by the MRF method to segment the image into different regions. This information can be gathered from the use of image data containing similar types of regions [5].
2) The multispectral images are of the same dimensions. Interpolation methods may be used to resize the images to one size. The classified image result is the dimensions as the interpolated images.
3) The individual sensor reliability factors for each image are known prior to blind image fusion. This may be provided by the sensor manufacturer or can be approximated by the individual sensor image classification accuracy [5].

B. The Two Approaches

1) Approach 1: Simultaneous Fusion of Distinct Restorations: We consider for simplicity the classification of a single noisy blurred image, although the method can easily be extended to the situations in which other sensor images are available. The technique is comprised of two stages. The first stage of the technique involves the blind restoration of the image using the NAS-RIF algorithm. We hypothesize that if we treat each restoration at a selected set of iterations as the output of a different image sensor band and fuse the results for classification, then there will be an overall regularizing effect on the output.

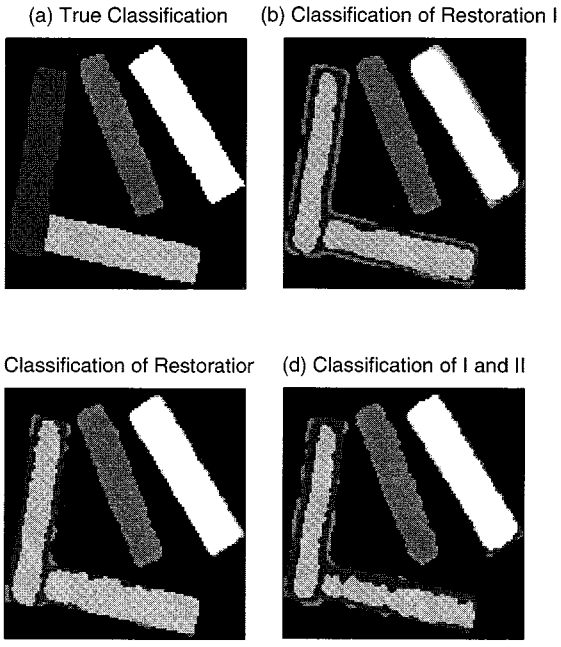

Fig. 8. Classification results for Approach 1 on photographic data. The five different grey-levels represent each of the different classes in the image. The white, light grey, medium grey, dark grey and black colors denote classes 1, 2, 3, 4, and 5, respectively.

The second stage fuses the various image estimates into a classified image using the MRF classification method. The fusion process takes into account the correlation in noise among the various image estimates. Fig. 3 gives an overview of the proposed architecture. We fuse several image estimates, which exhibit various degrees of blur removal and noise amplification,

TABLE III
PERCENTAGE CLASSIFICATION ACCURACIES FOR THE PHOTOGRAPHIC DATA. THE **boldface row** REPRESENT THE RESULTS OF THE PROPOSED APPROACH

Image(s) Fused	Overall	Class 1	Class 2	Class 3	Class 4	Class 5
Blurred Image (Fig. 7(b))	85.2	99.4	58.8	64.4	50.7	100.0
Rest. I (Fig. 7(c))	85.2	99.4	58.8	64.4	50.7	100.0
Rest. II (Fig. 7(d))	84.4	91.7	38.0	93.6	72.8	92.5
Rest. I & II	**86.4**	**95.0**	**46.6**	**95.6**	**61.0**	**93.8**

to produce a more reliable classified output. Assuming that we use the restorations at B different iterations of the NAS-RIF algorithm, each image estimate can be represented as

$$\begin{aligned}\hat{f}_k(i,j) &= g(i,j) * u_k(i,j) \\ &= f_k(i,j) + \hat{s}_k(i,j) + \hat{w}_k(i,j)\end{aligned} \tag{12}$$

where $k = 1, 2, \cdots, B$, and (12) follows directly from (5).

If we assume that the blurred image $g(i,j)$ has a moderate to high blurred signal-to-noise ratio, and that we obtain a series of B restorations such that $||h^{-1}(i,j) - u_k(i,j)|| < \epsilon$, where $\epsilon > 0$ is small and $||\cdot||$ represents the Euclidean norm, we can neglect $\hat{s}_k(i,j)$. Thus, we can consider the restored images to be corrupted primarily by $\hat{w}_k(i,j)$. Since the family of noise processes $\{\hat{w}_k(i,j)\}$ for $k = 1, 2, \cdots, B$ are all filtered versions of $w(i,j)$, it can be easily shown that if we assume $w(i,j)$ is zero mean Gaussian and white with variance σ_w^2,[10] the associated covariance matrix is given by

$$E\{\hat{\mathbf{w}}(i,j)\hat{\mathbf{w}}^T(i,j)\} = \sigma_w^2 \mathbf{R}_u \tag{13}$$

where $\hat{\mathbf{w}}(i,j) = [\hat{w}_1(i,j)\hat{w}_2(i,j)\cdots\hat{w}_B(i,j)]^T$, $E\{\cdot\}$ is the expectation operator, and $\mathbf{R}_u$ is the covariance matrix of $\{u_k(i,j)\}$ whose elements are given by

$$[\mathbf{R}_u]_{lm} = \sum_{\forall(i,j)} u_l(i,j)u_m(i,j) \tag{14}$$

for $l, m = 1, 2, \cdots, B$. We consider each restored image to be originating from a different band of the same sensor exhibiting correlated noise as described by the covariance matrix of (13),[11] and make use of the MRF classification method to fuse the results. The estimation of the parameters required for classification are discussed in Section V-B.

2) Approach 2: Restoration Stopping Criterion based on Fused Classifications: In our second approach, we investigate an error measure to determine a good stopping point for restoration. We desire a restored image which not necessarily provides a good classification on its own, but which when combined with other available sensor imagery produces an accurate classification.

We consider a three stage scheme in which restoration is performed in the first stage, and the data is fused and classified in the second stage. In the third stage, the results of the classification are fed back to the first stage to enhance the restoration. We make use of the smoothness properties[12] of the classified image to provide an estimate of the success of the blind image restoration stage. Often the visual quality of the restored image is not a good indicator of the reliability of higher-level processing tasks such as classification. Therefore, to improve the reliability of the fusion process, it is important to make use of information from the higher-level processing stage to enhance the restoration stage. Fig. 4 gives a summary of the proposed approach.

For simplicity we consider the fusion of a blurred image $g_1(i,j)$ and a noisy image of the same scene $g_2(i,j)$. At each iteration of the NAS-RIF algorithm, the restoration of $g_1(i,j)$ denoted $\hat{f}_{1k}(i,j)$ (where k is the current iteration) is passed through the MRF classification method. Classification of $\hat{f}_{1k}(i,j)$ and $g_2(i,j)$ is performed twice using two different β_{sp} values: $\beta_{sp} = 0$ and $\beta_{sp} = \beta_L$, where $\beta_L > 0$. We denote the classification results corresponding to $\beta_{sp} = 0$ and $\beta_{sp} = \beta_L$ by $\mathbf{C}_0^{(k)}$ and $\mathbf{C}_L^{(k)}$, respectively. The value of β_L is user-defined. Experience shows that any values between 0.5 and ten work well in simulations. A measure of the difference between $\mathbf{C}_0^{(k)}$ and $\mathbf{C}_L^{(k)}$ is used to determine the possibility of noise amplification effecting the fusion results. This measure is given by

$$\begin{aligned}&Err(\mathbf{C}_0^{(k)}, \mathbf{C}_L^{(k)}) \\ &\quad = \frac{\sum_{\forall(i,j)} \mathcal{I}(\mathbf{C}_0^{(k)}(i,j), \mathbf{C}_L^{(k)}(i,j))}{M} \times 100\end{aligned} \tag{15}$$

where M is the number of pixels in $\mathbf{C}_0$ or $\mathbf{C}_L$, and $\mathcal{I}(c_0, c_L) = 0$ if $c_0 = c_L$ and $\mathcal{I}(c_0, c_L) = 1$ otherwise. If this difference increases dramatically from one iteration to the next (a measure which is user-defined as explained in Section V-B), then the algorithm is terminated and the classified output is taken to be $\mathbf{C}_L^{(k-1)}$ where k is the current iteration.

As discussed in Section IV-A2, the fusion procedure involves the minimization of (8) which is comprised of two terms: a sensor-specific component involving U_{data} and a spatial context function U_{sp}. The β_{sp} parameter governs the relative importance of U_{sp}. If we consider the case where $\beta_{sp} = 0$, then the classification process is equivalent to pixel-level fusion as discussed in Section IV-A2. When $\beta_{sp} > 0$, the classification is equivalent to an intermediate-level fusion in which the classification of a pixel depends on that of its eight nearest neighbors. The larger the values of β_{sp} the stronger the effect of U_{sp} and the smoother and more regularized the classification.[13] It is reasonable to assume that if noise amplification significantly effects the classification, then the relative difference between $C_0^{(k)}$ and $C_L^{(k)}$ at iteration k will be large. The MRF classification method assumes knowledge of the variance of the additive noise for each restored image. Using similar approximations as

[10]This assumption can be justified using the central limit theorem, and presuming that the image undergoes additive noise from a number of independent sources.

[11]The reader should note that the elements of $\mathbf{R}_u$ are calculated using (14).

[12]The smoothness properties give a good indication of the extent of noise amplification and, therefore, somewhat indicate the degree of accuracy of the restored result.

[13]This is due to the fact that U_{sp} promotes uniformly classified neighborhood regions.

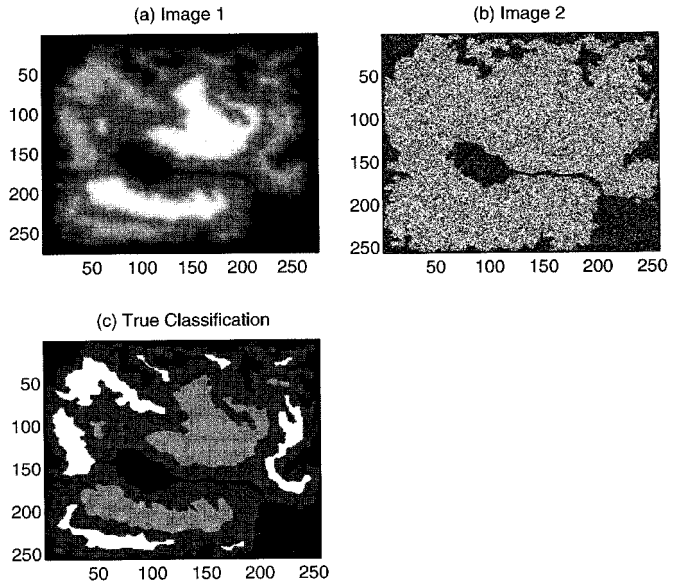

Fig. 9. Synthetic image data to test approach 2. Synthetic image data to test approach 2. (a) Blurred image, (b) noisy image degraded by multiplicative chi-squared noise, (c) true classification of the scene. The black, dark grey, light grey and white shades denote classes 1, 2, 3, and 4, respectively.

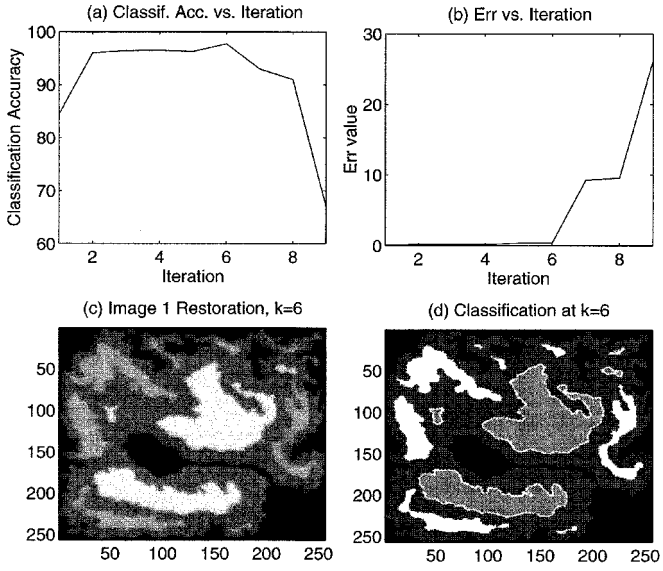

Fig. 10. Classification results for Approach 2 on synthetic data. (a) Classification accuracy versus iteration of the NAS-RIF algorithm. (b) The stopping criterion error as a function of iteration. (c) Restored image 1 at the sixth iteration of the NAS-RIF algorithm. (d) Classification result for the sixth iteration (corresponds to the result with the highest classification accuracy).

discussed at the end of Section IV-B1 in the formulation of (13) and (14), we can show that the variance of the associated additive noise of the restored image can be approximated as

$$\sigma^2_{\hat{w}} = \sigma^2_w \sum_{\forall(i,j)} u^2(i,j). \tag{16}$$

V. Simulation Results

A. Figure of Merit

We assess the success of the proposed algorithms using a popular quantitative figure of merit known as the classification accuracy (CA). CA is defined as:

$$CA \triangleq \frac{\text{Number of correctly classified image pixels}}{\text{Total number of pixels in the classified image}} \times 100, \tag{17}$$

or more formally

$$CA(\hat{\mathbf{C}}) = \frac{\mathcal{I}(\hat{\mathbf{C}}, \mathbf{C})}{M} \times 100 \tag{18}$$

where $\hat{\mathbf{C}}$ is the classification estimate, M is the number of pixels in $\hat{\mathbf{C}}$, $\mathbf{C}$ is the *true* classification, and $\mathcal{I}(\hat{c}, c) = 1$ if $\hat{c} = c$ and $\mathcal{I}(\hat{c}, c) = 0$ otherwise. The true classification $\mathbf{C}$ is known for synthetic data. For our simulations involving real data, we estimate it by passing an undistorted, but registered image of the scene through the MRF classification method. The result is then modified, if necessary, by a user to correct any pixels determined to be misclassified by the human eye.

B. Evaluation of Results

Each proposed blind fusion approach is tested on synthetic and real photographic data. For each technique we provide results for a synthetic image set and a photographic image set.

1) Results for Approach 1: Fig. 5 shows one set of synthetic data used to test our first approach. An undistorted synthetic image [shown in Fig. 5(a)] is blurred with a 7×7 Gaussian PSF[14]. White Gaussian noise is added to the resulting image to produce a blurred signal-to-noise ratio (BSNR) of 40 dB. The overall degraded image is shown in Fig. 5(b). The NAS-RIF algorithm is applied to the blurred and noisy image using an FIR filter size of 5×5[15] and assuming an accurate support size of 68×130. The restoration algorithm was terminated when excessive noise amplification was observed visually. The two image estimates prior to termination are used for classification. The images are shown in Fig. 5(c) and (d), and correspond to the sixth and seventh iterations of the NAS-RIF algorithm, respectively. We call the restoration at the sixth iteration, "restoration I,"and that at the 7th iteration, "restoration II." We can visually determine that restoration I shows some residual blurring while restoration II is clearer, but exhibits noise amplification.

The classification results are provided in Fig. 6 The scene was classified into four different kinds of regions delineated by four distinct grey levels. Fig. 6(a) shows the true classification, and Fig. 6(b) and (c) the classifications for restorations I and II, respectively. The resulting classification for our first approach which performs fusion on both of the restorations is shown in Fig. 6(d). The mean value of each class was taken directly from the image in Fig. 5(a). The noise was assumed to be the same

[14]A Gaussian PSF is commonly found in astronomical applications.

[15]This filter size has been found through experience to produce fast convergence for noisy images, [17].

for all classes and the covariance matrix was estimated using (13). The value of σ_w^2 was estimated from the blurred and noisy image, by measuring it from a block in the uniform background of the image.

Table II summarizes the classification accuracies. The MRF classification method was applied assigning $\beta_{sp} = 0.5$ (as was used in [5]). The two restorations show widely varying CA's. Classification of restoration I is shown to significantly improve CA over that of classifying the blurred image; classification of restoration II, however, reduces the CA over that of the blurred image. When we fuse the two images using our first approach we see that the classification accuracy is further improved over that of classifying restoration I solely. The results demonstrate how the proposed approach can improve the classification of degraded images. Table II also provides results for the classification of the degraded image with a second noisy (but unblurred image) with an SNR of 20 dB.[16] We see that our fusion approach still improves CA, but the improvement is less exaggerated due to the presence of the second noisy image.

We perform the same simulations on photographic color data of chalk, shown in Fig. 7. Five distinct classes are to be identified: four correspond to each of the colored chalk and a fifth to the background. The undistorted red band of the image in Fig. 7(a) is blurred to produce the effect of an out-of-focus camera. White Gaussian noise was added to the result to produce a BSNR of 40 dB. The restorations after the 1st and 11th iteration of the NAS-RIF algorithm are denoted "restoration I" [shown in Fig. 7(c)] and "restoration II" [shown in Fig. 7(d)], respectively. The NAS-RIF algorithm was run using a FIR filter size of 5×5, and a support size of 115×110 estimated from an unblurred green band of the image. The MRF classification method was applied to the blurred image, and the individual and combined restored images. The mean and variances (due to texture) of each of the classes were estimated from a similar, but unblurred version of the scene, and β_{sp} was set to 0.5. The overall variance of each class was estimated to be the sum of the variance due to the texture and the variance of the filtered noise $\tilde{w}(i,j)$ given by (13). The results are shown in Fig. 8. The approach has difficulty distinguishing between Class 2 and Class 4 because both colors of chalk exhibit similar luminance values in the red band [see Fig. 7(a)]. Incorporation of information in the green band, for example, can solve the ambiguity; such an example is shown in the next section.

The results demonstrate the potential of the fusion approach for robust classification. An improvement in CA of 1% to 2% is observed over that of classifying restoration I or II alone (i.e., without fusion). Thus, in some situations the added complexity of the fusion process gives a relatively small improvement. We believe that our proposed method is useful in situations in which accuracy is of primary concern. The suitability of the additional computational load is an application-dependent issue. If computational load is a concern, then the fusion scheme as we have presented it may not be a practical solution.

2) Results for Approach 2: Our second approach was tested on synthetic and photographic data to assess its potential for robust classification. Figs. 9 summarize the results for a synthetic data set. A synthetic image is degraded using a Gaussian PSF and noise to produce a BSNR of 50 dB; the result (denoted Image 1) is shown in Fig. 9(a). A second registered image of the

[16]This image is not displayed in any figures.

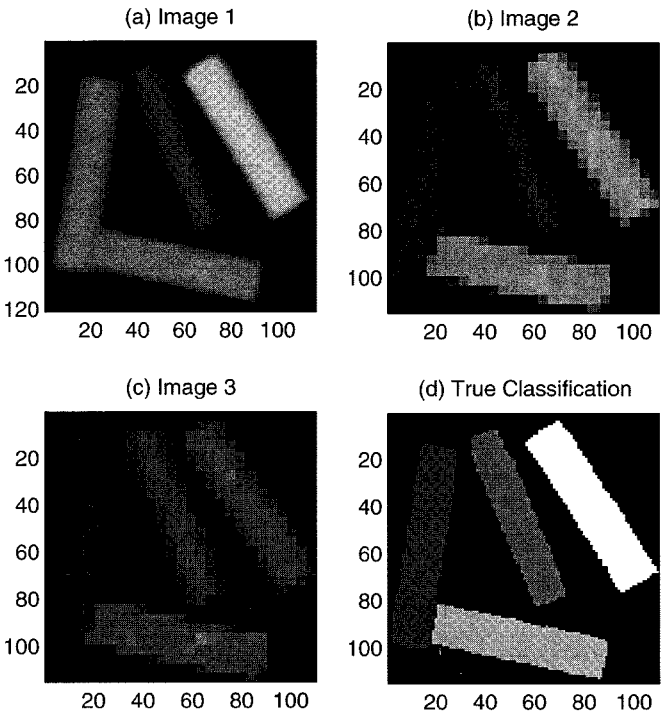

Fig. 11. Photographic color image data to test Approach 2. Photographic color image data to test Approach 2. (a) Blurred red band of the original image, (b) low resolution (by a factor of four) image of the green band, (c) low resolution (by a factor of four) image of the blue band, and (d) true classification results. The classifications are distinguished according to grey-level.

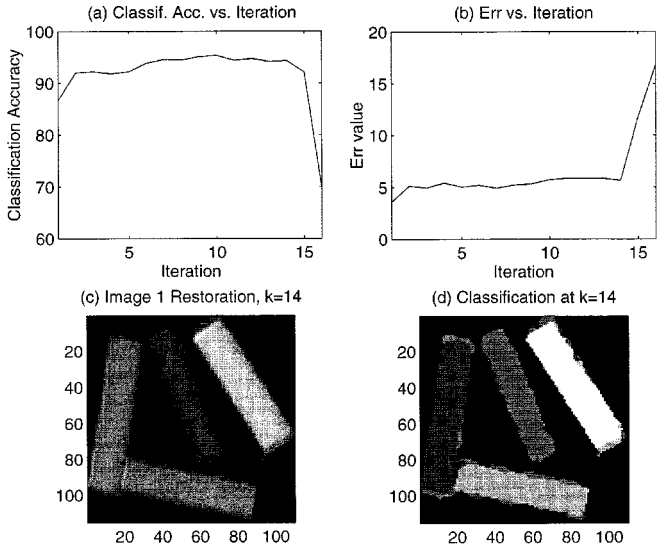

Fig. 12. Classification results for Approach 2 on photographic data. (a) Classification accuracy versus iteration of the NAS-RIF algorithm. (b) The stopping criterion error as a function of iteration. (c) Restored image 1 at the fourteenth iteration of the NAS-RIF algorithm. (d) Classification result for the 14th iteration.

TABLE IV
SUMMARY OF CHARACTERISTICS OF THE PROPOSED BLIND FUSION METHODS FOR IMAGE CLASSIFICATION

Classification Scenarios	Method 1	Method 2
How improved performance over non-fused classification is achieved	Fusion of different restorations regularizes the classification result by exploiting the complementarity in different restorations.	Automatic termination of the iterative restoration stage is dependent on the fused classification result so that an image estimate is selected which is appropriate for classification.
Disadvantages of the methods	Selection of the different restorations to fuse is an ad hoc procedure.	Only one restoration of the image band is used for classification so the complementarity of different restorations is not directly exploited for classification.
Effect of fusing the restored data with different sensor imagery	Improvement from the fusion of different restorations is reduced when other complementary imagery (from different bands) is included in the classification process.	The weakness of Method 2 against Method 1 is diminished when other sensor data is available for classification because the fusion of the different types of imagery regularizes the classification result.
Observations on algorithm performance	Performance is a function of the complementarity among the restorations to be fused. Both restorations should be fairly accurate, but should exhibit distinct properties (such as slight noise amplification & mild residual blurring) for the classification to be successful.	Improvement is a function of the behaviour of the iterative blind restoration algorithm employed. The more accurate the restoration stage, the more successful the classification output.
Behaviour for low SNRs	If the SNR of the degraded image data is low then it is often disadvantageous to include the restored data in the fusion process of classification.	Low SNR data causes unpredictable behaviour in the restoration stage which makes the termination point unreliable.

scene (denoted Image 2) is to be fused with the first to produce an overall classification of the scene. The second image is unblurred, but suffers from multiplicative chi-squared noise with eight degrees of freedom. The true classification of the scene is shown in Fig. 9(c). The four different grey levels denote the four classes in the scene.

Our second approach was applied to classify the images. An FIR filter size of 5×5, and an image support size of 256×256 accurately estimated from Image 2 were used for the NAS-RIF algorithm stage. Parameter settings of $\beta_L = 5$ and $Threshold = 10Err(C_0^{(k-1)}, C_L^{(k-1)})$ were used in the MRF classification and termination stages, respectively. The value of β_L was selected based on our previous investigation of the MRF classification method [11]. The classification results are shown in Fig. 10. As we can see, our proposed error metric of (15) indicates that iteration 6 of the NAS-RIF algorithm is the appropriate stopping point for an accurate classification. From Fig. 10(a), we see that indeed iteration 6 produces the optimal classification accuracy.

Similar results are obtained when our second approach is applied to the same photographic color data previously used to assess our first approach. Figs. 11 and 12 show the results. Again, we use an FIR filter size of 5×5, and an accurate image support size of 115×100 estimated from the unblurred green band. We select $\beta_L = 15$, and $Threshold = 2Err(C_0^{(k-1)}, C_L^{(k-1)})$ determined from observing the Err vs. iteration function shown in Fig. 12(b). We fuse the same blurred photographic red band image discussed in Section V-B1 and fuse it with low resolution (four times undersampled) versions of the green and blue bands. As we can see, the technique determines the 14th iteration to be an appropriate stopping point. Fig. 12(a) shows that this produces a CA of 93.3%. It should be noted that the algorithm was not terminated exactly at the peak CA. The peak of 95.3% was found at iteration 10, however, the termination point did produce a fairly reliable estimate. The improved accuracy of the second approach over to the first for the photographic data is due to the additional information provided by the low resolution green and blue bands of the image.

If we do not make use of the green and blue bands in our second approach, the algorithm terminates at iteration 3 which produces a classification accuracy of 85.1%. It should be emphasized that our results demonstrate that although the use of Method 2 provides an appropriate termination point, the results are not as successful as fusing two different restorations of the red band. As shown in Table III fusing the restorations at iterations 1 and 11 produces an accuracy of 86.4% which is higher due to the regularizing effect of the fusion process for classification.

C. Discussion

Although our first approach often improves CA, experience with simulation results reveals that the method is not always predictable or successful. For example, fusing a restoration which produces a high CA with one which produces a poor CA can result in a slight reduction in the overall fused CA. Selection of the appropriate restorations to fuse is an ad hoc procedure. As a rule of thumb, the authors observed that fusing complementary restorations (e.g. fusing an image which has mild noise amplification with one which exhibits residual blurring) improves CA over nonfused classification. In addition, we found that fusing a poor restoration (possibly exhibiting severe blurring and/or noise amplification) with a "good" restoration reduces the CA over that of a nonfused classification of the "good restoration."

The authors also observed that if complementary nonblurred, but possibly noisy information of the scene was available, then the degree of improvement of our first approach was diminished. The complementary information, which could be in the form of another image of the scene, raised the overall CA. However, the improvement in CA by fusing two or more restorations (instead of just one) with the additional imagery was reduced; for example, instead of a 1% improvement in CA, a 0.1% improvement was found.

Our second approach allows algorithm termination at a point of the NAS-RIF algorithm which produces a good classification accuracy, however, this is not always the most optimal result. In our simulations, we found that inclusion of complementary nonblurred imagery improved the performance of the algorithm (i.e., a more optimal classification was obtained).

The termination criteria is user-defined and depends on the kinds of images fused for classification. We found through simulations that threshold values of $10Err(\mathbf{C}_0^{(k)}, \mathbf{C}_L^{(k)})$ and $2Err(\mathbf{C}_0^{(k)}, \mathbf{C}_L^{(k)})$ worked well for high SNR synthetic and real image data, respectively. For low SNR's (i.e., approximately 20 dB and lower), the thresholds were often not reached. In such situations, it was found through experience that inclusion of the degraded image data did not improve the classification accuracy. Thus, it was better not to include the degraded imagery in the fusion process.

The results presented in this paper correspond to classification of high SNR blurry images. Simulations were also performed for degraded images with low SNR's. Our investigation revealed that the restorations of low SNR blurred images do not improve the classification accuracy and in many cases degrade the quality of the classification results. The authors have observed that this is due in part to the poor performance of the NAS-RIF algorithm for low SNR's. Noise amplification is almost immediately apparent during the restoration procedure. In such situation, the authors would recommend excluding the degraded data from the fusion process.

A summary of the main characteristics of each classification scenario is provided in Table IV. The authors believe that future work should involve combining the two approaches proposed in this paper to define a quantitative figure of merit to assess the potential of a set of restorations to produce an optimal fused CA. In addition, it would be useful to approximate a reliable statistical threshold for Method 2 by modeling the progression of noise amplification in the blind restoration process.

VI. Conclusions

Two architectures for the simultaneous blind deblurring and classification of imagery are proposed. The first approach attempts to regularize the classification process by fusing restorations of the blurred image exhibiting complementary information. The second technique determines a good blind restoration termination point to produce a reliable classification. The techniques are shown to produce promising results.

References

[1] M. A. Abidi and R. C. Gonzalez, *Data Fusion in Robotics and Machine Intelligence*. New York: Academic, 1992.

[2] D. Kundur and D. Hatzinakos, "Blind image deconvolution," *IEEE Signal Processing Mag.*, vol. 13, pp. 43–64, May 1996.

[3] A. K. Jain, *Fundamentals of Image Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[4] P. K. Varshney, "Scanning the issue," *Proc. IEEE*, vol. 85, pp. 3–5, Jan. 1997.

[5] A. H. S. Solberg, T. Text, and A. K. Jain, "Multisource classification of remotely sensed data: Fusion of landsat tm and sar images," *IEEE Trans. Geosci. Remote Sensing*, vol. 34, pp. 100–113, Jan. 1996.

[6] A. H. S. Solberg, A. K. Jain, and T. Taxt, "A Markov random field model for classification of multisource satellite imagery," *IEEE Trans. Geosci. Remote Sensing*, vol. 32, pp. 768–778, July 1994.

[7] M. C. Dodson, L. E. Pierce, and F. T. Ulaby, "Knowledge-based land-cover classification using ers-1/jers-1 sar composites," *IEEE Trans. Geosci. Remote Sensing*, vol. 31, pp. 83–99, Jan. 1993.

[8] J. A. Benediktsson, P. H. Swain, and O. K. Ersoy, "Neural network approaches versus statistical methods in classification of multisource remote sensing data," *IEEE Trans. Geosci. Remote Sensing*, vol. 28, pp. 540–552, July 1990.

[9] S. B. Serpico and F. Roli, "Classification of multisensor remote-sensing images by structured neural networks," *IEEE Trans. Geosci. Remote Sensing*, vol. 33, pp. 562–577, May 1995.

[10] A. K. Katsaggelos, *Digital Image Restoration*. New York: Springer-Verlag, 1991.

[11] D. Hatzinakos and D. Kunker, "Blind image fusion for surveillance," Dept. National Defense, Tech. Rep. W7714-6-9990, 1997.

[12] S. Haykin, *Blind Deconvolution*. Englewood Cliffs, NJ: Prentice-Hall, 1994.

[13] J. Biemond, R. L. Lagendijk, and R. M. Mersereau, "Iterative methods for image deblurring," *Proc. IEEE*, vol. 78, pp. 856–883, May 1990.

[14] D. Kundur and D. Hatzinakos, "Blind image restoration via recursive filtering using deterministic constraints," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing*, 1996.

[15] ——, "A novel blind deconvolution scheme for image restoration using recursive filtering," *IEEE Trans. Signal Processing*, vol. 46, Feb. 1998.

[16] J. Besag, "On the statistical analysis of dirty pictures," *J. R. Stat. Soc. B*, vol. 48, pp. 259–302, 1986.

[17] D. Kundur, "Blind deconvolution of still images using recursive inverse filtering," M.S. thesis, Univ. Toronto, Toronto, Ont., Canada, 1995.

Deepa Kundur (S'93–M'00) was born in Toronto, Ont., Canada. She received the B.A.Sc., M.A.Sc., and Ph.D. degrees, all from the Department of Electrical and Computer Engineering, University of Toronto, in 1993, 1995, and 1999, respectively.

She is currently a tenure track Assistant Professor with the Communications Group, University of Toronto. Her academic interests span the areas of digital watermarking for electronic commerce and content-based multimedia signal processing. Her research also encompasses blind image restoration and data fusion for the classification of remote sensing imagery. Her consulting experience includes the development and implementation of algorithms for the Defense Research Establishment Ottawa and Vaytek, Inc.

Dimitrios Hatzinakos (S'87–M'90–SM'98) received the Diploma degree from the University of Thessaloniki, Greece, in 1983, the M.A.Sc degree from the University of Ottawa, Ottawa, Ont., Canada, in 1986 and the Ph.D. degree from Northeastern University, Boston, MA, in 1990, all in electrical engineering.

In 1990, he joined the Department of Electrical and Computer Engineering, University of Toronto, Toronto, Ont., Canada, where he is now a tenured Associate Professor. His research interests are in the area of digital signal processing with applications to wireless communications, image processing, and multimedia. He has organized and taught many short courses on modern signal processing frameworks and applications devoted to continuing engineering education and given numerous seminars in the area of blind signal deconvolution. He is author/coauthor of more than 90 papers in technical journals and conference proceedings and he has contributed to four books in his areas of interest. His experience includes consulting to Electrical Engineering Consociates Ltd. and contracts with United Signals and Systems, Inc., Burns and Fry Ltd., Pipetronix Ltd., Defense Research Establishment Ottawa, and Vaytek, Inc.

Dr. Hatzinakos has been an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING since 1998 and the Guest Editor for the special issue of *Signal Processing* on signal processing technologies for short burst wireless communications. He was a member of the IEEE Statistical Signal and Array Processing Technical Committee from 1992 until 1995 and Technical Program Co-Chair of the 5th Workshop on Higher-Order Statistics in July 1997. He is a member of EURASIP, the Professional Engineers of Ontario) and the Technical Chamber of Greece.

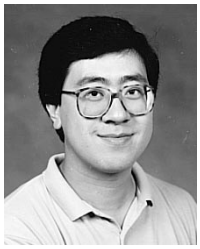

Henry Leung (S'88–M'90) received the B.Math. degree in applied mathematics from the University of Waterloo, Waterloo, Ont., Canada, in 1984, the M.Sc. degree in mathematics from the University of Toronto, Toronto, Ont., Canada, in 1985, and the M.Eng. and Ph.D. degrees in engineering physics and electrical engineering from McMaster University, Hamilton, Ont., in 1986 and 1991, respectively.

From 1990 to 1991, he was a Research Engineer with the Communications Research Laboratory, McMaster University. In 1991, he joined the Defense Research Establishment Ottawa, where he was involved in the design of automated systems for air and maritime multisensor surveillance. Since 1998, he has been with the Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alta., Canada, where he is now an Associate Professor. He is also an Adjunct Professor with the Department of Electrical Engineering, McMaster University. His research interests include chaos, computational intelligence, data fusion, nonlinear signal and image processing, multimedia, and wireless communications.