

Piecewise and Local Image Models for Regularized Image Restoration Using Cross-Validation

Scott T. Acton, *Senior Member, IEEE*, and Alan Conrad Bovik, *Fellow, IEEE*

Abstract—We describe two broad classes of useful and physically meaningful image models that can be used to construct novel smoothing constraints for use in the regularized image restoration problem. The two classes, termed *piecewise image models* (PIM's) and *local image models* (LIM's), respectively, capture unique image properties that can be adapted to the image and that reflect structurally significant surface characteristics. Members of the PIM and LIM classes are easily formed into regularization operators that replace differential-type constraints. We also develop an adaptive strategy for selecting the best PIM or LIM for a given problem (from among the defined class), and we explain the construction of the corresponding regularization operators. Considerable attention is also given to determining the regularization parameter via a cross-validation technique, and also to the selection of an optimization strategy for solving the problem. Several results are provided that illustrate the processes of model selection, parameter selection, and image restoration. The overall approach provides a new viewpoint on the restoration problem through the use of new image models that capture salient image features that are not well represented through traditional approaches.

Index Terms—Cross-validation, image restoration, local monotonicity, regularization.

I. INTRODUCTION

ONE OF THE classic problems in image processing is the restoration of a linearly degraded and noisy optical image to a pristine state. In this problem, an ideal image that has been digitized has been distorted by a lowpass blur arising from the motion of objects or of the sensor, defocus, or deficiencies in the optical system, for example. The problem of reconstituting the ideal image from the blurred observation can be greatly complicated by various factors such as nonlinearities in the blurring process, noninvertibility of the linear distortion (frequency destruction), and the addition of broadband or high-frequency noise to the blurred image. The noise may arise from, e.g., thermal effects, transmission noise in a channel, recording errors, or imperfections in the process of digitization. In this paper, we restrict our attention to the problem of reducing linear image blur in the presence of additive noise. The goals of sharpening and smoothing are conflicting, and any linear filtering approach is problematic: the reversion

of image blur requires a highpass sharpening of the image, whereas noise reduction ordinarily implies a degree of lowpass smoothing. Intermediate (bandpass techniques) that attempt to mediate this conflict, e.g., Wiener filtering, generally fail to satisfy the viewer.

Thus, a great number of optimization-based strategies have been developed that achieve better results; see [16] for a compendium of recent work. A fundamental guiding principle behind most of this work is the realization that the restoration problem is inherently ill-posed, hence strategies that regularize the problem, or make it well-posed, have been in the forefront of this research [16], [17]. The regularization approach poses the two conflicting goals, deblurring and noise eradication, as separate constraints to be simultaneously satisfied. A *data constraint* attempts to force a blurred version of the solution to resemble the degraded image that was acquired. In this way, the data constraint encourages deconvolution of the signal. A *smoothing constraint* penalizes solutions that deviate from smoothness. The smoothing constraint typically fills several roles: as a regularization operator that guarantees a unique, stable solution, as a model of the ideal image's local characteristics, and as a means for suppressing noise. Generally, the data and smoothing constraints are both expressed as functionals to be minimized over the space of images. Simultaneous minimization is most usually attained by linearly combining them into a single cost or energy functional to be minimized. The balance between the two constraints is dictated by the linear weight or *regularization parameter*.

The background literature on regularized image restoration has become quite large. Several studies examine techniques used to define the data constraint [16] and to evaluate the regularization parameter [6], [8], [13]–[15], [22], [24], [29]. Other works report progress on developing more effective methods of optimization for regularization [3], [4], [7], [9], [20], [31]. Recently, spatially and temporally adaptive methods of balancing the data constraint and the smoothing constraint have been explored [11], [12], [17], [18], [21]. Adaptive solutions, such as that of Kang and Katsaggelos [12], allow weighting that may change between pixel sites and that may also change between iterations.

Methods for defining the data constraint and methods for determining the regularization parameter have been well studied. Indeed, the choice of a data constraint that seeks deconvolution is not really open to argument, provided that the linear degradation assumption is accurate. Compared to the study of regularization parameter, less scrutiny has been given to the smoothing constraint. Indeed, it is nearly always taken to be a linear highpass operator, such as the integral (sum) of the

Manuscript received November 4, 1996; revised June 19, 1998. This material is based upon work supported in part by the U.S. Army Research Office under Grant DAAH04-95-1-0255. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jun Zhang.

S. T. Acton is with the Oklahoma Imaging Laboratory, School of Electrical and Computer Engineering, Oklahoma State University, Stillwater, OK 74078 USA (e-mail: sacton@master.ceat.okstate.edu).

A. C. Bovik is with the Laboratory for Vision Systems, Center for Vision and Image Sciences, Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX 78712-1084 USA.

Publisher Item Identifier S 1057-7149(99)03431-4.

Laplacian magnitude or the gradient magnitude over the image domain. Minimizing this kind of constraint presupposes a certain image model, or *characteristic property*, which usually does not hold everywhere in the image, e.g., near critical, information-bearing discontinuities. Recognizing this limitation, several researchers have proposed refined regularization strategies that enforce smoothness on a piecewise basis, using line processes to inhibit oversmoothing [4], [7], [9], [23], [31]. While improved efficacy in restoration has often been shown using these techniques (for low noise levels, generally), the introduction of a line process approach is concomitant to using an edge detector on the image; such strategies always have problems in a significant noise environment.

We take the position that insufficient attention has been paid to the development of more specialized smoothing constraints that attempt to model the image as something other than a function that has small derivatives almost everywhere. Instead, images should be modeled with descriptions that reflect more sophisticated surface properties that embody smoothness in a more interesting way. Of course, such a goal is quite ambitious owing to the considerable diversity of images and the surfaces from which they project. Therefore, the models used must also admit some degree of generality, although specificity may be attained in application-dependent circumstances.

In this paper, we propose two broad classes of useful image models that can be used to construct novel smoothing constraints for use in the regularized image restoration problem. Of course, “smoothness” is taken in a different sense in these classes. The two classes, referred to here, respectively, as the *piecewise image models* (PIM’s) and the *local image models* (LIM’s), each capture unique image properties. Members of the PIM and LIM classes can be easily formed into a regularization operator and used in place of differential-type constraints. In this paper, we propose and define the PIM and LIM classes, and for a given blurred and noisy image, we provide a method to select the proper PIM or LIM, to construct the corresponding regularization operator, and to set the regularization parameter. The overall approach provides a new avenue that uses novel image models for image restoration.

In the following section, the restoration problem is introduced and notation is made. The PIM’s and LIM’s are defined in Section III, along with the corresponding regularization operators. Selection of the regularization parameter and selection of the proper PIM/LIM for regularization are discussed in Section IV. Then, the minimization of the regularization functional is treated. Finally, parameter selection and image restoration results are provided in Section VI.

II. THE IMAGE RESTORATION PROBLEM

Image restoration algorithms of the type considered here seek to compute an image estimate $\hat{\mathbf{i}}$ from blurred, corrupted image data. The observed image \mathbf{g} is the result of degradation by a spatially invariant convolution operator \mathbf{B} and the addition of noise \mathbf{n} . So

$$\mathbf{g} = \mathbf{B}\mathbf{i} + \mathbf{n} \quad (1)$$

where \mathbf{B} is assumed to be block Toeplitz. The matrices \mathbf{g} , \mathbf{i} , and \mathbf{n} have the same size and contain a total of N elements (pixels).

Since the image restoration problem presented by (1) is ill-posed, an image estimate $\hat{\mathbf{i}}$ can be computed by a regularization procedure. The problem may be considered an optimization problem under two constraints. The first constraint forces consistency with the observed data, and essentially reverses the blurring process through deconvolution. Since this deconvolution operation is sensitive to noise and typically does not yield a smooth result, a second constraint is included that enforces “smoothness” on the solution. Typically, smoothness is interpreted as an integral (sum) of absolute spatial derivatives (differences) being small. In this paper, we replace this standard notion of smoothness with conformity to a specified image model in one of two classes: the PIM class or the LIM class.

In the classical Tikhonov approach [30], regularization is established by minimizing $\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}}$, to enforce smoothness, subject to the constraint $\|\mathbf{g} - \mathbf{B}\mathbf{h}\|_{\mathcal{D}} \leq \varepsilon$, to enforce consistency with the observed image data. Here, $\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}}$ increases with a decrease in the smoothness of \mathbf{h} . Or the two constraints may be simultaneously minimized, as suggested here, using (the Miller approach) [19]:

$$\|\mathbf{g} - \mathbf{B}\mathbf{h}\|_{\mathcal{D}} \leq \varepsilon \quad (2)$$

and

$$\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}} \leq E. \quad (3)$$

So, the two constraint problem may be formed using the Miller approach. The two terms may be combined into a single cost functional $E(\mathbf{h})$, termed the *energy functional*. Then, the quality of two solutions can be compared by evaluating their respective energies. The solution to the image restoration problem, once $E(\mathbf{h})$ is formed, corresponds to a minimization problem. The energy functional is given by

$$E(\mathbf{h}) = \|\mathbf{g} - \mathbf{B}\mathbf{h}\|_{\mathcal{D}} + \lambda \|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}} \quad (4)$$

with an optimal solution at

$$\hat{\mathbf{i}} = \arg \min_{\mathbf{h}} \{E(\mathbf{h})\}. \quad (5)$$

In (4), λ is the *regularization parameter*. The *data constraint* $\|\mathbf{g} - \mathbf{B}\mathbf{h}\|_{\mathcal{D}}$ enforces deconvolution by penalizing greater distances between the solution image \mathbf{h} degraded by \mathbf{B} and the observed image \mathbf{g} , defined by an appropriate distance norm $\|\cdot\|_{\mathcal{D}}$. This distance norm is selected according to *a priori* information about the noise process \mathbf{n} [25]. Note that it is not necessary to know the actual statistics (the variance) of the noise process to select an appropriate norm—only the type of noise distribution is needed. For example, the l_1 norm is optimal for Laplacian-distributed noise, where the l_2 norm is the best choice for Gaussian-distributed noise [25].

The term $\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}}$ is often referred to as the *regularization operator*. This operator typically enforces smoothness of the resulting solution by penalizing high-frequency content on a local basis. However, in the approach proposed here, $\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}}$ instead assesses an energy penalty for local deviation from a *characteristic property* defined by a PIM or LIM. The characteristic property may or may not conform to the traditional definition of smoothness. In fact, special-purpose PIM’s and LIM’s may be constructed that do not enforce

smoothness. This paper attempts to broaden the definition of the regularization operator beyond standard derivative-based kernels, enabling new constraints that enforce consistency with a specified image model. With the PIM's and LIM's, regularization can be specialized to retain certain structural features.

Along with the potential advantages of the specialized regularization operators come additional difficulties in solving the image restoration problem. First, regularization operators that depend continuously on the data must be constructed for each PIM and LIM (Section III). Due to the nonlinearity of these operators, the standard methods of selecting the regularization parameter cannot be employed, and a new approach must be adopted (Section IV). Furthermore, given more than one choice for potential regularization operators, the problem of selecting the most appropriate model must be addressed (Section IV). Finally, an efficient method that generates high quality solutions to the nonconvex energy functional $E(\mathbf{h})$ containing the nonlinear regularization operator must be suggested (Section V).

III. THE PIM/LIM MODELS AND REGULARIZATION OPERATORS

We divide the models used in regularization into two classes: the PIM's and the LIM's. The PIM's differ from the LIM's in that PIM's enforce the characteristic property on a piecewise basis, while LIM's enforce the characteristic property in each local neighborhood. With PIM's, the characteristic property is not enforced across region boundaries, given regions of sufficient size. The LIM properties hold across region boundaries. In this paper, two PIM's and two LIM's are described, along with the corresponding regularization operators. The four models introduced here by no means form an exhaustive set of image models for image restoration. The two PIM's described in this paper are the *piecewise constant* (PICO) image model and *piecewise linear* (PILI) image model, and the two LIM's are the *locally monotonic* (LOMO) image model and the *locally convex/concave* (LOCO) image model.

Here, each model will be defined and a regularization operator that can be used as a constraint in regularization will be expressed. Because the models depend strongly on local interactions, it is necessary to introduce a compact notation for local orientation. With respect to $h(x, y)$, let $d(x, y, k, l)$ be the l th difference in the k th direction. The four orientations are enumerated as follows: $k = 1$ denotes the south-north direction (only change in x), $k = 2$ denotes the east-west direction (only change in y), $k = 3$ denotes the southwest-northeast direction, and $k = 4$ denotes the southeast-northwest direction. For direction k , the l th difference between successive pixels is defined by the pixel that is $l + 1$ pixels away from (x, y) in the k th direction, less the pixel that is l pixels away. For example, $d(x, y, 1, 0) = h(x + 1, y) - h(x, y)$, $d(x, y, 1, 1) = h(x + 2, y) - h(x + 1, y)$, $d(x, y, 1, -1) = h(x, y) - h(x - 1, y)$, $d(x, y, 1, 2) = h(x + 3, y) - h(x + 2, y)$, $d(x, y, 3, 1) = h(x + 2, y - 2) - h(x + 1, y - 1)$, $d(x, y, 3, -4) = h(x - 4, y + 4) - h(x - 5, y + 5)$, etc. Pixel intensities can also be denoted by the distance from a given pixel $h(x, y)$. Let $h(x, y, d, l)$ represent

the pixel intensity l pixels away from (x, y) in direction d . For example, $h(x, y, d, 0) = h(x, y)$ for any d , $h(x, y, 1, 1) = h(x + 1, y)$, $h(x, y, 1, -1) = h(x - 1, y)$, $h(x, y, 4, 1) = h(x + 1, y + 1)$, $h(x, y, 4, -2) = h(x - 2, y - 2)$, etc.

A. Piecewise Image Models for Regularization

The two PIM's discussed here are formed from the characteristic properties of *piecewise constancy* and *piecewise linearity*. Of course, many other PIM's may be envisioned, for example, piecewise polynomials, piecewise exponential, etc. However, these two PIM's are the simplest and illustrate well the principle. Using these PIM's in the restoration problem as a type of smoothness constraint implies that the image being restored obeys or nearly obeys the PICO or PILI characteristic property, which is of course, not generally justifiable. The PICO property, in particular, should be applied only to images that were originally flat nearly everywhere (except at separating boundaries); such applications certainly exist. Both the PICO and the PILI models smooth interior image regions without sacrificing edges, as the PIM properties are not enforced across region boundaries. A difference does exist—where the PILI images can contain a variety of step and ramp edges (sharp and slowly varying intensity transitions), the PICO images are restricted to steplike edges separating constant regions. The more restrictive PICO model is most useful when restoring images from man-made scenes containing steplike edges, such as with document processing or in certain automated visual inspection tasks in manufacturing. Nevertheless, for each image restoration problem, we do not make *ad hoc* decisions on the selection of the proper model. Using the cross-validation technique described in Section IV, *the most effective model can be chosen automatically*.

The PIM definitions can be given simultaneously. First, a one-dimensional (1-D) definition of piecewise constancy and piecewise linearity will be stated. Then, the 1-D definitions can be used to form two-dimensional (2-D) definitions for use in image restoration. A 1-D signal x is *piecewise constant* (*piecewise linear*) of *degree* m , or PICO- m (PILI- m) if the length of the shortest constant (affine linear) subsequence in x is greater than or equal to m . If the definition holds, then each sample is part of a constant (linear) segment of length greater than or equal to m , where m determines the spatial extent of the neighborhood and is called the *degree* of the model. The lowest degree 1-D PICO (PILI) regression of interest is PICO-2 (PILI-3), since all signals are PICO-1 (PILI-2).

The 1-D definitions are more straightforward than the 2-D definitions, since the question of orientation must be addressed in two dimensions. An image \mathbf{h} is PICO- m (PILI- m) if \mathbf{h} is PICO- m (PILI- m) along every contiguous 1-D path for a set of prescribed orientations. A two-orientation PICO (PILI) definition enforces piecewise constancy (linearity) along image columns and rows, where the four-orientation definition includes the diagonal orientations and tends to limit streaking artifacts. Our empirical studies indicate that the four-orientation definition is most appropriate for the PIM's, due to the reduction of artifacts.

1) *The Piecewise Constant Model:* For each model, it is necessary to define a function that penalizes deviation from the characteristic property of the model. This function is used in the regularization operator $\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}}$. For piecewise constancy, $q(x, y)$ should penalize the deviation of $h(x, y)$ from its neighbors. We can use the absolute value of the difference signal, $d(x, y, k, l)$, to define $q(x, y)$. In $d(x, y, k, l)$, the directions k are defined by the orientations used in the PIM. The extent of the spatial neighborhood used (the parameter l) is a function of the model order m . To construct the regularization operator, the following function is computed at each image location (x, y) :

$$q(x, y) = \sum_{k=1}^4 \left\{ \prod_{j=-m+1}^0 \left[\sum_{l=j}^{j+m-1} |d(x, y, k, l)| \right] \right\}. \quad (6)$$

Note that (6) does not induce any particular norm, since the metric is applied as the final step in the construction of $\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}}$. The penalty in (6) has four separate subpenalties corresponding to potential violation of the PICO constraint in any of the four prescribed orientations. Within the penalty for each orientation, the product terms give a separate penalty for each non-PICO subsequence of length m ; only one of the possible subsequences [containing $h(x, y)$] needs to be constant, so a product is used rather than summation. A subsequence along a given path with an error of zero forces the error for all subsequences along that path to zero, since the penalties are multiplied. In forming $q(x, y)$, the use of the product in (6) is a fundamental difference in form between the PIM's and LIM's. The LIM's force each subsequence to obey the given property and do not allow a piecewise interpretation. Although the expression in (6) may appear to be computationally expensive, only $Nm(m-1)$ addition operations and $4Nm$ multiplication operations are required for the entire N -pixel image. By contrast, for an $m \times m$ Laplacian kernel, the equivalent regularization operator requires $N(m^2-1)$ addition operations and Nm^2 multiplication operations.

2) *The Piecewise Linear Model:* Piecewise linearity may also be used to enforce smoothness on corrupted signals. In the spirit of 1-D PILI models used to model statistical data [5] and in the spirit of topological models [27], we apply the PILI models in the image restoration process. For regularization, the PILI PIM allows intraregion image smoothing without degrading intensity discontinuities. Within each image region, the PILI model enforces smoothing while retaining intensity trends, which are approximated by linear functions. When extended to two dimensions, the PILI signal is in essence a piecewise planar surface. Although more liberal than the PICO model in terms of the preservation of both step and ramp edges, the PILI is less well-suited for extreme cases of additive noise. When using the PILI model, high-amplitude noise processes often contain local groupings of outliers that approximate linear segments and are retained erroneously in regularization. However, for lower-variance noise, the PILI-constrained image restoration results are often very good.

For the PILI PIM, the regularization operator incorporates a localized linear regression. Then, $q(x, y)$ may be computed by penalizing locations where $h(x, y)$ deviates from

piecewise linearity. Consider the 1-D function with ordinate y' , representing pixel intensity, and abscissa x' , representing the distance from pixel location (x, y) in direction d . Let $\mu(x, y, d, l)$ be defined as the least squares estimate for the slope for the linear regression of the m -length subsequence in direction d and displacement l with respect to location (x, y) in the image. $B(x, y, d, l)$ is the corresponding y -intercept. Since the x' values represent the displacement from the pixel at (x, y) , the y' -intercept represents the best estimate for $h(x, y)$ to minimize the least squares error (for that particular PILI- m subsequence). The 1-D linear regression is formed by

$$\begin{aligned} \mu(x, y, k, l) &= \frac{m \sum_{i=l, i \neq 0}^{l+m-1} h(x, y, k, i) i - \left(\sum_{i=l, i \neq 0}^{l+m-1} i \right) \left[\sum_{i=l, i \neq 0}^{l+m-1} h(x, y, k, i) \right]}{m \sum_{i=l, i \neq 0}^{l+m-1} i^2 - \left(\sum_{i=l, i \neq 0}^{l+m-1} i \right)^2} \end{aligned} \quad (7)$$

$$\bar{x}'(x, y, k, l) = \frac{\sum_{i=l, i \neq 0}^{l+m-1} i}{m-1} \quad (8)$$

$$\bar{y}'(x, y, k, l) = \frac{\sum_{i=l, i \neq 0}^{l+m-1} h(x, y, k, i)}{m-1} \quad (9)$$

and

$$B(x, y, k, l) = \bar{y}'(x, y, k, l) - \mu(x, y, k, l) \bar{x}'(x, y, k, l). \quad (10)$$

In (7)–(9), $i = 0$ is excluded from the summations to remove the effect of the pixel in question, $h(x, y)$. When $h(x, y)$ is a member of a linear subsequence in the k th direction with displacement l , it should equal $B(x, y, k, l)$. Then the regularization operator is defined by

$$q(x, y) = \sum_{k=1}^4 \left(\prod_{j=-m+1}^0 \left\{ \sum_{l=j}^{j+m-1} [h(x, y) - B(x, y, k, l)] \right\} \right). \quad (11)$$

Again, four separate penalties are applied for the four defined image orientations. Within each penalty for the 1-D paths centered at (x, y) , penalties are multiplied for each non-PILI subsequence containing $h(x, y)$. The penalties are multiplied, rather than added, because the PIM only requires the pixel at (x, y) to belong to one PILI subsequence in each direction. So, an error of zero for a particular subsequence in a particular direction will eliminate all possible penalties for that direction at that location (x, y) . The inside summation in (11), $\sum_{l=j}^{j+m-1} [h(x, y) - B(x, y, k, l)]$, gives the deviation of a particular length- m subsequence from its linear regression. Note that $q(x, y)$ need not be positive, since a norm will be applied ($\|\mathbf{q}(\mathbf{h})\|_{\mathcal{M}}$). In (11), $[h(x, y) - B(x, y, k, l)]$ could be

replaced by $\text{int}[h(x, y) - B(x, y, k, l)]$ (the nearest integer) to ensure that PILI- m solutions have penalties of exactly zero in the discrete-range case.

B. Local Image Models for Regularization

Instead of relaxing the characteristic property at region boundaries, the local image models (LIM's), enforce the property at *every* point in the image equally. Therefore, the properties must be sufficiently general to allow for nontrivial solutions; constancy and linearity obviously are not such properties! Indeed, it is harder to intuit global properties that can be described in terms of meaningful local properties. Two local characteristic properties of interest are studied here: *local monotonicity* and *local convexity/concavity*. Both the LOMO LIM and the LOCO model provide meaningful measurements of smoothing for the image restoration problem. In particular, the LOMO model provides a potent, general-purpose smoothing constraint that prohibits outliers but allows edge information to be preserved in a natural way. Many signals and images may be considered to have an approximate LOMO structure; indeed, this is why the median filter is of considerable interest as a signal/image smoothing device: it attempts to create LOMO signals/images from the input. The LOCO model is much more specialized; it is effective for restoring scenes that are inherently smooth and which contain regions of sustained trends in gray-level derivatives.

The definitions of LOMO and LOCO signals in both 1-D and 2-D are quite similar, and can again be given together: A 1-D signal \mathbf{x} is LOMO- m (LOCO- m) if every subsequence of \mathbf{x} of length $\leq m$ is monotonic (is either convex or concave). Note that since every 1-D signal is LOMO-2 (LOCO-3), so that LOMO-3 (LOCO-4) is the smallest property degree of interest. An image \mathbf{h} is LOMO- m (LOCO- m) if \mathbf{h} is monotonic (is either convex or concave) on every 1-D path of length $\leq m$ along a set of prescribed orientations. Since image streaking is not a problem with the LIM's, only two-orientation versions (along the rows and columns) are used. Using only two orientations (horizontal and vertical) with a PIM of degree m , an image streak of length $m - 1$ could be created (horizontally or vertically) in the restoration process, since the PIM's are enforced on a piecewise basis. Since the LIM's will enforce the model property without boundaries, an image streak is impossible, since the pixels of each side of the streak would be in violation of the LIM.

1) *The Locally Monotonic Model:* By incorporating the property of local monotonicity into the solution of the image restoration problem, we extend the previous 1-D work by Restrepo and Bovik [25], [26] and Sidiropoulos [28], who considered the more restricted problem of signal enhancement (noise without blur). The LOMO model preserves image structure by allowing monotonically increasing/decreasing intensity trends including every variety of idealized edge pattern. At the same time, outliers from noise are regarded as violating the LOMO property. Local monotonicity is a novel, very flexible definition of smoothness for digital signals.

To evaluate local monotonicity, we note that each length m subsequence must have successive differences that do not change sign to be LOMO- m . We can construct a regularization

operator that exploits this fact. The sign skeleton of the difference signal $d(x, y, k, l)$ can be used in constructing $q(x, y)$, the penalty for the deviation from local monotonicity at location (x, y) . The regularization operator for the LOMO LIM is computed using

$$q(x, y) = \sum_{k=1}^2 \left\{ \sum_{j=-m+1}^0 \left[\sum_{l=j}^{j+m-3} \left(\sum_{i=l+1}^{j+m-2} \{1 - s[d(x, y, k, l) \cdot d(x, y, k, i)]\} p(x, y, k, l, i) \right) \right] \right\} \quad (12)$$

where

$$p(x, y, k, l, i) = \frac{1}{2} \min \{ |d(x, y, k, l)|, |d(x, y, k, i)| \} \quad (13)$$

and the $[1 - s(\cdot)]$ term generates a zero penalty for absence of sign changes in the difference signal, where $s(\cdot)$ is a sigmoid function defined by

$$s(x) = \frac{\tanh(x) + 1}{2}. \quad (14)$$

The penalty enforced in (12) is proportional to the difference that is closest to zero, so that steplike edges are not treated inequitably. The term $p(x, y, k, l, i)$ gives the minimum change needed to correct the element in the subsequence which violates monotonicity on a local basis. A 1-D monotonic signal is always possible via *flattening*—making the signal constant in the nonmonotonic regions. If the penalty were based on the magnitude of the greatest difference, then edge points would receive greater penalties than nonedge points. In (12), the two innermost summations (with indices of l and i) add penalties wherever a subsequence of length m [containing $h(x, y)$] has a sign change in its difference signal. These penalties are summed, not multiplied, for each possible subsequence containing $h(x, y)$. This shows a fundamental difference between the LIM's and PIM's—with the LIM's, each subsequence must obey the LIM rule, whereas the PIM's are enforced on a piecewise basis.

2) *The Locally Convex/Concave Model:* As with the 1-D LOCO signal model [26], the 2-D LOCO LIM may be regarded as having a somewhat restricted application domain. The idea behind the LOCO characteristic property is that smoothness can be defined as a restriction on the number of changes in convexity within local neighborhoods. However, this model precludes steplike edges and can lead to oversmoothing. But, for images that naturally have a limited high-frequency content, the LOCO LIM can be quite powerful. The example provided in Section VI demonstrates such a case.

In the formulation of (4), a regularization operator $q(\mathbf{h})$ is required that penalizes deviations from local convexity/concavity. The LOMO version of $q(\mathbf{h})$ exploited the fact that interpixel differences within a subsequence of length m could not have sign changes. The LOCO LIM can use a similar property, enforcing the differences of the interpixel differences to be of the same sign or zero. In other words, the difference signal of a LOCO signal is LOMO! This restriction

will guarantee local convexity/concavity. In this case

$$q(x, y) = \sum_{k=1}^2 \left\{ \sum_{j=m-1}^0 \left[\sum_{l=j}^{j+m-3} \left(\sum_{i=l+1}^{j+m-2} \{1 - s[(d(x, y, k, l) - d(x, y, k, l+1))(d(x, y, k, i) - d(x, y, k, i+1))]\} \right) \cdot p(x, y, k, l, i) \right] \right\} \quad (15)$$

where

$$p(x, y, k, l, i) = \frac{1}{2} \min \{ |d(x, y, k, l) - d(x, y, k, l+1)|, |d(x, y, k, i) - d(x, y, k, i+1)| \} \quad (16)$$

and $s(\cdot)$ is the sigmoid defined in (14). Here, $p(\cdot)$ has the same effect as in (12), providing a penalty that is proportional to the distance from the nearest LOCO solution for the neighborhood of (x, y) . The penalty of (15) has the same form as the LOMO $q(x, y)$, except that the local monotonicity of the difference signal is evaluated. For each direction k , the summation indexed by j gives a potential penalty for each subsequence of length m containing $h(x, y)$. The innermost summations evaluate the local convexity of the subsequence by penalizing any sign change in the difference of the difference signal. So, (15) will penalize non-LOCO subsequences and will not penalize LOCO subsequences.

IV. CHOOSING THE MODEL AND THE REGULARIZATION PARAMETER VIA CROSS-VALIDATION

Although the PIM's and LIM's provide effective application-specific smoothing constraints for image processing problems, the selection of the regularization parameter λ to be used is naturally a significant issue, owing to the unusual and highly nonlinear format of the image models.

A. Selection of the Regularization Parameter

To form the regularization solution, the proper data constraint is selected corresponding to the additive noise distribution and the blurring operator. Then, the most suitable PIM or LIM may be selected, according to the technique described below in Section IV-B. Finally, the regularization parameter λ must be assigned properly to achieve an equitable balance between the two constraints.

Under the Miller formulation of the regularization problem, as in (2) and (3), the regularization parameter can be determined *a priori* using [18]

$$\lambda = \left(\frac{\varepsilon}{E} \right)^2. \quad (17)$$

If the variance of the additive noise is known, then the bound ε can be set. E , the bound on $\|q(\mathbf{h})\|_{\mathcal{M}}$ is difficult to obtain when a nonlinear regularization operator $q(\mathbf{h})$ is used, as with the PIM/LIM regularization operators. A constrained least squares (CLS) solution is given by Hunt [10]. Given *a priori* information on the additive noise, the following equation can be used to find a solution for λ :

$$\|g - B\mathbf{h}\|^2 = \|[I - BA(\lambda)]g\|^2 = \|\mathbf{n}\|^2 = N\sigma^2 \quad (18)$$

where

$$A(\lambda) = (B'B + \lambda q'q)^{-1} B'. \quad (19)$$

When the regularization operator q is nonlinear, $(B'B + \lambda q'q)$ is not well defined. So, the straightforward CLS method is inappropriate for the nonlinear case. Furthermore, the CLS method of [10] assumes that the error from the deblurring is independent, identically distributed (i.i.d.) with the error from the noise, which may not be a reasonable assumption. One possible solution is to linearize $A(\lambda)$; however, this approach is computationally expensive for images of practical sizes.

Other methods of deriving the regularization parameter include the mean-squared error (MSE) approach and the predicted mean-squared error (PMSE) approach [6]. By definition, the MSE method seeks to minimize

$$E[|e(\lambda)|^2] = \|\mathbf{z}\|^2 + E[\|\mathbf{h}\|^2] - 2E[\|\mathbf{z}'\mathbf{h}\|^2] \quad (20)$$

which reduces to minimizing $E[\|\mathbf{h}\|^2] - 2E[\|\mathbf{z}'\mathbf{h}\|^2]$, since \mathbf{z} does not depend on λ . When g, B , and q are known, the expectation operators can be dropped on $E[\|\mathbf{h}\|^2]$ and $E[\|\mathbf{z}'\mathbf{h}\|^2]$. It has been shown in [6] that minimizing (20) is equivalent to minimizing

$$\|q^{-1}[I - BA(\lambda)]^{3/2}g\|^2 = \sigma^2 \text{tr} \{q^{-2}[I - BA(\lambda)]^2\}. \quad (21)$$

The PMSE approach is similar to the MSE approach, except that $E[\|Be(\lambda)\|^2]$ is minimized. Galatsanos and Katsaggelos [6] show that this is equivalent to minimizing

$$\|[BA(\lambda)]^{1/2}[I - BA(\lambda)]g\|^2 = \sigma^2 \text{tr} \{BA(\lambda)[I - BA(\lambda)]\}. \quad (22)$$

Given a linear regularization operator q , the regularization parameter can be determined using (21) or (22). This is also prohibited when q is nonlinear, since $A(\lambda)$ is not well defined.

Reeves and Mersereau explored the use of "variable" regularization operators in [24]. Their application of *cross-validation* to the selection of a regularization parameter provides a suitable approach for the nonlinear regularization operators here. In the cross-validation process, the image is divided into an estimation set and a validation set. The pixels in the estimation set are used to obtain a result for a given set of parameters. At the same time, the validation set is used to evaluate the effectiveness of the parameter set and model used for regularization. The cross-validation method is attractive because the original uncorrupted, unblurred image is not needed to estimate the regularization parameter. Likewise, image prototypes and information on the corruptive noise process are not required in cross-validation.

First, the image is divided into P sets. Then, the parameter or model being evaluated can be imposed on all sets, except the chosen validation set. Let Q represent the set of all image coordinates (x, y) . The P sets are then given by the union: $\cup_{p=1}^P Q_p = Q$. Note that membership in the P sets is typically selected using a uniformly-distributed random number generator. In this way, the local dependencies between members of the same set should be minimized. An error measure computed using the validation set can then be

employed to test the quality of the solution derived with the given parameters. In traditional cross-validation, this process can be repeated P times, each time choosing a different set as the validation set, thus using each pixel for both estimation and validation. The drawback is that the image restoration problem has to be solved P times to complete the evaluation of P validation sets.

Let Q_v be the validation set. To perform cross-validation (using the l_2 norm case for both the data constraint and the smoothing constraint), the following energy functional is then minimized:

$$E^v(\mathbf{h}) = \sum_{(x,y) \in Q - Q_v} [g(x,y) - (\mathbf{B}\mathbf{h})_{(x,y)}]^2 + \lambda \|\mathbf{q}(\mathbf{h})\|^2 \quad (23)$$

where $(\mathbf{B}\mathbf{h})_{(x,y)}$ is the value of $\mathbf{B}\mathbf{h}$ at location (x,y) . Note that the data constraint (and therefore the regularization parameter) has no influence on the solution for the validation set pixels. Let the solution that minimizes (23) for validation set v be defined as \mathbf{h}_v . To reduce the computational burden of computing the cross-validation error (CV error), we employ the data division technique of [23]. In this case, only one validation set is used and only one minimization of (23) is required for each parameter value that is tested. If Q_v is the validation set, then (23) is minimized to compute \mathbf{h}_v , and the CV error is given by

$$V(\lambda) = \frac{1}{N} \sum_{(x,y) \in Q_v} [g(x,y) - (\mathbf{B}\mathbf{h}_v)_{(x,y)}]^2 \quad (24)$$

for an N -pixel image. Reeves and Mersereau indicate that the use of ten sets (so that the validation set contains 1/10 of the pixels) is a sufficient sampling to yield a reliable result [23]. Our experiments verify this assumption. Note that the data division validation error does not have to equal the full validation error, as long as both functions are minimized by approximately the same value of λ . So, the appropriate regularization parameter for the nonlinear PIM's/LIM's can be computed using (24). The selection of the regularization parameter, in this case, does not depend on *a priori* knowledge.

B. Selection of the PIM/LIM for Regularization

With the newly introduced classes of PIM's and LIM's and the possible introduction of other application specific models into these classes, the ability to select the proper model for a given image restoration problem is critical to the success of this paradigm. Previously, these models have been selected on a trial-and-error basis with only intuitive assumptions about the original image structure in mind [2]. Using cross-validation, the uncertainty of selecting the model heuristically can be improved upon.

A model is selected by finding the PIM or LIM that yields the lowest validation error. This method essentially allows the characteristic property assumption of each PIM or LIM to be tested on the pixels that are not members of the validation set. Of course, the selection of the regularization parameter still influences the quality of the result. In fact, the model and the regularization parameter may be selected

TABLE I
IMAGE EXAMPLES

| Example | Blurring Kernel | Dist. of Noise | σ | SNR |
|-----------|-----------------|----------------|----------|------|
| Cameraman | Uniform (1x9) | Gaussian | 10.0 | 15dB |
| Pebbles | Uniform (1x7) | Gaussian | 7.0 | 15dB |
| Books | Uniform (3x3) | Gaussian | 10.0 | 15dB |
| Straw | Uniform (1x11) | Laplacian | 6.7 | 10dB |

simultaneously, using a validation error that depends both on the model selected and on the regularization parameter. Let the characteristic set \mathcal{C} represent the model selected and the regularization operator $\mathbf{q}_{\mathcal{C}}(\cdot)$ enforce the corresponding smoothing constraint. For each model and regularization parameter tested, (23) is minimized to find \mathbf{h}_v , substituting the proper $\mathbf{q}_{\mathcal{C}}(\mathbf{h})$ for $\mathbf{q}(\mathbf{h})$. Then the data division CV error in (24) becomes a function of \mathcal{C} and λ . Remember that the regularization operator $\mathbf{q}_{\mathcal{C}}(\cdot)$ does depend on the model order m , which determines the spatial extent of the PIM or LIM. So, it is necessary to evaluate (24) for different model orders as well as over different models. It is our experience that the model degree $m > 4$ typically leads to a less desirable (oversmoothed) result. Therefore, only one or two reasonable model degrees need to be evaluated for each model.

V. MINIMIZATION OF THE REGULARIZATION FUNCTIONAL

The standard linear regularization operators can be constrained so that a convex optimization problem is formed [12], [14]. With a convex energy topology, simple gradient descent solutions may be formed to achieve restoration. The need for convexity is one of the major reasons that the Laplacian has often been used as the regularization operator. In the regularization approach presented here, there is no restriction to convex energy functions. With the nonlinear PIM and LIM constraints, convexity does not generally hold, hence steepest descent approaches to minimization will not avoid local minima in the regularization functional. Therefore, special attention has been given in this research to the optimization technique applied.

For a nonconvex optimization problem, stochastic simulated annealing (SA) can be used to locate a globally optimal solution [7]. However, the simulated annealing algorithm, even with a fast geometric annealing schedule [1], cannot be used for real-time or near real-time application. To limit the computational expense of image restoration, we minimize the regularization functional via *generalized deterministic annealing* (GDA), a combinatorial optimization technique that provides high-quality solutions for time-critical applications and which can be easily implemented on a parallel architecture [3]. GDA provides rapid, guaranteed convergence and the ability to escape undesirable local solutions, and hence is well suited for these problems.

GDA essentially estimates the stationary distribution of the SA transition probabilities at each annealing temperature using a mean field approach. The probability density for intensity l at location (x,y) is updated via (25), shown at the bottom of the next page, where I is the number of possible solutions for each pixel intensity (I is usually a limited subset of intensities centered at the current solution, e.g., $I = 9$) and

TABLE II
PIM/LIM IMAGE RESTORATION RESULTS

| Image | PIM/LIM | CV Error | MSE | ISNR |
|-----------|---------|----------|-------|-------|
| Cameraman | PICO | 127.1 | 390.3 | +1.4 |
| | PILI | 116.6 | 377.5 | +1.6 |
| | LOMO | 115.5 | 285.8 | +2.8 |
| | LOCO | 119.1 | 327.3 | +2.6 |
| Pebbles | PICO | 64.3 | 141.6 | +2.4 |
| | PILI | 63.2 | 149.4 | +2.1 |
| | LOMO | 64.5 | 151.7 | +1.9 |
| | LOCO | 60.8 | 111.6 | +3.4 |
| Books | PICO | 126.7 | 147.5 | +1.6 |
| | PILI | 128.4 | 207.8 | -0.05 |
| | LOMO | 128.1 | 192.6 | +0.4 |
| | LOCO | 127.5 | 180.0 | +0.6 |
| Straw | PICO | 56.8 | 298.7 | +3.2 |
| | PILI | 56.5 | 324.5 | +2.9 |
| | LOMO | 56.1 | 317.1 | +2.9 |
| | LOCO | 53.7 | 266.0 | +3.8 |

T is the annealing temperature (discussed below). The energy functional value $E(\mathbf{h}; h(x, y) = l)$ is computed by setting all pixel intensities to their mean field values (the expected value given the discrete probability densities) and then setting the pixel intensity at location (x, y) to $h(x, y) = l$.

The annealing process is started at a sufficiently high temperature T where all probabilities are equal (all $f_{x,y}(l) \approx (1/I)$). The temperature is reduced in a geometric manner until a singular distribution emerges at each pixel site [$f_{x,y}(l) \approx 1$ for some intensity l at each (x, y)]. The number of updates per temperature depends on Γ , the maximum possible change in the energy functional for a change in one pixel intensity, and on I , the number of possible solutions for each pixel. An upper bound on the number of iterations needed at temperature T to sufficiently estimate the stationary distribution is [3]

$$\tau(T) = \frac{\ln\left(\frac{1}{I^2}\right)}{\ln\left[\left(\frac{I-2}{2I}\right)e^{-T/\Gamma} + \frac{1}{2}\right]}. \quad (26)$$

This bound guarantees uniform convergence of each probability density within a bound of $1/I^2$.

Using these guidelines for GDA, image restoration with GDA requires $O(NI)$ updates using (25) for an image with N pixels at each temperature. In contrast, the SA algorithm needs $O(I^{2N})$ updates at each temperature. In experiments performed on Sun Sparc Ultra I, image restoration using a practical SA algorithm required over 9 h, while the GDA algorithm required 13 min for the same improvement in signal-to-noise ratio (SNR). Of course linear techniques (that use the fast Fourier transform, or FFT) will be vastly superior in

computational expense. For example, the CLS method uses only 18 s of processing time on the same architecture.

VI. RESULTS AND CONCLUSIONS

Here, four example image restoration cases have been provided, one for each PIM and LIM presented in the paper. For each image restoration case, the best results for each of the four models is reported. In the first case, the results are also compared to those given by a line process approach and by a linear CLS approach. The image restoration experiments use a variety of images obtained using different blurring processes and different additive noise processes, in order to demonstrate the flexibility of the PIM/LIM approach (see Table I). The table lists the image used, the blurring kernel applied in convolution, and the details of the additive noise process including the distribution, the standard deviation (σ) of the additive noise, and the SNR. The results shown reflect the optimal regularization parameter as computed via cross-validation.

Although attaining a high SNR does not necessarily imply that the subjective quality of the restored image is superior [11], [21], we seek results that at least improve, and that do not degrade the SNR. The SNR of a noisy image is defined by

$$\text{SNR} = 10 \cdot \log_{10}(\sigma_i^2 / \sigma_n^2) \quad (27)$$

where σ_i^2 is the variance of the blurred but uncorrupted image and σ_n^2 is the variance of the noise. To measure the quantitative improvement in the SNR, the improvement in SNR (ISNR) is calculated for each result as

$$\text{ISNR} = 10 \log_{10} \frac{\|\mathbf{g} - \mathbf{i}\|^2}{\|\hat{\mathbf{i}} - \mathbf{i}\|^2} \quad (28)$$

$$f_{x,y}(l) \leftarrow \frac{1}{I} \sum_{k=0}^{I-1} \left\{ \frac{1}{1 + \exp\left[\frac{E(\mathbf{h}; h(x, y) = l) - E(\mathbf{h}; h(x, y) = k)}{T}\right]} \right\} [f_{x,y}(k) + f_{x,y}(l)] \quad (25)$$

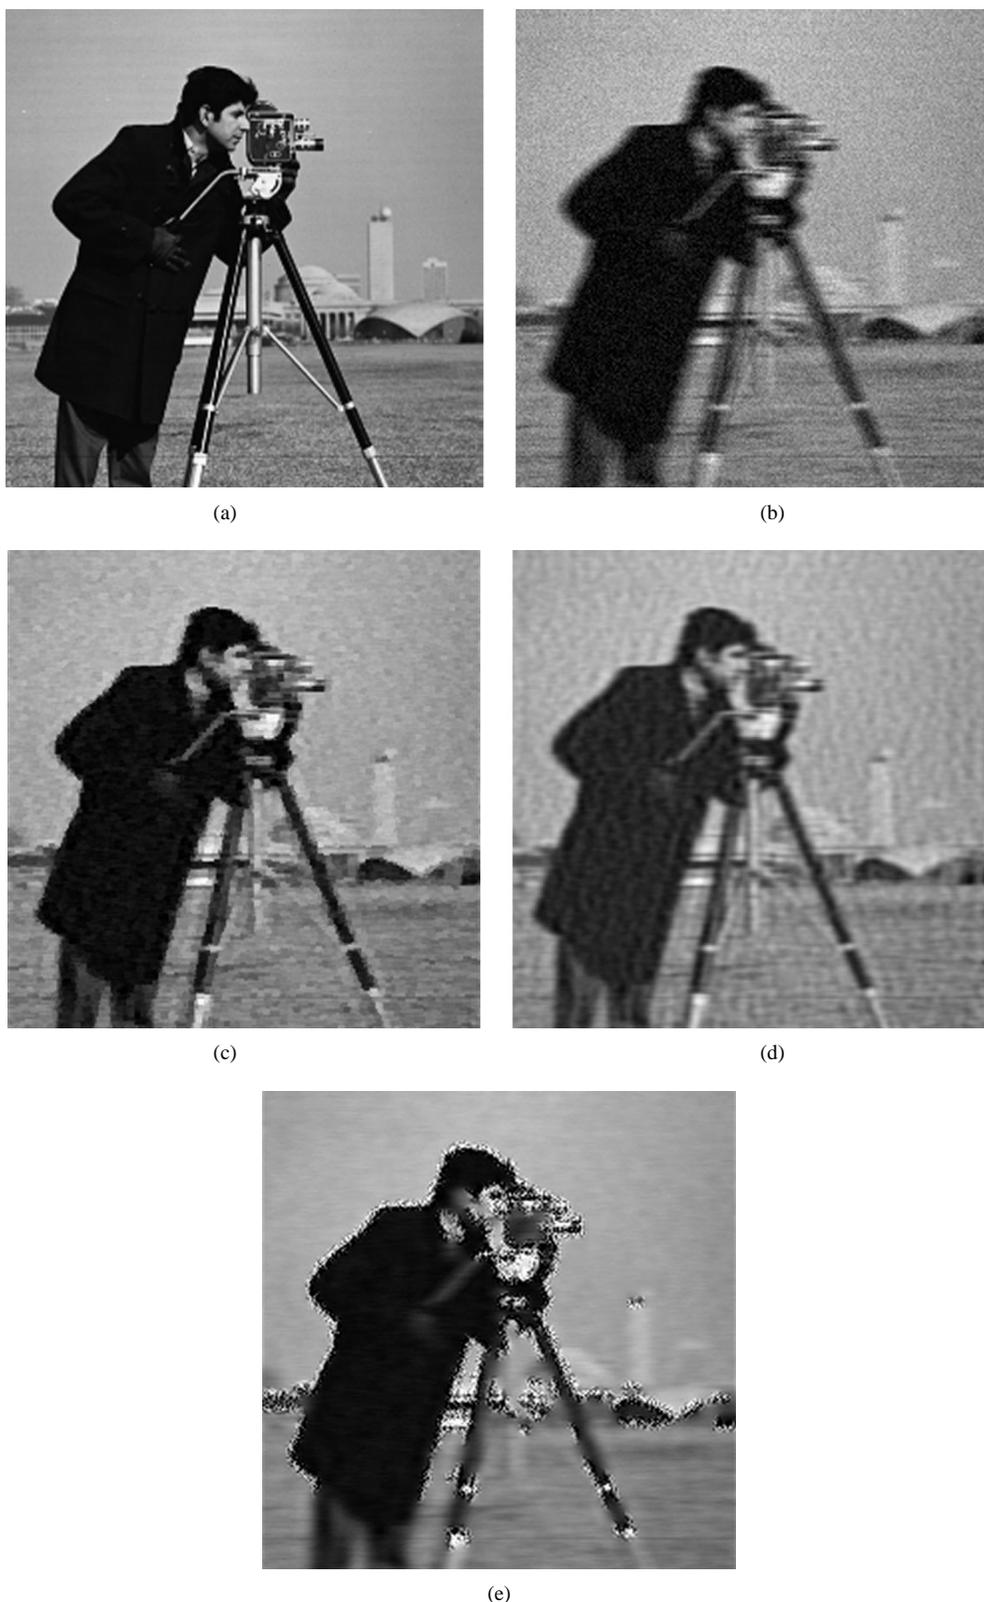


Fig. 1. Example of the LOMO LIM. (a) Original cameraman image. (b) Corrupted image (1×9 uniform blur with $\text{SNR} = 15$ dB Gaussian noise). (c) Restoration using the LOMO-3 regularization operator. (d) Restoration using CLS. (e) Restoration using line process model.

where \mathbf{g} is again the blurred, corrupted image, \mathbf{z} is the original image, and $\hat{\mathbf{z}}$ is the estimate produced by the image restoration algorithm.

For each image example, the CV error, MSE, and ISNR for each of the four PIM's and LIM's are tabu-

lated in Table II. Figs. 1–4 provide the original images, the blurred, noisy images, and the image restoration results for each case. Figs. 4–8 depict the selection of the regularization parameter using the cross-validation technique.

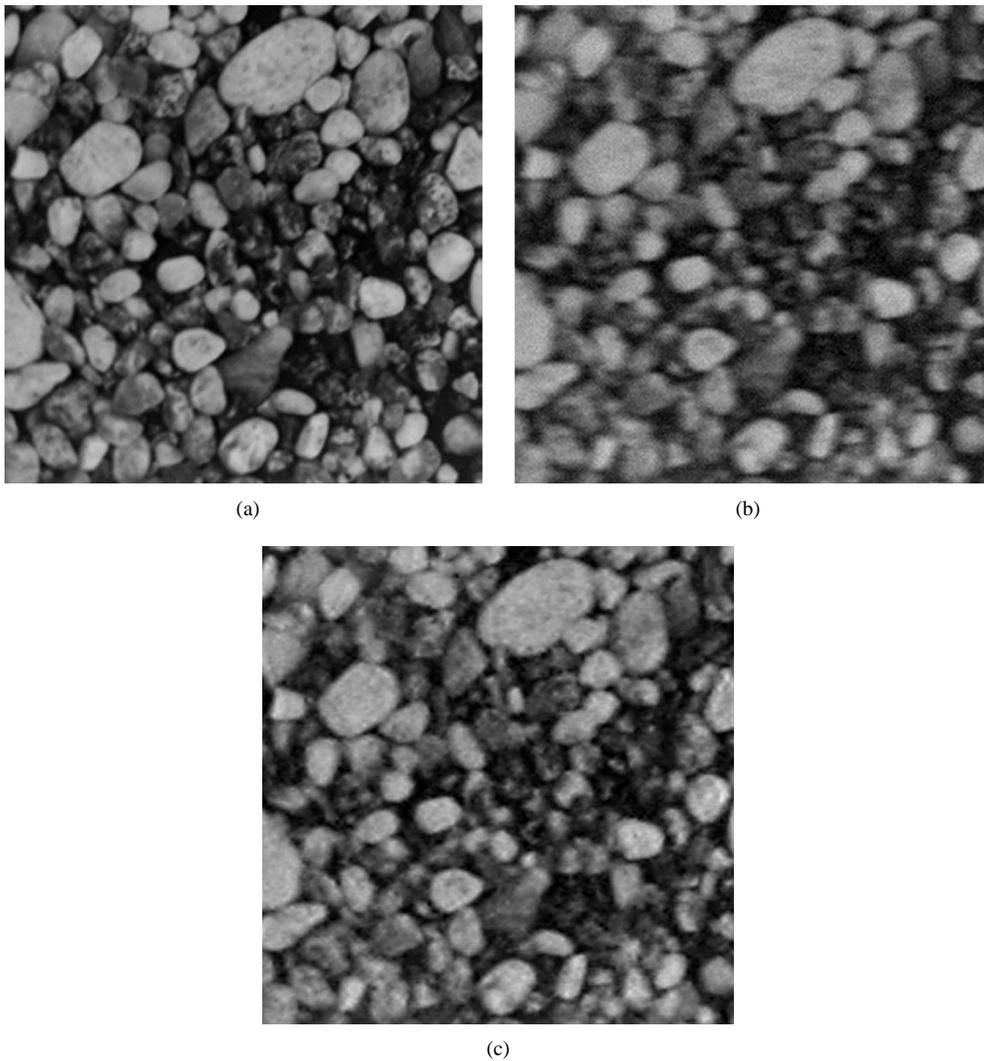


Fig. 2. Example of the LOCO LIM. (a) Original pebbles image. (b) Corrupted image (1×7 uniform blur with $\text{SNR} = 20$ dB Gaussian noise). (c) Restoration using the LOCO-4 regularization operator.

For complex natural images, the LOMO LIM provides a robust measure of smoothness. Using the LOMO-3 regularization operator, the cameraman image is restored after motion blur and severe degradation (Fig. 1). The 2.8 dB improvement in SNR noted in Table II is provided along with preservation of the sharp edges between the cameraman and the background. Among the four PIM’s and LIM’s, the LOMO model gives the lowest CV error, the lowest MSE, and the highest ISNR. Even with an original SNR of 15 dB, restoration is possible with the versatile LOMO model. As with each experiment performed, the regularization parameter was selected using the cross-validation technique. From Fig. 5, it may be observed that the CV error has a minimum value for the same λ as the true MSE of the result. This was confirmed for each PIM/LIM restoration experiment performed. So, without any statistical information on the image or on the degradation processes, the regularization parameter can be reliably chosen for the nonlinear regularization operators.

The LOMO LIM restoration results (using GDA) are compared to an FFT-based CLS method. In the frequency domain,

we have the following problem: $G(u, v) = B(u, v)I(u, v) + N(u, v)$. The CLS technique minimizes the response of the Laplacian to the solution, giving a restored image $H(u, v)$:

$$H(u, v) = \left[\frac{B^*(u, v)}{|B(u, v)|^2 + \lambda L(uv)} \right] G(u, v) \quad (29)$$

where $L(u, v)$ is the discrete Fourier transform (DFT) of the spatial domain Laplacian kernel, and λ is an unknown regularization parameter. λ is selected by iterating on (29) until the minimum value from the data constraint of (4) is discovered. Although the CLS method is computationally inexpensive, the edges in the original cameraman image have not been retained in the result [see Fig. 1(d)], and oversmoothing and ringing are evident.

To compare the LIM result to another nonlinear approach, a “line process” model [7] has been implemented. The line process uses the local image variance to control smoothing in the regularization process. When the variance exceeds a threshold, the regularization operator is not enacted. The

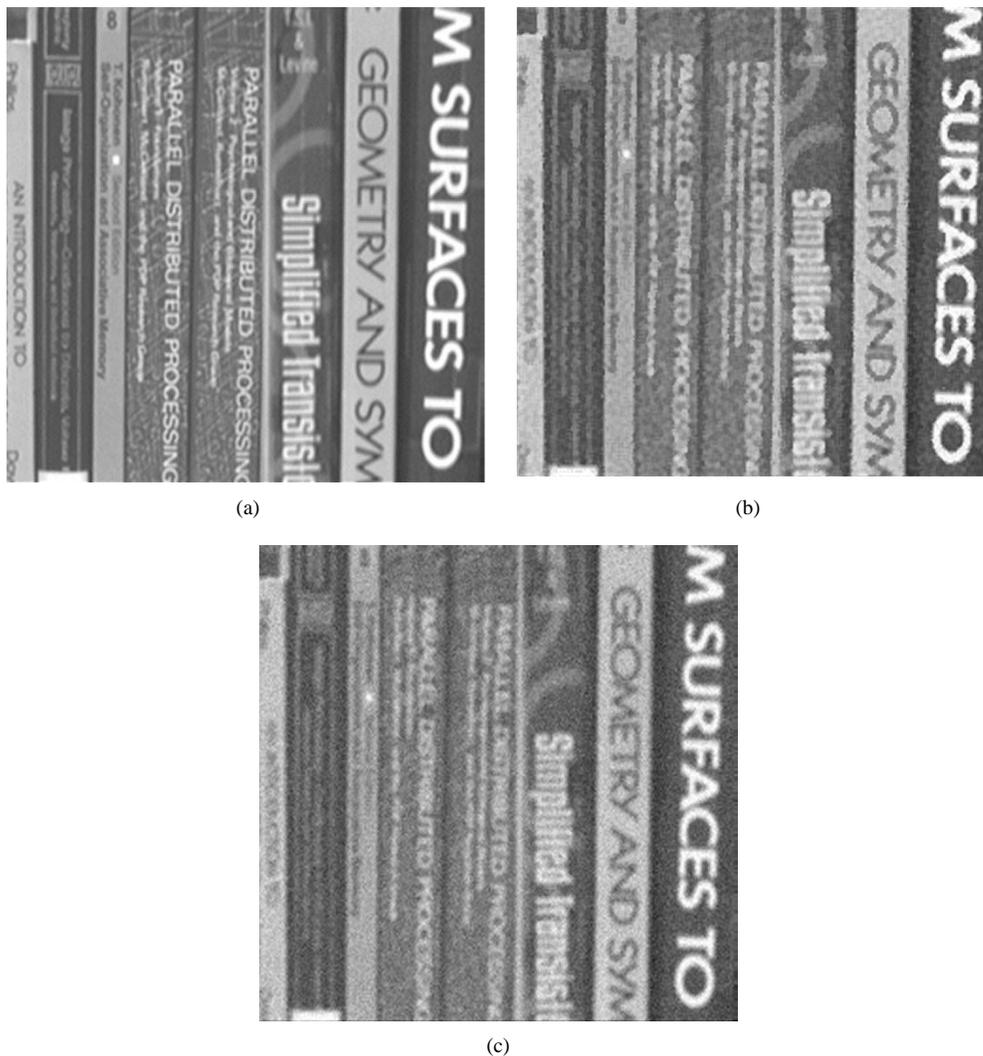


Fig. 3. Example of the PICO PIM. (a) Original “books” image. (b) Corrupted image (3×3 uniform blur with SNR = 15 dB Gaussian noise). (c) Restoration using the PICO-3 regularization operator.

image shown in Fig. 1(b) has been restored using the energy functional in (10)–(12) of [23]. The energy functional was minimized using simulated annealing. Although requiring over $30\times$ the computational resources as the GDA technique, the simulated annealing result with the line process model could not match the visual quality of the LOMO LIM. Near edges, the line process model has the tendency to retain outliers due to noise [see Fig. 1(e)].

The smooth stones in the pebbles image are a prime example of a locally convex/concave image [Fig. 2(a)]. The outliers from the 15 dB Gaussian noise are suppressed and the degradation from the motion blur is reversed in the regularization process, using the LOCO-4 regularization operator [Fig. 2(c)]. The MSE is halved and a 3.4 dB increase in the SNR is yielded (Table II), which represents the best results of the four models. The CV error closely follows the actual MSE behavior, allowing straightforward selection of the regularization parameter using cross-validation (Fig. 6).

The books image in Fig. 3(a) is a natural selection for the PICO model, since the surfaces in the scene are nearly

piecewise constant; so it is not surprising that the PICO model yields the lowest possible CV error of the four PIM’s/LIM’s presented here (Table II). After convolving the original image with a 3×3 uniform blurring kernel and adding 15 dB Gaussian noise, restoration using the PICO-3 regularization operator provided a sharp, smoothed image result [see Fig. 3(c)] without any *a priori* knowledge of the signal or noise statistics.

Fig. 4 furnishes an example of the PILI regularization operator. Although the observed image [Fig. 4(b)] is severely degraded by motion blur and 10 dB impulse noise, the linear straw pattern is recovered using the PILI PIM. This example also demonstrates the versatility of the overall paradigm, yielding success in the presence of Laplacian-distributed additive noise. Here, an l_1 norm is applied to the data constraint for optimality in the Laplacian case [25]. In the other examples presented, the l_2 norm is applied to optimally reduce the Gaussian-distributed noise. The PILI result for the “straw” image allows a 2.9 dB improvement in SNR (see Table II), while also preserving both ramp and step

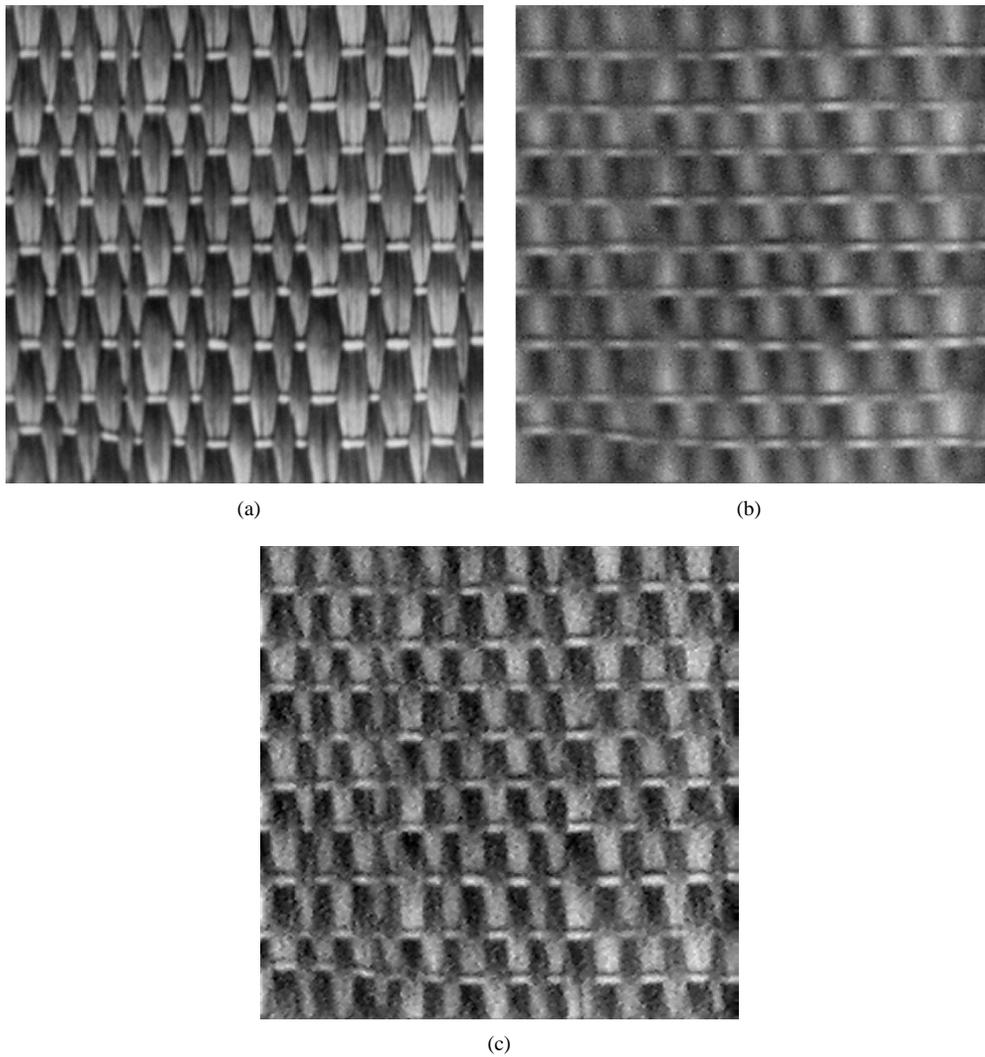


Fig. 4. Example of the PILI PIM. (a) Original straw image. (b) Corrupted image (1×11 uniform blur with SNR = 10 dB Laplacian noise). (c) Restoration using the PILI-3 regularization operator.

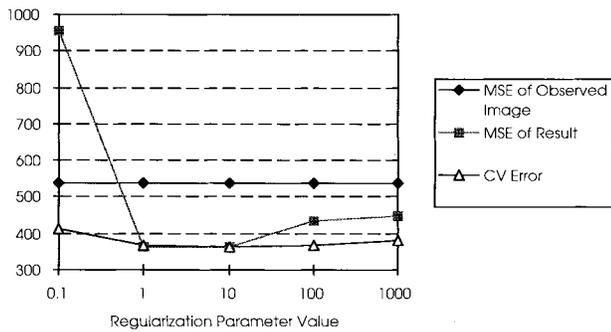


Fig. 5. Graph of MSE and CV error versus regularization parameter λ for LOMO restoration example in Fig. 1.

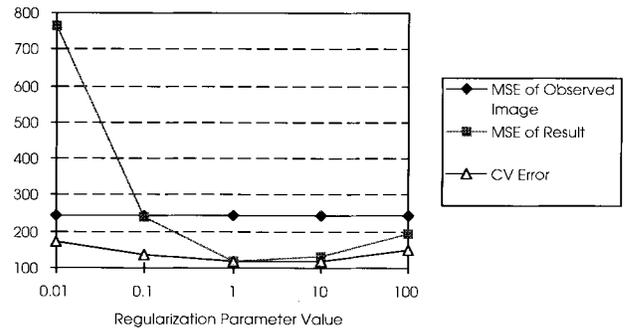


Fig. 6. Graph of MSE and CV error versus regularization parameter λ for LOCO restoration example in Fig. 2.

edges in the restoration result. But, as Table II indicates, the PILI PIM does not provide the lowest CV error of the four models. We speculate that the PILI model produces a higher CV error (and MSE) as compared to the LOCO model, for example, because the PILI model retains linear

groupings of outliers due to noise, as may be observed in Fig. 4(d).

The models of the PIM and LIM classes allow application-specific image restoration. The PICO, PILI, LOMO, and LOCO regularization operators extend the structure-preserving

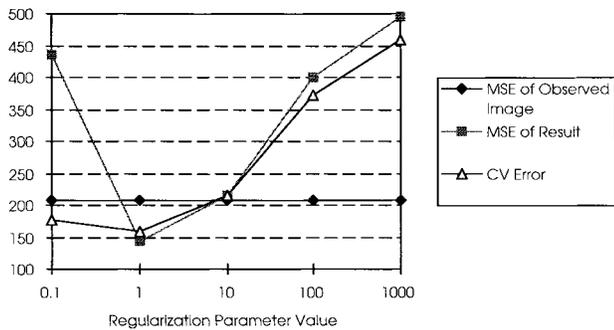


Fig. 7. Graph of MSE and CV error versus regularization parameter λ for PICO restoration example in Fig. 3.

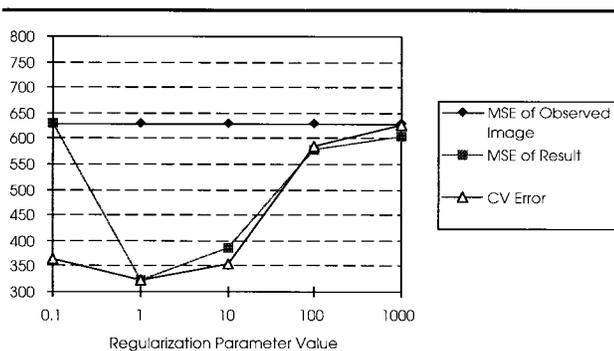


Fig. 8. Graph of MSE and CV error versus regularization parameter λ for PILI restoration example in Fig. 4.

abilities of the traditional regularization approach. For a given restoration problem, the most effective PIM or LIM can be selected via cross-validation, and the regularization parameter can also be determined, without knowledge of the original uncorrupted image. With the regularization functional and parameter in hand, a restored image can be computed rapidly using the GDA optimization technique.

In future studies, we plan to evaluate spatially adaptive regularization weighting as in [12] in conjunction with the PIM and LIM regularization operators. Although cross-validation has proven successful in determining the regularization parameter and selecting the proper model, further improvements to the computational efficiency of the cross-validation implementation is desired. We are also examining other potentially salient PIM's and LIM's, as well as a method to automatically construct the model for prototypical data.

REFERENCES

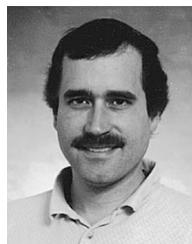
- [1] E. H. L. Aarts and J. Korst, *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. New York: Wiley, 1987.
- [2] S. T. Acton and A. C. Bovik, "Nonlinear image estimation using piecewise and local image models," *IEEE Trans. Image Processing*, vol. 7, pp. 979–991, July 1998.
- [3] ———, "Generalized deterministic annealing," *IEEE Trans. Neural Networks*, vol. 7, pp. 686–699, 1996.
- [4] G. Bilbro and W. E. Snyder, "Range image restoration using mean field annealing," *Advances in Neural Information Processing Systems I*, D. S. Touretzky, Ed. San Mateo, CA: Morgan Kaufmann, 1989.
- [5] F. L. Bookstein, "On a form of piecewise linear regression," *Amer. Stat.*, vol. 29, pp. 116–117, 1975.
- [6] N. P. Galatsanos and A. K. Katsaggelos, "Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation," *IEEE Trans. Image Processing*, vol. 1, pp. 322–336, 1992.
- [7] D. Geman and S. Geman, "Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-6, pp. 721–741, 1984.
- [8] P. Hall and D. Titterton, "Common structures of techniques for choosing smoothing parameters in regression problems," *J. R. Stat. Soc. B*, vol. 49, pp. 184–198, 1987.
- [9] H. Hiriyannaiah, G. Bilbro, W. Snyder, and R. C. Mann, "Restoration of piecewise-constant images by mean-field annealing," *J. Opt. Soc. Amer.*, vol. 6, no. 12, 1989.
- [10] B. R. Hunt, "The application of constrained least squares estimation to image restoration," *IEEE Trans. Comput.*, vol. C-22, pp. 805–812, 1973.
- [11] Y. Ichioka and N. Nakajima, "Iterative image restoration considering visibility," *J. Opt. Soc. Amer.*, vol. 71, pp. 983–988, 1983.
- [12] M. G. Kang and A. K. Katsaggelos, "A generalized formulation of the weighted smoothing functional for regularized image restoration," in *Proc. IEEE Int. Conf. Image Processing*, Austin, TX, Oct. 1994, pp. 695–699.
- [13] ———, "Simultaneous iterative image restoration and evaluation of the regularization parameter," *IEEE Trans. Signal Processing*, vol. 40, pp. 2329–2334, 1992.
- [14] ———, "General choice of the regularization functional in regularized image restoration," *IEEE Trans. Image Processing*, vol. 4, pp. 594–602, 1995.
- [15] A. K. Katsaggelos and M. G. Kang, "Iterative evaluation of the regularization parameter in regularized image restoration," *J. Vis. Commun. Image Represent.*, vol. 3, pp. 446–455, 1992.
- [16] A. K. Katsaggelos, Ed., *Digital Image Restoration*. Berlin, Germany: Springer-Verlag, 1991.
- [17] ———, "Iterative image restoration algorithms," *Opt. Eng.*, vol. 28, pp. 735–748, 1989.
- [18] R. L. Legendijk, J. Biemond, and D. E. Boeke, "Regularized iterative image restoration with ringing reduction," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, pp. 1804–1887, 1988.
- [19] K. Miller, "Least-squares method for ill-posed problems with a prescribed bound," *SIAM J. Math. Anal.*, vol. 1, pp. 52–74, 1970.
- [20] J. K. Paik and A. K. Katsaggelos, "Image restoration using a modified Hopfield network," *IEEE Trans. Image Processing*, vol. 1, no. 1, 1992.
- [21] S. S. Rajala and R. J. P. deFigueiredo, "Adaptive nonlinear image restoration by a modified Kalman filtering approach," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-29, pp. 1033–1042, 1981.
- [22] S. J. Reeves and A. C. Higdon, "Perceptual evaluation of the mean-square error choice of regularization parameter," *IEEE Trans. Image Processing*, vol. 4, pp. 107–111, 1995.
- [23] S. J. Reeves and R. M. Mersereau, "Automatic assessment of constraint sets in image restoration," *IEEE Trans. Image Processing*, vol. 1, pp. 119–122, 1992.
- [24] ———, "Optimal estimation of the regularization parameter and stabilizing functional for regularized image restoration," *Opt. Eng.*, vol. 29, pp. 446–454, 1990.
- [25] A. Restrepo (Palacios) and A. C. Bovik, "Locally monotonic regression," *IEEE Trans. Signal Processing*, vol. 41, pp. 2796–2810, 1993.
- [26] A. Restrepo (Palacios), "Locally monotonic regression and related techniques for signal smoothing and shaping," Ph.D. dissertation, Univ. Texas at Austin, Austin, TX, May 1990.
- [27] C. P. Rourke and B. J. Sanderson, *Introduction to Piecewise Linear Topology*. Berlin, Germany: Springer-Verlag, 1981.
- [28] N. Sidiropoulos, "Fast digital locally monotonic regression," *IEEE Trans. Signal Processing*, vol. 45, pp. 389–395, 1997.
- [29] A. M. Thompson, J. C. Brown, J. W. Kay, and D. M. Titterton, "A study of methods of choosing the smoothing parameter in image restoration by regularization," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, pp. 326–339, 1991.
- [30] A. N. Tikhonov and V. Y. Arsenin, *Solution of Ill-Posed Problem*. Washington, DC: Winston and Sons, 1977.
- [31] J. Zerubia and R. Chellapa, "Mean field approximation using compound Gauss-Markov random field for edge detection and image restoration," in *Proc. 1990 IEEE Int. Conf. Acoust., Speech, Signal Processing*, Albuquerque, NM, Apr. 1990, pp. 2193–2196.



Scott T. Acton (S'89–M'93–SM'99) received the B.S. degree in electrical engineering from the Virginia Polytechnic Institute and State University, Blacksburg, in 1988, and the M.S. and Ph.D. degrees in electrical engineering from the University of Texas at Austin in 1990 and 1993, respectively.

He has worked in industry for AT&T, Oakton, VA, MITRE Corporation, McLean, VA, and Motorola, Inc., Phoenix, AZ. Currently, he is an Associate Professor in the School of Electrical and Computer Engineering, Oklahoma State University (OSU), Stillwater, where he directs the Oklahoma Imaging Laboratory. The laboratory is sponsored by several organizations, including the Army Research Office (ARO), NASA, and Lucent Technologies. His research interests include multiscale image representations, diffusion algorithms, image morphology, and image restoration.

Dr. Acton was named the Halliburton Outstanding Young Faculty Member at OSU. He is an active participant in the IEEE, ASEE, SPIE, and Eta Kappa Nu. He is an associate editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING. He is the recipient of the 1996 Eta Kappa Nu Outstanding Young Electrical Engineer Award, a national award given annually since 1936.



Alan Conrad Bovik (S'80–M'81–SM'89–F'96) received the B.S., M.S., AND Ph.D. degrees in electrical and computer engineering in 1980, 1982, and 1984, respectively, all from the University of Illinois, Urbana-Champaign.

He is currently the General Dynamics Endowed Fellow and Professor in the Department of Electrical and Computer Engineering, The University of Texas at Austin, and Associate Director of the Center for Vision and Image Sciences. During the spring of 1992, he held a visiting position in the Division of Applied Sciences, Harvard University, Cambridge, MA. He is a registered Professional Engineer in the State of Texas and is a frequent consultant to industry and academic institutions. His current research interests include digital video, image processing, wavelets, 3-D microscopy, and computational aspects of biological visual perception. He has published over 250 technical articles in these areas and holds two U.S. patents.

Dr. Bovik is a recipient of the IEEE Signal Processing Society Meritorious Service Award (1998), the University of Texas Engineering Foundation Halliburton Award, and is a two-time Honorable Mention winner of the International Pattern Recognition Society Award for Outstanding Contribution (1988 and 1993). He has been involved in numerous professional society activities. Currently, he is serving on the Board of Governors of the IEEE Signal Processing Society. He is the Editor-in-Chief of the IEEE TRANSACTIONS ON IMAGE PROCESSING, and a member of the Editorial Boards of THE PROCEEDINGS OF THE IEEE and *Pattern Recognition*. He was the Founding General Chairman of the First IEEE International Conference on Image Processing, Austin, TX, November 1994.