An Efficient MRF Image-Restoration Technique Using Deterministic Scale-Based Optimization

Murali M. Menon

■ **A method for performing piecewise smooth restorations on images corrupted with high levels of noise has been developed. Based on a Markov Random Field (MRF) model, the method uses a neural network sigmoid nonlinearity between pixels in the image to produce a restoration with sharp boundaries while providing noise reduction. The model equations are solved with the Gradient Descent Gain Annealing (GDGA) method—an efficient deterministic search algorithm that typically requires fewer than 200 iterations for image restoration when implemented as a digital computer simulation. A novel feature of the GDGA method is that it *automatically* develops an annealing schedule by adaptively selecting the scale step size during iteration. The algorithm is able to restore images that have up to 71% of their pixels corrupted with non-Gaussian sensor noise. Results from simulations indicate that the MRF-based restoration remains useful at signal-to-noise ratios 5 to 6 dB lower than with the more commonly used median-filtering technique. These results are among the first such quantitative results in the literature.**

AN IMAGE-RESTORATION METHOD that reduces noise while preserving naturally occurring boundaries in a scene is presented. The method is useful as a preprocessor to enhance the performance of automatic target-recognition systems.

Target recognition is a process that can involve many stages, including measurement, preprocessing, detection, segmentation, feature extraction, and classification. For adequate recognition performance in a noisy environment, it is often important that the preprocessing stage be capable of restoring measured images. (We justify this statement in the section "Simulation Results.") The restoration should reduce the variability in the scene that results from measurement noise and clutter while preserving important features that make targets separable in the classification stage. Both objectives can be accomplished by using prior statistical knowledge of the measurement process and the clutter in the scene, or by using an empirical formulation of the desired restoration. (Details of using either a statistical or empirical formulation are contained in the following section.)

Using the latter approach, the work described in this article is based on an empirical image-restoration model that requires nearest neighbor pixels to have similar values (smoothing), without losing fidelity to the original measurement. The pixel interaction of the model smooths small pixel differences, but allows large differences to remain as a discontinuity (edge). If detailed statistical information concerning the measurement and scene is available, the information can be quantitatively incorporated into the

image-restoration model.

This article describes an image-restoration model that is based on a neural network formulation using Markov Random Fields (MRF), as described in the box, "Markov Random Fields." In the model, a neural network sigmoid function provides pairwise pixel interaction potentials. The function behaves quadratically for small differences but saturates for large differences. The MRF property of the model allows an image to be, in effect, decoupled into a large number of connected local neighborhoods, each of which can be processed independently. The local-neighbor information is propagated during iteration such that a global image restoration is effected when the system reaches a steady state. The restored image can be found by solving an optimization problem that depends on the pixel interaction potentials. The MRF property that allows each pixel update to depend only on a local neighborhood of pixels eases the computational burden. For the case of a Gaussian pixel interaction, the potential function is quadratic, leading to a simple optimization problem that involves the solution of a large set of linear equations. For sigmoid interaction potentials (the present work), a difficult high-dimensional nonlinear optimization problem results. Stochastic methods are commonly used to solve such problems, but such methods are often very slow and sensitive to the choice of annealing schedule. We propose the novel deterministic Gradient Descent Gain Annealing (GDGA) method for solving high-dimensional nonlinear optimization problems. This method is fast and *automatically* chooses an annealing schedule. GDGA is used to solve optimization problems resulting from the neural-network-based MRF image-restoration model. Previous deterministic annealing work, such as mean field annealing [1, 2], does not incorporate an automatic annealing schedule.

The utility of the MRF model in restoring images corrupted with varying levels of non-Gaussian measurement noise has been investigated. Model performance has been evaluated quantitatively in terms of target detection and recognition, and the performance has been compared to that of the commonly used median-filtering technique. The quantitative results reported in this article are among the first such results in the literature. Because of prohibitive computational requirements, few quantitative characterizations of image restoration algorithms have been performed. Most work in the literature has compared the restored imagery qualitatively, rather than determining the effect of the restoration stage on the *overall* system performance.

The same model can be applied to a large number of sensor measurements (Doppler, intensity, passive infrared, range, and video) by the adjustment of a single parameter. This feature is especially relevant for hardware implementation because it allows a single chip to be used for processing a wide variety of imagery. The model has a massively parallel architecture with local neighbor pixel interactions (four nearest neighbors) and can be implemented on a parallel-processing computer or a custom analog VLSI chip. Implementation of the model in analog VLSI would allow video-rate restoration of 512 × 512 pixel images.

Background

In this section the Bayesian formulation of image restoration is reviewed to show the formal connections to the restoration method that is the subject of this article. The Bayesian formulation relates the posterior probability that an estimate of the true image x^r is obtained given a measured image x^m and the prior probabilities:

$$P(x^r|x^m) = \frac{P(x^m|x^r)P(x^r)}{P(x^m)}. \qquad (1)$$

The term $P(x^m|x^r)$ incorporates prior knowledge of the measurement process, and $P(x^r)$ incorporates prior knowledge of the scene. The present model finds an estimate $\hat{x}^r$ that approximately maximizes $P(x^r|x^m)$ given the measurement x^m and prior knowledge of the scene in the form of $P(x^r)$. In the present model each pixel depends only on its four surrounding neighbors and the measured pixel as shown in Figure 1.

The probabilistic (Bayesian) formulation is equivalent to a physical system description in terms of an energy [3]:

$$E \propto -\log P(x). \qquad (2)$$

MARKOV RANDOM FIELDS

A SERIES OF EVENTS in time form a *Markov Chain* if the probability of the outcome of an event at time $t+1$ depends only on the outcome of the event at time t. This concept can also be applied to processes on a lattice. A *Markov Random Field (MRF)* defined on a lattice implies that the update of a pixel at site ij depends only on the values of pixels in a local neighborhood of sites N_{ij} (Figure A). In terms of conditional probabilities,

$$\begin{aligned} &P(X_{ij} = x_{ij} | X_{lk} = x_{lk}, \\ &\qquad lk \in \text{lattice}, lk \neq ij) \\ &= P(X_{ij} = x_{ij} | X_{lk} = x_{lk}, \\ &\qquad lk \in N_{ij}), \end{aligned}$$

where X_{ij} is the real-value distribution of a random variable associated with lattice site ij and x_{ij} is the specific value of the variable at that site. Thus the definition of an MRF on a lattice transforms a global problem into a more computationally tractable local problem.

It is also true that an MRF on a lattice has the following energy-based formulation:

$$P(x) = \frac{e^{-U(x)}}{Z},$$

where U is the global potential function for the entire lattice and Z is the partition function, which normalizes the probability $P(x)$ to a range from 0 to 1.

In the present work the energy $E(x)$ is defined over all independent pairs of sites p on the lattice:

$$U(x) = \sum_p E_p(x).$$

The specific energy of interaction for a pair of sites is given by a *sigmoid function:*

$$E_p(x) = \frac{1}{1 + e^{\beta x^2}},$$

where the gain ($\beta < 0$) defines the scale of the sigmoid, as shown in Figure B. A small magnitude of the gain produces a large-scale (broad) sigmoid, while a large magnitude of the gain produces a small-scale (narrow) sigmoid. For either case, the sigmoid function has the property that the response saturates after the input exceeds a certain level. For a high magnitude of the gain, note that the sigmoid saturates very quickly, even for small inputs. The gain in the sigmoid function is inversely proportional to the temperature of an energy-based formulation.

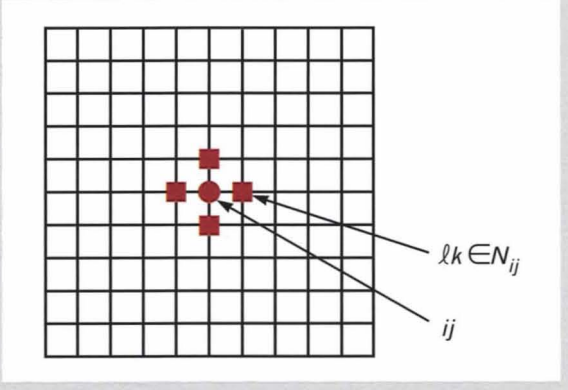

FIGURE A. Markov Random Field (MRF) defined on a lattice. In the figure, the update of a pixel at site ij depends only on the values of pixels at sites lk in a local neighborhood of sites N_{ij}.

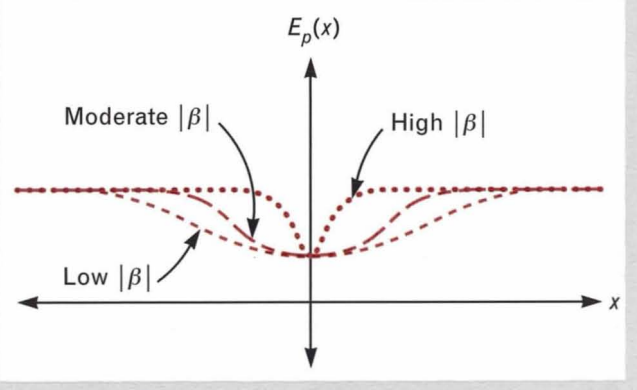

FIGURE B. Sigmoid function $E_p(x)$ for different magnitudes of the gain β. A small magnitude of the gain produces a large-scale (broad) sigmoid, while a large magnitude of the gain produces a small-scale (narrow) sigmoid.

Thus an MRF image processor may be specified by defining the energy function rather than the probabilities. This empirical approach is used in the present work.

Model Description

Equation 2 indicates that a minimization of the energy will result in a maximization of the probability $P(x^r | x^m)$. The total system energy can be expressed as the sum of a field term (which is due to the measured image) and a surround term (which is due to the neighbor interactions):

$$E = \lambda E^F + E^S. \tag{3}$$

The field coupling λ in Equation 3 is an adjustable parameter that determines the importance of the measurement term relative to the surround term: a small value of λ produces a highly smoothed image with little contribution from the measured image, whereas a large value essentially reproduces the measured image. The sigmoid function is used in both terms. The field term is given by

$$E^F = \Sigma_{ij} \frac{1}{1 + e^{\beta^F (\Delta_{ij}^m)^2}}, \tag{4}$$

where Δ_{ij}^m is the difference between the restored and measured pixels (i.e., $\Delta_{ij}^m = x_{ij}^r - x_{ij}^m$), and β^F is the saturation gain term. The sigmoid function is also used for the surround term:

$$E^S = \Sigma_p \frac{1}{1 + e^{\beta^S (\Delta_p^s)^2}}, \tag{5}$$

where Δ_p^s is the surround pair difference, i.e., $\Delta_p^s = x_{p1}^r - x_{p2}^r$, where p refers to all independent nearest neighbor pixel pairs in the image, and $p1$ and $p2$ refer to the members of a pair). Note that for an $M \times N$ lattice there are $(N - 1)M$ horizontal pairs and $(M - 1)N$ vertical pairs for a total of $2MN - M - N$ independent pairs.

The estimate of the original image $\hat{x}^r$ that minimizes the system energy is obtained with a deterministic search procedure. The present work uses the GDGA deterministic search (described in the subsection "Deterministic Solution") to decrease the saturation gain term β in Equations 4 and 5 from –0.001 to –10.0. (Note that the gain is negative in the present formulation.)

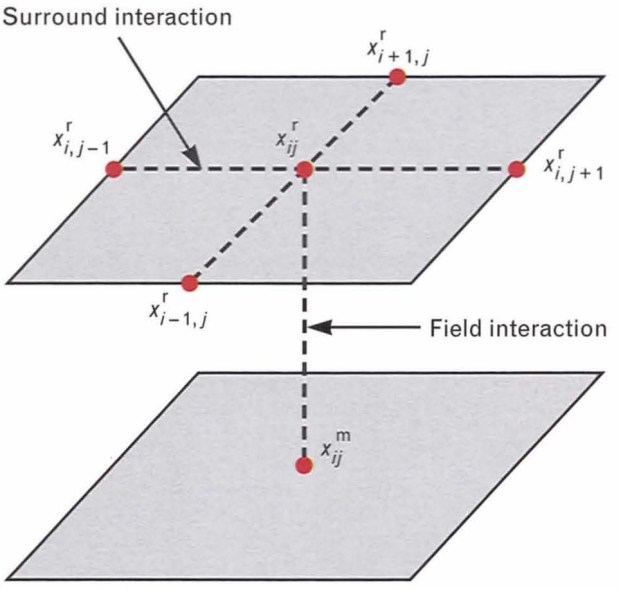

FIGURE 1. Nearest-neighbor architecture used in the Markov Random Field (MRF) image restoration.

The saturating aspect of the sigmoid function (from the neural network literature [4]) in Equation 5 allows the formation of sharp boundaries between dissimilar regions. The main advantage of using a sigmoid surround term is that sharp segmentations can be obtained without a separate "line process" [3], which would require solving $2MN - M - N$ extra equations. Hence the sigmoid term clearly reduces the computational load. For the same reason, a sigmoid function is used for the field, or measurement, term. The sigmoid function solves the problem of providing smoothing (noise reduction) while preserving naturally occurring boundary information in the scene.

Optimization Methods

Stochastic Solution

To solve the nonlinear optimization problem suggested by Equation 3, researchers have often attempted stochastic methods, which do not require the derivative of the energy with respect to the restored

state and, as a consequence, can be used for a wide range of optimization problems. Stochastic methods are also well suited for high-dimensional problems that are characterized by many acceptable solutions (restored states) all having approximately the same energy. Image restoration requires a solution with low energy, but does not need *the* global minimum.

The present work derives a stochastic solution by relating the statistical description of the problem to an energy-based representation. By making a correspondence to a physical system at thermal equilibrium, we can express the formulation in Equation 1 in terms of minimizing the energy of a system. The probability that a physical system in equilibrium with a heat bath at temperature T is in state i with energy E_i is given by the Boltzmann distribution:

$$P_T(x = i) = \frac{e^{\frac{-E_i}{k_B T}}}{Z(T)},$$

where k_B is the Boltzmann constant and $Z(T)$ is the partition function, which is simply the sum of the exponential term $E_i/(k_B T)$ over all possible states i.

It is assumed that the solutions to the optimization problem are equivalent to the states of a physical system and the cost of a solution corresponds to the energy of a state. Asymptotic convergence to a set of globally optimal solutions can be obtained provided that the different states are generated properly and the appropriate conditions are used to decide whether a given state should be accepted [5]. Stochastic methods for solving the optimization problem involve starting at a high temperature and annealing (i.e., reducing) the temperature until the system "freezes" to the minimum energy state. Ideally, the procedure would be implemented reversibly such that the system is always at thermal equilibrium and a true global minimum is reached rather than a metastable state.

Stochastic methods for solving nonlinear optimization problems typically use a simulated annealing method [6] combined with a Monte Carlo technique such as the Metropolis algorithm [7] or the Gibbs sampler [8]. For a comprehensive study that investigates the application of simulated annealing to image reconstruction, see Reference 3 by S. Geman and D. Geman. The problem with such stochastic solution techniques is that a good annealing schedule is difficult to determine, and the solution time can be prohibitive in terms of the number of iterations required because an equilibrium must be reached at each stage of annealing. At high temperatures a large temperature step is possible because the search covers a wide range of the state space. As the temperature is lowered, however, the system often reaches a critical point below which the state is "frozen," analogous to the phase diagram of real physical systems. If the critical point on the energy-versus-temperature curve were known, then large steps could be taken before the critical point were reached and small steps afterwards. Unfortunately, the "phase diagram" depends on the initial measurement, or field, term.

In practice a conservative annealing schedule is often used:

$$T \propto \frac{1}{\log k},$$

where k is the iteration number and T is the temperature. Such a schedule can require hundreds of thousands of iterations or more to produce an acceptable restoration. Automated Local Annealing (ALA) has been suggested to provide an automatic annealing schedule for neural networks [9], but the procedure is not directly applicable to an image-restoration formulation.

Another problem is that the computational expense of the stochastic method also depends on the number of allowable states per image pixel. An image with 8-bit pixels requires many more iterations for the full exploration of the state space as compared to, for example, a 4-bit image. Indeed, the solution of such nonlinear optimization problems remains a challenging research area.

Deterministic Solution

The large number of iterations that the stochastic approach requires in practice has motivated the use of a deterministic solution technique to solve the nonlinear image-restoration problem. The deterministic

approach attempts to minimize the system energy by iteratively updating pixel values across the lattice until a steady state is reached. In the approach, the use of high gain values for the iterative solution to Equation 3 produces a restored image with sharp boundaries. (Note: In the analogy of the physical system discussed earlier, a high gain value corresponds to a low temperature, or a small scale in that a small change in the input to the sigmoid function will produce a large change in the output.) The use of high gain values, however, will most likely lead to the procedure's being trapped in a local minimum. To remedy this problem we have developed the GDGA technique, which starts the solution procedure at a low gain (i.e., a high temperature, or large scale). The intermediate solution at low gain is then used as an initial condition to the problem at a higher gain, and the procedure is repeated until the final desired gain values are achieved. Solving a series of problems each at higher gain values is equivalent to temperature annealing in the stochastic approach. In addition, we have developed an *automatic* annealing (gain increase, or scale decrease) schedule that is described below.

An equation of motion based on the total energy from Equation 3 is defined by

$$\frac{\partial x^{\mathrm{r}}}{\partial t} = -\nabla_{x^{\mathrm{r}}} E\,, \tag{6}$$

where t represents a pseudo-time quantity. If Equations 4 and 5 are substituted for the total energy term in Equation 6, then the equation of motion for a single pixel at a lattice site ij is

$$\frac{\partial x^{\mathrm{r}}_{ij}}{\partial t} = \frac{-\partial}{\partial x^{\mathrm{r}}_{ij}} \sum_{lk} \left(\frac{\lambda}{1 + e^{\beta^{\mathrm{F}} (\Delta^{\mathrm{m}}_{lk})^2}} \right) - \frac{\partial}{\partial x^{\mathrm{r}}_{ij}} \sum_{p} \left(\frac{1}{1 + e^{\beta^{\mathrm{S}} (\Delta^{\mathrm{s}}_{p})^2}} \right),$$

where Σ_{lk} refers to all of the lattice sites in the image and Σ_p refers to all of the independent pixel pairs in the lattice. In the present work the MRF is given by the two horizontal and two vertical pairs associated with a given lattice site ij, resulting in a neighborhood N_{ij} given by

$$N_{ij} = \left\{ x_{i+1,j},\ x_{i-1,j},\ x_{i,j+1},\ x_{i,j-1} \right\}.$$

Hence, with this local neighborhood the update of a pixel at lattice site ij depends only on the pixel's four nearest neighbors.

The objective is to find the steady-state solution to Equation 6 that results in a state x^{r} that minimizes the system energy. The particular form of Equation 6, the equation of motion, guarantees that the steady-state solution minimizes the energy. This relationship can be shown by using the identity

$$\frac{\partial E}{\partial t} = \frac{\partial E}{\partial x^{\mathrm{r}}} \frac{\partial x^{\mathrm{r}}}{\partial t} \tag{7}$$

and substituting for $\partial x^{\mathrm{r}}/\partial t$ from Equation 6 into Equation 7, resulting in

$$\frac{\partial E}{\partial t} = -\left(\frac{\partial E}{\partial x^{\mathrm{r}}} \right)^2.$$

In the present work the GDGA deterministic technique is used to minimize the energy. This formulation is similar to the Graduated Non Convexity (GNC) approach of A. Blake and A. Zisserman [10] and the technique used by Y.G. LeClerc [11], and has some similarity to mean field annealing [1]. We have found GDGA to be substantially faster than the stochastic techniques described in the literature. The GDGA technique iteratively solves Equation 6 by calculating the gradient of the energy and updating the state (similar to an Euler solution of a system of coupled differential equations). In this approach the magnitude of the gain terms β^{F} and β^{S} in Equations 4 and 5 are increased from a value starting at 0.001. At small gain magnitudes the restoration acts to smooth the image because the energy terms are approximately locally quadratic with the pixel difference. (An energy term that is quadratic generates a larger penalty for larger pixel differences. Hence a smooth image, i.e., an image with equal pixel values, minimizes this energy.) Also, at small gain magnitudes all edges in the image are smoothed, resulting in a blurred image. As the magnitude of the gain is increased the natural-

ly occurring boundaries in the measured image start to appear, and eventually a sharp segmentation results.

The steady-state solution of Equation 6 at a given gain value is found by setting the gradient of the energy to zero and iteratively solving for the new pixel value. The deterministic technique is implemented with a fixed-point iteration around each pixel, in which the pixels are updated with a Jacobi (fully parallel) scheme [12]. In the technique, the gradient of the energy is set to zero, and the term x_{ij} is updated based on the old values of its neighbors:

$$x_{ij}^{r(new)} = \frac{\beta^S(A) + \lambda\beta^F x_{ij}^m g_{ij}^F}{\lambda\beta^F g_{ij}^F + \beta^S(B)},$$

where

$$A = g_{i,j-1}^S x_{i,j-1}^{r(old)} + g_{i,j+1}^S x_{i,j+1}^{r(old)} + g_{i-1,j}^S x_{i-1,j}^{r(old)} + g_{i+1,j}^S x_{i+1,j}^{r(old)}, \text{ and}$$

$$B = g_{i,j-1}^S + g_{i,j+1}^S + g_{i-1,j}^S + g_{i+1,j}^S . \tag{8}$$

In Equation 8 the nonlinear term g_{ij} is given by

$$g_{ij} = \frac{e^{\beta(\Delta_{ij})^2}}{\left(1 + e^{\beta(\Delta_{ij})^2}\right)^2}, \tag{9}$$

where Δ_{ij} is calculated based on the old pixel values. At this point the lattice could be updated, but a gain annealing schedule has not yet been specified. The GDGA technique *automatically* selects an annealing schedule by using feedback from the total system energy to select the gain step size. The strategy involves varying the step in β in order to maintain constant steps in energy. The β step is given by

$$\Delta\beta = \frac{\Delta E}{\frac{\partial E}{\partial \beta}}. \tag{10}$$

In Equation 10 a constant energy step ΔE is used, and the derivative is given by

$$\frac{\partial E}{\partial \beta} = \Sigma_{ij}(\Delta_{ij})^2 \frac{e^{\beta(\Delta_{ij})^2}}{\left(1 + e^{\beta(\Delta_{ij})^2}\right)}. \tag{11}$$

Note that in Equations 10 and 11 the gain term β refers to both the field and surround terms.

The GDGA technique starts at a small magnitude of gain and repeatedly applies Equation 8 until convergence, which typically requires fewer than 10 updates. Then Equation 10 is used to update the gain terms β^F and β^S and, with the new values, the pixels are again updated. The procedure is terminated when the magnitude of the gain becomes large—typically, a value of 10. For both simulated and real images, the GDGA algorithm is able to complete the restoration process (i.e., achieve sharp segmentation with noise removal) by using a total of 100 to 200 applications of the update equation (each application, or iteration, of Equation 8 updates all the pixels in the lattice). The restoration of a 128 × 128 pixel, 8-bit image requires less than 5 min on a SUN-4 workstation.

Further experiments have revealed that the auto-

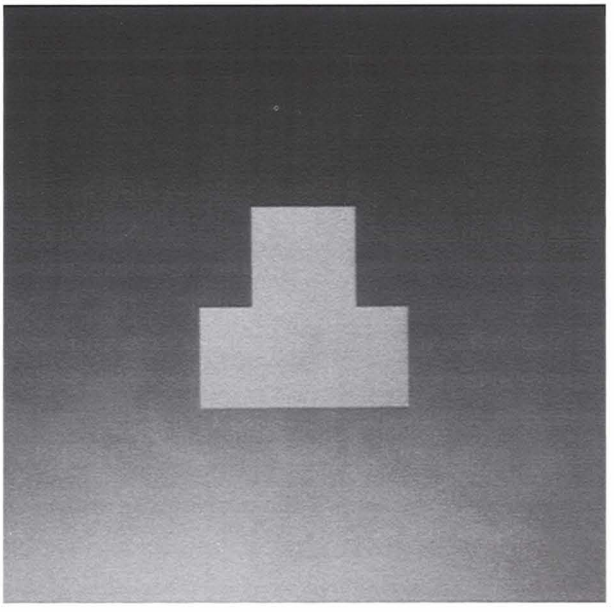

FIGURE 2. Original (noise free) image containing a linearly sloping background with a target at constant range.

FIGURE 3. The image of Figure 2 has been synthesized for a carrier-to-noise ratio (CNR) of 10 dB, corresponding to an image in which 20% of the pixel values are anomalous. (An anomaly is defined as a pixel value in the corrupted image that differs from the value in the original image by more than two range counts.) Each pixel in the image has 8 bits (256 gray levels) of resolution.

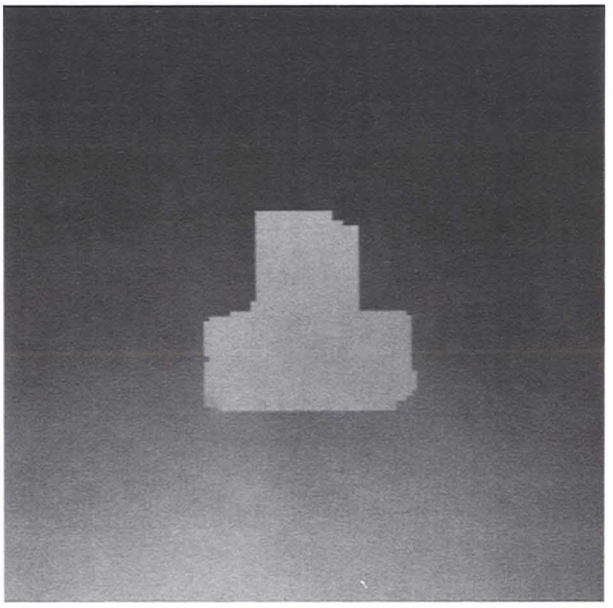

FIGURE 4. Results of MRF restoration on the synthetic range image of Figure 3. A four-nearest-neighbor MRF processor was used to restore the image. Note that the edges have been nearly perfectly preserved.

matic gain annealing process takes small steps in β at small magnitudes of gain and toward the end of the annealing takes large steps as the magnitude of β becomes larger. Thus the GDGA method adaptively adjusts the step size to make efficient use of each iteration. The adaptive nature of the algorithm is especially evident when comparing restorations of low- and high-noise imagery. Images with about 10% of the pixels corrupted with noise require fewer than 20 iterations for the entire restoration, while images with 70% noise require about 200 to 250 iterations.

Simulation Results

This section presents the qualitative and quantitative results obtained from applying the GDGA algorithm to the MRF image-restoration model described earlier. The algorithm was tested on a synthetic range image that had been corrupted with noise by a range-sensor measurement model described in the literature [13, 14]. The measurement model, which simulates a peak-detecting laser radar sensor that introduces anomalies into the measurement, was used to corrupt an image by relating the carrier-to-noise ratio (CNR) of the range sensor to the expected percent anomalies in the measurement. (An anomaly is defined as a pixel value in the corrupted image that differs from the value in the original image by more than two range counts.) In this work CNR values of 10 dB and 6 dB were used, corresponding to 20% and 71% anomalies, respectively, in the corrupted image. The measurement model does not assume a Gaussian distribution and is based on realistic sensor measurements.

The original (noise free) synthetic image shown in Figure 2 contains a simple shape at a constant pixel value against a background whose pixel values linearly increase from the top to the bottom of the image. In this work all of the input and restored images have 8 bits (256 gray levels) of resolution. Figure 3 shows the range image of Figure 2 after the image has been corrupted with 20% anomalies, and Figure 4 shows the result of the MRF restoration. Except for a few discrepancies at the boundary, the restoration is nearly perfect, especially in recovering the sloping background. Figure 5 shows the range image of Figure 2

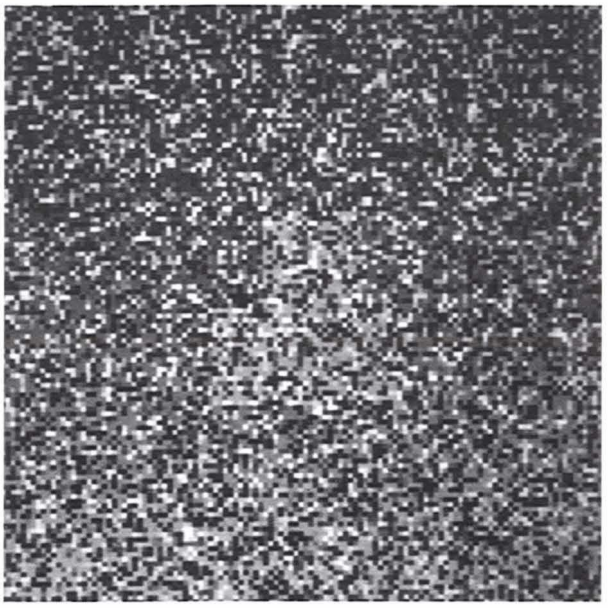

FIGURE 5. The image of Figure 2 has been synthesized for a CNR of 6 dB, corresponding to an image that has 71% anomalies. Each pixel in the image has 8 bits (256 gray levels) of resolution.

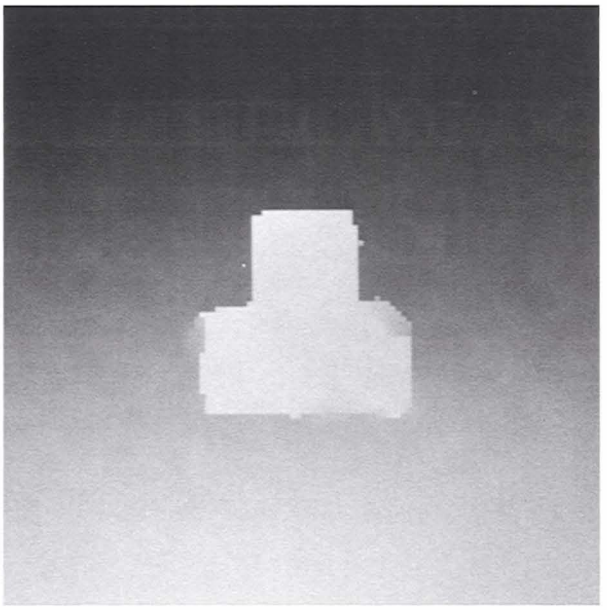

FIGURE 6. Results of MRF restoration on the synthetic range image of Figure 5. A four-nearest-neighbor MRF processor was used to restore the image. Note that the edges have been preserved.

with 71% anomalies, and Figure 6 shows the result of MRF restoration. Although the human visual system can barely recognize the original shape at this noise level, the MRF restoration is able to recover the underlying edges that define the target's shape, as shown in Figure 6. In both cases the MRF model produces a piecewise smooth restoration of the input image. The use of the sigmoid function facilitates noise reduction (smoothing) while preserving sharp discontinuities (edges).

Next we describe a quantitative analysis of the model's performance based on the average percent of anomalies in the restored image. This measure is relevant to target detection in a noisy environment. Figure 7 shows a statistical comparison between corrupted images with and without restoration. In the figure, each point represents an average over 10 runs at a fixed noise level. The average percent anomalies in the image gives an indication of the difficulty of target detection; i.e., the probability of detection is lower at higher anomaly percentages. For effective detection, the percent of anomalies must be less than about 10%. Thus the results show that MRF restoration is effective up to a CNR of about 6 dB. At this CNR the average percent of anomalies is 71% for the input image, 55% for the median-filtered image, and only 4.5% for the MRF-restored image. Clearly the MRF restoration provides superior detection performance in a high-noise environment.

We have also evaluated the utility of the MRF image-restoration model as a preprocessor for target recognition in a noisy environment. Eight different binary silhouettes—broadside views of different vehicles (Figure 8)—were used in this experiment. Each of the silhouettes was centered in a 128 × 128 pixel image to simulate range images, which were then corrupted with the range measurement model to produce images corresponding to realistic range measurements. Twelve such images were produced for each of the 8 silhouettes for a total of 96 images per each CNR value from 6 dB through 20 dB in 1-dB steps, and at 80 dB. Next, detection and segmentation were performed on these simulated range measurements to obtain noisy binary range slices in which only those pixels within a certain range are shown. The range slices were then used to train a Nearest Neighbor Classifier (NNC) [15] to separate the 8 different vehicles. We found that at very high CNR values (40 dB) the classifier

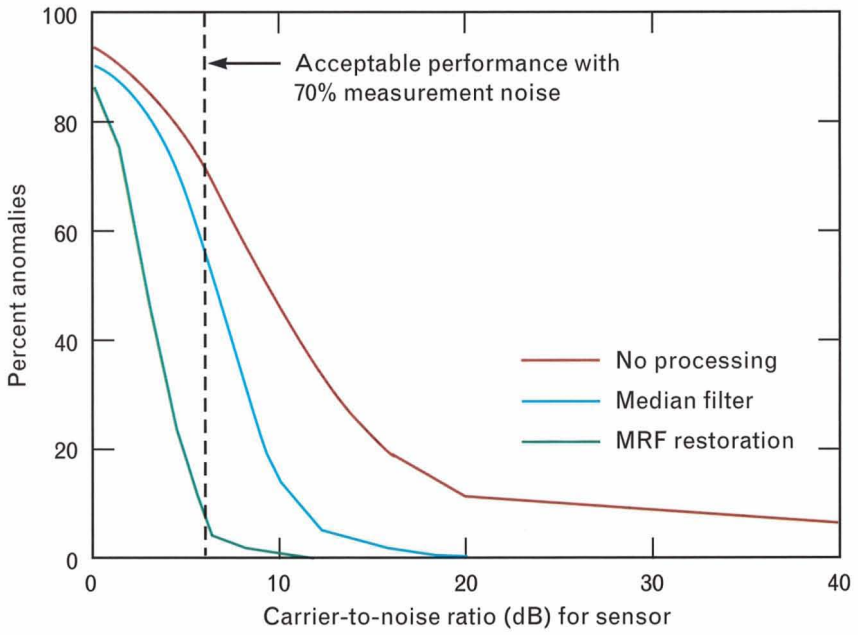

FIGURE 7. Average percent anomalies for range imagery as a function of sensor CNR. The results from MRF restoration are compared with median filtering and the case in which no processing has been performed to restore the image. Each data point represents an average over 10 runs at a given noise level. For effective target detection, the percent of anomalies must be less than about 10%. Thus the results indicate that MRF restoration is effective up to a CNR of about 6 dB. At this CNR value the average percent of anomalies is 71% for the input image, 55% for the median-filtered image, and only 4.5% for the MRF-restored image. Clearly, MRF restoration provides superior detection performance in a high-noise environment.

created a unique category for each silhouette. At lower CNR values (higher noise), however, the classifier formed extra categories because exemplars of the same target were sometimes classified into different categories. We repeated the above procedure twice: once using the median-filter technique on the corrupted images before the detection and segmentation steps, and once using MRF restoration. Figure 9 compares MRF restoration with iterated median filtering and with the case in which no processing had been performed to restore the image. The performance at each CNR value is defined as the fraction of the 96 examples that the NNC has classified correctly. Note that the MRF restoration is able to maintain an acceptable level of performance at a CNR that is 5 and 10 dB lower than with the median-filter and no-preprocessing case, respectively. Thus, with MRF restoration, a sensor can be operated at roughly 25% the power level required by the use of a median filter.

Hardware Implementation

For real-time image restoration, the MRF model can be implemented either digitally—on custom digital signal processing (DSP) chips or on a single-instruction multiple-data (SIMD) computer such as the Connection machine manufactured by Thinking Machines Corp.—or in an analog manner on a custom VLSI chip.

For digital implementation, 72 floating-point operations are required per pixel update. Typically, a pixel must be updated about 100 times over the course of the restoration. Thus a 256 × 256 image restoration would require 472 million floating-point operations, and a frame-rate restoration of the same image would require digital hardware that delivers 14 GFLOPS of performance. This performance level is at the leading edge of current digital processing tech-

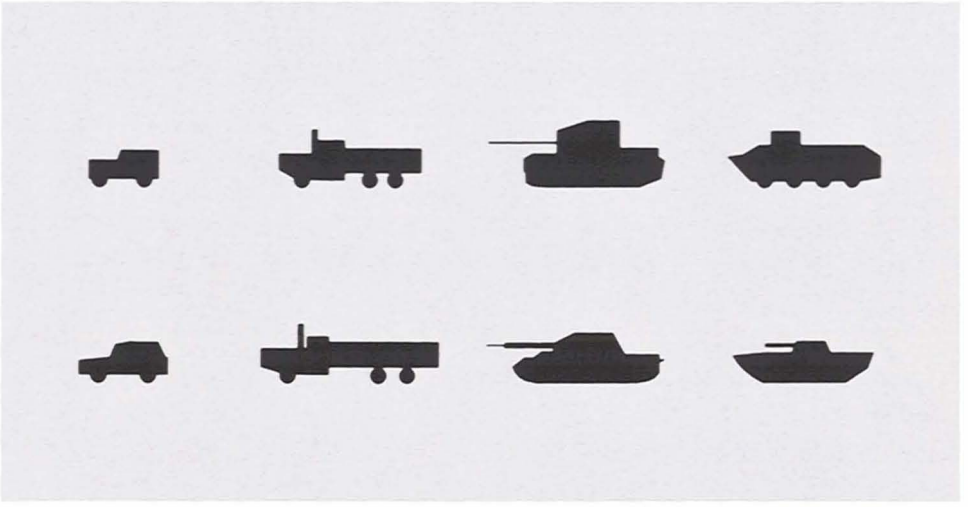

FIGURE 8. Binary silhouettes of 8 different vehicles that were used to evaluate the effect of MRF restoration on target recognition. (For the test results, see Figure 9.)

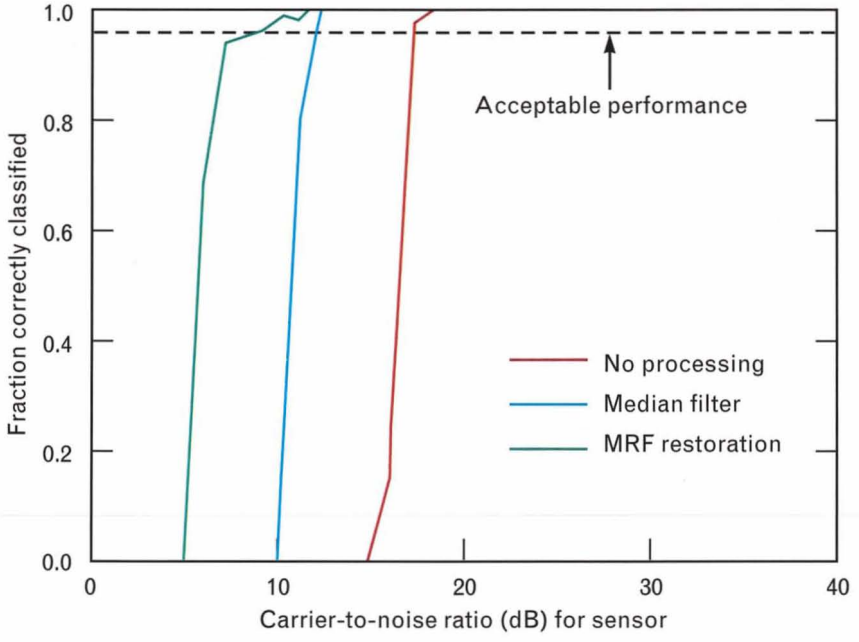

FIGURE 9. Fraction of correctly classified range slices as a function of sensor CNR. In the experiment, the binary silhouettes of 8 different vehicles (Figure 8) were used to create realistic noise-corrupted range measurements. Detection and segmentation were performed on these simulated range measurements to obtain noisy binary range slices in which only those pixels with a certain range are shown. A total of 12 such slices was produced for each of the 8 silhouettes at each CNR value. The 96 range slices at each CNR value were then used for training a Nearest Neighbor Classifier (NNC) [15] to separate the 8 different vehicles. Before being presented to the NNC, some of the range slices were restored by the MRF model and others by an iterated median filter. The results compare MRF restoration with the median-filter technique and with the case in which no processing has been performed.

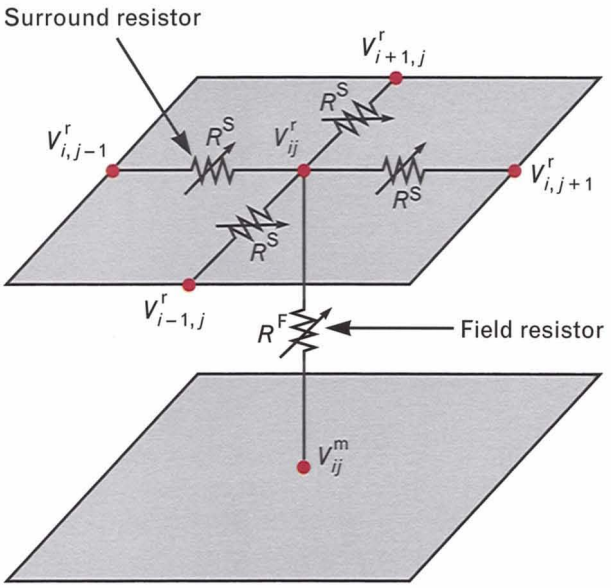

FIGURE 10. System architecture implemented as a resistive grid (compare with Figure 1). The voltages V^m and V^r represent the measured and restored images, respectively. Note the presence of both field and surround resistors.

nology. The advantage of an all-digital implementation is that it does not require the hardwiring of any of the system parameters.

An all-analog implementation requires the implementation of the energy function as an analog circuit. In the current example the system architecture can be represented as a resistive grid (Figure 10). The resistors are nonlinear in that their resistances are voltage dependent:

$$R(\Delta_p) \propto \frac{\left(1 + e^{\beta(\Delta_p)^2}\right)^2}{e^{\beta(\Delta_p)^2}},$$

where Δ_p refers to the voltage difference between a pair of adjacent sites. The nonlinear resistor is essentially the inverse of Equation 9, and the circuit consists of separate field and surround resistors, as shown in Figure 10. At steady state there is a current balance at every site and the voltages V^r correspond to the intensities in the restored image. The advantage of this implementation is that there is essentially a "processor" at every site, and the processing speed is limited only by the settling time of the analog circuit. The throughput is limited by the input/output onto and off the analog chip, and not the circuit. With current technology, images can be restored at a rate in the thousands of frames per second.

Summary

An efficient Markov Random Field (MRF) based method for performing piecewise smooth image restorations has been demonstrated. The underlying model uses a neural network sigmoid potential between pixel pairs to allow the formation of sharp boundaries between dissimilar regions in the presence of noise. A novel deterministic method—called Gradient Descent Gain Annealing (GDGA)—for solving the nonlinear coupled set of differential equations that the MRF model introduces was presented. The GDGA algorithm typically requires fewer than 200 iterations to restore an image, where the number of iterations is roughly proportional to the level of noise in the image. Computer simulations on noisy images have shown that restorations can be performed for very high noise levels (i.e., images that have up to 71% of their pixels corrupted with non-Gaussian sensor noise). Simulation results indicate that MRF restoration provides a 5-dB advantage in the carrier-to-noise ratio (CNR) over conventional iterated median filtering. Although the same model is currently used to restore images from different sensors, arbitrary potentials can be incorporated for the pixel interactions so that the system can be tailored to specific natural scenes and sensors. The system uses a massively parallel set of local neighborhoods (four nearest neighboring pixels) for efficient implementation on a parallel-processing computer or a custom analog VLSI chip.

Acknowledgments

The author wishes to thank members of the Opto-Radar Systems Group at Lincoln Laboratory for very helpful technical discussions and computer support throughout the course of this research. In particular, the author is indebted to William M. Wells III for many helpful discussions and suggestions, and for help in reviewing the existing image-restoration literature.

This work was sponsored by the Defense Advanced

Research Projects Agency.

REFERENCES

1. G.L. Bilbro and W.E. Snyder, "Range Image Restoration Using Mean Field Annealing," in *Advances in Neural Information Processing Systems I,* D.S. Touretzky, ed. (Morgan Kaufmann, San Mateo, CA, 1989), pp. 594–601.
2. D. Geiger and F. Girosi, "Parallel and Deterministic Algorithms for MRF's: Surface Reconstruction and Integration," *IEEE Trans. Pattern Anal. Mach. Intell.* **13,** 401 (1991).
3. S. Geman and D. Geman, "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images," *IEEE Trans. Pattern Anal. Mach. Intell.* **6,** 721 (1984).
4. D.E. Rumelhart, G.E. Hinton, and R.J. Williams, "Learning Internal Representations by Error Propagation," in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 2,* D.E. Rumelhart and J.L. McClelland, eds. (MIT Press, Cambridge, MA, 1986), p. 422.
5. E. Aarts and J. Korst, *Simulated Annealing and Boltzmann Machines* (John Wiley, NY, 1989), p. 130.
6. S. Kirkpatrick, C.D. Gelatt, Jr., and M.P. Vecchi, "Optimization by Simulated Annealing," *Science* **220,** 671 (1983).
7. N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller, "Equation of State Calculations by Fast Computing Machines," *J. Chem. Phys.* **21,** 1087 (1953).
8. J.L. Marroquin, "Probabilistic Solution of Inverse Problems," Ph.D. Thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA, 1985.
9. J. Leinbach, "Automatic Local Annealing," in *Advances in Neural Information Processing Systems I,* D.S. Touretzky, ed. (Morgan Kaufmann, San Mateo, CA, 1989), pp. 602–609.
10. A. Blake and A. Zisserman, *Visual Reconstruction* (MIT Press, Cambridge, MA, 1987), p. 131.
11. Y.G. LeClerc, "Constructing Simple Stable Descriptions for Image Partitioning," *Intl. J. Comput. Vision* **3,** 73 (May 1990).
12. S.D. Conte and C. de Boor, *Elementary Numerical Analysis: An Algorithmic Approach,* 3rd ed. (McGraw-Hill, NY, 1980), p. 226.
13. M.B. Mark, "Multipixel, Multidimensional Laser Radar System Performance," Ph.D. Thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA, 1986.
14. A.B. Gschwendtner, R.C. Harvey, and R.J. Hull, "Coherent IR Laser Technology," *Optical and Laser Remote Sensing,* D.K. Killinger and A. Mooradian, eds. (Springer-Verlag, NY, 1983), p. 327.
15. R.O. Duda and P.E. Hart, *Pattern Classification and Scene Analysis* (John Wiley, NY, 1973).

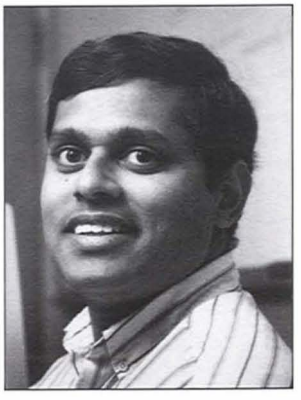

MURALI M. MENON
is currently a research staff member in the Opto-Radar Systems Group. He received a B.S., an M.S., and a Ph.D. degree in chemical engineering from Case Western Reserve University in Cleveland. His research interests include applied pattern recognition, signal processing, and image processing, with special interests in wavelets and artificial neural networks. He has spent the past six years at Lincoln Laboratory working on applications of neural networks for processing sensor data, including the design of automatic target recognition (ATR) systems.