This paper originally appeared in *Algorithms for Synthetic Aperture Radar Imagery VI*, E.G. Zelnio, editor, Proc. SPIE V 3721, SPIE, Orlando, (1999)

Markov random field segmentation methods for SAR target chips

Robert A. Weisenseel[a*], W. Clem Karl[a], David A. Castañon[a],
Gregory J. Power[b], Phil Douville[b]

[a]Boston University, Electrical and Computer Engineering Department
[b]Air Force Research Lab, Automatic Target Recognition Branch (AFRL/SNAT),
Building 620, 2241 Avionics Circle Wright-Patterson AFB, Ohio 45433-7321

ABSTRACT

DARPA's Moving and Stationary Target Acquisition and Recognition (MSTAR) program has shown that image segmentation of Synthetic Aperture Radar (SAR) imagery into target, shadow, and background clutter regions is a powerful tool in the process of recognizing targets in open terrain. Unfortunately, SAR imagery is extremely speckled. Impulsive noise can make traditional, purely intensity-based segmentation techniques fail. Introducing prior information about the segmentation image – its expected "smoothness" or anisotropy – in a statistically rational way can improve segmentations dramatically. Moreover, maintaining statistical rigor throughout the recognition process can suggest rational sensor fusion methods.

To this end, we introduce two Bayesian approaches to image segmentation of MSTAR target chips based on a statistical observation model and Markov Random Field (MRF) prior models. We compare the results of these segmentation methods to those from the MSTAR program. The technique we find by mapping the discrete Bayesian segmentation problem to a continuous optimization framework can compete easily with the MSTAR approach in speed, segmentation quality, and statistical optimality. We also find this approach provides more information than a simple discrete segmentation, supplying probability measures useful for error estimation.

Keywords: SAR, segmentation, MRF, statistical modeling, target recognition

1. INTRODUCTION

DARPA's Moving and Stationary Target Acquisition and Recognition (MSTAR) program is an effort to provide better automated tools for, as its name suggests, acquiring and recognizing targets, such as tanks or other military vehicles. One bonus of this program has been its freely distributed imagery and tools that have greatly advanced the image processing and recognition field, by providing a benchmark data set for comparing algorithms. Up until very recently, this data set has consisted entirely of Synthetic Aperture Radar (SAR) imagery.

One area of study within MSTAR is image segmentation as a feature extraction tool for target recognition. This work intersects with our own work on Reduced Signature Target Recognition. MSTAR has shown that image segmentation of SAR imagery into target, shadow, and background clutter regions can be a powerful tool in the process of recognizing targets, although most of its robustness is lost outside of open terrain. Unfortunately, SAR imagery is extremely speckled, as evidenced in the small set of examples of measured data in Figure 1. This impulsive noise can challenge traditional, purely intensity-based segmentation techniques. The resulting segmentations of such traditional approaches are often noisy. Introducing prior information about the segmentation image, such as its expected "smoothness" or anisotropy, in a statistically rational way can improve segmentations dramatically. Moreover, maintaining such statistical rigor throughout the recognition process can make fusing multiple sources of information straightforward and simplify the transition to more difficult reduced signatures.

In this paper, we introduce two Bayesian approaches to image segmentation of MSTAR target chips based on a statistical observation model and a Markov Random Field (MRF) prior model. We then compare the results of these segmentation approaches to those from the MSTAR program and from human segmentation. The continuous Bayesian MRF approach compares favorably to the MSTAR approach, opening a host of statistically based options for SAR image segmentation and information fusion.

*Correspondence: Email: lorax@bu.edu; WWW: http://mdsp.bu.edu/lorax_html

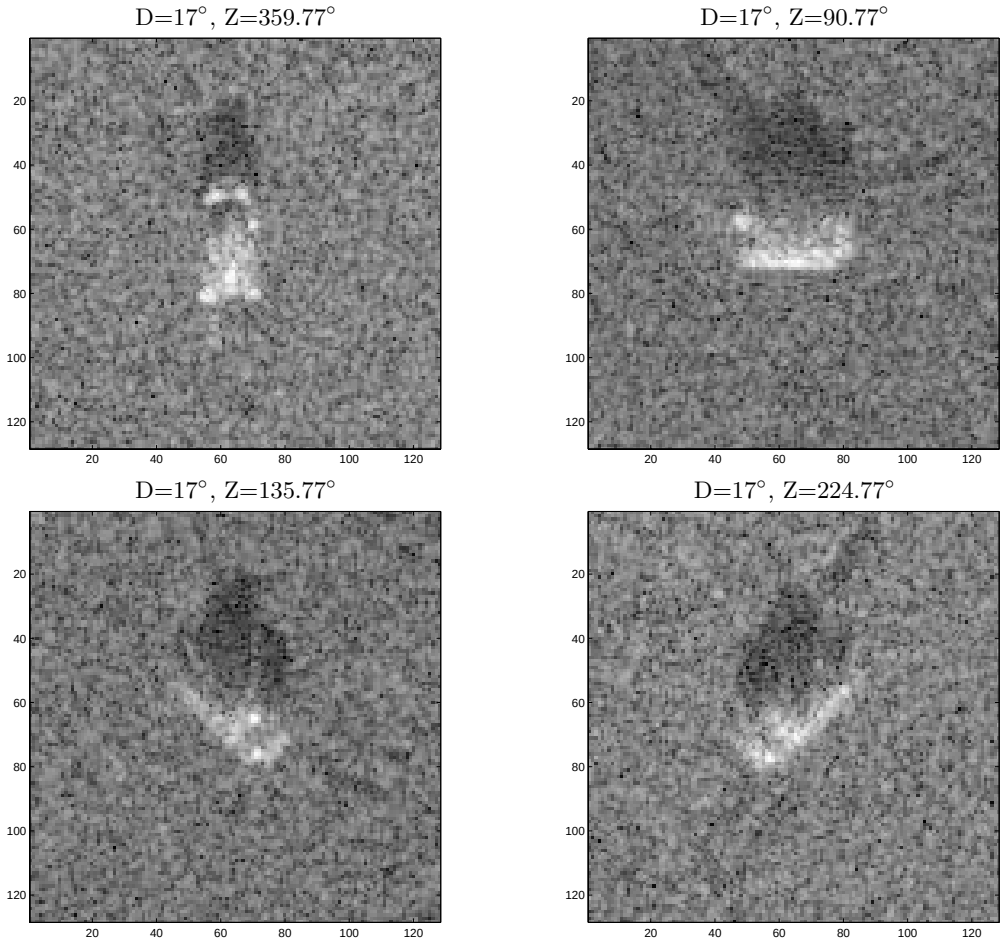

Figure 1. Example Logarithms of T72s from MSTAR's Public Data

2. BACKGROUND

There are a number of ways to segment SAR imagery. Possibly one of the simplest to implement is human segmentation, where humans sit in front of a computers and painstakingly draw the regions that they believe correspond to target, background, and shadow.[1] Examples of such segmentations appear in Figure 2. The point of the MSTAR program is to automate the target recognition process, since the sheer volume of data is too large compared to the number of available human analysts.

MSTAR has its own segmentation tools. These tools are well-developed and functional, but they are not based on a sound, robust statistical optimality criterion and may not be generalizable to a wider variety of signatures. MSTAR segmentations of measured imagery appear in Figure 3.

Another approach to segmentation is the one that we have chosen, based on Bayesian techniques. This method incorporates information from sources other than data. For example, we might want to incorporate some knowledge of a physical model for the system, beyond just observations.

We assume first that a given SAR image y consists of N pixels, each of which can take on positive real values. Also, we assume a segmentation label field x where each element, x_i, corresponds to a particular element, y_i, of y, and where each element can take on one of three values corresponding to target, background, and shadow labels. Mathematically,

$$y_i \in \Re^+ \text{ at pixel } i \qquad x_i = \begin{cases} 3 & \text{pixel } i = \text{target} \\ 2 & \text{pixel } i = \text{background clutter} \\ 1 & \text{pixel } i = \text{shadow} \end{cases} \tag{1}$$

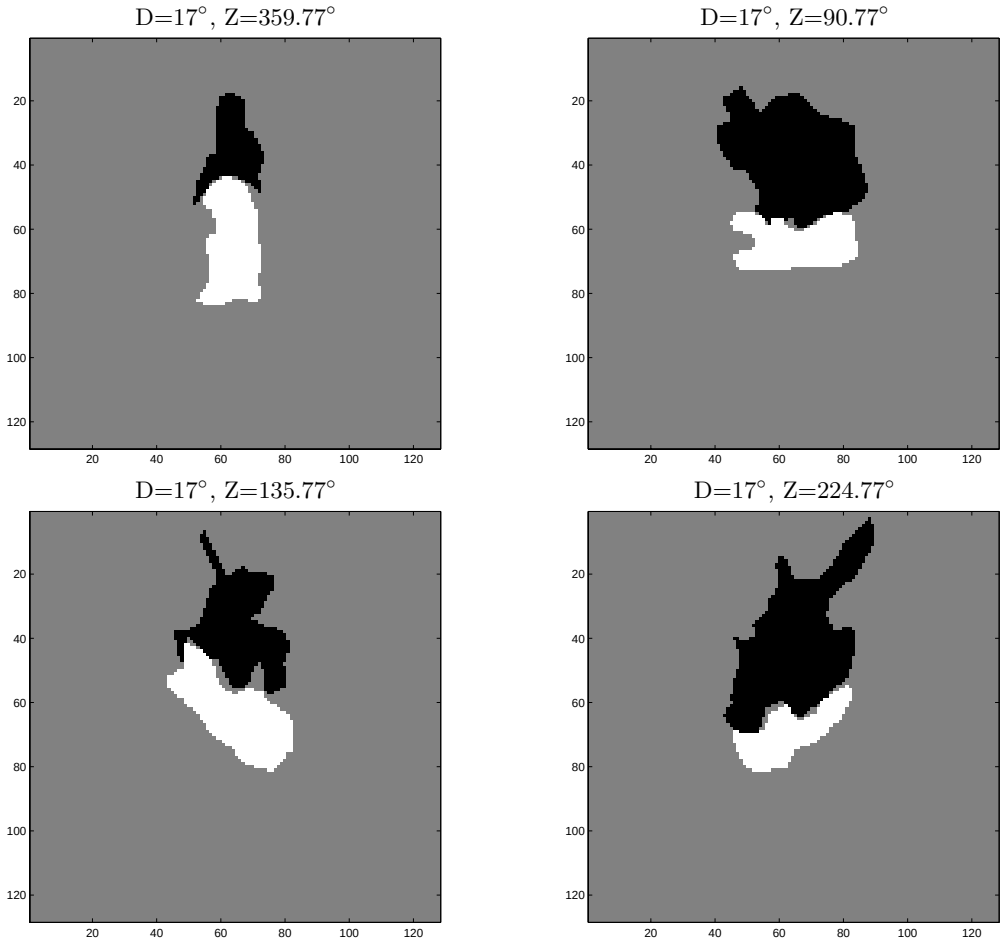

Figure 2. Example Human Segmentations of MSTAR Public Data

The Bayesian approach starts from the posterior probability density. One Bayesian approach that is appropriate to this problem is the Maximum A Posteriori (MAP) classifier. With x being the image where each pixel can take on three states, target, shadow, or background, and y being the SAR data image, the formula for determining the values of x is:

$$\hat{x} = \underset{x}{\text{argmax}} \; p(x|y) = \underset{x}{\text{argmax}} \; \frac{p(y|x)p(x)}{p(y)} = \underset{x}{\text{argmax}} \; p(y|x)p(x) \tag{2}$$

where $p(x|y)$ is the posterior probability distribution of the segmentation label image conditioned on the SAR data image, $p(y|x)$ is the likelihood of the data image conditioned on the segmentation label image, $p(x)$ is the prior probability distribution of the segmentation image, and $p(y)$ is the prior probability of the observation field. The first equality follows from Bayes rule, and the second from the fact that y is known, and therefore $p(y)$ is a constant. Thus, we just need three things: an observation probability model, a prior probability model, and an algorithm for maximizing over all possible segmentation images x.

3. METHOD

In our work, we use two different Markov Random Field (MRF) prior models to demonstrate our techniques.[2,3] In the first approach, we develop an anisotropic prior model that attempts to account for the tendency of shadow regions to fall directly downrange of target regions, as well as smoothing out small, isolated regions of shadow or target. This anisotropic model is somewhat complicated, and requires that we use simulated annealing to implement the

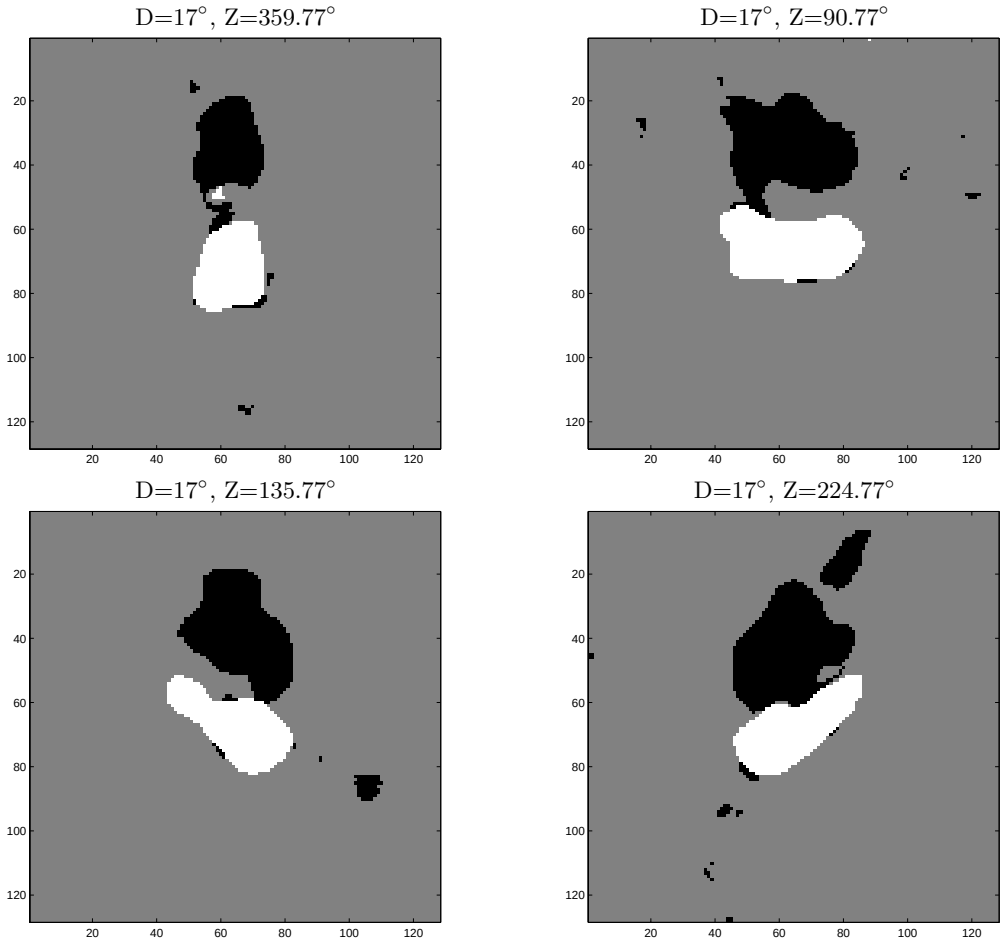

Figure 3. Example MSTAR Segmentations of Public Data

discrete optimization simply, at the expense of fast computation. The second approach uses an isotropic smoothness model that is readily transformable to a continuous, as opposed to discrete, space for x, allowing us to use fast gradient-based optimization methods. Both approaches use essentially the same observation model, which we will describe first.

3.1. Observation Model

For the observation model, we assume that each pixel of data, y_i, is conditionally independent of all other data pixels, conditioned on knowledge of the segmentation label at that pixel, x_i. Formally,

$$p(y|x) = p(y_1, \ldots, y_N|x_1, \ldots, x_N) = \prod_{i=1}^{N} p(y_i|x_i) \tag{3}$$

where N is the number of pixels in the image. We also assume that each data pixel is a conditionally identically distributed Weibull random variable,

$$p(y_i|x_i) = \frac{\beta(x_i)}{\alpha(x_i)} \left(\frac{y_i}{\alpha(x_i)}\right)^{\beta(x_i)-1} \exp\left\{-\left(\frac{y_i}{\alpha(x_i)}\right)^{\beta(x_i)}\right\} \quad \alpha, \beta > 0 \tag{4}$$

where $\alpha(x_i)$ and $\beta(x_i)$ are the scale and shape parameters, respectively, of the Weibull distribution, given that the observation is of segmentation label x_i. This label can be 1, 2, or 3 as stated in Equation 1.

We estimate the scale and shape parameters for each of the three distributions for the imagery from data, by assuming the human segmentations described in Section 2 were correct, and then using a maximum likelihood estimator for the parameters. For example, using all of the pixels labeled by a human as target from every measured training image, we estimated the parameters of the probability distribution of a particular observation pixel, assuming it came from a target region within the measured data training imagery:

$$\left[\hat{\alpha}(3), \hat{\beta}(3)\right] = \underset{\alpha,\,\beta}{\operatorname{argmax}} \; p(z|\alpha, \beta) \tag{5}$$

where z is that subset of all observation pixels y corresponding to a target region from every training image. We demonstrate the fit of the Weibull model to measured target, background, and shadow histograms in Figures 4(a)-(c), respectively.

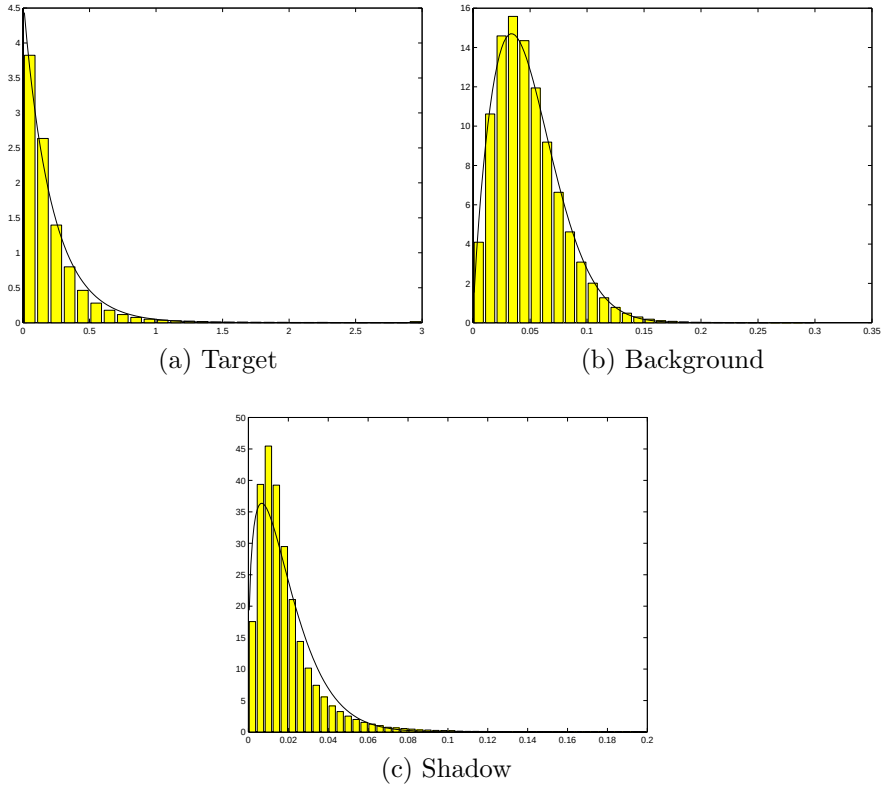

(a) Target (b) Background

(c) Shadow

Figure 4. Weibull Fit to Measured Data (Probability vs. Intensity)

3.2. Anisotropic, Discrete Model

3.2.1. Prior Model

In the first case we introduce spatial correlation of target, background, and shadow regions using an anisotropic MRF model. An MRF is a way of describing complex global structure through local pixel-to-pixel interactions in a statistically rigorous fashion and is related to Markov chains. We can extend the Markov chain concept to two dimensions with MRFs by saying a pixel depends only on its neighbors. The neighborhood we use is a four nearest neighborhood, shown visually in Figure 5.

By using the equivalence between MRFs and Gibbs distributions, we write the probability function for the MRF in Gibbs form:

$$p(x) = \frac{1}{\mathcal{Z}} \exp\left\{\alpha \frac{1}{2} \sum_{i=1}^{N} \sum_{j \in \eta_i} f_{i,j}(x_i, x_j)\right\} \tag{6}$$

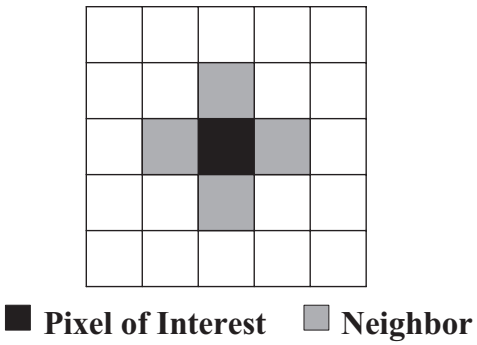

Figure 5. Neighborhood Structure of the Prior Model

where we use a single index for denoting each pixel of the image. The set of pixels η_i denotes the four nearest neighbors of pixel i. The variable α is just a weighting constant that can make the distribution "sharper" or "flatter," which correspondingly makes the segmentation field "smoother" or "noisier," respectively. $\mathcal{Z}$ is just a normalization constant.

The subscripts of the function $f_{i,j}(.,.)$, the potential function, make it possible for the MRF to be both inhomogeneous and anisotropic. In this work, we assume homogeneity, but not isotropy. For example, $f_{i,j}$ can be a different function depending on whether j is downrange of i or in one of the cross-range directions. By choosing a directional model, we incorporate prior knowledge about the structure of a SAR image segmentation field, i.e. that shadows usually fall downrange of targets and that there should be little or no background clutter between a target and its shadow. These functions appear in Table 1, with f_D denoting the potential function corresponding to pixel j being

$f_D(x_i, x_j)$		x_j = S	B	T
	S	0	-1	-2
x_i	B	-2	0	-1
	T	-1	-2	0

$f_U(x_i, x_j)$		x_j = S	B	T
	S	0	-2	-1
x_i	B	-1	0	-2
	T	-2	-1	0

$f_C(x_i, x_j)$		x_j = S	B	T
	S	0	-1	-1
x_i	B	-1	0	-1
	T	-1	-1	0

Table 1. Table of Potential Functions

downrange of pixel i, f_U denoting the potential function when j is up-range of i, and f_C denoting the potential for the case when j is in either of the cross-range directions from i. Note that f_U is just the transpose of f_D. The function f_C just implies that there is no penalty for two pixels of the same type neighboring in the cross-range direction, otherwise there is a penalty. The functions f_U and f_D are similar to f_C, but the cases of placing target downrange of shadow, background downrange of target, or background up-range of shadow are heavily penalized.

3.2.2. Simulated Annealing

To actually compute the MAP classifier for the anisotropic case, we combine the observation and prior models from Sections 3.1 and 3.2.1 into a function proportional to the posterior distribution with α as a weighting constant between the prior and observation distributions:

$$p(x|y) \propto e^{-H(x|y)} \tag{7}$$

$$H(x|y) = \sum_{i=1}^{N} \left[\ln \beta(x_i) - \beta(x_i) \ln \alpha(x_i) + (\beta(x_i) - 1) \ln y_i - \left(\frac{y_i}{\alpha(x_i)} \right)^{\beta(x_i)} - \alpha \sum_{j \in \eta_i} f_{i,j}(x_i, x_j) \right] \tag{8}$$

This equation shows the function we need to optimize, as described in Equation 2.The last thing we need is a means to maximize this function, and hence the posterior distribution, for all x.

We implement this maximization using a simulated annealing algorithm because the distribution of x is discrete.[4] Thus, we introduce a "temperature" parameter to the distribution:

$$p(x|y) \propto e^{-H(x|y)/T} \tag{9}$$

To obtain an optimal solution, T should start at a sufficiently high temperature and be reduced slowly according to a logarithmic cooling schedule that runs long enough to approach zero temperature,

$$T = \frac{c}{\ln(k+1)} \quad k = 1, 2, \ldots \tag{10}$$

where c is a specific temperature determined by the maximum possible change in $H(x|y)$ caused by changing a single pixel in the segmentation label image, x. However, to reduce the number of computations necessary, we initialize the system at a relatively low temperature with a pixel-by-pixel Maximum Likelihood (ML) solution using the likelihoods given in Equation 4 and then decrease the temperature with a faster logarithmic cooling schedule for 1000 iterations. We control the lower starting temperature and faster cooling schedule by setting c. Lastly, we reduce the temperature to zero and run the algorithm for 5 more iterations to force the system into the nearest local optimum.

While this approach is not guaranteed to achieve a globally optimal solution, the initialization makes this algorithm unlikely to do worse than the ML solution and the faster cooling schedule gives faster convergence. Because only local interactions are involved in this processing and the number of local operations required are small, this technique is well-suited to massive parallelization, either on multiple processors or on application specific ICs.

3.3. Isotropic, Continuous Model

3.3.1. Prior Model

We now present the case of a simpler isotropic discrete smoothness constraint that extends to a continuous modeling framework. The Gibbs distribution for this smoothness constraint is:

$$p(x) = \frac{1}{\mathcal{Z}} \exp\left\{ \alpha \frac{1}{2} \sum_{i=1}^{N} \sum_{j \in \eta_i} [\delta(x_i - x_j) - 1] \right\} \tag{11}$$

where $\mathcal{Z}$ is once again just a normalizing constant and $\delta(.)$ is just a discrete delta function. Note that this is the model obtained when the crossrange function from Table 1, $f_C(x_i, x_j)$, replaces of the downrange and uprange functions, f_D and f_U, in the anisotropic model of Section 3.2.1. A reader familiar with statistical mechanics might recognize this as a three-state Potts model. This model implies that a segmentation label for a particular pixel is more likely to be the same as its neighbors.

The simplest possible method for extending this model to a continuous segmentation field would be to approximate the delta function with a continuous function such as a narrow Gaussian. However, performing this approximation blindly could lead to unfortunate consequences. Particularly, such an approximation introduces a notion of distance that was not present in the discrete model. In the discrete model, target, background, and shadow labels were essentially equidistant, despite being numbered 3, 2, and 1 respectively. With the Gaussian approximation, background would be closer to target and shadow than shadow and target would be to each other. This change in distance would tend to favor introducing background in regions between target and shadow.

To maintain the equal distances between target, background, and shadow labels, we increase the dimension of the label space from one to three and assigned each of the three labels to the coordinate axes. Thus,

$$x_i = \begin{bmatrix} x_{i,1} \\ x_{i,2} \\ x_{i,3} \end{bmatrix} \quad x_i = \begin{cases} \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T & \text{pixel } i = \text{“pure” target} \\ \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^T & \text{pixel } i = \text{“pure” background} \\ \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^T & \text{pixel } i = \text{“pure” shadow} \end{cases} \tag{12}$$

Note that we can now rewrite Equation 11 in the equivalent form,

$$p(x) = \frac{1}{\mathcal{Z}} \exp\left\{ -\alpha \frac{1}{2} \sum_{i=1}^{N} \sum_{j \in \eta_i} \frac{1}{2} (x_i - x_j)^T (x_i - x_j) \right\} \tag{13}$$

so that we appear to have a simple distance measure between two neighboring pixels. With this definition, we can relax the labeling scheme to allow a segmentation label to occupy a continuous set of values under the constraint that it lie in the simplex of this space:

$$x_{i,k} \geq 0 \quad \sum_{k=1}^{3} x_{i,k} = 1 \tag{14}$$

This equation corresponds to a 2-dimensional Brownian motion or Wiener process. Using this method vastly simplifies the mathematics associated with the optimization, at little or no cost in terms of the final results. The prior model is now a continuous function of the variable to be optimized, yet if we constrain the possible values of x_i to the coordinate directions, we obtain the original model.

With x now a continuous variable, we must modify the observation model correspondingly. In fact, we choose to change the observation model very little, leaving it as a conditionally independent convex combination of the Weibull distributions developed previously, i.e. a mixture model:

$$\prod_{i=1}^{N} p(y_i|x_i) = \prod_{i=1}^{N} \sum_{k=1}^{3} x_{i,k} p_k(y_i) \tag{15}$$

where,

$$p_1(y_i) = p(y_i|x_i = [\ 1 \quad 0 \quad 0\ ]^T) \quad p_2(y_i) = p(y_i|x_i = [\ 0 \quad 1 \quad 0\ ]^T) \quad p_3(y_i) = p(y_i|x_i = [\ 0 \quad 0 \quad 1\ ]^T) \tag{16}$$

with $p(y_i|x_i)$ defined appropriately from Equation 4 and the estimates obtained from training based on human segmented data, as described in Section 3.1. Note that this arrangement will make the posterior distribution convex, so we can be assured that we can reach a global minimum.

One useful characteristic of this approach is that the x vector field can have several interesting interpretations. One is that we can view the values in any particular pixel vector, x_i, as the probabilities associated with the three possible labels. Alternatively, we can interpret it as the vector of fractions of each type of label at a given pixel, e.g. $\left[\ \frac{1}{2} \quad \frac{1}{2} \quad 0\ \right]^T$ would have the interpretation of being half shadow and half background. From such a viewpoint, this approach gives much more useful information than is obtained from a simple segmentation.

3.3.2. Quasi-Newton Optimization

Based on the changes to the models we made in the previous section, we can now write the objective function to be optimized. Taking the negative logarithm of Equation 2, we obtain the following minimization,

$$\begin{aligned}
\hat{x}_{\text{MAP}} &= \underset{x}{\text{argmin}} \quad -\log p(y|x) - \log p(x) \\
&= \underset{x}{\text{argmin}} \quad \log \mathcal{Z} \sum_{i=1}^{N} \left[-\log \sum_{k=1}^{3} x_{i,k} p_k(y_i) + \alpha \frac{1}{4} \sum_{j \in \eta_i} (x_i - x_j)^T (x_i - x_j) \right] \\
&= \underset{x}{\text{argmin}} \quad f(x)
\end{aligned} \tag{17}$$

Given an objective function, we now can use a constrained optimization algorithm. Specifically we use a projected Quasi-Newton method:

$$x^{k+1} = \left[x^k - s^k H^k \nabla f(x^k) \right]^+ \tag{18}$$

where $[.]^+$ denotes the projection operation onto the constraint set, H^k is given by the inverse of the diagonal of the Hessian matrix, and s^k is given by an Armijo adaptive step-size rule.[5]

We initialized x^0 with a flat field corresponding to equal probabilities of the three possible segmentation labels. We used a maximum number of iterations of 100, because in practice we found that most of the improvement came within the first 50 iterations. We set α to lie in the range of 5 to 20. These parameter choices appear to give satisfactory results, although a great deal more effort could be put into selecting optimal parameters. This technique is also open to parallelization methods.

To complete the segmentation we need to restore the original discrete space. We accomplish this simply by thresholding the resulting 3-vector at each pixel by assigning it a label based on the nearest unit vector corresponding to a "pure" target, background, or shadow label, as described in Equation 12. For example, a particular pixel $x_i = [0.3, 0.4, 0.3]^T$ would be assigned the background label. Choosing the largest value meshes neatly with the interpretation of x_i as a list of probabilities mentioned at the end of Section 3.3.1, selecting the most probable candidate.

4. RESULTS

The resulting MRF model segmentations appear for the T72 tank at a depression angle of 17° in Figures 6 and 7 at azimuths of approximately 0°, 90°, 135°, 225°. The segmentations shown are a subset of the total performed. We performed T72 segmentations at every azimuth from 0° to 355° at a spacing of approximately 5°. We also ran the MAP segmentation algorithms on the BTR70 transport and BMP2 tank imagery at the same spacing in azimuth and the same depression angle. As stated in Section 3.3.1, the continuous framework can provide much more information

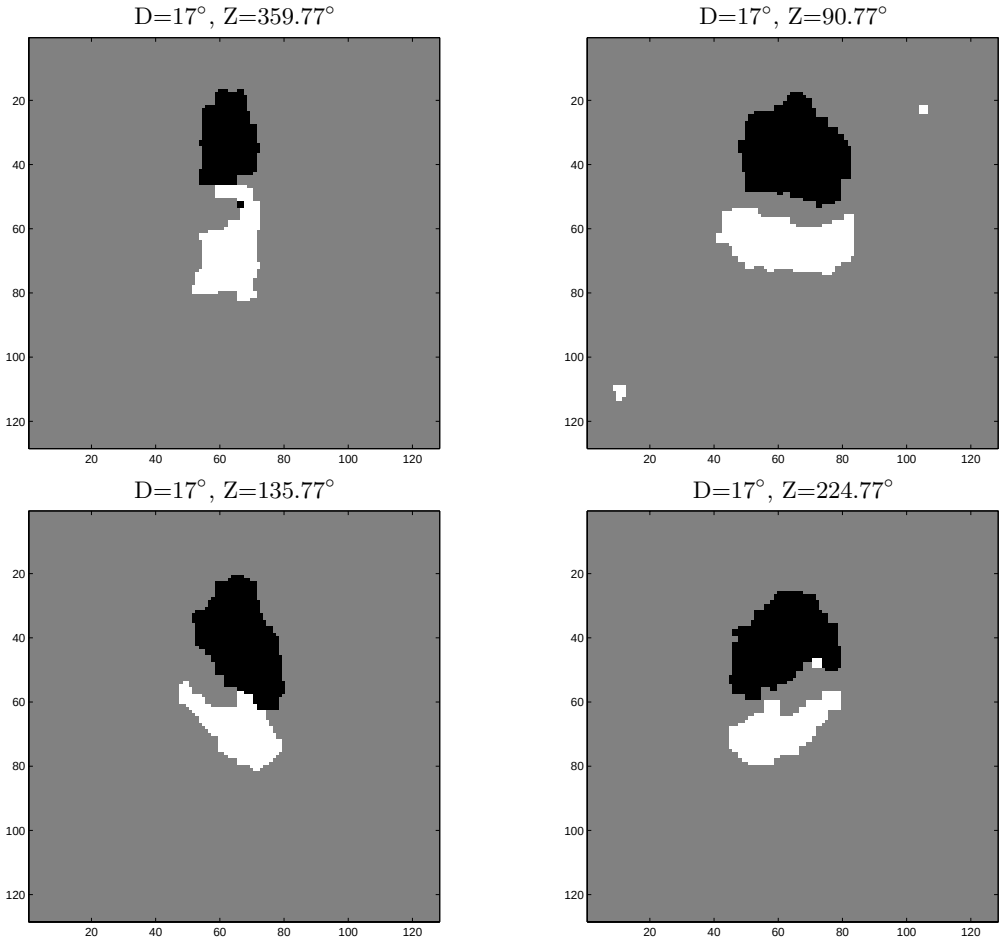

Figure 6. Anisotropic MRF Segmentations of measured T72s, 17° dep

than a simple segmentation. As an example, we present the three representative fraction images for the 0° case in Figure 8. The gun barrel return in the target image is much stronger than what we see in just a segmentation.

To provide an objective comparison between the human, MSTAR, and MRF segmentations we choose to use a 'goodness-of-fit' metric, percent pixels different or PP_d, that we develop and present in another paper at this conference.[1] This metric assumes that human segmentation is a reasonable basis for comparing techniques. By comparing how humans typically do when compared to one another when performing the same segmentation task, we can place an upper bound on how well an automated segmentation can reasonably be expected to do.

We can define the percent pixels different as follows,

$$PP_d = \frac{\sum_{i=1}^{N} |a_i - s_i|}{\max(A,S)} \quad a_i \in \{0,1\} \tag{19}$$

where N is the number of pixels, and a_i and s_i are arrays of binary values denoting a particular region of interest, e.g. target or shadow, defined by an automated segmentation tool or a human, respectively. A and S are just the

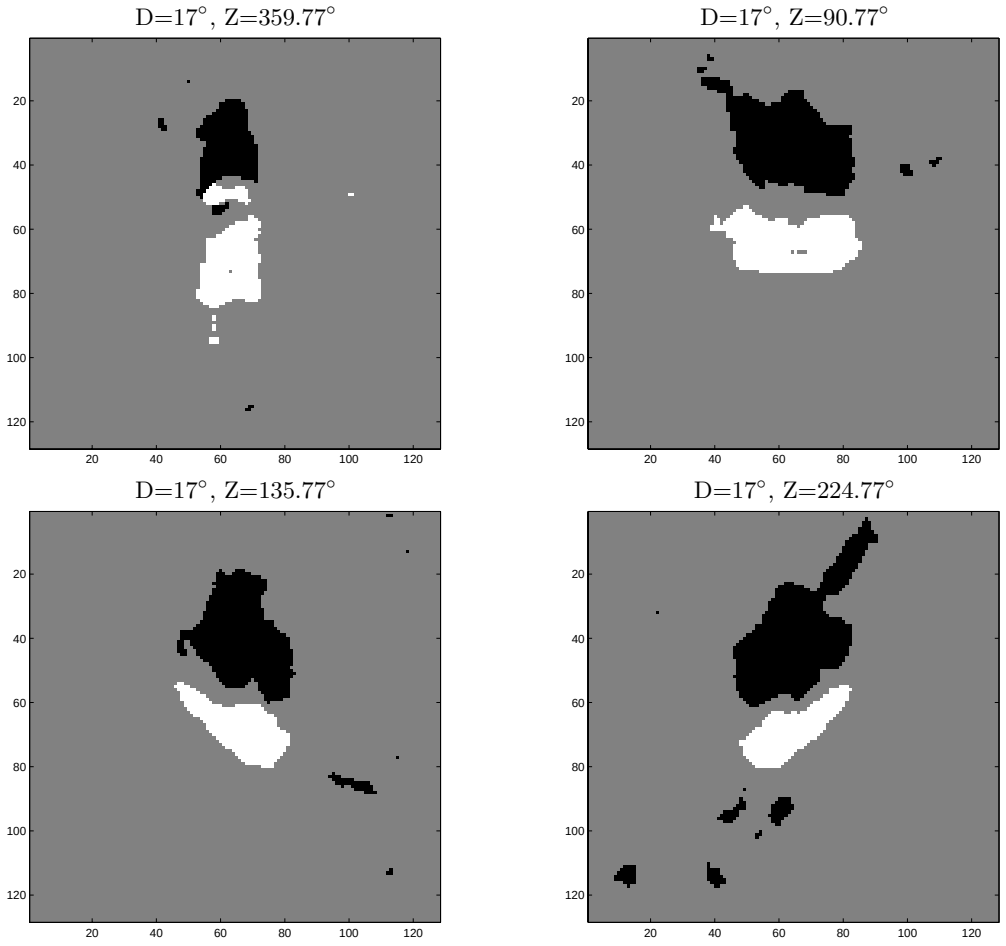

Figure 7. Isotropic MRF Segmentations of measured T72s, 17° dep

total areas of said regions, i.e.

$$A = \sum_{i=1}^{N} a_i \quad S = \sum_{i=1}^{N} s_i \tag{20}$$

The a array could just as easily be a different human analyst being compared to analyst h. If an automated system can come as close as two trained humans can do, then we assume that the automated system works well, within some bound.

We present typical PP_d's for various azimuths of the T72 segmentations in Table 2. We present the average PP_d's for each algorithms on each of the three vehicles tested in Table 3. From these results, it seems that the continuous optimal MRF scheme is not only fast, but it can perform at least as well as the MSTAR segmentation tools.

5. CONCLUSIONS

The results that we have presented clearly favor the continuous MRF approach. The classical gradient-based constrained optimization technique vastly outpaces simulated annealing. It also provides much higher quality segmentations, despite using a more simplistic model.

The continuation MRF segmentation method also performs at least as well as the MSTAR segmentation algorithm, if not slightly better in terms of segmentation quality. The fact that this technique is based on a clear statistical model gives it significant advantages when fusing results with other sensors or sources of information. In addition, the technique not only provides a segmentation map, but can also supply estimates of the probabilities of each possible segmentation label at every pixel. Such information could be extremely useful when estimating error.

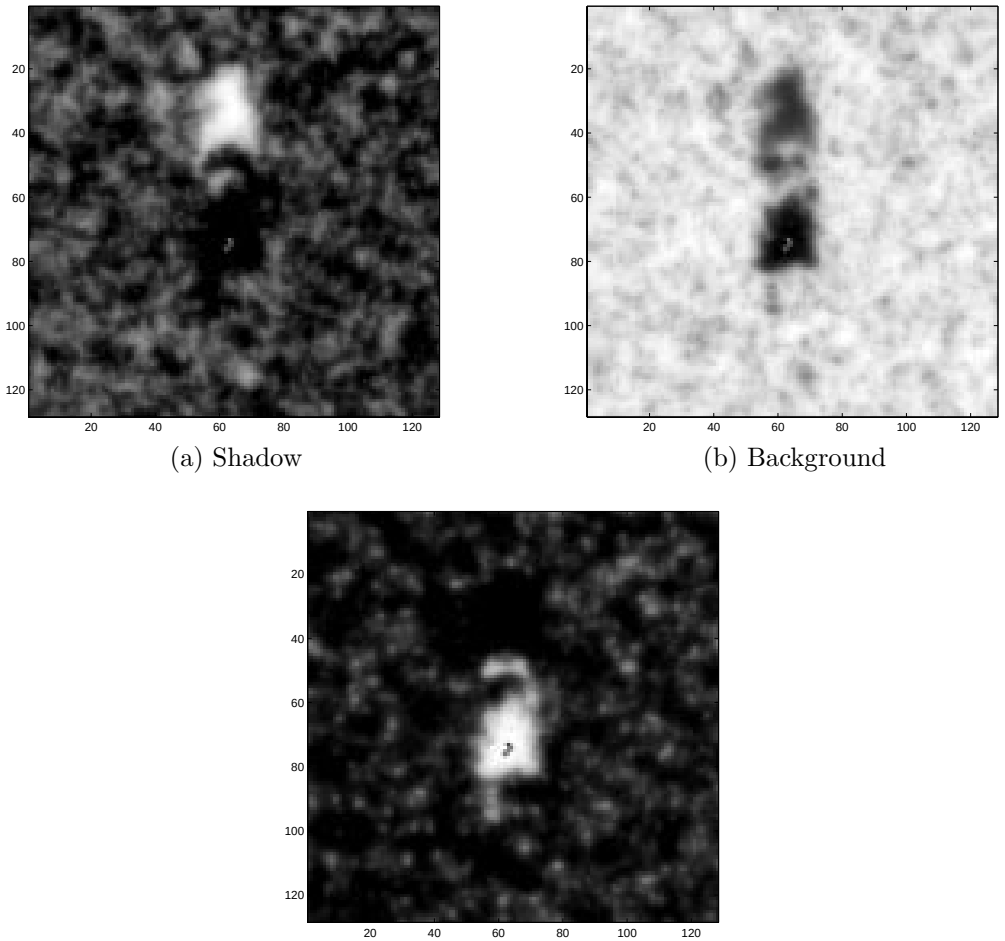

(a) Shadow

(b) Background

(c) Target

Figure 8. Fraction Images for D=17°, Z=359.77°

It may also be possible to extend this technique to more complicated models, including the important anisotropic case. Such advances will require further study, however, as the extension of the anisotropic model to a continuous formulation is neither easy, nor obvious.

6. ACKNOWLEDGMENTS

This work is partially supported by the Army Research Office under Grant ARO DAAG55-97-1-0013, the Air Force Office of Scientific Research under Grant F49620-96-1-0028 and under the Summer Graduate Student Research Program, and the National Institutes of Health under Grant NINDS 1 R01 NS34189.

We would like to thank Ben Booso for his assistance with the MSTAR segmentations, Doug Smith for his help with the human segmentations, and Lou Tamburino, Vince Velten, and Mark Minardi for their helpful advice.

REFERENCES

1. G.J. Power, R.A. Weisenseel, "ATR Subsystem Performance Measures using Manual Segmentation of SAR Target Chips", *Algorithms for Synthetic Aperture Radar Imagery VI*, E.G. Zelnio, editor, Proc. SPIE V 3721, SPIE, Orlando, 1999
2. G. Winkler, *Image Analysis, Random Fields, and Dynamic Monte Carlo Methods: a Mathematical Introduction*, Springer-Verlag, New York, (1995)
3. S.Z. Li, *Markov Random Field Modeling in Computer Vision*, Springer, New York, (1995)

	Shadow			Target				Shadow			Target		
Angle	M	D	C	M	D	C	Angle	M	D	C	M	D	C
0°	46	43	36	48	41	31	180°	37	39	35	42	40	33
15°	31	40	37	26	27	35	195°	28	24	26	33	41	24
30°	30	66	23	22	65	17	210°	20	44	45	37	42	30
45°	19	24	15	31	37	22	225°	29	42	39	33	42	27
60°	20	23	29	21	31	25	240°	22	31	20	24	29	26
75°	18	31	27	23	31	27	255°	35	40	43	29	43	30
90°	32	23	38	34	32	22	270°	50	40	69	23	38	15
105°	34	38	44	19	36	18	285°	13	13	25	25	21	21
120°	20	28	19	31	28	31	300°	18	21	27	25	30	21
135°	39	35	28	42	32	42	315°	15	18	24	23	27	20
150°	47	34	46	33	33	42	330°	23	32	23	39	36	26
165°	45	47	35	46	31	47	345°	23	29	32	23	26	17

Table 2. PPd's for the MSTAR (M), Discrete MRF (D), and Continuous MRF (C) Algorithms - T72, 15° intervals

	Shadow			Target		
Vehicle	M	D	C	M	D	C
T72	29	35	31	32	36	29
BMP2	30	32	30	44	40	34
BTR70	20	23	18	36	38	36

Table 3. Average PP_d's for MSTAR (M), Discrete MRF (D), and Continuous MRF (C) in Each Vehicle Class

4. S. Geman, D. Geman, "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images", IEEE Transactions on Pattern Analysis and Machine Intelligence, v PAMI-6 no 6, pp721-741, 1984
5. D.P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Belmont, MA, (1995)