Representations of thermodynamic variability in the automated understanding of FLIR scenes

Aaron D. Lanterman, Michael I. Miller, Donald L. Snyder

Center for Imaging Science (http://cis.wustl.edu)
Electronic Systems and Signals Research Laboratory, Campus Box 1127
Washington University, St. Louis, Missouri 63130-4899
adl@essrl.wustl.edu, http://essrl.wustl.edu/~adl

ABSTRACT

Grenander's pattern theory offers a unified approach to characterizing variability in complex systems. Automatic target recognition systems for forward-looking infrared sensors must be robust to three kinds of variability: 1) Geometric variability - Target appearances vary with their orientations and positions; 2) Image variability - Target appearances vary with their thermodynamic state, and natural backgrounds consist of widely varying textures; 3) Complexity/scene variability - The number of targets encountered will not be known in advance, and targets may enter or leave the scene at random times.

Pattern theoretic algorithms based on jump-diffusion processes which accommodate variabilities (1) and (3) have been proposed. The diffusions account for (1) by estimating positions and orientations, and the jumps account for (3) by adding and removing hypothesized targets and changing target types. Here we extend that work to better accommodate (2) by summarizing the thermodynamic state of targets with a parsimonious set of variables which become nuisance parameters in the Grenander/Bayesian formulation.

Keywords: pattern theory, jump-diffusion, automatic target recognition, infrared, FLIR

1 INTRODUCTION

Grenander's *pattern theory*[1–4] provides a fundamental framework for the representation and understanding of structures in complex systems. In the context of automatic target recognition, pattern theory provides a basis for the development of computational vision algorithms that are capable of recognizing and describing complex objects contained in natural and cluttered scenes, including both rigid targets and deformable shapes for which current representation and recognition techniques do not work well. The mathematical methodology is organized into the three principal components of image understanding: the *representation of complex scenes*, *image formation and sensor modeling*, and *computational search and optimization decision and/or recognition strategies*.

Section 2 introduces pattern theoretic ideas and illustrates them with some of our previous work with forward-looking infrared (FLIR) sensors. Section 3 discusses the need to extend this work to situations in which the radiant intensities of the scene components are not known *a priori*. A simple likelihood model which incorporates radiant nuisance parameters is presented in Section 4. In Section 5 we apply this likelihood to the estimation of rooftops in FLIR images, where we model the rooftops as parallelograms. A preliminary one-dimensional experiment illustrates the need to place constraints on the number of estimated targets via a "minimum description length" or "minimum complexity" strategy.

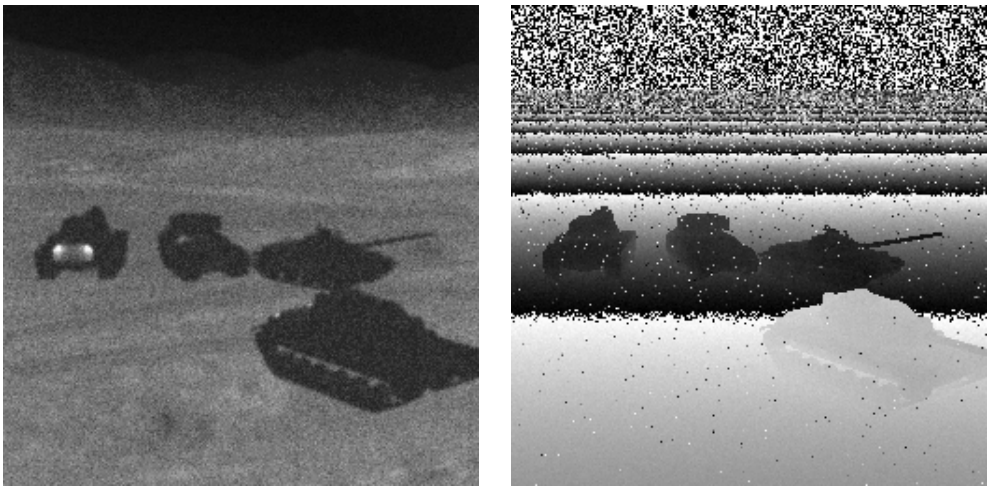

Figure 1: Simulated data sets. Left panel shows PRISM target models (courtesy Dr. Al Curran, Keweenaw Research Center, Michigan Tech. Univ.) superimposed over a real FLIR image, followed by blurring by a point-spread function and corruption by Poisson noise. Right panel shows a laser radar range image demonstrating range ambiguity, anomolous pixels, and range-dependent measurement errors.

2 SUMMARY OF PREVIOUS WORK

We have previously presented sensor likelihood models for FLIR intensity[5,6] and LADAR range images.[7,8] The sensor models have been specified according to the loglikelihood $L(D|\lambda)$ of the sensor detecting a particular set of data D given an ideal image λ. The FLIR model L_{IR}, based on the CCD camera model of Snyder, Hammoud, and White,[9] incorporates blurring by camera optics and photodetector noise characteristics. The LADAR model L_{LR}, based on work by Shapiro,[10,11] incorporates the range ambiguity, measurement errors which increase with range, and uniformly-distributed anomolous pixel measurements whose probability of occurance also increases with range. Figure 1 shows FLIR and LADAR range data simulated from these models. Others have derived sensor likelihood models for low-resolution narrowband sensor arrays and high resolution (range-profile or delay-doppler) radar imaging techniques. The formulation of loglikelihoods for new sensors often follows from knowledge of the physics of the sensor.

Statistical estimation algorithms designed to recover ideal images from collected data have met with tremendous success. Expectation-maximization algorithms for the estimation of λ from D are routinely used to recover images collected by the Hubble Space Telescope[12,13] and reconstruct tomographic images in emission tomography. Other EM algorithms have been derived for direction-of-arrival estimation[14] for narrowband sensor arrays and high-resolution radar imaging.[15,16]

The development of algorithms which can *understand the meaning* of an image λ has gone much more slowly and been met with far less success. In a practical sense, the image λ is an intermediate stage which is easier for a human to interpret than the raw data D. Let x represent the *task-specific information* we want to extract from the data; for instance, the existence (and possibly location and characteristics) of a previously undiscovered galaxy in a HST image, a tumor in an MRI scan, or a scud missile launcher in a synthetic aperature radar (SAR) image. We pass D through some sort of computing device which outputs λ and present λ to a human observer - whether it be an astronomer, physician, or pilot - who extracts x from λ with their training, their experience, and most of all, astonishingly powerful cognitive abilities which humans use almost continuously but whose mechanisms are still not well understood.

The *scene representation* aspect of pattern theory involves precise definition of the parameter space of x. Simple real-valued parameter vectors are often insufficiently rich to characterize complex scenes. In the ATR problem it is fruitful to define a target *template* at some fixed orientation and position; we can act upon this template by continuous *similarity transformations* such as rotation and translation to place it wherever it needs to be in the scene. Individual targets are *generators* which may be combined to form complex multi-target scenes. The parameter space of x is the space of all possible *configurations* of these targets.

For the scenes shown in Figure 1 consisting of rigid ground-based targets, the parameter space for each target is $SO(2) \times \Re^2 \times \mathcal{A}$, where $SO(2)$ is the rotation group acting around the vertical axis through the target body, $\Re^2$ is translation along the ground, and $\mathcal{A}$ is a set of template indices denoting target type, such as M2, M60, T62, etc. A ground-based scene with N targets has the cartesian product $[SO(2) \times \Re^2 \times \mathcal{A}]^N$ as its parameter space. In general the number of targets in the scene will not be known in advance, so the complete parameter space is actually a union of the various N-target parameter spaces $\bigcup_{N=0}^{\infty}[SO(2) \times \Re^2 \times \mathcal{A}]^N$.

Note that if we have airborne targets with full rigid degrees of freedom, then the parameter space for each target would be $SO(3) \times \Re^3 \times \mathcal{A}$, and $\mathcal{A}$ might index target types such as F-15, Mirage, MiG-29, etc. For moving bodies, the parameter space for an individual target is a function of time. Anatomical applications of pattern theory employ shape models with much higher-dimensional groups acting locally on parts of the shape.[17]

We place a prior density $\exp[P(x)]/Z$ on the space of configurations, which we specify here with a logprior $P(x)$ (where Z is a normalizer). Prior knowledge can come in many forms; for instance, moving bodies must obey Newton's laws of motion,[18] and most anatomical structures generally have smooth shapes.[19]

In the same sense that *sensor modeling* characterizes the stochastic mapping from λ to D, pattern theory models the *image formation* process which characterizes the (deterministic or stochastic) mapping from x to λ. We cascade these mechanisms with goal of designing algorithms which estimate the desired parameters x directly from the measured D. Notice from Figure 1 that the image formation mechanism for FLIR, call it $render_{IR} : x \mapsto \lambda_{IR}$, projects infrared intensities onto the detector via perspective projection and obscuration. The mapping for LADAR range, $render_{LR} : x \mapsto \lambda_{LR}$, measures the range to the target under the same perspective and obscuration effects. With image formation models, we can write the sensor loglikelihood in terms of the desired parameters. Sensor fusion is easily achieved by adding the loglikelihoods for the different sensors, for instance,

$$L(D|x) = L_{IR}(D_{IR}|render_{LR}(x)) + L_{LR}(D_{LR}|render_{IR}(x)) \tag{1}$$

For Bayesian estimation, we combine the likelihood and prior information to form a posterior density

$$p(x|D) = \exp[H(x|D)]/Z \tag{2}$$

where $H(x|D) = L(D|x) + P(x)$ is the logposterior.

This density is a highly complicated function of the configuration parameters. Finding the *maximum a posteriori* (MAP) estimate, the x which globally maximizes (2), is a daunting task. It may be more practical (and informative) to sample from the posterior and compute statistics from the resulting samples (recall that for real parameters, the minimum mean square error estimate is just the conditional mean $E[x|D]$). Such *random sampling* methods are abundant in the statistical literature.

Markov chain Monte Carlo techniques allow us to sample from distributions which would be difficult or impossible to simulate directly. In MCMC schemes, we simulate a stochastic process $X(\tau)$ with the property that the distribution of $X(\tau)$ converges in some sense to the distribution defined by (2) as $\tau \to \infty$. The Royal Statistical Society "Meeting on the Gibbs Sampler and Other Markov Chain Monte Carlo Methods"[20–23] provides a stimulating introduction to these methods.

Recall that the parameter space has continuous and discrete characteristics; the orientations, positions, etc. of the targets are continuous, but the varying number of targets and target types are discrete in nature. Any *compu-*

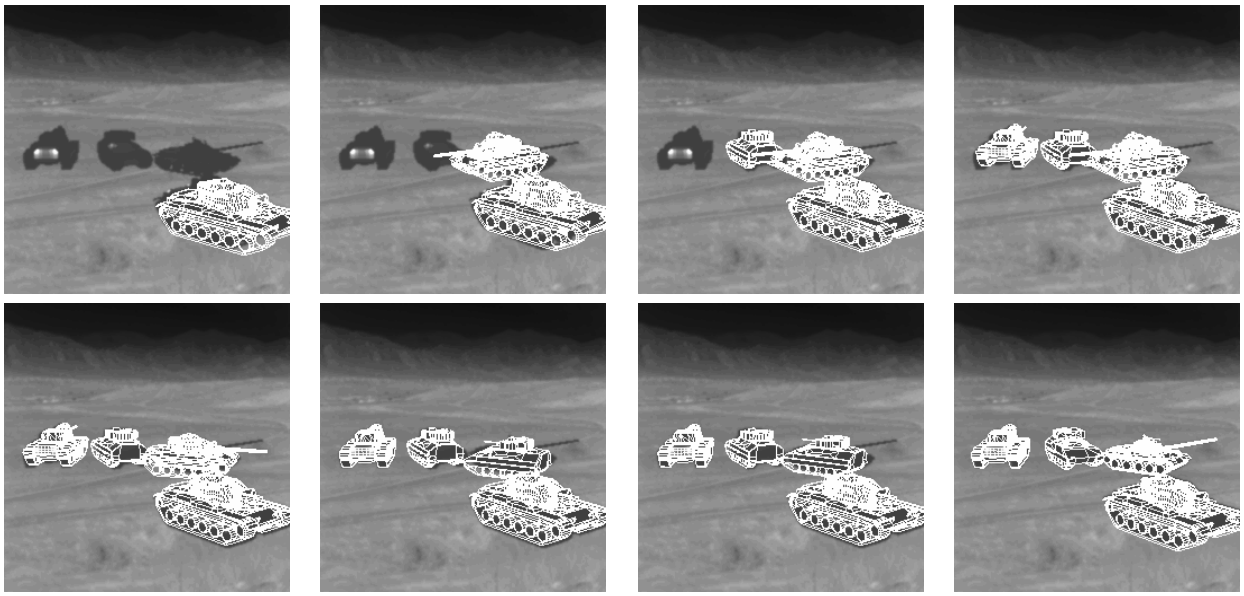

Figure 2: Selected iterations of a sample path of a jump-diffusion process for FLIR data.

tational search or optimization strategy for this parameter space must deal with these hybrid continuous/discrete aspects.

Algorithms based on *jump-diffusion processes* have been proposed for such problems. At exponentially distributed times, the algorithm "jumps" between disconnected areas of the parameter space, adding and removing hypothesized targets and changing their types. Between jumps, a *diffusion* process which drifts along the gradient of the logposterior refines the continuous parameters. A variety of jumping schemes can be used, corresponding to variations on traditional discrete-time Metropolis and Gibbs sampling algorithms. As an alternative to jump-diffusion, Green has proposed purely discrete-time reversible Metropolis-Hasting samplers which search through the continuous parameters via discrete steps.[24]

Figures (2) and (3) show jump-diffusion algorithms[6,8] executing inference on the data shown in Figure (1). The hypothesis is shown as a mesh outline. In these examples the parameterization was powerful enough, and the data strong enough, that the diffusion process appears like a deterministic gradient search, and there is so much probability mass at the "correct answer" that once the process finds the answer it is highly unlikely to leave it.

Jump-diffusion algorithms have also been implemented for tracking multiple aircraft,[25] finding mitochondria[26] in electronmicrographs, and performing inference in some more traditional statistical contexts, such as density estimation for Gaussian mixtures.[27]

3 THERMODYNAMIC NUISANCE PARAMETERS

It is instructive to contrast "noncooperative" military scenarios with "cooperative" biomedical applications. Pattern theoretic shape models have been constructed for estimating the number and shapes of amoebas in optical section microscopy and mitochondria in electronmicrographs. In the amoeba study, the luminescent intensity of the interior and exterior of the amoebas were assumed to have a known mean. In the electronmicrograph study, the texture characteristics of the various organelles were assumed known. In these applications it is reasonable

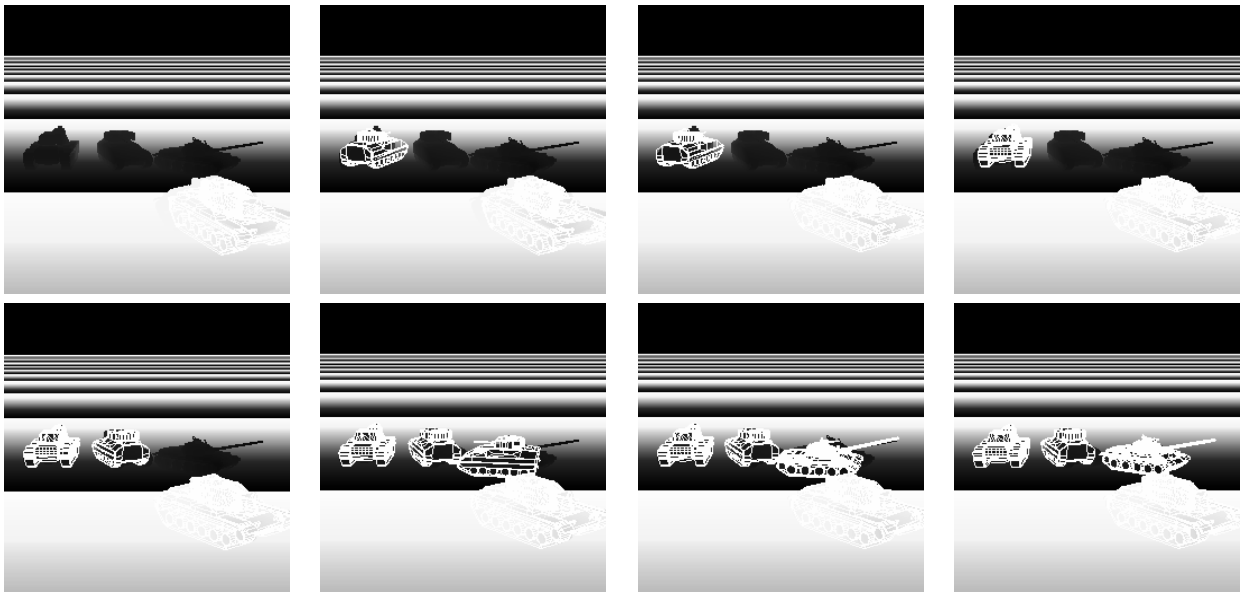

Figure 3: Selected iterations of a sample path of a jump-diffusion process for LADAR range data.

to assume that such information would be available from some auxillary hand-segmented training data. We do not have such a luxury in military applications. Even if previously collected imagery from the region of interest is available, changes in environmental conditions from one mission to the next might cause tremendous changes in radiant intensities, as illustrated by the Grayling I Field Experiment[28] data shown in Figure 4.

Consider the 8-12 micron images shown on the right. In the top image, the sky near the horizon appears brighter than the ground; in the bottom image of the same scene taken at a different time, the ground appears much brighter than the sky. Notice that the heaters in the scene show up more clearly in the 3-5 micron image shown in the upper left corner than they do in the associated 8-12 micron image in the upper right, but the 3-5 micron image in the lower left corner shows no detail at all. Different environmental conditions can provide tremendously different imagery.

Our previous work in ATR for FLIR, demonstrated in Figure 2, has assumed that the radiant intensities of the background and the targets were known *a priori*, with target radiances computed by Keweenaw Research Center's PRISM code.[29] A wide variety of IR simulation codes such as PRISM, GTSIG,[30] IRMA, have proven useful in simulating scenes for testing ATR algorithms, for training human operators, and for predicting the performance of both. Since these codes solve heat transfer equations to produce their result, they require detailed models of target thermodynamics describing internal heat sources and transfer characteristics. The difficulty involved in constructing such heating models has motivated the extraction of radiances from real imagery collected under various environmental and operating conditions and the development of interpolation techniques to predict signatures for other conditions for which data is unavailable. Validating any infrared prediction model, whether it is derived from simulation, empirical data, or both, can be difficult and expensive.

Assuming validated prediction models are available, they require a time series of environmental and operating conditions such as air temperature, solar irradiance, wind speed, engine RPM, etc. as input. Such information is unlikely to be available in practice. Even if validated target models were available, and the data such models require as input was available, we would still need to model the background radiance.

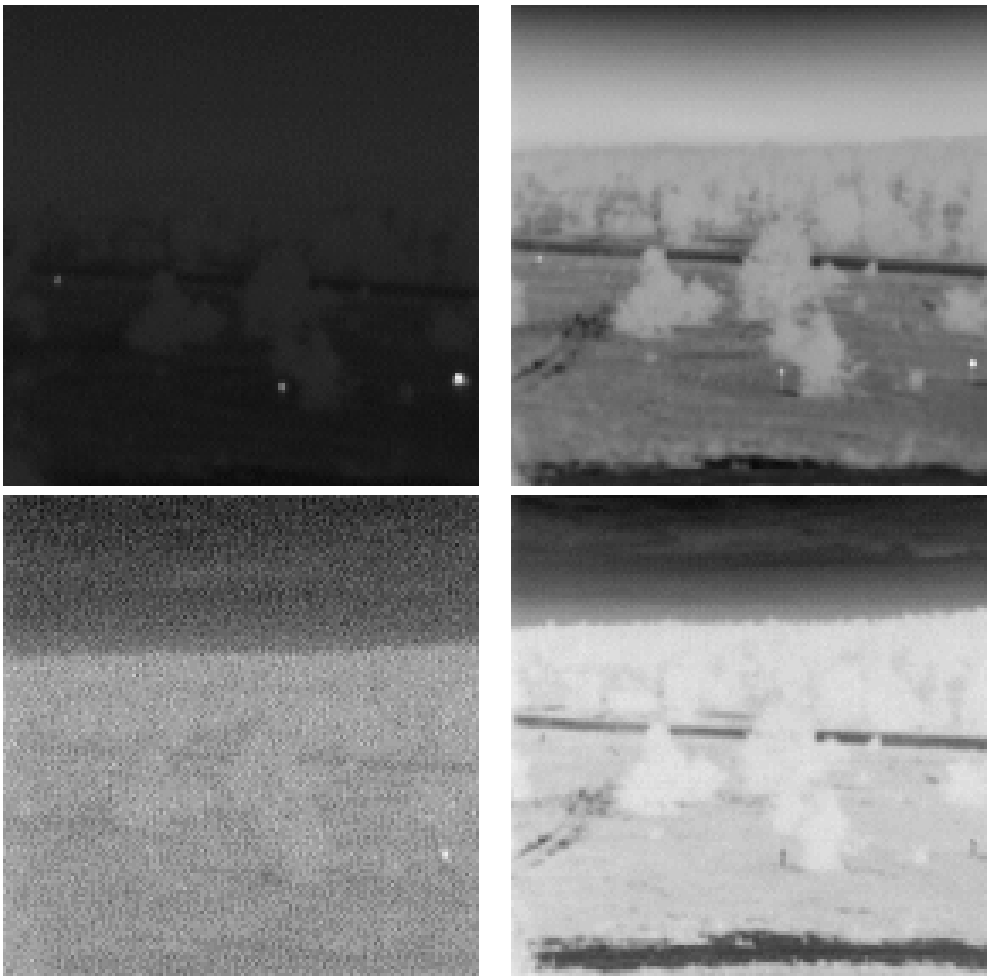

Figure 4: Data from the Grayling I Field Experiment, Smart Weapons Observability Enhancement Joint Test & Evaluation Program. Top and bottom rows are multispectral infrared images of the same scene taken at different times. The images in the left and right columns have passbands of 3-5 and 8-12 microns, repectively. Data courtesy Dr. Bob Guenther, Army Research Office.

4 SIMPLE LIKELIHOOD WITH NUISANCE PARAMETERS

Considering the above challenges, we are motivated to explore algorithms which require no prior knowledge of the intensities of the targets or the background. To best characterize that lack of knowledge, we will treat the intensities as unknown, nonrandom parameters.

4.1 A Poisson likelihood model

Denote by x the parameters we wish to estimate, such as the number of targets, their types, positions, and so on. Treat these as random parameters in the Bayesian formulation.

Suppose that the image array data $D(\cdot)$ is Poisson distributed with mean $\lambda(\cdot)$ which is proportional to the

radiant intensity of objects in the scene. The loglikelihood of the data is[31]

$$L(D|\lambda) = -\sum_i \lambda(i) + \sum_i D(i) \ln \lambda(i) \tag{3}$$

Of course, λ itself is a complicated, highly nonlinear function of the configuration parameters. Let us suppose that the desired parameters x specify a set of image regions R_j and that the intensity $\lambda(\cdot)$ is a constant λ_j across region R_j. A target may consist of several regions; for instance, the gun barrel and engine exhaust of a tank may be much hotter the tank's body, depending on if the gun was recently fired and if the engine is running.

The λ_j's are nuisance parameters. If we treat them as nonrandom variables parameterizing the loglikelihood, it is reasonable to replace the λ_j's in the loglikelihood with their ML estimates for a given realization of x.

In this case, the ML estimate $\hat{\lambda}_j$ of the intensity λ_j of a region R_j containing N_j pixels is simply the average of the data values $D(\cdot)$ in that region

$$\hat{\lambda}_j = \frac{1}{N_j} \sum_{k \in R_j} D(k) \tag{4}$$

Substituting these estimates into the loglikelihood (3) yields

$$L(D|\lambda) = -\sum_j N_j \hat{\lambda}_j + \sum_j \sum_{k \in R_j} D(k) \ln \hat{\lambda}_j \tag{5}$$

4.2 Possible extensions to the simple likelihood model

In real cameras, the light is passed through an optical system which blurs the ideal intensity $\lambda(\cdot)$ by a point-spread function $p(y|x)$. It is straightforward to adopt the EM algorithm derived for restoring blurry images corrupted by Poisson noise[33] to situations where there are known regions of unknown constant intensity. Instead of employing (4), iteration of the EM algorithm would yield estimates of $\hat{\lambda}_j$ which can be substituted into 3 as above. If the point-spread function is Gaussian or can be reasonably approximated as Gaussian, the convolutions in the EM algorithm can be rapidly computed on mesh architectures using difference equations.[34]

If it is unreasonable to assume that the background region is a constant intensity, we could still suppose that foreground target regions are of constant intensity λ_j, but that the background region has spatially varying intensity $\lambda_{BG}(\cdot)$. Certainly, we would expect it to vary smoothly. One possibility is to encourage smoothness background estimates with penalty methods such as Good's roughness,[34–36] written here in 2-D discretized form as

$$\alpha[\gamma(i+1,j) - \gamma(i,j)]^2 + \beta[\gamma(i,j+1) - \gamma(i,j)]^2 \tag{6}$$

where $\gamma = \sqrt{\lambda}$ and the α, β multipliers allow different amounts of smoothing in the horizontal and vertical directions, as is sometimes seen in FLIR imagery. This puts a nearest neighbor Markov random field structure on the background.

If we assume camera blurring is negligable, Good's roughness penalty may be implemented with a simple difference equation, with the differences broken across foreground/background boundaries. If we want to model the camera point-spread, the smoothness constraint can be incorporated in EM algorithm described above with simple one-step-late[37] approaches.

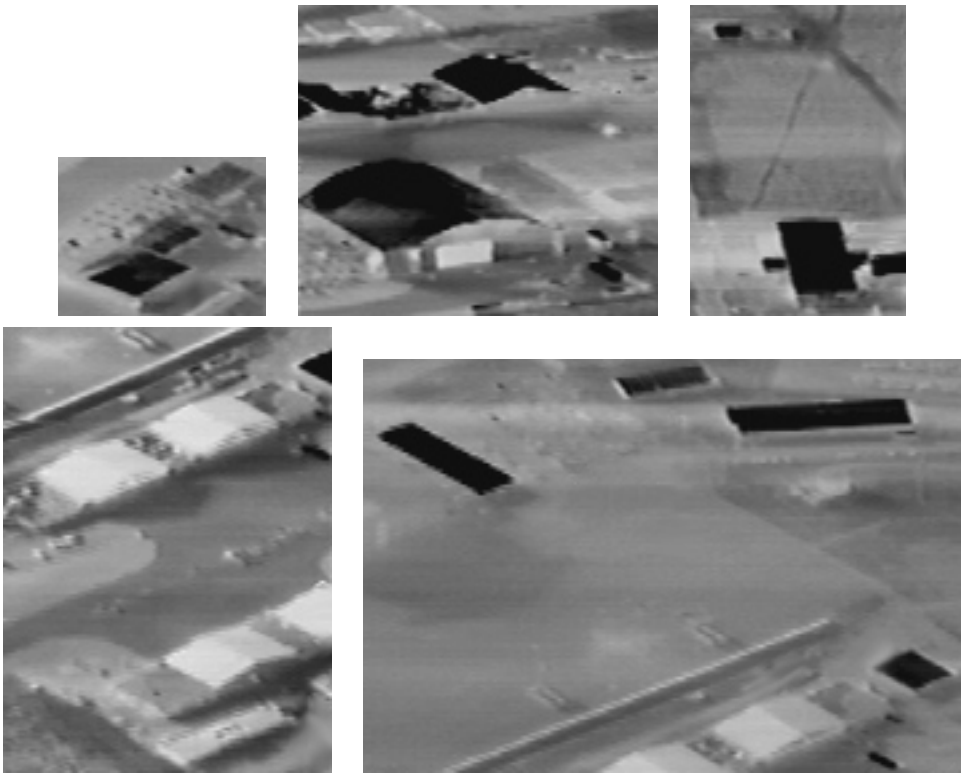

Figure 5: Data courtesy Budimer Zvolanek, McDonnell-Douglas.

5 SHAPE MODELS FOR ROOFTOPS

Rooftops often show up distinctly in FLIR images such as those shown in Figure 5. Sometimes they appear darker than the background, other times they appear lighter, often depending on the various materials used to roof buildings.

The camera projects 3-D scenes onto the 2-D detector via perspective projection. Here, we suppose that individual buildings take up a small enough span of the detector that the projection may locally be considered orthographic, so rectangular rooftops will appear as parallelograms on the detector. This suggests that we use a square as our template and transform it with the group $GL(2)$ (the general linear group of nonsingular 2X2 matrices) followed by the translation group $\Re^2$. The group $GL(2)$ permits scaling, rotation, and skewing of the square, allowing the full range of parallelograms to be expressed. The translation group lets us place the quadrilaterals wherever we need them in the scene.

The prior on parallelogram shapes could be derived from the sizes of targets the system expects to encounter. This would prevent the system from detecting features which are too small to be a target, and also discourage it from thinking of large expanses of background as targets.

5.1 One-dimensional example

We are currently implementing an inference algorithm for finding and estimating the shapes of rooftops and other parallelogram-type structures in FLIR images. To illustrate some important issues, we present a preliminary

experiment in one dimension conducted on a vertical slice through the upper middle image of Figure 5. The one-dimensional buildings were parameterized by their left and right edges. The prior on building width was taken to be uniform over an interval and zero outside of that interval. Building overlap was penalized substantially.

Instead of diffusions, a simple Gibbs sampler was employed to the adjust the building edge positions. At each step of the algorithm, the building positions were cycled through in a fixed order, with the posterior value of the current configuration compared to those obtained by moving the building edge one pixel to the left and one pixel to the right, with the result decided on probabilistically. At exponentially distributed times, a birth move was attempted. The extent of the data was split into 20 equally sized intervals. A random position was chosen in each interval, and the posterior was computed for buildings placed at each of those positions, with building width drawn from the prior. One candidate was probabilistically chosen based on the posterior, and this was randomly accepted or rejected based on the posterior with the additional building compared to the posterior without the additional building.

If this algorithm is run without any explicit constraints on the number of buildings permitted, the algorithm tries to pack the entire extent of the data with buildings, as shown in the top row of Figure 6. The upper left panel shows a time history of the process. The three most substantial image features - the ones that correspond to true structures in the original 2-D image - are found first, but the algorithm continues to pack in eleven more buildings.

This phenomenon was not observed in the work shown in Figure 2 since the background and target intensities were fixed. The algorithm had nothing to gain by birthing a target over the background. Here targets have variable intensity, and the intensity is allowed to be the same or very close to that of the background.

At first, the algorithm will find rooftops with intensities that are quite different than the background, since these yield the largest gains in likelihood. But even after it finds all the true rooftops (and perhaps a few other interesting features), the algorithm may continue to birth parallelogram after parallelogram, estimating intensities which are about the same as the main background, yielding only slight increases in loglikelihood, eventually filling in the background with extraneous objects.

Such pathological behavior can be observed in many model order estimation problems, such as estimation of auditory-nerve discharge rates[38] and determining the number of emitters in direction-of-arrival estimation.[39] The solution is to place constraints on the model order, for which minimum description length[40,41] or minimum complexity[42,43] criterial are well suited.

Here we tried a simple penalty proportional to the number of hypothesized targets, with the proportionality constant adjusted manually. The bottom row of Figure 6 shows the results of introducing this constraint. The algorithm finds three structures which correspond to the structures in the original 2-D image.

6 ACKNOWLEDGEMENTS

This work was supported by ARO DAAH04-95-1-0494, ONR/AASERT N00014-94-1-1135, and ARO/AASERT DAAH04-94-G-0209.

7 REFERENCES

[1] U. Grenander. *General Pattern Theory.* Oxford Univ. Press, 1994.

[2] U. Grenander. Advances in pattern theory: The 1985 Rietz lecture. *The Annals of Statistics*, 17(1):1–30, 1989.

[3] U. Grenander and D.M. Keenan. On the shape of plane images. *SIAM J. Appl. Math.*, 53(4):1072–1094,

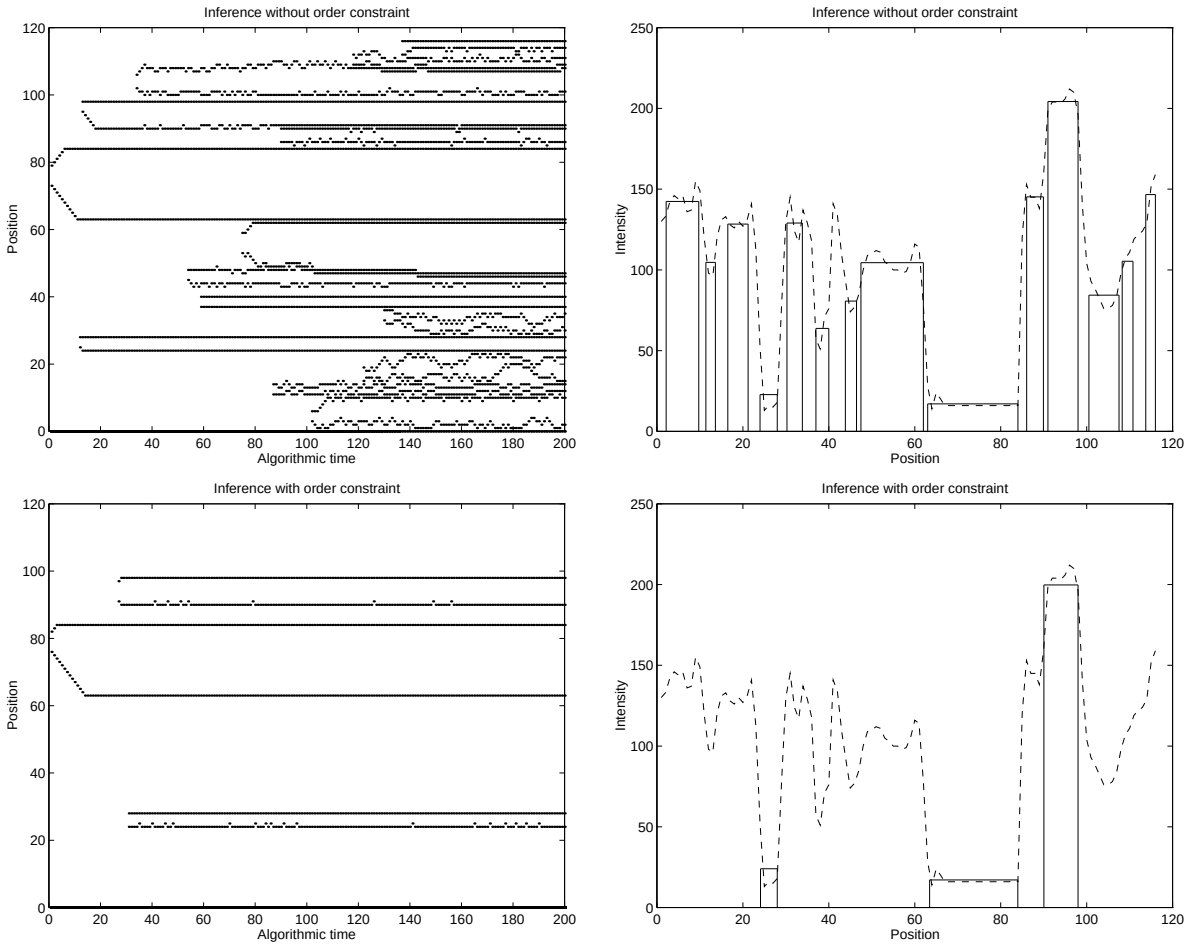

Figure 6: Left column: Evolution of the inference. Dots indicate rooftop boundaries. Right column: Data shown with dashed line; solid lines indicate hypothesized rooftops, derived from averaging samples, along with their estimated IR intensities. Top row illustrates experiment without an order constraint; bottom row illustrates experiment with an order constraint.

August 1993.

[4] D. Mumford. Pattern theory: a unifying perspective. In *Proceedings 1st European Congress of Mathematics.* Birkhauser, 1994.

[5] A.D. Lanterman, M.I. Miller, D.L. Snyder, and W.J. Miceli. Jump-diffusion processes for the automated understanding of flir scenes. In F.A. Sadjadi, editor, *Automatic Object Recognition IV*, volume 2234, pages 416–427, Orlando, FL, April 1994. SPIE.

[6] A.D. Lanterman, M.I. Miller, and D.L. Snyder. Implementation of jump-diffusion processes for understanding flir scenes. In F.A. Sadjadi, editor, *Automatic Object Recognition V*, volume 2485, pages 309–320, Orlando, FL, April 1995. SPIE.

[7] A.D. Lanterman, M.I. Miller, D.L. Snyder, and W.J. Miceli. The unification of detection, tracking, and recognition for millimeter wave and infrared sensors. 2562:150–161, August 1995.

[8] A.D. Lanterman, M.I. Miller, and D.L. Snyder. Automatic target recognition via the simulation of infrared scenes. In *Proc. of the 6th Annual Ground Target Modeling and Validation Conf.*, pages 195–203, Houghton, MI, August 1995. U.S. Army TACOM.

[9] D.L. Snyder, A.M. Hammoud, and R.L. White. Image recovery from data acquired with a charge-coupled-device camera. *Journal of the Optical Society of America A*, 10(5):1014–1023, May 1993.

[10] Jr. T. J. Green and J. H. Shapiro. Maximum-likelihood laser radar range profiling with the expectation-maximization algorithm. *Opt. Eng.*, 31:2343–2354, 1992.

[11] Jr. T. J. Green and J. H. Shapiro. Detecting objects in 3d laser radar range images. *Opt. Eng.*, 33:865–873, March 1994.

[12] M. Faisal, A. D. Lanterman, D. L. Snyder, and R. L. White. Implementation of a modified richardson-lucy method for image restoration on a massively parallel computer to compensate for space-variant point-spread of a charge-coupled-device camera. *J. Optical Society of America A*, in review, 1995.

[13] D. L. Snyder, C. W. Helstrom, A. D. Lanterman, M. Faisal, and R. L. White. Compensation for readout noise in ccd images. *J. Optical Society of America A*, 12:272–283, 1995.

[14] M.I. Miller and D. R. Fuhrmann. Maximum likelihood narrow-band direction finding and the EM algorithm. *IEEE Acoust. Speech and Signal Processing*, 38, No.9:560–577, 1990.

[15] D.L. Snyder, J.A. O'Sullivan, and M.I. Miller. The use of maximum-likelihood estimation for forming images of diffuse radar-targets from delay-doppler data. *IEEE Transactions on Information Theory*, 35(3):536–548, 1989.

[16] P. Moulin, J.A. O'Sullivan, and D.L. Snyder. A method of sieves for multiresolution spectrum estimation and radar imaging. *IEEE Transactions on Information Theory*, 38(2), March 1992.

[17] M.I. Miller, S. Joshi, D. R. Maffitt, J. G. McNally, and U. Grenander. Mitochondria, membranes and amoebae: 1,2 and 3 dimensional shape models. In Kanti Mardia, editor, *Statistics and Imaging*, volume II. Carfax Publishing C., Abingdon, Oxfordshire - England, 1994.

[18] A. Srivastava, N. Cutaia, M.I. Miller, J. A. O'Sullivan, and D. L. Snyder. Multi-target narrowband direction finding and tracking based on motion dynamics. In *Proceedings of the 30th Annual Allerton Conference on Communication, Control and Computing*, Urbana, Champaign, 1992. University of Illinois.

[19] S. Joshi, M. I. Miller, J. McNally, Y. Amit, and U. Grenander. Global shape models for optical sectioning microscopy. *Journal of the Optical Society*, October 1992.

[20] J. Besag and P. J. Green. Spatial statistics and bayesian computation. *J. Royal Statistical Society B*, 55(1):25–38, 1993.

[21] A.F.M. Smith and G.O. Roberts. Bayesian computation via the gibbs sampler and related markov chain monte carlo methods. *J. Royal Statistical Society B*, 55(1):3–37, 1993.

[22] W.R. Gilks, D.G. Clayton, D.J. Spiegelhalter, N.G. Best, A.J. McNeil, L.D. Sharples, and A.J. Kirby. Modelling complexity: Applications of gibbs sampling in medicine. *J. Royal Statistical Society B*, 55(1):39–52, 1993.

[23] Discussion on the meeting on the gibbs sampler and other markov chain monte carlo methods. *J. Royal Statistical Society B*, 55(1):53–101, 1993.

[24] P.J. Green. Reversible jump Markov chain Monte Carlo computation and bayesian model determination. *Biometrika*, 82(4):771–732, December 1995.

[25] M.I. Miller, A. Srivastava, and U. Grenander. Conditional-mean estimation via jump-diffusion processes in multiple target tracking/recognition. *IEEE Transactions on Signal Processing*, 43(11):2678–2690, November 1995.

[26] U. Grenander and M. I. Miller. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society B*, 56(3):549–603, 1994.

[27] D.B. Phillips and A.F.M. Smith. Bayesian model comparison via jump diffusions. Technical report series, *Department of Mathematics*, Imperial College of Science, Technology and Medicine, September 1994.

[28] J.P. Welsh and G.G. Koenig. Description of an approach to validation. In *Proc. of the 6th Annual Ground Target Modeling and Validation Conf.*, pages 62–76, Houghton, MI, August 1995. U.S. Army TACOM.

[29] *Prism 3.2 User's Manual.* Keweenaw Research Center, Michigan Tech. Univ., Houghton, MI, 1995.

[30] J.M. Catchart and Jr. A.D. Sheffer. Generation and application of high-resolution infrared computer imagery. *Optical Engineering*, 30(11):1745–1755, November 1991.

[31] D. L. Snyder and M. I. Miller. *Random Point Processes in Time and Space.* Springer-Verlag, 1991.

[32] H. L. Van Trees. *Detection, Estimation and Modulation Theory, Part I.* John Wiley and Sons, New York, 1968.

[33] D. L. Snyder, A. M. Hammoud, and R. L. White. Image recovery from data acquired with a charge-coupled-device camera. *J. Optical Society of America A*, 10:1014–1023, 1993.

[34] M.I. Miller and B. Roysam. Bayesian image reconstruction for emission tomography: Implementation of the EM algorithm and Good's roughness prior on massively parallel processors. *Proc. of the Natl. Acad. of Sci.*, 88:3223–3227, April 1991.

[35] C.S. Butler and M.I. Miller. Maximum a posteriori estimation for SPECT using regularization techniques on massively-parallel computers. *IEEE Transactions on Medical Imaging*, 12(1):84–89, March 1993.

[36] R. A. Tapia and J. R. Thompson. *Nonparametric probability density estimation.* Johns Hopkins University Press, Baltimore,Md., 1978.

[37] P. J. Green. Bayesian reconstructions from emission tomography data using a modified EM algorithm. *IEEE Trans. on Medical Imaging*, MI-9:84–93, 1990.

[38] K. E. Mark and M. I. Miller. Bayesian model selection and minimum description length estimation of auditory-nerve discharge rates. *Journal of the Acoustical Society of America*, 91, No. 2:989–1002, February 1992.

[39] M. Wax and T. Kailath. Detection of signals by information theoretic criteria. *IEEE Transactions on ASSP*, ASSP-33(2):387–392, 1985.

[40] J. Rissanen. A universal prior for integers and estimation by minimum description length. *The Annals of Statistics*, 11:416–431, 1983.

[41] G. Schwartz. Estimating the dimension of a model. *Annals of Statistics*, 6:461–464, 1978.

[42] J. Rissanen. Stochastic complexity and modeling. *The Annals of Statistics*, 14, no.3:1080–1100, 1986.

[43] J. Rissanen. Stochastic complexity. *Journal of the Royal Statistical Society B*, 49, no.3:223–239 and 252–265, 1987.