Multiscale approaches to moving target detection in image sequences

Thomas G. Allen
Mark R. Luettgen
ALPHATECH, Inc.
50 Mall Road
Burlington, Massachusetts 01803
E-mail: allen@alphatech.com

Alan S. Willsky
Massachusetts Institute of Technology
Room 35-439
Cambridge, Massachusetts 02139

Abstract. We discuss a novel multiscale approach to the detection of moving objects in a sequence of images. The approach is based in part on a multiframe generalization of an optical flow estimation algorithm previously developed by two of the authors. This algorithm provides an extremely efficient multiscale method for estimating optical flow in an image sequence and allows for the temporal accumulation of motion information. Moving objects in a sequence correspond to discontinuities in the true optical flow and, as a result, the residual image associated with the estimated optical flow can be used as a basis for detecting these objects. We propose an approach to detection based on morphological processing of the residual image, and illustrate its potential on real data.

Subject terms: adaptive wavelet transforms; multiscale estimation; automatic target detection; tracking; optical flow.

Optical Engineering 33(7), 2248–2254 (July 1994).

1 Introduction

A requirement for the processing of image sequence data—frequently taken at high temporal sampling rates—can be found in a steadily increasing array of both military and civilian applications, ranging from surveillance and tracking systems to intelligent vehicle highway systems to real-time medical imaging. The reasons for this are several, including the availability of new or enhanced sensing technology, the ever-increasing capabilities (and decreasing cost) of advanced computing systems, and the desire or need to confront increasingly ambitious applications.

Because of these developments, there are a number of critical challenges that must be met by the image processing and computer vision communities, including the apparently enormous computational complexity of image sequence processing problems and the particular technical challenges of different applications. In this paper, we describe an approach to meeting these challenges in one particular context, namely in the spatiotemporal estimation of motion in an image sequence and, more specifically, in the use of such a system to enhance the detection (and ultimately the tracking) of low-contrast moving targets in cluttered backgrounds.

A first point to note is that the estimation of motion—or more precisely what is known in the computer vision community as optical flow—can by itself be a computationally daunting task. For example, methods such as that developed by Horn and Schunck[1] require the solution of coupled partial differential equations (PDEs) for each successive image frame. Discretization of these PDEs for 512×512 images leads to linear systems of equations with over 500,000 variables. The extension of these approaches to the integration of motion information over a number of frames is even more complex and has generally been viewed as being prohibitively expensive computationally. In addition, in contexts such as target tracking, producing estimates of motion is not enough. Specifically, one must also produce an estimate of the accuracy in the optical flow estimates—for example, error variances—for these quantities to be meaningfully interpreted or fused with data from other sensors. Again, the computation of such quantities for standard optical flow formulations is far too complex to be considered.

Finally, there are several particular features of the target detection and tracking problem that add to the challenge. In particular, in standard surveillance systems data from imaging sensors such as infrared cameras or arrays are processed on a frame-by-frame basis producing "detections" analogous to plot reports produced by other sensors such as radar. From that point on, the data provided by the imaging systems are treated simply as a set of measurements of a target or targets taken over time, and only at this point are these "detection reports" integrated over time using various tracking algorithms such as Kalman filters. Although such techniques work well in contexts in which the targets are bright, they are fundamentally suboptimal, because the full imagery data are not coherently processed over time. Consequently, when one is confronted with targets with low contrast relative to

Paper AWT-13 received Jan. 10, 1994; revised manuscript received Mar. 17, 1994; accepted for publication Mar. 17, 1994. This paper is a revision of a paper presented at the SPIE conference on Wavelet Applications, Apr. 1994, Orlando, FL. The paper presented there appears (unrefereed) in SPIE Proceedings Vol. 2242.

the background, performance degrades significantly and possibly catastrophically.

Overcoming this problem, then, requires that we consider the problem of tracking the entire image field of view over time. If we can do this in a computationally reasonable way—including the computation of error statistics—we can then develop enhanced target detection schemes, either by integrating image energy over time to enhance target contrast or by detecting statistically significant anomalies in motion indicative of localized targets moving relative to a cluttered background (taking advantage of the fact that detecting a moving deer in the forest is far easier than detecting one that is standing still).

In this paper, we describe an approach to meeting these requirements. Our starting point is the method developed in Ref. 2 for estimating optical flow at a single instant in time based on two successive frames of image data. This approach uses a new class of multiscale statistical models that qualitatively and quantitatively lead to the same types of flow estimates obtained using other approaches (such as Horn and Schunck's) but with several critical differences and advantages. In particular, the resulting optical flow estimation algorithm is extremely fast. Furthermore, not only does it produce these estimates quickly, but in the process it also yields the desired error variance estimates. Moreover, these estimates and error variances are produced at a complete hierarchy of spatial scales, allowing one to trade off resolution versus accuracy and in fact to define the optimal scale for optimal flow estimation at each point in the image in the sense that estimation of finer scale detail at that point is not justified based on the level of uncertainty in the estimates predicted by the error variance calculation.

Starting from this optical flow estimation algorithm, we extend the approach in two ways. First, we describe a method for estimating optical flow over time rather than simply estimating it independently in each frame, thus allowing accuracy to improve over time as we gain information. Conceptually, what we describe is an extremely efficient temporal Kalman filter of extraordinarily high dimension—the state consists of a 2-D flow vector for each pixel—allowing us to solve problems that previously had been considered beyond the range of practicality.

The second extension that we describe here is the use of the measurement residual, i.e., the image of violations of the so-called "brightness constraint" (see Sec. 2). Areas in which these residuals are of statistically significant size (i.e., large with respect to the corresponding error standard deviation computed by our algorithm) indicate locations at which motion discontinuities caused by targets may be present. We propose an approach to detection based on morphological processing of the residual image, and illustrate its potential on real data.

This paper is organized as follows. In Sec. 2, we briefly review the multiscale signal processing framework developed in Refs. 3 to 5 along with the optical flow problem formulation and estimation algorithm as developed in Ref. 2. In Sec. 3, we describe our multiframe extension of the optical flow estimation algorithm that allows us to exploit temporal smoothness in the optical flow field, and we illustrate the extended algorithm on a real image sequence. That sequence contains a helicopter that is moving with respect to the background clutter, and in Sec. 4, we discuss our approach to target detection based on morphological processing of the optical flow residual image. Conclusions and areas of ongoing and future research are discussed in Sec. 5.

2 Multiscale Model-Based Algorithms for Optical Flow Calculation

The optical flow field is the 2-D distribution of apparent velocities on the imaging surface corresponding to the variation of brightness patterns in a series of image frames. The concept of optical flow is commonly defined through the brightness constraint equation[1]:

$$\frac{\partial}{\partial t}E(z_1,z_2,t)+\nabla E(z_1,z_2,t)x(z_1,z_2,t)=0 \ , \tag{1}$$

where $x(z_1,z_2,t)$ is the optical flow vector field, $E(z_1,z_2,t)$ is the image intensity at point (z_1,z_2) in the image at time t, and $\nabla E(z_1,z_2,t)$ is the gradient of the image intensity:

$$\nabla E(z_1,z_1,t)=\left[\frac{\partial}{\partial z_1}E(z_1,z_2,t),\frac{\partial}{\partial z_2}E(z_1,z_2,t)\right] \ . \tag{2}$$

The brightness constraint equation supplies a single constraint on the optical flow field at each point. However, there are two unknowns to recover at each pixel—the two components of $x(z_1,z_2,t)$—and hence the problem of determining these at each pixel using only Eq. (1) is ill-posed. Additional information must be supplied to completely specify the flow field; the usual approach is to regularize the problem by postulating some additional structure. One of the first, and most well known, regularization procedures is that proposed by Horn and Schunck[1] in which a global smoothness constraint (SC) is imposed on the optical flow field. The optical flow field is found in this case by solving the following optimization problem:

$$\hat{x}_{SC}=\text{argmin}_x\int\int R^{-1}\left[\frac{\partial}{\partial t}E(z_1,z_2,t)+\nabla E(z_1,z_2,t)x(z_1,z_2,t)\right]^2 + \|\nabla x(z_1,z_2,t)\|^2 \ \mathrm{d}z_1 \ \mathrm{d}z_2 \ , \tag{3}$$

where R is a parameter specifying the relative importance of the two constraints. We refer to the solution to the optimization problem in Eq. (3) as the SC solution. It can be shown using the calculus of variations that the solution of Eq. (3) satisfies a pair of coupled PDEs that when discretized lead to a sparse, but large, set of linear equations. Solution techniques encountered in the optical flow literature include Gauss-Seidel methods, successive over-relaxation (SOR) techniques, and multigrid methods.[1,6,7] The problem with these methods is that they are notoriously expensive computationally, and this was a principal motivation for the development of a multiscale optical flow estimation algorithm developed in Ref. 2 and briefly reviewed next.

The multiscale regularization (MR) algorithm is a method of computing optical flow inspired, in part, by the structure of the smoothness constraint. As discussed in Ref. 8, the optical flow problem formulation in Eq. (3) has an equivalent formulation in an estimation-theoretic context. Specifically, it corresponds to a statistical model in which the error in the brightness constraint is assumed to be spatially white and in which the two components of the optical flow are modeled as independent random fields, each of which has a zero-mean, spatially white gradient. As a result, the smoothness constraint essentially corresponds to modeling each component of the optical flow as a spatial Brownian motion, i.e., as a statistically self-similar, fractal process with a $1/|f|^2$ generalized spectrum.[9] Given this, a natural idea is that of using alternate prior statistical models—corresponding to alternative penalty terms to that in Eq. (3)—that possess the same type of fractal characteristics but that lead to computationally more attractive problem formulations. In Ref. 2 we propose the use of a fractal-like class of prior models recently introduced in Refs. 3 to 5. As shown there and in Refs. 10 and 11, this model class is extremely rich, and also leads to very efficient estimation and likelihood calculation algorithms.

These models represent the flow field at a set of scales of increasing resolution, $m=0, \ldots, M$. In particular, at scale m, we define a set of optical flow vectors:

$$x_m(i,j),\ i,j=1,2,\ldots,2^m \ . \tag{4}$$

At the coarsest level ($m=0$), the optical flow consists of a single vector aggregating the entire flow field. Each subsequent scale increases the number and resolution of optical flow vectors by a factor of 4, until at the finest level ($m=M$) the flow field has the same resolution as the raw image data. This series of scales is defined on a quadtree structure, as illustrated in Fig. 1, wherein each node at level m has four descendants at level $m+1$, $m=0, \ldots, M-1$.

The quadtree is not merely a data structure used in implementation of specific algorithms. Much more fundamentally, it defines the underlying structure of the multiscale optical flow model—the statistical dependencies between pixel values that will be exploited by an optical flow algorithm. The multiscale model is a downward (coarse-to-fine) model on the tree in which the value of the state at each node is a linear interpolation of its parent, plus an additional detail term. The measurements at each node in this scale recursive model are linear transformations of the node state, corrupted by Gaussian noise. For optical flow estimation, the measurements will be available only at the finest level and these measurements correspond precisely to the brightness constraint, as discussed in more detail in Ref. 2. The algorithm developed there uses these measurements to compute at each node the best estimate of the optical flow vector at each pixel

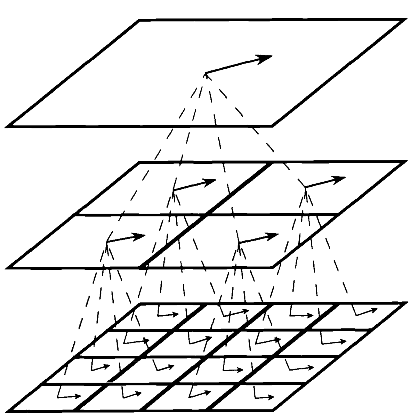

Fig. 1 Multiresolution model quadtree structure.

(in a least-squares sense) and leads to root-mean-square error performance that is comparable to that of smoothness constraint-based formulations such as that in Ref. 1, but with order of magnitude reductions in computational cost. Refer to Ref. 2 for detailed algorithm development, complexity analysis, and examples.

3 Multiframe Extension of the MR Algorithm

Applying the MR algorithm to a sequence of images is accomplished through a series of independent applications of the algorithm. The optical flow field is computed for the first frame, discarded, computed for the second frame, discarded, and so on, as if each field was independent of all prior and subsequent fields. It is logical to expect that the flow estimates in a given frame would be useful in improving subsequent estimates and we describe next a multiframe extension of the MR algorithm that allows us to do just that.

There are two basic approaches to temporal modeling of optical flow, Eulerian and Lagrangian.[12] The Eulerian model is viewer based, i.e., temporal coherence is imposed on the flow vectors as opposed to the underlying physical scene. In contrast, the Lagrangian model is object based, and so imposes the temporal coherence on the surface elements within the scene, not directly on the optical flow.

A first-order Eulerian dynamic model can be given by

$$\frac{\partial}{\partial t} x(i,j,t) = w(i,j,t) \ , \tag{5}$$

where $w(i,j,t)$ is a zero-mean space-time white noise with covariance q^{-1}. A first-order Lagrangian dynamic model is similar, but employs the total derivative with respect to time in place of the partial derivative:

$$\frac{\mathrm{d}}{\mathrm{d}t} x(i,j,t) = w(i,j,t) \ . \tag{6}$$

The Eulerian model is straightforward to implement, as the discrete counterpart of Eq. (5) is simply

$$x(i,j,k) = x(i,j,k-1) + w_t(i,j,k) \ , \tag{7}$$

$$w_t(i,j,k) \sim N(0,q_t I) \ . \tag{8}$$

Implementation of the Lagrangian model is far more difficult, requiring a correspondence between the image coordinates and the moving surface elements.[12] Such a correspondence can be calculated by tracking the surface elements within the frame.[13]

In Ref. 12, it is argued that under standard optical flow conditions the Eulerian method is a good approximation to the Lagrangian method. This is partially because the spatial optical flow derivatives have relatively small magnitudes, because of the smoothness-type constraints. In addition, the optical flow itself will have small magnitudes, because large object motions cannot be estimated from locally computed gradients. Therefore, the partial derivative in Eq. (5) is approximately equal to the total derivative in Eq. (6).

Our multiframe generalization of the MR algorithm is based on an Eulerian approach applied at the finest scale. That is, we assume

$$x_M(i,j,k) = x_M(i,j,k-1) + w_t(i,j,k) \ , \tag{9}$$

$$w_t(i,j,k) \sim N(0,q_t I) \ . \tag{10}$$

Our approach to including at time k the optical flow information from time $k-1$ involves two steps. First, we predict ahead the MR algorithm estimate of $x_M(i,j,k-1)$ using the (trivial) dynamics in Eq. (9)—the estimate of $x_M(i,j,k)$ based on the MR estimate of $x_M(i,j,k-1)$ is the MR estimate of $x_M(i,j,k-1)$—and we use the associated Lyapunov equation to update the smoothed error covariance provided by the MR algorithm. The predicted estimate is then combined with the MR estimate of $x_M(i,j,k)$ based on the measurement available at (i,j) at time k, which comes from the brightness constraint equation. The combination of the two estimates is based on viewing them as independent estimates of the same random variable, $x_M(i,j,k)$, and hence can be computed using standard linear least-squares fusion equations (see e.g., Ref. 3). The fused estimate and error covariance then are used as a basis (in place of the estimate based solely on the measurement at time k) for completing the rest of the MR algorithm calculations as usual.

We illustrate the multiframe extension of the MR algorithm on a sequence of real images of a helicopter against a natural background. Each frame in the sequence is an 8-bit gray-scale image of 480×480 pixels. The selection of this sequence was motivated by three factors: (1) the objects and background are real, not simulated; (2) the background is reasonably complex, with both varied vegetation and terrain; and (3) the target is difficult, if not impossible, to extract from the background using nontemporal methods. Figure 2 illustrates a frame in this sequence.

As described previously, the MR algorithm not only provides the optical flow at a resolution equivalent to the image itself, it also supplies the optical flow and associated 2×2 error covariances at increasingly coarse levels of resolution. As discussed in Ref. 2, the trace of these multiscale covariances can be employed to determine an optimal level of resolution for each region of image flow, and Fig. 3 illustrates the optimal resolution as a function of position. Given that each location in the flow field has an optimal representation resolution, and that the flow field itself is multiresolution, it is possible to construct the optimal representation optical flow. For each location in the image this is simply the optical flow estimate at the optimal representation resolution. Figure 4 contains the optimal representation optical flow for frame 2.

Figure 5 displays the optimal representation resolution at frame 2 for the multiframe MR algorithm (computed with a value of $q_t = 0$). Comparing these results with those of the

Fig. 3 Optimal representation resolution.

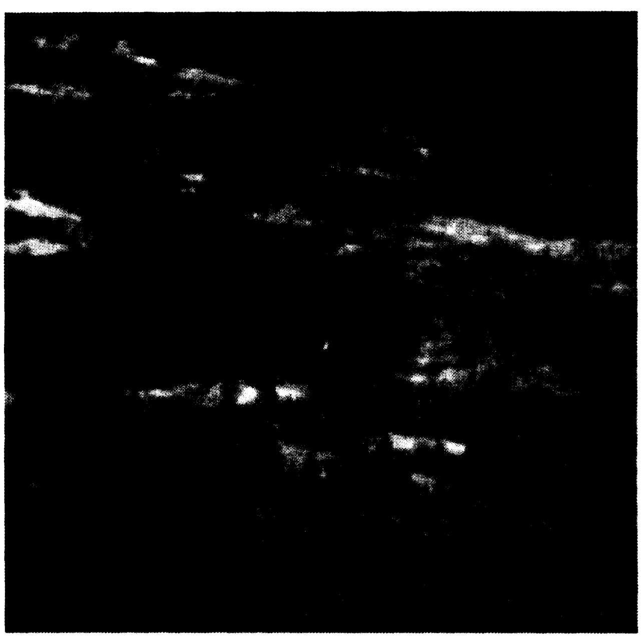

Fig. 2 Raw helicopter image.

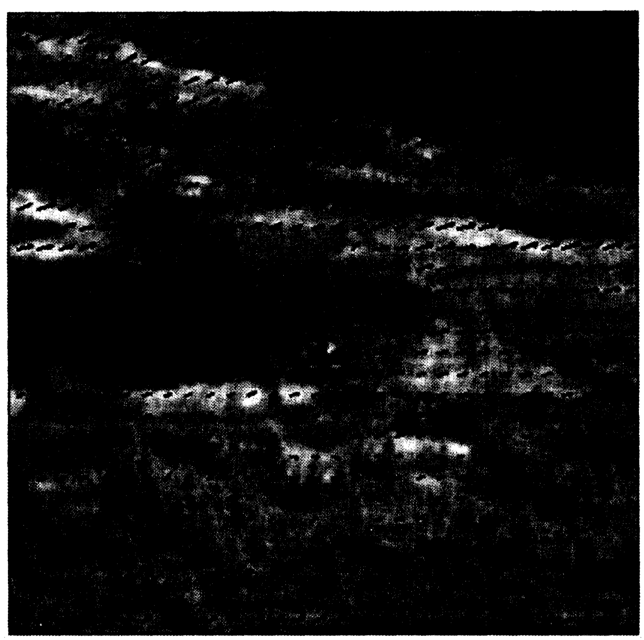

Fig. 4 Optimal representation optical flow.

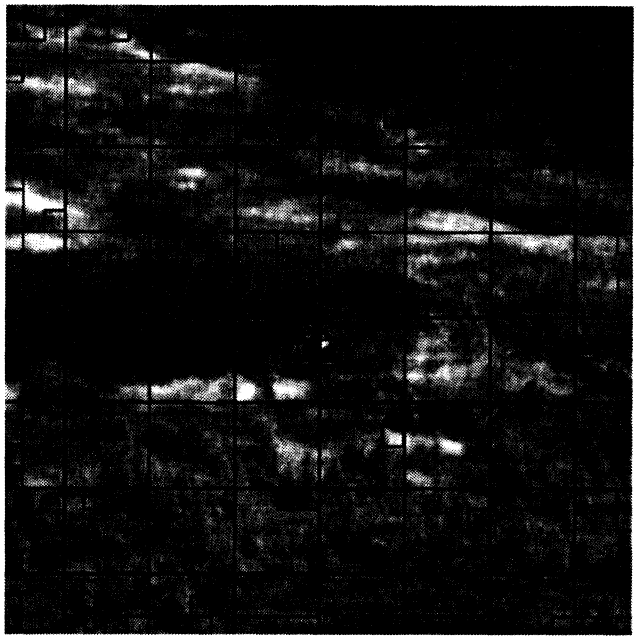

Fig. 5 Optimal representation resolution: multiframe MR algorithm.

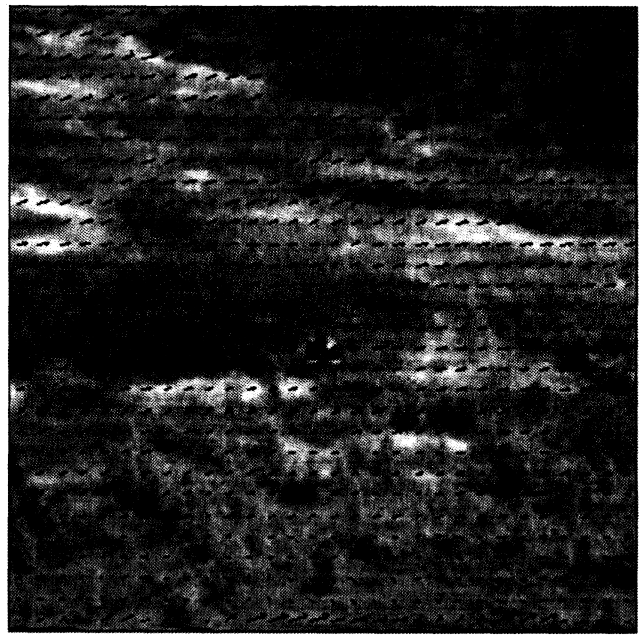

Fig. 6 Optimal representation optical flow: multiframe MR algorithm.

standard MR algorithm in Fig. 3 shows that the inclusion of multiframe temporal information improves the representation resolution in three ways: (1) the region around the helicopter fuselage resolved at level 5 has been increased in size, (2) the resolution at the center of the helicopter has been increased from level 5 to level 6, and (3) the large background patches have been partially reduced by level 4 to level 3 (homogenous background regions should be resolved at relatively coarse levels). The optimal representation resolution optical flow for the multiframe MR algorithm is shown in Fig. 6. The smoothness of this flow appears superior to that of the MR results in Fig. 4, particularly at the bottom of the image.

4 Moving Target Detection

There are several approaches to the development of moving target detection algorithms that exploit the optical flow estimates. One approach would be to segment the image based on the optical flow field.[14–16] Such a method is attractive in that it can be combined with other methods of image segmentation, such as texture estimation, to improve the separability of the scene. The approach to target detection that we take here is based on exploitation of the artificial smoothness of the flow imposed by the global smoothness constraint. Specifically, the smoothness constraint tends to obscure abrupt changes in the field resulting from the motion of small objects. The measurement data itself, i.e., the temporal and spatial derivatives, do not necessarily support the smoothed flow estimates. It is therefore possible to locate regions of discontinuities in the flow field by locating where the measurement data disagrees with the flow estimates. This is a definition of the measurement residual at each pixel, given by

$$\upsilon(i,j) = \frac{\partial}{\partial t}E(i,j) + \nabla E(i,j)\hat{x}_M(i,j) \ , \tag{11}$$

where $\hat{x}_M$ is the optimal representation resolution optical flow.

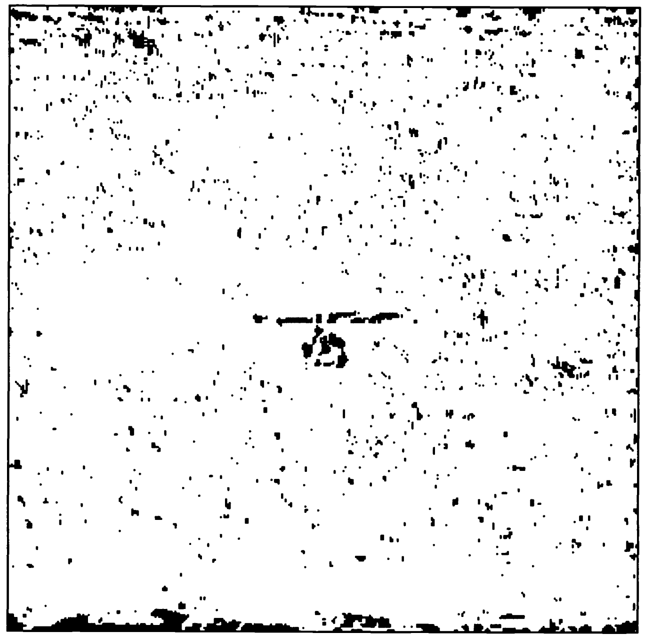

Fig. 7 Brightness constraint residuals: threshold = 5.

The residuals computed for the multiframe MR algorithm at frame 2 are displayed in Fig. 7. These are actually the absolute values of the residuals (because only magnitude is important for target detection), thresholded to create the binary image. The helicopter is quite visible. More importantly, the helicopter is the only element in the image with coherent structure (except for edge artifacts). It is particularly encouraging to notice the distinctness of the helicopter rotors. The rotors are less visible in the individual raw frames than the fuselage, but have significant temporal gradients, demonstrating the importance of employing optical flow information in target detection.

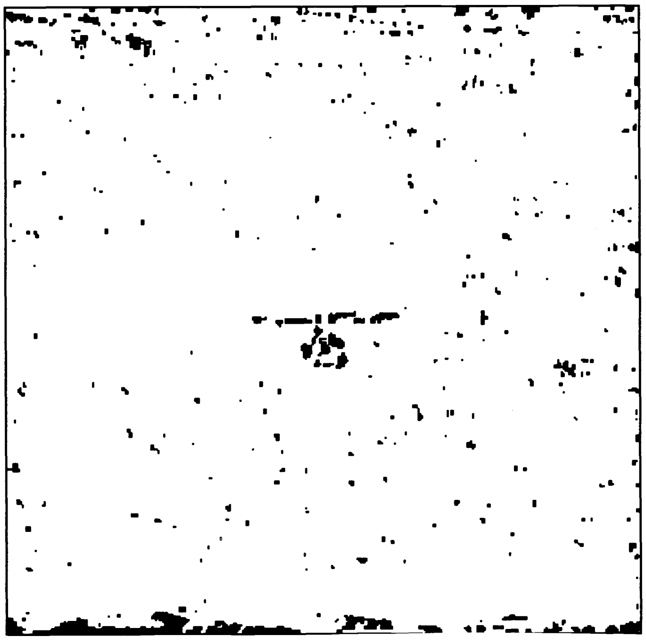

Fig. 8 Morphological processing: opening.

As an attempt to further increase the image of the helicopter, several simple binary morphological operations were applied to the thresholded image. Figure 8 shows the results of opening the image with a 2×2 square structuring element.[17] This result is only an example of operations that could be performed on the thresholded residuals. For instance, an obvious addition to the morphological transformations is an intermediate operation to remove isolated pixels, or small clusters of pixels. Such a process would eliminate much of the randomly structured residuals, leaving areas where a significant number of frame elements have moved in a coherent fashion.

5 Conclusions

The major objective of this effort was to determine whether multiresolution-based image analysis algorithms could meet the requirements for practical image analysis and information extraction in an advanced surveillance and tracking system. Our results demonstrate the considerable promise that these algorithms have for practical implementation and integration into advanced surveillance systems. Specifically, the multiscale algorithm we have proposed provides a direct mechanism for spatiotemporally coherent processing of image sequence data for enhanced target detection, by producing motion estimates over an entire image frame. This can be of considerable value for the detection and tracking of low observable targets for which methods based on standard single-frame detection algorithms fail because of low single-frame SNR. Moreover, the error statistics that are available with these estimates will be useful for the optimal fusion of image-derived estimates into a multitarget and possibly multisensor tracking system, and for system performance assessment. The multiscale motion estimates allow for the possibility of discriminating between areas in which fine-scale target motion can be discerned and areas in which only coarser background motion can be estimated. This is of potential use both for fully automated algorithms, in which this information can be used to provide "hand-over" information to a target tracking algorithm when fine-scale target motion is detected, and for operator cueing. Thus, multiresolution information provides the critical service of filtering the vast amount of image data for the operator or pilot and indicating the regions in which he or she should focus attention. Finally, the multiscale algorithm provides measurement residuals in which discontinuities in the motion field—because of targets moving relative to the background—are enhanced. These residual images provide an enhanced data set for automatic target detection and possibly identification.

Acknowledgments

This work was supported in part by Naval Air Warfare Center (NAWC) Weapons Division contract N60530-92-C-0064. The authors would like to thank Dr. Gary Hewer of NAWC for his support of this work.

References

1. B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.* **17**, 185–203 (1981).
2. M. Luettgen, W. C. Karl, and A. S. Willsky, "Efficient multiscale regularization with applications to the computation of optical flow," *IEEE Trans. Image Process.* **3**(1), 41–64 (1994).
3. K. C. Chou, A. S. Willsky, and A. Benveniste, "Multiscale recursive estimation, data fusion, and regularization," *IEEE Trans. Auto. Control* **39**(3), 464–478 (1994).
4. K. C. Chou, A. S. Willsky, and R. Nikoukhah, "Multiscale systems, Kalman filters, and Riccati equations," *IEEE Trans. Auto. Control* **39**(3), 479–492 (Apr. 1994).
5. M. Basseville, A. B. Benveniste, K. C. Chou, S. A. Golden, R. Nikoukhah, and A. S. Willsky, "Modeling and estimation of multiscale stochastic processes," *IEEE Trans. Inform. Theory* **38**(2), 766–784 (1992).
6. J. Prince and E. McVeigh, "Motion estimation from tagged MR image sequences," *IEEE Trans. Med. Imaging* **11**(2), 238–249 (1992).
7. D. Terzopoulos, "Image analysis using multigrid relaxation methods," *IEEE Trans. Pattern Anal. Mach. Intell.* **8**(2), 129–139 (1988).
8. A. Rougée, B. C. Levy, and A. S. Willsky, "An estimation-based approach to the reconstruction of optical flow," Report No. LIDS-P-1663, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA (Apr. 1987).
9. R. Szeliski, *Bayesian Modeling of Uncertainty in Low-level Vision,* Kluwer-Academic, Norwell, MA (1989).
10. M. Luettgen, W. C. Karl, A. S. Willsky, and R. R. Tenney, "Multiscale representations of Markov random fields," *IEEE Trans. Signal Process.* **41**(12), 3377–3396 (1993).
11. M. Luettgen and A. Willsky, "Likelihood calculation for a class of multiscale stochastic models, with application to texture discrimination," Report No. LIDS-P-2186, MIT Laboratory for Information and Decision Systems (Apr. 1994), to appear in *IEEE Trans. Image Process.*
12. T. M. Chin, "Dynamic estimation in computational vision," PhD Thesis, Electrical Engineering and Computer Science, Massachusetts Institute of Technology (1991).
13. M. J. Black and P. Anandan, "A model for the detection of motion over time," in *Proc. 3rd Int. Conf. on Computer Vision,* pp. 33–37, IEEE Press, Osaka, Japan (1990).
14. D. W. Murray and B. F. Buxton, "Scene segmentation from visual motion using global optimization," *IEEE Trans. Pattern Anal. Mach. Intell.* **9**(2), 220–228 (1987).
15. B. G. Schunck, "Image flow segmentation and estimation by constraint line clustering," *IEEE Trans. Pattern Anal. Mach. Intell.* **11**(10), 1010–1027 (1989).
16. W. B. Thompson, "Combining motion and contrast for segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.* **2**(6), 543–549 (1980).
17. R. M. Haralick, S. R. Sternberg, and X. Zhuang, "Image analysis using mathematical morphology," *IEEE Trans. Pattern Anal. Mach. Intell.* **9**(4), 532–550 (1987).

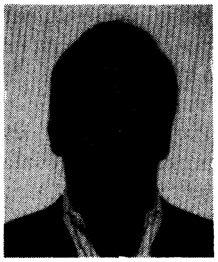

Thomas G. Allen received his SB in 1982 and his SM in 1984, both in aeronautics and astronautics, from the Massachusetts Institute of Technology. From 1982 to 1984 he held a Charles Stark Draper Fellowship. While at Draper, his research involved distributed estimation in surveillance systems. Since 1984 he has been on the technical staff at ALPHATECH, Inc., where he is currently a senior engineer. His research interests include multiobject tracking, parallel processing, image processing, and multiresolution algorithms. Mr. Allen is a member of the IEEE, the ACM, Sigma Xi, and Tau Beta Pi.

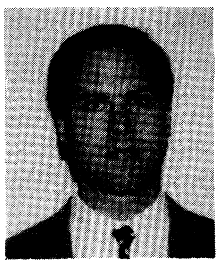

Mark R. Luettgen received the SB, SM, and PhD degrees in electrical engineering from the Massachusetts Institute of Technology in 1988, 1990, and 1993, respectively. From 1988 to 1990 he was with the Loral Corp., where he developed atmospheric optical propagation models and optical source detection and trajectory estimation algorithms. Since 1993 he has been a staff member at ALPHATECH, Inc. His research interests are detection and estimation theory, statistical signal processing, array signal processing optimization, and multiresolution methods. Dr. Luettgen is a member of Eta Kappa Nu and Sigma Xi.

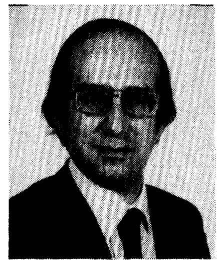

Alan S. Willsky received SB and PhD degrees from the Massachusetts Institute of Technology (MIT) in 1969 and 1973, respectively. He joined the MIT faculty in 1973 and is now professor of electrical engineering. From 1974 to 1981 Dr. Willisky was assistant director of the MIT Laboratory for Information and Decision Systems. He founded and is a member of the board of directors of ALPHATECH, Inc. He was program chairman for the 17th IEEE Conference on Decision and Control, an associate editor of several journals, a member of the Board of Governors and Vice President for Technical Affairs of the IEEE Control Systems Society, program chairman for the 1981 Bilateral Seminar on Control Systems held in China, and special guest editor of the 1992 special issue of the *IEEE Transactions on Information Theory* on wavelet transforms and multiresolution signal analysis. He is a distinguished member of the IEEE Control Systems Society. Dr. Willsky has given several plenary lectures at major scientific meetings, authored the research monograph *Digital Signal Processing and Control and Estimation Theory,* and coauthored the undergraduate text *Signals and Systems.* His awards include the 1975 Donald P. Eckman Award from the American Automatic Control Council, the 1979 Alfred Noble Prize from the ASCE, and the 1980 Browder J. Thompson Memorial Prize Award from the IEEE. Dr. Willsky's research interests are in the development and application of advanced methods of estimation and statistical signal and image processing. Methods he developed have been successfully applied in failure detection in high-performance aircraft, advanced surveillance and tracking systems, electrocardiogram analysis, computerized tomography, and remote sensing.