Machine Vision and Applications (1996) 9:32–42

Machine Vision and Applications

Motion estimation and segmentation

Tina Yu Tian, Mubarak Shah

Computer Vision Laboratory, Computer Science Department, University of Central Florida, Orlando, FL 32816, USA

Abstract. In the general structure-from-motion (SFM) problem involving several moving objects in a scene, the essential first step is to segment moving objects independently. We attempt to deal with the problem of optical flow estimation and motion segmentation over a pair of images. We apply a mean field technique to determine optical flow and motion boundaries and present a deterministic algorithm. Since motion discontinuities represented by line process are embedded in the estimation of the optical flow, our algorithm provides accurate estimates of optical flow especially along motion boundaries and handles occlusion and multiple motions. We show that the proposed algorithm outperforms other well-known algorithms in terms of estimation accuracy and timing.

Key words: Optical flow – Motion boundaries – Line process – Markov random fields – Mean field techniques

1 Introduction

A major goal of vision is to infer physical properties of objects present in a scene, such as their 3D structure and motion, given a sequence of images. This is the so-called structure-from-motion (SFM) problem in computer vision. Much effort has been devoted to the egomotion recovery problem, which considers only a single structure in the scene so that the image-to-scene reconstruction can be studied separately from the problem of segmentation. The segmentation problem arises as soon as more than one moving structure is present in the scene. Thus an essential first step towards solving the SFM problem in the more general case of a scene containing multiple moving objects is the segmentation of the image into regions that are likely to correspond to different objects (Adiv 1985).

The segmentation of independently moving objects in image sequences becomes more difficult when the camera is moving, and the motion parameters of the camera and each moving object are unknown. Although some image properties (e.g., color) can be used to segment the scene, perceptual studies (Braddick 1974; Julesz 1971) have shown that the human visual system can segment a scene into distinct objects based on motion information alone, as shown by our ability to detect camouflaged creatures as soon as they move. Consequently, we focus on motion-based segmentation.

Conventional, gradient-based motion estimation approaches impose a global smoothness constraint on the image, which results in inaccurate optical flow estimates near motion boundaries. Some researchers, such as Konrad and Dubois (1992), Harris and Koch et al. (1990) and Heitz and Bouthemy (1993), introduced line processes and a piecewise smoothness constraint into the regulation formula to improve the motion estimates. The problem is reduced to minimizing a nonconvex cost functional, for which two optimization methods can be used; namely, simulated annealing (Konrad and Dubois 1992), and the iterated condition mode (ICM) (Heitz and Bouthemy 1993). Simulated annealing converges rather slowly to the global minimum, and ICM gets trapped into a local minimum easily unless a good initial estimate is provided. Geiger and Girosi (1991) show that the mean field technique is more appealing for minimizing the nonconvex cost functional, since it converges faster to the global minimum.

In this paper, we deal with the problem of motion estimation and segmentation over a pair of images. We apply mean field techniques to determine the optical flow and motion boundaries, and we present a deterministic algorithm. The theoretical framework is based on a global bayesian decision associated with Markov random field (MRF) models and mean field techniques from statistical mechanics (Geiger and Girosi 1991). Since motion discontinuities are embedded in the flow estimation via line processes, (the sites placed midway between each horizontal or vertical pair of pixels), our algorithm provides accurate estimates of optical flow especially along motion boundaries. Our algorithm also handles occlusion and multiple motions. Without the line process, the pixels near the occluding boundary would mutually affect the optical flow estimation.

The rest of the paper is organized as follows. A brief description of the existing schemes is given in Sect. 2. The

Correspondence to: T.Y. Tian

problem for computing optical flow and detecting motion boundaries is formulated based on MRFs in Sect. 3. In Sect. 4, a deterministic, iterative algorithm is presented for solving the problem. To deal with the problem of large motion, an adaptive multigrid approach is employed in Sect. 5. Experimental results are provided in Sect. 6. Finally, a summary is given in Sect. 7.

2 Related work

The methods for estimating motion and its boundaries can be divided into the following three categories: boundaries can be detected either prior to, simultaneously with, or following the computation of the optical flow field.

2.1 Boundary detection prior to motion computation

Reichardt and Poggio (1980), Hildreth (1984), Spoerri and Ullman (1987) propose methods that detect discontinuities prior to the computation of the flow field. Reichardt and Poggio (1980) use flicker detectors, which are stimulated by the incoherent motion of moving of objects aggainst the background near the edges of the moving objects, to signal motion boundaries. Hildreth (1984) used the normal flow component to detect motion boundaries; he based this on the fact that if two adjacent objects move differently, the normal flow components change in sign and/or magnitude across the boundary. Spoerri and Ullman (1987) extend Hildreth's work and present three methods for segmenting moving objects independently: the bimodality method, the Kolmogorov-Smirnov method, and the dynamic occlusion method. The first two methods were motivated by the fact that the potential displacements or the normal flow components cluster around two points in a local histogram at a motion boundary. The dynamic occlusion method is based on the fact that the appearance and disappearance of thin bars are associated with motion boundaries. After the detection of motion boundaries by these methods, an accurate estimate of optical flow can be obtained. Nevertheless, these methods are sensitive to noise and are not very robust.

2.2 Boundary detection following motion computation

A common approach is to compute optical flow first and then detect the motion boundaries (with, for example an edge detector). This approach requires an accurate and dense optical flow. The two main approaches, intensity-based methods (Horn and Schunck 1981) and feature-based methods (Anandan 1989), have been proposed for determining the optical flow. The intensity-based methods exploit the relationship between the velocity and the spatial and temporal gradients in the image brightness, and assume that the optical flow varies smoothly in the image. This constraint is valid everywhere except at the motion boundaries, where it results in an incorrect optical flow estimate. This in turn causes the edge detector (Thompson et al. 1985) or other techniques (Adiv 1985; Murray and Buxton 1987) to produce some erroneous motion boundaries. Feature-based methods usually produce a sparse optical flow, which needs to be smoothly interpolated to produce a dense flow field. Although the motion boundaries are unknown at this stage, the interpolation scheme causes the motion boundaries to be smoothed out. Consequently, the motion boundaries may either be detected but delocalized, or be smoothed out.

Based on the fact that dynamic occlusion occurs at object boundaries, Mutch and Thompson (1985) propose a correspondence algorithm for motion analysis that detects occlusion by looking for significant regions where matches could not be found in the previous or subsequent frame. In Thompson et al. (1985), 1. Smooth each scalar component of the optical flow field with a symmetric Gaussian kernel 2. Compute the laplacian transform of the smoothed function 3. Subtract the two components to combine into a component-wise laplacian transform of the original flow field 4. Search the combined field for directional reversals, which indicate motion discontinuities. Schunck (1989) uses an iterative procedure. He computes the optical flow field originally by constraint line clustering, and to reduce the noisy boundaries, he interleaves the application of an edge detector to compute the motion boundaries with a smoothing of the computed flow field between the motion boundaries.

Some researchers suggest that the flow field be segmented with 2D or 3D motion models (Adiv 1985; Bouthemy and Rivero 1987; Murray and Buxton 1987, Wang and Adelson 1993). Assuming that the scene consists of piecewise planar surfaces, Adiv (1985) uses the Hough transform and a split-merge procedure to segment the flow field into components, in which each component is consistent with the rigid motion of a roughly planar patch. Based on the same assumption, Murray and Buxton (1985) model the flow field as spatial and temporal MRFs. They solve for the scene segmentation with a maximum a posteriori (MAP) formulation and global optimization by simulated annealing. Bouthemy and Rivero (1987) test for likelihood ratios with 2D motion models and construct regions according to motion homogeneity criteria in a split-and-merge manner. The segmentation begins with a constant motion model and is followed by a more elaborate one (possibly linear). Wang and Adelson (1993) estimate affine parameters within each subregion of the optical flow field, group the motion estimates from patches that cover the same object by k-means clustering, then use hypothesis testing on the motion models to reassign the regions.

2.3 Detecting boundaries and simultaneously computing motion

Since accurate optical flow along motion boundaries cannot be obtained without boundary information, it is preferable to compute the optical flow field and its motion boundaries in parallel as shown by Harris et al. (1990), Heitz and Bouthemy (1993), Konrad and Dubois (1992), Vlontzos and Geiger (1992). The framework for this type of method is based on a global bayesian decision associated with MRF models, incorporating the line process to represent motion discontinuity.

Poggio et al. (1988) propose a method that uses information from various vision modules such as stereo, motion,

texture, and color, to refine the estimation of surface discontinuities for better surface reconstruction. They formulate the problem in the MRF model with a line process representing a surface discontinuity. They used the simulated annealing technique to find the global minimum. Within the motion module, they computed the optical flow and motion boundaries separately with the two methods (i.e., computing motion boundaries either before or after computing optical flow). Thus, Poggio and colleagues' approach faces the same problems as the other methods.

Konrad and Dubois (1992), Harris and Koch et al. (1990), and Heitz and Bouthemy (1993) introduce line processes and a piecewise smoothness constraint into the regulation formulation to improve the motion estimates and compute motion boundaries. The problem is reduced to minimizing a nonconvex cost function. Konrad and Dubois (1992) use simulated annealing, at a huge computational expense, particularly when several constraints are incorporated. Heitz and Bouthemy (1993) use the ICM technique, which gets trapped into a local minimum easily unless a good initial estimate is provided. Harris et al. (1990) implement the line process with analog processing elements and provide a method to compute optical flow and motion boundaries with a very large scale integration (VLSI) technique. Unlike their work, our goal is to develop a deterministic numerical method for motion estimation.

Vlontzos and Geiger (1992) propose an algorithm for optical flow estimation that uses a mean-field technique. In their method, the sparse optical flow field is computed by a window matching scheme, then the dense flow field is computed by interpolating and smoothing the sparse flow field except at the motion boundaries. In our method, the problem is formulated directly in terms of intensity, avoiding the intermediate step of computing the flow by correlation. Zhang and Hanauer (1993) apply the mean-field theory to the *displacement* field estimation. They formulate the problem in terms of intensity values at two consecutive frames and three coupled MRF fields, namely the displacement field, the line field, and the segmentation field. However, including both the line field and the segmentation field into the formulation introduces redundant computations, since motion boundaries obtained from the line field also yield segmentation results. We introduce only two coupled MRF fields, the optical flow field and the line field, to formulate the problem in terms of intensity gradients. Thus our algorithm is simpler and computationally more efficient. Abdelqader and Rajala (1993) also propose a mean-field annealing approach to motion estimation in which only the displacement field is involved in the MRF formulation. Our coupled MRF model includes the flow field and the horizontal and vertical line fields. Thus motion estimation and segmentation are solved simultaneously.

Other methods have been proposed for motion detection and segmentation. Irani and Peleg (1992) do not compute the optical flow or the displacement field. Instead, they approximate the motion of objects by 2D parametric models, and compute the motion parameters of the dominant object, which is then eliminated from subsequent analysis. They repeat the process for the remaining regions to detect other objects and their motions. Black and Anandan (1991) present a method to estimate motion velocity incrementally and segment the independently moving objects over multiple image frames. They exploit the MRF formulation and the incremental stochastic minimization algorithm to overcome the disadvantage of simulated annealing. Jepson and Black (1993) use a probabilistic mixture model to represent multiple motions within an image patch and employ the expectation & maximization (EM) algorithm to compute a maximum likelihood estimate for the various motion parameters.

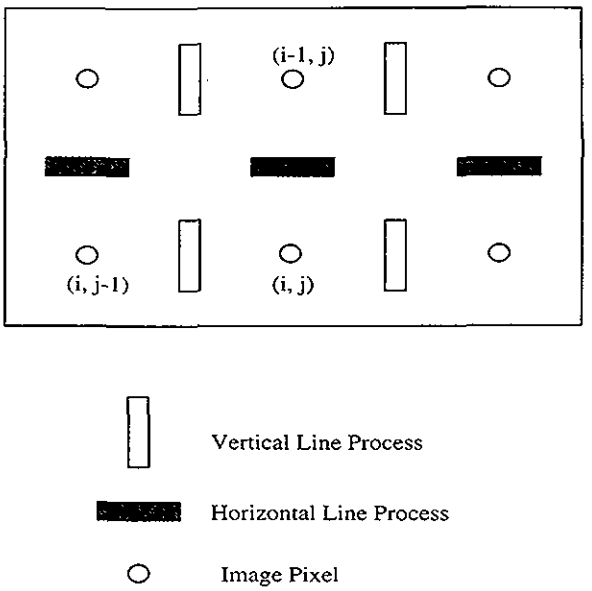

Fig. 1. The line fields in a dual lattice. ▯ Vertical line field, ▬ horizontal line field, ○ image pixel

3 Determining optical flow and motion boundaries with MRFs

3.1 Problem formulation based on MRFs

For rigid motion, since the optical flow field $\mathbf{U}$, where $\mathbf{U} = (u, v)^T$, consists of patches of vectors of similar orientation and length with potential discontinuities at motion boundaries, the probability distribution for $\mathbf{U}$ at site i depends only upon the neighborhood N_i. Hence, the flow field $\mathbf{U}$ can be modeled as a 2D vector MRF. Moreover, another field, the line process, located in the dual lattice, is used to represent the possible location of motion boundaries. Line process fields are defined as the sites placed midway between each horizontal or vertical pair of pixels. The dual lattice is another lattice coupled with the flow field lattice such that for each site of the flow field lattice there are two sites in the dual lattice, one corresponding to the horizontal line and the other corresponding to the vertical line, as shown in Fig. 1. The horizontal (vertical) line process represents motion discontinuities between pixels along the vertical (horizontal) direction. The horizontal line process, $d_{h_{ij}}$, connects pixel (i, j) to pixel $(i-1, j)$, while the vertical line process, $d_{v_{ij}}$, connects pixel (i,j) to pixel $(i, j-1)$ in the lattice.

The global bayesian estimation associated with MRF models provides a mathematical framework to recover motion and detect motion boundaries from the observed image sequence. The problem we are considering is to estimate the

2D optical flow field U and line process fields, d_h and d_v, from the image gradient ∇I, where $\nabla I = (I_x, I_y, I_t)^t$.

There are two possible types of errors in the estimation process. One type of error is in the derivative computations due to camera and quantization noise, aliasing, imprecision in the derivative filters, etc. If we describe this type of uncertainty as white gaussian noise n_1 with a standard deviation σ_1, then we have:

$$P(\nabla I \mid u, v) = C_1^{-1} \exp\left[\frac{1}{2\sigma_1^2} \sum_{i,j} \left(u_{i,j} I_{x_{i,j}} + v_{i,j} I_{y_{i,j}} + I_{t_{i,j}}\right)^2\right], \quad (1)$$

where u_{ij} and v_{ij} are the optical flow field defined at site (i, j) of the lattice and C_1 is a normalization constant. The term in the exponential in Eq. 1 is based on the brightness constancy equation of Horn and Schunck (1981).

Another type of error is caused by changes in lighting or reflectance, or the presence of multiple motions, in which case the brightness constancy constraint is unsatisfied. This type of error can similarly be modeled independently of n_1 as white gaussian noise n_2 with a standard deviation σ_2. Within the bayesian decision approach, a prior probability, $P(u, v, d_h, d_v)$, can be defined to impose a piecewise smoothness constraint on the optical flow field:

$$P(u, v, d_h, d_v) = C_2^{-1} \exp\left\{ -\sum_{i,j} \frac{1}{2\sigma_2^2}\left[\|\mathbf{U}_{i,j} - \mathbf{U}_{i-1,j}\|^2 \left(1 - d_{h_{i,j}}\right) + \|\mathbf{U}_{i,j} - \mathbf{U}_{i,j-1}\|^2 \left(1 - d_{v_{i,j}}\right)\right] + \gamma^h d_{h_{i,j}} + \gamma^v d_{v_{i,j}} \right\}, \quad (2)$$

where C_2 is the normalization constant, and:

$$\|\mathbf{U}_{i,j} - \mathbf{U}_{i-1,j}\|^2 = \left(u_{i,j} - u_{i-1,j}\right)^2 + \left(v_{i,j} - v_{i-1,j}\right)^2,$$
$$\|\mathbf{U}_{i,j} - \mathbf{U}_{i,j-1}\|^2 = \left(u_{i,j} - u_{i,j-1}\right)^2 + \left(v_{i,j} - v_{i,j-1}\right)^2.$$

The first two terms in the exponential part of Eq. 2 contain the sum of local interactions of the flow field and line process fields between the nearest neighbors in terms of the smoothness constraint to solve the well known aperture problem. If the gradient of the u flow or v flow in the horizontal or vertical direction is very high at site (i, j), the corresponding line process is likely to signal a discontinuity ($d_{v_{ij}} = 1$ or $d_{h_{ij}} = 1$) to increase $P(u, v, d_h, d_v)$. The third term in the exponential part of Eq. 2 takes into account the penalty paid each time a discontinuity is created, and is needed to prevent the creation of discontinuities everywhere.

By Bayes' theorem,

$$P(u, v, d_h, d_v \mid \nabla I) = \frac{P(\nabla I \mid u, v, d_h, d_v) P(u, v, d_h, d_v)}{P(\nabla I)},$$

where $P(\nabla I)$ is the probability of the image gradients. We assume that it is a constant, without a priori knowledge of the distribution of the image gradients. Since the error in the derivative computation is independent of the line process d_h and d_v, which are taken care of in the probability $P(u, v, d_h, d_v)$, the following relation holds: $P(\nabla I \mid u, v, d_h, d_v) = P(\nabla I \mid u, v)$, where $P(\nabla I \mid u, v)$ is given in Eq. 1. Now, we have a posterior probability for the variables u, v, d_h, and d_v, given the image gradient ∇I:

$$P(u, v, d_h, d_v \mid \nabla I) = C^{-1} \exp\left\{ -\sum_{i,j} \frac{1}{2\sigma_1^2} \left(u_{i,j} I_{x_{i,j}} + v_{i,j} I_{y_{i,j}} + I_{t_{i,j}}\right)^2 + \frac{1}{2\sigma_2^2}\left[\|\mathbf{U}_{i,j} - \mathbf{U}_{i-1,j}\|^2 \left(1 - d_{h_{i,j}}\right) + \|\mathbf{U}_{i,j} - \mathbf{U}_{i,j-1}\|^2 \left(1 - d_{v_{i,j}}\right)\right] + \gamma^h d_{h_{i,j}} + \gamma^v d_{v_{i,j}} \right\}. \quad (3)$$

The posterior probability can be written in the following form:

$$P\left(u, v, d_h, d_v \mid \nabla I\right) = Z^{-1} \exp\left[-V\left(u, v, d_h, d_v\right)\right], \quad (4)$$

where Z is the normalization constant, called a partition function:

$$Z = \sum_{\{u,v,d_h,d_v\}} \exp\left[-V\left(u, v, d_h, d_v\right)\right],$$

where $\sum_{\{u,v,d_h,d_v\}}$ is the sum over all the possible configurations $\{u, v, d_h, d_v\}$ of the system. V is a cost function of the following form, which is similar to the weak membrane model:

$$V\left(u, v, d_h, d_v\right) = \sum_{i,j} \left\{ \frac{1}{2\sigma_1^2} \left(u_{i,j} I_{x_{i,j}} + v_{i,j} I_{y_{i,j}} + I_{t_{i,j}}\right)^2 + \frac{1}{2\sigma_2^2}\left[\|\mathbf{U}_{i,j} - \mathbf{U}_{i-1,j}\|^2 \left(1 - d_{h_{i,j}}\right) + \|\mathbf{U}_{i,j} - \mathbf{U}_{i,j-1}\|^2 \left(1 - d_{v_{i,j}}\right)\right] + \gamma^h d_{h_{i,j}} + \gamma^v d_{v_{i,j}} \right\}. \quad (5)$$

The cost function consists of the brightness-constancy constraint, the piecewise smoothness constraint, and the penalty term for creating a motion discontinuity. Finding the MAP estimate of fields u, v, d_h, and d_v is reduced to the problem of minimizing the energy function V.

3.2 Energy minimization by mean field techniques

In this section, we address the minimization problem of the energy function associated with line processes. The global minimum of the nonconvex cost function V can be obtained by using simulated annealing, or a deterministic technique called the graduated nonconvexity (GNC) algorithm. In solving the problems of image segmentation and surface reconstruction, Geiger and Girosi (1991) use mean field techniques from statistical mechanics to find the deterministic equations with solutions that approximate the MAP estimator. They show that the mean field essentially represents the minimum variance bayesian estimator, and therefore can be used as a measure of field value. We propose to use mean field techniques to obtain an estimate of the global minimum of our energy function.

A well-known result (Geiger and Girosi 1991) from statistical mechanics and probability theory indicates that all mean values of the system can be obtained from the partition function Z, so we must compute the function Z. By introducing β, we have:

$$Z = \sum_{\{u,v,d_h,d_v\}} \exp[-\beta V(u,v,d_h,d_v)] .$$

The introduction of β allows us to apply continuation methods similar to simulated annealing. From Eq. 5, the partition function becomes:

$$Z = \sum_{\{u,v\}} \exp\left\{ -\beta \sum_{i,j} \left[\frac{1}{2\sigma^2} \left(u_{i,j} I_{x_{i,j}} + v_{i,j} I_{y_{i,j}} + I_{t_{i,j}}\right)^2 + \gamma^h + \gamma^v \right] \right\} \sum_{\{d_h,d_v\}} \exp\left\{ -\beta \sum_{i,j} \left[\left(1 - d_{h_{i,j}}\right) G^h_{i,j} + \left(1 - d_{v_{i,j}}\right) G^v_{i,j} \right] \right\}, \quad (6)$$

where $G^h_{ij} = \frac{1}{2\sigma_2^2} \|\mathbf{U}_{i,j} - \mathbf{U}_{i-1,j}\|^2 - \gamma^h$, and

$G^v_{ij} = \frac{1}{2\sigma_2^2} \|\mathbf{U}_{i,j} - \mathbf{U}_{i,j-1}\|^2 - \gamma^v$.

The line process term in Eq. 6 is the partition function of two independent spin systems, d_h and d_v, in an external field, G^h or G^v, with no interaction between neighboring sites. Each spin contributes to the partition function independently. Therefore, the line process can be successfully averaged out to yield an effective cost function depending only upon the fields u and v. The partition function can then be rewritten as:

$$Z = \sum_{\{u,v\}} \exp\left[-\beta V_{\text{eff}}(u,v)\right] ,$$

where $V_{\text{eff}}(u,v) = \sum_{i,j} \left\{ \frac{1}{2\sigma_1^2} \left(u_{i,j} I_{x_{i,j}} + v_{i,j} I_{y_{i,j}} + I_{t_{i,j}}\right)^2 - \frac{1}{\beta} \ln\left[\left(1 + e^{-\beta G^h_{i,j}}\right) \left(1 + e^{-\beta G^v_{i,j}}\right) \right] + \gamma^h + \gamma^v \right\}$.

Using the mean field technique (Geiger and Girosi 1991), we can derive the following mean field equations for the line process fields:

$$d_{h_{ij}} = \left\{1 + \exp\left[\beta\left(\gamma^h - \|\mathbf{U}_{i,j} - \mathbf{U}_{i-1,j}\|^2/2\sigma_2^2\right)\right]\right\}^{-1} , \quad (7)$$

$$d_{v_{ij}} = \left\{1 + \exp\left[\beta\left(\gamma^v - \|\mathbf{U}_{i,j} - \mathbf{U}_{i,j-1}\|^2/2\sigma_2^2\right)\right]\right\}^{-1} . \quad (8)$$

For simplicity we assume $\gamma^h = \gamma^v = \gamma$, where γ is one of the predefined parameters. Notice that the mean values of the line process fields can vary continuously from 0 to 1.

We pointed out earlier that the estimate of u and v obtained by minimizing the energy function V coincides with the maximum a posteriori estimate of Eq. 3 after summing out the line fields. This corresponds to minimizing V_{eff} with respect to u and v, and yields the following differential equations:

$$\partial V_{\text{eff}}(u,v)/\partial u_{ij} = 0 , \quad (9)$$

$$\partial V_{\text{eff}}(u,v)/\partial v_{ij} = 0 . \quad (10)$$

From Eqs. 9 and 10, we obtain the following equations:

$$\left(u_{i,j} I_{x_{i,j}} + v_{i,j} I_{y_{i,j}} + I_{t_{i,j}}\right) I_{x_{i,j}} + \frac{\sigma_1^2}{\sigma_2^2} \Big\{ \left(1 - d_{h_{i,j}}\right)\left(u_{i,j} - u_{i-1,j}\right) - \left(1 - d_{h_{i+1,j}}\right)\left(u_{i+1,j} - u_{i,j}\right) + \left(1 - d_{v_{i,j}}\right)\left(u_{i,j} - u_{i,j-1}\right) - \left(1 - d_{v_{i,j+1}}\right)\left(u_{i,j+1} - u_{i,j}\right) \Big\} = 0 , \quad (11)$$

$$\left(u_{i,j} I_{x_{i,j}} + v_{i,j} I_{y_{i,j}} + I_{t_{i,j}}\right) I_{y_{i,j}} + \frac{\sigma_1^2}{\sigma_2^2} \Big\{ \left(1 - d_{h_{i,j}}\right)\left(v_{i,j} - v_{i-1,j}\right) - \left(1 - d_{h_{i+1,j}}\right)\left(v_{i+1,j} - v_{i,j}\right) + \left(1 - d_{v_{i,j}}\right)\left(v_{i,j} - v_{i,j-1}\right) - \left(1 - d_{v_{i,j+1}}\right)\left(v_{i,j+1} - v_{i,j}\right) \Big\} = 0 . \quad (12)$$

We solve the simultaneous Eqs. 11 and 12 to find the mean field solution of $u_{i,j}$ and $v_{i,j}$ in terms of $d_{h_{i,j}}$ and $d_{v_{i,j}}$.

First we define the local average of discontinuity at pixel (i,j) as:

$$\bar{d}_{i,j} = \frac{1}{4}\left[4 - d_{h_{i,j}} - d_{h_{i+1,j}} - d_{v_{i,j}} - d_{v_{i,j+1}}\right] , \quad (13)$$

and the local averages of u and v at pixel (i,j) as:

$$\bar{u}_{i,j} = \frac{1}{4}\left[\left(1 - d_{h_{i,j}}\right) u_{i-1,j} + \left(1 - d_{v_{i,j}}\right) u_{i,j-1} + \left(1 - d_{h_{i+1,j}}\right) u_{i+1,j} + \left(1 - d_{v_{i,j+1}}\right) u_{i,j+1}\right] , \quad (14)$$

$$\bar{v}_{i,j} = \frac{1}{4}\left[\left(1 - d_{h_{i,j}}\right) v_{i-1,j} + \left(1 - d_{v_{i,j}}\right) v_{i,j-1} + \left(1 - d_{h_{i+1,j}}\right) v_{i+1,j} + \left(1 - d_{v_{i,j+1}}\right) v_{i,j+1}\right] . \quad (15)$$

We also assume $\lambda = 4\sigma_1^2/\sigma_2^2$. Then Eqs. 11 and 12 become:

$$\left(u_{ij} I_{x_{ij}} + v_{ij} I_{y_{ij}} + I_{t_{ij}}\right) I_{x_{ij}} + \lambda\left(\bar{d}_{ij} u_{ij} - \bar{u}_{ij}\right) = 0 , \quad (16)$$

$$\left(u_{ij} I_{x_{ij}} + v_{ij} I_{y_{ij}} + I_{t_{ij}}\right) I_{y_{ij}} + \lambda\left(\bar{d}_{ij} v_{ij} - \bar{v}_{ij}\right) = 0 . \quad (17)$$

Rewriting these above equations, we get the following more compact equations:

$$\left(\lambda\bar{d}_{ij} + I^2_{x_{ij}}\right) u_{ij} + I_{x_{ij}} I_{y_{ij}} v_{ij} = \lambda\bar{u}_{ij} - I_{x_{ij}} I_{t_{ij}} . \quad (18)$$

$$I_{x_{ij}} I_{y_{ij}} u_{ij} + \left(\lambda\bar{d}_{ij} + I^2_{y_{ij}}\right) v_{ij} = \lambda\bar{v}_{ij} - I_{y_{ij}} I_{t_{ij}} . \quad (19)$$

Solving the simultaneous Eqs. 18 and 19, we get:

$$u_{ij} = \frac{(\lambda\bar{d}_{ij} + I^2_{y_{ij}})\bar{u}_{ij} - I_{x_{ij}} I_{y_{ij}} \bar{v}_{ij} - \bar{d}_{ij} I_{x_{ij}} I_{t_{ij}}}{\bar{d}_{ij}(\lambda\bar{d}_{ij} + I^2_{x_{ij}} + I^2_{y_{ij}})} ;$$

$$v_{ij} = \frac{-I_{x_{ij}} I_{y_{ij}} \bar{u}_{ij} + (\lambda\bar{d}_{ij} + I^2_{x_{ij}})\bar{v}_{ij} - \bar{d}_{ij} I_{y_{ij}} I_{t_{ij}}}{\bar{d}_{ij}(\lambda\bar{d}_{ij} + I^2_{x_{ij}} + I^2_{y_{ij}})} .$$

Finally, the iterative formulae are obtained as follows:

$$u^{(n+1)}_{ij} = \left[\bar{u}^{(n)}_{ij} - a^{(n)}_{ij} I_{x_{ij}}\right] / \bar{d}^{(n)}_{ij} , \quad (20)$$

$$v^{(n+1)}_{ij} = \left[\bar{v}^{(n)}_{ij} - a^{(n)}_{ij} I_{y_{ij}}\right] / \bar{d}^{(n)}_{ij} , \quad (21)$$

where

$$a^{(n)}_{ij} = \left[I_{x_{ij}} \bar{u}^{(n)}_{ij} + I_{y_{ij}} \bar{v}^{(n)}_{ij} + \bar{d}^{(n)}_{ij} I_{t_{ij}}\right] / \left[\lambda\bar{d}^{(n)}_{ij} + I^2_{x_{ij}} + I^2_{y_{ij}}\right]$$

and the superscript denotes the iteration number.

During an updating cycle, the new approximation of the optical flow, $u^{(n+1)}$ and $v^{(n+1)}$, can be determined from the estimated brightness derivatives, I_x, I_y, and I_t, and from the local average of the previous flow, $\bar{u}^{(n)}$ and $\bar{v}^{(n)}$, and line process estimates, $\bar{d}^{(n)}$, respectively. The line process can be updated from the previously estimated mean values of $u^{(n)}$ and $v^{(n)}$. Note that the forms of Eqs. 7, 8, 20, and 21 are very suitable for a fast, parallel, and iterative scheme for a solution.

It is interesting to note that Horn and Schunck's (1981) method becomes a special case of our scheme when the line process fields are set to zero. In their method, a smoothness constraint is imposed over the whole image, so considerable error occurs in the vicinity of the object boundaries. In our scheme, however, the line process prevents the smoothness constraint from being imposed across motion boundaries. Assume that the scene contains two moving objects, say object 1 and Object 2. When the local averages of the optical flow at a pixel belonging to object 1 are computed according to Eqs. 14 and 15, its neighbors belonging to object 2 will be inhibited by the line process from participating in the computation if the current pixel and/or its neighbors belong to motion boundaries. Without the line process, the pixels near the occluding boundary would mutually affect the estimation of the corresponding optical flows.

4 The iterative algorithm

4.1 Algorithm

The iterative algorithm for computing optical flow and motion boundaries that uses Eqs. 7, 8, 20 and 21, derived in Sect. 3, is given in algorithm 1. In this algorithm, we use a three-point approximation (Battiti et al. 1991) to calculate the spatial and temporal brightness derivatives I_x, I_y, and I_t, since a three-point approximation of derivatives provides a better estimate than the two-point forward difference formula ($O(\Delta x)^2$ error versus $O(\Delta x)$ error). We also apply the additional constraint that motion discontinuities usually coincide with intensity edges. If no intensity edge is present, we increase the threshold ξ by increasing γ and/or σ_2 by a factor of ten.

Algorithm 1

1. Calculate the spatial and temporal brightness derivatives I_x, I_y, and I_t.
2. Apply Canny's edge detector and initialize the line processes, d_h and d_v, by edges.
3. Compute the local average value of discontinuity ($\bar{d}_{i,j}$), and optical flow ($\bar{u}_{i,j}$,$\bar{v}_{i,j}$) by Eqs. 13–15.
4. Calculate the new approximation of optical flow, (u, v), by with Eqs. 20 and 21.
5. If the pixel is not at an intensity discontinuity, increase threshold ξ by a factor of ten. Otherwise, leave $\xi = \sigma_2\sqrt{2\gamma}$ unchanged. Next, compute the new approximations of d_h and d_v with Eqs. 7 and 8.
6. Calculate an error norm for the last iteration step. (The average difference value between the new and old approximations of optical flow over all the image points has been used.)
7. Test the termination condition. If the error norm is less than a set threshold, then the algorithm is terminated.
8. Update the optical flow and line fields with the new approximation values, and then go to step 3.

end

4.2 Parameters

There are several parameters, namely σ_1, σ_2, γ, and β, involved in the algorithm. The parameter σ_1 is the standard deviation of the error in the intensity gradient, and σ_2 is the standard deviation of the error when the brightness constancy constraint is not satisfied, as in the case of multiple motions or changes of reflectance or lighting. The new parameter introduced in our formula, $\lambda = 4\sigma_1^2/\sigma_2^2$, plays a role similar to that of the regularization parameter in Horn and Schunck's approach. The parameter λ controls the balance between the "trust" in the data and the smoothness term. However, it is more meaningful to define λ as the ratio of variance of the two types of errors, namely the gradient error and the error due to brightness constancy constraint. If σ_1 is relatively large, we cannot get accurate intensity derivatives, and λ becomes a larger value that increases the weight of the smoothness term; if σ_2 is large, the brightness constancy constraint does not hold, λ becomes a smaller value which decreases the weight of the smoothness term and limits its smoothing to some extent. The parameters σ_1 and σ_2 can be estimated for each given image sequence with the mathematical methods described by Geiger and Girosi (1991). In our experiments, however, no attempt was made to estimate the actual value of σ_1 or σ_2. From Eqs. 20 and 21, one can see that λ (the ratio of σ_1 to σ_2) affects the optical flow estimation. Thus we always set $\mu = 1/2\sigma_2^2$ to 1, and visually adjust λ for each given image sequence.

For the parameter γ, from Eqs. 7 and 8, one can see that $\sqrt{\gamma/\mu}$ is the threshold for creating a motion discontinuity, where $\mu = 1/2\sigma_2^2$. Let us denote $\sqrt{\gamma/\mu}$ by ξ. We notice that if gradient of U is above ξ, $d_{h_{ij}}$ (or $d_{v_{ij}}$) is greater than 0.5, thus there is less smoothing, whereas if the gradient is below ξ, more smoothing is applied. The parameter ξ affects the quality and the amount of the motion boundary estimates. To find a good pattern of motion boundaries, ξ must be between 1 and 3 in our experiments; the exact value of ξ depends on the magnitude of the relative motion of objects.

The parameter β is introduced as an extra degree of freedom; continuation methods (e.g., simulated annealing) can be applied by varying β. Since the temperature of the system is defined as $T = 1/\beta$, varying β in the minimization process is equivalent to varying T. Sometimes, a finite β solution may be more desirable or robust. Our experiments show that varying β does not make much difference; thus β must always be set to 5.

5 The adaptive multigrid approach

The basic assumption made in solving Eqs. 11 and 12 with the discretized versions Eqs. 20 and 21 in the previous sec-

tion is that the spatial and temporal sampling steps are small compared to the motion amplitudes. If the object moves very fast, the brightness changes rapidly, and the gradients I_x, I_y, and I_t fail to be accurately estimated by a two-point or a three-point discrete approximation, since in this case, the step is not infinitesimal.

The multiscale pyramid scheme can be used to deal with this problem. Since the high spatial frequencies are attenuated at the lower resolutions, the spatial and temporal gradients are smoother and more accurate, unless the discretization error dominates. Analyzed by Battiti et al. (1991) the relative error in the flow field consists of derivative estimation error and quantization error, given by:

$$E = \frac{C}{\sigma^2}\left[(\Delta_t I)^2 - (\Delta_x I)^2 - (\Delta_y I)^2\right] + \left\{1/\left[(\Delta_x I)^2 + (\Delta_y I)^2\right] + 1/(\Delta_t I)^2\right\}^{1/2} \quad (22)$$

where $\Delta_x I = 2I_x \Delta x \quad \Delta_y I = 2I_y \Delta y \quad \Delta_t I = 2I_t \Delta t$,

and Δx, Δy, and Δt are the spatial and temporal sampling steps. I_x, I_y, and I_t are obtained by three-point approximation of the first derivatives. C is suggested to be $2\pi^2/3$ by Battiti et al. (1991), and σ is the standard derivation in the distribution of intensity values. The first term of Eq. 22 corresponds to the derivative estimation error, while the second term corresponds to the quantization error.

The quantization error is the largest at the coarsest scale while the derivative estimation error is the largest at the finest scale, and the total error reaches the minimum on some middle scale. Based on this observation, Battiti et al. (1991) propose an adaptive multigrid algorithm to compute the flow field. The algorithm starts by estimating the flow field at a reasonably coarse scale. This approximation is then updated on successive finer scales only on some regions of the image where the estimated flow error is greater than a predefined threshold, T_{err}. In this coarse-to-fine strategy, regions of the image involving various motion amplitudes are processed at the appropriate scale, avoiding corruption of good estimates due to poor derivative estimation at a finer scale.

In Battiti's approach, however, since motion boundaries are not taken into account, the flow field is smoothed continuously across the motion boundaries in each scale. To deal with this problem, our approach is embedded in an adaptive multigrid algorithm:

Algorithm 2

1. Build the Gaussian Pyramid (Burt 1984) associated with three successive images. The number of levels depends on the maximal displacement of pixels. An inhibition flag associated with each point of the image is initialized as 0 (noninhibited).
2. Apply Algorithm 1 to the noninhibited pixels at the current scale.
3. Calculate the relative error of flow, E, with Eq. 22 for each pixel. Check whether E is below the selected threshold. If it is, the estimate of the optical flow can be considered accurate enough, then the inhibition flags of the imme- diate four children of this pixel are set to 1 (inhibited).
4. Apply bilinear interpolation to interpolate the optical flow values to the next finer scale, where they are used as initial estimates of optical flow. Then go to Step 2.

end

6 Results

The real image sequence shown in Fig. 2a,b used for this test contains two toys, Dale and Scrooge (Disney World characters), moving in the opposite directions. The measured displacement is about 1 pixel/frame for each toy. The edge map, shown in Fig. 2c, obtained by a Canny edge detector was used to initialize the line processes. Figure 2d,e shows the optical flow and motion boundaries obtained from algorithm 1, respectively. As we can see after applying our algorithm, spurious edges that do not correspond to motion discontinuities have been smoothed out. and only the true motion boundaries remain.

Our algorithm was also tested on a relatively complicated scene, the Hamburg taxi sequence Fig. 3a. This sequence contains several moving objects; a taxi is turning the corner, a car in the lower left is moving from left to right, and a van in the lower right is moving right to left. The image velocities of the moving objects are approximately 1, 3, and 3 pixels/frame. We only used the lower-left portion of the image to study the approach described in this paper. We applied the adaptive multigrid algorithm to this sequence due to large motions. A three-level gaussian pyramid is used for this sequence. Since the motion of the taxi is small, the computation for the taxi portion continues from the coarsest level to the finest level because the estimated error of the optical flow is always above a predefined threshold, and we must keep on refining the optical flow to reduce the error. However, for the car portion, computation terminates at the intermediate level since the estimated error of optical flow due to derivative estimation error and/or quantization error is below the threshold at the intermediate level. If the computation were continued further for the car portion, the optical flow would be corrupted at the finer level due to poor derivative estimation at this level. The gaussian pyramid for frame 1, initial values for the line process, and the resulting optical flow at three levels are shown in Fig. 4. It is clear that the optical flow has been refined through multiscales. The optical flow and motion boundaries are displayed in Fig. 3c,d. We observe that there is no oversmoothing due to the introduction of the line process, and the adaptive multigrid algorithm does give better and faster results even for image sequences with different ranges of optical flow velocity.

6.1 Comparison

Barron et al. (1994) present a comprehensive evaluation and comparison of existing optical flow methods. Here we compare our method with Anandan's (1989) method, Horn and Schunck's (1981) method, and Fleet and Jepson's (1990) method on the same inputs, using the performance measures of Barron et al. The random dot synthetic square image sequence shown in Fig. 5 is used for quantitative performance comparisons. The optical flow error is measured in terms

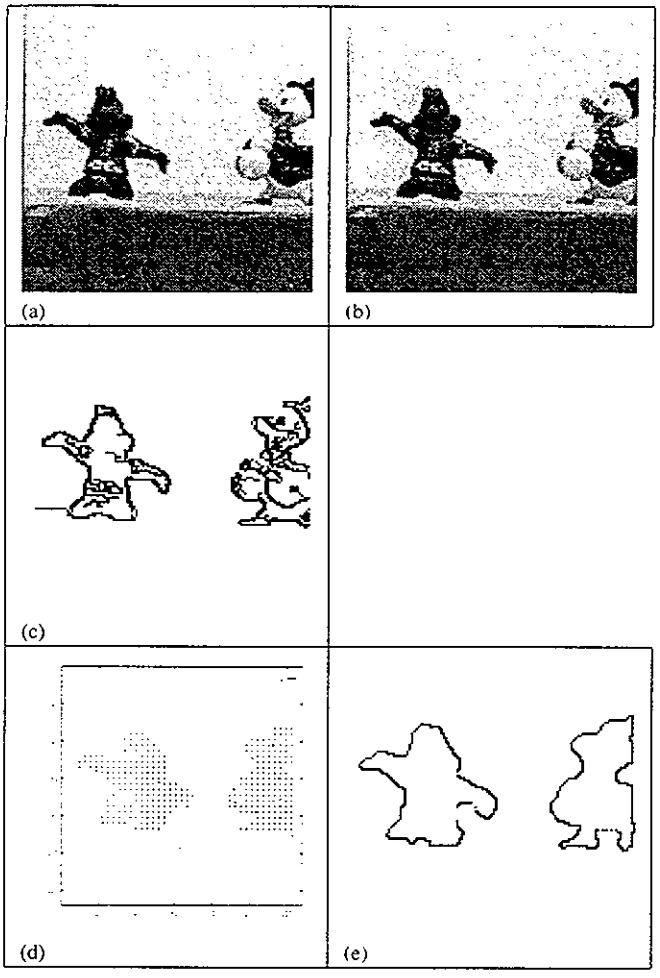

2

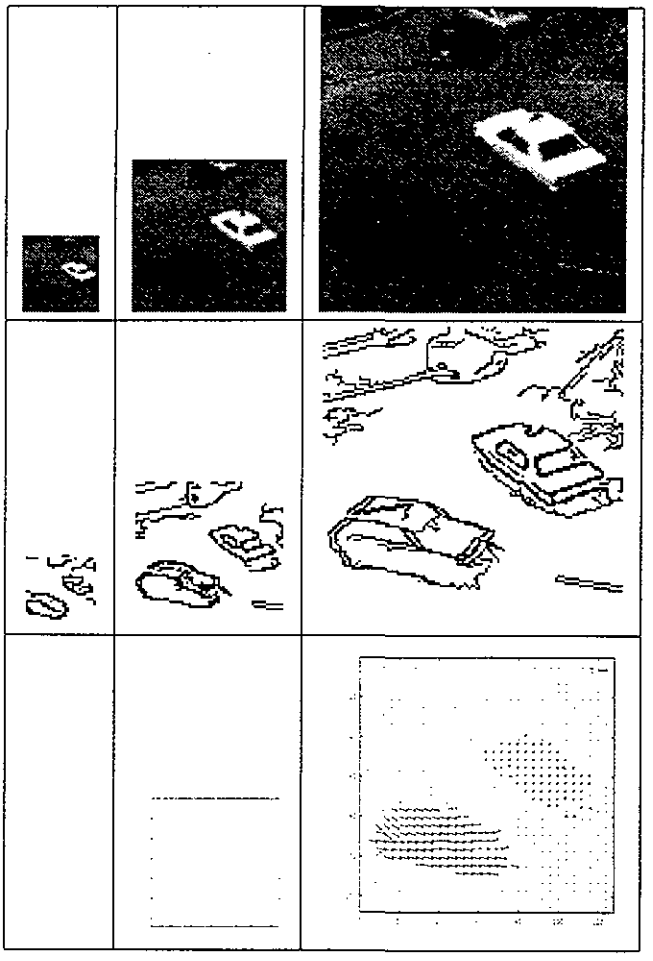

4

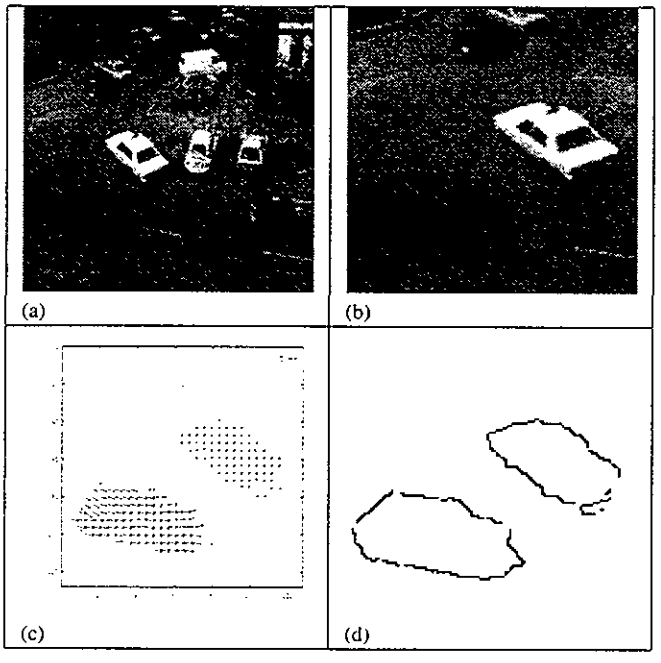

3

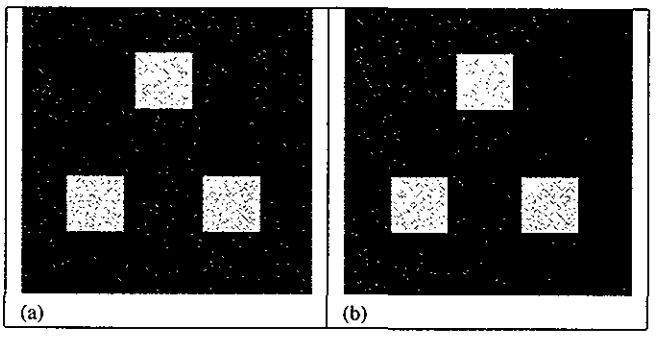

5

Fig. 2a–e. Results for toy images with $\lambda = 25$, $\gamma = 5$, $\mu = 1/2\sigma_2^2 = 1$, and $\beta = 5$. **a** frame 1; **b** frame 2; **c** initial value for the line processes, obtained by the Canny edge detector; **d** optical flow; **e** motion boundaries

Fig. 3a–d. Results for Hamburg Taxi with $\lambda = 100$, $\gamma = 7$, $\mu = 1/2\sigma_2^2 = 1$, $\beta = 5$, $T_{\text{err}} = 0.4$, and the sampling rate = 4. **a** frame 1; **b** the lower left portion of frame 1; **c** optical flow; **d** motion boundaries

Fig. 4. Top: Gaussian pyramid for frame 1; middle: initial values for line process at each scale, obtained by Canny edge detector. Bottom: optical flow at each scale

Fig. 5a,b. A random dot, synthetic square image with displacement of 1 pixel/frame

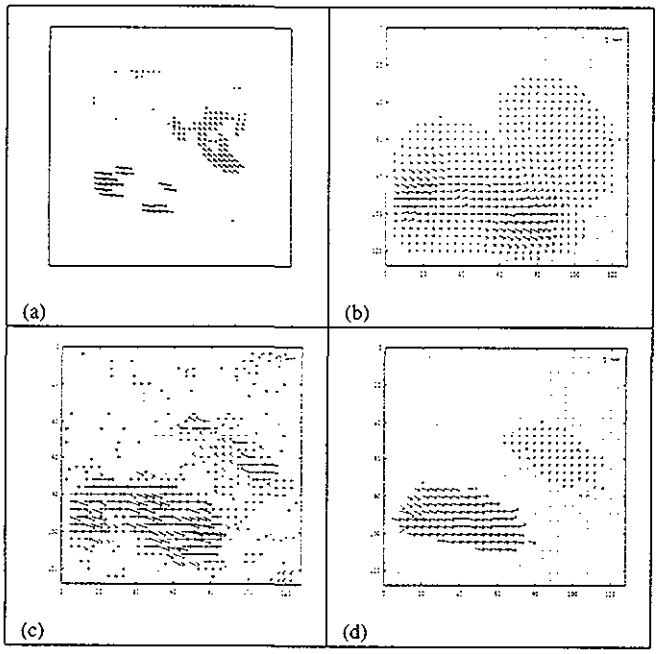

6

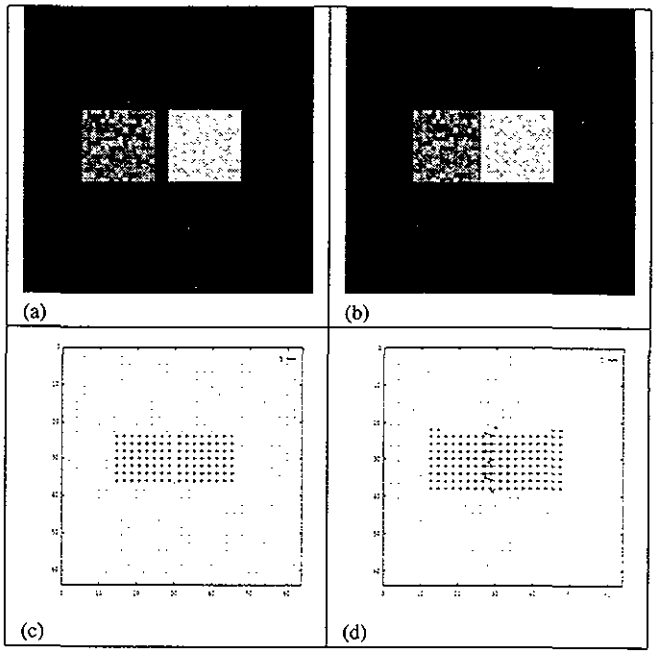

7

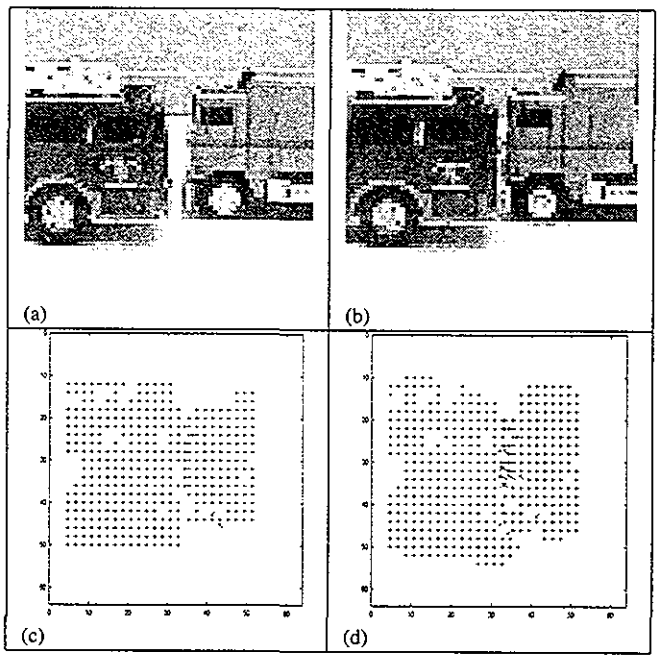

8

Fig. 6a–d. Comparison of the optical flow for Hamburg taxi sequence: **a** Fleet and Jepson's results; **b** Horn and Schunck's results; **c** Anandan's results; **d** our results

Fig. 7a–d. Occlusion and multiple motions – results for the synthetic sequence with two squares moving in opposite directions with displacement 1 pixel/frame: **a** frame 1; **b** frame 3; **c** the optical flow computed by our algorithm; **d** the optical flow computed by Horn and Schunck's algorithm

Fig. 8a–d. Occlusion and multiple motions – results for a sequence of toy vehicles. **a** frame 1; **b** frame 3; **c** the optical flow computed by our algorithm; **d** the optical flow computed by Horn and Schunck's algorithm

of the angle between the correct velocity and the corresponding estimated velocity. Table 1 shows the mean error and the standard deviation of the error for the flow computed by these methods, and the density of the computed flow field as well. It is surprising to see that Fleet and Jepson's method does not work well for the random dot images that involve texture and high-frequency data. However, their method works well with uniform, nontextured image data. In these experiments we have used Fleet and Jepson's program obtained by anonymous ftp (`ftp@csd.uwo.ca`). To be sure that we have used Fleet and Jepson's program properly, we also experimented with the translating, uniform square image sequence, which is the same as the sequence shown in Fig. 5 except that the background and objects have constant gray levels, 0 and 255, respectively. We used 21 frames of the square images as the input. Their program provided quite accurate results with an average error and standard deviation of error equal to 0.034° and 0.028°, respectively, which are very close to (0.07°, 0.08°), the error measurements of Fleet and Jepson's method for their uniform square sequence, as reported by Barron et al. (1994). The same program was then applied to the random dot sequence, which provides the sparse optical flow field with average error and standard deviation of error, and density of the optical flow equal to 7.38°, 11.79°, and 13.2%, respectively.

Figure 6 illustrates the optical flow field obtained by multigrid Horn and Schunck's method, Anandan's method, Fleet and Jepson's method, and our method for the Hamburg taxi sequence shown in Fig. 3. Fleet and Jepson's method provides the sparse optical flow field shown in Fig. 6a. In Horn and Schunck's and Anandan's results shown in Fig. 6b,c, the optical flow field has been smoothed across the motion boundaries. Some spurious optical flow vectors are present at the stationary regions. Neither Horn and Schunck's nor Ananadan's approach takes motion boundaries into account. Therefore, the results appear to be oversmoothed. However, our method explicitly takes motion boundaries into account, and provides better results as shown

Table 1. Performance comparison of optical flow methods on the random dot, square images.

Method	Average error	Standard deviation	Density
Fleet and Jepson	7.38°	11.79°	13.2%
Horn and Schunck	2.89°	9.01°	100%
Anandan	2.16°	8.47°	100%
Proposed method	*0.61°*	*3.70°*	*100%*

in Fig. 6d. For the Hamburg taxi sequence, the execution time on a Sun Sparc 4 for Fleet and Jepson's algorithm, Horn and Schunck's algorithm, Anandan's algorithm, and our algorithm are 21.8 min, 1.5 min, 14.2 min, and 4.4 min, respectively. Our algorithm, with the overhead of computing line process, is a factor of three times slower than Horn and Schunck's algorithm. However, this deterministic approach is computationally more efficient than the other methods for minimizing the cost function involving the line process, such as simulated annealing, ICM, etc. Also, the adaptive multigrid approach not only takes care of large motion, but also greatly reduces the computation time.

6.2 Occlusion and multiple motions

Introducing line process fields into the regularization formulation allows us to handle multiple motions. As an example, we first generated a synthetic sequence with two squares moving in opposite directions with a displacement of 1 pixel/frame. In this sequence one square occludes the other square in frame 3. The image sequence is shown in Fig. 7a,b. The optical flow obtained by Horn and Schunck's algorithm is shown in Fig. 7d, where the optical flow has globally been smoothed without the motion boundaries being taken into account. Our algorithm explicitly takes care of motion boundaries to ensure that smoothing is not imposed across step-intensity edges or occlusion boundaries in an attempt to handle occlusion. The optical flow computed by our single-scale algorithm is shown in Fig. 7c, and it is much better than Horn and Schunck's algorithm near the motion boundary.

Next, the real image sequences shown in Fig. 8a,b were used for this experiment, which contains two toy vehicles moving in the opposite directions. The measured displacement is about 1–2 pixels/frame for each vehicle. The results for the optical flow obtained from our single-scale algorithm and Horn and Schunck's algorithm are shown in Fig. 8c,d. Our algorithm provides better results, particularly along motion boundaries.

7 Summary

In this paper, we address the problem of optical flow estimation and motion segmentation. Line processes and a piecewise smoothness constraint are introduced to formulate the problem based on Markov random models, and the minimization solution is obtained by mean field techniques. The line process prevents the smoothness constraint from being imposed across motion boundaries. Since motion discontinuities represented by the line process are embedded into the optical flow estimation, our algorithm provides accurate estimates of optical flow especially along motion boundaries. It handles occlusion and multiple motions. Experimental results indicate that our approach provides good estimates of optical flow and motion boundaries. Using the performance measures proposed by Barron et al. (1994), we show that the proposed algorithm outperforms other well-known algorithms in terms of estimation accuracy and timing.

Our algorithm requires the edge detector to provide complete motion boundaries, along with some possible spurious edges, as initial values for the line processes. The spurious edges are smoothed out by our algorithm. No robust initialization is applied beyond the Canny edge detector. If the Canny edge detector produces incomplete motion boundaries for a given image sequence, our algorithm performs similarly to Horn and Schunck's algorithm on this sequence.

We are currently working on extending this framework to track coherently moving regions over stream of images. Furthermore, we can apply the SFM algorithm to each segmented moving object to obtain the motion parameters and a depth map.

Acknowledgement. The authors thank the anonymous referees who provided very helpful comments. The research reported in this paper was partially supported by National Science Foundation grants CDA 9122006 and IRI-9220768.

References

1. Abdelqader I, Rajala S (1993) Motion estimation from noisy image data. International Conference on Acoustics, Speech, and Signal Processing, Minneapolis, MN, USA. IEEE New York, pp 209–212
2. Adiv G (1985) Determining three-dimensional motion and structure from optical flow generated by several moving objects. Patt Anal Machine Intell 7:384–401
3. Anandan P (1989) A computational framework and an algorithm for the measurement of visual motion. Int J Comput Vis 2:283–310
4. Barron JL, Fleet DJ, Beauchemin SS (1994) Performance of optical flow techniques. *Int J Comput Vis* 12:43–77
5. Battiti R, Amaldi E, Koch C (1991) Computing optical flow across multiple scales: an adaptive coarse-to-fine strategy. Int J Comput Vis 6:133–145
6. Black M, Anandan P (1991) Robust dynamic motion estimation over time. IEEE Conference on Computer Vision and Pattern Recognition, Lahaina, Maui, Hawaii. IEEE Computer Society Press, Los Alamitos, California, pp 296–302
7. Bouthemy P, Rivero J (1987) A hierachical likelihood approach for region segmentation according to motion-based criteria. International Conference on Computer Vision, London, England. IEEE Computer Society Press, Washington DC, pp 463–467
8. J O Braddick (1974) A short-range process in apparent motion. Vis Res 14:519–527
9. Burt P (1984) The pyramid as a structure for efficient computation. Multiresolution image processing and analysis, Springer, Berlin Heidelberg New York pp 6–35
10. Fleet D, Jepson A (1990) Computation of component image velocity from local phase information. Int J Comput Vis, 5:77–104
11. Geiger D, Girosi F (1991) Parallel and deterministic algorithms from MRF's: surface reconstruction. Patt Anal Machine Intell 13:401–412
12. Harris J, Koch C, Staats E, Luo J (1990) Analog hardware for detecting discontinuities in early vision. Int J Comput Vis, 4:211–223
13. Heitz F, Bouthemy P (1993) Multimodal estimation of discontinous optical flow using Markov random fields. Patt Anal Machine Intell, 15:1217–1232
14. Hildreth EC (1984) Measurement of visual motion. MIT Press, Cambridge, Mass

15. Horn B, Schunck G (1981) Determining optical flow. Artif Intell 17:185–203
16. Irani M, Rousso B, Peleg S (1992) Detecting and tracking multiple moving objects using temporal integration. European Conference on Computer Vision, Santa Margherita, Ligure, Italy. Springer, Berlin Heidelberg New York, pp 282–287
17. Jepson A, Black M (1993) Mixture models for optical flow computation. IEEE Conference on Computer Vision and Pattern Recognition, New York. IEEE Computer Society Press, Los Alamitos, California, pp 760–761
18. Julesz B (1971) Foundations of Cyclopean perception. University of Chicago Press, Chicago
19. Konrad J, Dubois E (1992) Bayesian estimation of motion vector fields. Patt Anal Machine Intell 14:910–927
20. Murray D, Buxton B (1987) Scene segmentation from visual motion using global optimation. Patt Anal Machine Intell 9:220–228
21. Mutch K, Thompson W (1985) Analysis of accretion and deletion at boundaries in dynamic scenes. Patt Anal Machine Intell 7:133–138
22. Poggio T, Gamble E, Little J (1988) Parallel integration of vision modules. Science 242:436–440
23. Reichardt W, Poggio T (1980) Figure-ground discrimination by relative movement in the visual system of the fly. Part 1: experimental results. Biol cybernetics 35:81–100
24. Schunck B (1989) Image flow segmentation and estimation by constraint line clustering. Patt Anal Machine Intell 11:1010–1027
25. Spoerri A, Ullman S (1987) The early detection of motion boundary. International Conference on Computer Vision, London, England. IEEE Computer Society Press, Washington DC, pp 209–218
26. Thompson W, Mutch K, Berzins V (1985) Dynamic occlusion analysis in optical flow fields. Patt Anal Machine Intell 7:374–383
27. Vlontzos J, Geiger D (1992) A MRF approach to optical flow estimation. IEEE Conference on Computer Vision and Pattern Recognition, Champaign, Illinois. IEEE Computer Society Press, Los Alamitos, California, pp 853–856
28. Wang J, Adelson E (1993) Layer representation for motion analysis. IEEE Conference on Computer Vision and Pattern Recognition, New York. IEEE Computer Society Press, Los Alamitos, California, pp 361–366
29. Zhang J, Hanauer J (1993) The mean field theory for image motion estimation. International Conference on Acoustics, Speech, and Signal Processing, Minneapolis, MN, USA. IEEE, New York, USA, pp 197–200