DTU Library

Estimation of Dense Image Flow Fields in Fluids

Larsen, Rasmus; Conradsen, Knut; Ersbøll, Bjarne Kjær

[Link back to DTU Orbit](#)

Estimation of Dense Image Flow Fields in Fluids

Rasmus Larsen, Knut Conradsen, and Bjarne Kjær Ersbøll

Abstract— The estimation of flow fields from time sequences of satellite imagery has a number of important applications. For visualization of cloud or sea ice movements in sequences of crude temporal sampling, a satisfactory nonblurred temporal interpolation can be performed only when the flow field or an estimate thereof is known. Estimated flow fields in weather satellite imagery might also be used on an operational basis as inputs to short-term weather prediction. In this paper, we describe a method for the estimation of dense flow fields. Local measurements of motion are obtained by analysis of the local energy distribution, which is sampled by using a set of three-dimensional (3-D) spatio-temporal filters. The estimated local energy distribution also allows us to compute a confidence measure of the estimated local normal flow. The algorithm, furthermore, utilizes Markovian random fields in order to integrate the local estimates of normal flows into a dense flow field by using measures of spatial smoothness. To obtain smoothness, we will constrain first-order derivatives of the flow field. The performance of the algorithm is illustrated by the estimation of the flow fields corresponding to a sequence of Meteosat thermal images. The estimated flow fields are used in a temporal interpolation scheme.

Index Terms—Fluid flow, Markov random field, Meteosat optical flow.

I. INTRODUCTION

INDEPENDENTLY moving objects, rotation, dilation, and shear in image sequences combine to produce complex velocity fields. Therefore, valid velocity estimation is restricted to local computations. This ensures that, for sufficiently smooth velocity fields, the estimation can be based on translational image velocity.

Coherent image translation is the basis for several computational methods. The main methods include correlation-based methods [1]–[3], differential methods [4], [5], energy-based methods [6]–[8], and phase-based methods [9], [10].

Restricting measurements to small spatio-temporal neighborhoods, however, often results in the measurements being based on one-dimensional (1-D) intensity structures (edges and/or lines). In this case, we can only determine the component of the velocity orthogonal to the intensity contour reliably. This is known as the aperture problem [11].

In general this problem exists in neighborhoods of the image sequence that have a 1-D structure only, as well as neighborhoods that have no structure at all, that is, in homogeneous areas. On the other hand, for image sequence neighborhoods that exhibit two dimensional (2-D) spatial structures, such as

intensity corners or various textured regions, we can reliably extract the true velocity. In order to identify the type of neighborhood, several approaches have been reported. In a correlation-based approach, Anandan [12] used the curvature of the match surface and Nagel and Enkelmann [5] investigated the use of second-order spatial derivatives to identify the neighborhood. The use of quadrature-type filters tuned to different spatio-temporal frequencies has been used to identify the type of the neighborhood in energy—as well as phase-based approaches [10], [13].

Because the aperture problem results in flow fields that are not fully constrained, an assumption of (piecewise) smoothness of the velocity field must be applied in order to obtain a dense velocity field. One way of doing this is by applying a restriction that forces the spatial derivatives to be small. These restrictions are referred to as smoothness constraints [4]. Methods utilizing this type of smoothness constraints include the work of [5], [14]–[16]. Other approaches based on spatial filtering also have been reported. Simpson and Gobat [17], for example, used a vector median filter to obtain a smooth velocity field.

It is evident that local velocity estimation algorithms that are able to distinguish between the different natures of the neighborhood and thus the estimated velocity—component velocity or not—should be more successful than algorithms that are not. In Section III-A, we will consider a method to extract velocity estimates as well as related directional probabilities of the estimates, based on the local energy distribution. This is especially important because of the distortion of features, due to physical processes, that can occur in fluids (e.g., [18]).

In Section III-B, we will describe the implementation of the smoothness constraint. This smoothness constraint is formulated as a prior probability distribution for the velocity field that assigns high probability to fields that have small first-order spatial derivatives and low probability to fields that have large spatial derivatives. We will, furthermore, suggest an observation model that carefully relates the local estimates of normal velocity to a particular realization of the velocity field. Here it is important to recognize that local estimates are normal flows and that standard smoothing operations are not well suited in this case. Finally, we will combine the prior distribution and the observation model into a posterior distribution by using Bayes' theorem.

The algorithm is illustrated by the estimation of smooth cloud velocity fields in a series of Meteosat satellite images. Other work on cloud-motion estimation has been reported by [1] and [2]. The estimated flow fields will be used in two temporal interpolation schemes. First, intermediate frames between each of the originals are generated. Second, the
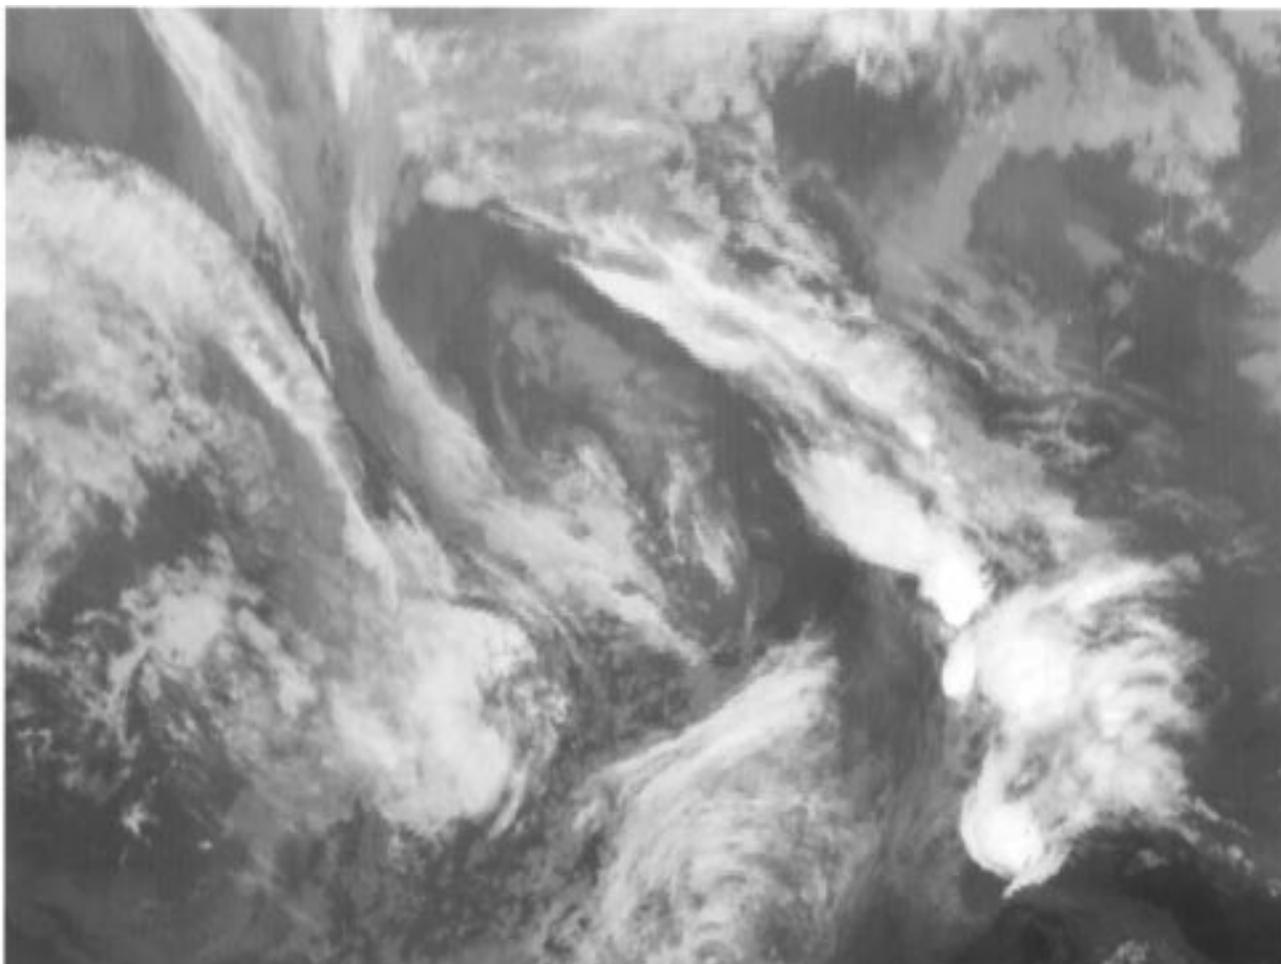
Fig. 1. Infrared channel from the Meteosat 5 satellite mapped to a polar stereographic projection showing the cloud cover over Europe at August 24, 1994, 05:00 GMT. The images are provided by the Danish Meteorological Institute.

estimated flow field is used to generate a frame to replace a missing frame in the sequence.

II. Data

The algorithms described in this article will be applied to a sequence of images recorded by the Meteosat 5 satellite. The images are from the infrared channel (10.5–12.5 μm). Preliminary processing performed by the Danish Meteorological Institute consists of mapping to a polar stereographic projection and resampling to a 576×768 equirectangular 7-km grid. The center of the grid is at 48.4° N latitude and 8.2° E longitude. The images are recorded with a time interval of 30 min. We will be using a 384×512 sequence of subimages centered at August 24, 1994, 05:00 GMT. The center image of the sequence is shown in Fig. 1.

III. Methods

This section is divided into two subsections. In Section III-A, we will consider the local velocity estimation by using a set of spatio-temporal directional quadrature filter pairs. After this, in Section III-B, we will formulate an algorithm for integrating these local estimates to a dense velocity field

by using smoothness constraints based on first-order spatial derivatives of the velocity field.

A. *Local Velocity Estimation*

Because motion estimation in image sequences can be viewed as identification of patterns repeating themselves over time, it is natural to try to describe the motion analysis in the Fourier domain. Let us consider a neighborhood containing a 1-D intensity structure (e.g., a line) that translates coherently through time. In the spatio-temporal domain, this corresponds to a neighborhood of iso-grey-level planes. Let these planes be given by their unit normal vector $k = (k_1, k_2, k_3)^T$. We will refer to this vector as the spatio-temporal orientation vector. The nonzero Fourier coefficients of this neighborhood are concentrated near a line in the Fourier domain defined by k.

The relationship between the spatio-temporal orientation vector and the normal flow vector $\mu = (\mu, \nu)^T$ is illustrated in Fig. 2. In this figure, a line translating with constant velocity through space is shown at four time instants. These lines span a plane in space-time, and k is the normal vector of this plane. In the figure $\mu^* = (\mu, \nu, 1)^T$ and the line orientation vector $e = (-k_2, k_1, 0)^T$. Because μ^* is perpendicular to k as well as e, its direction is given by the outer product of these vectors.
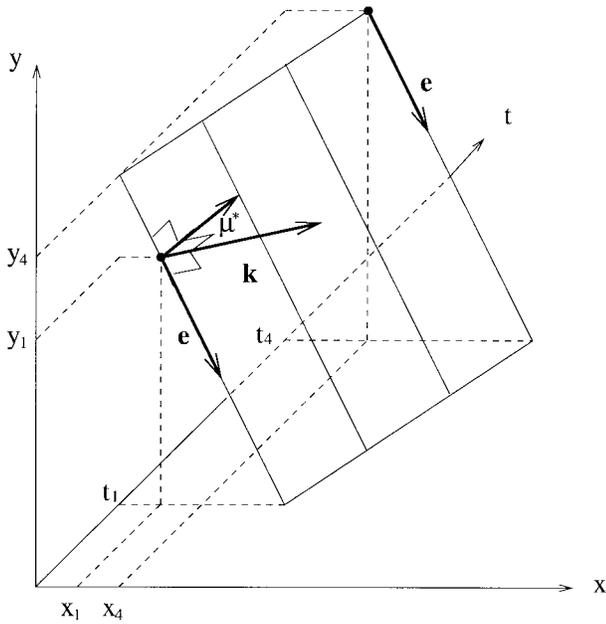
Fig. 2. Relationship between the spatio-temporal orientation vector and the normal flow. Refer to the text for a detailed explanation.

Finally, this outer product should be normalized so that the temporal coordinate equals 1. This results in the following relationship:

$$\boldsymbol{\mu} = (\mu, \nu)^T = \frac{-k_3}{k_1^2 + k_2^2} (k_1, k_2)^T. \tag{1}$$

Now, in order to estimate the spatio-temporal orientation vector, we will sample the Fourier domain by using a set of spatio-temporal filters. Using a method developed by Knutsson [8], this can be achieved by applying a set of directional quadrature filter pairs symmetrically distributed over the spatio-temporal orientation space. The directional quadrature filter pairs consist of a real even part and an odd imaginary part. The phase of the transfer function for the real part is shifted 90°, relative to the imaginary part. By squaring and adding the two filter responses we obtain a phase-independent estimate of the spectral density of the corresponding image structure.

As an approximation to such a set of filters, we will employ a set of P Gabor filters [19] with center frequencies given by the vertices of a diametrical symmetric regular polyhedron [8] (e.g., using the set of six filter directions, $\mathbf{n}_p$, $p = 1, \cdots, P(= 6)$ listed in Table I. The pth Gabor filter consists of a Gaussian function shifted to the point $\boldsymbol{k}_p = (k_{p1}, k_{p2}, k_{p3})^T$ in frequency space. The corresponding convolution mask consists of an odd (real) and an even (imaginary) part, which are easily computed by taking the Fourier transforms

$$q_p^e(z) = \frac{1}{(2\pi)^{3/2}\sigma^3} \cos(\boldsymbol{k}_p^T z) \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right)$$

$$q_p^o(z) = \frac{1}{(2\pi)^{3/2}\sigma^3} \sin(\boldsymbol{k}_p^T z) \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \tag{2}$$

$p = 1, \cdots, P$. Although Gabor filters are not quadrature pairs, they provide a reasonably good approximation for sufficiently

TABLE I
SPATIO-TEMPORAL DIRECTIONS OF THE GABOR FILTERS ARE GIVEN BY THESE COEFFICIENTS, WHERE $a = 2/\sqrt{10 + 2\sqrt{5}}$ AND $b = (1 + \sqrt{5})/\sqrt{10 + 2\sqrt{5}}$

p	1	2	3	4	5	6
x	a	a	b	-b	0	0
y	0	0	a	a	b	-b
t	b	-b	0	0	a	a

small bandwidths [20]. If the bandwidth in octaves is measured at one standard deviation of the Gaussian envelope in the frequency domain, it is given by (e.g., [13])

$$B = \log_2\left[\frac{\|\boldsymbol{k}_p\| + \frac{1}{\sigma}}{\|\boldsymbol{k}_p\| - \frac{1}{\sigma}}\right] \Rightarrow \|\boldsymbol{k}_p\| = \frac{1}{\sigma}\frac{2^B + 1}{2^B - 1} \tag{3}$$

$p = 1, \cdots, P$. Thus, the direction of $\boldsymbol{k}_p$ is given by the choice of $\mathbf{n}_p$, and the length $\|\boldsymbol{k}_p\|$ is given by the choice of B and σ. The choice of these parameters should be made taking into account the maximum expected velocity in the image sequence. A maximum velocity of $v_{\max}$ corresponds to a spatio-temporal frequency of $2\pi(v_{\max}^2 + 1)^{-1/2}$. Here the product of spatio-temporal frequency and spatio-temporal period is 2π. Setting the extent of the filter to one standard deviation of the Gaussian envelope in the frequency domain and by using (3), we find that

$$v_{\max} < \sqrt{[\pi\sigma(2^B - 1)]^2 - 1}. \tag{4}$$

Finally, the energy distribution of the Fourier domain as estimated by the set of quadrature filter pairs may be represented by the tensor (e.g., [8])

$$\boldsymbol{T} = \sum_{p=1}^{P} \sqrt{(q_p^e)^2 + (q_p^o)^2}\, \boldsymbol{n}_p \boldsymbol{n}_p^T = \sum_{p=1}^{P} q_p \boldsymbol{n}_p \boldsymbol{n}_p^T \tag{5}$$

where q_p is the output from the pth quadrature filter pair, and $\boldsymbol{n}_p$ is the unit normal vector defining the direction of the filter.

In order to find the direction of maximum spectral density, we must find the unit vector $\boldsymbol{k}$ that maximizes $\boldsymbol{k}^T \boldsymbol{T} \boldsymbol{k}$. This vector is the eigenvector corresponding to the largest eigenvalue of $\boldsymbol{T}$ (e.g., [21]). So for the coherently translating 1-D intensity structure, which has an effectively 1-D Fourier domain, the spatio-temporal orientation vector is found by an eigenanalysis of $\boldsymbol{T}$. Because the Fourier domain is 1-D, $\boldsymbol{T}$ has only one nonzero eigenvalue.

Now, if the translating structure has a 2-D intensity structure (e.g., a grey-level corner), the spatio-temporal domain is described by two spatio-temporal orientations, each of which gives rise to a nonzero eigenvalue of $\boldsymbol{T}$. The eigenvectors corresponding to these nonzero eigenvalues each correspond to a normal flow by using (1).

Given the true flow, we may determine the difference between the projection of the true flow onto either of these estimated normal flows and the normal flows themselves

$$d_k(x_i, y_i) = \left| [\boldsymbol{u}(x_i, y_i) - \boldsymbol{\mu}_k(x_i, y_i)]^T \cdot \frac{\boldsymbol{\mu}_k(x_i, y_i)}{\|\boldsymbol{\mu}_k(x_i, y_i)\|} \right| \tag{6}$$

$k = 1, 2$, where $\boldsymbol{u}(x_i, y_i)$, $\boldsymbol{\mu}_k(x_i, y_i)$ are the true flow and the estimated normal flows taken at the position (x_i, y_i). It is

the (weighted) sum of squares of these distances that should be minimized across the image in order to obtain an estimate of the velocity field.

Deviation from the assumption of coherent translation, imperfectly designed filters and noise will result in nonzero Fourier coefficients not being contained in one single line or plane. In this case, all three of the tensor eigenvalues will be nonzero. This allows us to extract information about the quality of the estimates of constraint lines we get from the eigenvectors corresponding to the two largest eigenvalues. Because imperfect conditions result in a nonzero third eigenvalue, we suggest using a confidence measure for each of the linear constraints, based on the difference of the corresponding eigenvalue and the smallest eigenvalue. Furthermore, a normalization of this difference should be made. This is evident as a noise-free high-step edge measures the motion just as well as a lower step does. We propose the following confidence measure for each of the linear constraints, given by the eigenvectors corresponding to the two largest eigenvalues

$$w_k(x_i, y_i) = \frac{\lambda_k(x_i, y_i) - \lambda_3(x_i, y_i)}{\lambda_k(x_i, y_i)}, \qquad k = 1, 2 \quad (7)$$

where $\lambda_1(x_i, y_i) \geq \lambda_2(x_i, y_i) \geq \lambda_3(x_i, y_i) \geq 0$ denote the eigenvalues of the tensor at position (x_i, y_i). This confidence measure approaches zero when the difference of the corresponding eigenvalue and the smallest eigenvalue approaches zero, and it attains its maximum value of one, when the smallest eigenvalue is zero, and the corresponding eigenvalue is the largest, or is equal to the largest eigenvalue, respectively.

B. Integration of Local Measurements

As mentioned in Section I, we will apply an assumption of smoothness with the purpose of fully constraining the velocity field by forcing the spatial derivatives of the velocity field to be small. Since Horn and Schunk's original paper [4], this has been investigated by several authors (e.g., [22] and [23]). One way of formulating such a smoothness constraint is by use of Markovian random fields [16], [24]. We do this by using the Bayesian paradigm [25]. First, we will formulate a prior distribution for the velocity field, based on the spatial derivatives of the field. If the spatial derivatives are implemented using the following finite differences:

$$\boldsymbol{u}_x(x_i, y_i) = \boldsymbol{u}(x_i + 1, y_i) - \boldsymbol{u}(x_i, y_i)$$
$$\boldsymbol{u}_y(x_i, y_i) = \boldsymbol{u}(x_i, y_i + 1) - \boldsymbol{u}(x_i, y_i) \quad (8)$$

where (x_i, y_i), $i \in \{1, 2, \cdots, N\}$ are the pixel positions, then the prior distribution of the flow field may be described by a Gibbs distribution $p(\{\boldsymbol{u}\}) = (1/Z) \exp(-\beta U_1)$, where Z is a normalization constant and the energy term is given by

$$U_1 = \sum_{i=1}^{N} \|\boldsymbol{u}_x(x_i, y_i)\|^2 + \sum_{i=1}^{N} \|\boldsymbol{u}_y(x_i, y_i)\|^2. \quad (9)$$

This probability distribution assigns high probability to fields that exhibit small derivatives and low probability to fields with high spatial derivatives. We will need a first-order neighborhood to implement U_1. In Fig. 3(a), this neighborhood is
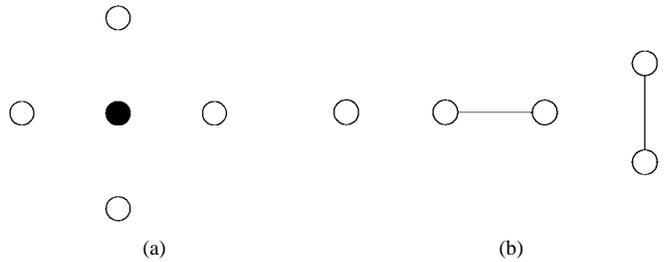

(a)   (b)

Fig. 3. Neighborhood systems and cliques for the pixel process. (a) The neighborhood configurations for (interior) points for a first-order Markovian random field. The points marked with the symbol ○ are the neighbors of the point marked ●. (b) The cliques corresponding to the neighborhood configuration in (a).

shown. In Fig. 3(b), the cliques necessary to implement the energy function are shown.

Having constructed this prior distribution for the flow field, we will now concern ourselves with an observation model. The observation model relates the local observations or measurements of velocity to any particular realization of the prior distribution. This is done by a conditional Gibbs distribution

$$P(\text{OBS}|\{\boldsymbol{u}\}) = \frac{1}{Z} \exp(-\alpha U_0)$$
$$= \frac{1}{Z} \exp\left[-\alpha \sum_{i=1}^{N} \sum_{k=1}^{2} w_k(x_i, y_i) \, d_k(x_i, y_i)^2 \right] \quad (10)$$

where $d_k(x_i, y_i)$ is the difference between the projection of the true flow onto the normal flow, given by the kth eigenvector and the normal flow itself at pixel (x_i, y_i), as described by (6). $w_k(x_i, y_i)$ is the confidence measure corresponding to this normal flow, given by (7). Z is a normalization constant. By using a Gibbs energy function that punishes large deviations in the projection of the true flow onto the observed normal flows, we allow smoothing in the direction not constrained by the normal flows, while smoothing in the direction of the normal flows is punished. Furthermore, the use of the confidence measures w_k, derived in the previous section as weights, allows us to take into consideration the quality of our measurements.

The prior distribution and the observation model are combined into a posterior distribution by using Bayes' theorem. The energy function of the posterior distribution thus becomes

$$U = \alpha \sum_{i=1}^{N} \sum_{k=1}^{2} w_k(x_i, y_i) \, d_k(x_i, y_i)^2$$
$$+ \beta \sum_{i=1}^{N} \left[\|\boldsymbol{u}_x(x_i, y_i)\|^2 + \|\boldsymbol{u}_y(x_i, y_i)\|^2 \right]. \quad (11)$$

In this energy function, we can control the properties of the estimated motion field. The smoothness is controlled by β and the faith in the observed or measured normal flows is controlled by α.

We can now apply a maximization scheme to the posterior distribution in order to obtain the maximum a *posteriori* estimate of the velocity field. Note that maximizing the posterior probability is equivalent to minimizing U in (11). Thus,
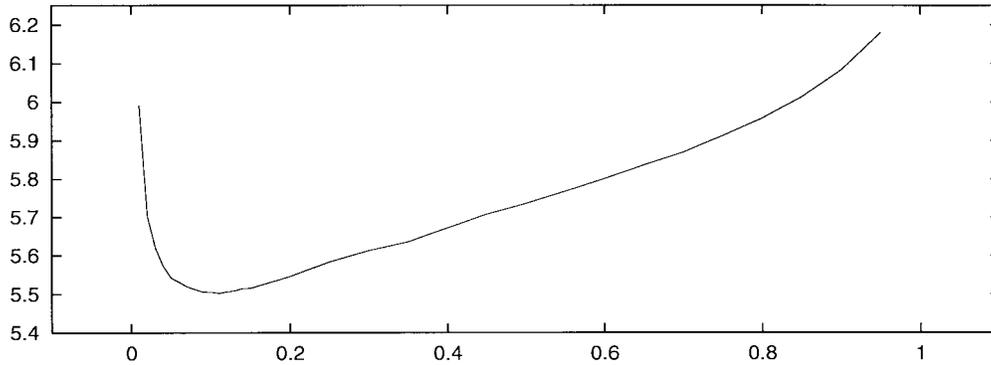
Fig. 4. The criterion for choosing α is shown as a function of α. The global minimum of 5.5 is attained for $\alpha = 0.11$.

without loss of generality, we can set $\beta = 1 - \alpha$ where $\alpha \in [0, 1]$. This has been implemented by using the iterated conditional modes (ICM) scheme by Besag [26]. The objective function is quadratic in the variables and the corresponding Hessian matrix is positive-definite for positive confidence weights. This ensures that we have one and only one local minimum, and that this is also the global minimum. In this case, ICM will converge to this minimum. We will start the algorithm by initializing all vectors to those of the previous image in the sequence, or in the case of the first image, to **0**.

It should be noted that the prior distributions listed here are suitable for a globally smooth field, i.e., such fields that we observe in fluids. Relevant examples in this case are the cloud movement recorded by weather satellites or ground-based radar stations as well as ocean current patterns that may also be observed from various spacecraft. In other situations, we are faced with independently moving well-defined objects, and in these cases, additional modeling is necessary (e.g., the inclusion of a line process [24] or a labeling procedure [27] in order to be able to estimate motion boundaries).

IV. RESULTS

This section is divided into three subsections. First, we will illustrate the local estimation of normal flows and how these normal flows may be integrated into a smooth flow field by using the proposed prior distribution. Second, we will use an estimated flow field to perform a temporal interpolation, and third, we will investigate how the estimated flow field may be used in order to generate replacement frames in the case of missing frames in an image sequence of a time varying scene.

A. Estimation of the Flow Field

As described in Section III-A, we will describe the local Fourier domain by the tensor given in (5). It should be noted here that the Fourier method of estimating the local velocities assumes that patterns of thermal gradient are moving without change over the time period spanning the actual filter size [18].

We will use a set of six Gabor filter pairs to sample the Fourier domain. The spatio-temporal directions of the Gabor filters are shown in Table I. Setting the bandwidth to one octave to approximate quadrature filter pairs, and by assuming a maximum (normal) velocity of 100 ms^{-1}, as prescribed in [2], we find by using (4) that the standard
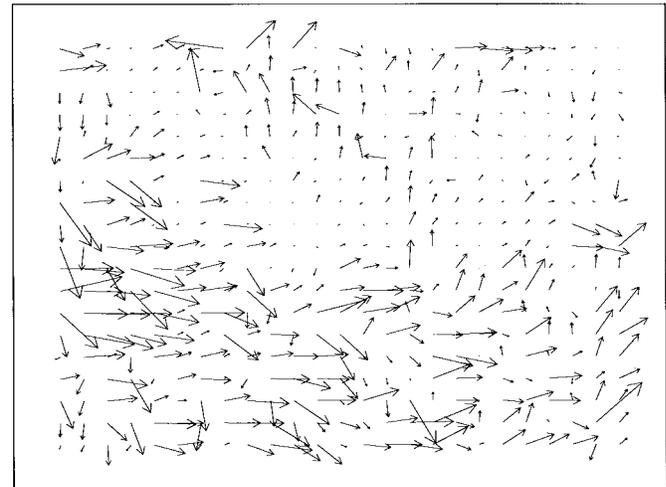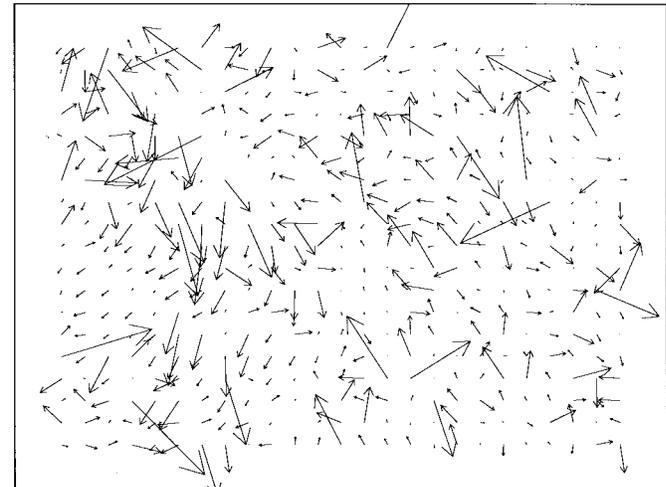
Fig. 5. Normal flows corresponding to the (a) largest eigenvalue and (b) the second largest eigenvalue computed from the Meteosat sequence corresponding to the 05:00 GMT frame.

deviation of the Gaussian envelope should be approximately 2.5 when operating on the second level of a Gaussian pyramid [28] (i.e., corresponding to a pixel size of 28 km/pixel). By truncating the Gaussian envelope of the Gabor filters at three standard deviations, we arrive at the filter size $17 \times 17 \times$
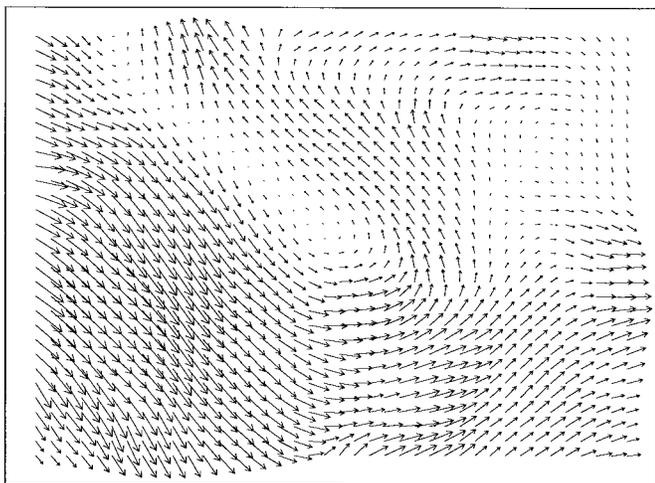
Fig. 6. Velocity field computed from the Meteosat sequence corresponding to the 05:00 GMT frame using $\alpha = 0.11$.

17. Because the Gabor filter kernels are separable [7], the 12 three-dimensional (3-D) convolutions may be performed by 84 1-D convolutions. Thus, reducing the computational load by a factor of 40.

The parameters α and $\beta(= 1 - \alpha)$ of (11) allow us to control the final velocity field. We will choose the optimal α for a given image at time t as the one that minimizes the mean value of the absolute differences between the pixel intensities of a predicted and the actual image at time $t + 1$. The predicted image is calculated by moving the pixels of the image at time t according to the estimated flow field at time t and resampling using nearest neighbor interpolation. The mean value of the absolute differences between the pixel intensities of a predicted and the actual image at time $t + 1$ is shown in Fig. 4. For the Meteosat sequence corresponding to the 05:00 GMT frame we find that α should be chosen to $\alpha = 0.11$. The estimated normal flows for this image are shown in Fig. 5 and the estimated flow field is shown in Fig. 6.

B. Temporal Interpolation in a Meteosat Satellite Sequence

An obvious application of the flow field estimation is temporal interpolation. We will illustrate how this may be done by using the flow field from Fig. 6. We will generate an intermediate image between the original images recorded at 05:00 and 05:30 GMT. We extrapolate from each of the two original images by using the corresponding flow field estimates. As we emphasize, while not blurring the image, we will use a nearest neighbor interpolation scheme to sample the extrapolated images to the original grid. Finally, we take the average of the two extrapolations. This is our intermediate image.

The interpolated images are shown in Fig. 7(a). A naïve approach to an interpolation would be to simply take the average of the two originals. The result of this is shown in Fig. 7(b). We can see that the naïve approach results in a much more blurred result than the method based on the estimated flow field. This intermediate image on the other hand displays a very satisfactory temporal interpolation.
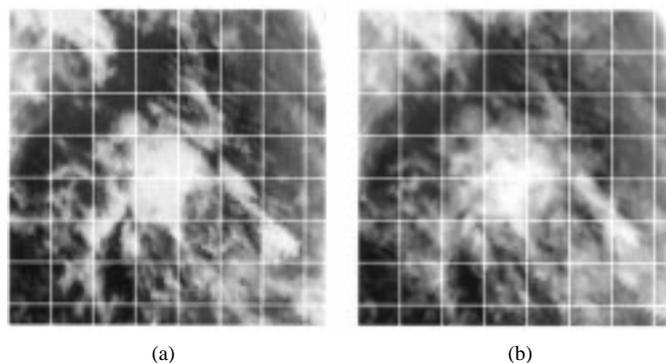

(a)   (b)

Fig. 7. (a) Segment of the artificially generated intermediate image corresponding to 5:15 GMT. (b) Simple average of the two original images recorded at 05:00 and 05:30 GMT. Note the echoing of all the moving contours. It is evident that our approach results in a more satisfying, sharper intermediate image without echoes.
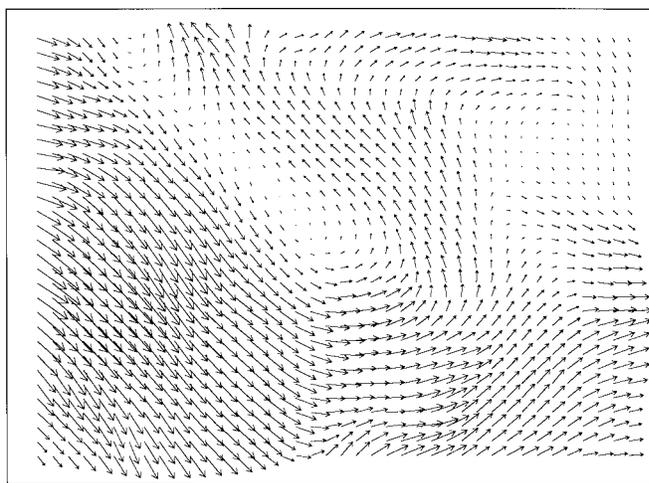

Fig. 8. Velocity field computed from the Meteosat sequence corresponding to the 05:00 GMT frame, estimated using the average of the 04:30 and the 05:30 GMT images in the place of the 05:00 image.

C. Missing Frame Replacement by Temporal Interpolation

At least twice every day we have missing frames in the Meteosat sequences. This is very disturbing to the eye when viewing the sequence as a film. We will show how the algorithm described above may be used to generate a satisfactory replacement for the missing frame. First, we need to estimate the flow field at the time of the missing frame. When performing the filtering operation, we will replace the missing frame with the average of the previous and the next images. In order to evaluate how this affects the flow estimation, we have removed the original 05:00 GMT image from the sequence, substituted it with the average of the 04:30 and 05:30 GMT images, and computed the flow field. The resulting field is shown in Fig. 8. Now we can generate the replacement 05:00 image by warping the 04:30 image according to this flow field.

Fig. 9 shows the capability for the missing frame replacement of the algorithm. In the top row, a 120×120 pixel segment from the Meteosat sequence corresponding to times 04:30, 05:00, and 05:30 is shown. In the bottom row, we see the replacement 05:00 image. This image is generated by estimating the flow field in a sequence where the 05:00 image
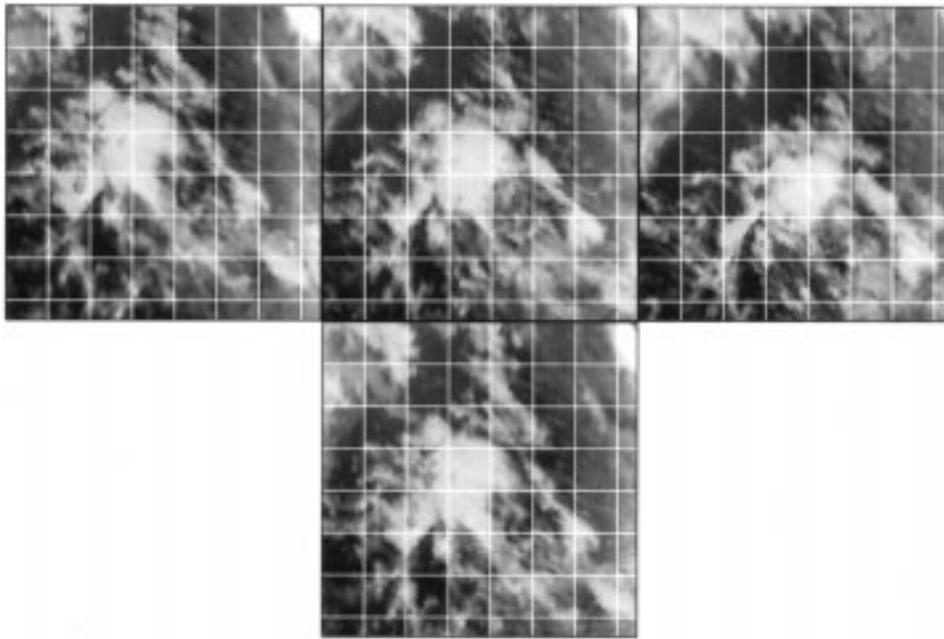
Fig. 9.   This figure shows the capability for missing frame replacement of the algorithm. In the top row, a 120 × 120 pixel segment from the Meteosat sequence corresponding to times 04:30, 05:00, and 05:30 is shown. In the bottom row, we see the replacement 05:00 image.

is replaced by the pixelwise average of the 04:30 and the 05:30 images and subsequently by using the 04:30 flow field to project the 04:30 image to 05:00. To ease the identification of cloud features for the reader, a white grid is overlaid. It is evident that position as well as shape of the clouds in the replacement image correspond well to the original 05:00 image.

V. Discussion

A. Data

The purpose of the exercises described above has been to perform a nonblurred temporal interpolation of the sequence. It should be noted that the extracted flow fields do not necessarily assimilate the corresponding wind fields. This is true because only some clouds move with the wind. Other problems are that we are only observing the top layer of the clouds and that the wind speed varies with height [29].

The use of the Fourier technique assumes that the patterns of thermal gradient are moving without (or with little) change over the time period spanned by the filters. Other physical processes that have an effect are diffusion and different types of heat transfer (sources/sinks) other than advection [3], [18].

B. Estimation of Normal Flows

The spatio-temporal Gabor filtering described in Section III-A provides a robust estimation of normal flows. In addition to identifying the zero, one, or two normal flows at each pixel, it provides us with a measure of the confidence we should attach to the corresponding normal flow.

The filter design (and choice of scale) has been based on an assumption of a maximum velocity based on the physical phenomenon that we are observing. In situations where this is not possible and in order obtain better resolution of the velocity estimates, we may utilize a pyramid representation of the image sequence. Preliminary studies of this have been carried out by Heeger [7]. Other studies that exploit the use of a coarse to fine approach include [30].

C. Integration Using Markovian Random Fields

The Markovian random field approach to integrate the measured normal flows has proven to be a powerful and successful technique. We have used restrictions on the first-order spatial derivatives of the flow field to reconstruct the flow field. The diversity of the admissible fields has proven sufficient for the estimation of flow fields in the case shown. The experiments indicate that it is necessary to restrict the first-order derivatives in order to propagate velocities across regions of sparse measurements.

The smoothness constraints have been chosen from a mathematical and not a physical point of view. Smoothness constraints involving terms that are more easily interpreted may be constructed. The sum of the norm of the first-order spatial derivatives may for instance be separated into three forms of deformation: dilation, shear, and rotation [31]. Knowledge of the physical phenomena we are observing may then be translated into different weights for the three terms. Also the estimation may be improved by including restrictions on temporal variation.

D. Temporal Interpolation

By using the flow field when performing the temporal interpolations, we have achieved much more satisfying results than the simple alternatives shown. The flow field estimations and the interpolations shown have, however, been contingent

on the intensity variation across the image being due to clouds only.

In case of intensity variation, due to variation in the terrain, sea/land-, land/ice-boundaries, etc., these areas should be screened prior to applying this algorithm. If these areas are not screened, they would create false normal flows. Algorithms for cloud screening, however, do exist (e.g., [32]–[34]).

E. Comparison with Standard Technique

In order to illustrate the power of the flow-field estimation technique described in this article for the estimation of fluid motion in spacecraft data, we will make a comparison with a standard technique for motion estimation by Anandan [12]. We have chosen this method because it also results in a 100% dense flow field, and because it too recognizes and utilizes the directional confidences in the estimated local velocities.

Anandan's algorithm consists of a hierarchical matching strategy based on a Laplacian pyramid representation of the images. At each level, the velocity at a pixel is estimated as the displacement that yields the best match between image regions at different times. The best match is defined as the one that minimizes the sum-of-squared differences (SSD) within a 5×5 window with Gaussian weights. The search space for the local velocity estimates is determined as the 3×3 neighborhood of the displacement computed at the immediate coarser level. Subpixel accuracy is obtained by fitting a quadratic surface to the SSD values about the minimum SSD value found with integer values. Furthermore, directional confidence measures are derived by extracting the principle curvatures of the SSD surface at the minimum. Then the smoothness constraint of Horn and Schunk [4] is employed. The sum of the squared first-order spatial derivatives of the flow field and the deviations from the local estimates weighted by the directional confidences is minimized over the entire field. The algorithm is initialized at the coarsest level by centering the search area at the position corresponding to a displacement equal to **0**.

In this comparison, we also start by subsampling the original images by a factor of four. By using four levels of the Laplacian pyramid, we are able to contain a maximum velocity of $100\text{-ms}^{-1} = 6.4$ pixels/frame in the search area at the coarsest level. We use a software implementation of Anandan's algorithm used in and made available by Barron *et al.* [35]. In Fig. 10, the estimated flow field corresponding to August 24, 05:00 GMT of the Meteosat sequence using Anandan's method is shown. It is evident that there are qualitative differences between the flow field estimated by our method shown in Fig. 6 and that using Anandan's method. The latter field does not possess the degree of continuity that would be expected of a natural phenomenon as the one we are observing. We may quantify this difference by using the same measure of goodness as we did when estimating the smoothness parameter of our method. For the optimal α we found that the mean value of the absolute differences between the prediction of the next image in the sequence using the estimated flow field was 5.5. For all other α's in the range [0.01, 0.99] this measure is never above 6.2. For Anandan's method, we find the same quantity
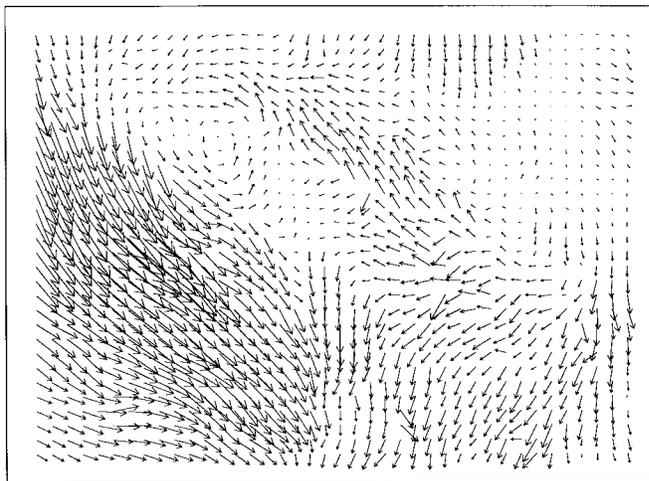

Fig. 10. Velocity field computed from the Meteosat sequence corresponding to the 05:00 GMT frame by use of the Anandan algorithm [12].

to be 8.3, which is significantly higher. Thus, showing the competitiveness of the algorithm described in this paper.

VI. CONCLUSION

A technique for estimation of smooth flow fields based on observations of the local distribution of energy measured by a set of Gabor-type filters and a prior distribution involving restrictions on the first-order derivatives using Markovian random fields has been described. The estimation technique is formulated using the Bayesian paradigm, and the maximum *a posteriori* estimate is found by use of the iterated conditional modes algorithm. The parameters of the model consist of the choice of standard deviation of the Gaussian envelope of the Gabor filters and of a parameter that balances the trust in the local observations of flow and the smoothness constraint that is inherent in the Markov model. The first of these parameters is chosen by setting a maximum velocity. For the application shown a physical limit is assumed. The latter parameter is estimated by using an exterior criterion based on the ability to predict the next image in the sequence by use of the estimated flow field.

The estimated flow fields have been used successfully in two temporal interpolation schemes. First we have used the flow fields to generate intermediate frames in a sequence of Meteosat images, and second, we have used the estimated flow fields to generate replacements for missing frames in the sequence. Both schemes are superior to the simple alternatives shown. Furthermore we have made a comparison between our algorithm and a standard method for estimating smooth flow fields. For the case study of this article our algorithm turned out to be superior to the standard algorithm.

the reviewers are thanked, as their comments on the original submission have led to a more complete and consistent paper.

REFERENCES

[1] J. A. Leese and C. A. Novak, "An automated technique for obtaining cloud motion from geosynchronous satellite data using cross correlation," *J. Appl. Meteor.,* vol. 10, pp. 118–132, Feb. 1971.

[2] J. Schmetz and K. Holmlund, "Operational cloud-motion winds from Meteosat infrared images," *J. Appl. Meteor.,* vol. 32, pp. 1206–1225, July 1993.

[3] D. D. Wahl and J. J. Simpson, "Satellite derived estimates of the normal and tangential components of near-surface flow," *Int. J. Remote Sensing,* vol. 12, pp. 2529–2571, 1991.

[4] B. K. P. Horn and B. G. Schunk, "Determining optical flow," *Artif. Intell.,* vol. 17, pp. 185–203, 1981.

[5] H.-H. Nagel and W. Enkelmann, "An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. PAMI-8, pp. 565–593, Sept. 1986.

[6] E. H. Adelson and J. R. Bergen, "Spatiotemporal energy models for the perception of motion," *J. Opt. Soc. Amer. A,* vol. 2, no. 2, pp. 284–299, Feb. 1985.

[7] D. J. Heeger, "Optical flow from spatio-temporal filters," in *Proc. 1st Int. Conf. Computer Vision,* London, U.K., June 1987, pp. 181–190.

[8] H. Knutsson, "Representing local structure using tensors," in *Proc. 6th Scandinavian Conf. Image Anal.,* M. Pietikäinen and J. Röning, Eds. Oulu, Finland: Pattern Recognition Soc. Finland, 1989, pp. 244–251.

[9] D. J. Fleet and A. D. Jepson, "Computation of component image velocity from local phase information," *Int. J. Comput. Vis.,* vol. 5, no. 1, pp. 77–104, 1990.

[10] D. J. Fleet, *Measurement of Image Velocity.* Norwell, MA: Kluwer, 1992

[11] D. Marr and S. Ullman, "Directional selectivity and its use in early visual processing," in *Proc. R. Soc. London,* 1981, vol. 211, pp. 151–180.

[12] P. Anandan, "A computational framework and an algorithm for the measurement of visual motion," *Int. J. Comput. Vis.,* vol. 2, pp. 283–310, 1989.

[13] L. Haglund, "Adaptive multidimensional filtering," Ph.D. dissertation, Dept. Elect. Eng., Linköping Univ., Linköping, Sweden, p. 158, 1992, p. 158.

[14] D. Terzopoulos, "Regularization of inverse visual problems involving discontinuities," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. PAMI-8, pp. 413–424, July 1986.

[15] J. Marroquin, S. Mitter, and T. Poggio, "Probabilistic solution of ill-posed problems in computational vision," *J. Amer. Stat. Assoc.,* vol. 82, pp. 76–89, Mar. 1987.

[16] J. Konrad and E. Dubois, "Bayesian estimation of motion vector fields," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. 14, pp. 910–927, Sept. 1992.

[17] J. J. Simpson and J. I. Gobat, "Robust velocity estimates, stream functions, and simulated lagrangian drifters from sequential spacecraft data," *IEEE Trans. Geosci. Remote Sensing,* vol. 32, pp. 479–493, May 1994.

[18] D. D. Wahl and J. J. Simpson, "Physical processes affecting the objective determination of near-surface velocity from satellite data," *J. Geophys. Res.,* vol. 95, pp. 13 511–16 528, 1990.

[19] D. Gabor, "Theory of communication," *J. Inst. Elect. Eng.,* vol. 93, pp. 429–457, 1946.

[20] D. J. Fleet and K. Langley, "Computational analysis of non-Fourier motion," Robotics Perception Lab., Dept. Comput. Inform. Sci., Queens Univ., Kingston, Ont., Canada, Tech. Rep. RPL-TR-9309, p. 45, 1993.

[21] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis,* 2nd ed. New York: Wiley, 1984.

[22] E. C. Hildreth, "Computations underlying the measurement of visual motion," *Artif. Intell.,* vol. 23, pp. 309–354, 1984.

[23] H.-H. Nagel, "On the estimation of optical flow: Relations between different approaches and some new results," *Artif. Intell.,* vol. 33, p. 299, 1987.

[24] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. PAMI-6, pp. 721–741, 1984.

[25] J. Besag, "Toward Bayesian image analysis," *J. Appl. Stat.,* vol. 16, no. 3, pp. 395–407, 1989.

[26] ——, "On the statistical analysis of dirty pictures," *J. Roy. Stat. Soc. B,* vol. 48, no. 3, pp. 259–302, 1986.

[27] P. Bouthemy and E. François, "Motion segmentation and qualitative dynamic scene analysis from an image sequence," *Int. J. Comput. Vis.,* vol. 10, no. 2, pp. 157–182, Apr. 1993.

[28] P. J. Burt, "Fast filter transforms for image processing," *Comput. Graph. Image Process.,* vol. 16, pp. 20–51, 1981.

[29] A. P. Cracknell and L. W. B. Hayes, *Introduction to Remote Sensing.* London, U.K.: Taylor & Francis, 1991.

[30] J. J. Bergen, P. Anandan, K. J. Hanna, and R. Hingorani, "Hierarchical model-based motion estimation," in *Proc. Second Europ. Conf. Comput. Vision,* 1992, pp. 237–252.

[31] B. Jähne, *Digital Image Processing.* Berlin, Germany: Springer-Verlag, 1991.

[32] R. Larsen, J. D. Hansen, K. Conradsen, and B. K. Ersbøll, "Temporal interpolation in Meteosat images," Inst. Math. Modeling, Tech. Univ. Denmark, Tech. Rep., 1991.

[33] J. J. Simpson and C. Humphrey, "An automated cloud screening algorithm for daytime AVHRR imagery," *J. Geophys. Res.,* vol. 96, pp. 13 459–13 481, 1990.

[34] T. C. Gallaudet and J. J. Simpson, "Automated cloud screening of AVHRR imagery using split-and-merge clustering," *Remote Sens. Environ.,* vol. 38, pp. 77–121, 1991.

[35] J. L. Barron, D. J. Fleet, and S. S. Beauchemin, "Performance of optical flow techniques," *Int. J. Comput. Vis.,* vol. 12, pp. 43–77, 1994.

Rasmus Larsen was born in Brørup, Denmark, on March 14, 1966. He received the M.Sc. (Eng.) degree in 1991 and the Ph.D. degree in statistical image analysis from the Technical University of Denmark, Lyngby, in 1994.

In 1995, he was a visiting scholar at Scripps Institution of Oceanography, University of California, San Diego. He is now an Assistant Research Professor at the Department of Mathematical Modeling, Technical University of Denmark. His research interests include estimation of fluid flow, statistical classification, and transformation of multivariate multitemporal data sets.

Knut Conradsen was born in Aabenraa, Denmark, on August 7, 1945. He received the M.Sc (Eng.) degree in mathematics from the University of Copenhagen, Copenhagen, Denmark, in 1970.

Since 1970, he has worked at The Institute of Mathematical Statistics and Operations Research (since 1993 has been a part of the Department of Mathematical Modeling, Technical University of Denmark). He was an Assistant Professor until 1974, an Associate Professor until 1990, a Reader from 1990 to 1992, and has since been a Full Professor of Statistical Image Analysis. Presently, he is also Vice-Rector. Since 1988, he has been Corresponding Editor of the *Bulletin of the Institute of Mathematical Statistics,* and since 1993, he has been Associate Editor of *Pattern Recognition.*

Bjarne Kjær Ersbøll was born in Hornsyld, Denmark, on June 9, 1957. He received the M.Sc. (Eng.) degree from the Technical University of Denmark, Lyngby, in 1983, and the Ph.D. degree in statistical image analysis from the Technical University of Denmark in 1990.

From 1983 to 1988, he worked at The Institute of Mathematical Statistics and Operations Research (since 1993, he was a part of the Department of Mathematical Modeling, Technical University of Denmark) as a Research Assistant, as an Assistant Professor from 1988 to 1992, and since 1992, as an Associate Professor in Mutivariate Statistics. During his professional career, he has worked with aspects of using statistics as a tool for solving image analysis problems, some examples being remote sensing, automatic analysis of fingerprints, and industrial quality control.