IMM
DEPARTMENT OF MATHEMATICAL MODELLING

Technical University of Denmark
DK-2800 Lyngby – Denmark

J. No. FFF0
4/1/1995
RL

ESTIMATION OF DENSE IMAGE FLOW FIELDS IN FLUIDS

Rasmus Larsen, Knut Conradsen and Bjarne Kjær Ersbøll

TECHNICAL REPORT

IMM-REP-95-04

IMM

ESTIMATION OF DENSE IMAGE FLOW FIELDS IN FLUIDS

Rasmus Larsen, Knut Conradsen and Bjarne Kjær Ersbøll

TECHNICAL REPORT

IMM-REP-95-04

IMM

Estimation of Dense Image Flow Fields in Fluids

Rasmus Larsen,* Knut Conradsen and Bjarne Kjær Ersbøll
Institute of Mathematical Modelling
Technical University of Denmark Building 321, DK–2800 Lyngby

Please address correspondence to Rasmus Larsen.

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 non blurred temporal interpolation can be performed only when the flow field or an estimate there-of 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 article 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 using a set of 3-D spatio-temporal filters. The estimated local energy distribution also allows us to compute a certainty measure of the estimated local flow. The algorithm furhtermore utilizes Markovian random fields in order to incorporate smoothness across the field. To obtain smothness we will constrain first as well as second 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.

* E–mail: rl@imm.dtu.dk Phone: +45 4588 1433 FAX: +45 4588 1397
WWW: http://www.imm.dtu.dk/documents/users/rl/homepage.html

1 Introduction

Independently moving objects, rotation, dilation, 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 (Wahl & Simpson, 1990), differential methods (Horn & Schunk, 1981; Nagel & Enkelmann, 1986), energy-based methods (Adelson & Bergen, 1985; Heeger, 1987; Knutsson, 1989) and phase-based methods (Fleet & Jepson, 1990; Fleet, 1992).

Restricting measurements to small spatio-temporal neighborhoods, however, often results in the measurements being based on one-dimensional 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 (Marr & Ullman, 1981).

In general this problem exists in neighborhoods of the image sequence that have a one dimensional 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 spatial structures, such as intensity corners or various textured regions, we can reliably extract the true velocity. In order to identify the type of neighbourhood several approaches have been reported. In a correlation-based approach Anandan (1989) used the curvature of the match surface, Nagel & Enkelmann (1986) investigated the use of second order spatial derivatives to identify the neighbourhood. The use of quadrature type filters tuned to different spatio-temporal frequencies has been used to identify the type of the neighbourhood in energy- as well as phase-based approaches (Haglund, 1992; Fleet, 1992)

Because the aperture problem results in flow fields that are not fully constrained an assumption of 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 force the spatial derivatives to be small. These restictions are referred to as smoothness constraints (Horn & Schunk, 1981). Methods utilizing this type of smoothness constraints include the work of (Nagel & Enkelmann, 1986; Terzopoulos, 1986; Marroquin, Mitter, & Poggio, 1987; Konrad & Dubois, 1992). Other approaches based on spatial filtering also have been reported. Simpson & Gobat (1994) 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 3.1 we will consider a method to extract velocity estimates as well as related directional certainties of the estimates based on the local energy distribution. This technique also will allow for the detection of unreliable results due to deviations from the assumption of coherent translational motion. This is especially important because of the distortion of features due to physical processes that can occur in fluids (e.g. Wahl & Simpson (1990)).

In Section 3.2 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 and second 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

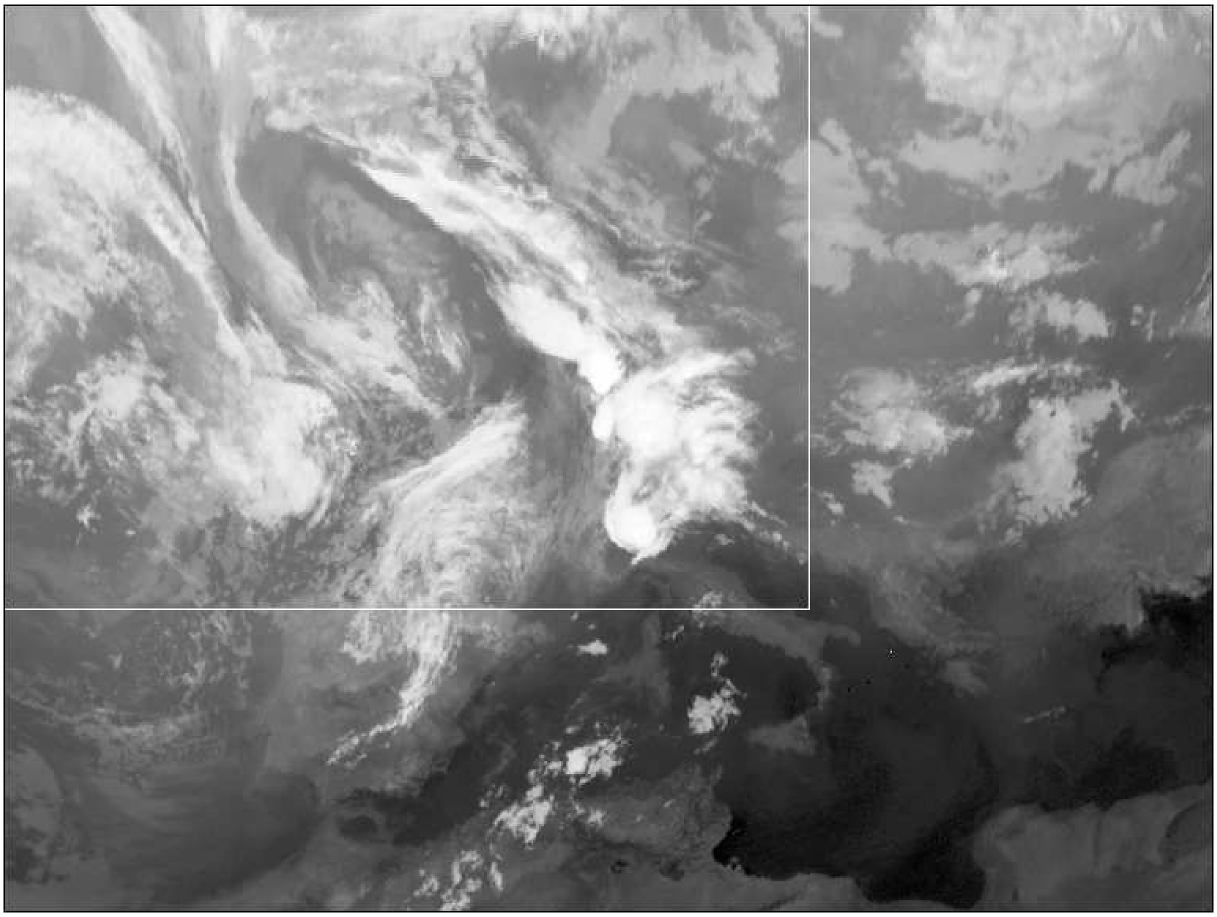

Figure 1: This is the infrared channel from the Meteosat 5 satellite mapped to a polar stereographic projection showing the cloud cover over Europe at August 24, 1994, 05:00 GMT. We will apply the flow estimation algorithms to the subimage marked here. The images a provided by the Danish Meteorological Institute.

the velocity field. Finally, we will combine the prior distribution and the observation model into a posterior distribution using Bayes' theorem.

The algorithm is illustrated by the estimation of smooth cloud velocity fields in a series of Meteosat satellite images. The estimated flow fields will be used in two temporal interpolation schemes. First, intermediate frames between each of the originals are generated, second, the estimated flow field is used to generate a frame to replace a drop-out in the sequence.

2 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 consist of mapping to a polar stereographic projection, interpolation to a equirectangular 7 km grid, and subsection to 576×768 images. 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 minutes. We will be using a 384×512 subsequence centered at August 24, 1994, 05:00 GMT. The center image of the sequence is shown in Figure 1.

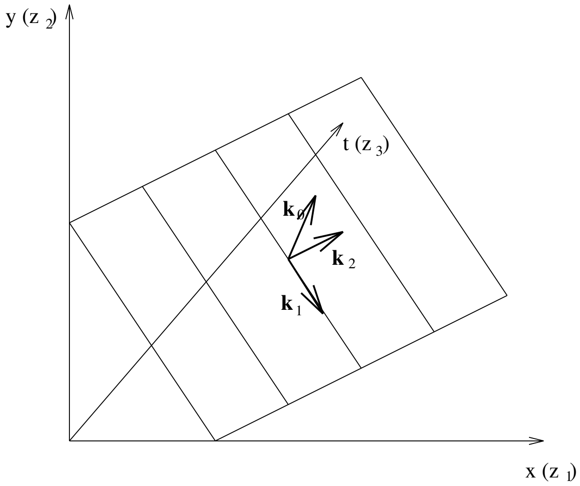

Figure 2: $\boldsymbol{k}$ is the unit normal vector of the iso-grey level planes generated by a translating line. $\boldsymbol{k}_1$ is the unit direction vector of the translating lines, and is therefore perpendicular to $\boldsymbol{k}$. $\boldsymbol{k}_2$ is a vector perpendicular to $\boldsymbol{k}$ as well as $\boldsymbol{k}_1$. The normal flow is found as the projection onto the (z_1, z_2) of $\boldsymbol{k}_2$ multiplied so its temporal coordinate equals 1.

3 Methods

This section is divided into two subsections. In Section 3.1 we will consider the local velocity estimation using a set of spatio-temporal directional quadrature filter pairs. After which in Section 3.2 we will formulate an algorithm for integrating these local estimates to a dense velocity field using smoothness constraints based on first and second order spatial derivatives of the velocity field.

3.1 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 neighbourhood containing a one dimensional intensity structure (e.g., a line) that translates coherently through time. In the spatio-temporal domain this corresponds to a neighbourhood of iso-grey level planes. Let these planes be given by their unit normal vector $\hat{\boldsymbol{k}} = (k_1, k_2, k_3)^T$. We will refer to this vector as the spatio-temporal orientation vector. The non-zero Fourier coefficents of this neighbourhood are concentrated to the line defined by $\boldsymbol{k}$. Furthermore the correspondence between this vector and the normal flow of the intensity structure is illustrated in Figure 2. The normal flow is given by

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

Now, in order to estimate this line, we will sample the Fourier domain using a set of spatio-temporal filters. Using a method developed by Knutsson (1989) 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.

In order to sample the Fourier domain we will employ a set of Gabor filters (Gabor, 1946) tuned to frequencies distributed evenly across all spatio-temporal orientations, i.e. the center frequencies of the filters are the vertices of a diametrical symmetric regular polyhedron (Knutsson, 1989). 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. By dividing the filter into an odd and an even part we get the two transfer functions

$$Q_p^e(\boldsymbol{k}) = \frac{1}{(2\pi)^3}\left[\exp\left(\frac{(\boldsymbol{k}-\boldsymbol{k}_p)^2\sigma^2}{2}\right) + \exp\left(\frac{(\boldsymbol{k}+\boldsymbol{k}_p)^2\sigma^2}{2}\right)\right] \tag{2}$$

$$Q_p^o(\boldsymbol{k}) = \frac{1}{(2\pi)^3}\left[\exp\left(\frac{(\boldsymbol{k}-\boldsymbol{k}_p)^2\sigma^2}{2}\right) - \exp\left(\frac{(\boldsymbol{k}+\boldsymbol{k}_p)^2\sigma^2}{2}\right)\right] \tag{3}$$

The corresponding convolution masks are easily computed by taking the Fourier transforms

$$q_p^e(\boldsymbol{z}) = \frac{1}{(2\pi)^{3/2}\sigma^3}\cos(\boldsymbol{k}_p^T\boldsymbol{z})\exp\left(-\frac{\boldsymbol{z}^2}{2\sigma^2}\right) \tag{4}$$

$$q_p^o(\boldsymbol{z}) = \frac{1}{(2\pi)^{3/2}\sigma^3}\sin(\boldsymbol{k}_p^T\boldsymbol{z})\exp\left(-\frac{\boldsymbol{z}^2}{2\sigma^2}\right) \tag{5}$$

Although Gabor filters are not quadrature pairs, they provide a reasonably good approximation for sufficiently small bandwidths (Fleet & Langley, 1993). If the bandwidth in octaves is measured at one standard deviation of the Gaussian envelope it is given by

$$B = \log_2\left[\frac{||\boldsymbol{k}|| + \sigma}{||\boldsymbol{k}|| - \sigma}\right] \tag{6}$$

Finally, the energy distribution of the Fourier domain as estimated by the set of quadrature filter pairs may be represented by the tensor

$$\boldsymbol{T} = \sum_p q_p \boldsymbol{n}_p \boldsymbol{n}_p^T \tag{7}$$

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 $\boldsymbol{T}$ (Anderson, 1984). So for the coherently translating one-dimensional intensity structure, which has an effectively one dimensional Fourier domain, the spatio-temporal orientation vector is found by an eigen analysis of $\boldsymbol{T}$. Because the Fourier domain is one dimensional $\boldsymbol{T}$ has only one non-zero eigenvalue.

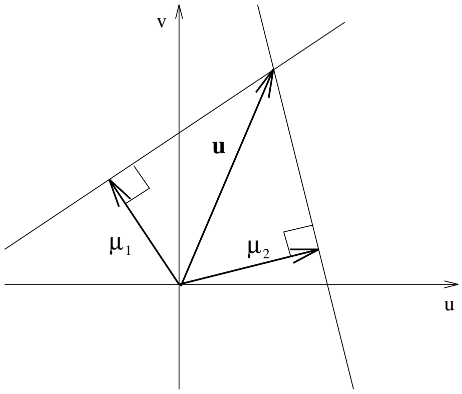

Figure 3: The true flow, $\boldsymbol{u}$, is constrained by the linear constraints given by the two normal flows, $\boldsymbol{\mu}_1$, and $\boldsymbol{\mu}_2$.

Now, if the translating structure has a two-dimensional intensity structure (e.g., a grey level corner) the spatio-temporal domain is described by two spatio-temporal orientations, each of which give rise to a non-zero eigenvalue of $\boldsymbol{T}$. The eigenvectors corresponding to these non-zero eigenvalues each translates into a normal flow by using Equation (1). These two normal flows each constrain the true flow in one direction as can be seen in Figure 3.

We can furthermore determine the perpendicular distance of the true flow to either of the constraint lines, this is given by

$$d_k(x_i, y_i) = \|(\boldsymbol{u}(x_i, y_i) - \boldsymbol{\mu}_k(x_i, y_i))^T \cdot \frac{\boldsymbol{\mu}_k(x_i, y_i)}{\|\boldsymbol{\mu}_k(x_i, y_i)\|}\|, \qquad k = 1, 2 \tag{8}$$

where $\boldsymbol{u}(x_i, y_i)$, $\boldsymbol{\mu}_k(x_i, y_i)$ are the true flow and the 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 the velocity field.

Deviation from the assumption of coherent translation, imperfectly designed filters and noise will result in non-zero Fourier coefficients not being contained in one single line or plane. In this case all three of the tensor eigenvalues will be non-zero. This allows us to extract information of 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 non-zero third eigenvalue we suggest using a confidence measure for each of the linear constraints based on the difference of the corresponding eigenvalue and the least 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.

$$w_k(x_i, y_i) = \frac{\lambda_k(x_i, y_i) - \lambda_3(x_i, y_i)}{\lambda_k(x_i, y_i)}, \quad k = 1, 2 \tag{9}$$

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 least eigenvalue approaches zero, and it attains its maximum value of one, when the least eigenvalue is zero, and the corresponding eigenvalue is the largest, or is equal to the largest eigenvalue, respectively.

3.2 Integration of local measurements

As mentioned in the Introduction 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 & Schunk (1981)'s original paper this has been investigated be several authors (e.g. (Hildreth, 1984; Nagel, 1987). One way of formulating such a smoothness constraint is by use of markovian random fields (Geman & Geman, 1984; Konrad & Dubois, 1992). We do this using the Bayesian paradigm (Besag, 1989). 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 these finite differences

$$\begin{aligned}
\boldsymbol{u}_x(x_i, y_i) &= \boldsymbol{u}(x_i{+}1, y_i) - \boldsymbol{u}(x_i, y_i) \\
\boldsymbol{u}_y(x_i, y_i) &= \boldsymbol{u}(x_i, y_i{+}1) - \boldsymbol{u}(x_i, y_i) \\
\boldsymbol{u}_{xx}(x_i, y_i) &= \boldsymbol{u}(x_i{+}1, y_i) - 2\boldsymbol{u}(x_i, y_i) + \boldsymbol{u}(x_i{-}1, y_i) \\
\boldsymbol{u}_{xy}(x_i, y_i) &= \boldsymbol{u}(x_i{+}1, y_i{+}1) - \boldsymbol{u}(x_i, y_i{+}1) - \boldsymbol{u}(x_i{+}1, y_i) + \boldsymbol{u}(x_i, y_i) \\
\boldsymbol{u}_{yy}(x_i, y_i) &= \boldsymbol{u}(x_i, y_i{+}1) - 2\boldsymbol{u}(x_i, y_i) + \boldsymbol{u}(x_i, y_i{-}1)
\end{aligned} \tag{10}$$

where (x_i, y_i), $i \in \{0, 1, \ldots, N-1\}$ are the pixel positions, then the prior distribution of the flow field may be described by a Gibbs distribution $p(\{\boldsymbol{u}\}) = \frac{1}{Z}\exp(-\beta_1 U_1 - \beta_2 U_2)$, where Z is a normalization constant and the two energy terms are given by

$$U_1 = \sum_{i=0}^{N-1} \|\boldsymbol{u}_x(x_i, y_i)\|^2 + \sum_{i=0}^{N-1} \|\boldsymbol{u}_y(x_i, y_i)\|^2. \tag{11}$$

$$U_2 = \sum_{i=0}^{N-1} \|\boldsymbol{u}_{xx}(x_i, y_i)\|^2 + \sum_{i=0}^{N-1} 2\|\boldsymbol{u}_{xy}(x_i, y_i)\|^2 + \sum_{i=0}^{N-1} \|\boldsymbol{u}_{yy}(x_i, y_i)\|^2, \tag{12}$$

where $\|\cdot\|$ is the Euclidean norm. This probability distribution assigns high probability to fields that exhibit small derivatives and low probability to field with high spatial derivatives. We will need a first order neighbourhood to implement U_1, and a third order neighbourhood to implement U_2. In Figures 4(a) and 4(b) these neighbourhoods are shown. In Figures 4(c) and 4(d) the cliques necessary to implement the energy functions are shown.

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

$$P(\boldsymbol{y}|\boldsymbol{u}) = \frac{1}{Z}\exp(-\alpha U_0) = \frac{1}{Z}\exp(-\alpha \sum_{i=0}^{N-1} \sum_{k=1}^{2} w_k(x_i, y_i) d_k(x_i, y_i)^2) \tag{13}$$

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 Equation (8). $w_k(x_i, y_i)$ is the certainty measure corresponding to this normal flow given by Equation (9). Z is a normalization constant. By using a Gibbs energy function that punishes

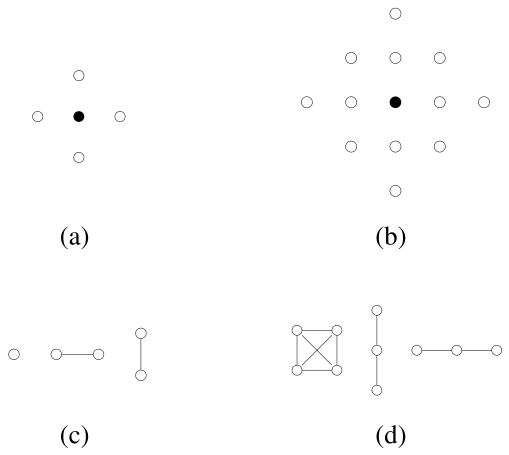

Figure 4: Neighbourhood systems and cliques for the pixel process. The neighbourhood configurations for (interior) points for (a) a first order and (b) a third order Markovian random field, respectively. The points marked with the symbol ∘ are the neighbors of the point marked •. (c) The cliques corresponding to the neighbourhood configuration in (a). (d) Some of the additional cliques corresponding to the neighbourhood configurations in (b). This neighbourhoods contains 24 cliques in all.

larger deviation 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 flow is punished. Furthermore the use of the confidence measures w_k derived in the previous section as weights allow us to take into consideration the quality of our measurements.

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

$$\begin{aligned} U \; = \; & \alpha \sum_{i=0}^{N-1} \sum_{k=1}^{2} w_k(x_i, y_i) d_k(x_i, y_i)^2 \\ & + \beta_1 \sum_{i=0}^{N-1} \left(\|\boldsymbol{u}_x(x_i, y_i)\|^2 + \|\boldsymbol{u}_y(x_i, y_i)\|^2 \right) \\ & + \beta_2 \sum_{i=0}^{N-1} \left(\|\boldsymbol{u}_{xx}(x_i, y_i)\|^2 + 2\|\boldsymbol{u}_{xy}(x_i, y_i)\|^2 + \|\boldsymbol{u}_{yy}(x_i, y_i)\|^2 \right) \end{aligned} \tag{14}$$

In this energy function we can control the properties of estimated motion field. The smoothness is controlled by β_1, and β_2, 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. This has been implemented using the Iterated Conditional Modes scheme by Besag (1986).

Filter no.	x	y	t
1	a	0	b
2	-a	0	b
3	b	a	0
4	b	-a	0
5	0	b	a
6	0	b	-a
7	-c	c	c
8	-c	c	-c
9	-c	-c	c
10	c	c	c

Table 1: The spatio-temporal directions of the Gabor filters are given by these coefficients, where $a = 2/(10+2\sqrt{5})$, $b = (1+\sqrt{5})/(10+2\sqrt{5})$, and $c = 1\sqrt{3}$.

4 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 using the proposed prior distribution. Secondly we will use an estimated flow field to perform a temporal interpolation and thirdly we will investigate how the estimated flow field may be used in order to generate replacement frames in the case of drop-outs in an image sequence of a time varying scene.

4.1 Estimation of the Flow Field

As described in Section 3.1 we will describe the local Fourier domain by the tensor given in Equation (7). We will use a set of ten Gabor filter pairs to sample the Fourier domain. The spatio-temporal directions of the Gabor filters are shown in table 1 (Haglund, 1992).

By setting the standard deviation of the Gaussian envelope equal to 2.5 and truncating the envelope at three standard deviations we arrive at the filter size, $17 \times 17 \times 17$. Following Fleet & Langley (1993) we set the bandwidth measured at one standard deviation to 0.8 octaves. This results in a spatio-temporal frequency of the filters of 1.48 from Equation (6). Since the Gabor filter kernels are separable (Heeger, 1987) the 20 3-D convolutions may be performed by 140 1-D convolutions. Thus reducing the computational load by a factor 40.

With Gabor filtering and the computation of normal flows using Equation (1) on the 05:00 GMT image of the Meteosat sequence we obtain the normal flows corresponding to the largest eigenvalues as can be seen in Figure 5. The computations are carried out on a spatially subsampled version of the image sequence. In order to remove unreliable measurements we have applied two threshholds. First, neighbourhoods of non-translational velocity are removed by only considering normal flows whose corresponding eigenvalue is at least twice as large as the smallest eigenvalues, i.e. $w_k > 0.5$, secondly by demanding that the local energy, i.e. the sum of the eigenvalues, exceeds 10 % of the maximum energy taking over the entire image we remove measurements of problematic signal-to-noise ratio.

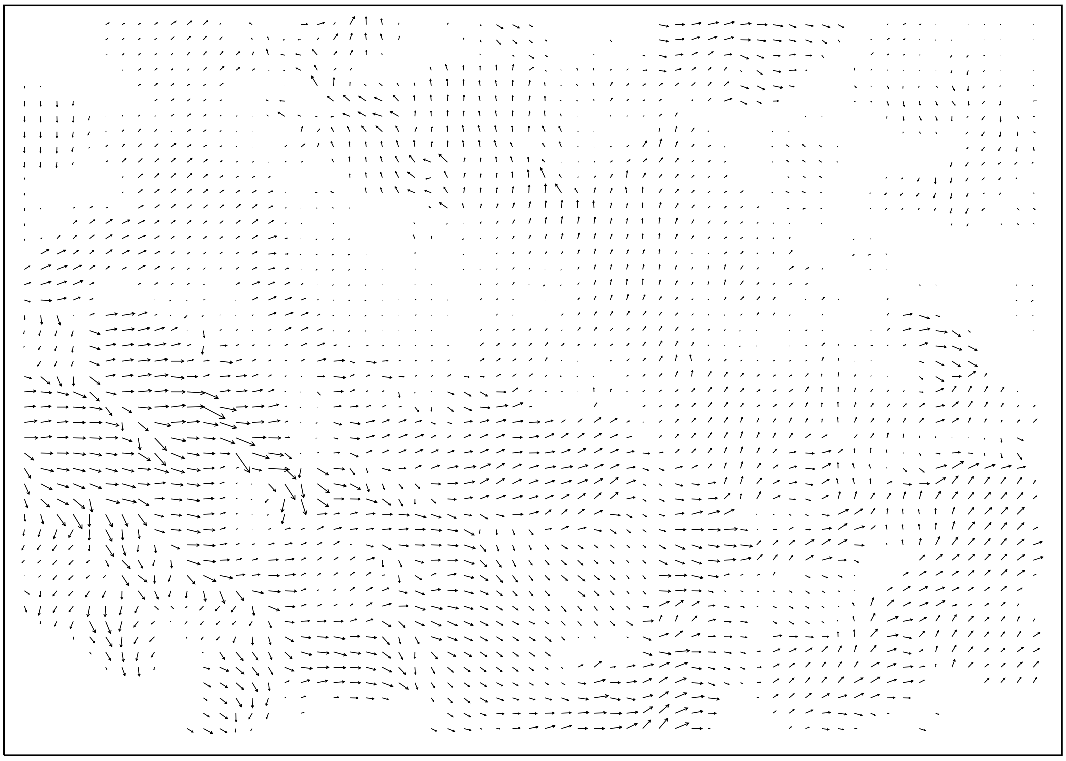

Figure 5: These are the normal flows computed from the Meteosat sequence corresponding to the 5:00 GMT frame.

α	β_1	β_2	U_0	U_1	U_2
0.999	0.001	0.000	0.0005267	280.1	0
0.100	0.900	0.000	3.171	213.9	0
0.999	0.000	0.001	0.004539	0	430.1
0.999	0.0005	0.0005	0.001917	525.8	296.2

Table 2: Energies

It is evident from Figure 5 that the estimated normal flow field is unsuited as an estimate of the velocity field, and unsuited as a means of performing non-blurring temporal interpolation. The normal flows are dependent on the orientation of the underlying image structure. To get a proper estimate of the velocity field, we will use the integration technique described in Section 3.2.

The parameters α, β_1, and β_2 of Equation 15 allows us to control the final velocity field. By increasing α while keeping the β's the same, we would expect the velocity field to become more rough, because this corresponds to increasing our faith in the measurements. On the other hand by increasing the β's while maintaing the same α, we would expect a smoother field. Also, a larger β_2 should allow for a more rapidly changing field as opposed to a larger β_1, because β_1 tends to force small first order derivatives whereas β_2 tends to force the second order derivatives to be small.

In order to illustrate this, we have estimated a series of flow fields using different parameter settings. In Figure 6 the flow field corresponding to $\alpha = 0.999$, $\beta_1 = 0.001$, and $\beta_2 = 0$ is shown. We can see that we have captured the motion as perceived from the image in Figure 1. This is especially clear in the vicinity of the rotational pattern up and left from the image center. It may also be noted that although the contours in this area are almost perpendicular to the

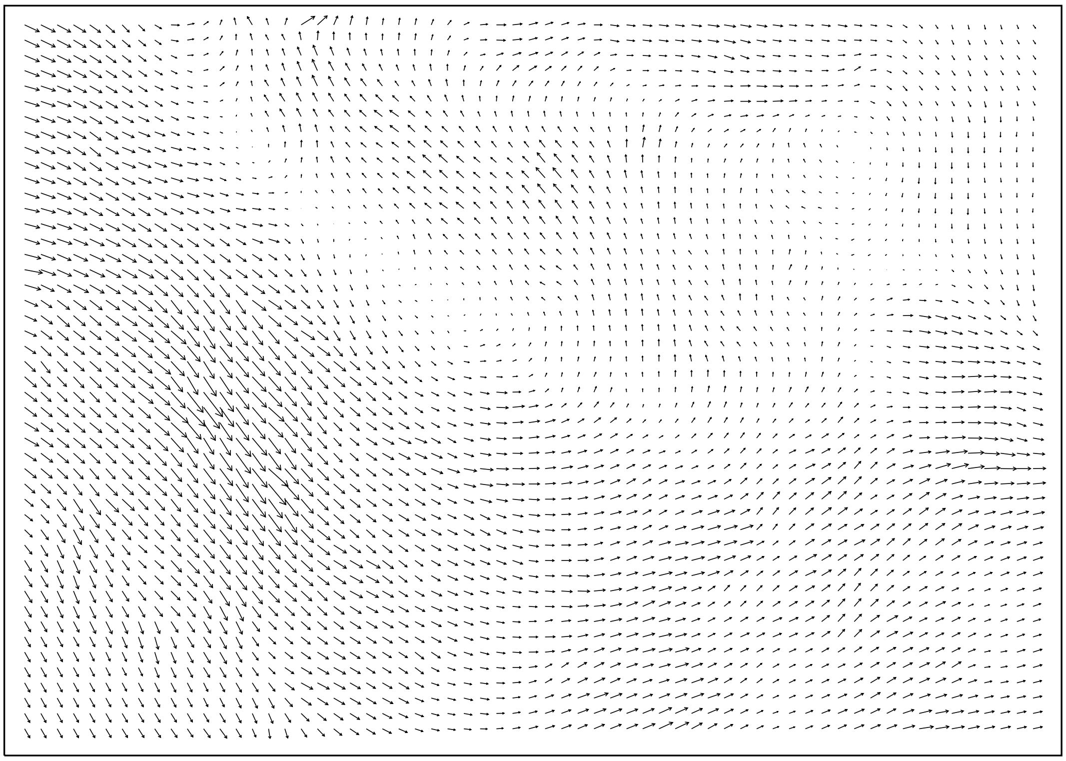

Figure 6: This is the velocity field is computed from the Meteosat sequence corresponding to the 5:00 GMT frame using $\alpha = 0.999$, $\beta_1 = 0.001$, and $\beta_2 = 0$.

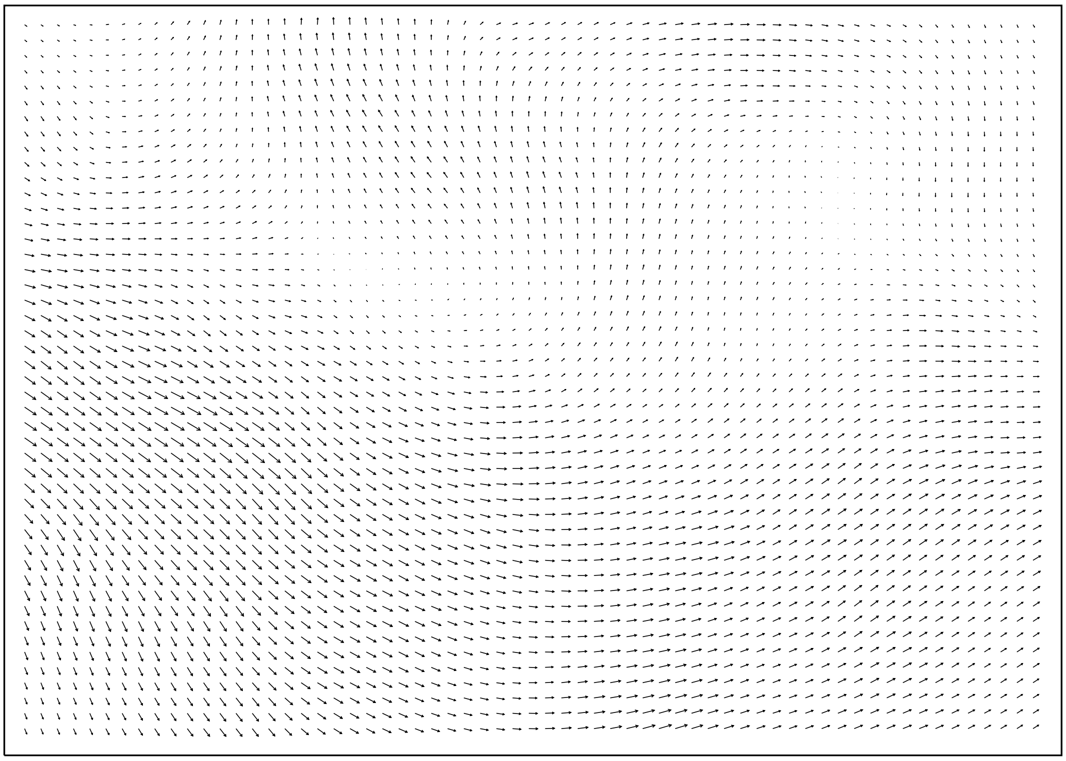

Figure 7: Analogous to Figure 6, except $\alpha = 0.1$, $\beta_1 = 0.9$, and $\beta_2 = 0$.

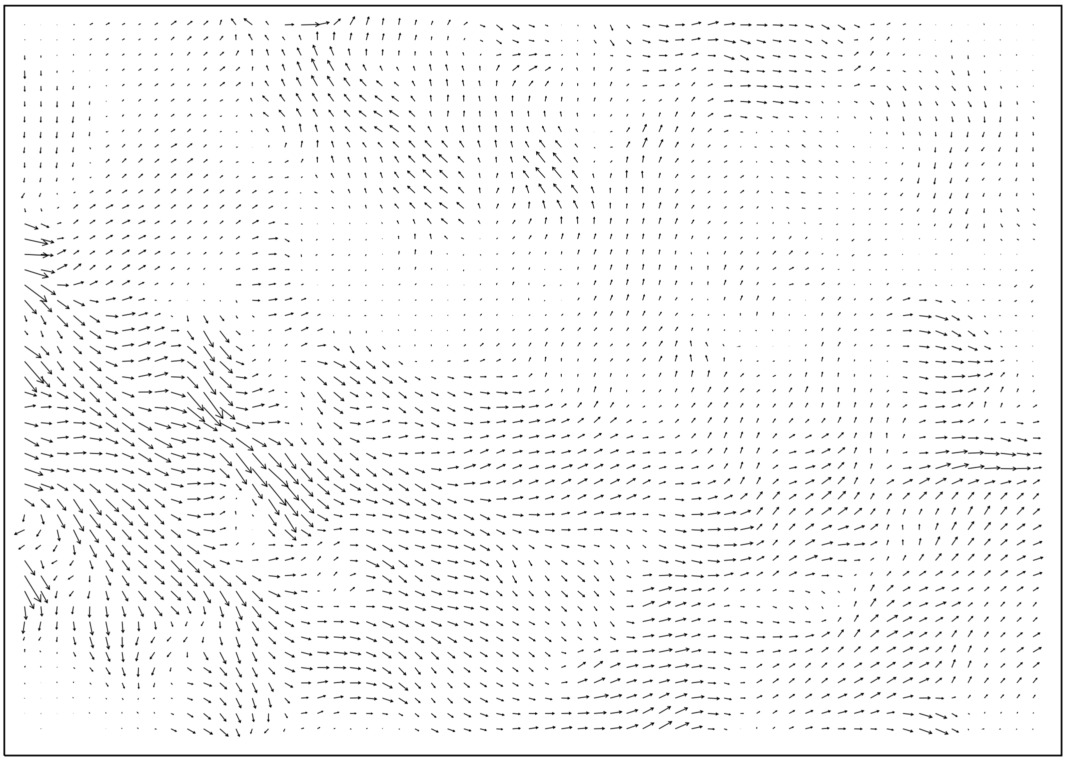

Figure 8: Analogous to Figure 6, except $\alpha = 0.999$, $\beta_1 = 0$, and $\beta_2 = 0.001$.

direction of motion, resulting in very small normal flows, see Figure 5, the estimated flow still captures the true motion. By computing the observation energy function (Equation (13)) of the estimated flow we can see to how high a degree we have violated the observations. These energies are tabulated in Table 2. For this flow field $U_0 = 0.0005267$. This is very small compared to the average estimated speed, which is 1.1 So the smoothing of the field is primarily done within the constraints given by the normal flows.

As expected, the estimated flow field becomes smoother as the $beta$ to α ratio increases. In Figure 7 we can see the flow field resulting from setting $\alpha = 0.1$, $\beta_1 = 0.9$, and $\beta_2 = 0$. This field is considerably smoother than the the one from the previous case. It may also be noted that U_0 has increased to 3.171, i.e. the observations are violated to higher degree.

In Figure 8 we have tried to estimate the flow field by resticting the second order derivatives only. The parameter settings are $\alpha = 0.999$, $\beta_1 = 0$, and $\beta_2 = 0.001$. We can see that restricting the second order derivatives allow for to much variation in the field. The algorithm has not been able to fill in the areas where no observations are available (e.g., the lower left corner), and the rotational pattern is not reconstructed.

4.2 Temporal Interpolation in a Meteosat Satellite Sequence

An obvious application of the the flow field estimation temporal interpolation. We will illustrate how this may done using the flow field from Figure 6. We will generate an intermediate images between the the original images recordes at 5:00 GMT and 5:30 GMT. We extrapolate from the each of the two original images using the corresponding flow field estimates. As we put an emphasis on not blurring the image, we will use a nearest neighbour interpolation scheme, as

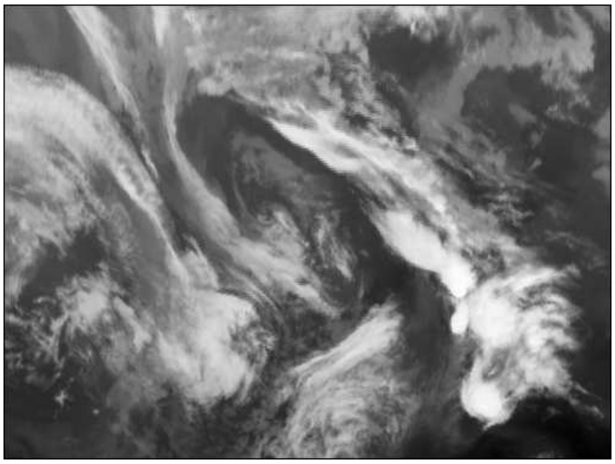

Figure 9: This is the artificially generated intermediate image corresponding to 5:15 GMT.

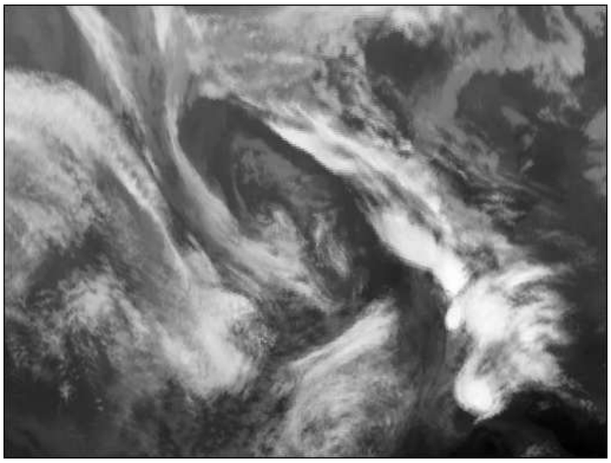

Figure 10: This is the simple average of the two original images recorded at 5:00 and 5:30 GMT. Note the echoing of all the moving contours.

opposed to for instance bilinear or cubic interpolation, to sample the extrapolated images to the original grid. Finally, we take the average of the two extrapolations. This is our intermediate image.

In Figure 9 the interpolated images is shown. A naïve approach to an interpolation would be to simply take the average of the two originals, the result of this is shown in Figure 10. 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 saticfactory temporal interpolation.

4.3 Drop-out Replacement by Temporal Interpolation

At least twice every day we have drop-outs in the Metosat 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 drop-outs. First we need to estimate the flow field at the time of the missing image. When performing the filtering operation we will replace the missing image 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 5:00 GMT image from the sequence, substituted it with the average of the 4:30 and 5:30 GMT images and

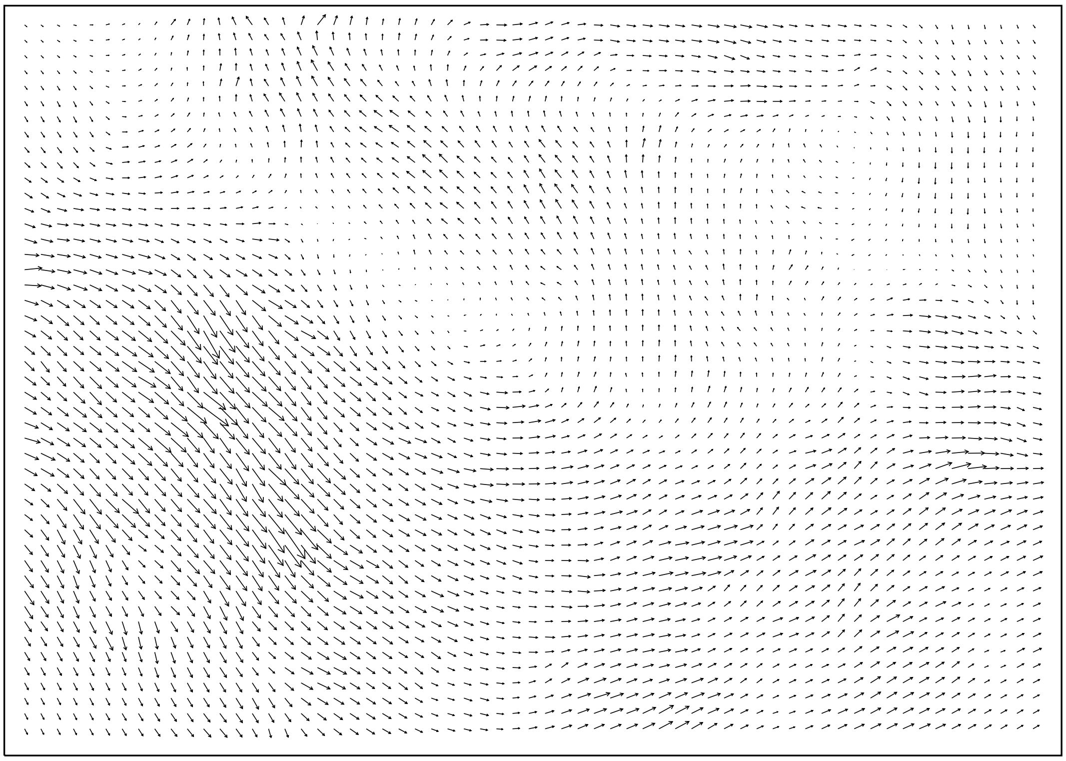

Figure 11: This is the velocity field is computed from the Meteosat sequence corresponding to the 5:00 GMT frame, estimated using the average of the 4:30 and the 5:30 GMt images in the place of the 5:00 image.

computed the flow field. The resulting field is shown in Figure 11. When compared to the flow field estimated using the original 5:00 GMT frame in Figure 6 changes are obvious only in the upper left corner, where the measurements are sparse, as can be seen in Figure 5. Now we can generate the replacement 5:00 image by warping the 4:30 image according to this flow field.

In Figure 12(a) and Figure 12(b) the replacement and the original 5:00 GMT images are shown. The replacement corresponds well to the original. When viewing the new sequence as a film the replaced images can not be distinguished from the originals.

5 Discussion

5.1 Data

The purpose of the exercises described above have been to perform a non-blurred temporal interpolation of the sequence. It should be noted that the extracted flow fields not necessarily assimilate the corresponding wind fields. This is true because only some clouds move with the wind. Another problem is that we are only observing the top layer of the clouds, and the winds may vary with height (Cracknell & Hayes, 1991).

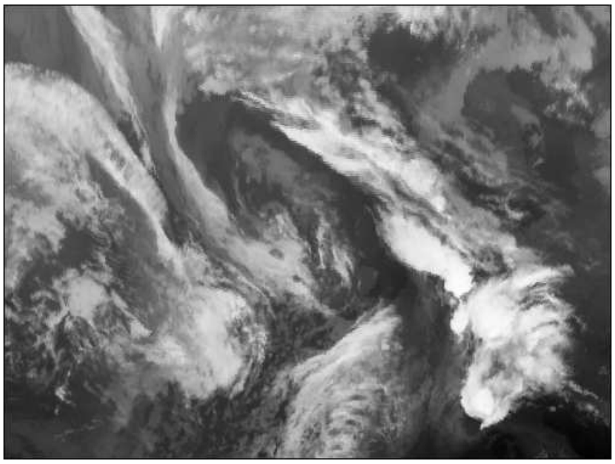

(a)

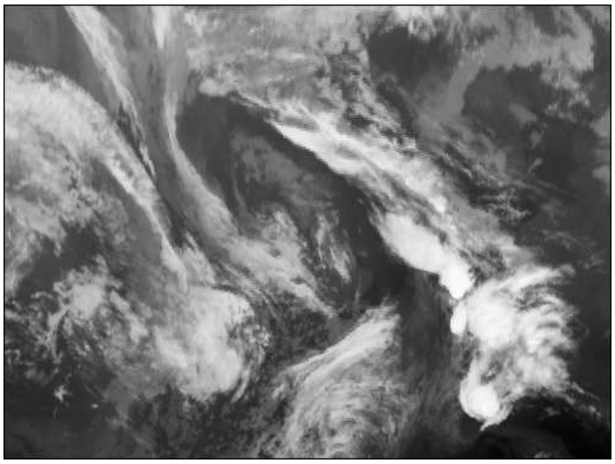

(b)

Figure 12: (a) The replacement image and (b) the original 5:00 GMT image

5.2 Estimation of Normal Flows

The spatio-temporal Gabor filtering described in Section 3.1 provides a robust estimation of normal flows. In addition to identifying the zero, one or two normal flows at each pizel, it provides us with a measure of the confidence we should attach to the corresponding normal flow. It is, however, only sensitive to spatio -temporal frequencies within a certain range. So for a specific set of filters the scale at which the computations are made should be carefully chosen from knowledge of the range of the velocities , and thus the spatio-temporal frequencies. In the examples shown above the filtering where carried out on the second level of a Gaussian pyramid (Burt, 1981), corresponding to a subsampling of a factor four.

To overcome this problem simultaneous filtering on a number of levels in a pyramid using the same set of Gabor filters will allow us to measure velocities over a larger range. Preliminary studies of this has been carried out by Heeger (1987).

5.3 Integration using Markovian Random Fields

The Markovian random field approach to integrate the measured normal flows has proven a powerful and succesful technique. We have used restictions of the first and second 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. This has not been possible by restricting the second order derivatives only.

The smoothness constraints have been chosen from a mathematical and not a physical point of view. Smoothness constraints involving parts that are more easily intepreted 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 (Jähne, 1991). Knowledge of the physical phenomena we are observing may then be translated into different weights for the three parts. Also the estimation may be improved be including restictions on temporal variation.

5.4 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., (Simpson & Humphrey, 1990; Gallaudet & Simpson, 1991)).

6 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 and second order derivatives has been described. Measures based on the local energy distribution to detect unreliable observations are also presented. These measures include deviation from coherent translation and the total energy present. The estimated flow fields have been used succesfully in two temporal interpolation scheme. First we have used the flow fields to generate intermediate frames in a sequence of Meteosat images, secondly, we have used the estimated flow fields to generate replacement frames in stead of drop-outs in the sequence. Both schemes are superior to the simple alternatives shown.

7 Acknowledgements

The authors wish to thank Senior Scientist Hans H. Valeur and Poul Nielsen, the Danish Meteorological Institute, for providing image data.