



Reliable Estimation of Dense Optical Flow Fields with Large Displacements

LUIS ALVAREZ

*Departamento de Informática y Sistemas, Universidad de Las Palmas, Campus de Tafira,
SP-35017 Las Palmas, Spain*
lalvarez@dis.ulpgc.es

JOACHIM WEICKERT

*Computer Vision, Graphics, and Pattern Recognition Group, Department of Mathematics and Computer Science,
University of Mannheim, D-68131 Mannheim, Germany*
Joachim.Weickert@ti.uni-mannheim.de

JAVIER SÁNCHEZ

*Departamento de Informática y Sistemas, Universidad de Las Palmas, Campus de Tafira,
SP-35017 Las Palmas, Spain*
jsanchez@dis.ulpgc.es

Abstract. In this paper we show that a classic optical flow technique by Nagel and Enkelmann (1986, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 8, pp. 565–593) can be regarded as an early anisotropic diffusion method with a diffusion tensor. We introduce three improvements into the model formulation that (i) avoid inconsistencies caused by centering the brightness term and the smoothness term in different images, (ii) use a linear scale-space focusing strategy from coarse to fine scales for avoiding convergence to physically irrelevant local minima, and (iii) create an energy functional that is invariant under linear brightness changes. Applying a gradient descent method to the resulting energy functional leads to a system of diffusion–reaction equations. We prove that this system has a unique solution under realistic assumptions on the initial data, and we present an efficient linear implicit numerical scheme in detail. Our method creates flow fields with 100 % density over the entire image domain, it is robust under a large range of parameter variations, and it can recover displacement fields that are far beyond the typical one-pixel limits which are characteristic for many differential methods for determining optical flow. We show that it performs better than the optical flow methods with 100 % density that are evaluated by Barron et al. (1994, Int. J. Comput. Vision, Vol. 12, pp. 43–47). Our software is available from the Internet.

Keywords: image sequences, optical flow, differential methods, anisotropic diffusion, linear scale-space, regularization, finite difference methods, performance evaluation

1. Introduction

The goal of optical flow computations is to estimate a so-called optical flow field which represents the apparent shift of greyvalue structures in the image plane. Recovering this displacement field is a key problem in computer vision and much research has been devoted to this field during the last two decades. For a survey

of these activities we refer to Mitiche and Boutheymy (1996), and performance evaluations of some of the most popular algorithms include papers of Barron et al. (1994), Jähne and Haussecker (in press), and Galvin et al. (1998).

One important class of optical flow methods consists of so-called differential methods. Often they are considered as useful only in the case of small displacement

fields. The goal of the present paper is to show that a combination of linear and nonlinear scale-space ideas may lead to a well-posed differential method that allows to recover the optical flow between two images with high accuracy, even in the case of large displacement fields.

We consider two images $I_1(x, y)$ and $I_2(x, y)$ (defined on \mathbb{R}^2 to simplify the discussion) which represent two consecutive views in a sequence of images. Under the assumption that corresponding pixels have equal grey values, the determination of the optical flow from I_1 to I_2 comes down to finding a function $\bar{h}(x, y) = (u(x, y), v(x, y))$ such that

$$I_1(x, y) = I_2(x + u(x, y), y + v(x, y)), \quad \forall (x, y) \in \mathbb{R}^2. \quad (1)$$

To compute $\bar{h}(x, y)$ the preceding equality is usually linearized yielding the so-called *linearized optical flow constraint*

$$I_1(\bar{x}) - I_2(\bar{x}) \approx \langle \nabla I_2(\bar{x}), \bar{h}(\bar{x}) \rangle \quad \forall \bar{x} \quad (2)$$

where $\bar{x} := (x, y)$. The linearized optical flow constraint is based on the assumption that the object displacements $\bar{h}(\bar{x})$ are small or that the image is slowly varying in space. In other cases, this linearization is no longer valid.

Frequently, instead of Eq. (1), the alternative equality

$$I_1(x - u(x, y), y - v(x, y)) = I_2(x, y), \quad \forall (x, y) \in \mathbb{R}^2 \quad (3)$$

is used. In this case the displacement $\bar{h}(x, y)$ is centred in the image $I_2(x, y)$.

The determination of optical flow is a classic ill-posed problem in computer vision (Bertero and Poggio, 1988), and it requires to be supplemented with additional regularizing assumptions. The regularization by Horn and Schunck (1981) reflects the assumption that the optical flow field varies smoothly in space. However, since many natural image sequences are better described in terms of piecewise smooth flow fields separated by discontinuities, much research has been done to modify the Horn and Schunck approach in order to permit such discontinuous flow fields (Aubert et al., to appear; Black and Anandan, 1991, 1996; Blanc-Féraud et al., 1993; Cohen, 1993; Deriche et al., 1995; Guichard and Rudin, 1996; Heitz and Boutheimy, 1993; Kumar et al., 1996; Mémin and Pérez, 1998; Nagel and Enkelmann, 1986; Nesi, 1993; Proesmans et al., 1994b;

Schnörr, 1994; Weickert, 1998). A survey of these approaches will be presented at the end of this section.

An important improvement of the Horn and Schunck method has been achieved by Nagel and Enkelmann (1986) (see also (Nagel, 1983)). They consider the following minimization problem:

$$E_{NE}(\bar{h}) = \int_{\mathbb{R}^2} (I_1(x - u(x, y), y - v(x, y)) - I_2(x, y))^2 dx dy + c \int_{\mathbb{R}^2} (\nabla u^T D(\nabla I_1) \nabla u + \nabla v^T D(\nabla I_1) \nabla v) dx dy \quad (4)$$

where c is a positive constant and $D(\nabla I_1)$ is a regularized projection matrix in the direction perpendicular of ∇I_1 :

$$D(\nabla I_1) = \frac{1}{|\nabla I_1|^2 + 2\lambda^2} \times \left\{ \begin{pmatrix} \frac{\partial I_1}{\partial y} \\ -\frac{\partial I_1}{\partial x} \end{pmatrix} \begin{pmatrix} \frac{\partial I_1}{\partial y} \\ -\frac{\partial I_1}{\partial x} \end{pmatrix}^T + \lambda^2 Id \right\}. \quad (5)$$

In this formulation, Id denotes the identity matrix. The advantage of this method is that it inhibits blurring of the flow across boundaries of I_1 at locations where $|\nabla I_1| \gg \lambda$.

In spite of its merits, however, this method still leaves room for improvements:

- (i) The Nagel–Enkelmann model uses an optical flow constraint which is centred in I_2 , while the projection matrix D in the smoothness term depends on I_1 . This inconsistency may create artifacts for large displacement fields.
- (ii) Refraining from a linearization of the optical flow constraint has the consequence that the energy functional (6) may be nonconvex. In this case popular algorithms such as gradient descent methods may get trapped in physically irrelevant local minima.
- (iii) Minimizers of the energy functional (6) are not invariant under linear brightness changes of the images I_1 and I_2 .

In the present paper we will address these points by introducing three improvements into the Nagel–Enkelmann model:

- (i) We design an energy functional that consistently centers both the optical flow constraint and the smoothness constraint in the same image.

- (ii) We encourage convergence to the global energy minimum by embedding the method into a linear scale-space framework that allows to focus down from coarse to fine scales in small steps.
- (iii) We introduce an adaptation of the parameters c and λ to the dynamic range of the images such that the resulting energy functional is invariant under linear brightness rescalings. This adaptation is particularly useful in the context of our scale-space focusing which alters the dynamic range of the images.

Applying the gradient descent method to our model leads to a coupled system of two diffusion-reaction equations, for which we establish the existence of a unique solution. Interestingly, these equations can be related to anisotropic diffusion filtering with a diffusion tensor. We present an efficient numerical scheme that is based on a linear implicit finite difference discretization. Afterwards, we discuss the role of the model parameters and demonstrate that our model allows very accurate recovery of optic flow fields for a large range of parameters. This is done by considering both synthetic image sequences, for which ground truth flow fields exist, as well as a real-world test sequence. Owing to the scale-space focusing, our model is particularly suited for recovering large displacement fields.

The paper is organized as follows: In Section 2 we describe our optical flow method that incorporates the three improvements, and we show that the Nagel–Enkelmann method and its modifications are closely related to anisotropic diffusion filtering. In Section 3 we present existence and uniqueness results for the nonlinear parabolic system that arises from using the gradient descent method for minimizing the energy functionals. Section 4 describes an efficient numerical discretization of this system based on a linear implicit finite difference scheme. Section 5 clarifies the role of the model parameters, and in Section 6 we present experimental results on synthetic and real-world image sequences. Finally, in Section 7 we conclude with a summary.

Related Work

Proesmans et al. (1994a, 1994b) studied a related approach that also dispenses with a linearization of the optical flow constraint in order to allow for larger displacements. Their method, however, requires six coupled partial differential equations and its nonlinear diffusion process uses a scalar-valued diffusivity instead

of a diffusion tensor. Their discontinuity-preserving smoothing is flow-driven while ours is image-driven. Another PDE technique that is similar in vein to the work of Proesmans et al. is a stereo method by Shah (1993). Other flow-driven regularizations with discontinuity-preserving properties include the work of Aubert et al. (to appear), Cohen (1993), Deriche et al. (1995), Hinterberger (1999), Kumar et al. (1996), Schnörr (1994), Weickert (1998), and Weickert and Schnörr (1999). Related stochastic regularization approaches have been studied by Black and Anandan (1991, 1996), Blanc-Féraud et al. (1993), Heitz and Bouthemy (1993), and Mémin and Pérez (1998). The image-driven anisotropic Nagel–Enkelmann approach has been subject to many subsequent studies. Examples include later work by Nagel (1987, 1990) as well as research by Schnörr (1991a, 1991b) and Snyder (1991). A multigrid realization of this method has been described by Enkelmann (1988), and a related pyramid framework is studied by Anandan (1989). An isotropic image-driven optic flow regularization is investigated by Alvarez et al. (1999). With respect to embeddings into a linear scale-space framework our method can be also be related to the optical flow approach of Florack et al. (1998). Their method differs from ours in that it is purely linear, applies scale selection mechanisms and does not use discontinuity-preserving nonlinear smoothness terms. Our focusing strategy for avoiding to end up in irrelevant local minima also resembles the *graduated non-convexity (GNC) algorithms* of Blake and Zisserman (1987). A preliminary version of our work has been presented at a conference (Alvarez et al., 1999), and a related optical flow method has been used by Hinterberger (1999) to generate a movie between two images.

2. The Model

In this section we consider three modifications of the Nagel–Enkelmann model in order to improve its performance in the case of large displacement fields. We also discuss relations between this method and anisotropic diffusion filtering.

2.1. Consistent Centering

We have seen that the energy functional (4) uses an optical flow constraint and a smoothness term that are centred in different images. Our experiments showed that this inconsistency may lead to artifacts when the

displacement field is large. As a remedy, we consider a modified energy functional where both the optical flow constraint and the smoothness constraint are related to I_1 :

$$\begin{aligned} E(\bar{h}) = & \int_{\mathbb{R}^2} (I_1(x, y) - I_2(x + u(x, y), y \\ & + v(x, y)))^2 dx dy \\ & + c \int_{\mathbb{R}^2} (\nabla u^T D(\nabla I_1) \nabla u \\ & + \nabla v^T D(\nabla I_1) \nabla v) dx dy. \end{aligned} \quad (6)$$

The associated Euler-Lagrange equations are given by the PDE system

$$\begin{aligned} c \operatorname{div} (D(\nabla I_1) \nabla u) + (I_1(\bar{x}) - I_2(\bar{x} + \bar{h}(\bar{x}))) \\ \times \frac{\partial I_2}{\partial x}(\bar{x} + \bar{h}(\bar{x})) = 0, \end{aligned} \quad (7)$$

$$\begin{aligned} c \operatorname{div} (D(\nabla I_1) \nabla v) + (I_1(\bar{x}) - I_2(\bar{x} + \bar{h}(\bar{x}))) \\ \times \frac{\partial I_2}{\partial y}(\bar{x} + \bar{h}(\bar{x})) = 0. \end{aligned} \quad (8)$$

In this paper, we are interested in solutions of the Eqs. (7) and (8) in the case of *large* displacement fields and images that are not necessarily slowly varying in space. Therefore, we do not use the linearized optic flow constraint (2) in the above system.

2.2. Relations to Anisotropic Diffusion Filtering

We obtain the solutions of the Euler-Lagrange Eqs. (7) and (8) by calculating the asymptotic state ($t \rightarrow \infty$) of the parabolic system

$$\begin{aligned} \frac{\partial u}{\partial t} = c \operatorname{div} (D(\nabla I_1) \nabla u) + (I_1(\bar{x}) - I_2(\bar{x} + \bar{h}(\bar{x}))) \\ \times \frac{\partial I_2}{\partial x}(\bar{x} + \bar{h}(\bar{x})), \end{aligned} \quad (9)$$

$$\begin{aligned} \frac{\partial v}{\partial t} = c \operatorname{div} (D(\nabla I_1) \nabla v) + (I_1(\bar{x}) - I_2(\bar{x} + \bar{h}(\bar{x}))) \\ \times \frac{\partial I_2}{\partial y}(\bar{x} + \bar{h}(\bar{x})). \end{aligned} \quad (10)$$

These equations do also arise when the steepest descent method is applied in order to minimize the energy (6).

Interestingly, this coupled system of diffusion-reaction equations reveals a diffusion tensor which resembles the one used for edge-enhancing anisotropic

diffusion filtering. Indeed, $D(\nabla I_1)$ has the eigenvectors $v_1 := \nabla I_1$ and $v_2 := \nabla I_1^\perp$. The corresponding eigenvalues are given by

$$\lambda_1(|\nabla I_1|) = \frac{\lambda^2}{|\nabla I_1|^2 + 2\lambda^2}, \quad (11)$$

$$\lambda_2(|\nabla I_1|) = \frac{|\nabla I_1|^2 + \lambda^2}{|\nabla I_1|^2 + 2\lambda^2}. \quad (12)$$

We observe, that $\lambda_1 + \lambda_2 = 1$ holds independently of ∇I_1 . In the interior of objects we have $|\nabla I_1| \rightarrow 0$, and therefore $\lambda_1 \rightarrow 1/2$ and $\lambda_2 \rightarrow 1/2$. At ideal edges where $|\nabla I_1| \rightarrow \infty$, we obtain $\lambda_1 \rightarrow 0$ and $\lambda_2 \rightarrow 1$. Thus, we have isotropic behaviour within regions, and at image boundaries the process smoothes anisotropically along the edge. This behaviour is very similar to edge-enhancing anisotropic diffusion filtering (Weickert, 1996), and it is also close in spirit to the modified mean-curvature motion considered in Alvarez et al. (1992). In this sense, one may regard the Nagel-Enkelmann method as an early predecessor of modern PDE techniques for image restoration.

One structural difference, however, should be observed: the optical flow Eqs. (9) and (10) use a temporally constant diffusion tensor, while the nonlinear diffusion tensor of anisotropic diffusion filtering is a function of the evolving image itself. Hence, the Nagel-Enkelmann model is anisotropic and space-variant, but it remains linear in its diffusion part. Related linear anisotropic diffusion filters have been pioneered by Iijima in the sixties and seventies in the context of optical character recognition; see Weickert et al. (1999) and the references therein. For a detailed treatment of anisotropic diffusion filtering we refer to Weickert (1998), an axiomatic classification of mean-curvature motion and related morphological PDEs for image analysis is presented in Alvarez et al. (1993), and recent collections of papers on PDE-based image smoothing methods include (Berger et al., 1996; Caselles et al., 1998; Romeny et al., 1997; Nielsen et al., 1999).

2.3. Recovering Large Displacements by Scale-Space Focusing

The energy functional (6) may be nonconvex due to its data term without linearization. In this case we cannot expect the uniqueness of solutions of the elliptic system (7) and (8). As a consequence, the asymptotic state of the parabolic system (9) and (10), which we use for approximating the optical flow, depends on the initial

data. Typically, we may expect that the algorithm converges to a local minimizer of the energy functional (6) that is located in the vicinity of the initial data. When we have small displacements in the scene, the natural choice is to take $u \equiv v \equiv 0$ as initialization of the flow. For large displacement fields, however, this may not work, and we need better initial data. To this end, we embed our method into a linear scale-space framework (Iijima, 1962; Weickert, 1999). Considering the problem at a coarse scale avoids that the algorithm gets trapped in physically irrelevant local minima. The coarse-scale solution serves then as initial data for solving the problem at a finer scale. Scale focusing has a long tradition in linear scale-space theory (see e.g. Bergholm (1987) for an early approach), and in spite of the fact that some theoretical questions remain open, it has not lost its popularity. For more details on linear scale-space theory we refer to Florack (1997), Iijima (1973, 1989), Lindeberg (1994) and Sporring et al. (1997). Using a scale-space approach enables us also to perform a finer and more reliable scale focusing as would be the case for related pyramid or multigrad approaches.

We proceed as follows. First, we introduce a linear scale factor in the parabolic PDE system in order to end up with

$$\begin{aligned} \frac{\partial u_\sigma}{\partial t} &= c \operatorname{div} (D (\nabla I_1^\sigma) \nabla u_\sigma) \\ &+ (I_1^\sigma(\bar{x}) - I_2^\sigma(\bar{x} + \bar{h}_\sigma(\bar{x}))) \\ &\times \frac{\partial I_2^\sigma}{\partial x}(\bar{x} + \bar{h}_\sigma(\bar{x})), \end{aligned} \quad (13)$$

$$\begin{aligned} \frac{\partial v_\sigma}{\partial t} &= c \operatorname{div} (D (\nabla I_1^\sigma) \nabla v_\sigma) \\ &+ (I_1^\sigma(\bar{x}) - I_2^\sigma(\bar{x} + \bar{h}_\sigma(\bar{x}))) \\ &\times \frac{\partial I_2^\sigma}{\partial y}(\bar{x} + \bar{h}_\sigma(\bar{x})) \end{aligned} \quad (14)$$

where $I_1^\sigma := G_\sigma * I_1$, $I_2^\sigma := G_\sigma * I_2$, $\bar{h}_\sigma(\bar{x}) := (u_\sigma(\bar{x}), v_\sigma(\bar{x}))$, and $G_\sigma * I_j$ represents the convolution of I_j with a Gaussian of standard deviation σ .

The convolution with a Gaussian blends the information in the images and allows us to recover a connection between the objects in I_1 and I_2 . In our application, this global support property that is characteristic for linear diffusion scale-spaces is very important. It makes them favourable over morphological scale-spaces in the sense of Alvarez et al. (1993), since the latter ones cannot transport information between topologically disconnected objects.

We start with a large initial scale σ_0 . Then we compute the optical flow $(u_{\sigma_0}, v_{\sigma_0})$ at scale σ_0 as the asymptotic state of the solution of the above PDE system using as initial data $u \equiv v \equiv 0$. Next, we choose a number of scales $\sigma_n < \sigma_{n-1} < \dots < \sigma_0$, and for each scale σ_i we compute the optical flow $(u_{\sigma_i}, v_{\sigma_i})$ as the asymptotic state of the above PDE system with initial data $(u_{\sigma_{i-1}}, v_{\sigma_{i-1}})$. The final computed flow corresponds to the smallest scale σ_n . In accordance with the logarithmic sampling strategy in linear scale-space theory (Koenderink, 1984), we choose $\sigma_i := \eta^i \sigma_0$ with some decay rate $\eta \in (0, 1)$.

2.4. Invariance Under Linear Greyvalue Transformations

A remaining shortcoming of the modified model is that the energy $E(\bar{h})$ is not invariant under grey level transformation of the form $(I_1, I_2) \rightarrow (kI_1, kI_2)$, where k is a constant. Therefore, the choice of the parameters depends strongly on the image contrast. This is especially problematic when the method is embedded in the scale-space focusing strategy, since the amount of smoothing influences the contrast range in the regularized images $G_\sigma * I_1$ and $G_\sigma * I_2$.

We address this problem by normalizing the parameters c and λ in such a way that the energy $E(\bar{h})$ becomes invariant under grey level transformation of the form $(I_1, I_2) \rightarrow (kI_1, kI_2)$. We compute c and λ by means of two parameters α and $s \in (0, 1)$ that are calculated via

$$\begin{aligned} c &= \frac{\alpha}{\max_{\bar{x}} (|\nabla G_\sigma * I_1(\bar{x})|^2)}, \\ s &= \int_0^\lambda \mathcal{H}_{|\nabla G_\sigma * I_1|}(z) dz \end{aligned}$$

where $\mathcal{H}_{|\nabla G_\sigma * I_1|}(z)$ represents the normalized histogram of $|\nabla G_\sigma * I_1|$. We name s the *isotropy fraction*. When $s \rightarrow 0$, the diffusion operator becomes anisotropic at all locations, and when $s \rightarrow 1$, it leads to isotropic diffusion everywhere. So now $c = c(\alpha, \nabla G_\sigma * I_1)$, and $\lambda = \lambda(s, \nabla G_\sigma * I_1)$. With this normalization of c and λ , the energy $E(\bar{h})$ is invariant under grey level transformation of the form $(I_1, I_2) \rightarrow (kI_1, kI_2)$. In practical applications of our method it is thus sufficient to specify the parameters α and s instead of c and λ . The parameters c and λ are then automatically adjusted to the dynamic image range in each step of the focusing procedure.

3. Existence and Uniqueness of the Parabolic System

In this section we show the existence and uniqueness of solutions of the parabolic system (13) and (14) where $D(\nabla I_1^\sigma)$ is given by (5). The parameters c and λ can be arbitrary positive real numbers. In particular, they may be determined as described in the previous section. First we introduce an abstract framework where we study the above system. This framework is used to show the existence and uniqueness of the solutions afterwards.

3.1. Abstract Framework

For simplicity we assume that the images are defined on the entire space \mathbb{R}^2 . We assume that the input images I_1 and I_2 belong to the functional space $L^2(\mathbb{R}^2)$. Let $H = L^2(\mathbb{R}^2) \times L^2(\mathbb{R}^2)$, and let us denote by $A : \mathcal{D}(A) \subset H \rightarrow H$ the differential operator defined by

$$A(\bar{h}) = -c \begin{pmatrix} \operatorname{div} (D(\nabla I_1^\sigma) \nabla u_\sigma) \\ \operatorname{div} (D(\nabla I_1^\sigma) \nabla v_\sigma) \end{pmatrix},$$

where $\mathcal{D}(A)$ is the domain of the mapping A . If $I_1 \in L^2(\mathbb{R}^2)$ then $I_1^\sigma \in W^{1,\infty}(\mathbb{R}^2)$, so ∇I_1^σ is bounded and the eigenvalues of the diffusion tensor $D(\nabla I_1^\sigma)$ are strictly positive. Therefore, as $c > 0$, the operator $A(\bar{h})$ is a maximal monotone operator. For more details about maximal monotone operators we refer to Brezis (1973). Next, let us introduce the function $F : H \rightarrow H$ defined by

$$F(\bar{h}) = (I_1^\sigma - I_2^\sigma (Id + \bar{h})) \nabla I_2^\sigma (Id + \bar{h}).$$

Then the abstract evolution problem can be written as

$$\begin{cases} \frac{d\bar{h}_\sigma}{dt} + A\bar{h}_\sigma = F(\bar{h}_\sigma) \text{ in } H, & \forall t \in [0, T] \\ \bar{h}_\sigma(0) = \bar{h}^0 \text{ in } H. \end{cases} \quad (15)$$

Any classical solution $\bar{h}_\sigma \in C^1([0, T]; H) \cap C([0, T]; \mathcal{D}(A))$ of (15) is given by

$$\bar{h}_\sigma(t) = S(t)\bar{h}^0 + \int_0^t S(t-s)F(\bar{h}_\sigma(s)) ds, \quad (16)$$

where $\{S(t)\}_{t>0}$ is the contraction semi-group associated to the homogeneous problem.

Definition. We say that $h \in C([0, T]; H)$ is a generalized solution of (15) if it satisfies (16).

3.2. Existence and Uniqueness Result

In order to prove existence and uniqueness, we have to establish a lemma first.

Lemma 1. *Suppose that $I_1, I_2 \in L^2(\mathbb{R}^2)$, then F is Lipschitz-continuous, and the Lipschitz constant L depends on the functions I_1 and I_2 and on σ .*

Proof: First we note that if $I_1, I_2 \in L^2(\mathbb{R}^2)$, then we have in particular that $I_2^\sigma \in W^{1,\infty}(\mathbb{R}^2)$ and $I_1^\sigma \in L^\infty(\mathbb{R}^2)$. Let $\bar{h}_1, \bar{h}_2 \in H$. For the i -th component of $F(\bar{h}_1) - F(\bar{h}_2)$, $i = 1, 2$, we have the following point-wise estimate.

$$\begin{aligned} & |F_i(\bar{h}_1) - F_i(\bar{h}_2)| \\ &= |(I_1^\sigma - I_2^\sigma (Id + \bar{h}_1)) \partial_i I_2^\sigma (Id + \bar{h}_1) \\ &\quad - (I_1^\sigma - I_2^\sigma (Id + \bar{h}_2)) \partial_i I_2^\sigma (Id + \bar{h}_2)|, \\ &\leq |I_2^\sigma (Id + \bar{h}_1) \partial_i I_2^\sigma (Id + \bar{h}_1) \\ &\quad - I_2^\sigma (Id + \bar{h}_2) \partial_i I_2^\sigma (Id + \bar{h}_2)| \\ &\quad + |I_1^\sigma| \cdot |\partial_i I_2^\sigma (Id + \bar{h}_1) - \partial_i I_2^\sigma (Id + \bar{h}_2)|, \\ &\leq \frac{1}{2} |\partial_i (|I_2^\sigma|^2) (Id + \bar{h}_1) - \partial_i (|I_2^\sigma|^2) (Id + \bar{h}_2)| \\ &\quad + \|I_1^\sigma\|_\infty \cdot |\partial_i I_2^\sigma (Id + \bar{h}_1) - \partial_i I_2^\sigma (Id + \bar{h}_2)|, \\ &\leq \frac{1}{2} C_{\text{Lip}}(\partial_i (|I_2^\sigma|^2)) \cdot |\bar{h}_1 - \bar{h}_2| + \|I_1^\sigma\|_\infty \\ &\quad \cdot C_{\text{Lip}}(\partial_i I_2^\sigma) \cdot |\bar{h}_1 - \bar{h}_2| \\ &\leq \left(\frac{1}{2} C_{\text{Lip}}(\partial_i (|I_2^\sigma|^2)) + \|I_1^\sigma\|_\infty \cdot C_{\text{Lip}}(\partial_i I_2^\sigma) \right) \\ &\quad \cdot |\bar{h}_1 - \bar{h}_2|, \end{aligned}$$

where $C_{\text{Lip}}(f)$ denotes the Lipschitz constant of the function f . We finally deduce that

$$\begin{aligned} \|F(\bar{h}_1) - F(\bar{h}_2)\|_H &= \|F_1(\bar{h}_1) - F_1(\bar{h}_2)\|_{L^2} \\ &\quad + \|F_2(\bar{h}_1) - F_2(\bar{h}_2)\|_{L^2} \\ &\leq \sum_{i=1}^2 \left(\frac{1}{2} C_{\text{Lip}}(\partial_i (|I_2^\sigma|^2)) + \|I_1^\sigma\|_\infty \right. \\ &\quad \left. \cdot C_{\text{Lip}}(\partial_i I_2^\sigma) \right) \cdot \|\bar{h}_1 - \bar{h}_2\|_H. \end{aligned}$$

We conclude the proof of the lemma by setting

$$L = \sum_{i=1}^2 \left(\frac{1}{2} C_{\text{Lip}}(\partial_i (|I_2^\sigma|^2)) + \|I_1^\sigma\|_\infty \cdot C_{\text{Lip}}(\partial_i I_2^\sigma) \right).$$

This shows the assertion. \square

Now we can state the existence and uniqueness result for problem (13) and (14).

Theorem 1. *Suppose that $I_1, I_2 \in L^2(\mathbb{R}^2)$ then, for all $\bar{h}^0 \in H$, there exists a unique generalized solution $\bar{h}_\sigma(t) \in C([0, \infty); H)$ of (13) and (14).*

Proof: The assumptions on I_1 and I_2 allow us to apply Lemma 1. Assume that $\bar{h}_1(t)$ and $\bar{h}_2(t)$ are solutions of (16) for initial conditions $\bar{h}_1(0)$ and $\bar{h}_2(0)$, then we have, using the fact that A is monotone (which yields $\|S(t)f\|_H \leq \|f\|_H$), and the Lipschitz continuity of F the following estimate.

$$\begin{aligned} \|\bar{h}_1(t) - \bar{h}_2(t)\|_H &\leq \|\bar{h}_1(0) - \bar{h}_2(0)\|_H \\ &\quad + L \int_0^t \|\bar{h}_1(s) - \bar{h}_2(s)\|_H ds. \end{aligned}$$

Applying the Gronwall–Bellman lemma (Brezis, 1973) gives

$$\|\bar{h}_1(t) - \bar{h}_2(t)\|_H \leq e^{Lt} \cdot \|\bar{h}_1(0) - \bar{h}_2(0)\|_H,$$

which yields uniqueness of the solution if it exists. Now consider the Banach space defined by

$$E = \{\bar{h} \in C([0, \infty); H) : \sup_{t \geq 0} \|\bar{h}(t)\|_H e^{-Kt} < \infty\}$$

endowed with the norm $\|\bar{h}\|_E = \sup_{t \geq 0} \|\bar{h}(t)\|_H e^{-Kt}$. Let $\phi : E \rightarrow C([0, \infty); H)$ be defined by

$$\phi(\bar{h})(t) = S(t)\bar{h}^0 + \int_0^t S(t-s)F(\bar{h}(s)) ds.$$

If $K > L$, then $\phi(E) \subset E$, and ϕ is $\frac{L}{K}$ -Lipschitz since

$$\begin{aligned} \|\phi(\bar{h}_1) - \phi(\bar{h}_2)\|_E &= \sup_{t \geq 0} \|\phi(\bar{h}_1)(t) - \phi(\bar{h}_2)(t)\|_H e^{-Kt}, \\ &\leq \sup_{t \geq 0} \int_0^t L \|\bar{h}_1(s) - \bar{h}_2(s)\|_H ds e^{-Kt} \\ &\leq \sup_{t \geq 0} L \|\bar{h}_1 - \bar{h}_2\|_E \cdot e^{-Kt} \int_0^t e^{Ks} ds \\ &\leq \sup_{t \geq 0} \frac{L}{K} \|\bar{h}_1 - \bar{h}_2\|_E \cdot e^{-Kt} (e^{Kt} - 1) \\ &\leq \frac{L}{K} \|\bar{h}_1 - \bar{h}_2\|_E. \end{aligned}$$

We deduce that ϕ is a contraction, and by Banach's fixed point theorem there exists a unique \bar{h}_σ such that $\phi(\bar{h}_\sigma) = \bar{h}_\sigma$. This is the generalized solution of (15), and the proof is concluded. \square

Remark. We notice that our existence and uniqueness proof is based on rather weak assumptions on the initial images I_1 and I_2 . We only assumed square integrability. They do not have to be continuous and may even be corrupted by noise or quantization artifacts, as is common for real-world images. The behaviour of the solution when σ goes to 0 is a challenging mathematical problem. If I_1 and I_2 are not sufficiently smooth, we cannot expect a good asymptotic behaviour when σ tends to 0. This suggests that the original images I_1 and I_2 should always be preprocessed by some small amount of Gaussian smoothing. In our experiments we shall also observe that it can be advantageous to stop the focusing procedure when σ attains the order of the pixel size.

4. Numerical Scheme

Next we describe an efficient algorithm for our optical flow model. We discretize the parabolic system (13) and (14) by finite differences (see e.g. Morton and Mayers (1994) for an introduction to this subject). All spatial derivatives are approximated by central differences, and for the discretization in t direction we use a linear implicit scheme. Let $D(\nabla G_\sigma * I_1) =: \begin{pmatrix} a & b \\ b & c \end{pmatrix}$. Then our linear implicit scheme has the structure

$$\begin{aligned} &\frac{u_{i,j}^{k+1} - u_{i,j}^k}{\tau} \\ &= c \left(\frac{a_{i+1,j} + a_{i,j}}{2} \frac{u_{i+1,j}^{k+1} - u_{i,j}^{k+1}}{h_1^2} + \frac{a_{i-1,j} + a_{i,j}}{2} \right. \\ &\quad \times \frac{u_{i-1,j}^{k+1} - u_{i,j}^{k+1}}{h_1^2} + \frac{c_{i,j+1} + c_{i,j}}{2} \frac{u_{i,j+1}^{k+1} - u_{i,j}^{k+1}}{h_2^2} \\ &\quad + \frac{c_{i,j-1} + c_{i,j}}{2} \frac{u_{i,j-1}^{k+1} - u_{i,j}^{k+1}}{h_2^2} + \frac{b_{i+1,j+1} + b_{i,j}}{2} \\ &\quad \times \frac{u_{i+1,j+1}^{k+1} - u_{i,j}^{k+1}}{2h_1h_2} + \frac{b_{i-1,j-1} + b_{i,j}}{2} \\ &\quad \times \frac{u_{i-1,j-1}^{k+1} - u_{i,j}^{k+1}}{2h_1h_2} - \frac{b_{i+1,j-1} + b_{i,j}}{2} \\ &\quad \times \frac{u_{i+1,j-1}^{k+1} - u_{i,j}^{k+1}}{2h_1h_2} - \frac{b_{i-1,j+1} + b_{i,j}}{2} \\ &\quad \left. \times \frac{u_{i-1,j+1}^{k+1} - u_{i,j}^{k+1}}{2h_1h_2} \right) \\ &+ I_{2,x}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) (I_1(\bar{x}_{i,j}) - I_2(\bar{x}_{i,j} + \bar{h}_{i,j}^k)) \\ &+ u_{i,j}^k I_{2,x}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) + v_{i,j}^k I_{2,y}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) \\ &- u_{i,j}^{k+1} I_{2,x}^2(\bar{x}_{i,j} + \bar{h}_{i,j}^k) - v_{i,j}^{k+1} I_{2,y}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) \\ &\times I_{2,x}(\bar{x}_{i,j} + \bar{h}_{i,j}^k), \end{aligned} \tag{17}$$

$$\begin{aligned}
& \frac{v_{i,j}^{k+1} - v_{i,j}^k}{\tau} \\
&= c \left(\frac{a_{i+1,j} + a_{i,j}}{2} \frac{v_{i+1,j}^{k+1} - v_{i,j}^{k+1}}{h_1^2} + \frac{a_{i-1,j} + a_{i,j}}{2} \right. \\
&\quad \times \frac{v_{i-1,j}^{k+1} - v_{i,j}^{k+1}}{h_1^2} + \frac{c_{i,j+1} + c_{i,j}}{2} \frac{v_{i,j+1}^{k+1} - v_{i,j}^{k+1}}{h_2^2} \\
&\quad + \frac{c_{i,j-1} + c_{i,j}}{2} \frac{v_{i,j-1}^{k+1} - v_{i,j}^{k+1}}{h_2^2} + \frac{b_{i+1,j+1} + b_{i,j}}{2} \\
&\quad \times \frac{v_{i+1,j+1}^{k+1} - v_{i,j}^{k+1}}{2h_1h_2} + \frac{b_{i-1,j-1} + b_{i,j}}{2} \\
&\quad \times \frac{v_{i-1,j-1}^{k+1} - v_{i,j}^{k+1}}{2h_1h_2} - \frac{b_{i+1,j-1} + b_{i,j}}{2} \\
&\quad \times \frac{v_{i+1,j-1}^{k+1} - v_{i,j}^{k+1}}{2h_1h_2} - \frac{b_{i-1,j+1} + b_{i,j}}{2} \\
&\quad \left. \times \frac{v_{i-1,j+1}^{k+1} - v_{i,j}^{k+1}}{2h_1h_2} \right) \\
&+ I_{2,y}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) (I_1(\bar{x}_{i,j}) - I_2(\bar{x}_{i,j} + \bar{h}_{i,j}^k)) \\
&+ u_{i,j}^k I_{2,x}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) + v_{i,j}^k I_{2,y}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) \\
&- u_{i,j}^{k+1} I_{2,x}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) I_{2,y}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) \\
&- v_{i,j}^{k+1} I_{2,y}^2(\bar{x}_{i,j} + \bar{h}_{i,j}^k). \tag{18}
\end{aligned}$$

Although this scheme might look fairly complicated at first glance, it is actually straightforward to implement. The notations are almost self-explaining: for instance, τ is the time step size, h_1 and h_2 denote the pixel size in x and y direction, respectively, $u_{i,j}^k$ approximates u_σ in some grid point $\bar{x}_{i,j}$ at time $k\tau$, and $I_{1,x}$ is an approximation to $G_\sigma * \frac{\partial I_1}{\partial x}$. We calculate values of type $I_2(\bar{x}_{i,j} + \bar{h}_{i,j}^k)$ by linear interpolation.

The idea behind linear implicit schemes is to use implicit discretizations in order to improve stability properties, as long that they lead to linear systems of equations. Implicit means that unknowns of the new time level appear on both sides of the equation. For nonlinear equations, fully implicit methods would require a computationally expensive solution of nonlinear systems of equations. The computationally less expensive linear implicit methods avoid this by using suitable Taylor expansions, such that the resulting system of equations become linear. In our case we achieved this by using

the first-order Taylor expansion

$$\begin{aligned}
I_2(\bar{x}_{i,j} + \bar{h}_{i,j}^{k+1}) &\approx I_2(\bar{x}_{i,j} + \bar{h}_{i,j}^k) \\
&+ (u_{i,j}^{k+1} - u_{i,j}^k) I_{2,x}(\bar{x}_{i,j} + \bar{h}_{i,j}^k) \\
&+ (v_{i,j}^{k+1} - v_{i,j}^k) I_{2,y}(\bar{x}_{i,j} + \bar{h}_{i,j}^k)
\end{aligned}$$

in a fully implicit discretization, and by discretizing $G_\sigma * \frac{\partial I_2}{\partial x}$ and $G_\sigma * \frac{\partial I_2}{\partial y}$ in an explicit way. A consistency analysis shows that the preceding scheme is of second order in space and of first order in time.

We solve the resulting linear system of equations iteratively by a symmetric Gauß–Seidel algorithm. In order to explain its structure let us suppose that we want to solve a linear system $A\bar{w} = \bar{b}$ where $A = D - L - U$ and D is a diagonal matrix, L a strictly lower triangular matrix, and U a strictly upper triangular matrix. Then the symmetric Gauß–Seidel iterations are given by

$$\begin{aligned}
(D - L) \bar{w}^{(n+1/2)} &= \bar{b} + U \bar{w}^{(n)}, \\
(D - U) \bar{w}^{(n+1)} &= \bar{b} + L \bar{w}^{(n+1/2)}
\end{aligned}$$

where the upper index denotes the iteration index. The systems are solved directly using forward and backward elimination, respectively.

In an earlier version of our work (Alvarez et al., 1999) we have studied an explicit scheme that did not require to solve linear systems of equations. The linear implicit approach that we employ in the meantime, however, has led to a speed-up of one to two orders of magnitude, since it allows significantly larger time step sizes without creating stability problems.

5. Parameters

Our algorithm for computing the optical flow depends on a number of parameters that have an intuitive meaning:

- The regularization parameter α specifies the balance between the smoothing term and the optical flow constraint. Larger values lead to smoother flow fields by filling in information from image edges where flow measurements with higher reliability are available.
- The isotropy fraction s determines the contrast parameter λ via the cumulative histogram of the image gradient magnitude. Choosing e.g. $s := 0.7$ means that the smoothness term diffuses isotropically at 70% of all image locations, while 30% of all locations are assumed to belong to image edges, where

smoothing is performed anisotropically along the edge.

- The scale σ_0 denotes the standard deviation of the largest Gaussian. In general, σ_0 is chosen according to the maximum displacement expected.
- The decay rate $\eta \in (0, 1)$ for the computation of the scales $\sigma_m := \eta^m \sigma_0$. We may expect a good focusing if η is close to 1.
- The smallest scale is given by σ_n . It should be close to the inner scale of the image in order to achieve optimal flow localization.
- The time step size τ and the stopping time T for solving the system (13) and (14) at each scale σ_m are pure numerical parameters. We experienced that fixing $\tau := 10$ and $T := 500$ creates results that are sufficiently close to the asymptotic state. Using smaller values of τ or larger values of T slowed down the algorithm without improving the quality of the flow fields.

In the next section we will see that the results of our method are hardly affected by fairly large parameter variations. As a consequence default values can be used for most of the parameters.

6. Experimental Results

Figure 1 shows our first experiment. We use a synthetic image composed of four black squares on a white back-

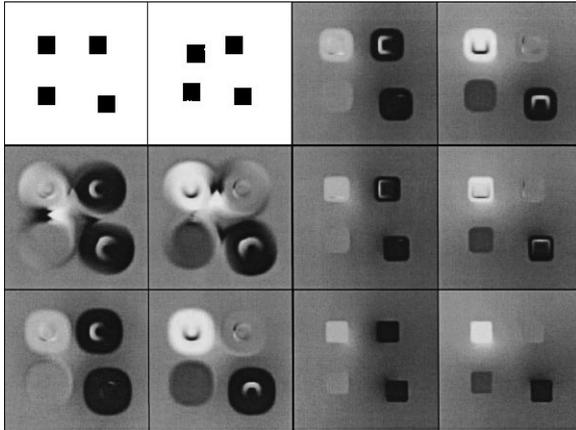


Figure 1. Computation of the optical flow for the square images with $\alpha = 0.6$, $s = 0.1$, and $\eta = 0.95$. (a) First row, left pair: Original image pair. (b) Second row, left pair: Optical flow components (u, v) for $\sigma_0 = 10$. (c) Third row, left pair: Optical flow result after focusing down to $\sigma_{12} = 5.7$. (d) First row, right pair: $\sigma_{25} = 2.9$. (e) Second row, right pair: $\sigma_{37} = 1.4$. (f) Third row, right pair: $\sigma_{50} = 0.8$.

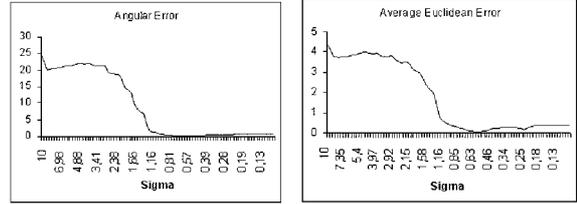


Figure 2. Left: Average angular error of the optic flow calculations for the squares in the first frame of Fig. 1. Right: Corresponding average Euclidean error.

ground. Each square moves in a different direction and with a different displacement magnitude: under the assumption that the x axis is oriented from left to right and the y axis from top to bottom, the left square on the top moves with $(u, v) = (5, 10)$, the right square on the top is displaced with $(u, v) = (-10, 0)$, the left square on the bottom is shifted by $(u, v) = (0, -5)$, and the right square on the bottom undergoes a translation by $(-10, -10)$. In order to visualize the flow field (u, v) we use two grey level images (u_{gl}, v_{gl}) defined by $u_{gl} := 128 + 12u$ and $v_{gl} := 128 + 12v$. In Fig. 1 we notice that the flow estimates improve significantly by focusing down from $\sigma_0 := 10$ to $\sigma_{50} := 0.8$: flow discontinuities evolve and the calculated flow fields approximate the true motion field more and more.

This qualitative observation is confirmed in the quantitative evaluations carried out in Fig. 2. The left plot shows the average angular errors in the four squares of the first frame. The angular error Ψ_e has been calculated in the same way as in Barron et al. (1994) using

$$\Psi_e := \arccos \left(\frac{u_c u_e + v_c v_e + 1}{\sqrt{(u_c^2 + v_c^2 + 1)(u_e^2 + v_e^2 + 1)}} \right) \quad (19)$$

where (u_c, v_c) denotes the correct flow, and (u_e, v_e) is the estimated flow. The right plot depicts the Euclidean error $\sqrt{(u_e - u_c)^2 + (v_e - v_c)^2}$ averaged over all pixels within the four squares of the first frame.

In both cases we observe that the error is reduced drastically by focusing down in scale-space until it reaches a very small value when the Gaussian width σ approaches the inner scale of the image. Further reduction of σ leads to slightly larger errors. It appears that this is caused by discretization and quantization effects. We evaluated the error only in the interior of the squares because of the constant background. The flow is not defined correctly in this area in the sense

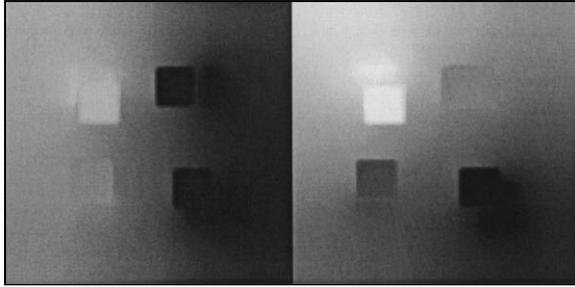


Figure 3. Result when the optical flow constraint and the smoothness constraint are centred in different images. Images and parameters are identical with those in Fig. 1.

that any displacement of the background is compatible with the image sequence.

Figure 3 demonstrates the importance of consistent centering of the optical flow constraint and the smoothness constraint. In this experiment the smoothness term follows the gradient of I_2 , while the optic flow term is centred in I_1 , i.e. we are looking for displacements from I_1 to I_2 . Since the boundary locations of I_1 and I_2 differ for large displacements, it is not surprising that the optical flow field that respects the discontinuities of I_2 is “leaking” at the boundaries of I_1 .

We notice that when an object moves across the image sequence, the background is partially occluded. This occlusion problem is illustrated in Fig. 4. In the

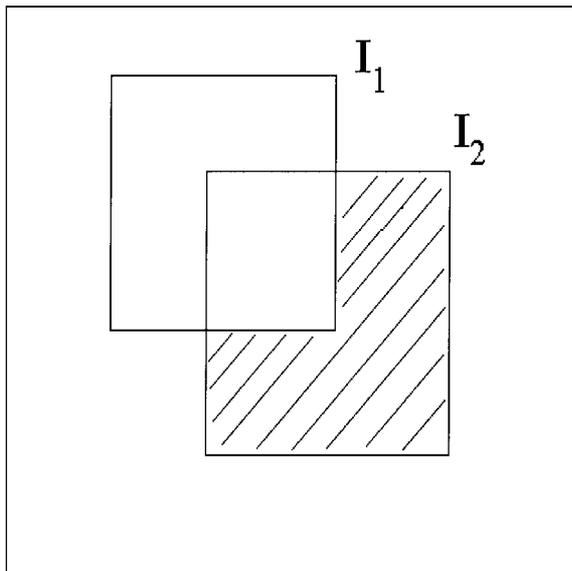


Figure 4. Illustration of the occlusion problem. A square is moving from I_1 to I_2 . The shadowed region in the image I_1 has no correspondence in I_2 .

direction of the object motion a region of the background is occluded, so the points of this region (the shadowed area of Fig. 4) have no correspondence in I_2 , and the optical flow constraint is no longer valid. In this background region some slight inhomogeneities appear as can be seen in Fig. 1. However, we observed that the smoothness term of the energy helps to reduce such effects.

For the following experiment we use the classical Hamburg taxi test sequence. These data are available at the ftp site <ftp://csd.uwo.ca> in the directory `pub/vision`. Instead of taking two consecutive frames—as is usually done—we consider the frames 15 and 19. The dark car at the left creates a largest displacement magnitude of approximately 12 pixels. In Figs. 5 and 6 we present the computed flow. The

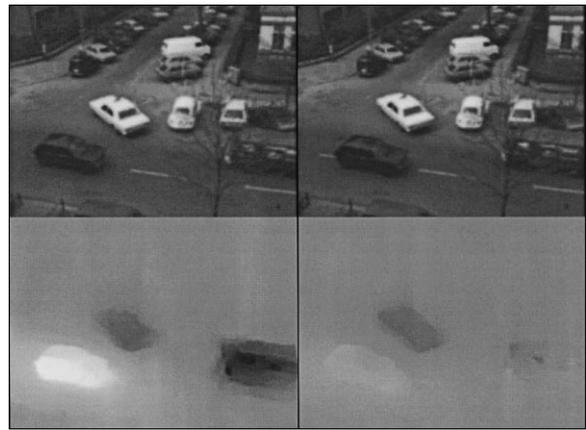


Figure 5. Computation of the optical flow for the taxi sequence (frames 15 and 19) with $\alpha = 0.6$, $s = 0.1$, $\sigma_0 = 10$, $\sigma_n = 0.8$, and $\eta = 0.95$.

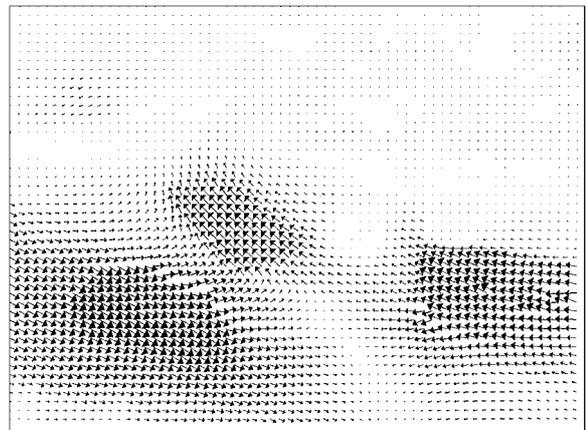


Figure 6. Vector plot of the optical flow from Fig. 5.

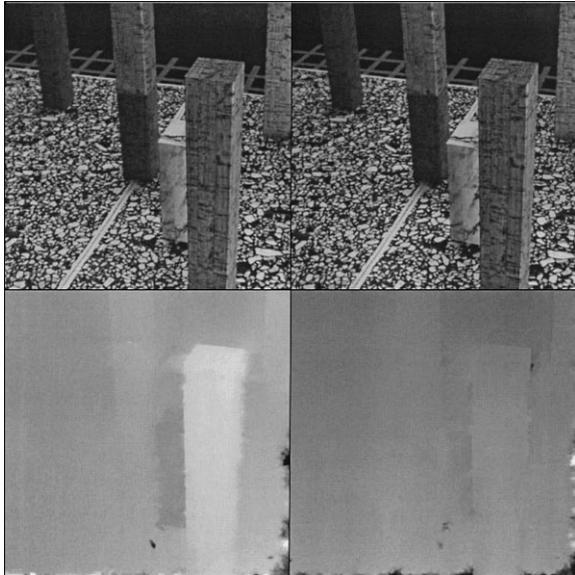


Figure 7. Optical flow results for the marbled block sequence (frames 20 and 25) with $\alpha = 0.6$, $s = 0.1$, $\sigma_0 = 10$, $\sigma_n = 0.8$, and $\eta = 0.95$.

computed maximal flow magnitude is 11.68, which is a good approximation of the actual displacement of the dark car. It is interesting to note that, although the movement of the pedestrian in the upper left part of the scene is difficult to recognize in the greyscale plot in Fig. 5, the vector plot in Fig. 6 shows that this motion is not suppressed.

In Figs. 7 and 8 we show the results of our method for the marbled block sequence. This sequence is copyright by H.-H. Nagel (KOGS/IAKS, University of Karlsruhe, Germany). It has been recorded and first evaluated by Michael Otte (Otte and Nagel, 1995) and it is available from <http://i21www.ira.uka.de/image-sequences>. The marbled block sequence described a real-world scene where the camera is moving. In our experiments we took the frame 20 and 25, and we used the same parameters as for the taxi scene.

Next we perform quantitative comparisons with classic optical flow techniques from the survey paper of Barron et al. (1994). This is done using their ground truth data as well as the evaluation utilities that are available from their ftp site <ftp://csd.uwo.ca> in the directory `pub/vision`. It should be noted that the results in Barron et al. (1994) have been achieved with test sequences where the displacements are small, while our method is designed for large displacement fields. Moreover, their methods also used a presmoothing in

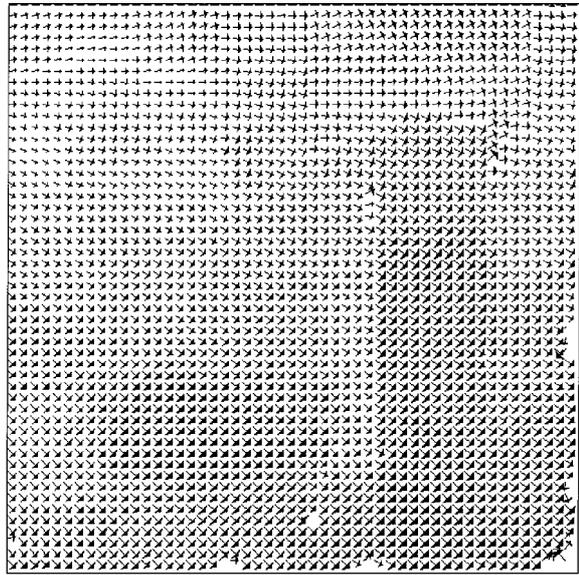


Figure 8. Vector plot of the optical flow from Fig. 7.

time which involves more than two frames, whereas we use only two frames. In spite of these limitations we are going to show that we can obtain competitive results with our method.

In the comparison we focus on those methods in Barron et al. (1994) that create flow fields with 100% density. For many subsequent tasks such as the inference of egomotion and surface structure this is a very desirable property. Local methods that yield a lower density may have to be supplemented with additional strategies for filling in information at locations where no results are available. Their practical performance may thus depend heavily on this postprocessing. Variational approaches with smoothness terms do not require such a postprocessing step as they automatically yield flow fields with 100% density.

In Figs. 9 and 10 we show the computed optical flow for the *Square2* sequence that depicts a square moving with velocity $(4/3, -4/3)$. Table 1 gives a comparison with the results of Barron et al. for some classic optic flow techniques that create flow fields with 100% density. It can be seen that our proposed technique reveals smaller errors than these methods. In particular, this also shows that our three modifications improve Nagel's method substantially. While the implementation of Nagel's method in Barron et al. (1994) gives an angular error of 34.57° , our method reveals an error of 10.97° . In this example Barron et al. assume that the background moves in the same direction

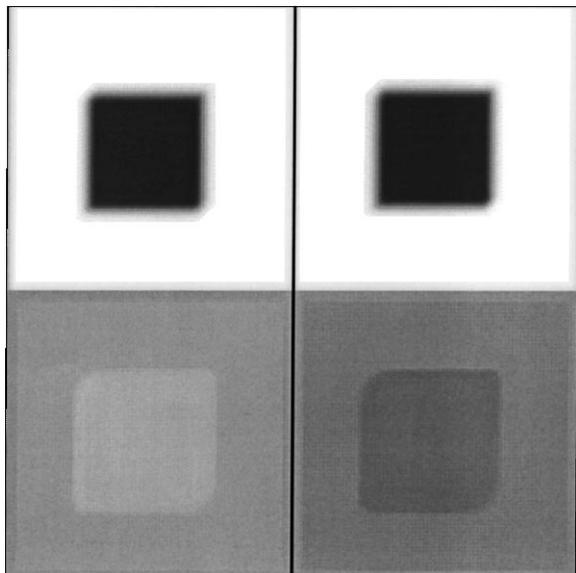


Figure 9. Computation of the optical flow for the *Square2* sequence with $\alpha = 0.6$, $s = 0.1$, $\sigma_0 = 10$, $\sigma_n = 1$, and $\eta = 0.95$.

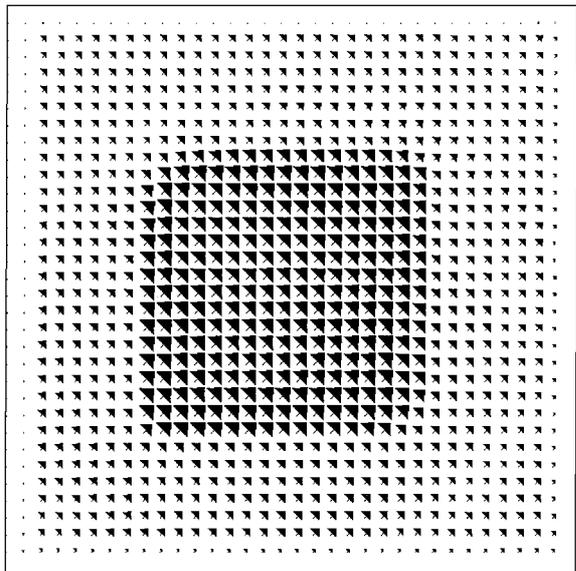


Figure 10. Vector plot of the optical flow from Fig. 9.

as the square. However, as the background is constant the displacement is not well defined in this area. If we focus our attention on the error of the computed flow within the interior of the square we obtain an average angular error of 0.85° . This shows that the computed flow is very accurate in the interior of the square.

Next we draw our attention to the most complex synthetic test sequence from Barron et al. (1994), the Yosemite sequence with cloudy sky. It contains dis-

Table 1. Comparison between the results from Barron et al. (1994) with 100% density and our method for the *Square2* sequence.

Technique	Aver. error ($^\circ$)	Stand. deviat. ($^\circ$)	Density (%)
Horn and Schunck (original)	47.21	14.60	100
Horn and Schunck (modified)	32.81	13.67	100
Nagel	34.57	14.38	100
Anandan (unthresholded)	31.46	18.31	100
Singh (step 1)	49.03	21.38	100
Singh (step 2)	46.12	18.64	100
Our method	10.97	9.60	100

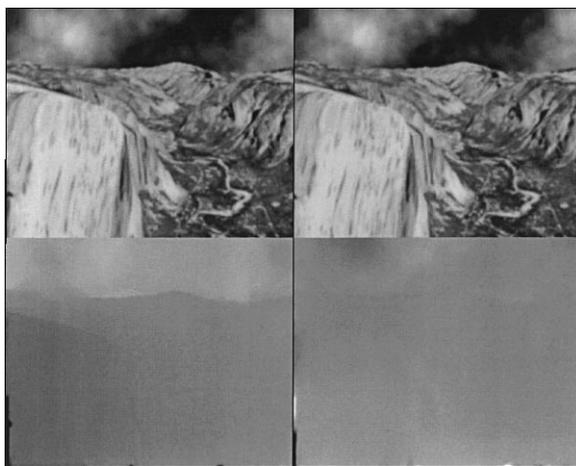


Figure 11. Computation of the optical flow for the Yosemite sequence with $\alpha = 0.6$, $s = 0.1$, $\sigma_0 = 5$, $\sigma_n = 1$, and $\eta = 0.95$.

placements of up to five pixels. Our optical flow results are shown in Figs. 11 and 12, and a juxtaposition with other methods can be found in Table 2. Again our technique outperforms all methods from Barron et al. (1994) which yield flow fields with 100% density. With an angular error of 5.53° it even reaches the estimation quality of typical methods with 30% density, and the standard deviation of 7.40° is lower than the standard deviation of all methods that have been evaluated in Barron et al. (1994): the best method (Lucas and Kanade with $\lambda_2 \geq 5.0$) had an average angular error of 3.22° with a standard deviation of 8.92° and a density of only 8.7%.

In order to evaluate the robustness of our algorithm with respect to the choice of parameters we present in Table 3 the errors for the Yosemite sequence taking different values of the parameters. To simplify the

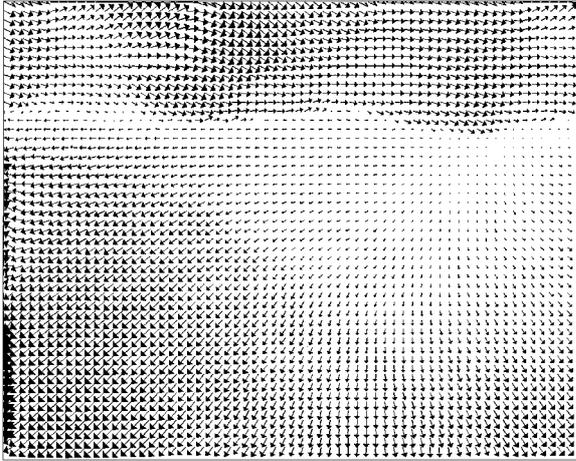


Figure 12. Vector plot of the optical flow from Fig. 11.

Table 2. Comparison between the results from Barron et al. (1994) with 100% density and our method for the Yosemite sequence.

Technique	Aver. error (°)	Stand. deviat. (°)	Density (%)
Horn and Schunck (original)	31.69	31.18	100
Horn and Schunck (modified)	9.78	16.19	100
Nagel	10.22	16.51	100
Anandan (unthresholded)	13.36	15.64	100
Uras et al. (unthresholded)	8.94	15.61	100
Singh (step 2)	10.03	13.13	100
Our method	5.53	7.40	100

presentation, we fixed the finest scale to $\sigma_n := 1$, and as numerical parameters we used $\tau := 10$ and $T := 500$. These parameters are almost independent of the image and can therefore be set to default values. Hence, we vary only the parameters α , s , η and σ_0 in Table 3.

First of all it can be seen that our method outperforms all methods in Barron et al. (1994) with 100% density not only in case of optimized parameters, but also for a rather large range of parameter settings. Let us now study the parameter influence in more detail.

One important observation from Table 3 is that the decay parameter η has an important influence of the result: values around 0.5, as are implicitly used by typical pyramid-based focusing algorithms, are by far not optimal. A slow focusing with $\eta = 0.95$ gives significantly better results. Our experience with other images suggests that η may be fixed to this value for all applications.

Table 3. Errors for the Yosemite sequence, using different parameters of the algorithm.

Smoothness α	Init. scale σ_0	Isotr. fract. s	Decay rate η	Angul. error (°)	Stand. dev. (°)
0.4	5	0.1	0.90	5.61	7.46
0.5	„	„	„	5.57	7.41
0.6	„	„	„	5.55	7.37
0.7	„	„	„	5.56	7.33
1.0	„	„	„	5.69	7.24
0.6	1	0.1	0.90	16.83	15.23
„	2.5	„	„	5.92	7.31
„	5	„	„	5.55	7.37
„	10	„	„	5.54	7.37
„	15	„	„	5.81	8.45
0.6	5	0.01	0.90	5.70	7.92
„	„	0.1	„	5.55	7.37
„	„	0.2	„	5.70	7.31
„	„	0.5	„	6.38	8.14
„	„	0.8	„	7.31	9.76
„	„	0.9	„	7.64	10.37
„	„	0.99	„	8.04	11.21
0.6	5	0.1	0.50	7.25	7.58
„	„	„	0.70	6.14	7.36
„	„	„	0.80	5.75	7.33
„	„	„	0.95	5.53	7.40
„	„	„	0.99	5.56	7.45

Choosing too a small value for the isotropy fraction s does hardly worsen the results, while for larger values the smoothness term becomes isotropic almost everywhere and approximates the Horn and Schunck scheme (Horn and Schunck, 1981). In order to avoid the resulting deteriorations, we propose to fix $s := 0.1$, which means that the method smooths anisotropically at 90% of all locations.

Regarding the smoothness parameter α , our method appeared to be rather robust with respect to over- and underestimations. We have thus used a fixed value of 0.6 for all experiments in the present paper.

As already mentioned, the initial scale σ_0 should be chosen such that it covers the largest expected displacements. We found that overestimations are less critical than underestimations. This also confirms the use of the focusing strategy. Too small values increase the danger of ending up in a physically irrelevant local minimum. Actually, σ_0 was basically the only parameter that we had to adapt in order to analyse different image

sequences. Since it has a clear physical interpretation, this adaptation was simple.

Remark. More detailed information about the experiments in this section can be found at the web site <http://serdis.dis.ulpgc.es/~lalvarez/research/demos>. In particular, some movies to illustrate the focusing strategy are presented. At this site we also provide a window oriented image processing software named XMetaWave (see González and Trujillo, 1994) which includes the algorithm that we have developed in this paper.

7. Conclusions

Usually, when computer vision researchers deal with variational methods for optical flow calculations, they linearize the optical flow constraint. Except for those cases where the images vary sufficiently slowly in space, linearization, however, does only work for small displacements. In this paper we introduced three improvements into a classical method by Nagel and Enkelmann where no linearization is used. We identified this method as two coupled linear anisotropic diffusion filters with a nonlinear reaction term. We showed that this parabolic system is well-posed from a mathematical viewpoint, and we presented a linear implicit finite difference scheme for its efficient numerical solution. In order to avoid that the algorithms converges to physically irrelevant local minima, we embedded it into a linear scale-space approach for focusing the solution from a coarse to a fine scale. A detailed quantitative analysis using test sequences with ground truth data showed the following results.

- The method can recover displacements of more than 10 pixels with good accuracy.
- It performs significantly better than Nagel's original method and all other methods with 100% density that are evaluated by Barron et al. (1994).
- The performance hardly deteriorated for quite a large range of parameters. This allows to use default parameter settings for many applications.

In spite of these favourable properties, there are still possibilities for further improvements and extensions of this algorithm. For instance, we expect that our method can benefit from results obtained from preceding frames. One can interpolate the flow in the previous

frames in order to have a first estimation of the flow in the current frame (Black and Anandan, 1991). If this first estimation is a reasonable initial guess for our algorithm, there is no need to introduce large scale corrections by using large values for σ_0 . A smaller initial scale σ_0 speeds up the focusing procedure and makes the algorithm faster. As an example for another extension possibility, we are currently investigating the use of our method for related matching problems such as 3D reconstructions from greyscale and colour stereo pairs (Alvarez et al., 2000; Alvarez and Sánchez, 2000). It is our hope that our method that combines anisotropic diffusion-reaction equations with linear scale-space techniques may serve as a motivation to study other combinations of linear and nonlinear scale-space approaches for solving computer vision problems.

Acknowledgement

This work has been supported by the European TMR network *Viscosity Solutions and their Applications*.

References

- Alvarez, L., Deriche, R., Sánchez, J., and Weickert, J. 2000. Dense disparity map estimation respecting image derivatives: A PDE and scale-space based approach. Report No. 3874, ROBOTVIS, INRIA, 2004 route des Lucioles, BP 93, 06902 Sophia-Antipolis Cedex, France.
- Alvarez, L., Esclarín, J., Lefébure, M., and Sánchez, J. 1999. A PDE model for computing the optical flow. In *Proc. XVI Congreso de Ecuaciones Diferenciales y Aplicaciones, C.E.D.Y.A. XVI*, Las Palmas de Gran Canaria, September 21–24, 1999, pp. 1349–1356.
- Alvarez, L., Guichard, F., Lions, P.-L., and Morel, J.-M. 1993. Axions and fundamental equations in image processing. *Arch. Rational Mech. Anal.*, 123:199–257.
- Alvarez, L., Lions, P.-L., and Morel, J.-M. 1992. Image selective smoothing and edge detection by nonlinear diffusion II. *SIAM J. Numer. Anal.*, 29:845–866.
- Alvarez, L. and Sánchez, J. 2000. 3-D geometry reconstruction using a color image stereo pair and partial differential equations. Report No. 6, Instituto Universitario de Ciencias y Tecnologías Cibernéticas, University of Las Palmas de Gran Canaria, Campus de Tafira, 35017 Las Palmas, Spain.
- Alvarez, L., Weickert, J., and Sánchez, J. 1999. A scale-space approach to nonlocal optical flow calculations. In *Lecture Notes in Computer Science, Vol. 1682: Scale-Space Theories in Computer Vision*, M. Nielsen, P. Johansen, O.F. Olsen, and J. Weickert (Eds.). Springer: Berlin, pp. 235–246.
- Anandan, P. 1989. A computational framework and an algorithm for the measurement of visual motion. *Int. J. Comput. Vision*, 2:283–310.
- Aubert, G., Deriche, R., and Kornprobst, P. to appear. Computing optical flow via variational techniques, in *SIAM J. Math. Anal.*
- Barron, J.L., Fleet, D.J., and Beauchemin, S.S. 1994. Performance of optical flow techniques. *Int. J. Comput. Vision*, 12:43–77.

- Berger, M.-O., Deriche, R., Herlin, I., Jaffré, J., and Morel, J.-M. (Eds.). 1996. ICAOS '96: Images, wavelets and PDEs. *Lecture Notes in Control and Information Sciences*, Vol. 219. Springer: London.
- Bergholm, F. 1987. Edge focusing. *IEEE Trans. Pattern Anal. Mach. Intell.*, 9:726–741.
- Bertero, M., Poggio, T.A., and Torre, V. 1988. Ill-posed problems in early vision. *Proc. IEEE*, 76:869–889.
- Black, M.J. and Anandan, P. 1991. Robust dynamic motion estimation over time. In *Proc. IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition, CVPR '91*, Maui, June 3–6, 1991, IEEE Computer Society Press, Los Alamitos, pp. 292–302.
- Black, M.J. and Anandan, P. 1996. The robust estimation of multiple motions: Parametric and piecewise smooth flow fields. *Computer Vision and Image Understanding*, 63:75–104.
- Blake, A. and Zisserman, A. 1987. *Visual Reconstruction*. MIT Press: Cambridge, MA.
- Blanc-Féraud, L., Barlaud, M., and Gaidon, T. 1993. Motion estimation involving discontinuities in a multiresolution scheme. *Optical Engineering*, 32(7):1475–1482.
- Brezis, H. 1973. *Opérateurs maximaux monotones et semi-groupes de contractions dans les espaces de Hilbert*. North Holland: Amsterdam.
- Caselles, V., Morel, J.M., Sapiro, G., and Tannenbaum, A. (Eds.). 1998. Special issue on partial differential equations and geometry-driven diffusion in image processing and analysis. *IEEE Trans. Image Proc.* 7(3).
- Cohen, I. 1993. Nonlinear variational method for optical flow computation. In *Proc. Eighth Scandinavian Conf. on Image Analysis, SCIA '93*, Tromsø, May 25–28, 1993, Vol. 1, pp. 523–530.
- Deriche, R., Kornprobst, P., and Aubert, G. 1995. Optical-flow estimation while preserving its discontinuities: A variational approach. In *Proc. Second Asian Conf. Computer Vision, ACCV '95*, Singapore, December 5–8, 1995, Vol. 2, pp. 290–295.
- Enkelmann, W. 1988. Investigation of multigrid algorithms for the estimation of optical flow fields in image sequences. *Computer Vision, Graphics and Image Processing*, 43:150–177.
- Florack, L. 1997. *Image Structure*. Kluwer: Dordrecht.
- Florack, L.M.J., Niessen, W.J., and Nielsen, M. 1998. The intrinsic structure of the optic flow incorporating measurement duality. *Int. J. Comput. Vision*, 27:263–286.
- Galvin, B., McCane, B., and Novins, K. 1998. Recovering motion fields: An analysis of eight optical flow algorithms. In *Proc. 1998 British Machine Vision Conference, BMVC '98*, Southampton, September 14–17, 1998.
- González, E., Trujillo, A., and XMegaWave. A window oriented image processing software. Departamento de Informática y Sistemas, Universidad de Las Palmas. <http://amiserver.dsis.ulpgc.es/xmwgus/>, 1994.
- Guichard, F. and Rudin, L. 1996. Accurate estimation of discontinuous optical flow by minimizing divergence related functionals. In *Proc. Third Int. Conf. on Image Processing, ICIP '96*, Lausanne, September 1996, Vol. 1, pp. 497–500.
- Romeny, B. ter Haar, Florack, L., Koenderink, J., and Viergever, M. (Eds.). 1997. Scale-space theory in computer vision. *Lecture Notes in Computer Science*, Vol. 1252. Springer: Berlin.
- Heitz, F. and Boutheimy, P. 1993. Multimodal estimation of discontinuous optical flow using Markov random fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15:1217–1232.
- Hinterberger, W. 1999. Generierung eines Films zwischen zwei Bildern mit Hilfe des optischen Flusses. M.Sc. Thesis, Industrial Mathematics Institute, University of Linz, Austria.
- Horn, B. and Schunck, B. 1981. Determining optical flow. *Artificial Intelligence*, 17:185–203.
- Iijima, T. 1962. Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bulletin of the Electrotechnical Laboratory* (in Japanese), 26:368–388.
- Iijima, T. 1973. *Pattern Recognition*. (in Japanese) Corona-sha.
- Iijima, T. 1989. *Series of Basic Information Technology, Vol. 6: The-ory of Pattern Recognition* (in Japanese). Morishita Publishing.
- Jähne, B. and Haussecker, H. 1998. Performance characteristics of low-level motion estimators in spatiotemporal images. In *Proc. Workshop on Evaluation and Validation of Computer Vision Algorithms*, Dagstuhl, March 16–20, 1998, R. Haralick, R. Klette, H.S. Stiehl, and M. Viergever (Eds.). Kluwer: Dordrecht. in press.
- Koenderink, J.J. 1984. The structure of images. *Biological Cybernetics*, 50:363–370.
- Kumar, A., Tannenbaum, A.R., and Balas, G.J. 1996. Optic flow: A curve evolution approach. *IEEE Trans. Image Proc.*, 5:598–610.
- Lindeberg, T. 1994. *Scale-Space Theory in Computer Vision*. Kluwer: Boston.
- Mémin, E. and Pérez, P. 1998. Dense estimation and object-based segmentation of the optical flow with robust techniques. *IEEE Trans. Image Proc.*, 7:703–719.
- Mitiche, A. and Boutheimy, P. 1996. Computation and analysis of image motion: A synopsis of current problems and methods. *Int. J. Comput. Vision*, 19:29–55.
- Morton, K.W. and Mayers, L.M. 1994. *Numerical Solution of Partial Differential Equations*. Cambridge University Press: Cambridge.
- Nagel, H.-H. 1983. Constraints for the estimation of displacement vector fields from image sequences. In *Proc. Eighth Int. Joint Conf. on Artificial Intelligence, IJCAI '83*, Karlsruhe, August 8–12, 1983, pp. 945–951.
- Nagel, H.-H. 1987. On the estimation of optical flow: Relations between different approaches and some new results. *Artificial Intelligence*, 33:299–324.
- Nagel, H.-H. 1990. Extending the 'oriented smoothness constraint' into the temporal domain and the estimation of derivatives of optical flow. In *Lecture Notes in Computer Science, Vol. 427: Computer Vision—ECCV '90*, O. Faugeras (Ed.). Springer: Berlin, pp. 139–148.
- Nagel, H.-H. and Enkelmann, W. 1986. An investigation of smoothness constraints for the estimation of displacement vector fields from images sequences. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8:565–593.
- Nesi, P. 1993. Variational approach to optical flow estimation managing discontinuities. *Image and Vision Computing*, 11:419–439.
- Nielsen, M., Johansen, P., Olsen, O.F., and Weickert, J. (Eds.). 1999. Scale-space theories in computer vision. *Lecture Notes in Computer Science*, Vol. 1682. Springer: Berlin.
- Otte, M. and Nagel, H.-H. 1995. Estimation of optical flow based on higher-order spatiotemporal derivatives in interlaced and non-interlaced image sequences. *Artificial Intelligence*, 78:5–43.
- Proesmans, M., Pauwels, E., and Van Gool, L. 1994a. Coupled geometry-driven diffusion equations for low-level vision. In *Geometry-Driven Diffusion in Computer Vision*, B.M. ter Haar Romeny (Ed.). Kluwer: Dordrecht, pp. 191–228.
- Proesmans, M., Van Gool, L., Pauwels, E., and Oosterlinck, A. 1994b. Determination of optical flow and its discontinuities using

- non-linear diffusion. In *Lecture Notes in Computer Science, Vol. 801: Computer Vision—ECCV '94*, J.-O. Eklundh (Ed.). Springer: Berlin, pp. 295–304.
- Schnörr, C. 1991a. Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class. *Int. J. Comput. Vision*, 6:25–38.
- Schnörr, C. 1991b. On functionals with greyvalue-controlled smoothness terms for determining optical flow. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15:1074–1079.
- Schnörr, C. 1994. Segmentation of visual motion by minimizing convex non-quadratic functionals. In *Proc. 12th Int. Conf. Pattern Recognition, ICPR 12*, Jerusalem, October 9–13, 1994, IEEE Computer Society Press, Los Alamitos, Vol. A, pp. 661–663.
- Shah, J. 1993. A nonlinear diffusion model for discontinuous disparity and half-occlusions in stereo. In *Proc. IEEE Comp. Soc. Conf. Computer Vision and Pattern Recognition, CVPR '93*, New York, June 15–17, 1993, IEEE Computer Society Press, Los Alamitos, pp. 34–40.
- Snyder, M.A. 1991. On the mathematical foundations of smoothness constraints for the determination of optical flow and for surface reconstruction. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13:1105–1114.
- Sporring, J., Nielsen, M., Florack, L., and Johansen, P. (Eds.). 1997. *Gaussian Scale-Space Theory*. Kluwer: Dordrecht.
- Weickert, J. 1996. Theoretical foundations of anisotropic diffusion in image processing. *Computing, Suppl.*, 11:221–236.
- Weickert, J. 1998. *Anisotropic Diffusion in Image Processing*. Teubner; Stuttgart.
- Weickert, J. 1998. On discontinuity-preserving optic flow. In *Proc. Computer Vision and Mobile Robotics Workshop, CVMR '98*, Santorini, September 17–18, 1998, S. Orphanoudakis, P. Trahanias, J. Crowley, and N. Katevas (Eds.). pp. 115–122.
- Weickert, J., Ishikawa, S., and Imiya, A. 1999. Linear scale-space has first been proposed in Japan. *J. Math. Imag. Vision*, 10:237–252.
- Weickert, J. and Schnörr, C. 1999. Räumlich-zeitliche Berechnung des optischen Flusses mit nichtlinearen flußabhängigen Glattheitsstermen. *Mustererkennung 1999*, W. Förstner, J.M. Buhmann, A. Faber, and P. Faber (Eds.). Springer: Berlin, pp. 317–324.