



Lucas/Kanade Meets Horn/Schunck: Combining Local and Global Optic Flow Methods

ANDRÉS BRUHN AND JOACHIM WEICKERT

*Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science, Saarland University,
Building 27, 66041 Saarbrücken, Germany*

bruhn@mia.uni-saarland.de

weickert@mia.uni-saarland.de

CHRISTOPH SCHNÖRR

*Computer Vision, Graphics and Pattern Recognition Group, Faculty of Mathematics and Computer Science,
University of Mannheim, 68131 Mannheim, Germany*

schnoerr@uni-mannheim.de

Received August 5, 2003; Revised April 22, 2004; Accepted April 22, 2004

First online version published in October, 2004

Abstract. Differential methods belong to the most widely used techniques for optic flow computation in image sequences. They can be classified into local methods such as the Lucas–Kanade technique or Bigün’s structure tensor method, and into global methods such as the Horn/Schunck approach and its extensions. Often local methods are more robust under noise, while global techniques yield dense flow fields. The goal of this paper is to contribute to a better understanding and the design of novel differential methods in four ways: (i) We juxtapose the role of smoothing/regularisation processes that are required in local and global differential methods for optic flow computation. (ii) This discussion motivates us to describe and evaluate a novel method that combines important advantages of local and global approaches: It yields dense flow fields that are robust against noise. (iii) Spatiotemporal and nonlinear extensions as well as multiresolution frameworks are presented for this hybrid method. (iv) We propose a simple confidence measure for optic flow methods that minimise energy functionals. It allows to sparsify a dense flow field gradually, depending on the reliability required for the resulting flow. Comparisons with experiments from the literature demonstrate the favourable performance of the proposed methods and the confidence measure.

Keywords: optic flow, differential techniques, variational methods, structure tensor, partial differential equations, confidence measures, performance evaluation

1. Introduction

Ill-posedness is a problem that is present in many image processing and computer vision techniques: Edge detection, for example, requires the computation of image derivatives. This problem is ill-posed in the sense of Hadamard,¹ as small perturbations in the signal may create large fluctuations in its derivatives (Yuille

and Poggio, 1986). Another example consists of optic flow computation, where the ill-posedness manifests itself in the nonuniqueness due to the aperture problem (Bertero et al., 1988): The data allow to compute only the optic flow component normal to image edges. Both types of ill-posedness problems appear jointly in so-called *differential methods* for optic flow recovery, where optic flow estimation is based on computing

spatial and temporal image derivatives. These techniques can be classified into *local* methods that may optimise some local energy-like expression, and *global* strategies which attempt to minimise a global energy functional. Examples of the first category include the Lucas–Kanade method (Lucas and Kanade, 1981; Lucas, 1984) and the structure tensor approach of Bigün and Granlund (1988) and Bigün et al. (1991), while the second category is represented by the classic method of Horn and Schunck (Horn and Schunck, 1981) and its numerous discontinuity-preserving variants (Alvarez et al., 1999; Aubert et al., 1999; Black and Anandan, 1991; Cohen, 1993; Heitz and Bouthemy, 1993; Kumar et al., 1996; Nagel, 1983; Nesi, 1993; Proesmans et al., 1994; Schnörr, 1994; Shulman and Hervé, 1989; Weickert and Schnörr, 2001). Differential methods are rather popular: Together with phase-based methods such as (Fleet and Jepson, 1990) they belong to the techniques with the best performance (Barron et al., 1994; Galvin et al., 1998). Local methods may offer relatively high robustness under noise, but do not give dense flow fields. Global methods, on the other hand, yield flow fields with 100% density, but are experimentally known to be more sensitive to noise (Barron et al., 1994; Galvin et al., 1998).

A typical way to overcome the ill-posedness problems of differential optic flow methods consists of the use of smoothing techniques and smoothness assumptions: It is common to smooth the image sequence prior to differentiation in order to remove noise and to stabilise the differentiation process. Local techniques use spatial constancy assumptions on the optic flow field in the case of the Lucas–Kanade method, and spatiotemporal constancy for the Bigün method. Global approaches, on the other hand, supplement the optic flow constraint with a regularising smoothness term. Surprisingly, the actual role and the difference between these smoothing strategies, however, has hardly been addressed in the literature so far. In a first step of this paper we juxtapose the role of the different smoothing steps of these methods. We shall see that each smoothing process offers certain advantages that cannot be found in other cases. Consequently, it would be desirable to combine the different smoothing effects of local and global methods in order to design novel approaches that combine the high robustness of local methods with the full density of global techniques. One of the goals of the present paper is to propose and analyse such an embedding of local methods into global approaches. This results in

a technique that is robust under noise and gives flow fields with 100% density. Hence, there is no need for a postprocessing step where sparse data have to be interpolated.

On the other hand, it has sometimes been criticised that there is no reliable confidence measure that allows to sparsify the result of a dense flow field such that the remaining flow is more reliable (Barron et al., 1994). In this way it would be possible to compare the real quality of dense methods with the characteristics of local, nondense approaches. In our paper we shall present such a measure. It is simple and applicable to the entire class of energy minimising global optic flow techniques. Our experimental evaluation will show that this confidence measure can give excellent results.

Our paper is organised as follows. In Section 2 we discuss the role of the different smoothing processes that are involved in local and global optic flow approaches. Based on these results we propose two *combined local-global (CLG) methods* in Section 3, one with spatial, the other one with spatiotemporal smoothing. In Section 4 nonlinear variants of the CLG method are presented, while a suitable multiresolution framework is discussed in Section 5. Our numerical algorithm is described in Section 6. In Section 7, we introduce a novel confidence measure for all global optic flow methods that use energy functionals. Section 8 is devoted to performance evaluations of the CLG methods and the confidence measure. A summary and an outlook to future work is given in Section 9. In the Appendix, we show how the CLG principle has to be modified if one wants to replace the Lucas–Kanade method by the structure tensor method of Bigün and Granlund (1988) and Bigün et al. (1991).

1.1. Related Work

In spite of the fact that there exists a very large number of publications on motion analysis (see e.g. (Mitiche and Bouthemy, 1996; Stiller and Konrad, 1999) for reviews), there has been remarkably little work devoted to the integration of local and global optic flow methods. Schnörr (Schnörr, 1993) sketched a framework for supplementing global energy functionals with multiple equations that provide local data constraints. He suggested to use the output of Gaussian filters shifted in frequency space (Fleet and Jepson, 1990) or

local methods incorporating second-order derivatives (Tretiak and Pastor, 1984; Uras et al., 1988), but did not consider methods of Lucas–Kanade or Bigün type.

Our proposed technique differs from the majority of global regularisation methods by the fact that we also use spatiotemporal regularisers instead of spatial ones. Other work with spatiotemporal regularisers includes publications by Murray and Buxton (1987), Nagel (1990), Black and Anandan (1991), Elad and Feuer (1998), and Weickert and Schnörr (2001).

While the noise sensitivity of local differential methods has been studied intensively in recent years (Bainbridge-Smith and Lane, 1997; Fermüller et al., 2001; Jähne, 2001; Kearney et al., 1987; Ohta, 1996; Simoncelli et al., 1991), the noise sensitivity of global differential methods has been analysed to a significantly smaller extent. In this context, Galvin et al. (1998) have compared a number of classical methods where small amounts of Gaussian noise had been added. Their conclusion was similar to the findings of Barron et al. (1994): the global approach of Horn and Schunck is more sensitive to noise than the local Lucas–Kanade method.

A preliminary shorter version of the present paper has been presented at a conference (Bruhn et al., 2002). Additional work in the current paper includes (i) the use of nonquadratic penalising functions, (ii) the application of a suitable multiresolution strategy, (iii) the proposal of a confidence measure for the entire class of global variational methods, (iv) the integration of the structure tensor approach of Bigün and Granlund (1988) and Bigün et al. (1991) and (v) a more extensive experimental evaluation.

2. Role of the Smoothing Processes

In this section we discuss the role of smoothing techniques in differential optic flow methods. For simplicity we focus on spatial smoothing. All spatial smoothing strategies can easily be extended into the temporal domain. This will usually lead to improved results (Weickert and Schnörr, 2001).

Let us consider some image sequence $g(x, y, t)$, where (x, y) denotes the location within a rectangular image domain Ω , and $t \in [0, T]$ denotes time. It is common to smooth the image sequence prior to differentiation (Barron et al., 1994; Kearney et al., 1987), e.g. by convolving each frame with some Gaussian $K_\sigma(x, y)$ of standard deviation σ :

$$f(x, y, t) := (K_\sigma * g)(x, y, t), \quad (1)$$

The low-pass effect of Gaussian convolution removes noise and other destabilising high frequencies. In a subsequent optic flow method, we may thus call σ the *noise scale*.

Many differential methods for optic flow are based on the assumption that the grey values of image objects in subsequent frames do not change over time:

$$f(x+u, y+v, t+1) = f(x, y, t), \quad (2)$$

where the displacement field $(u, v)^\top(x, y, t)$ is called *optic flow*. For small displacements, we may perform a first order Taylor expansion yielding the *optic flow constraint*

$$f_x u + f_y v + f_t = 0, \quad (3)$$

where subscripts denote partial derivatives. Evidently, this single equation is not sufficient to uniquely compute the two unknowns u and v (*aperture problem*): For nonvanishing image gradients, it is only possible to determine the flow component parallel to $\nabla f := (f_x, f_y)^\top$, i.e. normal to image edges. This so-called *normal flow* is given by

$$w_n = -\frac{f_t}{|\nabla f|} \frac{\nabla f}{|\nabla f|}. \quad (4)$$

Figure 1(a) depicts one frame from the famous Hamburg taxi sequence.² We have added Gaussian noise, and in Fig. 1(b)–(d) we illustrate the impact of presmoothing the image data on the normal flow. While some moderate presmoothing improves the results, great care should be taken not to apply too much presmoothing, since this would severely destroy important image structure.

In order to cope with the aperture problem, Lucas and Kanade (1981) and Lucas (1984) proposed to assume that the unknown optic flow vector is constant within some neighbourhood of size ρ . In this case it is possible to determine the two *constants* u and v at some location (x, y, t) from a weighted least square fit by minimising the function

$$E_{LK}(u, v) := K_\rho * ((f_x u + f_y v + f_t)^2). \quad (5)$$

Here the standard deviation ρ of the Gaussian serves as an *integration scale* over which the main contribution of the least square fit is computed.

A minimum (u, v) of E_{LK} satisfies $\partial_u E_{LK} = 0$ and $\partial_v E_{LK} = 0$. This gives the linear system

$$\begin{pmatrix} K_\rho * (f_x^2) & K_\rho * (f_x f_y) \\ K_\rho * (f_x f_y) & K_\rho * (f_y^2) \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} -K_\rho * (f_x f_t) \\ -K_\rho * (f_y f_t) \end{pmatrix} \quad (6)$$




Figure 1. From left to right, and from top to bottom: (a) Frame 10 of the Hamburg taxi sequence, where Gaussian noise with standard deviation $\sigma_n = 10$ has been added. The white taxi turns around the corner, the left car drives to the right, and the right van moves to the left. (b) Normal flow magnitude without presmoothing. (c) Normal flow magnitude, presmoothing with $\sigma = 1$. (d) Ditto, presmoothing with $\sigma = 5$. (e) Lucas-Kanade method with $\sigma = 0$, $\rho = 7.5$. (f) Ditto, $\sigma = 0$, $\rho = 15$. (g) Optic flow magnitude with the Horn-Schunck approach, $\sigma = 0$, $\alpha = 10^5$. (h) Ditto, $\sigma = 0$, $\alpha = 10^6$.

which can be solved provided that its system matrix is invertible. This is not the case in flat regions where the image gradient vanishes. In some other regions, the smaller eigenvalue of the system matrix may be close to 0, such that the aperture problem remains present and the data do not allow a reliable determination of the full optic flow. All this results in nondense flow fields. They constitute the most severe drawback of local gradient methods: Since many computer vision

applications require dense flow estimates, subsequent interpolation steps are needed. On the other hand, one may use the smaller eigenvalue of the system matrix as a confidence measure that characterises the reliability of the estimate. Experiments by Barron et al. (1994) indicated that this performs better than the trace-based confidence measure in Simoncelli et al. (1991).

Figure 1(e) and (f) show the influence of the integration scale ρ on the final result. In these images we have

displayed the entire flow field regardless of its local reliability. We can see that in each case, the flow field has typical structures of order ρ . In particular, a sufficiently large value for ρ is very successful in rendering the Lucas–Kanade method robust under noise.

In order to end up with dense flow estimates one may embed the optic flow constraint into a regularisation framework. Horn and Schunck (Horn and Schunck, 1981) have pioneered this class of global differential methods. They determine the unknown *functions* $u(x, y, t)$ and $v(x, y, t)$ as the minimisers of the global energy functional

$$E_{HS}(u, v) = \int_{\Omega} ((f_x u + f_y v + f_t)^2 + \alpha(|\nabla u|^2 + |\nabla v|^2)) dx dy \quad (7)$$

where the smoothness weight $\alpha > 0$ serves as *regularisation parameter*: Larger values for α result in a stronger penalisation of large flow gradients and lead to smoother flow fields.

Minimising this convex functional comes down to solving its corresponding Euler–Lagrange equations (Courant and Hilbert, 1953; Elsgolc, 1961). They are given by

$$0 = \Delta u - \frac{1}{\alpha}(f_x^2 u + f_x f_y v + f_x f_t), \quad (8)$$

$$0 = \Delta v - \frac{1}{\alpha}(f_x f_y u + f_y^2 v + f_y f_t). \quad (9)$$

with reflecting boundary conditions. Δ denotes the spatial Laplace operator:

$$\Delta := \partial_{xx} + \partial_{yy}. \quad (10)$$

The solution of these diffusion–reaction equations is not only unique (Schnörr, 1991), it also benefits from the *filling-in effect*: At locations with $|\nabla f| \approx 0$, no reliable local flow estimate is possible, but the regulariser $|\nabla u|^2 + |\nabla v|^2$ fills in information from the neighbourhood. This results in dense flow fields and makes subsequent interpolation steps obsolete. This is a clear advantage over local methods.

It has, however, been criticised that for such global differential methods, no good confidence measures are available that would help to determine locations where the computations are more reliable than elsewhere (Barron et al., 1994). It has also been observed that they may be more sensitive to noise than local differential methods (Barron et al., 1994; Galvin et al., 1998).

An explanation for this behaviour can be given as follows. Noise results in high image gradients. They serve as weights in the data term of the regularisation functional (7). Since the smoothness term has a constant weight α , smoothness is relatively less important at locations with high image gradients than elsewhere. As a consequence, *flow fields are less regularised at noisy image structures*. This sensitivity under noise is therefore nothing else but a side-effect of the desired filling-in effect. Figure 1(g) and (h) illustrate this behaviour. Figure 1(g) shows that the flow field does not reveal a uniform scale: It lives on a fine scale at high gradient image structures, and the scale may become very large when the image gradient tends to zero. Increasing the regularisation parameter α will finally also smooth the flow field at noisy structures, but at this stage, it might already be too blurred in flatter image regions (Fig. 1(h)).

3. A Combined Local–Global Method

We have seen that both local and global differential methods have complementary advantages and shortcomings. Hence it would be interesting to construct a hybrid technique that constitutes the best of two worlds: It should combine the robustness of local methods with the density of global approaches. This shall be done next. We start with spatial formulations before we extend the approach to the spatiotemporal domain.

3.1. Spatial Approach

In order to design a *combined local–global (CLG) method*, let us first reformulate the previous approaches. Using the notations

$$\mathbf{w} := (u, v, 1)^{\top}, \quad (11)$$

$$|\nabla \mathbf{w}|^2 := |\nabla u|^2 + |\nabla v|^2, \quad (12)$$

$$\nabla_3 f := (f_x, f_y, f_t)^{\top}, \quad (13)$$

$$J_{\rho}(\nabla_3 f) := K_{\rho} * (\nabla_3 f \nabla_3 f^{\top}) \quad (14)$$

it becomes evident that the Lucas–Kanade method minimises the quadratic form

$$E_{LK}(\mathbf{w}) = \mathbf{w}^{\top} J_{\rho}(\nabla_3 f) \mathbf{w}, \quad (15)$$

while the Horn–Schunck technique minimises the functional

$$E_{HS}(\mathbf{w}) = \int_{\Omega} (\mathbf{w}^{\top} J_0(\nabla_3 f) \mathbf{w} + \alpha |\nabla \mathbf{w}|^2) dx dy. \quad (16)$$

This terminology suggests a natural way to extend the Horn–Schunck functional to the desired CLG functional. We simply replace the matrix $J_0(\nabla_3 f)$ by the structure tensor $J_{\rho}(\nabla_3 f)$ with some integration scale $\rho > 0$. Thus, we propose to minimise the functional

$$E_{CLG}(\mathbf{w}) = \int_{\Omega} (\mathbf{w}^{\top} J_{\rho}(\nabla_3 f) \mathbf{w} + \alpha |\nabla \mathbf{w}|^2) dx dy. \quad (17)$$

Its minimising flow field (u, v) satisfies the Euler–Lagrange equations

$$0 = \Delta u - \frac{1}{\alpha} (K_{\rho} * (f_x^2) u + K_{\rho} * (f_x f_y) v + K_{\rho} * (f_x f_t)), \quad (18)$$

$$0 = \Delta v - \frac{1}{\alpha} (K_{\rho} * (f_x f_y) u + K_{\rho} * (f_y^2) v + K_{\rho} * (f_y f_t)). \quad (19)$$

It should be noted that these equations are hardly more complicated than the original Horn–Schunck Eqs. (8) and (9). All one has to do is to evaluate the terms containing image data at a nonvanishing integration scale. The basic structure with respect to the unknown functions $u(x, y, t)$ and $v(x, y, t)$ is identical. It is therefore not surprising that the well-posedness proof for the Horn–Schunck method that was presented in (Schnörr, 1991) can also be extended to this case.

3.2. Spatiotemporal Approach

The previous approaches used only *spatial* smoothness operators. Rapid advances in computer technology, however, makes it now possible to consider also spatiotemporal smoothness operators. Formal extensions in this direction are straightforward. In general, one may expect that spatiotemporal formulations give better results than spatial ones because of the additional denoising properties along the temporal direction. In the presence of temporal flow discontinuities smoothing along the time axis should only be used moderately. However, even in this case one can observe the beneficial effect of temporal information.

A spatiotemporal variant of the Lucas–Kanade approach simply replaces convolution with 2-D Gaus-

sians by spatiotemporal convolution with 3-D Gaussians. This still leads to a 2×2 linear system of equations for the two unknowns u and v .

Spatiotemporal versions of the Horn–Schunck method have been considered by Elad and Feuer (1998), while discontinuity preserving global methods with spatiotemporal regularisers have been proposed in different formulations in Black and Anandan (1991), Murray and Buxton (1987), Nagel (1990), Weickert and Schnörr (2001).

Combining the temporal extended variant of both the Lucas–Kanade and the Horn–Schunck method we obtain a spatiotemporal version of our CLG functional given by

$$E_{CLG3}(\mathbf{w}) = \int_{\Omega \times [0, T]} (\mathbf{w}^{\top} J_{\rho}(\nabla_3 f) \mathbf{w} + \alpha |\nabla_3 \mathbf{w}|^2) dx dy dt \quad (20)$$

where convolutions with Gaussians are now to be understood in a spatiotemporal way and

$$|\nabla_3 \mathbf{w}|^2 := |\nabla_3 u|^2 + |\nabla_3 v|^2. \quad (21)$$

Due to the different role of space and time the spatiotemporal Gaussians may have different standard deviations in both directions. Let us denote by J_{nm} the component (n, m) of the structure tensor $J_{\rho}(\nabla_3 f)$. Then the Euler–Lagrange equations for (20) are given by

$$\Delta_3 u - \frac{1}{\alpha} (J_{11} u + J_{12} v + J_{13}) = 0, \quad (22)$$

$$\Delta_3 v - \frac{1}{\alpha} (J_{12} u + J_{22} v + J_{23}) = 0. \quad (23)$$

One should note that they have the same structure as (18)–(19), apart from the fact that spatiotemporal Gaussian convolution is used, and that the spatial Laplacean Δ is replaced by the spatiotemporal Laplacean

$$\Delta_3 := \partial_{xx} + \partial_{yy} + \partial_{tt}. \quad (24)$$

The spatiotemporal Lucas–Kanade method is similar to the approach of Bigün and Granlund (1988) and Bigün et al. (1991). In the Appendix we show how the latter method can be embedded in a global energy functional.

4. Nonquadratic Approach

So far the underlying Lucas–Kanade and Horn–Schunck approaches are *linear* methods that are based on *quadratic* optimisation. It is possible to replace them by *nonquadratic* optimisation problems that lead to *nonlinear* methods. From a statistical viewpoint this can be regarded as applying methods from robust statistics where outliers are penalised less severely than in quadratic approaches (Hampel et al., 1986; Huber, 1981). In general, nonlinear methods give better results at locations with flow discontinuities. Robust variants of the Lucas–Kanade method have been investigated by Black and Anandan (1996) and by Yacoob and Davis (1999), respectively, while a survey of the numerous convex discontinuity-preserving regularisers for global optic flow methods is presented in Weickert and Schnörr (2001).

In order to render our approach more robust against outliers in both the data and the smoothness term we propose the minimisation of the following functional:

$$E_{CLG3-N}(\mathbf{w}) = \int_{\Omega \times [0, T]} (\psi_1(\mathbf{w}^\top J_\rho(\nabla_3 f) \mathbf{w}) + \alpha \psi_2(|\nabla_3 \mathbf{w}|^2)) dx dy dt \quad (25)$$

where $\psi_1(s^2)$ and $\psi_2(s^2)$ are nonquadratic penalisers. Encouraging experiments with related continuous energy functionals have been performed by Hinterberger et al. (2002). Suitable nonquadratic penalisers can be derived from nonlinear diffusion filter design, where preservation or enhancement of discontinuities is also desired (Weickert, 1998). In order to guarantee well-posedness for the remaining problem, we focus only on penalisers that are convex in s . In particular, we use a function that has been proposed by Charbonnier et al. (1994):

$$\psi_i(s^2) = 2\beta_i^2 \sqrt{1 + \frac{s^2}{\beta_i^2}}, \quad i \in 1, 2 \quad (26)$$

where β_1 and β_2 are scaling parameters. Under some technical requirements, the choice of convex penalisers ensures a unique solution of the minimisation problem and allows to construct simple globally convergent algorithms.

The Euler–Lagrange equations of the energy functional (25) are given by

$$0 = \operatorname{div} (\psi_2'(|\nabla_3 \mathbf{w}|^2) \nabla_3 u) - \frac{1}{\alpha} \psi_1'(\mathbf{w}^\top J_\rho(\nabla_3 f) \mathbf{w}) (J_{11} u + J_{12} v + J_{13}), \quad (27)$$

$$0 = \operatorname{div} (\psi_2'(|\nabla_3 \mathbf{w}|^2) \nabla_3 v) - \frac{1}{\alpha} \psi_1'(\mathbf{w}^\top J_\rho(\nabla_3 f) \mathbf{w}) (J_{21} v + J_{22} u + J_{23}). \quad (28)$$

with

$$\psi_i'(s^2) = \frac{1}{\sqrt{1 + \frac{s^2}{\beta_i^2}}}, \quad i \in 1, 2 \quad (29)$$

One should note that for large values of β_i the nonlinear case comes down to the linear one since $\psi_i'(s^2) \approx 1$.

5. Multiresolution Approach

All variants of the CLG method considered so far are based on a linearisation of the grey value constancy assumption. As a consequence, u and v are required to be relatively small so that the linearisation holds. Obviously, this cannot be guaranteed for arbitrary sequences. However, there are strategies that allow to overcome this limitation. These so called *multiscale focusing* or *multiresolution* techniques (Anandan, 1989; Black and Anandan, 1996; Mémin and Pérez, 1998; Mémin and Pérez, 2002) incrementally compute the optic flow field based on a sophisticated coarse-to-fine strategy:

Starting from a coarse scale the resolution is refined step by step. However, the estimated flow field at a coarser level is *not used* as initialisation at the next finer scale. In particular for energy functionals with a global minimum, such a proceeding would only lead to an acceleration of the convergence, since the result would not change. Instead, the coarse scale motion is used to warp (correct) the original sequence before going to the next finer level. This compensation for the already computed motion results in a hierarchy of modified problems that only require to compute small displacement fields, the so called *motion increments*. Thus it is not surprising that the final displacement field obtained by a summation of all motion increments is much more accurate regarding the linearisation of the grey value constancy assumption.

Let $\delta \mathbf{w}^m$ denote the motion increment at resolution level m , where $m = 0$ is the coarsest level with initialisation $\mathbf{w}^0 = (0, 0, 0)^\top$. Then $\delta \mathbf{w}^m$ is obtained by optimisation of the following spatiotemporal energy functional:

$$\begin{aligned} E_{CLG3-N}^m(\delta \mathbf{w}^m) &= \int_{\Omega \times [0, T]} (\psi_1(\delta \mathbf{w}^{m\top} J_\rho(\nabla_3 f(\mathbf{x} + \mathbf{w}^m)) \delta \mathbf{w}^m) \\ &\quad + \alpha \psi_2(|\nabla_3(\mathbf{w}^m + \delta \mathbf{w}^m)|^2)) \, d\mathbf{x} \end{aligned}$$

where $\mathbf{w}^{m+1} = \mathbf{w}^m + \delta \mathbf{w}^m$ and $\mathbf{x} = (x, y, t)$. One should note that warping the original sequence does only affect the data term. Since the smoothness assumption applies to the complete flow field, $\mathbf{w}^m + \delta \mathbf{w}^m$ is used as argument of the penaliser.

If we denote the structure tensor of the corrected sequence by $J_\rho^m = J_\rho(\nabla_3 f(\mathbf{x} + \mathbf{w}^m))$, the corresponding Euler–Lagrange equations are given by

$$\begin{aligned} 0 &= \operatorname{div}(\psi_2'(|\nabla_3(\mathbf{w}^m + \delta \mathbf{w}^m)|^2) \nabla_3 \delta u^m) \\ &\quad - \frac{1}{\alpha} \psi_1'(\delta \mathbf{w}^{m\top} J_\rho^m \delta \mathbf{w}) (J_{11}^m \delta u + J_{12}^m \delta v + J_{13}^m), \end{aligned} \quad (30)$$

$$\begin{aligned} 0 &= \operatorname{div}(\psi_2'(|\nabla_3(\mathbf{w}^m + \delta \mathbf{w}^m)|^2) \nabla_3 \delta v^m) \\ &\quad - \frac{1}{\alpha} \psi_1'(\delta \mathbf{w}^{m\top} J_\rho^m \delta \mathbf{w}) (J_{21}^m \delta v + J_{22}^m \delta u + J_{23}^m). \end{aligned} \quad (31)$$

6. Algorithmic Realisation

6.1. Spatial and Spatiotemporal Approach

Let us now discuss a suitable algorithm for the CLG method (18) and (19) and its spatiotemporal variant. To this end we consider the unknown functions $u(x, y, t)$ and $v(x, y, t)$ on a rectangular pixel grid of size h , and we denote by u_i the approximation to u at some pixel i with $i = 1, \dots, N$. Gaussian convolution is realised in the spatial/spatiotemporal domain by discrete convolution with a truncated and renormalised Gaussian, where the truncation took place at 3 times the standard deviation. Symmetry and separability has been exploited in order to speed up these discrete convolutions. Spatial derivatives of the image data have been approximated using a sixth-order approximation with the stencil $(-1, 9, -45, 0, 45, -9, 1)/(60h)$. Temporal derivatives are either approximated with a sim-

ple two-point stencil or the fifth-order approximation $(-9, 125, -2250, 2250, -125, 9)/(1920h)$.

Let us denote by J_{nmi} the component (n, m) of the structure tensor $J_\rho(\nabla f)$ in some pixel i . Furthermore, let $\mathcal{N}(i)$ denote the set of (4 in 2-D, 6 in 3-D) neighbours of pixel i . Then a finite difference approximation to the Euler–Lagrange equations (18)–(19) is given by

$$0 = \sum_{j \in \mathcal{N}(i)} \frac{u_j - u_i}{h^2} - \frac{1}{\alpha} (J_{11i} u_i + J_{12i} v_i + J_{13i}), \quad (32)$$

$$0 = \sum_{j \in \mathcal{N}(i)} \frac{v_j - v_i}{h^2} - \frac{1}{\alpha} (J_{21i} u_i + J_{22i} v_i + J_{23i}) \quad (33)$$

for $i = 1, \dots, N$. This sparse linear system of equations may be solved iteratively. The *successive overrelaxation (SOR) method* (Young, 1971) is a good compromise between simplicity and efficiency. If the upper index denotes the iteration step, the SOR method can be written as

$$\begin{aligned} u_i^{k+1} &= (1 - \omega) u_i^k + \omega \\ &\quad \frac{\sum_{j \in \mathcal{N}^-(i)} u_j^{k+1} + \sum_{j \in \mathcal{N}^+(i)} u_j^k - \frac{h^2}{\alpha} (J_{12i} v_i^k + J_{13i})}{|\mathcal{N}(i)| + \frac{h^2}{\alpha} J_{11i}}, \end{aligned} \quad (34)$$

$$\begin{aligned} v_i^{k+1} &= (1 - \omega) v_i^k + \omega \\ &\quad \frac{\sum_{j \in \mathcal{N}^-(i)} v_j^{k+1} + \sum_{j \in \mathcal{N}^+(i)} v_j^k - \frac{h^2}{\alpha} (J_{21i} u_i^{k+1} + J_{23i})}{|\mathcal{N}(i)| + \frac{h^2}{\alpha} J_{22i}} \end{aligned} \quad (35)$$

where

$$\mathcal{N}^-(i) := \{j \in \mathcal{N}(i) \mid j < i\}, \quad (36)$$

$$\mathcal{N}^+(i) := \{j \in \mathcal{N}(i) \mid j > i\} \quad (37)$$

and $|\mathcal{N}(i)|$ denotes the number of neighbours of pixel i that belong to the image domain.

The *relaxation parameter* $\omega \in (0, 2)$ has a strong influence on the convergence speed. For $\omega = 1$ one obtains the well-known *Gauß–Seidel method*. We usually use values for ω between 1.9 and 1.99. This numerically inexpensive overrelaxation step results in a speed-up by one order of magnitude compared with the Gauß–Seidel approach. We initialised the flow components for the first iteration by 0. The specific choice

of the initialisation is not critical since the method is globally convergent.

It should be noted that the iteration scheme does not require many computations per step, since one may compute expressions of type $\frac{h^2}{\alpha} J_{nmi}$ before entering the iteration loop. Moreover, any practical implementation requires only a single vector of size N for storing each of the two flow components u and v : Since the components are updated sequentially, there is no need for two vectors for the iteration levels k and $k + 1$.

Our CLG method has been implemented in ANSI C. For computing the optic flow between two image frames of size 316×252 on a 1533 GHz Athlon PC, one iteration takes 4 CPU milliseconds. In the 3-D case using 15 frames of the same sequence, one iteration takes 70 CPU milliseconds. For our performance evaluations in Section 8, we used 1000 iterations in the 2-D case and 200 iterations in the 3-D case. Since the iterative process converges fast in the beginning and slows down afterwards, one may get perceptually similar solutions already after significantly less iterations. The memory requirement was 5.9 MB in the 2-D example, and 63 MB in the 3-D case.

some pixel i by $\psi'_{1i} := \psi'_1(\mathbf{w}_i^\top J_{\rho i} \mathbf{w}_i)$ respectively $\psi'_{2i} := \psi'_2(|\nabla_3 \mathbf{w}_i|^2)$. Then the obtained nonlinear system of equations reads

$$0 = \sum_{j \in \mathcal{N}(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} \frac{u_j - u_i}{h^2} - \frac{\psi'_{1i}}{\alpha} (J_{11i} u_i + J_{12i} v_i + J_{13i}), \quad (38)$$

$$0 = \sum_{j \in \mathcal{N}(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} \frac{v_j - v_i}{h^2} - \frac{\psi'_{1i}}{\alpha} (J_{21i} u_i + J_{22i} v_i + J_{23i}) \quad (39)$$

for $i = 1, \dots, N$. One should keep in mind that the nonlinearity results from the dependency of ψ'_{1i} and ψ'_{2i} on u_i and v_i . In order to remove this nonlinearity, an outer fixed point iteration is applied. Keeping ψ'_{1i} and ψ'_{2i} fixed allows to solve the resulting linear equation system, e.g. by using the SOR method. In this case the iteration step is given by

$$u_i^{k+1} = (1 - \omega) u_i^k + \omega \frac{\sum_{j \in \mathcal{N}^-(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} u_j^{k+1} + \sum_{j \in \mathcal{N}^+(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} u_j^k - \psi'_{1i} \frac{h^2}{\alpha} (J_{12i} v_i^k + J_{13i})}{\sum_{j \in \mathcal{N}(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} + \psi'_{1i} \frac{h^2}{\alpha} J_{11i}}, \quad (40)$$

$$v_i^{k+1} = (1 - \omega) v_i^k + \omega \frac{\sum_{j \in \mathcal{N}^-(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} v_j^{k+1} + \sum_{j \in \mathcal{N}^+(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} v_j^k - \psi'_{1i} \frac{h^2}{\alpha} (J_{21i} u_i^{k+1} + J_{23i})}{\sum_{j \in \mathcal{N}(i)} \frac{\psi'_{2i} + \psi'_{2j}}{2} + \psi'_{1i} \frac{h^2}{\alpha} J_{22i}}. \quad (41)$$

By considering so called multigrid techniques instead of the proposed SOR method even real-time performance can be achieved with the CLG approach (Bruhn et al., 2003). However, since such advanced numerical schemes have to be developed explicitly for a given problem, their implementation is far more difficult than in the case of the SOR method.

6.2. Nonlinear and Multiresolution Approach

In the nonlinear case, the discretisation of the Euler-Lagrange Eqs. (27) and (28) is straightforward. For the sake of clarity let us denote the derivatives of the penalising functions ψ_1 and ψ_2 at

After solving the equation system ψ'_{1i} and ψ'_{2i} are updated using the new values of u_i and v_i . This procedure is repeated until a fixed point in ψ'_{1i} and ψ'_{2i} is reached.

In the case of the multiresolution framework the proceeding is similar. In order to discretise the Euler-Lagrange equations (30) and (31) at level m we define ψ_1 and ψ_2 at some pixel i by $\psi_1^m := \psi_1(\delta \mathbf{w}_i^m \mathbf{w}_i^m)$ respectively $\psi_2^m := \psi_2(|\nabla_3 \mathbf{w}_i^m + \delta \mathbf{w}_i^m|^2)$. The obtained nonlinear system of equations then reads

$$0 = \sum_{j \in \mathcal{N}(i)} \frac{\psi_1^m + \psi_2^m}{2} \frac{\delta u_j^m - \delta u_i^m}{h^2} - \frac{\psi_1^m}{\alpha} (J_{11i}^m \delta u_i^m + J_{12i}^m \delta v_i + J_{13i}^m), \quad (42)$$

$$0 = \sum_{j \in \mathcal{N}(i)} \frac{\psi'_{2i}{}^m + \psi'_{2j}{}^m}{2} \frac{\delta v_j^m - \delta v_i^m}{h^2} - \frac{\psi'_{1i}{}^m}{\alpha} (J_{21i}^m \delta u_i^m + J_{22i}^m \delta v_i + J_{23i}^m). \quad (43)$$

for $i = 1, \dots, N$. Again, a fixed point iteration is applied to remove the nonlinearity of $\psi'_{1i}{}^m$ and $\psi'_{2i}{}^m$ in δu_i^m and δv_i^m . Although iteration steps for the SOR method are not given explicitly, they can easily be derived from Eqs. (42) and (43). When a fixed point in $\psi'_{1i}{}^m$ and $\psi'_{2i}{}^m$ is reached the original sequence $f(\mathbf{x})$ is warped to $f(\mathbf{x} + \mathbf{w}^{m+1})$ by means of a backward registration based on bilinear interpolation. Structure tensor entries have also to be recomputed before the multiresolution strategy continues on the next finer level given by $m+1$.

7. A Confidence Measure for Energy-Based Methods

While global optic flow methods typically yield dense flow fields, it is clear that the flow estimates can-

not have the same reliability at all locations. Local methods, on the other hand, have natural confidence measures that help to avoid computing flow values at locations where there is not enough information for a reliable estimate. It would thus be interesting to find a confidence measure that allows to assess the reliability of a dense optic flow field. Barron et al. (1994) have identified the absence of such good measure as one of the main drawbacks of energy-based global optic flow techniques: Simple heuristics such as using $|\nabla f|$ as a confidence measure did not work well.

In order to address this problem, we propose a measure that may be applied to any energy-based global differential method for computing the optic flow: Since the energy functional E penalises deviations E_i from all pixels i in the image domain, it appears natural to use E_i for assessing the local reliability of the computation. All we have to do is to consider the cumulative histogram of the contributions E_i with $i = 1, \dots, N$. As an approximation to the p per cent locations with




Figure 2. (a) Top left: Frame 10 of the synthetic office sequence. (b) Top right: Degraded by Gaussian noise with $\sigma_n = 20$. (c) Bottom left: Ground truth optic flow field. (d) Bottom right: Computed optic flow field using the 2-D CLG method for the noisy sequence.

the highest reliability, we look for the p per cent locations where the contribution E_i is lowest. There are very efficient algorithms available for

this purpose; see e.g. (Press, 1992, Section 8.5). In the next section we shall observe that this simple criterion may work well over a large range of densities.

8. Experiments

8.1. Evaluation of the CLG Method

Figure 2 shows our first experiment. It depicts a zoom into a synthetic office scene where divergent motion is dominating. This test sequence is

available from the web site www.cs.otago.ac.nz/research/vision/. It has been created by Galvin et al. (1998). We have added Gaussian noise with zero mean and standard deviation $\sigma_n = 20$ to this sequence, and we used the 2-D CLG method (i.e. with spatial regularisation) for computing the flow field. Figure 2(d) shows that the recovered flow field is not very sensitive to Gaussian noise and that it coincides well with the ground truth flow field in Fig. 2(c).

These qualitative results are confirmed by the quantitative evaluations in Table 1, where we compare the average angular errors of the Lucas–Kanade, Horn–Schunck, and the CLG method for different noise levels and optimised smoothing parameters σ , ρ , and α .




Figure 3. (a) Top left: Frame 8 of the Yosemite sequence severely degraded by Gaussian noise with $\sigma_n = 40$. (b) Top right: Ground truth flow field. (c) Middle left: Computed flow field for $\sigma_n = 0$. (d) Middle right: Ditto for $\sigma_n = 10$. (e) Bottom left: $\sigma_n = 20$. (f) Bottom right: $\sigma_n = 40$.

Table 1. Average angular errors computed for the *office* sequence with varying standard deviations σ_n of Gaussian noise. 2-D implementations of the methods of Lucas/Kanade (LK), Horn/Schunck (HS) and the combined local-global approach (CLG) are compared.

σ_n	LK	HS	CLG
0	5.71°	4.36°	4.32°
10	6.79°	6.17°	5.89°
20	8.43°	8.30°	7.75°
40	11.47°	11.76°	10.73°

We computed the angular error via

$$\arccos\left(\frac{u_c u_e + v_c v_e + 1}{\sqrt{(u_c^2 + v_c^2 + 1)(u_e^2 + v_e^2 + 1)}}\right) \quad (44)$$

where (u_c, v_c) denotes the correct flow, and (u_e, v_e) is the estimated flow (cf. also (Barron et al., 1994)).

Table 1 shows that for small noise levels, Horn–Schunck performs better than Lucas–Kanade. This indicates that the filling-in effect of the Horn–Schunck algorithm is very useful here. For higher noise levels, Lucas–Kanade becomes somewhat more robust than Horn–Schunck, since the former does not reduce smoothing at noisy structures. The CLG method appears to be able to pick up the best of two worlds: It may benefit from filling-in effects in flat regions without renouncing robustness against noise. Table 1 shows that this combined effect leads to results that may be better than both the Lucas–Kanade and the Horn–Schunck method.

Another example demonstrating the robustness of the 2-D CLG method under Gaussian noise is shown in Fig. 3. It depicts the results for the synthetic Yosemite sequence with cloudy sky. This sequence, which is available from <ftp://csd.uwo.ca> under the directory `pub/vision`, combines divergent motion with the translational motion of the sky. It has been used by Barron et al. (1994) for evaluating a number of optic flow algorithms. Also in this example we can observe that the flow computations using the CLG method do hardly suffer from severe degradations by Gaussian noise.

Let us now investigate the sensitivity of the CLG method with respect to parameter variations. This is done in Tables 2 and 3 for the *Yosemite* sequence with and without clouds. The modified variant without cloudy sky is available from the web site <http://www.cs.brown.edu/people/black/images.html>. We observe that the average angular error does hardly deteriorate when smoothness param-

Table 2. Stability of the 2-D CLG method under variations of the smoothing parameters. Two of the three parameters have been set to their optimal value, while the other *one* may deviate from its optimum by a factor 2. The data refer to the *Yosemite* sequence with and without clouds degraded by Gaussian noise with $\sigma_n = 10$. AAE = average angular error.

Yosemite with clouds				Yosemite without clouds			
α	ρ	σ	AAE	α	ρ	σ	AAE
475	4.550	1.770	9.31°	1000	4.550	1.950	4.57°
633	"	"	9.23°	1666	"	"	4.44°
950	"	"	9.18°	2000	"	"	4.43°
1425	"	"	9.24°	3000	"	"	4.54°
1900	"	"	9.37°	4000	"	"	4.79°
950	2.275	1.770	9.25°	2000	2.275	1.950	4.46°
"	3.033	"	9.21°	"	3.033	"	4.44°
"	4.550	"	9.18°	"	4.550	"	4.43°
"	6.825	"	9.24°	"	6.825	"	4.49°
"	9.100	"	9.39°	"	9.100	"	4.62°
950	4.550	0.885	13.65°	2000	4.550	0.975	7.48°
"	"	1.180	10.58°	"	"	1.300	5.39°
"	"	1.770	9.18°	"	"	1.950	4.43°
"	"	2.655	10.24°	"	"	2.975	5.60°
"	"	3.540	12.30°	"	"	3.800	7.09°

Table 3. Stability of the 2-D CLG method under variations of the smoothing parameters. All three parameters have been set to a deviation of factor 2 from its optimum value. The data refer to the *Yosemite* sequence with and without clouds degraded by Gaussian noise with $\sigma_n = 10$. AAE = average angular error.

Yosemite with clouds				Yosemite without clouds			
α	ρ	σ	AAE	α	ρ	σ	AAE
475	2.275	0.885	14.72°	1000	2.275	0.975	8.37°
1900	2.275	0.885	13.48°	4000	2.275	0.975	7.27°
475	9.100	0.885	13.94°	1000	2.275	0.975	7.58°
1900	9.100	0.885	13.49°	4000	2.275	0.975	7.16°
475	2.275	3.540	11.92°	1000	9.100	3.800	6.48°
1900	2.275	3.540	13.71°	4000	9.100	3.800	8.57°
475	9.100	3.540	11.96°	1000	9.100	3.800	6.59°
1900	9.100	3.540	13.65°	4000	9.100	3.800	8.64°

eters are used that differ from their optimal setting by as much as a factor 2. Only the noise scale σ , that is responsible for the presmoothing of the original sequence, is slightly more sensitive. This observation is in accordance with the results presented in Table 3 where in addition to the noise scale also the remaining

Table 4. Results for the 2-D and 3-D CLG method using the *Yosemite* sequence with and without cloudy sky. Gaussian noise with varying standard deviations σ_n was added, and the average angular errors and their standard deviations were computed.

Yosemite with clouds			Yosemite without clouds		
σ_n	2-D CLG	3-D CLG	σ_n	2-D CLG	3-D CLG
0	$7.14^\circ \pm 9.28^\circ$	$6.18^\circ \pm 9.19^\circ$	0	$2.64^\circ \pm 2.27^\circ$	$1.79^\circ \pm 2.34^\circ$
10	$9.19^\circ \pm 9.62^\circ$	$7.25^\circ \pm 9.39^\circ$	10	$4.45^\circ \pm 2.94^\circ$	$2.53^\circ \pm 2.75^\circ$
20	$10.17^\circ \pm 10.50^\circ$	$8.62^\circ \pm 9.97^\circ$	20	$6.93^\circ \pm 4.31^\circ$	$3.47^\circ \pm 3.37^\circ$
40	$15.82^\circ \pm 11.53^\circ$	$11.21^\circ \pm 11.19^\circ$	40	$11.30^\circ \pm 7.41^\circ$	$5.34^\circ \pm 3.81^\circ$

two smoothness parameters differ from their optimal setting. The stability under parameter variations may be regarded as another experimental confirmation of the well-posedness of the CLG approach. Moreover, this also indicates that the method performs sufficiently robust in practice even if non-optimised default parameter settings are used.

In Table 4 we have studied the effect of replacing spatial smoothing steps by spatiotemporal ones for both *Yosemite* sequences. As one may expect, both the quality of the optic flow estimates and their robustness under Gaussian noise improve when temporal coherence is taken into account.

In Section 4 we have presented nonlinear variants of our spatial and spatiotemporal CLG approaches. Our next experiment compares these nonlinear versions to their linear counterparts for a variety of sequences. Generally, flow discontinuities do not cover more than a few percent of the estimated flow field, so only moderate improvements should be expected. These considerations are confirmed by Table 5 where the computed average angular errors are listed. A qualitative example of such a nonlinear variant is given in Fig. 4. The *Marble* sequence that was used for this purpose can be downloaded at the following internet address: http://i21www.ira.uk.de/image_sequences. It

Table 5. Average angular error for linear and nonlinear variants of our CLG method using various sequences.

Sequence	2-D		3-D	
	Linear	Nonlin.	Linear	Nonlin.
Yosemite (clouds)	7.14°	6.03°	6.18°	5.18°
Yosemite (no clouds)	2.64°	2.31°	1.79°	1.46°
Office	4.33°	4.13°	3.60°	3.24°
Marble	5.30°	5.14°	2.06°	1.70°

is easy to see that the flow discontinuities are much better preserved using the nonlinear variant. Values for the scaling parameters β_i within the nonlinear penalising functions have been optimised and are between $5 \cdot 10^{-2}$ and $5 \cdot 10^{-3}$ for all sequences.

Comparisons with other methods from the literature that yield dense flow fields are shown in Tables 6 and 7. We observe that for the *Yosemite* variant *with* clouds both the linear 2-D and the 3-D version of our method already perform favourably compared to other techniques that do not use multiscale focusing strategies. In the case of the *Yosemite* sequence *without* clouds nonlinear extensions are required to achieve competitive results with respect to recent techniques. Since the *Yosemite* sequence contains large displacements up to 5 pixels per frame, one should expect further improvements if the proposed multiresolution framework is applied. This is confirmed by average angular errors of 1.04° and 4.17° for the sequence without respectively with cloudy sky, thus outperforming all results from the literature reported so far. A qualitative evaluation of these numbers is given in Fig. 5, where the computed flow fields by the spatial and the spatiotemporal variant of our nonlinear multiresolution approach are presented for both *Yosemite* sequences. As one can see, the obtained flow estimates coincide very well with the corresponding ground truths.

So far we have considered only synthetic image sequences for our experimental evaluation. In order to show that the proposed approach yields also realistic optic flow estimates for real-world data, the *Ettlinger Tor* traffic sequence by Nagel is used for our last experiment. The sequence consists of 50 frames of size 512×512 and is available from http://i21www.ira.uka.de/image_sequences/. In contrast to our previous experiments, where always the complete set of frames was used for our




Figure 4. (a) *Top left*: Frame 10 of the *Marble* sequence. (b) *Top right*: Ground truth magnitude. (c) *Bottom left*: Computed flow field magnitude using linear 2-D CLG method. (d) *Bottom right*: Ditto for the nonlinear 3-D CLG variant.

spatiotemporal version, this time only four frames are considered. Thus the dislocation of motion boundaries remains within a reasonable scope, while the algorithm is able to benefit from the spatiotemporal denoising property at the same time.

The computed flow fields as well as its magnitude are shown in Fig. 6. Although interlacing artifacts are present in all frames, the estimated flow field is very realistic and the motion boundaries are rather sharp. In particular, this allows its use for the purpose of image segmentation. Obviously, only a simple thresholding step would be required.

8.2. Evaluation of the Confidence Measure

Let us now evaluate the quality of our energy-based confidence measure. To this end we have depicted in Fig. 7(a) the 20% quantile of locations where the 3-D CLG method has lowest contributions to the energy.

A comparison with Fig. 7(b)—which displays the result of a theoretical confidence measure that would be optimal with respect to the average angular error—demonstrates that the energy-based confidence method leads to a fairly realistic thinning of flow fields. In particular, we observe that this confidence criterion is very successful in removing the cloudy sky regions. These locations are well-known to create large angular errors in many optic flow methods (Barron et al., 1994). A number of authors have thus only used the modified *Yosemite* sequence without cloudy sky, or they have neglected the flow values from the sky region for their evaluations (Bab-Hadiashar and Suter, 1998; Black and Anandan, 1996; Black and Jepson, 1996; Farnebäck, 2000, 2001; Ju et al., 1996; Karlholm, 1998; Lai and Vemuri, 1998; Szeliski and Coughlan, 1994). As we have seen one may get significantly lower angular errors than for the full sequence with cloudy sky.

Table 6. Comparison between the results from the literature with 100% density and our results. All data refer to the *Yosemite* sequence *without* cloudy sky. Multiscale means that some focusing strategy using linear scale-space or pyramids has been applied. Spatiotemporal Information indicates the use of several frames for presmoothing and/or derivative computation. AAE = average angular.

Technique	Multiscale	Spatiotemporal information	Spatiotemporal constraint	AAE
Black and Anandan (1996)	✓	–	–	4.56°
Black (1994)	✓	–	✓	3.52°
Lauze et al. (2004)	–	–	–	2.82°
2-D CLG linear	–	–	–	2.64°
Szeliski and Coughlan (1994)	–	–	✓	2.45°
3-D CLG linear	–	✓	✓	2.31°
Black and Jepson (1996)	✓	–	–	2.29°
Ju et al. (1996)	–	–	–	2.16°
Bab-Hadiashar and Suter (1998)	–	✓	–	2.05°
2-D CLG nonlinear	–	–	–	1.79°
2-D CLG nonlinear multires.	✓	–	–	1.62°
Mémin and Pérez (2002)	✓	–	–	1.58°
3-D CLG nonlinear	–	✓	✓	1.46°
Farneback (2000)	–	✓	–	1.40°
Farneback (2001)	–	✓	–	1.14°
3-D CLG nonlinear multires.	✓	✓	✓	1.02°

Table 7. Comparison between the results from the literature with 100% density and our results. All data refer to the *Yosemite* sequence *with* cloudy sky. Multiscale means that some focusing strategy using linear scale-space or pyramids has been applied. Spatiotemporal Information indicates the use of several frames for presmoothing and/or derivatives computation. AAE = average angular.

Technique	Multiscale	Spatiotemporal information	Spatiotemporal constraint	AAE
Horn/Schunck, original (Barron et al., 1994)	–	✓	–	31.69°
Singh, step 1 (Barron et al., 1994)	–	–	–	15.28°
Anandan (Barron et al., 1994)	–	–	–	13.36°
Singh, step 2 (Barron et al., 1994)	–	–	–	10.44°
Nagel (Barron et al., 1994)	–	✓	–	10.22°
Horn/Schunck, modified (Barron et al., 1994)	–	✓	–	9.78°
Uras et al., unthresholded (Barron et al., 1994)	–	✓	–	8.94°
2-D CLG linear	–	–	–	7.09°
3-D CLG linear	–	✓	✓	6.24°
2-D CLG nonlinear	–	–	–	6.03°
Alvarez et al. (2000)	✓	–	–	5.53°
Mémin and Pérez (1998)	✓	–	–	5.38°
3-D CLG nonlinear	–	✓	✓	5.18°
2-D CLG nonlinear multires	✓	–	–	4.86°
Mémin and Pérez (1998)	✓	–	–	4.69°
3-D CLG nonlinear multires	✓	✓	✓	4.17°




Figure 5. From left to right, and from top to bottom: (a) Frame 8 of the Yosemite sequence without clouds. (b) Corresponding frame of the sequence with clouds. (c) Ground truth between frame 8 and 9 for the sequence without clouds. (d) Ditto for the sequence with clouds. (e) Computed flow field by our 2-D CLG multiresolution approach for the sequence without clouds. (f) Ditto for the sequence with clouds. (g) Computed flow field by our 3-D CLG multiresolution approach for the sequence without clouds. (h) Ditto for the sequence with clouds.

A quantitative evaluation of our confidence measure is given in Table 8. Here we have used the energy-based confidence measure to sparsify the dense flow field such that the reduced density coincides with densities of well-known optic flow methods. Most of them have been evaluated by Barron et al. (1994). We observe that the sparsified 3-D CLG method performs very favourably: It has a far lower angular error than all corresponding methods with the same density. In

several cases there is an order of magnitude between these approaches. At a flow density of 2.4%, an average angular error of 0.76° is reached. To our knowledge, these are the best values that have been obtained for this sequence in the entire literature. It should be noted that these results have been computed from an image sequence that suffers from quantisation errors since its grey values have been stored in 8-bit precision only.




Figure 6. (a) Left: Computed flow field between frame 5 and 6 of the interlaced *Ettlinger Tor* traffic sequence by Nagel. (b) Right: Flow field magnitude.




Figure 7. Confidence criterion for the *Yosemite* sequence. (a) Left: Locations with the lowest contributions to the energy (20% quantile). The non-black grey values depict the optic flow magnitude. (b) Right: Locations where the angular error is lowest (20% quantile).

In Table 8 we also observe that the angular error decreases *monotonically* under sparsification over the entire range from 100% down to 2.4%. This in turn indicates an interesting finding that may seem counter-intuitive at first glance: *Regions in which the filling-in effect dominates give particularly small angular errors.* In such flat regions, the data term vanishes such that a smoothly extended flow field may yield only a small local contribution to the energy functional. If there were large angular errors in regions with such low energy contributions, our confidence measure would not work well for low densities. This also confirms the observation that $|\nabla f|$ is not necessarily a good confidence measure (Barron et al., 1994): Areas with large gradients may represent noise or occlusions, where reliable flow information is difficult to obtain. The filling-in effect, however, may create more reliable information in

flat regions by averaging less reliable information that comes from all the surrounding high-gradient regions. The success of our confidence measure also confirms our previous findings that it is beneficial to supplement local methods with a global regulariser. A more extensive experimental evaluation of the energy based confidence measure is presented in Bruhn (2001).

9. Summary and Conclusions

In this paper we have analysed the smoothing effects in local and global differential methods for optic flow computation. As a prototype of local methods we used the least-square fit of Lucas and Kanade (1981) and Lucas (1984), while the Horn and Schunck approach (Horn and Schunck, 1981) was our representative for a global method.

Table 8. Comparison between the “nondense” results from Barron et al. (1994), Weber and Malik (1995), Ong and Spann (1999) and our results for the *Yosemite* sequence with cloudy sky. AAE = average angular error. CLG = average angular error of the 3-D CLG method with the same density. The sparse flow field has been created using our energy-based confidence criterion. The table shows that using this criterion clearly outperforms all results in the evaluation of Barron et al.

Technique	Density (%)	AAE (°)	CLG (°)
Singh, step 2, $\lambda_1 \leq 0.1$	97.7	10.03	6.04
Ong/Spann	89.9	5.76	5.26
Heeger, level 0	64.2	22.82	3.00
Weber/Malik	64.2	4.31	3.00
Horn/Schunck, original, $ \nabla f \geq 5$	59.6	25.33	2.72
Ong/Spann, thresholded	58.4	4.16	2.66
Heeger, combined	44.8	15.93	2.07
Lucas/Kanade, $\lambda_2 \geq 1.0$	35.1	4.28	1.71
Fleet/Jepson, $\tau = 2.5$	34.1	4.63	1.67
Horn/Schunck, modified, $ \nabla f \geq 5$	32.9	5.59	1.63
Nagel, $ \nabla f \geq 5$	32.9	6.06	1.63
Fleet/Jepson, $\tau = 1.25$	30.6	5.28	1.55
Heeger, level 1	15.2	9.87	1.15
Uras et al., $\det(H) \geq 1$	14.7	7.55	1.14
Singh, step 1, $\lambda_1 \leq 6.5$	11.3	12.01	1.07
Waxman et al., $\sigma_f = 2.0$	7.4	20.05	0.95
Heeger, level 2	2.4	12.93	0.76

We saw that the smoothing steps in each of these methods serve different purposes and have different advantages and shortcomings. As a consequence, we proposed a combined local-global (CLG) approach that incorporates the advantages of both paradigms: It is highly robust under Gaussian noise while giving dense flow fields. In order to improve the performance of our method further we considered spatiotemporal variants, nonlinear penalising functions that are well-known from robust statistics and multiscale focusing strategies that allow for an correct handling of large image displacements. Experiments have shown that the CLG method is not very sensitive under parameter variations. This method may serve as an example of how one can supplement local methods with a regulariser such that dense flow fields are obtained. As is shown in the Appendix, it can also be extended to the embedding of Bigün’s structure tensor method into a global energy functional.

We have also proposed a simple confidence measure that allows to sparsify the dense flow fields of energy-

based global methods, such that the most reliable local estimates can easily be found. This enables fair comparisons of the quality of local and global approaches. Our evaluations have shown that the proposed confidence measure may give excellent results over a large range of densities. Last but not least, its success has triggered a surprising finding: For global energy-based optic flow methods, flat regions in which the filling-in effect dominates may offer particularly reliable flow estimates. This explains the common observation that the image gradient magnitude is not a good confidence measure for global variational methods.

While we have already taken efforts to use efficient numerical methods, we have certainly not reached the end of the road yet. Therefore, we are currently investigating multigrid implementations of our technique and we are studying parallelisation possibilities on low latency networks.

Appendix: Extension to the Structure Tensor Method

The CLG approach can be extended in a straightforward way to the embedding of basically any local differential method into a global energy functional. Let us illustrate this general principle by focusing on another popular local method: the structure tensor approach of Bigün and Granlund (1988) and Bigün et al. (1991).

In Section 3.2 we have used a spatiotemporal variant of the Lucas–Kanade technique for the temporal extension of our CLG functional. This method is closely related to the approach of Bigün and Granlund (1988) and Bigün et al. (1991). While Lucas and Kanade make use of a least square fit to overcome the aperture problem, Bigün et al. follow a slightly different strategy: They minimise the quadratic form

$$E_{BG}(\tilde{w}) = \tilde{w}^\top J_\rho(\nabla_3 f) \tilde{w} \quad (45)$$

where $\tilde{w} := (\tilde{u}, \tilde{v}, r)^\top$, and the normalisation constraint

$$\tilde{w}^\top \tilde{w} = 1. \quad (46)$$

has to be fulfilled. This is achieved by searching for the eigenvector \tilde{w} that corresponds to the smallest eigenvalue of the structure tensor $J_\rho(\nabla_3 f)$. Normalising its third component to 1 yields $u = \frac{\tilde{u}}{r}$ and $v = \frac{\tilde{v}}{r}$ as the first two components.

In order to combine the local approach of Bigün et al. with some global differential optic flow technique, the method can be reformulated as an unconstrained minimisation of the local energy

$$E_{BG}(w) = \frac{w^\top J_\rho(\nabla_3 f) w}{|w|^2}. \quad (47)$$

This reformulation allows a comparison of (15) and (47), which shows that both types of least square fits differ only by the normalisation factor $\frac{1}{|w|^2}$. The CLG functional obtained by the embedding of Bigün's method in the spatiotemporal variant of Horn and Schunck has the following structure:

$$\begin{aligned} E_{CLG3-B}(w) &= \int_{\Omega \times [0, T]} \left(\frac{w^\top J_\rho(\nabla_3 f) w}{|w|^2} + \alpha |\nabla_3 w|^2 \right) dx dy dt. \end{aligned} \quad (48)$$

Its corresponding Euler–Lagrange equations are given by

$$\begin{aligned} 0 &= \Delta_3 u - \frac{1}{\alpha(u^2 + v^2 + 1)^2} (J_{11}(uv^2 + u) \\ &\quad + J_{12}(v^3 - u^2v + v) + J_{13}(v^2 - u^2 + 1) \\ &\quad - J_{22}uv^2 - 2J_{23}uv - J_{33}u), \quad (49) \\ 0 &= \Delta_3 v - \frac{1}{\alpha(u^2 + v^2 + 1)^2} (-J_{11}u^2v \\ &\quad + J_{12}(u^3 - uv^2 + u) - 2J_{13}uv + J_{22}(u^2v + v) \\ &\quad + J_{23}(u^2 - v^2 + 1) - J_{33}v). \quad (50) \end{aligned}$$

These nonlinear equations are somewhat more complicated than their linear Lucas–Kanade counterpart (22)–(23).

In order to encourage discontinuity-preserving optic flow fields, one can also introduce nonquadratic penalisers into the functional (48). This yields

$$\begin{aligned} E_{CLG3-BN}(w) &= \int_{\Omega \times [0, T]} \left(\psi_1 \left(\frac{w^\top J_\rho(\nabla_3 f) w}{|w|^2} \right) \right. \\ &\quad \left. + \alpha \psi_2(|\nabla_3 w|^2) \right) dx dy dt. \end{aligned} \quad (51)$$

Such a strategy may be regarded as an alternative to the discontinuity-preserving structure tensor methods in Brox and Weickert (2002) and Nagel and Gehrke (1998). It gives dense flow fields without additional postprocessing steps.

Acknowledgments

Our optic flow research is partly funded by the *Deutsche Forschungsgemeinschaft (DFG)* under the project WE 2602/3-1. This is gratefully acknowledged.

Notes

1. A problem is called *well-posed* in the sense of Hadamard, if it has a unique solution that depends continuously on the data. If one of these conditions is violated, it is called *ill-posed*.
2. The taxi sequence is available from <ftp://csd.uwo.ca> under the directory `pub/vision`

References

- 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*, Las Palmas de Gran Canaria, Spain. pp. 1349–1356.
- Alvarez, L., Weickert, J., and Sánchez, J. 2000. Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision*, 39(1):41–56.
- Anandan, P. 1989. A computational framework and an algorithm for the measurement of visual motion. *International Journal of Computer Vision*, 2:283–310.
- Aubert, G., Deriche, R., and Kornprobst, P. 1999. Computing optical flow via variational techniques. *SIAM Journal on Applied Mathematics*, 60(1):156–182.
- Bab-Hadiashar, A. and Suter, D. 1998. Robust optic flow computation. *International Journal of Computer Vision*, 29(1):59–77.
- Bainbridge-Smith, A. and Lane, R.G. 1997. Determining optical flow using a differential method. *Image and Vision Computing*, 15(1):11–22.
- Barron, J.L., Fleet, D.J., and Beauchemin, S.S. 1994. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77.
- Bertero, M., Poggio, T.A., and Torre, V. 1988. Ill-posed problems in early vision. *Proceedings of the IEEE*, 76(8):869–889.
- Bigün, J. and Granlund, G.H. 1988. Optical flow based on the inertia matrix in the frequency domain. In *Proc. SSAB Symposium on Picture Processing*, Lund, Sweden.
- Bigün, J., Granlund, G.H., and Wiklund, J. 1991. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790.
- Black, M.J. 1994. Recursive non-linear estimation of discontinuous flow fields. In *Computer Vision, ECCV'94*, J.-O. Eklundh (Ed.), vol. 800 of *Lecture Notes in Computer Science*, Springer: Berlin, pp. 138–145.
- Black, M.J. and Anandan, P. 1991. Robust dynamic motion estimation over time. In *Proc. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society Press: Maui, HI, 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(1):75–104.
- Black, M.J. and Jepson, A. 1996. Estimating optical flow in segmented images using variable-order parametric models with local deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(10):972–986.
- Brox, T. and Weickert, J. 2002. Nonlinear matrix diffusion for optic flow estimation. In *Pattern Recognition*, L. Van Gool, (Ed.), vol. 2449 of *Lecture Notes in Computer Science*, Springer: Berlin, pp. 446–453.
- Bruhn, A. 2001. Regularization in motion estimation. Master's thesis, Department of Mathematics and Computer Science, University of Mannheim, Germany.
- Bruhn, A., Weickert, J., Feddern, C., Kohlberger, T., and Schnörr, C. 2003. Real-time optic flow computation with variational methods. In *Computer Analysis of Images and Patterns*, N. Petkov and M.A. Westberg (Eds.), vol. 2756 of *Lecture Notes in Computer Science*, Springer: Berlin, pp. 222–229.
- Bruhn, A., Weickert, J., and Schnörr, C. 2002. Combining the advantages of local and global optic flow methods. In *Pattern Recognition*, L. Van Gool, (Ed.), vol. 2449 of *Lecture Notes in Computer Science*, Springer: Berlin, pp. 454–462.
- Charbonnier, P., Blanc-Féraud, L. Aubert, G., and Barlaud, M. 1994. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proc. 1994 IEEE International Conference on Image Processing*, vol. 2. IEEE Computer Society Press: Austin TX, pp. 168–172.
- Cohen, I. 1993. Nonlinear variational method for optical flow computation. In *Proc. Eighth Scandinavian Conference on Image Analysis*, vol. 1, Tromsø, Norway, pp. 523–530.
- Courant, R. and Hilbert, D. 1953. *Methods of Mathematical Physics*, vol. 1. Interscience: New York.
- Elad, M. and Feuer, A. 1998. Recursive optical flow estimation—adaptive filtering approach. *Journal of Visual Communication and Image Representation*, 9(2):119–138.
- Elsgolc, L.E. 1961. *Calculus of Variations*. Pergamon: Oxford.
- Farneback, G. 2000. Fast and accurate motion estimation using orientation tensors and parametric motion models. In *Proc. 15th International Conference on Pattern Recognition*, vol. 1, Barcelona, Spain, pp. 135–139.
- Farneback, G. 2001. Very high accuracy velocity estimation using orientation tensors, parametric motion, and simultaneous segmentation of the motion field. In *Proc. Eighth International Conference on Computer Vision*, vol. 1, IEEE Computer Society Press: Vancouver, Canada, pp. 171–177.
- Fermüller, C., Shulman, D., and Aloimonos, Y. 2001. The statistics of optical flow. *Computer Vision and Image Understanding*, 82(1):1–32.
- Fleet, D.J. and Jepson, A.D. 1990. Computation of component image velocity from local phase information. *International Journal of Computer Vision*, 5(1):77–104.
- Galvin, B., McCane, B., Novins, K., Mason, D., and Mills, S. 1998. Recovering motion fields: An analysis of eight optical flow algorithms. In *Proc. 1998 British Machine Vision Conference*, Southampton, England.
- Hampel, F.R., Ronchetti, E.M., Rousseeuw, P.J., and Stahel, W.A. 1986. *Robust Statistics: The Approach Based on Influence Functions*. MIT Press: Cambridge, MA.
- Heitz, F. and Bouthemy, P. 1993. Multimodal estimation of discontinuous optical flow using Markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(12):1217–1232.
- Hinterberger, W., Scherzer, O., Schnörr, C., and Weickert, J. 2002. Analysis of optical flow models in the framework of calculus of variations. *Numerical Functional Analysis and Optimization*, 23(1/2):69–89.
- Horn, B. and Schunck, B. 1981. Determining optical flow. *Artificial Intelligence*, 17:185–203.
- Huber, P.J. 1981. *Robust Statistics*. Wiley: New York.
- Jähne, B. 2001. *Digitale Bildverarbeitung*. Springer: Berlin.
- Ju, S., Black, M., and Jepson, A. 1996. Skin and bones: Multi-layer, locally affine, optical flow and regularization with transparency. In *Proc. 1996 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society Press: San Francisco, CA, pp. 307–314.
- Karlholm, J. 1998. *local signal models for image sequence analysis*. PhD thesis, Linköping University, Sweden, Dissertation No. 536.
- Kearney, J.K., Thompson, W.B., and Boley, D.L. 1987. Optical flow estimation: An error analysis of gradient-based methods with local optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(2):229–244.
- Kumar, A., Tannenbaum, A.R., and Balas, G.J. 1996. Optic flow: A curve evolution approach. *IEEE Transactions on Image Processing*, 5(4):598–610.
- Lai, S.-H. and Vemuri, B.C. 1998. Reliable and efficient computation of optical flow. *International Journal of Computer Vision*, 29(2):87–105.
- Lauze, F.-B., Kornprobst, P., Lenglet, C., Deriche, R., and Nielsen, M. 2004. Sur quelques méthodes de calcul de flot optique à partir du tenseur de structure: Synthèse et contribution. In *Proc. 14th French Conference on Pattern Recognition and Artificial Intelligence*, Toulouse, France.
- Lucas, B. and Kanade, T. 1981. An iterative image registration technique with an application to stereo vision. In *Proc. Seventh International Joint Conference on Artificial Intelligence*, Vancouver, Canada, pp. 674–679.
- Lucas, B.D. 1984. Generalized image matching by the method of differences. PhD thesis, School of Computer Science, Carnegie-Mellon University, Pittsburgh, PA.
- Mémin, E. and Pérez, P. 1998. Dense estimation and object-based segmentation of the optical flow with robust techniques. *IEEE Transactions on Image Processing*, 7(5):703–719.
- Mémin, E. and Pérez, P. 1998. A multigrid approach for hierarchical motion estimation. In *Proc. 6th International Conference on Computer Vision*, Bombay, India, pp. 933–938.
- Mémin, E. and Pérez, P. 2002. Hierarchical estimation and segmentation of dense motion fields. *International Journal of Computer Vision*, 46(2):129–155.
- Mitiche, A. and Bouthemy, P. 1996. Computation and analysis of image motion: A synopsis of current problems and methods. *International Journal of Computer Vision*, 19(1):29–55.
- Murray, D.W. and Buxton, B.F. 1987. Scene segmentation from visual motion using global optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(2):220–228.
- Nagel, H.-H. 1983. Constraints for the estimation of displacement vector fields from image sequences. In *Proc. Eighth International Joint Conference on Artificial Intelligence*, vol. 2, Karlsruhe, West Germany, pp. 945–951.
- Nagel, H.-H. 1990. Extending the 'oriented smoothness constraint' into the temporal domain and the estimation of derivatives of

- optical flow. In *Computer Vision—ECCV '90*, vol. 427 of *Lecture Notes in Computer Science*, O. Faugeras (Ed.), Springer: Berlin, pp. 139–148.
- Nagel, H.-H. and Gehrke, A. 1998. Spatiotemporally adaptive estimation and segmentation of OF-fields. In *Computer Vision—ECCV '98*, vol. 1407 of *Lecture Notes in Computer Science*, H. Burkhardt and B. Neumann, (Eds.), Springer: Berlin, pp. 86–102.
- Nesi, P. 1993. Variational approach to optical flow estimation managing discontinuities. *Image and Vision Computing*, 11(7):419–439.
- Ohta, N. 1996. Uncertainty models of the gradient constraint for optical flow computation. *IEICE Transactions on Information and Systems*, E79-D(7):958–962.
- Ong, E.P. and Spann, M. 1999. Robust optical flow computation based on least-median-of-squares regression. *International Journal of Computer Vision*, 31(1):51–82.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., and Flannery, B.P. 1992. *Numerical Recipes in C*, 2nd ed. Cambridge University Press: Cambridge, UK.
- Proesmans, M., Van Gool, L., Pauwels, E., and Oosterlinck, A. 1994. Determination of optical flow and its discontinuities using non-linear diffusion. In J.-O. Eklundh, (Ed.), *Computer Vision—ECCV '94*, volume 801 of *Lecture Notes in Computer Science*. Springer: Berlin, pp. 295–304.
- Schnörr, C. 1991. Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class. *International Journal of Computer Vision*, 6(1):25–38.
- Schnörr, C. 1993. On functionals with greyvalue-controlled smoothness terms for determining optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15:1074–1079.
- Schnörr, C. 1994. Segmentation of visual motion by minimizing convex non-quadratic functionals. In *Proc. Twelfth International Conference on Pattern Recognition*, vol. A, IEEE Computer Society Press: Jerusalem, Israel, pp. 661–663.
- Shulman, D. and Hervé, J. 1989. Regularization of discontinuous flow fields. In *Proc. Workshop on Visual Motion*, IEEE Computer Society Press: Irvine, CA, pp. 81–90.
- Simoncelli, E.P., Adelson, E.H., and Heeger, D.J. 1991. Probability distributions of optical flow. In *Proc. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society Press: Maui, HI, pp. 310–315.
- Stiller, C. and Konrad, J. 1999. Estimating motion in image sequences. *IEEE Signal Processing Magazine*, 16:70–91.
- Szeliski, R. and Coughlan, J. 1994. Hierarchical spline-based image registration. In *Proc. 1994 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society Press: Seattle, WA, pp. 194–201.
- Tretiak, O. and Pastor, L. 1984. Velocity estimation from image sequences with second order differential operators. In *Proc. Seventh International Conference on Pattern Recognition*, Montreal, Canada, pp. 16–19.
- Uras, S., Girosi, F., Verri, A., and Torre, V.A., 1988. computational approach to motion perception. *Biological Cybernetics*, 60:79–87.
- Weber, J. and Malik, J. 1995. Robust computation of optical flow in a multi-scale differential framework. *International Journal of Computer Vision*, 14:67–81.
- Weickert, J. 1998. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart.
- Weickert, J. and Schnörr, C. 2001. A theoretical framework for convex regularizers in PDE-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264.
- Weickert, J. and Schnörr, C. 2001. Variational optic flow computation with a spatio-temporal smoothness constraint. *Journal of Mathematical Imaging and Vision*, 14(3):245–255.
- Yacoob, Y. and Davis, L.S. 1999. Temporal multi-scale models for flow and acceleration. *International Journal of Computer Vision*, 32(2):1–17.
- Young, D.M. 1971. *Iterative Solution of Large Linear Systems*. Academic Press: New York.
- Yuille, A.L. and Poggio, T.A. 1986. Scaling theorems for zero crossings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):15–25.