



Highly Accurate Optic Flow Computation with Theoretically Justified Warping

NILS PAPENBERG

Institute for Mathematics, University of Lübeck, Wallstrasse 40, 23560 Lübeck, Germany
papenber@math.uni-luebeck.de (<http://www.math.uni-luebeck.de/safir>)

ANDRÉS BRUHN, THOMAS BROX, STEPHAN DIDAS 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 (<http://www.mia.uni-saarland.de>)
brox@mia.uni-saarland.de (<http://www.mia.uni-saarland.de>)
didas@mia.uni-saarland.de (<http://www.mia.uni-saarland.de>)
weickert@mia.uni-saarland.de (<http://www.mia.uni-saarland.de>)

Received September 21, 2004; Revised May 4, 2005; Accepted May 4, 2005

First online version published in January, 2006

Abstract. In this paper, we suggest a variational model for optic flow computation based on non-linearised and higher order constancy assumptions. Besides the common grey value constancy assumption, also gradient constancy, as well as the constancy of the Hessian and the Laplacian are proposed. Since the model strictly refrains from a linearisation of these assumptions, it is also capable to deal with large displacements. For the minimisation of the rather complex energy functional, we present an efficient numerical scheme employing two nested fixed point iterations. Following a coarse-to-fine strategy it turns out that there is a theoretical foundation of so-called warping techniques hitherto justified only on an experimental basis. Since our algorithm consists of the integration of various concepts, ranging from different constancy assumptions to numerical implementation issues, a detailed account of the effect of each of these concepts is included in the experimental section. The superior performance of the proposed method shows up by significantly smaller estimation errors when compared to previous techniques. Further experiments also confirm excellent robustness under noise and insensitivity to parameter variations.

Keywords: optical flow, differential methods, gradient constancy, warping, performance evaluation

1. Introduction

Optic flow estimation certainly belongs to one of the most crucial and well-investigated tasks in computer vision. This task is of surprisingly general nature encountered not only in motion estimation but also in 3D-reconstruction and image registration, for short, everywhere where correspondences between pixels have to be calculated. Starting with the classical works of

Horn and Schunck (1981) as well as Lucas and Kanade (1981), the last two decades have seen a tremendous improvement in the quality of optic flow estimation techniques. For example, the quadratic regulariser in the Horn and Schunck model has been replaced by smoothness constraints allowing piecewise smooth results and discontinuities in the flow field (Alvarez et al., 1999; Aubert et al., 1999; Black and Anandan, 1991; Cohen, 1993; Deriche et al., 1995; 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). Some of the underlying ideas are in spirit related to methods for joint motion estimation and motion segmentation (Farneäck, 2001; Mémin and Pérez, 2002; Cremers, 2003), and to optical flow methods motivated from robust statistics where outliers are penalised less severely (Black and Anandan, 1991, 1996). The problem of large displacements has been tackled by coarse-to-fine strategies (Anandan, 1989; Black and Anandan, 1996; Mémin and Pérez, 1998, 2002; Bruhn et al., 2005) as well as by non-linearised models (Nagel and Enkelmann, 1986; Alvarez et al., 2000). Furthermore, results have been enhanced merely by augmenting spatial approaches with an additional temporal dimension (Murray and Buxton, 1987; Nagel, 1990; Black and Anandan, 1991; Elad and Feuer, 1998; Weickert and Schnörr, 2001; Farneäck, 2001).

Undoubtedly, in the field of optic flow estimation considerable progress has been made thanks to many newly contributed ideas. But also efforts towards a deeper understanding of how methods actually work in detail and how parameter changes effect the output instigate the design of advanced and highly effective models. Moreover, the abundance of sophisticated methods from numerical optimisation theory became utilisable to the optic flow community as soon as models have been cast in a variational form.

The energy functional presented in this paper puts into effect the idea of using a variational formulation. Besides the standard assumption of grey value constancy, common to almost all optic flow schemes, the model incorporates constancy assumptions based on the image gradient and the Hessian. This makes the proposed technique more robust against illumination changes. The constancy of the gradient has been used earlier in Uras et al. (1988) and Tistarelli (1994) in order to overcome the aperture problem (Bertero, 1988) of local optic flow schemes. For variational approaches it has been mentioned in Schnörr (1994), yet the implementation of this concept in a variational model has not been tried in practice so far.

The second important feature of the model proposed in this paper is the formulation of all constancy assumptions in their original, nonlinear form. This results in a more accurate modelling of the optic flow, especially in the presence of large displacements. As of now such a nonlinear formulation has been employed only in the case of the grey value constancy assumption alone

(Nagel and Enkelmann, 1986; Alvarez et al., 2000). Here, also the constancy assumptions based on spatial derivatives are integrated into the model without linearisation.

In order to deal with this nonlinear model in a numerically efficient way that avoids possible local optima, we take advantage of a coarse-to-fine strategy. In combination with two nested fixed point iterations, one obtains not only a minimisation scheme that addresses the occurring nonlinearities, but one can also show that so-called warping techniques as used in Anandan (1989), Black and Anandan (1996) and Mémin and Pérez (1998, 2002) can be theoretically justified as numerical approximations of a single variational model containing non-linearised constancy assumptions. Up to now, the warping technique has been motivated from an algorithmic point of view only.

Several experiments give evidence for the superior performance of the proposed method. Comparison to the best results reported in the current literature reveals that our optic flow results are up to twice as accurate. Moreover, the method shows an appealing robustness w. r. t. noise and variation of the parameters together with reasonable computation times of only a few seconds on contemporary hardware.

Since our approach is based on the combination of several important concepts, most of which have also been successfully employed before, we also investigate the impact of each of these concepts in experiments, where the complexity of the model is increased step by step.

The remainder of this article is structured as follows: Section 2 introduces the variational model and discusses all underlying constancy assumptions. The numerical minimisation procedure is then described in Section 3. This leads to the theoretical foundation of warping approaches as a minimisation method for nonlinear data terms which is subject of Section 4. Section 5 reports on the experimental evaluation, whereas a short summary concludes the article in Section 6.

This paper extends work previously published at a conference (Brox et al., 2004). Substantial differences are, among other things, the introduction of additional constancy assumptions based on higher order derivatives, their integration into the numerical scheme, and a more extensive experimental evaluation in general. In particular, as the proposed method consists of several different concepts, the individual influence of each of these concept on the quality of the results is

investigated. This is sought to address the question why the proposed method reveals such a good performance.

2. Model

The proposed optic flow model is based on a variational approach, i.e. all model assumptions are formulated by means of an energy optimisation problem. In this section we introduce the energy functional that is to be minimised. In order to motivate this functional we consider $I : \mathbb{R}^3 \supset \Omega \rightarrow \mathbb{R}$ as an image sequence with two spatial and one temporal dimension at points $\mathbf{x} = (x_1, x_2, t)^\top$. We wish to determine the flow field $\mathbf{u} = (u_1, u_2, 1)^\top$, $u_i : \Omega \rightarrow \mathbb{R}^2$, $i = 1, 2$, representing the displacement between two frames at times t and $t + 1$. In the following, all model assumptions present in our approach are introduced individually.

2.1. Constancy Assumptions

Estimating motion entails the solution of a correspondence problem. That is, what pixel in one frame corresponds to what pixel in the other frame. In order to find these correspondences one needs to define some property or quantity that is not affected by the displacement. The common assumption is that the pixels' grey values stay constant during motion. In this paper it is shown that also other constancy assumptions are possible and can reveal certain advantages in comparison to the basic grey value constancy assumption.

Since it is advantageous for deriving the minimisation scheme, we mainly focus on constancy assumptions that can be expressed in the following form:

$$\|\mathcal{L} \circ I(\mathbf{x} + \mathbf{u}) - \mathcal{L} \circ I(\mathbf{x})\|_{l_2}^2 = 0 \quad (1)$$

where \mathcal{L} stands for a linear differential operator, like the identity operator in the case of the classical grey value constancy assumption, and $\|\cdot\|_{l_2}$ indicates the l_2 -norm as the square root of the sum of the squared vector/matrix components.

It is to be noted that the expression within the norm is nonlinear in \mathbf{u} . Since this causes problems when seeking to minimise the energy, this expression is mostly linearised by a first order Taylor expansion. The linear formulation of the constancy assumptions has the advantage that the algorithm for minimising the related energy inherits this linearity and hence is easier to solve. However, the linear approximation is sufficiently

accurate only if the gradient in the image changes linearly along the displacement, which is usually not the case, in particular in the presence of large displacements.

Actually, the original nonlinear formulation is the correct one. The only problem is that the minimisation of the corresponding energy functional requires a much more sophisticated numerical treatment. How such a minimization scheme can be designed is one of the topics of the present paper. In particular, we investigate the following non-linearised constancy assumptions.

- **Constancy of the grey value**

Generally the grey value of the pixel is assumed not to change during the movement:

$$\forall \mathbf{x} \in \Omega : \quad D_1(I, \mathbf{u}, \mathbf{x}) := \|I(\mathbf{x} + \mathbf{u}) - I(\mathbf{x})\|_{l_2}^2 = 0.$$

In a linear formulation this assumption leads to the well-known optic flow constraint

$$\forall \mathbf{x} \in \Omega : \quad I_{x_1}u_1 + I_{x_2}u_2 + I_t = 0,$$

as it has been formulated, for instance, in the classical algorithms of Horn and Schunck (1981) and Lucas and Kanade (1981). Subscripts denote partial derivatives.

Although this constancy assumption works fine in many cases, algorithms that rely only on this prerequisite cannot deal with image sequences with either local or global change in illumination. For image sequences where such settings appear, other constancy assumptions that are invariant against brightness changes are to be applied. Invariance can be ensured, for instance, by considering (spatial) derivatives.

- **Constancy of the gradient**

A global change in illumination both shifts and/or scales the grey values of an image sequence. Shifting the grey values will not affect the gradient. Contrary, scaling of the grey values changes only the length of the gradient vector, but not its direction. As a consequence, we assume that the spatial gradients of an image sequence can be considered as constant during motion:

$$\begin{aligned} \forall \mathbf{x} \in \Omega : \quad D_2(I, \mathbf{u}, \mathbf{x}) \\ := \|\nabla I(\mathbf{x} + \mathbf{u}) - \nabla I(\mathbf{x})\|_{l_2}^2 = 0, \end{aligned}$$

where $\nabla = (\partial_{x_1}, \partial_{x_2})^\top$ denotes the spatial gradient. A linearised formulation of this assumption can be found in Uras et al. (1988) and reads:

$$\forall \mathbf{x} \in \Omega : \begin{aligned} I_{x_1 x_1} u_1 + I_{x_1 x_2} u_2 + I_{x_1 t} &= 0, \\ I_{x_2 x_1} u_1 + I_{x_2 x_2} u_2 + I_{x_2 t} &= 0. \end{aligned}$$

- **Constancy of the Hessian**

It is not difficult to consider also higher order derivatives for the formulation of constancy assumptions. One choice including second order derivatives is the Hessian matrix \mathcal{H}_2 :

$$\forall \mathbf{x} \in \Omega : \begin{aligned} D_3(I, \mathbf{u}, \mathbf{x}) \\ := \|\mathcal{H}_2 I(\mathbf{x} + \mathbf{u}) - \mathcal{H}_2 I(\mathbf{x})\|_{l_2}^2 &= 0. \end{aligned}$$

Not all constancy assumptions based on derivatives are equally well-suited to estimate different types of motion. Both the gradient and the Hessian contain directional information. This has positive effects on the estimation of translational and divergent motions, since one obtains more than one constraint, which can improve robustness of the estimation technique. On the other hand, directional motion can be subject to changes when objects rotate, for example. Although a rather large rotation is needed to cause significant problems, constancy assumptions that include directional information are not optimal in such cases.

- **Constancy of the Laplacian**

One feature invariant against directional changes is the trace of the Hessian, which comes down to the Laplacian. The Laplacian can also be used to formulate a constancy assumption:

$$\forall \mathbf{x} \in \Omega : \begin{aligned} D_4(I, \mathbf{u}, \mathbf{x}) \\ := \|\Delta I(\mathbf{x} + \mathbf{u}) - \Delta I(\mathbf{x})\|_{l_2}^2 &= 0, \end{aligned}$$

where $\Delta = \partial_{x_1 x_1} + \partial_{x_2 x_2}$ denotes the spatial Laplacian.

Despite the fact that the following additional invariants cannot be rewritten in the form stated in (1), they are worth a brief mentioning.

- **Constancy of the gradient norm**

The spatial gradient in an arbitrary pixel of an image sequence can be decomposed into its norm and the directional part. While the directional information is changed by rotation, the norm of the gradient remains unaltered. Therefore, this quantity is useful for

the formulation of another constancy assumption:

$$\forall \mathbf{x} \in \Omega : \begin{aligned} D_5(I, \mathbf{u}, \mathbf{x}) \\ := (\|\nabla I(\mathbf{x} + \mathbf{u})\| - \|\nabla I(\mathbf{x})\|)^2 &= 0. \end{aligned}$$

Further invariants can be derived from the spatial Hessian.

- **Constancy of the norm of the Hessian**

$$\forall \mathbf{x} \in \Omega : \begin{aligned} D_6(I, \mathbf{u}, \mathbf{x}) := (\|\mathcal{H}_2 I(\mathbf{x} + \mathbf{u})\| \\ - \|\mathcal{H}_2 I(\mathbf{x})\|)^2 &= 0 \end{aligned}$$

- **Constancy of the determinant of the Hessian**

$$\forall \mathbf{x} \in \Omega : \begin{aligned} D_7(I, \mathbf{u}, \mathbf{x}) := (\det \mathcal{H}_2 I(\mathbf{x} + \mathbf{u}) \\ - \det \mathcal{H}_2 I(\mathbf{x}))^2 &= 0 \end{aligned}$$

Of course, it is possible to formulate constancy assumptions that utilise derivatives of order larger than two, but for the sake of brevity we will not discuss them in this paper. One should also note that higher order derivatives can be expected to yield a larger sensitivity to noise. Furthermore, the part of the image where the derivatives become zero, and hence do not provide any information, increases with the order of the derivatives.

A short summary of the proposed data terms is given in Table 1.

2.2. Smoothness Assumptions

A model for motion analysis that relies solely on constancy assumptions is in general not capable to determine the optic flow uniquely, especially in homogeneous areas. This problem is known as the aperture problem. Furthermore, it is reasonable to introduce a certain dependency between neighbouring pixels in order to deal with outliers caused by noise, occlusions, or other local violations of the constancy assumption.

This is achieved by augmenting the model with a smoothness assumption for the flow field. Horn and Schunck proposed in their model the following regulariser term (Horn and Schnuk, 1981):

$$E_{\text{smooth}}(\mathbf{u}) = \int_{\Omega} |\nabla u_1|^2 + |\nabla u_2|^2 \mathbf{d}\mathbf{x}. \quad (2)$$

However, such a smoothness assumption does not respect discontinuities in the flow field. In order to be able to capture also locally non-smooth motion, it is

Table 1. Overview of the presented constancy assumptions, their nonlinear formulation and the kind of motion they are used for.

	Constancy assumption	Data term	Motion
D_1	Grey value	$\ I(\mathbf{x} + \mathbf{u}) - I(\mathbf{x})\ _{l_2}^2$	all
D_2	Gradient	$\ \nabla I(\mathbf{x} + \mathbf{u}) - \nabla I(\mathbf{x})\ _{l_2}^2$	translational divergent
D_3	Hessian	$\ \mathcal{H}_2 I(\mathbf{x} + \mathbf{u}) - \mathcal{H}_2 I(\mathbf{x})\ _{l_2}^2$	translational divergent
D_4	Laplacian	$\ \Delta I(\mathbf{x} + \mathbf{u}) - \Delta I(\mathbf{x})\ _{l_2}^2$	all
D_5	Norm of gradient	$(\ \nabla I(\mathbf{x} + \mathbf{u})\ - \ \nabla I(\mathbf{x})\)^2$	all
D_6	Norm of hessian	$(\ \mathcal{H}_2 I(\mathbf{x} + \mathbf{u})\ - \ \mathcal{H}_2 I(\mathbf{x})\)^2$	all
D_7	Determinant of hessian	$(\det \mathcal{H}_2 I(\mathbf{x} + \mathbf{u}) - \det \mathcal{H}_2 I(\mathbf{x}))^2$	all

necessary to allow outliers in the smoothness assumption. This can be achieved by a non-quadratic penaliser function:

$$E_{\text{smooth}}(\mathbf{u}) = \int_{\Omega} \Psi(|\nabla u_1|^2 + |\nabla u_2|^2) \mathbf{d}\mathbf{x}. \quad (3)$$

In this way a piecewise smooth flow field is modelled. The penaliser function Ψ is chosen as $\Psi(s^2) = \sqrt{s^2 + \varepsilon^2}$ which yields the total variation regulariser proposed in Rudin et al. (1992) and Acar and Vogel (1994) for image denoising leading to a pseudo L_1 -minimisation. The quantity ε is not an additional parameter, but rather ensures the differentiability of Ψ in $s = 0$. It is chosen reasonably small, $\varepsilon = 0.001$, say.

Various other smoothness assumptions have been proposed and classified in the literature, e. g. the reader is referred to Weickert and Schnörr (2001) for further information.

Instead of solely assuming spatial smoothness, one can also introduce a spatio-temporal smoothness assumption (Nagel, 1990; Weickert and Schnörr, 2001). This is achieved by simply replacing the spatial gradient operator ∇ in the smoothness assumption by the spatio-temporal gradient $\nabla_3 = (\partial_{x1}, \partial_{x2}, \partial_t)$. While the method containing the spatial smoothness later on is marked as “2D”, we distinguish the method with the spatio-temporal smoothness assumption as “3D”. Note that due to the non-quadratic penalizer, spatio-temporal smoothness still allows discontinuities along the temporal axis.

2.3. Energy

A reasonable energy functional for optic flow estimation consists of a *data term*, which integrates the constancy assumptions, and a *smoothness term* introducing a smoothness assumption. In most cases it is not possible that both assumptions can be fulfilled at

the same time. Hence, there is a competition between both assumptions leading to the optimal compromise for the minimiser of the energy. The relative importance of each assumption is steered by a parameter α . This leads to the first coarse description of the energy functional

$$E(\mathbf{u}) = E_{\text{Data}}(\mathbf{u}) + \alpha E_{\text{Smooth}}(\mathbf{u}).$$

The smoothness term has already been defined in Section 2.2. For the construction of the data term, the constancy assumptions introduced earlier in Section 2.1 can serve as building blocks. The data term is designed such that it penalises the global deviation from the ideal model assumptions D_1, \dots, D_4 . Since we have several criteria at our disposal, the data term may take all of the model assumptions into account. Thus it consists of a combination of the quantities D_i weighted with positive real numbers γ_i ,

$$E_{\text{Data}}(\mathbf{u}) = \int_{\Omega} \sum_i \gamma_i D_i(I, \mathbf{u}, \mathbf{x}) \mathbf{d}\mathbf{x}.$$

This allows an adaptation to the image sequence at hand. In order to make the data term more robust with respect to outliers (Hampel et al., 1986; Huber, 1981) we further apply the same non-quadratic penaliser function $\Psi : \mathbb{R} \rightarrow \mathbb{R}$ as in the smoothness term:

$$E_{\text{Data}}(\mathbf{u}) = \int_{\Omega} \Psi \left(\sum_i \gamma_i D_i(I, \mathbf{u}, \mathbf{x}) \right) \mathbf{d}\mathbf{x}.$$

3. Numerical Solution

The energy functional defined in the last section describes an elaborated model for accurately estimating optic flows in various settings. However, the difficulty how to optimise this functional still remains unsolved. The energy functional contains several nonlinear terms, which make a minimisation non-trivial. In this section we introduce an iteration scheme which provides such a minimisation.

In a first step we restrict ourselves to a slightly simplified model where we only take the grey value constancy into account. In the final part of this section the scheme is extended to include also data terms of higher order. The simplified model reduces to

$$E(\mathbf{u}) = \int_{\Omega} \Psi((I(\mathbf{x} + \mathbf{u}) - I(\mathbf{x}))^2) \, d\mathbf{x} + \alpha \int_{\Omega} \Psi(\|\nabla u_1\|^2 + \|\nabla u_2\|^2) \, d\mathbf{x}. \quad (4)$$

The corresponding Euler-Lagrange equations that have to be satisfied by a function \mathbf{u} mimimising the energy (4) can be written as

$$\begin{aligned} \Psi'(I_z^2) \cdot I_z I_{x_1} \\ - \alpha \operatorname{div}(\Psi'(\|\nabla u_1\|^2 + \|\nabla u_2\|^2) \nabla u_1) &= 0 \\ \Psi'(I_z^2) \cdot I_z I_{x_2} \\ - \alpha \operatorname{div}(\Psi'(\|\nabla u_1\|^2 + \|\nabla u_2\|^2) \nabla u_2) &= 0 \end{aligned} \quad (5)$$

where

$$\begin{aligned} I_z &= I(\mathbf{x} + \mathbf{u}) - I(\mathbf{x}), \\ I_{x_1} &= \partial_{x_1} I(\mathbf{x} + \mathbf{u}), \\ I_{x_2} &= \partial_{x_2} I(\mathbf{x} + \mathbf{u}). \end{aligned} \quad (6)$$

Note that I_z is not a temporal derivative but a difference that is sought to be minimised. Since the energy functional (4) is nonlinear and not even convex, the solution of this equation system is quite a challenge. One way to handle nonlinear equations is to derive a fixed point scheme and to determine the solution iteratively. The actual iteration scheme as well as the proper choice of the initialisation are equally important.

3.1. Iteration Scheme

For the formulation of the iteration procedure we choose a fully implicit scheme in the smoothness term

and a semi-implicit scheme in the data term. Both approaches are preferable to explicit schemes, since they usually ensure faster convergence to the solution and better stability. Let $\mathbf{u}^k = (u_1^k, u_2^k, 1)^\top$ and let I_*^k denote abbreviations defined in (6) but with the iteration variable \mathbf{u}^k instead of \mathbf{u} , then \mathbf{u}^{k+1} can be obtained as the solution of

$$\begin{aligned} \Psi'((I_z^2)^{k+1}) \cdot I_z^{k+1} I_{x_1}^k \\ - \alpha \operatorname{div}(\Psi'(\|\nabla u_1^{k+1}\|^2 + \|\nabla u_2^{k+1}\|^2) \nabla u_1^{k+1}) &= 0 \\ \Psi'((I_z^2)^{k+1}) \cdot I_z^{k+1} I_{x_2}^k \\ - \alpha \operatorname{div}(\Psi'(\|\nabla u_1^{k+1}\|^2 + \|\nabla u_2^{k+1}\|^2) \nabla u_2^{k+1}) &= 0 \end{aligned} \quad (7)$$

where Ψ' denotes the derivative of Ψ with respect to its argument s^2 . This fixed point iteration is nonlinear in \mathbf{u} due to the structure of the components $I(\mathbf{x} + \mathbf{u}^{k+1})$ and $\Psi'(\cdot)$. Hence for further simplification towards linear equations, additional steps are necessary. First the terms of the form $I(\mathbf{x} + \mathbf{u}^{k+1})$ are linearised via Taylor expansion

$$\begin{aligned} I_z^{k+1} &= I(\mathbf{x} + \mathbf{u}^{k+1}) - I(\mathbf{x}) \\ &\approx I(\mathbf{x} + \mathbf{u}^k) + I_{x_1}^k du_1^k + I_{x_2}^k du_2^k - I(\mathbf{x}) \\ &= I_{x_1}^k du_1^k + I_{x_2}^k du_2^k + I_z^k, \end{aligned}$$

where we split the unknown iteration variable \mathbf{u}^{k+1} into the known variable \mathbf{u}^k and an unknown update $d\mathbf{u}^k = (du_1^k, du_2^k, 0)^\top$. This leads to a system

$$\begin{aligned} (\Psi')_{\text{Data}}^k \cdot ((I_{x_1}^k du_1^k + I_{x_2}^k du_2^k + I_z^k) I_{x_1}^k) \\ - \alpha \operatorname{div}((\Psi')_{\text{Smooth}}^k \nabla(u_1^k + du_1^k)) &= 0 \\ (\Psi')_{\text{Data}}^k \cdot ((I_{x_1}^k du_1^k + I_{x_2}^k du_2^k + I_z^k) I_{x_2}^k) \\ - \alpha \operatorname{div}((\Psi')_{\text{Smooth}}^k \nabla(u_2^k + du_2^k)) &= 0 \end{aligned} \quad (8)$$

with the abbreviations

$$\begin{aligned} (\Psi')_{\text{Data}}^k &:= \Psi'((I_{x_1}^k du_1^k + I_{x_2}^k du_2^k + I_z^k)^2), \\ (\Psi')_{\text{Smooth}}^k &:= \Psi'(\|\nabla(u_1^k + du_1^k)\|^2 \\ &\quad + \|\nabla(u_2^k + du_2^k)\|^2). \end{aligned} \quad (9)$$

Thanks to this linearisation the system (8) can be considered as a system of fixed point equations in the newly introduced variable $d\mathbf{u}^k$. The flow \mathbf{u}^k is known from the previous iteration step. Note that the system of equations is still nonlinear with respect to $d\mathbf{u}^k$, but now this nonlinearity stems only from the terms of the form Ψ' . We cope with this nonlinearity through an

approach similar to the one above and we introduce a second, inner, fixed point iteration that will allow to determine the increment \mathbf{du}^k used in the outer iteration. To this end, let $\mathbf{du}^{k,0} = 0$ our initialisation and let $\mathbf{du}^{k,l}$ denote the iteration variables at step l . Further, let $(\Psi')_{\text{Data}}^{k,l}$ and $(\Psi')_{\text{Smooth}}^{k,l}$ stand for the robustness factor and the diffusivity defined in (9). Then $\mathbf{du}^{k,l+1}$ is the solution of the *linear* system

$$\begin{aligned} & (\Psi')_{\text{Data}}^{k,l} \cdot \left((I_{x_1}^k du_1^{k,l+1} + I_{x_2}^k du_2^{k,l+1} + I_z^k) I_{x_1}^k \right) \\ & - \alpha \operatorname{div} \left((\Psi')_{\text{Smooth}}^{k,l} \nabla (u_1^k + du_1^{k,l+1}) \right) = 0 \\ & (\Psi')_{\text{Data}}^{k,l} \cdot \left((I_{x_1}^k du_1^{k,l+1} + I_{x_2}^k du_2^{k,l+1} + I_z^k) I_{x_2}^k \right) \\ & - \alpha \operatorname{div} \left((\Psi')_{\text{Smooth}}^{k,l} \nabla (u_2^k + du_2^{k,l+1}) \right) = 0. \end{aligned} \quad (10)$$

The fixed point of this system is employed to increment the outer iteration. After computation of the new outer iteration variable \mathbf{u}^{k+1} , we are able to obtain terms of $I(\mathbf{x} + \mathbf{u}^k)$ via interpolation. Such a calculation has to be performed only once just at the beginning of the inner iteration loop.

One should note that since the utilised penaliser function Ψ is convex, the increment \mathbf{du}^k is the solution of a convex optimisation problem. It is therefore uniquely defined and can be obtained by the suggested iteration procedure. Discretisation of the system (9) with finite differences leads to a sparse linear system that can be solved by common numerical methods such as successive over-relaxation (SOR) (Young, 1971).

3.2. Initialisation

The energy $E(\mathbf{u})$ to be minimised is nonlinear and non-convex. Hence the solution obtained by the proposed iteration scheme heavily depends on the initialisation. In order to avoid getting trapped in local minima, we suggest to employ a coarse-to-fine strategy.

For this purpose a complete image pyramid is generated where the image is successively downsampled by an arbitrary but fixed constant $\eta \in (0, 1)$ in accordance to Shannon's sampling theorem. For usual pyramids $\eta = 0.5$ is taken. Choosing larger η results in smoother transitions between levels of the pyramid and potentially leads to better results.

By the iteration scheme described above we are solving the minimisation problem successively on each level of the pyramid starting on the coarsest possible grid. There the scheme is initialised with the vanishing

flow $\mathbf{u}^0 = 0$. The solution of the outer iteration scheme is prolonged at each step to the next finer level, where it is employed for the initialisation of the outer fixed point iteration working on this level. The underlying assumption of such an approach is that on the coarsest level the energy indeed has a unique minimum already close to the global minimum of the original problem. The refinement at finer scales is then expected to yield this global minimum. A theoretical investigation on this topic can be found in Lefébure and Cohen (2001).

3.3. Extension to Higher Order Data Terms

The current minimisation scheme refers only to the model based on the grey value constancy. It is now shown that this scheme can also be extended to those higher order data terms that can be written in the form given in (1). For the sake of brevity and simplicity only the relevant terms are presented in tabulated form, since the essential idea of the minimisation scheme is the same for each of these data terms. For better comparison, we also display the terms belonging to the grey value constancy assumption already treated above.

Grey value constancy assumption:

Abbreviation:

$$I_z = I(\mathbf{x} + \mathbf{u}) - I(\mathbf{x})$$

Data term:

$$\Psi(I_z^2)$$

Euler Lagrange equation:

$$\Psi'(I_z^2) \cdot \nabla I(\mathbf{x} + \mathbf{u}) \cdot I_z$$

Iteration scheme:

$$\Psi'((I_z^{k+1})^2) \cdot \nabla I^k \cdot I_z^{k+1}$$

Taylor expansion:

$$I_z^{k+1} = I_z^k + (\nabla I^k)^\top \cdot \mathbf{du}^k$$

Gradient constancy assumption:

Abbreviation:

$$I_{x_i z} = I_{x_i}(\mathbf{x} + \mathbf{u}) - I_{x_i}(\mathbf{x}), \quad i = 1, 2$$

Data term:

$$\Psi \left(\sum_{i=1}^2 I_{x_i z}^2 \right)$$

Euler Lagrange equation:

$$\Psi' \left(\sum_{i=1}^2 I_{x_i z}^2 \right) \cdot \mathcal{H}_2 I(\mathbf{x} + \mathbf{u}) \cdot (I_{x_1 z}, I_{x_2 z})^\top$$

Iteration scheme:

$$\Psi' \left(\sum_{i=1}^2 (I_{x_i z}^{k+1}) \right) \cdot \mathcal{H}_2 I^k \cdot (I_{x_1 z}^{k+1}, I_{x_2 z}^{k+1})^\top$$

Taylor expansion:

$$I_{x_i z}^{k+1} = I_{x_i z}^k + (\nabla I_{x_i}^k)^\top \cdot \mathbf{du}^k, \quad i = 1, 2$$

Hessian constancy assumption:

Abbreviation:

$$I_{x_i x_j z} = I_{x_i x_j}(\mathbf{x} + \mathbf{u}) - I_{x_i x_j}(\mathbf{x}), \quad i, j = 1, 2$$

Data term:

$$\Psi \left(\sum_{i=1,2} I_{x_i x_j z}^2 \right)$$

Euler Lagrange equation:

$$\Psi' \left(\sum_{i=1,2} I_{x_i x_j z}^2 \right) \cdot \sum_{i,j=1}^2 \nabla I_{x_i x_j}(\mathbf{x} + \mathbf{u}) \cdot I_{x_i x_j z}$$

Iteration scheme:

$$\Psi' \left(\sum_{i=1,2} (I_{x_i x_j z}^{k+1})^2 \right) \cdot \sum_{i,j=1}^2 \nabla I_{x_i x_j}^k \cdot I_{x_i x_j z}^{k+1}$$

Taylor expansion:

$$I_{x_i x_j z}^{k+1} = I_{x_i x_j z}^k + (\nabla I_{x_i x_j}^k)^\top \cdot \mathbf{du}^k, \quad i, j = 1, 2$$

Laplacian constancy assumption:

Abbreviation:

$$\Delta I_z = \Delta I(\mathbf{x} + \mathbf{u}) - \Delta I(\mathbf{x})$$

Data term:

$$\Psi((\Delta I_z)^2)$$

Euler Lagrange equation:

$$\Psi'((\Delta I_z)^2) \cdot \nabla \Delta I(\mathbf{x} + \mathbf{u}) \cdot \Delta I_z$$

Iteration scheme:

$$\Psi'((\Delta I_z^{k+1})^2) \cdot \nabla \Delta I^k \cdot \Delta I_z^{k+1}$$

Taylor expansion:

$$\Delta I_z^{k+1} = \Delta I_z^k + (\nabla \Delta I^k)^\top \cdot \mathbf{du}^k$$

4. Relation to Warping

In the previous section, a numerical scheme has been presented that allows to minimise an energy functional with non-linearised constancy assumptions. In this section, it will be shown that our minimisation scheme can serve as justification for so-called coarse-to-fine warping techniques, which have been quite successful in optic flow computation, yet their motivation was so far based only on algorithmic considerations.

In former approaches (Anandan, 1989; Black and Anandan, 1996; Mémin and Pérez, 1998, 2002; Bruhn et al., 2005) warping methods have been used to compensate the disadvantage of linear data terms not being capable to deal adequately with large displacements. This has been achieved by the coarse-to-fine strategy. Since the displacement \mathbf{u} is already computed on a coarse level where the displacements are still small, and on the corresponding finer grid only the (also small) increment \mathbf{du} has to be obtained, the linearisation of the constancy assumption is a good approximation of the original assumption. By this methodology it became possible to successfully handle large displacements for the first time.

A second approach to deal with large displacements is based on the usage of nonlinear data terms as this has been proposed in Nagel and Enkelmann (1986) and Alvarez et al. (2000) and in the present paper. This approach offers the advantage to incorporate large displacements directly in the optic flow model. A shortcoming of this model, on the other hand, is that the corresponding nonlinear energy functional can have multiple local minima. In Alvarez et al. (2000) this problem has been addressed by a scale-space approach.

Table 2. Impact of different constancy assumptions on the quality of the optic flow field. Results of the 2D variant for the *Yosemite* sequence with clouds.

Yosemite with clouds									
Constancy	Term	γ_1	γ_2	γ_3	γ_4	σ	α	AAE	STD
Grey value	D_1	1	0	0	0	0.50	25	4.88°	7.60°
Gradient	D_2	0	1	0	0	0.85	7	2.56°	7.07°
Hessian	D_3	0	0	1	0	1.30	4	2.88°	6.55°
Laplacian	D_4	0	0	0	1	1.60	4	2.75°	8.05°

AAE = average angular error. STD = standard deviation.

In a similar way, we propose here a more efficient downsampling strategy. Both techniques aim at an initialisation of the flow field at the finest scale that is as close to the global minimum as possible.

The similarities in formulation and numerical realisation suggest a yet unexplained common background of warping methods on one hand, and methods with non-linearised constancy assumptions on the other hand. In fact we are able to demonstrate the equivalence of both schemes from the numerical point of view. If we restrict ourselves to spatial smoothness for the sake of simplicity, (8) reads

$$\begin{aligned} & (\Psi')_{\text{Data}}^k \cdot I_z^{k+1} I_x^k \\ & - \alpha \operatorname{div}((\Psi')_{\text{Smooth}}^k \nabla(u_1^k + du_1^k)) = 0 \\ & (\Psi')_{\text{Data}}^k \cdot I_z^{k+1} I_y^k \\ & - \alpha \operatorname{div}((\Psi')_{\text{Smooth}}^k \nabla(u_2^k + du_2^k)) = 0. \end{aligned}$$

For fixed k this system is identical with the Euler-Lagrange equations in Mémin and Pérez (2002). Just as in our approach, only the increment \mathbf{du} of the flow is calculated. Thus it is shown that the warping technique actually minimises an energy functional which is based on non-linearised constancy assumptions, whereas the minimisation relies on a coarse-to-fine strategy combined with two nested fixed point iterations.

5. Experiments

For evaluation purposes experiments with both synthetic and real-world image data were performed. The presented angular errors were computed via

$$\arccos \left(\frac{(u_1)_c(u_1)_e + (u_2)_c(u_2)_e + 1}{\sqrt{((u_1)_c^2 + (u_2)_c^2 + 1)((u_1)_e^2 + (u_2)_e^2 + 1)}} \right) \quad (11)$$

where the subscripts c and e denote the correct resp. the estimated flow (cf. Barron et al., 1994). As proposed in Barron et al. (1994) we preprocessed each image sequence by convolution with a Gaussian kernel before computation. In the following experiments the standard deviation of this Gaussian is referred to as parameter σ .

5.1. Synthetic Image Data

We start our experimental evaluation with a comparison of the four constancy assumptions $D_1 - D_4$ regarding their impact on the quality of the computed flow field. To this end, we consider the 2D variant of our algorithm with only one constancy assumption enabled at a time. This can be accomplished by setting the corresponding weights γ_i of the remaining constancy assumptions to zero. The famous *Yosemite* sequence with clouds created by Lynn Quam serves as test scene. This synthetic sequence combines divergent and translational motion under varying illumination and may be obtained from <ftp://csd.uwo.ca> under the directory `/pub/vision`.

In Fig. 1 the computed flow fields for all four different constancy assumptions are displayed. Moreover, the ground truth flow field is presented to allow for a visual comparison of the results. While the divergent motion of the flight is estimated accurately independently of the underlying constancy model, one can clearly observe a superior performance of the higher order data terms in those areas where illumination changes are present: the sky region. In this region the algorithm based on the grey value constancy assumption obviously fails to produce meaningful results. This behaviour is fully in accordance with our theoretical considerations from Section 2. A further confirmation of this observation is given by the results in Table 2. The listed average angular errors clearly indicate a sig-




Figure 1. (a) *Top left*: Frame 8 of the Yosemite sequence with clouds. (b) *Top right*: Ground truth with clouds. (c) *Middle left*: Computed flow field by our 2D method with constancy assumption D_1 (Grey value). (d) *Middle right*: Ditto with D_2 (Gradient). (e) *Bottom left*: Ditto with D_3 (Hessian). (f) *Bottom right*: Ditto with D_4 (Laplacian).

nificant better performance for those data terms that involve higher order derivatives.

After showing the good performance of different higher order constancy assumptions under varying illumination, we simplify our model before continuing

with the remaining experiments. Obviously, it is desirable to combine at least one higher order constancy assumption with the classical grey value constancy. For this reason we select the gradient constancy assumption to represent the class of higher order assumptions

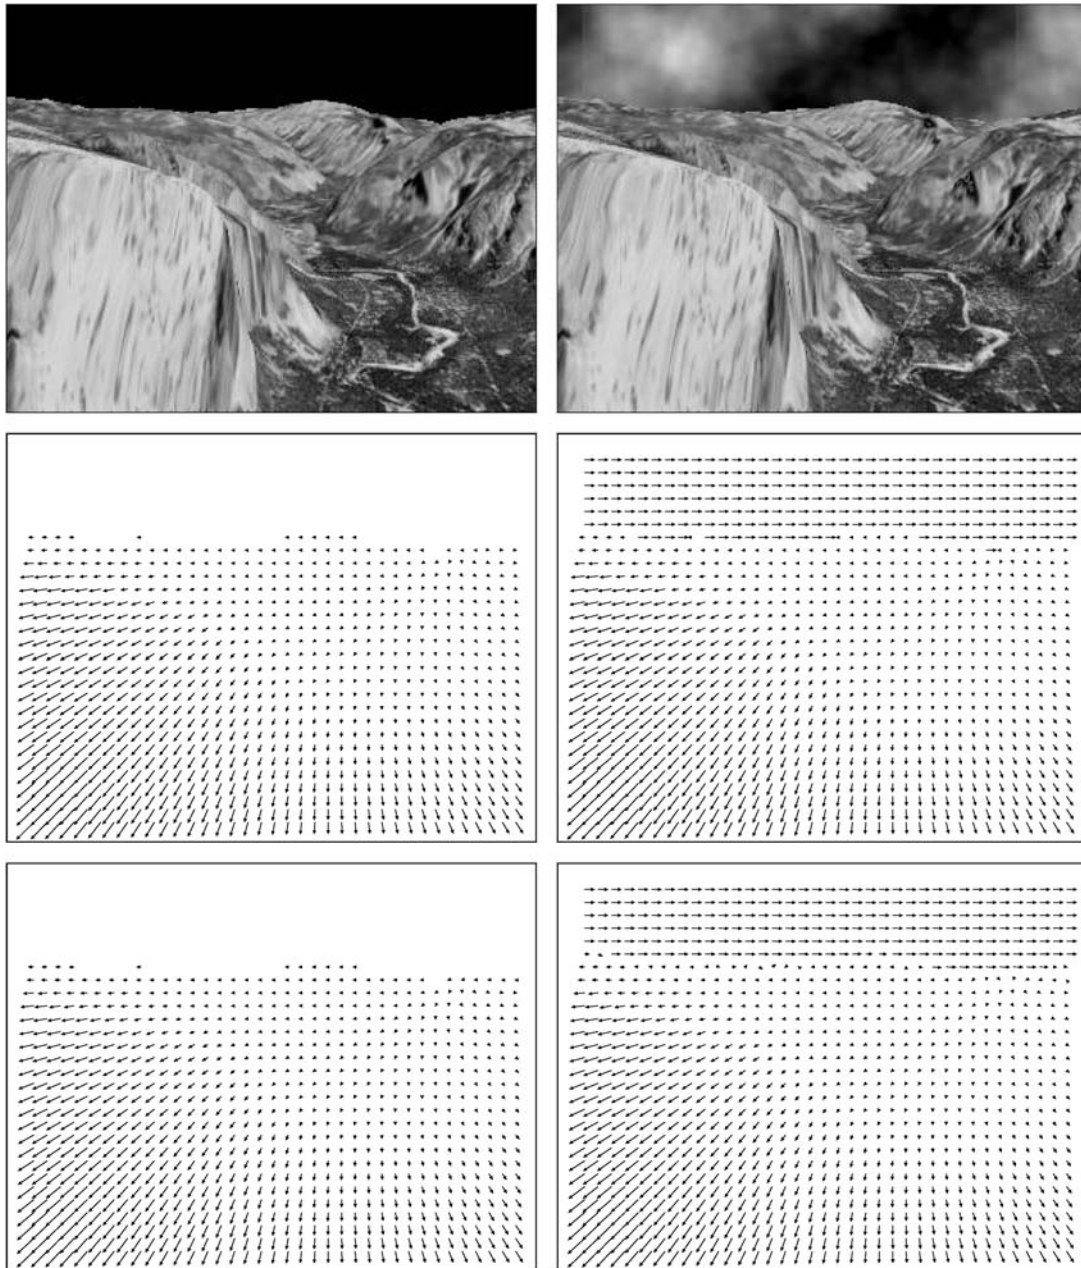


Figure 2. (a) *Top left*: Frame 8 of the *Yosemite* sequence without clouds. (b) *Top right*: Corresponding frame of the sequence with clouds. (c) *Middle left*: Ground truth without clouds. (d) *Middle right*: Ground truth with clouds. (e) *Bottom left*: Computed flow field by our 3D method for the sequence without clouds. (f) *Bottom right*: Ditto for the sequence with clouds.

in our energy functional and set $\gamma_3 = 0$ and $\gamma_4 = 0$. In order to reduce the number of parameters in the data term further, we set $\gamma_1 = 1$ fixed.

In our second experiment we compare the obtained results for the *Yosemite* sequence with clouds and its variant without clouds ([http://](http://www.cs.brown.edu/people/black/images.html)

www.cs.brown.edu/people/black/images.html) to the best results from the literature reported so far. This is done in Table 3. As one can see, our variational approach outperforms all other methods. Regarding the sequence with clouds, to the best of our knowledge, we achieve results that are more than twice as accurate

Table 3. Comparison between results from the literature with 100% density and our results for the *Yosemite* sequence with and without cloudy sky.

Yosemite with clouds			Yosemite without clouds		
Technique	AAE	STD	Technique	AAE	STD
Horn/Schunck, orig. (Barron et al. 1994)	31.69°	31.18°	Black (1994)	3.52°	3.25°
Singh, step 1 (Barron et al. 1994)	15.28°	19.61°	Szeliski and Coughlan (1994)	2.45°	3.05°
Anandan (Barron et al. 1994)	13.36°	15.64°	Black–Jepson (1996)	2.29°	2.25°
Singh, step 2 (Barron et al. 1994)	10.44°	13.94°	Ju et al. (1996)	2.16°	2.00°
Nagel (Barron et al. 1994)	10.22°	16.51°	Bab-Hadiashar–Suter (1998)	2.05°	2.92°
Horn–Schunck, mod. (Barron et al. 1994)	9.78°	16.19°	Lai–Vemuri (1998)	1.99°	1.41°
Uras et al. (Barron et al. 1994)	8.94°	15.61°	Our method (2D)	1.64°	1.43°
Alvarez et al. (2000)	5.53°	7.40°	Mémin–Pérez (2002)	1.58°	1.21°
Mémin–Pérez (1998)	4.69°	6.89°	Farneback (2000)	1.40°	2.57°
Bruhn et al. (2005)	4.17°	7.72°	Farneback (2001)	1.14°	2.14°
Our method (2D)	2.44°	6.90°	Bruhn et al. (2005)	1.02°	1.12°
Our method (3D)	1.78°	7.00°	Our method (3D)	0.98°	1.17°

AAE = average angular error. STD = standard deviation. 2D = spatial smoothness assumption. 3D = spatio-temporal smoothness assumption.

Table 4. Impact of the different model assumptions on the quality of the flow field for the *Yosemite* sequence with clouds.

Technique	Yosemite with clouds					AAE	STD
	DT-NQ	ST-NQ	3D	GRAD	DT-NL		
Classic Horn-Schunck	–	–	–	–	–	7.17°	9.09°
NQ smoothness term	–	✓	–	–	–	6.36°	8.34°
NQ data and smoothness term	✓	✓	–	–	–	5.97°	7.79°
NQ smoothness term 3D	–	✓	✓	–	–	5.66°	7.96°
NQ data and smoothness term 3D	✓	✓	✓	–	–	5.37°	7.81°
Gradient constancy	✓	✓	–	✓	–	3.50°	7.84°
Gradient constancy 3D	✓	✓	✓	✓	–	2.76°	7.58°
Non-linearised constancy	✓	✓	–	✓	✓	2.44°	6.90°
Non-linearised constancy 3D	✓	✓	✓	✓	✓	1.78°	7.00°

DT-NQ = Non-quadratic data term. ST-NQ = Non-quadratic smoothness term. 3D = spatio-temporal smoothness assumption. GRAD = Additional gradient constancy assumption. DT-NL = Non-linearised constancy assumption. AAE = average angular error. STD = standard deviation.

as all results from the literature. For the sequence without clouds, angular errors below one degree are reached for the first time with a method that offers full density. The corresponding flow fields presented in Fig. 2 give a qualitative impression of these raw numbers: they match the ground truth very well. Not only the discontinuity between the two types of motion is preserved, also the translational motion of the clouds is estimated accurately. The reason for this behaviour lies in the assumptions clearly stated in the energy functional: while the choice of the smoothness

term allows discontinuities, the gradient constancy assumption is able to handle brightness changes — like in the area of the clouds.

In our third experiment we study how the individual model assumptions influence the quality of the computed flow field. Table 4 shows that, starting from the classical approach of Horn and Schunck, each extension of the optic flow model implies a significant improvement in the average angular error. In a first step the introduction of a non-quadratic smoothness term allows the model to capture the motion discontinuities

Table 5. Results for our method with spatio-temporal smoothness assumption 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	AAE	STD	σ_n	AAE	STD
0	1.78°	7.00°	0	1.98°	1.15°
10	2.49°	6.56°	10	1.27°	1.26°
20	3.21°	6.46°	20	1.65°	1.40°
30	3.87°	6.60°	30	2.09°	1.59°
40	4.49°	6.87°	40	2.44°	1.78°

AAE = average angular error. STD = standard deviation.

more accurately. The extension of this model to the spatio-temporal domain leads to further improvements due to the availability of more data. Applying a non-quadratic function also to the data term addresses problems at the boundaries of the image sequence, where occlusions occur and therefore outliers in the data compromise the correct estimation of the flow field. Further augmentation by the gradient constancy assumption allows for the correct estimation of the motion even in the sky region, where local brightness changes are present. Finally, the usage of the non-linearised constancy assumptions enhances the accuracy of the results in general and in particular in the areas with large displacements.

Since we are using constancy assumptions that are based on higher order derivatives the question on the performance under noise arises. This question shall be answered in our fourth experiment. We added Gaussian noise of zero mean and different standard deviations to the *Yosemite* sequence with and without clouds prior to computing the flow field with our 3D method. The obtained results are presented in Table 5. They show that our approach even yields excellent flow estimates when severe noise is present.

In our fifth experiment we investigate the robustness of the free parameters in our approach: the weight γ_2 between the grey value and the gradient constancy assumption, the regularisation parameter α that steers the smoothness of the resulting flow field and σ , the standard deviation of the Gaussian kernel used for pre-smoothing. To this end, we computed results with parameter settings that deviate by a factor 2 in both directions from the optimum setting. The outcome listed in Table 6 shows that the method is also very robust under parameter variations.

Table 6. Parameter variation for our method with spatio-temporal smoothness assumption. Yosemite with clouds.

γ_2	σ	α	AAE	STD
100	1.0	60	1.78°	7.00°
50	1.0	60	1.82°	6.89°
200	1.0	60	1.90°	7.23°
100	1.0	30	2.57°	8.40°
100	1.0	120	2.09°	7.48°
100	0.5	60	2.21°	7.10°
100	2.0	60	2.23°	6.85°

AAE = average angular error. STD = standard deviation.

Although our paper does not focus on fast computation but on high accuracy, the implicit minimisation scheme presented here is also reasonably fast, especially if the reduction factor η is lowered or if the iterations are stopped before full convergence. The convergence behaviour and computation times for the *Yosemite* sequence with cloudy sky can be found in Table 7, both for the spatial and the spatio-temporal version. Computations have been performed on a 3.06 GHz Intel Pentium 4 processor executing C/C++ code.

As one may wonder whether it is always reasonable to assume spatio-temporal smoothness, Table 8 reveals the improvement in the AAE when using spatio-temporal smoothness. For this comparison we tested our method with two further publicly available sequences, the cropped *Street* sequence (available at www.cs.otago.ac.nz/research/vision) and the *Marble* sequence (available at http://i21www.ira.uka.de/image_sequences). Obviously, the spatio-temporal smoothness assumption has in general a quite positive impact. Temporal blurring

Table 7. Computation times and convergence for *Yosemite* sequence with clouds.

Reduction factor η	Outer fixed point iter.	Inner fixed point iter.	SOR iter.	Computation time	AAE
2D - spatial method					
0.95	77	10	5	14s	2.44°
0.90	38	2	5	1.9s	2.46°
0.85	25	2	5	1.2s	2.63°
0.80	18	2	5	0.93s	2.83°
0.75	14	2	5	0.78s	3.09°
0.70	12	2	5	0.66s	3.38°
0.65	10	2	5	0.58s	3.82°
0.60	8	2	3	0.45s	4.78°
Reduction factor η	Outer fixed point iter.	Inner fixed point iter.	SOR iter.	Computation time/frame	AAE
3D - spatio-temporal method					
0.95	77	10	5	19s	1.78°
0.95	77	2	5	4.5s	1.82°
0.90	38	2	5	2.3s	1.98°
0.85	25	2	5	1.6s	2.13°
0.80	18	2	5	1.2s	2.49°
0.75	14	2	5	1.0s	3.20°

artifacts are largely avoided by the discontinuity preserving model which only demands piecewise smoothness. This conjecture is further confirmed by the experiments in Bruhn et al. (2005).

5.2. Real-World Image Data

Finally we evaluate the performance of our method on real-world image data. To this end, we consider the *Karl Wilhelm Street* traffic sequence created by Nagel. Together with other challenging traffic scenes it is available from http://i21www.ira.uka.de/image_sequences/. It shows a surveillance video of an intersection and consists of 1033 frames of size 702×566 . One should note that this sequence has been recorded in interlaced mode and thus requires the handling of typical interlacing artifacts.

Figure 6 depicts the computed flow field between frame 860 and 861 as well as the corresponding magnitude plot. Evidently, our estimation gives very realistic results: the individual motion of cars, cyclists and pedestrians is well-captured by the proposed approach. Moreover, the motion of partially covered objects, such as the car that approaches the intersection from the right, are estimated precisely, as well. In this context we would like to emphasise that not all objects

Table 8. Comparison between our method with spatial (2D) and spatio-temporal (3D) smoothness.

Sequence	AAE 2D	AAE 3D
Yosemite with clouds	2.44°	1.78°
Yosemite without clouds	1.64°	0.98°
Street	4.93°	4.54°
Marble	4.74°	1.91°

AAE = average angular error.

are moving. For instance, the white van in the upper right and the pedestrians at the left crossing are stationary. As a consequence, the estimation of zero flow, as computed by our algorithm, is absolutely correct. Although the sequence suffers from the aforementioned interlacing artifacts, the flow boundaries are relatively sharp. Once more, this is a direct consequence of using non-quadratic functions such as the proposed total variation. From such a sharp and precise estimation it is just a small step towards a segmentation of all moving objects. Due to the accuracy of the computed flow field, the application of a rather simple segmentation technique such as thresholding would already be sufficient to achieve a very appealing result.

In our second experiment concerning real-world data we use the *Rheinhafen* traffic sequence that was




Figure 3. (a) Left: Frame 8 of the Yosemite sequence with clouds degraded by Gaussian noise of $\sigma n = 40$. (b) Right: Computed flow field by our 3D method.




Figure 4. (a) Left: Frame 10 of a cropped version of the Street sequence (b) Right: Computed flow field by our 3D method.




Figure 5. (a) Left: Frame 15 of the Marble sequence (b) Right: Computed flow field by our 3D method.

created by Nagel as well. It consists of 1000 frames of size 688×565 and is available for public download at the same internet address as the previously mentioned *Karl Wilhelm Street* scene. Again interlacing artifacts

are present in all frames. Figure 7 shows both the computed flow field between frame 1130 and 1131 and its magnitude plot (the sequence starts with frame 1000). As before the estimated flow field is very precise. Even



Figure 6. (a) Left: Computed flow field between frame 860 and 861 of the *Karl Wilhelm Street* traffic sequence. (b) Right: Computed magnitude of the optical flow field.



Figure 7. (a) Left: Computed flow field between frame 860 and 861 of the *Karl Wilhelm Street* traffic sequence. (b) Right: Computed magnitude of the optical flow field.

geometric phenomena such as perspective distortion are respected in the magnitude plot. This is indicated by the decreasing displacements of the bright van in the image foreground. The boundaries are again very sharp and would allow for a straightforward segmentation.

6. Conclusion

In this work, we have introduced a continuous, rotationally invariant model for optic flow computation based on nonlinear constancy assumptions. Besides the common grey value constancy assumption, analogues based on spatial derivatives, such as constancy of the gradient, Hessian, and Laplacian, have been incorporated into our formulation. Furthermore,

these constancy assumptions are not linearised as it is done in most other optic flow estimation methods. Additionally, the model contains a spatio-temporal TV regulariser in order to allow for discontinuities in the calculated flow field, and applies a robust function also to the data term. The variational framework offers a good opportunity of combining all these various concepts within a single energy functional.

The paper has also introduced a sound minimisation scheme for this energy functional. Starting with the nonlinear model, which does not impose any restrictions on the flow field in sharp contrast to linear approaches, the nonlinearity is resolved by two nested fixed point iterations. A coarse-to-fine strategy thereby avoids local minima appearing in this non-convex op-

timisation problem. An additional outcome of this paper, which emerges from the numerical scheme, is the explanation of the widely-used warping technique as a numerical approximation of the continuous energy functional with non-linearised constancy assumptions. This elucidates the success the warping technique is known for. We have further indicated that this important concept of postponing linearisations to the numerical scheme can also be transferred to constancy assumptions of higher order. It has been shown that the combination of various successful concepts in a single model and an elaborated numerical scheme for its optimisation is rewarded by excellent results and high robustness with respect to noise or parameter variations.

Acknowledgments

Our optic flow research has been partly funded by the projects WE 2602/1-1 and WE 2602/3-1 of the *Deutsche Forschungsgemeinschaft (DFG)* and the *DFG-Graduiertenkolleg 623 Leistungsgarantien für Rechnersysteme*. This is gratefully acknowledged. We further want to thank Bernhard Burgeth and the anonymous reviewers for useful comments regarding this paper.

References

- Acar, R. and Vogel, C.R. 1994. Analysis of bounded variation penalty methods for ill-posed problems. *Inverse Problems*, 10:1217–1229.
- 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.
- 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. *Proc. of the IEEE*, 76(8):869–889.
- 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*, pp. 138–145. Springer, Berlin.
- 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*, Maui, HI, June, IEEE Computer Society Press, 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. et al. 2004. High accuracy optical flow estimation based on a theory for warping. In *Computer Vision, Proc. 8th European Conference on Computer Vision*, T. Pajdla and J. Matas (Ed.) vol. 3024 of *Lecture Notes in Computer Science*, Springer, Prague, Czech Republic, pp. 25–36.
- Bruhn, A., Weickert, J., and Schnörr, C. 2005. Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 61(3):211–231.
- Cohen, I. 1993. Nonlinear variational method for optical flow computation. In *Proc. Eighth Scandinavian Conference on Image Analysis*, Tromsø, Norway, vol. 1, pp. 523–530.
- Cremers, D. 2003. A multiphase levelset framework for variational motion segmentation. In *Scale Space Methods in Computer Vision*, L. D. Griffin and M. Lillholm (Ed.) vol. 2695 of *Lecture Notes in Computer Science*, Springer, Berlin, pp. 599–614.
- Deriche, R., Kornprobst, P., and Aubert, G. 1995. Optical-flow estimation while preserving its discontinuities: a variational approach. In *Proc. Second Asian Conference on Computer Vision*, Singapore, vol. 2, pp. 290–295.
- Elad, M. and Feuer, A. 1998. Recursive optical flow estimation – adaptive filtering approach. *Journal of Visual Communication and Image Representation*, 9(2):119–138.
- Farneback, G. 2000. Fast and accurate motion estimation using orientation tensors and parametric motion models. In *Proc. 15th International Conference on Pattern Recognition*, Barcelona, Spain, vol. 1, 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*, Vancouver, Canada, July IEEE Computer Society Press, vol. 1, pp. 171–177.
- Hampel, F.R., et al. 1986. *Robust Statistics: The Approach Based on Influence Functions*. MIT Press, Cambridge, MA.
- Heitz, F. and Boutheimy, P. 1993. Multimodal estimation of discontinuous optical flow using Markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(12):1217–1232.
- Horn, B. and Schunck, B. 1981. Determining optical flow. *Artificial Intelligence*, 17:185–203.
- Huber, P.J. 1981. *Robust Statistics*. Wiley, New York.
- Ju, S., Black, M., and Jepson, 1996. A. 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*, San Francisco, CA, June IEEE Computer Society Press, pp. 307–314.

- 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.
- Lefébure, M. and Cohen, L.D. 2001. Image registration, optical flow and local rigidity. *Journal of Mathematical Imaging and Vision*, 14(2):131–147.
- 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.
- 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. Sixth International Conference on Computer Vision*, Bombay, India, Jan. Narosa Publishing House, 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.
- 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*, Karlsruhe, West Germany, vol. 2, 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*, O. Faugeras (Ed.) vol. 427 of *Lecture Notes in Computer Science*, 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 image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593.
- Nesi, P. 1993. Variational approach to optical flow estimation managing discontinuities. *Image and Vision Computing*, 11(7):419–439.
- Proesmans, M., et al. 1994. Determination of optical flow and its discontinuities using non-linear diffusion. In *Computer Vision - ECCV '94*, J.-O. Eklundh (Ed.) vol. 801 of *Lecture Notes in Computer Science*, Springer, Berlin, pp. 295–304.
- Rudin, L.I., Osher, S., and Fatemi, E. 1992. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268.
- Schnörr, C. 1994. Bewegungssegmentation von Bildfolgen durch die Minimierung konvexer nicht-quadratischer Funktionale. In *Mustererkennung 1994*, W. Kropatsch and H. Bischof (Ed.) Springer, Berlin, pp. 178–185.
- Schnörr, C., 1994. Segmentation of visual motion by minimizing convex non-quadratic functionals. In *Proc. Twelfth International Conference on Pattern Recognition*, Jerusalem, Israel, Oct. IEEE Computer Society Press, vol. A, pp. 661–663.
- Shulman, D. and Hervé, J. Regularization of discontinuous flow fields. In *Proc. Workshop on Visual Motion*, Irvine, CA, IEEE Computer Society Press, pp. 81–86.
- Szeliski, R. and Coughlan, J. 1994. Hierarchical spline-based image registration. In *Proc. 1994 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Seattle, WA, June IEEE Computer Society Press, pp. 194–201.
- Tistarelli, M. 1994. Multiple constraints for optical flow. In *Computer Vision - ECCV '94*, J.-O. Eklundh (Ed.) vol. 800 of *Lecture Notes in Computer Science*, Springer, Berlin, pp. 61–70.
- Uras, S., et al. 1988. A computational approach to motion perception. *Biological Cybernetics*, 60:79–87.
- Weickert, J. and Schnörr, C. 2001. A theoretical framework for convex regularizers in PDE-based computation of imagemotion. *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.
- Young, D. M. *Iterative Solution of Large Linear Systems*. Academic Press, New York, 1971.