Discontinuity-Preserving Computation of Variational Optic Flow in Real-Time

Andrés Bruhn[1], Joachim Weickert[1], Timo Kohlberger[2], and Christoph Schnörr[2]

[1] Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science, Building 27.1,
Saarland University, 66041 Saarbrücken, Germany
{bruhn, weickert}@mia.uni-saarland.de
[2] Computer Vision, Graphics, and Pattern Recognition Group,
Department of Mathematics and Computer Science,
University of Mannheim, 68131 Mannheim, Germany
{kohlberger, schnoerr}@uni-mannheim.de

Abstract. Variational methods are very popular for optic flow computation: They yield dense flow fields and perform well if they are adapted such that they respect discontinuities in the image sequence or the flow field. Unfortunately, this adaptation results in high computational complexity. In our paper we show that it is possible to achieve real-time performance for these methods if bidirectional multigrid strategies are used. To this end, we study two prototypes: i) For the anisotropic image-driven technique of Nagel and Enkelmann that results in a linear system of equations we derive a regular full multigrid scheme. ii) For an isotropic flow-driven approach with total variation (TV) regularisation that requires to solve a nonlinear system of equations we develop a full multigrid strategy based on a full approximation scheme (FAS). Experiments for sequences of size 160×120 demonstrate the excellent performance of the proposed numerical schemes. With frame rates of 6 and 12 dense flow fields per second, respectively, both implementations outperform corresponding modified explicit schemes by two to three orders of magnitude. Thus, for the first time ever, real-time performance can be achieved for these high quality methods.

Keywords: computer vision, optical flow, differential techniques, variational methods, multigrid methods, partial differential equations.

1 Introduction

In computer vision, the estimation of motion information from image sequences is one of the key problems. Typically, one is thereby interested in finding the displacement field between two consecutive frames, the so-called *optic flow*. In this context, variational methods play a very important role, since they allow for both a precise and dense estimation of the results. Such techniques are based on the minimisation of a suitable energy functional that consists of two terms: A data term that imposes temporal constancy on certain image features, e.g. on the grey value, and a smoothness term that regularises the often non-unique solution of the data term by an additional smoothness assumption.

R. Kimmel, N. Sochen, J. Weickert (Eds.): Scale-Space 2005, LNCS 3459, pp. 279–290, 2005.

Let us consider some image sequence $f(x, y, t)$, where (x, y) denotes the location within a rectangular image domain Ω, and $t \in [0, T]$ denotes time. Then, the assumption of a constant grey value over time can be formulated as

$$f(x + u, y + v, t + 1) - f(x, y, t) = 0. \tag{1}$$

Performing a Taylor expansion and dropping all higher order terms one obtains its linearised form, the so-called *optic flow constraint* (OFC)

$$f_x u + f_y v + f_t = 0. \tag{2}$$

Here, the function $(u(x, y, t), v(x, y, t))^\top$ is the wanted displacement field and subscripts denote partial derivatives. As classified in [33], there are basically three different types of strategies to regularise the evidently non-unique solution of this data term: *Homogeneous* regularisation that assumes overall smoothness and does not adapt to semantically important image or flow structures [20], *image-driven* regularisation that assumes piecewise smoothness and respects discontinuities in the image [1, 26] and *flow-driven* regularisation that assumes piecewise smoothness and respects discontinuities in the flow field; see e.g. [11, 29, 33]. Moreover, when considering image and flow-driven regularisation, one can distinguish between *isotropic* and *anisotropic* smoothness terms. While isotropic regularisers do not impose any smoothness at discontinuities, anisotropic ones permit smoothing along the discontinuity but not across it.

Although recent developments [7, 9, 25] have shown that variational methods are among the best techniques for computing the optic flow in terms of error measures [3], they are often considered to be too slow for real-time applications. In particular the computational costs for solving the resulting linear and nonlinear system of equations by using standard iterative solvers are regarded as too high. In [8] we have already demonstrated for variational methods with homogeneous regularisation that *bidirectional multigrid strategies* [5, 6, 19, 31, 35] do allow for real-time performance. These techniques that create a sophisticated hierarchy of equation systems with excellent error reduction properties belong to the fastest numerical schemes for solving linear or nonlinear systems of equations. In this paper we show that by using such methods also real-time performance for variational techniques with image- or flow-driven regularisation becomes possible. One should note that in this case the development of suitable multigrid strategies is much more difficult due to the anisotropy or nonlinearity of the underlying regularisation strategies. To the best of our knowledge our paper is the first one to report real-time performance for such variational optic flow methods on standard hardware.

Paper Organisation. Our paper is organised as follows. In Section 2 we give a short review on two variational techniques that serve as prototypes for image- and flow-driven regularisation. Section 3 shows how these problems can be discretised, while efficient bidirectional multigrid schemes for solving the resulting linear and nonlinear systems of equations are proposed in Section 4. In Section 5 we present an experimental evaluation that includes experiments with different real-world sequences as well as performance benchmarks for both prototypes. A summary in Section 6 concludes this paper.

Related Work. In the literature on variational optic flow methods, coarse-to-fine strategies are quite common to speed up the computation (see e.g. Anandan [2], Luettgen et

al. [24]). They are based on a successive refinement of the problem whereby coarse grid solutions serve as initial guesses on finer grids. However, from a numerical viewpoint such *unidirectional* schemes are not the end of the road. They are clearly outperformed by *bidirectional* multigrid methods that revisit coarser levels in order to obtain useful correction steps. While there is at least some literature on these highly efficient schemes for variational optic flow techniques with homogeneous and image-driven regularisation (Glazer [18], Terzopoulos [30], Zini et a. [37], El Kalmoun and Rüde [13], Enkelmann [15], Ghosal and Vaněk [17]), only the work of Borzi et al. [4] is known to the authors where an optic flow problem was solved by means of a nonlinear bidirectional multigrid scheme (FAS). Also for other tasks in image processing and computer vision multigrid methods have been used successfully. In the context of photometric stereo and image biniarisation Kimmel and Yavneh [22] developed an algebraic multigrid method, while Chan et al. [10] researched geometric multigrid schemes for variational deconvolution with TV regularisation. For TV denoising Vogel [32] proposed the use of a linear multigrid method within a nonlinear fixed-point iteration, while, very recently, Frohn-Schnauf et al. [16] investigated a nonlinear multigrid scheme (FAS) for the same task.

2 Prototypes for Variational Methods

2.1 The Method of Nagel and Enkelmann

As prototype for the class of optic flow methods with *image-driven* regularisation we consider the *anisotropic* technique of Nagel and Enkelmann [26]. Their method accounts for the problem of discontinuities by smoothing only along a projection of the flow gradient, namely its component orthogonal to the local image gradient. As a consequence, flow fields are obtained that avoid smoothing across discontinuities in the image data. The energy functional associated to this anisotropic form of regularisation is given by

$$E(u,v)=\int_\Omega \Big((f_x u + f_y v + f_t)^2 + \alpha(\nabla u^\top D(\nabla f)\nabla u + \nabla v^\top D(\nabla f)\nabla v) \Big)\, dxdy, \quad (3)$$

where $\nabla := (\partial x, \partial y)^\top$ denotes the spatial gradient and $D(\nabla f)$ is a projection matrix perpendicular to ∇f that is defined as

$$D(\nabla f) = \frac{1}{|\nabla f| + 2\epsilon^2} \begin{pmatrix} f_y^2 + \epsilon^2 & -f_x f_y \\ -f_x f_y & f_x^2 + \epsilon^2 \end{pmatrix} =: \begin{pmatrix} a & b \\ b & c \end{pmatrix}. \quad (4)$$

In this context ϵ serves as regularisation parameter that prevents the matrix $D(\nabla f)$ from getting singular. Following the calculus of variations [14], the minimisation of this convex functional comes down to solving its Euler–Lagrange equations that are given by

$$0 = f_x^2\, u + f_x f_y\, v + f_x f_t \;-\; \alpha\, \mathcal{L}_{\mathrm{NE}} u, \quad (5)$$

$$0 = f_x f_y\, u + f_y^2\, v + f_y f_t \;-\; \alpha\, \mathcal{L}_{\mathrm{NE}} v \quad (6)$$

with the *linear* differential operator

$$\mathcal{L}_{\mathrm{NE}} z(x,y) = \mathrm{div}\ (D(\nabla f(x,y))\nabla z(x,y)) \quad (7)$$

and reflecting Neumann boundary conditions.

2.2 The TV-Based Regularisation Method

In contrast to image-driven regularisation methods, *flow-driven* techniques reduce smoothing where edges in the *flow field* occur during computation. Our prototype for this class of variational optic flow techniques is an *isotropic* method that penalises deviations from the smoothness assumption with the L_1 norm of the flow gradient magnitude. This corresponds to total variation regularisation [28] and can be related to statistically robust error norms [21]. Thereby large deviations are penalised less severely than in the frequently used quadratic setting (L_2 norm). As a consequence, large gradient features such as edges are better preserved. The energy functional for this approach is given by

$$E(u,v) = \int_\Omega \left((f_x u + f_y v + f_t)^2 + \alpha \sqrt{|\nabla u|^2 + |\nabla v|^2 + \epsilon^2} \right) dxdy, \tag{8}$$

where ϵ serves as small regularisation parameter. A related functional that approximates TV regularisation is proposed in [34], while variational approaches for rotationally not invariant versions of TV regularisation have been researched in [11, 12, 23]. At first glance, the corresponding Euler-Lagrange equations that are given by

$$0 = f_x^2\, u + f_x f_y\, v + f_x f_t \; - \; \frac{\alpha}{2}\, \mathcal{L}_{\mathrm{TV}}(u,v), \tag{9}$$

$$0 = f_x f_y\, u + f_y^2\, v + f_y f_t \; - \; \frac{\alpha}{2}\, \mathcal{L}_{\mathrm{TV}}(v,u) \tag{10}$$

have a very similar structure than those in (5)-(6). However,

$$\mathcal{L}_{\mathrm{TV}}(z(x,y), \tilde{z}(x,y)) = \mathrm{div}\ (D(\nabla z(x,y), \nabla \tilde{z}(x,y))\ \nabla z(x,y)) \tag{11}$$

is evidently a *nonlinear* differential operator in z and $\tilde{z}$, since

$$D(\nabla z, \nabla \tilde{z}) = \frac{1}{\sqrt{|\nabla z|^2 + |\nabla \tilde{z}|^2 + \epsilon^2}}\, I =: \begin{pmatrix} a & b \\ b & c \end{pmatrix}, \tag{12}$$

where I is the identity matrix, $b = 0$ and $c = a$. As we will see later, this nonlinearity of the differential operator $\mathcal{L}_{\mathrm{TV}}$ has serious impact on the resulting discrete system of equations and on the derived multigrid strategy.

3 Discretisation

3.1 General Discretisation Aspects

Let us now discuss a suitable discretisation for the Euler-Lagrange equations (5)-(6) and (9)-(10). To this end we consider the unknown functions $u(x,y,t)$ and $v(x,y,t)$ on a rectangular pixel grid with cell size $\mathbf{h} = (h_x, h_y)^\top$, and we denote by $u^{\mathbf{h}}_{i,j}$ the approximation to u at some pixel i,j with $i = 1,...,N_x$ and $j = 1,...,N_y$. Spatial derivatives of the image data are approximated using a fourth-order approximation with the stencil $(1,-8,0,8,-1)/(12h)$, while temporal derivatives are computed with a simple two-point stencil. In order to discretise the divergence expressions in the differential operators L_{NE} and L_{TV} we propose the following finite difference approximations:

Table 1. Discretisations of averaging and differential operators

One-sided averaging	$M_x^{\pm,\mathbf{h}}\,(z_{i,j}) := \frac{z_{i\pm1,j}+z_{i,j}}{2}$	(13)
	$M_y^{\pm,\mathbf{h}}\,(z_{i,j}) := \frac{z_{i,j\pm1}+z_{i,j}}{2}$	(14)
One-sided differences	$D_x^{\pm,\mathbf{h}}\,(z_{i,j}) := \pm\frac{z_{i\pm1,j}-z_{i,j}}{h_x}$	(15)
	$D_y^{\pm,\mathbf{h}}\,(z_{i,j}) := \pm\frac{z_{i,j\pm1}-z_{i,j}}{h_y}$	(16)
Central differences	$D_x^{\mathbf{h}}\,(z_{i,j}) := \frac{z_{i+1,j}-z_{i-1,j}}{2h_x}$	(17)
	$D_y^{\mathbf{h}}\,(z_{i,j}) := \frac{z_{i,j+1}-z_{i,j-1}}{2h_y}$	(18)
Squared differences	$D_x^{2,\mathbf{h}}\,(z_{i,j}) := \frac{1}{2}\left(D_x^{+,\mathbf{h}}(z_{i,j})\right)^2 + \frac{1}{2}\left(D_x^{-,\mathbf{h}}(z_{i,j})\right)^2$	(19)
	$D_y^{2,\mathbf{h}}\,(z_{i,j}) := \frac{1}{2}\left(D_y^{+,\mathbf{h}}(z_{i,j})\right)^2 + \frac{1}{2}\left(D_y^{-,\mathbf{h}}(z_{i,j})\right)^2$	(20)
Gradient magnitude	$\left\lvert D^{2,\mathbf{h}}\,(z_{i,j})\right\rvert := \sqrt{D_x^{2,\mathbf{h}}(z_{i,j}) + D_y^{2,\mathbf{h}}(z_{i,j})}$	(21)

$$\partial_x\,(\,a(x,y)\;\partial_x z(x,y)) \approx D_x^{-,\mathbf{h}}\left(M_x^{+,\mathbf{h}}(a_{i,j})\;D_x^{+,\mathbf{h}}(z_{i,j})\right), \tag{22}$$

$$\partial_x\,(\,b(x,y)\;\partial_y z(x,y)) \approx D_x^{\mathbf{h}}\quad\left(\qquad b_{i,j}\;\;D_y^{\mathbf{h}}\quad(z_{i,j})\right), \tag{23}$$

$$\partial_y\,(\,b(x,y)\;\partial_x z(x,y)) \approx D_y^{\mathbf{h}}\quad\left(\qquad b_{i,j}\;\;D_x^{\mathbf{h}}\quad(z_{i,j})\right), \tag{24}$$

$$\partial_y\,(\,c(x,y)\;\partial_y z(x,y)) \approx D_y^{-,\mathbf{h}}\left(M_y^{+,\mathbf{h}}(c_{i,j})\;D_y^{+,\mathbf{h}}(z_{i,j})\right), \tag{25}$$

where the coefficients a, b and c are entries of the matrices $D(\nabla f)$ and $D(\nabla u, \nabla v)$ as shown in (4) and (12). Details on the required averaging and differential operators within the approximations are given in Table 1.

3.2 The Method of Nagel and Enkelmann

We are now in the position to write down the discrete Euler-Lagrange equations for the method of Nagel and Enkelmann. They are given by

$$0 = f_{x\,i,j}^{2,\mathbf{h}}\;u_{i,j}^{\mathbf{h}} + f_{x\,i,j}^{\mathbf{h}} f_{y\,i,j}^{\mathbf{h}}\;v_{i,j}^{\mathbf{h}} + f_{x\,i,j}^{\mathbf{h}} f_{t\,i,j}^{\mathbf{h}} \;-\; \alpha\; L_{\mathrm{NE\,i,j}}^{\mathbf{h}}\;u_{i,j}^{\mathbf{h}}, \tag{26}$$

$$0 = f_{x\,i,j}^{\mathbf{h}} f_{y\,i,j}^{\mathbf{h}}\;u_{i,j}^{\mathbf{h}} + f_{y\,i,j}^{2,\mathbf{h}}\;v_{i,j}^{\mathbf{h}} + f_{y\,i,j}^{\mathbf{h}} f_{t\,i,j}^{\mathbf{h}} \;-\; \alpha\; L_{\mathrm{NE\,i,j}}^{\mathbf{h}}\;v_{i,j}^{\mathbf{h}}, \tag{27}$$

for $i = 1, .., N_x$ and $j = 1, .., N_y$, where $L_{\mathrm{NE\,i,j}}^{\mathbf{h}}$ denotes the discrete version of the linear operator $\mathcal{L}_{\mathrm{NE}}$ at some pixel i, j. These $2N_xN_y$ equations constitute a *linear* system for the unknowns $u_{i,j}^{\mathbf{h}}$ and $v_{i,j}^{\mathbf{h}}$. One should note that there are two different types of coupling between the equations. The pointwise coupling between $u_{i,j}^{\mathbf{h}}$ and $v_{i,j}^{\mathbf{h}}$ in the data term and the anisotropic neighbourhood coupling via the operator $L_{\mathrm{NE\,i,j}}^{\mathbf{h}}$ in the smoothness term (for $u_{i,j}^{\mathbf{h}}$ and $v_{i,j}^{\mathbf{h}}$ separately).

3.3 The TV-Based Regularisation Method

Analogously, we discretise the Euler Lagrange equations for the TV-based regularisation method. The obtained *nonlinear* system of equations then reads

$$0 = f^{2,\mathbf{h}}_{x\,i,j}\, u^{\mathbf{h}}_{i,j} + f^{\mathbf{h}}_{x\,i,j} f^{\mathbf{h}}_{y\,i,j}\, v^{\mathbf{h}}_{i,j} + f^{\mathbf{h}}_{x\,i,j} f^{\mathbf{h}}_{t\,i,j} - \frac{\alpha}{2}\, L^{\mathbf{h}}_{\mathrm{TV\ i,j}}(u^{\mathbf{h}}_{i,j}, v^{\mathbf{h}}_{i,j})\, u^{\mathbf{h}}_{i,j}, \quad (28)$$

$$0 = f^{\mathbf{h}}_{x\,i,j} f^{\mathbf{h}}_{y\,i,j}\, u^{\mathbf{h}}_{i,j} + f^{2,\mathbf{h}}_{y\,i,j}\, v^{\mathbf{h}}_{i,j} + f^{\mathbf{h}}_{y\,i,j} f^{\mathbf{h}}_{t\,i,j} - \frac{\alpha}{2}\, L^{\mathbf{h}}_{\mathrm{TV\ i,j}}(u^{\mathbf{h}}_{i,j}, v^{\mathbf{h}}_{i,j})\, v^{\mathbf{h}}_{i,j}, \quad (29)$$

for $i = 1, .., N_x$ and $j = 1, .., N_y$. Here, the finite difference approximation of $\mathcal{L}_{\mathrm{TV}}(u, v)$ and $\mathcal{L}_{\mathrm{TV}}(v, u)$ results in the product of a common nonlinear operator $L^{\mathbf{h}}_{\mathrm{TV\ i,j}}(u^{\mathbf{h}}_{i,j}, v^{\mathbf{h}}_{i,j})$ and the pixel $u^{\mathbf{h}}_{i,j}$ and $v^{\mathbf{h}}_{i,j}$, respectively. Evidently, this constitutes a third way of coupling.

4 Multigrid

4.1 Basic Concept

In general, the obtained linear and nonlinear systems of equations are solved by using non-hierarchical iterative schemes; e.g. variants of the Jacobi or the Gauß-Seidel method [27, 36]. However, such techniques are not suitable for equation systems that are only coupled via a small local neighbourhood: It may take thousands of iterations to transport local information between unknowns that are not coupled directly. A Fourier analysis of the error confirms this observation: While high frequency components (small wavelength, local impact) are attenuated efficiently, lower frequency components (large wavelength, global impact) remain almost un-dampened. In order to overcome this problem multigrid methods are based on a sophisticated strategy. They make use of correction steps that compute the error (not a coarser version of the fine grid solution) on a coarser grid. Thus, lower frequency components of the error reappear as higher ones and allow for an efficient attenuation with standard iterative methods. In the following we explain this strategy in detail for both the linear and the nonlinear case by the example of a basic bidirectional two-grid cycle.

4.2 The Linear Two-Grid Cycle

For the sake of clarity, let us reformulate the linear equation system of the method of Nagel and Enkelmann (26)-(27) as

$$A^{\mathbf{h}} x^{\mathbf{h}} = f^{\mathbf{h}}. \quad (30)$$

Here $x^{\mathbf{h}}$ denotes the concatenated vector $((u^{\mathbf{h}})^\top, (v^{\mathbf{h}})^\top)^\top$, $A^{\mathbf{h}}$ is a symmetric positive definite matrix and $f^{\mathbf{h}}$ stands for the right hand side.

I) Multigrid methods starts by performing several iterations with a basic iterative solver. This is the so-called presmoothing relaxation step, where high frequency components of the error are removed. If we denote the result after these iterations by $\tilde{x}^{\mathbf{h}}$, the error is given by

$$e^{\mathbf{h}} = x^{\mathbf{h}} - \tilde{x}^{\mathbf{h}}. \quad (31)$$

II) Evidently, one is interested in finding $e^{\mathbf{h}}$ in order to correct the approximated solution $\tilde{x}^{\mathbf{h}}$. Although $e^{\mathbf{h}}$ cannot be computed directly, the linearity of $A^{\mathbf{h}}$ allows its computation via

$$A^{\mathbf{h}}e^{\mathbf{h}} = A^{\mathbf{h}}(x^{\mathbf{h}} - \tilde{x}^{\mathbf{h}}) = A^{\mathbf{h}}x^{\mathbf{h}} - A^{\mathbf{h}}\tilde{x}^{\mathbf{h}} = f^{\mathbf{h}} - A^{\mathbf{h}}\tilde{x}^{\mathbf{h}} = r^{\mathbf{h}}, \tag{32}$$

where $r^{\mathbf{h}}$ is called residual. Since high frequencies of the error have already been removed, we can speed up the computation by solving this equation system at a coarser resolution with grid cell size $\mathbf{H} = (H_x, H_y)^\top$:

$$A^{\mathbf{h}}e^{\mathbf{h}} = r^{\mathbf{h}} \quad \rightarrow \quad A^{\mathbf{H}}e^{\mathbf{H}} = r^{\mathbf{H}}. \tag{33}$$

One should note that at this point, a transfer of the original equation system to a coarser grid makes no sense: Unlike the error, the solution very probably contains (desired) high frequency components. A restriction of these components would severely deteriorate the approximative solution (aliasing).

III) After we have solved the residual equation system on the coarse grid with a method of our choice, we transfer the solution back to the fine grid and correct our approximation by the computed error

$$\tilde{x}^{\mathbf{h}}_{\text{new}} = \tilde{x}^{\mathbf{h}} + e^{\mathbf{h}}. \tag{34}$$

IV) In general, the transfer of the computed correction from a coarse grid by means of interpolation introduces some new high frequency components. To this end, a so-called postsmoothing relaxation step is performed, where once again some iteration of the basic iterative solver are applied.

4.3 The Nonlinear (FAS) Two-Grid Cycle

Also in this case, let us start with a reformulation of the nonlinear equation system resulting from the TV-based regularisation method (28)-(29) as

$$A^{\mathbf{h}}(x^{\mathbf{h}}) = f^{\mathbf{h}} \tag{35}$$

where $A^{\mathbf{h}}(x^{\mathbf{h}})$ is a nonlinear operator. The FAS strategy [5] works as follows:

I) We perform a presmoothing relaxation step with a nonlinear basic solver.

II) However, since $A^{\mathbf{h}}(x^{\mathbf{h}})$ is a nonlinear operator, the way of deriving a suitable coarse grid correction is significantly different from the linear case. The (implicit) relation between the error and the residual is given by

$$A^{\mathbf{h}}(\tilde{x}^{\mathbf{h}} + e^{\mathbf{h}}) - A^{\mathbf{h}}(\tilde{x}^{\mathbf{h}}) = f^{\mathbf{h}} - A^{\mathbf{h}}(\tilde{x}^{\mathbf{h}}) = r^{\mathbf{h}}. \tag{36}$$

In order to compute the desired correction we transfer the following nonlinear equation system to the coarse grid

$$A^{\mathbf{h}}(\tilde{x}^{\mathbf{h}} + e^{\mathbf{h}}) = r^{\mathbf{h}} + A^{\mathbf{h}}(\tilde{x}^{\mathbf{h}}) \quad \rightarrow \quad A^{\mathbf{H}}(\boxed{\tilde{x}^{\mathbf{H}}} + e^{\mathbf{H}}) = r^{\mathbf{H}} + \boxed{A^{\mathbf{H}}(\tilde{x}^{\mathbf{H}})}. \tag{37}$$

Here, frames visualise the additional terms compared to the linear case.

III) After we have solved the nonlinear residual equation system on the coarse grid, we subtract $\tilde{x}^{\mathbf{H}}$ from the solution in order to obtain $e^{\mathbf{H}}$. Its transfer to the fine grid then allows to perform the correction step.

IV) We perform a postsmoothing relaxation step with a nonlinear basic solver.

4.4 Advanced Multigrid Strategies

In order to increase the computational efficiency, the presented two-grid cycles are generally applied in a hierarchical way. While *V–cycles* make one recursive call of a two-grid cycle per level, faster converging *W–cycles* perform two. Nevertheless, multiple of such advanced cycles are required to reach the desired accuracy. Refining the original problem step by step (unidirectional coarse-to-fine approach) and solving the resulting linear or nonlinear equation system at each level by using some bidirectional V– or W–cycles, the multigrid strategy with the best performance is obtained: *full multigrid* [6]. For both the linear and nonlinear case we have developed such full multigrid schemes. Let us now sketch some implementation details.

4.5 Implementation Details

For the method of Nagel and Enkelmann we implemented a full multigrid scheme with four W–cycles per level each one based on one pre- and one postsmoothing iteration. In order to overcome the problematic anisotropic coupling we made use of a Gauß-Seidel method with alternating line relaxation (ALR) [35] as basic solver. For our second prototype, the TV-based regularisation method, we designed a FAS full multigrid scheme with two W–cycles per level each one based on two pre- and two postsmoothing iterations. In this case we embedded a Gauß-Seidel method with coupled point relaxation (CPR) [8] and frozen coefficients [16]. In order to allow for a complete multigrid hierarchy we thereby considered the use of non-dyadic intergrid transfer operators. As proposed in [8] they were realised by constant interpolation and simple averaging. Coarser versions of the linear and nonlinear operators were created by a discretisation coarse grid approximation (DCA) [35].

5 Experiments

In our first experiment we compare the efficiency of different numerical schemes for the discussed prototypes (Nagel and Enkelmann with $\alpha = 1000$ and $\epsilon = 10^{-2}$, TV-based

Table 2. Performance benchmark on a standard desktop computer with 3.06 GHz Pentium 4 CPU. Run times refer to the computation of a single flow field from the 160×120 dancing sequence

(a) Linear : Image-driven anisotropic regularisation (Nagel-Enkelmann)

Solver	Iterations	Time $[s]$	FPS $[s^{-1}]$	Speedup
Mod. Explicit Scheme ($\tau = 0.1666$)	36558	47.053	0.021	1
Gauß-Seidel (ALR)	607	3.608	0.277	13
Full Multigrid	1	0.171	5.882	275

(b) Nonlinear : Flow-driven isotropic regularisation (TV)

Solver	Iterations	Time $[s]$	FPS $[s^{-1}]$	Speedup
Mod. Explicit Scheme ($\tau = 0.0025$)	10631	30.492	0.033	1
Gauß-Seidel (CPR)	2679	6.911	0.145	4
FAS - Full Multigrid	1	0.082	12.172	372

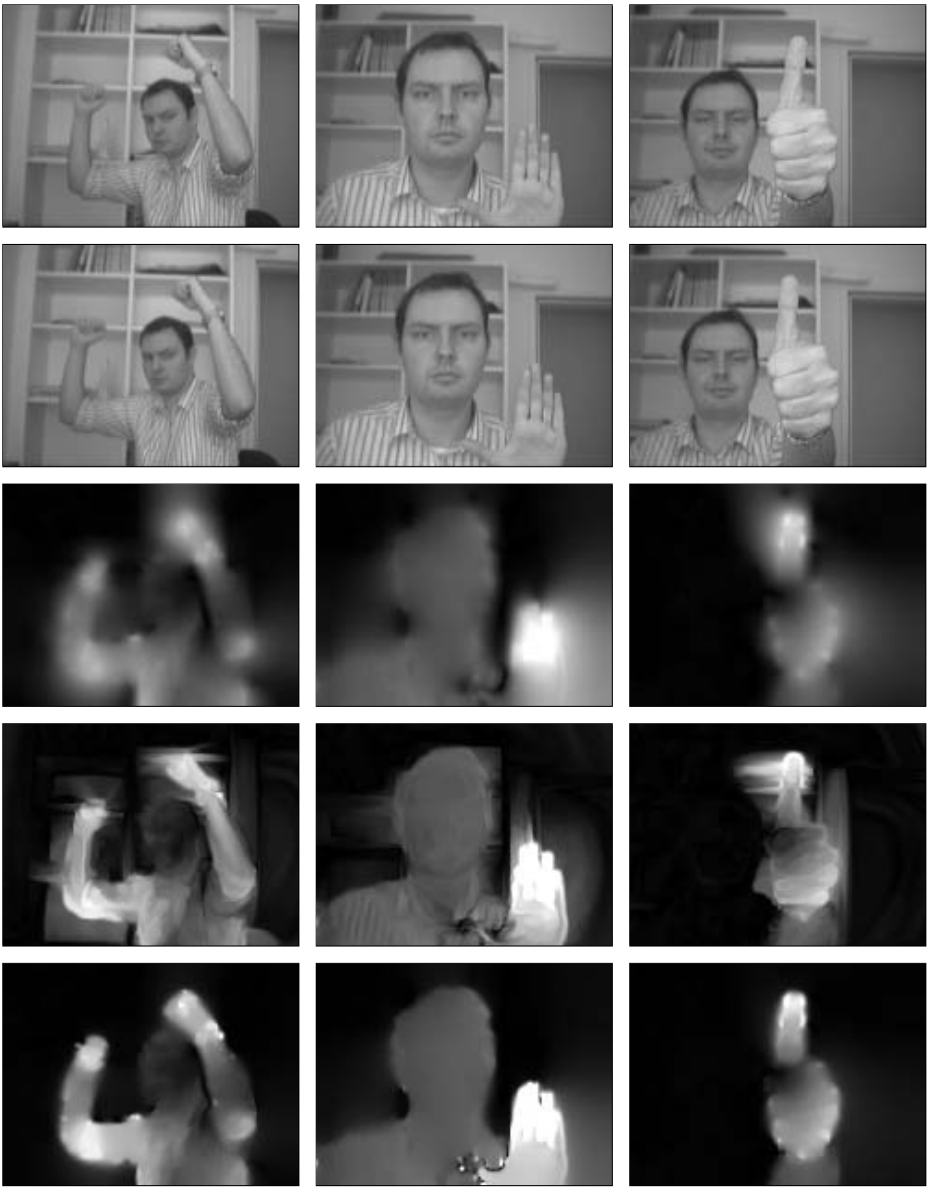

Fig. 1. *Left to right:* dancing sequence, waving sequence, rotating thumb sequence. *Top to bottom:* first frame, second frame, our CLG multigrid implementation from [8], our Nagel-Enkelmann multigrid implementation, our TV-based regularisation FAS multigrid implementation. *Brightness code:* The magnitude of a flow vector is encoded by its brightness. Brigther pixels stand for larger displacements. Color versions of the flow fields are available at `http://www.mia.uni-saarland.de/bruhn/scsp05/flowfields/`

regularisation method with $\alpha = 10$ and $\epsilon = 10^{-2}$). Apart from our full multigrid schemes we also implemented stand-alone versions of their basic solvers, namely the Gauß-Seidel methods with alternating line relaxation (ALR) and the Gauß-Seidel method with coupled point relaxation (CPR). Moreover, we considered a modified explicit scheme [34] that allows for larger time step sizes τ than ordinary explicit schemes (e.g. than gradient descent methods). For our evaluation we used a 160×120 real-world sequence in which a person dances in front of the camera. The iterations were stopped when the relative error $e_{rel} := \|x - \tilde{x}_n\|_2 / \|x\|_2$ dropped below 10^{-2}, where x denotes the correct solution and $\tilde{x}_n$ stands for the computed result after n iterations/cycles.

Table 2 shows the excellent performance of the proposed numerical schemes. In the linear case the presented full multigrid method outperforms the modified explicit scheme by two to three orders of magnitude. By allowing for the computation of six dense flow fields per second it is also more than one order of magnitude more efficient than its underlying basic solver. In the nonlinear case, the obtained speedups are even better. This time, the proposed FAS full multigrid method outperforms both the modified explicit scheme and the underlying basic solver by two to three orders of magnitude. Thereby, frame rates of twelve dense flow fields per second clearly show that also in this case real-time performance is well within our computational reach.

In our second experiment we compare the quality of both methods to that of a variational approach with homogeneous regularisation. To this end, we have computed flow fields for three different real-world sequences: for the previously used *Dancing Sequence* (complex motion), the *Waving Sequence* (translations and discontinuities) and the *Rotating Thumb Sequence* (rotation). The depicted colour plots in Figure 1 make the qualitative progress in the field of real-time variational optic flow computation explicit: One can easily see, that image- and flow-driven results are of much higher quality, since the underlying methods allow for a preservation of motion boundaries and discontinuities. Moreover, one can observe that the nonlinear flow-driven method is able to overcome the problem of oversegmentation that lies in the nature of image-driven techniques.

6 Summary and Conclusions

In this paper we have demonstrated that real-time optic flow computation on standard hardware is possible for variational optic flow techniques with both image- and flow-driven regularisation. This was accomplished by using highly efficient bidirectional full multigrid methods that solved the resulting linear and nonlinear systems of equations at different scales. In our experiments the proposed numerical schemes not only outperformed frequently used non-hierarchical solvers by two to three orders of magnitude, they also allowed for a very accurate estimation of the results. This shows that high quality optic flow computation and real-time performance are not opposing worlds. They can be combined if state-of-the-art numerical schemes are used. In our future work we plan to investigate different parallelisation strategies for the presented methods. This would allow us to process even larger sequences in real-time.

Acknowledgements

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

References

1. L. Alvarez, J. Esclarín, M. Lefébure, and J. Sánchez. A PDE model for computing the optical flow. In *Proc. XVI Congreso de Ecuaciones Diferenciales y Aplicaciones*, pages 1349–1356, Las Palmas de Gran Canaria, Spain, Sept. 1999.
2. P. Anandan. A computational framework and an algorithm for the measurement of visual motion. *International Journal of Computer Vision*, 2:283–310, 1989.
3. J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, Feb. 1994.
4. A. Borzi, K. Ito, and K. Kunisch. Optimal control formulation for determining optical flow. *SIAM Journal on Scientific Computing*, 24(3):818–847, 2002.
5. A. Brandt. Multi-level adaptive solutions to boundary-value problems. *Mathematics of Computation*, 31(138):333–390, Apr. 1977.
6. W. L. Briggs, V. E. Henson, and S. F. McCormick. *A Multigrid Tutorial*. SIAM, Philadelphia, second edition, 2000.
7. T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optic flow estimation based on a theory for warping. In T. Pajdla and J. Matas, editors, *Computer Vision – ECCV 2004*, volume 3024 of *Lecture Notes in Computer Science*, pages 25–36. Springer, Berlin, 2004.
8. A. Bruhn, J. Weickert, C. Feddern, T. Kohlberger, and C. Schnörr. Variational optical flow computation in real-time. *IEEE Transactions on Image Processing*, 2005. to appear.
9. A. Bruhn, J. Weickert, and C. Schnörr. Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 61(3):1–21, 2005. in press.
10. R. H. Chan, T. F. Chan, and W. L. Wan. Multigrid for differential-convolution problems arising from image processing. In G. Golub, S. H. Lui, F. Luk, and R. Plemmons, editors, *Proc. Workshop on Scientific Computing*, pages 58–72, Hong Kong, Sept. 1997.
11. I. Cohen. Nonlinear variational method for optical flow computation. In *Proc. Eighth Scandinavian Conference on Image Analysis*, volume 1, pages 523–530, Tromsø, Norway, May 1993.
12. R. Deriche, P. Kornprobst, and G. Aubert. Optical-flow estimation while preserving its discontinuities: a variational approach. In *Proc. Second Asian Conference on Computer Vision*, volume 2, pages 290–295, Singapore, Dec. 1995.
13. M. El Kalmoun and U. Rüde. A variational multigrid for computing the optical flow. In T. Ertl, B. Girod, G. Greiner, H. Niemann, H.-P. Seidel, E. Steinbach, and R. Westermann, editors, *Vision, Modelling and Visualization*, pages 577–584. IOS Press, 2003.
14. L. E. Elsgolc. *Calculus of Variations*. Pergamon, Oxford, 1961.
15. W. Enkelmann. Investigation of multigrid algorithms for the estimation of optical flow fields in image sequences. *Computer Vision, Graphics and Image Processing*, 43:150–177, 1987.
16. C. Frohn-Schnauf, S. Henn, and K. Witsch. Nonlinear multigrid methods for total variation denosing. *Computating and Visualization in Sciene*, 2004. to appear.
17. S. Ghosal and P. Č. Vaněk. Scalable algorithm for discontinuous optical flow estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(2):181–194, Feb. 1996.
18. F. Glazer. Multilevel relaxation in low-level computer vision. In A. Rosenfeld, editor, *Multiresolution Image Processing and Analysis*, pages 312–330. Springer, Berlin, 1984.

19. W. Hackbusch. *Multigrid Methods and Applications*. Springer, New York, 1985.
20. B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
21. P. J. Huber. *Robust Statistics*. Wiley, New York, 1981.
22. A. Kimmel and I. Yavneh. An algebraic multigrid approach for image analysis. *SIAM Journal on Scientific Computing*, 24(4):1218–1231, 2003.
23. A. Kumar, A. R. Tannenbaum, and G. J. Balas. Optic flow: a curve evolution approach. *IEEE Transactions on Image Processing*, 5(4):598–610, Apr. 1996.
24. M. R. Luettgen, W. C. Karl, and A. S. Willsky. Efficient multiscale regularization with applications to the computation of optical flow. *IEEE Transactions on Image Processing*, 3(1):41–64, 1994.
25. E. Mémin and P. Pérez. A multigrid approach for hierarchical motion estimation. In *Proc. 6th International Conference on Computer Vision*, pages 933–938, Bombay, India, Jan. 1998.
26. H.-H. Nagel and W. Enkelmann. 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, 1986.
27. J. M. Ortega and W. C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*, volume 30 of *Classics in Applied Mathematics*. SIAM, Philadelphia, 2000.
28. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
29. C. Schnörr. Segmentation of visual motion by minimizing convex non-quadratic functionals. In *Proc. Twelfth International Conference on Pattern Recognition*, volume A, pages 661–663, Jerusalem, Israel, Oct. 1994. IEEE Computer Society Press.
30. D. Terzopoulos. Image analysis using multigrid relaxation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(2):129–139, Mar. 1986.
31. U. Trottenberg, C. Oosterlee, and A. Schüller. *Multigrid*. Academic Press, San Diego, 2001.
32. C. R. Vogel. A multigrid method for total variation-based image denosing. *Computation and Control IV*, 20:323–331, 1995.
33. J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, Dec. 2001.
34. J. Weickert and C. Schnörr. Variational optic flow computation with a spatio-temporal smoothness constraint. *Journal of Mathematical Imaging and Vision*, 14(3):245–255, May 2001.
35. P. Wesseling. *An Introduction to Multigrid Methods*. Wiley, Chichester, 1992.
36. D. M. Young. *Iterative Solution of Large Linear Systems*. Academic Press, New York, 1971.
37. G. Zini, A. Sarti, and C. Lamberti. Application of continuum theory and multi-grid methods to motion evaluation from 3D echocardiography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(2):297–308, Mar. 1997.