J Math Imaging Vis (2011) 39: 230–244
DOI 10.1007/s10851-010-0237-z

Can Variational Models for Correspondence Problems Benefit from Upwind Discretisations?

Michael Breuß · Henning Zimmer · Joachim Weickert

Published online: 6 November 2010

Abstract Optic flow and stereo reconstruction are important examples of correspondence problems in computer vision. Correspondence problems have been studied for almost 30 years, and energy-based methods such as variational approaches have become popular for solving this task. However, despite the long history of research in this field, only little attention has been paid to the numerical approximation of derivatives that naturally occur in variational approaches.

In this paper we show that strategies from hyperbolic numerics can lead to a significant quality gain in computational results. Starting from a basic formulation of correspondence problems, we take on a novel perspective on the mathematical model. Switching the roles of known and unknown with respect to image data and displacement field, we use the arising hyperbolic colour equation as a basis for a refined numerical approach. For its discretisation, we propose to use one-sided differences in the correct direction identified via a smooth predictor solution. The one-sided differences that are first-order accurate are blended with higher-order central schemes. Thereby the blending mechanism interpolates between the following two situations: The one-sided method is employed at image edges which often coincide with edges in the displacement field. In smooth image parts the higher-order scheme is used. We apply our new scheme to several prototypes of variational models for optic flow and stereo reconstruction, where we achieve significant qualitative improvements compared to standard discretisations.

Keywords Finite difference methods · Adaptive methods · Hyperbolic colour equation · Variational correspondence problems · Variational optic flow · Variational stereo

M. Breuß (✉) · H. Zimmer · J. Weickert
Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science, Saarland University, Campus E1.1, 66123 Saarbrücken, Germany
e-mail: breuss@mia.uni-saarland.de

H. Zimmer
e-mail: zimmer@mia.uni-saarland.de

J. Weickert
e-mail: weickert@mia.uni-saarland.de

1 Introduction

Numerous computer vision applications, such as optic flow [13] or stereo reconstruction [17], require to solve a correspondence problem. This comes down to computing a displacement field which is the mapping that matches pixels of two given images. By use of the displacement field, non-trivial information about the depicted scenes can be obtained. In image sequence analysis the displacement field is called optic flow field and gives information about the apparent motion in a moving scene. In the stereo context, the absolute value of vectors in this field is called disparity and is needed to recover the depth information of a static scene. For an introduction to these and other related computer vision topics, see e.g. [17, 30].

Variational Approaches A successful class of techniques for solving correspondence problems are variational approaches that find the displacement field as the minimiser of an energy functional. Those methods have been studied for almost three decades, starting from the optic flow approach of Horn and Schunck [16]. During this period of time, many efforts have been spent to improve the quality of models. Some influential publications can be found in [3, 6, 7, 21, 23, 25, 28, 31–33, 36].

In order to apply such continuous-scale models to sampled digital images, one has to discretise the occurring image derivatives. This task offers a certain degree of freedom

in the choice of the derivative approximation. Surprisingly, this issue has hardly been studied for variational approaches to correspondence problems. If the discretisation is discussed at all, most approaches employ "standard" central finite difference approximations [3, 7, 36]. On the other hand, for variational approaches to problems like image restoration, more advanced approximation schemes have been considered for a long time [22, 26].

Our Idea In this paper we explore the use of sophisticated discretisations of *hyperbolic partial differential equations (HDEs)*, cf. [12, 18, 19, 29], that are usually relevant for describing gas or fluid dynamics. We show that the elliptic or parabolic PDEs arising in correspondence problems incorporate HDEs by considering the physics behind a transport process: Given an initial set of grey values (first image) and the velocity of transport (displacement), one can compute the grey values at a later time (second image). One realises that the role of *known* and *unknown* is switched compared to correspondence problems where the displacement is the unknown. This is a novel perspective on variational approaches to correspondence problems.

Our Contribution in Detail In this paper we make use of the mentioned relation between HDEs and correspondence problems. Exchanging the roles of known and unknown, we identify a *hyperbolic colour equation* as an important model component. This motivates us to consider one-sided upwind discretisations of image derivatives, since these are known to be useful in the context of HDEs. By a dedicated experiment in Sect. 3.2 we confirm that they can help to improve the estimation of the displacement field.

In order to obtain a reasonable compromise between such a first-order upwind approach and a good quality approximation in smooth regions, we borrow an idea from the numerics of HDEs: Like in so-called high-resolution schemes [18, 19, 29], we use low-order (upwind) difference approximations of image derivatives at image discontinuities but rely on high-order (central) differences in smooth regions. Note, that we only use the basic idea of the high-resolution schemes to obtain an adaptive method. Our adaptation is based on a smoothness measure that is specifically tailored to correspondence problems. Also in the choice of the high-order scheme, our procedure differs significantly from the usual high-resolution approach for HDEs.

Let us stress that our aim is not to contribute yet another state-of-the-art model for solving correspondence problems, but to advocate more suitable discretisations in order to obtain the best possible quality for a given model. By studying several prototypes of variational frameworks for optic flow and stereo reconstruction, we show that our approach can be beneficial for variational approaches to correspondence problems in general.

Related Work Our paper is the first journal paper concerned with the construction of a sophisticated numerical scheme for the hyperbolic colour equation that appears in variational models for correspondence problems. In this, we significantly extend our recent conference paper [35]. The most important extensions are: (i) We show a detailed experimental investigation of the mechanism that leads to improved results. (ii) We give a thorough discussion of the low- and high-order discretisations of second-order and mixed partial derivatives that appear in recent optic flow and stereo methods. (iii) We give a much more detailed account of the hyperbolic colour equation within variational models.

Paper Organisation In Sect. 2 we sketch the basics of variational approaches to correspondence problems in computer vision. After that, we discuss in Sect. 3 the arising hyperbolic colour equation and the effects of several discretisations of it. We also introduce the new adaptive numerical method there. In Sect. 4, we extend our approach to several prototypes of correspondence problems. The paper is finished by conclusions in Sect. 5.

2 The Variational Approach to Correspondence Problems

We introduce the setting by describing a classic and readily extendable variational model for correspondence problems in computer vision. For simplicity, we consider a 1D signal sequence $f(x,t)$ where $x \in \Omega$ denotes the position in the interval $\Omega \subset \mathbb{R}$ and $t \geq 0$ denotes time. For correspondence problems, at least two frames $f(x,t)$ and $f(x,t+1)$ of the signal evolution are given.

In order to compute the unknown displacement function $u(x)$ that maps f from time t to $t+1$, we consider the minimisation of the energy functional

$$E(u) = \int_\Omega \left[(f_x u + f_t)^2 + \alpha u_x^2\right] \mathrm{d}x, \tag{1}$$

where subscripts denote partial derivatives. This model is identical to a 1D version of the classic optic flow model of Horn and Schunck [16].

The term $(f_x u + f_t)^2$ is called *data term* and models how well the displacement u matches the signal sequence f. It is obtained as follows: We impose that the signal values are invariant under their displacement, i.e.

$$f(x+u, t+1) = f(x,t). \tag{2}$$

Especially in the context of 2D images, this basic assumption is called *brightness constancy assumption* [16]. Equation (2) is nonlinear in u which makes solving for u a difficult task. This is the motivation to use a first-order Taylor

series expansion to simplify this problem, which gives the so-called *linearised brightness constancy assumption*

$$f_x u + f_t = 0, \tag{3}$$

where we skipped the arguments of the functions. Using a quadratic penalisation of (3) then yields the data term from (1). The data term (3) allows to compute the solution

$$u = -\frac{f_t}{f_x}, \tag{4}$$

if $f_x \neq 0$. This is the so-called *normal flow*. However, in the presence of noisy signals, and for obtaining a solution in flat signal regions, additional assumptions are needed. These are especially crucial in the 2D case where the data term alone does not allow to compute an unique solution at all (*aperture problem* [4]).

One classical additional assumption for tackling the mentioned problems is the use of a *smoothness term* in conjunction with the data term [16]. The smoothness term models the assumption of a smoothly varying displacement field by penalising large derivatives of u. In this way it also allows to smoothly fill in the displacement field in regions where the data term is not sufficient. In our energy functional (1), the term u_x^2 is the smoothness term, and its contribution to the energy is steered by a smoothness weight $\alpha > 0$.

In order to actually compute a minimiser u of the energy (1), the calculus of variations [11] states that u necessarily has to fulfil the *Euler-Lagrange equation*

$$f_x(f_x u + f_t) - \alpha u_{xx} = 0, \tag{5}$$

with homogeneous Neumann boundary conditions.

Important Aspects As all effects of importance for us can be studied at hand of (1), we stick to it for a large part of the discussion. However, this model can easily be generalised to energy functionals

$$E(u) = \int_\Omega [M(f,u) + \alpha V(u_x)]\,\mathrm{d}x \tag{6}$$

with more complex data terms M and smoothness terms V, and to higher spatial dimensions. In the section devoted to numerical experiments we also consider such more advanced models.

3 The Colour Equation and Its Discretisation

Interpreting (3) as a PDE for the temporal evolution of f leads to a transport process in the form of a *hyperbolic colour equation*, cf. [19]. The name of this equation is derived from a specific application in fluid dynamics. There, for visualising flowlines in a fluid, coloured tracer particles can be given into the flow. The PDE governing the transport of the particles is the colour equation. Additionally, it is often assumed that the fluid moves always in the same way, i.e. the velocity of the fluid is in a steady state. Therefore, the velocity function of the colour is varying in space (as the displacement u in correspondence problems), but not in time.

As observed, in hyperbolic problems the roles of known and unknown are switched compared to correspondence problems. However, we will show in this paper that the same disretisation principles apply in both cases.

Typically, the colour equation is a PDE given in the framework of an initial value problem or an initial-boundary value problem. In our setting, the initial state $f(x,t)$ is evolved in time. The role of the other given state $f(x,t+1)$ will be (i) to determine the displacement direction and (ii) to provide data for accurate discretisations.

3.1 Discretisation Basics

For solving the Euler-Lagrange equation (5) numerically, we have to discretise the signal f, the displacement u, and their derivatives f_x, f_t and u_{xx}. For this we sample them on a spatio-temporal discrete grid. This gives the approximations $f_i^k \approx f(x_i,t_k)$ and $u_i \approx u(x_i)$, where $x_i := (i-\frac{1}{2})h$ and $t_k := k\tau$ with a spatial grid size h and a time step size τ. In this paper we only consider two frames f_i^k and f_i^{k+1}, and a temporal sampling of step size $\tau = 1$.

Now we turn to the discretisation of the occurring derivatives and the numerical boundary conditions. To this end we use the concept of finite differences, cf. [24]. As notation for the approximation of partial derivatives we use $f_d(x_i,t_k) \approx (f_d)_i^k$, where $d \in \{x, xx, t\}$, to denote the corresponding finite difference discretisation.

Temporal Discretisation For the time derivative we use the forward difference

$$(f_t)_i^k := \frac{1}{\tau}(f_i^{k+1} - f_i^k), \tag{7}$$

as this is the only reasonable choice, given the two frames f_i^k and f_i^{k+1}.

Discretisation of First-Order Spatial Derivatives The approximation of f_x offers different possibilities to define $(f_x)_i^l$, for $l \in \{k, k+1\}$. Basic choices are *forward*, *backward* and *central* differences:

$$\begin{aligned} \mathcal{D}_x^+ f_i^l &:= \frac{1}{h}(f_{i+1}^l - f_i^l), \\ \mathcal{D}_x^- f_i^l &:= \frac{1}{h}(f_i^l - f_{i-1}^l), \\ \mathcal{D}_x^0 f_i^l &:= \frac{1}{2h}(f_{i+1}^l - f_{i-1}^l) \end{aligned} \tag{8}$$

where the finite difference operators $\mathcal{D}^+$, $\mathcal{D}^-$ and $\mathcal{D}^0$ denote forward, backward and central differences, respectively.

Note that the approximation error of the one-sided differences is in $\mathcal{O}(h)$, whereas their central counterparts only involve an error of $\mathcal{O}(h^2)$. This, together with the unbiased stencil orientation, explains why they are a popular "standard" choice in image processing.

In order to increase the accuracy in computations for correspondence problems one may use *averaged differences*. These take into account differences from both time levels k and $k+1$. More specifically, we use second-order averaged central differences defined as

$$\mathcal{D}_x^0 f_i^{k+\frac{1}{2}} := \frac{1}{2}\left(\mathcal{D}_x^0 f_i^k + \mathcal{D}_x^0 f_i^{k+1}\right). \tag{9}$$

In the remainder of this paper such a central difference approximation will be referred to as a "standard" derivative approximation. We found empirically that this way to compute centred differences gives slightly better results than using data only from one of the frames.

Discretisation of Second-Order Spatial Derivatives Finally we have to approximate the second-order spatial derivative of the displacement function. As this choice is not crucial we propose a simple central approximation

$$(u_{xx})_i := \mathcal{D}_x^-\left(\mathcal{D}_x^+ u_i\right) = \frac{1}{h^2}(u_{i+1} - 2u_i + u_{i-1}). \tag{10}$$

Numerical Boundary Conditions Especially in the context of derivative computations, we have to pay attention to the boundary conditions employed in the numerics. As discrete boundary conditions we use homogeneous Neumann boundary conditions, in accordance with the procedure when computing the Euler-Lagrange equation (5). From an implementation point of view, these boundary conditions are realised by mirroring the signal values at the boundaries of the signal domain. More precisely, for a signal of length n, i.e., $(f_1, f_2, \ldots, f_{n-1}, f_n)$, we define the dummy values $f_0 := f_1$ and $f_{n+1} := f_n$ at the boundaries.

Multiscale Approach A differential model as given by an Euler-Lagrange equation is inherently local. This is reflected in the discretisation, as finite difference stencils for discretised derivatives just include the direct vicinity of a considered pixel. On the other hand, there is usually the need to establish correspondences between pixels over larger distances than covered by a stencil. Therefore, we employ a multiscale approach that is known as *warping*; see e.g. [7] for a detailed discussion. As a consequence, displacements are in practice always in the order of one pixel at all resolution levels.

Let us stress, that this strategy not only allows to capture large displacements, it also influences the discretisation. If image structures are displaced significantly between two frames, a discretisation like in (9) would be inaccurate as it would combine derivatives of completely different locations. By employing the multiscale approach such discretisations are readily justified.

3.2 Why the Colour Equation is Important

We now present an experiment which shows that an appropriate choice of $(f_x)_i^k$ is crucial for computing reasonable displacements u. Consider the two frames of a signal sequence in Fig. 1(a). There the signal is displaced by exactly one position to the right in its middle part and stays unchanged otherwise. This is illustrated in the ground truth displacement displayed in Fig. 1(b).

Note that this example comprises smooth as well as discontinuous signal and displacement regions. This makes it rather indicative. While the set-up of the experiment is simple, it is already of practical importance: The signal can be considered as one horizontal scanline from an orthoparallel 2D stereo problem which we will discuss in Sect. 4.

Let us note that the example also exhibits the so-called *occlusion problem*. This arises if a foreground object is shifted and occludes parts of the background. Thus, one cannot find any correspondence for the regions that are visible only in the first frame and are then occluded in the second frame. In our example this happens at point 9. Therefore, any computed displacements u will be corrupted at this occlusion. However, while we comment in this way on the expected computational results, the occlusion problem is not a topic in the focus of this paper. In practical computations, the occlusion problem is dealt with separately; see e.g. [8] and the references therein.

In Fig. 1(c)–(e) we depict the computed displacements using different discretisations for f_x. The displacements were obtained as the solution of the tridiagonal linear system of equations that arises from the discretised Euler-Lagrange equation (5). The smoothness weight was set here to the small value of $\alpha = 10^{-4}$ in order to show clearly the influence of the discretisation of the colour equation.

When comparing the displacements in Fig. 1(c)–(e), the effect of the discretisation of the colour equation becomes obvious. Central differences only perform well in the smooth signal regions, i.e. at the left and right boundaries. At discontinuities they produce severe oscillations. One-sided differences perform either favourably or fail totally. Obviously, the correct orientation matters here.

When using the correct one-sided difference scheme, the displacement almost coincides with the ground truth, except at one point. As indicated above, this is not a fault of the method, but is caused by the occlusion at the jump in the displacement. Note that as the displacement of an occluded point is in general undefined, we assigned in the ground truth to occluded points the displacement of their right neighbour.

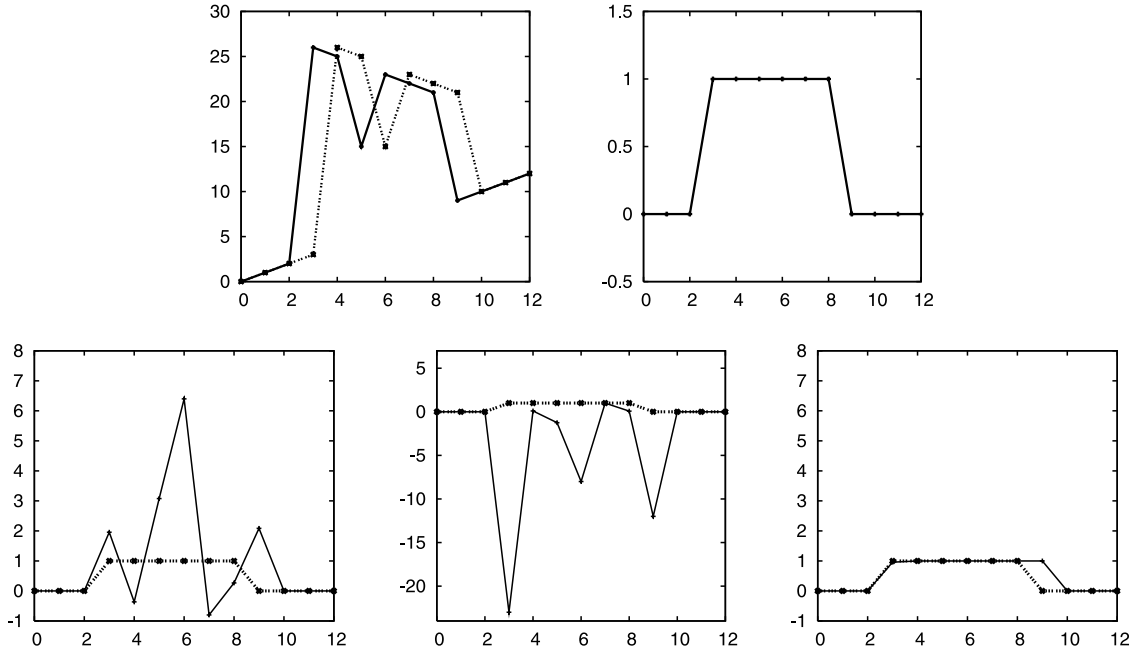

Fig. 1 *Top row*: (**a**) Signal at time k (*solid*) and $k+1$ (*dotted*). (**b**) Ground truth displacement. *Bottom row*: (**c**) Displacement computed using standard *central* differences averaged between level k and $k+1$ (*solid*), compared to the ground truth (*dotted*). (**d**) Same for one-sided *forward* differences. (**e**) Same for one-sided *backward* differences. Arrangements are from left to right

The observed behaviour in our experiment is in accordance with the theory of numerical methods for HDEs [18, 19]. There, so called *upwind schemes* are a widely used concept for the discretisation of transport equations. The term 'upwind' refers to correctly oriented one-sided differences. The correct orientation of an upwind stencil means in our case opposite to the displacement direction; see our experiment.

In the hyperbolic theory, central difference approximations as in (9) are known to lead to oscillations. They can even be unconditionally unstable, see e.g. [14, 29] for discussions. However, since in correspondence problems only one time step is performed, this instability is 'only' observable in terms of oscillations near strong gradients.

Let us remark that the stability of explicit schemes for HDEs involves a condition on the time step size called CFL condition [10, 19, 29]. This CFL condition can be translated into the assumption, that in one time step grey value information can only be transported up to one pixel. Therefore, the CFL condition for HDEs is related to the assumption of small displacements in correspondence problems. As discussed earlier, the CFL condition is always satisfied because we employ the multiscale approach that results in displacements of up to one pixel at all resolution levels.

3.3 The New Scheme

As we have seen, the low-order upwind differences perform well at signal discontinuities. However, in smooth regions the higher order of accuracy of central differences will give a better resolution of the displacement function. Hence, a natural idea is to combine the two types of schemes by using the high-order central approximation in smooth signal parts and upwinding at discontinuities.

This idea has been successfully used for the construction of so-called *high-resolution (HR) methods* [19, 29] for HDEs. They use a nonlinear blending of low- and high-order approximations, steered by a smoothness measure. In the classic framework for HDEs, the blending is performed by a limiter function that can be constructed to result in total variation stability of the HR scheme. This construction results in a non-uniform order but effectively in a discretisation of high quality, which explains the naming of HR methods.

The adaptation of this methodology to our variational framework gives an *adaptive high-resolution-type* (*HRT*) discretisation scheme for correspondence problems which is presented in the following. As motivated above, our main goal in the construction of the HRT method will be to obtain a better resolution at edges of u and around them.

Before proceeding with the scheme description, let us give some comments on similarities and differences of our method to the classic HR schemes, as we do not apply an off-the-shelf-approach in this paper. While it is very useful to consider the hyperbolic colour equation as a distinct important part to be re-interpreted and discretised, the final aim is to compute the displacement u. Especially, while a non-oscillatory resolution of edges in u is obviously important as seen by the experiment in Sect. 3.2, we do not have to spend

too much attention to structural properties of a discretisation, such as its total variation stability. This would only be important for a long-time integration of f, while in our case f is already given and only two time levels in f need to be considered.

Measuring Smoothness First we discuss how to determine the smooth and discontinuous regions of a signal. As indicated, this will be needed in order to steer the blending of the two considered schemes. Therefore, we introduce a *smoothness measure*

$$\Theta_i := \left|\mathcal{D}_x^- f_i^k - \mathcal{D}_x^+ f_i^k\right| + \left|\mathcal{D}_x^- f_i^{k+1} - \mathcal{D}_x^+ f_i^{k+1}\right|, \tag{11}$$

that is close to 0 in smooth regions where backward and forward differences of f_i^k and f_i^{k+1} are almost identical, and large at discontinuities of f_i.

Note that here one of the differences to a usual set-up of a high-resolution method for HDEs becomes obvious: In correspondence problems one already has the final state of the evolving signal f at hand, and so we can base our smoothness measure on both f^k and f^{k+1}. By the multiscale approach, such a combination of data from different frames will not result in mixing data from entirely different locations.

Determining the Upwind Directions Next we need to determine the appropriate upwind directions for discretising the colour equation. This is not straight forward, since the upwind direction depends on the direction of the displacement field, and this is exactly the unknown we aim to compute.

As a possible remedy, we propose to compute a *predictor solution* $\tilde{u}$ whose sign determines the upwind direction. The predictor is computed using the high-order standard approximation f_x^H of the derivative f_x. It is given by the averaged central difference approximation

$$\left(f_x^H\right)_i := \mathcal{D}_x^0 f_i^{k+\frac{1}{2}}. \tag{12}$$

In order to avoid oscillations as occurring in the experiment in Sect. 3.2, we use a comparatively large smoothness weight in this computation ($\alpha = 1$).

With the help of the predictor solution $\tilde{u}$, the *low-order upwind approximation* f_x^L of f_x is defined as

$$\left(f_x^L\right)_i := \begin{cases} \mathcal{D}_x^- f_i^k, & \text{if } \tilde{u}_i > 0, \\ \mathcal{D}_x^+ f_i^k, & \text{if } \tilde{u}_i < 0, \\ (f_x^H)_i, & \text{if } \tilde{u}_i = 0. \end{cases} \tag{13}$$

Revisiting the experiment from Fig. 1, one confirms that this definition agrees with the results obtained there.

We compute in (13) the upwind stencils employing frame k. This is theoretically justified since the linearised brightness constancy assumption (3) involves a linearisation at that time level. If we would switch the roles of the frames k and $k+1$, then the method would use at pixel i the upwind stencil pointing in the opposite direction.

The Blending Function Now we define the *blending function* $\Phi(\Theta_i)$ which realises the switch between high-order and low-order approximations in accordance to the value of Θ_i.

The idea is that it shall be close to 1 in smooth signal regions, which will yield a high-order approximation there. At discontinuities it shall be close to 0 which will lead to a low-order upwind approximation that is better suited there. For the actual choice of $\Phi(\Theta_i)$ we propose

$$\Phi(\Theta_i) := \begin{cases} 1 - \Theta_i, & \text{if } 0 \le \Theta_i < 1, \\ 0, & \text{else.} \end{cases} \tag{14}$$

The blending is performed here in a different way than in the usual set-up of high-resolution schemes for HDEs, cf. [19, 29] for detailed discussions of the latter setting. Other blending functions—especially those that are standard in the field of HDEs—do not lead to better results. We tested this, but as it turns out to be of no particular use we do not comment on it in more detail here.

The High-Resolution-Type (HRT) Discretisation Scheme Now everything is prepared to define the adaptive HRT discretisation. It reads as

$$(f_x)_i^k := \left(f_x^L\right)_i + \Phi(\Theta_i)\left[\left(f_x^H\right)_i - \left(f_x^L\right)_i\right], \tag{15}$$

using the function $\Phi(\Theta_i)$ to blend between the high-order derivative approximation f_x^H and its low-order counterpart f_x^L.

4 Evaluation of the HRT Scheme

Now we elaborate on the developed methodology by applying it within several prototypes of correspondence problems in computer vision. The purpose of our broad selection is to illustrate the benefits of our general concept. We consider: (i) The classic optic flow model of Horn and Schunck [16] which we have already discussed in a 1D version in Sect. 2, (ii) the more recent optic flow method of Brox et al. [7], and (iii) the variational stereo approach of Slesareva et al. [28]. We will briefly review the models, sketch how to extend our HRT discretisation scheme for these cases, and discuss computational results in detail.

4.1 Optic Flow: Basics

For optic flow computation we are given a 2D image sequence $f(x, y, t)$ where $(x, y)^\top \in \Omega_2$ denotes the location within a rectangular image domain $\Omega_2 \subset \mathbb{R}^2$ and $t \geq 0$ denotes time. The sought flow field $(u, v)^\top$ that gives the displacements from time t to $t+1$ is found as the minimiser of the 2D energy functional

$$E(u, v) = \int_{\Omega_2} [M(f, u, v) + \alpha V(\nabla u, \nabla v)] \, \mathrm{d}x \, \mathrm{d}y, \tag{16}$$

where $\nabla := (\partial_x, \partial_y)^\top$ denotes the spatial gradient operator.

4.2 Optic Flow: The Classic Method of Horn and Schunck

The brightness constancy assumption in the 2D optic flow case is given by

$$f(x + u, y + v, t + 1) = f(x, y, t). \tag{17}$$

After a first order Taylor linearisation and a quadratic penalisation one ends up with the data term of Horn and Schunck [16]:

$$M(f, u, v) = (f_x u + f_y v + f_t)^2. \tag{18}$$

Horn and Schunck [16] proposed in addition the quadratic smoothness term

$$V(\nabla u, \nabla v) = |\nabla u|^2 + |\nabla v|^2. \tag{19}$$

The corresponding Euler-Lagrange equations are

$$f_x \left(f_x u + f_y v + f_t\right) - \alpha \left(u_{xx} + u_{yy}\right) = 0, \tag{20}$$

$$f_y \left(f_x u + f_y v + f_t\right) - \alpha \left(v_{xx} + v_{yy}\right) = 0. \tag{21}$$

Discretisation in 2D In order to discretise the occurring 2D images and the flow field, we sample them on a 2D spatio-temporal discrete grid. For the images, this yields the approximation $f^k_{i,j} \approx f(x_i, y_j, t_k)$. Here, $x_i := (i - \frac{1}{2}) h_x$ and $y_j := (j - \frac{1}{2}) h_y$ for spatial grid sizes h_x and h_y in x- and y-direction, respectively. The discretisation of the flow fields works accordingly.

The discretised Euler-Lagrange equations (20) and (21) now lead to a penta-diagonal linear system of equations. Due to its sparsity, it can be solved by well-known iterative solvers [34]. We employ the successive overrelaxation (SOR) method.

4.2.1 *The HRT Scheme for the Method of Horn and Schunck*

Now we adapt the HRT discretisation scheme from Sect. 3.3 to the 2D optic flow case and the model of Horn and Schunck.

First of all we need distinct smoothness measures Θ_x, Θ_y for the x- and the y-direction, respectively. For Θ_x we use the according expression (11) from the 1D case, and Θ_y is obtained by using y- instead of x-differences.

The derivative approximations of f_x are obtained from (8). We only need to replace f_{i+l} by $f_{i+l,j}$, for $l \in \{-1, 0, 1\}$. The approximations of f_y can be easily obtained from the x-derivatives by switching the role of i and j. For f_t and u_{xx}, u_{yy}, v_{xx}, v_{yy} we use the corresponding 2D extension of (7) and (10), respectively.

4.2.2 *Numerical Experiments for the Method of Horn and Schunck*

In our first numerical experiment, we compute the flow field for a simple synthetic sequence we have created, see Fig. 3. The sequence depicts a rectangle that is displaced by one pixel to the right and one pixel to the bottom. This motion is encoded in the ground truth flow field in Fig. 3(c). To visualise flow fields, we use a colour code where colour encodes the direction and brightness the magnitude of the flow, cf. Fig. 2. Figure 3 also compares the results obtained with two different derivative approximations: (i) A standard scheme and (ii) our proposed adaptive HRT scheme. To measure the quality of the flow fields, we use the *average angular error* (*AAE*) measure [2] defined as

$$\mathrm{AAE}(u, v, \hat{u}, \hat{v}) := \frac{1}{n_x n_y} \cdot \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \arccos\left(\frac{u_{i,j}\hat{u}_{i,j} + v_{i,j}\hat{v}_{i,j} + 1}{(u^2_{i,j} + v^2_{i,j} + 1)(\hat{u}^2_{i,j} + \hat{v}^2_{i,j} + 1)}\right), \tag{22}$$

where $(\hat{u}, \hat{v})^\top$ denotes the ground truth flow field, and n_x, n_y denote the number of pixels in x- and y-direction, respectively. Note that the AAE measures the average angle between two spatio-temporal flow vectors. Thus, if the estimated vector is a multiple of the true vector, the angle is *not* equal to zero, as often assumed incorrectly. In Fig. 3

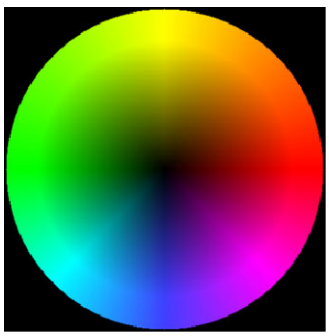

Fig. 2 (Color online) Colour code of the displacement vectors

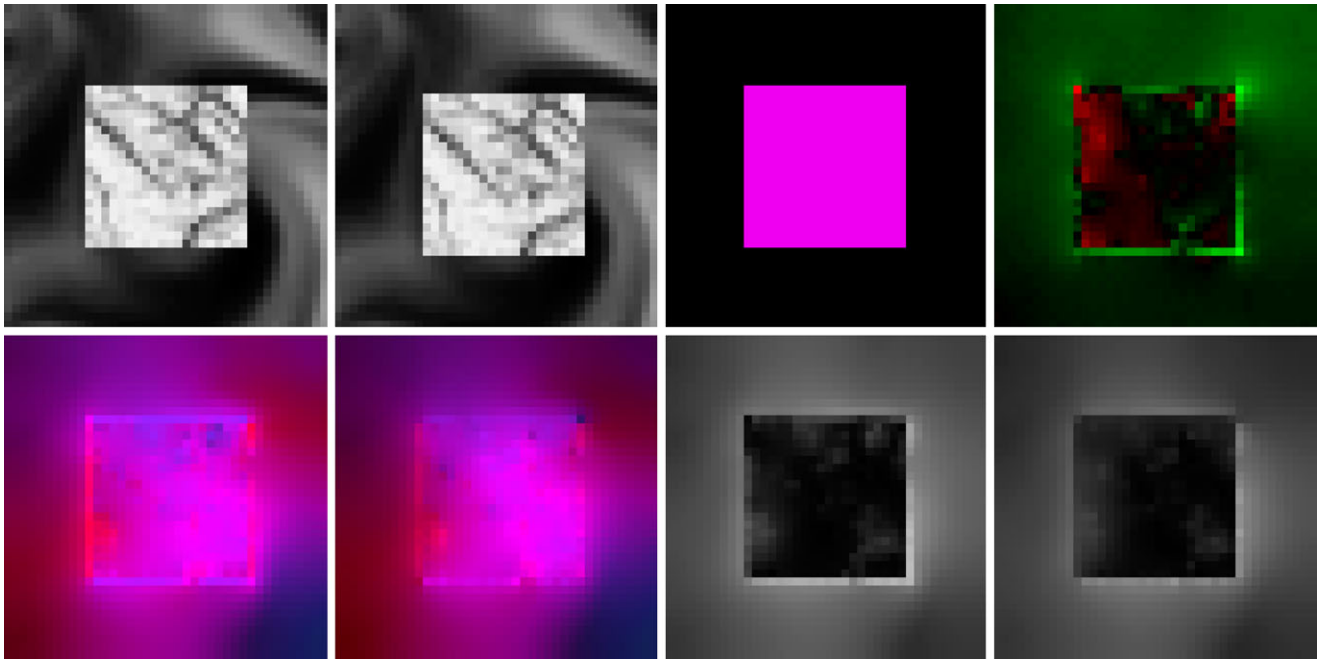

Fig. 3 (Color online) Results for the method of Horn and Schunck on our *Rectangle* sequence. We compare a standard derivative approximation to our adaptive HRT scheme. *First row*: (**a**) First frame. (**b**) Second frame. (**c**) Ground truth. (**d**) Difference of the error maps in (**g**) and (**h**). Green marks improvements, red impairments of our HRT scheme compared to the standard scheme. *Second row*: (**e**) Flow field with a standard derivative approximation. (**f**) Same with our adaptive HRT scheme. **(g)** Error map for (**e**). (**h**) Error map for (**f**)

Table 1 Error measures (AAE) for several sequences using the method of Horn and Schunck. We compare a standard derivative approximation to our proposed adaptive HRT scheme

	Rectangle	Marble	Yosemite	Street
Standard	31.93°	9.11°	10.72°	9.38°
HRT	28.40°	8.50°	9.53°	9.00°

we additionally show two error maps that visualise the AAE (brighter pixels correspond to larger errors). To ease comparison of these maps, we visualise in Fig. 3(d) the difference of the error maps. Here, green encodes improvements of the HRT over the standard scheme and red encodes impairments. The brightness of the pixels corresponds to the magnitude of the error. Inspecting the error maps and their difference, the expected benefits of the HRT scheme become obvious: Especially at the lower and right boundary of the rectangle, i.e. at regions with a large image discontinuity, the HRT scheme reduces the error. This observation is validated by the AAE measures that we show in Table 1.

This table also shows the AAE for more complex sequences, like the *Marble*,[1] the *Yosemite without clouds*,[2] and the *Street* sequence.[3] For these sequences, a comparison of error maps resulting from a standard derivative approximation and our adaptive scheme is shown in Fig. 4. It turns out that the HRT scheme allows to decrease the errors in regions with strong discontinuities. For *Marble* this is the case at the ground floor, for *Yosemite* we see an improvement at the lower left boundary, and for *Street* the error decreases at the leaves of the tree.

We wish to note that for this and also for the upcoming experiments, the predictor solution $(\tilde{u}, \tilde{v})$ is now computed with the same smoothness weight α as the final result. This becomes possible as the value of α is usually large enough for standard test sequences.

Finally, we want to evaluate if the upwind directions were correctly estimated using the predictor solution. In Fig. 5 we visualise the errors by greyscale maps. Here, black pixels mark locations where the upwind directions were correctly estimated, i.e. the sign of the predictor solution equals the sign of the ground truth flow. The grey pixels mark locations with an error in the direction estimation. The brightness of the pixels corresponds to the magnitude of the ground truth flow at this location. The latter is motivated by the observation that errors in the upwind direction estimation are on the one hand hardly avoidable, but also less harmful, if the flow is very small. Inspecting the error maps in Fig. 5, we

[1] Available at http://i21www.ira.uka.de/image_sequences.

[2] Available at http://www.cs.brown.edu/~black/images.html.

[3] Available at http://of-eval.sourceforge.net.

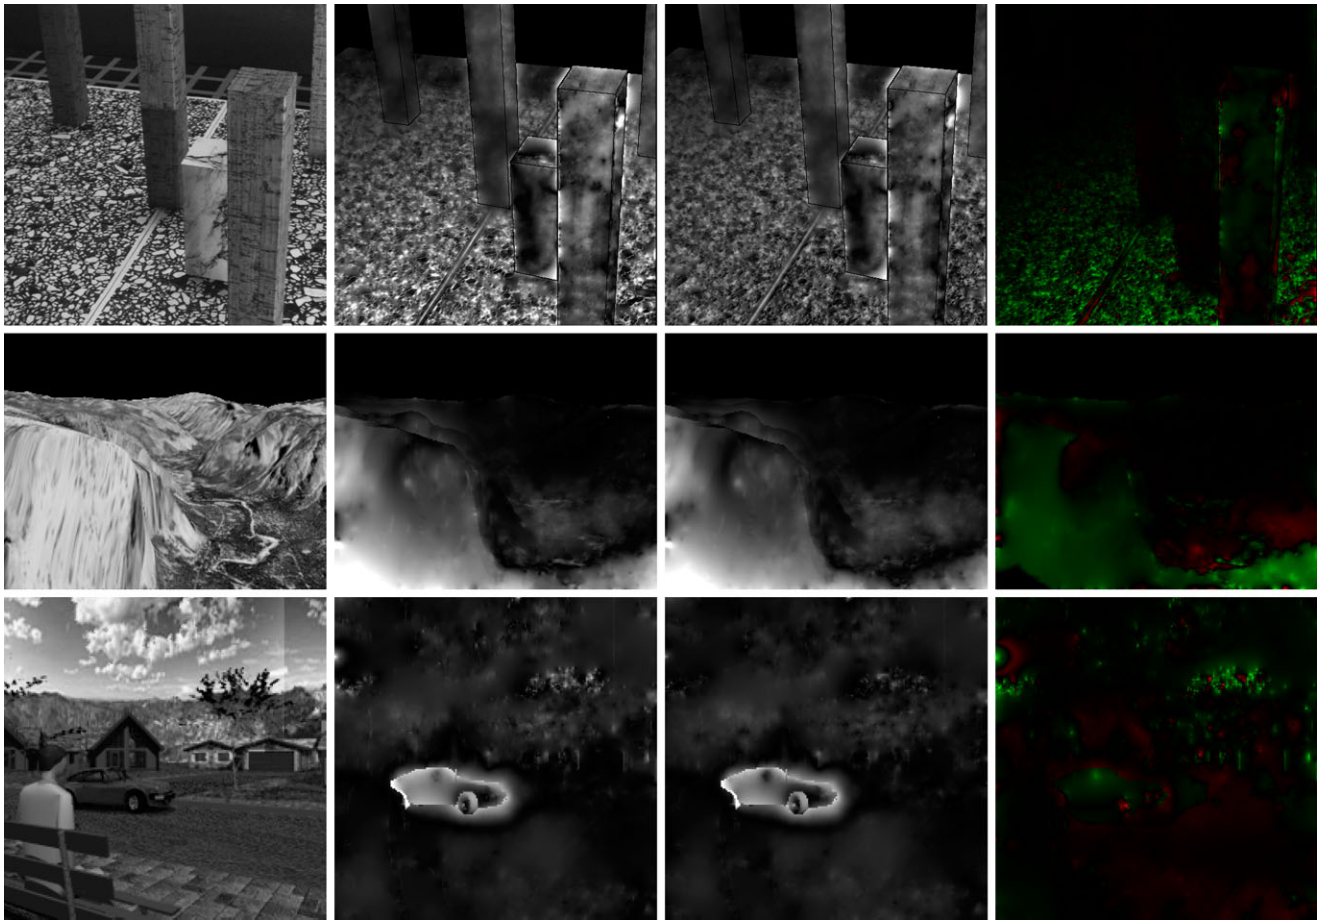

Fig. 4 (Color online) Results for the method of Horn and Schunck on several sequences. We compare a standard derivative approximation to our proposed adaptive HRT scheme. *From top to bottom*: *Marble*, *Yosemite without clouds* and *Street* sequence. *From left to right*: First frame, error map with a standard derivative approximation, same for our adaptive HRT scheme, difference of the two error maps (green marks improvements, red impairments of our HRT scheme compared to the standard scheme)

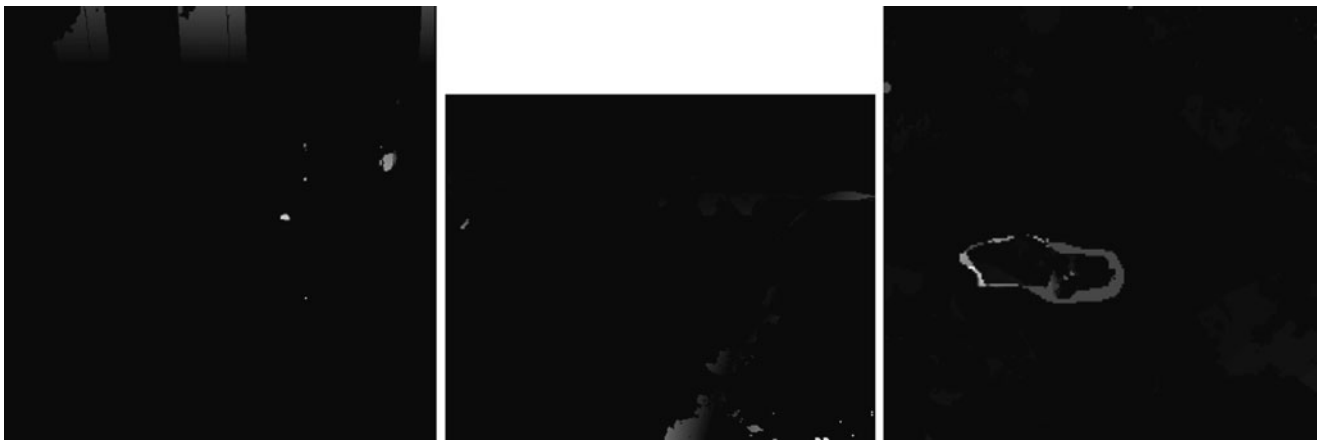

Fig. 5 (Color online) Error in the estimated upwind directions for the sequences shown in Fig. 4. *Black pixels* mark correct direction. *Grey pixels* mark errors, where the brightness encodes to the magnitude of the ground truth flow at the corresponding location. *From left to right*: *Marble*, *Yosemite without clouds* and *Street* sequence

verify that the computed predictor allows for a very reliable estimation of the correct upwind directions. Note that we do not show an error map for our previous experiment in Fig. 3 as the error is zero there. For the upcoming experiments, we also refrain from showing error maps for the upwind direction as this is difficult when using the proposed multiscale framework where flow increments are computed at each level.

4.3 Optic Flow: The Method of Brox et al.

A more recent and also more accurate optic flow method is the one of Brox et al. [7]. It extends the already presented approach of Horn and Schunck in several ways as briefly sketched in the following.

Data Term and Smoothness Term Brox et al. propose the data term $M(f,u,v)$ given by

$$\Psi_M\big(|f(x+u,y+v,t+1)-f(x,y,t)|^2 + \gamma|\nabla f(x+u,y+v,t+1)-\nabla f(x,y,t)|^2\big). \tag{23}$$

Let us review its innovations compared to the data term of Horn and Schunck (18). (i) To handle large displacements, the linearisation in the data term is postponed to the numerical part. There, a coarse-to-fine multiscale warping framework is used that computes small flow increments on each warping level via a linearised approach. The sum of all these increments then gives the final flow field. (ii) In addition to the brightness constancy assumption, also the gradient constancy assumption is imposed. It models the assumption that image gradients are invariant under their displacement, i.e., $\nabla f(x+u,y+v,t+1)=\nabla f(x,y,t)$, which renders the approach robust under varying illumination conditions. The contribution of the gradient constancy assumption to the data term is steered by the parameter $\gamma>0$. (iii) Finally, a robust penaliser function $\Psi(s^2)$ is used. This function is preferably positive, increasing, subquadratic and strictly convex in s. Whereas the first properties ensure robustness w.r.t. outliers caused by noise or occlusions, the latter guarantees that a unique minimum of the underlying energy exists. Brox et al. propose $\Psi_M(s^2):=\sqrt{s^2+\varepsilon^2}$, with a small regularisation parameter $\varepsilon=0.001$. This results in a modified differentiable L_1 penalisation.

The *smoothness term* uses the same penaliser as the data term:

$$V(\nabla u,\nabla v)=\Psi_V(|\nabla u|^2+|\nabla v|^2), \tag{24}$$

where $\Psi_V(s^2)=\Psi_M(s^2)$. This comes down to total variation (TV) penalisation [26], and yields a discontinuity-preserving behaviour.

Euler-Lagrange Equations We first introduce in accordance to [7] the abbreviations

$$f_{**}:=\partial_{**}f(x+u,y+v,t+1), \tag{25}$$

$$f_z:=f(x+u,y+v,t+1)-f(x,y,t), \tag{26}$$

$$f_{*z}:=\partial_* f(x+u,y+v,t+1)-\partial_* f(x,y,t), \tag{27}$$

where $**\in\{x,y,xx,xy,yy\}$, $*\in\{x,y\}$, and the variable z is used to emphasise the use of temporal differences in contrast to temporal derivatives. Using these abbreviations, the Euler-Lagrange equations for the method of Brox et al. are given by

$$\Psi'_M\big(f_z^2+\gamma(f_{xz}^2+f_{yz}^2)\big)\cdot(f_x f_z+\gamma(f_{xx}f_{xz}+f_{xy}f_{yz})) - \alpha\,\mathrm{div}\big(\Psi'_V\big(|\nabla u|^2+|\nabla v|^2\big)\nabla u\big)=0, \tag{28}$$

$$\Psi'_M\big(f_z^2+\gamma(f_{xz}^2+f_{yz}^2)\big)\cdot(f_y f_z+\gamma(f_{yy}f_{yz}+f_{xy}f_{xz})) - \alpha\,\mathrm{div}\big(\Psi'_V\big(|\nabla u|^2+|\nabla v|^2\big)\nabla v\big)=0. \tag{29}$$

To solve above equations, we follow the strategy proposed in [7]. Here, two nested fixpoint iterations are performed, which reduce the problem to the solution of a series of systems of linear equations. For an efficient solution of the latter, we use the multigrid framework proposed by Bruhn et al. [9].

4.3.1 The HRT Scheme for the Method of Brox et al.

Inspecting the Euler-Lagrange equations (28) and (29), we realise that due to the use of the gradient constancy assumption, also the second order and mixed derivatives f_{xx}, f_{yy}, f_{xy}, f_{xz} and f_{yz} occur.

First of all, this requires to define a smoothness measure for the mixed xy-direction. Given the smoothness measures Θ_x and Θ_y, we define the mixed expression as $\Theta_{xy}:=\Theta_x+\Theta_y$.

Second Order Derivative Approximations More involved are the high-order and the (one-sided) low-order approximations of the second order derivatives. These are now briefly presented, relying on the finite difference operators defined in (8).

High-Order The high-order approximations of f_{xx} and f_{yy} are defined in accordance to (10).

The mixed derivative $f_{xy}=\partial_y f_x$ is approximated in the finite difference case as

$$(f_{xy})^k_{i,j}:=\mathcal{D}^0_y\big(\mathcal{D}^0_x f^k_{i,j}\big)=\mathcal{D}^0_y\left(\frac{f^k_{i+1,j}-f^k_{i-1,j}}{2h_x}\right) = \frac{f^k_{i+1,j+1}-f^k_{i-1,j+1}-(f^k_{i+1,j-1}-f^k_{i-1,j-1})}{4h_xh_y}. \tag{30}$$

An averaged version taking into account both time levels is then obtained via

$$(f_{xy})_{i,j}^{k+\frac{1}{2}} := \frac{1}{2}\big(\mathcal{D}_y^0\big(\mathcal{D}_x^0 f_{i,j}^k\big) + \mathcal{D}_y^0\big(\mathcal{D}_x^0 f_{i,j}^{k+1}\big)\big). \tag{31}$$

Similarly, we define f_{xz} as

$$\begin{aligned}(f_{xz})_{i,j}^k &:= \mathcal{D}_z^+\big(\mathcal{D}_x^0 f_{i,j}^k\big)\\ &= \frac{1}{2h_x}\big(f_{i+1,j}^{k+1} - f_{i-1,j}^{k+1} - (f_{i+1,j}^k - f_{i-1,j}^k)\big),\end{aligned} \tag{32}$$

where $\mathcal{D}_z^+ f_{i,j}^k := f_{i,j}^{k+1} - f_{i,j}^k$ denotes the temporal difference. Analogously we proceed for f_{yz}.

Low-Order In the low-order upwind case, the sign of the predictor shall decide which one-sided difference should be used: We consider $\tilde{u}$ for x-derivatives and $\tilde{v}$ for y-derivatives, respectively.

For approximating f_{xx} we make use of the corresponding upwind difference for f_x:

$$\tilde{u} > 0 \quad : \quad (f_{xx})_{i,j}^k := \mathcal{D}_x^-\big(\mathcal{D}_x^- f_{i,j}^k\big), \tag{33}$$

$$\tilde{u} < 0 \quad : \quad (f_{xx})_{i,j}^k := \mathcal{D}_x^+\big(\mathcal{D}_x^+ f_{i,j}^k\big). \tag{34}$$

For $\tilde{u} = 0$, we use the corresponding high-order approximation. For f_{yy}, we proceed accordingly, taking into account the predictor $\tilde{v}$.

For the mixed derivative f_{xy} we have to use the two predictors $\tilde{u}$ and $\tilde{v}$. If a predictor is equal to zero, we use the corresponding high-order approximation, and if it is non-zero, its sign determines which one-sided upwind approximation to use. This leads to the case distinction summarised in Table 2. If $\tilde{u} = 0$ and $\tilde{v} = 0$ holds, we again use the high-order approximation of f_{xy}.

For f_{xz} we use the same approach as presented above in the high-order case but just use one-sided upwind differences for approximating f_x. This gives

$$\begin{aligned}\tilde{u} > 0 \quad &: \quad (f_{xz})_{i,j}^k := \mathcal{D}_z^+\big(\mathcal{D}_x^- f_{i,j}^k\big),\\ \tilde{u} < 0 \quad &: \quad (f_{xz})_{i,j}^k := \mathcal{D}_z^+\big(\mathcal{D}_x^+ f_{i,j}^k\big).\end{aligned} \tag{35}$$

Accordingly we proceed for f_{yz}.

Table 2 Upwind-type (one-sided) discretisations of the mixed derivative f_{xy}

Case	Discretisation of $(f_{xy})_{i,j}^k$
$\tilde{u} = 0, \tilde{v} > 0$	$\mathcal{D}_y^-\big(\mathcal{D}_x^0 f_i^k\big)$
$\tilde{u} = 0, \tilde{v} < 0$	$\mathcal{D}_y^+\big(\mathcal{D}_x^0 f_i^k\big)$
$\tilde{u} > 0, \tilde{v} = 0$	$\mathcal{D}_y^0\big(\mathcal{D}_x^- f_i^k\big)$
$\tilde{u} < 0, \tilde{v} = 0$	$\mathcal{D}_y^0\big(\mathcal{D}_x^+ f_i^k\big)$
$\tilde{u} > 0, \tilde{v} > 0$	$\mathcal{D}_y^-\big(\mathcal{D}_x^- f_i^k\big)$
$\tilde{u} > 0, \tilde{v} < 0$	$\mathcal{D}_y^+\big(\mathcal{D}_x^- f_i^k\big)$
$\tilde{u} < 0, \tilde{v} > 0$	$\mathcal{D}_y^-\big(\mathcal{D}_x^+ f_i^k\big)$
$\tilde{u} < 0, \tilde{v} < 0$	$\mathcal{D}_y^+\big(\mathcal{D}_x^+ f_i^k\big)$

4.3.2 *Numerical Experiments for the Method of Brox et al.*

We now show experiments for our adaptive HRT scheme used within the method of Brox et al. [7]. Due to the postponed linearisation, the robust data term and the discontinuity-preserving smoothness term, this method achieves reasonable results for more difficult sequences, like the ones from the popular Middlebury database [1].[4]

In Fig. 6, we show results for the *Urban3* sequence. We see that also for the method of Brox et al., the HRT scheme allows to improve the results at locations with strong discontinuities (marked in the images). In this context, we also refer to Fig. 7 where we show a plot of the gradient magnitude

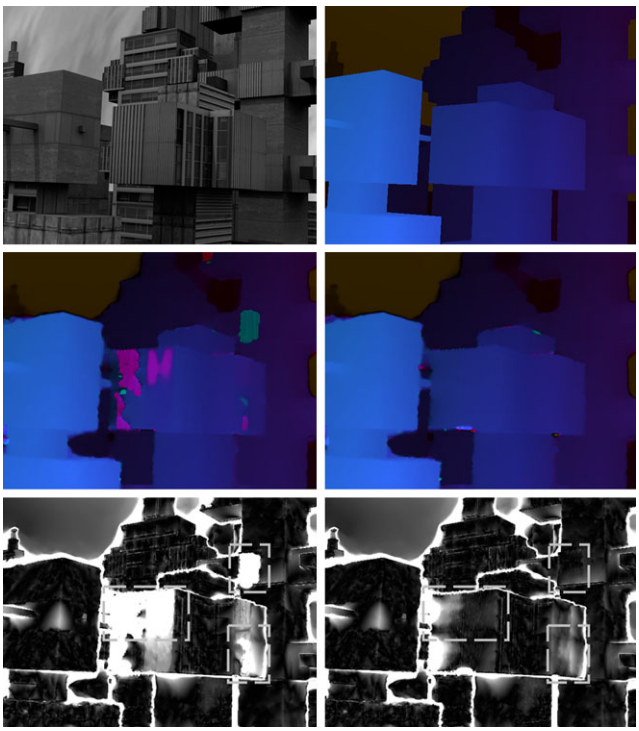

Fig. 6 (Color online) Results for the method of Brox et al. on the *Urban3* sequence. We compare a standard derivative approximation to our proposed adaptive HRT scheme. *First row*: (**a**) Reference frame (frame 10). (**b**) Ground truth flow field. *Second row*: (**c**) Flow field with a standard derivative approximation. (**d**) Same with our adaptive HRT scheme. *Third row*: (**e**) Error map with a standard derivative approximation. (**f**) Same with our adaptive HRT scheme. Boxes indicate regions of significantly better results with the HRT scheme

[4] Available at http://vision.middlebury.edu/flow/data/.

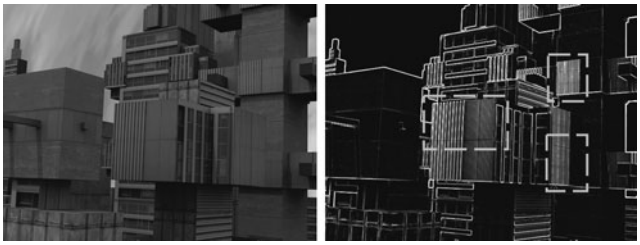

Fig. 7 Plot of the gradient magnitude for the *Urban3* sequence. *Left* (**a**): Reference frame. *Right* (**b**): Corresponding gradient magnitude, scaled to the range from 0 to 255. The *boxes* indicate regions where the HRT scheme performs significantly better than the standard approach, cf. Fig. 6. We observe that these are regions featuring strong gradients

Table 3 Error measures (AAE) for several Middlebury sequences and the method of Brox et al. We compare a standard derivative approximation scheme to a pure upwind scheme ($\Phi(\Theta_i) = 0$), and our proposed adaptive HRT scheme

	Standard	Upwind	HRT
Rubberwhale	3.69°	4.87°	3.55°
Hydrangea	2.18°	2.39°	2.16°
Dimetrodon	1.94°	3.06°	1.88°
Grove 2	2.49°	3.13°	2.46°
Grove 3	6.17°	6.75°	6.21°
Urban 2	2.84°	3.77°	2.74°
Urban 3	5.71°	4.58°	4.11°
Venus	4.79°	4.58°	4.45°

to support the latter observation. The qualitative improvement is confirmed by the corresponding AAE measures in Table 3. This table also lists other Middlebury sequences and gives errors for an upwind scheme only using one-sided low-order approximations ($\Phi(\Theta_i) = 0$). Analysing the results in Table 3 shows: (i) For 7 out of the 8 test sequences, the HRT scheme improves the quality over the standard scheme. Only for *Grove 3*, the standard scheme is slightly better. (ii) For the considered complex sequences, the blending between high-order and low-order approximations of the HRT scheme gives significantly better results than a pure upwind scheme. Note that for some sequences, e.g. *Dimetrodon*, the upwind scheme produces significantly worse results. This occurs in sequences with large smooth areas in the images. There, central derivative approximations are best suited. Consequently, an upwind scheme that uses the less appropriate first-order approximation gives bad results.

Let us also stress in this context, that the basic method of Brox et al. gives already results of low errors. Therefore, improvements in the order documented here are considerable.

4.4 Stereo Vision: Basics

The task of stereo reconstruction is also a correspondence problem, similar to optic flow.

In the stereo context, we are given an image pair $f_l(x, y)$, $f_r(x, y)$, denoting the left and right view of a static scene, respectively. The absolute value of the displacement field $(u, v)^\top$ between f_l and f_r is called *disparity* d. As the disparity is directly related to the depth of the corresponding scene point, it is a fundamental part of 3D reconstruction methods, cf. [15].

In contrast to optic flow, the displacements in the stereo context cannot be arbitrary. In fact, the corresponding point of a pixel in the first image has to lie on a specific line, the *epipolar line* [15], in the second image. For simplicity, we restrict ourselves to a basic, but often considered scenario: If the two cameras are *orthoparallel* to each other, or if the image pair has been rectified beforehand, the displacements are purely horizontal. This allows to reformulate the stereo problem as an optic flow problem with zero vertical displacement ($v = 0$).

4.5 Stereo Vision: The Method of Slesareva et al.

The variational stereo method of Slesareva et al. [28] is based on the optic flow approach of Brox et al., but enforces corresponding pixels to lie on the epipolar lines. In our orthoparallel scenario, the method of Slesareva et al. can thus be obtained from the optic flow method of Brox et al. (see Sect. 4.3) by setting $v = 0$.

In order to use the notation from the optic flow case, we consider the left and right images as two snapshots of an image sequence taken at time t and $t + 1$, respectively. Formally, $f_l(x, y) \equiv f(x, y, t)$ and $f_r(x, y) \equiv f(x, y, t + 1)$. This yields the energy functional

$$E(u) = \int_{\Omega_2} [M(f, u) + \alpha V(\nabla u)] \mathrm{d}x\, \mathrm{d}y, \tag{36}$$

whose minimiser u gives the sought disparity by $d = |u|$. The data term $M(f, u)$ reads as

$$\Psi_M\big(|f(x + u, y, t + 1) - f(x, y, t)|^2 + \gamma |\nabla f(x + u, y, t + 1) - \nabla f(x, y, t)|^2\big), \tag{37}$$

and the smoothness term is given by

$$V(\nabla u) = \Psi_V(|\nabla u|^2). \tag{38}$$

To minimise the energy (36), one solves the single Euler-Lagrange equation

$$\Psi'_M\big(f_z^2 + \gamma(f_{xz}^2 + f_{yz}^2)\big) \cdot (f_x f_z + \gamma(f_{xx} f_{xz} + f_{xy} f_{yz})) - \alpha \operatorname{div}\big(\Psi'_V\big(|\nabla u|^2\big)\nabla u\big) = 0, \tag{39}$$

using the abbreviations

$$f_{**} := \partial_{**} f(x + u, y, t + 1), \tag{40}$$

$$f_z := f(x+u, y, t+1) - f(x, y, t), \tag{41}$$

$$f_{*z} := \partial_* f(x+u, y, t+1) - \partial_* f(x, y, t), \tag{42}$$

where $** \in \{x, xx, xy\}$ and $* \in \{x, y\}$. The solution of (39) is in accordance to the solution of (28).

The adaption of the HRT scheme works in accordance to the optic flow case described in Sect. 4.3.1.

4.5.1 *Numerical Experiments for the Method of Slesareva et al.*

Our final experiments show that our adaptive HRT scheme is also beneficial for variational stereo. As test data, we used stereo pairs from the Middlebury stereo page [27].[5] To measure the quality of the disparity estimates, we use the *bad pixel error* (BPE) [27]. It gives the percentage of pixels that deviate more than a threshold δ_d from the ground truth $\hat{u}$, yielding the definition

$$\mathrm{BPE}(u, \hat{u},) := \frac{100}{n_x n_y} \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} T(|u_i - \hat{u}_i| > \delta_d), \tag{43}$$

where $T(b) = 1$ if $b =$ true, and 0 else. As proposed in [27], we set $\delta_d = 1$.

In Fig. 8, we show results for the *Plastic* pair. Again, the HRT scheme improves the results at locations with large discontinuities, which are marked in the bad pixel maps in Fig. 8(e) and (f). The corresponding BPE measures are summarised in Table 4 that, as before, lists also other Middlebury pairs and errors for an upwind scheme. Similar to the optic flow case, the HRT scheme gives the best results, compared to a standard and a pure upwind scheme.

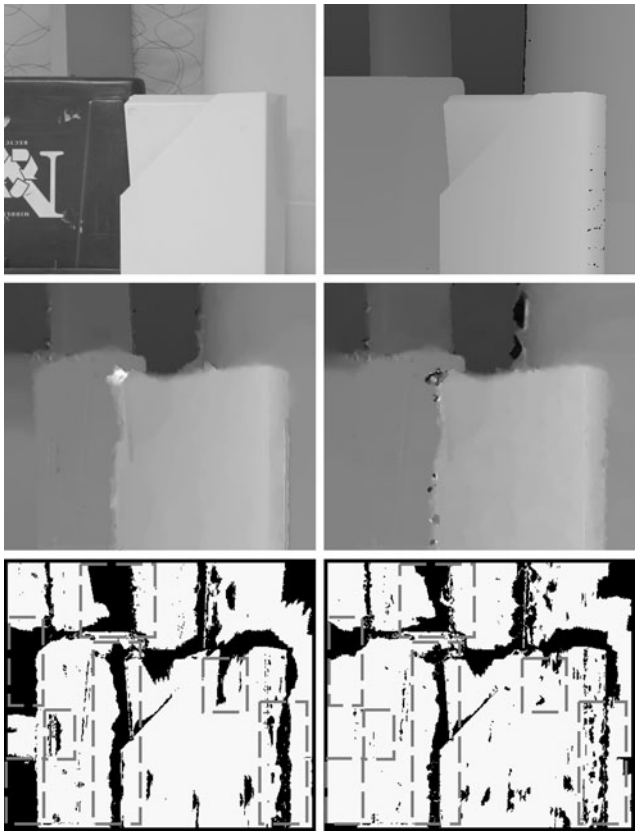

Fig. 8 Results for the method of Slesareva et al. on the *Plastic* pair. We compare a standard derivative approximation to our proposed adaptive HRT scheme. *First row*: (**a**) Left image. (**b**) Ground truth. *Second row*: (**c**) Disparity with a standard derivative approximation scheme. (**d**) Same with our adaptive HRT scheme. *Third row*: (**e**) Bad pixel map with a standard derivative approximation scheme (bad pixels are black). (**f**) Same with our adaptive HRT scheme

Table 4 Error measures (BPE) for several Middlebury image pairs and the method of Slesareva et al. We compare a standard derivative approximation scheme to a pure upwind scheme ($\Phi(\Theta_i) = 0$), and our proposed adaptive HRT scheme

	Standard	Upwind	HRT
Plastic	25.85%	21.35%	18.85%
Tsukuba	7.76%	7.59%	7.58%
Venus	3.06%	2.78%	2.77%
Teddy	17.45%	16.94%	16.75%
Cones	15.34%	15.34%	15.32%

5 Summary and Conclusion

Our paper is the first approach that exploits the structural relationship between data terms in variational approaches for correspondence problems and hyperbolic differential equations. This has led to novel sophisticated numerical schemes for the approximation of spatial image derivatives in correspondence problems. It relies on the idea to switch the role of known and unknown data in the data term, which leads to a hyperbolic colour equation. This equation is discretised with specific upwind schemes that are appropriate for the application to correspondence problems.

Note that our goal was not to introduce novel, more accurate *models*, which has been done in numerous publications in the last three decades. Our goal was to introduce a new class of *better discretisations*. They can be useful for *all* approaches that formulate correspondence problems in terms of differential expressions. In order to illustrate this general benefit, we have applied it to three prototypical methods: The optic flow approaches of Horn and Schunck [16] and of Brox et al. [7], and the stereo method of Slesareva et al. [28]. Our experiments demonstrate that the novel discretisations allow to improve the quality of results in a similar order as is usually obtained by model improvements.

[5] Available at http://vision.middlebury.edu/stereo/data/.

Although we have focused on variational models, we are convinced that these numerical ideas are more general and can also be useful for other differential methods for correspondence problems in computer vision. Part of our ongoing research is thus concerned with testing if our proposed schemes also yield better discretisations for local methods. These methods do not use a smoothness term but assume constancy of the flow field in some local neighbourhood. Popular examples for such methods are the Lucas-Kanade method [20], the structure tensor approach of Bigün et al. [5] and their numerous variants.

Acknowledgements Henning Zimmer gratefully acknowledges funding by the *International Max-Planck Research School (IMPRS)*.

References

1. Baker, S., Roth, S., Scharstein, D., Black, M.J., Lewis, J.P., Szeliski, R.: A database and evaluation methodology for optical flow. In: Proc. 2007 IEEE International Conference on Computer Vision. Rio de Janeiro, Brazil. IEEE Computer Society Press, Los Alamitos (2007)
2. Barron, J.L., Fleet, D.J., Beauchemin, S.S.: Performance of optical flow techniques. Int. J. Comput. Vis. **12**(1), 43–77 (1994)
3. Ben-Ari, R., Sochen, N.: Variational stereo vision with sharp discontinuities and occlusion handling. In: Proc. 2007 IEEE International Conference on Computer Vision. Rio de Janeiro, Brazil. IEEE Computer Society Press, Los Alamitos (2007)
4. Bertero, M., Poggio, T.A., Torre, V.: Ill-posed problems in early vision. Proc. IEEE **76**(8), 869–889 (1988)
5. Bigün, J., Granlund, G.H., Wiklund, J.: Multidimensional orientation estimation with applications to texture analysis and optical flow. IEEE Trans. Pattern Anal. Mach. Intell. **13**(8), 775–790 (1991)
6. Black, M.J., Anandan, P.: The robust estimation of multiple motions: parametric and piecewise smooth flow fields. Comput. Vis. Image Underst. **63**(1), 75–104 (1996)
7. Brox, T., Bruhn, A., Papenberg, N., Weickert, J.: High accuracy optical flow estimation based on a theory for warping. In: Pajdla, T., Matas, J. (eds.) Computer Vision—ECCV 2004, Part IV. Lecture Notes in Computer Science, vol. 3024, pp. 25–36. Springer, Berlin (2004)
8. Brown, M., Burschka, D., Hager, G.: Advances in computational stereo. IEEE Trans. Pattern Anal. Mach. Intell. **25**(8), 993–1008 (2003)
9. Bruhn, A., Weickert, J., Kohlberger, T., Schnörr, C.: A multigrid platform for real-time motion computation with discontinuity-preserving variational methods. Int. J. Comput. Vis. **70**(3), 257–277 (2006)
10. Courant, R., Friedrichs, K., Lewy, H.: Über die partiellen Differenzengleichungen der mathematischen Physik. Math. Ann. **100**(1), 32–74 (1928)
11. Elsgolc, L.: Calculus of Variations. Pergamon Press, Oxford (1962)
12. Evans, L.C.: Partial Differential Equations. Oxford University Press, Oxford (1998)
13. Fleet, D.J., Weiss, Y.: Optical flow estimation. In: Paragios, N., Chen, Y., Faugeras, O. (eds.) Handbook of Mathematical Models in Computer Vision, Chap. 15, pp. 239–258. Springer, Berlin (2006)
14. Godlewski, E., Raviart, P.-A.: Hyperbolic Systems of Conservation Laws. Mathématiques et Applications. Ellipses, Paris (1991)
15. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press, Cambridge (2000)
16. Horn, B., Schunck, B.: Determining optical flow. Artif. Intell. **17**, 185–203 (1981)
17. Klette, R., Schlüns, K., Koschan, A.: Computer Vision: Three-Dimensional Data from Images. Springer, Singapore (1998)
18. LeVeque, R.J.: Numerical Methods for Conservation Laws. Birkhäuser, Basel (1992)
19. LeVeque, R.J.: Finite Volume Methods for Hyperbolic Problems. Cambridge University Press, Cambridge (2002)
20. Lucas, B., Kanade, T.: 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 (1981)
21. Mansouri, A.R., Mitiche, A., Konrad, J.: Selective image diffusion: application to disparity estimation. In: Proc. 1998 IEEE International Conference on Image Processing, vol. 3. Chicago, IL, pp. 284–288 (1998)
22. Marquina, A., Osher, S.: Explicit algorithms for a new time dependent model based on level set motion for nonlinear deblurring and noise removal. SIAM J. Sci. Comput. **22**(2), 387–405 (2000)
23. Mémin, E., Pérez, P.: Hierarchical estimation and segmentation of dense motion fields. Int. J. Comput. Vis. **46**(2), 129–155 (2002)
24. Morton, K.W., Mayers, L.M.: Numerical Solution of Partial Differential Equations. Cambridge University Press, Cambridge (1994)
25. Nir, T., Bruckstein, A.M., Kimmel, R.: Over-parameterized variational optical flow. Int. J. Comput. Vis. **76**(2), 205–216 (2008)
26. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. Physica D **60**, 259–268 (1992)
27. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. Int. J. Comput. Vis. **47**(1–3), 7–42 (2002)
28. Slesareva, N., Bruhn, A., Weickert, J.: Optic flow goes stereo: A variational method for estimating discontinuity-preserving dense disparity maps. In: Kropatsch, W., Sablatnig, R., Hanbury, A. (eds.) Pattern Recognition. Lecture Notes in Computer Science, vol. 3663, pp. 33–40. Springer, Berlin (2005)
29. Toro, E.F.: Riemann Solvers and Numerical Methods for Fluid Dynamics, 2nd edn. Springer, Berlin (1999)
30. Trucco, E., Verri, A.: Introductory Techniques for 3-D Computer Vision. Prentice Hall, Englewood Cliffs (1998)
31. Wedel, A., Cremers, D., Pock, T., Bischof, H.: Structure- and motion-adaptive regularization for high accuracy optic flow. In: Proc. 2009 IEEE International Conference on Computer Vision. Kyoto, Japan. IEEE Computer Society Press, Los Alamitos (2009)
32. Werlberger, M., Pock, T., Bischof, H.: Motion estimation with non-local total variation regularization. In: Proc. 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. San Francisco, CA, USA. IEEE Computer Society Press, Los Alamitos (2010)
33. Xu, L., Jia, J., Matsushita, Y.: Motion detail preserving optical flow estimation. In: Proc. 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. San Francisco, CA, USA. IEEE Computer Society Press, Los Alamitos (2010)
34. Young, D.M.: Iterative Solution of Large Linear Systems. Dover, New York (2003)
35. Zimmer, H., Breuß, M., Weickert, J., Seidel, H.-P.: Hyperbolic numerics for variational approaches to correspondence problems. In: Tai, X.-C., et al. (eds.) Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science, vol. 5567, pp. 636–647. Springer, Berlin (2009)
36. Zimmer, H., Bruhn, A., Weickert, J., Valgaerts, L., Salgado, A., Rosenhahn, B., Seidel, H.-P.: Complementary optic flow. In: Cremers, D., Boykov, Y., Blake, A., Schmidt, F.R. (eds.) Energy Min-

imization Methods in Computer Vision and Pattern Recognition—EMMCVPR. Lecture Notes in Computer Science, vol. 5681, pp. 207–220. Springer, Berlin (2009)

Michael Breuß received a Diploma in Mathematics from the University of Hamburg, Germany, in 1998. After obtaining there also his PhD in 2001 and holding for some time a position as an Assistant Professor, he worked as a Post-Doctoral Researcher at the Université Bordeaux 1, France, and at the Technical University Braunschweig, Germany. There, he finished his Habilitation in Mathematics in 2006. Since April 2006, he is an Assistant Professor in the Mathematical Image Analysis Group at the Saarland University, Germany. He also held a temporary position as a Professor at the Johannes Gutenberg University in Mainz, Germany. His research interests are centred around numerical methods for partial differential equations in image processing and computer vision.

Henning Zimmer received a BSc and a MSc degree in Computer Science from Saarland University, Germany, in 2006 and 2007, respectively. In October 2010, he joined the Mathematical Image Analysis Group as a PhD student. His scientific interests focus on optic flow and variational stereo reconstruction.

Joachim Weickert is Professor of Mathematics and Computer Science at Saarland University, Saarbrücken, Germany, where he heads the Mathematical Image Analysis Group. He graduated and obtained his PhD from the University of Kaiserslautern, Germany, in 1991 and 1996. He worked as Post-Doctoral Researcher at the University Hospital of Utrecht, The Netherlands, and the University of Copenhagen, Denmark, and as Assistant Professor at the University of Mannheim, Germany.

Joachim Weickert has developed many models and efficient algorithms for image processing and computer vision using partial differential equations and optimisation principles. In particular he has contributed to diffusion filtering, optical flow computation, processing of tensor fields, and image compression. His scientific work covers more than 220 refereed publications. He has served in the editorial boards of nine international journals and given more than 120 invited talks. In 2010 he has received the Gottfried Wilhelm Leibniz Prize which is the highest German science award.