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

Reihe Informatik
15/2000

Variational Optic Flow Computation with a Spatio-Temporal Smoothness Constraint

Joachim Weickert and Christoph Schnörr

Technical Report 15/2000
Computer Science Series
July 2000

This report is a revision of the Technical Report 5/1999, *Optic Flow Calculations with Nonlinear Smoothness Terms Extended into the Temporal Domain*, March 1999.

The technical reports of the CVGPR Group are listed at the web site `http://www.cvgpr.uni-mannheim.de/publications.html`

Variational Optic Flow Computation with a Spatio-Temporal Smoothness Constaint

Joachim Weickert and Christoph Schnörr

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

{Joachim.Weickert,Christoph.Schnoerr}@uni-mannheim.de
http://www.cvgpr.uni-mannheim.de/{weickert,schnoerr}

Abstract

Nonquadratic variational regularization is a well-known and powerful approach for the discontinuity-preserving computation of optic flow. In the present paper, we consider an extension of flow-driven spatial smoothness terms to spatio-temporal regularizers. Our method leads to a rotationally invariant and time symmetric convex optimization problem. It has a unique minimum that can be found in a stable way by standard algorithms such as gradient descent. Since the convexity guarantees global convergence, the result does not depend on the flow initialization. An iterative algorithm is presented that is not difficult to implement. Qualitative and quantitative results for synthetic and real-world scenes show that our spatio-temporal approach (i) improves optic flow fields significantly, (ii) smoothes out background noise efficiently, and (iii) preserves true motion boundaries. The computational costs are only 50 % higher than for a pure spatial approach applied to all subsequent image pairs of the sequence.

Keywords: optic flow, differential techniques, variational methods, spatio-temporal regularization, partial differential equations, finite differences, performance evaluation

1998 ACM Computing Classification: I.4.8, I.4.3, I.4.4.

1 Introduction

Variational methods in image processing and computer vision have attracted a lot of interest in recent years. They offer the advantage of providing a clear mathematical formalism for all model assumptions. Minimizing the resulting energy functionals gives solutions that are optimal with respect to the specified assumptions.

One of the earliest application areas of variational methods within computer vision is the estimation of optic flow [15]. The optic flow field of an image sequence describes the displacement of brightness patterns over time. Applications of optic flow range from vision-based robot navigation to second-generation video compression. Numerous methods for calculating optic flow have been proposed in the last two decades;

see e.g. the survey papers of Mitiche and Bouthemy [19], and Stiller and Konrad [29]. Performance evaluations of some of the most popular algorithms have been carried out by Barron *et al.* [5] and Galvin *et al.* [12]. These papers also showed that variational optic flow methods belong to the techniques that perform well. In contrast to many other optic flow methods, they offer the advantage of creating flow fields with 100 % density, such that no postprocessing by interpolation becomes necessary.

Variational optic flow methods have been pioneered by Horn and Schunck [15] and improved by Nagel [21] and many others. Approaches of this type calculate optic flow as the minimizer of an energy functional, which consists of a data term and a smoothness term. The data term involves optic flow constraints such as the assumption that corresponding pixels in different frames should reveal the same grey value. The smoothness term usually requires that the optic flow field should vary smoothly in space. Such a term may be modified in an *image-driven* way in order to suppress smoothing at image boundaries; see e.g. [1, 2, 21, 26, 28]. Recently also *flow-driven* modifications have been proposed which reduce smoothing at flow discontinuities [4, 9, 10, 17, 24, 27, 33]. These nonlinear methods have already led to rather good results in spite of the fact that the smoothness term imposed only *spatial* smoothness of the flow field. They work locally in time and do not make use of the temporal coherence within the sequence.

The goal of this paper is to investigate an extension of spatial flow-driven smoothness terms to *spatio-temporal flow-driven* regularizations. Such an extension makes consequent use of the available data, and it leads to equations which are hardly more complicated than in the pure spatial case. Our experiments on synthetic and real-world sequences, however, show that this approach leads to significantly more robust results.

Our paper is organized as follows. In Section 2 we review optic flow approaches with spatial smoothness terms, and Section 3 describes our novel method using a spatio-temporal smoothness constraint. A simple numerical algorithm is derived in Section 4, and Section 5 analyses the performance of our approach by applying it to synthetic and real-world image sequences. The paper is concluded with a summary in Section 6. We have presented a shorter preliminary version of this work at a national symposium [35].

Related work. While spatial smoothness assumptions are common in the optic flow literature, spatio-temporal approaches are significantly less frequent.

An interesting extension of a smoothness constraint into the temporal domain has been proposed by Nagel [22]. He derived the model for a spatio-temporal oriented smoothness constraint, but did not present any experiments. Nagel's constraint was *image-driven*, since it reduces smoothing across *image* discontinuities. Our approach is *flow-driven* due to the more direct constraint that smoothing at *flow* discontinuities should be reduced. For a more detailed account on image- and flow-driven regularizers and a well-posedness framework for both types we refer to [36].

Other temporal smoothness assumptions that have been studied in the literature include the work of Murray and Buxton [20], Black and Anandan [7], and Black [6]. Their assumptions lead to *nonconvex* optimization problems which may have many local minima and for which it is difficult to find algorithms that are both efficient and converge to a global minimum. Algorithms that converge to a global minimum (such as simulated annealing [18]) are computationally very expensive, while methods that are more efficient (such as graduated non-convexity algorithms [8]) may get trapped in local minima. Our method leads to a nonquadratic *convex* optimization problem. It has a unique global minimum that can be found in a reliable way by using standard techniques from convex optimization, for instance gradient descent methods. Since their convergence is global, every arbitrary flow initialization leads to the same solution: the global minimum of the functional. This property is an important quality of a robust algorithm.

Another difference between the methods considered in [6, 7, 20] and our approach is that our method uses a genuinely *continuous* formulation and derives a discrete algorithm afterwards by discretizing the corresponding partial differential equations. The continuous formulation has the advantage of being rotationally invariant. Applying the well-established theory of discretization methods allows us to derive a numerically consistent scheme. It guarantees that rotational invariance is fulfilled up to an error of order $O(\frac{1}{N^2})$ where N denotes the number of pixels in x or y direction. Results of this type cannot be established in genuinely discrete formulations. The discrete models in [6, 7, 20], for instance, approximate continuous processes that are *not* rotationally invariant.

The approaches of Black and Anandan [7] and Black [6] use a model which applies incremental minimization over time. Such a technique is highly useful for tasks such as robotics where images have to be processed online and only information from the past is accessible. Our approach, however, is designed for batch mode since it is symmetric with respect to past and future. Methods of this type are useful if a video sequence is to be processed offline. In this case there is information available both from the past and from the future and there is no reason for using only a part of it. Temporal symmetry also guarantees that the first frame is processed in the same way as the last one.

2 Spatial Smoothness Terms

Let us denote an image sequence by some real-valued function $f(x, y, z)$ where (x, y) denotes the location within some rectangular image domain Ω and $z \in [0, T]$ is the time. Many variational optic flow calculations determine the optic flow vector $(u, v)^T$ based on two assumptions:

1. Corresponding features are supposed to maintain their intensity over time. A differential formulation of this brightness constancy assumption leads to the *optic flow constraint (OFC) equation*

$$f_x u + f_y v + f_z = 0, \tag{1}$$

   where the subscripts denote partial derivatives. Numerous generalizations exist where multiple constraint equations are used, or different "conserved quantities" (replacing intensity) are considered; see e.g. [3, 11, 30, 31].
   Evidently, the single equation (1) is not sufficient to determine the two unknown functions u and v uniquely *(aperture problem)*. In order to obtain a unique flow field, a second constraint is needed.

2. Such a second constraint may impose that the flow field should vary (piecewise) smoothly in space. This can be achieved if

$$\int_{\Omega} \Psi\left(|\nabla u|^2 + |\nabla v|^2\right)\, dx\, dy \tag{2}$$

is small, where $\Psi : \mathbb{R} \to \mathbb{R}$ is an increasing differentiable function and $\nabla := (\partial_x, \partial_y)^T$ denotes the 2D nabla operator. This assumption is called *smoothness constraint*. In the sequel we shall assume that $\Psi(s^2)$ is convex in s, and that there exist constants $c_1, c_2 > 0$ with $c_1 s^2 \le \Psi(s^2) \le c_2 s^2$ for all s. In this case the optic flow problem becomes well-posed. Examples for Ψ will be presented at the end of this section.

In order to satisfy both the optic flow and the smoothness constraint as good as possible, they are assembled into a single energy functional to be minimized:

$$E(u,v) := \int_{\Omega} \Big((f_x u + f_y v + f_z)^2 + \alpha\, \Psi(|\nabla u|^2 + |\nabla v|^2)\Big)\, dx\, dy \tag{3}$$

where the *regularization parameter* $\alpha > 0$ specifies the weight of the second summand *(smoothness term, regularizer)* relative to the first one *(data term)*. Larger values for α lead to smoother flow fields.

Using steepest descend for the minimization of (3) gives the diffusion–reaction system

$$u_t = \operatorname{div}\left(\Psi'\left(|\nabla u|^2 + |\nabla v|^2\right) \nabla u\right) - \tfrac{1}{\alpha} f_x (f_x u + f_y v + f_z), \tag{4}$$

$$v_t = \operatorname{div}\left(\Psi'\left(|\nabla u|^2 + |\nabla v|^2\right) \nabla v\right) - \tfrac{1}{\alpha} f_y (f_x u + f_y v + f_z), \tag{5}$$

where Ψ' is the derivative of Ψ with respect to its argument, and div denotes the 2D divergence operator, i.e. $\operatorname{div}\binom{a}{b} := \partial_x a + \partial_y b$. The diffusion time t is an artificial evolution parameter which should not be mixed up with the time z of the image sequence $f(x,y,z)$. For $t \to \infty$, the solution (u,v) gives the minimum of $E(u,v)$. It is unique since $\Psi(s^2)$ is convex in s.

The diffusivity in both equations is given by $\Psi'(|\nabla u|^2 + |\nabla v|^2)$. It steers the activity of the smoothing process: diffusion is strong at locations where the diffusivity is large, and smoothing is reduced at places where the diffusivity is small. We shall now consider some examples which demonstrate how the choice of Ψ influences the smoothing process.

1. Horn and Schunck [15] considered the linear case $\Psi(s^2) = s^2$. This corresponds to the constant diffusivity $\Psi'(s^2) = 1$. Therefore, the smoothing activity of the Horn and Schunck method does not depend on the flow variation $s^2 = |\nabla u|^2 + |\nabla v|^2$. As a consequence, the flow is also smoothed across motion boundaries. This explains a well-known drawback of this method: a blurry flow field which is ignorant of the true motion boundaries.

2. Many modifications have been proposed to alleviate this problem. Nagel [21] for instance reduced diffusion across image boundaries with large $|\nabla f|$. Thus, this

method considers an *image-driven* smoothness term for the flow field. In many cases this modification outperforms the Horn and Schunck approach. In specific situations, however, image discontinuities may not coincide with flow discontinuities: strongly textured rigid objects, for example, have numerous texture edges which are not motion boundaries. Then an image-driven smoothness term would lead to an oversegmentation and it would be desirable to replace it by one which respects flow discontinuities instead of image discontinuities.

3. A *flow-driven* smoothness term can be constructed by using a nonlinear convex regularizer $\Psi(s^2)$ which creates a decreasing diffusivity $\Psi'(s^2)$. This ensures that the smoothing is reduced at locations where the flow magnitude is small. In the context of optic flow, such methods have been considered by Schnörr [27] and Weickert [33]. One may for instance consider the regularizer

$$\Psi(s^2) := \varepsilon s^2 + (1-\varepsilon)\lambda^2 \sqrt{1 + s^2/\lambda^2} \tag{6}$$

with $0 < \varepsilon \ll 1$ and $\lambda > 0$. It leads to the diffusivity

$$\Psi'(s^2) = \varepsilon + \frac{1-\varepsilon}{\sqrt{1 + s^2/\lambda^2}}. \tag{7}$$

We observe that λ can be regarded as a *contrast parameter*: If the flow variation $s^2 = |\nabla u|^2 + |\nabla v|^2$ is large compared to λ^2, then the diffusivity is close to 0, and for $s^2 \ll \lambda^2$ the diffusivity tends to 1. Choosing a very small value for λ relates this method to total variation regularization, a powerful denoising technique permitting discontinuous solutions [25]. The parameter ε is only required for proving well-posedness. In practical applications it can be fixed to some small value, e.g. $\varepsilon := 10^{-6}$.

4. Other flow-driven smoothness terms from the literature [4, 9, 10, 17] replace the regularizer $\Psi(|\nabla u|^2 + |\nabla v|^2)$ by $\Psi(|\nabla u|^2) + \Psi(|\nabla v|^2)$. This leads to two diffusion–reaction equations where the joint diffusivity $\Psi'(|\nabla u|^2 + |\nabla v|^2)$ is replaced by $\Psi'(|\nabla u|^2)$ and $\Psi'(|\nabla v|^2)$, respectively. Hence, the coupling between the two equations becomes weaker and flow discontinuities may be formed at different locations for u and v. It should also be mentioned that in general such models are not rotationally invariant.

3 Spatio-Temporal Smoothness Terms

The methods that we have discussed so far work locally in time: two frames are sufficient to calculate the optic flow field. In general, however, we have much more data at our disposal, namely the entire image sequence. It would thus be consequent if we use this full information for computing the optic flow field. In this way we may expect more robust results.

Using the knowledge from the previous section it is not difficult to extend the smoothness constraint into the temporal domain. Instead of calculating the optic flow (u, v) as the minimizer of the two-dimensional integral (3) for each time frame z, we now minimize a single three-dimensional integral whose solution is the optic flow for *all* frames $z \in [0, T]$:

$$E(u,v) := \int\limits_{\Omega\times[0,T]} \Big((f_x u + f_y v + f_z)^2 \, + \, \alpha \, \Psi(|\nabla_\theta u|^2 + |\nabla_\theta v|^2) \Big) \, dx \, dy \, dz \tag{8}$$

where $\nabla_\theta := (\partial_x, \partial_y, \partial_z)^T$ denotes the spatio-temporal nabla operator. The corresponding steepest descent equations are given by

$$u_t \;=\; \nabla_\theta \cdot \left(\Psi' \left(|\nabla_\theta u|^2 + |\nabla_\theta v|^2 \right) \, \nabla_\theta u \right) - \tfrac{1}{\alpha} f_x (f_x u + f_y v + f_z), \tag{9}$$

$$v_t \;=\; \nabla_\theta \cdot \left(\Psi' \left(|\nabla_\theta u|^2 + |\nabla_\theta v|^2 \right) \, \nabla_\theta v \right) - \tfrac{1}{\alpha} f_y (f_x u + f_y v + f_z). \tag{10}$$

In contrast to the two-dimensional diffusion–reaction system (4)–(5) we now have a three-dimensional problem. In the present paper we study this process for the nonlinear regularizer given in (6).

The diffusion part in (9)–(10) has the same structure as nonlinear diffusion filters for regularizing three-dimensional vector-valued images. Such methods have first been applied by Gerig *et al.* [13] in the context of medical imaging. The latter approach, however, uses diffusivities from [23] which may create ill-posed processes. This cannot happen in our case, where convex smoothness terms in the energy functional create well-posed diffusion–reaction processes. The well-posedness proof is a straightforward extension of the results in [27] to the spatio-temporal case. For a more detailed discussion of nonlinear diffusion filtering we refer to [14, 32].

4 Numerical Aspects

We approximate the 2-D diffusion–reaction system (4)–(5) and its 3-D counterpart (9)–(10) by finite differences. Derivatives in x, y and z are approximated by central differences, and for the discretization in t direction we use a slightly modified explicit (Euler forward) scheme.

Each iteration step proceeds as follows. Let τ be the step size in t direction and let f_{xi}, f_{yi} and f_{zi} denote central difference approximations of f_x, f_y and f_z in some pixel i, respectively. Let the flow components for the first iteration be initialized by 0. The $(k{+}1)$-th iteration calculates the unknown flow components u_i^{k+1} and v_i^{k+1} using known values from level k:

$$\frac{u_i^{k+1} - u_i^k}{\tau} \;=\; \sum_j a_{ij}^k \, u_j^k - \tfrac{\tau}{\alpha} f_{xi} \, (f_{xi} u_i^{k+1} + f_{yi} v_i^k + f_{zi}), \tag{11}$$

$$\frac{v_i^{k+1} - v_i^k}{\tau} \;=\; \sum_j a_{ij}^k \, v_j^k - \tfrac{\tau}{\alpha} f_{yi} \, (f_{xi} u_i^k + f_{yi} v_i^{k+1} + f_{zi}), \tag{12}$$

The matrix entries a_{ij}^k result from a standard discretization of the divergence expressions:

$$\sum_j a_{ij}^k \, u_j^k \;:=\; \sum_{j \in \mathcal{N}(i)} \frac{{\Psi'}_j^k + {\Psi'}_i^k}{2} \, (u_j^k - u_i^k), \tag{13}$$

where $\mathcal{N}(i)$ denotes the set of (4 in 2-D, 6 in 3-D) neighbours of pixel i, and Ψ'^k_i approximates the diffusivity $\Psi'(|\nabla u|^2 + |\nabla v|^2)$ in pixel i at time level k. In (11) and (12), one expression in the reaction term is approximated at time level $k+1$ in order to improve stability. Note that this scheme can still be solved explicitly for u_i^{k+1} and v_i^{k+1}:

$$u_i^{k+1} = \frac{u_i^k + \tau \sum_j a_{ij}^k \, u_j^k - \frac{\tau}{\alpha} f_{xi} \, (f_{yi} v_i^k + f_{zi})}{1 + \frac{\tau}{\alpha} f_{xi}^2}, \tag{14}$$

$$v_i^{k+1} = \frac{v_i^k + \tau \sum_j a_{ij}^k \, v_j^k - \frac{\tau}{\alpha} f_{yi} \, (f_{xi} u_i^k + f_{zi})}{1 + \frac{\tau}{\alpha} f_{yi}^2}. \tag{15}$$

We used the time step size $\tau = 1/4$ in the 2-D case and $\tau = 1/6$ in the 3-D case. The iterations were stopped when the Euclidean norm of the relative residue dropped below 0.001.

The explicit scheme (11),(12) has been chosen for simplicity reasons. In order to gain absolute stability, it is also possible to replace it by a slightly more complicated semi-implicit approximation, for instance the additive operator splitting (AOS) scheme considered in [33].

It should be noted that the 3-D scheme requires only about 50 % more computing time than a corresponding 2-D scheme that is applied to all subsequent frame pairs of an image sequence: 2-D diffusion within a 4-neighbourhood is replaced by 3-D diffusion within a 6-neighbourhood. The main difference is an increased memory requirement, since, in the 3-D case, the whole sequence is processed simultaneously. For the typical test sequences in computer vision, this does not lead to problems when modern PCs or workstations are used: on a computer with 512 MB memory one can expect to be able to process sequences with sizes up to $256 \times 256 \times 128$.

5 Experiments

In this section we illustrate the behaviour of our method by applying it to three test image sequences. We compare pure 2D processing (eqns. (4)–(5)) with 3D processing (eqns. (9)–(10)).

Figure 1 depicts one frame from a hallway scene where a person is moving towards the camera. The calculated optic flow results are shown in Figure 2. For pure spatial regularization we observe that outliers dominate, and that it is difficult to achieve good motion segmentation by thresholding the optic flow vectors. Spatio-temporal regularization, on the other hand, creates a more homogeneous motion field within the contour of the person, and motion segmentation is much more realistic.

Figure 3 shows the results of our comparison for the famous Hamburg taxi sequence. It is available via anonymous ftp from the site `ftp://csd.uwo.ca` under the directory `pub/vision`. Also for this sequence one observes that spatio-temporal processing leads to more realistic motion segmentation. Moreover, it is less sensitive to noise than a pure spatial processing. It is worth emphasizing that in this and in the previous example, the same parameters have been used for both methods, and no presmoothing of any kind has been applied. All smoothing effects are thus caused by the regularizers. Figure

4 shows that the better smoothing behaviour of spatio-temporal regularizers may also reduce temporal aliasing problems: for the pedestrian in the taxi scene, pure spatial processing gives an optic flow field that points in the wrong direction, while spatio-temporal processing creates coherent optic flow vectors pointing in the correct direction.

After these qualitative comparisons, let us now turn our attention to a quantitative validation. To this end we consider a synthetic street sequence for which ground truth flow data are available. We obtained it from

`http://www.cs.otago.ac.nz/research/vision/Downloads/`

It has been created by Galvin *et al.* [12] for evaluating eight optic flow algorithms, and it is one of a few nontrivial test sequences with ground truth data, where motion boundaries are important. We used the full sequence from the web, consisting of 20 frames of size 200 × 200 pixels. An interesting detail is depicted in Figure 5(a), and the corresponding ground truth flow field and the calculated ones are given in Figures 5(b),(c),(d), respectively. For assessing the performance of our method, we calculated the angular error

$$\Psi_e := \arccos\left(\frac{u_c u_e + v_c v_e + 1}{\sqrt{(u_c^2 + v_c^2 + 1)(u_e^2 + v_e^2 + 1)}}\right) \tag{16}$$

where (u_c, v_c) denotes the correct flow, and (u_e, v_e) is the estimated flow (cf. also [5]). In order to make our method comparable with the other approaches, we applied some presmoothing by convolving the images with a Gaussian with standard deviation σ. Preprocessing steps of this type are common for evaluating optic flow algorithms [5].

With optimized parameters for σ, α and λ we obtained an average angular error of 6.62° for the spatial approach, and 4.85° for the spatio-temporal approach. The best method that has been reported by Galvin *et al.* [12] was a thresholded version of the Lucas–Kanade algorithm [16]. It achieved an average angular error of 5°, but the density of its flow field was only 32 %, while our method creates flow fields with 100 % density. The best full density method in [12] was an algorithm by Proesmans *et al.* [24] with an average angular error of 7°. This shows that our spatio-temporal method has very good performance. Also for less optimal parameter settings, the results remained competitive.

6 Conclusions and Further Work

We have presented a nonlinear spatio–temporal regularization approach for the computation of piecewise smooth optic flow. It leads to a convex nonquadratic optimization problem which has a unique minimum that can be recovered by a globally convergent gradient descent algorithm. The model has a rotationally invariant continuous formulation, it is symmetric in time and it avoids smoothing over spatial and temporal flow discontinuities. Qualitative and quantitative comparisons showed a significant improvement over pure 2D processing at low additional computational costs.

It appears that the limited computer memory was the main reason why spatio-temporal optic flow regularizers have been used so rarely in the past. Since this is no

Figure 1: *(a) Left:* Frame 8 of a hallway sequence of size $256 \times 256 \times 16$ pixels. A person is moving towards the camera. *(b) Right:* Detail.

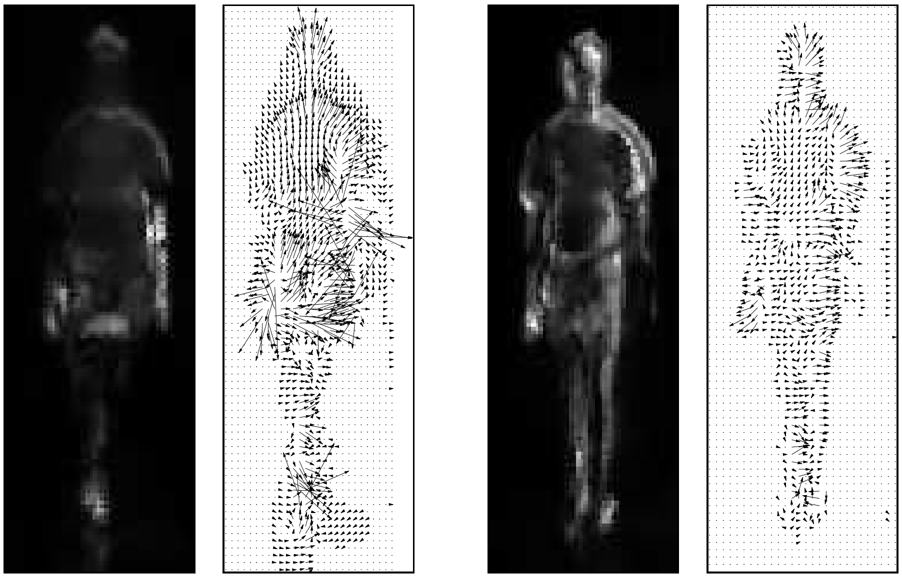

Figure 2: Computed flow fields using the hallway sequence. *From left to right: (a)* Grey value plot of the optic flow magnitude for 2-D processing (eqns. (4)–(5)). Note how outliers dominate the image such that other regions get scaled down. *(b)* Vector plot of the optic flow field for 2-D processing, subsampled by a factor 2. For better visibility, vectors w with $|w| < 0.2$ pixels have not been drawn. *(c)* Optic flow magnitude for 3-D processing (eqns. (9)–(10)). *(d)* Vector plot for 3-D processing, subsampled by a factor 2, and thresholded at 0.2 pixels. The proposed extension of adaptive smoothing to the temporal axis gives a much more coherent and complete result.

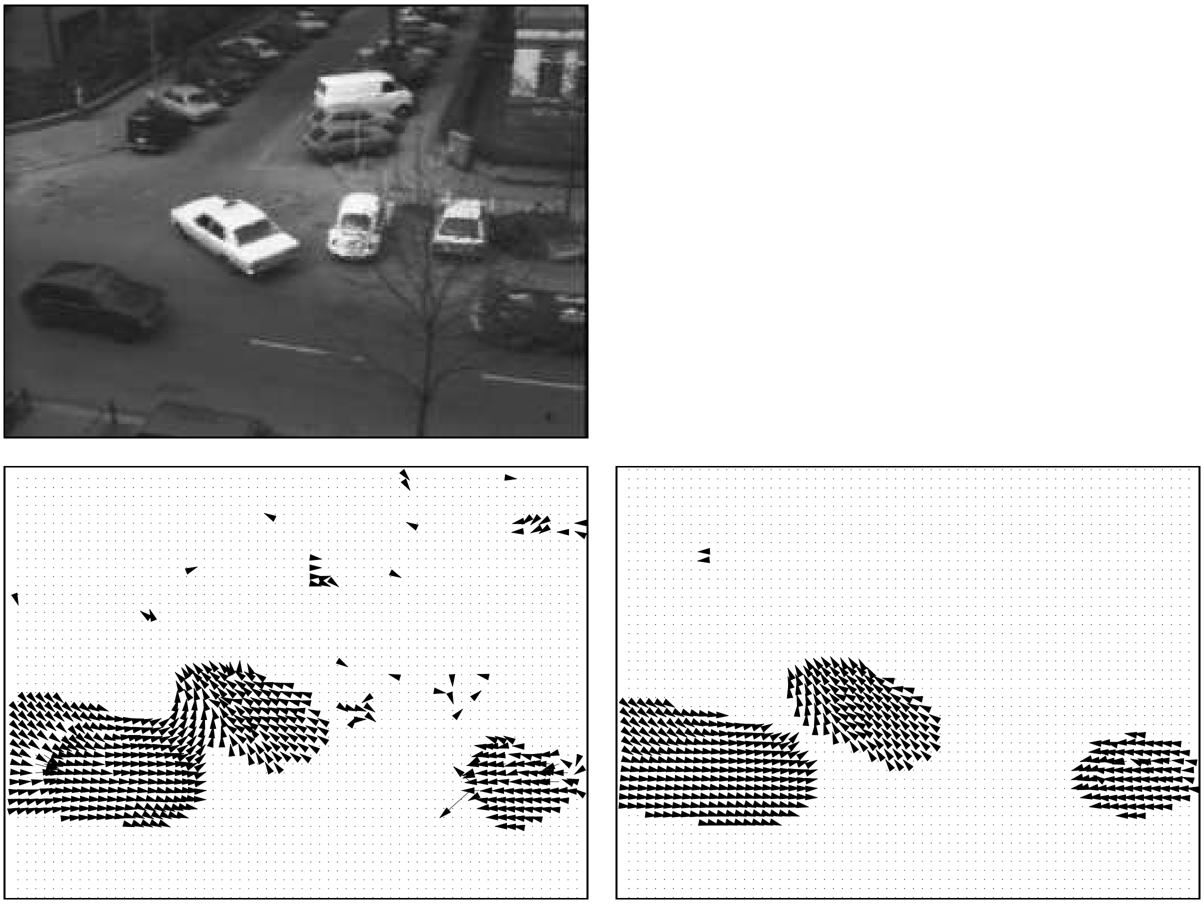

Figure 3: *(a) Top:* Frame 10 of the well–known Hamburg taxi scene (20 frames of size 256×190). *(b) Bottom left:* Optic flow field for 2-D processing, subsampled by factor 4 and thresholded at 0.2 pixels for better visibility. *(c) Bottom right:* Result for 3-D processing. Spatio-temporal regularization improves the vector fields significantly, smoothes out background noise, and preserves true motion boundaries.

longer a problem, it is likely that these methods will gain more importance in the future. It should also be noted that the spatio-temporal extension that we studied here is of course not limited to the specific nonlinear flow-driven regularizer that we used in this paper. It is a general strategy for exploiting the entire image sequence data for reliable optic flow estimation within a variational framework.

Based on these encouraging results we are currently investigating the design of highly efficient optic flow algorithms for sequential and parallel computer architectures. Some of these techniques will be based on our recent research on efficient algorithms for variational image restoration and nonlinear diffusion filtering [34].

Acknowledgements. We thank Ole Fogh Olsen and Mads Nielsen (Department of Computer Science, University of Copenhagen) for providing the hallway sequence. J.W. also thanks Jens Arnspang (Department of Computer Science, University of Copenhagen) for interesting discussions on optic flow. Part of this work has been supported

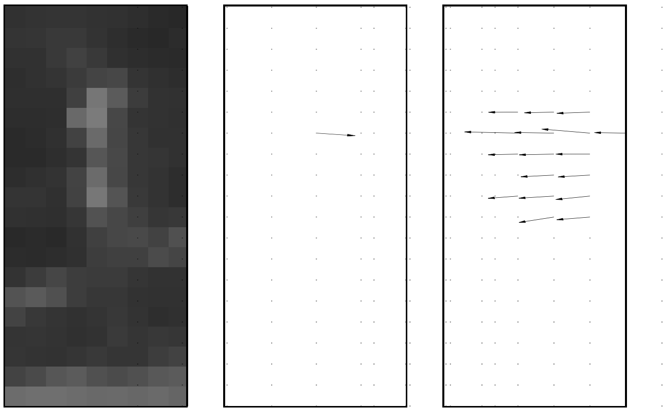

Figure 4: *(a) Left:* Magnification showing the pedestrian in the upper left part of the taxi scene (9×20 pixels). The pedestrian moves to the left. *(b) Middle:* Optic flow field for 2-D processing, subsampled by a factor 2, and thresholded at 0.2 pixels. The 2-D processing smoothes out the noisy local motion data (normal flow). Temporal aliasing, however, creates an erroneous motion to the right. *(c) Right:* Optic flow field for 3-D processing, subsampled by a factor 2. The 3D processing computes a coherent flow field for the "rigid part" of the pedestrian. The motion is in the correct direction. Regions with moving limbs are interpreted as noise at such a small spatial scale.

by the EU-TMR Project VIRGO.

References

[1] L. Alvarez, J. Esclarín, M. Lefébure and J. Sánchez, *A PDE model for computing the optical flow*, Proc. XVI Congreso de Ecuaciones Diferenciales y Aplicaciones (C.E.D.Y.A. XVI, Las Palmas de Gran Canaria, Sept. 21–24, 1999), 1349–1356, 1999.

[2] L. Alvarez, J. Weickert, J. Sánchez, *Reliable estimation of dense optical flow fields with large displacements*, Int. J. Comput. Vision, in press.

[3] J. Arnspang, *Notes on local determination of smooth optic flow and the translational property of first order optic flow*, Report DIKU 88/1, Dept. of Computer Science, University of Copenhagen, Universitetsparken 1, 2100 Copenhagen, Denmark, 1988.

[4] G. Aubert, R. Deriche, P. Kornprobst, *Computing optical flow via variational techniques*, SIAM J. Appl. Math, Vol. 60, 156–182, 1999.

[5] J.L. Barron, D.J. Fleet, S.S. Beauchemin, *Performance of optical flow techniques*, Int. J. Comput. Vision, Vol. 12, 43–77, 1994.

[6] M.J. Black, *Recursive non-linear estimation of discontinuous flow fields*, J.-O. Eklundh (Ed.), Computer vision – ECCV '94, Volume I, Lecture Notes in Computer Science, Vol. 800, Springer, Berlin, 138–145, 1994.

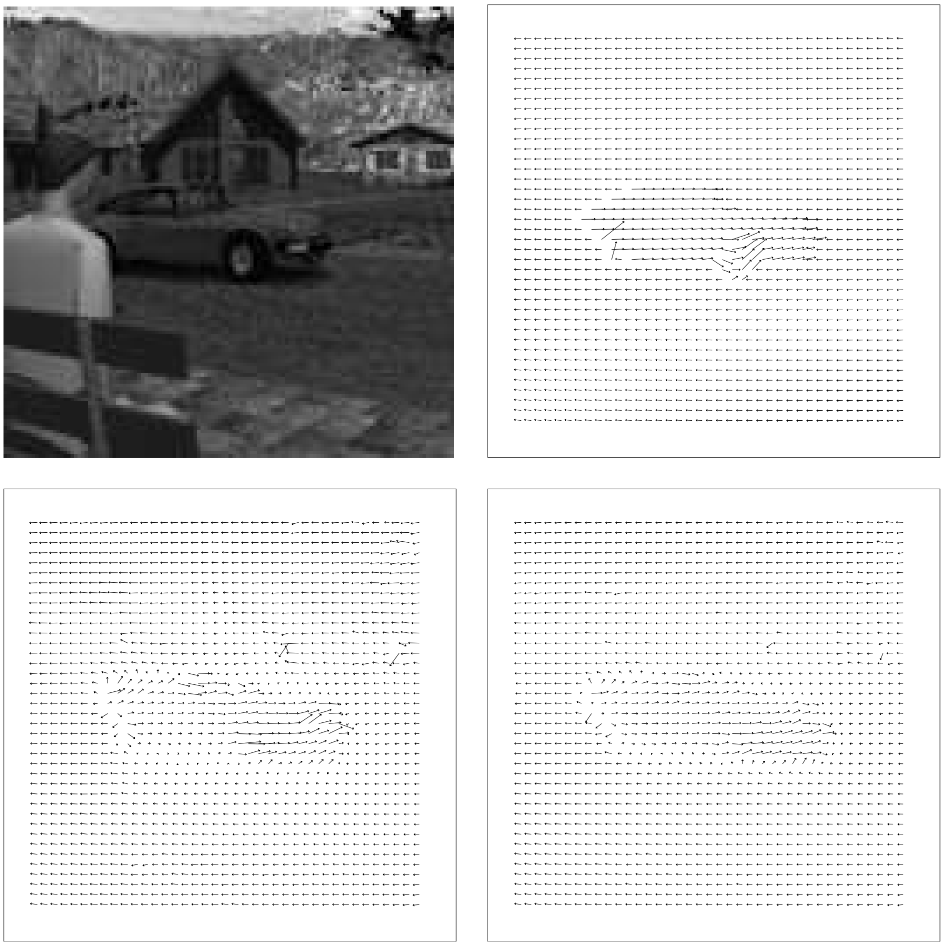

Figure 5: *(a) Top left:* Detail from a synthetic street scene (128×128 pixels). *(b) Top right:* Exact optic flow field, subsampled by a factor 3 and scaled by 2 for improving visibility. *(c) Bottom left:* Optic flow field for 2-D processing ($\sigma = 0.8$, $\alpha = 800$, $\lambda = 0.04$). *(d) Bottom right:* Result for 3-D processing ($\sigma = 0.6$, $\alpha = 250$, $\lambda = 0.05$).

[7] M.J. Black, P. Anandan, *Robust dynamic motion estimation over time*, Proc. IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition (CVPR '91, Maui, June 3–6, 1991), IEEE Computer Society Press, Los Alamitos, 292–302, 1991.

[8] A. Blake, A. Zisserman, *Visual reconstruction*, MIT Press, Cambridge (Mass.), 1987.

[9] I. Cohen, *Nonlinear variational method for optical flow computation*, Proc. Eighth Scandinavian Conf. on Image Analysis (SCIA '93, Tromsø, May 25–28, 1993), Vol. 1, 523–530, 1993.

[10] R. Deriche, P. Kornprobst, G. Aubert, *Optical-flow estimation while preserving its discontinuities: A variational approach*, Proc. Second Asian Conf. Computer Vision (ACCV '95, Singapore, December 5–8, 1995), Vol. 2, 290–295, 1995.

[11] D.J. Fleet, A.D. Jepson, *Computation of component image velocity from local phase information*, Int. J. Comput. Vision, Vol. 5, 77–104, 1990

[12] B. Galvin, B. McCane, K. Novins, D. Mason, S. Mills, *Recovering motion fields: An analysis of eight optical flow algorithms*, Proc. 1998 British Machine Vision Conference (BMVC '98, Southampton, September 14–17, 1998).

[13] G. Gerig, O. Kübler, R. Kikinis, F.A. Jolesz, *Nonlinear anisotropic filtering of MRI data*, IEEE Trans. Medical Imaging, Vol. 11, 221–232, 1992.

[14] B.M. ter Haar Romeny (Ed.), *Geometry-driven diffusion in computer vision*, Kluwer, Dordrecht, 1994.

[15] B. Horn, B. Schunck, *Determining optical flow*, Artif. Intell., Vol. 17, 185–203, 1981.

[16] B. Lucas, T. Kanade, *An iterative image registration technique with an application to stereo vision*, Proc. Seventh Int. Joint Conf. on Artificial Intelligence (IJCAI '81, Vancouver, August 1981), 674–679, 1981.

[17] A. Kumar, A.R. Tannenbaum, G.J. Balas, *Optic flow: a curve evolution approach*, IEEE Trans. Image Proc., Vol. 5, 598–610, 1996.

[18] P.J.M. van Laarhoven, E.H.L. Aarts, *Simulated annealing: theory and applications*, Reidel, Dordrecht, 1988.

[19] A. Mitiche, P. Bouthemy, *Computation and analysis of image motion: a synopsis of current problems and methods*, Int. J. Comput. Vision, Vol. 19, 29–55, 1996.

[20] D.W. Murray, B.F. Buxton, *Scene segmentation from visual motion using global optimization*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 9, 220–228, 1987.

[21] H.H. Nagel, *Constraints for the estimation of displacement vector fields from image sequences*, Proc. Eighth Int. Joint Conf. on Artificial Intelligence (IJCAI '83, Karlsruhe, August 8–12, 1983), 945–951, 1983.

[22] H.H. Nagel, *Extending the 'oriented smoothness constraint' into the temporal domain and the estimation of derivatives of optical flow*, O. Faugeras (Ed.), Computer vision – ECCV '90, Lecture Notes in Computer Science, Vol. 427, Springer, Berlin, 139–148, 1990.

[23] P. Perona, J. Malik, *Scale space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 12, 629–639, 1990.

[24] M. Proesmans, L. Van Gool, E. Pauwels, A. Oosterlinck, *Determination of optical flow and its discontinuities using non-linear diffusion*, J.-O. Eklundh (Ed.), Computer vision – ECCV '94, Volume II, Lecture Notes in Computer Science, Vol. 801, Springer, Berlin, 295–304, 1994.

[25] L.I. Rudin, S. Osher, E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D, Vol. 60, 259–268, 1992.

[26] C. Schnörr, *Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class*, Int. J. Comput. Vision, Vol. 6, 25–38, 1991.

[27] C. Schnörr, *Segmentation of visual motion by minimizing convex non-quadratic functionals*, Proc. 12th Int. Conf. Pattern Recognition (ICPR 12, Jerusalem, Oct. 9–13, 1994), Vol. A, IEEE Computer Society Press, Los Alamitos, 661–663, 1994.

[28] M.A. Snyder, *On the mathematical foundations of smoothness constraints for the determination of optical flow and for surface reconstruction*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 13, 1105–1114, 1991.

[29] C. Stiller, J. Konrad, *Estimating motion in image sequences*, IEEE Signal Proc. Magazine, Vol. 16, 70–91, 1999.

[30] S. Uras, F. Girosi, A. Verri, V. Torre, *A computational approach to motion perception*, Biol. Cybern., Vol. 60, 79–87, 1988.

[31] J. Weber, J. Malik, *Robust computation of optical flow in a multi-scale differential framework*, Int. J. Comput. Vision, Vol. 14, 67–81, 1995.

[32] J. Weickert, *Anisotropic diffusion in image processing*, Teubner-Verlag, Stuttgart, 1998.

[33] J. Weickert, *On discontinuity-preserving optic flow*, S. Orphanoudakis, P. Trahanias, J. Crowley, N. Katevas (Eds.), Proc. Computer Vision and Mobile Robotics Workshop (CVMR '98, Santorini, Sept. 17–18, 1998), 115–122, 1998.

[34] J. Weickert, J. Heers, C. Schnörr, K.J. Zuiderveld, O. Scherzer, H.S. Stiehl, *Fast parallel algorithms for a broad class of nonlinear variational diffusion approaches*, Real-Time Imaging, in press.

[35] J. Weickert, C. Schnörr, *Räumlich-zeitliche Berechnung des optischen Flusses mit nichtlinearen flußabhängigen Glattheitstermen*, W. Förstner, J.M. Buhmann, A. Faber, P. Faber (Eds.), Mustererkennung 1999, Springer, Berlin, 317–324, 1999.

[36] J. Weickert, C. Schnörr, *A theoretical framework for convex regularizers in PDE-based computation of image motion*, Report 13/2000, Computer Science Series, Dept. of Mathematics and Computer Science, University of Mannheim, 68131 Mannheim, Germany, 2000.