Variational Image Motion Computation: Theoretical Framework, Problems and Perspectives

Christoph Schnörr and Joachim Weickert

Computer Vision, Graphics, and Pattern Recognition Group
Dept. of Math. and Comp. Science, University of Mannheim
68131 Mannheim, Germany
{Christoph.Schnoerr,Joachim.Weickert}@ti.uni-mannheim.de

Abstract. Many differential methods for the recovery of the optic flow field from an image sequence can be expressed in terms of a variational problem where the optic flow minimizes some energy. Typically, these energy functionals consist of two terms: a data term, which requires e.g. that a brightness constancy assumption holds, and a regularizer that encourages global or piecewise smoothness of the flow field. In this paper we present a systematic classification of rotation invariant convex regularizers by exploring their connection to diffusion filters for multichannel images. This taxonomy provides a unifying framework for data-driven and flow-driven, isotropic and anisotropic, as well as spatial and spatio-temporal regularizers. While some of these techniques are classic methods from the literature, others are derived here for the first time. We conclude by pointing out some important research issues related to variational optic flow computation.

Keywords: optic flow, differential methods, variational regularization, diffusion filtering

1 Introduction

Motion computation in image sequences is key problem in computer vision. An important concept in this context is optic flow, the displacement field of corresponding pixels in subsequent frames of an image sequence. It provides not only a basis for motion-based image sequence interpretation but also for second generation video algorithms and numerous applications in, for instance, experimental fluid dynamics, meteorology, and medical imaging. Surveys of the state-of-the-art in motion computation can be found in papers by Mitiche and Bouthemy [24], and Stiller and Konrad [39]. For a performance evaluation of some of the most popular algorithms we refer to Barron *et al.* [4].

An important class for image motion computation constitute variational approaches, as pioneered by Horn and Schunck [18]. These approaches use a global smoothness term in order to compute an optic flow field from sparse local motion information. The use of a global smoothness constraint has often been criticized in the literature, in particular by researchers investigating local approaches to optic flow. These authors advocate to compute optic flow only at locations where the aperture problem can be solved by local window averaging, and the resulting sparse local motion information should be interpolated somehow at a later processing stage. It should be noted that variational approaches just make this implicit way of smoothing more explicit and simultaneously solve the interpolation problem.

The smoothness term usually requires that the optic flow field should vary smoothly in space [18]. Such a term may be modified in an *image-driven* way in order to suppress smoothing at or across image boundaries [1, 26]. As an alternative, *flow-driven* modifications have been proposed which reduce smoothing across flow discontinuities [5, 8, 9, 20, 31, 34, 42]. Most smoothness terms require only *spatial* smoothness. *Spatio-temporal* smoothness terms have been considered to a much smaller extent [25, 27, 45].

Minimizing continuous energy functionals leads in a natural way to partial differential equations (PDEs): applying gradient descent, for instance, yields a system of coupled diffusion–reactions equations for the two flow components. The fastly emerging use of PDE-based image restoration methods [14, 30], such as nonlinear diffusion filtering and total variation denoising, has motivated many researchers to apply similar ideas to estimate optic flow [1, 3, 8, 9, 17, 20, 29, 31, 34, 42]. A systematic framework that links the diffusion and optic flow paradigms, however, has not been studied so far.

Furthermore, from the framework of diffusion filtering it is also well-known that anisotropic filters with a diffusion tensor have more degrees of freedom than isotropic ones with scalar-valued diffusivities. These additional degrees of freedom can be used to obtain better results in specific situations [41]. However, similar nonlinear anisotropic regularizers have not been considered in the optic flow literature so far.

In the present paper, we present a theoretical framework for a broad class of regularization methods for optic flow estimation. For the reasons explained below (Section 2), we focus on models that allow a formulation in terms of convex and rotation invariant continuous energy functionals. We consider image-driven and flow-driven models, isotropic and anisotropic ones, as well as models with spatial and spatio-temporal smoothing terms (Section 3). Our taxonomy includes not only many existing models, but also interesting novel ones. For more details and proofs of our results we refer to [47]. Finally, we briefly discuss several open research problems in Section 5.

2 Variational Approach

2.1 Basic structure

Let us consider a real-valued image sequence $f(x, y, \theta)$, where (x, y) denotes the location within the image domain $\Omega \in \mathbb{R}^2$, and the time parameter $\theta \in [0, T]$ specifies the frame. The optic flow field $(u_1(x, y, \theta), u_2(x, y, \theta))$ describes the displacement between two subsequent frames θ and $\theta + 1$. If the spatial and temporal sampling is sufficiently fine, the common assumption of conserved grey values yields:

$$f_x u_1 + f_y u_2 + f_\theta = 0, \tag{1}$$

where the subscripts x, y and θ denote partial derivatives. This so-called *optic flow constraint (OFC)* forms the basis of many differential methods for estimating the optic flow. Numerous generalizations to multiple constraint equations and/or different "conserved quantities" (replacing intensity) exist; see e.g. [11, 40]. In the present paper, however, we confine ourselves to (1).

The basic idea underlying variational approaches is to recover the optic flow as a minimizer of some energy functional of type

$$E(u_1, u_2) := \int_\Omega \Big((f_x u_1 + f_y u_2 + f_\theta)^2 \; + \; \alpha \, V(\nabla f, \nabla u_1, \nabla u_2) \Big) \, dx \, dy \tag{2}$$

where $\nabla := (\partial_x, \partial_y)^T$ denotes the spatial nabla operator, and $u := (u_1, u_2)^T$. The first term in the energy functional is a *data term* requiring that the OFC be fulfilled, while the second term penalizes deviations from (piecewise) smoothness. The *smoothness term* $V(\nabla f, \nabla u_1, \nabla u_2)$ is called *regularizer*, and the positive smoothness weight α is the *regularization parameter*.

It is a classic result from the calculus of variations that – under mild regularity conditions – a minimizer (u_1, u_2) of some energy functional

$$E(u_1, u_2) := \int_\Omega G(x, y, u_1, u_2, \nabla u_1, \nabla u_2) \, dx \, dy \tag{3}$$

necessarily satisfies the so-called *Euler–Lagrange equations*

$$\partial_x G_{u_{ix}} + \partial_y G_{u_{iy}} - G_{u_i} = 0, \quad \partial_n u_i = 0 \;\text{ on }\; \partial\Omega, \quad i = 1, 2 \,,$$

which can be regarded as the steady state ($t \to \infty$) of the diffusion–reaction system

$$\partial_t u_i \;=\; \partial_x G_{u_{ix}} + \partial_y G_{u_{iy}} - G_{u_i} \,, \quad i = 1, 2 \,. \tag{4}$$

A classification of regularizers $V(\nabla f, \nabla u_1, \nabla u_2)$ and the underlying diffusion processes according to eqn. (4) will be presented in Section 3.

2.2 Continuous problem formulation, rotational invariance, and convexity

Modeling the optic flow recovery problem in terms of *continuous* energy functionals such as (2) offers the advantage of having a formulation that is as independent of the pixel grid as possible. A correct continuous model can be *rotation invariant*, and the use of well-established numerical methods shows how this rotation invariance can be approximated in a mathematically consistent way.

From both a theoretical and practical point of view, it can be attractive to use energy functionals that are *convex*. They have a unique minimum, and this global minimum can be found in a stable way by using standard techniques from convex optimization, for instance gradient descent methods. Having a unique minimum allows to use globally convergent algorithms, where 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. *Nonconvex* energy functionals, on the other hand, may have many local minima, and it is difficult to find algorithms that are both efficient and converge to a global minimum. Typical algorithms which converge to a global minimum (such as simulated annealing [21]) are computationally very expensive, while methods which are more efficient (such as graduated non-convexity algorithms [6]) may get trapped in a local minimum.

A discussion of approaches lacking these properties (e.g., [3, 5, 7–9, 16, 20, 23, 25, 31]) is not within the scope of the present paper.

3 Convex Regularizers: Theoretical Framework

In this section, we present a systematic classification of regularizers for optic flow computation by exploring their connection to multichannel diffusion filtering. Our taxonomy includes not only interesting re-formulations of existing models but also novel anisotropic flow–driven regularizers which have not been considered in the literature so far.

3.1 Classification

Homogeneous regularization. In 1981 Horn and Schunck [18] pioneered the field of regularization methods for optic flow computations. They used the regularizer

$$V_H(\nabla f, \nabla u_1, \nabla u_2) \;:=\; |\nabla u_1|^2 + |\nabla u_2|^2. \tag{5}$$

Schnörr [32] has established well-posedness by showing that this functional has a unique minimizer that depends continuously on the input data f. Recently, Hinterberger [17] proved similar well-posedness results for a related model with a different data term.

According to eqn. (4), we observe that the underlying diffusion process in the Horn and Schunck approach is the linear diffusion equation

$$\partial_t u_i = \Delta u_i = \operatorname{div}\,(g\,\nabla u_i) \tag{6}$$

with $g := 1$ and $i = 1, 2$. This equation is well-known for its regularizing properties and has been extensively used in the context of Gaussian scale-space; see [38] and the references therein. It smoothes, however, in a completely homogeneous way, since its diffusivity g equals 1 everywhere. As a consequence, it also blurs across semantically important flow discontinuities. This is the reason why the Horn and Schunck approach creates rather blurry optic flow fields. The regularizers described in the sequel are attempts to overcome this limitation.

Isotropic image–driven regularization. It seems plausible that motion boundaries are a subset of the image boundaries. Thus, a simple way to prevent smoothing at motion boundaries consists of introducing a weight function into the Horn and Schunck regularizer that becomes small at image edges. This modification yields

$$V_{II}(\nabla f, \nabla u_1, \nabla u_2) \ := \ g(|\nabla f|^2) \ \left(|\nabla u_1|^2 + |\nabla u_2|^2\right), \tag{7}$$

where g is a decreasing, strictly positive function. This regularizer has been proposed and theoretically analyzed by Alvarez *et al.* [1].

The underlying diffusion process is

$$\partial_t u_i = \text{div } \left(g(|\nabla f|^2)\, \nabla u_i\right) \qquad i = 1, 2\ . \tag{8}$$

It uses a scalar-valued diffusivity g that depends on the image gradient. Since the diffusivity does not depend on the flow itself, the diffusion process is linear.

Homogeneous regularization arises as a special case of (7) when $g(|\nabla f|^2) := 1$ is considered.

Anisotropic image–driven regularization. An early anisotropic modification of the Horn and Schunck functional is due to Nagel [26]; see also [2, 28, 32, 33, 37]. The basic idea is to reduce smoothing across image boundaries, while encouraging smoothing along image boundaries. This is achieved by considering the regularizer

$$V_{AI}(\nabla f, \nabla u_1, \nabla u_2) \ := \ \nabla u_1^T D(\nabla f) \nabla u_1 + \nabla u_2^T D(\nabla f) \nabla u_2. \tag{9}$$

$D(\nabla f)$ is a regularized projection matrix perpendicular to ∇f:

$$D(\nabla f) := \frac{1}{|\nabla f|^2 + 2\lambda^2} \ \left(\nabla f^{\perp} \nabla f^{\perp T} + \lambda^2 I\right), \tag{10}$$

where I denotes the unit matrix.

The underlying diffusion process is

$$\partial_t u_i = \text{div } \left(D(\nabla f)\, \nabla u_i\right) \qquad i = 1, 2 \tag{11}$$

The usage of a diffusion tensor $D(\nabla f)$ instead of a scalar-valued diffusivity allows a direction-dependent smoothing behaviour. This behaviour is very similar to edge-enhancing anisotropic diffusion filtering [41]. In contrast to edge-enhancing anisotropic diffusion, however, Nagel's optic flow technique is linear because the diffusion tensor depends on the image f but not on the unknown flow. Well-posedness for this model has been established in [32].

Homogeneous and isotropic image-driven regularizers are special cases of (9), where $D(\nabla f) := I$ and $D(\nabla f) := g(|\nabla f|^2)I$ are chosen.

Isotropic flow–driven regularization. Image-driven regularization methods may create oversegmentations for strongly textured objects: in this case we have much more image boundaries than motion boundaries. In order to reduce smoothing only at motion boundaries, one may consider using a purely flow-driven regularizer. This, however, is at the expense of refraining from quadratic optimization problems. In earlier work [34, 42], the authors considered regularizers of type

$$V_{IF}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \Psi\left(|\nabla u_1|^2 + |\nabla u_2|^2\right), \tag{12}$$

where $\Psi(s^2)$ is a differentiable and increasing function that is convex in s.

The underlying diffusion process is

$$\partial_t u_i = \operatorname{div}\left(\Psi'(|\nabla u_1|^2 + |\nabla u_2|^2)\, \nabla u_i\right) \qquad i = 1, 2\;, \tag{13}$$

where Ψ' denotes the derivative of Ψ with respect to its argument. In general, the diffusion process is nonlinear now. Since the nonlinear diffusivity is decreasing in its argument, smoothing at flow discontinuities is inhibited.

For the specific choice $\Psi(s^2) := s^2$, homogeneous regularization with diffusivity $\Psi'(s^2) = 1$ is recovered again.

Anisotropic flow–driven regularization. To complete our taxonomy, we consider anisotropic flow-driven regularizers next. To our knowledge, this has not been done in the literatur so far. The motivation comes from nonlinear diffusion filtering, where it was shown that anisotropic models with a diffusion tensor instead of a scalar-valued diffusivity offer advantages for images with noisy edges or interrupted structures [43].

We propose the following novel class of regularizers:

$$V_{AF}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \operatorname{tr}\Psi\left(\nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T\right). \tag{14}$$

For this definition to make sense, we have to extend the scalar-valued function $\Psi(s)$ to a matrix-valued function $\Psi(J)$. Let J denote some symmetric $n \times n$ matrix with orthonormal eigenvectors v_1,...,v_n and corresponding eigenvalues μ_1,...,μ_n. Then we define:

$$\Psi(J) \;:=\; \sum_i \Psi(\mu_i)\, v_i v_i^T. \tag{15}$$

The argument in (14),

$$J := \nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T\;, \tag{16}$$

is a symmetric and positive semidefinite 2×2 matrix. Hence, there exist two orthonormal eigenvectors v_1, v_2 with corresponding nonnegative eigenvalues μ_1, μ_2. These eigenvalues specify the contrast of the vector-valued image (u_1, u_2) in the directions v_1 and v_2, respectively. This concept has been introduced by Di Zenzo for edge analysis of multichannel images [10]. It can be regarded as a generalization of the structure tensor [12], and it is related to the first fundamental form in differential geometry [19].

The diffusion process underlying (14) can be shown to be [47]

$$\partial_t u_i = \operatorname{div}\left(\Psi'\left(\nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T\right)\, \nabla u_i\right) \qquad i = 1, 2\,, \tag{17}$$

where Ψ' is a matrix-valued function (cf. (15)).

There is an interesting connection between anisotropic and isotropic flow-driven regularization [47]. Isotropic flow-driven regularization (12) is recovered by just exchanging the trace operator and the penalty function Ψ:

$$V_{IF}(\nabla f, \nabla u_1, \nabla u_2) \;=\; \Psi\left(\operatorname{tr}\left(\nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T\right)\right)\;. \tag{18}$$

3.2 Unifying framework

Table 1 gives an overview of the smoothness investigated so far.

Table 1. Classification of regularizers for optic flow models.

	isotropic	anisotropic
image-driven	$g(\lvert\nabla f\rvert^2)\sum_{i=1}^{2}\lvert\nabla u_i\rvert^2$	$\sum_{i=1}^{2}\nabla u_i^T D(\nabla f)\nabla u_i$
flow-driven	$\Psi\left(\sum_{i=1}^{2}\lvert\nabla u_i\rvert^2\right)$	$\operatorname{tr}\,\Psi\left(\sum_{i=1}^{2}\nabla u_i \nabla u_i^T\right)$

This result suggests the following general regularizer comprising all these models [47]:

$$V(\nabla f, \nabla u) := (1-\beta)\,\Psi\left(\operatorname{tr}\,\nabla u^T D(\nabla f)\nabla u\right) \;+\; \beta\operatorname{tr}\Psi\left(\nabla u\, D(\nabla f)\nabla u^T\right) \tag{19}$$

Here, we used the compact notation $\nabla u := (\nabla u_1, \nabla u_2)$, and parameter $\beta \in [0,1]$ determines the anisotropy.

4 Extensions

4.1 Spatio-temporal regularizers

Using our results from the previous section it is straightforward to extend the smoothness constraint into the temporal domain. Instead of calculating the optic flow (u_1, u_2) as the minimizer of the two-dimensional integral (2) for each time frame θ, we now minimize a single three-dimensional integral whose solution is the optic flow for *all* frames $\theta \in [0, T]$:

$$E(u) \;:=\; \int\limits_{\Omega\times[0,T]} \left((f_x u_1 + f_y u_2 + f_\theta)^2 \;+\; \alpha\, V(\nabla_\theta f, \nabla_\theta u)\right) dx\, dy\, d\theta \tag{20}$$

where $\nabla_\theta := (\partial_x, \partial_y, \partial_\theta)^T$ denotes the spatio-temporal nabla operator.

The corresponding diffusion–reaction systems of spatio–temporal energy functionals have the same structure as the pure spatial ones that we investigated so far. The only difference is that the spatial nabla operator ∇ has to be replaced by its spatio-temporal analogue ∇_θ. Thus, one has to solve 3D diffusion–reaction systems instead of 2D ones.

Our model includes as special cases the anisotropic image-driven smoothness constraint of Nagel [27], and the isotropic flow-driven spatio-temporal regularizers studied by the authors in [45].

4.2 A design principle for anisotropic regularizers

There is an interesting relation between anisotropic flow-driven regularizers and isotropic flow-driven ones: the anisotropic regularizer $\operatorname{tr}\Psi(J)$ can be expressed by means of the eigenvalues μ_1, μ_2 of J as

$$V_{AF}(\nabla f, \nabla u_1, \nabla u_2) \;=\; \Psi(\mu_1) + \Psi(\mu_2), \tag{21}$$

while its isotropic counterpart $\Psi(\operatorname{tr} J)$ can be written as

$$V_{IF}(\nabla f, \nabla u_1, \nabla u_2) \;=\; \Psi(\mu_1 + \mu_2). \tag{22}$$

This observation motivates us to formulate the following design principle for rotationally invariant anisotropic flow-driven regularizers:

Design Principle (Rotationally Invariant Anisotropic Regularizers)
Assume that we are given some isotropic regularizer $\Psi(\sum_i |\nabla u_i|^2)$ with a non-quadratic function Ψ, and a decomposition of its argument

$$\sum_i |\nabla u_i|^2 = \sum_j \rho_j, \tag{23}$$

where the ρ_j are rotationally invariant expressions. Then the regularizer $\sum_j \Psi(\rho_j)$ is rotationally invariant and anisotropic.

An early example is given by the regularizer proposed by Schnörr [34, 35]:

$$V_{AFS}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \Psi\left(\operatorname{div}^2 u\right) + \Psi\left(\operatorname{rot}^2 u\right) + \Psi\left(\operatorname{sh}^2 u\right) \tag{24}$$

with $u := (u_1, u_2)^T$, $\operatorname{rot} u := u_{2x} - u_{1y}$, and $\operatorname{sh} u := \sqrt{(u_{2y} - u_{1x})^2 + (u_{1y} + u_{2x})^2}$. Applying the design principle, one can derive this expression from the identity

$$|\nabla u_1|^2 + |\nabla u_2|^2 = \tfrac{1}{2}\left(\operatorname{div}^2 u + \operatorname{rot}^2 u + \operatorname{sh}^2 u\right). \tag{25}$$

Using the regularizer (24) in the functional (2) leads to a highly anisotropic diffusion–reaction system (see [47]) which has no counterpart in the literature on multichannel diffusion filtering.

Based on the design principle proposed above, we are currently investigating the existence of other useful anisotropic regularizers.

5 Open problems

In this section, we describe some important research issues related to variational optical flow computation.

Efficient algorithms. Although the numerical solution of variational approaches poses no problems on current PCs, considerable research efforts are still necessary in order to achieve computation times close to the frame rate.

Two lines of research are attractive in this context (cf. [46]). The first one concerns the application of domain decomposition and multi-level approaches to the variational models described above. In this context, the notion "multi-level" refers to a research area of numerical mathematics (cf., e.g., [15]) and should not be confused with coarse-to-fine techniques used to cope with the large-motion problem (see [36] and the references therein, for example). The second line of research concerns the application of efficient numerical schemes developed for diffusion filtering [44] to the diffusion equations derived from variational approaches (see Section 2).

Another important question is how to modify these numerical algorithms so as to obtain optimal performance on a cluster of low-cost PCs running in parallel.

Computation of input data, amount of smoothing. In order to solve the variational problem (2), the input data f_x, f_y, f_θ have to be computed in a preprocessing step. Usually this involves some low-pass filter like the Gaussian, for example. On the other hand, the global smoothness term in (2) also implicitly takes into account spurious deviations of the input data. This becomes most noticeable in the case of spatio-temporal regularization [45].

Apparently, the relationship and dependency between these two types of regularization has not been thoroughly investigated so far.

Beyond convexity. In Sections 2.2 and 3, we emphasized the twofold relevance of convexity in for variational optimization. Firstly, it defines a framework for modeling various regularizers for optical flow computation and, secondly, it provides a mathematically clear basis for robust and efficient algorithms.

Concerning the design of computer vision systems, one surely has to abandon the former aspect as one moves towards higher level processing stages. At these stages, for example, one wishes not only to achieve piecewise smoothing of image data but a partitioning of the image section into (hopefully) meaningful subsets. Thus, some discrete-valued decision variables have to be introduced leading to difficult combinatorial optimization problems.

In this context, we are well aware of many *non-convex* approaches for variational optical flow computation (see the citations in Section 2.2), many of which aim at coping with these more involved problems. It should be noted, however, that these approaches give up *both* aspects of convexity mentioned above.

At present, it's not clear to us whether this necessarily has to hold true. Taking a look at the field of combinatorics [22], for example, shows that powerful problem relaxations exist making even the hardest combinatorial problems manageable. A key observation thereby is that these problems can often be approximated by convex optimization problems and efficiently be solved by corresponding polynomial-time algorithms. Furthermore, results with respect to classical graph problems [13] show that the approximation quality obtained in this way can be remarkably good.

In our future work, we will try to pursue this line of research in the context of early computational vision.

6 Conclusion

The goal of this paper was to derive a diffusion theory for optic flow functionals. Our corresponding classification of regularizers helped us not only to classify existing methods within a unifying framework, but also to identify gaps, where no models are available in the current literature.

One important novelty along these lines was the derivation of regularizers that can be related to anisotropic diffusion filters with a matrix-valued diffusion tensor. This also enabled us to propose a general design principle for anisotropic regularizers.

We hope that our systematic taxonomy provides a unifying platform for algorithms for the entire class of convex variational optic flow methods. Our future plans are to use such a platform for a detailed performance evaluation of the different methods in this paper. Furthermore, we will focus on the open points described in Section 5.

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*, To appear in International Journal of Computer Vision.
3. G. Aubert, R. Deriche, P. Kornprobst, *Computing optical flow via variational techniques*, SIAM J. Appl. Math, Vol. 60, 156–182, 1999.
4. J.L. Barron, D.J. Fleet, S.S. Beauchemin, *Performance of optical flow techniques*, Int. J. Comput. Vision, Vol. 12, 43–77, 1994.
5. M.J. Black, P. Anandan, *The robust estimation of multiple motions: Parametric and piecewise smooth flow fields*, Computer Vision and Image Understanding, Vol. 63, 75–104, 1996.
6. A. Blake, A. Zisserman, *Visual reconstruction*, MIT Press, Cambridge (Mass.), 1987.
7. L. Blanc–Féraud, M. Barlaud, T. Gaidon, *Motion estimation involving discontinuities in a multiresolution scheme*, Optical Engineering, Vol. 32, No. 7, 1475–1482, 1993.
8. 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.
9. 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.

10. S. Di Zenzo, *A note on the gradient of a multi-image*, Computer Vision, Graphics, and Image Processing, Vol. 33, 116–125, 1986.
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. W. Förstner, E. Gülch, *A fast operator for detection and precise location of distinct points, corners and centres of circular features*, Proc. ISPRS Intercommission Conf. on Fast Processing of Photogrammetric Data (Interlaken, June 2–4, 1987), 281–305, 1987.
13. M.X. Goemans, D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, Vol. 42:1115–1145, 1995.
14. B. ter Haar Romeny, L. Florack, J. Koenderink, M. Viergever (Eds.), *Scale-space theory in computer vision*, Lecture Notes in Computer Science, Vol. 1252, Springer, Berlin, 1997.
15. W. Hackbusch, editor. *Multigrid Methods V: Proc. Fifth European Multigrid Conference*, Stuttgart, Oct. 1-4 1998. Springer.
16. F. Heitz, P. Bouthemy, *Multimodal estimation of discontinuous optical flow using Markov random fields*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 15, 1217–1232, 1993.
17. W. Hinterberger, *Generierung eines Films zwischen zwei Bildern mit Hilfe des optischen Flusses*, M.Sc. thesis, Industrial Mathematics Institute, University of Linz, Austria, 1999.
18. B. Horn, B. Schunck, *Determining optical flow*, Artificial Intelligence, Vol. 17, 185–203, 1981.
19. E. Kreyszig, *Differential geometry*, University of Toronto Press, Toronto, 1959.
20. A. Kumar, A.R. Tannenbaum, G.J. Balas, *Optic flow: a curve evolution approach*, IEEE Trans. Image Proc., Vol. 5, 598–610, 1996.
21. P.J.M. van Laarhoven, E.H.L. Aarts, *Simulated annealing: theory and applications*, Reidel, Dordrecht, 1988.
22. L. Lovász, A. Schrijver. Cones of matrices and set–functions and 0–1 optimization. *SIAM J. Optimization*, Vol. 1, No. 2:166–190, 1991.
23. E. Mémin, P. Pérez, *Dense estimation and object-based segmentation of the optical flow with robust techniques*, IEEE Trans. Image Proc, Vol. 7, 703–719, 1998.
24. 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.
25. 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.
26. 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.
27. 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.
28. H.-H. Nagel, W. Enkelmann, *An investigation of smoothness constraints for the estimation of displacement vector fields from images sequences*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 8, 565–593, 1986.
29. P. Nesi, *Variational approach to optical flow estimation managing discontinuities*, Image and Vision Computing, Vol. 11, 419–439, 1993.

30. M. Nielsen, P. Johansen, O.F. Olsen, J. Weickert (Eds.), *Scale-space theories in computer vision*, Lecture Notes in Computer Science, Springer, Berlin, Vol. 1682, 1999.
31. 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, Lecture Notes in Computer Science, Vol. 801, Springer, Berlin, 295–304, 1994.
32. 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.
33. C. Schnörr, *On functionals with greyvalue-controlled smoothness terms for determining optical flow*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 15, 1074–1079, 1993.
34. 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.
35. C. Schnörr, R. Sprengel, B. Neumann, *A variational approach to the design of early vision algorithms*, Computing, Suppl. 11, 149–165, 1996.
36. E.P. Simoncelli. Bayesian multiscale differential optical flow. In B. Jähne, H. Haußecker, and P. Geißler, editors, *Handbook of Computer Vision and Applications*, volume 2, chapter 14, pages 297–422. Academic Press, San Diego, 1999.
37. 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.
38. J. Sporring, M. Nielsen, L. Florack, P. Johansen (Eds.), *Gaussian scale-space theory*, Kluwer, Dordrecht, 1997.
39. C. Stiller, J. Konrad, *Estimating motion in image sequences*, IEEE Signal Proc. Magazine, Vol. 16, 70–91, 1999.
40. J. Weber, J. Malik, *Robust computation of optical flow in a multi-scale differential framework*, Int. J. Comput. Vision, Vol. 14, 67–81, 1995.
41. J. Weickert, *Anisotropic diffusion in image processing*, Teubner, Stuttgart, 1998.
42. 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.
43. J. Weickert, *Coherence-enhancing diffusion of colour images*, Image and Vision Computing, Vol. 17, 199–210, 1999.
44. J. Weickert, B. ter Haar Romeny, M.A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE T.–Image Proc.*, Vol. 7, No. 3:398–410, 1998.
45. 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.
46. J. Weickert, J. Heers, C. Schnörr, K.-J. Zuiderveld, O. Scherzer, and H.-S. Stiehl. Fast parallel algorithms for a broad class of nonlinear variational diffusion approaches. *Real–Time Imaging*, 2000. In press.
47. J. Weickert, C. Schnörr. A theoretical framework for convex regularizers in PDE–based computation of image motion. Computer Science Series 13/2000, Dept. Math. and Comp. Science, University of Mannheim, Germany, June 2000. Submitted.