

Variational Optic Flow Computation: From Continuous Models to Algorithms

Joachim Weickert, Andrés Bruhn, Nils Papenberg, and Thomas Brox
Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science
Saarland University, Building 27
66041 Saarbrücken, Germany
{weickert, bruhn, papenberg, brox}@mia.uni-saarland.de

Abstract

Variational methods belong to the most successful techniques for computing the displacement field in image sequences. In this paper we analyse the different terms in the energy functional and sketch some of our recent contributions in this area.

1. Introduction

Already in 1981, Horn and Schunck introduced the first variational method for computing the displacement field (*optic flow*) in an image sequence [15]. This method is based on two assumptions that are characteristic for many variational optic flow methods: a brightness constancy assumption and a smoothness assumption. These assumptions enter a continuous energy functional whose minimiser yields the desired optic flow field. Performance evaluations such as [5, 11] showed that variational methods belong to the better performing techniques. It is thus not surprising that a lot of research has been carried out in order to improve these techniques even further: These amendments include refined model assumptions with discontinuity-preserving constraints [2, 10, 13, 21, 22, 25, 30] or spatiotemporal regularisation [6, 20, 31], improved data terms with modified constraints [3, 9, 21, 26] or nonquadratic penalisation [6, 14, 18, 29], and efficient multigrid algorithms [7, 12, 27, 32] for minimising these energy functionals.

The goal of the present paper is to analyse the data term and the smoothness term in detail and to survey some recent results on variational optic flow computation in our group. The paper is organised as follows: In Section 2 we sketch the general structure of these techniques. While Section 3 analyses the data term in more detail, a discussion of the different possibilities for smoothness constraints is given in Section 4. Algorithmic aspects are outlined in Section 5, and experiments are presented in Section 6.

2. General Structure

Let $f(x_1, x_2, x_3)$ denote some scalar-valued image sequence, where (x_1, x_2) is the location and x_3 denotes time. Often f is obtained by preprocessing some initial image sequence f_0 by convolving it with a Gaussian K_σ of standard deviation σ :

$$f = K_\sigma * f_0. \quad (1)$$

Let us assume that $D^k f$ describes the set of all partial (spatial and temporal) derivatives of f of order k , and that the optic flow field $u(x_1, x_2, x_3) = (u_1(x_1, x_2, x_3), u_2(x_1, x_2, x_3), 1)$ gives the displacement rate between subsequent frames. In the present paper we consider variational methods that are based on the minimisation of the continuous energy functional

$$E(u) = \int_{\Omega} \underbrace{M(D^k f, u)}_{\text{data term}} + \alpha \underbrace{S(\nabla f, \nabla u)}_{\text{regulariser}} dx \quad (2)$$

where the integration domain Ω is either a spatial or a spatiotemporal domain. In the spatial case we have $x := (x_1, x_2)^\top$ and $\nabla := \nabla_2 := (\partial_{x_1}, \partial_{x_2})^\top$, and in the spatiotemporal case we use the notations $x := (x_1, x_2, x_3)^\top$ and $\nabla := \nabla_3 := (\partial_{x_1}, \partial_{x_2}, \partial_{x_3})^\top$. The optic flow field $u(x_1, x_2, x_3)$ is obtained as a function that minimises $E(u)$. The energy functional $E(u)$ penalises all deviations from model assumptions. Typically it consists of a *data term* $M(D^k f, u)$ which expresses e.g. a brightness constancy assumption, and a *regulariser* $S(\nabla f, \nabla u)$ with $\nabla u := (\nabla u_1, \nabla u_2)^\top$ that penalises deviations from (piecewise) smoothness. The weight $\alpha > 0$ serves as *regularisation parameter*: Larger values correspond to more simplified flow fields.

It should be noted that such continuous energy functionals may be formulated in a rotationally invariant way. Results from numerical analysis show that consistent discretisations approximate this invariance under rotations arbitrarily well if the sampling is sufficiently fine. Moreover, if

the energy functional is convex, a unique minimiser exists that can be found in a relatively simple way by globally convergent algorithms. Variational optic flow methods are *global* methods: If there is not sufficient local information, the data term $M(D^k f, u)$ is so small that it is dominated by the smoothness term $\alpha S(\nabla f, \nabla u)$ which fills in information from more reliable surrounding locations. Thus, in contrast to local methods, the *filling-in effect* of global variational approaches always yields dense flow fields and no subsequent interpolation steps are necessary: Everything is automatically accomplished within a single variational framework.

3. Data Terms

Many differential methods for optic flow are based on the assumption that the grey values of image objects in subsequent frames do not change over time. Thus, if $(x_1(x_3), x_2(x_3))$ denotes the movement of some image structure, we obtain the following *optic flow constraint (OFC)* by applying the chain rule:

$$0 = \frac{df(x_1(x_3), x_2(x_3), x_3)}{dx_3} = f_{x_1} u_1 + f_{x_2} u_2 + f_{x_3}, \quad (3)$$

where $f_{x_i} := \partial_{x_i} f$. Note that the optic flow field satisfies $(u_1, u_2, 1)^\top = (\partial_{x_3} x_1, \partial_{x_3} x_2, 1)^\top$. In order to use Equation (3) within the energy functional (2), we penalise all deviations from zero by using the quadratic data term [15]

$$M_1(D^1 f, u) := (u^\top \nabla_3 f)^2. \quad (4)$$

This term can be modified in several ways:

1. One may assume that the optic flow is constant within some neighbourhood of order ρ . This leads to [17]

$$M_2(D^1 f, u) := K_\rho * ((u^\top \nabla_3 f)^2). \quad (5)$$

This data term offers advantages when noise is present.

2. Higher robustness under noise can also be achieved by penalising outliers less severely than a quadratic regulariser does: One may use a penaliser $\Psi(s^2)$ that is convex in s and increases less rapidly than quadratic functions, e.g. the regularised TV penaliser [24]

$$\Psi(s^2) = \sqrt{\varepsilon^2 + s^2}. \quad (6)$$

This modification transforms M_1 and M_2 into

$$M_3(D^1 f, u) := \Psi((u^\top \nabla_3 f)^2), \quad (7)$$

$$M_4(D^1 f, u) := \Psi(K_\rho * ((u^\top \nabla_3 f)^2)). \quad (8)$$

Instead of imposing constancy of the image brightness f along the path $(x_1(x_3), x_2(x_3))$, we may impose constancy

of the spatial brightness gradient $(f_{x_1}, f_{x_2})^\top$ along such a path [28]. This gives two equations:

$$u^\top \nabla_3 f_{x_1} = 0, \quad (9)$$

$$u^\top \nabla_3 f_{x_2} = 0. \quad (10)$$

Squaring and adding them produces the data term

$$M_5(D^2 f, u) := \sum_{i=1}^2 (u^\top \nabla_3 f_{x_i})^2. \quad (11)$$

In a similar way, imposing constancy of the (spatial) Hessian of f gives

$$M_6(D^3 f, u) := \sum_{i=1}^2 \sum_{j=1}^2 (u^\top \nabla_3 f_{x_i x_j})^2, \quad (12)$$

and constancy of the (spatial) Laplacian $\Delta_2 f$ yields

$$M_7(D^3 f, u) := (u^\top \nabla_3 (\Delta_2 f))^2. \quad (13)$$

There is no general rule which of these data terms should be preferred. While higher-order derivatives are more sensitive to noise, the data terms M_5 , M_6 and M_7 may offer advantages over M_1 when the brightness is not constant. On the other hand, M_1 and M_7 are more appropriate than M_5 and M_6 when non-translatory motion dominates. Thus the choice of the “best” data term will always depend on the specific problem.

4. Smoothness Terms

A taxonomy of the different possibilities to design smoothness constraints has been presented in [30]. It exploits the connection between regularisation methods and diffusion filtering: Minimising the energy functional (2) by means of steepest descend, we obtain a system of diffusion–reaction equations, where the diffusion term results from the regulariser $S(\nabla f, \nabla u)$, and the reaction term is induced by the data term $M(D^k f, u)$:

$$\partial_t u_1 = \partial_{x_1} S_{u_1, x_1} + \partial_{x_2} S_{u_1, x_2} - \frac{1}{\alpha} \partial_{u_1} M, \quad (14)$$

$$\partial_t u_2 = \partial_{x_1} S_{u_2, x_1} + \partial_{x_2} S_{u_2, x_2} - \frac{1}{\alpha} \partial_{u_2} M \quad (15)$$

where S_{u_i, x_j} denotes the partial derivative of S with respect to $\partial_{x_j} u_i$. The parameter t in this system of partial differential equations (PDEs) is a pure numerical parameter that should not be confused with the time x_3 of the image sequence. For $t \rightarrow \infty$, the steady state of the diffusion–reaction system is given by the Euler–Lagrange equations

$$0 = \partial_{x_1} S_{u_1, x_1} + \partial_{x_2} S_{u_1, x_2} - \frac{1}{\alpha} \partial_{u_1} M, \quad (16)$$

$$0 = \partial_{x_1} S_{u_2, x_1} + \partial_{x_2} S_{u_2, x_2} - \frac{1}{\alpha} \partial_{u_2} M. \quad (17)$$

Table 1: Taxonomy of optic flow regularisers (see [30]).

Name of Regulariser	$S(\nabla f, \nabla u)$
homogeneous [15]	$\sum_{i=1}^2 \nabla u_i ^2$
image-driven, isotropic [2]	$g(\nabla f ^2) \sum_{i=1}^2 \nabla u_i ^2$
image-driven, anisotropic [21]	$\sum_{i=1}^2 \nabla u_i^\top D(\nabla f) \nabla u_i$
flow-driven, isotropic [25]	$\Psi \left(\sum_{i=1}^2 \nabla u_i ^2 \right)$
flow-driven, anisotropic [30]	$\text{tr} \Psi \left(\sum_{i=1}^2 \nabla u_i \nabla u_i^\top \right)$

They constitute necessary conditions that a minimiser of $E(u)$ has to satisfy.

Let us now have a closer look at the impact of the regulariser. The simplest regulariser is the *homogeneous* regularisation of Horn and Schunck [15]. This quadratic regulariser of type $S(\nabla u) = |\nabla u_1|^2 + |\nabla u_2|^2$ penalises all deviations from smoothness of the flow field. It can be related to linear diffusion with a constant diffusivity. Thus, the flow field is blurred in a homogeneous way such that motion discontinuities may lose sharpness and get dislocated. It is thus not surprising that people have tried to construct a variety of discontinuity-preserving regularisers. Depending on the structure of the resulting diffusion term, we can classify a regulariser $S(\nabla f, \nabla u)$ as image-driven or flow-driven, and isotropic or anisotropic.

For *image-driven* regularisers, S is not only a function of the flow gradient ∇u but also of the image gradient ∇f . This function is chosen in such a way that it respects discontinuities in the image data. If only the gradient *magnitude* $|\nabla f|$ matters, the method is called *isotropic*. It can avoid smoothing at image edges. An *anisotropic* technique depends also on the *direction* of ∇f . Typically it reduces smoothing across edges of f (i.e. along ∇f), while smoothing along edges of f is still permitted. Image-driven regularisers can be related to linear diffusion processes.

Flow-driven regularisers take into account discontinuities of the unknown flow field u by preventing smoothing at or across flow discontinuities. If the resulting diffusion process uses a scalar-valued diffusivity that only depends on $|\nabla u|^2 := |\nabla u_1|^2 + |\nabla u_2|^2$, it is an *isotropic* process. Cases where also the direction of ∇u_1 and ∇u_2 matters are named *anisotropic*. Flow-driven regularisers lead to nonlinear diffusion processes.

Table 1 gives an overview of the different regularisers. As a rule of thumb, one can expect that flow-driven regularisers offer advantages over image-driven ones for




Figure 1: Example of a full multigrid implementation for four levels (from [7]). Starting from a coarse scale the solution is refined step by step.

highly textured sequences, where the numerous texture edges create an oversegmentation of the flow field. Moreover, anisotropic methods may give somewhat better results than isotropic ones, since the latter ones are too “lazy” at noisy discontinuities. More details can be found in [30].

5. Algorithms

For the numerical minimisation of the energy functional (2), two strategies are used very frequently:

In the first strategy, one discretises the parabolic diffusion–reaction system (14), (15) and recovers the optic flow field as the steady-state solution for $t \rightarrow \infty$. The simplest numerical scheme would be an explicit (Euler forward) finite difference scheme. More efficient methods include semi-implicit approaches that offer better stability properties at the expense of the need to solve linear systems of equations.

Alternatively, one can directly discretise the elliptic Euler-Lagrange equations (16), (17). This also requires to solve large linear or nonlinear systems of equations. Efficient methods for this task include *successive overrelaxation (SOR)* methods, *preconditioned conjugate gradient (PCG)* algorithms and *multigrid* techniques. Figure 1 shows an example of a full multigrid cycle with 4 levels. It has been used in [7, 8] for finding the minimum of a variational approach with data term M_2 and a homogeneous regulariser. On a 3.06 GHz PC, it was possible to compute up to 40 dense flow fields of size 200×200 pixels within a single second. This shows that computational efficiency is no problem for variational optic flow methods, when state-of-the-art numerical methods are used.

It should be noted that for convex energy functionals, there is no danger that any of these two methods gets trapped in a local minimum, since only one minimum exists and the method is globally convergent.

6. Experiments

We start our experiments by evaluating the impact of the data term. This is done in Table 2 where we

Table 2: Impact of the data term on the quality of the optic flow field. We used a spatial energy functional with homogeneous regularisation, and computed the average angular error (AAE) for the Yosemite sequence with clouds. The parameters σ and α have been optimised.

Constancy	Data Term	σ	α	AAE
Brightness	M_1	1.3	500	7.17°
Gradient	M_5	2.1	20	5.91°
Hessian	M_6	2.7	1.8	6.46°
Laplacian	M_7	2.5	3.0	6.18°

used the Yosemite sequence with clouds. This synthetic sequence and its ground truth flow field are available from <ftp://csd.uwo.ca> under the directory `pub/vision`. The experiments in Table 2 show that it can be worthwhile to replace the commonly used brightness constancy constraint by constraints that involve higher derivatives.

The influence of the regulariser is studied in Figure 2, which depicts a zoom into Nagel’s *Marble* sequence (i21www.ira.uka.de/image-sequences) together with the results for five spatial regularisers. As expected, homogeneous regularisation is fairly blurry, flow-driven regularisers offer advantages over image-driven ones in textured regions, and anisotropic regularisers perform better than isotropic ones.

Figure 3 presents a comparison between spatial and spatiotemporal energy functionals. It demonstrates that the additional assumption of temporal smoothness may lead to significantly improved results.

In Table 3 we juxtapose the angular errors of a number of optic flow methods. It shows that the spatiotemporal method in [29] – which combines the data term M_4 with an isotropic flow-driven regulariser – is one of the two best performing algorithms.

7. Summary and Extensions

In this paper we have outlined some basic design principles for variational optic flow methods, sketched their numerical implementation and studied their performance. Due to space limitations, we had to restrict ourselves to some of the most important features. There are several possibilities to improve the performance of these methods even further: One may for instance use non-linearised data terms [3, 6, 21], multilevel strategies that encourage convergence towards a global minimiser when nonconvex functionals are applied [3, 4, 18], and consider more sophisticated functionals in order to cope with occlusion problems [1, 23]. On the numerical side, parallelisation strategies can be investi-




Figure 2: (a) **Top left:** Detail from Frame 16 of the *Marble* sequence (128×128 pixels). (b) **Top right:** Optic flow magnitude for homogeneous regularisation. (c) **Middle left:** Image-driven isotropic regularisation (d) **Middle right:** Image-driven anisotropic regularisation. (e) **Bottom left:** Flow-driven isotropic regularisation (f) **Bottom right:** Flow-driven anisotropic regularisation. From [30].

gated, e.g. domain decomposition methods [16].

Acknowledgements. Our research has partly been funded by the *Deutsche Forschungsgemeinschaft (DFG)* and the *Graduiertenkolleg “Leistungsgarantien für Rechnerysteme”*. This is gratefully acknowledged.

References

- [1] L. Alvarez, R. Deriche, T. Papadopoulo, and J. Sánchez. Symmetrical dense optical flow estimation with occlusion detection. In A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, editors, *Computer Vision*




Figure 3: (a) **Left:** Detail of Frame 8 of the *Copenhagen hallway* sequence of size $256 \times 256 \times 16$ pixels. A person is moving towards the camera. (b) **Middle:** Computed flow field for a spatial approach with data term M_1 (brightness constancy) and isotropic flow-driven regularisation. (c) **Right:** Result for a spatiotemporal approach. From [31].

Table 3: Comparison between the results from the literature with 100 % density and our results using a 3-D functional with data term M_4 and an isotropic flow-driven regulariser. All data refer to the *Yosemite* sequence with cloudy sky. Multiscale means that some focusing strategy using linear scale-space or pyramids has been applied. AAE = average angular error. Adapted from [29].

Technique	multiscale	AAE
Horn/Schunck, original [5]	no	31.69°
Singh, step 1 [5]	no	15.28°
Anandan [5]	no	13.36°
Singh, step 2 [5]	no	10.44°
Nagel [5]	no	10.22°
Horn/Schunck, modified [5]	no	9.78°
Uras <i>et al.</i> , unthresholded [5]	no	8.94°
Alvarez/Weickert/Sánchez [3]	yes	5.53°
Mémin/Pérez (IEEE TIP) [18]	yes	5.38°
Weickert/Bruhn/Schnörr [29]	no	5.18°
Mémin/Pérez (ICCV '98) [19]	yes	4.69°

– *ECCV 2002*, volume 2350 of *Lecture Notes in Computer Science*, pages 721–736. Springer, Berlin, 2002.

- [2] L. Alvarez, J. Esclarín, M. Lefébure, and J. Sánchez. A PDE model for computing the optical flow. In *Proc. XVI Congreso de Ecuaciones Diferenciales y Aplicaciones*, pages 1349–1356, Las Palmas de Gran Canaria, Spain, September 1999.
- [3] L. Alvarez, J. Weickert, and J. Sánchez. Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision*, 39(1):41–56, August 2000.
- [4] P. Anandan. A computational framework and an algorithm for the measurement of visual motion. *International Journal of Computer Vision*, 2:283–310, 1989.
- [5] J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, February 1994.
- [6] M. J. Black and P. Anandan. Robust dynamic motion estimation over time. In *Proc. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 292–302, Maui, HI, June 1991. IEEE Computer Society Press.
- [7] A. Bruhn, J. Weickert, C. Feddern, T. Kohlberger, and C. Schnörr. Real-time optic flow computation with variational methods. In N. Petkov and M. A. Westenberg, editors, *Computer Analysis of Images and Patterns*, volume 2756 of *Lecture Notes in Computer Science*, pages 222–229. Springer, Berlin, 2003.
- [8] A. Bruhn, J. Weickert, C. Feddern, T. Kohlberger, and C. Schnörr. Variational optic flow computation in real-time. Technical Report 89, Dept. of Mathematics, Saarland University, Saarbrücken, Germany, June 2003.
- [9] A. Bruhn, J. Weickert, and C. Schnörr. Combining the advantages of local and global optic flow methods. In L. Van Gool, editor, *Pattern Recognition*, volume 2449 of *Lecture Notes in Computer Science*, pages 454–462. Springer, Berlin, 2002.
- [10] I. Cohen. Nonlinear variational method for optical flow computation. In *Proc. Eighth Scandinavian Conference on Image Analysis*, volume 1, pages 523–530, Tromsø, Norway, May 1993.
- [11] B. Galvin, B. McCane, K. Novins, D. Mason, and S. Mills. Recovering motion fields: an analysis of eight optical flow algorithms. In *Proc. 1998 British Machine Vision Conference*, Southampton, England, September 1998.

- [12] F. Glazer. Multilevel relaxation in low-level computer vision. In A. Rosenfeld, editor, *Multiresolution Image Processing and Analysis*, pages 312–330. Springer, Berlin, 1984.
- [13] F. Heitz and P. Bouthemy. Multimodal estimation of discontinuous optical flow using Markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(12):1217–1232, December 1993.
- [14] W. Hinterberger, O. Scherzer, C. Schnörr, and J. Weickert. Analysis of optical flow models in the framework of calculus of variations. *Numerical Functional Analysis and Optimization*, 23(1/2):69–89, May 2002.
- [15] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [16] T. Kohlberger, C. Schnörr, A. Bruhn, and J. Weickert. Domain decomposition for parallel variational optic flow computation. In B. Michaelis and G. Krell, editors, *Pattern Recognition*, volume 2781 of *Lecture Notes in Computer Science*, pages 196–203, Berlin, 2003. Springer.
- [17] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proc. Seventh International Joint Conference on Artificial Intelligence*, pages 674–679, Vancouver, Canada, August 1981.
- [18] E. Mémin and P. Pérez. Dense estimation and object-based segmentation of the optical flow with robust techniques. *IEEE Transactions on Image Processing*, 7(5):703–719, May 1998.
- [19] E. Mémin and P. Pérez. A multigrid approach for hierarchical motion estimation. In *Proc. Sixth International Conference on Computer Vision*, pages 933–938, Bombay, India, January 1998. Narosa Publishing House.
- [20] H.-H. Nagel. Extending the ‘oriented smoothness constraint’ into the temporal domain and the estimation of derivatives of optical flow. In O. Faugeras, editor, *Computer Vision – ECCV ’90*, volume 427 of *Lecture Notes in Computer Science*, pages 139–148. Springer, Berlin, 1990.
- [21] H.-H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
- [22] P. Nesi. Variational approach to optical flow estimation managing discontinuities. *Image and Vision Computing*, 11(7):419–439, September 1993.
- [23] M. Proesmans, L. Van Gool, E. Pauwels, and A. Oosterlinck. Determination of optical flow and its discontinuities using non-linear diffusion. In J.-O. Eklundh, editor, *Computer Vision – ECCV ’94*, volume 801 of *Lecture Notes in Computer Science*, pages 295–304. Springer, Berlin, 1994.
- [24] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [25] C. Schnörr. Segmentation of visual motion by minimizing convex non-quadratic functionals. In *Proc. Twelfth International Conference on Pattern Recognition*, volume A, pages 661–663, Jerusalem, Israel, October 1994. IEEE Computer Society Press.
- [26] C. Schnörr. Unique reconstruction of piecewise smooth images by minimizing strictly convex non-quadratic functionals. *Journal of Mathematical Imaging and Vision*, 4:189–198, 1994.
- [27] D. Terzopoulos. Image analysis using multigrid relaxation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(2):129–139, March 1986.
- [28] S. Uras, F. Girosi, A. Verri, and V. Torre. A computational approach to motion perception. *Biological Cybernetics*, 60:79–87, 1988.
- [29] J. Weickert, A. Bruhn, and C. Schnörr. Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. Technical Report 82, Dept. of Mathematics, Saarland University, Saarbrücken, Germany, April 2003.
- [30] J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, December 2001.
- [31] J. Weickert and C. Schnörr. Variational optic flow computation with a spatio-temporal smoothness constraint. *Journal of Mathematical Imaging and Vision*, 14(3):245–255, May 2001.
- [32] G. Zini, A. Sarti, and C. Lamberti. Application of continuum theory and multi-grid methods to motion evaluation from 3D echocardiography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(2):297–308, March 1997.