Combining the Advantages of Local and Global Optic Flow Methods

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

[1] Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science,
Building 27.1, Saarland University, 66041 Saarbrücken, Germany,
{bruhn,weickert}@mia.uni-saarland.de

[2] Computer Vision, Graphics, and Pattern Recognition Group,
Department of Mathematics and Computer Science,
University of Mannheim, 68131 Mannheim, Germany,
{schnoerr}@uni-mannheim.de

Abstract. Differential methods are frequently used techniques for optic flow computations. They can be classified into local methods such as the Lucas–Kanade technique or Bigün's structure tensor method, and into global methods such as the Horn–Schunck approach and its modifications. Local methods are known to be more robust under noise, while global techniques yield 100% dense flow fields. No clear attempts to combine the advantages of these two classes of methods have been made in the literature so far.

This problem is addressed in our paper. First we juxtapose the role of smoothing processes that are required in local and global differential methods for optic flow computation. This discussion motivates us to introduce and evaluate a novel method that combines the advantages of local and global approaches: It yields dense flow fields that are robust against noise. Finally experiments with different sequences are performed demonstrating its excellent results.

Keywords: visual motion, differential techniques, variational methods, structure tensor, partial differential equations.

1 Introduction

Differential methods belong to the most widely used techniques for optic flow estimation in image sequences. They are based on the computation of spatial and temporal image derivatives. Differential techniques can be classified into *local* methods that may optimize some local energy-like expression, and *global* strategies which attempt to minimize a global energy functional. Examples of the first category include the Lucas–Kanade method [9], the structure tensor approach of Bigün et al. [3] and its space–variant version by Nagel and Gehrke [12], but also techniques using second order derivatives such as [16]. Global approaches comprise the classic method of Horn and Schunck [6] and discontinuity-preserving

L. Van Gool (Ed.): DAGM 2002, LNCS 2449, pp. 454–462, 2002.

variants such as [10]. Together with phase-based methods [4], differential methods belong to the techniques with the best performance [2,5]. Local methods may offer relatively high robustness under noise, but do not give dense flow fields. Global methods, on the other hand, yield flow fields with 100 % density, but are experimentally known to be more sensitive to noise [2,5].

Almost all differential optic flow methods make use of smoothing techniques and smoothness assumptions: The actual role and the difference between these smoothing strategies, however, has hardly been addressed in the literature so far. In a first step of this paper we juxtapose the role of the different smoothing steps of these methods. We shall see that each smoothing process offers certain advantages that cannot be found in other cases. Consequently, it would be desirable to combine the different smoothing effects of local and global methods in order to design novel approaches that combine the high robustness of local methods with the high density of global techniques. One of the goals of the present paper is to propose and analyse such an embedding of local methods into global approaches. This results in a technique that is robust under noise and gives flow fields with 100 % density. Hence, there is no need for a postprocessing step where sparse data have to be interpolated.

Our paper is organized as follows. In Section 2 we discuss the role of the different smoothing processes that are involved in local and global optic flow approaches. Based on these results we propose two *combined local-global (CLG) methods* in Section 3, one with spatial, the other one with spatiotemporal smoothing. Section 4 is devoted to performance evaluations of the CLG method. Our paper is concluded with a summary in Section 5.

Related Work. Schnörr et al. [14] sketched a framework for supplementing global energy functionals with multiple equations that provide local data constraints. The local method in their experiments used the output of Gaussian filters shifted in frequency space [4]. Methods of Lucas–Kanade or Bigün type have not been considered in this context.

Our proposed technique differs from the majority of global regularization methods by the fact that we also use spatiotemporal regularizers instead of spatial ones. This relates our method to earlier work with spatiotemporal regularizers such as [11,17].

While the noise sensitivity of local differential methods has been studied intensively in recent years [1,7,8,13,15], the noise sensitivity of global differential methods has been analysed to a significantly smaller extent. In this context, Galvin et al. [5] have compared a number of classical methods where small amounts of Gaussian noise had been added. Their conclusion was similar to the findings of Barron et al. [2]: the global approach of Horn and Schunck is more sensitive to noise than the local Lucas–Kanade method.

2 Role of the Smoothing Processes

In this section we discuss the role of smoothing techniques in differential optic flow methods. For simplicity we focus on spatial smoothing. All spatial smoothing strategies can easily be extended into the temporal domain. This will usually lead to improved results.

Let us consider some image sequence $g(x, y, t)$, where (x, y) denotes the location within a rectangular image domain Ω, and $t \in [0, T]$ denotes time. It is common to smooth the image sequence prior to differentiation [2,8], e.g. by convolving each frame with some Gaussian $K_\sigma(x, y)$ of standard deviation σ:

$$f(x, y, t) \ := \ (K_\sigma * g)(x, y, t), \tag{1}$$

The low-pass effect of Gaussian convolution removes noise and other destabilizing high frequencies. In a subsequent optic flow method, we may thus call σ the *noise scale*. While some moderate presmoothing improves the results, great care should be taken not to apply too much presmoothing, since this would severely destroy important image structure.

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. For small displacements this yields the *optic flow constraint*

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

where the displacement field $(u, v)^\top(x, y, t)$ is called *optic flow* and subscripts denote partial derivatives. Evidently, this single equation is not sufficient to uniquely compute the two unknowns u and v *(aperture problem)*: For nonvanishing image gradients, it is only possible to determine the flow component parallel to $\nabla f := (f_x, f_y)^\top$, i.e. normal to image edges, the so-called *normal flow*. In order to cope with the aperture problem, Lucas and Kanade [9] proposed to assume that the unknown optic flow vector is constant within some neighbourhood of size ρ. In this case it is possible to determine the two *constants* u and v at some location (x, y, t) from a weighted least square fit by minimizing the function

$$E_{LK}(u, v) \ := \ K_\rho * \left((f_x u + f_y v + f_t)^2\right). \tag{3}$$

Here the standard deviation ρ of the Gaussian serves as an *integration scale* over which the main contribution of the least square fit is computed. Therefore the effect to the neighborhood is limited by the value of ρ. As a result structures of the same order do occur. In particular, a sufficiently large value for ρ is very successful in rendering the Lucas–Kanade method robust under noise.

A minimum (u, v) of E_{LK} satisfies $\partial_u E_{LK} = 0$ and $\partial_v E_{LK} = 0$. This gives the linear system

$$\begin{pmatrix} K_\rho * (f_x^2) & K_\rho * (f_x f_y) \\ K_\rho * (f_x f_y) & K_\rho * (f_y^2) \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \ = \ \begin{pmatrix} -K_\rho * (f_x f_t) \\ -K_\rho * (f_y f_t) \end{pmatrix} \tag{4}$$

which can be solved provided that its system matrix is invertible. This is not the case in flat regions where the image gradient vanishes. In some other regions, the smaller eigenvalue of the system matrix may be close to 0, such that the aperture problem remains present and the data do not allow a reliable determination of the full optic flow. All this results in nondense flow fields. They constitute the most severe drawback of local gradient methods: Since many computer vision applications require dense flow estimates, subsequent interpolation steps are required.

In order to end up with dense flow estimates one may embed the optic flow constraint into a regularization framework. Horn and Schunck [6] have pioneered this class of global differential methods. They determine the unknown *functions* $u(x, y, t)$ and $v(x, y, t)$ as the minimizers of the global energy functional

$$E_{HS}(u, v) \;=\; \int_{\Omega} \left((f_x u + f_y v + f_t)^2 \,+ \alpha \left(|\nabla u|^2 + |\nabla v|^2 \right) \right) dx\, dy \qquad (5)$$

where the smoothness weight $\alpha > 0$ serves as *regularization parameter:* Larger values for α result in a stronger penalization of large flow gradients and lead to smoother flow fields. The unique minimizer of this convex functional benefits from the *filling-in effect:* At locations with $|\nabla f| \approx 0$, no reliable local flow estimate is possible, but the regularizer $|\nabla u|^2 + |\nabla v|^2$ fills in information from the neighbourhood. This results in dense flow fields and makes subsequent interpolation steps obsolete. This is a clear advantage over local methods.

It has been observed that global methods may be more sensitive to noise than local differential methods [2,5]. An explanation for this behaviour can be given as follows. Noise results in high image gradients. They serve as weights in the data term of the regularization functional (5). Since the smoothness term has a constant weight α, smoothness is relatively less important at locations with high image gradients than elsewhere. As a consequence, *flow fields are less regularized at noisy image structures.* This sensitivity under noise is therefore nothing else but a side-effect of the desired filling-in effect. Increasing the regularization parameter α will finally also smooth the flow field at noisy structures, but at this stage, it might already be too blurred in flatter image regions.

3 A Combined Local–Global Method

We have seen that both local and global differential methods have complementary advantages and shortcomings. Hence it would be interesting to construct a hybrid technique that constitutes the best of two worlds: It should combine the robustness of local methods with the density of global approaches. This shall be done next. We start with spatial formulations before we extend the approach to the spatiotemporal domain. In order to design a *combined local–global (CLG) method*, let us first reformulate the previous approaches. Using the notations

$$\begin{aligned} w &:= (u, v, 1)^\top, & |\nabla w|^2 &:= |\nabla u|^2 + |\nabla v|^2, \\ \nabla_3 f &:= (f_x, f_y, f_t)^\top, & J_\rho(\nabla_3 f) &:= K_\rho * (\nabla_3 f\, \nabla_3 f^\top) \end{aligned}$$

it becomes evident that the Lucas–Kanade method minimizes the quadratic form

$$E_{LK}(w) = w^\top J_\rho(\nabla_3 f)\, w, \qquad (6)$$

while the Horn–Schunck technique minimizes the functional

$$E_{HS}(w) = \int_{\Omega} \left(w^\top J_0(\nabla_3 f)\, w + \alpha |\nabla w|^2 \right) \, dx\, dy. \tag{7}$$

This terminology suggests a natural way to extend the Horn–Schunck functional to the desired CLG functional. We simply replace the matrix $J_0(\nabla_3 f)$ by the structure tensor $J_\rho(\nabla_3 f)$ with some integration scale $\rho > 0$. Thus, we propose to minimize the functional

$$E_{CLG}(w) = \int_{\Omega} \left(w^\top J_\rho(\nabla_3 f)\, w + \alpha |\nabla w|^2 \right) \, dx\, dy. \tag{8}$$

Its minimizing flow field (u, v) satisfies the Euler–Lagrange equations

$$\Delta u - \tfrac{1}{\alpha} \left(K_\rho * (f_x^2)\, u + K_\rho * (f_x f_y)\, v + K_\rho * (f_x f_t) \right) = 0, \tag{9}$$

$$\Delta v - \tfrac{1}{\alpha} \left(K_\rho * (f_x f_y)\, u + K_\rho * (f_y^2)\, v + K_\rho * (f_y f_t) \right) = 0, \tag{10}$$

where Δ denotes the Laplacean.

A spatiotemporal variant of the Lucas–Kanade approach is due to Bigün et al. [3]. It replaces convolution with 2-D Gaussians by spatiotemporal convolution with 3-D Gaussians. This still leads to a 2×2 linear system of equations for the two unknowns u and v.

A spatiotemporal version of our CLG functional is given by

$$E_{CLG3}(w) = \int_{\Omega \times [0,T]} \left(w^\top J_\rho(\nabla_3 f)\, w + \alpha |\nabla_3 w|^2 \right) \, dx\, dy\, dt \tag{11}$$

The Euler–Lagrange equations in the spatiotemporal setting have the same structure as (9)–(10), apart from the fact that spatiotemporal Gaussian convolution is used, and that the spatial Laplacean is replaced by the spatiotemporal one due to $|\nabla_3 w|^2$. In general, the spatiotemporal Gaussians may have different standard deviations in space and time.

4 Experiments

For our experiments a standard finite difference discretization of the Euler–Lagrange equations (9)–(10) is used. The resulting sparse linear system of equations is solved iteratively by an SOR scheme. Apart from the first iteration, where additional convolutions with K_ρ are computed, the CLG method is as fast as the Horn and Schunck algorithm.

Figure 1 shows our first experiment. It depicts a flight through to the Yosemite National Park where divergent motion is dominating. The original synthetic sequence was created by Lynn Quam. A modified variant without clouds is available from `http://www.cs.brown.edu/people/black/images.html`.

We have added Gaussian noise with zero mean and different standard deviation to this sequence, and we used the 3-D CLG method for computing the

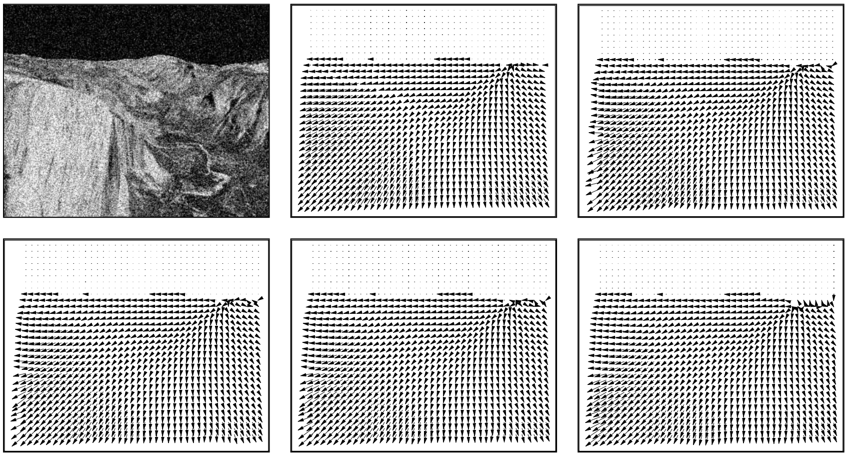

Fig. 1. *(a) Top left:* Frame 8 of the Yosemite sequence severely degraded by Gaussian noise with $\sigma_n = 40$. *(b) Top middle:* Ground truth flow field. *(c) Top right:* Computed flow field for $\sigma_n = 0$. *(d) Bottom left:* Ditto for $\sigma_n = 10$. *(e) Bottom middle:* $\sigma_n = 20$. *(f) Bottom right:* $\sigma_n = 40$.

flow field. Figure 1(c) shows that the recovered flow field is not very sensitive to Gaussian noise and that it coincides well with the ground truth flow field in Figure 1(b). These qualitative results are confirmed by the quantitative evaluations in Table 1, where we have studied the effect of replacing spatial smoothing steps by spatiotemporal ones. As one may expect, both the quality of the optic flow estimates and their robustness under Gaussian noise improve when temporal coherence is taken into account. The angular error has been computed as proposed in [2].

Another example demonstrating the excellent results of our CLG technique with spatiotemporal regularization is the Marble sequence. This non-synthetical sequence created by Nagel and Otte is available at the following internet adress: `http://i21www.ira.uka.de/image_sequences`

In Figure 2 (b) the ground truth flow field is shown, whereby the grey value at a pixel is related to the length of its displacement vector. As one can see the flow field estimated by our 3-D CLG technique in Figure 2 (c) matches the ground truth very well. This impression is confirmed by an average angular error

Table 1. Results for the 2-D and 3-D CLG method using the *Yosemite* sequence without clouds. Gaussian noise with varying standard deviations σ_n was added, and the average angular errors and their standard deviations were computed.

	σ_n=0	σ_n=10	σ_n=20	σ_n=40
2-D CLG	$2.64° \pm 2.27°$	$4.45° \pm 2.94°$	$6.93° \pm 4.31°$	$11.30° \pm 7.41°$
3-D CLG	$1.79° \pm 2.34°$	$2.53° \pm 2.75°$	$3.47° \pm 3.37°$	$5.34° \pm 3.81°$

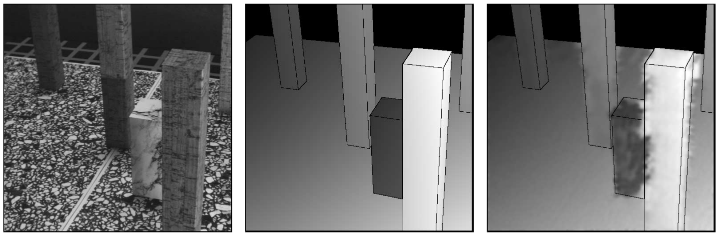

Fig. 2. *From left to right:* (a) Frame 16 of the Marble sequence. (b) Ground truth flow field (length). (c) Computed flow field (length).

Table 2. Stability of the 2-D CLG method under parameter variations. The data refer to the *marble* sequence without noise. AAE = average angular error.

σ	ρ	α	AAE
1.30	1.8	1000	5.70°
1.73	"	"	5.45°
2.60	"	"	5.30°
3.90	"	"	5.52°
5.20	"	"	6.05°

σ	ρ	α	AAE
2.60	0.9	1700	5.31°
"	1.2	"	5.30°
"	1.8	"	5.30°
"	2.7	"	5.30°
"	3.6	"	5.31°

σ	ρ	α	AAE
2.60	1.8	500	5.40°
"	"	666	5.34°
"	"	1000	5.30°
"	"	1500	5.33°
"	"	2000	5.39°

of 2.06°. A comparison to our result for the 2-D variant of 5.30° demonstrates a strong improvement by spatiotemporal regularization one more time.

Let us now investigate the sensitivity of the CLG method with respect to parameter variations. This is done in Table 2 for the *marble* sequence. We observe that the average angular error does hardly deteriorate when two parameters are fixed, while the other one varies by a factor 4. This stability under parameter variations may be regarded as another experimental confirmation of the well-posedness of the CLG approach. Moreover, this also indicates that the method performs sufficiently robust in practice even if non-optimized default parameter settings are used.

5 Summary and Conclusions

In this paper we have analysed the smoothing effects in local and global differential methods for optic flow computation. As prototypes of local methods we used the spatial least-square fit of Lucas and Kanade [9] and the spatiotemporal structure tensor method of Bigün et al. [3], while the Horn and Schunck approach [6] was our representative for a global method. We saw that the smoothing steps in each of these methods serve different purposes and have different advantages and shortcomings. As a consequence, we proposed a combined local-global (CLG) approach that incorporates the advantages of both paradigms: It is highly robust under Gaussian noise while giving dense flow fields. Experiments have also shown that the CLG method is not very sensitive under parameter variations. Nonlinear CLG techniques will be described in forthcoming papers.

Acknowledgements

Our research has been partly funded by the *Deutsche Forschungsgemeinschaft (DFG)* under the project SCHN 457/4-1.

References

1. A. Bainbridge–Smith and R. G. Lane. Determining optical flow using a differential method. *Image and Vision Computing*, 15(1):11–22, Jan. 1997.
2. J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, Feb. 1994.
3. J. Bigün, G. H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, Aug. 1991.
4. D. J. Fleet and A. D. Jepson. Computation of component image velocity from local phase information. *International Journal of Computer Vision*, 5(1):77–104, Aug. 1990.
5. 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, Sept. 1998.
6. B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
7. B. Jähne. *Digital Image Processing*. Springer, Berlin, 1997.
8. J. K. Kearney, W. B. Thompson, and D. L. Boley. Optical flow estimation: an error analysis of gradient-based methods with local optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(2):229–244, Mar. 1987.
9. 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, Aug. 1981.
10. H.-H. Nagel. Constraints for the estimation of displacement vector fields from image sequences. In *Proc. Eighth International Joint Conference on Artificial Intelligence*, volume 2, pages 945–951, Karlsruhe, West Germany, August 1983.
11. 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.
12. H.-H. Nagel and A. Gehrke. Spatiotemporally adaptive estimation and segmentation of OF-fields. In H. Burkhardt and B. Neumann, editors, *Computer Vision – ECCV '98*, volume 1407 of *Lecture Notes in Computer Science*, pages 86–102. Springer, Berlin, 1998.
13. N. Ohta. Uncertainty models of the gradient constraint for optical flow computation. *IEICE Transactions on Information and Systems*, E79-D(7):958–962, July 1996.
14. C. Schnörr, R. Sprengel, and B. Neumann. A variational approach to the design of early vision algorithms. *Computing Suppl.*, 11:149–165, 1996.
15. E. P. Simoncelli, E. H. Adelson, and D. J. Heeger. Probability distributions of optical flow. In *Proc. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 310–315, Maui, HI, June 1991. IEEE Computer Society Press.

16. O. Tretiak and L. Pastor. Velocity estimation from image sequences with second order differential operators. In *Proc. Seventh International Conference on Pattern Recognition*, pages 16–19, Montreal, Canada, July 1984.
17. 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.