

Symmetrical Dense Optical Flow Estimation with Occlusions Detection

Luis Alvarez¹, Rachid Deriche², Théo Papadopoulos², and Javier Sánchez¹

¹ Universidad de Las Palmas, Spain

{`lalvarez, jsanchez`}@dis.ulpgc.es

² INRIA Sophia-Antipolis, France

{`der, papadop`}@sophia.inria.fr

Abstract. Traditional techniques of dense optical flow estimation don't generally yield symmetrical solutions: the results will differ if they are applied between images I_1 and I_2 or between images I_2 and I_1 . In this work, we present a method to recover a dense optical flow field map from two images, while explicitly taking into account the symmetry across the images as well as possible occlusions and discontinuities in the flow field. The idea is to consider both displacements vectors from I_1 to I_2 and I_2 to I_1 and to minimise an energy functional that explicitly encodes all those properties. This variational problem is then solved using the gradient flow defined by the Euler-Lagrange equations associated to the energy. In order to reduce the risk to be trapped within some irrelevant minimum, a focusing strategy based on a multi-resolution technique is used to converge toward the solution. Promising experimental results on both synthetic and real images are presented to illustrate the capabilities of this symmetrical variational approach to recover accurate optical flow.

1 Introduction and Motivation

A large number of methods have been proposed in the computer vision community to address the important problem of motion analysis from image sequences. In this paper, we are interested in computing the 2D optical flow field which is a specific type of motion defined as the velocity field obtained from the temporal changes of the intensity values of the image sequence. 2D optical flow estimation has been extensively addressed in the literature, and the most common methods are correlation, gradient, spatiotemporal filtering, Fourier phase and energy based approaches. We refer the interested reader to [6,17] for an excellent analysis and evaluation of important optical flow methods. Many of these approaches use the classical constraint equation that relates the gradient of brightness to the components u and v of the local flow to estimate the optical flow. Because this problem is ill-posed, additional constraints are usually required. The most used one is to add a quadratic constraint on the gradient magnitude of the flow to impose some form of smoothness on the flow field as done originally by Horn and Schunk [16]. Other formulations have been proposed using different smoothness constraints, but despite their interest, they clearly lacked robustness

to the presence of occlusions and discontinuities. In order to estimate the optical flow more accurately, one have to explicitly take into account the problem of occlusions and discontinuities. See the following works, mostly based on variational approaches [20,19,10,24,15,14,22,13,5,18,11,3,31]. Due to the fact that the functional to be minimised is generally not convex, some focusing strategy embedding the method in a multi-resolution scheme or a linear scale-space have been successfully applied to reduce the risk to get trapped in some irrelevant minima [20,4,3].

The method presented here is inspired from this kind of variational framework that has also proven recently to be very useful in many other image processing and computer vision tasks [25,30,8,1]. We start explicitly with the method described in [4] and modify it using ideas similar to those developed in [9] to obtain an algorithm that will give the same solution if applied between images A and B or between images B and A. Contrarily to [9], we deal with discontinuities of the flow field and explicitly take into account the possibility of occlusions.

The article is organised as follows: in section 2, we will introduce some notations and basic concepts on which the approach will be based, section 3 describes the method which is demonstrated on some examples in section 4.

2 Notations and Previous Work

In this paper, images are represented as functions that map some coordinate space \mathcal{D} to an intensity space \mathcal{I} . We consider that two images, I_1 and I_2 , of the same scene are given:

$$I_1 : \mathcal{D}_1 \mapsto \mathcal{I} \quad \text{and} \quad I_2 : \mathcal{D}_2 \mapsto \mathcal{I} .$$

Note that it is assumed here that the two images share the same range \mathcal{I} . This assumption can be made without loss of generality as it is always possible to preprocess the images in order to normalize them. Preprocessing also allows us to limit ourselves on SSD-like (Sum of Squared Differences) criteria to compare the values of I_1 and I_2 . The basic problem addressed by this paper is establishing dense correspondences between those two images. To do so, two unknown functions, h_1 and h_2 , are introduced:

$$h_1 : \Omega_1 \subset \mathcal{D}_1 \mapsto \Omega_2 \quad \text{and} \quad h_2 : \Omega_2 \subset \mathcal{D}_2 \mapsto \Omega_1 .$$

For each point $\mathbf{m}_1 = (x_1, y_1)$ of Ω_1 (resp. $\mathbf{m}_2 = (x_2, y_2)$ of Ω_2), the function h_1 (resp. h_2) give its corresponding point $h_1(\mathbf{m}_1) = [u_1(x_1, y_1), v_1(x_1, y_1)]^T$ in Ω_2 (resp. the point $h_2(\mathbf{m}_2) = [u_2(x_2, y_2), v_2(x_2, y_2)]^T$ in Ω_1). Note that we have been careful to define these functions on subsets of \mathcal{D}_1 and \mathcal{D}_2 so that we can take account of occlusions in both images which are defined by the sets $\mathcal{D}_1 \setminus \Omega_1$ and $\mathcal{D}_2 \setminus \Omega_2$ respectively. Consequently, a good matching should satisfy the following two properties:

$$h_2 \circ h_1 = Id_{\Omega_1} \quad \text{and} \quad h_1 \circ h_2 = Id_{\Omega_2} , \quad (1)$$

where Id_{Ω_i} denotes the identity over the domain Ω_i . These equations mean that there is a one-to-one correspondence between the sets Ω_1 and Ω_2 (ie that the matched points are corresponding to each other) or equivalently that $h_1 = h_2^{-1}$ and $h_2 = h_1^{-1}$.

2.1 Dense Optical Flow

Given the previous notations, optical flow computation procedures have been defined as variational scheme that minimise an energy functional of the form:

$$E_1(h_1) = E_1^M(h_1) + \alpha E_1^R(h_1) , \quad (2)$$

where α is a weighting factor and $E_1^M(h_1)$ and $E_1^R(h_1)$ represents the matching and regularisation costs, respectively. The matching term can be written as:

$$E_1^M(h_1) = \frac{1}{\max \|\nabla I_1\|^2} \int_{\Omega_1} (I_1(\mathbf{m}_1) - I_2(h_1(\mathbf{m}_1)))^2 d\mathbf{m}_1$$

As announced previously, $E_1^M(h_1)$ corresponds to a SSD-like criterion, which is allowed by our assumptions. The factor in front of the integral is used to normalize the relative strength of the terms $E_1^M(h_1)$ and $E_1^R(h_1)$. Various different proposals have been made for the regularisation term. To prevent the flow map to be smoothed by the algorithm across images boundaries, we have chosen to use Nagel-Enkelmann operator [20] , first because its efficiency has been demonstrated numerous times in the context of optical flow estimations [3,6,12,19,21,20,26,27] and because of its simplicity since the underlying second order differential operator is linear.

$$E_1^R(h_1) = \int_{\Omega_1} \Phi(D(\|\nabla I_1\|), \nabla h_1) d\mathbf{m}_1 ,$$

with:

$$\Phi(D(\nabla I_i), \nabla h_i) = \text{trace}(\nabla(h_i)^t D(\nabla I_i) \nabla(h_i)) ,$$

where $\nabla(h_i)$ denotes the Jacobian matrix of the function $h_i(\mathbf{m}_i)$. Since $h_i(\mathbf{m}_i) = [u_i(\mathbf{m}_i), v_i(\mathbf{m}_i)]^T$, the previous formula can also be written as:

$$\Phi(D(\nabla I_i), \nabla h_i) = \nabla u_1^t D(\nabla I_1) \nabla u_1 + \nabla u_2^t D(\nabla I_1) \nabla u_2$$

$D(\nabla I_1)$ is a regularised projection matrix in the direction perpendicular to ∇I_1

$$D(\nabla I_1) = \frac{1}{\|\nabla I_1\|^2 + 2\nu^2} \left\{ \begin{bmatrix} \frac{\partial I_1}{\partial y_1} \\ -\frac{\partial I_1}{\partial x_1} \end{bmatrix} \begin{bmatrix} \frac{\partial I_1}{\partial y_1} \\ -\frac{\partial I_1}{\partial x_1} \end{bmatrix}^t + \nu^2 Id \right\} . \quad (3)$$

In this formulation, Id denotes the identity matrix and ν is a parameter used to control the desired level of isotropy.

It is interesting to note that the matrix $D(\nabla I_1)$ plays the role of a diffusion tensor. Its eigenvectors are $v_1 := \nabla I_1$ and $v_2 := \nabla I_1^\perp$, and the corresponding eigenvalues are given by

$$\lambda_1(\|\nabla I_1\|) = \frac{\nu^2}{\|\nabla I_1\|^2 + 2\nu^2}, \quad (4)$$

$$\lambda_2(\|\nabla I_1\|) = \frac{\|\nabla I_1\|^2 + \nu^2}{\|\nabla I_1\|^2 + 2\nu^2}. \quad (5)$$

In the interior of objects we have $\|\nabla I_1\| \rightarrow 0$, and therefore $\lambda_1 \rightarrow 1/2$ and $\lambda_2 \rightarrow 1/2$. At ideal edges where $\|\nabla I_1\| \rightarrow \infty$, we obtain $\lambda_1 \rightarrow 0$ and $\lambda_2 \rightarrow 1$. Thus, we have isotropic behaviour within regions, and at image boundaries the process smoothes anisotropically along the edge. This behaviour is very similar to edge-enhancing anisotropic diffusion filtering [29,28], and it is also close in spirit to the modified mean-curvature motion considered in [2,1].

At this point, two remarks can be made:

- Since Ω_1 is usually unknown, the domain \mathcal{D}_1 is used instead in the previous integrals.
- The criterion depicted above is clearly non-symmetrical in the images I_1 and I_2 so that the minimisation of $E_1(h_1)$ and that of the similarly defined $E_2(h_2)$ will not yield functions h_1 and h_2 that satisfy the constraints (1). In other words, the two procedures will not give the same matching which is unfortunate.

In the following section, we will discuss a method to reintroduce the symmetry across the two images while allowing simultaneously the detection of occlusions. Section 4 describes results obtained with this method.

3 Symmetric Optical Flow

Let us first assume that the function h_2 is known. The first step is to extend the energy depicted in equation (2) by adding a term $E_1^S(h_1)$ to impose the first of the two constraints (1):

$$E_1'(h_1, h_2) = E_1^M(h_1) + \alpha E_1^R(h_1) + \beta E_1^S(h_1, h_2), \quad (6)$$

where β is another weighting factor. $E_1^S(h_1)$ can be written as:

$$E_1^S(h_1, h_2) = \int_{\Omega_1} \Psi \left(\|h_2(h_1(\mathbf{m}_1)) - Id_{\Omega_1}(\mathbf{m}_1)\|^2 \right) d\mathbf{m}_1,$$

where Ψ is a function that will be used to provide some robustness in the method. This function will be detailed in section 3.3. For the time being, we can assume

that $\Psi(s) = s$. A large value of $\|h_2(h_1(\mathbf{m}_1)) - Id_{\Omega_1}(\mathbf{m}_1)\|$ means a lack of symmetry in the matching (h_2 being given). Assuming that h_2 is correct, when the solution for h_1 is reached, the most probable reasons for such a situation is an occlusion problem or an error in the estimated matching. Notice that taking only the first of the two constraints (1) just imposes that h_1 is injective and that h_2 is mapping from $h_1(\Omega_1)$ onto Ω_1 .

It is possible to define $E'_2(h_1, h_2)$ in a similar fashion, with the roles of $h_i, i = 1..2$ and the $\Omega_i, i = 1..2$ reversed. A completely symmetric criterion is then obtained easily by:

$$E'(h_1, h_2) = E'_1(h_1, h_2) + E'_2(h_1, h_2) .$$

3.1 A Multi-resolution Scheme Approach to Escape from Local Minima

In general, the Euler–Lagrange equation associated to the energy E' will have multiple solutions. As a consequence, the asymptotic state of the parabolic equation that will be used to minimize E' , will depend on the initial data (h_1^0, h_2^0) . Typically, we may expect the algorithm to converge toward a local minimum of the energy functional that is located in the vicinity of that initial data. To reduce the risk of being trapped into irrelevant local minima, we embed our method into a multi-resolution framework. At the coarsest-scale the problem is much smoother, so that many irrelevant local minima disappear. Using the coarse-scale solution as initialisation for finer scales helps in getting close to the most relevant global minimum. A detailed analysis of the usefulness of such a focusing strategy in the context of a related optic flow problem can be found for instance in [4].

3.2 Minimising the Energy

The previous energy function is minimised as the asymptotic solution of a gradient flow based on the Euler–Lagrange partial differential equation corresponding to the energy E' . In order to look for the minimum of this energy, we proceed in an iterative way as follows: first we begin with an initial approximation (h_1^0, h_2^0) and then obtain (h_1^{n+1}, h_2^{n+1}) from (h_1^n, h_2^n) as an iteration of a gradient descent method applied to the energy:

$$E'(h_1^{n+1}, h_2^{n+1}) = E'_1(h_1^{n+1}, h_2^n) + E'_2(h_1^n, h_2^{n+1})$$

The initial approximation is obtained by a classical correlation at the coarsest level of resolution. For the gradient descent, applying Euler-Lagrange, and using the notation $I^h = I \circ h$ for compactness, we obtain:

$$\begin{aligned} \frac{\partial E'}{\partial h_1} = & - \left(I_1 - I_2^{h_1} \right) (\nabla I_2)^{h_1} - \alpha \operatorname{div} (D(\|\nabla I_1\|) \nabla h_1) + \\ & \beta \Psi' \left(\|h_2 \circ h_1 - Id_{\Omega_1}\|^2 \right) (\nabla h_2)^{h_1} \cdot (h_2 \circ h_1 - Id_{\Omega_1}) \end{aligned}$$

$$\begin{aligned} \frac{\partial E'}{\partial h_2} = & - \left(I_2 - I_1^{h_2} \right) (\nabla I_1)^{h_2} - \alpha \operatorname{div} (D(\|\nabla I_1\|) \nabla h_2) + \\ & \beta \Psi' \left(\|h_1 \circ h_2 - \operatorname{Id}_{\Omega_2}\|^2 \right) (\nabla h_1)^{h_2} \cdot (h_1 \circ h_2 - \operatorname{Id}_{\Omega_2}) \end{aligned}$$

Using coordinates, the updates for the evolution of the PDE are:

$$\begin{pmatrix} \frac{\partial u_1}{\partial t} \\ \frac{\partial v_1}{\partial t} \end{pmatrix} = - \frac{\partial E'}{\partial h_1} \quad \begin{pmatrix} \frac{\partial u_2}{\partial t} \\ \frac{\partial v_2}{\partial t} \end{pmatrix} = - \frac{\partial E'}{\partial h_2}$$

Given a time step dt , these equations give the updates du_1, dv_1, du_2, dv_2 that must be applied to (h_1^n, h_2^n) to obtain (h_1^{n+1}, h_2^{n+1}) .

3.3 Choosing the Function Ψ

As said above, $\Psi(\cdot)$ can be chosen to improve the robustness of the algorithm. We have tested two possibilities:

$$\Psi_1(s) = s, \quad \Psi_2(s) = \frac{s}{\gamma} e^{1 - \frac{s}{\gamma}},$$

where γ is a threshold for which we consider there is a good matching. Note that the function Ψ_1 corresponds to standard least-squares (remember that the function Ψ is applied to the square of the error values), whereas Ψ_2 is similar in essence to least-trimmed-squares which can be used to deal with outliers. Indeed, with least-trimmed-squares, only the smallest residuals are summed up into the criterion. This is what achieves the function Ψ_2 (see figure 1): when the error is too large, the function Ψ_2 cancels its contribution to the criterion whereas for small errors it has a linear behaviour. These two constraints along with the fact that a variational approach is used so that the function chosen must be differentiable, imposes the general shape of the Ψ function. Ψ_2 is such a function (other choices obeying to the same constraints can be made).

In preliminary experiments, we have noticed that the locations where the flows obtained as the solutions of E_1 and E_2 are the least symmetrical with each other can be associated with the occlusions of the two images. We will use this property to recover those locations. To do so, the domain of integration will be extended to the whole images (ie \mathcal{D}_1 and \mathcal{D}_2 respectively), but the symmetry errors will be weighted by the Ψ_2 function. This weighting ensures that the large errors in symmetry will not affect the criterion. As a consequence the sub-domain Ω_1 can be found as the domain where the un-weighted errors are too large. A better choice of the function Ψ is under consideration [7].

4 Experimental Results

We now show results on both synthetic and real images. In both cases, the theoretical flow is available, and a comparison of the obtained and the expected




Fig. 1. $\Psi_2(s^2)$

results will be made. For both the experiments, the algorithm has been run for 200 iterations with a time step $dt = 10$. In [4], the non-symmetric method is compared with the other (the ones used in [6]) classical dense optical flow methods that give a density of matches of 100%. Typically, the non-symmetric method yields better results than the classical ones. Consequently, in this section, the comparison is limited to the non-symmetric and symmetric methods which have been depicted in the previous sections. This is the most relevant option.

4.1 Synthetic Images

We have first tried our algorithm with the set of synthetic images depicted in figure 2. Occlusion with these images is important so that they were processed with the robust version, ie $\Psi = \Psi_2$. Figure 3 show the theoretical result that should be obtained.




Fig. 2. The initial input images

Figure 4 shows the results obtained with our algorithm. A pyramid of three levels of resolution have been used with $\alpha = 1$, $\beta = 0.5$, $\gamma = 5$ and $\nu = 0.1$. Notice how the smoothness constraint, managed to give reasonable results (of course there are multiple solutions, but the algorithm choose one of minimal cost with respect to the smoothness constraint) in homogeneous regions while being able to respect the discontinuities. Figure 5 shows the computed occlusion




Fig. 3. The true optic flow. On the left (resp. right), the u (resp. v) component of the true flow. On the top (resp bottom) is the true function h_1 (resp. h_2).

images. These were obtained by selecting the image points where the symmetry conditions are violated by more than a fixed threshold ($\gamma = 5$). As it can be seen, these locations correspond fairly well to the occlusions. However, these have been slightly eroded on the background side of the occluding region, as the smoothing term of the partial differential equation has been in action at those locations.

In addition, table 1 shows the two error measures for both the symmetric and the non-symmetric cases and for both the left-to-right and the right-to-left flows. The two measures are the average of Euclidean norm of the difference between the true and the computed flow and the average angle between those two flows as defined in [6]. These measures are averaged only on the squares as the flow is undefined on the background with these images. It can be seen that errors are roughly divided by two with symmetrisation.

Finally, figure 6 shows the residual symmetry errors $\|h_1 \circ h_2 - Id\|$ and $\|h_2 \circ h_1 - Id\|$ over the images. Except in the occluded regions, these errors are very close to zero (black) everywhere, which shows that the algorithm indeed obtained a symmetric matching. The maximum symmetry error (outside of the occluded regions) is of about 1.0 for both $\|h_1 \circ h_2 - Id\|$ and $\|h_2 \circ h_1 - Id\|$ with the symmetrical method and of about 25.5 with the non-symmetric one (since the background is textureless there is no way the non-symmetric method can give symmetrical flows).




Fig. 4. The computed optic flow. On the left (resp. right), the u (resp. v) component of flow. On the top (resp. bottom) is shown h_1 (resp. h_2).




Fig. 5. The computed occlusions. On the left (resp. right), the occlusion mask computed for the first (resp. second) image. On the top (resp. bottom) is the true (resp. the computed) occlusions.

Table 1. The mean errors in norm and angle between the computed and the true flow in the non-symmetric and the symmetric cases. LR (resp. RL) means left to right (resp. right to left) flow.

| | | Euclidean norm Error | Angular Error |
|-----------------|----|----------------------|---------------|
| Non symmetrical | LR | 0.16 | 0.50 |
| | RL | 0.18 | 0.58 |
| Symmetrical | LR | 0.081 | 0.18 |
| | RL | 0.084 | 0.19 |




Fig. 6. The symmetry errors $\|h1 \circ h2 - Id\|$ (left) and $\|h2 \circ h1 - Id\|$ (right).

4.2 Real Images

The algorithm depicted above has also been run on various pairs of real images. One particularly useful pair is that provided in the paper [23] as it provides the ground truth optical flow. These images have been kindly provided by the KOGS/IAKS group of Karlsruhe University. Figures 7 and 8 show the input images and the two components of the true flow. As there is very little occlusion with these images, they were processed with the variant $\Psi = \Psi_1$.




Fig. 7. The initial input images

Figure 9 give the results obtained with this set of data both with (on the bottom) and without (on the top) the symmetrisation constraints. Three levels of resolution have been used with $\alpha = 0.3$, $\beta = 0.1$ and $\nu = 0.3$. As it can be seen, the symmetrisation constraints had several effects:

- The flow estimated on the floor is smoother.
- The computed flow on the column situated at the middle of the image is more homogeneous.
- The large error at the places marked with the letters A and C has disappeared at the cost of a much smaller one at the bottom right.
- The vertical component of the flow in the background is much more consistent with the ground truth solution (see the location marked with the letter B). In this very low-textured zone, the addition of the symmetry constraint resulted in a much better solution than the one obtained with just the regularisation constraint.




Fig. 8. The true optical flow. On the left (resp. right), the u (resp. v) component of the true flow.

In addition, figure 10 shows the two error measures: the Euclidean norm of the difference between the true and the computed flow, and the angle between those two flows as defined in [6]. Table 2 summarises the means of these two errors. Here again, errors are roughly divided by two with the symmetric method.

Table 2. The mean errors in norm and angle between the computed and the true flow.

| | Euclidean norm Error | Angular Error |
|-----------------|----------------------|---------------|
| Non symmetrical | 0.37 | 9.4 |
| Symmetrical | 0.17 | 5.2 |

Finally, figure 11 shows the the residual symmetry errors $\|h1 \circ h2 - Id\|$ and $\|h2 \circ h1 - Id\|$ over the images. Here again, these errors are very close to zero




Fig. 9. The estimated optic flow with (top) and without (bottom) the symmetry constraint. On the left (resp. right) is shown the u (resp. v) component of the computed optical flow.

everywhere so that it has been necessary to show the images with a gamma correction of 2, in order to see the residual errors. Small problems remain along the edges of the marbled floor where the texture contours have affected the symmetry constraint. The maximal symmetry errors are of about 0.8 for $\|h1 \circ h2 - Id\|$ and 1.4 for $\|h2 \circ h1 - Id\|$ with the symmetric method and respectively of 1.5 and 1.9 with the non-symmetric one.

5 Conclusion

In this paper, we have presented a PDE based optical flow estimation that preserve the discontinuity of the flow field along the image contours in such a way that it is symmetrical with respect to both images. This is done by computing simultaneously the flow from image 1 to image 2 and from image 2 to image 1, while adding explicit terms in the PDE that constrain those two flows to be compatible. The places where this compatibility cannot be achieved correspond to occlusions, so that the process allows for the detection of sufficiently large occlusions. The experimental results on synthetic and real images demonstrate the validity of the approach and the importance of the symmetry term, not only for the accuracy of the results but also as a tool to recover occluded regions.

It is important to note that the ideas developed can be generalized to many similar situations. For example, they can be very easily extended to deal with




Fig. 10. The errors on the optical flow with (top) and without (bottom) the symmetry constraint. On the left (resp. right) is shown the angular (resp. Euclidian) error of the computed optical flow.




Fig. 11. The symmetry errors $\|h1 \circ h2 - Id\|$ (left) and $\|h2 \circ h1 - Id\|$ (right). These images are displayed with a gamma correction of 4, otherwise they would have appeared totally black.

stereo images (ie when the matching is constrained to satisfy some given epipolar geometry). This is an ongoing research. More details will be given in a forthcoming research report.

Acknowledgments. This work was partially supported by the spanish research project TIC 2000-0585 (MCYT) and by the European IST project COGVISYS 3E010361.

References

1. L. Alvarez, R. Deriche, J. Weickert, and J. Sánchez. Dense disparity map estimation respecting image discontinuities: A PDE and scale-space based approach. In *IAPR International Workshop on Machine Vision Applications*, Tokyo, Japan, Nov. 2000. The University of Tokyo. A revised version has been accepted for publication in 2002 in *Journal of Visual Communication and Image Representation*.
2. L. Alvarez, P. Lions, and J. Morel. Image selective smoothing and edge detection by nonlinear diffusion (II). *SIAM Journal of Numerical Analysis*, 29:845–866, 1992.
3. L. Alvarez, J. Weickert, and J. Sánchez. A scale-space approach to nonlocal optical flow calculations. In M. Nielsen, P. Johansen, O. Olsen, and J. Weickert, editors, *Scale-Space Theories in Computer Vision*, volume 1682 of *Lecture Notes in Computer Science*, pages 235–246. Springer-Verlag, 1999.
4. L. Alvarez, J. Weickert, and J. Sánchez. Reliable estimation of dense optical flow fields with large displacements. *The International Journal of Computer Vision*, 39(1):41–56, Aug. 2000.
5. G. Aubert, R. Deriche, and P. Kornprobst. Computing optical flow via variational techniques. *SIAM Journal of Applied Mathematics*, 60(1):156–182, 1999.
6. J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *The International Journal of Computer Vision*, 12(1):43–77, 1994.
7. M. Black and P. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *The International Journal of Computer Vision*, 19(1):57–91, 1996.
8. V. Caselles, J. Morel, G. Sapiro, and A. Tannenbaum. Introduction to the special issue on partial differential equations and geometry-driven diffusion in image processing and analysis. *IEEE Transactions on Image Processing*, 7(3):269–273, 1998.
9. G. E. Christensen. Consistent linear-elastic transformations for image matching. In A. K. et al, editor, *Proceedings of IPMI*, volume 1613 of *LNCS*, pages 224–237, Berlin, Heidelberg, 1999. Springer-Verlag.
10. I. Cohen. Nonlinear variational method for optical flow computation. In *Scandinavian Conference on Image Analysis*, volume 1, pages 523–530, 1993.
11. R. Deriche, P. Kornprobst, and G. Aubert. Optical flow estimation while preserving its discontinuities: A variational approach. In *Proceedings of the 2nd Asian Conference on Computer Vision*, volume 2, pages 71–80, Singapore, Dec. 1995.
12. W. Enkelmann. Investigation of multigrid algorithms for the estimation of optical flow fields in image sequences. *Computer Vision, Graphics, and Image Processing*, 43:150–177, 1988.
13. S. Ghosal and P. Vaněk. A fast scalable algorithm for discontinuous optical flow estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(2):181–194, Feb. 1996.

14. F. Guichard and L. Rudin. Accurate estimation of discontinuous optical flow by minimizing divergence related functionals. In *Proceedings of the International Conference on Image Processing*, volume I, pages 497–500, 1996.
15. 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, Dec. 1993.
16. B. Horn and B. Schunk. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
17. B. McCane, K. Novins, D. Crannitch, and B. Galvin. On benchmarking optical flow. *Computer Vision and Image Understanding*, 84(1):126–143, 2002.
18. E. Memin and P. Perez. Multiresolution Markov random field and multigrid algorithm for discontinuity preserving estimation of the optical flow. In *Proceedings of the International Society for Optical Engineering*, San Diego, USA, July 1995.
19. H. Nagel. Constraints for the estimation of displacement vector fields from image sequences. In *International Joint Conference on Artificial Intelligence*, pages 156–160, 1983.
20. H. Nagel and W. Enkelmann. An investigation of smoothness constraint for the estimation of displacement vector fields from images sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
21. H.-H. Nagel. On the estimation of optical flow: relations between different approaches and some new results. *Artificial Intelligence Journal*, 33:299–324, 1987.
22. P. Nési. Variational approach to optical flow estimation managing discontinuities. *Image and Vision Computing*, 11(7):419–439, Sept. 1993.
23. M. Otte and H. Nagel. Optical flow estimation: Advances and comparisons. In J.-O. Eklundh, editor, *Proceedings of the 3rd European Conference on Computer Vision*, volume 800 of *Lecture Notes in Computer Science*, pages 51–70. Springer-Verlag, 1994.
24. M. Proesmans, L. Van Gool, E. Pauwels, and A. Oosterlinck. Determination of Optical Flow and its Discontinuities using Non-Linear Diffusion. In *Proceedings of the 3rd ECCV, II*, number 801 in *Lecture Notes in Computer Science*, pages 295–304. Springer-Verlag, 1994.
25. G. Sapiro. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press, 2001.
26. C. Schnörr. Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class. *The International Journal of Computer Vision*, 6(1):25–38, 1991.
27. M. Snyder. On the mathematical fundations of smoothness constraints for the determination of optical flow and for surface reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11), Nov. 1995.
28. J. Weickert. *Anisotropic Diffusion in Image Processing*. PhD thesis, University of Kaiserslautern, Germany, Laboratory of Technomathematics, Jan. 1996.
29. J. Weickert. Theoretical foundations of anisotropic diffusion in image processing. *Computing Supplement*, 11:221–236, 1996.
30. J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner-Verlag, Stuttgart, 1998.
31. J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *The International Journal of Computer Vision*, 45(3):245–264, Dec. 2001.