Efficient Mechanism for Discontinuity Preserving in Optical Flow Methods

Nelson Monzón, Javier Sánchez, and Agustín Salgado

Department of Computer Science
University of Las Palmas de Gran Canaria
35017 Las Palmas de Gran Canaria, Spain
{nmonzon,jsanchez,asalgado}@ctim.es
http://www.ctim.es

Abstract. We propose an efficient solution for preserving the motion boundaries in variational optical flow methods. This is a key problem of recent TV-L^1 methods, which typically create rounded effects at flow edges. A simple strategy to overcome this problem consists in inhibiting the smoothing at high image gradients. However, depending on the strength of the mitigating function, this solution may derive in an *ill-posed* formulation. Therefore, this type of approaches is prone to produce instabilities in the estimation of the flow fields. In this work, we modify this strategy to avoid this inconvenience. Then, we show that it provides very good results with the advantage that it yields an unconditionally stable scheme. In the experimental results, we present a detailed study and comparison between the different alternatives.

Keywords: Optical Flow, Motion Estimation, TV-L^1, Variational Method, Discontinuity-preserving

1 Introduction

One of the main problems in variational optical flow methods is the preservation of flow discontinuities. Typically, the solution in these methods is obtained as the minimization of a continuous functional, which makes it difficult to separate different moving regions. In particular, TV-L^1 methods are successful in creating piecewise-smooth motion fields. However, these approaches generate rounded shapes near the borders of the objects. In order to avoid these problems, some methods have introduced decreasing functions in order to stop the diffusion at image boundaries. This idea originally comes from [1] and is often used in many recent methods, such as in [13] or [14].

The most important problem of these inhomogeneous diffusion schemes is that they easily produce instabilities in the computed flow fields. Depending on the value of the parameters, the method may turn ill-posed. We can observe this situation in Fig. 1: depending on the parameters, we may obtain smooth solutions, similar to the Brox *et al.* method [5], or solutions with well-preserved discontinuities. In the last image, instabilities appear in the form of blobs with

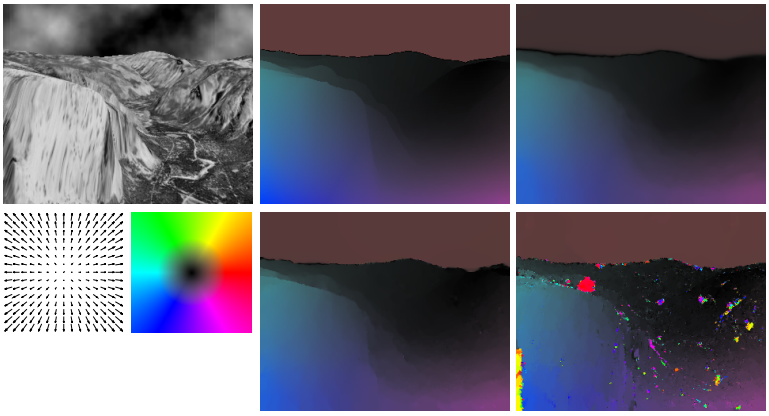

Fig. 1. Instability problem. Top row: the *Yosemite with clouds* sequence; the ground truth; and the solution obtained with the method of Brox *et al.* Bottom row: the color scheme used to represent the motion fields; example of well-preserved discontinuities in the optical flow; and instabilities due to wrong parameter setting.

large flow values. Therefore, the problem is to determine the value of the parameter that yields accurate results without introducing instabilities.

The decreasing function normally depends on the gradient of the image and a parameter that determines its decay rate. This parameter should be chosen carefully in order to avoid instabilities. Many state-of-the-art methods use a default value, which is typically very conservative in practice. In [9], the authors analyse this strategy and show that it provides promising results, but it is difficult to guess the correct configuration.

The aim of this work is to overcome these drawbacks. We propose a simple and efficient mechanism to avoid the ill-posed problem. We attain this by introducing a small constant that assures a minimum isotropic diffusion. This simple strategy turns the method very stable at the same time that it preserves the good features of the former approach.

In the experimental results, we analyse and compare both strategies. We show that this strategy outperforms the basic scheme. It provides similar results when the parameter is correctly chosen and still remains stable for a large range of values. Thus, we obtain a more reliable method that also allows preserving the flow discontinuities.

Related Work: Since the seminal work of Horn and Schunck [7], many works have dealt with the problem of discontinuities in the flow field. This method hardly preserves the motion boundaries, as shown in the online work [8]. One of the former approaches is due to Nagel and Enkelman [10]. In this case, the regularization process is steered by a diffusion tensor that depends on the image gradient. It diffuses anisotropically at image contours and isotropically in homogeneous regions. Black and Anandan [4] introduced robust functionals in the regularization term, which showed to produce piecewise motion fields. In a

similar way, Cohen [6] proposed to use a TV scheme in the regularization strategy, producing similar results. The method by Alvarez *et al.* [1] introduced a decreasing function to inhibit the smoothing at image contours. The generalization in the use of L^1 functionals was proposed in [5] and [15]. These two methods have been analysed in [11] and [12], respectively. The idea of using decreasing functions in TV-L^1 approaches is simple and is often used in many recent works, like in [13] or [14].

In Sect. 2, we explain our optical flow model. Then, in Sect. 3, we minimize the energy functional and explain the numerical details for its implementation. The experimental results, in Sect. 4, show the performance of the different methods and the benefits of the new proposal. Finally, the conclusions in Sect. 5.

2 Optical Flow Model

Given two images in a sequence, $I_1, I_2 : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$, the optical flow, $\mathbf{w} = (u(\mathbf{x}), v(\mathbf{x}))^T$, establishes the correspondences between the pixels of both images, with $\mathbf{x} = (x, y)^T \in \Omega$. Our energy model relies on the Brox *et al.* model [5] and reads as:

$$\begin{aligned} E(\mathbf{w}) = & \int_{\Omega} \Psi\left(\left(I_2(\mathbf{x}+\mathbf{w}) - I_1(\mathbf{x})\right)^2\right) \mathbf{dx} \\ & + \gamma \int_{\Omega} \Psi\left(\left|\nabla I_2(\mathbf{x}+\mathbf{w}) - \nabla I_1(\mathbf{x})\right|^2\right) \mathbf{dx} \\ & + \alpha \int_{\Omega} \Phi\left(\nabla I_1, \nabla \mathbf{w}\right) \mathbf{dx}, \end{aligned} \tag{1}$$

with $\Psi(s^2) = \sqrt{s^2 + \epsilon^2}$ and $\epsilon := 0.001$ a small constant. The behavior of the smoothing strategy depends on $\Phi(\cdot)$. Some typical examples are the following: Horn and Schunck [7], $\Phi\left(\nabla I_1, \nabla \mathbf{w}\right) = |\nabla u|^2 + |\nabla v|^2$; Alvarez *et al.* [1], $\Phi\left(\nabla I_1, \nabla \mathbf{w}\right) = f(\nabla I_1)\left(|\nabla u|^2 + |\nabla v|^2\right)$; Brox *et al.*, $\Phi\left(\nabla I_1, \nabla \mathbf{w}\right) = \Psi(|\nabla u|^2 + |\nabla v|^2)$; or Xu *et al.* [14], $\Phi\left(\nabla I_1, \nabla \mathbf{w}\right) = f(\nabla I_1)\left(|\nabla u| + |\nabla v|\right)$.

$f(\cdot)$ is a decreasing function that inhibits the regularization at object contours. Some alternatives are

$$f\left(\nabla I_1\right) = e^{-\lambda |\nabla I_1|^{\kappa}}, \; f\left(\nabla I_1\right) = \frac{1}{1 + \lambda \left|\nabla I_1\right|^2}. \tag{2}$$

We will be using the exponential function in our experiments. In [9], the authors analyse its behavior with respect to λ and κ. After their experimental results, we may conclude that $\kappa := 1$ is a good compromise between stability and accuracy.

Using the continuous L^1 functional, our smoothing function can be expressed as $\Phi\left(\nabla I_1, \nabla \mathbf{w}\right) = \Psi\left(f(\nabla I_1)\left(|\nabla u|^2 + |\nabla v|^2\right) + \epsilon^2\right)$. The problem with this functional is that it easily produces instabilities. This problem arises because

$f(\cdot)$ vanishes for large values of the gradient of the image, cancelling the regularization term. A simple and efficient mechanism, to overcome this problem, is to introduce a small constant, β, in order to avoid this cancellation. This can be easily achieved with $f(\nabla I_1) = e^{-\lambda|\nabla I_1|} + \beta$, as

$$\Phi\left(\nabla I_1, \nabla \mathbf{w}\right) = \Psi\left(\left(e^{-\lambda|\nabla I_1|} + \beta\right)\left(|\nabla u|^2 + |\nabla v|^2\right)\right), \tag{3}$$

with $\beta := 0.0001$. A similar strategy has been used in [2], but the authors use separate derivatives for each component of the optical flow.

3 Minimizing the Energy Functional

The minimum of the energy functional (1) can be found by solving the associated Euler-Lagrange equations, which are given by

$$\begin{aligned}
0 = &\Psi'_D \cdot (I_2(\mathbf{x}+\mathbf{w}) - I_1(\mathbf{x})) \cdot I_{2,x}(\mathbf{x}+\mathbf{w}) \\
&+ \gamma\, \Psi'_G \cdot ((I_{2,x}(\mathbf{x}+\mathbf{w}) - I_{1,x}(\mathbf{x})) \cdot I_{2,xx}(\mathbf{x}+\mathbf{w}) \\
&+ (I_{2,y}(\mathbf{x}+\mathbf{w}) - I_{1,y}(\mathbf{x})) \cdot I_{2,xy}(\mathbf{x}+\mathbf{w})) \\
&- \alpha\, \mathrm{div}\left(\Phi' \cdot f(\nabla I_1)\nabla u\right), \\
0 = &\Psi'_D \cdot (I_2(\mathbf{x}+\mathbf{w}) - I_1(\mathbf{x})) \cdot I_{2,y}(\mathbf{x}+\mathbf{w}) \\
&+ \gamma\, \Psi'_G \cdot ((I_{2,x}(\mathbf{x}+\mathbf{w}) - I_{1,x}(\mathbf{x})) \cdot I_{2,xy}(\mathbf{x}+\mathbf{w}) \\
&+ (I_{2,y}(\mathbf{x}+\mathbf{w}) - I_{1,y}(\mathbf{x})) \cdot I_{2,yy}(\mathbf{x}+\mathbf{w})) \\
&- \alpha\, \mathrm{div}(\Psi'_S \cdot f(\nabla I_1)\nabla v),
\end{aligned} \tag{4}$$

with $\Psi'_D := \Psi'\left((I_2(\mathbf{x}+\mathbf{w}) - I_1(\mathbf{x}))^2\right)$, $\Psi'_G := \Psi'\left(|\nabla I_2(\mathbf{x}+\mathbf{w}) - \nabla I_1(\mathbf{x})|^2\right)$ and $\Psi'_S := \Psi'\left(f(\nabla I_1)\left(|\nabla u|^2 + |\nabla v|^2\right)\right)$. The partial derivatives of the images are denoted by subscripts.

In order to solve this system, we discretize the equations using centered finite differences. Then, the system of equations is solved by means of an iterative approximation, such as the SOR method. Due to the nonlinear nature of these formulas, the resolution of these equations requires two fixed point iterations, in order to converge to a steady state. The warping of I_2 is approximated using Taylor series and bicubic interpolation.

These equations are embedded in a multiscale strategy that allows recovering large displacements. Starting from the coarsest scale, we obtain a solution to the above system, and then this solution is progressively refined in the finer scales. Details on the discretization of this scheme are given in [5] or, more extensively, in [12].

4 Experimental Results

In this section we compare the results of the Brox *et al.* method, $f\left(\nabla I_1\right) = 1$, the basic exponential function, $f\left(\nabla I_1\right) = e^{-\lambda|\nabla I_1|}$, and the new proposal, $f\left(\nabla I_1\right) = e^{-\lambda|\nabla I_1|} + \beta$. The parameters of these experiments are set according to [12].

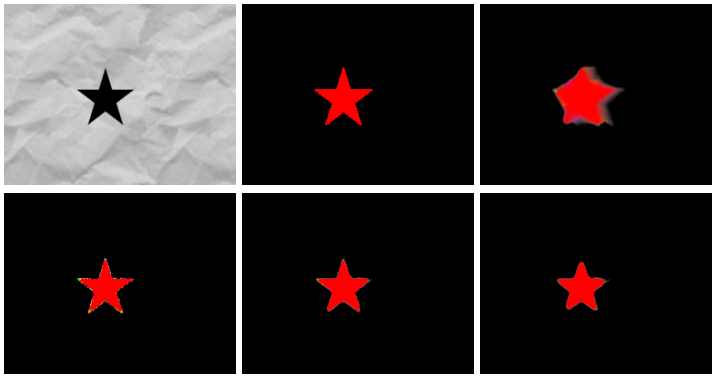

Fig. 2. *Star* sequence. First row: the original image, the ground truth and the optical flow obtained with the Brox *et al.* method. Second row: results for $f(\nabla I_1) = e^{-\lambda|\nabla I_1|}$ and $f(\nabla I_1) = e^{-\lambda|\nabla I_1|} + \beta$, with $\beta := 0.00001$ and $\beta := 0.001$, respectively.

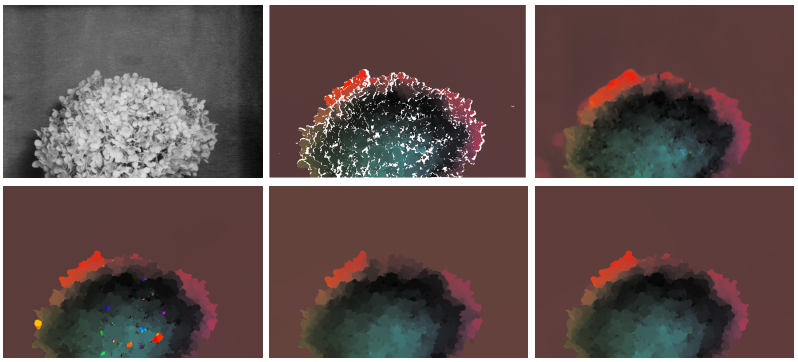

Fig. 3. *Hydrangea* sequence. First row: the original image, the ground truth and the solution of the Brox *et al.* method. Second row: results for $f(\nabla I_1) = e^{-\lambda|\nabla I_1|}$ and $f(\nabla I_1) = e^{-\lambda|\nabla I_1|} + \beta$, with $\beta := 0.00001$ and $\beta := 0.0001$.

Figure 2 shows an example of a black star that moves fifteen pixels horizontally. The Brox *et al.* method cannot completely stop the diffusion at discontinuities. In fact, we can see that it has many difficulties to deal with this type of geometric shapes. The solution of the exponential method is much better, especially at motion boundaries. However, if the parameter is not correctly chosen, some instabilities appear at the star contours. This problem disappears when we use $f(\nabla I_1) = e^{-\lambda|\nabla I_1|} + \beta$. The small constant avoids instabilities at the same time that it preserves discontinuities.

In Figs. 3 and 4, we show the results for *Hydrangea* and *Grove2* sequences, from the Middlebury benchmark database [3]. We can observe that the results of the basic exponential method are promising, since the preservation of discon-

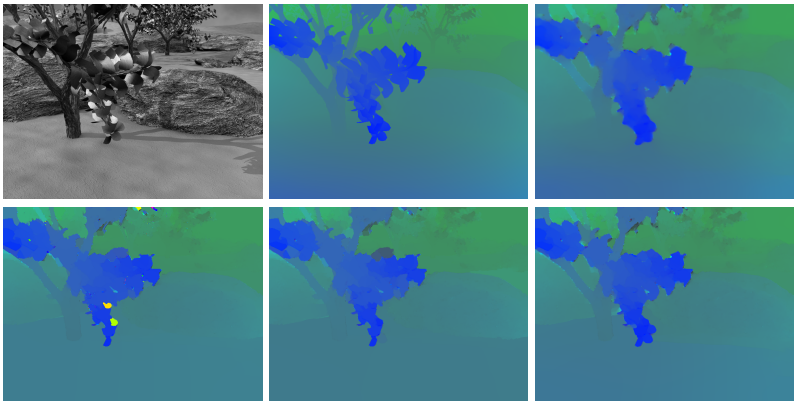

Fig. 4. *Grove2* sequence. First row: the original image, the ground truth and the solution of the Brox *et al.* method. Second row: results for $f(\nabla I_1) = e^{-\lambda|\nabla I_1|}$ and $f(\nabla I_1) = e^{-\lambda|\nabla I_1|} + \beta$, with $\beta := 0.00001$ and $\beta := 0.0001$.

tinuities is accurate. Nevertheless, the number of outliers is important. These are removed using the β constant.

The decreasing function allows segmenting the shape of the geometric sequence from the background motion. It provides very good results for both convex and concave shapes in general. However, we appreciate instabilities at the contour of the geometric images and in many small places of the Middlebury dataset. In this method, it is difficult to find an optimal value for λ. Using the new proposal effectively eliminates the instabilities.

In Fig. 5, we compare the Average Angular Error (AAE) evolution for both approaches. The two variants significantly improve the outcome of the Brox *et al.* method ($\lambda := 0$). The evolution is very similar at the beginning in both methods. However, when λ increases, the pure exponential method becomes unstable. In contrast, the other approach have a smoother evolution, yielding very good results for larger values of λ. We even observe that using the small isotropic constant provides slightly better results for the *Star* and *Grove2* sequences.

Table 1 shows the best AAE results for the *Star*, *Yosemite with clouds* and Middlebury sequences. We observe an important improvement in the *Star* and *Yosemite* sequences with respect to the Brox *et al.* method. The results for the Middlebury sequences also improve. The solutions of both exponential methods are similar because the minimum error is usually given for the same configuration (see the graphics). Nevertheless, the second alternative does not produce instabilities and is more stable for a large range of λ values.

5 Conclusion

We proposed an efficient strategy for preserving the motion boundaries in variational optical flow methods. The use of decreasing functions with TV approaches

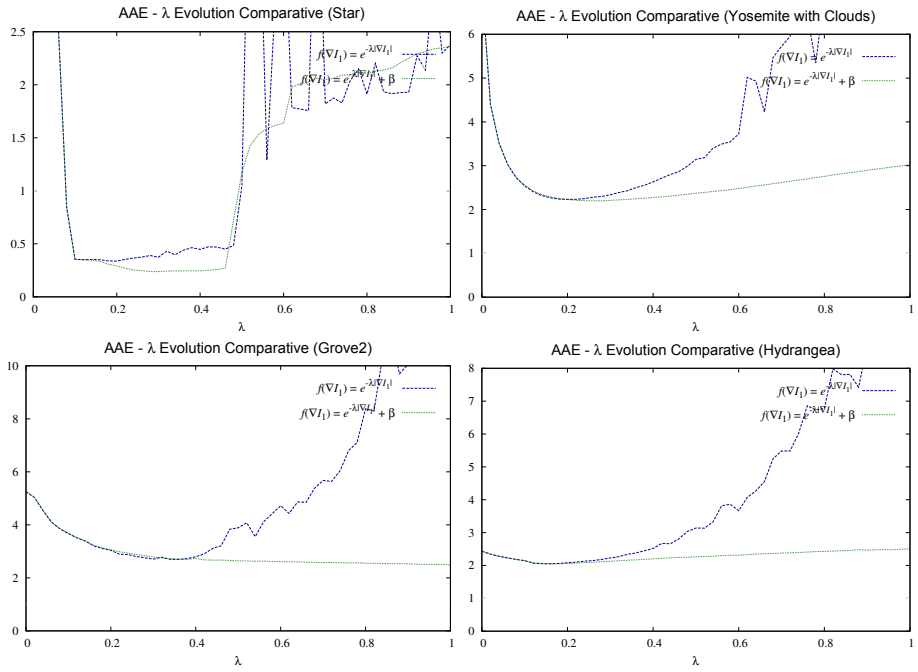

Fig. 5. AAE evolution for the *Star* and *Yosemite* sequences (top row), and *Grove2* and *Hydrangea* (bottom row).

Table 1. AAE results for the *Star*, *Yosemite with clouds* and Middlebury sequences.

Sequence	**Brox *et al.***	$f(\nabla I_1) = e^{-\lambda\lvert\nabla I_1\rvert}$	$f(\nabla I_1) = e^{-\lambda\lvert\nabla I_1\rvert} + \beta$
Star	2.919^o	0.271^o	0.222^o
Yosemite with clouds	2.367^o	1.977^o	1.976^o
Hydrangea	2.076^o	2.035^o	2.027^o
Grove2	2.198^o	2.109^o	2.111^o

allows mitigating the diffusion at contours. We have shown that the instability problems are effectively removed with our proposal at the same time that we obtain very good accuracy at motion boundaries. Another advantage of these strategies is that they are very easy to implement from the basic Brox *et al.* method. Furthermore, they provide much better results for simple sequences, such as the geometric test images, where the dominant gradient clearly separates the different motions. Eliminating the instability problems turns this method very interesting for real applications.

Acknowledgments. This work has been partially supported by the Spanish Ministry of Science and Innovation through the research project TIN2011-25488 and the University of Las Palmas de Gran Canaria grant ULPGC011-006.

References

1. Luis Álvarez, Julio Esclarín, Martin Lefébure, and Javier Sánchez. A PDE model for computing the optical flow. In *XVI Congreso de Ecuaciones Diferenciales y Aplicaciones, C.E.D.Y.A. XVI*, pages 1349–1356, Las Palmas de Gran Canaria, Spain, 1999.
2. Alper Ayvaci, Michalis Raptis, and Stefano Soatto. Sparse occlusion detection with optical flow. *International Journal of Computer Vision*, 97(3):322–338, May 2012.
3. Simon Baker, Daniel Scharstein, J. P. Lewis, Stefan Roth, Michael J. Black, and Richard Szeliski. A database and evaluation methodology for optical flow. In *International Conference on Computer Vision*, pages 1–8, 2007.
4. Michael J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding*, 63(1):75 – 104, 1996.
5. T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In T. Pajdla and J. Matas, editors, *European Conference on Computer Vision (ECCV)*, volume 3024 of *Lecture Notes in Computer Science*, pages 25–36, Prague, Czech Republic, May 2004. Springer.
6. Isaac Cohen. Nonlinear Variational Method for Optical Flow Computation. In *Proceedings of the 8th Scandinavian Conference on Image Analysis*, pages 523–530, Tromso, Norway, 1993. IAPR.
7. Berthold K. P. Horn and Brian G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
8. Enric Meinhardt-Llopis, Javier Sánchez, and Daniel Kondermann. Horn-Schunck Optical Flow with a Multi-Scale Strategy. *Image Processing On Line*, 2013:151–172, 2013. `http://dx.doi.org/10.5201/ipol.2013.20`.
9. Nelson Monzón, Javier Sánchez, and Agustín Salgado. Optic flow: Improving discontinuity preserving. In *EUROCAST'13: Proceedings of the 14th international conference on Computer aided systems theory*, pages 114–116, Berlin, Heidelberg, 2013. Springer-Verlag.
10. H H Nagel and W Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transanctions on Pattern Analysis and Machine Intelligence*, 8:565–593, September 1986.
11. Javier Sánchez, Enric Meinhardt-Llopis, and Gabriele Facciolo. TV-L1 Optical Flow Estimation. *Image Processing On Line*, 2013:137–150, 2013. `http://dx.doi.org/10.5201/ipol.2013.26`.
12. Javier Sánchez, Nelson Monzón, and Agustín Salgado. Robust Optical Flow Estimation. *Image Processing On Line*, 2013:242–260, 2013. `http://dx.doi.org/10.5201/ipol.2013.21`.
13. Andreas Wedel, Daniel Cremers, Thomas Pock, and Horst Bischof. Structure- and motion-adaptive regularization for high accuracy optic flow. In *Proceedings of IEEE International Conference on Computer Vision*, pages 1663–1668, September 2009.
14. Li Xu, Jiaya Jia, and Yasuyuki Matsushita. Motion Detail Preserving Optical Flow Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1293–1300, June 2010.
15. C. Zach, T. Pock, and H. Bischof. A Duality Based Approach for Realtime TV-L1 Optical Flow. In Fred A. Hamprecht, Christoph Schnörr, and Bernd Jähne, editors, *Pattern Recognition*, volume 4713 of *Lecture Notes in Computer Science*, chapter 22, pages 214–223. Springer, Berlin, Heidelberg, 2007.