



CTIM
CENTRO DE TECNOLOGÍAS
DE LA IMAGEN

CTIM Technical Report

ISSN 2254-2353

**Discontinuity Preserving in TV-L1
Optical Flow Methods**

*Nelson Monzón López, Agustín J. Salgado de
la Nuez and Javier Sánchez Pérez*

No. 5

Las Palmas de Gran Canaria
04 November 2014

Discontinuity Preserving in TV- L^1 Optical Flow Methods

Nelson Monzón, Agustín Salgado and Javier Sánchez

Centro de Tecnologías de la Imagen
Departamento de Informática y Sistemas
Universidad de Las Palmas de Gran Canaria
35017 Las Palmas de Gran Canaria, Spain

Abstract

We analyze the discontinuity preserving problem in TV- L^1 optical flow methods. This type of methods typically creates rounded effects at flow boundaries, which usually do not coincide with object contours. A simple strategy to overcome this problem consists in inhibiting the diffusion at high image gradients. In this work, we first introduce a general framework for TV regularizers in optical flow and relate it with some standard approaches. Our survey takes into account several methods that use decreasing functions for mitigating the diffusion at image contours. Consequently, this kind of strategies may produce instabilities in the estimation of the optical flows. Hence, we study the problem of instabilities and show that it actually arises from an *ill-posed* formulation. From this study, it is possible to come across with different schemes to solve this problem. One of these consists in separating the pure TV process from the mitigating strategy. This has been used in another work and we demonstrate here that it has a good performance. Furthermore, we propose two alternatives to avoid the instability problems: (i) we study a fully automatic approach that solves the problem based on the information of the whole image; (ii) we derive a semi-automatic approach that takes into account the image gradients in a close neighborhood adapting the parameter in each position. In the experimental results, we present a detailed study and comparison between the different alternatives. These methods provide very good results, especially for sequences with a few dominant gradients. Additionally, a surprising effect of these approaches is that they can cope with occlusions. This can be easily achieved by using strong regularizations and high penalizations at image contours.

Keywords: Optical Flow, Total Variation, TV- L^1 , Variational Method, Diffusion Process.

1 Introduction

The estimation of accurate motion fields is one of the fundamental challenges in computer vision. Commonly known as optical flow, it consists in estimating the apparent displacement of the pixels through an image sequence. Although there are many strategies to calculate the flow, variational methods are among the most accurate techniques.

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. Flow boundaries are normally associated with the contours of the objects in the scene. However, the converse is not true, since adjacent objects, moving in the same direction, may belong to the same optical flow region, with no motion discontinuities between them. Introducing a simple mechanism for inducing the boundaries of the flow field from the objects is not easy. Moreover, the problem becomes more challenging since it is difficult to differentiate between object contours and textures, if we rely on the information of the image gradients.

Many ideas for preserving the optical flow contours come from the field of image denoising and regularization. For instance, Perona and Malik [21] proposed an anisotropic model for

regularizing an image by preserving its discontinuities. This method is based on decreasing functions that inhibit the diffusion at high image gradients. Rudin, Osher and Sethian [23], on the other hand, proposed to minimize the Total Variation (TV) of an image with an attachment to the original image. This leads to a diffusion equation that reduces the image noise, yielding sharp edges. Some improvements on the Perona and Malik model are given in Black *et al.* [10]. This establishes the relation between this kind of anisotropic diffusion processes, robust statistics and the minimization of energy functionals. It generalizes the use of robust functionals in order to deal with outliers. A review on different strategies for diffusion filtering in image regularization and restoration can be seen in [36, 37]. The author introduces the theory underlying the use of diffusion tensors in image filtering, e.g., the structure tensor [7]. Another source of inspiration for discontinuity preserving in optical flow is related with the bilateral filtering, introduced by Tomasi and Manduchi [32]. In this case, the idea is to regularize an image using the information of the pixels that are near the actual position and have similar intensities or colors.

Since the seminal work of Horn and Schunck [16], many works have appeared dealing with the problem of discontinuities. One of the former approaches is due to Nagel and Enkelmann [20]. In this case, the regularization process is steered by a diffusion tensor that depends on the image gradient: in homogeneous regions, it smoothes isotropically, like in [16]; at object boundaries, where the gradient of the image is high, the regularization is mainly carried out along the image contours. Proesmans *et al.* [22] introduced the Perona and Malik [21] anisotropic scheme in the estimation of optical flow. They also introduced a symmetric coherence model that helps to detect the discontinuities of the flow. In a similar way, Black and Anandan [8, 9] used this type of anisotropic regularization. They established the relation between robust statistics and anisotropic diffusion. They also extended the use of robust functionals to the whole energy terms, turning the method more robust against outliers. They showed that this strategy deals with image noise at the same time that it preserves flow edges. On another hand, a TV scheme, like in [23], was introduced in optical flow by Cohen [15].

The method by Álvarez *et al.* [2] introduces a decreasing function to inhibit the smoothing at image contours. Nevertheless, they did not use any robustification function in the data term, so it is more sensitive to image noise. Aubert *et al.* [4] explicitly propose to use an L^1 functional in the data term and any robustification function for smoothing. Alvarez *et al.* [3] uses the Nagel-Enkelmann diffusion tensor, together with a nonlinear brightness formulation and a linear scale-space for the estimation of large displacements.

The generalization in the use of continuous L^1 functionals was proposed in [11, 14] and subsequent works. In fact, this has already been proposed before, e.g., in [4], but the former introduced a term based on the attachment of the image gradients, which is invariant to constant brightness changes. On the other hand, non-continuous L^1 functionals have also been used in Zach *et al.* [45] that relies on a dual formulation, which yields a very efficient numerical scheme. The work in [35] increases the robustness to illumination changes using the textural part of the images, somehow similar to the gradient term of Brox *et al.* [11]. Although the former two approaches are similar, they provide much different results, as can be seen in the online works [31] and [30], respectively.

Some examples in the use of diffusion tensors with robustification functions are given in [47, 46]. In this case, the authors introduce a motion tensor in the data term and a regularization tensor in the smoothness term, which are designed in a similar way, taking into account not only the variation of image intensities but also the variation of the image gradients. The latter tensor uses a quadratic penaliser for the diffusion along the contours, while a Perona-Malik diffusivity is used for mitigating the diffusion across flow edges.

TV- L^1 methods have several drawbacks: (i) they create rounded shapes near the borders or corners of the objects; (ii) typically, the edges are dislocated and usually do not coincide with the image contours; (iii) they produce staircase effects, yielding piecewise but planar

motion regions. The first two inconveniences are due to the fact that the regularization process does not depend on the image information but on the flow field. In order to avoid these, some methods have introduced decreasing functions in order to stop the diffusion at image boundaries. This idea originally comes from [2] and has recently been used in several methods, such as in [34, 43]. It has also been used in Werlberger *et al.* [40], where a diffusion tensor steers the regularization in the direction of the image gradient and its orthogonal direction. It incorporates a decreasing function in the direction of the gradient, so it inhibits the smoothing across edges.

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 become ill-posed if the smoothing term is canceled. Most of the aforementioned strategies assign an empirical value to these parameters. Unfortunately, the parameters that better preserve discontinuities are those that risk to produce instabilities. One way to avoid the ill-posed problem is to introduce a small constant that assures a minimum isotropic behavior, like in Monzón *et al.* [19] or Ayvaci *et al.* [5]. However, the value of this constant depends on the regularization factor, so that if this is large, then discontinuities will not be respected. In this work we show that this strategy outperforms the basic approach, but it is possible to find better solutions.

The idea of bilateral filtering has been introduced in optical flow by Yoon and Kweon [44] and Xiao *et al.* [42]. In this case, the authors propose to regularize the flow field depending on the proximity and similarity of the intensities and flow values. In fact, this has to be seen as an extended trilinear filtering. They also used the information of occlusions to manipulate the range of the filters. Bilateral filtering has been used more recently, in combination with a TV- L^1 approach, in [39].

There exist other strategies for improving the definition of discontinuities. For instance, some authors propose to compute the optical flow at the same time that the objects in the scene are segmented. In this way, the segmentation provides more information about the edges. Some examples in this line are [17, 28, 29, 33]. In [27], the authors combine median filtering and bilateral filtering, in a post-processing step, to improve the flow field at edges. Other approaches mix TV- L^1 strategies with descriptor matching, such as [12, 43], which estimate the motion of some sparse features that are later introduced in the optimization process. This allows to deal with small moving objects. Moreover, since some features are typically associated with edges, this may help to define discontinuities. The methods based on segmentation, post-filtering or feature based matching, are difficult to analyze, especially concerning the problem of discontinuity preserving. These usually mix different aspects that affect the definition of flow boundaries, from standard robust regularization approaches to other *ad-hoc* processes.

The aim of this work is to analyze the behavior of TV- L^1 optical flow methods. In particular, we study the effect of these methods at discontinuities. We compare three schemes based on Brox *et al.* [11]. These schemes rely on regularization strategies that use decreasing functions to inhibit the smoothing at high image gradients. This is a simple solution that solves the main drawbacks of TV- L^1 methods, i.e., the problems of rounded and dislocated contours. At the same time, it preserves the advantages, such as creating piecewise and continuous motion regions, dealing with outliers and textured areas. Furthermore, this strategy allows to increase the smoothing strength, so we can obtain more continuous flows inside the regions. This work is a continuation of our previous conference articles Monzón *et al.* [18, 19] and Sánchez *et al.* [25], where we achieved very promising results.

The preservation of discontinuities, using the above scheme, depends on three main factors: (i) the regularization parameter, that balances the weight of smoothing with respect to the attachment to the image data; (ii) the gradient of the flow; and, (iii) the gradient of the image. The bigger the regularization parameter, the smoother the solution we obtain. On the contrary, with strong gradients of the image or the flow, the diffusion is mitigated.

Instabilities usually appear when the regularization is too small or when the gradient of the image is high.

The decreasing function depends on the gradient of the image and an additional parameter that determines its decay rate. This parameter should be chosen carefully in order to avoid instabilities. Many state-of-the-art methods make use of a default value, which is typically very conservative in practice. In this work we analyze the influence of this parameter and show that, only in simple situations, a default value is worthwhile. We propose two mechanisms for automatically determining its value, which in general provide better results. These approaches eliminate the effect of instabilities, efficiently solving the ill-posed problem.

On the one hand, we study a global approach, which computes an optimal value that avoids stopping the diffusion as seen in [25]. It yields a unique value for the whole image. In the experiments we will see that this approach attains the best solutions for simple sequences, but its results are poorer when there is a broad variety of gradients. On the other hand, we propose a semi-automatic approach that is better adapted to local variations of the gradient. In this case, the method adapts the input value in order to avoid instabilities: whenever it finds a possible problematic situation, it adapts the parameter to allow a minimum diffusion. In this way, we can choose a more restrictive value for the parameter.

In this work, we also show that this type of methods can easily deal with occlusions. This can be achieved by increasing the value of the regularization parameter, so that the smoothing outweighs the attachment to the intensity values. However, such situations tend to create planar motions.

In Sect. 2, we introduce a general framework for this kind of strategies. Then, we analyze the problem of instabilities in Sect. 3. After this analysis, we deduce the global and local alternatives in Sect. 4. The experimental results, in Sect. 5, deepens in the study of the different methods and the new proposals. Finally, a summary of the main ideas and conclusions in Sect. 6.

2 Theoretical Framework

Given two images in a sequence, $I_1, I_2 : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$, of gray level values, the optical flow, $\mathbf{w} = (u(\mathbf{x}), v(\mathbf{x}))^T$, puts in correspondence the pixels of the first image with the pixels of the second, with $\mathbf{x} = (x, y)^T \in \Omega$. Functions $u(\mathbf{x})$ and $v(\mathbf{x})$ are the horizontal and vertical displacements, respectively.

Variational optical flow techniques rely on a global energy functional. The minimization of this functional provides a solution to the optical flow problem between two frames. Standard approaches typically include an attachment and a regularization term, as follows:

$$E(\mathbf{w}) = \int_{\Omega} \mathcal{D}(I_1, I_2, \nabla I_1, \nabla I_2, \mathbf{w}) \, \mathbf{d}\mathbf{x} + \alpha \int_{\Omega} \mathcal{R}(\nabla I_1, \nabla u, \nabla v) \, \mathbf{d}\mathbf{x}. \quad (1)$$

Many different alternatives have been proposed in the literature for each of these terms. In our case, we use L^1 norms in a similar way as in the Brox *et al.* model [11, 13]. Then, our attachment term reads as

$$\begin{aligned} \mathcal{D}(I_1, I_2, \nabla I_1, \nabla I_2, \mathbf{w}) := & \Psi \left((I_2(\mathbf{x} + \mathbf{w}) - I_1(\mathbf{x}))^2 \right) \\ & + \gamma \Psi \left(|\nabla I_2(\mathbf{x} + \mathbf{w}) - \nabla I_1(\mathbf{x})|^2 \right), \end{aligned} \quad (2)$$

with $\Psi(s^2) = \sqrt{s^2 + \epsilon^2}$ and $\epsilon := 0.001$ a small constant. Henceforth, we will use this scheme for the attachment term. This will remain unaltered in the experimental results and we will enable or disable the second term in (2) for convenience.

Table 1: Examples of rotational invariant regularization schemes, $\mathcal{R}(\nabla I_1, \nabla u, \nabla v) := \Phi(\nabla u^T \mathbf{D}(\nabla I_1) \nabla u + \nabla v^T \mathbf{D}(\nabla I_1) \nabla v)$

Method	$\Phi(s)$	$\mathbf{D}(\nabla I_1)$
Horn-Schunck [16]	s	\mathbf{I}
Brox <i>et al.</i> [11]	$\sqrt{s + \epsilon^2}$	\mathbf{I}
Alvarez <i>et al.</i> [2]	s	$g(\nabla I_1) \mathbf{I}$
Xu <i>et al.</i> [43]	\sqrt{s}	$g(\nabla I_1) \mathbf{I}$
Monzón <i>et al.</i> [19]	$\sqrt{s + \epsilon^2}$	$(g(\nabla I_1) + \beta) \mathbf{I}$
Nagel-Enkelmann [20]	s	$\frac{\nabla I_1^\perp \nabla I_1^{\perp T} + \beta^2 \mathbf{I}}{ \nabla I_1 ^2 + 2\beta^2}$
Sánchez <i>et al.</i> [24]	$\sqrt{s + \epsilon^2}$	$\frac{\nabla I_1^\perp \nabla I_1^{\perp T} + \beta^2 \mathbf{I}}{ \nabla I_1 ^2 + 2\beta^2}$
Álvarez <i>et al.</i> [1]	s	$\frac{(C_1 g(\nabla I_1) + C_2) \nabla I_1^\perp \nabla I_1^{\perp T} + \beta^2 g(\nabla I_1) \mathbf{I}}{ \nabla I_1 ^2 + 2\beta^2}$

The regularization term is responsible for the continuity of the computed optical flow. It can be generalized using a tensor matrix as

$$\mathcal{R}(\nabla I_1, \nabla u, \nabla v) := \Phi(\nabla u^T \mathbf{D}(\nabla I_1) \nabla u + \nabla v^T \mathbf{D}(\nabla I_1) \nabla v) \quad (3)$$

or

$$\mathcal{R}(\nabla I_1, \nabla u, \nabla v) := \Phi(\nabla u^T \mathbf{D}(\nabla I_1) \nabla u) + \Phi(\nabla v^T \mathbf{D}(\nabla I_1) \nabla v), \quad (4)$$

with $\Phi(\cdot)$ a convex function and \mathbf{D} a 2×2 matrix that allows to steer the diffusion process depending on the structure of the image. The first alternative yields a rotational invariant regularizer, while the second produces a system of decoupled diffusion equations, as shown in Weickert and Schnörr [38]. In that work, the authors established a taxonomy of isotropic and anisotropic regularization schemes, steered by the image or flow information. For the sake of simplicity, we concentrate on the rotational invariant formulation, given in (3).

2.1 Regularization Strategies

The behavior of the smoothing strategy depends on the choices for Φ and \mathbf{D} . The previous generalization may be adapted to many state-of-the-art variational methods. Table 1 summarizes several rotational invariant methods that fit in the model given in (3).

In this table, \mathbf{I} stands for the identity matrix. The second choice corresponds to the Brox *et al.* spatial model [11, 13], which has been thoroughly analyzed in [31] – we use the implementation of this work, available online.

In Table 2, there is a survey of non-rotational methods that can be framed in (4). Other methods, such as the stochastic models proposed in [8, 17, 26], or bilateral filter approaches, such as [39], cannot be related with these formulations. Neither the rotational invariant methods, explained in [47, 46], can be framed in this generic framework, because they need to compute the eigenvectors from a more sophisticated regularization tensor.

$g(\cdot)$ in these tables represents a decreasing function that inhibits regularization at object contours. Some typical alternatives [21] are

$$g(\nabla I_1) = e^{-\lambda |\nabla I_1|^\kappa}, \quad g(\nabla I_1) = \frac{1}{1 + \lambda |\nabla I_1|^2}. \quad (5)$$

Table 2: Examples of non-rotational invariant regularization schemes, $\mathcal{R}(\nabla I_1, \nabla u, \nabla v) := \Phi(\nabla u^T \mathbf{D}(\nabla I_1) \nabla u) + \Phi(\nabla v^T \mathbf{D}(\nabla I_1) \nabla v)$

Method	$\Phi(s)$	$\mathbf{D}(\nabla I_1)$
Cohen [15]	\sqrt{s}	\mathbf{I}
Wedel <i>et al.</i> [34]	\sqrt{s}	$g(\nabla I_1) \mathbf{I}$
Ayvaci <i>et al.</i> [5]	\sqrt{s}	$(g(I_{1,x}) + \beta, g(I_{1,y}) + \beta) \mathbf{I}$
Werlberger <i>et al.</i> [41]	\sqrt{s}	$g(\nabla I_1) \frac{\nabla I_1 \nabla I_1^T}{ \nabla I_1 ^2} + \frac{\nabla I_1^\perp \nabla I_1^{\perp T}}{ \nabla I_1 ^2}$

We will be using the exponential in our experiments. In [18], the authors analyze 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, so we fix this parameter in the tests.

We observe that many of the methods in Tables 1 and 2 depend on decreasing functions to enhance discontinuities. One of the drawbacks of these functions is that, depending on the λ and the gradient of the image, it may cancel the smoothing term, turning the method ill-posed. Some methods avoid this problem by introducing a constant, β , that assures a minimum isotropic diffusion. This strategy has been used, for instance, in Monzón *et al.* [19] or [5]. One of the benefits of decreasing functions is that we can increase the smoothing strength, in order to obtain more continuous flows, at the same time that it respects the contours of the objects. In fact, there exists a relation between the smoothing parameter and the decreasing function, that will be analyzed below.

The Nagel-Enkelmann operator has been traditionally used in quadratic energy models, such as in [3, 20]. This operator allows anisotropic diffusion at the borders of the objects and isotropic regularization in homogeneous regions. In [24], the authors implement the operator inside a TV- L^1 functional. Unfortunately, this scheme seems to be unstable and the parameters have to be carefully selected. We can see in the tables that this operator has been used both in rotational and non-rotational methods. It has also been adapted in several works with decreasing functions to mitigate the diffusion across the boundaries of the objects.

2.2 Minimizing the Energy Functional

The minimum of the energy functional (1) can be found by solving the associated Euler-Lagrange equations. The data term is given in (2) and the smoothing term in (3). Then, the Euler-Lagrange equations are given by

$$\begin{aligned}
0 &= \Psi'_D \cdot (I_2(\mathbf{x} + \mathbf{w}) - I_1(\mathbf{x})) \cdot I_{2,x}(\mathbf{x} + \mathbf{w}) \\
&\quad + \gamma \Psi'_G \cdot ((I_{2,x}(\mathbf{x} + \mathbf{w}) - I_{1,x}(\mathbf{x})) \cdot I_{2,xx}(\mathbf{x} + \mathbf{w}) \\
&\quad + (I_{2,y}(\mathbf{x} + \mathbf{w}) - I_{1,y}(\mathbf{x})) \cdot I_{2,xy}(\mathbf{x} + \mathbf{w})) \\
&\quad - \alpha \operatorname{div}(\Phi' \cdot \mathbf{D}(\nabla I_1) \nabla u), \\
0 &= \Psi'_D \cdot (I_2(\mathbf{x} + \mathbf{w}) - I_1(\mathbf{x})) \cdot I_{2,y}(\mathbf{x} + \mathbf{w}) \\
&\quad + \gamma \Psi'_G \cdot ((I_{2,x}(\mathbf{x} + \mathbf{w}) - I_{1,x}(\mathbf{x})) \cdot I_{2,xy}(\mathbf{x} + \mathbf{w}) \\
&\quad + (I_{2,y}(\mathbf{x} + \mathbf{w}) - I_{1,y}(\mathbf{x})) \cdot I_{2,yy}(\mathbf{x} + \mathbf{w})) \\
&\quad - \alpha \operatorname{div}(\Phi' \cdot \mathbf{D}(\nabla I_1) \nabla v),
\end{aligned} \tag{6}$$

with $\Psi'(s^2) = \Phi'(s^2) = \frac{1}{2\sqrt{s^2 + \epsilon^2}}$. In order to simplify these equations, we have used the following notation:

$$\begin{aligned}\Psi'_D &:= \Psi' \left((I_2(\mathbf{x} + \mathbf{w}) - I_1(\mathbf{x}))^2 \right), \\ \Psi'_G &:= \Psi' \left(|\nabla I_2(\mathbf{x} + \mathbf{w}) - \nabla I_1(\mathbf{x})|^2 \right), \\ \Phi' &:= \Phi' \left(\nabla u^T \mathbf{D}(\nabla I_1) \nabla u + \nabla v^T \mathbf{D}(\nabla I_1) \nabla v \right).\end{aligned}\tag{7}$$

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 warpings of I_2 are approximated using Taylor expansions.

These equations are embedded in a multiscale strategy that allows to recover large displacements. Starting from the coarsest scales, we obtain a solution to the above system, and then upgrade the value of the optical flow for the next finer scale. We use *motion increments*, $\mathbf{h}^{k+1} = \mathbf{h}^k + \mathbf{d}\mathbf{h}^k$, so that, in each scale, we compute each increment, $\mathbf{d}\mathbf{h}^k$, and the final optical flow is obtained as an accumulative value for all increments.

Details on the discretization of this scheme are given in [11, 13] or, more extensively, in [31]. The main difference, with respect to the basic model, is the implementation of the matrix-value diffusion process.

3 Ill-posedness of the Anisotropic Model

Hereafter, we assume that $\mathbf{D}(\nabla I_1) := g(\nabla I_1) \mathbf{Id} = e^{-\lambda|\nabla I_1|} \mathbf{Id}$. Expanding the smoothing term in (6), we obtain the following relations

$$\begin{aligned}S_u(\mathbf{w}) &= \alpha \operatorname{div}(\Phi' \cdot \mathbf{D}(\nabla I_1) \nabla u) = \alpha \operatorname{div}(\Phi' \cdot e^{-\lambda|\nabla I_1|} \nabla u), \\ S_v(\mathbf{w}) &= \alpha \operatorname{div}(\Phi' \cdot \mathbf{D}(\nabla I_1) \nabla v) = \alpha \operatorname{div}(\Phi' \cdot e^{-\lambda|\nabla I_1|} \nabla v).\end{aligned}\tag{8}$$

The diffusivity, including α , is given by

$$\mathcal{D}(\alpha, \lambda, \nabla I_1, \nabla \mathbf{w}) = \alpha \cdot \Phi' \cdot e^{-\lambda|\nabla I_1|} = \frac{\alpha e^{-\lambda|\nabla I_1|}}{\sqrt{e^{-\lambda|\nabla I_1|} (|\nabla u|^2 + |\nabla v|^2) + \epsilon^2}}.\tag{9}$$

The ill-posed problem appears when this diffusivity becomes zero. This is equivalent to cancel the regularization term in the energy functional (1). Therefore, this problem depends on the values of α , λ and ∇I_1 . For example, this term nullifies if λ or ∇I_1 are very large, or when α is too small.

In order to avoid these problems, we may impose a lower bound on the diffusivity, so that, in the worse case, $\mathcal{D}(\mathbf{x}) := \xi$, with ξ a small constant, e.g., $\xi := 1 \times 10^{-4}$,

$$\frac{\alpha e^{-\lambda|\nabla I_1|}}{\sqrt{e^{-\lambda|\nabla I_1|} (|\nabla u|^2 + |\nabla v|^2) + \epsilon^2}} = \xi.\tag{10}$$

Solving the previous equation, we obtain

$$\lambda_{1,2} = -\ln \left(\xi \frac{\left(|\nabla u|^2 + |\nabla v|^2 \right) \xi \pm \sqrt{\left(|\nabla u|^2 + |\nabla v|^2 \right)^2 \xi^2 + 4\alpha^2 \epsilon^2}}{2\alpha^2 |\nabla I_1|} \right).\tag{11}$$

This equation yields two solutions, one real and another imaginary, which can be discarded. If we analyze further, we find that the denominator of (10) is bounded in the limits

$$0 < \sqrt{e^{-\lambda|\nabla I_1|}(|\nabla u|^2 + |\nabla v|^2) + \epsilon^2} \leq \sqrt{|\nabla u|^2 + |\nabla v|^2 + \epsilon^2} \approx |\nabla \mathbf{w}|, \quad (12)$$

with the upper bound approximately equal to the flow magnitude. In this sense, the denominator does not pose any problem because it will never be very large: in practical situations, the magnitude of the flow will be much smaller than the image size.

The main problem arrives when $\alpha e^{-\lambda|\nabla I_1|} \cong 0$. One solution, to avoid this situation, is to impose the following constraint

$$\alpha e^{-\lambda|\nabla I_1|} = \xi. \quad (13)$$

This equation only depends on the input data ∇I_1 and α . Thus, we may adapt λ at the beginning of the process, to assure this condition. If we solve the above equation, then we obtain

$$\lambda = \frac{-\ln(\xi) + \ln(\alpha)}{|\nabla I_1|}. \quad (14)$$

ξ must be chosen so that the diffusivity is big enough to avoid the ill-posedness. This represents the smallest value that starts the regularization. In the experimental results, we analyze the behavior of this constant and heuristically find the most appropriate value.

Notice that both conditions, (11) and (14), successfully remove the ill-posed problem, but the second is much easier to implement.

4 Determining the value of λ

Above we deduced the value of λ that avoids stopping the diffusion process. For this, we introduced the ξ constant, which must be big enough so that there always exists a smoothing process. If ξ is large, λ may turn negative, and the regularization becomes isotropic, similar to Brox *et al.* [11].

In this section, we derive two schemes to automatically avoid the ill-posed problem. The idea is to adapt the value of λ when we detect a possible instability. Taking into account (11) and (14), this problem can be easily detected by checking

$$\left(\frac{\alpha e^{-\lambda|\nabla I_1|}}{\sqrt{e^{-\lambda|\nabla I_1|}(|\nabla u|^2 + |\nabla v|^2) + \epsilon^2}} < \xi \right) \text{ or } \left(\alpha e^{-\lambda|\nabla I_1|} < \xi \right), \quad (15)$$

if we choose (11) or (14), respectively.

We may think of two possible alternatives: on the one hand, we may impose this condition in the whole image, so that we determine the value of λ that better fits to the more restrictive situation, i.e., that corresponding to the maximum gradient; on the other hand, we may assure this condition in every pixel, looking at the gradient in each position.

Based on this reasoning, and selecting the second condition in (15) for simplicity, we propose two schemes for determining the value of λ :

Global approach. The first is based on a *global* approach seen in [25]. We calculate λ as a constant using the maximum gradient of the image as

$$\lambda_g := \frac{-\ln(\xi) + \ln(\alpha)}{\max_{\mathbf{x} \in \Omega} \{|\nabla I_1(\mathbf{x})|\}}. \quad (16)$$

This will provide a unique value, which is obtained from the more restrictive situation. The advantage of this approach is that it is fully automatic. However, one limitation of this

strategy is that it will not respect the contours when the gradient is far from the maximum gradient. In this sense, this approach may be very conservative.

Local approach. The second alternative is based on a *local* strategy. In this case, λ is calculated as

$$\lambda_l(\mathbf{x}) := \begin{cases} \frac{-\ln(\xi) + \ln(\alpha)}{|\nabla I_1(\mathbf{x})|} & \text{if } \max_{\mathbf{x}_i \in \mathcal{N}(\mathbf{x})} \{\alpha e^{-\lambda_0 |\nabla I_1(\mathbf{x}_i)|}\} < \xi \\ \lambda_0 & \text{Otherwise} \end{cases}, \quad (17)$$

where λ_0 is a constant that is used as a reference value and $\mathcal{N}(\mathbf{x})$ is a close neighborhood around \mathbf{x} . If the second condition in (15) is satisfied, then we correct the value of λ as in (14). λ_0 can be calculated from the input images or introduced as a parameter. If this is conveniently selected, this approach allows to respect a broad variety of image contours, at the same time that it eliminates the possible instabilities.

5 Experimental Results

In this section, we analyze the methods depicted in Table 3, which includes the Brox *et al.* [11] method, a continuous version of the Xu *et al.* [43] smoothing strategy (*Exponential*) and the Monzón *et al.* [19] approach (*Exponential β*). We also take into account the two alternatives proposed in Sect. 4, that we will refer as *Exponential λ -Global* and *Exponential λ -Local*, respectively.

The first part of these experiments will study each method separately. Then, we will carry out a thorough comparison between the different approaches. Finally, we will present some numerical results for studying the accuracy of the methods, using the best parameter setting and default parameters.

We use synthetic sequences from two main datasets: a set of images with geometric figures, like circles, squares, stars or rectangles, presenting pronounced discontinuities and large displacements; and four sequences from the Middlebury benchmark database [6]. These sequences are RubberWhale, Grove2, Hydrangea and Yosemite.

The aim of the experiments is to analyze the influence of the regularization in the preservation of discontinuities. Therefore, our study will basically concentrate on the α and λ parameters. In the majority of the experiments, γ has been set to zero. The remaining parameters have been fixed as follows: $\eta := 0.75$, $\varepsilon := 0.001$ and *inner_iterations* := 1; *outer_iterations* := 38 for the Yosemite and Middlebury sequences, and *outer_iterations* := 15 for the geometric ones. N_{scales} is automatically calculated so that the image size, at the coarsest scale, is around 16×16 pixels. These parameters are set according to [31].

The optical flows are represented using the color scheme shown in Fig. 1. The color represents the orientation and the intensity the magnitude, similar to the vector field on the left.

Table 3: Survey methods

Method	$\mathcal{R}(\nabla I_1, \nabla u, \nabla v)$
Brox <i>et al.</i> [11]	$\sqrt{ \nabla u ^2 + \nabla v ^2 + \epsilon^2}$
<i>Exponential</i> [43]	$\sqrt{e^{-\lambda \nabla I_1 }(\nabla u ^2 + \nabla v ^2) + \epsilon^2}$
<i>Exponential β</i> [19]	$\sqrt{(e^{-\lambda \nabla I_1 } + \beta)(\nabla u ^2 + \nabla v ^2) + \epsilon^2}$
<i>Exponential λ-Global</i> [25] (16)	$\sqrt{e^{-\lambda_g \nabla I_1 }(\nabla u ^2 + \nabla v ^2) + \epsilon^2}$
<i>Exponential λ-Local</i> (17)	$\sqrt{e^{-\lambda_l \nabla I_1 }(\nabla u ^2 + \nabla v ^2) + \epsilon^2}$

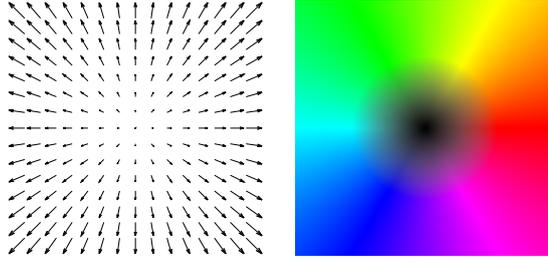


Figure 1: Color scheme for motion representation

5.1 Analysis of the *Exponential* Method

Figures 2 and 3 show the optical flows obtained for the geometric and Middlebury sequences, with α constant and different λ values. We observe that a small value underestimates the magnitude of the optical flows. This is typical because the regularization parameter is high.

When the value of λ is properly chosen, then the resulting optical flow is very accurate. Discontinuities are correctly preserved, except at the corner of the geometric figures. This is probably due to the gradient and the numerical scheme used. On the other hand, we observe that occlusions are also correctly handled. The use of a large α , and the strict stopping effect of λ , allows to fill the occluded regions with the value from the background. This is achieved because, on the one hand, it overweights the regularization with respect to the attachment term and, on the other hand, it does not mix the information with the flow inside the figures.

Finally, we appreciate how outliers appear when λ is large. These are typically located at flow discontinuities and, in the case of the star, also inside the figure, probably as the result of smoothing with the outliers at contours. A similar behavior can be observed in the Middlebury sequences.

The graphics depicted in Fig. 4 show the evolution of the AAE with respect to λ . The values shown in the horizontal axis correspond to the λ values used in Figs. 2 and 3.

We observe that AAE errors decrease in all the graphics. For $\lambda := 0$ the result is equivalent to the Brox *et al.* method. Thus, the exponential function clearly improves the basic model. This is more important in the geometric figures, where there is an important jump in accuracy after a suitable value. Then, the errors increase for bigger λ values and the solutions become unstable. The best results are normally obtained for large values of α , which also delay the instability problems. In the case of Grove2, the best result is obtained for a small α .

Looking at these graphics, the λ optimum is located approximately in the range $[0.15, 0.25]$ for the square sequence. Values below 0.15 has a large AAE and outliers appear when λ is above 0.25. In the star sequence, the same happens in the range $[0.08, 0.2]$. The best results are located in the range $[0.12, 0.4]$ and $[0.02, 0.4]$ for the Grove2 and Yosemite sequences, respectively. Outliers appear when λ is above 0.4, although the results continue to be coherent.

After these experiments, we may conclude that a default value for λ may be chosen around 0.17. We note that the input images have been normalized between 0 and 255, so this default value correctly fits to this range of values.

5.2 Analysis of the *Exponential* β Method

Figures 5 and 6 show several examples of the *Exponential* β method. We have used large values for α and λ , and set γ to zero. In Fig 5, we use the geometric sequences of the star, the double rectangle and the square. In the first one, we observe that, increasing the value of β , the rounded effects of Brox *et al.* appears at the star edges. In the double rectangle, the improvement achieved with a small value of β is very flattering, eliminating the outliers

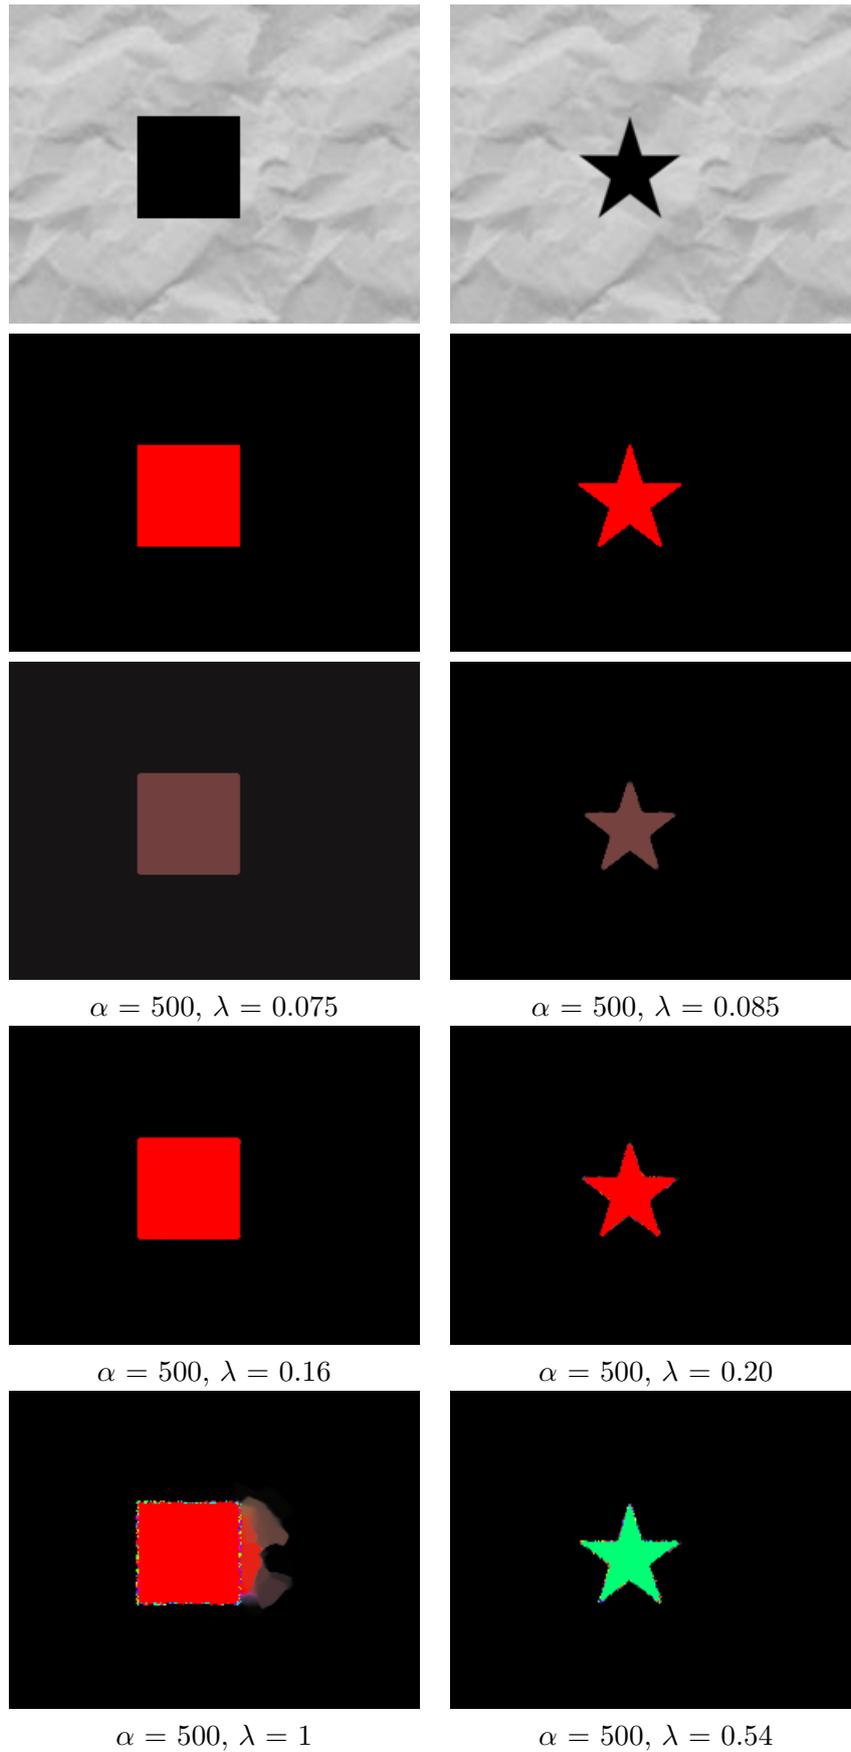


Figure 2: *Exponential* method. From top to bottom, the first frame of each sequence, the ground truth and the optical flows obtained.

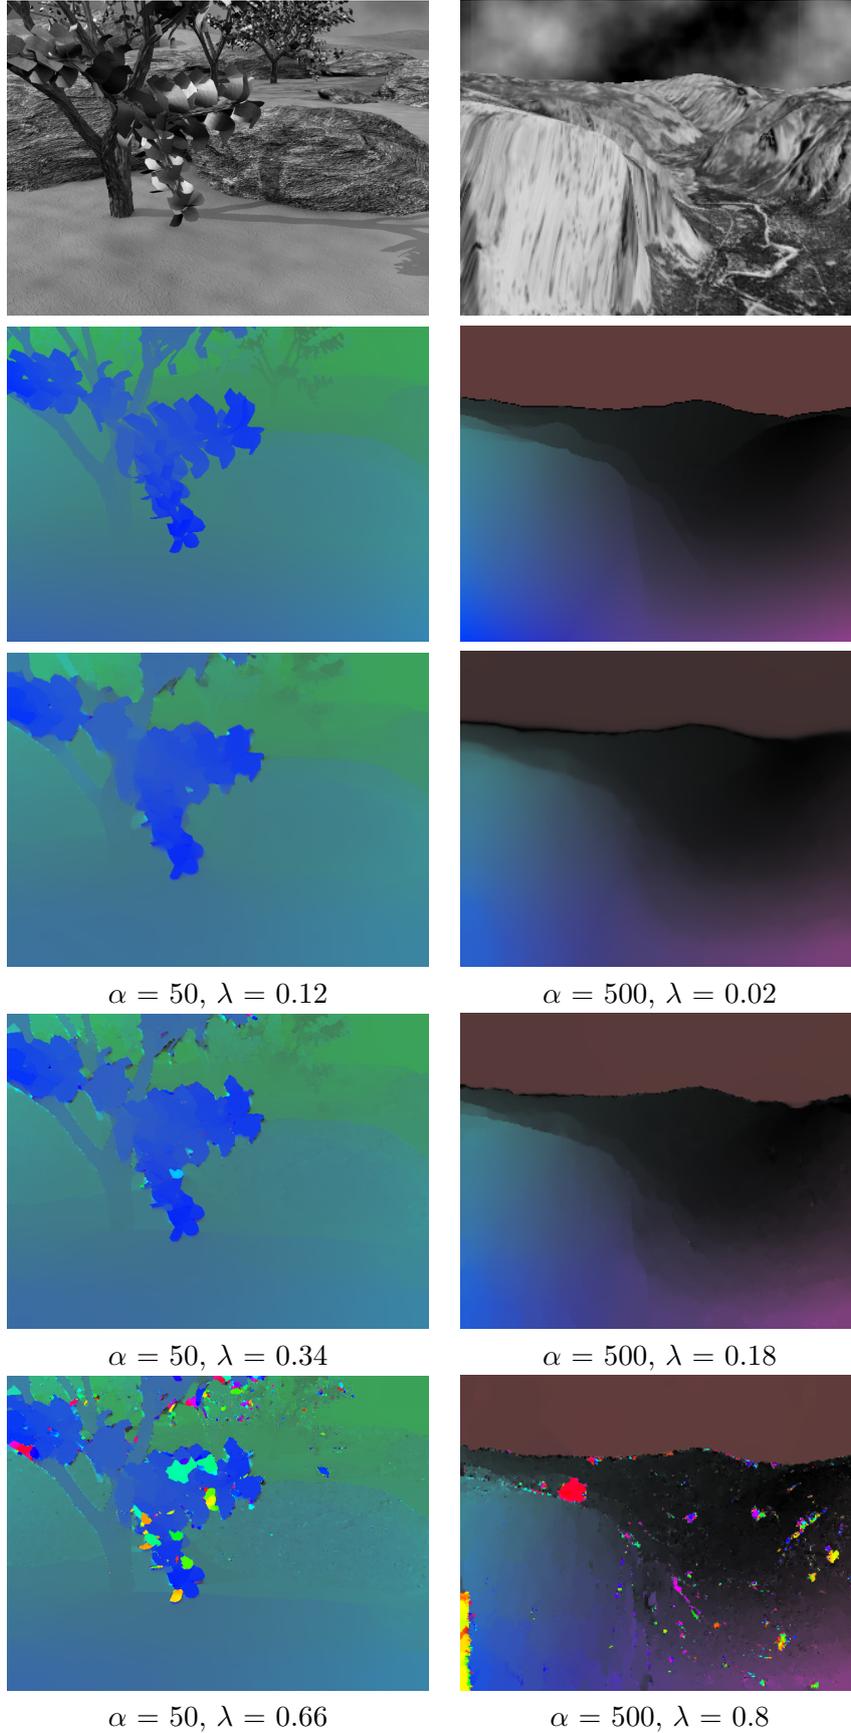


Figure 3: *Exponential* method. From top to bottom, the first frame of each sequence, the ground truth and the optical flows obtained.

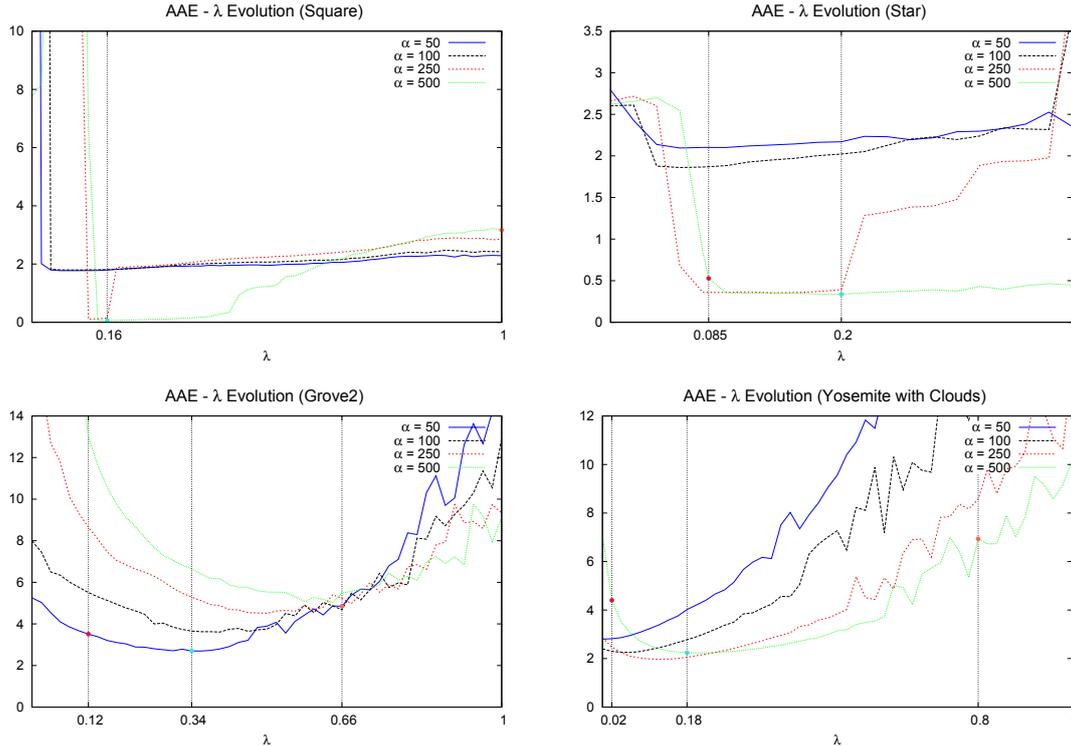


Figure 4: λ evolution using the *Exponential* method

of the exponential method. However, when $\beta := 0.001$, the resulting flow is very smoothed, disappearing the clear separation between both rectangles. This situation is more aggressive in the square sequence due to the use of a larger α .

In Fig. 6 we show the results for Hydrangea, Grove2 and RubberWhale. These sequences have an important number of strong gradient variations between the leafs of Hydrangea and Grove2 and some objects in RubberWhale.

In this figure, we can observe that the result of the *Exponential* method is promising, since the preservation of discontinuities is accurate. Nevertheless, the number of outliers is important. These are removed with the *Exponential* β method using a small β . By contrast, if this parameter is bigger, the motion field is very smoothed. In RubberWhale, discontinuities are correctly detected, but some regions are flattened, like the motion at the wheel.

The graphics depicted in Fig. 7 show the evolution of the AAE with respect to β . The blue line represents a configuration that uses little regularization and where the exponential function intercede slightly in the result. The green line shows the behavior of the method using medium values for α and λ . Finally, the red line highlight the result using extreme values for these two parameters, incrementing the regularization and the importance of the decreasing function. The first column depicts the graphics of the geometric sequences while, the second, shows the evolution of the AAE in Middlebury.

We observe that, in general, when we use the *Exponential* solution ($\beta := 0$), a large parameter configuration offers high errors. In the Middlebury sequences, the large configuration shows that the error decreases considerably with low values of β and, according its influence grows, the error increases. The double rectangle has a similar effect, but the estimation worsens faster and more aggressively. The solutions of the star, on the other hand, remain stable until β reaches 0.004, decreasing the error. Afterwards, it begins to grow strongly.

The intermediate configuration shows a behavior similar to the above. The square sequence is initially very stable, but for $\beta > 0.0001$ the error increases considerably. Finally,

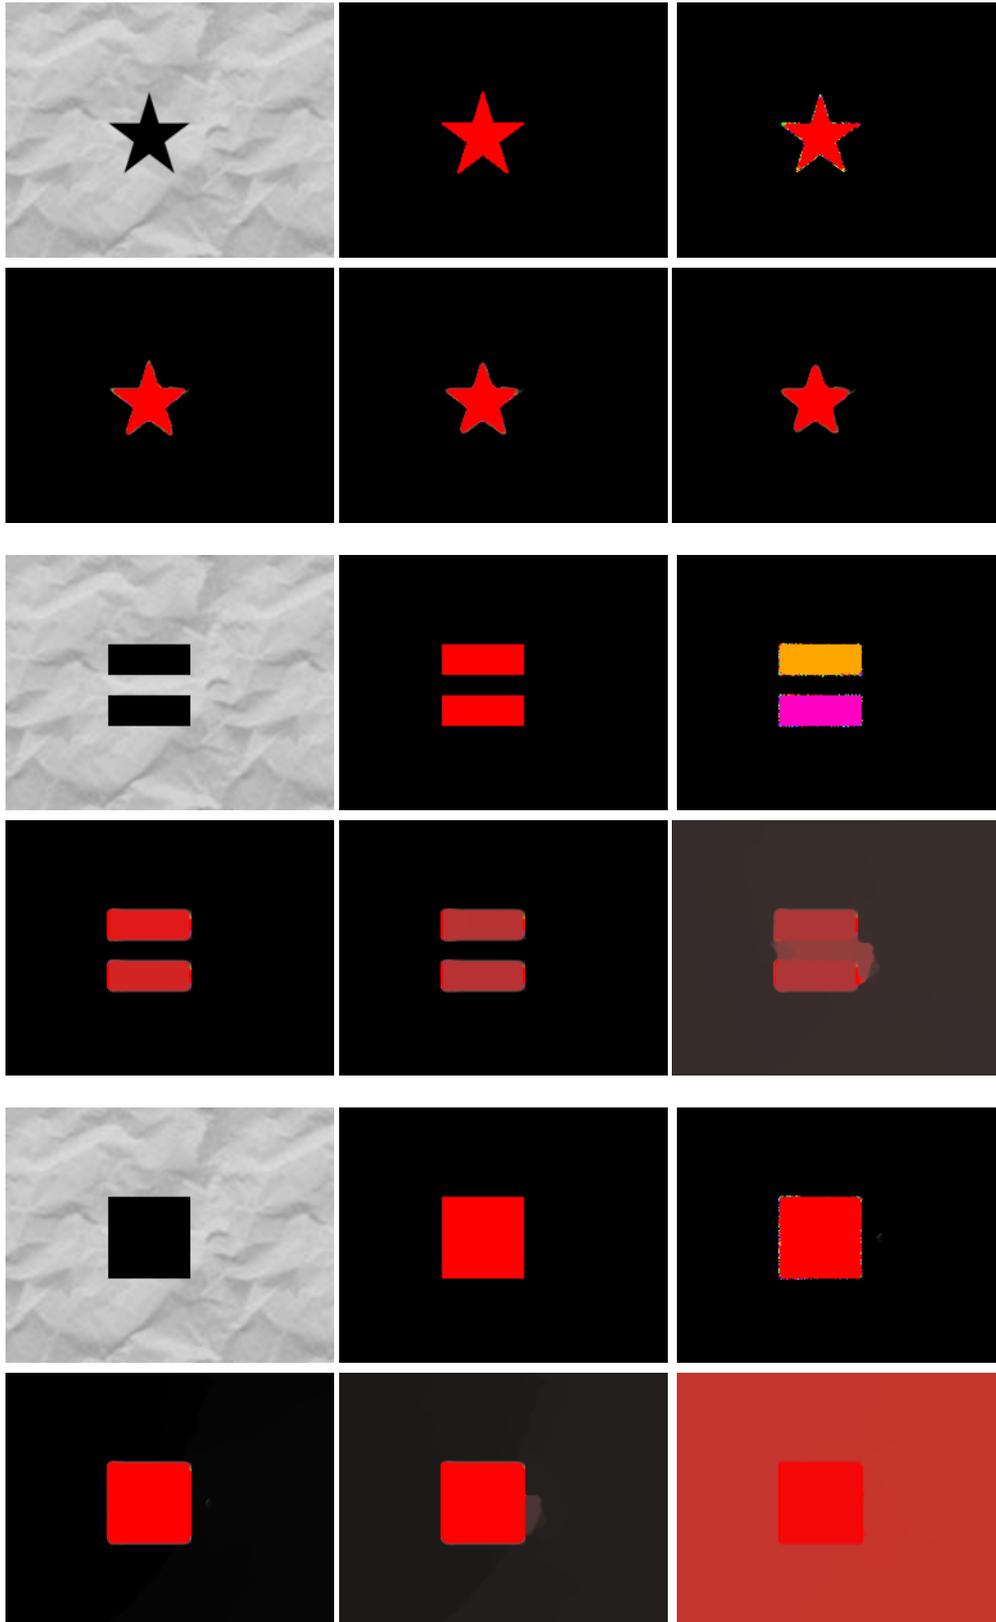


Figure 5: *Exponential* β method. Example with the star, double rectangle and square sequences. The first row shows the original image, the ground truth and the optical flow obtained with the basic *Exponential* method. The second row shows the results for $\beta := \{0.00001, 0.0001, 0.001\}$, respectively.

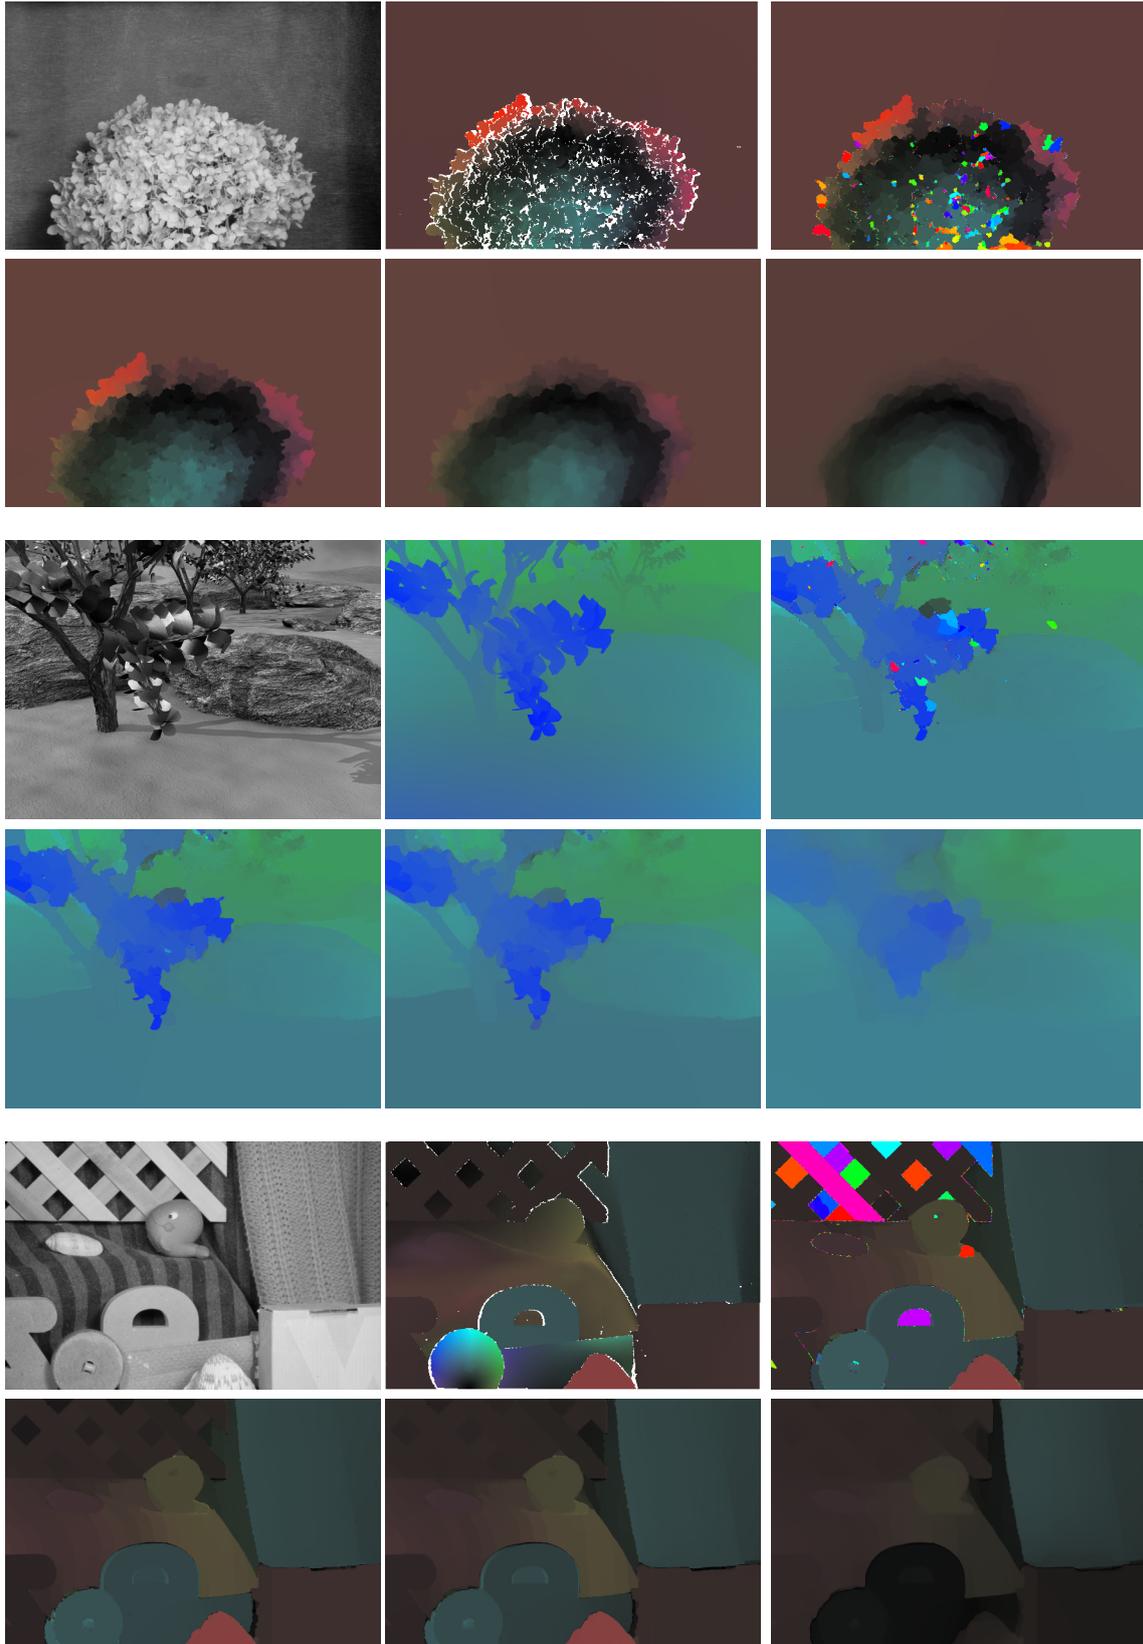


Figure 6: *Exponential* β method. Example with the Hydrangea, Grove2 and RubberWhale sequences. The first row shows the original image, the ground truth and the optical flow obtained with the basic *Exponential* method. The second row shows the results for $\beta := \{0.00001, 0.0001, 0.001\}$, respectively.

we see that the blue line is very stable in general. This means that, for regular values of α and λ , this method does not provide an improvement with respect to the *Exponential* method. On the other hand, it does not worsen the results, except for the double rectangle sequence.

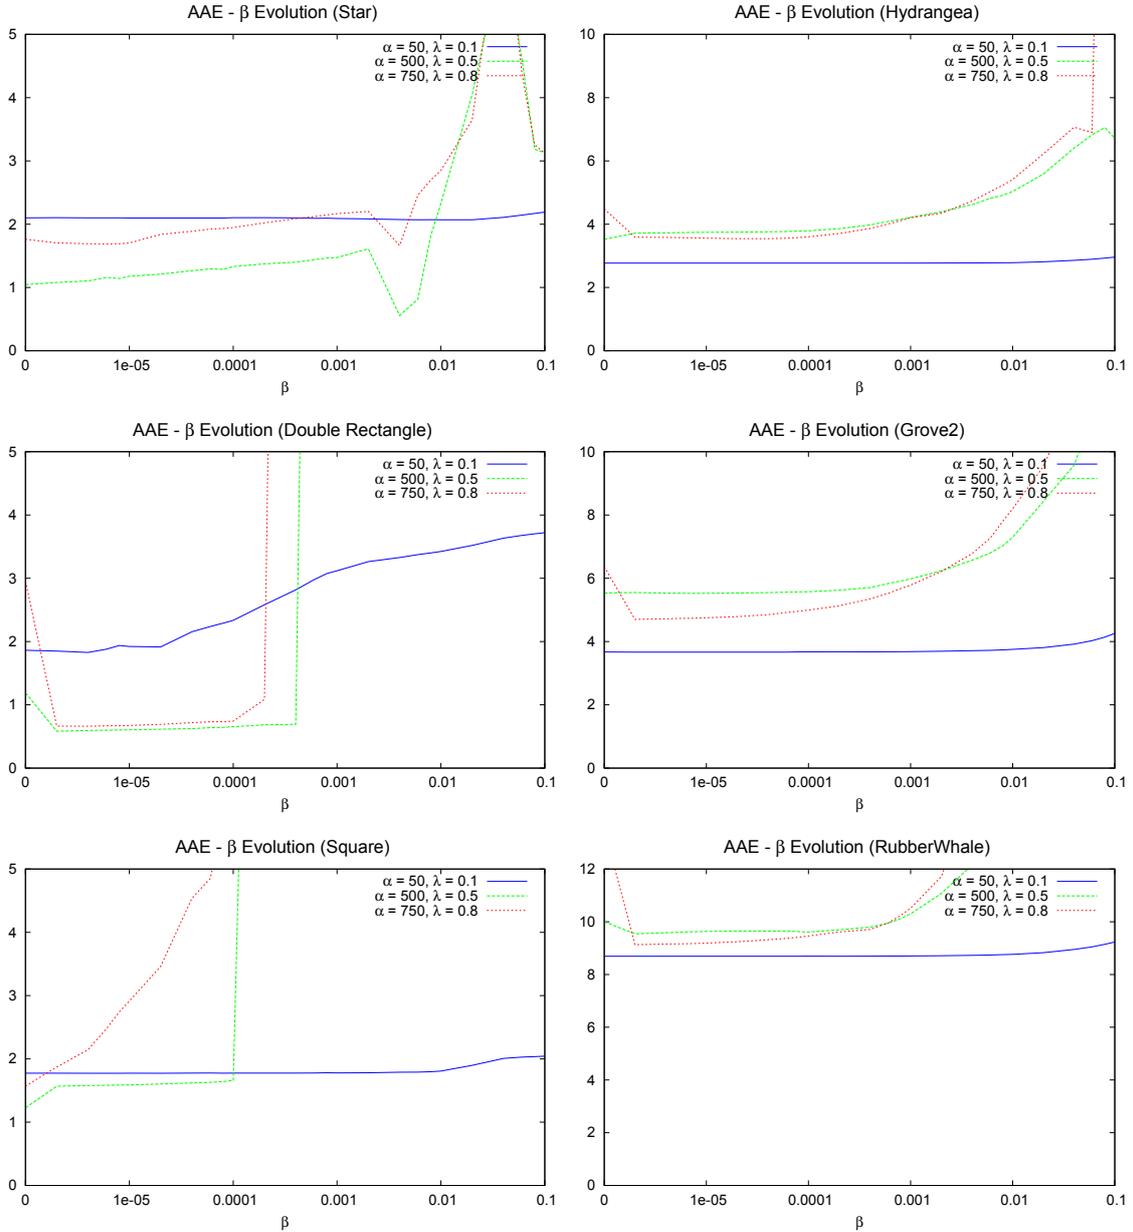


Figure 7: *Exponential* β method. AAE evolution with respect to β . First column shows the graphics for the geometric sequences and, the second, the graphics for the Middlebury sequences.

5.3 Analysis of the *Exponential* λ -Global and λ -Local Methods

In this section, we analyze the behavior of the variants of the *Exponential* method, as seen in Sect. 4: the *Exponential* λ -Global method, given in (16); and the *Exponential* λ -Local method, given in (17).

It is desirable that λ could be computed automatically. This is the idea behind the *local* and *global* approaches. The main difference between them is that, in the *global* variant, λ is calculated using the maximum gradient of the whole image, whereas, in the *local* one, the

parameter is adapted according to the neighborhood around each pixel.

λ -*Global* correctly preserves the flow discontinuities when the corresponding image gradients are similar in the whole image. However, when the gradients in the discontinuities are much smaller than the maximum gradient, this method works poorly. In this case, the discontinuities tend to be smoothed due to the difference with respect to the maximum gradient. The λ -*Local* approach tries to avoid this problem, adopting the same idea but locally. The diffusion process is controlled in a different way depending on the local gradient.

5.3.1 λ -Global strategy

In this section, we analyze the influence of the ξ parameter in the *global* strategy. This parameter helps avoid the instability problems that may appear in the *Exponential* method. In the following tests, we use large values for α and λ , in order to study the influence of ξ when outliers appear in the solutions. Figures 8 and 9 show the results obtained for the geometric and Hydrangea sequences, respectively.

In Fig. 8, we see that the λ -*Global* strategy provides more accurate solutions than the *Exponential* method. Small values of ξ still yields instabilities in the optical flow, however, when ξ is increased, the outliers progressively disappear. For larger values, the solutions become smoother and, finally, the flow discontinuities are not preserved.

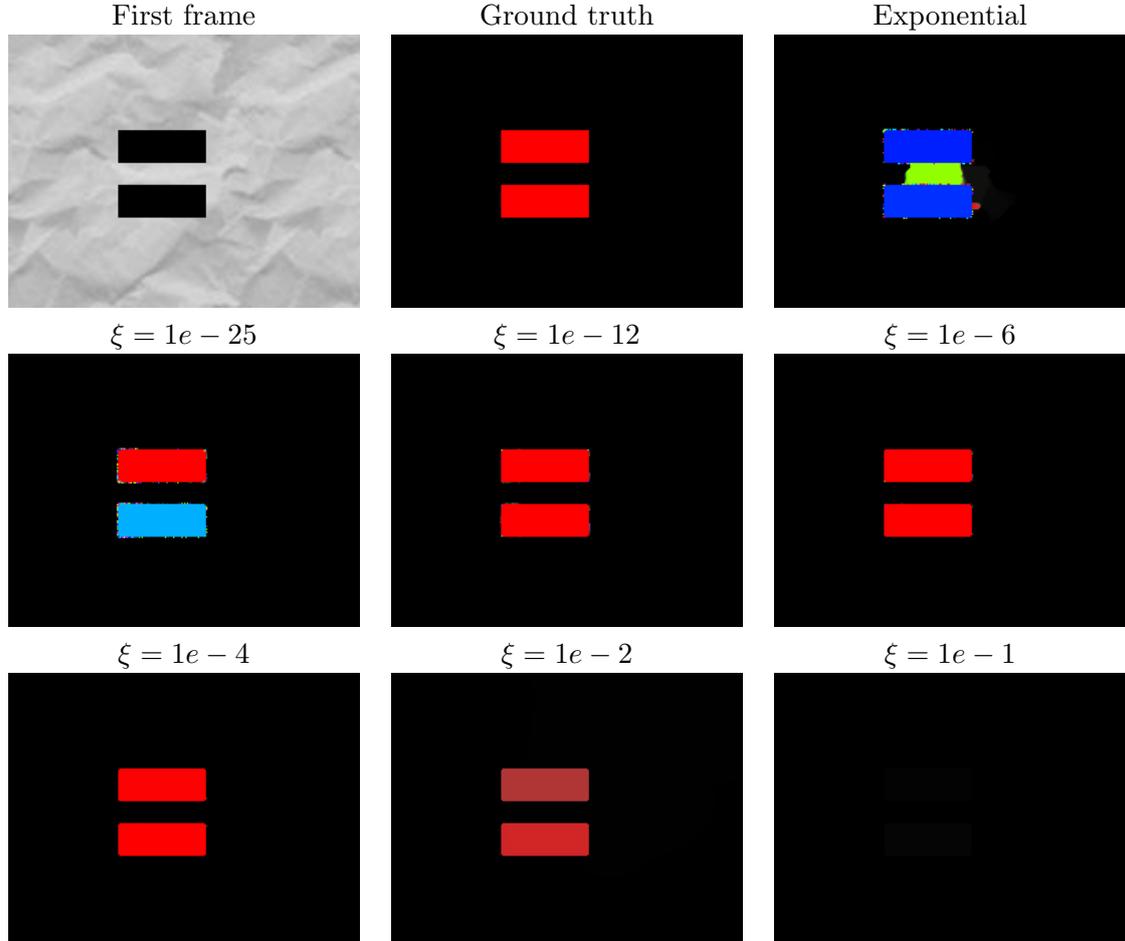


Figure 8: λ -*Global* strategy. ξ evolution with $\lambda = 1$. From left to right, top to bottom: the first frame of the sequence, the ground truth, the solution given by the *Exponential* method; the estimations obtained using the *global* strategy for different values of ξ .

Figure 9 depicts the same evolution with respect to the Hydrangea sequence. The behavior is similar to the previous sequence and the instabilities disappear when we increase the

value of ξ . However, we observe a stronger regularization of the motion field. This is due to the big difference between the maximum gradient and the gradients of other discontinuities, that yields the quantization of the optical flow in a few planar regions.

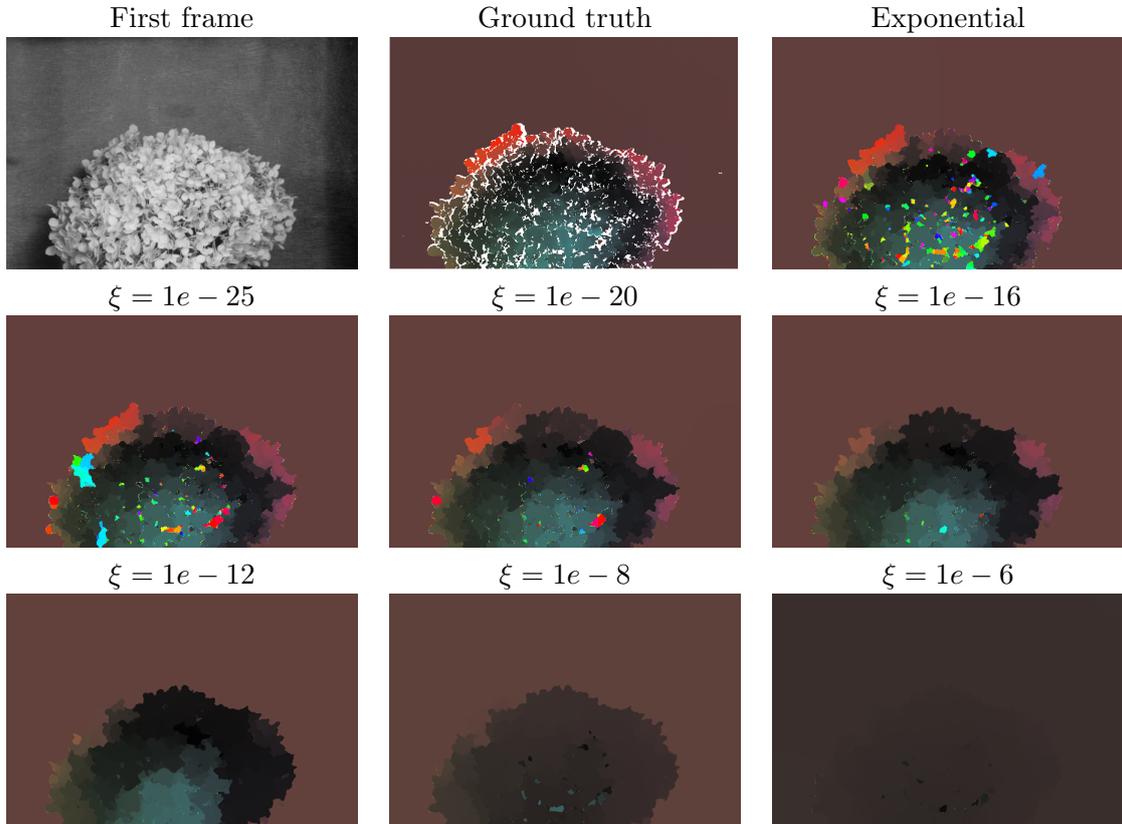


Figure 9: λ -Global strategy. ξ evolution with $\lambda = 1$. From left to right, top to bottom: the first frame of the sequence, the ground truth, the solution given by the *Exponential* method; the estimations obtained using the *global* strategy for different values of ξ .

5.3.2 λ -Local strategy

In the λ -Local strategy, a different λ is applied to each image pixel using the local maximum gradient. In Figs. 10 and 11, we show the results obtained with the *Exponential* method and the *local* approach, for several values of ξ , using the double rectangle and Hydrangea sequences, respectively.

In the geometric sequences, the maximum gradient is located in the main object contours. For this reason, the behavior of the λ -Local and λ -Global strategies seems to be similar (compare Figs. 8 and 10). Nevertheless, there are two important differences: on the one hand, we observe that the outliers are not completely removed until a bigger ξ ; on the other hand, discontinuities of the optical flows are better preserved for larger values.

In the Hydrangea sequence, the differences between the two strategies are more evident. The λ -Local strategy better preserves the flow discontinuities, as shown in Fig. 11. This approach seems to be more robust with respect to a big variety of image gradients. The use of local gradients, instead of the maximum, allows to control more precisely the amount of diffusivity in each region of the image.

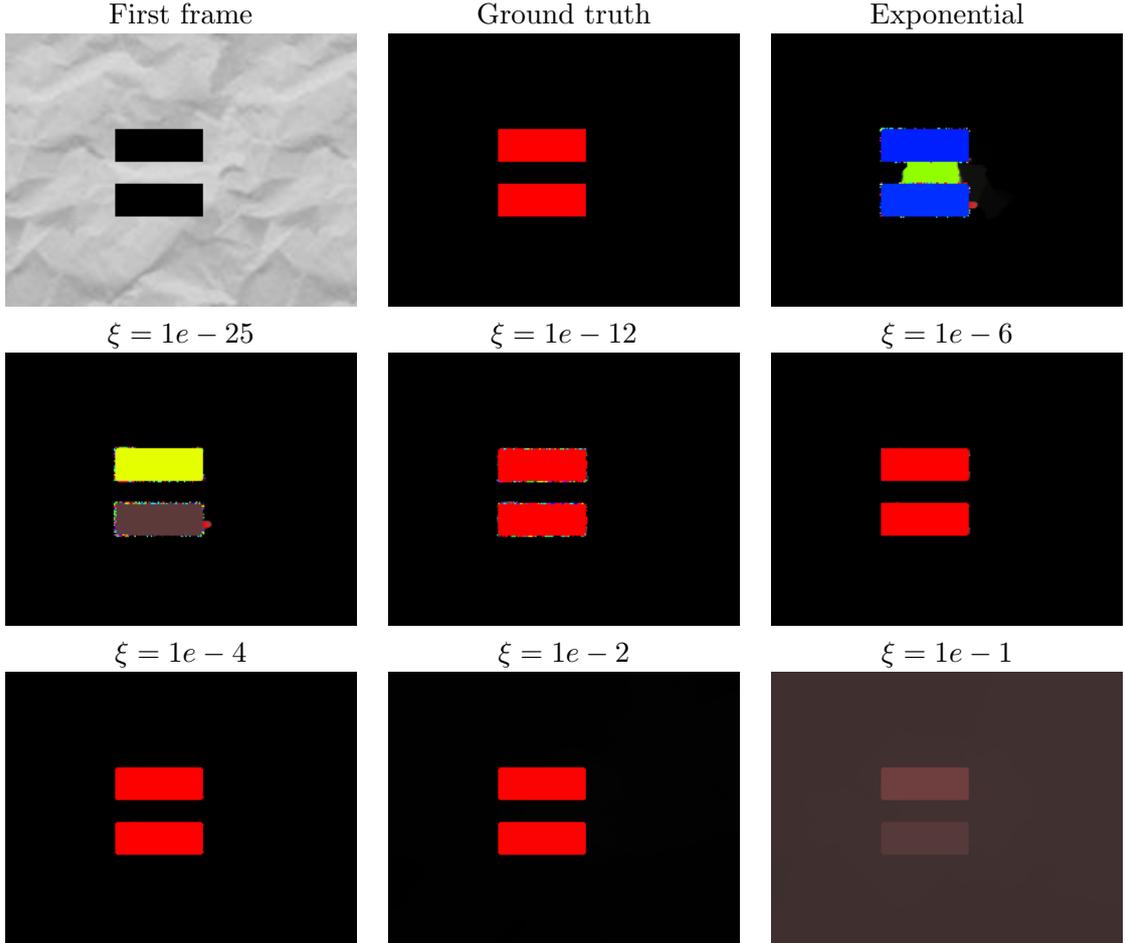


Figure 10: λ -Local strategy. ξ evolution with $\lambda = 1$. From left to right, top to bottom: the first frame of the sequence, the ground truth, the solution given by the *Exponential* method; the estimations obtained using the *local* strategy for different values of ξ .

5.4 Comparison of the Methods

Once analyzed the different approaches, we now compare them and highlight their main features. Figures 12, 13, 14, 15, 16 and 17 show the solutions obtained with the different methods, using several geometric and Middlebury sequences.

We show three results for each method in these figures: first, a typical solution obtained with a regular configuration of the α and λ parameters; then, a good solution looking for better parameters; and, finally, a limiting situation where we observe the appearance of outliers or very regularized solutions (large values of α and λ). In these experiments, the gradient term has been cancelled ($\gamma = 0$).

In general, the Brox *et al.* method can not completely stop the diffusion at discontinuities. In fact, it has many difficulties to deal with simple geometric shapes (see Figs. 12, 13 and 14). Some typical problems are that contours of the motion regions are normally dislocated, the magnitude of the optical flow inside the figures is underestimated or occlusions have a strong impact on the final solution. These problems are much more significant with concave shapes (see the star sequence). These also happen in the Middlebury sequences. However, the differences with respect to the exponential methods are not so noticeable. In the Rubber-Whale sequence (Fig. 17), we observe smooth transitions at the contours of regions moving in opposite directions.

In the results of the *Exponential* method, discontinuities are, in general, correctly preserved. The shapes of the geometric sequences are clearly segmented from the background

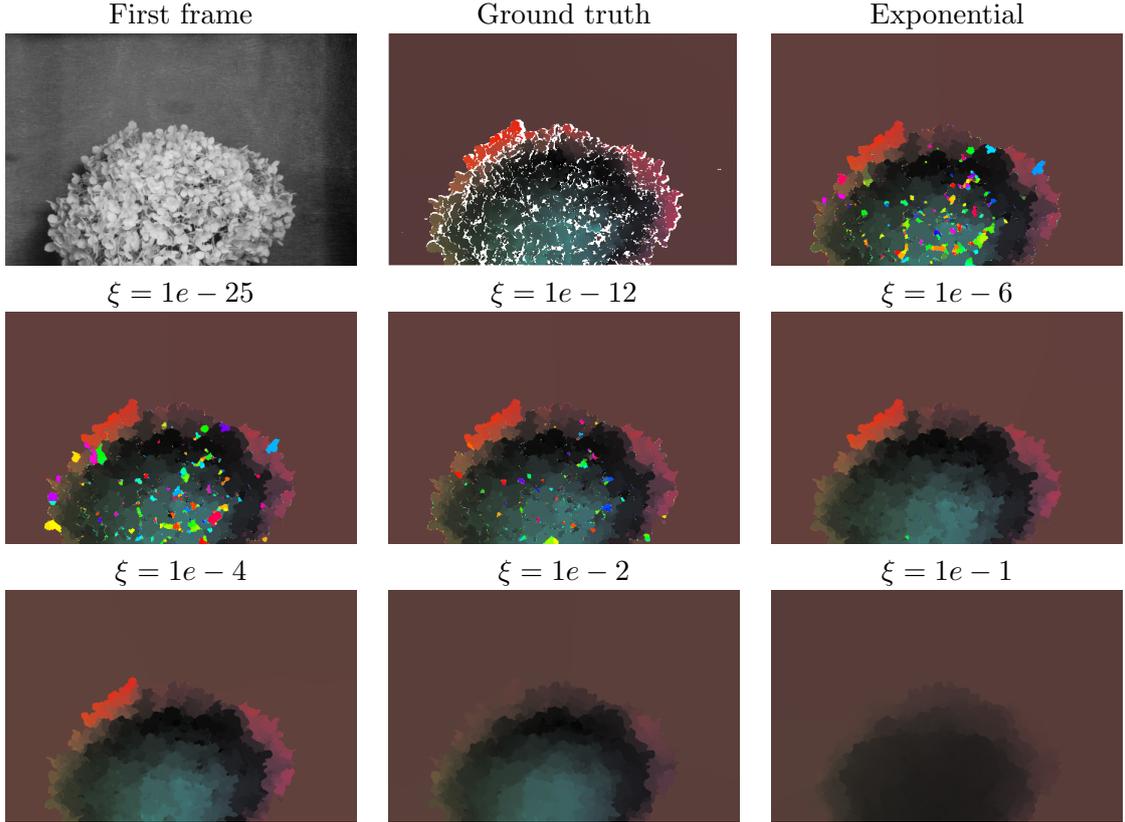


Figure 11: λ -Local strategy. ξ evolution with $\lambda = 1$. From left to right, top to bottom: the first frame of the sequence, the ground truth, the solution given by the *Exponential* method; the estimations obtained using the *local* strategy for different values of ξ .

motion. We observe that the effect of occlusions has a limited impact in the final solutions. Moreover, a large regularization parameter allows to fill the information inside these regions with the correct values. It also provides very good results for both convex and concave shapes. The transitions between regions almost disappear for large values of λ . 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 .

The *Exponential β* approach effectively eliminates the instabilities of the *Exponential* method. In the images, we see that typical solutions are very similar to the *Exponential* method, but when we increase the value of α , the method behaves like the Brox *et al.* approach. Although it is very stable, it cannot cope with occlusions due to the small isotropic regularization.

The λ -Global and λ -Local strategies obtain similar results to the *Exponential* method. The main benefit is that they do not introduce any instability: in the experiments we see how the outliers that appear in the *Exponential* solution have been removed. We do not appreciate significant differences between these two strategies in the geometric sequences. This makes sense because the maximum image gradient is associated with the object discontinuities. However, in the Middlebury sequences the λ -Local strategy offers better results. In fact, the solutions seem to be the same as the *Exponential* method but without outliers.

5.4.1 λ Evolution

Next, we analyze the accuracy of the exponential methods with respect to λ , for small, medium and large α values. Figures 18 and 19 show the results for the geometric and Middlebury sequences, respectively.

In Fig. 18, we observe that the different variants significantly improve the outcome of the Brox *et al.* method. The first column shows that, in the case of the circle and square, the methods are very stable. This stability differs greatly in the star and the double rectangle sequences, where the AAE evolution is more unstable after reaching the minimum error, specially in the *Exponential* method. In the second and third columns, we observe that the accuracy is surprisingly very high in all the sequences. Nevertheless, as the value of λ increases, the AAE augments noticeably, although more slowly for larger values of α . The behavior of the *Exponential* variant is very unstable in the star and the double rectangle. This effect is less pronounced using *Exponential β* and *Exponential λ -Local*.

In Fig. 19, we notice that a small α does not always improve the Brox *et al.* results, as it happens in Yosemite, RubberWhale and Hydrangea. In fact, we observe that, for the last two sequences, larger values of α hardly improves the result of Brox *et al.* At the beginning of the graphics, the exponential variants have a similar behavior. However, when λ increases, the *Exponential* method becomes unstable, while the other exponential approaches have a smoother evolution. As α augments, the stability of the methods gets better.

Curiously, the error evolution of *Exponential β* and *Exponential λ -Local* are very similar, reflecting, in many cases, a continuous improvement as λ increases. This similarity can guide us to decide which approach is more interesting, because, although both of them present very close results, *Exponential β* is easier to implement.

Finally, it is interesting to note that, *λ -Global* usually finds a value of λ that is close to the best solution in the majority of occasions. Theoretically, the result represented by the green dot should match with the best error of the pure *Exponential* function. This does not always happen because the exponential method uses the same λ in each scale while the *λ -Global* strategy recalculates this parameter.

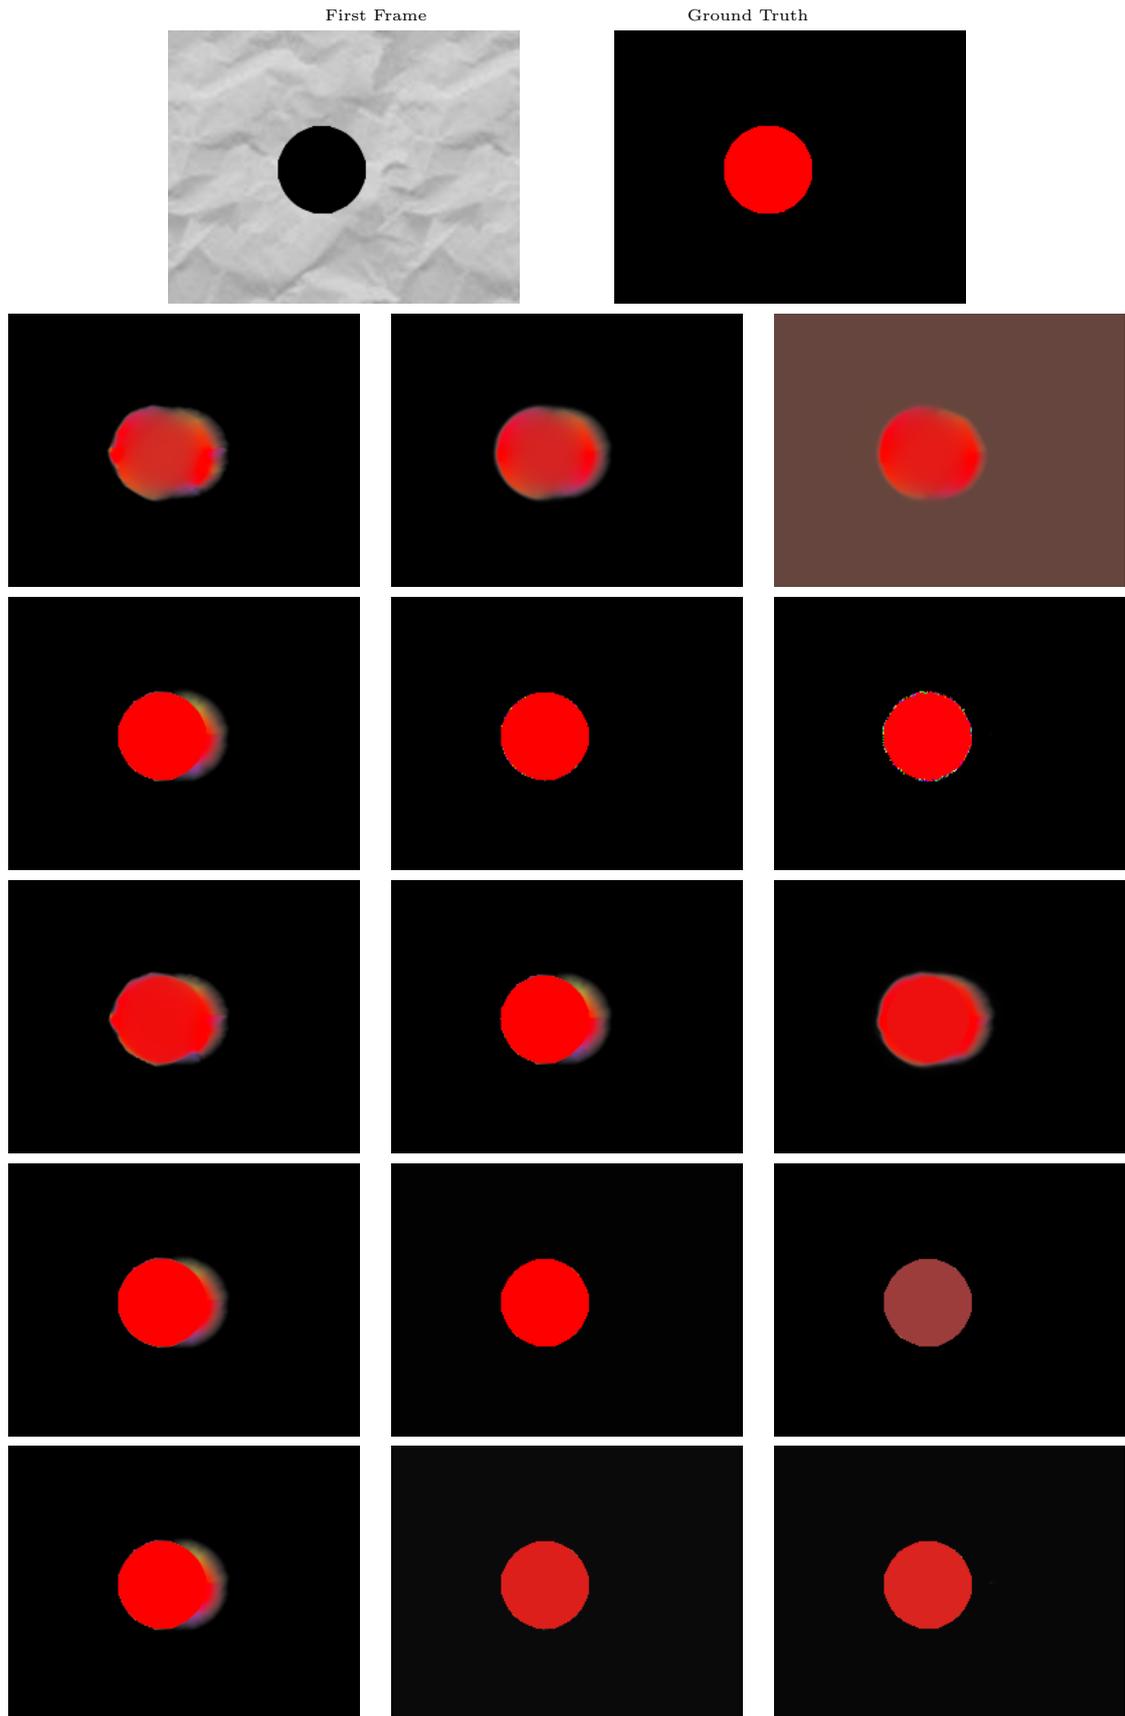


Figure 12: From top to bottom: Solutions of the Brox *et al.*, *Exponential*, *Exponential β* , *Exponential λ -Global* and *Exponential λ -Local* methods, respectively. The columns show increasing values for α and λ .

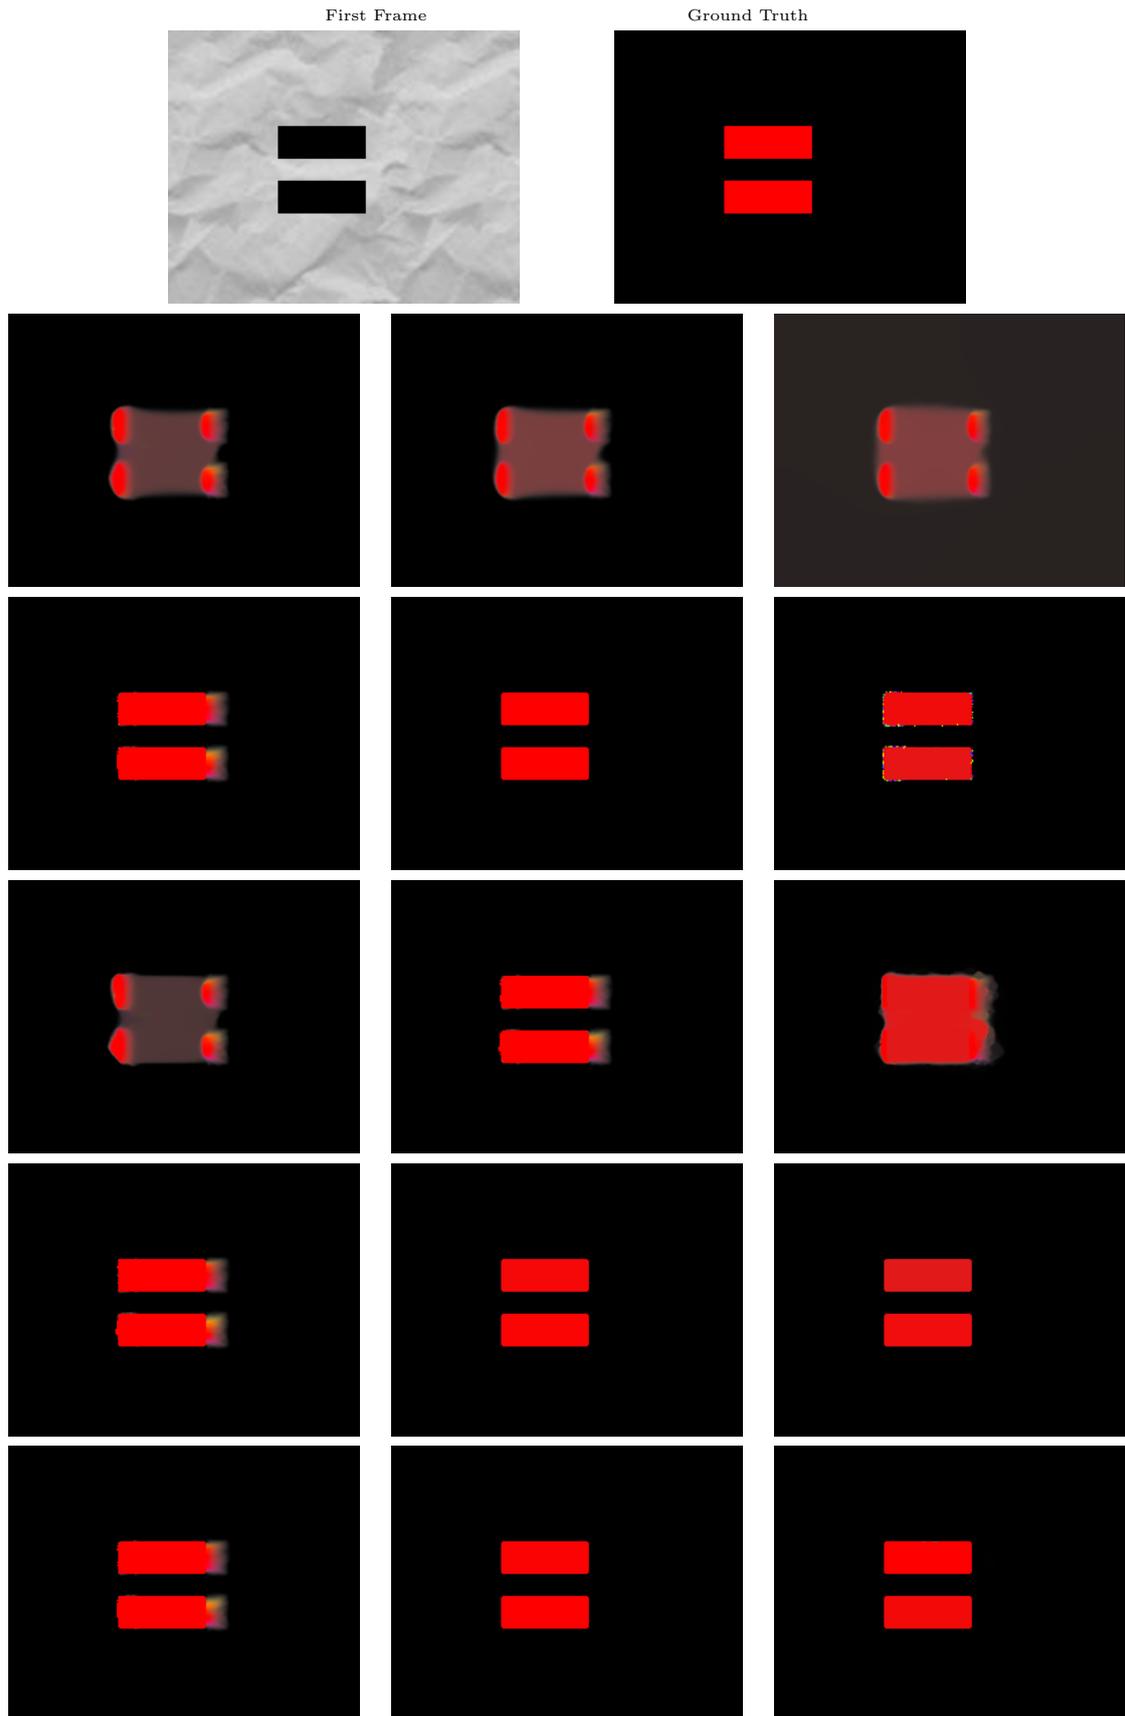


Figure 13: From top to bottom: Solutions of the Brox *et al.*, *Exponential*, *Exponential β* , *Exponential λ -Global* and *Exponential λ -Local* methods, respectively. The columns show increasing values for α and λ .

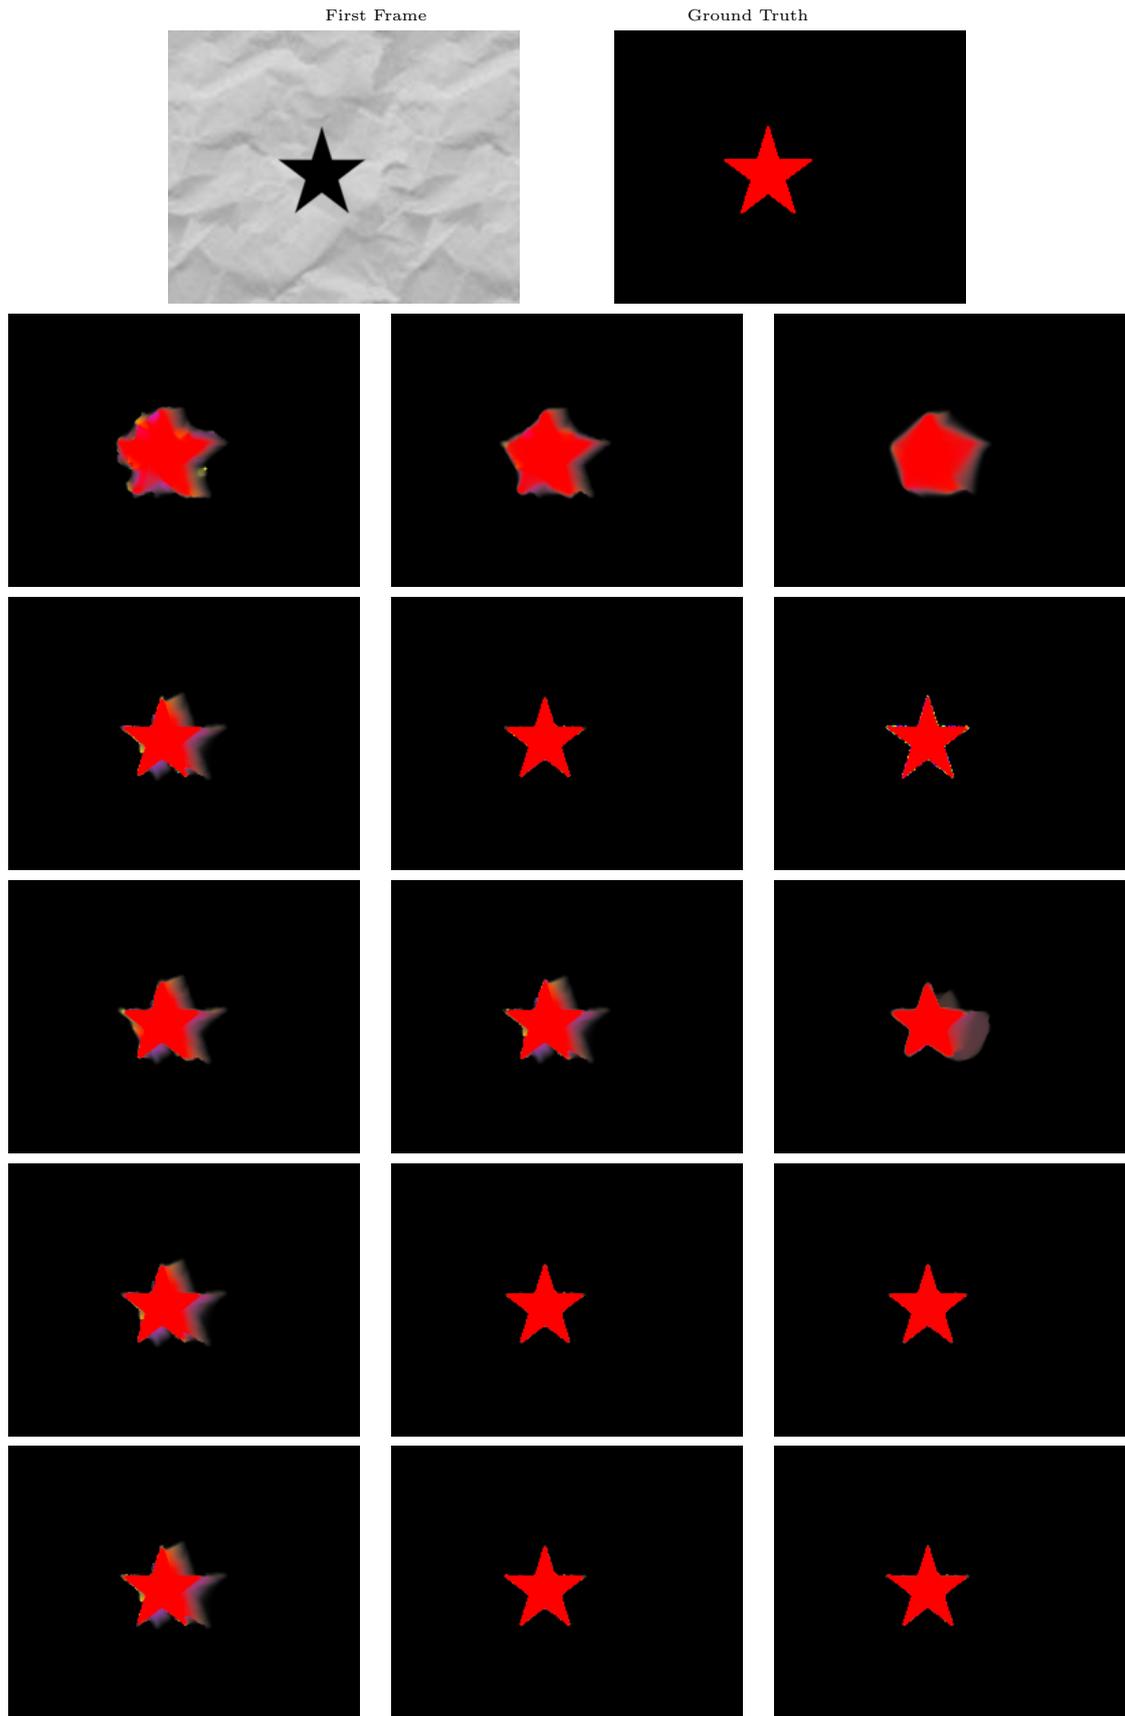


Figure 14: From top to bottom: Solutions of the Brox *et al.*, *Exponential*, *Exponential β* , *Exponential λ -Global* and *Exponential λ -Local* methods, respectively. The columns show the results for increasing values of α and λ .

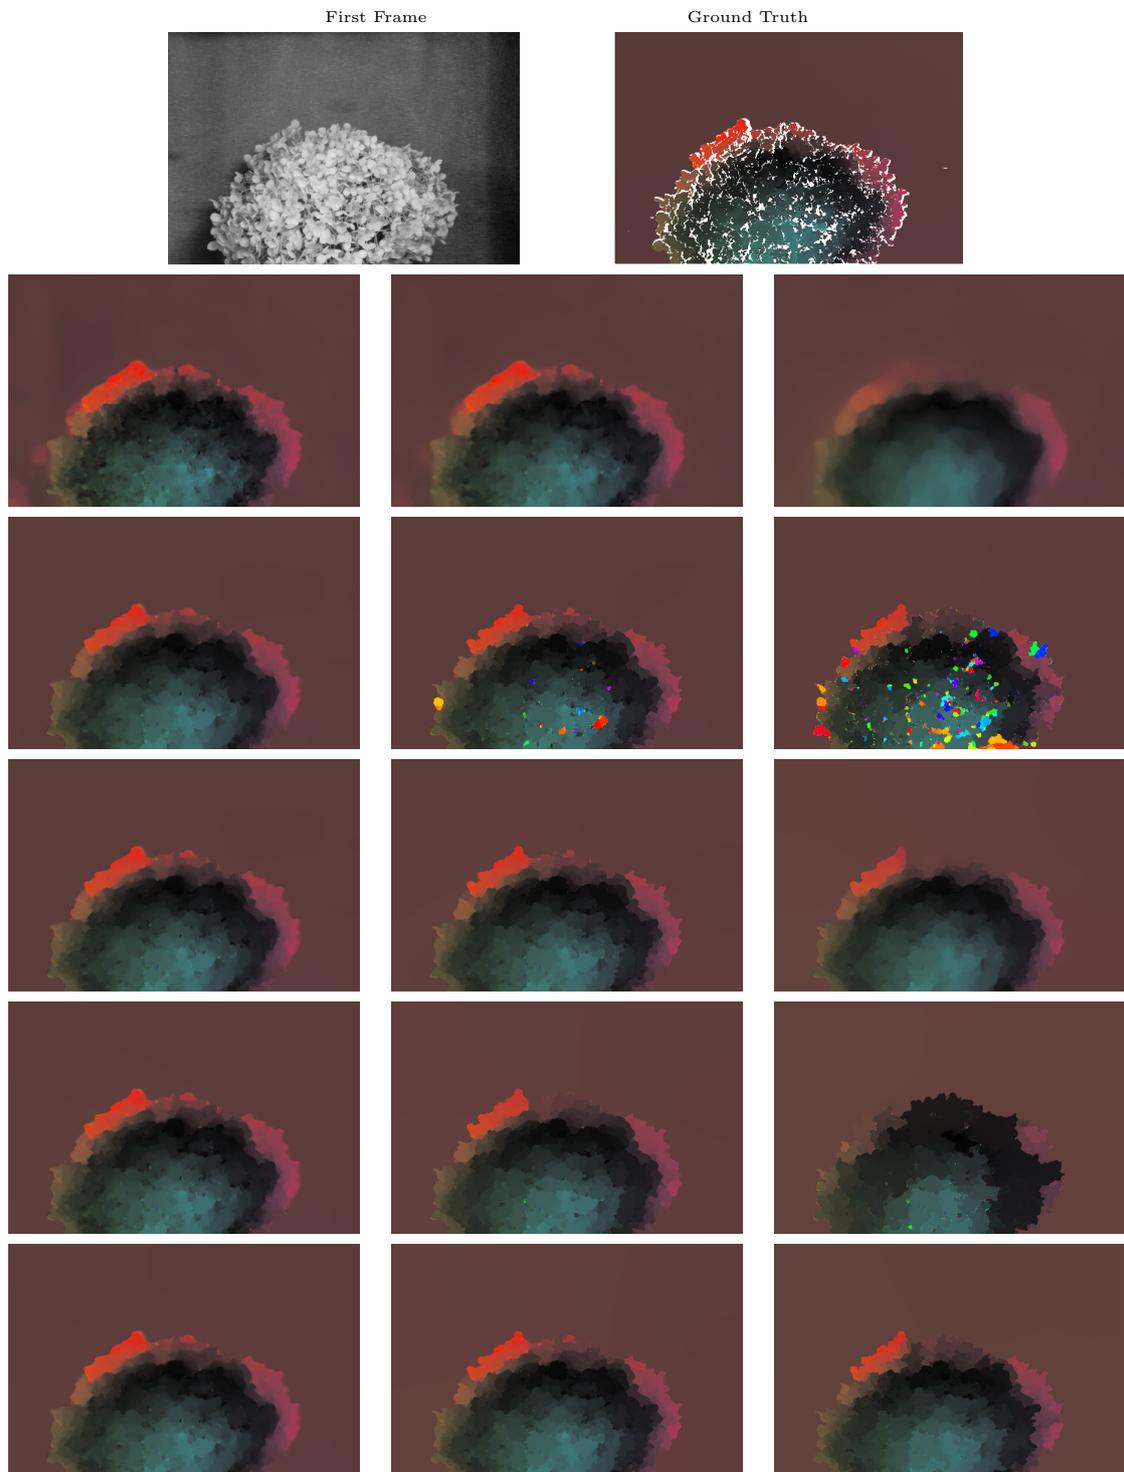


Figure 15: From top to bottom: Solutions of the Brox *et al.*, *Exponential*, *Exponential β* , *Exponential λ -Global* and *Exponential λ -Local* methods, respectively. The columns show the results for increasing values of α and λ .

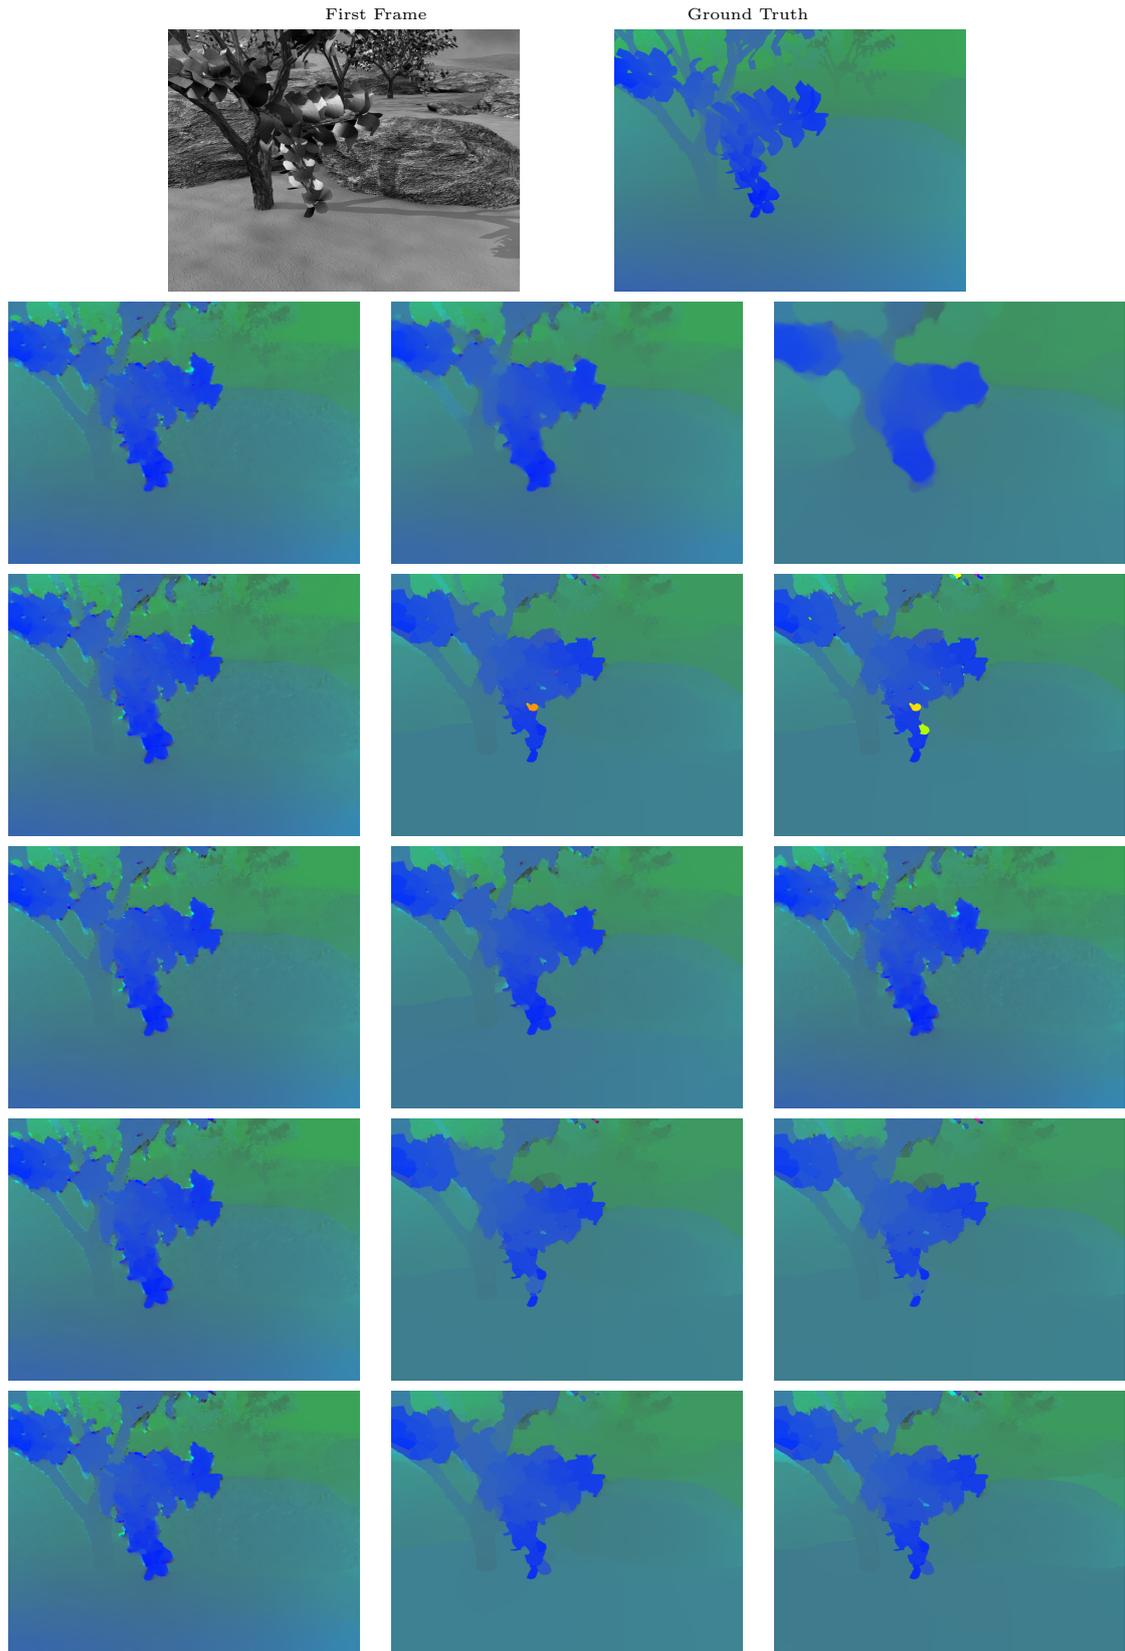


Figure 16: From top to bottom: Solutions of the Brox *et al.*, *Exponential*, *Exponential β* , *Exponential λ -Global* and *Exponential λ -Local* methods, respectively. The columns show the results for increasing values of α and λ .

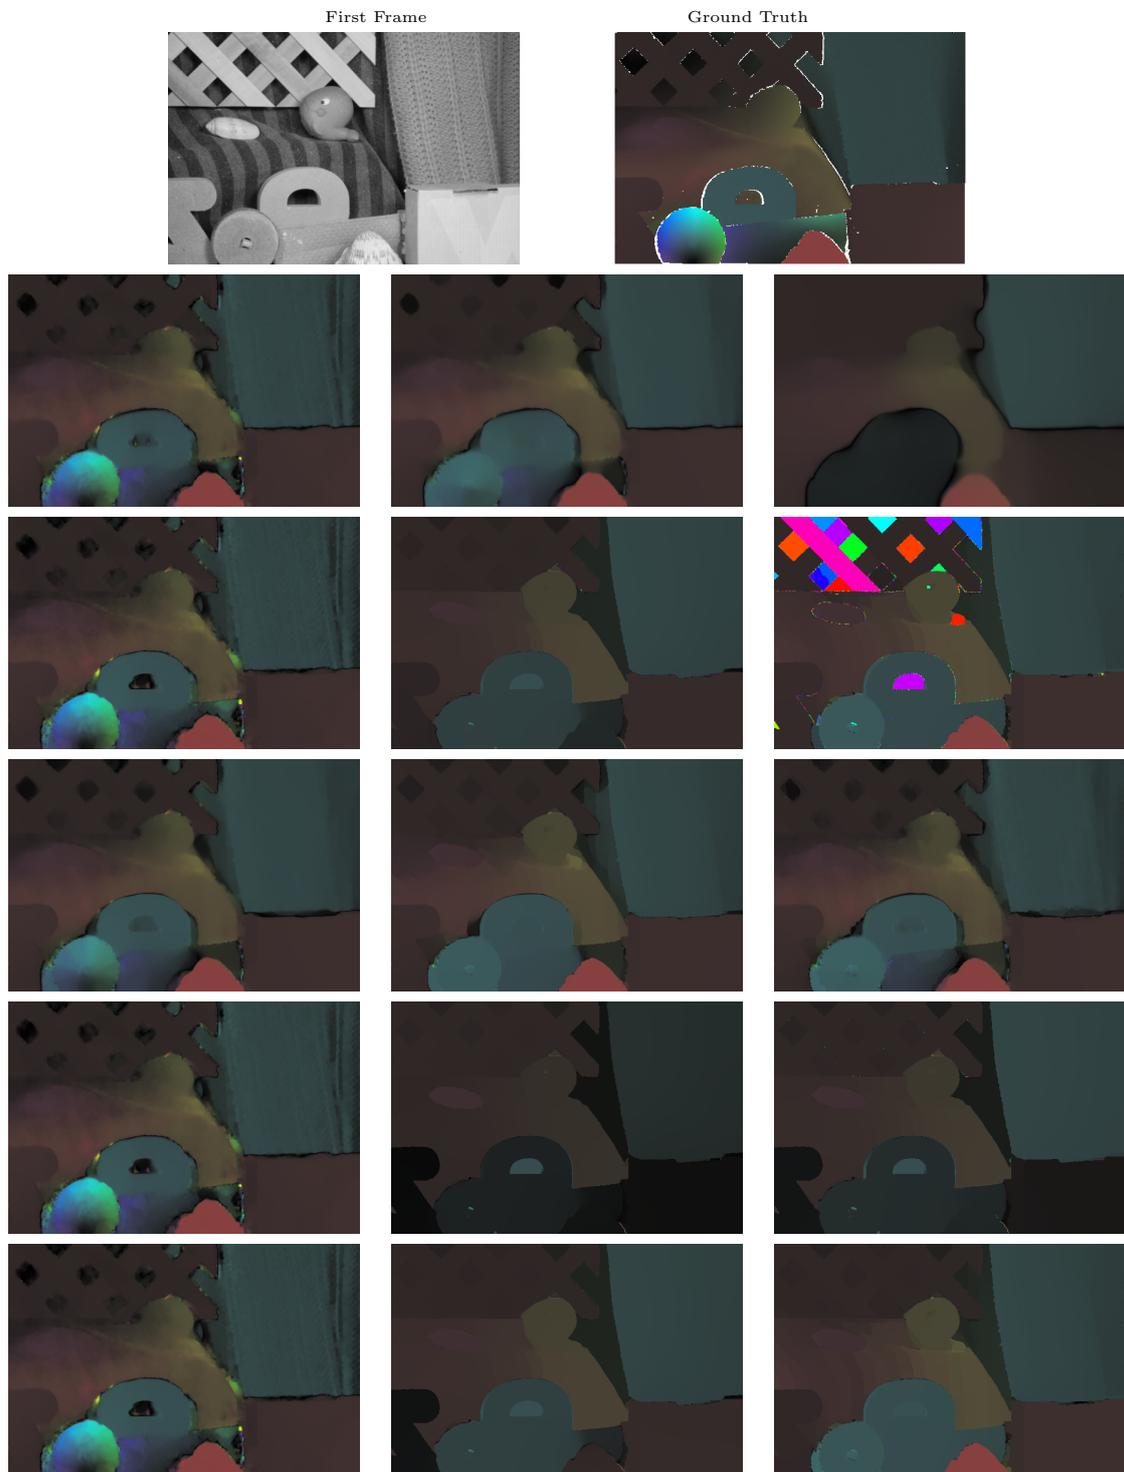


Figure 17: From top to bottom: Solutions of the Brox *et al.*, *Exponential*, *Exponential β* , *Exponential λ -Global* and *Exponential λ -Local* methods, respectively. The columns show the results for increasing values of α and λ .

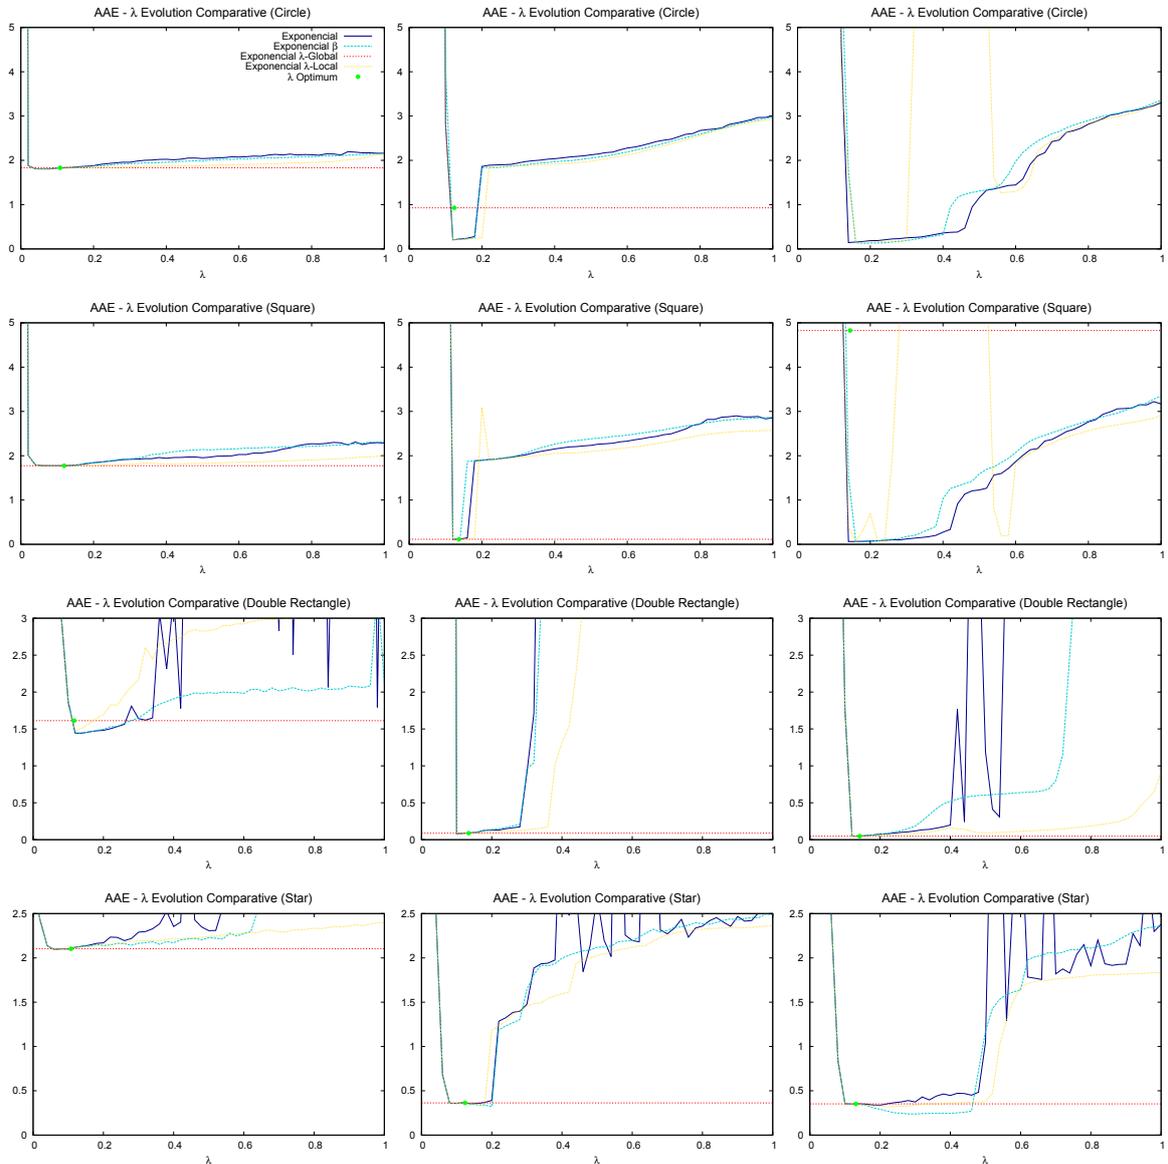


Figure 18: λ evolution for the geometric sequences: the first column shows the results using a small regularization parameter ($\alpha := 50$); the second, a medium value ($\alpha := 250$); and, the third, a large value ($\alpha := 500$).

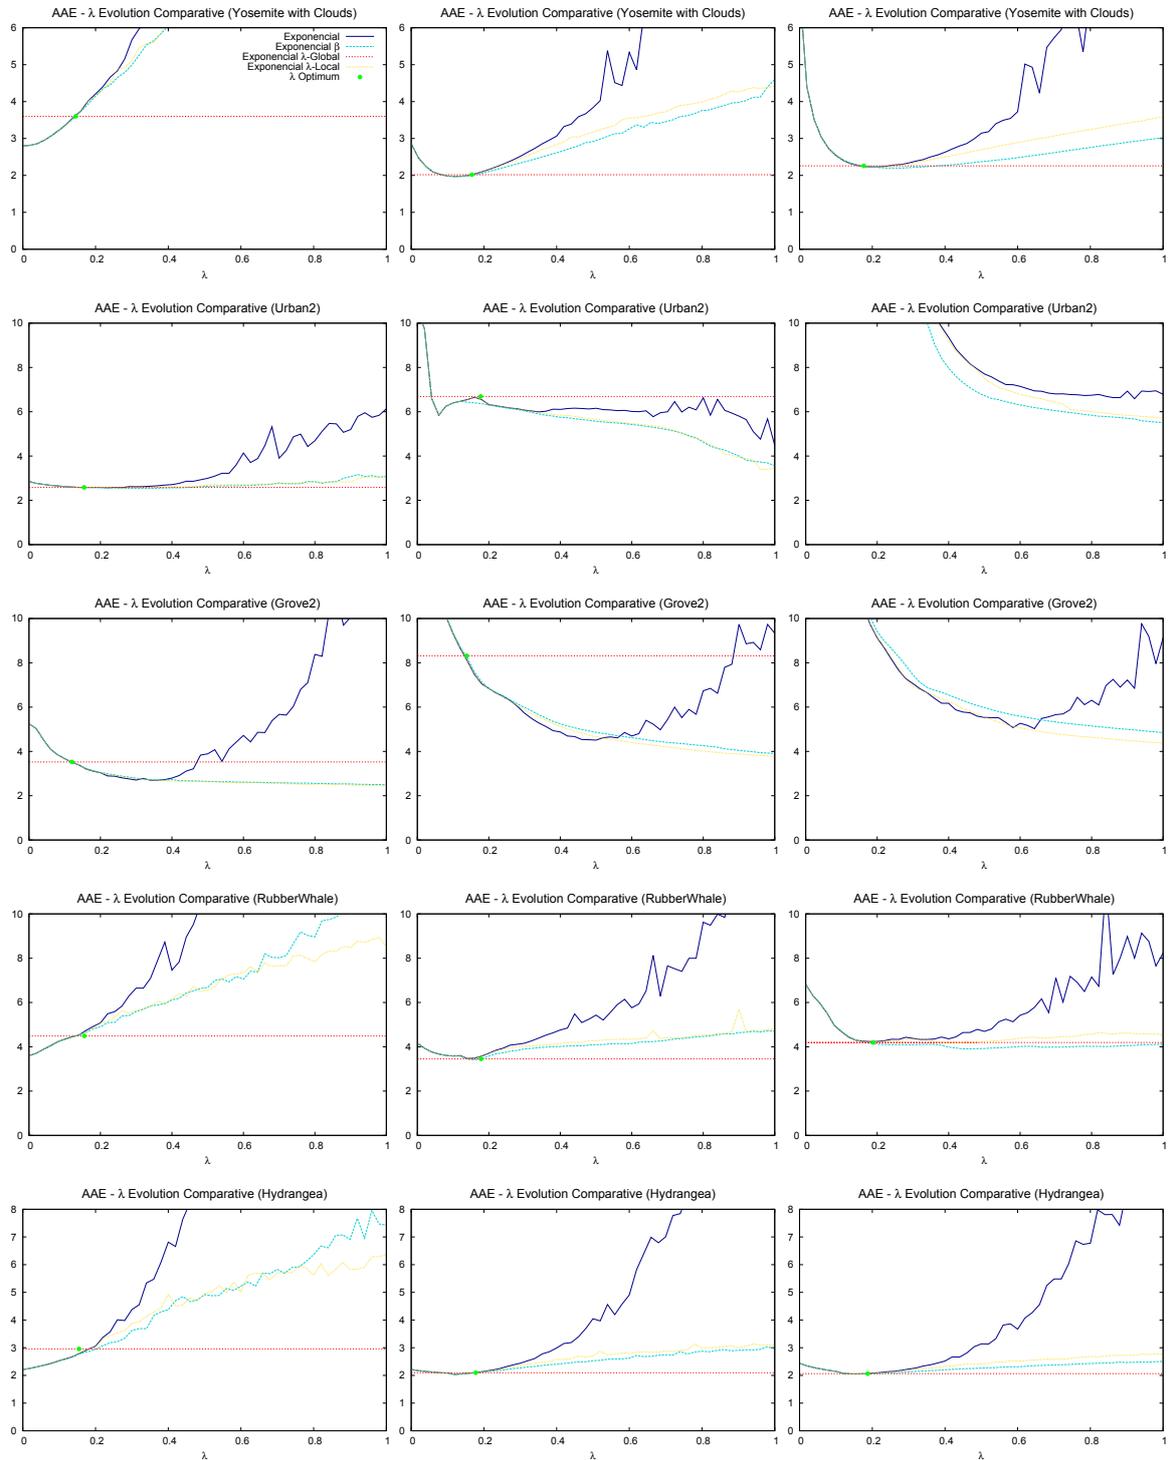


Figure 19: λ evolution for the Middlebury sequences: the first column shows the results using a small regularization parameter ($\alpha := 50$); the second, a medium value ($\alpha := 250$); and, the third, a large value ($\alpha := 500$).

5.5 Numerical Results

In this section we analyze the accuracy of the methods for the best and default parameter configurations. First, we depict in Tables 4 and 5 the AAE results for the best settings, using the geometric and Middlebury sequences, respectively. Correspondingly, Table 6 shows the AAE results of the geometric and Middlebury sequences using a default configuration. These parameters can be seen in Table 7.

In every case, we have looked for the best α , γ and λ that provide the best numerical results. The values for the remaining parameters have been fixed as in Sect. 5.

Table 4: Best AAE results for the geometric sequences and Yosemite with clouds.

		α	γ	λ	AAE
Circle	<i>Brox et al.</i>	8	1	-	2.679^o
	<i>Exponential</i>	58700	40	0.22	0.090^o
	<i>Exponential β</i>	5960	16	0.12	0.046^o
	<i>Exponential λ-Global</i>	56771	97	-	0.086^o
	<i>Exponential λ-Local</i>	5919	45	0.05	0.074^o
Double Rectangle	<i>Brox et al.</i>	14	1	-	4.295^o
	<i>Exponential</i>	60111	27	0.22	0.015^o
	<i>Exponential β</i>	5316	17	0.1	0.049^o
	<i>Exponential λ-Global</i>	67693	51	-	0.016^o
	<i>Exponential λ-Local</i>	59290	30	0.25	0.012^o
Square	<i>Brox et al.</i>	9	0	-	2.801^o
	<i>Exponential</i>	57100	24	0.25	0.016^o
	<i>Exponential β</i>	5655	20	0.2	0.021^o
	<i>Exponential λ-Global</i>	42590	70	-	0.018^o
	<i>Exponential λ-Local</i>	59980	90	0.2	0.017^o
Star	<i>Brox et al.</i>	5	0	-	2.919^o
	<i>Exponential</i>	60990	75	0.22	0.271^o
	<i>Exponential β</i>	5193	35	0.2	0.222^o
	<i>Exponential λ-Global</i>	58900	495	-	0.276^o
	<i>Exponential λ-Local</i>	6929	37	0.25	0.242^o
Yosemite	<i>Brox et al.</i>	147	15	-	2.367^o
	<i>Exponential</i>	221	10	0.1	1.977^o
	<i>Exponential β</i>	250	10	0.15	1.976^o
	<i>Exponential λ-Global</i>	251	9	-	2.003^o
	<i>Exponential λ-Local</i>	218	9	0.12	1.965^o

In the geometric sequences (Table 4), we observe an important improvement with respect to the *Brox et al.* method. However, these results were obtained using a large value of α , which allows to create planar motion regions and correctly deal with occlusions. This type of configuration fits very well with these simple sequences, since the shapes are clearly segmented from the background and the motions are constant inside the regions. In any case, this demonstrates the capabilities of these strategies. The results for Yosemite also improve the *Brox et al.* solution. The AAE decreases in about 17% for the *Exponential λ -Local* approach.

The results for the Middlebury sequences are moderate (see Table 5). The improvement is about 25% for Urban2, using the *Exponential λ -Local* method, and 10% for Venus, using *Exponential β* . In the rest of sequences, the amelioration is not so evident. Looking at Table 6, we see that the results, using default parameters, are similar to the previous results.

Table 5: Best AAE results for the Middlebury benchmark database.

		α	γ	λ	AAE
Urban2	<i>Brox et al.</i>	9	2	-	2.438^o
	<i>Exponential</i>	10	1	0.104	2.362^o
	<i>Exponential β</i>	21	4	0.1	2.422^o
	<i>Exponential λ-Global</i>	10	1	-	2.376^o
	<i>Exponential λ-Local</i>	15	2	0.12	2.400^o
Urban3	<i>Brox et al.</i>	2	1	-	3.544^o
	<i>Exponential</i>	4	1	0.22	2.697^o
	<i>Exponential β</i>	6	1	0.28	2.767^o
	<i>Exponential λ-Global</i>	3	1	-	3.058^o
	<i>Exponential λ-Local</i>	6	0	0.5	2.622^o
RubberWhale	<i>Brox et al.</i>	415	135	-	3.453^o
	<i>Exponential</i>	353	55	0.08	3.293^o
	<i>Exponential β</i>	255	35	0.1	3.307^o
	<i>Exponential λ-Global</i>	210	21	-	3.445^o
	<i>Exponential λ-Local</i>	354	55	0.08	3.291^o
Hydrangea	<i>Brox et al.</i>	125	10	-	2.142^o
	<i>Exponential</i>	156	6	0.12	2.035^o
	<i>Exponential β</i>	292	12	0.12	2.027^o
	<i>Exponential λ-Global</i>	167	5	-	2.058^o
	<i>Exponential λ-Local</i>	201	8	0.12	2.031^o
Grove2	<i>Brox et al.</i>	8	1	-	2.198^o
	<i>Exponential</i>	20	1	0.12	2.109^o
	<i>Exponential β</i>	20	1	0.12	2.111^o
	<i>Exponential λ-Global</i>	18	1	-	2.110^o
	<i>Exponential λ-Local</i>	20	1	0.12	2.109^o
Grove3	<i>Brox et al.</i>	39	5	-	5.972^o
	<i>Exponential</i>	13	1	0.05	5.709^o
	<i>Exponential β</i>	17	1	0.1	5.715^o
	<i>Exponential λ-Global</i>	19	1	-	5.732^o
	<i>Exponential λ-Local</i>	19	1	0.08	5.704^o
Dimetrodon	<i>Brox et al.</i>	690	339	-	1.588^o
	<i>Exponential</i>	280	120	0.08	1.653^o
	<i>Exponential β</i>	210	85	0.08	1.655^o
	<i>Exponential λ-Global</i>	205	57	-	1.804^o
	<i>Exponential λ-Local</i>	280	120	0.08	1.653^o
Venus	<i>Brox et al.</i>	14	4	-	4.434^o
	<i>Exponential</i>	5	0	0.2	4.178^o
	<i>Exponential β</i>	4	0	0.08	3.998^o
	<i>Exponential λ-Global</i>	20	4	-	4.097^o
	<i>Exponential λ-Local</i>	7	1	0.09	4.085^o

Table 6: AAE results using default parameters in the geometric and Middlebury test sequences

	<i>Brox et al.</i>	<i>Exponential</i>	<i>Exponential β</i>	<i>Exponential λ-Global</i>	<i>Exponential λ-Local</i>
Double Rectangle	4.400 ^o	2.188 ^o	2.289 ^o	1.907 ^o	1.917 ^o
Star	3.070 ^o	2.493 ^o	2.468 ^o	2.330 ^o	2.359 ^o
Square	2.939 ^o	2.500 ^o	2.436 ^o	2.218 ^o	2.228 ^o
Circle	2.657 ^o	2.849 ^o	2.762 ^o	2.201 ^o	2.174 ^o
Dimetrodon	1.804 ^o	1.792 ^o	1.792 ^o	2.009 ^o	1.989 ^o
Grove2	2.298 ^o	2.444 ^o	2.458 ^o	2.264 ^o	2.257 ^o
Grove3	6.059 ^o	6.492 ^o	6.500 ^o	6.165 ^o	6.134 ^o
Hydrangea	2.282 ^o	2.491 ^o	2.486 ^o	2.686 ^o	2.546 ^o
RubberWhale	3.727 ^o	3.625 ^o	3.635 ^o	3.911 ^o	3.815 ^o
Urban2	2.491 ^o	2.460 ^o	2.456 ^o	2.388 ^o	2.367 ^o
Urban3	5.633 ^o	4.353 ^o	4.355 ^o	3.860 ^o	4.000 ^o
Venus	4.490 ^o	4.384 ^o	4.376 ^o	4.212 ^o	4.196 ^o

Table 7: Default parameter configuration used in the geometric and Middlebury test sequences.

	α	γ	λ
<i>Brox et al.</i>	17	4	-
<i>Exponential</i>	35	8	0.1
<i>Exponential β</i>	35	8	0.1
<i>Exponential λ-Global</i>	12	2	-
<i>Exponential λ-Local</i>	12	2	0.09

6 Conclusions

In this work we have reviewed several techniques for preserving discontinuities in TV- L^1 optical flow methods. For this purpose, we have introduced a general framework that may harbor many regularization strategies. Our study focuses on the use of decreasing functions for mitigating the diffusion at contours. We have analyzed the ill-posed problem that arises when decreasing functions are used in the regularization term. This has led us to propose two efficient strategies to overcome this situation.

In the experimental results, we have thoroughly studied each of the methods and alternatives. This has helped us to empirically determine the set of default parameters for which the methods work efficiently. We have also compared their behavior and shown that the instability problems are effectively removed with our proposals.

One advantage of these discontinuity preserving 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. We may conclude that using the image and flow information for inhibiting the smoothing is a desired property. Nevertheless, the methods in this survey do not treat this information in an optimum manner.

Eliminating the instability problems, with any of our proposals, turns this method very interesting for real applications. The global approach computes the parameter that better adapts to the maximum gradient in the image. As seen in the experiments, if this gradient is representative of the moving objects, then it will provide the best solution. However, this strategy will provide an average performance if the maximum gradient is much bigger than those associated with the object motions. In any case, it is interesting for automatically determining the exponential parameter and avoiding instabilities.

Some of these inconveniences are solved with the local strategy. It does not compute the parameter automatically but it adapts its value in order to avoid instabilities. This is carried out by examining the local gradient information. In this way, it allows to respect more contours and better adapt to more varying motions. However, when the parameter stands for a very small gradient, then it yields an over-segmented motion field.

As we have seen in the experiments, the use of decreasing functions also allows dealing with occlusions. This can be achieved by increasing the strength of the regularization term, at the same time that we use a high penalization factor at discontinuities. This can be seen in the experiments with the geometric sequences. This effect can also be obtained in more complex sequences, however, this special configuration has a negative impact on other parts of the motion field.

The experimental results have shown the limitations of these strategies. Even though the improvement is important in many sequences, we have seen examples where it does not improve with respect to the basic TV scheme. This is the case, for instance, in the RubberWhale or Hydrangea sequence. These sequences present an important variety of gradients in both the image and the flow.

In future works, we will incorporate more methods to this survey. The methods that steer the diffusion based on the image gradient will probably exhibit a similar behavior. Perhaps, this is the reason why several recent methods have combined the estimation of optical flow with a process for motion segmentation.

Probably, a good method for obtaining precise discontinuities should rely on the combination of the information of the image and flow gradients; but, most importantly, it should be invariant to their magnitudes, i.e., flow discontinuities should be present whenever there exists any motion, independently if they correspond to low contrast or slow moving objects. We will also explore this idea further in the future.

Acknowledgments

This work has been partially supported by the Spanish Ministry of Science and Innovation through the research project TIN2011-25488 and the grant ULPGC011-006.

References

- [1] Luis Álvarez, Rachid Deriche, Théodore Papadopoulo, and Javier Sánchez. Symmetrical dense optical flow estimation with occlusions detection. *International Journal of Computer Vision*, 75(3):371–385, 2007.
- [2] 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.
- [3] Luis Álvarez, Joachim Weickert, and Javier Sánchez. Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision*, 39(1):41–56, 2000.
- [4] G. Aubert, R. Deriche, and P. Kornprobst. Computing optical flow via variational techniques. *SIAM Journal on Applied Mathematics*, 60:156–182, 1999.
- [5] Alper Ayvaci, Michalis Raptis, and Stefano Soatto. Sparse occlusion detection with optical flow. *International Journal of Computer Vision*, 97(3):322–338, May 2012.
- [6] 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.
- [7] Josef Bigun and Gösta H Granlund. Optimal orientation detection of linear symmetry. In *Proc. First Int’l Conf. Comput. Vision*, pages 433–438, 1987.
- [8] Michael J. Black and P. Anandan. A framework for the robust estimation of optical flow. In *Proceedings of Fourth International Conference on Computer Vision*, pages 231–236, 1993.
- [9] 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.
- [10] Michael J. Black, G. Sapiro, D. H. Marimont, and D. Heeger. Robust anisotropic diffusion. *Trans. Img. Proc.*, 7(3):421–432, March 1998.
- [11] 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 *LNCS*, pages 25–36, Prague, Czech Republic, May 2004. Springer.
- [12] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(3):500–513, 2011.
- [13] Andrés Bruhn and Joachim Weickert. Towards ultimate motion estimation: Combining highest accuracy with real-time performance. In *International Conference on Computer Vision (ICCV)*, volume 1, pages 749–755, Washington, DC, USA, October 2005. IEEE Computer Society. <http://dx.doi.org/10.1109/ICCV.2005.240>.

- [14] Andrés Bruhn, Joachim Weickert, and Christoph Schnörr. Lucas /kanade meets horn/schunck: combining local and global optic flow methods. *International Journal of Computer Vision*, 61:211–231, February 2005.
- [15] 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.
- [16] Berthold K. P. Horn and Brian G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [17] Étienne Mémin and Patrick 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.
- [18] N. Monzón, J. Sánchez, and A. 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.
- [19] Nelson Monzón, Javier Sánchez, and Agustín Salgado. Efficient mechanism for discontinuity preserving in optical flow methods. In Abderrahim Elmoataz, Olivier Lezoray, Fathallah Nouboud, and Driss Mammass, editors, *Image and Signal Processing*, volume 8509 of *Lecture Notes in Computer Science*, pages 425–432. Springer International Publishing, 2014.
- [20] 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, September 1986.
- [21] Pietro Perona and Jitendra Malick. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:629–629, July 1990.
- [22] 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 third European conference on Computer Vision (Vol. II)*, ECCV '94, pages 295–304, Secaucus, NJ, USA, 1994. Springer-Verlag New York, Inc.
- [23] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, November 1992.
- [24] Javier Sánchez, Agustín Salgado, and Nelson Monzón. Optical flow estimation with consistent spatio-temporal coherence models. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, pages 366–370. Institute for Systems and Technologies of Information, Control and Communication, 2013.
- [25] Javier Sánchez, Agustín Salgado, and Nelson Monzón. Preserving Accurate Motion Countours with Reliable Parameter Selection. pages 209–213. IEEE International Conference on Image Processing (ICIP2014), 2014.
- [26] Deqing Sun, J. P. Lewis, and Michael J. Black. Learning optical flow. In *European Conference on Computer Vision (ECCV)*, pages 83–97, 2008.
- [27] Deqing Sun, Stefan Roth, and Michael J. Black. Secrets of optical flow estimation and their principles. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 0, pages 2432–2439, Los Alamitos, CA, USA, 2010. IEEE Computer Society.

- [28] Deqing Sun, E. Sudderth, and Michael J. Black. Layered Image Motion with Explicit Occlusions, Temporal Consistency, and Depth Ordering. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, volume 23, pages 2226–2234, 2010.
- [29] Deqing Sun, E. Sudderth, and Michael J. Black. Layered segmentation and optical flow estimation over time. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1768–1775. IEEE, 2012.
- [30] Javier Sánchez Pérez, 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>.
- [31] Javier Sánchez Pérez, Nelson Monzón López, and Agustín Salgado de la Nuez. Robust Optical Flow Estimation. *Image Processing On Line*, 3:252–270, 2013.
- [32] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *ICCV*, pages 839–846, 1998.
- [33] Markus Unger, Manuel Werlberger, Thomas Pock, and Horst Bischof. Joint motion estimation and segmentation of complex scenes with label costs and occlusion modeling. In *CVPR*, pages 1878–1885. IEEE, 2012.
- [34] 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.
- [35] Andreas Wedel, Thomas Pock, Christopher Zach, Horst Bischof, and Daniel Cremers. An improved algorithm for tv-l1 optical flow. In *Statistical and Geometrical Approaches to Visual Motion Analysis*, pages 23–45. Springer Berlin Heidelberg, 2009.
- [36] Joachim Weickert. A review of nonlinear diffusion filtering. In *Proceedings of the First International Conference on Scale-Space Theory in Computer Vision*, SCALE-SPACE '97, pages 3–28, London, UK, UK, 1997. Springer-Verlag.
- [37] Joachim Weickert. *Anisotropic diffusion in image processing*, volume 1. Teubner Stuttgart, 1998.
- [38] Joachim Weickert and Christoph Schnörr. A theoretical framework for convex regularizers in pde-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, 2001.
- [39] M. Werlberger, T. Pock, and H. Bischof. Motion estimation with non-local total variation regularization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2464–2471, San Francisco, CA, USA, June 2010.
- [40] Manuel Werlberger, Werner Trobin, Thomas Pock, Andreas Wedel, Daniel Cremers, and Horst Bischof. Anisotropic huber-l1 optical flow. In *BMVC*. British Machine Vision Association, 2009.
- [41] Manuel Werlberger, Werner Trobin, Thomas Pock, Andreas Wedel, Daniel Cremers, and Horst Bischof. Anisotropic huber-l1 optical flow. In *Proceedings of the British Machine Vision Conference (BMVC)*, London, UK, September 2009. to appear.
- [42] Jiangjian Xiao, Hui Cheng, Harpreet Sawhney, Cen Rao, and Michael Isardi. Bilateral filtering-based optical flow estimation with occlusion detection. In *European Conference on Computer Vision*, volume I, pages 211–224, 2006.

- [43] 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.
- [44] Kuk-Jin Yoon and In-So Kweon. Adaptive support-weight approach for correspondence search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):650–656, 2006.
- [45] 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, Berlin, Heidelberg, 2007.
- [46] Henning Zimmer, Andrés Bruhn, and Joachim Weickert. Optic flow in harmony. *International Journal of Computer Vision*, 93(3):368–388, July 2011.
- [47] Henning Zimmer, Andrés Bruhn, Joachim Weickert, Levi Valgaerts, Agustín Salgado, Bodo Rosenhahn, and Hans-Peter Seidel. Complementary optic flow. In *Proceedings of the 7th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition*, EMMCVPR '09, pages 207–220, Berlin, Heidelberg, 2009. Springer-Verlag.



Centro de Tecnologías de la Imagen
Universidad de Las Palmas de Gran Canaria
<http://www.ctim.es>