

A Rotation-Invariant Regularization Term for Optical Flow Related Problems

Roberto P.Palomares, Gloria Haro and Coloma Ballester

Pompeu Fabra University, DTIC, 08018

Abstract. This paper proposes a new regularization term for optical flow related problems. The proposed regularizer properly handles rotation movements and it also produces good smoothness conditions on the flow field while preserving discontinuities. We also present a dual formulation of the new term that turns the minimization problem into a saddle-point problem that can be solved using a primal-dual algorithm. The performance of the new regularizer has been compared against the Total Variation (TV) in three different problems: optical flow estimation, optical flow inpainting, and optical flow completion from sparse samples. In the three situations the new regularizer improves the results obtained with the TV as a smoothing term.

1 Introduction

The optical flow problem, also called motion estimation, is a key problem in computer vision. Its aim is to recover the apparent motion field of two consecutive frames. It is a classic ill-posed problem, related to the aperture problem. Thus, additional constraints are made in order to regularize the problem and single out a solution. Horn and Schunk (HS) [1] proposed to compute the flow field as a variational problem, where the searched vector field u corresponds to the minimum of an energy functional. Their energy model has two parts: a fidelity data term imposing the brightness constancy assumption, and a regularization term. The HS model incorporates quadratic terms in both parts (regularizers and fidelity data term), so it does not allow discontinuities in the flow field. Many models, during the last thirty years, have been developed to avoid the previous problem. Several robust estimators have replaced the original quadratic norm [2], either in the data term [3] or regularization term. For the regularization term, anisotropic diffusion (image-driven) [4–7], second order smoothness assumptions [8] and isotropic diffusion as the Total Variation (TV) [9, 10] have been proposed. For static scenes, the epipolar geometry can be used as a weak prior [11] or to define an over-parameterized optical flow whose parameters are regularized instead of the flow [12].

Optical flow regularization terms are also useful for motion inpainting, where the optical flow in the missing region is completed by looking for a smooth flow that matches the motion in the known region. Motion inpainting has been applied for completing the optical flow in regions of low confidence due to occlusions,

transparencies, etc [13]. It has also been used for video stabilization [14] and for video inpainting [15].

Total Variation is one of the most used regularizers due to its nice properties such as its ability to recover the image discontinuities, and the existence of efficient and robust numerical schemes with guaranteed convergence. However, it suffers from staircasing effects and does not properly handle rotations.

Our contribution in this paper is a new regularization term which is invariant to infinitesimal flow rotations and which is able to preserve discontinuities/jumps in the flow field. The proposed regularizer does not increase the energy of the functional when the flow field is a rotation movement and, moreover, it keeps some nice properties of the TV, providing a smooth flow field that preserves discontinuities. Real scenes usually contain rotation movements, sometimes at an infinitesimal level, and therefore the use of the proposed term could help to obtain accurate and realistic optical flows. It has been tested in three different optical flow related problems: motion estimation, motion inpainting, and motion reconstruction from sparse samples. In the three kinds of applications it outperforms the results obtained by the TV.

The paper is organized as follows: In Sect. 2 we give the theoretical motivation for the new regularization term. In Sect. 3 we describe a variational model for the optical flow estimation and how to minimize its energy. In Sect. 4 we provide details of the numerical implementation. Some experiments are presented in Sect. 5. Finally, in Sect. 6 we present the main conclusions.

2 Motivation for the Regularizer

Let $U : \mathcal{R}^3 \rightarrow \mathcal{R}^3$ be the velocity vector field of a fluid in \mathcal{R}^3 . For U smooth enough, we consider a first-order Taylor approximation of U ,

$$U(\mathbf{x} + \mathbf{h}) = U(\mathbf{x}) + J_{\mathbf{x}}U(\mathbf{x})\mathbf{h} + o(|\mathbf{h}|) \quad \text{for } \mathbf{h} \rightarrow 0 \quad (1)$$

where $\mathbf{x} = (x, y, z)$, $\mathbf{h} = (h_1, h_2, h_3)$ and $J_{\mathbf{x}}U(\mathbf{x})$ is the Jacobian matrix at the point \mathbf{x} . For \mathbf{h} small enough, let us approximate

$$U(\mathbf{x} + \mathbf{h}) \approx U(\mathbf{x}) + J_{\mathbf{x}}U(\mathbf{x})\mathbf{h}. \quad (2)$$

The first term at the right-hand side of the above equation represents a translation. The second term provides information on the rotation and deformation movement (scaling and shearing). A matrix can be decomposed in two parts, the symmetric and the antisymmetric part. In our case, let us denote by $D = \frac{1}{2}(J_{\mathbf{x}}U(\mathbf{x}) + J_{\mathbf{x}}U(\mathbf{x})^T)$ the symmetric part of $J_{\mathbf{x}}U(\mathbf{x})$ and by $C = \frac{1}{2}(J_{\mathbf{x}}U(\mathbf{x}) - J_{\mathbf{x}}U(\mathbf{x})^T)$ the antisymmetric one. The symmetric part measures the area change ratio (divergence) while the antisymmetric part describes the infinitesimal rotation of the vector field (curl) [16]. It is easy to see that the previous argument is valid for \mathcal{R}^2 . With this idea in mind, we propose to use as a motion regularizer a measure that does not penalize infinitesimal rotations and,

as we will see, still produces good smoothness conditions and is able to preserve discontinuities/jumps in the flow field.

We propose the following regularization term

$$E_{D_u}(u) \equiv \int_{\Omega} \left\| \frac{1}{2} (Du + Du^T) \right\|_F dx \quad (3)$$

where, for smooth functions, $Du = \begin{pmatrix} u_{1x} & u_{1y} \\ u_{2x} & u_{2y} \end{pmatrix}$, being $u = (u_1, u_2) : \Omega \rightarrow R^2$ the optical flow; $\|\cdot\|_F$ denotes the Frobenius norm. Let us recall that the Frobenius norm can be defined as $\|A\|_F = \langle A, A \rangle_F^{1/2}$, where $\langle A, B \rangle_F = \text{Trace}(A^t B)$.

For $u \in \mathcal{L}^1(\Omega, \mathcal{R}^2)$, $E_{D_u}(u)$ can be defined by its dual representation as

$$E_{D_u}(u) = \sup_{\|\xi\|_F \leq 1} \int_{\Omega} \langle u_1, \text{div}(\xi_{11}, \xi_{12}) \rangle + \langle u_2, \text{div}(\xi_{12}, \xi_{22}) \rangle dx, \quad (4)$$

where the supremum is taken over all $\xi \in C_c^1(\Omega; \text{Sym}(R^{2 \times 2}))$ s.t. pointwise $\|\xi(x)\|_F \leq 1$. We can restrict ξ to be a symmetric matrix due to the following Lemma.

Lemma 1. *If ξ and C are symmetric and antisymmetric 2×2 matrices, respectively, then $\langle C, \xi \rangle_F = 0$.*

For u , smooth enough, (4) can be written as

$$E_{D_u}(u) = \sup_{\|\xi\|_F \leq 1} \int_{\Omega} \left\langle \frac{1}{2} (Du + Du^T), \xi \right\rangle_F dx. \quad (5)$$

The proposed regularization term is related to the symmetric part of the following term based on the Frobenius norm which, for u smooth enough and using dual variables, can be written as

$$\int_{\Omega} \|Du\|_F dx = \sup_{\|\xi\|_F \leq 1} \int_{\Omega} \langle Du, \xi \rangle_F dx, \quad (6)$$

with $\xi \in C_c^1(\Omega; R^{2 \times 2})$. This term has been called $\text{TV}_{\ell^2}(u)$ in the paper [17].

Let us remark that, if u is a rotation flow such as $u(x_1, x_2) = (-x_2, x_1)$, then $E_{D_u}(u)$ vanishes. Moreover, as in [17], for a translation of an object in a static background, our term handles different directions of translation equally. In a different context to ours, the authors in [18] use the idea of regularizing the symmetric part of the deformation.

3 Optical Flow Functional

In this section we present the proposed model to estimate the optical flow, and how to minimize its energy.

3.1 The Model

To show the benefits of our term, we build up from a well-known optical flow model [10] that uses TV as the regularization term and an L^1 data term

$$\min_u \int_{\Omega} (|\nabla u_1| + |\nabla u_2| + \lambda |I_1(x+u) - I_0|) dx, \quad (7)$$

where I_0, I_1 are two consecutive frames and $u = (u_1, u_2)$ is the estimated optical flow between them.

We replace the Total Variation as regularizing term by the new flow-rotation-invariant-regularizer defined by (3). Then, in order to compute the optical flow $u = (u_1, u_2)$ between two consecutive frames I_0 and I_1 of a video sequence, we propose to minimize the following energy

$$\int_{\Omega} \left\| \frac{1}{2} (Du + Du^T) \right\|_F dx + \lambda \int_{\Omega} |I_1(x+u) - I_0(x)| dx. \quad (8)$$

The use of L^1 type-norm measures has proven a good performance in front of L^2 norms. Unfortunately, it increases the difficulty when minimizing the functional due to its non differentiability. We introduce an auxiliary variable v representing the optical flow, as in [10], and we penalize its deviation from u . Thus, we minimize the energy

$$\int_{\Omega} \left\| \frac{1}{2} (Du + Du^T) \right\|_F dx + \lambda \int_{\Omega} |I_1(x+v) - I_0(x)| dx + \frac{1}{2\theta} \int_{\Omega} (u-v)^2 dx \quad (9)$$

with respect to u and v .

Notice that by minimizing this energy we do not impose regularization of the skew symmetric part of the Jacobian. Indeed, given any function f of Bounded Variation, the symmetric part of the Jacobian of the deformation $(f(x_2), -f(x_1))$ vanishes. Therefore, in order to prevent an irregular behavior of this part of the flow field, one can add to the functional, e.g., an additional classical TV term. In any case, the experiments show that the proposed term alone keeps some nice properties of the TV providing a smooth flow field that preserves discontinuities.

This energy can be minimized by an alternating minimization procedure. On the other hand, to minimize (9) with respect to v , we linearize $I_1(x+v)$ around a given optical flow map u_0 using first order Taylor approximation. Therefore, the expression in the fidelity data term can be approximated by

$$\rho(v) := I_1(x+u_0) + \langle \nabla I_1(x+u_0), (v-u_0) \rangle - I_0(x). \quad (10)$$

Then, our functional (9) becomes

$$\int_{\Omega} \left\| \frac{1}{2} (Du + Du^T) \right\|_F dx + \lambda \int_{\Omega} |\rho(v)| dx + \frac{1}{2\theta} \int_{\Omega} (u-v)^2 dx. \quad (11)$$

3.2 Minimizing the Energy

The energy (11) can be minimized by alternating steps updating either u or v in every iteration. The minimization procedure is

1. For v fixed, minimize (11) with respect to u .

Chambolle and Pock proposed a primal-dual algorithm to minimize the ROF model [19]. It is based on a dual formulation of the TV [20]. Following the ideas of [19], we reformulate the optical flow model as a min-max problem incorporating dual variables. Then, our minimization problem (11) can be solved as a saddle-point problem. For v fixed, we solve

$$\min_u \max_{\xi} \int_{\Omega} \langle \frac{1}{2} (Du + Du^T), \xi \rangle dx + \int_{\Omega} \frac{1}{2\theta} (u - v)^2 dx \quad (12)$$

where the dual variables are $\xi = \begin{pmatrix} \xi_{11} & \xi_{12} \\ \xi_{12} & \xi_{22} \end{pmatrix}$ and satisfy $\|\xi\|_F \leq 1$. Let us notice that, using previous Lemma 1, we can restrict ξ to be a symmetric matrix.

Proposition 1. *The solution of (12) is given by the following iterative scheme*

$$\xi_{11}^{n+1} = \frac{\xi_{11}^n + \tau \bar{u}_{1x}^n}{\max(1, \|\xi\|_2)} \quad (13)$$

$$\xi_{22}^{n+1} = \frac{\xi_{22}^n + \tau \bar{u}_{2y}^n}{\max(1, \|\xi\|_2)} \quad (14)$$

$$\xi_{12}^{n+1} = \frac{\xi_{12}^n + \frac{\tau}{2} (\bar{u}_{1y}^n + \bar{u}_{2x}^n)}{\max(1, \|\xi\|_2)} \quad (15)$$

$$u_1^{n+1} = u_1^n - \sigma \left(\frac{(u_1^n - v_1)}{\theta} - \operatorname{div}(\xi_{11}^n, \xi_{12}^n) \right) \quad (16)$$

$$u_2^{n+1} = u_2^n - \sigma \left(\frac{(u_2^n - v_2)}{\theta} - \operatorname{div}(\xi_{12}^n, \xi_{22}^n) \right) \quad (17)$$

$$\bar{u}_1^{n+1} = 2u_1^{n+1} - u_1^n \quad (18)$$

$$\bar{u}_2^{n+1} = 2u_2^{n+1} - u_2^n \quad (19)$$

where u is the primal variable and ξ is the dual variable.

2. For u fixed, minimize with respect to v the following functional

$$\int_{\Omega} \lambda |\rho(v)| + \int_{\Omega} \frac{1}{2\theta} (u - v)^2 dx. \quad (20)$$

Since (20) does not depend on spatial derivatives on v , a simple thresholding step gives us an explicit solution [10].

Proposition 2. *The minimum of (20) with respect to v is*

$$v = u + \begin{cases} \lambda \theta \nabla I_1 & \text{if } \rho(u) < -\lambda \theta |\nabla I_1|^2 \\ -\lambda \theta \nabla I_1 & \text{if } \rho(u) > \lambda \theta |\nabla I_1|^2 \\ -\rho(u) \frac{\nabla I_1}{|\nabla I_1|^2} & \text{if } |\rho(u)| \leq \lambda \theta |\nabla I_1|^2 \end{cases} \quad (21)$$

The whole minimization algorithm is presented in Algorithm 1.




Fig. 1. Results on the MPI Sintel training set “Bandage-2” (clean version). From top to bottom and from left to right: (a) First frame, (b) optical flow ground truth, (c) estimated optical flow using the TV- L^1 method, (d) estimated optical flow using an L^1 data term and the proposed regularizer.

4 Implementation

The minimization of (11) to estimate the optical flow is embedded into a coarse-to-fine multi-level approach in order to be able to deal with large motion fields. The numerical algorithm is summarized in Algorithm 1. The image gradient is computed using a five-point stencil as in [5]. Image warpings use bicubic interpolation. Our code is written in C. To do a fair comparison against the TV- L^1 optical flow model of [10], we took the implementation of [21], which is also written in C, and changed their numerical scheme by a primal-dual approach as in our case. The algorithm parameters are initialised with the same default settings. Both time-steps are set to $\tau = \sigma = 0.125$ to ensure the convergence. As stopping criterion both optical flow methods use the infinite-norm between two consecutive values of u with a threshold of 0.01. The coupling parameter θ is equal to 0.3. Input images have been normalized between $[0, 1]$. The fidelity data term weight λ is set to 40. Let us remark that the parameters value have been fixed to ensure a good performance for all the sequences of the Middlebury Dataset. All the experiments use the previous parameters even if the images come from another database.

5 Experiments

In this section we provide two sets of experiments. The first one is designed to verify the good performance of the new regularization term over different types of movements. The second one shows the properties of the presented term to recover rotation movements. We use both real and synthetic images, and two databases: the Middlebury flow benchmark [22] and some images from the MPI Sintel Flow Dataset [23].

Input : Two consecutive frames I_1, I_2
Output: Flow field \mathbf{u}

Compute down-scaled images I_1^s, I_2^s for $s = 1, \dots, N_{\text{scales}}$;
Initialize $\mathbf{u}^{N_{\text{scales}}} = \xi^{N_{\text{scales}}} = \mathbf{v}^{N_{\text{scales}}} = 0$;
for $s \leftarrow N_{\text{scales}}$ **to** 1 **do**
 for $w \leftarrow 1$ **to** N_{warps} **do**
 Compute $I_1^s(x + v_0(x)), I_1^{sx}(x + v_0(x)), I_1^{sy}(x + v_0(x))$;
 using bicubic interpolation;
 while $n < N_{\text{max}}$ **or** $tol < error$ **do**
 Compute v via Prop. 2;
 Compute ξ and u via Prop. 1;
 end
 end
end

Algorithm 1: Coarse-to-fine multi-level approach to compute optical flow



Fig. 2. Results for a pure rotation movement (3 degree anti-clockwise). *From left to right:* (a) First frame. (b) Optical flow ground truth. (c) Estimated optical flow using the TV regularizer. (d) Estimated optical flow using the proposed regularizer.

5.1 Global accuracy

In this section we compute the optical flow from a set of images from standard datasets.

Middlebury DataSet - It contains a training set where ground truth is available. We use it to verify the accuracy of our regularization term against the well-established TV term. In order to compare our regularizer to other high-order regularization methods, we have implemented TGV_2 (see [24]). The parameters α_0 and α_1 have been chosen to ensure a good performance for all the sequences of the Middlebury Dataset. Table 1 shows how our method improves the optical flow maps for almost all the images by measuring the Average Angular Error (AEE) and Average End-point Error (EPE).

MPI Sintel Flow DataSet - MPI Sintel dataset has 23 training sequences. For every frame, there are two different images, “clean” and “final”. The difference is that the second set adds complexity to the first one by incorporating atmospheric effects, depth of field blur, motion blur, color correction and other details. We evaluate the frames 32-38 from the training set “Bandage-2” (see Fig. 1). We have chosen this sequence because the dragon’s movement is almost a pure rotation. The quantitative results are in Table 2.

Table 1. Error measures in the Middlebury dataset with public ground truth flow.

<i>Middlebury</i>	Dim.	Hyd.	Rub.	Gro2	Gro3	Urb2	Urb3	Ven.
EPE-TV	0.1537	0.2286	0.1916	0.1496	0.6808	0.3709	0.6034	0.3563
EPE-TGV	0.1405	0.2913	0.1868	0.1471	0.6255	0.3832	0.5639	0.3263
EPE-Ours	0.1393	0.2598	0.1716	0.1438	0.6113	0.3518	0.4815	0.3595
AAE-TV	2.8458°	2.6528°	6.0472°	2.2356°	6.5780°	2.9230°	5.3784°	5.9934°
AAE-TGV	2.6014°	3.3339°	6.0431°	2.1670°	6.0920°	3.0723°	5.4091°	5.2621°
AAE-Ours	2.5336°	3.0934°	5.6276°	2.1302°	5.9385°	2.7349°	4.7173°	6.1516°

Table 2. Global error measures in images of the sequence Bandage-2 from MPI. The first and second set of results correspond, respectively, to the “clean” and “final” frames.

<i>Bandage – 2</i>	I-32	I-33	I-34	I-35	I-36	I-37	I-38
EPE-TV	0.3012	0.2635	0.2281	0.1969	0.1819	0.1831	0.1744
EPE-Ours	0.2875	0.2528	0.2200	0.1871	0.1738	0.1744	0.1646
AAE-TV	4.1854°	4.3195°	4.4798°	4.6716°	4.5370°	4.5531°	4.4934°
AAE-Ours	3.9796°	4.0286°	4.2245°	4.3852°	4.3138°	4.3766°	4.2783°
EPE-TV	0.4871	0.4125	0.3514	0.3081	0.2788	0.2538	0.2304
EPE-Ours	0.4793	0.3833	0.3380	0.2999	0.2723	0.2464	0.2208
AAE-TV	6.7334°	6.6018°	6.7360°	6.9075°	6.6736°	6.3551°	6.0901°
AAE-Ours	6.5131°	6.0037°	6.4608°	6.5416°	6.5060°	6.1108°	5.8036°

Table 3. Local error measures in images of Bandage-2 from MPI dataset. The first and second set of results correspond, respectively, to the clean and final frames.

<i>Bandage – 2</i>	I-32.	I-33.	I-34.	I-35	I-36	I-37	I-38.
EPE-TV	0.6651	0.5622	0.4984	0.4209	0.4026	0.4326	0.4052
EPE-Ours	0.6095	0.5213	0.4651	0.3847	0.3631	0.3936	0.3693
AAE-TV	5.4301°	5.6568°	6.1474°	6.4521°	6.7121°	6.8091°	6.6867°
AAE-Ours	5.1992°	5.4078°	5.9587°	6.1164°	6.2567°	6.4884°	6.3163°
EPE-TV	0.9592	0.7842	0.7031	0.6330	0.5383	0.5389	0.4877
EPE-Ours	0.9661	0.7426	0.6839	0.6086	0.5209	0.5157	0.4575
AAE-TV	8.3658°	8.9096°	9.5192°	10.0930°	9.1798°	8.8827°	8.1718°
AAE-Ours	8.6947°	8.7647°	9.6918°	9.8830°	9.1128°	8.6662°	7.8128°




Fig. 3. Color representation of the level sets of the L^2 norm of the optical flow estimations shown in Fig 2. *From left to right:* (a) ground truth, (b), TV regularizer, and (c) proposed regularized.

5.2 Rotation movements

Our regularizer shows good performance in general movements, but it is specifically designed to appropriately handle rotation movements. This section contains several experiments to demonstrate the performance of our term against the TV when dealing with rotation movements in two different types of applications: motion estimation and motion inpainting. It is easy to observe how in all the figures the rotation movements look more realistic with the new regularizer and present smoother transitions as well.

Optical flow estimation - The following results compare the optical flow estimated by the TV- L^1 model [10] against the estimation obtained by solving the functional (11), i.e. an L^1 data term plus our proposed optical flow regularizer.

First, we evaluate our optical flow algorithm on a synthetic pair of images. The purpose of this synthetic sequence is to test the robustness to rotation movements. Fig. 2 shows a pure anti-clockwise rotation of 3 degrees. In contrast to the TV- L^1 model, our method does not present piecewise constant zones and it is more accurate. Table 4 shows the Average Angular Error (AEE) and Average End-point Error (EPE) for both methods (TV- L^1 and our approach). Fig. 3 shows a color representation of the L^2 norm of the optical flow and some of its level lines. The level lines of the ground truth flow are circles since the movement is a pure rotation. The level lines of the intensity of the estimated optical flow with the proposed regularizer are much closer to circles than the ones obtained from the flow estimated with the TV regularizer.

For the two databases with ground truth, MPI and Middlebury, we evaluate some regions that contain almost a pure rotation. The column “Local” of Table 5 shows the local errors for the RubberWhale sequence and Table 3 shows the local errors around the dragon’s head in the MPI sequence. Fig. 4 shows how in the Army sequence the rotation movement obtained with our regularizer is more realistic than the one obtained with the TV. Fig. 5 shows how our method better recovers the contour of the dragon’s head and it is possible to observe little details as the eyes or the mouth.




Fig. 4. Army sequence from the Middlebury dataset with a rotation movement. *From left to right* (a) Optical flow estimated with the TV regularizer. (b) Optical flow estimated with the proposed regularizer (c) Details of the estimations using the TV (top) and the proposed regularizer (bottom).

In Fig. 7 we present another real example which consists of a video sequence of a double windmill captured by a camera with difficult light conditions. Each windmill rotates in opposite directions as is shown in the first row of Fig. 7. The last two rows of the Fig. 7 show the difference between the smooth transition of our term (third row), which agrees with the almost radially symmetric movement of a windmill, as opposite to the piecewise constant zones that appear due to the TV term (second row). Fig. 6 shows this effect: it displays a color representation of the L^2 norm of the optical flow at each point and some of its level lines. Being a rotation movement, the level lines of the motion intensity should be circles, as in the synthetic example shown in Fig. 3. As seen in Fig. 6, the level lines of the optical flow norm obtained with our regularizer are closer to circles than the ones obtained with the TV, meaning that our regularization term recovers the rotation movement in a more realistic way than the TV.

Table 4. Error measures for the synthetic images in Fig. 2 of a pure rotation.

<i>Pure Rotation</i>	Synthetic
EPE-TV	0.0204
EPE-Ours	0.0122
AAE-TV	0.7108°
AAE-Ours	0.4351°

Optical flow inpainting and interpolation from sparse samples - We have designed two proof of concept experiments to illustrate that the proposed regularization term is able to properly recover rotation movements. In particular, we show how it performs in reconstructing missing optical flow data and compare it to the TV. Two different cases of missing data have been addressed: data missing in a hole (we denote it as *optical flow inpainting*), and data missing along the whole image with the exception from some sparse locations (*optical*

Table 5. Local optical flow estimation error (column “Local”), inpainting error (column “Inpainting”) and interpolation error (column “Interpolation”) for RubberWhale, from Middlebury dataset. Local errors mean that we only measure the error in a square around the the toy wheel and the dragon’s head respectively.

<i>RubberWhale</i>	Local	Inpainting	Interpolation
EPE-TV	0.8571	0.0611	0.4669
EPE-Ours	0.4374	0.0559	0.3088
AAE-TV	11.4892°	1.8610°	8.4497°
AAE-Ours	5.9242°	1.8118°	6.9612°

flow interpolation from sparse samples). In both cases we complete the missing data by a diffusion of the known values. The resulting diffusion depends on the smoothness term used in the minimization problem for recovering the data.

If we use the proposed regularizer, the minimization problem is

$$\min_u \int_W \left\| \frac{1}{2} (Du + Du^T) \right\|_F dx \quad \text{with } u|_{\partial W} = u_0|_{\partial W}, \quad (22)$$

where W is the missing data domain and u_0 is the known flow field in the complement of the set W . Also, the minimization problem based on the TV is

$$\min_u \int_W (|\nabla u_1| + |\nabla u_2|) dx \quad \text{with } u|_{\partial W} = u_0|_{\partial W}. \quad (23)$$

W is defined in the following ways for the two kind of applications:

- **Optical flow inpainting** - Squared patches are removed from the flow field (and u is initialized to zero in these regions). The motion in the known regions, u_0 , is the optical flow ground truth. Regions where the ground truth is unknown also form part of the inpainting domain W , together with the squared patches.
- **Optical flow interpolation from sparse samples** - We remove the optical flow values in the 95% of the pixels (uniformly distributed along the image). Therefore, W is made of this 95% of pixels. Again, u_0 is defined from the optical flow ground truth.

For testing both regularization terms we use the RubberWhale sequence from the Middlebury dataset [22] because it contains almost a pure rotation. For both cases of missing data we solve the minimization problems (22) and (23) using a primal-dual approach. The resulting numerical scheme for (22) is the one explained in Prop. 1 without the terms depending on the parameter θ and with a final step that sets the boundary conditions (known optical flow in ∂W). Fig. 8, Fig. 9 and Table 5 refer to this experiment. The last row of Fig. 8 shows the regions where the optical flow is missing for both kind of applications. Fig. 9 shows the reconstructions obtained with the two regularizers in both experiments. Table 5 shows the quantitative results, showing that the optical flow completion based on the proposed regularizer presents a lower reconstruction error than the one based on the TV.




Fig. 5. Dragon's head zoom. *First row: Optical flow with an L^1 data term and (a) the TV regularizer, (b) the proposed regularizer. Second row: EPE. (c) TV. (d) The proposed regularizer.*

6 Conclusion

We have proposed a new regularization term for optical flow models with the properties of invariance to infinitesimal rotations and the ability to preserve discontinuities/jumps in the flow field. The proposed regularizer has been tested in three different kind of problems related to the optical flow: motion estimation, motion inpainting, and motion interpolation from sparse samples. For the optical flow estimation we combine it with an L^1 data fidelity term, as in [10], and the proposed variational problem is solved using a dual formulation. The numerical experiments show that the proposed regularization term combined with an L^1 data term improves the TV- L^1 model for motion estimation. For the other two problems, optical flow reconstruction in two different cases of missing regions, we show that the missing information is better recovered with a functional based on the proposed regularizer, compared to the TV regularizer. As future work we plan to study the combination of the new regularizer with more advanced data terms robust to illumination changes, occlusions and fast movements.

Acknowledgement. We acknowledge partial support by MICINN project, reference MTM2012-30772, by GRC reference 2009 SGR 773 funded by the Generalitat de Catalunya, and by the ERC Advanced Grant INPAINTING (Grant agreement no.: 319899). The second author acknowledges partial support to the Ramón y Cajal program of the MINECO.




Fig. 6. Color representation of the level sets of the L^2 norm of the optical flow estimations shown in Fig 7. *From left to right:* (a) TV regularizer. (b) The proposed regularizer.




Fig. 7. Double windmill sequence. *From top to bottom:* (a) Original frames, (b) TV- L^1 optical flow and (c) optical flow estimated with an L^1 data term and the proposed regularizer.




Fig. 8. RubberWhale sequence from the Middlebury dataset with a rotation movement. *First row:* (a) First frame. (b) Ground Truth. *Second row:* (c) Inpainting mask (in white). (d) Missing optical flow (white pixels) and known values at sparse locations.




Fig. 9. Optical flow reconstruction in missing regions. *First row (from left to right):* Inpainting results. (a) TV regularizer. (b) Proposed regularizer. *Second row (from left to right):* Interpolation from sparse samples (c) TV regularizer. (d) Proposed regularizer. (e) Details, TV regularizer (top) and the proposed regularizer (bottom).

References

1. Horn, B. K. P., Schunck, B. G.: Determining optical flow. *Artificial Intelligence*, **17(1-3)** (1981)185–203
2. Black, Michael J., Anandan, P.: The Robust Estimation of Multiple Motions: Parametric and Piecewise-smooth Flow Fields. *Computer Vision and Image Understanding* **63 (1)**(1996) 75–104
3. Zimmer, H., Bruhn, A., Weickert, J.: Optic Flow in Harmony. *International Journal of Computer Vision* **93(3)**, (2011) 368–388.
4. Nagel, H H., Enkelmann, W.: An Investigation of Smoothness Constraints for the Estimation of Displacement Vector Fields from Image Sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8 (5)** (1986) 565–593
5. Wedel, A., Cremers D., Pock, T., Bischof, H.: Structure-and Motion-adaptative Regularization for High Accuracy Optic Flow *IEEE 12th International Conference on Computer Vision* (2009) 1663–1668
6. Werlberger, M., Trobin, W., Pock, T., Wedel, A., Cremers, D., Bischof, H.: Anisotropic Huber-L1 Optical Flow. *Proceedings of the British Machine Vision Conference* (2009)
7. Weickert, J., Schnörr, C.: A Theoretical Framework for Convex Regularizers in PDE-Based Computation of Image Motion. *International Journal of Computer Vision* **45(3)** (2001)
8. Trobin, W., Pock, T., Cremers, D., Bischof, H.: An Unbiased Second-Order Prior for High-Accuracy Motion Estimation. *DAGM* (2008) 396–405
9. Rudin, L. I., Osher, S., Fatemi, E.: Nonlinear Total Variation Based Noise Removal Algorithms. *Journal of Physics D: Applied Physics* **60** (1992) 259–268
10. Zach, C., Pock, T., Bischof, H.: A Duality Based Approach for Realtime TV-L1 Optical Flow. *DAGM* (2007) 214–223
11. Wedel, A., Pock, T. , Braun, J., Franke, U., Cremers, D.: Duality TV-L1 flow with fundamental matrix prior *23rd International Conference Image and Vision Computing New Zealand*, (2008) 1–6
12. Rosman, G., Shem-tov, S., Bitton, D., Nir, T., Adiv, G., Kimmel, R., Feuer, A., Bruckstein, A. M.: Over-Parameterized Optical Flow Using a Stereoscopic Constraint, Scale Space and Variational Methods in *Computer Vision* **6667** (2012) 761-772
13. Kondermann, C., Kondermann, D., Garbe, C.: Postprocessing of optical flows via surface measures and motion inpainting. *DAGM* (2008) 355–364
14. Matsushita, Y., Ofek, E., Ge, W., Tang, X., Shum, H.-Y.: Full-frame video stabilization with motion inpainting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28 (7)** (2006) 1150–1163
15. Shiratori, T, Matsushita, Y.: Video Completion by Motion Field Transfer.: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2006) 411–418
16. Chorin, A J., Marsden, J E.: *A mathematical introduction to fluid mechanics: Texts in applied mathematics*, Springer (1990)
17. Strelakovsky, E., Chambolle, A., Cremers, D.: Convex Relaxation of Vectorial Problems with Coupled Regularization. *SIAM Journal on Imaging Sciences* (2014)
18. Berkels, B., Rätz, A., Rumpf, M., Voigt, A.: Extracting Grain Boundaries and Macroscopic Deformations from Images on Atomic Scale. *Journal of Scientific Computing*. **35(1)** (2007), 1–23

19. Chambolle, A., Pock, T.: A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision* **40** (1)(2011) 120–145
20. Chambolle, A.: An Algorithm for Total Variation Minimization and Applications. *Journal of Mathematical Imaging and Vision* **20** (2004) 89–97
21. Sánchez Pérez, J., Meinhardt-Llopis, E., Facciolo, G.: TV-L1 Optical Flow Estimation. *Image Processing On Line (IPOL)* (2013) 137–150
22. Baker, S., Scharstein, D., Lewis, J. P., Roth, S., Black, M. J., Szeliski, R.: A Database and Evaluation Methodology for Optical Flow. *International Journal of Computer Vision* **92** (2010)
23. Butler, D. J., Wulff, J., Stanley, G. B. Black, M. J.: A naturalistic open source movie for optical flow evaluation. *European Conference on Computer Vision* (2012) 611–625
24. Bredies, K.: Recovering piecewise smooth multichannel images by minimization of convex functionals with total generalized variation penalty. *Lecture Notes in Computer Science* (2014).