

# FALDOI: a new minimization strategy for large displacement variational optical flow

Roberto P. Palomares · Enric Meinhardt-Llopis ·  
Coloma Ballester · Gloria Haro

Received: date / Accepted: date

**Abstract** We propose a large displacement optical flow method that introduces a new strategy to compute a good local minimum of any optical flow energy functional. The method requires a given set of discrete matches, which can be extremely sparse, and an energy functional which locally guides the interpolation from those matches. In particular, the matches are used to guide a structured coordinate-descent of the energy functional around these keypoints. It results in a two-step minimization method at the finest scale which is very robust to the inevitable outliers of the sparse matcher and able to capture large displacements of small objects. Its benefits over other variational methods that also rely on a set of sparse matches are its robustness against very few matches, high levels of noise and outliers. We validate our proposal using several optical flow variational models. The results consistently outperform the coarse-to-fine approaches and achieve good qualitative and quantitative performance on the standard optical flow benchmarks.

**Keywords** Optical flow · Variational methods · Coordinate descent · Sparse matches

## 1 Introduction

Optical flow is the apparent motion field between two consecutive frames of a video. More generally, it can be defined as a dense correspondence field between an arbitrary pair of images. There are two large families of methods for computing image correspondences: local and global methods. Local methods establish a point correspondence by minimizing a distance measure between the matching neighborhoods [11, 60]. They provide a sparse correspondence field since not all the image points are discriminative enough to guarantee a single correspondence. On the other hand, global or variational methods [24, 2, 9, 62, 64, 50] provide a dense solution by minimizing a global energy. Recent work on optical flow estimation [4, 14, 18, 27, 35, 59, 43, 30] is mostly focused on solving the major challenges that appear in realistic scenarios and outdoor scenes, such as large displacements, motion discontinuities, illumination changes, and occlusions.

In this article we introduce a new strategy to compute good local minima of any optical flow energy functional which allows to capture large displacements. The method relies on a discrete set of matches between the two input images, which can be extremely sparse and contaminated

---

The first, second and third authors acknowledge partial support by MICINN project, reference MTM2012-30772, by TIN2015-70410-C2-1-R (MINECO/FEDER, UE) and by GRC reference 2014 SGR 1301, Generalitat de Catalunya, and the fourth author by the European Research Council (advanced grant Twelve Labours), Office of Naval research (ONR grant N00014-14-1-0023).

---

R.P. Palomares, C. Ballester and G. Haro  
Universitat Pompeu Fabra  
E-mail: roberto.palomares@upf.edu, coloma.ballester@upf.edu, gloria.haro@upf.edu

E. Meinhardt-Llopis  
École Normale Supérieure de Cachan  
E-mail: enric.meinhardt@cmla.ens-cachan.fr

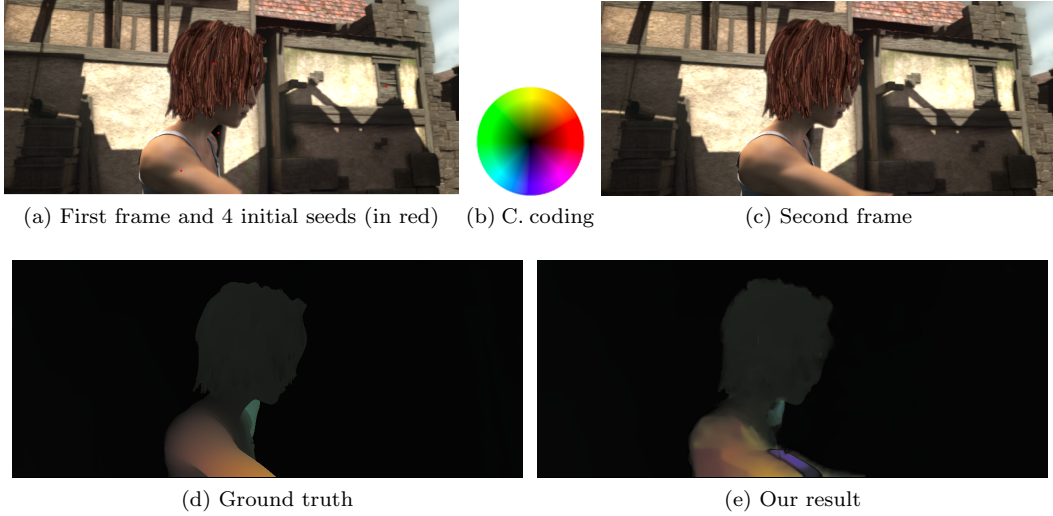


Fig. 1: Example illustrating that it is enough to have a single matching seed per each region with an expected smooth motion; 4 seeds in this case, which have been computed using SIFT. The first row shows the original pair of frames  $I_t$  and  $I_{t+1}$ , where in  $I_t$  we have superimposed in red the position of the 4 seeds (head, arm, dress and background). Last row shows our result and the ground truth. The energy functional used is the classical  $TV_{\ell_2-L1}$ .

by outliers, and is used as a guide to find a local minimum of the chosen energy functional. The novelties include how these matches are used in the optimization problem. The proposed method is a two-step minimization process of the optical flow energy. The first step computes a good and dense local minimum by propagating the initial matches with a region growing strategy. The propagation is driven by the minimization of the energy on patches and the order of update is established by the value of the local version of the energy, *i.e.* the energy restricted to a square patch of the image domain. It can be thought as a grouped coordinate descent of the target functional, where the variables are defined by the optical flow values in a patch, and the functional is minimized on the patch while fixing the rest of flow values. Coordinate descent methods provide efficient and fast optimization schemes even for huge-scale problems [39]. The basic idea is to minimize a multi-variable objective function by optimizing over a single coordinate while holding others fixed. The sweep pattern is the name used to define how the alternating minimization process moves along all the coordinates. In order to obtain a good minimum and accelerate convergence, we follow an adaptive choice of the sweep pattern driven by a seed growing algorithm based on the value of the energy after local minimization in a patch centered at that coordinate. This information is stored in a priority queue from which the coordinate is selected and the adaptive sweep pattern is iteratively created. The result of the first step is a dense optical flow which is further refined in the second step: A global minimization of the energy taking as initial condition the flow obtained in the first stage. Both minimization steps work with the full resolution of the image directly.

Our method allows to choose the energy functional. For instance, by choosing it with desirable properties such as robustness to illumination changes, occlusions and motion discontinuities, those properties are inherited by our method. We can compare our method to the common coarse-to-fine (also called multi-scale) approach or to the strategy of including the information of sparse matches by incorporating an extra term in the energy that penalizes deviations from the flows given by the matches [10, 49, 55]; both of them also find local minima of arbitrary optical flow energy functionals. We find that we consistently outperform the multi-scale strategy for a variety of input sparse matches and energy functionals. We present qualitative and quantitative results on several datasets such as Middlebury [5], MPI-Sintel [12], KITTI 2012 [21], and KITTI 2015 [34] datasets. The performance is better than LDOF [10] and comparable to DeepFlow [55] while being more robust, less dependent on the density of seeds and based only on a single energy with two

terms (no need of extra parameters). In our case, the ability to find large displacements requires only that *at least one* match is correctly given for each object in motion, as shown in the experiment of Figure 1. Moreover, we find that our minimization strategy is very robust to the presence of many incorrect keypoints in the input, as illustrated in Figure 3, where there are only 2 correct matches and 508 outliers. In contrast, the alternative strategy of including sparse matches in the energy functional with an additional term is not robust to a small number of matches (see Figure 10) nor to a large percentage of wrong matches or high levels of noise (see Table 6).

The remainder of the paper is organized as follows. In Section 2 we revise previous work on optical flow estimation. Section 3 presents our large displacement optical flow method (which we denote by FALDOI, as an acronym made of the first letter of the words "large displacement optical flow method by an astute initialization"). Section 4 presents results from several energy functionals and an analysis of the properties and performance of our strategy depending on the density of the initial set of discrete matches. Finally, the conclusions are summarized in Section 5.

## 2 Related work

The seminal work of Horn and Schunk [24] has inspired many subsequent works also based on an energy-based formulation of the optical flow problem and focused on different kind of limitations and their improvements. Initial progress was devoted to the use of more robust data terms and Total-Variation-based regularizers in order to obtain sharper solutions [9, 8, 51, 62]. The correct preservation of motion discontinuities is a key issue that has motivated many proposals in the smoothness term: Starting from the use of decreasing functions of the image gradients [14, 45, 58], diffusion tensors [38, 57, 63], coupled regularization of the flow channels [50, 41] or second-order regularization [42], to the more recent non-local regularization terms [28, 52, 56]. The classical hypothesis for defining the data term has been the brightness constancy assumption but this is very limiting in realistic scenarios where illumination changes as well as occlusions may appear. Robustness against additive illumination changes can be obtained by using the gradient constancy assumption [40]; while advanced data terms based on patch measures present more general invariances, in particular, the Normalized Cross Correlation [47, 54, 56] is invariant to linear brightness changes and the Census transform is invariant under monotonically increasing rescalings [23, 37, 46]. Patch-based data terms are used in the state-of-the-art methods, including stereo and optical flow problems. As they are more informative and allow to better characterise the local image structure they result in more accurate flow estimations. On the other hand, occlusions represent an important problem for optical flow estimation methods and some works include a characterization of the occlusion areas in the energy, [1, 3, 6, 18], while others estimate the occlusion areas based on a triangulation of the image [27].

Energy-based methods are called global methods since they find correspondences by minimizing an energy defined on the whole image. Traditionally, global methods include a linearization of the warped image in the data term to make the optimization problem more tractable. The linear approximation is only valid for small displacements, that is why these methods are usually embedded in a coarse-to-fine warping scheme in order to better capture large displacements. However, they still fail to correctly handle large motions of small objects (not present at coarser scales). On the other hand, local methods are much better adapted to capture large movements. The work of Álvarez *et al.* [2] was probably the first to note that the standard coarse-to-fine approach may not be enough and proposed to modify it by using a linear scale-space focusing strategy from coarse to fine scales in order to avoid convergence to incorrect local minima. More recently, Steinbrücker *et al.* [48] proposed an algorithm that does not require the coarse-to-fine warping strategy. It was one of the first attempts to capture large motions with variational methods, where the data term and the regularizer are decoupled by a quadratic relaxation technique and the optimization problem is directly solved at the finest scale by alternating two global minimizations while decreasing the decoupling parameter. One of the minimization problems is convex; the other one is non-convex but it is a point-wise optimization problem, so its solution is found using an exhaustive search for every pixel, which can be performed in parallel. Another early work is the proposal by Brox and Malik [10], that incorporates sparse matches into the variational model by adding an extra term

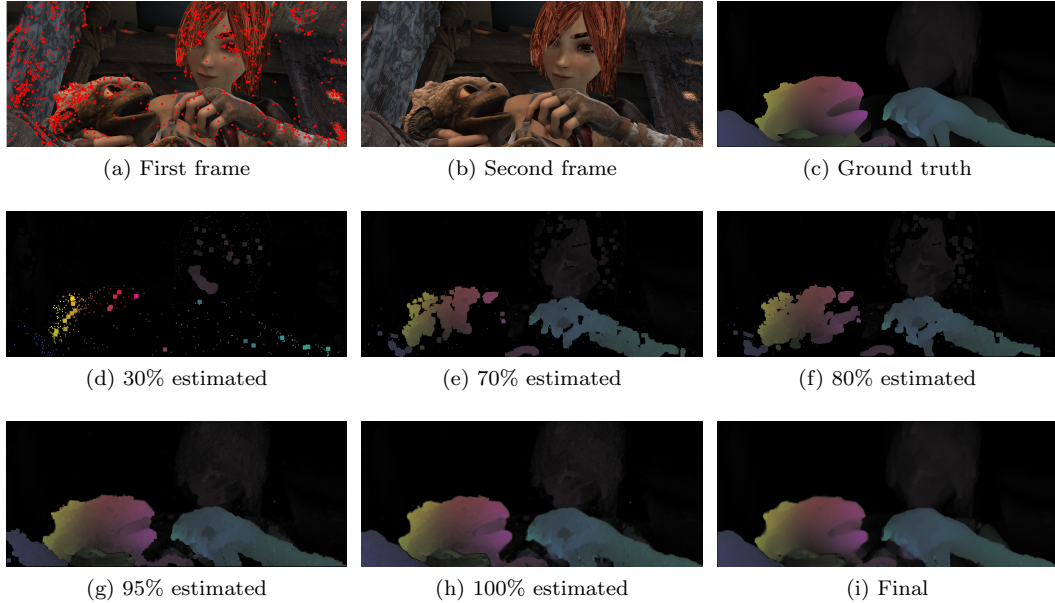


Fig. 2: Example of a sparse-to-dense evolution of the estimated optical flow during the local step, starting from a set of initial seeds provided by SIFT (shown in red in (a)). The final estimate (after the global minimization step) is also shown. The optical flow values are represented using the color coding scheme shown in Fig. 1. The energy functional used is the classical  $TV_{\ell_2}-L1$ .

that penalizes deviations of the estimated optical flow from the flow given by the matches obtained by HOG descriptors [15] at some sparse locations. The extended variational method is solved with a coarse-to-fine strategy as usual. DeepFlow [55] and SparseFlow [53] follow the same strategy as proposed by Brox and Malik [10], the difference among them lying in the way they find the matches that define the matching term in the variational approach. The descriptor matching proposed by Weinzaepfel *et al.* [55] is inspired by non-rigid 2D warping and deep convolutional networks and permits a more dense and deformable matching compared to the popular HOG/SIFT-like descriptors [15, 33]. The matching algorithm proposed by Timofte *et al.* [53] is also robust to non-rigid deformations and is based on the compressed sensing theory. The proposal by Xu *et al.* [58] also uses a coarse-to-fine approach but the initial flow in each level of the pyramid is modified. At each level, different candidate flows are considered: the flow propagated from the previous level, sparse feature matching (computed by SIFT), and dense nearest neighbor patch matching. Then, for each pixel, the optimal flow is selected among the different candidates by solving a discrete optimization problem.

Other works that do not use a coarse-to-fine strategy are based on correspondences obtained by a nearest neighbor algorithm; one is based on a purely local method [7] while another [14] refines the initial dense (and noisy) correspondence field using a motion segmentation and the minimization of a global energy. Some recent methods are based on a sparse-to-dense approach; they start from a sparse set of correspondences that capture large displacements and these are densified by edge-aware interpolation techniques [30, 43]. Later, the densified flow is refined by minimizing a global energy at the finest scale. Other works [35, 59] also include a final refinement of the flow by minimizing a global continuous energy in the original image scale directly; either using a discrete inference problem in a conditional random field as a first step [35], or a piecewise adaptive parametric model embedded in an energy function that combines both continuous and discrete variables [59]. The recent proposal of Fortun *et al.* [18] solves a discrete optimization of a global energy in the original scale where the different flow candidates are obtained in a previous step by patch-based parametric motion estimation; this method also makes a special treatment of occlusions. On the other hand, an almost purely data-based dense optical flow is indeed possible,

by working with approximate nearest neighbor fields with a hierarchical search strategy and an advanced outlier filtering [4].

Here, we also propose to get rid of the coarse-to-fine strategy and directly minimize the energy at the finest scale with the help of some sparse correspondences that capture the large displacement motions. The main difference with respect to the previous works [30, 43] is the way of combining the sparse matches with the variational approach: the initial matches are densified by a region growing process driven by the minimization of a local version of the target energy functional. Thus, the same variational tool is used both for the densification step and for the final global optimization problem. In contrast, the sparse-to-dense interpolation of the sparse matches in [30, 43] is based on a regularity measure that takes into account occlusion and boundary constraints. In this way, since in our case we are always taking into account a data term and a regularizer, together with a smart local minimization of the energy guided by the best matches, our sparse-to-dense intermediate step is more robust to outliers in the initial set of sparse correspondences and is able to correctly recover a dense smooth motion in a certain region with just a single correspondence in it and without the need of using edge information (see Figures 2 and 3). The outliers in this experiment were randomly introduced to test our claim against a huge percentage of them. Nevertheless, the robustness of our iterated algorithm to usual outliers is shown in Figure 4, the different Tables and remaining experiments and it relies on pruning strategies based on saliency and forward-backward optical flow consistency.

### 3 Proposed minimization strategy

In this section we present a minimization strategy that can be applied to any optical flow energy functional and which is founded on estimating a good local minimum with the help of a discrete set of matches. It is able to benefit both from the sparse techniques, which handle arbitrarily large displacements, and from the continuous optimization of a variational formulation, which yields dense flow fields with subpixel accuracy. The basic idea of the method is to assume that at least some matches are correct, and propagate the correct information from those seeds driven by the minimization of the energy around them. This is illustrated in Figure 1, where each region of smooth movement has at least one correct seed.

The sparse set of initial correspondences (we will refer to them as *seeds*) is used as a reference or guide to recover a dense flow field. This is done by iteratively growing the seeds by a local (patch-based) minimization of the functional in a proper order (detailed in Sect. 3.1.1 and Sect. 3.1.2). This dense optical flow is then refined by a global minimization of the energy. The algorithm always works at the finest scale of the image. A pseudo-code of the whole strategy is presented in Algorithm 1 while Algorithm 3 presents a slightly modified version which is explained in Sect. 3.1.2.

In order to detail the proposed approach, we introduce some notation and assumptions. Let us denote two consecutive image frames of a video sequence as  $I_t, I_{t+1} : \Omega \rightarrow \mathbb{R}$ . As usual, we assume that the image domain  $\Omega$  is a rectangle in  $\mathbb{R}^2$ . In order to compute the optical flow  $\mathbf{u} : \Omega \rightarrow \mathbb{R}^2$  between  $I_t$  and  $I_{t+1}$ , we use a discrete set of matches  $\mathcal{M} = \{(\mathbf{x}_i, \mathbf{y}_i)\}$ ,  $i = 1, \dots, N$ , and an energy  $E(\mathbf{u})$ , defined from  $I_t$  and  $I_{t+1}$ . Minimizing  $E$  on an appropriate set, a minimum  $\mathbf{u}$  of  $E$  represents an optical flow between  $I_t$  and  $I_{t+1}$ . We assume that the discrete matches in  $\mathcal{M}$  have been computed with a sparse matching algorithm in such a way that  $\mathbf{x}_i$  is thought as belonging to the image domain of  $I_t$  and  $\mathbf{y}_i$  as belonging to the image domain of  $I_{t+1}$ . From these discrete matches, we compute the initial set of seeds, denoted here by  $S$ , by defining  $\mathbf{u}(\mathbf{x}_i) = \mathbf{y}_i - \mathbf{x}_i$ ,  $i = 1, \dots, N$ . Each seed  $\mathbf{p}$  in the finite set of initial seeds  $S$  stores the corresponding pixel  $\mathbf{x}^{\mathbf{p}}$ , its related optical flow  $\mathbf{u}^{\mathbf{p}}$  and the local energy of the functional around it (see Algorithm 1).

#### 3.1 Computing a good local minimum

This step can be interpreted as an adaptive grouped coordinated descent approach driven by the lowest local values of the energy on patches centered at pixels. It is inspired from the match propagation principle [31], where a set of initial sparse matches, the seeds, are propagated to

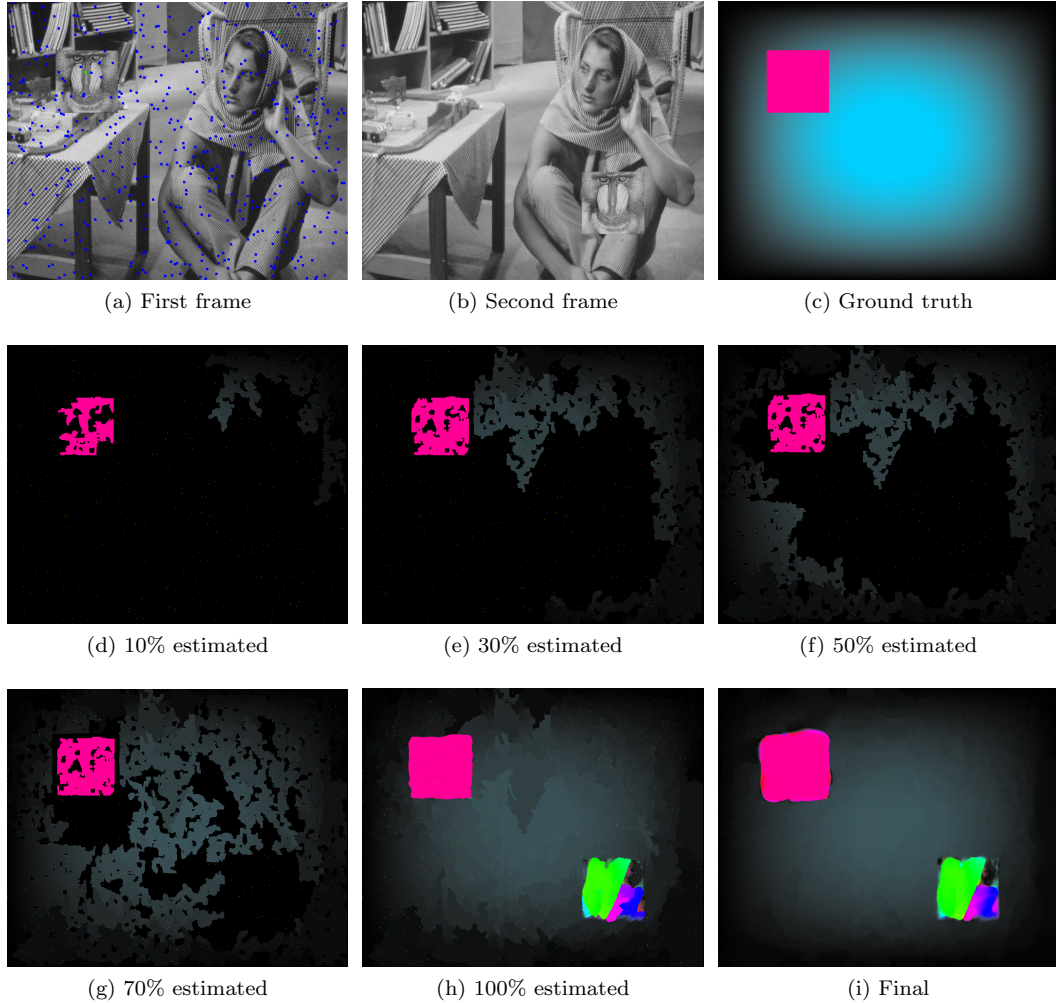


Fig. 3: Sparse-to-dense evolution of the fixed coordinates for a toy sequence made of the composition of a large displacement (moving the *Baboon* small image) and a smooth movement of the background. The initial set of sparse matches contains only 2 correct matches (obtained by SIFT and shown in green) and 508 outliers (shown in blue). Even in the presence of a huge number of outliers, the method recovers the correct motion in all the image with the exception of the occluded region since the chosen functional ( $TV_{\ell_2-L1}$ ) does not take into account occlusions. For visualization purposes and better observe the sweep pattern in the optical flow estimation, in this figure we have increased the color contrast. The optical flow values are represented using the color coding scheme shown in Fig. 1.

neighboring pixels using a similar technique to the region growing strategy. This principle was used in the work of [26] in order to obtain a quasi-dense disparity map, assuming that the seeds and their neighboring pixels may present similar disparity values. Moreover, it shares ideas with the coordinate descent methods, which are optimization techniques in multi-variable functions. They minimize the objective function solving a sequence of one variable minimization problems. Each subproblem improves the estimate of the solution by minimizing along a selected coordinate while all other coordinates are fixed. Generally, each coordinate is visited several times to reach a minimum. The sweep pattern is the name used to define how the alternating minimization process moves along all the coordinates. If there is a fixed order to visit the coordinates, this is called path-wise coordinate descent [19] or cyclic coordinate. In our case, the election of the sweep pattern

has a fundamental role during the minimization process; we will follow an adaptive choice of the sweep pattern driven by a seed growing algorithm based on the value of the local minimization of the energy in a patch centered at that coordinate (or pixel in our case). The sweep process is managed through a priority queue where the potential candidates (optical flow) for each coordinate are inserted. Each candidate presents a related energy that is used to determine its position in the queue: Candidates with less associated local energy will be at the top of the priority queue. The initial seeds are inserted with zero energy.

In the following we will first present the baseline algorithm for the local minimization step, where every pixel is visited just once. Then we explain how it is iterated in order to revisit the pixels several times and gain more robustness against occlusions and resistant outliers of the sparse matcher.

### 3.1.1 Baseline algorithm: *faldoi*

Initially, the seeds are inserted to the priority queue (containing, as said before, the value of the optical flow and the energy on a local patch centered at the pixels) with zero energy. Along the minimization process, new candidates will be added to the queue. This collection of potential candidates for each coordinate will be sorted based on their local energies. Whenever an element is removed from the queue to fix a coordinate, we are selecting the candidate with the lowest associated local energy. The aforementioned process is repeated until the priority queue is empty and a dense optical flow is obtained from the initial seeds. There may be several candidates for the same pixel in the queue; when a candidate is extracted from the queue to fix it, if its corresponding pixel has already been fixed (by a candidate with lowest energy) nothing is done. Once, all the pixels have been fixed, the rejected candidates are extracted from the queue until it is emptied.

Each time that a coordinate is extracted from the queue to fix it, a local minimization of the energy  $E(u)$  is solved on a square patch centered around the pixel previously fixed. Then, the estimated optical flow values of its neighbors are inserted as potential candidates into the priority queue with the energy (after local minimization) of the patch centered at the fixed coordinate.

Whenever the energy is minimized in a local patch  $\mathcal{P}$  centered at the coordinate that has been fixed we need to provide an initial flow in the unknown values in  $\mathcal{P} \cap W$  (where  $W$  is the set of locations where the optical flow has not been fixed yet). The initialization is just an interpolation of the already fixed optical flow values in the patch through the Laplace equation resulting from the minimization of the following Dirichlet energy

$$\min_{u_i} \int_{\mathcal{P}} |\nabla u_i|^2 d\mathbf{x} \quad \text{s.t.} \quad u_i = u_i^0 \text{ in } \mathcal{P} \cap W^C$$

where  $i = 1, 2$ ,  $\mathbf{u} = (u_1, u_2)$ , and  $\mathbf{u}^0 = (u_1^0, u_2^0)$  contains the optical flow of the already fixed coordinates in  $\mathcal{P} \cap W^C$  ( $W^C$  denotes the complementary set of  $W$ ). The Euler-Lagrange equation derived from the previous energy is the Laplace equation, which is solved by gradient descent with Neumann boundary conditions. Other intra-patch interpolations are allowed (e.g., based on the bilateral filter or even a constant interpolation based on the flow values on  $\mathcal{P} \cap W^C$ ).

The details of the energy minimization process are given in Appendix A and Algorithm 4. In practice, we consider patches of  $11 \times 11$  pixels (although bigger patches can be used, it increases the computational cost which is proportional to the size of the patch) and we perform 10 iterations of the minimization process, in every local patch (line 7 of Algorithm 2), of a version of  $E(u)$  where the warped image has been linearized with a single warping.

**Algorithm 1:** faldoi (large displacement optical flow with astute initialization)

---

**Input** : Images  $A, B$   
**Input** : Functional  $E$   
**Input** : Image matcher  $M$   
**Input** : Patch size  $w$   
**Output:** Flow  $\mathbf{u}$

```

1  $\mathbf{u} \leftarrow \text{NULL}$  // initialize flow field with empty data
2  $Q \leftarrow \emptyset$  // initialize an empty priority queue
3  $(\mathbf{x}_i, \mathbf{y}_i)_{i=1, \dots, N} \leftarrow M(A, B)$  // compute discrete matches
4 for  $i = 1, \dots, N$  do // add the matches as seeds with zero energy
5    $Q.\text{push}(0, \mathbf{x}_i, \mathbf{y}_i - \mathbf{x}_i)$ 
6  $\mathbf{u} \leftarrow \text{basic-faldoi-growing}(A, B, \mathbf{u}, E, Q, w)$  // basic faldoi growing
7  $\mathbf{u} \leftarrow \text{flow-refinement}(A, B, \Omega, E, \mathbf{u})$  // Minimize  $E_{A,B}(\mathbf{u})$  over  $\Omega$  (Algorithm 4)

```

---

**Algorithm 2:** basic-faldoi-growing (densify an incomplete flow)

---

**Input** : Images  $A, B$   
**Input** : Flow  $\mathbf{u}_0$   
**Input** : Functional  $E$   
**Input** : Priority queue  $Q$   
**Input** : Patch size  $w$   
**Output:** Flow  $\mathbf{u}$

```

1 while  $Q.\text{num\_elements}() > 0$  do
2    $e, \mathbf{x}, \mathbf{v} \leftarrow Q.\text{pop}()$  // get the candidate with lowest energy
3   if  $\mathbf{u}(\mathbf{x}) = \text{NULL}$  then
4      $\mathbf{u}(\mathbf{x}) \leftarrow \mathbf{v}$  // fix the field for this candidate
5    $\mathcal{P}_x \leftarrow \text{extract-patch}(\Omega, \mathbf{x}, w)$  // extract patch of size  $w \times w$  around  $x$ 
6    $\mathbf{u} \leftarrow \text{interpolate}(\mathbf{u}, \mathcal{P}_x)$  // fill-in missing values
7    $\mathbf{u} \leftarrow \text{flow-refinement}(A, B, \mathcal{P}_x, \mathcal{E}, \mathbf{u})$  // Minimize  $E_{A,B}(\mathbf{u})$  over  $\mathcal{P}_x$  (Algorithm 4)
8    $e \leftarrow E_{A,B}(\mathbf{u}, \mathcal{P}_x)$  // compute the energy of the solution
9   for  $\mathbf{y} \in \mathcal{N}(\mathbf{x})$  do // traverse the neighbors of  $\mathbf{x}$ 
10    if  $\mathbf{u}(\mathbf{y}) = \text{NULL}$  then
11       $Q.\text{push}(e, \mathbf{y}, \mathbf{u}(\mathbf{y}))$  // push this candidate value

```

---

### 3.1.2 Iterated faldoi

Given the fact that, generally, in coordinate descent methods each coordinate has to be visited several times in order to reach a minimum, we propose to revisit the coordinates and thus perform several iterations of the baseline algorithm presented in Sect. 3.1.1 with a forward-backward pruning of the flow values between two consecutive iterations in order to gain more robustness against resistant outliers.

More precisely, between two consecutive iterations of the coordinate descent completing a sweep pattern, we perform a pruning based on a forward-backward consistency check. The goal is to remove wrong matches, specially in occlusion areas and also due to self-similarity. The latter may appear due to outliers in the initial seeds that have expanded by the region growing scheme. Our algorithm computes both the forward,  $\mathbf{u}^F$ , and the backward,  $\mathbf{u}^B$ , optical flows between any two frames  $I_t$  and  $I_{t+1}$ . Then, a forward-backward consistency check of these flows is performed and the forward flow values at points  $\mathbf{x} \in \Omega$  not verifying  $\|\mathbf{u}^F(\mathbf{x}) + \mathbf{u}^B(\mathbf{x} + \mathbf{u}^F(\mathbf{x}))\| < \epsilon$  are removed, where  $\epsilon > 0$  is a small constant (it is set to 2 in our experiments). In the same way, the backward flow at points  $\mathbf{x} \in \Omega$  not verifying  $\|\mathbf{u}^B(\mathbf{x}) + \mathbf{u}^F(\mathbf{x} + \mathbf{u}^B(\mathbf{x}))\| < \epsilon$  are removed.

In the iterated faldoi, the sweeps following the first one are slightly different from the baseline sweep of Sect. 3.1.1 because they start not from a sparse set of optical flow values but from a dense

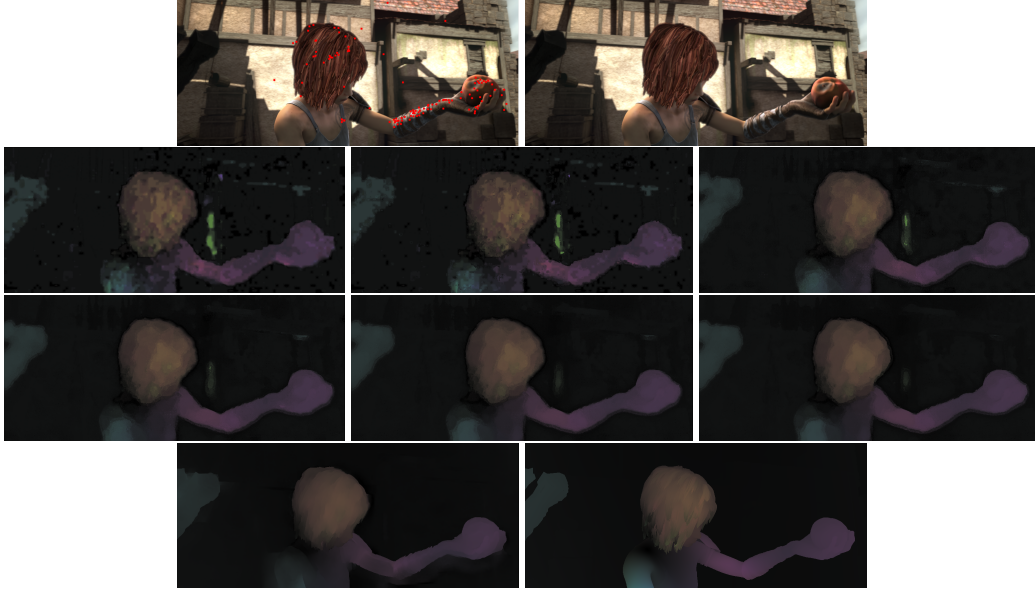


Fig. 4: Iterated faldoi strategy applied to the  $NLTV-CSAD$  energy corresponding to  $MAX\_IT=6$  in Algorithm 3. The evolution of the optical flow estimate is shown from left to right and top to bottom. The first line shows the two frames with the initial seeds superimposed in red in the first frame. The second and third lines display the six iterations of the sweep pattern in the local minimization step. The last line shows the final estimate (global minimization) and the ground truth. Let us notice how the motion boundaries are progressively refined, the piecewise constant appearance disappears, and the outlier is considerably reduced, thus adding robustness to outliers. In this experiment the seeds have been computed using Deep Matching, which gives more outliers than SIFT (used for the experiment with the same pair of images in Figure 2).

motion field with holes that arise after the forward-backward pruning. The two main differences with respect to the main sweep are:

1. *Initialization of the queue.* The seeds that survive to the forward-backward consistency check are inserted into a new priority queue with zero local energy. The rest of optical flow values that also survive to the forward-backward pruning are added with its associated local energy.
2. *Initialization of the optical flow of a patch before local minimization.* Every time a local patch is minimized an initial flow in the patch is needed. We make a distinction between the pixels that passed the forward-backward consistency test and those that did not. In the pixels that passed the test the initial flow is the most updated flow value by the last local minimization in that pixel. For that, we use an auxiliary flow which is updated after every local minimization. On the other hand, in the pixels that did not pass the test we fill-in the initial flow by iterating a bilateral filtering in the local patch. This intra-patch initialization could be also used in the first baseline sweep of Sect. 3.1.1. Nevertheless, in that very first interpolation, for efficiency reasons we opted for the rougher Laplace interpolation explained above.

A pseudo-code of the iterated faldoi strategy is presented in Algorithm 3. Figure 4 shows the importance of the iterations of the sweep pattern in the local minimization step. The optical flow is improved from one iteration to the following, in particular, the motion boundaries are progressively better recovered and the effect of the piecewise constant flow is lost. Moreover, this iterated faldoi strategy adds robustness to outliers present in the initial seeds due to the forward-backward consistency check between two consecutive iterations. The initial outlier regions are reduced, after six iterations to a single region, which is completely lost after the global iteration step. In this experiment the seeds have been computed using Deep Matching, which gives more outliers than SIFT (which was used for the experiment with the same pair of images in Figure 2).

Let us remark that in this example, we did not include the initial pruning of seeds based on the low saliency of patches that is used, e.g., in [43].

### 3.2 Global Minimization

The last step is a global minimization of  $E(u)$ , taking as initial condition the dense solution provided by the previous step, the local minimization guided by the initial set of matches. Let us remark that we only minimize the global energy at the finest scale, that is, we avoid the multi-scale approach. During both steps, local and global, we minimize the same energy functional following the same numerical scheme, which is detailed in Algorithm 4 and Appendix A. For the global minimization we perform 4 warpings of the energy.

---

**Algorithm 3:** iterated-faldoi (iterated version of faldoi + pruning of wrong matches)

---

**Input** : Images  $A, B$   
**Input** : Functional  $E$   
**Input** : Image matcher  $M$   
**Input** : Patch size  $w$   
**Output**: Flow  $u$

```

1  $u^F \leftarrow NULL$  // initialize forward flow field with empty data
2  $u^B \leftarrow NULL$  // initialize backward flow field with empty data
3  $Q^F \leftarrow \emptyset, Q^B \leftarrow \emptyset$  // initialize empty priority queues
4  $(x_i^F, y_i^F)_{i=1, \dots, N} \leftarrow M(A, B)$  // compute discrete forward matches
5  $(x_i^B, y_i^B)_{i=1, \dots, M} \leftarrow M(B, A)$  // compute discrete backward matches
6  $(x_i^F, y_i^F)_{i=1, \dots, N} \leftarrow \text{saliency-pruning}(A, (x_i^F, y_i^F)_{i=1, \dots, N})$  // remove matches in
7  $(x_i^B, y_i^B)_{i=1, \dots, M'} \leftarrow \text{saliency-pruning}(B, (x_i^B, y_i^B)_{i=1, \dots, M})$  // flat areas
8 for  $i = 1, \dots, N'$  do // add the matches as seeds with zero energy
9    $Q^F.push(0, x_i^F, y_i^F - x_i^F)$ 
10 for  $i = 1, \dots, M'$  do // add the matches as seeds with zero energy
11    $Q^B.push(0, x_i^B, y_i^B - x_i^B)$ 
12  $n \leftarrow \text{MAX-IT}$  // pre-specified number of iterations
13 while  $n > 0$  do
14    $u^F \leftarrow \text{basic-faldoi-growing}(A, B, u^F, E, Q^F, w)$  // basic faldoi growing
15    $u^B \leftarrow \text{basic-faldoi-growing}(B, A, u^B, E, Q^B, w)$  // basic faldoi growing
16    $(u^F, u^B) \leftarrow \text{fb-pruning}(u^F, u^B)$  // forward-backward consistency check
17   for each  $u_i^F \in u^F$  do // add the forward candidates with their energies
18      $e \leftarrow E_{A,B}(u_i^F, \mathcal{P}_{x_i^F})$ 
19      $Q^F.push(e, x_i^F, u_i^F)$ 
20   for each  $u_i^B \in u^B$  do // add the backward candidates with their energies
21      $e \leftarrow E_{A,B}(u_i^B, \mathcal{P}_{x_i^B})$ 
22      $Q^B.push(e, x_i^B, u_i^B)$ 
23    $n \leftarrow n - 1$ 
24  $u \leftarrow \text{flow-refinement}(A, B, \Omega, E, u^F)$  // Minimize  $E_{A,B}(u)$  over  $\Omega$  (Algorithm 4)

```

---

## 4 Results

Our proposal assumes that two ingredients are given: a discrete set of correspondences between two frames of a video (the seeds) and an optical flow energy functional, namely,  $E(u)$ . Section 4.1 presents the different energies we have used in this paper and Section 4.2 briefly discusses several possibilities for the initial sparse correspondences. Later on, in Sect. 4.3 we provide a quantitative

and qualitative comparison among the several possibilities for the energy terms, as well as a comparison between our method and several state-of-the-art methods, including methods based on the combination of sparse matches and variational techniques [10, 30, 43, 55].

#### 4.1 Different possibilities for the energy

The proposed framework is independent from the energy functional  $E(\mathbf{u})$  and we validate it by using several energies. Following most of the optical flow variational approaches in the literature, the different possibilities will share the common feature of being made of two terms: a data fidelity term  $E_D(\mathbf{u})$ , and a regularization term  $E_R(\mathbf{u})$ ,

$$E(\mathbf{u}) = E_D(\mathbf{u}) + \beta E_R(\mathbf{u}), \quad (1)$$

where  $\mathbf{u} : \Omega \rightarrow \mathbb{R}^2$ ,  $\Omega$  is the image domain,  $I_t, I_{t+1} : \Omega \rightarrow \mathbb{R}^d$  are two consecutive frames ( $d = 1$  for gray level images and  $d = 3$  for color images), and  $\mathbf{u} = (u_1, u_2)$  represents the motion field between them.

Let us also remark that our method allows the inclusion of other terms as, *e.g.*, a third term dealing with the occlusions, as in [6, 3], which will not be considered here for the sake of simplicity in the presentation.

With the aim of having a robust data cost, specially under illumination changes, we use a convex and continuous approximation of the data cost based on the Census transform [23, 46, 61], approximated by a sum of centralized absolute differences, denoted as *CSAD* [54], and defined as:

$$E_D^1(\mathbf{u}) = \int_{\Omega} \int_{\Omega} |I_t(\mathbf{x}) - I_t(\mathbf{y}) - I_{t+1}(\mathbf{x} + \mathbf{u}) + I_{t+1}(\mathbf{y} + \mathbf{u})| \chi(\mathbf{x} - \mathbf{y}) d\mathbf{y} d\mathbf{x}, \quad (2)$$

where  $\chi$  denotes the characteristic function of a square of size  $P \times P$  centered at the origin ( $P = 7$  in our experiments, as in [54]).

On the other hand, we also consider the classical point-wise  $L^1$  data term that imposes the well-known brightness constancy assumption, namely,

$$E_D^2(\mathbf{u}) = \int_{\Omega} |I_{t+1}(\mathbf{x} + \mathbf{u}) - I_t(\mathbf{x})| d\mathbf{x}. \quad (3)$$

It is important to correctly preserve the motion boundaries in order to obtain an accurate optical flow. Motion boundaries are usually aligned with image boundaries, this aspect has motivated the use of edge detectors (*e.g.* [17, 29]) in previous optical flow methods [30, 35, 43]. Instead, we opt for using the Non-Local Total Variation (*NLTV*) regularizer. *NLTV* was used for optical flow estimation in [56], where the authors show its ability to better recover motion boundaries, in particular in low-textured areas, occluded regions, and in small objects (when used in a coarse-to-fine scheme). In our case, a non-local regularizer which better captures motion discontinuities is also very useful in the local minimization step where the initial correspondences are densified. We apply the *NLTV* regularizer to each flow channel independently:

$$E_R^1(\mathbf{u}) = \int_{\Omega} \int_{\Omega} \omega(\mathbf{x}, \mathbf{y}) (|u_1(\mathbf{x}) - u_1(\mathbf{y})| + |u_2(\mathbf{x}) - u_2(\mathbf{y})|) d\mathbf{y} d\mathbf{x}, \quad (4)$$

where  $\omega(\mathbf{x}, \mathbf{y}) = \frac{1}{\mathcal{W}(\mathbf{x})} e^{-\frac{\Delta_c(\mathbf{x}, \mathbf{y})}{\sigma_c}} e^{-\frac{\Delta_s(\mathbf{x}, \mathbf{y})}{\sigma_s}}$ ,  $\Delta_c(\mathbf{x}, \mathbf{y})$  denotes the Euclidean distance between the color values at  $\mathbf{x}$  and  $\mathbf{y}$  in the Lab space,  $\Delta_s(\mathbf{x}, \mathbf{y})$  is the Euclidean distance between points  $\mathbf{x}$  and  $\mathbf{y}$ ,  $\mathcal{W}(\mathbf{x}) = \int_{\Omega} \omega(\mathbf{x}, \mathbf{y}) d\mathbf{y}$  is a normalizing constant, and  $\sigma_c, \sigma_s > 0$  are constant parameters (set to  $\sigma_c = 2$  and  $\sigma_s = 2$ ).

As before, with the purpose of testing the proposed minimization procedure against different energies, we consider other regularization terms such as the classical coupled Total Variation (called  $TV_{\ell_2}(u)$  in [50]), namely,

$$E_R^2(\mathbf{u}) = \int_{\Omega} \|\nabla \mathbf{u}\|_2 d\mathbf{x} = \int_{\Omega} \sqrt{|\nabla u_1(\mathbf{x})|^2 + |\nabla u_2(\mathbf{x})|^2} d\mathbf{x}, \quad (5)$$

Appendix A details the optimization algorithms for all the energies. Let us now just say that the data and the regularization terms are decoupled and standard methods such as primal-dual, thresholding or median schemes are used.

#### 4.2 Different possibilities for the initial set of seeds

In this work, the initial seeds are computed from one of the sparse matchers in the literature. There are several methods that provide a set of sparse matches between two different images containing common objects, representing two views of a scene or, as in our case, two frames of a video. Some of them are based in the estimation of distinctive point location and matching [33,36]. Being based on local properties, they can be used to estimate arbitrarily large displacements. In our experiments, we will use SIFT [33] or Deep Matching [55]. Although SIFT is very effective, it has some disadvantages when dealing with small textureless objects or with non rigid deformations. The recent Deep Matching algorithm handles these problems and generates a more dense set of matchings. We have compared these two different methods, SIFT and Deep Matching, to compute the initial set of matches (the seeds). The first column of Figures 5 and 6 show some examples of the set of seeds obtained by the SIFT and Deep Matching algorithms. We use the SIFT implementation of [44], with its default parameters. Also, we use the Deep Matching implementation with the default parameters presented in [55].

With the aim of avoiding outliers in homogeneous areas, we initially remove the seeds having a low local saliency, which is determined by the minimum eigenvalue of the autocorrelation matrix computed locally. If the saliency is below a threshold (set to 0.045 in our experiments), these seeds are removed. This pruning is also used by other authors (see, *e.g.*, [43]).

In any case, we claim that our optical flow strategy is able to recover the dense motion regardless of the number of initial seeds; the only condition is to have at least one correct seed in every region  $\Omega_i \subset \Omega$  where the motion is smooth. As a proof of concept, in Fig. 1 we use a single seed per  $\Omega_i$ , which has been extracted from the ground truth flow. The estimated optical flow compares favourably to the ground truth even in these cases where the initial set of matches is extremely sparse. In contrast, other sparse-to-dense methods like [43] that produce state-of-the-art results are not capable to estimate the optical flow in this challenging situation. Another example is shown in Fig. 5 where similar optical flow results are obtained independently of the cardinality of the set of seeds, which have been computed using either the SIFT (Fig. 5, second row) or DeepMatching [55] (Fig. 5, third row) algorithms. Let us notice that DeepMatching produces in general more matches than SIFT. On the other hand, Fig. 6 shows a situation where the lack of seeds in some regions  $\Omega_i$  (second row, seeds provided by SIFT matches) produces an incorrect flow estimation in these areas. A much better estimation is obtained when there is at least one seed in these regions, as it happens when we use as seeds the matches provided by DeepMatching (Fig. 6, third row).

#### 4.3 Experimental results

The proposed method has been tested on several publicly available databases: Middlebury [5], MPI-Sintel [12], KITTI 2012 [21], and KITTI 2015 [34], as well as on proof of concept examples, chosen in order to obtain a better understanding of the fundamental parts of the proposed strategy. Let us remark that all results have been obtained by using the grayscale versions of the original color frames. The color version is only used to compute the seeds in case they come from the Deep Matching algorithm and to compute the local regularization weights in case of the Non-Local Total Variation as regularization term. All the results in this section have been obtained with the iterated faldoi minimization strategy presented in Sect. 3.1.2. We have fixed the parameters for all the experiments, more details are given in Appendix B.

First, we present a comparison of our strategy against the multi-scale approach using different functionals. In order to assure that the proposed framework is a real alternative to the coarse-to-fine warping strategy and, therefore, also valid for sequences that do not have large displacements of small objects, experiments on the Middlebury optical flow data set [5] have been performed for

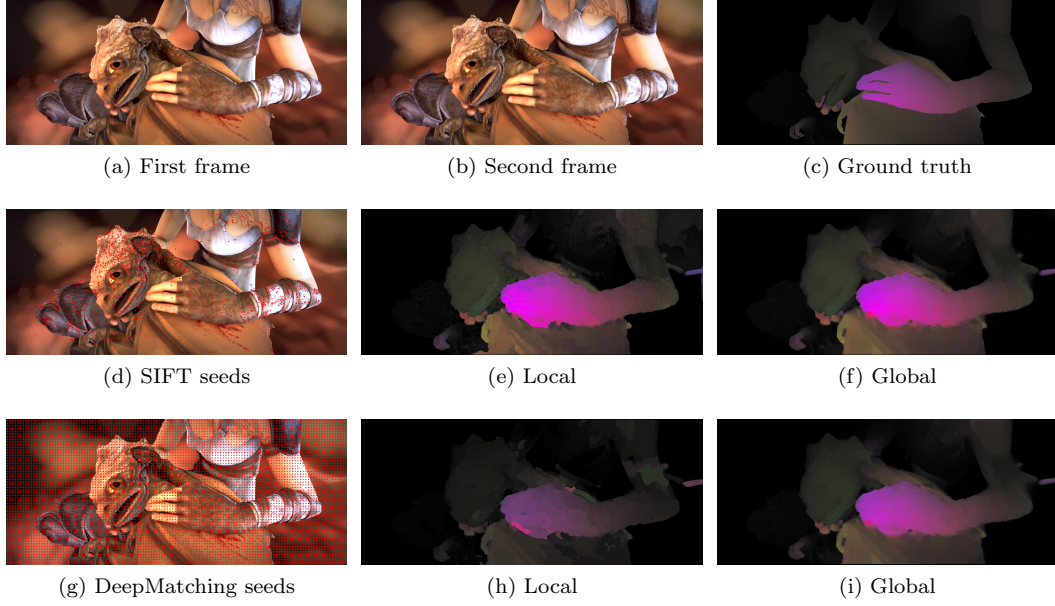


Fig. 5: Sequences with enough initial seeds at both trackers to recover a correct dense flow. The energy functional used is the classical  $TV_{\ell_2}-L1$ .



Fig. 6: Sequence where only the second tracker (Deep Matching [55]) gets enough seeds to estimate a correct flow. The energy functional used is the classical  $TV_{\ell_2}-L1$ .

both approaches. Table 1 shows how our framework achieves better results in all the samples in the dataset and for both energy functionals,  $TV_{\ell_2}-L1$  and  $NLTV-CSAD$ , even if the difference is slight in some cases. Some qualitative results are shown for images of the MPI Sintel database that contain large displacements in Figure 8 where the initial seeds have been computed using DeepMatching. The first row shows, from left to right, two consecutive frames and the optical flow ground truth (color coded). The image (d) of the second row shows the ground truth occlusion map (the occluded points are shown in white). The multi-scale results obtained using the  $TV_{\ell_2}-L1$  and

Middlebury	Dim.	Hyd.	Rub.	Gro2	Gro3	Urb2	Urb3	Ven
$TV_{\ell_2}$ -L1, multi-scale	0.1537	0.2286	0.1916	0.1496	0.6808	0.3709	0.6034	0.3563
$TV_{\ell_2}$ -L1, our min.	0.1243	0.2078	0.1876	0.1397	0.5945	0.3599	0.4354	0.3109
$NLTV$ -CSAD, multi-scale	0.1286	0.2003	0.1509	0.1405	0.5652	0.4315	0.5763	0.3647
$NLTV$ -CSAD, our min.	0.1075	0.1984	0.1477	0.1389	0.5414	0.4028	0.5521	0.2961

Table 1: Error measures in the Middlebury dataset with public ground truth. The first two lines correspond to the  $TV_{\ell_2}$ -L1 energy functional while the last two lines are based on the  $NLTV$ -CSAD functional. The results obtained minimizing with the classical coarse-to-fine scheme are shown in the first and third line and the ones obtained minimizing with our proposal are shown in the second and fourth line.

Functional, minimization	EPE all	EPE mat.	EPE unmat.	s0-10	s10-40	s40+
$TV_{\ell_2}$ -L1, multi-scale	6.6453	5.1181	15.8828	2.9261	8.0172	45.3955
$NLTV$ -CSAD, multi-scale	8.4243	7.0860	16.8037	2.8129	9.3389	50.4899
$TV_{\ell_2}$ -L1, our min. (SIFT)	6.7307	5.4784	14.8439	3.2928	8.3336	47.2940
$NLTV$ -CSAD, our min. (SIFT)	7.3337	6.1058	14.7972	4.4182	10.1687	49.8888
$TV_{\ell_2}$ -L1, our min. (Deep)	4.7688	3.5043	13.5503	2.3861	6.4305	32.3350
$TV_{\ell_2}$ -CSAD, our min. (Deep)	4.0158	2.7901	12.5350	2.0757	5.4551	31.5814
$NLTV$ -CSAD, our min. (Deep)	4.0060	2.7845	12.2698	2.3115	5.2989	30.8537

Table 2: Results for several methods (coarse-to-fine and our proposal), several energies ( $TV_{\ell_2}$ -L1 and  $NLTV$ -CSAD) and several initial sets of seeds (SIFT seeds and Deep Matching seeds) on a randomly selected subset of the MPI-Sintel *Final* training set.

the  $NLTV$ -CSAD energies are shown in (e) and (f), respectively, while the corresponding results obtained by our minimization strategy (iterated faldoi) are shown in (h) and (i), respectively. Finally, image (g) displays the result of our minimization faldoi strategy applied to the  $TV_{\ell_2}$ -CSAD energy functional. As it can be observed, the use of an advanced data term based on patches reduces the outlier area and better recovers the human shape. After adding the non-local regularizer the motion boundaries are more accurate and the outlier is removed.

Table 2 provides a quantitative comparison of different estimations of the optical flow obtained with different energy functionals ( $TV_{\ell_2}$ -L1 and  $NLTV$ -CSAD) and different set of seeds (SIFT seeds and Deep Matching seeds). We use a subset of 10% of frames randomly selected from the final training set of the MPI Sintel database. Let us notice that the energy based on the non-local Total Variation and a smooth approximation of the Census transform (CSAD) is the one which achieves the best results, both quantitative and qualitatively. On the other hand, Figure 7 shows the energy values using both approaches, multi-scale and our proposal (with Deep Matching seeds) over the same functionals. Notice how in general our method gives a lower or similar energy than the multi-scale method.

Table 3 and Figure 9 show a quantitative and a qualitative comparison, respectively, among our method and several state-of-the-art methods, including methods based on the combination of sparse matches and variational techniques [10, 55] via an extra coupling term and that can also be adapted to any energy. Our approach outperforms LDOF [10] in both Clean and Final versions of the MPI-Sintel dataset. Compared to DeepFlow [55] we get better results in the Clean version and in the Final version we get an EPE which is only a tenth worse, thus we can say that the results are comparable in the Sintel database. Our approach gives better results than S2D-Matching [30], which is another sparse-to-dense technique that uses edge information and occlusion estimation in the densification process. The images in Fig. 9 have been taken from the MPI-Sintel webpage. Compared to [10, 35, 55] and EpicFlow [43] (also a sparse-to-dense technique that achieves good positions in the Sintel table) our result better recovers the silhouettes of the girl (elbows, armpit, right hip and lower part of the legs) and the birds without using precomputed edge information as [43], although it produces some halos around the birds and the top-left corner of the image.

Table 4 shows a quantitative comparison among our method and several methods on the KITTI 2012 database [21]. As expected, the performance of our approach is better than LDOF [10], and

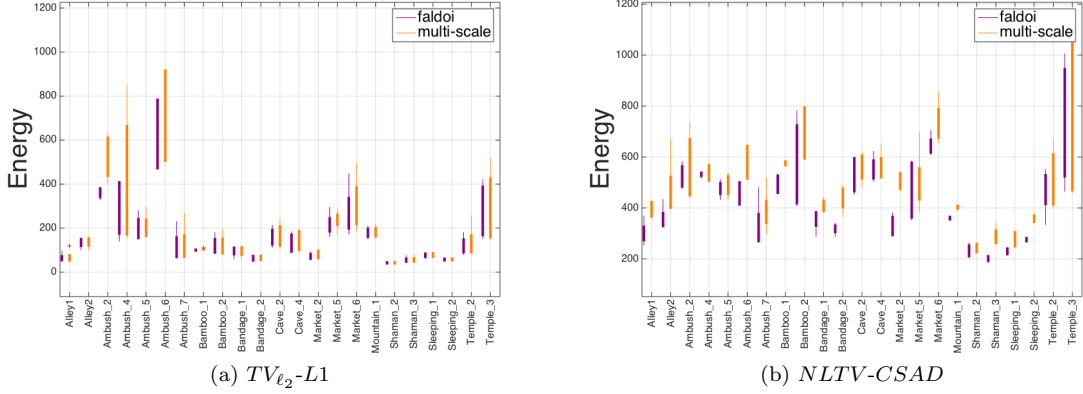


Fig. 7: Energy results for  $TV_{\ell_2}-L1$  and  $NLTV-CSAD$  using the multi-scale approach and our proposal on a randomly selected subset of the MPI-Sintel *Final* training set. It is the standard boxplot representation. Red dots are outliers, each box shows the first, second and third quartiles and the whiskers length corresponds to 1.5 interquartile range (IQR).

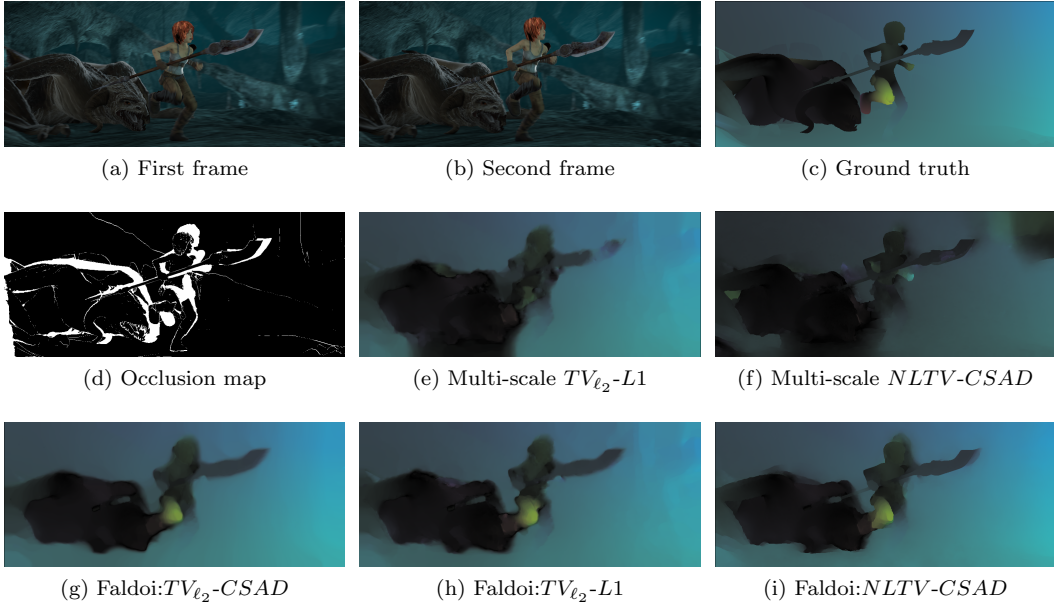


Fig. 8: Comparison of our strategy against the multi-scale approach using different functionals. The initial seeds have been computed using DeepMatching.

comparable to CRTflow [16], EpicFlow [43] and DeepFlow [55]. KITTI 2015 [34] can be used to evaluate optical flow and scene flow algorithms. Table 5 only contains results for optical flow methods.

In the approach of Brox et al. [10] or Weinzaepfel et al. [55] the matches are precomputed and then added as a constraint to the energy term. Thanks to these matches the motion of small objects that disappear at the coarser scales is recovered. However, these approaches need a minimum density of sparse matches over the area of the small object in order to correctly capture large displacements, even if the matches are weighted strong enough and enough iterations are performed. By contrast, our proposal only needs one single seed per area motion to recover the whole motion field. This is illustrated in Figure 10 where our proposal is able to recover the four large displacements

	EPE all	EPE mat.	EPE unmat.	s0-10	s10-40	s40+
<i>Final</i>						
<i>FlowFields</i> <sup>[4]</sup> [4]	5.810	2.621	31.799	1.157	3.739	33.890
<i>DiscreteFlow</i> <sup>[8]</sup> [35]	6.077	2.937	31.685	1.074	3.832	36.339
<i>EpicFlow</i> <sup>[10]</sup> [43]	6.285	3.060	32.564	1.135	3.727	38.021
<i>DeepFlow</i> <sup>[16]</sup> [55]	7.212	3.336	38.781	1.284	4.107	44.118
<i>Our approach</i> <sup>[21]</sup>	7.337	3.580	37.904	1.487	4.355	43.526
<i>S2D – Matching</i> <sup>[27]</sup> [30]	7.872	3.918	40.093	1.172	4.695	48.782
<i>LDOf</i> <sup>[49]</sup> [10]	9.116	5.037	42.344	1.485	4.839	57.296
<i>Clean</i>						
<i>DiscreteFlow</i> <sup>[4]</sup> [35]	3.567	1.108	23.626	0.703	2.277	20.906
<i>FlowFields</i> <sup>[5]</sup> [4]	3.748	1.056	25.700	0.546	2.110	23.602
<i>EpicFlow</i> <sup>[8]</sup> [43]	4.115	1.360	26.595	0.712	2.117	25.859
<i>Our approach</i> <sup>[14]</sup>	4.927	1.542	32.535	1.047	2.647	29.719
<i>DeepFlow</i> <sup>[20]</sup> [55]	5.377	1.771	34.751	0.960	2.730	33.701
<i>S2D-Matching</i> <sup>[38]</sup> [30]	6.510	2.792	36.785	0.622	3.012	44.187
<i>LDOf</i> <sup>[51]</sup> [10]	7.563	3.432	41.170	0.936	2.908	51.696

Table 3: Results on MPI-Sintel test set (9th of Jun. of 2016). The first and second set of results correspond, respectively, to the *Final* and *Clean* frames. For our results we used the *NLTV-CSAD* energy.

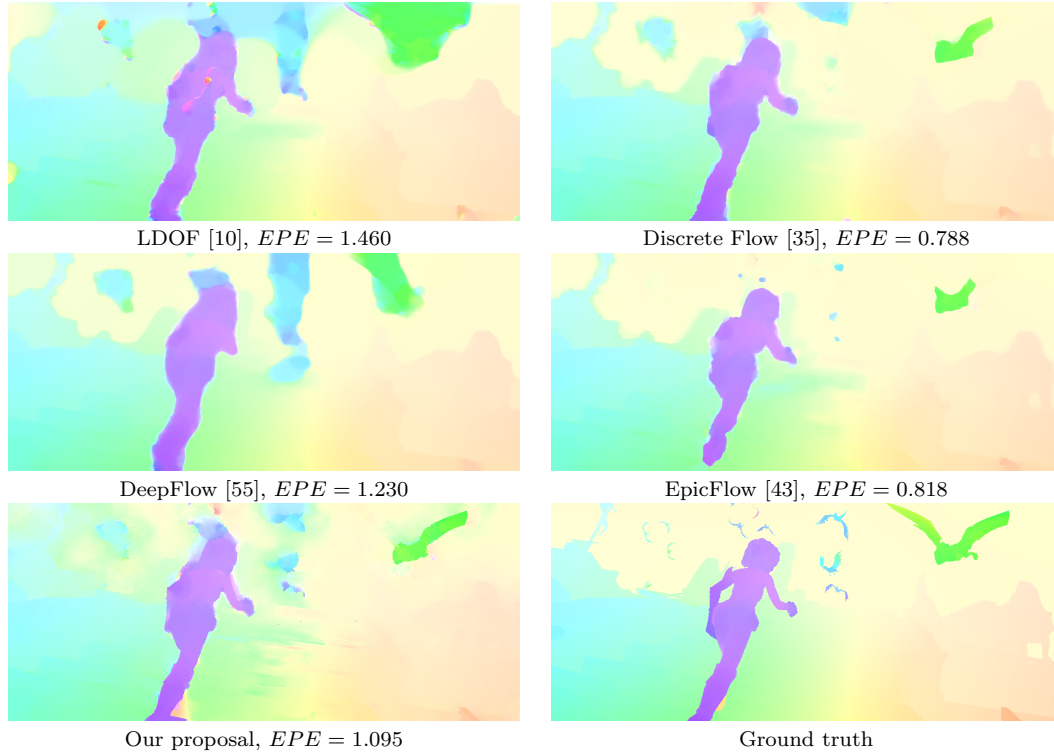


Fig. 9: Qualitative and quantitative comparison of different optical flow methods in a frame of the MPI-Sintel database (clean test). In our proposal we use the *NLTV-CSAD* energy with Deep Matching seeds.

	EPE-noc	EPE-all	Out-Noc3	Out-All 3
<i>FlowFields</i> <sup>[22]</sup> [4]	1.4	3.5	5.77	14.01
<i>DiscreteFlow</i> <sup>[27]</sup> [35]	1.3	3.6	6.23	16.63
<i>DeepFlow</i> <sup>[31]</sup> [55]	1.5	5.8	7.22	17.79
<i>EpicFlow</i> <sup>[32]</sup> [43]	1.5	3.8	7.88	17.08
<i>Our approach</i> <sup>[35]</sup>	1.7	5	8.81	19.93
<i>CRT flow</i> <sup>[37]</sup> [16]	2.7	6.5	9.43	18.72
<i>LDOF</i> <sup>[59]</sup> [10]	5.6	12.4	21.93	31.39

Table 4: Results on Kitti 2012 test (14th of Sep. of 2016). Out-Noc3 (resp. Out-all3) refers to the percentage of pixels where the estimated optical flow presents an error above 3 pixels in non-occluded areas (resp. all pixels). EPE-noc is the EPE over non-occluded areas and EPE-all is over all the image. For our results we used the *NLTV-CSAD* energy.

	Fl-bg	Fl-bf	Fl-all
<i>PatchBatch</i> <sup>[14]</sup> [20]	19.98	30.24	21.69
<i>DiscreteFlow</i> <sup>[17]</sup> [35]	21.53	26.68	22.38
<i>CPM – Flow</i> <sup>[19]</sup> [25]	22.32	27.79	23.23
<i>EpicFlow</i> <sup>[23]</sup> [43]	25.81	33.56	27.10
<i>Our approach</i> <sup>[24]</sup>	27.08	31.51	27.81
<i>DeepFlow</i> <sup>[26]</sup> [55]	27.96	35.28	29.18
<i>HS</i> <sup>[31]</sup> [52]	39.90	53.59	42.18

Table 5: Results on Kitti 2015 test (14th of Sep. of 2016). Fl refers the percentage of optical flow outliers. bg (resp. fg) refers of percentage of outliers only over background regions (resp. foreground regions) and all means over all ground truth pixels. For our results we used the *NLTV-CSAD* energy.

$\sigma$	10	20	30	40
<i>DeepFlow</i> without matches (pure multi-scale)	0.7468	1.0353	1.3832	1.4808
<i>DeepFlow</i> [55]	0.7766	1.5796	2.6128	4.3918
<i>EpicFlow</i> [43]	1.1418	1.6654	2.3386	3.2328
<i>Ours</i> ( $TV_{\ell_2}$ - $L1$ ) with SIFT matches	0.9486	1.5634	2.1090	2.9830

Table 6: Results over a set of five samples from the MPI-Sintel clean dataset with different standard deviation ( $\sigma$ ) levels of Gaussian noise.

with just a single seed in each region while DeepFlow does not succeed. We also include the result obtained with LDOF [10], but in this case it has been obtained with their own seeds (using their originally binary code).

Recent optical flow datasets as MPI-Sintel [12] contain different and challenging effects, such as illumination changes, large displacements, blur, etc. In MPI-Sintel these effects have been artificially created to provide naturalistic video sequences. However, the evaluation on these datasets does not take into account the robustness to the shot noise that appears in any real sequence, being noise one of the main limitations to any imaging system. Thus, we evaluate the robustness of several approaches to noise. To this goal, we corrupt the clean images from [12] with additive white Gaussian noise of standard deviation  $\sigma$ , for several values of  $\sigma$ . Sparse-to-dense techniques are very dependent of the initial seeds that are used to obtain a dense optical flow. It is important to note that our proposal works even with an extremely sparse set of initial seeds. This fact allows us to choose the best matching method according to the image peculiarities, without caring that much about the density of the correspondences. In particular, for highly noisy images SIFT correspondences are more robust than DeepMatching ones. Table 6 shows a comparative of our method with different optical flow estimation methods for different levels of Gaussian noise. We use the  $TV-L_1$  functional with SIFT matches. As shown in the table, the multi-scale method provides

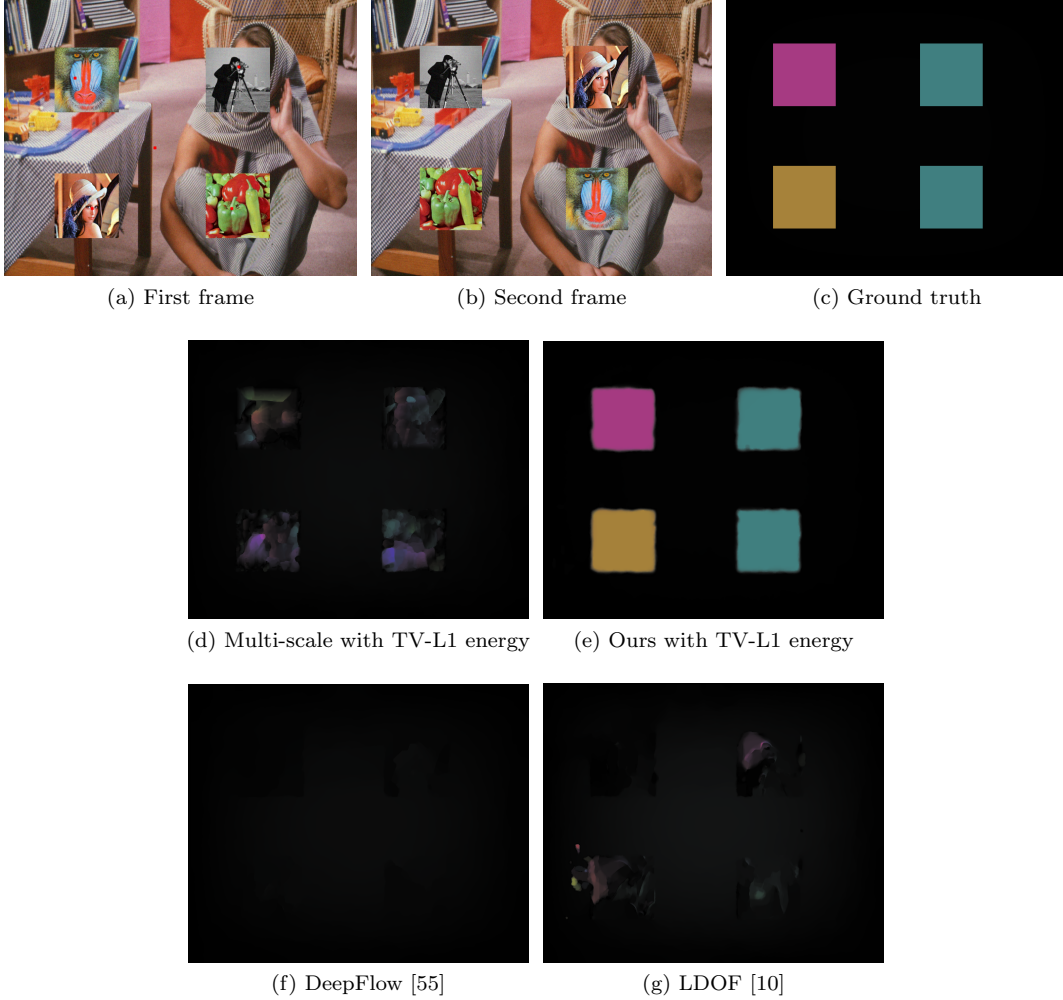


Fig. 10: Large displacement on a composition made of the images *Barbara*, *Baboon*, *Cameraman*, *Lena* and *Peppers*. It contains large diagonal translations and a slight deformation of the background. Five initial seeds have been manually selected and are shown in (a). The figures show the differences among Deepflow [55] which is a multi-scale strategy using seeds similar to LDOF approach, the multi-scale approach for the TV-L1 energy, and our proposal for minimizing the same energy. The last image shows the LDOF [10] result computed with its own seeds.

the best optical flow estimation in noisy images. For noise levels of standard deviation greater or equal to 10 and 20 respectively, our method produces better results than EpicFlow and DeepFlow.

## 5 Conclusions

We have proposed and provided an in-depth analysis of a method based on a variational formulation of the optical flow problem. The method works at the original scale of the image and finds a good local minimum of any energy functional using an adaptive coordinate descent strategy guided by a sparse set of initial matches. This is a general technique that consistently outperforms the multi-scale strategy for the same energy functional. With respect to alternative techniques that also include sparse matches in any energy functional, our performance is comparable to DeepFlow [55] and superior to LDOF [10] while being more robust to a low density of matches, high levels of noise and outliers in the matches. The only requirement is that at least one correct match is given

for each object in motion. For best overall results, we propose to use an energy with advanced data and regularization terms. Namely, we chose a smooth variant of the Census transform with a non-local  $TV$  regularization, providing robustness to illumination changes and occlusions while handling motion discontinuities. We present accurate quantitative and qualitative results that are comparable with state-of-the-art methods. As future work we plan to model the occlusions in the functional to reduce the halo effect in occluded regions.

## A Minimizing the energy

The numerical minimization algorithm for the general energy (1) is obtained in this paper by decoupling both terms. We linearize the image  $I_{t+1}$  near a given optical flow  $\mathbf{u}_0 = (u_{0,1}, u_{0,2})$  and make the following approximation  $I_{t+1}(\mathbf{x} + \mathbf{u}(x)) \approx I_{t+1}^{lin}(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ , where

$$I_{t+1}^{lin}(\mathbf{x} + \mathbf{u}(\mathbf{x})) = I_{t+1}(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) + I_{t+1}^x(\mathbf{x} + \mathbf{u}_0(\mathbf{x}))(u_1 - u_{0,1})(\mathbf{x}) + I_{t+1}^y(\mathbf{x} + \mathbf{u}_0(\mathbf{x}))(u_2 - u_{0,2})(\mathbf{x}),$$

and  $I_{t+1}^x, I_{t+1}^y$  denote the partial derivatives of  $I_{t+1}$  with respect to  $x$  and  $y$  respectively. Let us recall that the two data terms  $E_D$  that we have considered in Sect. 4.1 depend on  $I_t(\mathbf{x})$  and  $I_{t+1}(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ ; we will denote as  $E_{D,lin}$  the same data term but depending on  $I_t(\mathbf{x})$  and  $I_{t+1}^{lin}(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ . In order to decouple the fidelity term  $E_{D,lin}(u)$  and the regularization term  $E_R(\mathbf{u})$  in (1), we introduce an auxiliary variable  $\mathbf{v}$  representing the optical flow and we penalize its deviation from  $\mathbf{u}$ . Thus, the energy to minimize is

$$J(\mathbf{u}, \mathbf{v}) = J_{D,lin}(\mathbf{v}) + \beta J_R(\mathbf{u}) + \frac{1}{2\theta} \int_{\Omega} \|\mathbf{u} - \mathbf{v}\|^2, \quad (6)$$

depending on the two variables  $\mathbf{u}, \mathbf{v}$ , where  $\theta > 0$ . The decoupled energy (6) can be minimized by an alternating minimization procedure; alternatively fixing one variable and minimizing with respect to the other one. Sect. 4.1 presents the different possibilities for the energy.

1. For  $\mathbf{v}$  fixed, let us consider each of the two different regularization terms,  $J_R^1(\mathbf{u})$  and  $J_R^2(\mathbf{u})$ , presented in Sect. 4.1.
  - 1.1. In the case of  $J_R^1(\mathbf{u})$ , we reformulate the problem as a min-max problem incorporating the dual variables. Then, our minimization problem can be solved as a saddle-point problem. Following the notation of Osher *et al.* [22], for  $\mathbf{v} = (v_1, v_2)$  fixed, we solve

$$\int_{\Omega} \int_{\Omega} \omega(\mathbf{x}, \mathbf{y})(u_i(\mathbf{x}) - u_i(\mathbf{y}))p(\mathbf{x}, \mathbf{y})d\mathbf{y}d\mathbf{x} + \frac{1}{2\theta} \int_{\Omega} (u_i - v_i)^2 d\mathbf{x}, \quad (7)$$

for  $i = 1, 2$ , and  $p$  is the dual variable defined on  $\Omega \times \Omega$ . Let us explain it in detail. First, it is necessary to extend the notion of derivatives to a non-local framework. The non-local derivative can be written as

$$\partial_y u_i(\mathbf{x}) = \frac{u_i(\mathbf{x}) - u_i(\mathbf{y})}{d(\mathbf{x}, \mathbf{y})} \quad (8)$$

where  $d(\mathbf{x}, \mathbf{y})$  is a positive measure between two points  $\mathbf{x}, \mathbf{y}$ . By taking  $d(\mathbf{x}, \mathbf{y})$  such that  $w(\mathbf{x}, \mathbf{y}) = d(\mathbf{x}, \mathbf{y})^{-2}$ , the non-local gradient  $\nabla_w u_i(\mathbf{x}, \mathbf{y})$  is defined as the vector of all partial derivatives:

$$\nabla_w u_i(\mathbf{x}, \mathbf{y}) = (u_i(\mathbf{x}) - u_i(\mathbf{y}))\sqrt{w(\mathbf{x}, \mathbf{y})} \quad \mathbf{x}, \mathbf{y} \in \Omega. \quad (9)$$

Now, by writing  $\mathbf{p} := p(\mathbf{x}, \mathbf{y})$  for  $(\mathbf{x}, \mathbf{y}) \in \Omega \times \Omega$ , the non-local divergence  $\text{div}_w \mathbf{p}(\mathbf{x})$  is defined as the adjoint of the non-local gradient:

$$\text{div}_w \mathbf{p}(\mathbf{x}) = \int_{\Omega} (p(\mathbf{x}, \mathbf{y}) - p(\mathbf{y}, \mathbf{x}))\sqrt{w(\mathbf{x}, \mathbf{y})}d\mathbf{y}. \quad (10)$$

**Proposition 1** *The solution of (7) is given by the following iterative scheme*

$$p(\mathbf{x}, \mathbf{y})^{n+1} = \frac{p(\mathbf{x}, \mathbf{y})^n + \tau(\bar{u}_i^n(\mathbf{x}) - \bar{u}_i^n(\mathbf{y})\sqrt{w(\mathbf{x}, \mathbf{y})})}{1 + \tau|\nabla_w u_i(\mathbf{x}, \mathbf{y})|} \quad (11)$$

$$u_i^{n+1}(\mathbf{x}) = u_i^n(\mathbf{x}) - \sigma \left( \frac{(u_i^n(\mathbf{x}) - v_i(\mathbf{x}))}{\theta} - \text{div}_w \mathbf{p}(\mathbf{x}) \right) \quad (12)$$

$$\bar{u}_i^{n+1}(\mathbf{x}) = 2u_i^{n+1}(\mathbf{x}) - u_i^n(\mathbf{x}) \quad (13)$$

where  $u_i$  is the primal variable and  $\mathbf{p}$  is the dual variable.

- 1.2. In the case of  $J_R^2(\mathbf{u})$  we use the primal-dual algorithm that Chambolle proposed to minimize the ROF model [13] and which is based on a dual formulation of the TV. Then, our minimization problem can be solved as a saddle-point problem. For  $\mathbf{v}$  fixed, we solve

$$\min_{\mathbf{u}} \max_{\xi} \int_{\Omega} \langle D\mathbf{u}, \xi \rangle d\mathbf{x} + \int_{\Omega} \frac{1}{2\theta} \|\mathbf{u} - \mathbf{v}\|^2 d\mathbf{x} \quad (14)$$

where the dual variables are  $\xi = \begin{pmatrix} \xi_{11} & \xi_{12} \\ \xi_{21} & \xi_{22} \end{pmatrix}$  and satisfy  $\|\xi\|_F \leq 1$ .

**Proposition 2** *The solution of (14) is given by the following iterative scheme*

$$\xi_{i1}^{n+1} = \frac{\xi_{i1}^n + \tau \bar{u}_{ix}^n}{\max(1, \|\xi\|_2)}, \quad \xi_{i2}^{n+1} = \frac{\xi_{i2}^n + \tau \bar{u}_{iy}^n}{\max(1, \|\xi\|_2)}, \quad (15)$$

$$u_i^{n+1} = u_i^n - \sigma \left( \frac{(u_i^n - v_i)}{\theta} - \text{div}(\xi_{i1}^n, \xi_{i2}^n) \right), \quad (16)$$

$$\bar{u}_i^{n+1} = 2u_i^{n+1} - u_i^n, \quad (17)$$

where  $i = 1, 2$ .

2. For  $\mathbf{u}$  fixed, let us consider each of the two different data terms,  $J_D^1(\mathbf{v})$  and  $J_D^2(\mathbf{v})$ , presented in Sect. 4.1.

- 2.1. Case  $J_D^1(\mathbf{v})$ , Li and Osher [32] present a simple algorithm to find the optimal value of the function  $E(x) = \sum_i^n w_i |x - a_i| + F(x)$  when the  $w_i$  are non-negative and  $F$  is strictly convex. If  $F$  is also differentiable and  $F'$  is bijective, it is possible to obtain an explicit formula in terms of the median. For  $\mathbf{u}$  fixed, we solve

$$\int_{\Omega} C(\mathbf{v}, \mathbf{x}) d\mathbf{x} + \frac{1}{2\theta} \int_{\Omega} \|\mathbf{u} - \mathbf{v}\|^2 d\mathbf{x}. \quad (18)$$

Following the ideas of [54], we solve the discrete version of this problem. Due to the isotropy of the quadratic term, the optimal solution of  $C(\mathbf{v}, \mathbf{x})$  can be obtained solving a one dimensional problem. In particular, setting  $\mathbf{v} = \hat{\mathbf{v}} + \delta \frac{\nabla I(\mathbf{x} + \mathbf{v}_o)}{|\nabla I(\mathbf{x} + \mathbf{v}_o)|} + \delta \frac{\nabla^+ I(\mathbf{x} + \mathbf{v}_o)}{|\nabla^+ I(\mathbf{x} + \mathbf{v}_o)|}$  being  $\nabla^+ I$  an orthogonal vector to the gradient, where  $\delta$  is our new variable. Then, we minimize over  $\delta$

$$\frac{1}{2\tau} \delta^2 + \lambda \int_{\Omega} |\nabla I_{t+1}(\mathbf{x} + \hat{\mathbf{v}}_o)| |G(\hat{\mathbf{v}}) + \delta| d\mathbf{y} \quad (19)$$

where

$$G_{\mathbf{y}}(\hat{\mathbf{v}}) = \frac{I_t(\mathbf{x}) - I_t(\mathbf{y}) - I_{t+1}(\mathbf{x} + \hat{\mathbf{v}}_o) + I_{t+1}(\mathbf{y} + \hat{\mathbf{v}}_o) + (\hat{\mathbf{v}} - \hat{\mathbf{v}}_o)^T \nabla I_{t+1}(\mathbf{x} + \hat{\mathbf{v}}_o)}{|\nabla I_{t+1}(\mathbf{x} + \hat{\mathbf{v}}_o)|}$$

**Proposition 3** *The minimum of (19) with respect to  $\delta$  is*

$$\delta^* = \text{median}\{b_1, \dots, b_n, a_0, \dots, a_n\} \quad (20)$$

where  $b_i = -G_i(\hat{\mathbf{v}})$  and  $a_i = (n - 2i)\lambda |\nabla I_{t+1}(\mathbf{x} + \mathbf{v}_o)|$  for all the discrete neighbors  $i$  (corresponding to  $\mathbf{y}$  above), where  $n$  is the number of points in the discrete neighborhood.

2.2. Case  $J_D^2(\mathbf{v})$ . Notice that this term is a particular case of the previous data term. The functional to minimize

$$\int_W \lambda |\rho(\mathbf{v})| + \frac{1}{2\theta} \int_W \|\mathbf{u} - \mathbf{v}\|^2 d\mathbf{x}, \quad (21)$$

where  $\rho(\mathbf{v}) = I_{t+1}(\mathbf{x}, \mathbf{v}_o) + \langle \nabla I_{t+1}(\mathbf{x} + \mathbf{v}_o), (\mathbf{v} - \mathbf{v}_o) \rangle - I_t(\mathbf{x})$ , does not depend on spatial derivatives on  $\mathbf{v}$ . Then, a simple thresholding step gives an explicit solution [62].

**Proposition 4** *The minimum of (21) with respect to  $\mathbf{v}$  is*

$$\mathbf{v} = \mathbf{u} + \begin{cases} \lambda\theta\nabla I_{t+1} & \text{if } \rho(\mathbf{u}) < -\lambda\theta|\nabla I_{t+1}|^2 \\ -\lambda\theta\nabla I_{t+1} & \text{if } \rho(\mathbf{u}) > \lambda\theta|\nabla I_{t+1}|^2 \\ -\rho(\mathbf{u}) \frac{\nabla I_{t+1}}{|\nabla I_{t+1}|^2} & \text{if } |\rho(\mathbf{u})| \leq \lambda\theta|\nabla I_{t+1}|^2 \end{cases} \quad (22)$$

---

**Algorithm 4:** Global Minimization to compute the final optical flow.

---

**Input** : Two consecutive frames  $I_1, I_2$ , and an initial optical flow  $\mathbf{u}_0$

**Output:** Flow field  $\mathbf{u}$

---

```

1 Initialize  $\xi = \mathbf{v} = \mathbf{0}$ ;
2 Initialize  $\mathbf{u} = \mathbf{u}_0$ ;
3 for  $w \leftarrow 1$  to  $N_{warps}$  do
4   Compute  $I_{t+1}(\mathbf{x} + \mathbf{v}_0(\mathbf{x}))$ ,  $I_{t+1}^x(\mathbf{x} + \mathbf{v}_0(\mathbf{x}))$ ,  $I_{t+1}^y(\mathbf{x} + \mathbf{v}_0(\mathbf{x}))$ , using bicubic interpolation;
5   while  $n < N_{\max}$  or  $tol < error$  do
6     Compute  $\mathbf{v}$  via Prop. 3 or Prop. 4 ;
7     Compute  $\xi$  or  $p$  and  $\mathbf{u}$  via Prop. 1 or Prop. 2 ;
8   end
9 end
```

---

## B Implementation details

Our code is written in C. The numerical scheme to solve the functional  $E(u)$ , in both steps, is based on the implementation of [41]. Image warpings use bicubic interpolation. The image gradient is computed using centered-derivatives. Input images have been normalized between  $[0, 1]$ . The algorithm parameters are initialized with the same default setting for all the experiments. Both time steps are set to  $\tau = \sigma = 0.125$  to ensure convergence. As stopping criterion, the optical flow method uses the infinite-norm between two consecutive values of  $u$  with a threshold of 0.01. The coupling parameter  $\theta$  is set to 0.3. The smoothness term weight  $\beta$  is set to  $1/40$  for the  $TV_{\ell_2}$ -L1 functional and  $\beta = \frac{N-1}{80}$  for the  $NLTV$ -CSAD one, as suggested by [54], where  $N$  is the cardinality of the neighborhood considered in the CSAD term (we use a neighborhood of  $7 \times 7$  in the data term and then  $N = 49$ ) and we fixed  $\sigma_c = 2$  and  $\sigma_s = 2$  for the spatial and color domain of the  $NLTV$  term. For the iterated faldoi strategy we set  $MAX\_IT$  to 3. The size of the patch  $\mathcal{P}$  in the local minimization is  $11 \times 11$ . The complexity of our algorithm is  $\mathcal{O}(n)$ , where  $n$  is the number of pixels of the image frame. The basic faldoi algorithm takes around 20 seconds for  $TV_{\ell_2}$ -L1 energy and around 10 minutes for  $NLTV$ -CSAD over an *MPI - Sintel* image. Notice that our algorithm using  $NLTV$ -CSAD energy is very slow, especially because our implementation does not use parallelized code. As  $NLTV$ -CSAD can be easily parallelized, it should take the same time for both functionals using a GPU implementation.

**Acknowledgements** The authors would like to thank Gabriele Facciolo for all the discussions with him about the subject of this paper

## References

1. Alvarez, L., Deriche, R., Papadopoulos, T., Sánchez, J.: Symmetrical dense optical flow estimation with occlusions detection. In: European Conference on Computer Vision, pp. 721–735 (2002)
2. Alvarez, L., Weickert, J., Sánchez, J.: Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision* **39**(1), 41–56 (2000)
3. Ayvaci, A., Raptis, M., Soatto, S.: Sparse occlusion detection with optical flow. *International Journal of Computer Vision* **97**(3), 322–338 (2012)
4. Bailer, C., Taetz, B., Stricker, D.: Flow fields: Dense correspondence fields for highly accurate large displacement optical flow estimation. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 4015–4023 (2015)
5. Baker, S., Scharstein, D., Lewis, J.P., Roth, S., Black, M.J., Szeliski, R.: A database and evaluation methodology for optical flow. *International Journal on Computer Vision* **92**(1), 1–31 (2011)
6. Ballester, C., Garrido, L., Lazcano, V., Caselles, V.: A TV-L1 Optical Flow Method with Occlusion Detection, *Lectures Notes in Computer Science*, vol. 7476, p. 3140 (2012)
7. Bao, L., Yang, Q., Jin, H.: Fast edge-preserving patchmatch for large displacement optical flow. *Image Processing, IEEE Transactions on* **23**(12), 4996–5006 (2014)
8. Black, M.J., Anandan, P.: The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer vision and image understanding* **63**(1), 75–104 (1996)
9. Brox, T., Bruhn, A., Papenberger, N., Weickert, J.: High accuracy optical flow estimation based on a theory for warping. In: European Conference on Computer Vision, vol. 3024, pp. 25–36 (2004)
10. Brox, T., Malik, J.: Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **33**(3), 500–513 (2011)
11. Buades, A., Facciolo, G.: Reliable multi-scale and multi-window stereo matching. *SIAM J. Imaging Sciences* **8**(2), 888–915 (2015)
12. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: European Conference on Computer Vision, pp. 611–625 (2012)
13. 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), 120–145 (2011)
14. Chen, Z., Jin, H., Lin, Z., Cohen, S., Wu, Y.: Large displacement optical flow from nearest neighbor fields. In: *Computer Vision and Pattern Recognition, IEEE Conf. on*, pp. 2443–2450 (2013)
15. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 886–893 (2005)
16. Demetz, O., Hafner, D., Weickert, J.: Morphologically invariant matching of structures with the complete rank transform. *International Journal of Computer Vision* **113**(3), 220–232 (2015)
17. Dollár, P., Zitnick, C.L.: Structured forests for fast edge detection. In: *Computer Vision (ICCV), IEEE International Conference on*, pp. 1841–1848 (2013)
18. Fortun, D., Bouthemy, P., Kervrann, C.: Aggregation of local parametric candidates with exemplar-based occlusion handling for optical flow. *Computer Vision and Image Understanding*
19. Friedman, J., Hastie, T., Hfling, H., Tibshirani, R.: Pathwise coordinate optimization. *Tech. rep., Annals of Applied Statistics* (2007)
20. Gadot, D., Wolf, L.: Patchbatch: a batch augmented loss for optical flow. *CoRR* **abs/1512.01815** (2015). URL <http://arxiv.org/abs/1512.01815>
21. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2012)
22. Gilboa, G., Osher, S.: Nonlocal operators with applications to image processing. *Multiscale Modeling & Simulation* **7**(3), 1005–1028 (2008)
23. Hafner, D., Demetz, O., Weickert, J.: Why is the census transform good for robust optic flow computation? In: *Scale-Space and Variational Methods in Computer Vision*, pp. 210–221 (2013)
24. Horn, B.K.P., Schunck, B.G.: Determining optical flow. *Artificial Intelligence* **17**, 185–203 (1981)
25. Hu, Y., Song, R., Li, Y.: Efficient coarse-to-fine patchmatch for large displacement optical flow. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
26. Kannala, J., Brandt, S.: Quasi-dense wide baseline matching using match propagation. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8 (2007)
27. Kennedy, R., Taylor, C.J.: Optical flow with geometric occlusion estimation and fusion of multiple frames. In: *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pp. 364–377 (2015)
28. Krähenbühl, P., Koltun, V.: Efficient nonlocal regularization for optical flow. In: *European Conference on Computer Vision (ECCV)*, pp. 356–369 (2012)
29. Leordeanu, M., Sukthankar, R., Sminchisescu, C.: Efficient closed-form solution to generalized boundary detection. In: *European Conference on Computer Vision (ECCV)*, pp. 516–529 (2012)
30. Leordeanu, M., Zanfir, A., Sminchisescu, C.: Locally affine sparse-to-dense matching for motion and occlusion estimation. In: *Computer Vision, IEEE International Conf. on*, pp. 1721–1728 (2013)
31. Lhuillier, M., Quan, L.: Match propagation for image-based modeling and rendering. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(8), 1140–1146 (2002)
32. Li, Y., Osher, S.: A new median formula with applications to pde based denoising. *Commun. Math. Sci.* **7**(3), 741–753 (2009). URL <http://projecteuclid.org/euclid.cms/1256562821>
33. Lowe, D.G.: Object recognition from local scale-invariant features. In: *Proc. of the International Conf. on Computer Vision* (1999)

34. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
35. Menze, M., Heipke, C., Geiger, A.: Discrete optimization for optical flow. In: Pattern Recognition, pp. 16–28 (2015)
36. Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas J., S.F., Kadir, T., Van Gool, L.: An Algorithm for Total Variation Minimization and Applications. *International journal of computer vision* **65**(1-2), 43–72 (2005)
37. Müller, T., Rabe, C., Rannacher, J., Franke, U., Mester, R.: Illumination-robust dense optical flow using census signatures. In: Proc. of the DAGM Conf. on Pattern Recognition, pp. 236–245 (2011)
38. Nagel, H.H., Enkelmann, W.: An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **8**(5), 565–593 (1986)
39. Nesterov, Y.: Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization* **22**(2), 341–362 (2012)
40. Papenberg, N., Bruhn, A., Brox, T., Didas, S., Weickert, J.: Highly accurate optic flow computation with theoretically justified warping. *International Journal of Computer Vision* **67**(2), 141–158 (2006)
41. Palomares, R., Haro, G., Ballester, C.: A rotation-invariant regularization term for optical flow related problems. In: Computer Vision–ACCV, pp. 304–319. Springer (2014)
42. Ranftl, R., Bredies, K., Pock, T.: Non-local total generalized variation for optical flow estimation. In: Computer Vision–ECCV 2014, pp. 439–454 (2014)
43. Revaud, J., Weinzaepfel, P., Harchaoui, Z., Schmid, C.: EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow. In: Computer Vision and Pattern Recognition (2015)
44. Rey Otero, I., Delbracio, M.: Anatomy of the SIFT Method. *Image Processing On Line* **4**, 370–396 (2014). DOI 10.5201/ipol.2014.82
45. Sánchez, J., Salgado, A., Monzón, N.: Preserving accurate motion contours with reliable parameter selection. In: Image Processing (ICIP), IEEE International Conference on, pp. 209–213 (2014)
46. Stein, F.: Efficient computation of optical flow using the census transform. In: DGAM, pp. 79–86 (2004)
47. Steinbrücker, F., Pock, T., Cremers, D.: Advanced data terms for variational optic flow estimation. In: Proc. Vision, Modeling and Visualization (2009)
48. Steinbrücker, F., Pock, T., Cremers, D.: Large displacement optical flow computation without warping. In: International Conference on Computer Vision, pp. 1609–1614 (2009)
49. Stoll, M., Volz, S., Bruhn, A.: Adaptive integration of feature matches into variational optical flow methods. In: K.M. Lee, Y. Matsushita, J.M. Rehg, Z. Hu (eds.) Computer Vision – ACCV 2012: 11th Asian Conference on Computer Vision, Revised Selected Papers, Part III, pp. 1–14 (2013)
50. Strekalovskiy, E., Chambolle, A., Cremers, D.: Convex relaxation of vectorial problems with coupled regularization. *SIAM J. Imaging Sciences* **7**(1), 294–336 (2014)
51. Sun, D., Roth, S., Black, M.J.: Secrets of optical flow estimation and their principles. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 2432–2439 (2010)
52. Sun, D., Roth, S., Black, M.J.: A quantitative analysis of current practices in optical flow estimation and the principles behind them. *International Journal of Computer Vision* **106**(2), 115–137 (2014)
53. Timofte, R., Van Gool, L.: Sparse flow: Sparse matching for small to large displacement optical flow. In: Applications of Computer Vision (WACV), IEEE Winter Conf. on, pp. 1100–1106 (2015)
54. Vogel, C., Roth, S., Schindler, K.: An evaluation of data costs for optical flow. In: Pattern recognition, pp. 343–353 (2013)
55. Weinzaepfel, P., Revaud, J., Harchaoui, Z., Schmid, C.: DeepFlow : Large displacement optical flow with deep matching. *International Conference on Computer Vision* (2013)
56. Werlberger, M., Pock, T., Bischof, H.: Motion estimation with non-local total variation regularization. In: Computer Vision and Pattern Recognition, IEEE Conf. on, pp. 2464–2471 (2010)
57. Werlberger, M., Trobin, W., Pock, T., Wedel, A., Cremers, D., Bischof, H.: Anisotropic huber-l1 optical flow. In: BMVC, 2, p. 3 (2009)
58. Xu, L., Jia, J., Matsushita, Y.: Motion detail preserving optical flow estimation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **34**(9), 1744–1757 (2012)
59. Yang, J., Li, H.: Dense, accurate optical flow estimation with piecewise parametric model. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1019–1027 (2015)
60. Yoon, K.J., Kweon, I.S.: Adaptive support-weight approach for correspondence search. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **28**(4), 650–656 (2006)
61. Zabih, R., Woodfill, J.: Non-parametric local transforms for computing visual correspondence. In: European Conference on Computer Vision (ECCV), pp. 151–158 (1994)
62. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime tv-l1 optical flow. In: Proc. of the 29th DAGM Conference on Pattern Recognition, pp. 214–223 (2007)
63. Zimmer, H., Bruhn, A., Weickert, J.: Optic flow in harmony. *International Journal of Computer Vision* **93**(3), 368–388 (2011)
64. Zimmer, H., Bruhn, A., Weickert, J., Valgaerts, L., Salgado, A., Rosenhahn, B., Seidel, H.P.: Complementary optic flow. In: Energy minimization methods in computer vision and pattern recognition, pp. 207–220 (2009)