



HAL
open science

Aggregation of local parametric candidates with exemplar-based occlusion handling for optical flow

Denis Fortun, Patrick Bouthemy, Charles Kervrann

► To cite this version:

Denis Fortun, Patrick Bouthemy, Charles Kervrann. Aggregation of local parametric candidates with exemplar-based occlusion handling for optical flow. *IEEE Transactions on Image Processing*, 2014, pp.27. hal-01001758v1

HAL Id: hal-01001758

<https://inria.hal.science/hal-01001758v1>

Submitted on 21 Jul 2014 (v1), last revised 22 Dec 2015 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Aggregation of Local Parametric Candidates with Exemplar-based Occlusion Handling for Optical Flow ¹

DENIS FORTUN, PATRICK BOUTHEMY AND CHARLES KERVRANN
Inria, Centre Rennes - Bretagne Atlantique, Rennes, France

July 16, 2014

Abstract

Handling all together large displacements, motion details and occlusions remains an open issue for reliable computation of optical flow in a video sequence. We propose a two-step aggregation paradigm to address this problem. The idea is to supply local motion candidates at every pixel in a first step, and then combine them to determine the global optical flow field in a second step. We exploit local parametric estimations combined with patch correspondences and we experimentally demonstrate that they are sufficient to produce highly accurate motion candidates. The aggregation step is designed as the discrete optimization of a global regularized energy. The occlusion map is estimated jointly with the flow field throughout the two steps. We propose a generic exemplar-based approach for occlusion filling with motion vectors. We achieve state-of-the-art results in computer vision benchmarks, with particularly significant improvements in the case of large displacements and occlusions.

Keywords: optical flow, occlusion, large displacement, local parametric motion, aggregation framework.

1 Introduction

Optical flow is a key information when addressing important problems in computer vision such as moving object segmentation, object tracking, egomotion computation, obstacle detection or action recognition. The challenge for an optical flow estimation method is to deal with a large variety of image contents and motion types. Optical flow has been historically evaluated on sequences exhibiting small displacements and smooth motion fields, like in the Yosemite sequence [7]. Once initial issues were solved, other challenges were addressed [44], and new situations have been proposed by new benchmarks [4, 19]. Various and sometimes opposite motion conditions must be handled together, as illumination changes, large areas of smooth motion, motion details, large displacements, motion discontinuities, occluded regions (i.e., points disappearing in the next image).

Optical flow methods first rely on a data constancy assumption, e.g., applied to image intensity or spatial intensity gradient. Then, it is combined with a spatial, or sometimes space-time, coherency

¹This work was realized as part of the Quaero program, funded by OSEO, French State agency for innovation. It was also partly supported by Institut Curie, CNRS UMR 144 and the France-BioImaging project granted by the "Investissement d'Avenir" program.

constraint on the expected velocity field. Existing approaches can be broadly classified into *local* and *global* methods.

Local spatial coherency arises when considering a parametric motion model, e.g., local translation [42], 4-parameter sub-affine model, affine model, 8-parameter quadratic model [48], in a given neighborhood or an appropriate local region. Optimization requires that the neighborhood is sufficiently textured or contains interest points such as corners, to supply accurate and reliable velocity vectors.

In contrast, global methods express the flow field coherency by imposing a global smoothness constraint in addition to the data constancy term, known as the regularization term of the global energy as pioneered by [31, 47]. Global methods overcome uncertainty yielded by local supports in uniform intensity regions by diffusing motion from informative to non informative regions via the global regularization constraint. The optimization problem of seminal model [31] was optimally solvable, but the estimation was affected by oversmoothing and was limited to small displacements.

Numerous modifications of this original model, starting with [9, 30], have been designed to resolve these two crucial issues, namely, handling of large displacements and preservation of motion discontinuities. It was usually achieved by introducing a multi-resolution and incremental coarse-to-fine framework along with piecewise smoothing or robust estimation. The data-driven term of the global optimization has also received attention. Image features like image gradient [15], texture component [58] or Census transform [28], and matching criteria like Normalized Cross Correlation (NCC) [60], convey invariance properties to overcome limitations of the classical intensity constancy assumption. Several data-driven terms were considered in [36] and managed by a locally adaptive fusion scheme. However, intricate optimization issues came with the increasing complexity of the modeling.

Existing local methods are far from being able to compete with global models in terms of accuracy in computer vision benchmarks. However, several works based on joint estimation and segmentation of the motion field have shown that when appropriate segmented regions are found, affine models can be very accurate representations [54, 56]. However, the alternate optimization schemes involved are sensitive to the initialization of the region supports.

In this paper, we define a new method for optical flow computation called *AggregFlow* which exhibits several distinctive features. First, we advocate the systematic computation of affine motion models over a set of size-varying square patches combined with patch-based pairings. Indeed, we experimentally demonstrate that the sets of motion vectors computed that way comprise at least one accurate motion vector for each pixel. On this basis, we build an optical flow estimation method composed of a first step computing local parametric candidates followed by a second step aggregating these candidates to produce the global flow field. The motion vector candidates are independently estimated on local supports without segmentation step. The aggregation is performed by a discrete optimization algorithm which selects one candidate at each pixel while ensuring piecewise smoothing of the resulting flow field.

Secondly, we address the occlusion problem in an original way by blending it with the motion estimation issue through the two steps of *AggregFlow*. Motion candidates are extended in occluded areas with an exemplar-based technique. The estimated parametric model of the dominant motion in the image also contributes to create supplementary motion candidates. We extract local occlusion cues in the first step of *AggregFlow* and exploit them to guide the joint estimation of the occlusion map and motion field in the aggregation step. Motion estimation in occluded regions is performed with a generic global exemplar-based approach. Specifically, we properly deal with large displacements producing large occluded regions.

Our method can thus be viewed as a novel and efficient combination of local and global ap-

proaches for occlusion-aware optical flow computation. The main original features and contributions of our method AggregFlow are listed below:

- Motion candidates are locally estimated by a general parametric patch-based method which ensures relevant and accurate motion vectors at every point among all the computed candidates.
- Feature matching is integrated in an original and efficient way in the two-step aggregation framework.
- We define a generic exemplar-based method for occlusion filling with motion vectors.
- We propose a joint motion and occlusion estimation framework based on a sparse model guided by a local occlusion confidence map.
- AggregFlow outperforms existing methods on the MPI Sintel benchmark which involves large displacements and occlusions, and it is competitive in the Middlebury benchmark composed of videos depicting smaller movements.

A preliminary approach without any occlusion handling and dedicated to a specific application was presented in [26].

The paper is organized as follows. Section 2 describes related work. In Section 3, we present the parametric computation of motion candidates and the local detection of occlusions. Section 4 is devoted to the aggregation stage. In Section 5, we report experimental results demonstrating the performance of AggregFlow. Section 6 contains concluding remarks.

2 Related work

Hereunder, we briefly review the literature on optical flow computation while focusing on issues related to our contributions.

2.1 Feature correspondences and large displacements

The integration of feature correspondences in dense motion estimation has been investigated in several recent works. A first class of methods integrates feature correspondences in a global energy model. Variational methods [14, 16, 29, 59] include an additional term to a classical global energy to impose the flow to be close to pre-computed correspondences. Giving a fixed weight to the correspondences, this approach is sensitive to matching errors. To overcome this problem, [14, 59] focused on improving the matching step. Another class of methods use correspondences to reduce the search space for discrete optimization and provide a coarse initialization for subsequent refinement [21, 46, 64]. The main motivation of the attempts based on feature matching is to get rid of the drawbacks of the coarse-to-fine scheme imposed by variational optimization, in particular the loss of large displacements of small objects.

Our patch correspondence is related to [21, 46, 64] in the sense that it is used in the candidates generation process. However, our candidates are not coarse approximations to be refined in a global subsequent step and we do not adopt any global variational optimization.

2.2 Occlusions

Occlusions play a crucial role for motion estimation [53], especially under large displacements, since no motion measurements are available in occluded areas. Therefore, a proper occlusion handling must distinguish between *occlusion detection*, segmenting the image into occluded and non-occluded regions, and *occlusion filling*, applying a specific treatment to motion estimation in occluded regions. Occlusion detection has been mostly undertaken as a subsequent operation to motion computation, by thresholding a consistency measure issued from the estimated motion field, like geometric forward-backward motion mismatch [33], mapping unicity [64] or data constancy violation [61]. Several flows and image criteria have been combined in a learning framework [32]. The main limitation of the latter is that accuracy of occlusion detection is highly dependent on the quality of the initial motion estimation. To overcome this problem, other approaches estimate the occlusion map jointly with the motion field [3, 33, 37, 50]. Our occlusion detection falls in the latter category.

The problem of filling occluded regions with estimated velocity vectors when the occlusion map is known is closely related to the image inpainting problem. Inpainting methods can be coarsely divided into two classes, diffusion-based methods [8, 20] and exemplar-based methods [23, 39]. A synthesis of these two approaches has been investigated in [17] in a variational framework. Occlusion filling is usually tackled by diffusion-based (or geometry-oriented) methods, propagating motion from non-occluded regions to occluded regions via partial derivative equation (PDE) resolution [3, 33, 50, 64]. In exemplar-based image inpainting, the missing part is filled by copying pixels of the observed images. The framework is non local in the sense that similar pixels can be sought anywhere in the image. We adapt this strategy to occlusion filling with motion vectors.

2.3 Parametric motion estimation

The use of a parametric model has been widely investigated in motion estimation [10, 22, 26, 43, 48, 54]. Applied on the whole image domain, affine or quadratic models are adequate to estimate the dominant image motion induced by the camera motion [48]. For accurate dense motion estimation, parametric approximations are only valid locally. Local regions are usually defined as square patches centered on each pixel [10, 42], possibly with an adaptation of the patch size [52], or its position [34]. It has the merit of being easy to implement with a low computational cost, but it is clearly outperformed by sophisticated extensions of [31] introduced in modern global optical flow methods.

As aforementioned, more complex region shapes can be estimated by joint motion segmentation and estimation. Existing approaches can be divided in two classes. A first class of methods relies on an independent image color segmentation and tries to fit parametric motion in each region [11, 12, 27, 62, 67], possibly with the help of an independent global variational estimation [11, 62]. The drawback is that image color segmentation may lead to an over-segmentation of the motion field. The second class of methods jointly estimates region supports and parametric motion models for each region [22, 49, 54, 56]. It is achieved by minimizing a global energy with respect to supports and motion parameters of the regions. However, the global energy is highly non-convex and consequently difficult to minimize and particularly sensitive to the initialization of the optimization procedure.

The motion field produced by `AggregFlow` is composed of affine motion vectors estimated in square patches without any motion segmentation. `AggregFlow` implicitly selects the best patch size and position when selecting the best motion candidate for each pixel in the second step.

2.4 Motion discontinuities

In the variational setting, the problem of preserving discontinuities has been addressed by modifying the regularization term [47]. The seminal work of [31] used a quadratic penalty function on the gradient magnitude of motion vectors. The first attempt to preserve discontinuities was investigated in [30] where a binary map of local motion discontinuities was introduced and estimated jointly with the motion field using two interwoven Markov Random Fields (MRF). The regularization is thus canceled on motion discontinuities. Subsequent improvement has then been reached with the use of robust penalty functions in the regularization term [10, 43]. The robust L_1 norm is often retained in variational settings owing to its convexity [15, 60].

2.5 Discrete optimization and aggregation paradigm

Discrete optimization is an alternative to variational methods and is often able to find good local minima for non differentiable and non-convex energy functionals. To combine the subpixel accuracy of the continuous variational approach and the efficiency of discrete minimization, the authors of [40] built a discrete motion space from motion fields delivered by several global variational estimations with different parameter settings. An energy function is then optimized by successive fusions of global proposals, which are efficiently performed by a binary graph-cut method. In [25], we followed a similar approach but with a semi-local patch-based variational estimation of candidate motion vectors. In [1], a set of candidate motion vectors is computed at each pixel using phase correlation in overlapping patches. The candidates are then linearly combined to create a global motion field. Recent works [21, 46] also exploit discrete graph-cut optimization in a two-step paradigm. However, the principle is different than ours. Indeed, the motion candidate generation step only aims at finding dominant displacements and the aggregation provides a coarse initialization for a subsequent global refinement. Discrete optimization is also associated with a variational framework in [64] as an intermediate stage between scales of a coarse-to-fine framework, in order to limit the loss of details of the flow. Another aggregation-related work is the image colorization method of [18]. Color candidates are obtained with patch correspondences, and a candidate is selected at each pixel by minimizing a global energy in a variational setting.

3 Local motion candidates and occlusion cues

We describe in this section the first step of our method *AggregFlow*. It exploits local information to supply motion candidates and occlusion cues. A set of motion vector candidates is generated at every pixel by a combination of patch correspondences and local parametric motion model estimations. A specific treatment is applied to occluded regions by exemplar-based extension of the motion candidates set. We also exploit the dominant motion in the image due to camera motion. Motion candidates and occlusion cues form the input of the second stage of *AggregFlow* described in Section 4.

Our approach can be viewed as a new way to address the problem of choosing the local neighborhood for parametric estimation. Rather than adapting the regions *a priori* or jointly with the motion field, we operate in two steps: 1) estimation of motion candidates on several supports at every pixel, 2) implicit selection of the best support through the selection of the optimal candidate at each pixel within the aggregation step. In the sequel, we denote two consecutive image frames as $I_1, I_2 : \Omega \rightarrow \mathbb{R}$, with Ω denoting the image domain.

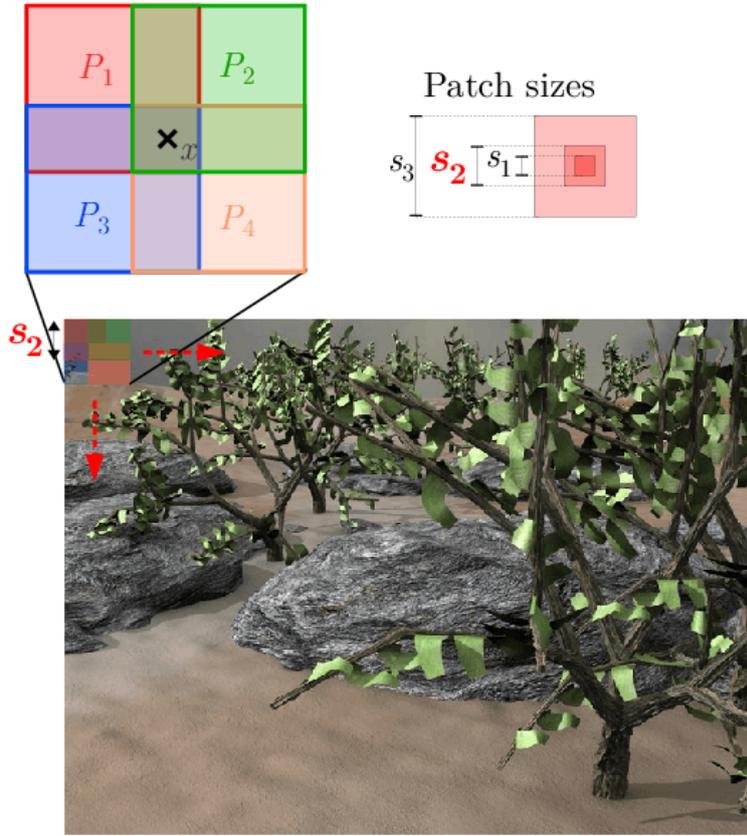


Figure 1: Four patches of set $\mathcal{P}_{s_2, \alpha}$ for a given size s_2 of the set $\mathcal{S} = \{s_1, s_2, s_3\}$, and overlapping ratio $\alpha = 0.3$. The pixel x is contained in the patches P_1, \dots, P_4 . Motion estimation in each of these patches provide motion candidates for x .

3.1 Local parametric motion candidates

3.1.1 Set of overlapping patches in I_1

The local supports for motion candidates computation are overlapping square patches of different sizes. Let us denote $\mathcal{P}_{s, \alpha}$ the patch set for a fixed patch size s and an overlapping ratio $\alpha \in [0, 1]$ indicating the proportion of surface shared by neighboring patches (see illustration of Fig. 1). Let $\mathcal{S} = \{s_1, \dots, s_n\}$ be a set of n patch sizes, we then define $\mathcal{P}_{\mathcal{S}, \alpha} = \bigcup_{s \in \mathcal{S}} \mathcal{P}_{s, \alpha}$. To capture different motion scales, the patch sizes must cover a large range of values. In all our experiments, we will use $\mathcal{S} = \{16, 44, 104\}$. Due to the overlap and the number of patch sizes ($n > 1$), one given pixel $x \in \Omega$ belongs to several patches. The motion vectors are estimated independently in each patch in two sub-steps described below: patch correspondences and affine motion estimations.

3.1.2 Patch correspondences

For each patch $P_1 \in \mathcal{P}_{\mathcal{S}, \alpha}$, we first determine the set $\mathcal{M}_N(P_1)$ of the N most similar patches to P_1 in I_2 . Let us put forward that we do not aim at keeping at this stage the best correspondence only but at selecting N relevant correspondences to subsequently constitute motion candidates.

The matching step is generic and could be achieved with any arbitrary feature matching algorithm. We use a combination of the saturation and value channels of the HSV color space to gain partial robustness to illumination changes [66] and we use the Sum of Absolute Distances (SAD) to compare patches. To avoid that the set $\mathcal{M}_N(P_1)$ uselessly contains too close patches, we impose a minimal distance between two patches of $\mathcal{M}_N(P_1)$. Hence, for each established pair of corresponding patches $P_{1,2} = (P_1, P_2)$ with $P_2 \in \mathcal{M}_N(P_1)$, we get the translation vector $\mathbf{w}_{P_{1,2}} \in \mathbb{Z}^2$ shifting P_1 onto P_2 .

3.1.3 Affine motion refinement

The displacements estimated by patch correspondences are integer-pixel translational approximations. To reach subpixel accuracy and to allow for more complex motion, we refine the first sub-step of coarse translational motion vector $\mathbf{w}_{P_{1,2}}$ with the estimation of a local affine motion model in every pair $P_{1,2}$. Denoting Ω_{P_1} the pixel domain of P_1 , the affine motion model $\delta\mathbf{w}_{P_{1,2}} : \Omega_{P_1} \rightarrow \mathbb{R}^2$ between P_1 and P_2 is defined at a pixel $x = (x_1, x_2)^\top$ as:

$$\delta\mathbf{w}_{P_{1,2}}(x) = (a_1 + a_2x_1 + a_3x_2, a_4 + a_5x_1 + a_6x_2)^\top. \quad (1)$$

The parameter vector $\boldsymbol{\theta}_{P_{1,2}} = (a_1, a_2, a_3, a_4, a_5, a_6)^\top$ of the affine model is estimated assuming brightness constancy:

$$\hat{\boldsymbol{\theta}}_{P_{1,2}} = \arg \min_{\boldsymbol{\theta}_{P_{1,2}}} \int_{\Omega_{P_1}} \phi(P_2(x + \mathbf{w}_{P_{1,2}} + \delta\mathbf{w}_{P_{1,2}}(x)) - P_1(x)) dx \quad (2)$$

where the penalty function $\phi(\cdot)$ is chosen as the robust Tukey's function. The problem (2) is solved with the publicly available Motion2D software² [48], which implements a multi-resolution incremental minimization scheme involving an IRLS (Iteratively Reweighted Least Squares) technique.

3.1.4 Final set of motion candidates

The above described two-step estimation is repeated for every patch of $\mathcal{P}_{\mathcal{S},\alpha}$ and generates a set of candidate motion vectors $\mathcal{C}(x)$ at each pixel $x \in \Omega$ defined as follows:

$$\mathcal{C}(x) = \{\mathbf{w}_{P_{1,2}}(x) + \delta\mathbf{w}_{P_{1,2}}(x) : P_1 \in \mathcal{P}_{\mathcal{S},\alpha}(x), P_2 \in \mathcal{M}_N(P_1)\}, \quad (3)$$

where $\mathcal{P}_{\mathcal{S},\alpha}(x) = \{P \in \mathcal{P}_{\mathcal{S},\alpha} : x \in P\}$.

Let us make a few comments on the estimation scheme for computing motion candidates. A coarse motion estimation followed by a refinement step has been investigated in several previous works [21, 41, 46], but it has always been dedicated to global motion fields. In our case, the refinement is local and adapted to each patch correspondence. Classical local motion estimation methods based on [42] also rely on square patches, but assign the computed motion vector only to the center point of each patch. On the opposite, parametric motion estimation in segmented regions as in [22] apply to regions of arbitrary shape. Our patch distribution can be considered as an intermediate level between these two extremes. Indeed, we use square patches as in [42] and thus avoid the complex segmentation step. However, we exploit the whole vector field issued from the affine model estimated in each patch. As a consequence, every pixel inherits several motion candidates from the affine motion estimations performed in patches of different positions and sizes which the given pixel belongs to. Finally, in contrast to several other methods using

²<http://www.irisa.fr/vista/Motion2D/>

feature correspondences [16, 21, 59], we do not select one single patch correspondence but we keep the N best ones.

The interest of the local set of motion candidates supplied by AggregFlow is three-fold. First, the correspondence sub-step enables to capture large displacements even for small patch sizes. Thus, it allows us to correctly deal with small structures undergoing large displacements in contrast to coarse-to-fine schemes. Second, by considering a large variety of patches, we get rid of the predefined choice of the local neighborhood encountered in parametric motion estimation. The selection of the proper patch via its corresponding motion candidate is transferred to the aggregation stage. Third, introducing patches of several sizes enables to tackle motion of different scales.

3.2 Motion candidates in occluded areas

The generation of motion candidates described in Section 3.1 does not differentiate between occluded and non-occluded pixels. For a given pixel x , if all the patches of $\mathcal{P}_{\mathcal{S},\alpha}(x)$ mainly contain occluded pixels, there is no chance to correctly estimate a relevant motion candidate at x in that way. Therefore, we compute motion candidates in occluded regions in a specific manner.

Let us define the occlusion map $o : \Omega \rightarrow \{0, 1\}$

$$o(x) = \begin{cases} 1 & \text{if } x \text{ is occluded,} \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

The occluded regions are denoted $\mathcal{O} = \{x \in \Omega : o(x) = 1\}$. The computation of map o will be addressed in Section 3.4 and Section 4, and we assume for now that o is known.

3.2.1 Occlusion filling with motion vectors

When occluded regions are known, occlusion filling with motion vectors is conceptually closely related to image inpainting, since it recovers motion in regions where motion is by definition *not observable*: The occluded pixels do not appear in the next image and consequently have no corresponding points. Classical methods for motion-based occlusion filling operate in a variational framework by cancelling the data term and letting the diffusion process of the regularization propagate the optical flow in occluded regions [3, 64]. The diffusion-based class of inpainting methods [8] acts similarly. They perform well in case of thin missing areas or cartoon-like images, but they are usually outperformed by exemplar-based inpainting methods [23] for large missing regions. In order to deal with large occlusions produced by large displacements, we follow the inpainting analogy and we overcome the problem of local motion candidates estimation in occluded areas by designing an exemplar-based scheme. In the first step of AggregFlow, the motion candidates set is thus augmented by copy-paste operations.

3.2.2 Exemplar-based candidates extension

We rely on the assumption that motion at an occluded pixel $x \in \mathcal{O}$ is similar to the motion of a close non-occluded pixel $m_o(x) \in \Omega \setminus \mathcal{O}$ belonging to the same object or the same background part. To provide relevant motion candidates at x , we copy motion candidates from $\mathcal{C}(m_o(x))$ to $\mathcal{C}(x)$. The search domain $\mathcal{V}_o \subset \Omega \setminus \mathcal{O}$ for $m_o(x)$ is constrained to be close to the occlusion boundaries. Figure 2(e) represents the occluded regions \mathcal{O} (in white) and the search domain \mathcal{V}_o (in red), and Fig. 2(f)

superimposes the two sets on I_1 . Searching for pixel $m_o(x)$ for $x \in \mathcal{O}$ is actually easier for motion-based occlusion filling than for image inpainting. Indeed, occluded regions are not completely uninformative, while inpainted regions are, since we have access to the information supplied by image I_1 even in \mathcal{O} . Thus, as $m_o(x)$ is expected to belong to the same object as x , we use color similarity to find the match in I_1 :

$$m_o(x) = \arg \min_{y \in \mathcal{V}_o} D(I_1, x, y), \quad (5)$$

where $D(I_1, x, y)$ is the distance between patches centered respectively in x and y . As in Section 3.1, we resort to a SAD in the HSV space.

An extended candidate set $\mathcal{C}_+(x)$ is created for occluded pixels by adding to the initial set $\mathcal{C}(x)$ the motion candidates of their matched pixel $m_o(x)$:

$$\mathcal{C}_+(x) = \mathcal{C}(x) \cup \mathcal{C}(m_o(x)), \quad \forall x \in \mathcal{O}. \quad (6)$$

By convention, $\forall x \in \Omega \setminus \mathcal{O}$, $\mathcal{C}_+(x) = \mathcal{C}(x)$.

3.2.3 Occlusions due to camera motion

A particular class of occluded (or disappearing) regions occurs at image borders in the case of large camera motion (Fig. 3). We cope with this issue by estimating the dominant image motion due to camera motion. To do so, we use again the robust parametric estimation described in Section 3.1, but now, we apply it to the whole image [48], to retrieve the dominant motion. We found in our experiments that the quadratic model was more adequate to accurately cope with large and sometimes complex camera motion. The resulting parametric motion field $\mathbf{w}_{cam} : \Omega \rightarrow \mathbb{R}^2$ is added to the motion candidates, and we end up with the final set of motion candidates \mathcal{C}_f :

$$\mathcal{C}_f(x) = \mathcal{C}_+(x) \cup \{\mathbf{w}_{cam}(x)\}, \quad \forall x \in \Omega. \quad (7)$$

The camera motion candidates are mostly useful for occluded pixels, but it can sometimes provide relevant motion candidates in unoccluded regions of the background as well, so that we finally add it to all pixels in Ω .

3.3 Best candidate flow

To validate our method for computing motion candidates, we have exploited sequences from MPI Sintel and Middlebury datasets [4, 19] provided with ground truth. We create the *Best Candidate Flow* (BCF) by selecting at each pixel x the candidate motion vector of $\mathcal{C}_f(x)$ closest to the ground-truth vector. In order to evaluate our occlusion module, we distinguish between the BCF determined with the candidates extension described in the preceding section (or full BCF) and the BCF without it. Parameters involved in the local motion computation are set to $\mathcal{S} = \{16, 44, 104\}$, $\alpha = 0.75$, $N = 2$.

Illustrations of the accuracy of the BCF are provided in Fig. 2 and Fig. 3 on sequences of the MPI Sintel benchmark with large occluded regions. Besides, we make a specific focus on the improvements obtained with the candidates extensions. The difference between BCF without candidates extension and the full BCF is clearly visible for occluded pixels and testifies the importance of the exemplar-based and camera motion candidates extensions. Overall, the full BCF is very close to the ground-truth motion field revealing the performance of the local parametric motion computation in the first

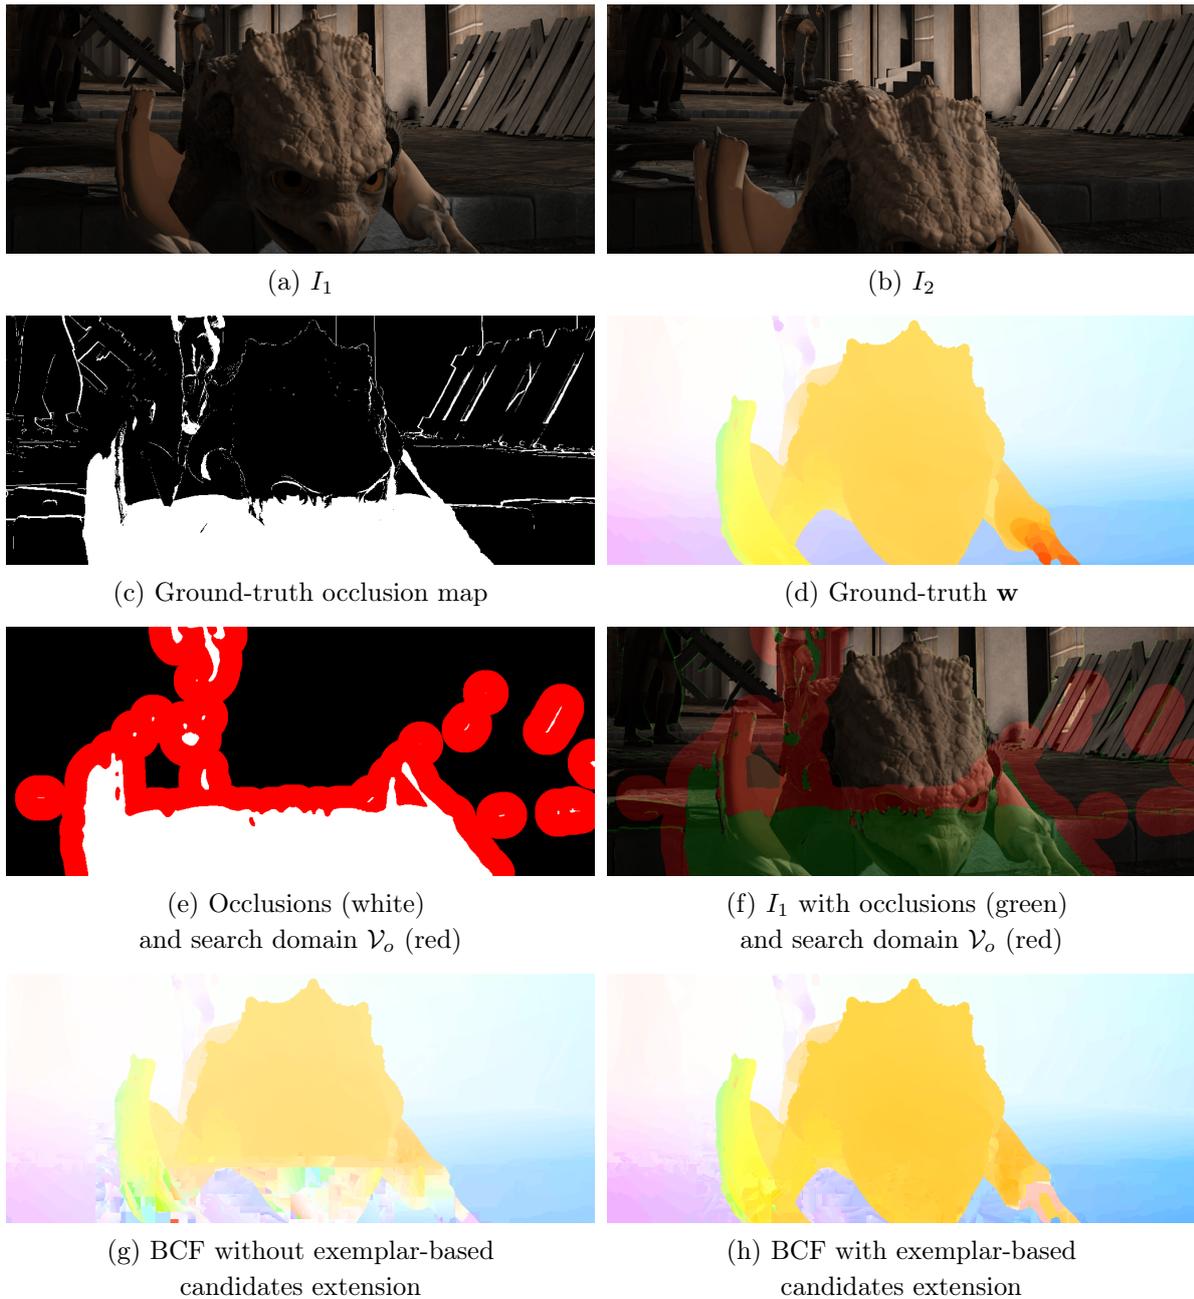


Figure 2: Illustration of the performance improvement with exemplar-based candidates extension. First row: two successive input images. Second row: ground-truth occlusion map and motion field. Third row: representation of the search domain \mathcal{V}_o (displayed here after median filtering of the occlusion map for the sake of visibility only). Fourth row: Best Candidate Flow obtained respectively without and with the exemplar-based candidates extension.

step of *AggregFlow*. Indeed, we report in Table 1 the objective evaluation given by the Endpoint Error (EPE) scores for the full BCF and BCF without candidates extensions, on the sequences provided with ground-truth in the datasets MPI SINTEL and MIDDLEBURY. We also compare them

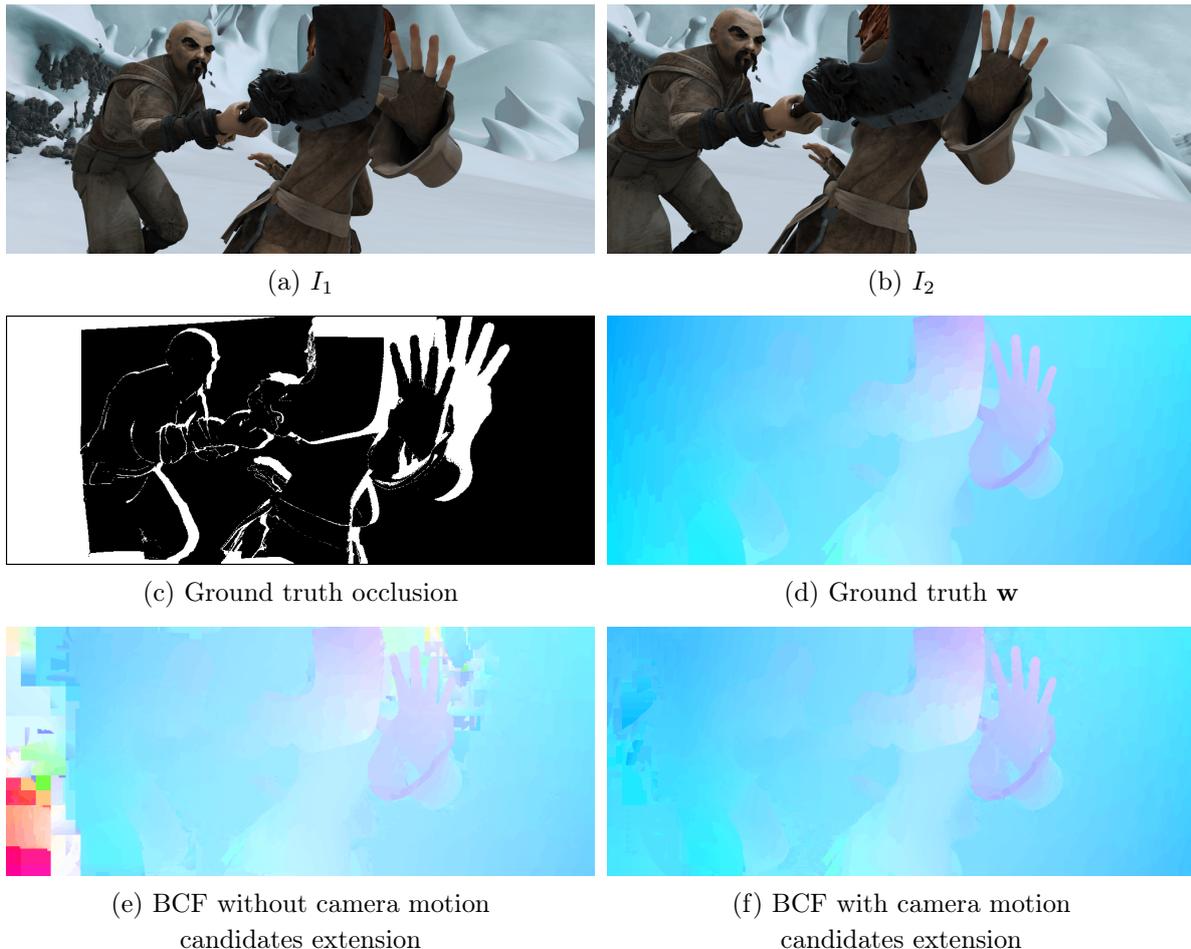


Figure 3: Performance improvement with camera motion candidates extension. First row: two successive input images. Second row: ground-truth occlusion map and motion field. Third row: Best Candidate Flow obtained respectively without and with the camera motion candidates extension.

with those of motion fields supplied by [59, 64], as obtained with publicly available code. Both BCFs outperform state-of-the-art methods [59, 64] in the two benchmarks. Accuracy is further significantly improved with full BCF, especially for the MPI SINTTEL sequences where large displacements and wide occluded regions are present. It demonstrates that the combination of local affine estimations in square patches with patch correspondences as described in Section 3.1, is quite relevant and sufficient to recover very accurate motion fields. The challenge now is to select the best velocity vector among the motion candidates at every pixel.

3.4 Occlusion confidence map

In Section 3.2, the occlusion map o was assumed to be known, and we addressed the motion-based occlusion filling problem by recovering motion candidates for occluded pixels from non-occluded areas. Occlusion detection, that is the determination of o , will be performed through the two steps of AggregFlow. In the first step, we compute a coarse occlusion confidence map, which will be used in the aggregation to guide the estimation. Our procedure is simple and exploits the patch distribution

Table 1: EPE-all scores of motion fields on sequences with ground-truth from MPI MPI SINTTEL and MIDDLEBURY datasets.

	MPI SINTTEL	MIDDLEBURY
Full BCF	0.792	0.0710
BCF w/o candidates extension	1.851	0.0833
DeepFlow [59]	4.691	0.386
MDP-Flow2 [64]	4.006	0.223

$\mathcal{P}_{\mathcal{S},\alpha}$ and the correspondences used for motion candidates estimation. Nevertheless, from a more general point of view, the coarse occlusion confidence map could be designed differently, e.g., in the framework of [35].

We first perform a coarse occlusion detection at the patch level. We consider the smallest patch size s_1 of the set \mathcal{S} defined in Section 3.1 and detect the occluded patches of the set $\mathcal{P}_{s_1,\alpha}$. A common and simple occlusion detection consists in checking the consistency of forward and backward estimated motion vectors [32, 33, 46]. We apply the same principle to patches of $\mathcal{P}_{s_1,\alpha}$. Simplifying the notations of Section 3.1 for the sake of readability, let us denote T_P^f the forward displacement between a patch $P \subset I_1$ and its matched patch $M_P \subset I_2$, and T_P^b the backward displacement between M_P and its matched patch in I_1 . The forward-backward consistency criterion states that the patch P is occluded if $\|T_P^f + T_P^b\| > \nu$, where ν is a threshold. We then infer a patch-based occlusion map o_P as follows:

$$o_P(x) = \begin{cases} 1 & \text{if } \exists P \in \mathcal{P}_{s_1,\alpha}(x) \text{ such that } P \text{ is occluded} \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Let us now consider the point set \mathcal{X}_{o_P} composed of the centers of each occluded patch: $\mathcal{X}_{o_P} = \{x \in \Omega : o_P(x) = 1, x \text{ is the center pixel of } P\}$. We use the density of the point set as an indicator of the presence of occlusions. We apply a Parzen density estimation on $\mathcal{X}_{o_P} = \{x_1, \dots, x_{N_P}\}$, with N_P the number of occluded patches:

$$\omega_o(x) = \frac{1}{N_P} \sum_{i=1}^{N_P} \frac{1}{\sigma} K\left(\frac{x - x_i}{\sigma}\right), \quad (9)$$

where σ is the bandwidth parameter and we choose K to be a Gaussian kernel. We set $\sigma = s_1$. The occlusion confidence map ω_o is thus built as a probability density of the occlusion state. Figure 4 shows an example of o_P and ω_o . The map ω_o will be exploited in the aggregation stage to guide a sparsity-constrained occlusion detection.

The output of AggregFlow first step are the motion candidates set $\mathcal{C}_f(x)$ and the occlusion confidence map ω_o . They will be exploited in the aggregation stage described in the next section to generate the final motion and occlusion fields.

4 Discrete aggregation

The analysis of the Best Candidate Flow in subsection 3.3 has shown that the set of candidates at each pixel contains at least one motion vector very close to the ground truth. Therefore, we conceive the aggregation as the selection of the best candidate at every pixel. To this end, we formulate the

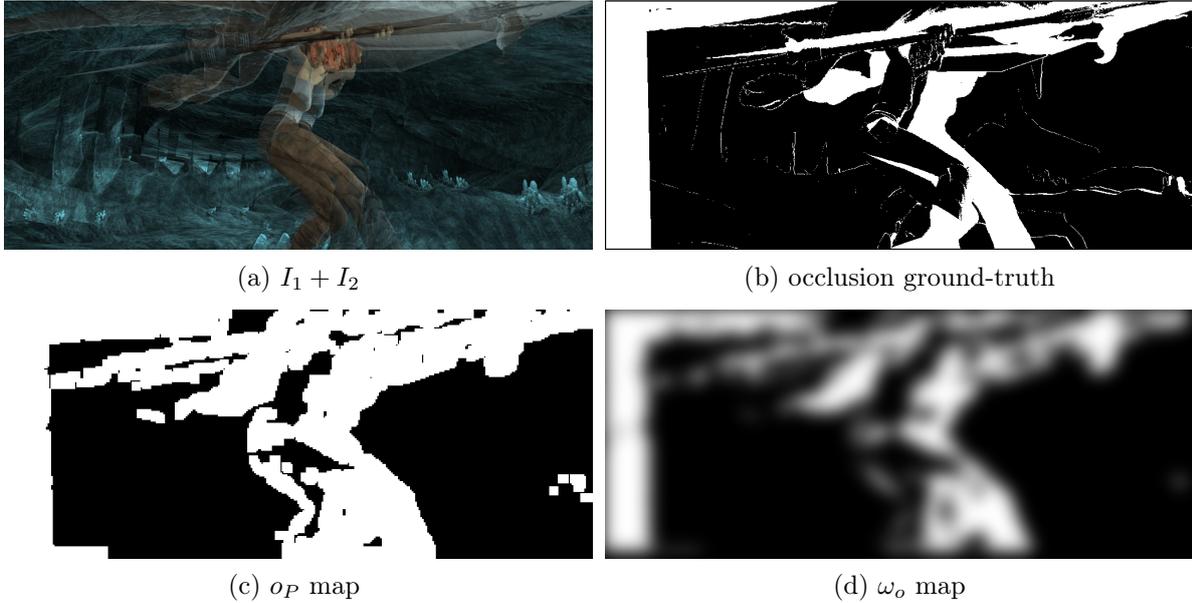


Figure 4: Patch-based occlusion detection. First row: Overlap of the two successive input images and occlusion ground-truth. Second row: Corresponding computed patch-based occlusion map o_P and occlusion confidence map ω_o .

aggregation as a discrete optimization problem, where the discrete finite motion vector space at each pixel x is composed of the motion candidates $\mathcal{C}_f(x)$. The occlusion map will be estimated jointly with the motion field while exploiting the occlusion confidence map ω_o . The aggregation step amounts to minimizing the global energy function $E(\mathbf{w}, o)$:

$$\{\hat{\mathbf{w}}, \hat{o}\} = \arg \min_{\{\mathbf{w}, o\}} E(\mathbf{w}, o) \quad \text{s.t. } \mathbf{w}(x) \in \mathcal{C}_f(x), o(x) \in \{0, 1\}.$$

In the following, we detail the design of $E(\mathbf{w}, o)$ and the optimization strategy we have adopted.

4.1 Global energy

The aggregation energy is composed of four terms:

$$E(\mathbf{w}, o) = E_{data}(\mathbf{w}, o, I_1, I_2) + E_{occ}(o, \omega_o) + E_{reg}^{\mathbf{w}}(\mathbf{w}) + E_{reg}^o(o). \quad (10)$$

4.1.1 Data term E_{data}

The data term accounts for the relations between motion, occlusion and input images. At non-occluded pixels, i.e., $o(x) = 0$, we rely on the usual constancy assumption of image intensity and of spatial image gradient, and we robustly penalize the deviation from the data constraints. The potential ρ_{vis} associated to non-occluded (or visible) pixels is given by:

$$\rho_{vis}(x, \mathbf{w}) = \phi(I_2(x + \mathbf{w}(x)) - I_1(x)) + \gamma\phi(\nabla I_2(x + \mathbf{w}(x)) - \nabla I_1(x)), \quad (11)$$

where ϕ is the L_1 norm and γ balances intensity and gradient constancy potentials. Resorting to discrete optimization allows us to use the non-linearized brightness constancy equation. Thus,

coarse-to-fine scheme is not required to cope with large displacements, and we avoid drawbacks related to the loss of small objects with large displacements.

At occluded pixels, no correspondence can be established by definition, and consequently none image feature constancy constraint can be exploited. Therefore, coherently with the motion candidate extension of the first step, we define an exemplar-based data term for occluded pixels, encoded in the potential ρ_{occ} :

$$\rho_{occ}(x, \mathbf{w}, m) = \|\mathbf{w}(x) - \mathbf{w}(m(x))\|^2, \quad (12)$$

where $m(x)$ is the visible pixel matched with pixel x as obtained in (5). The motion vector of an occluded pixel is thus expected to be similar to the motion vector of its matched non-occluded pixel. The data term is finally formed by incorporating the selection of either the visible or the occlusion potential using the occlusion map:

$$E_{data}(\mathbf{w}, o, I_1, I_2) = \sum_{x \in \Omega} (1 - o(x)) \rho_{vis}(x, \mathbf{w}) + \lambda_1 o(x) \rho_{occ}(x, \mathbf{w}, m). \quad (13)$$

In contrast to other occlusion filling methods which only cancel the visibility term ρ_{vis} in occluded areas and fill the occlusions with motion vectors by diffusion [3, 50, 64], ρ_{occ} acts as a valid data term at occluded pixels.

Concerning the occlusion recovery (i.e., the optimization w.r.t. o), the data term favors the selection of the occluded label at pixels where the data constancy term is strongly violated. The continuous approach of [3] operates in a similar way. In [3], the data constancy deviation is balanced by an estimated continuous residual intensity field, from which occluded points are retrieved by thresholding. In contrast, our occlusion map is binary by nature, and strongly prevents the influence of irrelevant data-constancy constraints on motion estimation in occluded areas.

4.1.2 Occlusion term E_{occ}

The data term (13) favours the detection of occluded pixels and must be counterbalanced by another term penalizing occlusion occurrence defined by:

$$E_{occ}(o, \omega_o) = \lambda_2 \sum_x \omega_o(x) o(x), \quad (14)$$

where ω_o is the occlusion confidence map computed in the first stage. The penalty of occlusion occurrence can be interpreted as a sparsity constraint on the binary occlusion field o . A sparsity constraint for occlusion detection was also proposed in [3] in a continuous setting, and in [50] for a binary occlusion variable, but without confidence map.

If we set $\forall x \in \Omega, \omega_o(x) = 1$, which would be similar to what is done in [3, 50], the data-driven occlusion detection would boil down to the data term (13) and (14) would be a pure sparse prior constraint. The detection of the occlusion map would be then too tightly coupled with the currently estimated motion field. We would face a chicken-and-egg problem, where o is determined by \mathbf{w} , which also depends on o . The consequence on the alternate optimization scheme would be a rapid trap into a local minimum.

Illustrations are given in Fig. 5. The results of two variational methods without occlusion handling [16, 59] are displayed in Fig. 5 (e,f). In both cases, the motion field in the occluded region, highlighted by the red bounding box, is wrongly estimated because no occlusion detection is performed. If the occlusion map is initialized to $o(x) = 0, \forall x \in \Omega$, the occlusion terms of our energy (10) are canceled in the very first iteration of the alternate optimization, which results in a

similar behaviour to the methods [16, 59]. If $\forall x \in \Omega, \omega_o(x) = 1$, the convergence remains trapped in the initial local minimum, as displayed in Fig. 5 (g,h). The reason is that the occlusion map is determined by the motion field and cannot deviate from the output of the first iteration. The role of the confidence map ω_o is then to act as an additional evidence for occlusion detection, relaxing the coupling between \mathbf{w} and o . The guidance of ω_o enables to deviate from the output of the first iteration and to converge to the result shown in Fig. 5 (i,j).

4.1.3 Regularization terms $E_{reg}^{\mathbf{w}}$ and E_{reg}^o

The term $E_{reg}^{\mathbf{w}}(\mathbf{w})$ enforces piecewise smoothness of the motion field:

$$E_{reg}^{\mathbf{w}}(\mathbf{w}) = \lambda_3 \sum_{\langle x,y \rangle} \beta(x) \phi(\|\mathbf{w}(x) - \mathbf{w}(y)\|^2) \quad (15)$$

where $\langle x, y \rangle$ denotes the two-site clique issued from the 8-neighborhood system. The weights $\beta(x)$ are given by $\beta(x) = \exp(-\|\nabla I_1^0(x)\|^2/\tau^2)$ to modulate the regularization according to the intensity edge strength. To limit the influence of noise and textured regions on the weights, we consider a smoothed version I_1^0 of I_1 obtained with the L_0 smoothing of [63], favouring piecewise constant images and preserving only the abrupt edges.

It is also important to impose smoothness of the occlusion map with the term E_{reg}^o :

$$E_{reg}^o(o) = \lambda_4 \sum_{\langle x,y \rangle} (1 - \delta(o(x) = o(y))), \quad (16)$$

where δ designates the Kronecker function equal to 1 if its argument is true. The term $E_{reg}^o(o)$ completes the exemplar-based occlusion filling described in Section 4.1.1 with diffusion-based occlusion filling.

4.2 Optimization

The optimization problem (10) is solved by alternating minimization w.r.t. \mathbf{w} and o . The initial value of o is given by the coarse patch-based occlusion detection o_P defined in (8). The matching variable m attached to the exemplar-based candidates extension is initialized with m_o defined in (5) and is recomputed after each update of the occlusion map. Convergence was empirically observed after three iterations in most cases. To avoid unnecessary computational cost, we fix the number of iterations to 3 for all sequences. Table 5 gives an overview of AggregFlow. Hereafter, we give details on the minimization procedure concerning \mathbf{w} and o .

Once $\hat{\mathbf{w}}$ is fixed, the energy to optimize w.r.t. o amounts to:

$$\begin{aligned} \min_o \sum_{x \in \Omega} (1 - o(x)) \rho_{vis}(x, \hat{\mathbf{w}}) + \lambda_1 \sum_{x \in \Omega} o(x) \rho_{occ}(x, \hat{\mathbf{w}}, m) \\ + \lambda_2 \sum_x \omega_o(x) o(x) + \lambda_4 \sum_{\langle x,y \rangle} (1 - \delta(o(x) = o(y))). \end{aligned} \quad (17)$$

Since the pairwise term is submodular, the problem (17) can be solved exactly with standard graph cut method [13].

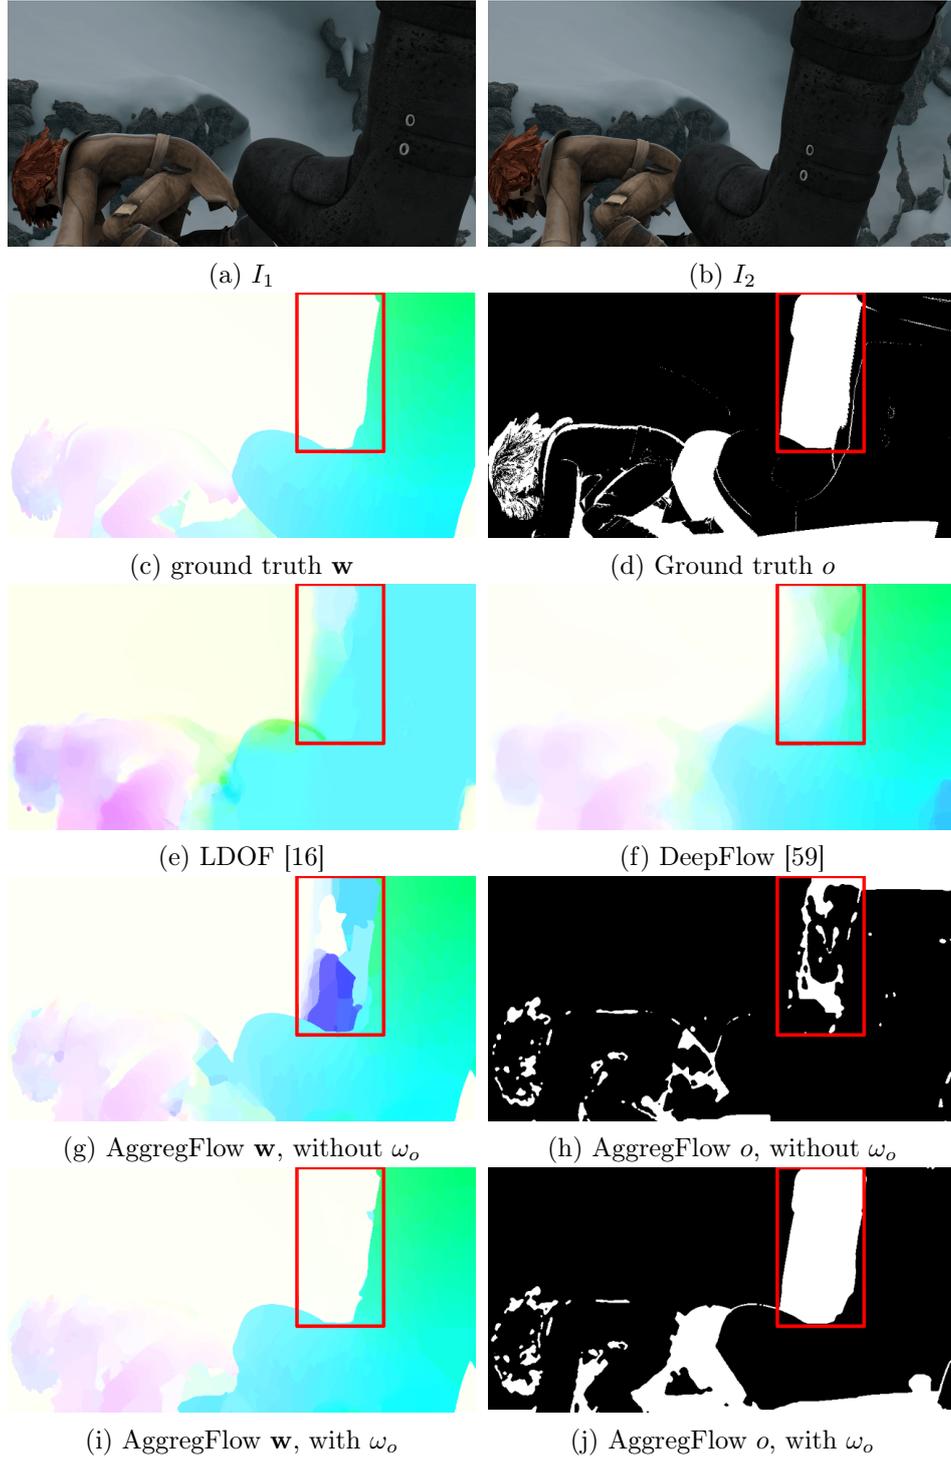


Figure 5: Influence of the occlusion confidence map ω_o on motion and occlusion estimation. (e),(f): results of variational methods [16, 59] without occlusion handling. (g),(h): similar behaviour of our method without occlusion confidence map and impact on the occlusion detection. (i),(j): output of AggregFlow when integrating the occlusion confidence map.

The optimization w.r.t. \mathbf{w} with $\hat{\delta}$ fixed is more difficult. The reduced energy function writes:

$$\begin{aligned} \min_{\mathbf{w}} \sum_{x \in \Omega} (1 - \hat{\delta}(x)) \rho_{vis}(x, \mathbf{w}) + \lambda_1 \hat{\delta}(x) \rho_{occ}(x, \mathbf{w}, m) \\ + \lambda_3 \sum_{\langle x, y \rangle} \beta(x) \phi(\|\mathbf{w}(x) - \mathbf{w}(y)\|^2). \end{aligned} \quad (18)$$

The motion label space has the specificity to be huge (the size of the candidates set $\mathcal{C}_f(x)$ can exceeds 200), and to be spatially varying (each set of motion candidates is specific to each pixel). Message passing methods like belief propagation [24] and TRW-S [38] can be applied to spatially varying label sets, as investigated in [57] for stereo, but we found these methods to be too slow for our minimization problem (18). An alternative is to resort to graph-cut move-making methods [13], generalized in [40] to spatially varying label sets. In this setting, each *move* is a binary optimization problem defined on an auxiliary variable selecting between two global proposals. Due to the spatial variability of the proposals and their independence, the submodularity of the regularization potential of (18) cannot be ensured, and only suboptimal moves can be achieved using QPBO [51].

Our aggregation problem differs from the one of [40] since our motion candidates are locally determined. In contrast, [40] exploits global flow fields that can be directly used as proposals in the move-making process. Thus, we have to build global flow field proposals at each iteration from the local motion candidates computed in patches. The important point is to ensure to some degree spatial smoothness of the proposal. To this end, at each iteration, we choose a size s_i and we take a subset of $\mathcal{P}_{s_i, \alpha}$ formed by non-overlapping patches. Then, we retain for every pixel x the motion candidate from $\mathcal{C}_f(x)$ corresponding to the patch where pixel x lies. We build as many global proposals as necessary to explore the motion candidate space.

Another issue arises from the non-local interaction involved in the exemplar-based term $\rho_{occ}(x, \mathbf{w}, m)$. To make the optimization problem tractable, we transform $\rho_{occ}(x, \mathbf{w}, m)$ to a pixel-wise term at each *move-making* iteration by fixing the exemplar-based constraint $\mathbf{w}(m(x))$ to its value at the previous iteration. At a given *move-making* iteration i , denoting $\hat{\mathbf{w}}^{(i-1)}$ the value of \mathbf{w} at iteration $i - 1$, the potential becomes:

$$\rho_{occ}(x, \mathbf{w}, m) = \left\| \mathbf{w}(x) - \hat{\mathbf{w}}^{(i-1)}(m(x)) \right\|^2. \quad (19)$$

In the next section, we analyse the performance of AggregFlow with experiments on challenging image sequences.

5 Experimental results

5.1 Implementation details

All the patch correspondences involved in AggregFlow are computed with the PatchMatch algorithm [6] based on the minimal C++ code provided by the authors³. A weighted median filtering with bilateral weights [65] is performed as a post-processing step to enhance motion edges as advocated in [55]. For the discrete minimization, we use available QPBO and max-flow code⁴. After extensive experimental tests, the aggregation parameters have been set to $\lambda_1 = 5$, $\lambda_2 = 50$, $\lambda_3 = 500$, $\lambda_4 = 20$ for for all the image sequences of the MPI Sintel benchmark and to $\lambda_1 = 2$, $\lambda_2 = 10$, $\lambda_3 = 250$, $\lambda_4 = 4.5$ for all the image sequences of the Middlebury dataset. As a representative

³http://gfx.cs.princeton.edu/pubs/Barnes_2009_PAR/index.php

⁴<http://pub.ist.ac.at/vnk/software.html>

t

Table 2: Overview of AggregFlow.

<p>1. Local step</p> <p>1.1. Generate the motion candidates sets $\mathcal{C}(x)$ (3)</p> <p>1.2. Compute patch-based occlusion map o_P (8) Derive the occlusion confidence map ω_o from o_P (9)</p> <p>1.3. Compute the matching variables $m_o(x)$ (5) Extend motion candidates in occluded regions to obtain \mathcal{C}_f (7)</p> <p>Output of the 1st step: \mathcal{C}_f, ω_o</p> <p>2. Global aggregation Initialize $o = o_P$ and $m = m_o$ Iterate:</p> <p>2.1. Estimate \mathbf{w} (18)</p> <p>2.2. Estimate o (17)</p> <p>2.3. Update m (5)</p> <p>Output of the 2nd step: \mathbf{w}, o</p> <p>3. Post-processing : weighted median filtering on \mathbf{w}</p>
--

example (the one used to compare methods), the computation time for the *Urban2* sequence of the Middlebury benchmark is 27 minutes on a Intel Xeon laptop with 2.20GHz clock speed and 64Gb RAM. Nevertheless, the first step of AggregFlow can be massively parallelized, which should lead to a far less computation cost with a GPU implementation for instance. Most of the computation time is consumed in the patch correspondence sub-step for the largest patch size (106×106 pixels). The determination of the matching variable m is performed with patches of size 11×11 .

5.2 Quantitative results on computer vision benchmarks

We have evaluated AggregFlow on the two most representative benchmarks for optical flow: MPI Sintel flow dataset⁵ [19] and Middlebury flow dataset⁶ [4], which offer different and complementary challenges. We have retained the Endpoint Error measure (EPE) for quantitative evaluation. Results of [64] and [59] reported in Table 6, Fig.6 and Fig.7 have been obtained with the public codes provided by the authors^{7,8}.

MPI Sintel flow dataset Sequences of the most recent MPI Sintel benchmark [19] are characterized by long-range motion, motion blur, non-rigid motion, and wide occluded areas. Methods are evaluated on two versions of the sequences named *Clean* and *Final*. The Final version adds motion and defocus blur along with atmospheric effects like fog on some sequences. We reproduce in Tables 3 and 4 public results of the top ranked published methods, which are presently (at paper submission date, June 4, 2014) available on the MPI Sintel website. Results are analyzed through several indicators: “EPE all” is the average EPE on all the sequences; “EPE matched” and “EPE unmatched” restrict the error measure respectively to regions that remain visible in adjacent frames (non-occluded pixels) and to regions that are visible only in one of two adjacent frames (occluded pixels); “d0-10” denotes EPE over regions closer than 10 pixels to the nearest occlusion boundary, and thus reveals the ability to recover motion discontinuities; “s40+” denotes EPE over regions with velocities larger than 40 pixels per frame. Methods are ranked regarding their EPE all. Visual comparison with results supplied by [59] and [64] on training sequences (i.e., MPI Sintel sequences provided with ground truth) is available in Fig.6.

As for the *Clean* subset, our method AggregFlow ranks first over the published methods. The most significant improvement is obtained on the unmatched category, which emphasizes the efficiency of our occlusion framework. AggregFlow is ranked second for the d0-10 metric, which demonstrates its capacity to recover motion discontinuities as confirmed by results displayed in Fig.6. First, it is due to the robust affine estimation of the motion candidates able to capture locally dominant motion in case of two or even several motions present inside patches, which preserves motion discontinuities. It is also made successful by the efficient occlusion module, which allows us to moderate the need for motion field regularization. Indeed, missing information in occluded regions is usually tackled by imposing high regularization with the result of oversmoothing the rest of the motion field (see motion fields computed with DeepFlow [59] in Fig.6). In case of very large displacements (s40+ metric), all the first five methods (AggregFlow, [59, 64, 5, 41]) somehow integrate feature matching in their motion estimation process to capture the largest deformations. The top rank of AggregFlow demonstrates the efficiency of the aggregation framework for integrating feature matching.

⁵<http://sintel.is.tue.mpg.de/>

⁶<http://vision.middlebury.edu/flow/>

⁷<http://www.cse.cuhk.edu.hk/leojia/projects/flow/>

⁸<http://lear.inrialpes.fr/src/deepmatching/>

Table 3: Results on the MPI Sintel Clean test subset.

	EPE all	EPE matched	EPE unmatched	d0-10	s40+
AggregFlow	4.754	1.694	29.685	3.705	31.184
DeepFlow [59]	5.377	1.771	34.751	4.519	33.701
MDP-Flow2 [64]	5.837	1.869	38.158	3.210	39.459
EPPM [5]	6.494	2.675	37.632	4.997	39.152
S2D-Matching [41]	6.510	2.792	36.785	5.523	44.187
Classic+NLP [55]	6.731	2.949	37.545	5.573	45.290
FC-2Layers-FF [54]	6.781	3.053	37.144	5.841	45.962
MLDP-OF [45]	7.297	3.260	40.183	5.581	51.146

Table 4: Results on the MPI Sintel Final test subset.

	EPE all	EPE matched	EPE unmatched	d0-10	s40+
DeepFlow [59]	7.212	3.336	38.781	5.650	44.118
AggregFlow	7.329	3.696	36.929	5.538	44.858
S2D-Matching [41]	7.872	3.918	40.093	5.975	48.782
FC-2Layers [54]	8.137	4.261	39.723	6.537	51.349
MLDP-OF [45]	8.287	4.165	41.905	6.345	50.540
Classic+NLP [55]	8.291	4.287	40.925	6.520	51.162
EPPM [5]	8.377	4.286	41.695	6.556	49.083
MDP-Flow2 [64]	8.445	4.150	43.430	5.703	50.507

As for the *Final* version AggregFlow is ranked second in terms of EPE all. The slight decreasing in performance compared to the Clean subset is mainly due to errors caused by the added fog effect in the two *ambush* sequences. As emphasized in [5], local intensity-based displacement computation tends to capture the motion of the fog rather than the movement of objects appearing in transparency. As our candidates estimation is local, it is subject to this limitation. Global variational approaches are able to diffuse motion estimates in these regions and are consequently better suited for this kind of situations. Despite this shortcoming, our method still yields significant improvement in unmatched regions and on motion discontinuities. One solution to improve results in fog regions would be to incorporate a more sophisticated feature correspondence technique as the ones proposed in [41, 59].

Middlebury dataset The Middlebury benchmark is composed of sequences with small displacements, where the main challenge is to be able to recover both complex smooth deformation, motion discontinuities and motion details. Table 5 reproduces public results presently available (at paper submission date, June 4, 2014) for the same methods as those taken for comparison on the MPI Sintel benchmark. Visual comparative results are displayed in Fig.7. It can be observed that the EPE values, together with the differences between methods, are much lower than for the MPI Sintel dataset. The average EPE score computed over the considered methods is equal to 6.22 for the MPI Sintel Clean subset and to 0.327 for the Middlebury dataset, with respective variance of 0.613 and 0.0025. We also provide the average rank over the 8 test sequences for each method which is

Table 5: Results on the Middlebury benchmark for the same set of methods.

	EPE all	Avg. rank
MDP-Flow2 [64]	0.245	7.8
FC-2Layers-FF [54]	0.283	19.3
Classic+NL [55]	0.319	27.1
EPPM [5]	0.329	32.6
AggregFlow	0.339	35.9
MLDP-OF [45]	0.349	32.6
S2D-Matching [41]	0.347	34.6
DeepFlow [59]	0.416	48.8

Table 6: Results on the MPI Sintel training subset. Scores correspond to the EPE all metric

	AggregFlow	AggregFlow w/o occlusion	DeepFlow [59]	MDP-Flow2 [64]
<i>ambush_2</i>	5.632	9.456	14.743	12.083
<i>ambush_4</i>	11.923	16.515	14.647	15.570
<i>ambush_5</i>	5.042	5.500	8.333	6.591
<i>ambush_6</i>	5.854	6.251	9.928	8.466
<i>market_5</i>	9.957	11.958	15.056	12.816
<i>market_6</i>	3.626	4.547	6.606	5.384
<i>cave_2</i>	6.029	8.228	10.082	8.347
<i>cave_4</i>	3.706	4.185	4.234	3.815
<i>temple_3</i>	5.875	8.314	11.895	9.011
Average	6.002	8.417	10.614	9.120

the metric used for global ranking on the Middlebury website.

On the whole Middlebury benchmark, AggregFlow is ranked 38 over 97 tested methods in terms of average rank on the results (evaluated with the average endpoint error on the sequence) obtained for the eight test sequences. Notwithstanding, it is still very close to the ranked two MDP-Flow2 method [64] in terms of EPE metric, knowing that the top ranked published method OFLAF [36] has an average rank of 6.8 and an EPE all of 0.197 (OFLAF method was not tested on the MPI Sintel benchmark). Visual results reported in Fig.7 confirm the tightness of performance. In particular, the preservation of motion discontinuities with AggregFlow is more satisfying than with the DeepFlow method [59]. These results also show that AggregFlow is competitive for recovering motion details in addition to the large velocities of the MPI Sintel benchmark.

5.3 Occlusion handling

As aforementioned the impact of our occlusion framework on optical flow estimation is demonstrated by the EPE unmatched metric scores obtained on the MPI Sintel benchmark (Tables 3 and 4). Results of Fig.6 reveal the superiority of AggregFlow in coping with occluded regions. Since the occlusion framework is composed of several elements, we detail the influence of each one in the following. The efficiency of the motion candidates extension in occluded regions has already been highlighted in Section 3.3 and Table 1 through the analysis of the Best Candidate Flow.

To evaluate the occlusion model of the aggregation step, we report in Table 6 results obtained on a selection of training sequences of the MPI Sintel benchmark with the largest displacements. We distinguish between the full AggregFlow method, and AggregFlow without the occlusion-related terms in (10), that is, by setting $\lambda_1 = 0$, $\lambda_2 = 0$ and $\lambda_3 = 0$. The improvement due to the occlusion terms is clearly significant since the average EPE is 8.417 for AggregFlow without occlusions and 6.002 for full Aggregflow. It can also be noticed that even without handling occlusion AggregFlow still performs better than competing methods. The role of the occlusion confidence map involved in the sparsity constraint (14) has already been shown in Section 4.1.2 and Fig.5.

Recovered occlusion maps are displayed in Fig.6 and Fig.7. For the large occluded regions of Figure 6 for which ground truth is available, the estimated occlusion map is correct in most cases. A specific behaviour is particularly prominent in the *market_5* example, where occlusions are overdetected. It is due to the modeling assumption stating that occluded regions correspond to large violations of the data constancy equation. Regions of illumination changes may thus be detected as occlusions. While it leads strictly speaking to wrong occlusion detection, it can still be beneficial to motion estimation by implicitly treating illumination changes.

6 Conclusion

We have presented a new two-step optical flow estimation method called AggregFlow. It articulates the computation of local motion candidates and their global aggregation while jointly recovering occlusion maps. The framework is generic, and both the local and global steps could be adapted for specific purposes. We demonstrated the added value of combining patch correspondences and patch-based affine motion estimation to produce highly accurate motion candidates, advocating the relevance of patch-based parametric motion estimation, provided size and position of the patches are appropriately defined. The integration of multiple patch correspondences in the candidates generation process allows us to deal with local matching ambiguities. We formulated the aggregation step as a discrete optimization problem, selecting the best motion candidate at every pixel while preserving motion discontinuities and achieving occlusion recovery. The occlusion scheme acts in both steps of AggregFlow. An exemplar-based occlusion term is incorporated in the global aggregation energy. Incidentally, it could be integrated in other estimation paradigms as well, e.g., in variational approaches. Occlusion cues derived from the computed motion candidates are exploited in the sparse modeling of occlusions. Overall, AggregFlow achieves state-of-the-art results on the MPI Sintel benchmark. The most significant improvements are reached in occluded regions and for large displacements.

Extensions of the method could tackle remaining matching errors in the patch correspondence and in the exemplar search substeps. A more elaborate and discriminative distance than the pixel-based L_1 distance could be envisioned for patch matching. Future work could also deal with a GPU implementation to largely improve computation efficiency.

References

- [1] A. Alba, E. Arce-Santana, and M. Rivera. Optical flow estimation with prior models obtained from phase correlation. In *ISVC*, Las Vegas, November 2010.
- [2] P. Arias, G. Facciolo, V. Caselles, and G. Sapiro. A variational framework for exemplar-based image inpainting. *Int. J. Comp. Vis.*, 93(3):319–347, 2011.

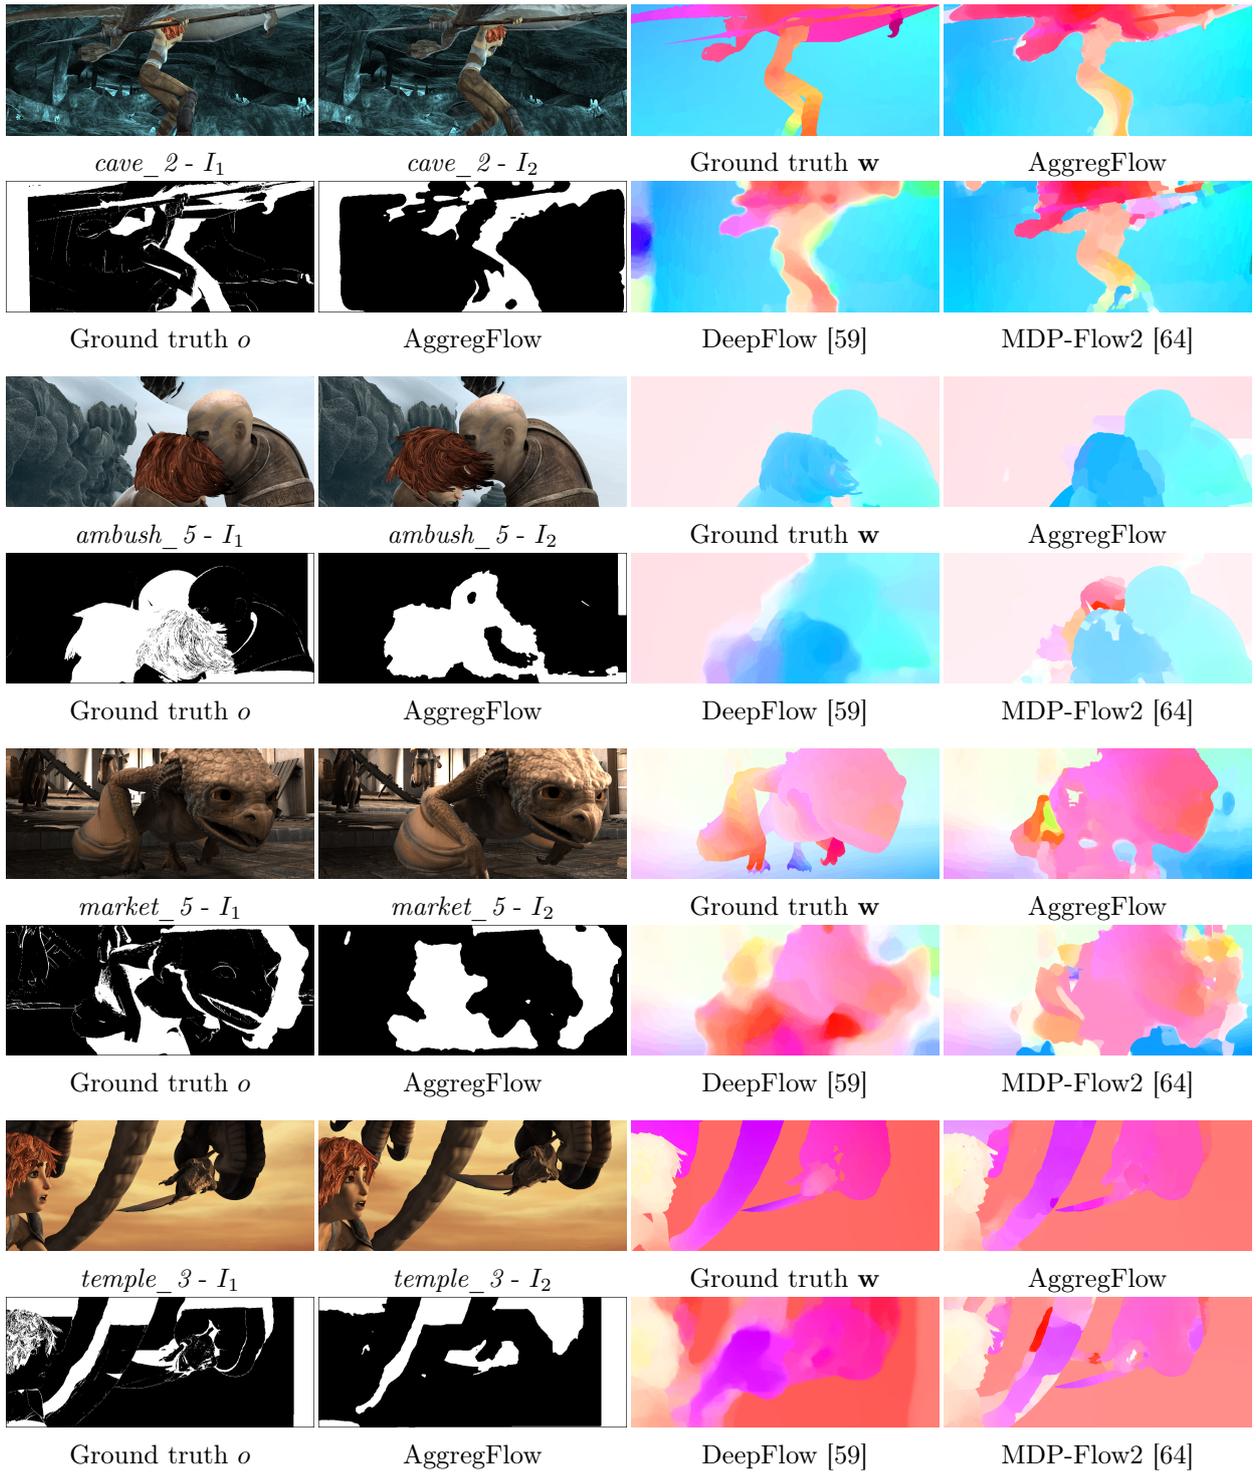


Figure 6: Comparative evaluation with [59] and [64] on several sequences of the MPI Sintel dataset. Every first row from left to right: successive input images, ground truth motion field, motion field computed with AggregFlow. Every second row from left to right: ground truth occlusion, occlusion map computed with AggregFlow, motion fields computed with DeepFlow [59] and MDP-Flow2 [64].

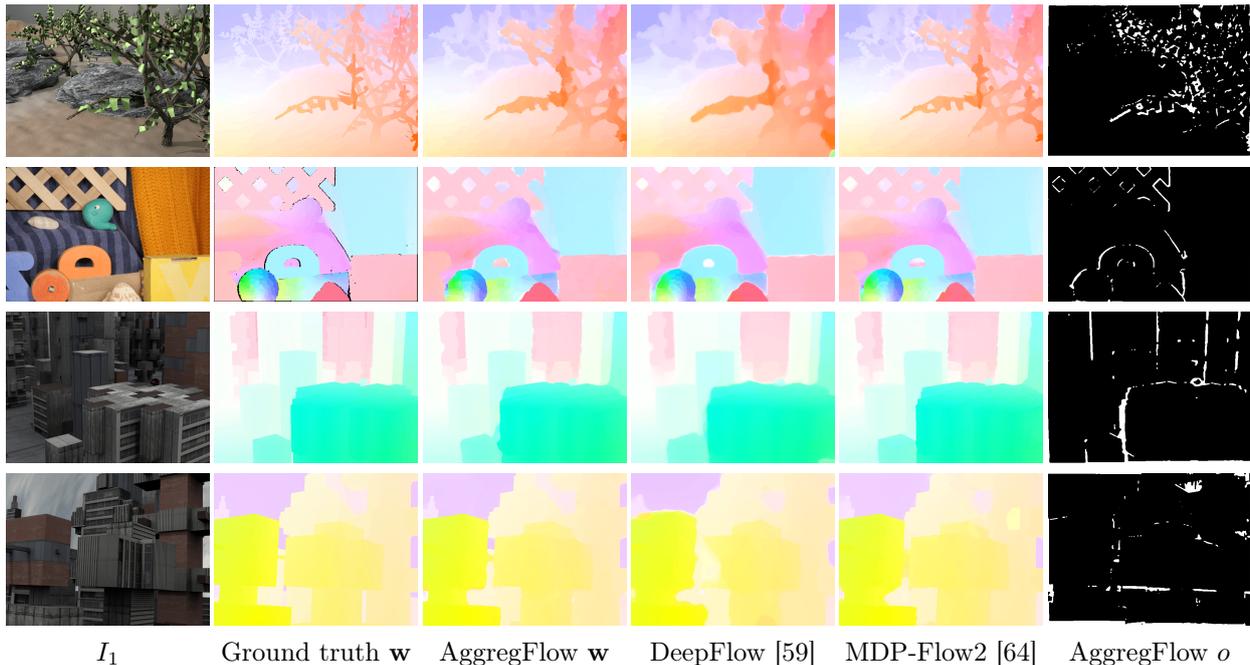


Figure 7: Comparative evaluation with [59] and [64] on several sequences of the Middlebury dataset. From top to bottom: sequences *grove3*, *rubberwhale*, *urban2*, *urban3*. In each row from left to right: first input image; ground truth motion field; motion field computed resp. with AggregFlow, DeepFlow [59] and MDP-Flow2 [64]; occlusion map computed with AggregFlow.

- [3] A. Ayvaci, M. Raptis, and S. Soatto. Sparse occlusion detection with optical flow. *Int. J. Comp. Vis.*, 97(3):322–338, 2012.
- [4] S. Baker, D. Scharstein, JP Lewis, S. Roth, M.J. Black, and R. Szeliski. A database and evaluation methodology for optical flow. *Int. J. Comp. Vis.*, 92(1):1–31, 2011.
- [5] L. Bao, Q. Yang and H. Jin. Fast edge-preserving PatchMatch for large displacement optical flow. In *CVPR*, Columbus, June 2014.
- [6] C. Barnes, E. Shechtman, A. Finkelstein, and D.B. Goldman. PatchMatch: A randomized correspondence algorithm for structural image. In *SIGGRAPH*, New Orleans, August 2009.
- [7] J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *Int. J. Comp. Vis.*, 12(1):43–77, 1994.
- [8] M. Bertalmio, G. Sapiro, V. Caselles and C. Ballester. Image inpainting. In *SIGGRAPH*, New Orleans, July 2000.
- [9] M.J. Black and P. Anandan. A framework for the robust estimation of optical flow. In *ICCV*, Berlin, May 1993.
- [10] M.J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Comp. Vis. Image Understanding*, 63(1):75–104, 1996.
- [11] M.J. Black and A.D. Jepson. Estimating optical flow in segmented images using variable-order parametric models with local deformations. *IEEE Trans. Patt. Anal. Mach. Intell.*, 18(10):972–986, 1996.

- [12] M. Bleyer, C. Rhemann and M. Gelautz. Segmentation-based motion with occlusions using graph-cut optimization. In *DAGM*, Berlin, Sept. 2006.
- [13] Y. Boykov, O. Veksler, R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Trans. Patt. Anal. Mach. Intell.*, 23(11):1222–1239, 2001.
- [14] J. Braux-Zin, R. Dupont and I. Bartoli. A general dense image matching framework combining direct and feature-based costs. In *ICCV*, Sidney, Dec. 2013.
- [15] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, Prague, May 2004.
- [16] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE Trans. Patt. Anal. Mach. Intell.*, 33(3):500–513, 2011.
- [17] A. Bugeau, M. Bertalmío, V. Caselles and G. Sapiro. A comprehensive framework for image inpainting. *IEEE Trans. Image Process.*, 19(10):2634–2645, 2010.
- [18] A. Bugeau, V.T. Ta and N. Papadakis. Variational exemplar-based image colorization. *IEEE Trans. Image Process.*, 23(1):298–307, 2014.
- [19] D.J. Butler, J. Wulff, G.B. Stanley and M.J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, Florence, Oct. 2012.
- [20] T. Chan, S.H. Kang, and J.H. Shen. Eulers elastica and curvature based inpaintings. *SIAM J. Applied Math.*, 63(2), 564–592, 2002.
- [21] L. Chen, H. Jin, Z. Lin, S. Coben, Y. Wu. Large displacement optical flow from nearest neighbor fields. In *CVPR*, Portland, June 2013.
- [22] D. Cremers and S. Soatto. Motion competition: A variational approach to piecewise parametric motion segmentation. *Int. J. Comp. Vis.*, 62(3):246–265, 2005.
- [23] A. Criminisi, P. Pérez and K. Tomaya. Region filling and object removal by exemplar-based image inpainting. In *IEEE Trans. Image Process.*, 13(9):1200–1212, 2004.
- [24] P.F. Felzenszwalb and D.P. Huttenlocher. Efficient belief propagation for early vision. *Int. J. Comp. Vis.*, 70(1):41–54, 2006.
- [25] D. Fortun and C. Kervrann. Semi-local variational optical flow estimation. In *ICIP*, Orlando, Sept. 2012.
- [26] D. Fortun, P. Bouthemy, P. Paul-Gilloteaux, and C. Kervrann. Aggregation of patch-based estimations for illumination-invariant optical flow in live cell imaging. In *ISBI*, San-Francisco, April 2013.
- [27] M. Gelgon and P. Bouthemy. A region-level motion-based graph representation and labeling for tracking a spatial image partition. *Patt. Recognition*, 33(4):725–740, 2000.
- [28] D. Hafner, O. Demetz, and J. Weickert. Why Is the Census Transform Good for Robust Optic Flow Computation?. In *SSVM*, March 2012.
- [29] P. Héas and E. Mémin. Three-dimensional motion estimation of atmospheric layers from image sequences. *IEEE Trans. Geosc. Remote Sensing*, 46(8):2385–2396, 2008.
- [30] F. Heitz and P. Bouthemy. Multimodal estimation of discontinuous optical flow using Markov random fields. *IEEE Trans. Patt. Anal. Mach. Intell.*, 15(12):1217–1232, 1993.
- [31] B.K.P. Horn and B.G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

- [32] A. Humayun, O. Mac Aodha, and G.J. Brostow. Learning to find occlusion regions. In *CVPR*, Colorado Springs, June 2011.
- [33] S. Ince and J. Konrad. Occlusion-aware optical flow estimation. *IEEE Trans. Image Process.*, 17(8):1443–1451, 2008.
- [34] P.M. Jodoin and M. Mignotte. Optical-flow based on an edge-avoidance procedure. *Comp. Vis Image Understanding*, 113(4):511–531, 2009.
- [35] C. Kervrann, J. Boulanger, T. Pécot, P. Perez, and J. Salamero. Multiscale neighborhood-wise decision fusion for redundancy detection in image pairs. *SIAM J. Multiscale Modeling & Simulation*, 9(4):1829–1865, 2011.
- [36] T.H. Kim, H. Lee, and K.M. Lee. Optical flow via locally adaptive fusion of complementary data costs. In *ICCV*, Sydney, Dec. 2013.
- [37] V. Kolmogorov and R. Zabih. Computing visual correspondences with occlusion using graph cuts. In *ICCV*, Vancouver, 2001.
- [38] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Trans. Patt. Anal. Mach. Intell.*, 28(10):1568–1583, 2006.
- [39] N. Komodakis and G. Tziritas. Image completion using efficient belief propagation via priority scheduling and dynamic pruning. *IEEE Trans. Image Process.*, 16(11):2649–2661, 2007.
- [40] V. Lempitsky, S. Roth, and C. Rother. Fusionflow: Discrete-continuous optimization for optical flow estimation. In *CVPR*, Anchorage, 2008.
- [41] M. Leordeanu, A. Zanfir, and C. Sminchisescu. Locally affine sparse-to-dense matching for motion and occlusion estimation. In *ICCV*, Sydney, December 2013.
- [42] B.D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *IJCAI*, Vancouver, 1981.
- [43] E. Mémin and P. Pérez. Dense estimation and object-based segmentation of the optical flow with robust techniques. *IEEE Trans. Image Process.*, 7(5):703–719, 1998.
- [44] A. Mitiche and P. Bouthemy. Computation and analysis of image motion: a synopsis of current problems and methods. *Int. J. Comp. Vis.*, 19(1):29–55, July 1996.
- [45] M. Mohamed, H. Rashwan, B. Mertsching, M. Garcia, and D. Puig. Illumination-robust optical flow approach using local directional pattern. In *IEEE Trans. Circ. Sys. Video Technol.*, to appear, 2014.
- [46] M.G. Mozerov. Constrained optical flow estimation as a matching problem. *IEEE Trans. Image Process.*, 22(5):2044–2055, 2013.
- [47] H.H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Trans. Patt. Anal. Mach. Intell.*, 8(5):565–593, 1986.
- [48] J.M. Odobez and P. Bouthemy. Robust multiresolution estimation of parametric motion models. *J. of Vis. Comm. Image Repres.*, 6(4):348–365, 1995.
- [49] J.M. Odobez and P. Bouthemy. Direct incremental model-based image motion segmentation for video analysis. *Signal Process.*, 66(2):143–155, 1998.
- [50] N. Papadakis, R. Yildizoglu, J.F. Aujol and V. Caselles. High-dimension multilabel problems: Convex or nonconvex relaxation? *SIAM J. on Imaging Sciences*, 6(4):2603–2639, 2013.

- [51] C. Rother, V. Kolmogorov, V. Lempitsky and M. Szummer. Optimizing binary MRFs via extended roof duality. In *CVPR*, Minneapolis, June 2007.
- [52] T. Senst, V. Eiselen and T. Sikora. Robust local optical flow for feature tracking. *IEEE Trans. Circ. Sys. Video Technol.*, 22(9):1377–1387, 2012.
- [53] A.N. Stein and M. Hebert. Occlusion boundaries from motion: Low-level detection and mid-level reasoning. *Int. J. Comp. Vis.*, 82:325–357, 2009.
- [54] D. Sun, E. Sudderth and M. Black. Layered segmentation and optical flow estimation over time. In *CVPR*, Portland, June 2012.
- [55] D. Sun, S. Roth, and M. Black. A quantitative analysis of current practices in optical flow estimation and the principles behind them. *IEEE Trans. Patt. Anal. Mach. Intell.*, 106(2):115–137, 2014.
- [56] M. Unger, M. Werlberger, T. Pock and H. Bischof. Joint motion estimation and segmentation of complex scenes with label costs and occlusion modeling. In *CVPR*, Providence, June 2012.
- [57] J. Ulén and C. Olsson. Simultaneous fusion moves for 3D-label stereo. In *EMMCVPR*, Lund, August 2013.
- [58] A. Wedel, T. Pock, C. Zach, H. Bischof and D. Cremers. An improved algorithm for TV-L 1 optical flow. In *Dagstuhl Visual Motion Analysis Workshop*, 2008.
- [59] P. Weinzaepfel, J. Revaud, Z. Harchaoui, C. Schmid. Large displacement optical flow with deep matching. In *ICCV*, Sydney, December 2013.
- [60] M. Werlberger, T. Pock, and H. Bischof. Motion estimation with non-local total variation regularization. In *CVPR*, San Francisco, June 2010.
- [61] J. Xiao, H. Cheng, H. Sawhney, C. Rao, and M. Isardi. Bilateral filtering-based optical flow estimation with occlusion detection. In *ECCV*, Graz, 2006.
- [62] L. Xu, J. Chen and J. Jia. A segmentation-based variational model for accurate optical flow estimation. In *ECCV*, Marseille, October 2008.
- [63] L. Xu, C. Lu, Y. Xu and J. Jia. Image smoothing via L_0 gradient minimization. *ACM Trans. Graphics*, 30(6):174–186, 2011.
- [64] L. Xu, J. Jia, and Y. Matsushita. Motion detail preserving optical flow estimation. *IEEE Trans. Patt. Anal. Mach. Intell.*, 34(9):1744–1757, 2012.
- [65] L. Xu, Z. Dai, and D. Jia. Scale invariant optical flow. in *ECCV*, Firenze, Sept. 2012.
- [66] H. Zimmer, A. Bruhn and J. Weickert. Optic flow in harmony. *Int. J. Comp. Vis.*, 93(3):368–388, 2011.
- [67] C.W. Zitnick, N. Jovic and S.B. Sin. Consistent segmentation for optical flow estimation. In *ICCV*, Beijing, October 2005.