

Occlusion-Aware Optical Flow Estimation

Serdar Ince and Janusz Konrad, *Fellow, IEEE*

Abstract—Optical flow can be reliably estimated between areas visible in two images, but not in occlusion areas. If optical flow is needed in the whole image domain, one approach is to use additional views of the same scene. If such views are unavailable, an often-used alternative is to extrapolate optical flow in occlusion areas. Since the location of such areas is usually unknown prior to optical flow estimation, this is usually performed in three steps. First, occlusion-ignorant optical flow is estimated, then occlusion areas are identified using the estimated (unreliable) optical flow, and, finally, the optical flow is corrected using the computed occlusion areas. This approach, however, does not permit interaction between optical flow and occlusion estimates. In this paper, we permit such interaction by proposing a variational formulation that jointly computes optical flow, implicitly detects occlusions and extrapolates optical flow in occlusion areas. The extrapolation mechanism is based on anisotropic diffusion and uses the underlying image gradient to preserve structure, such as optical flow discontinuities. Our results show significant improvements in the computed optical flow fields over other approaches, both qualitatively and quantitatively.

Index Terms—Anisotropic diffusion, disparity estimation, motion estimation, occlusions, optical flow.

I. INTRODUCTION

OPTICAL flow algorithms [1]–[6] are among the best methods for the estimation of disparity in stereo images and motion in video sequences. However, occlusion areas [Fig. 1(a)], resulting from scene structure and/or object motion, pose significant challenges. In this paper, by the term “occlusion area” we refer to an area in one image [area A in image I_L , Fig. 1(a)] that disappears in the other image (I_R). Note that a disappearing area becomes an appearing area (also known as “uncovered area”) if the direction of arrows is reversed. When considering motion, the direction of arrows is related to time (forward *versus* backward), while in stereo it is related to the order of views (left-to-right *versus* right-to-left). Although the work presented here is generic and applies to both stereo and motion, our specific examples are in stereo, and, thus, we will also refer to optical flow as disparity.

It is important to note that for any pair of images optical flow is undefined in an occlusion area since, by definition,

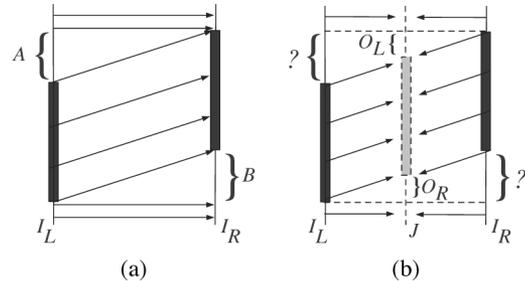


Fig. 1. Illustration of occlusion effects on a horizontal cross-section of two images depicting position change of a simple object (black): (a) area A from I_L is being occluded in I_R by the object, while area B is being uncovered and (b) intermediate image J that can be reconstructed if areas O_L and O_R are identified and proper correspondence with either I_L or I_R is established.

a corresponding area cannot be found in the other image [in Fig. 1(a) points in A have no match in I_R]. However, in certain applications it is desirable to identify optical flow *everywhere* in an image. For example, when rendering a virtual view, depth (disparity) needs to be known for all pixels to be rendered [Fig. 1(b)], while in video frame-rate conversion motion of all pixels is needed. Solutions to this problem have been proposed using additional images [7]–[9]. In one interesting approach, a data-matching term in optical flow formulation is adaptively disabled in order to select the best prediction among several images [9]. We will use a similar disabling strategy in the approach proposed here. From now on, however, we are concerned with the estimation of optical flow from two images only.

Most optical-flow estimation errors arise in occlusion areas due to forced, but unreliable, intensity matching, and are further aggravated by smoothing of optical flow across object boundaries (adjacent to occlusion area). Without explicitly detecting occlusions, methods have been proposed to deal with oversmoothing, such as image-adaptive, isotropic diffusion [10], anisotropic diffusion (image or flow-adaptive) [11], [12], [5], [13], [14], and isotropic diffusion with nonquadratic regularizers [4], [6]. These methods lead to discontinuity-preserving vector fields, but do not account for occlusions. A step towards accounting for occlusions was taken by jointly estimating forward and backward optical-flow fields under flow-field regularization [15], [16]. In one case, the resulting flow diffusion was isotropic but adapted to the difference between forward and backward optical-flow estimates (large difference indicating occlusions) [15], while in the other case the diffusion was anisotropic and based on underlying image gradients [16].

The above methods lead to discontinuity-preserving vector fields, but still produce erroneous results at object boundaries since incorrect intensity matches are allowed despite occlusions. One remedy is to extrapolate (inpaint), rather than estimate, optical flow in occlusion areas [17], [13]. Ideally, one would

Manuscript received August 9, 2007; revised March 25, 2008. First published June 24, 2008; last published July 11, 2008 (projected). This work was supported in part by the National Science Foundation under Grant ECS-0219224 and in part by the National Institutes of Health under Grant 1R21HD050655-01. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Thomas S. Denney.

S. Ince was with Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA. He is now with IntelliVid Corporation, Cambridge, MA 02138 USA (e-mail: ince@alum.bu.edu).

J. Konrad is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA (e-mail: jkonrad@bu.edu).

Digital Object Identifier 10.1109/TIP.2008.925381

first identify occlusion areas and then compute optical flow accounting for these areas. However, it is unclear how to find occlusion areas without first computing optical flow. Hence, in practice, occlusion-unaware optical flow is computed first, resulting in incorrect vectors in occlusion areas (a match is forced despite lack of correspondence). Then, occlusion areas are identified, and, finally, the optical flow is corrected there [13]. This three-step approach is deficient since incorrect vectors from occlusion areas affect their neighbors in visible areas due to spatial regularization typically used. Moreover, this approach does not bootstrap optical flow estimates to improve occlusion detection results.

More recently, methods have been proposed that jointly estimate disparity and occlusions. One such method is based on a visibility constraint but treats the disparity correction in occlusion areas as a postprocessing step [17]. Moreover, this corrective step assumes that depth in occlusion areas and surrounding pixels is constant along epipolar lines. Although often valid, this assumption fails for image backgrounds with varying depth. Another method jointly estimates bidirectionally consistent (forward/backward) motion fields and occlusion labels using Markov random fields in a Bayesian framework [18]. However, the occlusion detection mechanism relies on an intensity, rather than flow-field, mismatch, and no vector correction is performed in occlusion areas. Graph cut methods [19], [20] have been used in disparity/motion estimation under occlusions, as well. These methods explicitly define an occlusion term in the formulation, however the intensity matching and occlusion terms usually do not interact, i.e., disable each other, unlike in the approach proposed here.

A joint approach somewhat similar to the one proposed here was recently developed by Xiao *et al.* [21]. Although not formulated through a single cost function under optimization, the approach is implemented in a loop and can be considered a joint approach. However, the method employs a bilateral filter and locally adjusts filter strength by using precomputed occlusion labels in each iteration, whereas we use anisotropic diffusion and let the joint formulation drive the diffusion process by using “soft” occlusion information (no explicit occlusion labeling step). Moreover, the occlusion detection relies on intensity mismatch, that is unreliable under image noise and illumination changes [22], while we use an optical-flow mismatch.

In this paper, we deal with the deficiencies of prior approaches by proposing a variational formulation that jointly estimates optical-flow vectors, implicitly detects occlusions and extrapolates optical flow in occlusion areas. The evolving occlusions force vector extrapolation (*via* diffusion) by automatically disabling intensity matching at occluded pixels, but permit standard estimation at visible pixels. By using anisotropic diffusion driven by underlying image gradient, the interaction between occlusion-area and some visible-area optical-flow vectors is inhibited. At the same time, this joint formulation solved iteratively facilitates interaction between optical-flow vectors and occlusion labels, thus leading to more coherent solutions. Evaluation of the proposed approach on synthetic and real images demonstrates its efficacy, both qualitatively and quantitatively.

The paper is organized as follows. In Section II, we describe image-driven anisotropic optical-flow diffusion. In Section III, we describe our joint formulation, and in Section IV, we present experimental results.

II. OPTICAL FLOW EXTRAPOLATION *via* IMAGE-DRIVEN ANISOTROPIC DIFFUSION

Consider rendering intermediate view J based on views I_L and I_R [Fig. 1(b)]. In a simple scenario, disparity between I_L and I_R would be computed first, and then intensities of I_L and I_R would be disparity-compensated onto J to create the intermediate view. However, since area A undergoes occlusion in I_R , disparity cannot be computed here, which makes rendering the corresponding area O_L (visible in I_L) impossible. Similarly, disparity remains undefined in the uncovered area B , and, thus, O_R (visible in I_R) cannot be reconstructed. In order to reconstruct O_L and O_R , a proper correspondence with I_L and I_R , respectively, needs to be established *via* disparity. Also, as recently suggested, motion-based object segmentation could be improved had vectors been known in occlusion areas [23].

An analogous problem was solved for images by Bertalmio *et al.* [24]. They proposed an algorithm to fill in, or *inpaint*, missing areas in an image using surrounding intensities and their gradients. Their approach consists of two steps: first, extending available gradients into a missing area, and then applying anisotropic diffusion to propagate the available intensities into this area. The image gradients are extended first so that the subsequent anisotropic diffusion preserves them (e.g., intensity/color discontinuities). In the case of disparity inpainting, image intensities are known in occlusion areas, and, thus, the first step is not needed under the assumption that image and disparity discontinuities coincide. Therefore, we propose to extrapolate disparities using anisotropic diffusion driven by image gradient.

Let $\mathbf{x} = (x, y)^T \in \Omega$ be a spatial position in image I defined on Ω . Also, let $\{\mathbf{d}(\mathbf{x})\}_{\mathbf{x} \in \Omega}$ be a disparity field to be computed; $\mathbf{d} = [u, v]^T$ with u and v being, respectively, horizontal and vertical components of the disparity vector \mathbf{d} . Finally, let $O \subset \Omega$ be an occlusion area in image I . In order to inpaint disparities \mathbf{d} , we exploit the underlying image structure by carrying out the following minimization:

$$\min_{u,v} \int \int_O (F_{\mathbf{x}}(u, I) + F_{\mathbf{x}}(v, I)) d\mathbf{x} \quad (1)$$

where $F_{\mathbf{x}}$ is defined as follows:

$$F_{\mathbf{x}}(u, I) = \nabla^T u(\mathbf{x}) \begin{bmatrix} g(|I^x(\mathbf{x})|) & 0 \\ 0 & g(|I^y(\mathbf{x})|) \end{bmatrix} \nabla u(\mathbf{x}) \quad (2)$$

$g(\cdot)$ is a monotonically decreasing function, and I^x, I^y are horizontal and vertical derivatives of I , respectively.

The above minimization leads to anisotropic diffusion; disparities in occlusion areas are diffused while accounting for the underlying image gradient. Assuming that the gradient magnitude within an object is small, an iterative algorithm imple-

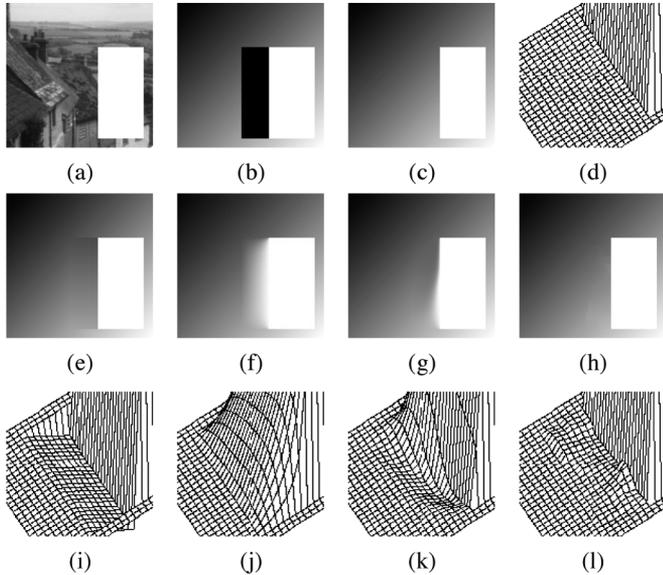


Fig. 2. Comparison of disparity extrapolation methods on computer-generated images: (a) I_L ; (b) partial disparity map with ground-truth occlusions (black); ground-truth disparity as (c) intensity image and (d) 3-D surface; and the extrapolated disparity based on (e) depth constancy along epipolar line; (f) isotropic diffusion; (g) standard inpainting; (h) proposed approach; and (i-l) corresponding 3-D surfaces of disparity extrapolated in occlusion areas.

menting (1) will diffuse disparities inside each object only. The edge-stopping function $g(\cdot)$ will prevent diffusion across object boundaries because gradient is usually large there. The cost function (2) is very similar to the one proposed by Perona and Malik [25]. However, while in Perona and Malik's case image intensity undergoes smoothing and at the same time drives the edge-stopping function $g(\cdot)$ to control anisotropy during diffusion, in our case horizontal and vertical disparity components are being (separately) smoothed but anisotropy is controlled by the underlying image intensity. This approach has also been used by others to regularize disparities [13], [14].

In Fig. 2, we compare results of the proposed image-driven anisotropic disparity diffusion with those of other extrapolation methods on a synthetic image. We assume that we are given a stereo pair [left image shown in Fig. 2(a)] and a partially estimated horizontal (1-D) disparity map with occlusions marked in black [Fig. 2(b)]. We would like to extrapolate the disparity map in the occlusion area and closely approach the ground-truth disparity, shown in Fig. 2(c) as an intensity image and in Fig. 2(d) as a 3-D surface. A simple extrapolation using depth constancy along epipolar line [13] leads to a patch-like result shown in Fig. 2(e). On the other hand, isotropic diffusion [Fig. 2(f)] is overly smooth. The results of standard inpainting and the proposed approach are shown in Fig. 2(g) and (h), respectively. The last row in Fig. 2 shows the same results in form of a 3-D surface around the occlusion area. Although standard inpainting preserves structure much better than depth constancy and isotropic diffusion, there is still unwanted smoothing especially at the bottom of the white rectangle. The image-driven anisotropic diffusion, however, produces an extrapolated disparity field with a clear discontinuity and is barely distinguishable from ground-truth.

III. JOINT OPTICAL FLOW ESTIMATION/INPAINTING

As shown in the previous section, image-driven anisotropic diffusion can be an effective tool in extrapolation of optical flow in occlusion areas. However, we assumed that occlusion areas are known as is optical flow outside occlusion areas, and, thus, the inpainting is basically a postprocessing step. Now, we propose a new approach that combines optical-flow estimation, occlusion detection and optical-flow extrapolation in a single formulation.

Let $I_L : \Omega_L \rightarrow R^+$, $I_R : \Omega_R \rightarrow R^+$, and let \mathbf{x} belong either to Ω_L or Ω_R . We would like to compute two disparity fields: $\{\mathbf{d}_L(\mathbf{x}) = [u_L(\mathbf{x}), v_L(\mathbf{x})]^T\}_{\mathbf{x} \in \Omega_L}$ and $\{\mathbf{d}_R(\mathbf{x}) = [u_R(\mathbf{x}), v_R(\mathbf{x})]^T\}_{\mathbf{x} \in \Omega_R}$ that, for pixels visible in both images, minimize some metric of the following photometric errors:

$$\begin{aligned} \rho_{LR}(\mathbf{x}) &= I_L(\mathbf{x}) - I_R(\mathbf{x} + \mathbf{d}_L(\mathbf{x})) \\ \rho_{RL}(\mathbf{x}) &= I_R(\mathbf{x}) - I_L(\mathbf{x} + \mathbf{d}_R(\mathbf{x})). \end{aligned} \quad (3)$$

At the same time, in order to distinguish occluded and visible image areas, we propose to use the disparity mismatch (geometric constraint)

$$\begin{aligned} \epsilon_L(\mathbf{x}) &= \|\mathbf{d}_L(\mathbf{x}) + \mathbf{d}_R(\mathbf{x} + \mathbf{d}_L(\mathbf{x}))\| \\ \epsilon_R(\mathbf{x}) &= \|\mathbf{d}_R(\mathbf{x}) + \mathbf{d}_L(\mathbf{x} + \mathbf{d}_R(\mathbf{x}))\| \end{aligned} \quad (4)$$

where $\|\cdot\|$ denotes Euclidean norm. Both $\epsilon_L(\mathbf{x})$ and $\epsilon_R(\mathbf{x})$ are expected to be small for visible pixels and larger for occluded and uncovered pixels. Note that although photometric errors (3) could have been used as occlusion indicators, they are less robust to noise and intensity variations [22].

In order to model data-matching, disparity and occlusion constraints, we propose three pairs of energy functions (six in total). Combined using adjustable weights, these functions will allow us, during minimization, to trade off intensity matching accuracy for disparity and occlusion smoothness.

Since photometric constraints, expressed through errors (3), do not hold in occlusion areas (ρ_{LR}^2 and ρ_{RL}^2 are large), we need to disable their impact on the overall cost function whenever ϵ_L or ϵ_R is large (indicating occlusion). We can accomplish this by multiplying ρ_{LR}^2 and ρ_{RL}^2 by a weight function inversely proportional to ϵ_L and ϵ_R (4), respectively. The larger the disparity mismatch, the smaller will be the contribution of this photometric error to the overall cost function. We propose a monotonically decreasing weight function $D(z) = 1/(1 + Kz^2)$, with constant $K > 0$ controlling function's slope. $D(\epsilon_L(\mathbf{x}))$ and $D(\epsilon_R(\mathbf{x}))$ approach zero as the disparities $\mathbf{d}_L, \mathbf{d}_R$ are less and less capable of compensating each other. We define the first pair of energy functions as follows:

$$\begin{aligned} E_L^P &= \int \int_{\Omega_L} D(\epsilon_L(\mathbf{x})) [\rho_{LR}(\mathbf{x})]^2 d\mathbf{x} \\ E_R^P &= \int \int_{\Omega_R} D(\epsilon_R(\mathbf{x})) [\rho_{RL}(\mathbf{x})]^2 d\mathbf{x}. \end{aligned} \quad (5)$$

These energies differ from the usual optical flow formulation by their ability to disable the impact of photometric error when the disparities do not compensate each other. This is essential because these areas are most likely occluded and the intensity

matching term is not beneficial. On the contrary, it may lead to false solutions. A similar idea was used in [9], where a prediction term was disabled in favor of another prediction term (multiple images available), while here the prediction term is disabled in favor of diffusion term (two images only).

We embed the idea of image-driven disparity extrapolation through the second pair of energy functions

$$\begin{aligned} E_L^S &= \int \int_{\Omega_L} (F_{\mathbf{x}}(u_L, I_L) + F_{\mathbf{x}}(v_L, I_L)) d\mathbf{x} \\ E_R^S &= \int \int_{\Omega_R} (F_{\mathbf{x}}(u_R, I_R) + F_{\mathbf{x}}(v_R, I_R)) d\mathbf{x} \end{aligned} \quad (6)$$

with $F_{\mathbf{x}}$ defined in (2). Note that energies (5) and (6) jointly lead to edge-preserving regularization (no disparity smoothing across strong intensity gradients) when $D(\cdot)$ is close to 1, but result in disparity inpainting when $D(\cdot)$ is zero (i.e., possible occlusion area), since the data-matching terms are disabled.

The energies (5) can be easily made arbitrarily small by choosing vector fields with sufficiently large ϵ_L and ϵ_R (4) for all \mathbf{x} . In order to prevent this, we propose an explicit occlusion model through the following energies:

$$\begin{aligned} E_L^O &= \int \int_{\Omega_L} (1 - D(\epsilon_L(\mathbf{x}))) d\mathbf{x} \\ E_R^O &= \int \int_{\Omega_R} (1 - D(\epsilon_R(\mathbf{x}))) d\mathbf{x}. \end{aligned} \quad (7)$$

Note that $1 - D(\epsilon_L(\mathbf{x}))$ and $1 - D(\epsilon_R(\mathbf{x}))$ approach 1 as the geometric errors $\epsilon_L(\mathbf{x})$ and $\epsilon_R(\mathbf{x})$ grow, and can be thought of as occlusion indicators in I_L and I_R , respectively. The above energy terms, by introducing a penalty at each occlusion point, keep the total area of occlusions from growing indefinitely. Otherwise, all image points declared as occluded would result in a low-energy, but degenerate, solution. Since minimization of these terms encourages $D(z)$ to be close to 1, the computed vector fields are encouraged to be as close inverses of each other as possible (except occlusion areas).

In order to perform joint disparity estimation, implicit occlusion detection and disparity extrapolation, we combine the above energy terms and carry out the following two minimizations *simultaneously*:

$$\begin{aligned} \min_{\mathbf{d}_L} E_L, \quad E_L &= E_L^P + \eta E_L^S + \mu E_L^O \\ \min_{\mathbf{d}_R} E_R, \quad E_R &= E_R^P + \eta E_R^S + \mu E_R^O \end{aligned} \quad (8)$$

where η and μ are regularization factors. Since E_L and E_R are functionals of both \mathbf{d}_L and \mathbf{d}_R (through $D(\epsilon_L)$ and $D(\epsilon_R)$), the minimization is performed in an interleaved fashion; one iteration of minimization is performed on E_L and the estimated \mathbf{d}_L is used in one iteration on E_R , and vice versa. A derivation of Euler–Lagrange equations can be found in the Appendix.

One may wonder why not minimize the sum $E_L + E_R$ with respect to \mathbf{d}_L and \mathbf{d}_R instead. It turns out that such a formulation leads locally (around occlusion areas) to contradictory constraints. Consider $\mathbf{x} \notin A \subset \Omega_L$ but such that $\mathbf{x} \in B \subset \Omega_R$ [Fig. 1(a)]; intensity at \mathbf{x} in I_L is visible in I_R but intensity at \mathbf{x} in I_R is not visible in I_L . When minimizing E_L with respect to \mathbf{d}_L at \mathbf{x} (8), a reliable disparity estimate can be

obtained since an intensity match can be established between I_L and I_R . When minimizing $E_L + E_R$ also with respect to \mathbf{d}_L at \mathbf{x} , the following energy terms need to be considered: $E_L^P + \eta E_L^S + \mu E_L^O + E_R^P + \mu E_R^O$ (note that E_R^S is independent of \mathbf{d}_L). Clearly, in addition to E_L , also the sum $E_R^P + \mu E_R^O$ is minimized, thus biasing the estimate away from the correct solution obtained when minimizing E_L only. This bias occurs since $E_R^P + \mu E_R^O$ is a function of *both* \mathbf{d}_L and \mathbf{d}_R , and is particularly severe for \mathbf{x} visible in I_L but uncovered in I_R since $\mathbf{d}_R(\mathbf{x})$ is unreliable (obtained by anisotropic diffusion). The impact of $\mathbf{d}_R(\mathbf{x})$ occurs through $\rho_{RL}^2(\mathbf{x})$ which may be large since \mathbf{x} is in an uncovered area of I_R (no match in I_L), and through $\epsilon_R(\mathbf{x})$ which may also be large for the same reason ($\mathbf{d}_R(\mathbf{x})$ points to a *visible* area in I_L from where \mathbf{d}_L is unlikely to point back to \mathbf{x}). In order to reduce both terms, \mathbf{d}_L diverges from the correct estimate. Such erroneous solutions occur only around occlusion/uncovered areas, which we confirmed experimentally. However, in separate, but alternating, minimizations of E_L with respect to \mathbf{d}_L and of E_R with respect to \mathbf{d}_R , no such effects take place. Although we have no proof of convergence of alternating minimizations (8), a multiresolution implementation (see below) and regularizers (6) help avoid divergence, and we did not experience it in our experiments. Note that a similar minimization strategy was used by Alvarez *et al.* [16].

IV. IMPLEMENTATION AND EXPERIMENTAL RESULTS

We discretized the resulting partial differential evolution equations using finite differences (see [25] for the discretization of anisotropic diffusion). We used an explicit discretization scheme for its simplicity, and a small time step ($\Delta t = 1.5 \times 10^{-5}$) to assure stability of calculations. All sub-pixel (noninteger position) values, e.g., $I_R(\mathbf{x} + \mathbf{d}_L(\mathbf{x}))$, were computed using bicubic interpolation. We used a hierarchical implementation to avoid local minima. Images were prefiltered with a Gaussian filter and downsampled so that at the lowest resolution the maximum disparity did not exceed 1–2 pixels. The estimation was started at the lowest resolution and the result propagated to the next higher resolution by interpolation. We used rectified stereo pairs, i.e., $\mathbf{d} = [u \ 0]^T$; the vertical disparity component was set to zero in estimation, therefore leading to a simplified version of the algorithm.

In order to carry out evaluation of the proposed algorithm, we introduce two different weighting functions D in our energy formulation: $D_1(z) = 1/(1 + K_1 z^2)$ which weights the photometric error in (5) and $D_2(z) = 1/(1 + K_2 z^2)$ which keeps the total area of occlusions from growing indefinitely (7). As shown in Table I, for different values of K_1 and K_2 , and different functional forms of $g(z)$ our formulation may be simplified to the original optical flow [1], edge-preserving optical flow [5], or symmetric optical flow [16] estimation, the latter one forcing the two disparity fields to be close inverses of each other. The symmetric optical flow algorithm includes the occlusion-limiting term (7) but does not disable the data-matching term in (5). This is of interest for state-of-the-art video coding based on the discrete wavelet transform (DWT) as it is able to ensure a close invertibility of vector fields, important for such coders [26]. Also, note that for $K_1 = K_2 = 0$ and $g(z) = 1$,

TABLE I
OPTICAL FLOW (OF) ESTIMATION ALGORITHMS TESTED

Algorithm	K_1	K_2	$g(z)$
Original OF [1]	0	0	1
Edge-preserving OF [5]	0	0	monotonically-decreasing
Symmetric OF [16]	0	>0	monotonically-decreasing
Proposed algorithm	>0	>0	monotonically-decreasing

TABLE II
ABSOLUTE ERROR PER PIXEL IN COMPUTED DISPARITY FIELDS

	Image #1 (Fig. 3)		Image #2 (Fig. 4)	
	u_L	u_R	u_L	u_R
Original OF	4.57	4.67	1.63	1.44
Edge-preserving OF	1.55	1.51	0.81	0.52
Symmetric OF	1.61	1.83	0.60	0.45
Proposed algorithm	0.58	0.53	0.35	0.36

TABLE III
ABSOLUTE DISPARITY ERROR PER PIXEL FOR u_L ON TEST IMAGE FROM FIG. 4 AT DIFFERENT LEVELS OF ZERO-MEAN WHITE GAUSSIAN NOISE

Resulting PSNR(dB)	Original OF	Edge-preserving OF	Symmetric OF	Proposed algorithm
No noise	1.63	0.81	0.60	0.35
27.01	1.66	0.91	0.74	0.50
24.09	1.69	1.08	0.81	0.60
23.12	1.69	1.00	0.87	0.64
20.35	1.80	1.18	0.96	0.72

TABLE IV
ABSOLUTE DISPARITY ERROR PER PIXEL FOR THE TEST IMAGE FROM FIG. 4 AND DIFFERENT PARAMETER VALUES. IN EACH EXPERIMENT, ONE PARAMETER IS ADJUSTED WHILE OTHER PARAMETERS ARE UNCHANGED

$\eta = 6000, \mu = 2000$			$K = 10, \mu = 2000$			$K = 10, \eta = 6000$		
K	u_L	u_R	η	u_L	u_R	μ	u_L	u_R
3	0.52	0.46	1000	0.54	0.45	100	1.00	1.16
7	0.47	0.43	3000	0.43	0.40	1000	0.53	0.47
10	0.35	0.36	6000	0.35	0.36	2000	0.35	0.36
12	0.37	0.36	9000	0.37	0.37	3000	0.44	0.43

minimizations in (8) reduce to two original optical flow algorithms executed in parallel. In all experiments, whenever K_1 and K_2 are nonzero we use the value of 10, while $\eta = 6000$ and $\mu = 2000$.

First, we tested the four approaches on two synthetic sequences. Fig. 3 shows an unusually shaped object that is displaced horizontally by 15 pixels over a stationary background. The original images and corresponding ground-truth occlusion area for I_L are in the top row of Fig. 3. The ground-truth disparity map for I_L and its four estimates, presented as intensity, as well as the recovered occlusions are shown in the remaining two rows. The second synthetic sequence (Fig. 4) is more challenging; two circles displace in opposite directions. There are three occlusion regions between images and a significant portion of occlusions is due to one object covering the other. Table II shows the absolute error per pixel for the estimated disparities.

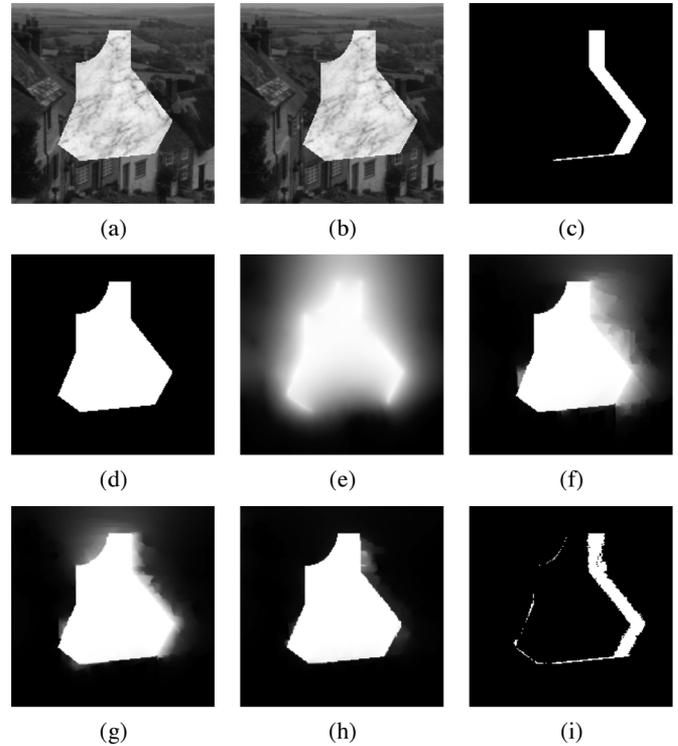


Fig. 3. Results for a computer-generated pair of images: (a) I_L ; (b) I_R ; ground-truth: (c) occlusions for I_L and (d) disparity for I_L , and disparities for I_L computed using progressively more complex formulations: (e) original OF; (f) edge-preserving OF; (g) symmetric OF; (h) proposed method; and (i) likely occlusion areas obtained by thresholding $1 - D(\epsilon_L(x))$. In disparity images, black and white colors represent 0 and 15 pixels of disparity, respectively.

Clearly, the proposed method [Figs. 3(h) and 4(h)] outperforms the original [Figs. 3(e) and 4(e)] and edge-preserving optical flow [Figs. 3(f) and 4(f)] algorithms, both subjectively and numerically. Note a significant improvement offered by edge-preserving regularization compared to the original optical flow algorithm. The symmetric optical flow algorithm [Figs. 3(g) and 4(g)] offers some subjective and numerical advantage over the edge-preserving optical flow but since it enforces forward/backward vector consistency at occluded pixels, the improvement is limited. Had the occlusion areas been very small, the symmetric optical flow would have improved the results significantly [16]. In our images, however, disparity mismatch over large occlusion areas affects visible pixels through diffusion and results in disparity errors. Still, this method is of interest for DWT-based video coding due to close mutually inverse properties of the resulting vector fields [26].

We also compared the four approaches in the presence of noise; we added zero-mean white Gaussian noise to the test image from Fig. 3. The absolute disparity error per pixel for u_L is shown in Table III for different levels of noise. Clearly, the proposed method performs well under noise as well. This can be explained by the adaptive nature of the algorithm; since disparities at noisy pixels usually lead to significant geometric errors (4), the contribution from these pixels is disabled in (5). It should be also noted that the hierarchical scheme used, which includes a prefiltering step, acts as a noise suppressor and helps all tested methods deal with noise.

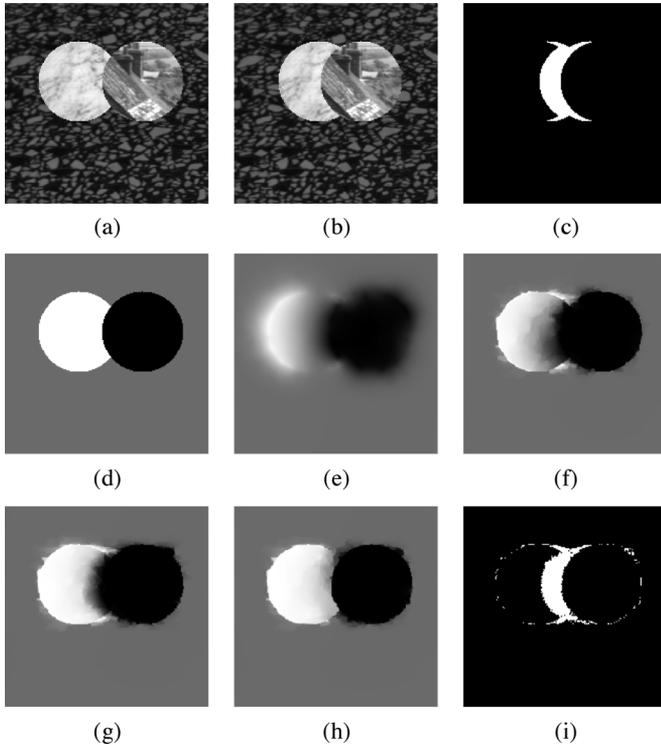


Fig. 4. Results for another pair of computer-generated images; see caption of Fig. 3 for description. In the disparity images, black, gray, and white colors represent -10 , 0 , and 10 pixels of disparity, respectively.

Finally, we tested the algorithms on camera-acquired images: *Exit* (Fig. 5) and *Michel* (Fig. 6). We present results for the symmetric optical-flow and the proposed algorithm only. For *Exit*, the improvements are clear in occlusion area to the right of the person closest to the camera, visible especially in close-up images [Fig. 5(e) and (f)]; in the symmetric optical flow result, there is a clear spillover of disparities from person's body into the background (pixels in the occlusion area fail to find correspondence in the other image). However, such errors are largely corrected by the proposed method, because the matching term (5) is disabled in occlusion areas. The estimated occlusion areas are shown in Fig. 5(g). Similar improvements can be observed for *Michel* [Fig. 6(c)–(f)]. Note the large occlusion areas, e.g., behind the head, that lead to incorrect large disparities for symmetric optical flow [Fig. 6(c) and (e)], but are corrected by the proposed method [Fig. 6(d) and (f)]. In this example, we used $\mu = 5000$, because the occlusion area is much larger; unless the occlusion count is penalized, it will grow beyond reasonable limits.

A. Parameter Selection

As in other methods, an important issue is parameter selection. In the proposed formulation, three parameters K , η and μ influence the results. We chose $\eta = 6000$ and $\mu = 2000$ experimentally. While a larger η would force an even smoother disparity field, a larger μ would further reduce the number of (implicitly) estimated occlusion pixels. We chose K to be 10 because when ϵ , i.e., mismatch between vector fields, is larger than 1 pixel, $D(\epsilon)$ falls below 0.1, a small enough value to significantly reduce contribution of the intensity matching term.

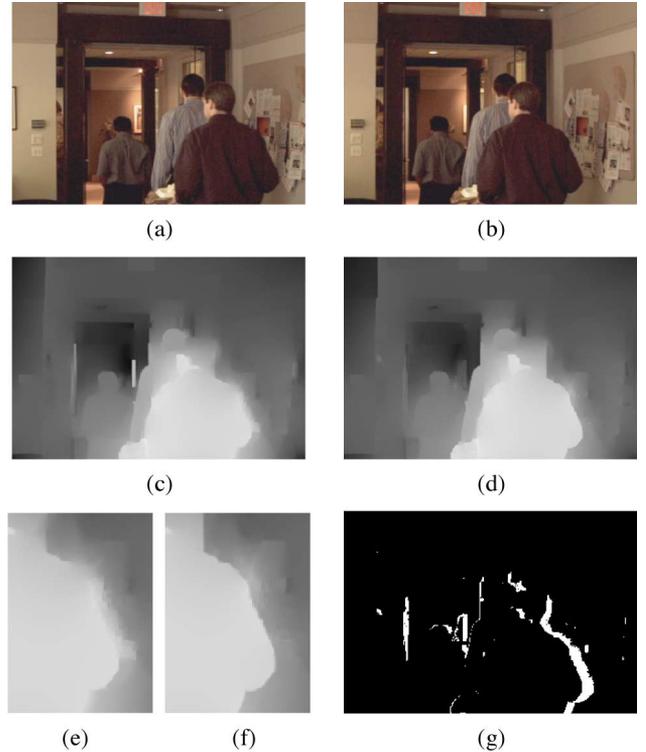


Fig. 5. Experimental results for *Exit* image pair (property of Mitsubishi Electric Research Laboratories) : (a) I_L ; (b) I_R ; estimated disparity for I_R : (c) symmetric-OF; and (d) proposed method; and (e)–(f) close-ups of results from (c)–(d), (g) likely occlusion areas obtained by thresholding $1 - D(\epsilon_R(\mathbf{x}))$.

Note that there exist methods such as expectation maximization [27], min-max principle [28] and unbiased risk estimator [29] that can be used to automatically select parameter values.

In order to demonstrate that a very precise selection of parameters is not necessary in our method, Table IV shows absolute disparity error per pixel for the test image from Fig. 4 while changing either K , or η , or μ . It can be seen that as parameters are increased threefold, the error changes at most by 10%–20%. We also tested a relatively very small value of $\mu = 100$; almost half of the pixels were marked as occlusion. This was to be expected since, as we mentioned earlier, this leads to mismatched disparities \mathbf{d}_L and \mathbf{d}_R and, consequently, to disabling of the photometric error (5).

V. CONCLUSIONS AND FUTURE WORK

We presented a variational framework for joint optical flow estimation, occlusion detection, and optical flow extrapolation based on two images only. The new formulation calculates two closely symmetric flow fields and also inpaints the optical flow in occlusion areas. This formulation improves state-of-the-art approaches in three ways. First, the image-driven anisotropic diffusion fills-in flow vectors in occlusion areas respecting image structure (intensity discontinuities), therefore providing *plausible* solutions (unlike constant-disparity extrapolation). Second, the joint formulation permits interaction between optical flow and occlusions during estimation, thus allowing mutual corrections (unlike in the case of three-step approaches). Third, by disabling the data-matching term in occlusion areas, the estimation bias of optical flow vectors is eliminated since

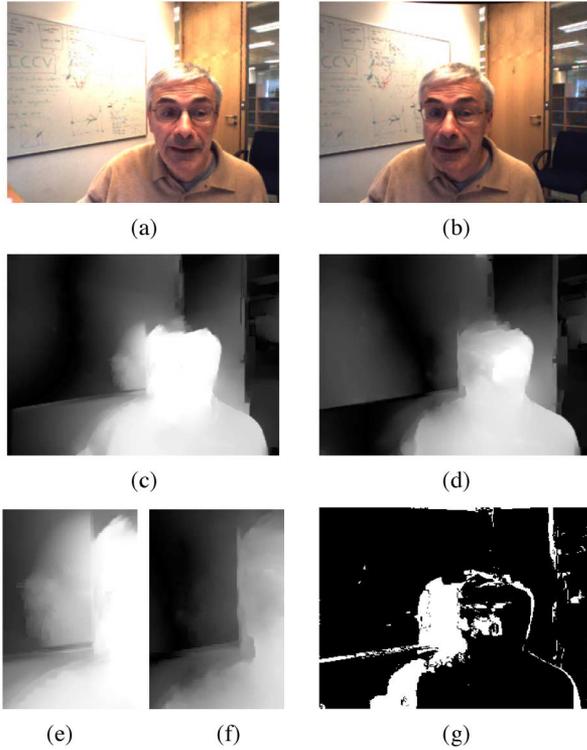


Fig. 6. Experimental results for *Michel* image pair (property of Microsoft Research Cambridge, U.K.): see caption of Fig. 5 for description.

the system relies exclusively on anisotropic diffusion (in visible areas data-driven flow and diffusion work together).

The proposed algorithm has shown significant improvement over original and edge-preserving optical flow formulations both subjectively and numerically. Moreover, the symmetric variant of the proposed algorithm may be interesting for DWT-based video coding because of the particular relationship between the resulting vector fields (close mutual inverses) [26]. We observed that the proposed approach brings only about 40% of additional computational load to the standard optical flow algorithm due to the additional interpolation operations stemming from terms such as \tilde{v}^x . We also observed that for larger occlusion areas the number of iterations must be increased so that diffusion can fill-in the areas using neighboring values. One shortcoming of the method is evident in highly textured images because image-driven flow diffusion is inhibited due to the high local intensity gradient; a common problem of image-driven regularizers. This problem is being currently addressed.

APPENDIX A

Let $\Phi_L = \{u_L, v_L\}$, $\Phi_R = \{u_R, v_R\}$ be sets of disparity field components that we seek by minimizations in (8). Assuming that $\Omega_L = \Omega_R = \Omega$ in all energy terms (5)–(7), we can rewrite cost functionals in (8) as follows:

$$E_L = \int \int_{\Omega} e_L(\mathbf{x}) d\mathbf{x}, \quad e_L(\mathbf{x}) = e_L^P(\mathbf{x}) + \eta e_L^S(\mathbf{x}) + \mu e_L^O(\mathbf{x})$$

$$E_R = \int \int_{\Omega} e_R(\mathbf{x}) d\mathbf{x}, \quad e_R(\mathbf{x}) = e_R^P(\mathbf{x}) + \eta e_R^S(\mathbf{x}) + \mu e_R^O(\mathbf{x})$$

where

$$e_L^P(\mathbf{x}) = D_L(\mathbf{x})[\rho_{LR}(\mathbf{x})]^2$$

$$e_R^P(\mathbf{x}) = D_R(\mathbf{x})[\rho_{RL}(\mathbf{x})]^2$$

$$e_L^S(\mathbf{x}) = F_{\mathbf{x}}(u_L, I_L) + F_{\mathbf{x}}(v_L, I_L)$$

$$e_R^S(\mathbf{x}) = F_{\mathbf{x}}(u_R, I_R) + F_{\mathbf{x}}(v_R, I_R)$$

$$e_L^O(\mathbf{x}) = (1 - D_L(\mathbf{x}))$$

$$e_R^O(\mathbf{x}) = (1 - D_R(\mathbf{x}))$$

and $D_L(\mathbf{x}) \triangleq D(\epsilon_L(\mathbf{x}))$, $D_R(\mathbf{x}) \triangleq D(\epsilon_R(\mathbf{x}))$. We will minimize E_L and E_R simultaneously by assuming that \mathbf{d}_L is constant when computing \mathbf{d}_R and vice versa. This will lead to interleaved descent equations, i.e., one iteration of \mathbf{d}_L using the values of \mathbf{d}_R from previous iteration and vice versa.

Using the calculus of variations, two Euler–Lagrange equations (one for each unknown in Φ) for each E , can be found in the form of

$$e'_L(\omega_L) = (\partial e_L)/(\partial \omega_L) - (\partial)/(\partial x)(\partial e_L)/(\partial \omega_L^x) - (\partial)/(\partial y)(\partial e_L)/(\partial \omega_L^y) = 0$$

and

$$e'_R(\omega_R) = (\partial e_R)/(\partial \omega_R) - (\partial)/(\partial x)(\partial e_R)/(\partial \omega_R^x) - (\partial)/(\partial y)(\partial e_R)/(\partial \omega_R^y) = 0$$

where $e'_L(\omega_L)$ and $e'_R(\omega_R)$ are the first variations with respect to $\omega_L \in \Phi_L$, $\omega_R \in \Phi_R$, whereas ω^x and ω^y are derivatives with respect to x and y , respectively. Expanding each equation and omitting derivatives that are equal to zero (e.g., $\partial e_L^P/\partial u_L^x = 0$), we get two Euler–Lagrange equations for each E as follows:

$$\frac{\partial e_L^P}{\partial \omega_L} + \mu \frac{\partial e_L^O}{\partial \omega_L} - \eta \left(\frac{\partial}{\partial x} \frac{\partial e_L^S}{\partial \omega_L^x} + \frac{\partial}{\partial y} \frac{\partial e_L^S}{\partial \omega_L^y} \right) = 0$$

$$\frac{\partial e_R^P}{\partial \omega_R} + \mu \frac{\partial e_R^O}{\partial \omega_R} - \eta \left(\frac{\partial}{\partial x} \frac{\partial e_R^S}{\partial \omega_R^x} + \frac{\partial}{\partial y} \frac{\partial e_R^S}{\partial \omega_R^y} \right) = 0 \quad (9)$$

where, again, $\omega_L \in \Phi_L$ and $\omega_R \in \Phi_R$. Partial derivatives with respect to u_L, v_L, u_R, v_R can be computed as follows (\mathbf{x} was dropped for simplicity of notation)

$$\frac{\partial e_L^P}{\partial u_L} = \frac{\partial D_L}{\partial u_L} (\rho_{LR})^2 - 2D_L \tilde{I}_R^x \rho_{LR}$$

$$\frac{\partial e_L^P}{\partial v_L} = \frac{\partial D_L}{\partial v_L} (\rho_{LR})^2 - 2D_L \tilde{I}_R^y \rho_{LR}$$

$$\frac{\partial e_R^P}{\partial u_R} = \frac{\partial D_R}{\partial u_R} (\rho_{RL})^2 - 2D_R \tilde{I}_L^x \rho_{RL}$$

$$\frac{\partial e_R^P}{\partial v_R} = \frac{\partial D_R}{\partial v_R} (\rho_{RL})^2 - 2D_R \tilde{I}_L^y \rho_{RL}$$

$$\frac{\partial}{\partial x} \frac{\partial e_L^S}{\partial u_L^x} + \frac{\partial}{\partial y} \frac{\partial e_L^S}{\partial u_L^y} = \frac{\partial (2g(|I_L^x|)u_L^x)}{\partial x} + \frac{\partial (2g(|I_L^y|)u_L^y)}{\partial y}$$

$$\frac{\partial}{\partial x} \frac{\partial e_L^S}{\partial v_L^x} + \frac{\partial}{\partial y} \frac{\partial e_L^S}{\partial v_L^y} = \frac{\partial (2g(|I_L^x|)v_L^x)}{\partial x} + \frac{\partial (2g(|I_L^y|)v_L^y)}{\partial y}$$

$$\frac{\partial}{\partial x} \frac{\partial e_R^S}{\partial u_R^x} + \frac{\partial}{\partial y} \frac{\partial e_R^S}{\partial u_R^y} = \frac{\partial (2g(|I_R^x|)u_R^x)}{\partial x} + \frac{\partial (2g(|I_R^y|)u_R^y)}{\partial y}$$

$$\frac{\partial}{\partial x} \frac{\partial e_R^S}{\partial v_R^x} + \frac{\partial}{\partial y} \frac{\partial e_R^S}{\partial v_R^y} = \frac{\partial (2g(|I_R^x|)v_R^x)}{\partial x} + \frac{\partial (2g(|I_R^y|)v_R^y)}{\partial y}$$

$$\frac{\partial e_L^O}{\partial u_L} = -\frac{\partial D_L}{\partial u_L}, \quad \frac{\partial e_L^O}{\partial v_L} = -\frac{\partial D_L}{\partial v_L}$$

$$\frac{\partial e_R^O}{\partial u_R} = -\frac{\partial D_R}{\partial u_R}, \quad \frac{\partial e_R^O}{\partial v_R} = -\frac{\partial D_R}{\partial v_R}$$

where I_x^x and I_y^y are horizontal and vertical derivatives of I , while \tilde{I}_L^x and \tilde{I}_L^y are derivatives evaluated off \mathbf{x} , e.g., $\tilde{I}_L^x(\mathbf{x}) = I_L^x(\mathbf{x} + \mathbf{d}_R(\mathbf{x}))$. Furthermore, we have

$$\begin{aligned}\frac{\partial D_L}{\partial u_L} &= -2K \frac{(1 + \tilde{u}_R^x)\epsilon_{L,u} + \tilde{v}_R^x\epsilon_{L,v}}{(1 + K(\epsilon_{L,u})^2 + K(\epsilon_{L,v})^2)^2} \\ \frac{\partial D_L}{\partial v_L} &= -2K \frac{\tilde{u}_R^y\epsilon_{L,u} + (1 + \tilde{v}_R^y)\epsilon_{L,v}}{(1 + K(\epsilon_{L,u})^2 + K(\epsilon_{L,v})^2)^2} \\ \frac{\partial D_R}{\partial u_R} &= -2K \frac{(1 + \tilde{u}_L^x)\epsilon_{R,u} + \tilde{v}_L^x\epsilon_{R,v}}{(1 + K(\epsilon_{R,u})^2 + K(\epsilon_{R,v})^2)^2} \\ \frac{\partial D_R}{\partial v_R} &= -2K \frac{\tilde{u}_L^y\epsilon_{R,u} + (1 + \tilde{v}_L^y)\epsilon_{R,v}}{(1 + K(\epsilon_{R,u})^2 + K(\epsilon_{R,v})^2)^2}\end{aligned}$$

where

$$\begin{aligned}\epsilon_{L,u}(\mathbf{x}) &= u_L(\mathbf{x}) + u_R(\mathbf{x} + \mathbf{d}_L(\mathbf{x})) \\ \epsilon_{L,v}(\mathbf{x}) &= v_L(\mathbf{x}) + v_R(\mathbf{x} + \mathbf{d}_L(\mathbf{x})) \\ \epsilon_{R,u}(\mathbf{x}) &= u_R(\mathbf{x}) + u_L(\mathbf{x} + \mathbf{d}_R(\mathbf{x})) \\ \epsilon_{R,v}(\mathbf{x}) &= v_R(\mathbf{x}) + v_L(\mathbf{x} + \mathbf{d}_R(\mathbf{x}))\end{aligned}$$

are individual components of disparity errors (4) and $\tilde{u}_L^x, \tilde{v}_L^x, \tilde{u}_L^y, \tilde{v}_L^y$ are again derivatives evaluated off \mathbf{x} , e.g., $\tilde{u}_L^x(\mathbf{x}) = u_L^x(\mathbf{x} + \mathbf{d}_R(\mathbf{x}))$. Using an auxiliary time variable t , equations in (9) can be solved by discretizing gradient descent equations

$$\partial\omega_L/\partial t = -e'_L(\omega), \partial\omega_R/\partial t = -e'_R(\omega), \omega_L \in \Phi_L$$

and

$$\omega_R \in \Phi_R.$$

ACKNOWLEDGMENT

The authors would like to thank Prof. M. Gennert of Worcester Polytechnic Institute for a suggestion regarding occlusion terms (7).

REFERENCES

- [1] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.*, vol. 17, pp. 185–203, 1981.
- [2] B. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. Int. Joint Conf. Artificial Intelligence*, Aug. 1981, pp. 674–679.
- [3] R. March, "Computation of stereo disparity using regularization," *Pattern Recognit. Lett.*, vol. 8, pp. 181–187, Oct. 1988.
- [4] M. J. Black and P. Anandan, "The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields," *Comput. Vis. Image Understanding*, vol. 63, no. 1, pp. 75–104, 1996.
- [5] L. Alvarez, R. Deriche, J. Sánchez, and J. Weickert, "Dense disparity map estimation respecting image discontinuities: A PDE and scale-space based approach," *J. Vis. Commun. Image Representation*, vol. 13, pp. 3–21, 2002.
- [6] T. Brox, A. Bruhn, N. Papenbergh, and J. Weickert, "High accuracy optical flow estimation based on a theory for warping," in *Proc. European Conf. Computer Vision*, 2004, pp. 25–36.
- [7] R. Depommier and E. Dubois, "Motion estimation with detection of occlusion areas," in *Proc. IEEE Int. Conf. Acoustics Speech Signal Processing*, Mar. 1992, pp. III.269–III.272.
- [8] N. Grammalidis and M. G. Strintzis, "Disparity and occlusion estimation in multiocular systems and their coding for the communication of multiview image sequences," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 8, no. 3, pp. 328–344, Jun. 1998.
- [9] C. Strecha and L. Van Gool, "PDE-based multi-view depth estimation," in *Proc. 1st Int. Symp. 3D Data Processing Visualization and Transmission*, 2002, vol. 2, pp. 416–425.
- [10] L. Alvarez, J. Esclarin, M. Lefébure, and J. Sánchez, "A PDE model for computing the optical flow," in *Proc. XVI Congreso de Ecuaciones Diferenciales y Aplicaciones*, Sep. 1999, pp. 1349–1356.
- [11] H.-H. Nagel and W. Enkelmann, "An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, no. 5, pp. 565–593, Sep. 1986.
- [12] A.-R. Mansouri, A. Mitiche, and J. Konrad, "Selective image diffusion: Application to disparity estimation," in *Proc. IEEE Int. Conf. Image Processing*, Oct. 1998, vol. 3, pp. 284–288.
- [13] H. Kim and K. Sohn, "3D reconstruction from stereo images for interactions between real and virtual objects," *Signal Process., Image Commun.*, vol. 20, no. 1, pp. 61–75, 2005.
- [14] X. Huang and E. Dubois, "Three-view dense disparity estimation with occlusion detection," in *Proc. IEEE Int. Conf. Image Processing*, Sep. 2005, vol. 3, pp. 393–396.
- [15] M. Proesmans, L. Van Gool, E. Pauwels, and A. Oosterlinck, "Determination of optical flow and its discontinuities using non-linear diffusion," in *Proc. European Conf. Computer Vision*, 1994, vol. 2, pp. 295–304.
- [16] L. Alvarez, R. Deriche, T. Papadopoulou, and J. Sánchez, "Symmetrical dense optical flow estimation with occlusions detection," in *Proc. European Conf. Computer Vision*, 2002, vol. 1, pp. 721–736.
- [17] J. Sun, Y. Li, S. B. Kang, and H.-Y. Shum, "Symmetric stereo matching for occlusion handling," presented at the IEEE Conf. Computer Vision Pattern Recognition, 2005.
- [18] K. P. Lim, A. Das, and M. N. Chong, "Estimation of occlusion and dense motion fields in a bidirectional Bayesian framework," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 24, no. 5, pp. 712–718, May 2002.
- [19] V. Kolmogorov and R. Zabih, "Computing visual correspondence with occlusions using graph cuts," in *Proc. IEEE Int. Conf. Computer Vision*, 2001, vol. 2, pp. 508–515.
- [20] J. Xiao and M. Shah, "Motion layer extraction in the presence of occlusion using graph cuts," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 27, no. 10, pp. 1644–1659, Oct. 2005.
- [21] J. Xiao, H. Cheng, H. Sawhney, C. Rao, and M. Isnardi, "Bilateral filtering-based optical flow estimation with occlusion detection," presented at the European Conf. Computer Vision, 2006.
- [22] S. Ince and J. Konrad, "Geometry-based estimation of occlusions from video frame pairs," in *Proc. IEEE Int. Conf. Acoustics Speech Signal Processing*, Mar. 2005, vol. II, pp. 933–936.
- [23] P. Aguiar and J. Moura, "Figure-ground segmentation from occlusion," *IEEE Trans. Image Process.*, vol. 14, no. 8, pp. 1109–1124, Aug. 2005.
- [24] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Proc. 27th Annu. Conf. Computer Graphics and Interactive Techniques*, Jul. 2000, pp. 417–424.
- [25] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 7, pp. 629–639, Jul. 1990.
- [26] J. Konrad and N. Božinović, "Importance of motion in motion-compensated temporal discrete wavelet transforms," *Proc. SPIE Image and Video Communications and Process.*, vol. 5685, pp. 354–365, Jan. 2005.
- [27] A. P. Dempster, N. M. Laird, and D. P. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc. B*, vol. 39, no. 1, pp. 1–38, 1977.
- [28] M. Gennert and A. Yuille, "Determining the optimal weights in multiple objective function optimization," in *Proc. IEEE Int. Conf. Computer Vision*, 1988, vol. 2, pp. 87–89.
- [29] L. Ng and V. Solo, "A data-driven method for choosing smoothing parameters in optical flow problems," in *Proc. IEEE Int. Conf. Image Processing*, 1997, pp. 360–363.



Serdar Ince received the B.S. and M.S. degrees from the Middle East Technical University, Ankara, Turkey, and the Ph.D. degree from Boston University, Boston, MA, all in electrical engineering.

He is now at IntelliVid Corporation, Cambridge, MA. His research interests include image/video processing, stereoscopic and 3-D imaging, and video analytics.



Janusz Konrad (M'93–SM'98–F'08) received the M.Eng. degree from the Technical University of Szczecin, Poland, in 1980, and the Ph.D. degree from McGill University, Montréal, QC, Canada, in 1989.

From 1989 to 2000, he was with INRS-Télécommunications, Montréal. Since 2000, he has been with Boston University, Boston, MA.

Dr. Konrad is an Associate Technical Editor for the *IEEE Communications Magazine* and Associate Editor for the *EURASIP International Journal on Image and Video Processing*. He was an Associate Editor for the *IEEE TRANSACTIONS ON IMAGE PROCESSING* and the *IEEE TRANSACTIONS SIGNAL PROCESSING LETTERS*, member of the IMDSP Technical Committee of the IEEE Signal Processing Society, as well as the Technical Program Co-Chair of the ICIP 2000 and Tutorials Co-Chair of the ICASSP 2004. He is a co-recipient of the 2001 Signal Processing Magazine award for a paper co-authored with Dr. C. Stiller and the 2004–2005 EURASIP Image Communications Best Paper Award for paper co-authored with Dr. N. Božinović.