

Toward robust estimation of specular flow

Yair Adato¹

<http://www.cs.bgu.ac.il/~adato>

Todd Zickler²

<http://www.eecs.harvard.edu/~zickler/>

Ohad Ben-Shahar¹

www.cs.bgu.ac.il/~ben-shahar

¹ Department of Computer Science,
Ben Gurion University
Israel

² School of Engineering and Applied
Science, Harvard University
MA, USA

Abstract

Specular flow is an important class of optical flow whose utility in visual tasks has gained much interest in contemporary vision research. Unfortunately, however, reliably estimation of specular flow from image sequences is an open question that was never addressed formally before. In this paper we first argue that existing optical flow algorithms are incapable of reliable specular flow estimation due to their typical regularization criteria that conflict the unique and singular structure of specular flows. We show these discrepancies both qualitatively and using quantitative evaluation based on a first-of-its-kind benchmark dataset with ground truth specular flow data. We then suggest to generalize the popular optical flow variational framework using spatial weighting of different regularizers, and we propose new regularization terms that correspond better to the expected singularities of specular flows. Finally, we show how these contributions significantly improves specular flow estimation.

1 Introduction

Inference from visual scenes that contain highly specular objects is a particularly challenging computational problem. Most recent contributions to such problems have exploited the *specular flow* – the vector field that is induced on the image plane as a result of a relative motion between the camera, object, or environment – to facilitate diverse tasks such as shape inference [1, 2, 3, 4], 3D pose estimation [5] and detection of rigid objects [6]. However, while all these studies clearly demonstrate that specular flows can be exploited successfully (and theoretically rigorously) in various computational ways, they do not address the problem of how specular flows can be estimated from image data. At best, it was suggested to use *optical flow* algorithms on specular flow image sequences [7]. Indeed, optical flow is classically defined as ‘the apparent motion of brightness patterns over the image plane’ [8] and in this sense specular flows *are* optical flows. As we show here, however, both classical and contemporary optical flow algorithms fail to estimate specular flows accurately and reliably, thus render their subsequent utility impractical.

Understanding the unique properties of specular flow is key in identifying the reasons for the failure of optical flow algorithms and central to developing reliable specular flow estimators from images. While past studies of specular flows have not addressed their estimation from images [1, 3, 4] they do provide critical insights that are important for this problem as well. First, all of these studies demonstrate direct link between the behavior of moving

specular reflections and the curvature of the surface. For example, Adato et al. [10] derived the following closed form equation

$$\mathbf{u} = \frac{1}{2K} \cdot \begin{pmatrix} \frac{f_x f_{yy} - f_y f_{xy}}{2\|\nabla f\| \cdot (1 + \|\nabla f\|^2)} & -\frac{f_x f_{xy} + f_y f_{yy}}{(1 + \|\nabla f\|^2)^2} \\ \frac{f_x f_{xy} - f_y f_{xx}}{2\|\nabla f\| \cdot (1 + \|\nabla f\|^2)} & \frac{f_x f_{xx} + f_y f_{xy}}{(1 + \|\nabla f\|^2)^2} \end{pmatrix} \begin{pmatrix} \omega_\alpha \\ \omega_\beta \end{pmatrix} \quad (1)$$

in which $f(x, y)$ is the specular shape, $(\omega_\alpha, \omega_\beta)^T$ is a description of the relative object-environment motion, \mathbf{u} is the specular flow, and K is the Gaussian curvature of the surface. Clearly, the inverse relationship between K and the induced specular flow \mathbf{u} indicates that the magnitude of the specular flow grows dramatically as surface curvature decreases. Second, while specular flows exhibit *magnitude singularities* at (for orthographic observer [10]) or near (for perspective observer [11]) the projection of parabolic points, at these points they also exhibit *orientation singularities*. Formally, if $m(x, y)$ and $\theta(x, y)$ are the magnitude and orientation components of the specular flow at each point, respectively, $\alpha(s)$ is a parametric representation of the singularity curve and $\beta(t)$ is an image curve that intersects $\alpha(s)$ *non transversely* at time $t = 0$, then the one sided limits at that point of intersection satisfy

$$\lim_{t \rightarrow 0^+} m(s(t)) = \lim_{t \rightarrow 0^-} m(s(t)) = \infty \quad \text{and} \quad \lim_{t \rightarrow 0^+} \theta(s(t)) = \pi + \lim_{t \rightarrow 0^-} \theta(s(t)). \quad (2)$$

We call these coupled phenomena the *parabolic singularity*.

Singularities in traditional optical flows and specular flows are fundamentally different for one more reason. While the former always occur at surface boundaries and hence can be detected directly from the image by exploiting image gradients, specular flow singularities might also occur in the *interior* of surfaces (i.e., away from boundaries). Since the curvature events from which they emerge are not manifested in images as directly as boundaries, specular flow singularities are much harder to detect, trace, and handle.

In this paper we study specular flow estimation and propose computational foundations towards reliable and practical algorithms. We first analyze the weakness of existing optical flow algorithms vis-a-vis specular flow properties and place the blame for their failure on specular flow data on the widely used smoothness regularization term. We then generalize the popular optical flow framework to use classical regularizations when permissible, while employing other regularizations where smoothness is prohibited. To validate our results we establish a novel evaluation benchmark and a new database that contains image sequences of virtual and physical specular objects together with their corresponding specular flow ground truth. To our best knowledge, this is the first attempt to create such a dataset, which could supplement an important missing piece in existing optical flow benchmark sites [9].

2 Related work: Could optical flow algorithms do the job?

Perhaps the only work that incorporates some aspects of specular flow estimation is due to Roth and Black [12], whose reconstruction approach can also estimate the flow while assuming prior knowledge on the geometrical structure of the specular surfaces in the scene. To our best knowledge, no work has been done on the general problem of estimating specular flow directly from image sequences only. What *has* been researched extensively is the estimation of *optical* flow, which raises the possibility that these algorithms could address specular flow estimation as well (since specular flows are a class of optical flows). In the next section we will argue for the opposite.

The standard Middlebury optical flow benchmark site [9] clearly shows that optical flow algorithms have been improving progressively on traditional optical flow inputs. Interestingly, while the algorithmic methods are numerous, most share common foundations, and in particular, they follow the same seminal ideas from Horn and Schunck [13], such as the




Figure 1: The acquisition system and selected preliminary evaluation data. **(a)** High precision 3D printer (Connex 350, Objet Geometries Ltd., <http://www.objet.com>) used for fabrication of ground truth shapes. **(b-c)** Selected real objects before (up) and after (bottom) specular coating. **(d)** Sample ground truth specular flow separated to its orientation (up) and magnitude (bottom) components. **(e)** The specular flow acquisition device.

brightness constancy assumption and the global *piecewise smoothness* constraint which assists in coping with noise, sparse data, and the aperture problem [9]. Over the years, numerous studies have consistently improved the accuracy and robustness of optical flow algorithms by modifying Horn and Schunk’s energy functional, e.g., via robust penalty functions (e.g., [9]), by data terms endowed with a *gradient constancy* assumption [10], via data term normalization [14], or the incorporation of color information (e.g. [18, 24]), to name but a few. Numerous studies have also focused on improving the smoothness term, for example by incorporating motion segmentation [2], by modeling scene motion [10], by using spatio-temporal regularization [6, 7], or by combining anisotropic smoothing [16, 22, 24].

Given all these advances in optical flow algorithms, it is tempting to examine the degree to which they facilitate reliable *specular* flow estimation. However, we immediately observe that all existing algorithms are incompatible with fundamental properties of specular flows:

- All flavors of the smoothness term disagree with the fact that the specular flow magnitude can become very large or even unbounded in certain regions.
- Since existing smoothness terms are typically based on image gradient, they do not address well parabolic singularities, which are related to surface *curvature*.
- Specular flow is a distorted, shape-dependent reflection of physical scene motion and therefore may be modeled incorrectly by existing (e.g., linear) motion models.

Based on all of the above, it is reasonable to predict that existing state-of-the-art optical flow algorithms will perform poorly on specular flows. But to validate our predication, one should use ground truth data with which different algorithms can be compared. Unfortunately, as we discuss next, existing datasets cannot serve this purpose.

3 Specular flow benchmark and evaluation methodology

While the recently established Middlebury optical flow dataset [9] has become a standard *de facto* in optical flow evaluation, it is missing several important cases like specular and transparent scenes and cannot be used to evaluate specular flow estimation algorithms. Indeed, obtaining real image sequences with their corresponding specular flow ground truth is a very challenging task. Among other contributions, in this paper we present technical steps for establishing such a full scale benchmark database, and we present a preliminary version that includes several objects. While most of this discussion is included in the Supplementary Material, here we mention that establishing this database involves the creation of objects with ground truth shape using a state-of-the-art, high precision 3D printer (Fig. 1a), the acquisition of specular image sequences using a custom made device (Fig. 1e), and the computation

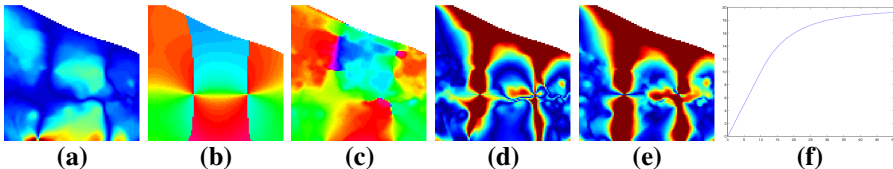


Figure 2: Two qualitatively different specular flows with similar AE rank (of $\sim 21^\circ$). **(a)** Magnitude map of both flows. **(b-c)** Orientation map of the first and the second flows, respectively. **(d-e)** Angular Error (AE) of the first and the second flows, respectively. Note how similar the error maps are, despite the qualitative differences in the orientation maps (panels b and c). **(f)** The error damping function from Eq. 3.

of ground truth specular flow data from the known shape using the generative equation of Adato et al. [1] (Eq. 1). The preliminary data sets created this way includes several objects which span a variety of geometrical configurations with varying degree of complexity (see Fig. 1). While a full-fledged online evaluation site is currently under construction, these preliminary data sets are already available publicly [19] and we hope they would promote much research on specular flow estimation.

In addition to these data sets, specular flow estimation calls for special attention with regard to evaluation methodology. As already discussed in the literature [4, 9], we too argue that the commonly used angular error (AE) and the end point (EP) measures are problematic in terms of evaluating the performance of flow estimation algorithms. In particular, two estimated flows can have similar AE or EP error level while in fact they are qualitatively different (see Fig. 2). Combined with the fact that vector fields are naturally represented by their magnitude and orientation components, which indeed carry distinct information and different significance in the context of motion description, we therefore suggest that performance evaluation of optical/specular flow algorithms should be based on *separate* evaluation of the orientation (O) and magnitude (M) components of the estimated flow, namely via an *Orientation Error* (OE) map and a *Magnitude Error* (ME) map. The OE and ME can be averaged to obtain the *Average Orientation Error* (AOE) and the *Average Magnitude Error* (AME), or a two dimensional error descriptor (*AOE, AME*). Although such an error description is not susceptible to total ordering, we argue that the advantages that such a description carries in analyzing the performance of algorithms greatly exceeds this limitation. Hence, throughout this paper we primarily use this proposed evaluation approach.

Since specular flow magnitude could grow unbounded, the raw ME map might include errors of infinite size. For practical reasons we therefore need to map the ME non linearly to a finite interval, for example via a stereographic projection or by using the following function:

$$m(e) \triangleq \begin{cases} e & e < \frac{1}{2}\chi \\ \frac{e^2}{\chi} \cdot \frac{1}{\rho(\chi) + (\frac{e}{\chi})^2} & e \geq \frac{1}{2}\chi \end{cases} \quad (3)$$

where χ is a bound on the practical magnitude allowed in the flow (see Sec. 5) and $\rho(\chi)$ is a parameter that ensures continuity (see Fig. 2f):

Using all of the above, we have evaluated various state-of-the-art optical flow algorithms on our novel specular flow data sets where the flows are induced by specular scenes/objects *with* and *without* parabolic lines. The results show that, even in the absence of parabolic lines, all tested techniques could not handle specular flows well. In particular, while the flow orientation was estimated relatively accurately, all existing algorithms fail to estimate the flow magnitude reliably, especially in regions where the flow changes rapidly (i.e., where the gradient of flow magnitude is relatively large, see Supplementary Material and also in Fig. 3). Empirical support for this observation is further obtained by evaluating the energy model of

each of the tested algorithms on the desired ground truth. In all cases the energy obtained this way is *significantly* higher compared to the energy of the estimated flow, indicating that the flow model employed by these methods is not suitable for specular flow.

Given these results, we continued to test existing algorithms on specular scenes containing objects *with* parabolic singularities. As predicted, the results of the existing algorithms not only include unacceptably large quantitative errors, but also global qualitative distortions that extend well beyond the parabolic singularities (Fig. 3 and Supplementary Material). Evidently, most algorithms either infer increasingly smaller magnitudes rather than larger ones or tend to estimate a smooth flow transition across the parabolic line instead of the expected 180° orientation discontinuity (see Fig. 5 for close ups of such regions).

We conclude that the existing optical flow algorithms estimate specular flows with major qualitative discrepancies and significant estimation errors compared to the ground truth, and therefore a new approaches are required for robust estimation of specular flows.

4 A generalized optical flow framework

To extend existing methods to handle specular flow we suggest to generalize the optical flow variational framework such that the advantages of existing regularizations is kept in most regions, while more appropriate regularization is employed where the flow violates the smoothness assumption. Assuming one can tell (exactly or approximately) the division of the image plane into regions of different properties, an optical flow model that switches (spatially) between the appropriate terms will have the following general form

$$E(u, v) = \int_{\Omega} \sum_i \eta_i(x, y) E_i(dI, u, v) dx dy \quad s.t. \quad \eta_i(x, y) \geq 0 \wedge \sum_i \eta_i(x, y) = 1 \quad (4)$$

where u, v is the sought-after flow, dI stands for the image derivatives (up to a desired order), E_i is a (data or regularization) term appropriate for regions of type i , and $\eta_i(x, y)$ is a spatial *confidence function* that describes the degree of compatibility of each point to this type of region. The normalization constraint $\sum_i \eta_i(x, y) = 1$ reflects a prior that region types are mutually exclusive (in the probabilistic sense), but if necessary it can be removed. Note that unlike in previous approaches, this new framework allows to reduce completely the influence of the data term in certain regions. Indeed, in this work we will use this capacity along the parabolic singularities. Note that if $\eta_i(x, y)$ cannot be known fully in advance, they too can participate in the optimization process, thereby facilitating the recovery of region division as a by-product.

Clearly, the above suggested formal framework generalizes existing optical flow algorithms (e.g., by setting $E_1 = E_{data}$, $E_2 = E_{smooth}$, $\eta_1(x, y) \equiv 1 - \alpha$ and $\eta_2(x, y) \equiv \alpha$, where $0 < \alpha < 1$ is the traditional regularization weight). In this paper, however, we would like to take advantage of its wider expressive power via the space-varying confidence functions in order to apply new and different regularizations in regions around parabolic singularities.

5 The parabolic singularity terms

As discussed in Sec. 1, at the parabolic singularity the specular flow exhibits both magnitude singularity and orientation discontinuity of 180° . Therefore, the flow structure at such points can be abstracted by the two conditions

$$\begin{aligned} (u^2 + v^2 - \chi^2)^2 &= 0 \\ I(x, y, t) &= I(x + u, y + v, t + 1) = I(x - u, y - v, t + 1) \end{aligned} \quad (5)$$

where χ is the maximum practical magnitude allowed in the estimated specular flow. Hence, combined with a robust penalty function ψ [■], we define the “parabolic term” regularizer as follows:

$$E_p \triangleq \Psi \left((I(x+u, y+v, t+1) - I(x-u, y-v, t+1))^2 \right) + \Psi((u^2 + v^2 - \chi^2)^2). \quad (6)$$

This term should apply to image points infinitesimally close to the parabolic singularity, i.e., generically, to a set of codimension 1. This set, represented via the corresponding confidence function, should be measured from image data, and later we discuss how one can obtain its initial gross estimation and then refine it during the optimization process.

One “disadvantage” of E_p is its aggressive encouragement of large-magnitude flow. Hence, by applying it to points too distant from the parabolic singularity (due to an inaccurate confidence function) we are still at risk of obtaining large scale distortions. Hence, we also define a weaker version of the parabolic term that incorporates the orientation structure only:

$$E_n \triangleq \Psi \left((I(x+u, y+v, t+1) - I(x-u, y-v, t+1))^2 \right). \quad (7)$$

This “neighborhood” term is constructive for a finite-size strip around the parabolic singularity since in practical terms it can be employed for all pixels whose distance from the singularity is smaller than their optical flow magnitude. This term has an added advantage – it is significantly easier to approximate its corresponding confidence function using standard computational tools.

Using these two terms, and following our suggested computational framework, Eq. 4, we now suggest to obtain more reliable specular flow estimation by minimizing the following energy functional

$$E(u, v) = \int (\eta_0(x, y)E_d + \eta_1(x, y)E_s + \eta_2(x, y)E_n + \eta_3(x, y)E_p) dx dy, \quad (8)$$

where E_d abbreviates E_{data} , E_s abbreviates E_{smooth} , $\eta_2(x, y)$ and $\eta_3(x, y)$ are the spatial functions describing the confidence that point (x, y) is *near* or *at* a parabolic singularity, respectively.

Having Eq. 8 as our objective function, one can derive the Euler-Lagrange equations as the necessary condition for its minimization:

$$\begin{aligned} \partial_u(\eta_0 E_d + \eta_2 E_n + \eta_3 E_p) - \eta_1(\partial_x(\partial_{u_x} E_s) + \partial_y(\partial_{u_y} E_s)) &= 0 \\ \partial_v(\eta_0 E_d + \eta_2 E_n + \eta_3 E_p) - \eta_1(\partial_x(\partial_{v_x} E_s) + \partial_y(\partial_{v_y} E_s)) &= 0. \end{aligned} \quad (9)$$

To express these terms more conveniently, we use first order Taylor approximation

$$I(x+u, y+v, t+1) - I(x-u, y-v, t+1) = 2I_x u + 2I_y v = 2\nabla I \cdot (u, v). \quad (10)$$

Note that Eq. 10, in conjunction with the desired minimization of E_n (Eq. 7) implies that the flow should be orthogonal to the image gradient around the parabolic singularities. Having all of the above, we can now derive a linearized version of $\partial_u E_n$ and $\partial_v E_n$.

$$\begin{aligned} \partial_u E_n &= \Psi'((2I_x u + 2I_y v)^2) \cdot 8(I_x^2 u + I_x I_y v) \\ \partial_v E_n &= \Psi'((2I_x u + 2I_y v)^2) \cdot 8(I_x I_y u + I_y^2 v), \end{aligned} \quad (11)$$

which can be used back in the Euler-Lagrange Eq. 9. Following the theory of calculus of variation we therefore conclude that a necessary condition for minimizing Eq. 8 is

$$\begin{aligned} \eta_0 \partial_u E_d &+ (\eta_2 + \eta_3) \Psi'((2I_x u + 2I_y v)^2) \cdot 8(I_x^2 u + I_x I_y v) \\ &+ \eta_3 u \Psi'((u^2 + v^2 - \chi^2)^2) - \eta_1(\partial_x(\partial_{u_x} E_s) + \partial_y(\partial_{u_y} E_s)) = 0 \\ \eta_0 \partial_v E_d &+ (\eta_2 + \eta_3) \Psi'((2I_x u + 2I_y v)^2) \cdot 8(I_x I_y u + I_y^2 v) \\ &+ \eta_3 v \Psi'((u^2 + v^2 - \chi^2)^2) - \eta_1(\partial_x(\partial_{v_x} E_s) + \partial_y(\partial_{v_y} E_s)) = 0 \end{aligned} \quad (12)$$

where one has the freedom to plug in any E_d and E_s of choice, or use the corresponding terms from any existing variational optical flow algorithm.

Clearly, to become constructive, the variational process discussed above requires the functions $\eta_2(x, y)$ and $\eta_3(x, y)$ which describe the confidence that point (x, y) is near or at a parabolic singularity. Obviously, unless one has complete knowledge of the specular scene, this prior knowledge must be extracted from the image sequence itself while applying proper heuristics. Here we propose to do so by considering the eigenvectors of the structure tensor, J_ρ , in the spatio-temporal domain, as also suggested earlier (e.g., [8, 13]).

Although local methods that exploit this tensor for optical flow estimation often fail to provide sufficient end results, here it will prove sufficient for obtaining the rough initial

estimation of the parabolic confidence function which is sufficient to bootstrap our specular flow estimation process. As mentioned above, the flow’s magnitude increases rapidly in the neighborhood of parabolic singularities. Thus, the temporal component of the smallest eigenvector (which implies the optical flow direction) must be relatively small compared to its spatial components. Here we will use this heuristics to initialize $\eta_2(x, y)$.

Unlike the neighborhood confidence function $\eta_2(x, y)$, the parabolic confidence function $\eta_3(x, y)$ should be estimated more carefully since its corresponding parabolic term ($E_{parabolic}$) encourages dramatic increase in the magnitude of the estimated flow, which could result in significant errors if applied away from parabolic singularities. Therefore, we suggest estimating the parabolic confident function as part of the minimization process by detecting the parabolic singularity directly from the calculated flow. In particular, $\eta_3(x, y)$ should be increased in regions where the covariant derivative of the estimated specular flow is large (and vice versa), an operation which can be executed at the end of each optimization iteration. Obviously, each such update should be followed by normalization of all $\eta_i(x, y)$ in order to preserve the condition $\sum_i \eta_i(x, y) = 1$.

6 Instantiation and numerical implementation

The algorithmic framework suggested in this paper is designed to generalize optical flow algorithms in order to extend their computational domain to include non traditional behaviors such as specular singularities. In this sense, one could plug in to our framework *any* variational optical flow algorithm (i.e., by setting the proper terms in Eq. 8) to end up with a new version of that algorithm that handles specular flows better. We emphasize that whatever data and smoothness terms are incorporated, the resultant computational scheme provides improved estimation capacity for specular flows *without* lost of performance on traditional optical flow sequences.

As a proof of concept, we implemented the proposed algorithm by using relative simple data and smoothness terms and show that even with this particular choice the obtained algorithm can estimate specular flow significantly better than state-of-the-art optical flow algorithms. More formally, to obtain an instant of our suggested framework we incorporated as data term a normalized version of the brightness constancy assumption [14, 24] with a robust non-quadratic penalizer function [9] and the most basic smoothness term proposals [13] with the same non-quadratic penalizer., i.e.,

$$\begin{aligned} E_d &= \psi \left(\rho_0 (I(x+u, y+v, t+1) - I(x, y, t))^2 \right) \\ E_s &= \psi (|\nabla u|^2 + |\nabla v|^2) \end{aligned} \quad (13)$$

where $\rho_0 = 1/(|\nabla I|^2 + \varepsilon^2)$ is the normalization factor and $\psi(s^2) \triangleq \sqrt{s^2 + \varepsilon^2}$. Incorporating both of these terms together with a normalized version of the parabolic singularity and the parabolic neighborhood terms (Eq. 6 and Eq. 7), we obtain the following energy functional

$$\begin{aligned} E(u, v) &= \int \psi(\rho_0(I(x+u, y+v, t+1) - I(x, y, t))^2) + \eta_1 \psi(|\nabla u|^2 + |\nabla v|^2) + \eta_3 \psi((u^2 + v^2 - \chi^2)^2) \\ &\quad + (\eta_2 + \eta_3) \psi(\rho_0(I(x+u, y+v, t+1) - I(x-u, y-v, t+1))^2) dx dy. \end{aligned} \quad (14)$$

With Eq. 14 as our flow model, our goal is finding a flow (u, v) that minimizes it, given a particular image sequence $I(x, y, t)$. To express the corresponding Euler-Lagrange equations we also employ the following notations (similar to Brox *et al.* [9])

$$\begin{aligned} I_* &\triangleq \partial_* I(x+u, y+v, t+1) & I_u &\triangleq I(x+u, y+v, t+1) - I(x, y, t) \\ I_{xu} &\triangleq \partial_x I(x+u, y+v, t+1) - \partial_x I(x, y, t) & I_{yu} &\triangleq \partial_y I(x+u, y+v, t+1) - \partial_y I(x, y, t) \end{aligned}$$

where $*$ \in (x, y, xx, xy, yy) . We use the index u to emphasize that the derivative is in the direction of the flow. With these notations, the necessary condition for the minimization of Eq. 14 becomes

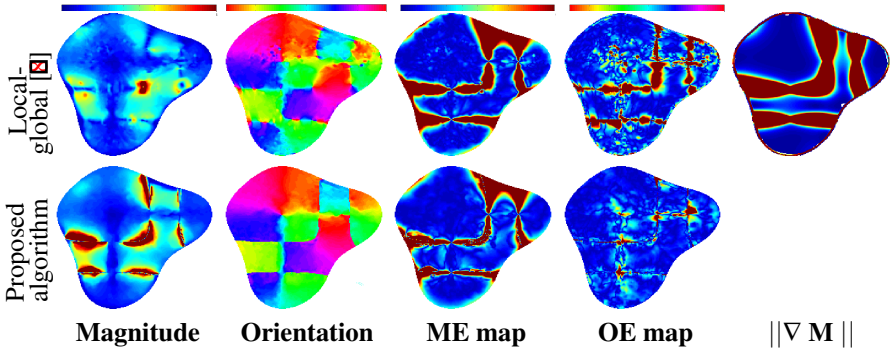


Figure 3: Performance of existing and the suggested algorithms on specular flow data. Compare result of the lower flow to the ground truth in Fig. 1. The $\|\nabla M\|$ column illustrates how previous algorithms tend to fail in their orientation estimation where the flow’s (ground truth) magnitude grows rapidly. Note the great improvement in our result. This is the case even if the flow contains no parabolic singularities (see Supplementary Material)

$$\begin{aligned}
 & \psi'(\rho_0(I_u)^2) \cdot \rho_0 I_x I_u - \eta_1 \cdot \text{div}(\psi'(|\nabla u|^2 + |\nabla v|^2) \cdot |\nabla u|) + \eta_3 \cdot u(\psi'((u^2 + v^2 - \chi^2)^2)) \\
 & \quad + (\eta_2 + \eta_3) \cdot \psi'(\rho_0(2I_x u + 2I_y v)^2) \cdot \rho_0(I_x^2 u + I_x I_y v) = 0 \\
 & \psi'(\rho_0(I_u)^2) \cdot \rho_0 I_y I_u - \eta_1 \cdot \text{div}(\psi'(|\nabla u|^2 + |\nabla v|^2) \cdot |\nabla v|) + \eta_3 \cdot v(\psi'((u^2 + v^2 - \chi^2)^2)) \\
 & \quad + (\eta_2 + \eta_3) \cdot \psi'(\rho_0(2I_x u + 2I_y v)^2) \cdot \rho_0(I_x I_y u + I_y^2 v) = 0 .
 \end{aligned} \tag{15}$$

Finally, the last non-linear equations can be handled by linearization while the effect of local minima is reduced by a coarse-to-fine strategy similar to Brox et al. [24] but with one important difference. For specular flows, the standard coarse-to-fine strategy is less effective since the interpolation required in the transition between scale levels of the pyramid is particularly harmful for parabolic singularities. The averaging of flow vectors of opposite orientations across the singularities typically results in mutual cancellation and the emergence of zero magnitude pattern in these regions. Hence, we use the covariant derivative of the estimated flow (in the direction of the flow) to locate these (pixel-wide) regions, and then set the flow direction to the direction normal the image gradient at that point, as suggested by Eq. 10.

7 Experimental results and future work

Using the proposed benchmark and the evaluation methodology, we compared the proposed approach to many classical [13, 15] and contemporary [7, 8, 18, 22, 24] optical flow algorithms. In all cases we used gray level images and running parameters were optimized manually for the most accurate results. Results with our approach were computed twice, first by estimating the spatial confidence function η_2 from the ground truth flow and then by estimating it directly from the given image sequences (as described in Sec. 5).

As shown in Figs. 3-4 and Table 1, existing optical flow algorithms fail to estimate well the specular flow, all exhibiting common types of errors. For space consideration, we show only the method that was the most accurate for any given image sequence (see Supplementary Material for more results). At the same time, and despite using only *basic* data and smoothness terms, our algorithm provides superior results on the same input, estimating the flow’s structure with better accuracy, both qualitatively and quantitatively.

Obviously, while here we outline an approach for better specular flow estimation, much could be done for further improvements. For example, using data and smoothness terms from contemporary optical flow algorithms (as opposed to the classical terms we used here), could

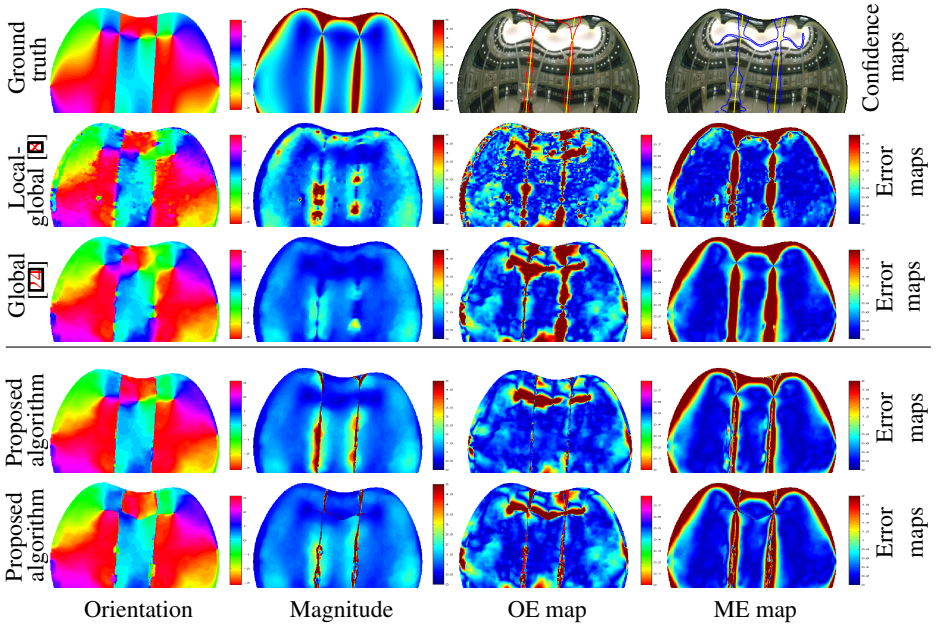


Figure 4: Comparison of the performance of two optical flow algorithms [8, 24] (top) to our algorithm (bottom) on real specular flow sequence. Confidence maps (top right) show ground truth parabolic singularities (yellow) and estimate regions of η_2 , when computed from ground truth flow (left/red) and from real images (right/blue).

clearly improve results. Similarly, a rigorous formulation of a time-evolving η_i is likely to enhance the framework as a whole. Finally, we mention again that part of our contribution is the first-ever formulation of a ground truth specular flow data set with which specular flow estimation algorithms can be evaluated. Although these data are already available to researchers [19], a full fledged online evaluation site would be an important research tool and is currently under construction.

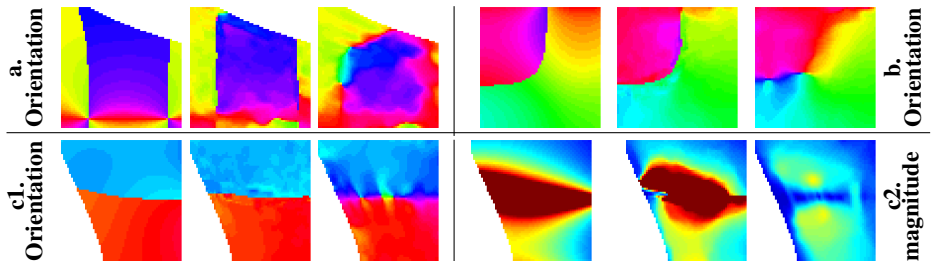


Figure 5: A close up of three regions of interest (ROIs a,b, and c) of an estimated specular flow. All four panels show ground truth on the left, result of our proposed algorithm in the center. The results by existing state-of-the-art algorithm [8, 24] on the right exhibit typical qualitative failures: large regions where the OE is significant (a), smooth transition instead a fixed orientation jump at the singularity (b), and inference of smaller magnitude instead of large one around parabolic singularities (c).

	a. Synthetic image sequence						b. Real image sequence						c. Real image sequence					
	AOE			AME			AOE			AME			AOE			AME		
	E	P	R	E	P	R	E	P	R	E	P	R	E	P	R	E	P	R
Local-Global [8]	17°	69°	13°	1.3	7.9	0.86	18°	42°	14°	1.0	4.5	0.6	18°	42°	14°	1.0	4.5	0.6
Global [24]	16°	74°	14°	1.4	8.1	0.98	17°	51°	13°	1.2	5.6	0.7	17°	51°	13°	1.2	5.6	0.7
Proposed	8°	33°	7°	0.8	4.0	0.54	11°	33°	10°	0.9	3.8	0.7	13°	34°	11°	1.0	4.0	0.6

Table 1: Quantitative comparative results of AOE and AME for one synthetic (a) and one real (b,c) specular sequences. Panels b and c correspond to the two lower rows from Fig. 4, which rely on neighborhood confidence functions calculated from the ground truth and the real image pair, respectively. AOE and AME are detailed for the entire image (E), the (challenging) parabolic singularity regions (P), and the complementary, regular part of the flow (R). Note how our approach outperforms the competition in virtually all cases.

Acknowledgements

This project is funded by the Israel Science Foundation under grant No. 1245/08 and the US National Science Foundation under grant IIS-0712956. O.B.S and Y.A. thank the generous support of the Frankel Fund and the Paul Ivanier Robotics Center at Ben-Gurion University.

References

- [1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Toward a theory of shape from specular flow. In *Proceedings of the IEEE International Conference on Computer Vision*, 2007.
- [2] T. Amiaz and N. Kiryati. Piecewise-smooth dense optical flow via level sets. *International Journal of Computer Vision*, 68(2):111–124, 2006.
- [3] S. Baker, D. Scharstein, J.P. Lewis, S. Roth, M. Black, and R. Szeliski. A database and evaluation methodology for optical flow. In *Proceedings of the IEEE International Conference on Computer Vision*, 2007.
- [4] J.L. Barron, D.J. Fleet, and S.S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, 1994.
- [5] J. Bigun, G.H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, 1991.
- [6] M.J. Black and P. Anandan. Robust dynamic motion estimation over time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 296–302, 1991.
- [7] T. Brox, A. Bruhn, N. Papenbergh, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *Proceedings of the European Conference on Computer Vision*, 2004.
- [8] A. Bruhn, J. Weickert, and C. Schnorr. Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 61(3):211–231, 2005.
- [9] G.D. Canas, Y. Vasilyev, Y. Adato, T. Zickler, S. Gortler, and O. Ben-Shahar. A linear formulation of shape from specular flow. In *Proceedings of the IEEE International Conference on Computer Vision*, 2009.
- [10] J.Y. Chang, R. Raskar, and A. Agrawal. 3d pose estimation and segmentation using specular cues. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1706–1713, 2009.

- [11] G. Farneback. Very high accuracy velocity estimation using orientation tensors, parametric motion, and simultaneous segmentation of the motion field. In *Proceedings of the IEEE International Conference on Computer Vision*, 2001.
- [12] B.K.P. Horn, editor. *Robot vision*. MIT Press, Cambridge, MA, 1986.
- [13] B.K.P. Horn and B.G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [14] S.H. Lai and B.C. Vemuri. Reliable and efficient computation of optical flow. *International Journal of Computer Vision*, 29(2):87–105, 1998.
- [15] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the Image Understanding Workshop*, pages 121–130, 1981.
- [16] H.H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
- [17] S. Roth and M.J. Black. Specular flow and the recovery of surface structure. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1869–1876, 2006.
- [18] P. Sand and S. Teller. Particle video: Long -range motion estimation using point trajectories. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2195–2202, 2006.
- [19] Specular flow benchmark. <http://www.cs.bgu.ac.il/~vision>.
- [20] Y. Vasilyev, Y. Adato, T. Zicker, and O. Ben-Shahar. Dense specular shape from multiple specular flows. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [21] S. Waldon and C.R. Dyer. Dynamic shading, motion parallax and qualitative shape. In *Proc. IEEE Workshop on Qualitative Vision*, pages 61–70, New York City, NY, USA, 1993.
- [22] Manuel Werlberger, Werner Trobin, Thomas Pock, Andreas Wedel, Daniel Cremers, and Horst Bischof. Anisotropic Huber-L1 optical flow. In *Proceedings of the British Machine Vision Conference*, 2009.
- [23] D. Zang, K. Doerschner, and P R. Schrater. Rapid inference of object rigidity and reflectance using optic flow. In *Proceedings of the International Conference on Computer Analysis of Images and Patterns*, pages 881–888, 2009.
- [24] H. Zimmer, A. Bruhn, J. Weickert, L. Valgaerts, A. Salgado, B. Rosenhahn, and H.P. Seidel. Complementary optic flow. In *EMMCVPR*, pages 207–220, 2009.