Motion Estimation & Optical Flow

Motion Estimation

■ The estimation of motion fields from image sequences is among the central problems in computer vision.

■ With increasing amount of image sequence data – more and more video-capable cameras, higher frame rates, videos on the internet – image sequence analysis is becoming increasingly important.

■ Compared to still images, video contains an enormous amount of information about the world surrounding us in the sense that structures can often be distinguished based on their temporal evolution.

■ Some applications of motion estimation are already integrated in camera software – panorama generation from several images, video stabilization to remove camera shake, etc.

■ Mathematically the problem of motion estimation from images is an ill-posed problem, which means that the question is not sufficiently specified to assure a unique solution.

The Correspondence Problem

Algorithmically, the key challenge in motion estimation is to solve the correspondence problem. Given two images, determine for each point in either image the corresponding partner in the other image. Many computer vision problems are inherently such correspondence problems:

■ Disparity estimation from stereo images: Determine a one-dimensional displacement for each pixel to determine the corresponding pixel in the other image. This displacement is inversely proportional to the depth of the respective point.

■ Multimodal registration: Given two medical images of an organ acquired with different sensors – for example CT (Computer Tomography) and MRI (Magnetic Resonance Imaging), or CT and PET (Positron-Emission Tomography) – compute an optimal alignment of these images.

■ Shape Matching: Given two object shapes (contours in 2D or surfaces in 3D), determine a correspondence between pairs of points from either shape.

Motion and Grouping
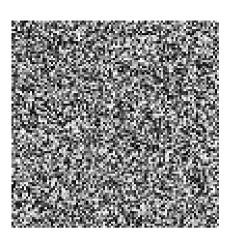
Regions of random brightness values moving with respect to oneanother.

Wallpaper regions moving with respect to oneanother.

Motion and Grouping
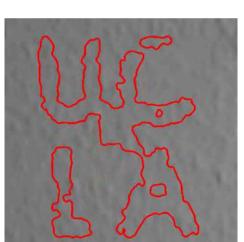
Automatic segmentation of the moving regions.

D. Cremers, A. L. Yuille, DAGM 2003

Motion and 3D Structure

Several images of a static scene filmed by a moving camera. Foreground objects move faster than background objects.

Motion and 3D Structure
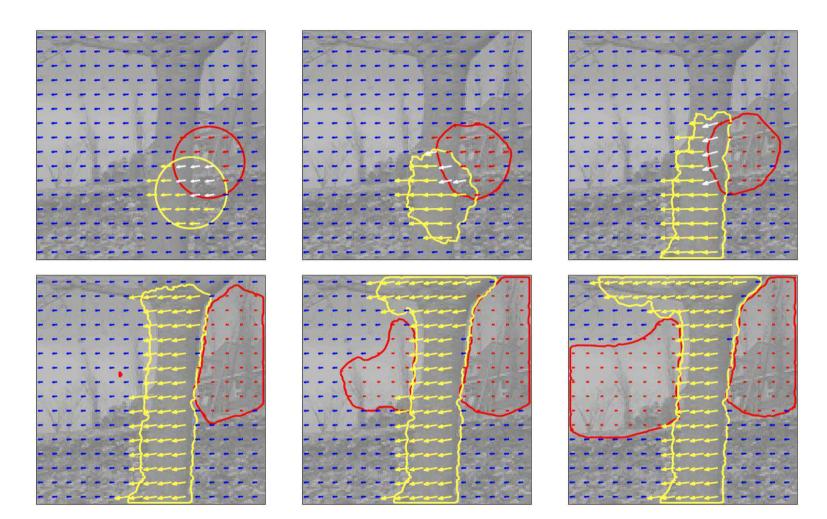
Motion-based segmentation into depth layers.

D. Cremers, S. Soatto, Int. Conf. on Computer Vision 2003.

Motion and 3D Structure

T. Schoenemann & D. Cremers,

Near Real-time Motion Segmentation, DAGM 2006.

Motion and Transparency

Image sequence showing semi-transparent superposition of two
motions.

Author: Michael Black

Applications of Motion Estimation

■ **Grouping and Segmentation**: Motion information allows to identify image regions as objects. This can also be done if semi-transparent structures overlap at a given location.

■ **Tracking**: Using motion information, objects can be tracked in a video sequence.

■ **Depth estimation**: Motion information allows to infer the distance of respective objects from the camera. In principle, one can recover the 3D geometry of the world from an image sequence.

■ **Time-to-Impact**: In the context of driver assistance, motion information allows to make predictions when an obstacle will be hit. As a consequence, one can initiate evasion maneuvers or breaking.

■ **Video compression**: Motion information allows to efficiently compress videos (MPEG encoding).

The Aperture Problem

In general, one cannot estimate motion in direction of constant brightness (for example along an image edge). This limitation is referred to as the aperture problem. For example: No matter how the horizontal stripe pattern behind the mask is displaced, we will only observe its vertical motion.

The Aperture Problem: Measurability

Region 1:   no motion

Region 2:   horizontal motion only

Region 3:   motion in both directions

The Brightness Constancy Assumption

■ Given an image sequence $I : \Omega \times [0, T] \to \mathbb{R}$, on the image plane $\Omega \subset \mathbb{R}^2$ and the time interval $[0, T]$, we wish to compute a motion field $v : \Omega \times [0, T] \to \mathbb{R}^2$, which assigns to each point $x \in \Omega$ at each time $t \in [0, T]$ a motion vector $v(x, t)$.

■ Let

$$x : [0, T] \to \Omega$$

denote the trajectory of an object point over time. The classical assumption in motion estimation state that the brightness of a moving point remains constant over time:

$$I(x(t), t) = \text{const.} \quad \forall t \in [0, T]$$

■ Assuming the brightness function to be differentiable, we can deduce that the total time derivative must vanish:

$$\frac{d}{dt} I(x(t), t) = \nabla I(x(t), t) \frac{dx(t)}{dt} + \frac{\partial I(x(t), t)}{\partial t} = 0 \quad \forall t \in [0, T]$$

The Optic Flow Constraint

■ The term $v(x, t) = \frac{dx}{dt}$ is nothing but the velocity of the moving point that we are looking for. Thus the assumptions of brightness constancy and differntiability lead to a relation between the desired velocity field $v(x, t)$ and the spatial and temporal image gradients:

$$\boxed{\nabla I^\top v + I_t = 0}$$

This equation is referred to as the differential brigthness constancy constraint or the optic flow constraint.

■ The optic flow constraint reflects the previously discussed aperture problem: It does not allow statements regarding motion along the level lines of constant intensity. More specifically, let $\tilde{v} = v + \eta$ be a modified motion field with η an arbitrary vector field normal to the image gradient ∇I. The $\tilde{v}$ also fulfills the optic flow constraint:

$$\nabla I^\top \tilde{v} + I_t = \nabla I^\top (v + \eta) + I_t = \nabla I^\top v + I_t = 0.$$

The OFC and the Aperture Problem

■ The aperture problem is reflected in the optic flow constraint because the constraint is invariant to changes in the motion field which are orthogonal to the local image gradient.

■ The central problem in motion estimation lies in the fact that the constraint coupling the velocity field $v(x, t)$ and the image gradients cannot be directly solved for v.

■ More specifically, the flow constraint provides the projection $v_\perp$ of the velocity vector v onto the image gradient ∇I. Dividing the OFC by $|\nabla I|$ leads to:

$$v_\perp \equiv \left(\frac{\nabla I}{|\nabla I|} \right)^\top v = -\frac{I_t}{|\nabla I|}$$

This component of the velocity normal to the level lines is called the normal flow. It is simply given by the negative ratio of temporal and spatial image gradient.

Example: Traffic Scene

$$I_1(x)$$

$$I_2(x)$$

$$I_t \approx |I_1(x) - I_2(x)|$$

$$|v_n| = \frac{|I_t|}{|\nabla I|} \approx \frac{|I_2 - I_1|}{\left|\nabla \frac{I_1 + I_2}{2}\right|}$$

Author: Daniel Cremers

Additional Assumptions

■ The optic flow constraint is necessary but not sufficient to uniquely determine a motion field. It only specifieds the normal component of the velocity field.

■ In order to eliminate the additional degree of freedom, we therefore need to make additional assumptions.

■ Two pioneering approaches:

◆ *Lucas and Kanade 1981*: Assume that the velocity in an entire window around each point is constant. If the window is "sufficiently" large one obtains a unique solution. (around 5300 citations in Jan 2012).

◆ *Horn and Schunck 1981*: A variational approach to motion estimation based on the assumption of spatial smoothness of the the flow field $v(x,t)$. Extensions to temporal smoothness are straight-forward. (around 6800 citations in Jan 2012). This paper is often considered the first variational method in computer vision.

Lucas and Kanade

For each point $(x, y) \in \Omega \subset \mathbb{R}^2$ and time $t \in [0, T]$, *Lucas und Kanade (1981)* separately determine a motion vector $v(x, y, t)$ by assuming that the motion field is constant in a certain neighborhood $U_\sigma(x, y) \subset \Omega$ around this point. The motion vector $v = (v_1, v_2)$ is determined in a least squares manner by minimizing the energy:

$$E(v) = \int_{U_\sigma(x,y)} \left(\nabla I^\top v + I_t \right)^2 dx'\, dy' = \int_{U_\sigma(x,y)} \left(I_x v_1 + I_y v_2 + I_t \right)^2 dx'\, dy'$$

The necessary condition for optimality is that the partial derivatives of this energy with respect to the two parameters v_1 and v_2 must vanish:

$$\frac{\partial E(v)}{\partial v_1} = \int_{U_\sigma(x,y)} I_x \left(I_x v_1 + I_y v_2 + I_t \right) dx'\, dy' \stackrel{!}{=} 0$$

$$\frac{\partial E(v)}{\partial v_2} = \int_{U_\sigma(x,y)} I_y \left(I_x v_1 + I_y v_2 + I_t \right) dx'\, dy' \stackrel{!}{=} 0$$

Lucas and Kanade

$$\frac{\partial E(v)}{\partial v_1} = \int_{U_\sigma(x,y)} I_x \left(I_x v_1 + I_y v_2 + I_t \right) dx' \, dy' \overset{!}{=} 0$$

$$\frac{\partial E(v)}{\partial v_2} = \int_{U_\sigma(x,y)} I_y \left(I_x v_1 + I_y v_2 + I_t \right) dx' \, dy' \overset{!}{=} 0$$

Since v_1 and v_2 are assumed constant over $U_\sigma(x, y)$ we can extract them from the integral and obtain a linear equation system of the form:

$$M v = b \qquad \Rightarrow \ v = M^{-1} b$$

mit

$$M = \int_{U_\sigma(x,y)} \nabla I \, \nabla I^\top \, dx' \, dy', \qquad \text{und} \ b = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = - \int_{U_\sigma(x,y)} \nabla I \, I_t \, dx' \, dy'$$

$\rightarrow$ For each point (x, y) determine v by inversion of a 2×2-matrix.

Lucas and Kanade

In the original approach of Lucas and Kanade all points of the window $U_\sigma(x, y)$ are treated equally. In practice, it is preferable to give more weight to the central pixels. The corresponding cost function is then:

$$E(v) = \int_\Omega G_\sigma(x - x')\left(\nabla I^\top v + I_t\right)^2 dx' \, dy' = G_\sigma * \left(\nabla I^\top v + I_t\right)^2,$$

where the squared optic flow constraint is weighted by some function G_σ (for example a Gaussian kernel). The corresponding linear equation system is given by

$$M_\sigma v = b_\sigma \qquad \text{where}$$

$$M_\sigma = G_\sigma * \left(\nabla I \, \nabla I^\top\right) = G_\sigma * \begin{pmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{pmatrix}, \quad \text{and} \quad b_\sigma = -G_\sigma * \left(\nabla I \, I_t\right).$$

The matrix M_σ is called structure tensor.

Lucas and Kanade: Solutions?

The goal of Lucas and Kanade was to dermine a unique velocity vector under the assumption of local constancy of the velocity. Depending on the local intensity structure there are three possible cases (see slide 12):

1. The brightness is entirely constant over U_σ, then the gradient ∇I is zero in the neighborhood, the matrix M is 0 and no velocity can be estimated. (Test: $\mathrm{trace}(M) < \epsilon$?)

2. All image gradients in the neighborhood U_σ are colinear. Then $\mathrm{rank}(M) = \mathrm{rank}(\nabla I \nabla I^\top) = 1$. The matrix M has only one non-zero eigenvalue. It is not invertible, but one can determine the normal flow: $v_n = -I_t/|\nabla I|$. (Test: $\det(M) < \epsilon$?)

3. The gradient ∇I in the window U_σ takes on multiple directions. Then we have $\mathrm{rang}(M) = 2$ or. $\det(M) \neq 0$ and we can determine the velocity vector v by matrix inversion.

Lucas/Kanade: Example

One of two images   Color-coded flow field

Color hue encodes direction, color brightness encodes magnitude

Author: Thomas Brox

Horn and Schunck

The approach of *Horn and Schunck (1981)* is considered the first variational approach in computer vision (cf. Snakes: 1988, Mumford-Shah: 1989). In addition to the optic flow constraint for each point, one assumes spatial smoothness of the velocity field $v(x)$:

$$E(v) = \int_\Omega \left(\nabla I^\top v + I_t \right)^2 dx\, dy \; + \; \lambda \int_\Omega |\nabla v(x)|^2 \, dx\, dy.$$

Increasing smoothness of the flow field can be imposed by increasing the weight $\lambda > 0$ of the regularizer. In contrast to standard notation, ∇v does not refer to not refer to the divergence of the flow field but to the gradients in each component:

$$|\nabla v(x)|^2 \equiv |\nabla v_1(x)|^2 + |\nabla v_2(x)|^2$$

In contrast to Lucas and Kanade, the approach of Horn and Schunck gives rise to a spatially dense flow field.

Euler-Lagrange Equations

Let $v = (v_1, v_2)$ be the flow field with components v_1 and v_2 in x- and y-direction. The minimizer of the Horn and Schunck functional

$$E(v) = \frac{1}{2} \int_\Omega \left(I_x v_1 + I_y v_2 + I_t\right)^2 dx\, dy + \frac{\lambda}{2} \int_\Omega |\nabla v_1(x)|^2 + |\nabla v_2(x)|^2 \, dx\, dy.$$

must fulfill the Euler-Lagrange equations:

$$\begin{cases} \dfrac{\partial E}{\partial v_1} = I_x \left(I_x v_1 + I_y v_2 + I_t\right) - \lambda \Delta v_1 = 0 \\[3mm] \dfrac{\partial E}{\partial v_2} = I_y \left(I_x v_1 + I_y v_2 + I_t\right) - \lambda \Delta v_2 = 0 \end{cases}$$

These equations are linear and can be solved with a Gauss-Seidel or Jacobi solver. The regularizer imposes smoothness of the computed flow field. It generates a fill in effect: Components of the velocity field which are not affected by the optic flow constraint are simply adopted from neighboring regions.

Horn/Schunck: Examples

One of two images

Color-coded flow field

Color encodes direction and magnitude

Author: Thomas Brox

Lucas/Kanade vs. Horn/Schunck

Advantages of Lucas/Kanade:

- Fast and simple computation,

- often acceptable and robust results,

Advantages of Horn/Schunck:

- dense flow fields,

- more general: allows for non-translational motion such as rotation,

- strict convexity assures unique solution,

- global fill-in effect, smoothness can be regulated by the parameter λ.

- further extensions: discontinuous flow fields, segmentation,...

Limitations of Both Approaches

- **Small motion assumption**: The optic flow constraint only holds infinitesimally and thus only applies to small velocity. In general brightness constancy implies:

$$I_1(x) = I_2(x + v(x)).$$

  Linearization (under the assumption that v is small) leads to the optic flow constraint. For larger motions it is no longer valid.

- **Brightness constancy**: The assumption of brightness constancy is not always valid: Light reflexes on shiny materials, multimodal image registration (where modalities like CT and PET assign different brightness values to the same structure), lighting variations over time, automatic gain control in the camera, etc.

- The approach of Horn and Schunck tends to oversmooth flow fields. In particular, it does not allow for disontinuities in the flow field.

- The above approaches are formulated for two images. In general we have sequences with many consecutive images.

Further Advances

Since the pioneering works of Lucas/Kanade and Horn/Schunck a multitude of publications on optic flow estimation have appeared. A paper which integrates a number of advances is *Brox et al., ECCV 2004*:

■ Discontinuity-preserving smoothness:

$$\int |\nabla v|^2 dx \quad \rightarrow \quad \int |\nabla v| dx$$

■ Coarse-to-fine warping scheme to allow for larger motion:

$$|\nabla I^\top v + I_t|^2 \quad \rightarrow \quad |I_1(x) - I_2(x + v)|^2$$

■ Robust non-quadratic data terms to allow for outliers:

$$|I_1(x) - I_2(x + v)|^2 \quad \rightarrow \quad |I_1(x) - I_2(x + v)|$$

■ Gradient constancy to account for global brightness changes.

Further Advances

Over the years, the Horn and Schunck approach was modified to the form:

$$E(\boldsymbol{v}) = E_{data}(\boldsymbol{v}) + \alpha\, E_{smooth}(\boldsymbol{v}),$$

mit

$$E_{data}(\boldsymbol{v}) = \int \psi\Big(\underbrace{\big|I(\boldsymbol{x}+\boldsymbol{v}) - I(\boldsymbol{x})\big|^2}_{\text{brightness constancy}} + \gamma\, \underbrace{\big|\nabla I(\boldsymbol{x}+\boldsymbol{v}) - \nabla I(\boldsymbol{x})\big|^2}_{\text{gradient constancy}} \Big)\, d\boldsymbol{x},$$

and

$$E_{smooth}(\boldsymbol{v}) = \int \psi\left(|\nabla_3 u|^2 + |\nabla_3 w|^2 \right)\, d\boldsymbol{x},$$

where

$$\boldsymbol{x} \equiv (x, y, t), \quad \boldsymbol{v} \equiv (u, w, 1), \quad \text{and } \nabla_3 \equiv (\partial_x, \partial_y, \partial_t),$$

and

$$\psi(s^2) = \sqrt{s^2 + \epsilon}.$$
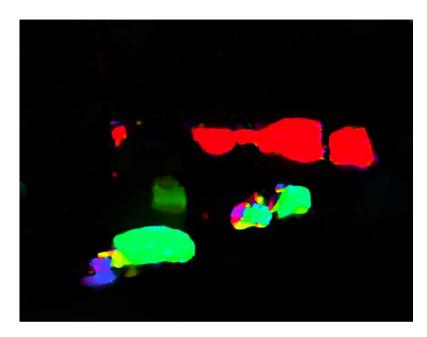
Discontinuity-preserving Flow Fields

One of two images

Color coded flow field

Color encodes motion direction and magnitude

Author: Thomas Brox

Experimental Comparison

first of two images

Lucas & Kanade '81

Horn & Schunck '81

Brox et al. '04

Author: Thomas Brox

Further Advances

In recent years, due to the Middlebury optical flow benchmark the topic of optical flow estimation has gained renewed interest. Three further improvements contained in the paper *Wedel et al., "Adaptive Regularization...", ICCV 2009* are:

■ Quadratic relaxation to decouple data term and regularizer:

$$\min_{v} \int \left| I(x + v) - I(x) \right| + |\nabla v| dx$$

$$\rightarrow \min_{v,u} \int \left| I(x + v) - I(x) \right| + |\nabla u| + \frac{1}{2\theta} |u - v|^2 \, dx$$

■ Data-dependent regularization which favors flow edges to coincide with image edges:

$$\int |\nabla v| dx \quad \rightarrow \quad \int_{\Omega} \exp\left(-\alpha |\nabla I_\sigma|^\alpha\right) |\nabla v| \, d^2 x$$

■ Rigid body regularization to impose rigid body motion rather than smoothness.

Optical flow evaluation results
Statistics: Average SD R2.5 R5.0 R10.0 A50 A75 A95
Error type: angle end-point interpolation normalized interpolation
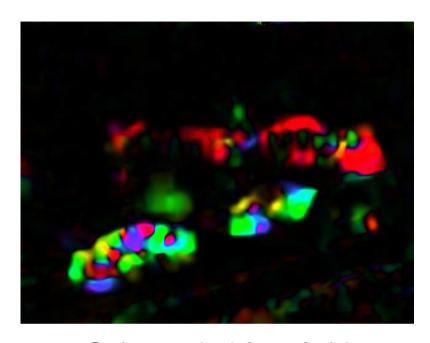
Motion Estimation & Optical Flow

- Motion Estimation
- The Correspondence Problem
- Motion and Grouping
- Motion and Grouping
- Motion and 3D Structure
- Motion and 3D Structure
- Motion and 3D Structure
- Motion and Transparency
- Applications of Motion Estimation
- The Aperture Problem
- The Aperture Problem: Measurability
- The Brightness Constancy Assumption
- The Optic Flow Constraint
- The OFC and the Aperture Problem
- Example: Traffic Scene
- Additional Assumptions
- Lucas and Kanade
- Lucas and Kanade
- Lucas and Kanade
- Lucas and Kanade: Solutions?
- Lucas/Kanade: Example
- Horn and Schunck
- Euler-Lagrange Equations
- Horn/Schunck: Examples
- Lucas/Kanade vs. Horn/Schunck

Average angle error	avg. rank	Army (Hidden texture) all	disc	untext	Mequon (Hidden texture) all	disc	untext	Schefflera (Hidden texture) all	disc	untext	Wooden (Hidden texture) all	disc	untext	Grove (Synthetic) all	disc	untext	Urban (Synthetic) all	disc	untext	Yosemite (Synthetic) all	disc	untext	Teddy (Stereo) all	disc	untext
TV-L1-improved [18]	3.8	3.36	9.63	2.62	2.82	10.7	2.23	6.50	15.8	2.73	3.80	21.3	1.76	3.34	4.38	2.39	5.97	18.1	5.67	3.57	4.92	3.43	4.01	9.84	3.44
F-TV-L1 [17]	5.1	5.44	12.5	5.69	5.46	15.0	4.03	7.48	16.3	3.42	5.08	23.3	2.81	3.42	4.34	3.03	4.05	15.1	3.18	2.43	3.92	1.87	3.90	9.35	2.61
Brox et al. [7]	6.0	4.80	14.4	4.29	4.05	13.5	3.71	6.63	16.0	7.26	5.22	22.7	3.22	4.56	6.09	3.40	3.97	17.9	3.41	2.07	3.76	1.18	5.14	11.9	4.28
Fusion [8]	6.4	4.43	13.7	4.08	2.47	8.91	2.24	3.70	9.68	3.12	3.68	19.8	2.54	4.26	5.16	4.31	6.32	16.8	6.15	4.55	5.78	3.10	7.12	13.6	7.86
SegOF [12]	6.9	5.85	13.5	3.98	7.40	14.9	8.13	8.55	17.3	9.01	6.50	18.1	5.14	3.90	4.53	4.81	6.57	21.7	6.81	1.65	3.49	1.08	3.71	9.23	3.63
Dynamic MRF [9]	7.0	4.58	12.4	4.14	3.25	13.9	2.27	6.02	16.8	2.36	4.39	22.6	2.51	3.61	4.55	3.46	6.81	22.2	6.78	2.41	3.48	3.69	9.26	17.8	10.2
CBF [14]	7.1	3.95	10.1	3.44	3.70	10.6	3.85	5.64	13.5	3.34	3.71	21.5	1.99	4.36	5.50	3.55	11.3	19.1	9.05	6.79	7.37	11.6	5.50	11.8	5.66
Learning Flow [13]	8.1	4.23	11.7	3.41	4.16	15.3	3.42	6.78	16.9	3.83	6.41	25.3	4.25	4.66	6.01	4.00	6.33	20.7	5.30	3.09	4.84	2.91	7.08	15.0	5.27
GraphCuts [16]	8.1	6.25	14.3	5.53	8.60	20.1	6.61	7.91	15.4	10.9	4.88	19.0	3.05	3.78	4.71	3.94	8.74	16.4	5.39	4.04	4.87	4.85	6.35	12.2	6.05
Second-order prior [10]	8.2	3.84	11.2	3.11	3.12	12.9	2.17	6.96	17.2	2.83	3.84	20.5	2.09	4.83	5.83	3.90	14.0	21.8	8.28	7.74	6.88	11.7	6.74	13.4	5.80
SPSA-learn [15]	9.5	6.84	16.7	6.74	8.47	19.4	7.49	12.5	23.1	13.1	8.40	25.8	7.08	3.87	4.66	4.10	6.32	18.8	6.89	2.56	3.85	1.79	7.29	12.5	7.47
2D-CLG [3]	10.5	10.1	22.6	7.59	9.84	16.9	11.1	16.9	28.2	18.8	14.1	31.1	13.1	3.86	4.62	4.53	5.98	21.2	5.97	1.76	3.14	1.46	6.29	12.9	5.81
GroupFlow [11]	10.8	8.00	18.6	8.09	11.1	23.7	10.3	12.6	25.6	12.8	5.84	20.3	4.39	4.69	5.81	3.67	9.29	22.4	10.1	2.11	3.99	2.29	5.75	10.0	7.39
Black & Anandan 2 [2]	11.3	7.83	18.7	6.41	9.70	21.9	8.60	13.7	23.7	18.1	10.9	30.0	9.44	4.60	5.55	5.06	7.85	17.6	6.38	2.61	4.44	2.15	8.58	14.3	8.54
Horn & Schunck [6]	13.1	8.01	19.9	8.38	9.13	23.2	7.71	14.2	25.9	14.6	12.4	30.6	11.3	4.64	5.64	4.60	8.21	24.4	8.45	4.01	5.41	1.95	9.16	17.5	8.86
Black & Anandan [1]	14.4	9.32	19.4	10.0	13.5	22.5	14.3	17.2	27.4	18.9	14.0	32.0	12.9	5.89	6.74	8.03	8.99	17.9	8.77	3.10	4.88	3.96	13.2	18.9	15.2
Pyramid LK [4]	17.2	13.9	20.9	21.4	24.1	23.1	30.2	20.9	29.5	21.9	22.2	34.6	25.0	18.7	23.1	20.2	21.2	24.5	21.0	6.41	7.02	10.8	25.6	31.5	34.5
MediaPlayer™ [5]	17.5	18.3	30.8	15.0	17.7	29.2	17.4	19.9	32.7	21.6	26.3	45.9	25.9	7.33	7.33	10.0	19.0	31.4	19.1	12.7	18.7	17.2	17.4	22.9	20.7

Move the mouse over the numbers in the table to see the corresponding images. Click to compare with the ground truth.
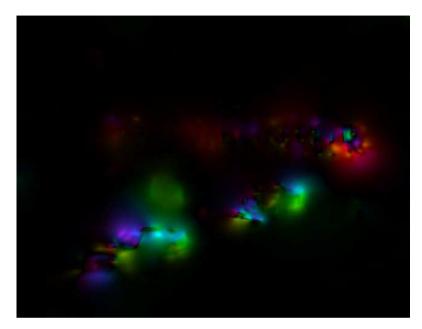
The Middlebury Benchmark (Dec 2009)
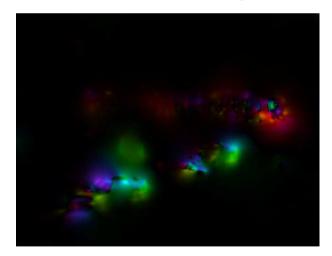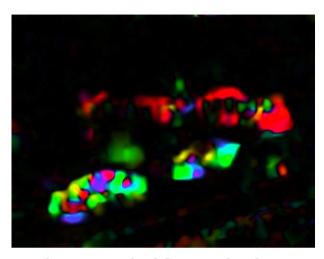
Motion Estimation & Optical Flow

- Motion Estimation
- The Correspondence Problem
- Motion and Grouping
- Motion and Grouping
- Motion and 3D Structure
- Motion and 3D Structure
- Motion and 3D Structure
- Motion and Transparency
- Applications of Motion Estimation
- The Aperture Problem
- The Aperture Problem: Measurability
- The Brightness Constancy Assumption
- The Optic Flow Constraint
- The OFC and the Aperture Problem
- Example: Traffic Scene
- Additional Assumptions
- Lucas and Kanade
- Lucas and Kanade
- Lucas and Kanade
- Lucas and Kanade: Solutions?
- Lucas/Kanade: Example
- Horn and Schunck
- Euler-Lagrange Equations
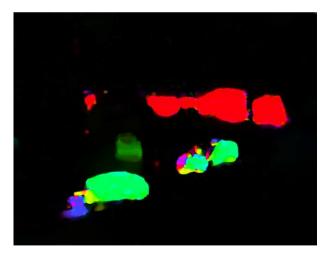- Horn/Schunck: Examples
- Lucas/Kanade vs. Horn/Schunck

Average end-point error	avg. rank	Army (Hidden texture) GT all	im0 disc	im1 untext	Mequon (Hidden texture) GT all	im0 disc	im1 untext	Schefflera (Hidden texture) GT all	im0 disc	im1 untext	Wooden (Hidden texture) GT all	im0 disc	im1 untext	Grove (Synthetic) GT all	im0 disc	im1 untext	Urban (Synthetic) GT all	im0 disc	im1 untext	Yosemite (Synthetic) GT all	im0 disc	im1 untext	Teddy (Stereo) GT all	im0 disc	im1 untext
Adaptive [26]	3.9	0.09	0.26	0.06	0.23	0.78	0.18	0.54	1.19	0.21	0.18	0.91	0.10	0.88	1.25	0.73	0.50	1.28	0.31	0.14	0.16	0.22	0.65	1.37	0.79
Spatially variant [22]	5.5	0.10	0.27	0.08	0.22	0.75	0.19	0.43	1.00	0.18	0.19	1.05	0.10	1.05	1.41	1.16	0.59	1.61	0.43	0.13	0.11	0.28	0.96	1.72	1.28
TV-L1-improved [20]	6.2	0.09	0.26	0.07	0.20	0.71	0.16	0.53	1.18	0.22	0.21	1.24	0.11	0.90	1.31	0.72	1.51	1.93	0.84	0.18	0.17	0.31	0.73	1.62	0.87
Multicue MRF [24]	6.5	0.11	0.26	0.11	0.19	0.53	0.17	0.24	0.49	0.19	0.24	1.13	0.15	0.79	1.10	0.72	1.47	1.60	0.85	0.28	0.19	0.71	0.78	1.53	1.09
F-TV-L1 [18]	7.8	0.14	0.35	0.14	0.34	0.98	0.26	0.59	1.19	0.26	0.27	1.36	0.16	0.90	1.30	0.76	0.54	1.62	0.36	0.13	0.15	0.20	0.68	1.56	0.66
DPOF [21]	7.9	0.15	0.30	0.11	0.34	1.01	0.25	0.29	0.59	0.26	0.26	0.94	0.20	0.80	1.13	0.63	0.90	1.85	0.66	0.27	0.22	0.54	0.65	1.20	0.99
Fusion [9]	8.2	0.11	0.34	0.10	0.19	0.69	0.16	0.29	0.66	0.23	0.20	1.19	0.14	1.07	1.42	1.22	1.35	1.49	0.86	0.20	0.20	0.26	1.07	2.07	1.39
Brox et al. [8]	9.1	0.12	0.37	0.11	0.31	0.97	0.28	0.48	1.11	0.48	0.28	1.28	0.18	1.13	1.57	1.11	1.02	2.02	0.60	0.10	0.13	0.11	0.93	2.00	1.07
Dynamic MRF [10]	10.0	0.12	0.34	0.11	0.22	0.89	0.16	0.44	1.13	0.20	0.24	1.29	0.14	1.11	1.52	1.13	1.54	2.37	0.93	0.13	0.12	0.31	1.27	2.33	1.66
SegOF [13]	10.3	0.15	0.36	0.10	0.57	1.16	0.59	0.68	1.24	0.64	0.32	0.86	0.26	1.18	1.50	1.47	1.63	2.09	0.96	0.08	0.13	0.12	0.70	1.50	0.69
CBF [15]	10.3	0.10	0.28	0.09	0.29	0.79	0.29	0.45	0.98	0.24	0.21	1.22	0.13	0.96	1.39	0.74	2.32	2.19	1.29	0.34	0.27	0.85	0.86	1.78	1.06
Second-order prior [11]	11.4	0.10	0.30	0.08	0.22	0.85	0.15	0.57	1.28	0.23	0.20	1.14	0.11	1.13	1.55	1.03	2.52	2.45	1.25	0.42	0.25	1.09	0.98	1.92	1.07
Learning Flow [14]	12.2	0.11	0.32	0.09	0.29	0.99	0.23	0.55	1.24	0.29	0.36	1.56	0.25	1.25	1.64	1.41	1.55	2.32	0.85	0.14	0.18	0.24	1.09	2.09	1.27
Filter Flow [23]	12.7	0.17	0.39	0.13	0.43	1.09	0.38	0.75	1.34	0.78	0.70	1.54	0.68	1.13	1.38	1.51	0.57	1.32	0.44	0.22	0.23	0.26	0.96	1.66	1.13
GraphCuts [17]	12.9	0.16	0.38	0.14	0.59	1.36	0.46	0.56	1.07	0.64	0.26	1.14	0.17	0.96	1.35	0.84	2.25	1.79	1.22	0.22	0.17	0.43	1.22	2.05	1.78
Black & Anandan 3 [7]	13.8	0.18	0.42	0.19	0.58	1.31	0.50	0.95	1.58	0.70	0.41	1.59	0.45	1.08	1.42	1.22	1.43	2.28	0.83	0.15	0.17	0.17	1.11	1.98	1.30
SPSA-learn [16]	14.3	0.18	0.45	0.17	0.57	1.32	0.51	0.84	1.50	0.72	0.52	1.64	0.49	1.12	1.42	1.39	1.75	2.14	1.06	0.13	0.13	0.19	1.32	2.08	1.73
GroupFlow [12]	15.1	0.21	0.51	0.21	0.79	1.69	0.72	0.86	1.64	0.74	0.30	1.07	0.26	1.29	1.81	0.82	1.94	2.30	1.36	0.11	0.14	0.19	1.06	1.96	1.35
2D-CLG [3]	16.8	0.28	0.62	0.21	0.67	1.21	0.70	1.12	1.80	0.99	1.07	2.06	1.12	1.23	1.52	1.62	1.54	2.15	0.96	0.10	0.11	0.16	1.38	2.26	1.83
Horn & Schunck [6]	17.9	0.22	0.55	0.22	0.61	1.53	0.52	1.01	1.73	0.80	0.78	2.02	0.77	1.26	1.58	1.55	1.43	2.59	1.00	0.16	0.18	0.15	1.51	2.50	1.88
Black & Anandan 2 [2]	17.9	0.21	0.52	0.17	0.65	1.52	0.58	0.93	1.54	1.03	0.76	1.97	0.73	1.40	1.80	1.54	2.04	2.22	1.12	0.15	0.17	0.26	1.68	2.64	2.06
Black & Anandan [1]	21.1	0.26	0.54	0.27	0.99	1.61	1.07	1.23	1.77	1.14	0.98	2.10	0.97	1.59	1.97	1.95	2.19	2.26	1.41	0.20	0.18	0.41	2.34	3.46	2.98
STOB [25]	21.6	0.30	0.70	0.36	1.09	1.77	1.21	1.25	1.98	1.03	1.56	2.26	1.71	1.54	1.82	2.14	2.02	2.79	1.36	0.17	0.16	0.26	2.43	3.18	3.31
FOLKI [19]	23.0	0.29	0.73	0.33	1.52	1.96	1.80	1.23	2.04	0.95	0.99	2.20	1.08	1.53	1.85	2.07	2.14	3.23	1.60	0.26	0.21	0.68	2.67	3.27	4.32
Pyramid LK [4]	26.0	0.39	0.61	0.61	1.67	1.78	2.00	1.50	1.97	1.38	1.57	2.39	1.78	2.94	3.72	2.58	3.33	2.74	2.43	0.30	0.24	0.73	3.80	5.08	4.86

The Middlebury Benchmark (Jan 2012)

Average endpoint error	avg. rank	Army (Hidden texture) GT all	Army im0 disc	Army im1 untext	Mequon (Hidden texture) GT all	Mequon im0 disc	Mequon im1 untext	Schefflera (Hidden texture) GT all	Schefflera im0 disc	Schefflera im1 untext	Wooden (Hidden texture) GT all	Wooden im0 disc	Wooden im1 untext	Grove (Synthetic) GT all	Grove im0 disc	Grove im1 untext	Urban (Synthetic) GT all	Urban im0 disc	Urban im1 untext	Yosemite (Synthetic) GT all	Yosemite im0 disc	Yosemite im1 untext	Teddy (Stereo) GT all	Teddy im0 disc	Teddy im1 untext
MDP-Flow2 [40]	6.6	0.09 9	0.23 4	0.07 5	0.16 2	0.52 1	0.13 3	0.22 2	0.46 3	0.17 6	0.17 15	0.93 23	0.09 8	0.65 5	0.98 6	0.43 5	0.29 1	0.91 2	0.26 3	0.11 7	0.13 10	0.17 8	0.51 10	1.11 12	0.72 9
Layers++ [38]	7.0	0.08 3	0.21 2	0.07 5	0.19 5	0.56 3	0.17 15	0.20 1	0.40 1	0.18 10	0.13 2	0.58 1	0.07 2	0.48 1	0.70 1	0.33 2	0.47 13	1.01 4	0.33 13	0.15 27	0.14 20	0.24 25	0.46 2	0.88 1	0.72 9
nLayers [62]	7.2	0.07 1	0.19 1	0.06 1	0.22 14	0.59 4	0.19 22	0.25 6	0.54 5	0.20 20	0.15 6	0.84 14	0.08 4	0.53 2	0.78 2	0.34 3	0.44 8	0.84 1	0.30 6	0.13 16	0.13 10	0.20 15	0.47 4	0.97 2	0.67 6
Sparse-NonSparse [59]	10.0	0.08 3	0.23 4	0.07 5	0.22 14	0.73 16	0.18 17	0.28 9	0.64 7	0.19 15	0.14 4	0.71 6	0.08 4	0.67 8	0.99 8	0.48 8	0.49 16	1.06 6	0.32 10	0.14 22	0.11 1	0.28 37	0.49 5	0.98 3	0.73 12
TC-Flow [48]	10.6	0.07 1	0.21 2	0.06 1	0.15 1	0.59 4	0.11 1	0.31 15	0.78 16	0.14 1	0.16 10	0.86 16	0.08 4	0.75 13	1.11 14	0.54 12	0.42 6	1.40 19	0.25 2	0.11 7	0.12 5	0.29 40	0.62 17	1.35 18	0.93 29
LSM [41]	10.7	0.08 3	0.23 4	0.07 5	0.22 14	0.73 16	0.18 17	0.28 9	0.64 7	0.19 15	0.14 4	0.70 4	0.09 8	0.66 6	0.97 5	0.48 8	0.50 17	1.06 6	0.33 13	0.15 27	0.12 5	0.29 40	0.50 7	0.99 5	0.73 12
Classic+NL [31]	11.8	0.08 3	0.23 4	0.07 5	0.22 14	0.74 19	0.18 17	0.29 12	0.65 12	0.19 15	0.15 6	0.73 7	0.09 8	0.64 4	0.93 4	0.47 6	0.52 19	1.12 9	0.33 13	0.16 34	0.13 10	0.29 40	0.49 5	0.98 3	0.74 15
IROF-TV [56]	13.0	0.09 9	0.25 12	0.08 14	0.22 14	0.77 20	0.19 22	0.30 14	0.70 14	0.19 15	0.18 20	0.93 23	0.11 20	0.73 10	1.04 9	0.56 15	0.44 8	1.69 32	0.31 8	0.09 3	0.11 1	0.12 3	0.50 7	1.08 10	0.73 12
MDP-Flow [26]	14.2	0.09 9	0.25 12	0.08 14	0.19 5	0.54 2	0.18 17	0.24 4	0.55 6	0.20 20	0.16 10	0.91 20	0.09 8	0.74 11	1.06 11	0.61 18	0.46 12	1.02 5	0.35 18	0.12 11	0.14 20	0.17 8	0.78 33	1.68 36	0.97 32
OFH [39]	15.6	0.10 17	0.25 12	0.09 23	0.19 5	0.69 12	0.14 6	0.43 24	1.02 29	0.17 6	0.17 15	1.08 31	0.08 4	0.87 21	1.25 20	0.73 24	0.43 7	1.69 32	0.32 10	0.10 4	0.13 10	0.18 11	0.59 16	1.40 21	0.74 15
OF-Mol [49]	15.7	0.08 3	0.23 4	0.07 5	0.28 33	0.99 38	0.20 28	0.28 9	0.64 7	0.19 15	0.16 10	0.80 10	0.09 8	0.75 13	1.12 15	0.50 10	0.52 19	1.09 8	0.33 13	0.16 34	0.13 10	0.30 44	0.56 14	1.08 10	0.78 17
Sparse Occlusion [57]	16.1	0.09 9	0.24 9	0.08 14	0.22 14	0.63 8	0.19 22	0.38 22	0.91 21	0.18 10	0.17 15	0.85 15	0.09 8	0.75 13	1.09 13	0.47 6	0.34 3	1.00 3	0.26 3	0.22 53	0.22 58	0.28 37	0.53 12	1.13 13	0.67 6
NL-TV-NCC [25]	16.2	0.10 17	0.26 17	0.08 14	0.22 14	0.72 15	0.15 8	0.35 19	0.85 19	0.16 4	0.15 6	0.70 4	0.09 8	0.79 16	1.16 17	0.51 11	0.78 30	1.38 17	0.48 27	0.16 34	0.15 31	0.26 30	0.55 13	1.16 14	0.55 3
TrajectoryFlow [60]	16.2	0.10 17	0.26 17	0.07 5	0.20 10	0.73 16	0.13 3	0.37 21	0.94 22	0.15 3	0.13 2	0.67 2	0.07 2	0.82 17	1.23 19	0.54 12	0.66 25	1.42 22	0.44 22	0.16 34	0.14 20	0.37 51	0.58 15	1.25 15	0.79 18
CostFilter [42]	17.1	0.10 17	0.27 21	0.08 14	0.20 10	0.63 8	0.15 8	0.22 2	0.45 2	0.18 10	0.19 22	0.88 19	0.12 24	0.60 3	0.90 3	0.28 1	0.75 29	1.19 11	0.50 29	0.21 52	0.24 60	0.40 53	0.46 2	1.02 7	0.62 4
SimpleFlow [52]	17.5	0.09 9	0.24 9	0.08 14	0.24 26	0.78 21	0.20 28	0.43 24	0.96 25	0.21 24	0.16 10	0.77 8	0.09 8	0.71 9	1.04 9	0.55 14	1.47 46	1.59 28	0.76 41	0.13 16	0.12 5	0.22 21	0.50 7	1.04 9	0.72 9
L1-Patches [63]	18.0	0.08 3	0.28 22	0.06 1	0.22 14	0.93 33	0.16 11	0.32 17	0.78 16	0.14 1	0.11 1	0.67 2	0.06 1	1.12 48	1.49 48	1.34 49	0.88 33	1.16 10	0.48 27	0.15 27	0.14 20	0.29 40	0.45 1	1.01 6	0.45 1
Adaptive [20]	20.1	0.09 9	0.26 17	0.06 1	0.23 24	0.78 21	0.18 17	0.54 32	1.19 36	0.21 24	0.18 20	0.91 20	0.10 18	0.88 23	1.25 20	0.73 24	0.50 17	1.28 15	0.31 8	0.14 22	0.16 37	0.22 21	0.65 21	1.37 20	0.79 18
DPOF [18]	21.7	0.12 34	0.33 34	0.08 14	0.26 29	0.80 25	0.20 28	0.24 4	0.49 4	0.20 20	0.19 22	0.83 13	0.13 27	0.66 6	0.98 6	0.40 4	1.11 39	1.41 21	0.57 34	0.25 59	0.14 20	0.55 59	0.51 10	1.02 7	0.54 2
Adapt-Window [34]	21.9	0.10 17	0.24 9	0.09 23	0.19 5	0.59 4	0.15 8	0.27 7	0.64 7	0.17 6	0.18 18	0.82 12	0.11 20	0.74 11	1.07 12	0.56 15	1.78 55	1.73 35	0.95 50	0.22 53	0.16 37	0.45 56	0.70 25	1.28 16	0.88 25
Complementary OF [21]	22.9	0.11 24	0.28 22	0.10 32	0.18 3	0.63 8	0.12 2	0.31 15	0.75 15	0.18 10	0.19 22	0.97 25	0.12 24	0.97 33	1.31 29	1.00 40	1.78 55	1.73 35	0.87 48	0.11 7	0.12 5	0.22 21	0.68 22	1.48 23	0.95 30
ACK-Prior [27]	23.1	0.11 24	0.25 12	0.09 23	0.18 3	0.59 4	0.13 3	0.27 7	0.64 7	0.16 4	0.15 6	0.78 9	0.09 8	0.82 17	1.14 16	0.71 21	1.90 57	1.90 42	0.99 53	0.23 57	0.17 43	0.49 58	0.77 31	1.44 22	0.91 27
ComplOF-FED-GPU [36]	23.3	0.11 24	0.29 26	0.10 32	0.21 13	0.78 21	0.14 6	0.32 17	0.79 18	0.17 6	0.19 22	0.99 26	0.11 20	0.89 24	1.29 24	0.73 24	1.25 41	1.74 37	0.64 38	0.14 22	0.13 10	0.30 44	0.64 19	1.50 25	0.83 21
Classic++ [32]	23.7	0.09 9	0.25 12	0.07 5	0.23 24	0.78 21	0.19 22	0.43 24	1.00 28	0.22 26	0.20 26	1.11 32	0.10 18	0.87 21	1.30 26	0.66 19	0.47 13	1.62 29	0.33 13	0.17 39	0.14 20	0.32 48	0.79 34	1.64 32	0.92 28
Aniso. Huber-L1 [22]	24.1	0.10 17	0.28 22	0.08 14	0.31 36	0.88 30	0.28 39	0.56 36	1.13 32	0.29 39	0.20 26	0.92 22	0.13 27	0.84 19	1.20 18	0.70 20	0.39 4	1.23 12	0.28 5	0.17 39	0.15 31	0.27 35	0.64 19	1.36 19	0.79 18
TriangleFlow [30]	26.3	0.11 24	0.29 26	0.09 23	0.26 29	0.95 36	0.17 15	0.47 30	1.07 30	0.18 10	0.16 10	0.87 18	0.09 8	1.07 40	1.47 45	1.10 43	0.87 32	1.39 18	0.57 34	0.15 27	0.19 53	0.23 24	0.63 18	1.33 17	0.84 22
TV-L1-improved [17]	28.0	0.09 9	0.26 17	0.07 5	0.20 10	0.71 14	0.16 11	0.53 31	1.18 35	0.22 26	0.21 30	1.24 37	0.11 20	0.90 25	1.31 29	0.72 22	1.51 48	1.93 44	0.84 45	0.18 43	0.17 43	0.31 46	0.73 27	1.62 31	0.87 24
LocallyOriented [55]	28.5	0.12 34	0.35 37	0.08 14	0.33 39	1.01 40	0.25 37	0.61 42	1.30 43	0.28 36	0.18 18	0.80 10	0.13 27	0.93 30	1.29 24	0.79 28	0.98 36	1.48 25	0.56 33	0.12 11	0.14 20	0.21 19	0.73 27	1.48 23	0.95 30
CBF [12]	28.8	0.10 17	0.28 22	0.09 23	0.34 40	0.80 25	0.37 41	0.43 24	0.95 24	0.26 32	0.21 30	1.14 33	0.13 27	0.90 25	1.27 23	0.82 30	0.41 5	1.23 12	0.30 6	0.23 57	0.19 53	0.39 52	0.76 30	1.56 27	1.02 33
CLG-TV [51]	30.1	0.11 24	0.29 26	0.09 23	0.32 37	0.86 29	0.30 40	0.55 34	1.17 34	0.28 36	0.25 35	1.05 28	0.17 36	0.92 29	1.30 26	0.79 28	0.47 13	1.72 34	0.35 18	0.17 39	0.17 43	0.25 28	0.74 29	1.57 29	0.88 25
F-TV-L1 [15]	30.4	0.14 41	0.35 37	0.14 46	0.34 40	0.98 37	0.26 38	0.59 41	1.19 36	0.26 32	0.27 38	1.36 44	0.16 35	0.90 25	1.30 26	0.76 27	0.54 21	1.62 29	0.36 20	0.13 16	0.15 31	0.20 15	0.68 22	1.56 27	0.66 5
Brox et al. [5]	30.8	0.11 24	0.32 32	0.11 38	0.27 32	0.93 33	0.22 34	0.39 23	0.94 22	0.24 30	0.24 32	1.25 38	0.13 27	1.10 45	1.39 39	1.43 53	0.89 35	1.77 39	0.55 32	0.10 4	0.13 10	0.11 2	0.91 36	1.83 39	1.13 39
Rannacher [23]	31.1	0.11 24	0.31 30	0.09 23	0.25 27	0.84 28	0.21 32	0.57 38	1.27 42	0.26 33	0.24 32	1.32 41	0.13 27	0.91 28	1.33 31	0.72 22	1.49 47	1.95 46	0.78 42	0.15 27	0.14 20	0.26 30	0.69 24	1.58 30	0.86 23
Second-order prior [8]	32.3	0.11 24	0.31 30	0.09 23	0.26 29	0.93 33	0.20 28	0.57 38	1.25 41	0.26 33	0.20 26	1.04 27	0.12 24	0.94 31	1.34 32	0.83 33	0.61 24	1.93 44	0.47 26	0.20 46	0.16 37	0.34 50	0.77 31	1.64 32	1.07 35
Fusion [6]	32.5	0.11 24	0.34 35	0.10 32	0.19 5	0.69 12	0.16 11	0.29 12	0.66 13	0.23 28	0.20 26	1.19 35	0.14 33	1.07 40	1.42 41	1.22 46	1.35 43	1.49 26	0.86 47	0.20 46	0.20 56	0.26 30	1.07 44	2.07 49	1.39 46
p-harmonic [29]	33.3	0.12 34	0.36 42	0.11 38	0.25 27	0.82 27	0.21 32	0.57 38	1.24 38	0.28 36	0.26 36	1.20 36	0.19 40	1.07 40	1.39 39	1.31 48	0.44 8	1.65 31	0.37 21	0.15 27	0.16 37	0.21 19	0.87 35	1.76 37	1.06 34
Bartels [43]	34.3	0.12 34	0.30 29	0.11 38	0.22 14	0.65 11	0.19 22	0.35 19	0.86 20	0.23 28	0.28 40	1.32 41	0.18 38	0.97 33	1.38 36	0.98 37	1.20 40	1.76 38	0.78 42	0.20 46	0.17 43	0.48 57	0.91 36	1.88 40	1.22 41
Dynamic MRF [7]	36.2	0.12 34	0.34 35	0.11 38	0.22 14	0.89 31	0.16 11	0.44 29	1.13 32	0.20 20	0.24 32	1.29 40	0.14 33	1.11 46	1.52 53	1.13 44	1.54 49	2.37 57	0.93 49	0.13 16	0.12 5	0.31 46	1.27 51	2.33 57	1.66 48
SegOF [10]	36.4	0.15 43	0.36 42	0.10 32	0.57 46	1.16 46	0.59 51	0.68 45	1.24 38	0.64 44	0.32 43	0.86 16	0.26 43	1.18 53	1.50 52	1.47 55	1.63 52	2.09 49	0.96 51	0.08 1	0.13 10	0.12 3	0.70 25	1.50 25	0.69 8
LDOF [28]	36.9	0.12 34	0.35 37	0.10 32	0.32 37	1.06 41	0.24 36	0.43 24	0.98 27	0.30 41	0.45 46	2.48 61	0.26 43	1.01 37	1.37 35	1.05 41	1.10 38	2.08 48	0.67 39	0.12 11	0.15 31	0.24 25	0.94 39	2.05 46	1.11 38

Motion in 3D: Scene Flow

The optical flow gives the motion in the image plane, sometimes referred to as the apparent motion. Based on stereo video one can jointly estimate depth maps and a dense 3D motion field called scene flow.
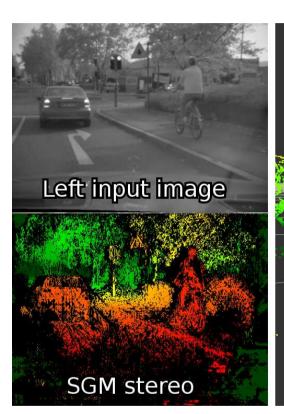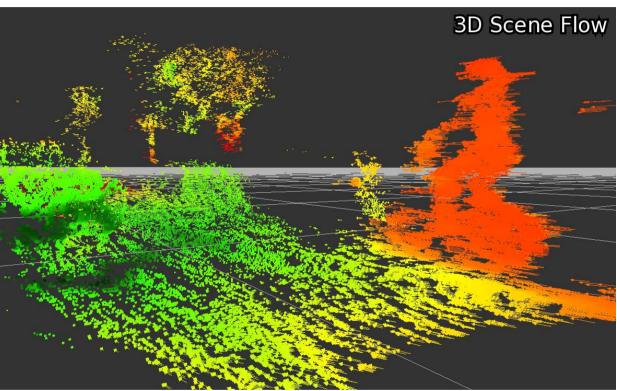

Wedel et al., "Stereoscopic Scene Flow Computation for 3D Motion Understanding", Int. J. of Computer Vision 2011.

Motion in 3D: Scene Flow

Wedel et al., "Stereoscopic Scene Flow Computation for 3D Motion Understanding", Int. J. of Computer Vision 2011.

Andreas Wedel
Daniel Cremers

Stereo Scene Flow
for 3D Motion
Analysis

Springer