



Chapter 10

Motion Estimation and Optical Flow

Computer Vision I: Variational Methods

Winter 2016/17

Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Prof. Daniel Cremers
Chair for Computer Vision and Pattern Recognition
Departments of Informatics & Mathematics
Technical University of Munich

Overview

- 1 Motion Estimation
- 2 The Aperture Problem
- 3 The Optical Flow Constraint
- 4 Regularity Assumption
- 5 Lucas and Kanade
- 6 Horn and Schunck
- 7 Comparison
- 8 Limitations of Classical Approaches
- 9 Brox et al. 2004
- 10 Wedel et al. 2009
- 11 Motion Segmentation
- 12 Scene Flow Estimation



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

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.



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

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 Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Moving regions of random brightness values



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Moving wallpaper regions



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Automatic segmentation of the moving regions.

Cremers, Yuille, DAGM 2003



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation



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



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison


Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation



Schoenemann & Cremers,

Near Real-time Motion Segmentation, DAGM 2006.



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation




Image sequence showing semi-transparent superposition of
two motions.

Author: Michael Black

- **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).



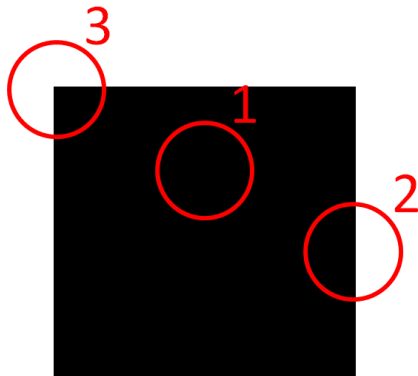
The Aperture Problem



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

Consider three observers each watching a local patch of a moving black box.



Observer 1: no motion can be observed

Observer 2: horizontal motion only

Observer 3: motion in both directions



The Brightness Constancy Assumption

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

Let

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

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 differentiability lead to a relation between the desired velocity field $v(x, t)$ and the spatial and temporal image gradients:

$$\nabla I^\top v + I_t = 0$$

This equation is referred to as the **differential brightness 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



$I_1(x)$



$I_2(x)$



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



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

Author: Daniel Cremers



Additional Assumptions: Spatial Regularity

The optic flow constraint is necessary but not sufficient to uniquely determine a motion field. It only specifies 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. (over 11100 citations in Jan 2016).
- *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. (over 11600 citations in Jan 2016). 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:

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

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





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

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

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:

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

where

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

For each point determine \mathbf{v} by **inversion of a 2×2 -matrix**.

Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation



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

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

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

The matrix M_σ is called **structure tensor**.

Lucas and Kanade: Solutions?



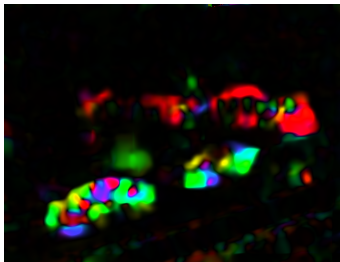
The goal of Lucas and Kanade was to determine a unique velocity vector under the assumption of local constancy of the velocity. Depending on the local intensity structure there are three possible cases (see slide 13):

- 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: $\text{trace}(M) < \epsilon$?)
- 2 All image gradients in the neighborhood U_σ are colinear. Then $\text{rank}(M) = \text{rank}(\nabla I \nabla I^T) = 1$. 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 $\text{rank}(M) = 2$ 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



Hue encodes direction, brightness encodes magnitude.

Author: Thomas Brox





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 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**.

Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

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

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

must fulfill the **Euler-Lagrange equations**:

$$\begin{cases} \frac{\partial E}{\partial v_1} = I_x (I_x v_1 + I_y v_2 + I_t) - \lambda \Delta v_1 = 0 \\ \frac{\partial E}{\partial v_2} = I_y (I_x v_1 + I_y v_2 + I_t) - \lambda \Delta v_2 = 0 \end{cases}$$


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



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation



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,...

Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

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 (for small v) 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 **discontinuities in the flow field**.
- The above approaches are formulated for two images. In general we have **sequences with many 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**:

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

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

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

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

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

- **Gradient constancy** to account for global brightness changes.



Further Advances

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

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

mit

$$E_{data}(\mathbf{v}) = \int \psi \left(\underbrace{|I_2(\mathbf{x} + \mathbf{v}) - I_1(\mathbf{x})|^2}_{\text{brightness constancy}} + \gamma \underbrace{|\nabla I_2(\mathbf{x} + \mathbf{v}) - \nabla I_1(\mathbf{x})|^2}_{\text{gradient constancy}} \right) d\mathbf{x},$$

and

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

where

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

and

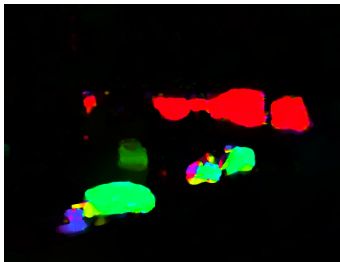
$$\psi(s^2) = \sqrt{s^2 + \epsilon}.$$



Discontinuity-preserving Flow Fields



One of two images



Color coded flow field



Color encodes motion direction and magnitude

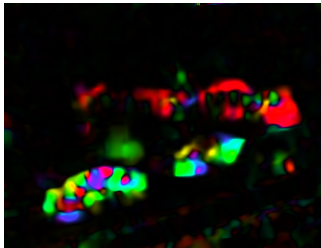
Author: Thomas Brox




Experimental Comparison




first of two images



Lucas & Kanade '81



Horn & Schunck '81



Brox et al. '04

Author: Thomas Brox



Further Advances

Due to various benchmarks like 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:

$$\begin{aligned} & \min_v \int |I_2(x+v) - I_1(x)| + |\nabla v| dx \\ \rightarrow & \min_{v,u} \int |I_2(x+v) - I_1(x)| + |\nabla u| + \frac{1}{2\theta} |u - v|^2 dx \end{aligned}$$

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

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

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



The Middlebury Benchmark



Motion Estimation

The Aperture Problem

The Optical Flow Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Optical flow evaluation results

Statistics: Average SD R0.5 R1.0 R2.0 A50 A75 A95
 Error type: endpoint angle interpolation normalized interpolation

Average endpoint error	Army (Hidden texture)			Mequon (Hidden texture)			Schefflera (Hidden texture)			Wooden (Hidden texture)			Grove (Synthetic)			Urban (Synthetic)			Yosemite (Synthetic)			Teddy (Stereo)				
	avg. rank	GT	im0	im1	GT	im0	im1	GT	im0	im1	GT	im0	im1	GT	im0	im1	GT	im0	im1	GT	im0	im1	GT	im0	im1	
		all	disc	untext	all	disc	untext	all	disc	untext	all	disc	untext	all	disc	untext	all	disc	untext	all	disc	untext	all	disc	untext	
Adaptive [20]	4.4	0.09	0.26	0.06	0.23	0.78	0.46	0.54	0.90	0.21	0.18	0.91	0.10	0.88	1.25	0.73	0.50	1.28	0.31	0.14	0.16	0.12	0.22	0.65	1.37	0.79
Complementary OF [21]	5.7	0.11	0.28	0.10	0.18	0.63	0.12	0.11	0.75	0.18	0.19	0.97	0.12	0.97	1.31	1.00	1.14	1.12	0.22	0.11	0.12	0.22	0.68	1.48	0.95	
Aniso. Huber-L1 [22]	5.8	0.10	0.28	0.08	0.31	0.88	0.28	0.14	0.11	0.29	0.20	0.92	0.13	0.84	1.20	0.70	0.39	1.23	0.28	0.17	0.15	0.27	0.64	1.36	0.79	
DPOF [18]	6.1	0.13	0.35	0.09	0.25	0.79	0.07	0.31	0.90	0.21	0.19	0.62	0.15	0.74	1.09	0.49	0.66	1.80	0.63	0.19	0.17	0.14	0.35	0.50	1.08	0.55
TV-L1-improved [17]	7.2	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	0.16	0.17	0.31	0.73	1.62	0.87
CBF [12]	7.8	0.10	0.28	0.09	0.34	1.00	0.37	0.43	0.95	0.26	0.21	1.14	0.13	0.90	1.27	0.82	0.41	1.23	0.30	0.23	0.22	0.19	0.39	0.76	1.56	1.02
Brox et al. [5]	8.4	0.11	0.32	0.11	0.27	0.93	0.22	0.39	0.94	0.24	0.24	1.25	0.13	1.10	1.39	1.43	0.89	1.77	0.55	0.10	0.13	0.11	0.91	1.83	1.13	
Rannacher [23]	8.5	0.11	0.31	0.09	0.25	0.84	0.21	0.57	1.27	0.26	0.24	1.32	0.13	0.91	1.33	0.72	1.49	1.95	0.78	0.15	0.14	0.26	0.69	1.58	0.86	
F-TV-L1 [15]	8.8	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	
Second-order prior [8]	9.0	0.11	0.31	0.09	0.26	0.93	0.20	0.57	1.25	0.26	0.20	1.04	0.12	0.94	1.34	0.83	0.61	1.93	0.47	0.20	0.16	0.12	0.34	0.77	1.64	1.07
Fusion [6]	9.4	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	0.18	0.20	0.26	1.07	1.42	1.39
Dynamic MRF [7]	11.1	0.12	0.34	0.10	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	1.83	1.66	
SegOF [10]	11.7	0.15	0.36	0.14	0.57	1.16	0.59	0.68	1.24	0.64	0.32	1.06	0.26	1.63	1.50	1.47	1.63	2.09	1.46	0.08	0.13	0.12	0.70	1.57	0.89	
Learning Flow [11]	13.3	0.11	0.32	0.09	0.29	1.00	0.23	0.55	1.24	0.29	0.36	1.56	0.25	1.25	1.64	2.11	1.55	1.72	0.85	0.14	0.10	0.18	0.24	1.09	1.89	1.27
Filter Flow [19]	14.3	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	0.52	1.32	0.44	0.22	0.20	0.23	0.26	0.96	1.66	1.12
GraphCuts [14]	14.5	0.16	0.38	0.15	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	2.17	0.22	0.22	0.17	0.14	0.43	1.22	1.77	1.78
Black & Anandan [4]	15.0	0.18	0.42	0.19	0.58	1.31	0.52	0.95	1.58	0.70	0.49	1.59	0.45	1.08	1.42	1.22	0.15	1.16	0.78	0.13	0.12	0.17	0.11	1.11	1.98	1.30
SPSA-learn [13]	15.7	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	1.94	1.06	0.13	0.13	0.19	1.32	1.89	1.73	
GroupFlow [9]	15.9	0.21	0.51	0.21	0.79	1.69	0.72	0.86	1.64	0.74	0.30	1.4	0.77	1.29	2.21	1.82	0.82	1.94	2.30	0.11	0.14	0.19	1.06	1.96	1.35	
2D-CLG [1]	17.4	0.28	0.62	0.21	0.67	2.01	0.70	1.12	2.1	0.99	1.07	2.2	0.61	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 [3]	18.6	0.22	0.55	0.22	0.61	1.53	0.52	1.01	2.0	0.80	0.73	2.0	0.77	1.26	2.0	1.55	1.43	2.59	1.00	0.16	0.14	0.18	0.15	1.51	2.50	1.88
TI-DOFE [24]	19.6	0.38	0.64	0.47	1.16	2.2	1.26	1.39	2.06	1.17	1.29	2.21	1.41	1.27	2.1	1.57	1.28	2.57	1.01	0.13	0.15	0.16	1.87	2.71	2.53	
FOLKI [16]	22.6	0.29	0.73	0.33	1.52	2.1	1.80	1.23	2.04	0.95	0.99	2.1	1.08	1.53	2.1	2.07	2.14	2.3	1.60	0.26	0.21	0.22	0.68	2.67	3.27	4.32
Pyramid LK [2]	23.7	0.39	0.61	0.61	1.67	2.4	2.00	1.50	2.4	1.97	1.57	2.4	1.78	2.94	3.72	2.98	3.33	2.74	2.43	0.30	0.24	0.24	0.73	3.80	5.08	4.88

Move the mouse over the numbers in the table to see the corresponding images. Click to compare with the ground truth.

Source: Baker et al., "A Database and Evaluation Methodology for Optical Flow", IJCV 2011.

The Middlebury Benchmark



Algorithm	Data term				Prior term					Optimization					Misc.			
	L1 norm	Other robust penalty Fn	Gradient/other features	Illum. modeling/norm/color	L1/TV norm	Other robust penalty Fn	Spatial weighting	Anisotropic weighting	Higher-order prior	Rigidity prior	Cont.-gradient descent	Cont.-variational/extremal	Cont.-other	Discr.-fusion	Discr.-reparameterization	Learning	Visibility/occlusion	Color
Adaptive (Wedel et al. 2009)	X		X	X	X	X	X		X			X						
Complementary OF (Zimmer et al. 2009)	X		X	X		X	X	X			X							X
Aniso. Huber-L1 (Werlberger et al. 2009)	X		X			X	X	X				X						
DPOF (Lei and Yang 2009)				X	X		X				X				X		X	X
TV-L1-improved (Wedel et al. 2008)	X			X	X							X						
CBF (Trobin et al. 2008)	X				X				X				X					X
Brox et al. (Brox et al. 2004)	X	X			X						X							X
F-TV-L1 (Wedel et al. 2008)	X				X				X			X						
Second-order prior (Trobin et al. 2008)	X								X			X						
Fusion (Lempitsky et al. 2008)		X	X			X	X			X			X					X
Dynamic MRF (Glocker et al. 2008)	X				X									X				
Seg OF (Xu et al. 2008)	X				X		X				X					X	X	X
Learning Flow (Sun et al. 2008)		X	X			X		X			X					X		
Filter Flow (Seitz and Baker 2009)	X			X	X		X		X			X						X
Graph Cuts (Cooke 2008)	X				X									X				X
Black & Anandan (Black and Anandan 1996)		X					X			X								
SPSA-learn (Li and Huttenlocher 2008)		X					X			X						X		X
Horn & Schunck (Horn and Schunck 1981)											X							

Motion Estimation

The Aperture Problem

The Optical Flow Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Source: Baker et al., IJCV 2011



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Hand moving toward camera

Motion Segmentation

The estimated motion is typically not the final goal of image analysis. In a real-world video one may be interested in segmenting the differently moving objects.

Motion-based segmentation is a **chicken-and-egg problem**: To reliably estimate motion, we need a certain support-area (ideally the entire region which moves coherently). Yet, in order to partition the image plane into coherently moving region, we need to know the motion at each pixel.

In *Cremers, Soatto, IJCV 2005*, we tackled this chicken-and-egg problem, proposing a method to **jointly estimate a segmentation and motion vectors v_i associated with each region Ω_i** :

$$E(\Omega_1, \dots, \Omega_n, v_1, \dots, v_n) = \sum_{i=1}^n \int_{\Omega_i} |\nabla I^\top v_i + I_t|^2 dx + \frac{\nu}{2} |\partial \Omega_i|.$$

This can be seen as a variation of the Mumford-Shah model, where rather than estimating the average brightness of each region we estimate its average motion.



Piecewise Parametric Motion Segmentation

While the above model allows to separate differently **translating** regions, in many real world scenarios objects undergo rigid motion giving rise to **rotational** or **zooming** flow fields.

The above model can be extended to allow a parametric motion for each region:

$$v_i(x) = S(x) p_i = \begin{pmatrix} x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \end{pmatrix} (a_i \ b_i \ c_i \ d_i \ e_i \ f_i)^\top,$$

with a vector $p_i \in \mathbb{R}^6$ defining an **affine motion** for region Ω_i .

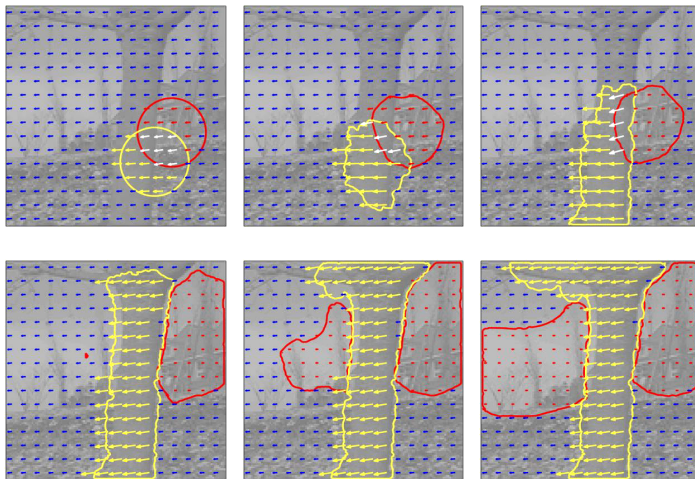
The variational approach

$$E(\Omega_1, \dots, \Omega_n, p_1, \dots, p_n) = \sum_{i=1}^n \int_{\Omega_i} |\nabla I^\top S(x) p_i + I_t|^2 dx + \frac{\nu}{2} |\partial \Omega_i|.$$

leads to a **piecewise parametric motion field**. In the two-region case, this can be solved at around 30 fps – see **Schoenemann, Cremers, DAGM 2006**.



Piecewise Parametric Motion Segmentation



Motion-based segmentation into depth layers.

Cremers, Soatto, IJCV 2005.



Piecewise Parametric Motion Segmentation



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Hand moving toward camera

Cremers, Soatto, IJCV 2005.

Piecewise Parametric Motion Segmentation



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Piecewise parametric motion segmentation with level sets

Cremers, Soatto, IJCV 2005.

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**.

Let $I(x, y, t)^\ell$ and $I(x, y, t)^r$ be the left and right images at pixel (x, y) and time t and d the stereo disparity for that pixel.

Wedel et al. '08 make several constancy assumptions:

$$I(x, y, t)^\ell = I(x + u, y + v, t + 1)^\ell$$

$$I(x + d, y, t)^r = I(x + d + d' + u, y + v, t + 1)^r,$$

where d' denotes the change in disparity (motion in z-direction). Enforcing consistency of the left and right images at time $t + 1$ leads to:

$$I(x + u, y + v, t + 1)^\ell - I(x + d + d' + u, y + v, t + 1)^r = 0.$$

Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Variational Scene Flow Estimation

Wedel et al. '08 combine these assumptions with a smoothness prior for **variational scene flow estimation**:

$$E(u, v, d') = E_{\text{data}} + E_{\text{smooth}},$$

with

$$\begin{aligned} E_{\text{data}} = & \int_{\Omega} \left| I(x+u, y+v, t+1)^{\ell} - I(x, y, t)^{\ell} \right| dx dy \\ & + \int_{\Omega} c(x, y) \left| I(x_d+d'+u, y+v, t+1)^r - I(x_d, y, t)^r \right| dx dy \\ & + \int_{\Omega} c(x, y) \left| I(x_d+d'+u, y+v, t+1)^r - I(x+u, y+v, t+1)^{\ell} \right| dx dy, \end{aligned}$$

and

$$E_{\text{smooth}} = \int_{\Omega} \sqrt{\lambda |\nabla u|^2 + \lambda |\nabla v|^2 + \gamma |\nabla d'|^2} dx dy,$$

where $c: \Omega \rightarrow \{0, 1\}$, $c(x, y) = 0$ if there is no disparity at (x, y) .



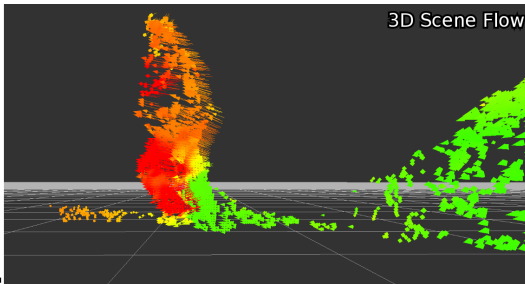
Motion in 3D: Scene Flow



Left image at time t



Left image at time $t+1$



Wedel, Brox, Vaudrey, Rabe, Franke, Cremers, "Stereoscopic Scene Flow Computation for 3D Motion Understanding", Int. J. of Computer Vision 2011.

Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches



Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

Motion in 3D: Scene Flow



Wedel, Brox, Vaudrey, Rabe, Franke, Cremers, "Stereoscopic Scene Flow Computation for 3D Motion Understanding", Int. J. of Computer Vision 2011.

Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

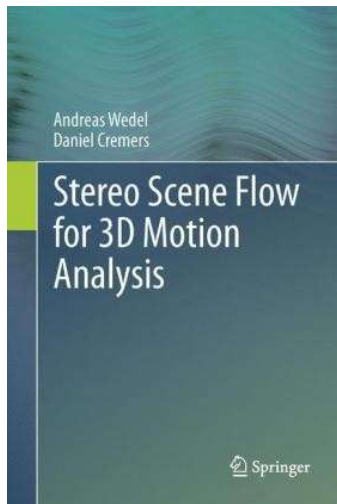
Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation

*Wedel, Brox, Vaudrey, Rabe, Franke, Cremers, "Stereoscopic
Scene Flow Computation for 3D Motion Understanding",
Int. J. of Computer Vision 2011.*



Motion Estimation

The Aperture Problem

The Optical Flow
Constraint

Regularity Assumption

Lucas and Kanade

Horn and Schunck

Comparison

Limitations of Classical
Approaches

Brox et al. 2004

Wedel et al. 2009

Motion Segmentation

Scene Flow Estimation