High Accuracy Optical Flow Estimation Based on a Theory for Warping

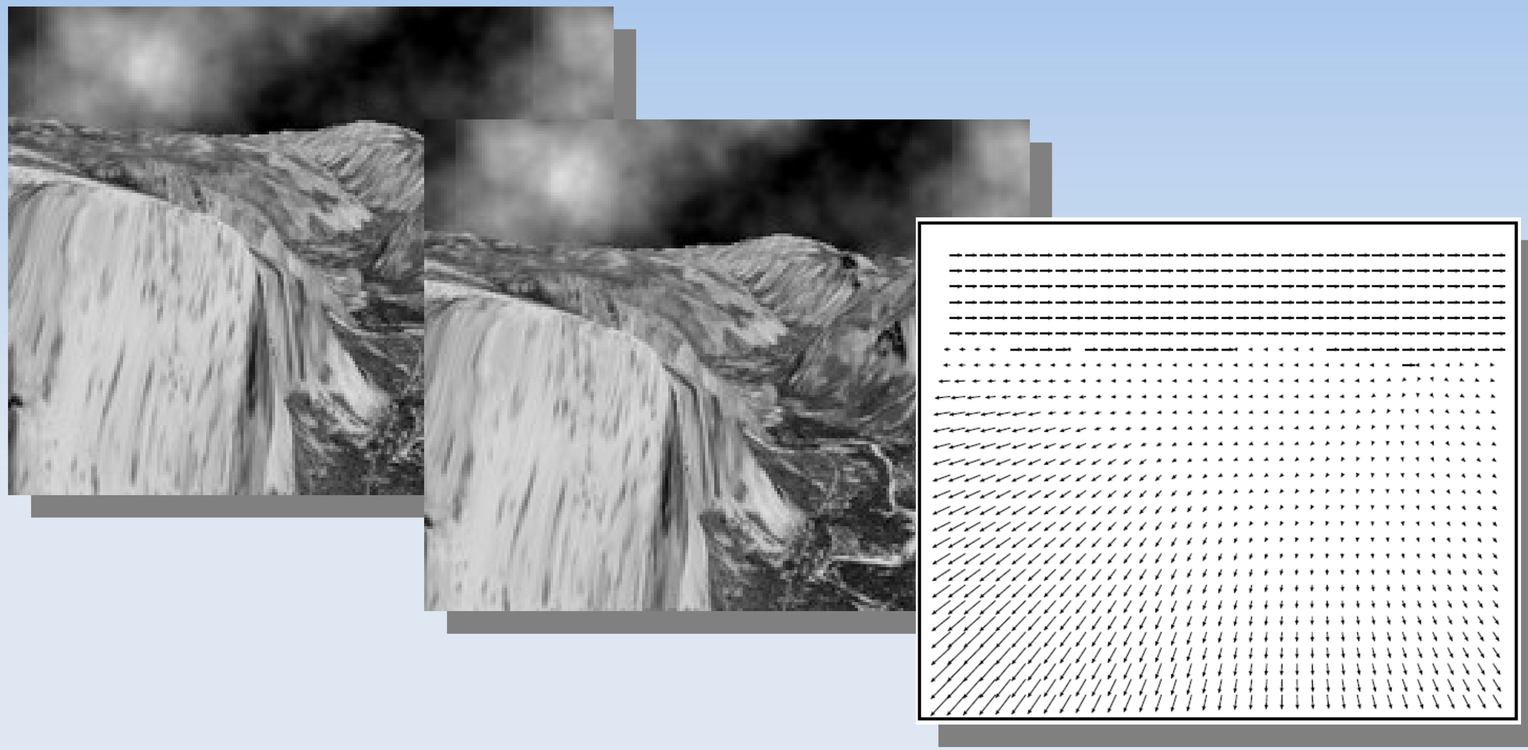

talk by

Sebastian Limbach

based on a prizewinning paper by

Thomas Brox, Andrés Bruhn, Nils Papenberg and Joachim Weickert

Outline

- Motivation
- A novel variational approach
- Coarse-to-fine warping
- Summary

Motivation

What is optical flow estimation?

- displacement field between two images
- correspondence problem
- understanding of details of existing methods and quality of new methods have increased dramatically

Motivation

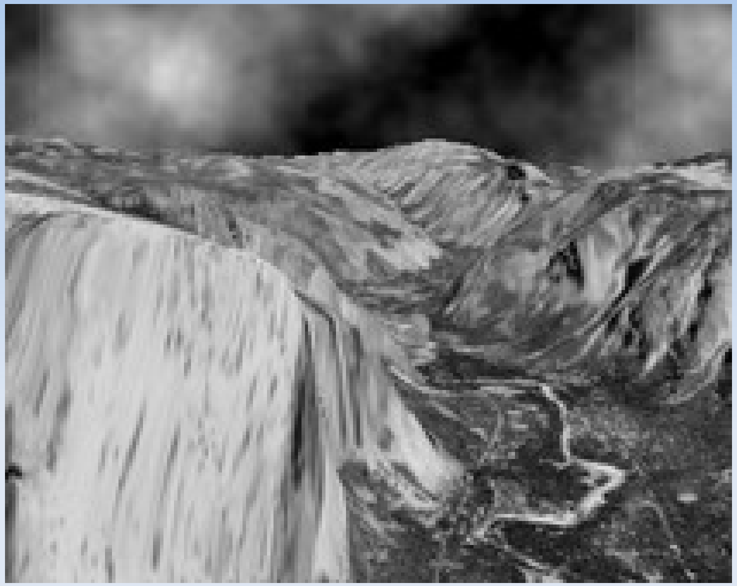

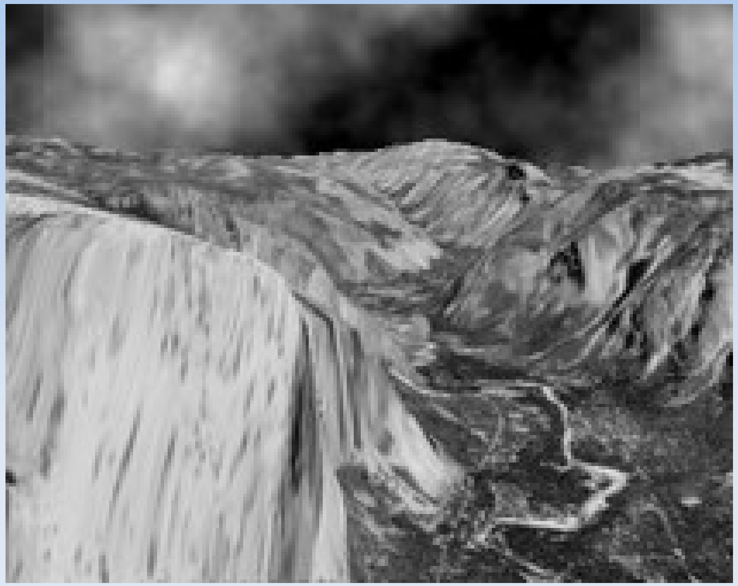

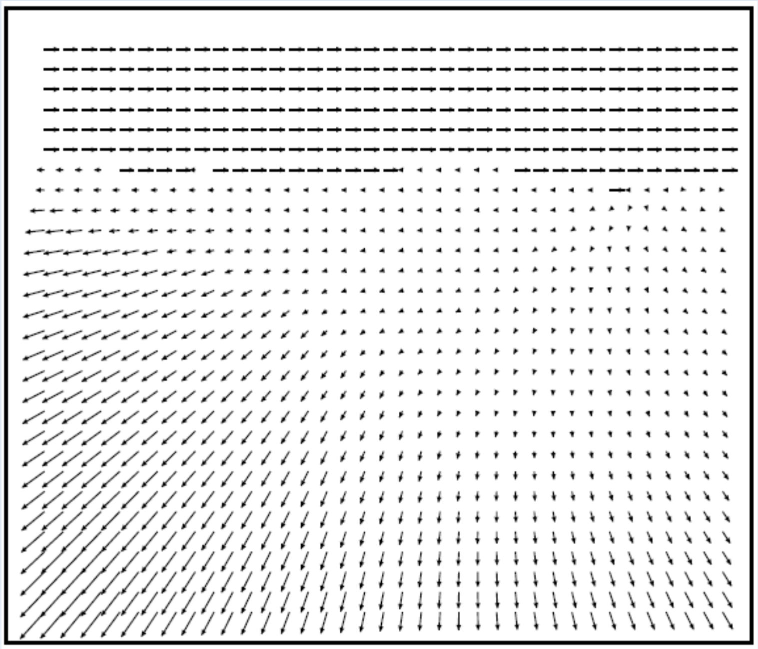

Motivation

Applications

- Motion planning e.g. collision avoidance
- Computer Vision
- Video compression (although mpeg uses much simpler algorithms)
- ...

Motivation

The new approach...

- ...combines concepts from several methods for optical flow estimation
- ...avoids many shortcomings of those methods
- ...outperforms all methods from the literature so far
- ...is very robust with respect to noise and parameter variations

The variational approach

Basic Ideas

- Given: Image sequence

$$I(x, y, t)$$

- Required: Displacement vectors

$$(u(x, y, t), v(x, y, t), 1)^T$$

The variational approach

Variational method

- different constraints on displacement vectors -> energy functional
- looking for functions u and v which minimize the energy functional

Constraints

- grey value constancy assumption

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

- linearisation by Taylor Series expansion yields the famous optical flow constraint

$$I_x u+I_y v+I_t=0$$

- no linearisation in this approach
- gradient constancy assumption

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

- again no linearisation

Constraints

Smoothness assumption

- no interaction between neighbouring pixels so far
- leads to problems where the flow field vanishes or cannot fully be determined (aperture problem)
- further assumption: (piecewise) smoothness of the flow field

Energy functional

Putting it all together

$$E_{Data}(u,v)=\int_{\Omega}\Psi\left(\left|I(\vec{x}+\vec{w})-I(\vec{x})\right|^2+\gamma\left|\nabla I(\vec{x}+\vec{w})-\nabla I(\vec{x})\right|^2\right)d\vec{x}$$

- The energy functional integrates over the whole image domain Ω
- $\vec{x}:=(x,y,t)^T$
- $\vec{w}:=(u,v,1)^T$

Energy functional

Putting it all together

$$E_{Data}(u,v)=\int_{\Omega}\Psi\left(\left|I(\vec{x}+\vec{w})-I(\vec{x})\right|^2+\gamma\left|\nabla I(\vec{x}+\vec{w})-\nabla I(\vec{x})\right|^2\right)d\vec{x}$$

- grey value constancy assumption

Energy functional

Putting it all together

$$E_{Data}(u,v)=\int_{\Omega}\Psi\left(\left|I(\vec{x}+\vec{w})-I(\vec{x})\right|^2+\gamma\left|\nabla I(\vec{x}+\vec{w})-\nabla I(\vec{x})\right|^2\right)d\vec{x}$$

- gradient constancy assumption
- $\gamma\geq 0$, weight between both assumptions

Energy functional

Putting it all together

$$E_{Data}(u,v)=\int_{\Omega}\Psi\left(\left|I(\vec{x}+\vec{w})-I(\vec{x})\right|^2+\gamma\left|\nabla I(\vec{x}+\vec{w})-\nabla I(\vec{x})\right|^2\right)d\vec{x}$$

- to get a more robust energy, an increasing, convex function is applied
- $\Psi(s^2)=\sqrt{s^2+\epsilon^2},\epsilon=0.001$

Energy functional

Putting it all together

$$E_{Smooth}(u, v) = \int_{\Omega} \Psi\left(\left|\nabla_3 u\right|^2 + \left|\nabla_3 v\right|^2\right) d\vec{x}$$

- smoothness assumption
- penalising the total variation of the flow field
- spatial / spatio-temporal gradient

$$\nabla := (\partial x, \partial y)^T, \nabla_3 := (\partial x, \partial y, \partial t)^T$$

Energy functional

Putting it all together

$$E(u,v)=E_{Data}+\alpha E_{Smooth}$$

- regularisation parameter α > 0
- goal is to find functions u and v that minimise this energy
- the calculus of variations states that minimising functions must fulfil the Euler-Lagrange equations

Minimisation

Some abbreviations

$$I_x := \partial_x I(\vec{x}+\vec{w})$$

$$I_y := \partial_y I(\vec{x}+\vec{w})$$

$$I_z := I(\vec{x}+\vec{w}) - I(\vec{x})$$

$$I_{xx} := \partial_{xx} I(\vec{x}+\vec{w})$$

$$I_{xy} := \partial_{xy} I(\vec{x}+\vec{w})$$

$$I_{yy} := \partial_{yy} I(\vec{x}+\vec{w})$$

$$I_{xz} := \partial_x I(\vec{x}+\vec{w}) - \partial_x I(\vec{x})$$

$$I_{yz} := \partial_y I(\vec{x}+\vec{w}) - \partial_y I(\vec{y})$$

Minimisation

Euler-Lagrange-Equations

$$\Psi'(I_z^2+\gamma(I_{xz}^2+I_{yz}^2))\cdot(I_x I_z+\gamma(I_{xx} I_{xz}+I_{xy} I_{yz}))$$
$$-\alpha\,\mathrm{div}(\Psi'(|\nabla_3 u|^2+|\nabla_3 v|^2)\nabla_3 u)=0$$
$$\Psi'(I_z^2+\gamma(I_{xz}^2+I_{yz}^2))\cdot(I_y I_z+\gamma(I_{yy} I_{yz}+I_{xy} I_{xz}))$$
$$-\alpha\,\mathrm{div}(\Psi'(|\nabla_3 u|^2+|\nabla_3 v|^2)\nabla_3 v)=0$$

- non-convex functions
- non-linear functions

Numerical Approximation

Fixed point iteration

$$\Psi'\left(\left(I_z^{k+1}\right)^2+\gamma\left(\left(I_{xz}^{k+1}\right)^2+\left(I_{yz}^{k+1}\right)^2\right)\right)\cdot\left(I_x^k I_z^{k+1}+\gamma\left(I_{xx}^k I_{xz}^{k+1}+I_{xy}^k I_{yz}^{k+1}\right)\right)$$
$$-\alpha\ \mathrm{div}\left(\Psi'\left(\left|\nabla_3 u^{k+1}\right|^2+\left|\nabla_3 v^{k+1}\right|^2\right)\nabla_3 u^{k+1}\right)=0$$

- index k indicates iteration step
- $w^k = (u^k, v^k, 1)^T$
- starting with initial values, we are looking for new values w^{k+1} in each step
- still non-linear

Numerical Approximation

Taylor Series Expansion

$$I_z^{k+1} \approx I_z^k + I_x^k du^k + I_y^k dv^k$$
$$I_{xz}^{k+1} \approx I_{xz}^k + I_{xx}^k du^k + I_{xy}^k dv^k$$
$$I_{yz}^{k+1} \approx I_{yz}^k + I_{xy}^k du^k + I_{yy}^k dv^k$$

- first order expansions for $I_{?z}$ are applied
- linearisation in the numerical scheme instead of the model assumptions
- unknowns splitted:

$$u^{k+1} = u^k + du^k, \quad v^{k+1} = v^k + dv^k$$

Numerical Approximation

Second fixed point iteration

- to remove non-linearity in ψ', a second fixed point iteration is applied
- applied to discrete image data we end up at a linear system of equations in the unknown increments $du^{k,l+1}$ and $dv^{k,l+1}$
- can be solved with common numerical methods (Gauss-Seidel, SOR iterations, ...)

Warping

Warping

- pyramid of images with a scaling factor η
- algorithm starts at the coarsest scale with initial values
- inner fixed point iteration on this scale gives dw^0 and therefore $w^1 = w^0 + dw^0$
- on the next finer scale the second frame is warped by w^1 using bilinear interpolation
- the increments du^1 and dv^1 are computed and one gets w^2 etc.

Warping

initial values

w^0

coarse

$I^0(x,y,t)$

$I^0(x,y,t+1)$

Warping

initial values

coarse

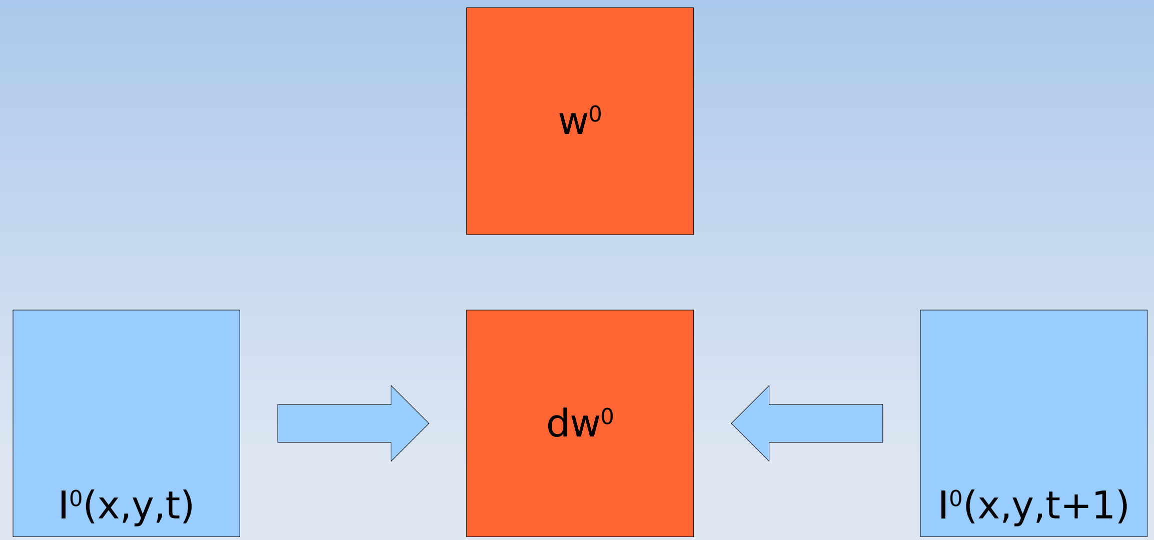

Warping

initial values

coarse

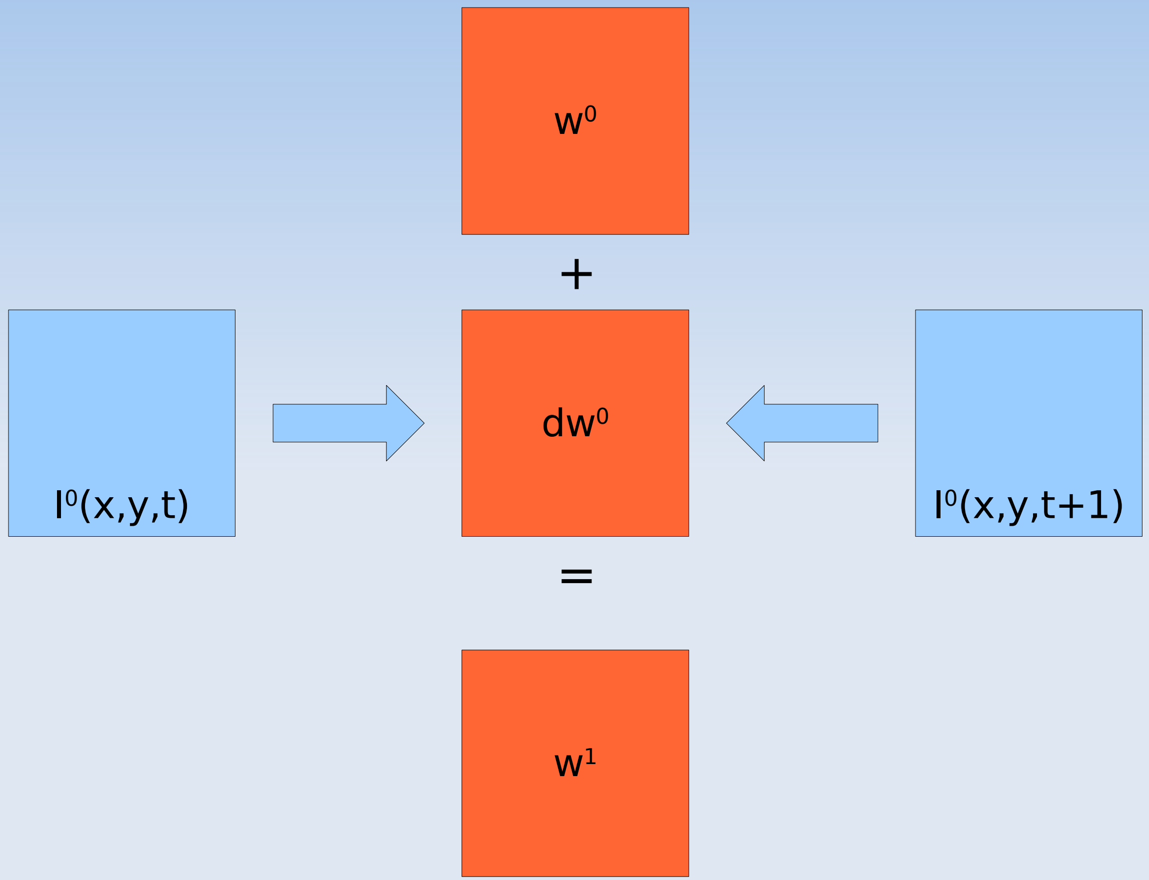

Warping

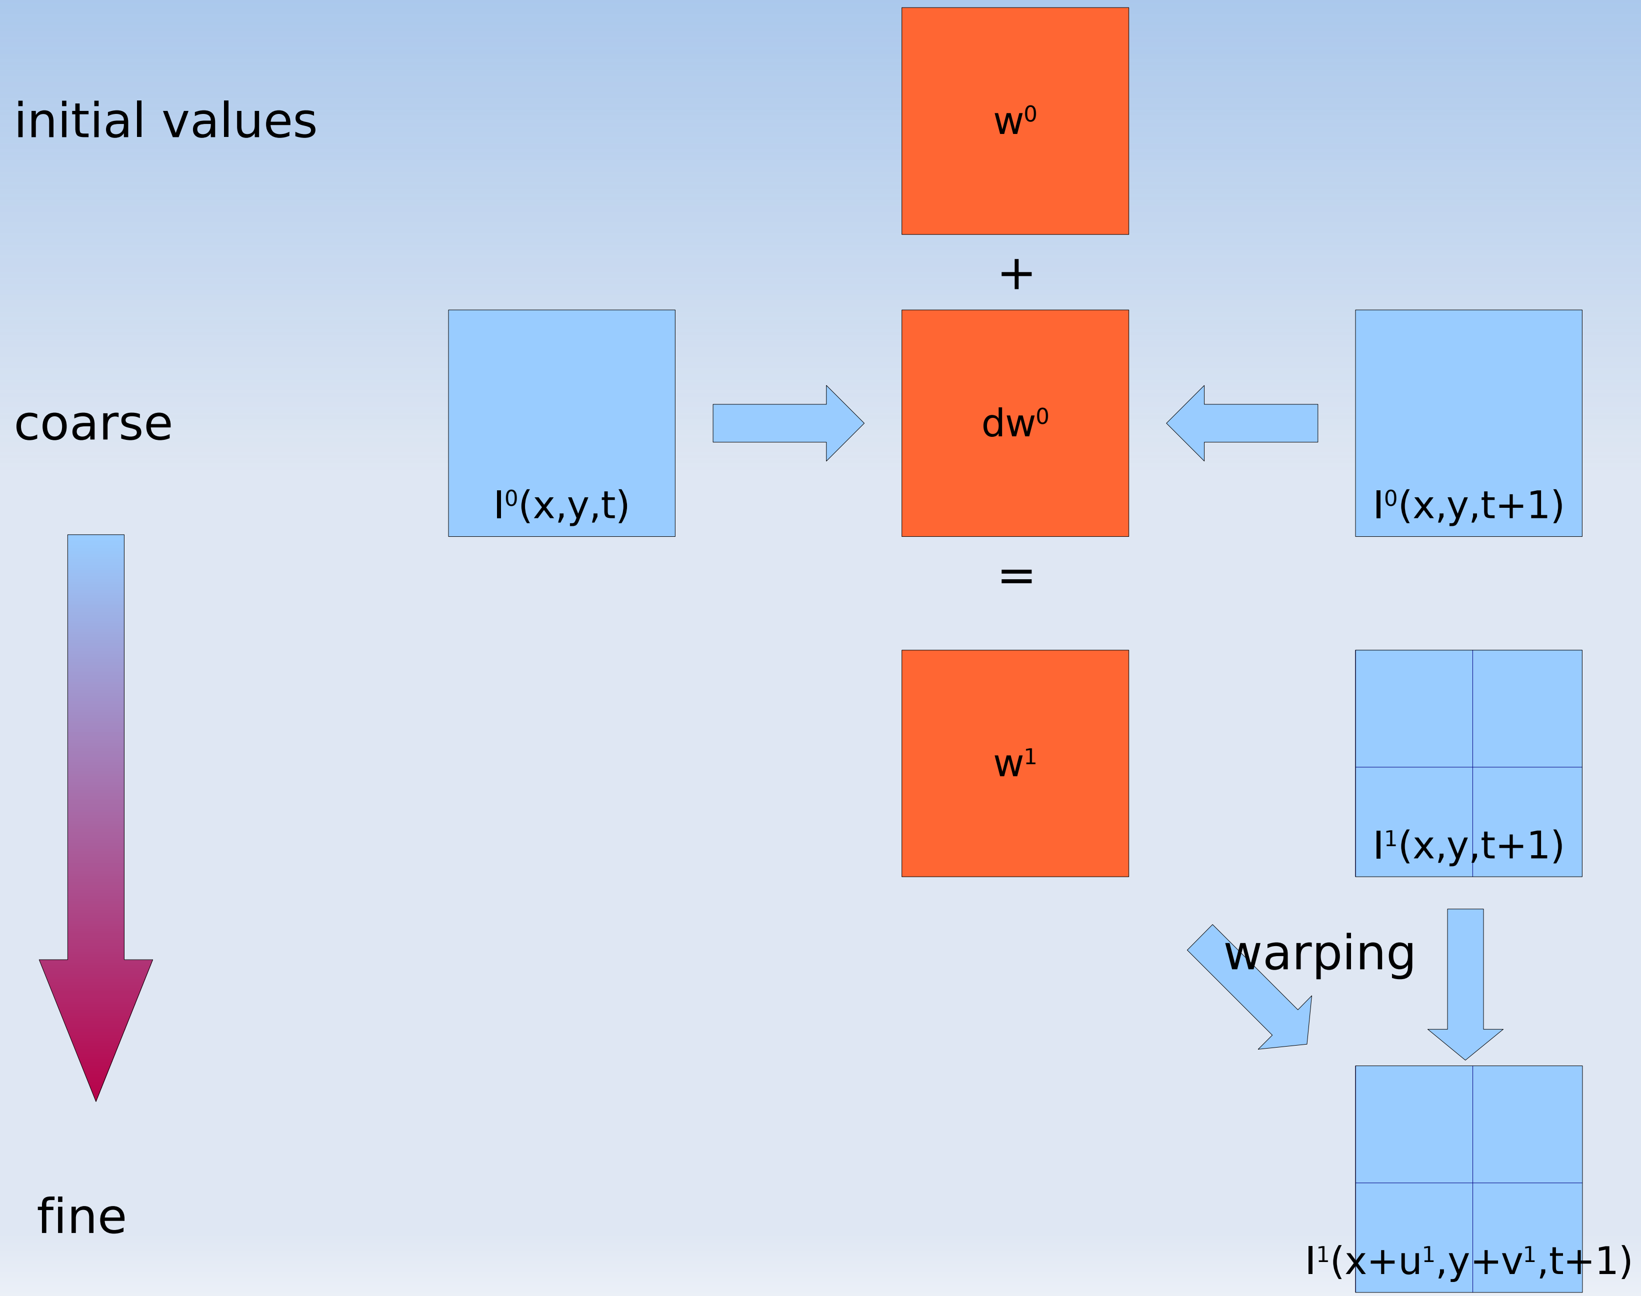
initial values
coarse
fine
w⁰
+
dw⁰
=
w¹
I⁰(x,y,t)
I⁰(x,y,t+1)
I¹(x,y,t+1)
warping
I¹(x+u¹,y+v¹,t+1)

Warping

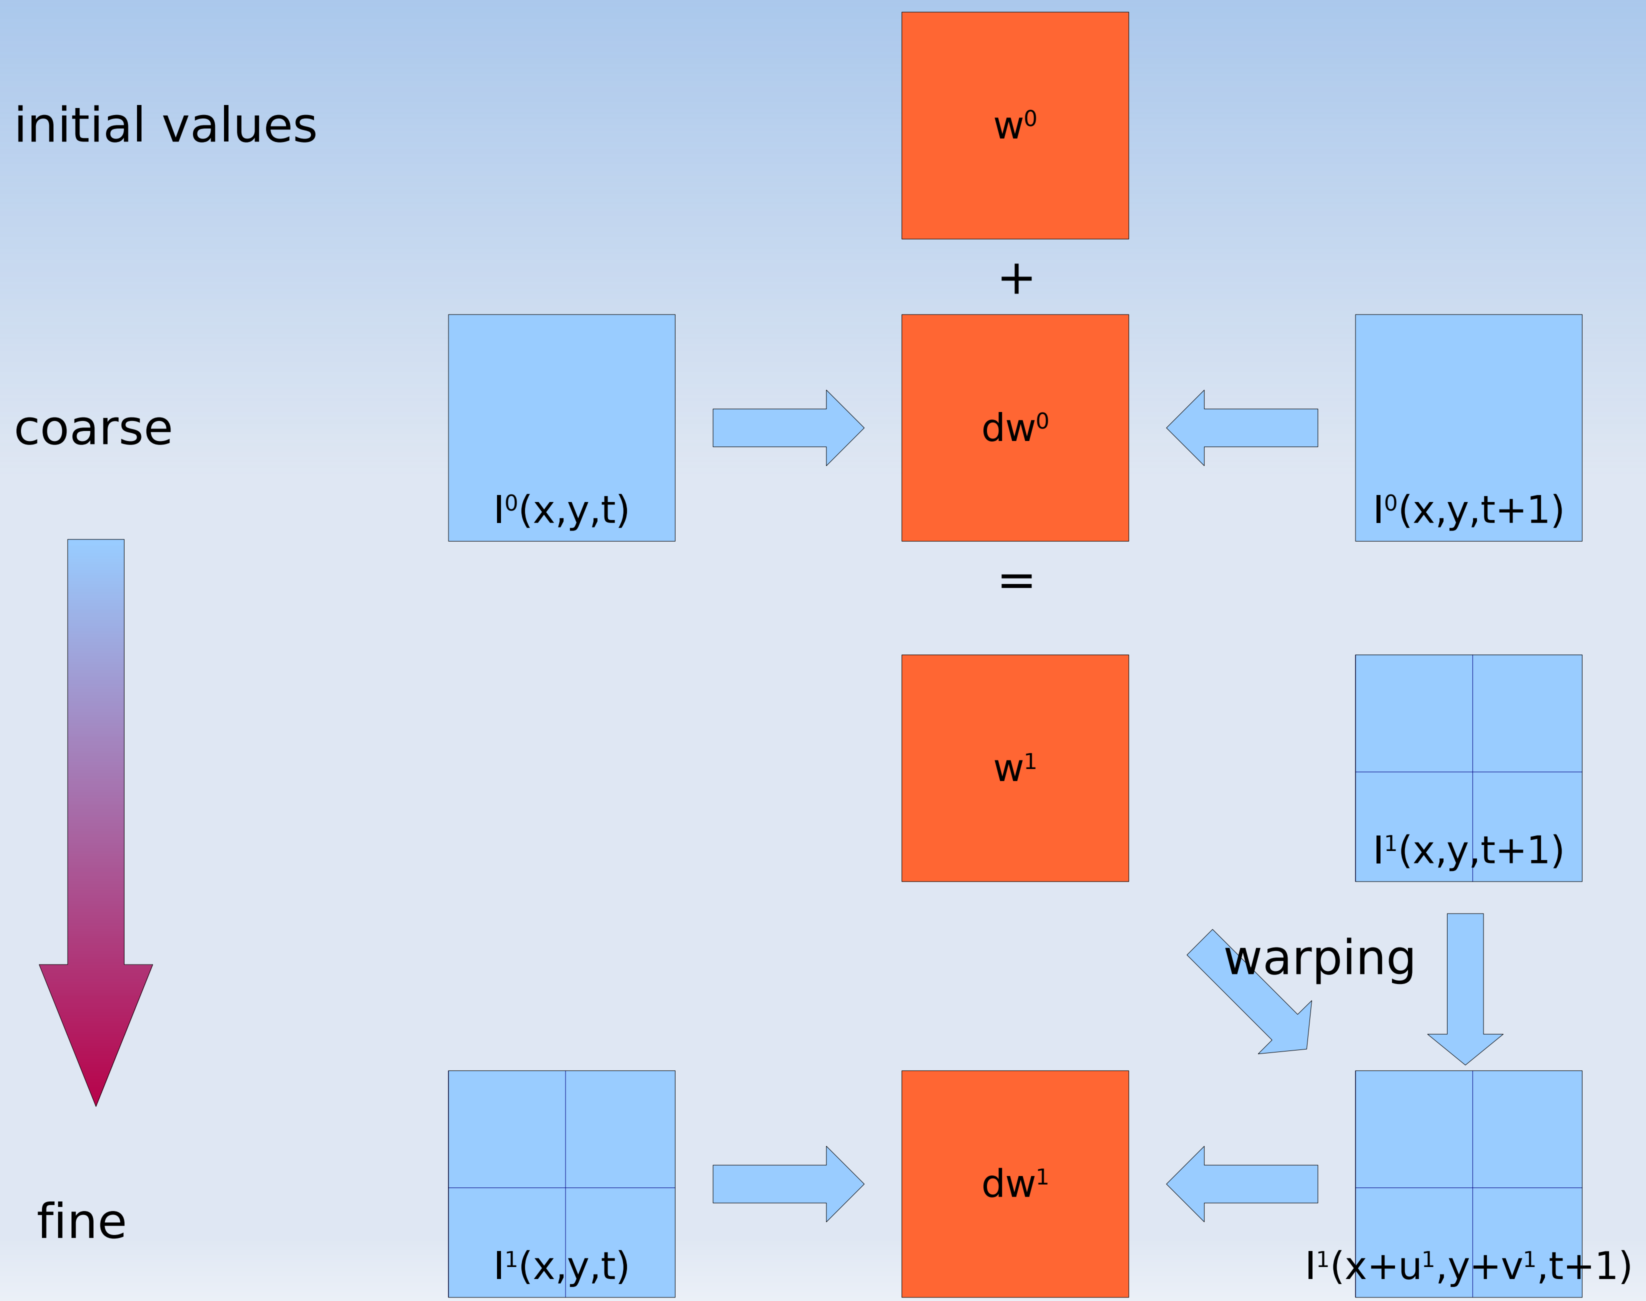
initial values
coarse
fine
w⁰
+
dw⁰
=
w¹
dw¹
I⁰(x,y,t)
I⁰(x,y,t+1)
I¹(x,y,t+1)
warping
I¹(x,y,t)
I¹(x+u¹,y+v¹,t+1)

Warping

What has been reached?

- by splitting the unknowns and using the coarse-to-fine warping we have to solve a difference problem at each scale
- we end up solving a series of convex problems instead of the non-convex initial problem
- this strategy is needed to avoid local minima
- the warping strategy is theoretically justified

Warping

Connection to previous approaches (1)

- many optic flow methods use the linearised optical flow constraint (vs. non-linearised OFC)
- convex energy functional, but the model cannot handle large displacements well (vs. non-convex energy functional)
- therefore the coarse-to-fine warping is applied -> only the (small) difference dw has to be computed
- at each scale, a separate linearised energy functional is minimised

Warping

Connection to previous approaches (2)

- similar results in both approaches
- both approaches result in equivalent Euler-Lagrange equations
- -> the “linearised OFC / warping” - approach minimises a non-linearised constancy assumption by means of fixed point iteration on w

Evaluation

Benchmark results (1)

Yosemite with clouds

Technique	AAE	STD
Nagel [5]	10.22°	16.51°
Horn–Schunck, mod. [5]	9.78°	16.19°
Uras *et al.* [5]	8.94°	15.61°
Alvarez *et al.* [2]	5.53°	7.40°
Weickert *et al.* [24]	5.18°	8.68°
Mémin–Pérez [16]	4.69°	6.89°
Our method (2D)	**2.46°**	**7.31°**
Our method (3D)	**1.94°**	**6.02°**

Yosemite without clouds

Technique	AAE	STD
Ju *et al.* [12]	2.16°	2.00°
Bab-Hadiashar–Suter [4]	2.05°	2.92°
Lai–Vemuri [13]	1.99°	1.41°
Our method (2D)	**1.59°**	**1.39°**
Mémin–Pérez [16]	1.58°	1.21°
Weickert *et al.* [24]	1.46°	1.50°
Farnebäck [10]	1.14°	2.14°
Our method (3D)	**0.98°**	**1.17°**

Table 1. Comparison between the results from the literature with 100 % density and our results for the *Yosemite* sequence with and without cloudy sky. AAE = average angular error. STD = standard deviation. 2D = spatial smoothness assumption. 3D = spatio-temporal smoothness assumption.

Source: T. Brox, A. Bruhn, N. Papenberg, J. Weickert: “High Accuracy Optical Flow Estimation Based on a Theory for Warping”

Evaluation

Benchmark results (2)

Yosemite with clouds		
σ_n	AAE	STD
0	1.94°	6.02°
10	2.50°	5.96°
20	3.12°	6.24°
30	3.77°	6.54°
40	4.37°	7.12°

Yosemite without clouds		
σ_n	AAE	STD
0	0.98°	1.17°
10	1.26°	1.29°
20	1.63°	1.39°
30	2.03°	1.53°
40	2.40°	1.71°

Table 2. Results for the *Yosemite* sequence with and without cloudy sky. Gaussian noise with varying standard deviations σ_n was added, and the average angular errors and their standard deviations were computed. AAE = average angular error. STD = standard deviation.

Source: T. Brox, A. Bruhn, N. Papenberg, J. Weickert: “High Accuracy Optical Flow Estimation Based on a Theory for Warping”

Evaluation

Benchmark results (3)

3D - spatio-temporal method

reduction factor η	outer fixed point iter.	inner fixed point iter.	SOR iter.	computation time/frame	AAE
0.95	77	5	10	23.4s	1.94°
0.90	38	2	10	5.1s	2.09°
0.80	18	2	10	2.7s	2.56°
0.75	14	1	10	1.2s	3.44°

Table 4. Computation times and convergence for Yosemite sequence with clouds.

Source: T. Brox, A. Bruhn, N. Papenberg, J. Weickert: “High Accuracy Optical Flow Estimation Based on a Theory for Warping”

Summary

The new variational approach

- uses non-linearised model assumptions
- linearisations in the numerical scheme
- coarse-to-fine warping applied for better approximation of the global minimum of the non-convex energy functional

Summary

Warping Methods

- if the optical flow constraint is linearised, the functional is easier to solve
- to deal with larger displacements, a coarse-to-fine warping is applied
- one can show that this strategy in fact solves a non-linearised optical flow constraint
- -> theoretical foundation for the warping strategy

References

- T. Brox, A. Bruhn, N. Papenberg, J. Weickert: "High Accuracy Optical Flow Estimation Based on a Theory for Warping", Proc. 8th European Conference on Computer Vision, vol. 4, pp. 25-36, Berlin, Heidelberg: Springer Verlag 2004.
- E. Mémin, P. Pérez: "Hierarchical estimation and segmentation of dense motion fields", International Journal of Computer Vision, vol. 46(2), pp. 129-155, 2002.
- A. Bruhn: "Correspondence Problems in Computer Vision", lecture notes, Summer Term 2007.
- B. Burgeth: "Image Processing and Computer Vision", lecture notes, Winter Term 2006/2007.