DAGM 2011 Tutorial on

Convex Optimization for Computer Vision

Part 3: Convex Solutions for Stereo and Optical Flow

Daniel Cremers
Computer Vision Group
Technical University of Munich

Graz University of Technology

Thomas Pock
Institute for Computer Graphics and Vision
Graz University of Technology

Frankfurt, August 30, 2011

Overview

1. Motion estimation

2. Stereo estimation

Motion estimation

- Motion estimation (optical flow) is a central topic in computer vision,
- Computes a 2D vector field, describing the motion of pixel intensities

Applications:

- Tracking
- Video compression, video interpolation
- 3D reconstruction

Motion estimation

- Motion estimation (optical flow) is a central topic in computer vision,
- Computes a 2D vector field, describing the motion of pixel intensities

Applications:

- Tracking
- Video compression, video interpolation
- 3D reconstruction

Motion estimation

- Motion estimation (optical flow) is a central topic in computer vision,
- Computes a 2D vector field, describing the motion of pixel intensities

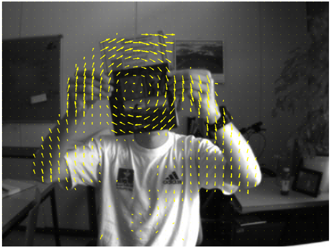

Applications:

- Tracking
- Video compression, video interpolation
- 3D reconstruction

Motion estimation

- Motion estimation (optical flow) is a central topic in computer vision,
- Computes a 2D vector field, describing the motion of pixel intensities

Applications:

- Tracking
- Video compression, video interpolation
- 3D reconstruction

Challenges

Motion estimation is still a very difficult problem

- Aperture problem

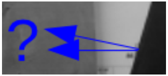

- No information in untextured areas
- Illumination changes, shadows, ...
- Large motion of small objects, occlusions, ...

Challenges

Motion estimation is still a very difficult problem

- Aperture problem

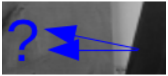

- No information in untextured areas
- Illumination changes, shadows, ...
- Large motion of small objects, occlusions, ...

Challenges

Motion estimation is still a very difficult problem

- Aperture problem
- No information in untextured areas

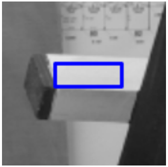

- Illumination changes, shadows, ...
- Large motion of small objects, occlusions, ...

Challenges

Motion estimation is still a very difficult problem

- Aperture problem
- No information in untextured areas

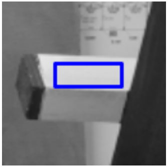

- Illumination changes, shadows, ...
- Large motion of small objects, occlusions, ...

Challenges

Motion estimation is still a very difficult problem

- Aperture problem
- No information in untextured areas
- Illumination changes, shadows, ...

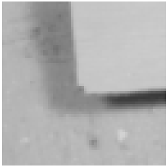

- Large motion of small objects, occlusions, ...

Challenges

Motion estimation is still a very difficult problem

- Aperture problem
- No information in untextured areas
- Illumination changes, shadows, ...

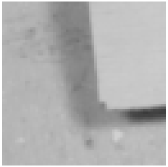

- Large motion of small objects, occlusions, ...

Challenges

Motion estimation is still a very difficult problem

- Aperture problem
- No information in untextured areas
- Illumination changes, shadows, ...
- Large motion of small objects, occlusions, ...

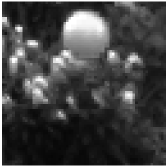

Challenges

Motion estimation is still a very difficult problem

- Aperture problem
- No information in untextured areas
- Illumination changes, shadows, ...
- Large motion of small objects, occlusions, ...

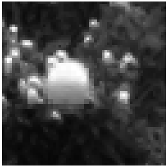

The correspondence problem

- Find corresponding points in successive frames

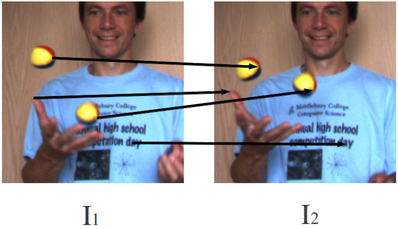

- Brightness (color) constancy assumption

$$I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}(\mathbf{x})) \approx 0$$

- $\mathbf{u}(\mathbf{x}) = (u_1(\mathbf{x}), u_2(\mathbf{x}))$ is the displacement vector
- Ambiguity: Many points with similar brightness (color)!
- Generalization: Constancy of image features (gradients, NCC, ...)

Variational motion estimation

- Generic variational model for motion estimation

$$\min_{\mathbf{u}} \underbrace{\mathcal{R}(\mathbf{u})}_{\text{Regularization term}} + \underbrace{\int_\Omega |I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}(\mathbf{x}))|^p \, \mathrm{d}\mathbf{x}}_{\text{Data term}}$$

Variational motion estimation

- Generic variational model for motion estimation

$$\min_{\mathbf{u}} \underbrace{\mathcal{R}(\mathbf{u})}_{\text{Regularization term}} + \underbrace{\int_\Omega |I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}(\mathbf{x}))|^p \, \mathrm{d}\mathbf{x}}_{\text{Data term}}$$

- Regularization term:
  - Should favor physically meaningful flow fields
  - Popular convex regularizers: Quadratic, total variation, ...

Variational motion estimation

- Generic variational model for motion estimation

$$\min_{\mathbf{u}} \underbrace{\mathcal{R}(\mathbf{u})}_{\text{Regularization term}} + \underbrace{\int_\Omega |I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}(\mathbf{x}))|^p \, d\mathbf{x}}_{\text{Data term}}$$

- Regularization term:
  - Should favor physically meaningful flow fields
  - Popular convex regularizers: Quadratic, total variation, ...
- Data term:
  - Highly non-convex $\rightarrow$ hard to minimize
  - Different strategies to deal with the non-convexity of the data term

Variational motion estimation

- Generic variational model for motion estimation

$$\min_{\mathbf{u}} \underbrace{\mathcal{R}(\mathbf{u})}_{\text{Regularization term}} + \underbrace{\int_{\Omega} |I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}(\mathbf{x}))|^p \, d\mathbf{x}}_{\text{Data term}}$$

- Regularization term:
  - Should favor physically meaningful flow fields
  - Popular convex regularizers: Quadratic, total variation, ...
- Data term:
  - Highly non-convex → hard to minimize
  - Different strategies to deal with the non-convexity of the data term
- Vast literature on motion estimation:
  - Window based optical flow: [Lucas, Kanade, 1981]
  - Variational optical flow: [Horn, Schunck, 1981]
  - Discontinuity preserving optical flow: [Shulman, Hervé '89]
  - Robust optical flow: [Black, Anadan, '93]
  - Highly accurate optical flow: [Brox, Bruhn, Papenberg, Weickert '04]
  - Real-time optical flow: [A. Bruhn, J. Weickert, T. Kohlberger, C. Schnörr '05]
  - Primal-dual optimization on the GPU: [Zach, Pock, Bischof '07]

Linearization of the image

- Perform a first order Taylor expansion of the function $I_2(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ at $\mathbf{x} + \mathbf{u}_0(\mathbf{x})$ [Horn, Schunck, 1981], [Lucas, Kanade, 1981]

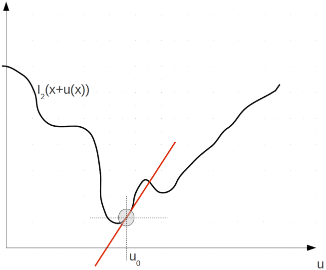

Linearization of the image

- Perform a first order Taylor expansion of the function $I_2(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ at $\mathbf{x} + \mathbf{u}_0(\mathbf{x})$ [Horn, Schunck, 1981], [Lucas, Kanade, 1981]

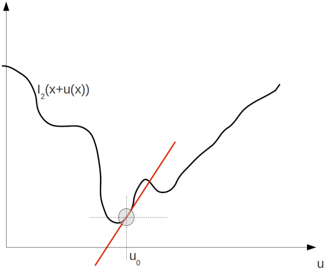

- $I_2(\mathbf{x} + \mathbf{u}(\mathbf{x})) \approx I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) + \langle \nabla I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle$
- Only valid close to $\mathbf{u}_0$, i.e. $\|\mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x})\| \leq \varepsilon$

Linearization of the image

- Perform a first order Taylor expansion of the function $I_2(\mathbf{x}+\mathbf{u}(\mathbf{x}))$ at $\mathbf{x}+\mathbf{u}_0(\mathbf{x})$ [Horn, Schunck, 1981], [Lucas, Kanade, 1981]

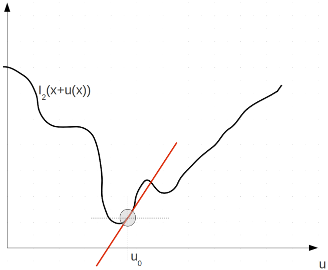

- $I_2(\mathbf{x}+\mathbf{u}(\mathbf{x})) \approx I_2(\mathbf{x}+\mathbf{u}_0(\mathbf{x})) + \langle \nabla I_2(\mathbf{x}+\mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle$
- Only valid close to $\mathbf{u}_0$, i.e. $\|\mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x})\| \leq \varepsilon$

- Leads to the classical optical flow constraint:

$$\rho(\mathbf{u}) = I_1(\mathbf{x}) - I_2(\mathbf{x}+\mathbf{u}_0(\mathbf{x})) - \langle \nabla I_2(\mathbf{x}+\mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle \approx 0$$

- Note: $\rho(\mathbf{u})$ is linear in $\mathbf{u}$ and hence $|\rho(\mathbf{u})|$ is convex!

TV-L^1 motion estimation

- It turns out that total variation regularization in combination with a L^1 data term performs well
- Total variation allows for motion discontinuities
- L^1 data term allows for outliers in the data term (occlusions, noise, ...)

$$\min_{\|\mathbf{u}-\mathbf{u}_0\|\leq\varepsilon} \alpha \int_\Omega |D\mathbf{u}| + \|\rho(\mathbf{u})\|_1$$

- Non-differentiable and hence difficult to solve

TV-L^1 motion estimation

- It turns out that total variation regularization in combination with a L^1 data term performs well
- Total variation allows for motion discontinuities
- L^1 data term allows for outliers in the data term (occlusions, noise, ...)

$$\min_{\|\mathbf{u}-\mathbf{u}_0\|\leq\varepsilon} \alpha \int_\Omega |D\mathbf{u}| + \|\rho(\mathbf{u})\|_1$$

- Non-differentiable and hence difficult to solve

- Smoothing and fixed-point iteration: [Brox, Bruhn, Papenberg, Weickert '04]
- Primal-dual optimization: [Chambolle, Pock, '10]

$$\min_{\|\mathbf{u}-\mathbf{u}_0\|\leq\varepsilon} \max_{\|\mathbf{p}\|_\infty\leq\alpha} -\int_\Omega \mathbf{u}\,\mathrm{div}\,\mathbf{p}\;\mathrm{d}\mathbf{x} + \|\rho(\mathbf{u})\|_1$$

- Allows to compute the exact solution

Second-order approximation of the data term

- Consider a more general non-convex data term of the form

$$\int_\Omega \phi(x, \mathbf{u}(x)) \, \mathrm{d}x$$

- Perform a second order Taylor expansion of the data term $\phi(x, \mathbf{u}(x))$ around $\mathbf{u}_0(x)$ [Werlberger, Pock, Bischof '10]

$$\phi(x, \mathbf{u}(x)) \approx \phi(x, \mathbf{u}_0(x)) + (\nabla\phi(x, \mathbf{u}_0(x)))^{\mathrm{T}} (\mathbf{u}(x) - \mathbf{u}_0(x)) + (\mathbf{u}(x) - \mathbf{u}_0(x))^{\mathrm{T}} (\nabla^2\phi(x, \mathbf{u}_0(x))) (\mathbf{u}(x) - \mathbf{u}_0(x)),$$

- To ensure convexity the Hessian $\nabla^2\phi(x, \mathbf{u}_0(x))$ has to be positive semidefinite
- We use the following diagonal approximation of the Hessian

$$\nabla^2\phi = \begin{bmatrix} (\phi_{xx}(x, u_0(x)))^+ & 0 \\ 0 & (\phi_{yy}(x, u_0(x)))^+ \end{bmatrix}$$

- Can be used with arbitrary data terms: SAD, NCC, ...
- Still only valid in a small neighborhood around $\mathbf{u}_0$
- Minimization using primal-dual schemes

Large displacements

- How can we compute large displacements?
- Integrate the algorithm in a coarse-to fine / warping framework

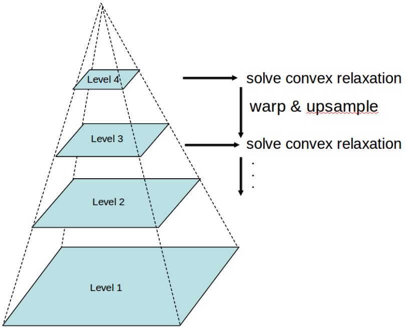

- Similar to multigrid schemes, speeds up the minimization process
- Does not give any guarantees!

Large displacement optical flow without warping

- Consider the following equivalent generic formulation [Steinbrücker, Pock, Cremers, '09]

$$\min_{\mathbf{u}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{u}) \, \mathrm{d}\mathbf{x} \quad \Longleftrightarrow \quad \min_{\mathbf{u},\mathbf{v}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{v}) \, \mathrm{d}\mathbf{x} \quad \text{s.t.} \quad \mathbf{u} = \mathbf{v}$$

Large displacement optical flow without warping

- Consider the following equivalent generic formulation [Steinbrücker, Pock, Cremers, '09]

$$\min_{\mathbf{u}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{u}) \, d\mathbf{x} \quad \Longleftrightarrow \quad \min_{\mathbf{u},\mathbf{v}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{v}) \, d\mathbf{x} \quad \text{s.t.} \quad \mathbf{u} = \mathbf{v}$$

- Quadratic penality approach to obtain a unconstrained formulation

$$\min_{\mathbf{u},\mathbf{v}} \mathcal{R}(\mathbf{u}) + \frac{1}{2\theta}\|\mathbf{u} - \mathbf{v}\|_2^2 + \int_\Omega \phi(\mathbf{v}) \, d\mathbf{x}$$

- Becomes equivalent to the constrained formulation for $\theta \to 0^+$

Large displacement optical flow without warping

- Consider the following equivalent generic formulation [Steinbrücker, Pock, Cremers, '09]

$$\min_{\mathbf{u}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{u})\ \mathrm{d}\mathbf{x} \quad \Longleftrightarrow \quad \min_{\mathbf{u},\mathbf{v}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{v})\ \mathrm{d}\mathbf{x} \quad \text{s.t.} \quad \mathbf{u} = \mathbf{v}$$

- Quadratic penality approach to obtain a unconstrained formulation

$$\min_{\mathbf{u},\mathbf{v}} \mathcal{R}(\mathbf{u}) + \frac{1}{2\theta}\|\mathbf{u}-\mathbf{v}\|_2^2 + \int_\Omega \phi(\mathbf{v})\ \mathrm{d}\mathbf{x}$$

- Becomes equivalent to the constrained formulation for $\theta \to 0^+$
- Observations:
  - Solution with respect to **u** reduces to an image denoising problem
  - Solution with respect to **v** reduces to pointwise non-convex problems

Large displacement optical flow without warping

- Consider the following equivalent generic formulation [Steinbrücker, Pock, Cremers, '09]

$$\min_{\mathbf{u}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{u})\, \mathrm{d}\mathbf{x} \quad \Longleftrightarrow \quad \min_{\mathbf{u},\mathbf{v}} \mathcal{R}(\mathbf{u}) + \int_\Omega \phi(\mathbf{v})\, \mathrm{d}\mathbf{x} \quad \text{s.t.} \quad \mathbf{u} = \mathbf{v}$$

- Quadratic penality approach to obtain a unconstrained formulation

$$\min_{\mathbf{u},\mathbf{v}} \mathcal{R}(\mathbf{u}) + \frac{1}{2\theta}\|\mathbf{u} - \mathbf{v}\|_2^2 + \int_\Omega \phi(\mathbf{v})\, \mathrm{d}\mathbf{x}$$

- Becomes equivalent to the constrained formulation for $\theta \to 0^+$
- Observations:
  - Solution with respect to **u** reduces to an image denoising problem
  - Solution with respect to **v** reduces to pointwise non-convex problems
- Annealing-type scheme: Alternating minimization for a sequence of decreasing parameters θ_i
- **Advantages:** No coarse-to-fine, no warping, arbitrary data terms
- **Disadvantage:** Results strongly depend on the sequence θ_i

Graz University of Technology

Varying illumination

- Note: optical flow $\neq$ motion estimation!

Varying illumination

- Note: optical flow $\neq$ motion estimation!
- Instead of (or in combination with) the original images $I_{1,2}$ use gradient images $\nabla I_{1,2}$ [Brox, Bruhn, Papenberg, Weickert '04]

Varying illumination

- Note: optical flow $\neq$ motion estimation!
- Instead of (or in combination with) the original images $I_{1,2}$ use gradient images $\nabla I_{1,2}$ [Brox, Bruhn, Papenberg, Weickert '04]
- Photometric invariants: [Mileva, Bruhn, Weickert '07]

Varying illumination

- Note: optical flow $\neq$ motion estimation!
- Instead of (or in combination with) the original images $I_{1,2}$ use gradient images $\nabla I_{1,2}$ [Brox, Bruhn, Papenberg, Weickert '04]
- Photometric invariants: [Mileva, Bruhn, Weickert '07]
- Structure-texture decomposition: Illumination changes and shadows correspond to large image features [Wedel, Pock, Zach, Bischof, Cremers '08]

Varying illumination

- Note: optical flow $\neq$ motion estimation!
- Instead of (or in combination with) the original images $I_{1,2}$ use gradient images $\nabla I_{1,2}$ [Brox, Bruhn, Papenberg, Weickert '04]
- Photometric invariants: [Mileva, Bruhn, Weickert '07]
- Structure-texture decomposition: Illumination changes and shadows correspond to large image features [Wedel, Pock, Zach, Bischof, Cremers '08]
- Addititve decomposition using the ROF model:

$$I = S + T \;, \quad S := \arg\min_u \mathrm{TV}(u) + \frac{\lambda}{2}\|u - I\|_2^2$$

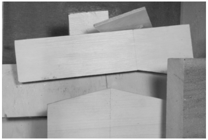

(a) I

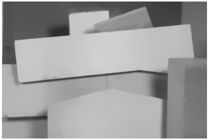

(b) S

(c) T

- Use texture component T to compute the optical flow

Modified optical flow constraint

- Recall the optical flow constraint

$$\rho(\mathbf{u}) = I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) - \langle \nabla I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle \approx 0$$

Modified optical flow constraint

- Recall the optical flow constraint

$$\rho(\mathbf{u}) = I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) - \langle \nabla I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle \approx 0$$

- We can modify the constraint [Shulman, Hervé '89]

$$\delta(\mathbf{u}, v) = I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) - \langle \nabla I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle - v(\mathbf{x}) \approx 0$$

- $v(\mathbf{x})$ is a smooth function modeling illumination changes
- Note that $\delta(\mathbf{u}, v)$ is still linear in $\mathbf{u}$ and v!

Modified optical flow constraint

- Recall the optical flow constraint

$$\rho(\mathbf{u}) = I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) - \langle \nabla I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle \approx 0$$

- We can modify the constraint [Shulman, Hervé '89]

$$\delta(\mathbf{u}, v) = I_1(\mathbf{x}) - I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) - \langle \nabla I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle - v(\mathbf{x}) \approx 0$$

- $v(\mathbf{x})$ is a smooth function modeling illumination changes
- Note that $\delta(\mathbf{u}, v)$ is still linear in $\mathbf{u}$ and v!

- Additional regularization needed for $v(\mathbf{x})$

$$\min_{\|\mathbf{u}-\mathbf{u}_0\| \leq \varepsilon, v} \alpha \int_\Omega |D\mathbf{u}| + \beta \int_\Omega |Dv| + \|\delta(\mathbf{u}(\mathbf{x}), v(\mathbf{x}))\|_1$$

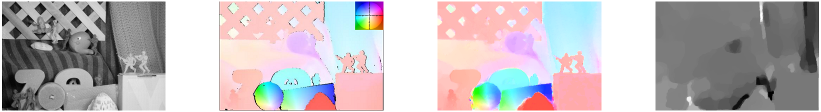

(a) Input (b) Ground truth (c) Estimated motion (d) Illumination

Overview

1. Motion estimation

2. Stereo estimation

Stereo

- If I_1 and I_2 come from a stereo camera or a moving camera that browses a static scene, the displacement can be restricted to 1D problems on the epipolar lines, [Slesareva, Bruhn, Weickert '05]
- Each stereo pair can be normalized such that the displacement is only horizontally
- The depth z can be computed from the displacement u via

$$z(x,y) = \frac{bf}{u(x,y)}$$

where b is the baseline and f is the focal length of the camera

Stereo

- If I_1 and I_2 come from a stereo camera or a moving camera that browses a static scene, the displacement can be restricted to 1D problems on the epipolar lines, [Slesareva, Bruhn, Weickert '05]
- Each stereo pair can be normalized such that the displacement is only horizontally
- The depth z can be computed from the displacement u via

$$z(x,y) = \frac{bf}{u(x,y)}$$

where b is the baseline and f is the focal length of the camera

- Optical flow constraint for stereo

$$\hat{\rho}(u) = I_1 - I_2(x + u_0(x,y), y) - \partial_x I_2(x + u_0(x,y), y)(u(x,y) - u_0(x,y)) \approx 0$$

- TV-L^1 based stereo

$$\min_{\|u-u_0\| \leq \varepsilon} \alpha \int_\Omega |Du| + \|\hat{\rho}(u(x,y))\|_1$$

Stereo

- If I_1 and I_2 come from a stereo camera or a moving camera that browses a static scene, the displacement can be restricted to 1D problems on the epipolar lines, [Slesareva, Bruhn, Weickert '05]
- Each stereo pair can be normalized such that the displacement is only horizontally
- The depth z can be computed from the displacement u via

$$z(x,y) = \frac{bf}{u(x,y)}$$

  where b is the baseline and f is the focal length of the camera

- Optical flow constraint for stereo

$$\hat{\rho}(u) = I_1 - I_2(x + u_0(x,y), y) - \partial_x I_2(x + u_0(x,y), y)(u(x,y) - u_0(x,y)) \approx 0$$

- TV-L^1 based stereo

$$\min_{\|u-u_0\| \leq \varepsilon} \alpha \int_\Omega |Du| + \|\hat{\rho}(u(x,y))\|_1$$

- Advantages
  - Highly accurate due to sub-pixel accuracy
  - fast to compute (real-time)
- Disadvantages
  - Does not compute the globally optimal solution (coarse-to-fine)

Application: range estimation in a driving car (with Daimler AG)

- Input images provided by a calibrated stereo rig

(a) Left image

(b) Right image

- Range image computed by the TV-L^1 based stereo algorithm

(a) Range image

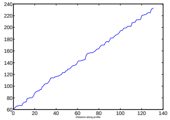

(b) Profile of street

- Total variation regularization leads to the staircasing effect!

Total generalized variation

- The total variation can be written (via the convex conjugate) as

$$\mathrm{TV}_\alpha(u) = \alpha \int_\Omega |Du| = \sup \left\{ \int_\Omega u \operatorname{div} v \, \mathrm{d}x \;\middle|\; v \in \mathcal{C}_c^1(\Omega, \mathbb{R}^d), \|v\|_\infty \le \alpha \right\},$$

Total generalized variation

- The total variation can be written (via the convex conjugate) as

$$\mathrm{TV}_\alpha(u) = \alpha \int_\Omega |Du| = \sup \left\{ \int_\Omega u \operatorname{div} v \, \mathrm{d}x \;\middle|\; v \in \mathcal{C}^1_c(\Omega, \mathbb{R}^d), \|v\|_\infty \leq \alpha \right\},$$

- In [Bredies, Kunisch, Pock, SIIMS'10], we proposed a generalization of the total variation to higher order smoothness.

$$\mathrm{TGV}^k_\alpha(u) = \sup \left\{ \int_\Omega u \operatorname{div}^k v \, \mathrm{d}x \;\middle|\; v \in \mathcal{C}^k_c(\Omega, \mathrm{Sym}^k(\mathbb{R}^d)), \right. \\ \left. \|\operatorname{div}^l v\|_\infty \leq \alpha_l, \; l = 0, \ldots, k-1 \right\},$$

Total generalized variation

- The total variation can be written (via the convex conjugate) as

$$\mathrm{TV}_\alpha(u) = \alpha \int_\Omega |Du| = \sup \left\{ \int_\Omega u \,\mathrm{div}\, v \;\mathrm{d}x \;\middle|\; v \in \mathcal{C}_c^1(\Omega, \mathbb{R}^d), \|v\|_\infty \le \alpha \right\},$$

- In [Bredies, Kunisch, Pock, SIIMS'10], we proposed a generalization of the total variation to higher order smoothness.

$$\mathrm{TGV}_\alpha^k(u) = \sup \left\{ \int_\Omega u \,\mathrm{div}^k\, v \;\mathrm{d}x \;\middle|\; v \in \mathcal{C}_c^k(\Omega, \mathrm{Sym}^k(\mathbb{R}^d)), \right.$$
$$\left. \|\mathrm{div}^l\, v\|_\infty \le \alpha_l,\; l = 0, \dots, k-1 \right\},$$

- For $k = 2$ it can be written as

$$\mathrm{TGV}_\alpha^2(u) = \inf_{\mathbf{w}} \alpha_1 \int_\Omega |Du - \mathbf{w}| + \alpha_0 \int_\Omega |D\mathbf{w}|$$

- TGV^2 can be used to reconstruct piecewise affine functions

Image restoration examples

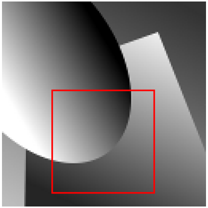

(a) Clean image

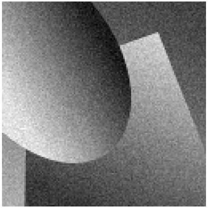

(b) Noisy image

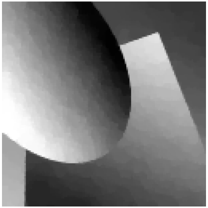

(c) TV

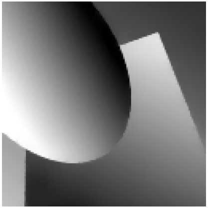

(d) TGV^2

Image restoration examples

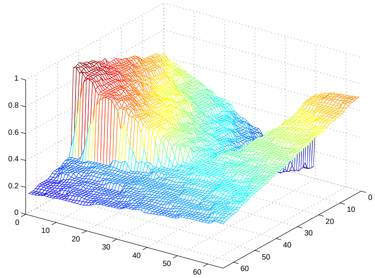

(a) TV

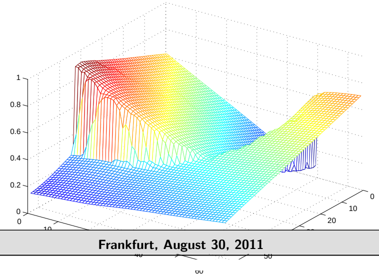

TGV based stereo

- Simply replace TV regularization by TGV regularization in the stereo model [Ranftl, Pock, Gehrig, Franke '11]

$$\min_{\|u-u_0\|\leq\varepsilon,\mathbf{w}} \alpha_1 \int_\Omega |Du - \mathbf{w}| + \alpha_0 \int_\Omega |D\mathbf{w}| + \|\hat{\rho}(u(x,y))\|_1$$

- Comparison on the stereo problem

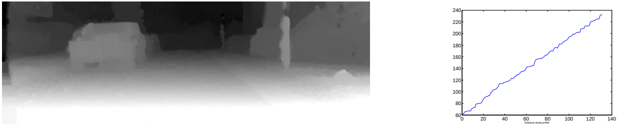

(a) TV

(b) TGV^2

Range estimation from a driving car

Summary and open questions

- Introduced the problem of motion estimation in computer vision
- Motion estimation is still a challenging problem, not near to be solved
- Highly non-convex data term leads to numerical difficulties
- A simple linearization approach works well in practice
- Can be used for stereo estimation
- TGV regularization avoids staircasing-artifacts

Summary and open questions

- Introduced the problem of motion estimation in computer vision
- Motion estimation is still a challenging problem, not near to be solved
- Highly non-convex data term leads to numerical difficulties
- A simple linearization approach works well in practice
- Can be used for stereo estimation
- TGV regularization avoids staircasing-artifacts

- Global Solutions for Motion and Stereo?