Convex Methods for Dense Motion, Stereo and 3D

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

1st Live Dense Reconstruction Workshop, ICCV 2011, Barcelona

Computer vision deals with inverse problems

- Projection of the 3D world onto the 2D image plane

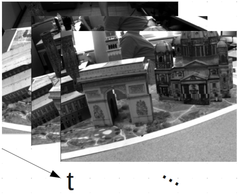

- Determine unknown model parameters based on observed data

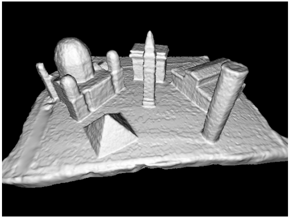

Computer vision deals with inverse problems

- Projection of the 3D world onto the 2D image plane

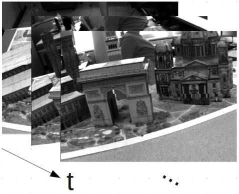

- Determine unknown model parameters based on observed data

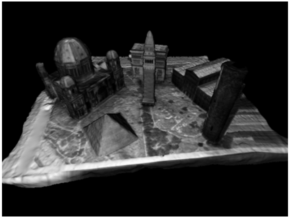

Computer vision is highly ambiguous

What you see ...

Computer vision is highly ambiguous

What you see ... is maybe not what it is!

[Fukuda's Underground Piano Illusion]

Energy minimization methods

- It is in general not possible to solve inverse problems directly
- Add some smoothness assumption to the unknown solution
- Leads to the energy minimization approach

Energy minimization methods

- It is in general not possible to solve inverse problems directly
- Add some smoothness assumption to the unknown solution
- Leads to the energy minimization approach

$$\boxed{\min_u \mathcal{R}(u) + \mathcal{D}(u, f)}$$

- Energy functional is designed such that low-energy states reflect the physical properties of the problem
- Minimizer provides the best (in the sense of the model) solution to the problem

Energy minimization methods

- It is in general not possible to solve inverse problems directly
- Add some smoothness assumption to the unknown solution
- Leads to the energy minimization approach

$$\boxed{\min_u \mathcal{R}(u) + \mathcal{D}(u, f)}$$

- Energy functional is designed such that low-energy states reflect the physical properties of the problem
- Minimizer provides the best (in the sense of the model) solution to the problem

- Different philosophies:
  - Continuous Setting: Variational methods
  - Discrete Setting: Markov/Conditional random fields

Variational motion estimation

- 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 et al. '04]
  - Real-time optical flow: [Bruhn et al. '05]
  - Real-time optical flow on the GPU: [Zach et al. '07]

Variational motion estimation

- 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 et al. '04]
  - Real-time optical flow: [Bruhn et al. '05]
  - Real-time optical flow on the GPU: [Zach et al. '07]

- Generic variational model for motion estimation

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

where $p, q \geq 1$, $I_{1,2}$ are the input images and $\mathbf{u} = (u_1, u_2)$ is the flow field

Variational motion estimation

- 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 et al. '04]
  - Real-time optical flow: [Bruhn et al. '05]
  - Real-time optical flow on the GPU: [Zach et al. '07]

- Generic variational model for motion estimation

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

  where $p, q \geq 1$, $I_{1,2}$ are the input images and $\mathbf{u} = (u_1, u_2)$ is the flow field
- Regularization term:
  - Should favor physically meaningful flow fields
  - Popular convex regularizers: Quadratic, TV, TGV, ...

Variational motion estimation

- 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 et al. '04]
  - Real-time optical flow: [Bruhn et al. '05]
  - Real-time optical flow on the GPU: [Zach et al. '07]
- Generic variational model for motion estimation

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

  where $p, q \geq 1$, $I_{1,2}$ are the input images and $\mathbf{u} = (u_1, u_2)$ is the flow field
- Regularization term:
  - Should favor physically meaningful flow fields
  - Popular convex regularizers: Quadratic, TV, TGV, ...
- Data term:
  - Highly non-convex → hard to minimize
  - Different strategies to deal with the non-convexity of the data term

Linearization

- 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]

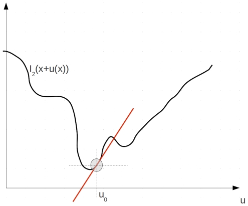

Linearization

- 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]

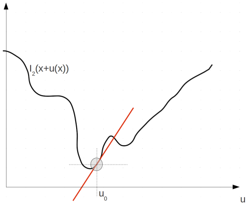

- $I_2(\mathbf{x}+\mathbf{u}(\mathbf{x})) \approx I_2(\mathbf{x}+\mathbf{u}_0(\mathbf{x})) + \langle \nabla I_2(\mathbf{x}+\mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x})\rangle$
- Only valid close to $\mathbf{u}_0$, i.e. $\|\mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x})\|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

Graz University of Technology

Linearization

- 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]

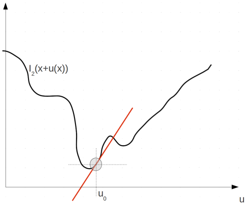

- $I_2(\mathbf{x} + \mathbf{u}(\mathbf{x})) \approx I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})) + \langle \nabla I_2(\mathbf{x} + \mathbf{u}_0(\mathbf{x})), \mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x}) \rangle$
- Only valid close to $\mathbf{u}_0$, i.e. $\|\mathbf{u}(\mathbf{x}) - \mathbf{u}_0(\mathbf{x})\|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

- 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 $|\rho(\mathbf{u})|_q^q$ is convex!

TV-L^1 motion estimation

- It turns out that total variation regularization in combination with a L^1 data term performs well
- In the discrete setting it is written as

$$\min_{\mathbf{u}} \alpha\|\nabla \mathbf{u}\|_1 + \|\rho(\mathbf{u})\|_1$$

- Total variation allows for motion discontinuities
- L^1 data term allows for outliers in the data term (occlusions, noise, ...)
- Thanks to the linearization, we have a convex problem
- 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
- In the discrete setting it is written as

$$\min_{\mathbf{u}} \alpha \|\nabla \mathbf{u}\|_1 + \|\rho(\mathbf{u})\|_1$$

- Total variation allows for motion discontinuities
- L^1 data term allows for outliers in the data term (occlusions, noise, ...)
- Thanks to the linearization, we have a convex problem
- Non-differentiable and hence difficult to solve

- Different approaches
  - Smoothing and fixed point iterations [Brox et al. '04]
  - Primal-dual optimization [Chambolle, Pock, '10]

Smoothing and fixed point iterations

- Replace ℓ_1 norms by a C^∞ approximation, e.g. $|t|_\varepsilon = \sqrt{t^2 + \varepsilon^2}$

$$\min_{\mathbf{u}} \alpha \|\nabla \mathbf{u}\|_{1,\varepsilon} + \|\rho(\mathbf{u})\|_{1,\varepsilon}$$

- Compute gradient (Euler Lagrange equation in the functional setting)

$$\nabla^T \mathrm{diag}(1/|\nabla \mathbf{u}|_\varepsilon)\nabla \mathbf{u} + \mathrm{diag}(1/|\rho(\mathbf{u})|_\varepsilon)\rho(\mathbf{u})\rho' = 0$$

- Solve via a fixed point iteration

$$\nabla^T \mathrm{diag}(1/|\nabla \mathbf{u}^k|_\varepsilon)\nabla \mathbf{u}^{k+1} + \mathrm{diag}(1/|\rho(\mathbf{u}^k)|_\varepsilon)\rho(\mathbf{u}^{k+1})\rho' = 0$$

- The method requires to solve a sequence of systems of linear equations
- Can be ill-conditioned for small ε and hence slow
- Sophisticated multigrid-algorithm achieves realtime performance [Bruhn et al. '05]

Primal-dual optimization

- Use concepts of convex optimization to rewrite the original problem as a convex-concave saddle-point problem

$$\min_{\mathbf{u}} \alpha\|\nabla \mathbf{u}\|_1 + \|\rho(\mathbf{u})\|_1 \quad \Longleftrightarrow \quad \min_{\mathbf{u}} \max_{\mathbf{p}} \langle \nabla \mathbf{u}, \mathbf{p}\rangle + \|\rho(\mathbf{u})\|_1 - \iota_{\|\mathbf{p}\|_\infty \leq \alpha}(\mathbf{p})$$

- Can be solved exactly using primal-dual optimization [Chambolle, Pock, '10]
- Simple iterative scheme

$$\begin{cases} \mathbf{u}^{k+1} = \mathrm{shrink}_{\rho(\mathbf{u})}\left(\mathbf{u}^k - \tau(\nabla^T \mathbf{p}^k)\right) \\ \mathbf{p}^{k+1} = \mathrm{proj}_{\|\mathbf{p}\|_\infty \leq \alpha}\left(\mathbf{p}^k + \sigma(\nabla(2\mathbf{u}^{k+1} - \mathbf{u}^k))\right) \end{cases},$$

- Guaranteed convergence rate of $\mathcal{O}(1/k)$
- Real-time performance on the GPU

Large displacements

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

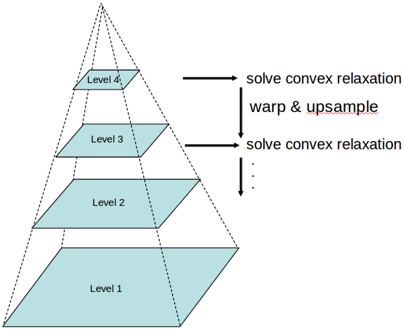

- Similar to multigrid schemes, speeds up the minimization process
- Does not give any guarantees!

[Werlberger et al. '11]

Varying illumination

- Recall the optical flow constraint

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

Varying illumination

- 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]

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

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

Varying illumination

- 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]

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

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

$$\min_{\mathbf{u},s} \alpha \|\nabla \mathbf{u}\|_1 + \beta \|\nabla s\|_1 + \|\delta(\mathbf{u}, s)\|_1$$

Varying illumination

- 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]

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

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

$$\min_{\mathbf{u},s} \alpha \|\nabla \mathbf{u}\|_1 + \beta \|\nabla s\|_1 + \|\delta(\mathbf{u}, s)\|_1$$

(a) Input

Varying illumination

- 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]

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

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

$$\min_{\mathbf{u},s} \alpha \|\nabla \mathbf{u}\|_1 + \beta \|\nabla s\|_1 + \|\delta(\mathbf{u}, s)\|_1$$

(a) Input

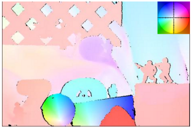

(b) Ground truth

Varying illumination

- 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]

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

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

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

$$\min_{\mathbf{u}, s} \alpha \|\nabla \mathbf{u}\|_1 + \beta \|\nabla s\|_1 + \|\delta(\mathbf{u}, s)\|_1$$


(a) Input

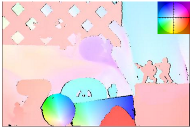
(b) Ground truth

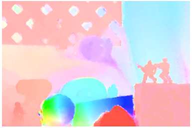
(c) Estimated motion

Varying illumination

- 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]

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

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

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

$$\min_{\mathbf{u},s} \alpha \|\nabla \mathbf{u}\|_1 + \beta \|\nabla s\|_1 + \|\delta(\mathbf{u}, s)\|_1$$

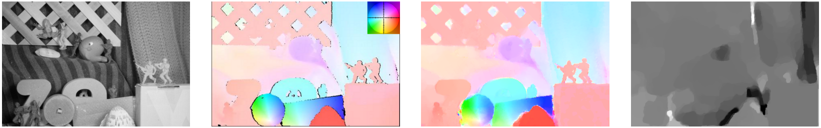

(a) Input (b) Ground truth (c) Estimated motion (d) Illumination

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

$$\min_u \alpha\|\nabla u\|_1 + \|\hat{\rho}(u)\|_1$$

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

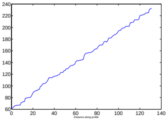

(b) Profile of street

- Total variation regularization leads to the staircasing effect!

Total generalized variation

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

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

Total generalized variation

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

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

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

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

Total generalized variation

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

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

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

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

- 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

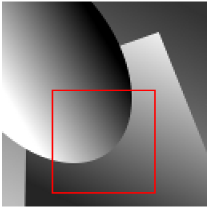

(a) Clean image

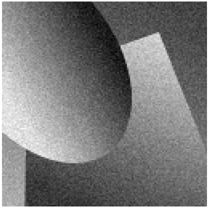

(b) Noisy image

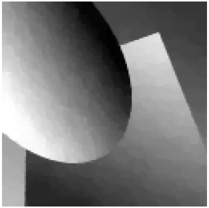

(c) TV

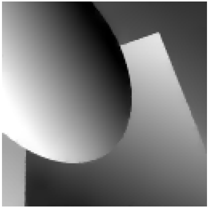

(d) TGV^2

Image restoration examples

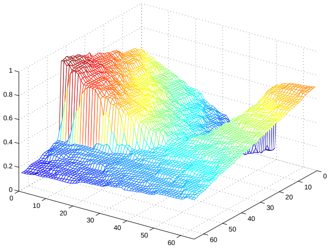

(a) TV

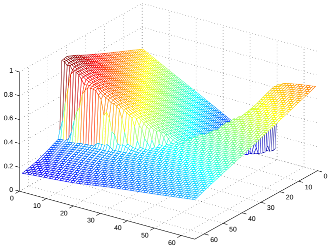

(b) TGV^2

TGV based stereo

- Simply replace TV regularization by TGV regularization in the stereo model

$$\min_{u,\mathbf{w}} \alpha_1 \|\nabla u - \mathbf{w}\|_1 + \alpha_0 \|\nabla \mathbf{w}\|_1 + \|\hat{\rho}(u)\|_1$$

- Comparison on the stereo problem

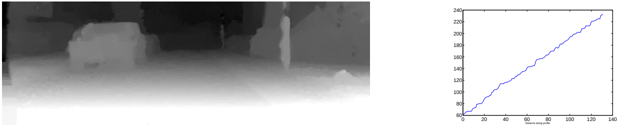

(a) TV

(b) TGV^2

Variational shape from focus (with Alicona Imaging)

- Record an image sequence by varying the focus of an imaging system

- Variational model that computes a piecewise smooth surface which maximizes a certain sharpness measure $\sigma(x, u)$ on the surface

$$\min_{u,\mathbf{w}} \alpha_1 \|\nabla u - \mathbf{w}\|_1 + \alpha_0 \|\nabla \mathbf{w}\|_1 - \sum_x \sigma(x, u)$$

- Computation takes a few seconds on the GPU

Global solutions of non-convex variational models

- Consider the following non-convex energy-functional

$$\min_u \int_\Omega f(x, u(x), \nabla u(x)) \, \mathrm{d}x$$

- We assume that $f(x, t, p)$ is convex in p but non-convex in t

Global solutions of non-convex variational models

- Consider the following non-convex energy-functional

$$\min_u \int_\Omega f(x, u(x), \nabla u(x))\,\mathrm{d}x$$

- We assume that $f(x, t, p)$ is convex in p but non-convex in t
- Example: TV-L^1 stereo

$$f(x, u(x), \nabla u(x)) = \alpha|\nabla u| + |I_1(x) - I_2(x + u(x))|$$

- Can we compute a global minimizer of this problem?

Global solutions of non-convex variational models

- Consider the following non-convex energy-functional

$$\min_u \int_\Omega f(x, u(x), \nabla u(x)) \, \mathrm{d}x$$

- We assume that $f(x, t, p)$ is convex in p but non-convex in t
- Example: TV-L^1 stereo

$$f(x, u(x), \nabla u(x)) = \alpha |\nabla u| + |I_1(x) - I_2(x + u(x))|$$

- Can we compute a global minimizer of this problem?
- In a discrete MRF setting, a solution has been proposed by [Ishikawa, '03] by a graph cut on a higher-dimensional graph

Global solutions of non-convex variational models

- Consider the following non-convex energy-functional

$$\min_u \int_\Omega f(x, u(x), \nabla u(x)) \, \mathrm{d}x$$

- We assume that $f(x, t, p)$ is convex in p but non-convex in t
- Example: TV-L^1 stereo

$$f(x, u(x), \nabla u(x)) = \alpha |\nabla u| + |I_1(x) - I_2(x + u(x))|$$

- Can we compute a global minimizer of this problem?
- In a discrete MRF setting, a solution has been proposed by [Ishikawa, '03] by a graph cut on a higher-dimensional graph
- What about the continuous setting?
- [Pock, Cremers, Bischof, Chambolle, SIIMS'10]

The approach of Alberti, Bouchitte and Dal Maso

- The calibration method of [Alberti, Bouchitte, Dal Maso, '03], was originally developed for the Mumford-Shah functional
- The basic idea is to consider the graph Γ_u of u instead of the function u
- Rewrite $E(u)$ by means of the flux of vector field ϕ through the graph Γ_u

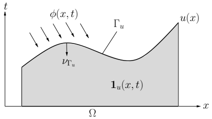

- The characteristic function $\mathbf{1}_u$ of the subgraph of a function $u \in \mathcal{BV}(\Omega \times \mathbb{R}, [0,1])$ is defined as

$$\mathbf{1}_u(x,t) = \begin{cases} 1, & \text{if } t < u(x), \\ 0, & \text{else.} \end{cases}$$

- The normal ν_{Γ_u} of the interface Γ_u is given by

$$\nu_{\Gamma_u} = \frac{(\nabla u, -1)}{\sqrt{|\nabla u|^2 + 1}}$$

A lower bound

- Suppose, the maximum flux of a vector field $\phi = (\phi^x, \phi^t)$ through the graph provides a lower bound to $E(u)$

$$E(u) \geq \sup_{\phi \in \mathcal{K}} \int_{\Gamma_u} \phi \cdot \nu_{\Gamma_u} \, \mathrm{d}\mathcal{H}^2.$$

A lower bound

- Suppose, the maximum flux of a vector field $\phi = (\phi^x, \phi^t)$ through the graph provides a lower bound to $E(u)$

$$E(u) \geq \sup_{\phi \in \mathcal{K}} \int_{\Gamma_u} \phi \cdot \nu_{\Gamma_u} \, \mathrm{d}\mathcal{H}^2.$$

- It turns out that equality holds for

$$\mathcal{K} = \left\{ \phi = (\phi^x, \phi^t) \,\middle|\, \phi^t(x,t) \geq f^*(x,t,\phi^x(x,t)) \right\}$$

A lower bound

- Suppose, the maximum flux of a vector field $\phi = (\phi^x, \phi^t)$ through the graph provides a lower bound to $E(u)$

$$E(u) \geq \sup_{\phi \in \mathcal{K}} \int_{\Gamma_u} \phi \cdot \nu_{\Gamma_u} \, \mathrm{d}\mathcal{H}^2.$$

- It turns out that equality holds for

$$\mathcal{K} = \left\{ \phi = (\phi^x, \phi^t) \;\middle|\; \phi^t(x,t) \geq f^*(x,t,\phi^x(x,t)) \right\}$$

- The integral can be extended to $\Omega \times \mathbb{R}$

$$E(u) = \sup_{\phi \in K} \int_{\Omega \times \mathbb{R}} \phi \cdot D\mathbf{1}_u,$$

- Convex relaxation and solution via primal-dual algorithm

Digital surface model of Graz (with Arnold Irschara)

Input

Digital surface model of Graz (with Arnold Irschara)

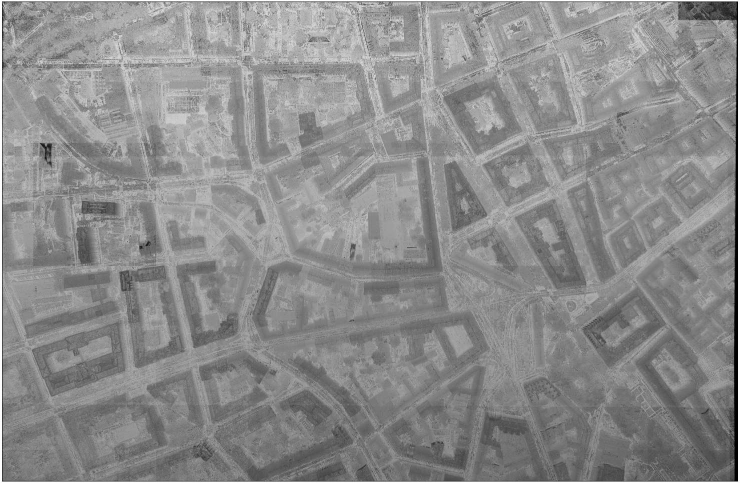

Data term only

Digital surface model of Graz (with Arnold Irschara)

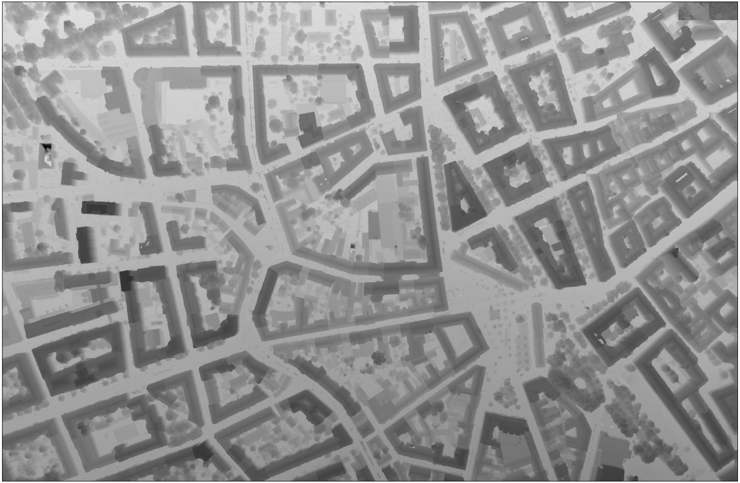

Convex variational approach

Digital surface model of Graz (with Arnold Irschara)

Online 3D reconstruction (with Gottfried Graber)

- Aim: Online computation of a 3D model as a camera browses the scene
- Use PTAM framework [Klein, Murray '07] to track the camera in real-time
- Volumetric range integration [Zach, Pock, Bischof '07]

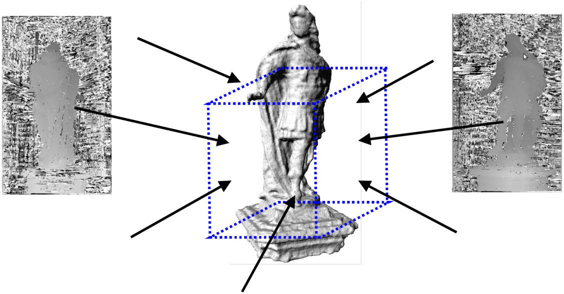

Online 3D reconstruction

- Surface is represented as the zero-level set of a signed distance function
- Convert range hypotheses to signed distance functions
- Each voxel (i, j, k) holds histogram $(h^b)_{b=1}^n$ of distance hypotheses d^b [Zach, '08]

Online 3D reconstruction

- Surface is represented as the zero-level set of a signed distance function
- Convert range hypotheses to signed distance functions
- Each voxel (i, j, k) holds histogram $(h^b)_{b=1}^n$ of distance hypotheses d^b [Zach, '08]
- Global volumetric TV-L^1 range integration

$$\min_u \|\nabla u\|_1 + \lambda \sum_{i,j,k} \sum_b h^b_{i,j,k} |u_{i,j,k} - d^b|$$

Online 3D reconstruction

- Surface is represented as the zero-level set of a signed distance function
- Convert range hypotheses to signed distance functions
- Each voxel (i, j, k) holds histogram $(h^b)_{b=1}^n$ of distance hypotheses d^b [Zach, '08]
- Global volumetric TV-L^1 range integration

$$\min_u \|\nabla u\|_1 + \lambda \sum_{i,j,k} \sum_b h^b_{i,j,k} |u_{i,j,k} - d^b|$$

- Can be solved exactly via the primal-dual algorithm

(a) 3D Scene

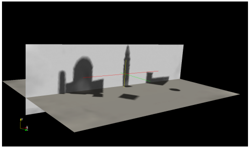

(b) Volume

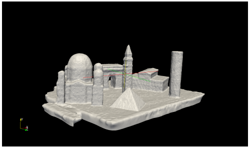

(c) Surface

→ See our poster and demo

Discussion

- I talked about convex methods to compute dense motion stereo and 3D

Discussion

- I talked about convex methods to compute dense motion stereo and 3D
- Simple linearization leads to real-time performance for motion and stereo, but can be unstable

Discussion

- I talked about convex methods to compute dense motion stereo and 3D
- Simple linearization leads to real-time performance for motion and stereo, but can be unstable
- Calibration method allows to compute high-quality results, but large computational costs

Discussion

- I talked about convex methods to compute dense motion stereo and 3D
- Simple linearization leads to real-time performance for motion and stereo, but can be unstable
- Calibration method allows to compute high-quality results, but large computational costs
- Volumetric range integration can deal with low-quality depth maps to produce high quality 3D models

Discussion

- I talked about convex methods to compute dense motion stereo and 3D
- Simple linearization leads to real-time performance for motion and stereo, but can be unstable
- Calibration method allows to compute high-quality results, but large computational costs
- Volumetric range integration can deal with low-quality depth maps to produce high quality 3D models

- Do we really need high quality depth maps?

Discussion

- I talked about convex methods to compute dense motion stereo and 3D
- Simple linearization leads to real-time performance for motion and stereo, but can be unstable
- Calibration method allows to compute high-quality results, but large computational costs
- Volumetric range integration can deal with low-quality depth maps to produce high quality 3D models

- Do we really need high quality depth maps?
- Probably depends on the application ...

Thank you for your attention!