Convex Relaxation Methods for Computer Vision

Daniel Cremers

Computer Science Department

TU Munich






with Kalin Kolev, Evgeny Strekalovskiy, Thomas Pock, Bastian Goldlücke, Antonin Chambolle & Jan Lellmann

3D Reconstruction from Multiple Views

Optimization in Computer Vision

Image segmentation:

Geman, Geman ’84, Blake, Zisserman ‘87, Kass et al. ’88,

Mumford, Shah ’89, Caselles et al. ‘95, Kichenassamy et al. ‘95,

Paragios, Deriche ’99, Chan, Vese ‘01, Tsai et al. ‘01, …

Multiview stereo reconstruction:

Non-convex energies

Faugeras, Keriven ’98, Duan et al. ‘04, Yezzi, Soatto ‘03,

Seitz et al. ‘06, Hernandez et al. ‘07, Labatut et al. ’07, …

Optical flow estimation:

Horn, Schunck ‘81, Nagel, Enkelmann ‘86, Black, Anandan ‘93,

Alvarez et al. ‘99, Brox et al. ‘04, Baker et al. ‘07, Zach et al. ‘07,

Sun et al. ‘08, Wedel et al. ’09, …

Non-convex versus Convex Energies

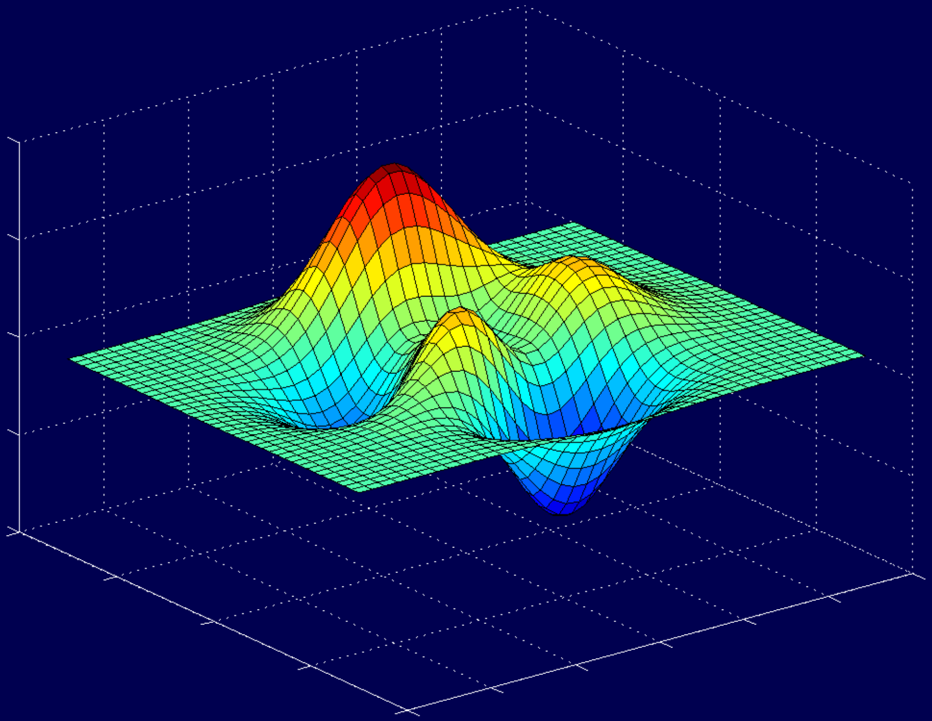

Non-convex energy

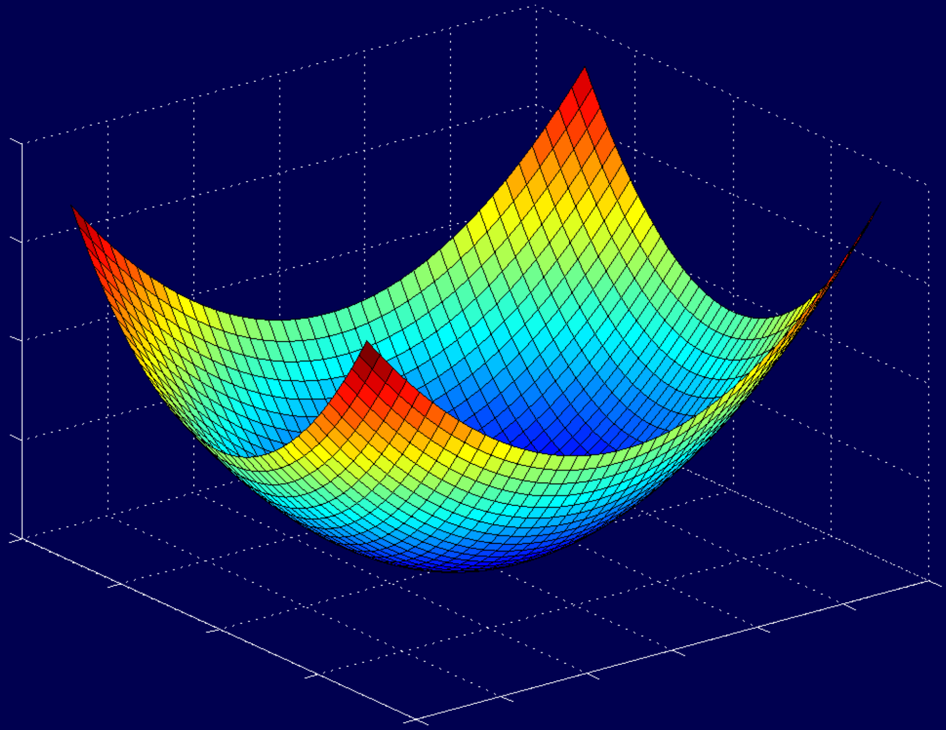

Convex energy

Some related work: *Brakke '95, Alberti et al. '01, Chambolle '01, Attouch et al. '06, Nikolova et al. '06, Cremers et al. '06, Bresson et al. '07, Lellmann et al. '08, Zach et al. '08, Chambolle et al. '08, Pock et al. '09, Zach et al. '09, Brown et al. '10, Bae et al. '10, Yuan et al. '10,…*

Multiview reconstruction

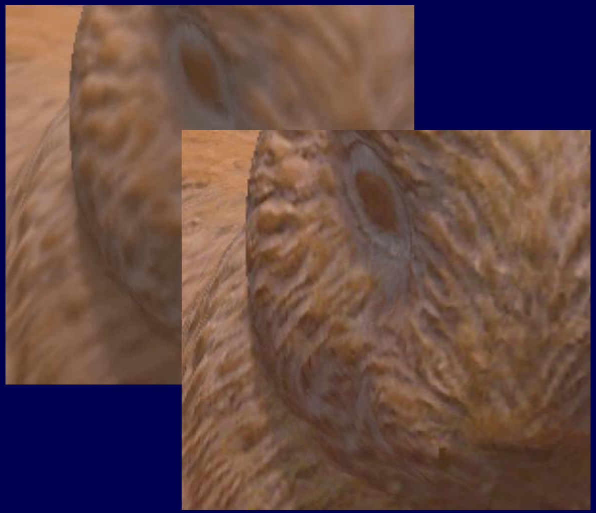

Super-res.textures

4D reconstruction

Stereo reconstruction

Segmentation

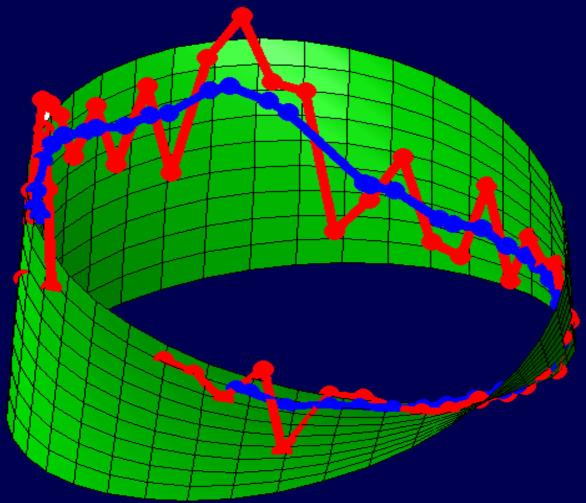

Manifold-valued functions

Multiview reconstruction

Super-res.textures

4D reconstruction

Stereo reconstruction

Segmentation

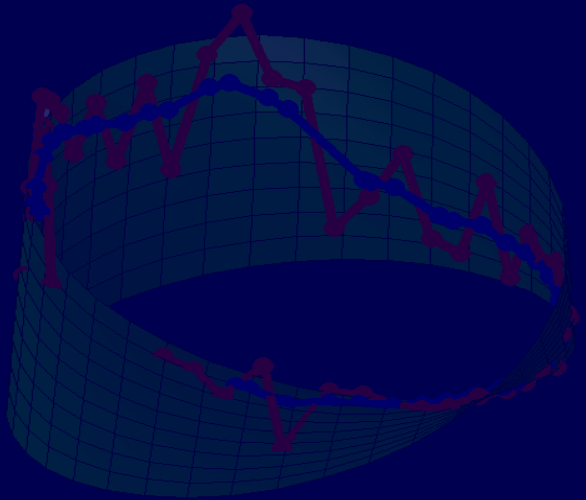

Manifold-valued functions

$$\rho : \left(V \subset \mathbb{R}^3\right) \to [0, 1]$$

$$E(S) = \int_S \rho(s)\, ds$$

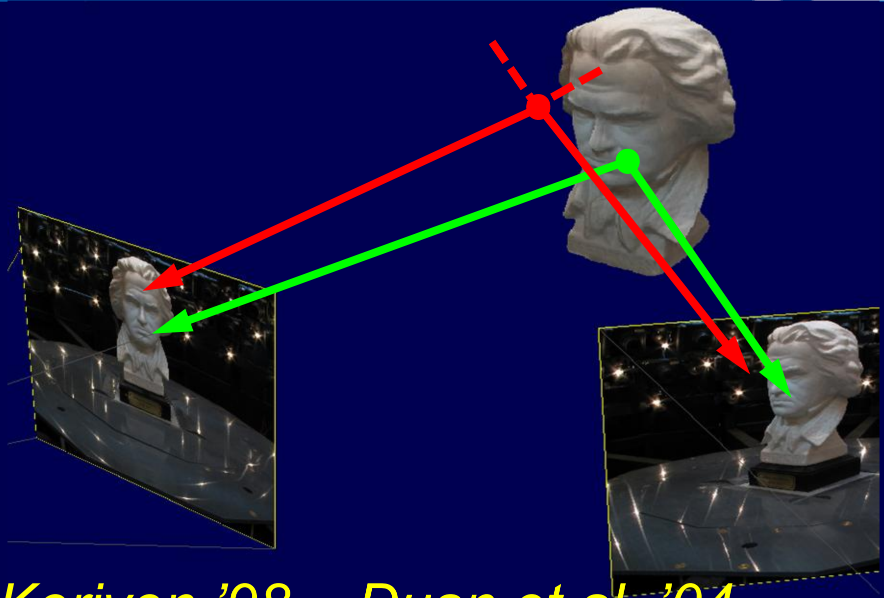

3D Reconstruction: *Faugeras, Keriven '98, Duan et al. '04*

Segmentation: *Kichenassamy et al. '95, Caselles et al. '95*

Optimal solution is the empty set: $\arg\min_S E(S) = \emptyset$

Resort:

- Local optimization: *Faugeras, Keriven TIP '98*
- Generative object/background modeling: *Yezzi, Soatto '03,…*
- Constrain search space: *Vogiatsis, Torr, Cipolla CVPR '05*
- Intelligent ballooning: *Boykov, Lempitsky BMVC '06*

Silhouette Consistent Reconstructions

$$\min_S \int_S \rho \, ds$$

$$\text{s. t.} \quad \pi_i(S) = S_i \;\; \forall\, i = 1, \dots, n$$

$$\pi_i : V \to \Omega_i$$

$$S_i \subset \Omega_i$$

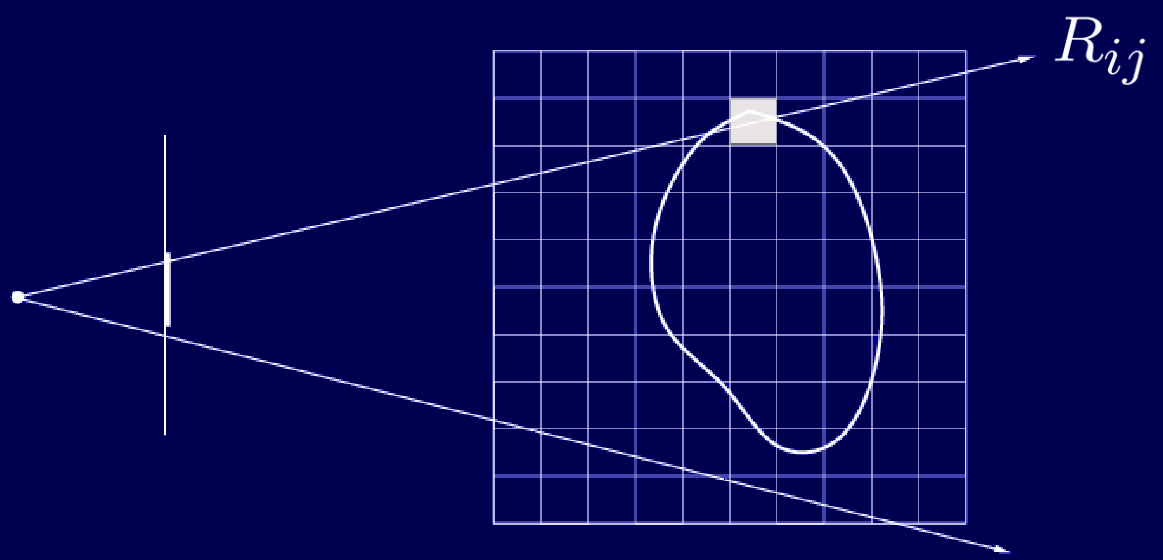

Kolev et al., IJCV 2009, Cremers, Kolev, PAMI 2011

Silhouette Consistent Reconstructions

$$\min_S \int_S \rho \, ds$$

$$\text{s. t.} \quad \pi_i(S) = S_i \;\; \forall\, i = 1, \ldots, n$$

$$u = 1_{\text{int}(S)}$$

$$\min_u \int_V \rho(x) |\nabla u(x)| \, dx$$

$$\Sigma = \begin{cases} \text{s. t.} \quad \sout{u : V \to \{0,1\}} \;\; u : V \to [0,1] \\ \int_{R_{ij}} u(x)\, dx \geq \delta \;\; \text{if } j \in S_i \\ \int_{R_{ij}} u(x)\, dx = 0 \;\; \text{if } j \notin S_i \end{cases}$$

Proposition: The set Σ of silhouette-consistent solutions is convex.

Kolev et al., IJCV 2009, Cremers, Kolev, PAMI 2011

Reconstruction of Fine-scale Structures

Image data courtesy of Yasutaka Furukawa.

Multiview reconstruction

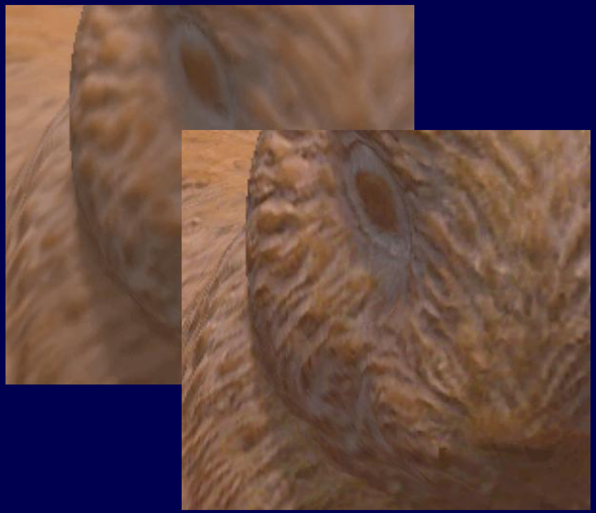

Super-res.textures

4D reconstruction

Stereo reconstruction

Segmentation

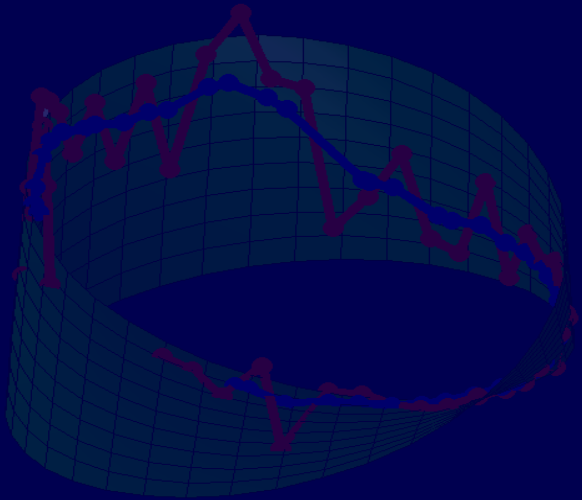

Manifold-valued functions

Surface Evolution to Optimum

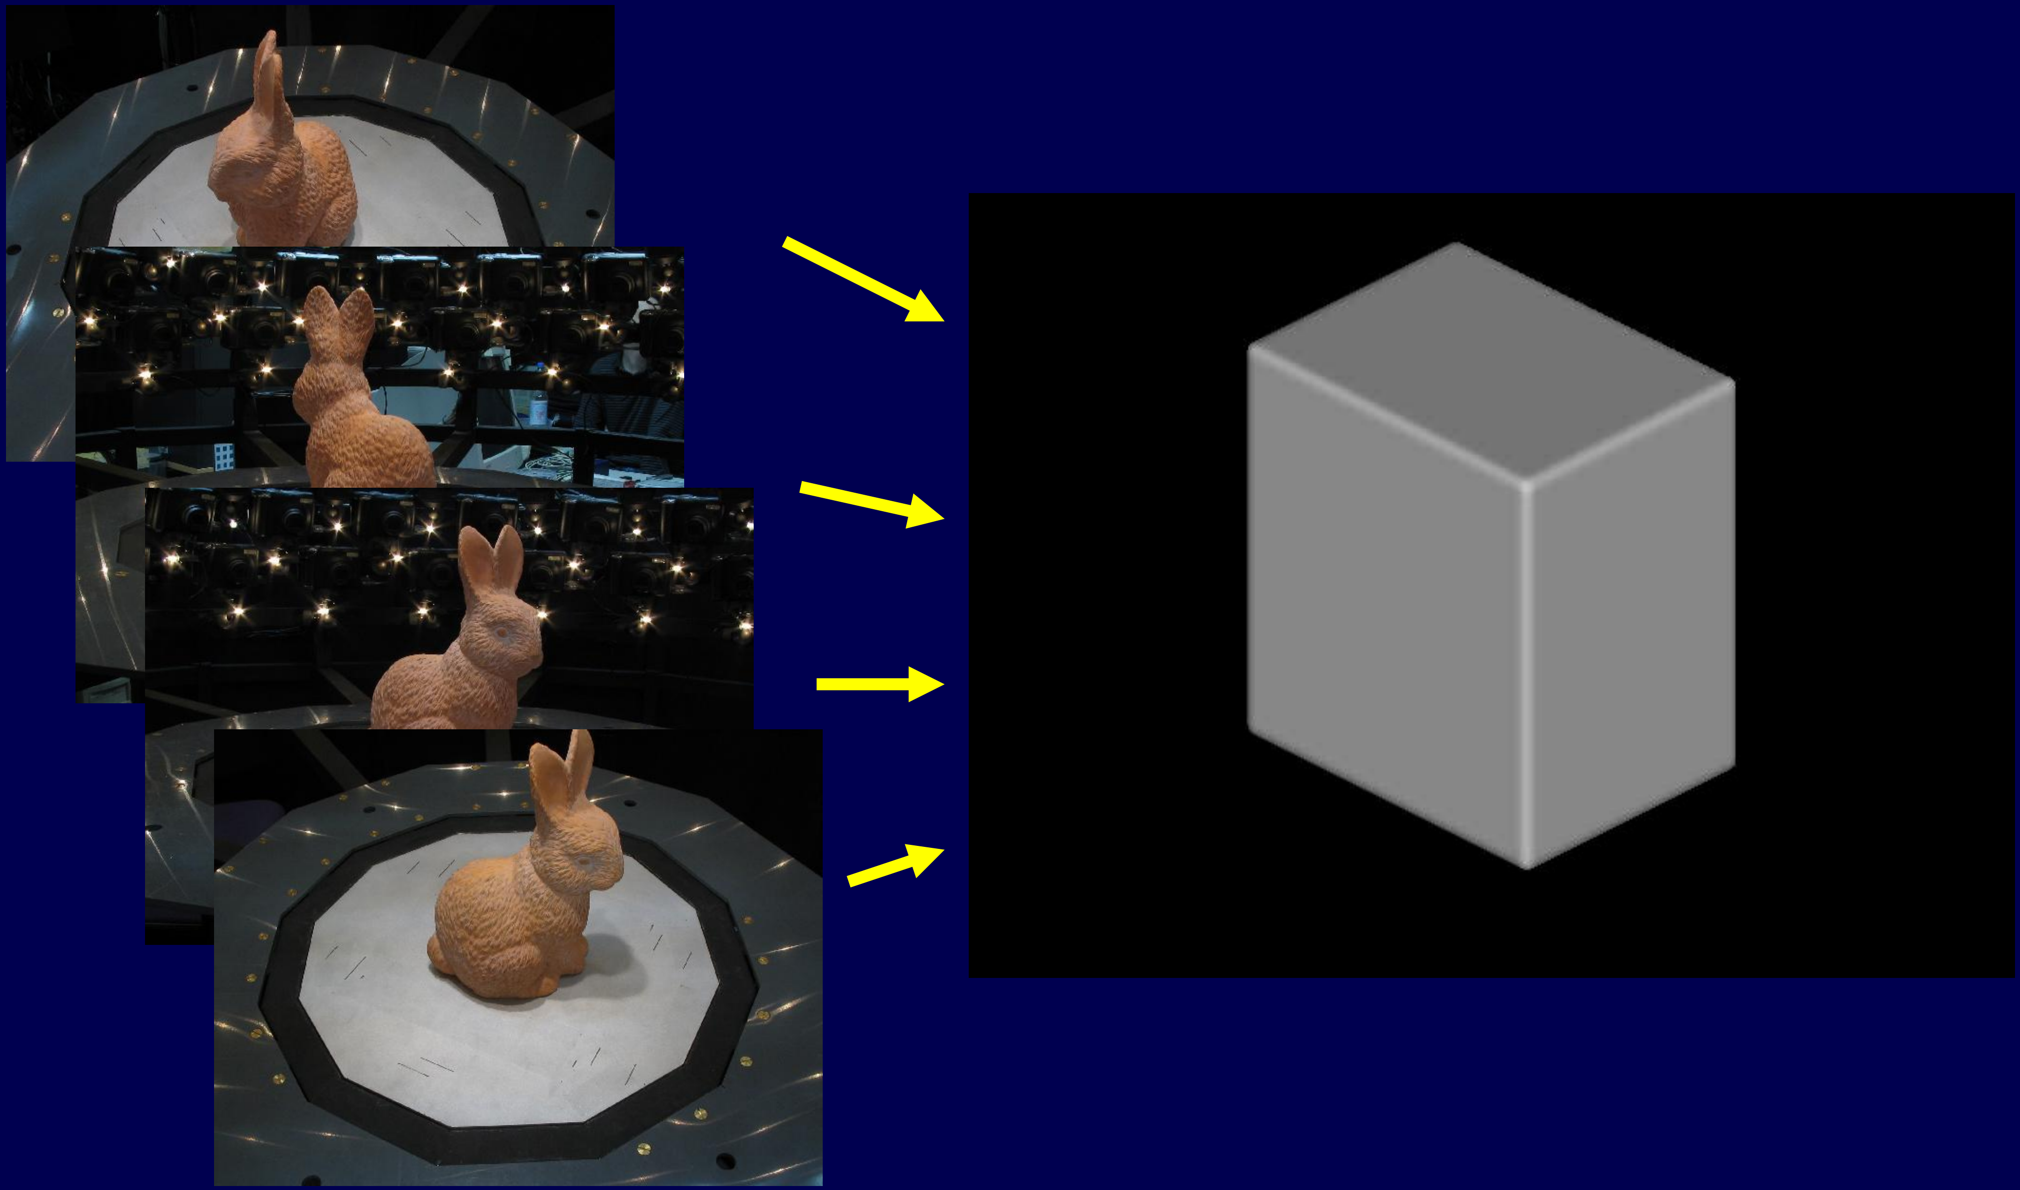

Super-Resolution Texture Map

Given all images $\mathcal{I}_i : \Omega_i \to \mathbb{R}^3$, determine the surface color $T : S \to \mathbb{R}^3$

$$\min_T \sum_{i=1}^{n} \int_{\Omega_i} \left(b * (T \circ \beta_i) - \mathcal{I}_i \right)^2 dx + \lambda \int_S \|\nabla_S T\| ds$$

blur & downsample   back-projection

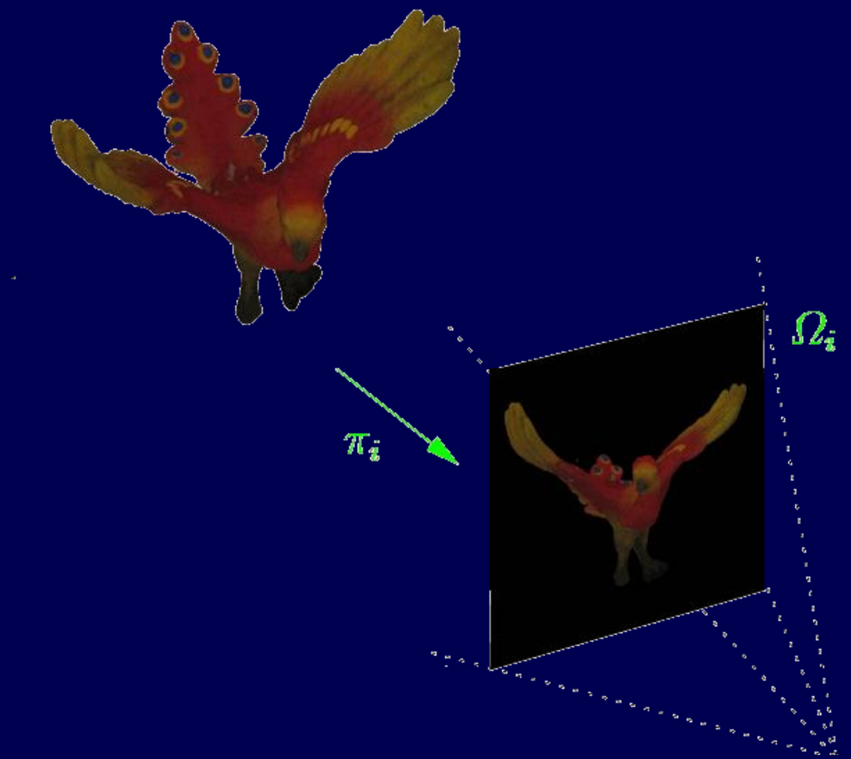

Goldlücke, Cremers, ICCV '09, DAGM '09, IJCV '13*

** Best Paper Award*

Super-Resolution Texture Map

Goldlücke, Cremers, ICCV ’09, DAGM ’09, IJCV ‘13*

** Best Paper Award*

Closeup of input image

Super-resolution texture

* Best Paper Award

Goldlücke, Cremers, ICCV '09, DAGM '09, IJCV '13*

Reconstructing the Niobids Statues

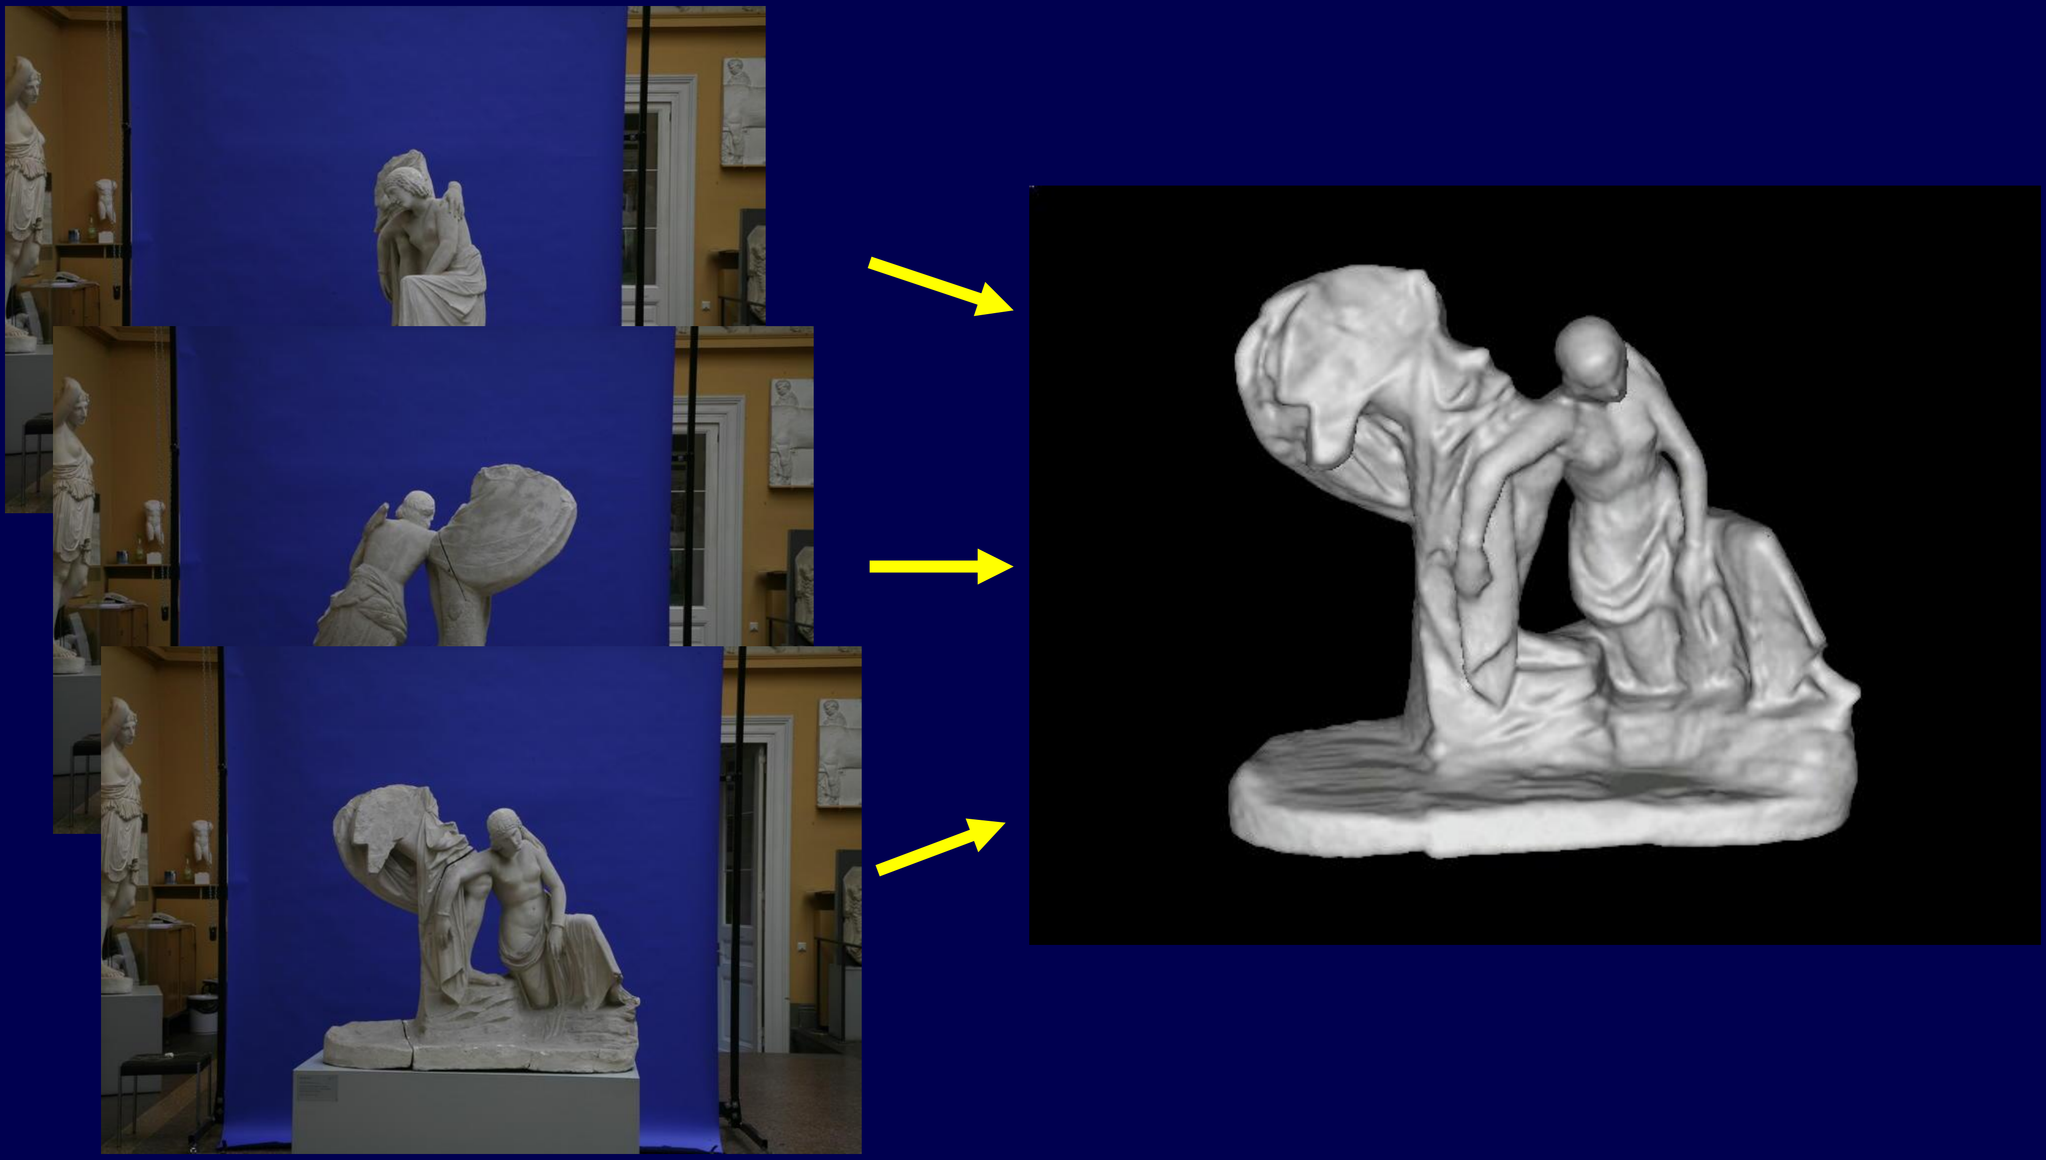

Kolev, Cremers, ECCV ’08, PAMI 2011

Overview

Multiview reconstruction

Super-res.textures

4D reconstruction

Stereo reconstruction

Segmentation

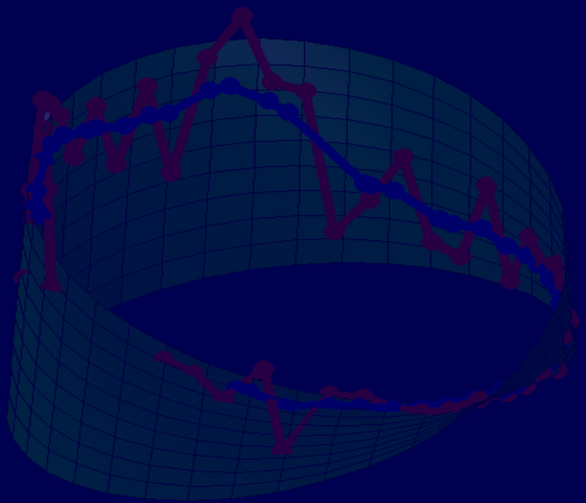

Manifold-valued functions

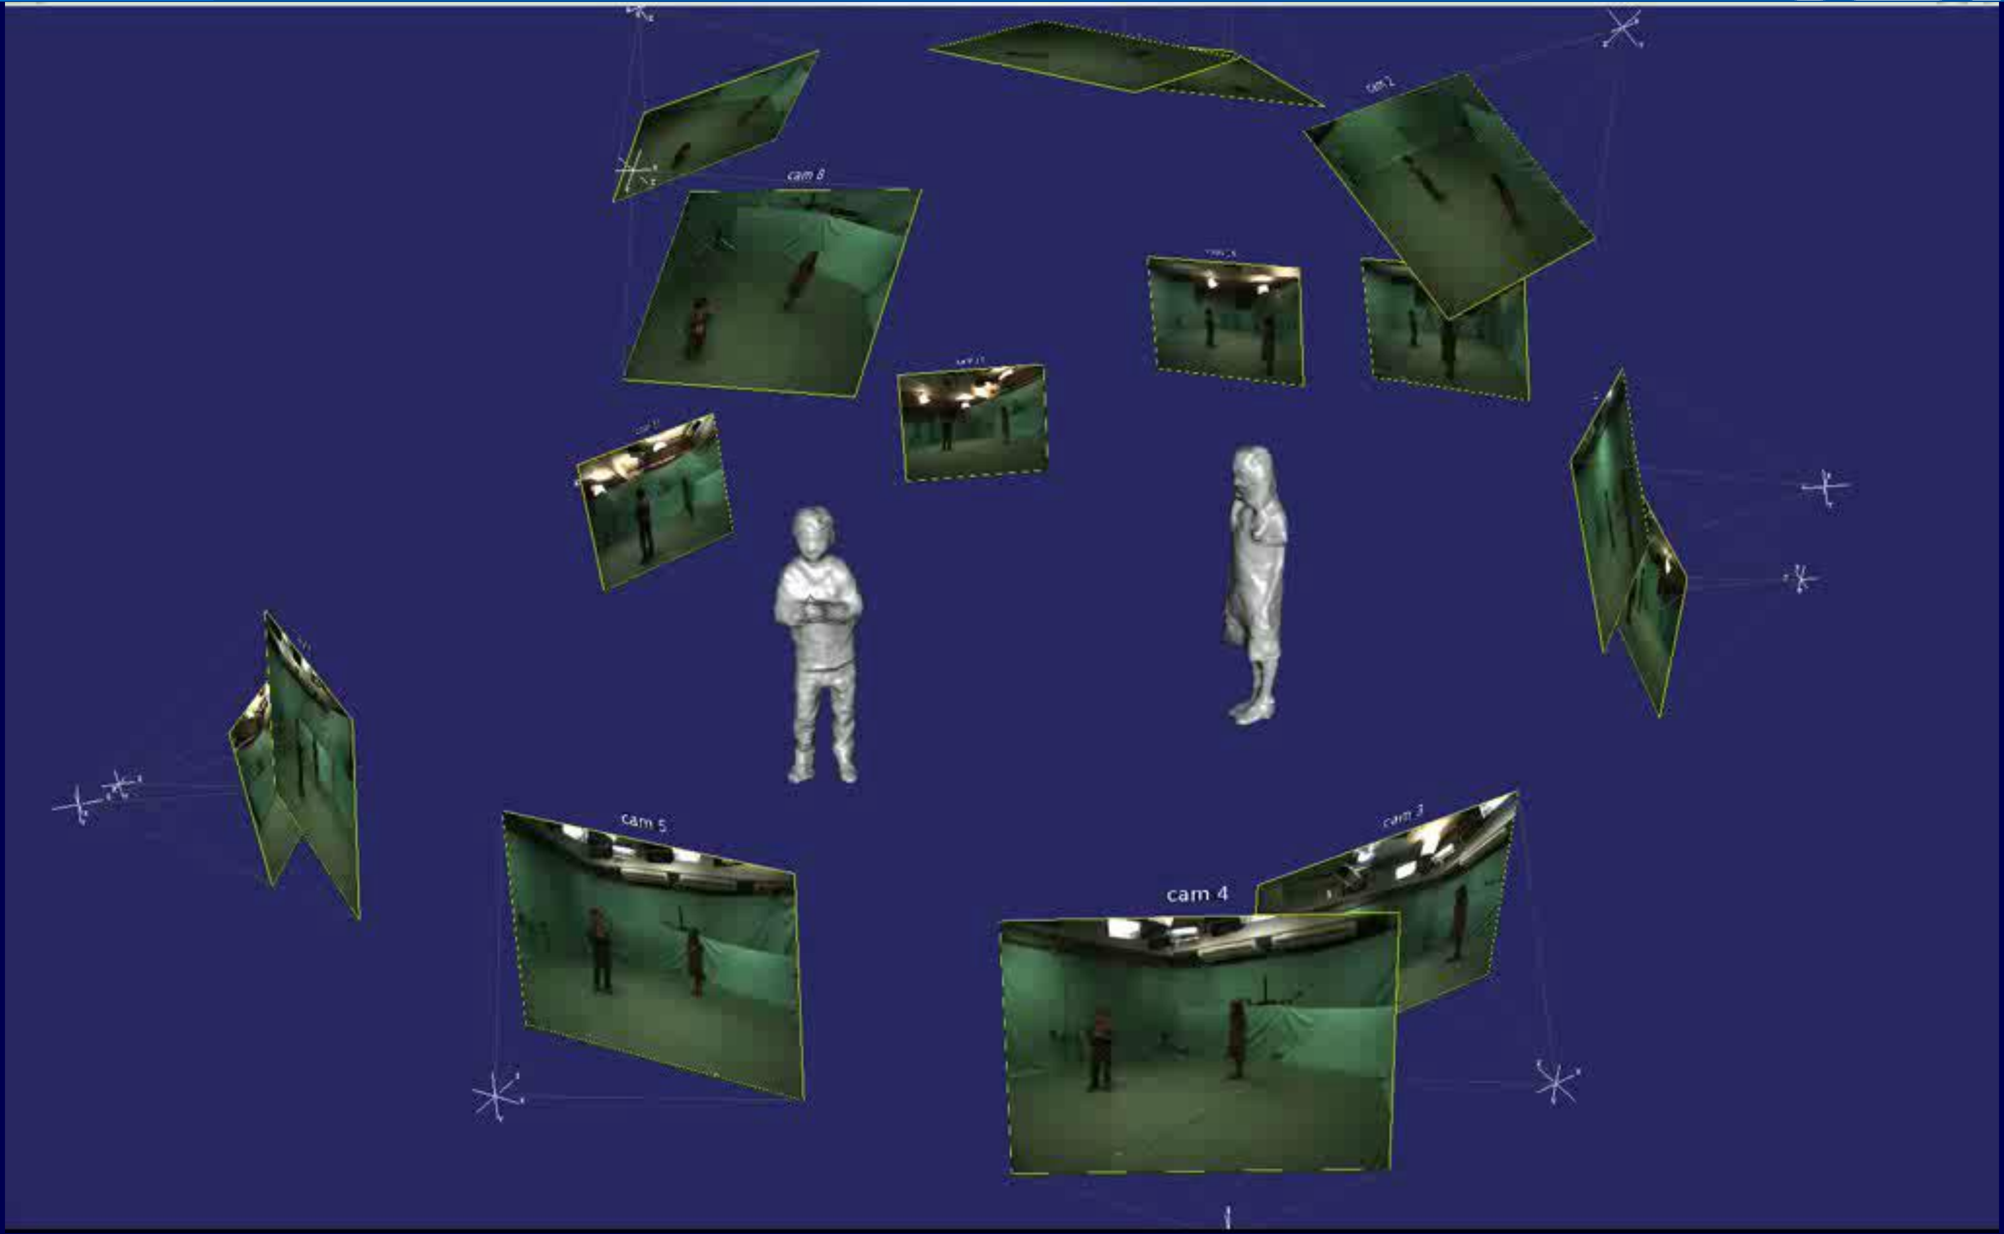

Oswald, Cremers, ICCV '13 4DMoD Workshop

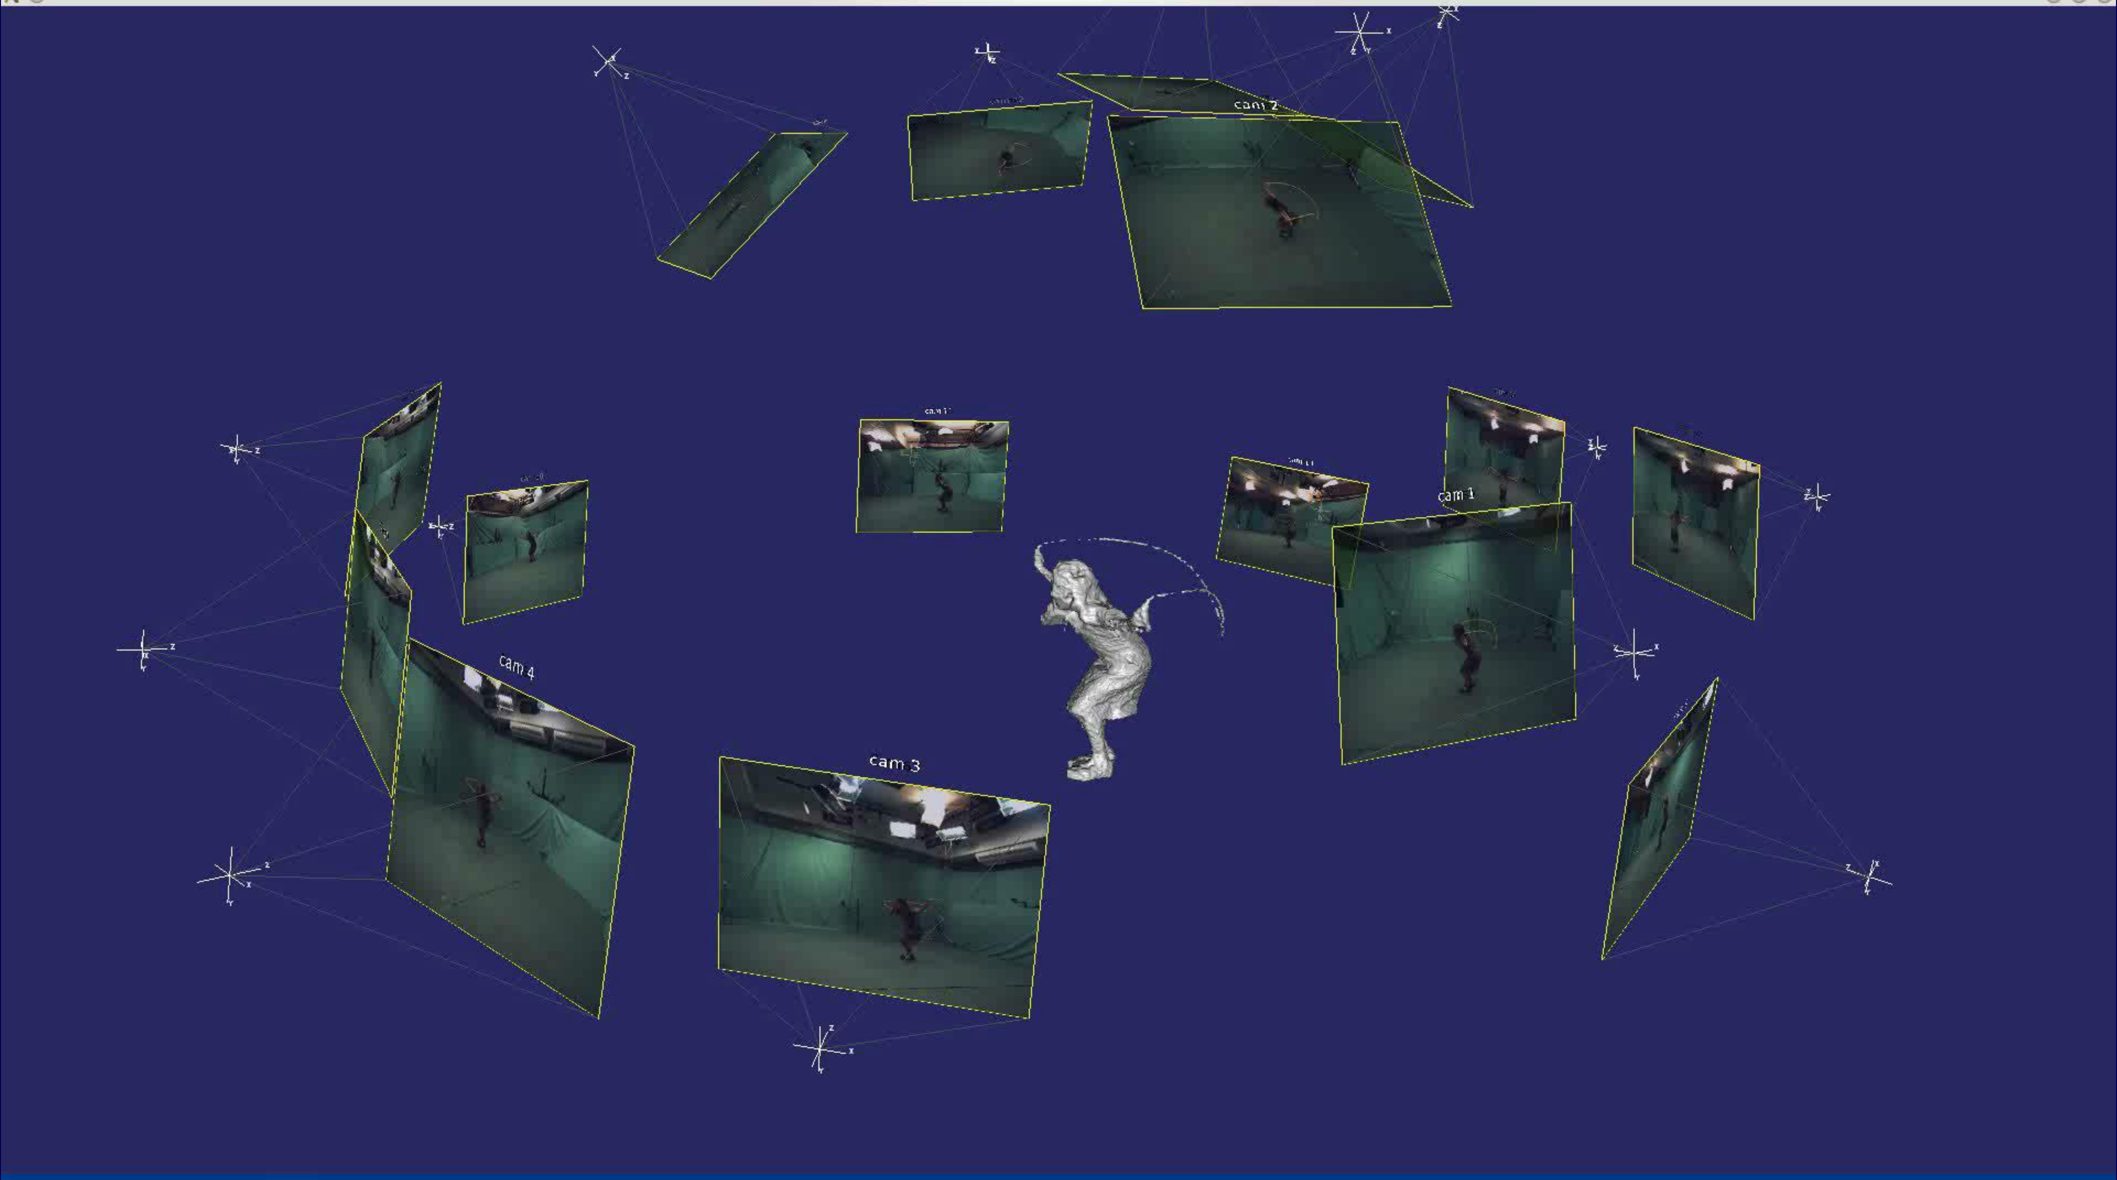

Oswald, Cremers, ICCV '13 4DMoD Workshop

Action Reconstruction

Action Reconstruction

Overview


Multiview reconstruction


Super-res.textures


4D reconstruction


Stereo reconstruction


Segmentation

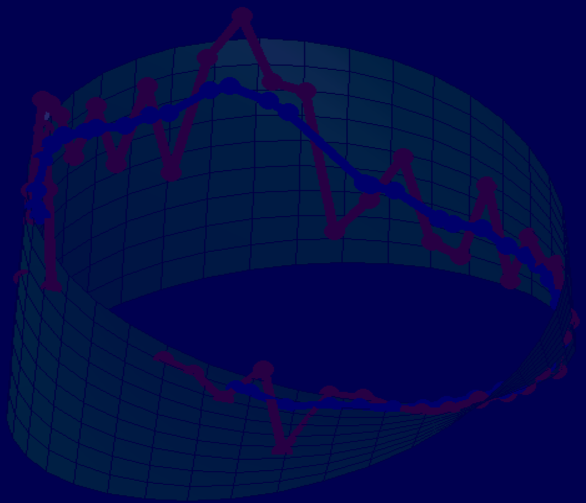
Manifold-valued functions

$$u : \Omega \to \Gamma = [\gamma_{min}, \gamma_{max}]$$

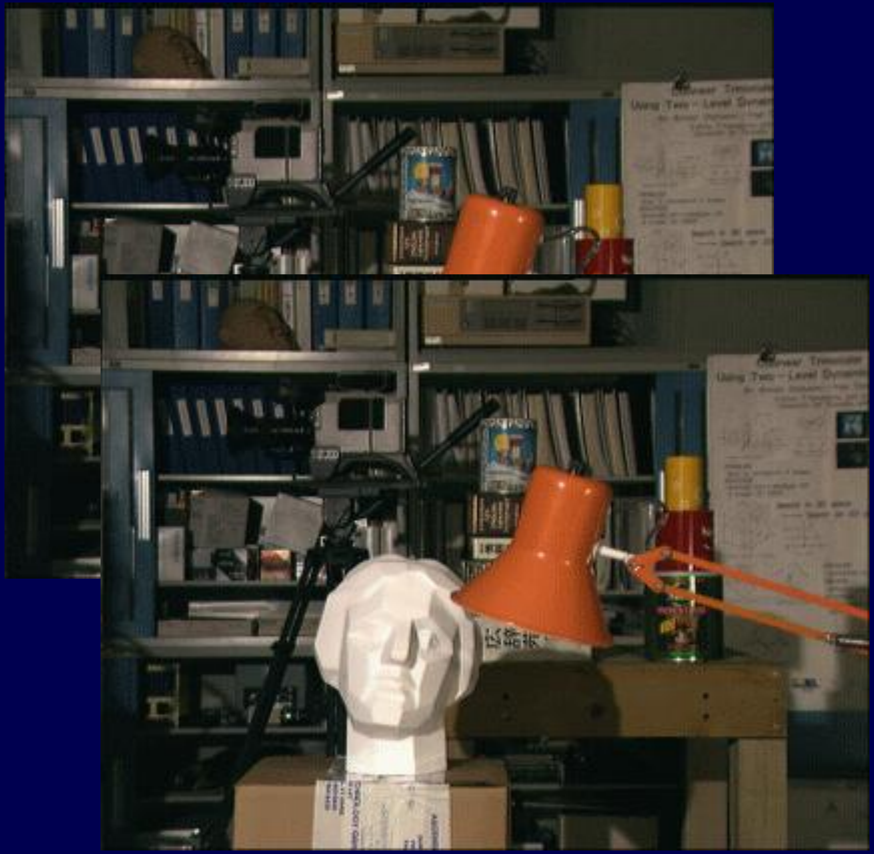

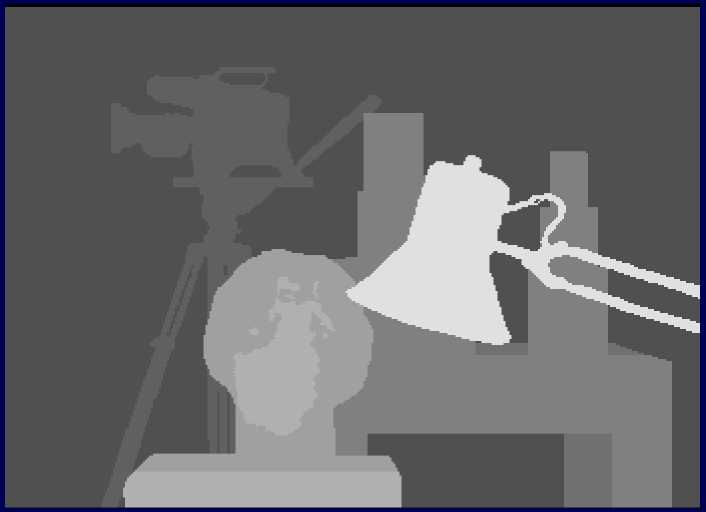

Example: Stereo

$$u : \Omega \to \Gamma = [\gamma_{min}, \gamma_{max}]$$

$$E(u) = \underbrace{\int_\Omega \rho(x, u(x))\, dx}_{\text{nonconvex data term}} + \underbrace{\int_\Omega |\nabla u(x)|\, dx}_{\text{label regularity}} \qquad (*)$$

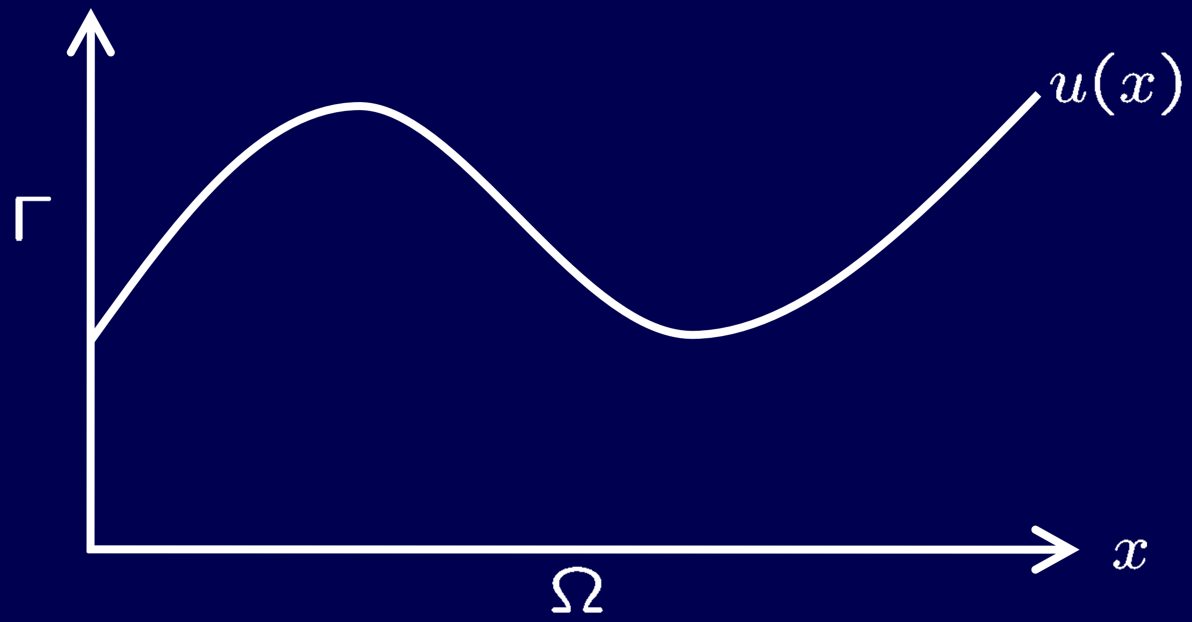

Pock , Schoenemann, Graber, Bischof, Cremers ECCV '08

$$u : \Omega \to \Gamma = [\gamma_{min}, \gamma_{max}]$$

$$E(u) = \int_\Omega \rho(x, u(x))\, dx + \int_\Omega |\nabla u(x)|\, dx \qquad (*)$$

Let $v : (\Sigma = \Omega \times \Gamma) \to \{0, 1\}$ $\quad v(x,\gamma) = 1_{u \geq \gamma}(x)$

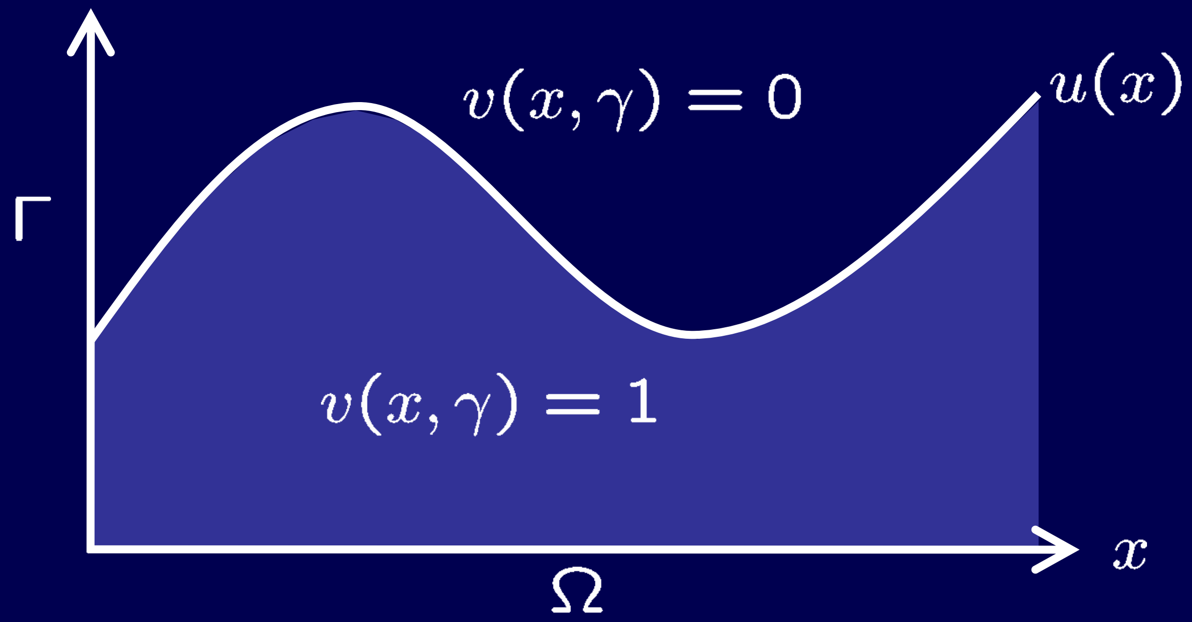

Pock , Schoenemann, Graber, Bischof, Cremers ECCV '08

Cartesian Currents and Relaxation

$$u : \Omega \to \Gamma = [\gamma_{min}, \gamma_{max}]$$

$$E(u) = \int_\Omega \rho(x, u(x))\, dx + \int_\Omega |\nabla u(x)|\, dx \qquad (*)$$

nonconvex functional

Let $v : (\Sigma = \Omega \times \Gamma) \to \{0, 1\}$ $\qquad v(x,\gamma) = 1_{u \geq \gamma}(x)$

Theorem: Minimizing $(*)$ is equivalent to minimizing

$$E(v) = \int_\Sigma \rho(x,\gamma)|\partial_\gamma v(x,\gamma)| + |\nabla v(x,\gamma)|\, dx\, d\gamma \qquad (**)$$

convex functional

Solve $(**)$ in relaxed space ($v : \Sigma \to [0,1]$) and threshold to obtain a globally optimal solution.

Pock , Schoenemann, Graber, Bischof, Cremers ECCV '08

Global Optima for Convex Regularizers

Let

$$E(u) = \int_\Omega f(x, u, \nabla u)\, dx$$

be continuous in $x \in \mathbb{R}^d$ and u, and convex in ∇u.

<u>Theorem:</u>

For any function $u \in W^{1,1}(\Omega; \mathbb{R})$ we have:

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

where ϕ is constrained to the convex set

$$\mathcal{K} = \Big\{ \phi = (\phi^x, \phi^t) \in C_0 \left(\Omega \times \mathbb{R}; \mathbb{R}^d \times \mathbb{R} \right) \; : \\ \phi^t(x,t) \geq f^*(x, t, \phi^x(x,t)) \; , \; \forall x, t \in \Omega \times \mathbb{R} \Big\} \, .$$

Pock, Cremers, Bischof, Chambolle, SIAM J. on Imaging Sciences ’10

Global Optima for Convex Regularizers

The functional $E(u)$ can be minimized by solving the relaxed saddle point problem

$$\min_v F(v) = \min_v \sup_{\phi \in \mathcal{K}} \int_{\Omega \times \mathbb{R}} \phi \cdot Dv,$$

Theorem:

The functional F fulfills a generalized coarea formula:

$$F(v) = \int_{-\infty}^{\infty} F(\mathbf{1}_{v \geq s})\, ds.$$

As a consequence, we have a thresholding theorem assuring that we can globally minimize the functional $E(u)$.

Pock, Cremers, Bischof, Chambolle, SIAM J. on Imaging Sciences '10

Given the saddle point problem

$$\min_{x \in C} \max_{y \in K} \langle Ax, y \rangle + \langle g, x \rangle - \langle h, y \rangle$$

with close convex sets C and K and linear operator A of norm L.

The iterative algorithm

$$\begin{cases} y^{n+1} = \Pi_K\left(y^n + \sigma(A\bar{x}^n - h)\right) \\ x^{n+1} = \Pi_C\left(x^n - \tau(A^* y^{n+1} + g)\right) \\ \bar{x}^{n+1} = 2\,x^{n+1} - x^n \end{cases}$$

converges with rate $O(1/N)$ to a saddle point for $\sigma\,\tau\,L^2 \leq 1$.

Pock, Cremers, Bischof, Chambolle, ICCV '09, Chambolle, Pock '10

Evolution to Global Minimum

One of two input images

Courtesy of Microsoft

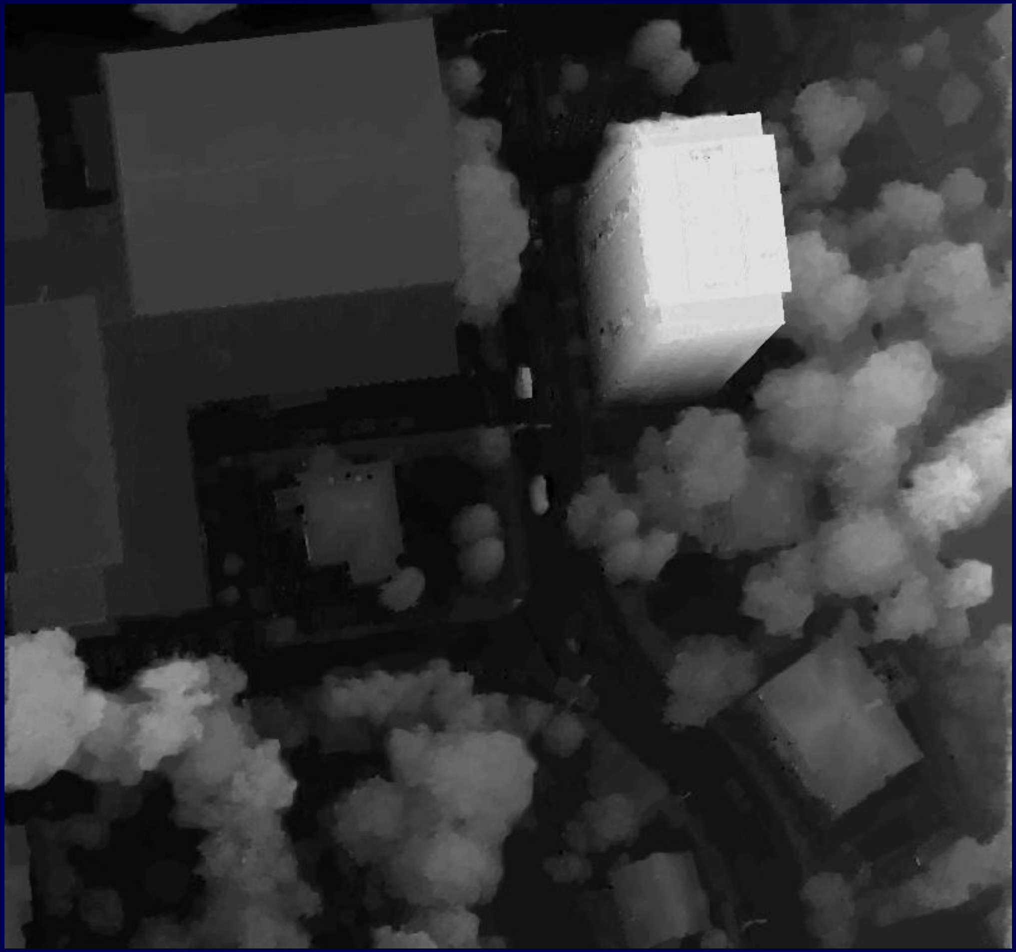

Depth reconstruction

Reconstruction from Aerial Images

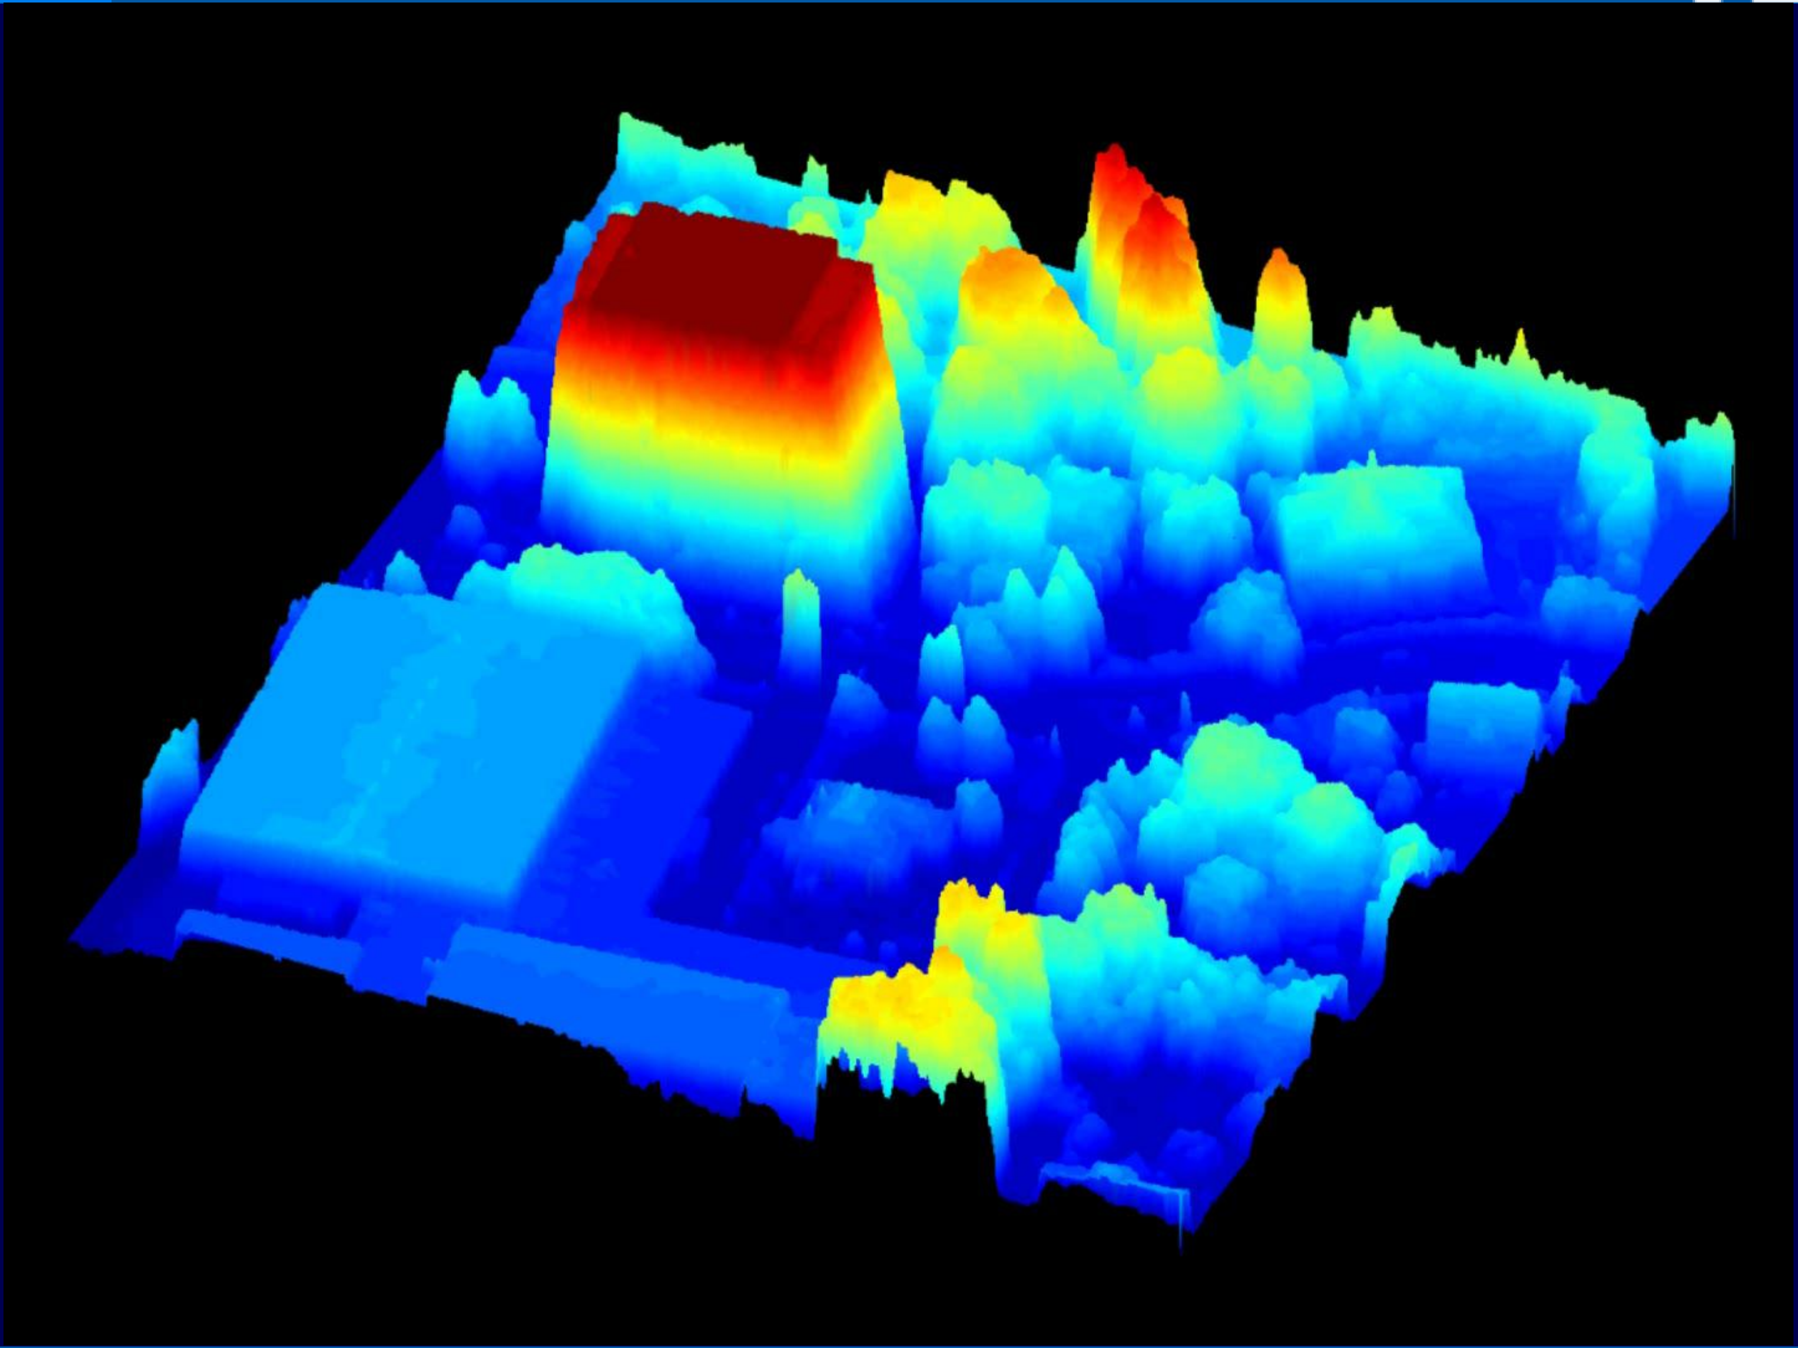

Overview


Multiview reconstruction


Super-res.textures


4D reconstruction


Stereo reconstruction


Segmentation

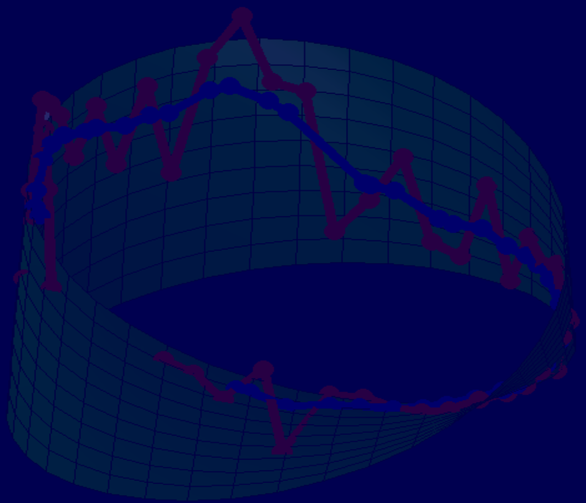
Manifold-valued functions

$$\min_{\Omega_0,\dots,\Omega_n} \frac{1}{2}\sum_i |\partial\Omega_i| + \sum_i \int_{\Omega_i} f_i(x)\,dx$$

$$\text{s.t.} \quad \bigcup_i \Omega_i = \Omega \subset \mathbb{R}^d, \quad \text{and } \Omega_i \cap \Omega_j = \emptyset \;\; \forall i \neq j$$

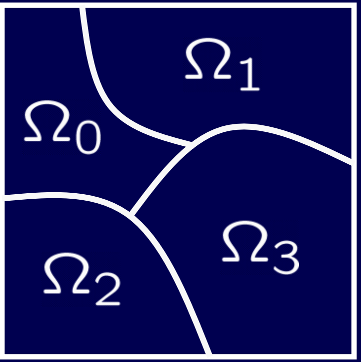

Potts '52, Blake, Zisserman '87, Mumford-Shah '89, Vese, Chan '02

Proposition: With $v_i = 1_{\Omega_i}$, this is equivalent to

$$\min_{v \in \mathcal{B}} \frac{1}{2}\sum_i \int_\Omega |Dv_i| + \int_\Omega v_i\, f_i\, dx = \min_{v \in \mathcal{B}} \sup_{p \in \mathcal{K}} \sum_i \int_\Omega v_i \operatorname{div} p_i\, dx + \int_\Omega v_i\, f_i\, dx$$

where $\mathcal{K} = \left\{ p = (p_1, \dots, p_n)^\top \in \mathbb{R}^{n \times d} : \left| p_i - p_j \right| \leq 1, \; \forall i < j \right\}$

Chambolle, Cremers, Pock '08, SIIMS '12, Pock et al. CVPR '09

Test Case: The Triple Junction

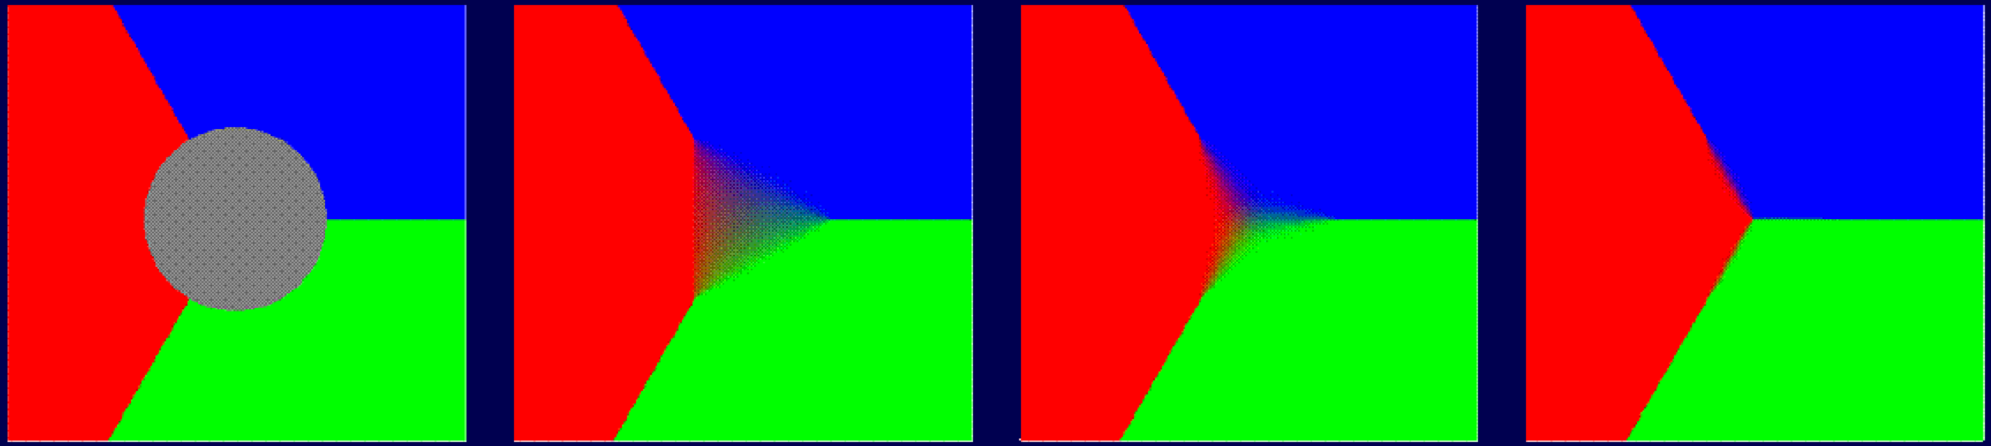

Input image *Lellmann et al. '08* *Zach et al. '08* our approach

<u>Proposition:</u> The proposed relaxation strictly dominates alternative relaxations.

Chambolle, Cremers, Pock '08, SIIMS '12, Pock et al. CVPR '09

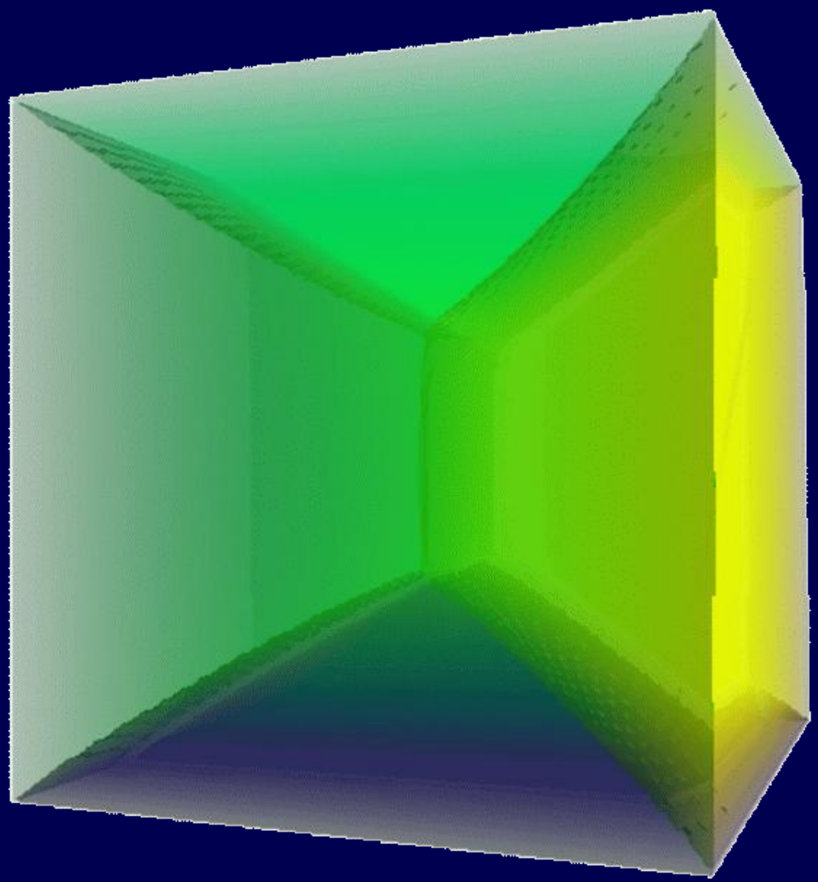

3D min partition inpainting

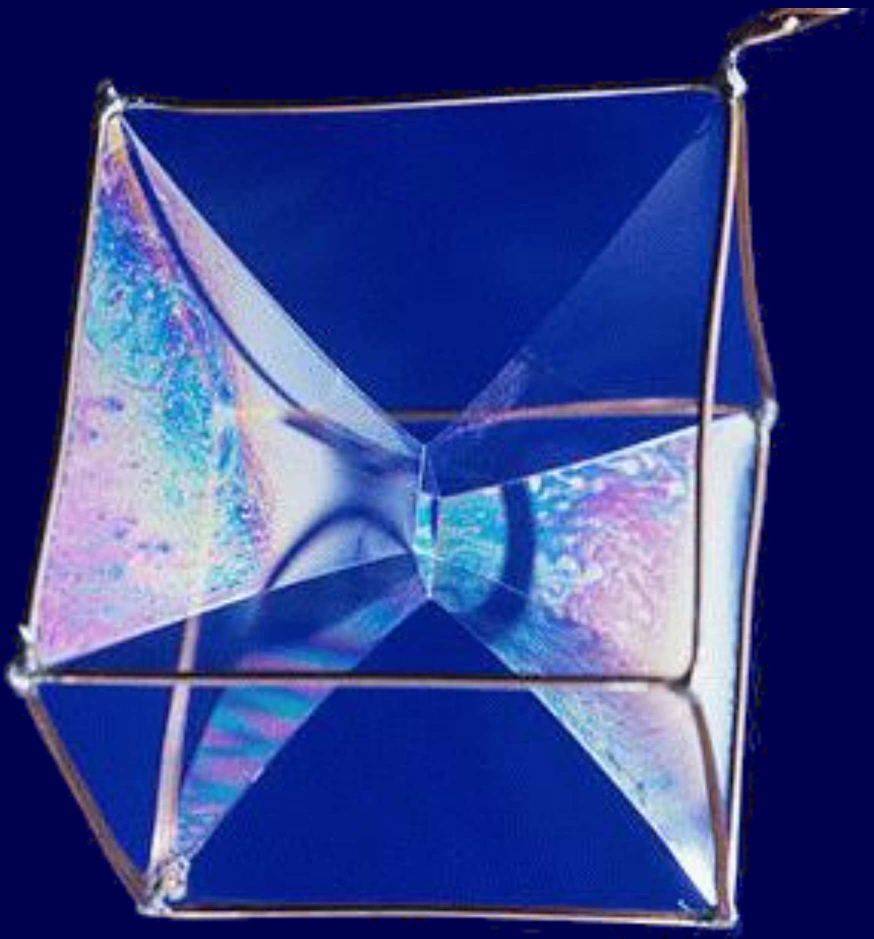

Photograph of a soap film

Chambolle, Cremers, Pock '08, SIIMS '12, Pock et al. CVPR '09

The Minimal Partition Problem

Input color image

10 label segmentation

Chambolle, Cremers, Pock ’08, SIIMS ‘12, Pock et al. CVPR ’09

Nieuwenhuis, Strekalovskiy, Cremers ICCV ’13

Idea: Impose a prior on the relative size of object parts

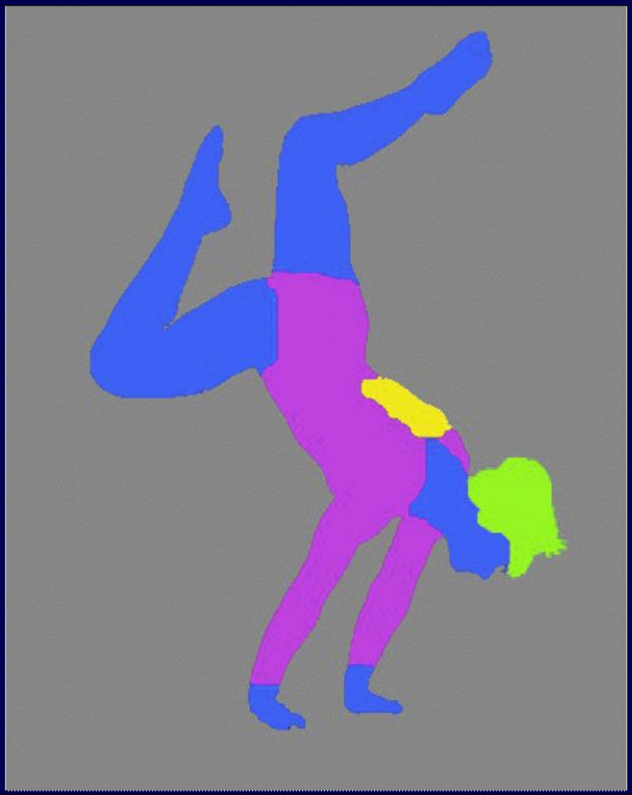

$$\int v_4 \, dx \leq 0.1 \cdot \int (1 - v_5) \, dx$$

Nieuwenhuis, Strekalovskiy, Cremers ICCV '13

with length regularity

with proportion prior

Nieuwenhuis, Strekalovskiy, Cremers ICCV '13

$$E(u) = \lambda \int_\Omega (f-u)^2 dx + \int_{\Omega \setminus S_u} |\nabla u|^2 dx + \nu \mathcal{H}^1(S_u) \qquad (*)$$

Mumford, Shah '89

For $u \in SBV(\Omega),\ \Omega \subset \mathbb{R}^n,\ (*)$ can be written as

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

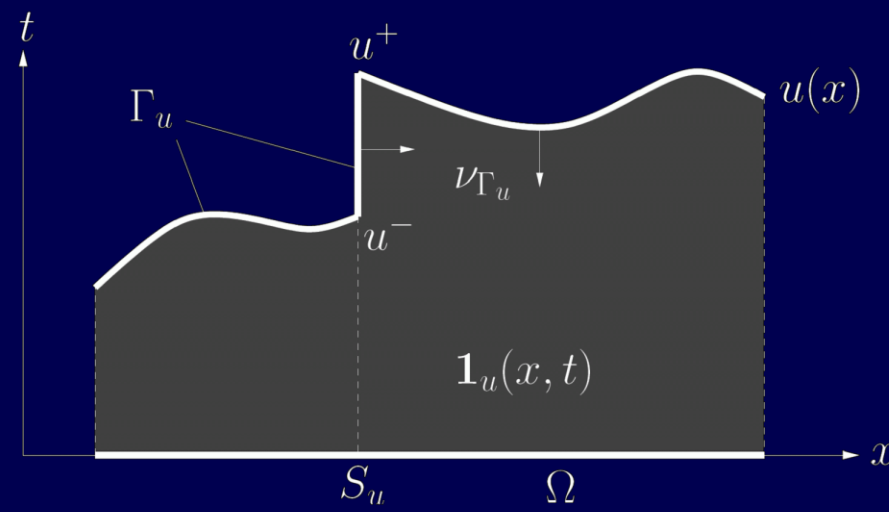

with a convex set

$$K = \Big\{ \varphi \in C_0(\Omega \times \mathbb{R}; \mathbb{R}^n \times \mathbb{R}) :$$

$$\varphi^t(x,t) \geq \frac{\varphi^x(x,t)^2}{4} - \lambda(t - f(x))^2, \quad \left| \int_{t_1}^{t_2} \varphi^x(x,s) ds \right| \leq \nu \Big\},$$

Alberti, Bouchitte, Dal Maso '04

Piecewise Smooth Approximation

Input image

piecewise constant

piecewise smooth

Pock, Cremers, Bischof, Chambolle ICCV '09

The Crack Tip & Open Boundaries

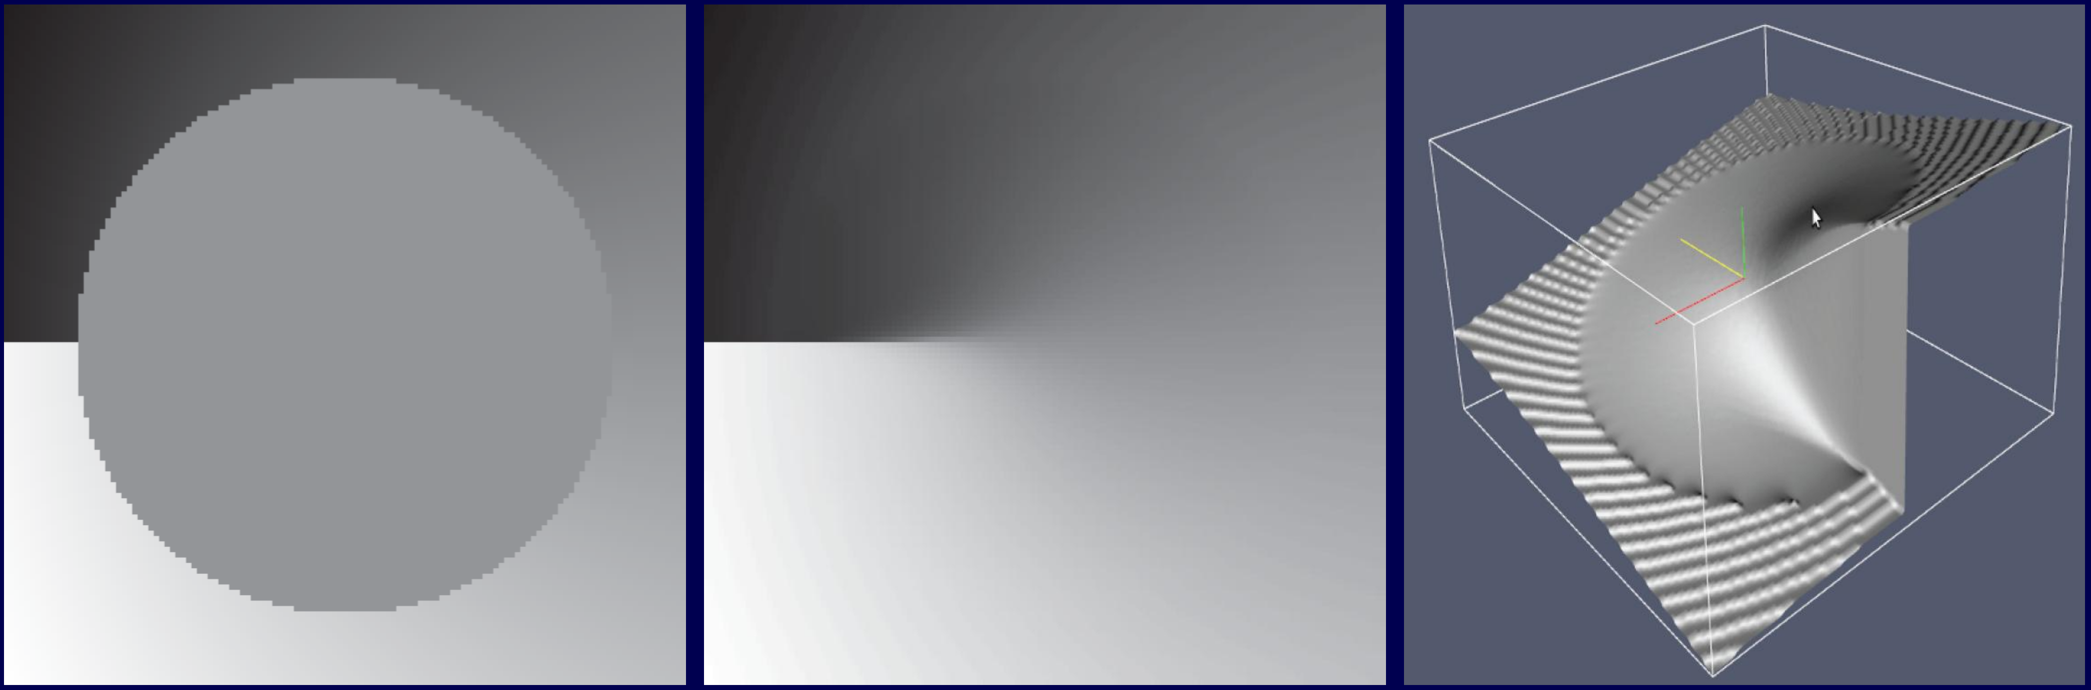

fixed boundary values | inpainted crack tip | surface plot

Pock, Cremers, Bischof, Chambolle ICCV '09

The Vectorial Mumford-Shah Problem

For $u \in L^1(\Omega, \mathbb{R}^k)$, we consider the functional

$$E(u) = \int_\Omega |f - u|^2 dx + \lambda \int_{\Omega \setminus S_u} \sum_{i=1}^k |\nabla u_i|^2 dx + \nu \mathcal{H}^1(S_u).$$

Proposition: For $v = 1_u = (1_{u_1}, \ldots, 1_{u_k})$, we have:

$$E(u) = \mathcal{F}(v) := \sup_{\sigma \in K} \sum_{i=1}^k \int_{\Omega \times \mathbb{R}} \sigma_i(x,t) \cdot Dv_i(x,t)$$

with the convex set:

$$K = \Big\{ \sigma \,\Big|\, (\sigma_i^x, \sigma_i^t) \in C_c^\infty(\Omega \times \mathbb{R}; \mathbb{R}^n \times \mathbb{R}),$$
$$\sigma_i^t(x, t_i) \geq \frac{1}{4\lambda} |\sigma_i^x(x, t_i)|^2 - (t_i - f_i(x))^2,$$
$$\sum_{j=1}^k \left| \int_{t_j}^{t_j'} \sigma_j^x(x, s)\, ds \right| \leq \nu, \quad \forall\, 1 \leq i \leq k,\, x \in \Omega,\, t_j < t_j' \Big\}.$$

Strekalovskiy, Chambolle, Cremers, CVPR '12

Input image

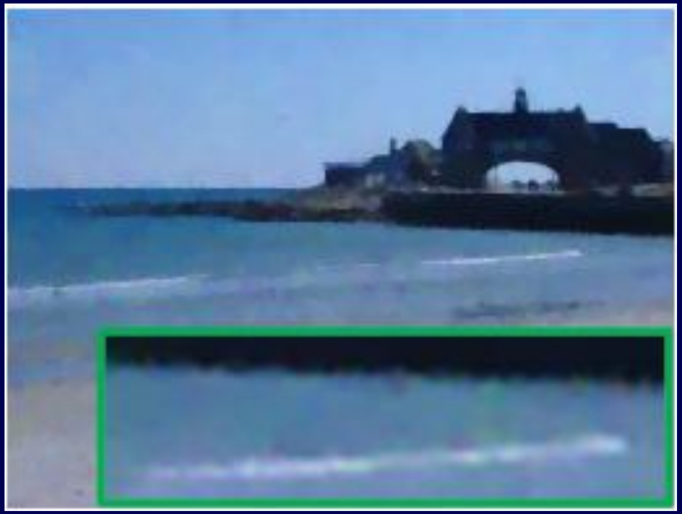

TV denoised

Vectorial Mumford-Shah

Strekalovskiy, Chambolle, Cremers, CVPR ‘12

Channelwise versus Vectorial

Input image

Channelwise MS

Vectorial MS

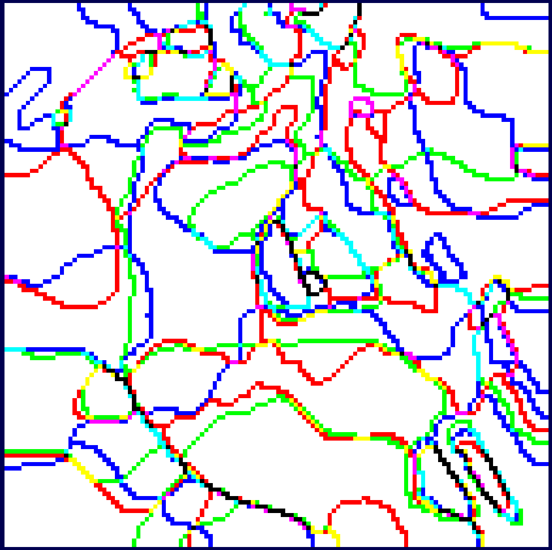

Jump set S_u

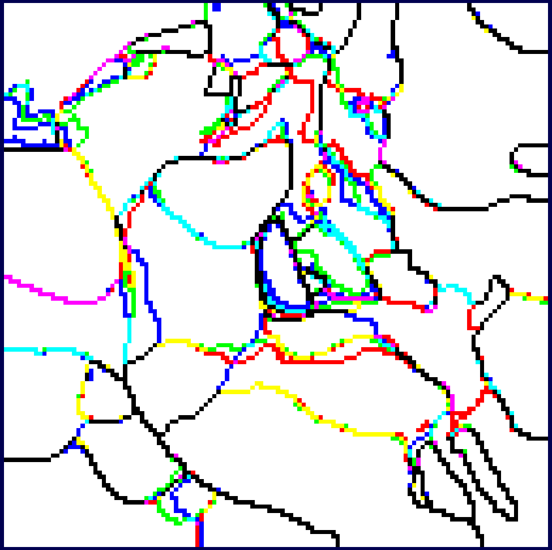

Jump set S_u

Channelwise versus Vectorial

Input image

Channelwise MS

Vectorial MS

Strekalovskiy, Chambolle, Cremers, CVPR '12

Multiview reconstruction

Super-res.textures

4D reconstruction

Stereo reconstruction

Segmentation

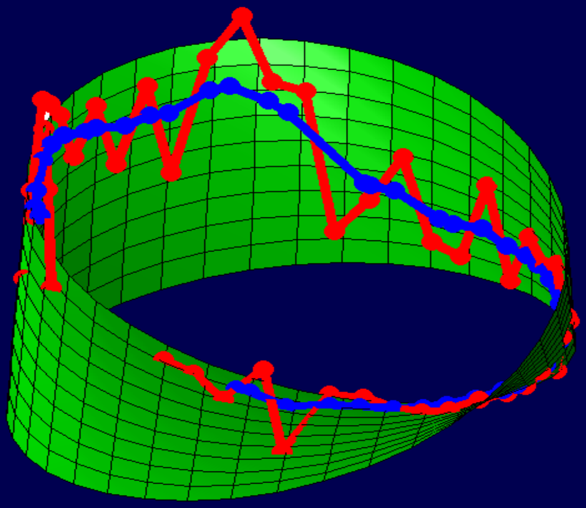

Manifold-valued functions

color image processing

$$\mathcal{M} = \mathbb{R}^3$$

optical flow estimation

$$\mathcal{M} = \mathbb{R}^2$$

normal field inpainting

$$\mathcal{M} = \mathcal{S}^2$$

Cremers, Strekalovskiy, Siims '12
Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Vectorial Total Variation ($\mathcal{M} = \mathbb{R}^k$)

Separate directions, uncoupled (*Blomgren, Chan, TIP '98*):

$$TV_S(u) := \sum_{i=1}^{k} TV(u_i) \;=\; \sup_{\xi:\Omega\to(\mathbb{E}^d)^k} \sum_{i=1}^{k} \int_\Omega u_i \,\text{div}\xi_i \, dx$$

Separate directions, coupled (*Sapiro, Ringach, TIP '96*):

$$TV_F(u) := \int_\Omega \|\nabla u\|_2 \, dx \;=\; \sup_{\xi:\Omega\to\mathbb{E}^{d\times k}} \sum_{i=1}^{k} \int_\Omega u_i \,\text{div}\xi_i \, dx$$

Shared direction, coupled (*Goldlücke et al., SIIMS '12*):

$$TV_F(u) := \int_\Omega \|\nabla u\|_{\sigma_1} \, dx = \sup_{\xi:\Omega\to\mathbb{E}^d, \eta:\Omega\to\mathbb{E}^k} \sum_{i=1}^{k} \int_\Omega u_i \,\text{div}(\eta_i \xi) \, dx$$

Consider the problem

$$\min_{u:\Omega\to\mathcal{M}} \int_\Omega s\big(x, u(x)\big)\, dx \;+\; TV_{\mathcal{M}}(u),$$

with a Riemannian manifold $\mathcal{M}$.

$$TV_{\mathcal{M}}(u) = \int_{\Omega\setminus S_u} |\nabla u|\, dx + \int_{S_u} d_{\mathcal{M}}(u^-, u^+)\, d\mathcal{H}^{d-1}.$$

geodesic distance on the manifold

Cremers, Strekalovskiy, Siims '12
Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Continuous labeling problem with all points of $\mathcal{M}$:

$$\min_{u':\Omega\to\mathcal{P}(\mathcal{M})} \sup_{p:\Omega\times\mathcal{M}\to\mathbb{R}^d} \int_\Omega \langle u', s\rangle dx + \int_\Omega \langle u', \text{Div}\, p\rangle dx$$

$$\text{s.t. } \|p(x,z_1) - p(x,z_2)\|_2 \leqslant d_{\mathcal{M}}(z_1,z_2), \ \forall z_1, z_2 \in \mathcal{M}, \forall x \in \Omega, \quad (*)$$

Proposition: The pairwise constraints $(*)$ are equivalent to

$$\|D_z p(x,z)\|_\sigma \leqslant 1, \quad \forall z \in \mathcal{M}, \forall x \in \Omega$$

with spectral norm $\|M\|_\sigma = \sup_{v\in T_z\mathcal{M}} \frac{\|\langle M, v\rangle_{T_z\mathcal{M}}\|_2}{\|v\|_{T_z\mathcal{M}}}$ for $M \in (T_z\mathcal{M})^d$

linear number of constraints, respects manifold structure

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Input signal $f : [0, 1] \to \mathbb{R}^2$

Denoising u using $TV_{\mathbb{R}^2}$

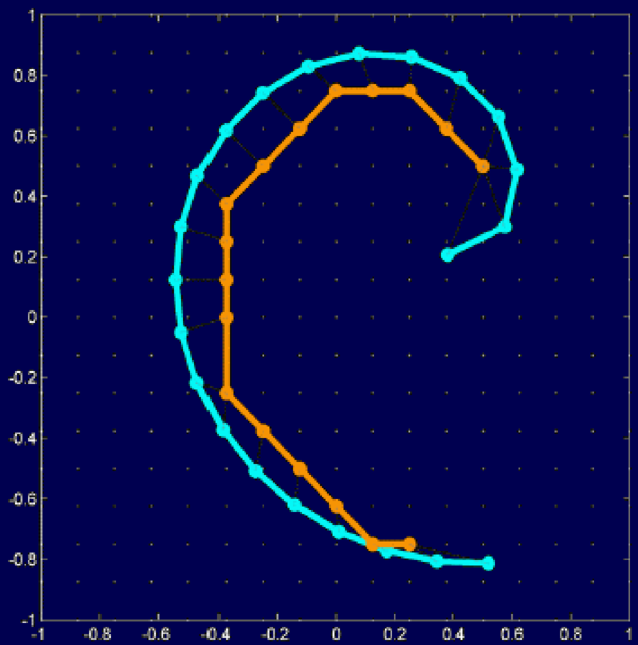

Finite Labeling

8-Neighborhood

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Total Variation for Functions with Values in a Manifold

$$\mathcal{M} = \mathbb{R}^2$$

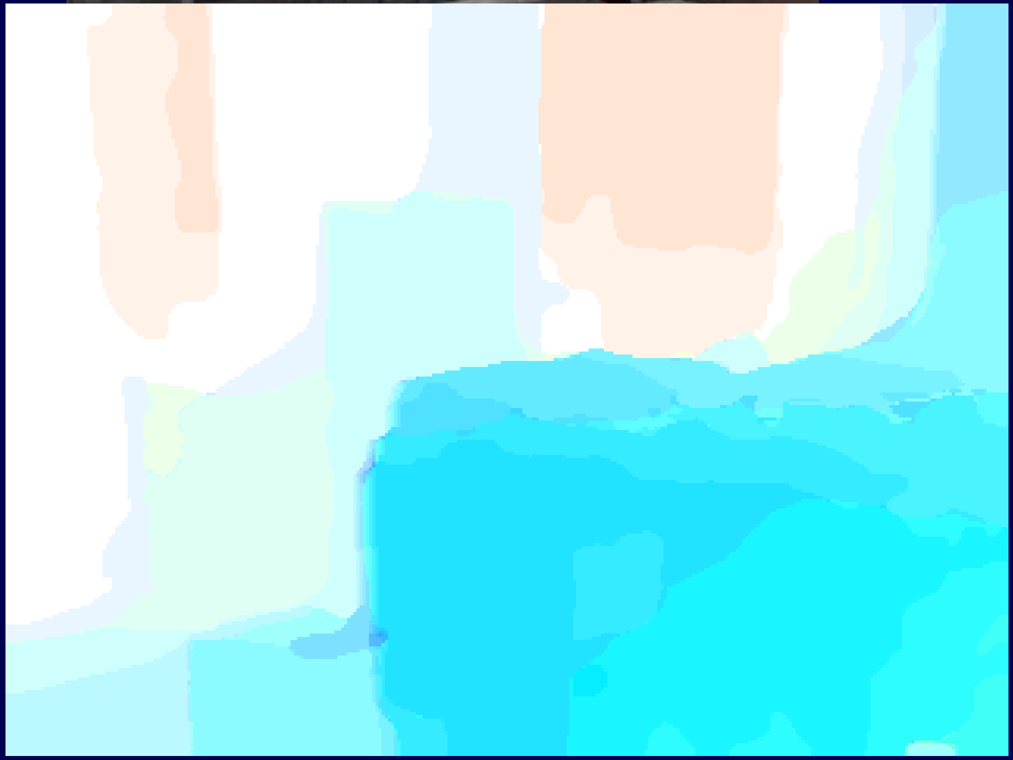

flow with finite labeling

flow with continuous labeling

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

$$\mathcal{M} = \mathcal{S}^2$$

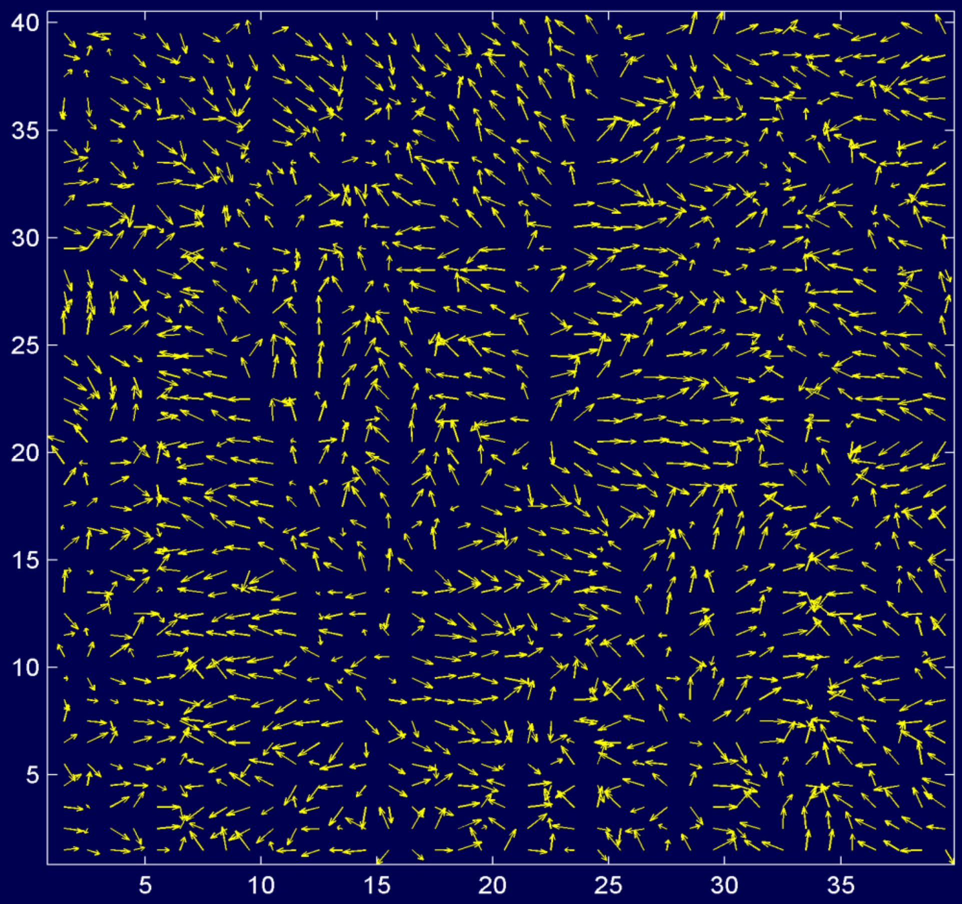

noisy normal field

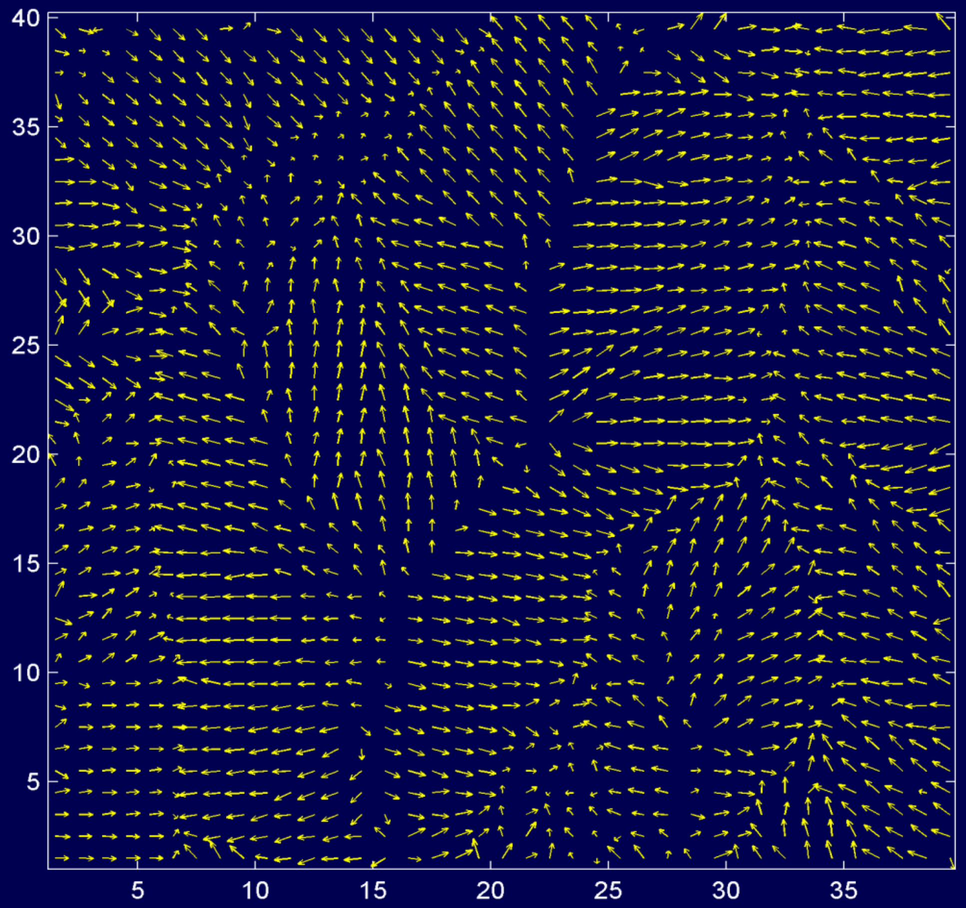

$TV_{\mathcal{S}^2}$ - denoised

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

$$\mathcal{M} = \mathcal{S}^2$$

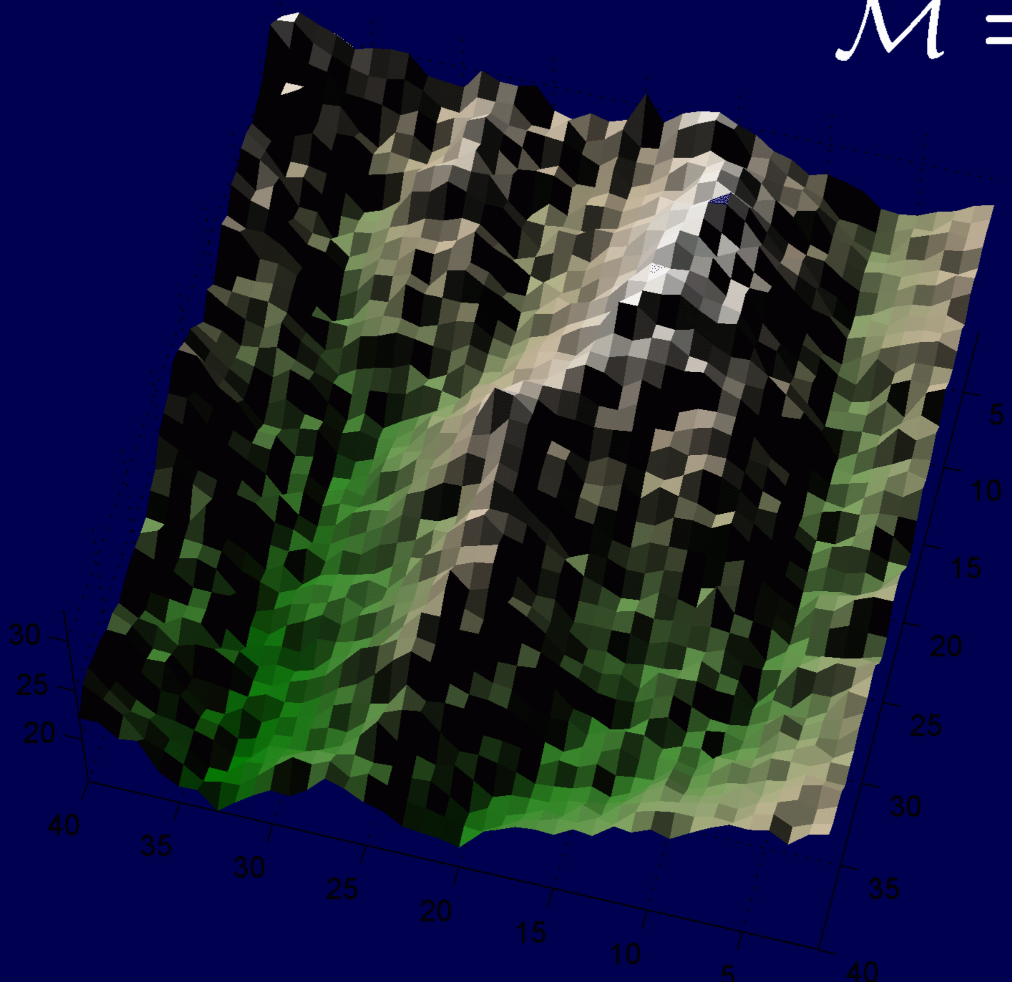

shading with noisy normal field

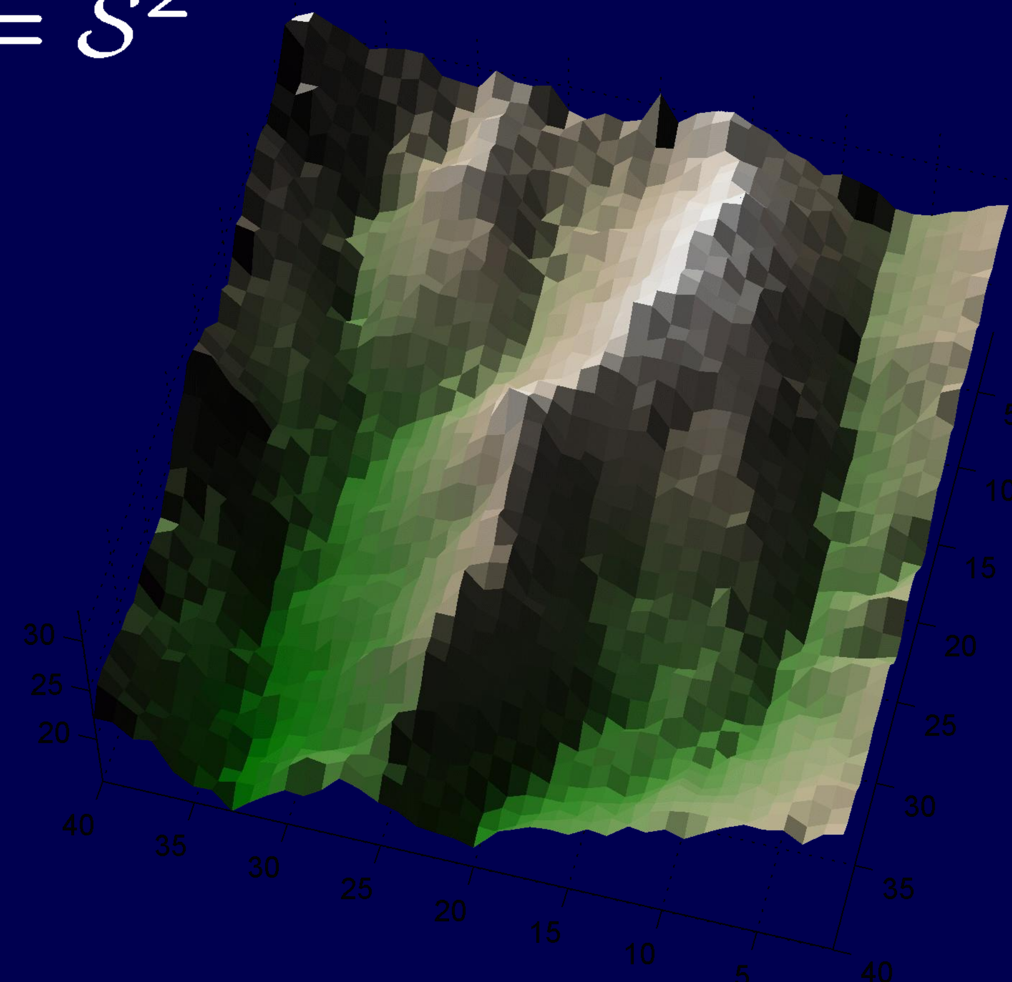

shading with denoised normals

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

$$\mathcal{M} = \mathcal{S}^2$$

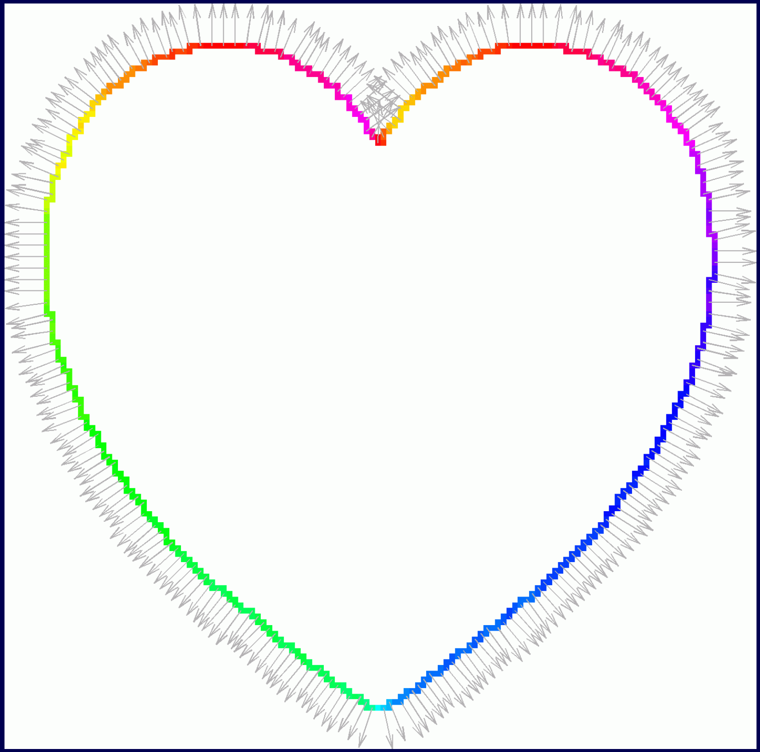

normals on the boundary

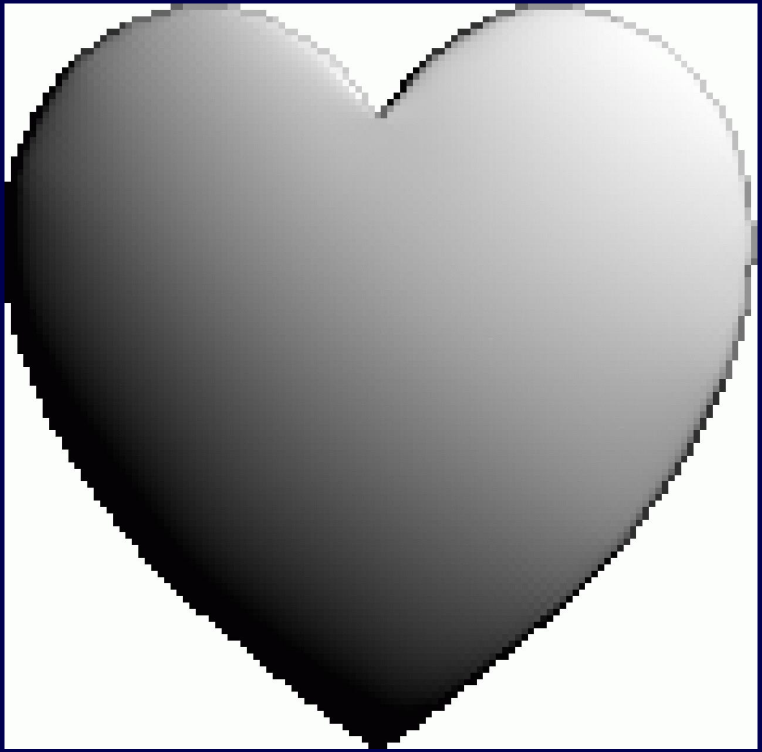

$TV_{\mathcal{S}^2}$ - inpainted normal field

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Conclusion

We can express image analysis problems in terms of convex functionals.

We can minimize these functionals using provably convergent primal-dual algorithms.

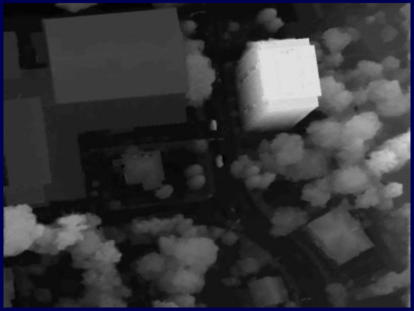

We can define relaxations for functions with values in a manifold using continuous labeling.

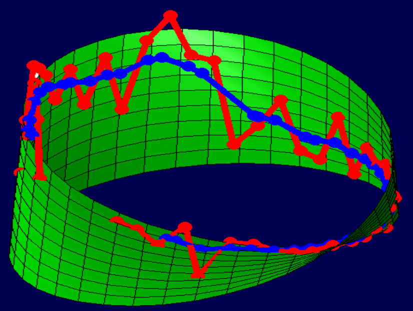

Solutions are independent of initialization and either optimal or within a bound of the optimum.