Convex Relaxation Methods for Computer Vision

Daniel Cremers

Computer Science & Mathematics

Technical University of Munich

     

Kalin Kolev, Evgeny Strekalovskiy, Jan Stühmer, Martin Oswald, Thomas Pock, Antonin Chambolle

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:

Faugeras, Keriven '98, Duan et al. '04, Yezzi, Soatto '03,
Seitz et al. '06, Hernandez et al. '07, Labatut et al. '07, …

Non-convex energies

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, …

Convex Relaxation Techniques

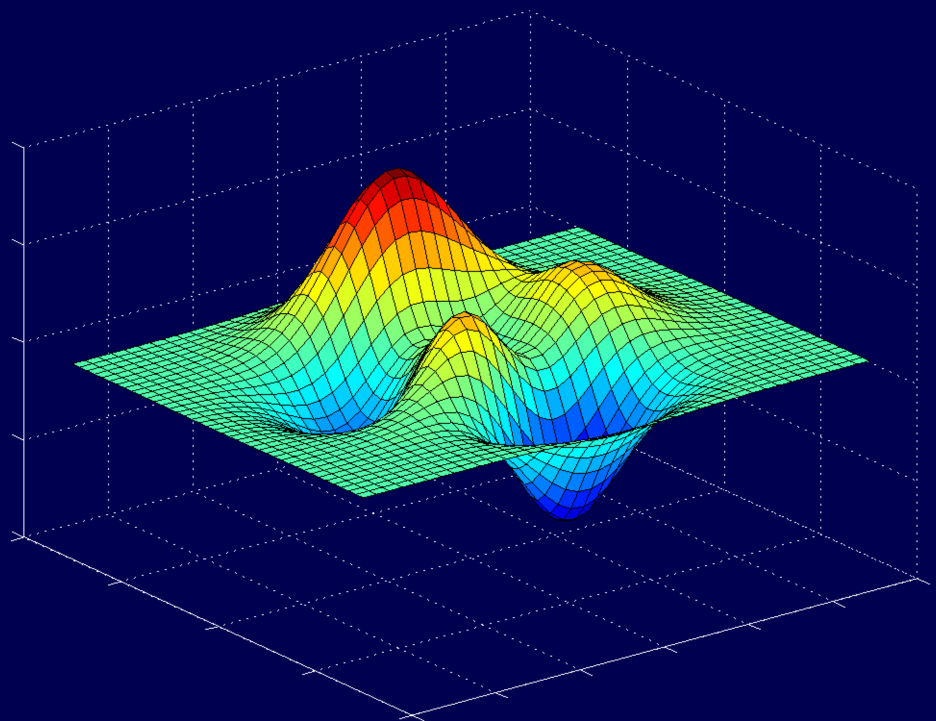

Non-convex energy

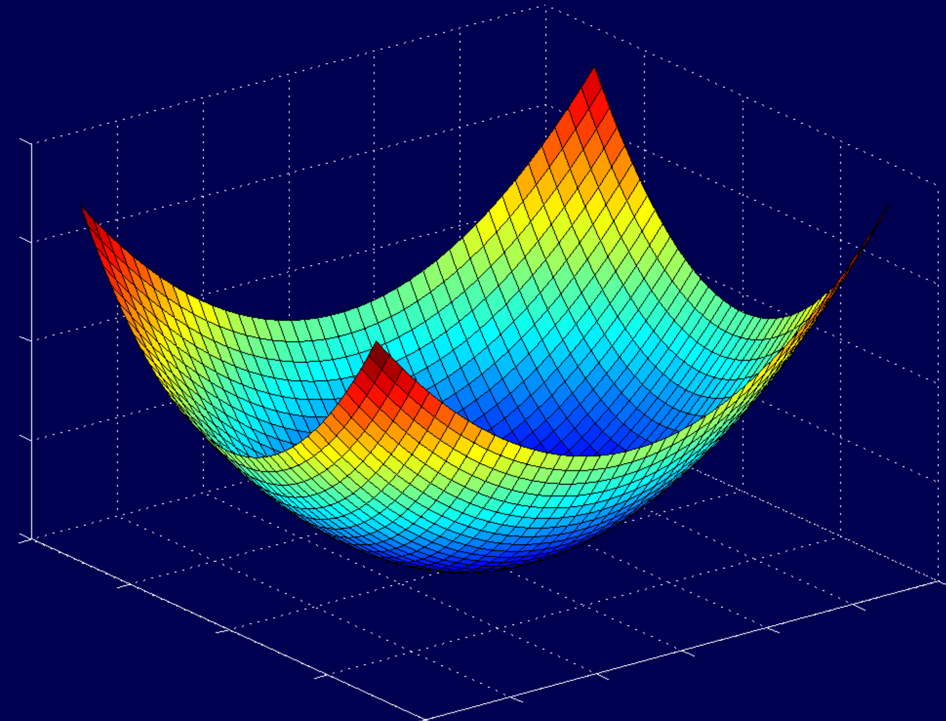

Convex energy

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,…*

Geometric Optimization via
Convex Relaxation

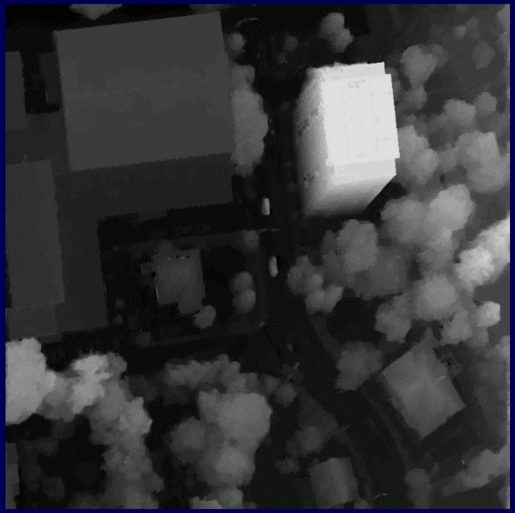

Convex multilabel optimization

Nonconvex regularizers

Geometric Optimization via
Convex Relaxation

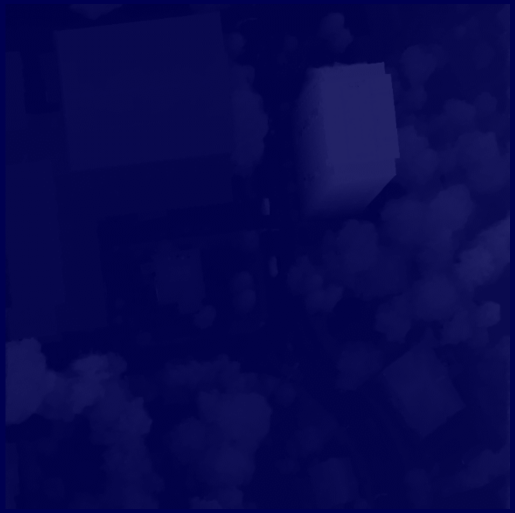

Convex multilabel optimization

Nonconvex regularizers

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

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

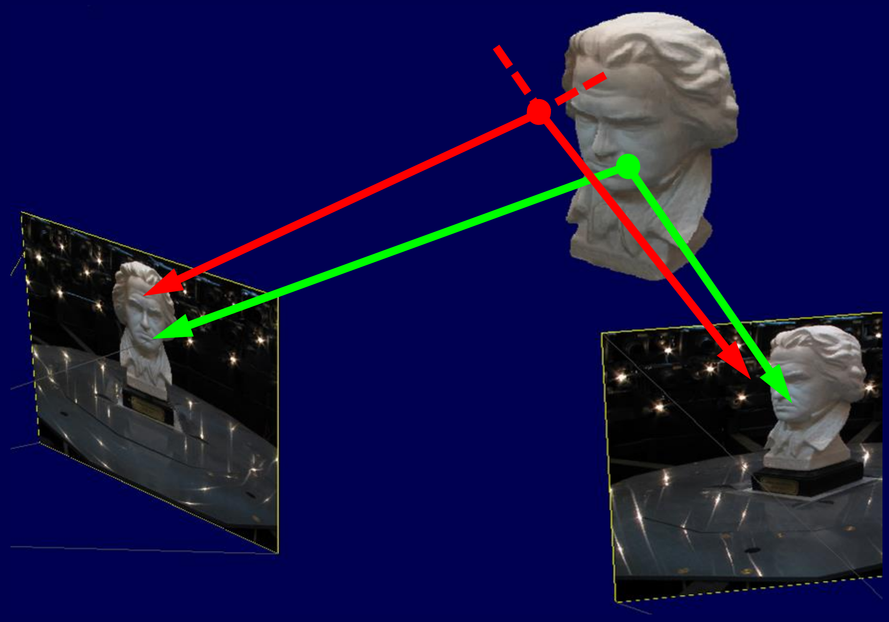

3D Reconstruction: *Faugeras, Keriven '98, Duan et al. '04*

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

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

Silhouette-Consistent Reconstruction

$$\min_S \int_S \rho \, dA \qquad \text{s.t. } \pi_i(S) = S_i \quad \forall i = 1, \dots, n$$

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

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

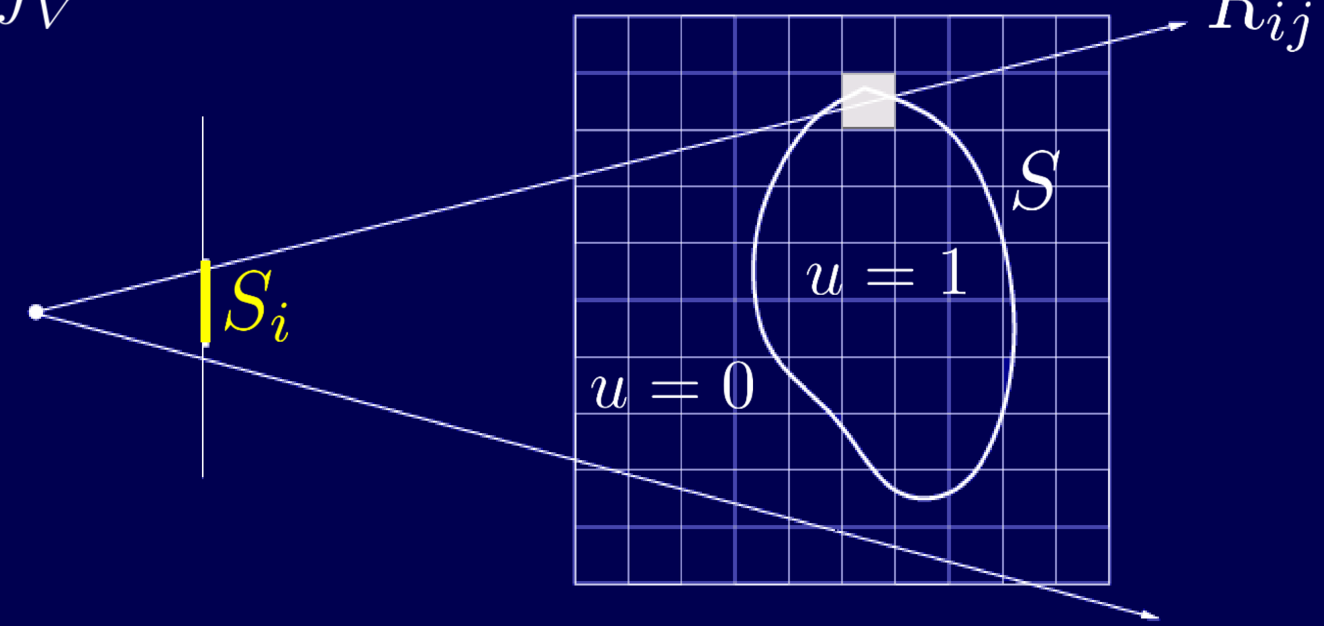

Kolev, Cremers, ECCV ’08, PAMI 2011

Silhouette-Consistent Reconstruction

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

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

$$\min_u \int_V \rho(x) \, |\nabla u| \, dx \quad = \min_{u \in \Sigma} \sup_{|\xi| \leq \rho} \int u \, \text{div}\xi \, dx$$

s.t. ~~$u : V \to \{0, 1\}$~~ $u : V \to [0, 1]$

$$\Sigma = \begin{cases} \displaystyle\int_{R_{ij}} u(x) \, dx \geq 1 & \text{if } j \in S_i \quad \forall i, j \\ \displaystyle\int_{R_{ij}} u(x) \, dx = 0 & \text{if } j \notin S_i \quad \forall i, j \end{cases}$$

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

Kolev, Cremers, ECCV '08, PAMI 2011

An Efficient Saddle Point Solver

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.

Proposition: The primal-dual algorithm

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

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

Reconstructing the Niobids Statues

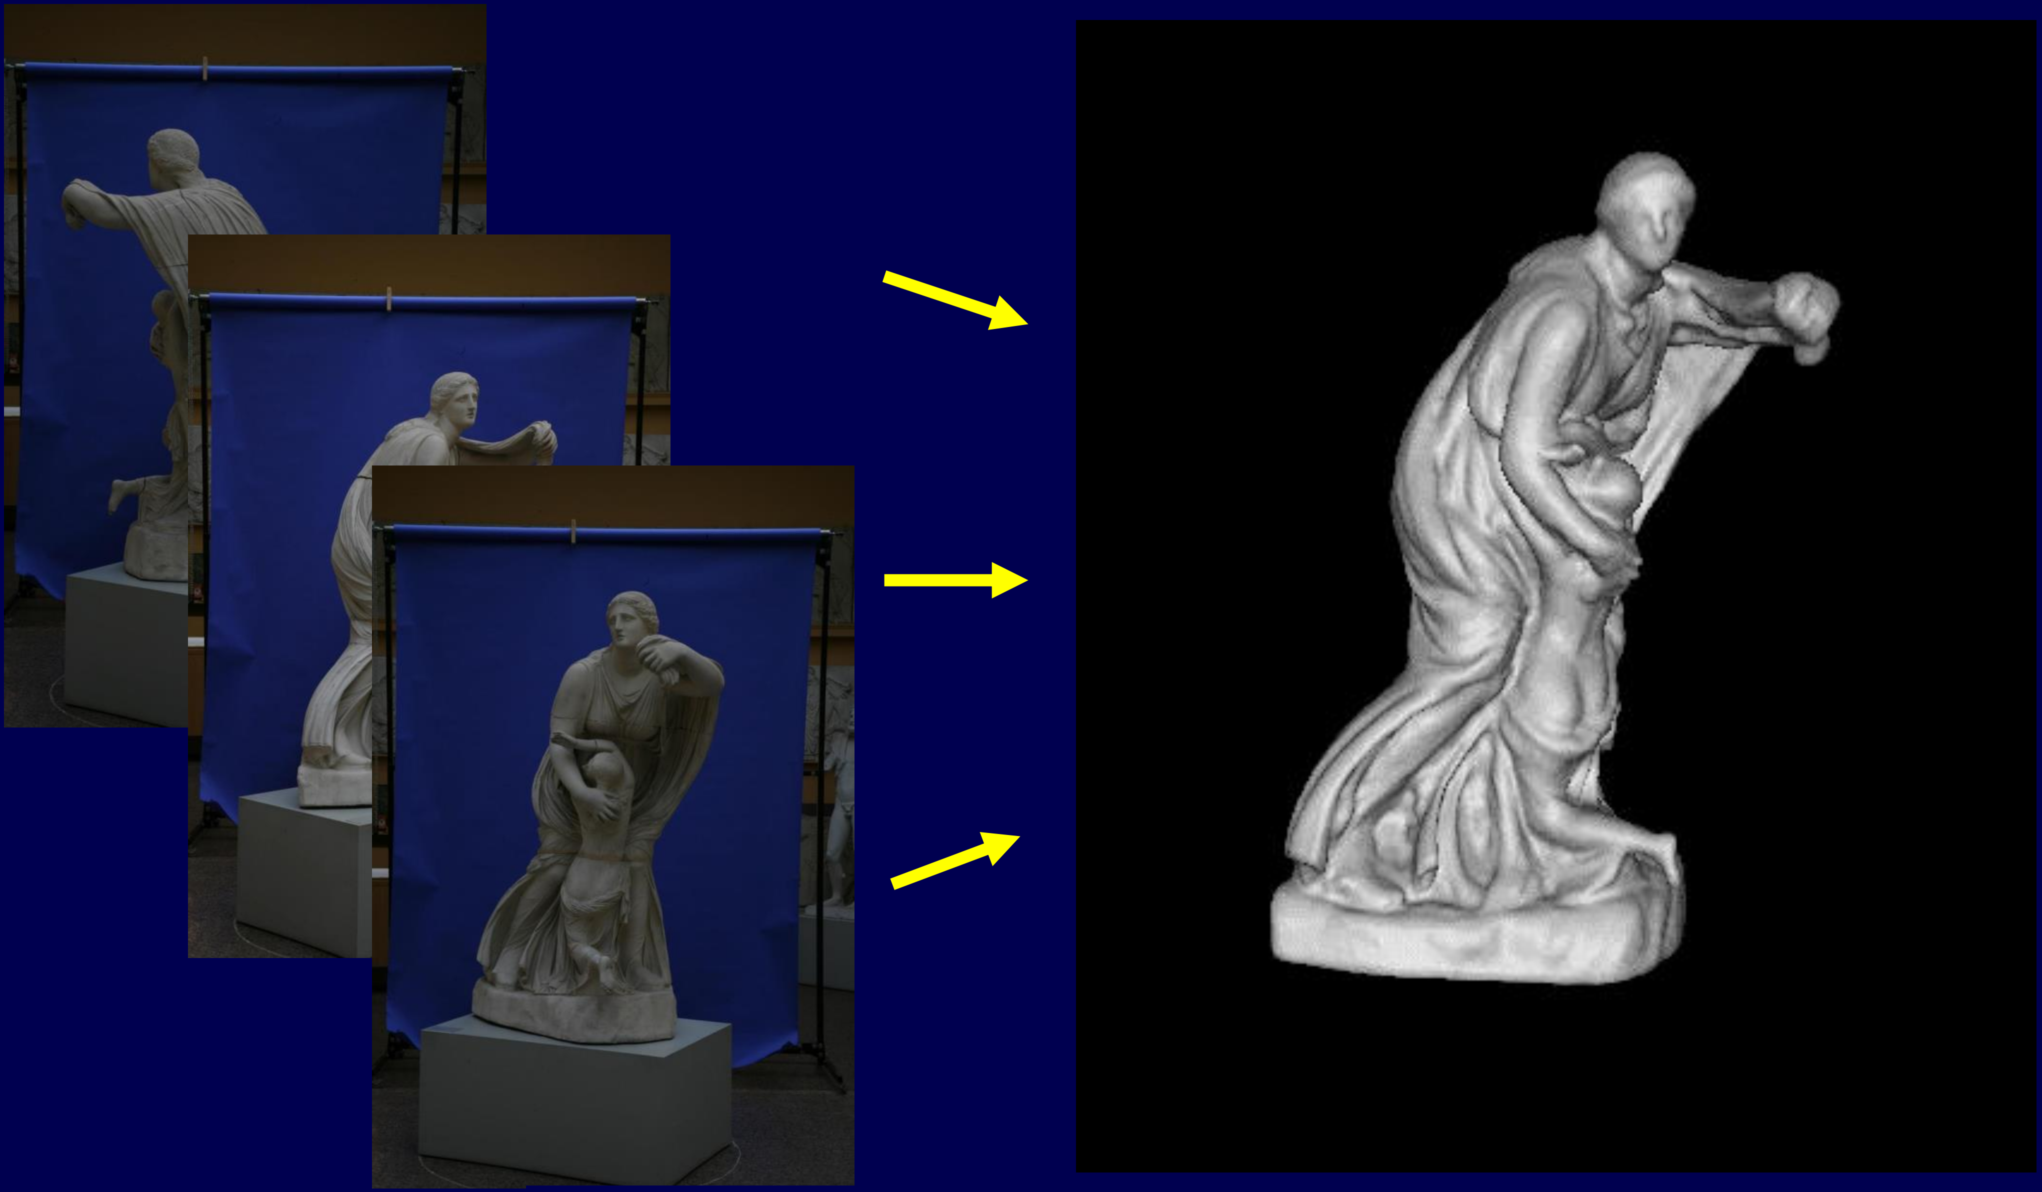

Kolev, Cremers, ECCV ’08, PAMI ‘11

Reconstructing Dynamic Scenes

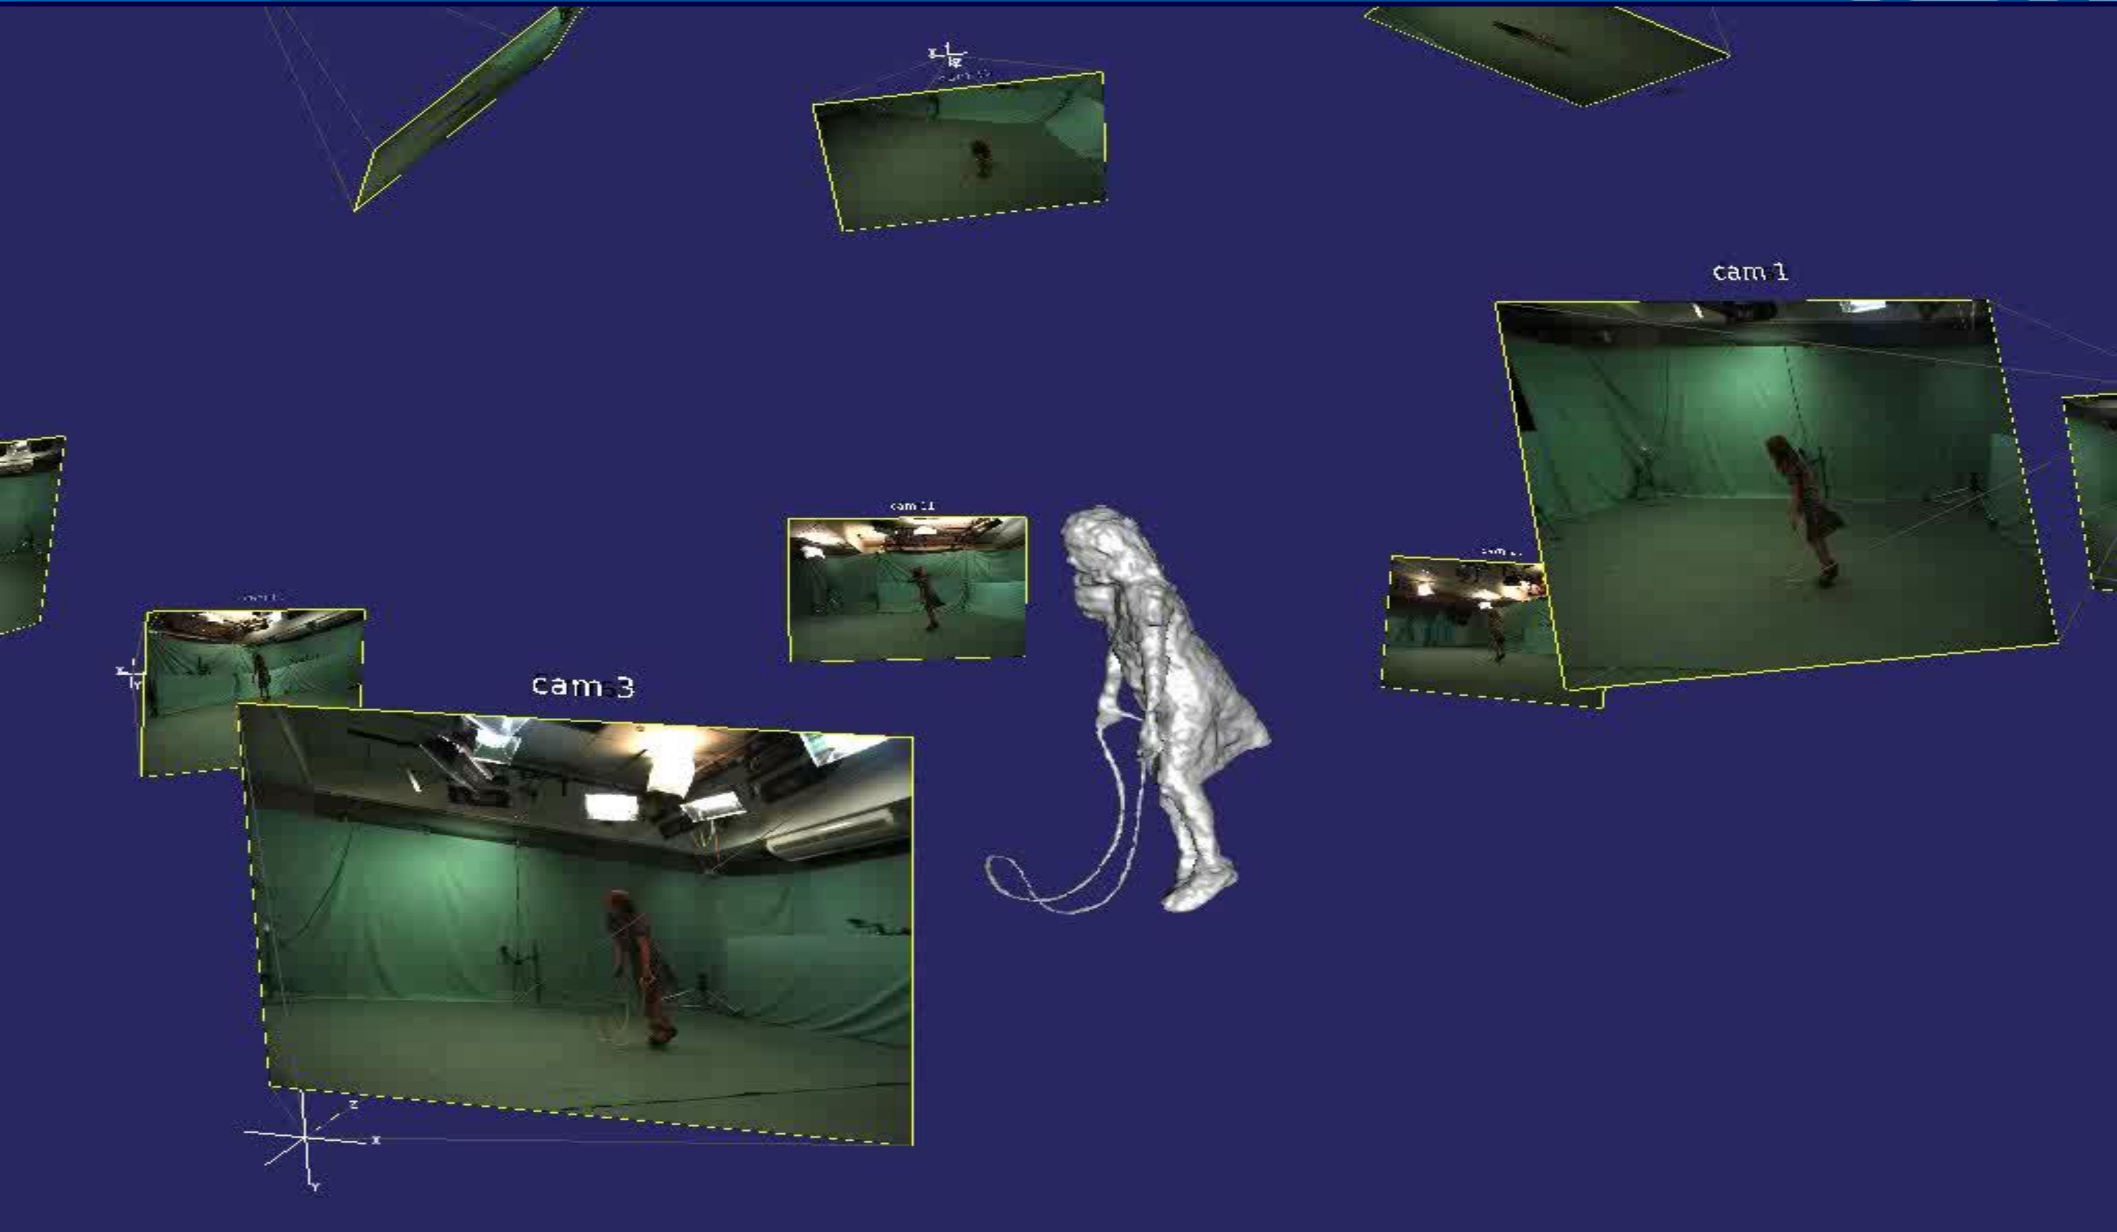

Oswald, Stühmer, Cremers, ECCV '14

Geometric Optimization via
Convex Relaxation

Convex multilabel optimization

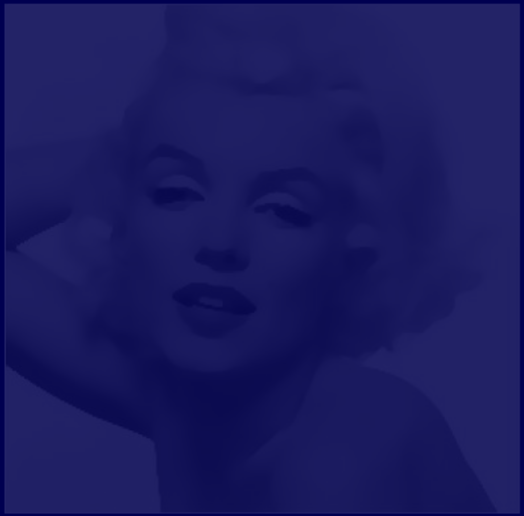

Nonconvex regularizers

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

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

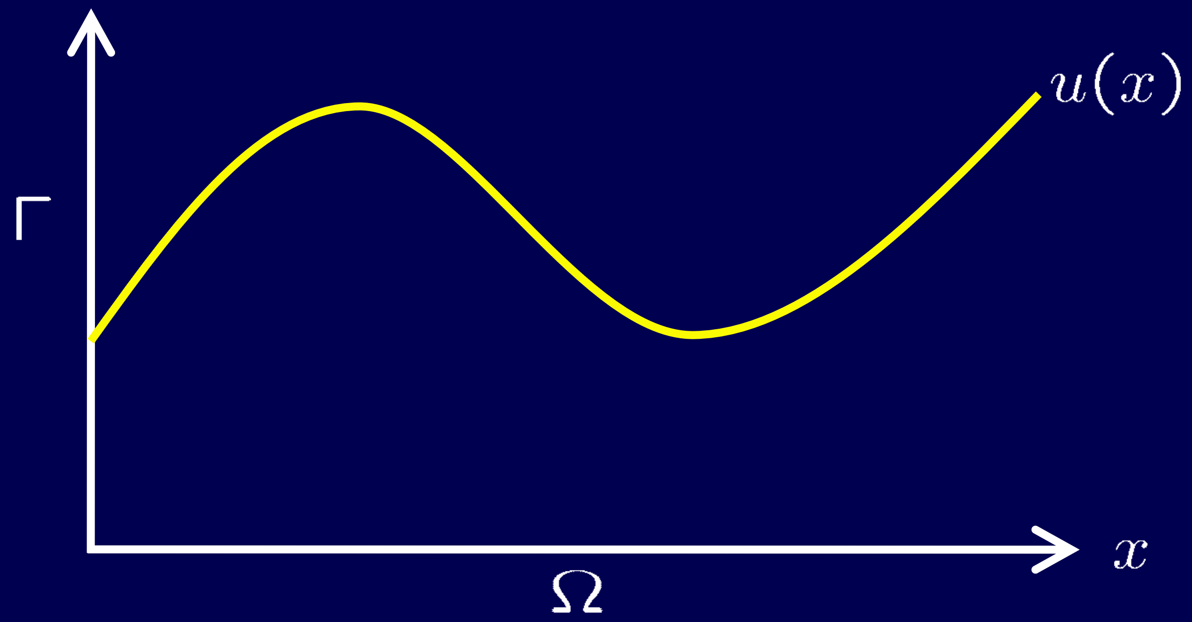

Ishikawa PAMI 2003

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

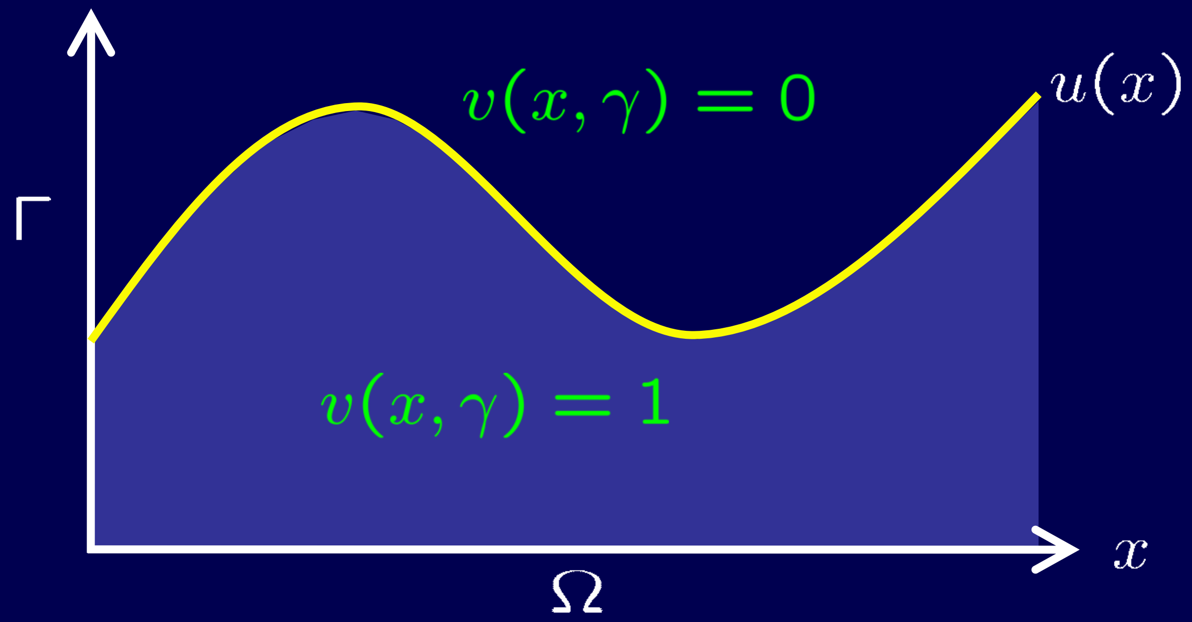

Pock , Schoenemann, Graber, Bischof, Cremers ECCV '08

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

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

nonconvex functional

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

Proposition 1: Minimizing $(*)$ is equivalent to minimizing $\delta\big(u(x) - \gamma\big)$

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

convex functional

Proposition 2: $(**)$ can be solved optimally by convex relaxation and thresholding.

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

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.

Theorem:

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

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

<u>Theorem:</u>

The functional F fulfills a generalized coarea formula:

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

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

One of two input images I_1, I_2

Courtesy of Microsoft

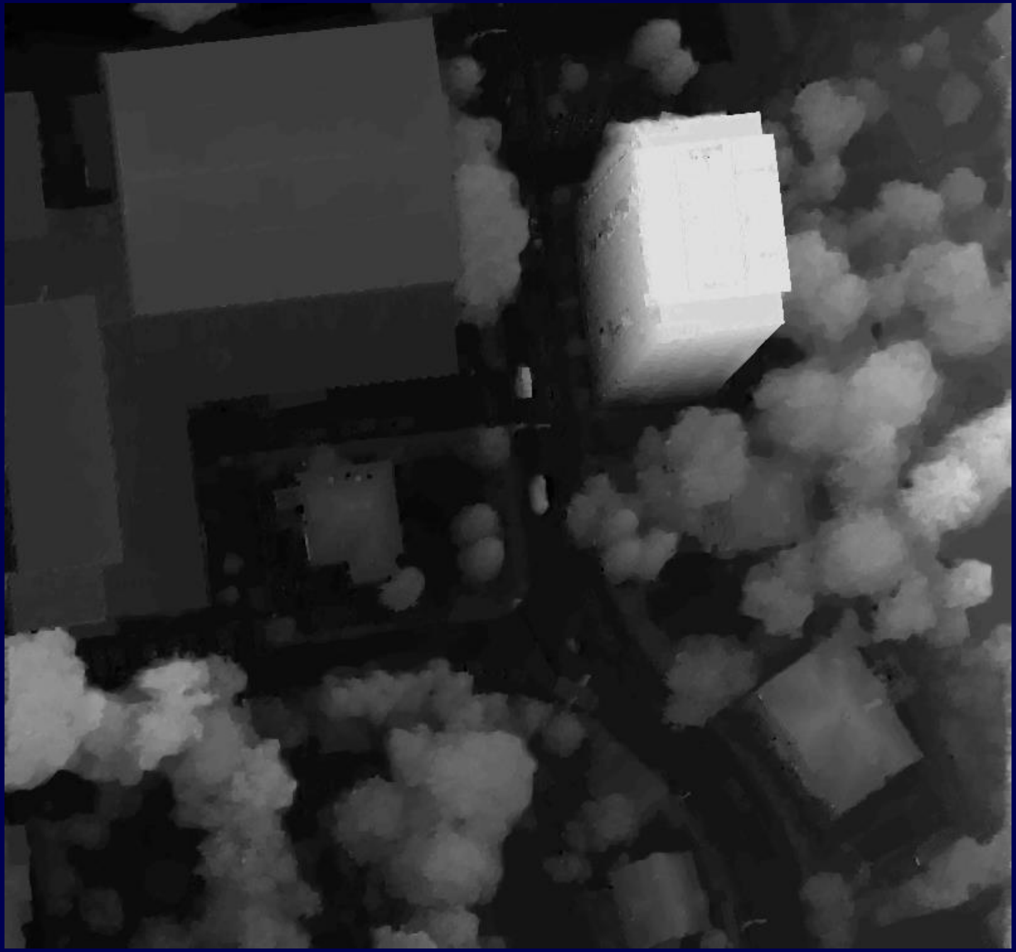

Depth reconstruction

$$\rho\Big(x, u(x)\Big) = \Big|I_1(x) - I_2(x+u)\Big|$$

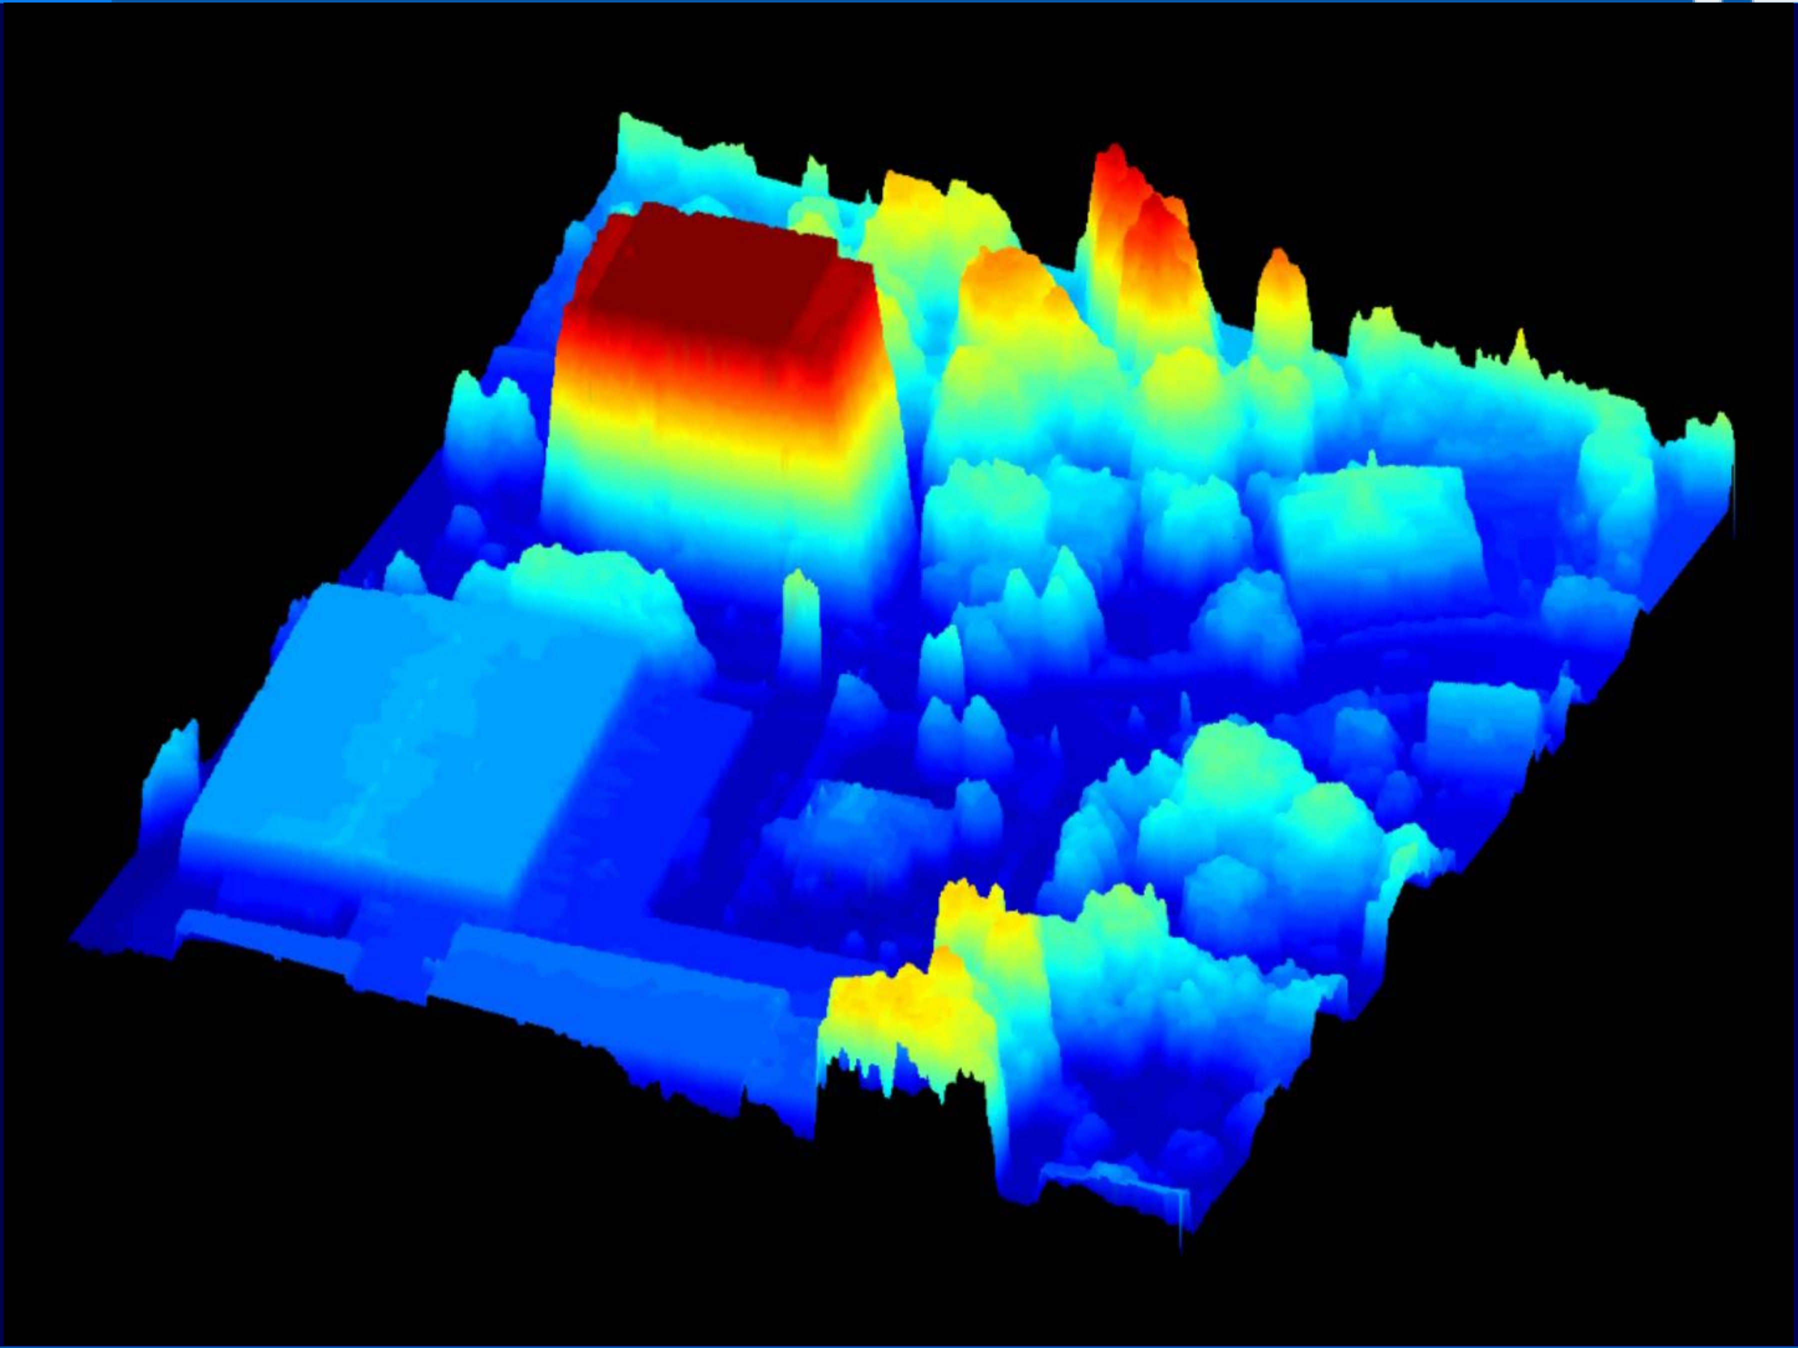

Overview

Geometric Optimization via
Convex Relaxation

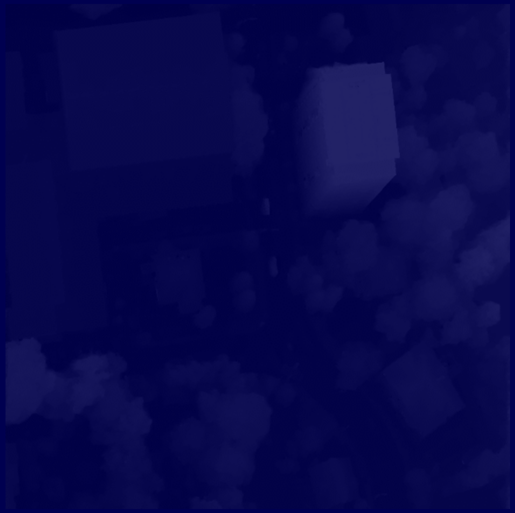

Convex multilabel optimization

Nonconvex regularizers

Nonconvex Regularizers

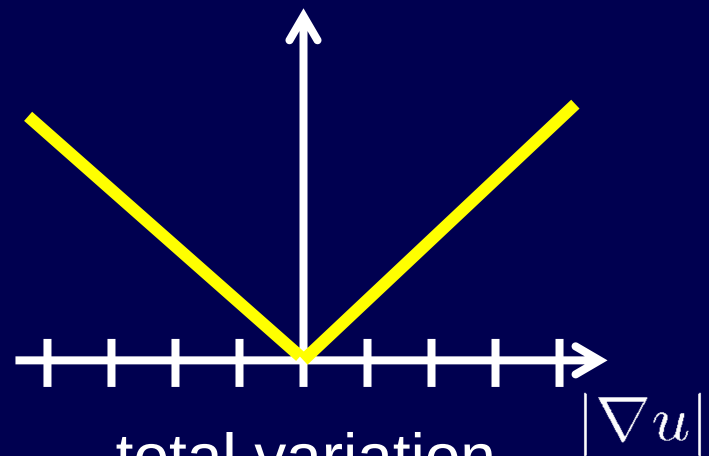

total variation

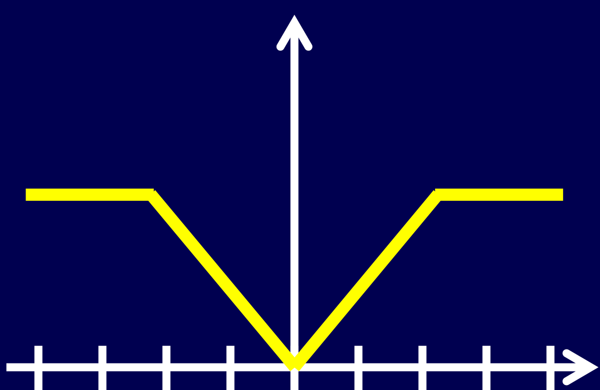

truncated linear

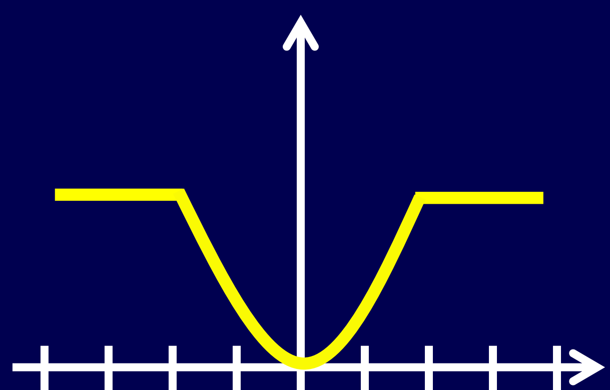

truncated quadratic
Mumford-Shah

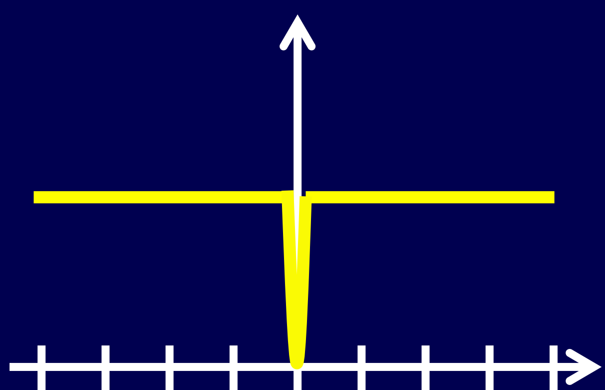

Potts
minimal partitions

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

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

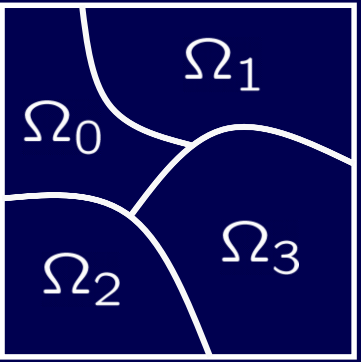

Potts '52, Blake, Zisserman '87, Mumford-Shah '89, Vese, Chan '02

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

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

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

Input color image

10 label segmentation

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

Piecewise Smooth Approximation

Input image

piecewise constant

piecewise smooth

Pock, Cremers, Bischof, Chambolle ICCV ’09

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

Color Mumford-Shah

Input image

Channelwise MS

Vectorial MS

Strekalovskiy, Chambolle, Cremers, CVPR '12

Summary

A variety of originally non-convex optimization problems can be solved by convex relaxation.

The relaxed problems can be solved efficiently with provably convergent primal-dual methods.

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

The lifting approach leads to a drastic increase in memory and runtime.