Variational Methods for Computer Vision

Part 3: Image Segmentation and Convex Relaxation Methods

Daniel Cremers

Computer Science & Mathematics

TU Munich

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

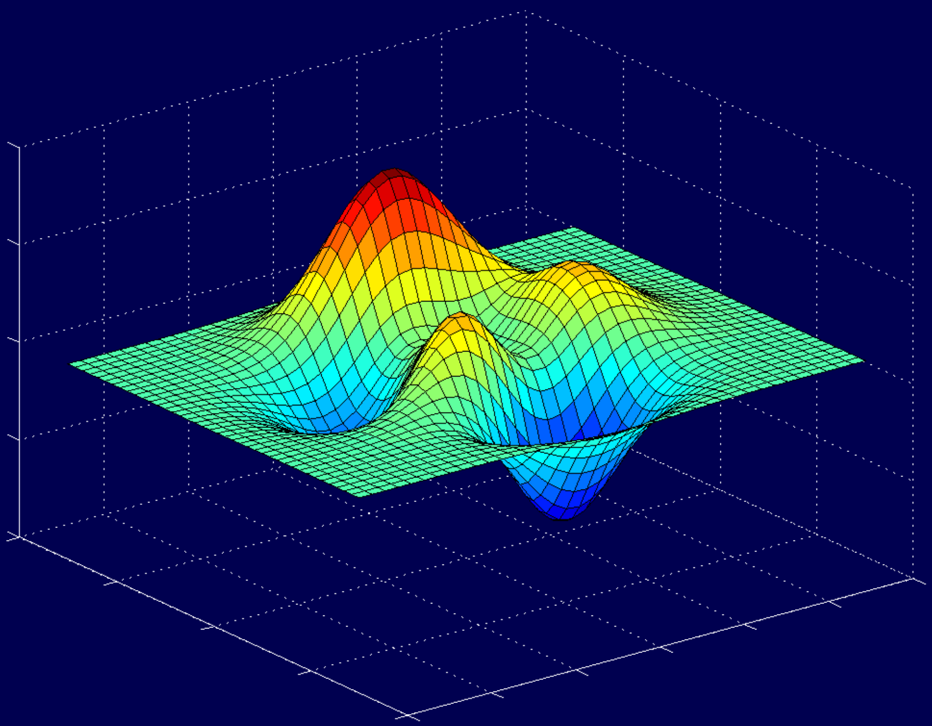

Non-convex energy

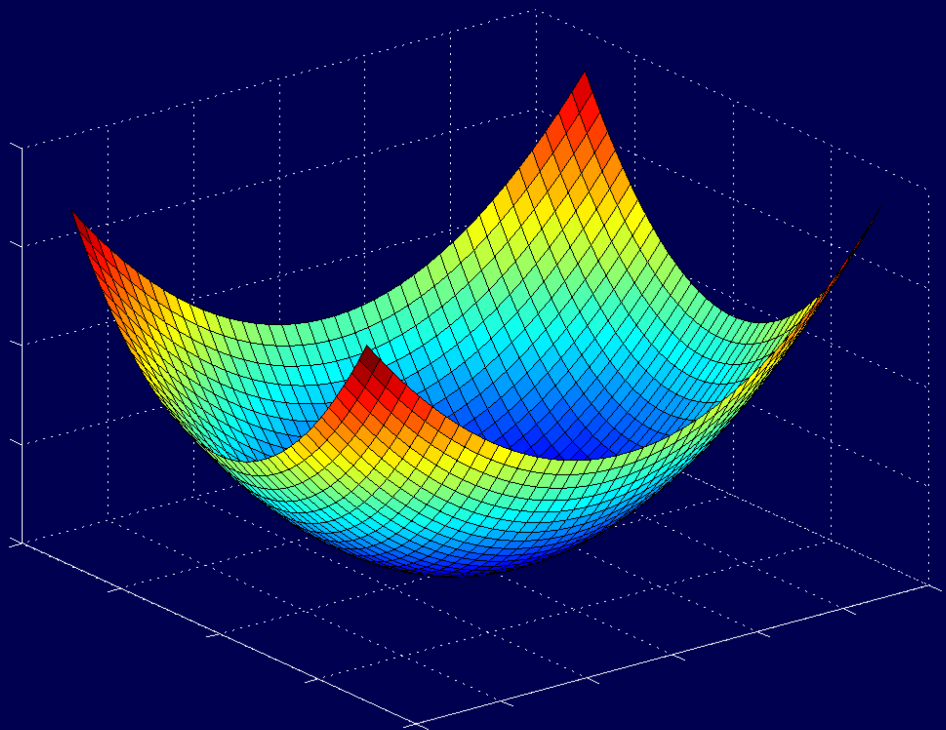

Convex energy

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

Let $\Omega \subset \mathbb{R}^d$ and $f, u : \Omega \to \mathbb{R}^k$.

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

Mumford, Shah '89, Blake, Zisserman '87
Ambrosio, Tortorelli '90, Vese, Chan '02

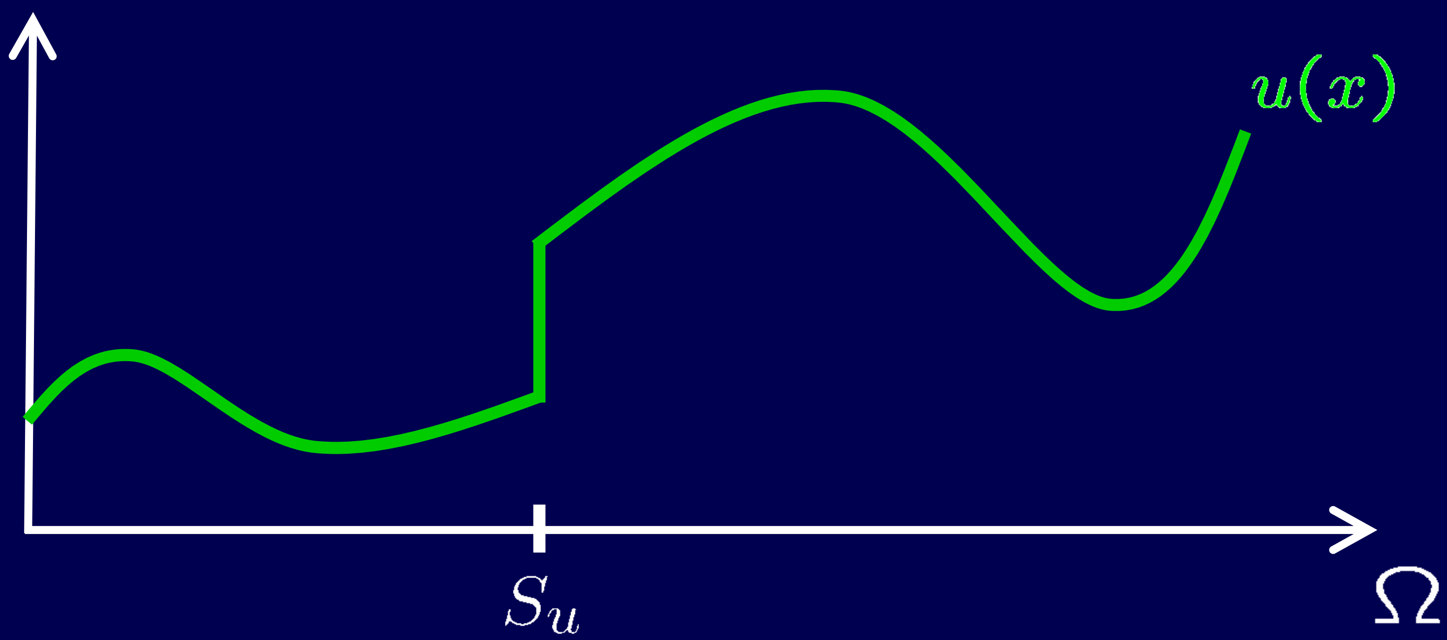

The Mumford-Shah Functional

Let $\Omega \subset \mathbb{R}^d$ and $f, u : \Omega \to \mathbb{R}^k$.

$$E(u) = \int_\Omega |f-u|^2 dx \; + \; \lambda \int_{\Omega \setminus S_u} |\nabla u|^2 dx \; + \; \nu \, |S_u|$$

Mumford, Shah '89, Blake, Zisserman '87
Ambrosio, Tortorelli '90, Vese, Chan '02

Piecewise constant approximation for $\lambda \to \infty$:

$$E(\{\Omega_i, \mu_i\}_i) = \sum_i \int_{\Omega_i} |f(x) - \mu_i|^2 \, dx \; + \; \frac{\nu}{2} \left| \partial \Omega_i \right|$$

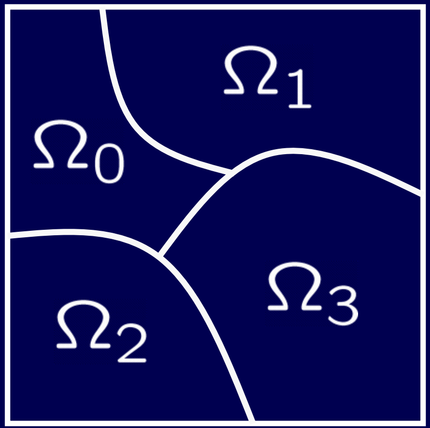

s.t. $\bigcup_i \Omega_i = \Omega$, and $\Omega_i \cap \Omega_j = \emptyset \;\; \forall i \neq j$

Mumford, Shah '89, Chan, Vese '01, Potts '52, Ising '25

Comparison of Regularizers

$$E(u) = \int_\Omega |f-u|^2 dx \; + \; \int_\Omega \psi(\nabla u)\, dx$$

$$\psi(s) = \min(|s|^2, \nu)$$

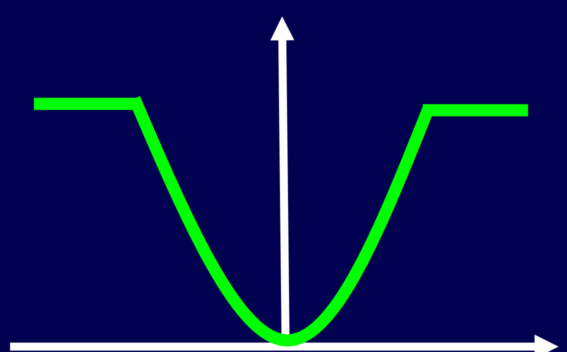

Truncated quadratic

$$\psi(s) = \begin{cases} 0, & \text{if } s=0 \\ \nu, & \text{else} \end{cases}$$

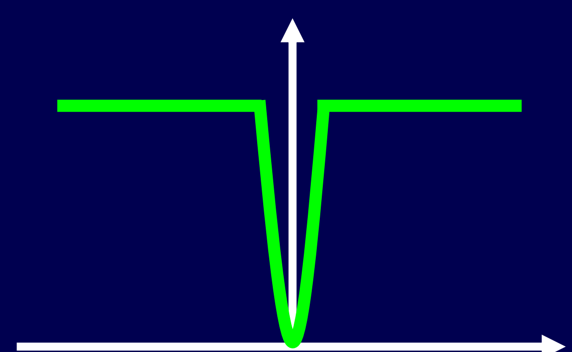

Potts model

$$\psi(s) = |s|$$

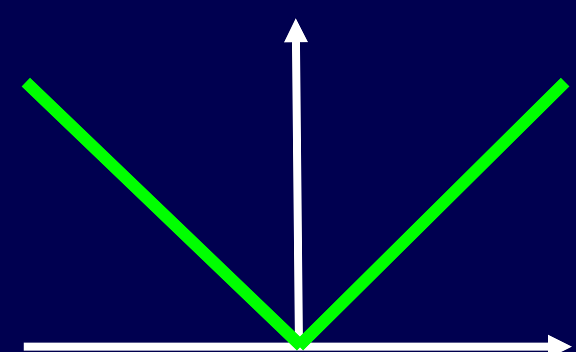

Linear / TV

Input image

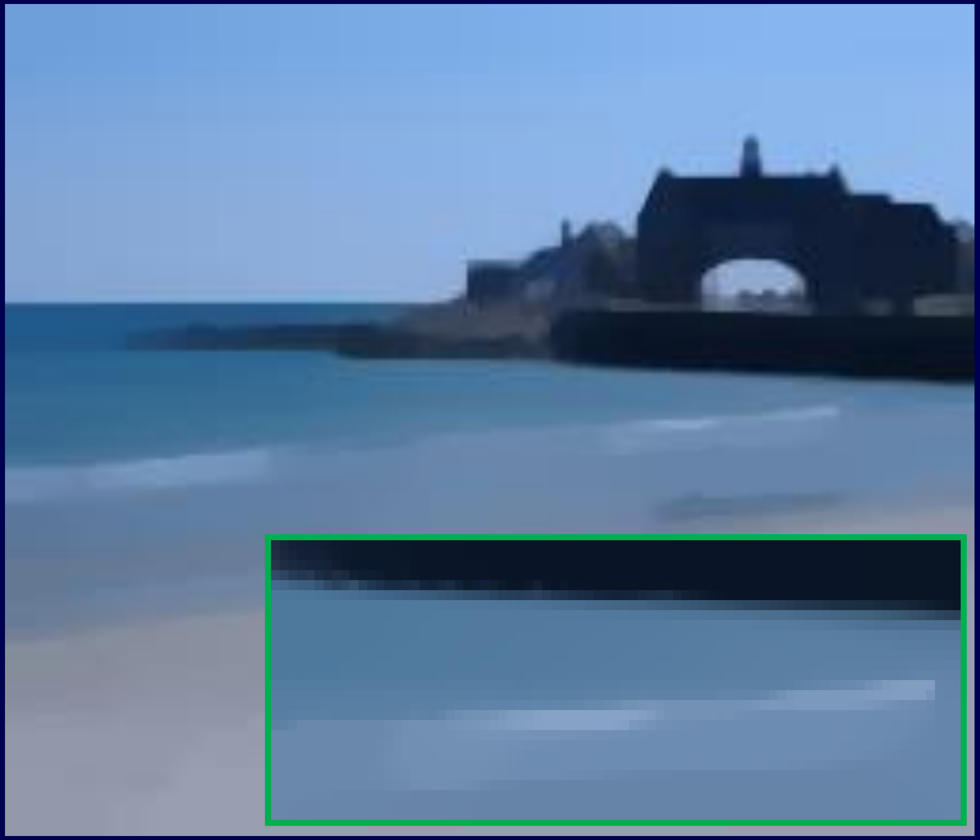

TV-denoised

+ Convex & fast to minimize
- Oversmoothing in flat regions (staircasing)
- Reduces contrast at edges

Convex multilabel optimization

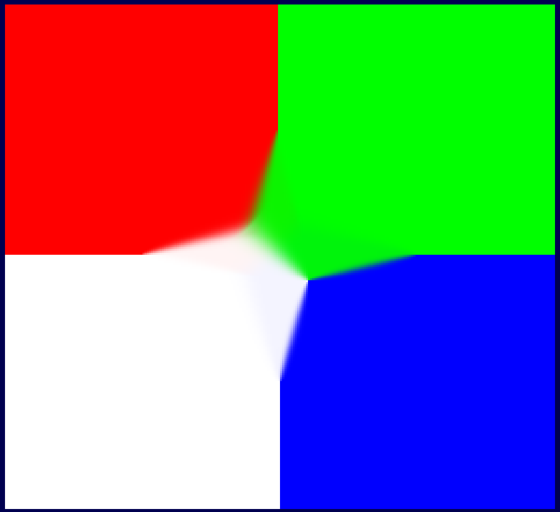

Minimal partitions

Semantic segmentation

Mumford-Shah

Convex multilabel optimization

Minimal partitions

Semantic segmentation

Mumford-Shah

Functional Lifting & Convex Relaxation

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

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

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

Functional Lifting & Convex Relaxation

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

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

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

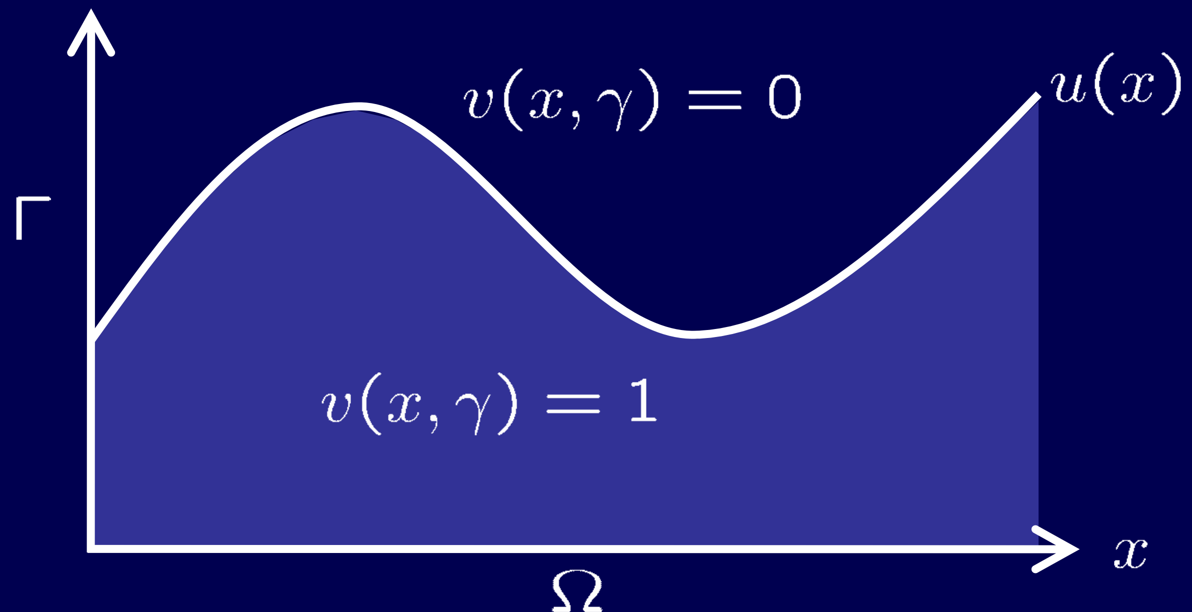

Ishikawa , PAMI ‘03

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

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

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

nonconvex functional

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

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

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

convex functional

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

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

Evolution to Global Minimum

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>

$E(u)$ can be minimized globally by solving the saddle point problem

$$\min_{v \in C} \sup_{\phi \in \mathcal{K}} \int_{\Omega \times \mathbb{R}} \phi \cdot D\, v,$$

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

An Efficient Saddle Point Solver

Given the saddle point problem

$$\min_{v \in C} \max_{\phi \in K} \langle Av, \phi \rangle + \langle g, v \rangle - \langle h, \phi \rangle$$

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

The iterative algorithm

$$\begin{cases} \phi^{n+1} &= \Pi_K\Big(\phi^n + \sigma(A\bar{v}^n - h)\Big) \\ v^{n+1} &= \Pi_C\Big(v^n - \tau(A^*\phi^{n+1} + g)\Big) \\ \bar{v}^{n+1} &= 2\,v^{n+1} - v^n \end{cases}$$

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

Reconstruction from Aerial Images

One of two input images

Courtesy of Microsoft

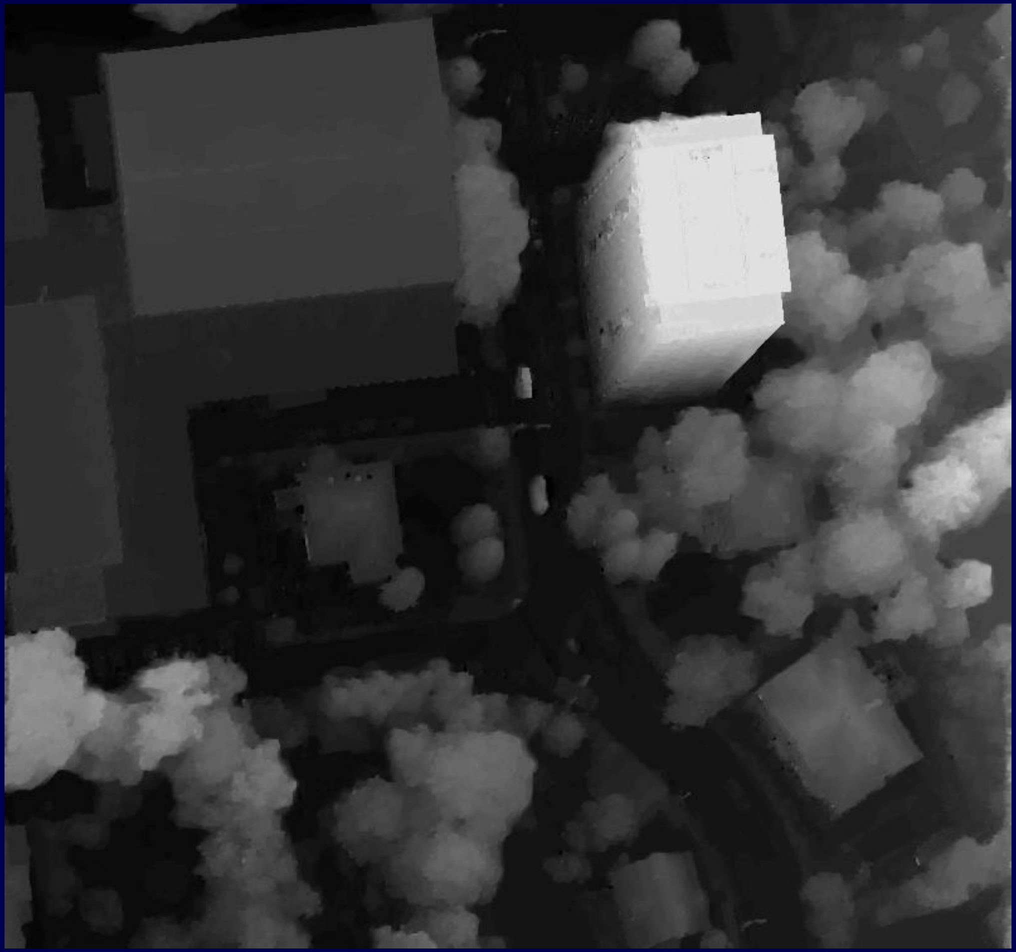

Depth reconstruction

Reconstruction from Aerial Images

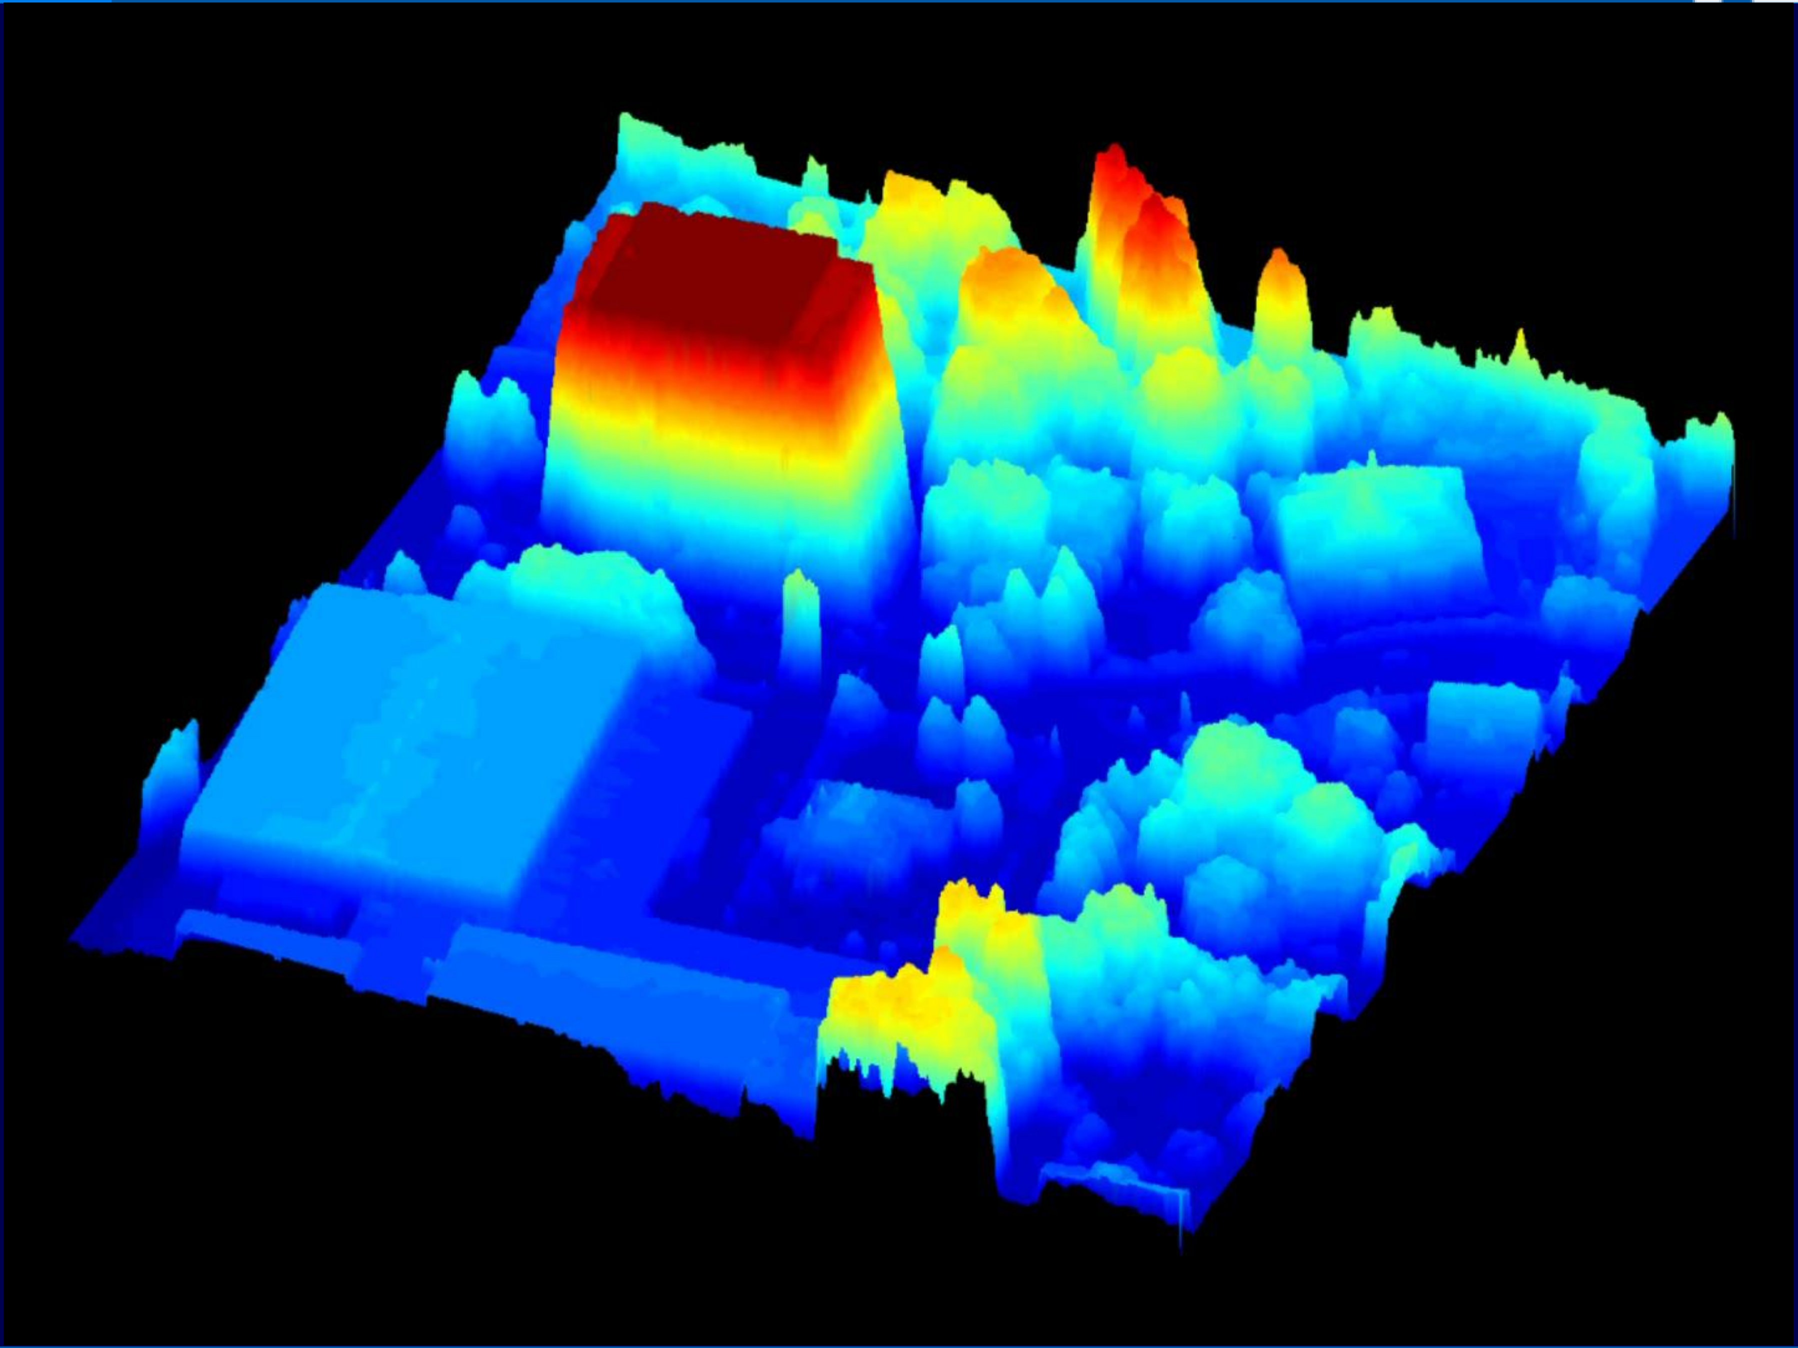

Convex multilabel optimization

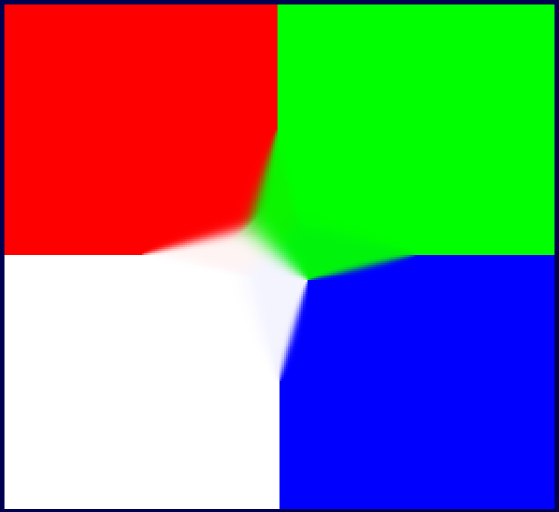

Minimal partitions

Semantic segmentation

Mumford-Shah

The Minimal Partition Problem

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

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

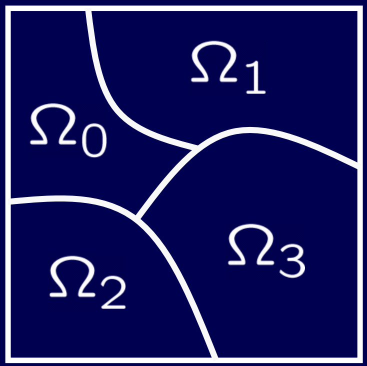

Potts '52, Blake, Zisserman '87, Mumford-Shah '89, Vese, Chan '02

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

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

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

Test Case: The Triple Junction

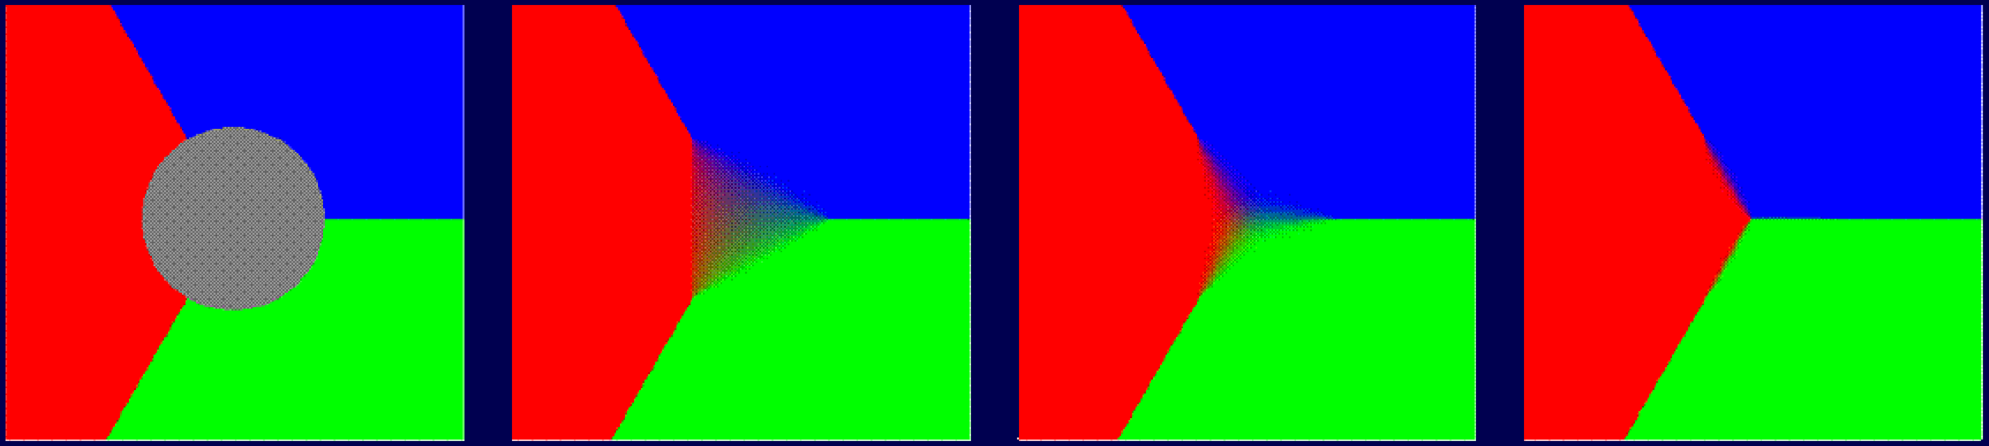

Input image | Lellmann et al. ’08 | Zach et al. ’08 | our approach

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

Chambolle, Cremers, Pock ’08, SIIMS ‘12

Four-Region Case

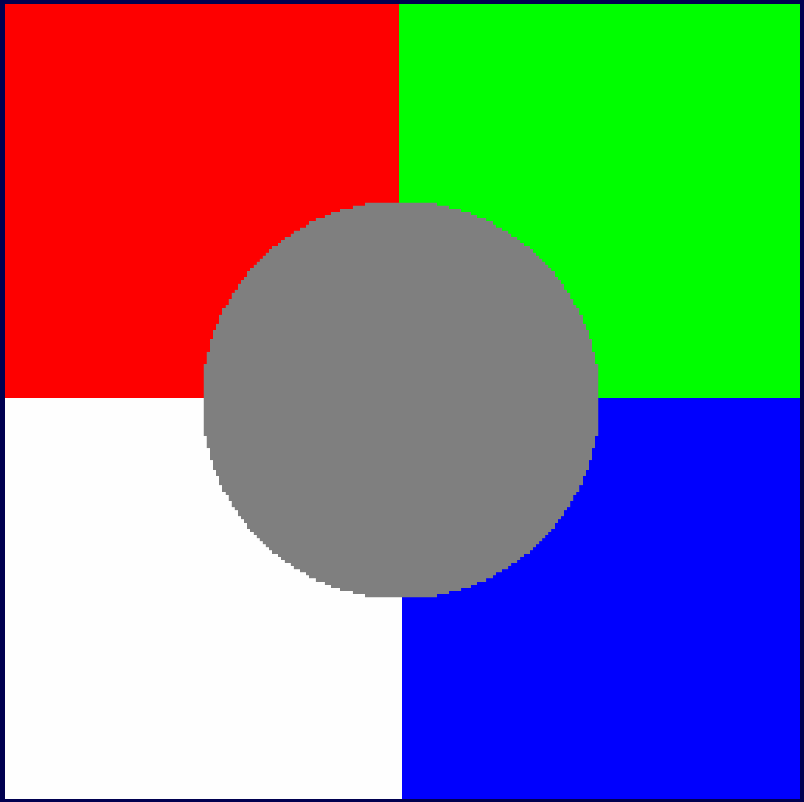

Input image

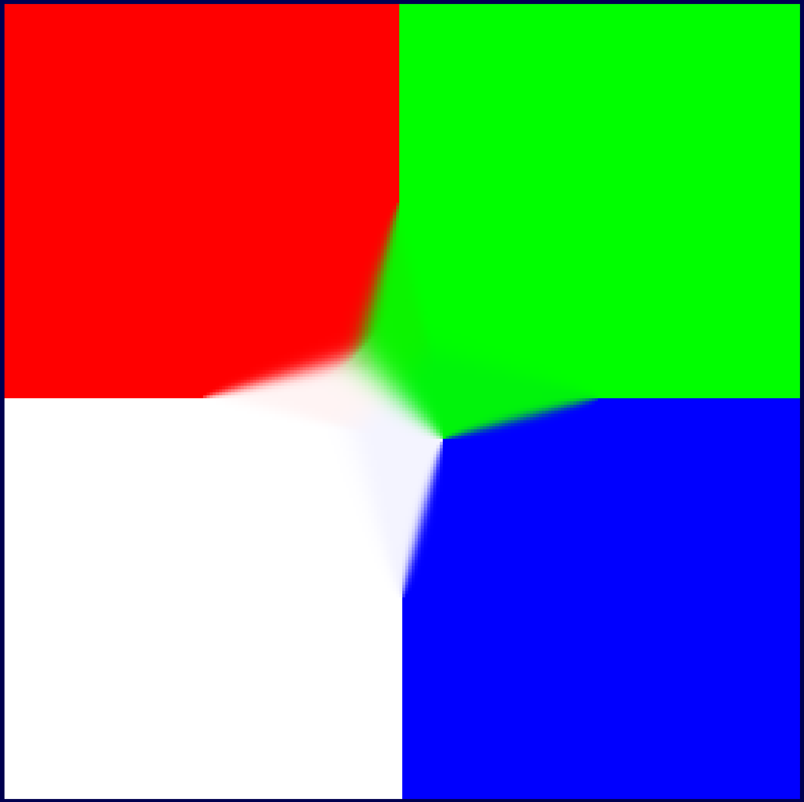

Inpainted

Chambolle, Cremers, Pock '08, SIIMS '12

Minimal Surfaces in 3D

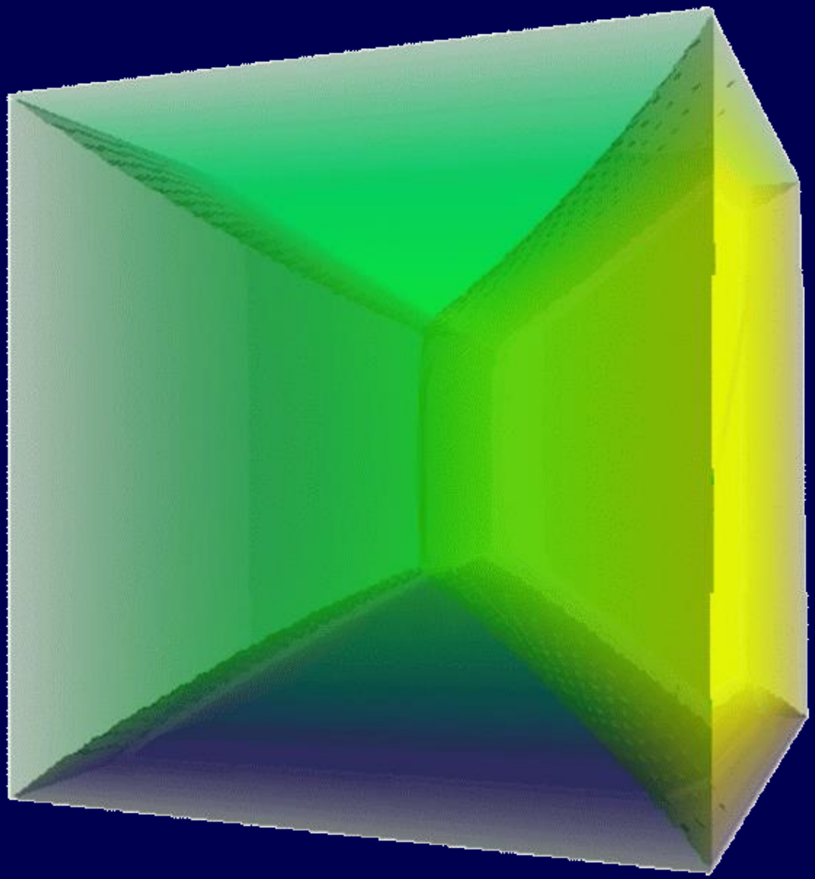

3D min partition inpainting

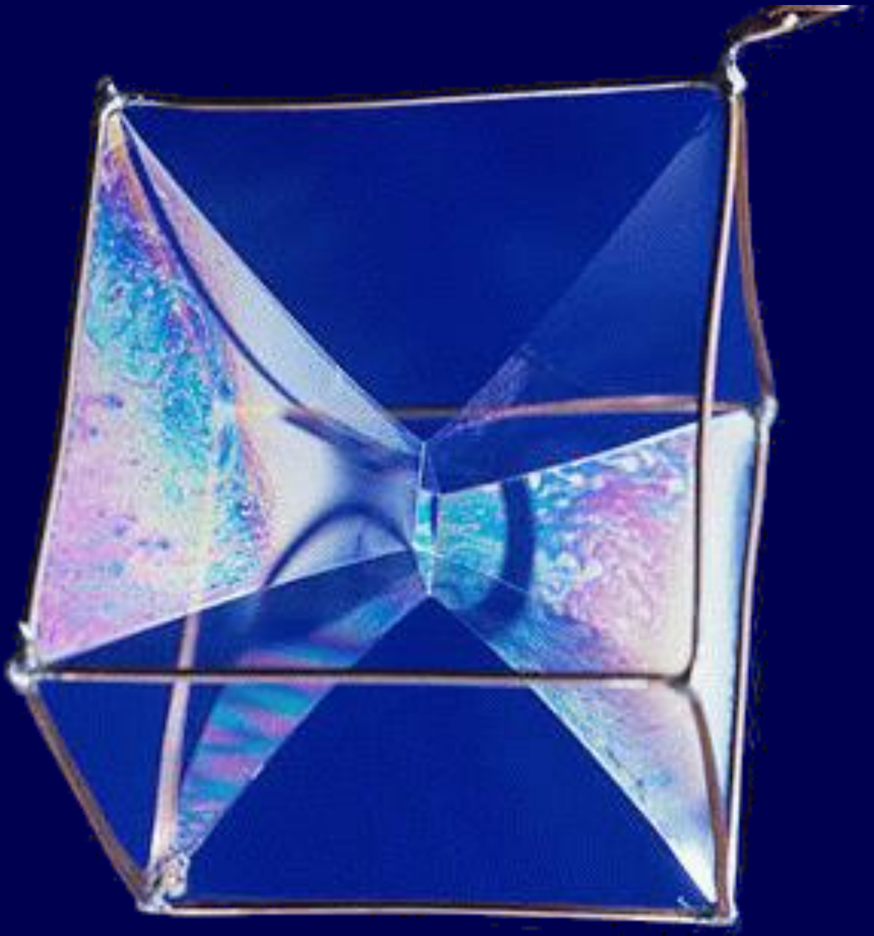

Photograph of a soap film

Chambolle, Cremers, Pock ’08, SIIMS ‘12

The Minimal Partition Problem

Input color image

10 label segmentation

Chambolle, Cremers, Pock ’08, SIIMS ‘12

Nieuwenhuis, Cremers, PAMI '12

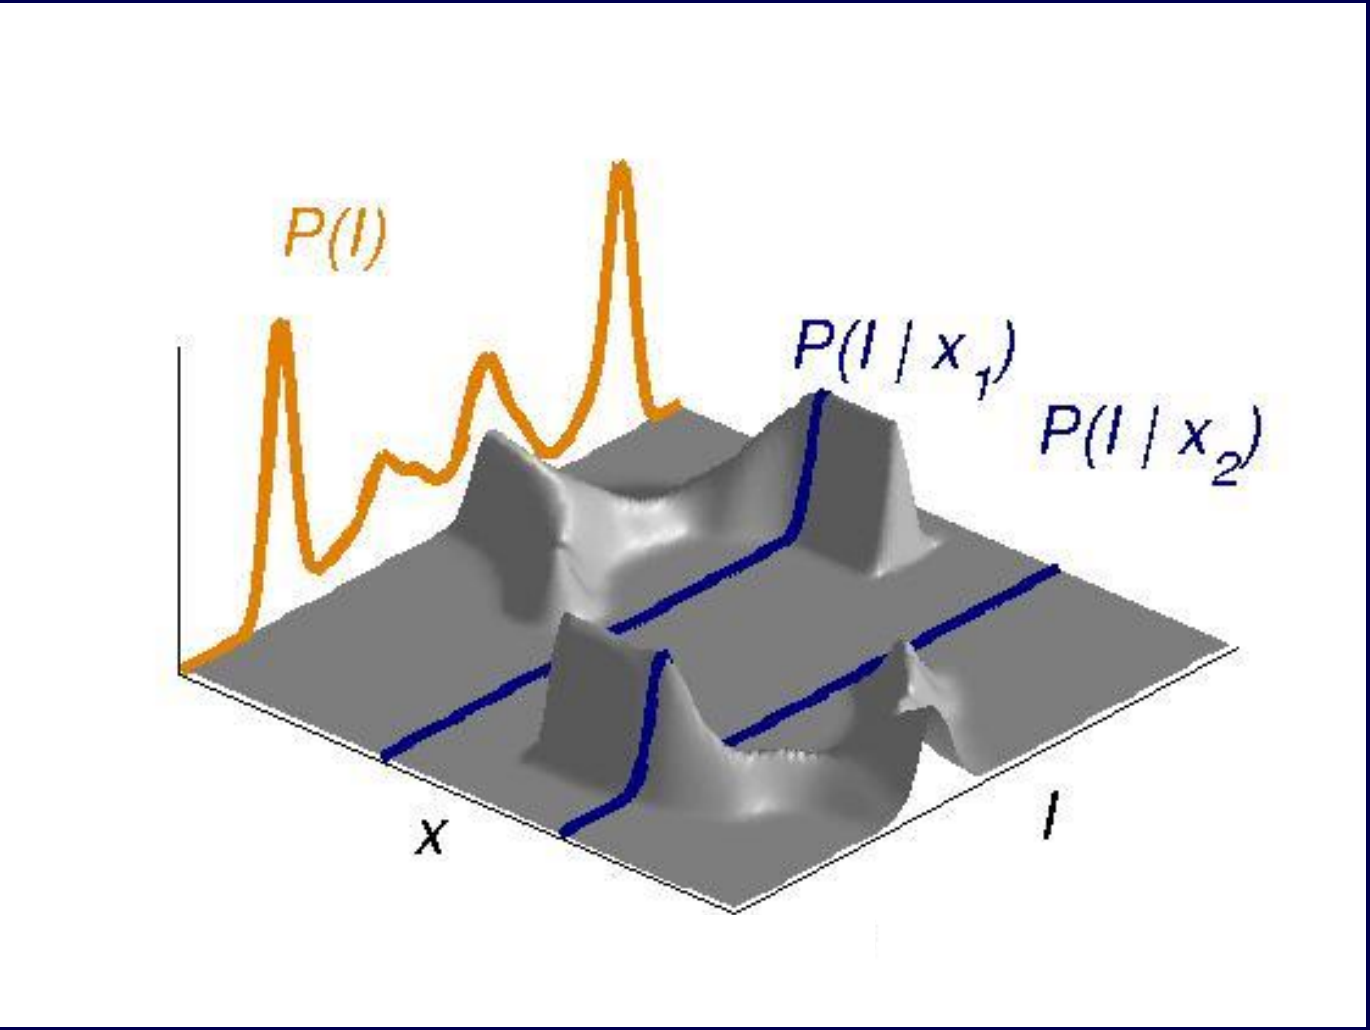

Nieuwenhuis, Cremers, PAMI '12

Interactive Segmentation

Nieuwenhuis, Cremers, PAMI '12

Nieuwenhuis, Cremers, PAMI '12

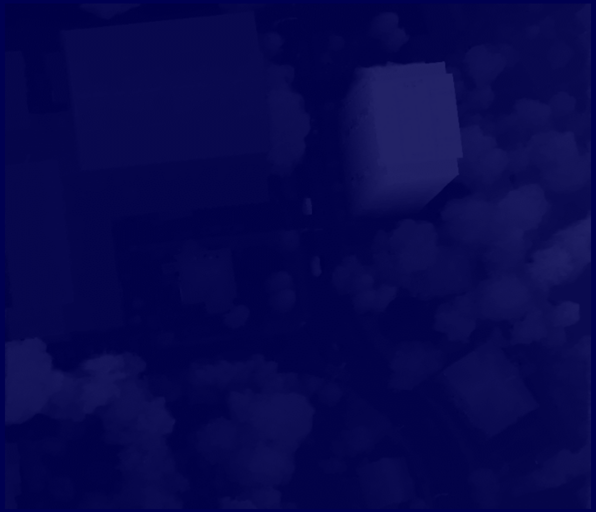

Convex multilabel optimization

Minimal partitions

Semantic segmentation

Mumford-Shah

Semantic Image Segmentation

Ladicki et al. ECCV ‘10, Souiai et al. EMMCVPR ‘13

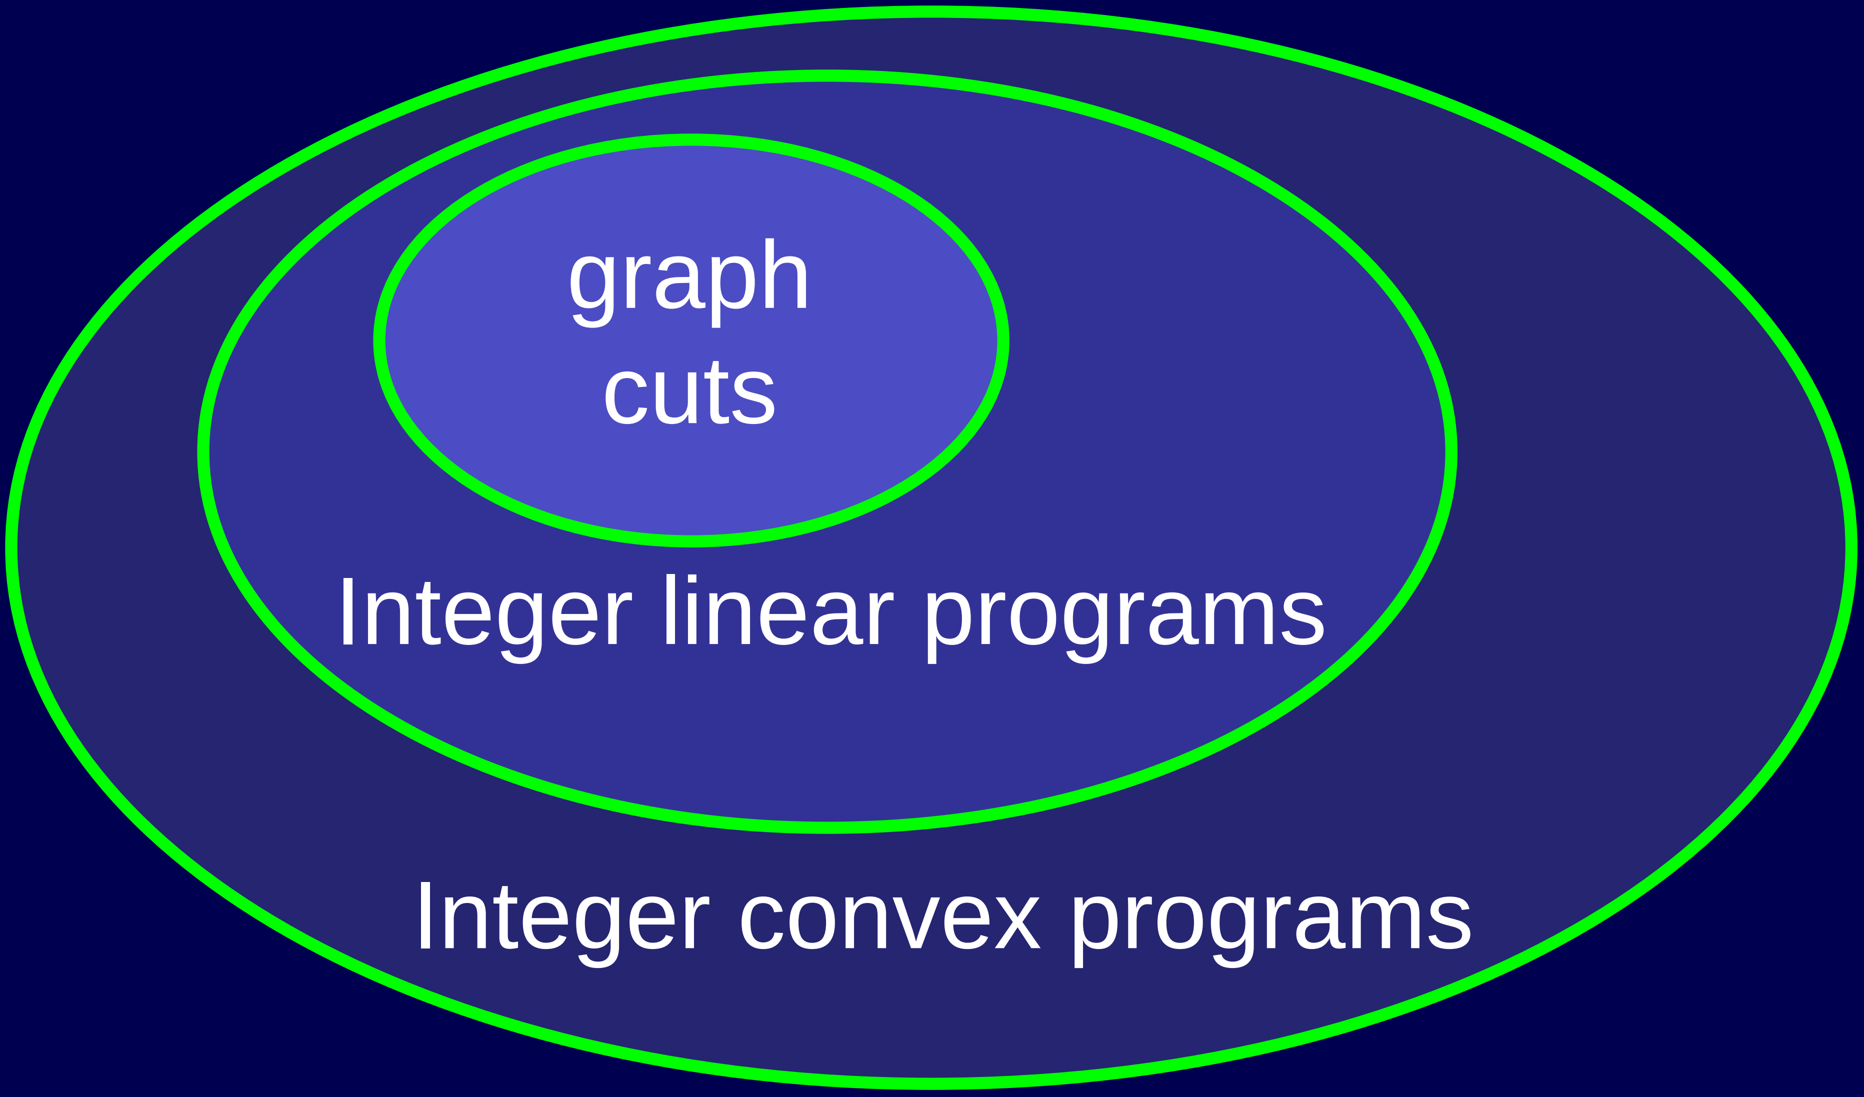

Klodt et al., ECCV '08, Nieuwenhuis et al. PAMI '13

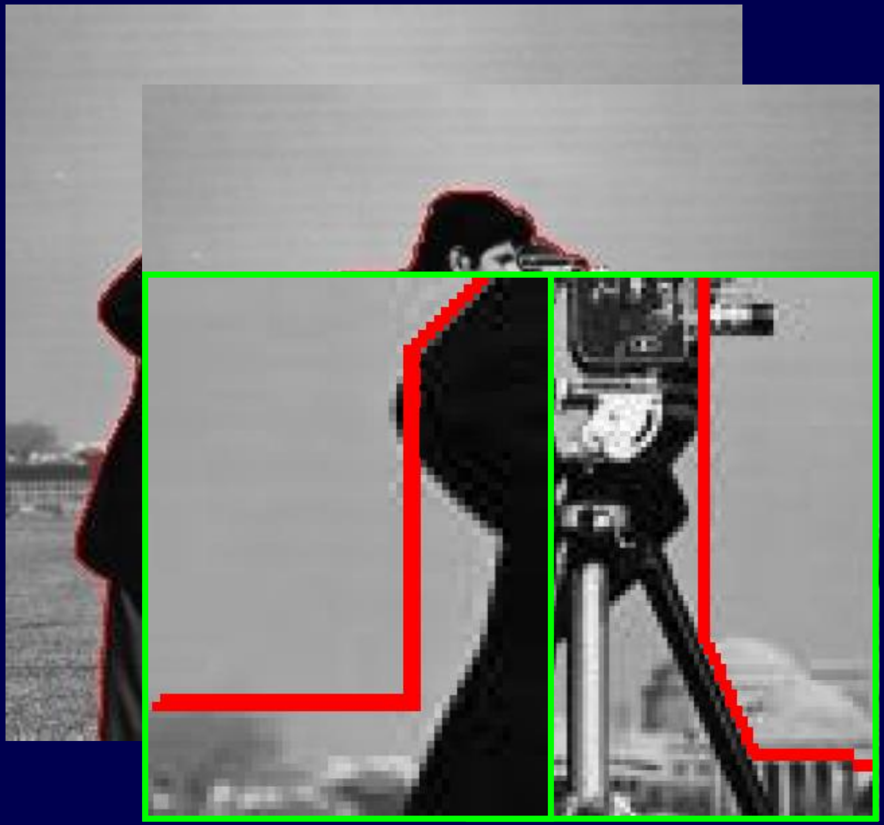

Discrete graph cut optimization
(4-connected grid)

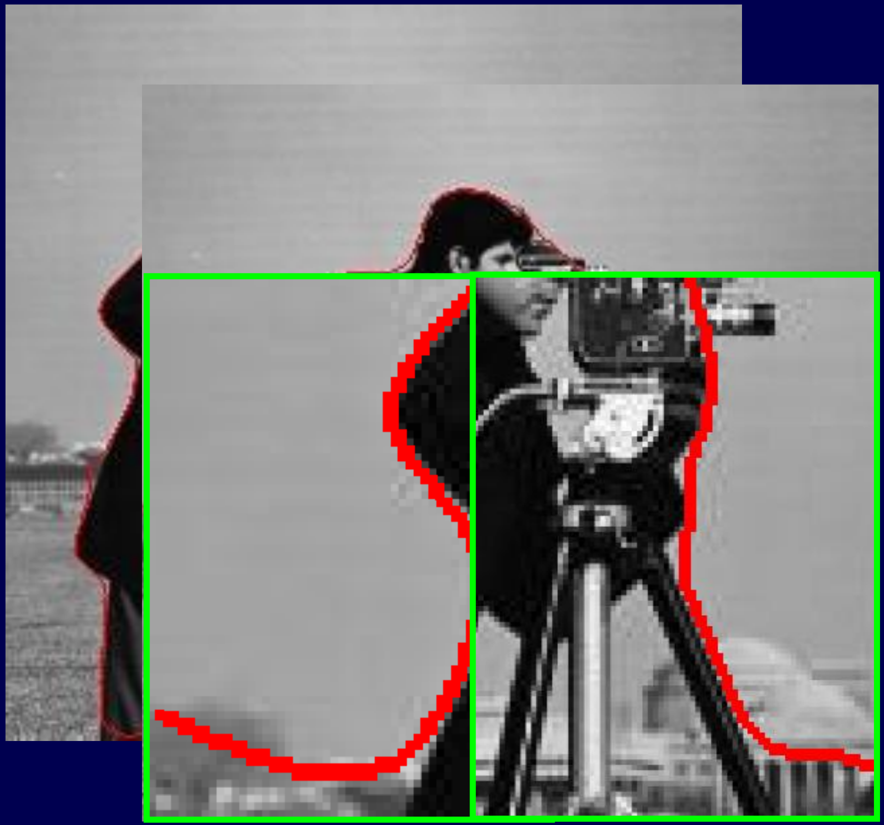

Continuous convex formulation
(4-connected grid)

Improvement: Larger neighborhoods
(*Boykov, Kolmogorov '03*, *Kirsanov, Gortler '04*)

A Minimal Surface: The Catenoid

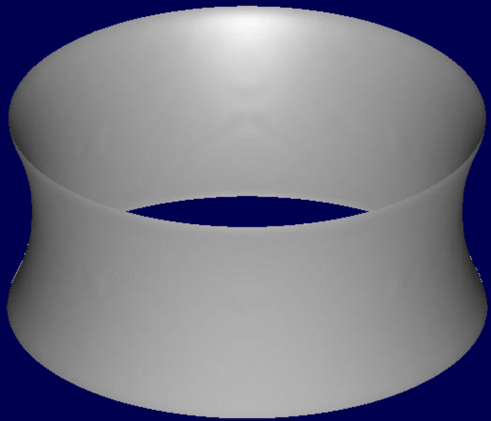

True solution

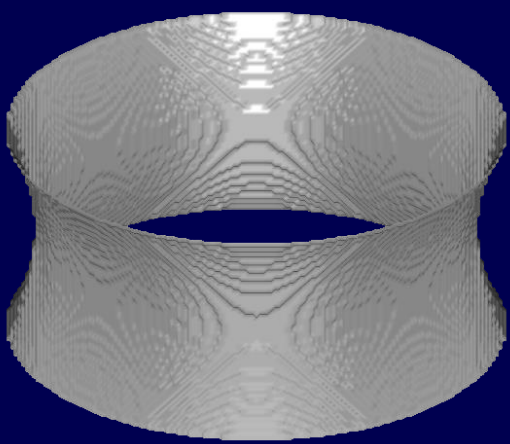

Convex formulation
(6-connected grid)

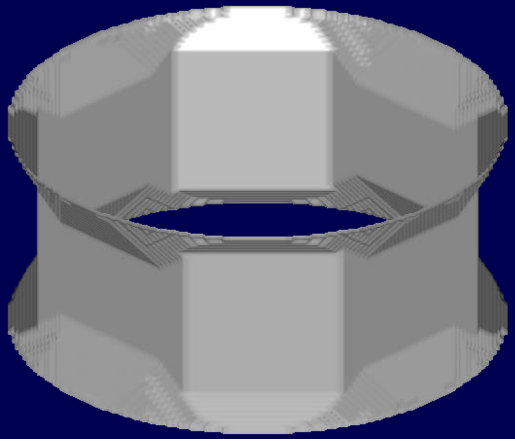

Graph cut
(26-connected grid)

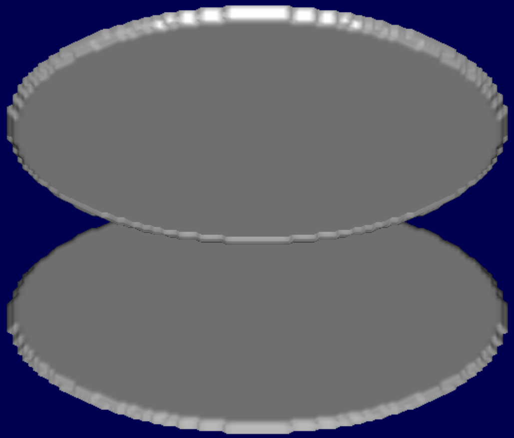

Graph cut
(6-connected grid)

Klodt et al., ECCV '08

Metrication Errors

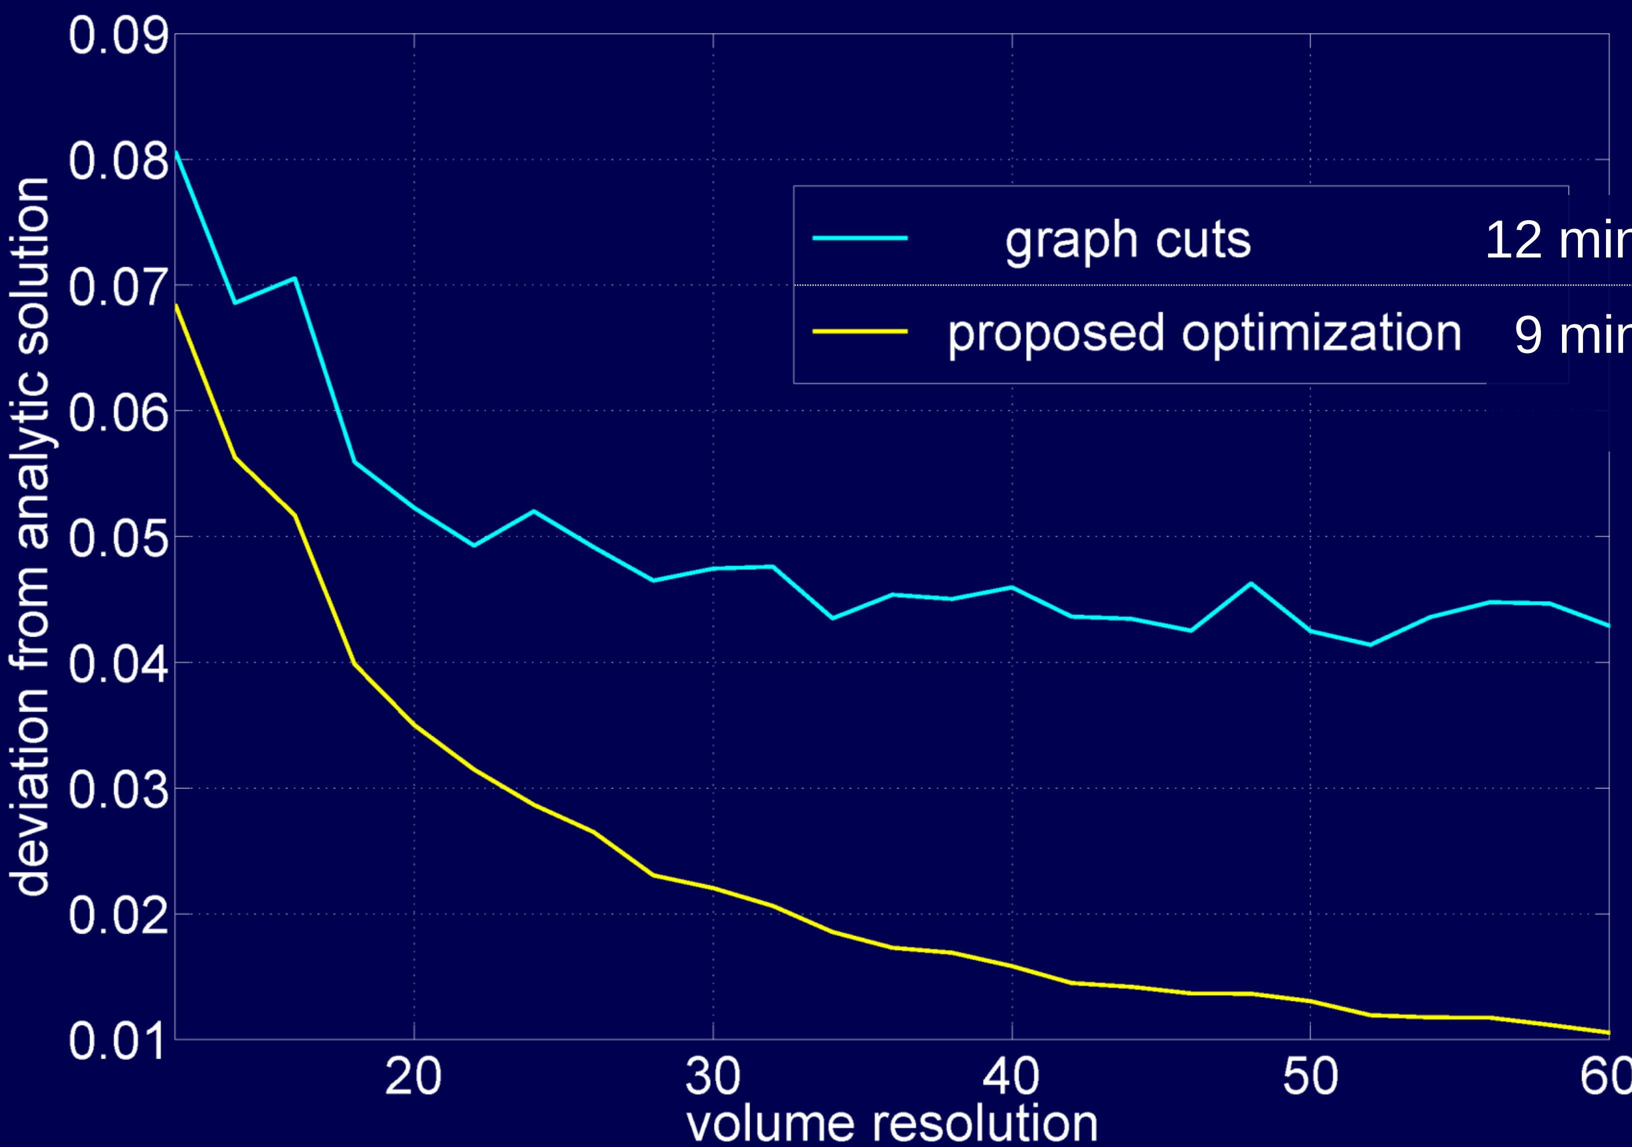

12 min. 32 sec. (CPU)

9 min. 34 sec. (CPU)

8 sec. (GPU)

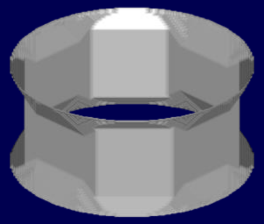

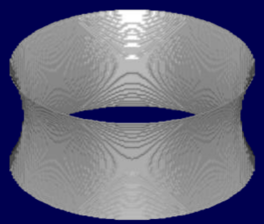

Klodt et al., ECCV '08

Semantic Image Segmentation

Input images

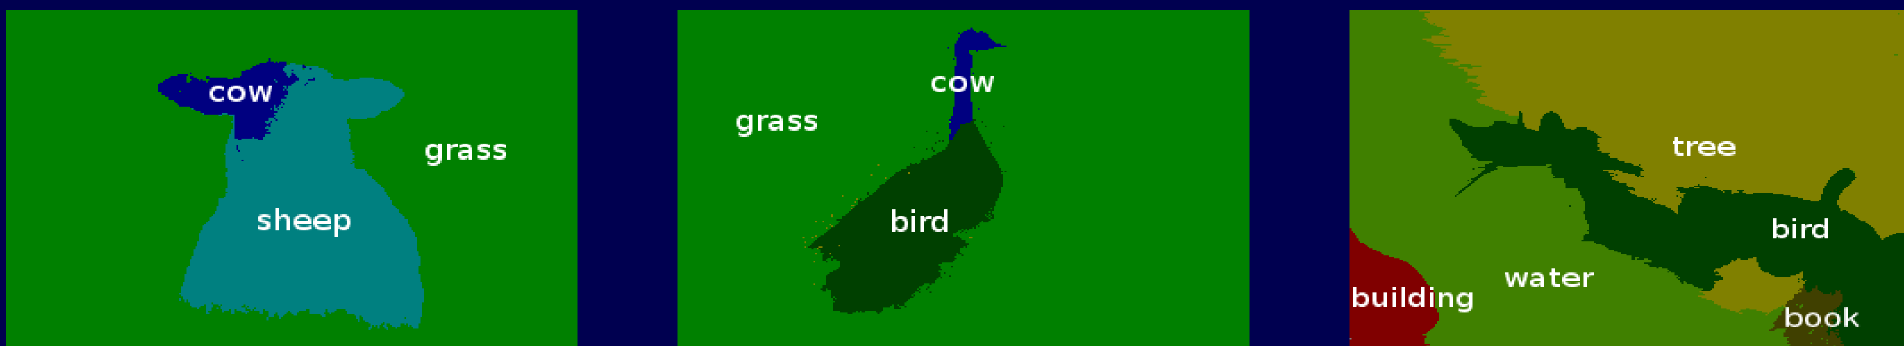

Segmentation with boundary length regularity

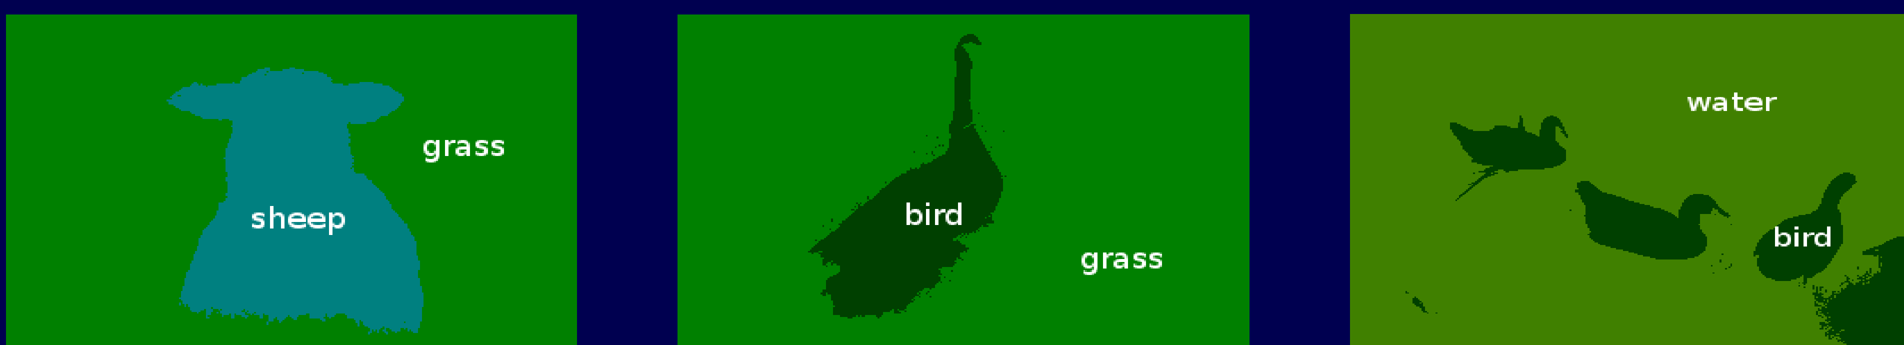

Segmentation with label configuration prior

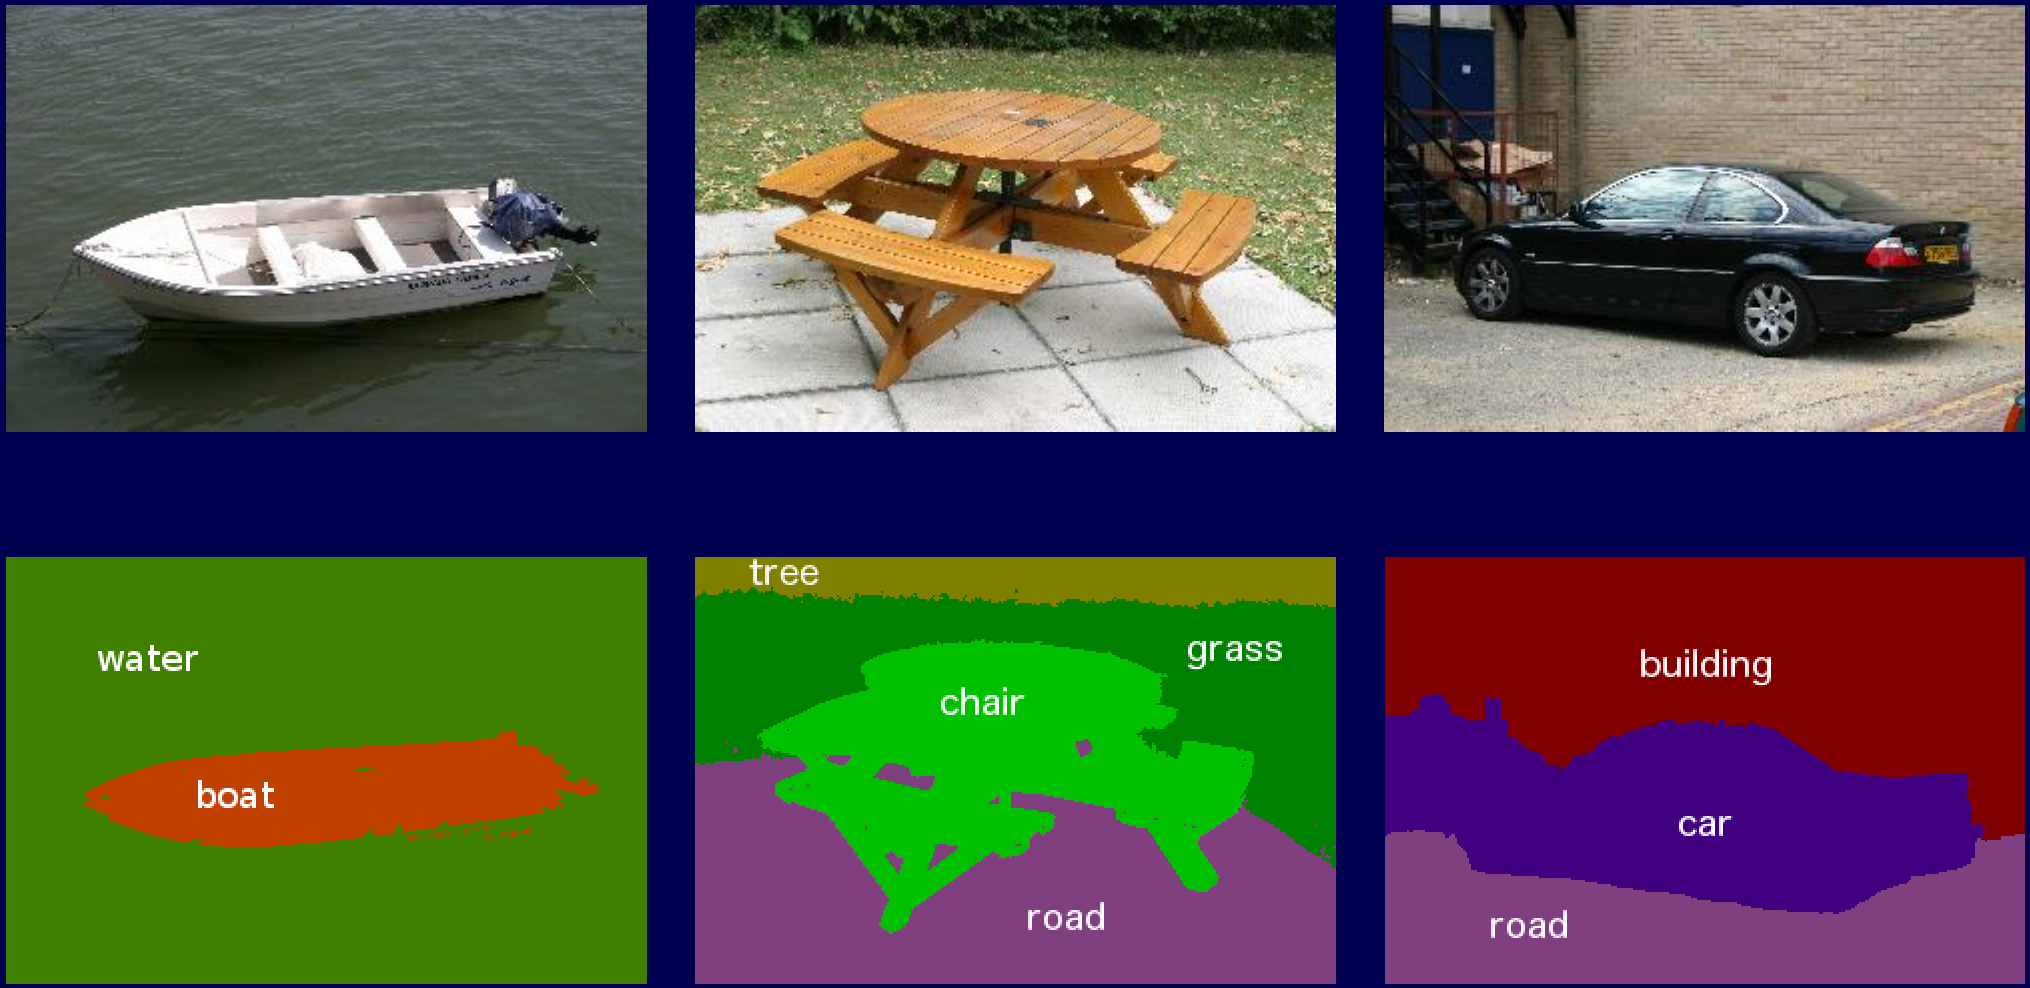

Souiai et al. EMMCVPR '13

General Ordering Constraints

Strekalovskiy, Cremers, ICCV 2011

Related discrete approach: *Liu et al. PAMI '10*

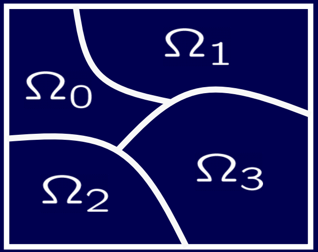

<u>Reminder:</u> With $v_i = 1_{\Omega_i}$, the minimal partition problem is:

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

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

Consider instead the more general convex set:

$$\mathcal{K}_d = \left\{ p \in \mathbb{R}^{n \times m} : \boxed{\langle p_j - p_i, \nu \rangle \leq d(i, j, \nu),} \; \forall i < j, \, \nu \in \mathbb{S}^{m-1} \right\}$$

Penalize transitions depending on label values i, j and orientation ν.

Strekalovskiy, Cremers, ICCV 2011

General Ordering Constraints

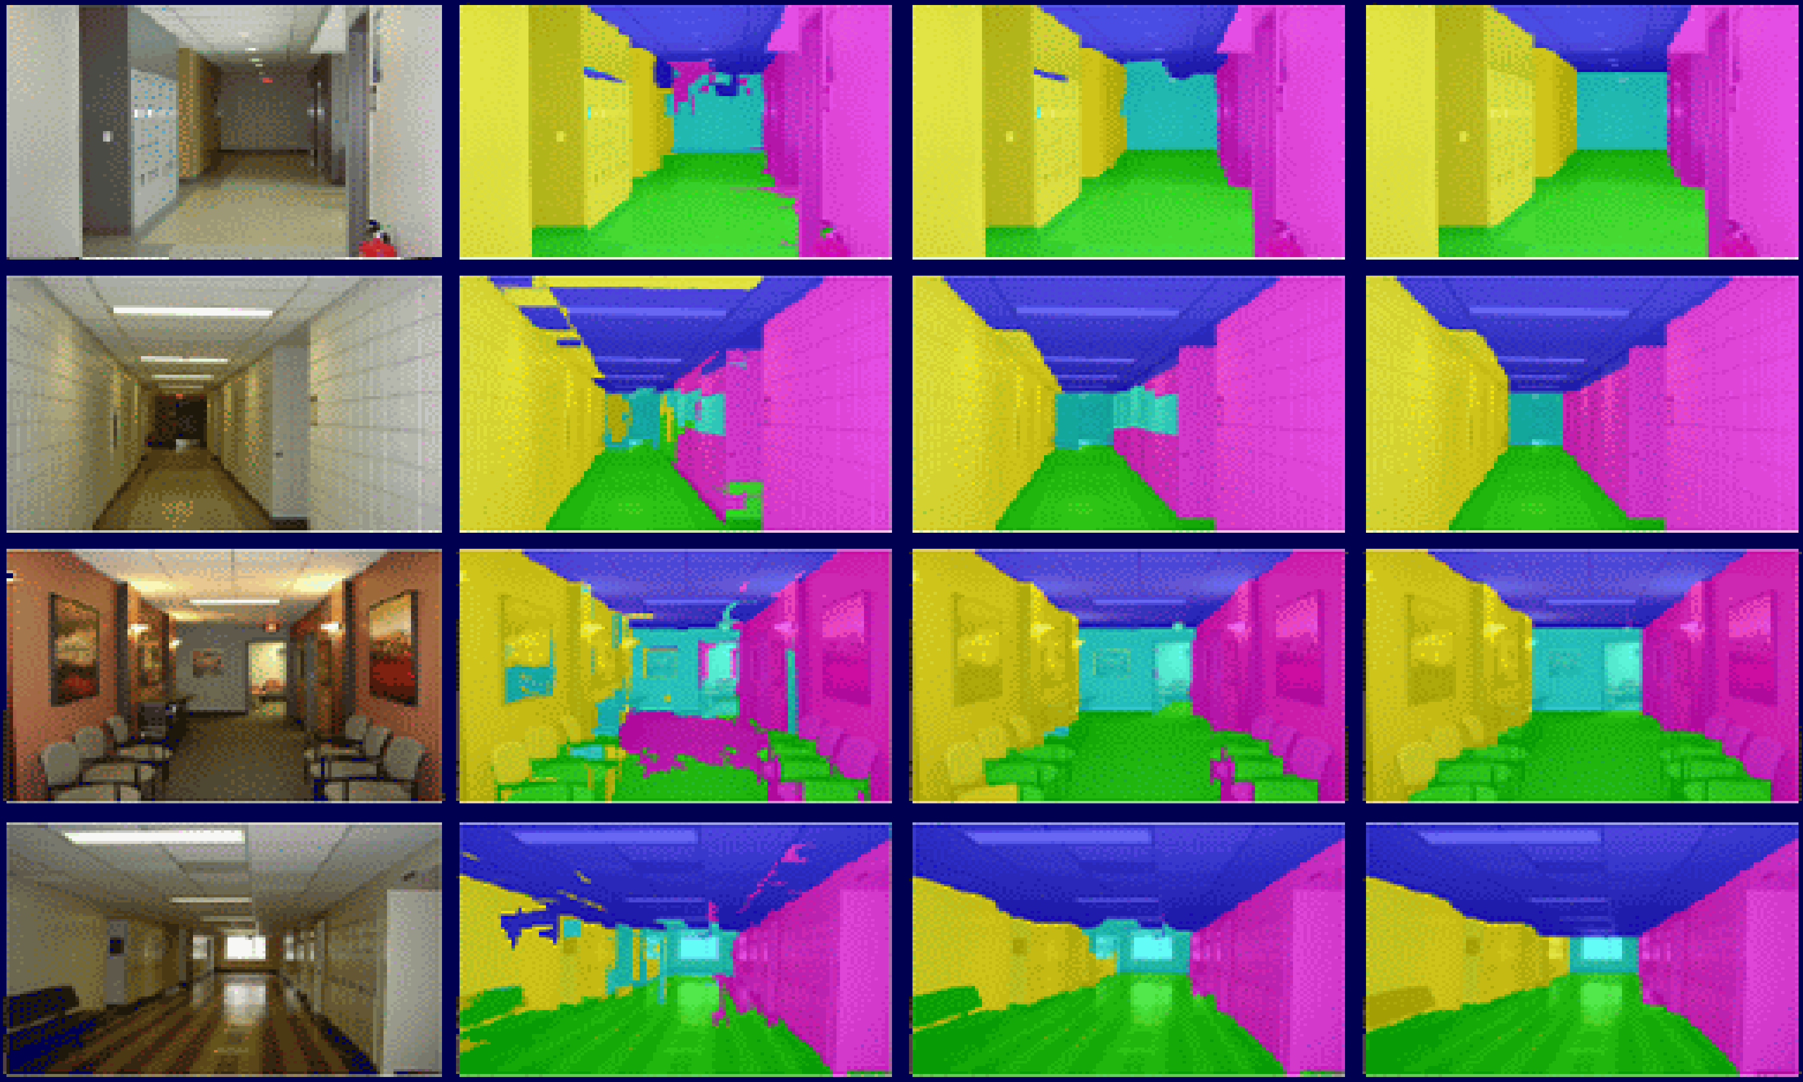

Input | Data term | Min. partition | Ordering

Strekalovskiy, Cremers, ICCV 2011

General Ordering Constraints

Strekalovskiy, Cremers, ICCV 2011

General Ordering Constraints

Input Min. partition Ordering

Strekalovskiy, Cremers, ICCV 2011


Convex multilabel optimization


Minimal partitions


Semantic segmentation


Mumford-Shah

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

Mumford, Shah '89

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

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

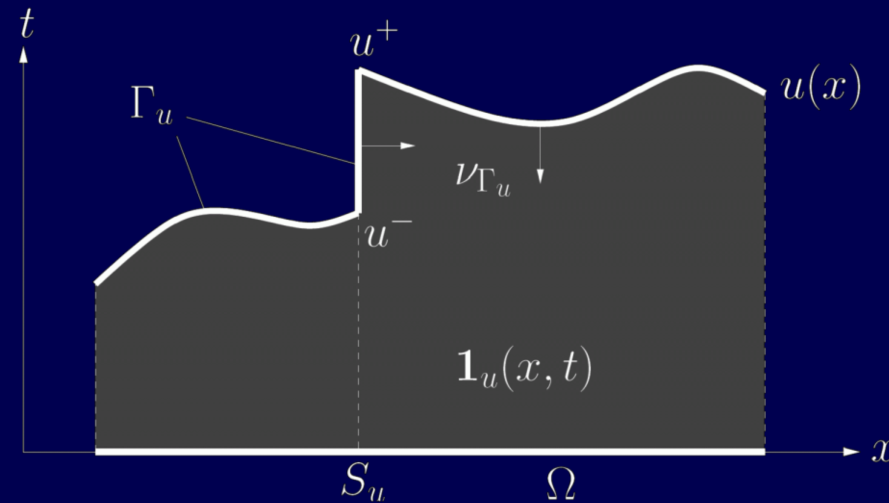

with a convex set

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

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

Alberti, Bouchitte, Dal Maso '04

Input image

piecewise constant

piecewise smooth

Pock, Cremers, Bischof, Chambolle ICCV ’09

Piecewise Smooth: Scalar Case

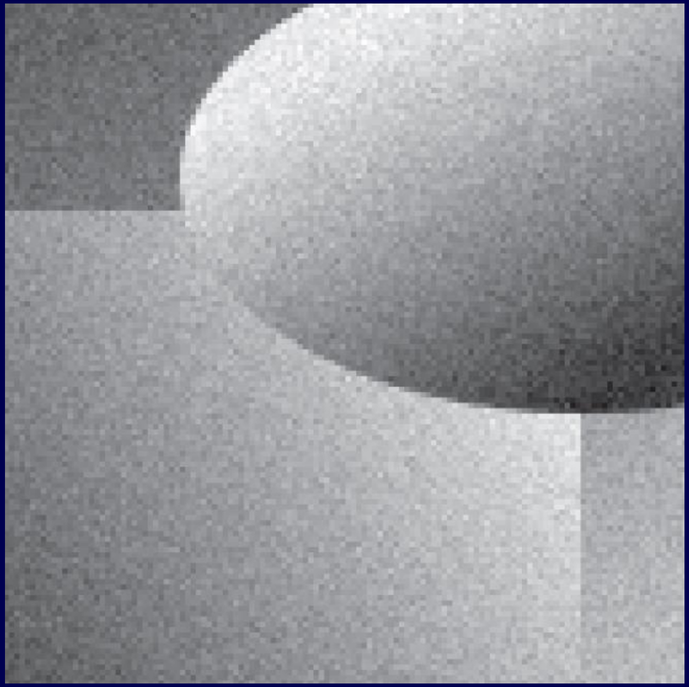
noisy input

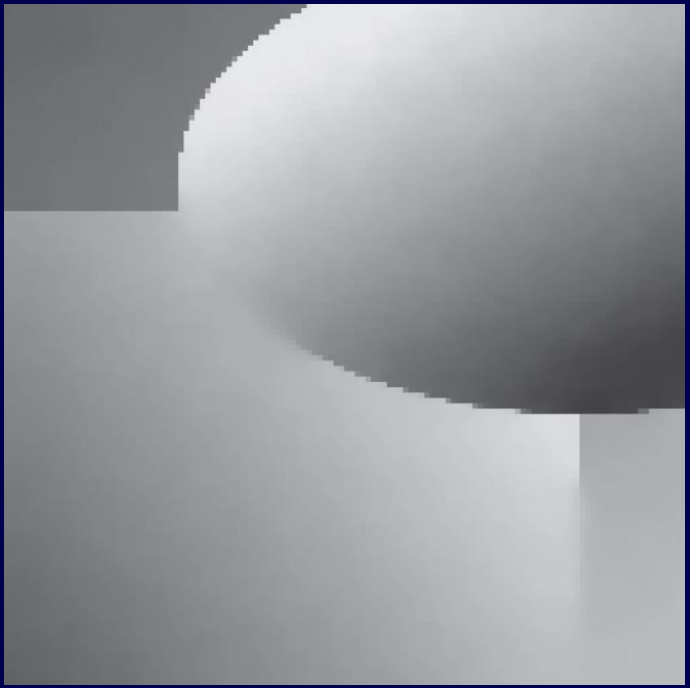
restoration

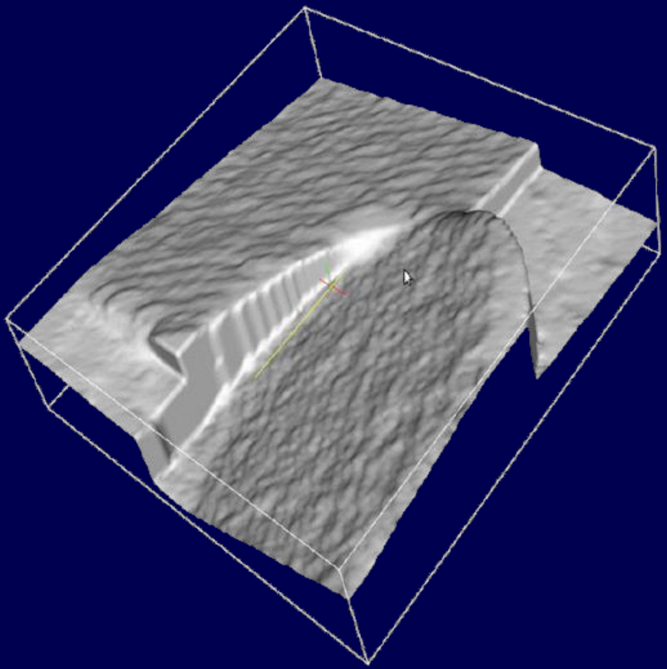
surface plot

Pock, Cremers, Bischof, Chambolle ICCV '09

fixed boundary values

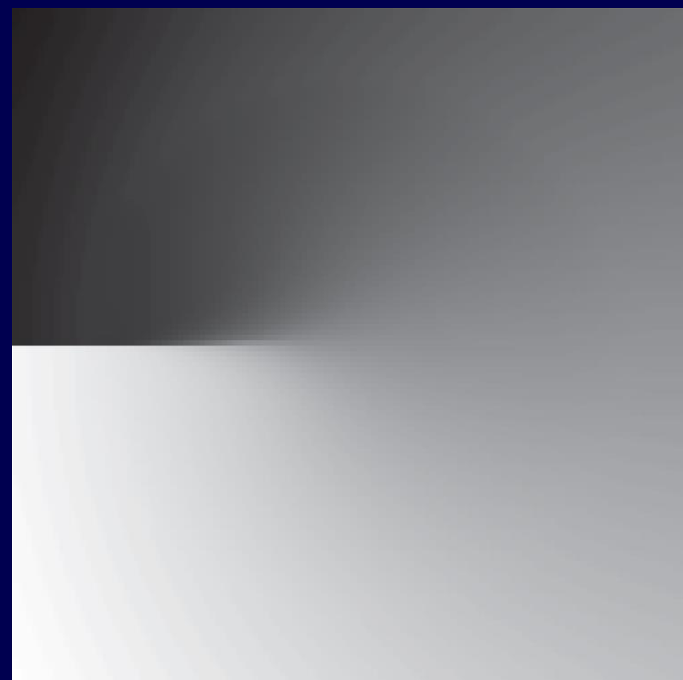

inpainted crack tip

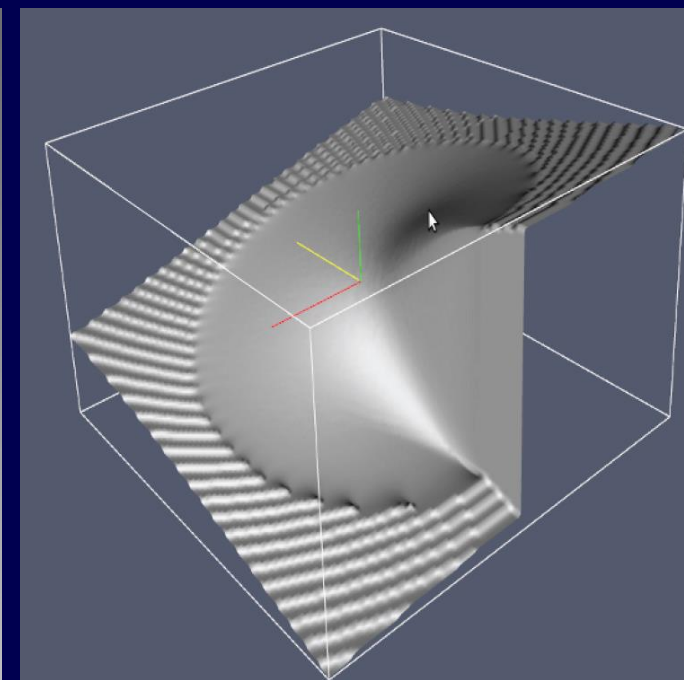

surface plot

Pock, Cremers, Bischof, Chambolle ICCV '09

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

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

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

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

with the convex set:

$$\mathcal{K} = \Big\{ \sigma \,\Big|\, (\sigma_i^x, \sigma_i^t) \in C_c^\infty(\Omega \times \mathbb{R}; \mathbb{R}^n \times \mathbb{R}),$$

$$\sigma_i^t(x, t_i) \geq \frac{1}{4\lambda} |\sigma_i^x(x, t_i)|^2 - (t_i - f_i(x))^2,$$

$$\boxed{\sum_{j=1}^k \left| \int_{t_j}^{t_j'} \sigma_j^x(x,s)\, ds \right| \leq \nu,} \quad \forall\, 1 \leq i \leq k,\; x \in \Omega,\; t_j < t_j' \Big\}.$$

$\mathcal{O}(n_1^2 \cdots n_k^2)$ constraints!

Strekalovskiy, Chambolle, Cremers, CVPR '12

Proposition: The constraint set

$\mathcal{O}(n_1^2 \cdots n_k^2)$ constraints

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

is equivalent to the constraint set

$\mathcal{O}(n_1^2 + \ldots + n_k^2)$ constraints

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

→ Same complexity as channel-wise processing.

Strekalovskiy, Chambolle, Cremers, CVPR '12

The Vectorial Mumford-Shah Problem

Input image

TV denoised

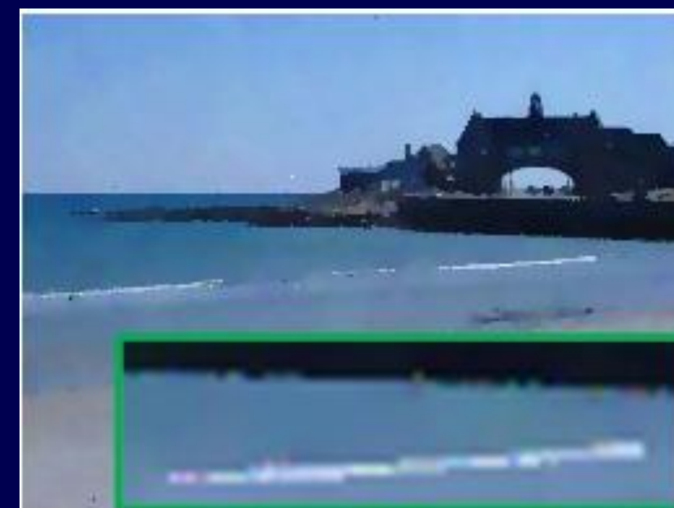

Vectorial Mumford-Shah

Strekalovskiy, Chambolle, Cremers, CVPR ‘12

Input image

Channelwise MS

Vectorial MS

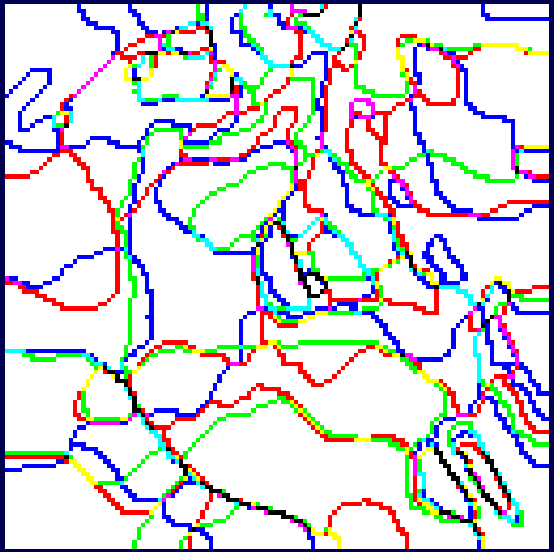

Jump set S_u

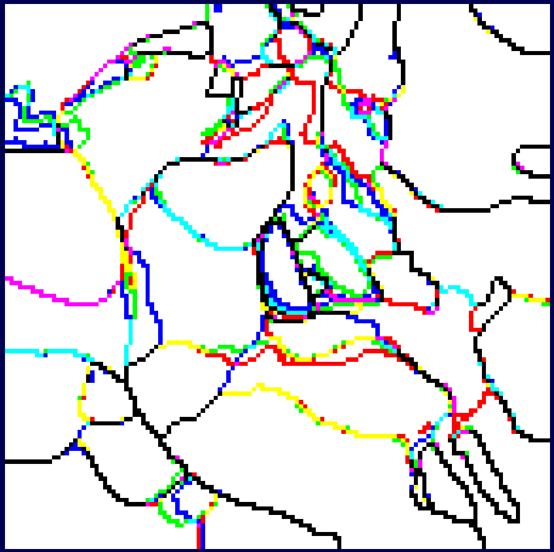

Jump set S_u

Channelwise versus Vectorial


Input image


Channelwise MS


Vectorial MS

Strekalovskiy, Chambolle, Cremers, CVPR ‘12

Input image

$\lambda = \infty, \nu = 0.05$

$\lambda = \infty, \nu = 0.1$

$\lambda = \infty, \nu = 0.2$

Strekalovskiy, Chambolle, Cremers, CVPR '12

Summary

Convex relaxations for real-valued estimation problems can be derived by discretizing the space of permissible values.

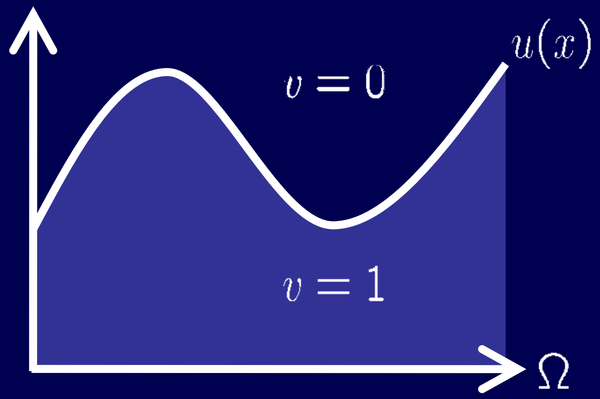

In the scalar-valued case we obtain provably optimal solutions for convex regularizers.

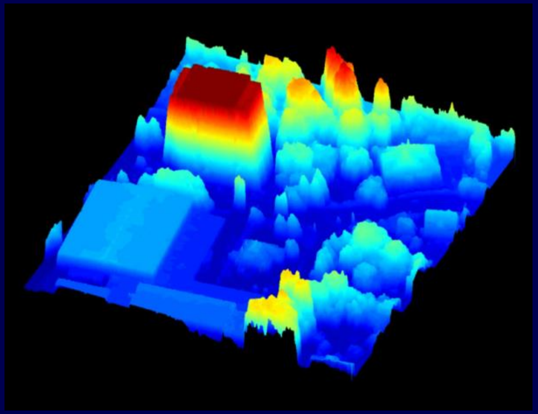

For nonconvex regularizers (Mumford-Shah and minimal partition) we get near-optimal solutions independent of the initialization.