Variational Methods in Computer Vision
ICCV Tutorial, 6.11.2011

Chapter 5

Convex Relaxations for Multi-label Problems

Daniel Cremers and Bastian Goldlücke
Computer Vision Group
Technical University, Munich

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

Overview

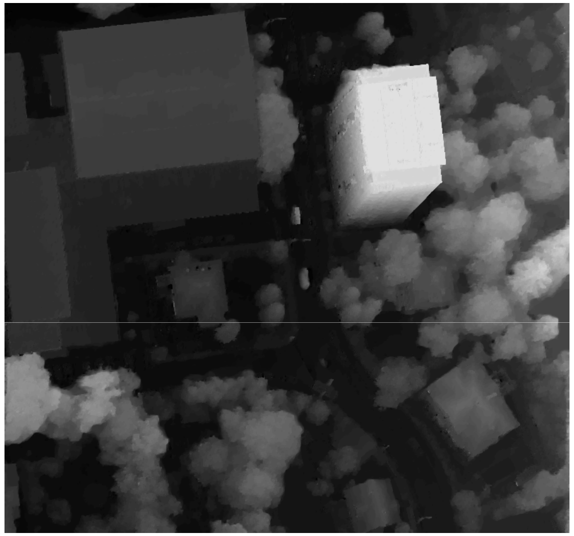

Multi-label optimization

Minimal Partition Problems

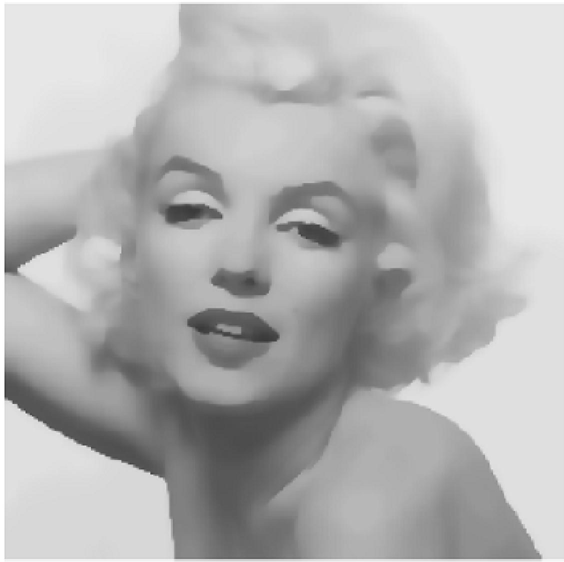

The Mumford-Shah Problem

Label Layout Constraints

Overview

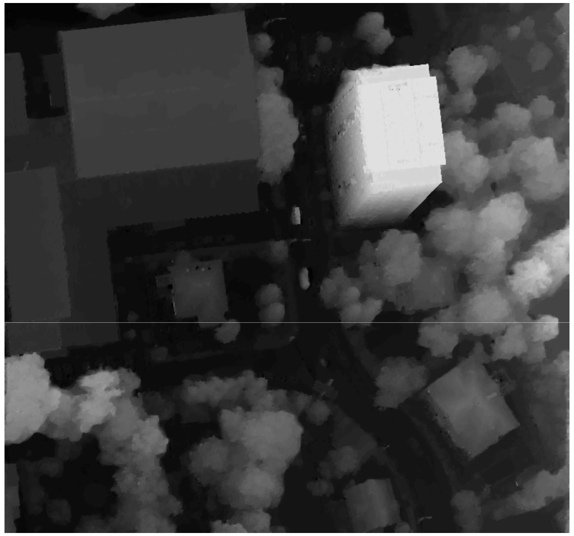

Multi-label optimization

Minimal Partition Problems

The Mumford-Shah Problem

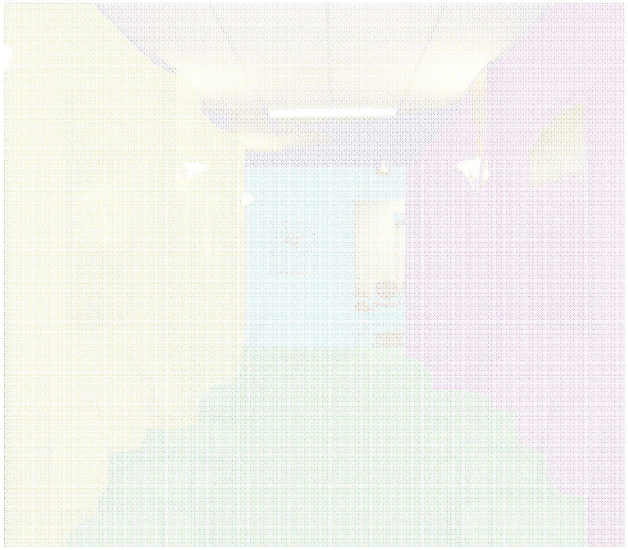

Label Layout Constraints

D. Cremers, B. Goldlücke, T. Pock

From binary to multi-label optimization

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

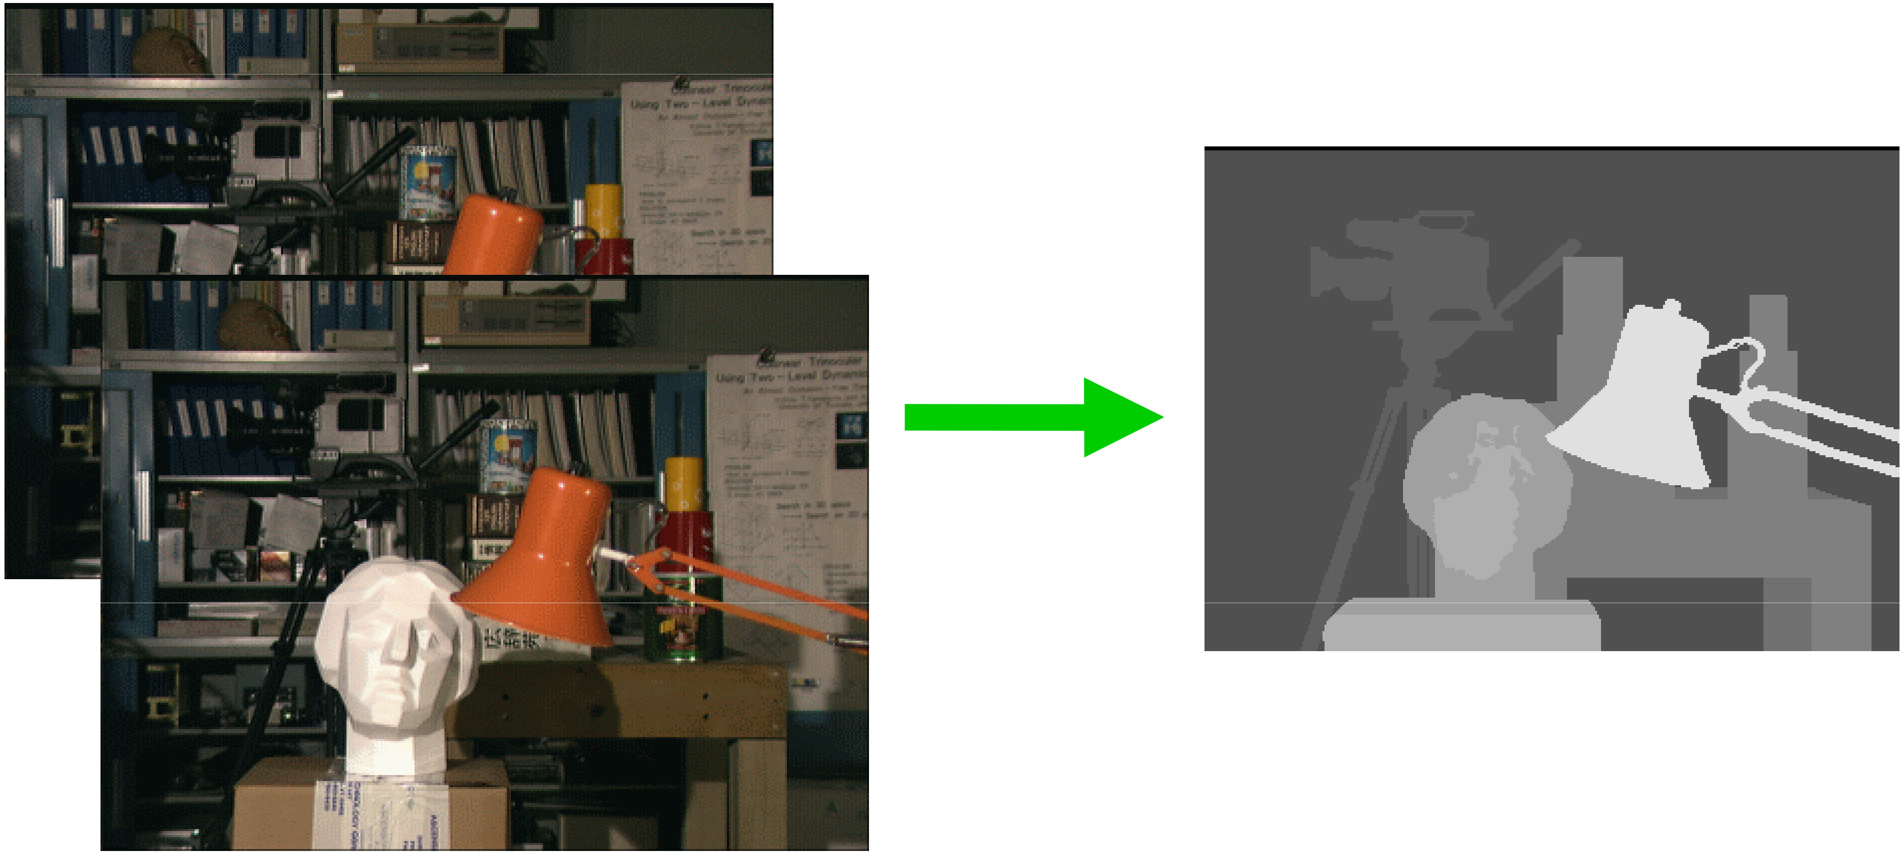

Example: Stereo

Cartesian currents and relaxation

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

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

nonconvex functional

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

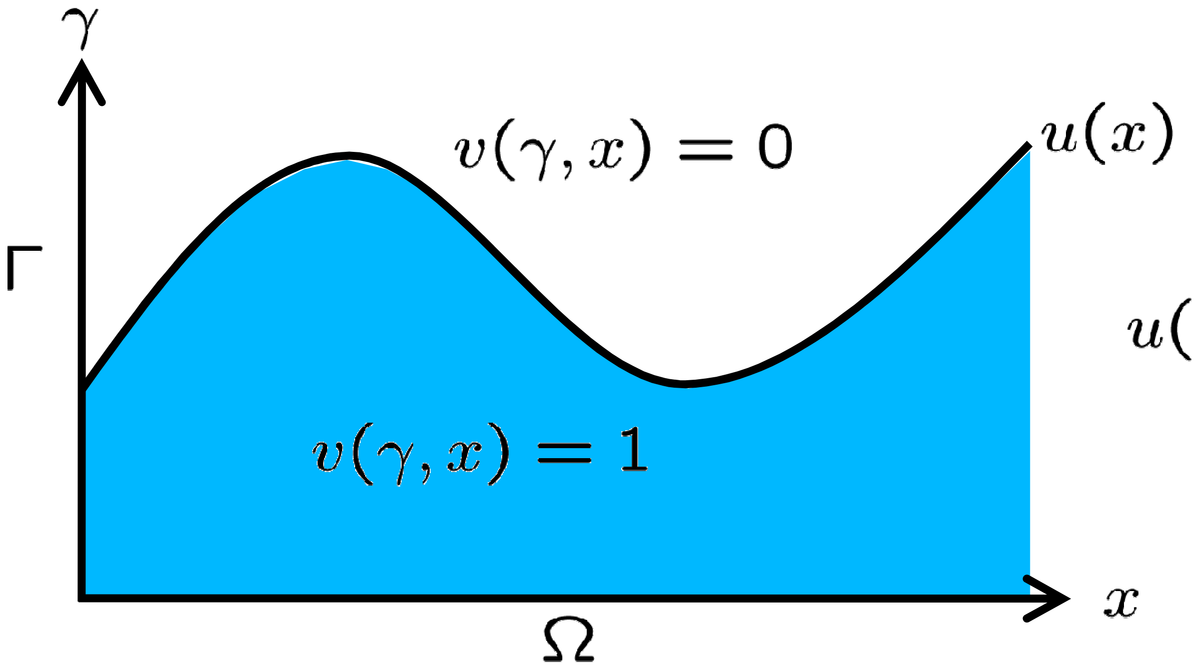

$$u(x) = \gamma_{min} + \int_\Gamma v(x,\gamma)\, d\gamma$$

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

Cartesian currents and relaxation

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

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

nonconvex functional

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

<u>Theorem:</u> Minimizing $(*)$ is equivalent to minimizing

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

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

D. Cremers, B. Goldlücke, T. Pock

Evolution to global optimum

Depth reconstruction from aerial images

One of two input images

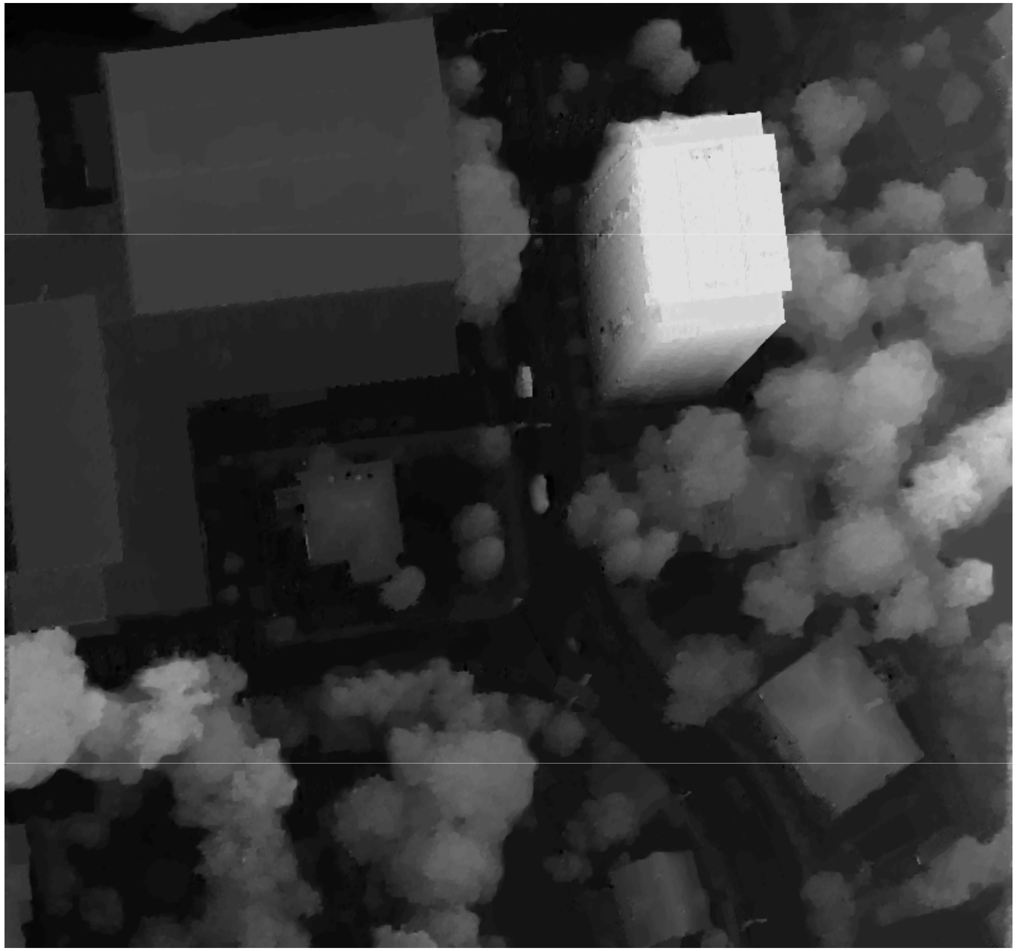

Depth reconstruction

Image data courtesy of Microsoft Graz

Extension to arbitrary convex regularizers

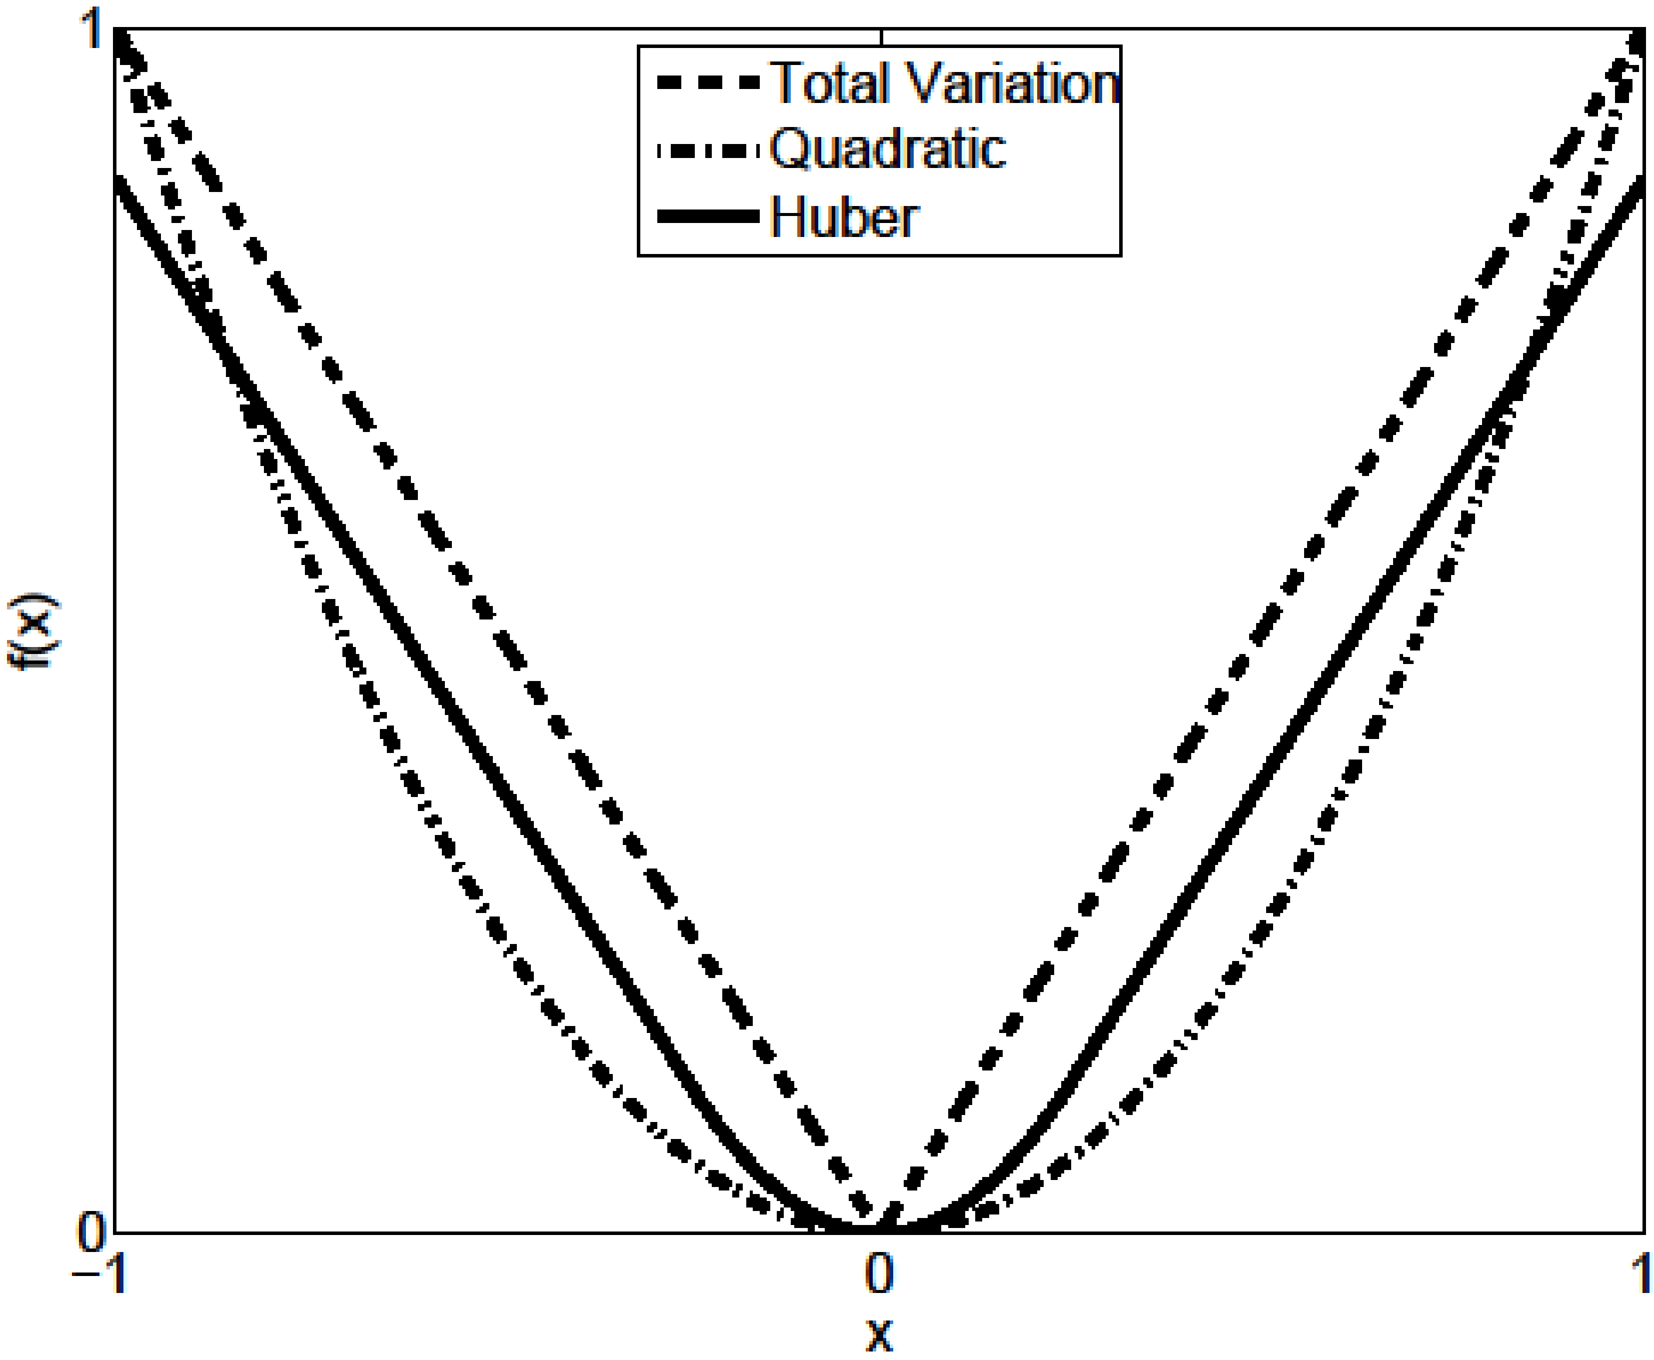

Global optima for functionals with convex regularizers

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 the flow 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 functionals with convex regularizers

Therefore 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

Global optima for functionals with convex regularizers

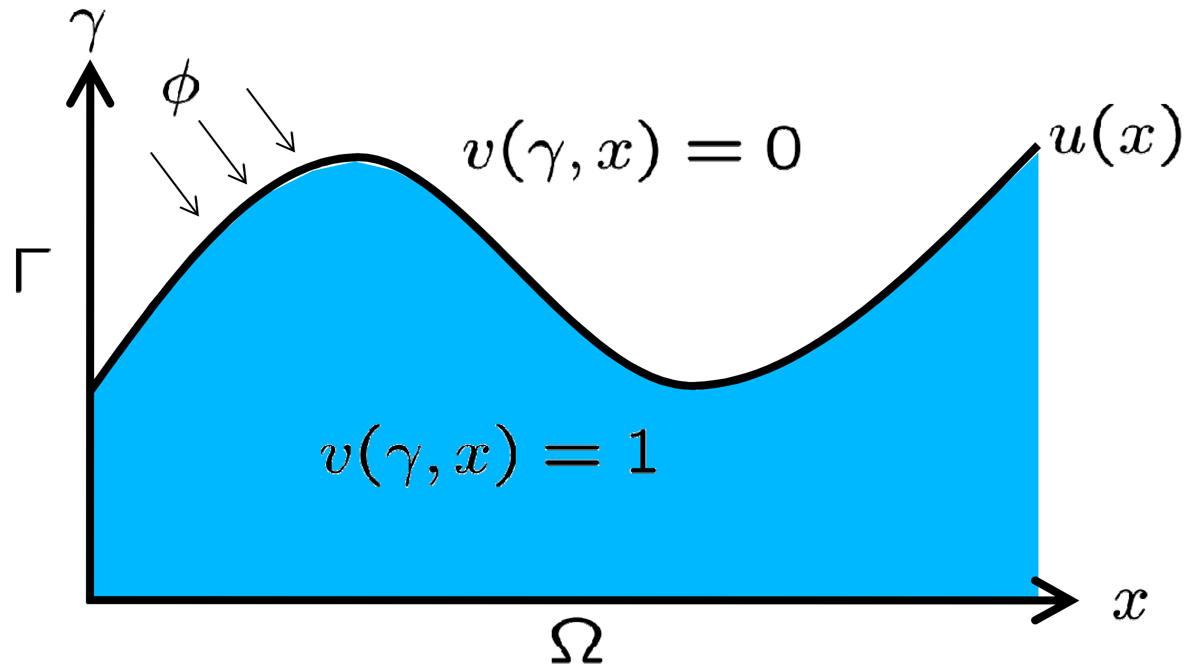

<u>Key observations:</u>

- Minimizing $E(u)$ gives rise to a maximization of flux.
- The various terms in $E(u)$ give rise to respective flow constraints.
- The computational challenges are the minimization of saddle point problems with convex constraints.

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

Outlier removal in range images: a comparison

$$f(x, u, \nabla u) = \min \left\{ \left(I(x) - u(x) \right)^2, \theta \right\} + R(|\nabla u|)$$

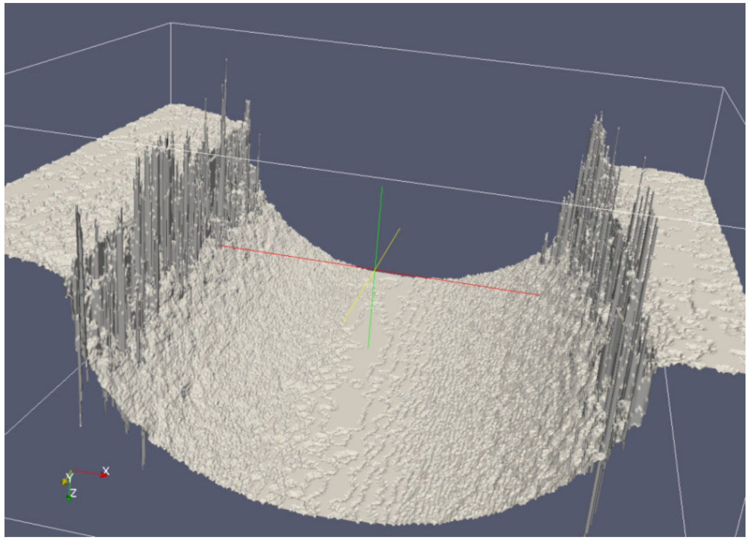

Original data $I(x)$

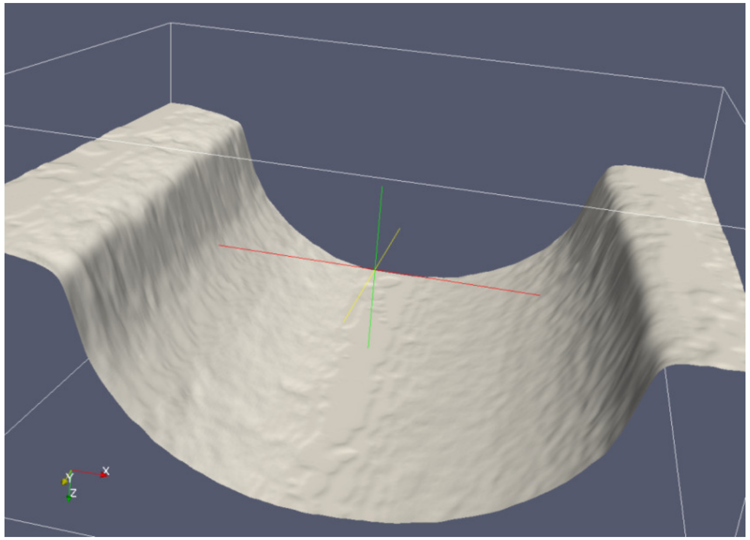

with quadratic regularizer

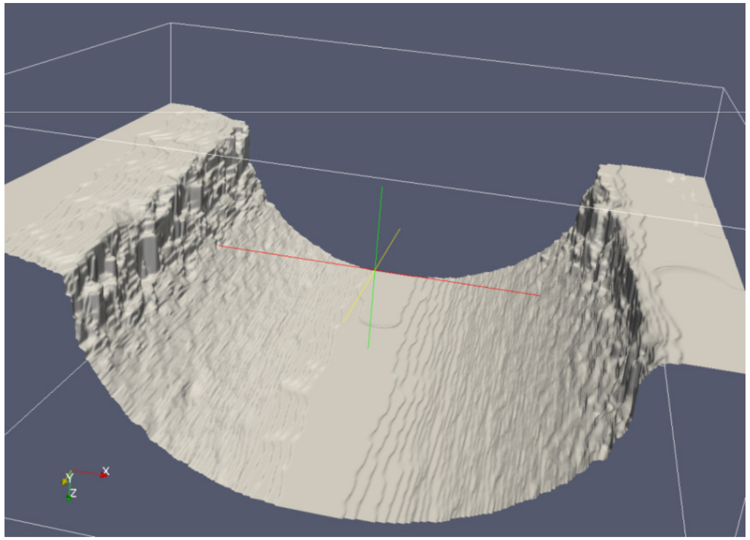

with TV regularizer

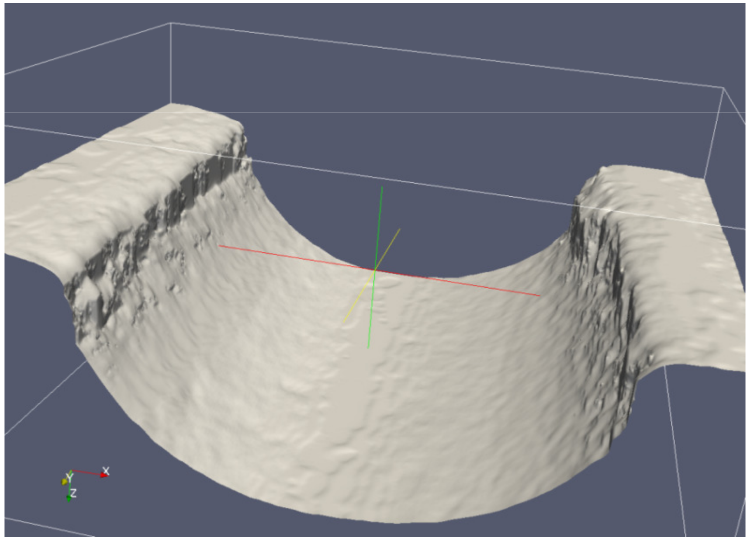

with Huber regularizer

Challenges: nonconvex regularizers & vector-valued

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

What if f is non-convex in ∇u ?

Robustness through truncated regularizers, image segmentation,…

$\longrightarrow$ This chapter!

What if u is vector-valued?

Optical flow, color image processing,…

$\longrightarrow$ Last chapter!

Overview

Multi-label optimization

Minimal Partition Problems

The Mumford-Shah Problem

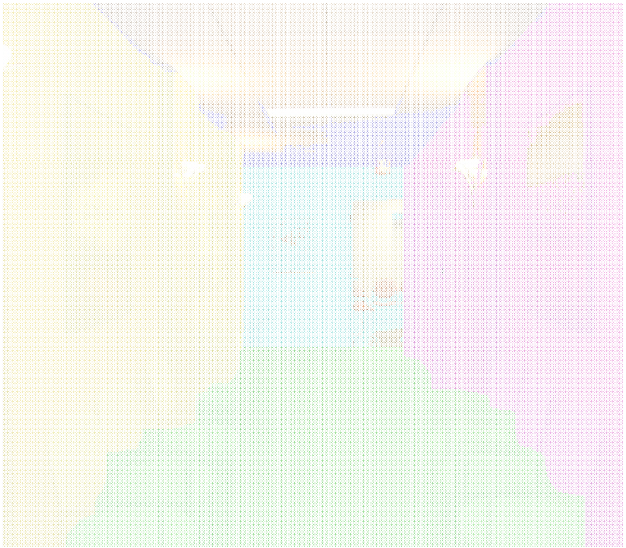

Label Layout Constraints

A convex formulation via paired calibrations

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

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

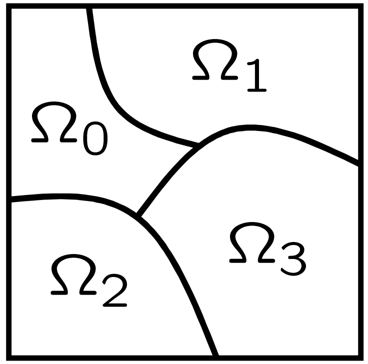

Potts '52, Mumford-Shah '89, Vese, Chan '02

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

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

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

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

Test case: the triple junction

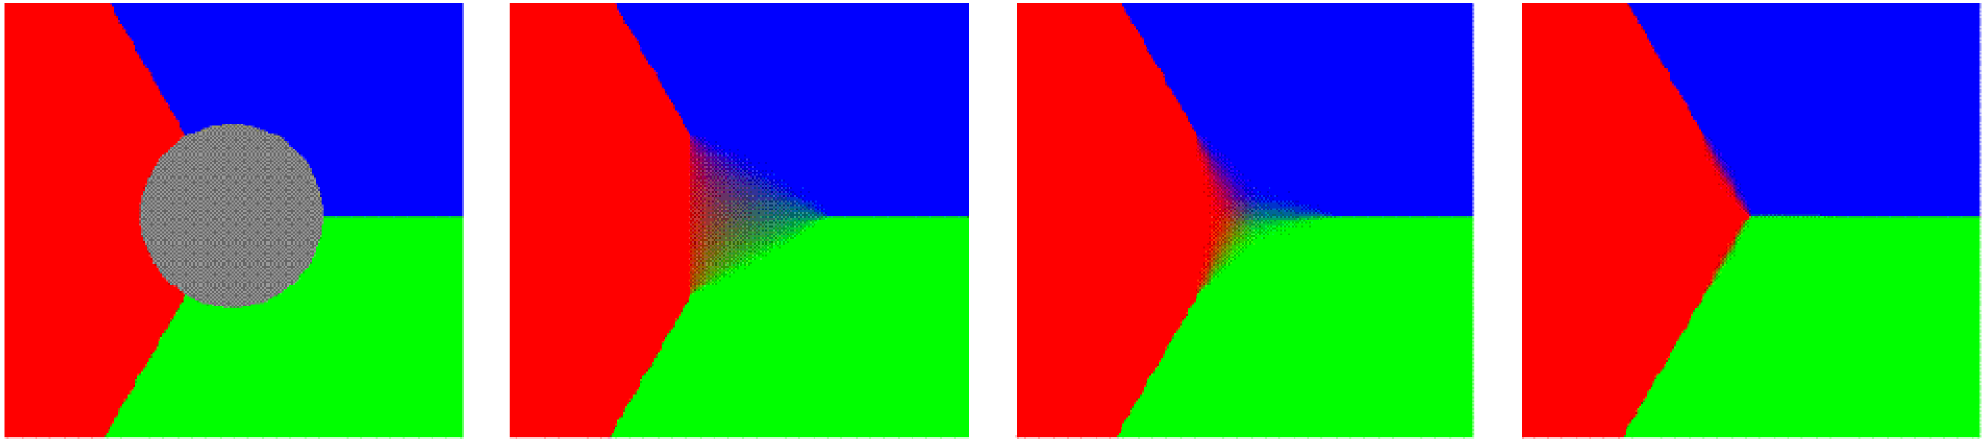

Input image | *Lellmann et al. '08* | *Zach et al. '08* | *our approach*

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

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

A four-region inpainting problem

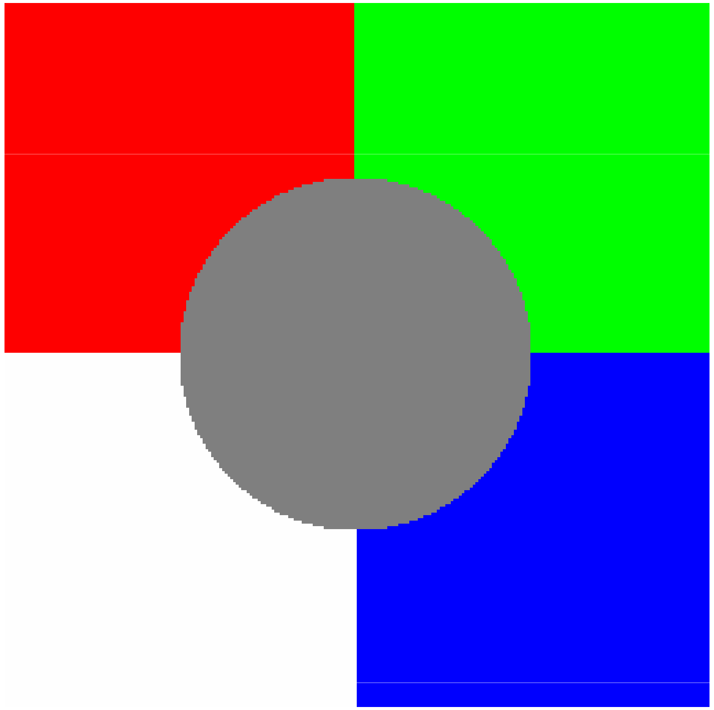

Input image

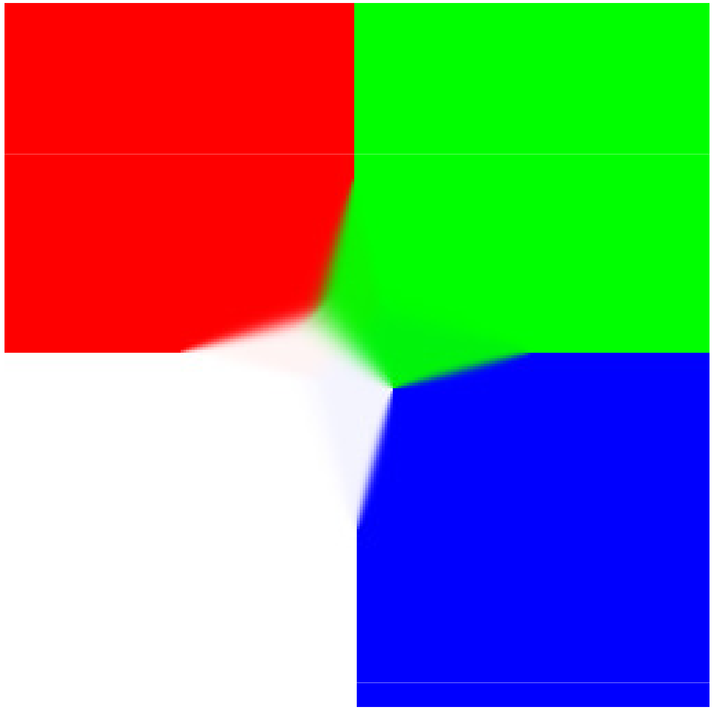

Inpainted

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

Simulating soap bubble configurations

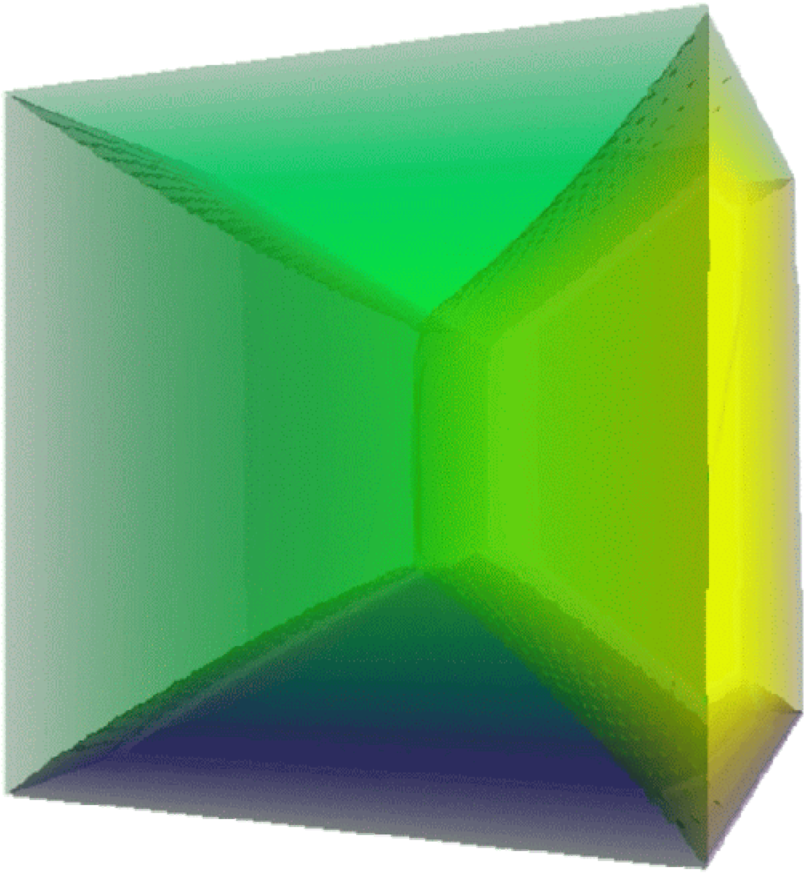

3D min partition inpainting

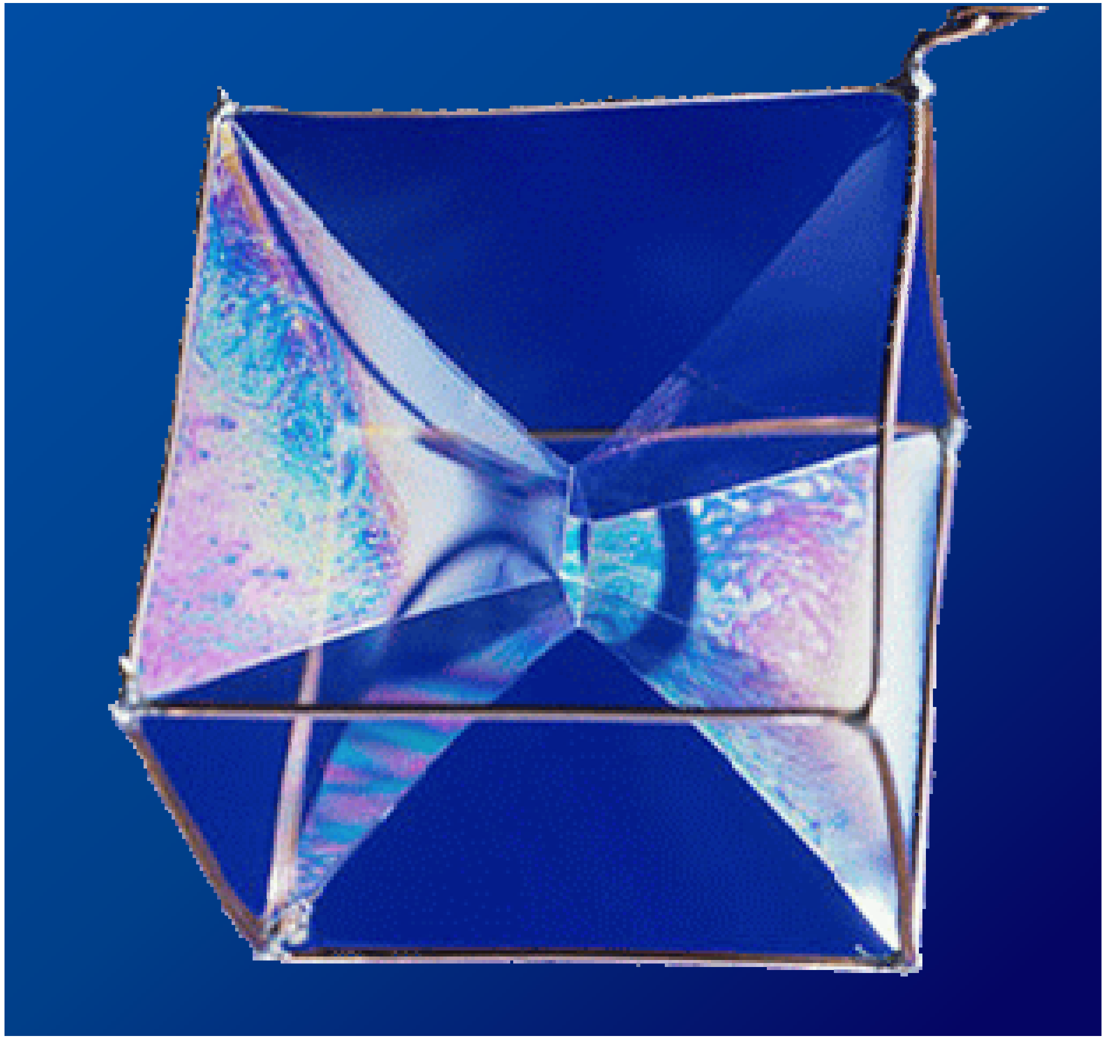

Soap film photo

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

Segmentation of white matter & gray matter

Input image

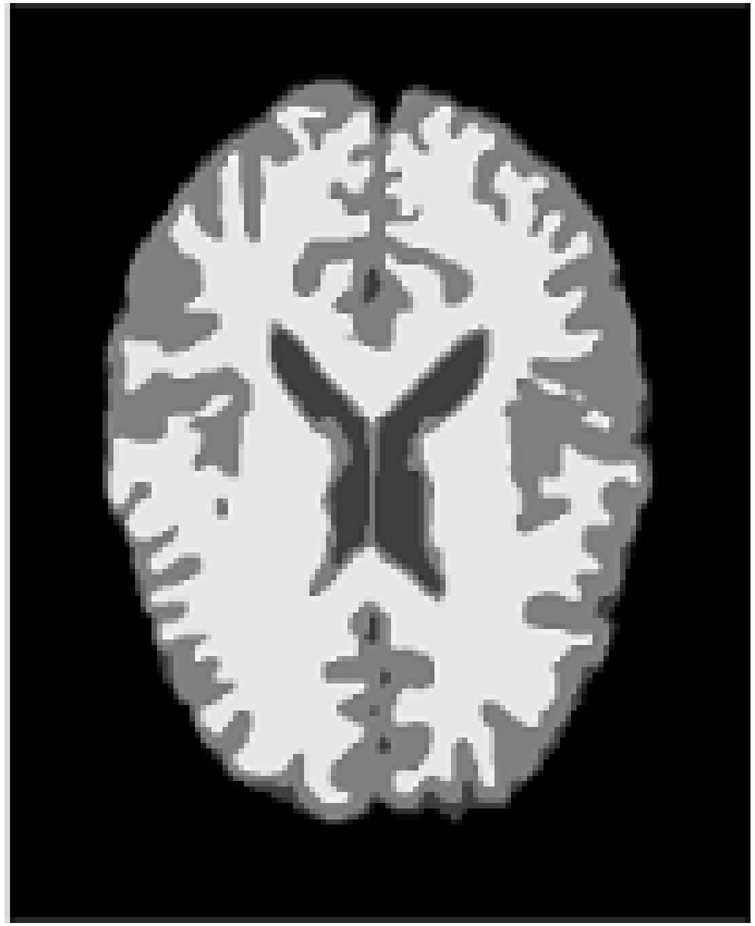

4-region segmentation

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

Image segmentation with multiple regions

Input image | 5 label segmentation | 10 label segmentation

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

Image segmentation with multiple regions

Input color image

10 label segmentation

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

Interactive segmentation with adaptive color models

Nieuwenhuis, Toeppe, Cremers EMMCVPR '11

Interactive segmentation with adaptive color models

Nieuwenhuis, Toeppe, Cremers EMMCVPR '11

Overview

Multi-label optimization

Minimal Partition Problems

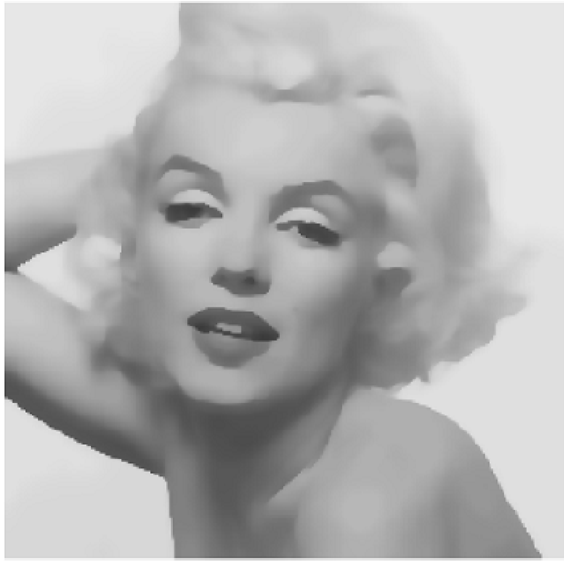

The Mumford-Shah Problem

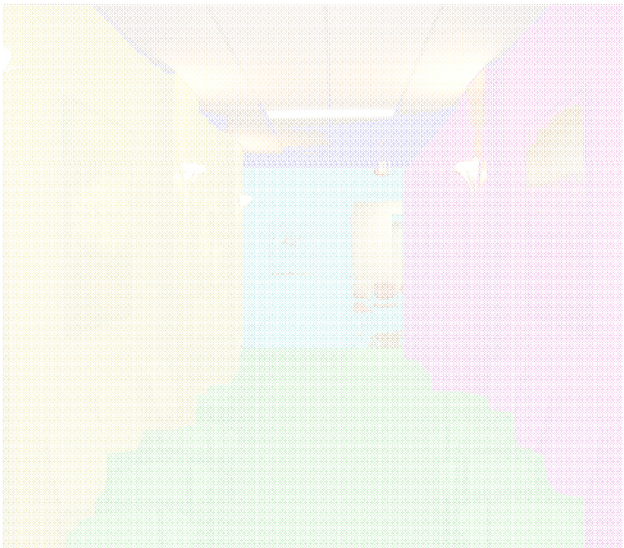

Label Layout Constraints

A convex formulation via Fenchel duality

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

Mumford, Shah '89

For $u \in SBV(\Omega)$, $(*)$ can be written as *(Alberti, Bouchitte, Dal Maso '04)*

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

with a convex set

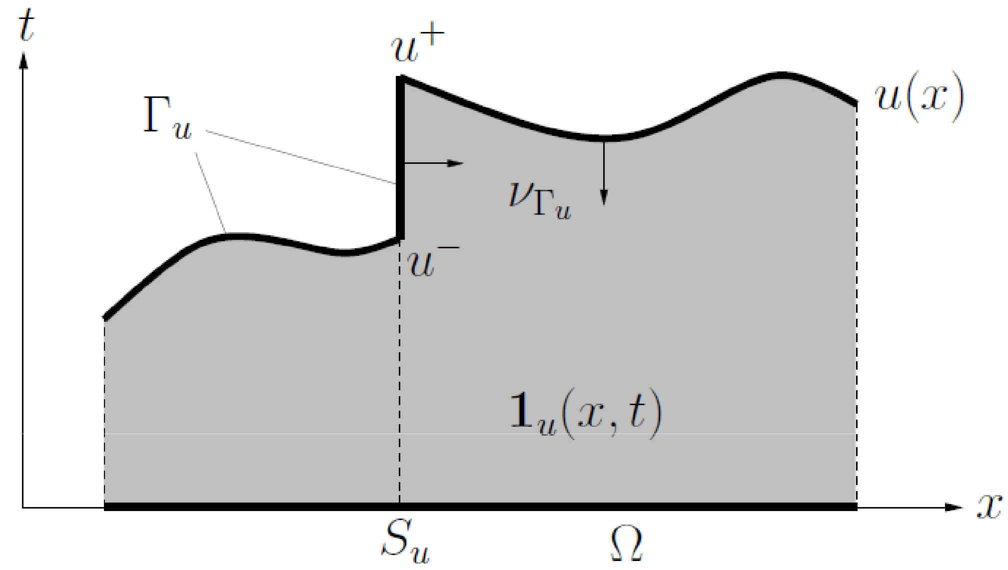

$$K = \left\{ \varphi \in C_0(\Omega \times \mathbb{R}; \mathbb{R}^2) : \right.$$

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

Pock, Cremers, Bischof, Chambolle ICCV '09

Piecewise constant vs. piecewise smooth


Input image


piecewise constant


piecewise smooth

Pock, Cremers, Bischof, Chambolle ICCV '09

Piecewise smooth approximation

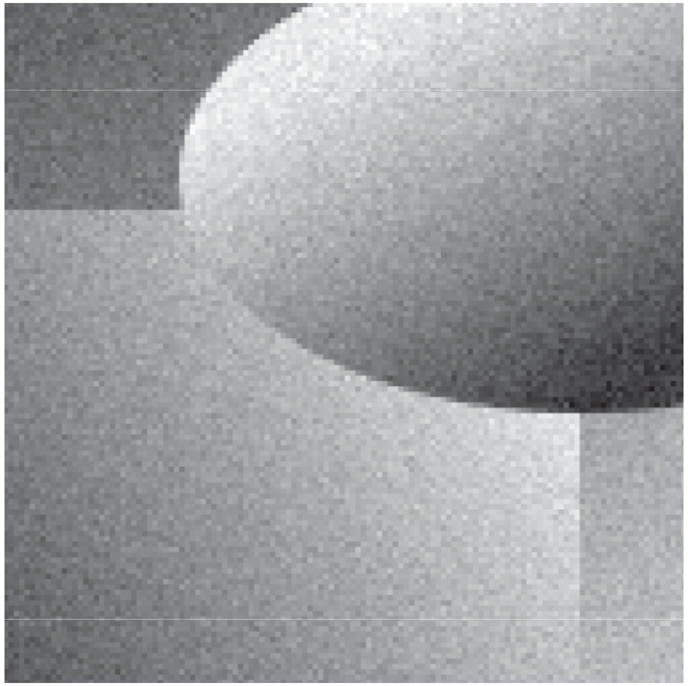

noisy input

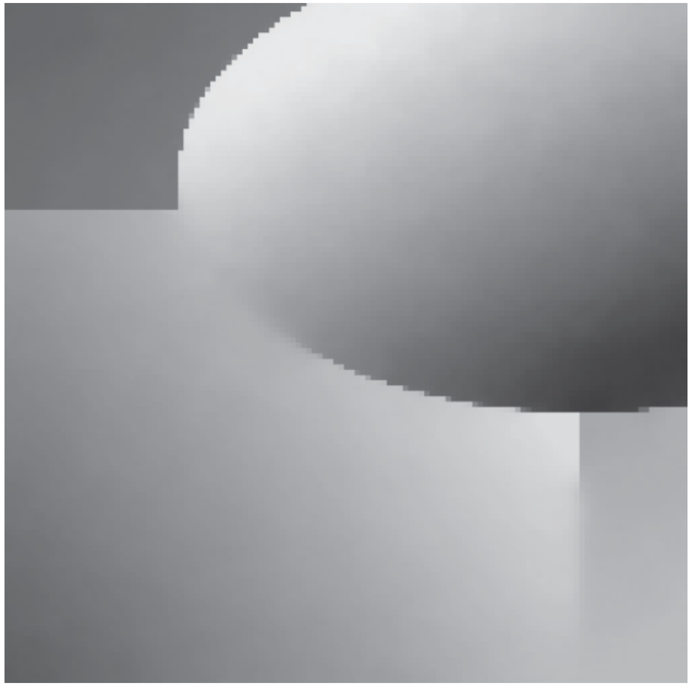

restoration

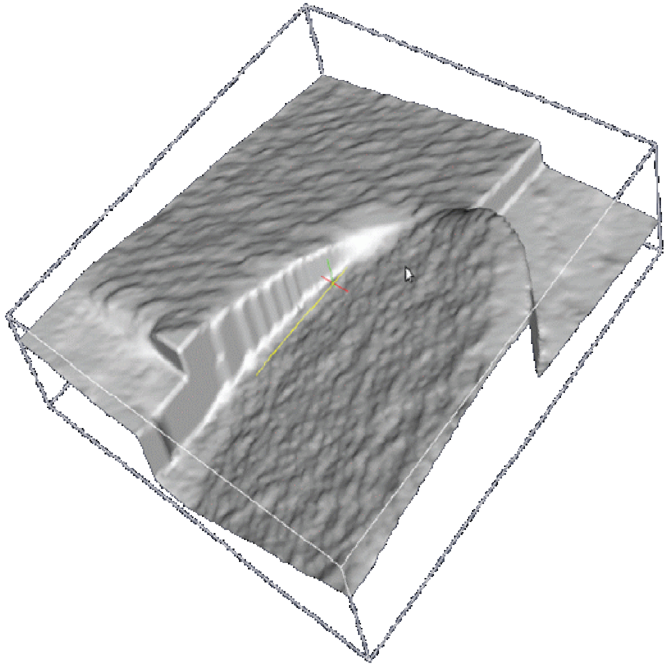

surface plot

Pock, Cremers, Bischof, Chambolle ICCV '09

Open boundaries and the crack tip problem

fixed boundary values

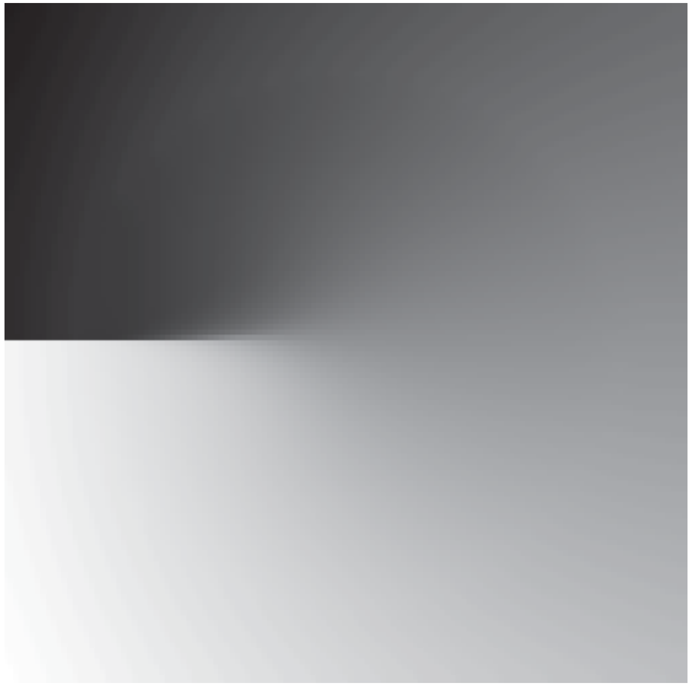

inpainted crack tip

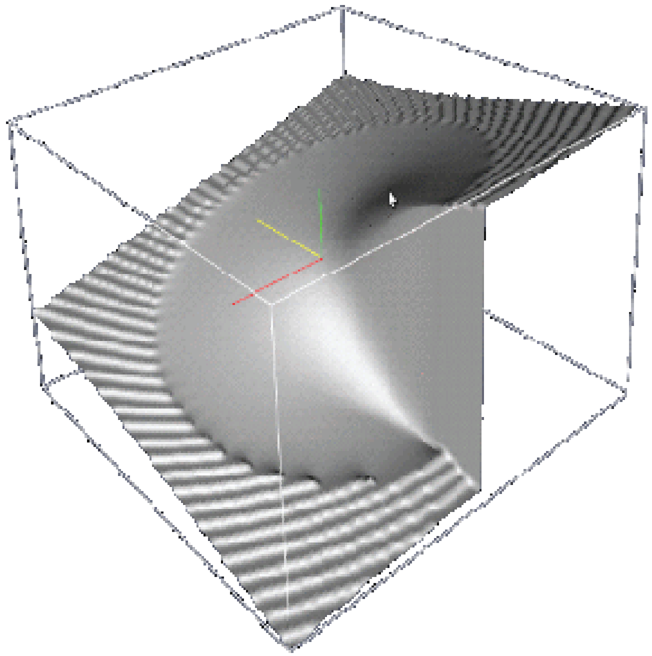

surface plot

Pock, Cremers, Bischof, Chambolle ICCV '09

Overview

Multi-label optimization

Minimal Partition Problems

The Mumford-Shah Problem

Label Layout Constraints

D. Cremers, B. Goldlücke, T. Pock

Imposing directional dependent penalties

Strekalovskiy, Cremers, ICCV 2011

Examples of label layouts

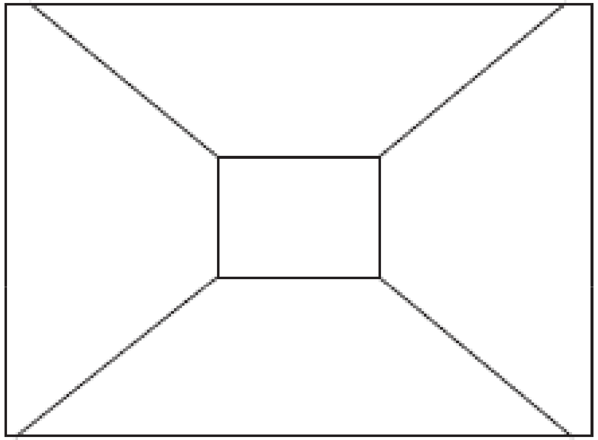

Five regions layout (Liu et al. [10])

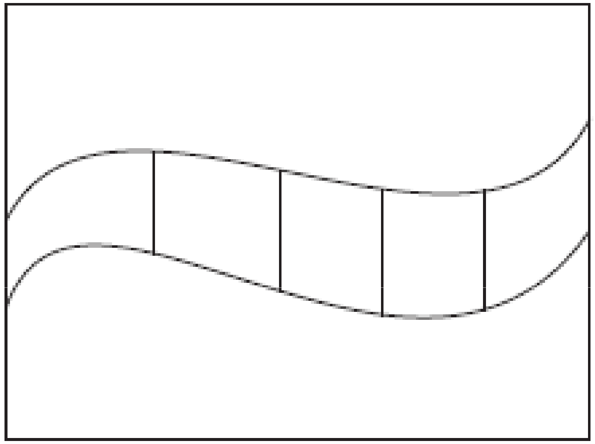

Tiered layout (Felzenszwalb et al. [4])

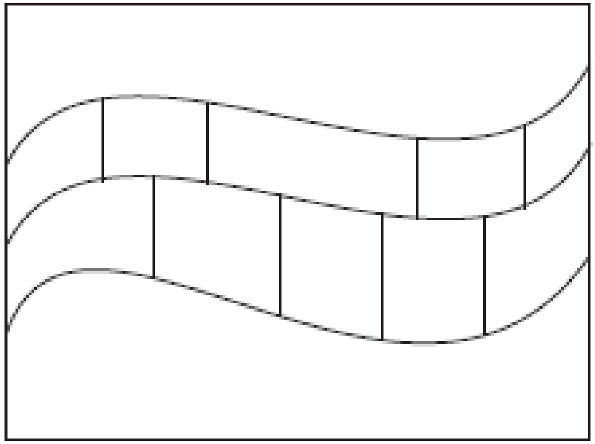

Four and more tiers

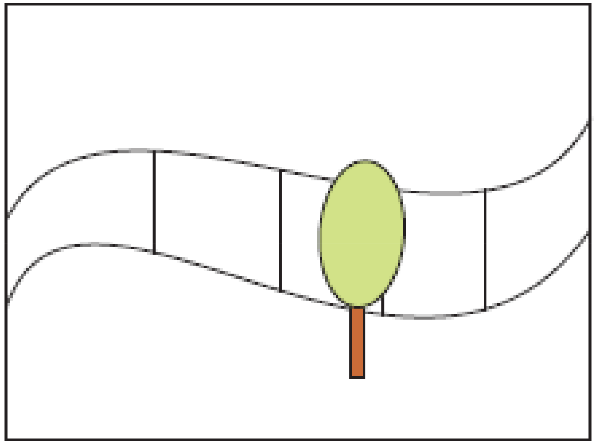

Floating objects

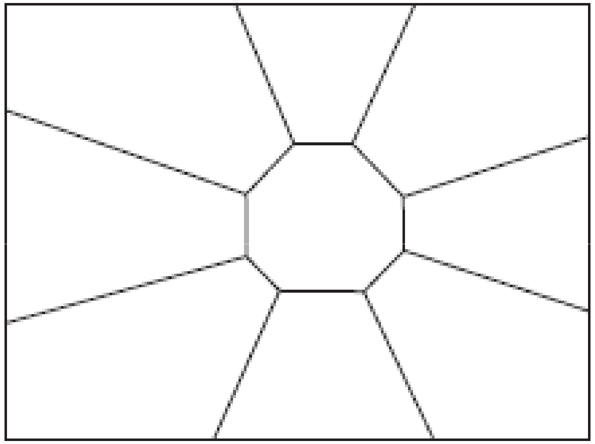

Convex shape prior

Strekalovskiy, Cremers, ICCV 2011

Generalization to directional dependent cost

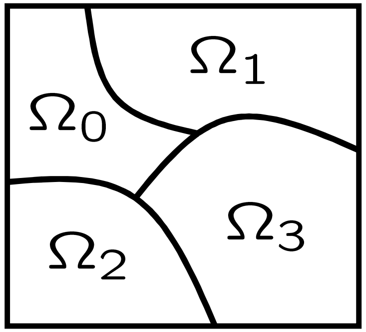

Reminder: With $v_i = 1_{\Omega_i}$, the segmentation problem is:

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

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

Consider instead the more general convex set:

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

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

Strekalovskiy, Cremers, ICCV 2011

The tiered layout

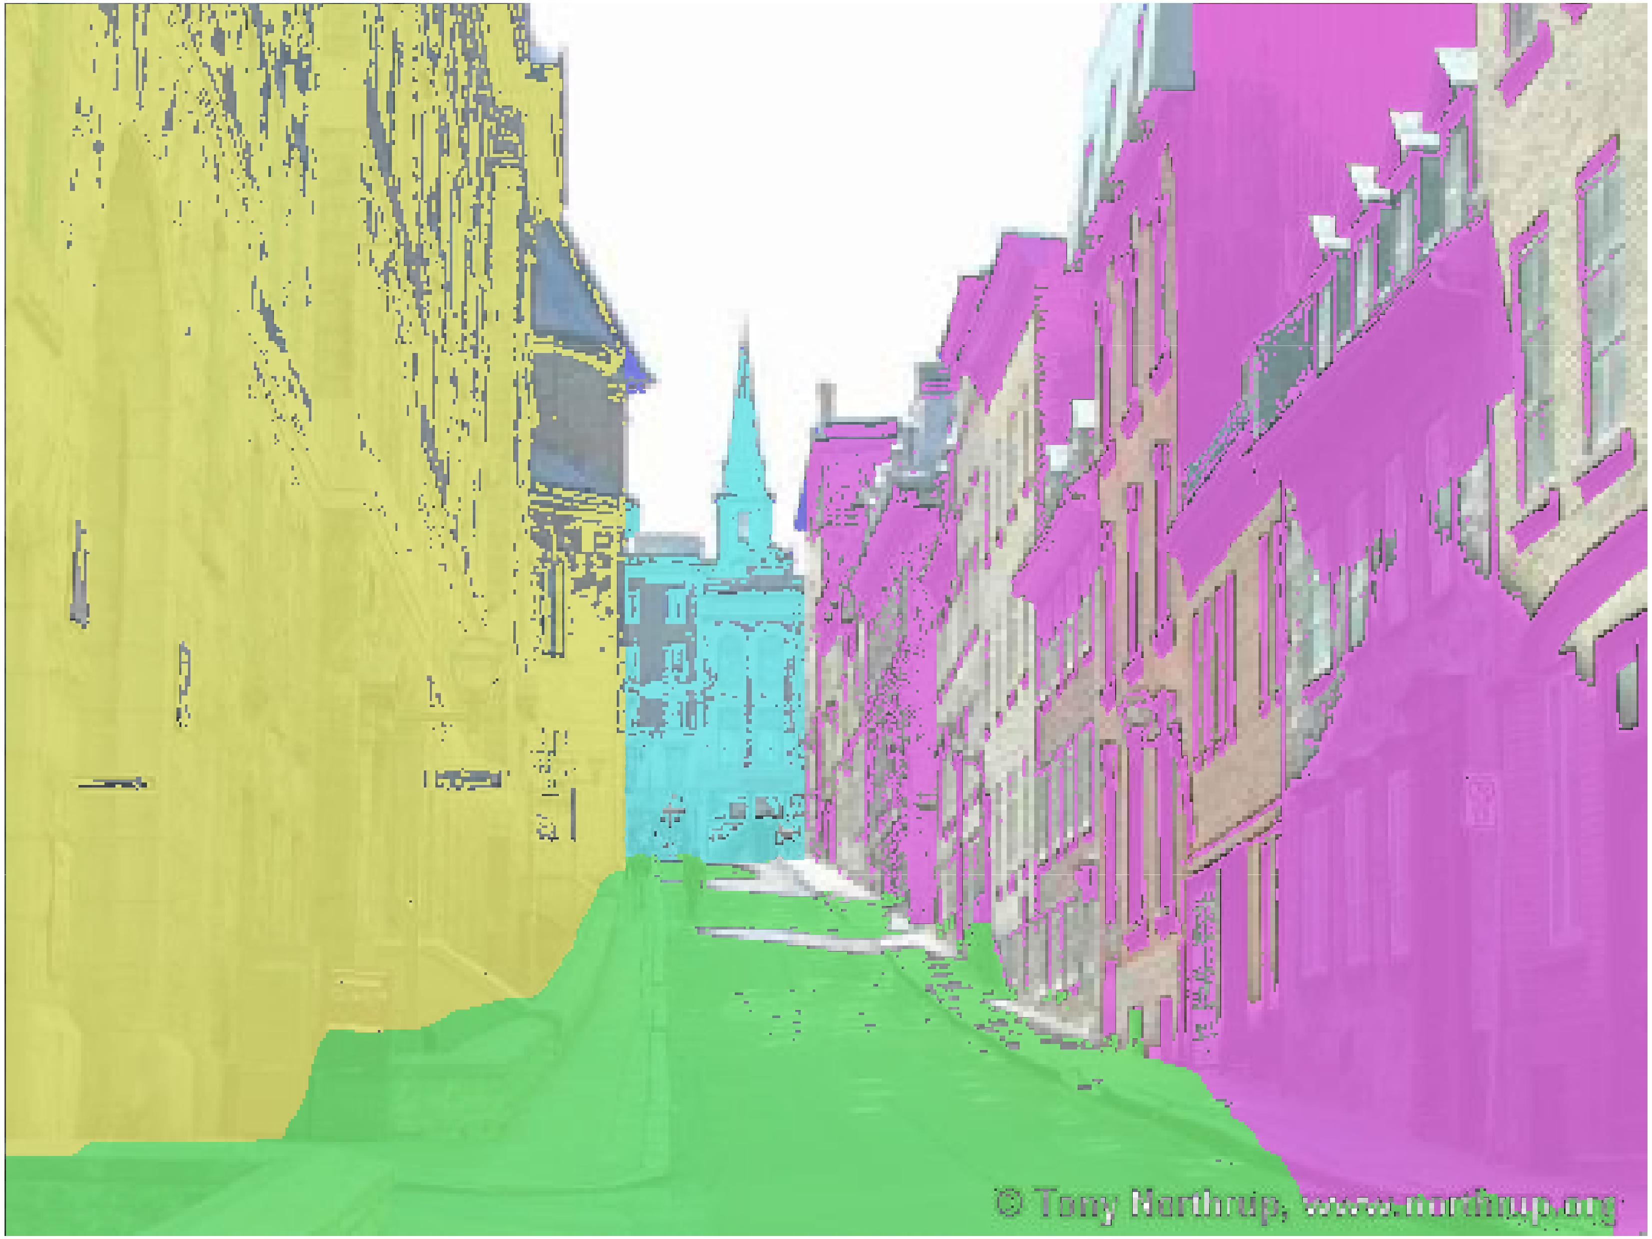

Strekalovskiy, Cremers, ICCV 2011

The five-regions layout

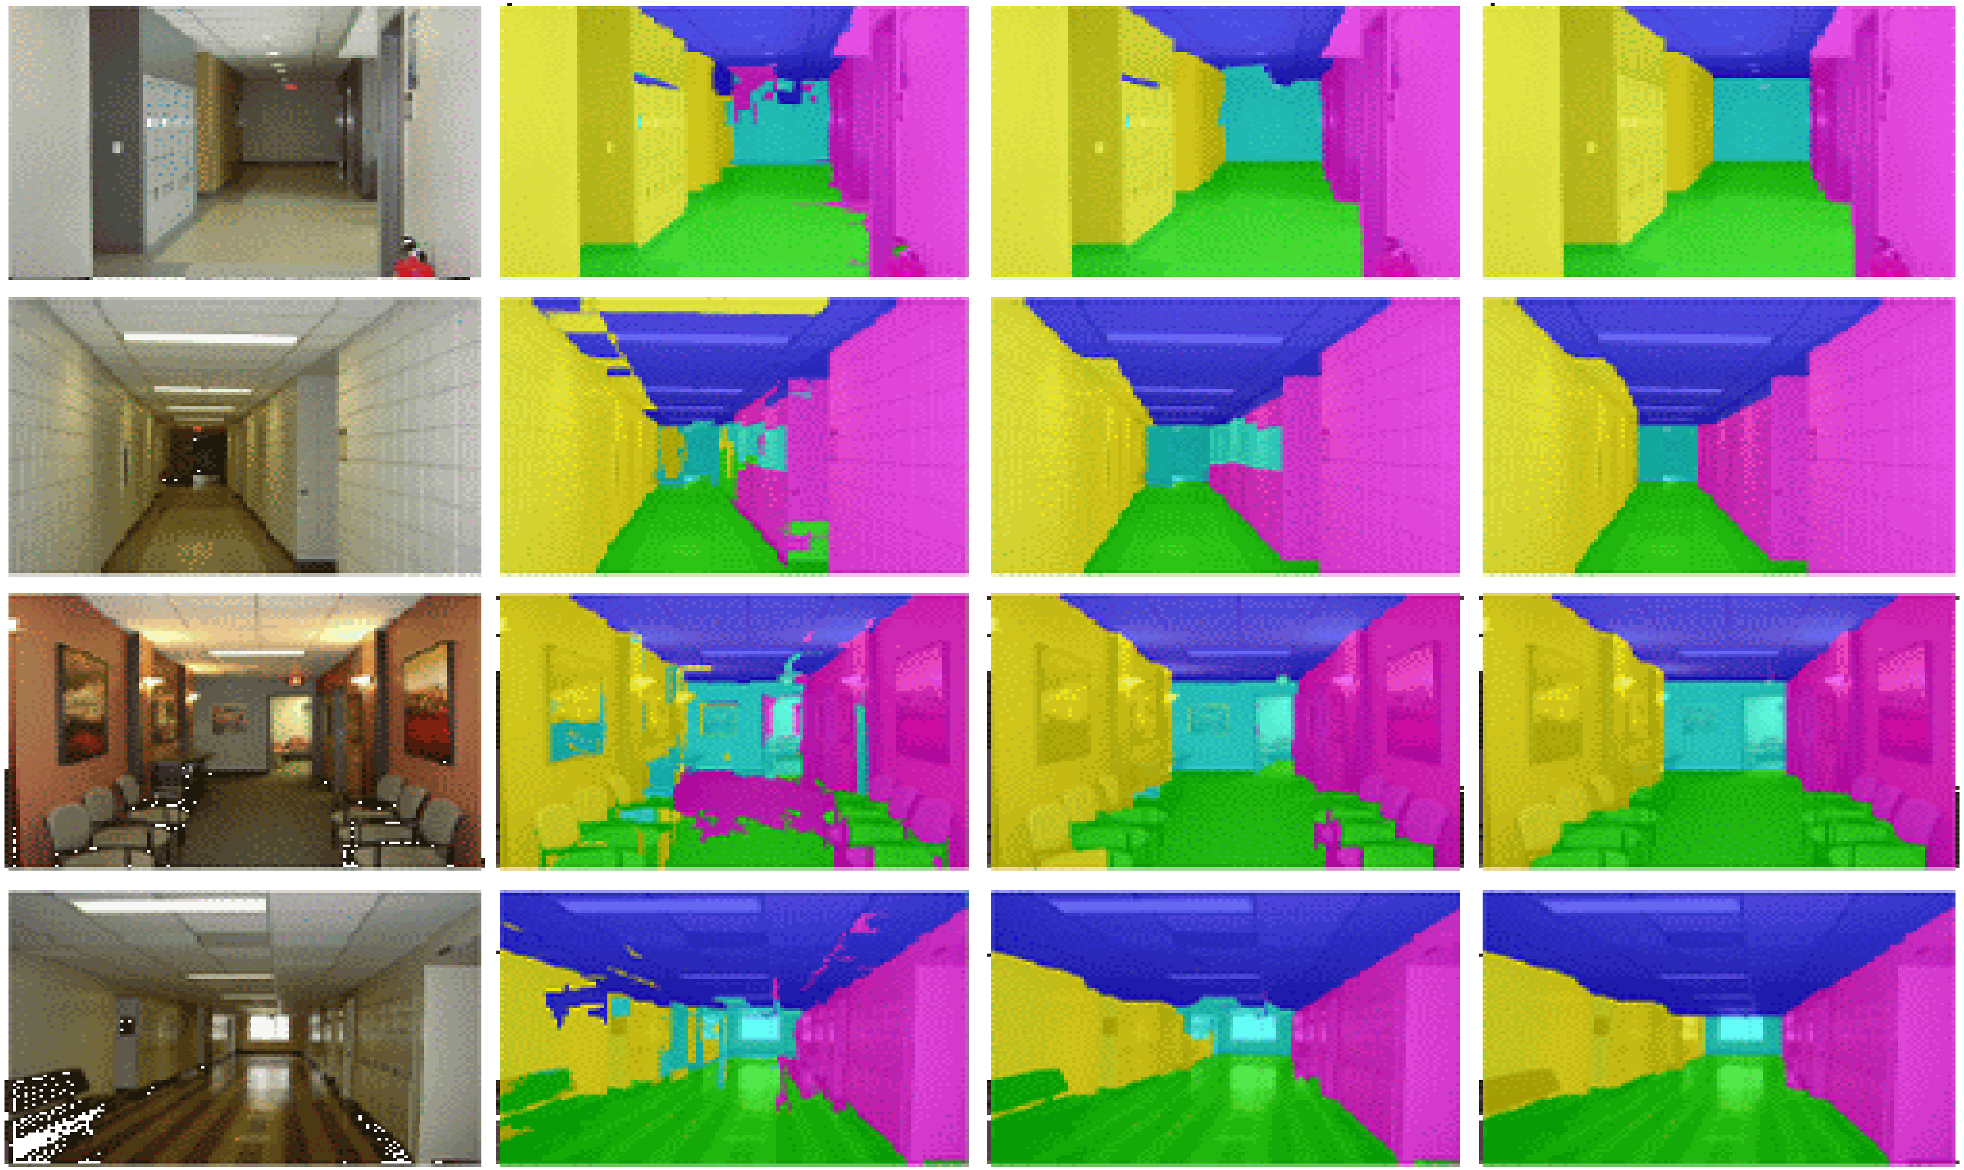

Input Data term Potts Ordering

Strekalovskiy, Cremers, ICCV 2011

D. Cremers, B. Goldlücke, T. Pock

The tiered layout

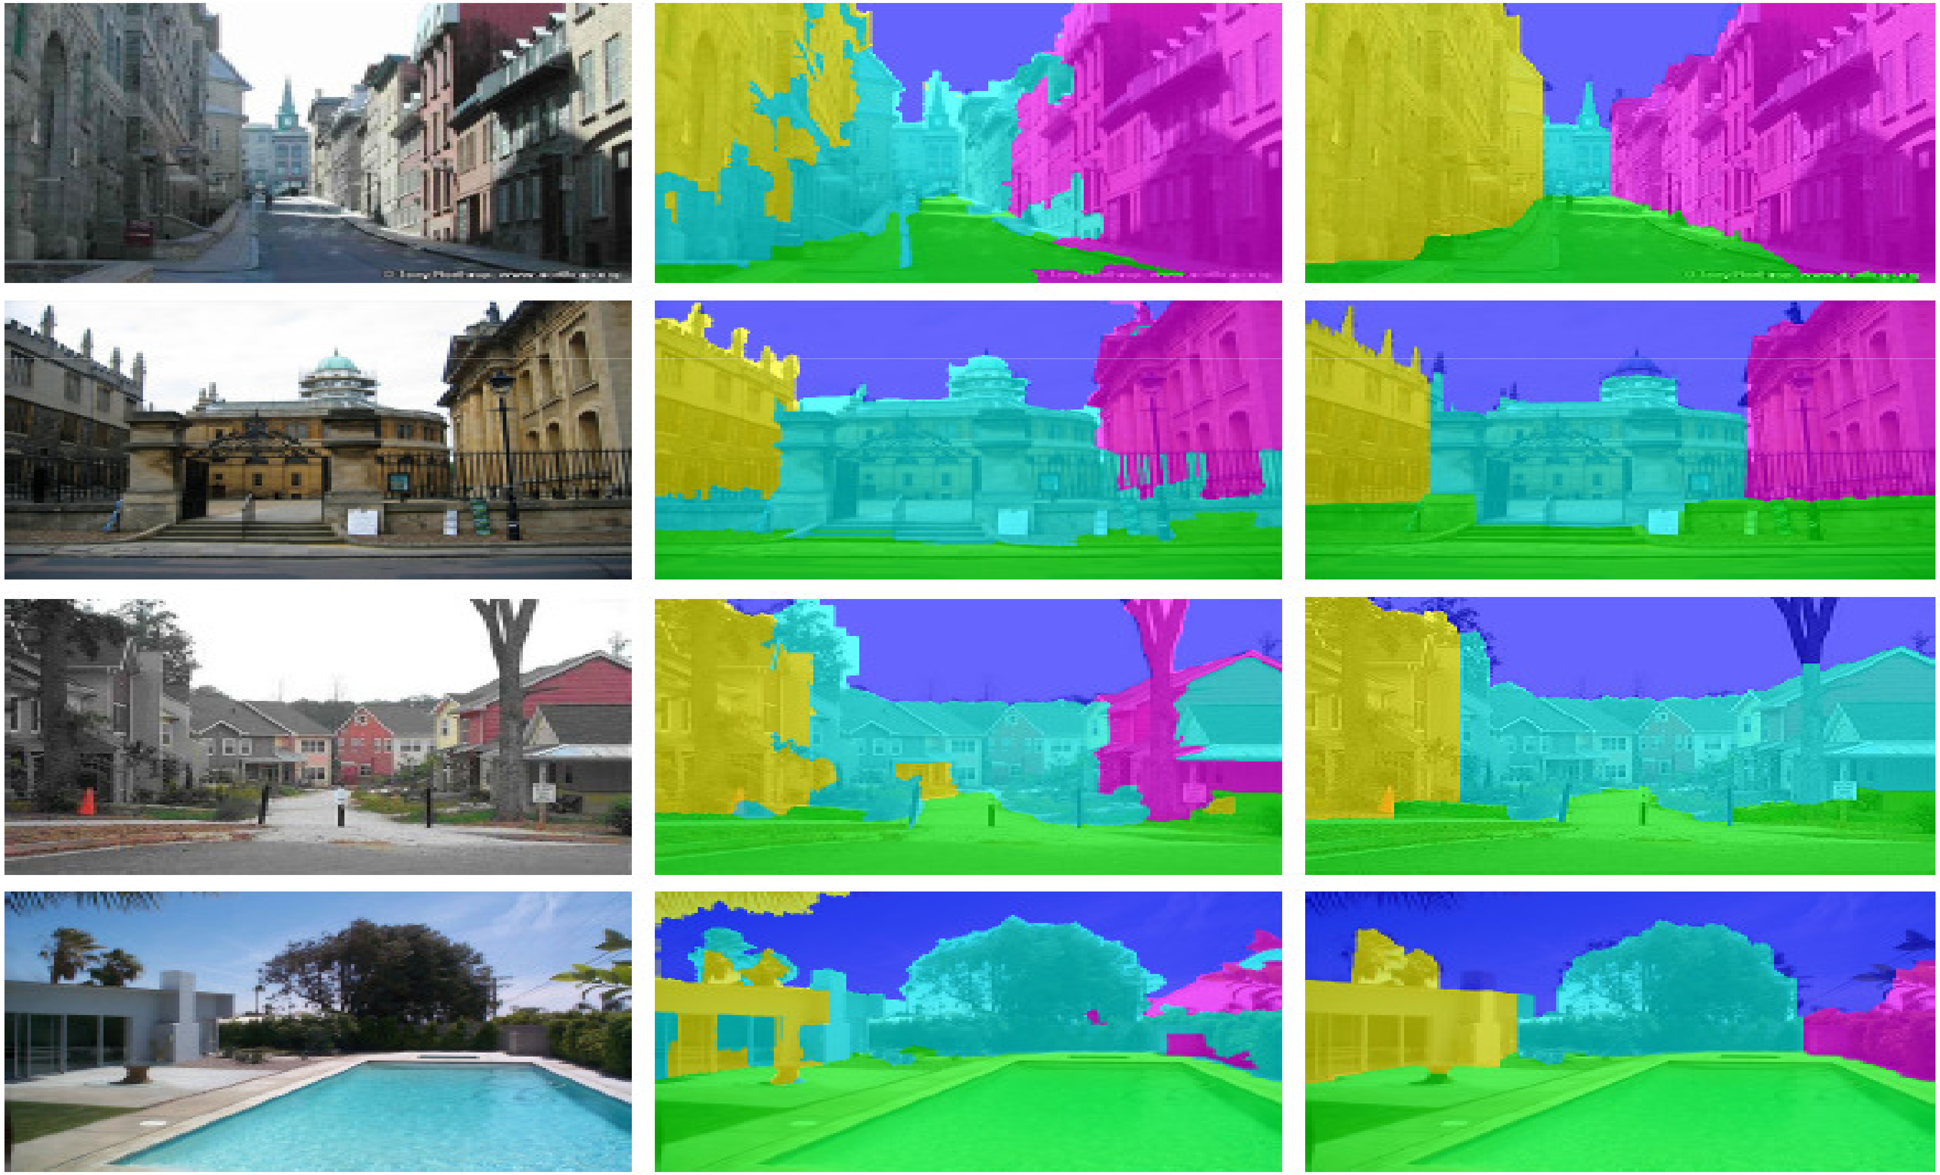

Input Potts Ordering

Strekalovskiy, Cremers, ICCV 2011

Encoding shape priors through label layouts

Input with seeds

Segmentation

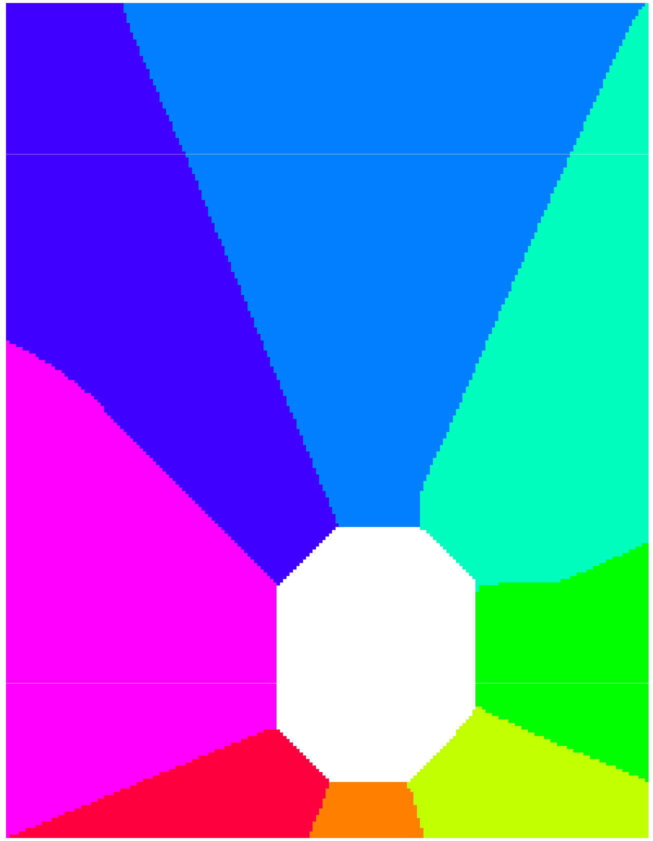

Computed labeling

Strekalovskiy, Cremers, ICCV 2011

Encoding shape priors through label layouts

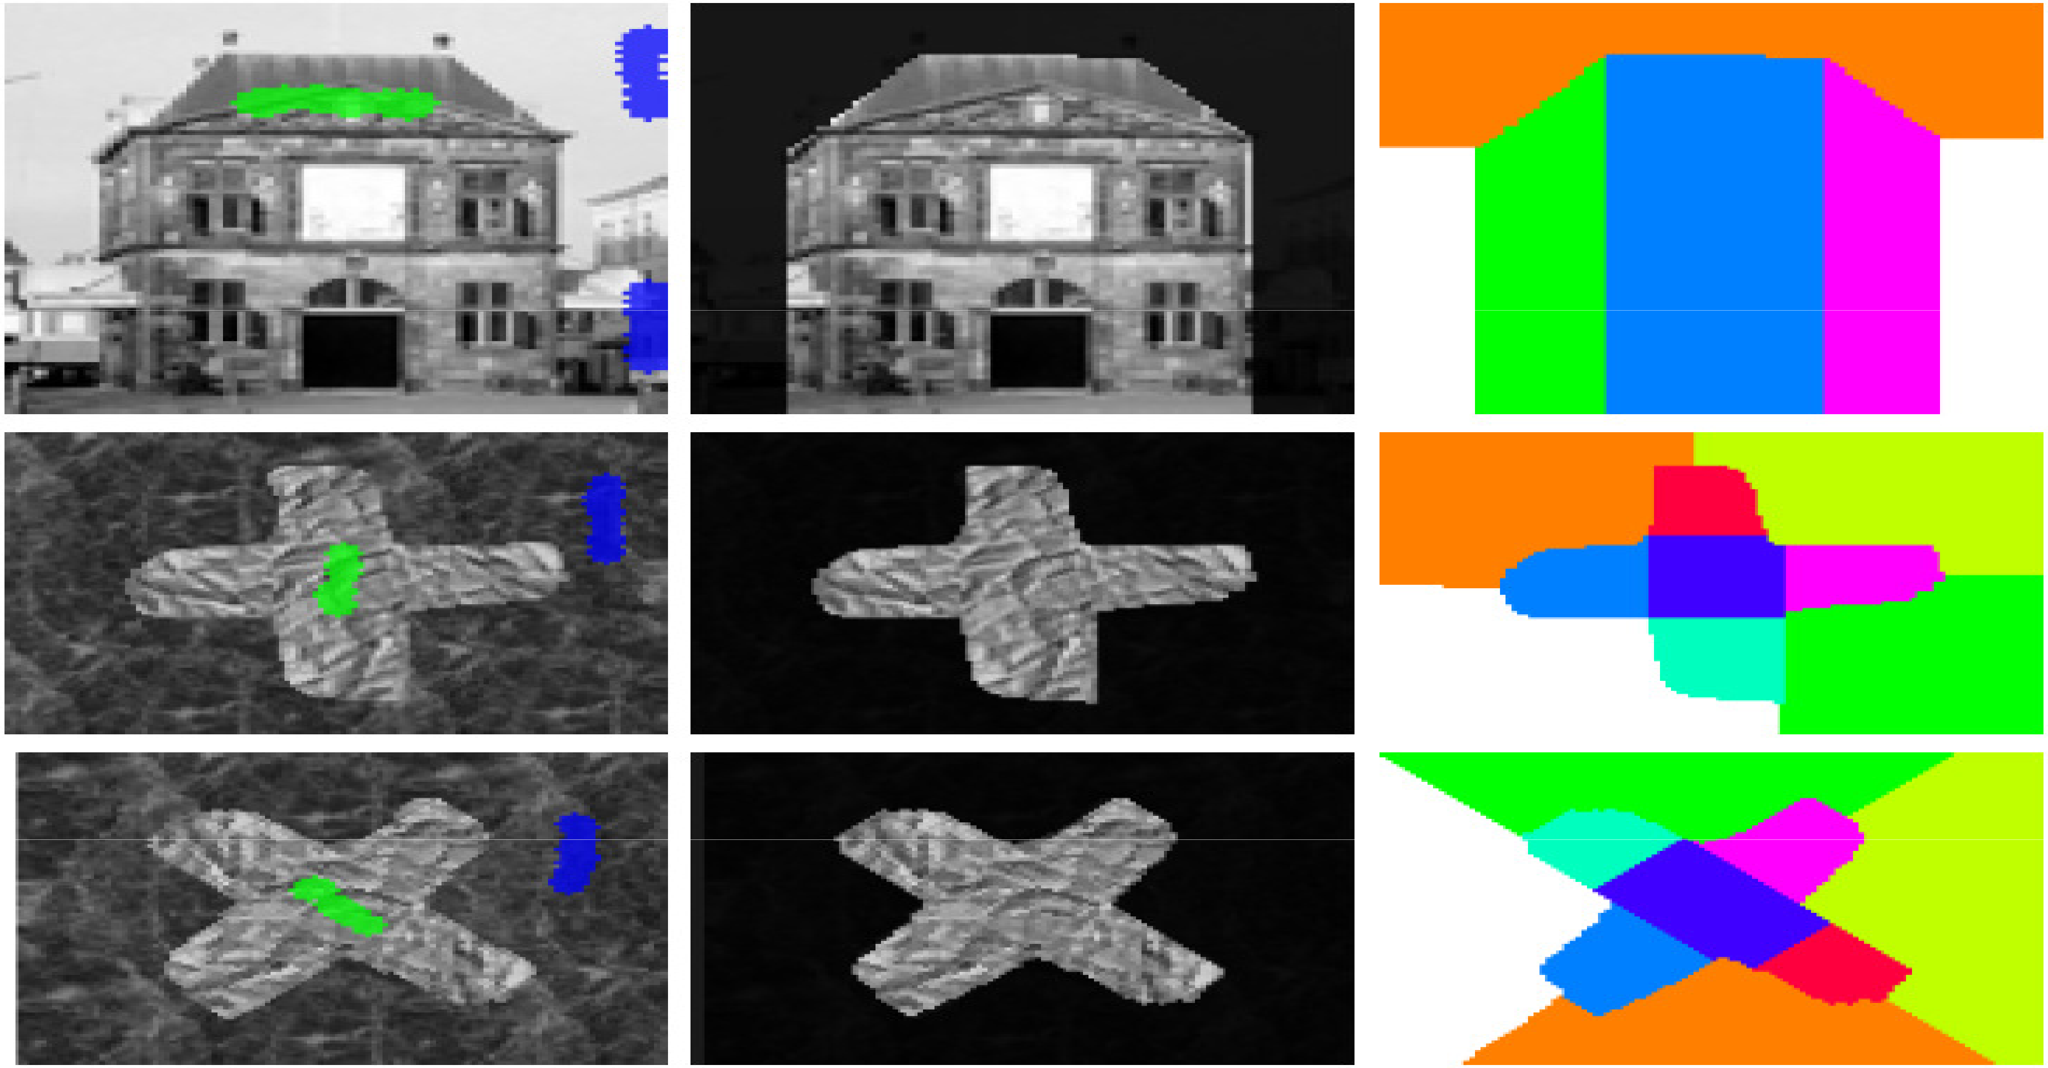

Input with seeds | Segmentation | Computed labeling

Strekalovskiy, Cremers, ICCV 2011

Summary

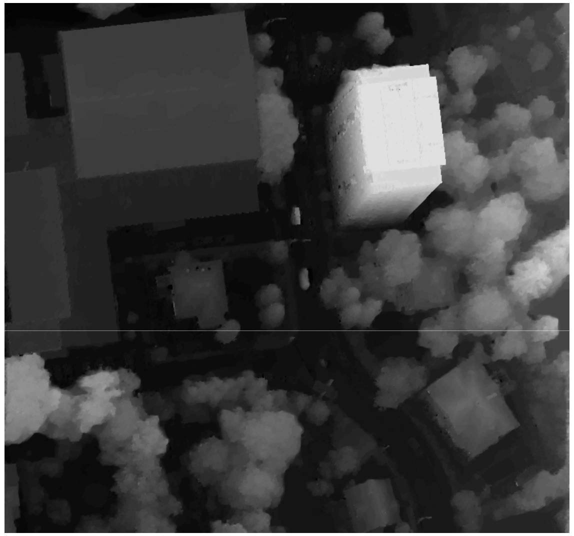

Multi-label optimization

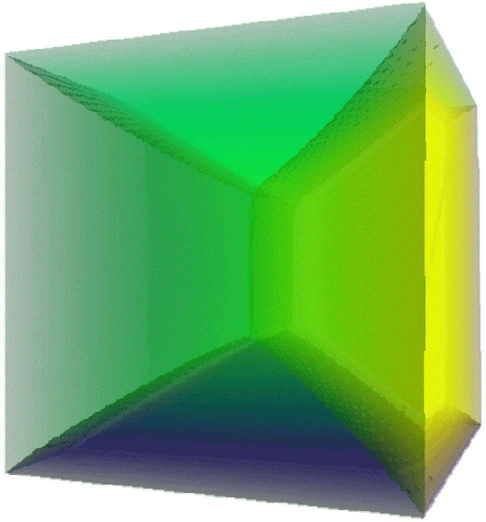

Minimal partition problems

Multi-region segmentation

Mumford-Shah

Label layout constraints

Imposing shape priors