Convex Relaxation Techniques for Functions with Values in a Riemannian Manifold

Daniel Cremers

Computer Science & Mathematics

Technische Universität München

Evgeny Strekalovskiy

Jan Lellmann

Computer Vision Challenges

Segmentation

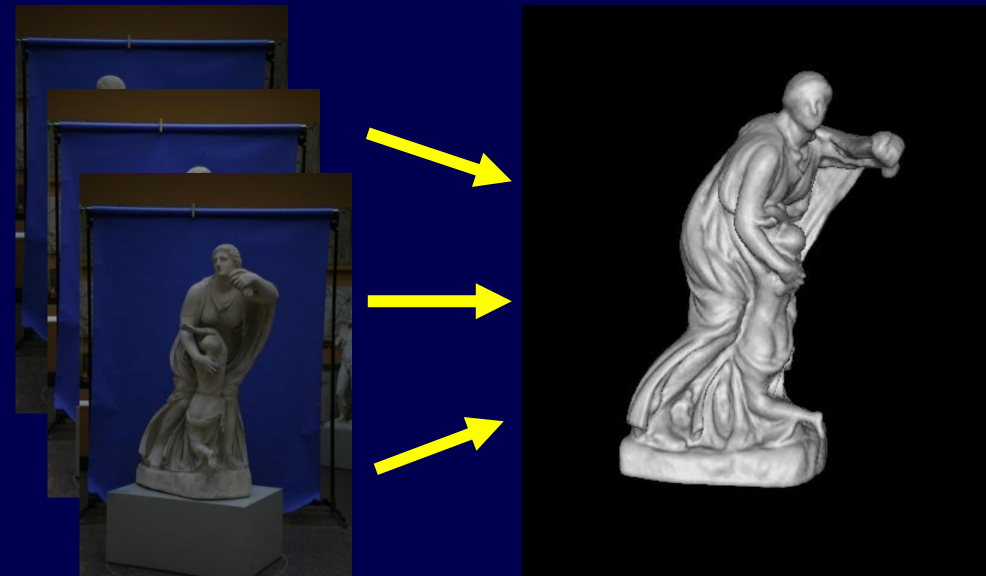

Multi-view Reconstruction

Space-time Reconstruction

Optical Flow

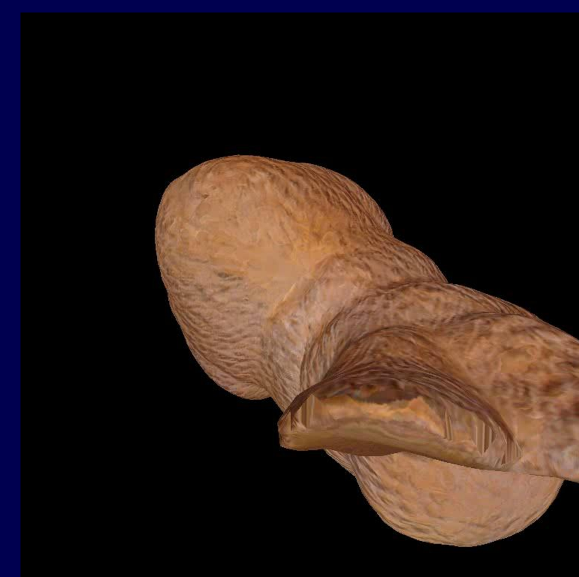

Super-resolution Texture

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

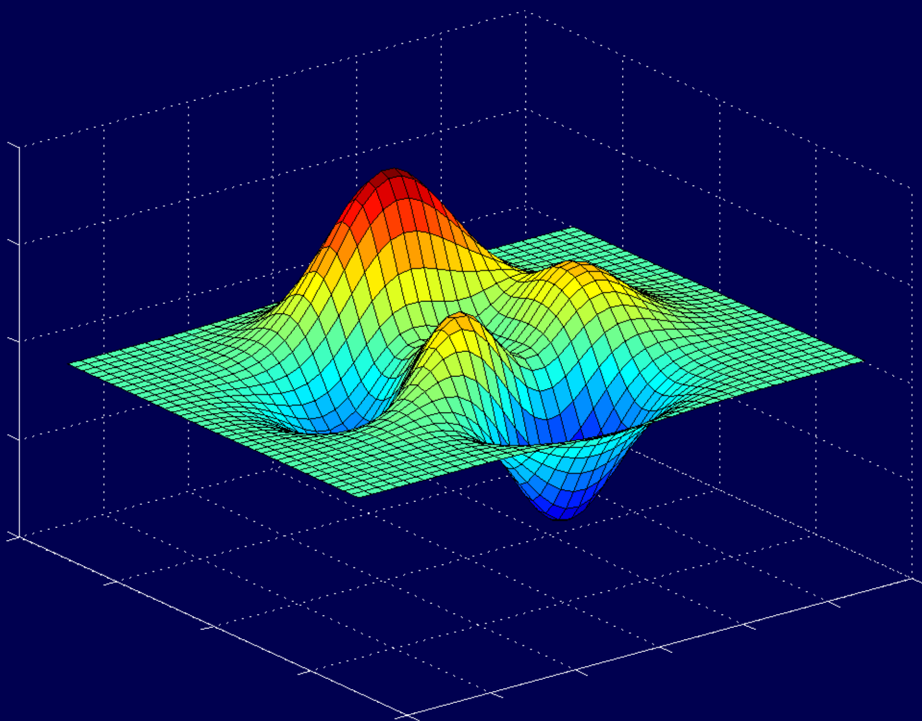

Non-convex energy

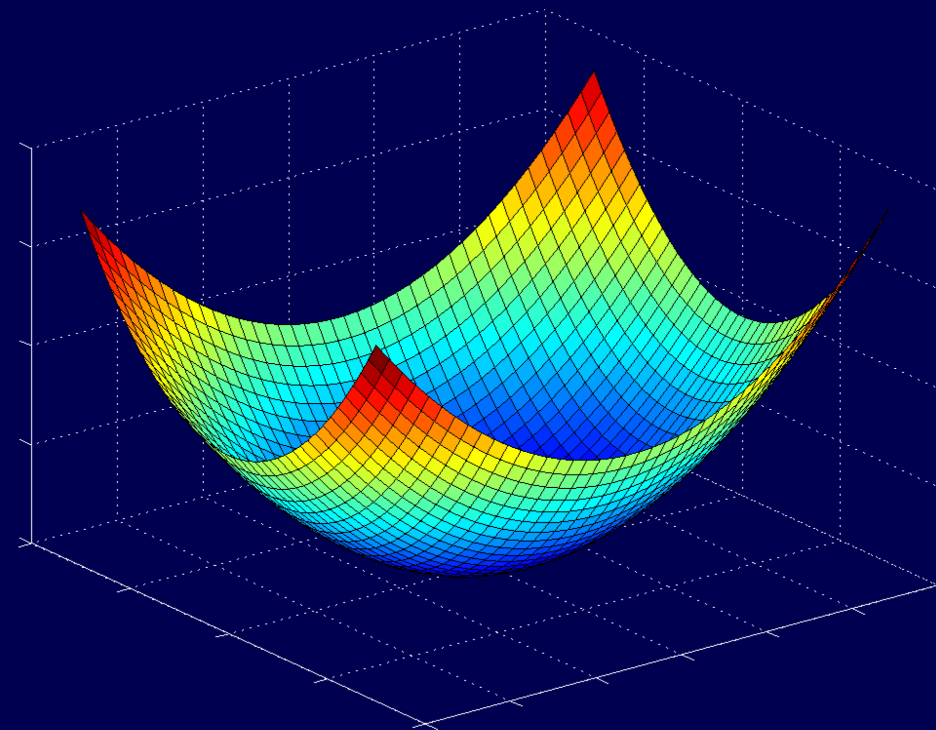

Convex energy

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,…*

color image processing

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

optical flow estimation

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

normal field inpainting

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

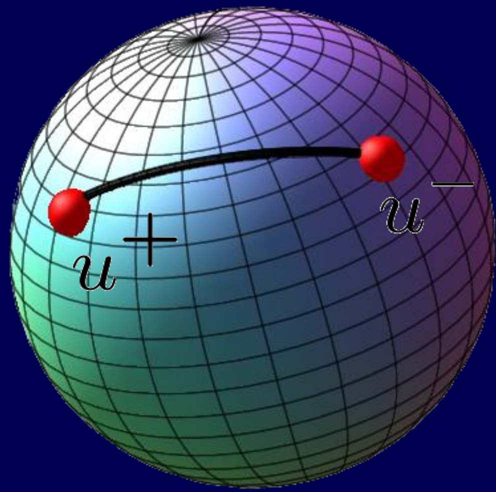

Problem statement

Convex regularizers

Nonconvex regularizers

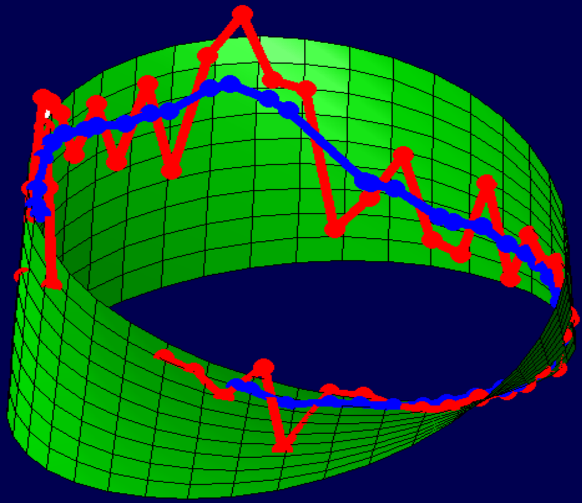

Manifold-valued functions

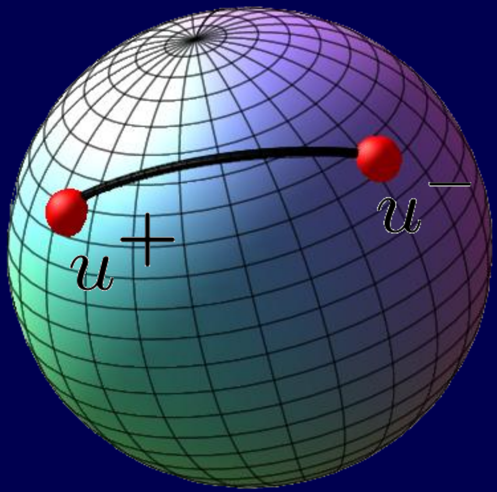

Problem statement

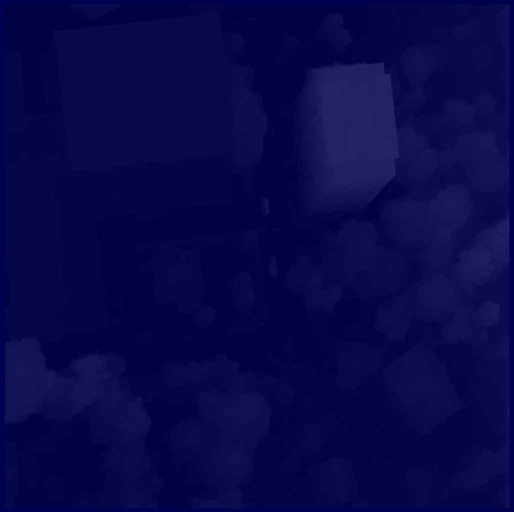

Convex regularizers

Nonconvex regularizers

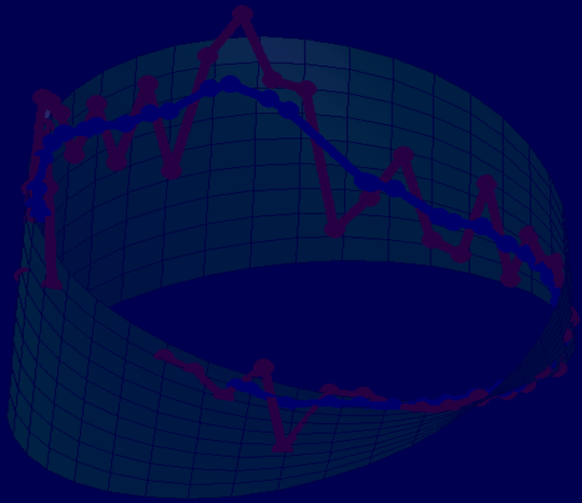

Manifold-valued functions

Consider the problem

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

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

Both the data term and the regularizer $R_{\mathcal{M}}$ may be non-convex.

Examples:

color denoising: $\mathcal{M} = \mathbb{R}^3, \;\; s(x,u) = |u - f|$

normal field denoising: $\mathcal{M} = \mathbb{S}^2, \;\; s(x,u) = |u - f|$

optical flow: $\mathcal{M} = \mathbb{R}^2, \;\; s(x,u) = |f_1 - f_2 \circ u|$

For a scalar-valued function $u \in BV(\Omega;\mathbb{R})$, the total variation is

$$TV(u) := \sup_{|\xi|\leq 1} \int_\Omega u \,\mathrm{div}\xi \, dx$$

For differentiable functions, integration by parts gives

$$TV(u) = \sup_{|\xi|\leq 1} \int_\Omega \nabla u \, \xi \, dx = \int_\Omega |\nabla u| \, dx$$

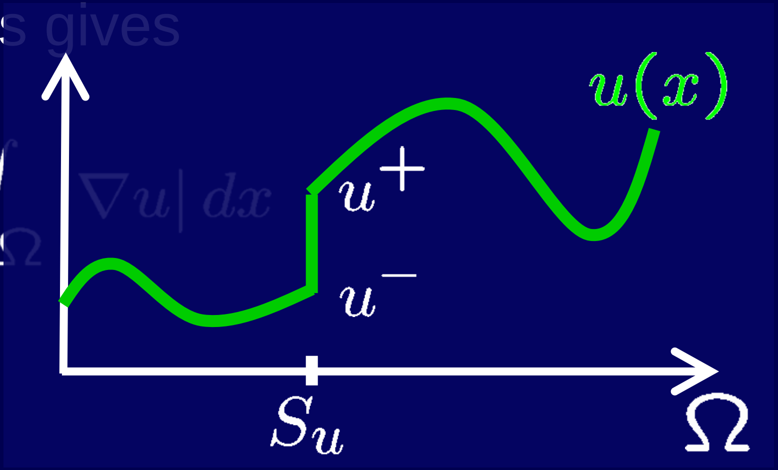

For $u \in BV(\Omega;\mathbb{R})$, we have

$$TV(u) = \int_{\Omega - S_u} |\nabla u| \, dx + \int_{S_u} |u^+ - u^-| \, d\mathcal{H}^{d-1} + |D^c(u)|$$

Herve, Shulman '89, Rudin, Osher Fatemi '92

Consider the problem

denoising: $s(x,u) = |u - f|$

flow: $s(x,u) = |f_1 - f_2 \circ u|$

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

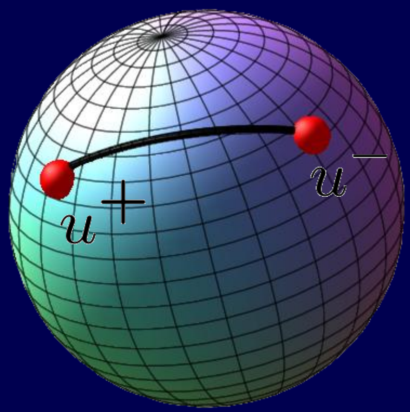

with a Riemannian manifold $\mathcal{M}$.

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

geodesic distance on the manifold

Giaquinta, Mucci, PAMQ '07, Cremers, Strekalovskiy, JMIV '12
Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Proposition 1:

For non-Euclidean manifolds, $TV_{\mathcal{M}}$ is generally nonconvex.

Proposition 2:

The discrete version of

$$\min_{u:\Omega\to\mathcal{S}^1} \int_\Omega s\big(x, u(x)\big)\, dx \;+\; TV_{\mathcal{S}^1}(u),$$

is NP-hard.

Proof:

For 3 equidistantly spaced labels, $TV_{\mathcal{S}^1}$ is equivalent to a Potts regularizer on three labels.

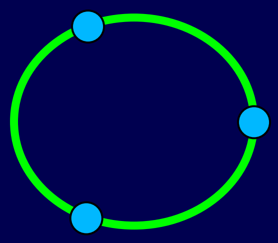

Cremers, Strekalovskiy, JMIV 2012.

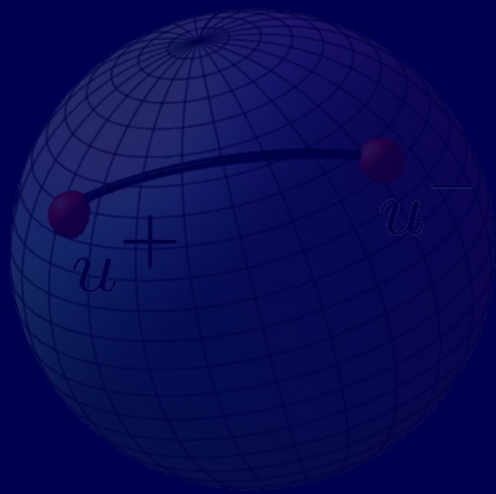

Problem statement


Convex regularizers


Nonconvex regularizers

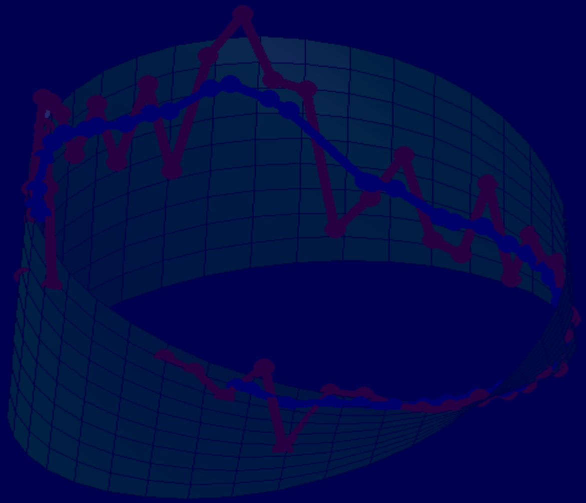
Manifold-valued functions

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

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

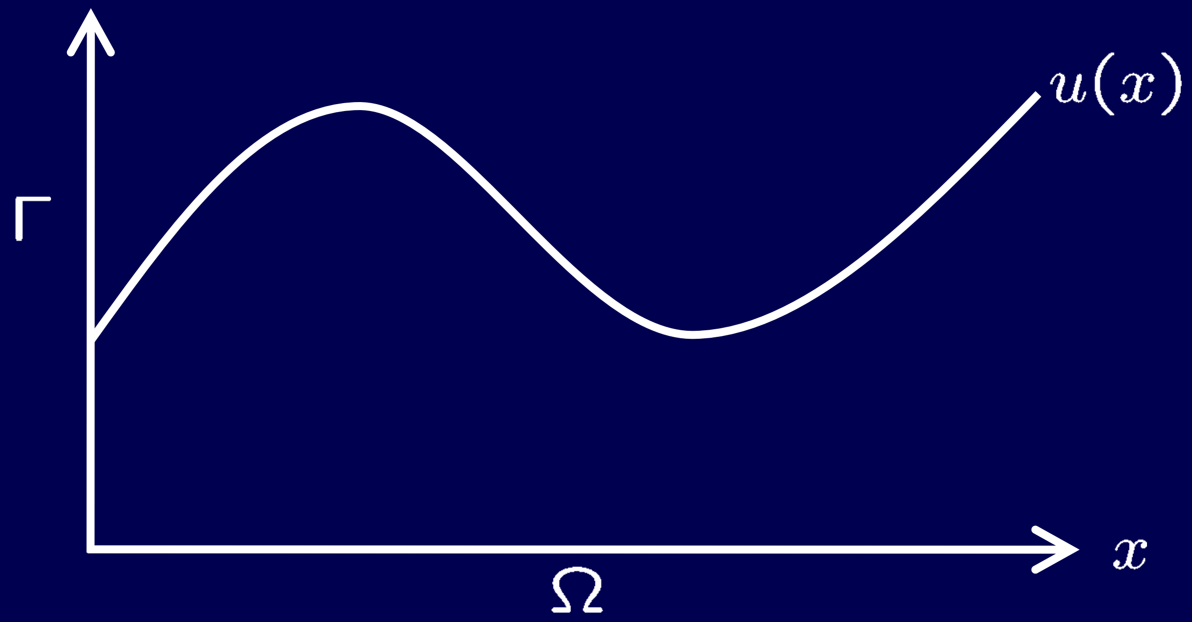

Pock , Schoenemann, Graber, Bischof, Cremers ECCV '08

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

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

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

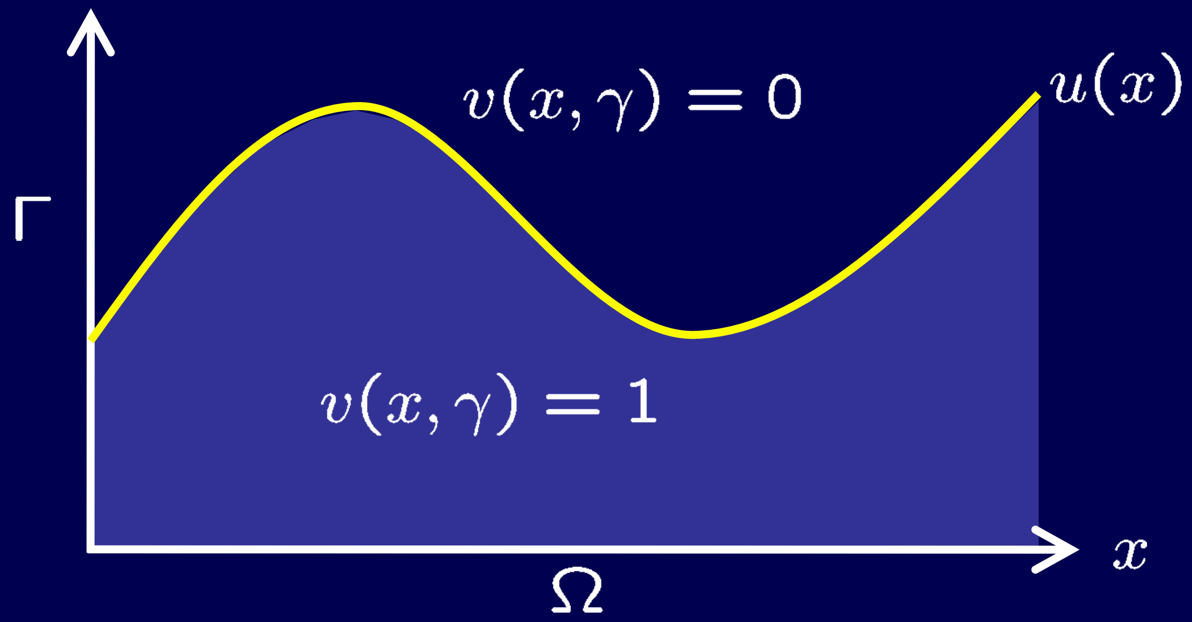

Pock , Schoenemann, Graber, Bischof, Cremers ECCV '08

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

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

nonconvex functional

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

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.

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

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

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.

The iterative 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

Reconstruction from Aerial Images

One of two input images

Courtesy of Microsoft

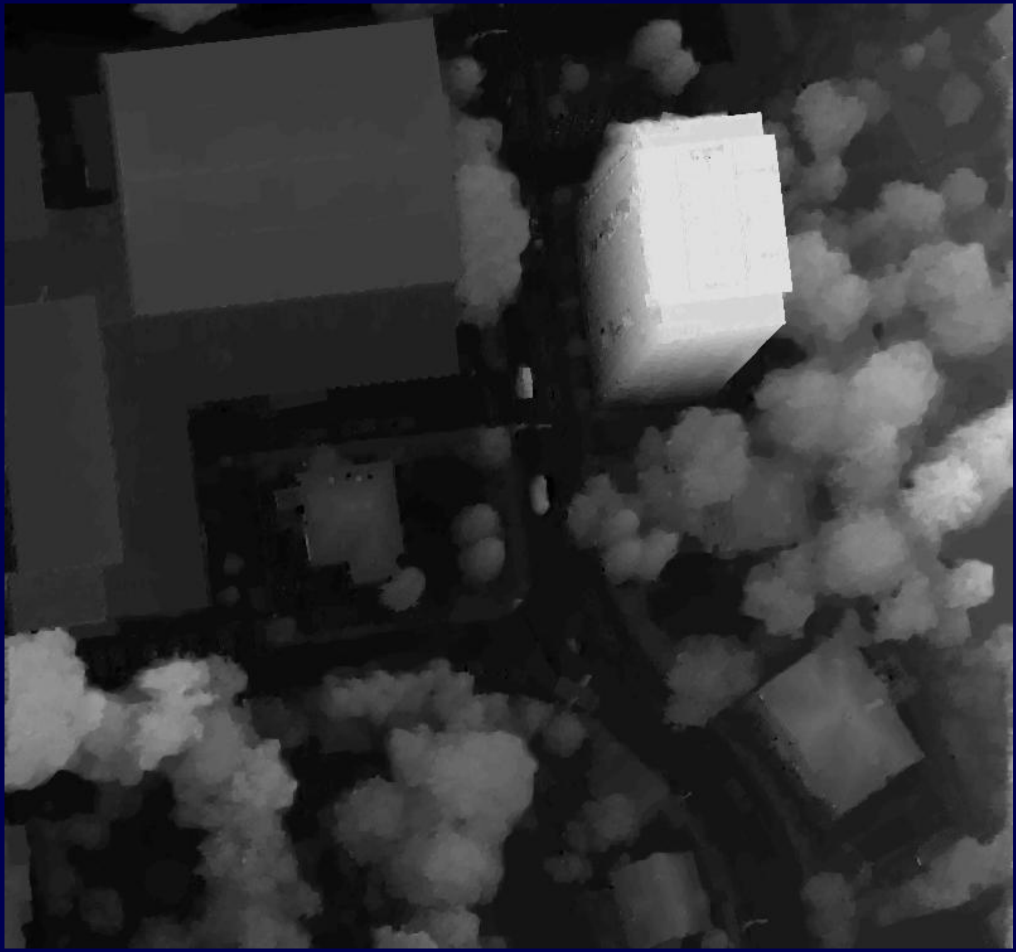

Depth reconstruction

Reconstruction from Aerial Images

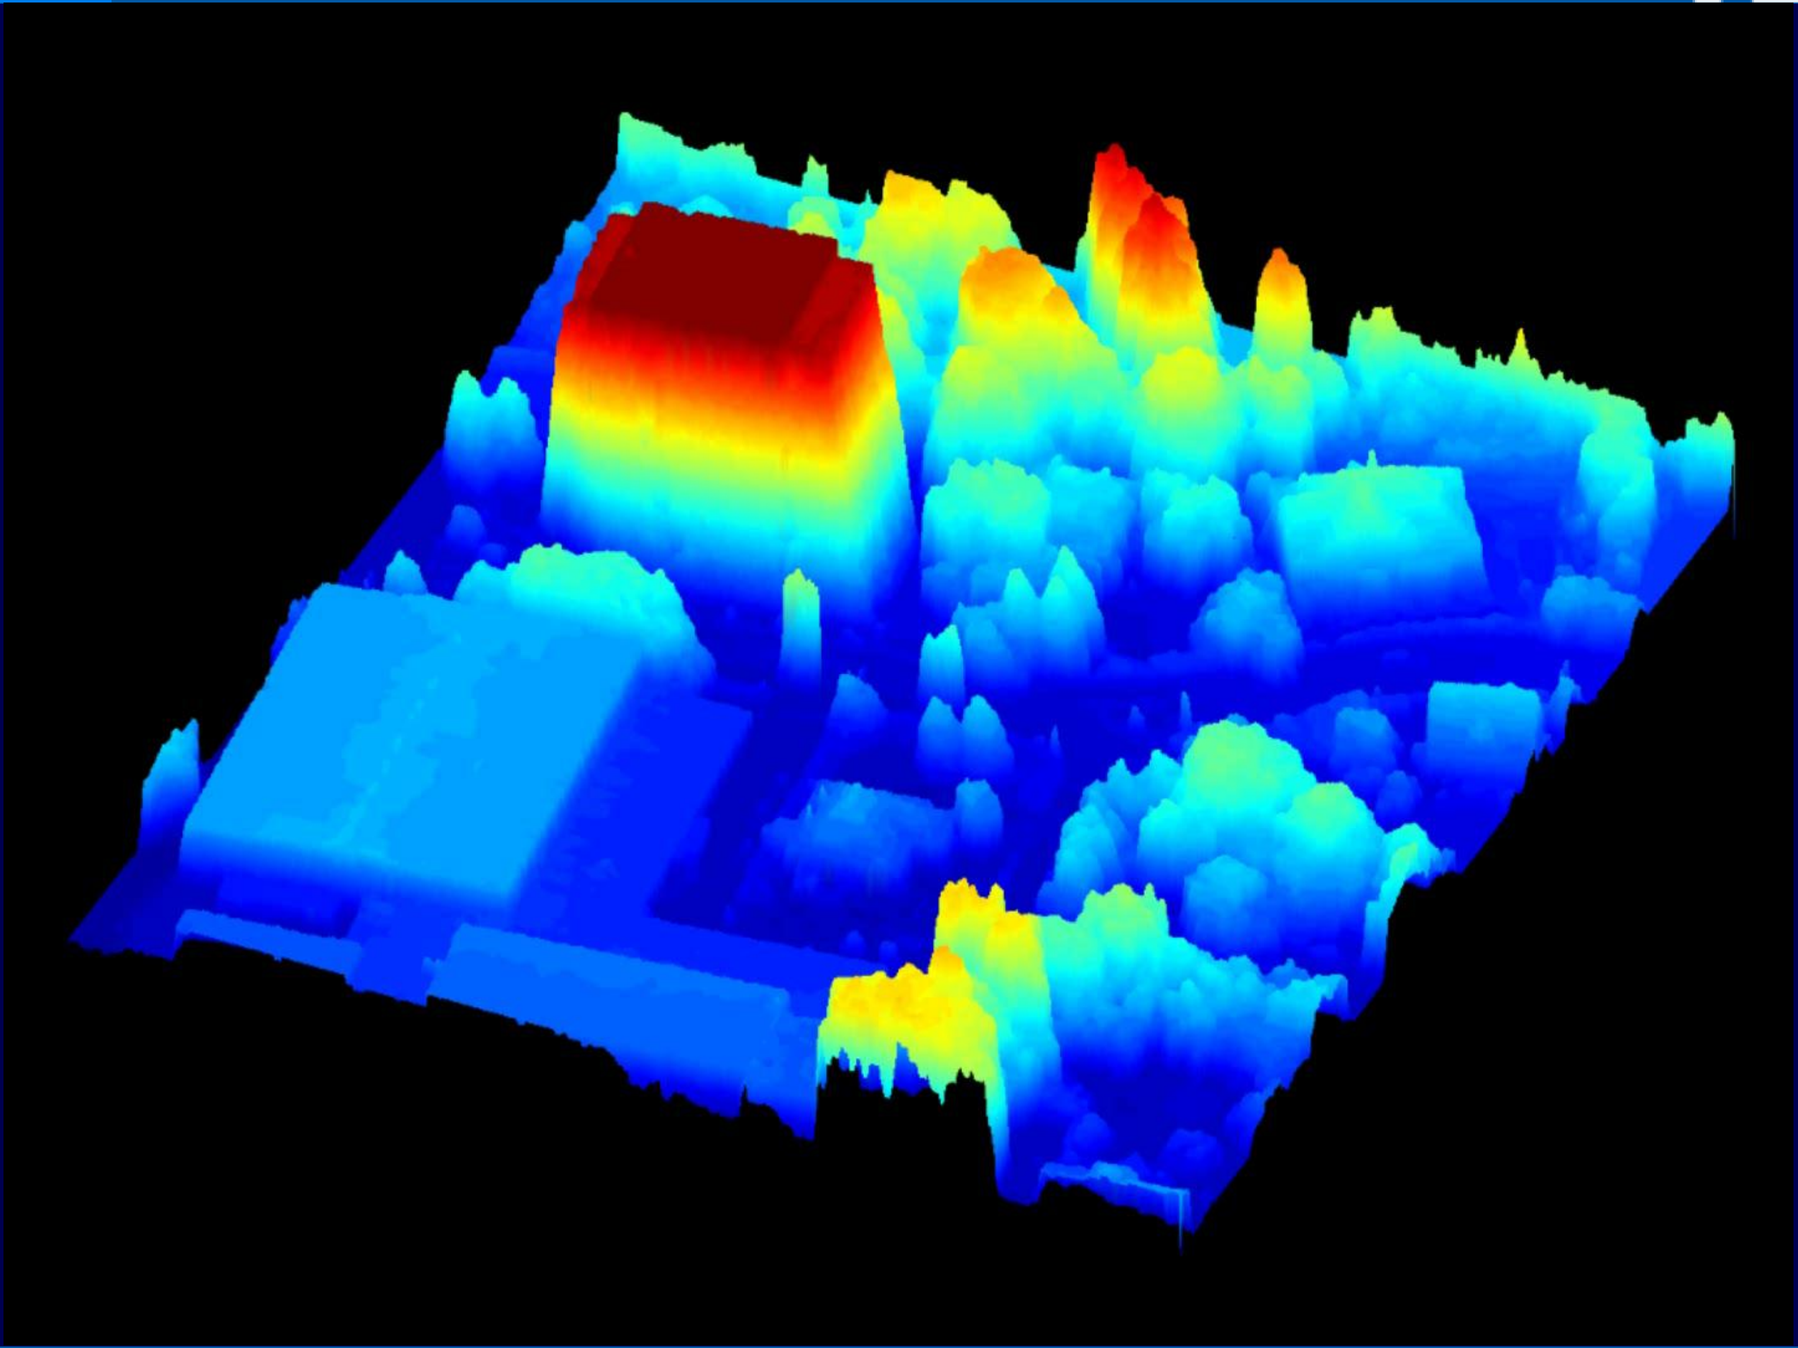

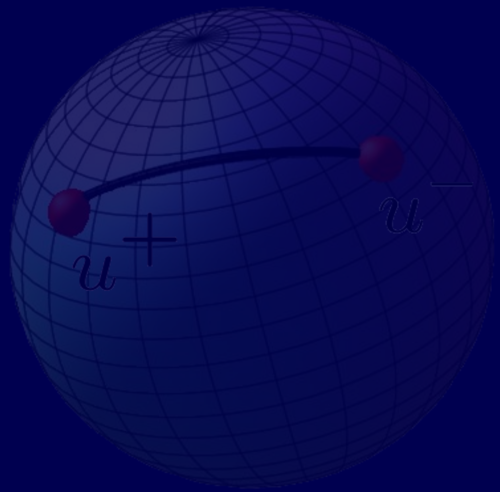

Problem statement

Convex regularizers

Nonconvex regularizers

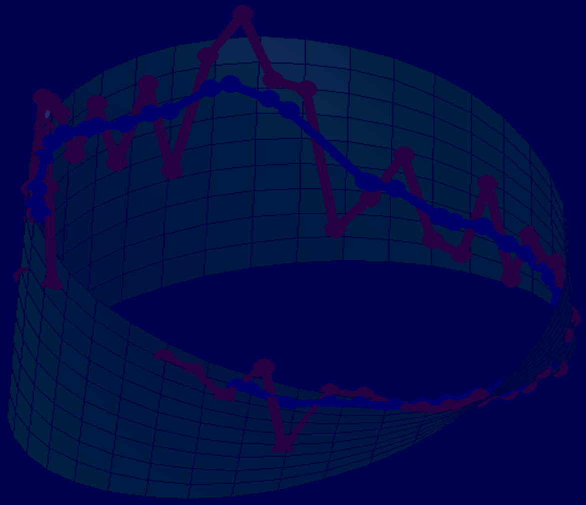

Manifold-valued functions

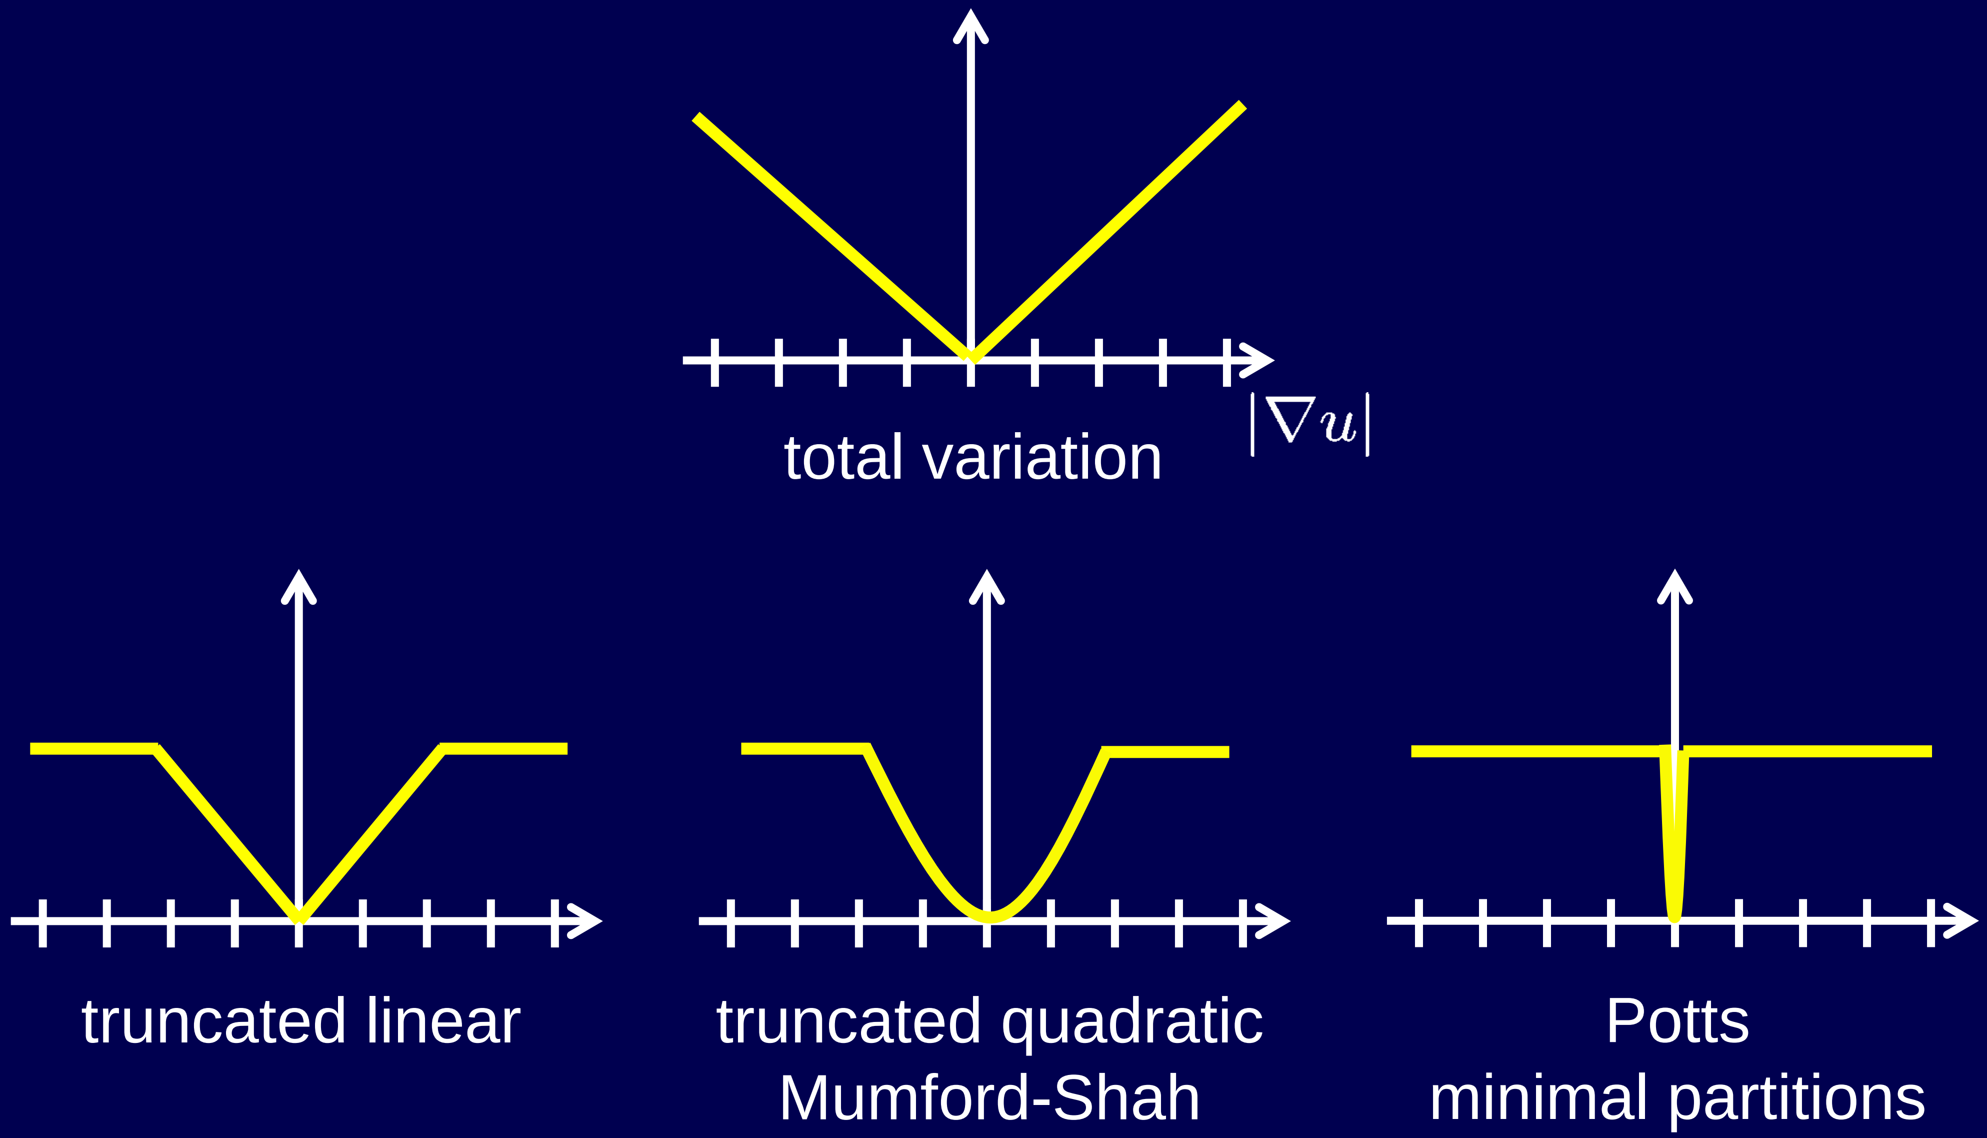
$|\nabla u|$
total variation
truncated linear
truncated quadratic
Mumford-Shah
Potts
minimal partitions

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

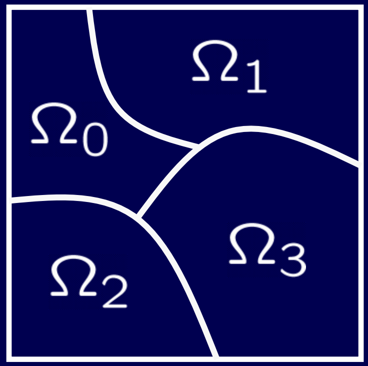

$$\text{s.t.} \quad \bigcup_i \Omega_i = \Omega \subset \mathbb{R}^d, \quad \text{and } \Omega_i \cap \Omega_j = \emptyset \;\; \forall i \neq j$$

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

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

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

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

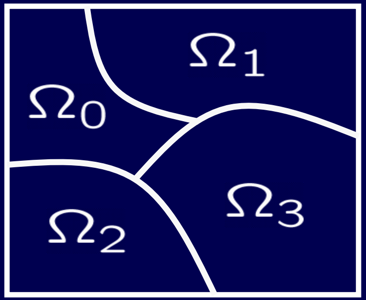

Reminder: With $v_i = 1_{\Omega_i}$, the segmentation problem is:

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

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

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 direction ν.

Strekalovskiy, Cremers, ICCV 2011

Strekalovskiy, Cremers, ICCV 2011

Strekalovskiy, Cremers, ICCV 2011

Piecewise Smooth Approximation

Input image

piecewise constant

piecewise smooth

Pock, Cremers, Bischof, Chambolle ICCV ’09

Color Mumford-Shah

Input image

Channelwise MS

Vectorial MS

Strekalovskiy, Chambolle, Cremers, CVPR '12

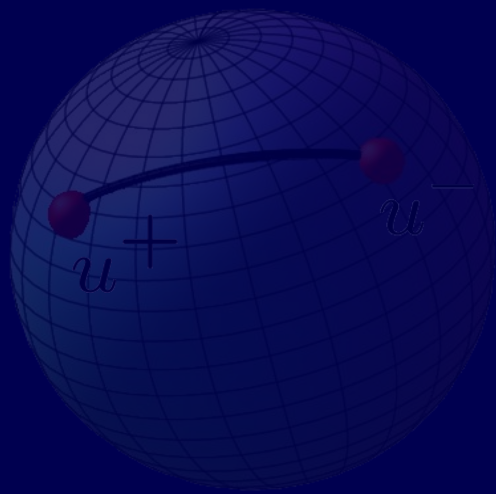

Problem statement

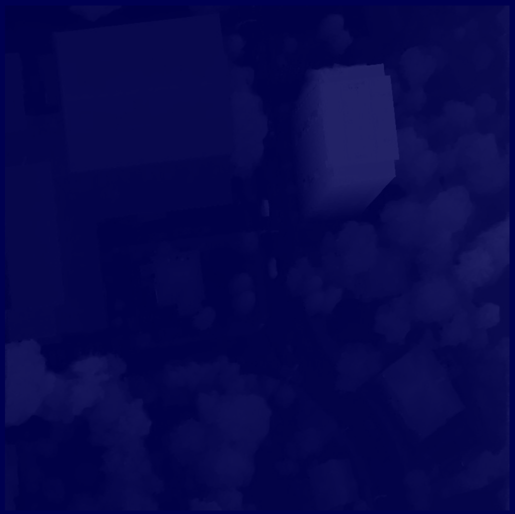

Convex regularizers

Nonconvex regularizers

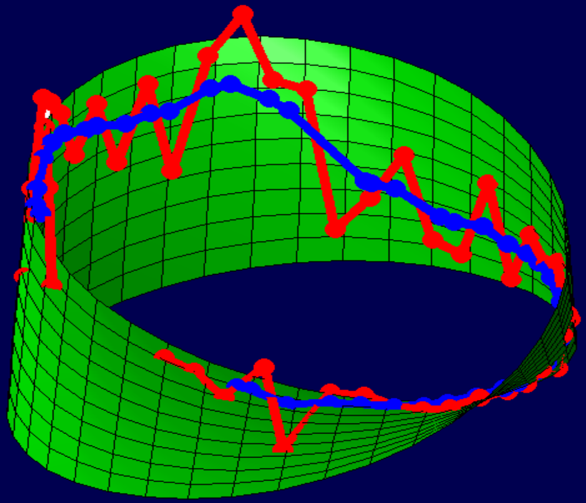

Manifold-valued functions

Consider the problem

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

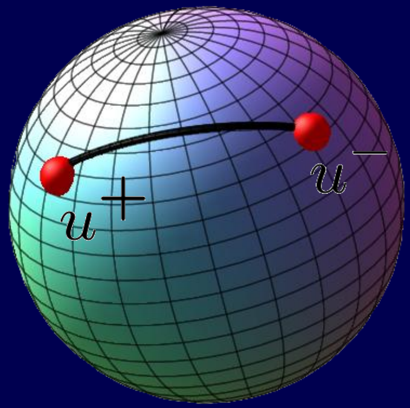

with a Riemannian manifold $\mathcal{M}$.

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

geodesic distance on the manifold

Giaquinta, Mucci, PAMQ '07, Cremers, Strekalovskiy, JMIV '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 local manifold structure

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

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

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

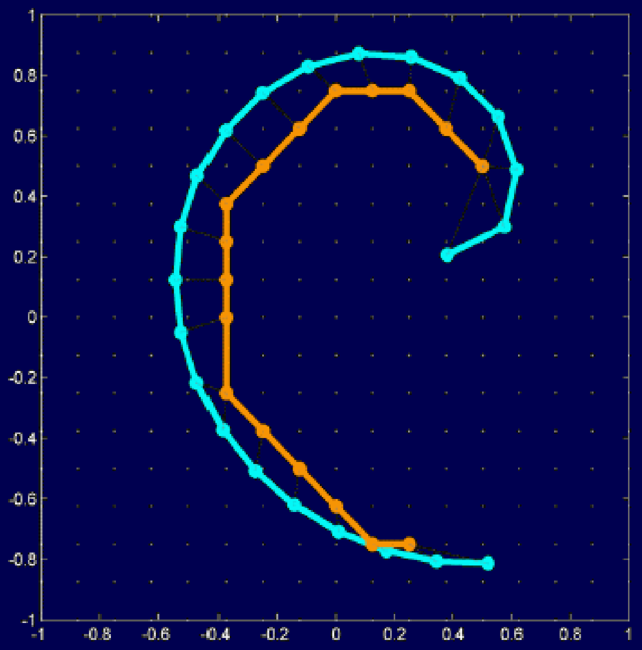

Finite Labeling

8-Neighborhood

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

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

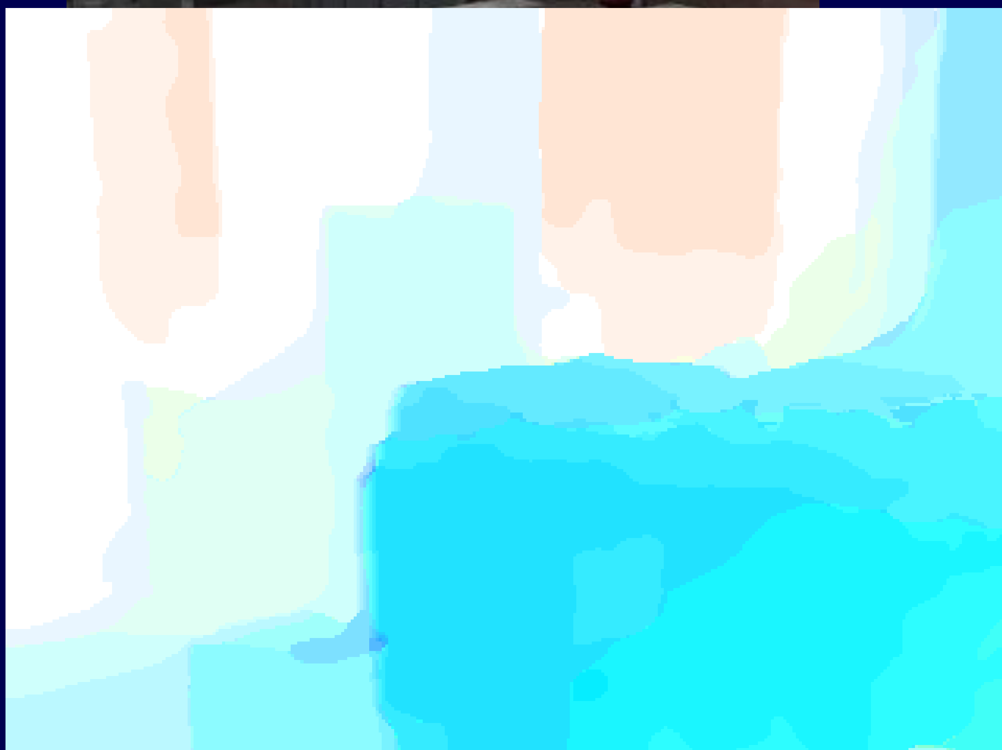

flow with finite labeling

flow with continuous labeling

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

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

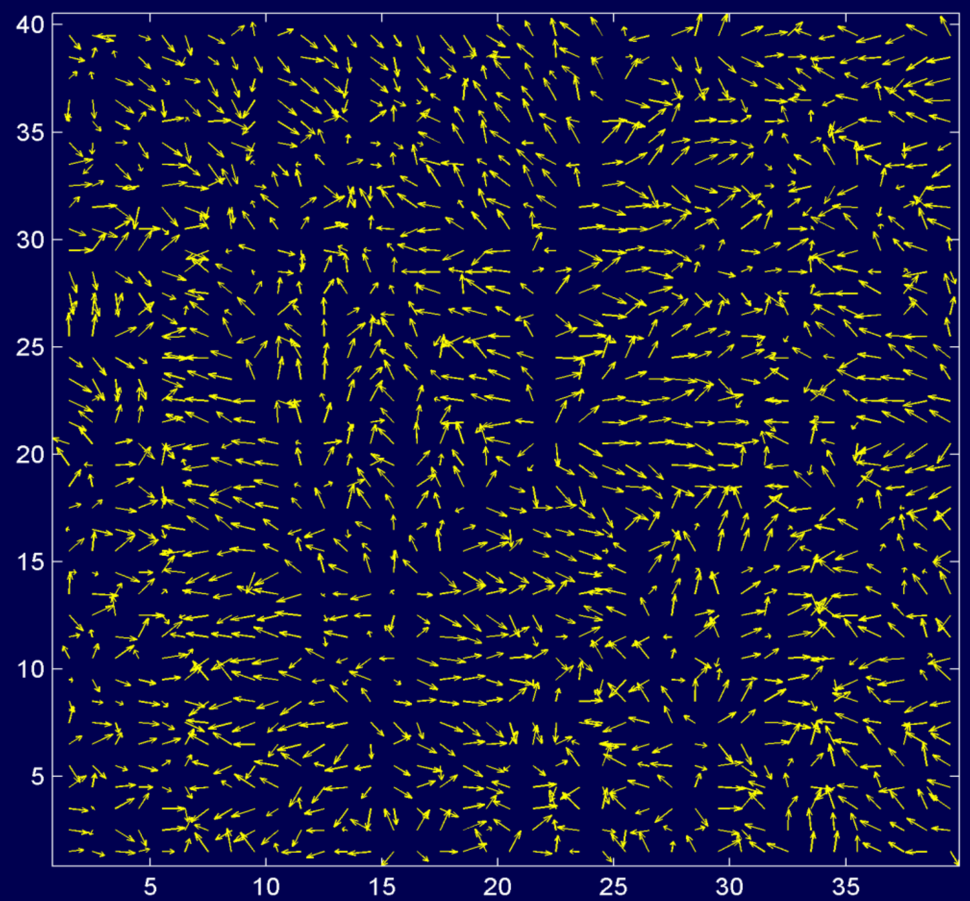

noisy normal field

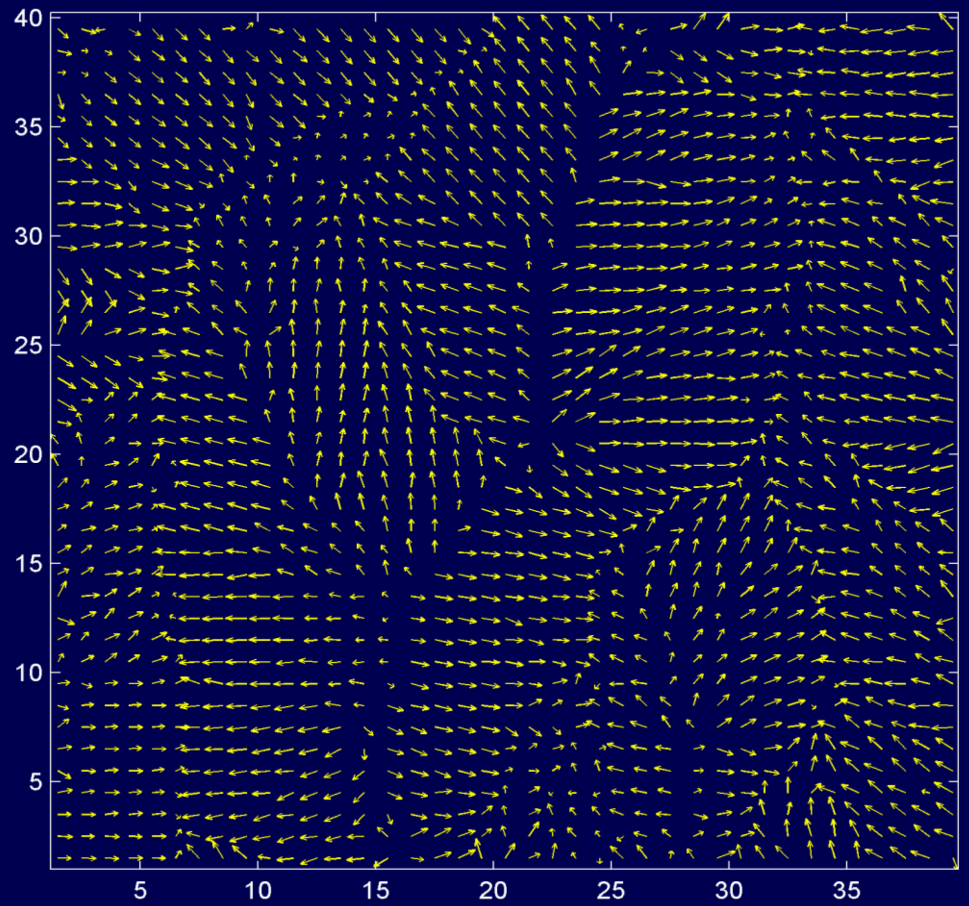

$TV_{\mathcal{S}^2}$ - denoised

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Example: Normal Field Denoising

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

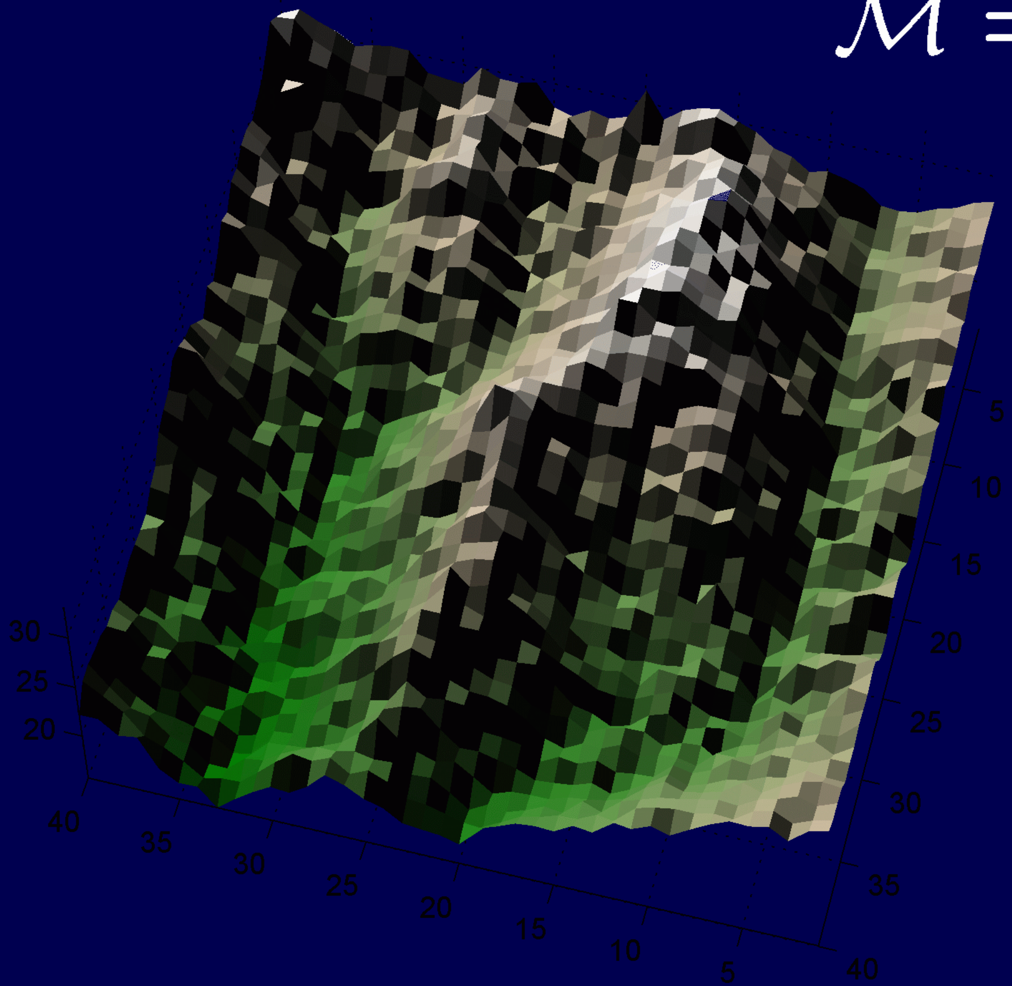

shading with noisy normal field

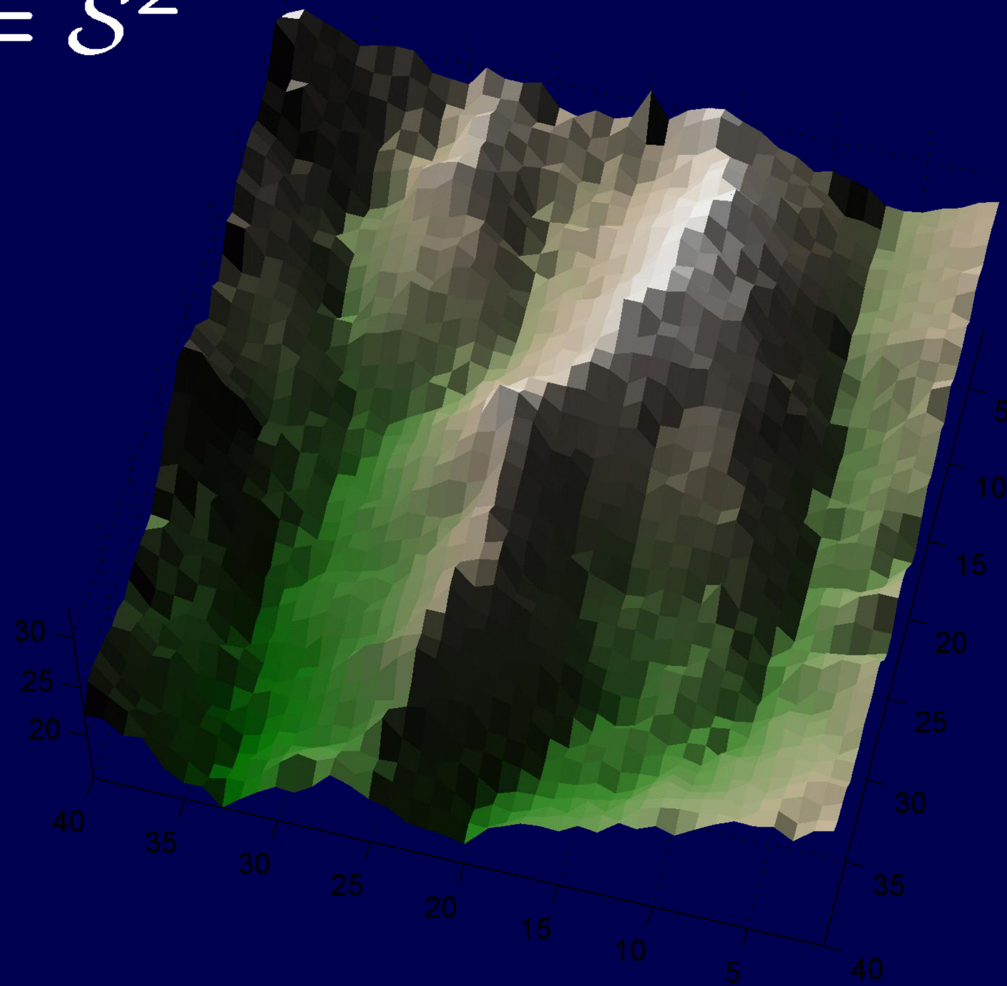

shading with denoised normals

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

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

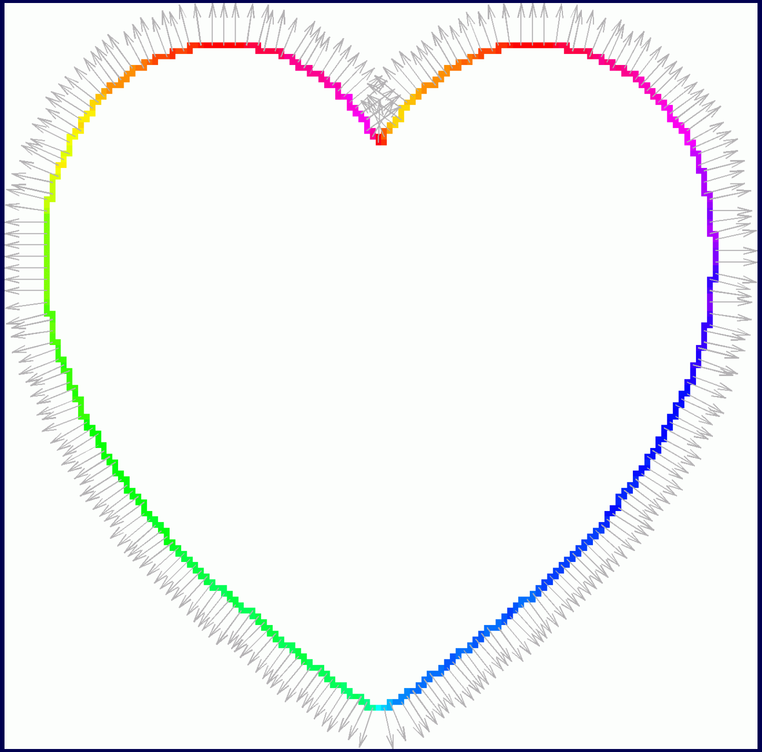

normals on the boundary

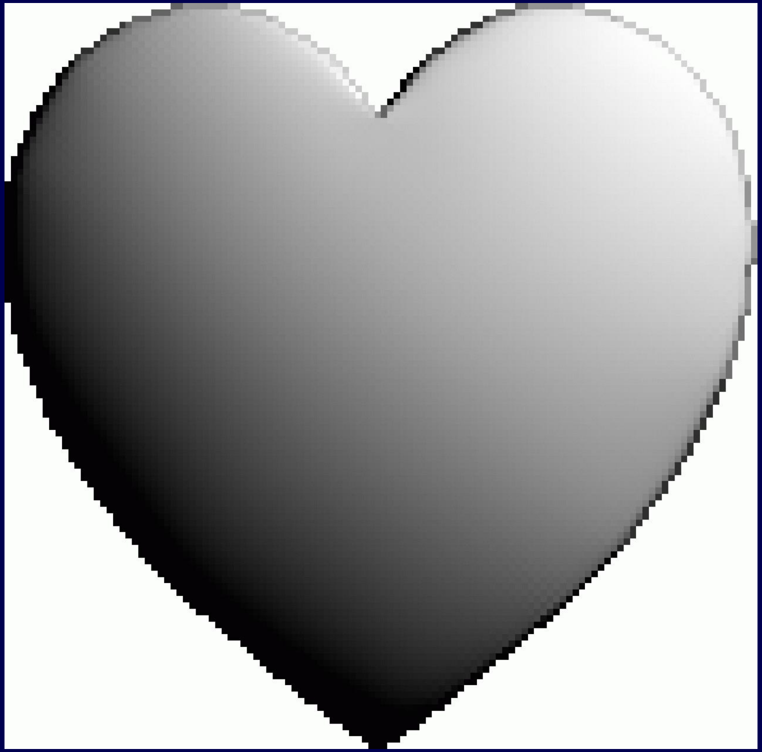

$TV_{\mathcal{S}^2}$ - inpainted normal field

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV '13

Example: Chromaticity Denoising

brightness $I : \Omega \to \mathbb{R}$

chromaticity $C = I/|I| : \Omega \to \mathcal{S}^2$

smoothed chromaticity

smoothed brightness & chromaticity

smoothed brightness

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV ‘13

$$\mathcal{M} = SO(3)$$

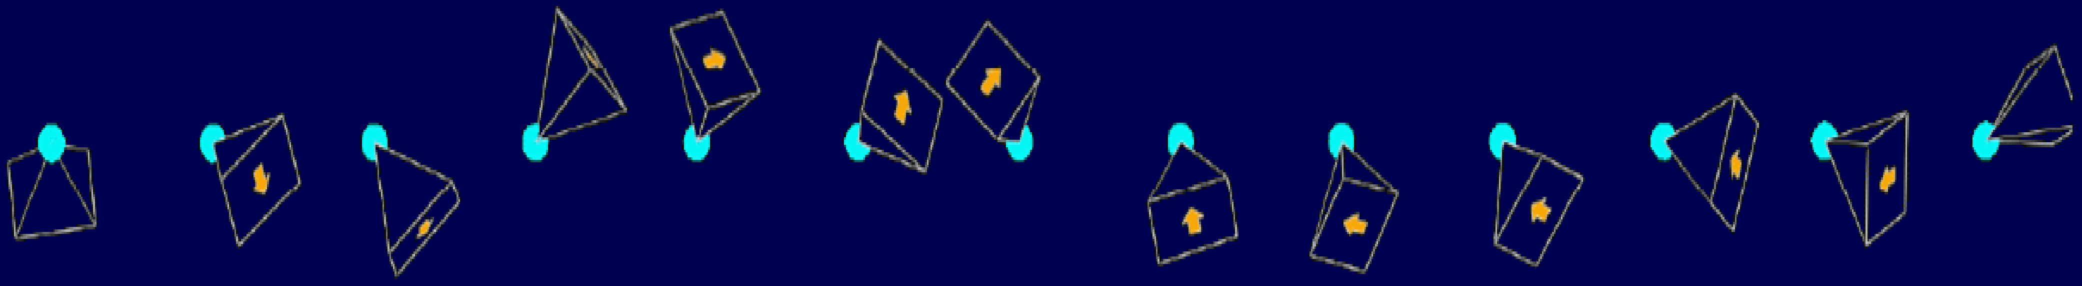

input camera rotation

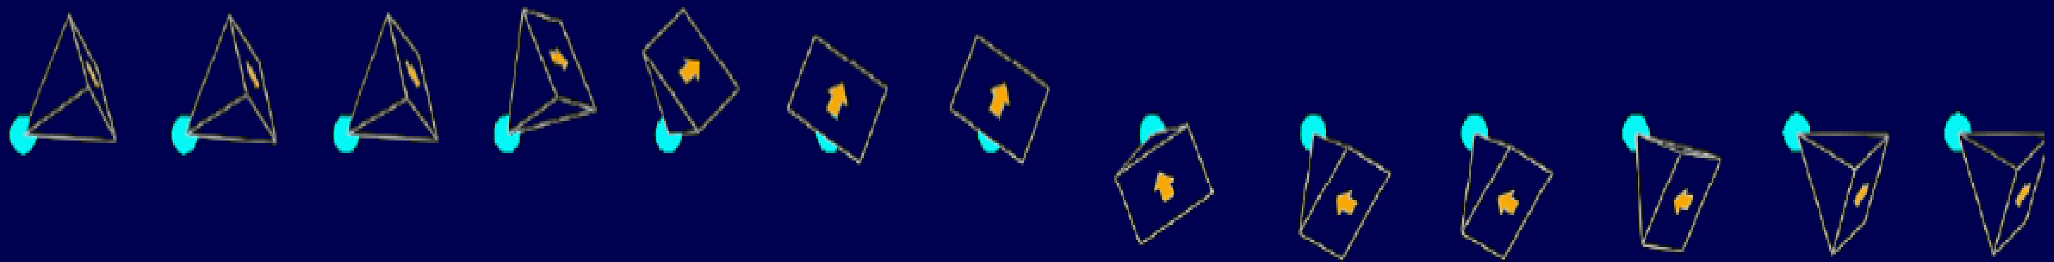

$SO(3)$-denoised camera rotation

Lellmann, Strekalovskiy, Kötter, Cremers, ICCV ‘13

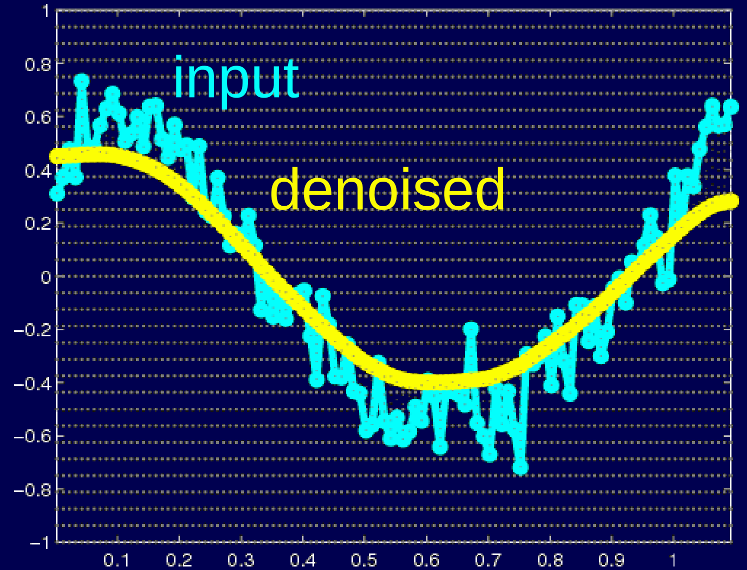

Quadratic regularizer

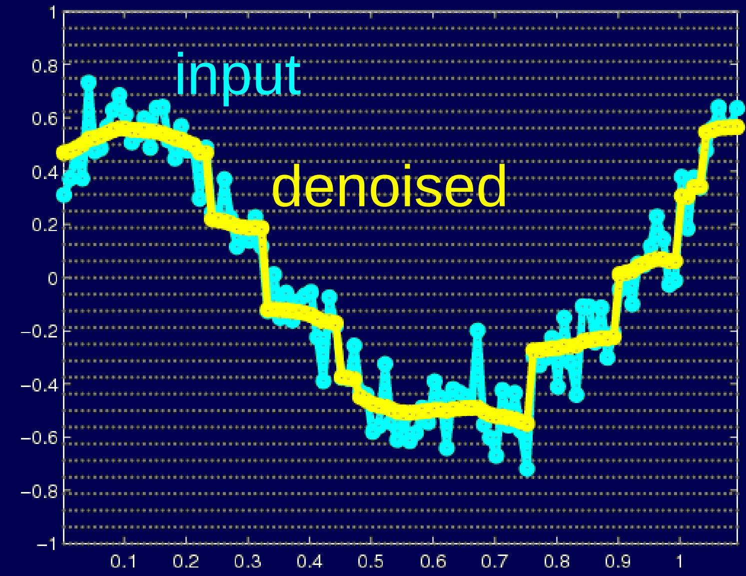

Mumford-Shah

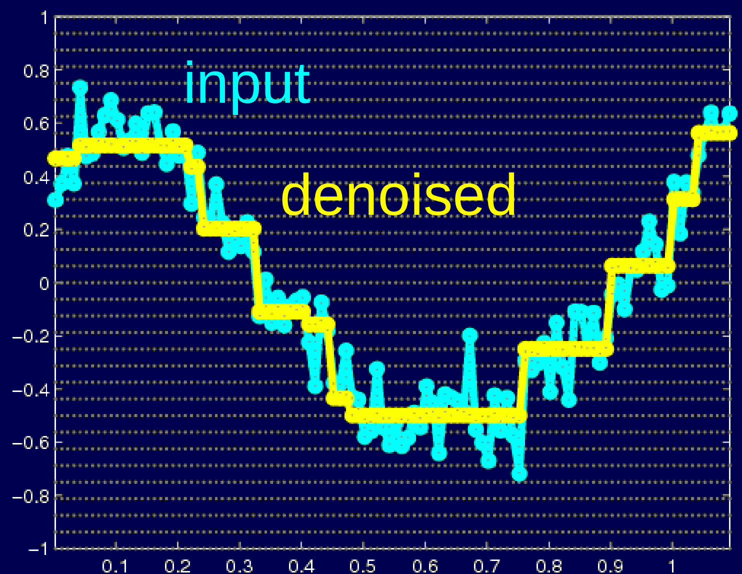

Potts model $\gamma = 0.05$

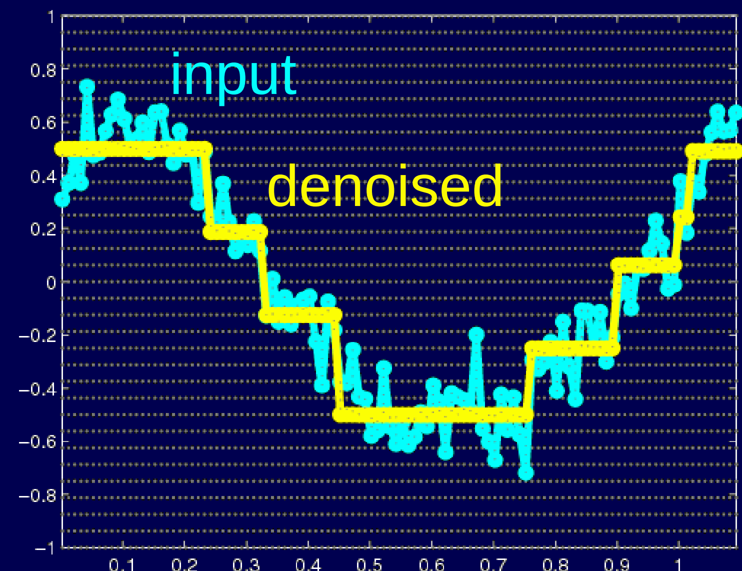

Potts model $\gamma = 0.1$

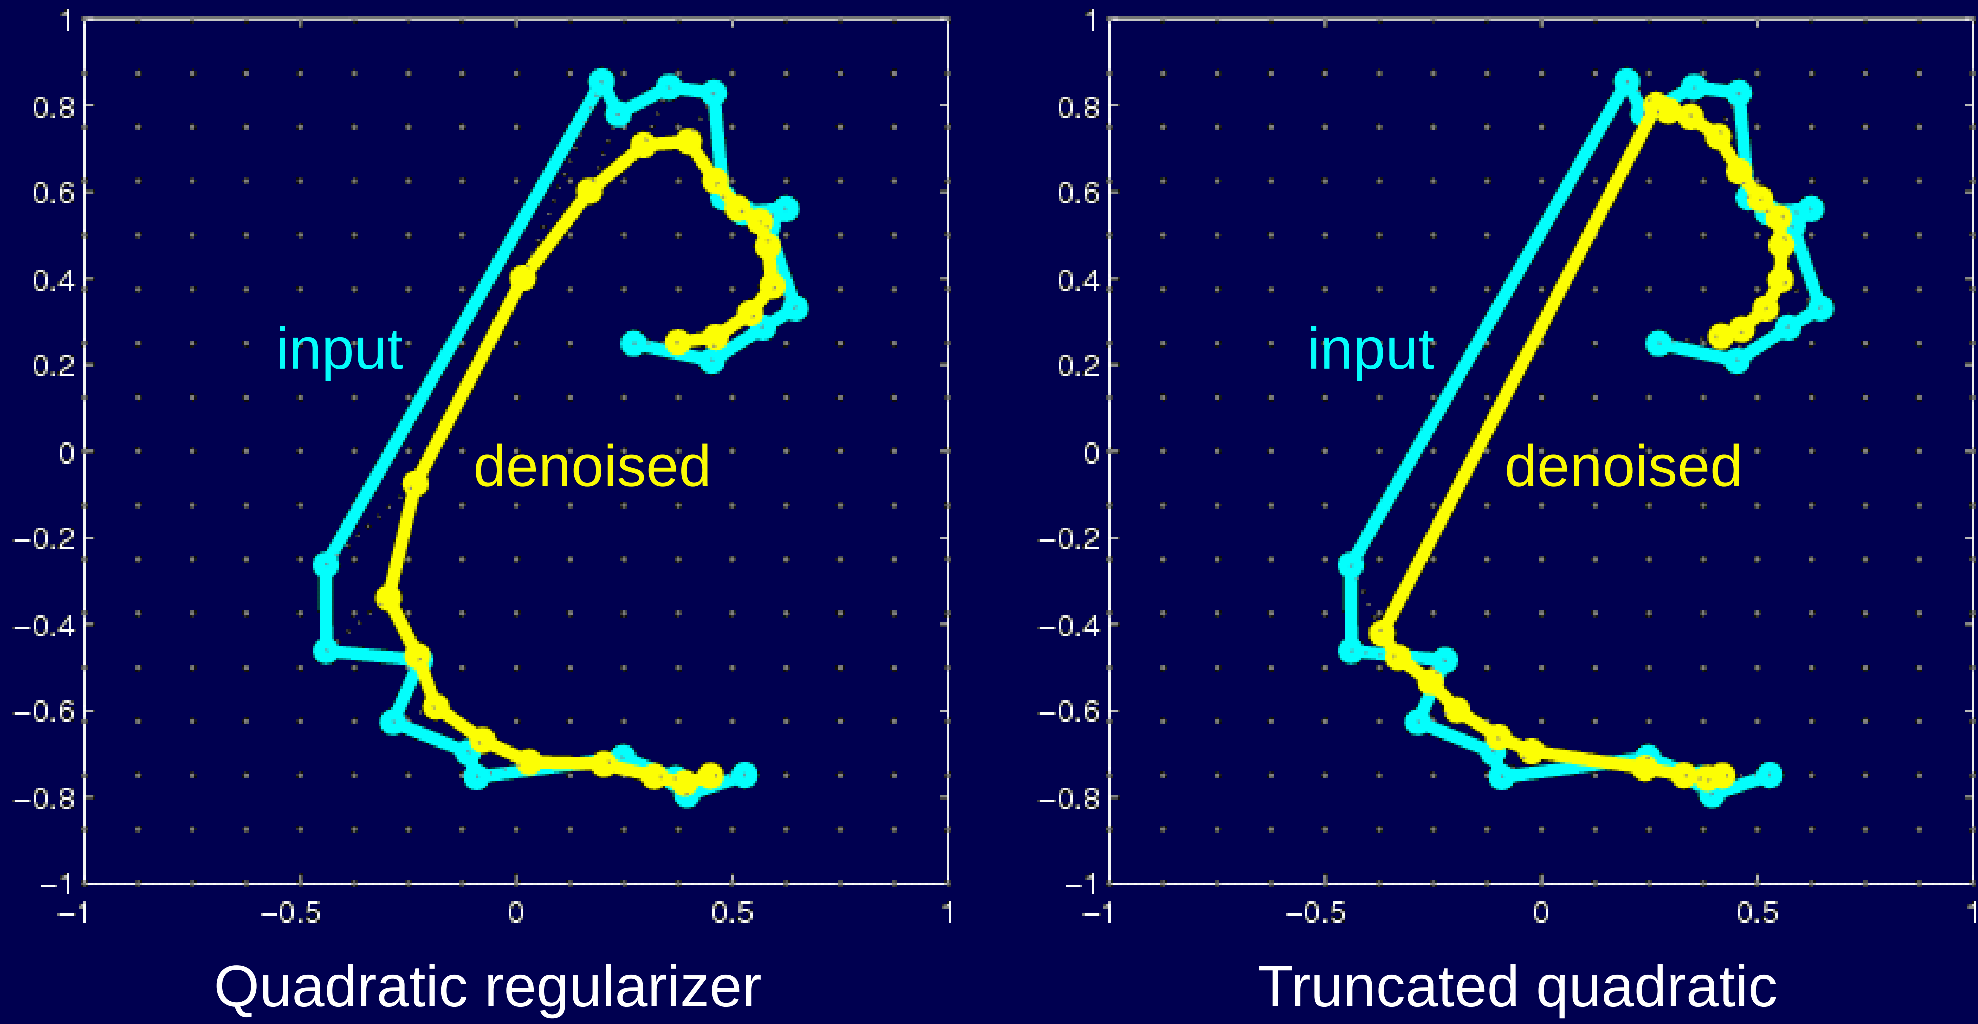
1
0.8
0.6
0.4
0.2
0
−0.2
−0.4
−0.6
−0.8
−1
−1
−0.5
0
0.5
1
input
denoised
Quadratic regularizer
input
denoised
Truncated quadratic

Summary

We proposed a convex relaxation for solving variational problems for functions with values in a Riemannian manifold, including denoising, optical flow or inpainting.

The approach can handle arbitrary Riemannian manifolds, non-convex data terms and a variety of convex and non-convex regularizers.

The continuous labeling approach provides less orientation-bias and grid-bias than existing finite labeling approaches (sublabel accuracy).