Variational models with finite and infinite label spaces

Jan Lellmann

DAMTP, University of Cambridge

Joint work with: F. Lenzen, C. Schnörr (IPA/HCI, University of Heidelberg)
C.-B. Schönlieb (University of Cambridge)
D. Cremers, E. Strekalovskiy (TU Munich)

Rehovot, February 2015

Overview

Finite label spaces

Motivation – Problem

- Finite labeling problem:

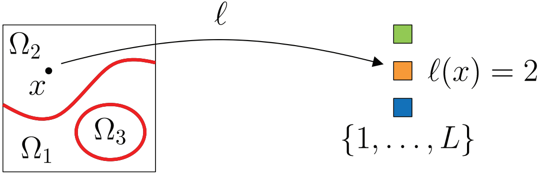

- Partition image domain Ω into L regions
- *Discrete* decision at each point in *continuous* domain Ω
- Variational Approach:

$$\min_{\ell} \underbrace{\int_{\Omega} s(\ell(x), x)dx}_{\text{local data fidelity}} + \underbrace{J(\ell)}_{\text{regularizer}}$$

Motivation – Multiclass Labeling

- **Applications:** Segmentation, denoising, 3D reconstruction, depth from stereo, inpainting, photo montage, optical flow,...

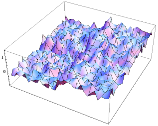
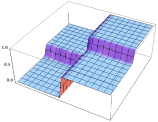


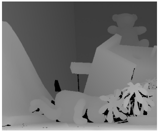

Model – Multi-Class Labeling

- Multi-class relaxation: [Lie et al. 06, Zach et al. 08, Lellmann et al. 09, Pock et al. 09]

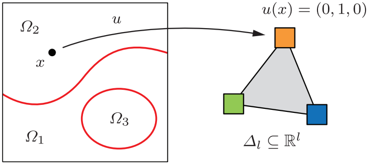

- Embed labels into $\mathbb{R}^L$ as $\mathcal{E} := \{e^1, \ldots, e^L\}$, relax integrality constraint to the unit simplex:

$$\Delta_L := \{x \in \mathbb{R}^L | x \geqslant 0, \sum_i x_i = 1\} = \operatorname{conv} \mathcal{E},$$

$$\min_{u \in \mathrm{BV}(\Omega, \Delta_L)} f(u), \quad f(u) := \int_\Omega \langle u(x), s(x) \rangle dx + \int_\Omega \Psi(Du)$$

- Advantages: No explicit parametrization, rotation invariance, *convex*

Model – Envelope Relaxation

- $J(\ell)$: Weight boundary length by *interaction potential* $d(i, j)$

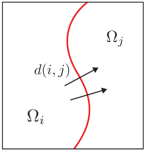



- $J(u)$ implicitly defined as local envelope for given d

[ChambolleCremersPock08,LellmannSchnoerr10]

$$J(u) := \sup_{v \in \mathcal{D}} \int_\Omega \langle u, \operatorname{Div} v \rangle = \int_\Omega \underbrace{\sigma_{\mathcal{D}_{\text{loc}}}(Du)}_{\Psi(Du)},$$

$$\mathcal{D} := \left\{ v \in (C_c^\infty)^{d \times L} \middle| v(x) \in \mathcal{D}_{\text{loc}} \; \forall x \in \Omega \right\},$$

$$\mathcal{D}_{\text{loc}} := \left\{ (v^1, \ldots, v^L) \in \mathbb{R}^{d \times L} \middle| \|v^i - v^j\| \leqslant d(i,j) \; \forall i, j \right\}.$$

Motivation – BV formulation

- Function space formulation avoids metrication artifacts:

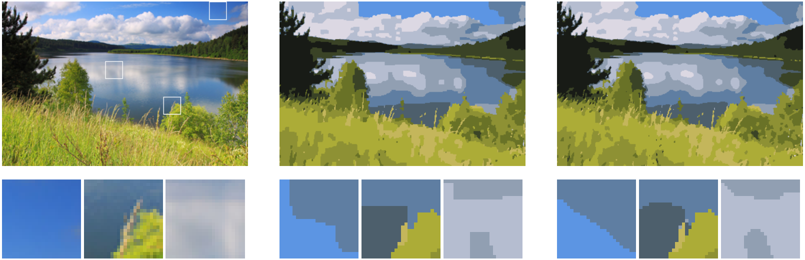

Overview

Optimality

Model – Rounding

- *Fractional* solutions may occur:

- **Goal:** Find *rounding scheme* $u^* \mapsto \bar{u}^* : \Omega \to \{e^1, \ldots, e^L\}$ such that

$$f(\underbrace{\bar{u}^*}_{\text{rounded relaxed solution}}) \leqslant C f(\underbrace{u^*_{\mathcal{E}}}_{\text{best integer solution}}).$$

  for some $C \geqslant 1$.

- 2-class case: $C = 1$ using generalized *coarea formula* (Choquet integral, Lovász extension, levelable function,...)

  [Strang83, ChanEsedogluNikolova06, Zach et al. 09, Olsson et al. 09]

Optimality – Example

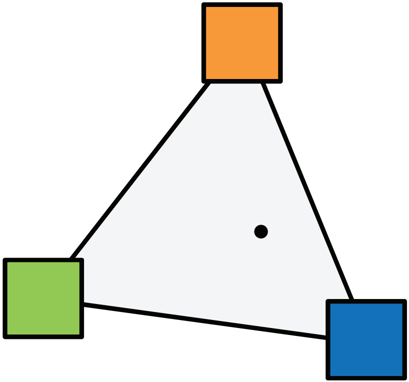

Optimality – Example

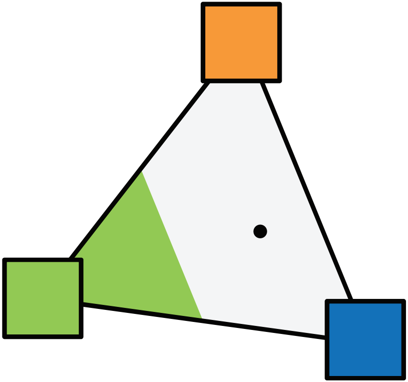

Optimality – Example

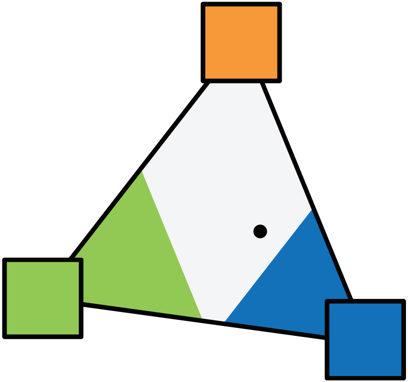

Optimality – Example

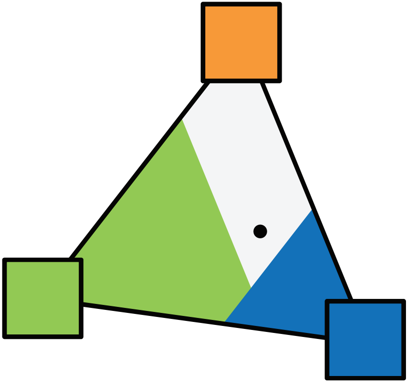

Optimality – Example

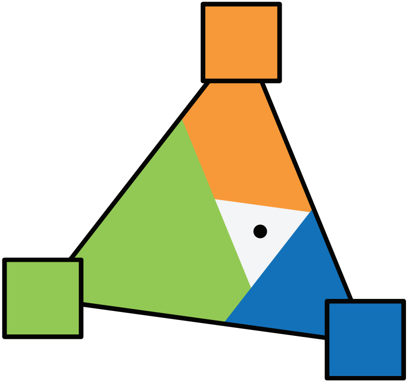

Optimality – Example

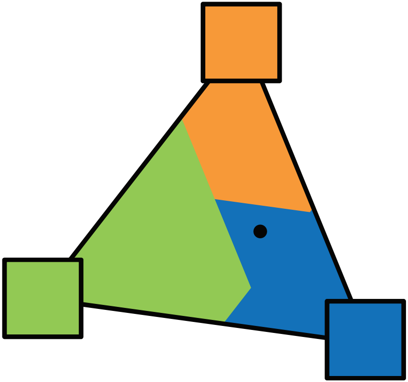

Optimality – Example

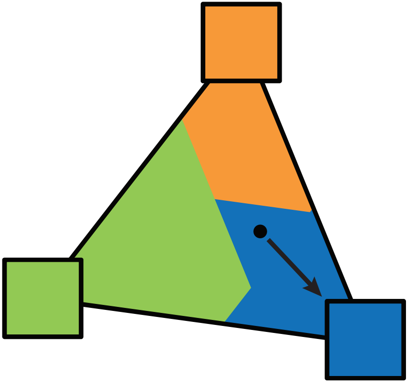

Optimality – Main Result

Theorem (Optimality [LellmannLenzenSchnoerr2011]**)**

Let $u \in \mathrm{BV}(\Omega, \Delta_L)$, $s \in L^\infty(\Omega)^L$, $s \geqslant 0$, d *metric. Then*

$$\mathbb{E}f(\bar{u}) \leqslant 2\frac{\max_{i\neq j} d(i,j)}{\min_{i \neq j} d(i,j)} f(u) \qquad \textit{and} \qquad \mathbb{E}f(\bar{u}^*) \leqslant 2\frac{\max_{i\neq j} d(i,j)}{\min_{i \neq j} d(i,j)} f(u^*_{\mathcal{E}}).$$

- Compatible with bounds for finite-dimensional multiway cut, α-expansion, LP relaxation [Dahlhaus et al. 94, KleinbergTardos 99, Boykov et al. 01, KomodakisTziritas 07]
- Formulated in BV, independent of discretization, true *a priori* bound
- Most likely almost tight: $2(L-1)/L$ lower bound for finite-dimensional case
- In practice much better, usually $C \approx 1.01 \ldots 1.05$

Experiments

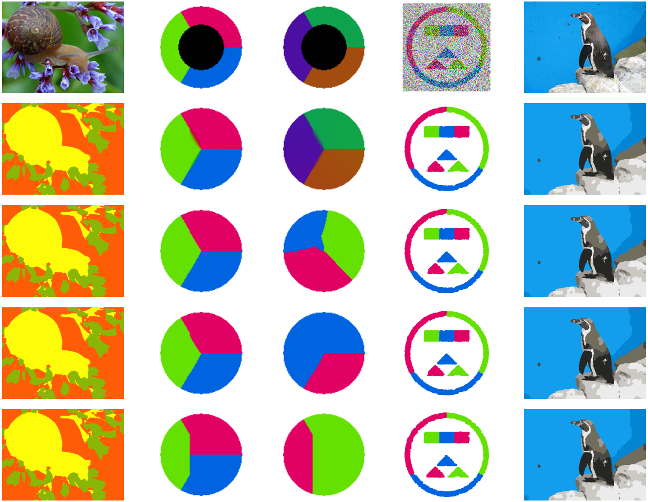

Summary

- **Setting:**
  - Tight convex relaxation of multiclass image labeling for **finite** label spaces
- **Bounds:**
  - *Probabilistic a priori* bound
  - Independent of data term
  - Compatible with finite-dimensional results

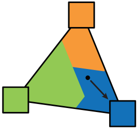

Overview

Infinite label spaces

Application – phase denoising

- Interferometric data has only the *phase* of a distance measurement

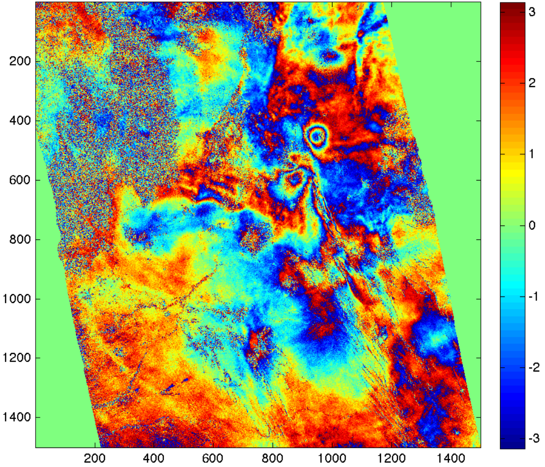

Applications

- S^d — angles, phase data, normals, directions,
- $SO(3)$ — 3-dimensional orientations,
- $\mathcal{T}$ — Torus, Moebius band, Klein bottle, ...
- ...also better results for $\mathbb{R}^n$–valued data (!)

Manifold Constraints

- The **range** of the data u is **constrained to a (Riemannian) manifold** $\mathcal{M}$: [LellmannStrekalovskiyKoetterCremers2013]

$$\min_{u:\Omega\to\mathcal{M}} \int_\Omega f(x,u(x))dx + J(u), \qquad \text{ROF: } \int_\Omega d_{\mathcal{M}}(u(x),I(x))^2\,dx + J(u).$$

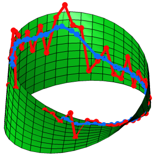

TV-based: Penalise jumps by *geodesic distance*: [Giaquinta,Mucci]

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

- Non-convex due to manifold constraint!

Labeling approach

- Straightforward approach: pick l points on the manifold and treat problem as a *labeling* problem:

$$\min_{u':\Omega\to\Delta_l} \sup_{p:\Omega\to\mathbb{R}^{d\times l}} \int_\Omega \langle u', s\rangle dx + \lambda \int_\Omega \langle u', \operatorname{Div} v\rangle dx,$$

$$\text{s.t. } \|p^{i_1}(x) - p^{i_2}(x)\|_2 \leqslant d(i_1, i_2)\ \forall i_1, i_2 \in \mathcal{J}, \forall x \in \Omega.$$

- Many constraints $\to$ reduce neighbourhood if necessary

Labeling

- ...ROF-denoising of 2D curve $u : [0, 1] \to \mathbb{R}^2$:

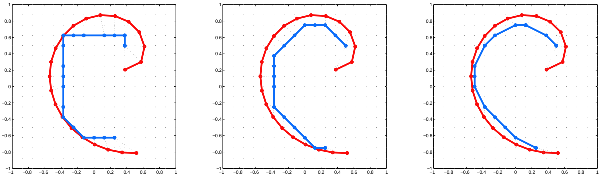

- "Problem": solutions are piecewise constant and biased towards sample points on the manifold...
- Manifold is treated as a set of points.
- We do not use the fact that we have a manifold structure!

Lipschitz constraints

- Labeling formulation with *all* points in the manifold, each $u'(x)$ is a *probability measure* on $\mathcal{M}$: [LellmannStrekalovskiyKoetterCremers2013]

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

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

- The last line is a *Lipschitz constraint*!

$$\ldots \text{s.t. } \|D_z p(x,\cdot)\|_\sigma \leqslant 1, \quad \forall z \in \mathcal{M}, \forall x \in \Omega,$$

$\|\cdot\|_\sigma$ spectral norm.

- Uses structure of the manifold
- *Linear* number of constraints!

Improved formulation

- ROF-denoising of 2D curve:

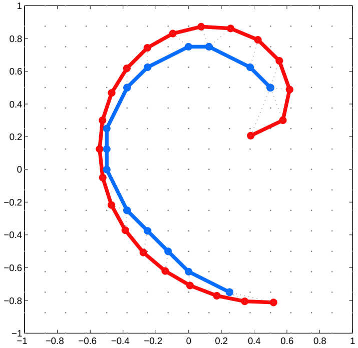

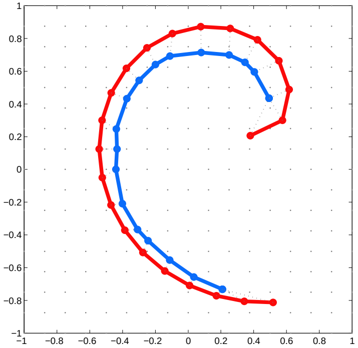

- We get full *sub-label accuracy*!

Applications

- Can easily be adapted to new manifolds such as...

Applications

- The Moebius strip: Rudin-Osher-Fatemi

$$\min_{u:\Omega\to\mathcal{M}} \frac{1}{2}\int_\Omega d_{\mathcal{M}}(u(x), I(x))^2\, dx + \lambda TV_{\mathcal{M}}(u).$$

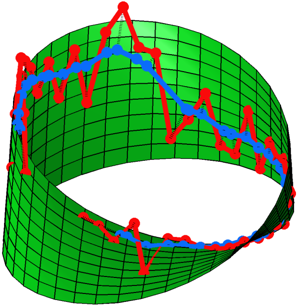

Applications

- Non-convex data terms: Rudin-Osher-Fatemi

$$\min_{u:\Omega\to\mathcal{M}} -\frac{1}{2}\int_\Omega d_{\mathcal{M}}(u(x),I(x))^2\,dx + \lambda TV_{\mathcal{M}}(u).$$

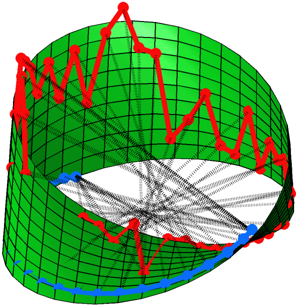

Applications

- ...plain and simple $\mathbb{R}^2$ (and it improves!)...

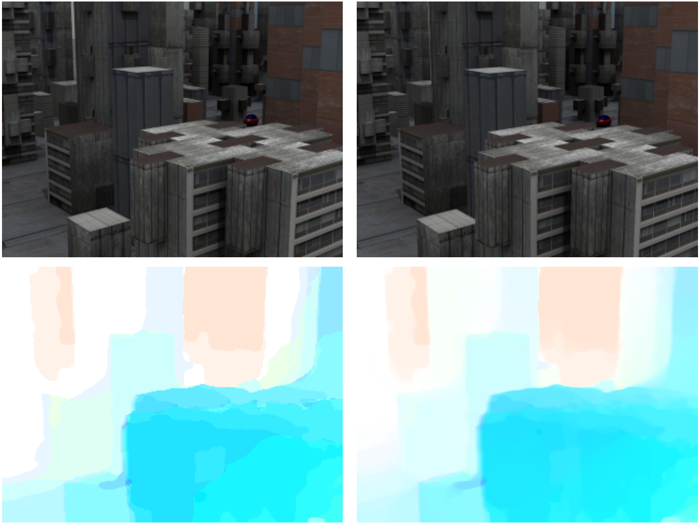

- ...the two-dimensional sphere $\mathcal{S}^2$ (surface normals I)...

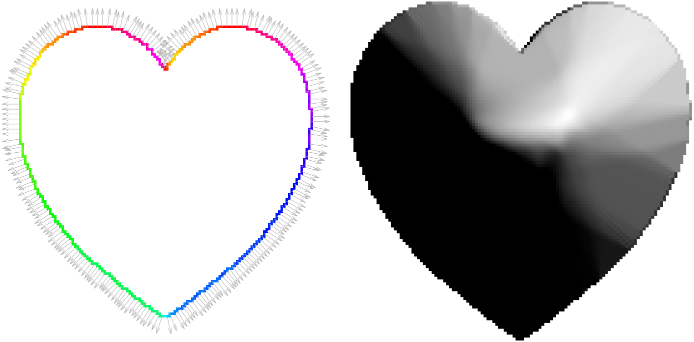

Applications

- ...the two-dimensional sphere $\mathcal{S}^2$ (surface normals II)...

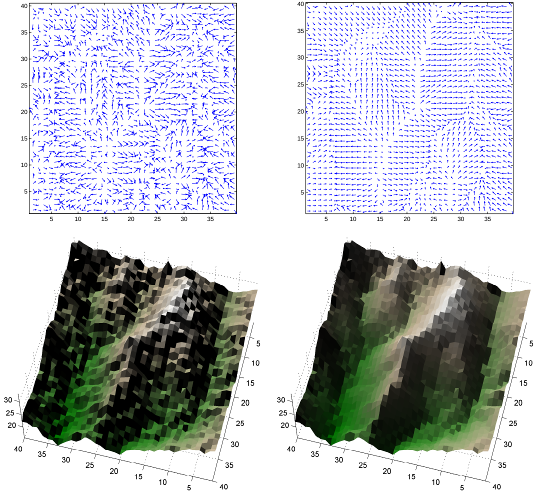

Applications

- ...the space $SO(3)$ of Euclidean rotations, represented as unit quaternions:

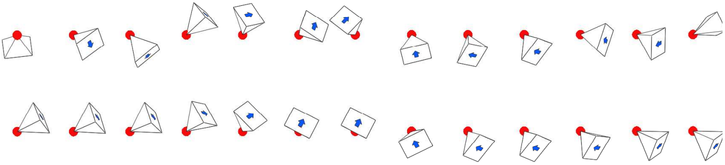

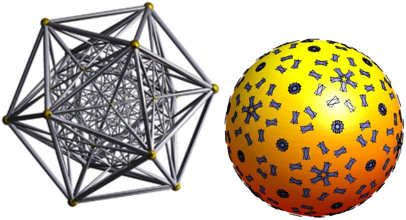

General convex regularization terms

- General convex regularization:

$$\min_{u:\Omega\to\mathcal{M}} D(u) + \int_{\Omega\setminus S_u} h(x,\nabla u)dx + \lambda \int_{S_u} \min\{\gamma, d_{\mathcal{M}}(u^-,u^+)\} d\mathcal{H}^{m-1}.$$

- Approximation: [StrekalovskiyChambolleCremers2012]

$$\min_{u':\Omega\to\mathbb{P}(\mathcal{M})} \max_{\substack{p:\Omega\times\mathcal{M}\to\mathbb{R}^m \\ q:\Omega\times\mathcal{M}\to\mathbb{R}}} \int_\Omega \langle u', s\rangle dx + \lambda \int_\Omega \langle u', \operatorname{Div} p - q\rangle dx$$

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

$$q(x,z) \geqslant h^*(x, D_z p(x,z)), \quad \forall z \in \mathcal{M}, \forall x \in \Omega,$$

General convex regularization terms

- TV

$$\min_{u:\Omega\to\mathcal{M}} D(u) + \lambda \int_{\Omega\setminus S_u} \|\nabla u(x)\| dx + \int_{S_u} \lambda d_{\mathcal{M}}(u^-, u^+) d\mathcal{H}^{m-1}$$

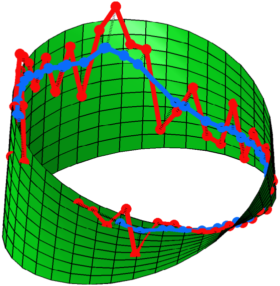

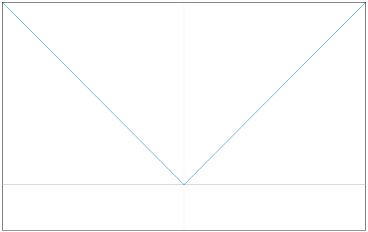

General convex regularization terms

- Truncated TV:

$$\min_{u:\Omega\to\mathcal{M}} D(u)+\lambda\int_{\Omega\setminus S_u}\|\nabla u(x)\|dx+\int_{S_u}\min\{\gamma,\lambda d_{\mathcal{M}}(u^-,u^+)\}d\mathcal{H}^{m-1}$$

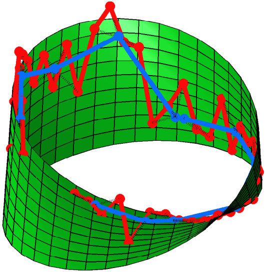

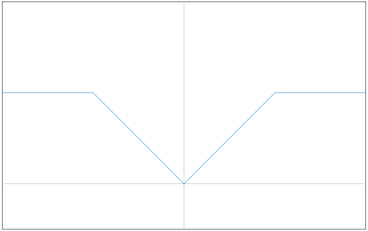

General convex regularization terms

- Quadratic (smoothness, less robust):

$$\min_{u:\Omega\to\mathcal{M}} D(u) + \lambda \int_\Omega \|\nabla u(x)\|^2 dx$$

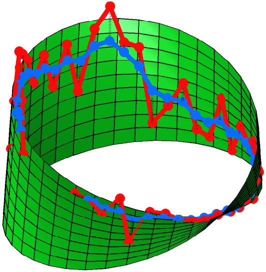

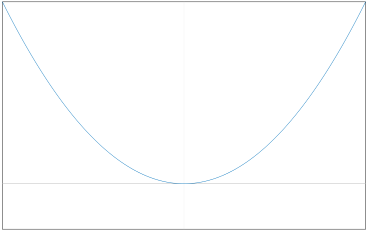

General convex regularization terms

- Approximation of Mumford-Shah/Truncated Quadratic:

$$\min_{u:\Omega\to\mathcal{M}} D(u) + \lambda \int_{\Omega\setminus S_u} \|\nabla u(x)\|^2 dx + \gamma \mathcal{H}^{m-1}(S_u)$$

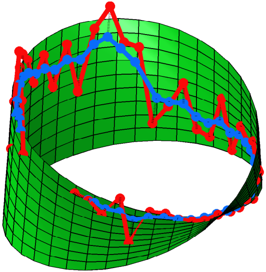

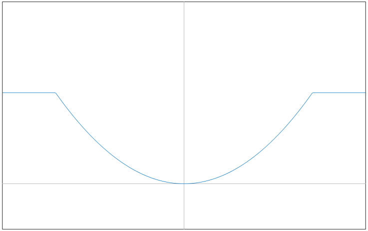

General convex regularization terms

TV on $\mathcal{S}^2$

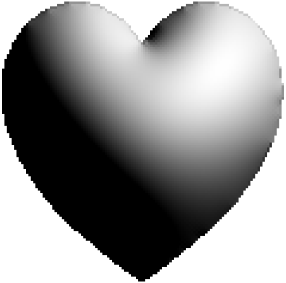

Quadratic on $\mathcal{S}^2$

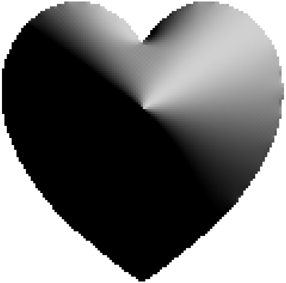

Quadratic on $\mathbb{R}^3$
$+$ normalization

Bregman

- Bregman iteration: convexity splitting

$$\min_u D(u) + 0 = \min_u D(u) + \underbrace{J(u)}_{convex} \underbrace{-J(u)}_{concave}$$

- *Linearize* the concave part:

$$u^{k+1} = \arg\min_u D(u) + J(u) - \left(J(u^k) + \langle v^k, u - u^k \rangle\right), \quad v^k \in \partial J(u^k)$$

- Only adds a linear term. On scalar data, gradually introduces details and converges to the input $\rightarrow$ stop if suitable solution found
  - scalar case: $u(x)$ are *real values*
  - here: $u(x)$ are *probability distributions*

Bregman

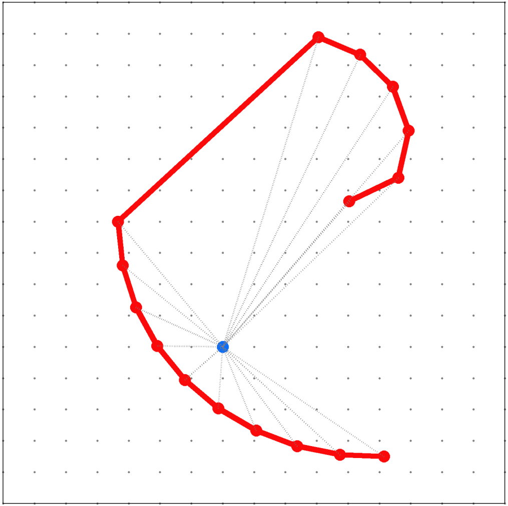

Bregman

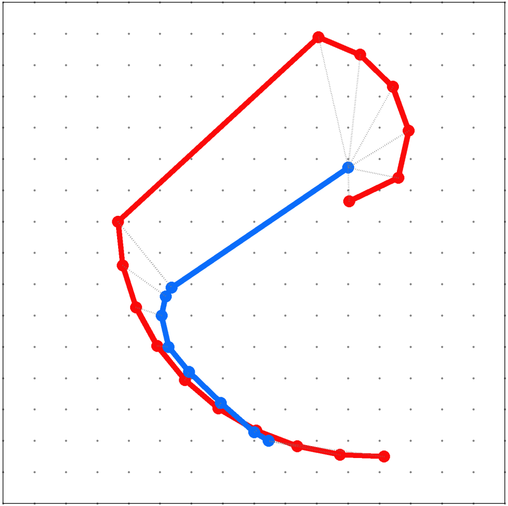

Bregman

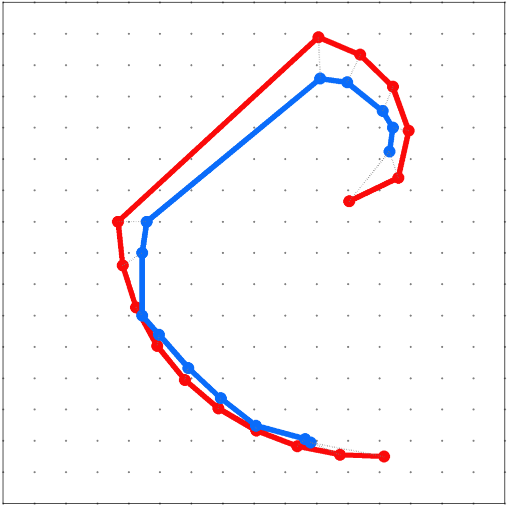

Bregman

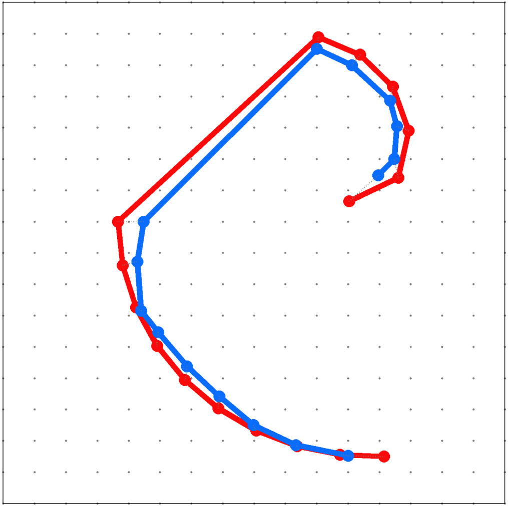

Bregman

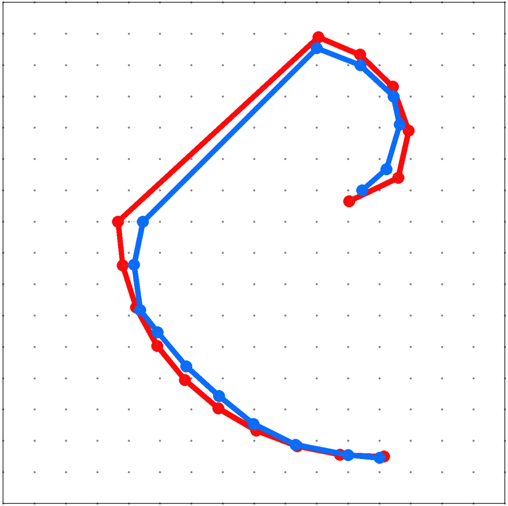

Bregman

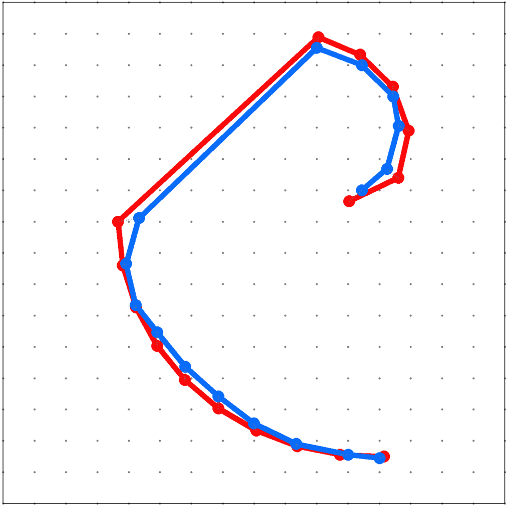

Bregman

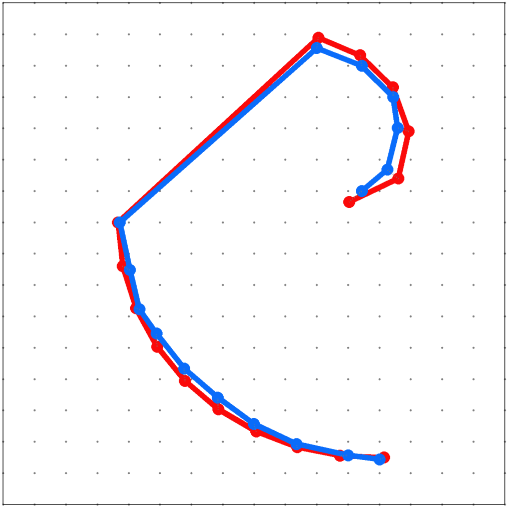

Bregman

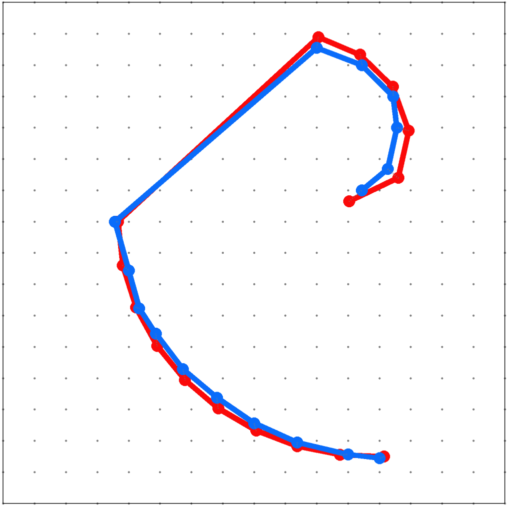

Bregman

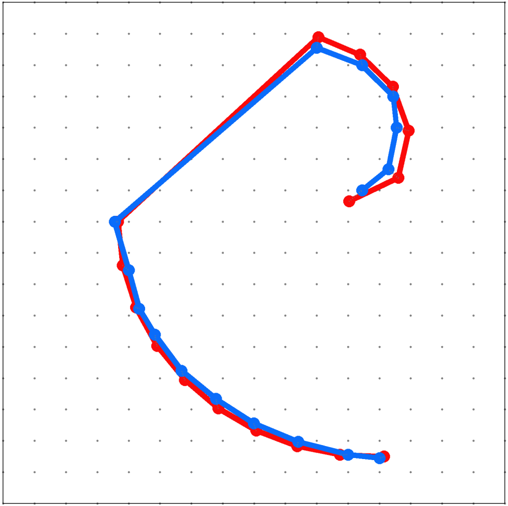

Bregman

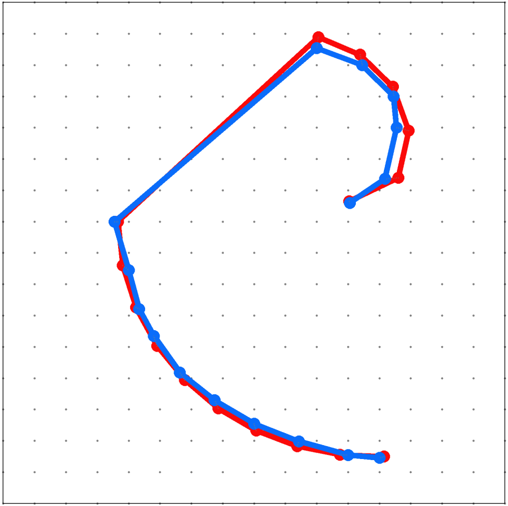

Conclusion II

- **Variational problems with infinite label spaces**
  - For values in a manifold: Sub-label accuracy, computationally more efficient
  - Easy to extend to arbitrary Riemannian manifolds
  - Many properties appear to transfer from the scalar case
  - Code available