A Mathematical Framework for Variational Models of Measure-Valued Data

MIC Institute of Mathematics and Image Computing
Universität zu Lübeck

Thomas Vogt
Joint work with Jan Lellmann
Universität zu Lübeck, Germany

Overview

Goal: Restoration and denoising of measure-valued images.

- Image data: $f\colon \Omega \to \mathcal{P}(X)$, $\Omega \subset \mathbb{R}^d$.
- X: Compact metric space (e.g., sphere $\mathbb{S}^2$).
- $\mathcal{P}(X)$: Probability measures on X.

Proposed variational model for restoration:

$$\inf_{u\colon \Omega \to \mathcal{P}(X)} \int_\Omega W_1(f(x), u(x))\, dx + \lambda \mathrm{TV}_{W_1}(u)$$

- W_1: Wasserstein-1-distance (see below).
- TV_{W_1}: Our proposed total variation seminorm for measure-valued functions (see below).

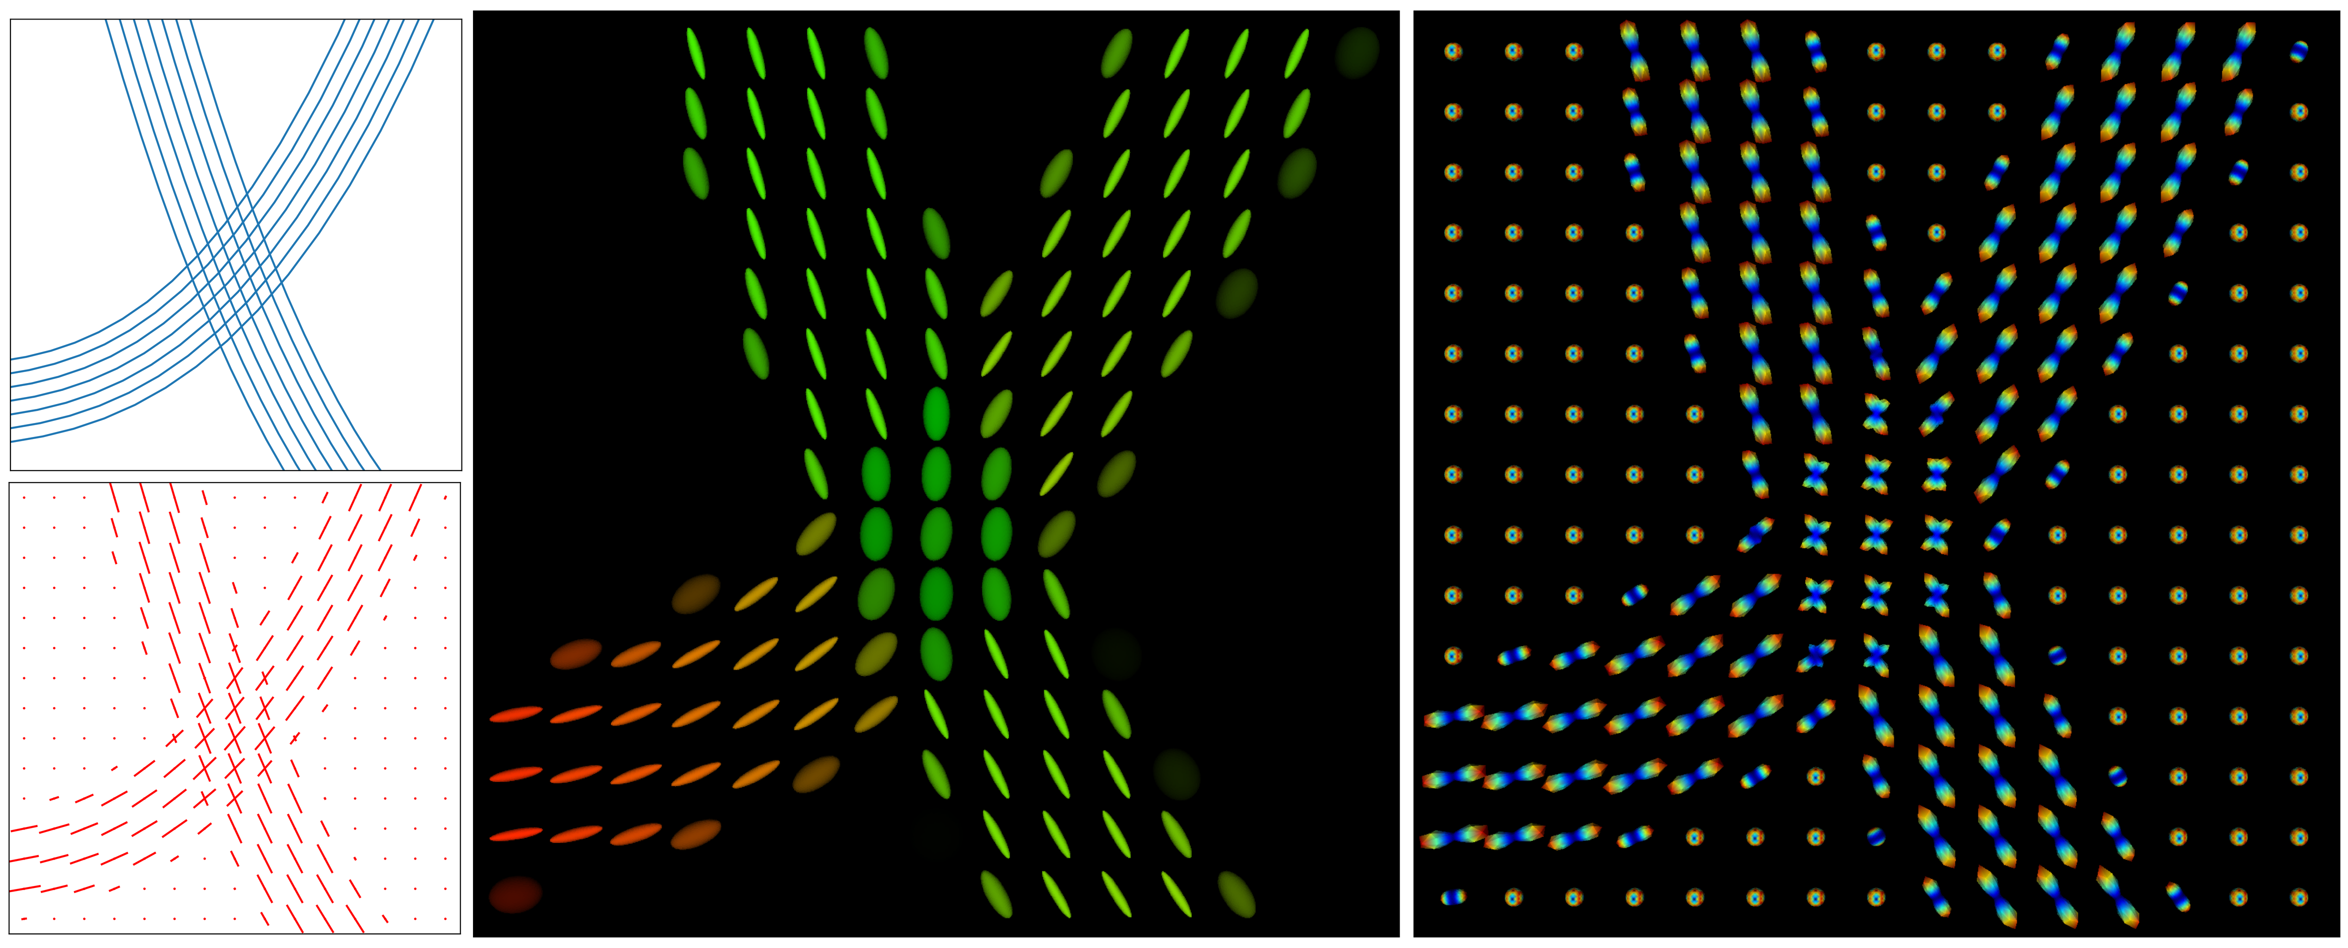

Figure: *Top left:* 2-D fiber phantom. *Bottom left:* Peak directions on a 15×15 grid, derived from the phantom and used for the generation of synthetic HARDI data. *Center:* The diffusion tensor (DTI) reconstruction approximates diffusion directions in a parametric way using tensors, visualized as ellipsoids. *Right:* The ODF reconstruction represents fiber orientation using probability measures at each point, which allows to accurately recover fiber crossings in the center region.

Motivation: Diffusion-Weighted Imaging

In medical applications, diffusion-weighted images (such as HARDI from MRI) describe the local *diffusivity of water*, containing information about the fiber architecture in tissues that exhibit fibrous microstructures (e.g., muscle fibres, axons in cerebral white matter). There exist several interpretations of diffusion-weighted signals:

Diffusion Tensor Imaging (DTI): A 3D-tensor encodes the diffusion behaviour in each voxel. This approach fails to encode the fibrous structure in regions of crossing and branching (see Figure above).

Orientation distribution functions (ODF) are estimated from the data in *Q-Ball Imaging (QBI)* and *Constraint Spherical Deconvolution (CSD)*: In each voxel $x \in \Omega$, the ODF $f(x) \in \mathcal{P}(\mathbb{S}^2)$ encodes the marginal probability of diffusion in a given direction, allowing conclusions to be drawn about crossing and branching of fibers at a scale smaller than the voxel size.

A Metric Space from Optimal Transport

Goal: Measure the distance between two concentrated probability measures by taking the metric distance between the points of concentration.

Definition: Wasserstein-1-metric from optimal transport:

$$W_1(\mu, \mu') := \inf\left\{ \int_{X\times X} d_X(x,y)\, d\gamma(x,y) \colon \gamma \in \mathcal{P}(X \times X), \pi_1\gamma = \mu, \pi_2\gamma = \mu' \right\}$$

- $\pi_i\gamma$: i-th marginal of γ, $\pi_1\gamma(A) := \gamma(A \times X)$, $\pi_2\gamma(A) := \gamma(X \times A)$.
- Kantorovich-Rubinstein duality:

$$W_1(\mu, \mu') = \sup\left\{ \int_X p\, d(\mu - \mu') : |p|_{\mathrm{Lip}(X)} \leq 1 \right\}$$

- $\mathrm{Lip}(X)$: Lipschitz-continuous functions on X.

The dual form allows for an **efficient implementation** of the Wasserstein distance using a first-order primal-dual algorithm.

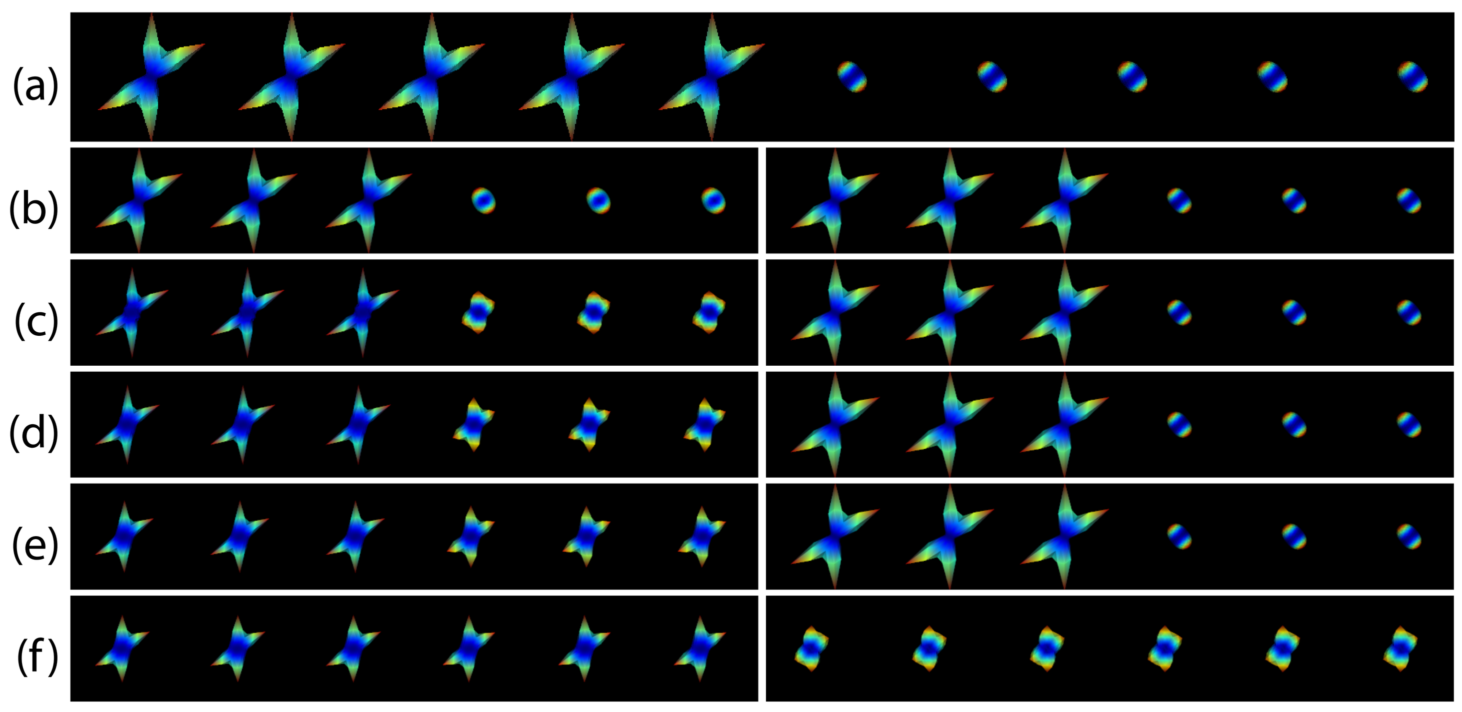

Figure: Synthetic 1D Q-ball image of bimodal and almost uniform ODFs (cf. [Weinmann et al.]): The (a) original data was denoised (b)–(f) using an L^2 (left) and a W_1 data term (right) for increasing values of λ (on the left-hand side $\lambda = 0.05, 0.55, 1.05, 1.55, 2.05$ and on the right-hand side $\lambda = 0.05, 1.35, 2.65, 3.95, 5.25$). Both models preserve the edge. However, as is known from classical ROF models, the L^2 data term produces a gradual transition (contrast loss) towards the constant image, while the W_1 data term exhibits a sudden phase transition.

From Optimal Transport to TV Regularization

Goal: Implement a TV-like regularizer for measure-valued images that penalises jumps according to the Wasserstein-1-distance.

Definition: For $u\colon \Omega \to \mathcal{P}(X)$, let

$$\mathrm{TV}_{W_1}(u) := \sup\left\{ \int_\Omega \langle -\operatorname{div} p(x), u(x)\rangle\, dx \;:\; p \in C_c^1(\Omega; \mathrm{Lip}(X)^d),\ |p(x)|_{\mathrm{Lip}} \leq 1 \right\}.$$

Proposition: Let U be compactly contained in Ω with C^1-boundary ∂U. Let $u^+, u^- \in \mathcal{P}(X)$ and let $u\colon \Omega \to \mathcal{P}(X)$ be defined as

$$u(x) = u^+ \text{ if } x \in U \text{ and } u^- \text{ if } x \in \Omega \setminus U.$$

Then $\mathrm{TV}_{W_1}(u) = \mathcal{H}^{d-1}(\partial U) \cdot W_1(u^+, u^-)$.

Contrary to previous TV-approaches, the proposed formulation has the desired property: The jump from u^+ to u^- is penalised by the Wasserstein distance times the length of the edge.

Comparison: (L2-)ROF Functional

A Rudin-Osher-Fatemi (ROF) model for square-integrable densities $L^2(X) \subset \mathcal{P}(X)$:

$$\inf_{u\colon \Omega \to \mathcal{P}(X)} \int_\Omega \int_X (f(x,z) - u(x,z))^2\, dz\, dx + \lambda \mathrm{TV}_{W_1}(u)$$

Compared with the proposed model (based on a W_1 data fidelity term), this setup exhibits the typical drawbacks of ROF models such as contrast loss across edges.

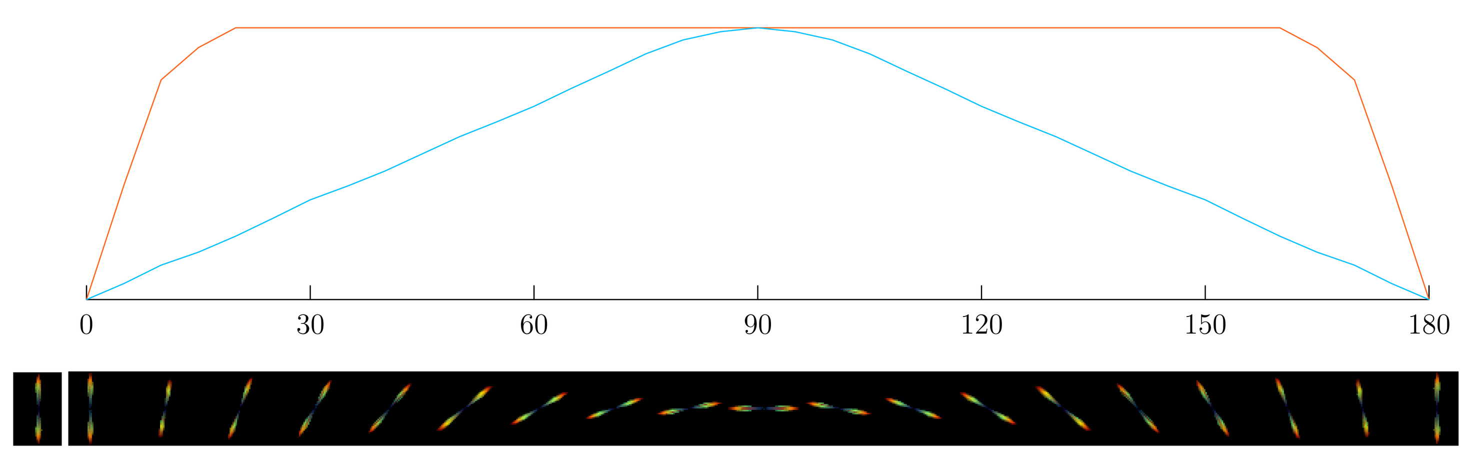

Figure: *Horizontal axis:* Angle of main diffusion direction relative to the reference diffusion profile in the bottom left corner. *Vertical axis:* (Normalized) distances of the ODFs in the bottom row to the reference ODF in the bottom left corner (L^1-distances in orange and W^1-distance in blue). L^1-distances do not reflect the linear change in direction, whereas the W^1-distance exhibits an almost-linear profile. L^p-distances for other values of p (such as $p = 2$) show a behavior similar to L^1-distances.

Literature

- Vogt, T. and Lellmann, J. Measure-Valued Variational Models with Applications to Diffusion-Weighted Imaging. arXiv: *1710.00798* (2017)
- Vogt, T. and Lellmann, J. An Optimal Transport-Based Restoration Method for Q-Ball Imaging. In: *Scale Space and Variational Methods in Computer Vision 2017, Kolding, Denmark*. 271–282 (2017)
- Weinmann, A., Demaret, L., Storath, M. J. Mumford-Shah and Potts Regularization for Manifold-Valued Data. *J Math Imaging Vis 55*. 428 (2016)