Combinatorial Optimization in Computer Vision (IN2245)

Frank R. Schmidt
Csaba Domokos

Winter Semester 2015/2016

Classical Image Segmentation


Given Image

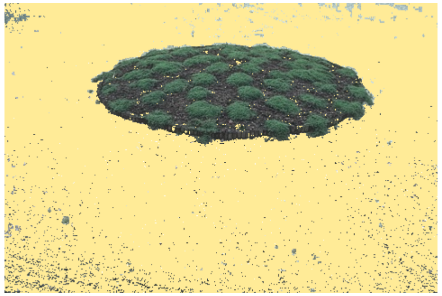
Data Term

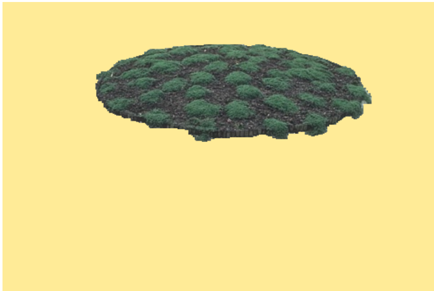
Data + Regularizer

The overall energy we like to minimize is

$$E(x) = \sum_{i=1}^{n} f_i(x_i) + \underbrace{\sum_{i=1}^{n} \sum_{j \in \mathcal{N}(i)} f_{ij} \cdot \delta(x_i, x_j)}_{\text{weighted length}} \quad , x \in \mathbb{B}^n.$$

This submodular energy can be globally minimized via graph-cut if $\delta_{\ell_1,\ell_2} = [\![\ell_1 \neq \ell_2]\!]$.

Medical and RGB Images

Differences to color-images are

- In general, grayscaled images (intensities of magnetic field etc.)
- Boundaries are often difficult to detect
- Color models for *foreground* and *background* tend to look *very similar*
- Interested in multi-label segmentation

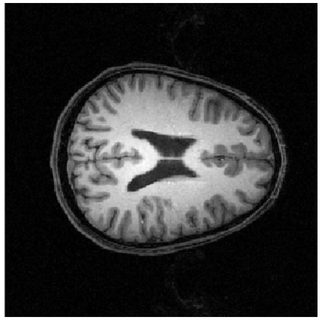

Brain

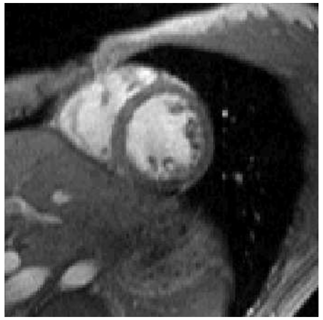

Heart

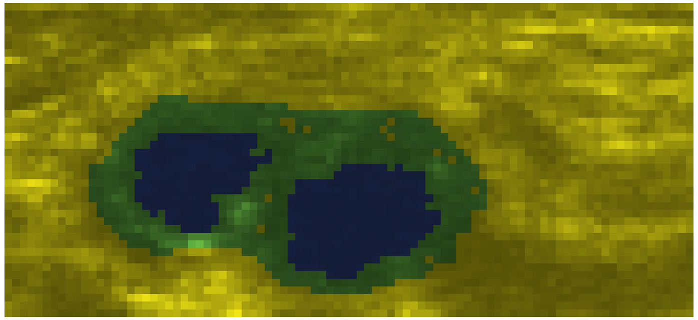

Carotid Artery

Medical Imaging and Multilabeling

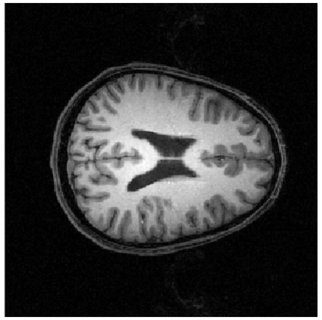

Brain

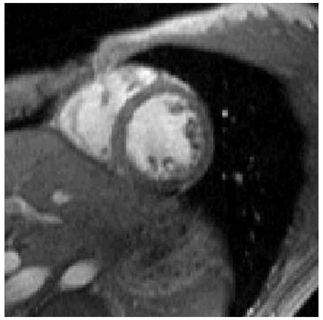

Heart

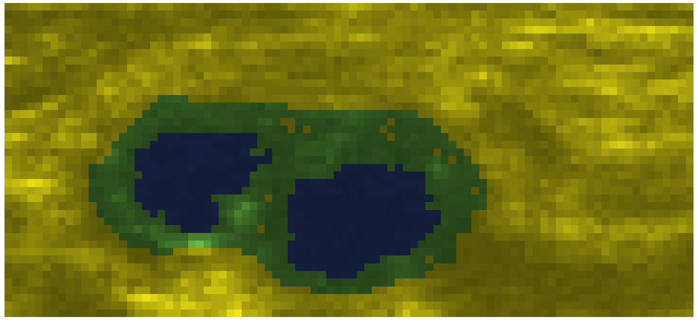

Carotid Artery

$$E(x) = \sum_{i=1}^{n} f_i(x_i) + \sum_{i=1}^{n} \sum_{j \in \mathcal{N}(i)} f_{ij} \cdot \delta(x_i, x_j) \qquad , x \in \{\mathbf{0}, \dots, \mathbf{k}\}^n$$

This problem is NP-hard for the Potts model $\delta(\ell_1, \ell_2) = [\![\ell_1 \neq \ell_2]\!]$ and $k > 1$.

Anatomy vs. Geometry

The term **anatomy** is derived from Greek:
ἀνά "upwards"
τέμνω "to cut"
and refers to the study about the internal structure of organisms.

Anatomy refers to the true nature of the internal structure.
Very often, we cannot detect this true nature.
Instead, we like to enforce certain geometrical properties.

The term **geometry** is also derived from Greek:
γῆ "earth"
μέτρον "measurement"
and refers to the study about **shape** and **relative positions** of objects.

While we are interested in an anatomical model, we usually enforce geometrical constraints.

Geometrical Constraints

If we want to segment a medical obeservation into its components, we can cast this as a multilabeling problem. Enforcing geometrical constraints is equivalent to restricting the set $\mathcal{L}^n$ of feasible labelings.

Given $x \in \mathcal{L}^n$ and $\ell \in \mathcal{L}$, we refer to $S_\ell := \{i | x_i = \ell\}$ as the **region of ℓ**.

Given two different regions S_α and S_β, one migth be interested in the following geometrical constraints:

$$S_\alpha \subset S_\beta \quad \text{(inclusion constraint)}$$
$$S_\alpha \supset\subset S_\beta \quad \text{(exclusion constraint)}$$
$$\text{dist}(S_\alpha, S_\beta) \geqslant d \quad \text{(minimal distance constraint)}$$
$$\text{dist}(S_\alpha, S_\beta) \leqslant d \quad \text{(maximal distance constraint)}$$

There may be different distance functions $\text{dist}(\cdot, \cdot)$ that we can use.

Probabilistic Interpretation

Using Bayes' rule, we can write any multilabeling problem as a probability maximization problem

$$\begin{aligned}\operatorname*{argmax}_{x \in \mathcal{L}^n} P(x|I) &= \operatorname*{argmax}_{x \in \mathcal{L}^n} \frac{P(I|x) \cdot P(x)}{P(I)} \\ &= \operatorname*{argmin}_{x \in \mathcal{L}^n} -\log(P(I|x)) - \log(P(x))\end{aligned}$$

The *likelihood* $P(I|x)$ tells us how well a segmentation $x \in \mathcal{L}^n$ fits to the observed image I. The *prior* $P(x)$ tells us how likely a certain segmentation $x \in \mathcal{L}^n$ is. It does not depend on the observation.

If we want to exclude a certain labeling $x \in \mathcal{L}$, its probability is "0" and we have $-\log(P(x)) = +\infty$.

Multisurface Segmentation

The vision group of Sonka proposed in 2006 a method for medical surface segmentation that bears a certain resemblance to the Ishikawa construction.

They considered **parametrized** surfaces, *i.e.*, mappings $f : \Omega_{xy} \to \mathbb{N}_0$.

Besides the vertical infinity costs, they also added additional infinity costs that assures that the surface does not change too much in the x- and y-direction.

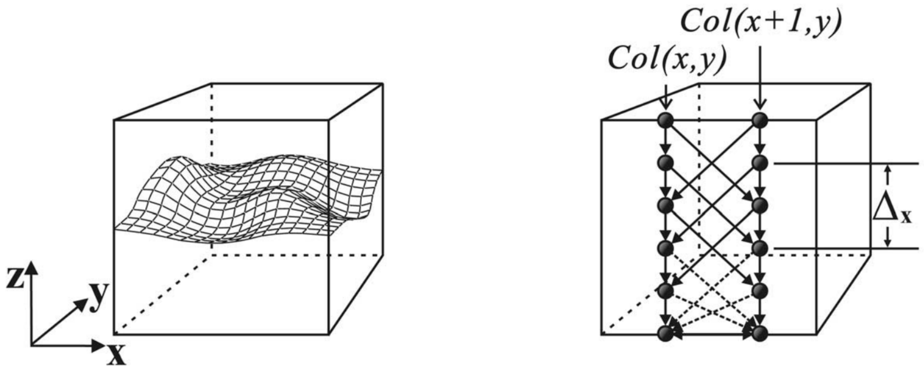

Data & Regularity Term

Since they only wanted to employ a data term $c_{(x,y,z)}$ for the surface that is related to the gradient of the image, they used the following data term

$$f_{(x,y,z)}(1) = \begin{cases} c_{(x,y,z)} & \text{if z=0} \\ c_{(x,y,z)} - c_{(x,y,z-1)} & \text{if z>0.} \end{cases}$$

These data terms proved to result in faster algorithms than the original idea of Ishikawa to use vertical edges beween neighboring layers.

By adding infinity edges

$$[(x,y,z),(x \pm 1, y, z - \Delta_x] \quad \text{and} \quad [(x,y,z),(x \pm 1, y, z + \Delta_x]$$
$$[(x,y,z),(x, y \pm 1, z - \Delta_y] \quad \text{and} \quad [(x,y,z),(x, y \pm 1, z + \Delta_y]$$

the surface can only vary by at least Δ_x or Δ_y in the x- resp. y-direction.

Closed Surface

In order to also handle closed surfaces, they proposed the technique of **unfolding** in order to obtain a surface that can be parametrized.

By doing so, they introduced in fact **cylindrical coordinates**.

A disadvantage of this approach is the explicit knowledge of the central axis of the medical object. Also, the data term has to be transformed into the cylindrical coordinates which may result in small inaccuracies.

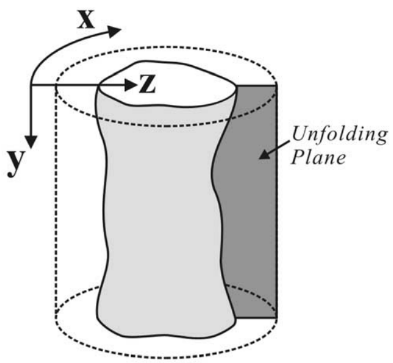

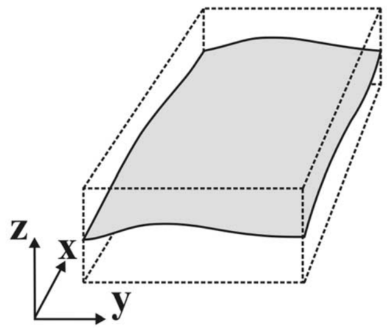

Multiple Surfaces

If one wants to handle multiple surfaces with different data term, the surface construction has to be repeated for each surface. Sonka's group assumed that all surfaces can be parametrized with the same unfolding technique.

Afterwards, additional edges between corresponding columns enforce the minimal and maximal distance constraint. This constraint does not depend on the final segmentation but on the chosen unfolding technique.

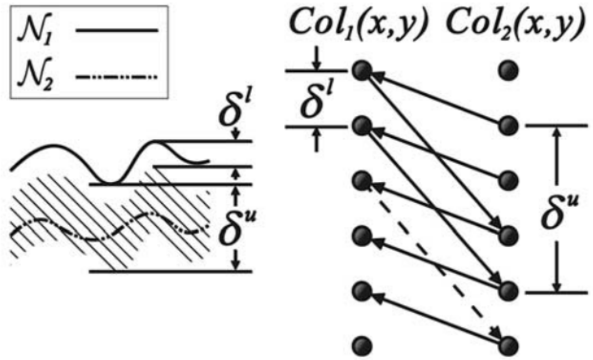

Results

The results are competitive, but the implementation is in general difficult

- The central structure should be easy to detect (pre-processing)
- Gradient information have to be recomputed with respect to the unfolding technique (pre-processing).
- The results are given in cylindrical coordinates. The actual image segmentation has to be derived from it (post-processing).

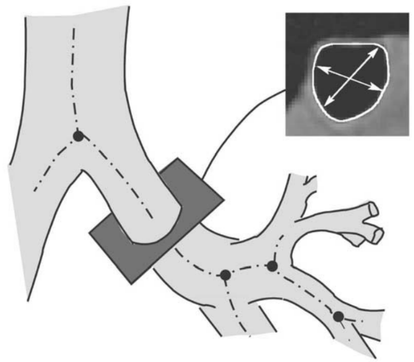

Limitations

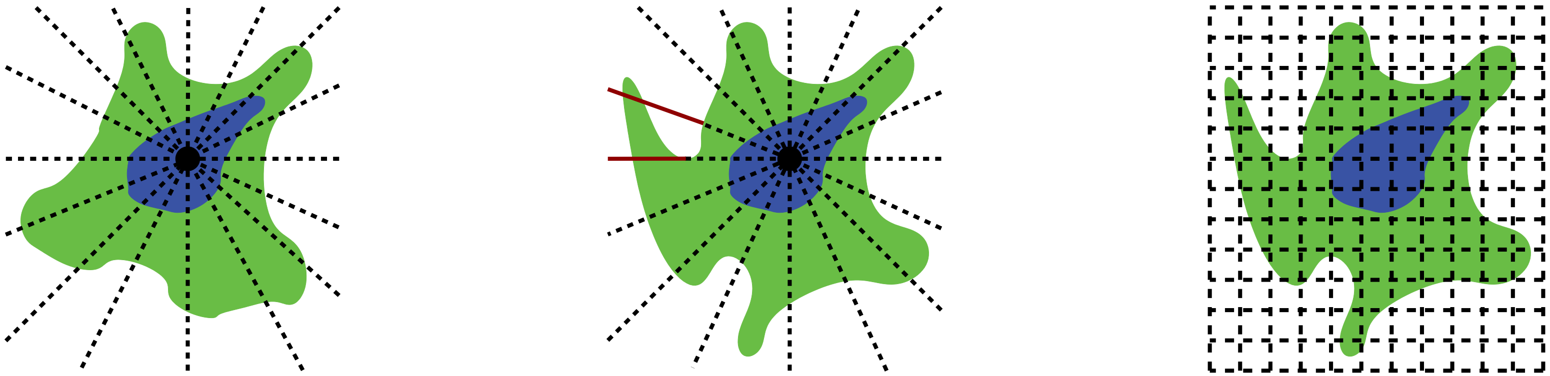

Every star-shaped region can be modelled, but the enforced distance constraints might be limited to the straight lines leaving the center.

Only star-shaped regions can be modelled.
This limits the set of feasible labelings.

We plan to expand the space of possible labelings and enforcing a minimal distance as a well as a maximal distance constraint.

Ishikawa Construction

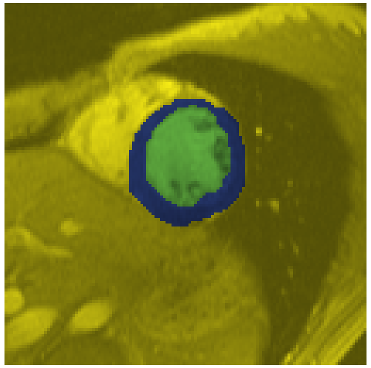

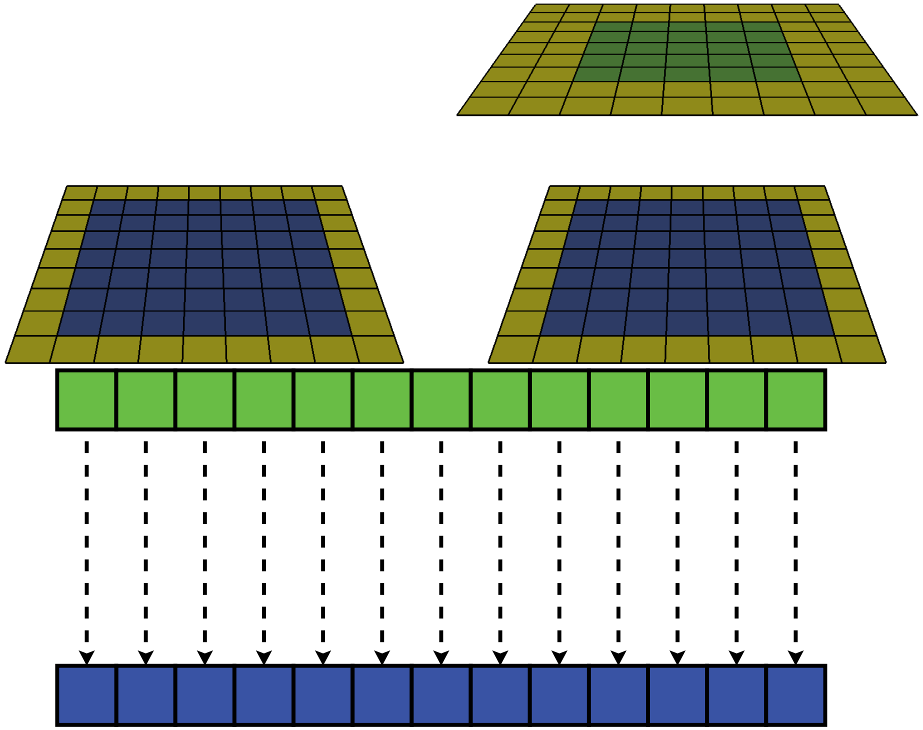

Given the labelspace $\mathcal{L} = \{0, \dots, k\}$, the involved variables are k copies of the n pixels in the image domain, resulting in $n \cdot k$ vertices.

Infinity-edges between these layers assure the **inclusion constraint** $S_{\ell+1} \subset S_\ell$.

Additional edges can encode any penalty $\delta(\ell_1, \ell_2) := d(|\ell_1 - \ell_2|)$
if $d(\cdot)$ is convex.

Minimal Distance Constraint

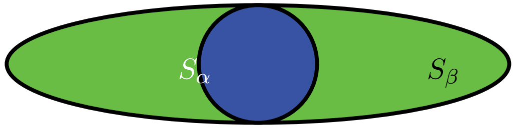

We address the nested multilabeling problem

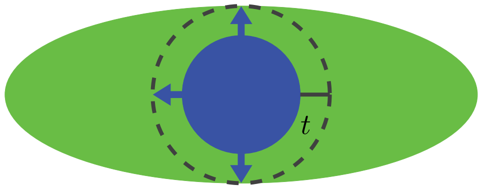

s.t. minimal tubular distance

The **tubular distance** t between $S_\alpha \subset S_\beta$ is the minimal value r such that

$$S_\alpha \oplus B_r \subset S_\beta$$

$$A \oplus B = \bigcup_{b \in B} \{a + b | a \in A\} \qquad B_r = \{x \in \mathbb{R}^d | \, \|x\| \leqslant r\}$$

Nested Labeling

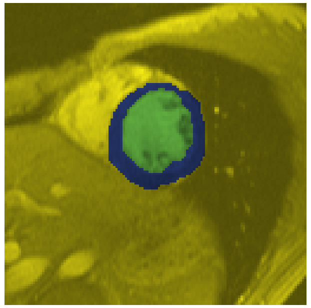

We are interested in the following label penalty

$$\delta(\ell_1, \ell_2) = \begin{cases} 0 & , |\ell_1 - \ell_2| = 0 \\ 1 & , |\ell_1 - \ell_2| = 1 \\ \infty & , \text{otherwise} \end{cases}$$

In practice, ∞ is a large constant K and one can formulate δ as

$$\delta(\ell_1, \ell_2) = \begin{cases} 0 & , |\ell_1 - \ell_2| = 0 \\ 1 + (\Delta - 1) \cdot K & , |\ell_1 - \ell_2| = \Delta \end{cases}$$

This penalty is convex in $|\ell_1 - \ell_2|$.

Nested Labeling (Ishikawa)

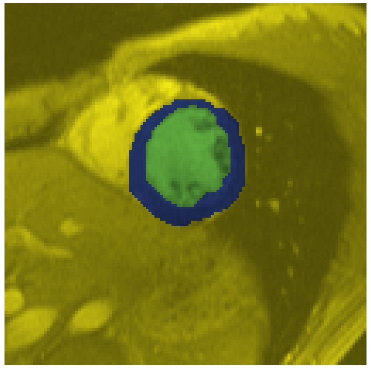

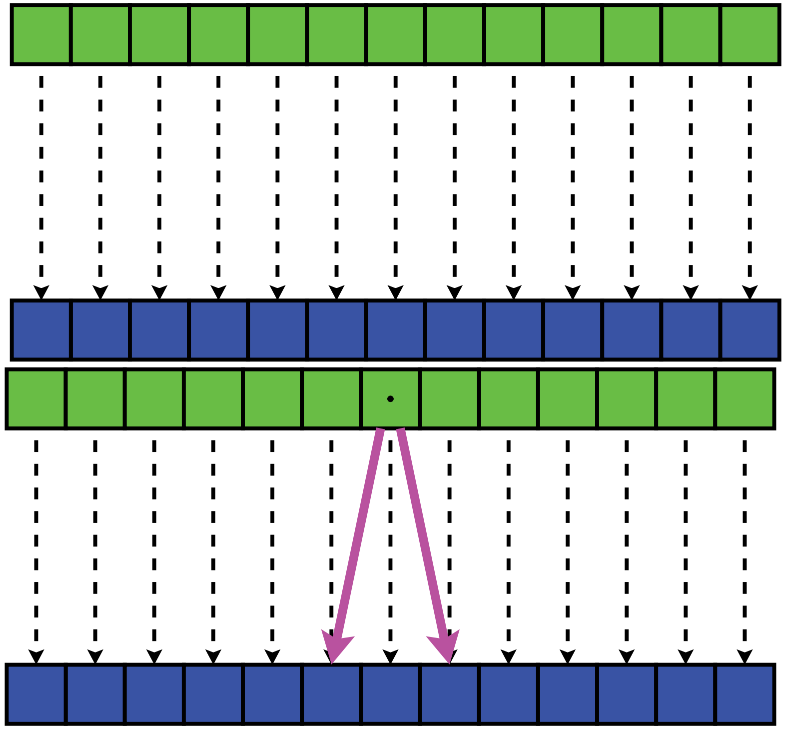

Given the labelspace $\mathcal{L} = \{0, \ldots, k\}$, the involved variables are k copies of the n pixels in the image domain, resulting in $n \cdot k$ vertices.

Infinity-edges between these layers assure the **inclusion constraint** $S_{\ell+1} \subset S_\ell$.

Additional infinity-edges enforce a *tubular distance* of **at least one pixel** between neighboring regions.

Minimal Distance Constraint

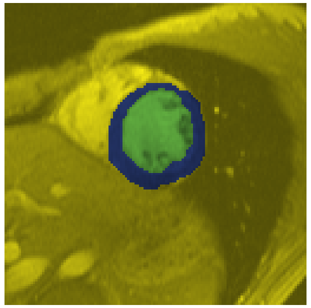

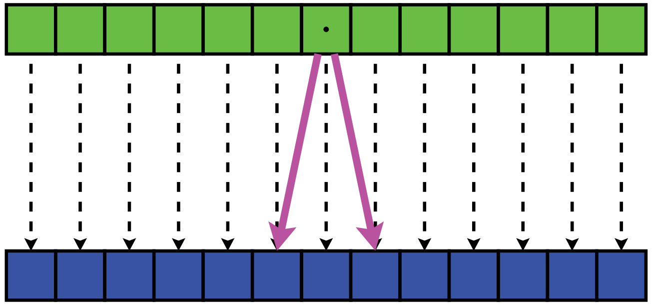

Increasing the neighborhood enforces larger minimal distance constraints.

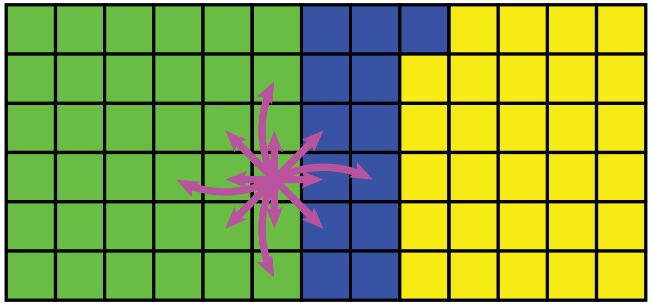

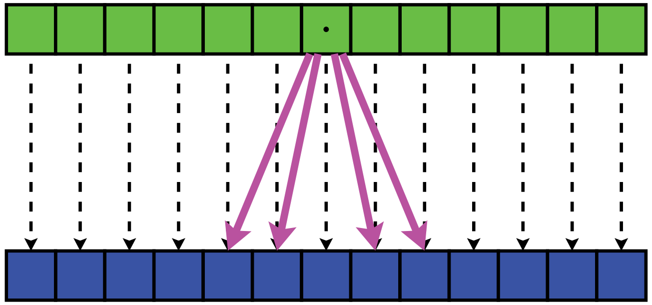

Heart Segmentation

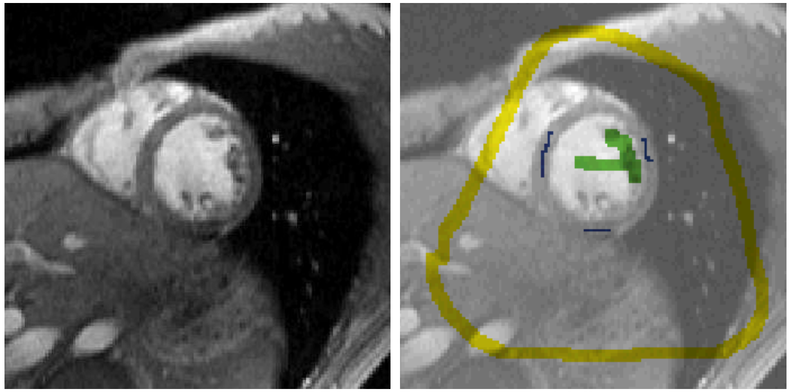
Image Seeds

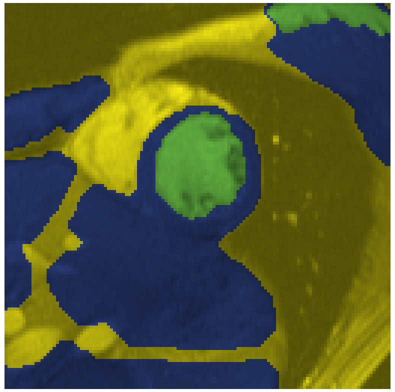
Minimal Distance

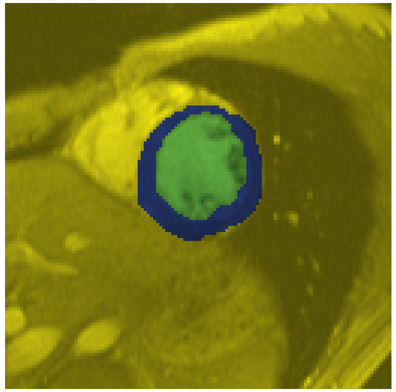
Maximal Distance

A. Delong, Y. Boykov: ICCV, 2009
F. R. Schmidt, Y. Boykov: ECCV, 2012

Literature

Distance Constraints

- Li, Wu, Chen, Sonka, *Optimal Surface Segmentation in Volumetric Images – A Graph-Theoretical Approach*, 2006, IEEE TPAMI 28, 119–134.
- Delong, Boykov, *Globally Optimal Segmentation of Mult-Region Objects*, 2009, IEEE ICCV, 285–292.
- Schmidt, Boykov, *Hausdorff Distance Constraint for Multi-Surface Segmentation*, 2012, ECCV, 598–611.