Combinatorial Optimization in Computer Vision (IN2245)

Frank R. Schmidt
Csaba Domokos

Winter Semester 2015/2016

Medical Imaging and Multilabeling

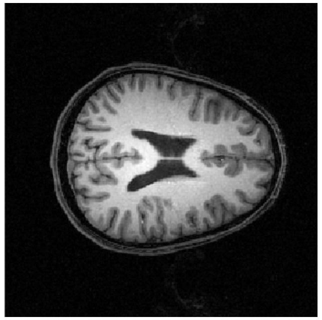
Brain

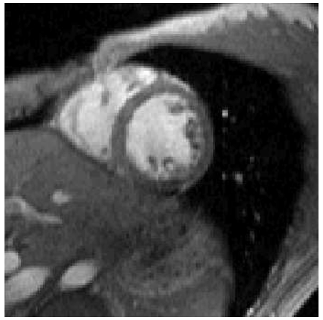
Heart

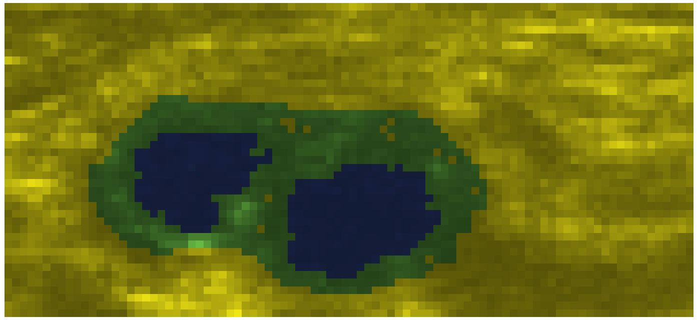
Carotid Artery

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

We like to use $\delta(\ell_1, \ell_2)$ in order to model certain 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:

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

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

Inclusion Constraint

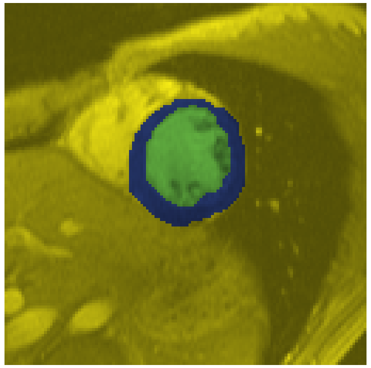

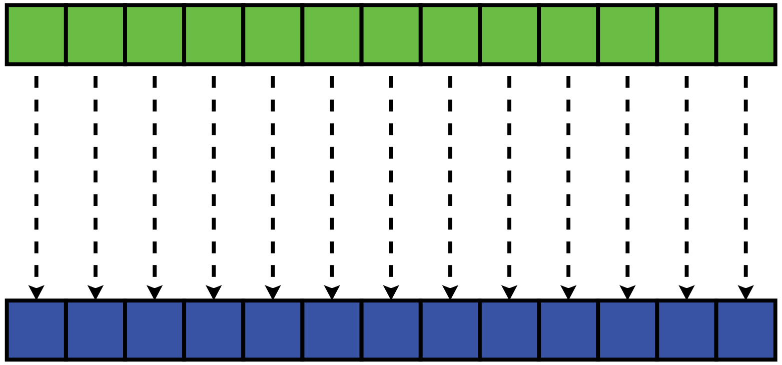

Given the labelspace $\mathcal{L} = \{0, \ldots, k\}$, the involved variables are k copies of the n pixels in the image domain, resulting in $n \cdot k$ vertices.

Infinity-edges between these layers assure the **inclusion constraint** $S_\alpha \subset S_\beta$.

In contrast to the classical Ishikawa consruction, we do not need to enforce them only among neighboring layers.

Minimal Distance Constraint

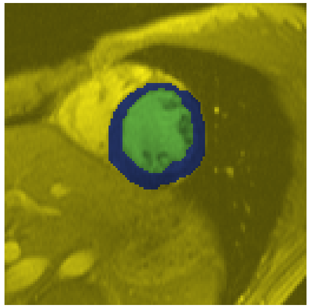

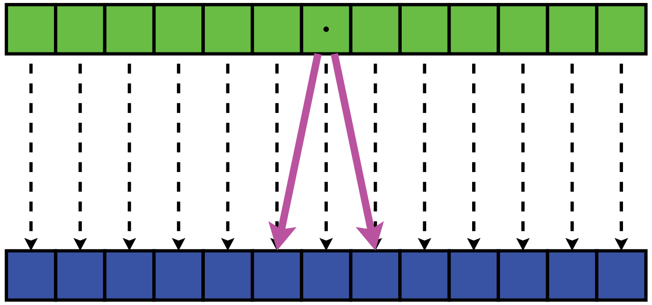

Increasing the neighborhood enforces larger minimal distance constraints.

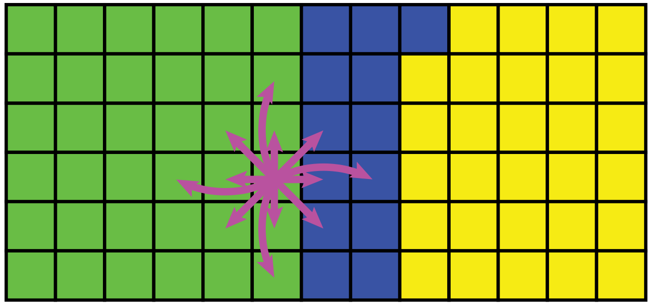

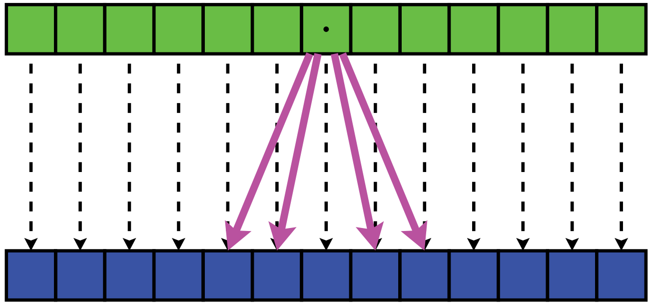

Maximal Distance Constraint

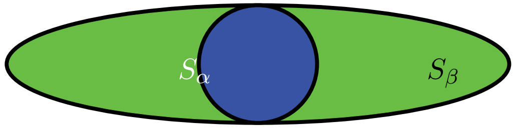

We address the nested multilabeling problem

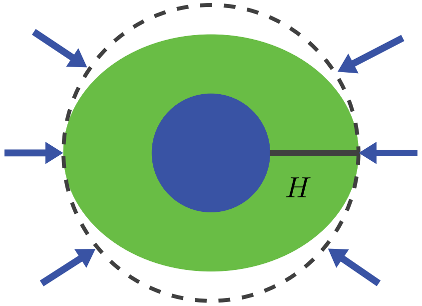

s.t. maximal Hausdorff distance

The **Hausdorff distance** H between S_α and S_β is the minimal value r such that

$S_\beta \oplus B_r \supset S_\alpha$, *i.e.*, $\mathrm{dist}_{\mathsf{HD}}(S_\alpha, S_\beta) = \max_{a \in S_\alpha} \min_{b \in S_\beta} \|a - b\|$

Hausdorff Distance Constraint

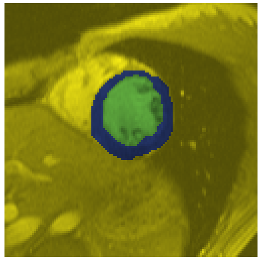

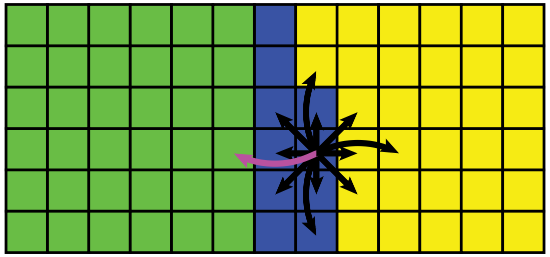

- We want to enforce a Hausdorff distance of at most δ pixels.

$$\mathrm{dist}_{\mathsf{HD}}(A, B) = \max_{a \in A} \min_{b \in B} \|a - b\|$$

- The involved energy can be written as:

$$p_{\mathsf{HD}}(\xi_{i,\ell}) = \infty \cdot \xi_{i,\ell} \prod_{j \in \mathcal{N}_\delta(i)} \bar{\xi}_{j,\ell+1} \leqslant \infty \cdot \xi_{i,\ell} \bar{\xi}_{j,\ell+1}$$

- It is enough if only one edge ensures this constraint.

Hausdorff Distance Constraint

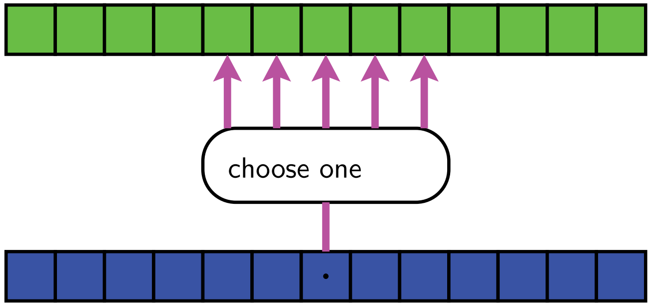

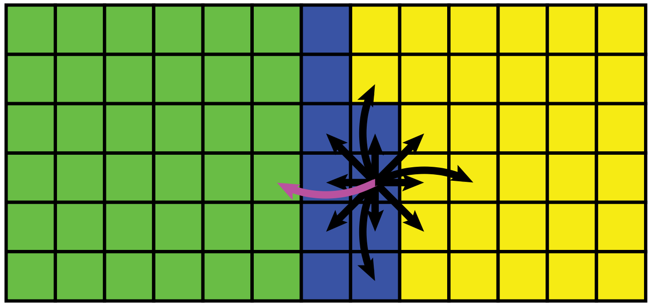

- Enforcing the Hausdorff distance constraint can be formulated within the NP class.
- If we have more than one involved edge, the term becomes

$$p(\xi_{i,\ell}, \xi_{j_1,\ell+1}, \xi_{j_2,\ell+1}) = \infty \cdot \xi_{i,\ell}\bar{\xi}_{j_1,\ell+1}\bar{\xi}_{j_2,\ell+1}.$$

- Since $p(1,\cdot,\cdot)$ is supermodular and $p(\cdot,0,\cdot)$ is submodular, it is difficult to seperate this energy in its submodular and supermodular components.

Submodular-Supermodular Procedure

One method to address sort of problems is the **submodular-supermodular procedure**.

Given: An energy E that is neither sub- nor supermodular.
Idea: Iteratively minimize a submodular upper envelope of E.

1. Let $k = 0$ and S^k an arbitrary "binary" labeling.
2. Find a submodular energy E^{k+1} such that $E^{k+1}(S) \geqslant E(S)$ for all S and $E^{k+1}(S^k) = E(S^k)$.
3. Let $S^{k+1} := \arg\min_S E^{k+1}(S)$ and increase k.
4. If $E^k(S^k) < E^{k-1}(S^{k-1})$ continue at Step 2.

Submodular-Supermodular Procedure

Given: An energy E that is neither sub- nor supermodular.
Idea: Iteratively minimize a submodular upper envelope of E.

1. Let $k = 0$ and S^k an arbitrary "binary" labeling.
   Start with a feasible segmentation.
2. Find a submodular energy E^{k+1} such that
   $E^{k+1}(S) \geqslant E(S)$ for all S and $E^{k+1}(S^k) = E(S^k)$.
   Enforce the maximal distance
   along the shortest path to S_ℓ^k.
3. Let $S^{k+1} := \arg\min_S E^{k+1}(S)$ and increase k.
4. If $E^k(S^k) < E^{k-1}(S^{k-1})$ continue at Step 2.

Guided by Shortest Distance

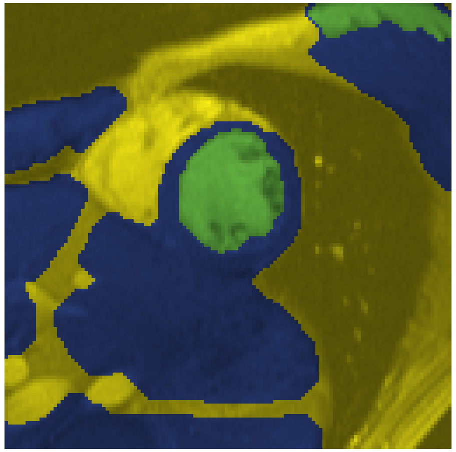

Initial Segmentation

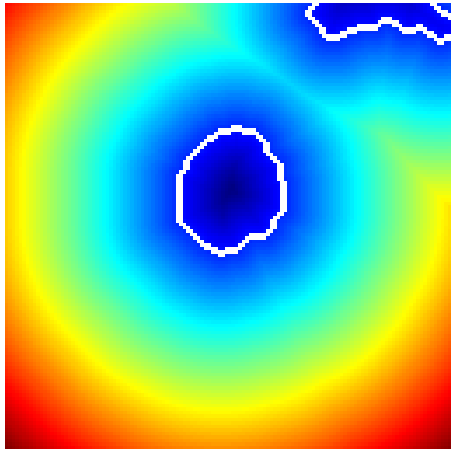

Signed Distance Map

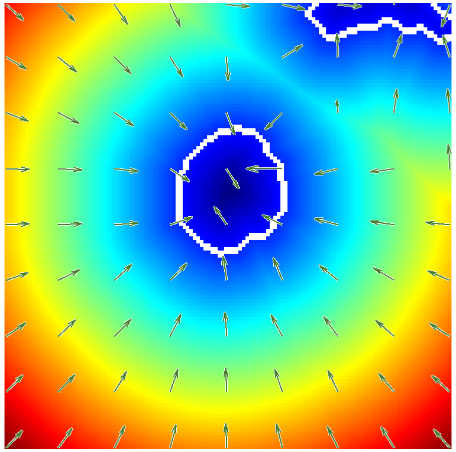

Hausdorff Constraints

In Practice

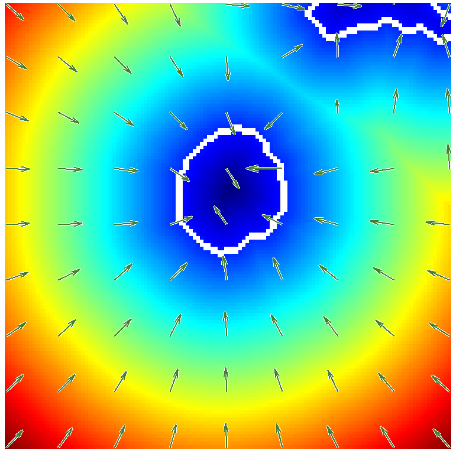

Hausdorff Constraints

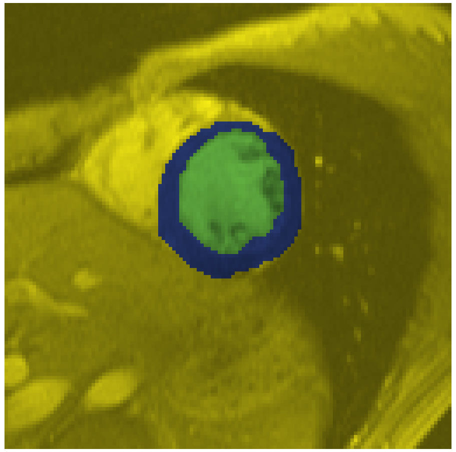

Segmentation Convergence

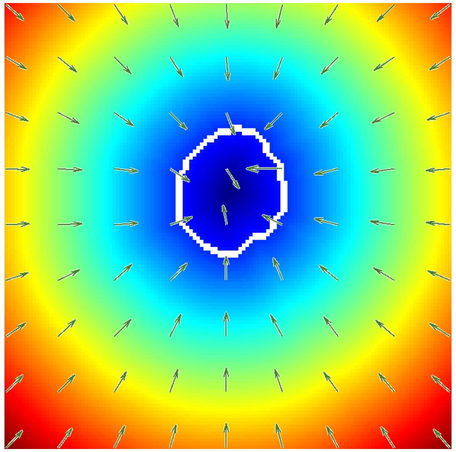

Hausdorff Constraints

Artery Segmentation

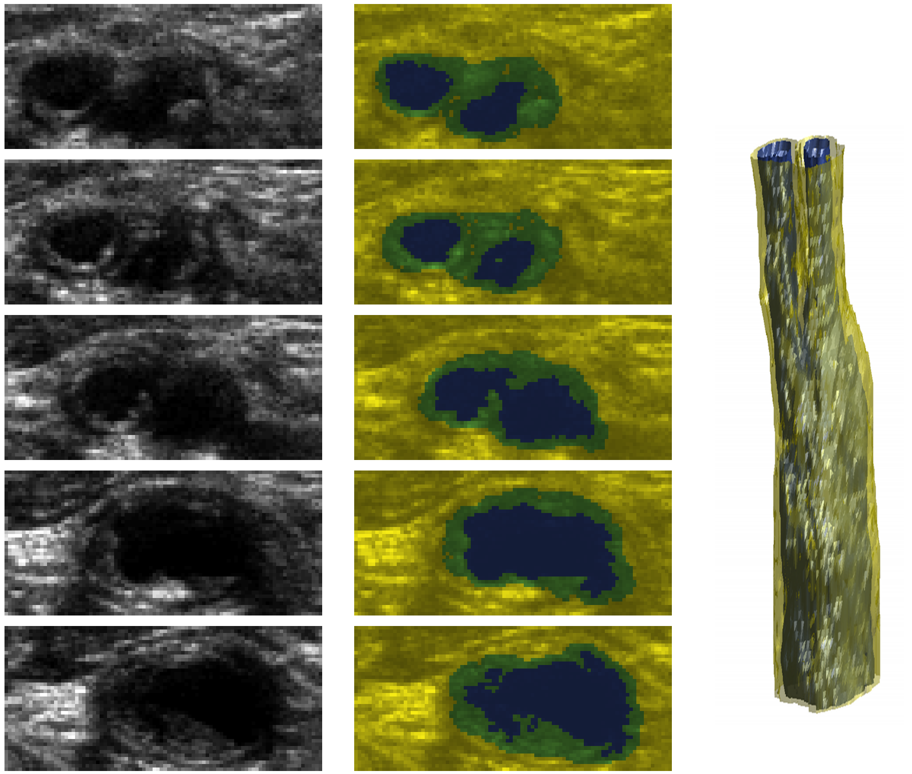

Images Distance Constraints 3D-Segmentation

Brain Segmentation

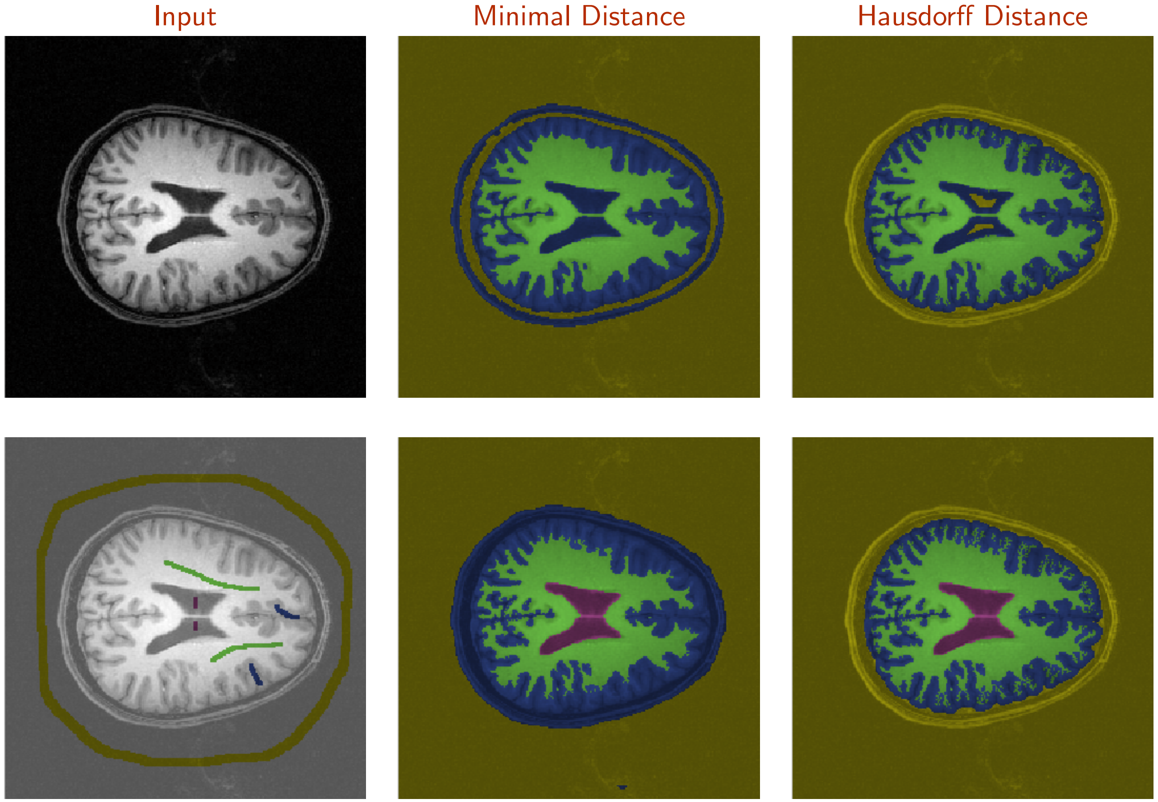

Exclusion Constraint

The exclusion constraint

$$S_\alpha \supset\subset S_\beta \qquad \text{(exclusion constraint)}$$

can be written as

$$\begin{aligned} S_\alpha \cap S_\beta &= \emptyset \\ \neg\left[x \in S_\alpha \wedge x \in S_\beta\right] \\ \neg\left[x \in S_\alpha\right] \vee \left[x \notin S_\beta\right] \\ \left[x \in S_\alpha\right] \Rightarrow \left[x \notin S_\beta\right] \\ S_\alpha &\subset \overline{S}_\beta \end{aligned}$$

Thus, the exclusion constraint is a special version of an inclusion constraint.

Enforcing One Exclusion Constraint

The inclusion constraint can be enforced with the two variables $\xi_{i,1} := [i \in S_\alpha]$, $\xi_{i,2} := [i \in S_\beta]$ and the infinity edge

$$\infty \cdot \xi_{i,1}\overline{\xi}_{i,2}.$$

The exclusion constraint can be enforced with the two variables $\xi_{i,1} := [i \in S_\alpha]$, $\xi_{i,2} := [i \in \overline{S}_\beta]$ and the infinity edge

$$\infty \cdot \xi_{i,1}\overline{\xi}_{i,2}.$$

Combining Inclusion and Exclusion

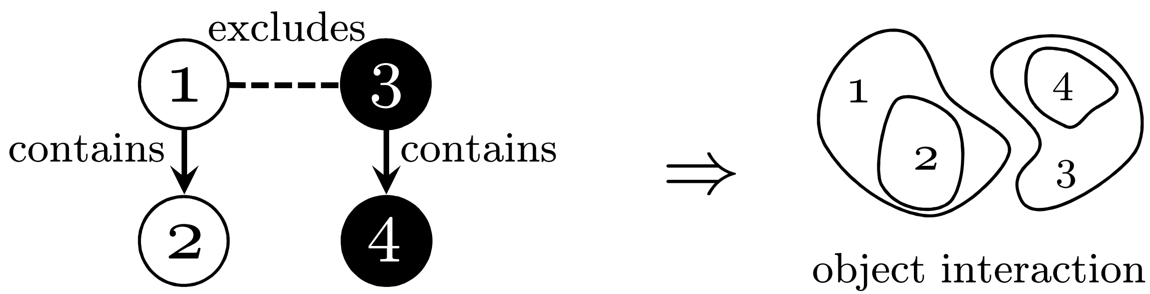

There are situations when the exclusion constraint can be combined with inclusion constraints. For this, certain auxiliary variables $\xi_{i,\ell}$ have to be negated (black in the figure).

There are situations when this leads to **frustrated cycles**.

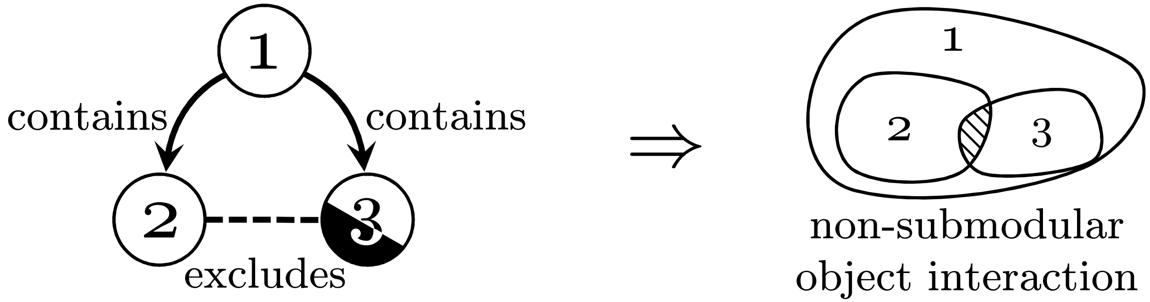

Scene Layout Estimation

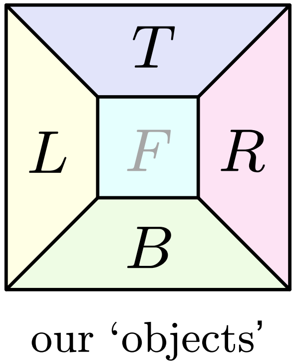

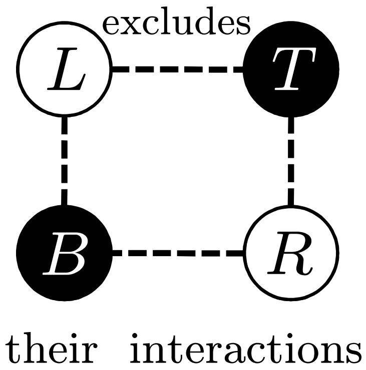

Hoiem et al. presented a method that associated to a pixel the probability that is belongs to one of the following five labels: **B**ottom, **T**op, **L**eft, **R**ight and **F**ront.

In order to avoid frustrated cycles one can model two of the four exclusion constraints within the label set $\{L, T, R, B\}$.

Results

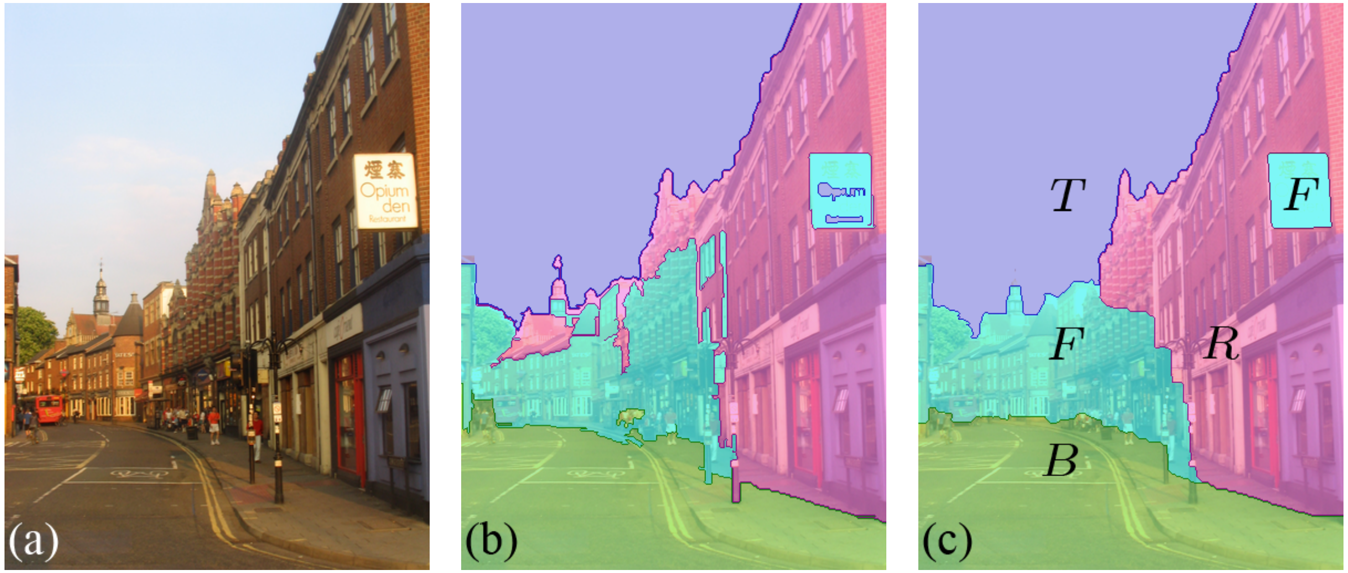

Input — Data Term — + Length

Results

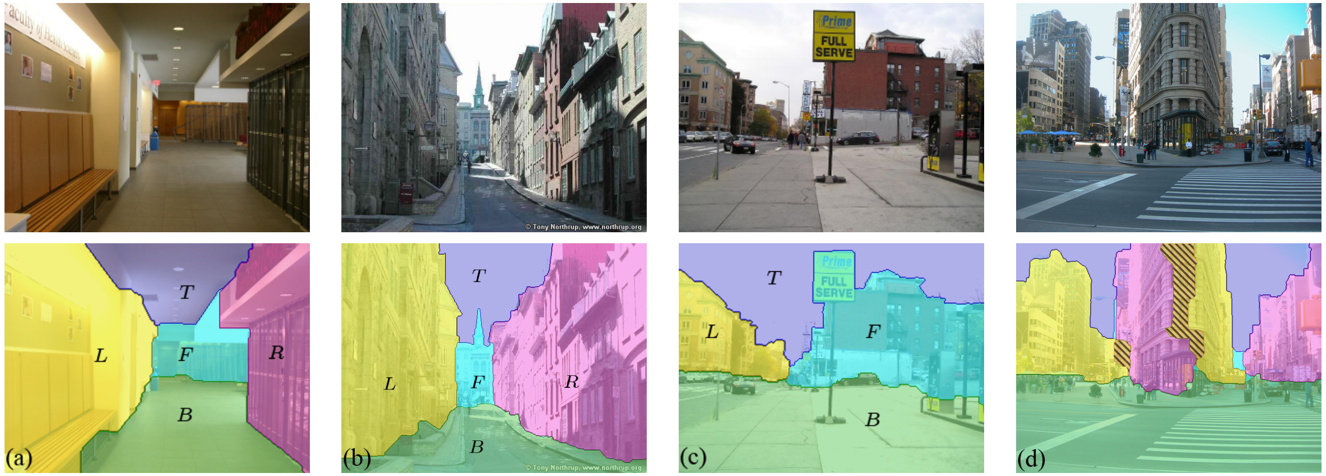

In a lot of situation, only enforcing two constraints might be enough. Nonetheless, there are situations when the exclusion of L and R are needed.

Literature

Maximal Distance Constraint

- Li, Wu, Chen, Sonka, *Optimal Surface Segmentation in Volumetric Images – A Graph-Theoretical Approach*, 2006, IEEE TPAMI 28, 119–134.
- Schmidt, Boykov, *Hausdorff Distance Constraint for Multi-Surface Segmentation*, 2012, ECCV, 598–611.
- Narasimhan, Bilmes, *A submodular-supermodular Procedure with Applications to Discriminative Structured Learning*, 2005 UAI, 404–412.

Exclusion Constraint

- Delong, Boykov, *Globally Optimal Segmentation of Mult-Region Objects*, 2009, IEEE ICCV, 285–292.
- Hoiem, Efros, Herbert, *Recovering Surface Layout from an Image*, 2007, IJCV 75(1).