

Probabilistic Graphical Models in Computer Vision (IN2329)

Csaba Domokos

Summer Semester 2015/2016

7. FastPD & Branch-and-MinCut

FastPD

Recall: Primal-dual LP for multi-label problem *


The (relaxed) primal LP:

$$\begin{aligned} \min_{x_{i:\alpha}, x_{ij:\alpha\beta} \geq 0} & \sum_{i \in \mathcal{V}} \sum_{\alpha \in \mathcal{L}} E_i(\alpha) x_{i:\alpha} + \sum_{(i,j) \in \mathcal{E}} w_{ij} \sum_{\alpha, \beta \in \mathcal{L}} d(\alpha, \beta) x_{ij:\alpha\beta} \\ \text{subject to} & \sum_{\alpha \in \mathcal{L}} x_{i:\alpha} = 1 \quad \forall i \in \mathcal{V} \\ & \sum_{\alpha \in \mathcal{L}} x_{ij:\alpha\beta} = x_{j:\beta} \quad \forall \beta \in \mathcal{L}, (i,j) \in \mathcal{E} \\ & \sum_{\beta \in \mathcal{L}} x_{ij:\alpha\beta} = x_{i:\alpha} \quad \forall \alpha \in \mathcal{L}, (i,j) \in \mathcal{E} \end{aligned}$$

The dual LP:

$$\begin{aligned} \max_{y_i, y_{ij:\alpha}, y_{ji:\beta}} & \sum_{i \in \mathcal{V}} y_i \\ \text{subject to} & y_i - \sum_{j \in \mathcal{V}: (i,j) \in \mathcal{E}} y_{ij:\alpha} \leq E_i(\alpha) \quad \forall i \in \mathcal{V}, \alpha \in \mathcal{L} \\ & y_{ij:\alpha} + y_{ji:\beta} \leq w_{ij} d(\alpha, \beta) \quad \forall (i,j) \in \mathcal{E}, \alpha, \beta \in \mathcal{L} \end{aligned}$$

Recall: Primal-dual schema *



Typically, primal-dual ϵ -approximation algorithms construct a sequence $(x^k, y^k)_{k=1, \dots, t}$ of primal and dual solutions until the elements x^t, y^t of the last pair are both **feasible** and **satisfy the relaxed primal complementary slackness conditions**, hence the condition $\langle c, x \rangle \leq \epsilon \langle b, y \rangle$ will be also fulfilled.

Recall: Complementary slackness conditions *

From now on, in case of Algorithm PD1, we only assume that $d(\alpha, \beta) = 0 \Leftrightarrow \alpha = \beta$, and $d(\alpha, \beta) \geq 0$ (i.e. semi-metric).

The *complementary slackness conditions* reduces to

$$\begin{aligned} y_i - \sum_{j \in \mathcal{V}: (i,j) \in \mathcal{E}} y_{ij:\alpha} & \geq \frac{E_i(x_i)}{\epsilon_1} \Rightarrow y_i \geq \frac{E_i(x_i)}{\epsilon_1} + \sum_{j \in \mathcal{V}: (i,j) \in \mathcal{E}} y_{ij:\alpha} \\ y_{ij:\alpha} + y_{ji:\beta} & \geq \frac{w_{ij} d(x_i, x_j)}{\epsilon_2} \end{aligned}$$


for specific values of $\epsilon_1, \epsilon_2 \geq 1$.

If $x_i = x_j = \alpha$ for neighboring pairs $(i, j) \in \mathcal{E}$, then

$$0 = w_{ij} d(\alpha, \alpha) \geq y_{ij:\alpha} + y_{ji:\alpha} \geq \frac{w_{ij} d(\alpha, \alpha)}{\epsilon_2} = 0,$$

therefore we get that $y_{ij:\alpha} = -y_{ji:\alpha}$.

Recall: Update primal and dual variables *



Dual variables update: Given the current active labels, any non-active label is raised, until it either reaches the active label, or attains the maximum raise allowed by the upper bound.


Primal variables update: Given the new heights, there might still be vertices whose active labels are not at the lowest height. For each such vertex i , we select a non-active label, which is below x_i , but has already reached the maximum raise allowed by the upper bound.

The optimal update of the α -heights can be simulated by pushing the **maximum amount of flow** through a directed graph $G' = (\mathcal{V} \cup \{s, t\}, \mathcal{E}', c, s, t)$.

Recall: Reassign rule *

Label α will be the new label of i (i.e. $x'_i = \alpha$) iff there exists *unsaturated path* (i.e. $f_{ij} < \text{cap}_{ij}$) between the source node s and node i . In all other cases, i keeps its current label (i.e. $x'_i = x_i$).

$$\begin{aligned} f_{ij} & < \text{cap}_{ij} \\ h'_i(\alpha) - h_i(\alpha) & < h_i(x_i) - h_i(\alpha) \\ h'_i(\alpha) & < h_i(x_i) = h'_i(x_i) \end{aligned}$$



$APF^{x',y'} \leq APF^{x,y}$, where $APF^{x,y}$ is defined as

$$\begin{aligned} APF^{x,y} &\triangleq \sum_{i \in \mathcal{V}} h_i(x_i) = \sum_{i \in \mathcal{V}} \left(E_i(x_i) + \sum_{j \in \mathcal{V}, (i,j) \in \mathcal{E}} \text{load}_{ij} \right) \\ &= \sum_{i \in \mathcal{V}} E_i(x_i) + \sum_{(i,j) \in \mathcal{E}} (y_{ij:x_i} + y_{ji:x_j}) \\ &\leq \sum_{i \in \mathcal{V}} E_i(x_i) + \sum_{(i,j) \in \mathcal{E}} w_{ij} d(x_i, x_j) = E(\mathbf{x}). \end{aligned}$$

This condition shows that the algorithm terminates (assuming integer capacities), due to the reassign rule, which ensures that a new active label has always lower height than the previous active label, i.e. $h'_i(x'_i) \leq h_i(x_i)$.

PD2

Parameterization of the PD2 algorithm

We now assume that d is a *metric*.

In fact, PD2 represents a family of algorithms parameterized by $\mu \in [\frac{1}{\epsilon_{\text{app}}}, 1]$. Algorithm PD2 $_{\mu}$ will achieve *complementary slackness conditions* with

$$\epsilon_1 \triangleq \frac{1}{\mu} \epsilon_{\text{app}} \geq \frac{1}{\epsilon_{\text{app}}} \epsilon_{\text{app}} \geq 1 \quad \text{and} \quad \epsilon_2 = \epsilon_{\text{app}}.$$

Algorithm PD1 always generates a feasible dual solution at any of its inner iterations, whereas PD2 $_{\mu}$ **may allow any such dual solution to become infeasible**.

Dual-fitting: PD2 $_{\mu}$ ensures that the (probably infeasible) final dual solution is "not too far away from feasibility", which practically means that if that solution is divided by a suitable factor, it will become feasible again.

Complementary slackness conditions *

Similarly to Algorithm PD1, the equalities will hold for $i \in \mathcal{V}$

$$y_i = \min_{\alpha \in \mathcal{L}} h_i(\alpha) = h_i(x_i) = E_i(x_i) + \sum_{i \in \mathcal{V}, (i,j) \in \mathcal{E}} y_{ij:x_i}.$$

PD2 $_{\mu}$ generates a series of intermediate pairs satisfying *complementary slackness conditions* for $\epsilon_1 \geq 1$ and $\epsilon_2 \geq \frac{1}{\mu} = \frac{1}{1/\epsilon_{\text{app}}} = \epsilon_{\text{app}}$:

$$\begin{aligned} \frac{E_i(x_i)}{\epsilon_1} + \sum_{i \in \mathcal{V}, (i,j) \in \mathcal{E}} y_{ij:x_i} &\leq E_i(x_i) + \sum_{i \in \mathcal{V}, (i,j) \in \mathcal{E}} y_{ij:x_i} = h_i(x_i) = y_i \quad \forall i \in \mathcal{V}. \\ \frac{w_{ij} d(x_i, x_j)}{\epsilon_2} &\leq \mu w_{ij} d(x_i, x_j) = \text{load}_{ij} \quad \forall (i, j) \in \mathcal{E}. \end{aligned}$$

Like PD1, PD2 $_{\mu}$ also maintains non-negativity of *active balance variables*.

Dual fitting

The dual solution of the last intermediate pair may be infeasible, since, instead of the feasibility condition $y_{ij:\alpha} + y_{ji:\beta} \leq w_{ij} d(\alpha, \beta)$, PD2 $_{\mu}$ maintains the conditions:

$$y_{ij:\alpha} + y_{ji:\beta} \leq 2\mu w_{ij} d_{\text{max}} \quad \forall (i, j) \in \mathcal{E}, \forall \alpha, \beta \in \mathcal{L}.$$

These conditions also ensure that the last dual solution \mathbf{y} , is not "too far away from feasibility". By replacing \mathbf{y} with $\mathbf{y}^{\text{fit}} = \frac{\mathbf{y}}{\mu \epsilon_{\text{app}}}$ we get that

$$y_{ij:\alpha}^{\text{fit}} + y_{ji:\beta}^{\text{fit}} = \frac{y_{ij:\alpha} + y_{ji:\beta}}{\mu \epsilon_{\text{app}}} \leq \frac{2\mu w_{ij} d_{\text{max}}}{\mu \epsilon_{\text{app}}} = \frac{2\mu w_{ij} d_{\text{max}}}{\mu 2d_{\text{max}}/d_{\text{min}}} = w_{ij} d_{\text{min}} \leq w_{ij} d(\alpha, \beta).$$

This means that \mathbf{y}^{fit} is **feasible**.

- 1: **function** DUAL_FIT(\mathbf{y})
- 2: **return** $\mathbf{y}^{\text{fit}} \leftarrow \frac{\mathbf{y}}{\mu \epsilon_{\text{app}}}$
- 3: **end function**

Update primal and dual variables

The main/only difference in the subroutine `Update_Duals_Primals`($\alpha, \mathbf{x}, \mathbf{y}$) is the definition of the capacities corresponding to the n -edges. More precisely, assuming an α -iteration, where $x_i = \beta \neq \alpha$ and $x_j = \gamma \neq \alpha$ for a given $(i, j) \in \mathcal{E}$:

$$\begin{aligned} \text{cap}_{ij} &= \mu w_{ij} (d(\beta, \alpha) + d(\alpha, \gamma) - d(\beta, \gamma)), \\ \text{cap}_{ji} &= 0. \end{aligned} \quad (1)$$

All the capacities in the flow must be non-negative. This motivates that d must be a metric.

By applying $\text{load}_{ij} = y_{ij:\beta} + y_{ji:\gamma} = \mu w_{ij} d(\beta, \gamma)$ one can get

$$\begin{aligned} y'_{ij:\alpha} &= y_{ij:\alpha} + \text{cap}_{ij} = y_{ij:\alpha} + \mu w_{ij} (d(\beta, \alpha) + d(\alpha, \gamma) - d(\beta, \gamma)) \\ &= y_{ij:\alpha} + y_{ij:\beta} + y_{ji:\alpha} + \mu w_{ij} d(\alpha, \gamma) - y_{ij:\beta} - y_{ji:\gamma} = \mu w_{ij} d(\alpha, \gamma) - y_{ji:\gamma}, \end{aligned}$$

which ensures that

$$\text{load}_{ij}^{x',y'} = y_{ij:\alpha} + y_{ji:\gamma} = (\mu w_{ij} d(\alpha, \gamma) - y_{ji:\gamma}) + y_{ji:\gamma} = \mu w_{ij} d(\alpha, \gamma).$$

Subroutine PreEdit_Duals($\alpha, \mathbf{x}, \mathbf{y}$) *

The role of this routine is to edit current solution \mathbf{y} , before the subroutine `Update_Duals_Primals`(α, \mathbf{x}), so that

$$\text{load}_{ij}^{x',y'} = y_{ij:\alpha} + y_{ji:\gamma} = \mu w_{ij} d(\alpha, \gamma).$$

- 1: **function** PREEDIT_DUALS($\alpha, \mathbf{x}, \mathbf{y}$)
- 2: **for all** $(i, j) \in \mathcal{E}$ with $x_i \neq \alpha, x_j \neq \alpha$ **do**
- 3: $y_{ij:\alpha} \leftarrow \mu w_{ij} d(\alpha, \gamma) - y_{ji:\gamma}$
- 4: $y_{ji:\alpha} \leftarrow y_{ji:\gamma} - \mu w_{ij} d(\alpha, \gamma)$
- 5: **end for**
- 6: **return** \mathbf{y}
- 7: **end function**

Equivalence of PD2 $_{\mu=1}$ and α -expansion

One can show that PD2 $_{\mu=1}$ indeed generates an ϵ_{app} solution.

If $\mu = 1$, then $\text{load}_{ij} = w_{ij} d(x_i, x_j)$. It can be shown that $APF^{x,y} = E(\mathbf{x})$, whereas in any other case $APF^{x,y} \leq E(\mathbf{x})$.

If $\mu < 1$, then the primal (dual) objective function necessarily decreases (increases) per iteration. Instead, APF constantly decreases.

Recall that APF is the sum of active labels' heights and PD2 $_{\mu=1}$ always tries to choose the *lowest* label among x_i and α . During an α -iteration, PD2 $_{\mu=1}$ chooses an x' that minimizes APF with respect to any other α -expansion \bar{x} of current solution \mathbf{x} .

Theorem 1. Let (x', y') denote the next primal-dual pair due to an α -iteration and let \bar{x} denote α -expansion of the current primal. Then

$$E(x') = APF^{x',y'} \leq APF^{\bar{x},y'} \leq E(\bar{x}).$$

$E(x') \leq E(\bar{x})$ shows that the α -expansion algorithm is equivalent to PD2 $_{\mu=1}$.

PD3

By modifying the Algorithm PD2 _{$\mu=1$} , we will get Algorithm PD3, which can be applied even if d is non-metric function.

Recall that PD2 _{$\mu=1$} maintains the *optimality criterion*: $\text{load}_{ij} \leq w_{ij}d(x_i, x_j)$.

Since d is not metric, we have **conflicting label-triplet** (α, β, γ) :

$$d(\beta, \gamma) > d(\beta, \alpha) + d(\alpha, \gamma).$$

Algorithm PD3_a: During the primal-dual variable update, in an α -iteration, when $x_i \neq \alpha$ and $x_j \neq \alpha$, i.e. in (1), we set $\text{cap}_{ij} = 0$.

It can be shown that for a *conflicting triplet*

$$\text{load}_{ij} = w_{ij}(d(\beta, \gamma) - d(\beta, \alpha)) \geq w_{ij}d(\alpha, \gamma).$$

Intuitively, PD3_a **overestimates the distance** between labels α, γ in order to restore the triangle inequality for the current conflicting label-triplet (α, β, γ) .

PD3_b *

We choose to set $\text{cap}_{ij} = +\infty$ and no further differences between PD3_b and PD2 _{$\mu=1$} exist.

This has the following important effect: the solution \mathbf{x}' produced at the current iteration, can never assign the pair of labels γ, β to the objects i, j respectively.

PD3_c *

PD3_c first adjusts the dual solution \mathbf{y} for any $(i, j) \in \mathcal{E}$:

$$\text{load}_{ij} \leq w_{ij}d(\alpha, \gamma) + d(\gamma, \beta).$$

After this initial adjustment, PD3_c proceeds exactly as PD2 _{$\mu=1$} , except for the fact that the term $d(\alpha, \beta)$ (1) is replaced by

$$\bar{d}(\beta, \gamma) \triangleq \frac{\text{load}_{ij}}{w_{ij}} \leq d(\beta, \alpha) + d(\alpha, \gamma) < d(\beta, \gamma).$$

Intuitively, PD3_c works in a complementary way to PD3_a algorithm, i.e. in order to restore the triangle inequality for the conflicting label-triplet (α, β, γ) , it chooses to **underestimate the distance** between labels (β, γ) (instead of overestimating the distance between either labels α, γ or α, β).

Results: Stereo matching *



| Distance $d(\alpha, \beta)$ | $\epsilon_{\text{app}}^{\text{PD1}}$ | $\epsilon_{\text{app}}^{\text{PD2}_{\mu=1}}$ | $\epsilon_{\text{app}}^{\text{PD3}_a}$ | $\epsilon_{\text{app}}^{\text{PD3}_b}$ | $\epsilon_{\text{app}}^{\text{PD3}_c}$ | ϵ_{app} |
|---------------------------------|--------------------------------------|--|--|--|--|-------------------------|
| $\mathbb{1}[\alpha \neq \beta]$ | 1.0104 | 1.0058 | 1.0058 | 1.0058 | 1.0058 | 2 |
| $\min(5, \alpha - \beta)$ | 1.0226 | 1.0104 | 1.0104 | 1.0104 | 1.0104 | 10 |
| $\min(5, \alpha - \beta ^2)$ | 1.0280 | - | 1.0143 | 1.0158 | 1.0183 | 10 |

Branch-and-MinCut

Introduction

We address the problem of **binary image segmentation**, where we also consider *non-local parameters* that are known *a priori*.

For example, one can assume prior knowledge about the **shape** of the foreground segment or the **color distribution** of the foreground and/or background.

Let us consider an *undirected graphical model* $G = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of pixels and \mathcal{E} consists of 8-connected pairs of pixels. We define the energy function $E: \{0, 1\}^{\mathcal{V}} \times \Omega \rightarrow \mathbb{R}$ for non-local parameter $\omega \in \Omega$:

$$E(\mathbf{y}, \omega) = C(\omega) + \sum_{i \in \mathcal{V}} F^i(\omega) \cdot y_i + \sum_{i \in \mathcal{V}} B^i(\omega) \cdot (1 - y_i) + \sum_{(i,j) \in \mathcal{E}} P^{ij}(\omega) \cdot |y_i - y_j|,$$

where $C(\omega)$ is a *constant energy* w.r.t. \mathbf{y} , and $F^i(\omega)$ and $B^i(\omega)$ are the *unary energies* defining the cost of assigning the pixel i to the foreground and to the background, respectively. $P^{ij}(\omega) \in \mathbb{R}_0^+$ is **non-negative** for each $(i, j) \in \mathcal{E}$ ensuring the tractability of $E(\mathbf{x}, \omega)$.

Globally optimal segmentation *

The segmentation is given by binary labeling $\mathbf{y} \in \mathbb{B}^{\mathcal{V}} = \{0, 1\}^{\mathcal{V}}$, where individual pixel labels are denoted by $y_i \in \mathbb{B}$ (1:foreground, 0:background). We assume that non-local parameter $\omega \in \Omega$ are taken from a **discrete set**.

Shape priors will be encoded as a product space of various poses and deformations of the *template*, while color priors will correspond to the set of parametric color distributions.

The goal is to achieve a **globally optimal** segmentation under non-local priors. The applied optimization method relies on two techniques: **graph cuts** and **branch-and-bound**.

Although a global minimum can be achieved, the worst case complexity of the method is large (essentially, the same as the exhaustive search over the space of non-local parameters).

An alternative way to solve the problem is to apply *alternating minimization*.

Lower bound

FastPD PD2 PD3 Branch-and-MinCut

$L(\Omega)$ denotes the lower bound for $E(\mathbf{y}, \omega)$ over $\mathbb{B}^V \times \Omega$:

$$\begin{aligned} & \min_{\mathbf{y} \in \mathbb{B}^V, \omega \in \Omega} E(\mathbf{y}, \omega) \\ &= \min_{\mathbf{y} \in \mathbb{B}^V, \omega \in \Omega} \left\{ C(\omega) + \sum_{i \in V} F^i(\omega) \cdot y_i + \sum_{i \in V} B^i(\omega) \cdot (1 - y_i) + \sum_{(i,j) \in \mathcal{E}} P^{ij}(\omega) \cdot |y_i - y_j| \right\} \\ &\geq \min_{\mathbf{y} \in \mathbb{B}^V} \left\{ \min_{\omega \in \Omega} C(\omega) + \sum_{i \in V} \min_{\omega \in \Omega} F^i(\omega) \cdot y_i + \sum_{i \in V} \min_{\omega \in \Omega} B^i(\omega) \cdot (1 - y_i) + \right. \\ &\quad \left. \sum_{(i,j) \in \mathcal{E}} \min_{\omega \in \Omega} P^{ij}(\omega) \cdot |y_i - y_j| \right\} \\ &= \min_{\mathbf{y} \in \mathbb{B}^V} \left\{ C_\Omega + \sum_{i \in V} F_\Omega^i(\omega) \cdot y_i + \sum_{i \in V} B_\Omega^i(\omega) \cdot (1 - y_i) + \sum_{(i,j) \in \mathcal{E}} P_\Omega^{ij}(\omega) \cdot |y_i - y_j| \right\} \\ &= L(\Omega). \end{aligned}$$

$C_\Omega, F_\Omega^i, B_\Omega^i, P_\Omega^{ij}$ denote the minima of $C(\omega), F^i(\omega), B^i(\omega), P^{ij}(\omega)$ over $\omega \in \Omega$ referred to as **aggregated energies**.

IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 25 / 38

Monotonicity

FastPD PD2 PD3 Branch-and-MinCut

Suppose $\Omega_1 \subset \Omega_2$, then the inequality $L(\Omega_1) \geq L(\Omega_2)$ holds.

Proof. Let us define $A(\mathbf{y}, \Omega)$ as

$$\begin{aligned} A(\mathbf{y}, \Omega) &\triangleq \min_{\omega \in \Omega} C(\omega) + \sum_{i \in V} \min_{\omega \in \Omega} F^i(\omega) \cdot y_i + \sum_{i \in V} \min_{\omega \in \Omega} B^i(\omega) \cdot (1 - y_i) \\ &\quad + \sum_{(i,j) \in \mathcal{E}} \min_{\omega \in \Omega} P^{ij}(\omega) \cdot |y_i - y_j|. \end{aligned}$$

Assume $\Omega_1 \subset \Omega_2$. Then, for any $\mathbf{y} \in \mathbb{B}^V$

$$\begin{aligned} & A(\mathbf{y}, \Omega_1) \\ &= \min_{\omega \in \Omega_1} C(\omega) + \sum_{i \in V} \min_{\omega \in \Omega_1} F^i(\omega) y_i + \sum_{i \in V} \min_{\omega \in \Omega_1} B^i(\omega) (1 - y_i) + \sum_{(p,q) \in \mathcal{E}} \min_{\omega \in \Omega_1} P^{pq}(\omega) |y_p - y_q| \\ &\geq \min_{\omega \in \Omega_2} C(\omega) + \sum_{i \in V} \min_{\omega \in \Omega_2} F^i(\omega) y_i + \sum_{i \in V} \min_{\omega \in \Omega_2} B^i(\omega) (1 - y_i) + \sum_{(i,j) \in \mathcal{E}} \min_{\omega \in \Omega_2} P^{ij}(\omega) |y_i - y_j| \\ &= A(\mathbf{y}, \Omega_2). \end{aligned}$$

IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 26 / 38

Monotonicity

FastPD PD2 PD3 Branch-and-MinCut

Proof. Continued

Note that $L(\Omega) = \min_{\mathbf{y} \in \mathbb{B}^V} A(\mathbf{y}, \Omega)$.

Let $\mathbf{y}_1 \in \operatorname{argmin}_{\mathbf{y} \in \mathbb{B}^V} A(\mathbf{y}, \Omega_1)$ and $\mathbf{y}_2 \in \operatorname{argmin}_{\mathbf{y} \in \mathbb{B}^V} A(\mathbf{y}, \Omega_2)$, then from the monotonicity, one gets:

$$L(\Omega_1) = A(\mathbf{y}_1, \Omega_1) \geq A(\mathbf{y}_1, \Omega_2) \geq A(\mathbf{y}_2, \Omega_2) = L(\Omega_2). \quad \square$$

IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 27 / 38

Computability and tightness

FastPD PD2 PD3 Branch-and-MinCut

Computability: the lower bound $L(\Omega)$ equals the minimum of a *regular function*, which can be globally minimized via graph-cuts.

Tightness: for a singleton $\Omega = \{\omega\}$ (i.e. $|\Omega| = 1$) the bound $L(\Omega)$ is **tight**, that is


$$L(\{\omega\}) = \min_{\mathbf{y} \in \mathbb{B}^V} E(\mathbf{y}, \omega).$$

IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 28 / 38

Best-first branch-and-bound optimization

FastPD PD2 PD3 Branch-and-MinCut



The discrete domain Ω can be hierarchically clustered and the binary tree of its subregions can be considered.


At each step the *active node* with the **smallest** lower bound is removed from the *active front*, while two of its *children* are added to the *active front* (due to *monotonicity property* they have higher or equal lower bounds).

IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 29 / 38

Best-first branch-and-bound optimization

FastPD PD2 PD3 Branch-and-MinCut



If the *active node* with the smallest lower bound turns out to be a **leaf** ω' and \mathbf{y}' is the corresponding optimal segmentation, then $E(\mathbf{y}', \omega') = L(\omega')$ due to the *tightness property*. Consequently, (\mathbf{y}', ω') is a **global minimum**.

Remark that in worst-case any optimization has to search exhaustively over Ω .

IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 30 / 38

Pseudo code of Branch-And-Mincut *

FastPD PD2 PD3 Branch-and-MinCut

```

1: Front  $\leftarrow \emptyset$  ▷ initializing the priority queue
2:  $[C_0, \{F_0^i\}, \{B_0^i\}, \{P_0^{ij}\}] \leftarrow \text{GetAggregPotentials}(\Omega_0)$ 
3:  $\text{LB}_0 \leftarrow \text{GetMaxFlowValue}(\{F_0^i\}, \{B_0^i\}, \{P_0^{ij}\}) + C_0$ 
4:  $\text{Front.InsertWithPriority}(\Omega_0, -\text{LB}_0)$ 
5: while true do ▷ advancing front
6:    $\Omega \leftarrow \text{Front.PullHighestPriorityElement}()$ 
7:   if  $\text{IsSingleton}(\Omega)$  then ▷ global minimum found
8:      $\omega \leftarrow \Omega$ 
9:      $[C, \{F^i\}, \{B^i\}, \{P^{ij}\}] \leftarrow \text{GetAggregPotentials}(\omega)$ 
10:     $\mathbf{x} \leftarrow \text{FindMinimumViaMincut}(\{F^i\}, \{B^i\}, \{P^{ij}\})$ 
11:    return  $(\mathbf{x}, \omega)$ 
12:  end if
13:   $[\Omega_1, \Omega_2] \leftarrow \text{GetChildrenSubdomains}(\Omega)$ 
14:   $[C_1, \{F_1^i\}, \{B_1^i\}, \{P_1^{ij}\}] \leftarrow \text{GetAggregPotentials}(\Omega_1)$ 
15:   $\text{LB}_1 \leftarrow \text{GetMaxFlowValue}(\{F_1^i\}, \{B_1^i\}, \{P_1^{ij}\}) + C_1$ 
16:   $\text{Front.InsertWithPriority}(\Omega_1, -\text{LB}_1)$ 
17:   $[C_2, \{F_2^i\}, \{B_2^i\}, \{P_2^{ij}\}] \leftarrow \text{GetAggregPotentials}(\Omega_2)$ 
18:   $\text{LB}_2 \leftarrow \text{GetMaxFlowValue}(\{F_2^i\}, \{B_2^i\}, \{P_2^{ij}\}) + C_2$ 
19:   $\text{Front.InsertWithPriority}(\Omega_2, -\text{LB}_2)$ 
20: end while

```

IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 31 / 38

Segmentation with shape priors

FastPD PD2 PD3 Branch-and-MinCut

The prior is defined by the set of exemplar binary segmentations $\{\mathbf{x}^\omega \mid \omega \in \Omega\}$, where Ω is a discrete set indexing the exemplar segmentations.

We define a joint prior over the segmentation and the non-local parameter:

$$E_{\text{prior}}(\mathbf{y}, \omega) = \sum_{i \in V} (1 - x_i^\omega) \cdot y_i + \sum_{i \in V} x_i^\omega \cdot (1 - y_i).$$

This encourages the segmentation \mathbf{y} to be close in the Hamming-distance ($d_H(\mathbf{a}, \mathbf{b}) = \frac{1}{N} \sum_{i=1}^N [a_i \neq b_i]$) to one of the exemplar shapes.

The segmentation energy may be defined by adding a standard *contrast-sensitive Potts-model* for $\lambda, \sigma > 0$:

$$E(\mathbf{y}, \omega) = E_{\text{prior}}(\mathbf{y}, \omega) + \lambda \sum_{(i,j) \in \mathcal{E}} \frac{e^{-\frac{|I_i - I_j|}{\sigma}}}{|i - j|} \cdot |y_i - y_j|,$$

where I_i denotes RGB colors of the pixel i .


IN2329 - Probabilistic Graphical Models in Computer Vision

7. FastPD & Branch-and-MinCut - 32 / 38

The shape prior is given by a set of templates, whereas each template can be located anywhere within the image.

$\Omega = \Delta \times \Theta$, where the set Δ indexes the set of all exemplar segmentations x_δ and Θ corresponds to translations.

Any exemplar segmentation x^ω for $\omega = (\delta, \theta)$ is then defined as some *exemplar segmentation* x_δ centered at the origin and then *translated* by the shift θ .



For Δ we use agglomerative bottom-up clustering resulting in a (binary) clustering tree $T_\Delta = \{\Delta = \Delta_0, \Delta_1, \dots, \Delta_N\}$.

To build a clustering tree for Θ , we recursively split along the “longer” dimension. This leads to a (binary) tree $T_\Theta = \{\Theta = \Theta_0, \Theta_1, \dots, \Theta_N\}$.

Branch operation

Each *nodeset* Ω_t in the combined tree is defined by a pair $\Delta_t \times \Theta_t$.

The **looseness** of a nodeset Ω_t is defined as the number of pixels that change their mask value under different shapes in Ω_t (i.e. neither background nor foreground):

$$\Lambda(\Omega_t) = |\{i \mid \exists \omega_1, \omega_2 : x_i^{\omega_1} = 0 \text{ and } x_i^{\omega_2} = 1\}|.$$

The tree is built in a recursive top-down fashion as follows.

We start by creating a root nodeset $\Omega_0 = \Delta_0 \times \Theta_0$. Given a nodeset $\Omega_t = \Delta_t \times \Theta_t$ we consider (recursively) two possible splits: 1) split along the shape dimension or 2) split along the shift dimension. The split that minimizes the sum of loosenesses is preferred.

The recursion stops when the leaf level is reached within both the *shape* and the *shift* trees.

Results *



Yellow: global minimum of E ; Blue: feature-based car detector; Red: global minimum of the combination of E with detection results (detection is included as a constant potential)

The prior set Δ was built by manual segmentation of 60 training images coming with the dataset.

Summary *

Primal-dual schema:

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} E_i(x_i) + \sum_{(i,j) \in \mathcal{E}} w_{ij} \cdot d(x_i, x_j)$$

- ◆ PD1: d is a semi-metric
- ◆ PD2: d is a metric (equivalent to α -expansion)
- ◆ PD3: d is a non-metric function

For **binary image segmentation** we learned a global optimal solution, based on *branch and bound* optimization, in the presence of (shape) prior information.

In the **next lecture** we will learn about **exact inference** (probabilistic and MAP) on tree structured factor graphs.

Literature *

FastPD

1. Nikos Komodakis and Georgios Tziritas. Approximate labeling via the primal-dual schema. Technical report, University of Crete, February 2005
2. Nikos Komodakis and Georgios Tziritas. Approximate labeling via graph-cuts based on linear programming. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(8):1436–1453, August 2007

Branch-and-MinCut

1. Victor Lempitsky, Andrew Blake, and Carsten Rother. Branch-and-mincut: Global optimization for image segmentation with high-level priors. *Journal of Mathematical Imaging and Vision*, 44(3):315–329, March 2012