

Energy Minimization

Raquel Urtasun

TTI Chicago

Feb 26, 2013

Disparity Estimation

- DSI: Disparity image



Scene



Ground truth

Stereo Estimation Methods

- Local methods
- Grow and seed methods: use a few good correspondences and grow the estimation from them
- Adaptive Window methods (AW)
- Global methods: define a Markov random field over
 - Pixel-level
 - Fronto-parallel planes
 - Slanted planes

- The energy is defined as

$$E(d_1, \dots, d_n) = \sum_i C(d_i) + \sum_i \sum_{j \in \mathcal{N}(i)} C(d_i, d_j)$$

where $x_i \in \{0, 1, \dots, D\}$ represents a variable for the disparity of the i -th pixel

- This optimization is in general NP-hard.

- The energy is defined as

$$E(d_1, \dots, d_n) = \sum_i C(d_i) + \sum_i \sum_{j \in \mathcal{N}(i)} C(d_i, d_j)$$

where $x_i \in \{0, 1, \dots, D\}$ represents a variable for the disparity of the i -th pixel

- This optimization is in general NP-hard.
- Global optima can be obtained in a few cases.

- The energy is defined as

$$E(d_1, \dots, d_n) = \sum_i C(d_i) + \sum_i \sum_{j \in \mathcal{N}(i)} C(d_i, d_j)$$

where $x_i \in \{0, 1, \dots, D\}$ represents a variable for the disparity of the i -th pixel

- This optimization is in general NP-hard.
- Global optima can be obtained in a few cases.

Semiglobal block matching [Hirschmueller08]

- The energy is defined as

$$E(d_1, \dots, d_n) = \sum_i C(d_i) + \sum_i \sum_{j \in \mathcal{N}(i)} C(d_i, d_j)$$

with the following pairwise term

$$C(d_i, d_j) = \begin{cases} 0 & \text{if } d_i = d_j \\ \lambda_1 & \text{if } |d_i - d_j| = 1 \\ \lambda_2 & \text{otherwise} \end{cases}$$

- It computes the costs in each direction

$$D_j(\mathbf{p}; d) = C(\mathbf{p}; d) + \min_{d'} \{D(\mathbf{p} - \mathbf{j}, d') + \rho_d(d - d')\}$$

Semiglobal block matching [Hirschmueller08]

- The energy is defined as

$$E(d_1, \dots, d_n) = \sum_i C(d_i) + \sum_i \sum_{j \in \mathcal{N}(i)} C(d_i, d_j)$$

with the following pairwise term

$$C(d_i, d_j) = \begin{cases} 0 & \text{if } d_i = d_j \\ \lambda_1 & \text{if } |d_i - d_j| = 1 \\ \lambda_2 & \text{otherwise} \end{cases}$$

- It computes the costs in each direction

$$D_j(\mathbf{p}; d) = C(\mathbf{p}; d) + \min_{d'} \{D(\mathbf{p} - \mathbf{j}, d') + \rho_d(d - d')\}$$

- And aggregate the costs

$$D(\mathbf{p}; d) = \sum_j L_j(\mathbf{p}, d)$$

Semiglobal block matching [Hirschmueller08]

- The energy is defined as

$$E(d_1, \dots, d_n) = \sum_i C(d_i) + \sum_i \sum_{j \in \mathcal{N}(i)} C(d_i, d_j)$$

with the following pairwise term

$$C(d_i, d_j) = \begin{cases} 0 & \text{if } d_i = d_j \\ \lambda_1 & \text{if } |d_i - d_j| = 1 \\ \lambda_2 & \text{otherwise} \end{cases}$$

- It computes the costs in each direction

$$D_j(\mathbf{p}; d) = C(\mathbf{p}; d) + \min_{d'} \{D(\mathbf{p} - \mathbf{j}, d') + \rho_d(d - d')\}$$

- And aggregate the costs

$$D(\mathbf{p}; d) = \sum_j L_j(\mathbf{p}, d)$$

- Then do winner take all

- The energy is defined as

$$E(d_1, \dots, d_n) = \sum_i C(d_i) + \sum_i \sum_{j \in \mathcal{N}(i)} C(d_i, d_j)$$

with the following pairwise term

$$C(d_i, d_j) = \begin{cases} 0 & \text{if } d_i = d_j \\ \lambda_1 & \text{if } |d_i - d_j| = 1 \\ \lambda_2 & \text{otherwise} \end{cases}$$

- It computes the costs in each direction

$$D_j(\mathbf{p}; d) = C(\mathbf{p}; d) + \min_{d'} \{D(\mathbf{p} - \mathbf{j}, d') + \rho_d(d - d')\}$$

- And aggregate the costs

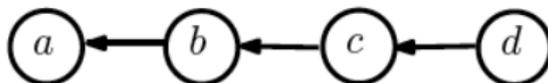
$$D(\mathbf{p}; d) = \sum_j L_j(\mathbf{p}, d)$$

- Then do winner take all

- Given distribution $p(y_1, \dots, y_n)$
- Inference: computing functions of the distribution
 - mean
 - marginal
 - conditionals
- Marginal inference in singly-connected graph (trees)
- Later: extensions to loopy graphs

[Source: P. Gebler]

- ▶ Consider Markov chain $(a, b, c, d \in \{0, 1\})$



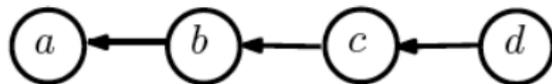
with distribution

$$p(a, b, c, d) = p(a | b)p(b | c)p(c | d)p(d)$$

- ▶ Task: compute the marginal $p(a)$

[Source: P. Gehler]

Variable Elimination



$$\begin{aligned} p(a) &= \sum_{b,c,d} p(a, b, c, d) \\ &= \sum_{b,c,d} p(a | b)p(b | c)p(c | d)p(d) \end{aligned}$$

- ▶ Naive: $2 \times 2 \times 2 = 8$ states to sum over
- ▶ Re-order summation:

$$p(a) = \sum_{b,c} p(a | b)p(b | c) \underbrace{\sum_d p(c | d)p(d)}_{\gamma_d(c)}$$

[Source: P. Gehler]

Variable Elimination

$$p(a) = \sum_{b,c} p(a | b)p(b | c) \underbrace{\sum_d p(c | d)p(d)}_{\gamma_d(c)}$$

$$p(a) = \sum_b p(a | b) \underbrace{\sum_c p(b | c)\gamma_d(c)}_{\gamma_c(b)}$$

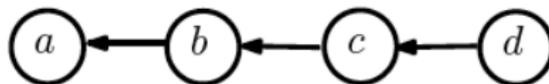
$$p(a) = \sum_b p(a | b)\gamma_c(b)$$

- ▶ We need $2 + 2 + 2 = 6$ calculations
- ▶ For a chain of length T scale linearly $n * 2$, cf naive approach 2^n

[Source: P. Gehler]

Finding Conditional Marginals

- ▶ Again:



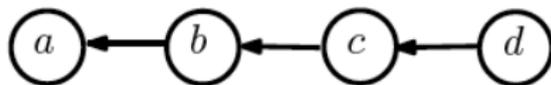
$$p(a, b, c, d) = p(a | b)p(b | c)p(c | d)p(d)$$

- ▶ Now find $p(d | a)$

$$\begin{aligned} p(d | a) &\propto \sum_{b,c} p(a | b)p(b | c)p(c | d)p(d) \\ &= \sum_c \underbrace{\sum_b p(a | b)p(b | c)p(c | d)p(d)}_{\gamma_b(c)} \\ &\stackrel{\text{def}}{=} \gamma_c(d) \text{ not a distribution} \end{aligned}$$

[Source: P. Gehler]

Finding Conditional Marginals



- ▶ Found that

$$p(d | a) = k\gamma_c(d)$$

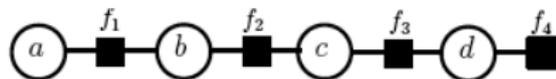
- ▶ and since $\sum_d p(d | a) = 1$

$$k = \frac{1}{\sum_d \gamma_c(d)}$$

- ▶ Again $\gamma_c(d)$ is not a distribution (but a message)

[Source: P. Gehler]

Now with factor graphs



$$p(a, b, c, d) = f_1(a, b)f_2(b, c)f_3(c, d)f_4(d)$$

$$\begin{aligned} p(a, b, c) &= \sum_d p(a, b, c, d) \\ &= f_1(a, b)f_2(b, c) \underbrace{\sum_d f_3(c, d)f_4(d)}_{\mu_{d \rightarrow c}(c)} \end{aligned}$$

$$p(a, b) = \sum_c p(a, b, c) = f_1(a, b) \underbrace{\sum_c f_2(b, c)\mu_{d \rightarrow c}(c)}_{\mu_{c \rightarrow b}(b)}$$

[Source: P. Gehler]

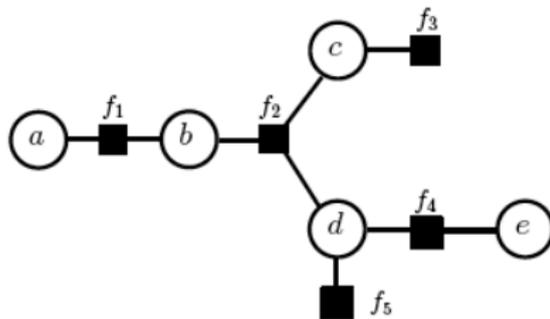
Inference in Chain Structured Factor Graphs

- Simply recurse further
- $\gamma_{m \rightarrow n}(n)$ carries the information beyond m
- We did not need the factors in general (next) we will see that making a distinction is helpful

[Source: P. Gehler]

General singly-connected factor graphs I

- ▶ Now consider a branching graph:



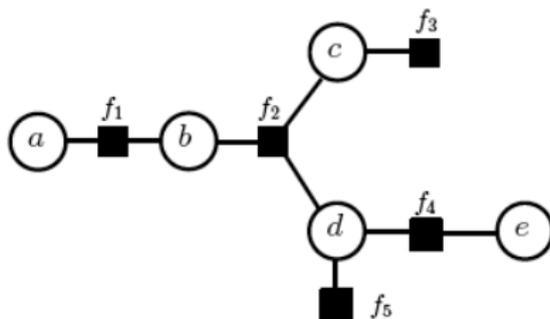
with factors

$$f_1(a, b)f_2(b, c, d)f_3(c)f_4(d, e)f_5(d)$$

- ▶ For example: find marginal $p(a, b)$

[Source: P. Gehler]

General singly-connected factor graphs II

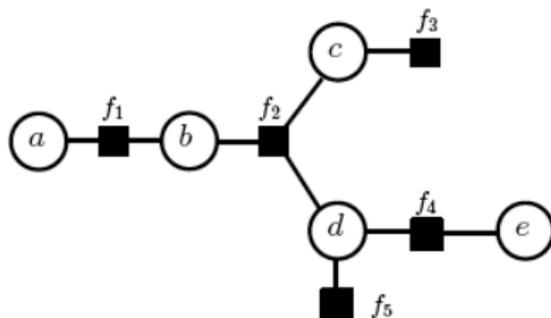


$$p(a, b) = f_1(a, b) \underbrace{\sum_{c, d, e} f_2(b, c, d) f_3(c) f_5(d) f_4(d, e)}_{\mu_{f_2 \rightarrow b}(b)}$$

$$\mu_{f_2 \rightarrow b}(b) = \sum_{c, d} f_2(b, c, d) \underbrace{f_3(c)}_{\mu_{c \rightarrow f_2}(c)} \underbrace{f_5(d) \sum_e f_4(d, e)}_{\mu_{d \rightarrow f_2}(d)}$$

[Source: P. Gehler]

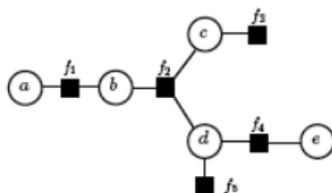
General singly-connected factor graphs III



$$\mu_{d \rightarrow f_2}(d) = \underbrace{f_5(d)}_{\mu_{f_5 \rightarrow d}(d)} \underbrace{\sum_e f_4(d, e)}_{\mu_{f_4 \rightarrow d}(d)}$$

[Source: P. Gehler]

General singly-connected factor graphs IV



- ▶ If we want to compute the marginal $p(a)$:

$$p(a) = \underbrace{\sum_b f_1(a, b) \mu_{f_2 \rightarrow b}(b)}_{\mu_{f_1 \rightarrow a}(a)}$$

- ▶ which we could also view as

$$p(a) = \sum_b f_1(a, b) \underbrace{\mu_{f_2 \rightarrow b}(b)}_{\mu_{b \rightarrow f_1}(b)}$$

[Source: P. Gehler]

Summary

- Once computed, messages can be re-used
- All marginals $p(c), p(d), p(c, d), \dots$ can be written as a function of messages
- We need an algorithm to compute all messages: Sum-Product algorithm

[Source: P. Gehler]

Sum-product algorithm overview

- Algorithm to compute all messages efficiently, assuming the graph is singly-connected
- It can be used to compute any desired marginals
- Also known as belief propagation (BP)

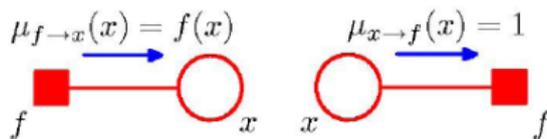
The algorithm is composed of

- 1 Initialization
- 2 Variable to Factor message
- 3 Factor to Variable message

[Source: P. Gehler]

1. Initialization

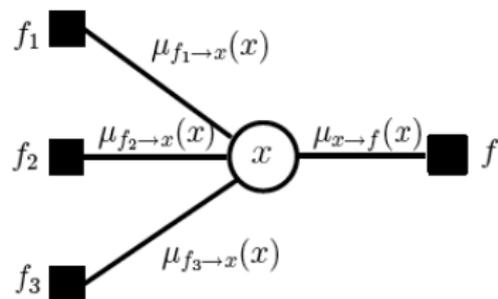
- Messages from extremal (simplicial) node factors are initialized to the factor (left)
- Messages from extremal (simplicial) variable nodes are set to unity (right)



[Source: P. Gehler]

2. Variable to Factor message

$$\mu_{x \rightarrow f}(x) = \prod_{g \in \{\text{ne}(x) \setminus f\}} \mu_{g \rightarrow x}(x)$$

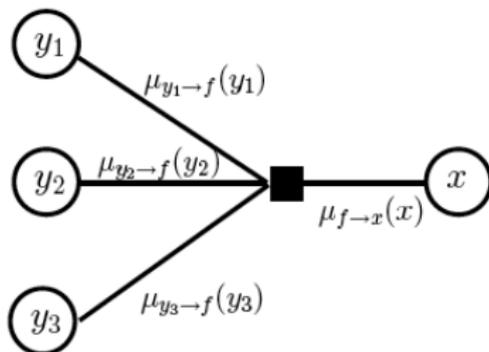


[Source: P. Gebler]

3. Factor to Variable message

- We sum over all states in the set of variables
- This explains the name for the algorithm (sum-product)

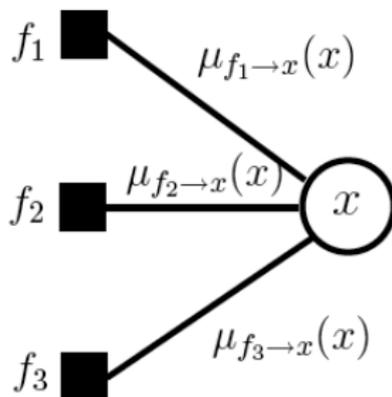
$$\mu_{f \rightarrow x}(x) = \sum_{y \in \mathcal{X}_f \setminus x} \phi_f(\mathcal{X}_f) \prod_{y \in \{\text{ne}(f) \setminus x\}} \mu_{y \rightarrow f}(y)$$



[Source: P. Gehler]

Marginal computation

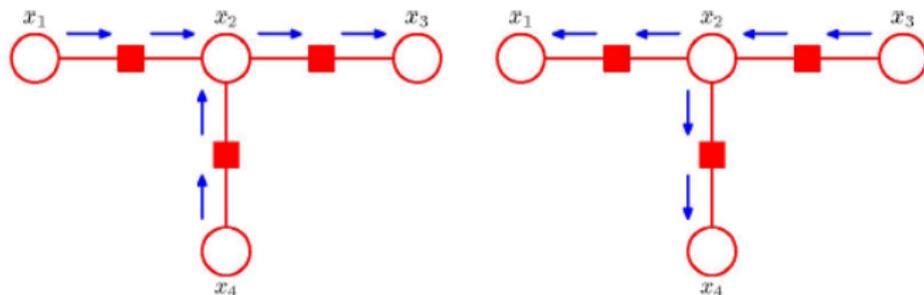
$$p(x) \propto \prod_{f \in \text{ne}(x)} \mu_{f \rightarrow x}(x)$$



[Source: P. Gehler]

Message Ordering

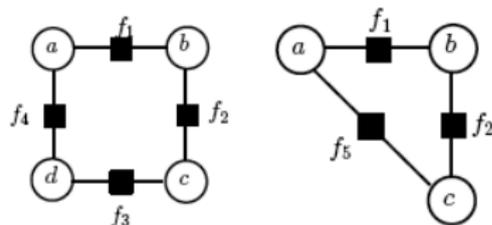
- ▶ Messages depend on previous computed messages
- ▶ Only extremal nodes/factors do not depend on other messages
- ▶ To compute all messages in the graph
 1. leaf-to-root: (pick root node, compute messages pointing towards root)
 2. root-to-leave: (compute messages pointing away from root)



[Source: P. Gehler]

Problems with loops

- Marginalizing over d introduces new link (changes graph structure – in contrast to singly connected graphs)



$$p(a, b, c, d) = f_1(a, b)f_2(b, c)f_3(c, d)f_4(d, a)$$

and marginal

$$p(a, b, c) = f_1(a, b)f_2(b, c) \underbrace{\sum_d f_3(c, d)f_4(d, a)}_{f_5(a, c)}$$

[Source: P. Gehler]

What to infer?

- ▶ Mean

$$\mathbb{E}_{p(x)}[x] = \sum_{x \in \mathcal{X}} xp(x)$$

- ▶ Mode

$$x^* = \operatorname{argmax}_{x \in \mathcal{X}} p(x)$$

- ▶ Conditional Distributions

$$p(x_i, x_j \mid x_k, x_l) \text{ or } p(x_i \mid x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$$

- ▶ Max-Marginals

$$x_i^* = \operatorname{argmax}_{x_i \in \mathcal{X}_i} p(x_i) = \dots \int dx_n \operatorname{argmax}_{x_i \in \mathcal{X}_i} \int_{(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)} p(x) dx_1$$

[Source: P. Gehler]

Computing the Partition Function

- ▶ The partition function ($p(x) = \frac{1}{Z} \prod_f \Phi_f(\mathcal{X}_f)$) (normalization constant) Z can be computed after the leaf-to-root step (no need for the root-to-leaf step) (choose any $x \in \mathcal{X}$)

$$Z = \sum_{\mathcal{X}} \prod_f \phi_f(\mathcal{X}_f) \quad (10)$$

$$= \sum_x \sum_{\mathcal{X} \setminus \{x\}} \prod_{f \in \text{ne}(x)} \prod_{f \notin \text{ne}(x)} \phi_f(\mathcal{X}_f) \quad (11)$$

$$= \sum_x \prod_{f \in \text{ne}(x)} \sum_{\mathcal{X} \setminus \{x\}} \prod_{f \notin \text{ne}(x)} \phi_f(\mathcal{X}_f) \quad (12)$$

$$= \sum_x \prod_{f \in \text{ne}(x)} \mu_{f \rightarrow x}(x) \quad (13)$$

[Source: P. Gehler]

- ▶ In large graphs, messages may become very small
- ▶ Work with log-messages instead $\lambda = \log \mu$
- ▶ Variable-to-factor messages

$$\mu_{x \rightarrow f}(x) = \prod_{g \in \{\text{ne}(x) \setminus f\}} \mu_{g \rightarrow x}(x)$$

then becomes

$$\lambda_{x \rightarrow f}(x) = \sum_{g \in \{\text{ne}(x) \setminus f\}} \lambda_{g \rightarrow x}(x)$$

[Source: P. Gehler]

- ▶ Work with log-messages instead $\lambda = \log \mu$
- ▶ Factor-to-Variable messages

$$\mu_{f \rightarrow x}(x) = \sum_{y \in \mathcal{X}_f \setminus x} \Phi_f(\mathcal{X}_f) \prod_{y \in \{\text{ne}(f) \setminus x\}} \mu_{y \rightarrow f}(y) \quad (16)$$

then becomes

$$\lambda_{f \rightarrow x}(x) = \log \left(\sum_{y \in \mathcal{X}_f \setminus x} \Phi(\mathcal{X}_f) \exp \left[\sum_{y \in \{\text{ne}(f) \setminus x\}} \lambda_{y \rightarrow f}(y) \right] \right) \quad (17)$$

[Source: P. Gehler]

- ▶ Log-Factor-to-Variable Message:

$$\lambda_{f \rightarrow x}(x) = \log \sum_{y \in \mathcal{X}_f \setminus x} \Phi_f(\mathcal{X}_f) \exp \sum_{y \in \{\text{ne}(f) \setminus x\}} \lambda_{y \rightarrow f}(y) \quad (18)$$

- ▶ large numbers lead to numerical instability
- ▶ Use the following equality

$$\log \sum_i \exp(v_i) = \alpha + \log \sum_i \exp(v_i - \alpha) \quad (19)$$

- ▶ With $\alpha = \max \lambda_{y \rightarrow f}(y)$

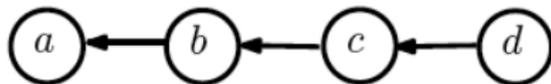
[Source: P. Gehler]

Finding the maximal state: Max-Product

- ▶ For a given distribution $p(x)$ find the most likely state:

$$x^* = \operatorname{argmax}_{x_1, \dots, x_n} p(x_1, \dots, x_n)$$

- ▶ Again use factorization structure to distribute the maximisation to local computations
- ▶ Example: chain



$$f(x_1, x_2, x_3, x_4) = \phi(x_1, x_2)\phi(x_2, x_3)\phi(x_3, x_1)$$

[Source: P. GeHLer]

Be careful: not maximal marginal states!

- ▶ The most likely state

$$x^* = \operatorname{argmax}_{x_1, \dots, x_n} p(x_1, \dots, x_n)$$

does not need to be the one for which the marginals are maximized:

- ▶ For all $i = 1, \dots, n$

$$x_i^* = \operatorname{argmax}_{x_i} p(x_i)$$

- ▶ Example:

	$x = 0$	$x = 1$
$y = 0$	0.3	0.4
$y = 1$	0.3	0.0

[Source: P. Gehler]

Example chain

$$\begin{aligned}\max_x f(x) &= \max_{x_1, x_2, x_3, x_4} \phi(x_1, x_2)\phi(x_2, x_3)\phi(x_3, x_4) \\ &= \max_{x_1, x_2, x_3} \phi(x_1, x_2)\phi(x_2, x_3) \underbrace{\max_{x_4} \phi(x_3, x_4)}_{\gamma(x_3)} \\ &= \max_{x_1, x_2} \phi(x_1, x_2) \underbrace{\max_{x_3} \phi(x_2, x_3)\gamma(x_3)}_{\gamma(x_2)} \\ &= \max_{x_1} \underbrace{\max_{x_2} \phi(x_1, x_2)\gamma(x_2)}_{\gamma(x_1)} \\ &= \max_{x_1} \gamma(x_1)\end{aligned}$$

[Source: P. Gehler]

Example chain

- ▶ Once computed the messages ($\gamma(\cdot)$) find the optimal values

$$x_1^* = \operatorname{argmax}_{x_1} \gamma(x_1)$$

$$x_2^* = \operatorname{argmax}_{x_2} \phi(x_1^*, x_2) \gamma(x_2)$$

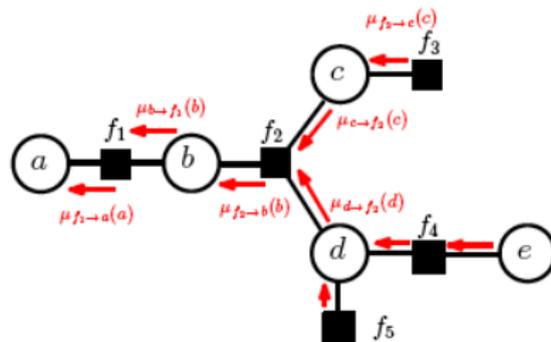
$$x_3^* = \operatorname{argmax}_{x_3} \phi(x_2^*, x_3) \gamma(x_3)$$

$$x_4^* = \operatorname{argmax}_{x_4} \phi(x_3^*, x_4) \gamma(x_4)$$

- ▶ this is called **backtracking** (an application of dynamic programming)
- ▶ can choose arbitrary start point

[Source: P. Gehler]

- Spot the messages:



$$\begin{aligned}
 \max_x f(x) &= \max_{a,b,c,d,e} f_1(a,b)f_2(b,c,d)f_3(c)f_4(d,e)f_5(d) \\
 &= \max_a \mu_{f_2 \rightarrow a}(a)
 \end{aligned}$$

[Source: P. Gehler]

Max-Product Algorithm

Pick any variable as root and

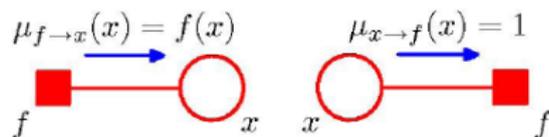
- 1 Initialisation (same as sum-product)
- 2 Variable to Factor message (same as sum-product)
- 3 Factor to Variable message

Then compute the maximal state

[Source: P. Gehler]

1. Initialization

- Messages from extremal node factors are initialized to the factor
- Messages from extremal variable nodes are set to unity



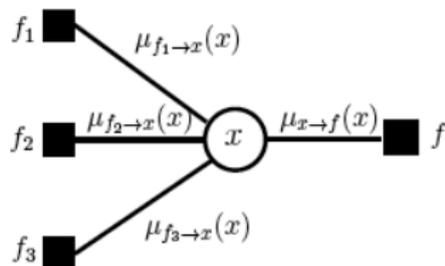
- Same as sum product

[Source: P. Gehler]

2. Variable to Factor message

- Same as for sum-product

$$\mu_{x \rightarrow f}(x) = \prod_{g \in \{\text{ne}(x) \setminus f\}} \mu_{g \rightarrow x}(x)$$

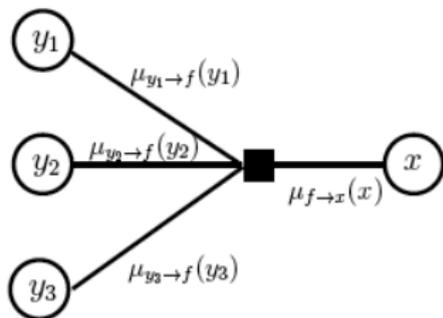


[Source: P. Gehler]

3. Factor to Variable message

- Different message than in sum-product
- This is now a max-product

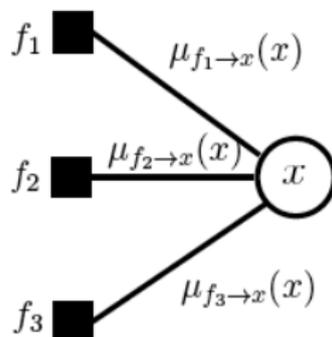
$$\mu_{f \rightarrow x}(x) = \max_{y \in \mathcal{X}_f \setminus x} \phi_f(\mathcal{X}_f) \prod_{y \in \{\text{ne}(f) \setminus x\}} \mu_{y \rightarrow f}(y)$$



[Source: P. Gehler]

Maximal state of Variable

$$x^* = \operatorname{argmax}_x \prod_{f \in \operatorname{ene}(x)} \mu_{f \rightarrow x}(x)$$



- This does not work with loops
- Same problem as the sum product algorithm

[Source: P. Gehler]

Algorithm 1: Belief Propagation on Trees

```
1:  $(\log Z, \mu) = \text{BeliefPropagation}(V, F, E, E)$ 
2: Input:
3:    $(V, F, E)$ , tree-structured factor graph,
4:    $E$ , energies  $E_F$  for all  $F \in F$ .
5: Output:
6:    $\log Z$ , log partition function of  $p(y)$ ,
7:    $\mu$ , marginal distributions  $\mu_F$  for all  $F \in F$ .
8: Algorithm:
9: Fix an element of  $V$  arbitrarily as tree root
10: Compute leaf-to-root order  $R$  as sequence of directed
11:   edges of  $E$ 
12: for  $i = 1, \dots, |R|$  do
13:   if  $(v, F) = R(i)$  is variable-to-factor edge then
14:     Compute  $q_{Y_i \rightarrow F}$  using (3.2)
15:   else
16:      $(F, v) = R(i)$  is factor-to-variable edge
17:     Compute  $r_{F \rightarrow Y_i}$  using (3.3)
18:   end if
19: end for
20: Compute  $\log Z$  by (3.4)
21: Compute root-to-leaf order  $R' = \text{reverse}(R)$ 
22: for  $i = 1, \dots, |R'|$  do
23:   if  $(v, F) = R'(i)$  is variable-to-factor edge then
24:     Compute  $q_{Y_i \rightarrow F}$  using (3.2)
25:     Compute  $\mu_F$  using (3.5)
26:   else
27:      $(F, v) = R'(i)$  is factor-to-variable edge
28:     Compute  $r_{F \rightarrow Y_i}$  using (3.3)
29:     Compute  $p(y)$  using (3.6)
30:   end if
31: end for
```

- Keep on doing this iterations, i.e., loopy BP
- The problem with loopy BP is that it is not guaranteed to converge
- Message-passing algorithms based on LP relaxations have been developed
- These methods are guaranteed to converge
- Perform much better in practice

Algorithm 2: Loopy Belief Propagation (sum-product)

```
1:  $(\log Z, \mu) = \text{SUMPRODUCTLOOPYBP}(V, \mathcal{F}, \mathcal{E}, E, \varepsilon, T)$ 
2: Input:
3:    $(V, \mathcal{F}, \mathcal{E})$ , factor graph,
4:    $E$ , energies  $E_F$  for all  $F \in \mathcal{F}$ ,
5:    $\varepsilon$ , convergence tolerance,
6:    $T$ , maximum number of iterations.
7: Output:
8:    $\log Z$ , approximate log partition function of  $p(y)$ ,
9:    $\mu$ , approximate marginal distributions  $\mu_F$  for all  $F \in \mathcal{F}$ .
10: Algorithm:
11:  $q_{Y_i \rightarrow F}(y_i) \leftarrow 0$ , for all  $(i, F) \in \mathcal{E}, y_i \in \mathcal{Y}_i$ 
12:  $\mu_F(y_F) \leftarrow 0$ , for all  $F \in \mathcal{F}, y_F \in \mathcal{Y}_F$ 
13: for  $t = 1, \dots, T$  do
14:   for  $(v, F) \in \mathcal{F}$  do
15:     for  $y_i \in \mathcal{Y}_i$  do
16:       Compute  $r_{F \rightarrow Y_i}(y_i)$  using (3.3)
17:     end for
18:   end for
19:   for  $(v, F) \in \mathcal{F}$  do
20:     for  $y_i \in \mathcal{Y}_i$  do
21:       Compute  $q_{Y_i \rightarrow F}(y_i)$  using (3.9) to (3.11)
22:     end for
23:   end for
24:   Compute approximate marginals  $\mu'$  using (3.12) to (3.17)
25:    $u \leftarrow \|\mu' - \mu\|_\infty$  {Measure change in beliefs}
26:    $\mu \leftarrow \mu'$ 
27:   if  $u \leq \varepsilon$  then
28:     break {Converged}
29:   end if
30: end for
31: Compute  $\log Z$  using (3.18)
```



Ways to get an approximate solution typically

- Dynamic programming approximations
- Sampling
- Simulated annealing
- Graph-cuts: imposes restrictions on the type of pairwise cost functions
- Message passing: iterative algorithms that pass messages between nodes in the graph. Which graph?

Inference with graph cuts

Submodular Functions

- A Pseudo-boolean function $f : \{0, 1\}^n \rightarrow \Re$ is **submodular** if

$$f(A) + f(B) \geq \underbrace{f(A \vee B)}_{OR} + \underbrace{f(A \wedge B)}_{AND} \quad \forall A, B \in \{0, 1\}^n$$

- Example: $n = 2$, $A = [1, 0]$, $B = [0, 1]$

$$f([1, 0]) + f([0, 1]) \geq f([1, 1]) + f([0, 0])$$

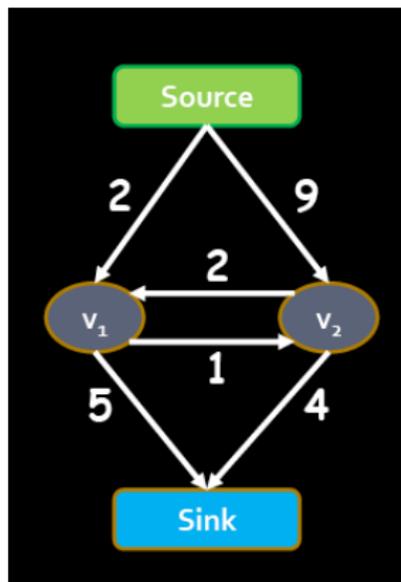
- Sum of submodular functions is submodular \rightarrow Easy to proof.
- Some energies in computer vision can be submodular

Minimizing submodular Functions

- Pairwise submodular functions can be transformed to st-mincut/max-flow [Hammer, 65].
- Very low running time $\sim \mathcal{O}(n)$

The ST-mincut problem

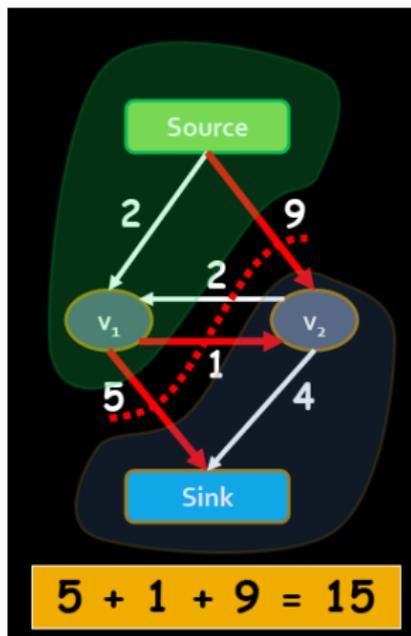
- Suppose we have a graph $G = \{V, E, C\}$, with vertices V , Edges E and costs C .



[Source: P. Kohli]

The ST-mincut problem

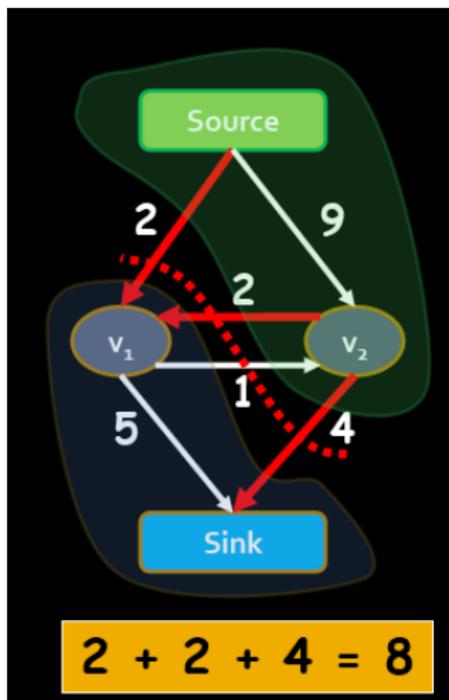
- An st-cut (S,T) divides the nodes between source and sink.
- The cost of a st-cut is the sum of cost of all edges going from S to T



[Source: P. Kohli]

The ST-mincut problem

- The st-mincut is the st-cut with the minimum cost

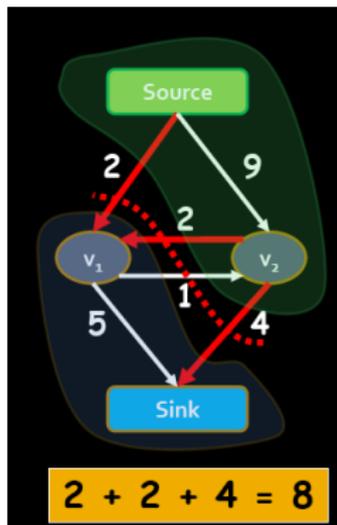


[Source: P. Kohli]

Back to our energy minimization

Construct a graph such that

- 1 Any st-cut corresponds to an assignment of x
- 2 The cost of the cut is equal to the energy of x : $E(x)$



[Source: P. Kohli]

St-mincut and Energy Minimization

$$E(\mathbf{x}) = \sum_i \theta_i(x_i) + \sum_{i,j} \theta_{ij}(x_i, x_j)$$

For all ij $\theta_{ij}(0,1) + \theta_{ij}(1,0) \geq \theta_{ij}(0,0) + \theta_{ij}(1,1)$



Equivalent (transformable)

$$E(\mathbf{x}) = \sum_i c_i x_i + \sum_{i,j} c_{ij} x_i(1-x_j) \quad c_{ij} \geq 0$$

[Source: P. Kohli]

How are they equivalent?

$$A = \theta_{ij}(0,0)$$

$$B = \theta_{ij}(0,1)$$

$$C = \theta_{ij}(1,0)$$

$$D = \theta_{ij}(1,1)$$

		x_j							
		0	1	0	1				
x_i	0	A	B	= A +	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td>0</td><td>0</td></tr> <tr><td>1</td><td>C-A</td></tr> </table>	0	0	1	C-A
	0	0							
1	C-A								
1	C	D	+	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td>0</td><td>0</td></tr> <tr><td>1</td><td>0</td></tr> </table>	0	0	1	0	
0	0								
1	0								

		x_j							
		0	1	0	1				
x_i	0	0	D-C	+	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td>0</td><td>B+C-A-D</td></tr> <tr><td>1</td><td>0</td></tr> </table>	0	B+C-A-D	1	0
	0	B+C-A-D							
1	0								
1	0	D-C							

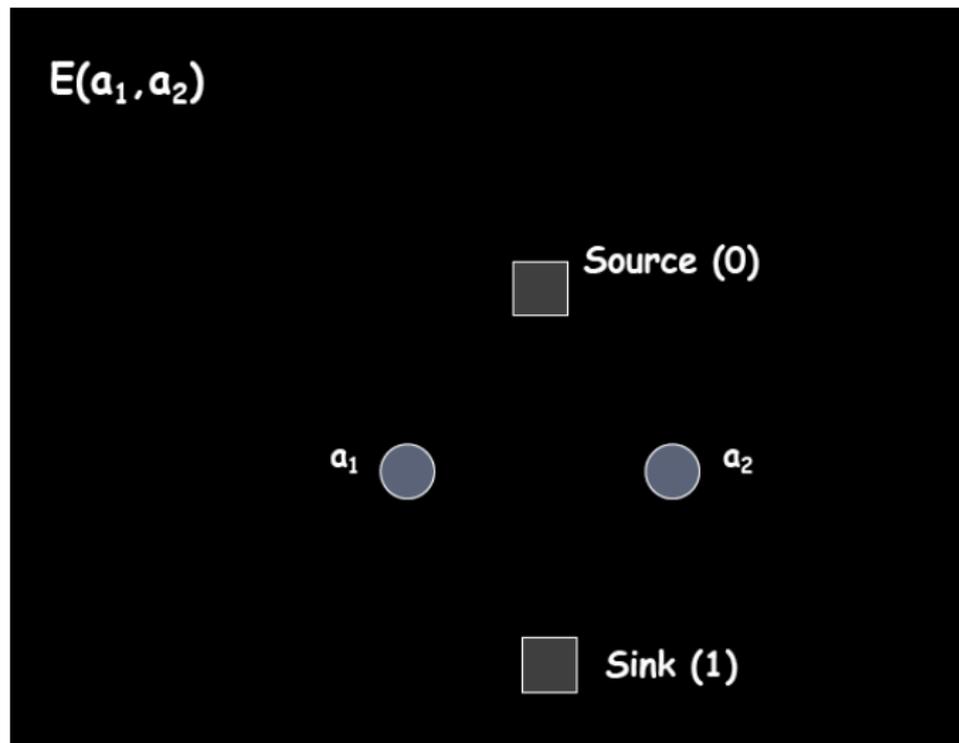
if $x_1=1$ add C-
A
 if $x_2 = 1$ add
D-C

$$\begin{aligned}
 \theta_{ij}(x_i, x_j) &= \theta_{ij}(0,0) \\
 &+ (\theta_{ij}(1,0) - \theta_{ij}(0,0)) x_i + (\theta_{ij}(1,0) - \theta_{ij}(0,0)) x_j \\
 &+ (\theta_{ij}(1,0) + \theta_{ij}(0,1) - \theta_{ij}(0,0) - \theta_{ij}(1,1)) (1-x_i) x_j
 \end{aligned}$$

$B+C-A-D \geq 0$ is true from the submodularity of θ_{ij}

[Source: P. Kohli]

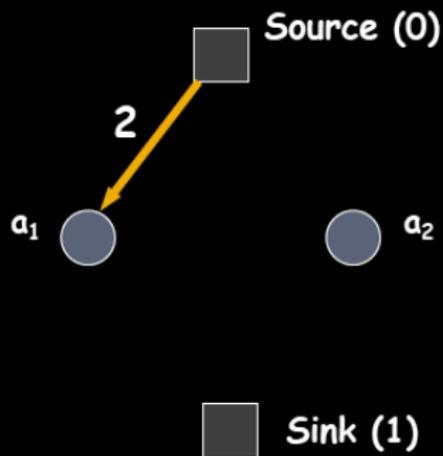
Graph Construction



[Source: P. Kohli]

Graph Construction

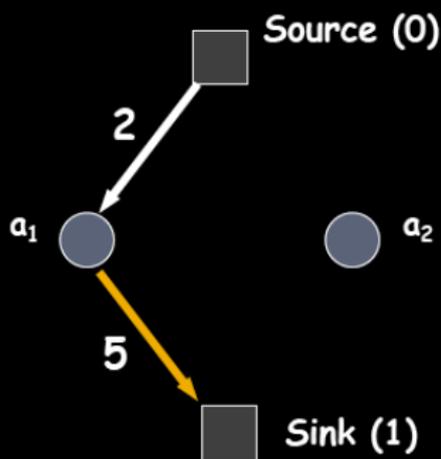
$$E(\mathbf{a}_1, \mathbf{a}_2) = 2\mathbf{a}_1$$



[Source: P. Kohli]

Graph Construction

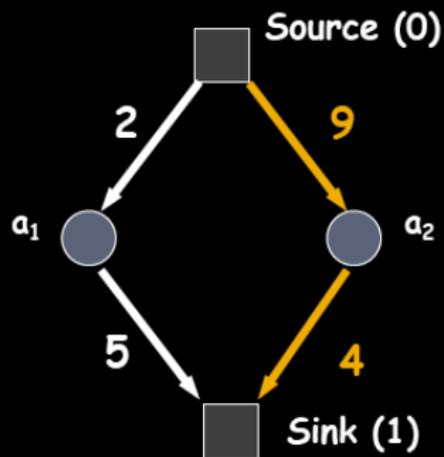
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1$$



[Source: P. Kohli]

Graph Construction

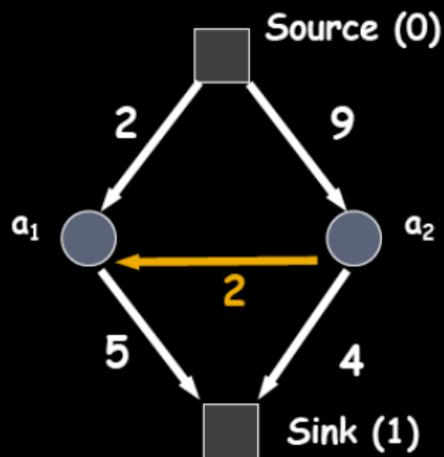
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2$$



[Source: P. Kohli]

Graph Construction

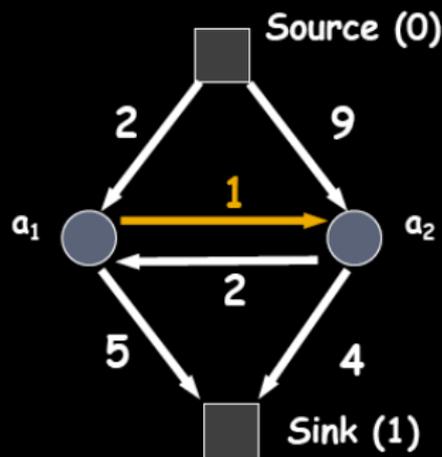
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2 + 2a_1\bar{a}_2$$



[Source: P. Kohli]

Graph Construction

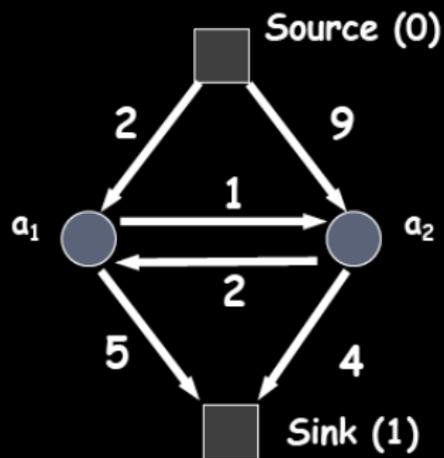
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2 + 2a_1\bar{a}_2 + \bar{a}_1a_2$$



[Source: P. Kohli]

Graph Construction

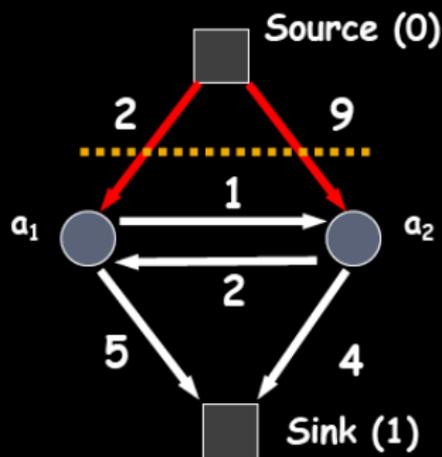
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2 + 2a_1\bar{a}_2 + \bar{a}_1a_2$$



[Source: P. Kohli]

Graph Construction

$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2 + 2a_1\bar{a}_2 + \bar{a}_1a_2$$



Cost of cut = 11

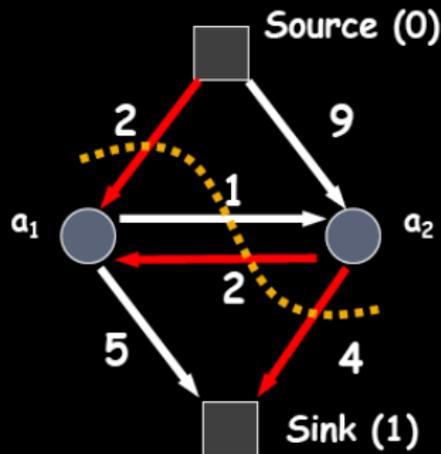
$$a_1 = 1 \quad a_2 = 1$$

$$E(1, 1) = 11$$

[Source: P. Kohli]

Graph Construction

$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2 + 2a_1\bar{a}_2 + \bar{a}_1a_2$$



st-minicut cost = 8

$$a_1 = 1 \quad a_2 = 0$$

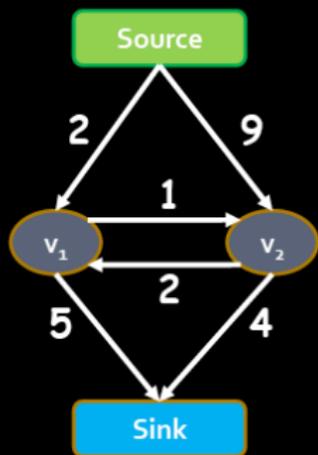
$$E(1, 0) = 8$$

[Source: P. Kohli]

How to compute the St-mincut?

Solve the dual **maximum flow** problem

Compute the maximum flow between Source and Sink s.t.



Edges: Flow < Capacity

Nodes: Flow in = Flow out

Min-cut \ Max-flow Theorem

In every network, the maximum flow equals the cost of the st-mincut

Assuming non-negative capacity

[Source: P. Kohli]

How does the code look like

```
Graph *g;
```

```
For all pixels p
```

```
    /* Add a node to the graph */
```

```
    nodeID(p) = g->add_node();
```

```
    /* Set cost of terminal edges */
```

```
    set_weights(nodeID(p), fgCost(p), bgCost(p));
```

```
end
```

```
for all adjacent pixels p,q
```

```
    add_weights(nodeID(p), nodeID(q), cost(p,q));
```

```
end
```

```
g->compute_maxflow();
```

```
label_p = g->is_connected_to_source(nodeID(p));
```

```
// is the label of pixel p (0 or 1)
```



Source (0)



Sink (1)

[Source: P. Kohli]

How does the code look like

```
Graph *g;
```

```
For all pixels p
```

```
/* Add a node to the graph */
```

```
nodeID(p) = g->add_node();
```

```
/* Set cost of terminal edges */
```

```
set_weights(nodeID(p), fgCost(p), bgCost(p));
```

```
end
```

```
for all adjacent pixels p,q
```

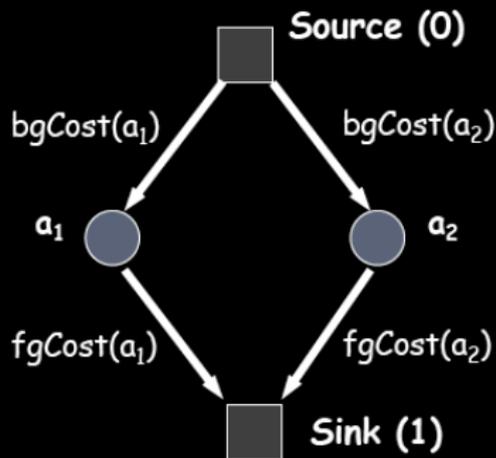
```
add_weights(nodeID(p), nodeID(q), cost(p,q));
```

```
end
```

```
g->compute_maxflow();
```

```
label_p = g->is_connected_to_source(nodeID(p));
```

```
// is the label of pixel p (0 or 1)
```



[Source: P. Kohli]

How does the code look like

```
Graph *g;
```

```
For all pixels p
```

```
/* Add a node to the graph */
```

```
nodeID(p) = g->add_node();
```

```
/* Set cost of terminal edges */
```

```
set_weights(nodeID(p), fgCost(p), bgCost(p));
```

```
end
```

```
for all adjacent pixels p,q
```

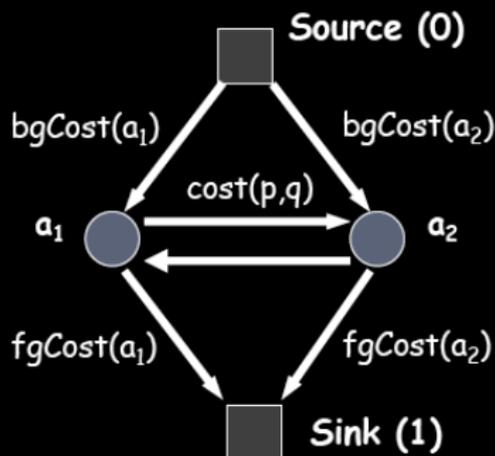
```
add_weights(nodeID(p), nodeID(q), cost(p,q));
```

```
end
```

```
g->compute_maxflow();
```

```
label_p = g->is_connected_to_source(nodeID(p));
```

```
// is the label of pixel p (0 or 1)
```



[Source: P. Kohli]

How does the code look like

```
Graph *g;
```

```
For all pixels p
```

```
/* Add a node to the graph */
```

```
nodeID(p) = g->add_node();
```

```
/* Set cost of terminal edges */
```

```
set_weights(nodeID(p), fgCost(p), bgCost(p));
```

```
end
```

```
for all adjacent pixels p,q
```

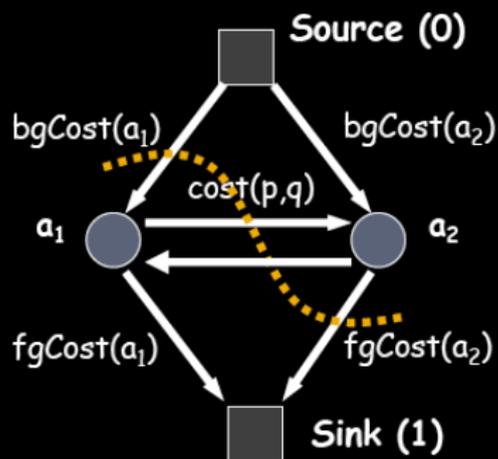
```
add_weights(nodeID(p), nodeID(q), cost(p,q));
```

```
end
```

```
g->compute_maxflow();
```

```
label_p = g->is_connected_to_source(nodeID(p));
```

```
// is the label of pixel p (0 or 1)
```



$$a_1 = \text{bg} \quad a_2 = \text{fg}$$

[Source: P. Kohli]

Graph cuts for multi-label problems

- Exact Transformation to QPBF [Roy and Cox 98] [Ishikawa 03] [Schlesinger et al. 06] [Ramalingam et al. 08]

So what is the problem?

$E_m(y_1, y_2, \dots, y_n)$	→	$E_b(x_1, x_2, \dots, x_m)$
$y_i \in L = \{l_1, l_2, \dots, l_k\}$		$x_i \in L = \{0, 1\}$
Multi-label Problem		Binary label Problem

such that:

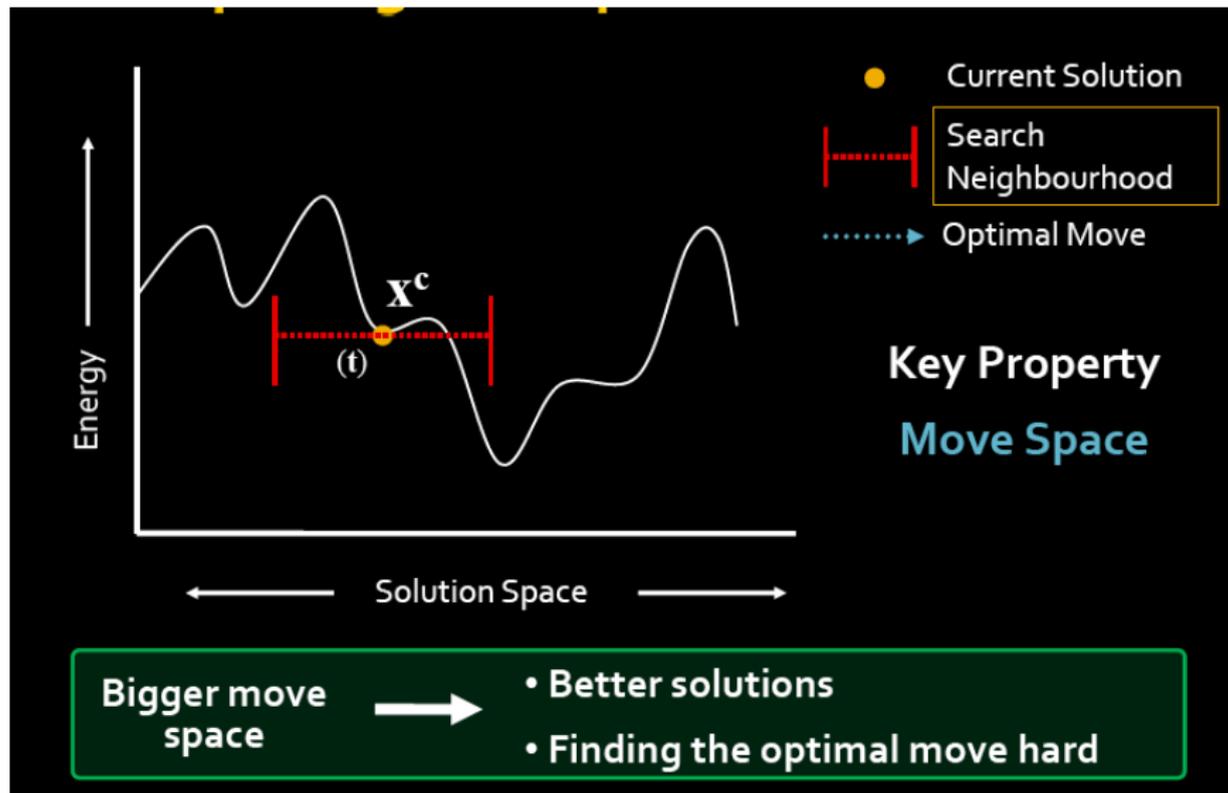
Let Y and X be the set of feasible solutions, then

1. One-One encoding function $T: X \rightarrow Y$
2. $\arg \min E_m(y) = T(\arg \min E_b(x))$

- Very high computational cost

[Source: P. Kohli]

Computing the Optimal Move

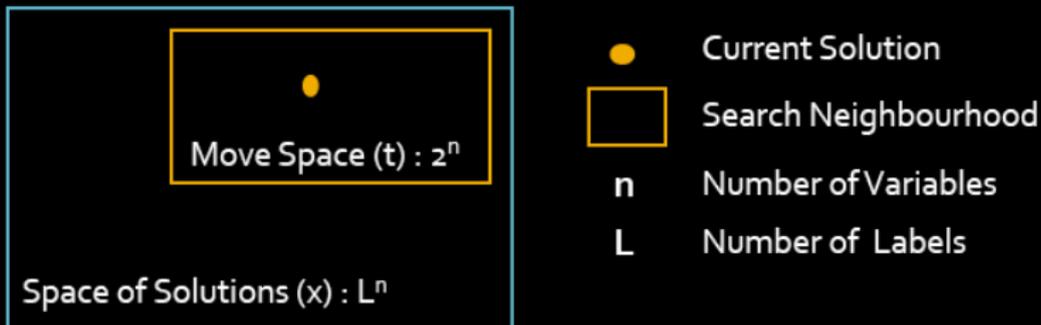


[Source: P. Kohli]

Minimizing Pairwise Functions

[Boykov Veksler and Zabih, PAMI 2001]

- Series of locally optimal moves
- Each move reduces energy
- Optimal move by minimizing submodular function



[Source: P. Kohli]

- Consider pairwise MRFs

$$E(f) = \sum_{\{p,q\} \in \mathcal{N}} V_{p,q}(f_p, f_q) + \sum_p D_p(f_p)$$

with \mathcal{N} defining the interactions between nodes, e.g., pixels

- D_p non-negative, but arbitrary.

- Consider pairwise MRFs

$$E(f) = \sum_{\{p,q\} \in \mathcal{N}} V_{p,q}(f_p, f_q) + \sum_p D_p(f_p)$$

with \mathcal{N} defining the interactions between nodes, e.g., pixels

- D_p non-negative, but arbitrary.
- This is the graph-cuts notation.
- Important to notice it's the same thing.

- Consider pairwise MRFs

$$E(f) = \sum_{\{p,q\} \in \mathcal{N}} V_{p,q}(f_p, f_q) + \sum_p D_p(f_p)$$

with \mathcal{N} defining the interactions between nodes, e.g., pixels

- D_p non-negative, but arbitrary.
- This is the graph-cuts notation.
- Important to notice it's the same thing.

Two general classes of pairwise interactions

- **Metric** if it satisfies for any set of labels α, β, γ

$$V(\alpha, \beta) = 0 \iff \alpha = \beta$$

$$V(\alpha, \beta) = V(\beta, \alpha) \geq 0$$

$$V(\alpha, \beta) \leq V(\alpha, \gamma) + V(\gamma, \beta)$$

- **Semi-metric** if it satisfies for any set of labels α, β, γ

$$V(\alpha, \beta) = 0 \iff \alpha = \beta$$

$$V(\alpha, \beta) = V(\beta, \alpha) \geq 0$$

Two general classes of pairwise interactions

- **Metric** if it satisfies for any set of labels α, β, γ

$$V(\alpha, \beta) = 0 \iff \alpha = \beta$$

$$V(\alpha, \beta) = V(\beta, \alpha) \geq 0$$

$$V(\alpha, \beta) \leq V(\alpha, \gamma) + V(\gamma, \beta)$$

- **Semi-metric** if it satisfies for any set of labels α, β, γ

$$V(\alpha, \beta) = 0 \iff \alpha = \beta$$

$$V(\alpha, \beta) = V(\beta, \alpha) \geq 0$$

Examples for 1D label set

- Truncated quadratic is a semi-metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|^2)$$

with K a constant.

- Truncated absolute distance is a metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|)$$

with K a constant.

Examples for 1D label set

- Truncated quadratic is a semi-metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|^2)$$

with K a constant.

- Truncated absolute distance is a metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|)$$

with K a constant.

- For multi-dimensional, replace $|\cdot|$ by any norm.

Examples for 1D label set

- Truncated quadratic is a semi-metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|^2)$$

with K a constant.

- Truncated absolute distance is a metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|)$$

with K a constant.

- For multi-dimensional, replace $|\cdot|$ by any norm.
- Potts model is a metric

$$V(\alpha, \beta) = K \cdot T(\alpha \neq \beta)$$

with $T(\cdot) = 1$ if the argument is true and 0 otherwise.

Examples for 1D label set

- Truncated quadratic is a semi-metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|^2)$$

with K a constant.

- Truncated absolute distance is a metric

$$V(\alpha, \beta) = \min(K, |\alpha - \beta|)$$

with K a constant.

- For multi-dimensional, replace $|\cdot|$ by any norm.
- Potts model is a metric

$$V(\alpha, \beta) = K \cdot T(\alpha \neq \beta)$$

with $T(\cdot) = 1$ if the argument is true and 0 otherwise.

Binary Moves

- $\alpha - \beta$ **moves** works for semi-metrics
- α **expansion** works for V being a metric

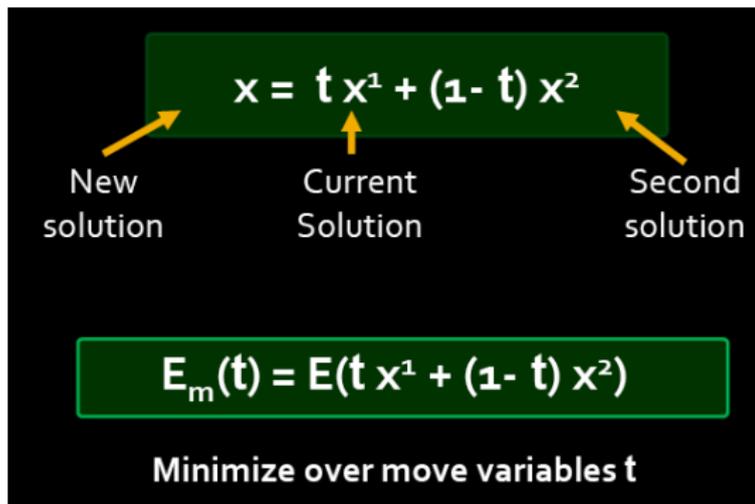
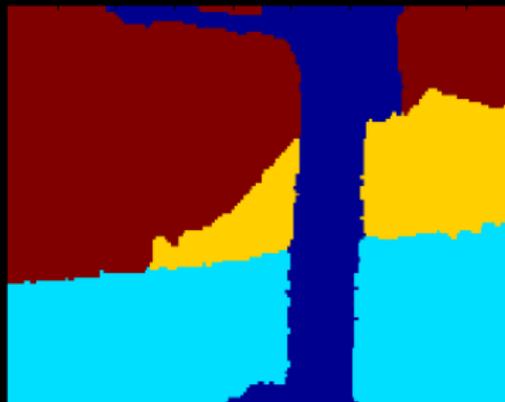


Figure: Figure from P. Kohli tutorial on graph-cuts

- For certain x^1 and x^2 , the move energy is sub-modular QPBF

- Variables labeled α , β can swap their labels

Swap Sky, House



[Source: P. Kohli]

- Variables labeled α, β can swap their labels
- Move energy is submodular if:
 - Unary Potentials: Arbitrary
 - Pairwise potentials: **Semi-metric**

$$\begin{array}{c} \theta_{ij}(l_a, l_b) \geq 0 \\ \theta_{ij}(l_a, l_b) = 0 \iff a = b \end{array}$$

Examples: **Potts model, Truncated Convex**

[Source: P. Kohli]

Expansion Move

- Variables take label α or retain current label



Status: Expanded Sky to Tree



[Source: P. Kohli]

- Variables take label α or retain current label
- Move energy is submodular if:
 - Unary Potentials: Arbitrary
 - Pairwise potentials: **Metric**

Semi metric
+
Triangle Inequality

$$\theta_{ij}(l_a, l_b) + \theta_{ij}(l_b, l_c) \geq \theta_{ij}(l_a, l_c)$$

Examples: **Potts model, Truncated linear**

Cannot solve truncated quadratic

[Source: P. Kohli]

More formally

- Any labeling can be uniquely represented by a partition of image pixels $\mathbf{P} = \{\mathcal{P}_l | l \in \mathcal{L}\}$, where $\mathcal{P}_l = \{p \in \mathcal{P} | f_p = l\}$ is a subset of pixels assigned label l .
- There is a one to one correspondence between labelings f and partitions \mathcal{P} .

More formally

- Any labeling can be uniquely represented by a partition of image pixels $\mathbf{P} = \{\mathcal{P}_l | l \in \mathcal{L}\}$, where $\mathcal{P}_l = \{p \in \mathcal{P} | f_p = l\}$ is a subset of pixels assigned label l .
- There is a one to one correspondence between labelings f and partitions \mathcal{P} .
- Given a pair of labels α, β , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an $\alpha - \beta$ **swap** if $\mathcal{P}_l = \mathcal{P}'_l$ for any label $l \neq \alpha, \beta$.

More formally

- Any labeling can be uniquely represented by a partition of image pixels $\mathbf{P} = \{\mathcal{P}_l | l \in \mathcal{L}\}$, where $\mathcal{P}_l = \{p \in \mathcal{P} | f_p = l\}$ is a subset of pixels assigned label l .
- There is a one to one correspondence between labelings f and partitions \mathcal{P} .
- Given a pair of labels α, β , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an $\alpha - \beta$ **swap** if $\mathcal{P}_l = \mathcal{P}'_l$ for any label $l \neq \alpha, \beta$.
- The only difference between \mathcal{P} and \mathcal{P}' is that some pixels that were labeled in \mathcal{P} are now labeled in \mathcal{P}' , and vice-versa.

More formally

- Any labeling can be uniquely represented by a partition of image pixels $\mathbf{P} = \{\mathcal{P}_l | l \in \mathcal{L}\}$, where $\mathcal{P}_l = \{p \in \mathcal{P} | f_p = l\}$ is a subset of pixels assigned label l .
- There is a one to one correspondence between labelings f and partitions \mathcal{P} .
- Given a pair of labels α, β , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an $\alpha - \beta$ **swap** if $\mathcal{P}_l = \mathcal{P}'_l$ for any label $l \neq \alpha, \beta$.
- The only difference between \mathcal{P} and \mathcal{P}' is that some pixels that were labeled in \mathcal{P} are now labeled in \mathcal{P}' , and vice-versa.
- Given a label l , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an α -**expansion** if $\mathcal{P}_\alpha \subset \mathcal{P}'_\alpha$ and $\mathcal{P}'_l \subset \mathcal{P}_l$.

More formally

- Any labeling can be uniquely represented by a partition of image pixels $\mathbf{P} = \{\mathcal{P}_l | l \in \mathcal{L}\}$, where $\mathcal{P}_l = \{p \in \mathcal{P} | f_p = l\}$ is a subset of pixels assigned label l .
- There is a one to one correspondence between labelings f and partitions \mathcal{P} .
- Given a pair of labels α, β , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an $\alpha - \beta$ **swap** if $\mathcal{P}_l = \mathcal{P}'_l$ for any label $l \neq \alpha, \beta$.
- The only difference between \mathcal{P} and \mathcal{P}' is that some pixels that were labeled in \mathcal{P} are now labeled in \mathcal{P}' , and vice-versa.
- Given a label l , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an α -**expansion** if $\mathcal{P}_\alpha \subset \mathcal{P}'_\alpha$ and $\mathcal{P}'_l \subset \mathcal{P}_l$.
- An α -**expansion** move allows any set of image pixels to change their labels to α .

More formally

- Any labeling can be uniquely represented by a partition of image pixels $\mathbf{P} = \{\mathcal{P}_l | l \in \mathcal{L}\}$, where $\mathcal{P}_l = \{p \in \mathcal{P} | f_p = l\}$ is a subset of pixels assigned label l .
- There is a one to one correspondence between labelings f and partitions \mathcal{P} .
- Given a pair of labels α, β , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an $\alpha - \beta$ **swap** if $\mathcal{P}_l = \mathcal{P}'_l$ for any label $l \neq \alpha, \beta$.
- The only difference between \mathcal{P} and \mathcal{P}' is that some pixels that were labeled in \mathcal{P} are now labeled in \mathcal{P}' , and vice-versa.
- Given a label l , a move from a partition \mathcal{P} (labeling f) to a new partition \mathcal{P}' (labeling f') is called an α -**expansion** if $\mathcal{P}_\alpha \subset \mathcal{P}'_\alpha$ and $\mathcal{P}'_l \subset \mathcal{P}_l$.
- An α -**expansion** move allows any set of image pixels to change their labels to α .

Example



Figure: (a) Current partition (b) local move (c) $\alpha - \beta$ -swap (d) α -expansion.

1. Start with an arbitrary labeling f
 2. Set `success := 0`
 3. For each pair of labels $\{\alpha, \beta\} \subset \mathcal{L}$
 - 3.1. Find $\hat{f} = \operatorname{argmin} E(f')$ among f' within one α - β swap of f
 - 3.2. If $E(\hat{f}) < E(f)$, set $f := \hat{f}$ and `success := 1`
 4. If `success = 1` goto 2
 5. Return f
-

1. Start with an arbitrary labeling f
2. Set `success := 0`
3. For each label $\alpha \in \mathcal{L}$
 - 3.1. Find $\hat{f} = \operatorname{argmin} E(f')$ among f' within one α -expansion of f
 - 3.2. If $E(\hat{f}) < E(f)$, set $f := \hat{f}$ and `success := 1`
4. If `success = 1` goto 2
5. Return f

Finding optimal Swap move

- Given an input labeling f (partition \mathcal{P}) and a pair of labels α, β we want to find a labeling \hat{f} that minimizes E over all labelings within one $\alpha - \beta$ -swap of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_{\alpha\beta} = (\mathcal{V}_{\alpha\beta}, \mathcal{E}_{\alpha\beta})$.

Finding optimal Swap move

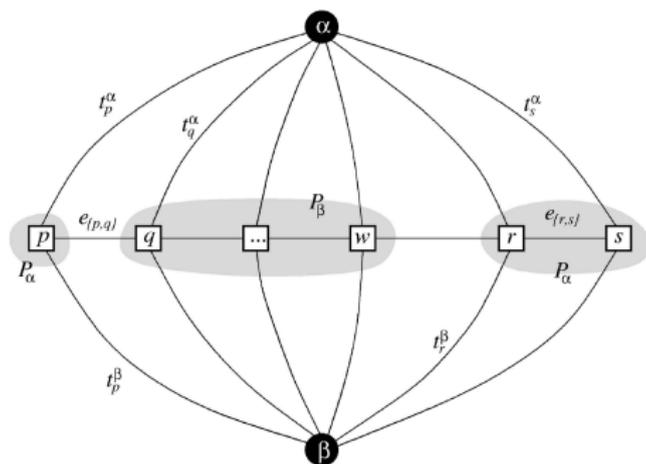
- Given an input labeling f (partition \mathcal{P}) and a pair of labels α, β we want to find a labeling \hat{f} that minimizes E over all labelings within one $\alpha - \beta$ -swap of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_{\alpha\beta} = (\mathcal{V}_{\alpha\beta}, \mathcal{E}_{\alpha\beta})$.
- The structure of this graph is dynamically determined by the current partition \mathcal{P} and by the labels α, β .

Finding optimal Swap move

- Given an input labeling f (partition \mathcal{P}) and a pair of labels α, β we want to find a labeling \hat{f} that minimizes E over all labelings within one $\alpha - \beta$ -swap of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_{\alpha\beta} = (\mathcal{V}_{\alpha\beta}, \mathcal{E}_{\alpha\beta})$.
- The structure of this graph is dynamically determined by the current partition \mathcal{P} and by the labels α, β .

Graph Construction

- The set of vertices includes the two terminals α and β , as well as image pixels p in the sets \mathcal{P}_α and \mathcal{P}_β (i.e., $f_p \in \{\alpha, \beta\}$).
- Each pixel $p \in \mathcal{P}_{\alpha\beta}$ is connected to the terminals α and β , called t -links.
- Each set of pixels $p, q \in \mathcal{P}_{\alpha\beta}$ which are neighbors is connected by an edge $e_{p,q}$



edge	weight	for
t_p^α	$D_p(\alpha) + \sum_{\substack{q \in \mathcal{N}_p \\ q \notin \mathcal{P}_{\alpha\beta}}} V(\alpha, f_q)$	$p \in \mathcal{P}_{\alpha\beta}$
t_p^β	$D_p(\beta) + \sum_{\substack{q \in \mathcal{N}_p \\ q \notin \mathcal{P}_{\alpha\beta}}} V(\beta, f_q)$	$p \in \mathcal{P}_{\alpha\beta}$
$e_{\{p,q\}}$	$V(\alpha, \beta)$	$\{p,q\} \in \mathcal{N}$ $p, q \in \mathcal{P}_{\alpha\beta}$

Computing the Cut

- Any cut must have a single t -link not cut.
- This defines a labeling

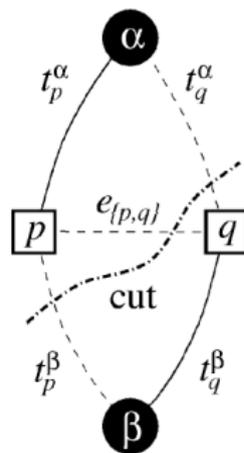
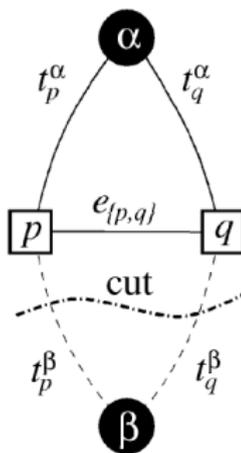
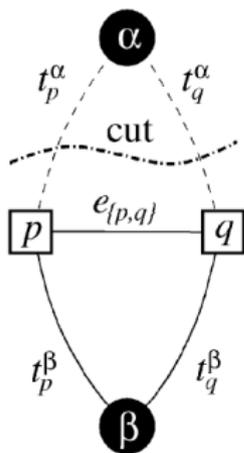
$$f_p^{\mathcal{C}} = \begin{cases} \alpha & \text{if } t_p^\alpha \in \mathcal{C} \text{ for } p \in \mathcal{P}_{\alpha\beta} \\ \beta & \text{if } t_p^\beta \in \mathcal{C} \text{ for } p \in \mathcal{P}_{\alpha\beta} \\ f_p & \text{for } p \in \mathcal{P}, p \notin \mathcal{P}_{\alpha\beta}. \end{cases}$$

- There is a one-to-one correspondences between a cut and a labeling.
- The energy of the cut is the energy of the labeling.
- See Boykov et al, "*fast approximate energy minimization via graph cuts*" PAMI 2001.

Properties

- For any cut, then

- (a) If $t_p^\alpha, t_q^\alpha \in \mathcal{C}$ then $e_{\{p,q\}} \notin \mathcal{C}$.
- (b) If $t_p^\beta, t_q^\beta \in \mathcal{C}$ then $e_{\{p,q\}} \notin \mathcal{C}$.
- (c) If $t_p^\beta, t_q^\alpha \in \mathcal{C}$ then $e_{\{p,q\}} \in \mathcal{C}$.
- (d) If $t_p^\alpha, t_q^\beta \in \mathcal{C}$ then $e_{\{p,q\}} \in \mathcal{C}$.



Finding the optimal α expansion

- Given an input labeling f (partition \mathcal{P}) and a label α we want to find a labeling \hat{f} that minimizes E over all labelings within one α -expansion of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_\alpha = (\mathcal{V}_\alpha, \mathcal{E}_\alpha)$.

Finding the optimal α expansion

- Given an input labeling f (partition \mathcal{P}) and a label α we want to find a labeling \hat{f} that minimizes E over all labelings within one α -expansion of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_\alpha = (\mathcal{V}_\alpha, \mathcal{E}_\alpha)$.
- The structure of this graph is dynamically determined by the current partition \mathcal{P} and by the label α .

Finding the optimal α expansion

- Given an input labeling f (partition \mathcal{P}) and a label α we want to find a labeling \hat{f} that minimizes E over all labelings within one α -expansion of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_\alpha = (\mathcal{V}_\alpha, \mathcal{E}_\alpha)$.
- The structure of this graph is dynamically determined by the current partition \mathcal{P} and by the label α .
- Different graph than the $\alpha - \beta$ swap.

Finding the optimal α expansion

- Given an input labeling f (partition \mathcal{P}) and a label α we want to find a labeling \hat{f} that minimizes E over all labelings within one α -expansion of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_\alpha = (\mathcal{V}_\alpha, \mathcal{E}_\alpha)$.
- The structure of this graph is dynamically determined by the current partition \mathcal{P} and by the label α .
- Different graph than the $\alpha - \beta$ swap.

Graph Construction

- The set of vertices includes the two terminals α and $\bar{\alpha}$, as well as all image pixels $p \in \mathcal{P}$.
- Additionally, for each pair of neighboring pixels p, q such that $f_p \neq f_q$ we create an auxiliary node $a_{p,q}$.

Graph Construction

- The set of vertices includes the two terminals α and $\bar{\alpha}$, as well as all image pixels $p \in \mathcal{P}$.
- Additionally, for each pair of neighboring pixels p, q such that $f_p \neq f_q$ we create an auxiliary node $a_{p,q}$.
- Each pixel p is connected to the terminals α and $\bar{\alpha}$, called t -links.

Graph Construction

- The set of vertices includes the two terminals α and $\bar{\alpha}$, as well as all image pixels $p \in \mathcal{P}$.
- Additionally, for each pair of neighboring pixels p, q such that $f_p \neq f_q$ we create an auxiliary node $a_{p,q}$.
- Each pixel p is connected to the terminals α and $\bar{\alpha}$, called t -links.
- Each set of pixels p, q which are neighbors and $f_p = f_q$, we connect with an n -link.

Graph Construction

- The set of vertices includes the two terminals α and $\bar{\alpha}$, as well as all image pixels $p \in \mathcal{P}$.
- Additionally, for each pair of neighboring pixels p, q such that $f_p \neq f_q$ we create an auxiliary node $a_{p,q}$.
- Each pixel p is connected to the terminals α and $\bar{\alpha}$, called t -links.
- Each set of pixels p, q which are neighbors and $f_p = f_q$, we connect with an n -link.
- For each pair of neighboring pixels such that $f_p \neq f_q$, we create a triplet $\{e_{p,a}, e_{a,q}, t_a^{\bar{\alpha}}\}$.

Graph Construction

- The set of vertices includes the two terminals α and $\bar{\alpha}$, as well as all image pixels $p \in \mathcal{P}$.
- Additionally, for each pair of neighboring pixels p, q such that $f_p \neq f_q$ we create an auxiliary node $a_{p,q}$.
- Each pixel p is connected to the terminals α and $\bar{\alpha}$, called t -links.
- Each set of pixels p, q which are neighbors and $f_p = f_q$, we connect with an n -link.
- For each pair of neighboring pixels such that $f_p \neq f_q$, we create a triplet $\{e_{p,a}, e_{a,q}, t_a^{\bar{\alpha}}\}$.
- The set of edges is then

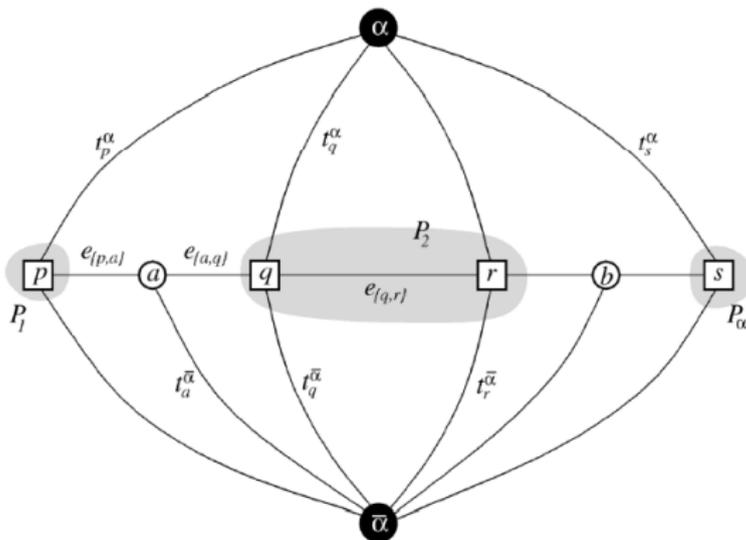
$$\mathcal{E}_\alpha = \left\{ \bigcup_{p \in \mathcal{P}} \{t_p^\alpha, t_p^{\bar{\alpha}}\}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p \neq f_q}} \mathcal{E}_{\{p,q\}}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p = f_q}} e_{\{p,q\}} \right\}$$

Graph Construction

- The set of vertices includes the two terminals α and $\bar{\alpha}$, as well as all image pixels $p \in \mathcal{P}$.
- Additionally, for each pair of neighboring pixels p, q such that $f_p \neq f_q$ we create an auxiliary node $a_{p,q}$.
- Each pixel p is connected to the terminals α and $\bar{\alpha}$, called t -links.
- Each set of pixels p, q which are neighbors and $f_p = f_q$, we connect with an n -link.
- For each pair of neighboring pixels such that $f_p \neq f_q$, we create a triplet $\{e_{p,a}, e_{a,q}, t_a^{\bar{\alpha}}\}$.
- The set of edges is then

$$\mathcal{E}_\alpha = \left\{ \bigcup_{p \in \mathcal{P}} \{t_p^\alpha, t_p^{\bar{\alpha}}\}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p \neq f_q}} \mathcal{E}_{\{p,q\}}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p = f_q}} e_{\{p,q\}} \right\}$$

Graph Construction



edge	weight	for
$t_p^{\bar{\alpha}}$	∞	$p \in \mathcal{P}_\alpha$
$t_p^{\bar{\alpha}}$	$D_p(f_p)$	$p \notin \mathcal{P}_\alpha$
t_p^α	$D_p(\alpha)$	$p \in \mathcal{P}$
$e_{\{p,a\}}$	$V(f_p, \alpha)$	$\{p, q\} \in \mathcal{N}, f_p \neq f_q$
$e_{\{a,q\}}$	$V(\alpha, f_q)$	
$t_a^{\bar{\alpha}}$	$V(f_p, f_q)$	
$e_{\{p,q\}}$	$V(f_p, \alpha)$	$\{p, q\} \in \mathcal{N}, f_p = f_q$

- There is a one-to-one correspondences between a cut and a labeling.

$$f_p^{\mathcal{C}} = \begin{cases} \alpha & \text{if } t_p^\alpha \in \mathcal{C} \\ f_p & \text{if } t_p^{\bar{\alpha}} \in \mathcal{C} \end{cases} \quad \forall p \in \mathcal{P}.$$

- The energy of the cut is the energy of the labeling.
- See Boykov et al, "*fast approximate energy minimization via graph cuts*" PAMI 2001.

Property 5.2. *If $\{p, q\} \in \mathcal{N}$ and $f_p \neq f_q$, then a minimum cut \mathcal{C} on \mathcal{G}_α satisfies:*

- (a) *If $t_p^\alpha, t_q^\alpha \in \mathcal{C}$ then $\mathcal{C} \cap \mathcal{E}_{\{p,q\}} = \emptyset$.*
- (b) *If $t_p^{\bar{\alpha}}, t_q^{\bar{\alpha}} \in \mathcal{C}$ then $\mathcal{C} \cap \mathcal{E}_{\{p,q\}} = t_a^{\bar{\alpha}}$.*
- (c) *If $t_p^{\bar{\alpha}}, t_q^\alpha \in \mathcal{C}$ then $\mathcal{C} \cap \mathcal{E}_{\{p,q\}} = e_{\{p,a\}}$.*
- (d) *If $t_p^\alpha, t_q^{\bar{\alpha}} \in \mathcal{C}$ then $\mathcal{C} \cap \mathcal{E}_{\{p,q\}} = e_{\{a,q\}}$.*

Global Minimization Techniques

Ways to get an approximate solution typically

- Dynamic programming approximations
- Sampling
- Simulated annealing
- Graph-cuts: imposes restrictions on the type of pairwise cost functions
- Message passing: iterative algorithms that pass messages between nodes in the graph. Which graph?

Now we can solve for the MAP (approximately) in general energies. We can solve for other problems than stereo