

Energy Minimization

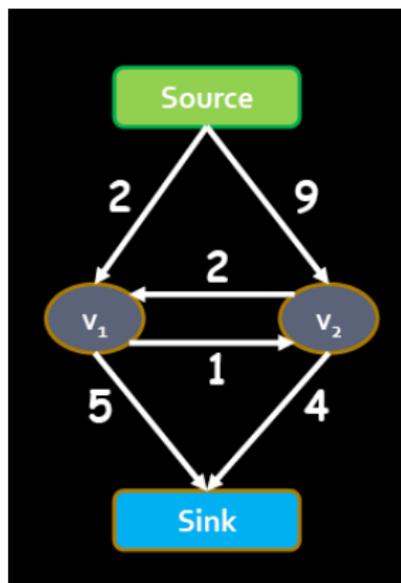
Raquel Urtasun

TTI Chicago

Feb 28, 2013

The ST-mincut problem

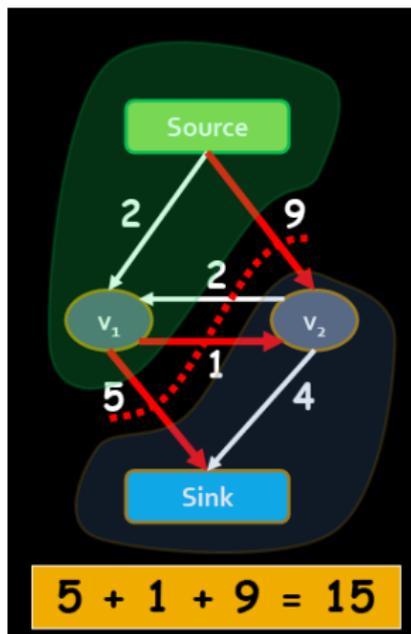
- Suppose we have a graph $G = \{V, E, C\}$, with vertices V , Edges E and costs C .



[Source: P. Kohli]

The ST-mincut problem

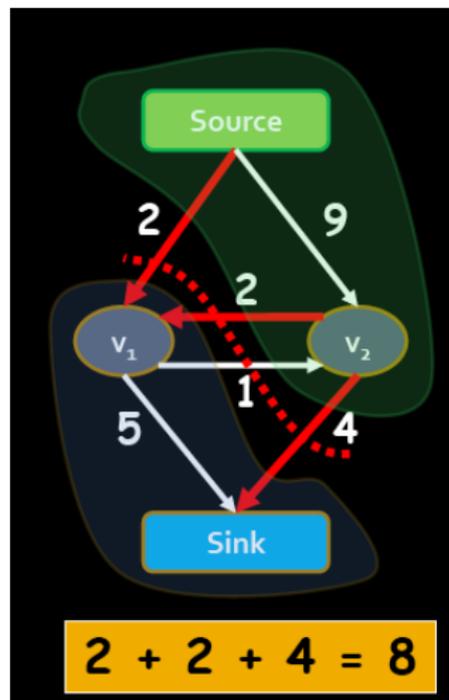
- An st-cut (S,T) divides the nodes between source and sink.
- The cost of a st-cut is the sum of cost of all edges going from S to T



[Source: P. Kohli]

The ST-mincut problem

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

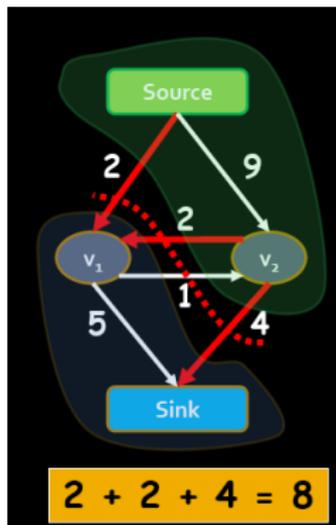


[Source: P. Kohli]

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

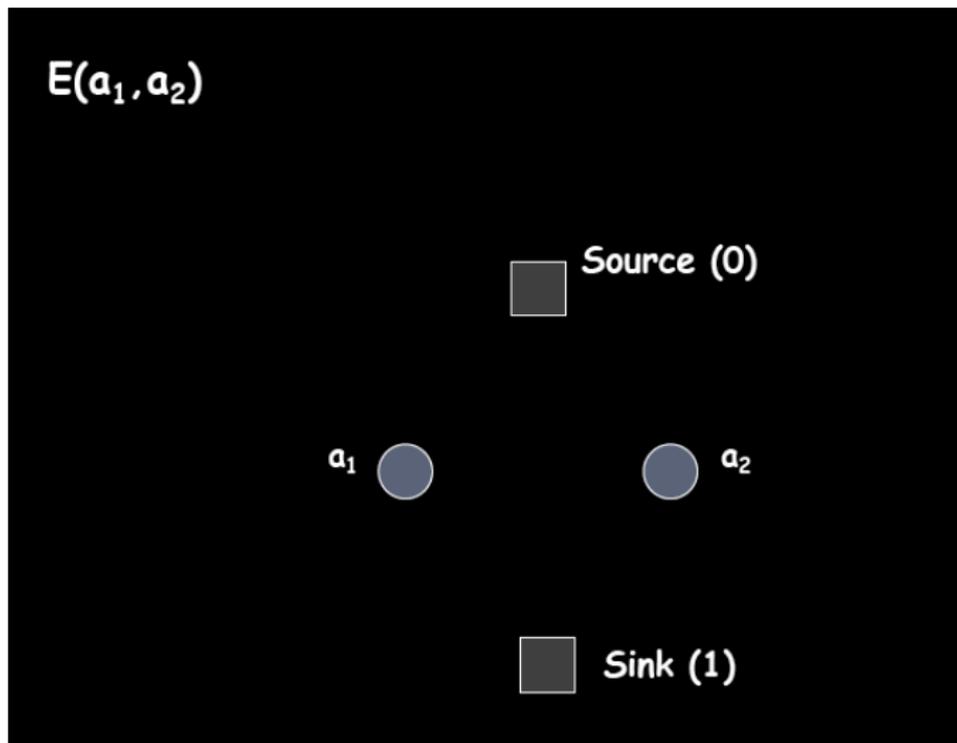
if $x_1=1$ add C-
A if $x_2 = 1$ add
D-C

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

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

[Source: P. Kohli]

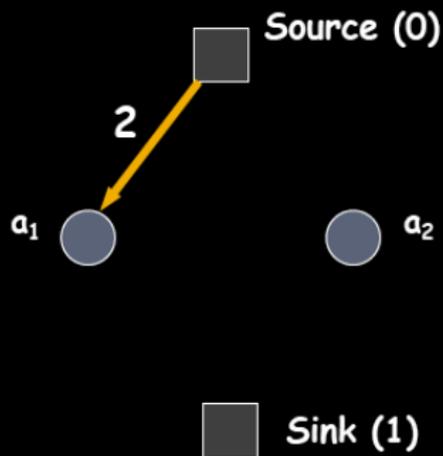
Graph Construction



[Source: P. Kohli]

Graph Construction

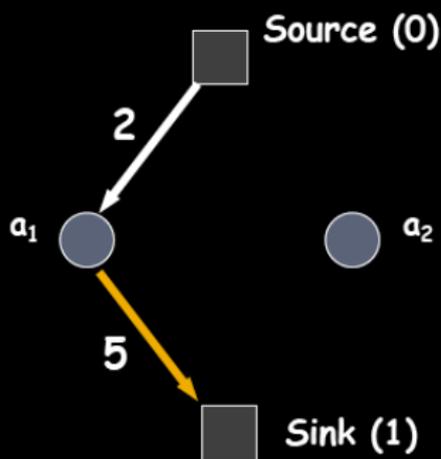
$$E(\mathbf{a}_1, \mathbf{a}_2) = 2\mathbf{a}_1$$



[Source: P. Kohli]

Graph Construction

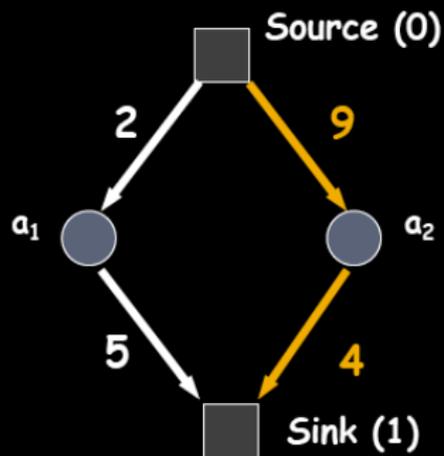
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1$$



[Source: P. Kohli]

Graph Construction

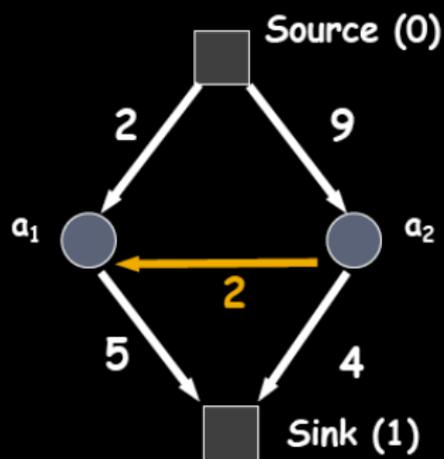
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2$$



[Source: P. Kohli]

Graph Construction

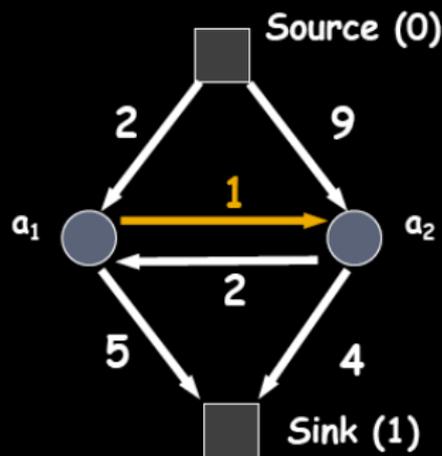
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2 + 2a_1\bar{a}_2$$



[Source: P. Kohli]

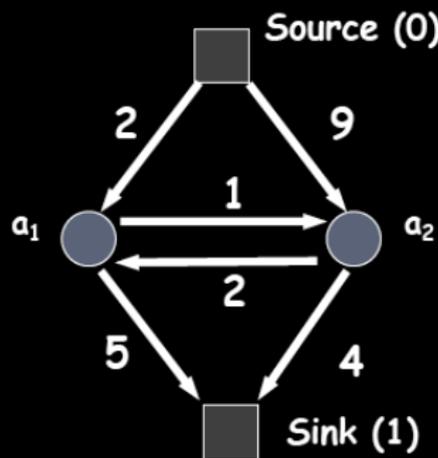
Graph Construction

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



[Source: P. Kohli]

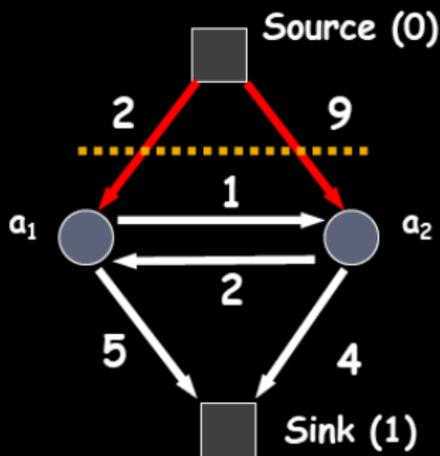
$$E(a_1, a_2) = 2a_1 + 5\bar{a}_1 + 9a_2 + 4\bar{a}_2 + 2a_1\bar{a}_2 + \bar{a}_1a_2$$



[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

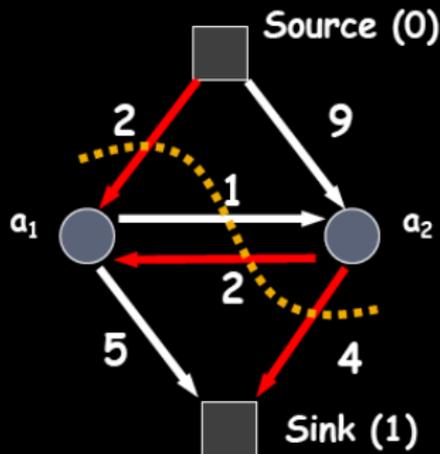
$$a_1 = 1 \quad a_2 = 1$$

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

[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

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

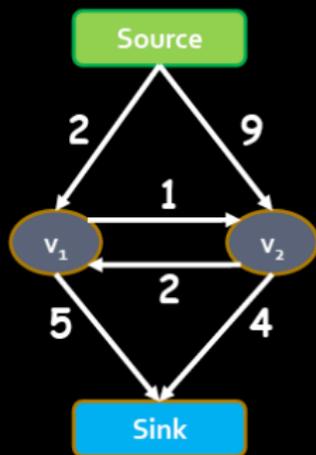
$$E(1, 0) = 8$$

[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
```

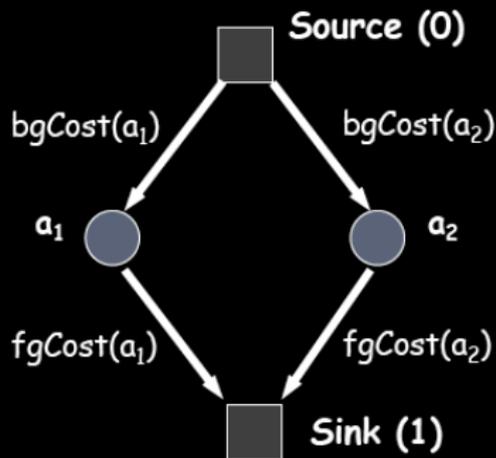
```
    add_weights(nodeID(p), nodeID(q), cost(p,q));
```

```
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
```

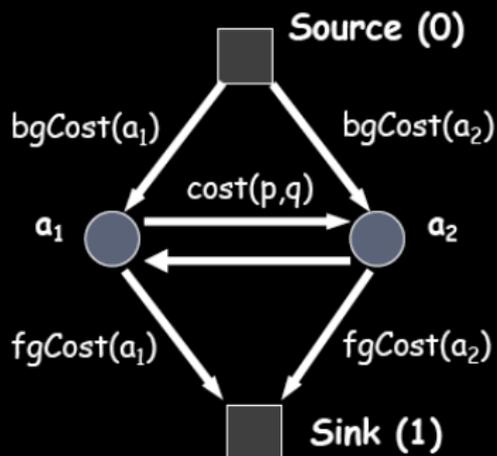
```
add_weights(nodeID(p), nodeID(q), cost(p,q));
```

```
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
```

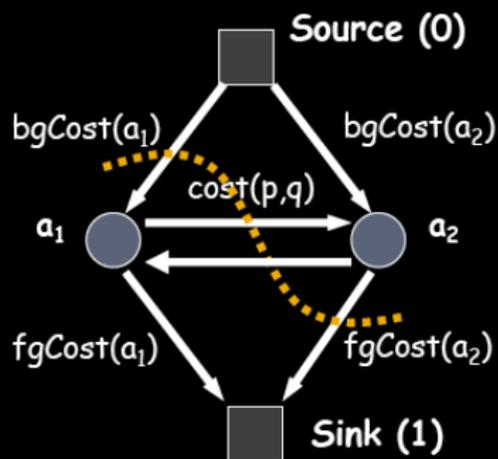
```
add_weights(nodeID(p), nodeID(q), cost(p,q));
```

```
end
```

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

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

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



$$a_1 = bg \quad a_2 = fg$$

[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:

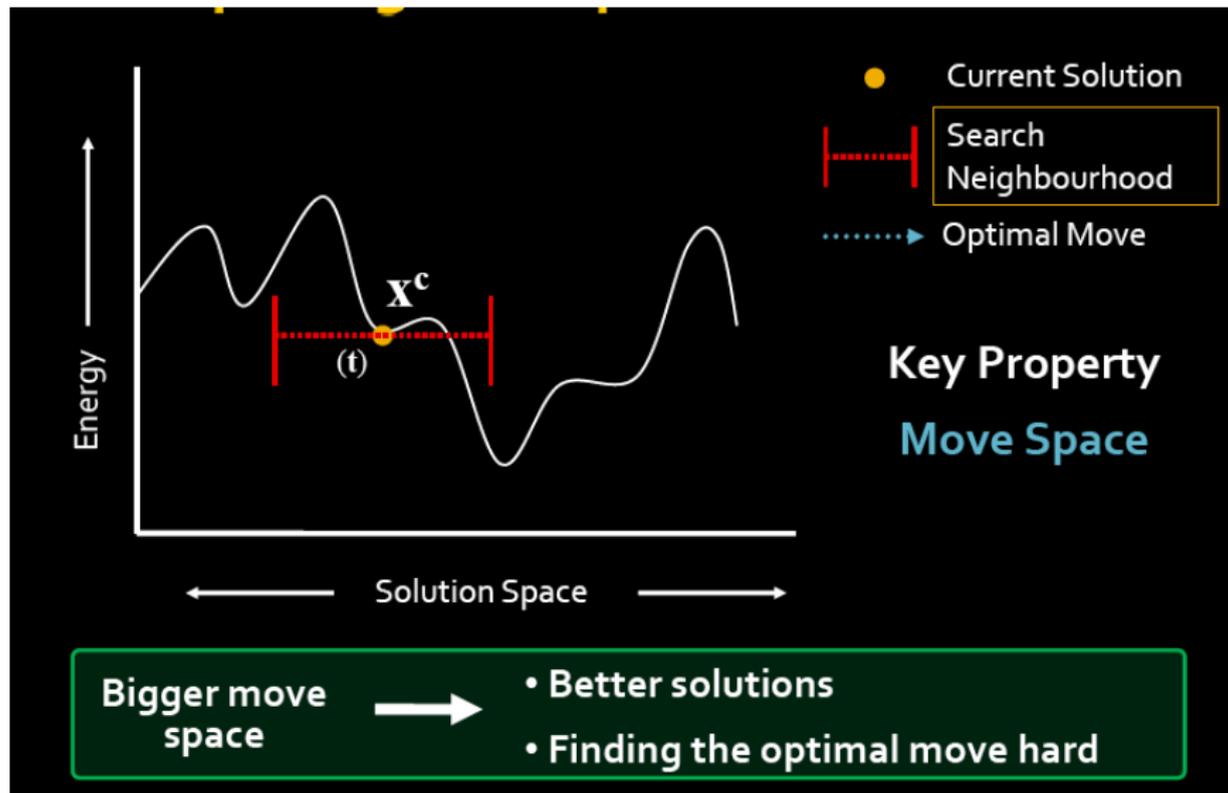
Let Y and X be the set of feasible solutions, then

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

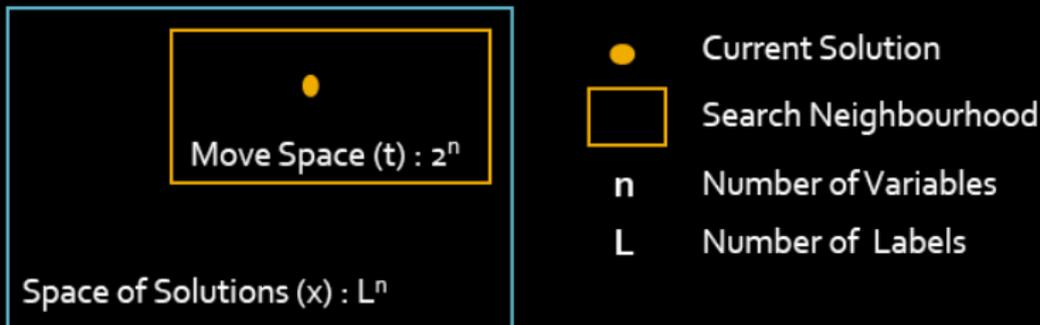


[Source: P. Kohli]

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

$$\begin{aligned}V(\alpha, \beta) = 0 &\leftrightarrow \alpha = \beta \\V(\alpha, \beta) &= V(\beta, \alpha) \geq 0 \\V(\alpha, \beta) &\leq V(\alpha, \gamma) + V(\gamma, \beta)\end{aligned}$$

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

$$\begin{aligned}V(\alpha, \beta) = 0 &\leftrightarrow \alpha = \beta \\V(\alpha, \beta) &= V(\beta, \alpha) \geq 0\end{aligned}$$

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

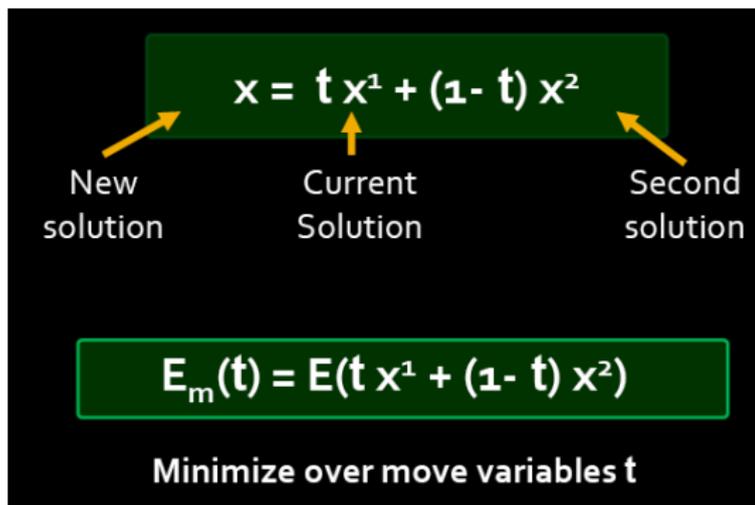
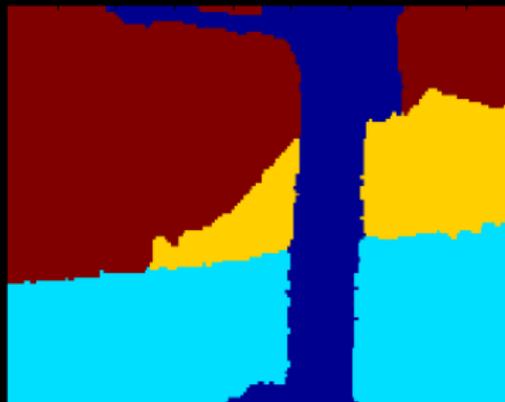


Figure: Figure from P. Kohli tutorial on graph-cuts

- 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

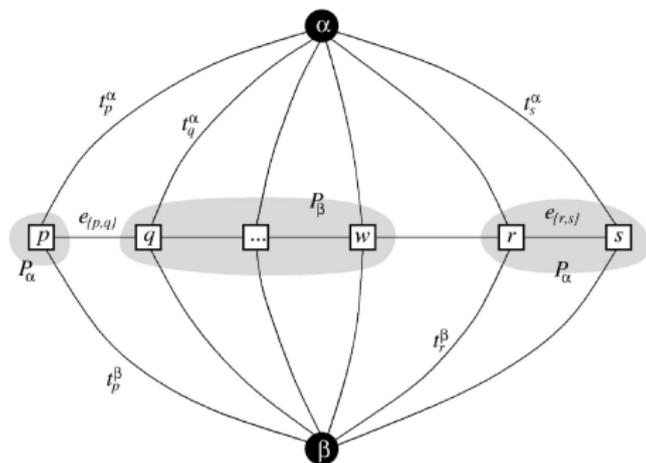
- Given an input labeling f (partition \mathcal{P}) and a pair of labels α, β we want to find a labeling \hat{f} that minimizes E over all labelings within one $\alpha - \beta$ -swap of f .
- This is going to be done by computing a labeling corresponding to a minimum cut on a graph $\mathcal{G}_{\alpha\beta}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

Computing the Cut

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

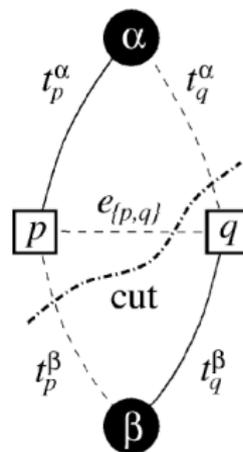
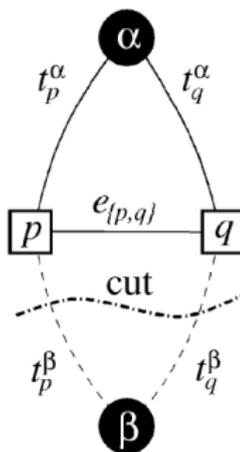
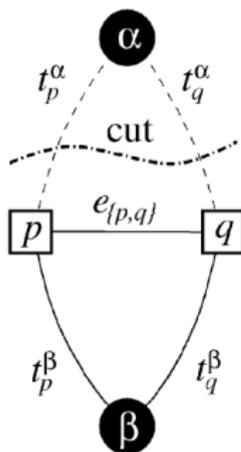
$$f_p^{\mathcal{C}} = \begin{cases} \alpha & \text{if } t_p^\alpha \in \mathcal{C} \text{ for } p \in \mathcal{P}_{\alpha\beta} \\ \beta & \text{if } t_p^\beta \in \mathcal{C} \text{ for } p \in \mathcal{P}_{\alpha\beta} \\ f_p & \text{for } p \in \mathcal{P}, p \notin \mathcal{P}_{\alpha\beta}. \end{cases}$$

- 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

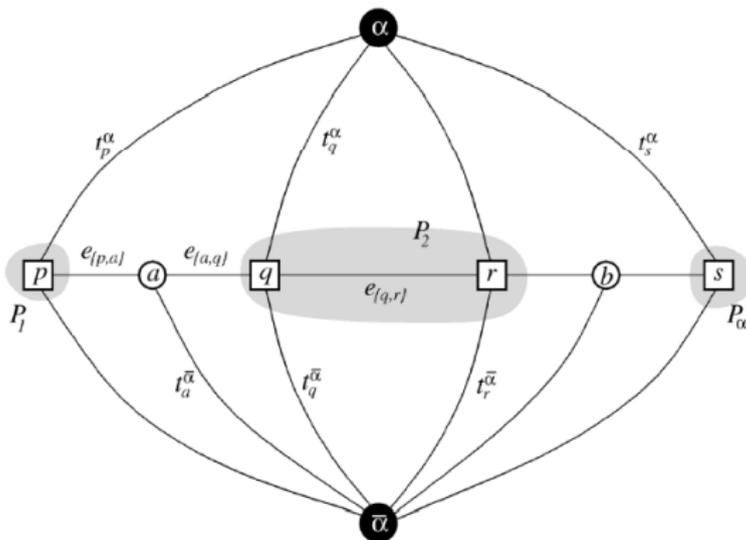
$$\mathcal{E}_\alpha = \left\{ \bigcup_{p \in \mathcal{P}} \{t_p^\alpha, t_p^{\bar{\alpha}}\}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p \neq f_q}} \mathcal{E}_{\{p,q\}}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p = f_q}} e_{\{p,q\}} \right\}$$

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.

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

Let's look at data/benchmarks

Benchmarks

Two benchmarks with very different characteristics



(Middlebury)



(KITTI)

Middlebury Stereo Evaluation – Version 2



- Laboratory
- Lambertian

Middlebury Stereo Evaluation – Version 2



- Laboratory
- Lambertian
- Rich in texture

Middlebury Stereo Evaluation – Version 2



- Laboratory
- Lambertian
- Rich in texture
- Medium-size label set

Middlebury Stereo Evaluation – Version 2



- Laboratory
- Lambertian
- Rich in texture
- Medium-size label set
- Largely fronto-parallel

Middlebury Stereo Evaluation – Version 2



- Laboratory
- Lambertian
- Rich in texture
- Medium-size label set
- Largely fronto-parallel

Middlebury Stereo Evaluation – Version 2



Error Threshold = 1		<u>Tsukuba</u> ground truth			<u>Venus</u> ground truth			<u>Teddy</u> ground truth			<u>Cones</u> ground truth		
Algorithm	Avg.												
CoopRegion [41]	8.8	<u>0.87</u> ₄	1.16 ₁	4.61 ₃	<u>0.11</u> ₄	0.21 ₃	1.54 ₇	<u>5.16</u> ₁₆	8.31 ₁₁	13.0 ₁₃	<u>2.79</u> ₁₇	7.18 ₄	8.01 ₂₃
AdaptingBP [17]	9.0	<u>1.11</u> ₁₉	1.37 ₇	5.79 ₁₉	<u>0.10</u> ₃	0.21 ₄	1.44 ₅	<u>4.22</u> ₈	7.06 ₆	11.8 ₉	<u>2.48</u> ₇	7.92 ₁₁	7.32 ₁₀
ADCensus [94]	7.3	<u>1.07</u> ₁₅	1.48 ₁₃	5.73 ₁₇	<u>0.09</u> ₂	0.25 ₇	1.15 ₃	<u>4.10</u> ₆	6.22 ₃	10.9 ₆	<u>2.42</u> ₅	7.25 ₅	6.95 ₆
SurfaceStereo [79]	18.2	<u>1.28</u> ₃₂	1.65 ₂₁	6.78 ₃₇	<u>0.19</u> ₁₈	0.28 ₁₀	2.61 ₃₂	<u>3.12</u> ₂	5.10 ₁	8.65 ₁	<u>2.89</u> ₂₁	7.95 ₁₃	8.26 ₃₀
GC+SegmBorder [57]	27.1	<u>1.47</u> ₄₅	1.82 ₃₂	7.86 ₅₈	<u>0.19</u> ₁₉	0.31 ₁₂	2.44 ₂₆	<u>4.25</u> ₉	5.55 ₂	10.9 ₇	<u>4.99</u> ₇₇	5.78 ₁	8.66 ₃₇
WarpMat [55]	20.8	<u>1.16</u> ₂₀	1.35 ₆	6.04 ₂₄	<u>0.18</u> ₁₇	0.24 ₆	2.44 ₂₆	<u>5.02</u> ₁₃	9.30 ₁₇	13.0 ₁₅	<u>3.49</u> ₃₉	8.47 ₂₂	9.01 ₄₄
RDP [102]	12.5	<u>0.97</u> ₁₀	1.39 ₉	5.00 ₉	<u>0.21</u> ₂₃	0.38 ₁₉	1.89 ₁₃	<u>4.84</u> ₁₀	9.94 ₁₉	12.6 ₁₁	<u>2.53</u> ₈	7.69 ₈	7.38 ₁₁
RVbased [116]	11.6	<u>0.95</u> ₉	1.42 ₁₁	4.98 ₈	<u>0.11</u> ₆	0.29 ₁₁	1.07 ₁	<u>5.98</u> ₂₁	11.6 ₃₁	15.4 ₂₇	<u>2.35</u> ₃	7.61 ₆	6.81 ₅
OutlierConf [42]	12.9	<u>0.88</u> ₆	1.43 ₁₂	4.74 ₅	<u>0.18</u> ₁₆	0.26 ₉	2.40 ₂₂	<u>5.01</u> ₁₂	9.12 ₁₆	12.8 ₁₂	<u>2.78</u> ₁₆	8.57 ₂₃	6.99 ₇

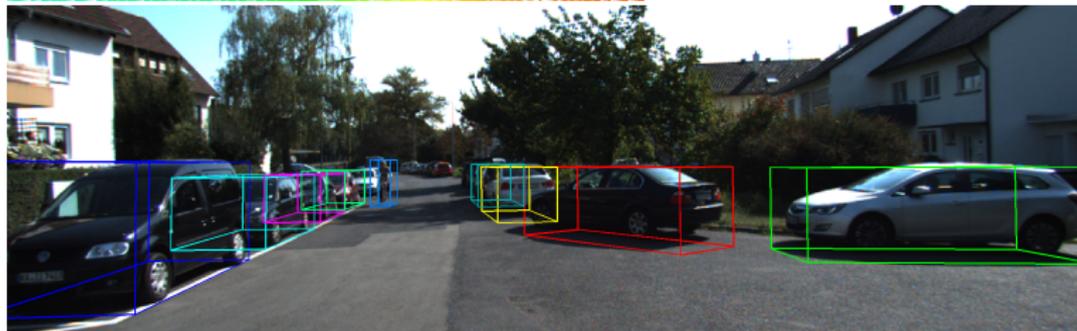
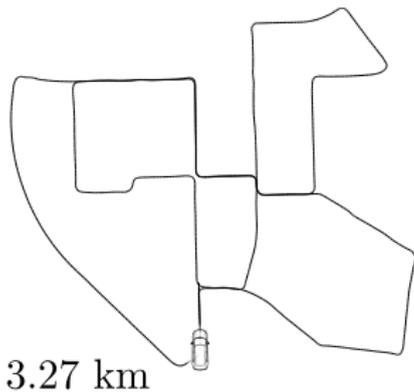
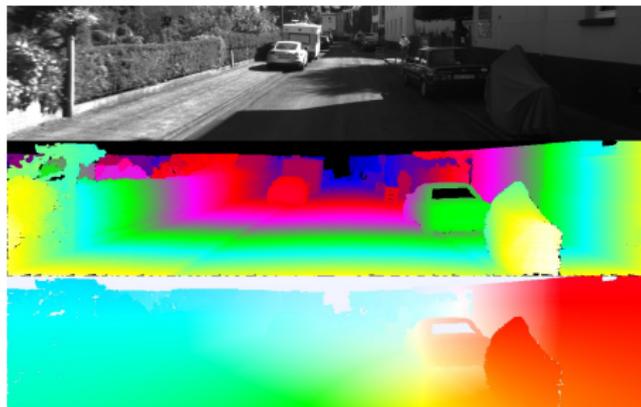
- Best methods < 3% errors (for all non-occluded regions)
- <http://vision.middlebury.edu/stereo/data/>

Benchmarks: KITTI Data Collection

- **Two stereo rigs** (1392×512 px, 54 cm base, 90° opening)
- **Velodyne** laser scanner, **GPS+IMU** localization
- **6 hours** at 10 frames per second!



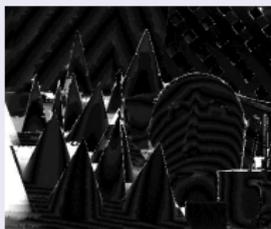
The KITTI Vision Benchmark Suite



Novel Challenges

Fast guided cost-volume filtering (Rhemann et al., CVPR 2011)

Middlebury, Errors: **2.7%**

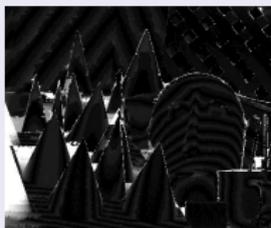


- Error threshold: 1 px (Middlebury) / 3 px (KITTI)

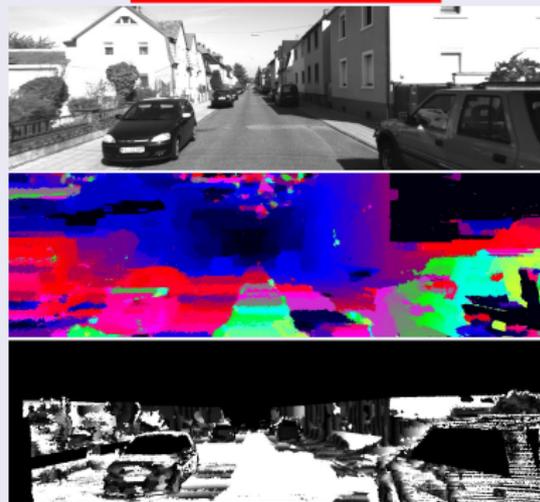
Novel Challenges

Fast guided cost-volume filtering (Rhemann et al., CVPR 2011)

Middlebury, Errors: **2.7%**



KITTI, Errors: **46.3%**



- Error threshold: 1 px (Middlebury) / 3 px (KITTI)

Novel Challenges

So what is the difference?

Middlebury



- Laboratory
- Lambertian

KITTI

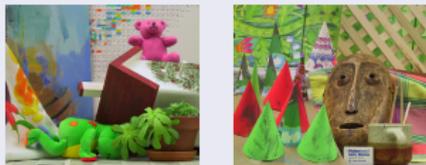


- Moving vehicle
- Specularities

Novel Challenges

So what is the difference?

Middlebury



- Laboratory
- Lambertian
- Rich in texture

KITTI

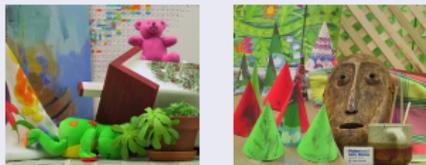


- Moving vehicle
- Specularities
- Sensor saturation

Novel Challenges

So what is the difference?

Middlebury



- Laboratory
- Lambertian
- Rich in texture
- Medium-size label set

KITTI



- Moving vehicle
- Specularities
- Sensor saturation
- Large label set

Novel Challenges

So what is the difference?

Middlebury



- Laboratory
- Lambertian
- Rich in texture
- Medium-size label set
- Largely fronto-parallel

KITTI



- Moving vehicle
- Specularities
- Sensor saturation
- Large label set
- Strong slants

Novel Challenges

So what is the difference?

Middlebury



- Laboratory
- Lambertian
- Rich in texture
- Medium-size label set
- Largely fronto-parallel

KITTI



- Moving vehicle
- Specularities
- Sensor saturation
- Large label set
- Strong slants

Stereo Evaluation

Rank	Method	Setting	Out-Noc	Out-All	Avg-Noc	Avg-All	Density	Runtime	Environment	Compare
1	PCBP		4.13 %	5.45 %	0.9 px	1.2 px	100.00 %	5 min	4 cores @ 2.5 Ghz (Matlab + C/C++)	<input type="checkbox"/>
Koichiro Yamaguchi, Tamir Hazar, David McAllester and Raquel Urtasun. Continuous Markov Random Fields for Robust Stereo Estimation , ECCV 2012.										
2	ISGM		5.16 %	7.19 %	1.2 px	2.1 px	94.70 %	8 s	2 cores @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Simon Herrmann and Reinhard Klette. Iterative Semi-Global Matching for Robust Driver Assistance Systems , ACCV 2012.										
3	SGM		5.83 %	7.08 %	1.2 px	1.3 px	85.80 %	3.7 s	1 core @ 3.0 Ghz (C/C++)	<input type="checkbox"/>
Heiko Hirschmüller. Stereo Processing by Semi-Global Matching and Mutual Information , IEEE Transactions on Pattern Analysis and Machine Intelligence 2008.										
4	SNCC		6.27 %	7.33 %	1.4 px	1.5 px	100.00 %	0.27 s	1 core @ 3.0 Ghz (C/C++)	<input type="checkbox"/>
N. Enecke and J. Eggert. A Two-Stage Correlation Method for Stereoscopic Depth Estimation , DICTA 2010.										
5	ITGV		6.31 %	7.40 %	1.3 px	1.5 px	100.00 %	7 s	1 core @ 3.0 Ghz (Matlab + C/C++)	<input type="checkbox"/>
Rene Ranftl, Stefan Gehrig, Thomas Pock and Horst Bischof. Pushing the Limits of Stereo Using Variational Stereo Estimation , IEEE Intelligent Vehicles Symposium 2012.										
6	BSSM		7.50 %	8.89 %	1.4 px	1.6 px	94.87 %	20.7 s	1 core @ 3.5 Ghz (C/C++)	<input type="checkbox"/>
Anonymous submission										
7	OCV-SGBM		7.64 %	9.13 %	1.8 px	2.0 px	86.50 %	1.1 s	1 core @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Heiko Hirschmüller. Stereo processing by semiglobal matching and mutual information , PAMI 2008.										
8	ELAS		8.24 %	9.95 %	1.4 px	1.6 px	94.55 %	0.3 s	1 core @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Andreas Geiger, Martin Roser and Raquel Urtasun. Efficient Large-Scale Stereo Matching , ACCV 2010.										
9	MS-DSI		10.68 %	12.11 %	1.9 px	2.2 px	100.00 %	10.3 s	>8 cores @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Anonymous submission										
10	SDM		10.98 %	12.19 %	2.0 px	2.3 px	63.58 %	1 min	1 core @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Jana Kostkova. Stratified dense matching for stereopsis in complex scenes , BMVC 2003.										
11	GCSF		12.06 %	13.26 %	1.9 px	2.1 px	60.77 %	2.4 s	1 core @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Jan Cech, Jordi Sanchez-Riera and Radu P. Horaud. Scene Flow Estimation by Growing Correspondence Seeds , CVPR 2011.										
12	GCS		13.37 %	14.54 %	2.1 px	2.3 px	51.06 %	2.2 s	1 core @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Jan Cech and Radim Sara. Efficient Sampling of Disparity Space for Fast And Accurate Matching , BenCOS 2007.										
13	CostFilter		19.96 %	21.05 %	5.0 px	5.4 px	100.00 %	4 min	1 core @ 2.5 Ghz (Matlab)	<input type="checkbox"/>
Christoph Rhemann, Asmaa Hosni, Michael Bleyer, Carsten Rother and Margrit Gelautz. Fast Cost-Volume Filtering for Visual Correspondence and Beyond , CVPR 2011.										
14	OCV-BM		25.39 %	26.72 %	7.6 px	7.9 px	55.84 %	0.1 s	1 core @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
G. Bradski. The OpenCV Library . Dr. Dobb's Journal of Software Tools 2000.										
15	GC+occ		33.50 %	34.74 %	8.6 px	9.2 px	87.57 %	6 min	1 core @ 2.5 Ghz (C/C++)	<input type="checkbox"/>
Vladimir Kolmogorov and Ramin Zabih. Computing Visual Correspondence with Occlusions using Graph Cuts , ICCV 2001.										

Global methods: define a Markov random field over

- Pixel-level
- Fronto-parallel planes
- Slanted planes

Plane MRFs

- First segment an image into small regions, i.e., superpixels
- Assume that the 3D world is composed of small frontal/slanted planes

- First segment an image into small regions, i.e., superpixels
- Assume that the 3D world is composed of small frontal/slanted planes
- Good representation if the superpixels are small and respect boundaries

$$E(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_i C(\mathbf{x}_i) + \sum_i \sum_{j \in \mathcal{N}_i} C(\mathbf{x}_i, \mathbf{x}_j)$$

with $\mathbf{x}_i \in \mathfrak{R}$ for the fronto-parallel planes, and $\mathbf{x}_i \in \mathfrak{R}^3$ for the slanted planes

- First segment an image into small regions, i.e., superpixels
- Assume that the 3D world is composed of small frontal/slanted planes
- Good representation if the superpixels are small and respect boundaries

$$E(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_i C(\mathbf{x}_i) + \sum_i \sum_{j \in \mathcal{N}_i} C(\mathbf{x}_i, \mathbf{x}_j)$$

with $\mathbf{x}_i \in \mathfrak{R}$ for the fronto-parallel planes, and $\mathbf{x}_i \in \mathfrak{R}^3$ for the slanted planes

- This are continuous variables. Is this a problem?

- First segment an image into small regions, i.e., superpixels
- Assume that the 3D world is composed of small frontal/slanted planes
- Good representation if the superpixels are small and respect boundaries

$$E(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_i C(\mathbf{x}_i) + \sum_i \sum_{j \in \mathcal{N}_i} C(\mathbf{x}_i, \mathbf{x}_j)$$

with $\mathbf{x}_i \in \mathfrak{R}$ for the fronto-parallel planes, and $\mathbf{x}_i \in \mathfrak{R}^3$ for the slanted planes

- This are continuous variables. Is this a problem?
- What can I do to solve this? Discretize the problem

- First segment an image into small regions, i.e., superpixels
- Assume that the 3D world is composed of small frontal/slanted planes
- Good representation if the superpixels are small and respect boundaries

$$E(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_i C(\mathbf{x}_i) + \sum_i \sum_{j \in \mathcal{N}_i} C(\mathbf{x}_i, \mathbf{x}_j)$$

with $\mathbf{x}_i \in \mathfrak{R}$ for the fronto-parallel planes, and $\mathbf{x}_i \in \mathfrak{R}^3$ for the slanted planes

- This are continuous variables. Is this a problem?
- What can I do to solve this? Discretize the problem
- The unitary are usually aggregation of cost over the local matching on the pixels in that superpixel

- First segment an image into small regions, i.e., superpixels
- Assume that the 3D world is composed of small frontal/slanted planes
- Good representation if the superpixels are small and respect boundaries

$$E(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_i C(\mathbf{x}_i) + \sum_i \sum_{j \in \mathcal{N}_i} C(\mathbf{x}_i, \mathbf{x}_j)$$

with $\mathbf{x}_i \in \mathfrak{R}$ for the fronto-parallel planes, and $\mathbf{x}_i \in \mathfrak{R}^3$ for the slanted planes

- This are continuous variables. Is this a problem?
- What can I do to solve this? Discretize the problem
- The unitary are usually aggregation of cost over the local matching on the pixels in that superpixel
- Pairwise is typically smoothness

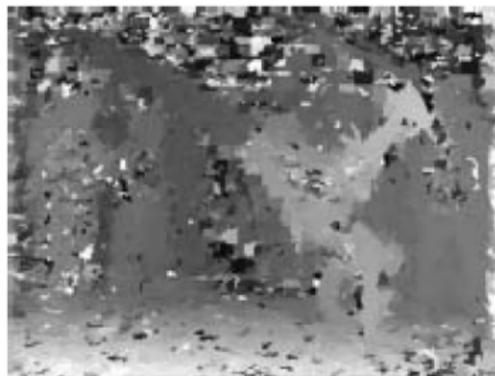
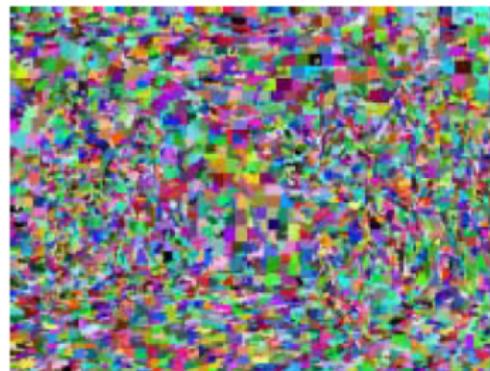
- First segment an image into small regions, i.e., superpixels
- Assume that the 3D world is composed of small frontal/slanted planes
- Good representation if the superpixels are small and respect boundaries

$$E(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_i C(\mathbf{x}_i) + \sum_i \sum_{j \in \mathcal{N}_i} C(\mathbf{x}_i, \mathbf{x}_j)$$

with $\mathbf{x}_i \in \mathfrak{R}$ for the fronto-parallel planes, and $\mathbf{x}_i \in \mathfrak{R}^3$ for the slanted planes

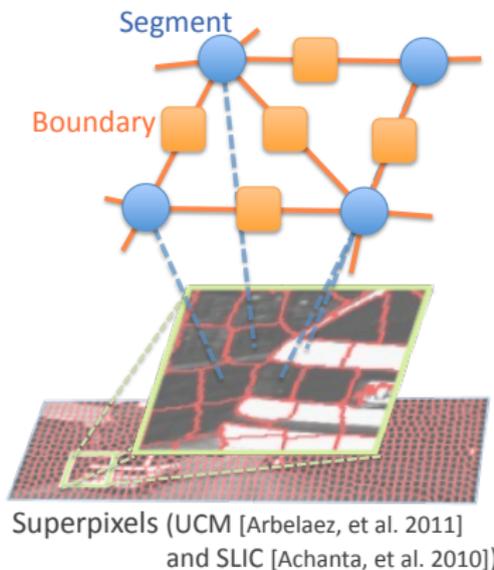
- These are continuous variables. Is this a problem?
- What can I do to solve this? Discretize the problem
- The unary are usually aggregation of cost over the local matching on the pixels in that superpixel
- Pairwise is typically smoothness

Slanted-plane MRFs



A more sophisticated occlusion model

- MRF on continuous variables (slanted planes) and discrete var. (boundary)
- Combines depth ordering (segmentation) and stereo



Segment variable $y_i = (\alpha_i, \beta_i, \gamma_i)$

Slanted 3D plane of segment

Continuous variable

Boundary variable o_{ij}

Relationship between segments

4 states



Occlusion



Hinge



Coplanar

Discrete variable

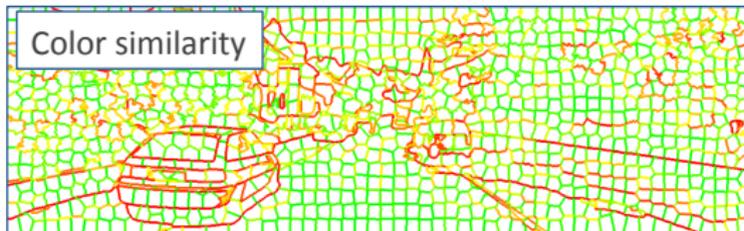
- Takes as input disparities computed by any local algorithm

Energy of PCBP-Stereo

- \mathbf{y} the set of slanted 3D planes, \mathbf{o} the set of discrete boundary variables

$$E(\mathbf{y}, \mathbf{o}) = E_{color}(\mathbf{o}) + E_{match}(\mathbf{y}, \mathbf{o}) + E_{compatibility}(\mathbf{y}, \mathbf{o}) + E_{junction}(\mathbf{o})$$

Similar color  Likely to be coplanar

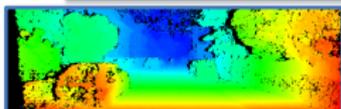


Energy of PCBP-Stereo

- \mathbf{y} the set of slanted 3D planes, \mathbf{o} the set of discrete boundary variables

$$E(\mathbf{y}, \mathbf{o}) = E_{color}(\mathbf{o}) + E_{match}(\mathbf{y}, \mathbf{o}) + E_{compatibility}(\mathbf{y}, \mathbf{o}) + E_{junction}(\mathbf{o})$$

Agreement with result of input disparity map



Computed by any matching method
(Modified semi-global matching)

$$\text{Truncated quadratic function } \phi_i^{\text{TP}}(\mathbf{p}, \mathbf{y}_i, K) = \min \left(|\mathcal{D}(\mathbf{p}) - \hat{d}_i(\mathbf{p}, \mathbf{y}_i)|, K \right)^2$$

Disparity map Slanted plane

On boundary

“Occlusion” – Foreground segment owns boundary



Energy of PCBP-Stereo

- \mathbf{y} the set of slanted 3D planes, \mathbf{o} the set of discrete boundary variables

$$E(\mathbf{y}, \mathbf{o}) = E_{color}(\mathbf{o}) + E_{match}(\mathbf{y}, \mathbf{o}) + E_{compatibility}(\mathbf{y}, \mathbf{o}) + E_{junction}(\mathbf{o})$$

- (1) Preference of boundary label (Coplanar > Hinge > Occlusion)

Impose penalty $\lambda_{occ} > \lambda_{hinge} > 0$

- (2) Boundary labels \longleftrightarrow match \longleftrightarrow Slanted planes

“Occlusion”	\longleftrightarrow	$\hat{d}_{front}(\mathbf{p}) > \hat{d}_{back}(\mathbf{p})$
“Hinge”	\longleftrightarrow	$\hat{d}_i(\mathbf{p}) = \hat{d}_j(\mathbf{p})$ on boundary
“Coplanar”	\longleftrightarrow	$\hat{d}_i(\mathbf{p}) = \hat{d}_j(\mathbf{p})$ in both segments



Energy of PCBP-Stereo

- \mathbf{y} the set of slanted 3D planes, \mathbf{o} the set of discrete boundary variables

$$E(\mathbf{y}, \mathbf{o}) = E_{color}(\mathbf{o}) + E_{match}(\mathbf{y}, \mathbf{o}) + E_{compatibility}(\mathbf{y}, \mathbf{o}) + E_{junction}(\mathbf{o})$$

Occlusion boundary reasoning [Malik 1987]

Penalize impossible junctions

Impossible cases

