Multi-label energy minimization via graph cuts

Stereo matching

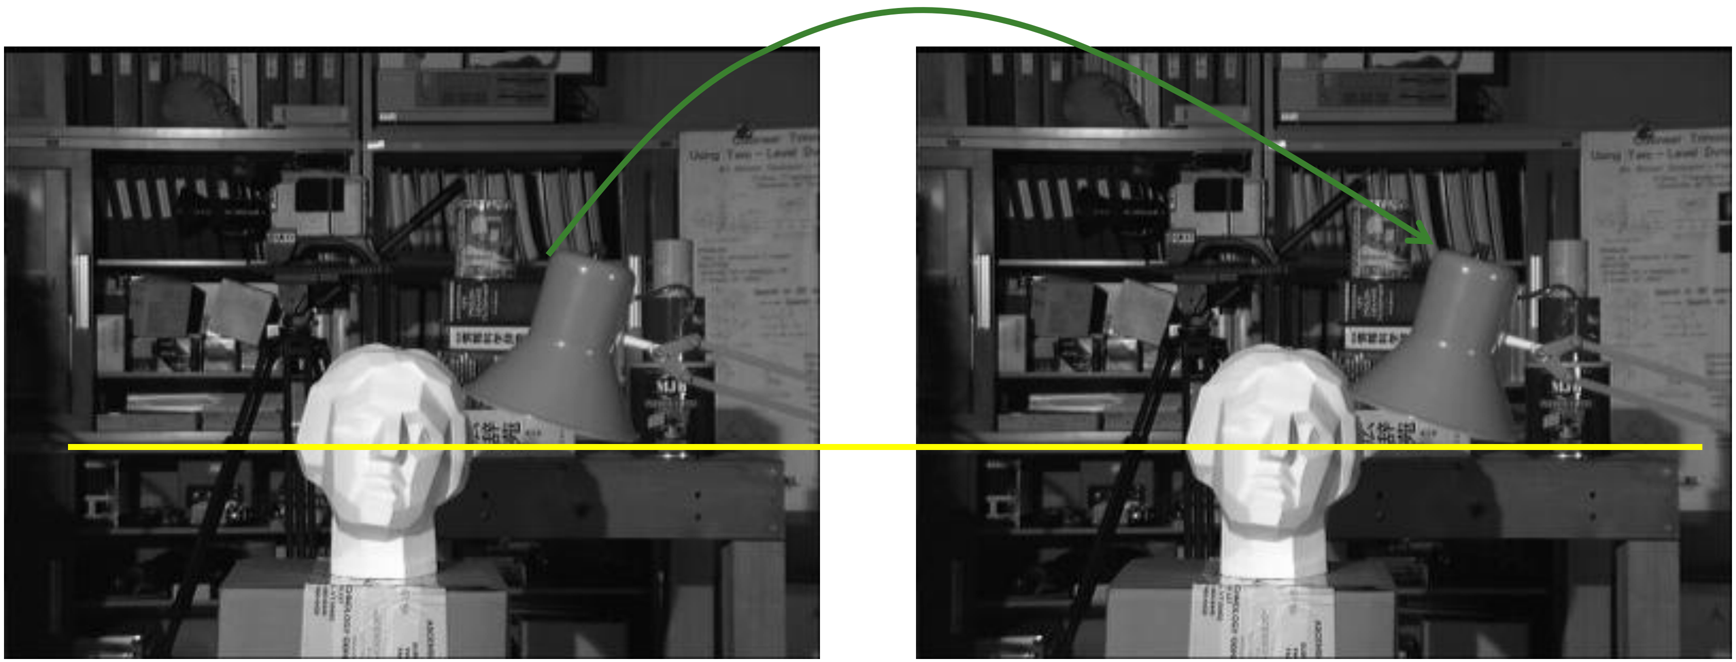

- Extract correspondences between similar images
- Images are typically assumed to be horizontally aligned (rectification)

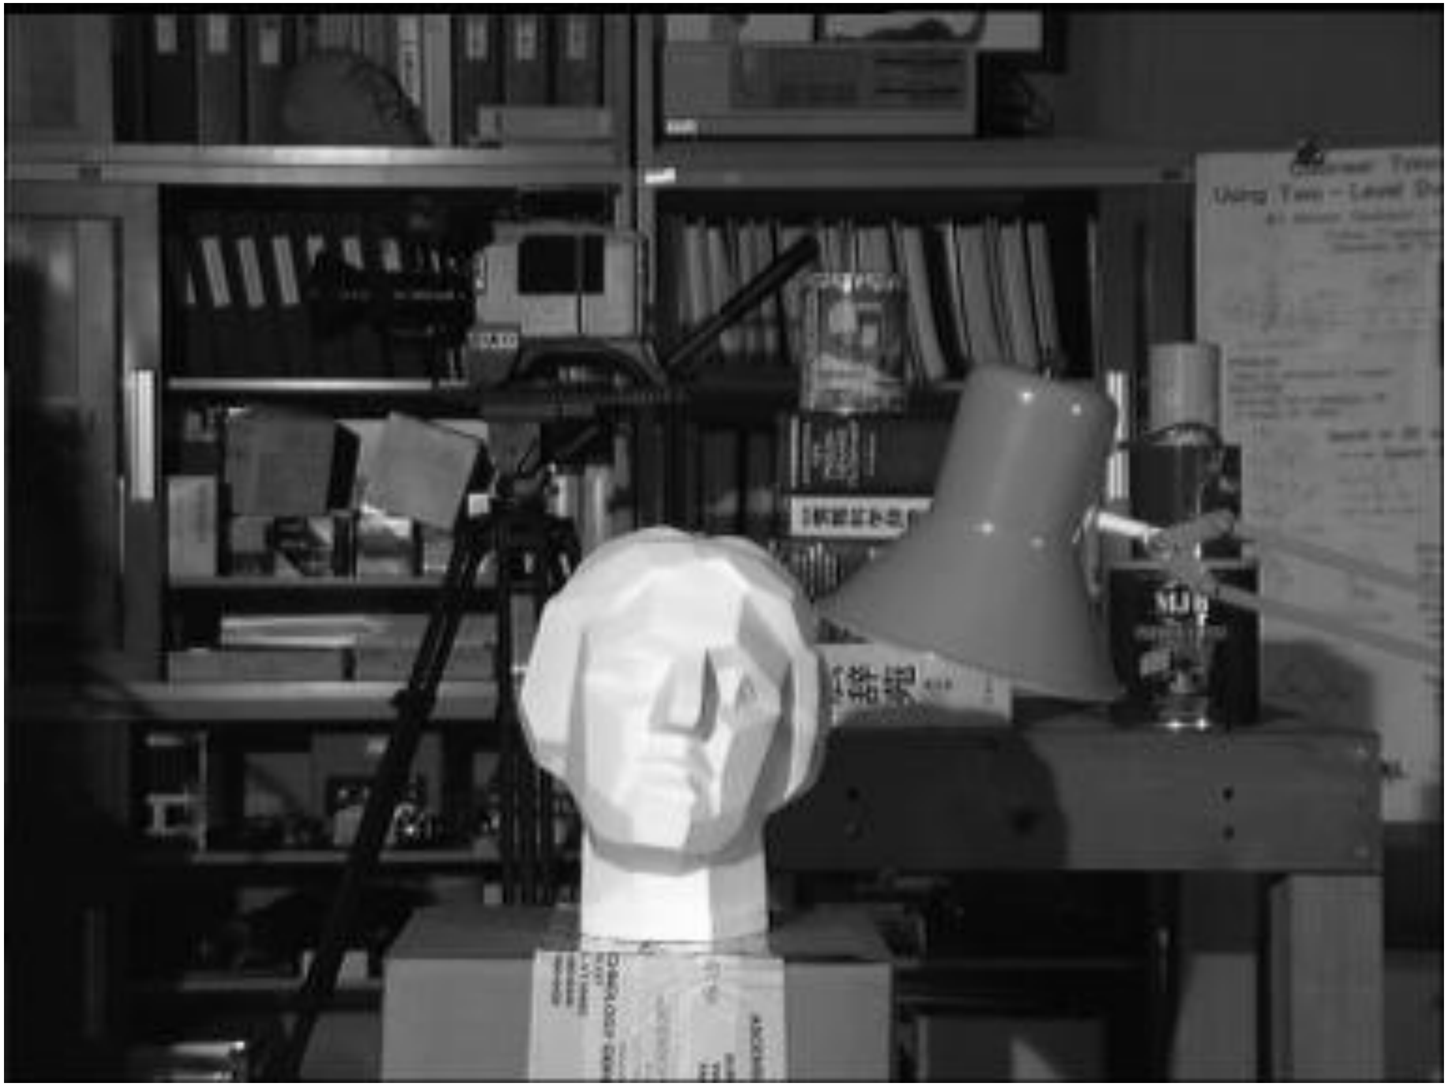

- Correspondences via horizontal shifts (called "disparities")

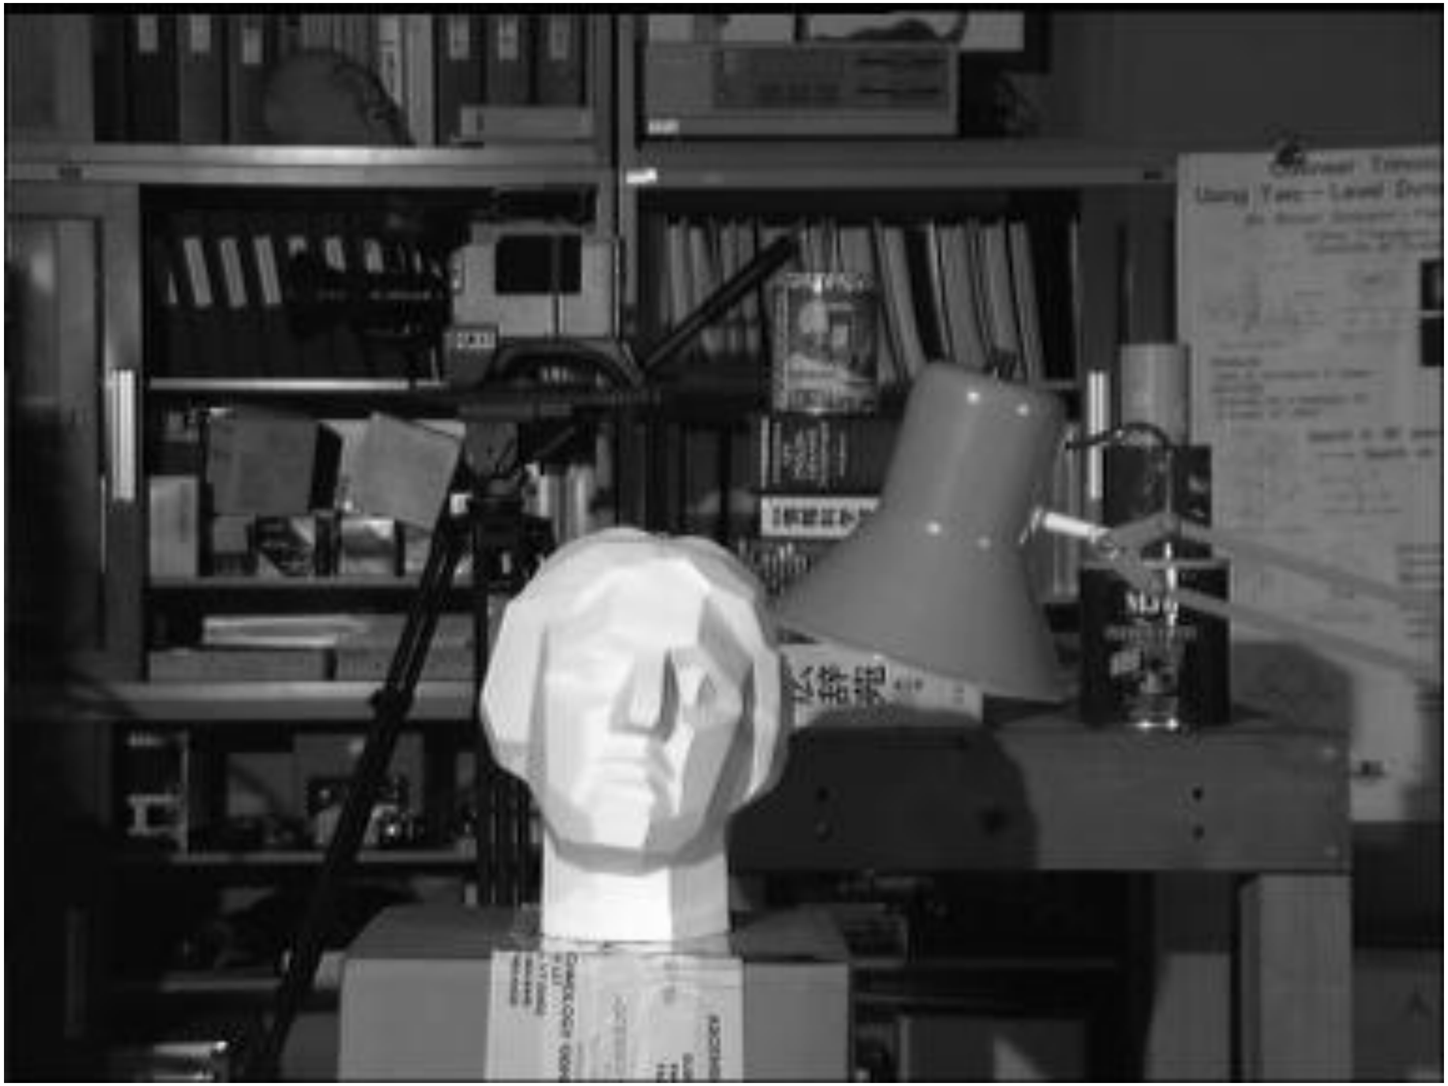

- Correspondences via horizontal shifts (called “disparities”)

Stereo matching

- Visualization of disparities (disparity map)
- Disparity inversely proportional to depth

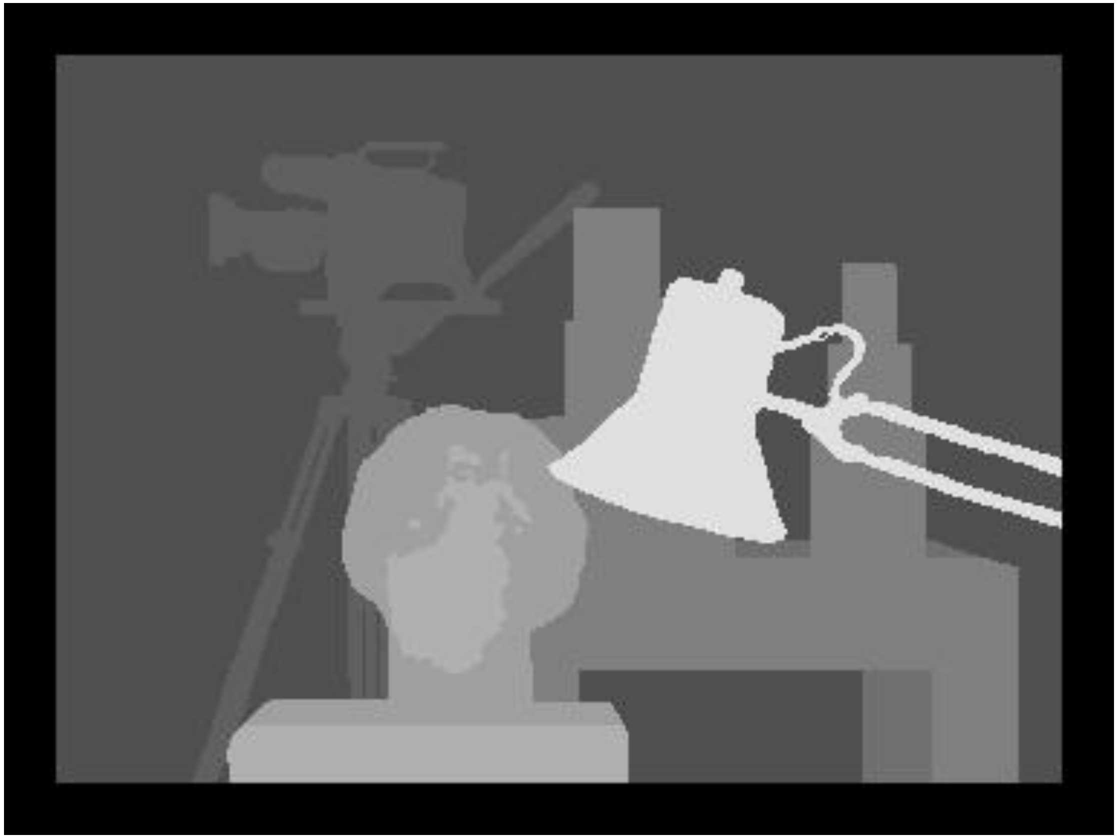

Window based approach

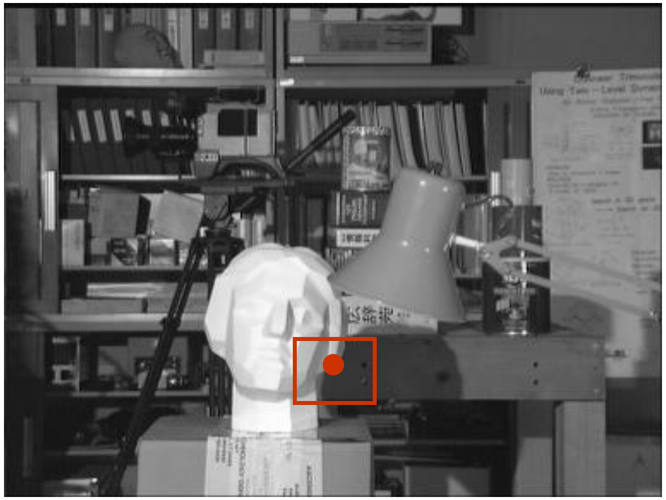

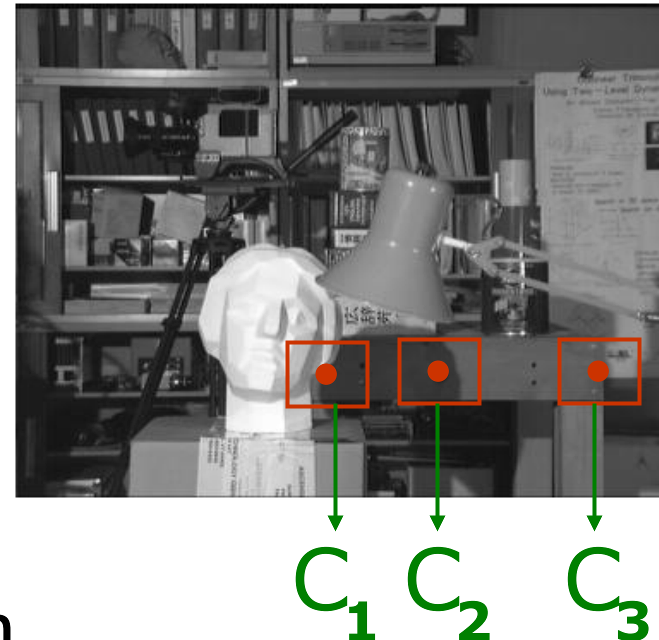

- Winner-takes-all approach
- Windows matched independently
- Small or large windows can be used
- With a simple trick, running time can be made independent of window size

Small vs large windows

small window

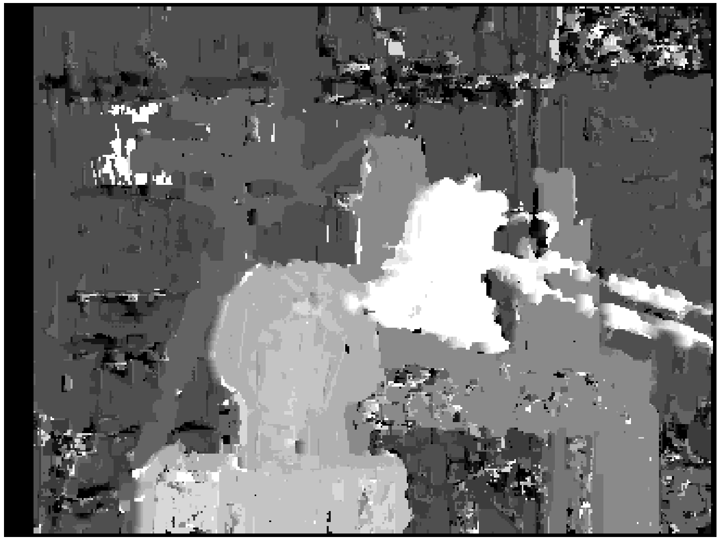

- better at boundaries
- noisy in low texture areas

large window

- better in low texture areas
- blurred boundaries

Scan-line approaches

- Match scan lines independently, i.e., introduce coherence only along scanlines
(what is the resulting MRF?)

- Better than window-based approach

- But still not good enough

- Problem: streaking artifacts

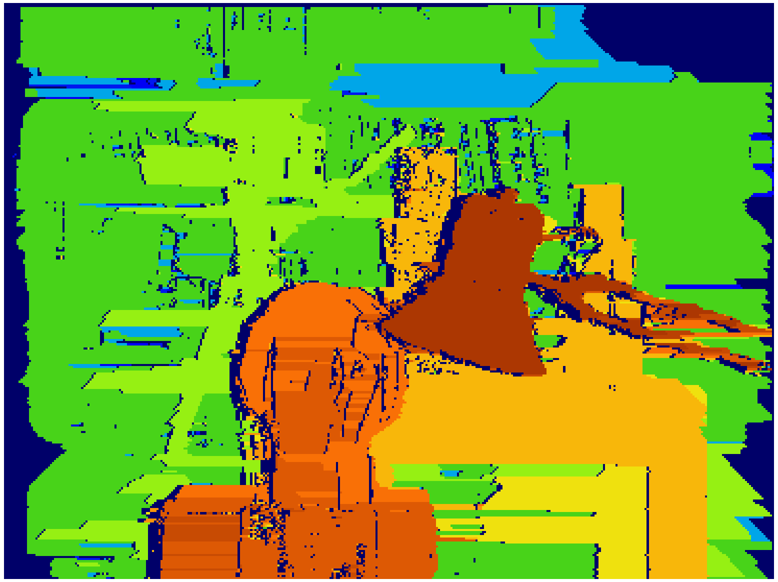

Graph-cut approach

- We will use a 2D grid for our MRF
- We will penalize disparity discontinuities both in horizontal or vertical direction
- Much better modeling
(spatial coherence along AND across scanlines)

Graph-cut approach

- Resulting MRF energy:

$$\underbrace{\sum_p (I_p - I_{p+d_p})^2}_{\text{Photo-consistency}} + \underbrace{\sum_{pq} w_{pq} |d_p - d_q|}_{\text{Spatial coherence}}$$

- How can we select the weights w_{pq}?

- Why not just apply loopy-BP in this case?

MRF optimization via graph-cuts

- Optimizing MRF energies of the following form:

$$\sum_p V_p(d_p) + \sum_{pq} w_{pq} \left| d_p - d_q \right|$$

- Belief propagation can not guarantee an optimal solution (loopy graph)
- We will use graph-cut based methods (exact global optimum in polynomial time)
- But how can this be reduced to a graph-cut problem?

MRF optimization via graph-cuts

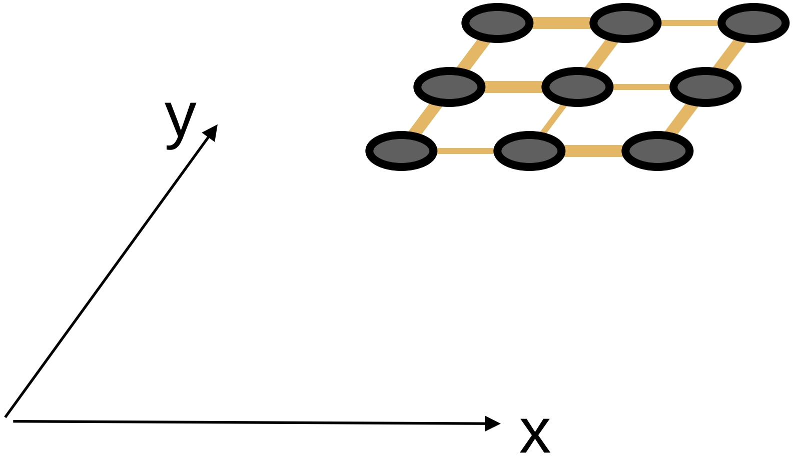

MRF optimization via graph-cuts

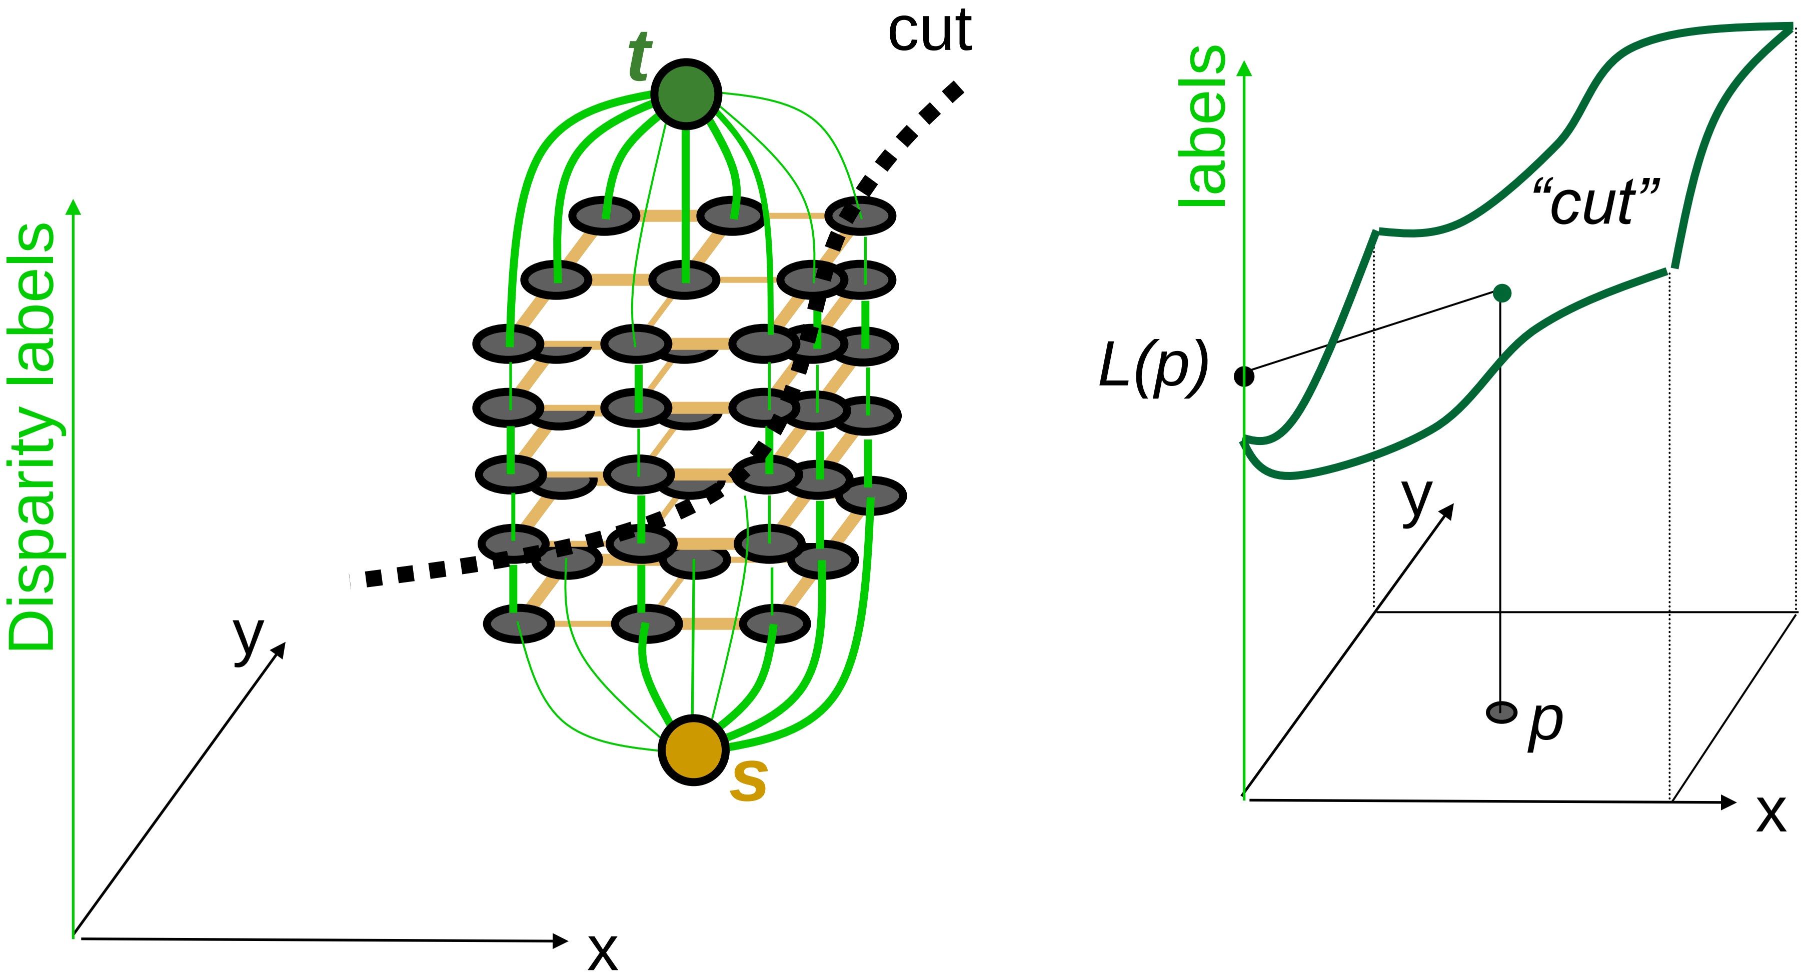

MRF optimization via graph-cuts

Let's concentrate on one pair of neighboring pixels (p,q)

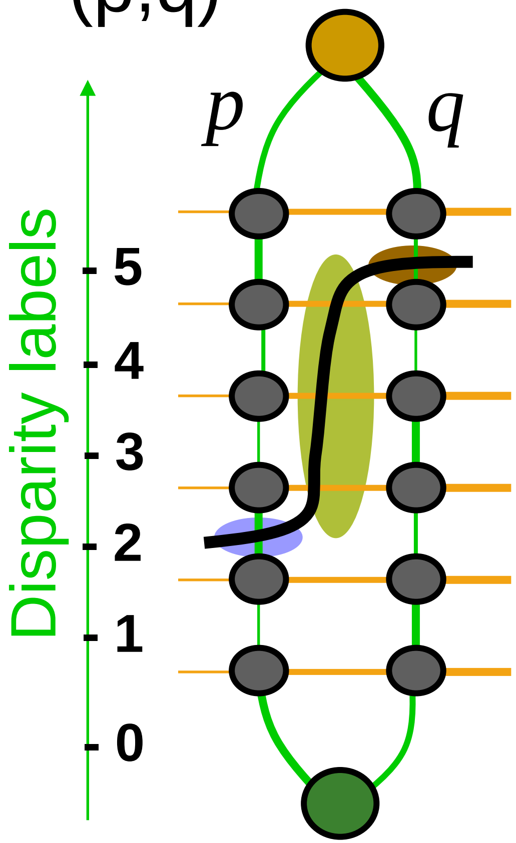

cost of vertical edges

$$E(d_p, d_q) = D_p(2) + D_q(5) + \ldots$$

$$+ \quad w_{pq} \cdot |3| \quad + \ldots$$

cost of horizontal edges

MRF optimization via graph-cuts

Let's concentrate on one pair of neighboring pixels (p,q)

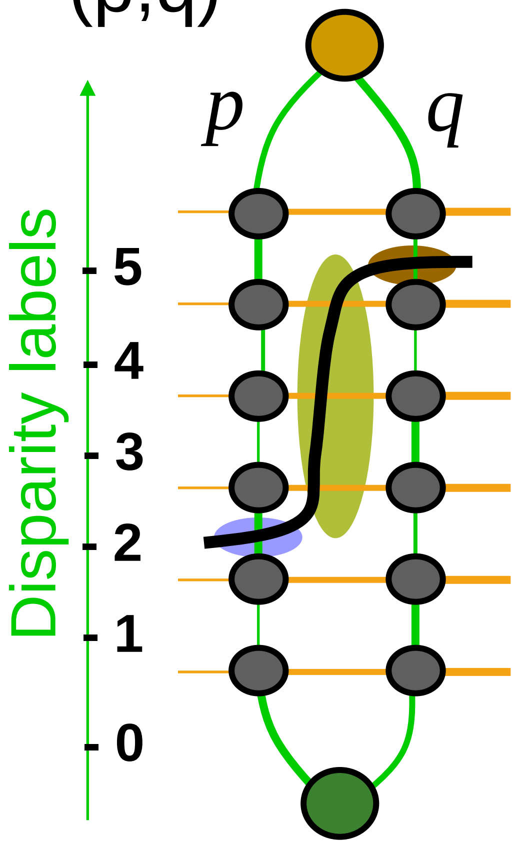

cost of vertical edges

$$E(d_p, d_q) = D_p(d_p) + D_q(d_q) + \dots$$

$$+ \quad w_{pq} \cdot | d_p - d_q | + \dots$$

cost of horizontal edges

MRF optimization via graph-cuts

The combined energy over the entire grid G is:

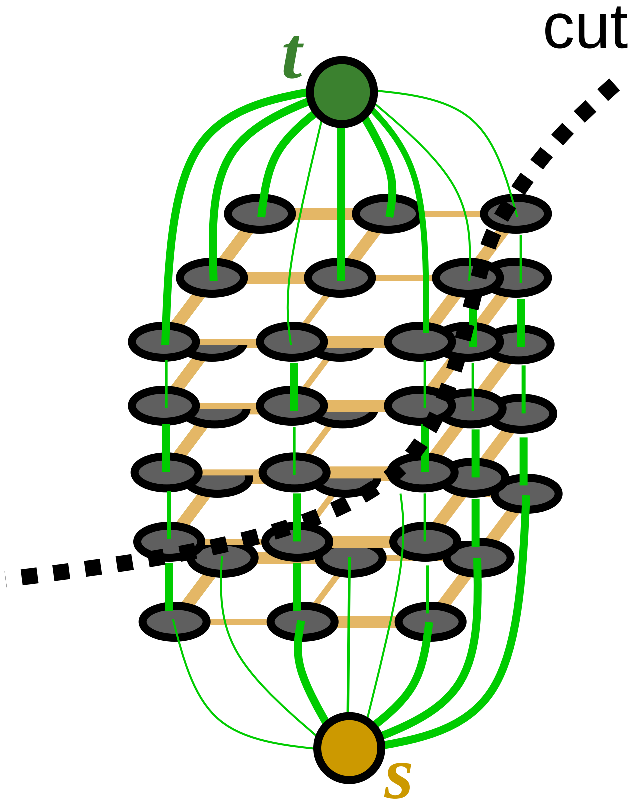

(**photo consistency**)
cost of vertical edges

$$E(\mathbf{d}) = \sum_{p \in G} D_p(d_p)$$

$$+ \sum_{\{p,q\} \in N} w_{pq} \cdot | d_p - d_q |$$

cost of horizontal edges
(**spatial consistency**)

Scan-line stereo vs. Multi-scan-line stereo

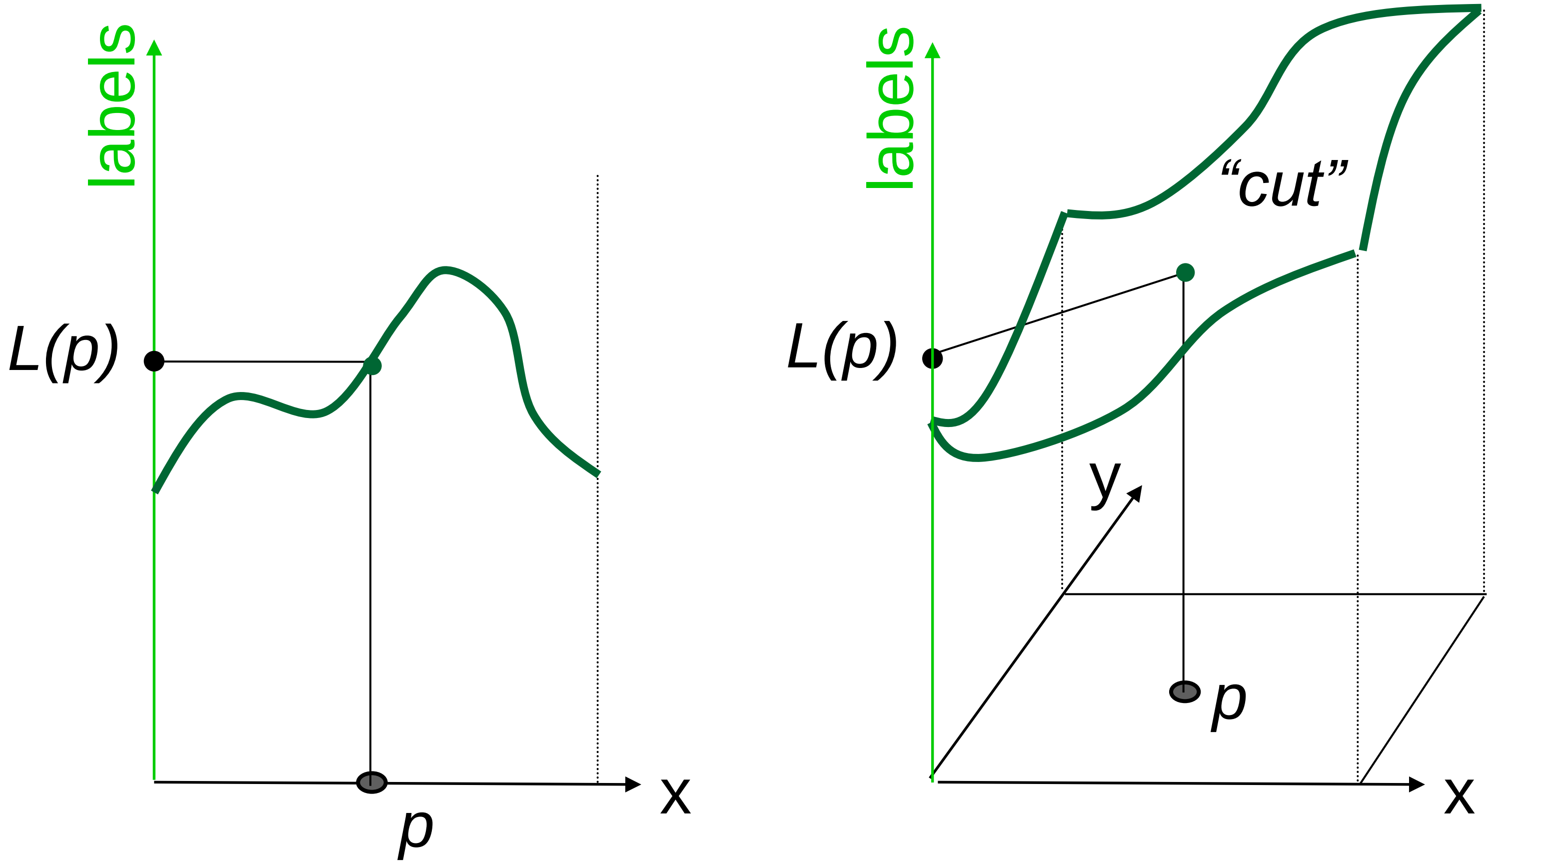

Dynamic Programming
(single scan line optimization)

s-t Graph Cuts
(multi-scan-line optimization)

Scan-line vs. graph-cut stereo

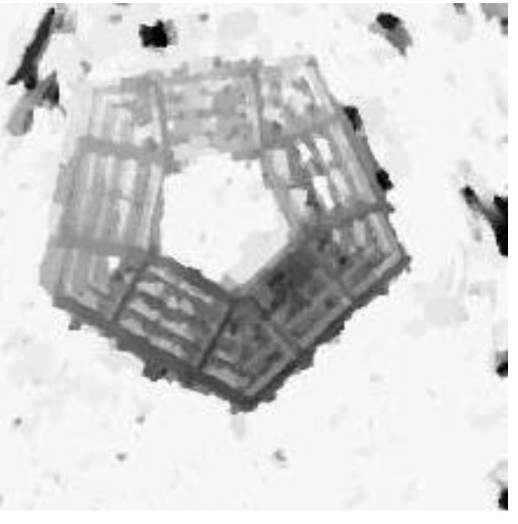

multi scan line stereo
(graph cuts)

single scan-line stereo
(DP)

Scan-line vs. graph-cut stereo

multi scan line stereo
(graph cuts)

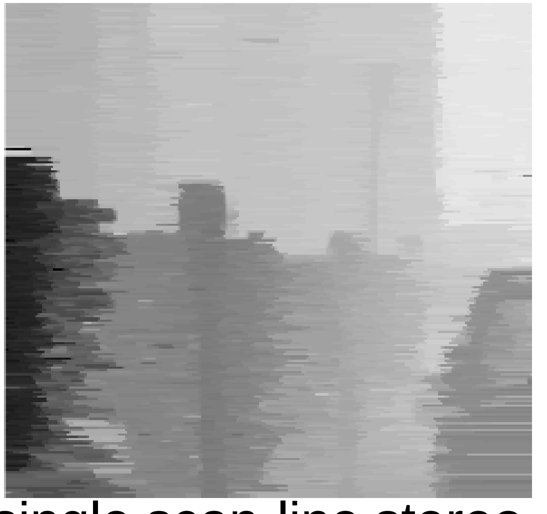

single scan-line stereo
(DP)

Multi-label energy minimization via graph-cuts (continued)

Optimizing multi-label MRFs via graph-cuts

- We saw how to use graph cuts to optimize:
  - binary MRFs
  - restricted class of multi-label MRFs
- Next
  - Using graph-cuts to optimize a more general class of MRFs
  - More general refers to the pairwise potentials

Optimizing multi-label MRFs via graph-cuts

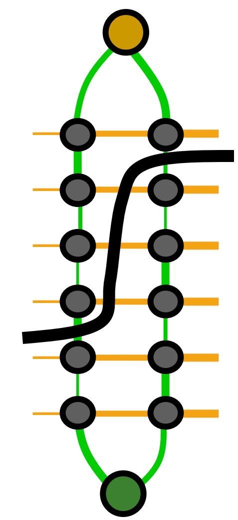

- Linear interactions
  - Can be optimized exactly using graph-cuts
  - We saw graph construction in previous class

- Actually a similar graph construction applies to "convex" interactions

  (sketch graph construction on board)

Why more general pairwise interactions?

- "convex" vs "discontinuity preserving" interactions

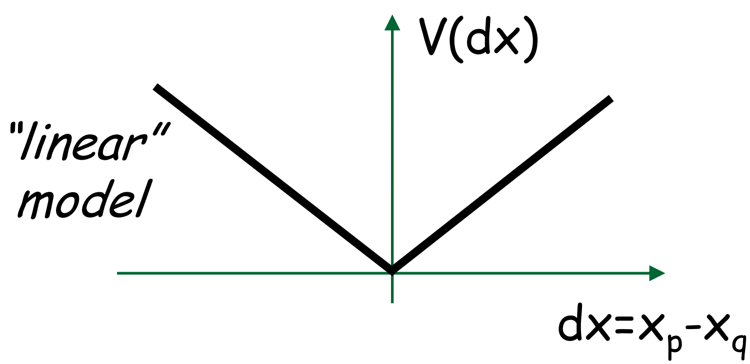

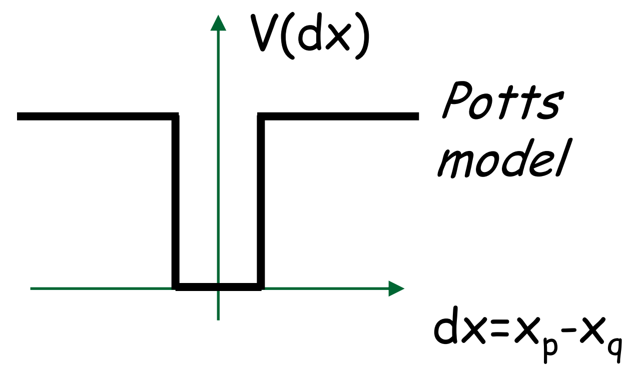

Why more general pairwise interactions?

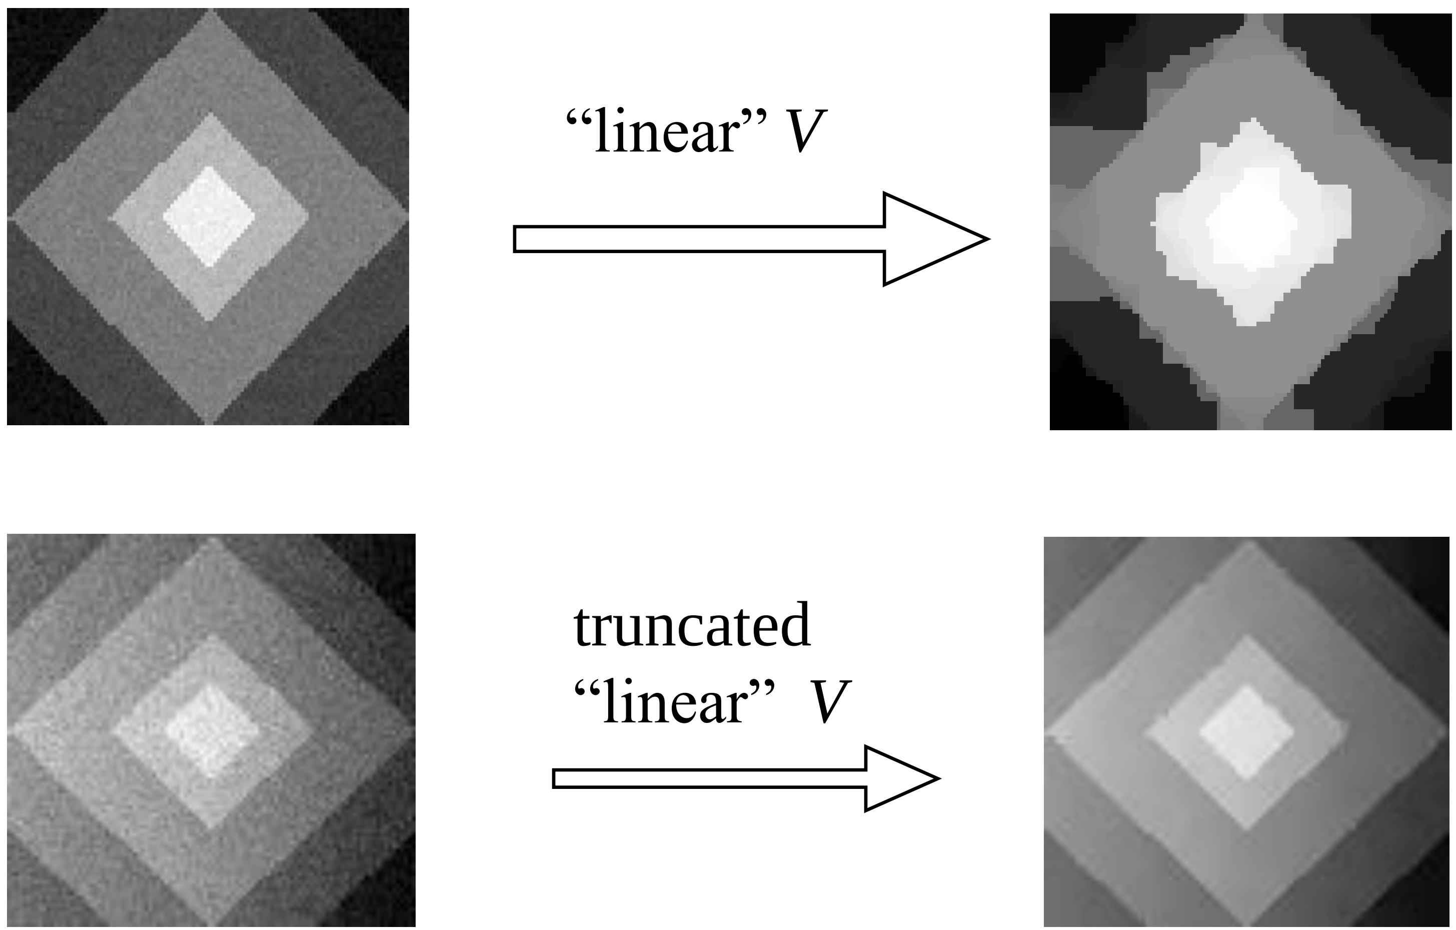

MRF hardness

- Hardness of optimizing an MRF depends critically on the type of its pairwise potentials (i.e. the distance function between labels)

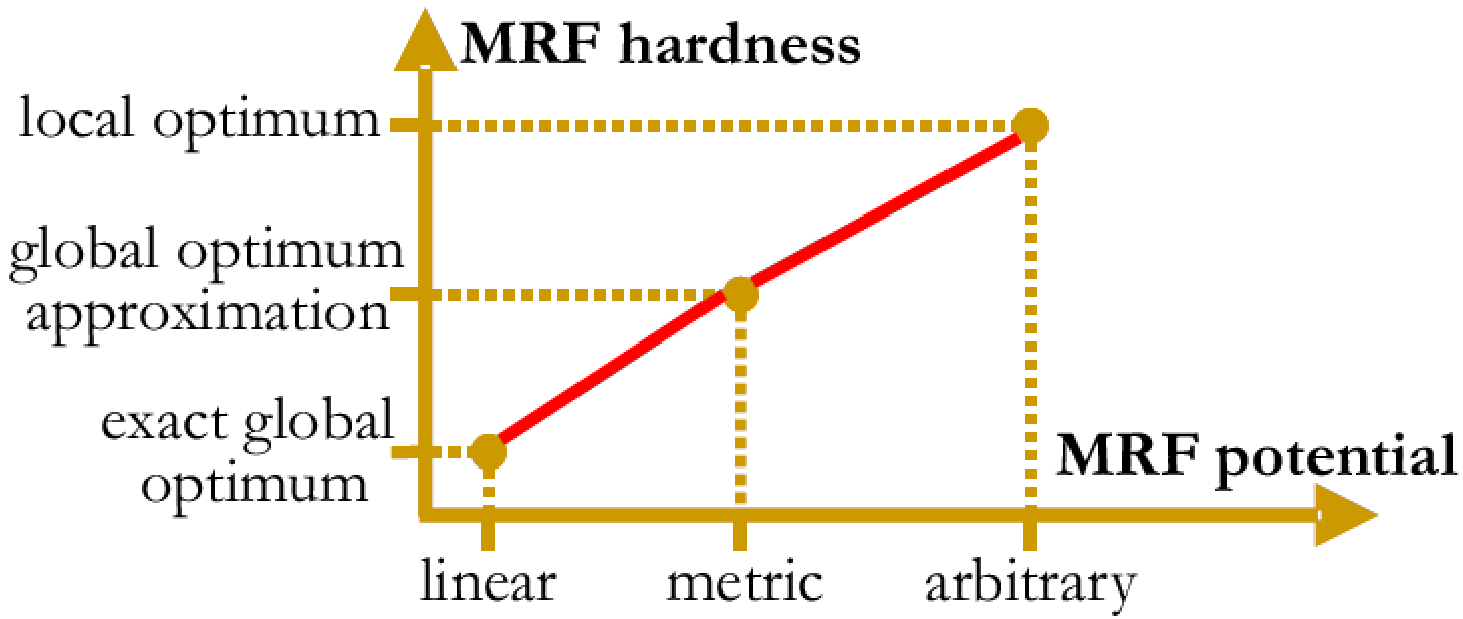

Handling more general MRFs via graph-cuts

- We will restrict the type of pairwise potentials
  - Pairwise potential is assumed to be a metric distance between labels (metric potential)
- We will derive an algorithm called **alpha-expansion**
- Approximate optimality guarantees

Alpha-expansion

- Idea: try to optimize multi-label MRF using binary MRFs
- We will have to use many binary MRFs
- Iterative algorithm
  - Binary MRF optimized at each step
- Improving our solution incrementally

Alpha-expansion

- As its name reveals, alpha-expansion relies on the so called **expansion moves**

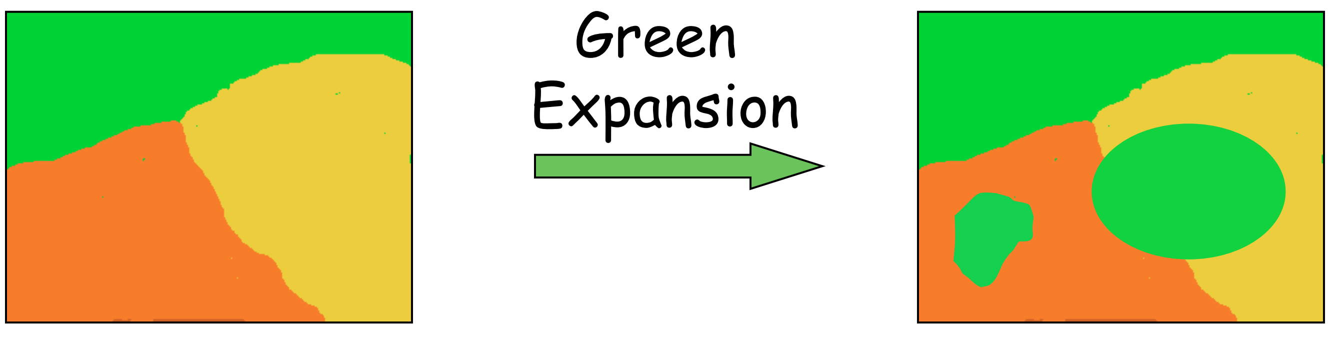

Alpha-expansion

- Optimal expansion
  - Given any label **a**, the optimal a-expansion move is the one that yields the minimum energy

- Given label **a**, how can we find the optimal a-expansion move with respect to the current configuration?

- **What does this problem reduce to?** (show on board)

Alpha-expansion for multi-label MRF

1. Start with any labeling

2. run through all labels and for each label **a**
    2a. compute optimal a-expansion move
    2b. if better energy found, accept the move

3. Stop if no label change, otherwise goto 2

Alpha-expansion for multi-label MRF

1. Start with any labeling

2. run through all labels and for each label **a**

    2a. **Binary MRF optimization**

    2b. if better energy found, accept the move

3. Stop if no label change, otherwise goto 2

Alpha-expansion example

initial solution

- -expansion
- -expansion
- -expansion
- -expansion
- -expansion
- -expansion
- -expansion

For each move we choose expansion that gives the largest decrease in the energy: **binary optimization problem**

Alpha-expansion

- How do we optimize the binary MRFs encountered during alpha-expansion?
- Well, we will reduce them to graph-cut problems
  (we already know how to do that for certain binary MRFs)

Alpha-expansion for multi-label MRF

1. Start with any labeling

2. run through all labels and for each label **a**
   2a. compute optimal a-expansion move
   2b. if better energy found, accept the move

3. Stop if no label change, otherwise goto 2

Alpha-expansion for multi-label MRF

1. Start with any labeling

2. run through all labels and for each label **a**
    2a. **Binary MRF optimization**
    2b. if better energy found, accept the move

3. Stop if no label change, otherwise goto 2

Alpha-expansion for multi-label MRF

1. Start with any labeling

2. run through all labels and for each label **a**
   - 2a. **s-t graph cut problem**
   - 2b. if better energy found, accept the move

3. Stop if no label change, otherwise goto 2

alpha-expansion thus reduces to solving a series of graph-cut problems