10/07/11

Segmentation: MRFs and Graph Cuts

Computer Vision

CS 143, Brown

James Hays

Many slides from Kristin Grauman and Derek

Image segmentation: toy example

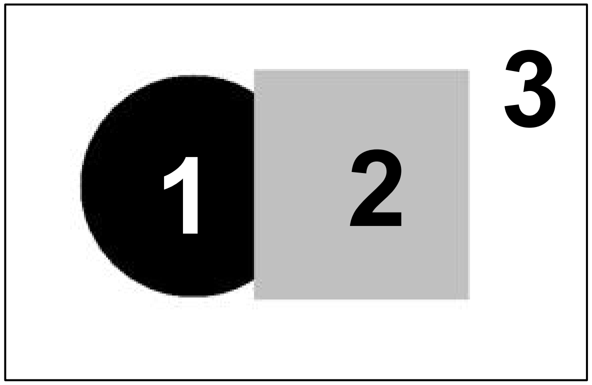

input image

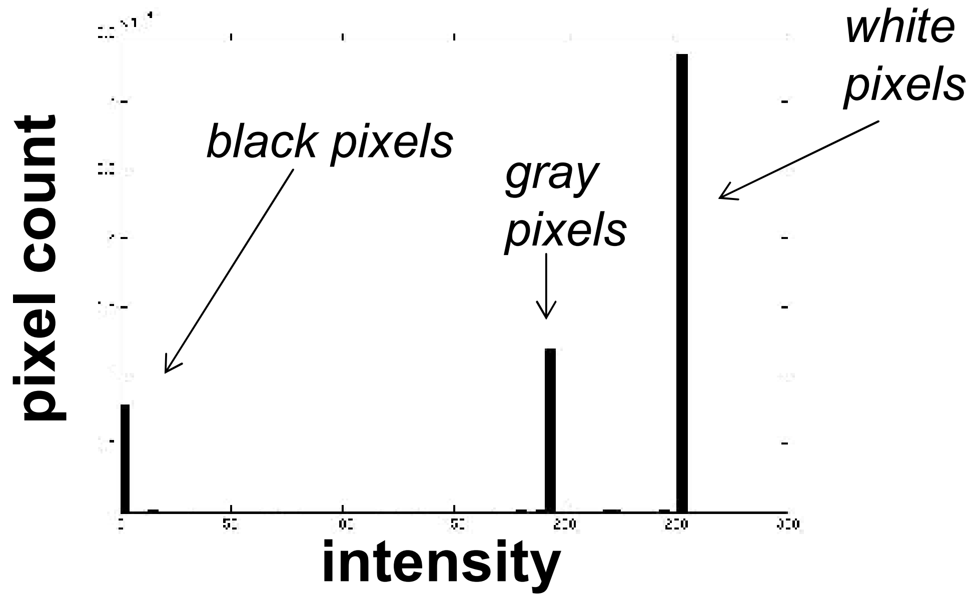

- These intensities define the three groups.
- We could label every pixel in the image according to which of these primary intensities it is.
  - i.e., *segment* the image based on the intensity feature.
- What if the image isn't quite so simple?

Kristen Grauman

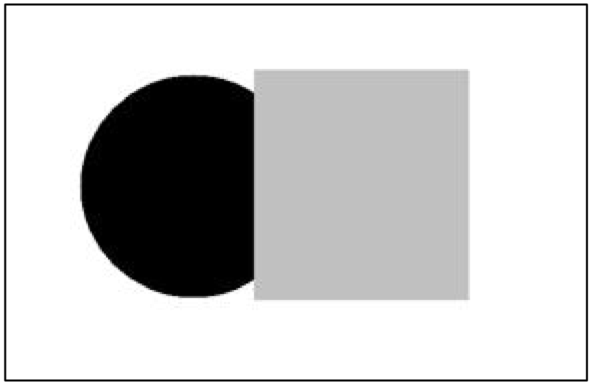

input image

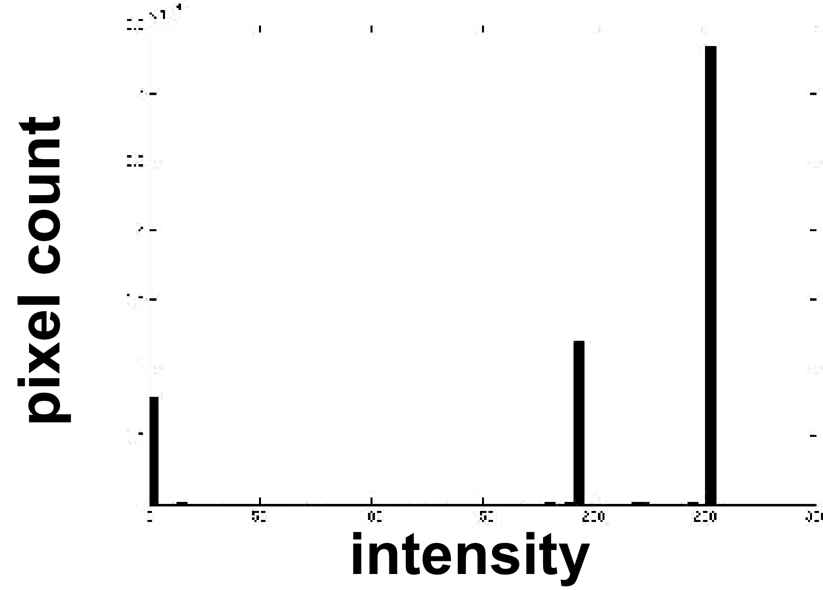

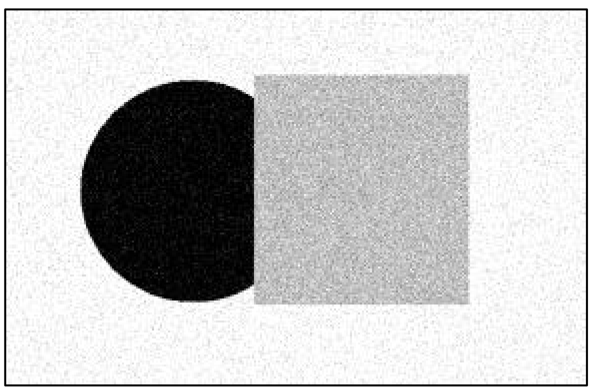

input image

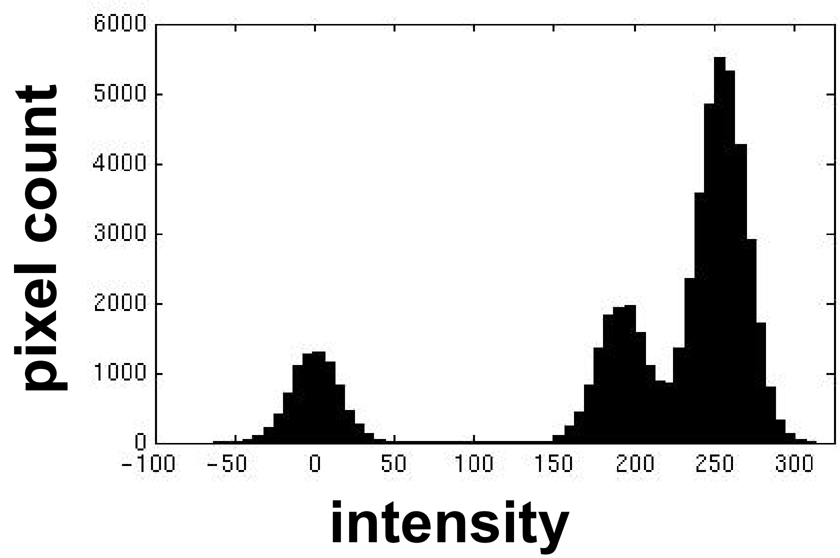

Kristen Grauman

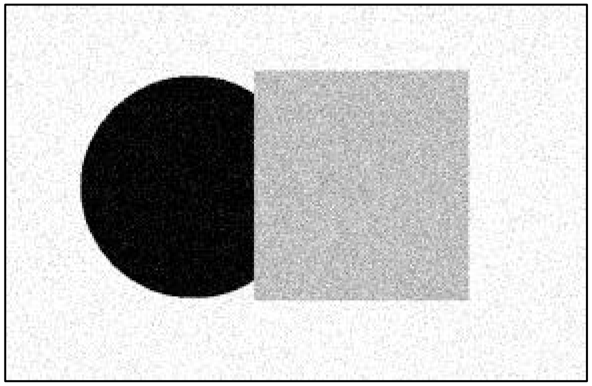

input image

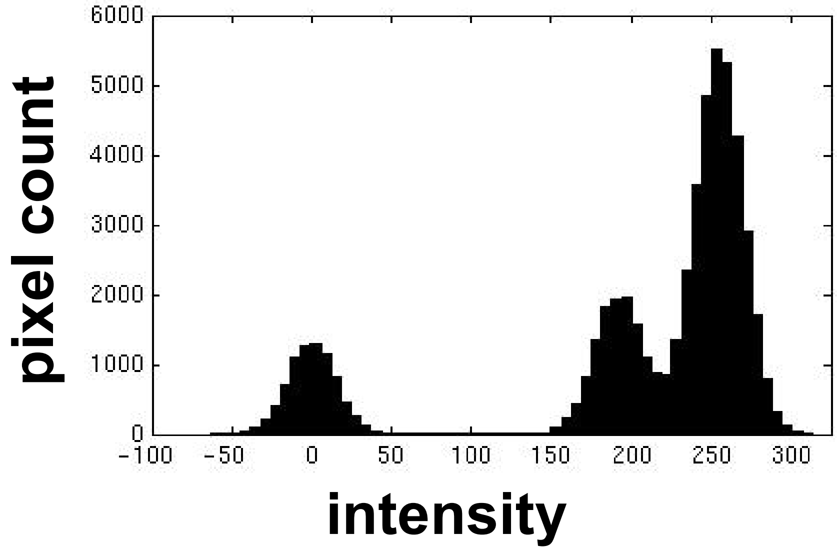

- Now how to determine the three main intensities that define our groups?
- We need to ***cluster.***

Kristen Grauman

Clustering

- With this objective, it is a "chicken and egg" problem:
  - If we knew the **cluster centers**, we could allocate points to groups by assigning each to its closest center.

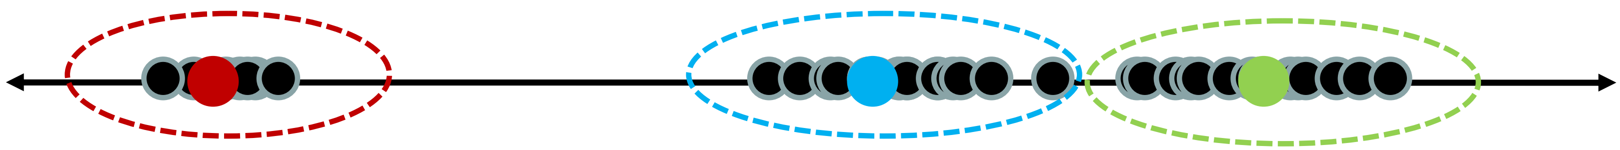

  - If we knew the **group memberships**, we could get the centers by computing the mean per group.

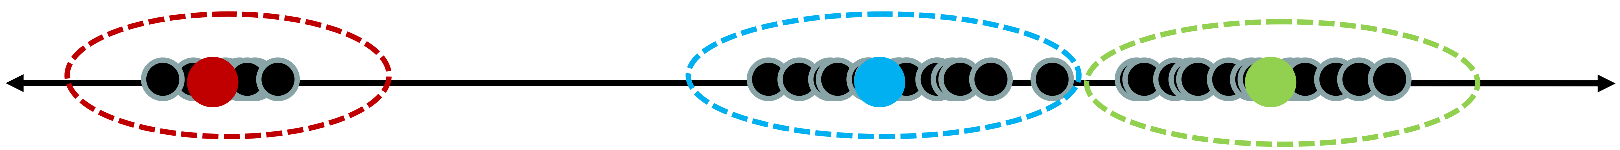

Kristen Grauman

Smoothing out cluster assignments

- Assigning a cluster label per pixel may yield outliers:

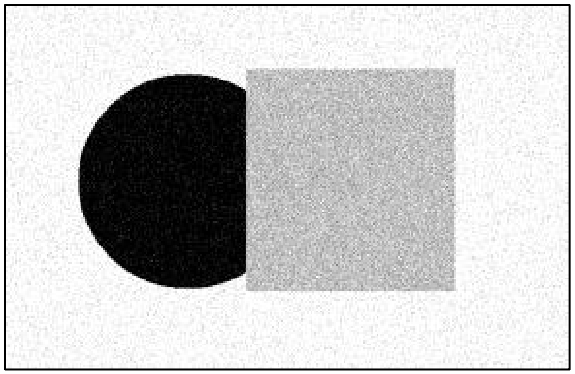

original

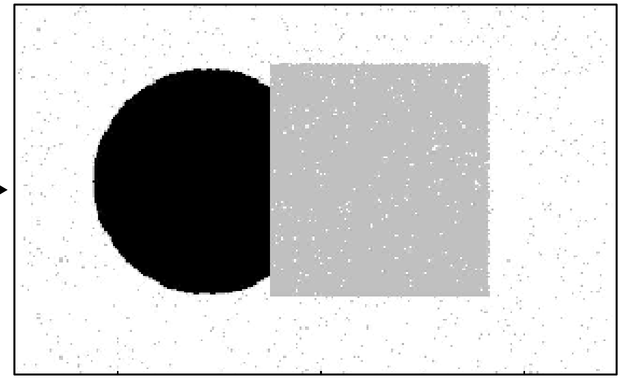

labeled by cluster center's intensity

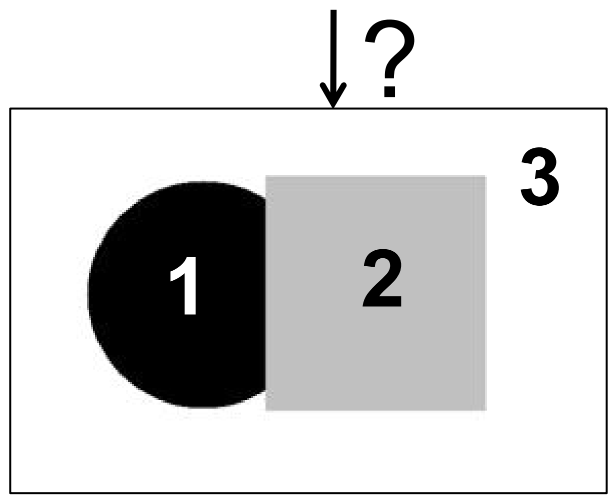

- How to ensure they are spatially smooth?

Kristen Grauman

Solution

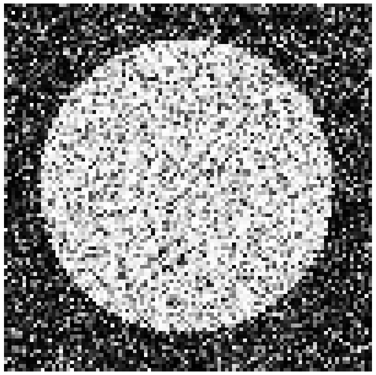

P(foreground | image)

Encode dependencies between pixels

Normalizing constant

$$P(\mathbf{y};\theta,data)=\frac{1}{Z}\prod_{i=1..N}f_1(y_i;\theta,data)\prod_{i,j\in edges}f_2(y_i,y_j;\theta,data)$$

Labels to be predicted

Individual predictions

Pairwise predictions

Slide: Derek Hoiem

Writing Likelihood as an "Energy"

$$P(\mathbf{y};\theta,data)=\frac{1}{Z}\prod_{i=1..N}p_1(y_i;\theta,data)\prod_{i,j\in edges}p_2(y_i,y_j;\theta,data)$$

$$Energy(\mathbf{y};\theta,data)=\sum_i\psi_1(y_i;\theta,data)+\sum_{i,j\in edges}\psi_2(y_i,y_j;\theta,data)$$

"Cost" of assignment y_i

"Cost" of pairwise assignment y_i, y_j

Slide: Derek Hoiem

Markov Random Fields

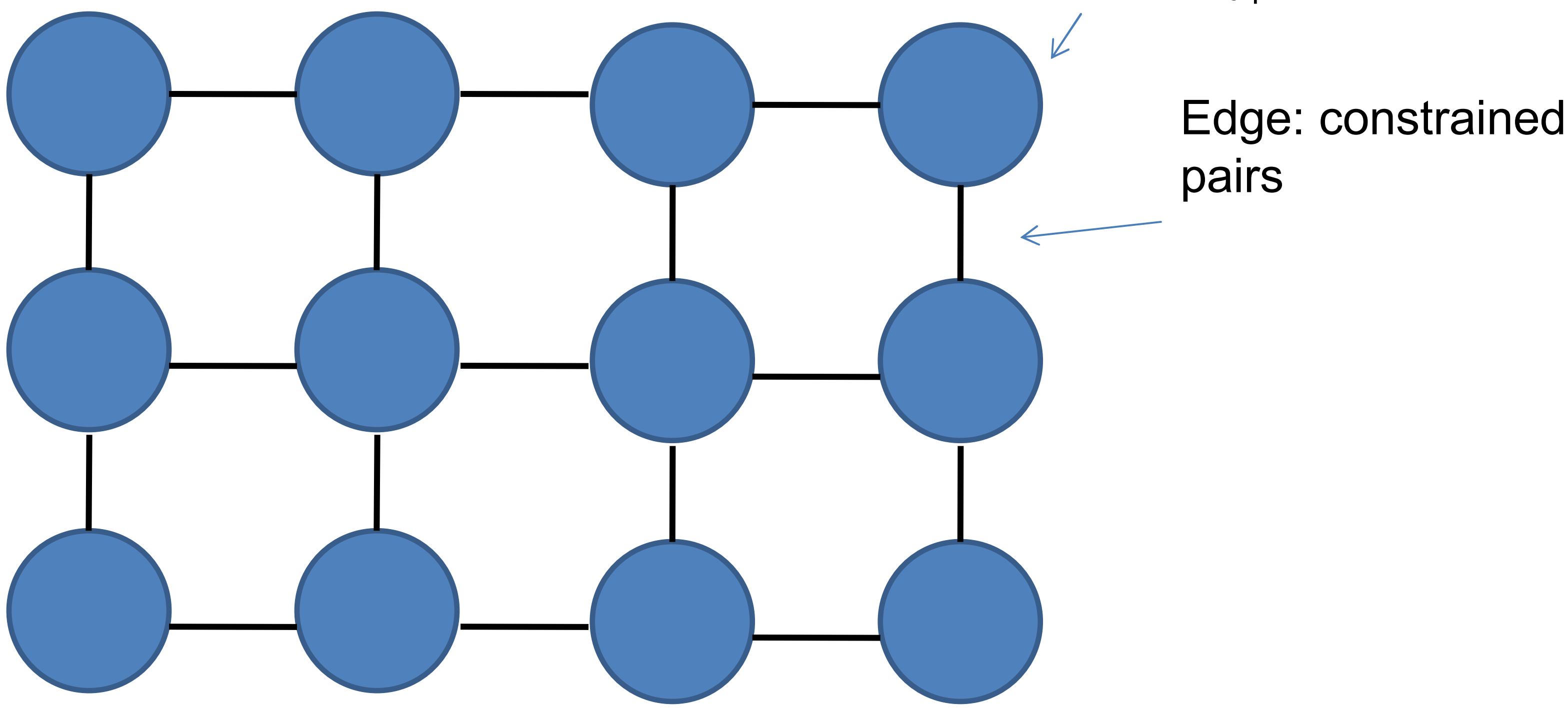

Cost to assign a label to each pixel

Cost to assign a pair of labels to connected pixels

$$Energy(\mathbf{y};\theta,data) = \sum_i \psi_1(y_i;\theta,data) + \sum_{i,j \in edges} \psi_2(y_i,y_j;\theta,data)$$

Slide: Derek Hoiem

Markov Random Fields

- Example: “label smoothing” grid

Unary potential

0: -logP(y_i = 0 ; data)
1: -logP(y_i = 1 ; data)

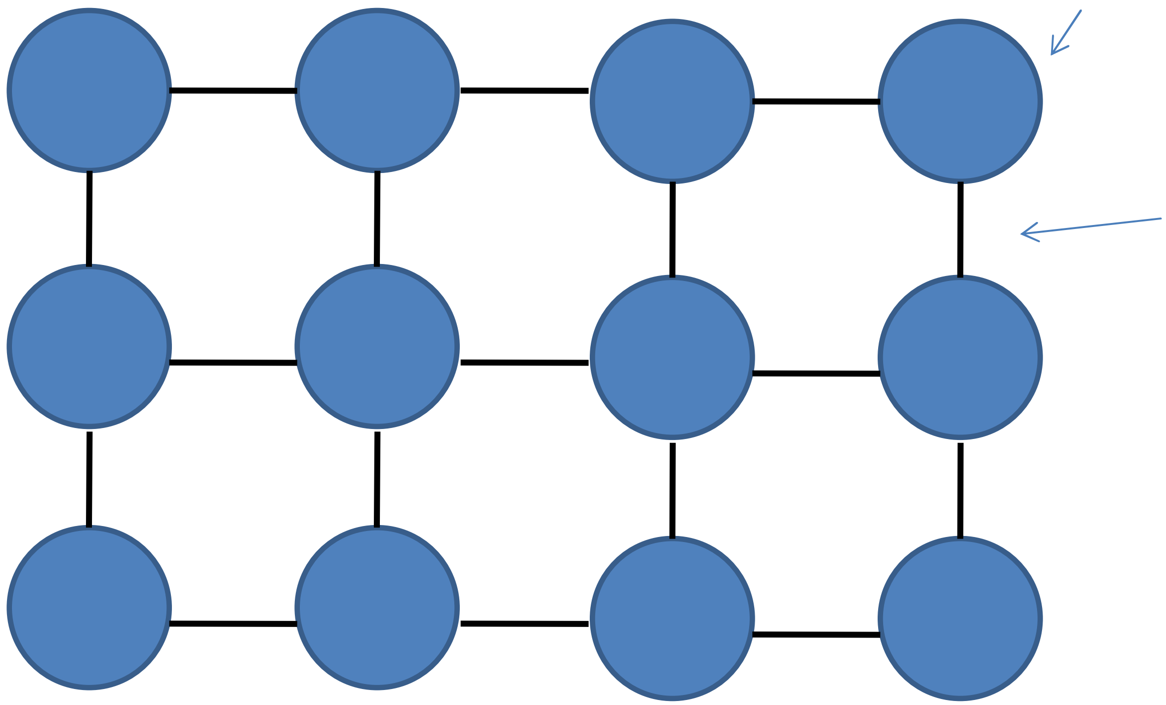

Pairwise Potential

	0	1
0	0	K
1	K	0

$$Energy(\mathbf{y};\theta,data)=\sum_i \psi_1(y_i;\theta,data)+\sum_{i,j\in edges}\psi_2(y_i,y_j;\theta,data)$$

Slide: Derek Hoiem

Solving MRFs with graph cuts

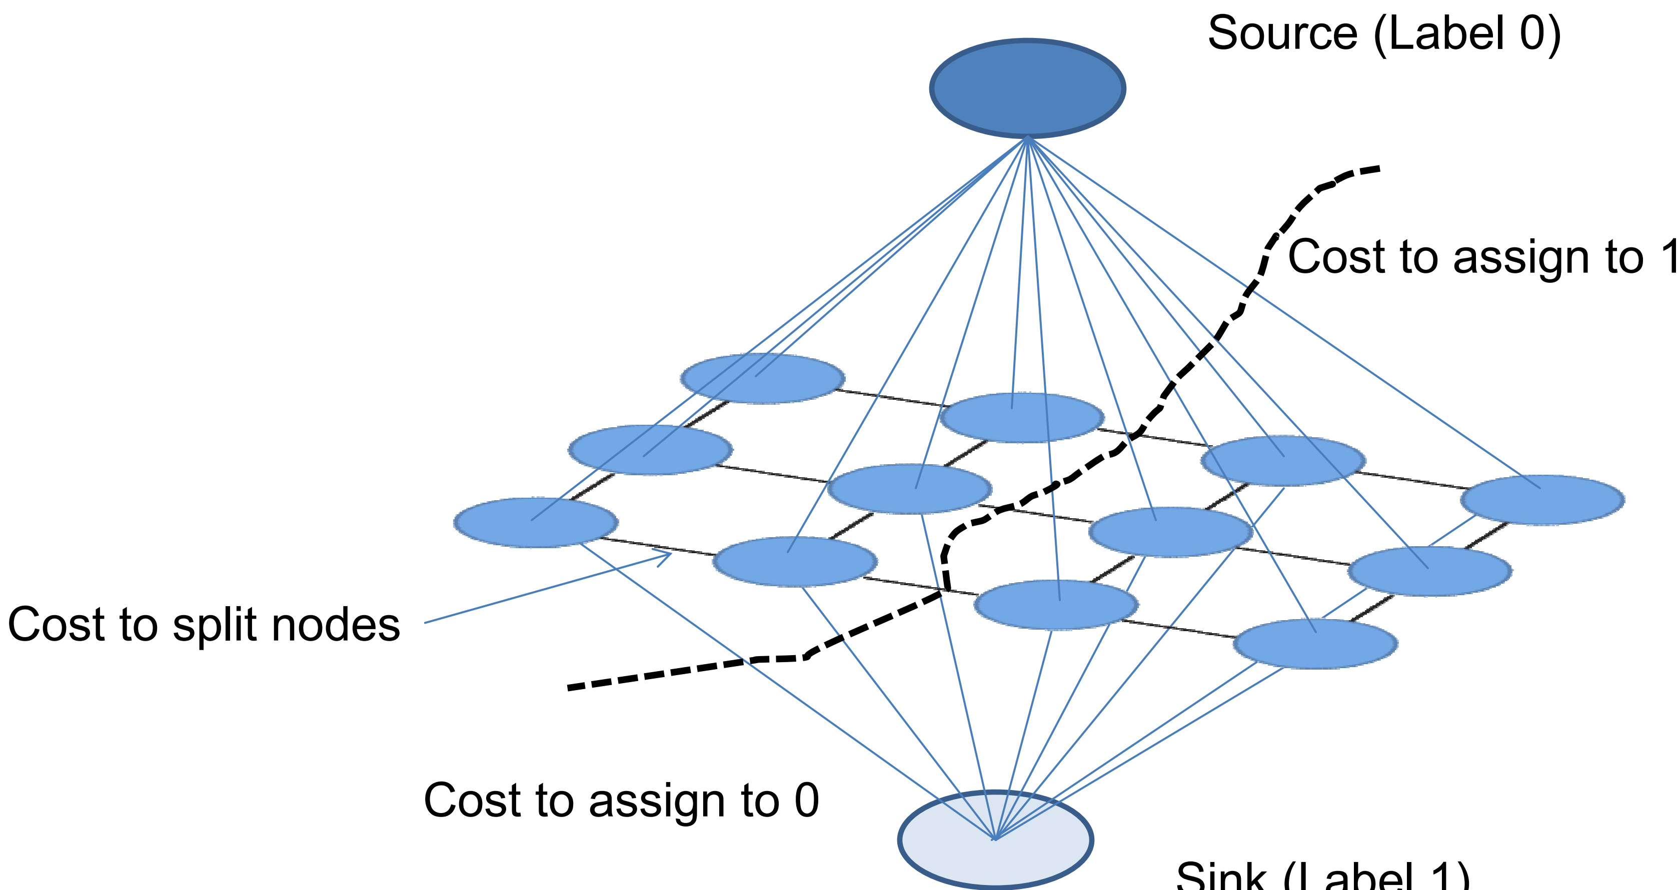

$$Energy(\mathbf{y};\theta,data)=\sum_{i}\psi_1(y_i;\theta,data)+\sum_{i,j\in edges}\psi_2(y_i,y_j;\theta,data)$$

Slide: Derek Hoiem

Solving MRFs with graph cuts

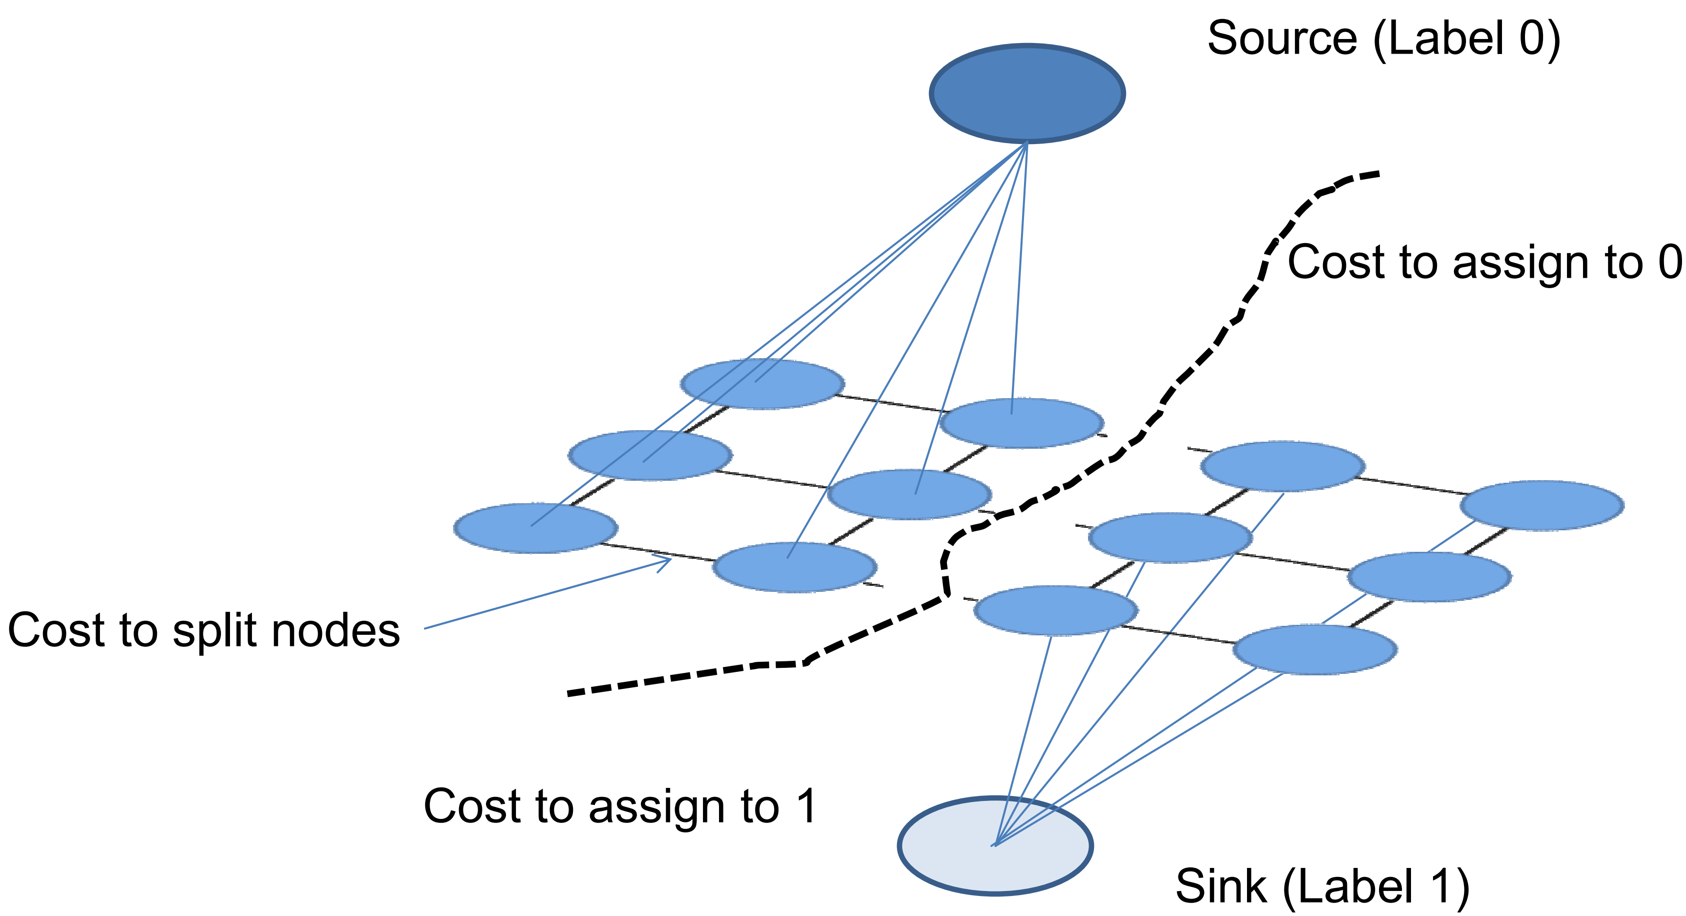

$$Energy(\mathbf{y};\theta,data)=\sum_{i}\psi_1(y_i;\theta,data)+\sum_{i,j\in edges}\psi_2(y_i,y_j;\theta,data)$$

Slide: Derek Hoiem