


Cutting Images: Graphs and Boundary Finding



“The Double Secret”, Magritte

Computational Photography
Derek Hoiem, University of Illinois

Last class: texture synthesis and transfer





Last class: in-painting



This Lecture: Finding Seams and Boundaries

Segmentation



This Lecture: Finding Seams and Boundaries

Retargeting




This Lecture: Finding Seams and Boundaries


Stitching



+



=





This Lecture: Finding seams and boundaries

Fundamental Concept: The Image as a Graph

- Intelligent Scissors: Good boundary = short path
- Graph cuts: Good region has low cutting cost

Semi-automated segmentation

User provides imprecise and incomplete specification of region – your algorithm has to read his/her mind.



Key problems

1. What groups of pixels form cohesive regions?
2. What pixels are likely to be on the boundary of regions?
3. Which region is the user trying to select?

What makes a good region?

- Contains small range of color/texture
- Looks different than background
- Compact




What makes a good boundary?

- High gradient along boundary
- Gradient in right direction
- Smooth




The Image as a Graph



Intelligent Scissors


Mortenson and Barrett (SIGGRAPH 1995)



Intelligent Scissors


Mortenson and Barrett (SIGGRAPH 1995)

A good image boundary has a short path through the graph.




Intelligent Scissors

- Formulation: find good boundary between seed points
- Challenges
 - Minimize interaction time
 - Define what makes a good boundary
 - Efficiently find it



Intelligent Scissors: method

1. Define boundary cost between neighboring pixels
2. User specifies a starting point (seed)
3. Compute lowest cost from seed to each other pixel
4. Get path from seed to cursor, choose new seed, repeat



Intelligent Scissors: method

1. Define boundary cost between neighboring pixels
 - a) Lower if edge is present (e.g., with `edge(im, 'canny')`)
 - b) Lower if gradient is strong
 - c) Lower if gradient is in direction of boundary



Gradients, Edges, and Path Cost



Gradient Magnitude




Path Cost



Edge Image


Intelligent Scissors: method

1. Define boundary cost between neighboring pixels
2. User specifies a starting point (seed)
 - Snapping



Intelligent Scissors: method

1. Define boundary cost between neighboring pixels
2. User specifies a starting point (seed)
3. Compute lowest cost from seed to each other pixel
 - Dijkstra's shortest path algorithm



Dijkstra's shortest path algorithm


Initialize, given seed s :

- Compute $\text{cost}_2(q, r)$ % cost for boundary from pixel q to neighboring pixel r
- $\text{cost}(s) = 0$ % total cost from seed to this point
- $\mathbf{A} = \{s\}$ % set to be expanded
- $\mathbf{E} = \{ \}$ % set of expanded pixels
- $\mathbf{P}(q)$ % pointer to pixel that leads to q

Loop while \mathbf{A} is not empty


1. $q =$ pixel in \mathbf{A} with lowest cost
2. Add q to \mathbf{E}
3. for each pixel r in neighborhood of q that is not in \mathbf{E}
 - a) $\text{cost_tmp} = \text{cost}(q) + \text{cost}_2(q, r)$
 - b) if (r is not in \mathbf{A}) OR ($\text{cost_tmp} < \text{cost}(r)$)
 - i. $\text{cost}(r) = \text{cost_tmp}$
 - ii. $\mathbf{P}(r) = q$
 - iii. Add r to \mathbf{A}

Example of Dijkstra's algorithm




Intelligent Scissors: method

1. Define boundary cost between neighboring pixels
2. User specifies a starting point (seed)
3. Compute lowest cost from seed to each other pixel
4. Get new seed, get path between seeds, repeat






Intelligent Scissors: improving interaction

1. Snap when placing first seed
2. Automatically adjust to boundary as user drags
3. Freeze stable boundary points to make new seeds



Where will intelligent scissors work well, or have problems?



Grab cuts and graph cuts

Magic Wand
(198?)



User
Input




Regions

Intelligent Scissors
Mortensen and Barrett (1995)



Boundary


GrabCut



Regions & Boundary


Result

Segmentation with graph cuts



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

Segmentation with graph cuts



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

Graph cuts segmentation

1. Define graph

- usually 4-connected or 8-connected

2. Set weights to foreground/background

- Color histogram or mixture of Gaussians for background and foreground

$$unary_potential(x) = -\log \left(\frac{P(c(x); \theta_{foreground})}{P(c(x); \theta_{background})} \right)$$


3. Set weights for edges between pixels

$$edge_potential(x, y) = k_1 + k_2 \exp \left\{ \frac{-\|c(x) - c(y)\|^2}{2\sigma^2} \right\}$$


4. Apply min-cut/max-flow algorithm

5. Return to 2, using current labels to compute foreground, background models

What is easy or hard about these cases for graphcut-based segmentation?



Easier examples



More difficult Examples

Camouflage &
Low Contrast



Initial
Rectangle



Initial
Result


Fine structure



Harder Case



Lazy Snapping (Li et al. SG 2004)



Limitations of Graph Cuts

- Requires associative graphs
 - Connected nodes should prefer to have the same label
- Is optimal only for binary problems

Other applications: Seam Carving

[Seam Carving – Avidan and Shamir \(2007\)](#)



[Demo: http://nparashuram.com/seamcarving/](http://nparashuram.com/seamcarving/)

Other applications: Seam Carving

- Find shortest path from top to bottom (or left to right), where cost = gradient magnitude

<http://www.youtube.com/watch?v=6NclJXTlugc>




[Seam Carving – Avidan and Shamir \(2007\)](#)

Other applications: stitching


[Graphcut Textures – Kwatra et al. SIGGRAPH 2003](#)



+



=



Other applications: stitching

[Graphcut Textures – Kwatra et al. SIGGRAPH 2003](#)



+



Ideal boundary:

1. Similar color in both images
2. High gradient in both images

Summary of big ideas

- Treat image as a graph
 - Pixels are nodes
 - Between-pixel edge weights based on gradients
 - Sometimes per-pixel weights for affinity to foreground/background
- Good boundaries are a short path through the graph (Intelligent Scissors, Seam Carving)
- Good regions are produced by a low-cost cut (GrabCuts, Graph Cut Stitching)

Take-home questions

1. What would be the result in “Intelligent Scissors” if all of the edge costs were set to 1?
2. How could you change boundary costs for graph cuts to work better for objects with many thin parts?

Next class

- Compositing and blending