Computational Photography

Prof. Feng Liu

Spring 2022

http://www.cs.pdx.edu/~fliu/courses/cs510/

05/05/2022

With slides by F. Durand, Y.Y. Chuang, R. Raskar, and C. Rother

Last Time

- Image segmentation
  - Normalized cut and segmentation

Today

- Segmentation
  - Interactive image segmentation

Magic Wand (Photoshop) | Intelligent Scissors, Mortensen and Barrett (1995) | GrabCut, Rother et al. 2004

User Input

Result

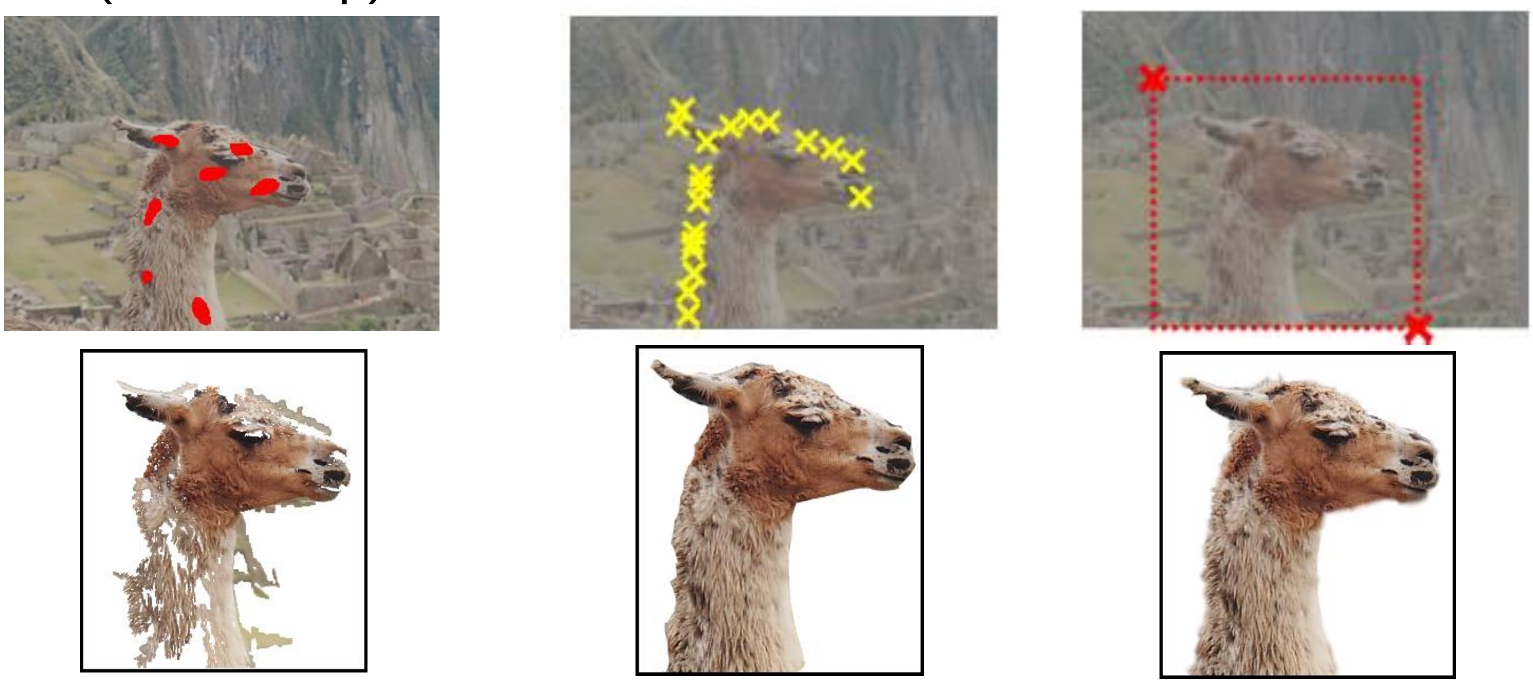

Image credit: Rother et al. 2004

Start

- Segmentation
  - Interactive image segmentation

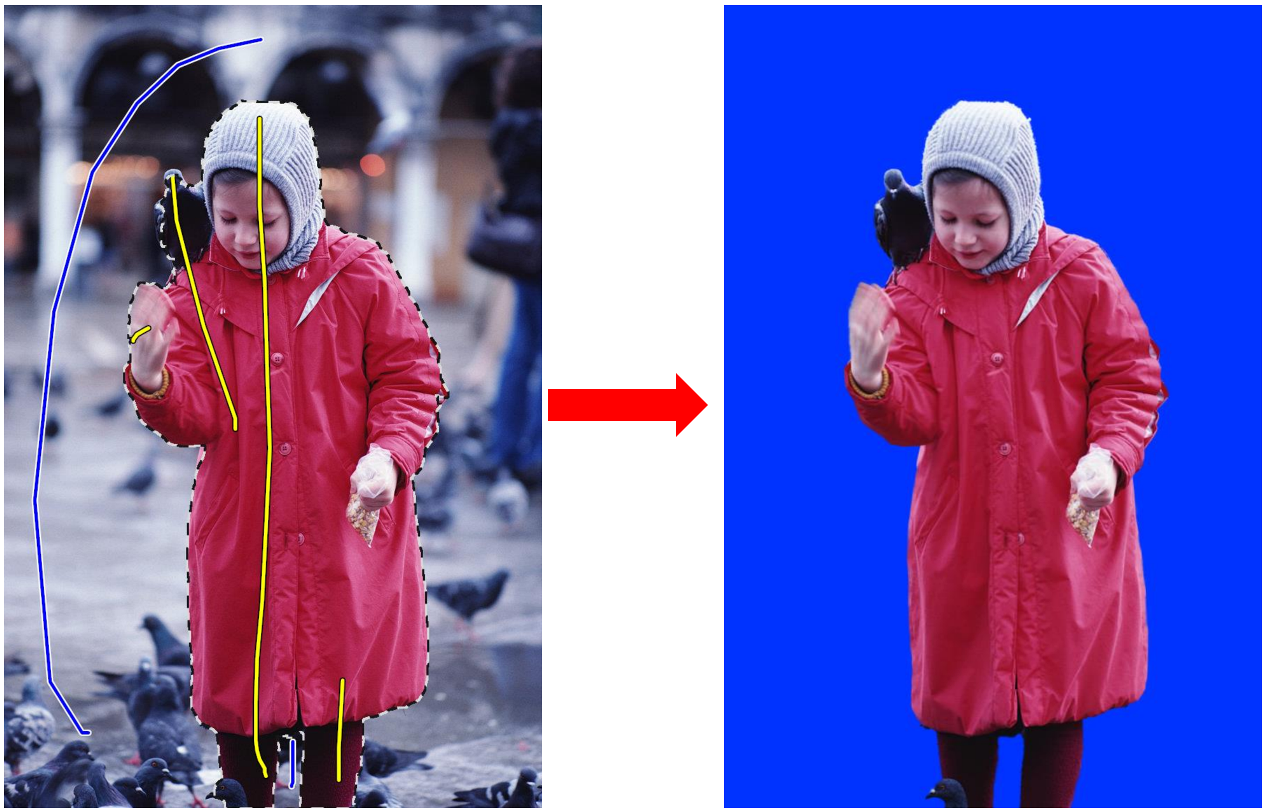

Segmentation by Graph Cut

- Interactive image segmentation using graph cut
- Binary label: foreground vs. background
- User labels some pixels
  - usually sparser
- Exploit
  - Statistics of known Fg & Bg
  - Smoothness of label
- Turn into discrete graph optimization
  - Graph cut (min cut / max flow)

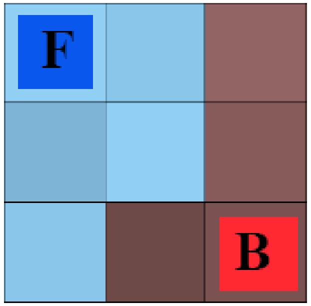

F	F	B
F	F	B
F	B	B

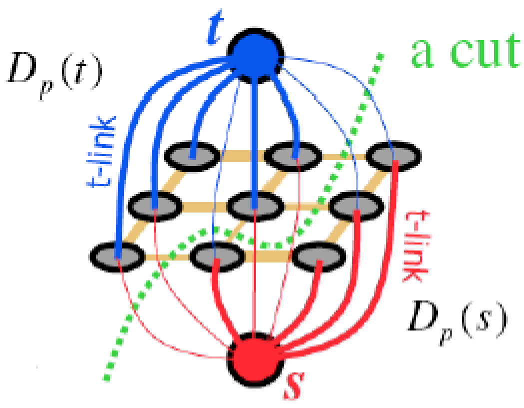

Slide credit: Y.Y. Chuang

Energy function

- Segmentation as Labeling
  - one value per pixel, F or B
- Energy(labeling) = data + smoothness
  - Very general situation
  - Will be minimized
- Data: for each pixel
  - Probability that this color belongs to F (resp. B)
- Smoothness (aka regularization): per neighboring pixel pair
  - Penalty for having different label
  - Penalty is down-weighted if the two pixel colors are very different
  - Similar in spirit to bilateral filter

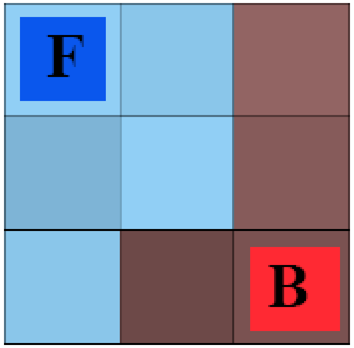

One labeling (ok, not best)

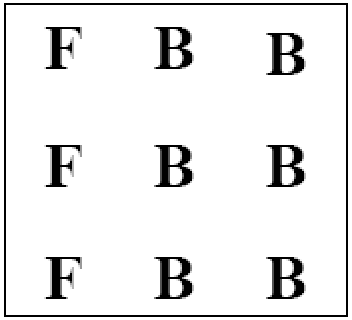

Data

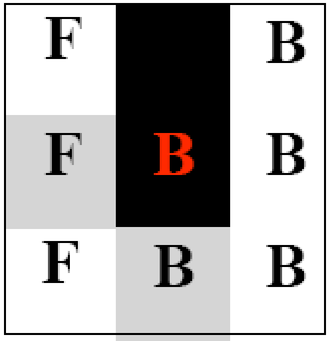

Smoothness

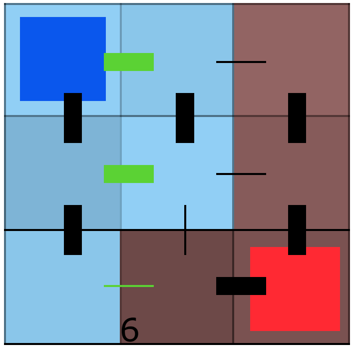

Slide credit: F. Durand

Data term

- A.k.a regional term
  (because integrated over full region)
- $D(L)=\Sigma_i$ -log $h[L_i](C_i)$
- Where *i* is a pixel
  L_i is the label at *i* (F or B),
  C_i is the pixel value
  $h[L_i]$ is the histogram of
  the observed Fg (resp Bg)
- Note the minus sign

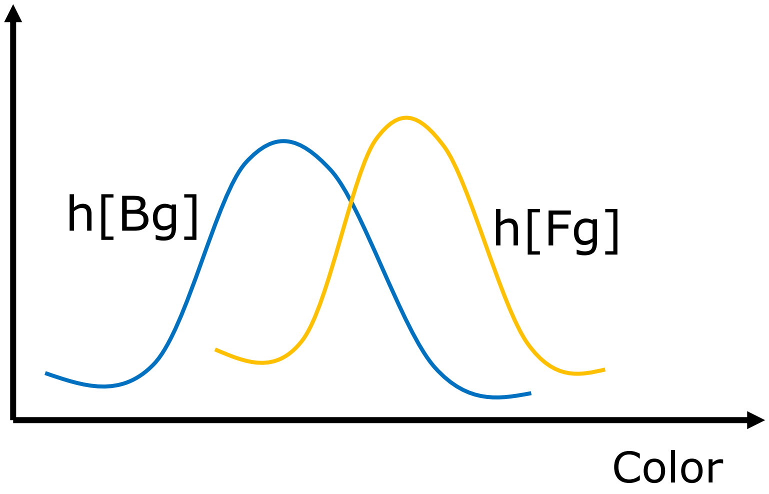

Slide credit: F. Durand

Hard constraints

- The user has provided some labels
- The quick and dirty way to include constraints into optimization is to replace the data term by a huge penalty if not respected.
- D(L_i)=0 if respected
- D(L_i)=K if not respected
  - e.g. K= #pixels

Smoothness term

- a.k.a boundary term, a.k.a. regularization
- $S(L)=\Sigma_{\{j,\ i\}\ \text{in}\ N}\ B(C_i,C_j)\ \delta(L_i\text{-}L_j)$
- Where i,j are neighbors
  - e.g. 8-neighborhood
    (but I show 4 for simplicity)
- $\delta(L_i\text{-}L_j)$ is 0 if $L_i=L_j$, 1 otherwise
- $B(C_i,C_j)$ is high when C_i and C_j are similar, low if there is a discontinuity between those two pixels
  - e.g. $\exp(-||C_i\text{-}C_j||^2/2\sigma^2)$
  - where σ can be a constant
    or the local variance
- Note positive sign

F	B	B
F	B	B
F	B	B

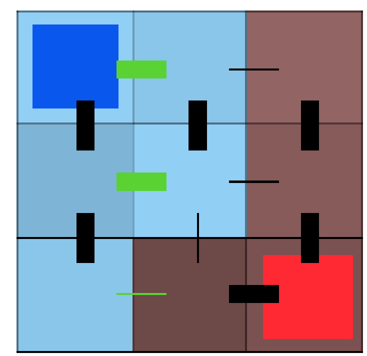

Slide credit: F. Durand

Optimization

- E(L)=D(L)+λ S(L)
- λ is a black-magic constant
- Find the labeling that minimizes E
- In this case, how many possibilities?
  - 2^9 (512)
  - We can try them all!
  - What about megapixel images?

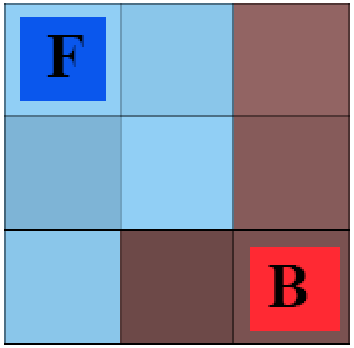

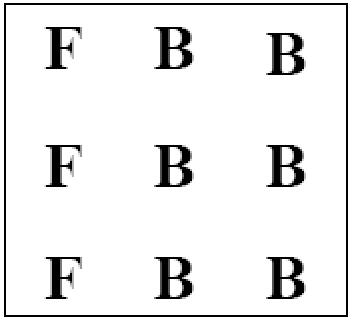

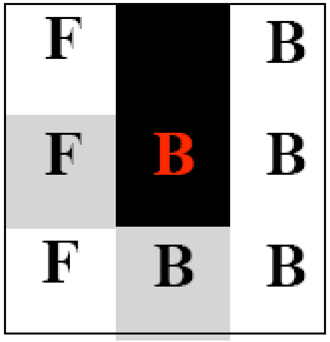

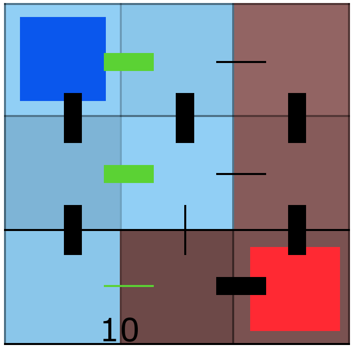

Slide credit: F. Durand

Labeling as a graph problem

- Each pixel = node
- Add two nodes F & B
- Labeling: link each pixel to either F or B

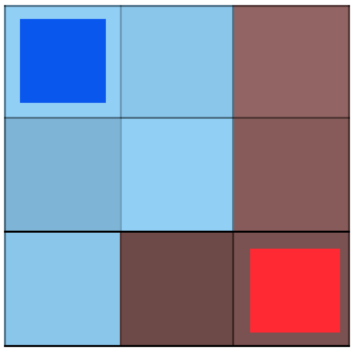

F	F	B
F	F	B
F	B	B

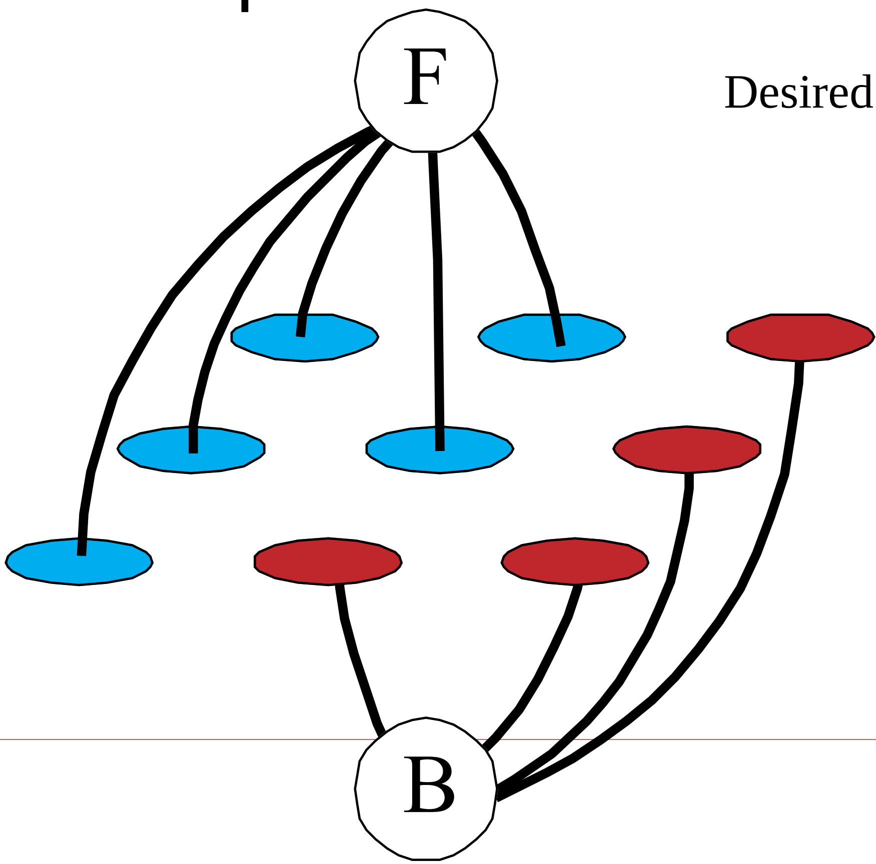

Desired result

Slide credit: F. Durand

Data term

- Put one edge between each pixel and F & B
- Weight of edge = minus data term
  - Don't forget huge weight for hard constraints
  - Careful with sign

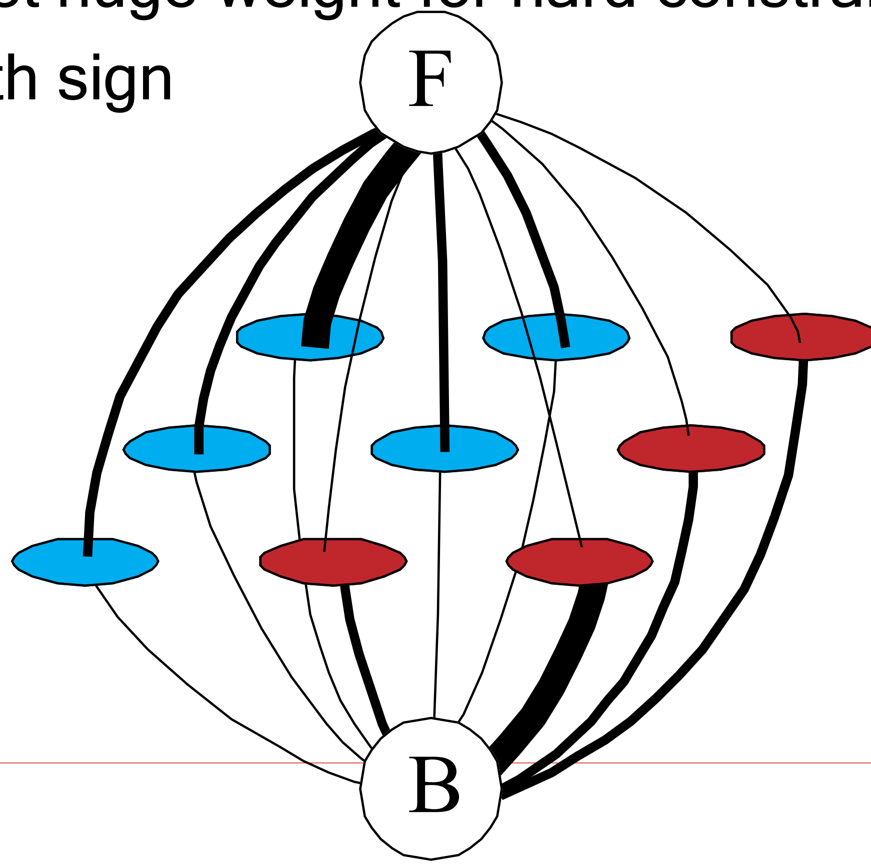

Slide credit: F. Durand

Smoothness term

- Add an edge between each neighbor pair
- Weight = smoothness term

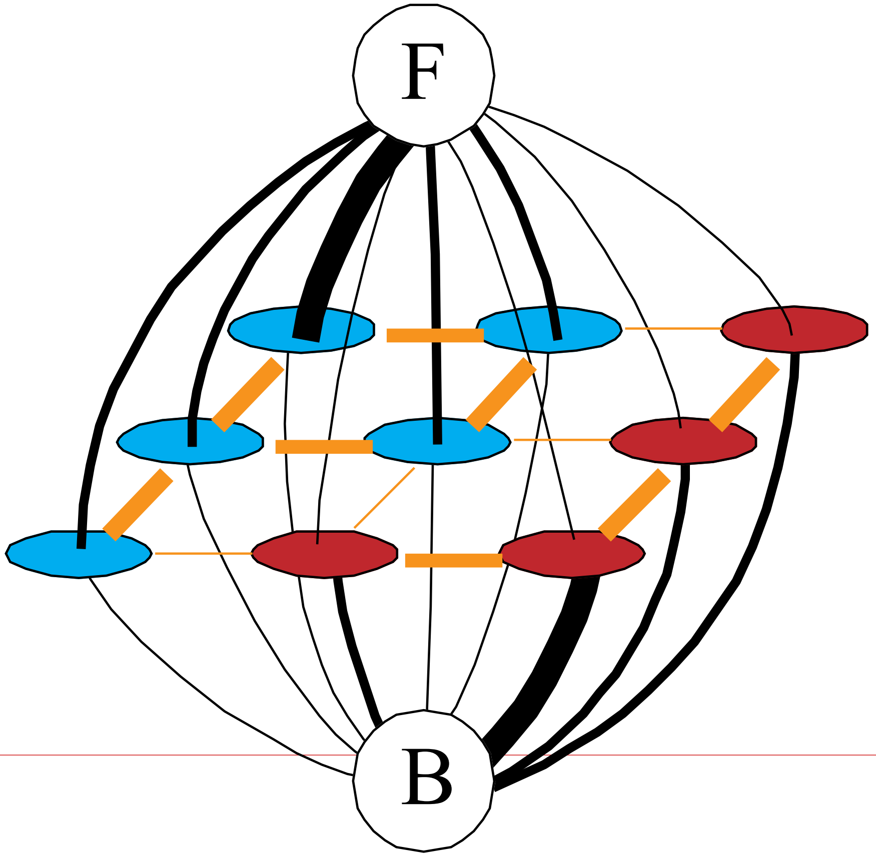

Slide credit: F. Durand

Min cut

- Energy optimization equivalent to min cut
- Cut: remove edges to disconnect F from B
- Minimum: minimize sum of cut edge weight

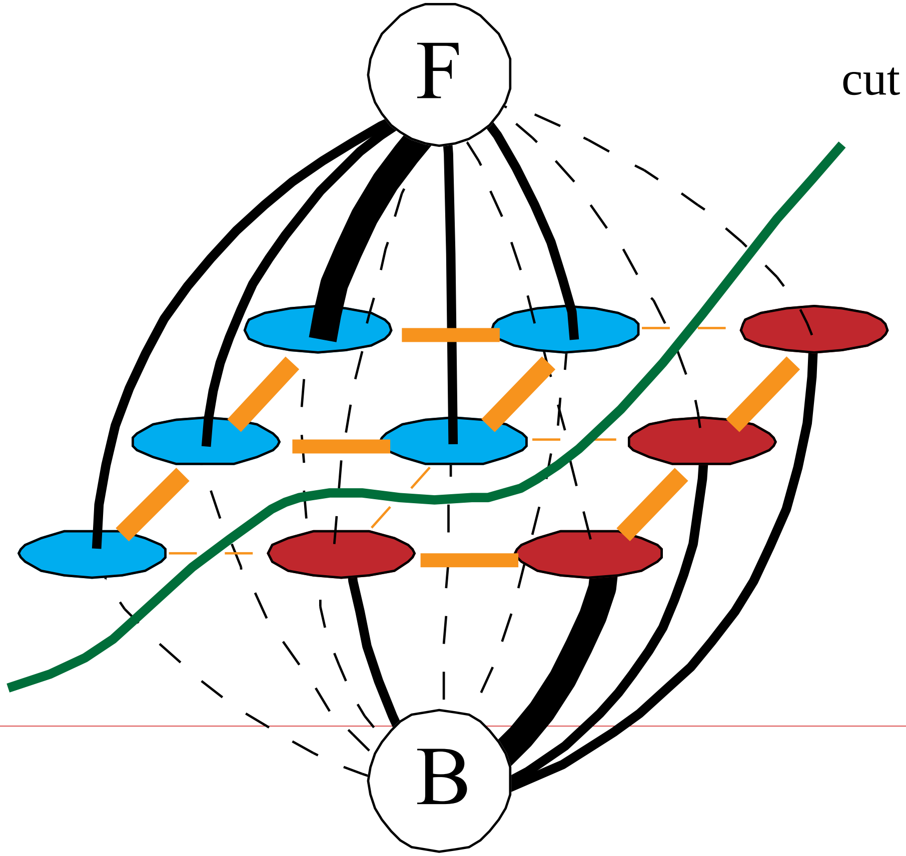

Slide credit: F. Durand

Min cut <=> labeling

- In order to be a cut:
  - For each pixel, either the F or G edge has to be cut
- In order to be minimal
  - Only one edge label per pixel can be cut (otherwise could be added)

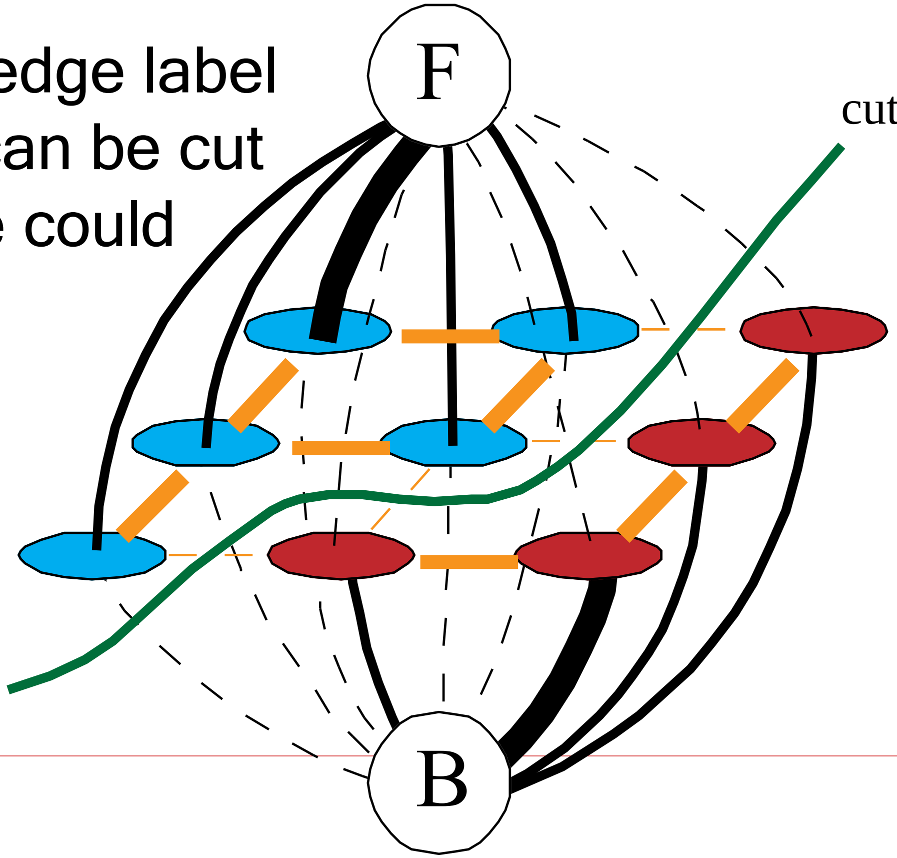

Slide credit: F. Durand

Computing a multiway cut

- With 2 labels: classical min-cut problem
  - Solvable by standard flow algorithms
    - polynomial time in theory, nearly linear in practice
    - Code: C++ from OpenCV
      - Matlab wrapper: http://www.wisdom.weizmann.ac.il/~bagon/matlab.html
  - More than 2 terminals: NP-hard [Dahlhaus *et al.*, STOC '92]
    Code: http://vision.ucla.edu/~brian/gcmex.html
- Efficient approximation algorithms exist
  - Within a factor of 2 of optimal
  - Computes local minimum in a strong sense
    - even very large moves will not improve the energy
  - Yuri Boykov, Olga Veksler and Ramin Zabih, [Fast Approximate Energy Minimization via Graph Cuts](), International Conference on Computer Vision, September 1999.

Slide credit: Y.Y. Chuang

GrabCut

Interactive Foreground Extraction using Iterated Graph Cuts

Carsten Rother
Vladimir Kolmogorov
Andrew Blake

Microsoft Research Cambridge-UK

Photomontage

SIGGRAPH2004

Framework

- **Input**: Image $\mathbf{x} \in \{\mathbf{R},\mathbf{G},\mathbf{B}\}^{\mathbf{n}}$
- **Output**: Segmentation $\mathbf{S} \in \{\mathbf{0},\mathbf{1}\}^{\mathbf{n}}$
- **Parameters:** Colour $\mathbf{\Theta}$,Coherence λ
- **Energy:** $E(\mathbf{\Theta},\mathbf{S},\mathbf{x},\lambda) = E_{Col} + E_{Coh}$
- **Optimization:** $\arg\min_{\mathbf{S},\mathbf{\Theta}} E(\mathbf{S},\mathbf{\Theta},\mathbf{x},\lambda)$

Graph Cuts

Boykov and Jolly (2001)

Image

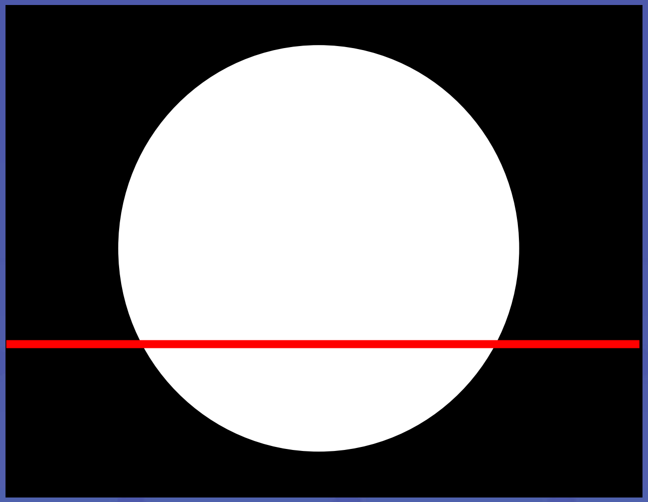

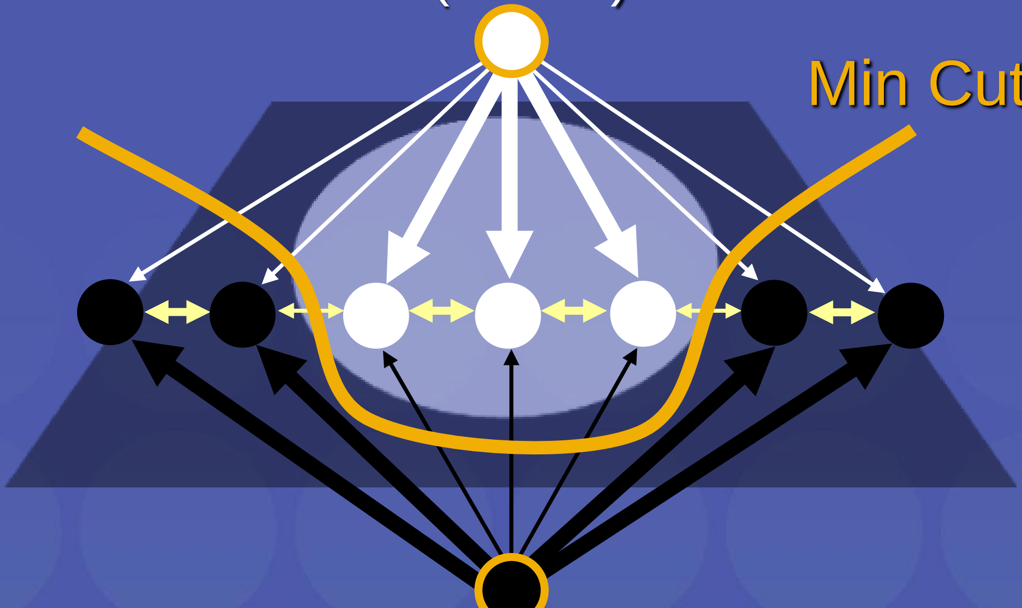

Cut: separating source and sink; Energy: collection of edges

Min Cut: Global minimal enegry in polynomial time

Iterated Graph Cut

User Initialisation

$$\arg\min_{\Theta} E(\mathbf{S}, \Theta, \mathbf{x}, \lambda)$$

$$\arg\min_{\mathbf{S}} E(\mathbf{S}, \Theta, \mathbf{x}, \lambda)$$

K-means for learning colour distributions

Graph cuts to infer the segmentation

Iterated Graph Cuts

Result

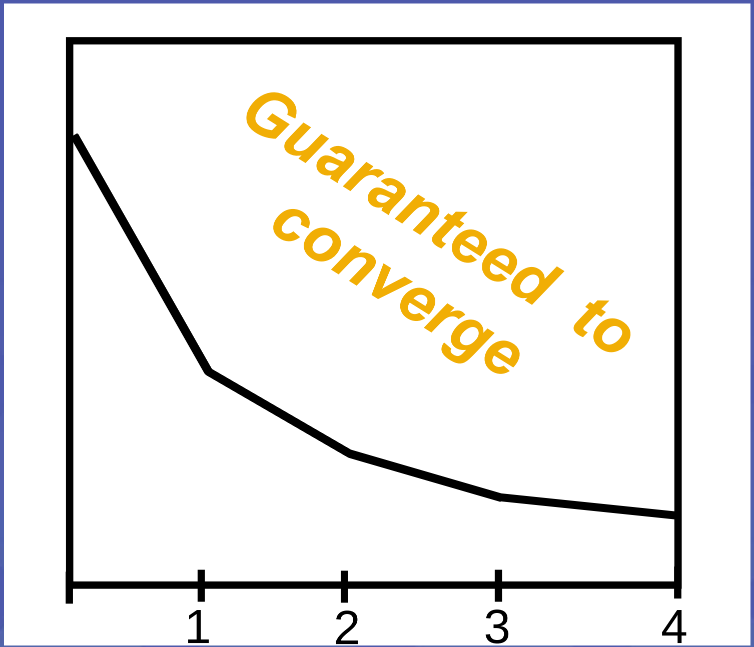

Energy after each Iteration

Colour Model

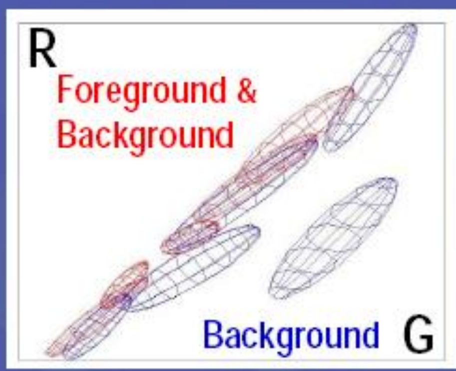

Iterated graph cut

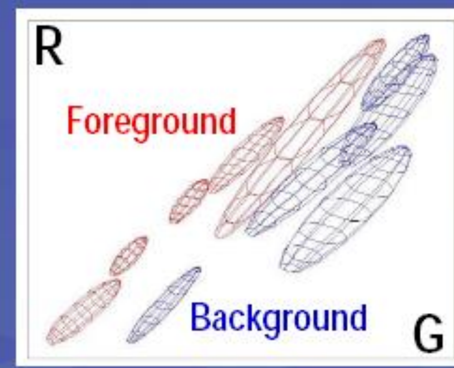

Gaussian Mixture Model (typically 5-8 components)

$$E_{Col}(\Theta, \mathbf{S}, \mathbf{x}) = \sum_{\mathbf{n}} \mathbf{D}(\mathbf{S_n}, \Theta, \mathbf{x_n})$$

Coherence Model

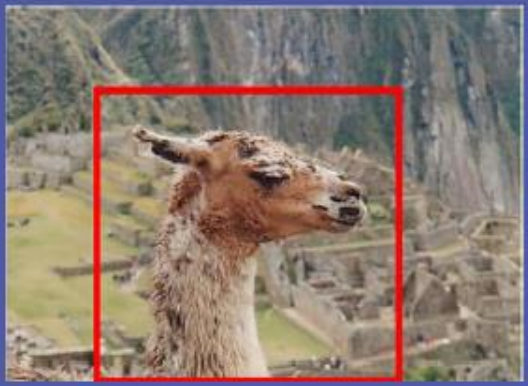

An object is a coherent set of pixels:

$$E_{coh}(\mathbf{S}, \mathbf{x}, \lambda) = \lambda \sum_{i,j \text{ adj.}} (S_i \neq S_j)\ exp\{-\frac{1}{2\sigma^2}||x_i - x_j||^2\}$$

$\lambda = 0$

$\lambda = 50$

$\lambda = 1000$

Blake et al. (2004): Learn Θ, λ jointly

Moderately straightforward examples

SIGGRAPH2004

… GrabCut completes automatically

Difficult Examples

SIGGRAPH2004

	Camouflage & Low Contrast	Fine structure	No telepathy
Initial Rectangle			
Initial Result			

Evaluation – Labelled Database

Available online: http://research.microsoft.com/vision/cambridge/segmentation/

Comparison

SIGGRAPH2004

	Boykov and Jolly (2001)	GrabCut
User Input		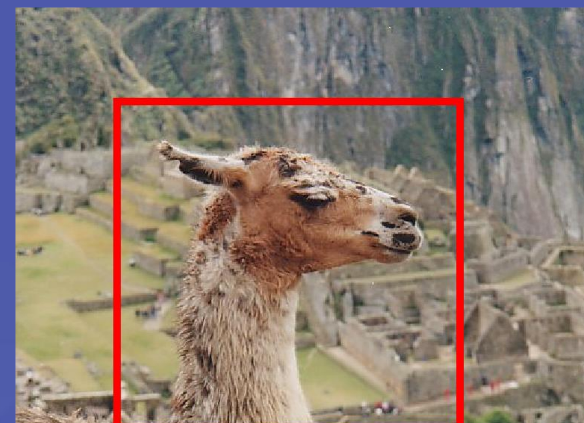
Result	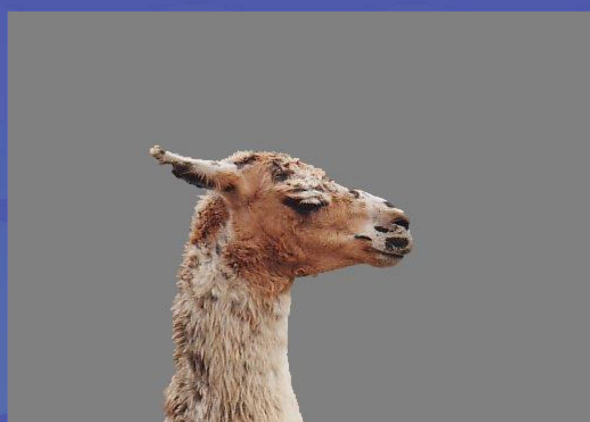	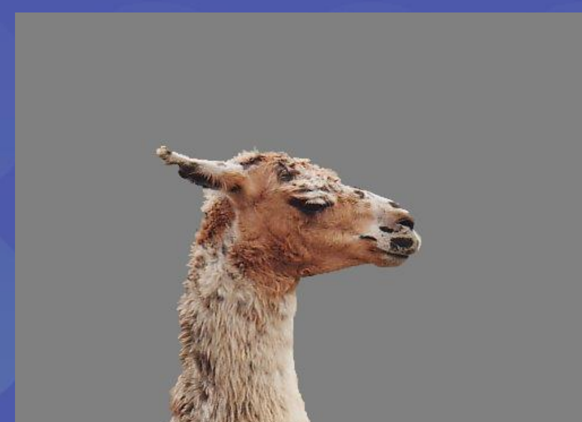
	Error Rate: 0.72%	**Error Rate: 0.72%**

Summary

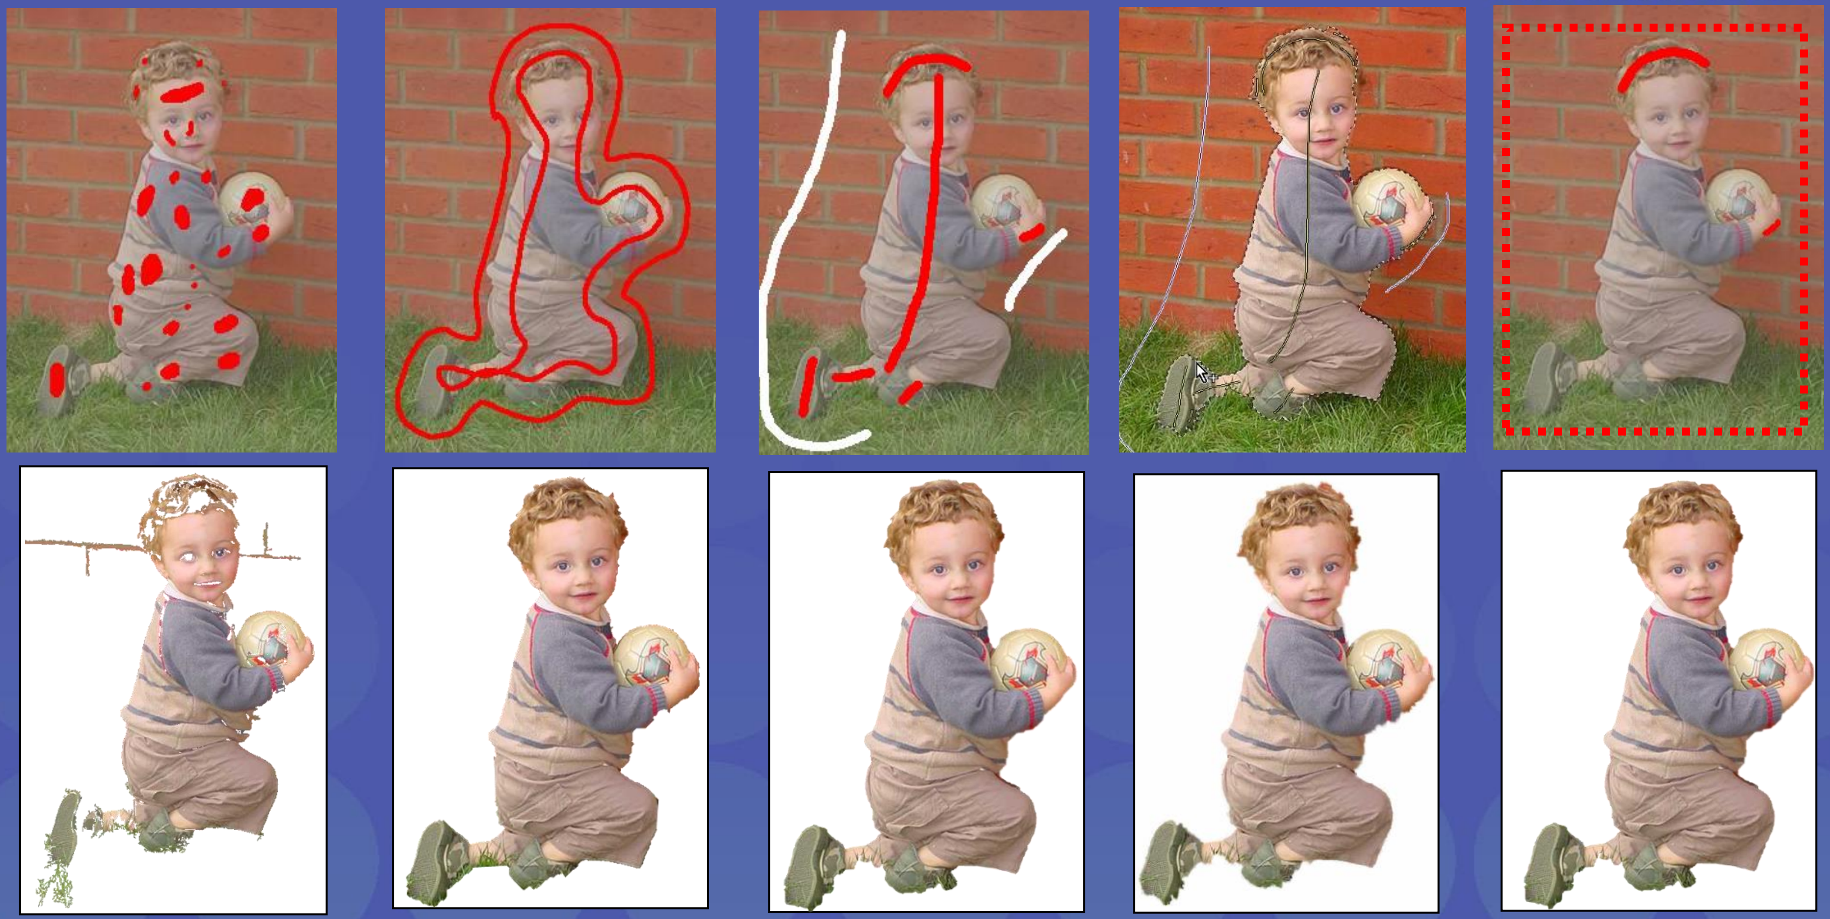

Magic Wand (198?)

Intelligent Scissors Mortensen and Barrett (1995)

Graph Cuts Boykov and Jolly (2001)

LazySnapping Li et al. (2004)

GrabCut Rother et al. (2004)

Interactive Digital Photomontage

- Combining multiple photos
- Find seams using graph cuts
- Combine gradients and integrate

Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David Salesin, Michael Cohen, “Interactive Digital Photomontage”, SIGGRAPH 2004

Slide credit: Y.Y. Chuang

Slide credit: Y.Y. Chuang

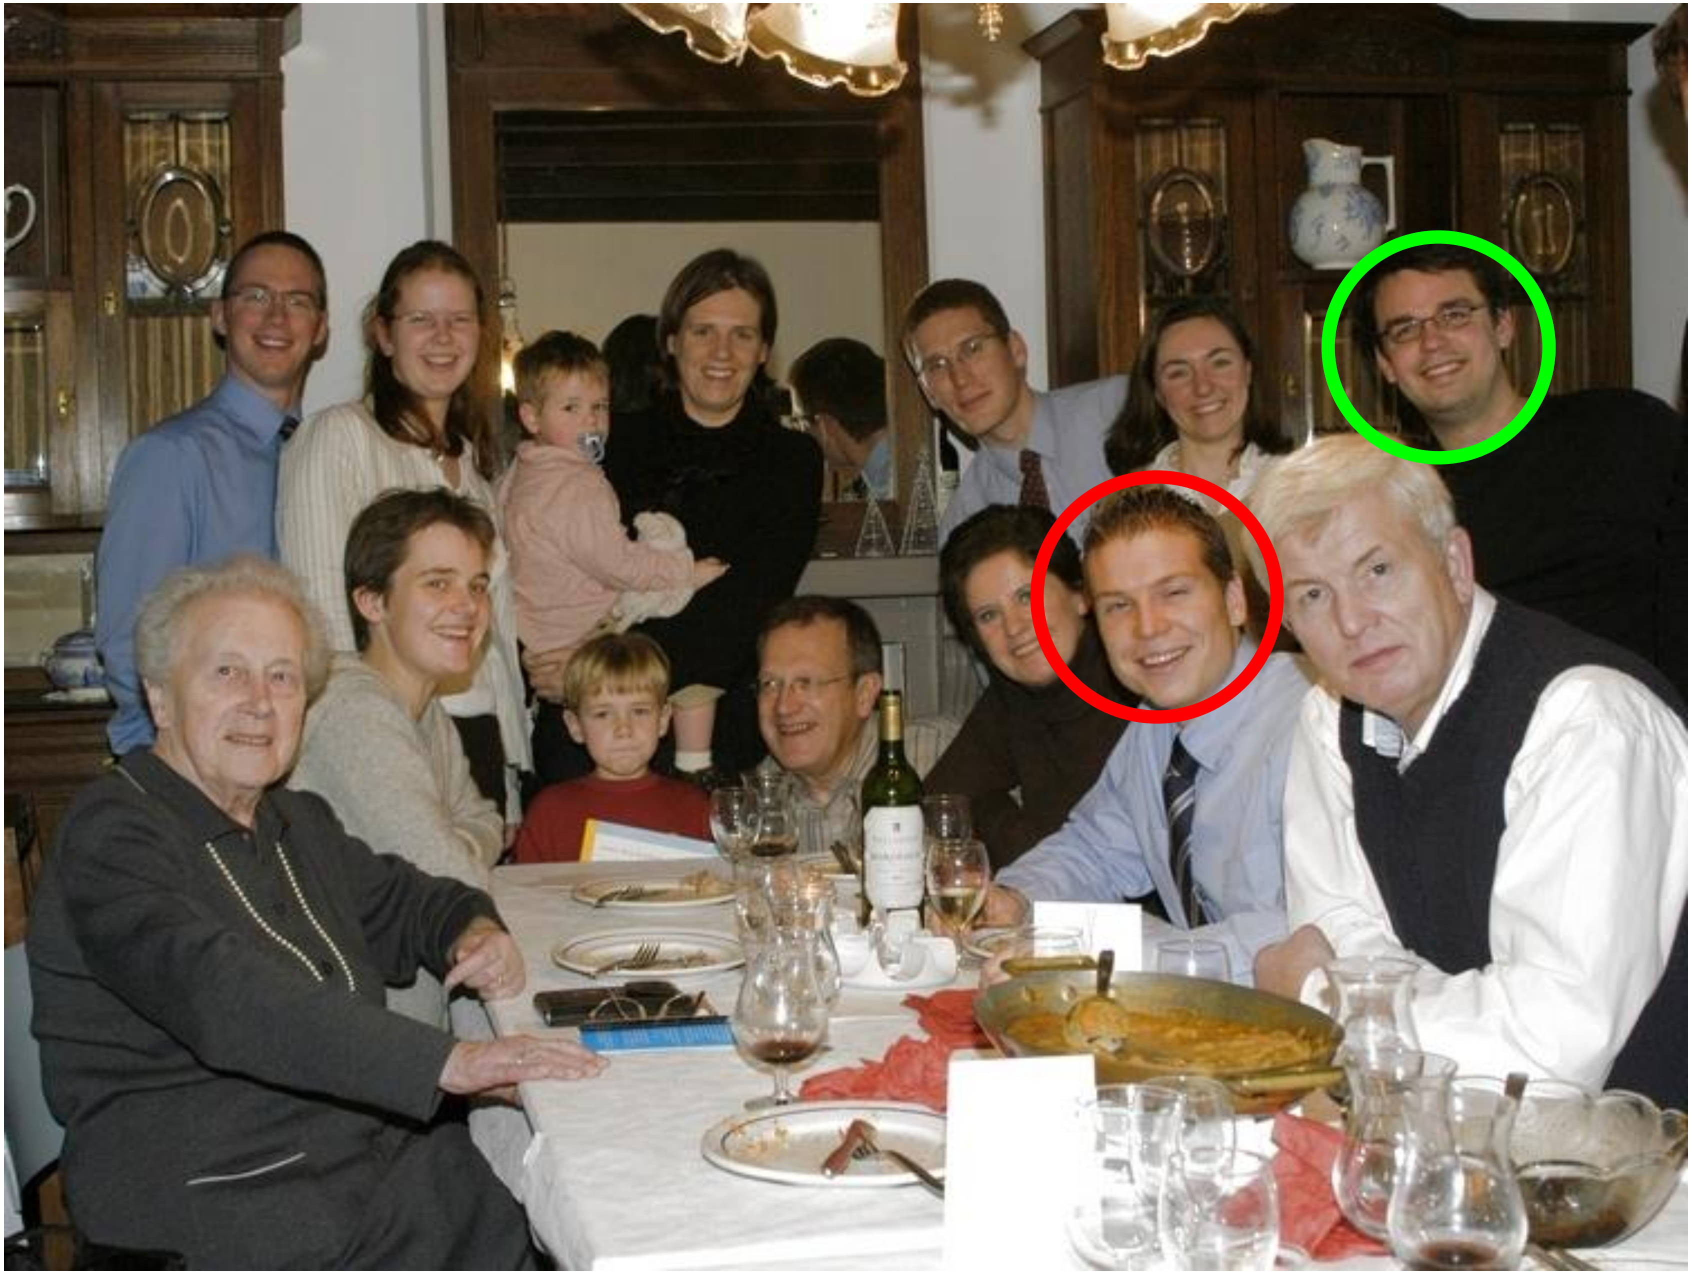

Slide credit: Y.Y. Chuang

Slide credit: Y.Y. Chuang

Slide credit: Y.Y. Chuang

Slide credit: Y.Y. Chuang

Slide credit: Y.Y. Chuang

set of originals

photomontage

Slide credit: Y.Y. Chuang

Source images | Brush strokes | Computed labeling

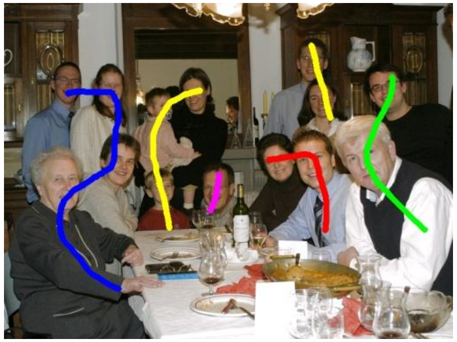

Composite

Slide credit: Y.Y. Chuang

Brush strokes

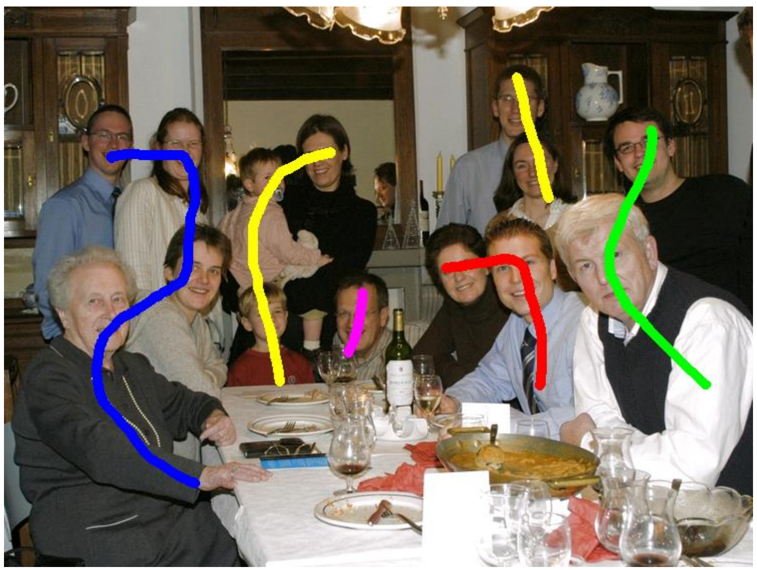

Computed labeling

Slide credit: Y.Y. Chuang

Student paper presentation

AutoCollage

C. Rother, L. Bordeaux, Y. Hamadi, and A. Blake
SIGGRAPH 2006

Presenter: Elmusraty, Gusaue S Abdullatef (Pref: Kai)

Student paper presentation

Rectangling Panoramic Images via Warping

Kaiming He, Huiwen Chang, and Jian Sun
SIGGRAPH 2013

Presenter: Loveless, Blake

Next Time

- Matting
- Student paper presentations
  - 05/10: Mcgowan, Travis
    - Learning to See in the Dark
      C. Chen, Q. Chen, J. Xu and V. Koltun. IEEE CVPR 2018
  - 05/10: McKinney, Drew
    - Video Tapestries with Continuous Temporal Zoom
      C. Barnes, D. Goldman, E. Shechtman, and A. Finkelstein
      SIGGRAPH 2010