More on Stereo

Outline

- Weighted window-based correlation
- Diffusion
- Energy minimization
- Graph cuts

Selecting Window Size

- Small window: more detail, but more noise
- Large window: more robustness, less detail
- Example:

Selecting Window Size

3 pixel window

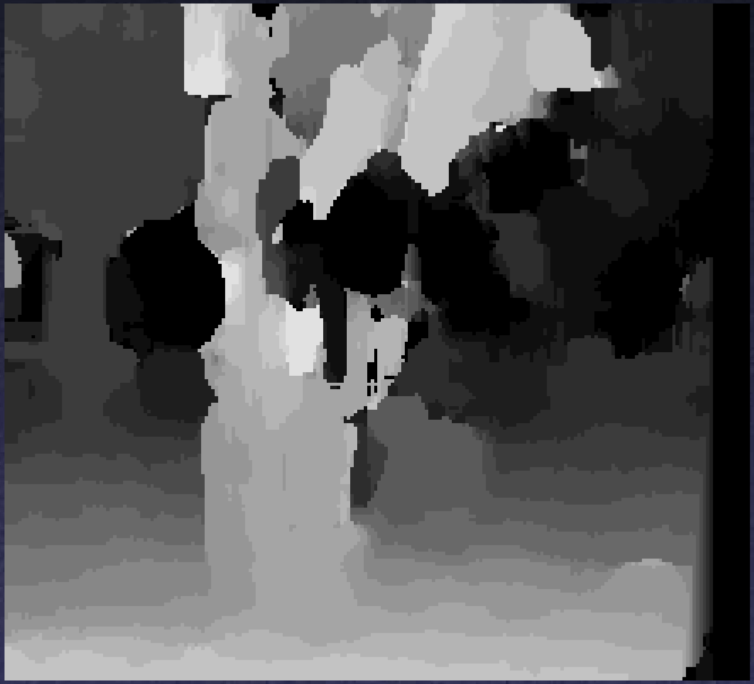

20 pixel window

Window-Based Correlation

- For each pixel
  - For each disparity
    - For each pixel in window
      - Compute difference
  - Find disparity with minimum SSD

Reverse Order of Loops

- For each disparity
  - For each pixel
    - For each pixel in window
      - Compute difference
- Find disparity with minimum SSD at each pixel

Non-Square Windows

- Compromise: have a large window, but higher weight near the center
- Example: Gaussian
- For each disparity
  - For each pixel
    - Compute *weighted* SSD

Diffusion

- For each disparity
  - For each pixel
    - Matching error E_i = (squared) difference in intensities
  - For *n* iterations:
    - For each pixel:
      - $E_i \leftarrow (1-4\lambda)\, E_i + \lambda\, \Sigma E_j$
- Sum is over four neighbors of each pixel

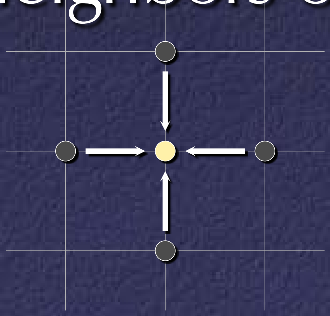

Non-Linear Diffusion

- To prevent blurring even more, only perform diffusion in ambiguous regions
- For each pixel, compute certainty
  - High certainty iff one disparity has low error, all others have high error
- For each pixel, only perform diffusion if certainty goes up

Certainty Metrics for Non-Linear Diffusion

- Winner margin: normalized difference between lowest and second-lowest error

$$C(i,j) = \frac{E_{min2} - E_{min}}{\sum_d E_d}$$

- Entropy:

$$C(i,j) = -\sum_d p(d) \log p(d), \quad p(d) = \frac{e^{-E_d}}{\sum_{d'} e^{-E_{d'}}}$$

Results

- Scharstein and Szeliski, 1996

3 pixel window

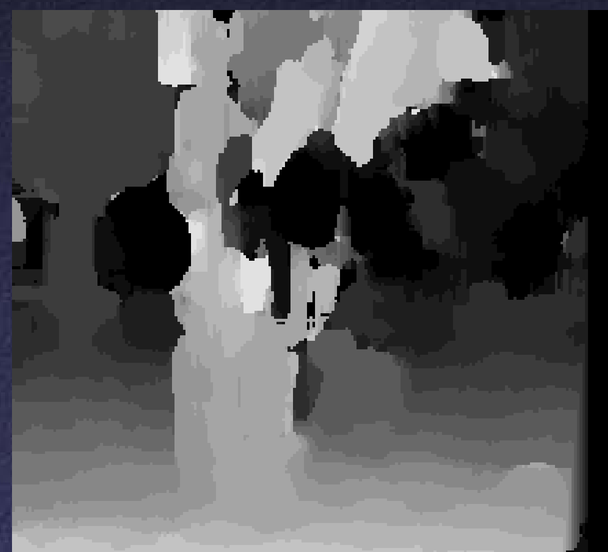

20 pixel window

Nonlinear diffusion

Energy Minimization

- Another approach to improve quality of correspondences
- Assumption: disparities vary (mostly) smoothly
- Minimize energy function:

$$E_{data} + \lambda E_{smoothness}$$

- E_{data}: how well does disparity match data
- $E_{smoothness}$: how well does disparity match that of neighbors – regularization

Energy Minimization

- If data and energy terms are nice (continuous, smooth, etc.) can try to minimize via gradient descent, etc.
- In practice, disparities only piecewise smooth
- Design smoothness function that doesn't penalize large jumps too much
  - Example: $E_{smoothness}(\alpha,\beta) = \min(|\alpha - \beta|, K)$

Energy Minimization

- Hard to find global minima of non-smooth functions
  - Many local minima
  - Provably NP-hard
- Practical algorithms look for approximate minima (e.g., simulated annealing)

Energy Minimization via Graph Cuts

- Boykov, Veksler, and Zabih, 2001
- Define a class of operations
  - e.g., change some of the disparities to α
- Look for operations that reduce energy
- Terminate when no operations of the class being considered reduce energy

Energy Minimization via Graph Cuts

- Different kinds of operations:

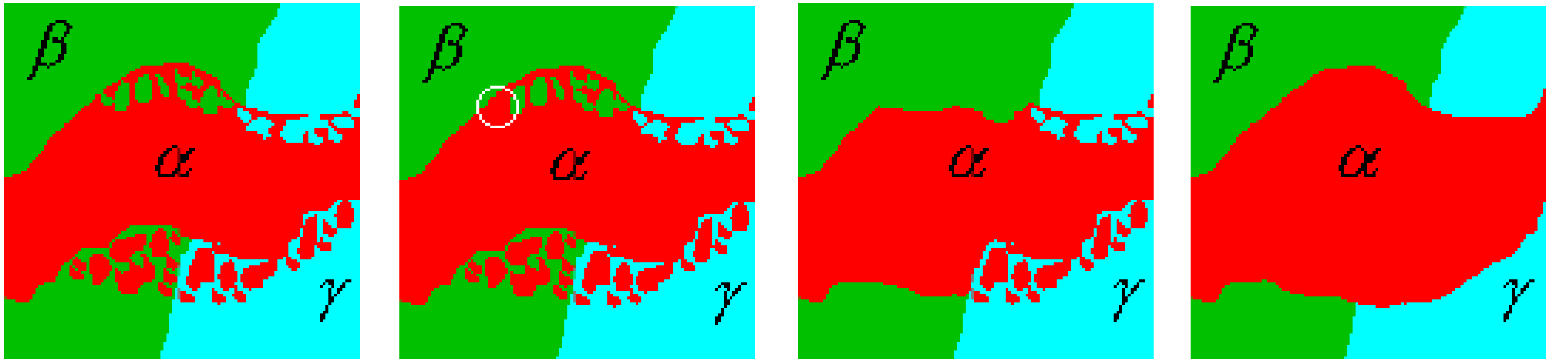

(a) initial labeling (b) standard move (c) α-β-swap (d) α-expansion

- Challenge: how to find operations that reduce energy the most

Energy Minimization via Graph Cuts

- Represent possible operations as cuts through graphs
- Graph cut: minimal subset of edges that separates two (given) nodes of graph
- Fast algorithms for computing minimal-cost cuts

Energy Minimization via Graph Cuts

- α-β swap: interchange α and β labels

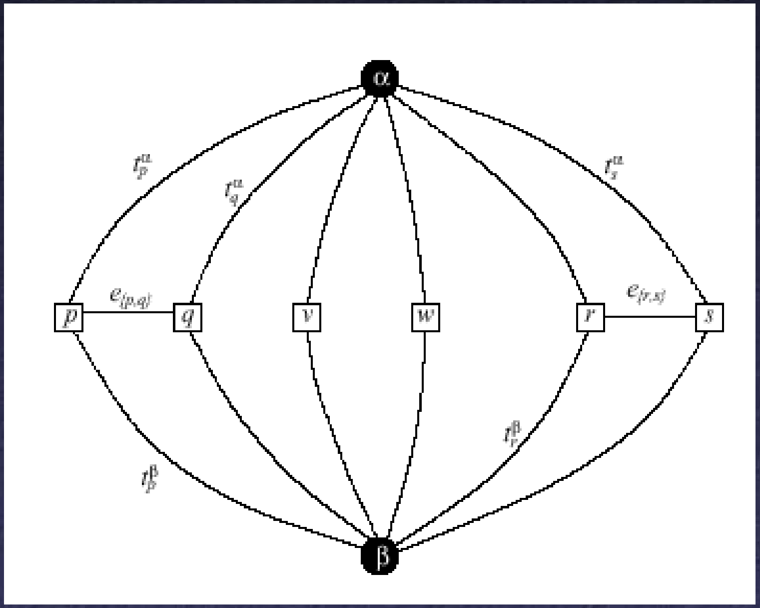

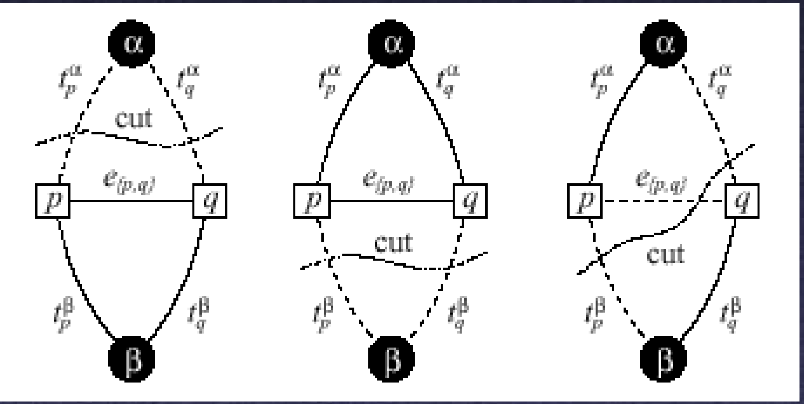

Energy Minimization via Graph Cuts

- α expansion: add pixels to α class

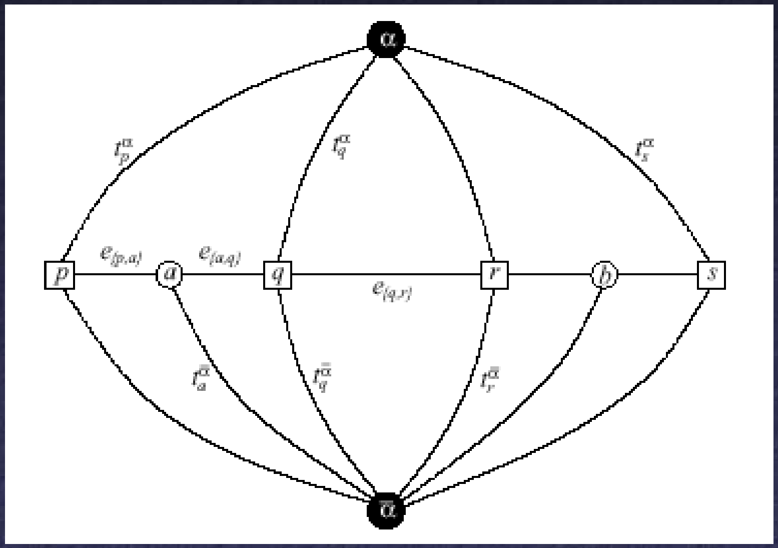

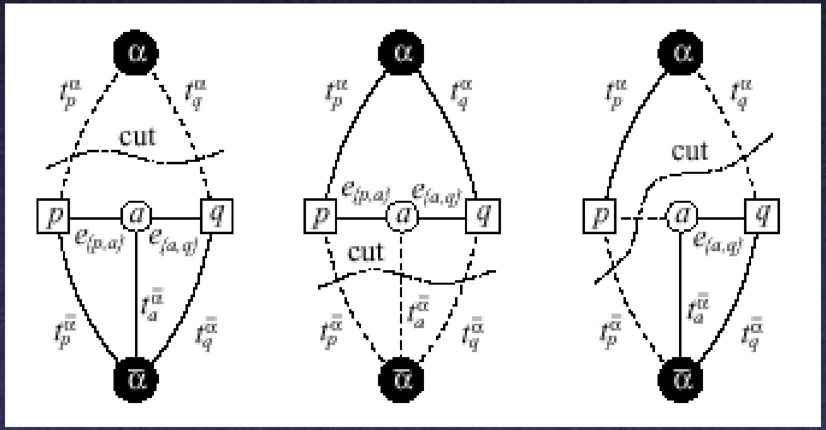

Results

Image

Ground truth

Swap algorithm

Expansion algorithm

Results

Image

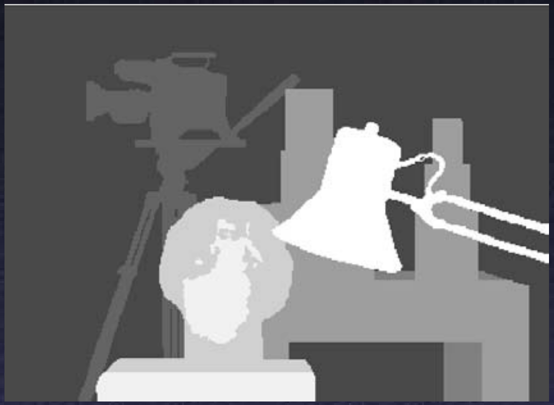

Ground truth

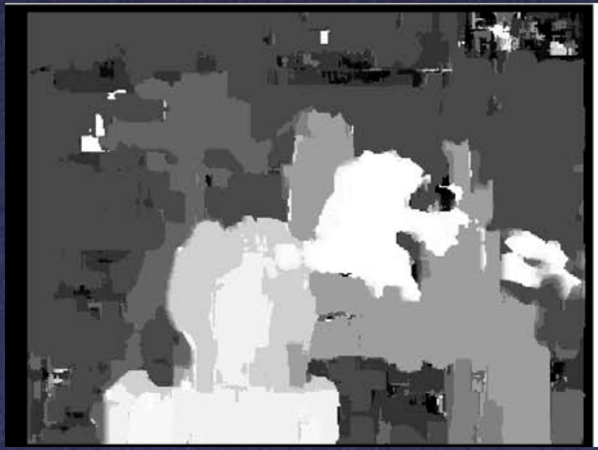

Normalized correlation

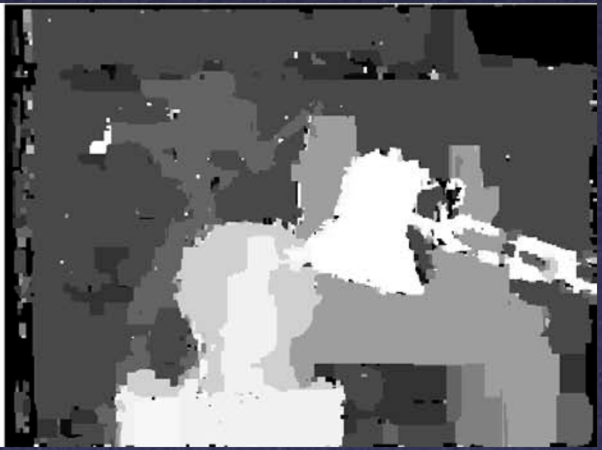

Simulated annealing

Energy Minimization for Image Smoothing

- Apply same principle: image should be close to original image, but piecewise smooth

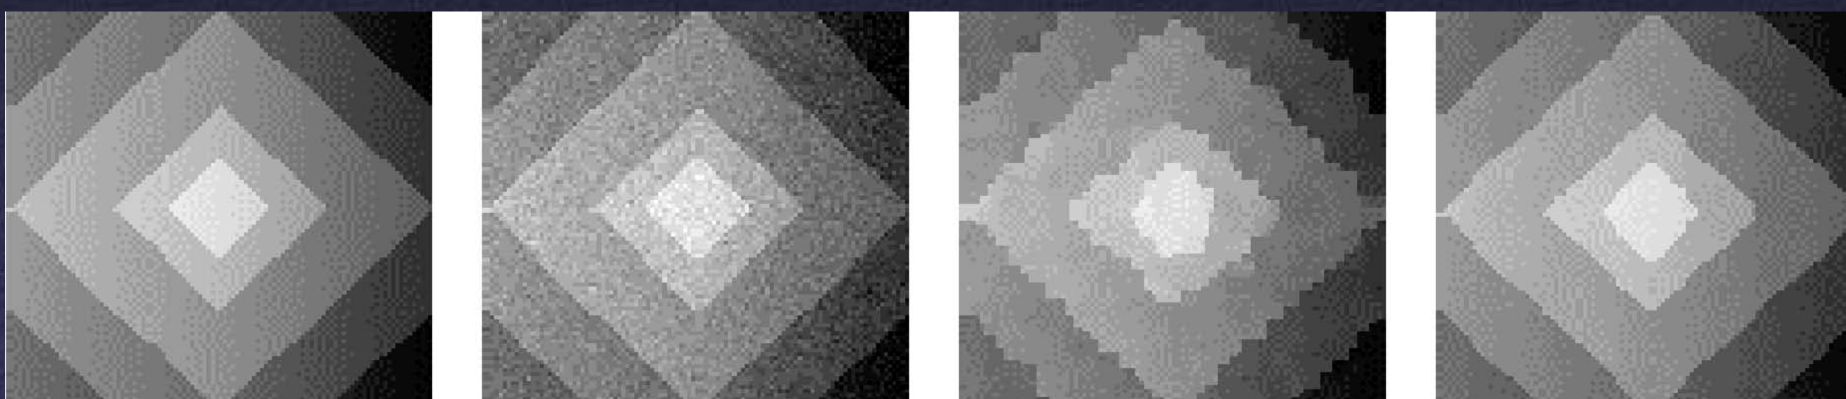

Original image | Noise added | Local energy minimum with one-pixel changes | Local energy minimum with α-expansion