Stereo
(Part 3)

Computer Vision I

CSE 252A

Lecture 14

Announcements

– Read Sections 7.5 and 20.1.2
– Homework 3 is due Dec 4, 11:59 PM

Features on same epipolar line

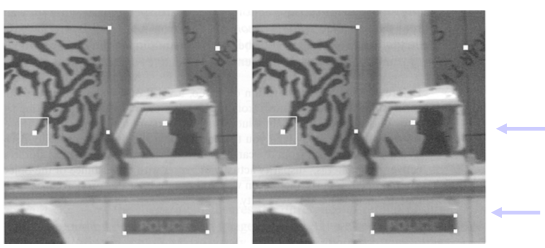

Summary of Stereo Constraints

CONSTRAINT	BRIEF DESCRIPTION
1-D Epipolar Search	Arbitrary images of the same scene may be rectified based on epipolar geometry such that stereo matches lie along one-dimensional scanlines. This reduces the computational complexity and also reduces the likelihood of false matches.
Monotonic Ordering	Points along an epipolar scanline appear in the same order in both stereo images, assuming that all objects in the scene are approximately the same distance from the cameras.
Image Brightness Constancy	Assuming Lambertian surfaces, the brightness of corresponding points in stereo images are the same.
Match Uniqueness	For every point in one stereo image, there is at most one corresponding point in the other image.
Disparity Continuity	Disparities vary smoothly (i.e. disparity gradient is small) over most of the image. This assumption is violated at object boundaries.
Disparity Limit	The search space may be reduced significantly by limiting the disparity range, reducing both computational complexity and the likelihood of false matches.
Fronto-Parallel Surfaces	The implicit assumption made by area-based matching is that objects have fronto-parallel surfaces (i.e. depth is constant within the region of local support). This assumption is violated by sloping and creased surfaces.
Feature Similarity	Corresponding features must be similar (e.g. edges must have roughly the same length and orientation).
Structural Grouping	Corresponding feature groupings and their connectivity must be consistent.

(From G. Hager)

Stereo matching

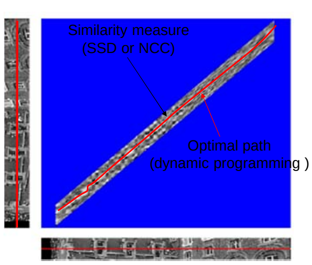

Constraints
• epipolar
• ordering
• uniqueness
• disparity limit
• disparity gradient limit

Trade-off
• Matching cost (data)
• Discontinuities (prior)

(From Pollefeys)

(Cox et al. CVGIP'96; Koch'96; Falkenhagen´97;
Van Meerbergen,Vergauwen,Pollefeys,VanGool IJCV'02)

Stereo Matching with Dynamic Programming

(Slides adapted from Jim Rehg at GA Tech)

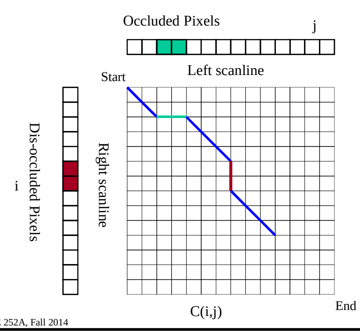

Dynamic programming yields the optimal path through grid. This is the best set of matches that satisfy the ordering constraint

Every pixel on each scanline will be labeled as matching, or occluded.

Dynamic Programming

- Efficient algorithm for solving sequential decision (optimal path) problems. Cost associated with each arc.

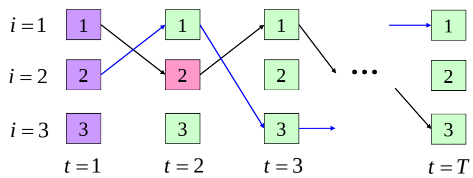

How many paths through this trellis? 3^T

Using Dynamic Programming, can find optimal path in O(M T) time (here M=3)

Dynamic Programming for Stereo

- Efficient algorithm for solving sequential decision (optimal path) problems.

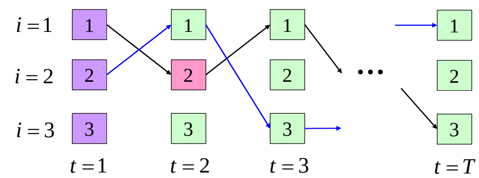

For Stereo,

t can denote pixel coordinates across an epipolar line in one image
i can denote the disparity to the other epipolar line

Dynamic Programming

Used with Hidden Markov Models, Viterbi Algorithm

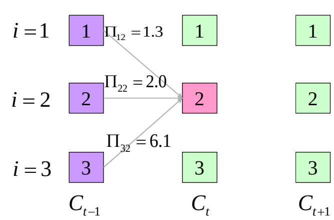

Suppose cost can be decomposed into stages:

$$\Pi_{ij} = \text{Cost of going from state } i \text{ to state } j$$

Dynamic Programming
Minimum Cost Path

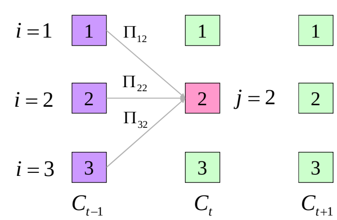

What is minimum cost of reaching node j at time t?

Dynamic Programming

Used with Hidden Markov Models, Viterbi Algorithm

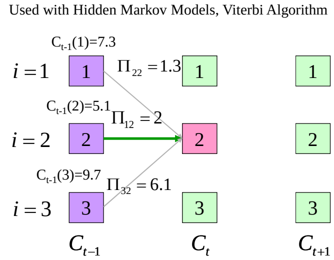

What is minimum cost of reaching node j at time t?

$$C_t(j) = \min_i \left(\Pi_{ij} + C_{t-1}(i) \right)$$

Minimum cost of path from t=0 to reach state j at time t.

Dynamic Programming

Used with Hidden Markov Models, Viterbi Algorithm

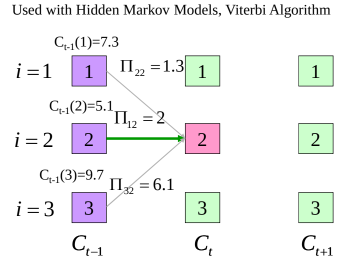

$$C_t(j) = \min_i \left(\Pi_{ij} + C_{t-1}(i) \right)$$

$$b_t(j) = \arg\min_i \left(\Pi_{ij} + C_{t-1}(i) \right)$$

So, $b_t(2) = 2$

$b_t(j)$ gives previous state along minimum cost path

Dynamic Programming

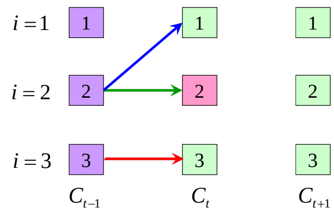

$C_t(j) = \min_i\left(\Pi_{ij} + C_{t-1}(i)\right)$

$b_t(j) = \arg\min_i\left(\Pi_{ij} + C_{t-1}(i)\right)$

So, $b_t(3) = 3$
$b_t(2) = 2$
$b_t(1) = 2$

$b_t(j)$ gives previous state along minimum cost path

Dynamic Programming

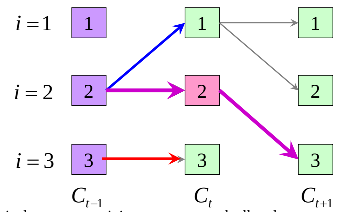

1. Iteratively, compute minimum cost to reach all nodes

$$C_t(j) = \min_i \Pi_{ij} + C_{t-1}(i)$$

2. Recursively, starting with the node at time t-max, select lowest cost terminal node, and backtrack along path

$$b_t(j) = \arg\min \Pi_{ij} + C_{t-1}(i)$$

Compute Optimal Path Costs

``A Maximum Likelihood Stereo Algorithm'', Cox, Hingorani, Rao, Maggs, Computer Vision & Image Understanding, 63, 3, pp. 542-567.

$$\text{Occlusion} = \left[\ln\left(\frac{P_D}{1-P_D}\frac{\phi}{|(2\pi)^d S_s^{-1}|^{\frac{1}{2}}}\right)\right]$$

```
for (i=1;i≤ N;i++){ C(i,0) = i*Occlusion }
for (i=1;i≤ M;i++) { C(0,i) = i*Occlusion}
for(i=1;i≤ N;i++){
   for(j=1;j≤ M;j++){
      min1 = C(i-1,j-1)+c(z_{1,i},z_{2,j});
      min2 = C(i-1,j)+Occlusion;
      min3 = C(i,j-1)+Occlusion;
      C(i,j) = cmin = min(min1,min2,min3);
      if(min1==cmin) M(i,j) = 1;
      if(min2==cmin) M(i,j) = 2;
      if(min3==cmin) M(i,j) = 3;
   }}
```

C(i,j): Cost of optimal path to match of pixels i and j
M(i,j): 'Pointer' to previous node along optimal path

Back tracking to get optimal path

```
p=N;
q=M;
while(p!=0 && q!=0){
   switch(M(p,q)){
      case 1:
         p matches q
         p--;q--;
         break;
      case 2:
         p is unmatched
         p--;
         break;
      case 3:
         q is unmatched
         q--;
         break;
   }}
```


Stereo Matching with Dynamic Programming

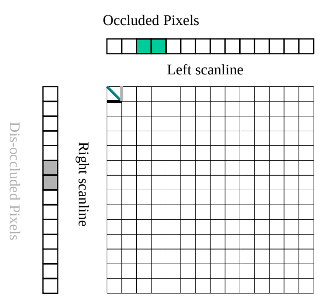

C(i,j) is minimum of
1. C(i-1,j-1) + match-cost of pixel L(i) & R(i)
2. C(i-1,j) +occlusion-penalty
3. C(i,j-1)+occlusion-penalty

Terminal

Stereo Matching with Dynamic Programming

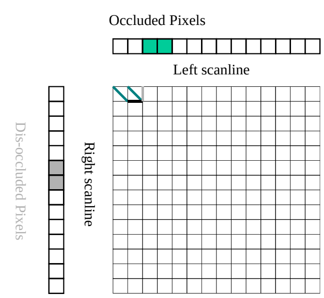

Scan across grid computing optimal cost for each node given its upper-left neighbors.

Terminal

Stereo Matching with Dynamic Programming

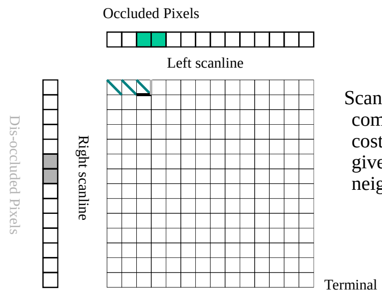

Scan across grid computing optimal cost for each node given its upper-left neighbors.

Stereo Matching with Dynamic Programming

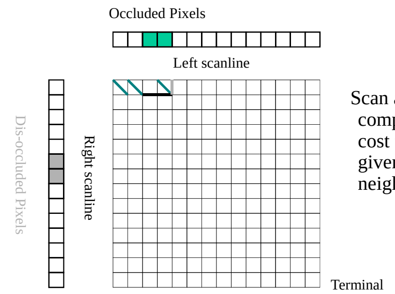

Scan across grid computing optimal cost for each node given its upper-left neighbors.

Stereo Matching with Dynamic Programming

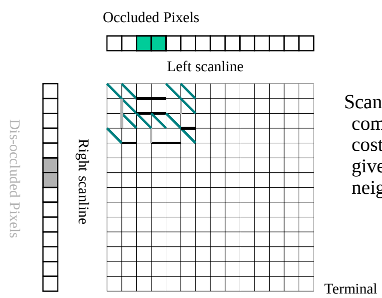

Scan across grid computing optimal cost for each node given its upper-left neighbors.

Stereo Matching with Dynamic Programming

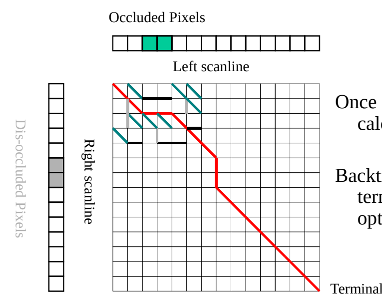

Once C(i,j) is completely calculated:

Backtrack from the terminal to get the optimal path.

Some Challenges & Problems

- Photometric issues:
  - specularities
  - strongly non-Lambertian BRDF’s
- Surface structure
  - lack of texture
  - repeating texture within horopter bracket
- Geometric ambiguities
  - as surfaces turn away, difficult to get accurate reconstruction (affine approximate can help)
  - at the occluding contour, likelihood of good match but incorrect reconstruction

Many variations

- Use features like edges, rather than intensities (Baker, Binford, 1981). See text.
- Subpixel interpolation
- Probabilistic framework
- Creases
- Occlusion penalities

Variations on Binocular Stereo

1. Trinocular Stereopsis
2. Helmholtz Reciprocity Stereopsis

Trinocular Epipolar Constraints

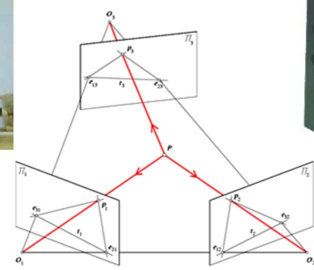

$$\begin{cases} p_1^T \mathcal{E}_{12} p_2 = 0 \\ p_2^T \mathcal{E}_{23} p_3 = 0 \\ p_3^T \mathcal{E}_{31} p_1 = 0 \end{cases}$$

These constraints are not independent!

$$e_{31}^T \mathcal{E}_{12} e_{32} = e_{12}^T \mathcal{E}_{23} e_{13} = e_{23}^T \mathcal{E}_{31} e_{21} = 0$$

Helmholtz reciprocity

$$\rho(\theta_{in}, \phi_{in} ; \theta_{out}, \phi_{out}) = \rho(\theta_{out}, \phi_{out} ; \theta_{in}, \phi_{in})$$

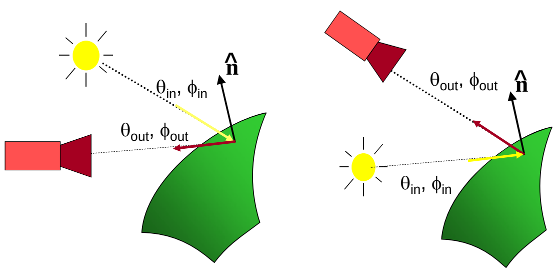

[Helmholtz, 1910], [Minnaert, 1941], [Nicodemus et al, 1977]

Point Source Illumination

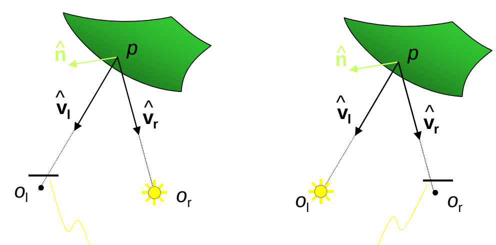

$$i_r = \rho(\hat{v}_l, \hat{v}_r) \frac{\hat{n} \cdot \hat{v}_l}{|o_l - p|^2}$$

=

Matching Constraint

$$\boxed{\left(i_l \underbrace{\frac{\hat{v}_l}{|o_l - p|^2}}_{\vec{w}_l} - i_r \underbrace{\frac{\hat{v}_r}{|o_r - p|^2}}_{\vec{w}_r} \right) \cdot \hat{n} = 0}$$

measured

computed from geometric calibration

Using Multiple Helmholtz Stereo Pairs

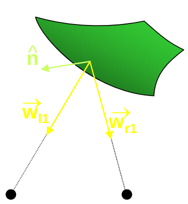

$$\begin{pmatrix} i_{l1}\vec{w}_{l1}^T - i_{r1}\vec{w}_{r1}^T \\ i_{l2}\vec{w}_{l2}^T - i_{r2}\vec{w}_{r2}^T \\ \vdots \end{pmatrix} \hat{n} = 0$$

- **Multiple views (at least three pairs) yield a matrix constraint equation.**
- **Matrix must be Rank 2.**
- **Search for depth where rank constraint is satisfied.**

Finding the Normal at each point

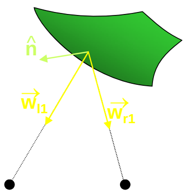

$$\begin{bmatrix} i_{l1}\vec{w}_{l1}^{T} - i_{r1}\vec{w}_{r1}^{T} \\ i_{l2}\vec{w}_{l2}^{T} - i_{r2}\vec{w}_{r2}^{T} \\ \vdots \end{bmatrix} \hat{n} = 0$$

- **Additionally, the surface normal $\hat{n}$ must lie in the kernel of the matrix**

Disparity and Normal Field

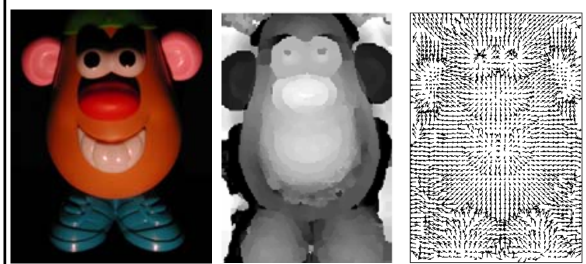

Experimental Setup

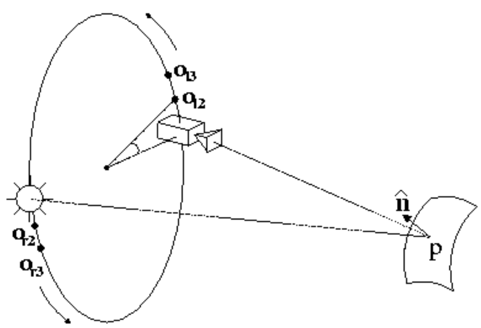

Experimental Aparatus

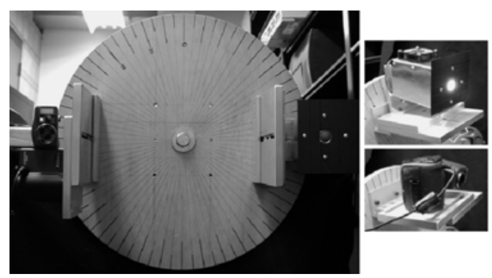

Helmholtz Stereopsis

Second Generation Rig

Bulldog: Disparity

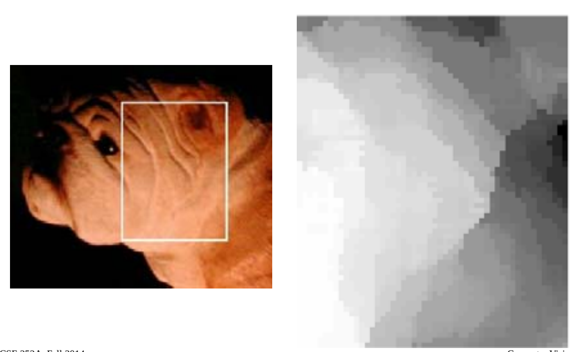

Bulldog: Normal Field

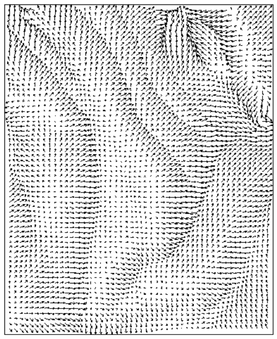

Plastic Baby Doll: Disparities

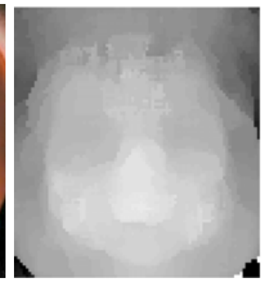

Plastic Baby Doll: Normal Field

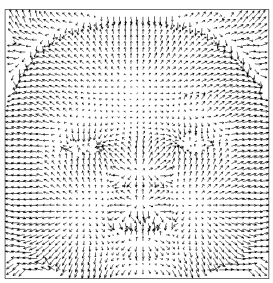

Surface after integrating normal field

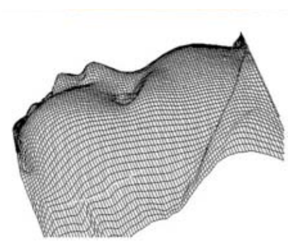

Reciprocal Images: Typical Dataset

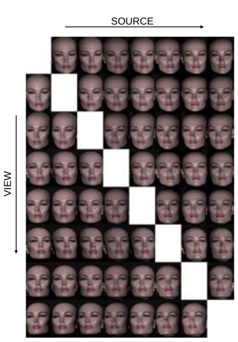

Reciprocal Images: Typical Dataset

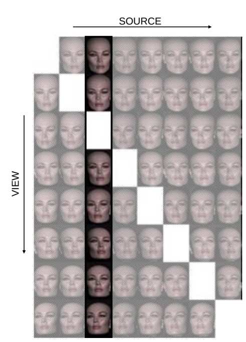

Conventional Stereo

- Constant brightness
- No structure in textureless regions

Reciprocal Images: Typical Dataset

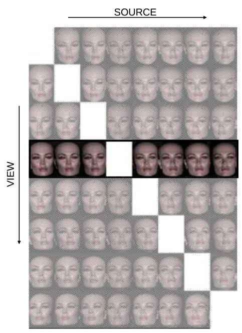

Conventional Stereo
- Constant brightness
- No structure in textureless regions

Photometric Stereo
- Needs reflectance model
- No direct depth estimates

Reciprocal Images: Typical Dataset

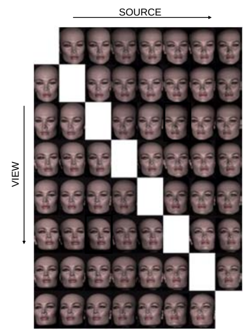

Conventional Stereo
- Constant brightness
- No structure in textureless regions

Photometric Stereo
- Needs reflectance model
- No direct depth estimates

Helmholtz Stereo
- No assumed reflectance
- Gives depth and surface normals

Metric Reconstruction

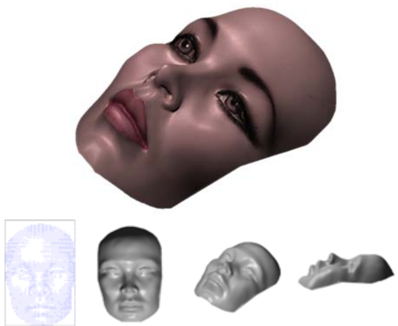

Comparison to other methods

Method	Assumed Reflectance	Surface Information Recovered	Recovers Constant Intensity	Active/ Passive	Depth Discontinuities	Handles Half-Occlusion	Robust to Cast Shadows
Photometric Stereopsis	Lambertian or Known	Surface Normals	Surface Normals	Active	No	NA	No
Multinocular Stereopsis	Lambertian	Depth	Nothing	Passive	Sometimes	Sometimes	Yes
Helmholtz Stereopsis	Arbitrary	Depth + Surface Normals	Surface Normals	Active	Yes	Yes	Yes

More on stereo …

The Middlebury Stereo Vision Research Page
http://cat.middlebury.edu/stereo/

Recommended reading

D. Scharstein and R. Szeliski.
A Taxonomy and Evaluation of Dense Two-Frame Stereo Correspondence Algorithms. IJCV 47(1/2/3):7-42, April-June 2002. PDF file (1.15 MB) - includes current evaluation. Microsoft Research Technical Report MSR-TR-2001-81, November 2001.

Myron Z. Brown, Darius Burschka, and Gregory D. Hager. Advances in Computational Stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8):993-1008, 2003.