
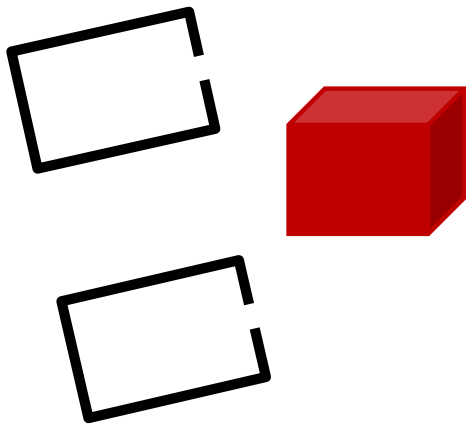


CSCI 497P/597P: Computer Vision

Scott Wehrwein

Stereo: Metrics, Rectification Planesweep Stereo



Announcements

- Reminder: Exam
 - out this morning
 - due Tuesday night
- P02
 - code due tonight
 - artifact due Tuesday night

Goals

- Understand the basic metrics used to compare patches (SSD, SAD, NCC)
- Understand how to rectify a pair of stereo images given their intrinsics and extrinsics.
- Understand the plane sweep stereo algorithm.

A Stereo Algorithm

1. For every pixel (x, y)

1. For every disparity d

1. Get patch from image 1 at (x, y)
2. Get patch from image 2 at $(x + d, y)$
3. Compute cost using your metric of choice

```
C = np.array(h,w,D)
for r in range(0,h):
    for c in range(0,w):
        for d in range(min_d, max_d):
            C[r,c,d] = metric(get_patch(im1,r,c), get_patch(im2,r,c+d))

disp = np.argmin(C, axis=2)
depth = f * b / disp
```

Metrics for Stereo Matching

$$\text{img}_1 \text{ patch} = W_1, \quad \text{img}_2 \text{ patch} = W_2$$

- SSD = sum of squared differences

$$\text{np.sum}((W_2 - W_1)**2)$$

- SAD = sum of absolute differences

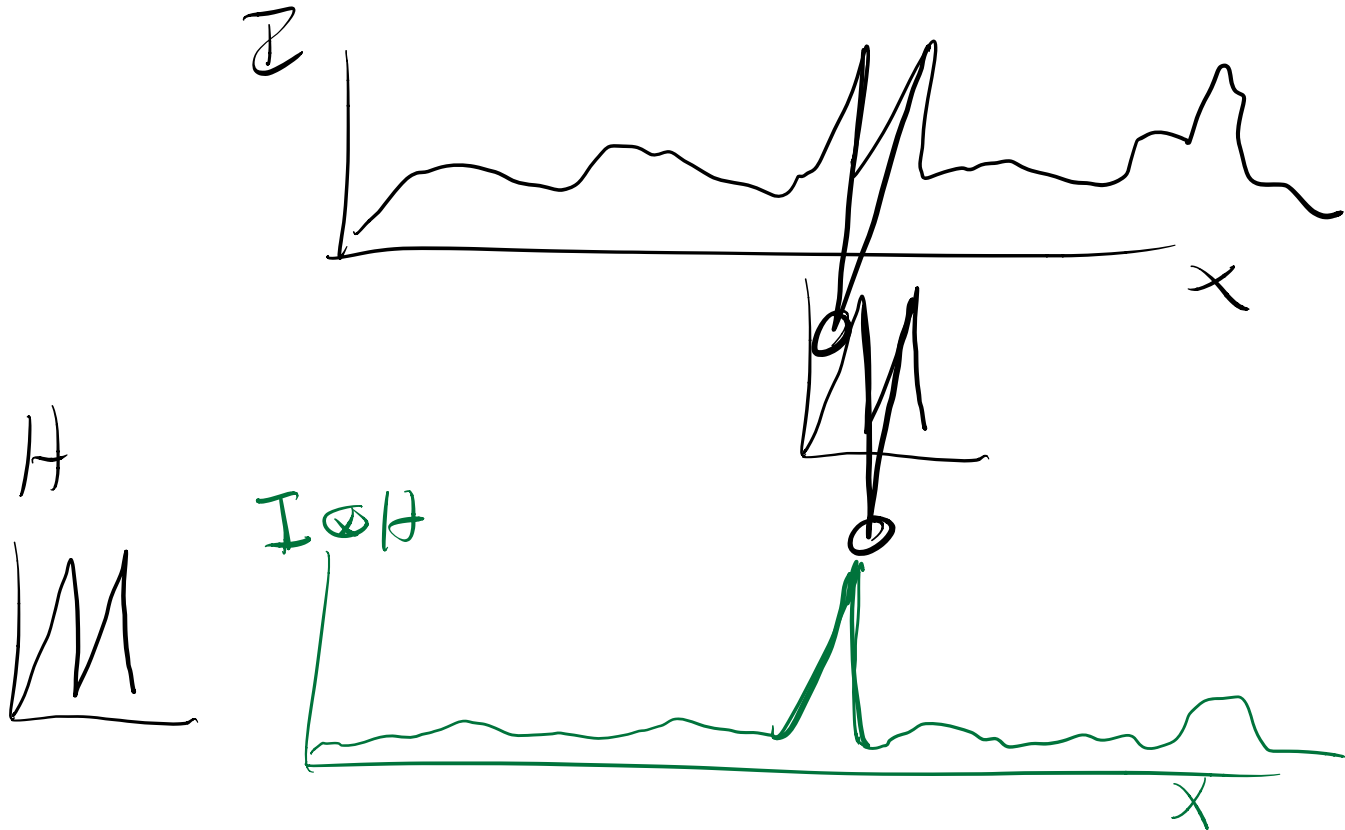
$$\text{np.sum}(\text{np.abs}(W_2 - W_1))$$

- NCC = normalized cross-correlation

– (more convolution cross correlation!)

Un-Normalized Cross Correlation


Insight: a cross-correlation filter is good at finding patches that **look like itself**.



Normalized Cross Correlation

Approach: apply a **patch** from one image as a **filter** across the other.

Trick: normalize patches before computing product to add invariance.



regions A, B, write as vectors \mathbf{a} , \mathbf{b}

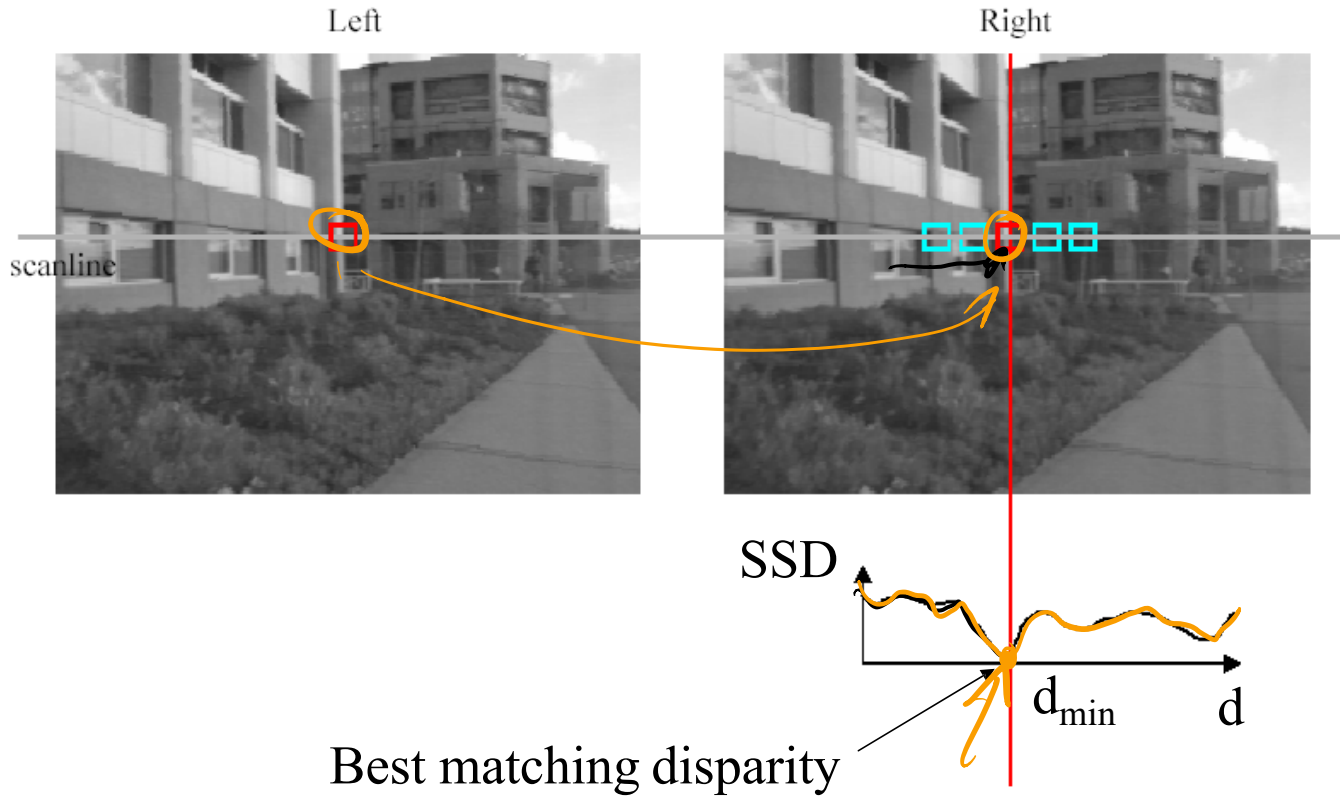
subtract the mean of each vector:


$$\mathbf{a} \rightarrow \mathbf{a} - \langle \mathbf{a} \rangle, \quad \mathbf{b} \rightarrow \mathbf{b} - \langle \mathbf{b} \rangle$$

$$\text{cross correlation} = \frac{\mathbf{a} \cdot \mathbf{b}}{|\mathbf{a}| |\mathbf{b}|}$$

Invariant to $I \rightarrow \alpha I + \beta$

Stereo matching based on SSD



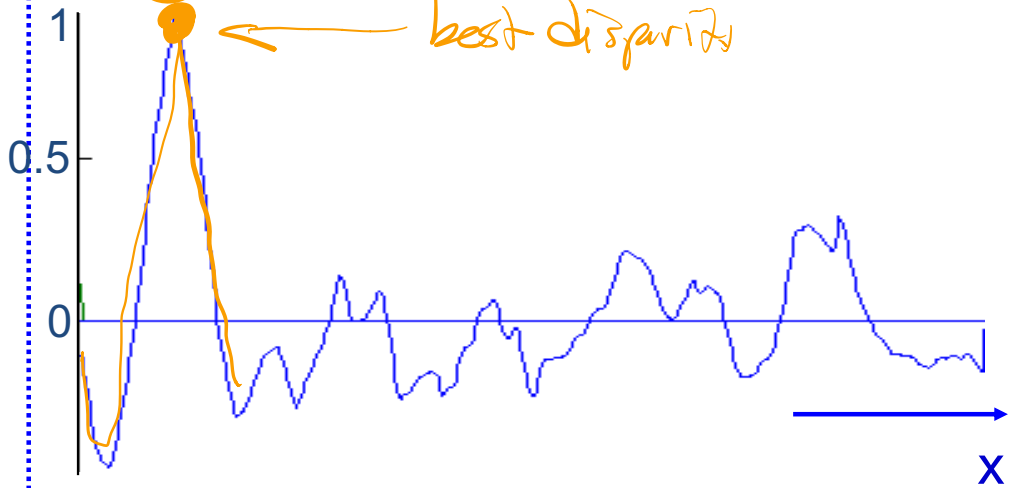




Stereo with NCC: The Good Case



left image band

right image band





Stereo with NCC: The Bad Case

target region



left image band



right image band



cross
correlation

Stereo results

- Data from University of Tsukuba
- Similar results on other images without ground truth




Scene



Ground truth

Results with window search



Window-based matching
(best window size)



Ground truth

Better methods exist...



Fancier method



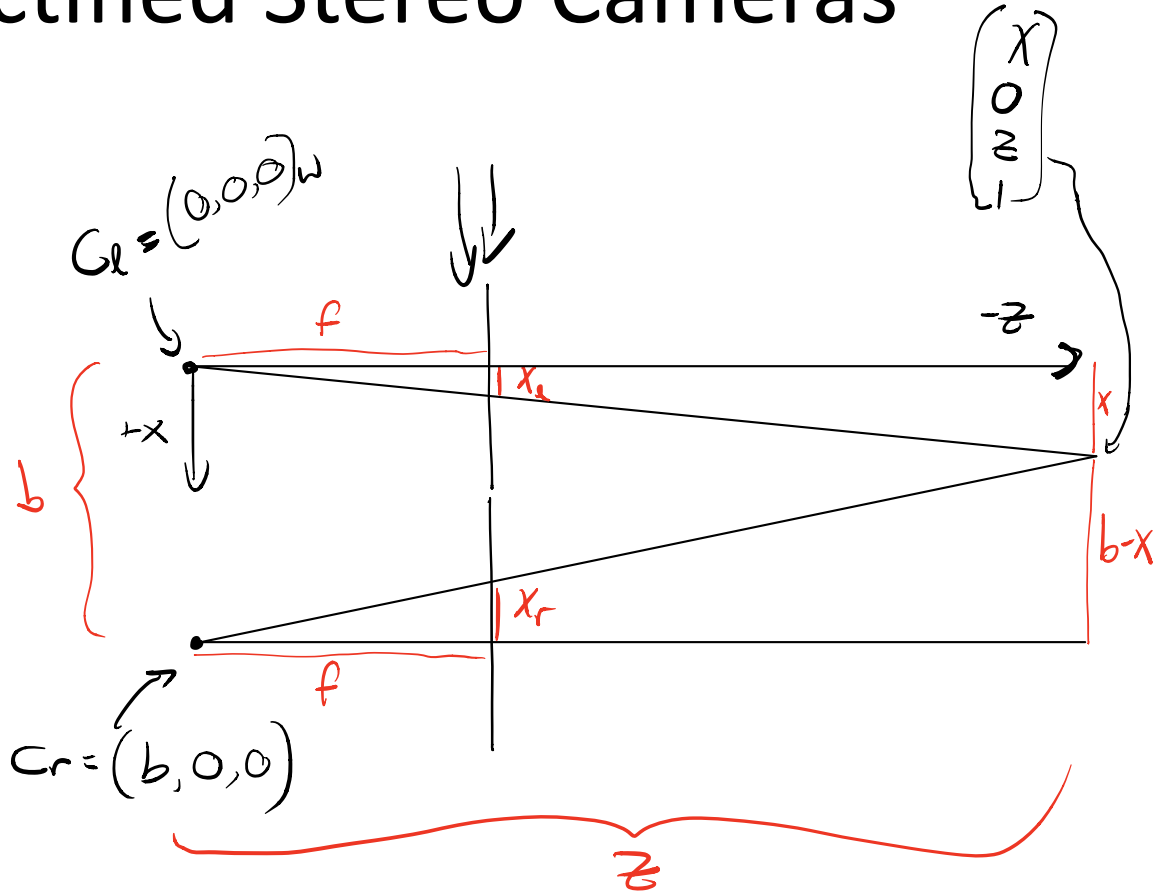
Ground truth

Boykov et al., [Fast Approximate Energy Minimization via Graph Cuts](#),
International Conference on Computer Vision, September 1999.

For the latest and greatest: <http://www.middlebury.edu/stereo/>

Rectified Stereo Cameras

- 2 cameras
- same f
- Same PP
- COP off by b
 b in x
↑ known



What if the cameras **aren't** rectified?


If cameras are **calibrated**, i.e., we know:

$$\begin{array}{cc} K_l & K_r \\ R_l & R_r \\ (\text{also } c_l) & t_r \end{array}$$


...then we can **rectify** the stereo pair

...or we can use **plane sweep stereo**


Rectifying cameras




Rectifying cameras




Rectifying cameras



Rectifying cameras



Rectifying cameras




Example: A rectified stereo pair



GOTO notes

Plane Sweep Stereo



A Stereo Algorithm

1. For every pixel (x, y)

1. For every disparity d

1. Get patch from image 1 at (x, y)
2. Get patch from image 2 at $(x + d, y)$
3. Compute cost using your metric of choice

```
C = np.array(h,w,d)
for r in range(0,h):
    for c in range(0,w):
        for d in range(-maxd, maxd):
            C[r,c,d] = metric(get_patch(im1,r,c), get_patch(im2,r,c+d))

disp = np.max(C, axis=2)
depth = f * b / disp
```


Plane Sweep Stereo Algorithm


1. For every disparity d


1. For every pixel (x, y)

1. Get patch from image 1 at (x, y)
2. Get patch from image 2 at $(x + d, y)$
3. Compute cost using your metric of choice


```
C = np.array(h,w,d)
for d in range(-maxd, maxd):
    for r in range(0,h):
        for c in range(0,w):
            C[r,c,d] = metric(get_patch(im1,r,c), get_patch(im2,r,c+d))

disp = np.max(C, axis=2)
depth = f * b / disp
```







Plane Sweep Stereo





- each plane defines an image \Rightarrow composite homography

Active stereo with structured light





Li Zhang's one-shot stereo



- Project “structured” light patterns onto the object
 - simplifies the correspondence problem
 - basis for active depth sensors, such as Kinect and iPhone X (using IR)

Other methods for getting depth

Laser scanning



Digital Michelangelo Project
<http://graphics.stanford.edu/projects/mich/>


- Optical triangulation
 - Project a single stripe of laser light
 - Scan it across the surface of the object
 - This is a very precise version of structured light scanning

Laser scanned models




The Digital Michelangelo Project, Levoy et al.

Laser scanned models




The Digital Michelangelo Project, Levoy et al.

Laser scanned models




The Digital Michelangelo Project, Levoy et al.

Laser scanned models



The Digital Michelangelo Project, Levoy et al.

Microsoft Kinect



Active stereo with structured light



<https://ios.gadgethacks.com/news/watch-iphone-xs-30k-ir-dots-scan-your-face-0180944/>