



# CSCI 497P/597P: Computer Vision

Scott Wehrwein

## Stereo Rectification Planesweep Stereo



# Announcements

- Reminder: Exam out Friday, due Monday night.

# P1 Artifacts: Hybrid

- 5-way tie (2 votes each)

# Garrett Claeys




# Jake Friberg



# Brady Geleynse



# Sean McCulloch



# Melissa Swift





# P1 Artifacts: Laplacian Forest Sweeney



# Scott's low-effort hybrid artifact



# Goals

- 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(-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
```

# Rectified Stereo Cameras

## Assumptions

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



# What if the cameras aren't rectified?

- Assume cameras are **calibrated**, i.e., we know:

$$K_l, K_r, R_l, R_r, t_l, t_r$$
$$\begin{matrix} \uparrow & \uparrow \\ C_l & C_r \end{matrix}$$

# Projective Geometry: Homogeneous Points

Which of the following 3-vectors does not represent the same projective point as the others?

A.  $[12, 8, 4]$   $(3 \ 2 \ 1)$

*Socratic.com*


→ B.  $[8, 6, 3]$   $(\frac{8}{3}, 2, 1)$

*CSC 497P*

C.  $[24, 16, 8]$   $(3 \ 2 \ 1)$


→ D.  $[15, 10, 5]$   $(3 \ 2 \ 1)$

# Geometric Interpretation



Homogeneous coords live in  $\mathbb{P}^2$   
2D Projective Space






$$P_i' = \begin{pmatrix} K & R & K^{-1} \end{pmatrix} P_i$$

$3 \times 3$     $3 \times 3$     $3 \times 3$   
 ↓   ↓   ↓  
 ↑   ↑   ↑  
 project back   rotate   unproject into camera

H


# What if the cameras aren't rectified?

- Assume cameras are **calibrated**, i.e., we know:







# Rectifying cameras




# Rectifying cameras



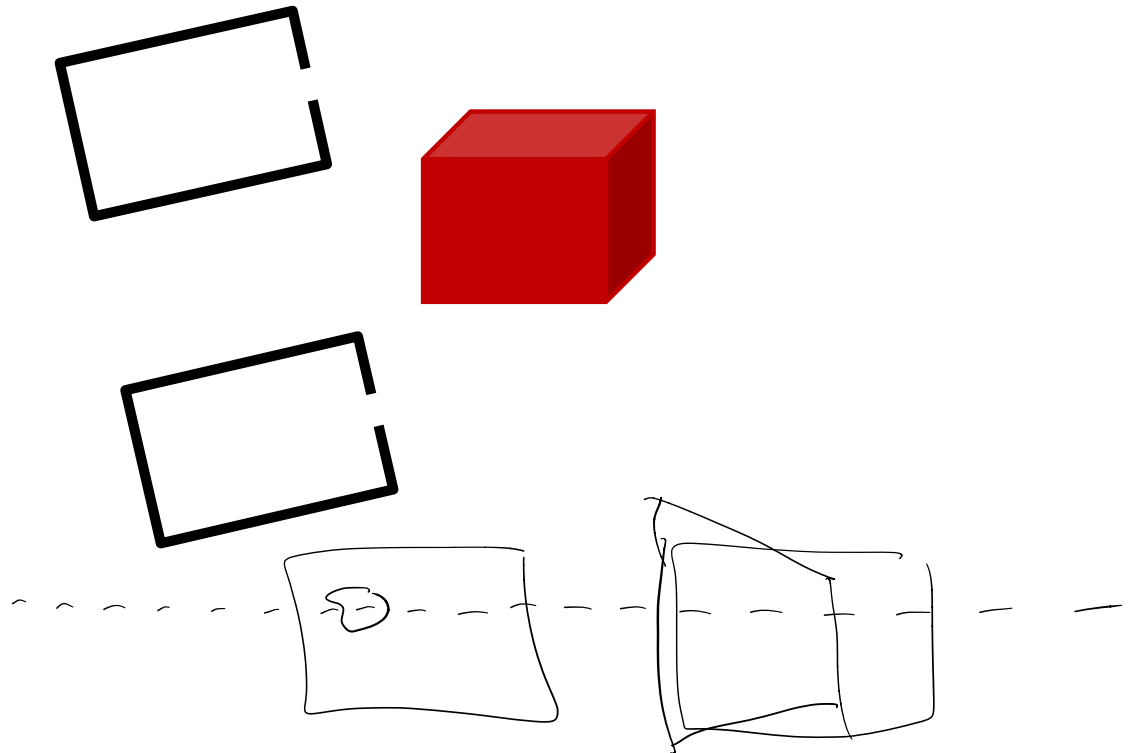
# Rectifying cameras



# Rectifying cameras



# Rectifying cameras





# 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
```

Rectified

$C_l \cdot$

$C_r \cdot$

Not Rectified

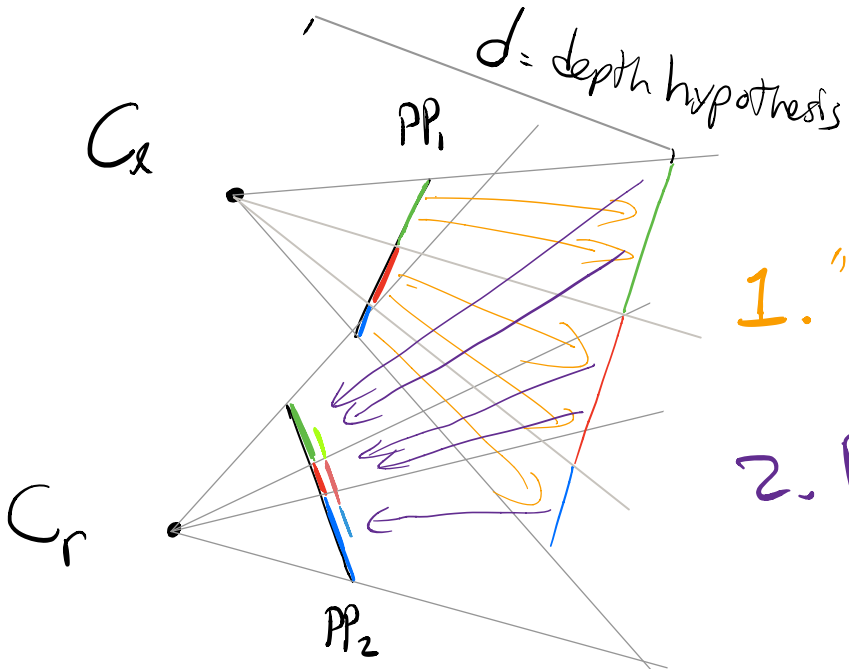


1. "Unproject":

$$\begin{bmatrix} X_w \\ Y_w \\ Z_w \\ 1 \end{bmatrix}$$

$$= \begin{matrix} \text{World} \\ \downarrow \\ -t_e \end{matrix} R_e^T \begin{matrix} \text{cam} \\ @ \\ d \\ \downarrow \\ \text{cam} \end{matrix} K_e^{-1} \begin{matrix} \text{img} \\ \downarrow \\ \begin{bmatrix} X_e \\ Y_e \\ 1 \end{bmatrix} \end{matrix}$$

← add



1. "Unproject"

2. Reproject

## 2. "Reproject":

$$\begin{bmatrix} x_r \\ y_r \\ 1 \end{bmatrix} = K_r \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix}$$

↑  
normalized

↑  
img<sub>r</sub>  
(2D homog)

↑  
cam<sub>r</sub>  
(2D homog)

$$\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} R_r & | & t_r \end{bmatrix} \begin{bmatrix} x_w \\ y_w \\ z_w \\ 1 \end{bmatrix}$$

↑  
cam<sub>r</sub> (3D homog)


↑  
world







# Plane Sweep Stereo



- each plane defines an image  $\Rightarrow$  composite homography



















