

Stereo II



CSC420

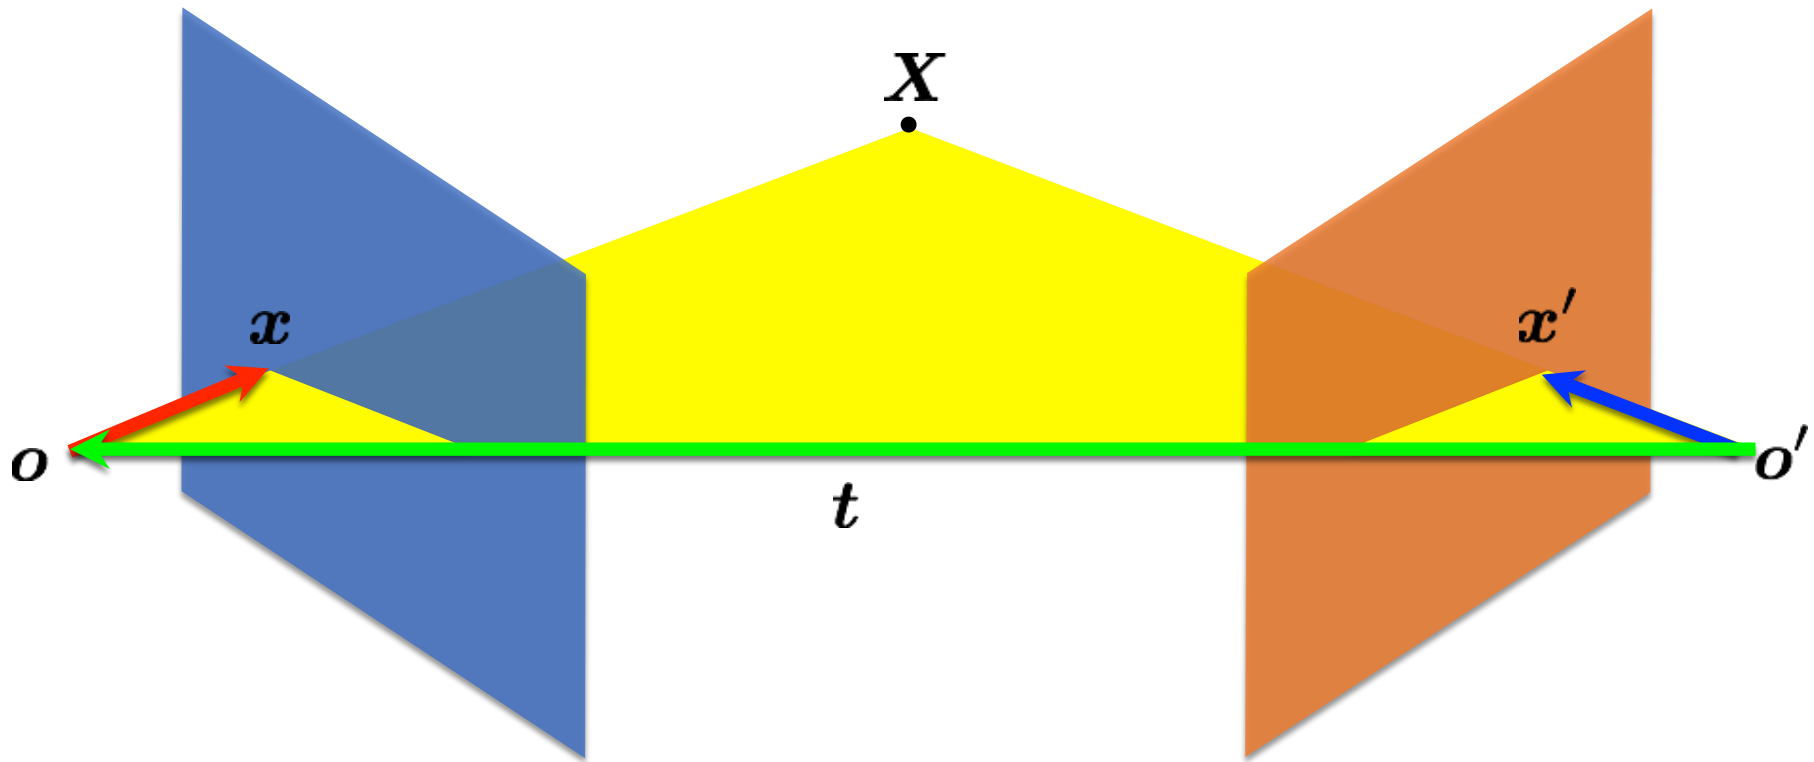
David Lindell

University of Toronto

cs.toronto.edu/~lindell/teaching/420

Slide credit: Babak Taati ← Ahmed Ashraf ← Sanja Fidler, Yannis Gkioulekas, Kris Kitani, Srinivasa Narasimhan

Recap



If these three vectors are coplanar $\mathbf{x}, \mathbf{t}, \mathbf{x}'$ then

$$(\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

Essential Matrix

rigid motion

$$\mathbf{x}' = \mathbf{R}(\mathbf{x} - \mathbf{t})$$

coplanarity

$$(\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})(\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})([\mathbf{t}_\times] \mathbf{x}) = 0$$

$$\mathbf{x}'^\top (\mathbf{R}[\mathbf{t}_\times]) \mathbf{x} = 0$$


$$\mathbf{x}'^\top \mathbf{E} \mathbf{x} = 0$$

Essential Matrix

[Longuet-Higgins 1981]

$$\hat{\mathbf{x}}'^{\top} \mathbf{E} \hat{\mathbf{x}} = 0$$

The essential matrix operates on image points expressed in **2D coordinates** expressed in the camera coordinate system

$$\hat{\mathbf{x}}' = \mathbf{K}'^{-1} \mathbf{x}'$$

$$\hat{\mathbf{x}} = \mathbf{K}^{-1} \mathbf{x}$$

camera point image point

Writing out the epipolar constraint in terms of image coordinates

$$\mathbf{K}'^{-\top} \mathbf{E} \mathbf{K}^{-1} \mathbf{x} = 0$$

$$\mathbf{x}'^{\top} (\mathbf{K}'^{-\top} \mathbf{E} \mathbf{K}^{-1}) \mathbf{x} = 0$$

$$\mathbf{x}'^{\top} \mathbf{F} \mathbf{x} = 0$$

properties of the F ~~E~~ matrix

Longuet-Higgins equation

$$x'^{\top} F x = 0$$

Epipolar lines

$$x^{\top} l = 0$$

$$l' = F x$$

$$x'^{\top} l' = 0$$

$$l = F^T x'$$

Epipoles

$$e'^{\top} F = 0$$

$$F e = 0$$

(points in **image** coordinates)

Breaking down the fundamental matrix

$$\mathbf{F} = \mathbf{K}'^{-\top} \mathbf{E} \mathbf{K}^{-1}$$

$$\mathbf{F} = \mathbf{K}'^{-\top} [\mathbf{t}_x] \mathbf{R} \mathbf{K}^{-1}$$

Depends on both intrinsic and extrinsic parameters

Breaking down the fundamental matrix

$$\mathbf{F} = \mathbf{K}'^{-\top} \mathbf{E} \mathbf{K}^{-1}$$

$$\mathbf{F} = \mathbf{K}'^{-\top} [\mathbf{t}_x] \mathbf{R} \mathbf{K}^{-1}$$

Depends on both intrinsic and extrinsic parameters

How would you solve for F ?

$$\mathbf{x}_m'^{\top} \mathbf{F} \mathbf{x}_m = 0$$

The 8-point algorithm

Assume you have M matched *image* points

$$\{\mathbf{x}_m, \mathbf{x}'_m\} \quad m = 1, \dots, M$$

Each correspondence should satisfy

$$\mathbf{x}'_m{}^\top \mathbf{F} \mathbf{x}_m = 0$$

How would you solve for the 3 x 3 \mathbf{F} matrix?

Assume you have M matched *image* points

$$\{\mathbf{x}_m, \mathbf{x}'_m\} \quad m = 1, \dots, M$$

Each correspondence should satisfy

$$\mathbf{x}'_m{}^\top \mathbf{F} \mathbf{x}_m = 0$$

How would you solve for the 3 x 3 \mathbf{F} matrix?

S V D

Assume you have M matched *image* points

$$\{\mathbf{x}_m, \mathbf{x}'_m\} \quad m = 1, \dots, M$$

Each correspondence should satisfy

$$\mathbf{x}'_m{}^\top \mathbf{F} \mathbf{x}_m = 0$$

How would you solve for the 3 x 3 \mathbf{F} matrix?

Set up a homogeneous linear system with 9 unknowns

$$\mathbf{x}_m'^\top \mathbf{F} \mathbf{x}_m = 0$$

$$\begin{bmatrix} x'_m & y'_m & 1 \end{bmatrix} \begin{bmatrix} f_1 & f_2 & f_3 \\ f_4 & f_5 & f_6 \\ f_7 & f_8 & f_9 \end{bmatrix} \begin{bmatrix} x_m \\ y_m \\ 1 \end{bmatrix} = 0$$

How many equation do you get from one correspondence?

$$\begin{bmatrix} x'_m & y'_m & 1 \end{bmatrix} \begin{bmatrix} f_1 & f_2 & f_3 \\ f_4 & f_5 & f_6 \\ f_7 & f_8 & f_9 \end{bmatrix} \begin{bmatrix} x_m \\ y_m \\ 1 \end{bmatrix} = 0$$

ONE correspondence gives you ONE equation

$$\begin{aligned} x_m x'_m f_1 + x_m y'_m f_2 + x_m f_3 + \\ y_m x'_m f_4 + y_m y'_m f_5 + y_m f_6 + \\ x'_m f_7 + y'_m f_8 + f_9 = 0 \end{aligned}$$

$$\begin{bmatrix} x'_m & y'_m & 1 \end{bmatrix} \begin{bmatrix} f_1 & f_2 & f_3 \\ f_4 & f_5 & f_6 \\ f_7 & f_8 & f_9 \end{bmatrix} \begin{bmatrix} x_m \\ y_m \\ 1 \end{bmatrix} = 0$$

Set up a homogeneous linear system with 9 unknowns

$$\begin{bmatrix} x_1 x'_1 & x_1 y'_1 & x_1 & y_1 x'_1 & y_1 y'_1 & y_1 & x'_1 & y'_1 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_M x'_M & x_M y'_M & x_M & y_M x'_M & y_M y'_M & y_M & x'_M & y'_M & 1 \end{bmatrix} \begin{bmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \\ f_5 \\ f_6 \\ f_7 \\ f_8 \\ f_9 \end{bmatrix} = \mathbf{0}$$

How many equations do you need?

Each point pair (according to epipolar constraint)
contributes only one scalar equation

$$\mathbf{x}_m'^\top \mathbf{F} \mathbf{x}_m = 0$$

Note: This is different from the Homography estimation
where each point pair contributes 2 equations.

We need at least 8 points

Hence, the 8 point algorithm!

How do you solve a homogeneous linear system?

$$\mathbf{A}\mathbf{X} = \mathbf{0}$$

How do you solve a homogeneous linear system?

$$\mathbf{A}\mathbf{X} = \mathbf{0}$$

Total Least Squares

minimize $\|\mathbf{A}\mathbf{x}\|^2$

subject to $\|\mathbf{x}\|^2 = 1$

How do you solve a homogeneous linear system?

$$\mathbf{A}\mathbf{X} = \mathbf{0}$$

Total Least Squares

minimize $\|\mathbf{A}\mathbf{x}\|^2$

subject to $\|\mathbf{x}\|^2 = 1$

S V D !

Eight-Point Algorithm

0. (Normalize points)
1. Construct the $M \times 9$ matrix **A**
2. Find the SVD of **A**
3. Entries of **F** are the elements of column of **V** corresponding to the least singular value
4. (Enforce rank 2 constraint on F)
5. (Un-normalize F)

Eight-Point Algorithm

0. (Normalize points)

1. Construct the $M \times 9$ matrix **A**


2. Find the SVD of **A**

3. Entries of **F** are the elements of column of **V** corresponding to the least singular value

4. (Enforce rank 2 constraint on F)

5. (Un-normalize F)

See Hartley-
Zisserman for why we
do this



Eight-Point Algorithm

0. (Normalize points)

1. Construct the $M \times 9$ matrix **A**

2. Find the SVD of **A**

3. Entries of **F** are the elements of column of **V** corresponding to the least singular value

4. (Enforce rank 2 constraint on F)

5. (Un-normalize F)



How do we do this?

Eight-Point Algorithm

0. (Normalize points)

1. Construct the $M \times 9$ matrix **A**

2. Find the SVD of **A**

3. Entries of **F** are the elements of column of **V** corresponding to the least singular value

4. (Enforce rank 2 constraint on F)

5. (Un-normalize F)

How do we do this?

S V D !

Enforcing rank constraints

Problem: Given a matrix F , find the matrix F' of rank k that is closest to F ,

$$\min_{\substack{F' \\ \text{rank}(F')=k}} \|F - F'\|^2$$

Solution: Compute the singular value decomposition of F ,

$$F = U\Sigma V^T$$

Form a matrix Σ' by replacing all but the k largest singular values in Σ with 0.

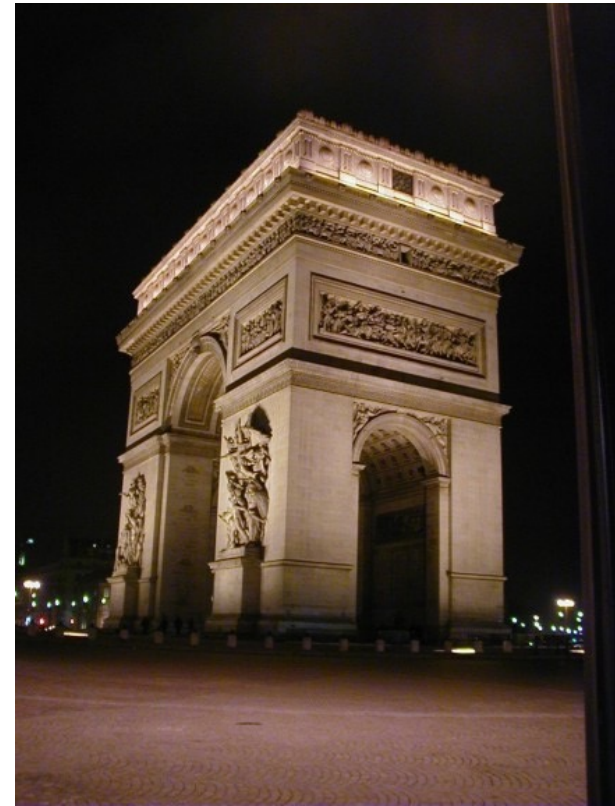
Then the problem solution is the matrix F' formed as,

$$F' = U\Sigma'V^T$$

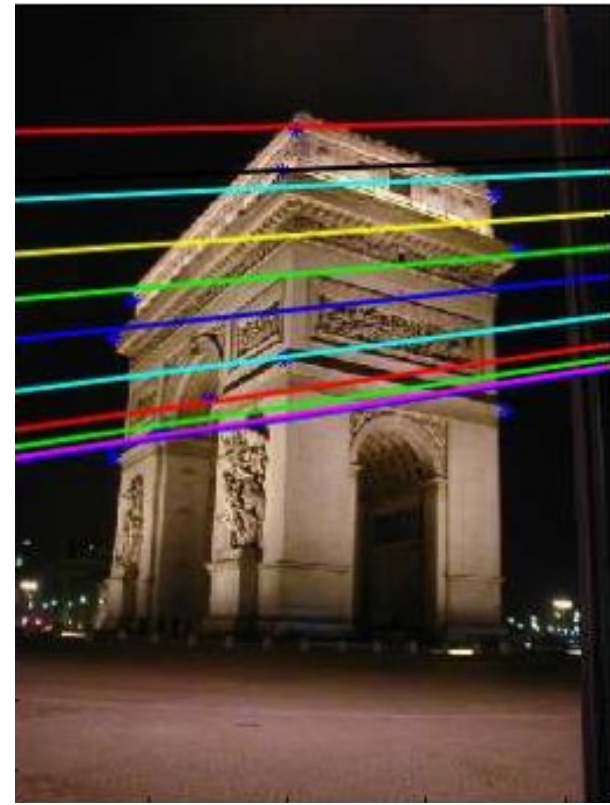
Eight-Point Algorithm

0. (Normalize points)
1. Construct the $M \times 9$ matrix **A**
2. Find the SVD of **A**
3. Entries of **F** are the elements of column of **V** corresponding to the least singular value
4. (Enforce rank 2 constraint on F)
5. (Un-normalize F)

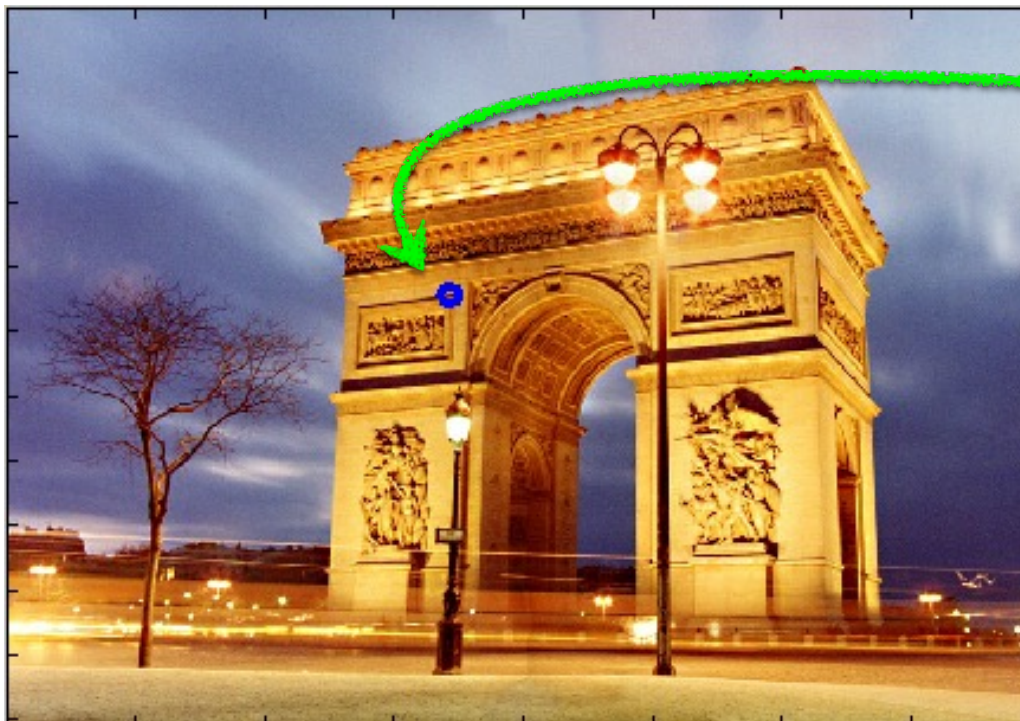
Example



epipolar lines



$$\mathbf{F} = \begin{bmatrix} -0.00310695 & -0.0025646 & 2.96584 \\ -0.028094 & -0.00771621 & 56.3813 \\ 13.1905 & -29.2007 & -9999.79 \end{bmatrix}$$

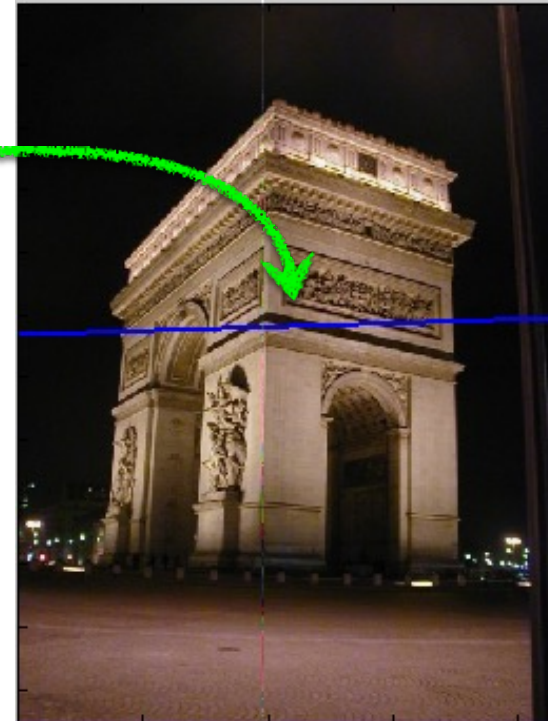
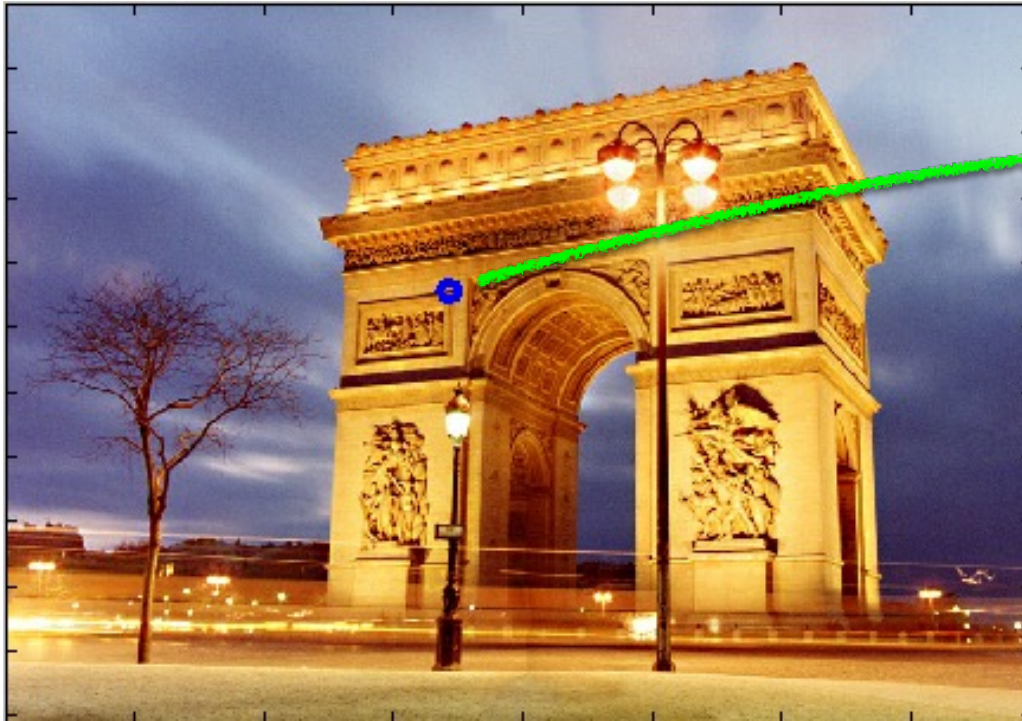


$$\mathbf{x} = \begin{bmatrix} 343.53 \\ 221.70 \\ 1.0 \end{bmatrix}$$

$$\begin{aligned} \mathbf{l}' &= \mathbf{F}\mathbf{x} \\ &= \begin{bmatrix} 0.0295 \\ 0.9996 \\ -265.1531 \end{bmatrix} \end{aligned}$$

$$l' = \mathbf{F}x$$

$$= \begin{bmatrix} 0.0295 \\ 0.9996 \\ -265.1531 \end{bmatrix}$$



Where is the epipole?



How would you compute it?



$$\mathbf{F}e = \mathbf{0}$$

The epipole is in the right null space of \mathbf{F}

How would you solve for the epipole?



$$\mathbf{F}e = \mathbf{0}$$

The epipole is in the right null space of \mathbf{F}

How would you solve for the epipole?

S V D !

Revisiting triangulation

How would you reconstruct 3D points?

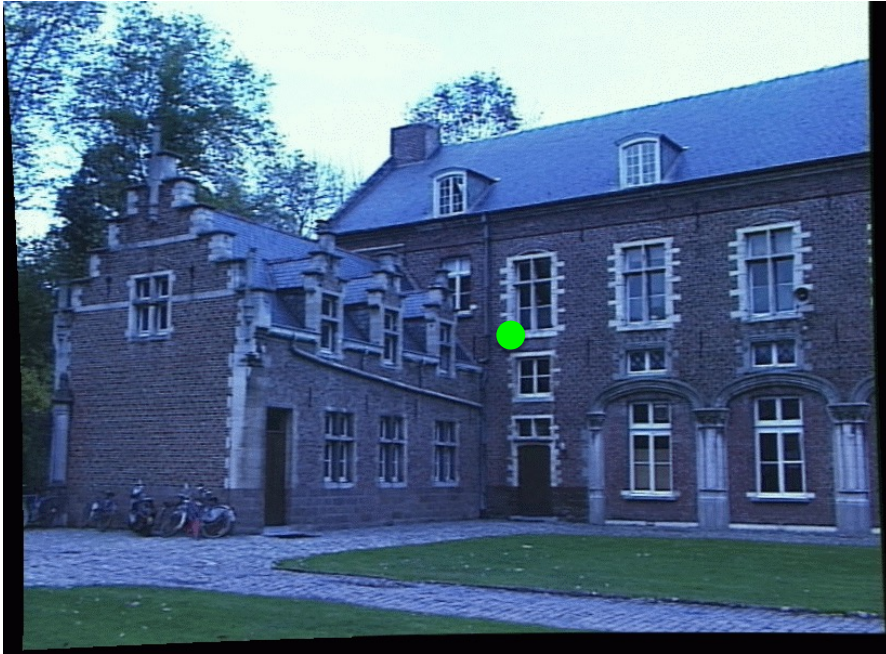


Left image



Right image

How would you reconstruct 3D points?



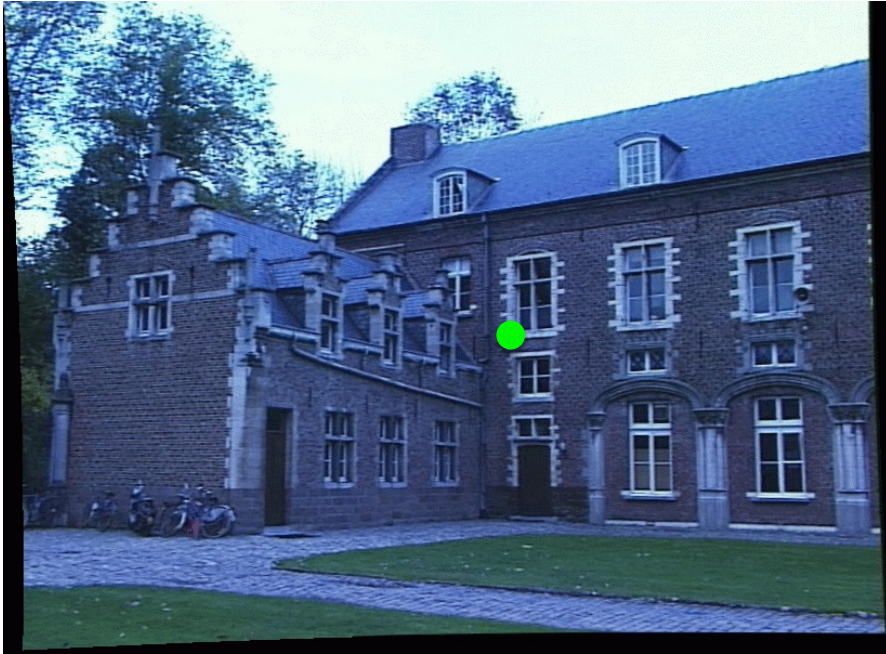
Left image



Right image

1. Select point in one image (how?)

How would you reconstruct 3D points?



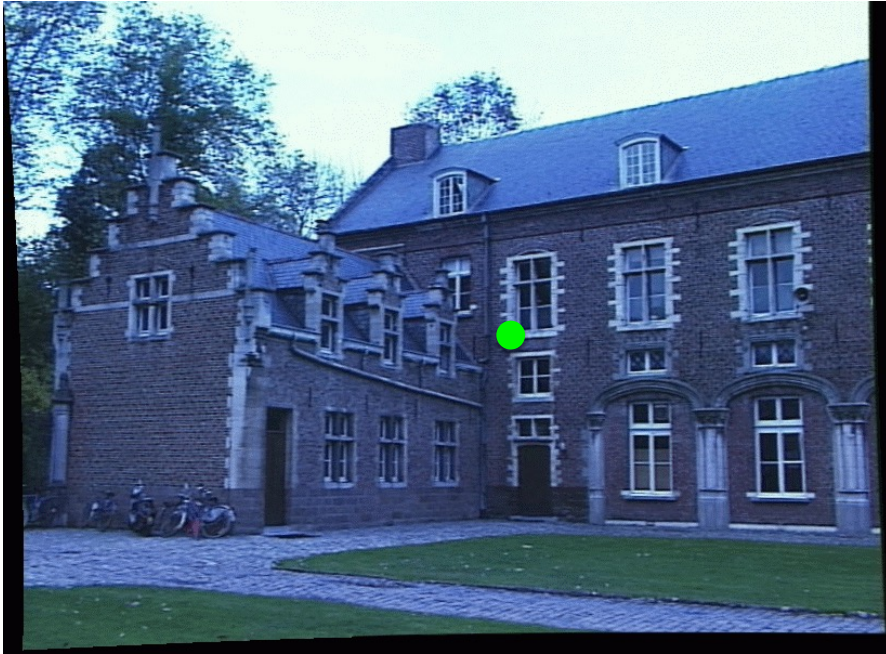
Left image



Right image

1. Select point in one image (how?)
2. Form epipolar line for that point in second image (how?)

How would you reconstruct 3D points?



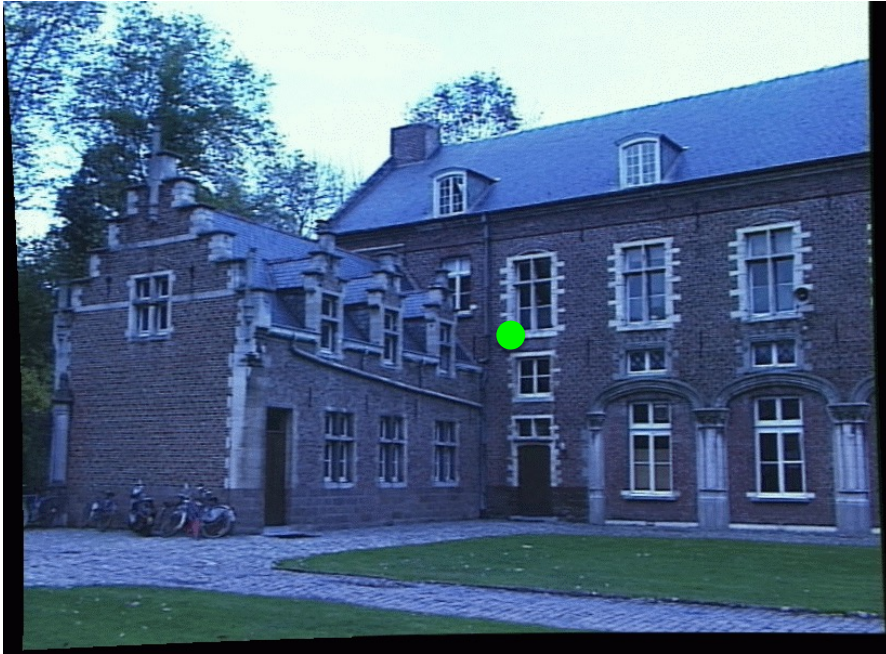
Left image



Right image

1. Select point in one image (how?)
2. Form epipolar line for that point in second image (how?)
3. Find matching point along line (how?)

How would you reconstruct 3D points?



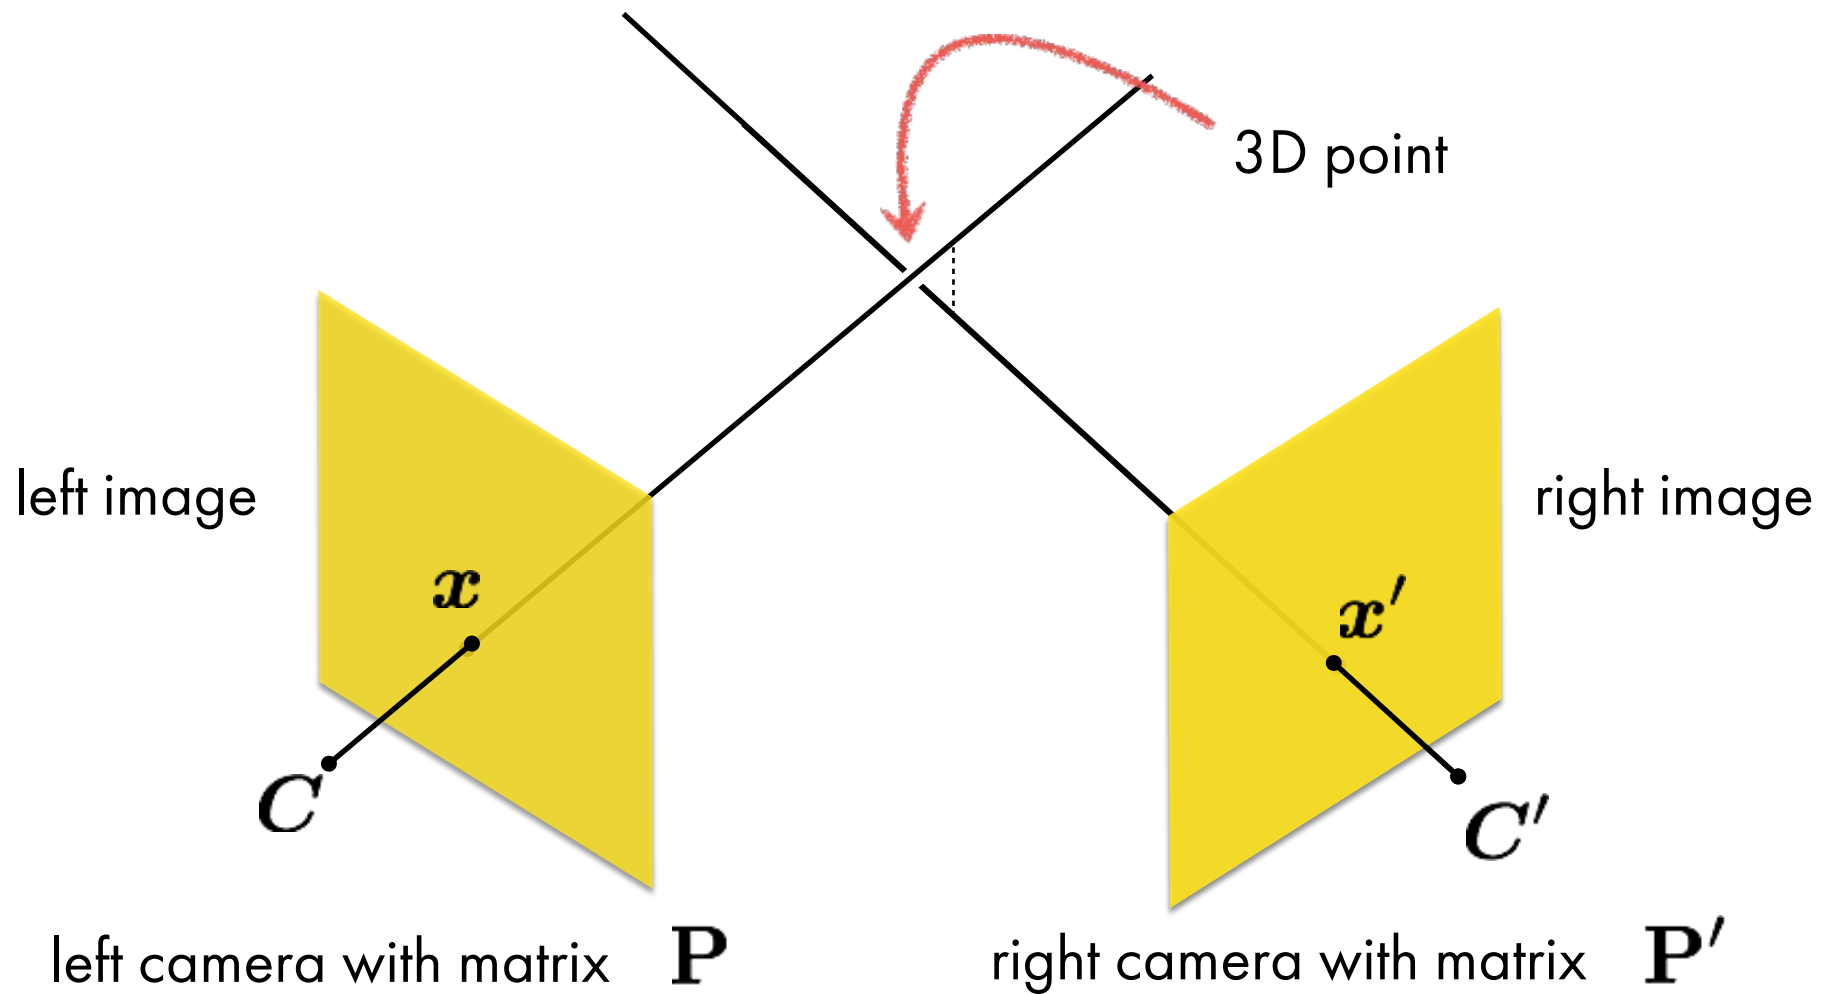
Left image



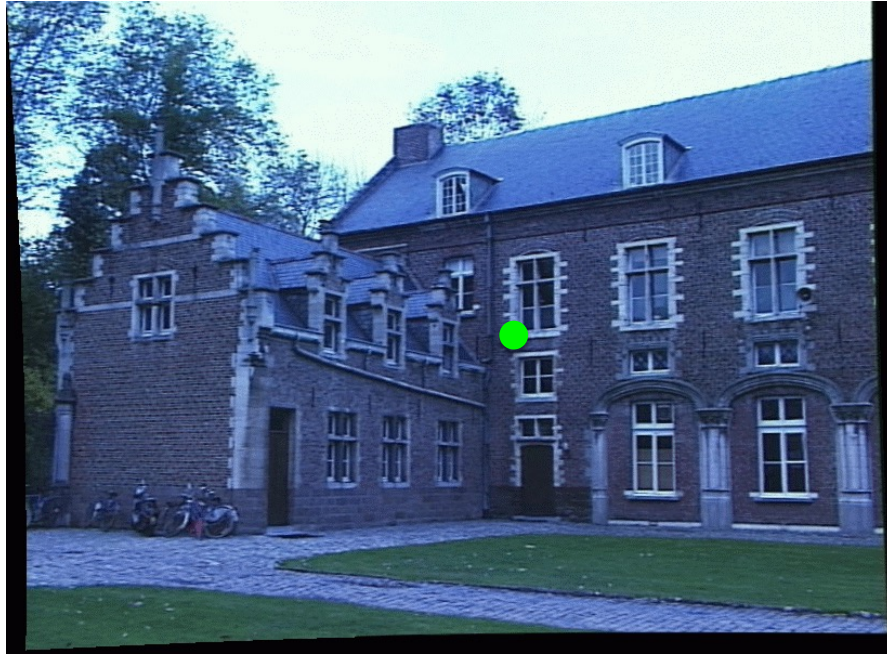
Right image

1. Select point in one image (how?)
2. Form epipolar line for that point in second image (how?)
3. Find matching point along line (how?)
4. Perform triangulation (how?)

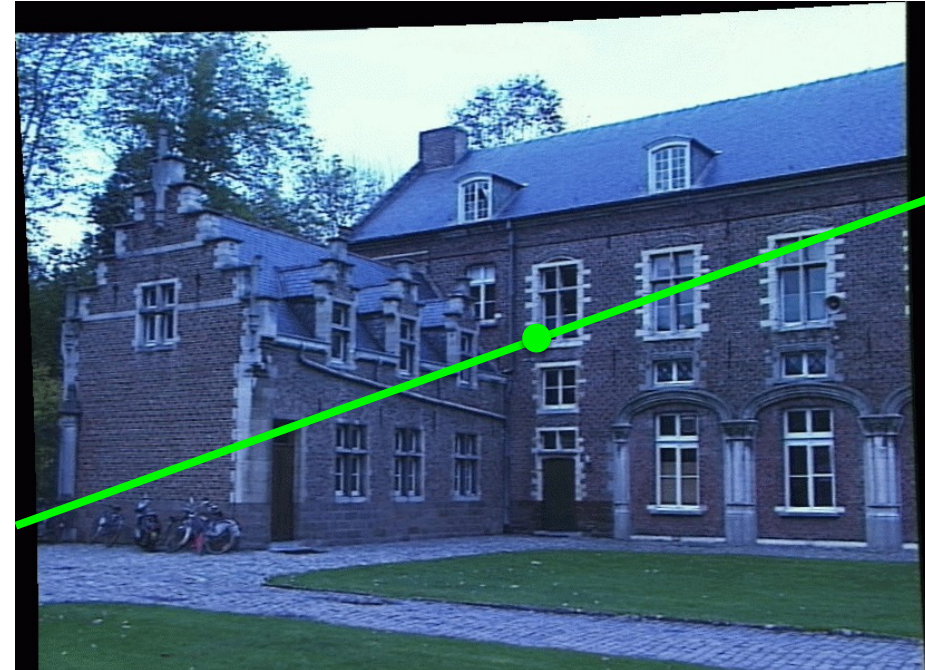
Triangulation



How would you reconstruct 3D points?



Left image



Right image

1. Select point in one image (how?)
2. Form epipolar line for that point in second image (how?)
3. Find matching point along line (how?)
4. Perform triangulation (how?)

What are the disadvantages of this procedure?

Stereo rectification



What's different between these two images?

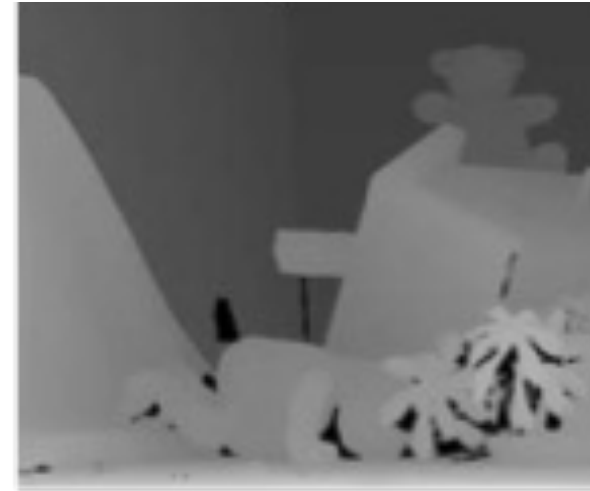




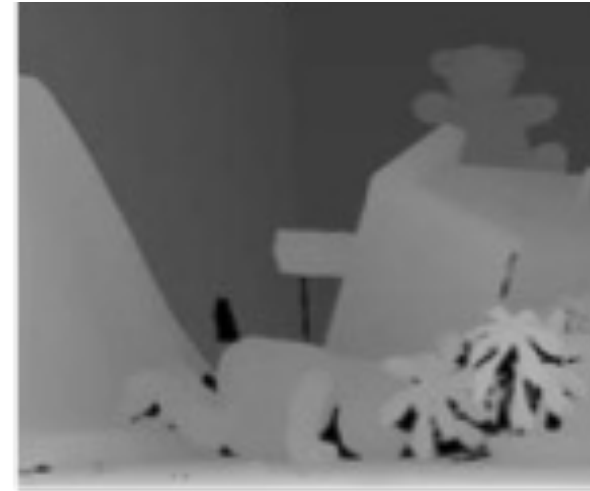


Objects that are close move more or less?

The amount of horizontal movement is inversely proportional to ...

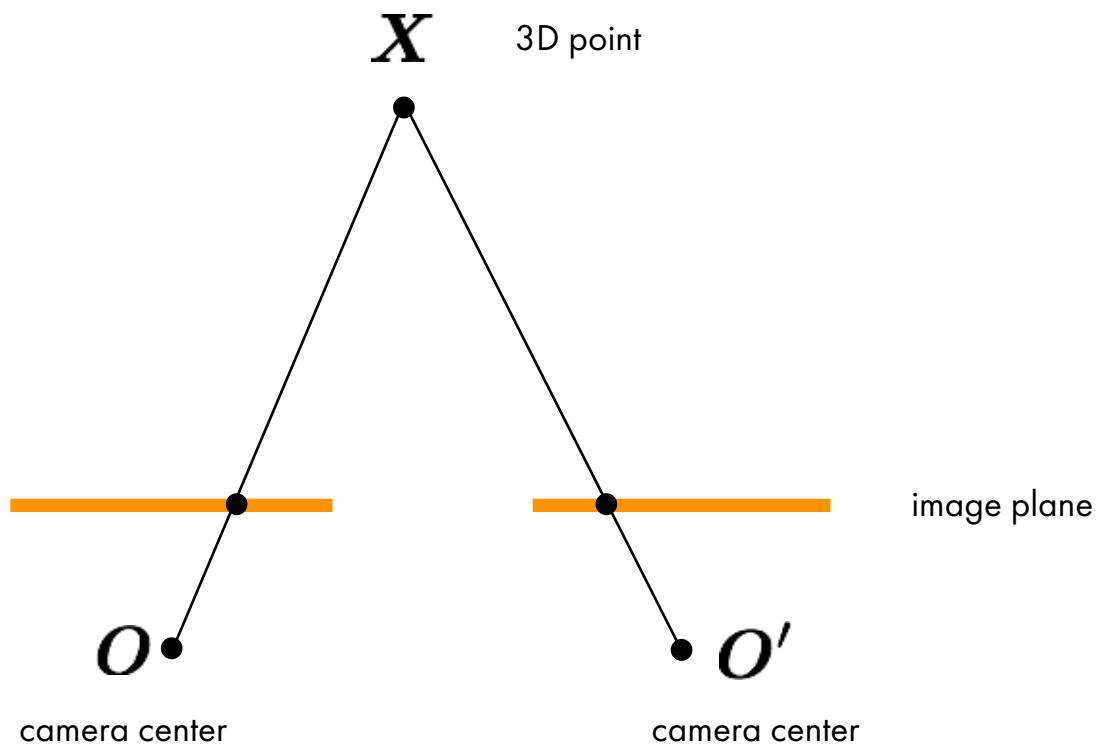


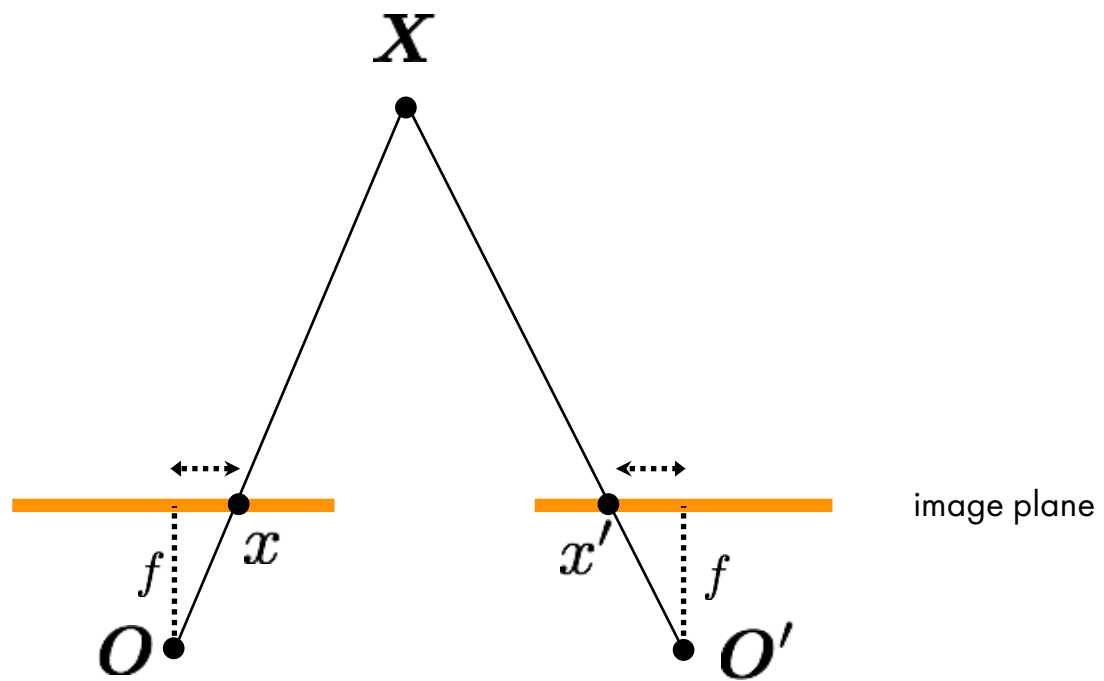
The amount of horizontal movement is inversely proportional to ...

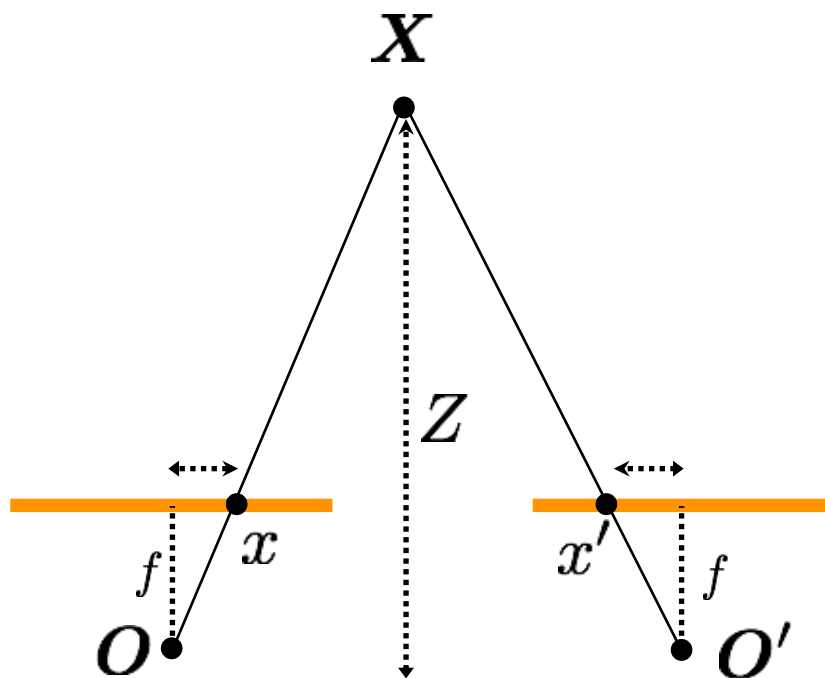


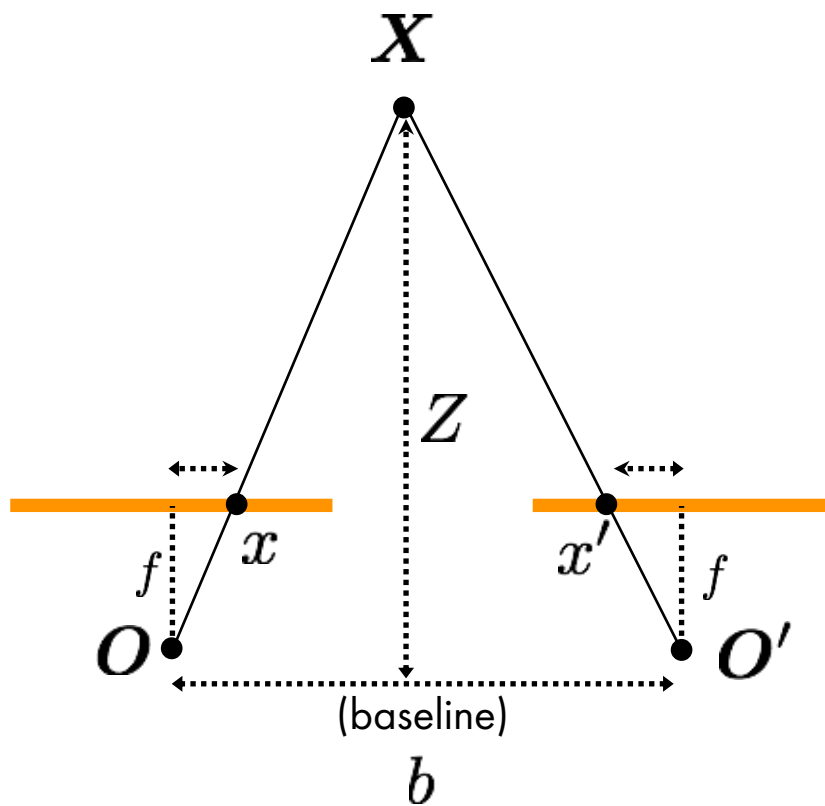
... the distance from the camera.

More formally...

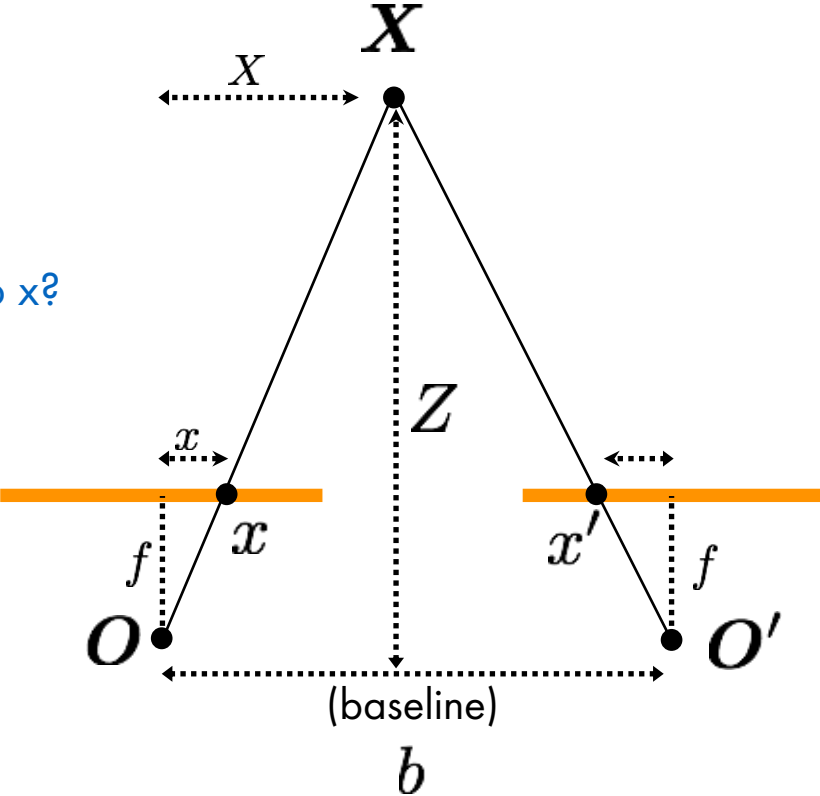




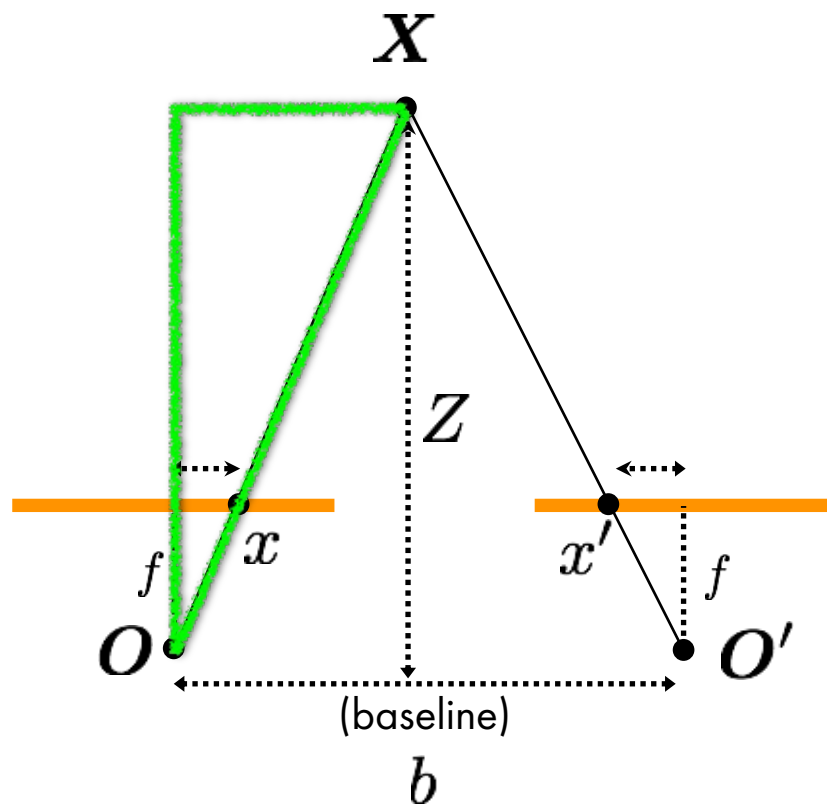




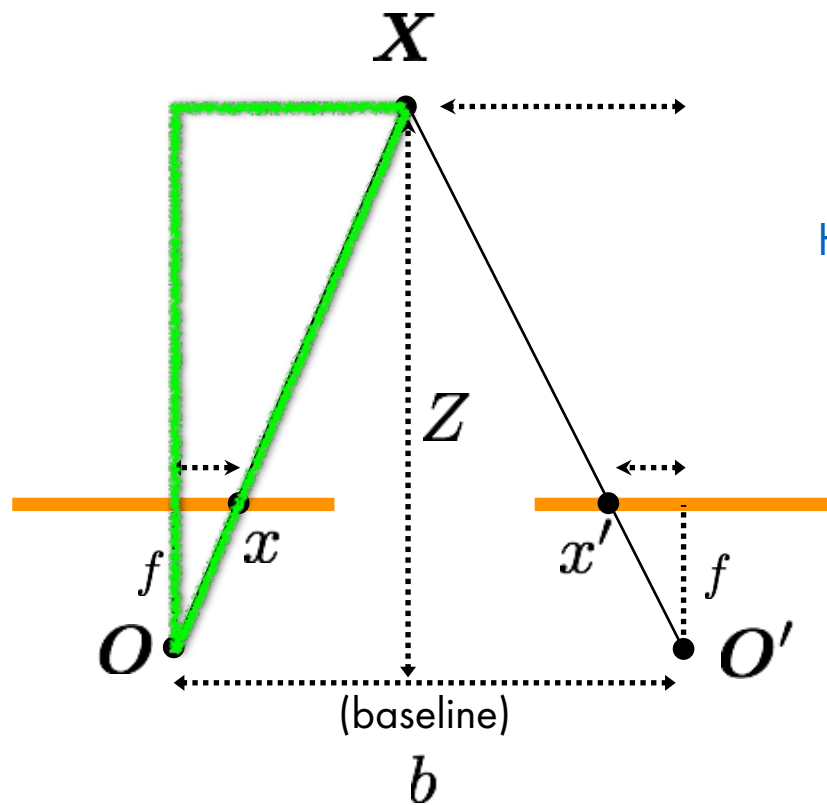
How is X related to x ?



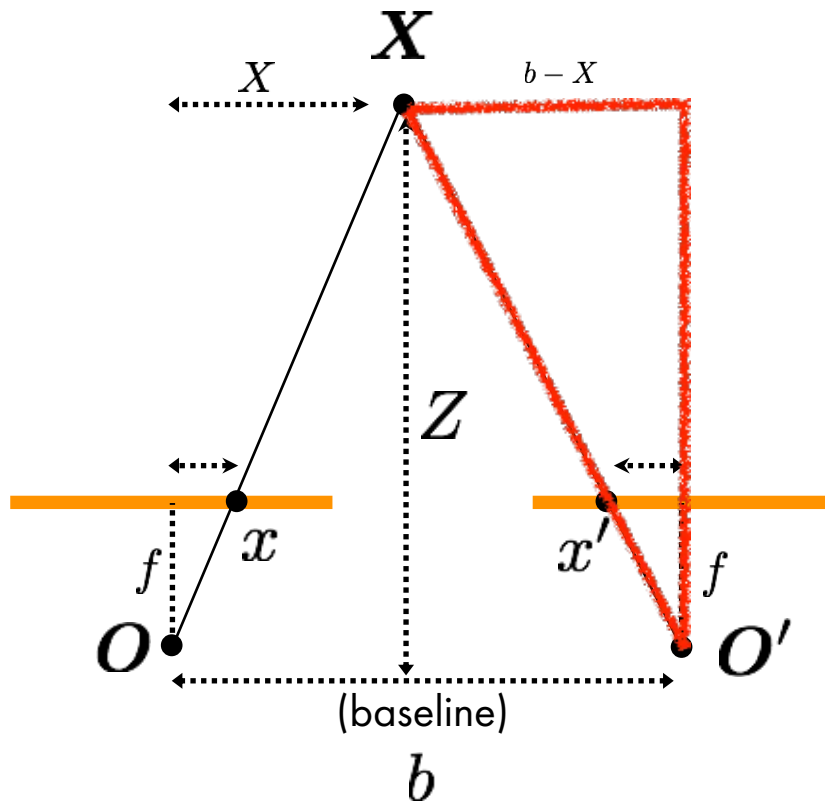
$$\frac{X}{Z} = \frac{x}{f}$$



$$\frac{X}{Z} = \frac{x}{f}$$

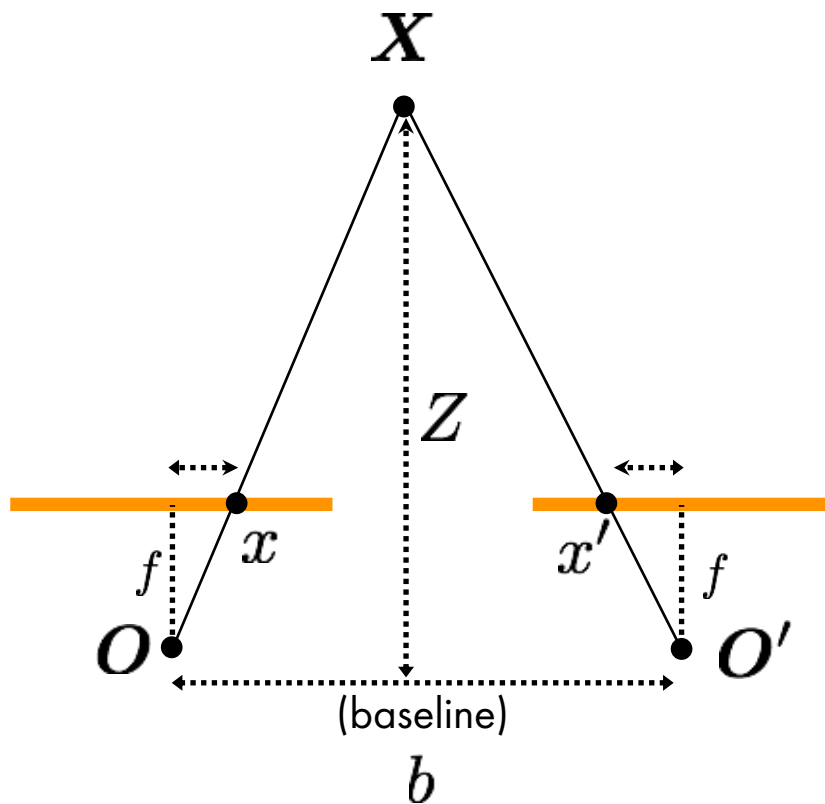


$$\frac{X}{Z} = \frac{x}{f}$$



$$\frac{b - X}{Z} = \frac{x'}{f}$$

$$\frac{X}{Z} = \frac{x}{f}$$



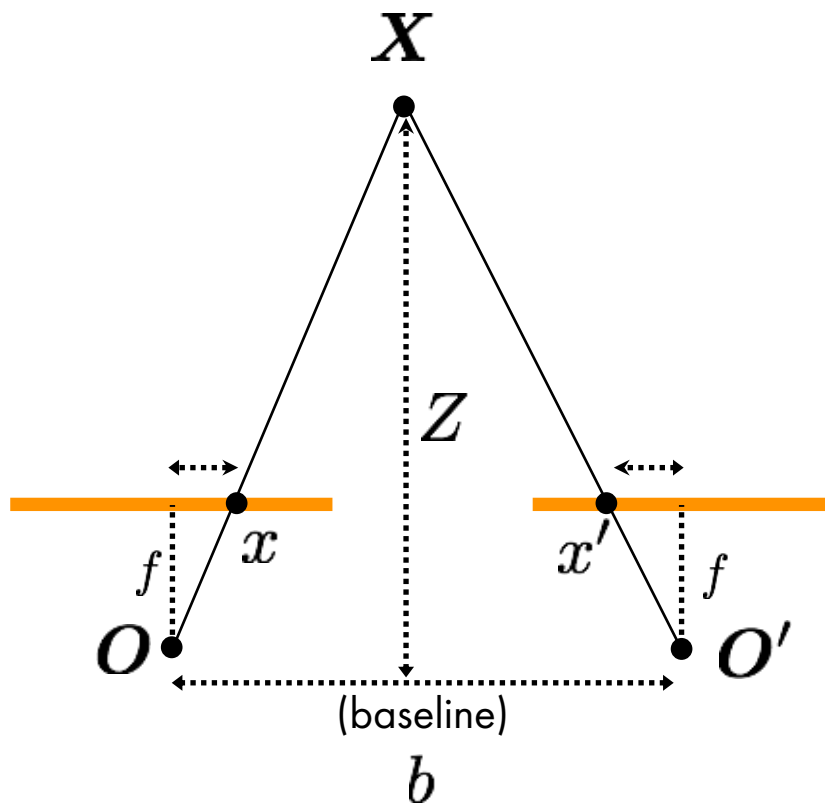
$$\frac{b - X}{Z} = \frac{x'}{f}$$

Disparity

$$d = x - x' \quad (\text{wrt to camera origin of image plane})$$

$$= \frac{bf}{Z}$$

$$\frac{X}{Z} = \frac{x}{f}$$



$$\frac{b - X}{Z} = \frac{x'}{f}$$

Disparity

$$d = x - x'$$

inversely proportional to
depth

$$= \frac{bf}{Z}$$



Subaru
Eyesight system

Pre-collision
braking





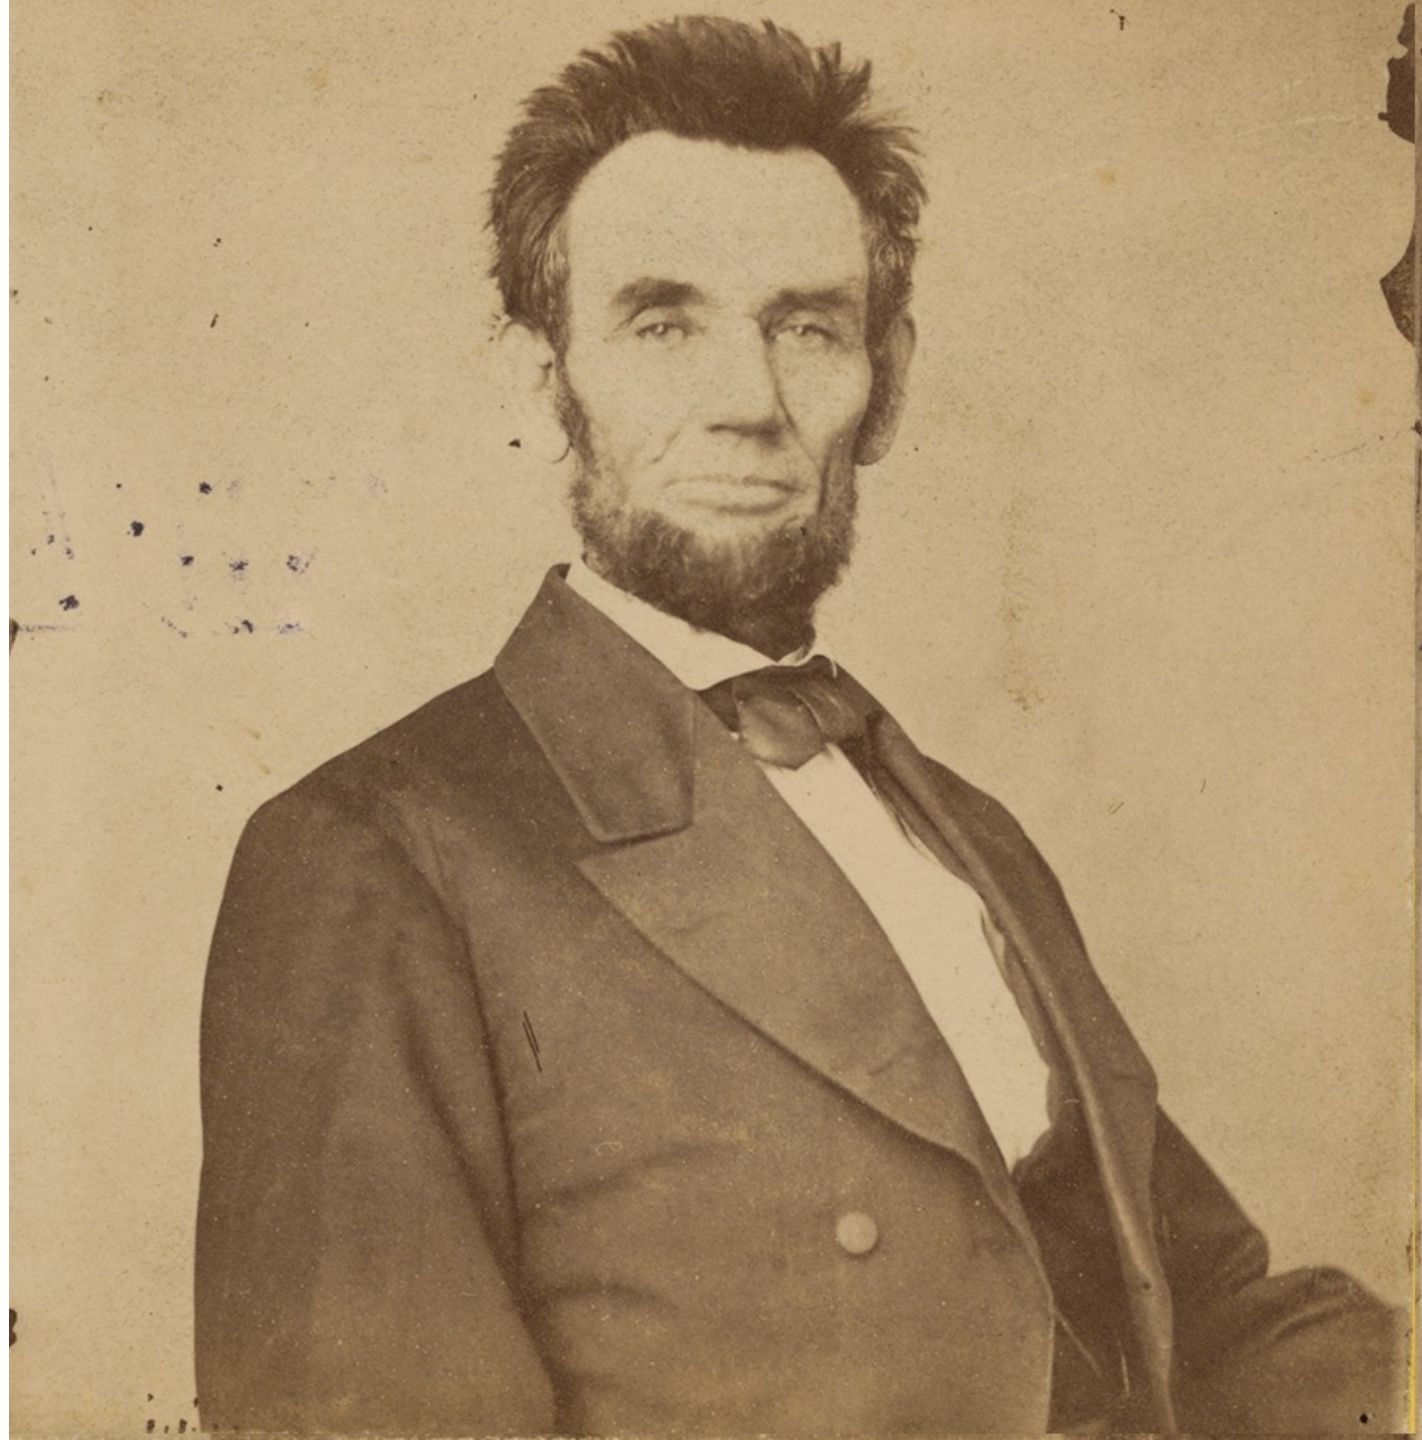
What other vision system uses disparity for depth sensing?

Stereoscopes: A 19th Century Pastime



HON. ABRAHAM LINCOLN, President of United States.







Public Library, Stereoscopic Looking Room, Chicago, by Phillips, 1923





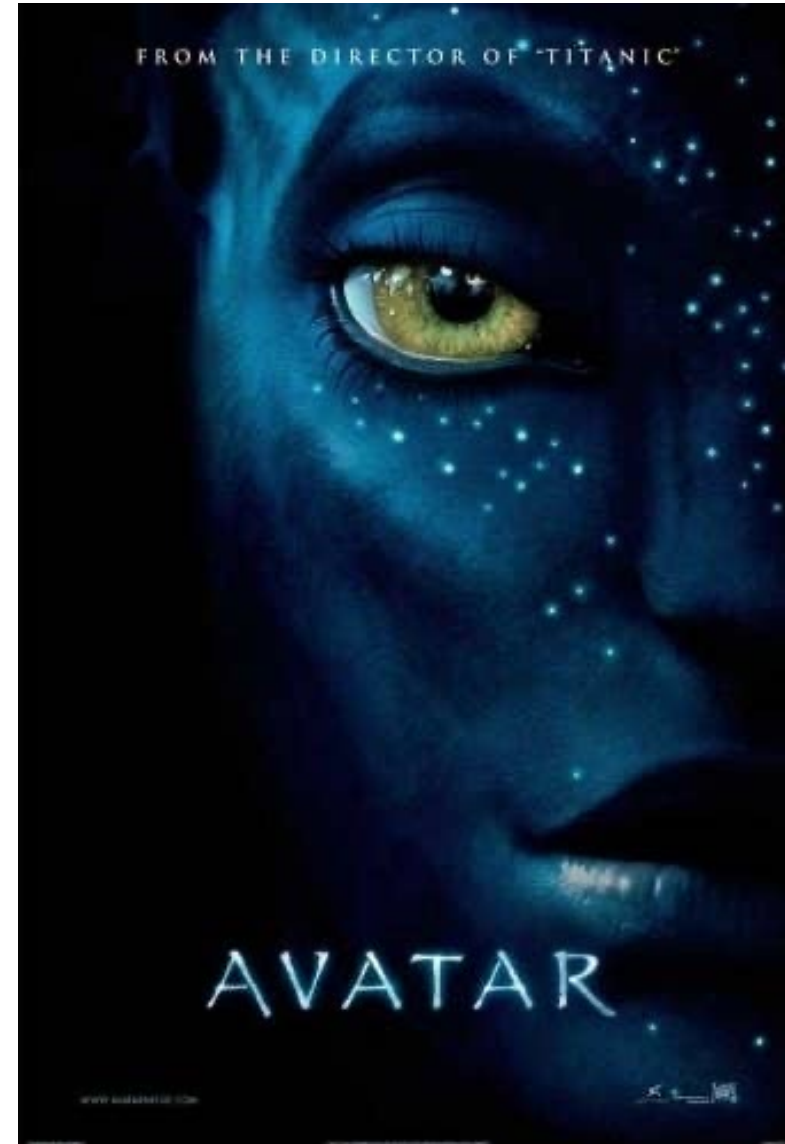
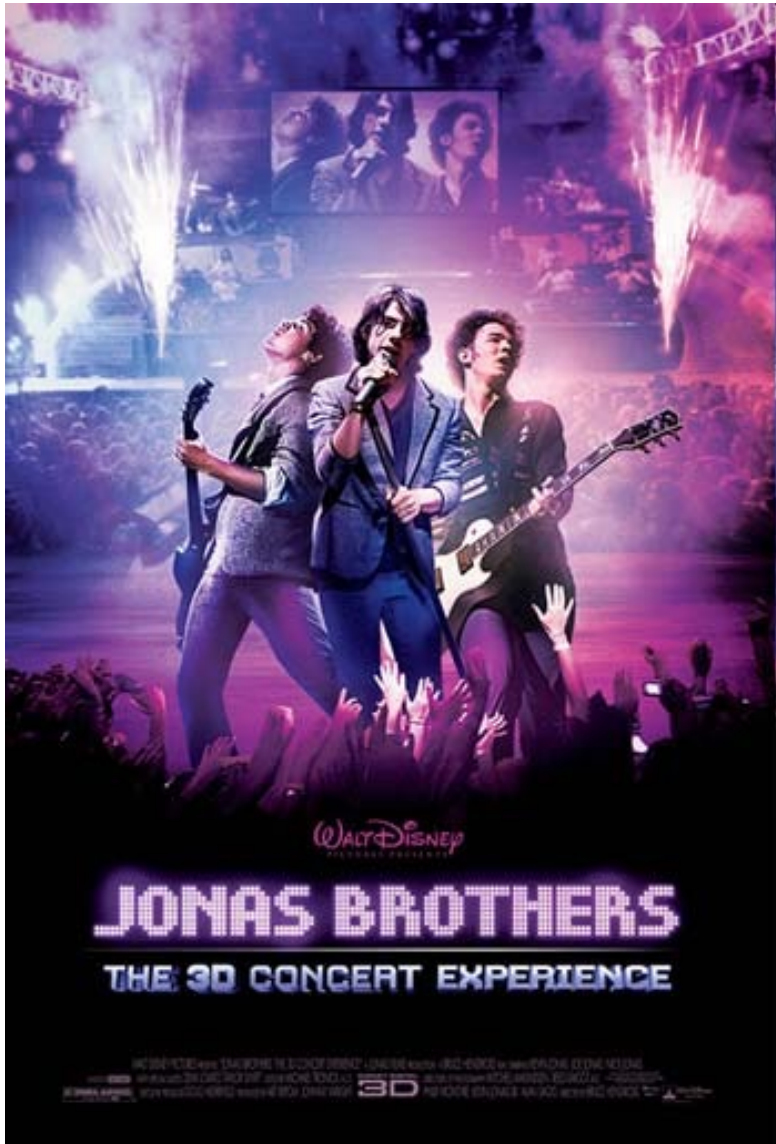
Teesta suspension bridge-Darjeeling, India



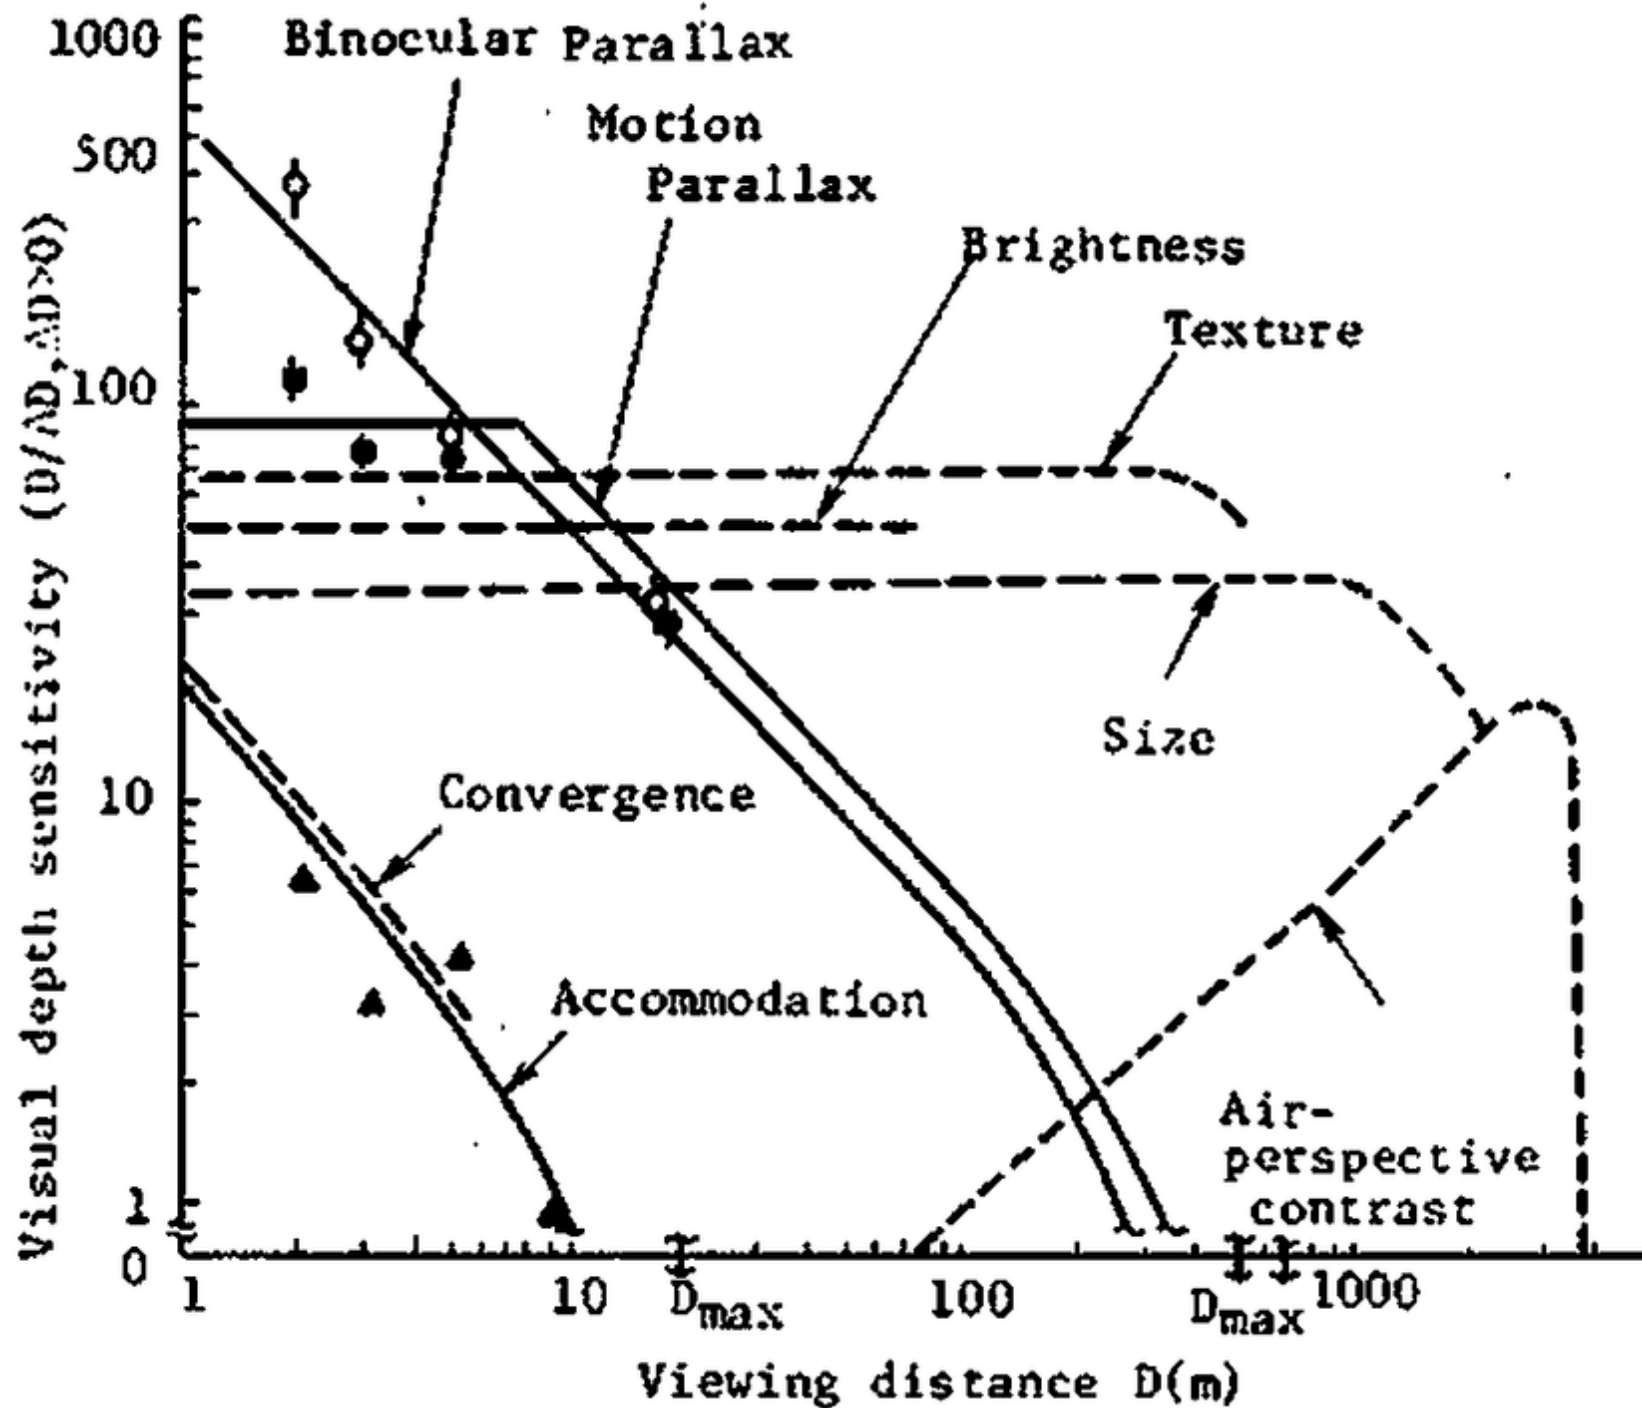


Mark Twain at Pool Table", no date, UCR Museum of Photography

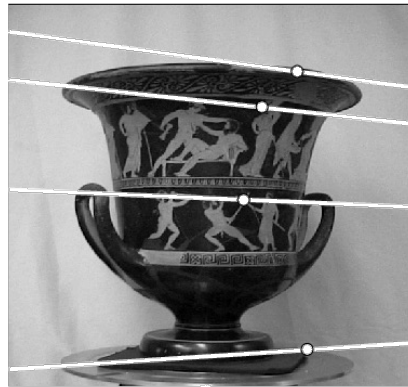
This is how 3D movies work



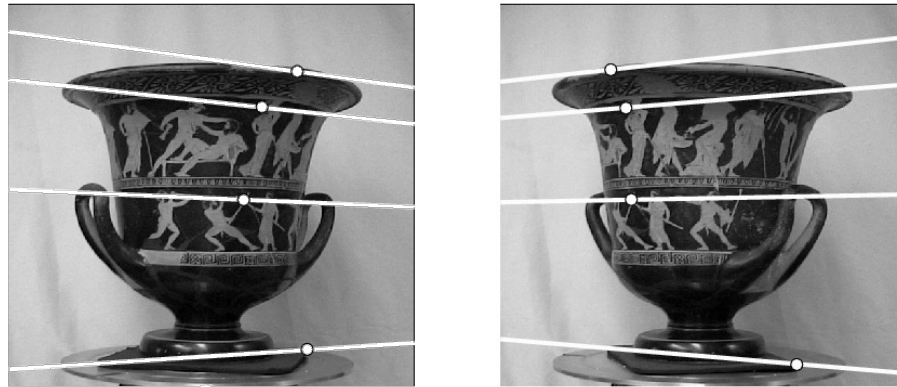
Is disparity the only depth cue the human visual system uses?



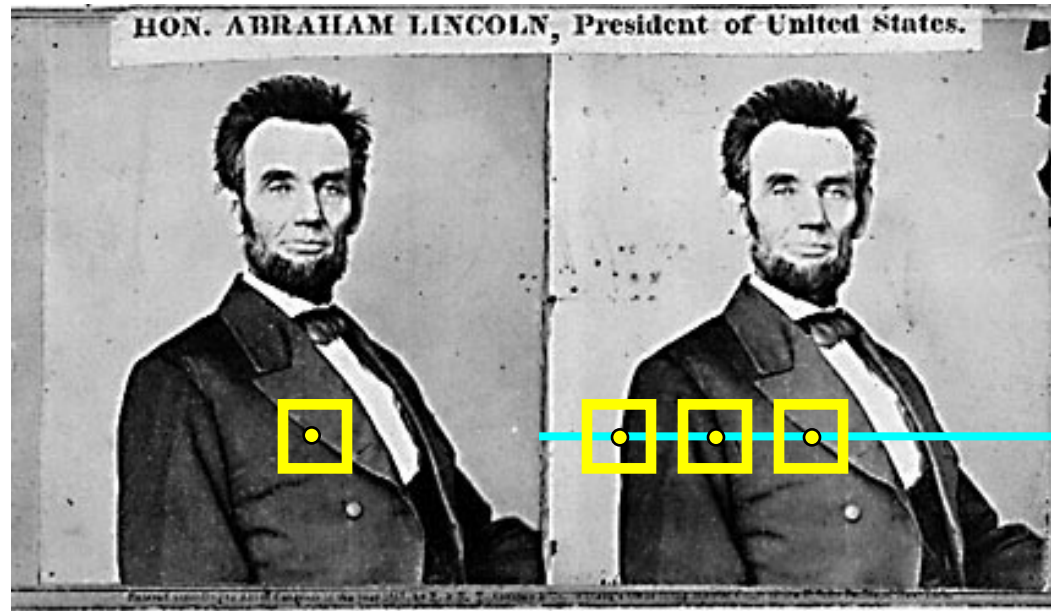
So can I compute depth from any two images
of the same object?



So can I compute depth from any two images
of the same object?

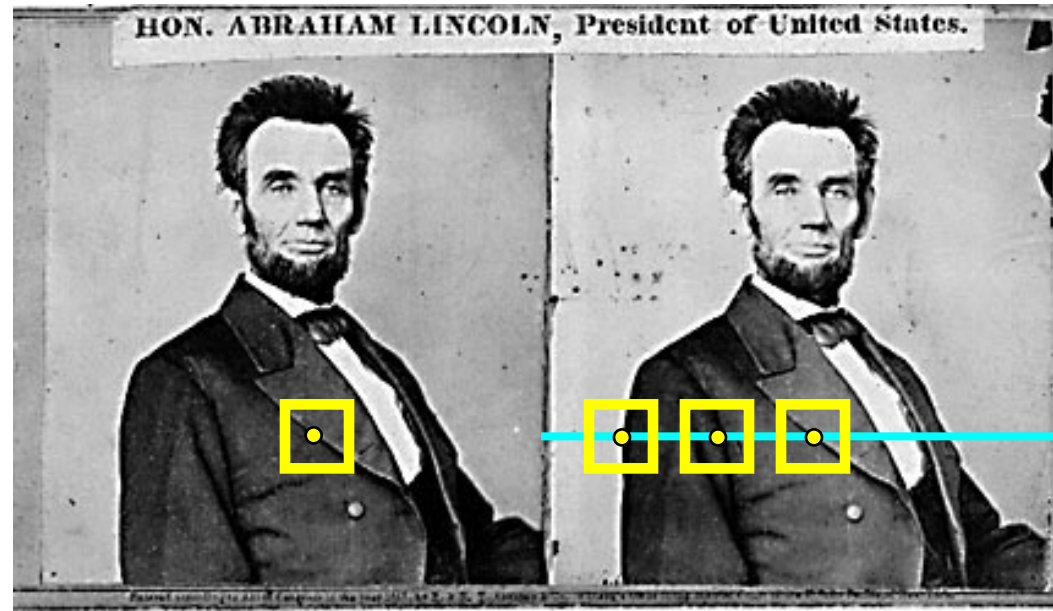


1. Need sufficient baseline
2. Images need to be 'rectified' first (make epipolar lines horizontal)

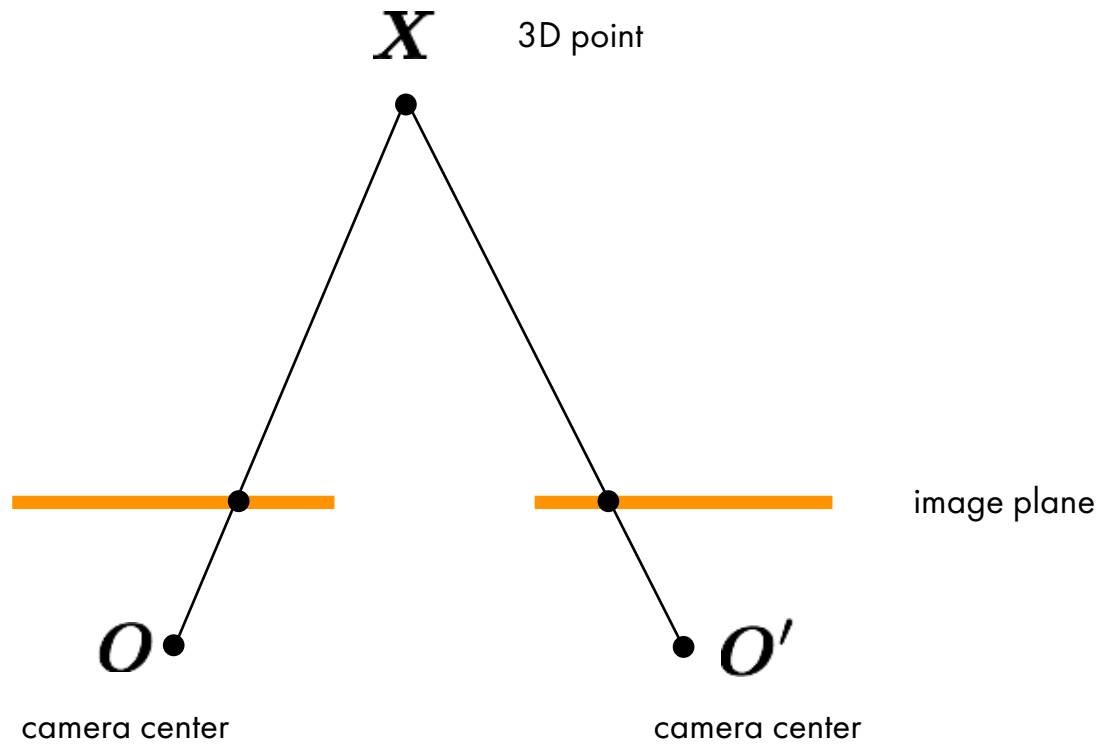


1. Rectify images
(make epipolar lines horizontal)
2. For each pixel
 - a. Find epipolar line
 - b. Scan line for best match
 - c. Compute depth from disparity

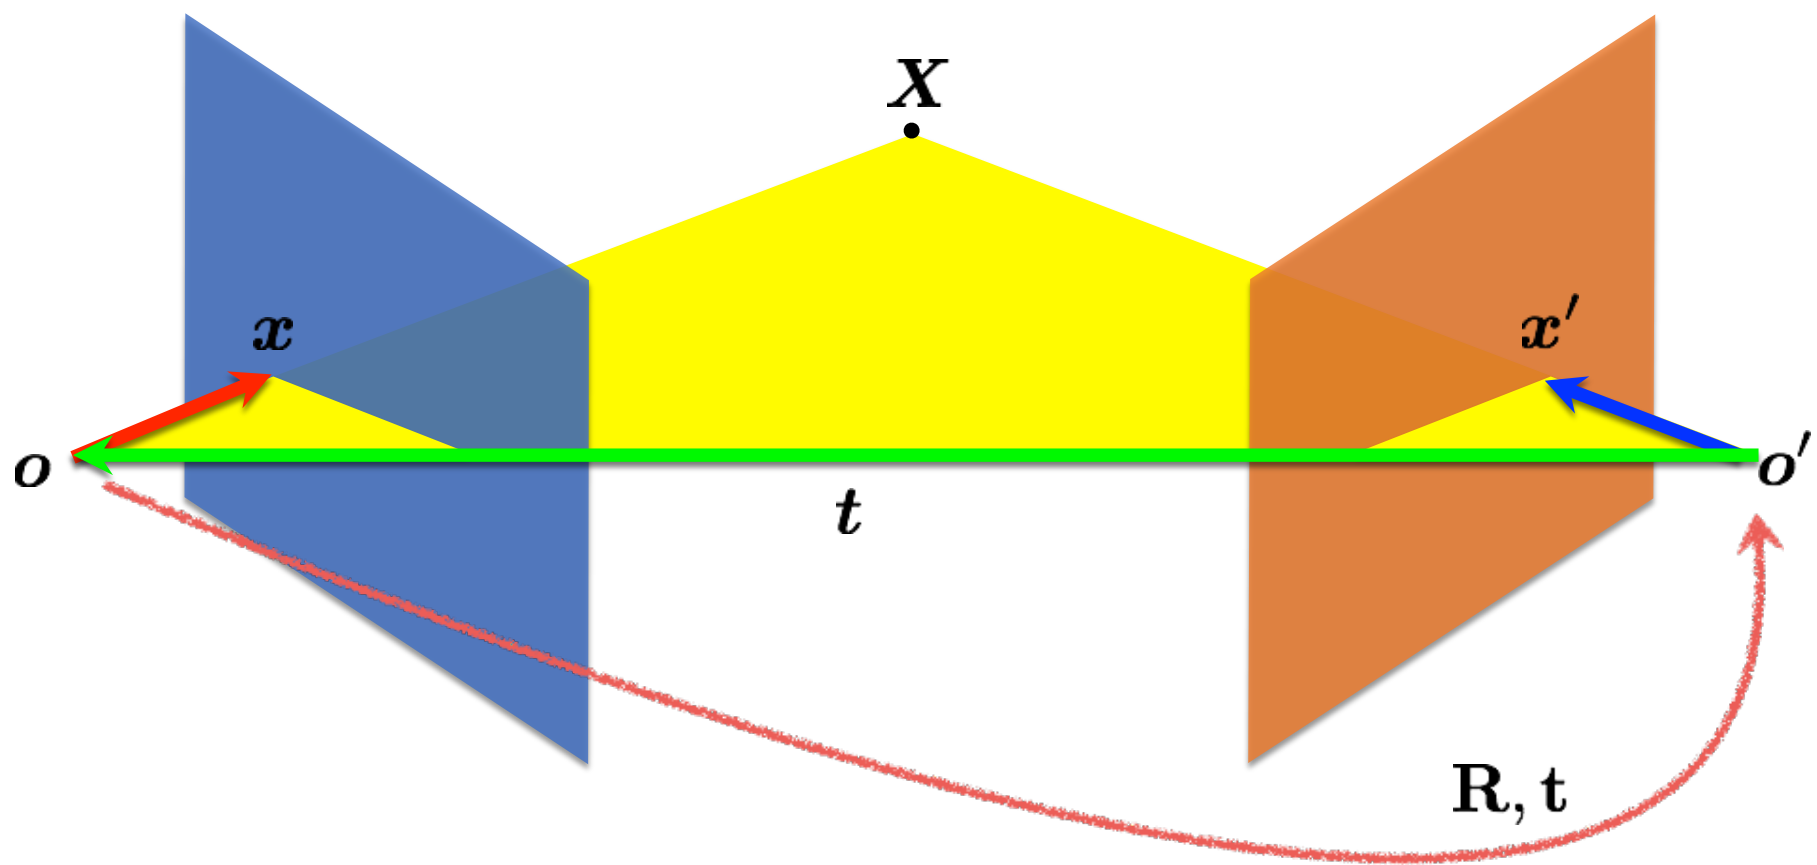
$$Z = \frac{bf}{d}$$



How can you make the epipolar lines horizontal?



What's special about these two cameras?

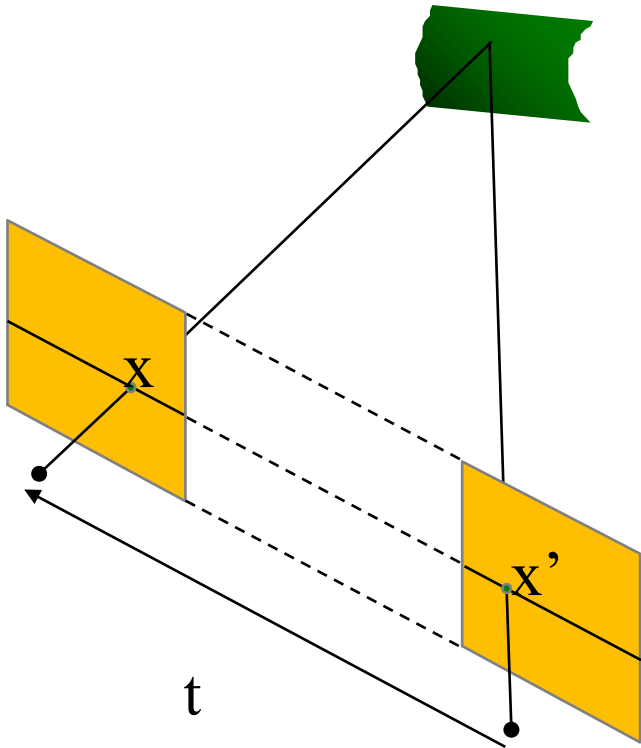


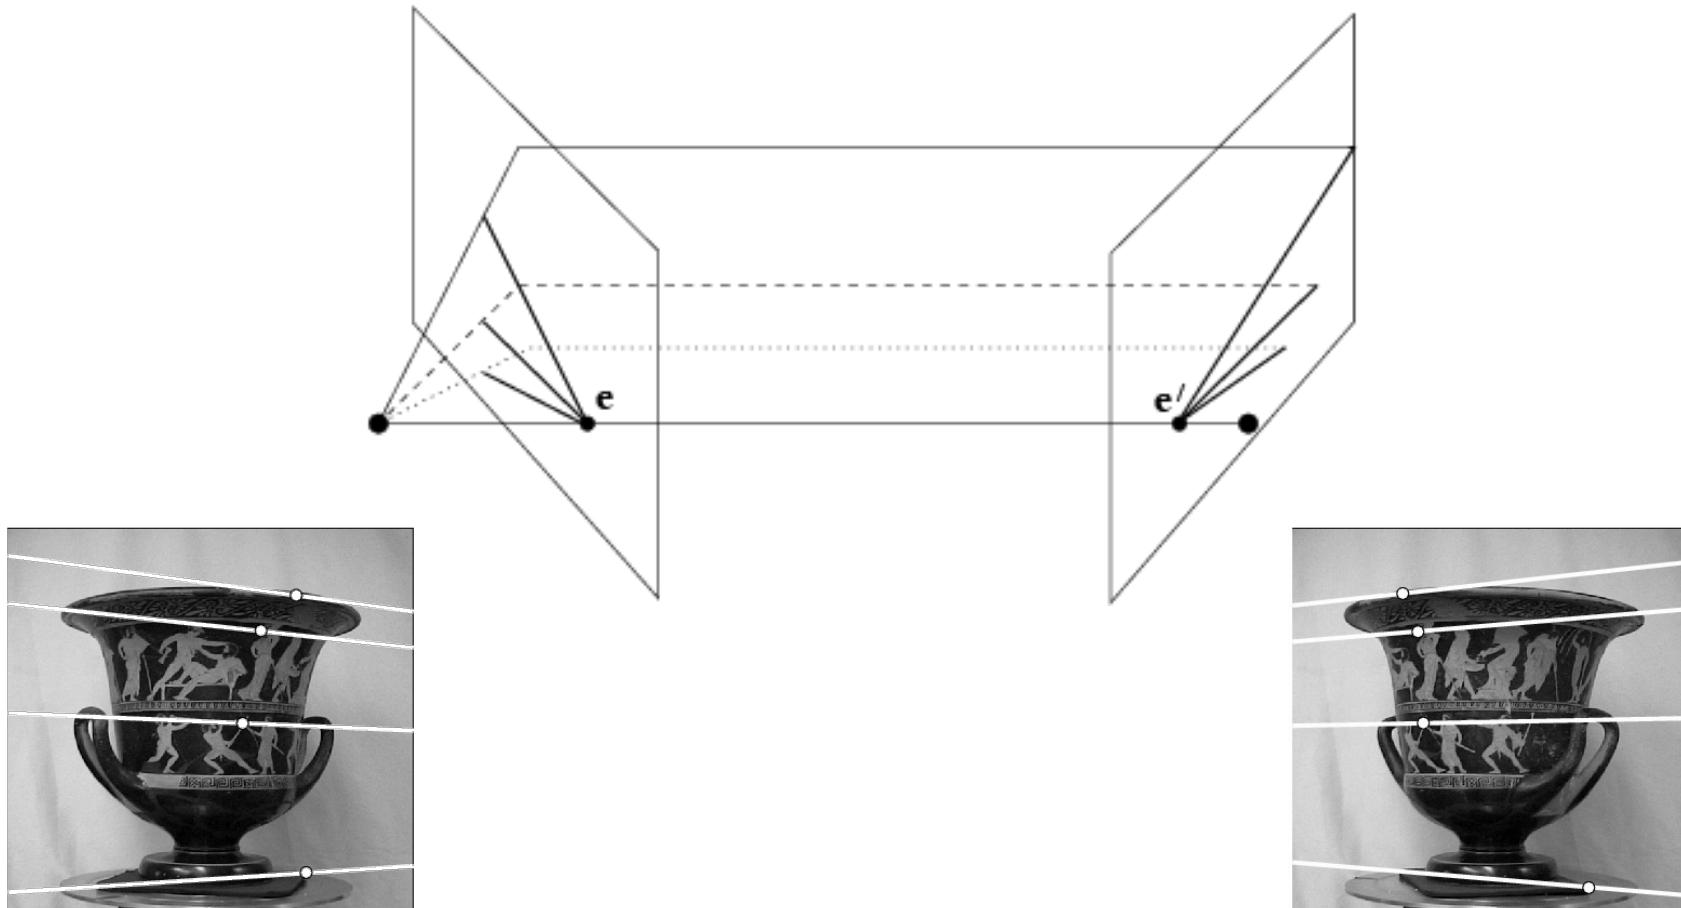
$$x' = \mathbf{R}(x - t)$$

When are epipolar lines horizontal?

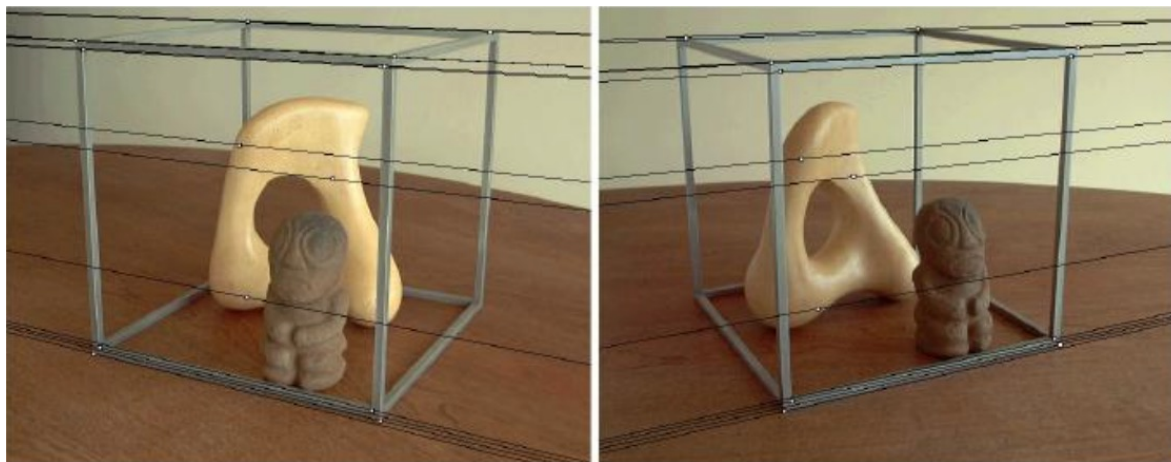
When this relationship holds:

$$R = I \quad t = (T, 0, 0)$$





It's hard to make the image planes exactly parallel

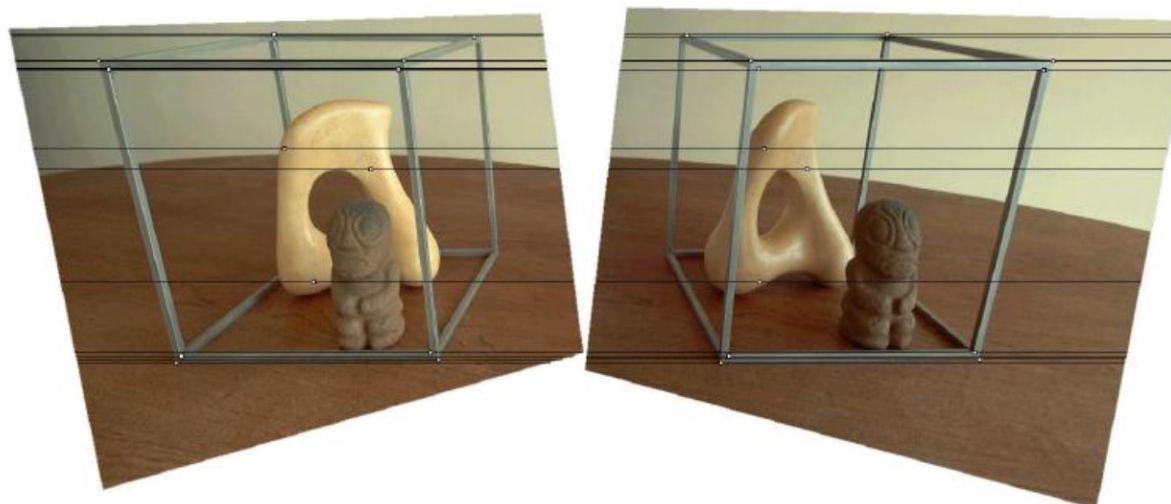


How can you make the epipolar lines horizontal?

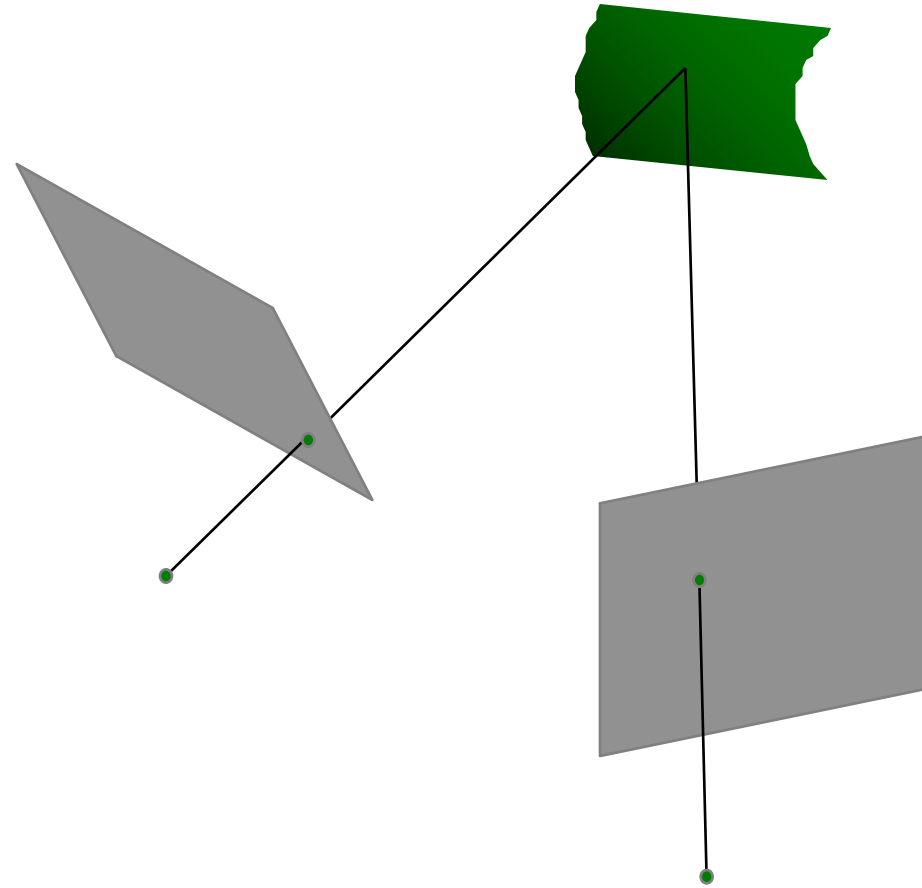




Use stereo rectification?



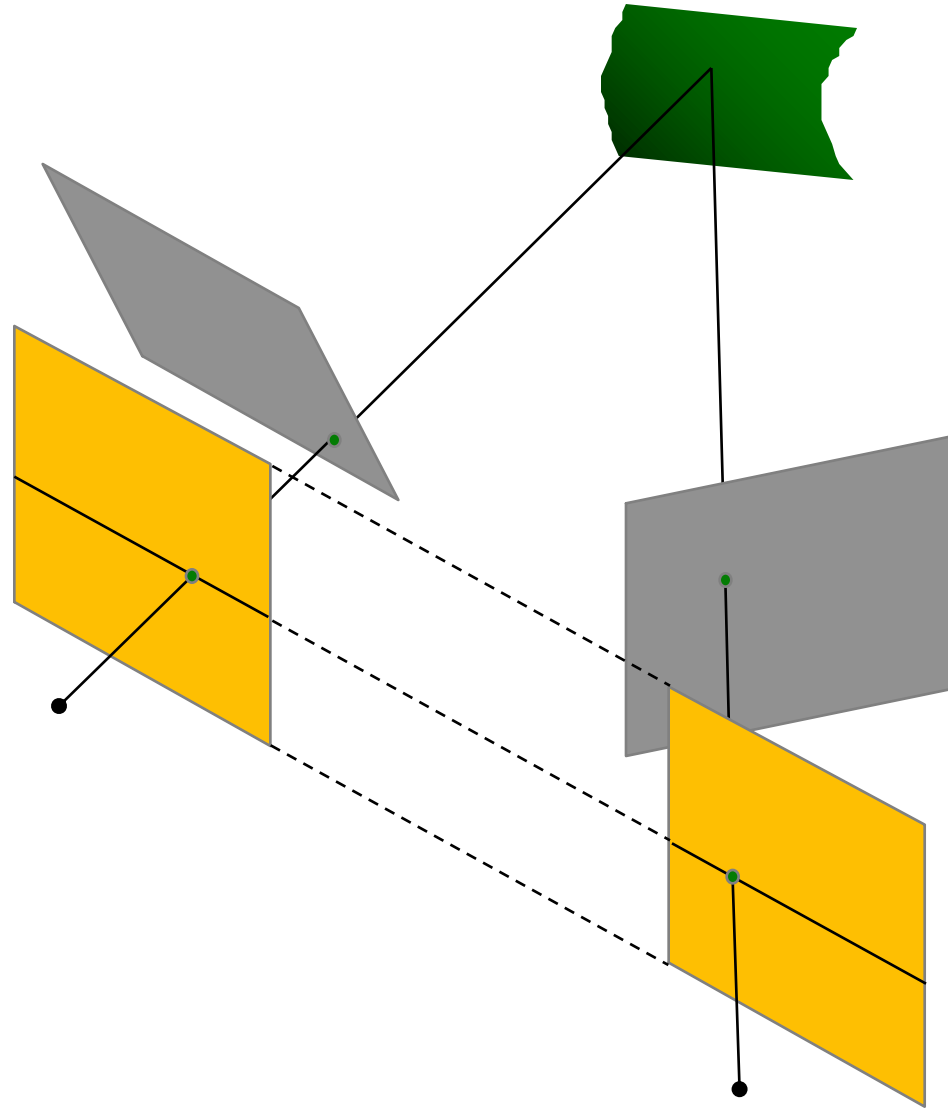
What is stereo rectification?



What is stereo rectification?

Reproject image planes
onto a common plane
parallel to the line
between camera centers

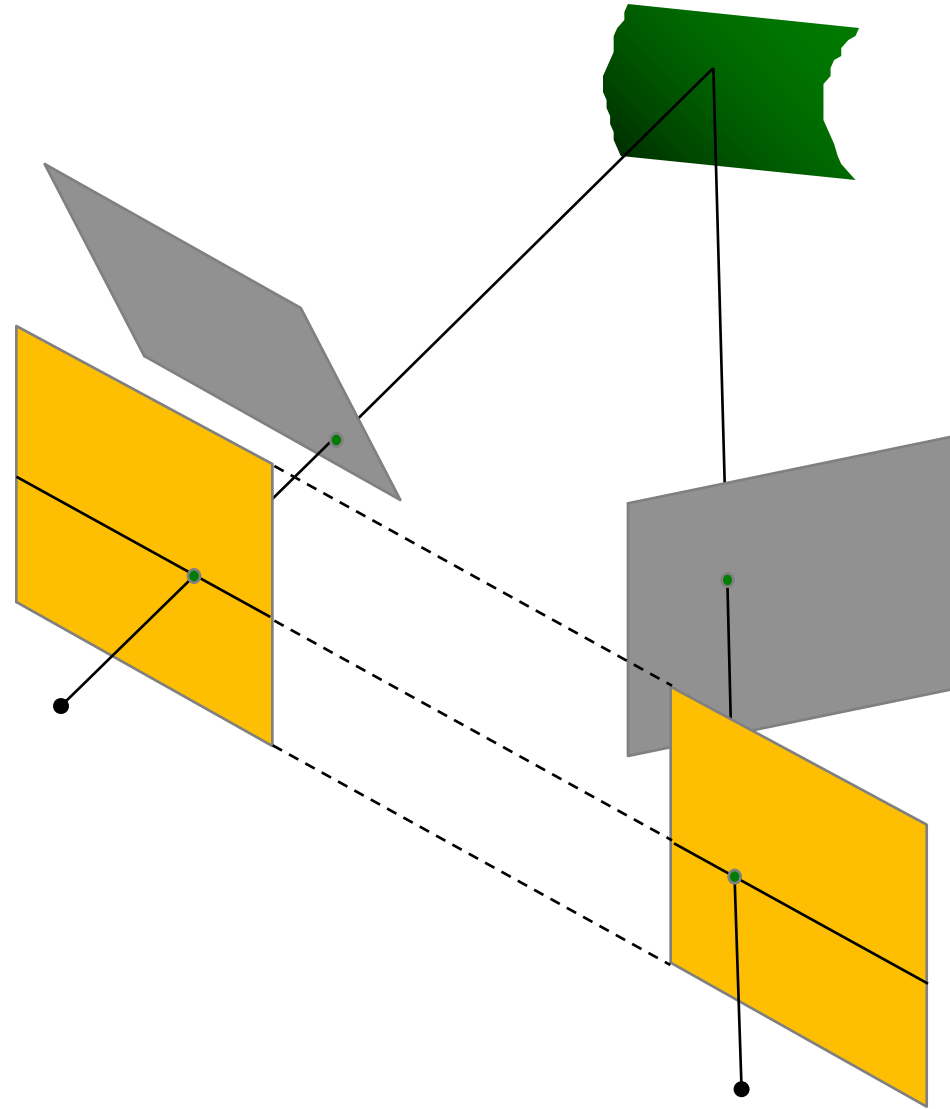
How can you do this?



What is stereo rectification?

Reproject image planes
onto a common plane
parallel to the line
between camera centers

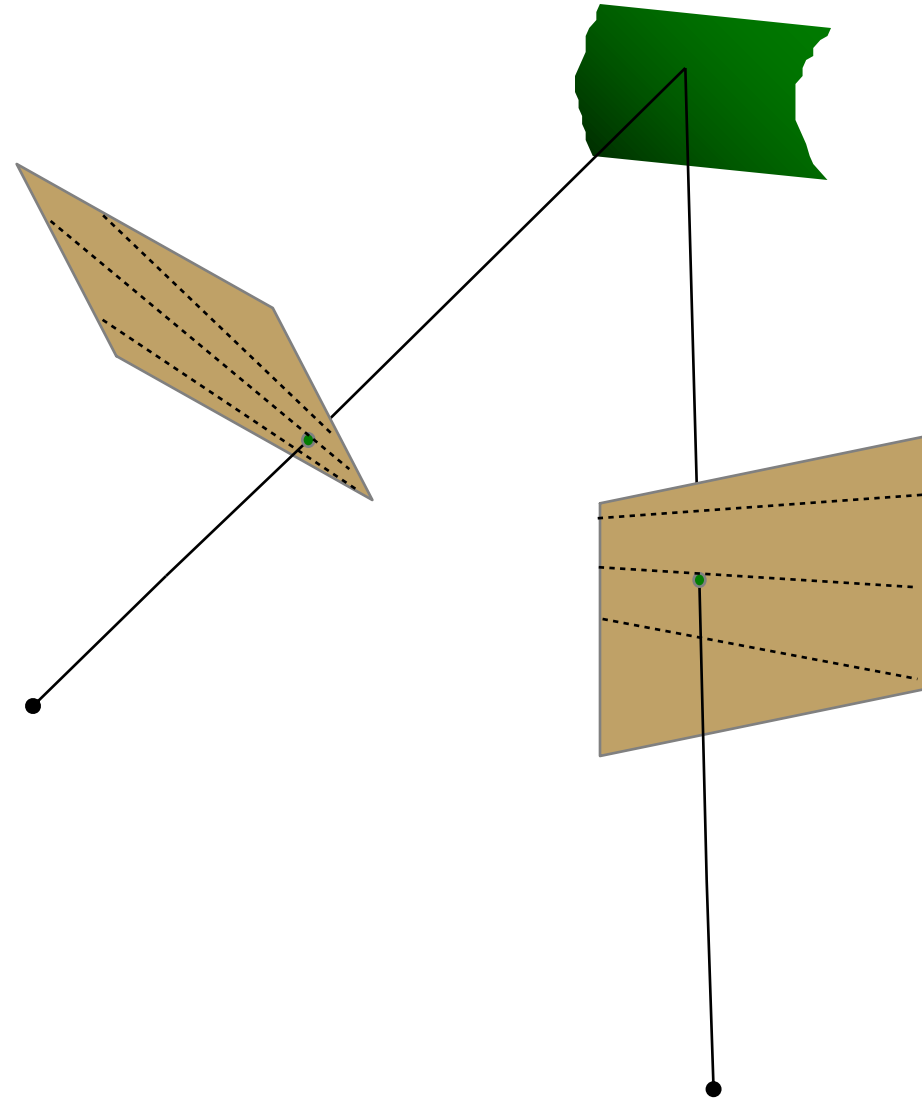
Need two
homographies (3x3
transform), one for each
input image reprojection



Stereo Rectification

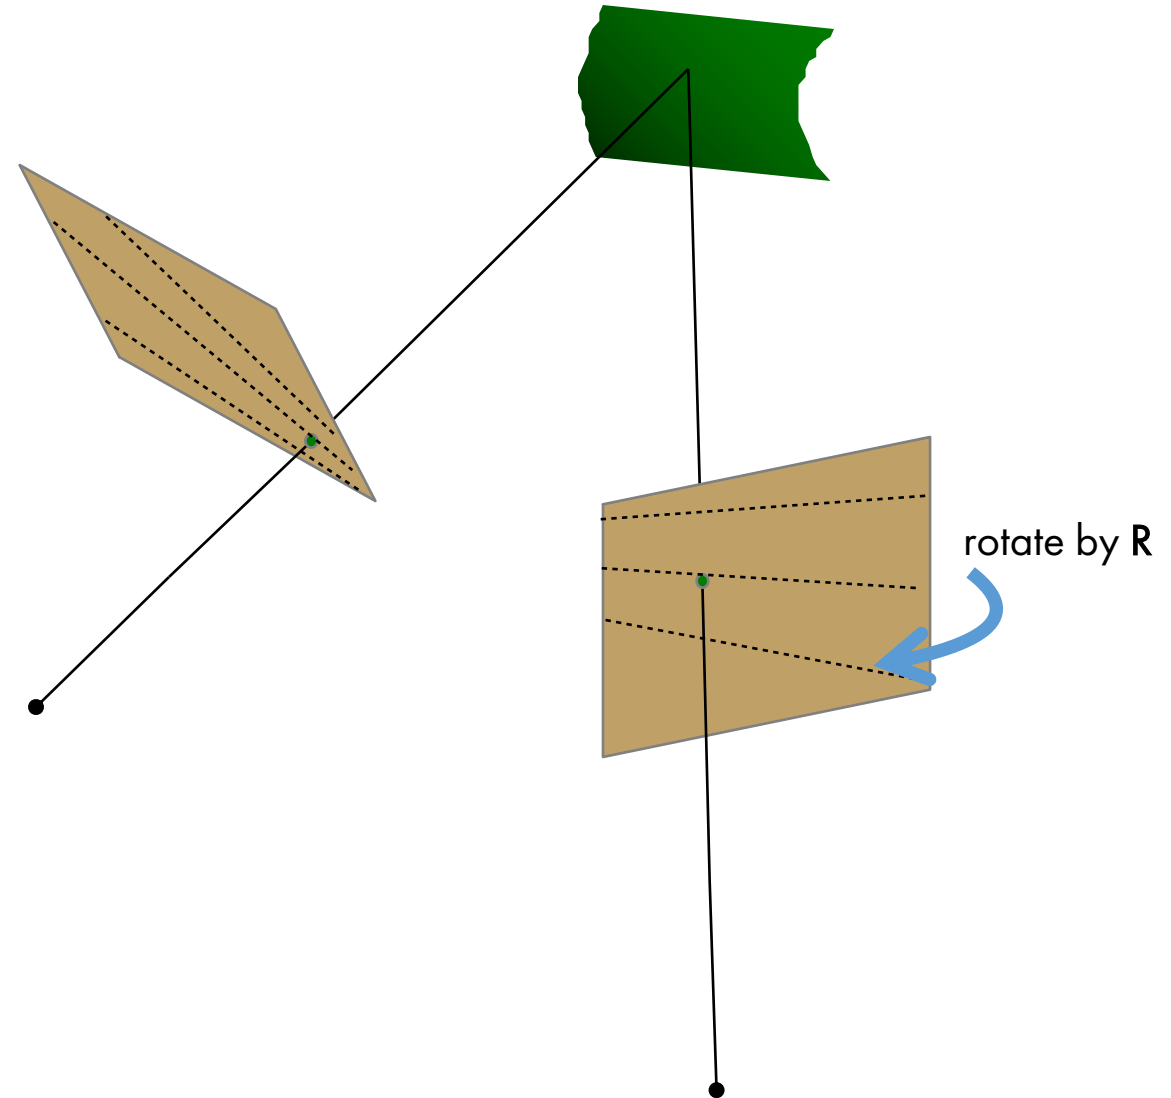
1. **Rotate** the right camera by **R**
(aligns camera coordinate system orientation only)
2. Rotate (**rectify**) the left camera so that the epipole is at infinity
3. Rotate (**rectify**) the right camera so that the epipole is at infinity
4. Adjust the **scale**

Stereo Rectification:



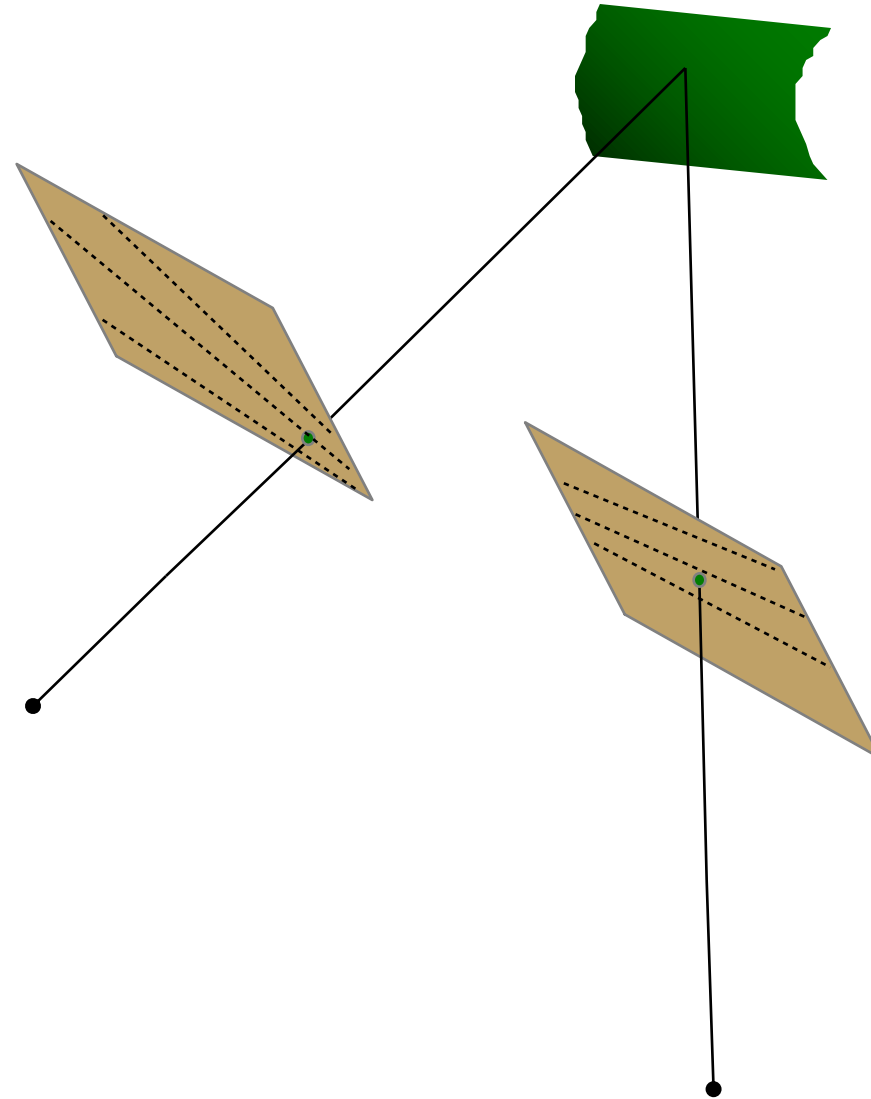
1. Compute \mathbf{E} to get \mathbf{R}
2. Rotate right image by \mathbf{R}
3. Rotate both images by \mathbf{R}_{rect}
4. Scale both images by \mathbf{H}

Stereo Rectification:



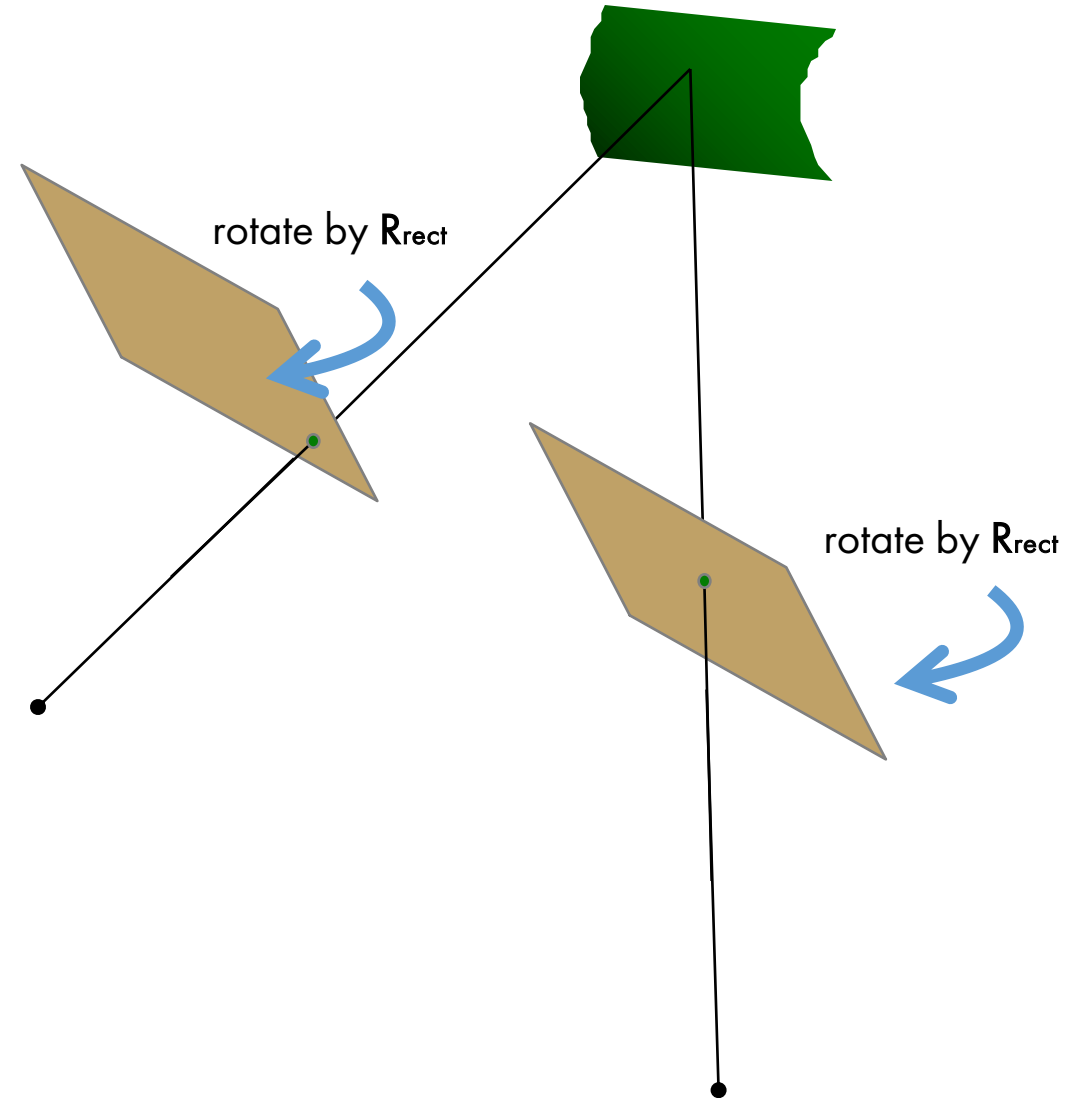
1. Compute \mathbf{E} to get \mathbf{R}
2. Rotate right image by \mathbf{R}
3. Rotate both images by \mathbf{R}_{rect}
4. Scale both images by \mathbf{H}

Stereo Rectification:



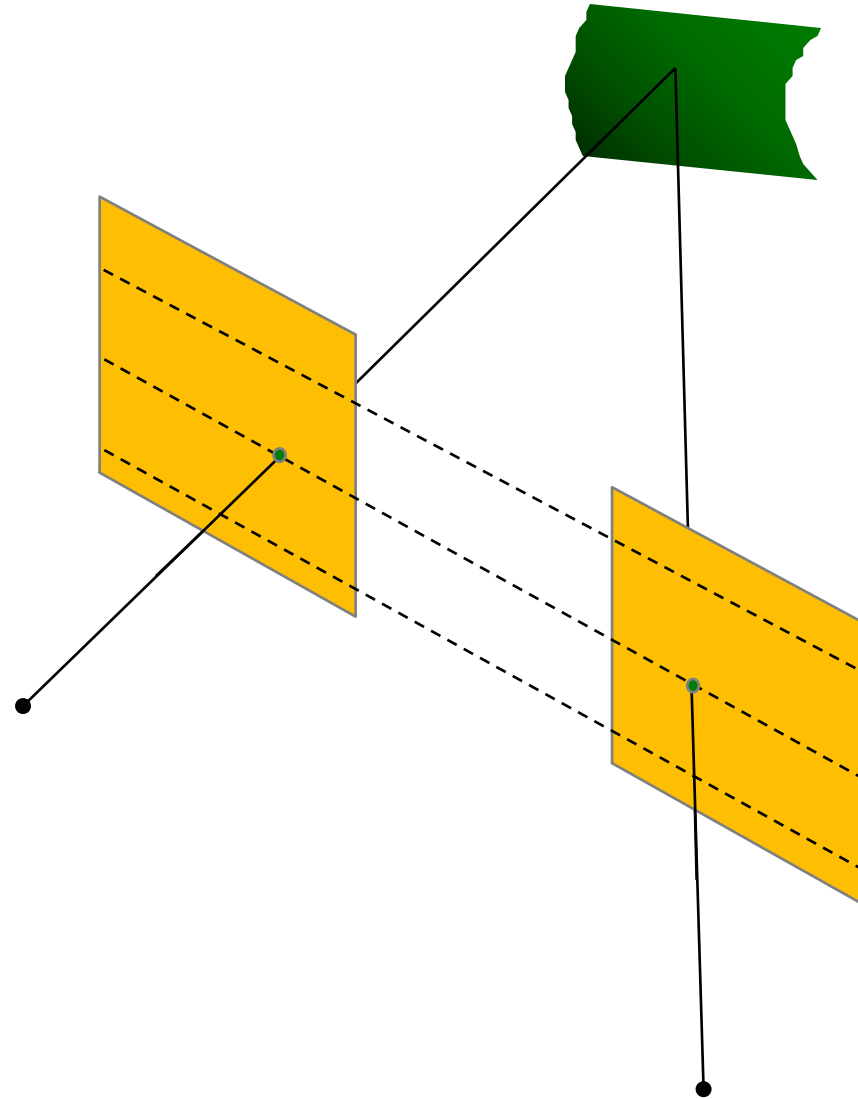
1. Compute \mathbf{E} to get \mathbf{R}
2. Rotate right image by \mathbf{R}
3. Rotate both images by \mathbf{R}_{rect}
4. Scale both images by \mathbf{H}

Stereo Rectification:



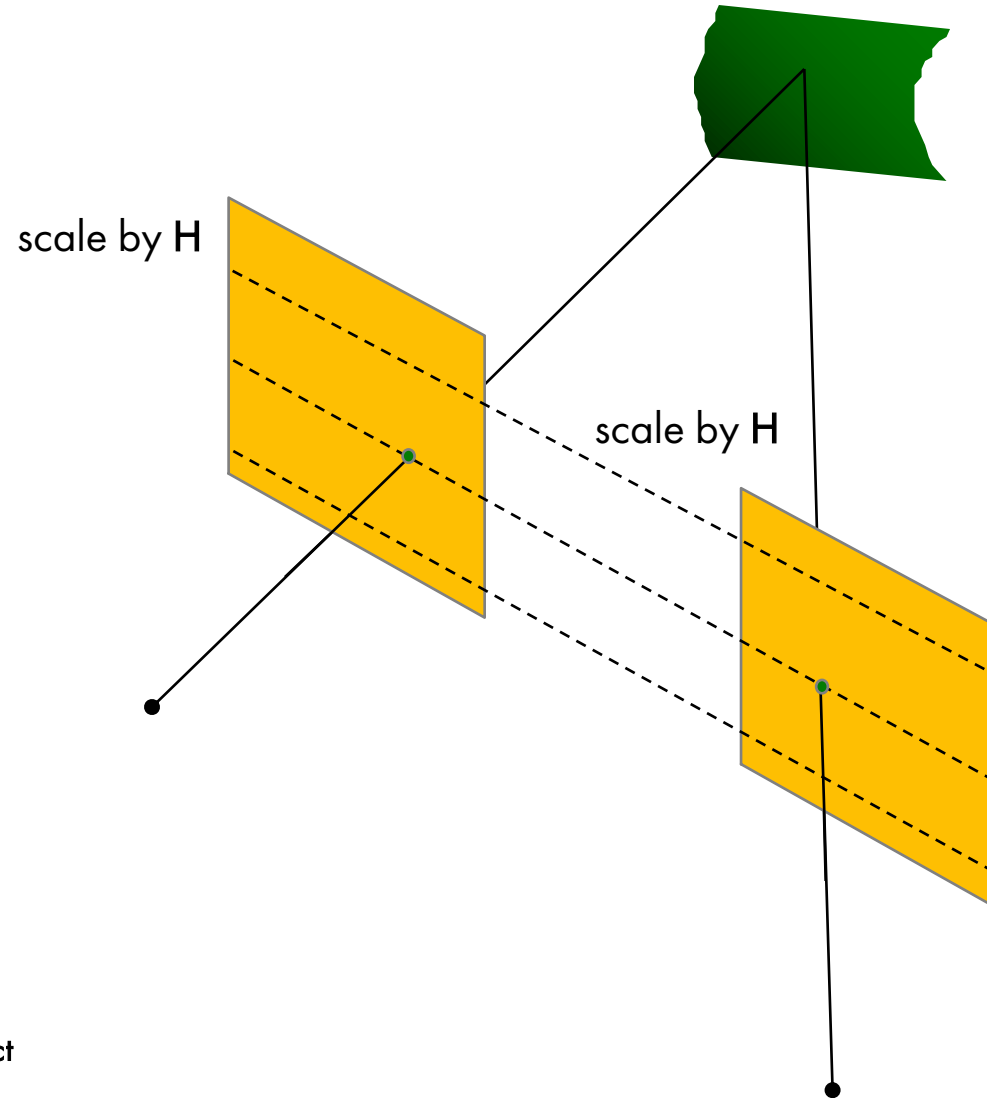
1. Compute E to get R
2. Rotate right image by R
3. Rotate both images by R_{rect}
4. Scale both images by H

Stereo Rectification:



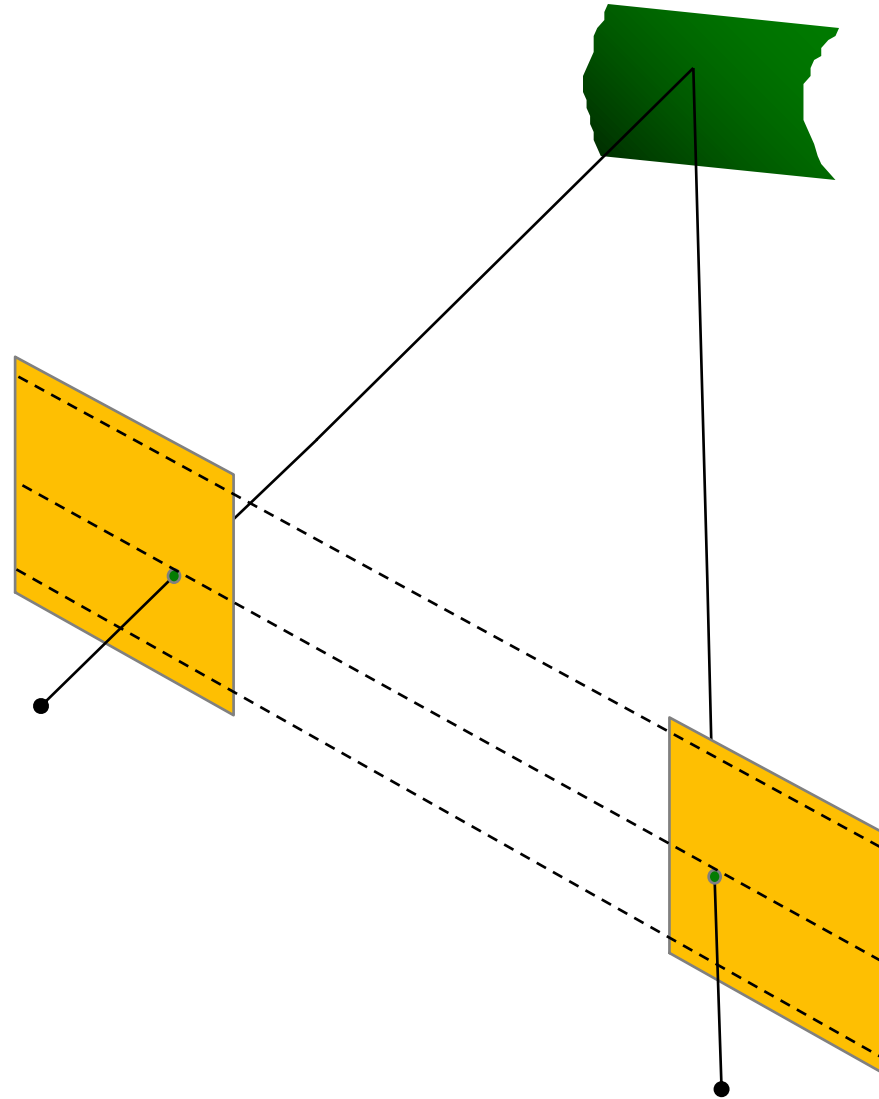
1. Compute \mathbf{E} to get \mathbf{R}
2. Rotate right image by \mathbf{R}
3. Rotate both images by \mathbf{R}_{rect}
4. Scale both images by \mathbf{H}

Stereo Rectification:



1. Compute \mathbf{E} to get \mathbf{R}
2. Rotate right image by \mathbf{R}
3. Rotate both images by \mathbf{R}_{rect}
4. Scale both images by H

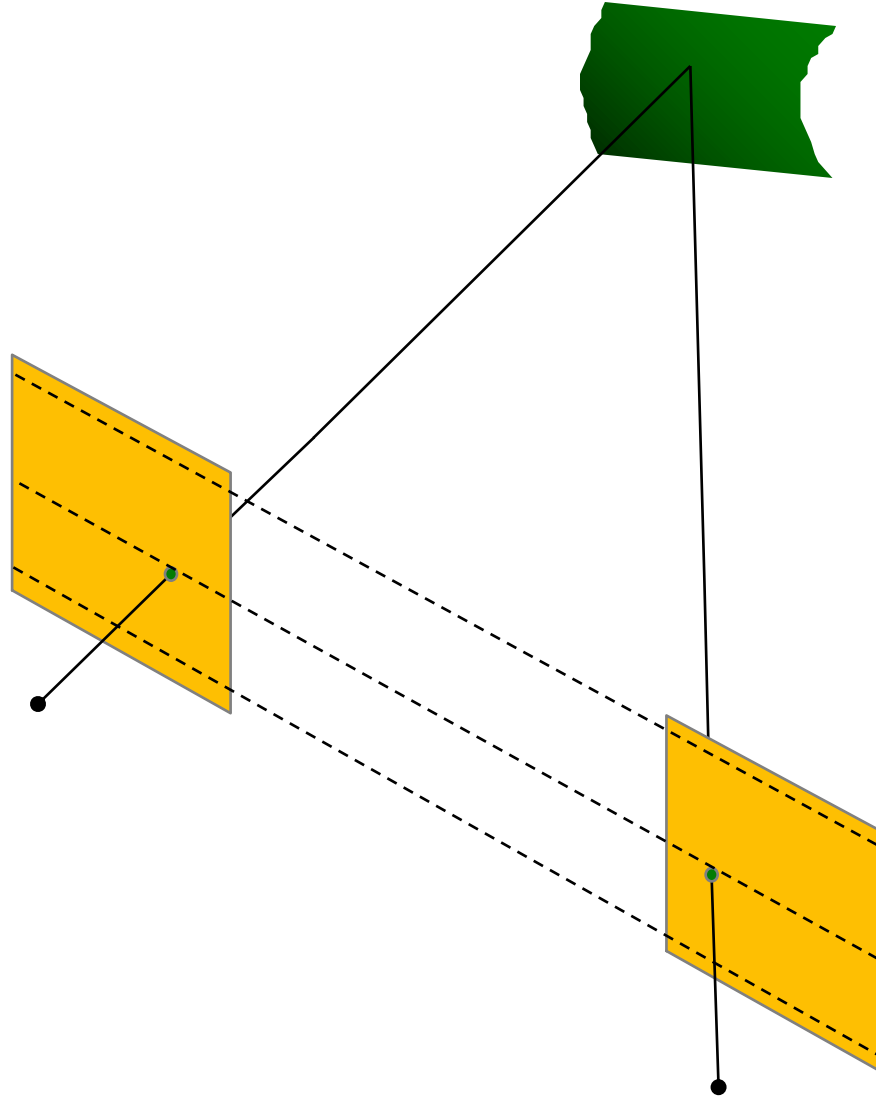
Stereo Rectification:



1. Compute \mathbf{E} to get \mathbf{R}
2. Rotate right image by \mathbf{R}
3. Rotate both images by \mathbf{R}_{rect}
4. Scale both images by \mathbf{H}

Stereo Rectification:

1. Compute E to get R
2. Rotate right image by R
3. Rotate both images by R_{rect}
4. Scale both images by H



Step 1: Compute \mathbf{E} to get \mathbf{R}

$$\text{SVD: } \mathbf{E} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \quad \text{Let } \mathbf{W} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

We get FOUR solutions:

$$\mathbf{E} = [\mathbf{R}|\mathbf{T}]$$

$$\mathbf{R}_1 = \mathbf{U}\mathbf{W}\mathbf{V}^\top \quad \mathbf{R}_2 = \mathbf{U}\mathbf{W}^\top\mathbf{V}^\top \quad \mathbf{T}_1 = U_3 \quad \mathbf{T}_2 = -U_3$$

two possible rotations

two possible translations

We get FOUR solutions:

$$\mathbf{R}_1 = \mathbf{U}\mathbf{W}\mathbf{V}^\top$$

$$\mathbf{T}_1 = U_3$$

$$\mathbf{R}_1 = \mathbf{U}\mathbf{W}\mathbf{V}^\top$$

$$\mathbf{T}_2 = -U_3$$

$$\mathbf{R}_2 = \mathbf{U}\mathbf{W}^\top \mathbf{V}^\top$$

$$\mathbf{T}_2 = -U_3$$

$$\mathbf{R}_2 = \mathbf{U}\mathbf{W}^\top \mathbf{V}^\top$$

$$\mathbf{T}_1 = U_3$$

Which one do we choose?

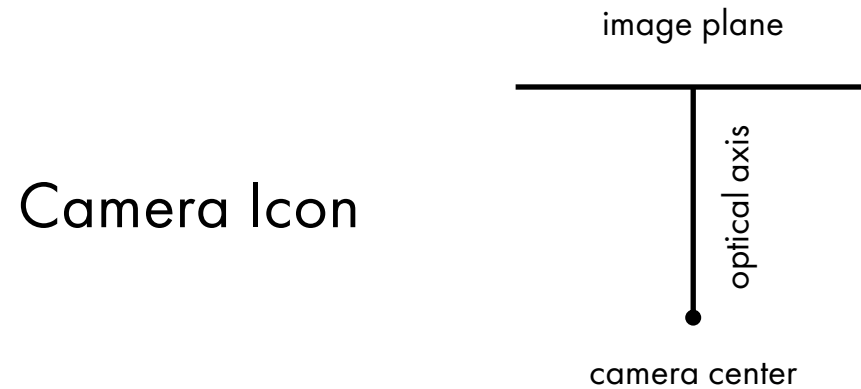
Compute determinant of R, valid solution must be equal to 1

(note: $\det(R) = -1$ means rotation and reflection)

Compute 3D point using triangulation, valid solution has positive Z value

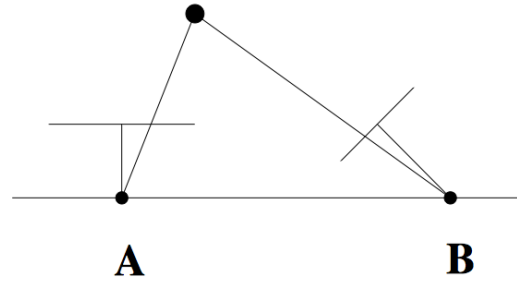
(Note: negative Z means point is behind the camera)

Let's visualize the four configurations...

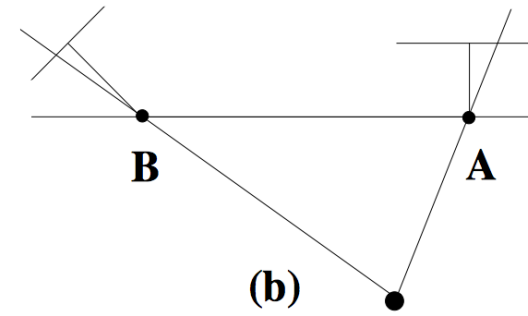


Find the configuration where the point is in front of both cameras

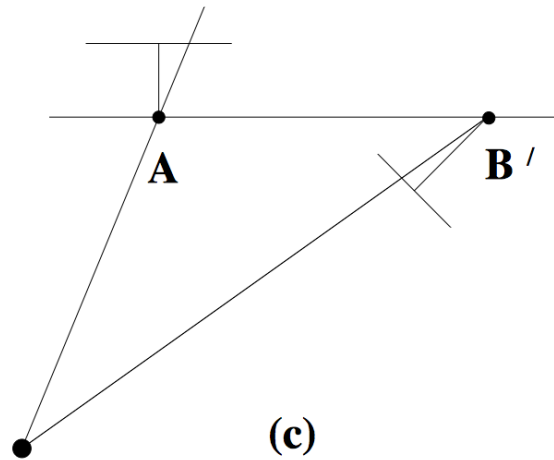
Find the configuration where the point is in front of both cameras



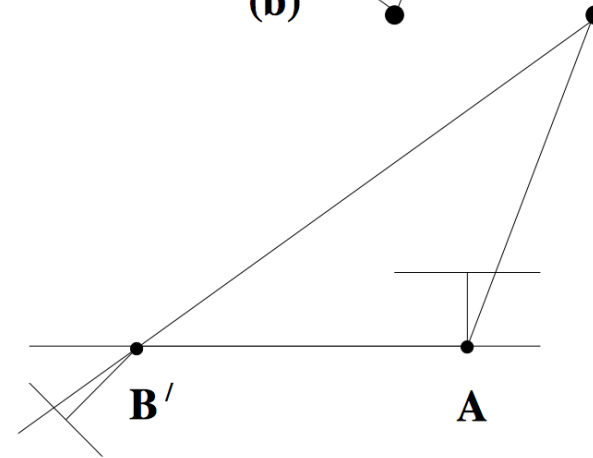
(a)



(b)

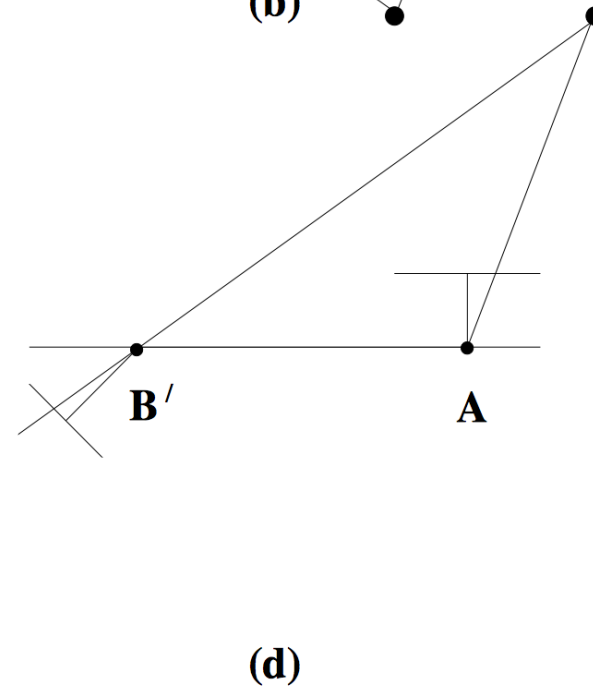
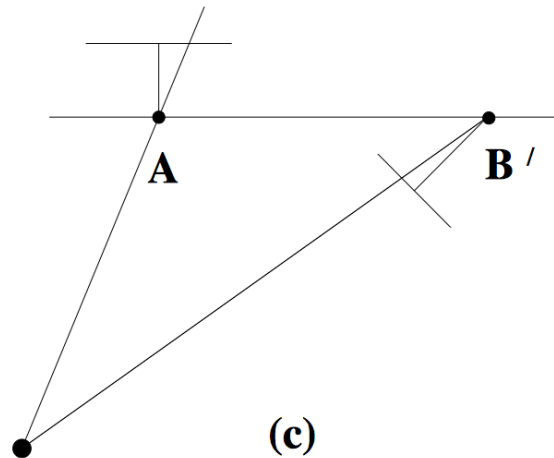
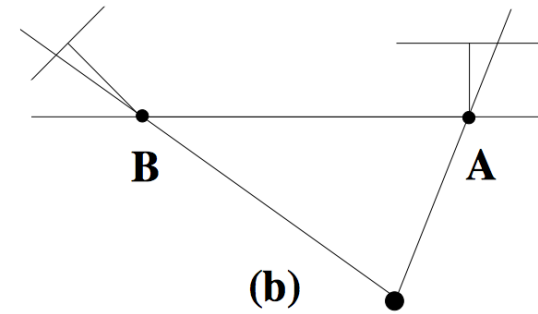
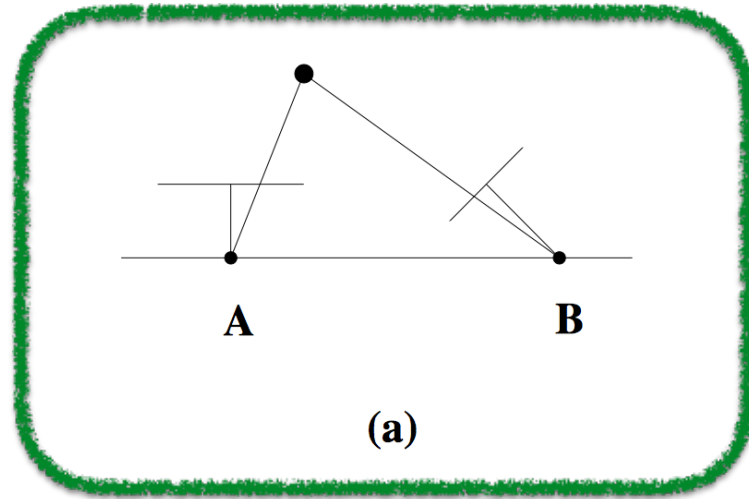


(c)

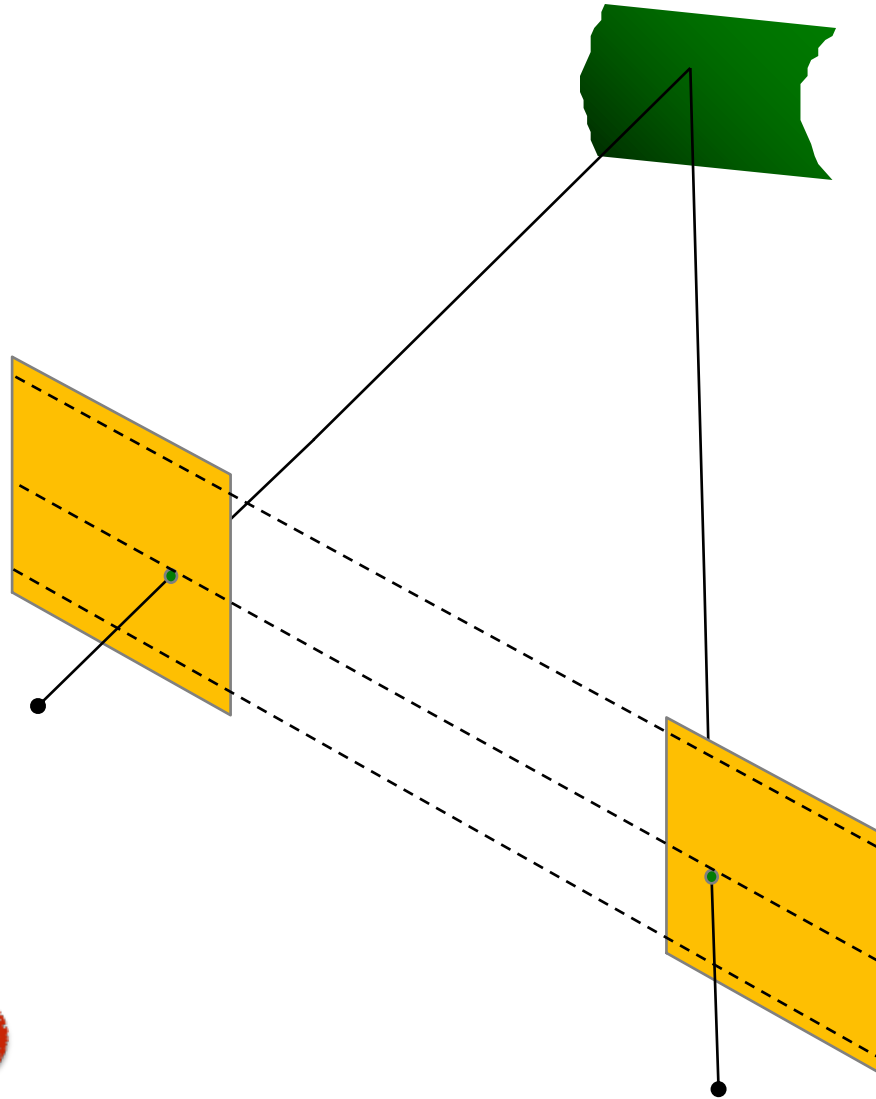


(d)

Find the configuration where the points is in front of both cameras



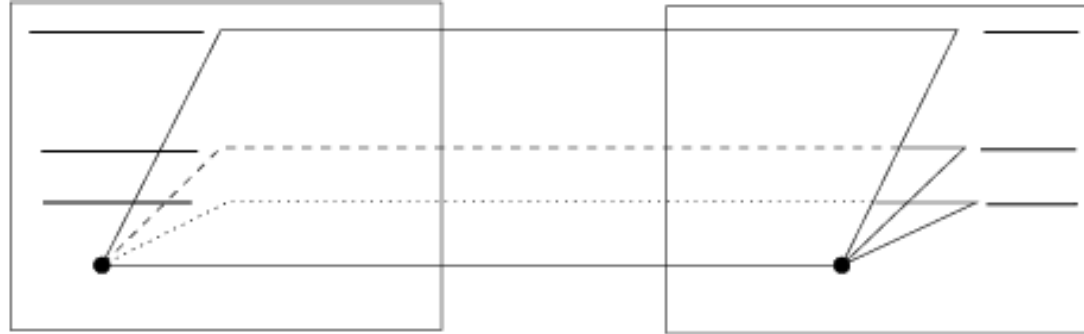
Stereo Rectification:



1. Compute E to get R
2. Rotate right image by R
3. Rotate both images by R_{rect}
4. Scale both images by H

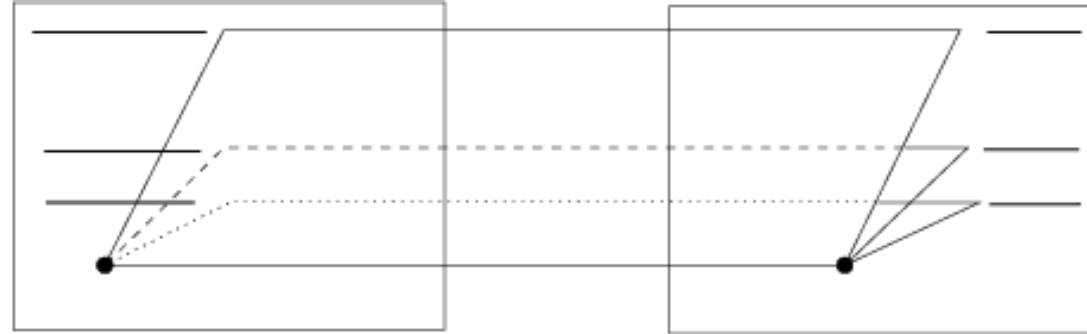
When do epipolar lines
become horizontal?

Parallel cameras



Where is the epipole?

Parallel cameras



epipole at infinity

Setting the epipole to infinity

(Building R_{rect} from E)

$$\text{Let } R_{\text{rect}} = \begin{bmatrix} \mathbf{r}_1^\top \\ \mathbf{r}_2^\top \\ \mathbf{r}_3^\top \end{bmatrix} \quad \text{Given: } \begin{array}{l} \text{epipole } \mathbf{e} \\ \text{(using SVD on } E\text{)} \\ \text{(translation from } E\text{)} \end{array}$$

$$\mathbf{r}_1 = \mathbf{e}_1 = \frac{T}{\|T\|} \quad \text{epipole coincides with translation vector}$$

$$\mathbf{r}_2 = \frac{1}{\sqrt{T_x^2 + T_y^2}} \begin{bmatrix} -T_y & T_x & 0 \end{bmatrix} \quad \begin{array}{l} \text{cross product of } \mathbf{e} \text{ and the} \\ \text{direction vector of the} \\ \text{optical axis} \end{array}$$

$$\mathbf{r}_3 = \mathbf{r}_1 \times \mathbf{r}_2 \quad \text{orthogonal vector}$$

If $\mathbf{r}_1 = \mathbf{e}_1 = \frac{\mathbf{T}}{\|\mathbf{T}\|}$ and $\mathbf{r}_2, \mathbf{r}_3$ orthogonal

then $R_{\text{rect}} \mathbf{e}_1 = \begin{bmatrix} \mathbf{r}_1^\top \mathbf{e}_1 \\ \mathbf{r}_2^\top \mathbf{e}_1 \\ \mathbf{r}_3^\top \mathbf{e}_1 \end{bmatrix} = \begin{bmatrix} ? \\ ? \\ ? \end{bmatrix}$

If $\mathbf{r}_1 = \mathbf{e}_1 = \frac{T}{||T||}$ and \mathbf{r}_2 \mathbf{r}_3 orthogonal

then $R_{\text{rect}} \mathbf{e}_1 = \begin{bmatrix} \mathbf{r}_1^\top \mathbf{e}_1 \\ \mathbf{r}_2^\top \mathbf{e}_1 \\ \mathbf{r}_3^\top \mathbf{e}_1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$

Where is this point located on the image plane?

If $\mathbf{r}_1 = \mathbf{e}_1 = \frac{T}{||T||}$ and \mathbf{r}_2 \mathbf{r}_3 orthogonal

then $R_{\text{rect}} \mathbf{e}_1 = \begin{bmatrix} \mathbf{r}_1^\top \mathbf{e}_1 \\ \mathbf{r}_2^\top \mathbf{e}_1 \\ \mathbf{r}_3^\top \mathbf{e}_1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$

Where is this point located on the image plane?

At x-infinity

Stereo Rectification Algorithm

1. Estimate \mathbf{E} using the 8 point algorithm (SVD)
2. Estimate the epipole \mathbf{e} (SVD of \mathbf{E})
3. Build \mathbf{R}_{rect} from \mathbf{e}
4. Decompose \mathbf{E} into \mathbf{R} and \mathbf{T}
5. Set $\mathbf{R}_1 = \mathbf{R}_{\text{rect}}$ and $\mathbf{R}_2 = \mathbf{R}\mathbf{R}_{\text{rect}}$
6. Rotate each left camera point (warp image)
$$[x' \ y' \ z'] = \mathbf{R}_1 [x \ y \ z]$$
7. Rectified points as $\mathbf{p} = f/z' [x' \ y' \ z']$
8. Repeat 6 and 7 for right camera points using \mathbf{R}_2



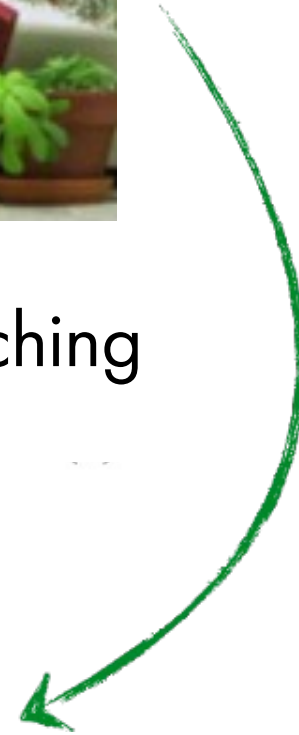
What can we do after rectification?

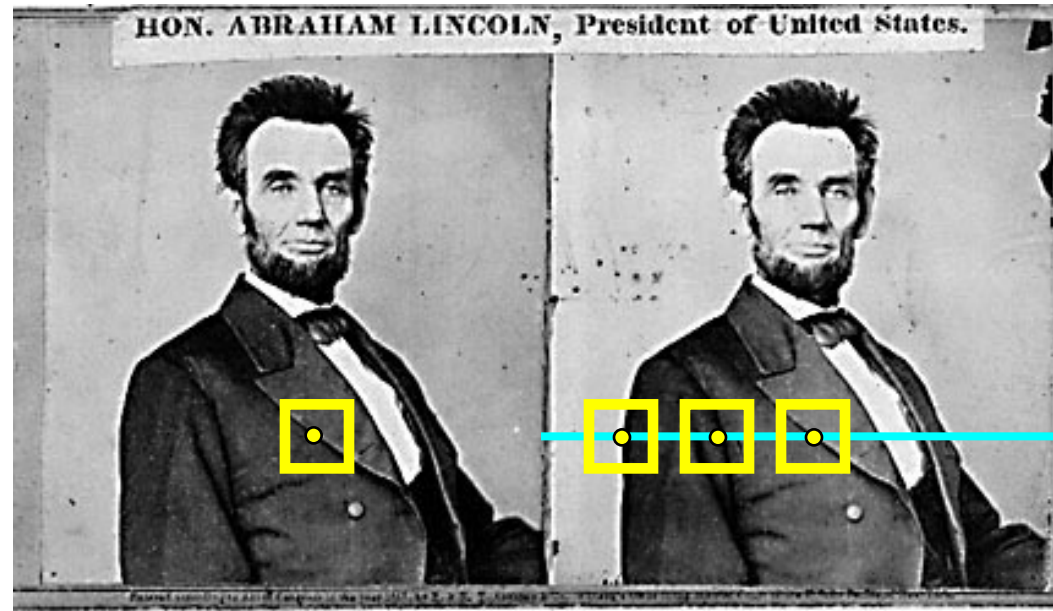


Stereo matching



Depth Estimation via Stereo Matching





1. Rectify images
(make epipolar lines horizontal)
2. For each pixel
 - a. Find epipolar line
 - b. Scan line for best match
 - c. Compute depth from disparity

How would
you do this?

$$Z = \frac{bf}{d}$$

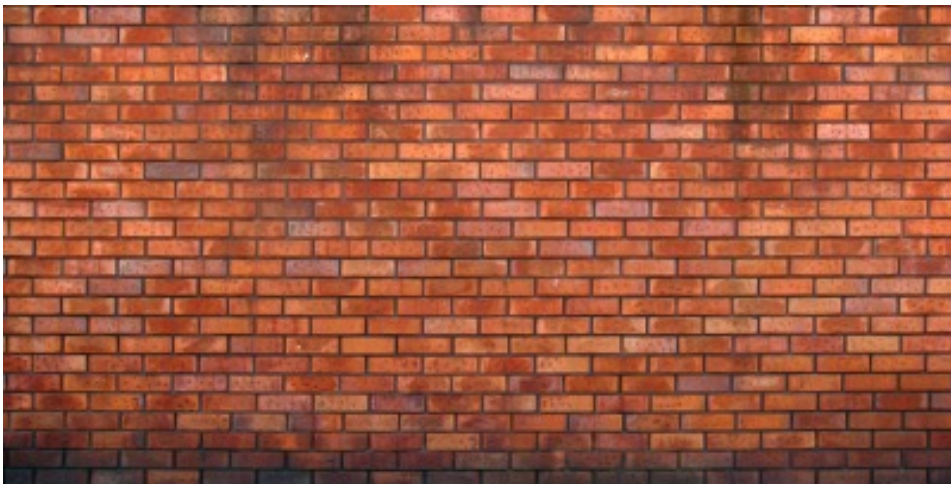
Reminder from filtering

How do we detect an edge?

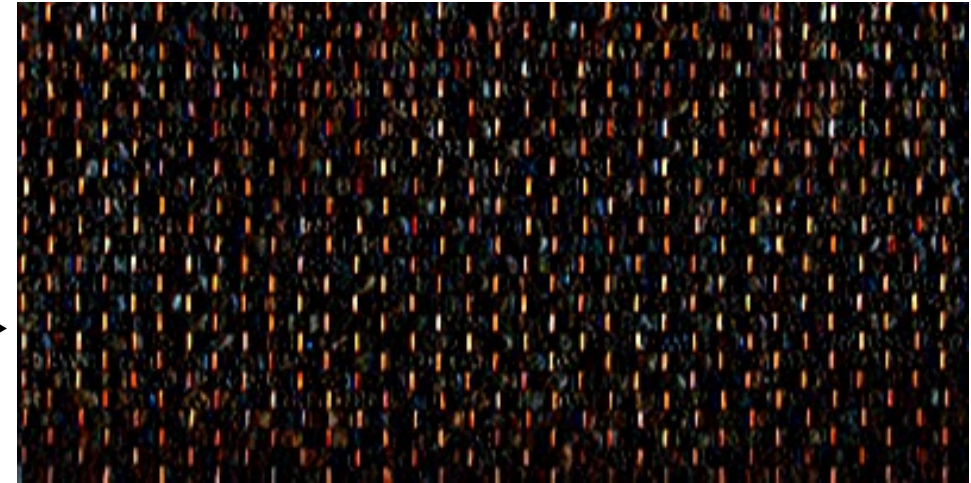
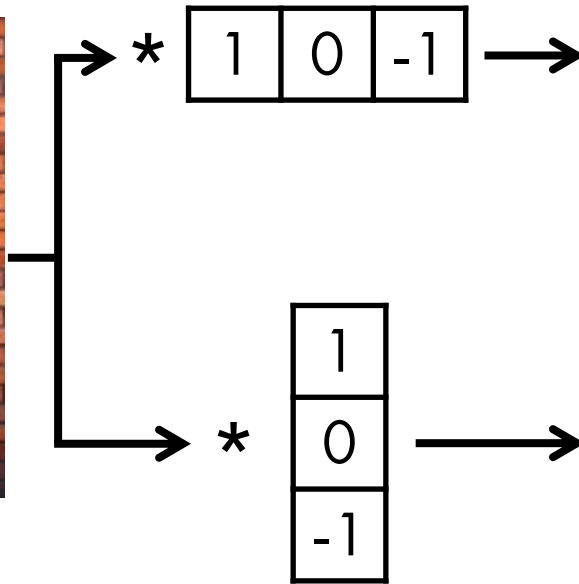
Reminder from filtering

How do we detect an edge?

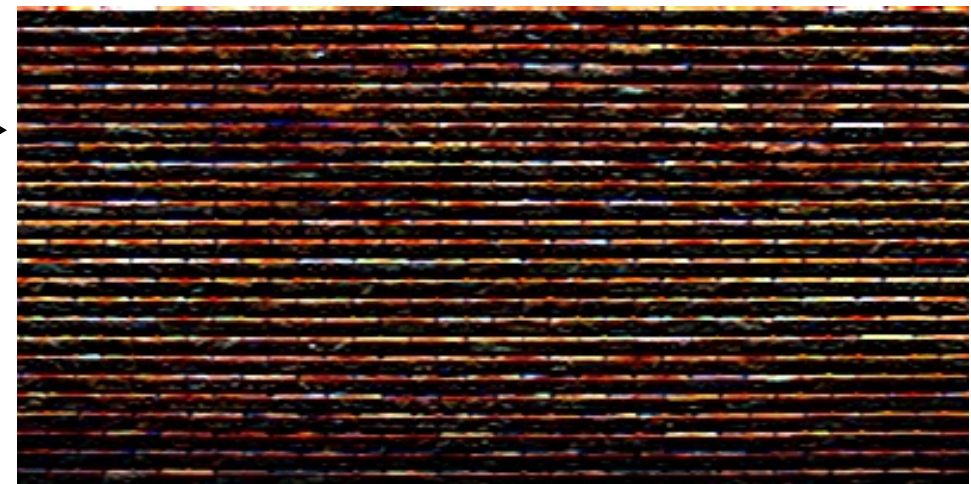
- We filter with something that looks like an edge.



original



horizontal edge filter

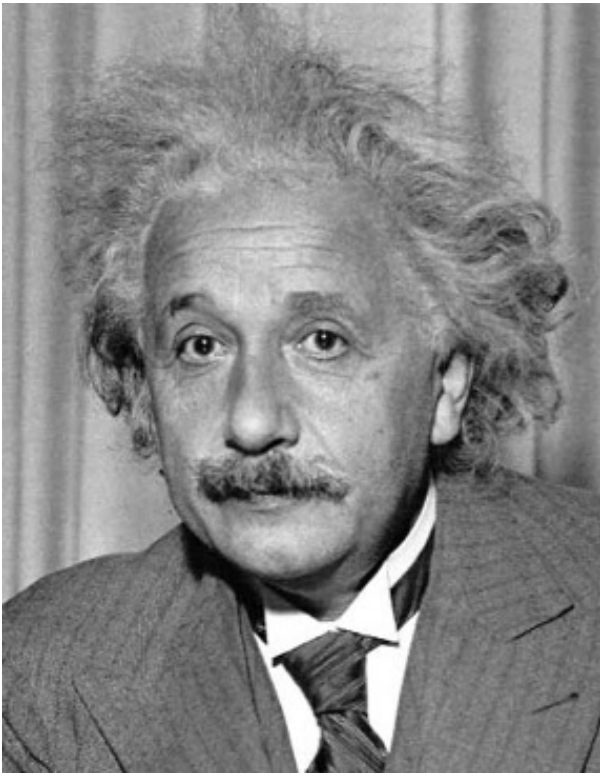


vertical edge filter

We can think of linear filtering as a way to evaluate how similar an image is locally to some template.

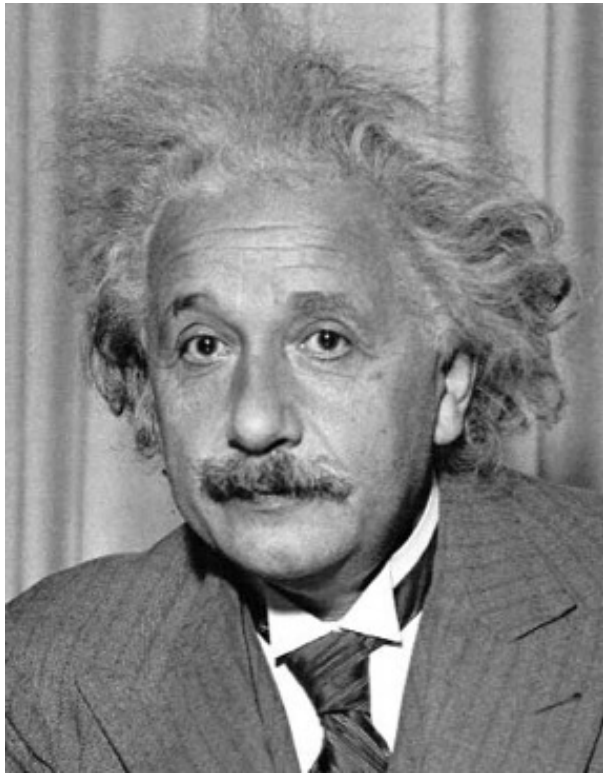
Find this template

How do we detect the template n the following image?




Find this template

How do we detect the template  in the following image?

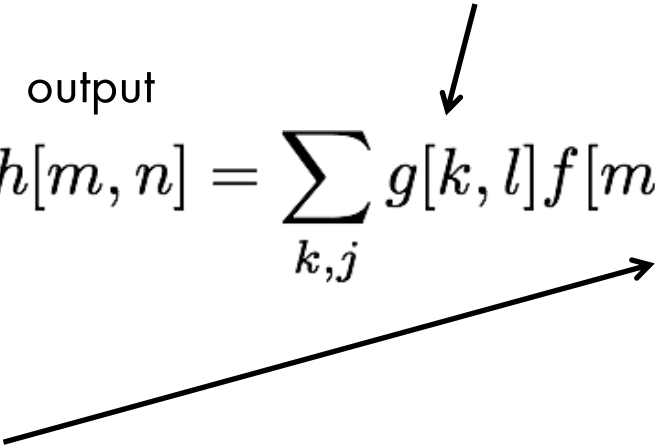


output

$$h[m, n] = \sum_{k, l} g[k, l] f[m + k, n + l]$$

filter 

image

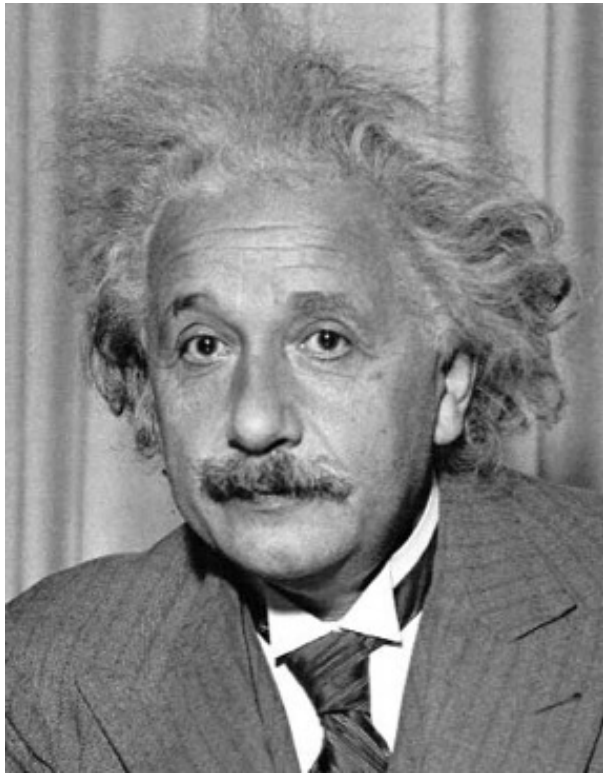


What will the output look like?

Solution 1: Filter the image using the template as filter kernel.


Find this template

How do we detect the template  in the following image?

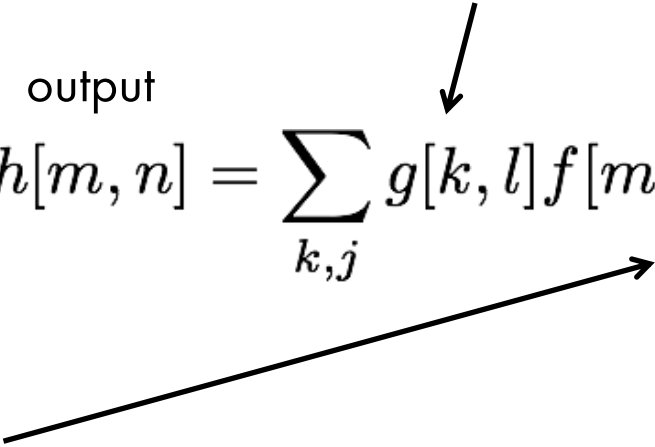


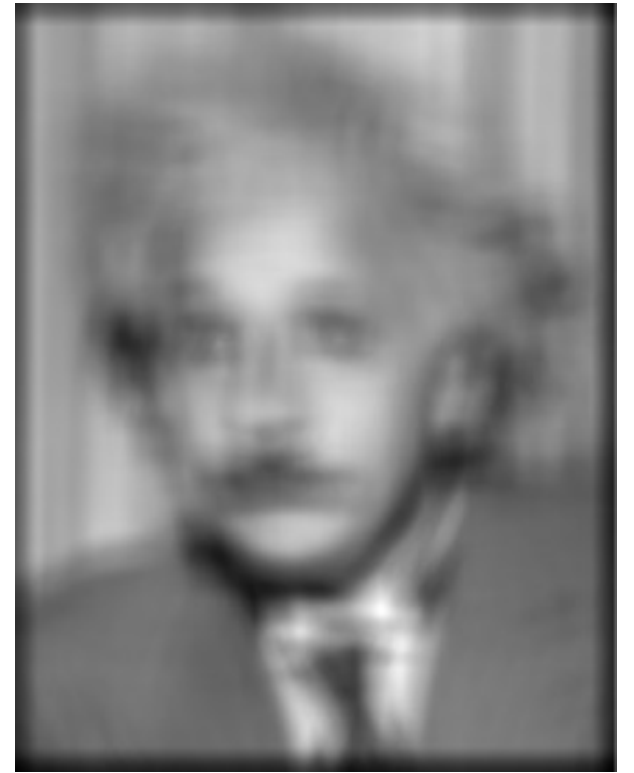
output

$$h[m, n] = \sum_{k, j} g[k, l] f[m + k, n + l]$$

filter 

image



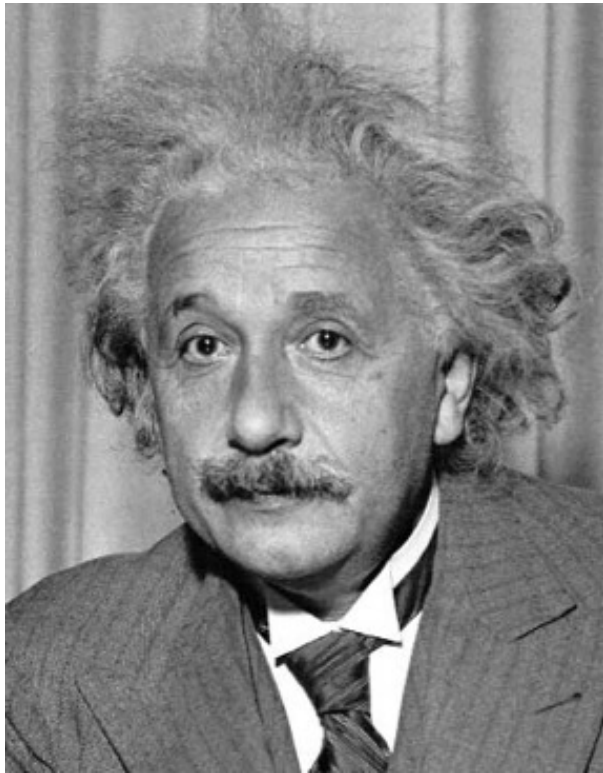


Solution 1: Filter the image using the template as filter kernel.

What went wrong?


Find this template

How do we detect the template  in the following image?

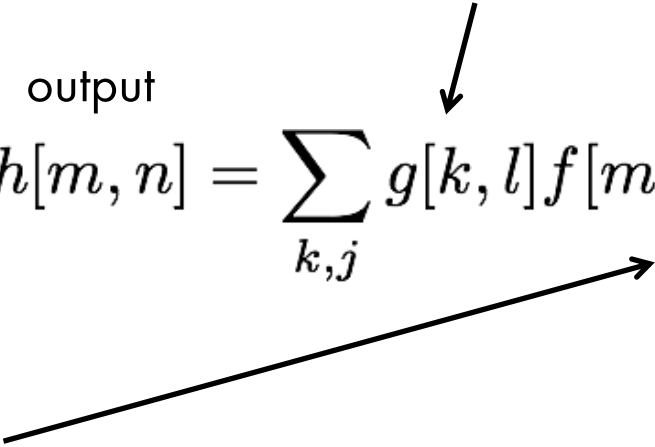


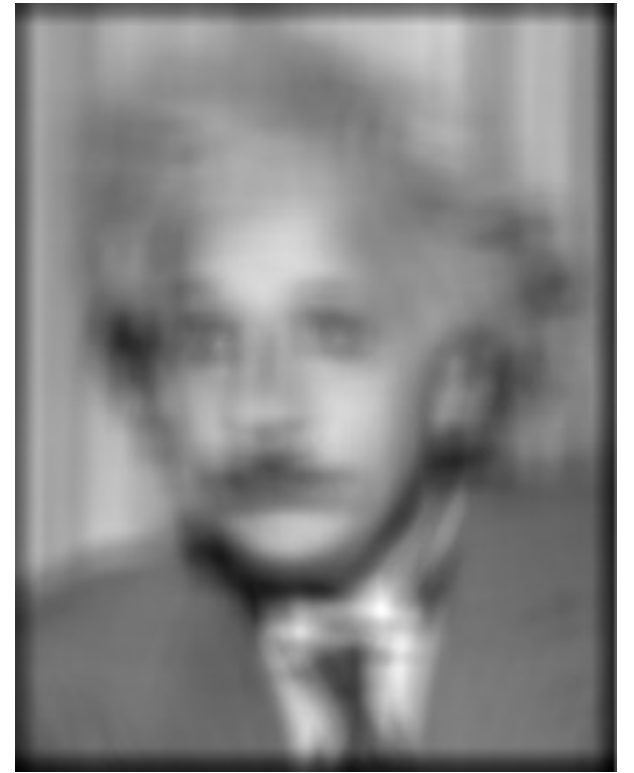
output

$$h[m, n] = \sum_{k, l} g[k, l] f[m + k, n + l]$$

filter 

image



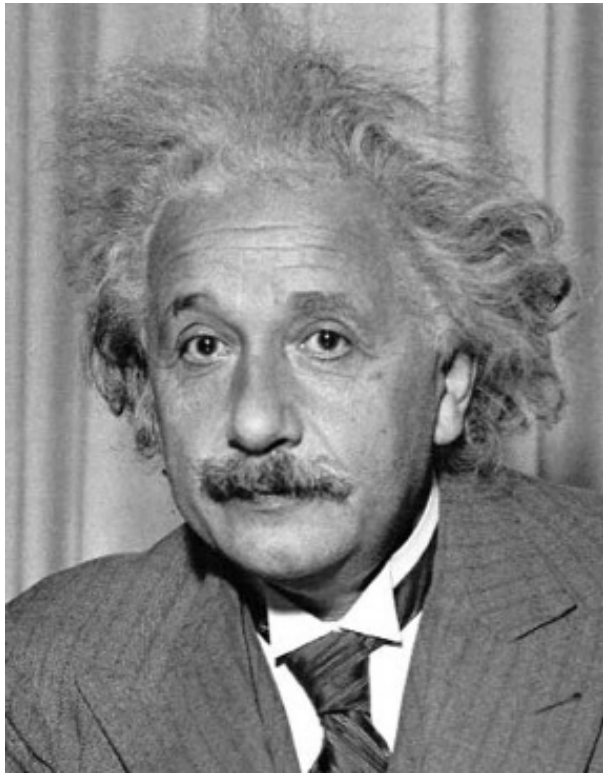


Solution 1: Filter the image using the template as filter kernel.

Increases for higher local intensities.


Find this template

How do we detect the template  in the following image?



output

$$h[m, n] = \sum_{k, l} (g[k, l] - \bar{g}) f[m + k, n + l]$$

filter  template mean

image

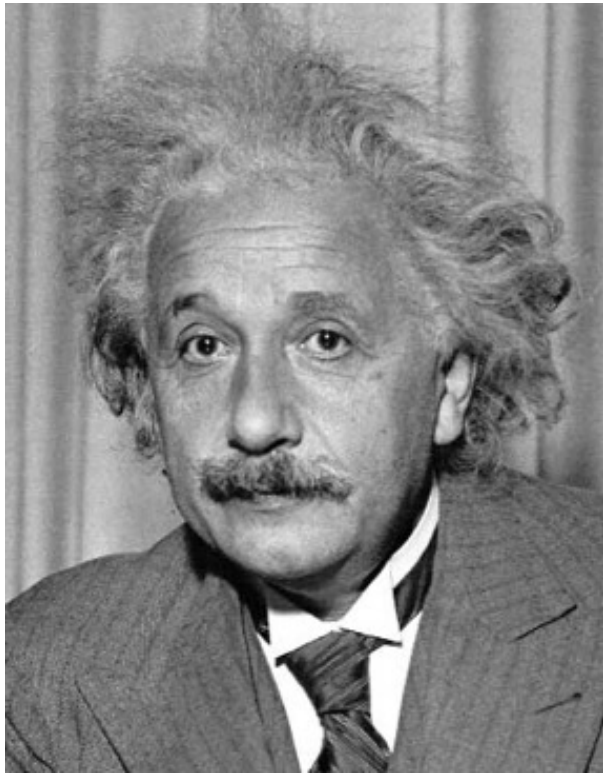
The diagram shows the equation for template matching. An arrow points from the word 'filter' to the term $g[k, l]$. Another arrow points from the words 'template mean' to the term \bar{g} . A third arrow points from the word 'image' to the term $f[m + k, n + l]$.

What will the output look like?


Solution 2: Filter the image using a zero-mean template.

Find this template

How do we detect the template  in the following image?



output

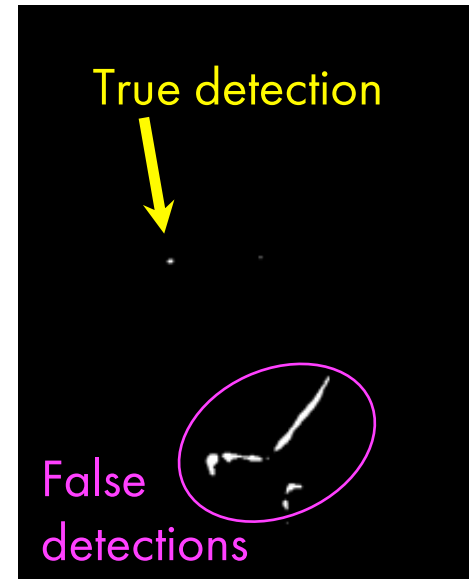
filter  template mean

$$h[m, n] = \sum_{k, l} (g[k, l] - \bar{g}) f[m + k, n + l]$$

image

thresholding

output

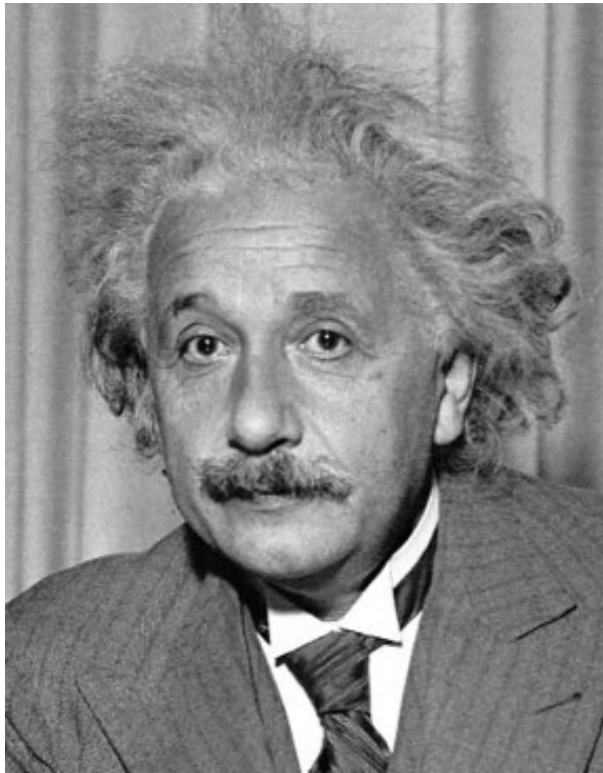



Solution 2: Filter the image using a zero-mean template.

What went wrong?

Find this template

How do we detect the template  in the following image?



filter  template mean

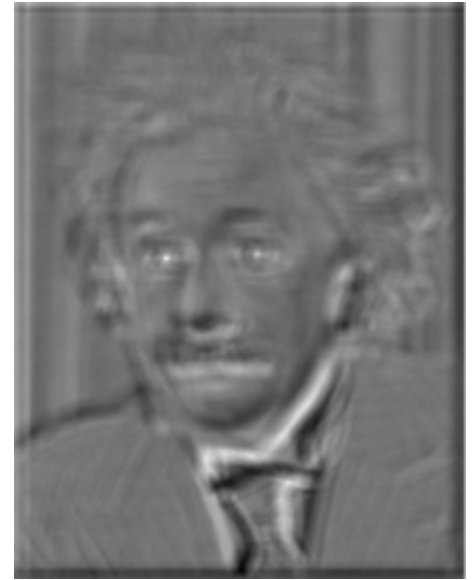
output

$$h[m, n] = \sum_{k, l} (g[k, l] - \bar{g}) f[m + k, n + l]$$

image

The diagram shows the equation with arrows indicating the components: 'filter' points to $g[k, l]$, 'template mean' points to \bar{g} , 'output' points to $h[m, n]$, and 'image' points to $f[m + k, n + l]$.

output

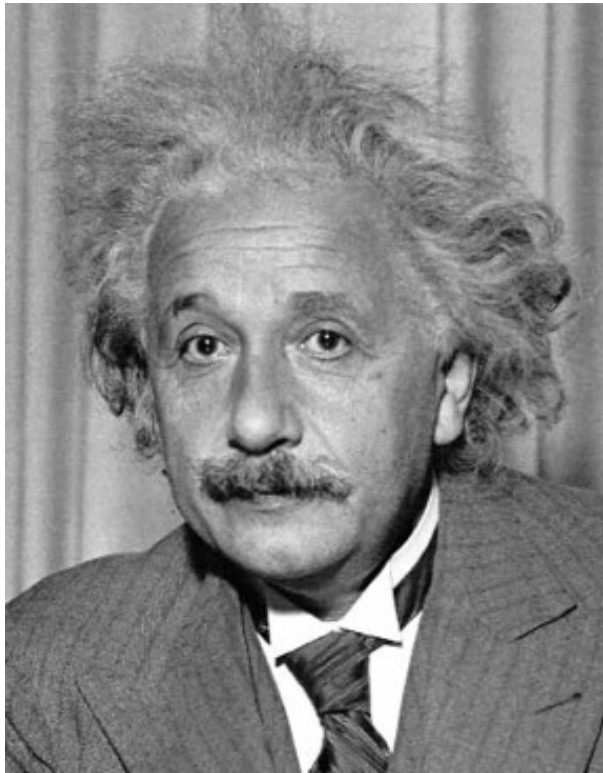


Not robust to high-contrast areas


Solution 2: Filter the image using a zero-mean template.

Find this template

How do we detect the template  in the following image?

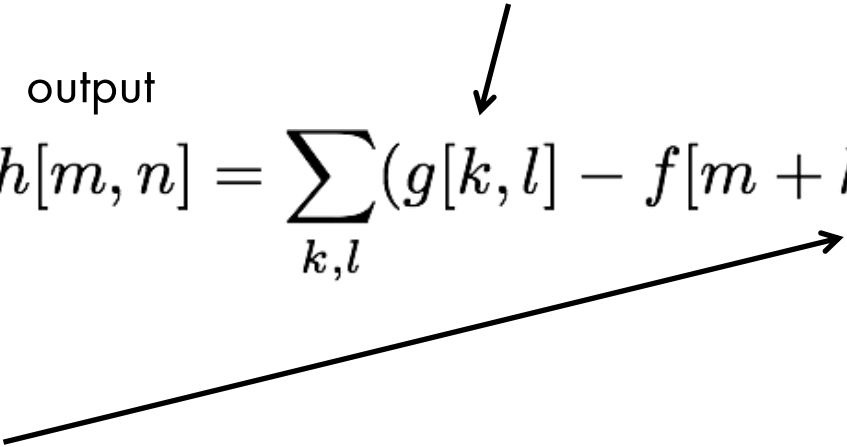


output

filter 

$$h[m, n] = \sum_{k, l} (g[k, l] - f[m + k, n + l])^2$$

image

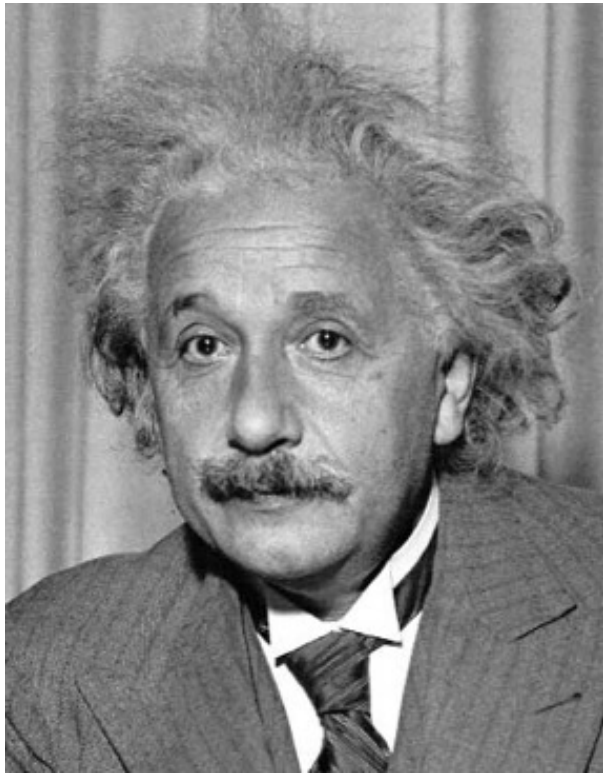


What will
the output
look like?


Solution 3: Use sum of squared differences (SSD).

Find this template

How do we detect the template  in the following image?



output

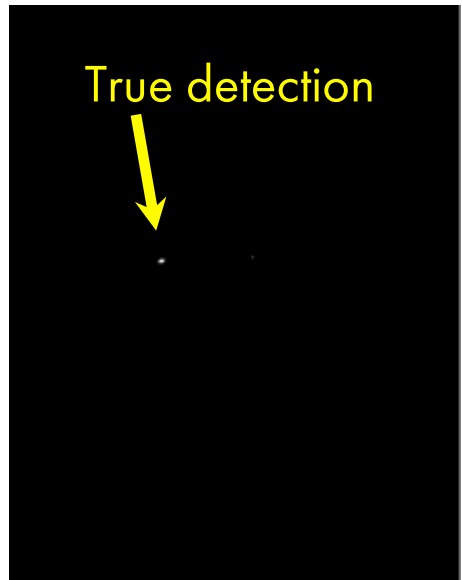
filter 

$$h[m, n] = \sum_{k, l} (g[k, l] - f[m + k, n + l])^2$$

image

thresholding

1-output

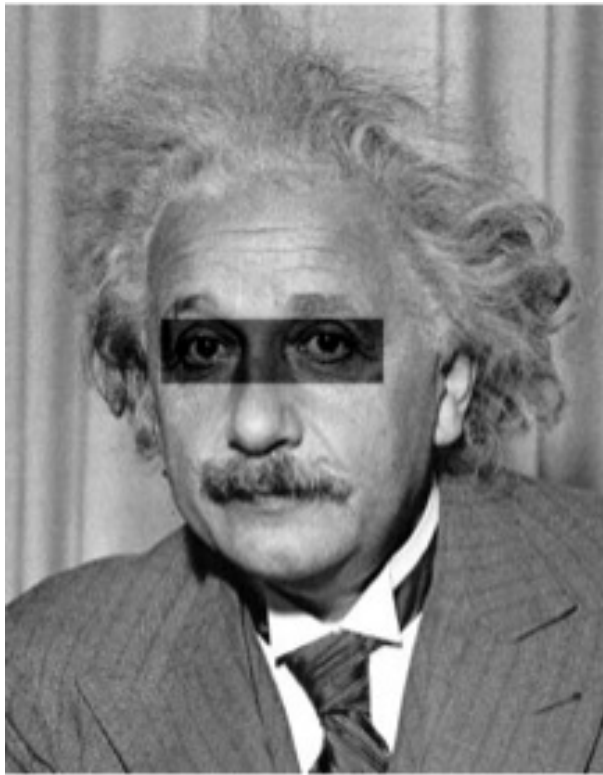


Solution 3: Use sum of squared differences (SSD).

What could go wrong?


Find this template

How do we detect the template  in the following image?



1-output



filter 

output

$$h[m, n] = \sum_{k, l} (g[k, l] - f[m + k, n + l])^2$$

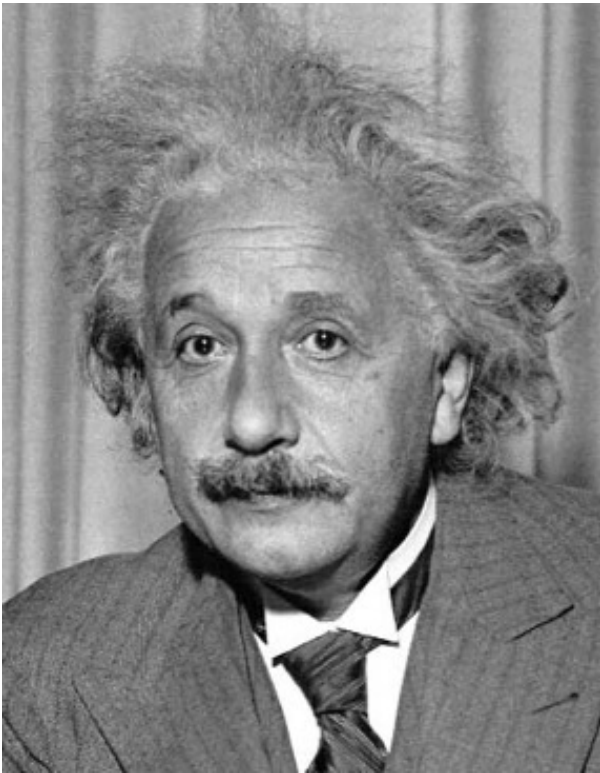
image

Not robust to local
intensity changes

Solution 3: Use sum of squared differences (SSD).

Find this template

How do we detect the template  in the following image?



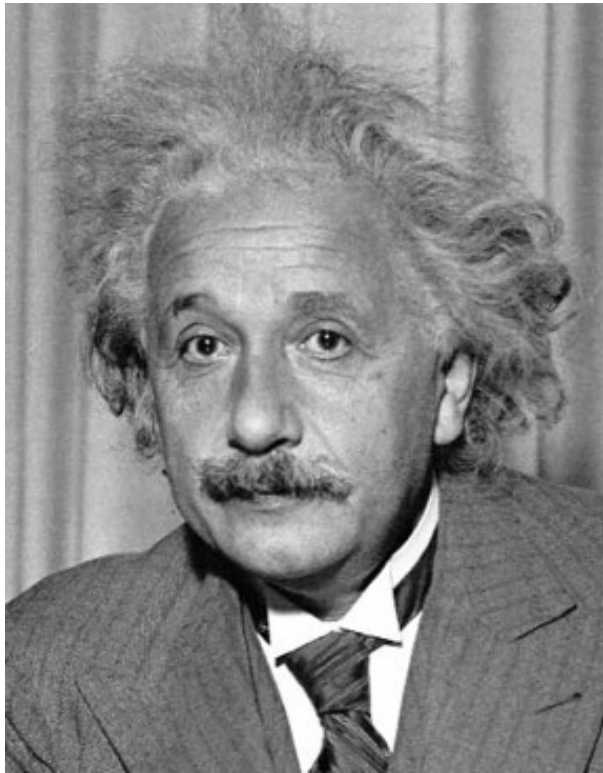
Observations so far:

- subtracting mean deals with brightness bias
- dividing by standard deviation removes contrast bias


Can we combine the two effects?

Find this template

How do we detect the template  in the following image?



What will
the output
look like?

filter  template mean

output

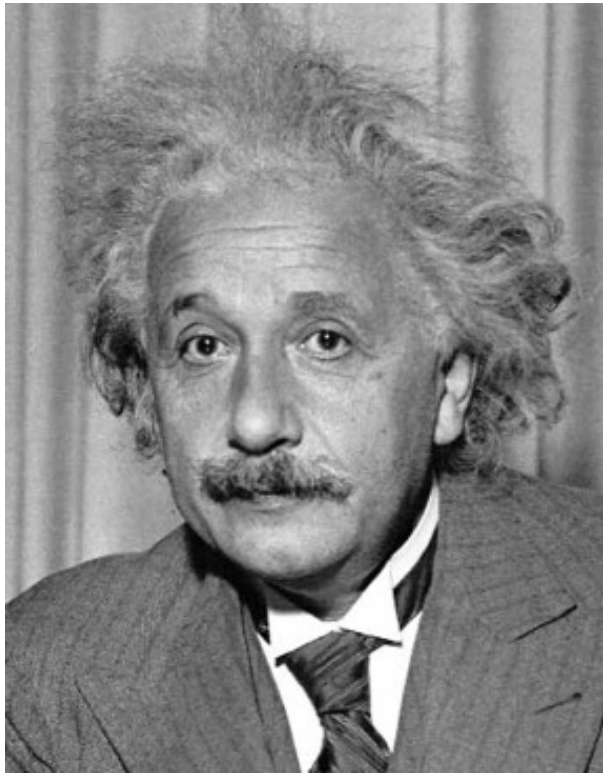
$$h[m, n] = \frac{\sum_{k,l} (g[k, l] - \bar{g})(f[m + k, n + l] - \bar{f}_{m,n})}{\sqrt{(\sum_{k,l} (g[k, l] - \bar{g})^2 \sum_{k,l} (f[m + k, n + l] - \bar{f}_{m,n})^2)}}$$

image local patch mean

Solution 4: Normalized cross-correlation (NCC).

Find this template

How do we detect the template  in the following image?



1-output



thresholding

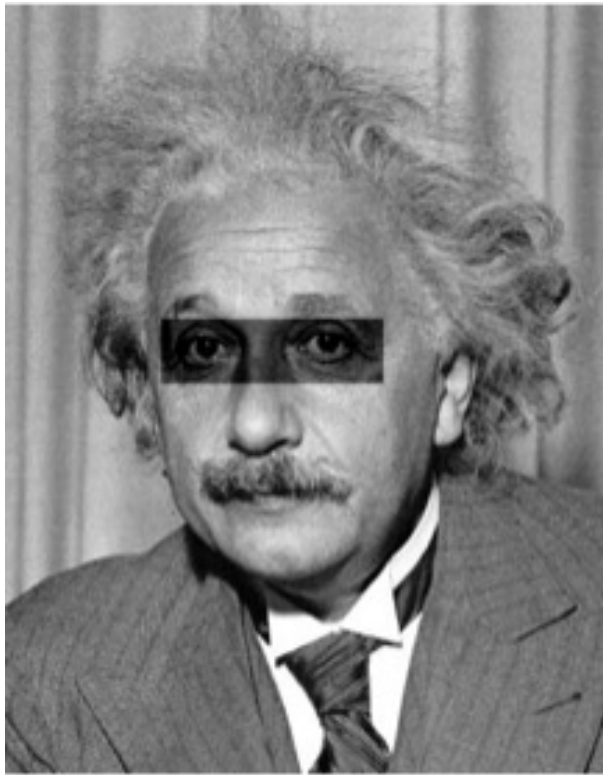
True detections



Solution 4: Normalized cross-correlation (NCC).

Find this template

How do we detect the template  in the following image?



1-output



thresholding

True detections



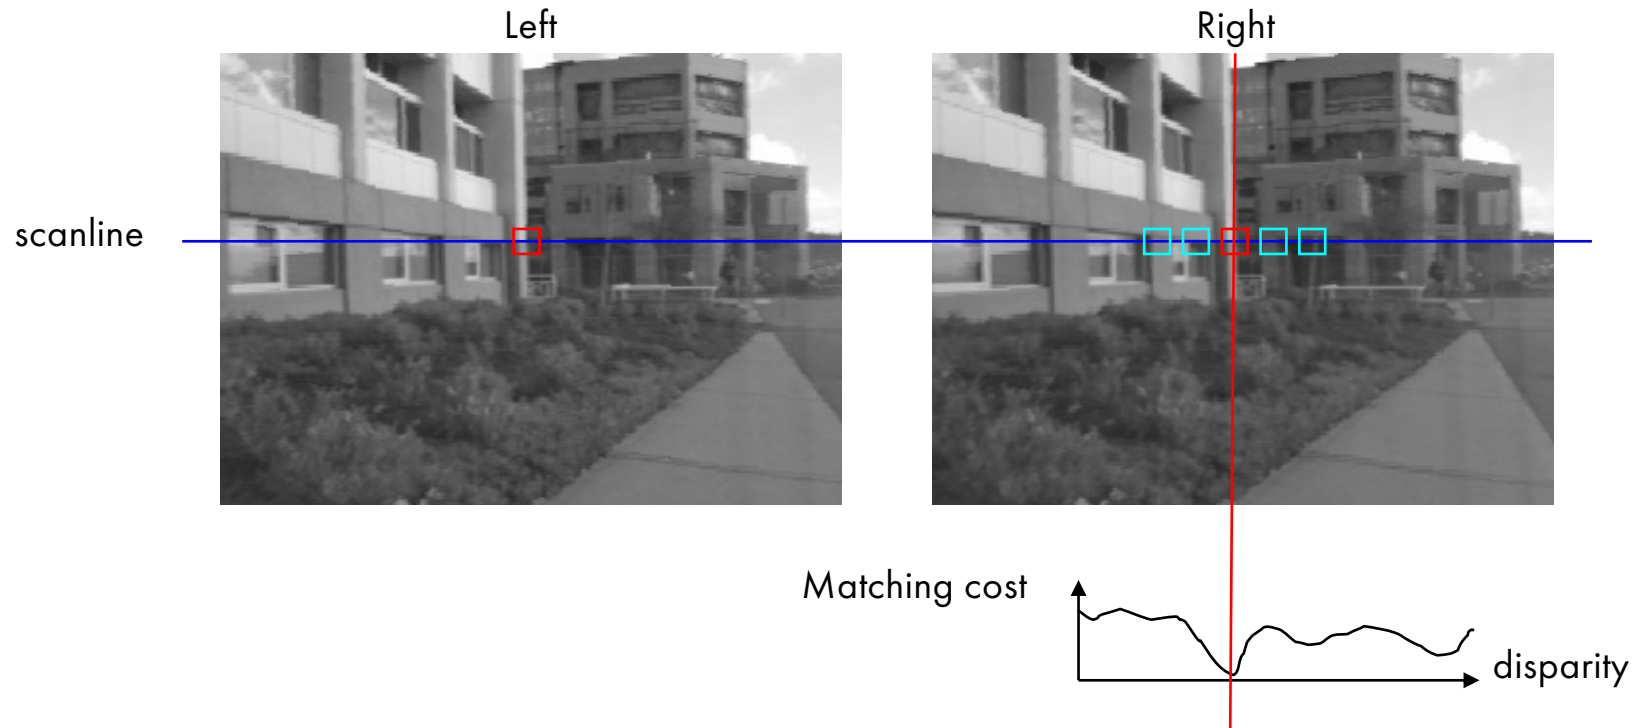
Solution 4: Normalized cross-correlation (NCC).

What is the best method?

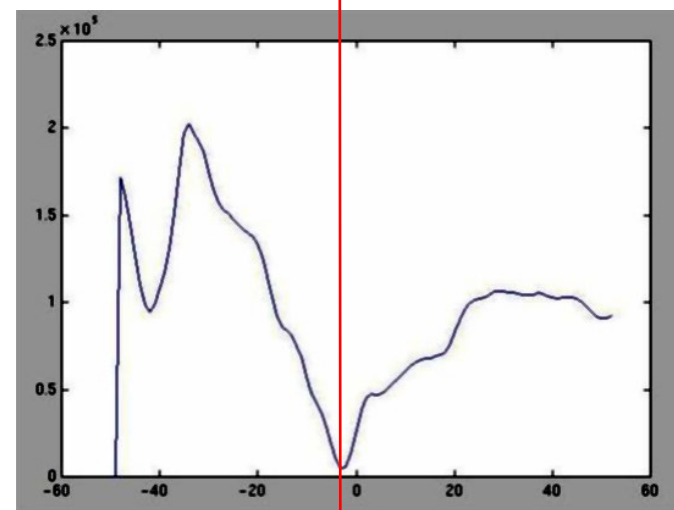
It depends on whether you care about speed or invariance.

- Zero-mean: Fastest, very sensitive to local intensity.
- Sum of squared differences: Medium speed, sensitive to intensity offsets.
- Normalized cross-correlation: Slowest, invariant to contrast and brightness.

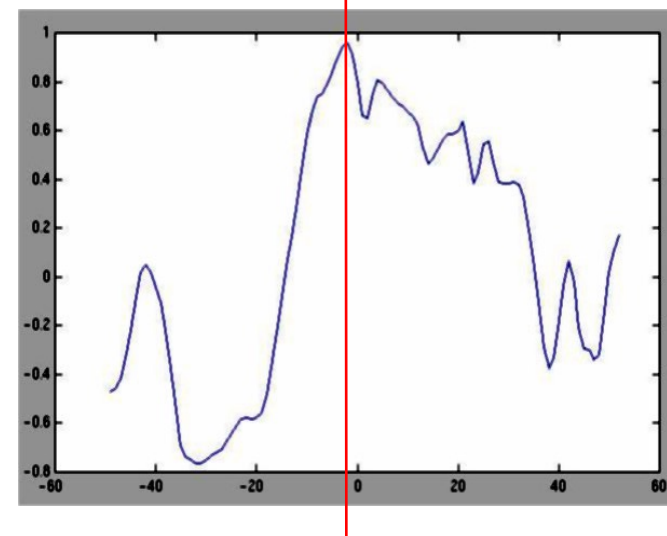
Stereo Block Matching



- Slide a window along the epipolar line and compare contents of that window with the reference window in the left image
- Matching cost: SSD or normalized correlation

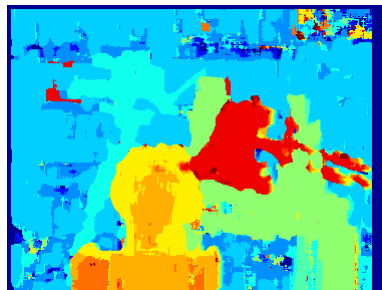


SSD

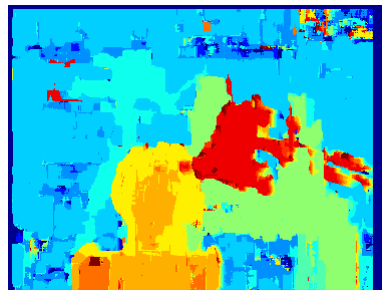


Normalized cross-correlation

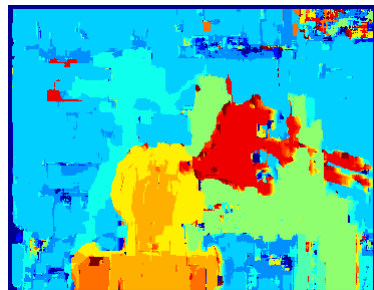
Similarity Measure	Formula
Sum of Absolute Differences (SAD)	$\sum_{(i,j) \in W} I_1(i,j) - I_2(x+i, y+j) $
Sum of Squared Differences (SSD)	$\sum_{(i,j) \in W} (I_1(i,j) - I_2(x+i, y+j))^2$
Zero-mean SAD	$\sum_{(i,j) \in W} I_1(i,j) - \bar{I}_1(i,j) - I_2(x+i, y+j) + \bar{I}_2(x+i, y+j) $
Locally scaled SAD	$\sum_{(i,j) \in W} I_1(i,j) - \frac{\bar{I}_1(i,j)}{\bar{I}_2(x+i, y+j)} I_2(x+i, y+j) $
Normalized Cross Correlation (NCC)	$\frac{\sum_{(i,j) \in W} I_1(i,j) \cdot I_2(x+i, y+j)}{\sqrt{\sum_{(i,j) \in W} I_1^2(i,j) \cdot \sum_{(i,j) \in W} I_2^2(x+i, y+j)}}$



SAD



SSD



NCC

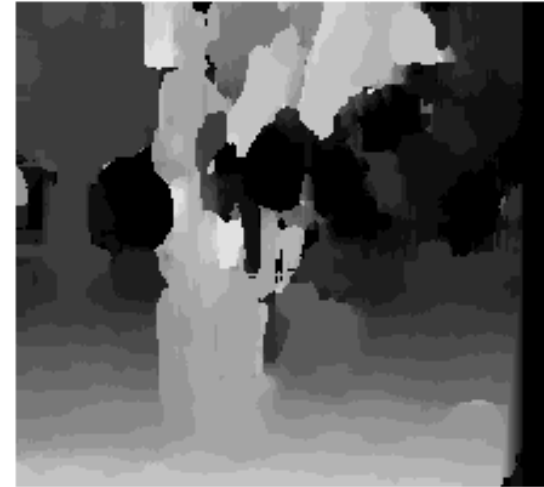


Ground truth

Effect of window size



$W = 3$



$W = 20$

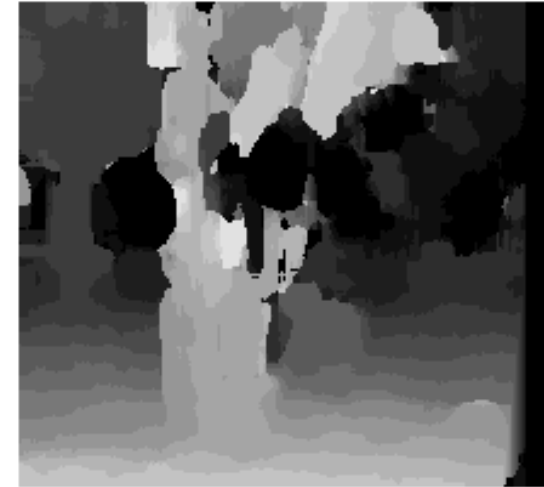
Effect of window size



$W = 3$

Smaller window

- + More detail
- More noise

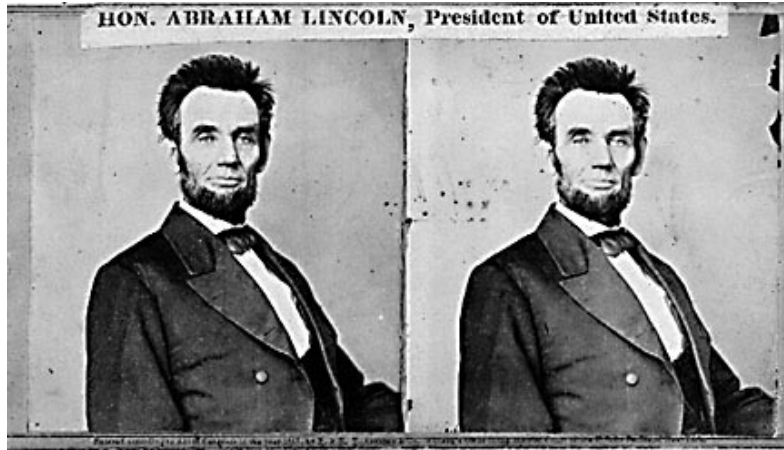


$W = 20$

Larger window

- + Smoother disparity maps
- Less detail
- Fails near boundaries

When will stereo block matching fail?



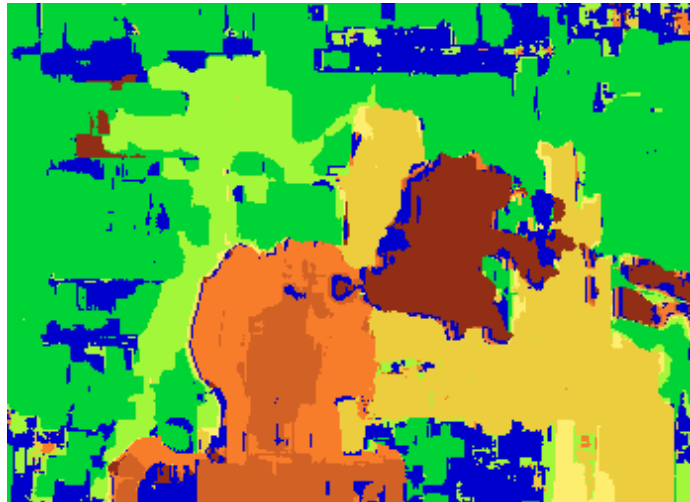
When will stereo block matching fail?



Improving stereo matching



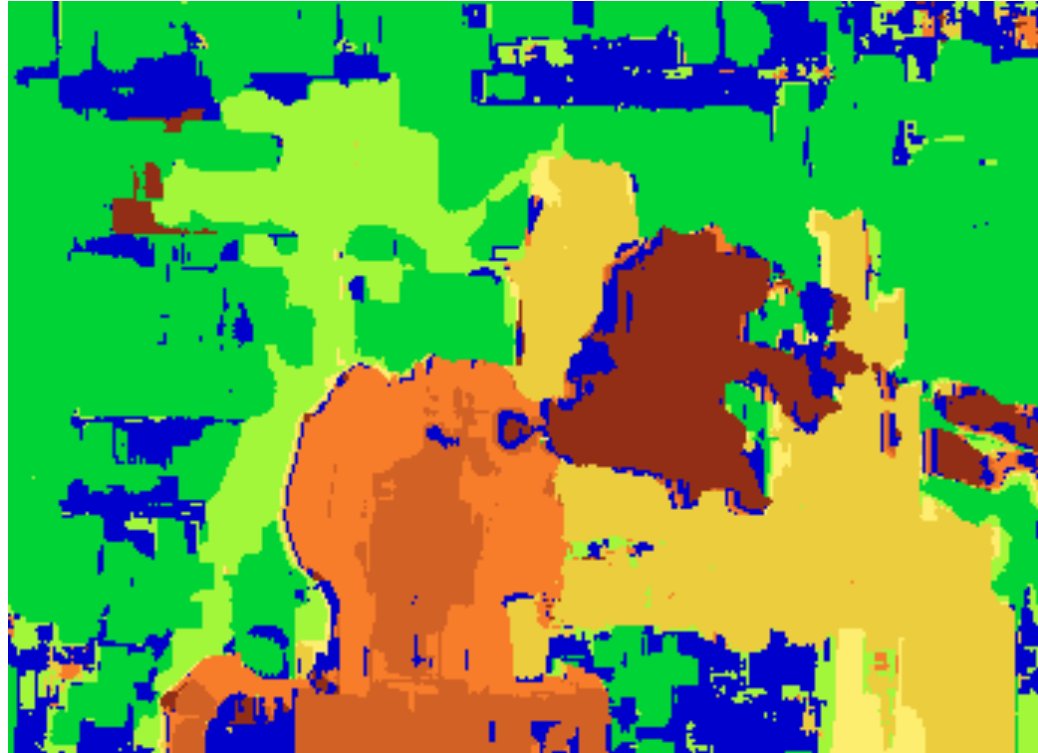
Block matching



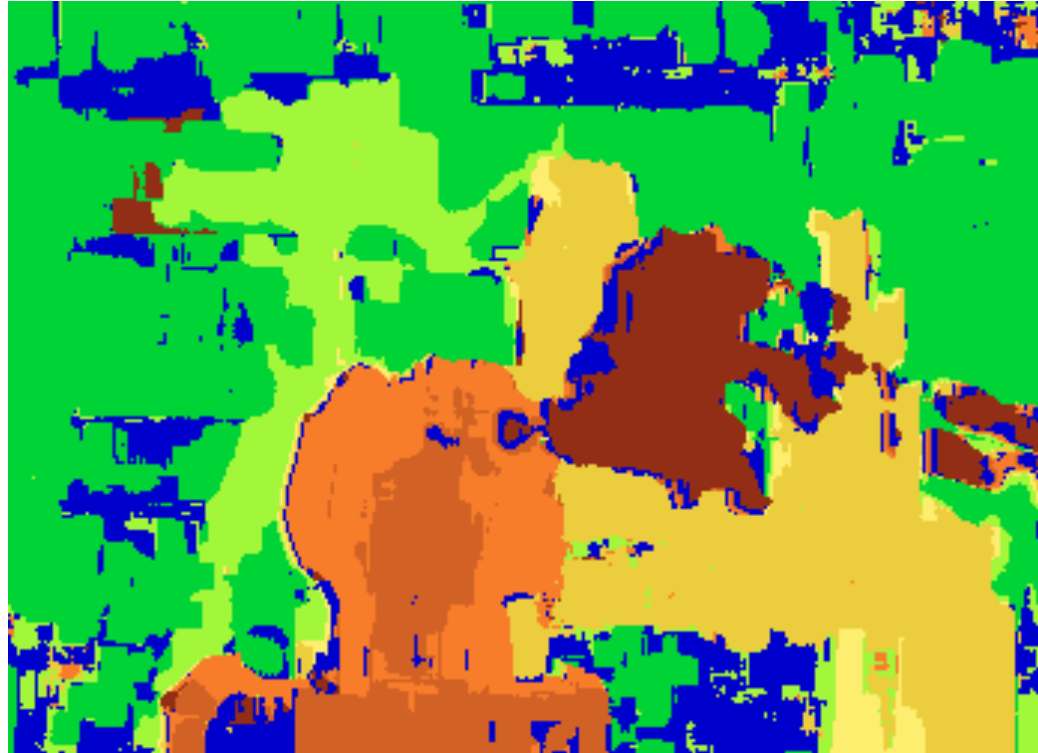
Ground truth



What are some problems with the result?



How can we improve depth estimation?



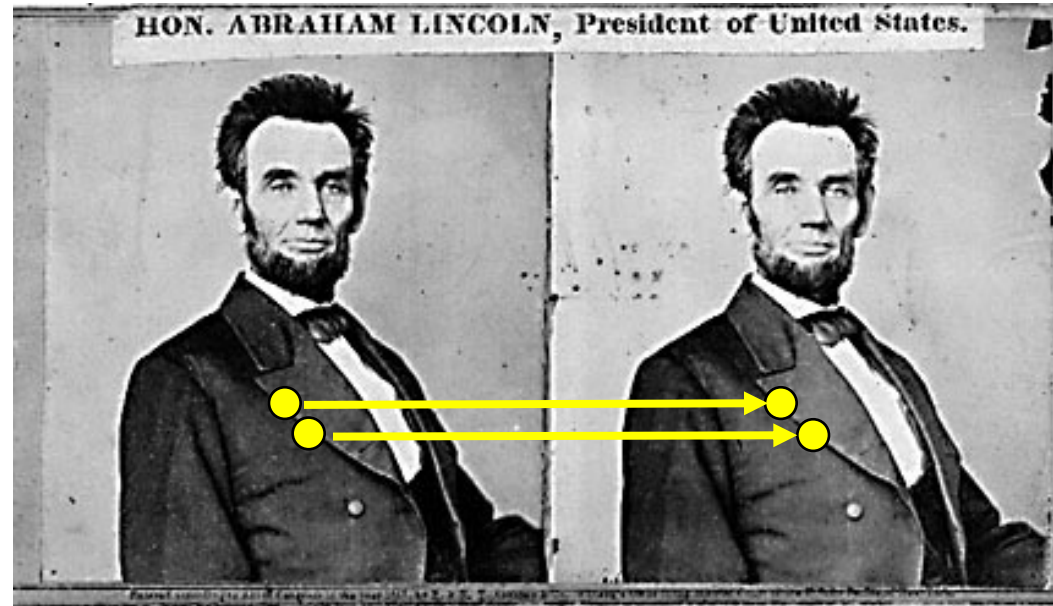
How can we improve depth estimation?

Too many discontinuities.
We expect disparity values to change slowly.

Let's make an assumption:
depth should change smoothly

Stereo matching as ...

Energy Minimization



What defines a good stereo correspondence?

1. Match quality

- Want each pixel to find a good match in the other image

2. Smoothness

- If two pixels are adjacent, they should (usually) move about the same amount

energy function
(for one pixel)

$$E(d) = \underbrace{E_d(d)}_{\text{data term}} + \lambda \underbrace{E_s(d)}_{\text{smoothness term}}$$

Want each pixel to find a good match in
the other image
(block matching result)

Adjacent pixels should (usually) move
about the same amount
(smoothness function)

$$E(d) = E_d(d) + \lambda E_s(d)$$

$$E_d(d) = \sum_{(x,y) \in I} C(x, y, d(x, y))$$

data term

SSD distance between windows centered
at $I(x, y)$ and $J(x + d(x, y), y)$

$$E(d) = E_d(d) + \lambda E_s(d)$$

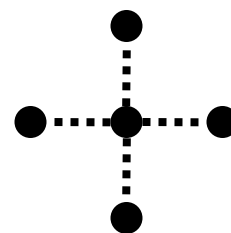
$$E_d(d) = \sum_{(x,y) \in I} C(x, y, d(x, y))$$

SSD distance between windows centered
at $I(x, y)$ and $J(x + d(x, y), y)$

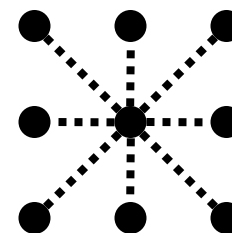
$$E_s(d) = \sum_{(p,q) \in \mathcal{E}} V(d_p, d_q)$$

smoothness term

\mathcal{E} : set of neighboring pixels



4-connected
neighborhood



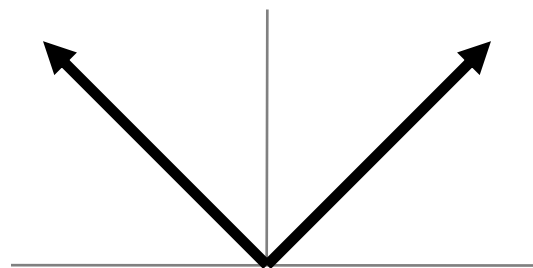
8-connected
neighborhood

$$E_s(d) = \sum_{(p,q) \in \mathcal{E}} V(d_p, d_q)$$

smoothness term

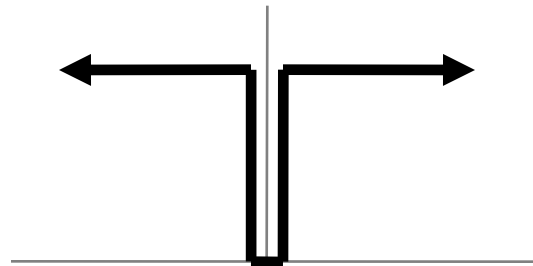
$$V(d_p, d_q) = |d_p - d_q|$$

L₁ distance



$$V(d_p, d_q) = \begin{cases} 0 & \text{if } d_p = d_q \\ 1 & \text{if } d_p \neq d_q \end{cases}$$

"Potts model"



One possible solution...

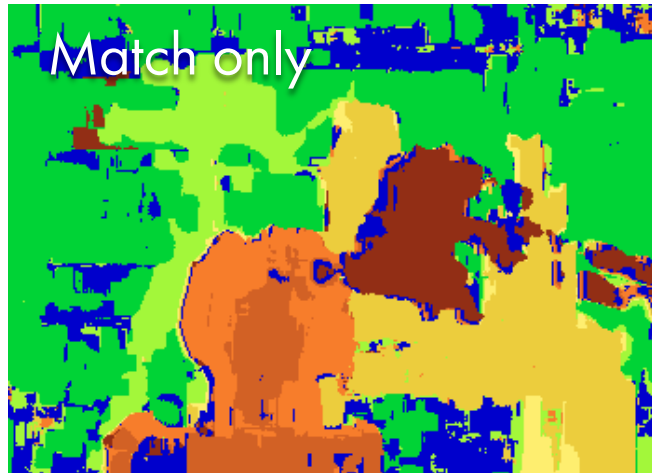
Dynamic Programming

$$E(d) = E_d(d) + \lambda E_s(d)$$

Can minimize this independently per scanline
using dynamic programming (DP) ●.....●.....●

$D(x, y, d)$: minimum cost of solution such that $d(x, y) = d$

$$D(x, y, d) = C(x, y, d) + \min_{d'} \{D(x - 1, y, d') + \lambda |d - d'|\}$$



Y. Boykov, O. Veksler, and R. Zabih, [Fast Approximate Energy Minimization via Graph Cuts](#), PAMI 2001

All of these cases remain difficult, what can we do?



Structured light

Use controlled ("structured") light to make correspondences easier

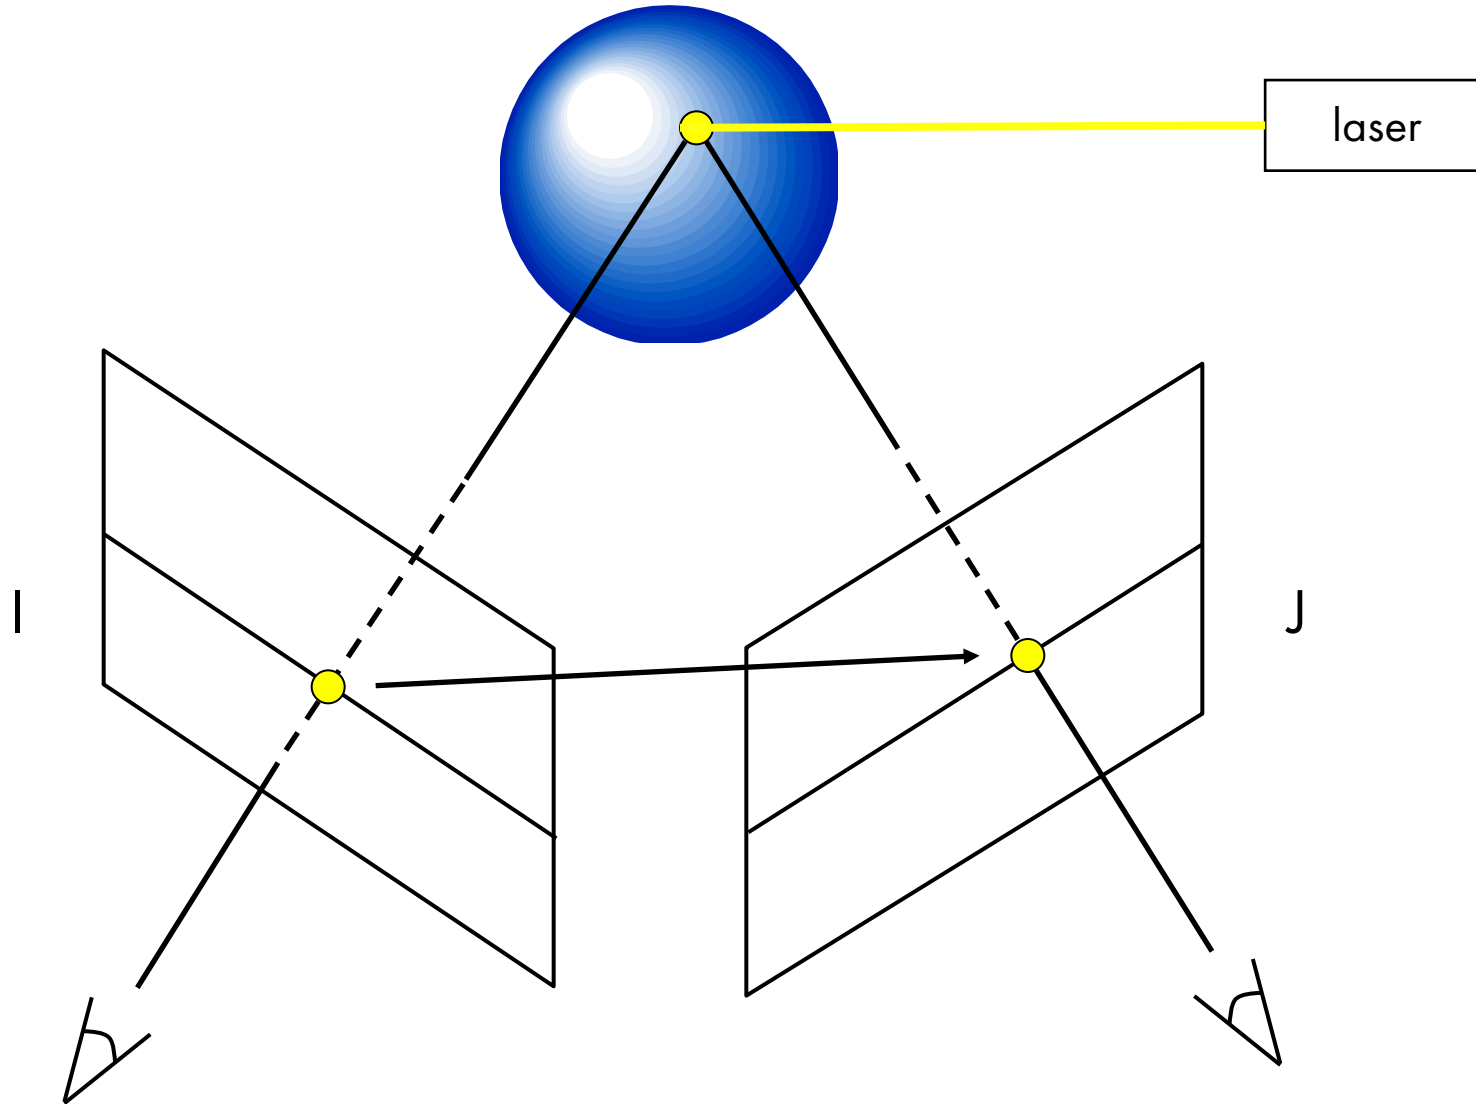
Disparity between laser points on the same scanline in the images determines the 3-D coordinates of the laser point on object



Use controlled ("structured") light to make correspondences easier

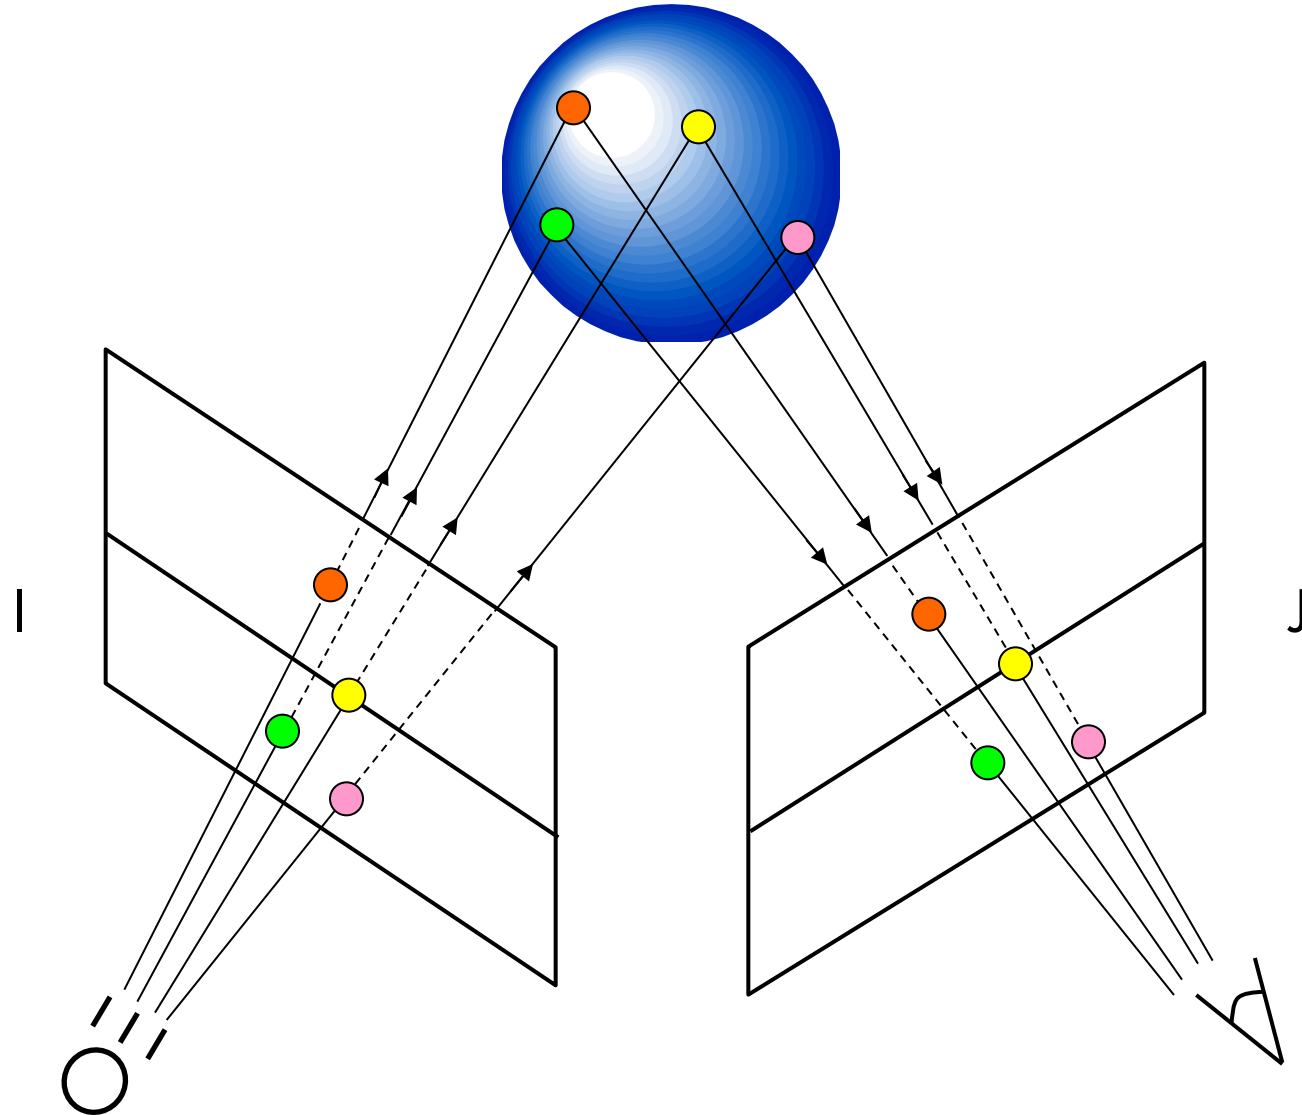


Structured light and two cameras

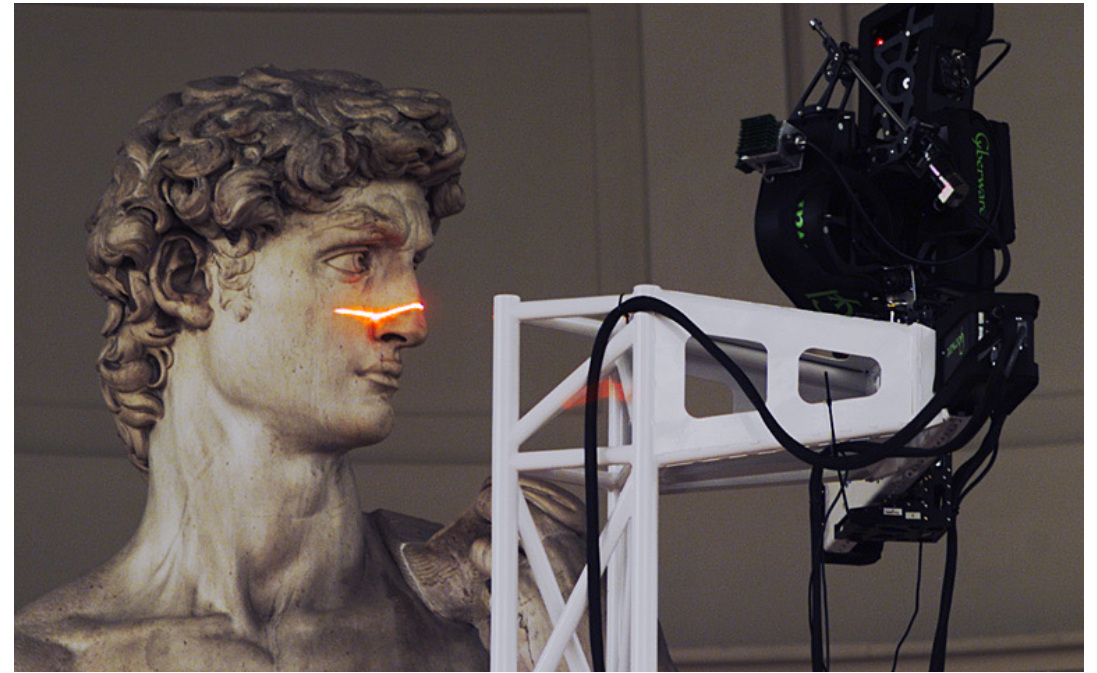
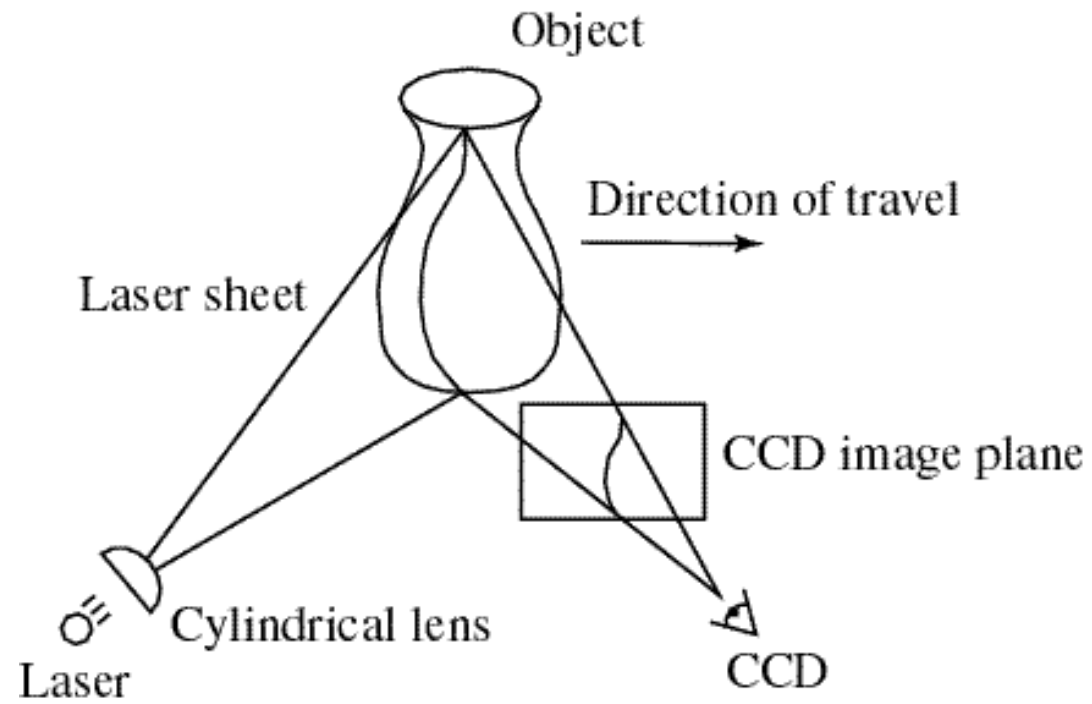


Structured light and one camera

Projector acts like
"reverse" camera



Example: Laser scanner



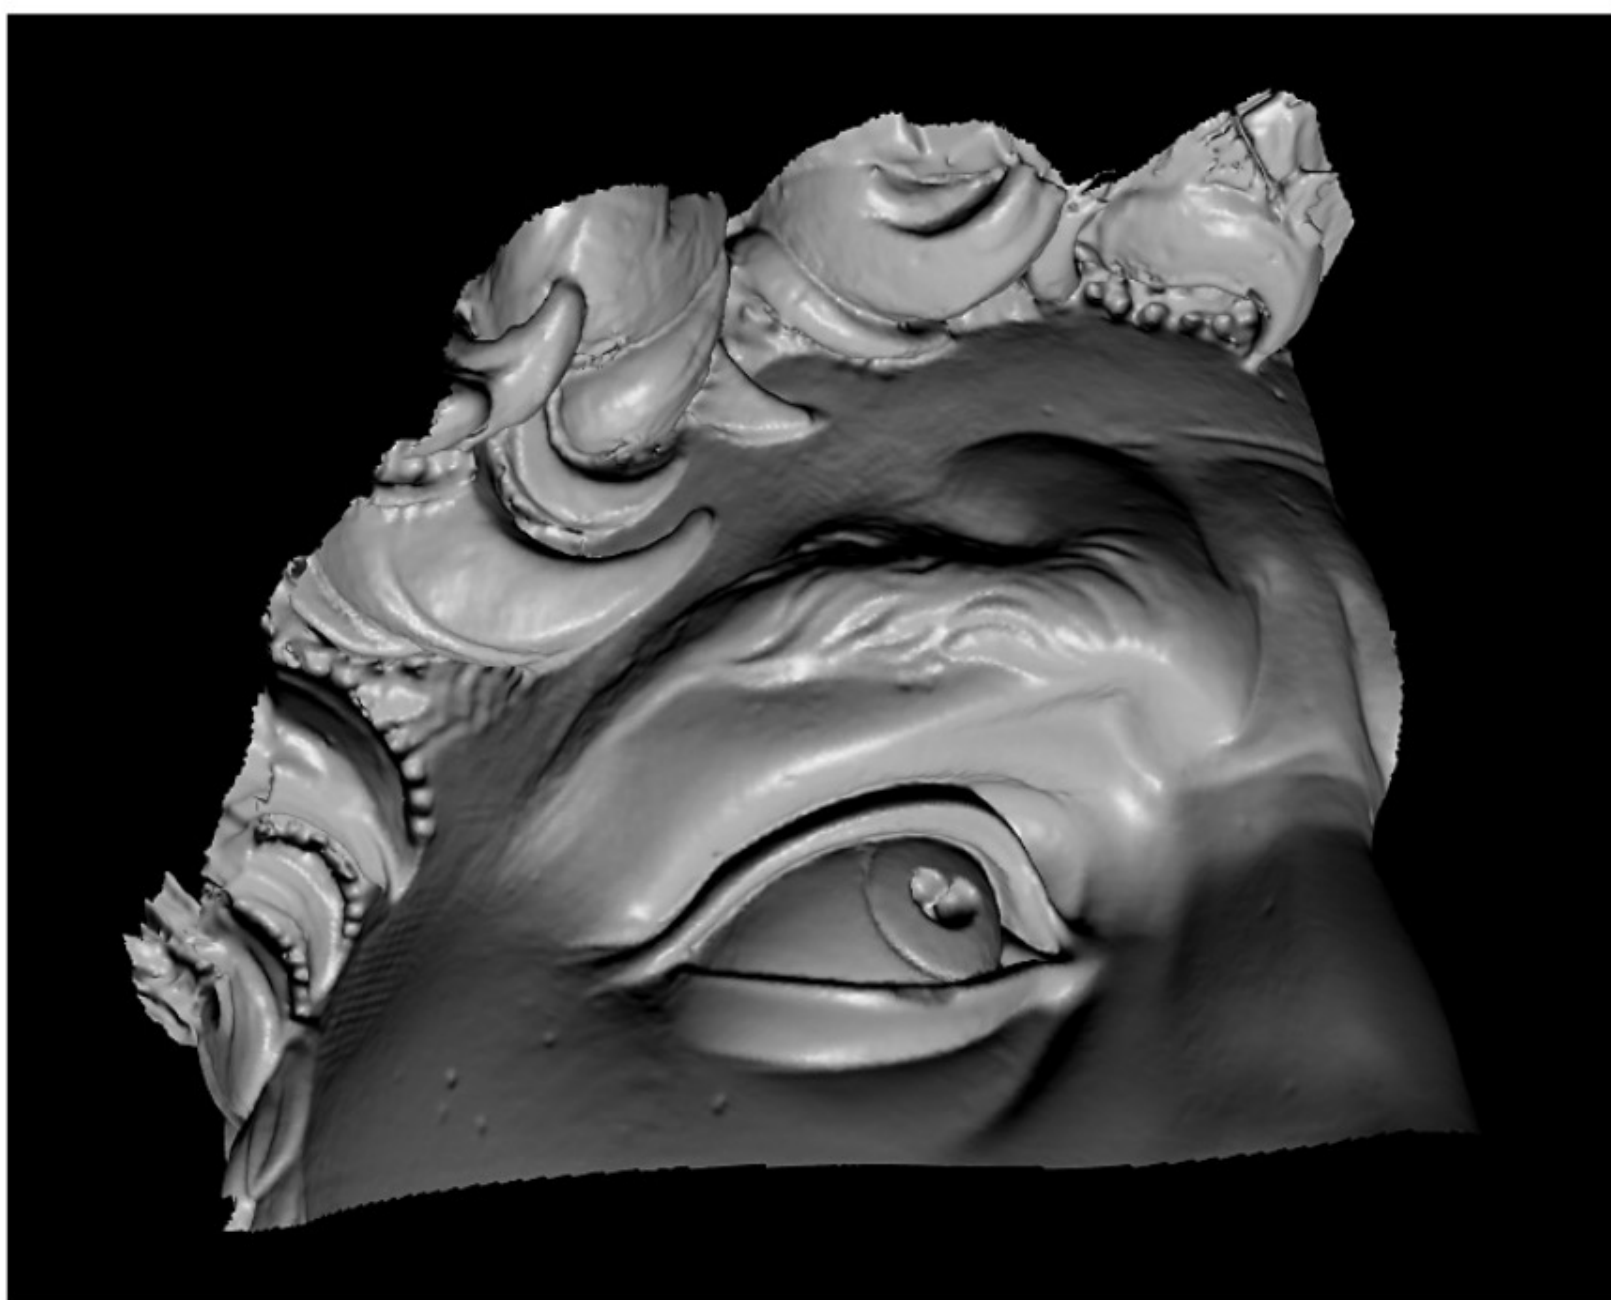
Digital Michelangelo Project
<http://graphics.stanford.edu/projects/mich/>



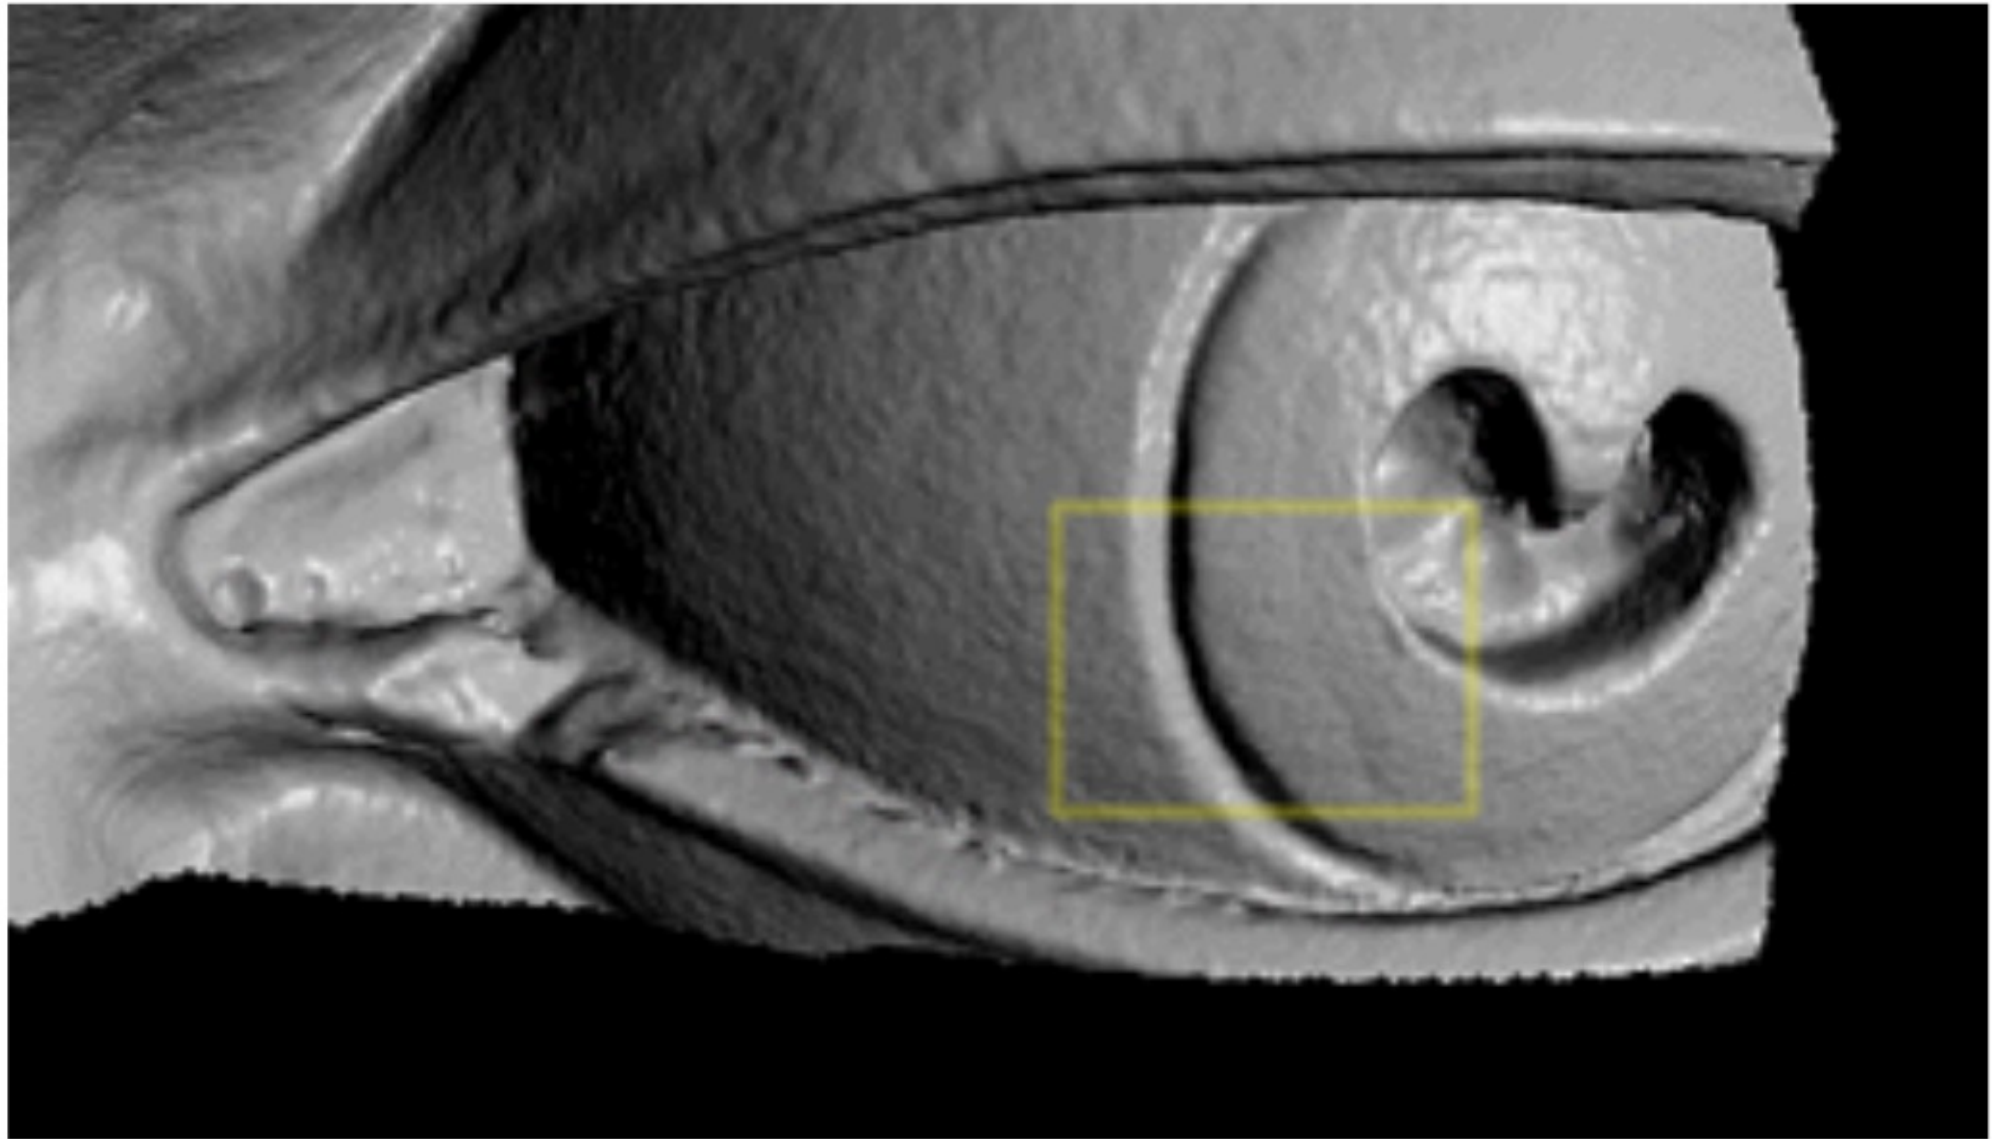
The Digital Michelangelo Project, Levoy et al.



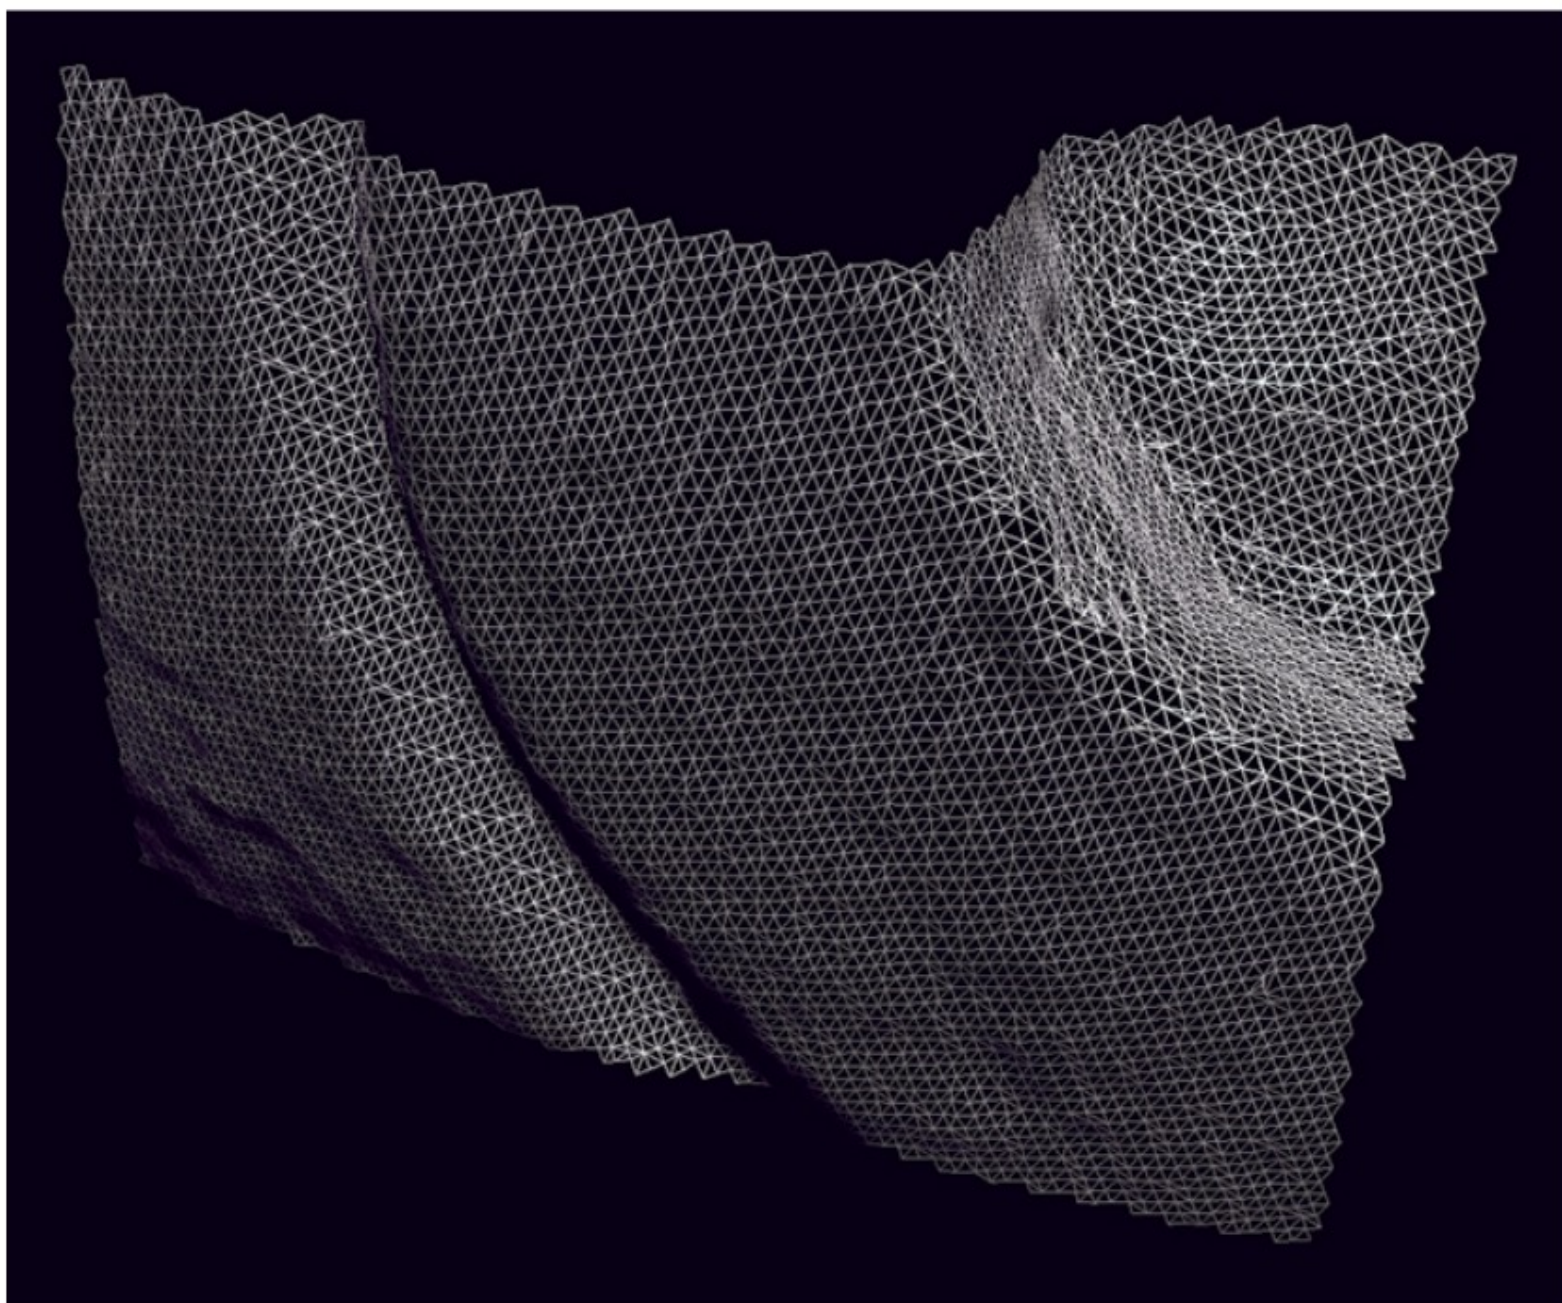
The Digital Michelangelo Project, Levoy et al.



The Digital Michelangelo Project, Levoy et al.



The Digital Michelangelo Project, Levoy et al.



The Digital Michelangelo Project, Levoy et al.

Summary – Stuff You Need To Know

Camera models

- intrinsic and extrinsic parameters
- camera matrix
- camera-to-world transformation/camera-to-camera transformation

Essential and Fundamental Matrices

- How E is derived and relates to F
- How to solve for E or F using the 8-point algorithm
- How to rectify both images to be parallel
- How are the epipoles computed?

Stereo Matching

- How does block matching work?
- How are 3D points computed once you have a rectified stereo camera?
- Given the intrinsics, disparity, and baseline, how is a 3D point computed?

References

Basic reading:

- Szeliski textbook, Section 8.1 (not 8.1.1-8.1.3), Chapter 11, Section 12.2.
- Hartley and Zisserman, Section 11.12.