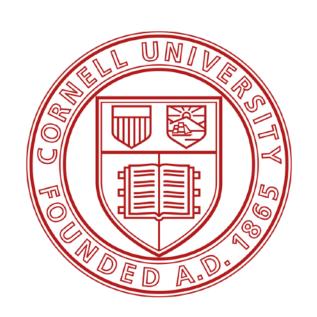
Lecture 19: Two-view geometry

CS 5670: Introduction to Computer Vision



Today

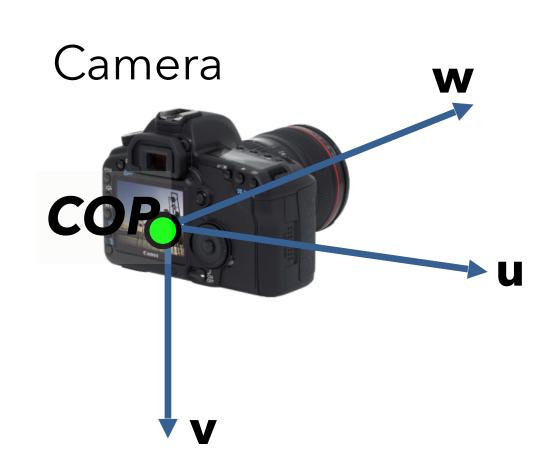
- Epipolar geometry
- Stereo matching
- Image alignment

PS5: making panoramas



Camera parameters

How can we model the geometry of a camera?



Three important coordinate systems:

- 1. World coordinates
- Camera coordinates
- Image coordinates

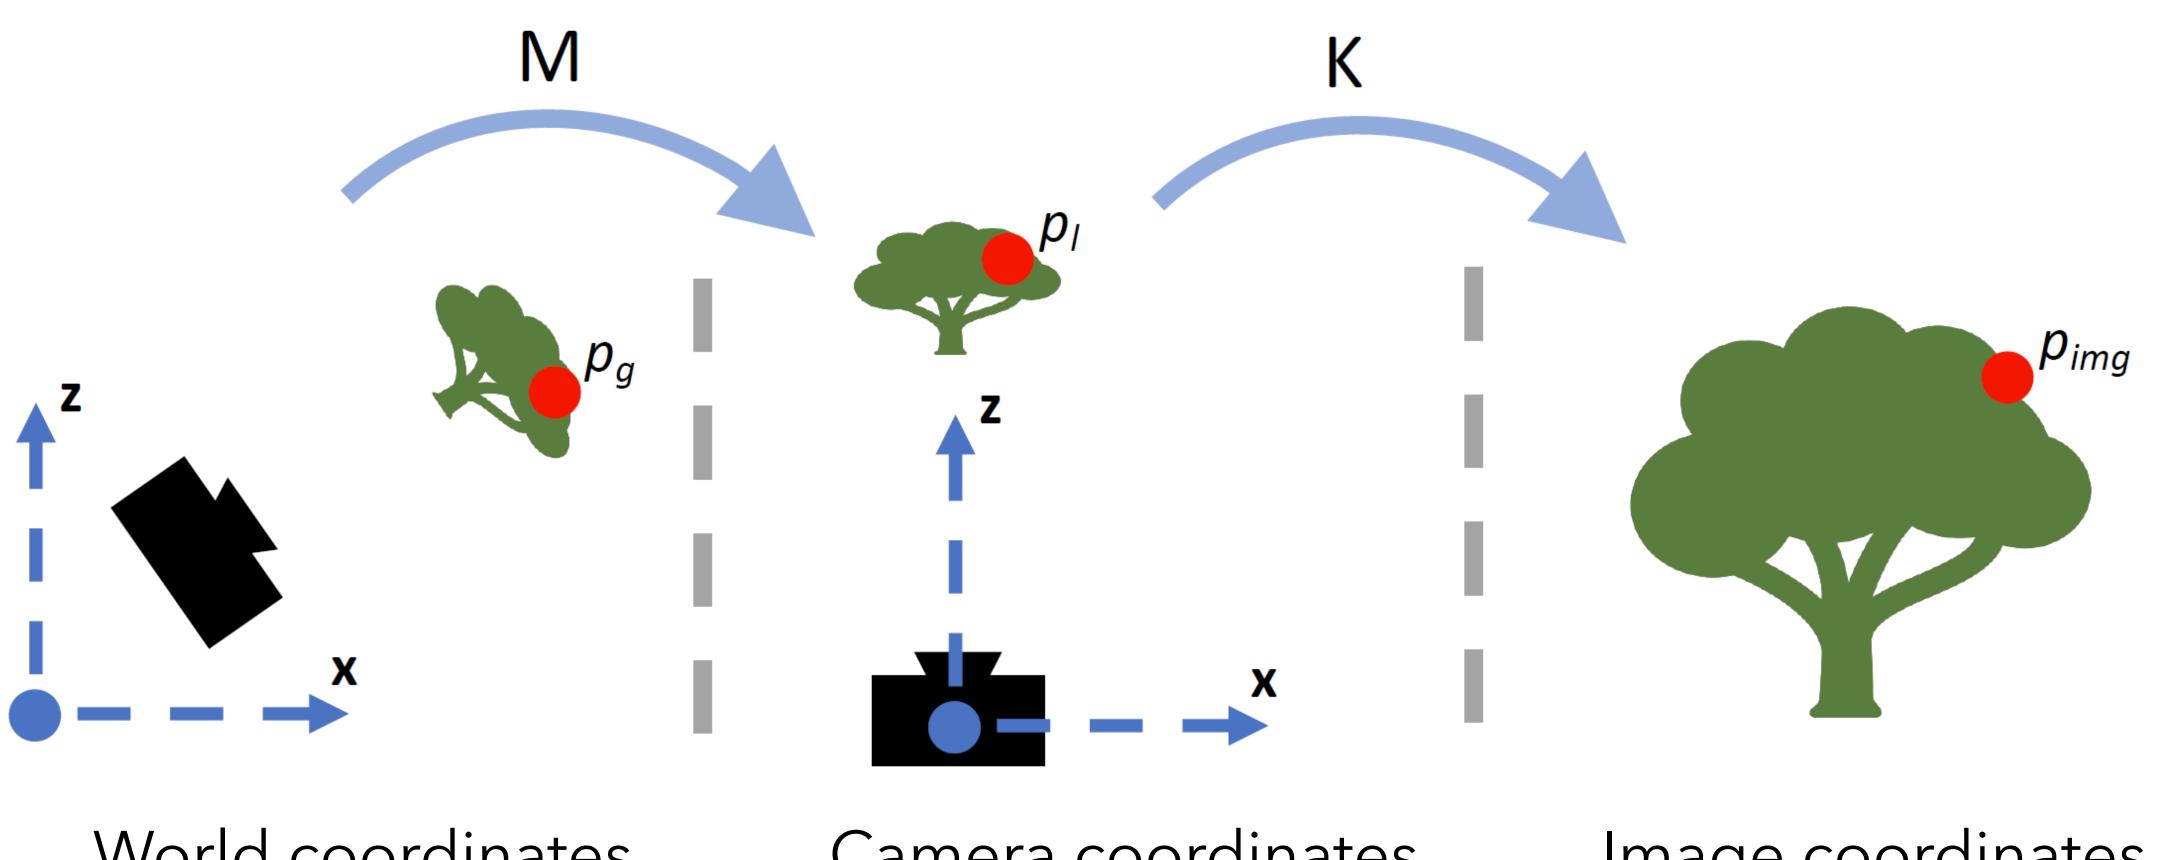


"The World"

How do we project a given world point (x, y, z) to an image point?

Source: N. Shavely

Coordinate frames



World coordinates

Camera coordinates

Image coordinates

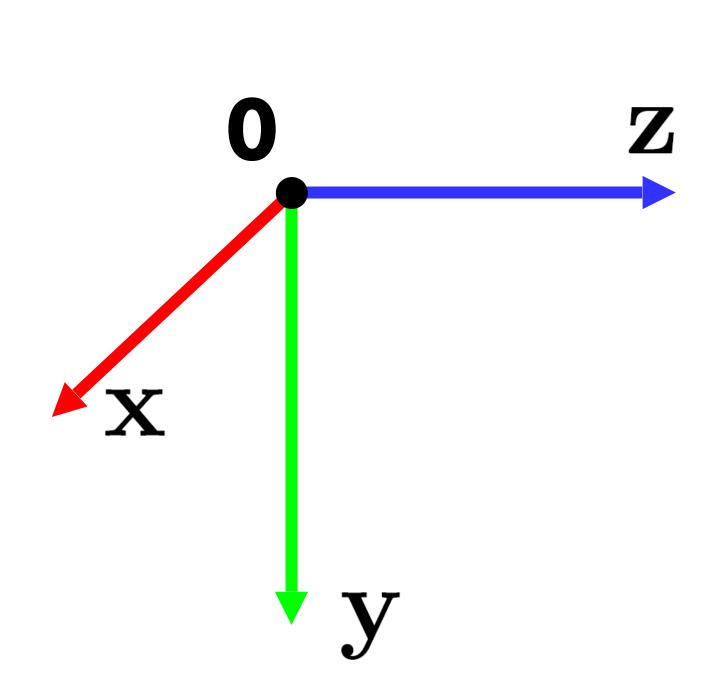
Source: N. Snavely Figure credit: Peter Hedman

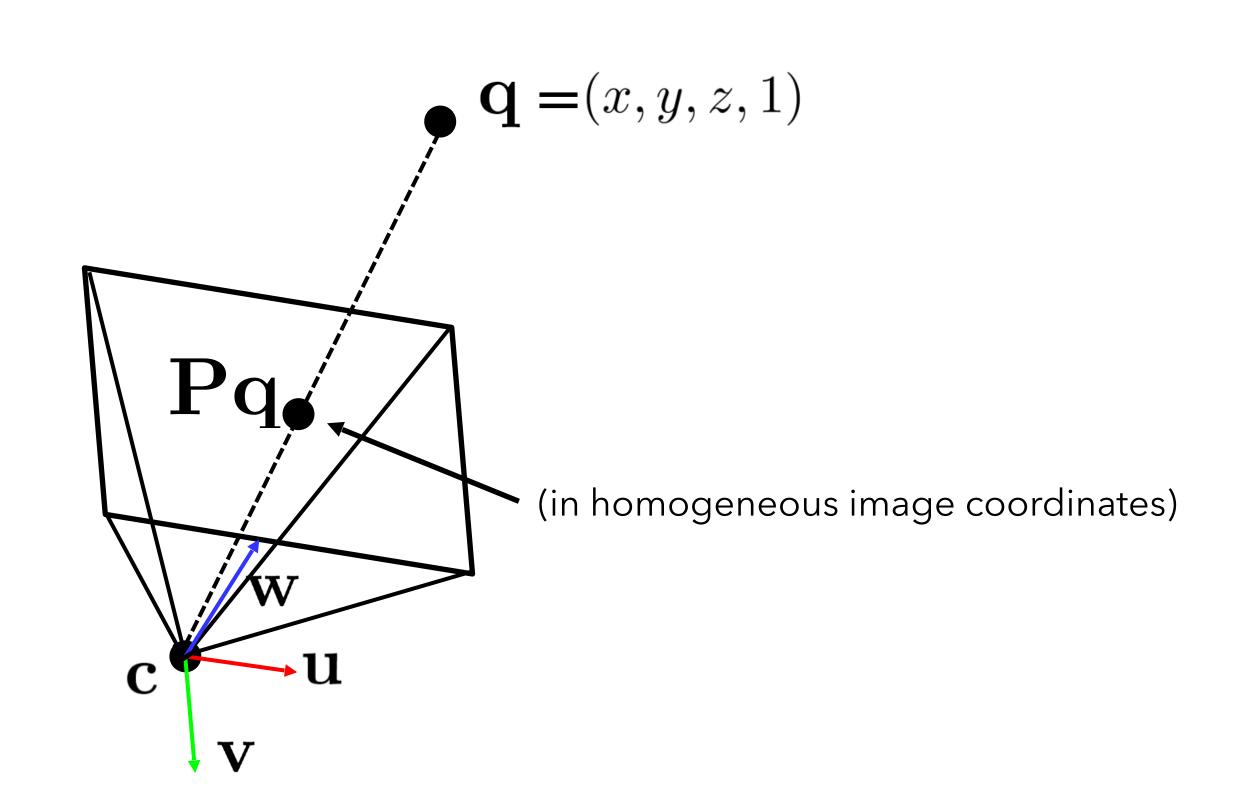
Camera parameters

To project a point (x, y, z) in world coordinates into a camera

- First transform (x, y, z) into camera coordinates
- Need to know
 - Camera position (in world coordinates)
 - Camera orientation (in world coordinates)
- Then project into the image plane to get image (pixel) coordinates
 - Need to know camera intrinsics

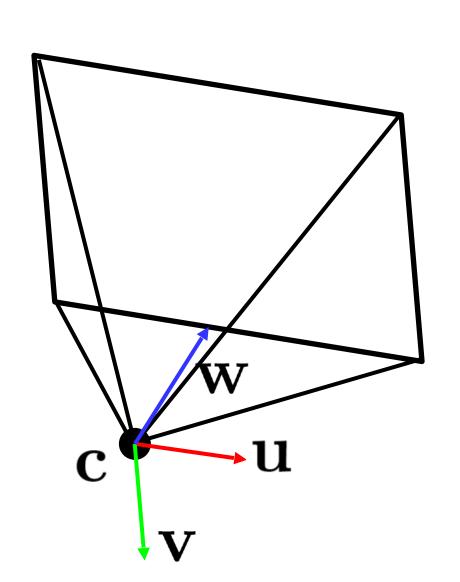
Projection matrix





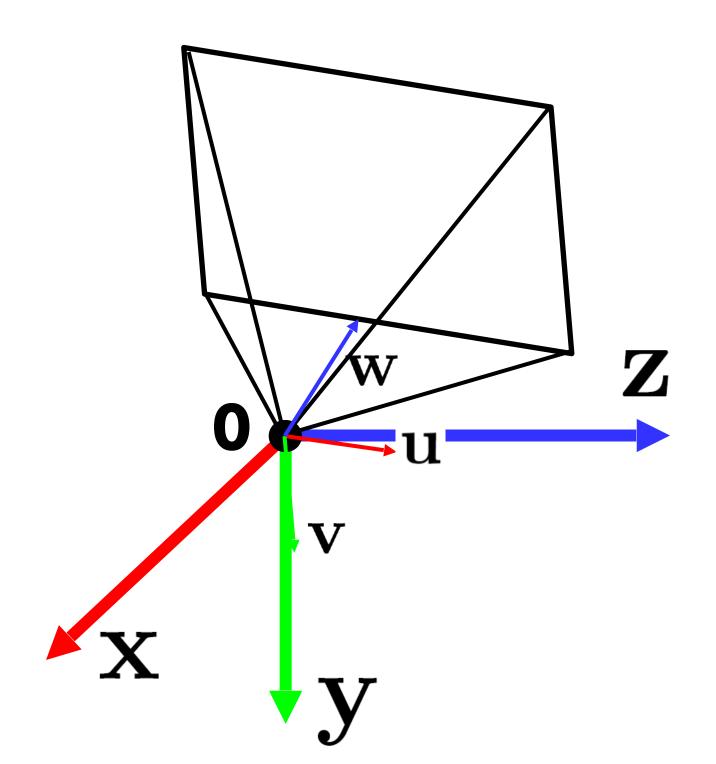
- How do we get the camera to "canonical form"?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)

Step 1: Translate by -c



8

- How do we get the camera to "canonical form"?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)



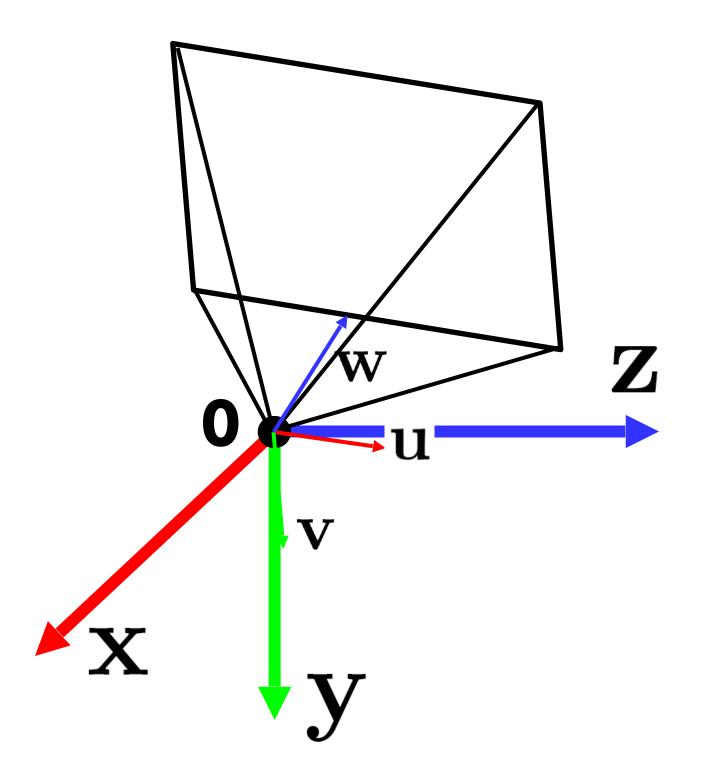
Step 1: Translate by -c

How do we represent translation as a matrix

$$\mathbf{\Gamma} = \begin{bmatrix} \mathbf{I}_{3\times3} & -\mathbf{c} \\ 0 & 0 & 0 \end{bmatrix}$$

.__ . . | 1: _ _ 1: _ _ ._ 9

- How do we get the camera to "canonical form"?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)

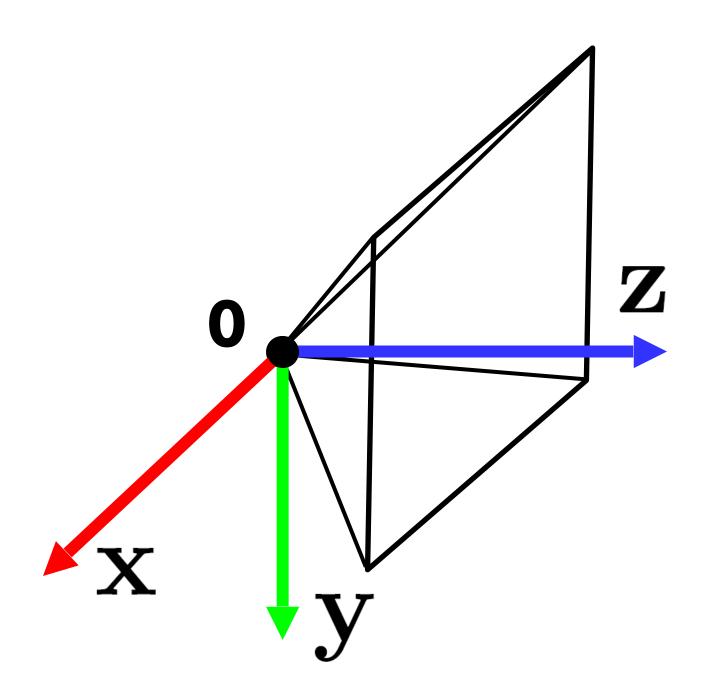


Step 1: Translate by -c

Step 2: Rotate by R

$$\mathbf{R} = \left[egin{array}{c} \mathbf{u}^T \ \mathbf{v}^T \ \mathbf{w}^T \end{array}
ight]$$
 3x3 rotation matrix

- How do we get the camera to "canonical form"?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)



Step 1: Translate by -c

Step 2: Rotate by R

$$\mathbf{R} = egin{bmatrix} \mathbf{u}^T \ \mathbf{v}^T \ \mathbf{w}^T \end{bmatrix}$$

(with extra row/column of [0 0 0 1])

Perspective projection

$$\begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

(converts from 3D rays in camera coordinate system to pixel coordinates)

in general,
$$\mathbf{K} = egin{bmatrix} f & s & c_x \\ 0 & lpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}$$
 (upper triangular matrix)

Q: aspect ratio (1 unless pixels are not square)

 $S:\mathbf{skew}$ (0 unless pixels are shaped like rhombi/parallelograms)

 (c_x,c_u) : principal point ((w/2,h/2) unless optical axis doesn't intersect projection plane at image center)

Typical intrinsics matrix

$$\mathbf{K} = \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

- 2D affine transform corresponding to a scale by f (focal length) and a translation by (c_x, c_y) (principal point)
- Maps 3D rays to 2D pixels

Projection matrix

$$\mathbf{P} = \mathbf{K} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{R} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3 \times 3} & -\mathbf{c} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
projection rotation translation

The **K** matrix converts 3D rays in the camera's coordinate system to 2D image points in image (pixel) coordinates.

This part converts 3D points in world coordinates to 3D rays in the camera's coordinate system. There are 6 parameters represented (3 for position/translation, 3 for rotation).

Projection matrix

$$\mathbf{P} = \mathbf{K} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{R} & 0 \\ 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3 \times 3} & -\mathbf{c} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
projection
$$\begin{bmatrix} \mathbf{R} & -\mathbf{Rc} \end{bmatrix}$$
(sometimes called t)
$$\mathbf{P} = \mathbf{K} \begin{bmatrix} \mathbf{R} & -\mathbf{Rc} \end{bmatrix}$$

Focal length

Can think of as "zoom"



24mm



50mm



200mm

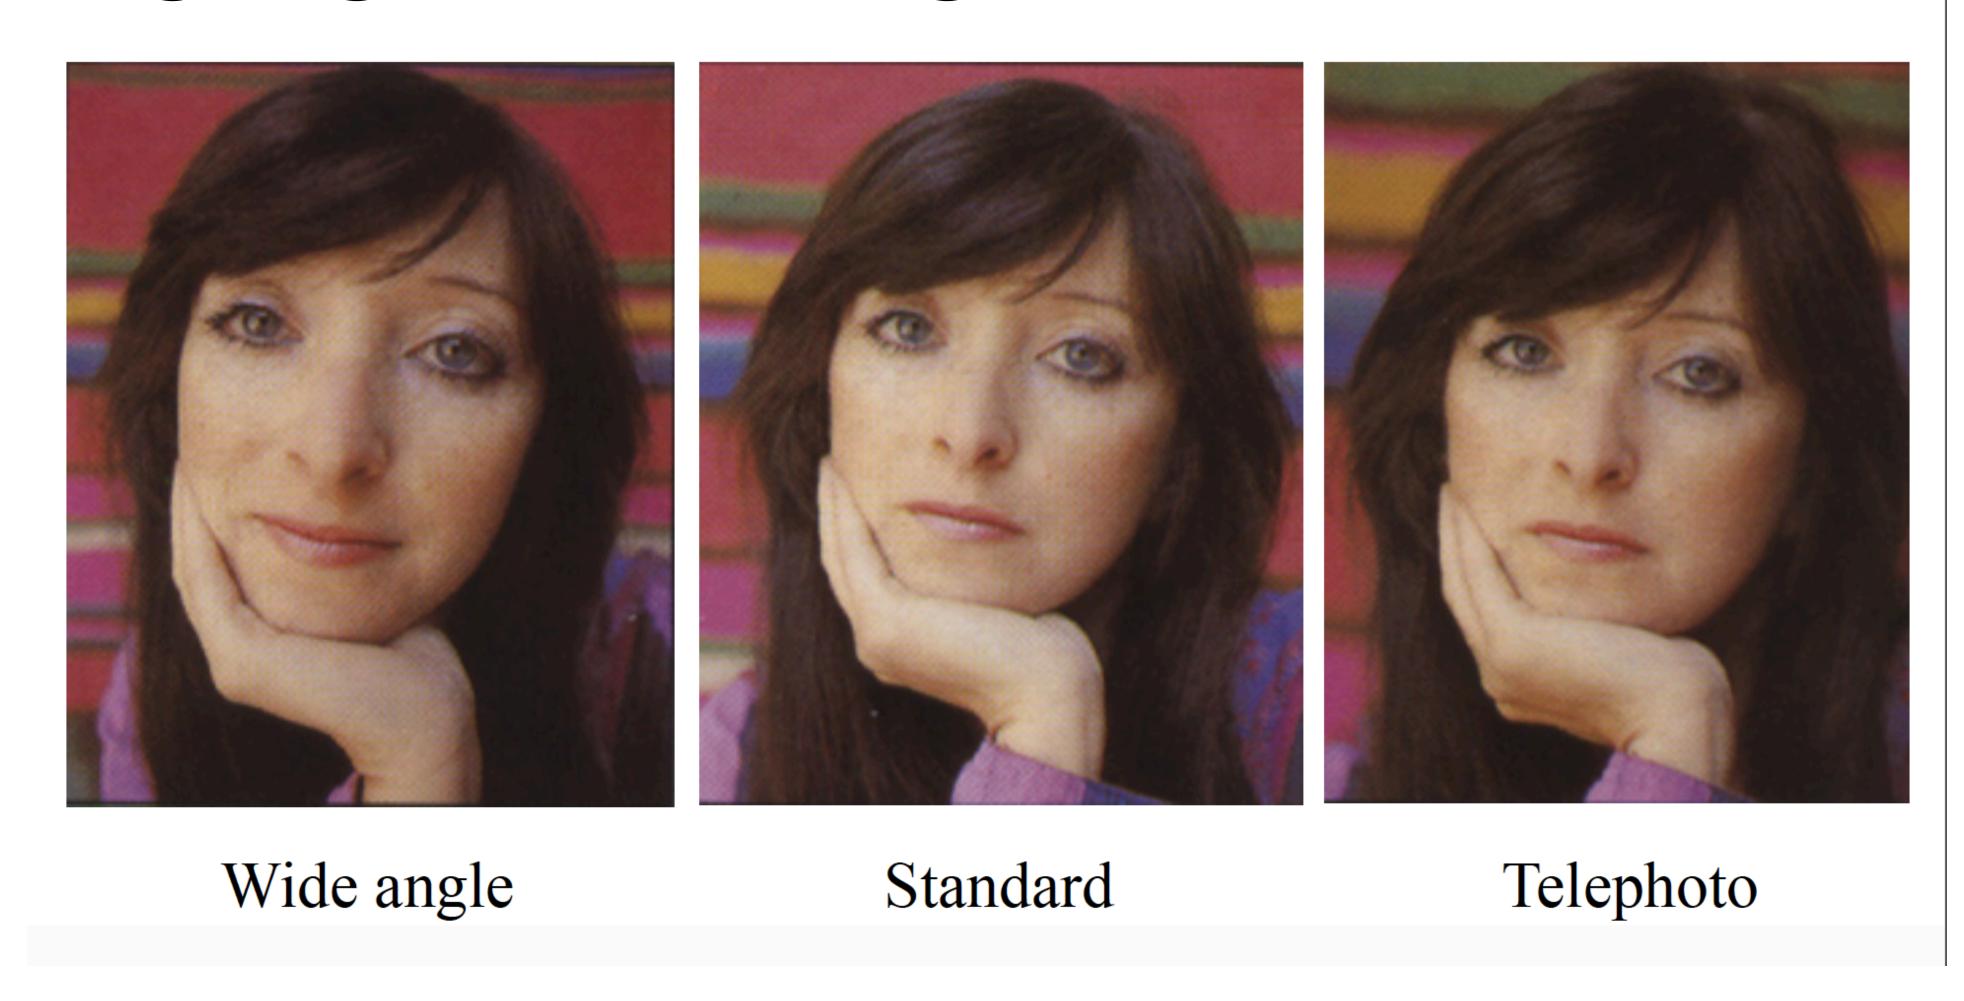


Also related to field of view

1 6

Source: N. Snavely

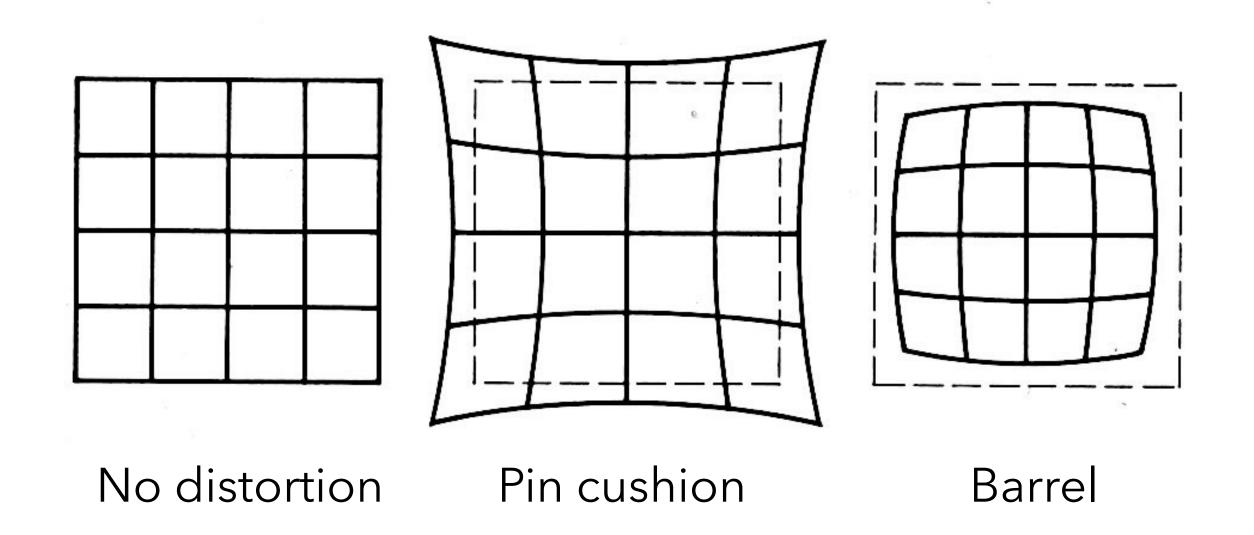
Changing focal length





http://petapixel.com/2013/01/11/how-focal-length-affects-your-subjects-apparent-weight-as-seen-with-a-cat/

Distortion



- Radial distortion of the image
 - Caused by imperfect lenses
 - Deviations are most noticeable for rays that pass through the edge of the lens





Modeling distortion

Project
$$x_n' = \hat{x}/\hat{z}$$
 to "normalized" $y_n' = \hat{y}/\hat{z}$ image coordinates
$$r^2 = x_n'^2 + y_n'^2$$
 Apply radial distortion
$$x_d' = x_n'(1 + \kappa_1 r^2 + \kappa_2 r^4)$$

$$y_d' = y_n'(1 + \kappa_1 r^2 + \kappa_2 r^4)$$
 Apply focal length translate image center
$$y' = fy_d' + y_c$$

- To model lens distortion
 - Use above projection operation instead of standard projection matrix multiplication

Correcting radial distortion

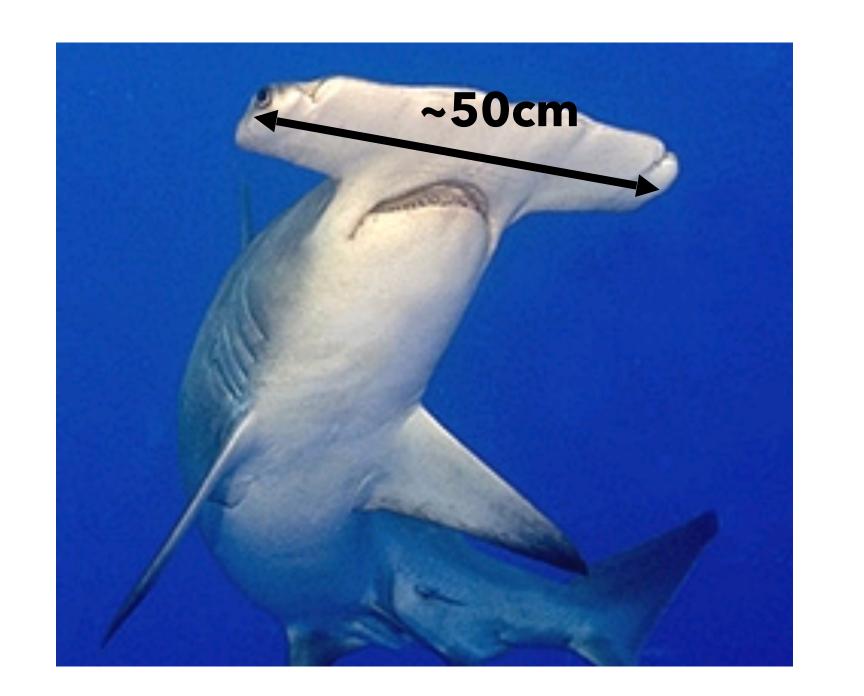




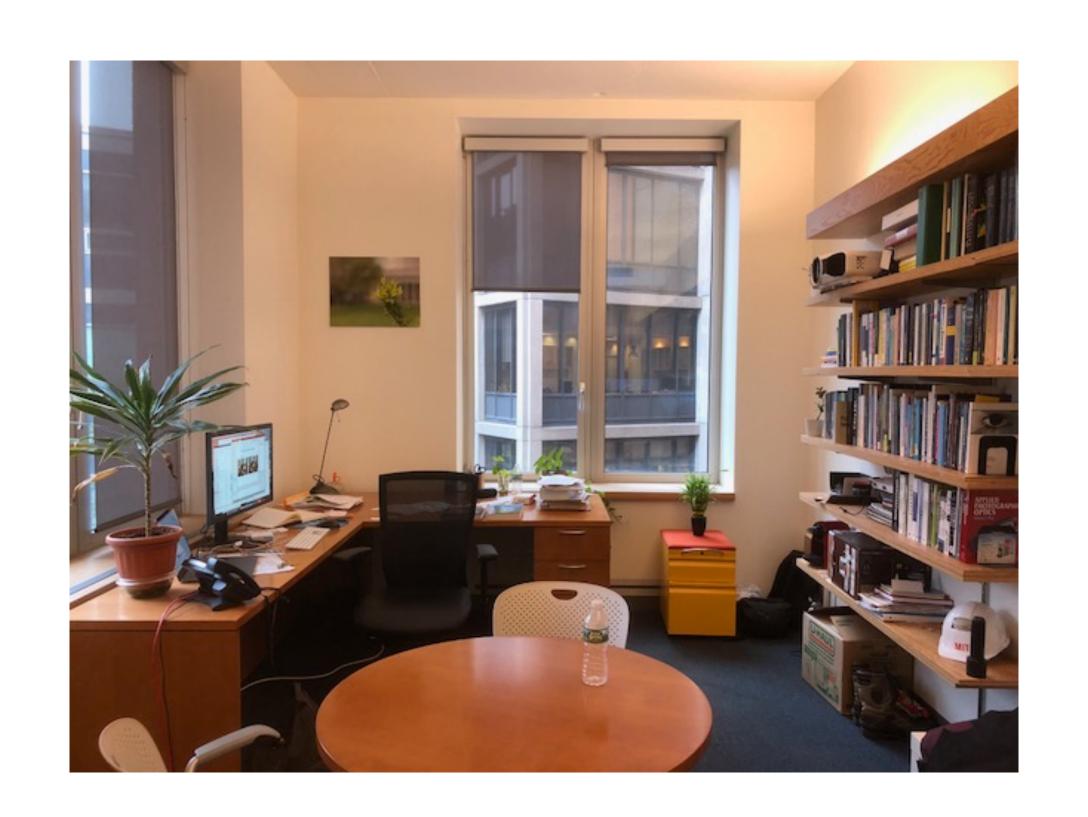
Estimating depth from multiple views

Stereo vision

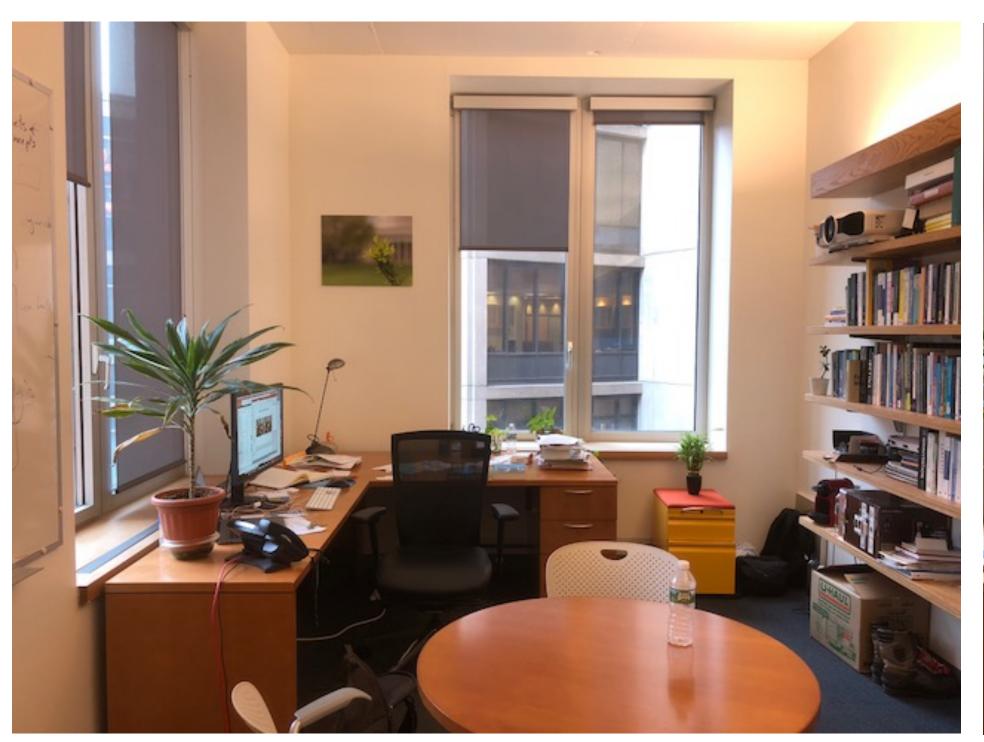




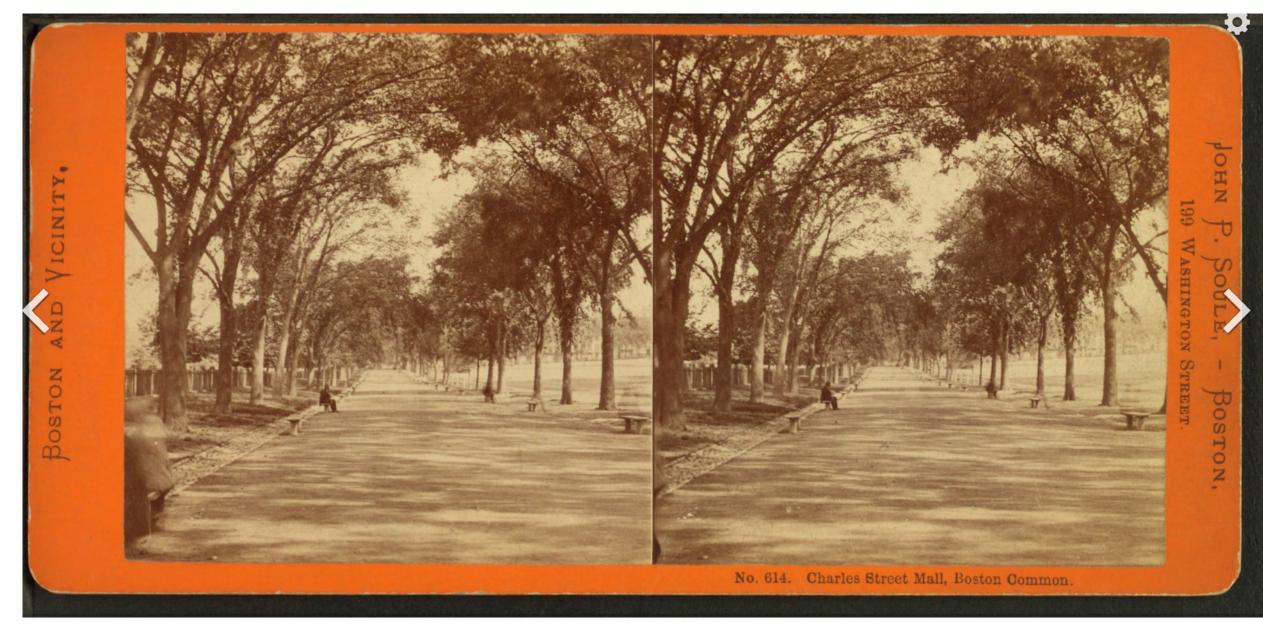
1 vs. 2 eyes



1 vs. 2 eyes





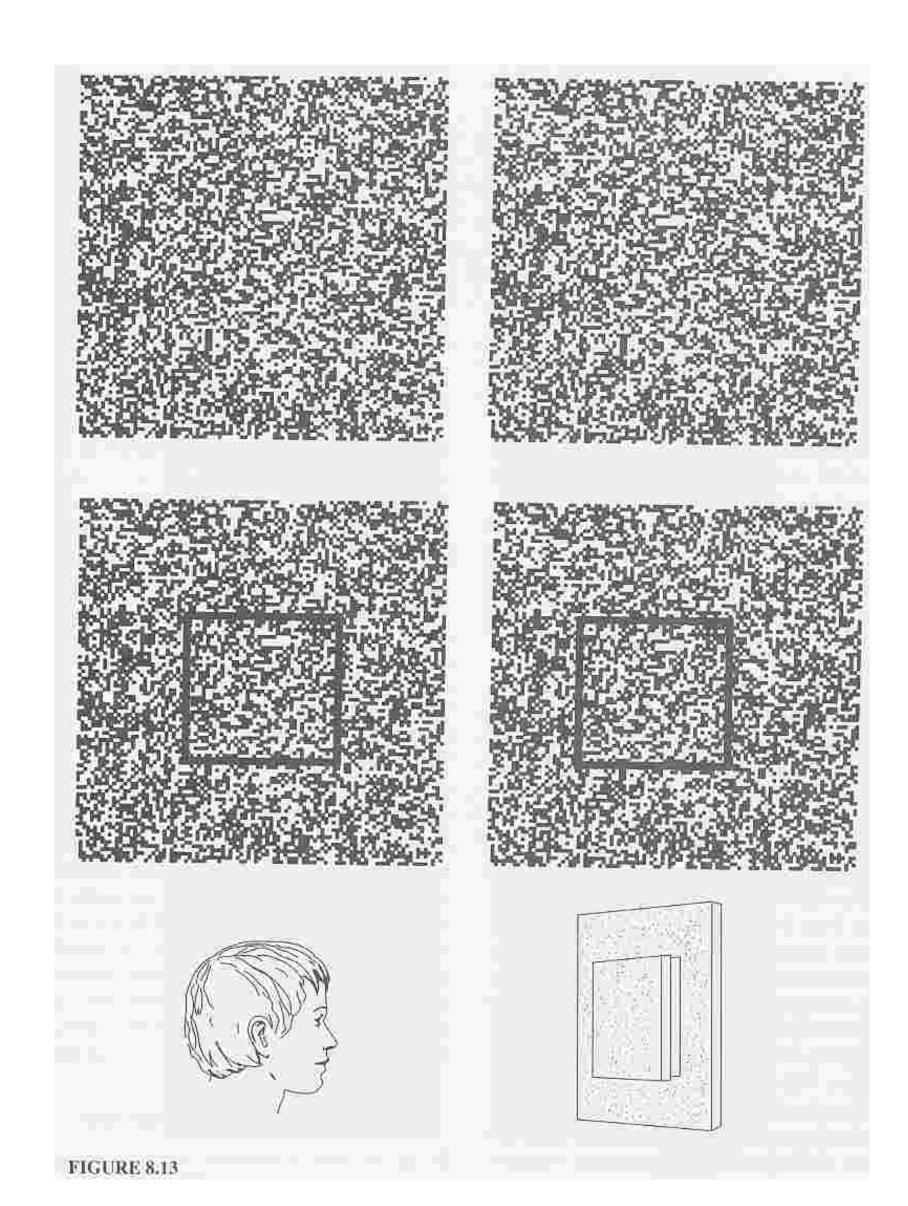


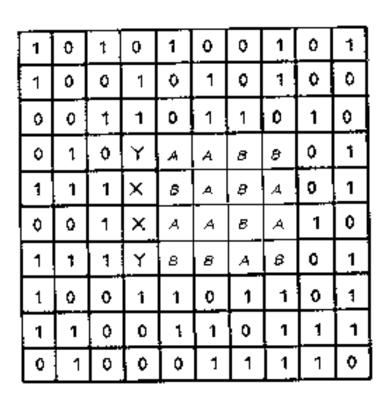
Stereoscopic card

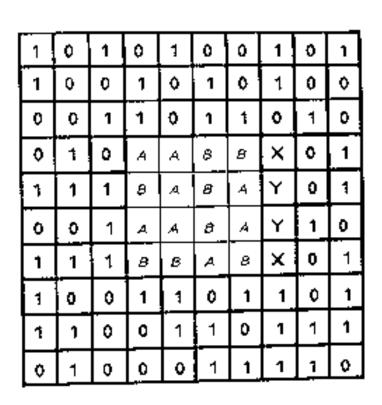


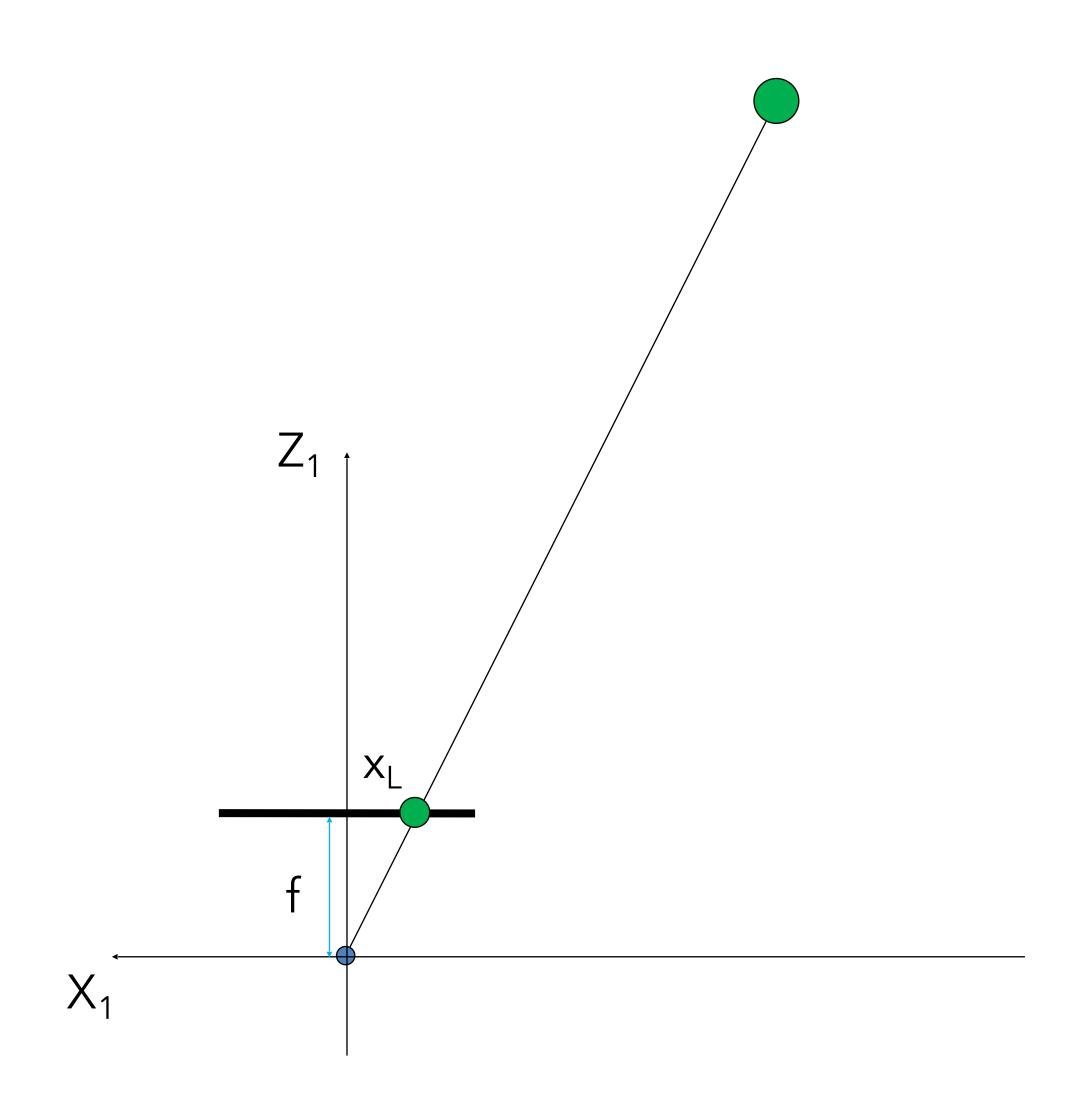
Brewster stereoscope

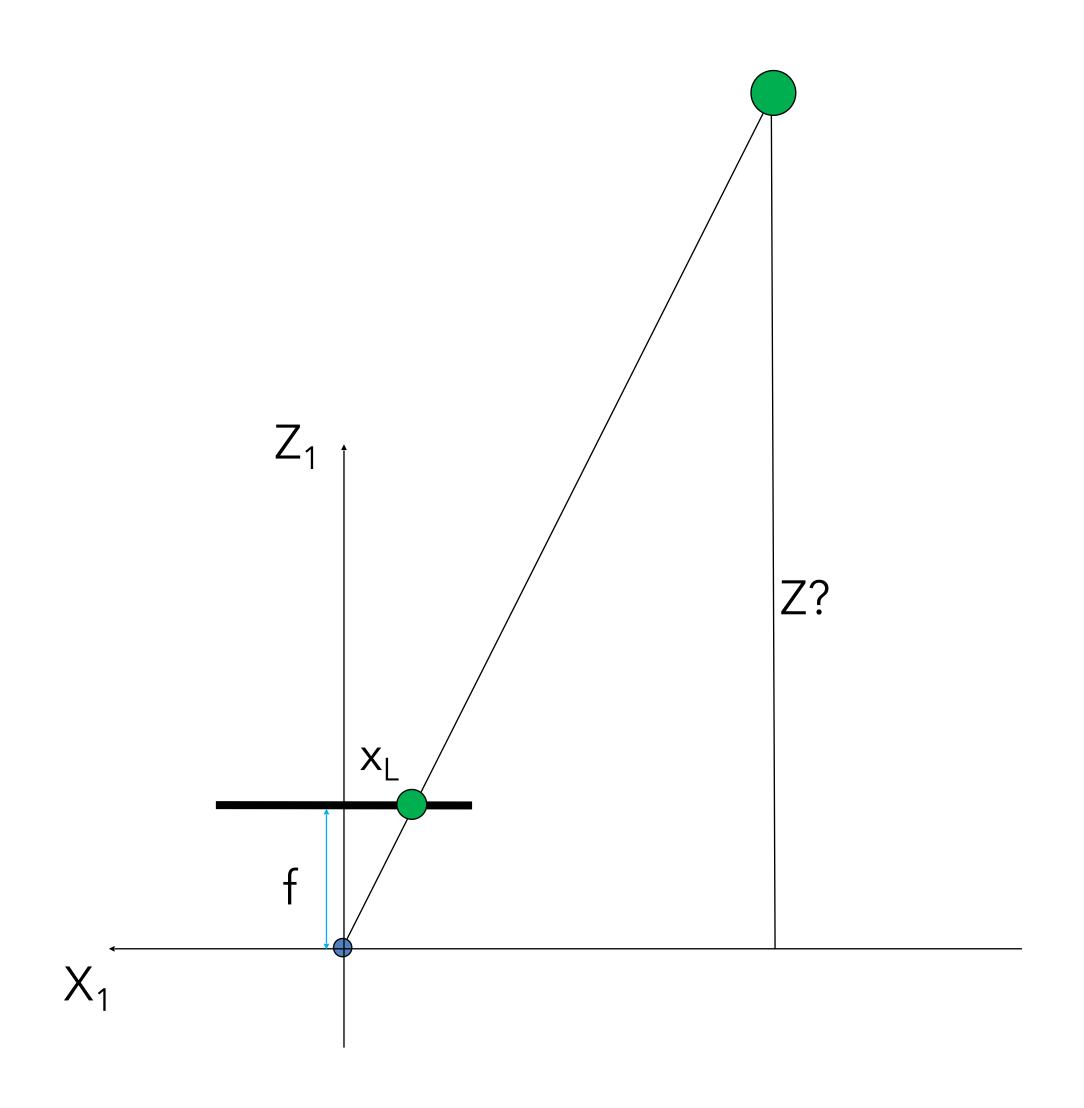
Depth without objects

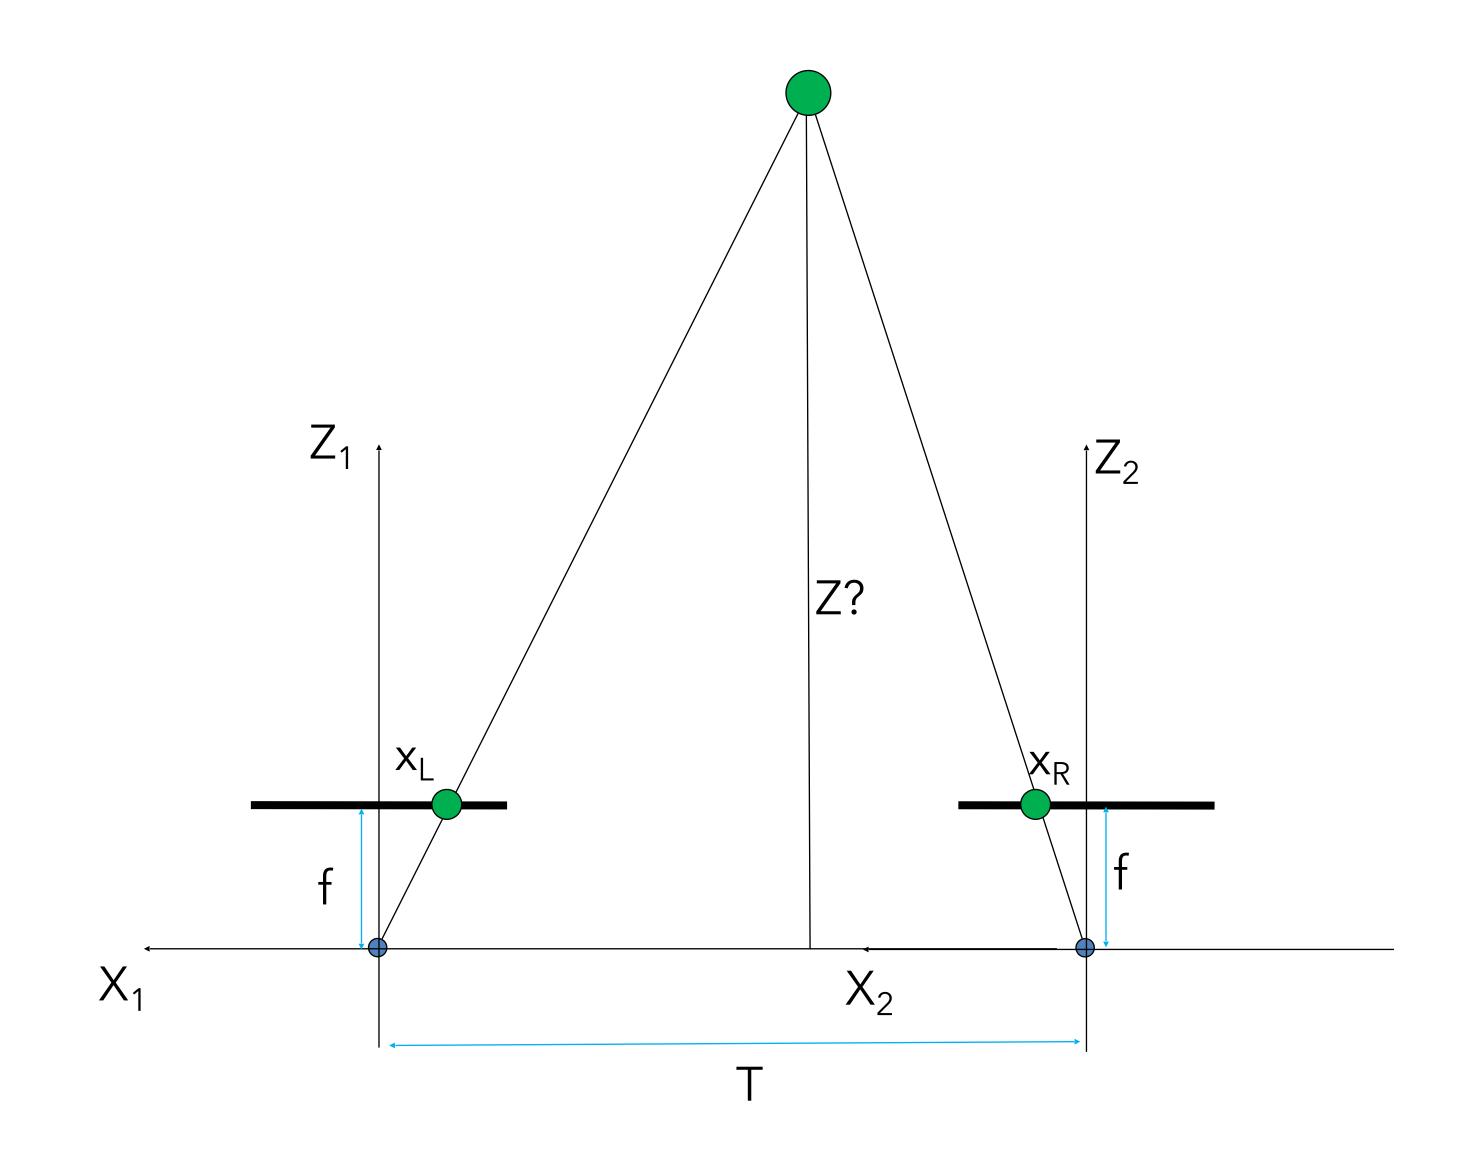


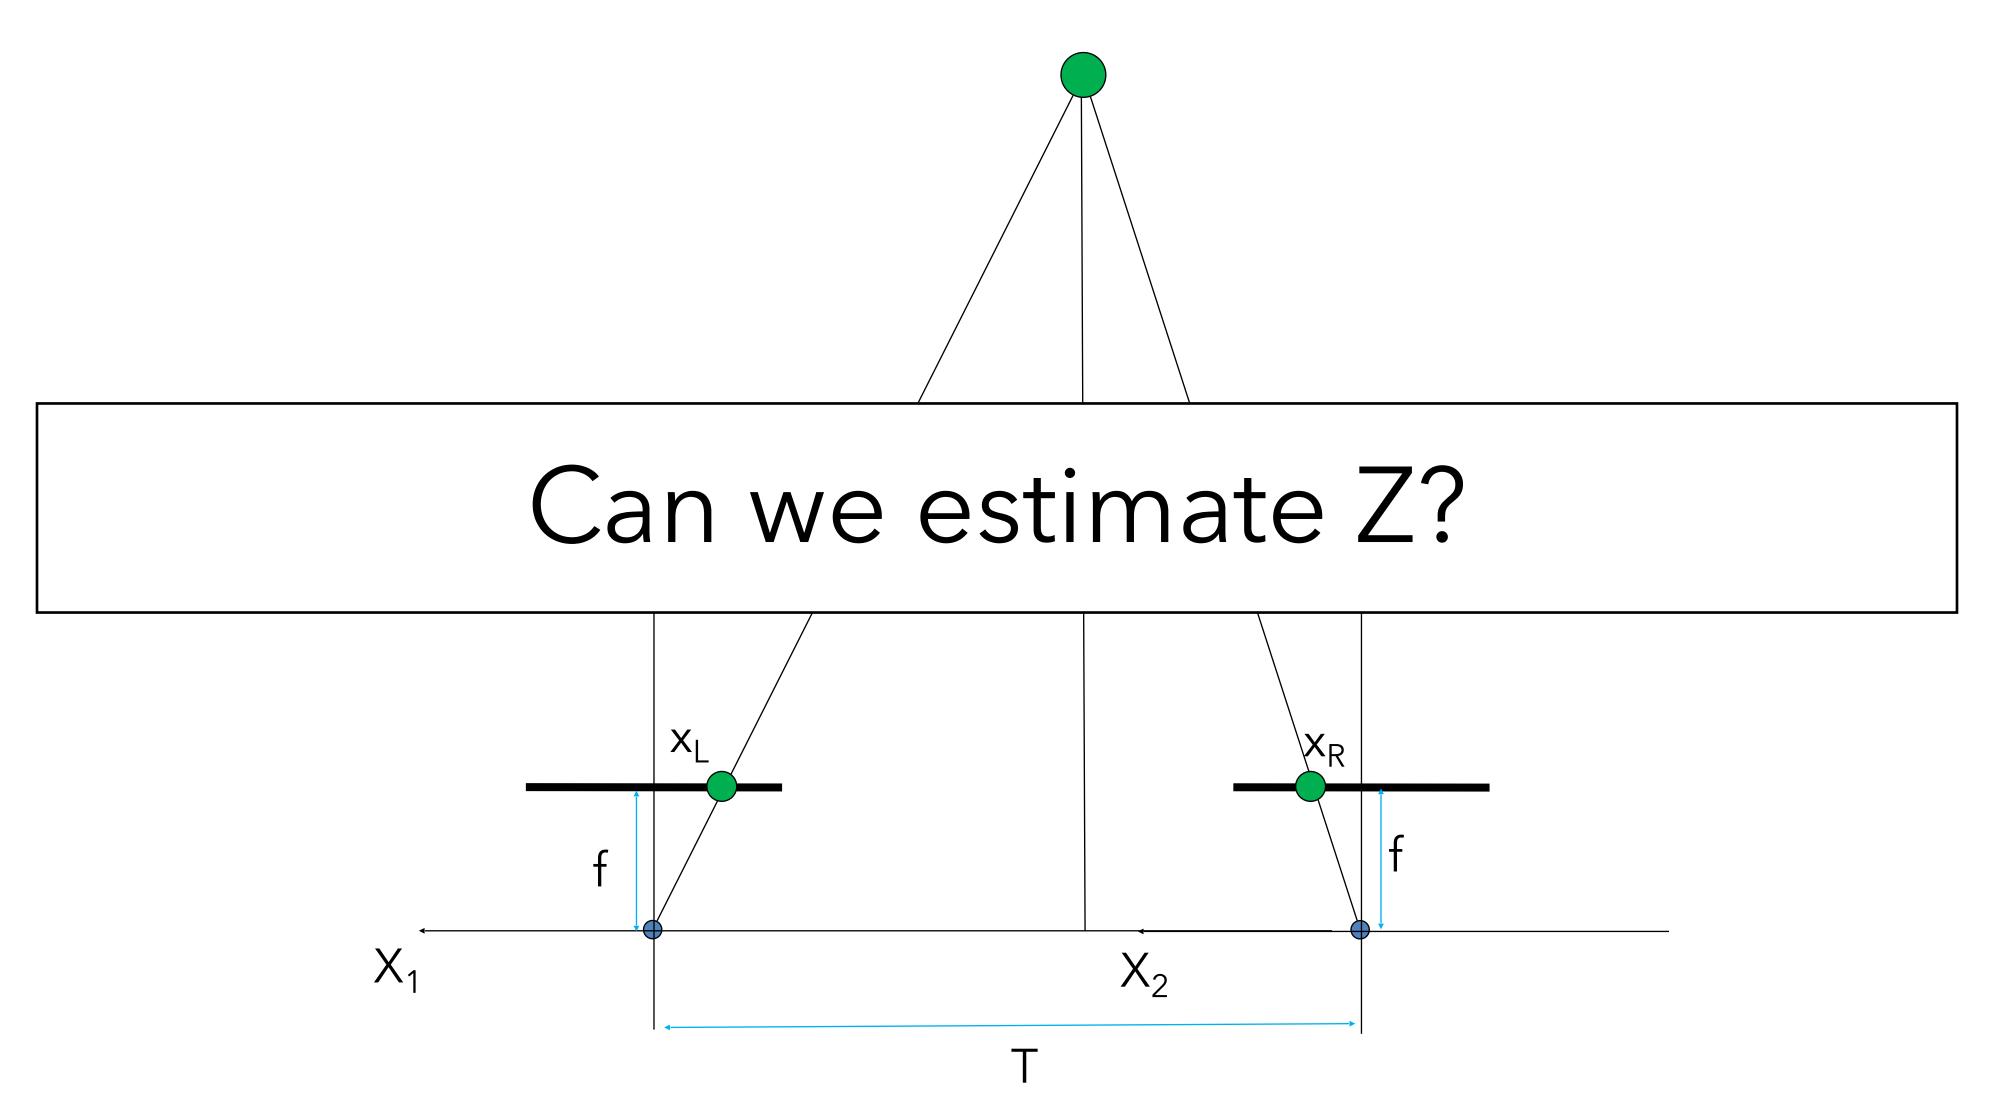


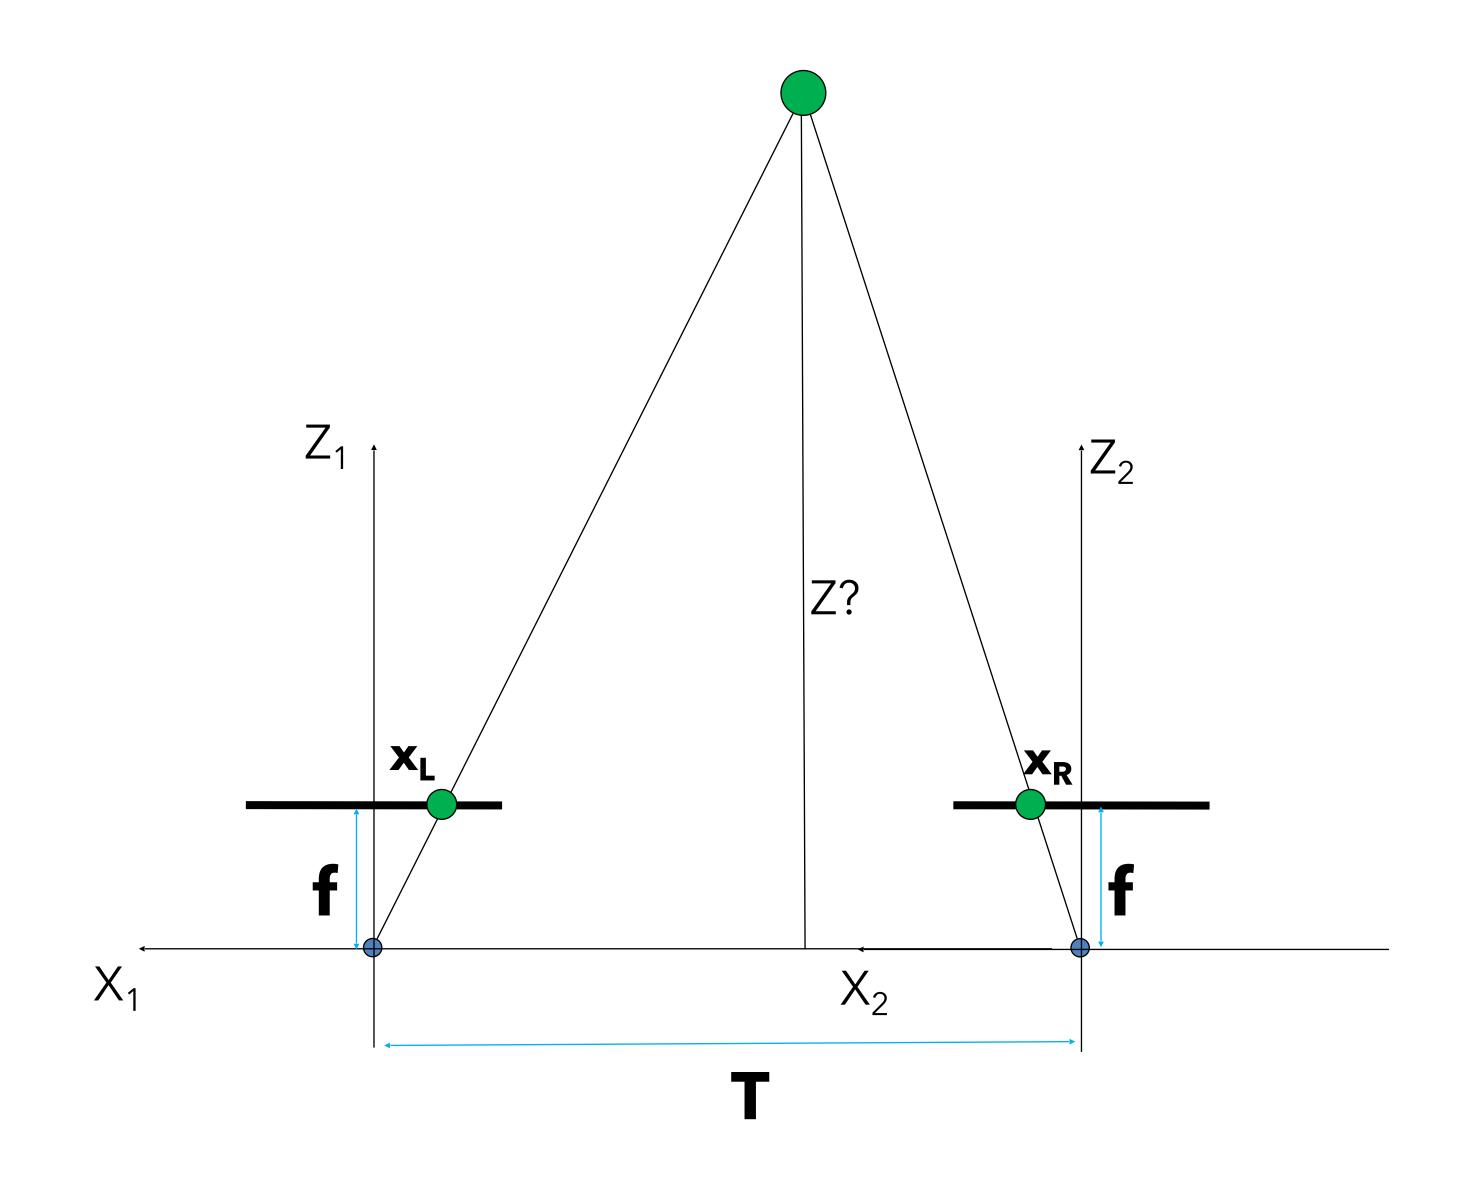


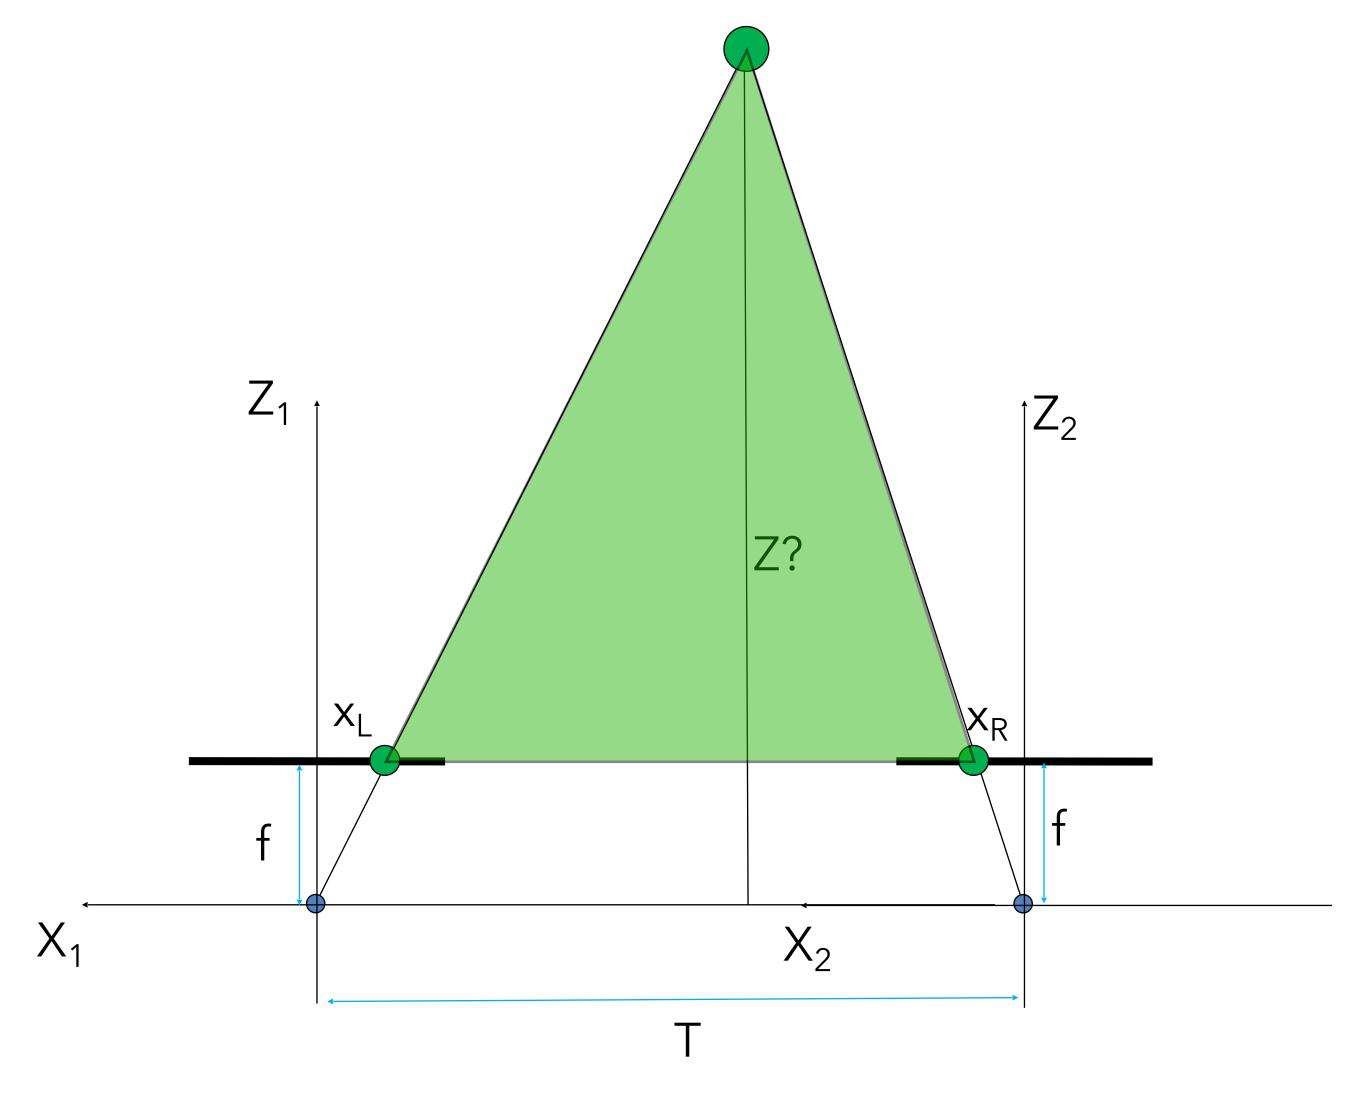








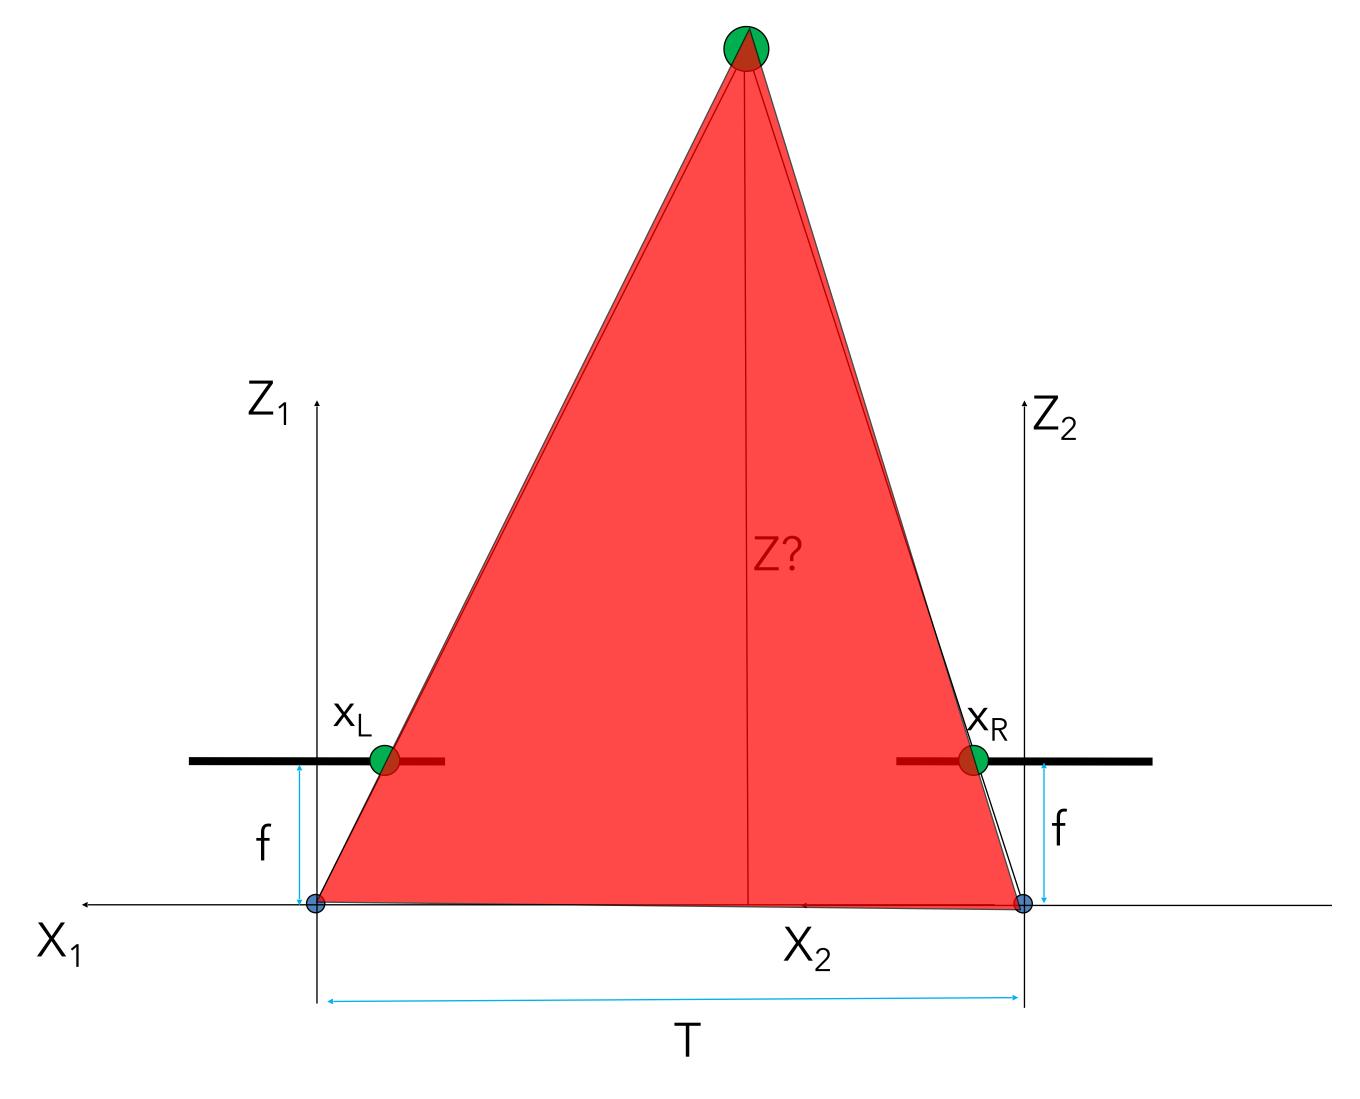




Similar triangles!

34

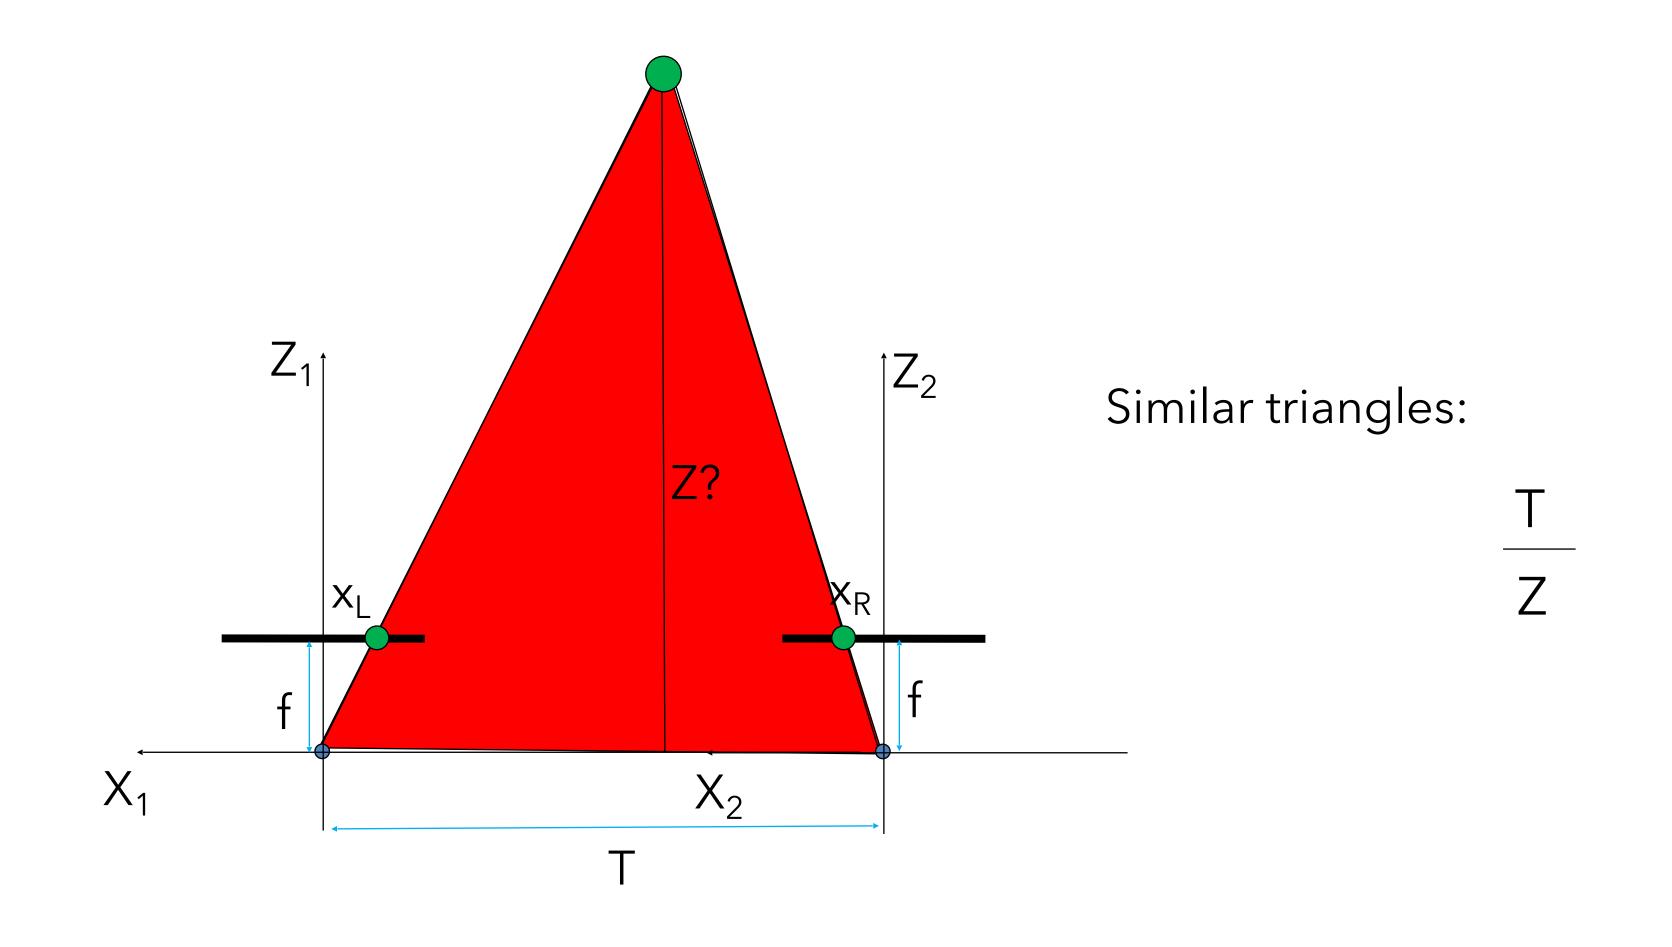
Source: Torralba, Isola, Freeman



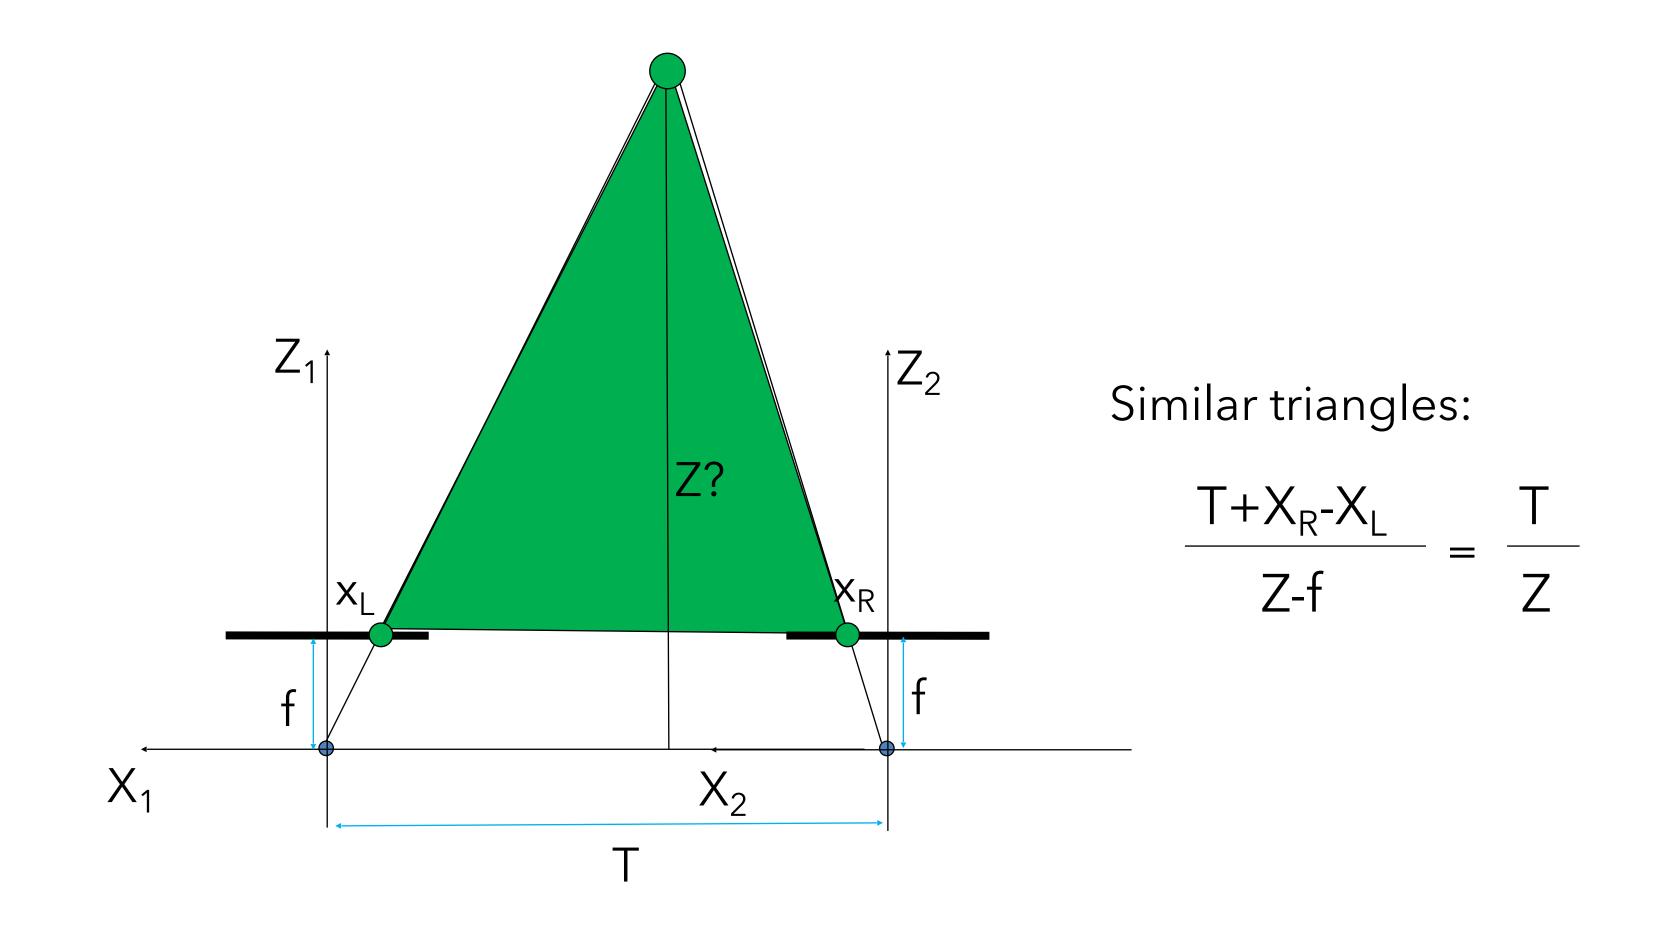
Similar triangles!

35

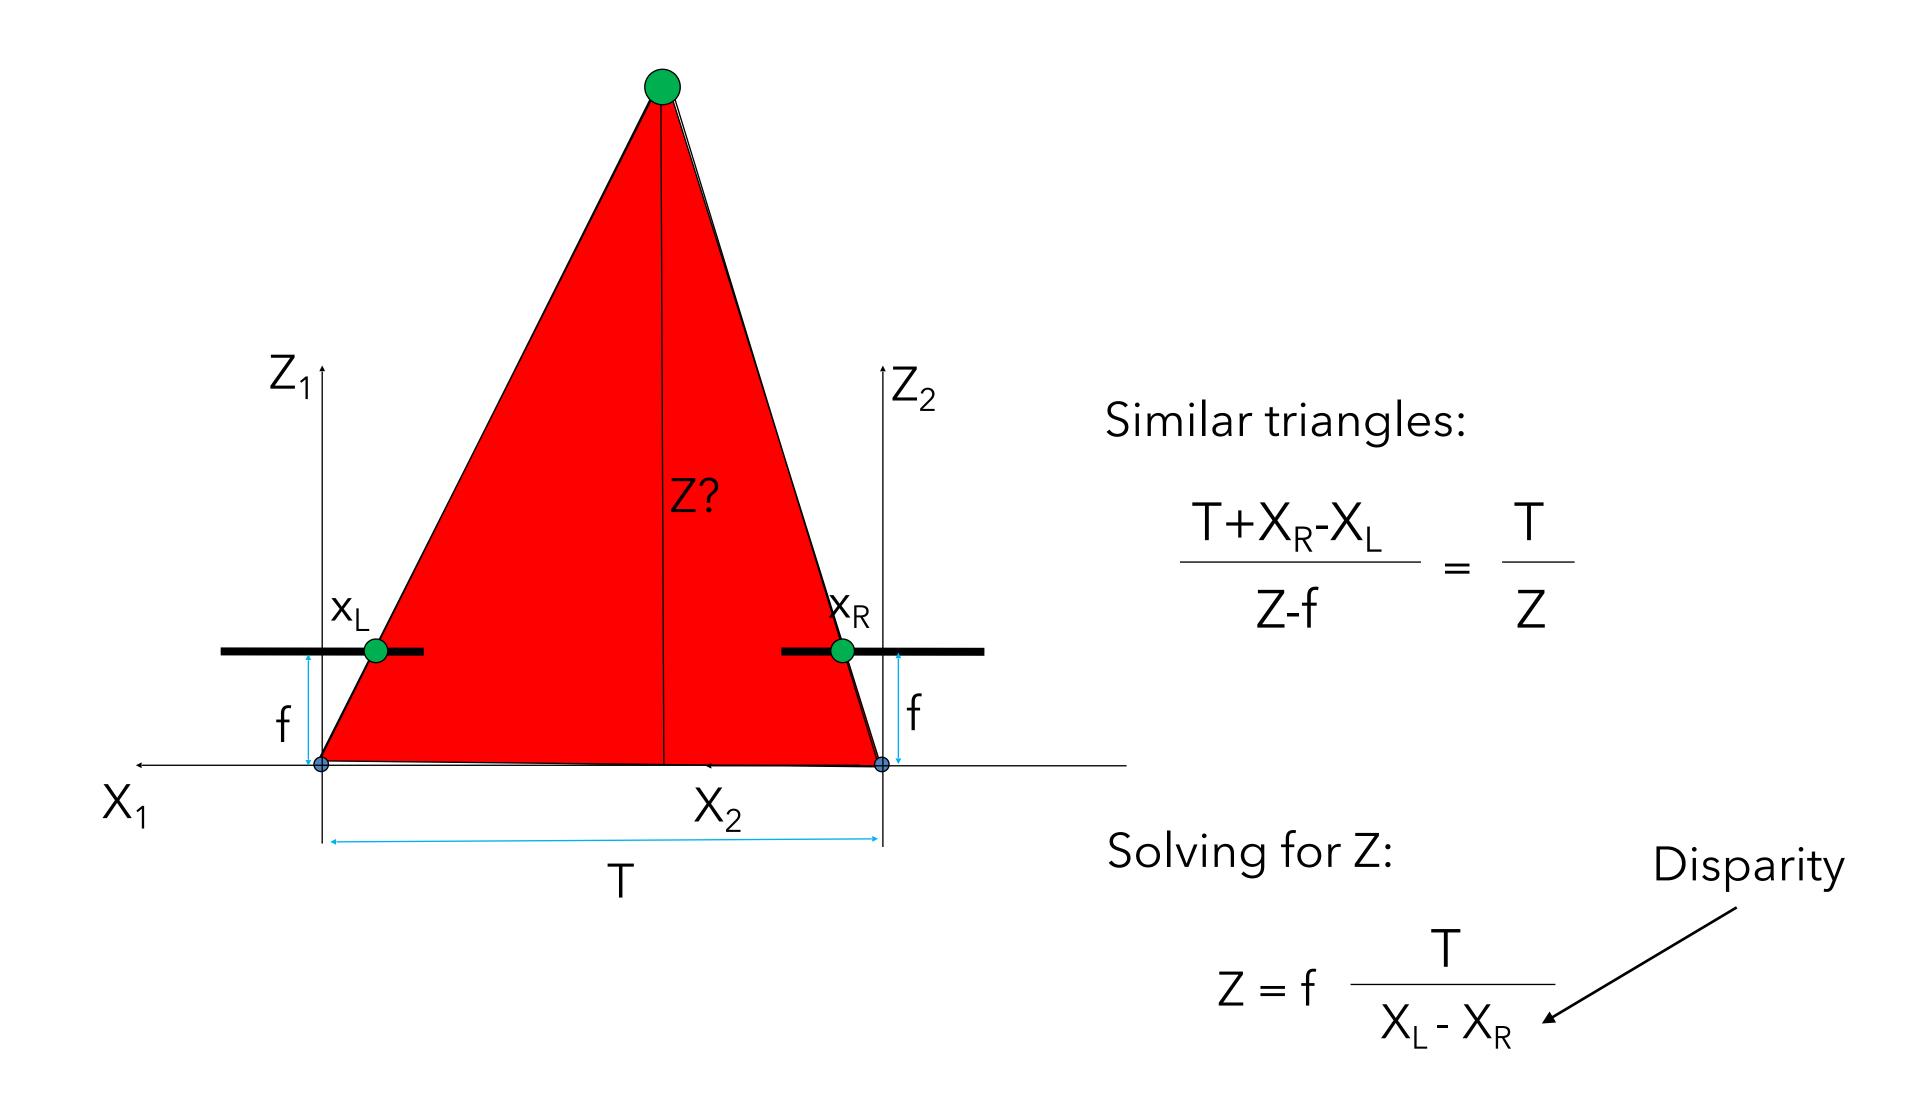
Source: Torralba, Isola, Freeman



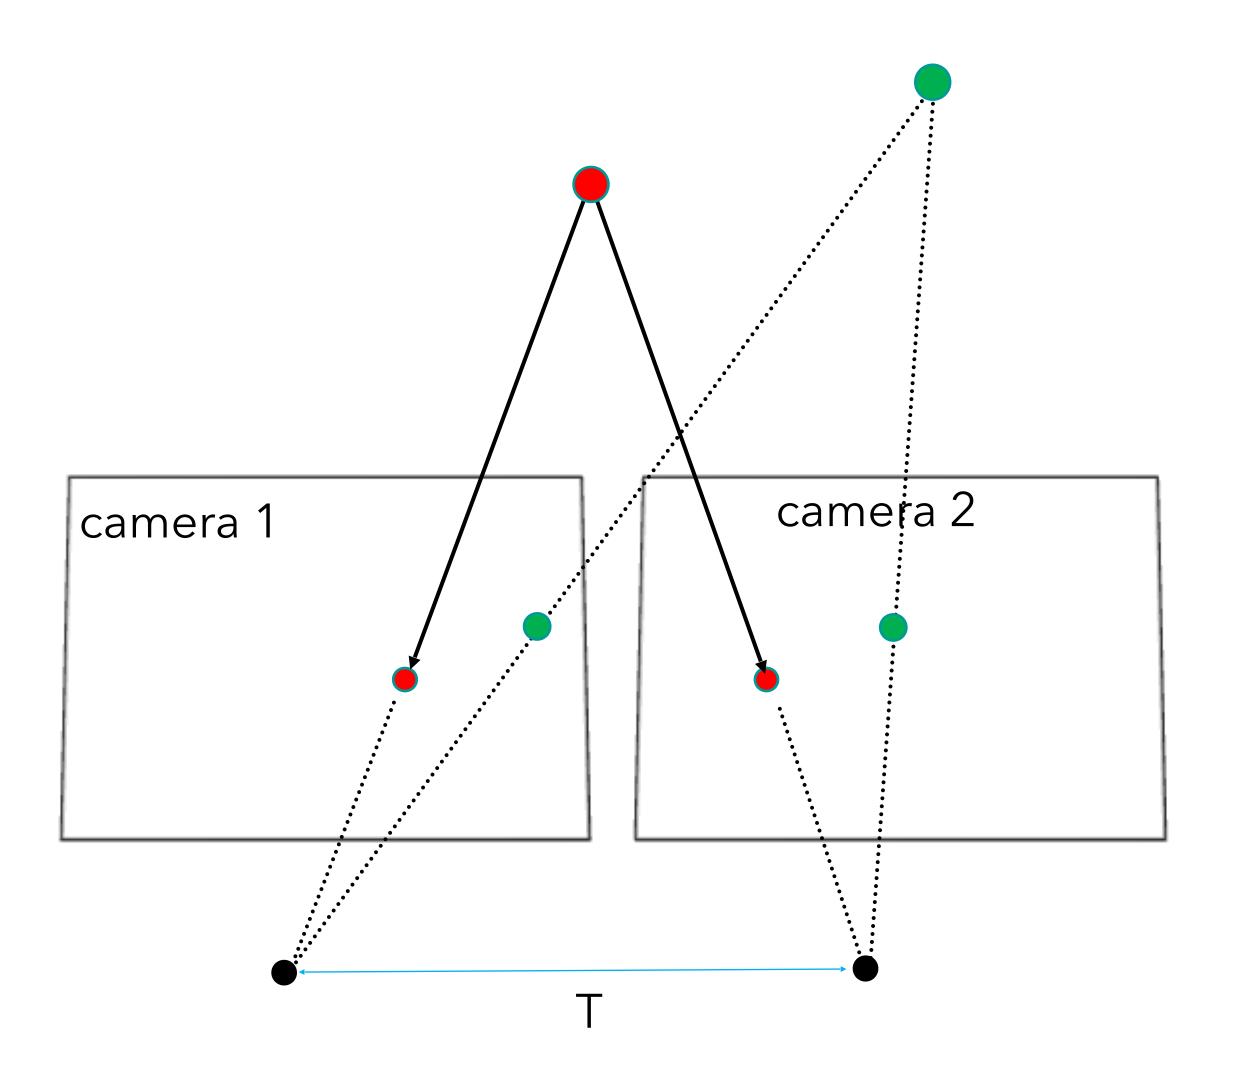
Geometry for a simple stereo system

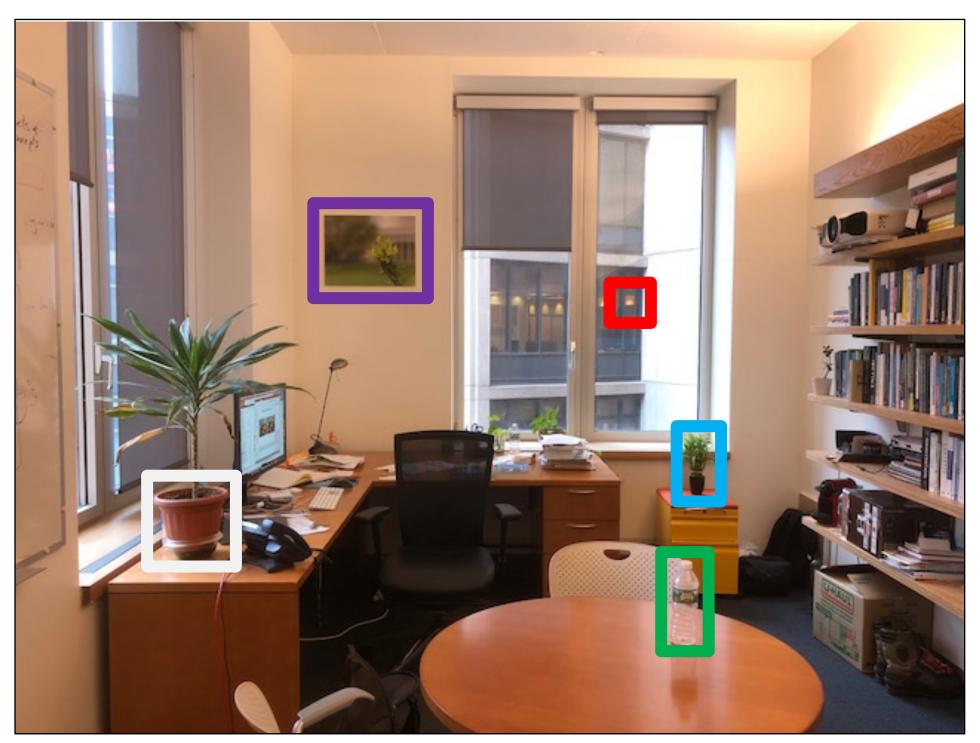


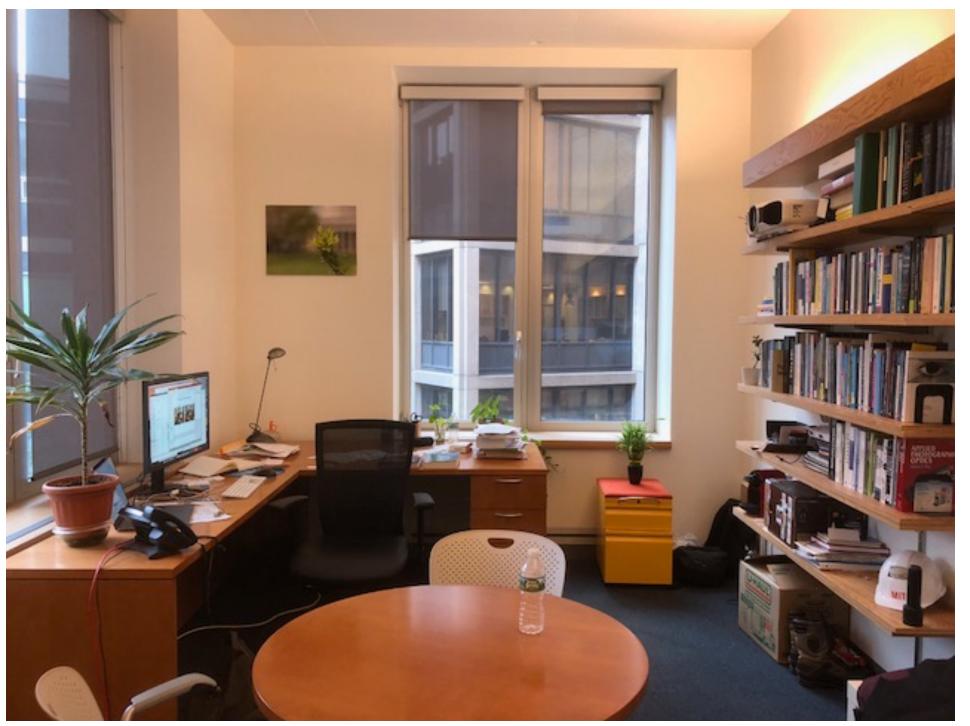
Geometry for a simple stereo system



In 3D



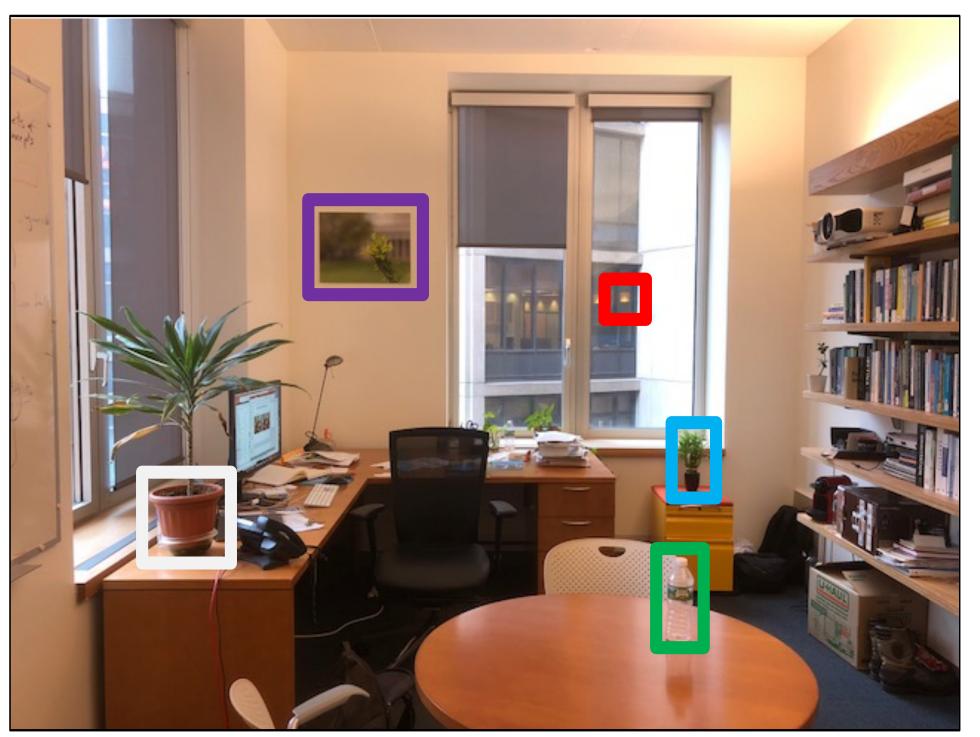


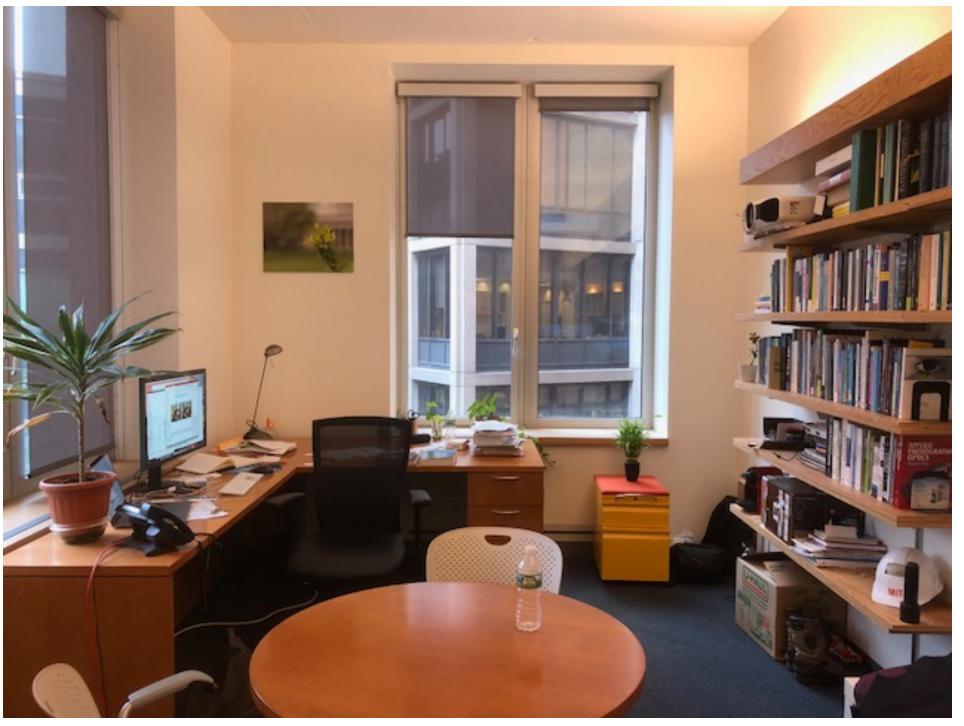


Left image

Right image

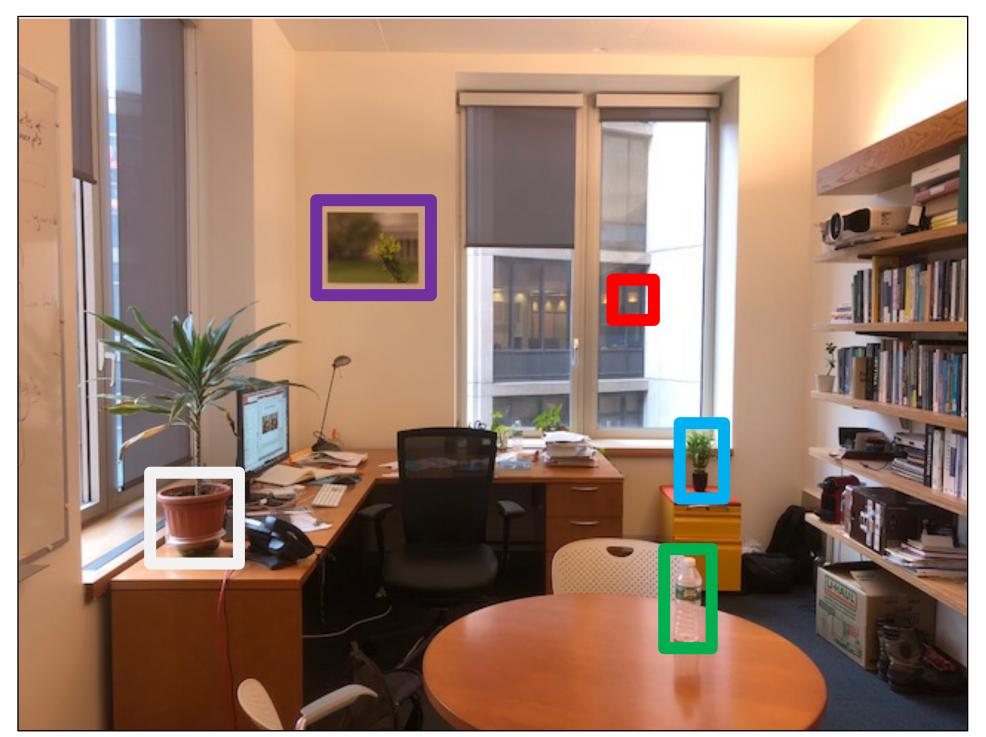
Second picture is ~1m to the right

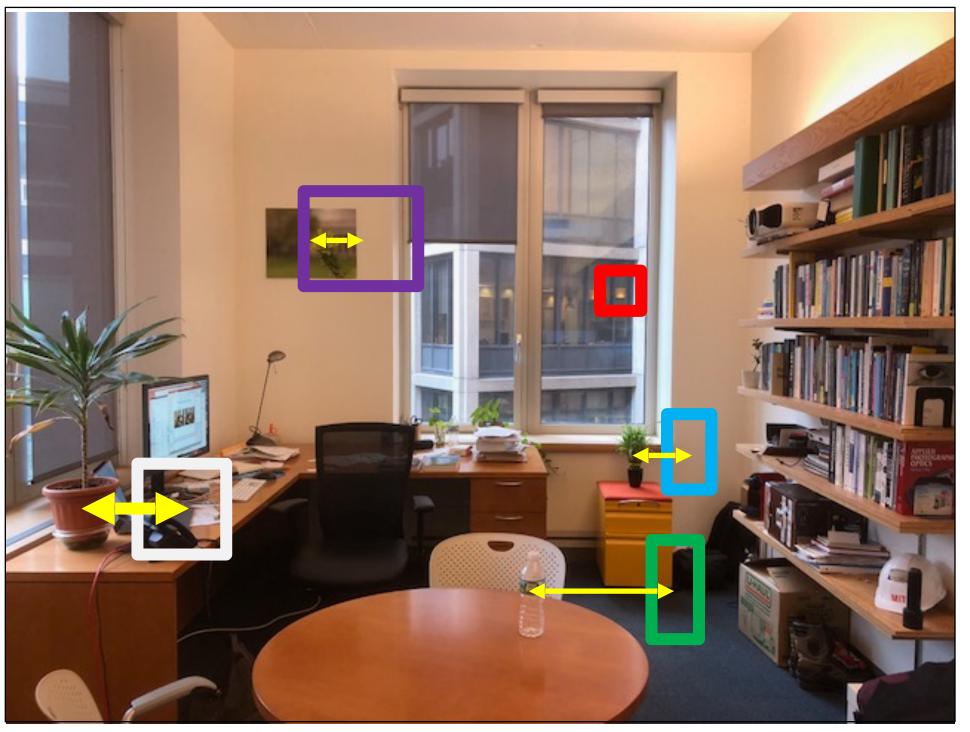




Left image

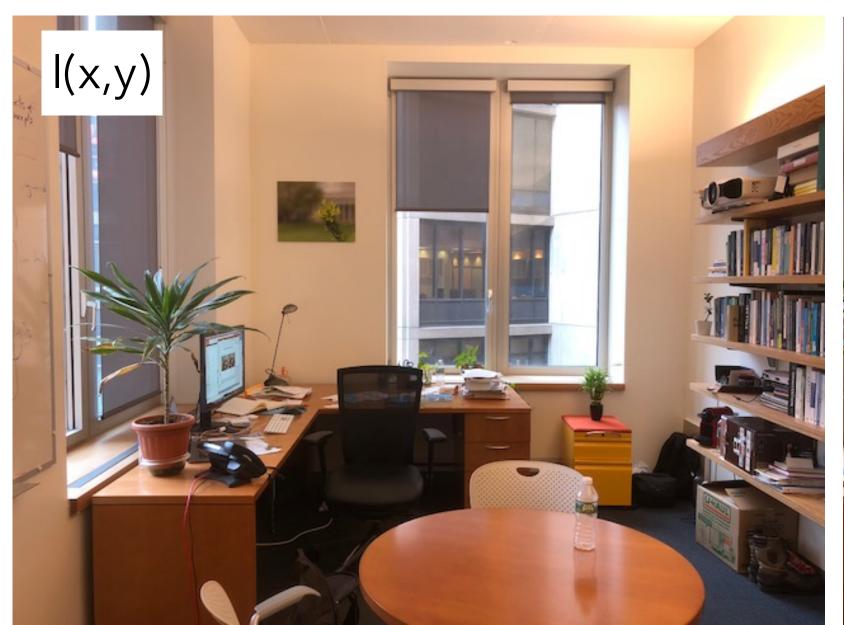
Right image

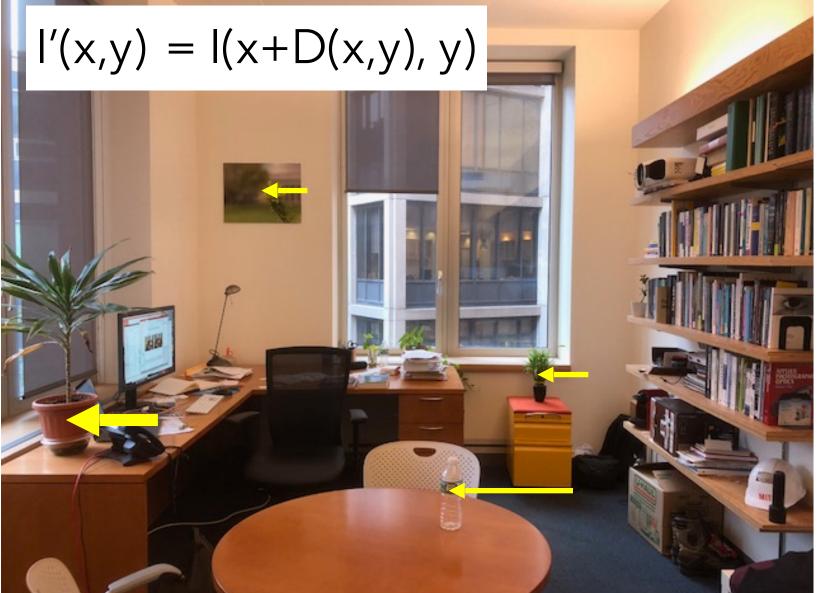




Left image

Right image





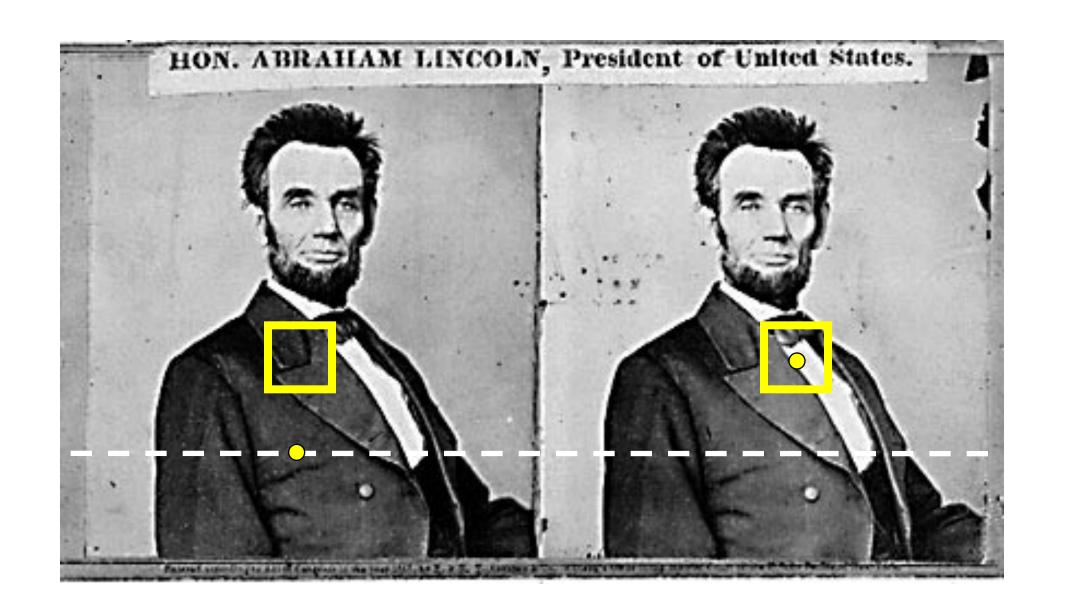
$$Z(x,y) = \frac{f}{D(x,y)}$$

Finding correspondences



We only need to search for matches along horizontal lines.

Basic stereo algorithm

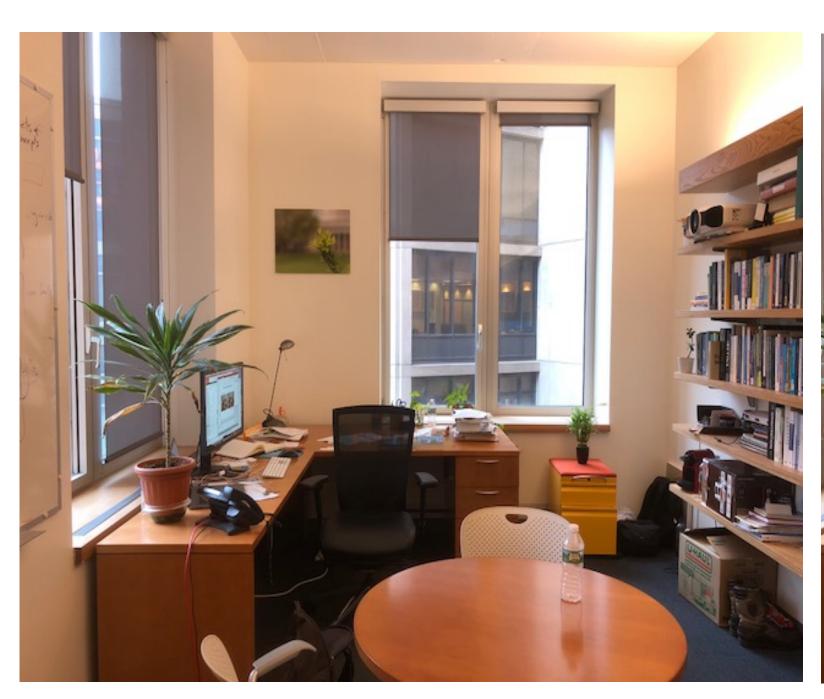


For each "epipolar line"

For each pixel in the left image

- compare with every pixel on same epipolar line in right image
- pick pixel with minimum match cost

Computing disparity





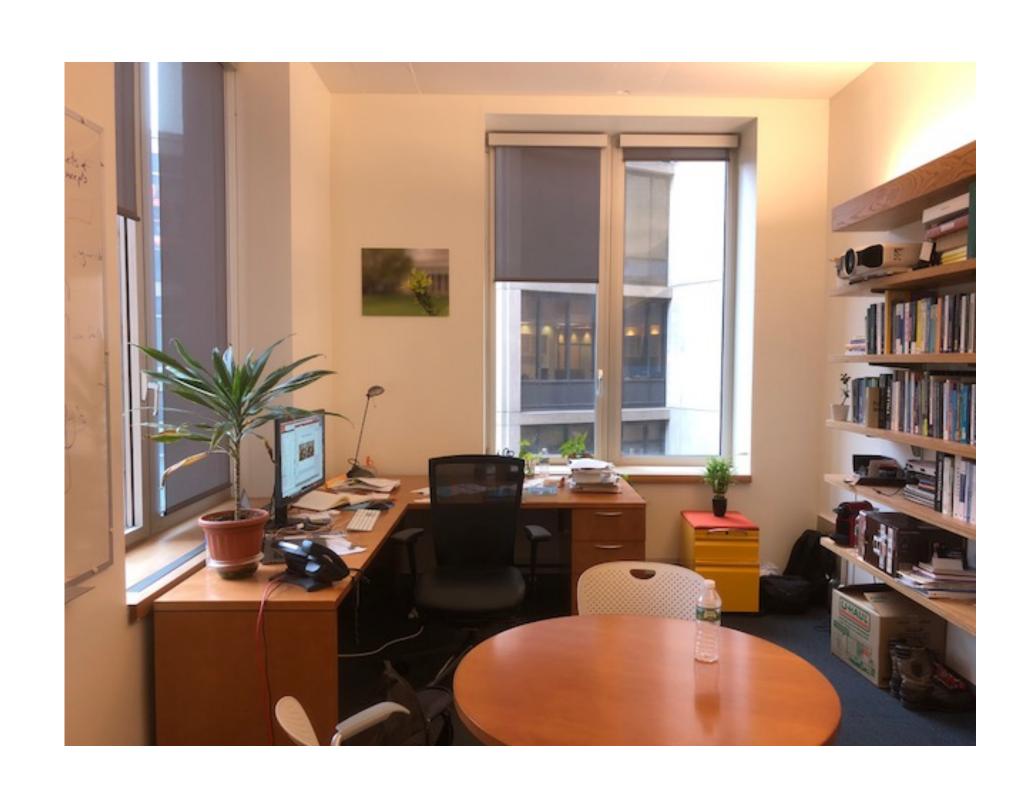
Computing disparity

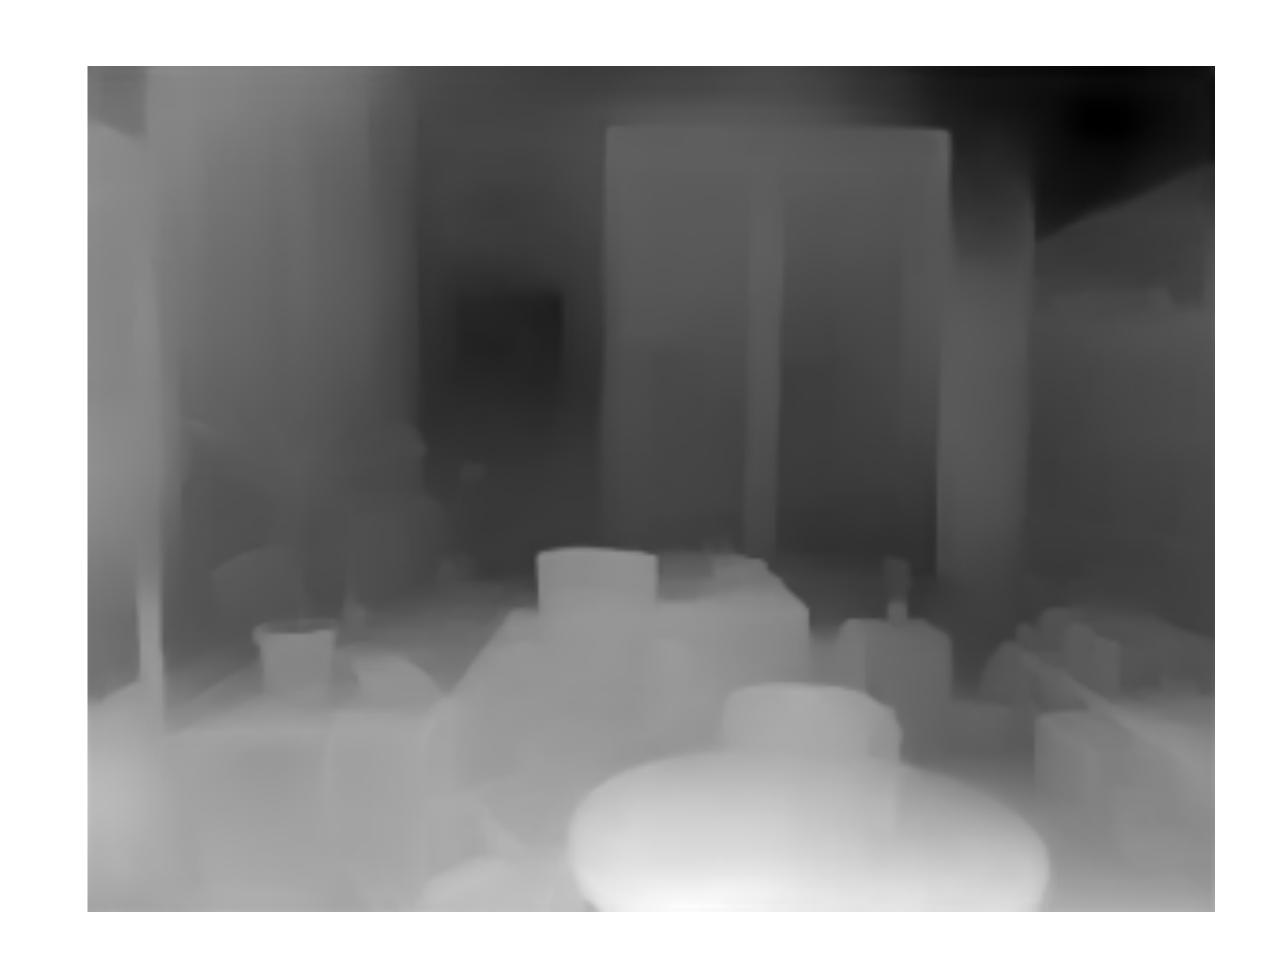




Semi-global matching [Hirschmüller 2008]

Can also learn depth from a single image





MegaDepth: Learning Single-View Depth Prediction from Internet Photos

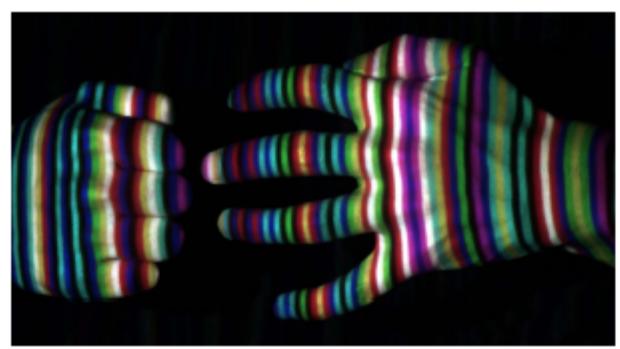
Zhengqi Li Noah Snavely
Department of Computer Science & Cornell Tech, Cornell University

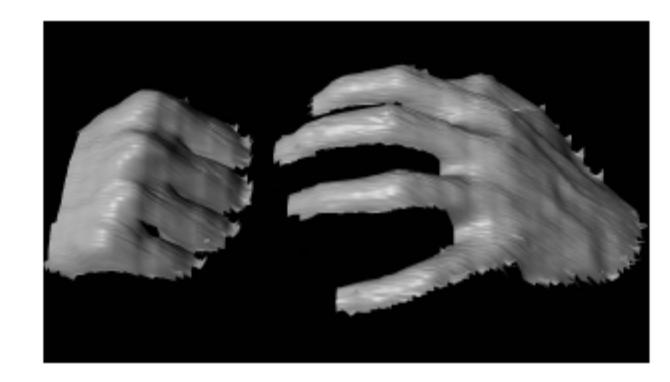
48

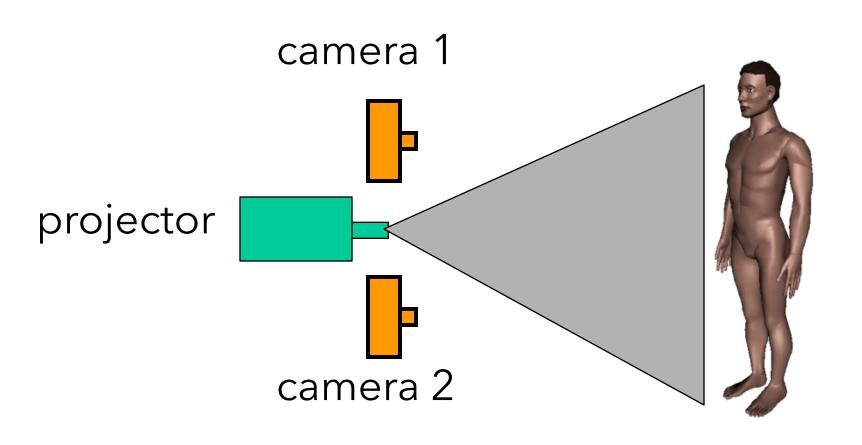
Source: Torralba, Isola, Freeman

Active stereo with structured light

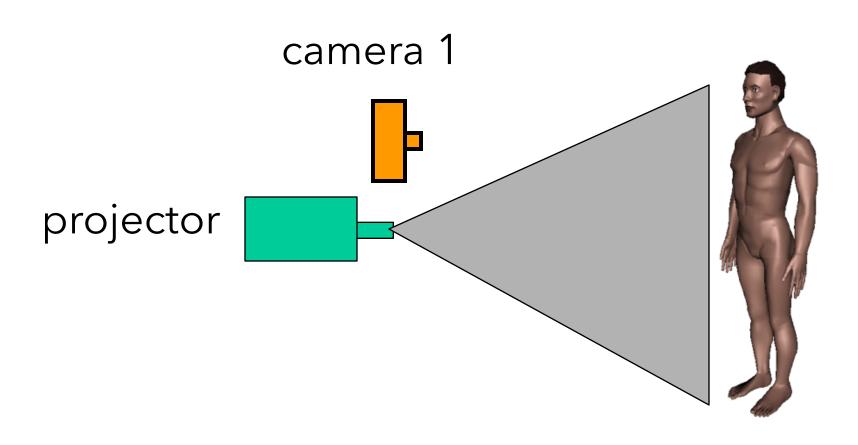








Easy-to-match pattern

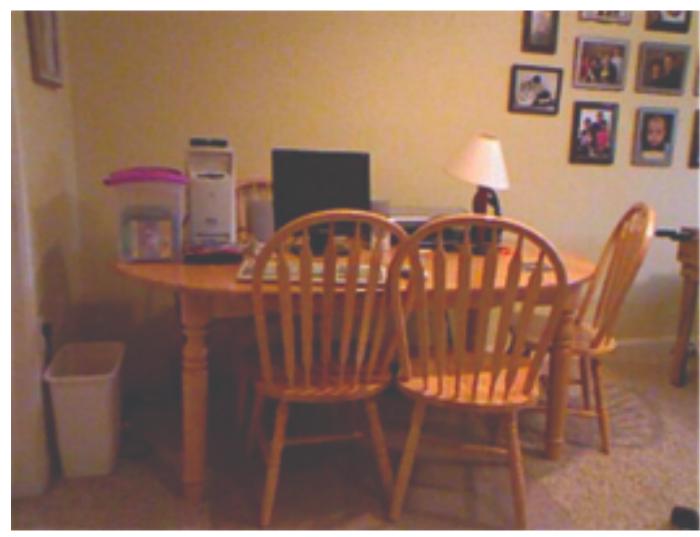


Do we really need the second camera?

RGB-D sensors

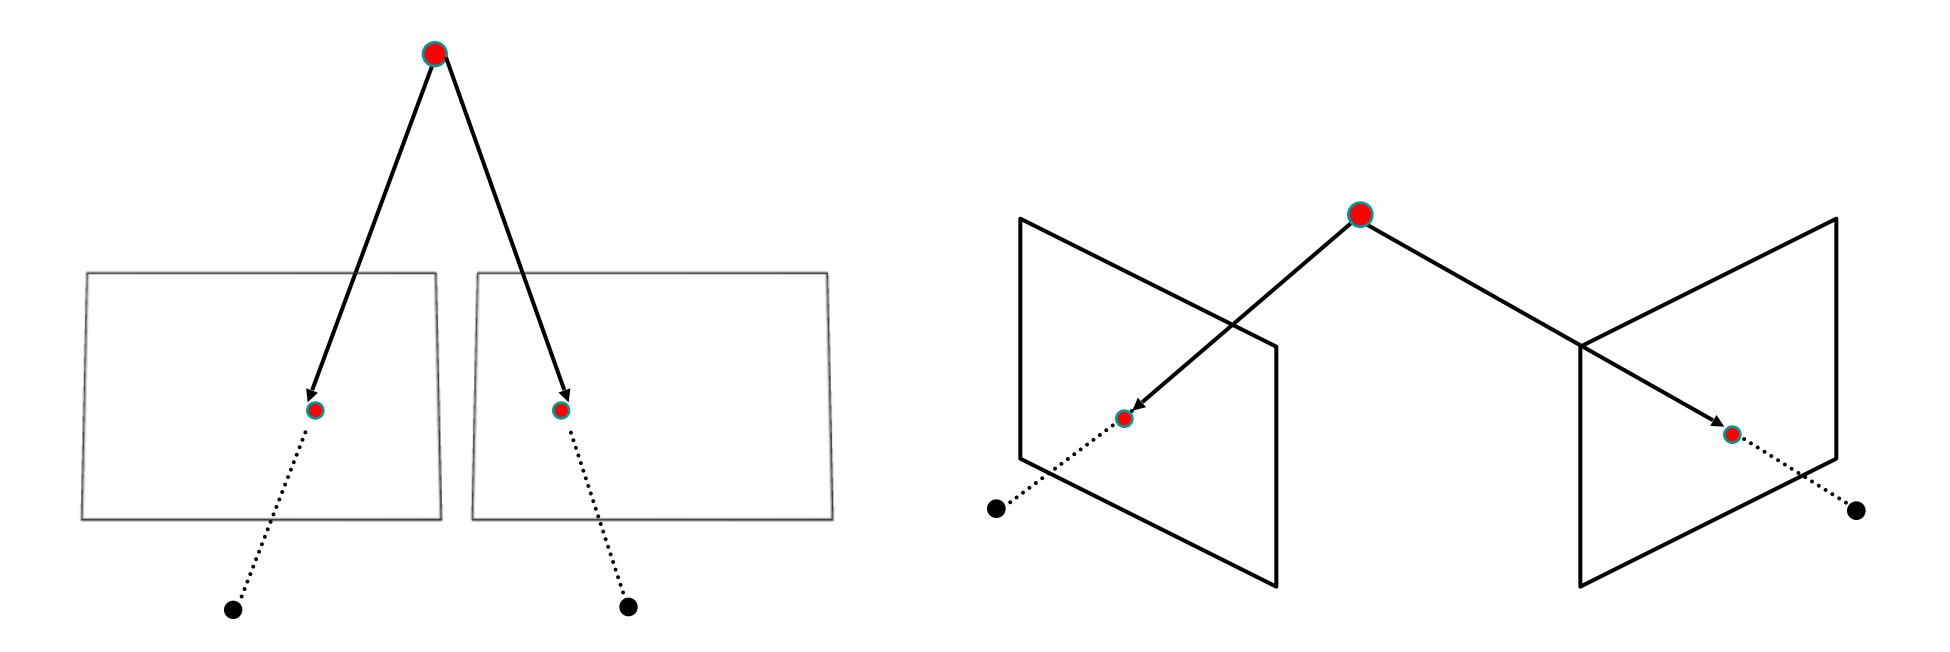


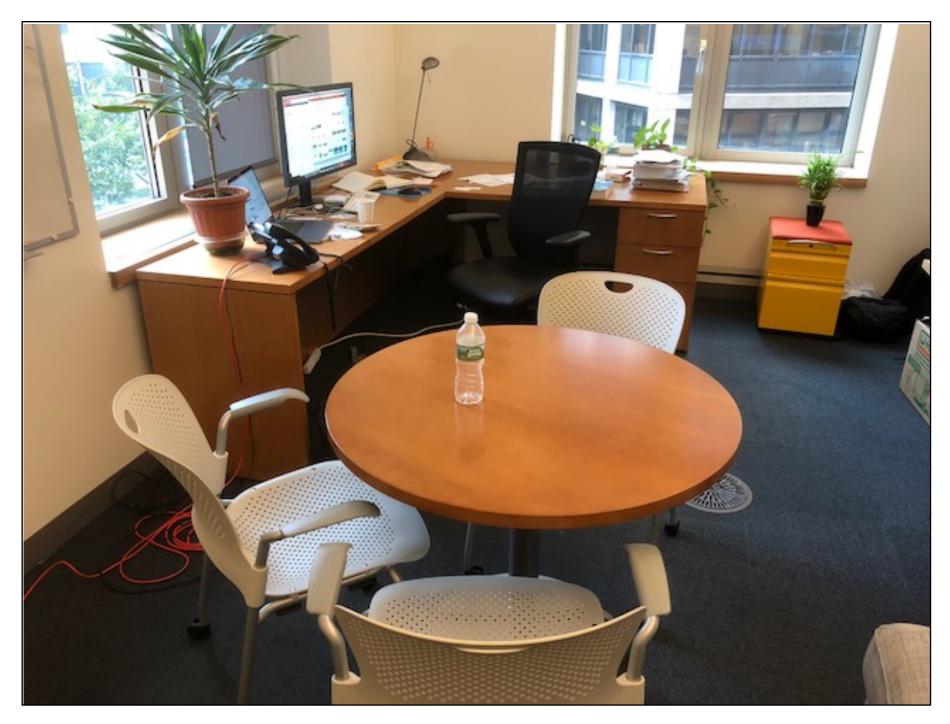


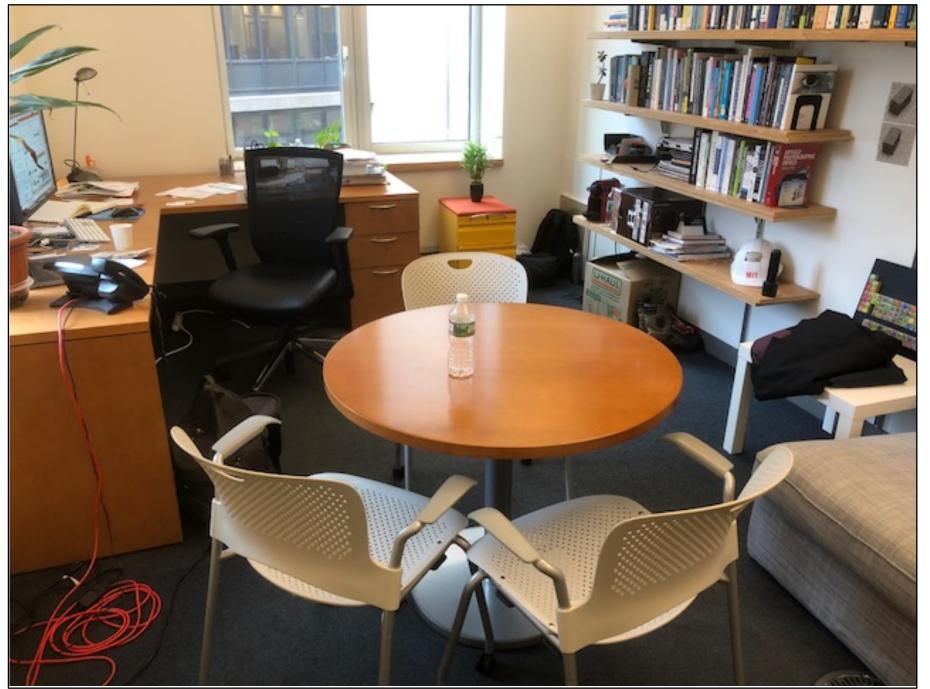


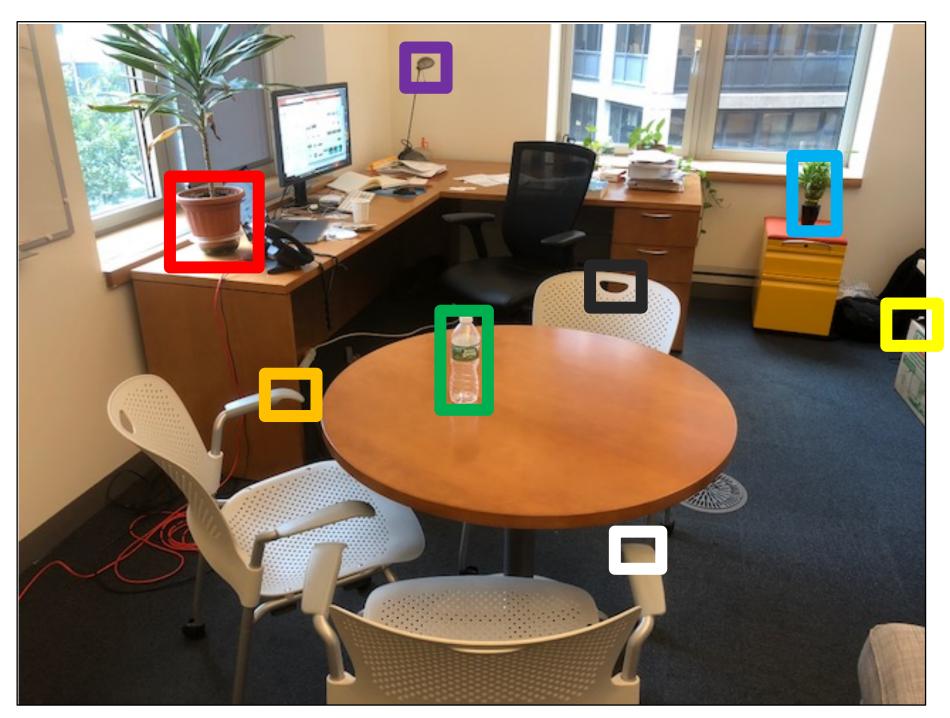
General case

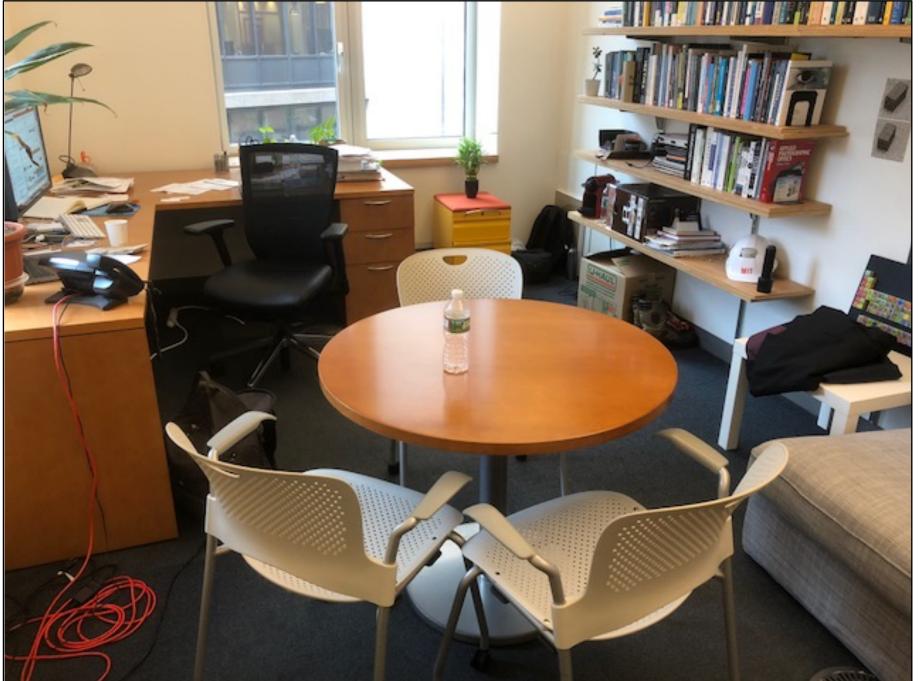
• The two cameras need not have parallel optical axes.



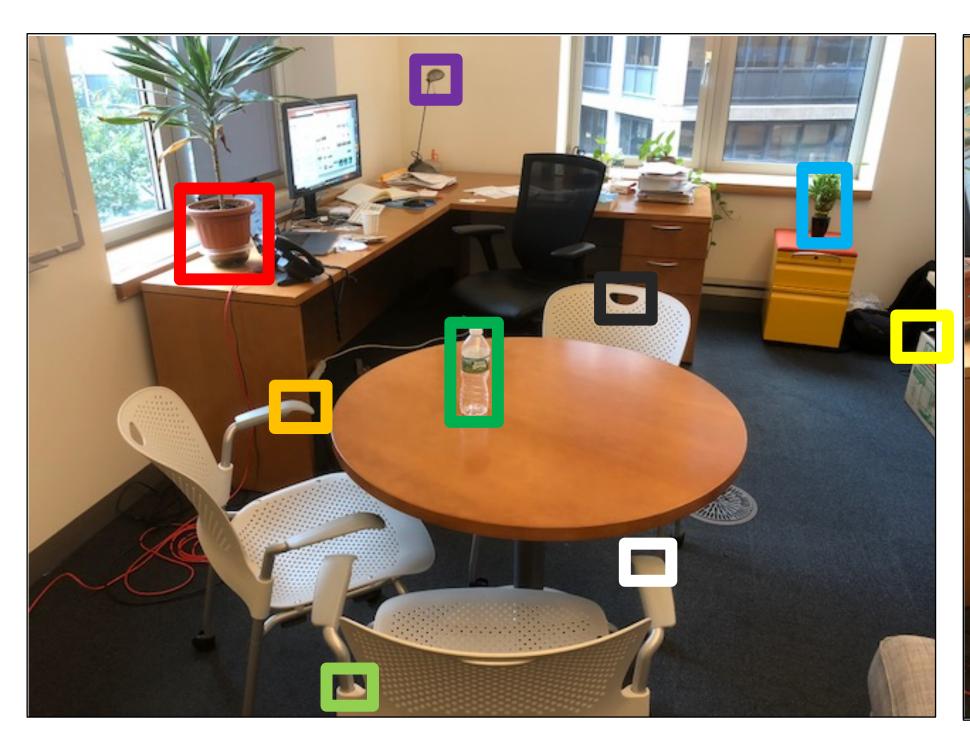


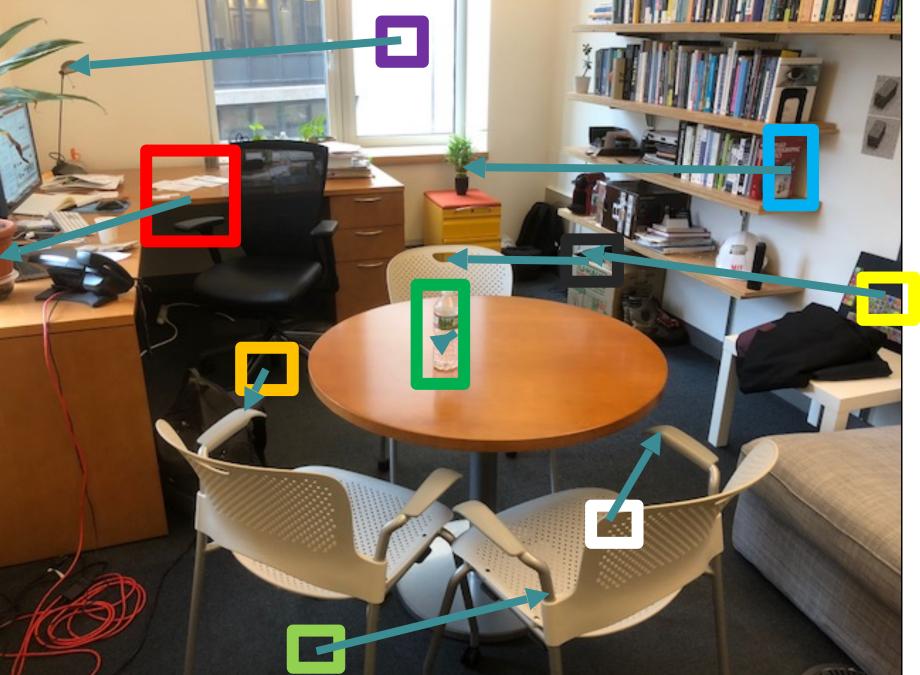




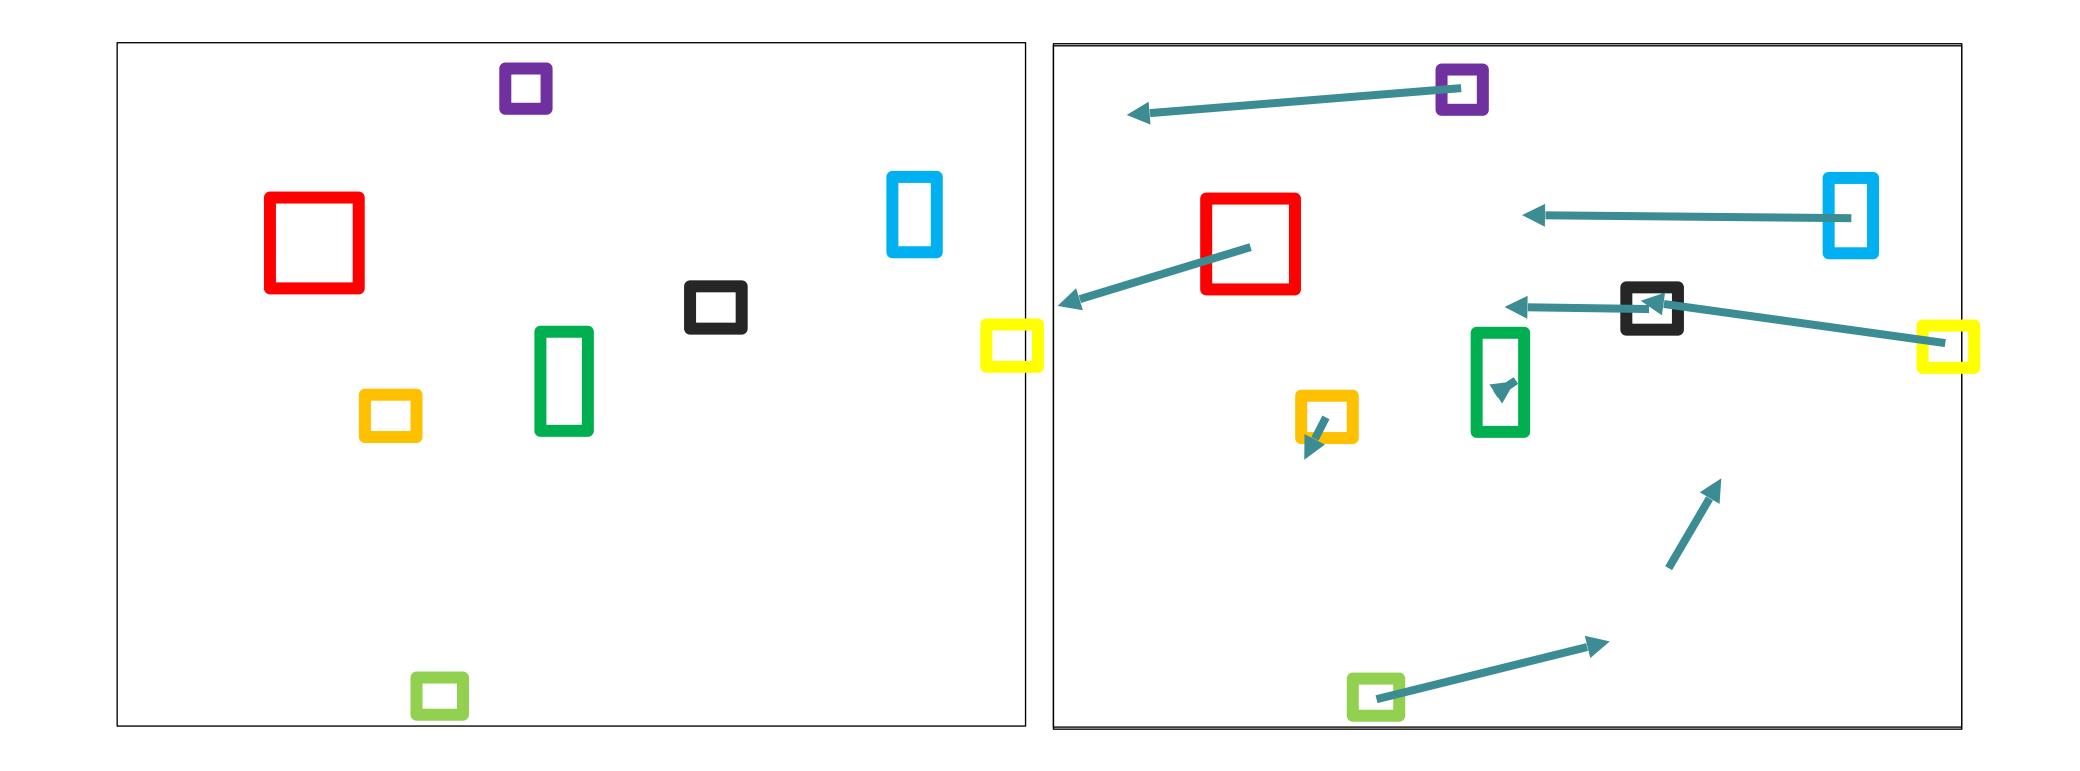


Do we need to search for matches only along horizontal lines?





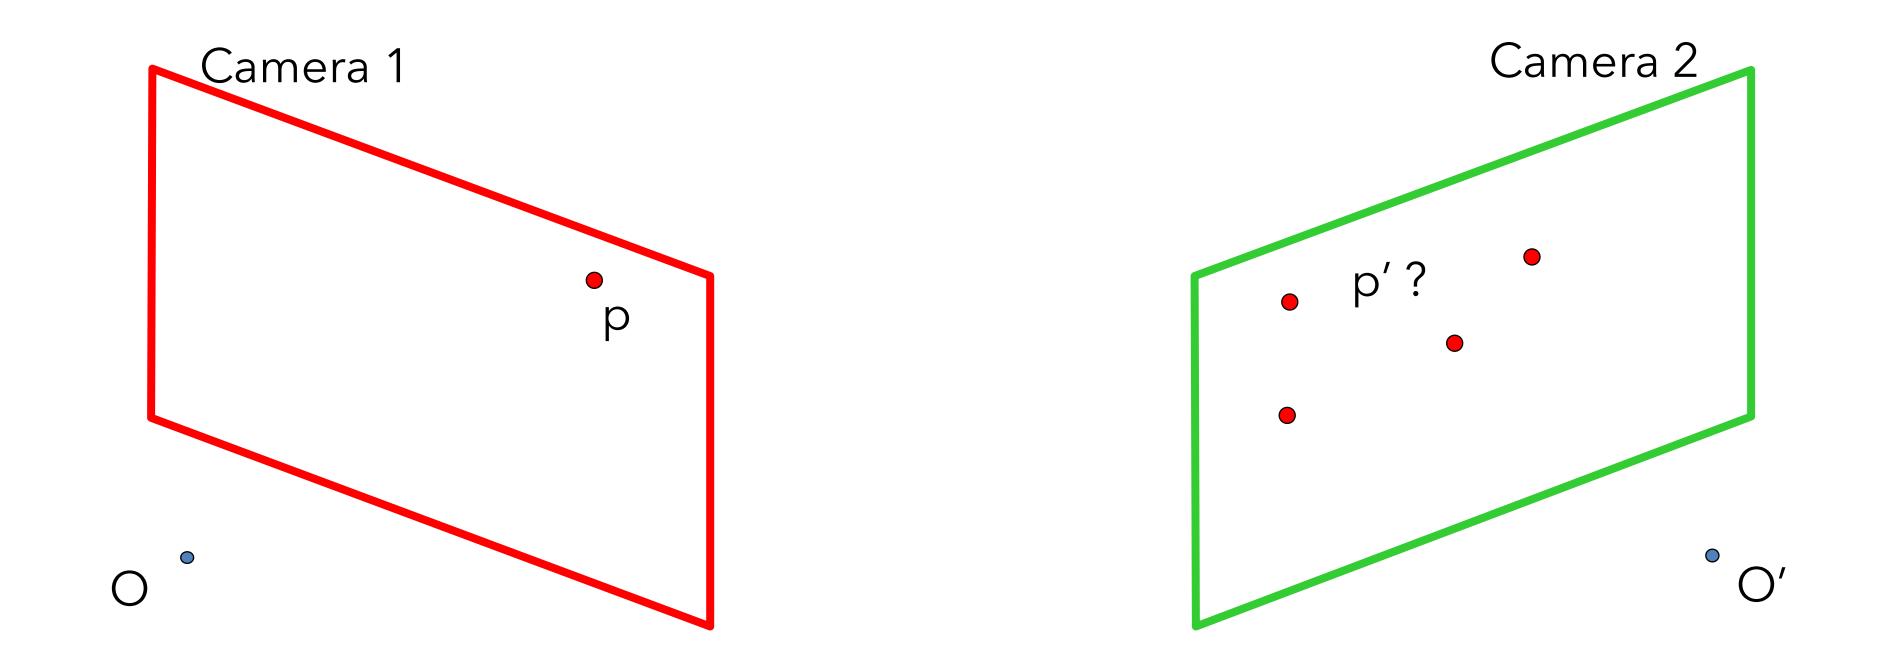
Do we need to search for matches only along horizontal lines?



Do we need to search for matches only along horizontal lines?

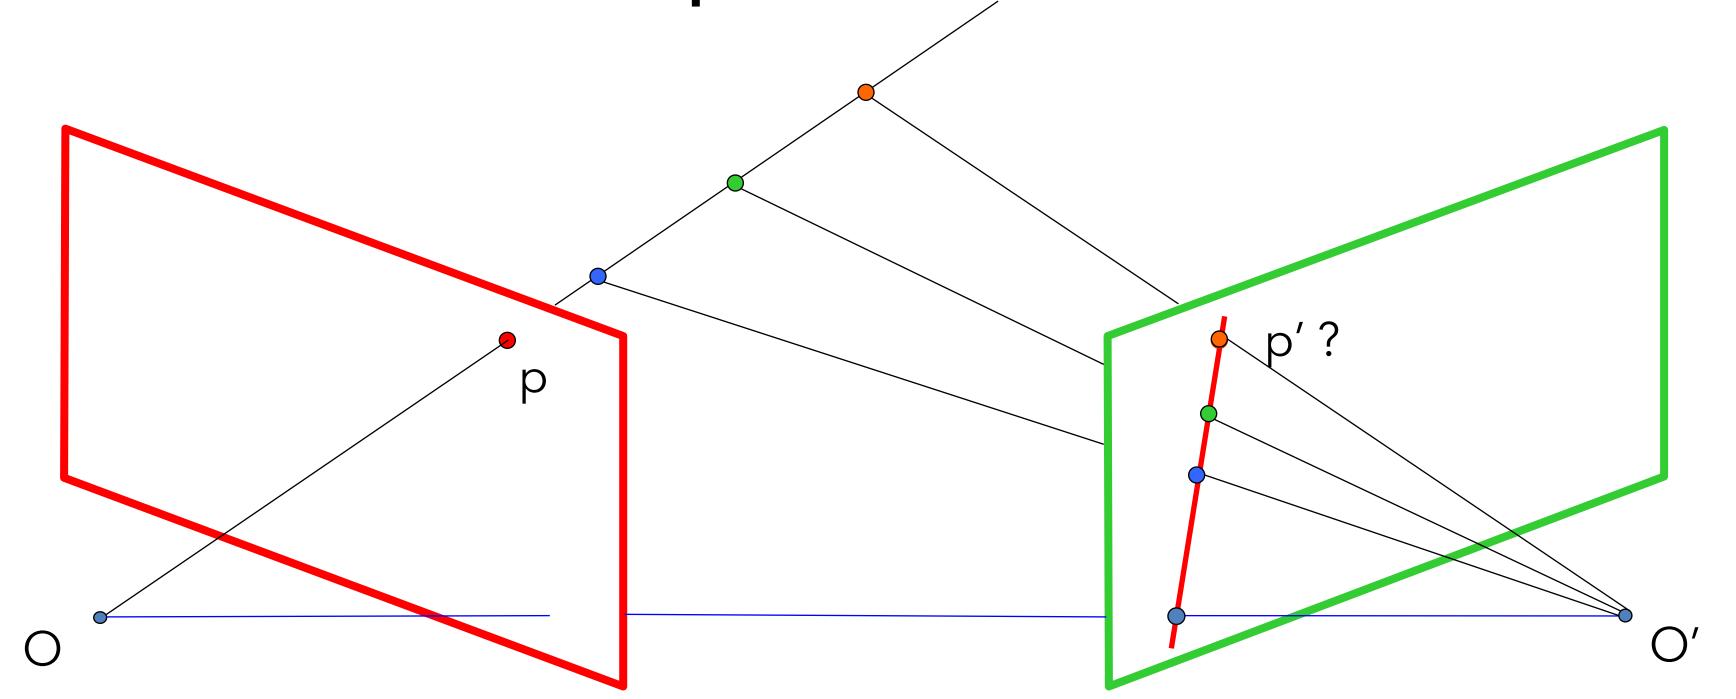
It looks like we might need to search everywhere... are there any constraints that can guide the search? 55

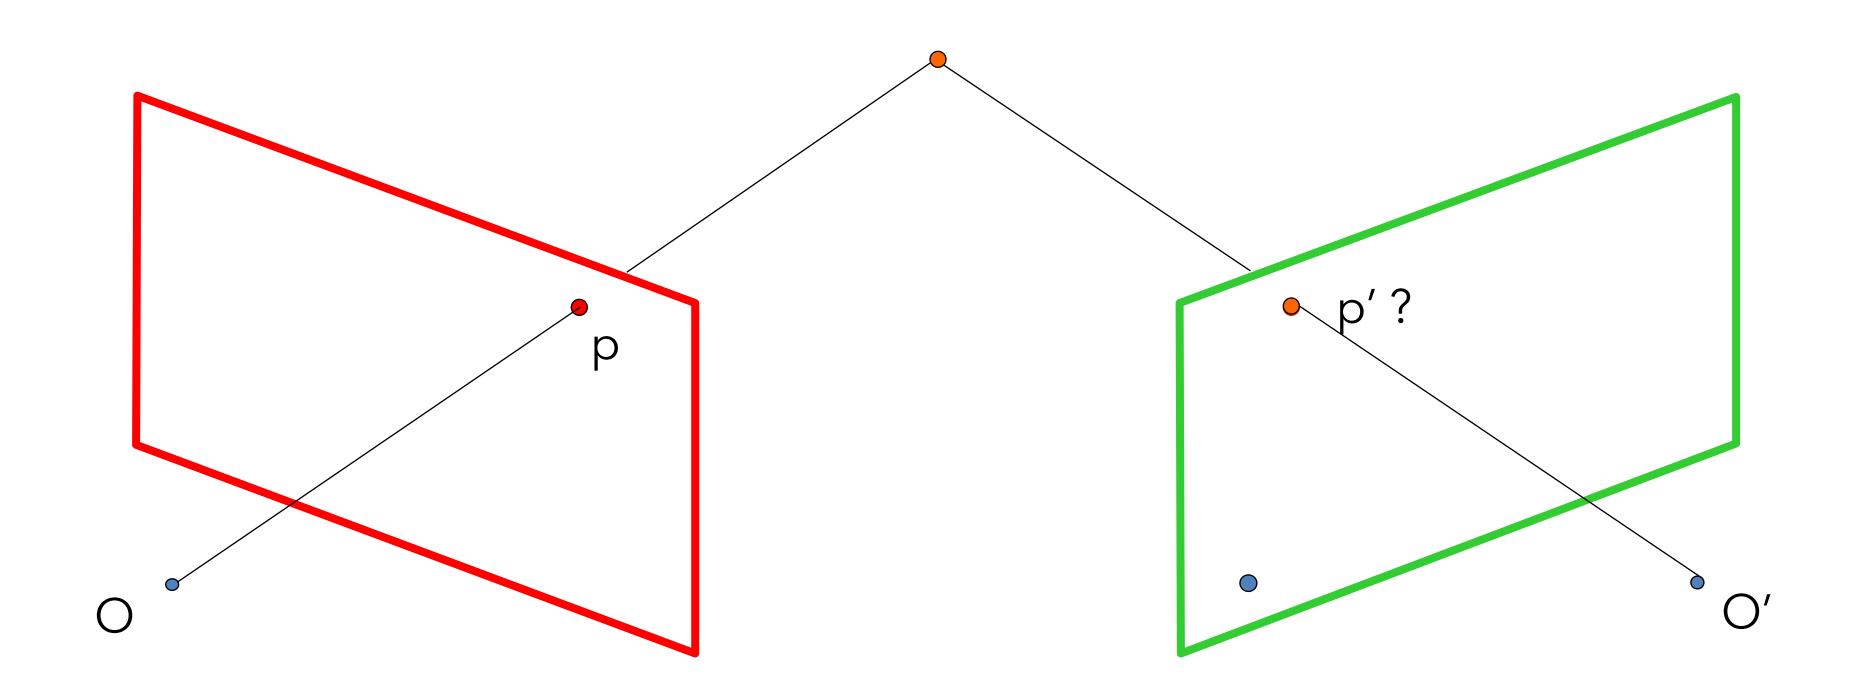
Stereo correspondence constraints

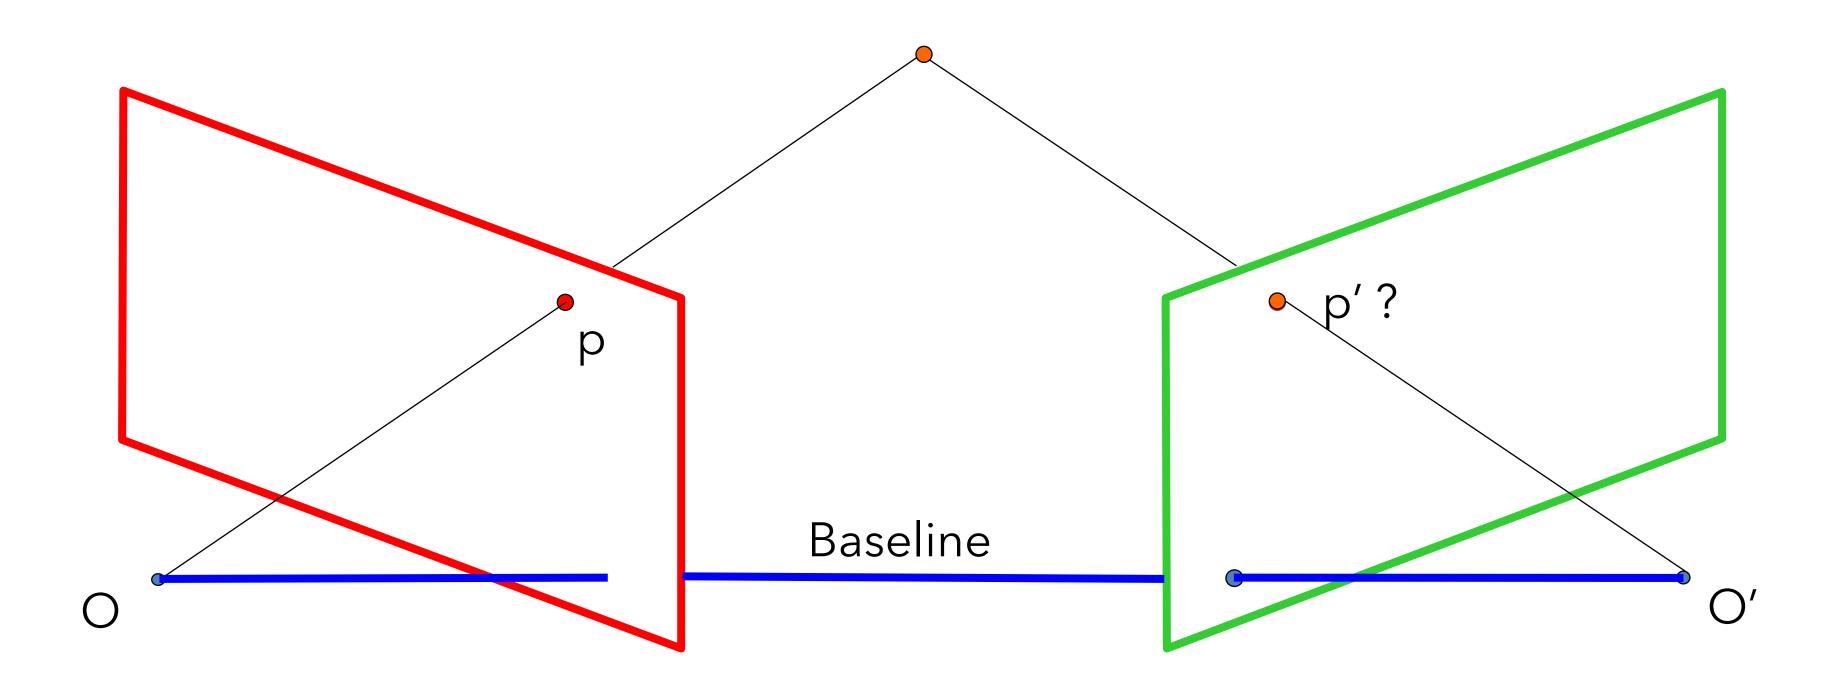


If we see a point in camera 1, are there any constraints on where we will find it on camera 2?

Stereo correspondence constraints

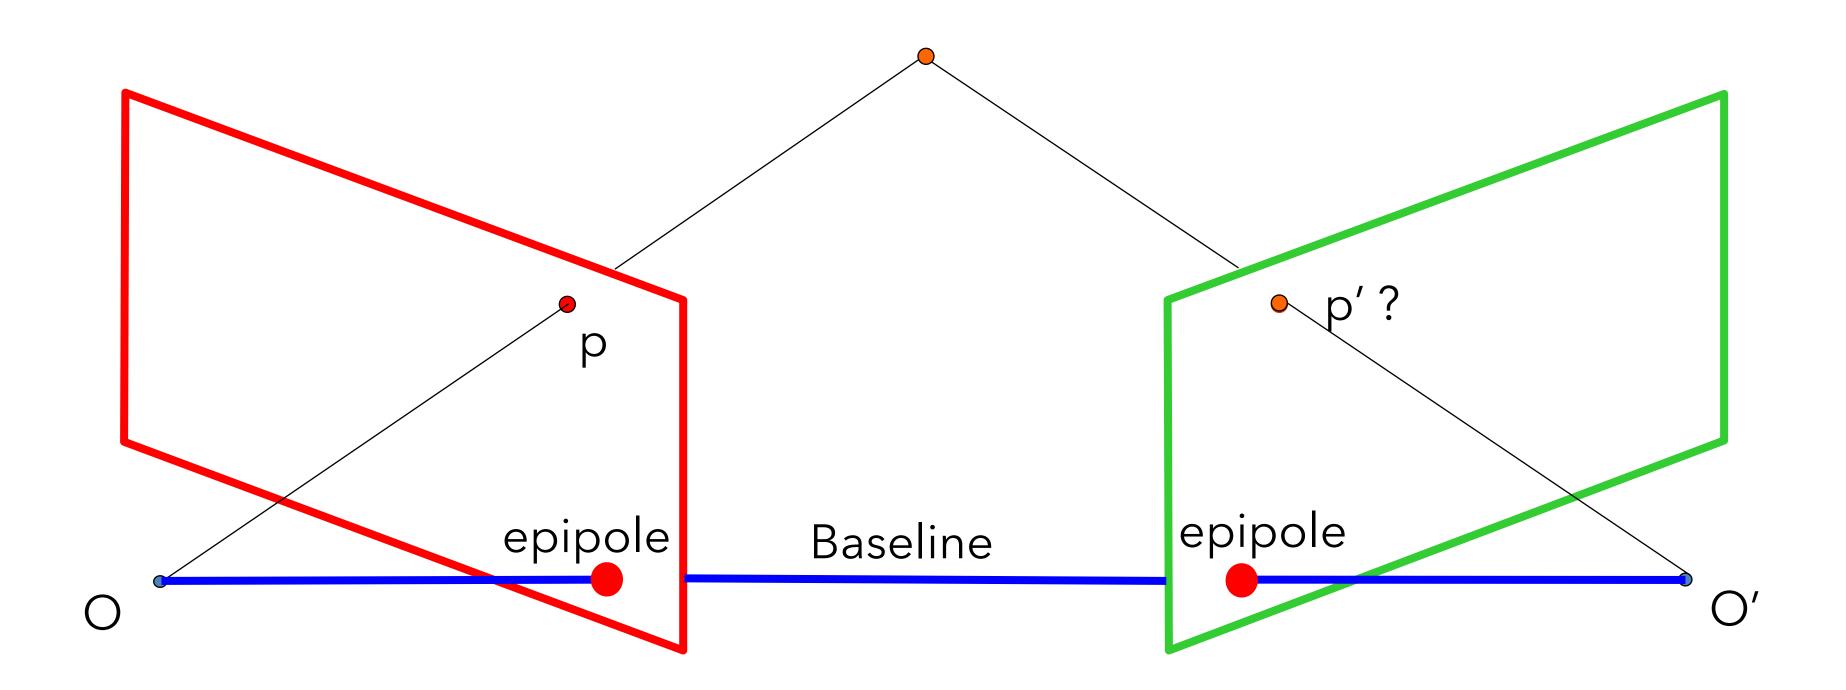






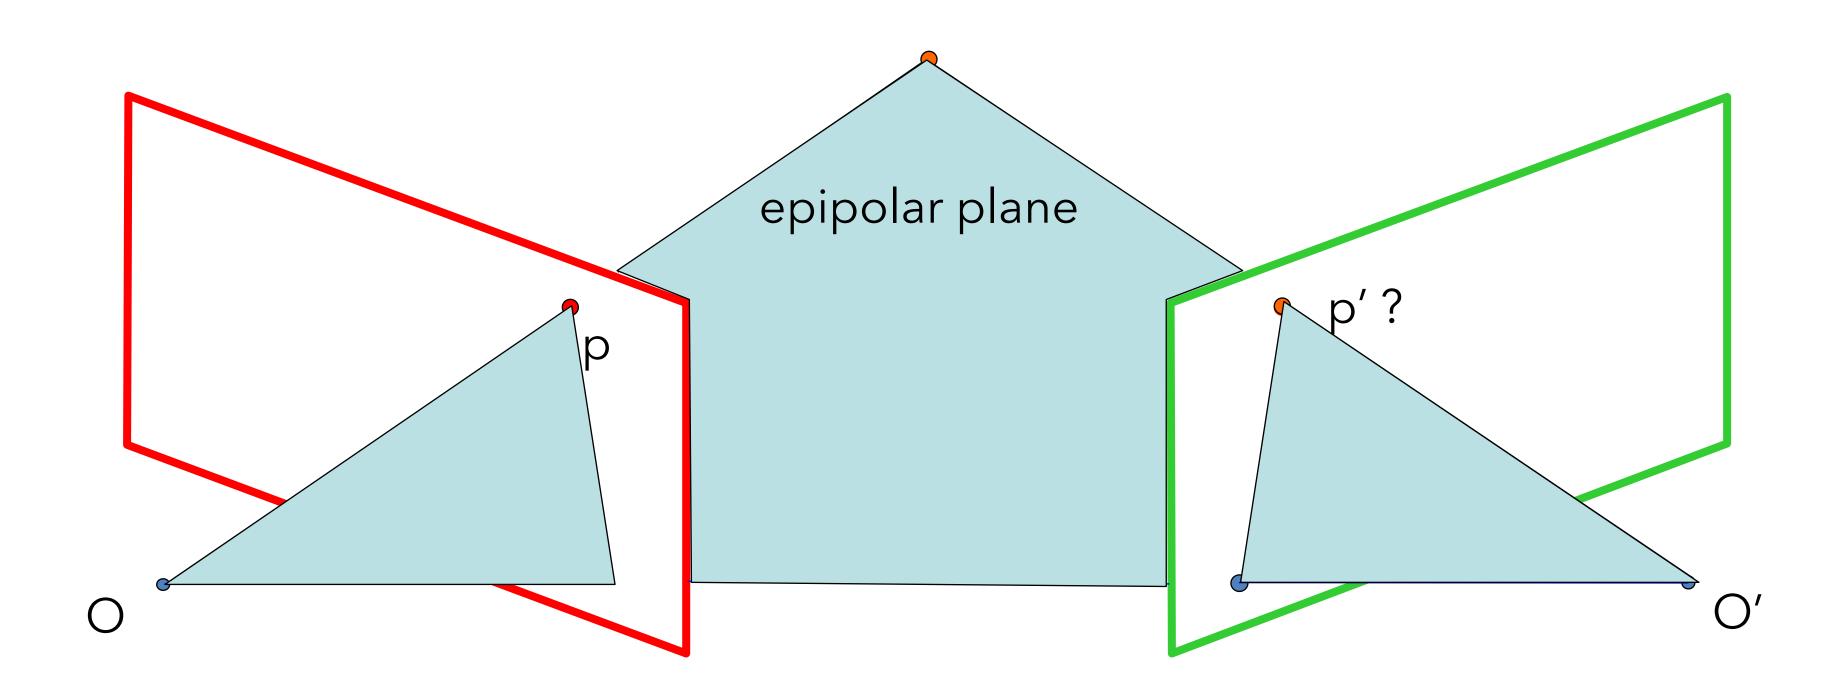
Baseline: the line connecting the two camera centers

Epipole: point of intersection of baseline with the image plane



Baseline: the line connecting the two camera centers

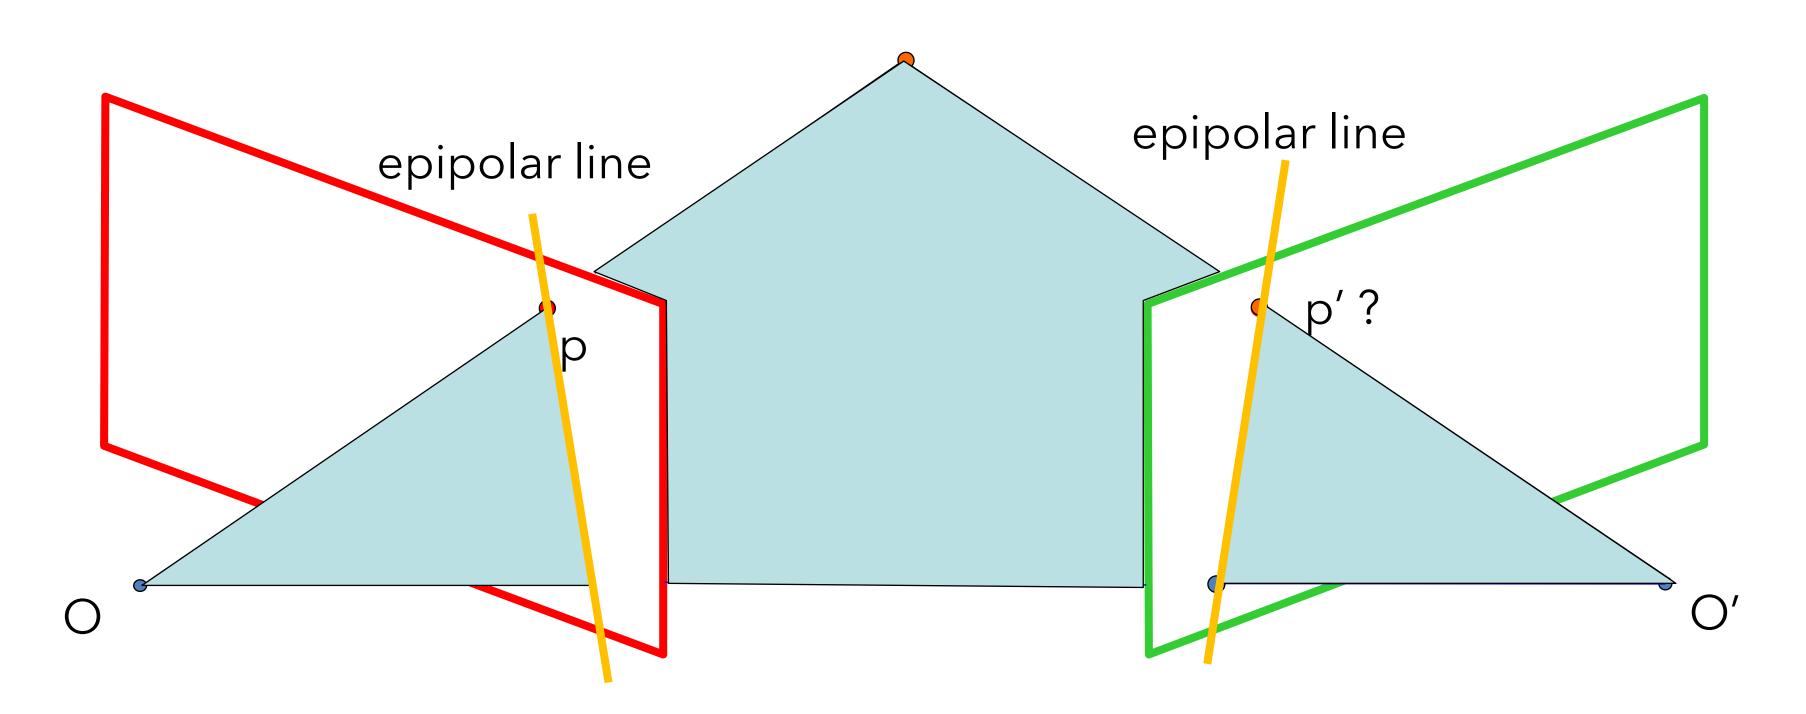
Epipole: point of intersection of baseline with the image plane



Baseline: the line connecting the two camera centers

Epipole: point of intersection of baseline with the image plane

Epipolar plane: the plane that contains the two camera centers and a 3D point in the world



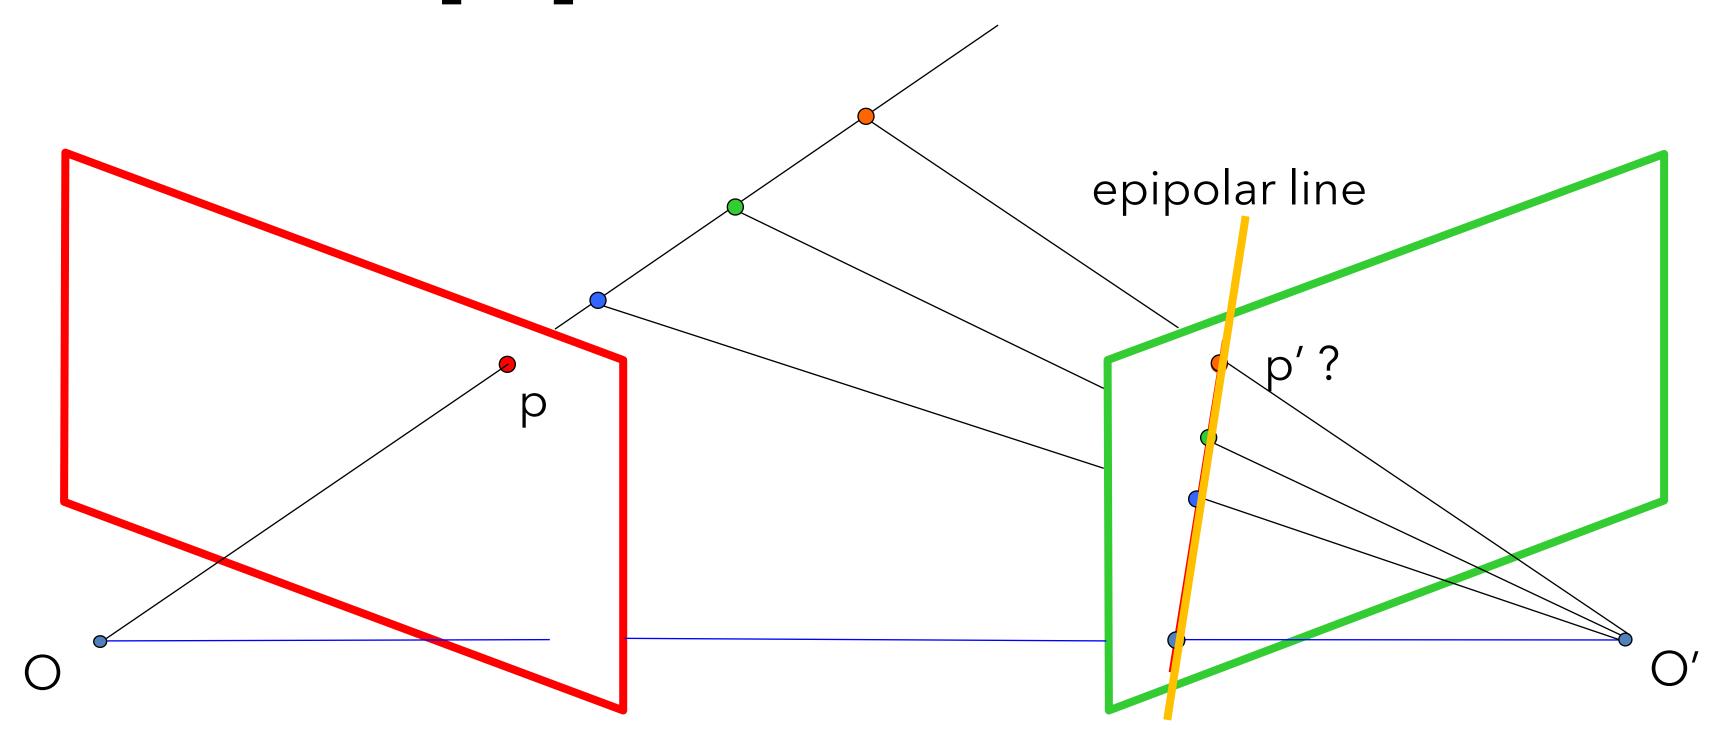
Baseline: the line connecting the two camera centers

Epipole: point of intersection of baseline with the image plane

Epipolar plane: the plane that contains the two camera centers and a 3D point in the world

Epipolar line: intersection of the epipolar plane with each image plane

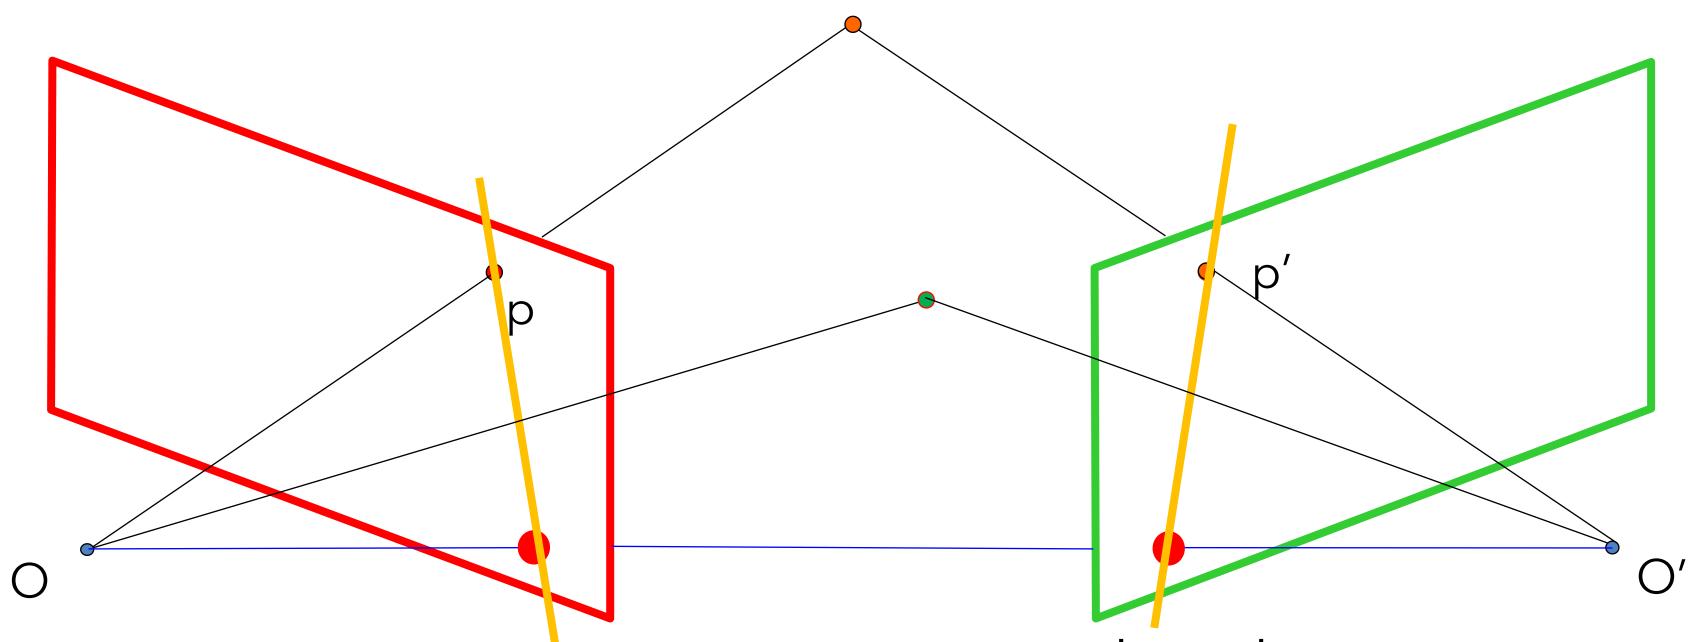
Epipolar constraint



We can search for matches across epipolar lines

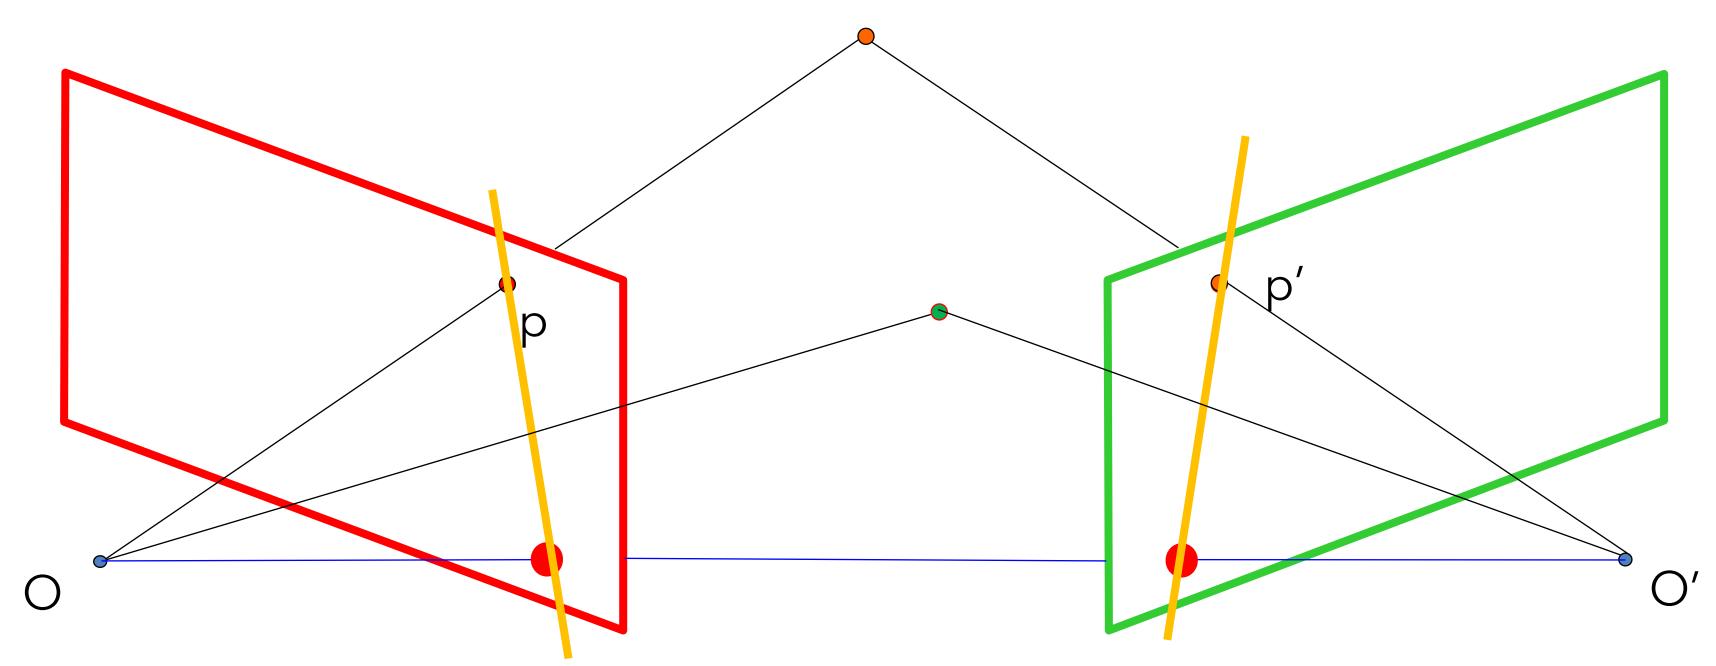
All epipolar lines intersect at the epipoles

Epipolar constraint



- If we observe a point in one image, its position in the other image is constrained to lie on the epipolar line.
- ullet How do we get this line? We want a function that, given a point ${f p}$ tells us the line:

$$f(\mathbf{p}) = [a, b, c]$$
 such that $ax' + by' + c = 0$
where $\mathbf{p}' = [x', y']$. In other words: $f(\mathbf{p})^{\mathsf{T}}\mathbf{p}' = 0$

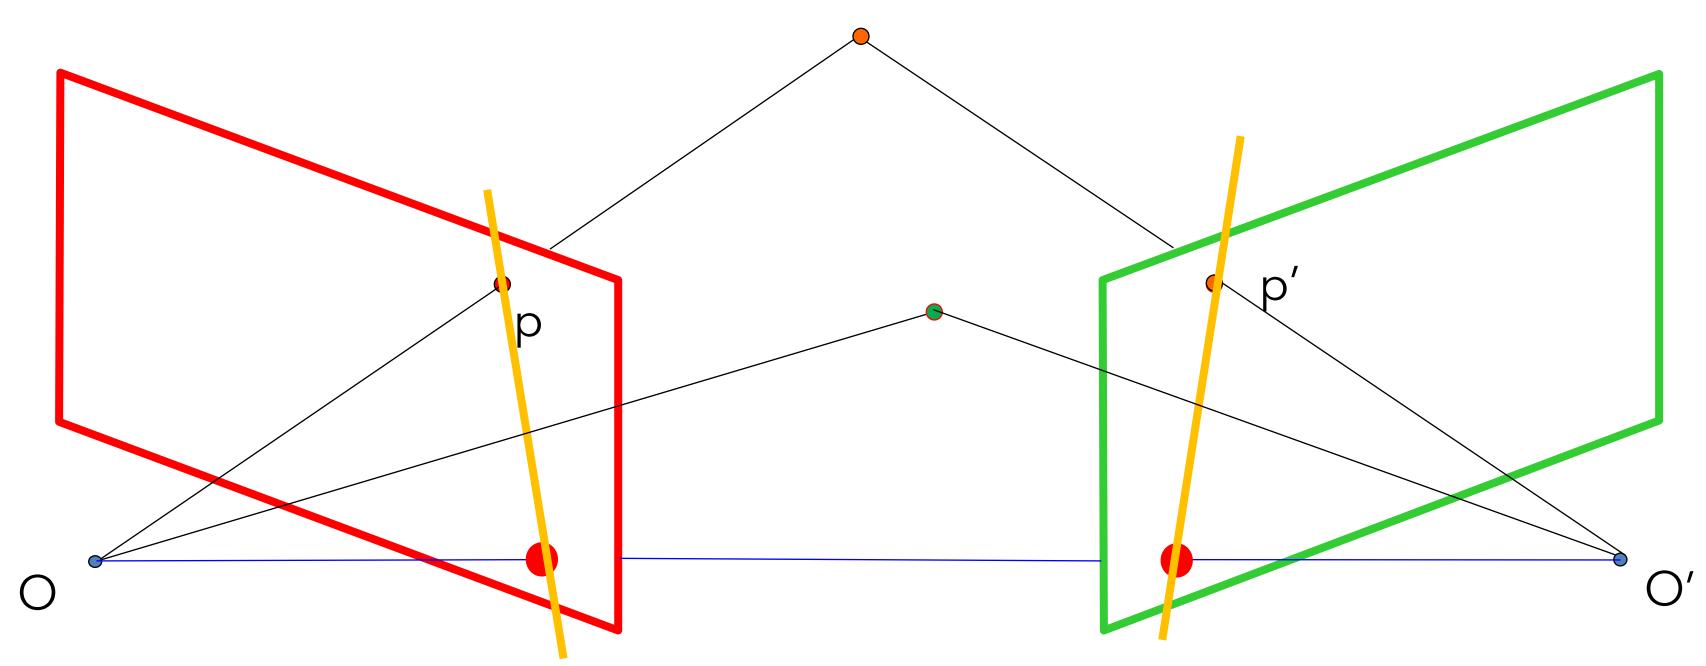


It can be shown that our function, f, can be written as a matrix multiplication in homogeneous coordinates!

$$f(\mathbf{p}) = \mathbf{p}^{\mathsf{T}} F = [a, b, c]$$

F: the fundamental matrix

p: image point in homogeneous coordinates



It can be shown that our function, f, can be written as a matrix multiplication in homogeneous coordinates!

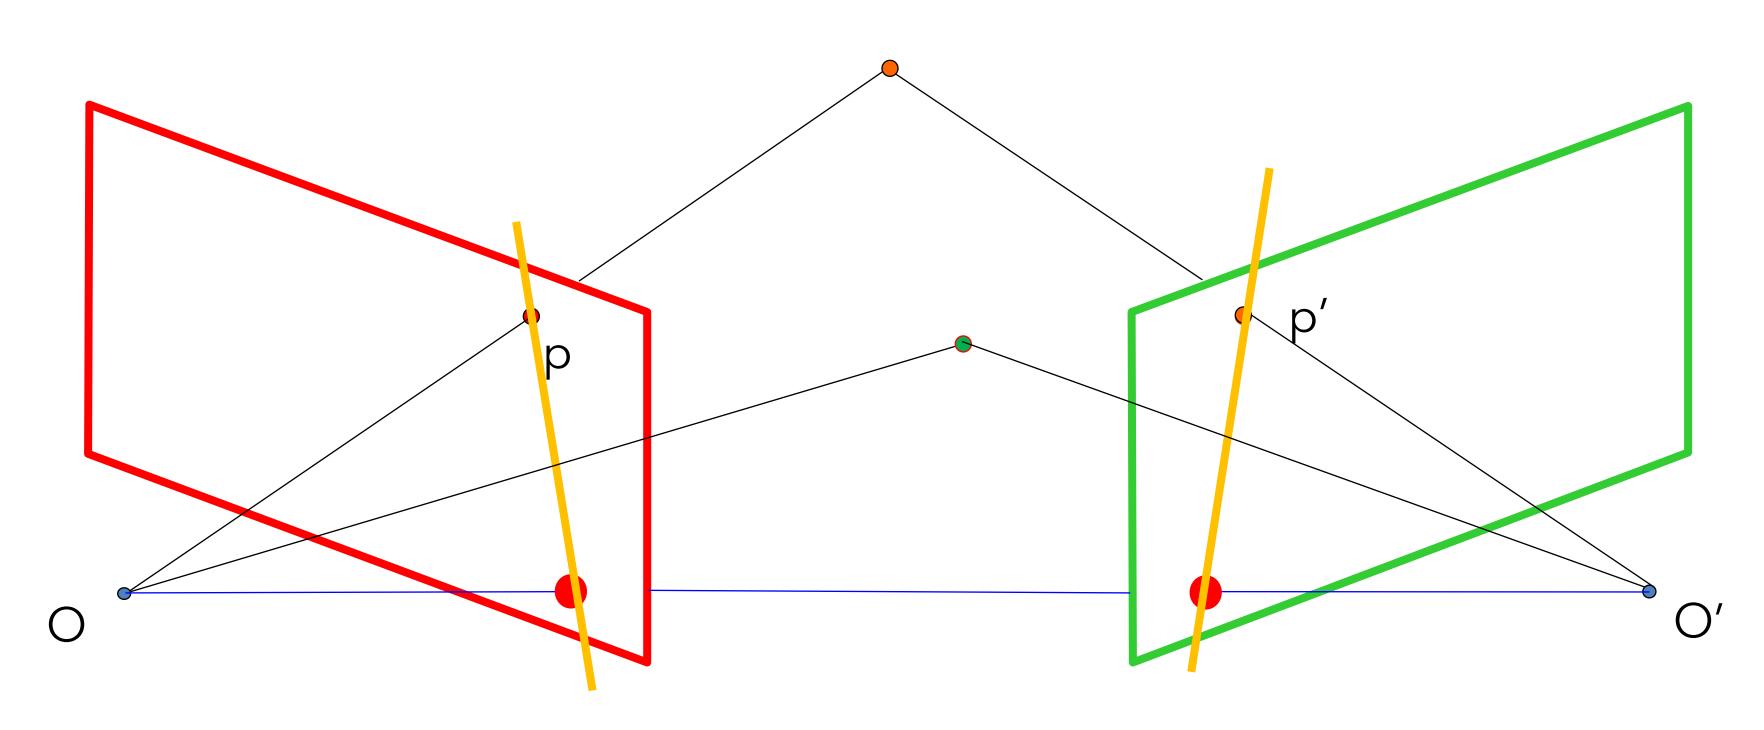
More concisely:

$$\mathbf{p}^{\mathsf{T}}F \mathbf{p}' = 0$$

F: the fundamental matrix

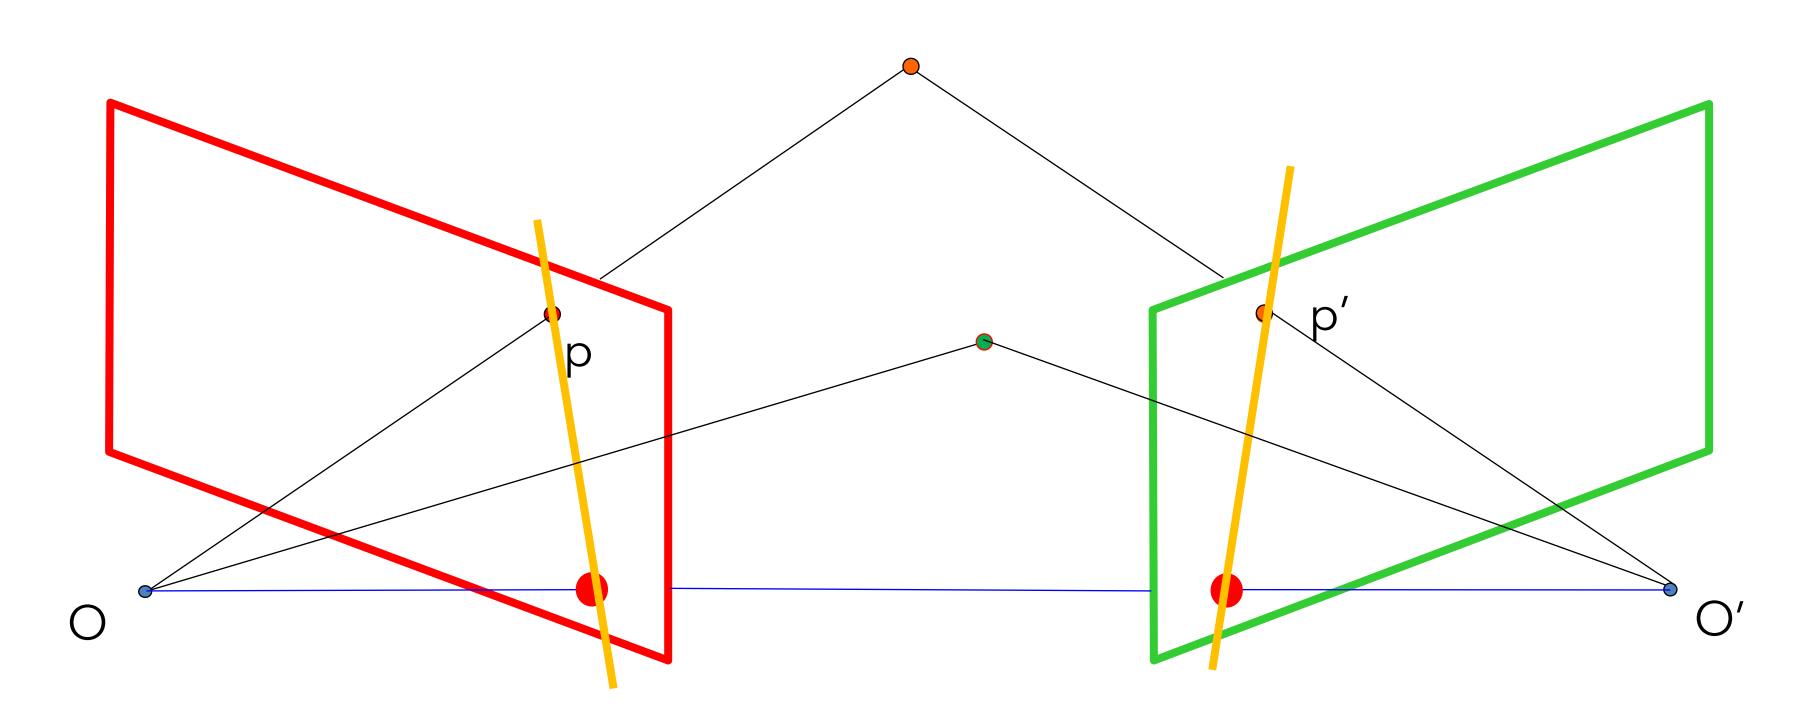
 $\mathbf{p}^{\mathsf{T}} F \mathbf{p}' = 0$ p, p': corresponding image points

Source: Torralba, Isola, Freeman



How this works:

$$\mathbf{p}^{\mathsf{T}}F \mathbf{p}' = 0$$



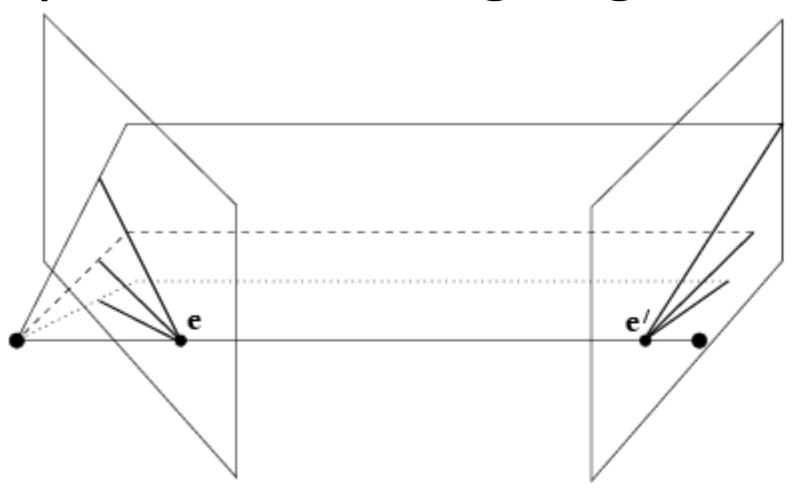
How this works:

$$(\mathbf{p}^{\mathsf{T}}F)\ \mathbf{p}' = 0$$
$$\mathbf{u}^{\mathsf{T}}\mathbf{p}' = 0$$

u: a line induced by p

p, p': image points in homogeneous coordinates

Example: converging cameras





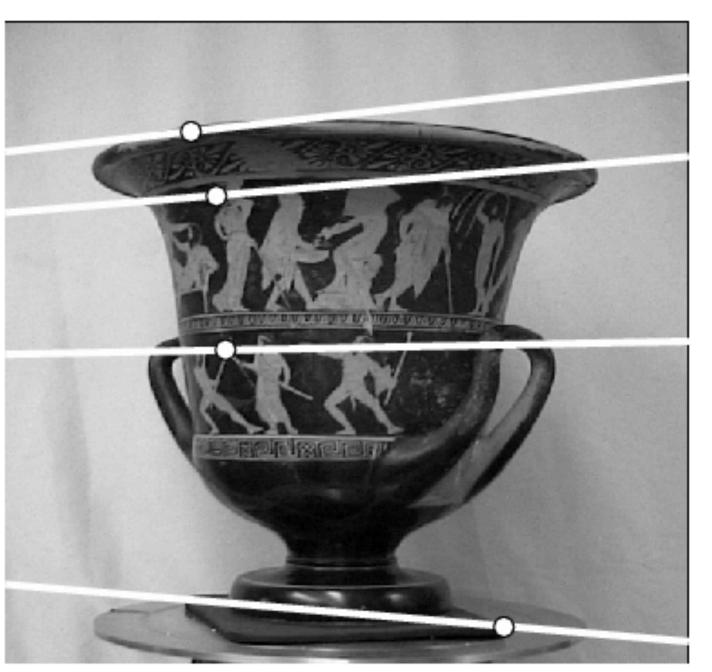
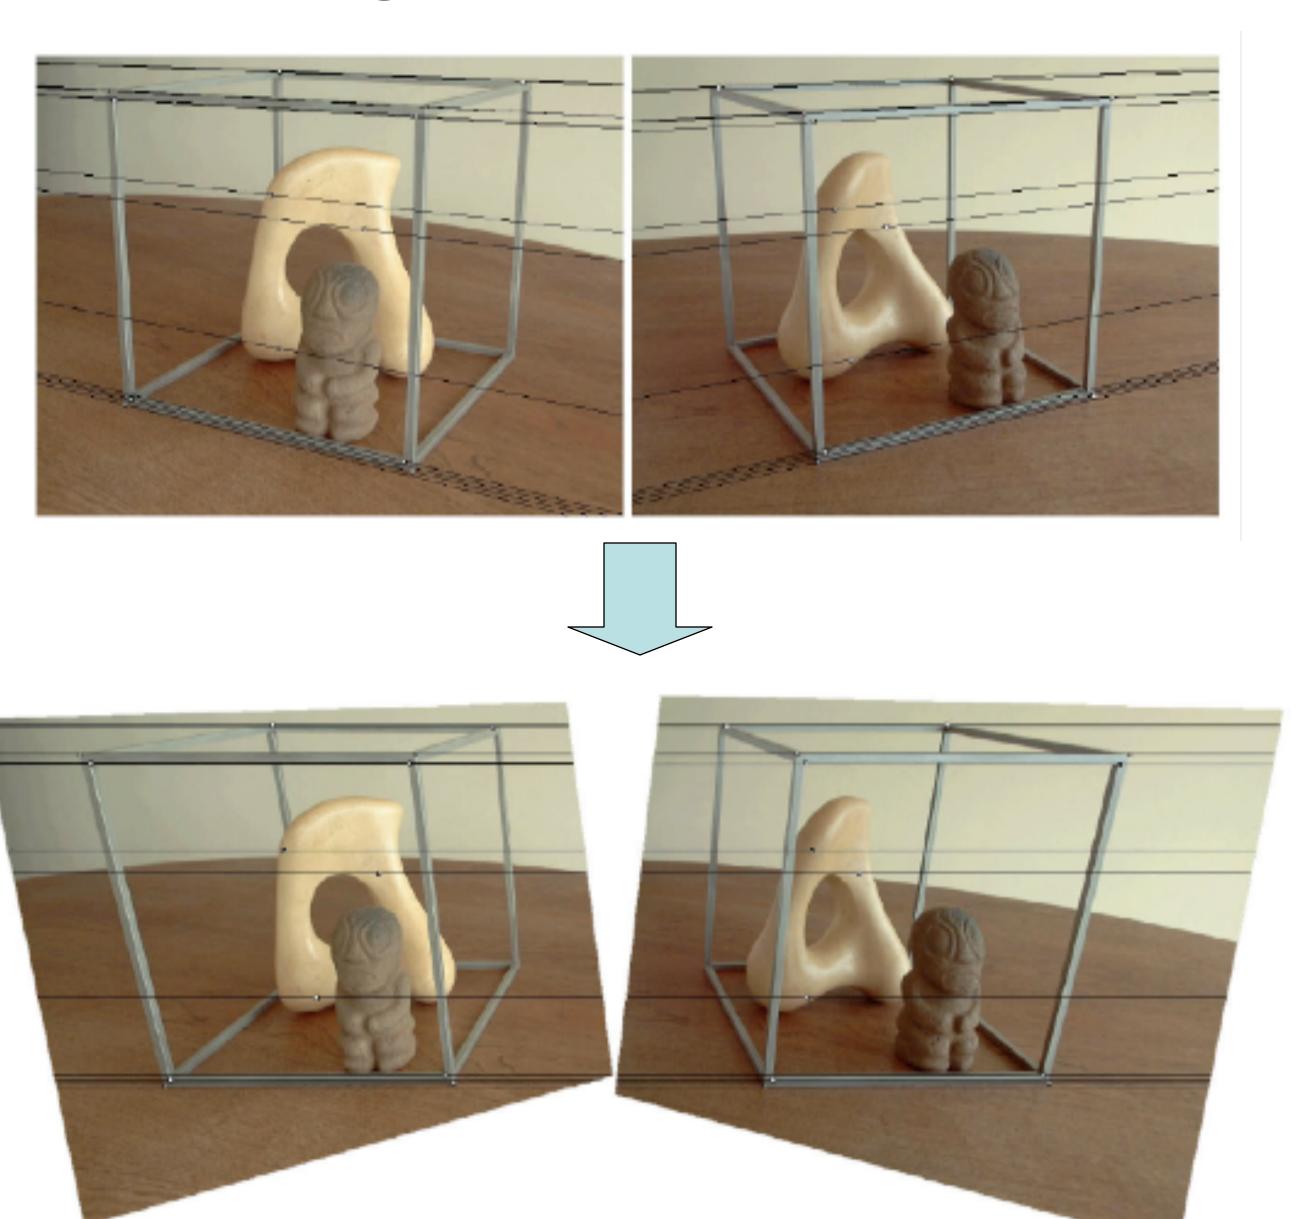


Image rectification



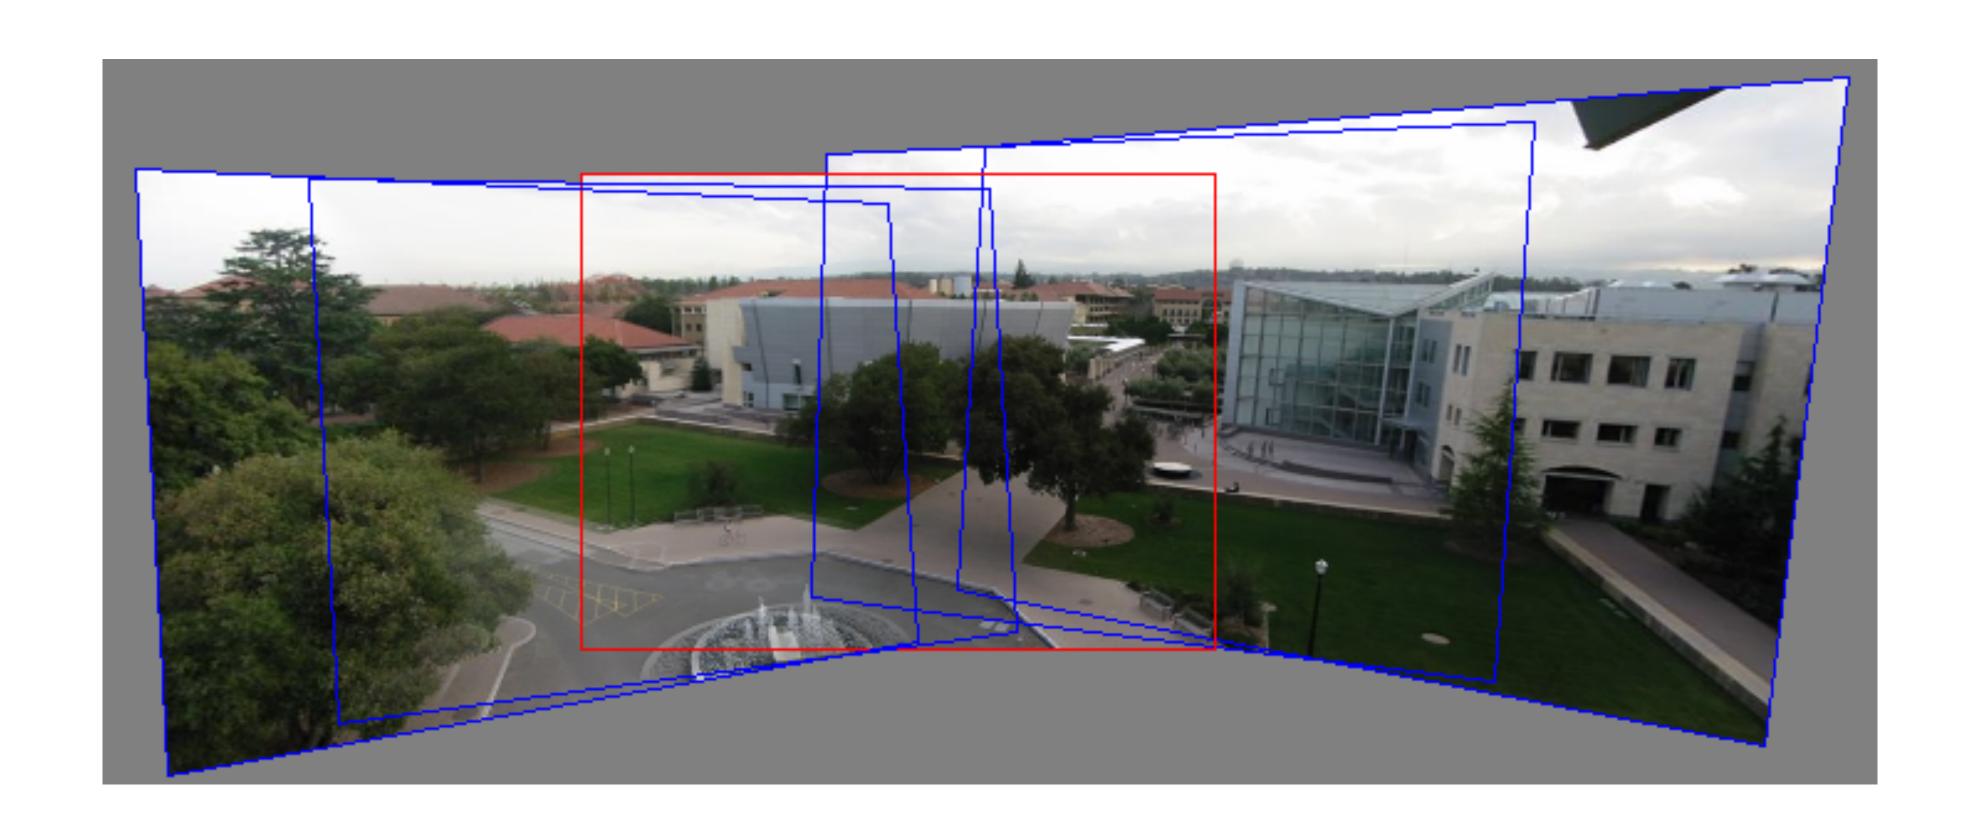
70

Example





Making panoramas



Making panoramas

What is the geometric relationship between these two images?

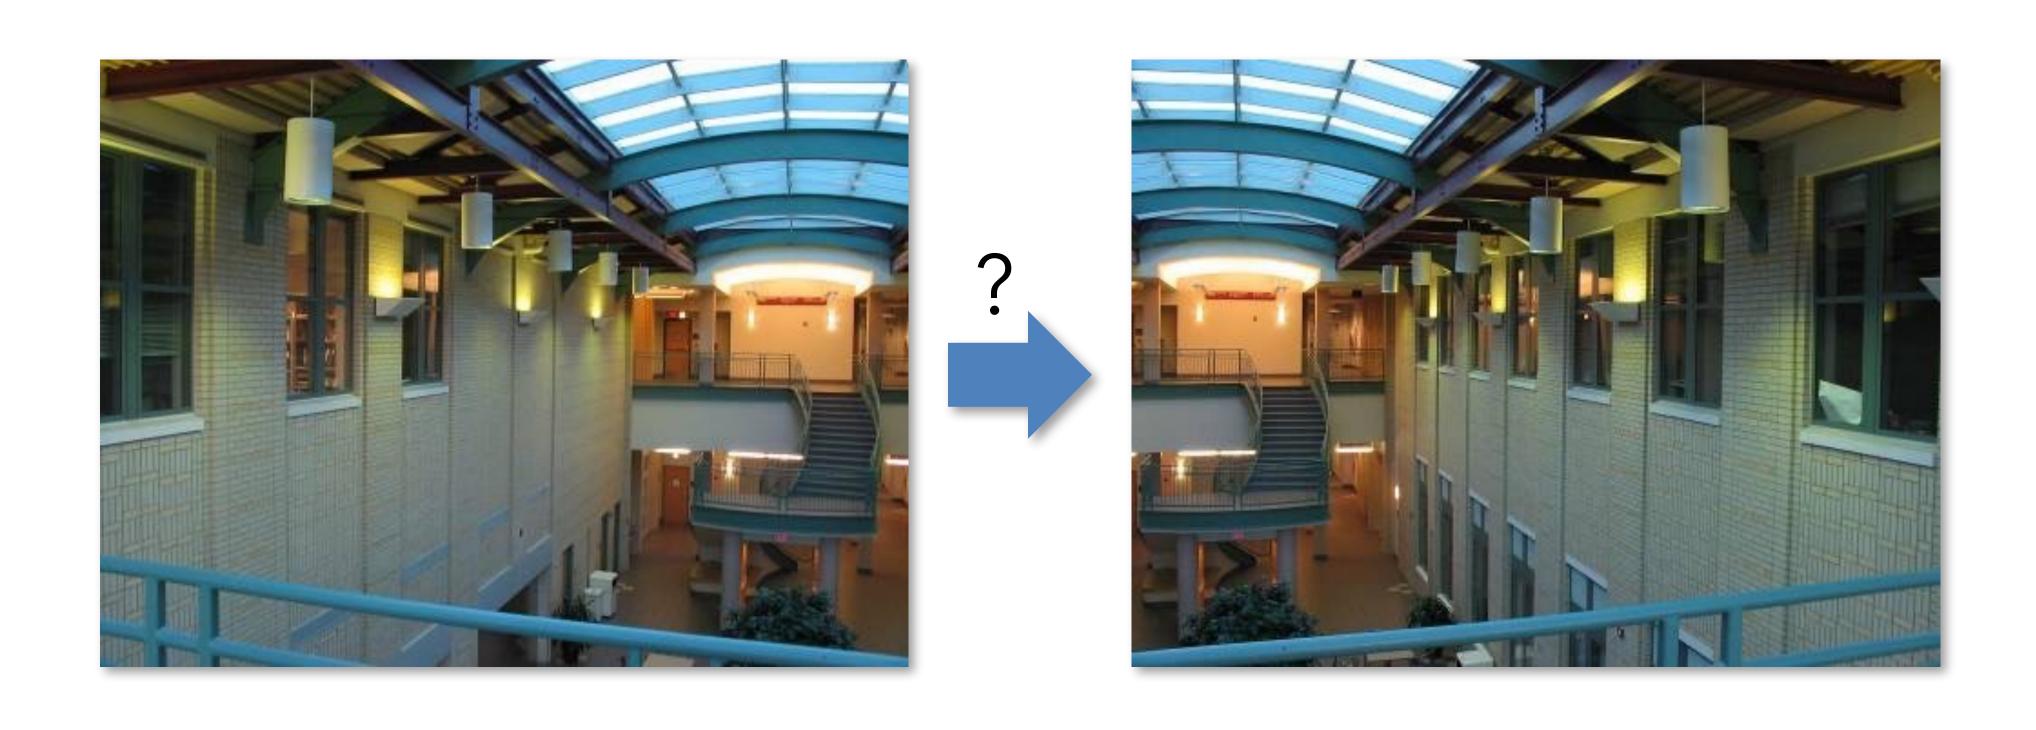


Image alignment



Why don't these image line up exactly?

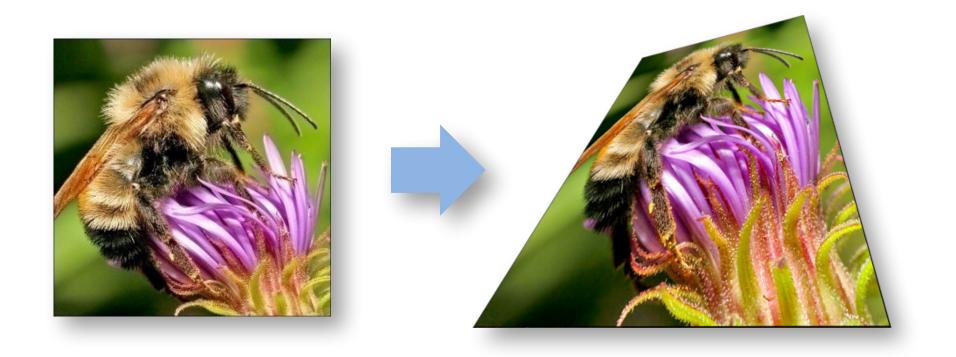
Source: N. Snavely

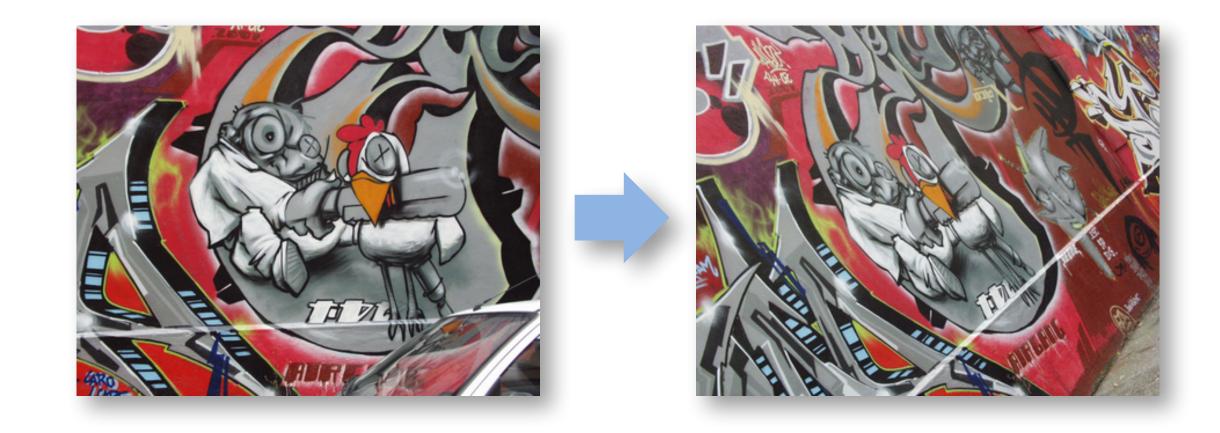
Recall: affine transformations

Projective Transformations aka Homographies aka Planar Perspective Maps

$$\mathbf{H} = egin{bmatrix} a & b & c \ d & e & f \ g & h & 1 \end{bmatrix}$$

Called a **homography**(or planar perspective map)





Homographies

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$
Note that this can be
0!

A "point at infinity"

$$\sim \begin{bmatrix} \frac{ax+by+c}{gx+hy+1} \\ \frac{dx+ey+f}{gx+hy+1} \\ 1 \end{bmatrix}$$

Homography

Example: two pictures taken by rotating the camera:







Homography

Example: two pictures taken by rotating the camera:



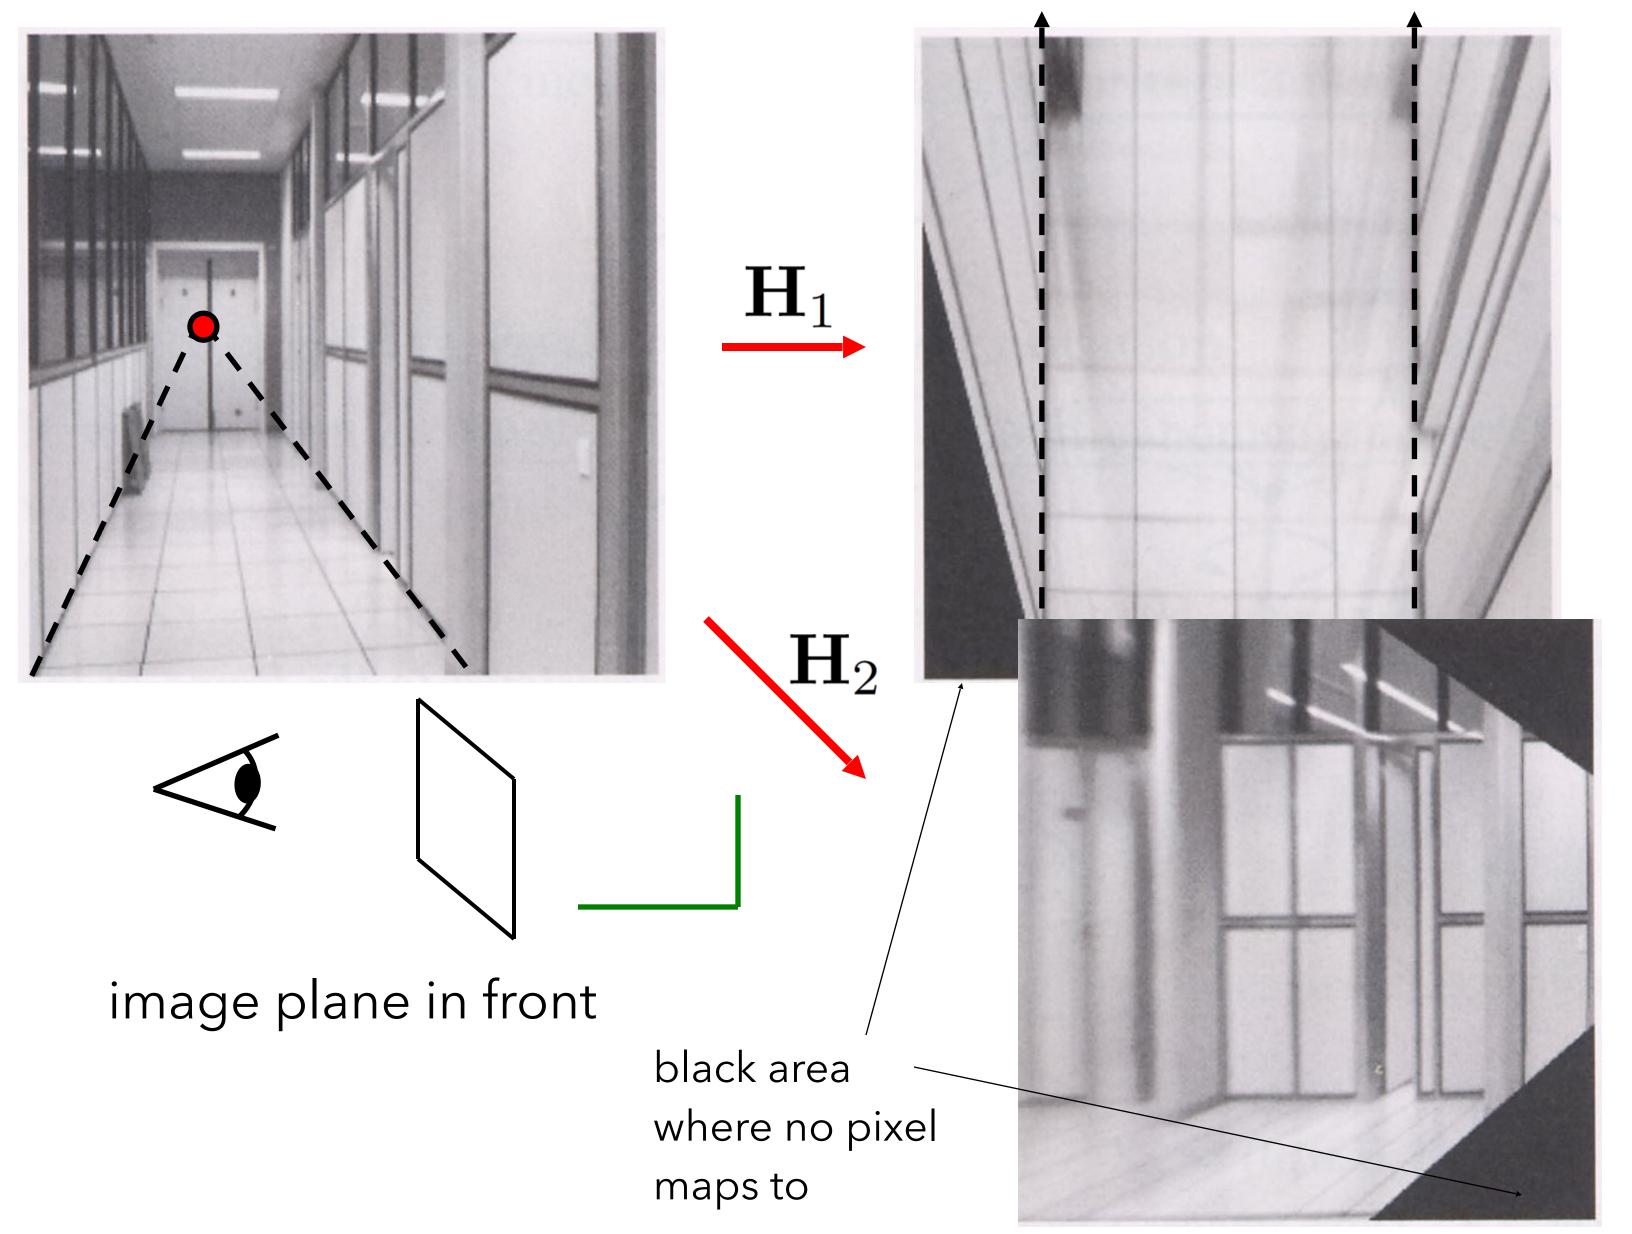


With a homography you can map both images into a single camera:



We'll see why in PS5!

Plane-to-plane homography



Source: N. Snavely

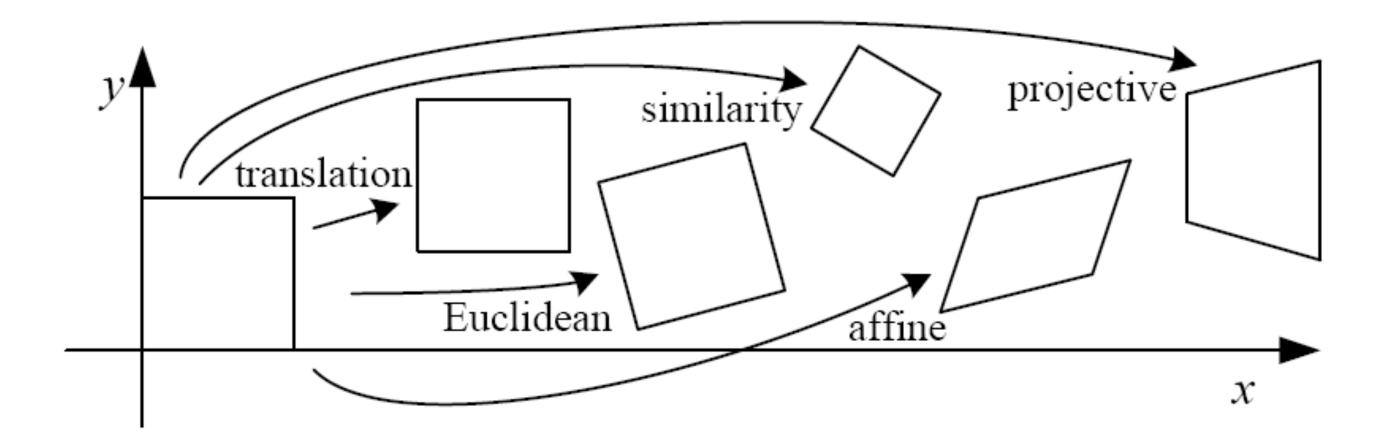
Homographies

- - Projective warps

• Homographies ...
$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

- Properties of projective transformations:
 - Origin does not necessarily map to origin
 - Lines map to lines
 - Parallel lines do not necessarily remain parallel
 - Closed under composition

2D image transformations



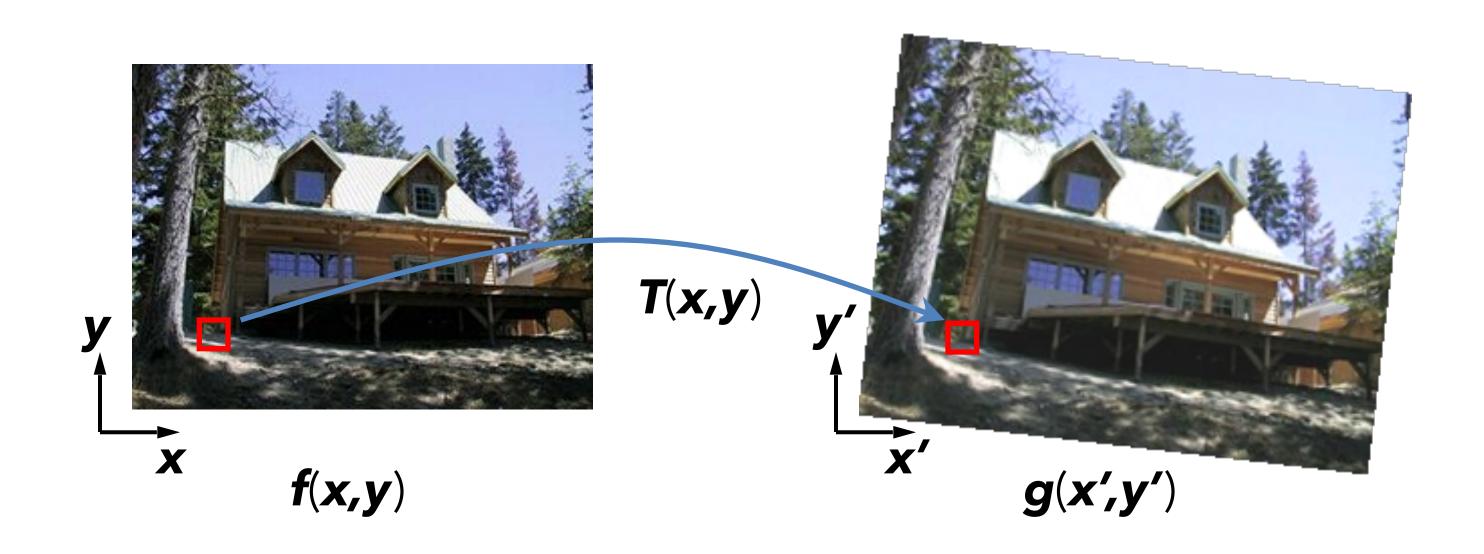
Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\left[egin{array}{c c} oldsymbol{I} & oldsymbol{t} \end{array} ight]_{2 imes 3}$	2	orientation + · · ·	
rigid (Euclidean)	$\left[egin{array}{c c} oldsymbol{R} & oldsymbol{t} \end{array} ight]_{2 imes 3}$	3	lengths +···	
similarity	$\left[\begin{array}{c c} sR & t\end{array}\right]_{2\times 3}$	4	angles + · · ·	
affine	$\left[egin{array}{c} oldsymbol{A} \end{array} ight]_{2 imes 3}$	6	parallelism + · · ·	
projective	$\left[egin{array}{c} ilde{m{H}} \end{array} ight]_{3 imes 3}$	8	straight lines	

Source: N. Snavely

How do we perform this warp?

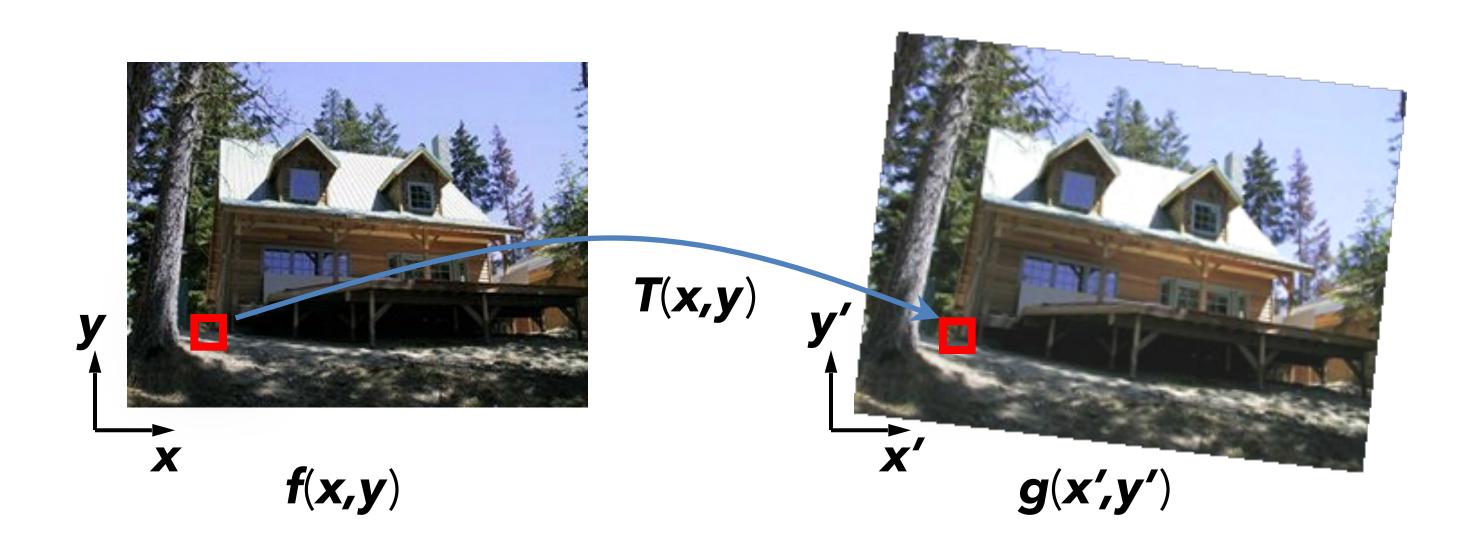
Image warping

Given a coordinate transformation (x',y') = T(x,y) and a source image f(x,y), how do we compute a transformed image g(x',y') = f(T(x,y))?



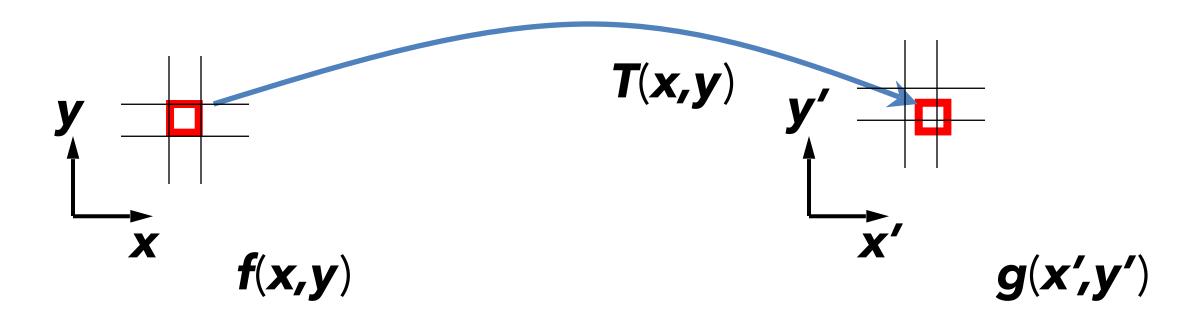
Forward warping

- Send each pixel f(x) to its corresponding location (x',y') = T(x,y) in g(x',y')
- What if a pixel lands "between" two pixels?



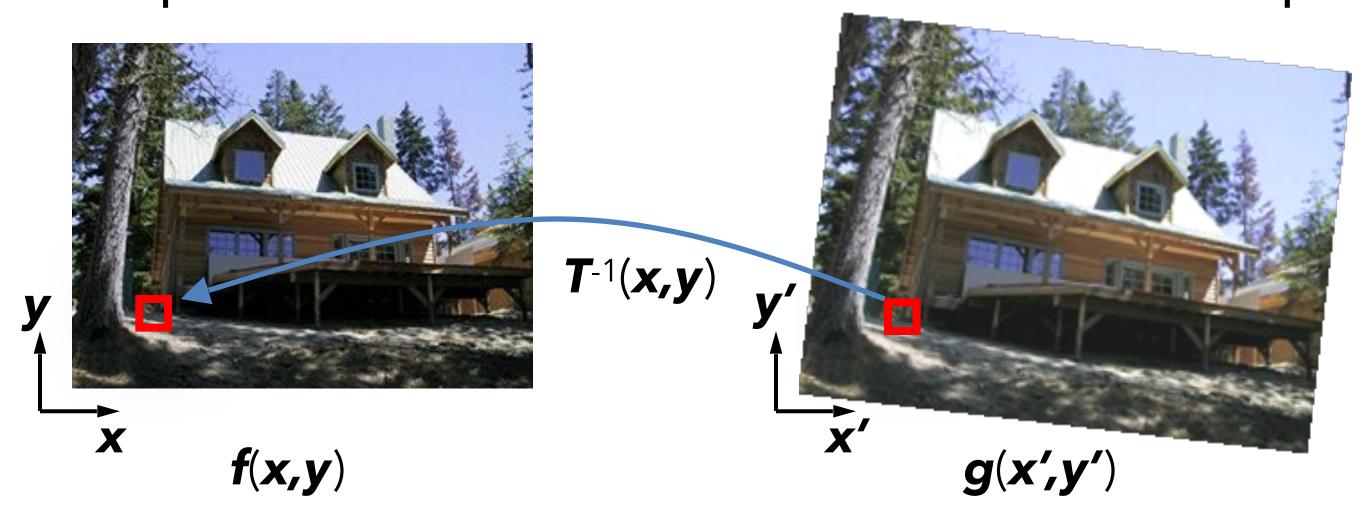
Forward warping

- Send each pixel f(x) to its corresponding location (x',y') = T(x,y) in g(x',y')
- What if a pixel lands "between" two pixels?
 - Answer: add "contribution" to several pixels, normalize later (splatting)
 - Can still result in holes



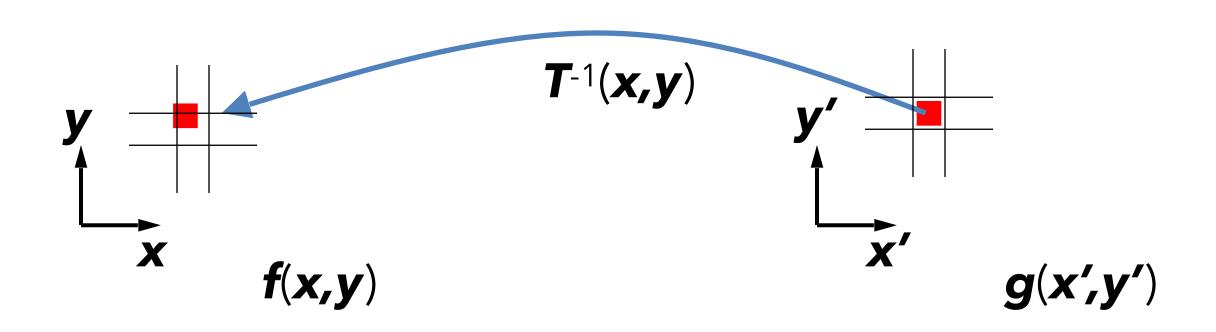
Backward (inverse) warping

- Get each pixel g(x',y') from its corresponding location $(x,y) = T^{-1}(x,y)$ in f(x,y)
 - Requires taking the inverse of the transform
 - What if pixel comes from "between" two pixels?



Backward (inverse) warping

- Get each pixel g(x') from its corresponding location x' = h(x) in f(x)
 - What if pixel comes from "between" two pixels?
 - Answer: resample color value from interpolated source image



Next class: estimating geometry from images