

# Lecture 17: Multi-view geometry

# Today

- Epipolar geometry
- Stereo matching
- Image alignment

# PS8: making panoramas



# Recall: homogeneous coordinates

Representing translations:

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

homogeneous image  
coordinates

$$\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} x + t_x \\ y + t_y \\ 1 \end{bmatrix}$$

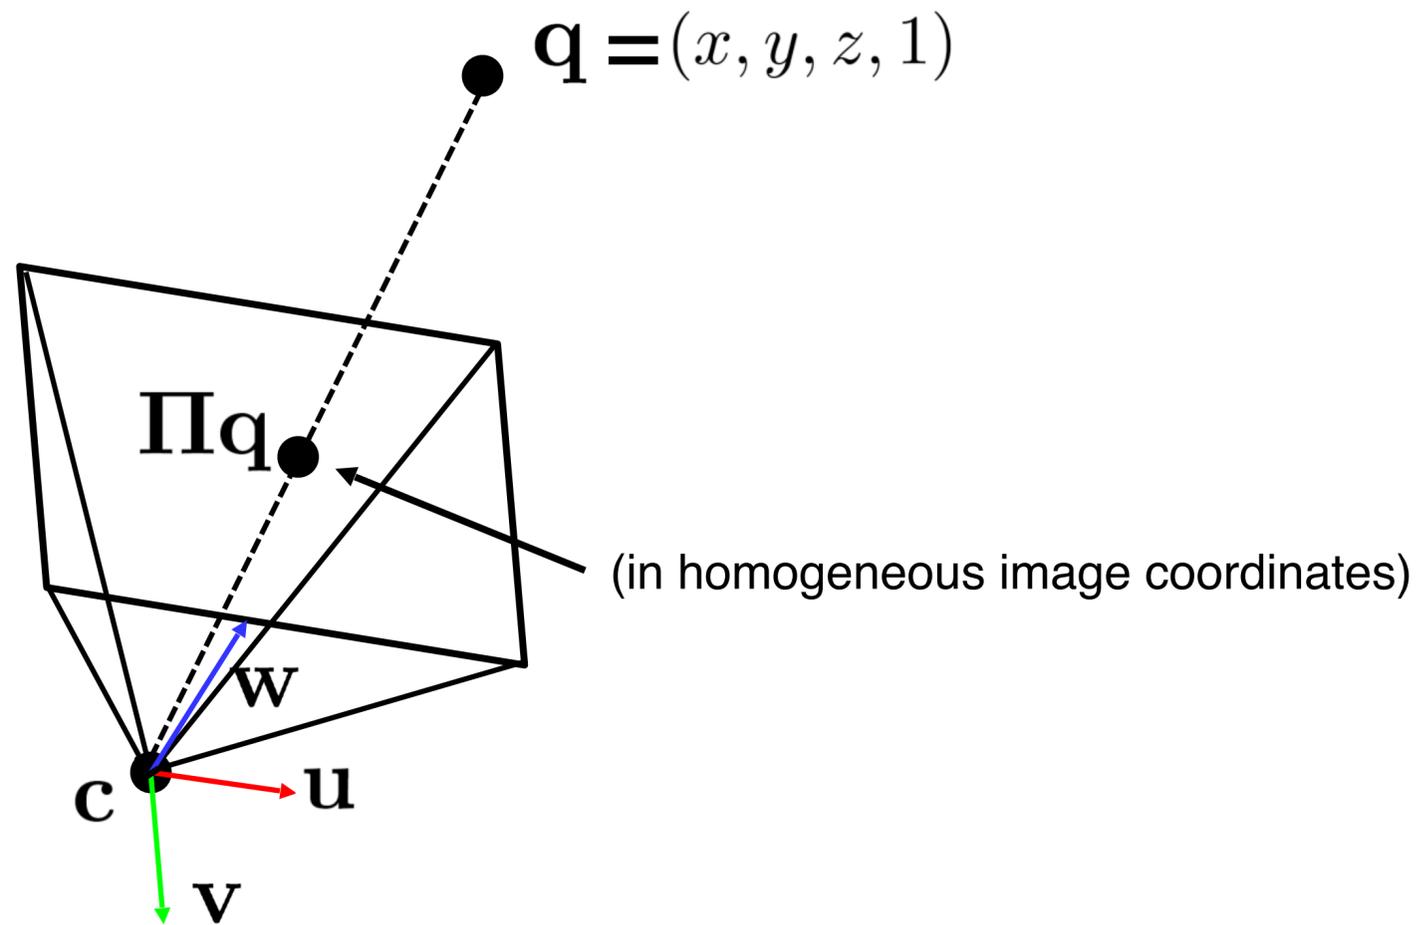
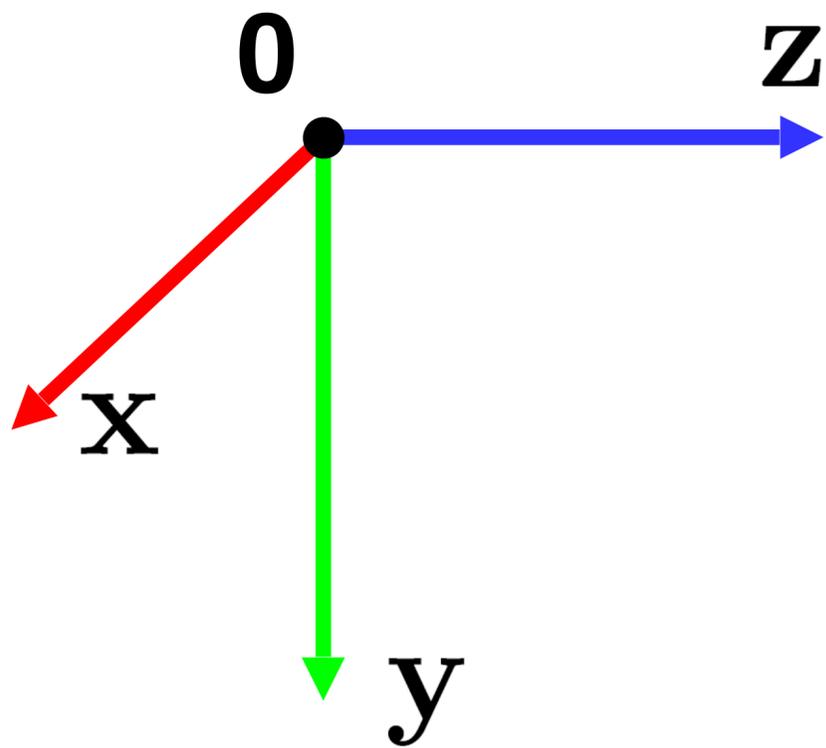
Converting back to image coordinates

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$

# Recall: camera parameters

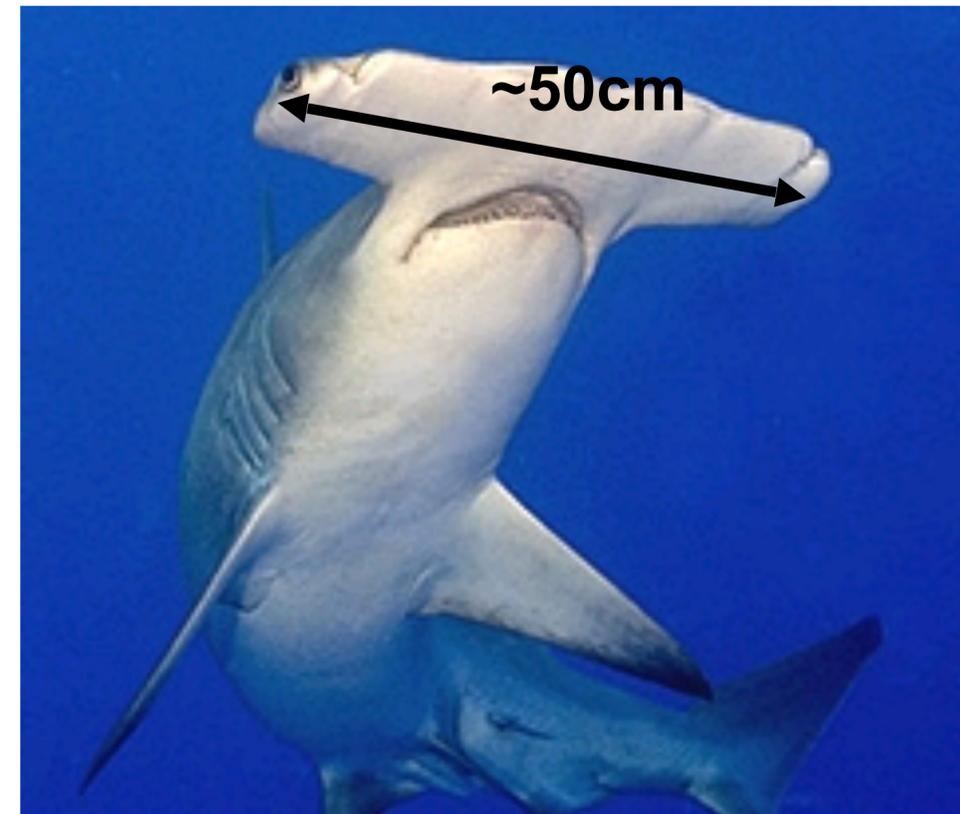
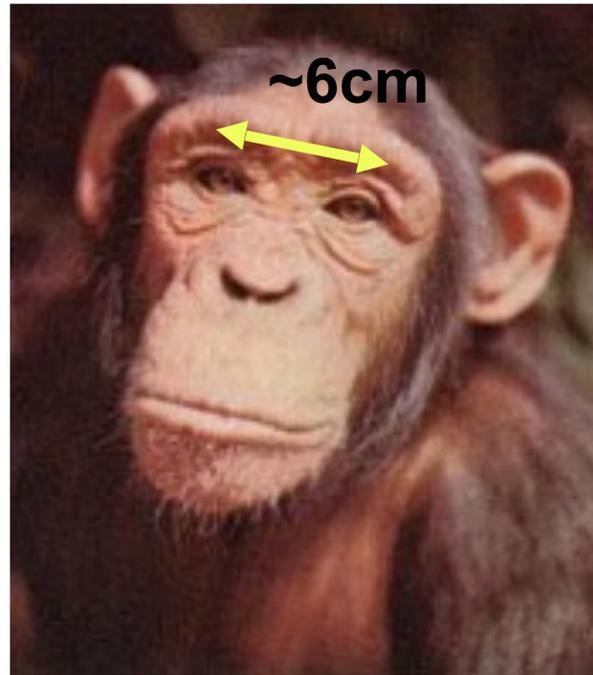
$$\mathbf{\Pi} = \underbrace{\begin{bmatrix} f & s & c_x \\ 0 & \alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}}_{\text{intrinsic}} \underbrace{\begin{bmatrix} \mathbf{R}_{3 \times 3} & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}}_{\text{rotation}} \underbrace{\begin{bmatrix} \mathbf{I}_{3 \times 3} & \mathbf{T}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}}_{\text{translation}}$$

# Projection matrix



# Estimating depth from multiple views

# Stereo vision

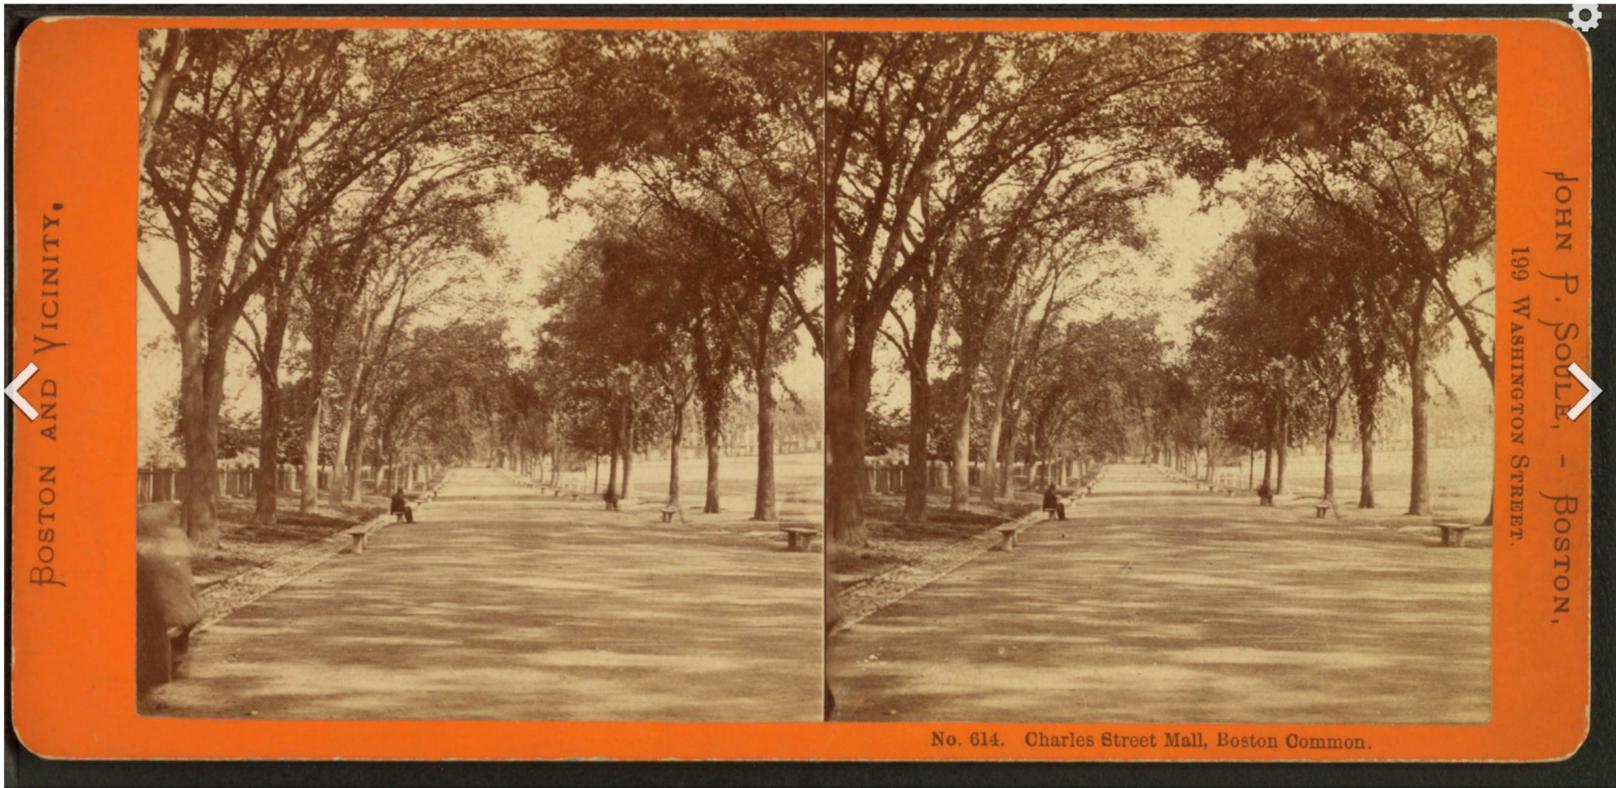


# 1 vs. 2 eyes



# 1 vs. 2 eyes





Stereoscopic card



Brewster stereoscope

# Depth without objects

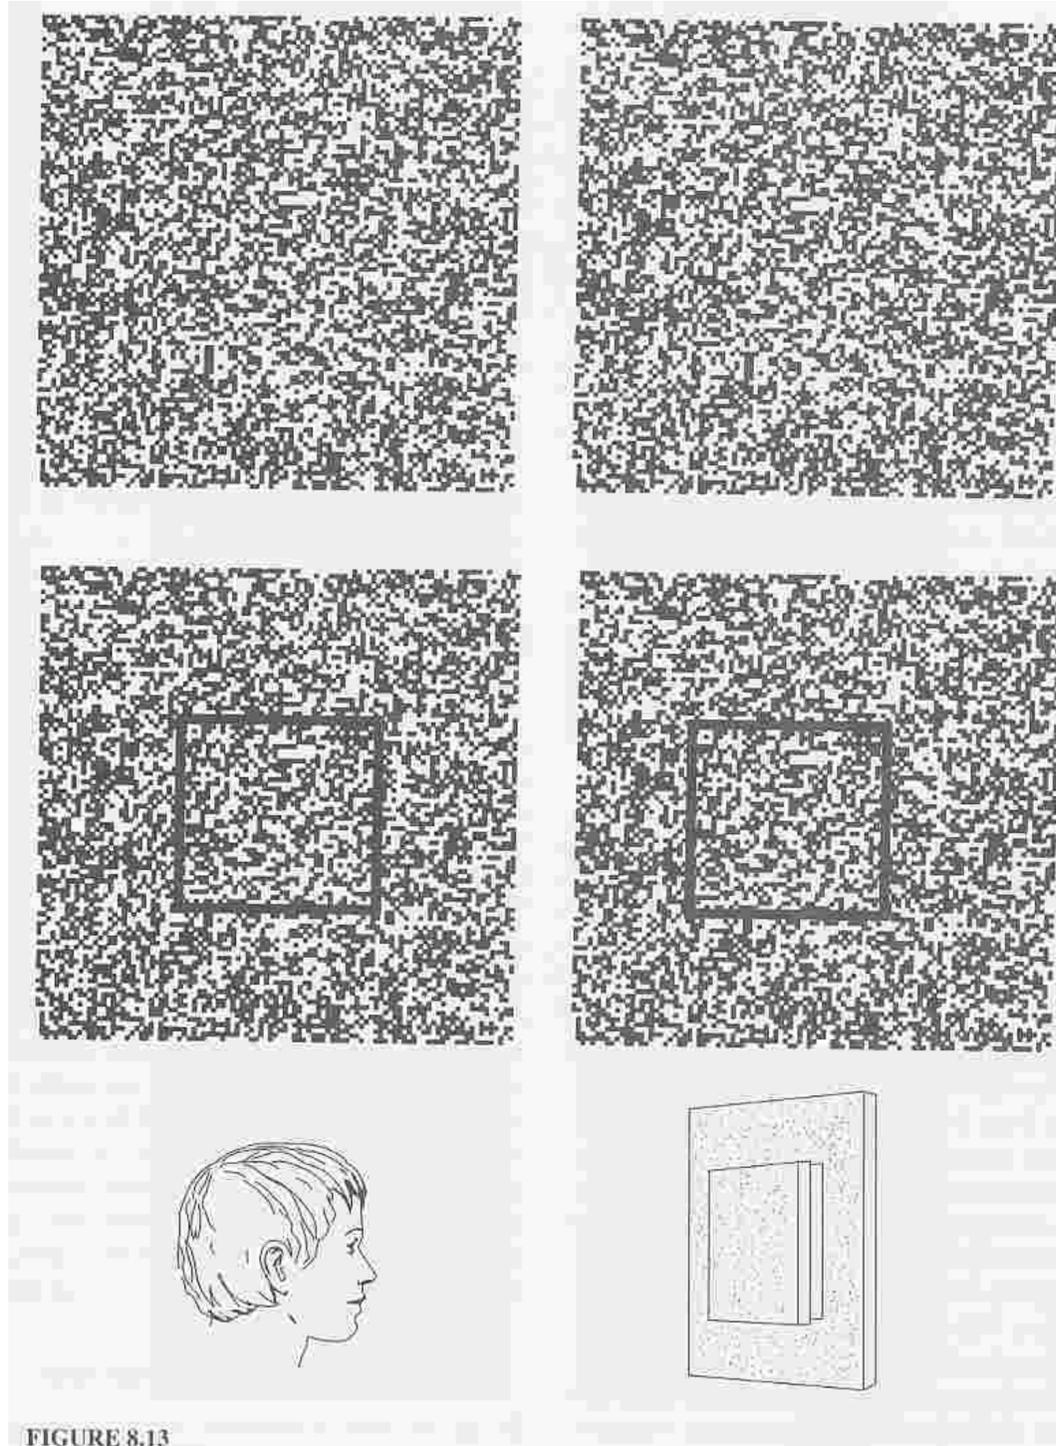
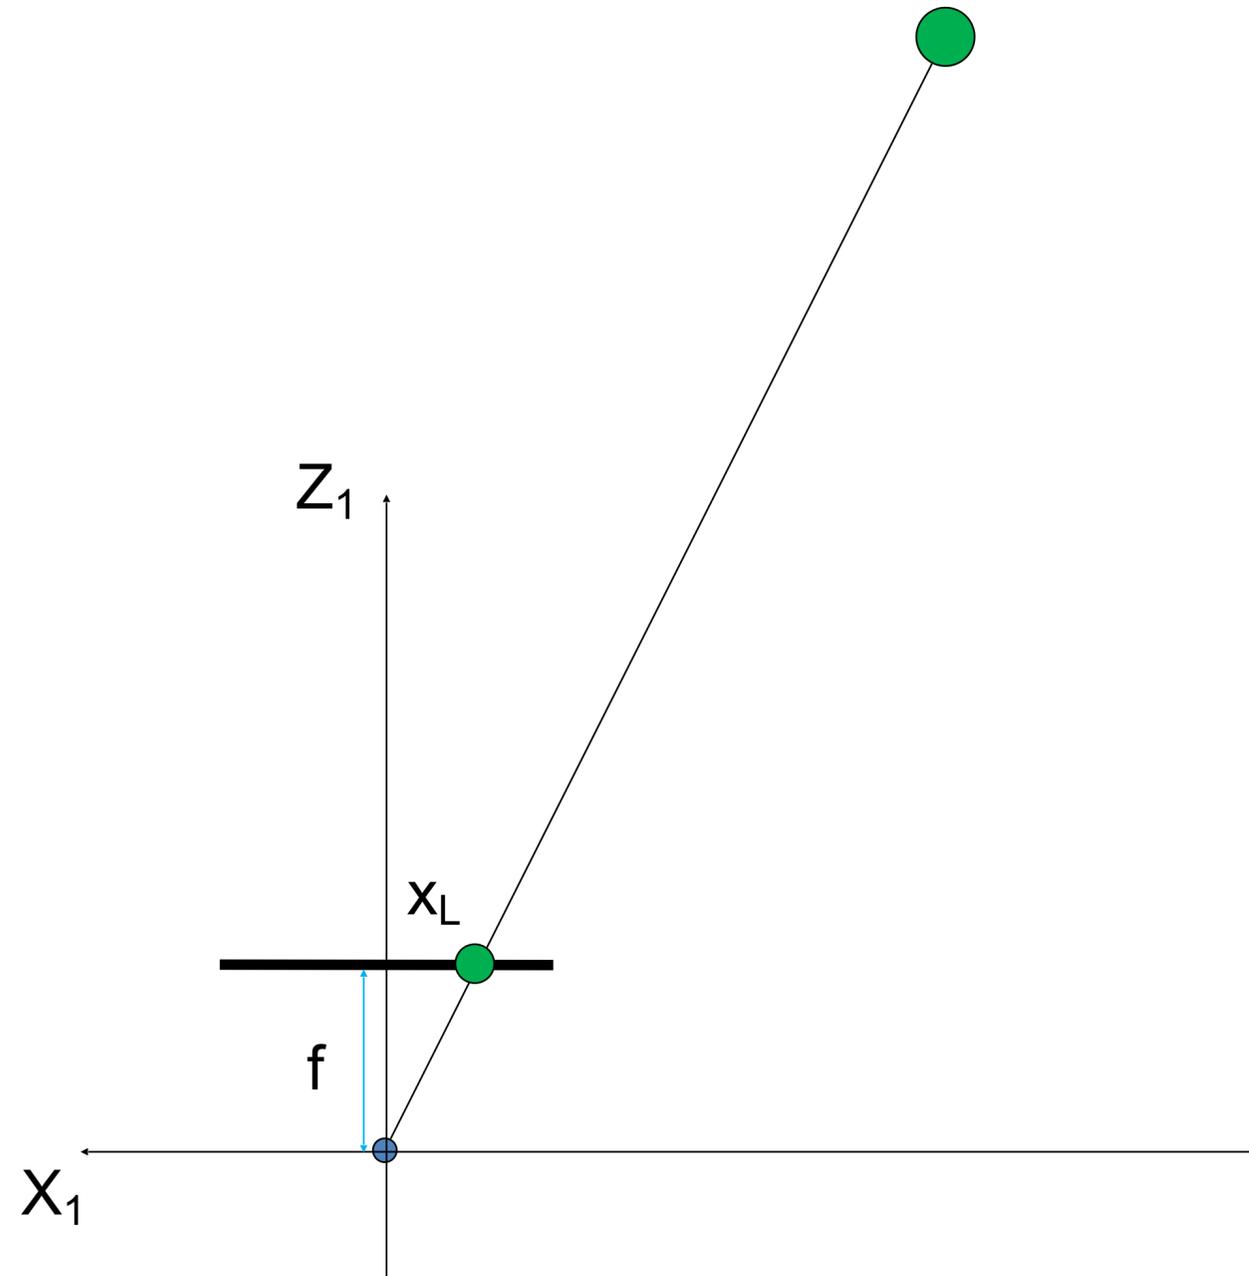


FIGURE 8.13

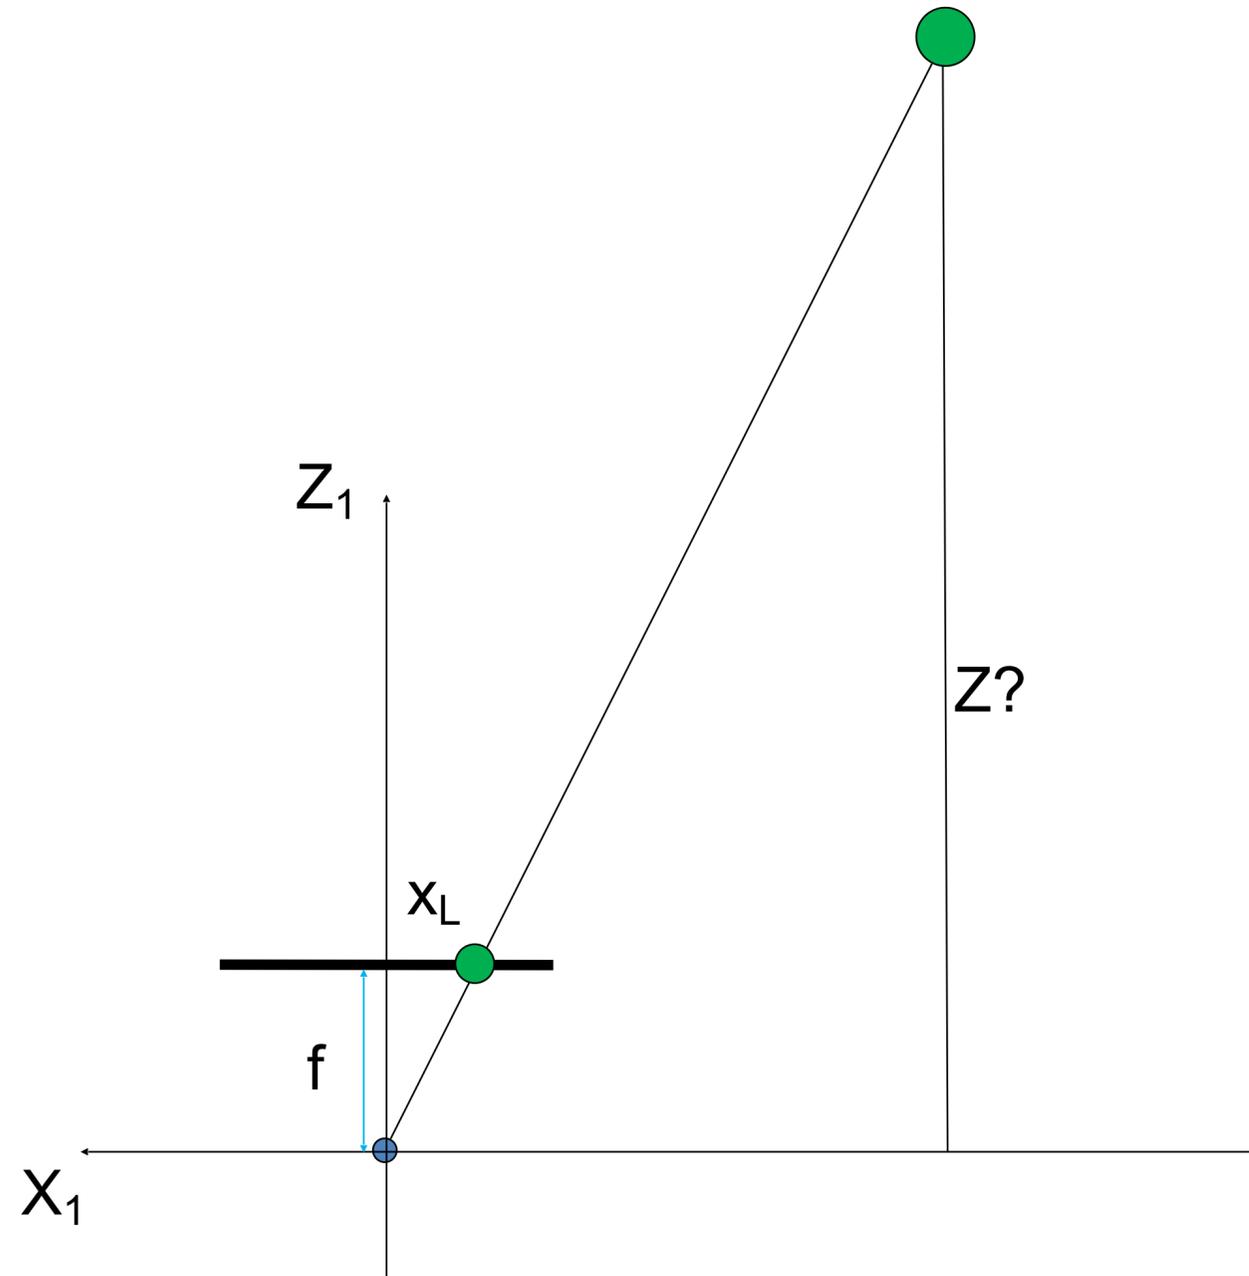
1	0	1	0	1	0	0	1	0	1
1	0	0	1	0	1	0	1	0	0
0	0	1	1	0	1	1	0	1	0
0	1	0	Y	A	A	B	B	0	1
1	1	1	X	B	A	B	A	0	1
0	0	1	X	A	A	B	A	1	0
1	1	1	Y	B	B	A	B	0	1
1	0	0	1	1	0	1	1	0	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

1	0	1	0	1	0	0	1	0	1
1	0	0	1	0	1	0	1	0	0
0	0	1	1	0	1	1	0	1	0
0	1	0	A	A	B	B	X	0	1
1	1	1	B	A	B	A	Y	0	1
0	0	1	A	A	B	A	Y	1	0
1	1	1	B	B	A	B	X	0	1
1	0	0	1	1	0	1	1	0	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

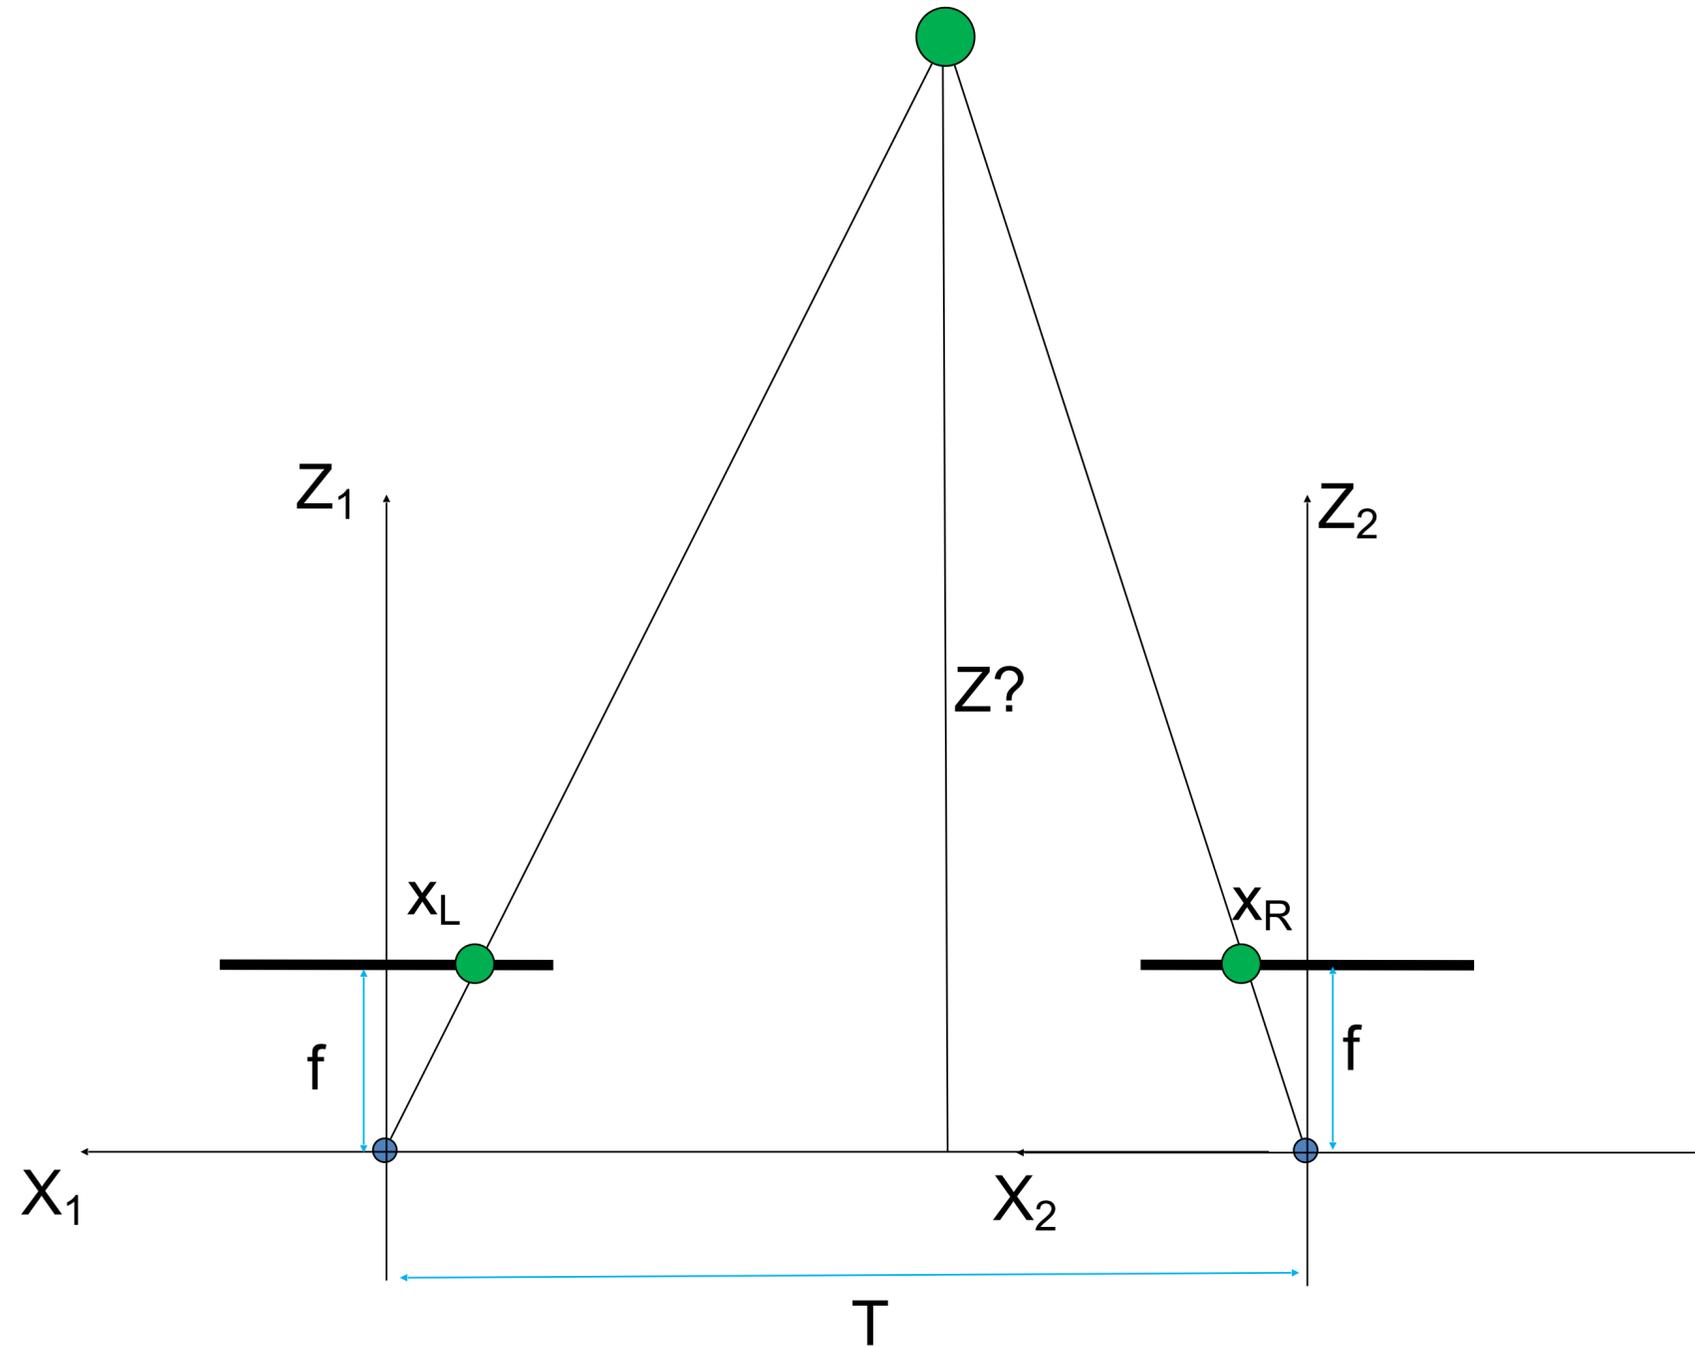
# Geometry for a simple stereo system



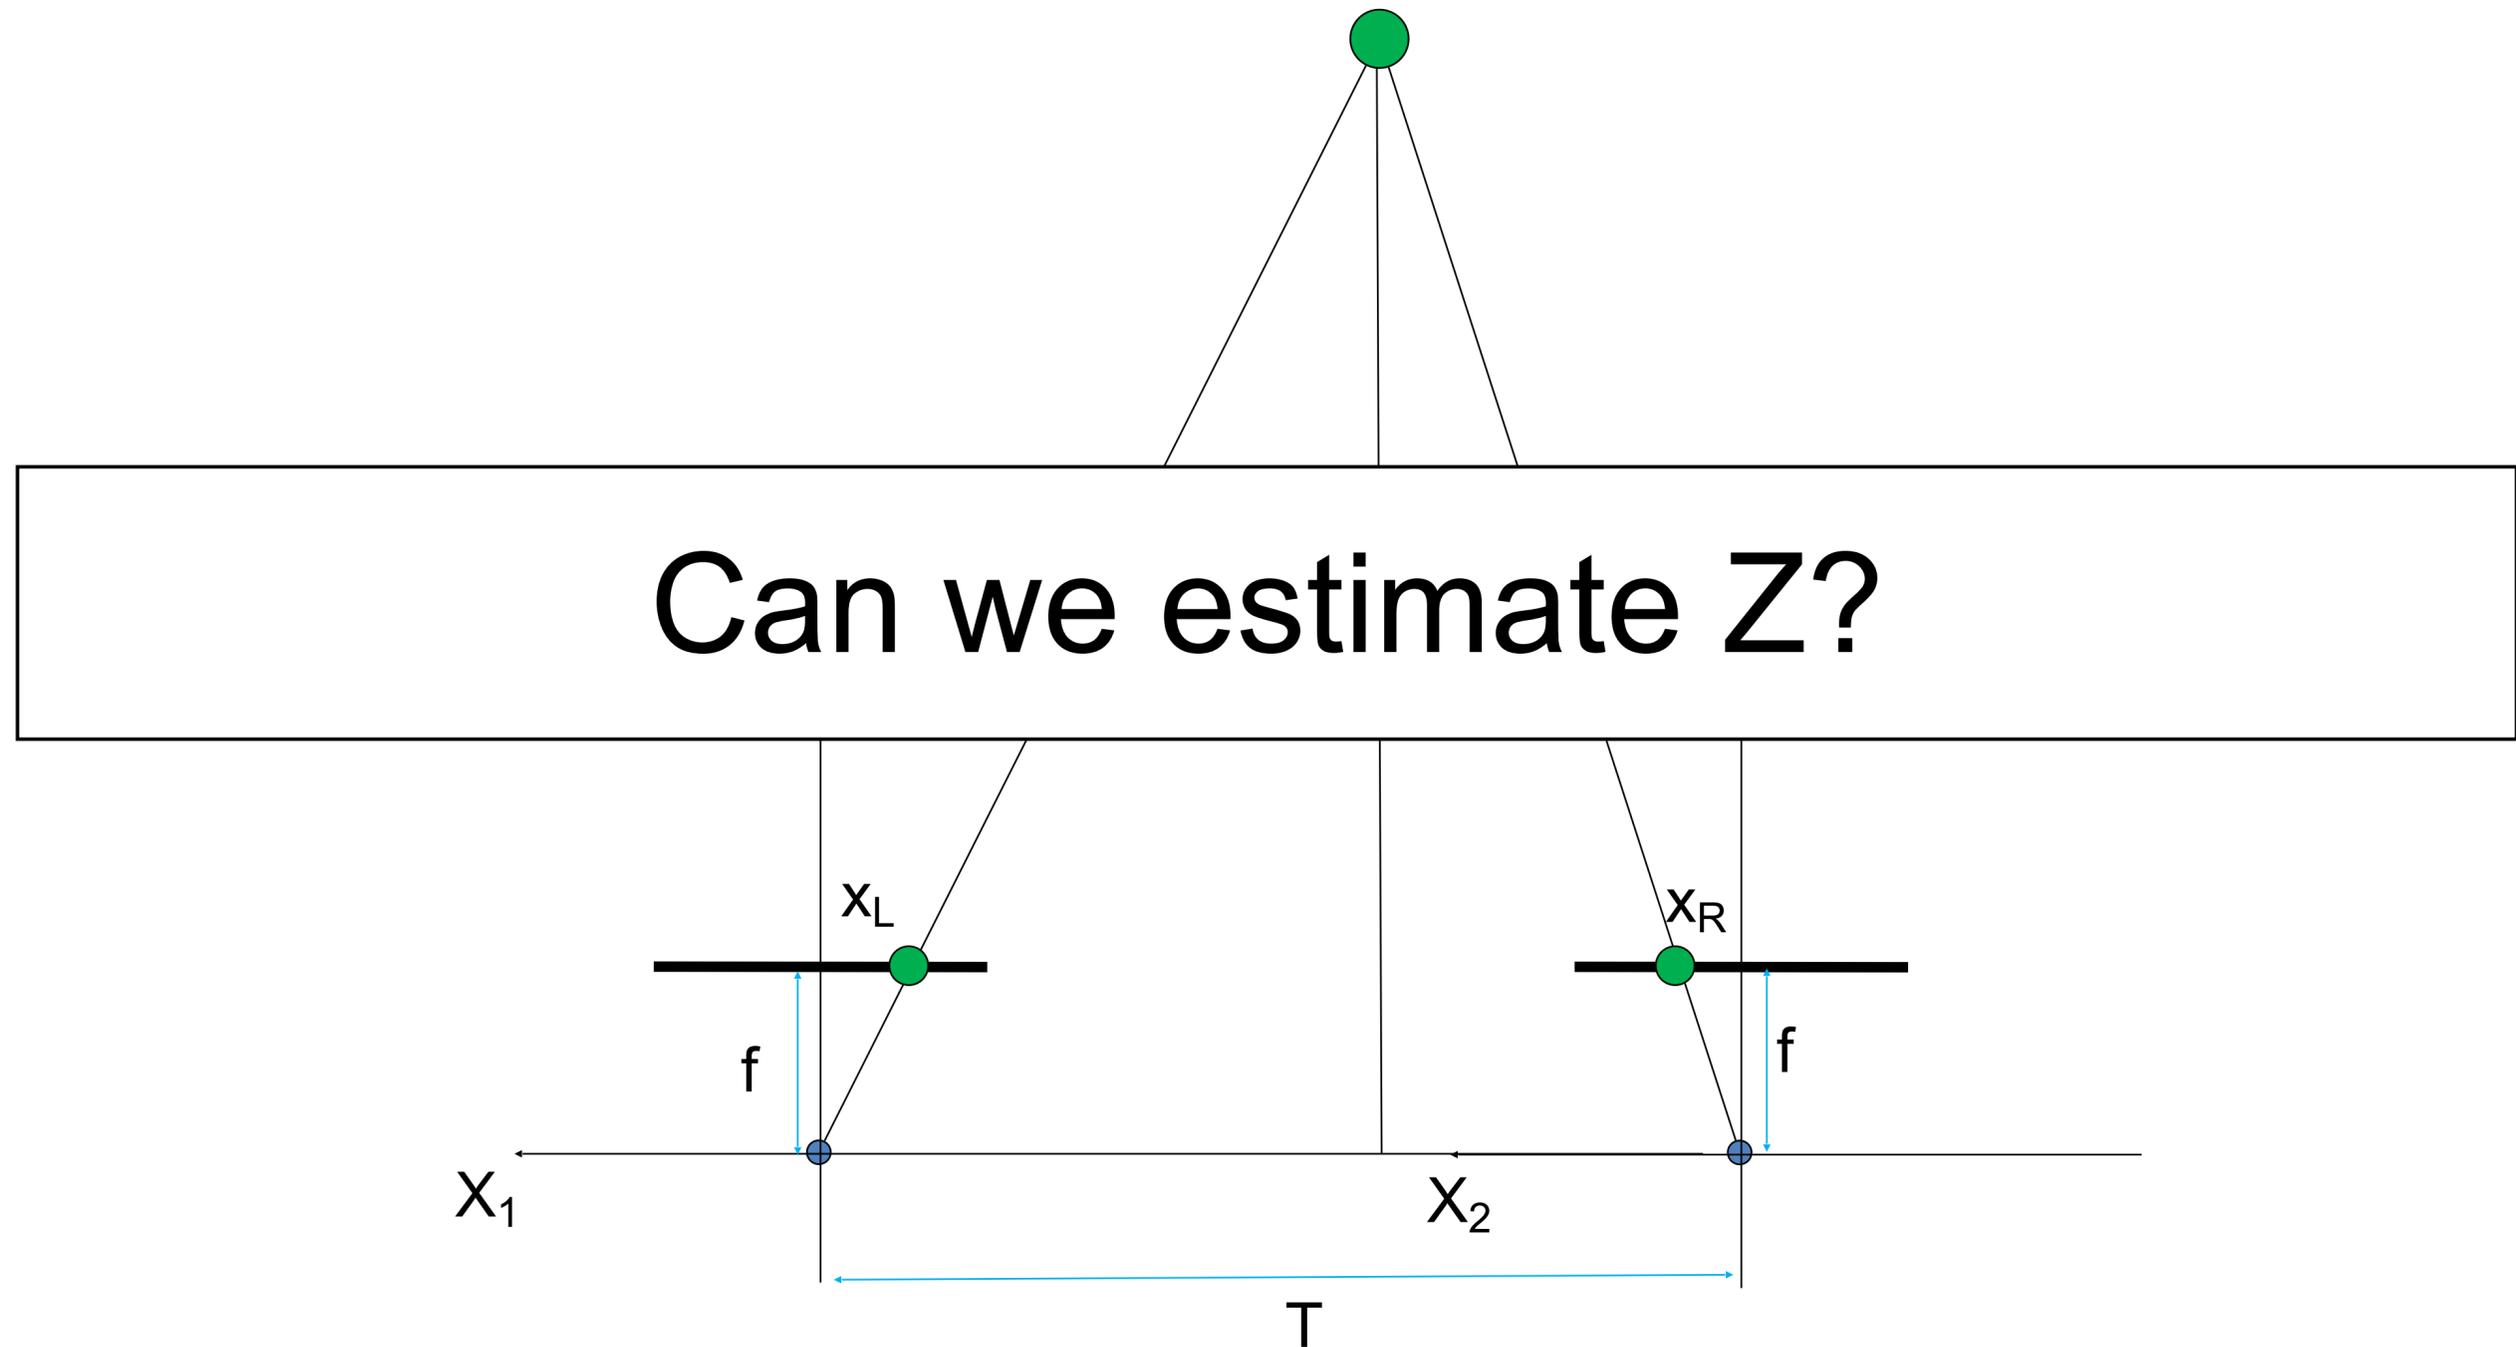
# Geometry for a simple stereo system



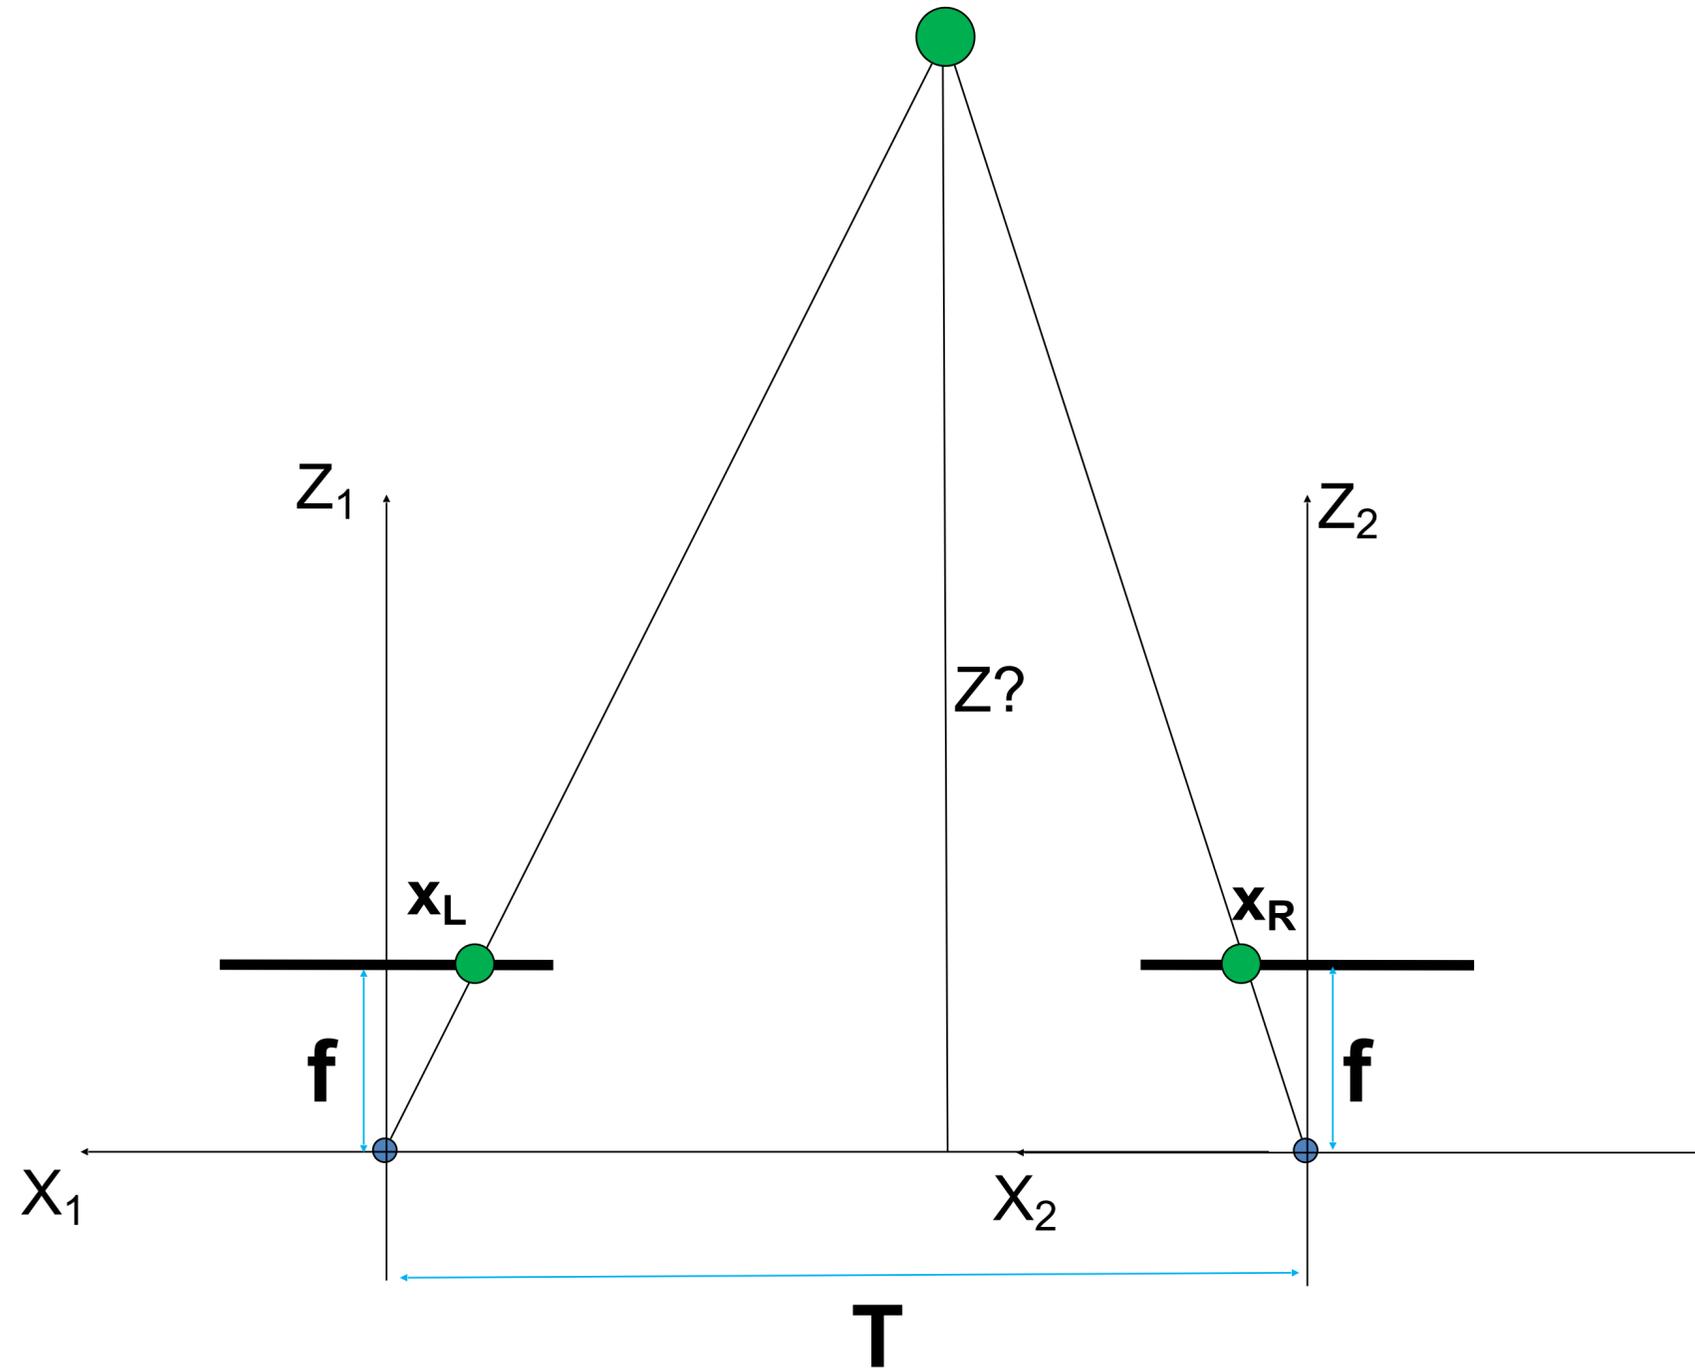
# Geometry for a simple stereo system



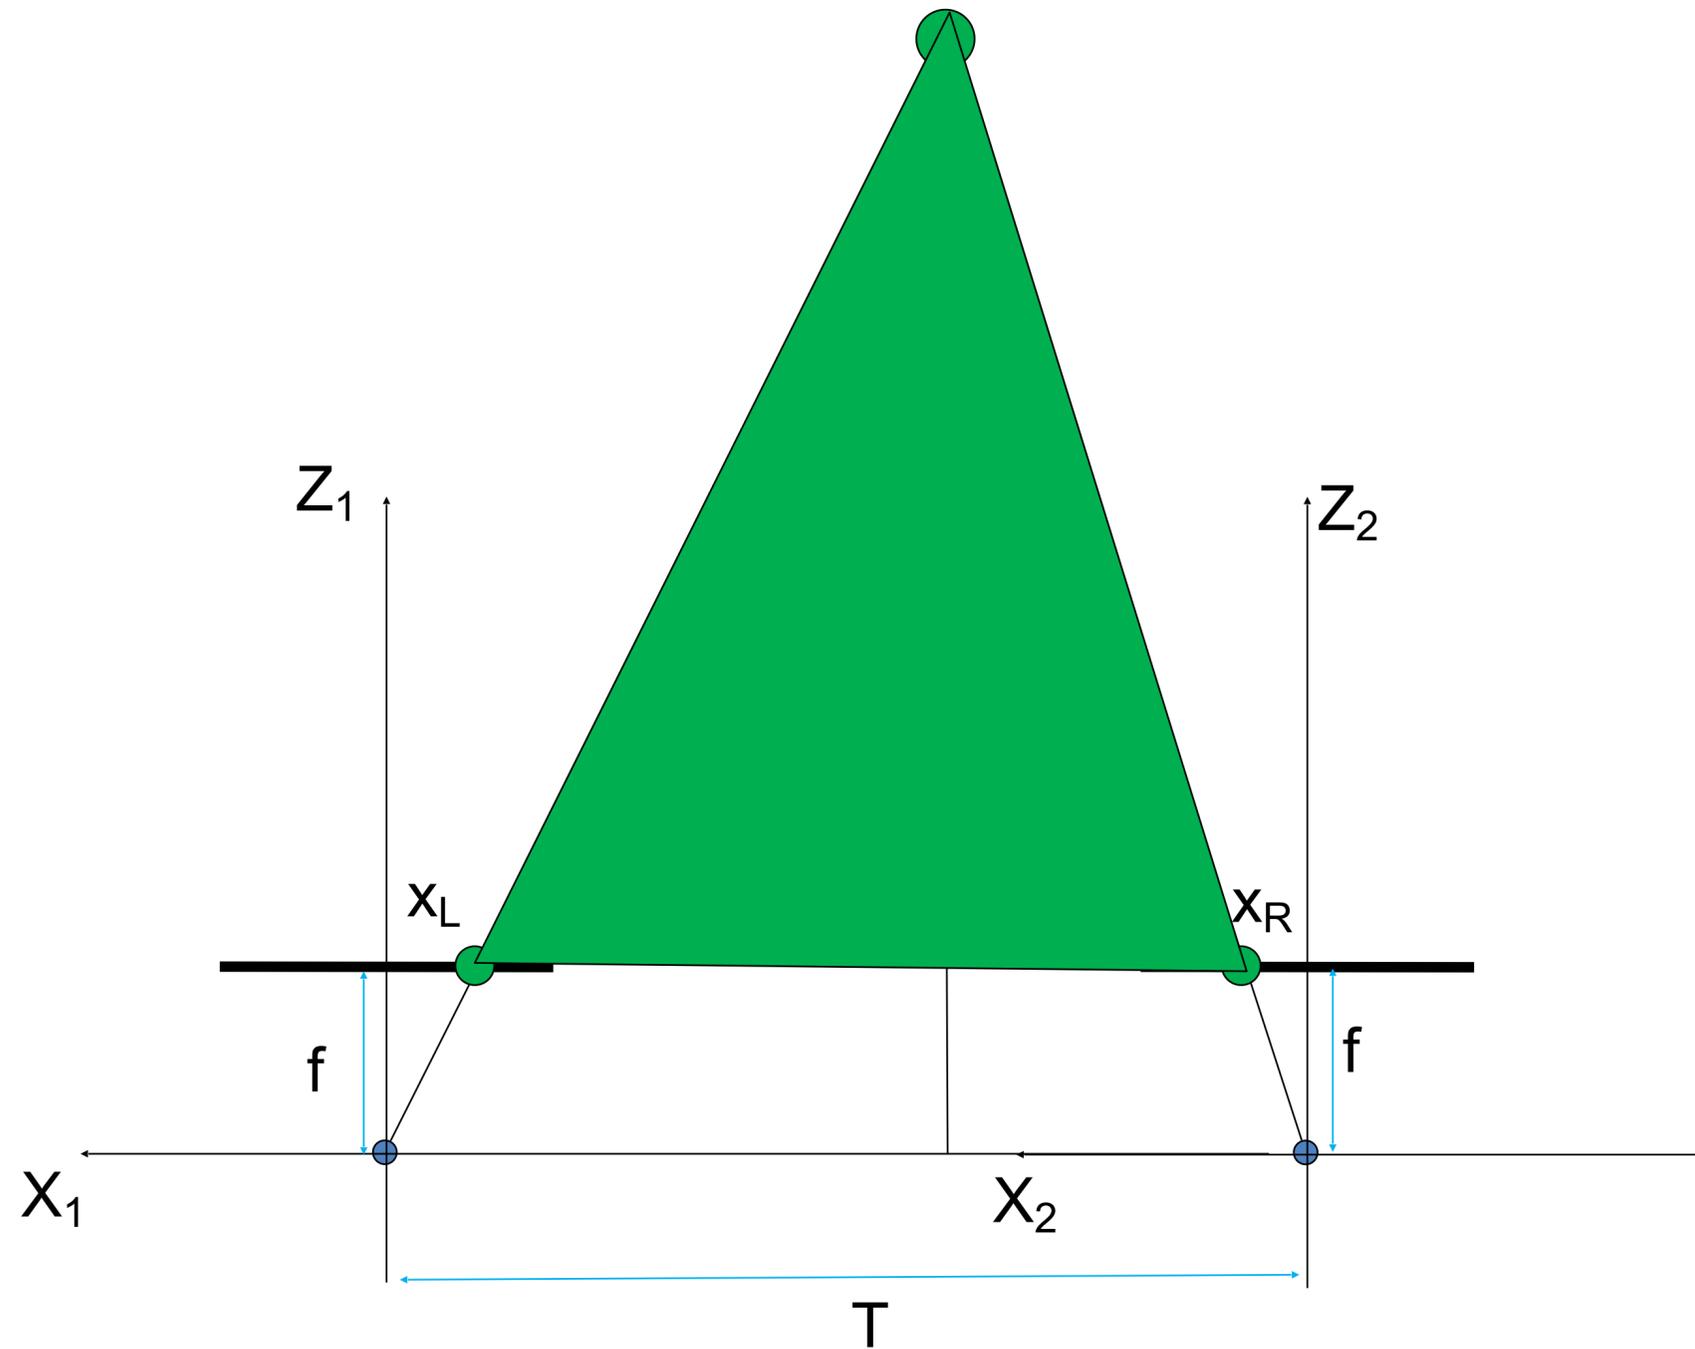
# Geometry for a simple stereo system



# Geometry for a simple stereo system

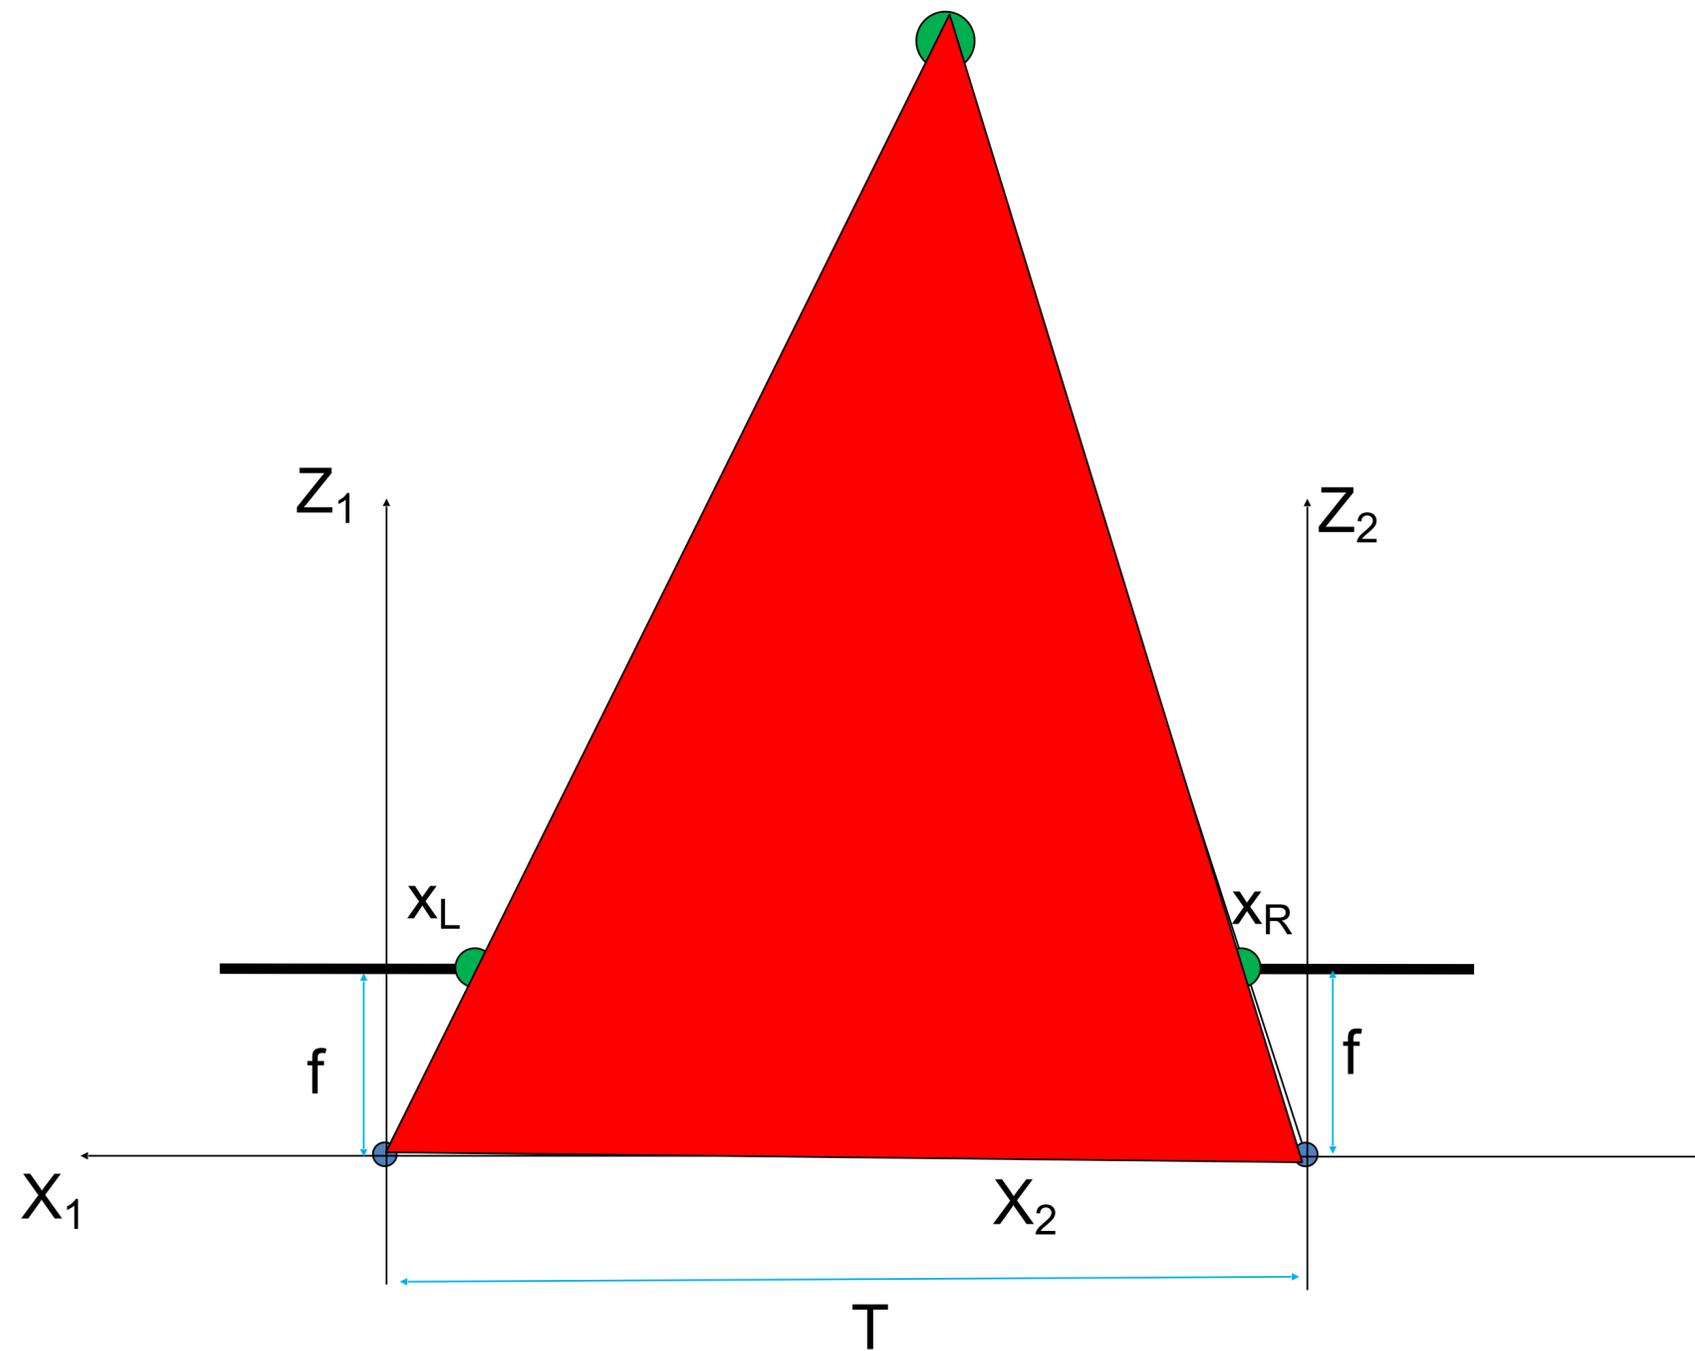


# Geometry for a simple stereo system



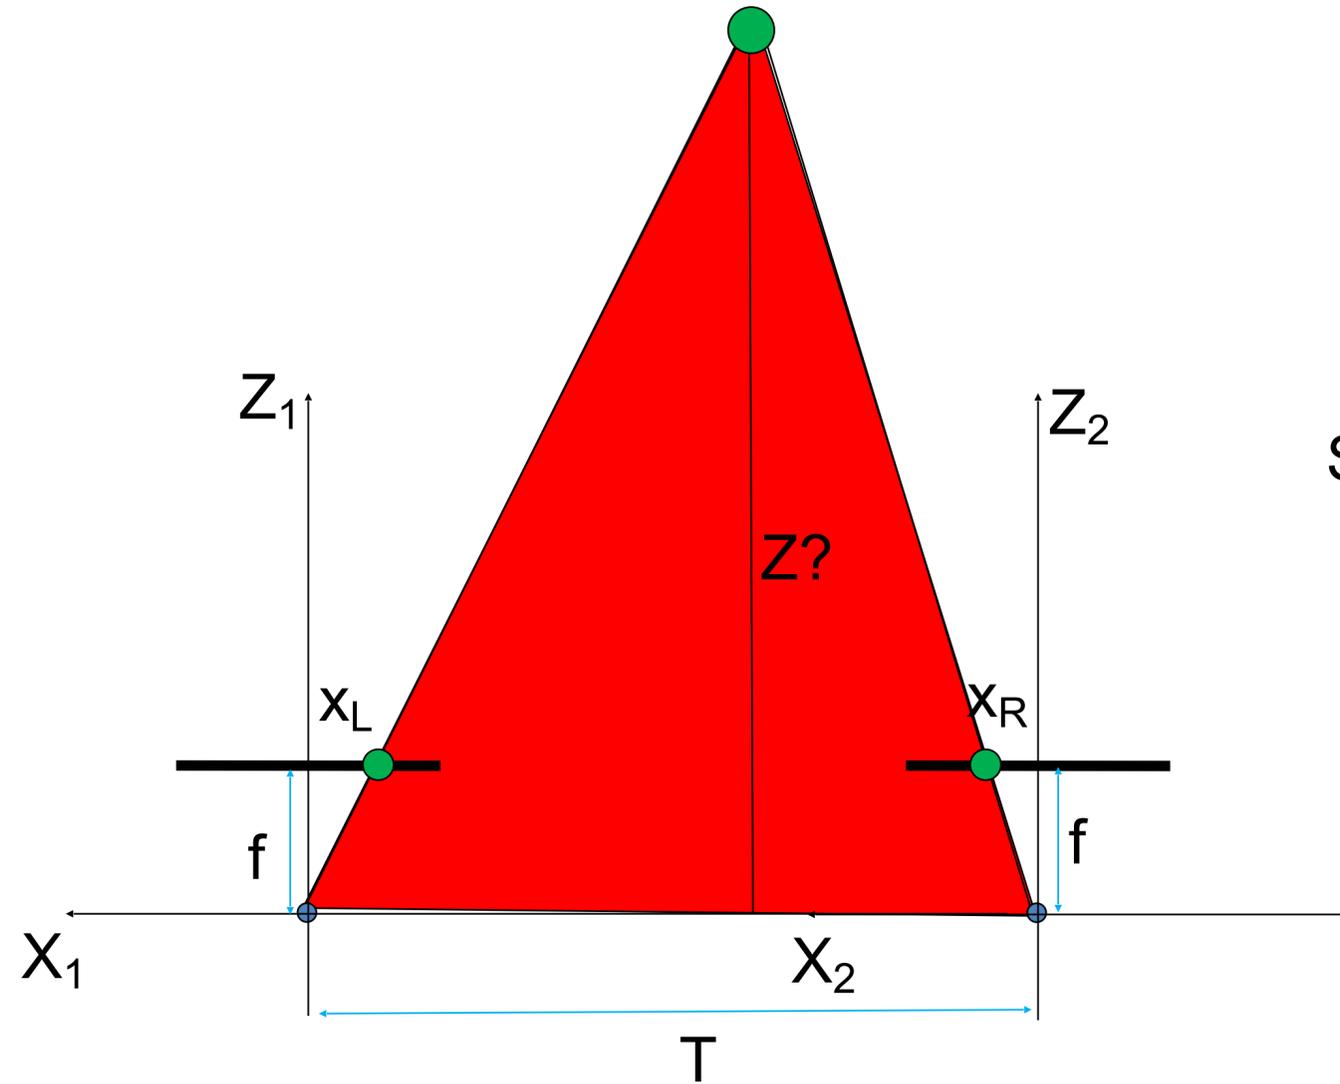
Similar triangles

# Geometry for a simple stereo system



Similar triangles

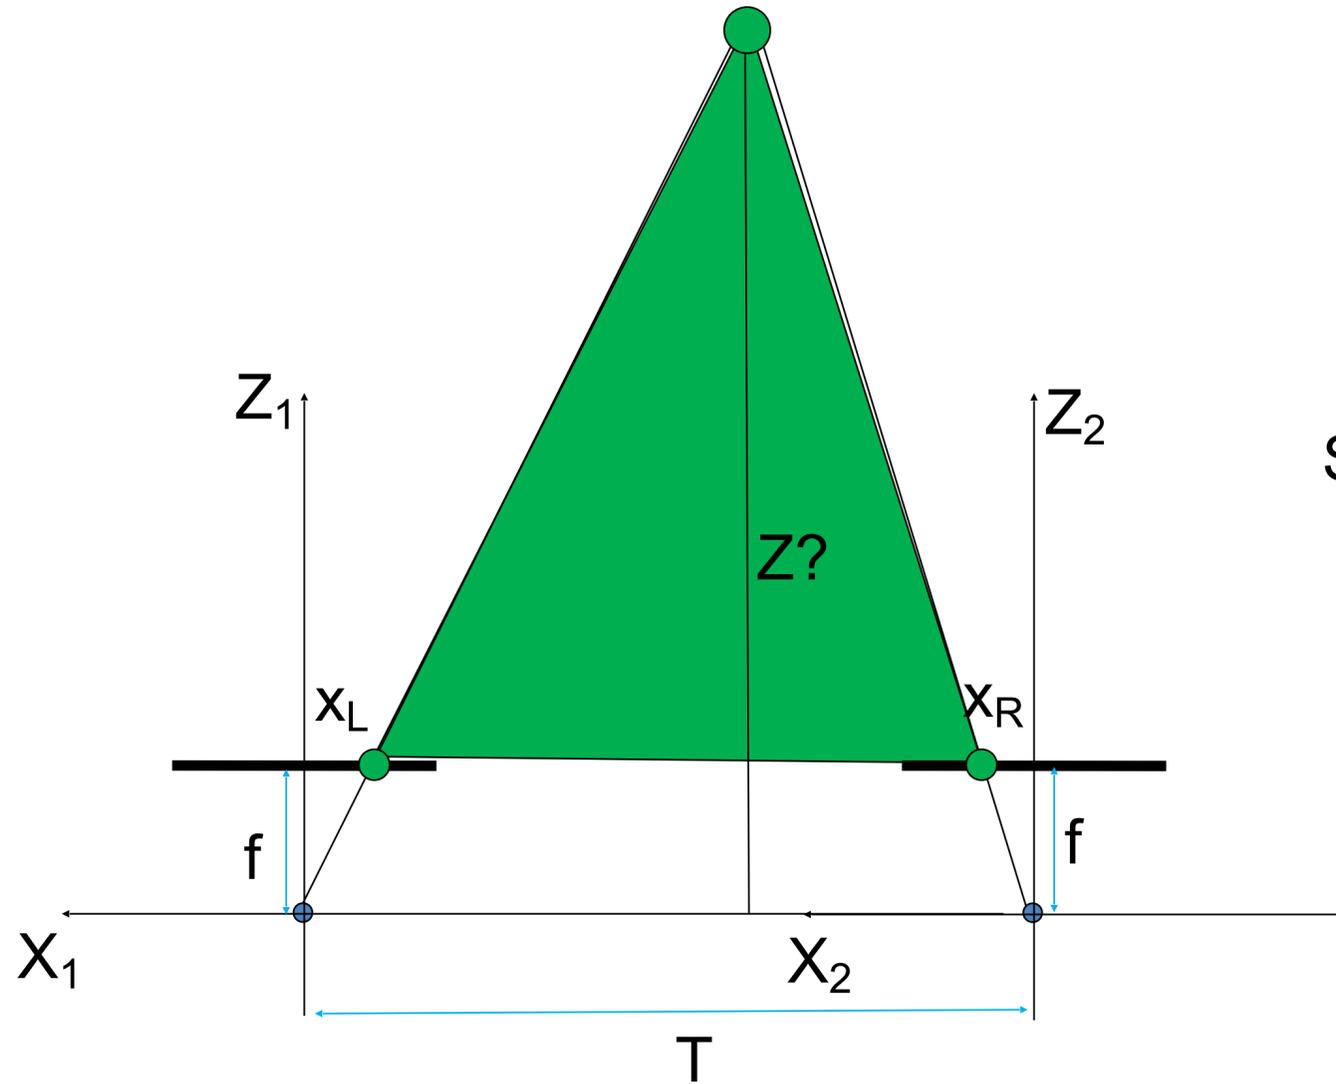
# Geometry for a simple stereo system



Similar triangles:

$$\frac{T}{Z}$$

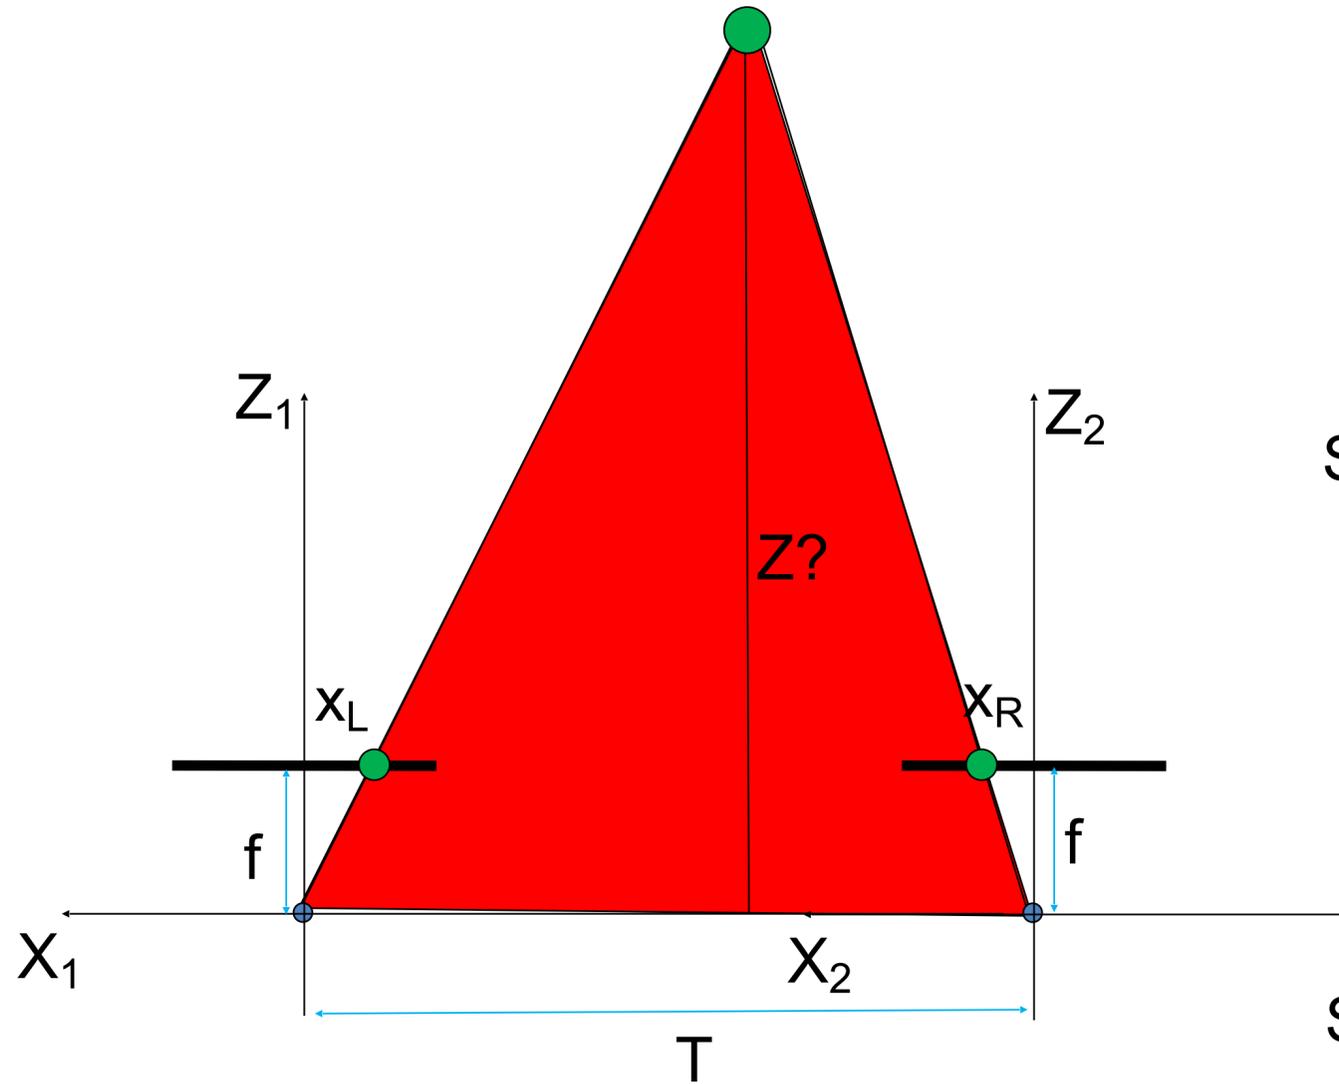
# Geometry for a simple stereo system



Similar triangles:

$$\frac{T + X_R - X_L}{Z - f} = \frac{T}{Z}$$

# Geometry for a simple stereo system



Similar triangles:

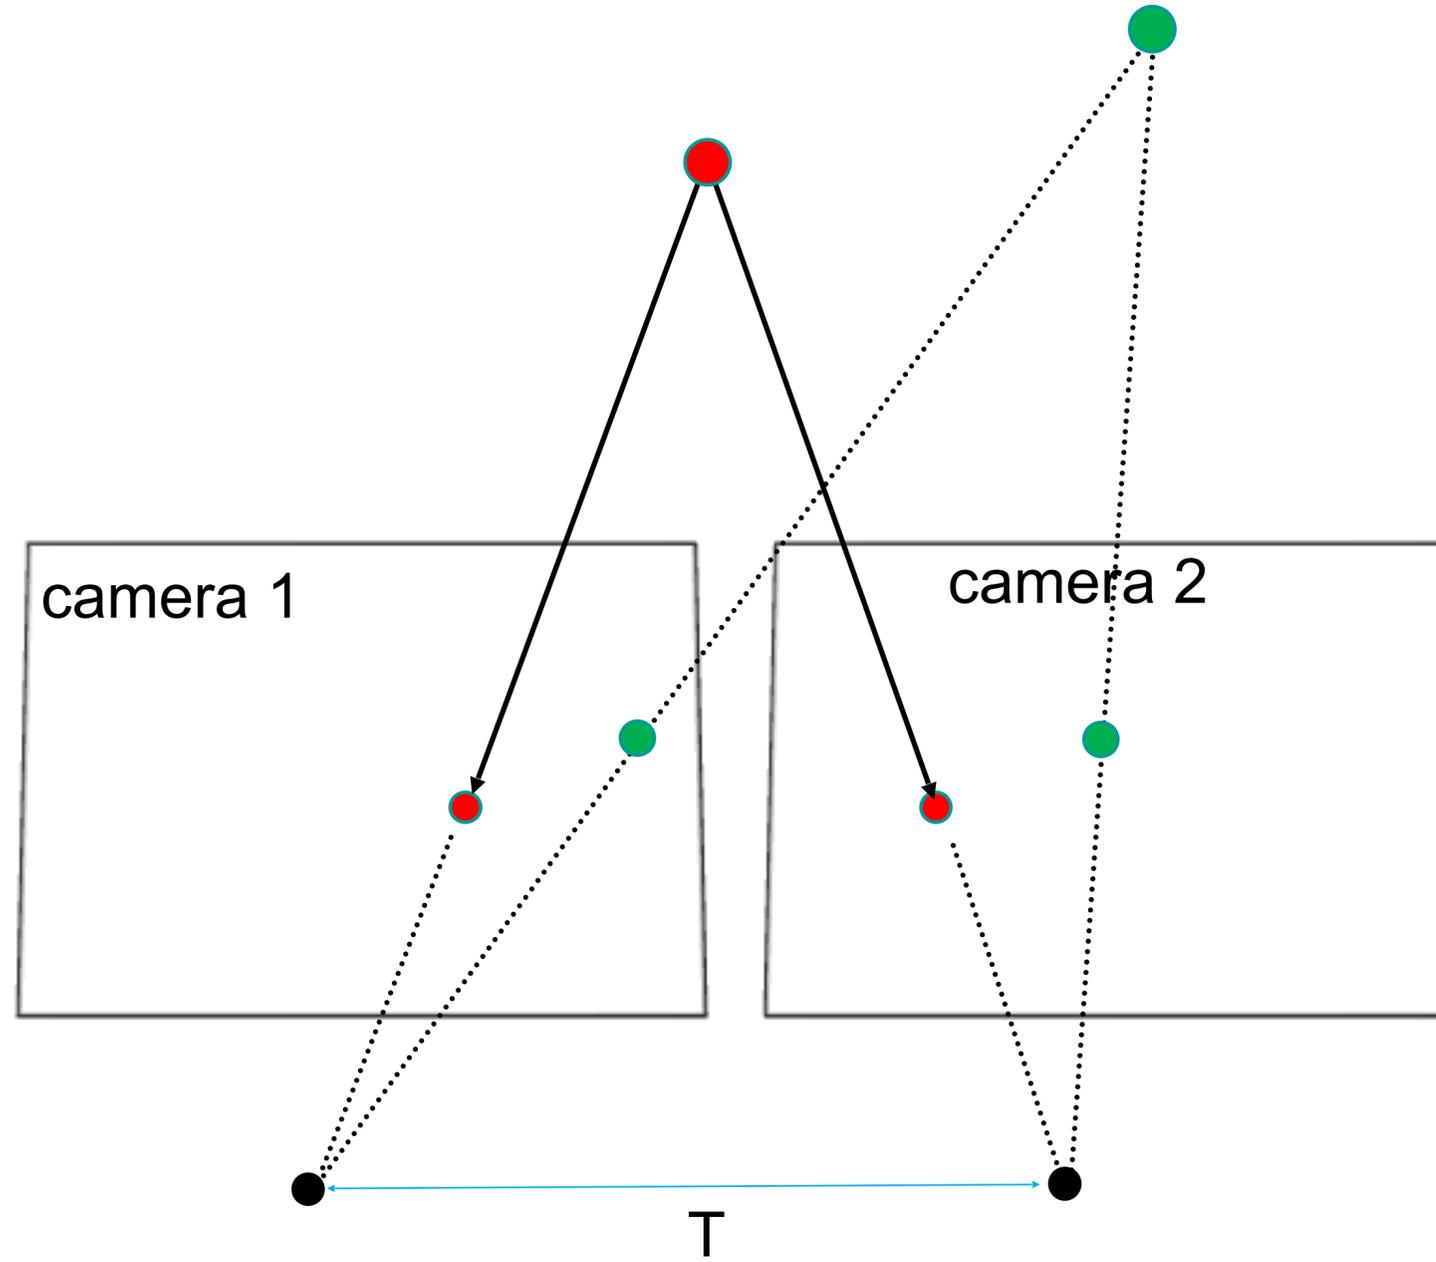
$$\frac{T + X_R - X_L}{Z - f} = \frac{T}{Z}$$

Solving for  $Z$ :

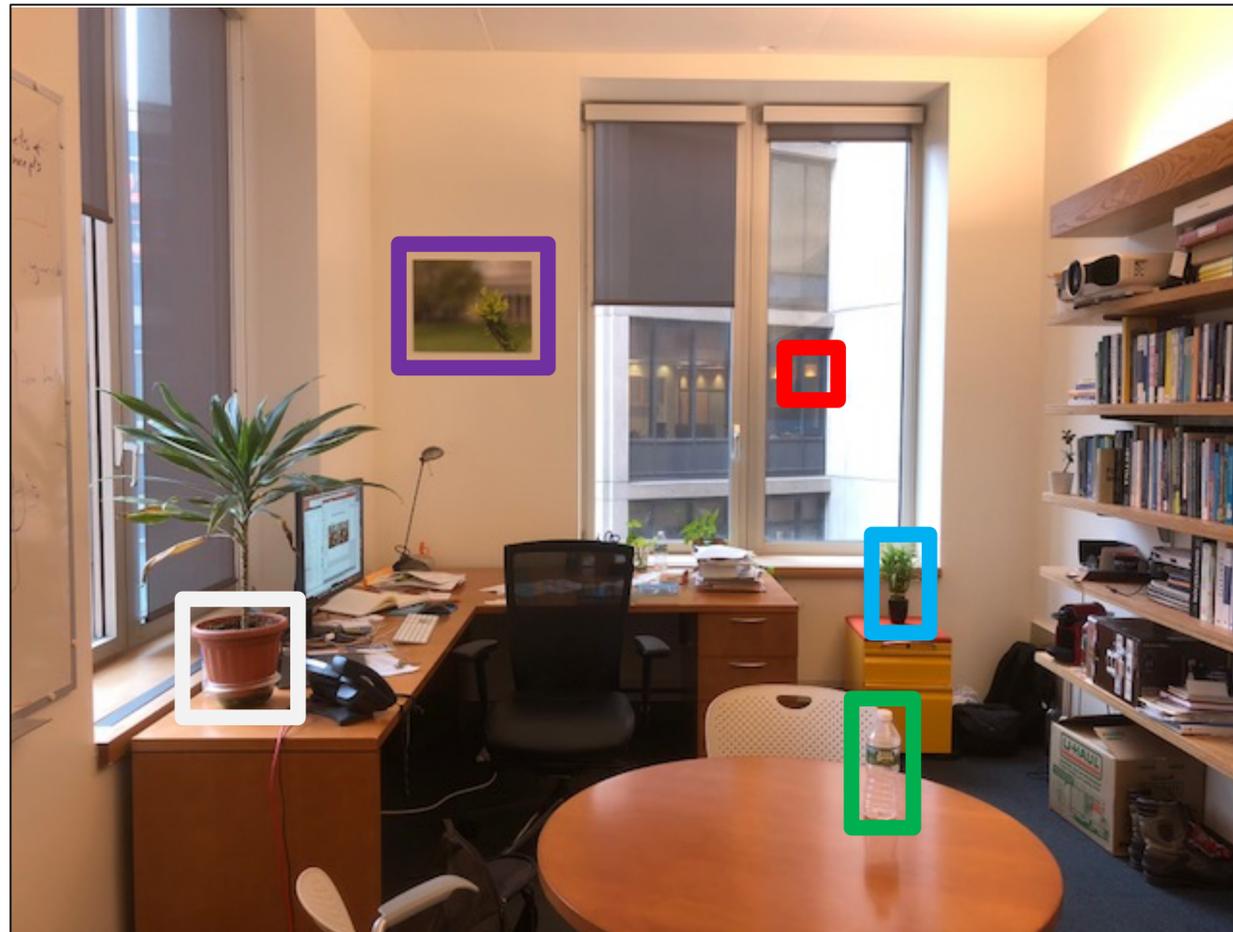
$$Z = f \frac{T}{X_R - X_L}$$

Disparity

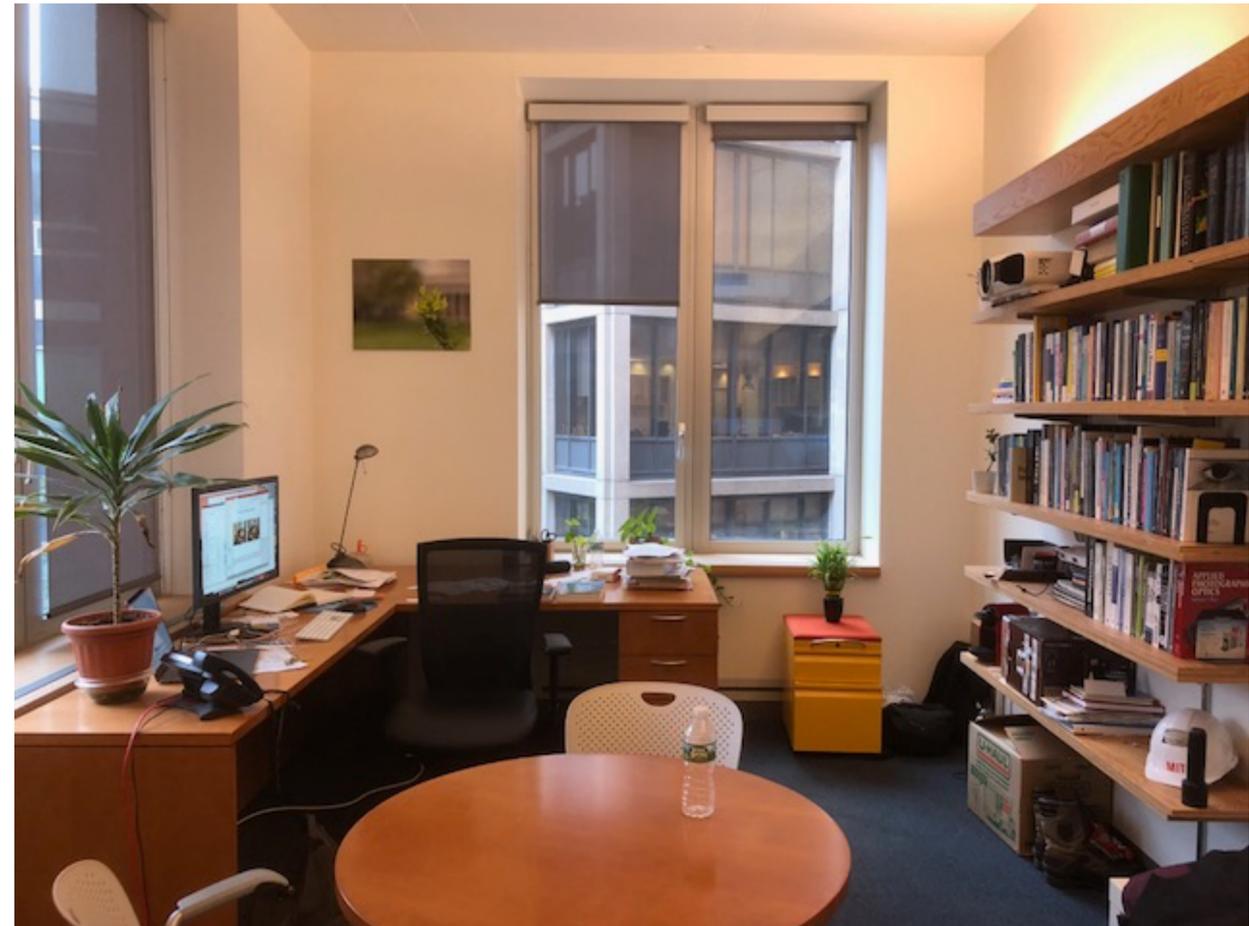
# In 3D



# Disparity map



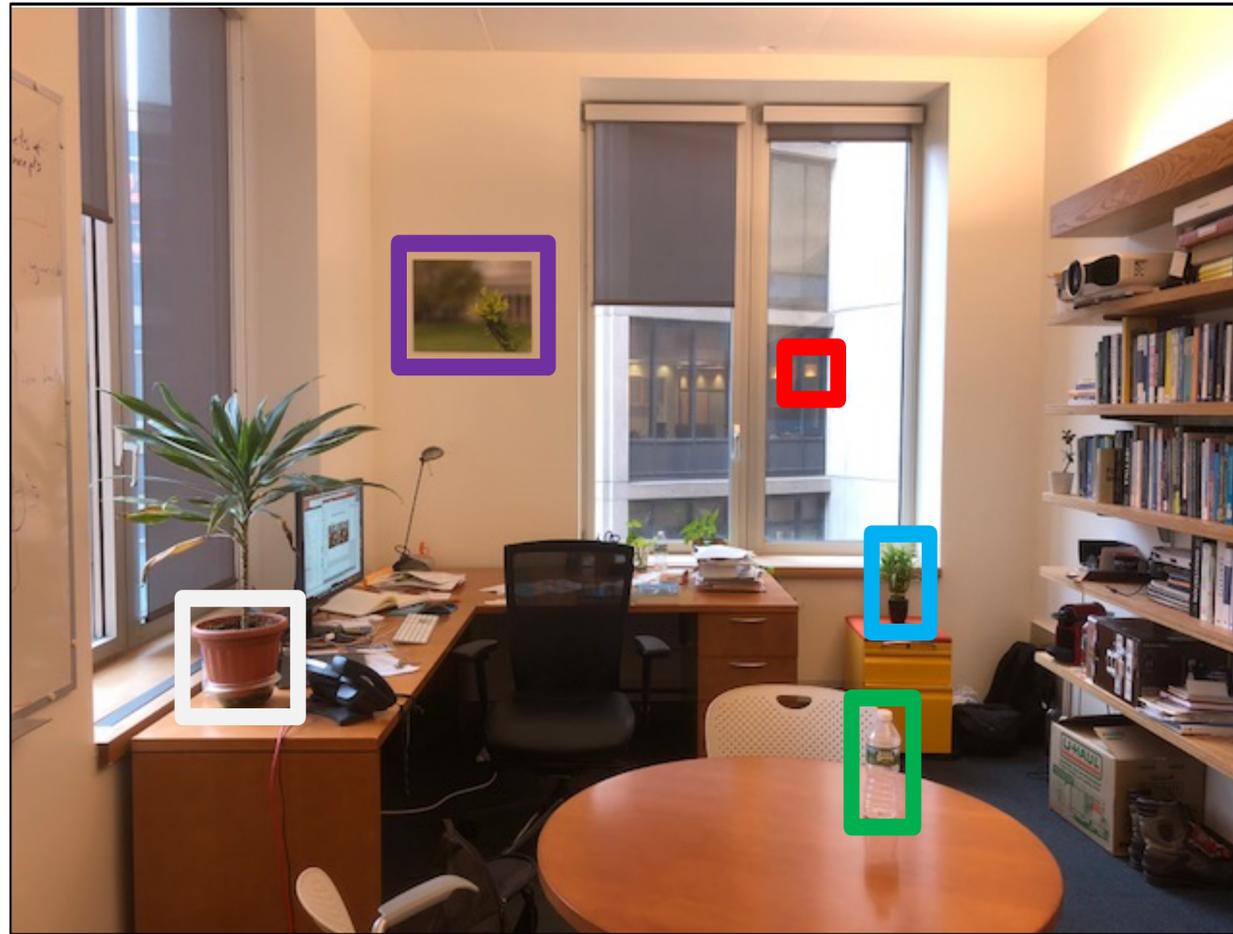
Left image



Right image

Second picture is ~1m to the right

# Disparity map

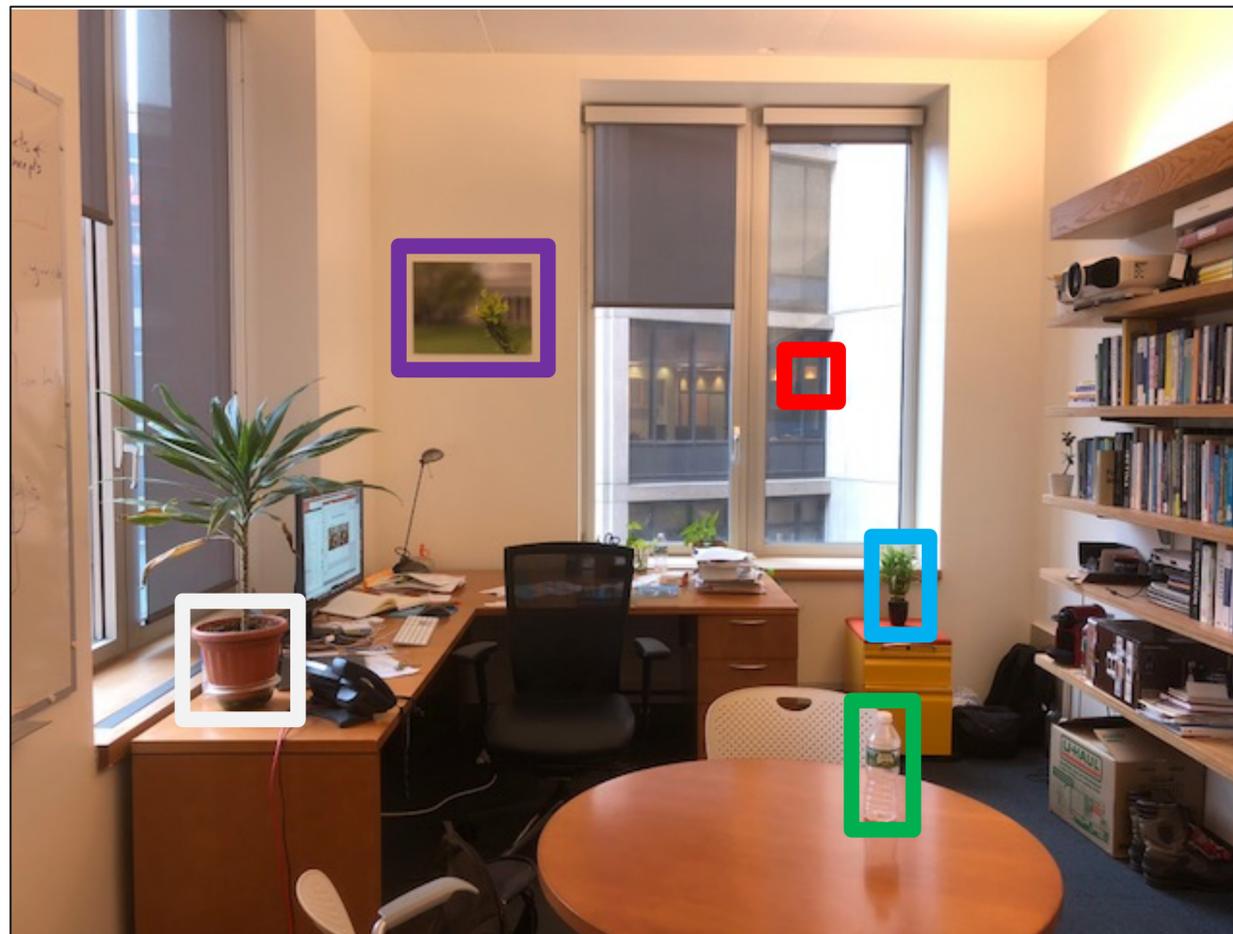


Left image

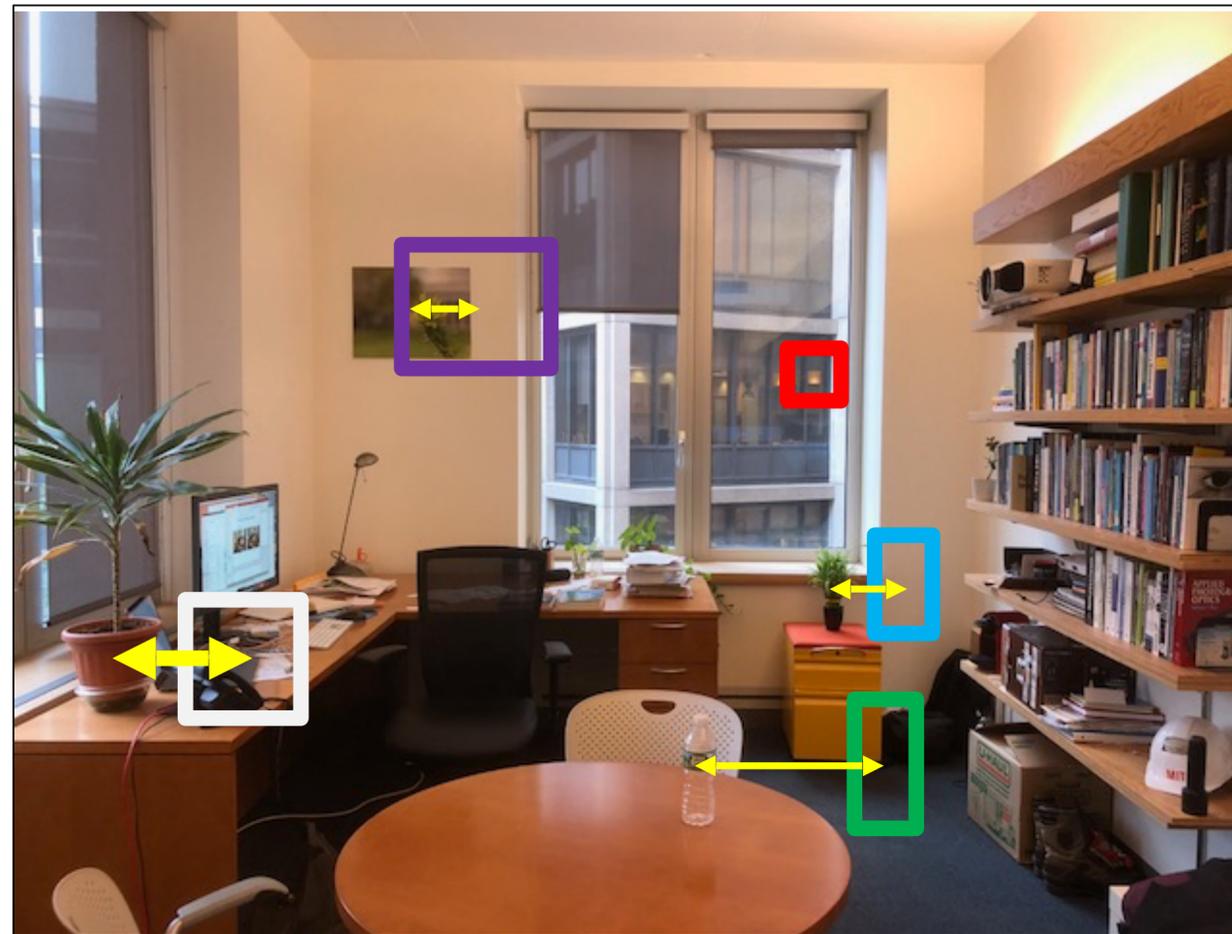


Right image

# Disparity map

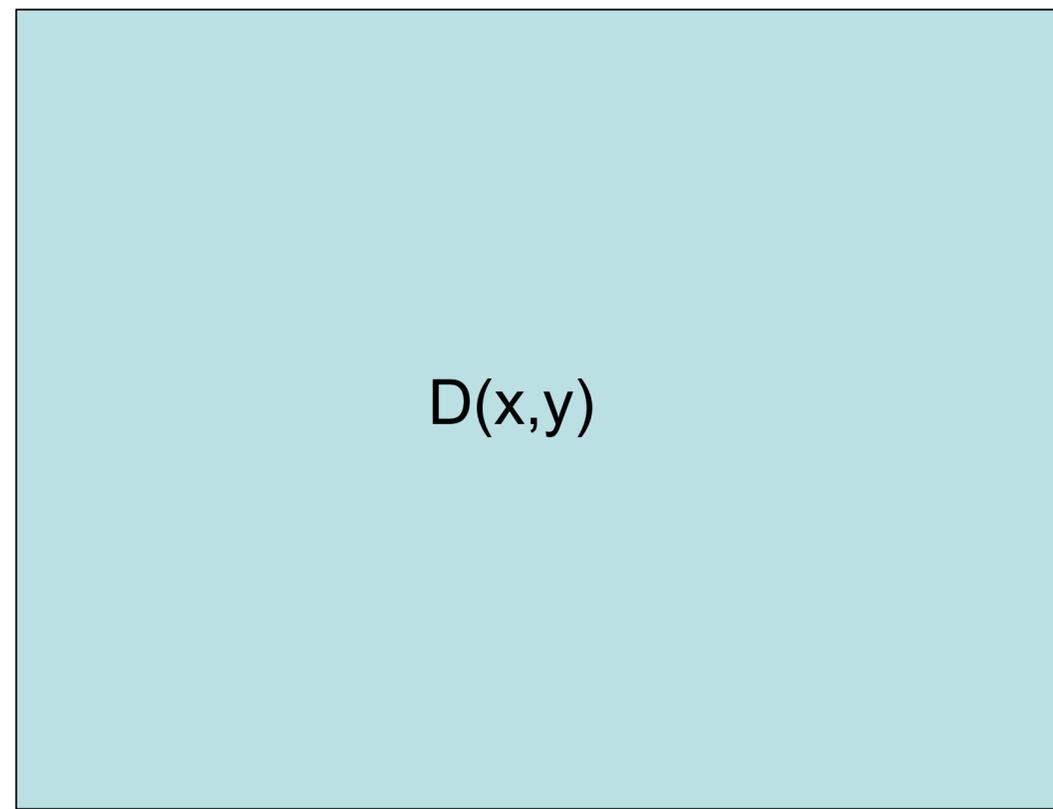
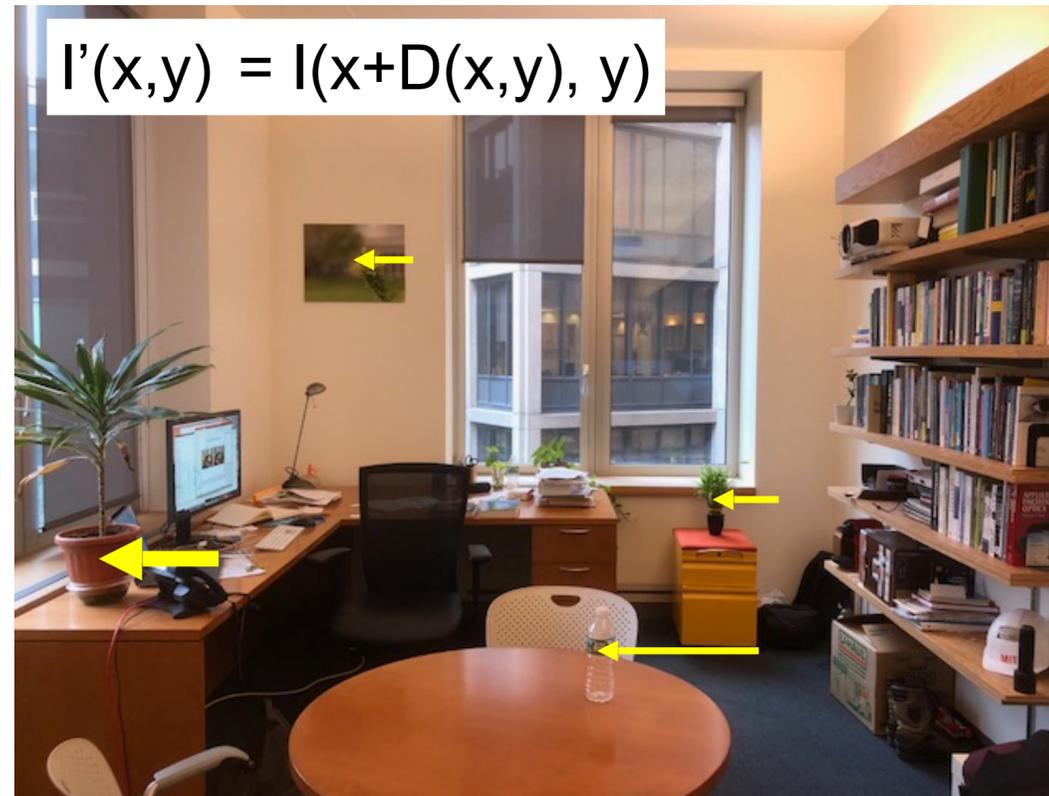
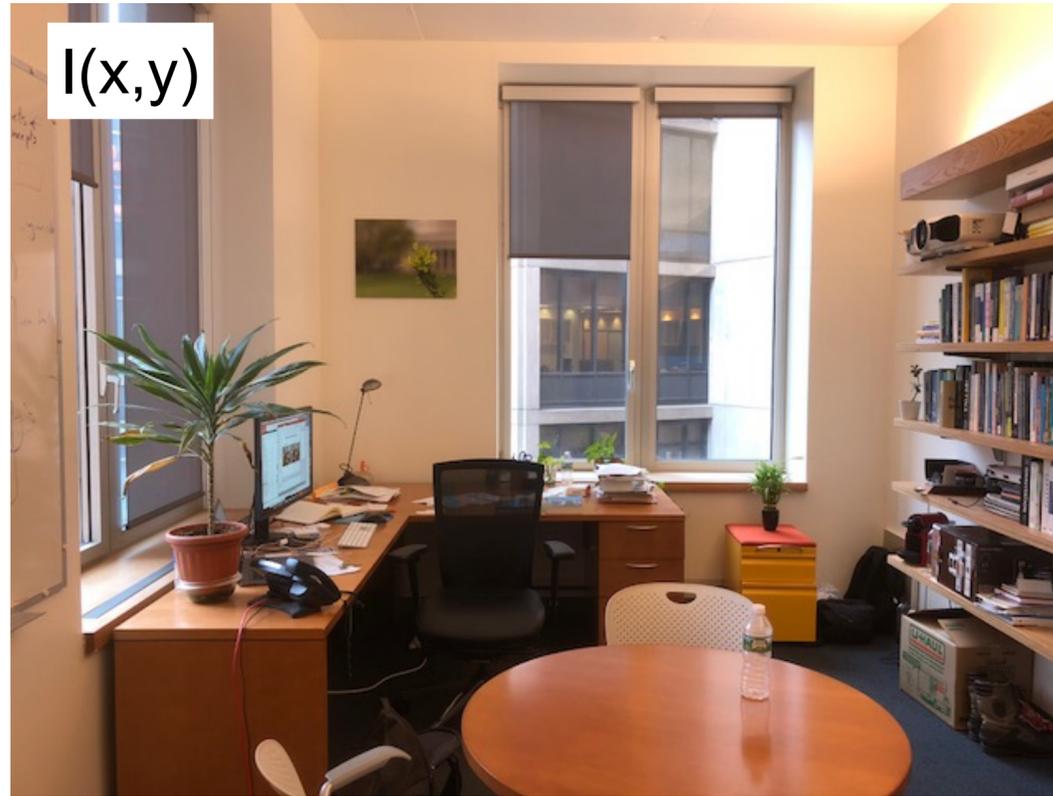


Left image



Right image

# Disparity map



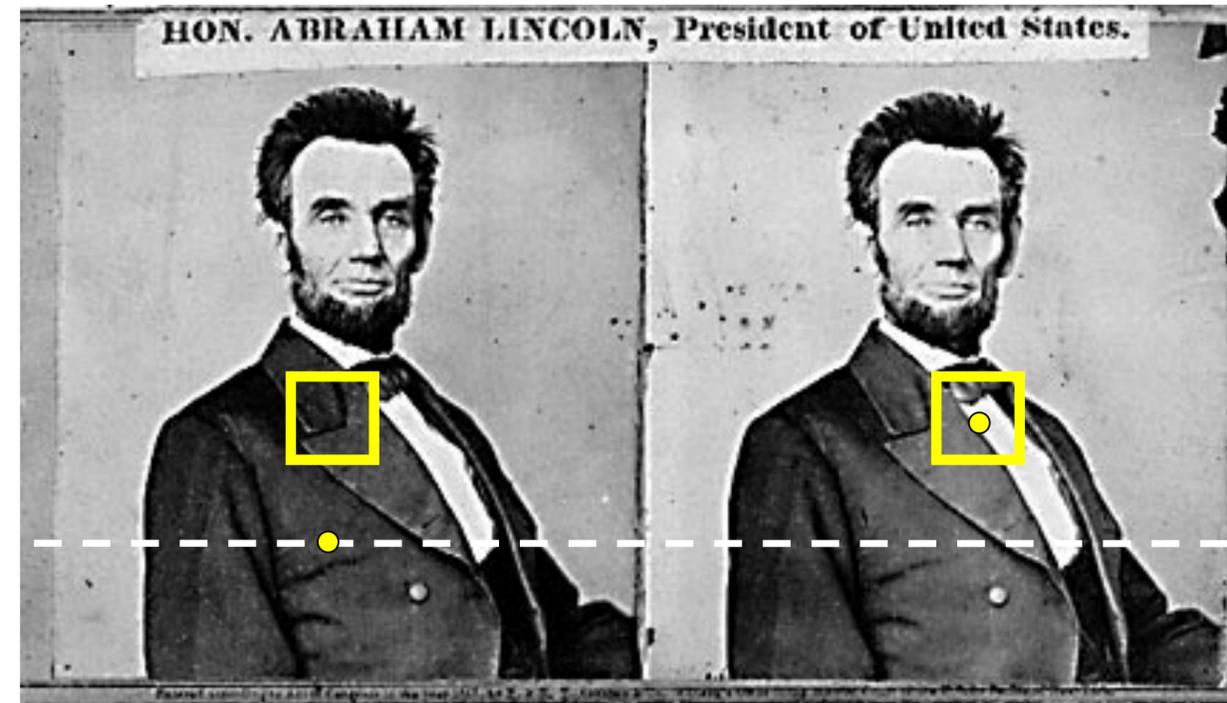
$$Z(x,y) = \frac{f}{D(x,y)}$$

# 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

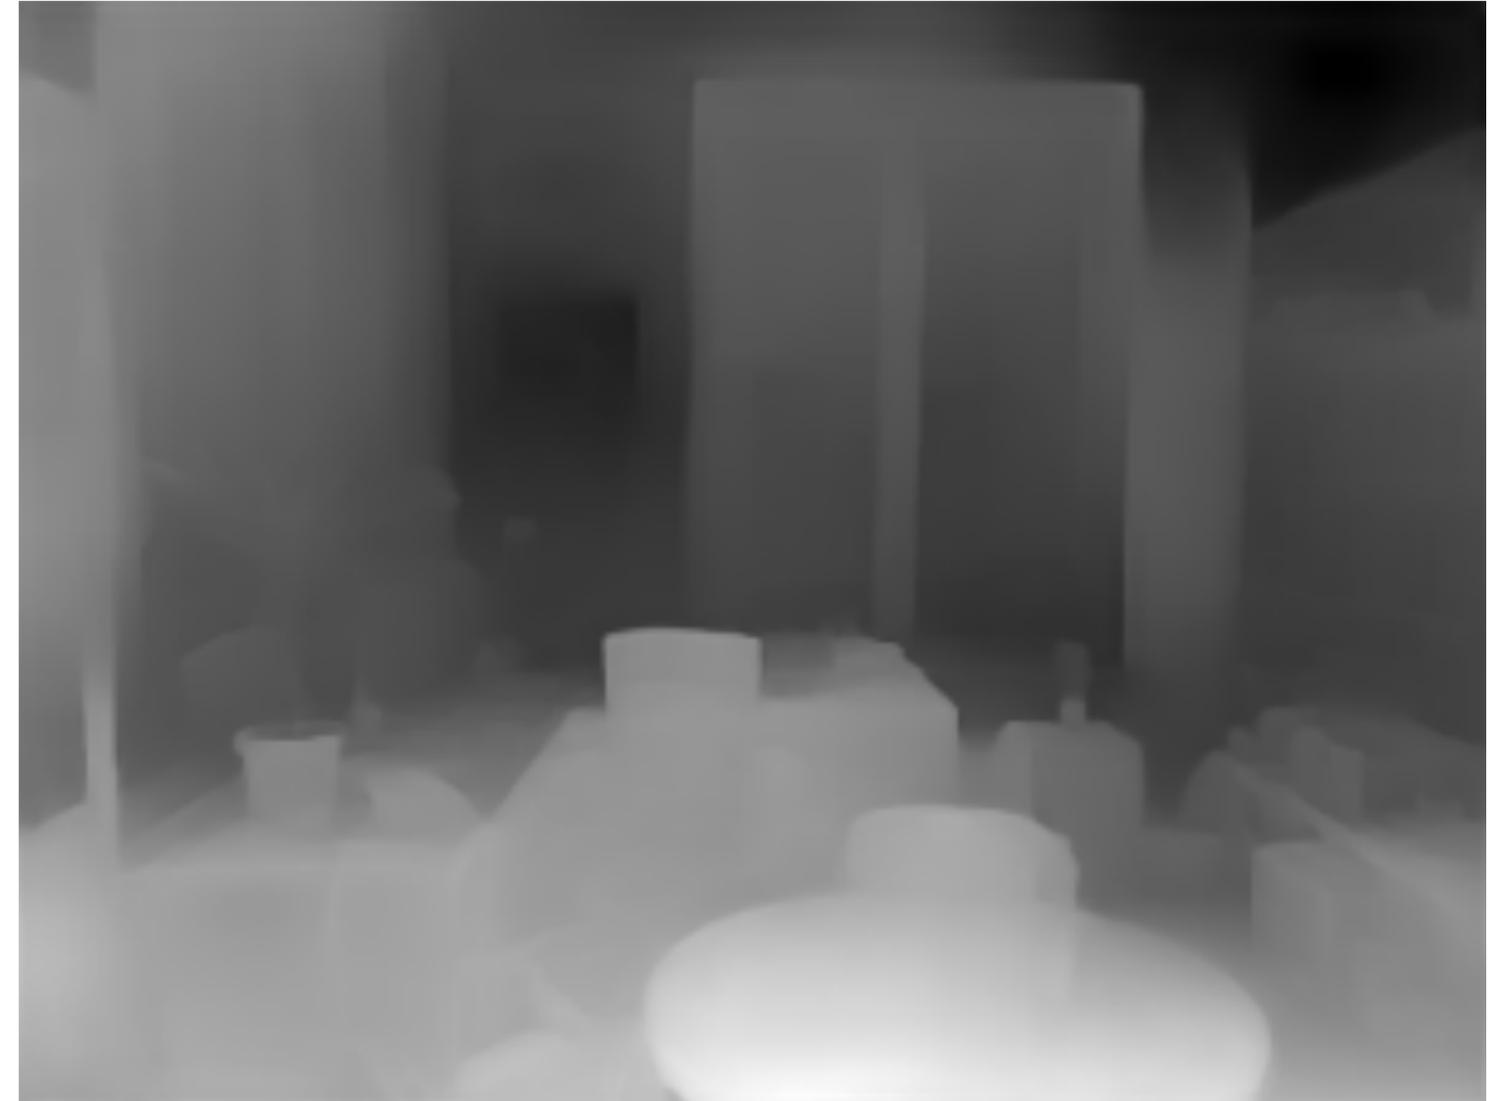


# Computing disparity



Semi-global matching [Hirschmüller 2008]

# Can also learn depth from a single image



**MegaDepth: Learning Single-View Depth Prediction from Internet Photos**

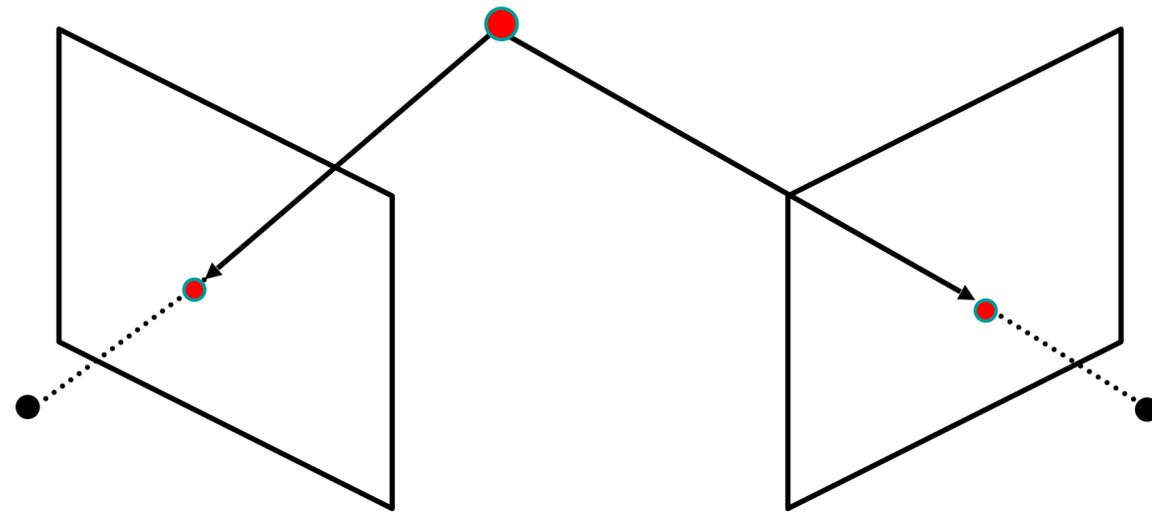
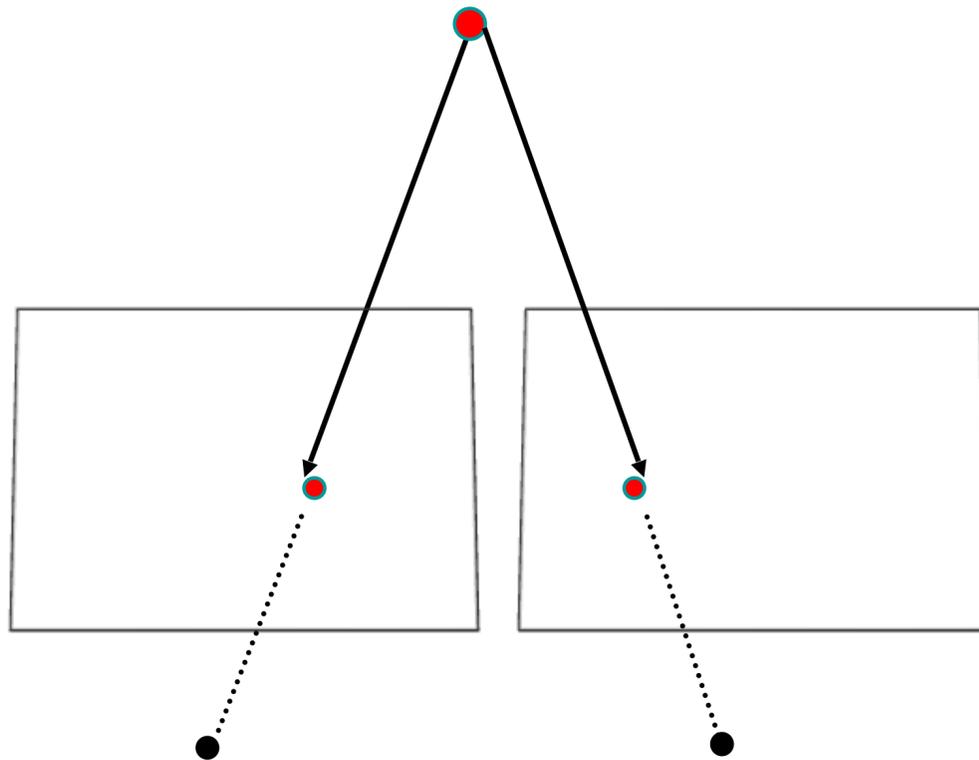
Zhengqi Li    Noah Snavely  
Department of Computer Science & Cornell Tech, Cornell University

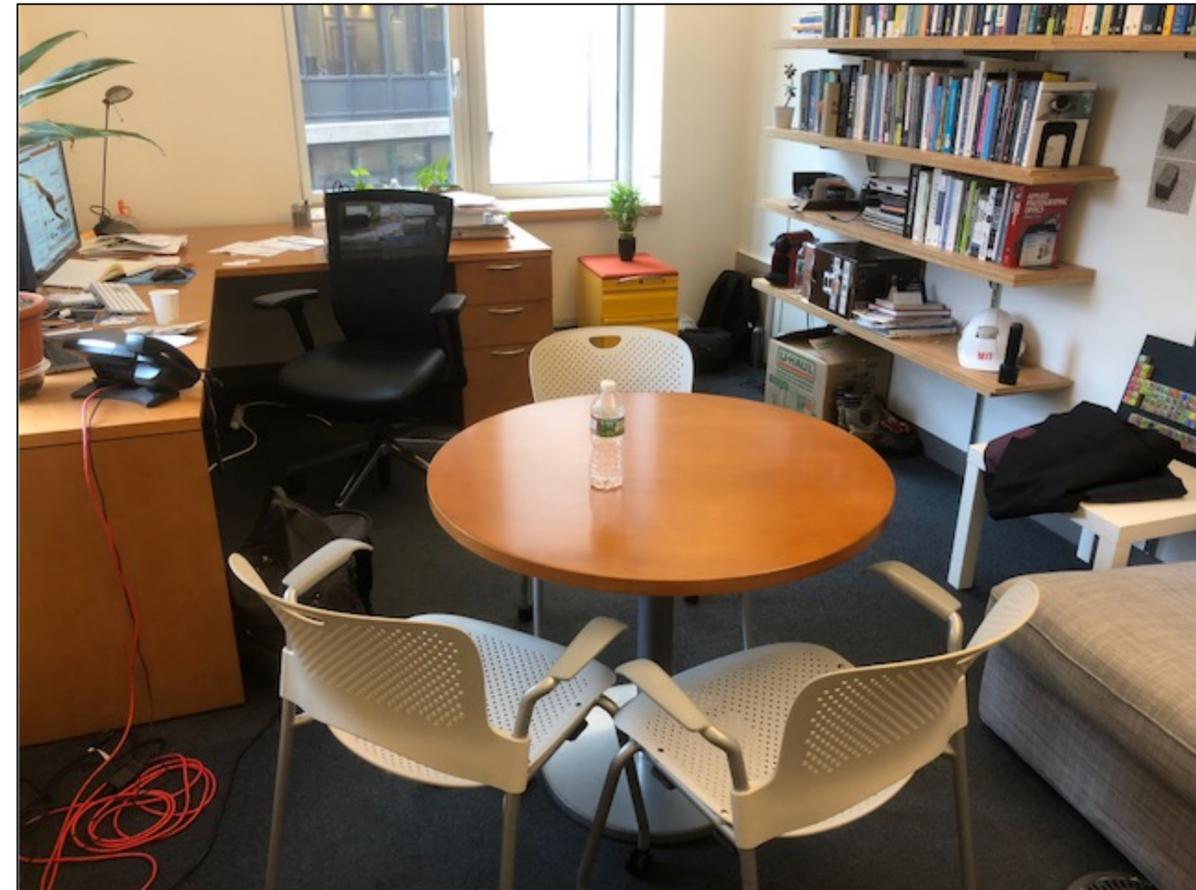
32

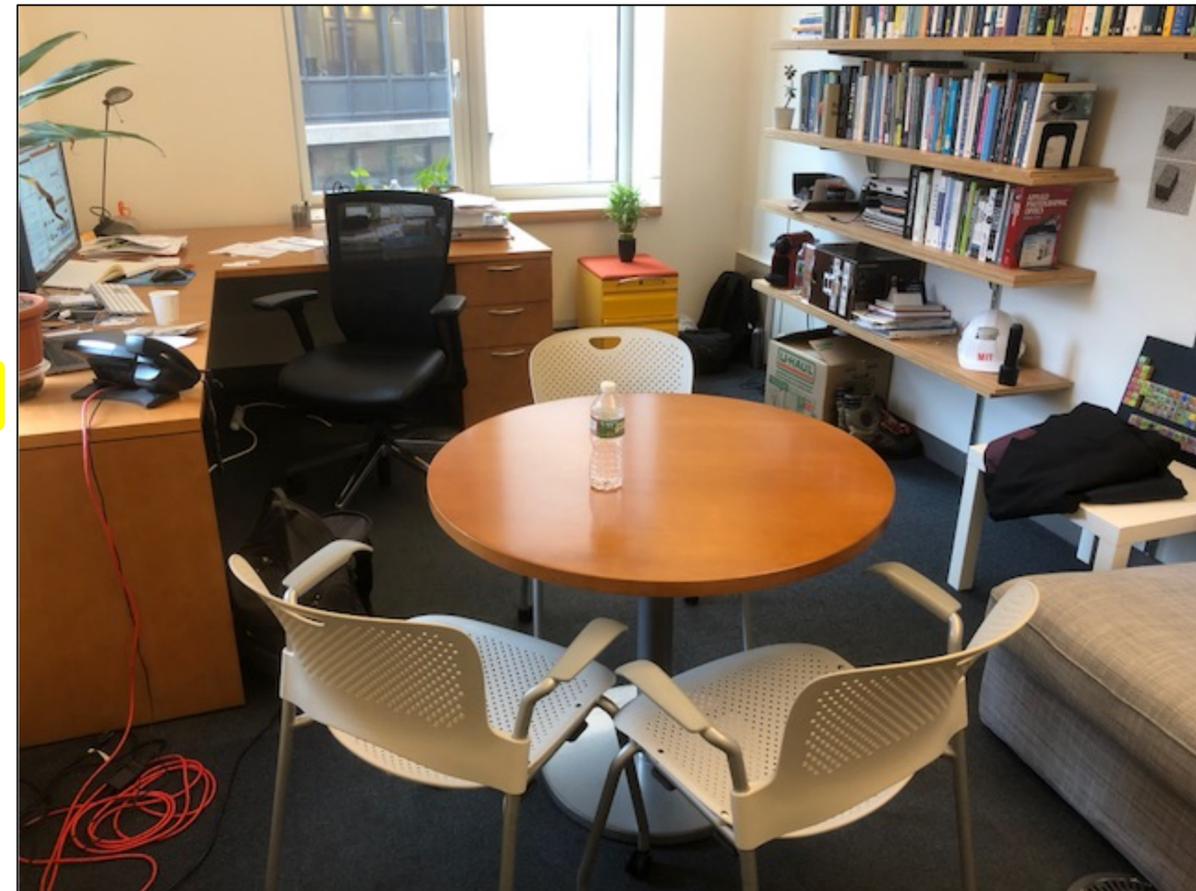
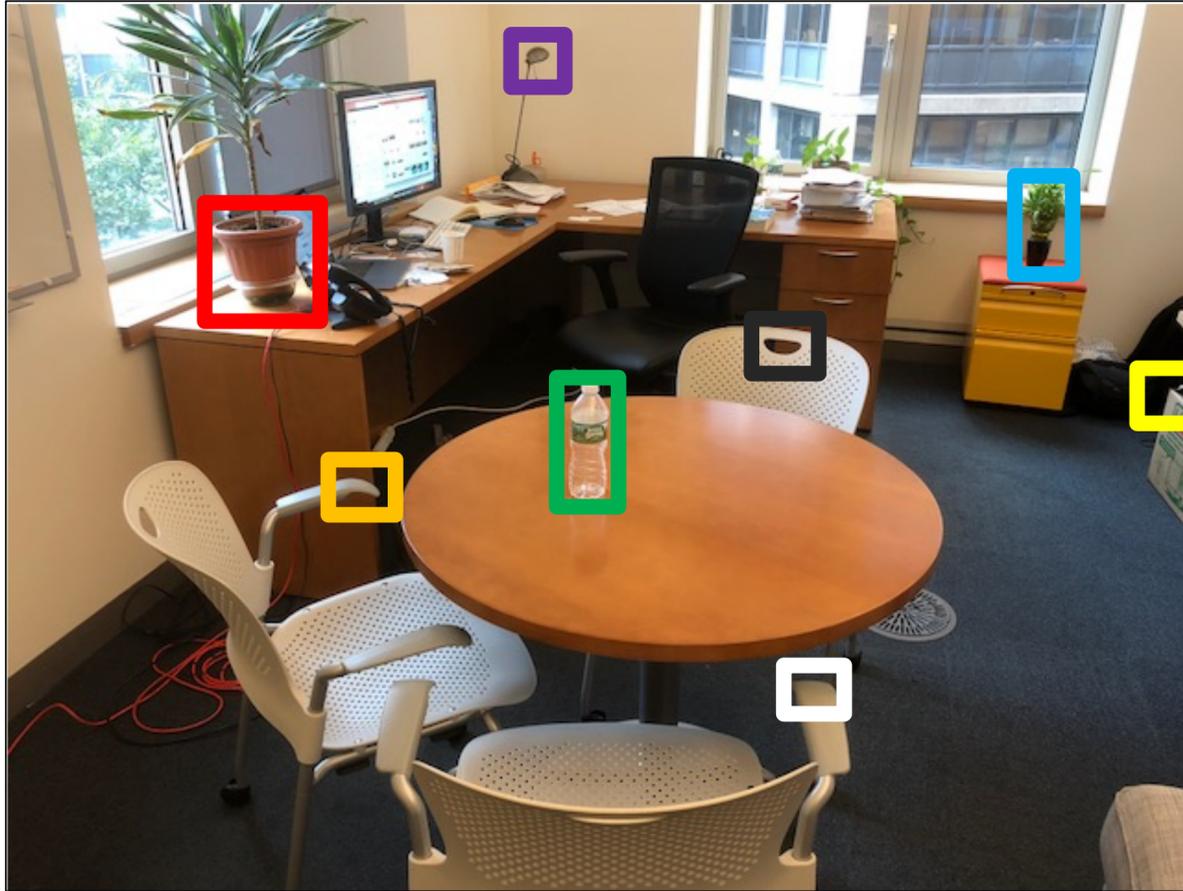
Source: Torralba, Isola, Freeman

# General case

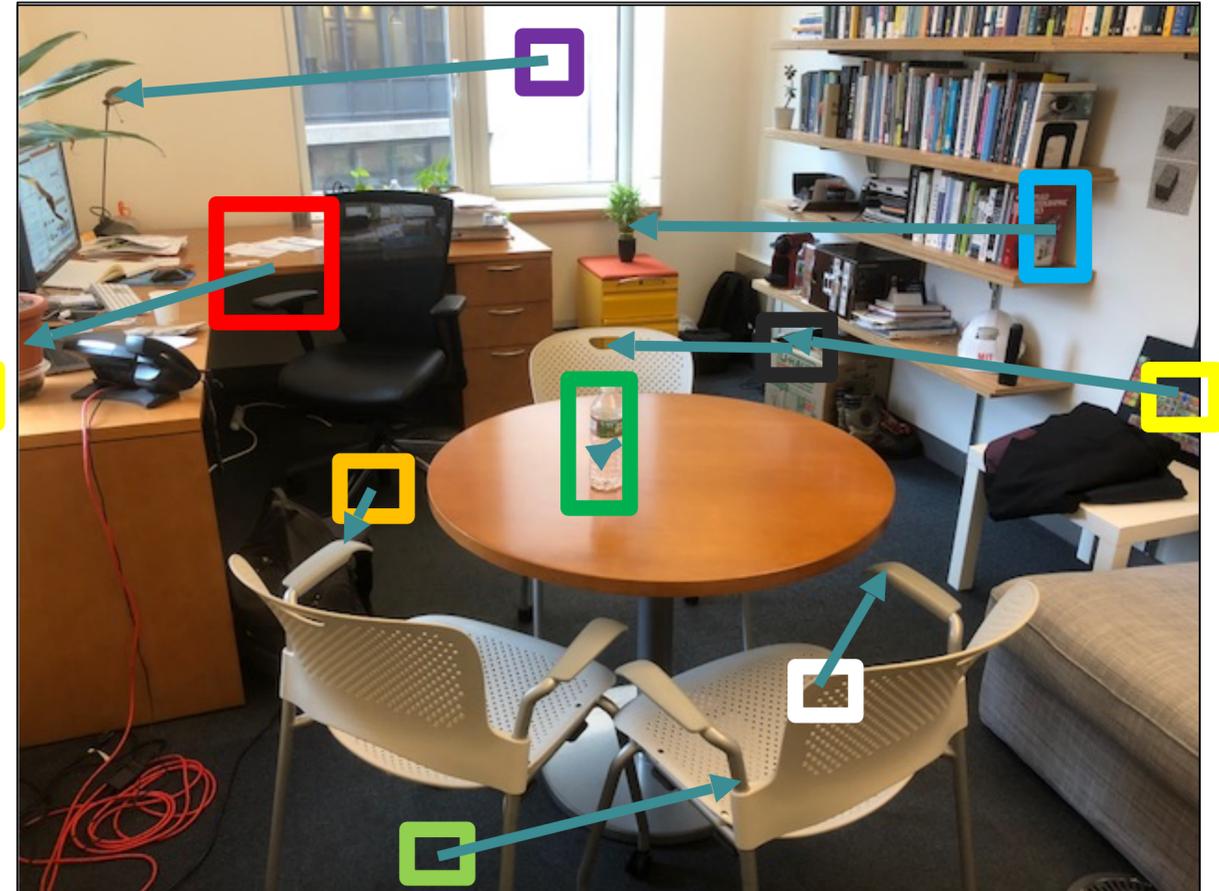
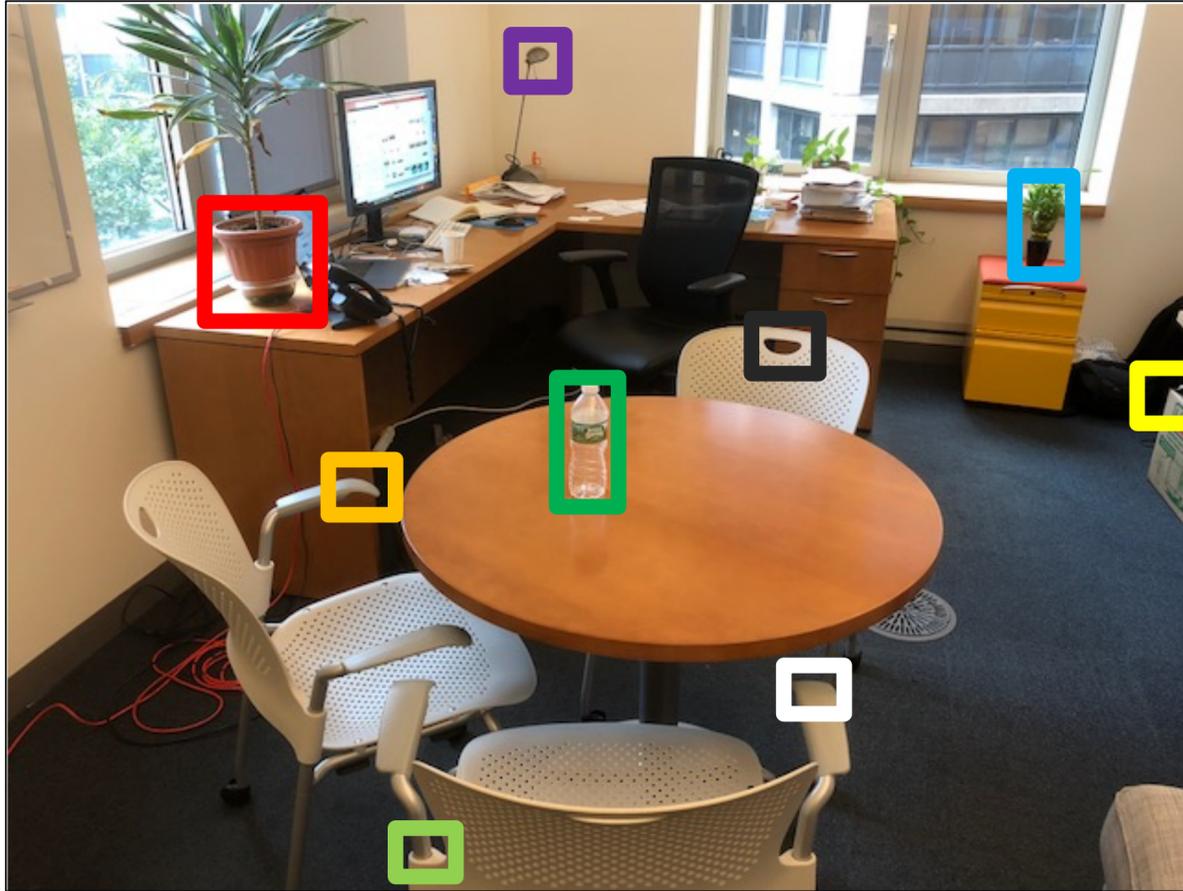
- The two cameras need not have parallel optical axes.



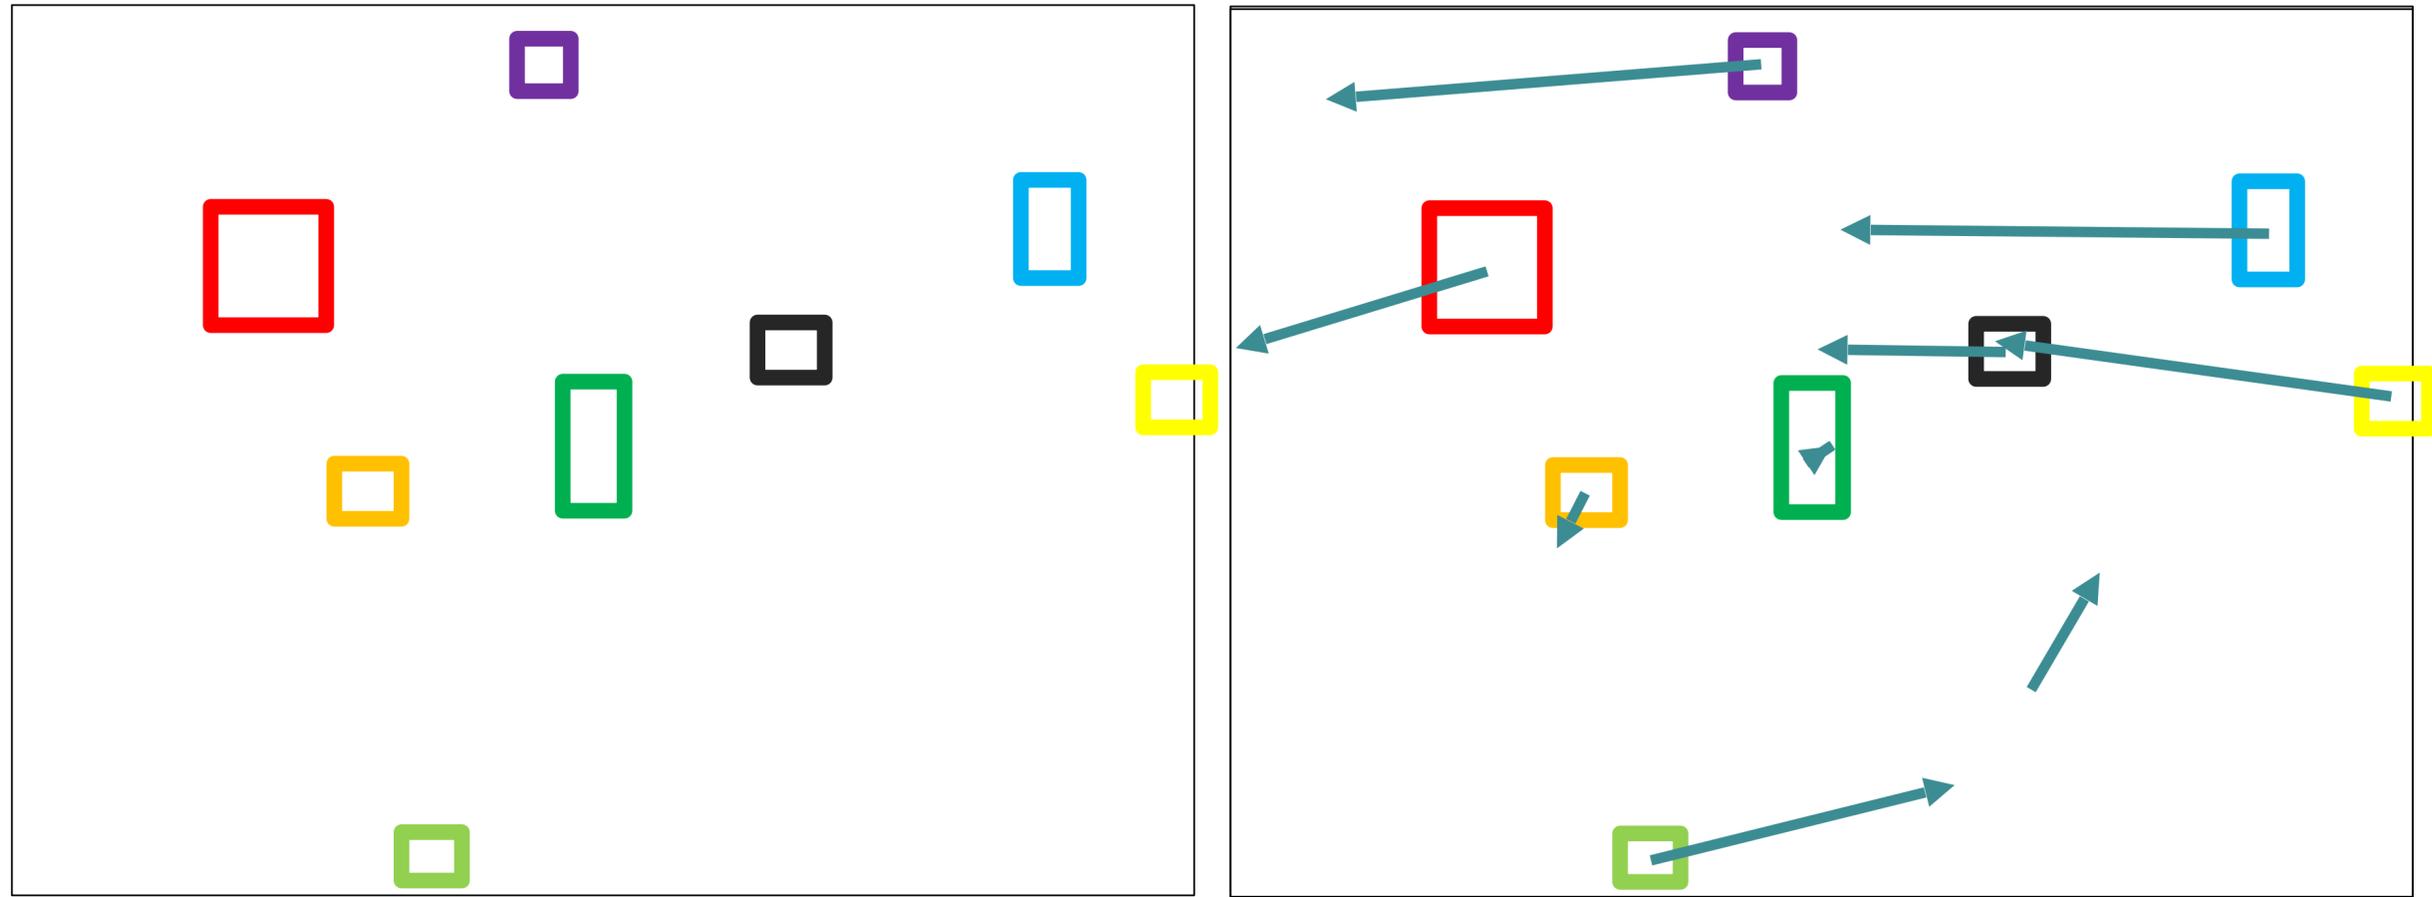




Do we need to search for matches only along horizontal lines?



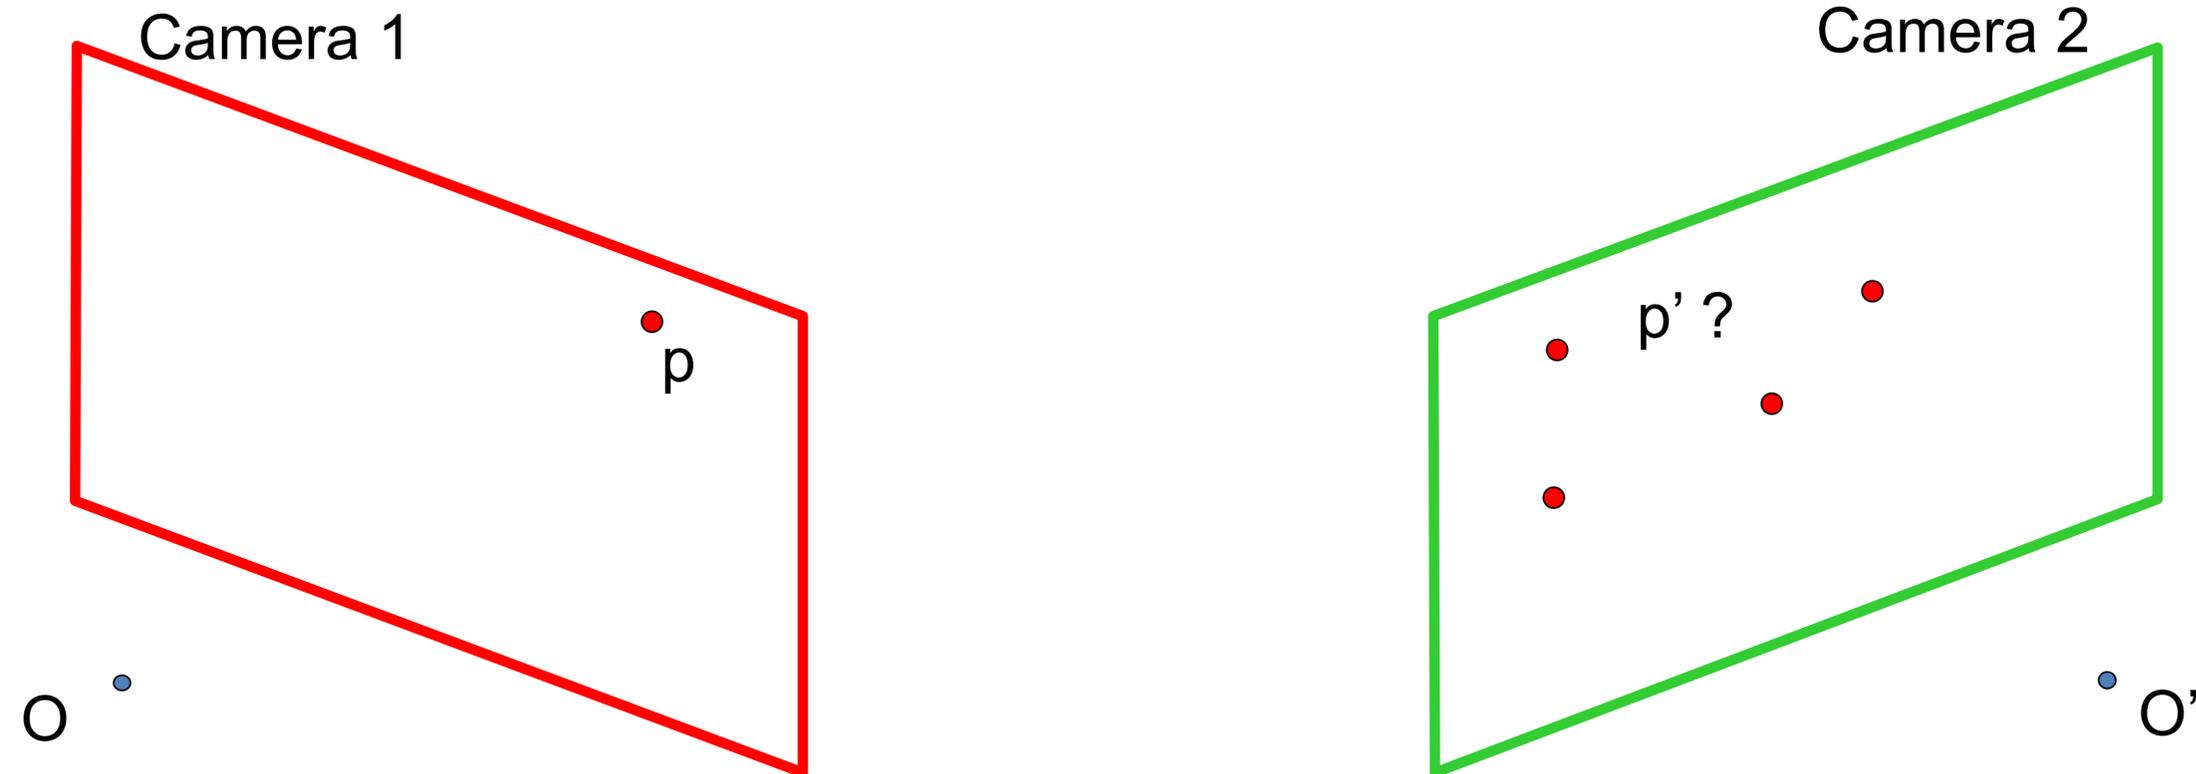
Do we need to search for matches only along horizontal lines?



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

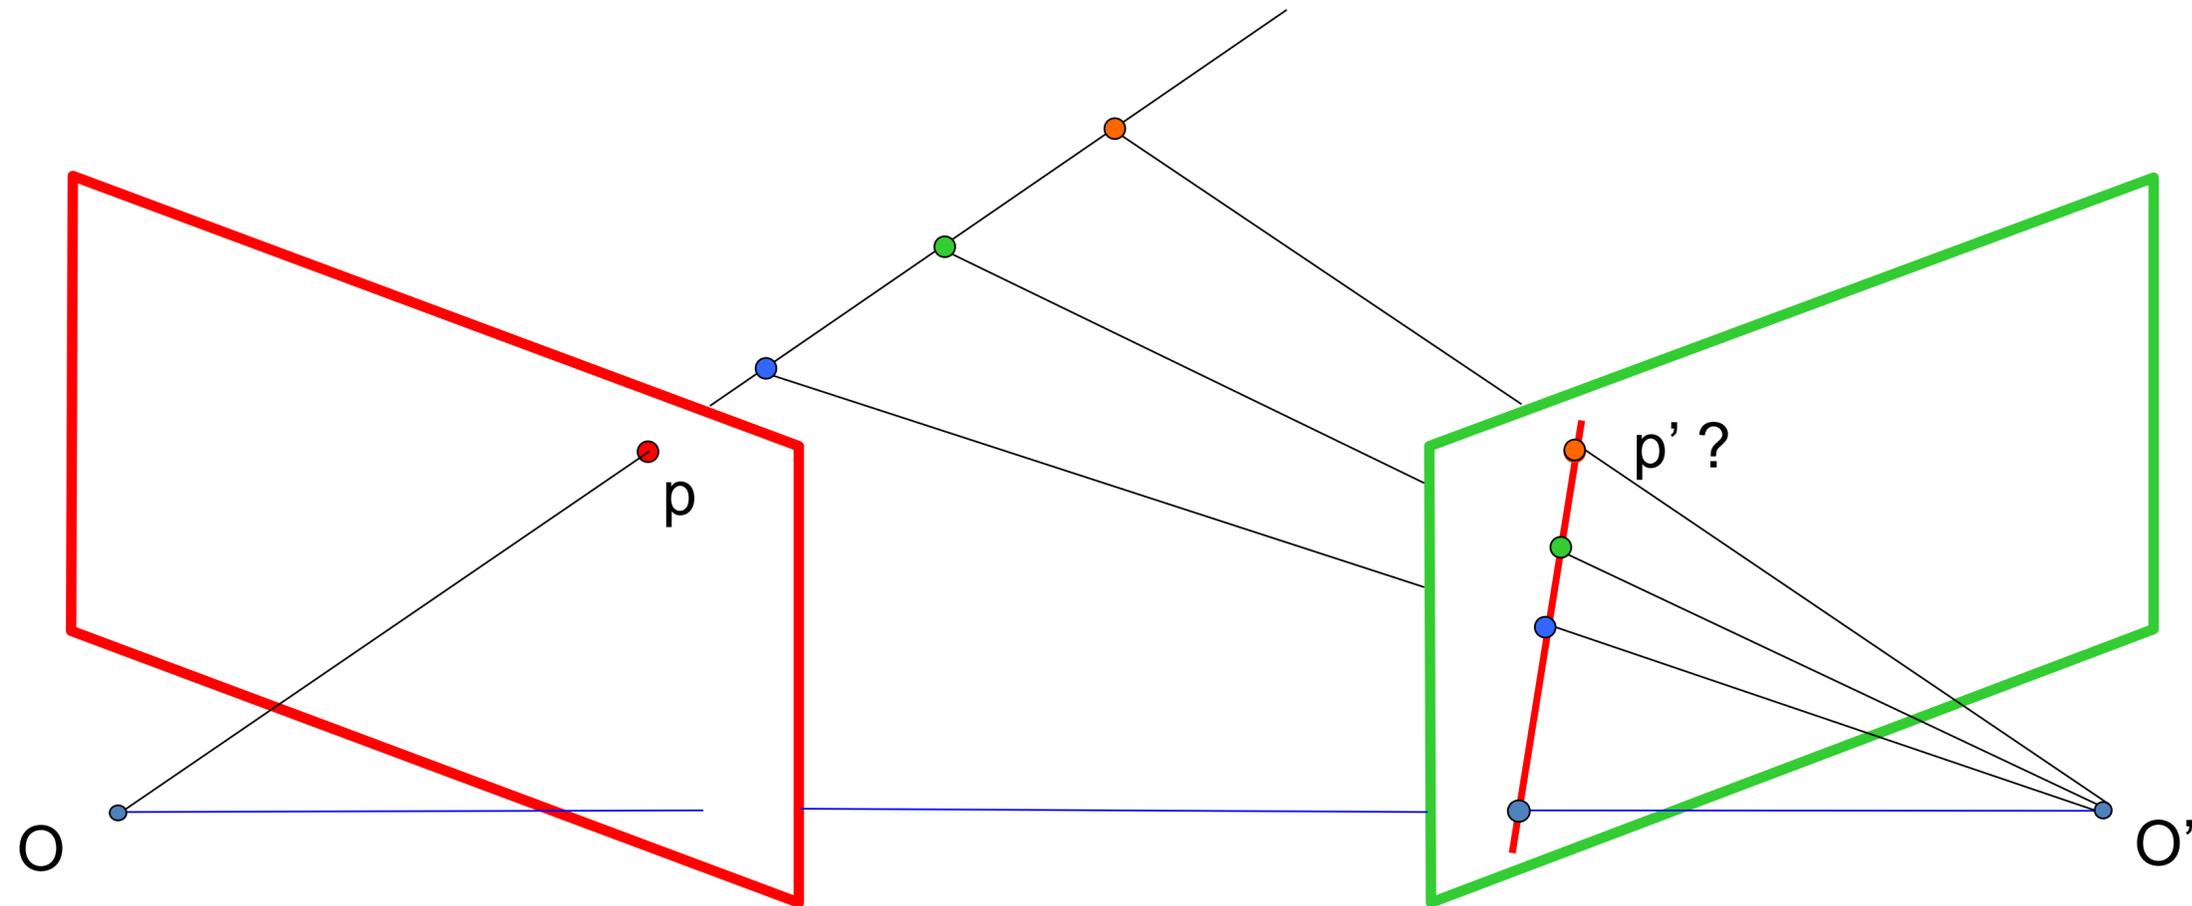
It looks like we might need to search everywhere... are there any constraints that can guide the search?

# Stereo correspondence constraints

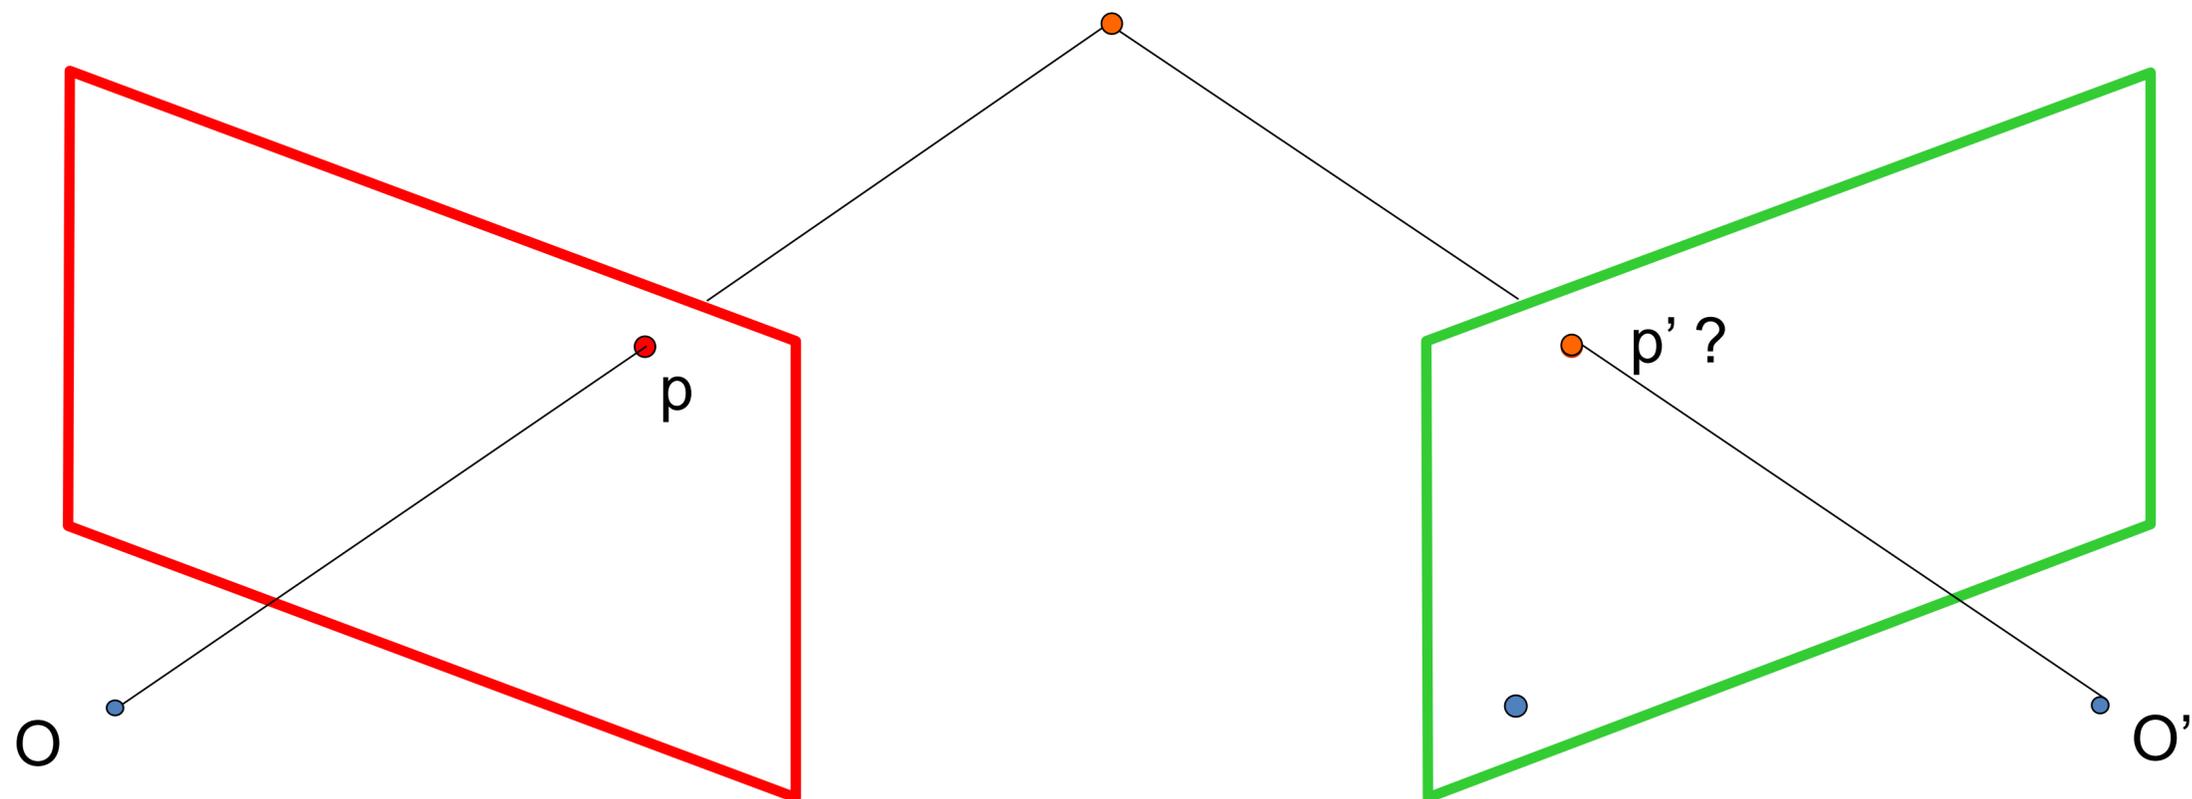


If we see a point in camera 1, are there any constraints on where we will find it on camera 2?

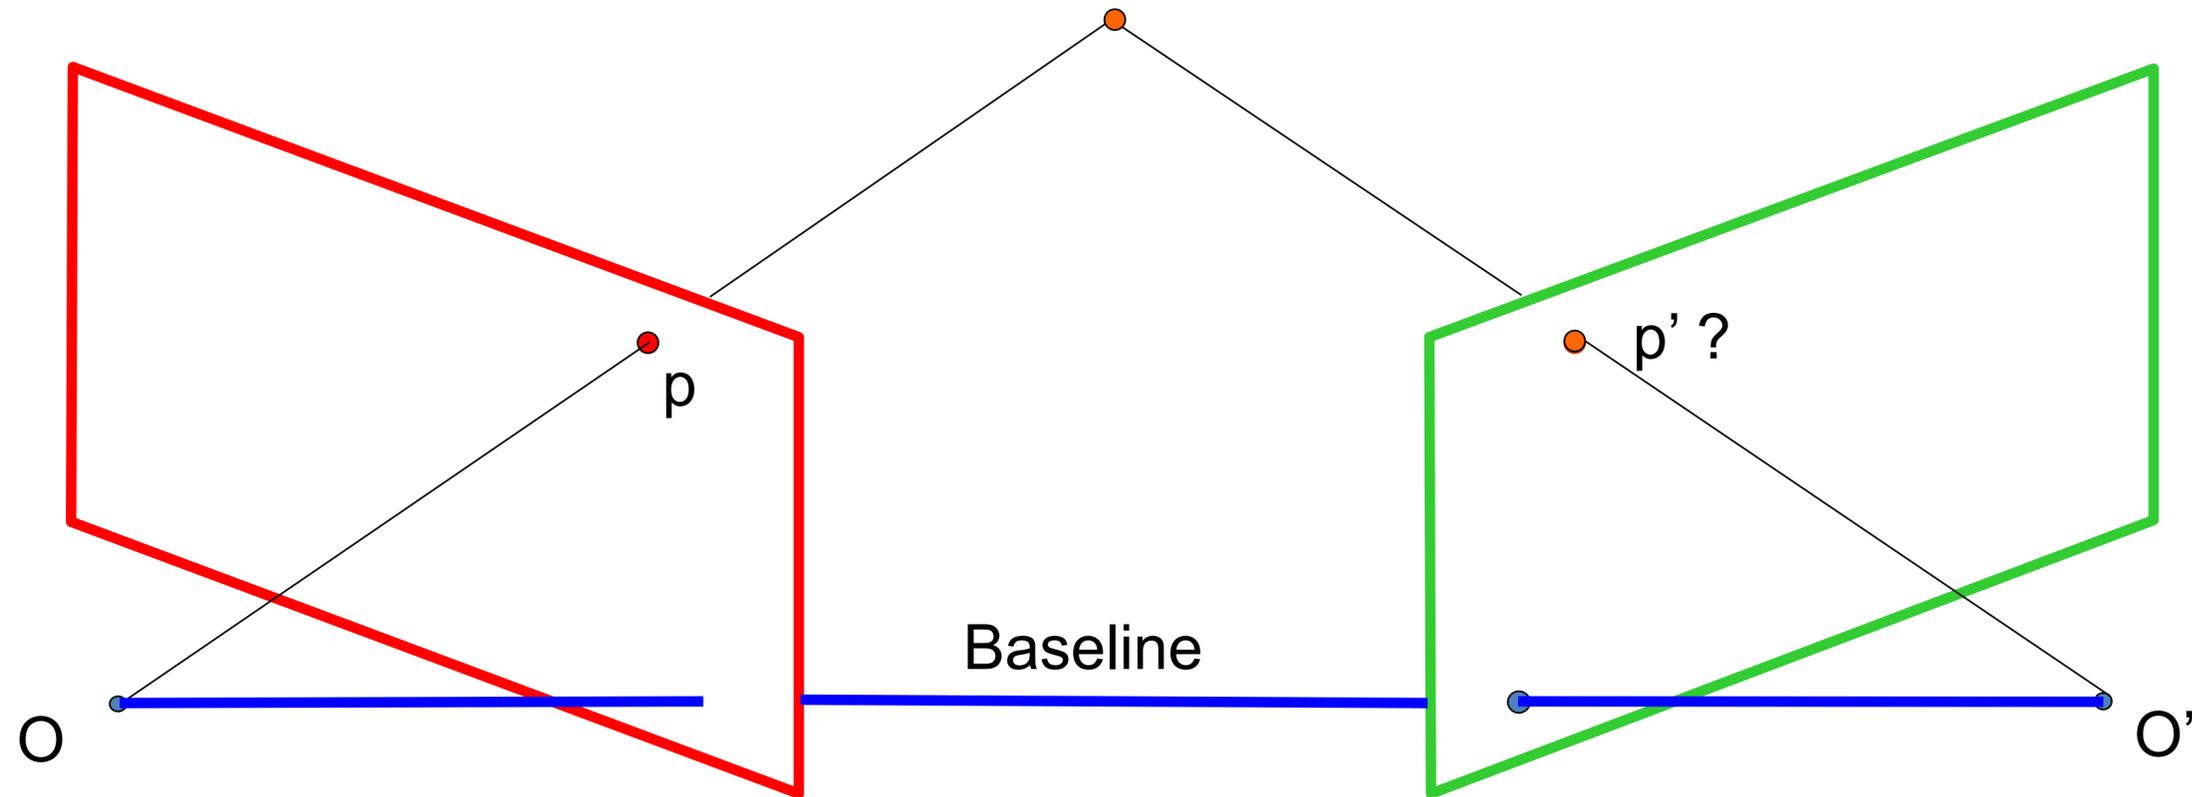
# Stereo correspondence constraints



# Some terminology



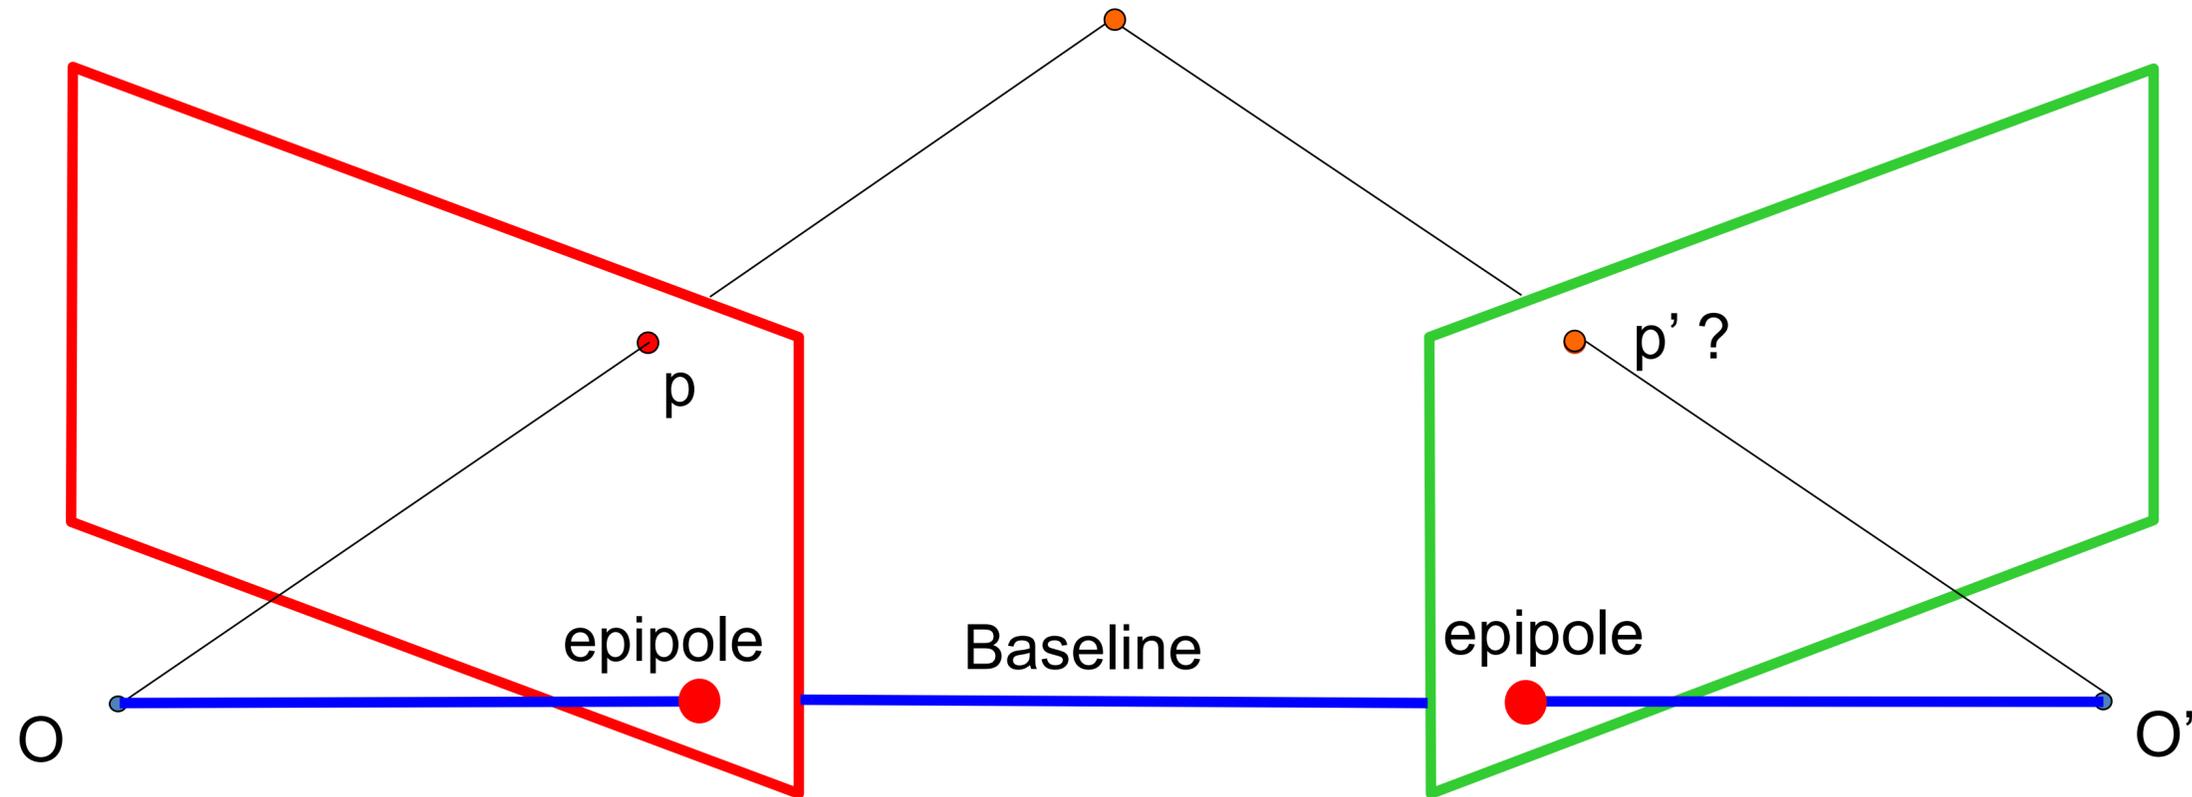
# Some terminology



**Baseline:** the line connecting the two camera centers

**Epipole:** point of intersection of *baseline* with the image plane

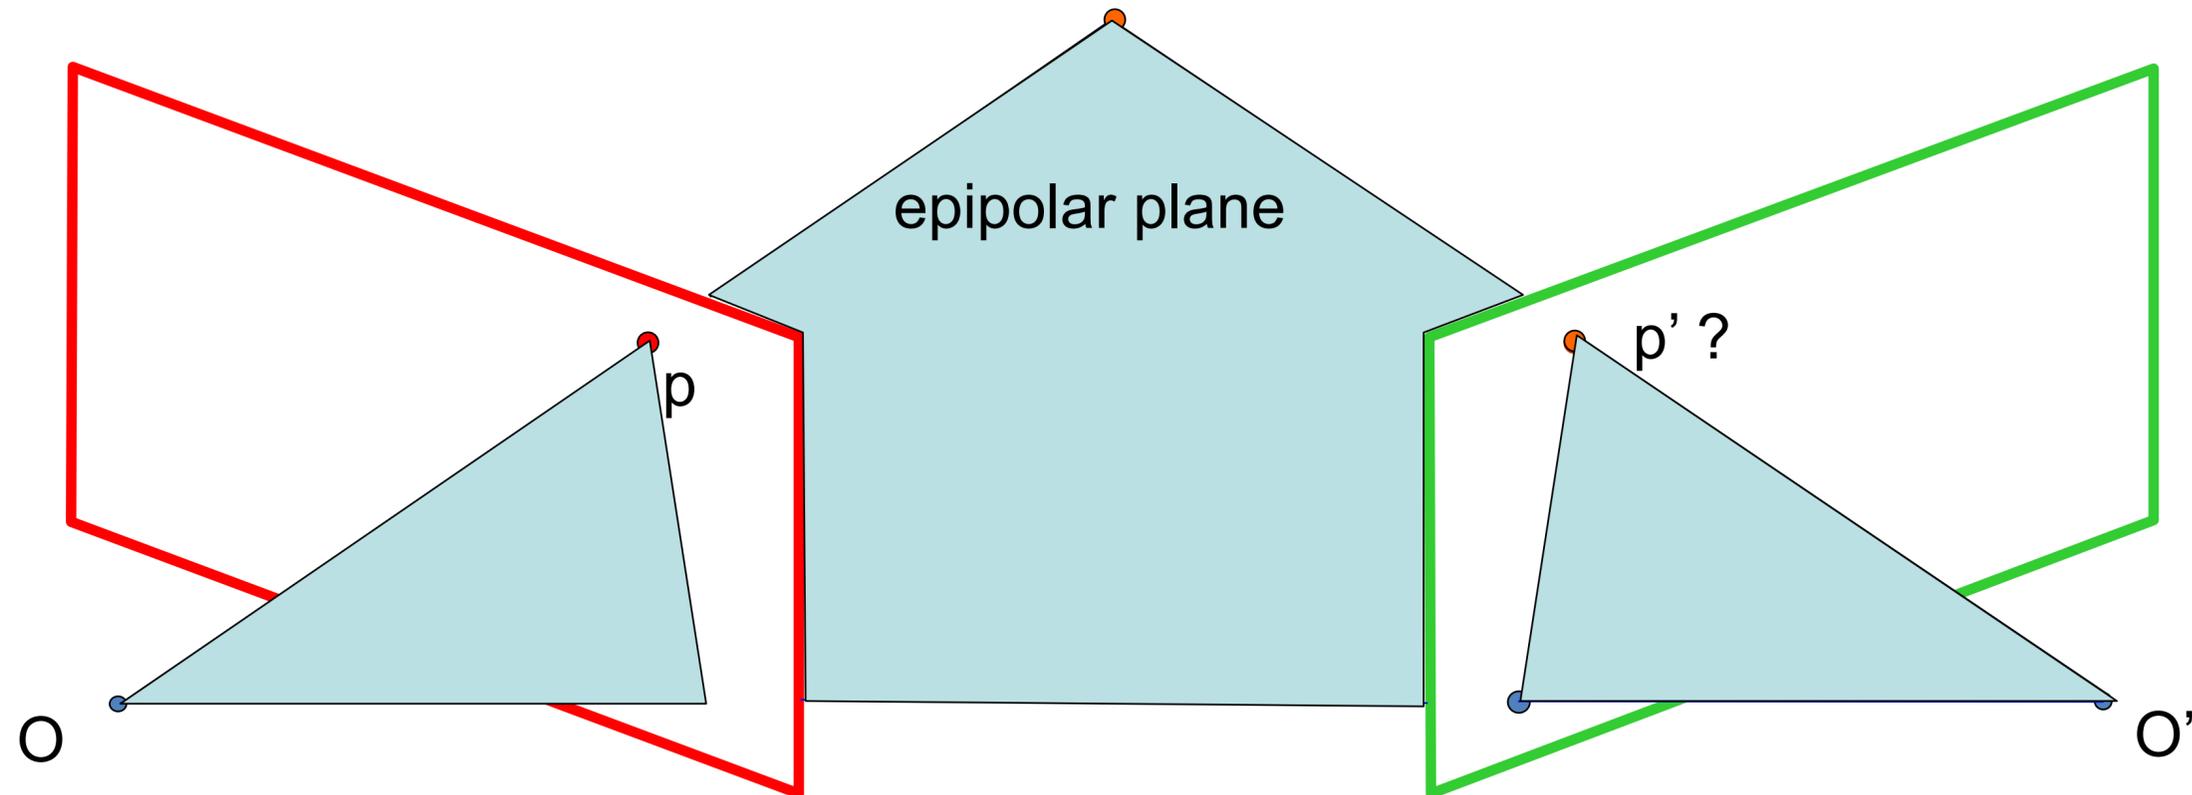
# Some terminology



**Baseline:** the line connecting the two camera centers

**Epipole:** point of intersection of *baseline* with the image plane

# Some terminology

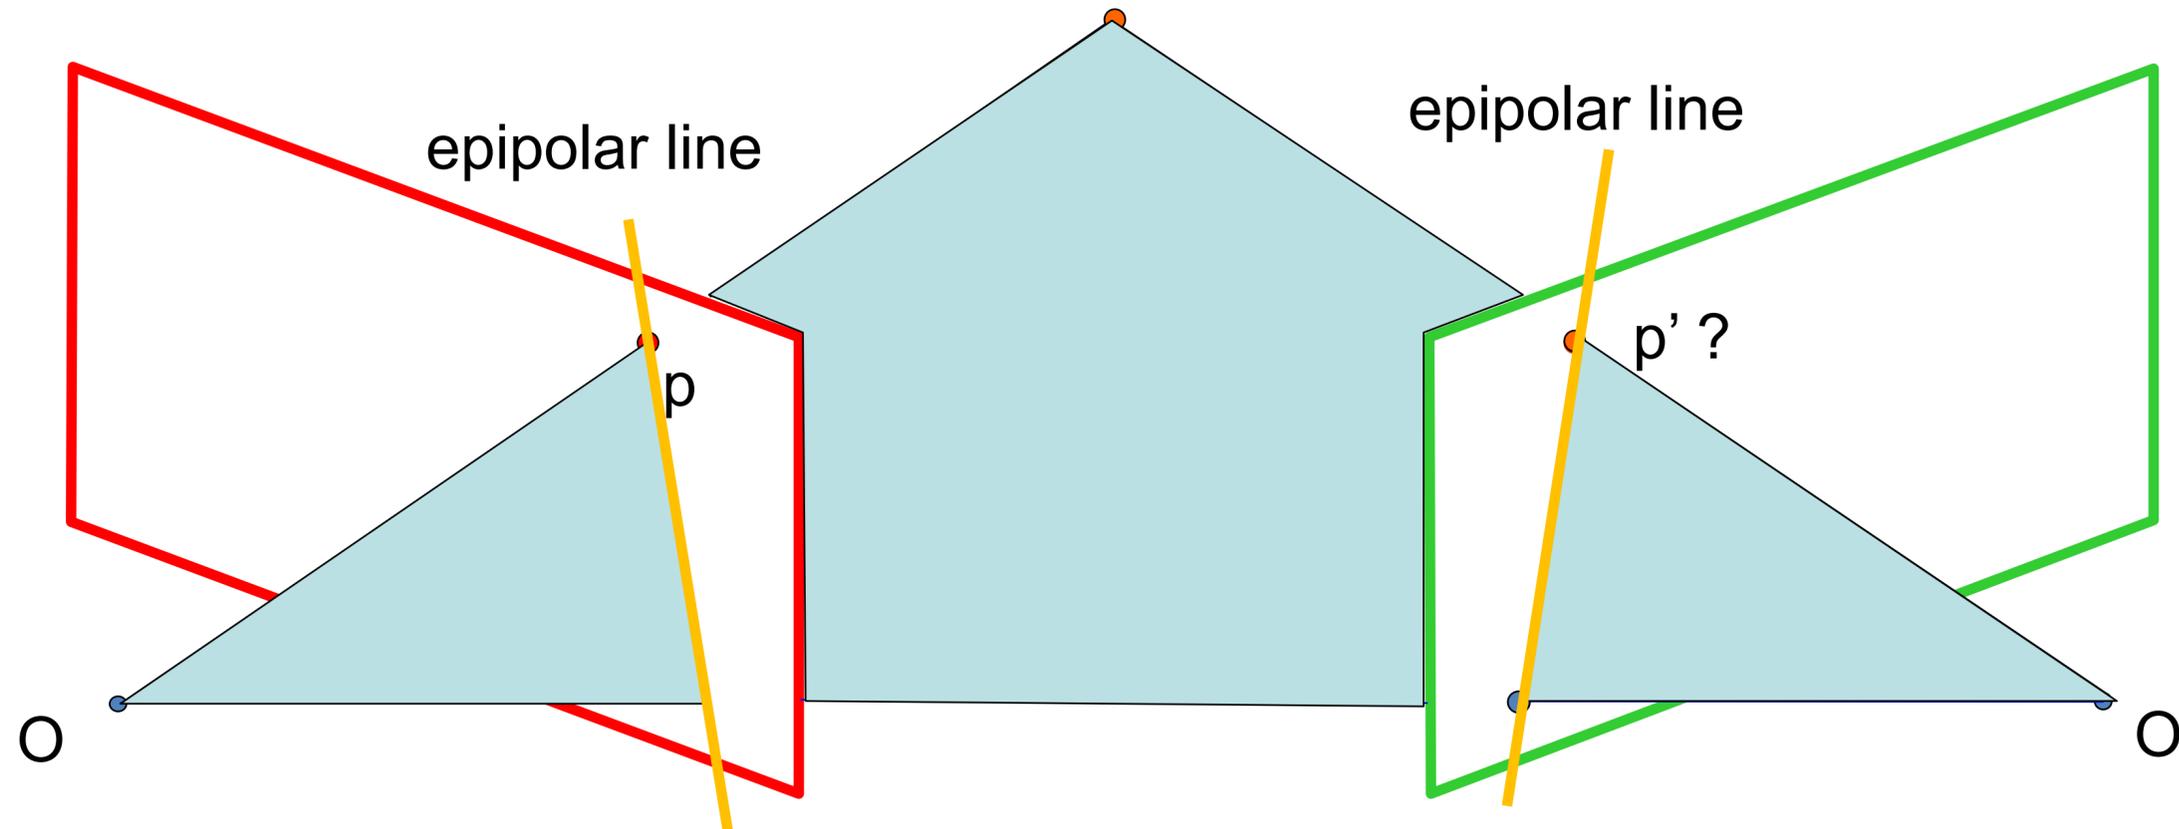


**Baseline:** the line connecting the two camera centers

**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

# Some terminology



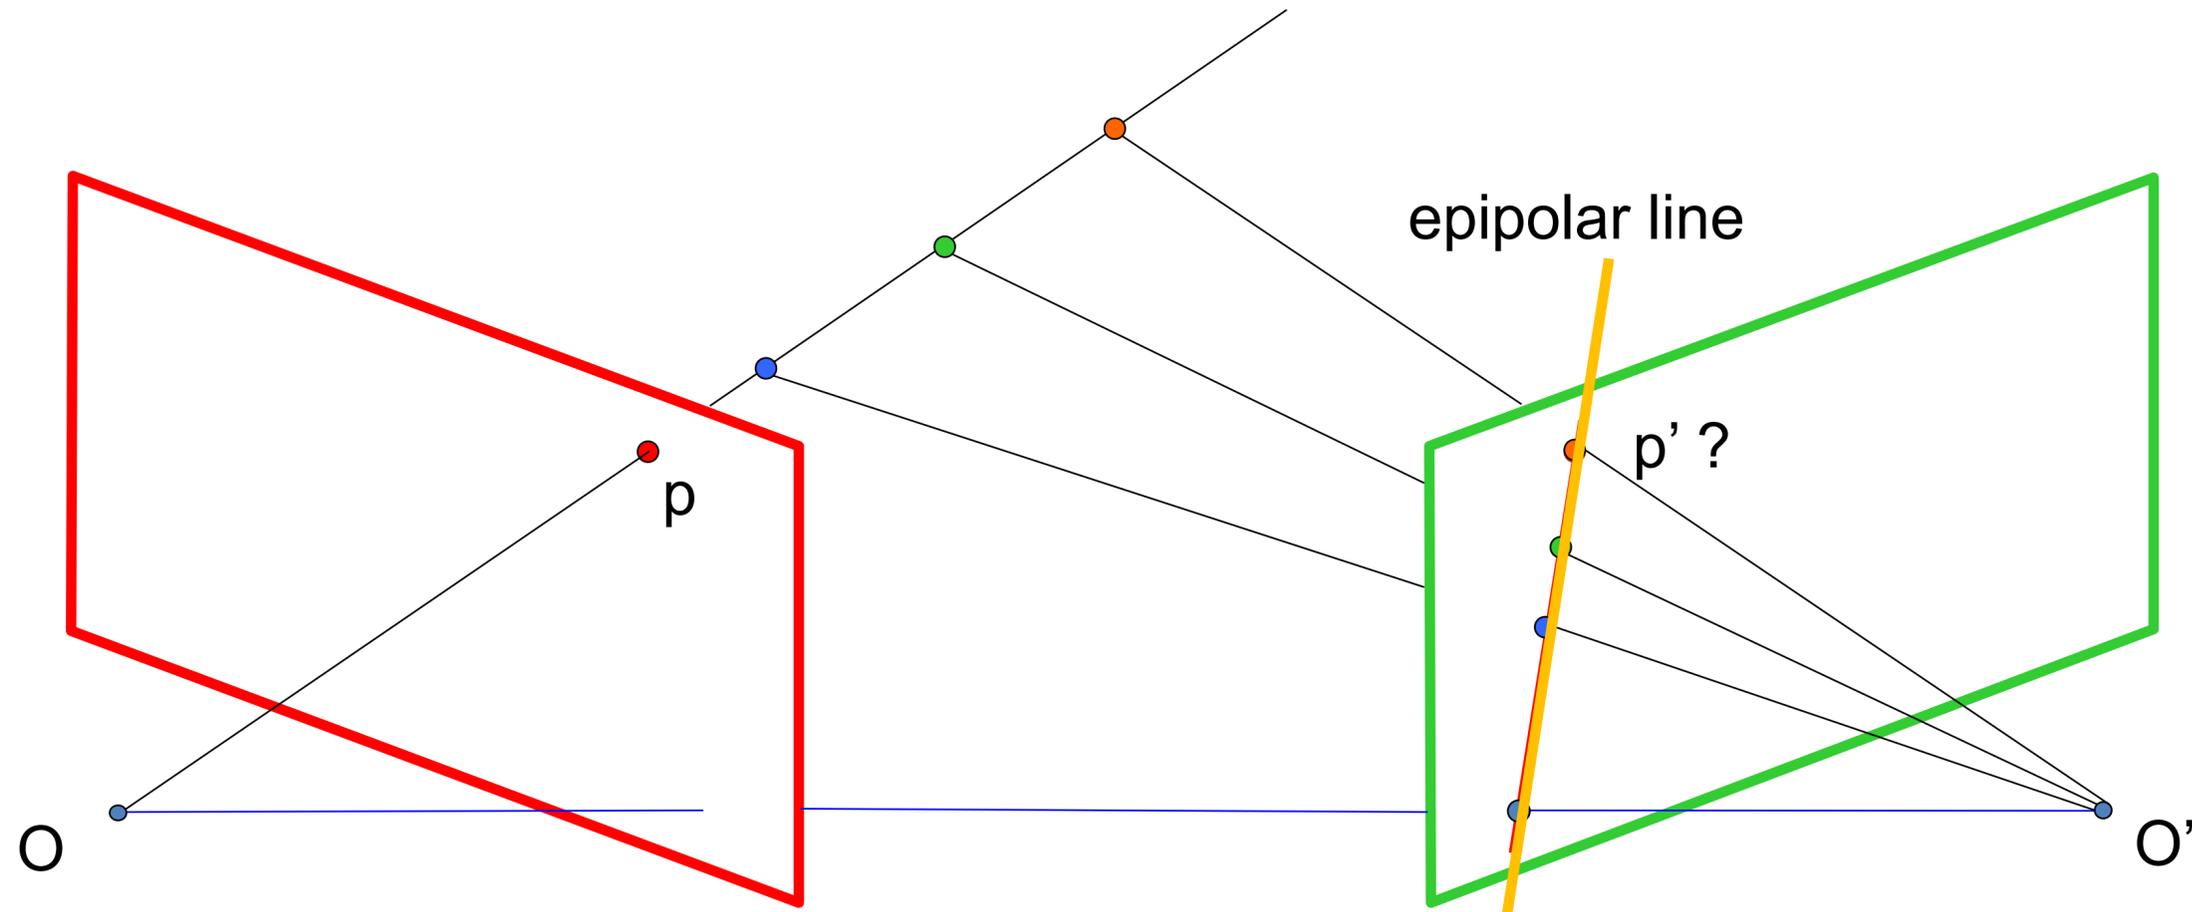
**Baseline:** the line connecting the two camera centers

**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

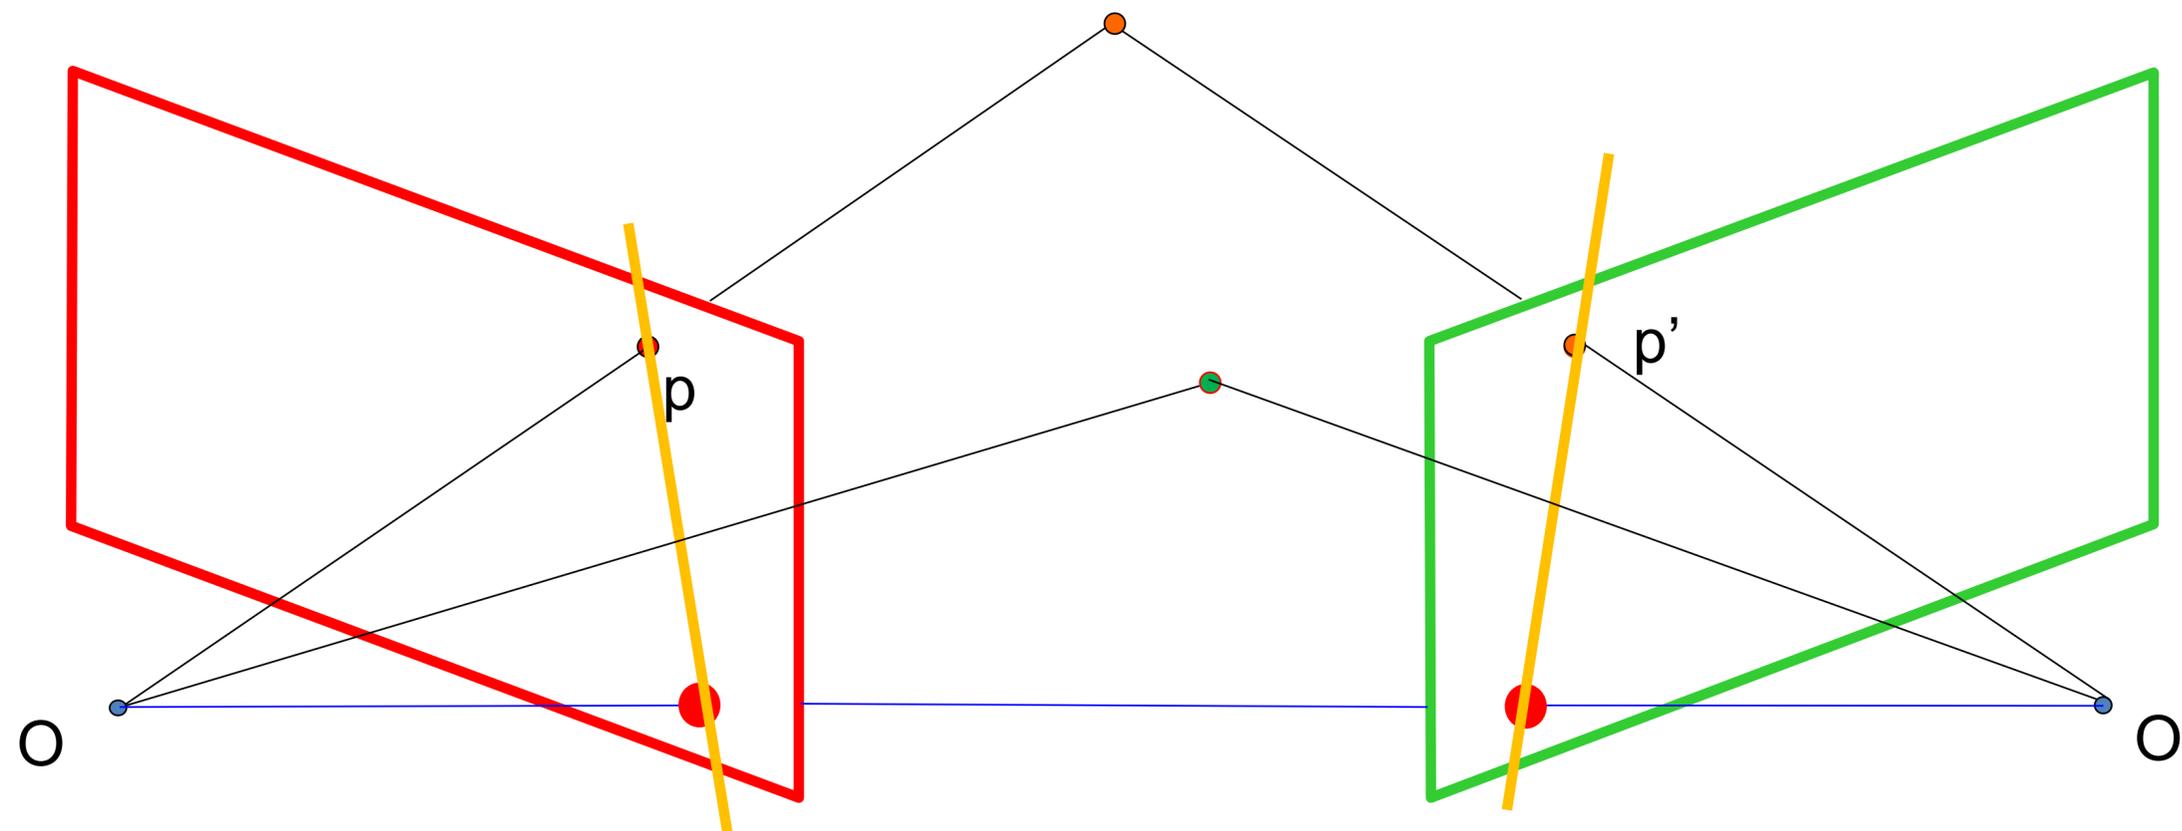
# Epipolar constraint



We can search for matches across epipolar lines

All epipolar lines intersect at the epipoles

# Epipolar constraint



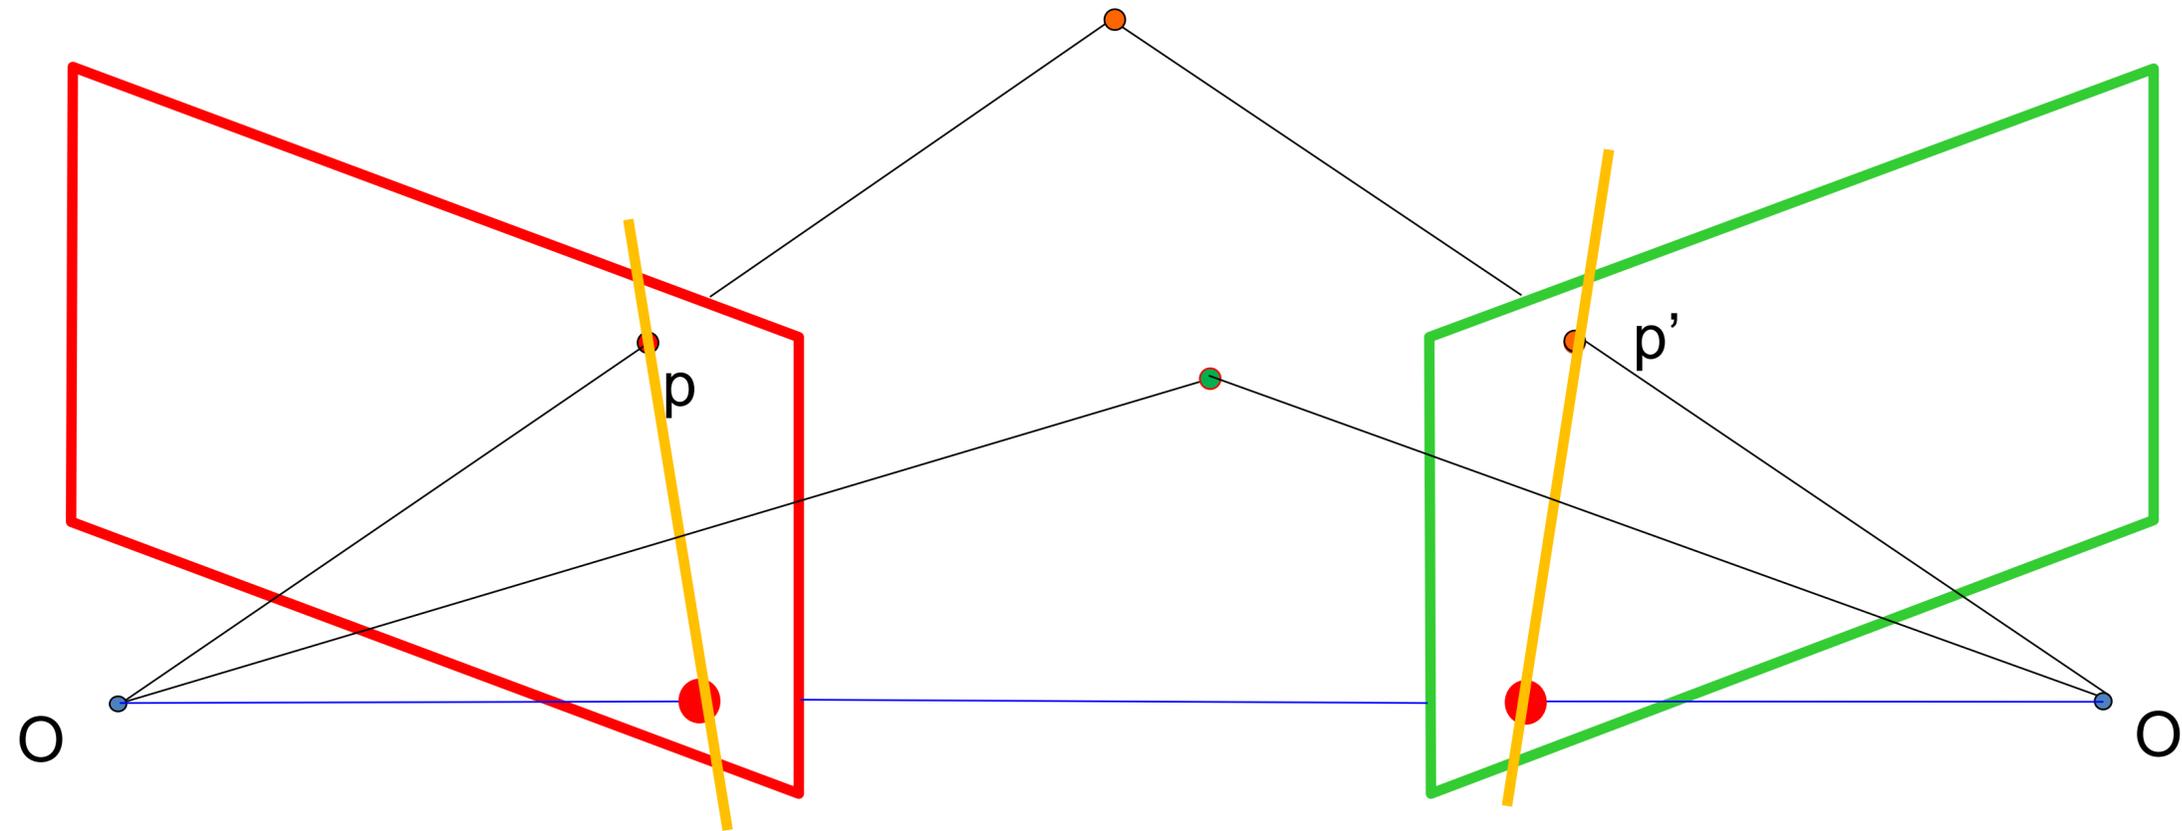
If we observe a point in one image, its position in the other image is constrained to lie on the epipolar line.

How do we get this line? We want a function that, given a point  $\mathbf{p}$  tells us what this line is:

$$f(\mathbf{p}) = [a, b, c] \quad \text{such that} \quad ax' + by' + c = 0$$

where  $\mathbf{p}' = [x', y']$ . In other words:  $f(\mathbf{p})^\top \mathbf{p}' = 0$

# The fundamental matrix



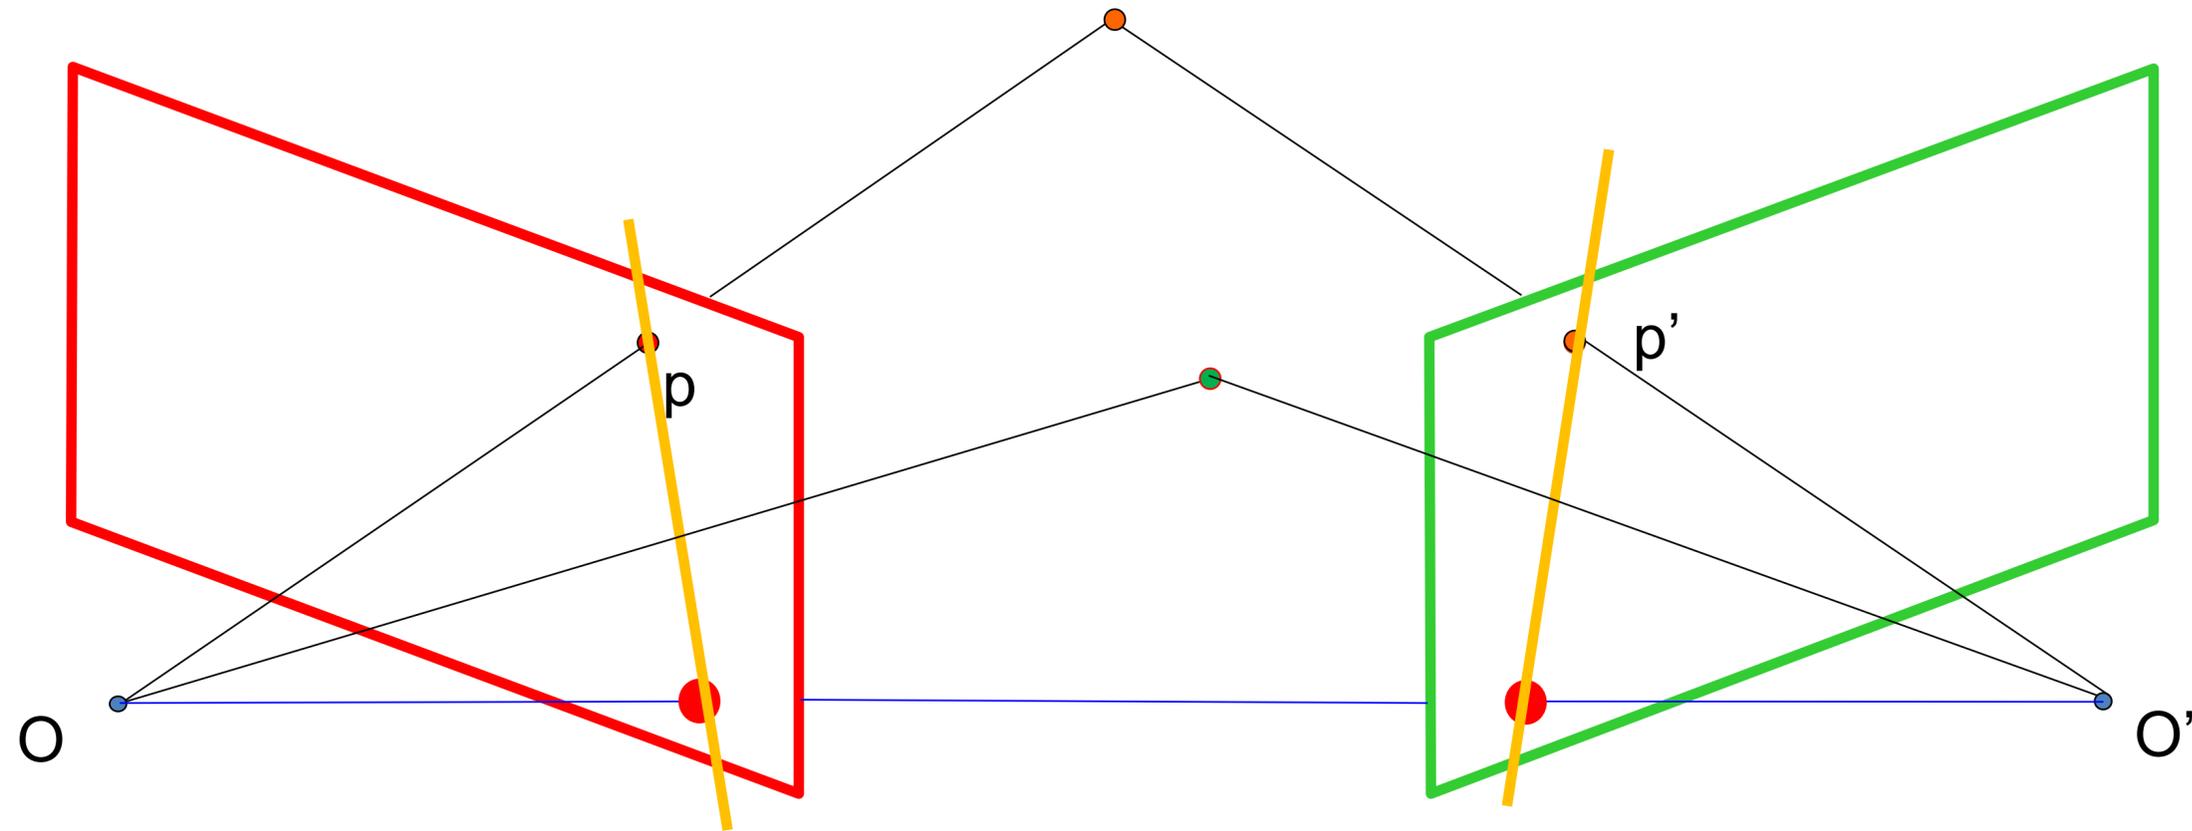
It can be shown that our function,  $f$ , can be written as a **matrix multiplication in homogeneous coordinates**.

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

F: fundamental matrix  
p image point in homogeneous coordinates

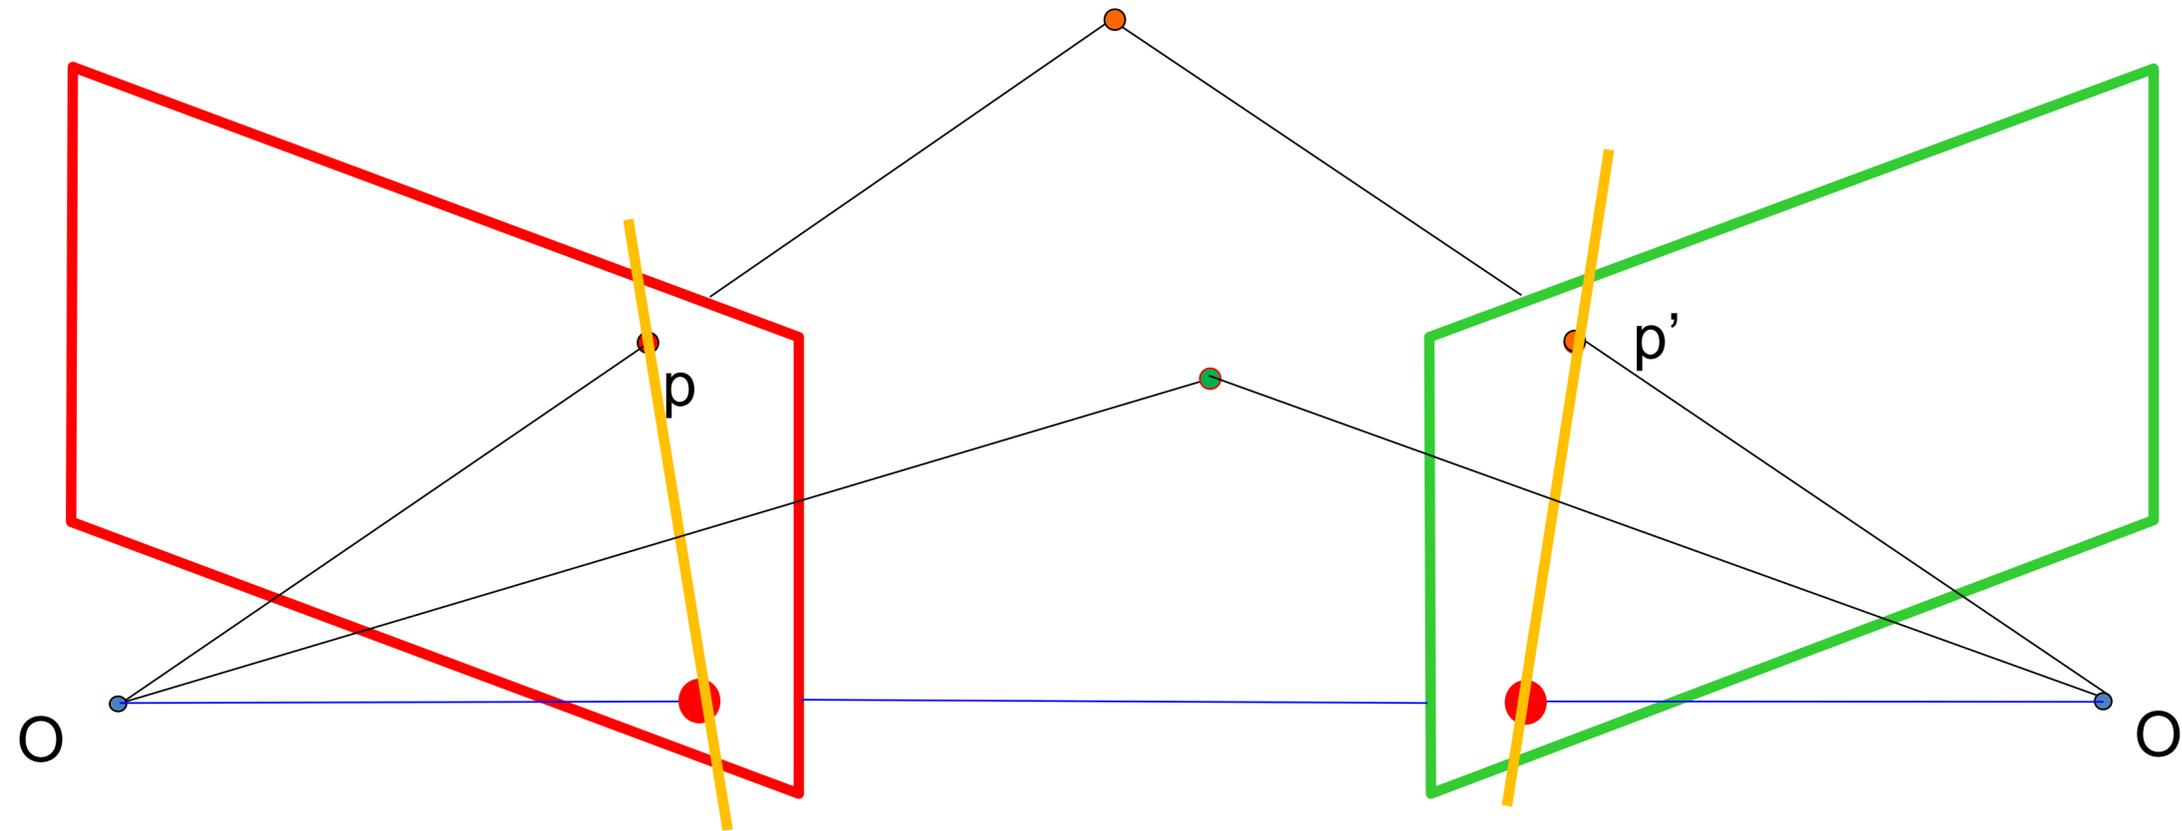


# The fundamental matrix



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

# The fundamental matrix



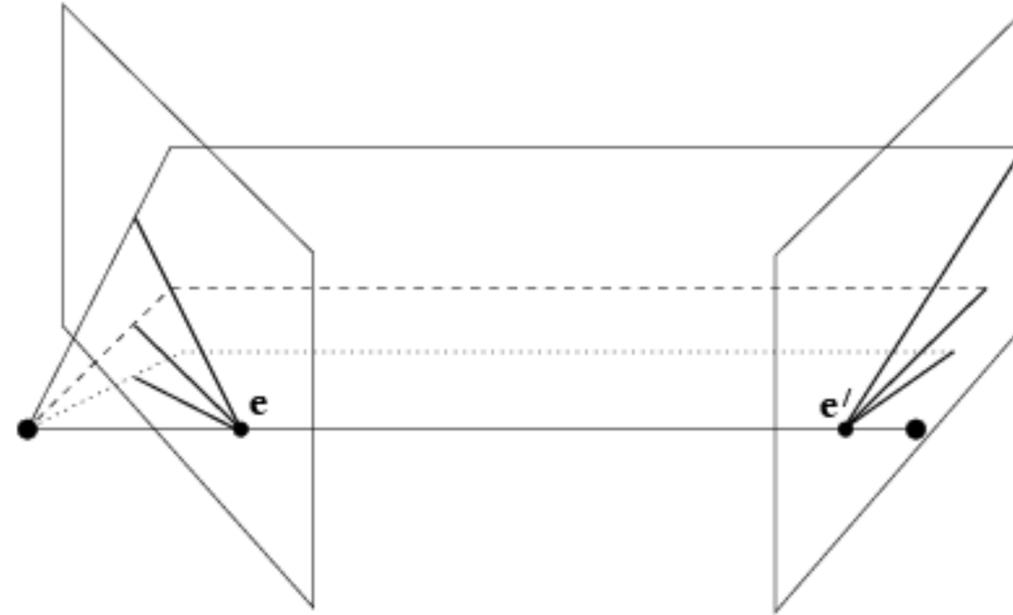
$$(\mathbf{p}^\top F) \mathbf{p}' = 0$$

$$\mathbf{u}^\top \mathbf{p}' = 0$$

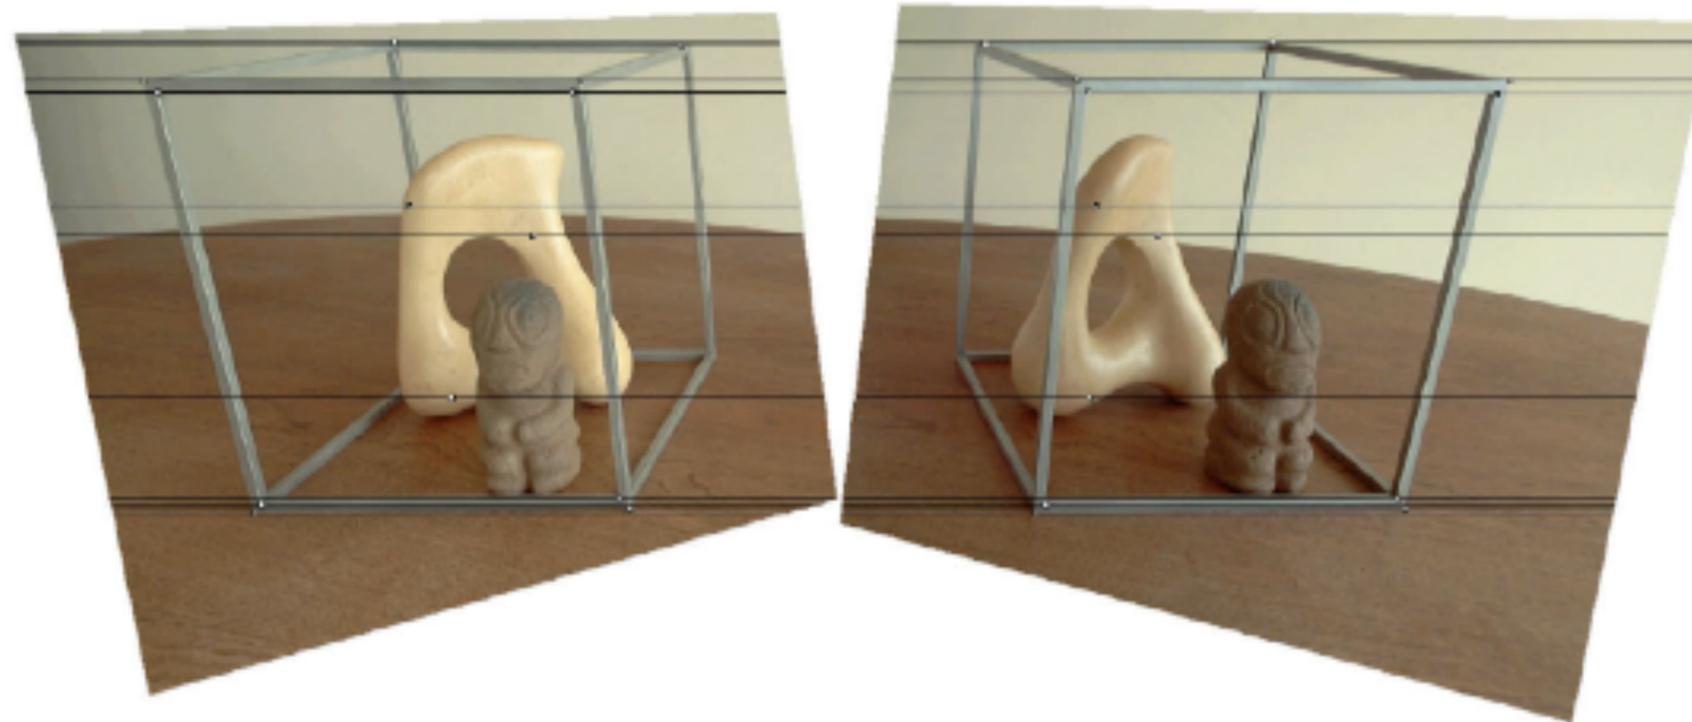
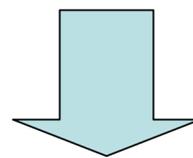
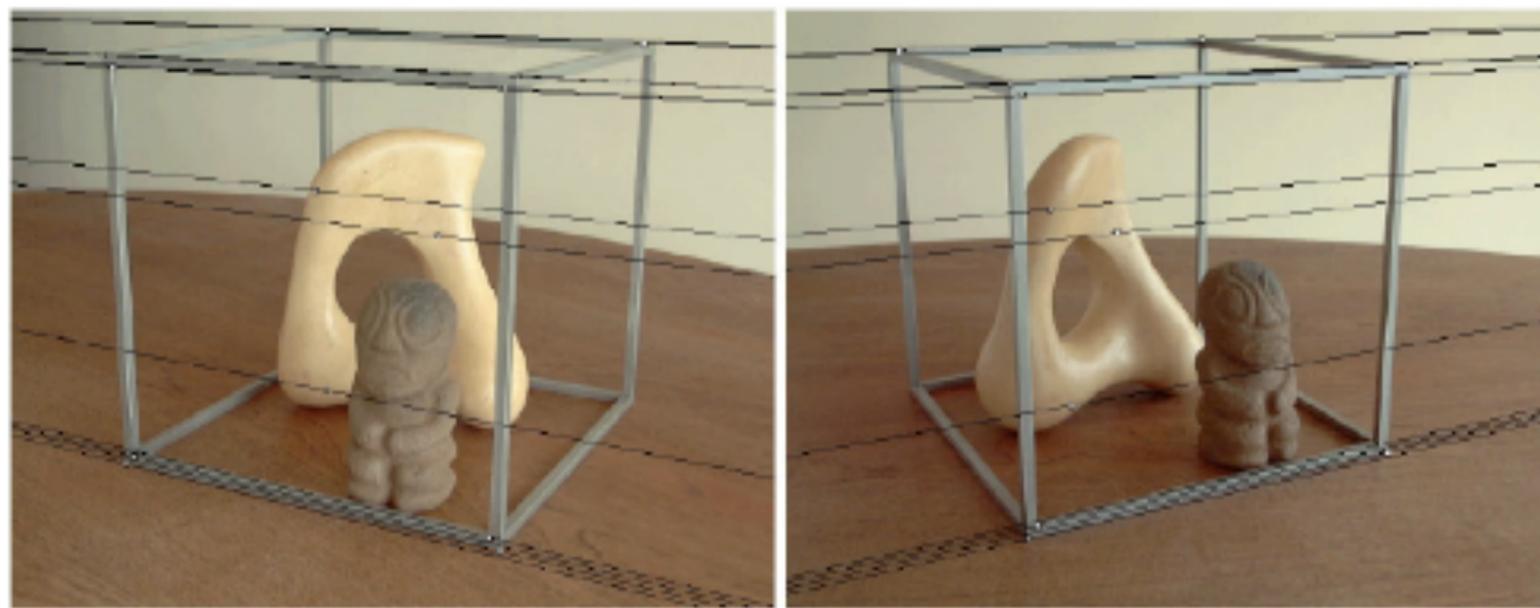
$\mathbf{u}$ : a line induced by  $\mathbf{p}$

$\mathbf{p}$ ,  $\mathbf{p}'$ : image points in homogeneous coordinates

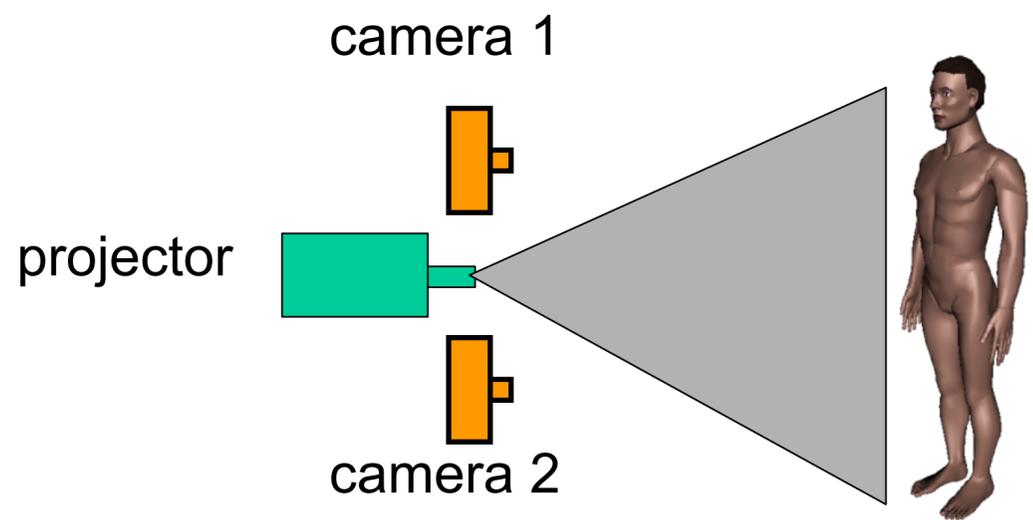
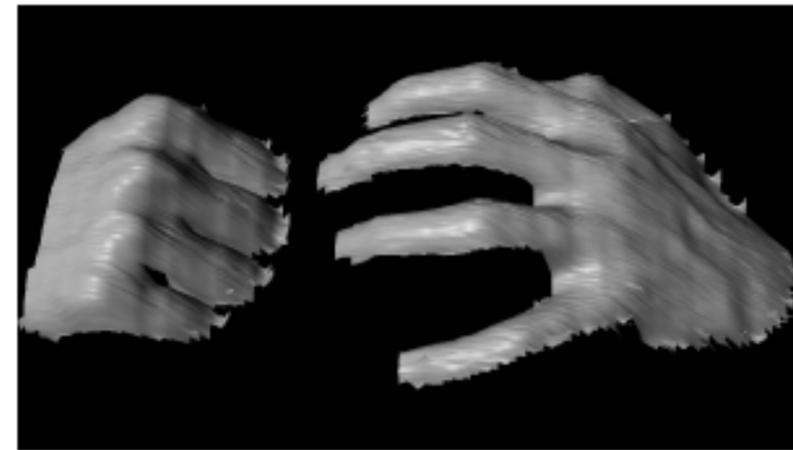
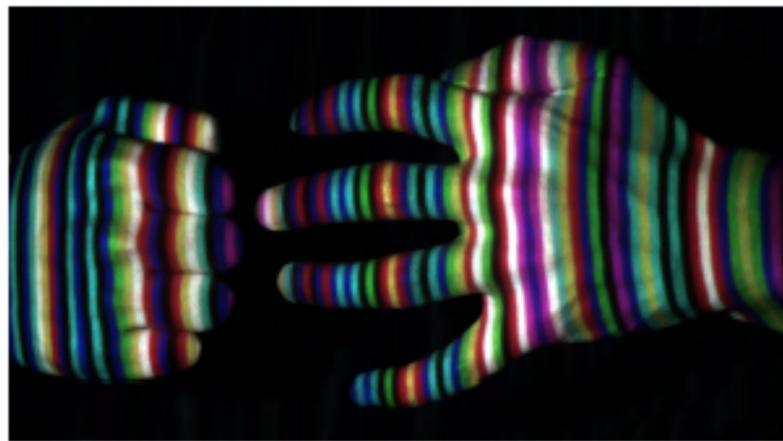
# Example: converging cameras



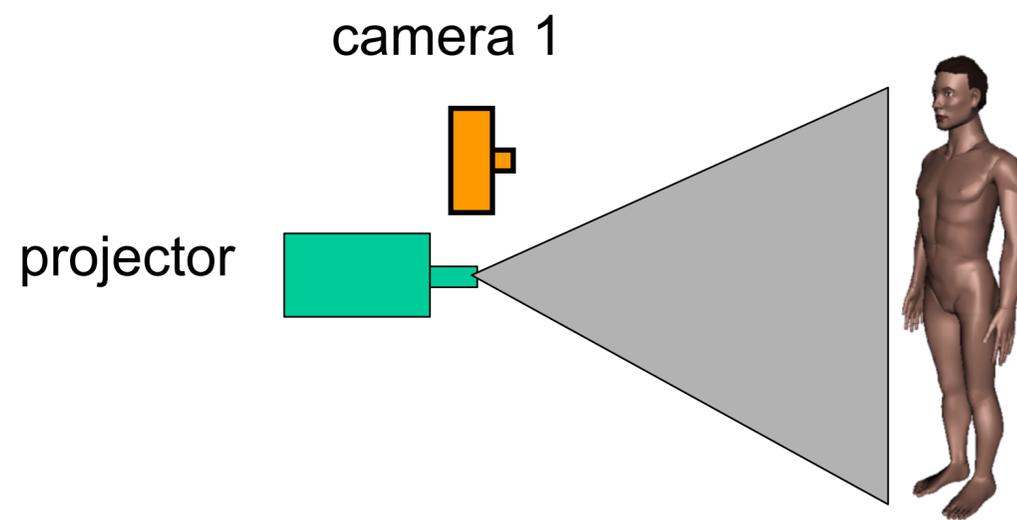
# Image rectification



# Active stereo with structured light

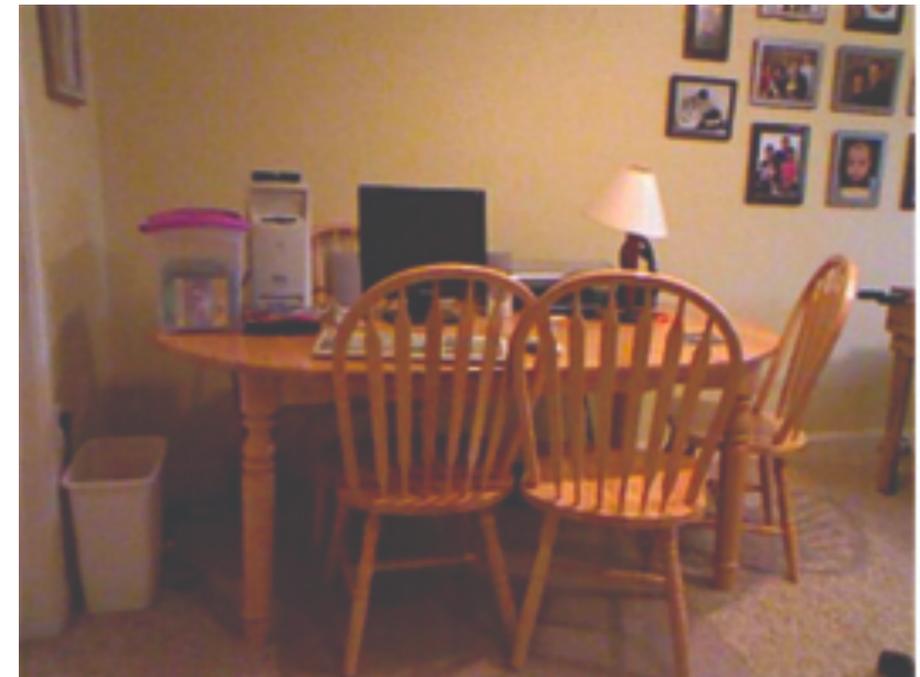


Easy-to-match pattern

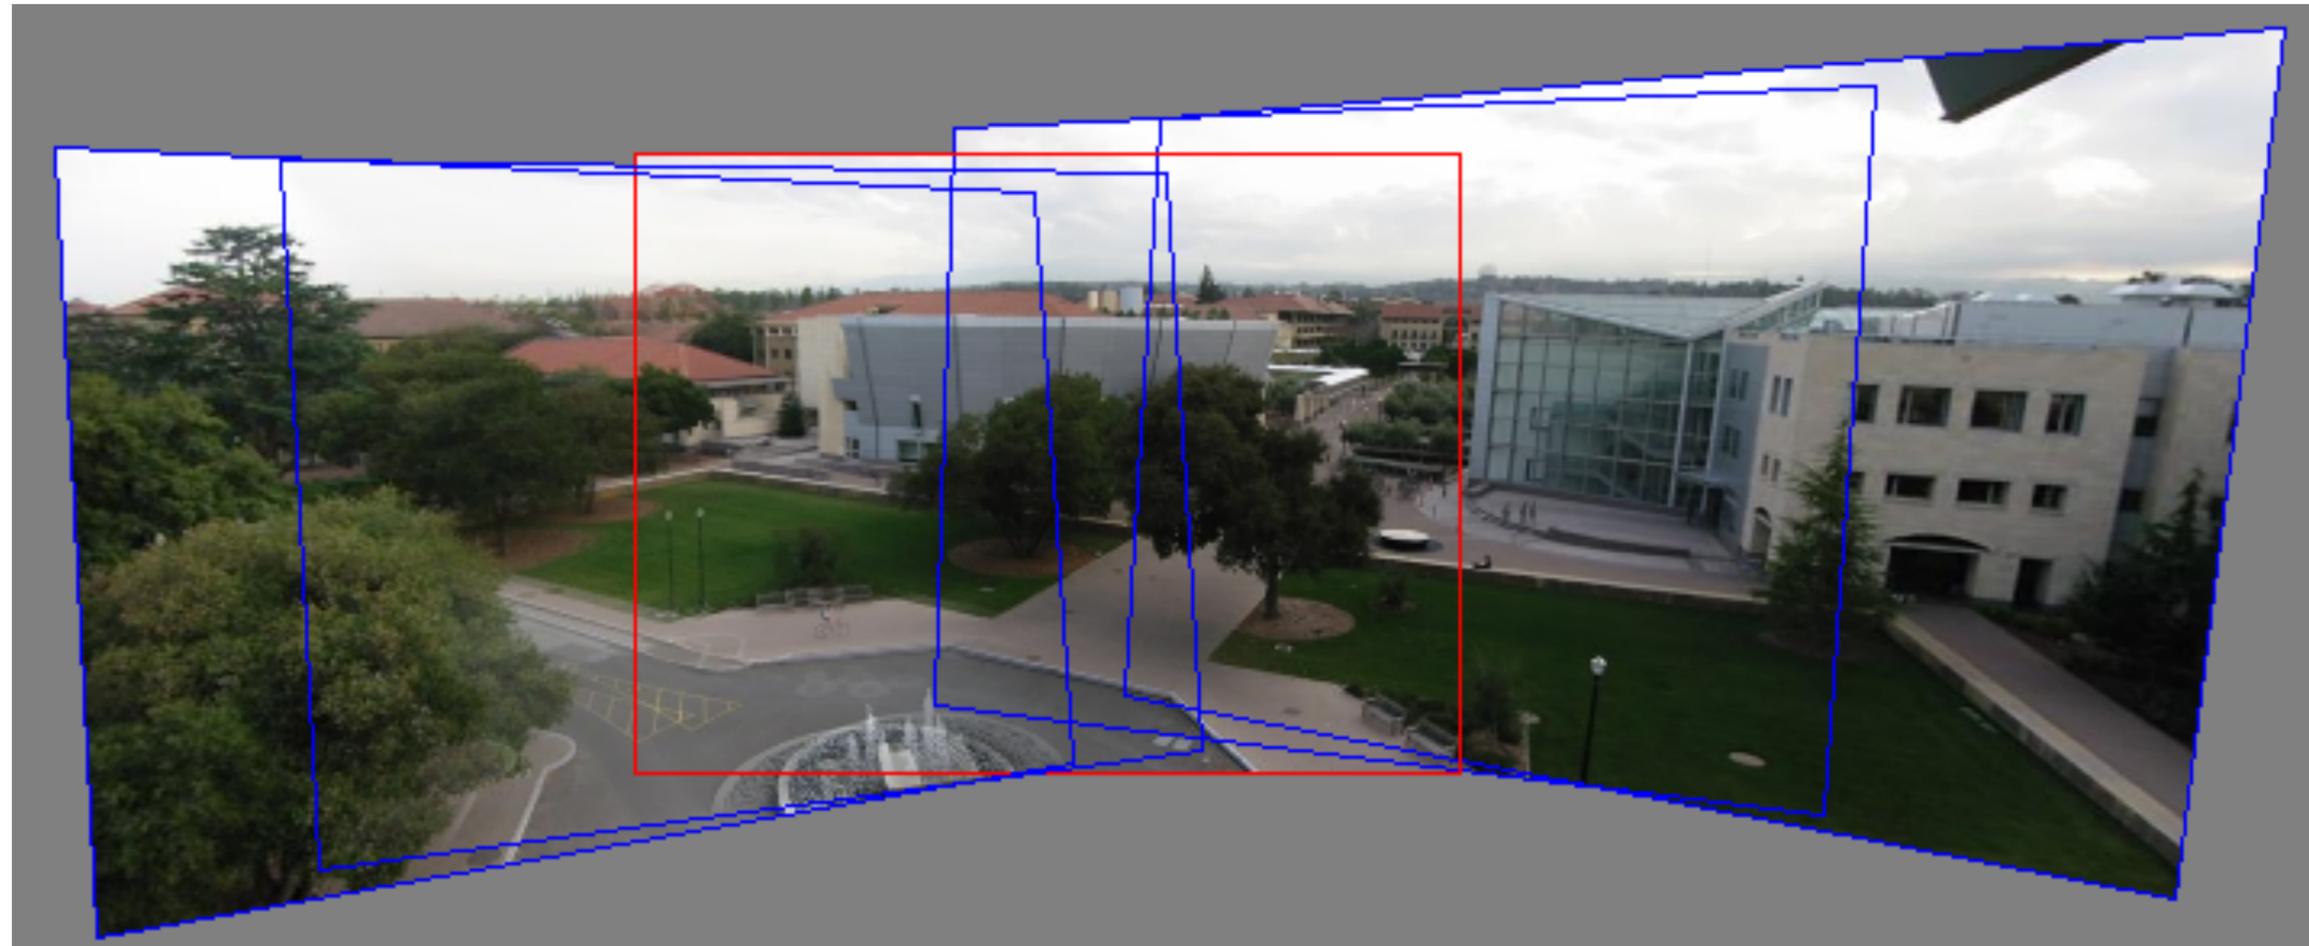


Do we really need the second camera?

# RGB-D sensors

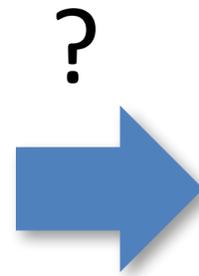


# Making panoramas



# Making panoramas

What is the geometric relationship between these two images?



# Image alignment



Why don't these images line up exactly?

# Recall: affine transformations

$$\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix}$$

affine transformation

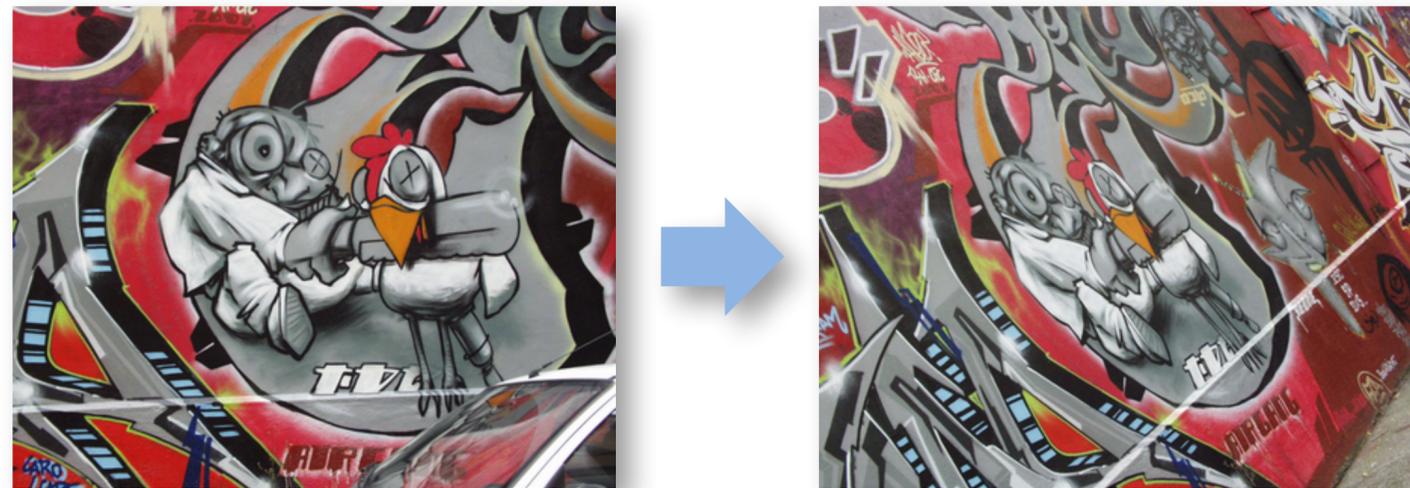
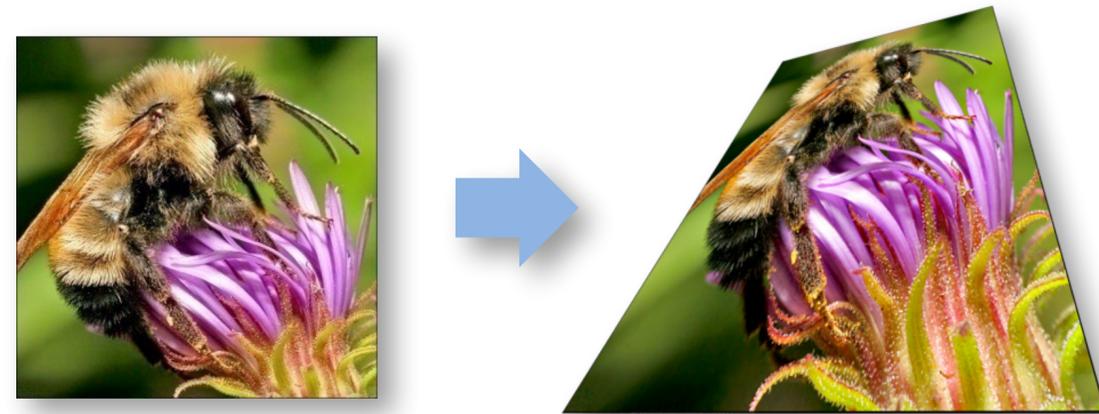


what happens when we  
change this row?

# Projective Transformations aka Homographies aka Planar Perspective Maps

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

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:



If we try to build a panorama by overlapping them:



# Homography

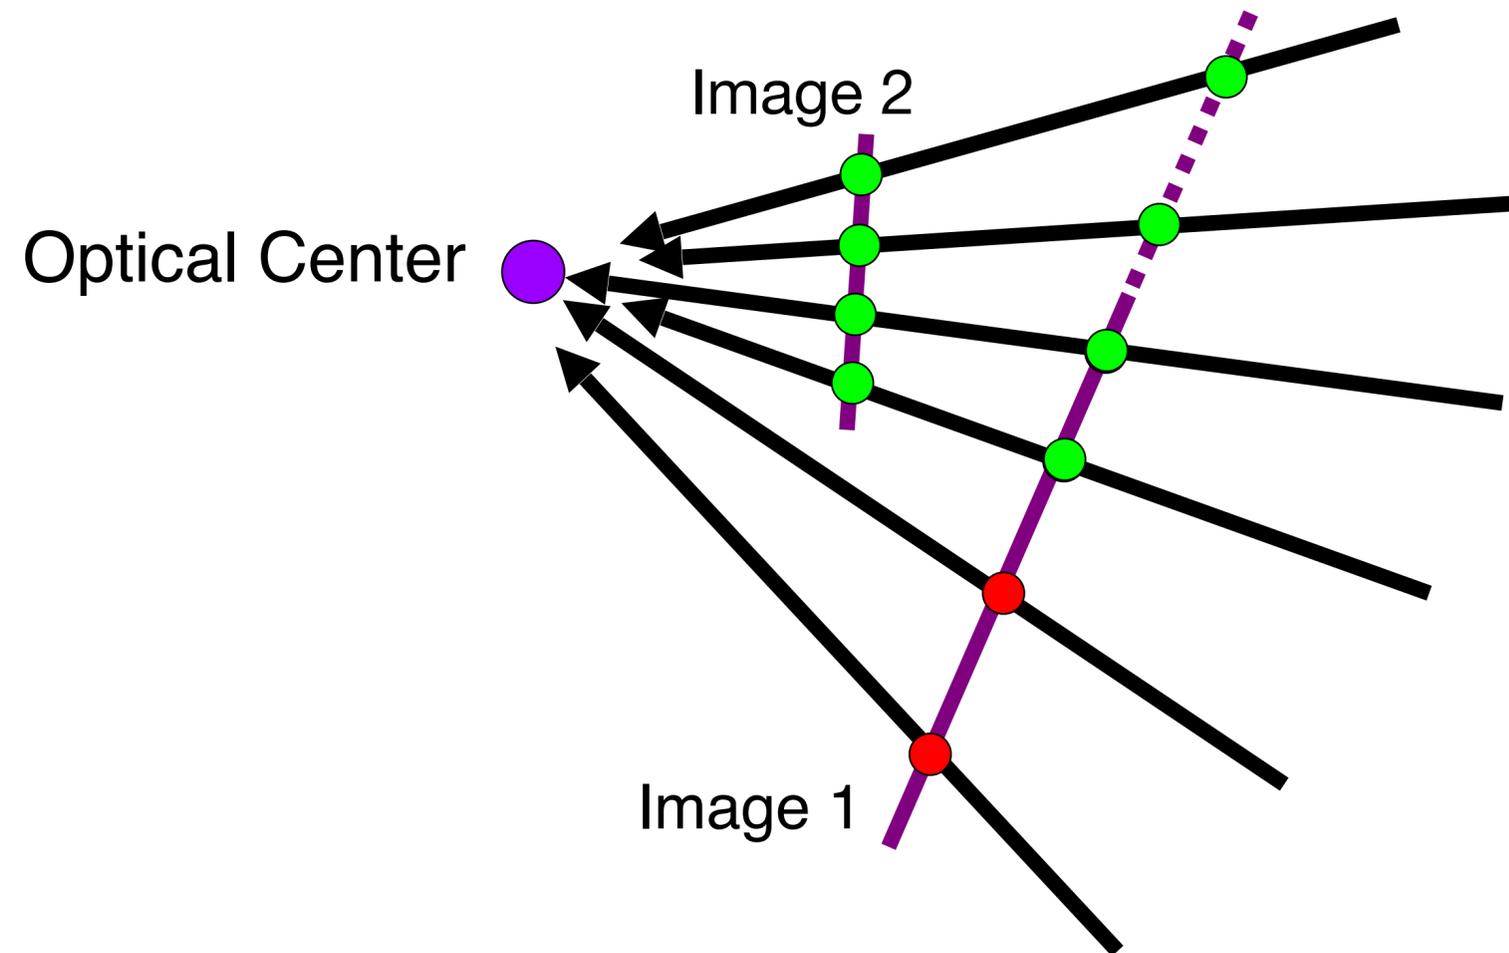
Example: two pictures taken by rotating the camera:



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



# Why does this work?



**Step 1:** Convert pixels in image 2 to rays in camera 2's coordinate system.

$$\begin{bmatrix} X_2 \\ Y_2 \\ Z_2 \end{bmatrix} = \mathbf{K}_2^{-1} \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix}$$

**Step 2:** Convert rays in camera 2's coordinates to rays in camera 1's coordinates.

$$\begin{bmatrix} X_1 \\ Y_1 \\ Z_1 \end{bmatrix} = \mathbf{R}_2^T \mathbf{K}_2^{-1} \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix}$$

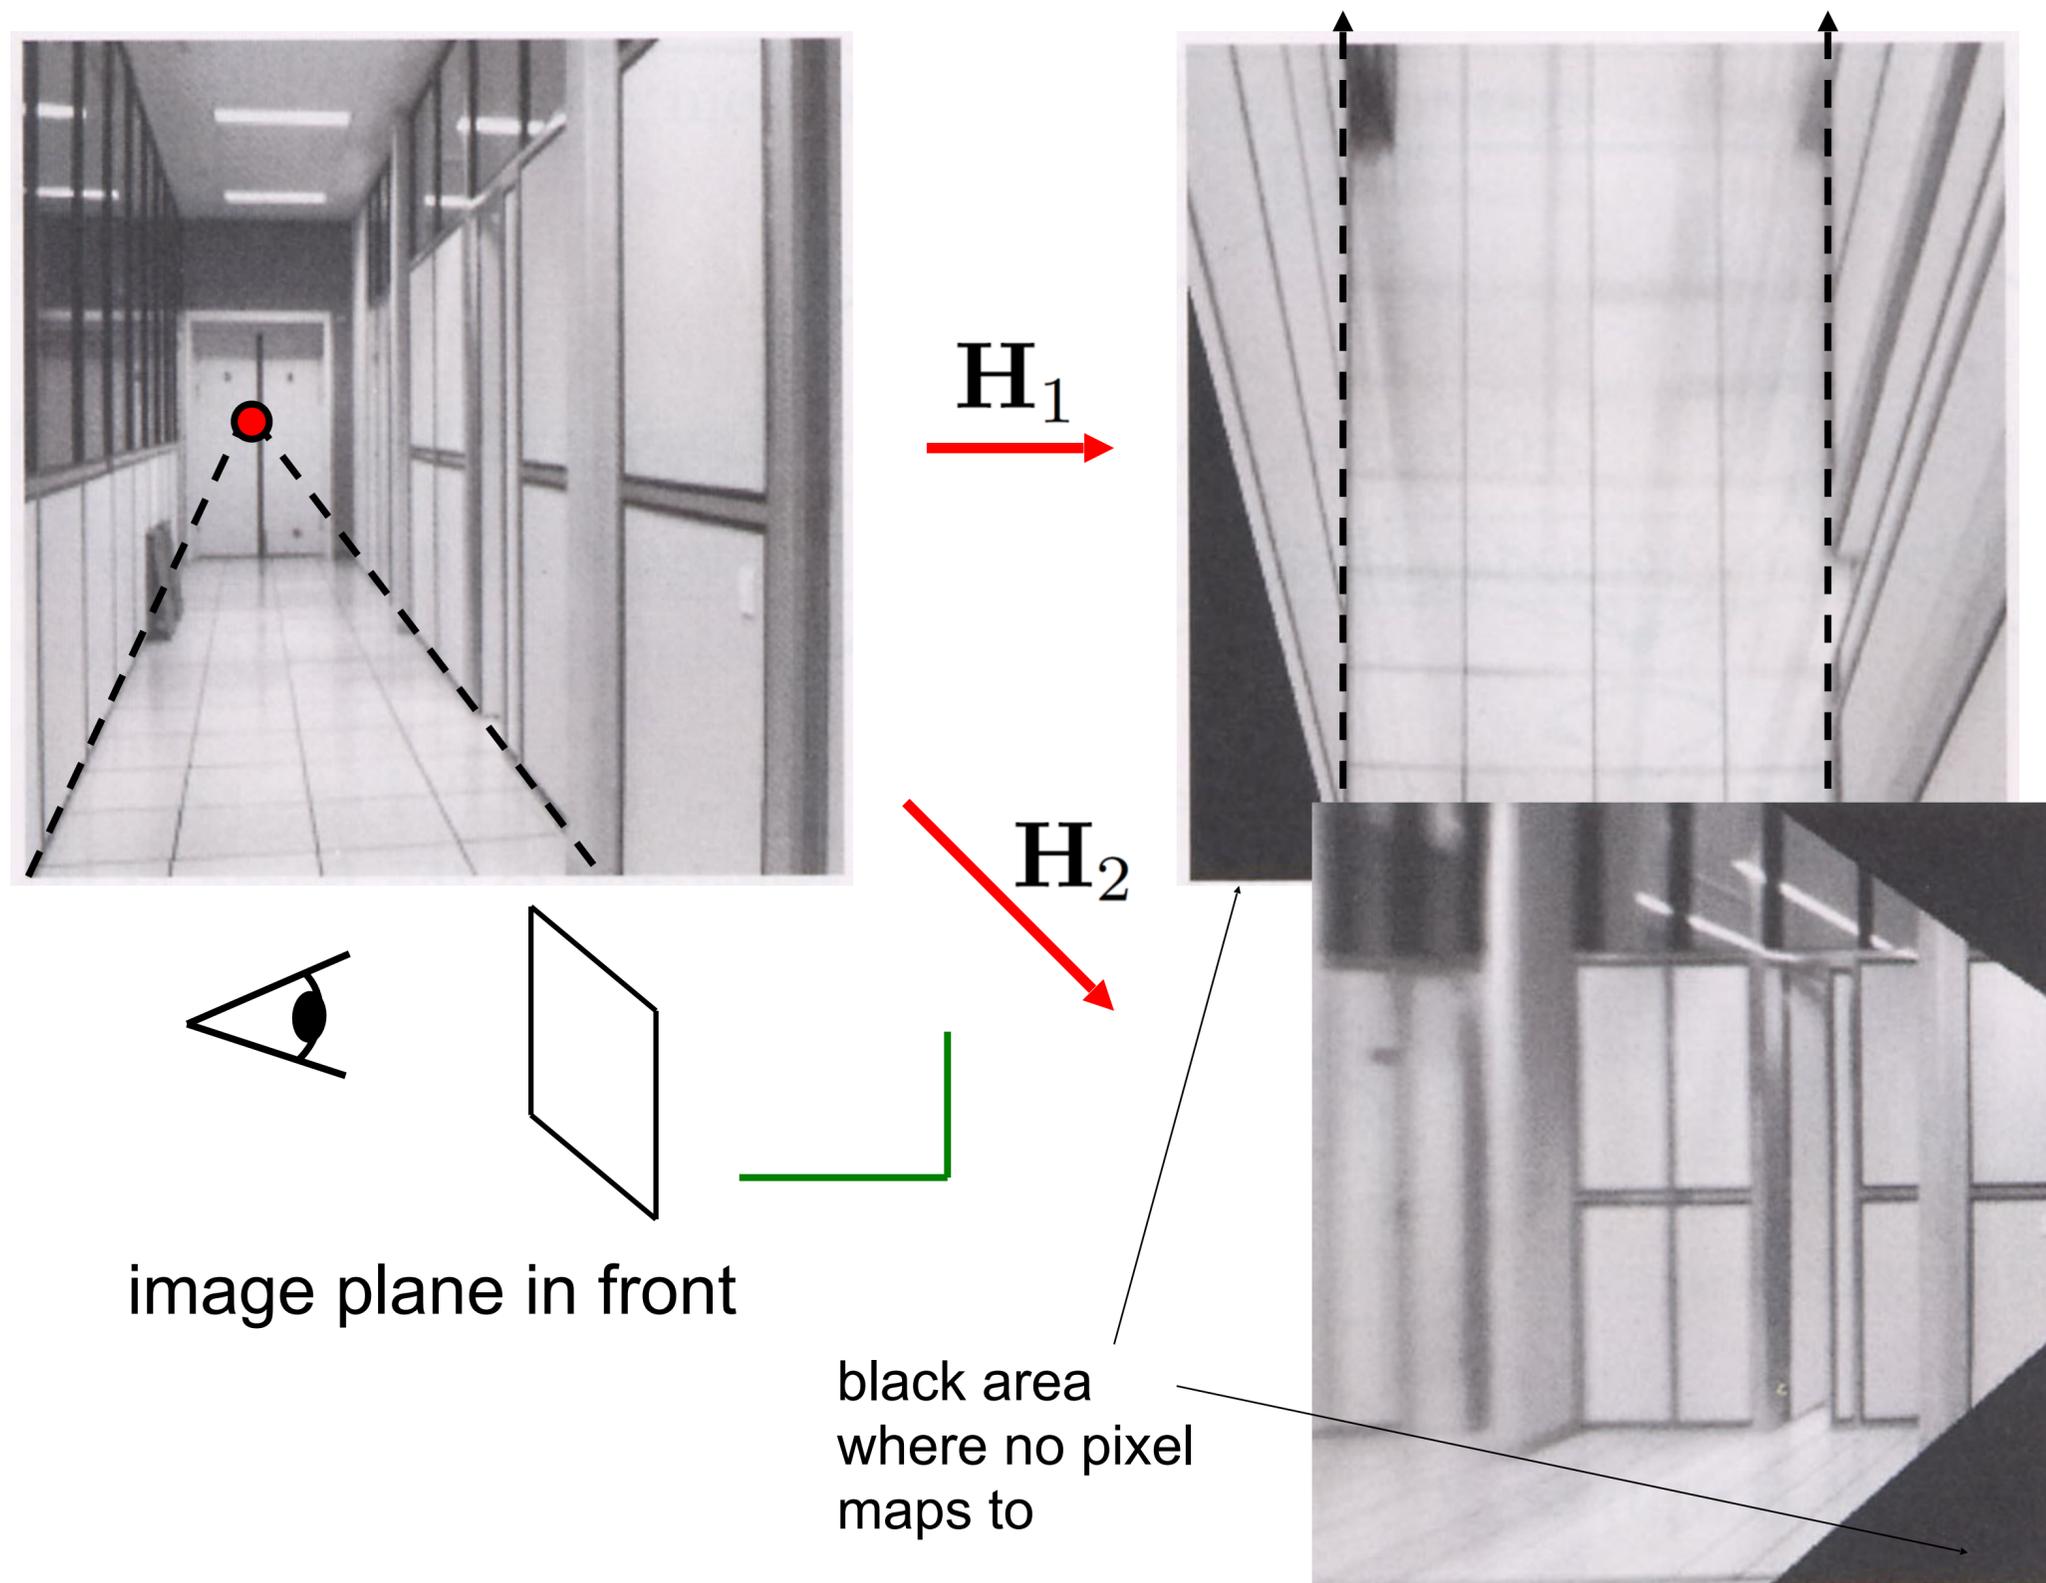
**Step 3:** Convert rays in camera 1's coordinates to pixels in image 1's coordinates.

$$\begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} \sim \underbrace{\mathbf{K}_1 \mathbf{R}_2^T \mathbf{K}_2^{-1}}_{\substack{\text{3x3 homography} \\ \uparrow}} \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix}$$

**How do we map points in image 2 into image 1?**

	image 1	image 2
intrinsics	$\mathbf{K}_1$	$\mathbf{K}_2$
extrinsics (rotation only)	$\mathbf{R}_1 = \mathbf{I}_{3 \times 3}$	$\mathbf{R}_2$

# Plane-to-plane homography



# Homographies

- Homographies ...

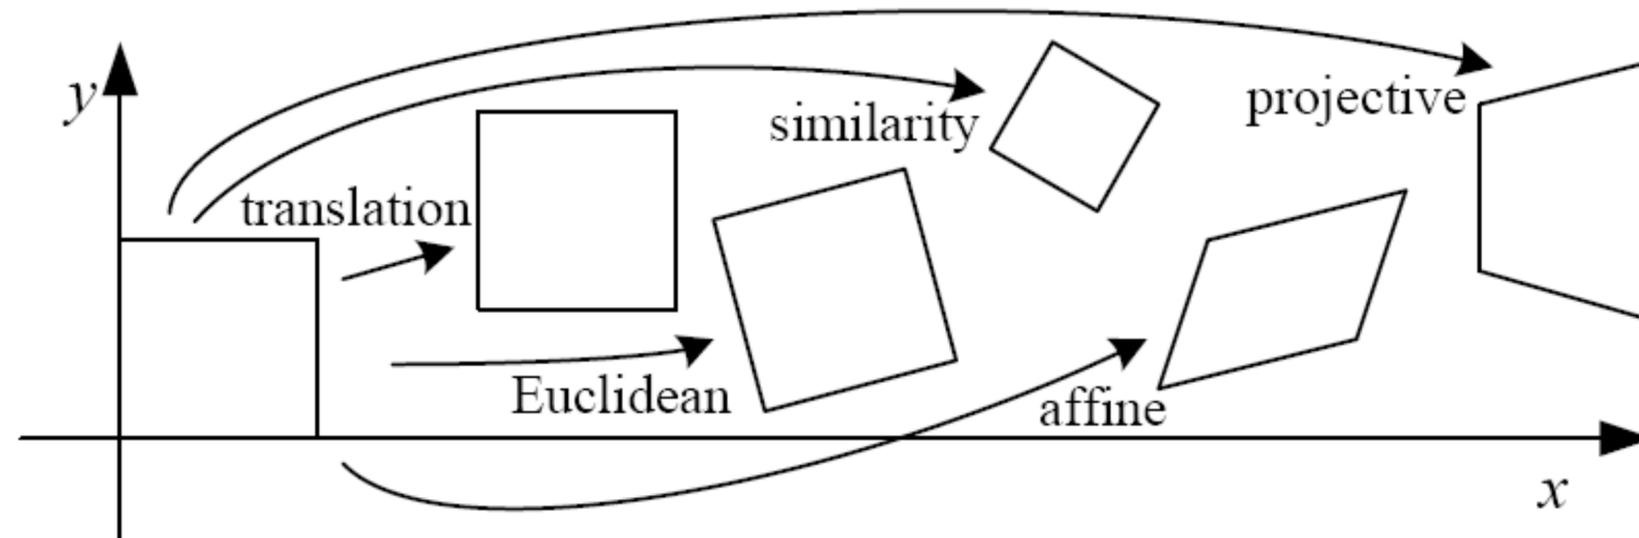
- Affine transformations, and
- Projective warps

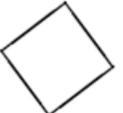
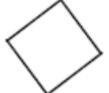
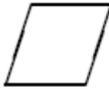
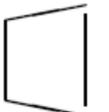
$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

- 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

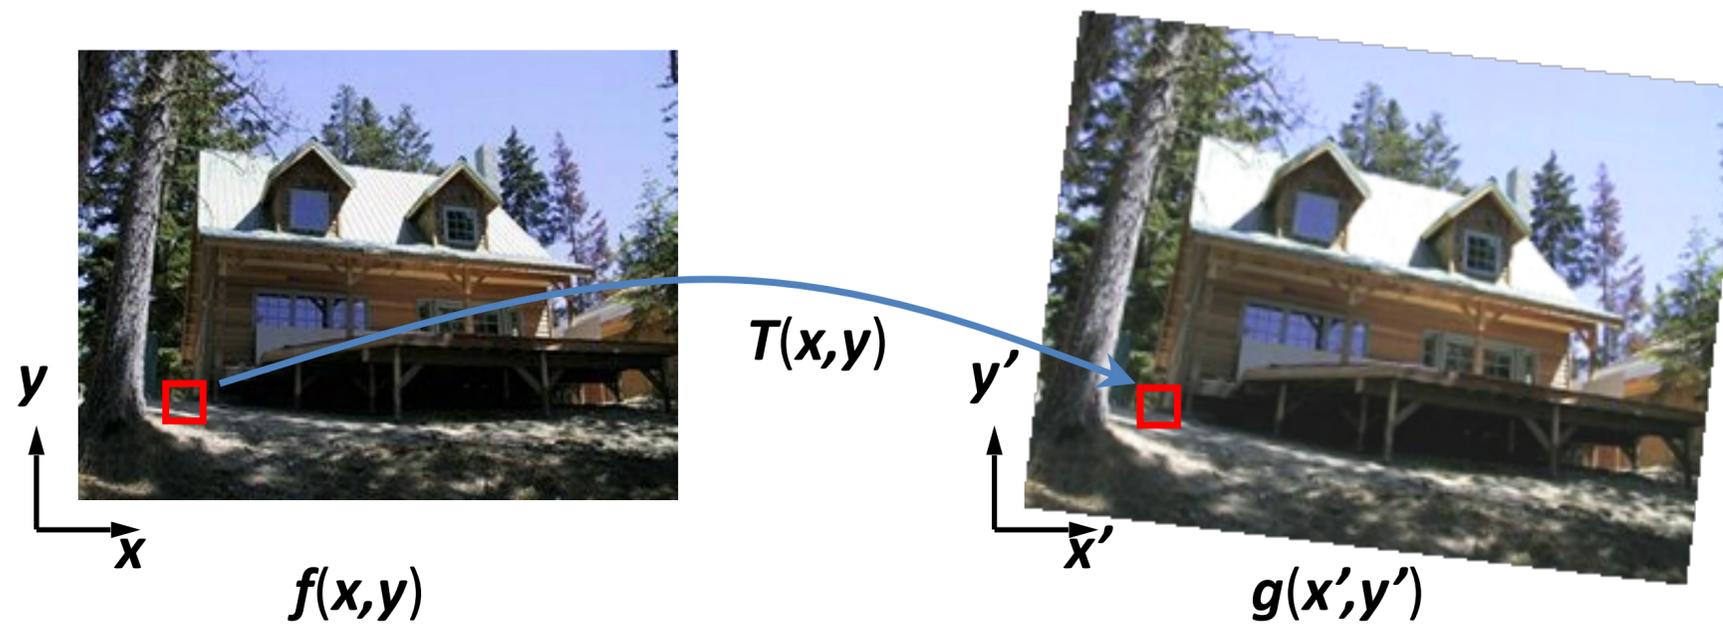


Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	2	orientation + ...	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	3	lengths + ...	
similarity	$\begin{bmatrix} s\mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	4	angles + ...	
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{2 \times 3}$	6	parallelism + ...	
projective	$\begin{bmatrix} \tilde{\mathbf{H}} \end{bmatrix}_{3 \times 3}$	8	straight lines	

How do we perform this warp?

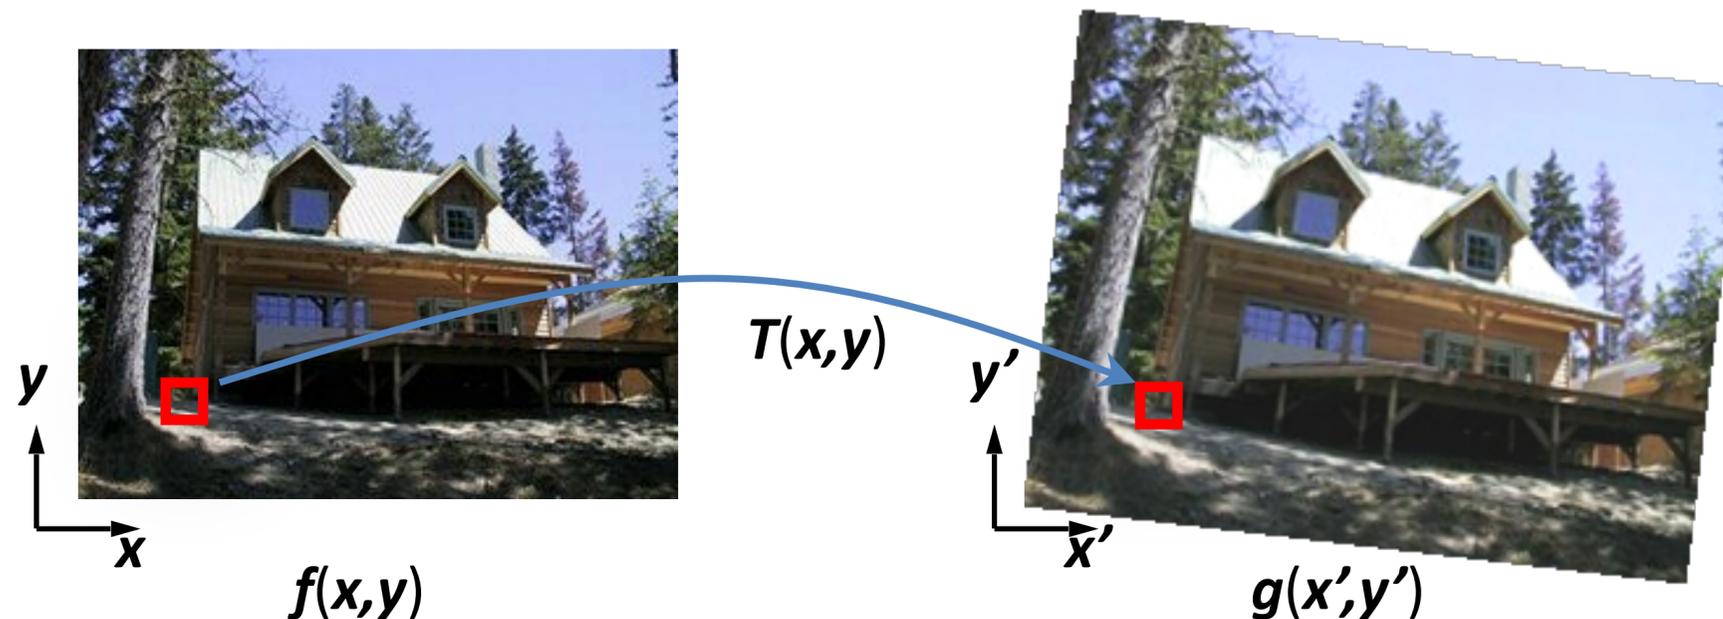
# Image warping

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



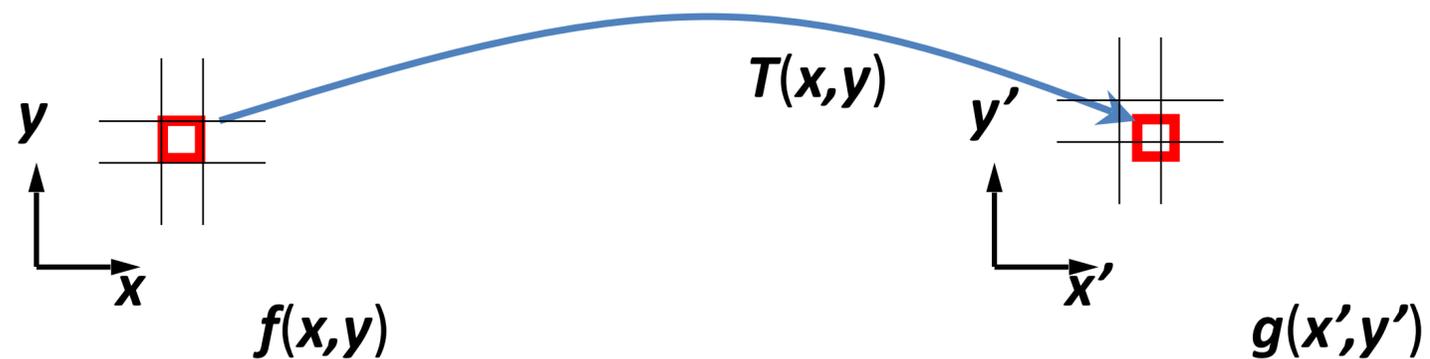
# Forward warping

- 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?



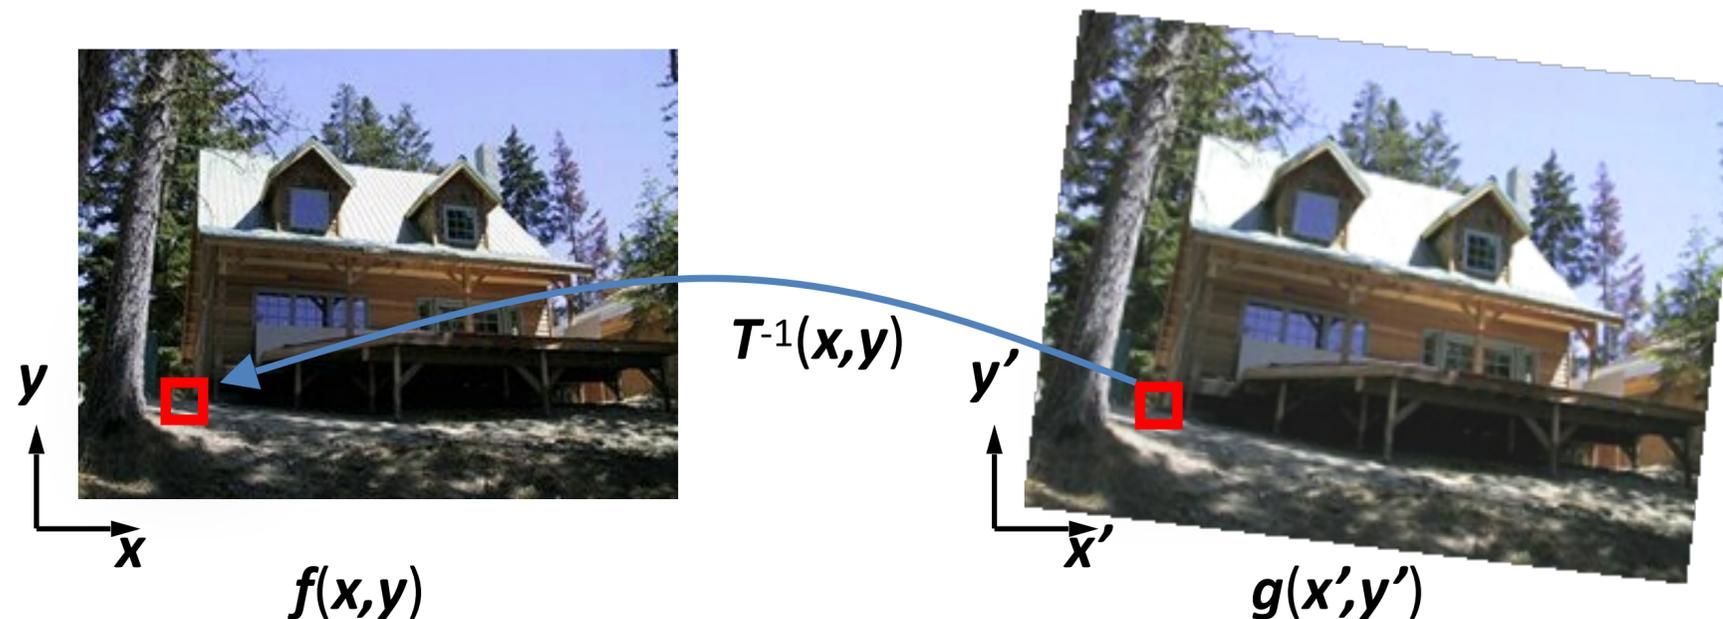
# Forward warping

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



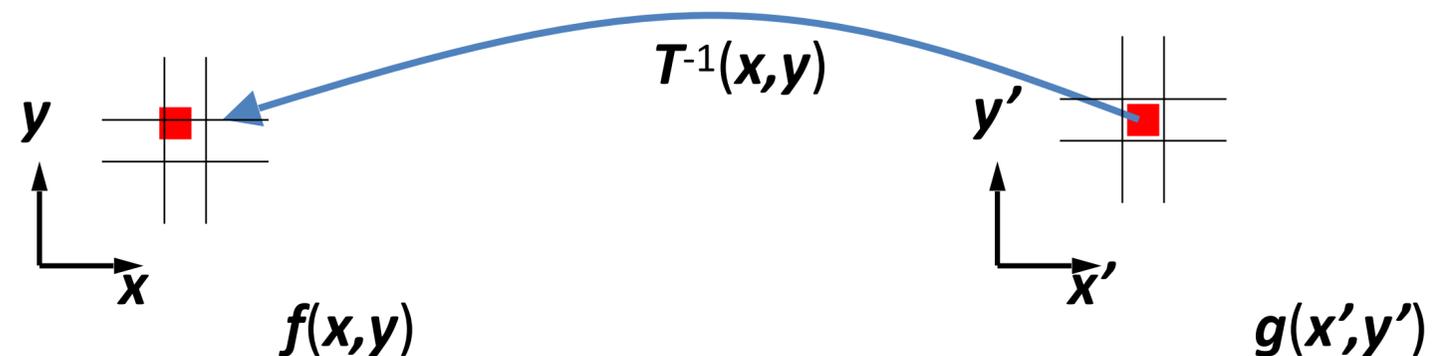
# Inverse warping

- 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?



# Inverse warping

- Get each pixel  $g(\mathbf{x}')$  from its corresponding location  $\mathbf{x}' = \mathbf{h}(\mathbf{x})$  in  $f(\mathbf{x})$
- What if pixel comes from “between” two pixels?
- Answer: *resample* color value from *interpolated* source image



**Next lecture: estimating geometry from images**