Review

- Previous section:
  - Feature detection and matching
  - Model fitting and outlier rejection

Project 2 questions?

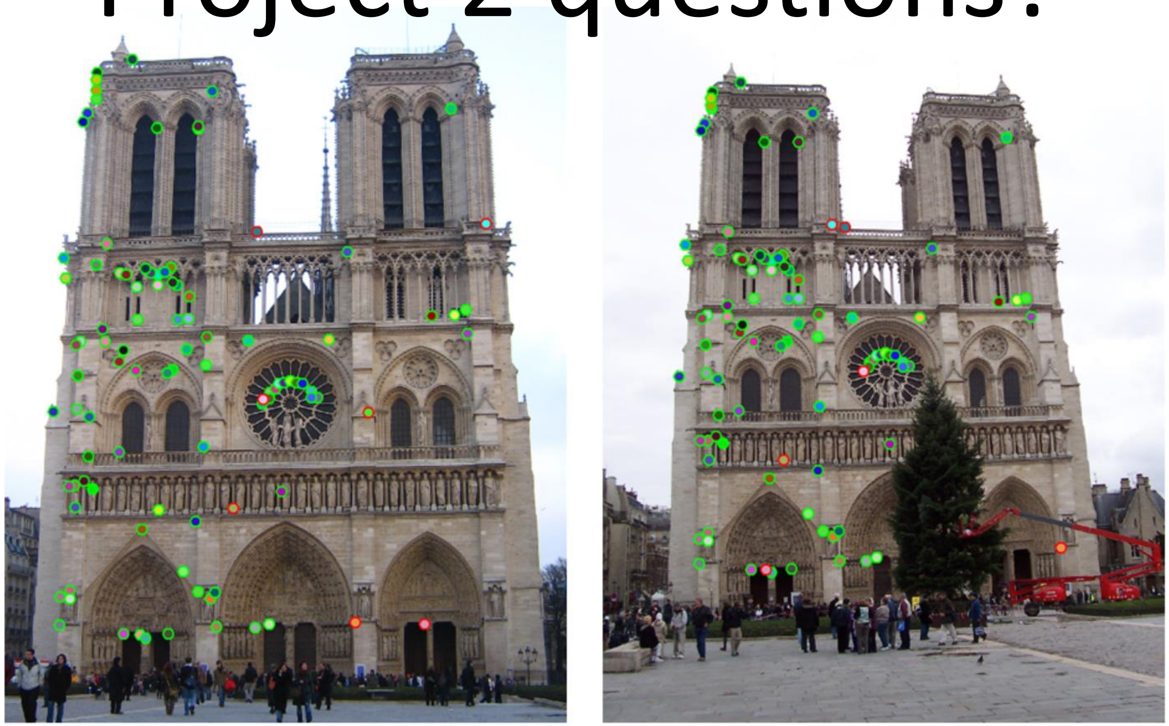

The top 100 most confident local feature matches from a baseline implementation of project 2. In this case, 93 were correct (highlighted in green) and 7 were incorrect (highlighted in red).

Project 2: Local Feature Matching

CS 143: Introduction to Computer Vision

Brief

- Due: 11:59pm on Monday, October 7th, 2013
- Stencil code: `/course/cs143/asgn/proj2/code/`
- Data: `/course/cs143/asgn/proj2/data/` includes 93 images from 9 different outdoor scenes.
- Html writeup template: `/course/cs143/asgn/proj2/html/`
- *Partial* project materials are also available in **proj2.zip (1.7 MB)**. Includes only the two test images shown above.
- Handin: `cs143_handin proj2`
- Required files: README, code/, html/, html/index.html

Review: Interest points

- Keypoint detection: repeatable and distinctive
  - Corners, blobs, stable regions
  - Harris, DoG, MSER

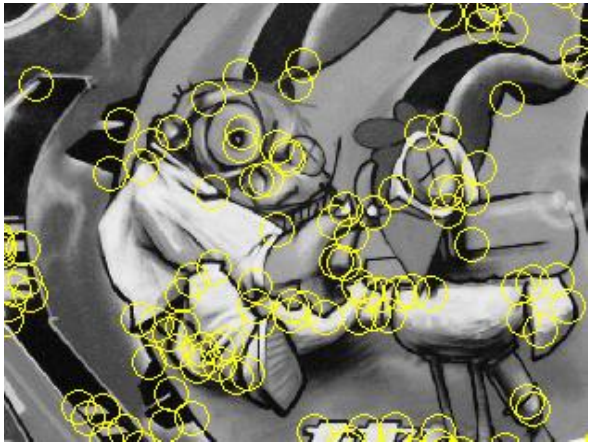

Harris Detector [Harris88]

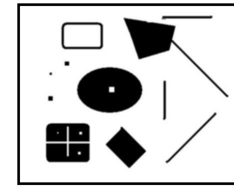

- Second moment matrix

$$\mu(\sigma_I, \sigma_D) = g(\sigma_I) * \begin{bmatrix} I_x^2(\sigma_D) & I_x I_y(\sigma_D) \\ I_x I_y(\sigma_D) & I_y^2(\sigma_D) \end{bmatrix}$$

1. Image derivatives (optionally, blur first)

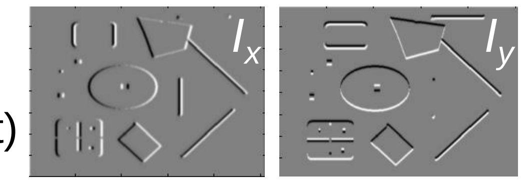

2. Square of derivatives

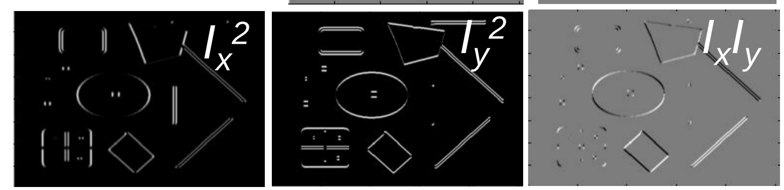

$$\det M = \lambda_1 \lambda_2$$

$$\operatorname{trace} M = \lambda_1 + \lambda_2$$

3. Gaussian filter $g(\sigma_I)$

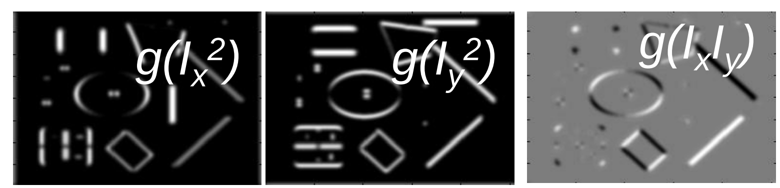

4. Cornerness function – both eigenvalues are strong

$$har = \det[\mu(\sigma_I, \sigma_D)] - \alpha[\operatorname{trace}(\mu(\sigma_I, \sigma_D))^2] =$$

$$g(I_x^2)g(I_y^2) - [g(I_x I_y)]^2 - \alpha[g(I_x^2) + g(I_y^2)]^2$$

5. Non-maxima suppression

Review: Local Descriptors

- Most features can be thought of as templates, histograms (counts), or combinations
- Most available descriptors focus on edge/gradient information
  - Capture texture information
  - Color rarely used

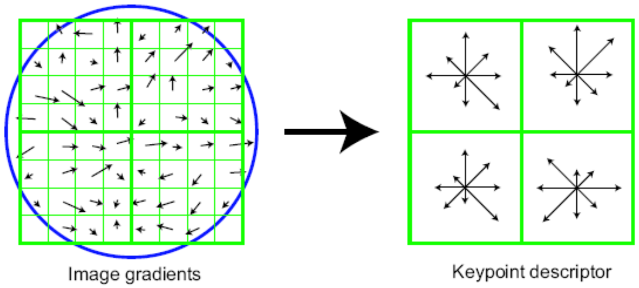

K. Grauman, B. Leibe

Review: Hough transform

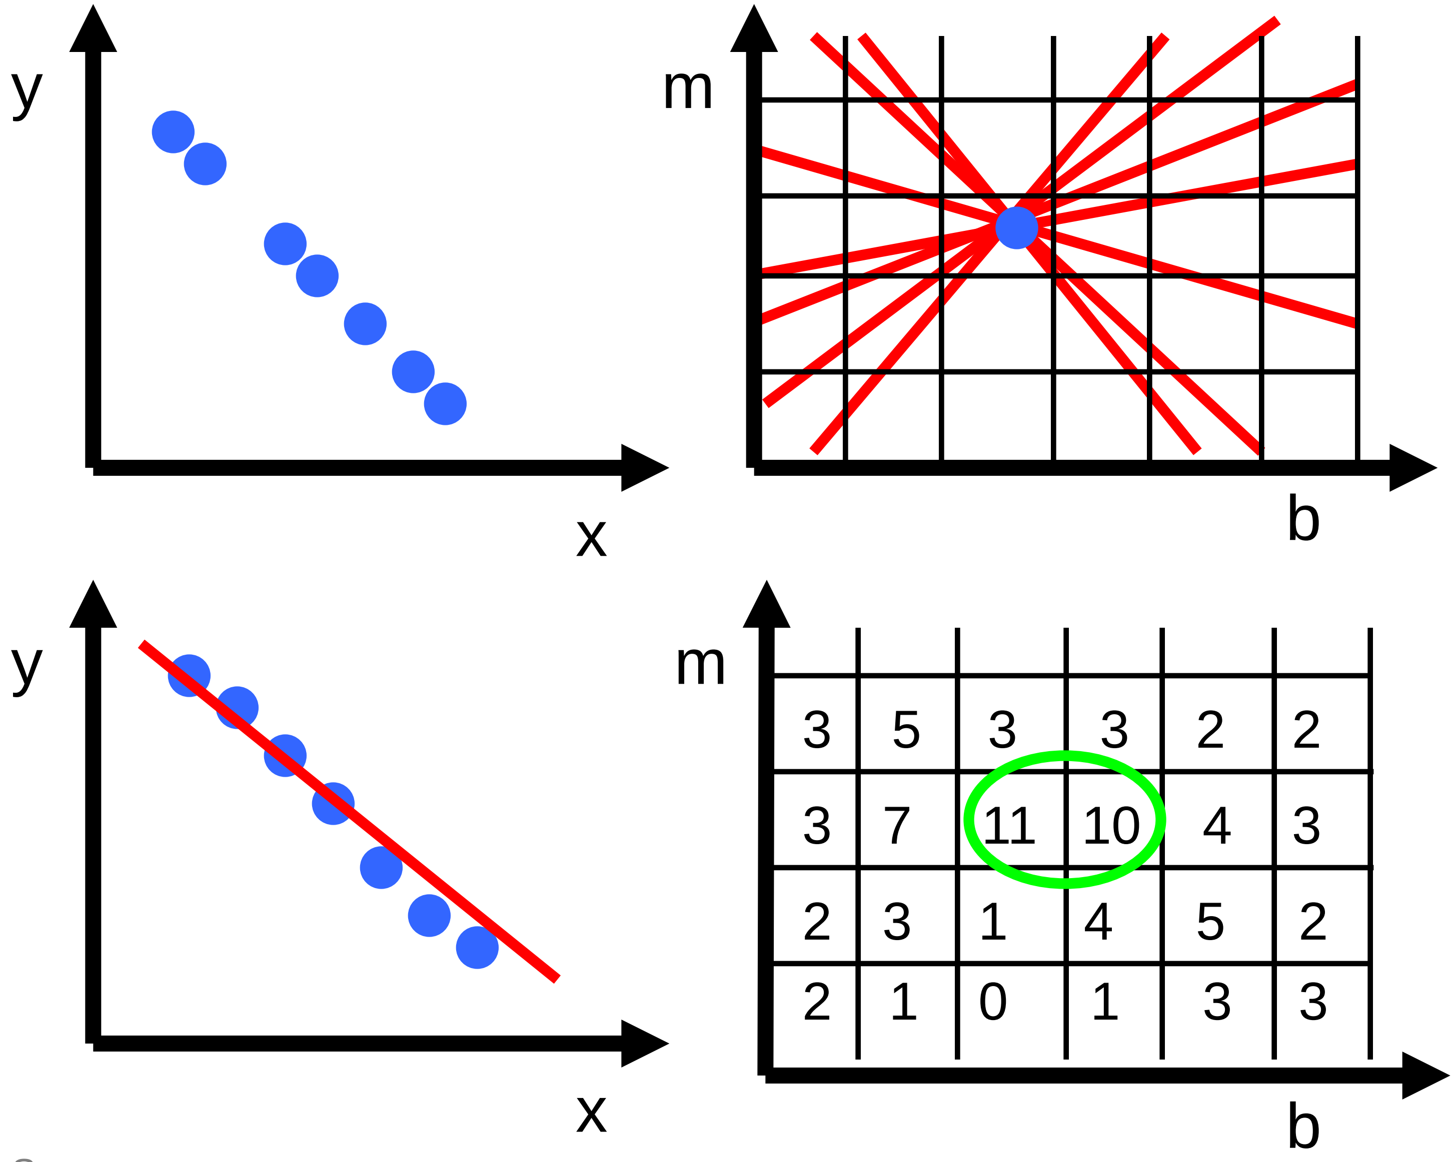

Slide from S. Savarese

Review: RANSAC

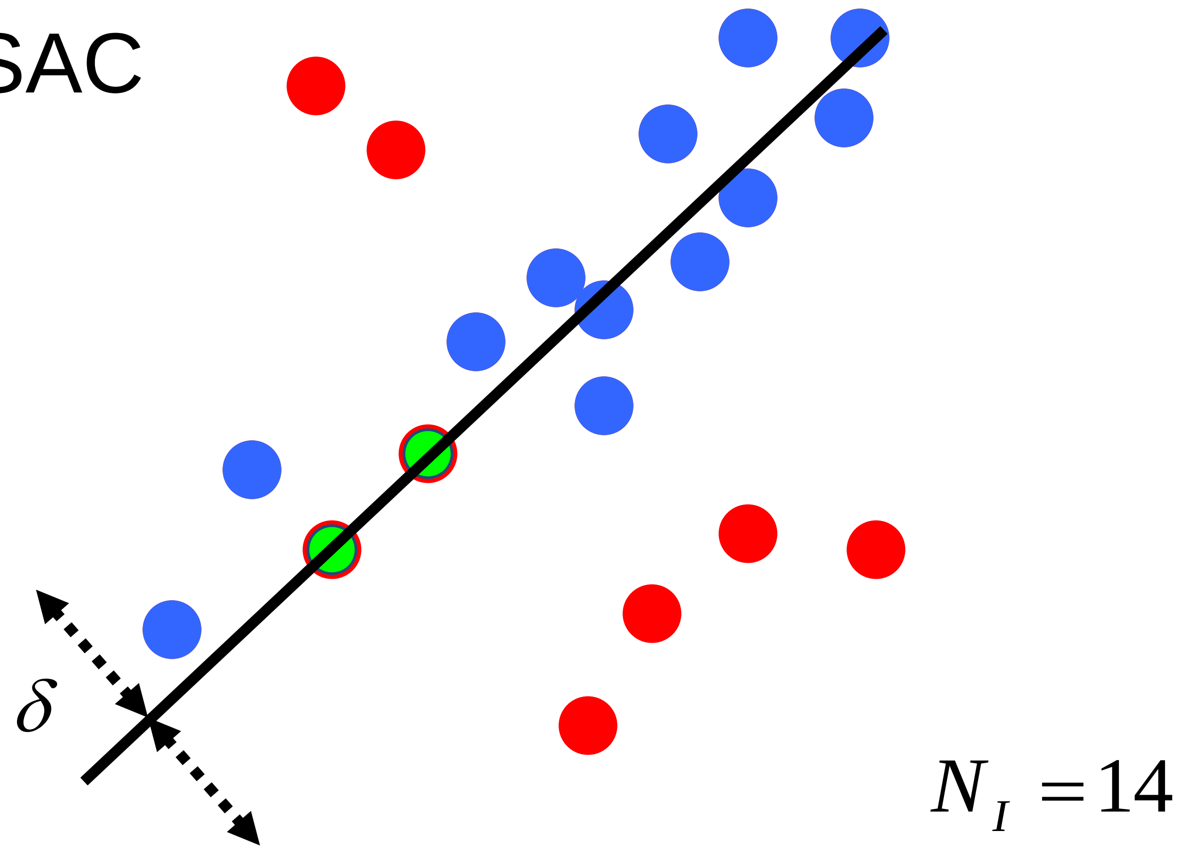

Algorithm:

1. **Sample** (randomly) the number of points required to fit the model (#=2)
2. **Solve** for model parameters using samples
3. **Score** by the fraction of inliers within a preset threshold of the model

Repeat 1-3 until the best model is found with high confidence

Review: 2D image transformations

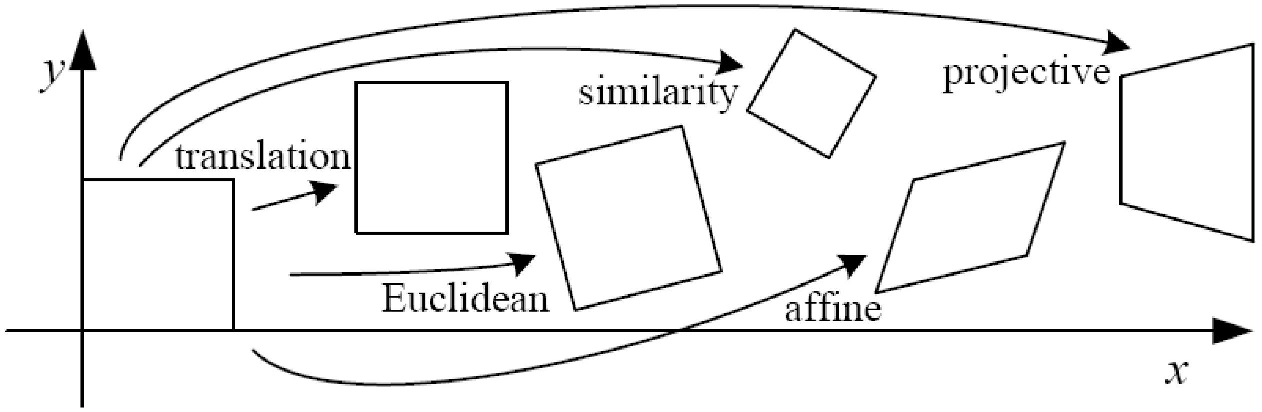

Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} \boldsymbol{I} \mid \boldsymbol{t} \end{bmatrix}_{2\times 3}$	2	orientation $+\cdots$	
rigid (Euclidean)	$\begin{bmatrix} \boldsymbol{R} \mid \boldsymbol{t} \end{bmatrix}_{2\times 3}$	3	lengths $+\cdots$	
similarity	$\begin{bmatrix} s\boldsymbol{R}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

Szeliski 2.1

Stereo: Epipolar geometry

CS143, Brown

James Hays

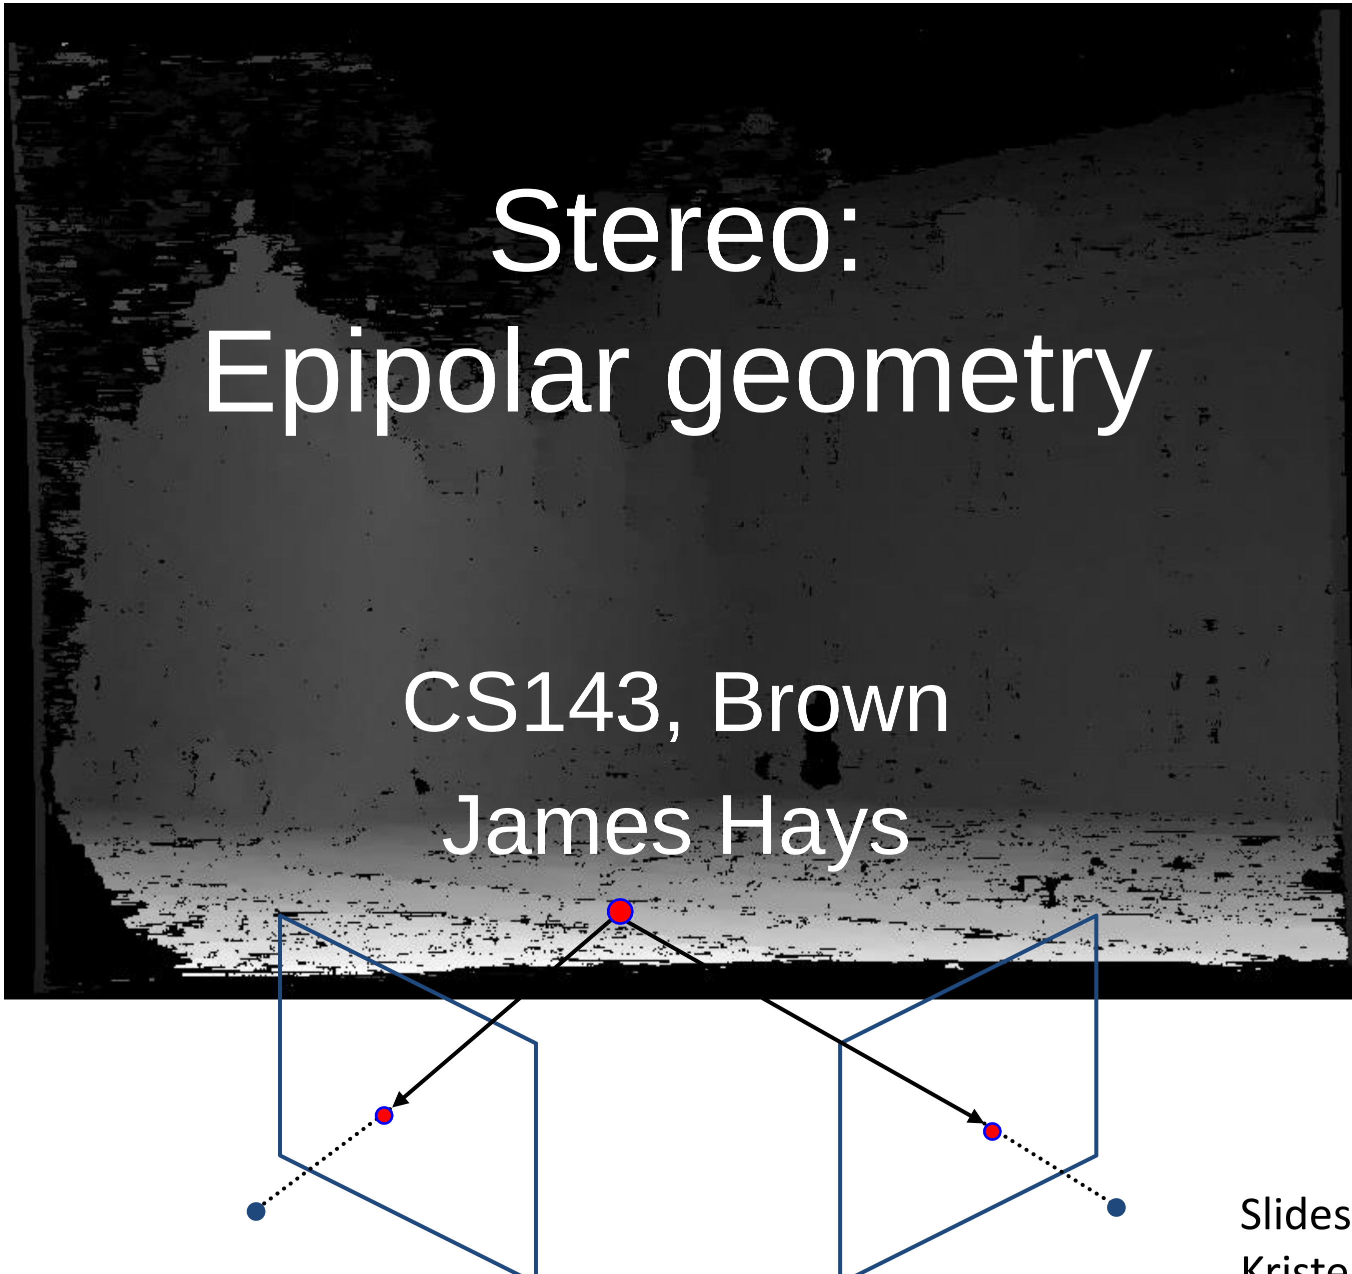

Slides by
Kristen Grauman

Multiple views

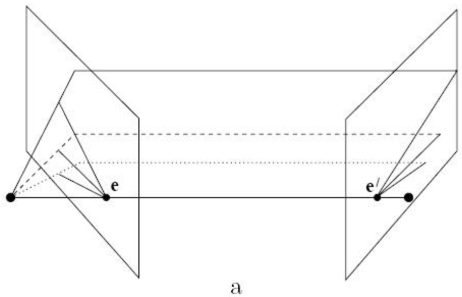

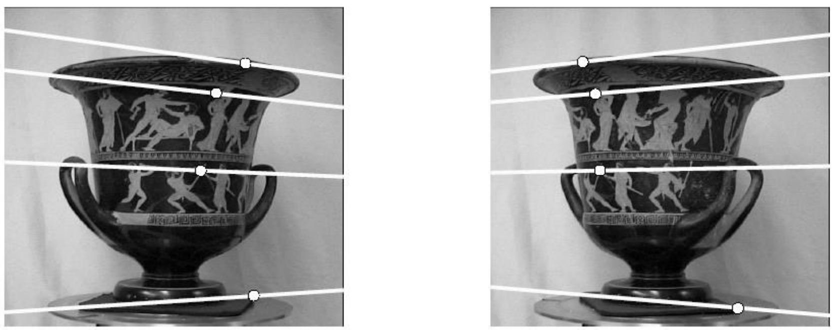

Hartley and Zisserman

stereo vision
structure from motion
optical flow

Why multiple views?

- Structure and depth are inherently ambiguous from single views.

Images from Lana Lazebnik

Why multiple views?

- Structure and depth are inherently ambiguous from single views.

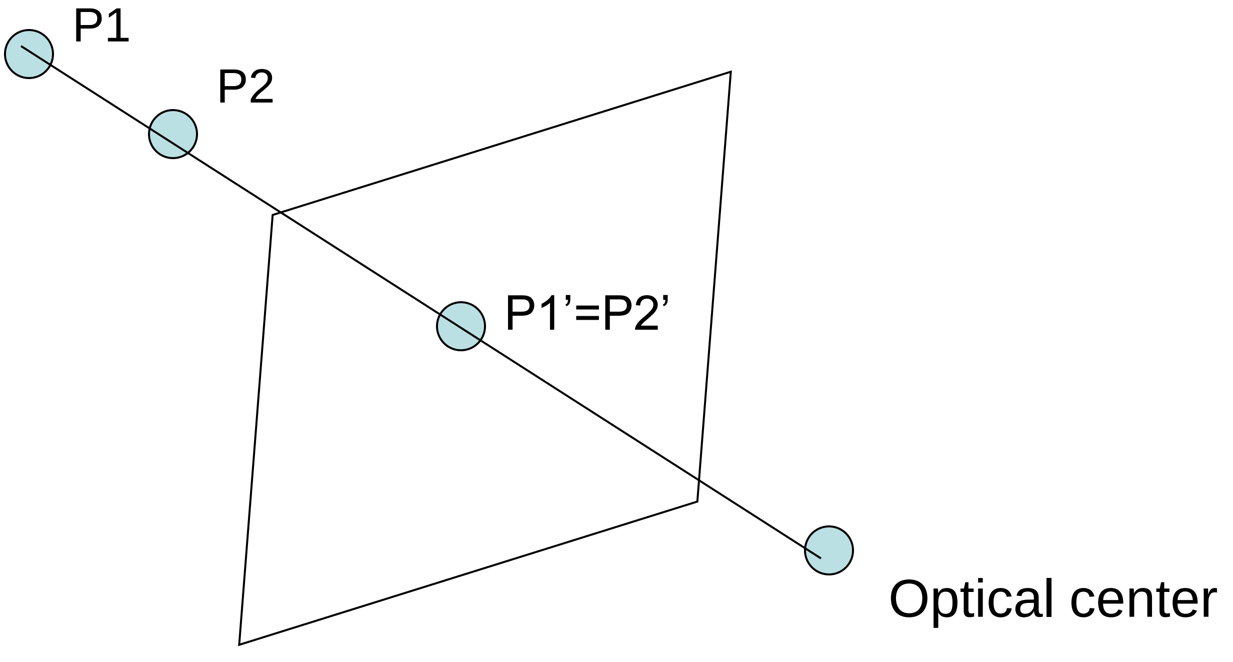

- What cues help us to perceive 3d shape and depth?

Shading

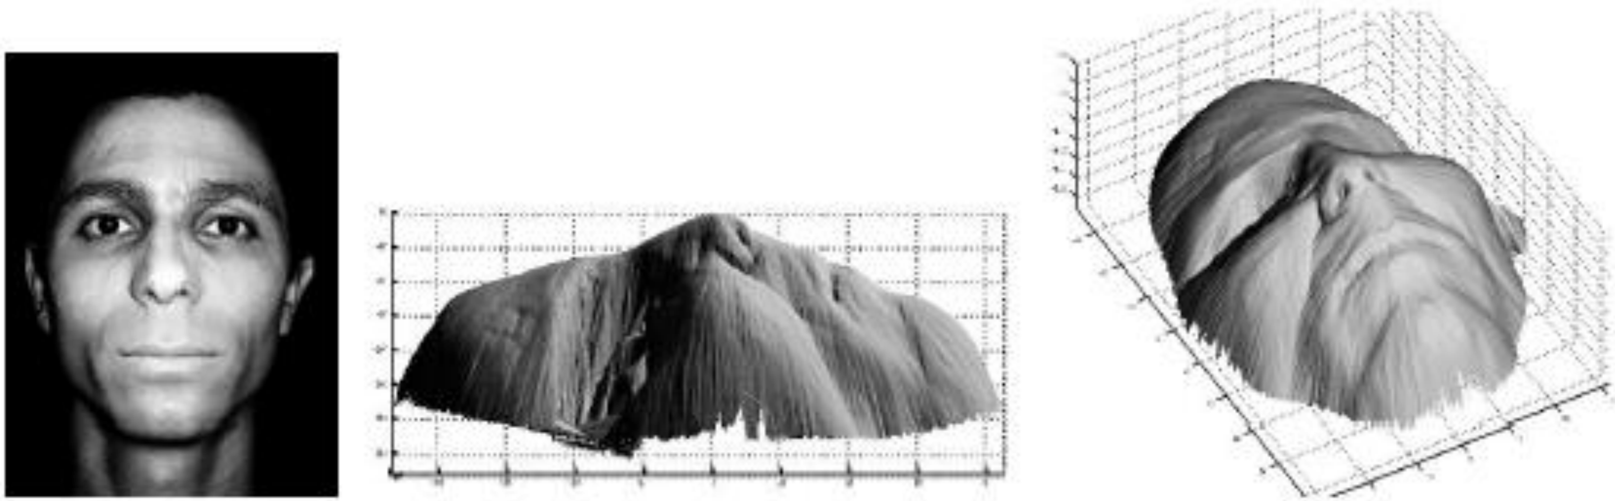

a) b) c)

[Figure from Prados & Faugeras 2006]

Focus/defocus

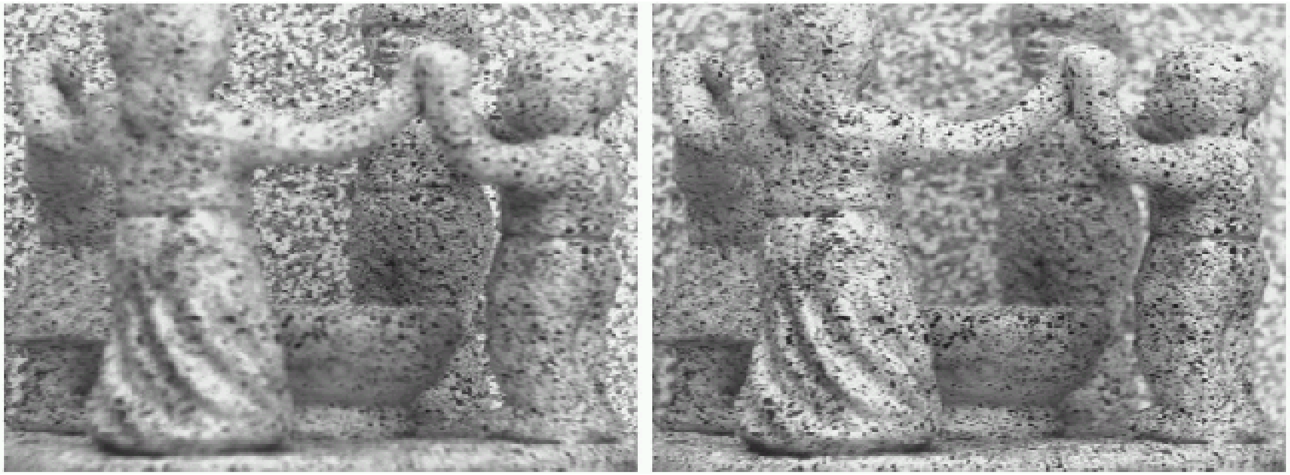

Images from same point of view, different camera parameters

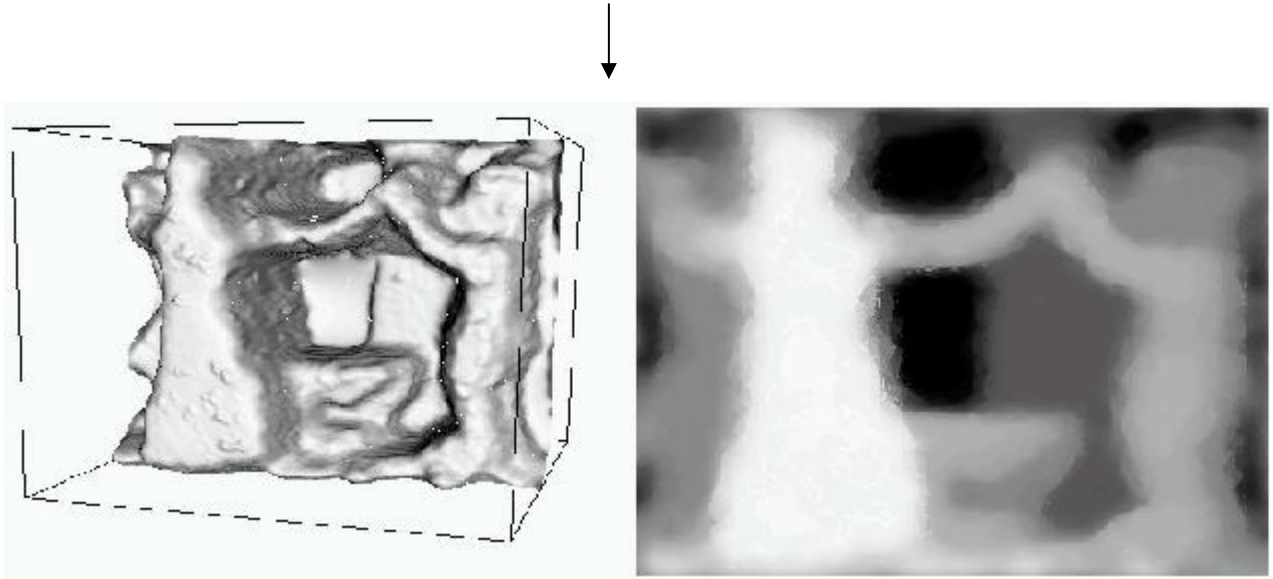

3d shape / depth estimates

[figs from H. Jin and P. Favaro, 2002]

Texture

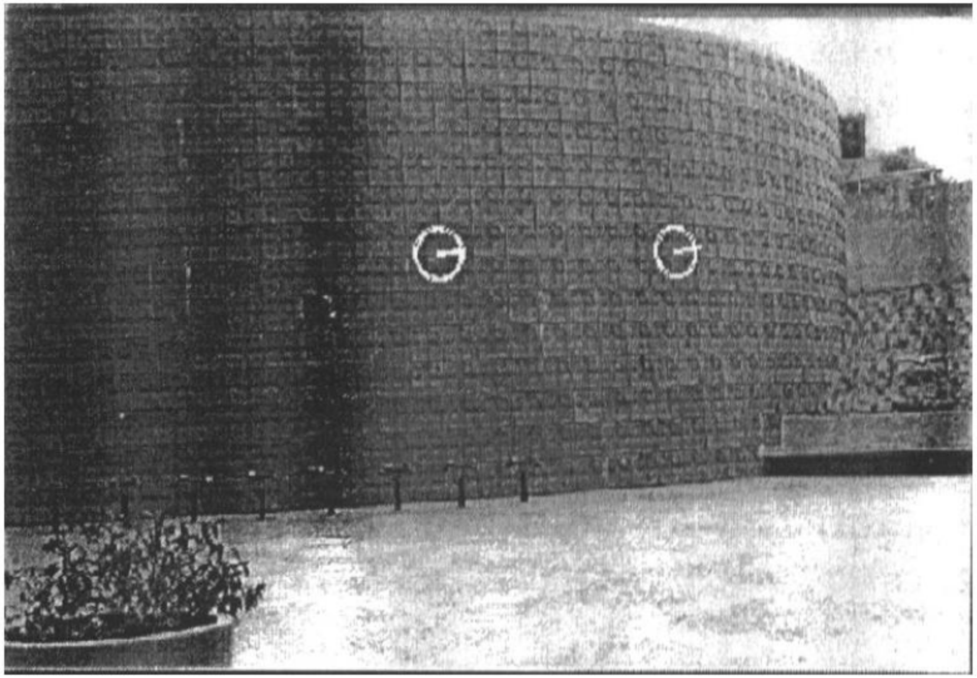

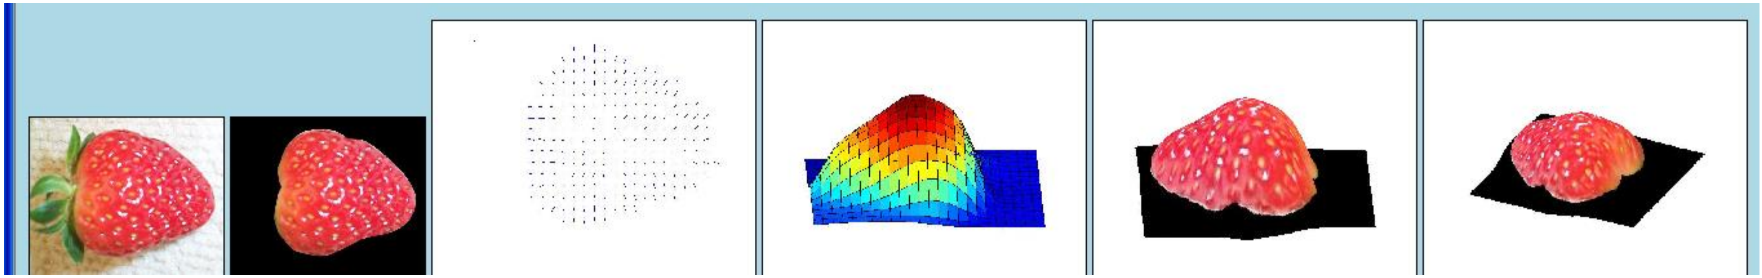

[From A.M. Loh. The recovery of 3-D structure using visual texture patterns. PhD thesis]

Perspective effects

Image credit: S. Seitz

Motion

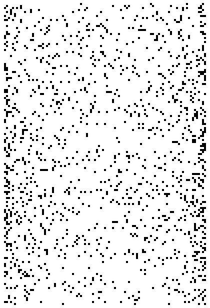

Figures from L. Zhang

http://www.brainconnection.com/teasers/?main=illusion/motion-shape

Occlusion

Rene Magritt'e famous painting *Le Blanc-Seing* (literal translation: "The Blank Signature") roughly translates as "free hand" or "free rein".

Estimating scene shape

- “Shape from X”: Shading, Texture, Focus, Motion…
- **Stereo**:
  - shape from “motion” between two views
  - infer 3d shape of scene from two (multiple) images from different viewpoints

Main idea:

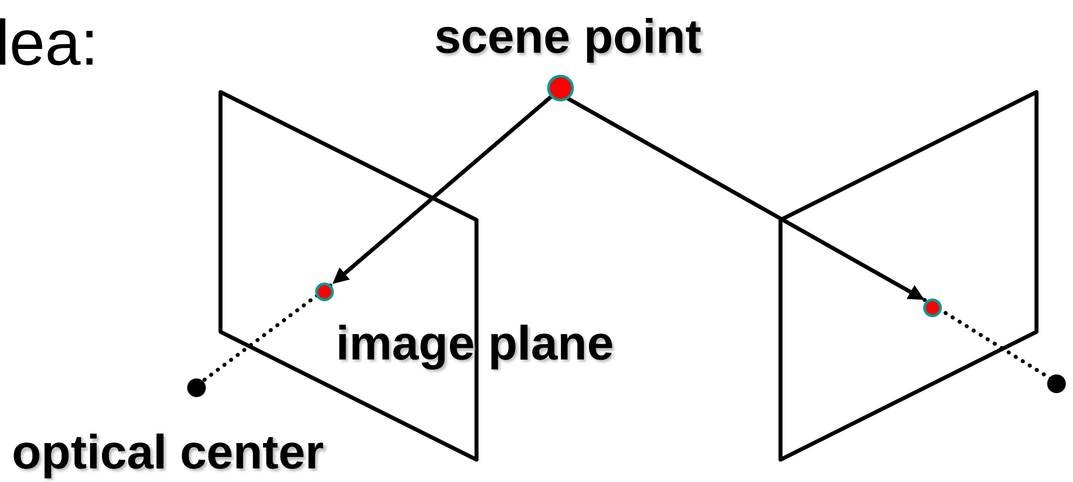

Outline

- Human stereopsis
- Stereograms
- Epipolar geometry and the epipolar constraint
  - Case example with parallel optical axes
  - General case with calibrated cameras

Human eye

Rough analogy with human visual system:

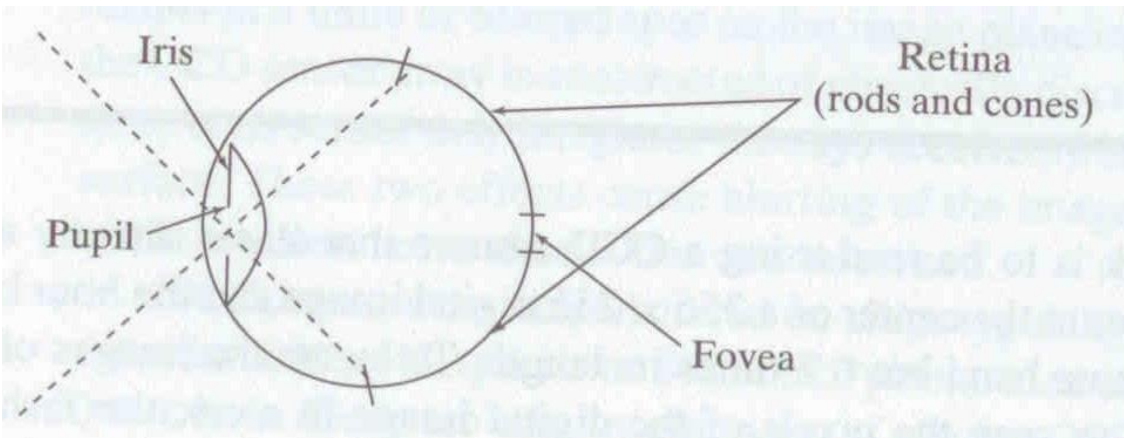

Pupil/Iris – control amount of light passing through lens

Retina - contains sensor cells, where image is formed

Fovea – highest concentration of cones

Fig from Shapiro and Stockman

Human stereopsis: disparity

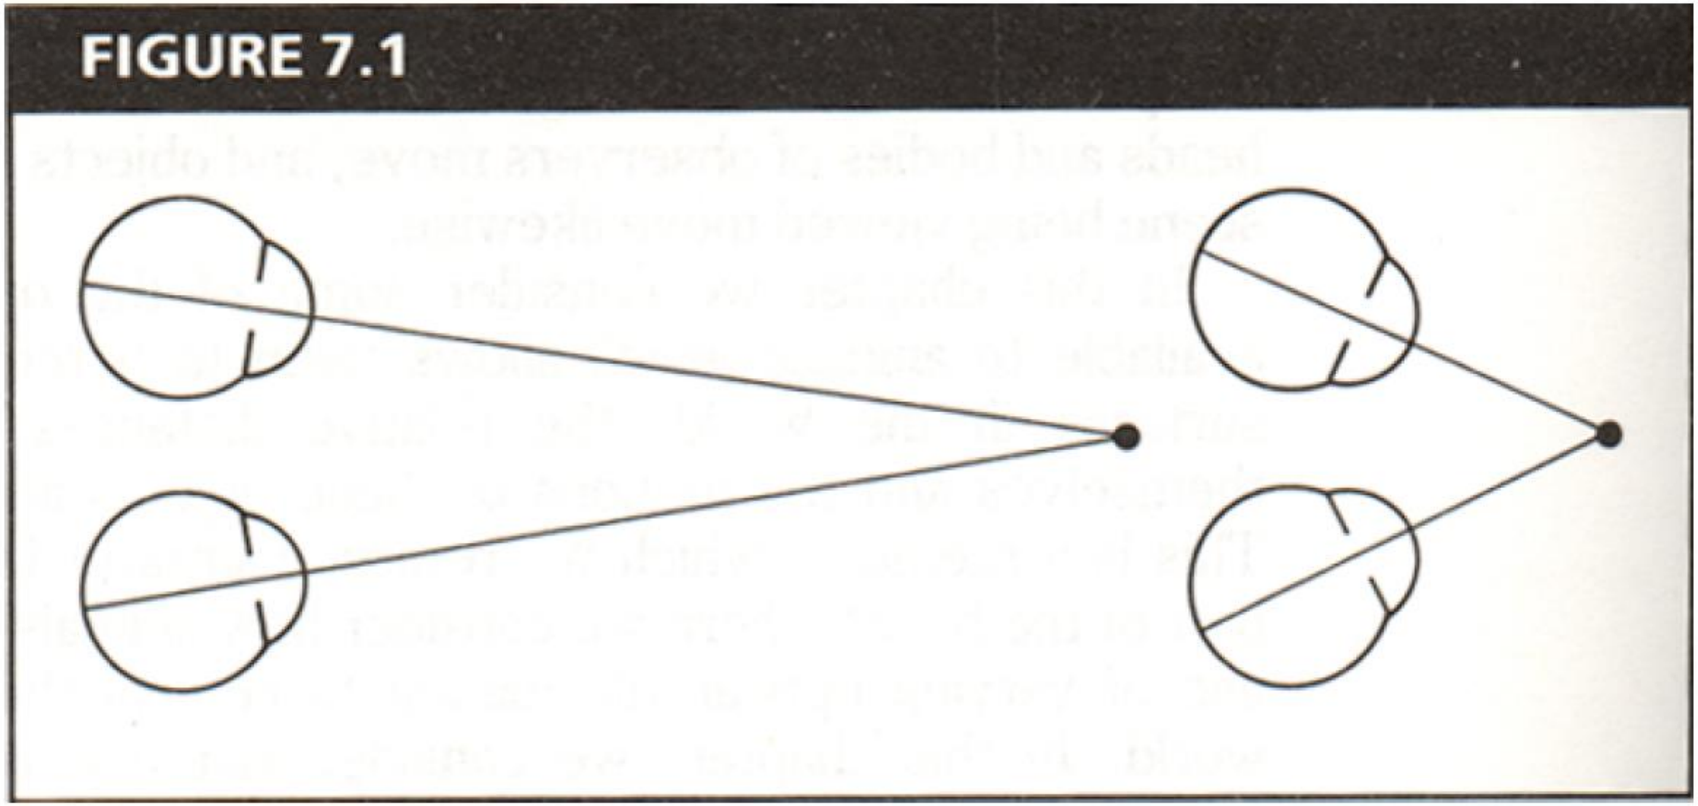

From Bruce and Green, Visual Perception, Physiology, Psychology and Ecology

Human eyes **fixate** on point in space – rotate so that corresponding images form in centers of fovea.

Human stereopsis: disparity

FIGURE 7.3

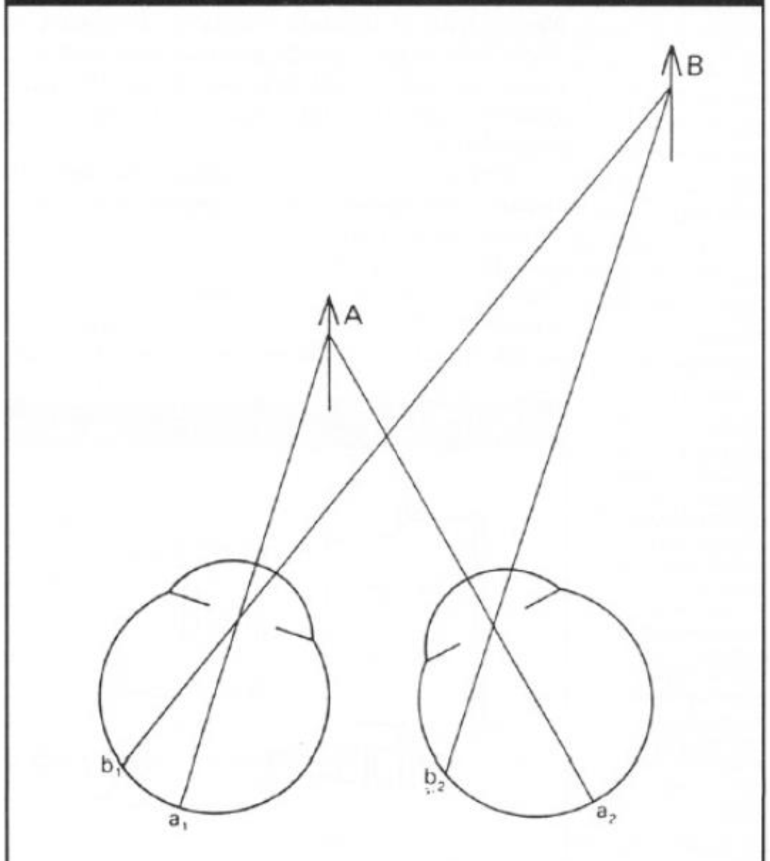

From Bruce and Green, Visual Perception, Physiology, Psychology and Ecology

Disparity occurs when eyes fixate on one object; others appear at different visual angles

Adapted from David Forsyth, UC Berkeley

Random dot stereograms

- Julesz 1960: Do we identify local brightness patterns before fusion (monocular process) or after (binocular)?

- To test: pair of synthetic images obtained by randomly spraying black dots on white objects

Random dot stereograms

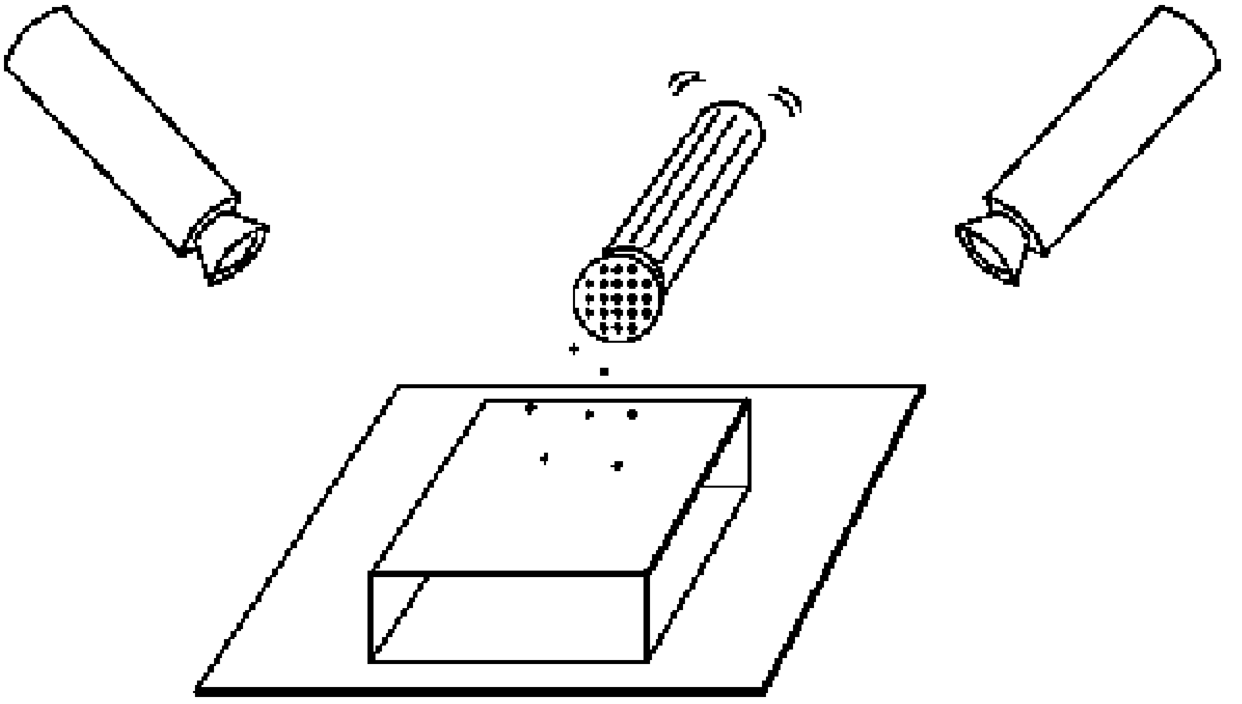

Random dot stereograms

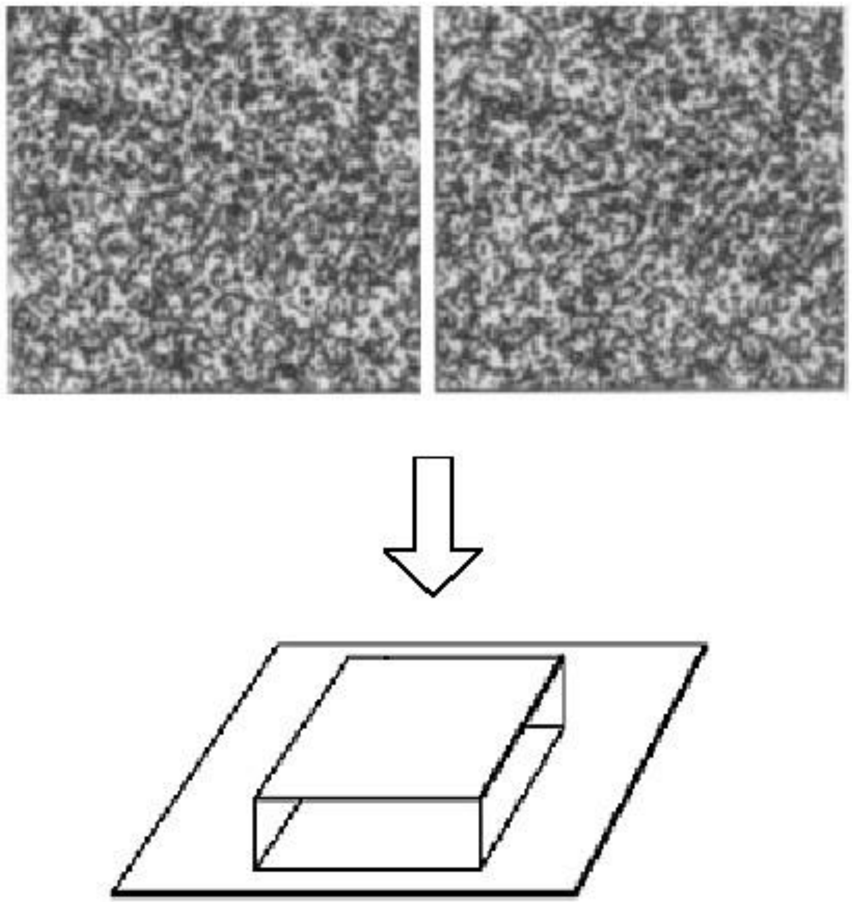

Random dot stereograms

- When viewed monocularly, they appear random; when viewed stereoscopically, see 3d structure.

- Conclusion: human binocular fusion not directly associated with the physical retinas; must involve the central nervous system
- Imaginary "*cyclopean retina*" that combines the left and right image stimuli as a single unit
- High level scene understanding not required for Stereo

Stereo photography and stereo viewers

Take two pictures of the same subject from two slightly different viewpoints and display so that each eye sees only one of the images.

Invented by Sir Charles Wheatstone, 1838

Image from fisher-price.com

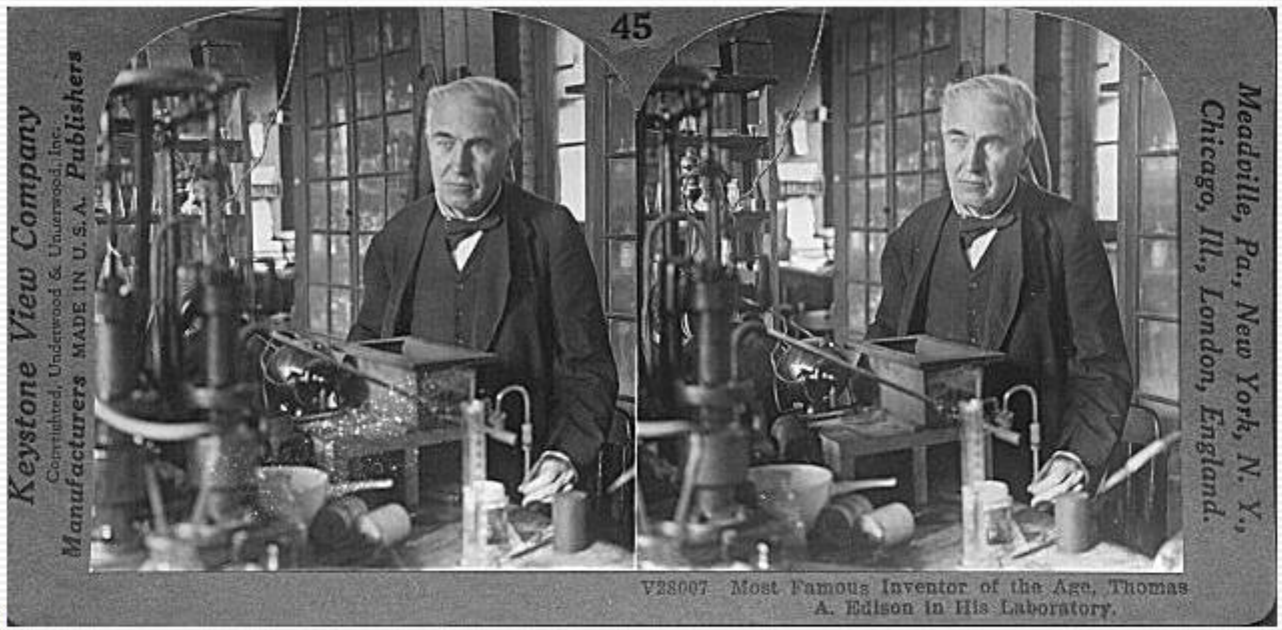

http://www.johnsonshawmuseum.org

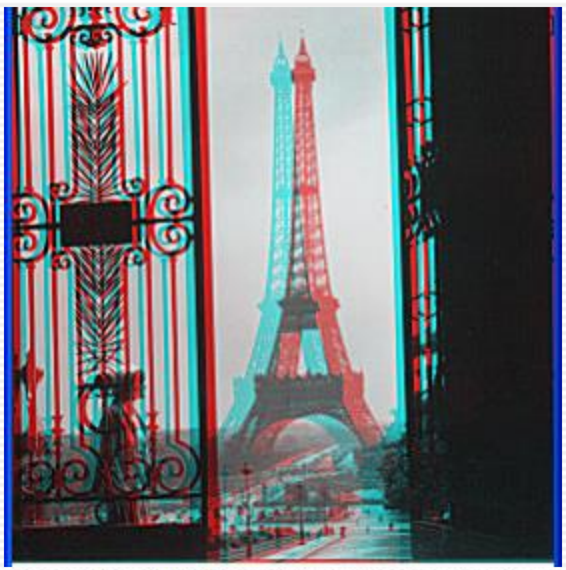

http://www.johnsonshawmuseum.org

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

http://www.well.com/~jimg/stereo/stereo_list.html

Autostereograms

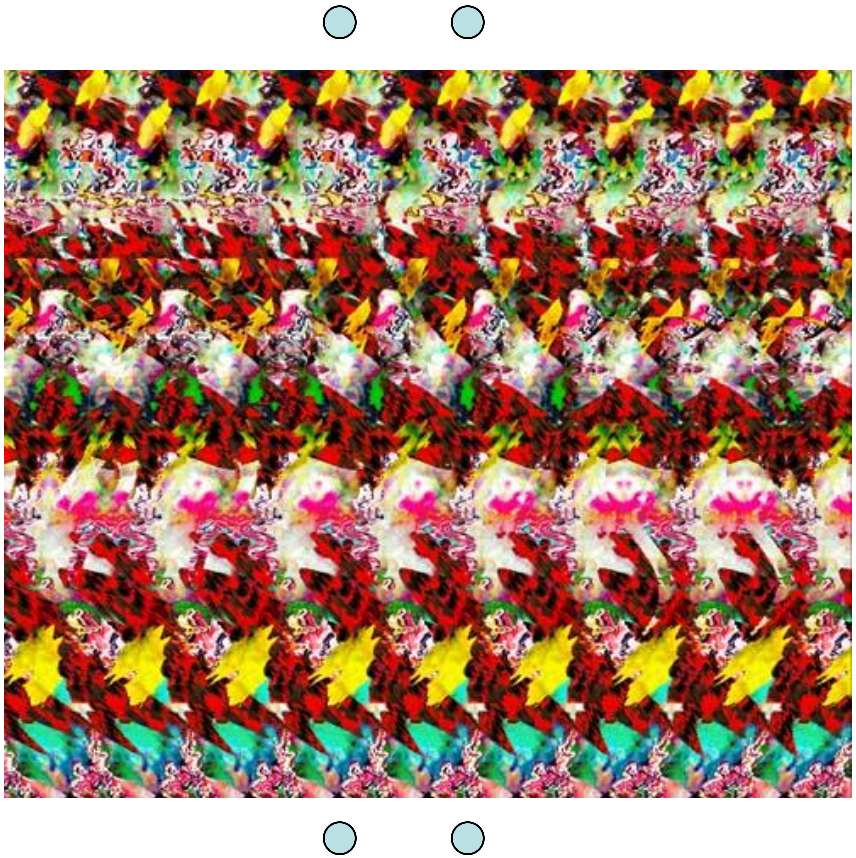

Exploit disparity as depth cue using single image.

(Single image random dot stereogram, Single image stereogram)

Images from magiceye.com

Autostereograms

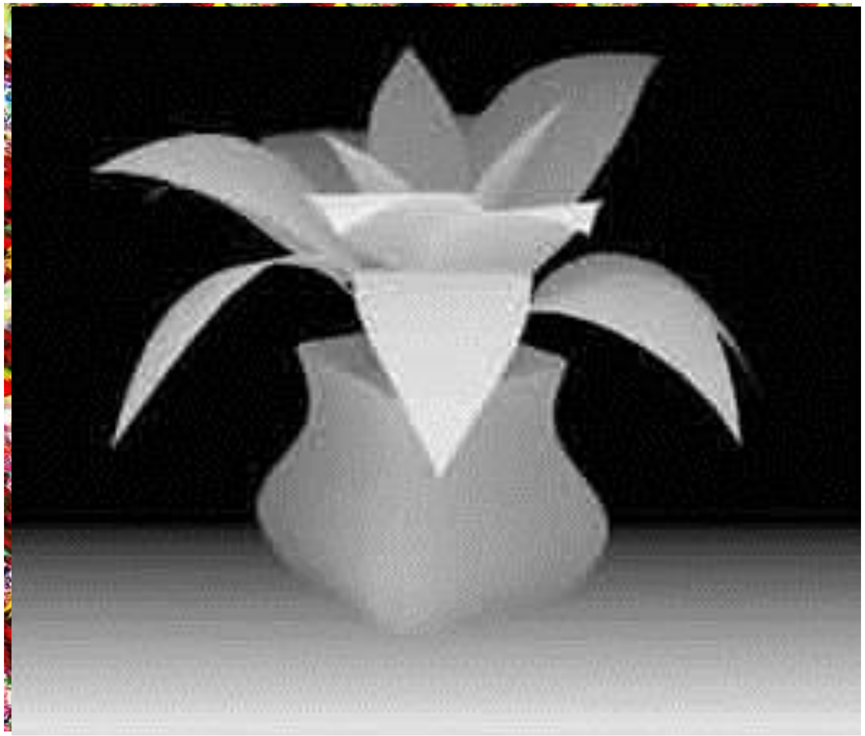

Images from magiceye.com

Estimating depth with stereo

- **Stereo**: shape from "motion" between two views
- We'll need to consider:
  - Info on camera pose ("calibration")
  - Image point correspondences

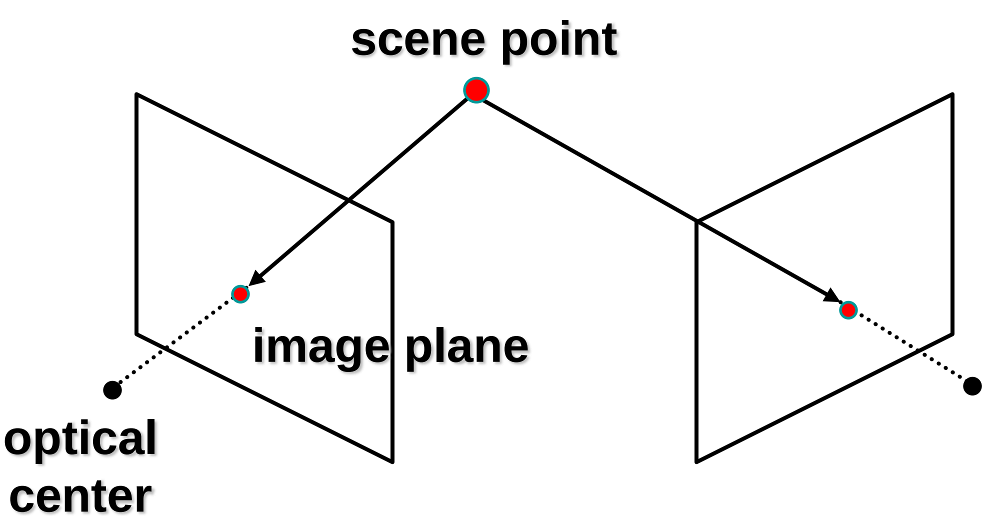

Stereo vision

Two cameras, simultaneous views

Single moving camera and static scene

Camera parameters

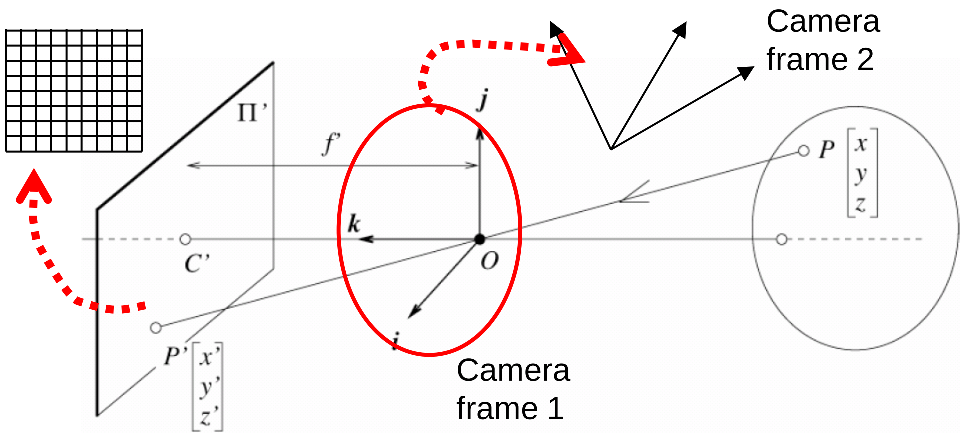

Extrinsic parameters:
Camera frame 1 ←→ Camera frame 2

Intrinsic parameters:
Image coordinates relative to camera ←→ Pixel coordinates

- *Extrinsic* params: rotation matrix and translation vector
- *Intrinsic* params: focal length, pixel sizes (mm), image center point, radial distortion parameters

We'll assume for now that these parameters are given and fixed.

Outline

- Human stereopsis
- Stereograms
- Epipolar geometry and the epipolar constraint
  - Case example with parallel optical axes
  - General case with calibrated cameras

Geometry for a simple stereo system

- First, assuming parallel optical axes, known camera parameters (i.e., calibrated cameras):

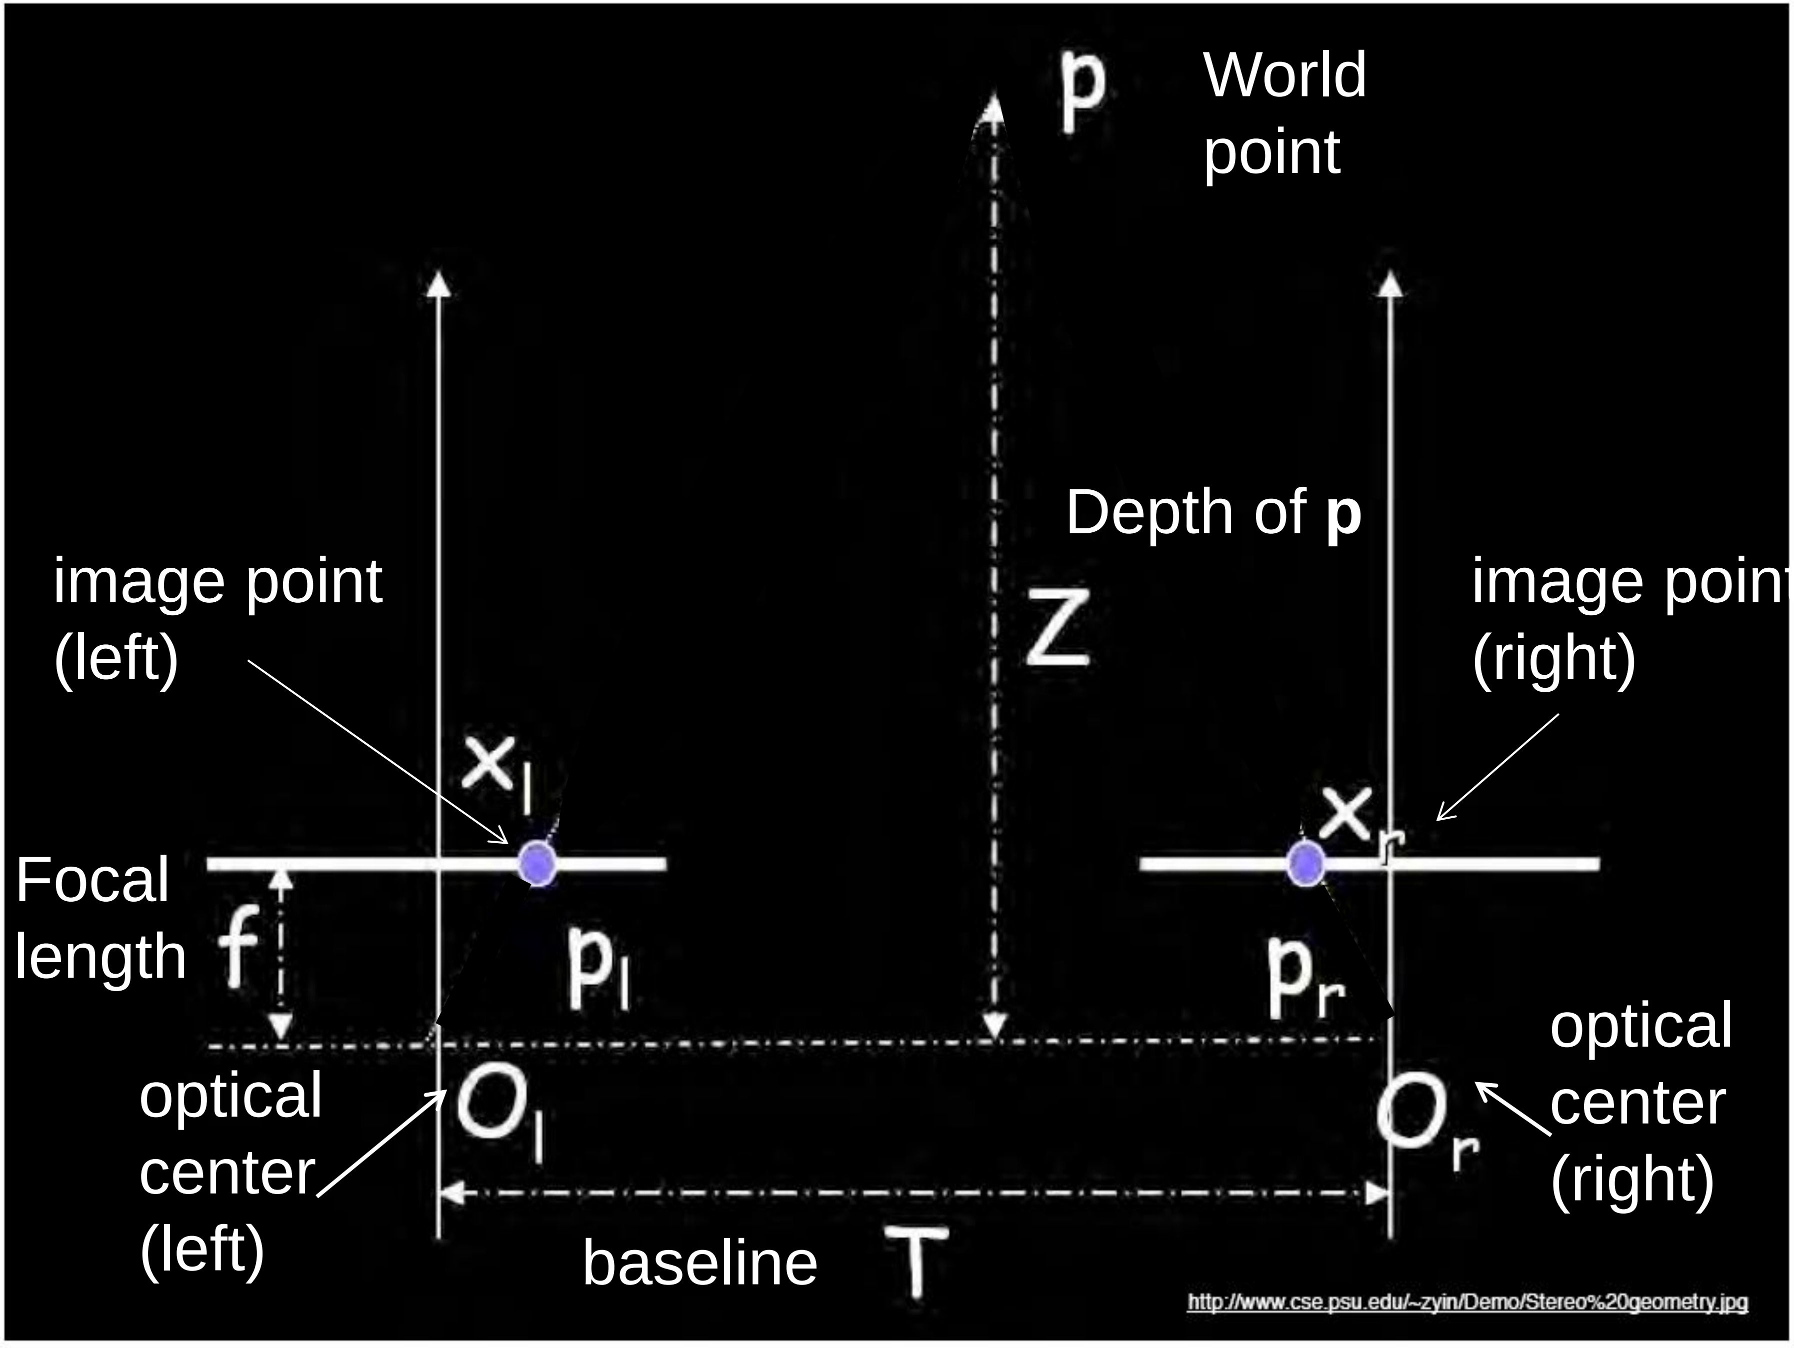
p
World
point
Depth of **p**
Z
image point
(left)
image point
(right)
x_l
x_r
Focal
length
f
p_l
p_r
optical
center
(left)
O_l
O_r
optical
center
(right)
baseline
T
http://www.cse.psu.edu/~zyin/Demo/Stereo%20geometry.jpg

Geometry for a simple stereo system

- Assume parallel optical axes, known camera parameters (i.e., calibrated cameras). **What is expression for Z?**

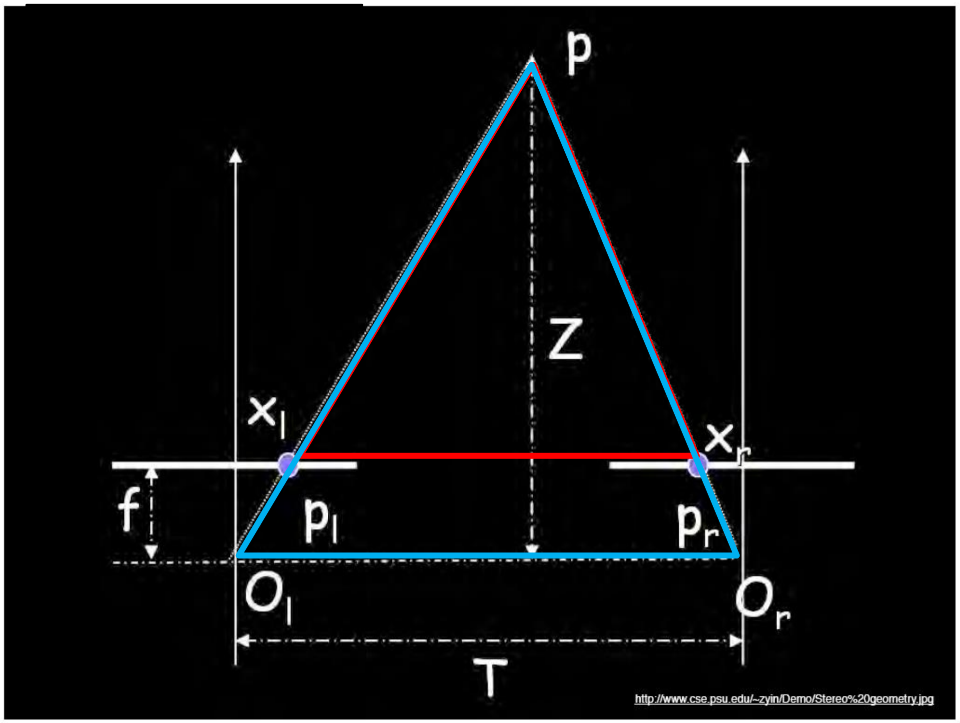

Similar triangles (p_l, P, p_r) and (O_l, P, O_r):

$$\frac{T + x_l - x_r}{Z - f} = \frac{T}{Z}$$

$$Z = f \frac{T}{x_r - x_l}$$

disparity

Depth from disparity

image $I(x,y)$

Disparity map $D(x,y)$

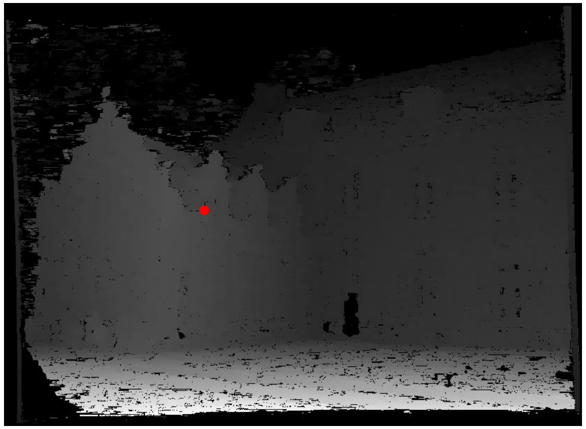

image $I'(x',y')$

$$(x',y')=(x+D(x,y), y)$$

So if we could find the **corresponding points** in two images, we could **estimate relative depth**…

Basic stereo matching algorithm

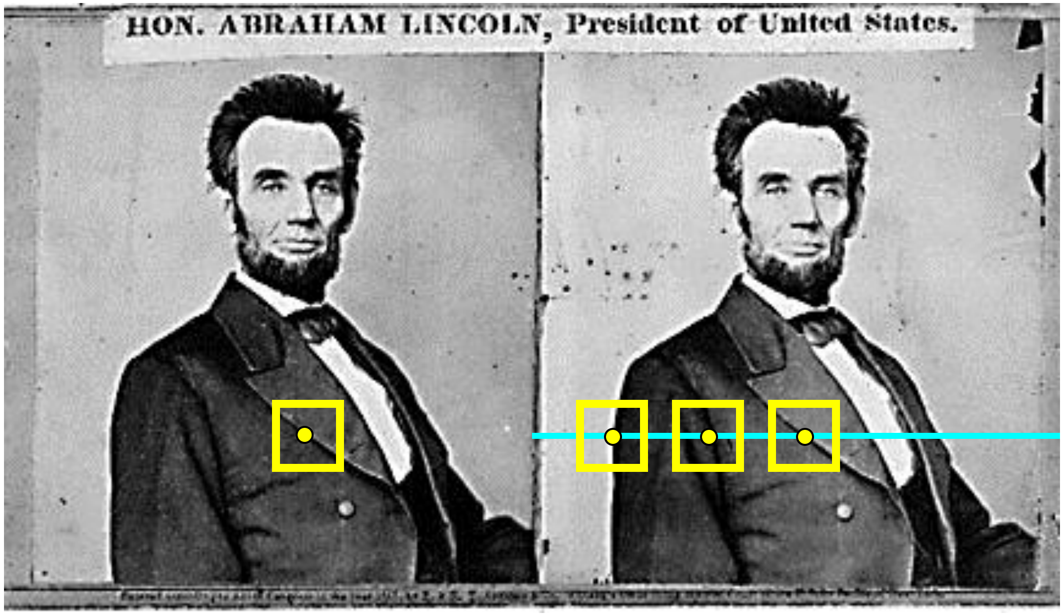

- If necessary, rectify the two stereo images to transform epipolar lines into scanlines
- For each pixel x in the first image
  - Find corresponding epipolar scanline in the right image
  - Examine all pixels on the scanline and pick the best match x’
  - Compute disparity x-x’ and set depth(x) = fB/(x-x’)

Correspondence search

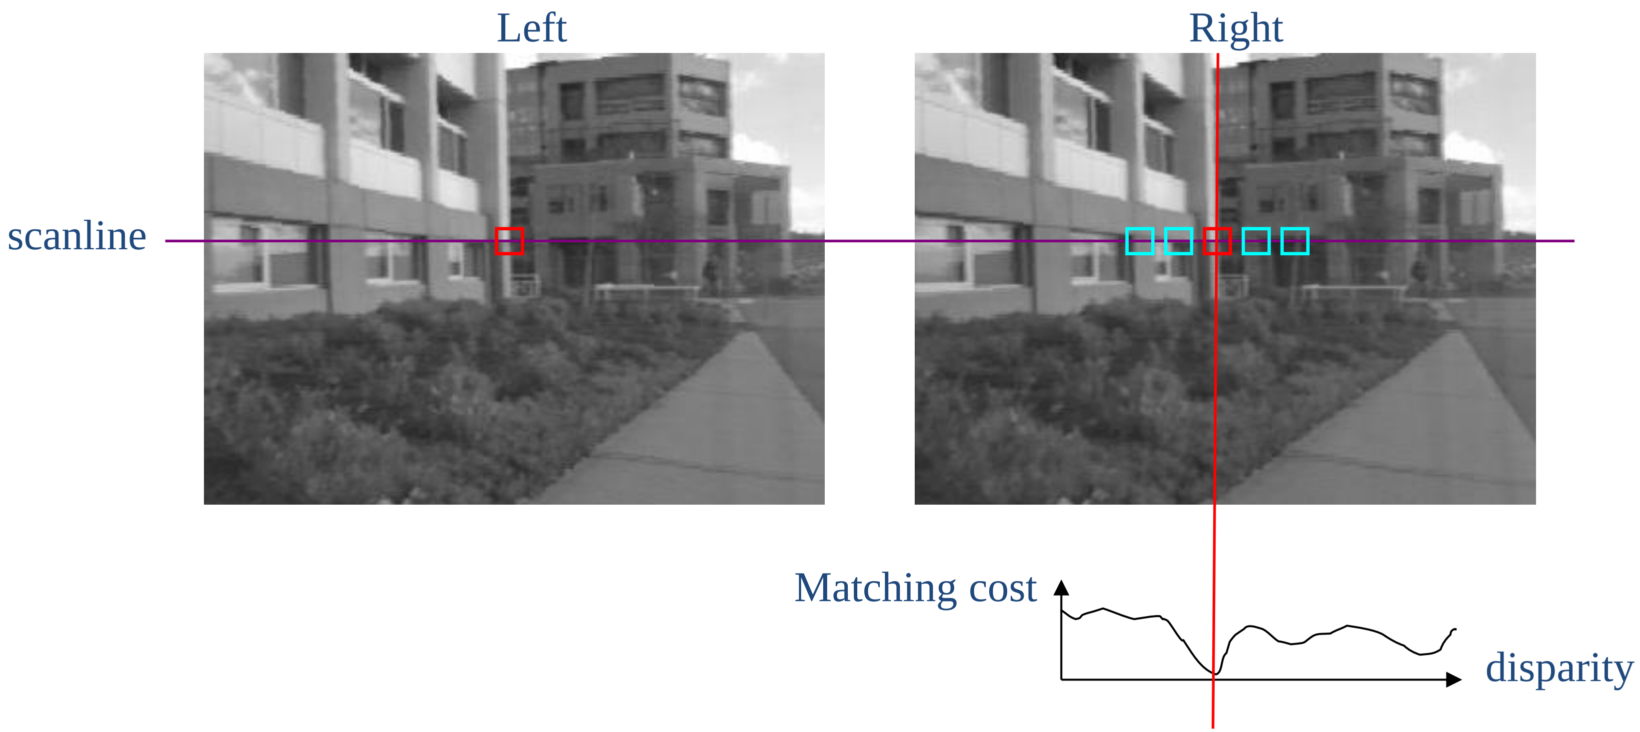

- Slide a window along the right scanline and compare contents of that window with the reference window in the left image
- Matching cost: SSD or normalized correlation

Correspondence search

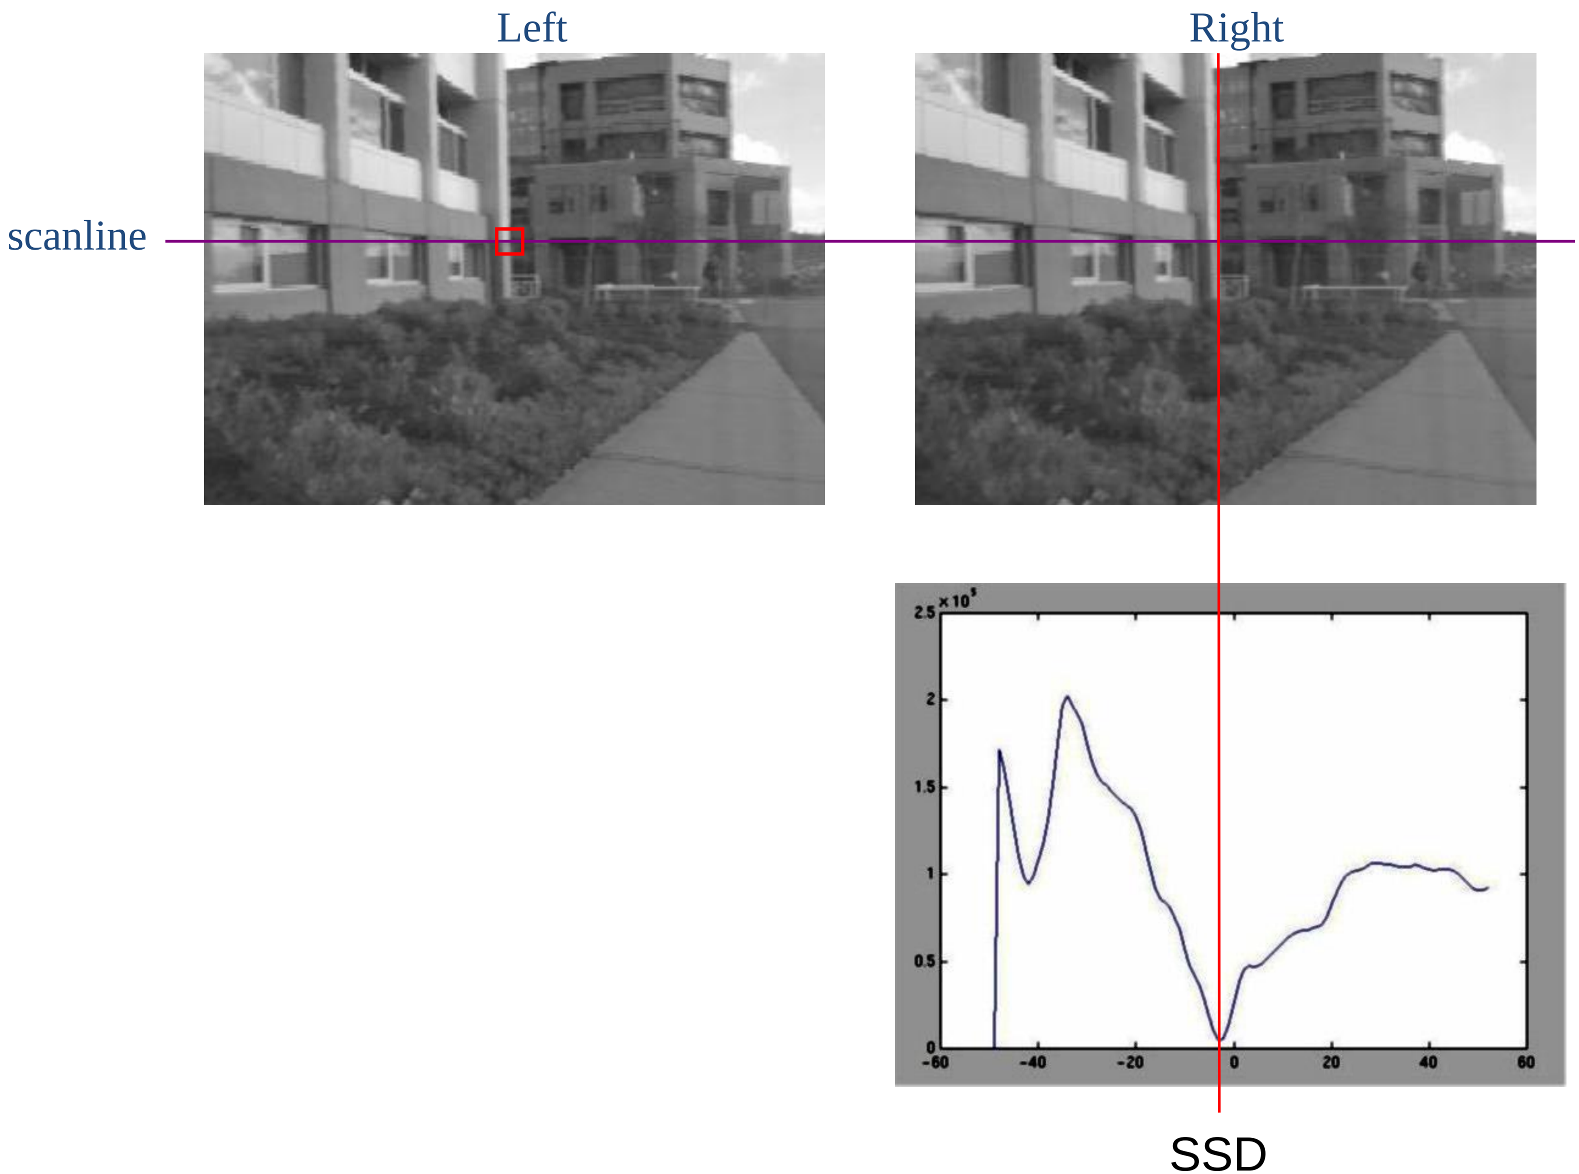

Correspondence search

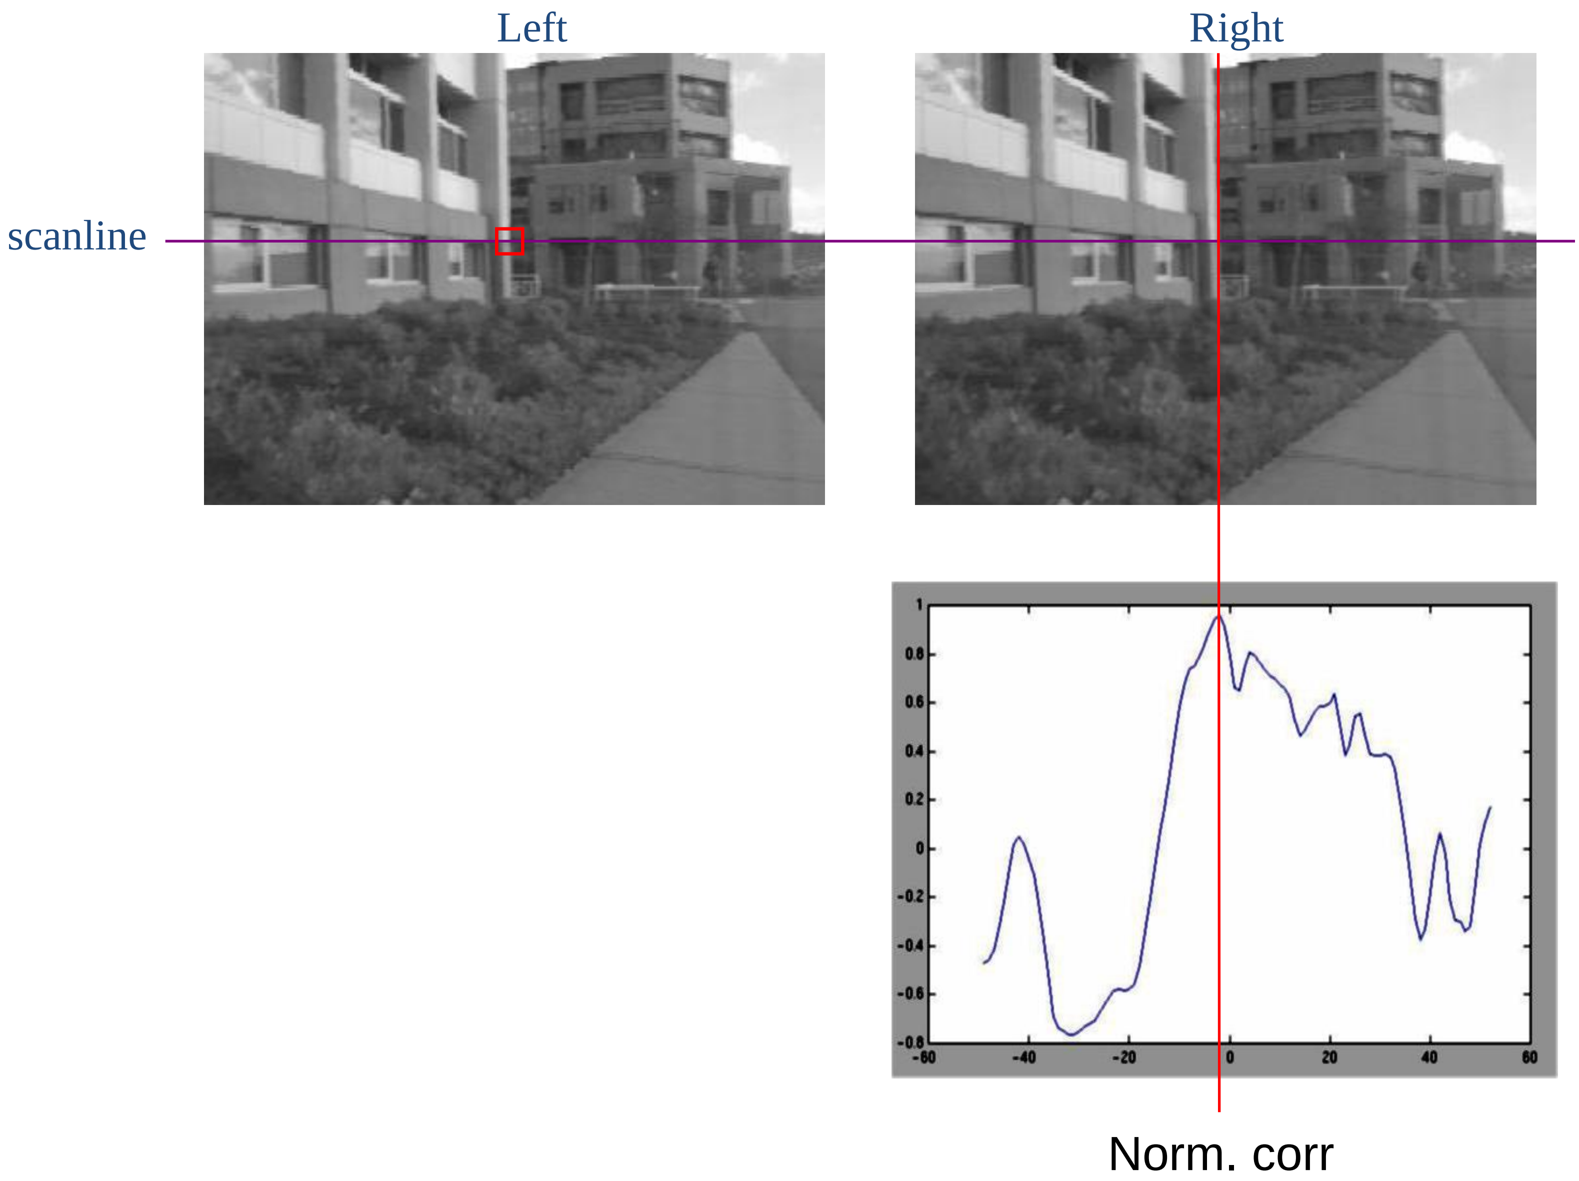

Effect of window size

W = 3

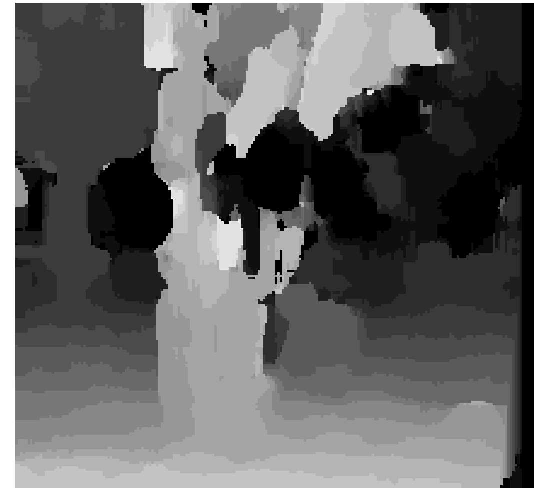

W = 20

- Smaller window
  - \+ More detail
  - – More noise
- Larger window
  - \+ Smoother disparity maps
  - – Less detail

Failures of correspondence search

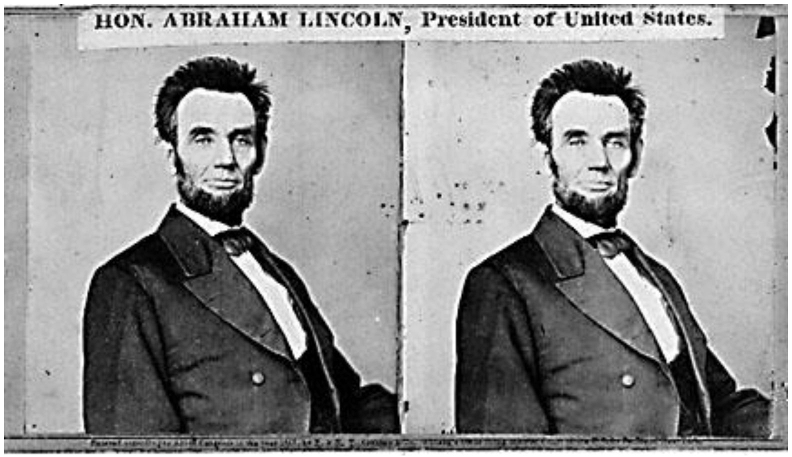

Textureless surfaces

Occlusions, repetition

Non-Lambertian surfaces, specularities

Results with window search

Data

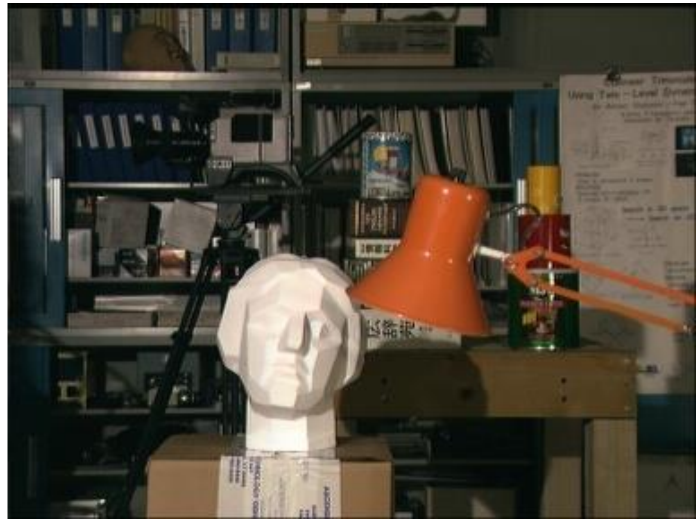

Window-based matching

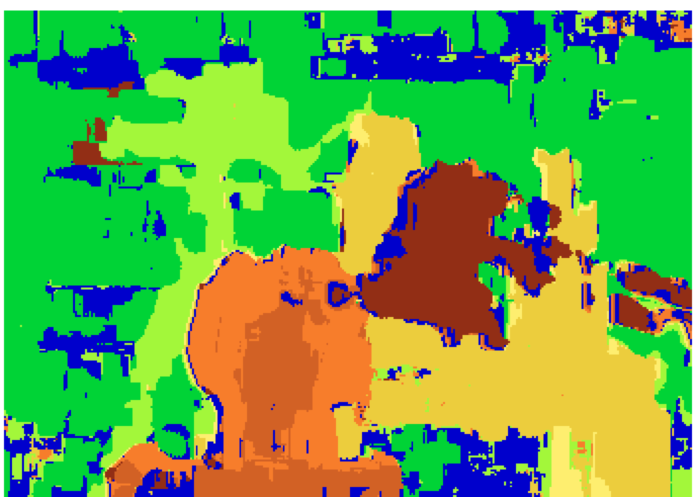

Ground truth

How can we improve window-based matching?

- So far, matches are independent for each point

- What constraints or priors can we add?

Summary

- Depth from stereo: main idea is to triangulate from corresponding image points.
- Epipolar geometry defined by two cameras
  - We've assumed known extrinsic parameters relating their poses
- Epipolar constraint limits where points from one view will be imaged in the other
  - Makes search for correspondences quicker

- **Terms**: epipole, epipolar plane / lines, disparity, rectification, intrinsic/extrinsic parameters, essential matrix, baseline

Coming up

- Stereo Algorithms
- Structure from Motion