CS4670: Computer Vision

Noah Snavely

Lecture 10: Cameras

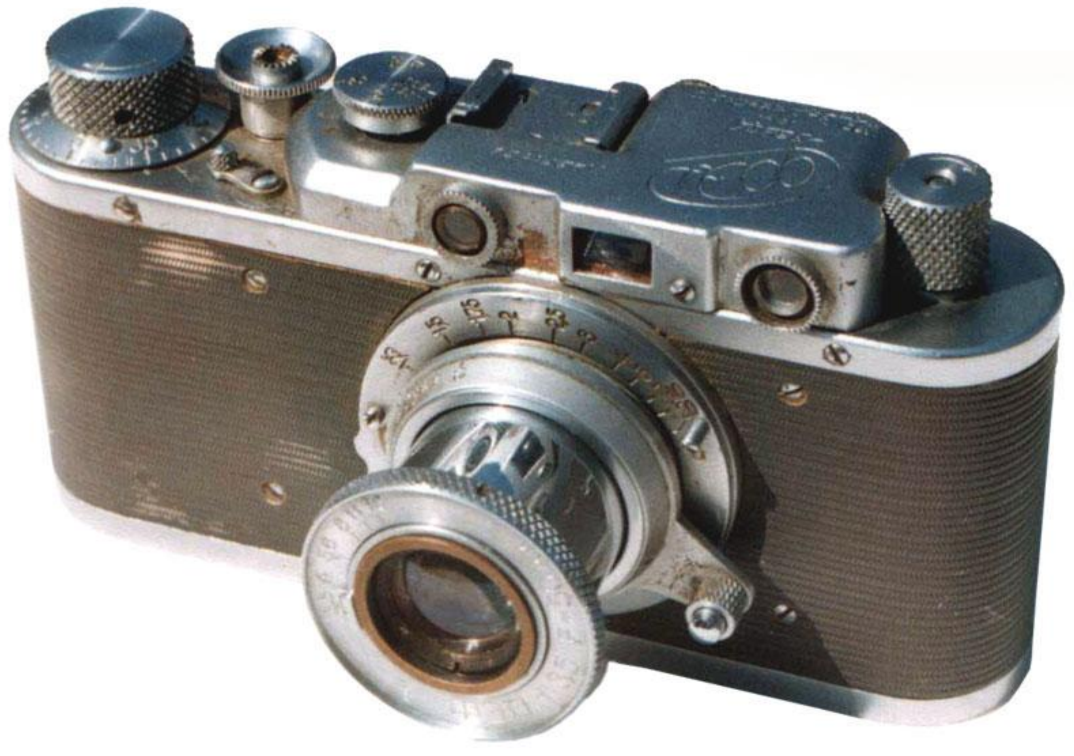

Source: S. Lazebnik

Reading

- Szeliski 2.1.3-2.1.6

Announcements

- Project 2a: extended to Wednesday, October 6, at 8:59pm
  - Dominant orientation: you can either compute the maximum eigenvector, or the smoothed gradient (as in the MOPs paper)
  - Scale-space detection not required
  - You can use any OpenCV routine that duplicates what you did for P1 (cvSmooth, cvSobel, cvFilter)
  - Black Harris images: you may need to multiply the image by some large constant to see the Harris values

RANSAC

- Back to linear regression
- How do we generate a hypothesis?

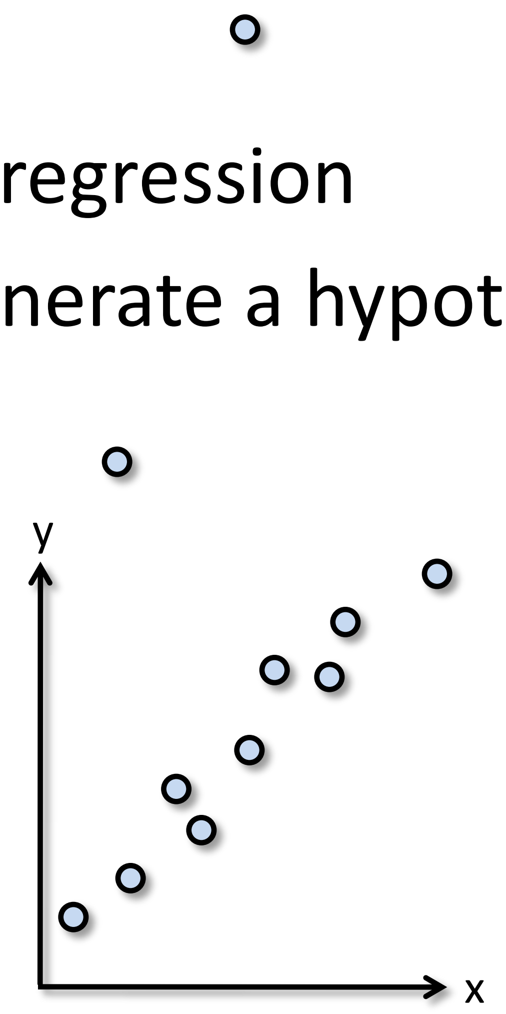

RANSAC

- Back to linear regression
- How do we generate a hypothesis?

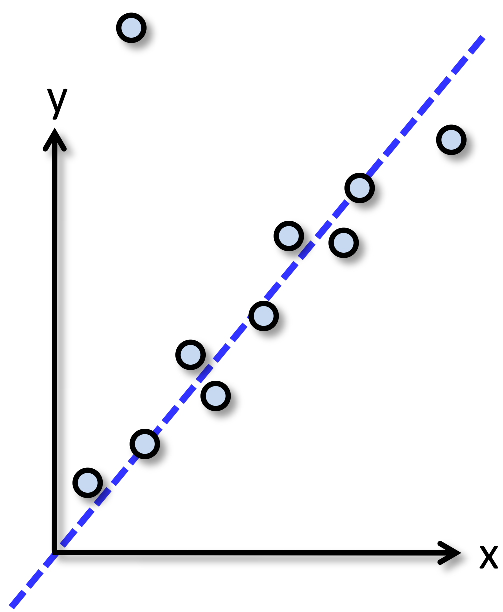

RANSAC

- General version:
  1. Randomly choose *s* samples
     - Typically *s* = minimum sample size that lets you fit a model
  2. Fit a model (e.g., line) to those samples
  3. Count the number of inliers that approximately fit the model
  4. Repeat *N* times
  5. Choose the model that has the largest set of inliers

How many rounds?

- If we have to choose *s* samples each time
  - with an outlier ratio *e*
  - and we want the right answer with probability *p*

	proportion of outliers *e*						
s	5%	10%	20%	25%	30%	40%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177

$p = 0.99$

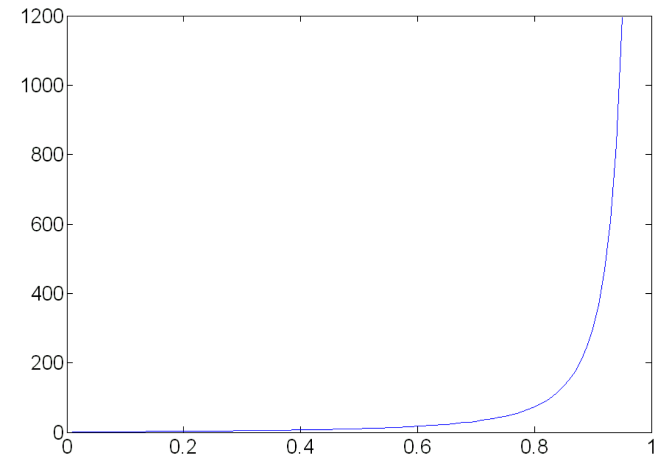

Source: M. Pollefeys

How big is *s*?

- For alignment, depends on the motion model
  - Here, each sample is a correspondence (pair of matching points)

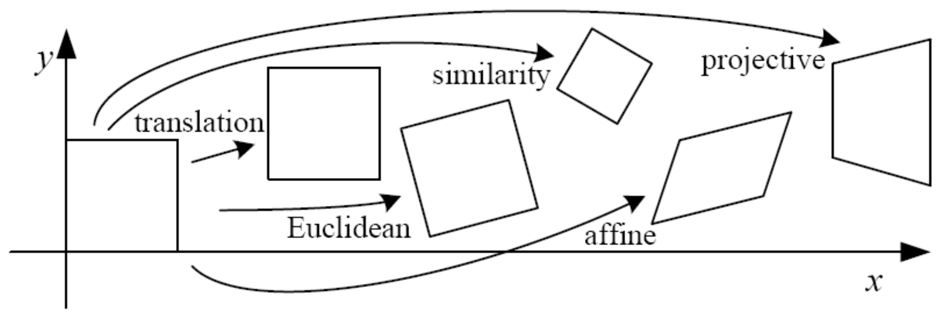

Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} I \mid t \end{bmatrix}_{2\times 3}$	2	orientation $+\cdots$	
rigid (Euclidean)	$\begin{bmatrix} R \mid t \end{bmatrix}_{2\times 3}$	3	lengths $+\cdots$	
similarity	$\begin{bmatrix} sR \mid t \end{bmatrix}_{2\times 3}$	4	angles $+\cdots$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

RANSAC pros and cons

- Pros
  - Simple and general
  - Applicable to many different problems
  - Often works well in practice
- Cons
  - Parameters to tune
  - Sometimes too many iterations are required
  - Can fail for extremely low inlier ratios
  - We can often do better than brute-force sampling

Final step: least squares fit

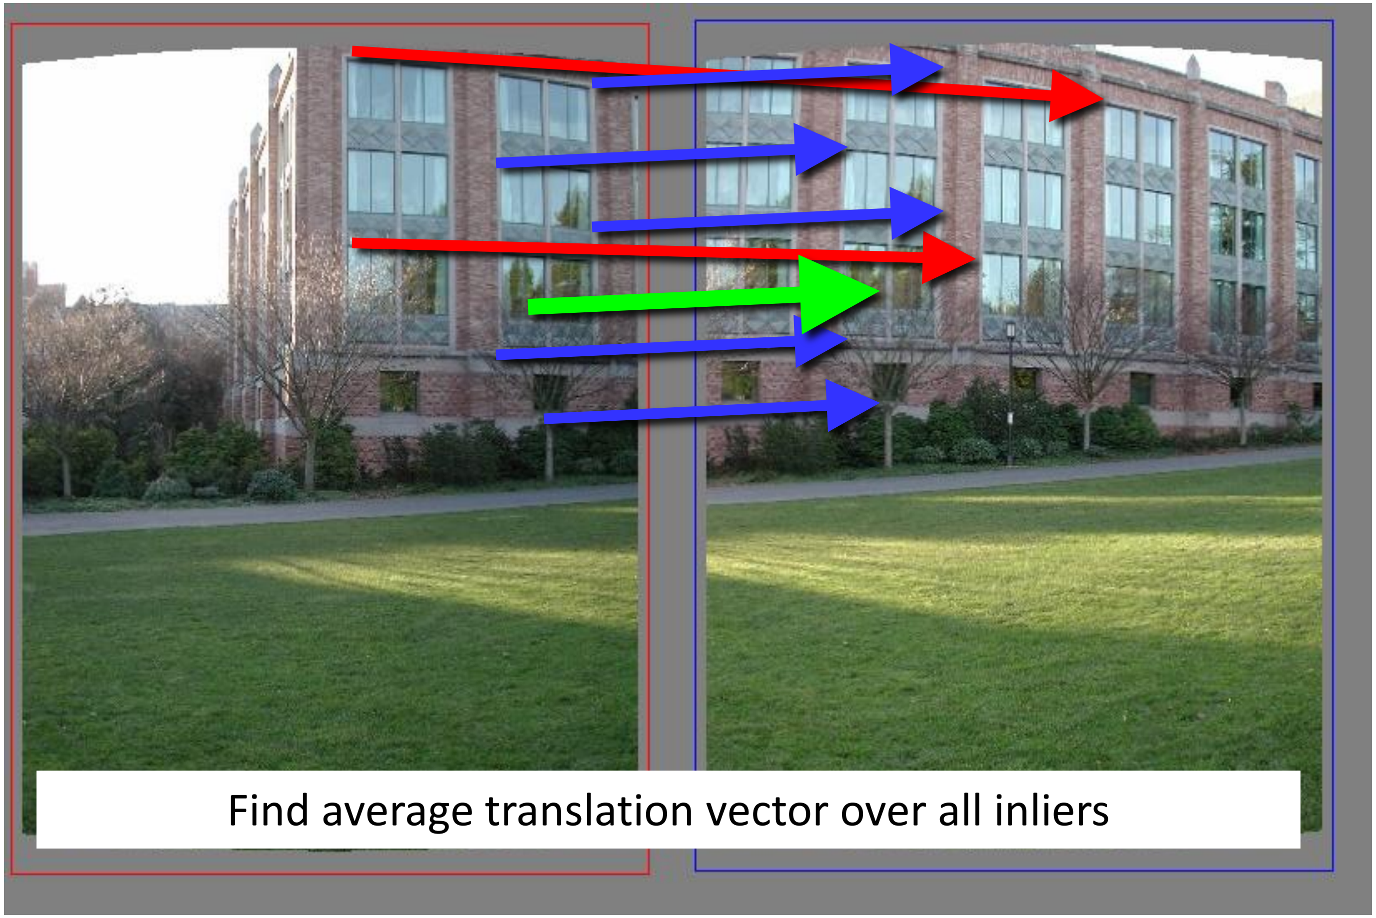

Find average translation vector over all inliers

RANSAC

- An example of a “voting”-based fitting scheme
- Each hypothesis gets voted on by each data point, best hypothesis wins

- There are many other types of voting schemes
  - E.g., Hough transforms…

Hough transform

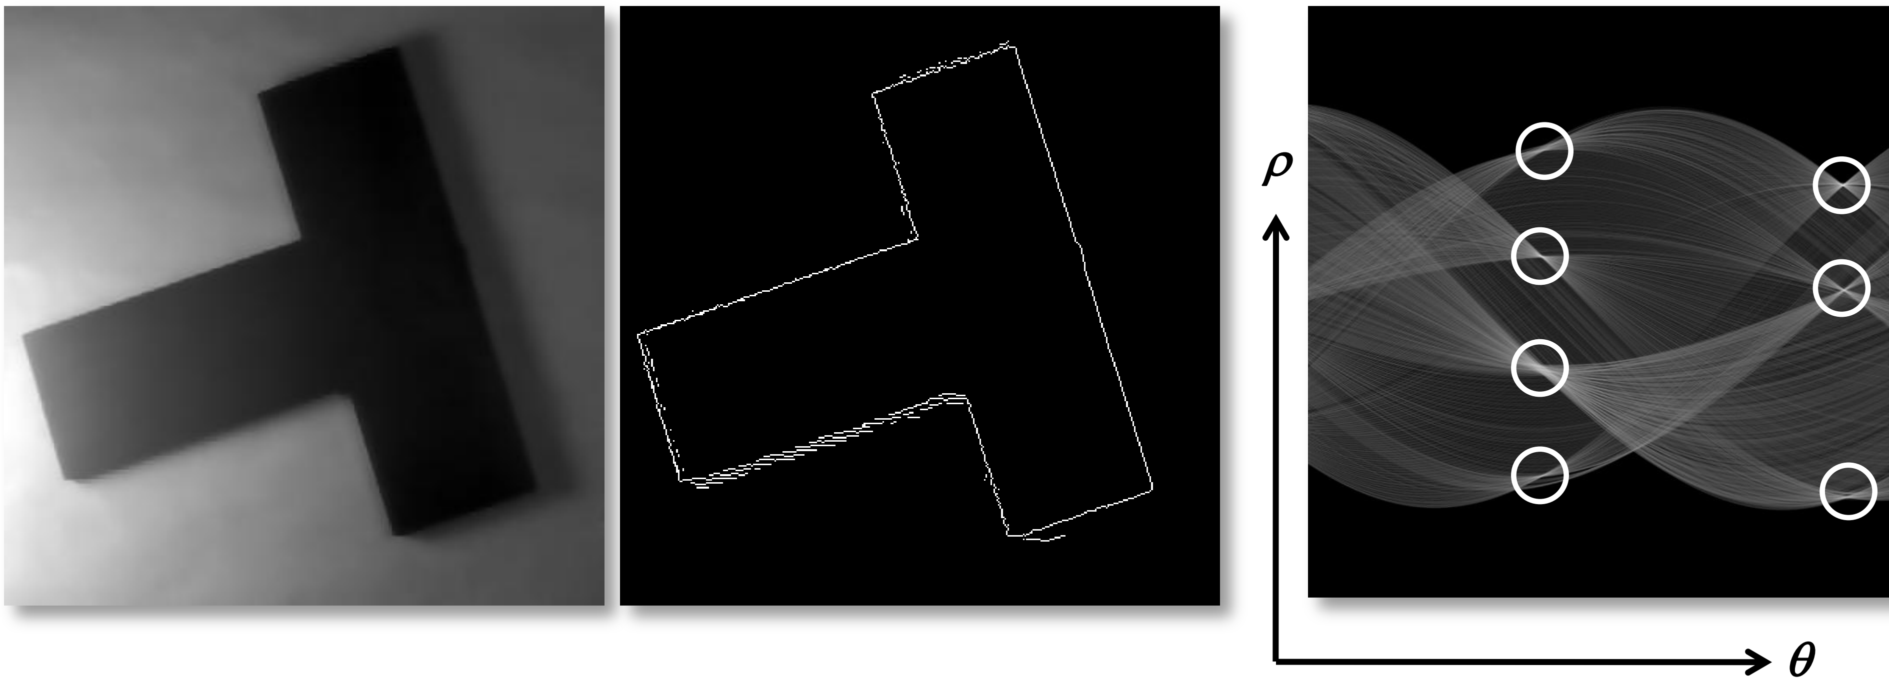

Panoramas

- Now we know how to create panoramas!
- Given two images:
  - Step 1: Detect features
  - Step 2: Match features
  - Step 3: Compute a homography using RANSAC
  - Step 4: Combine the images together (somehow)
- What if we have more than two images?

Can we use homographies to create a 360 panorama?

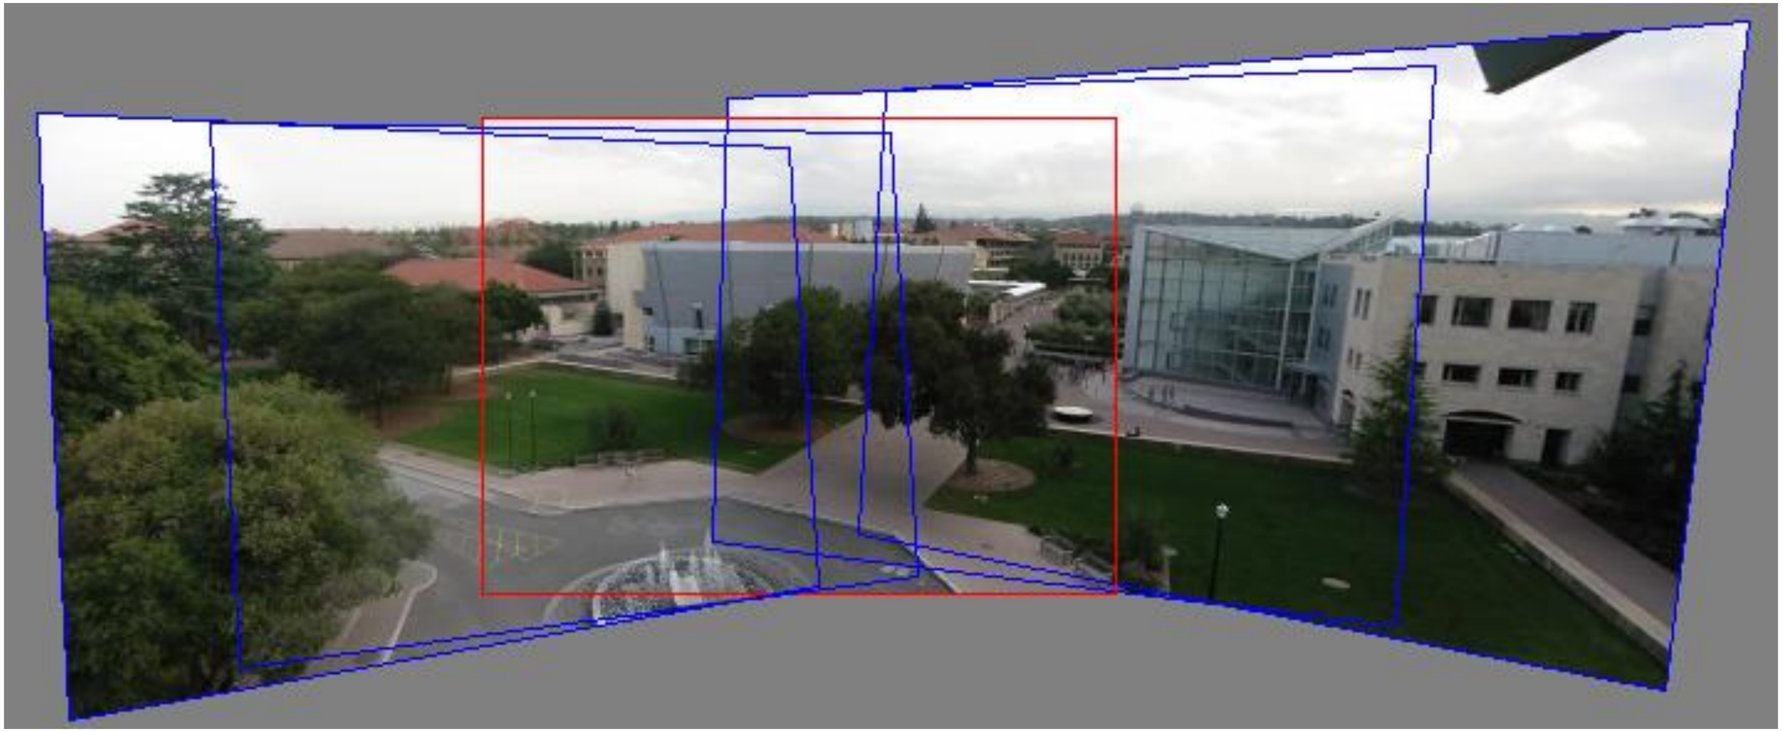

- In order to figure this out, we need to learn what a **camera** is

360 panorama

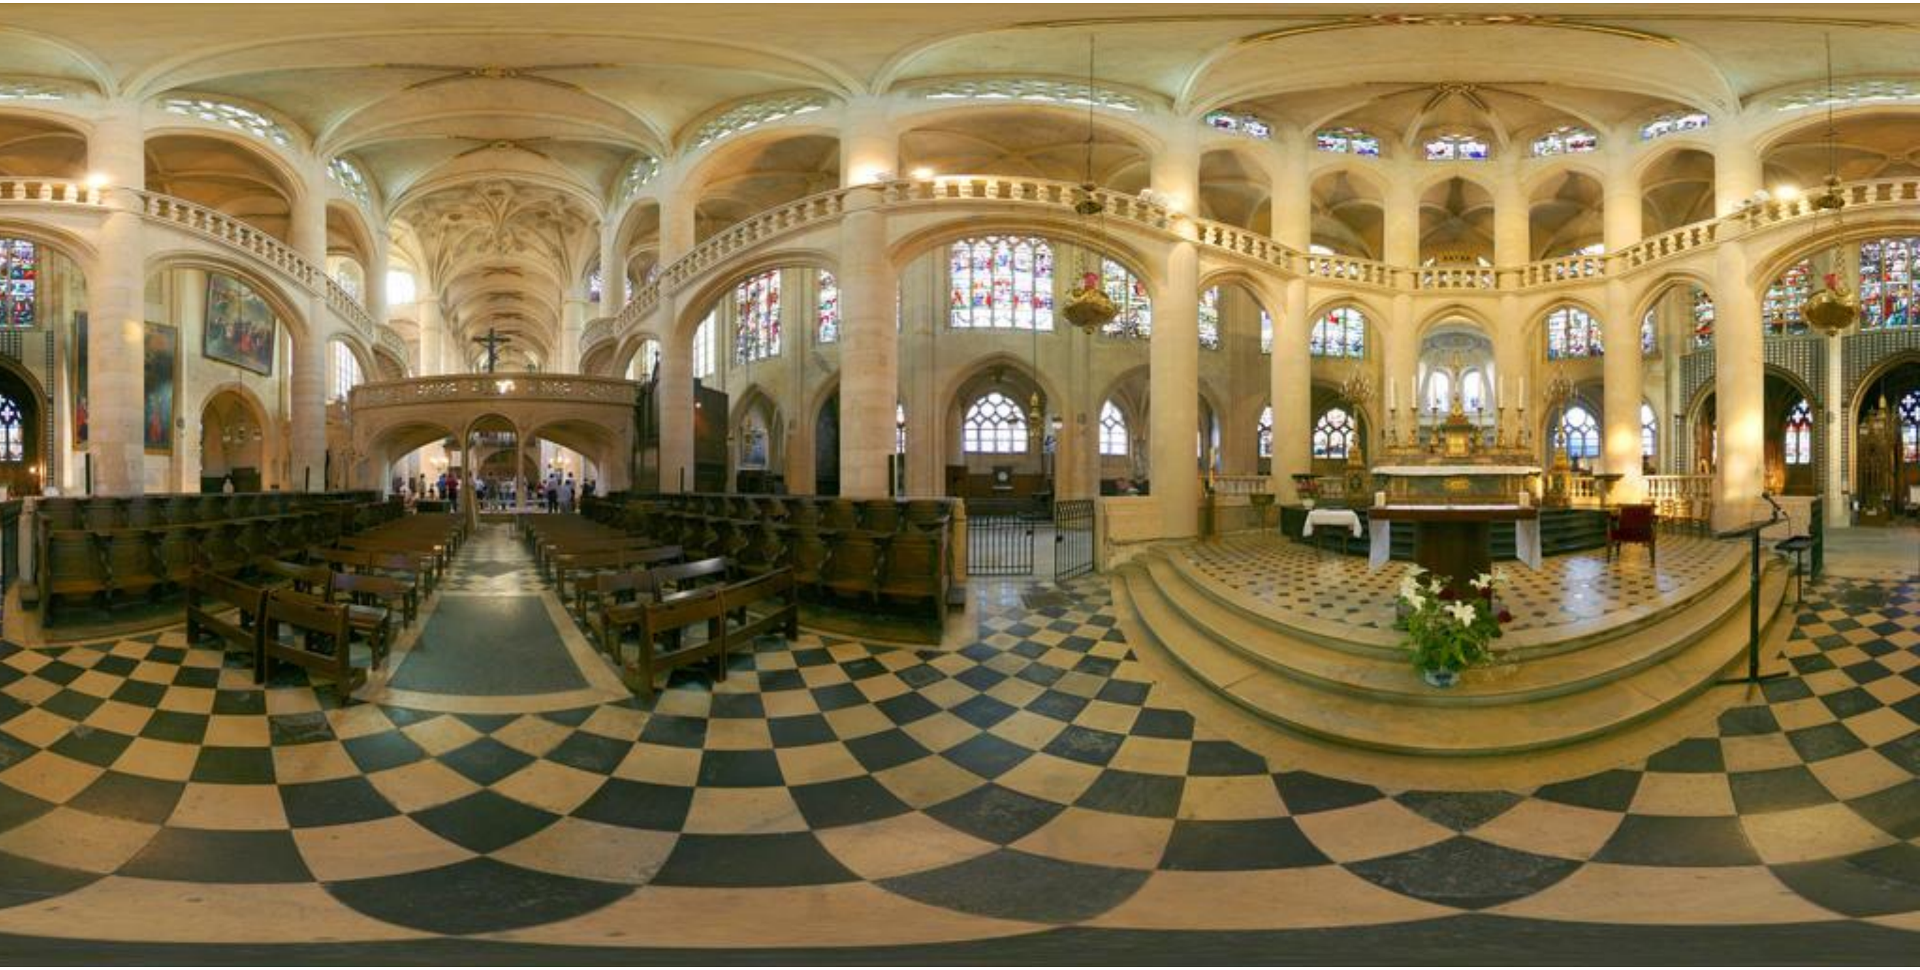

Questions?

Image formation

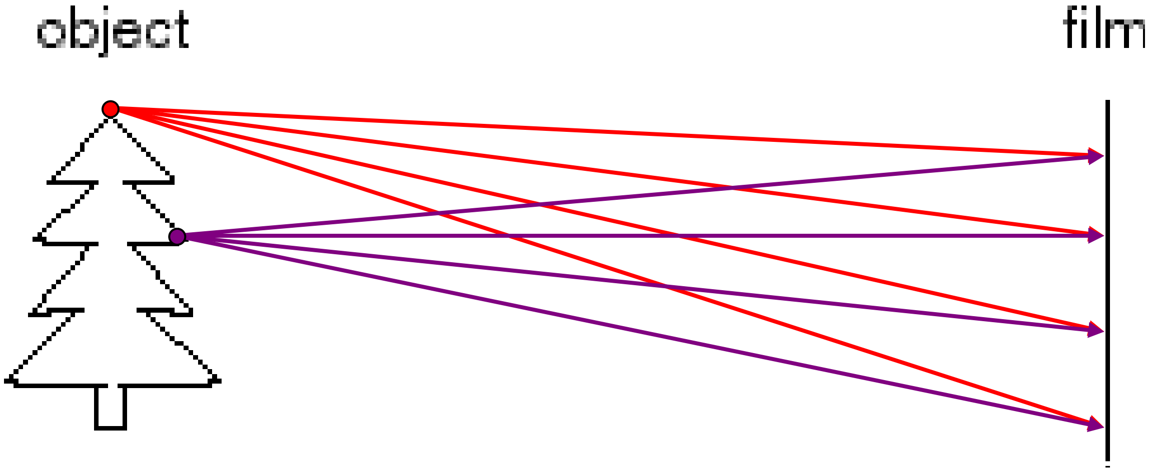

- Let’s design a camera
  - Idea 1: put a piece of film in front of an object
  - Do we get a reasonable image?

Pinhole camera

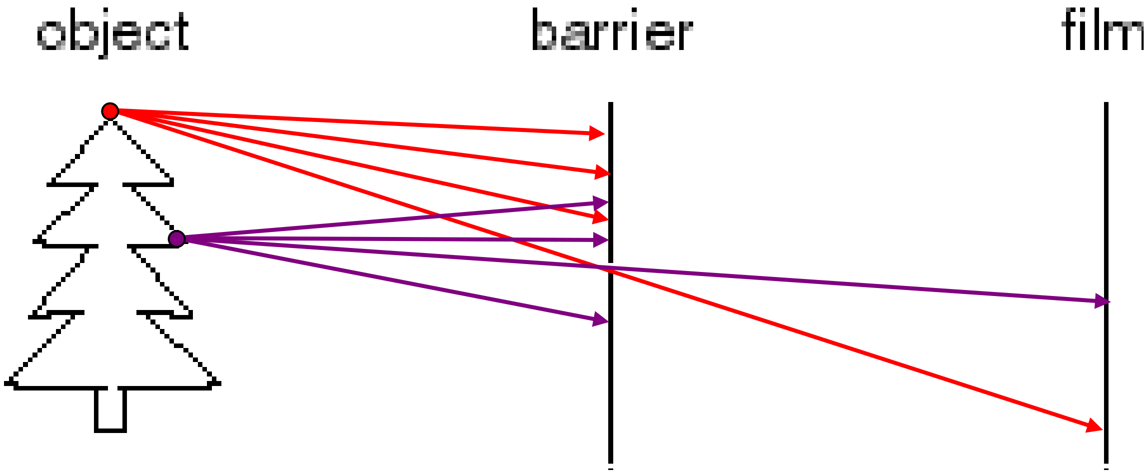

- Add a barrier to block off most of the rays
  - This reduces blurring
  - The opening known as the **aperture**
  - How does this transform the image?

Camera Obscura

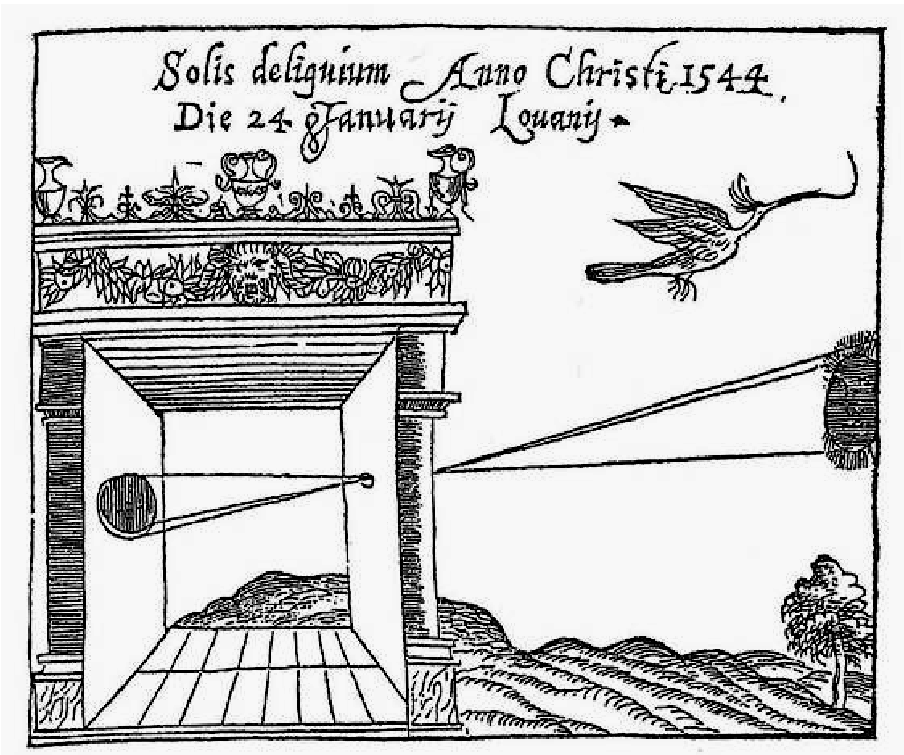

Gemma Frisius, 1558

- Basic principle known to Mozi (470-390 BC), Aristotle (384-322 BC)
- Drawing aid for artists: described by Leonardo da Vinci (1452-1519)

Source: A. Efros

Camera Obscura

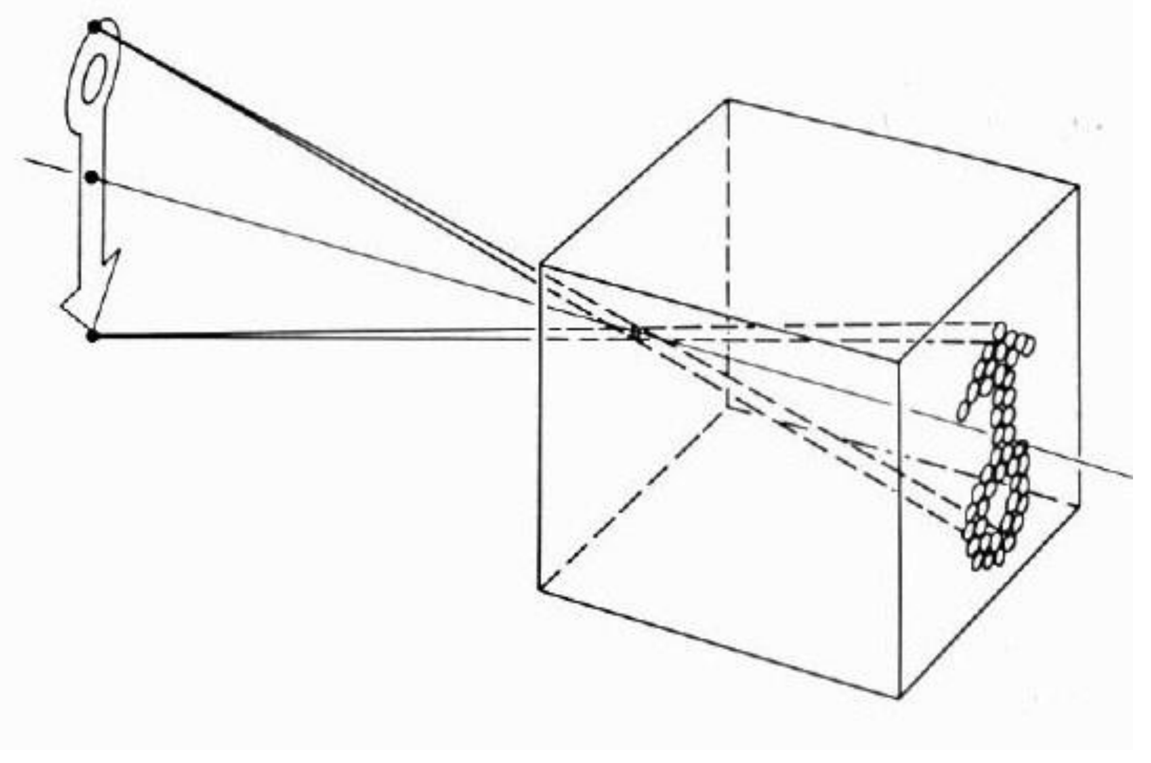

Home-made pinhole camera

Why so blurry?

http://www.debevec.org/Pinhole/

Slide by A. Efros

Pinhole photography

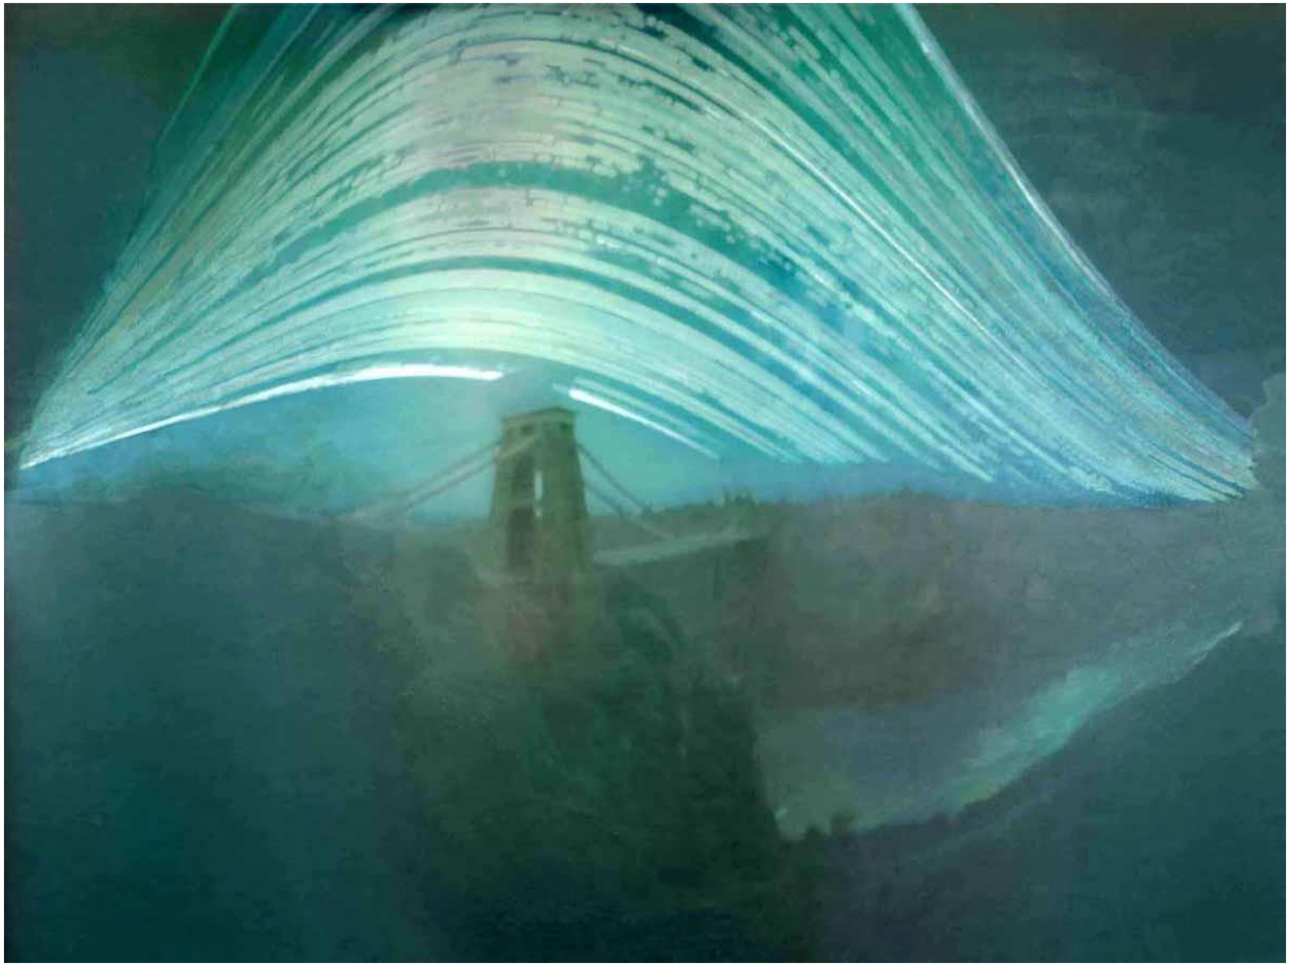

Justin Quinnell, The Clifton Suspension Bridge. December 17th 2007 - June 21st 2008

***6-month* exposure**

Shrinking the aperture

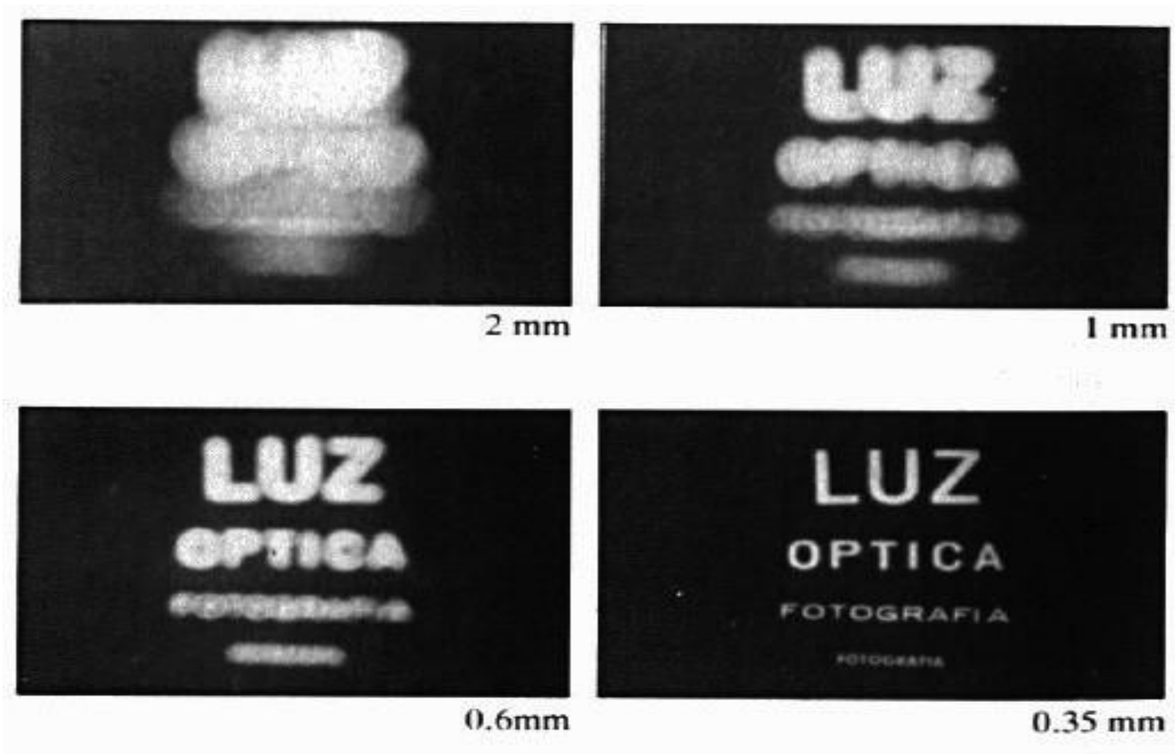

- Why not make the aperture as small as possible?
  - Less light gets through
  - *Diffraction* effects...

Shrinking the aperture

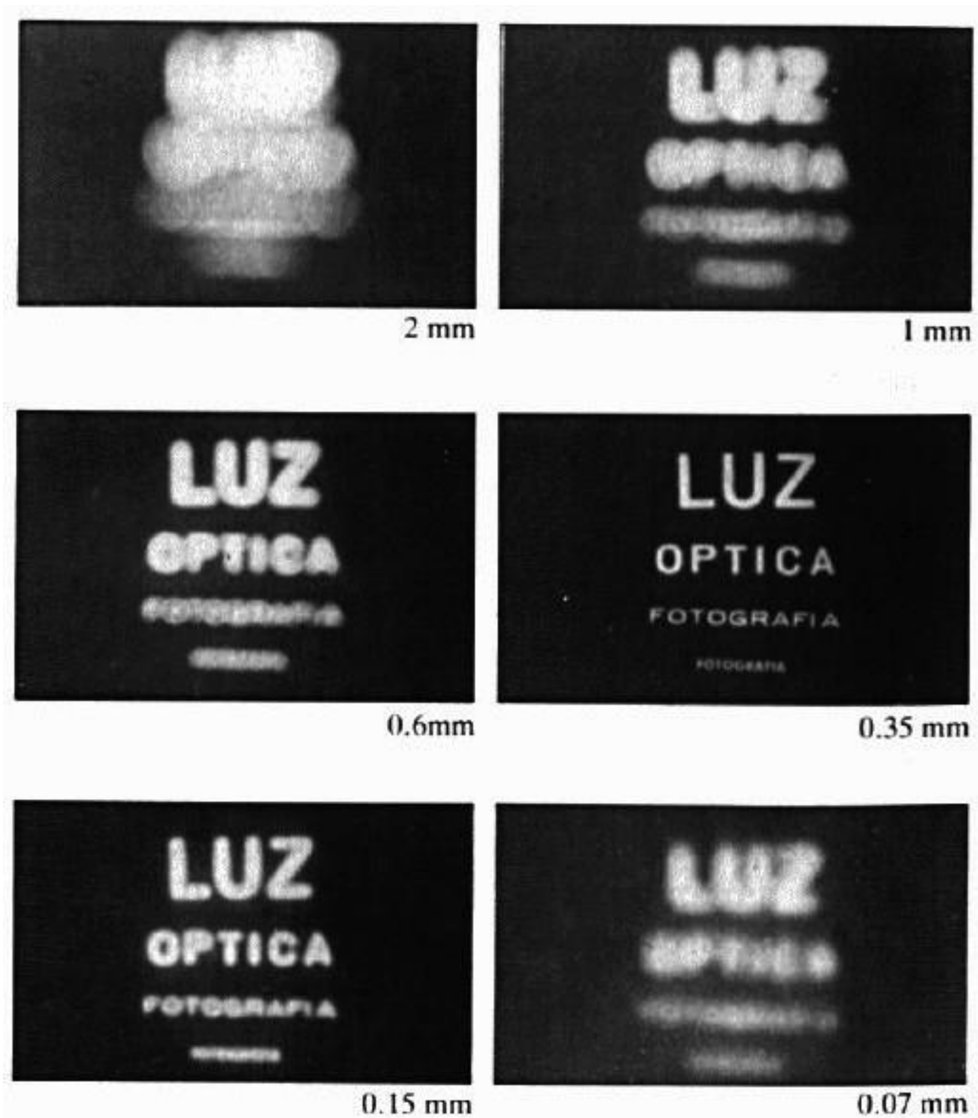

Adding a lens

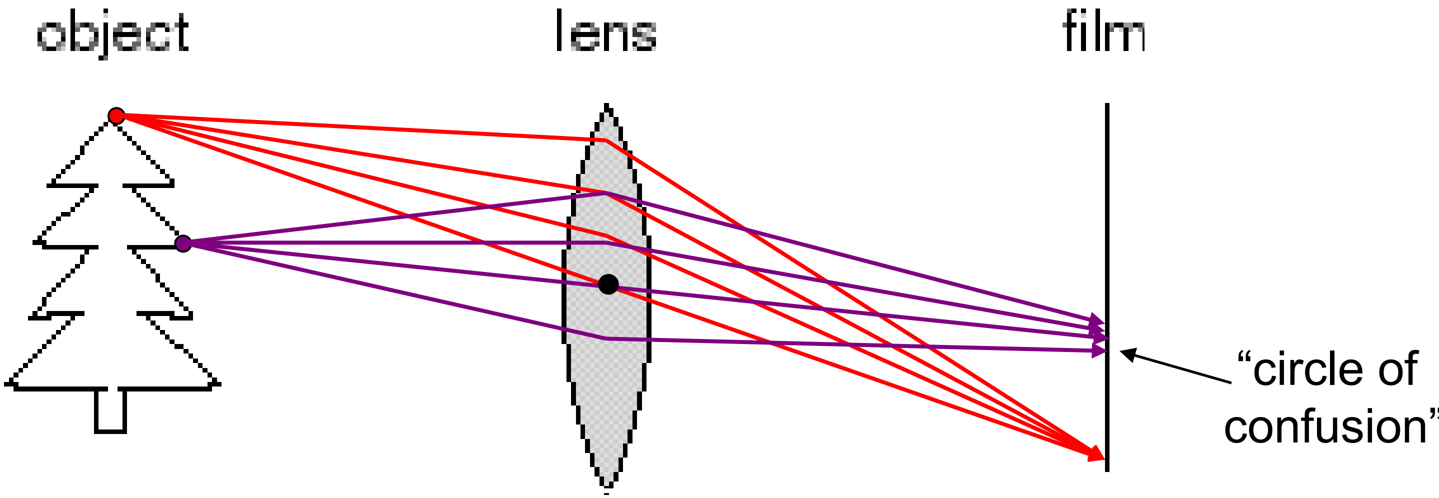

- A lens focuses light onto the film
  - There is a specific distance at which objects are “in focus”
    - other points project to a “circle of confusion” in the image
  - Changing the shape of the lens changes this distance