

6.869 Advances in Computer Vision

<http://people.csail.mit.edu/torralba/courses/6.869/6.869.computervision.htm>

Spring 2010

Lecture 1

Introduction



- Mondays/Wednesdays 1:00-2:30 pm
- Room 2-139
- Instructor: Antonio Torralba
- Email: torralba@csail.mit.edu
- TA: Joseph Lim

6.869 Advances in Computer Vision

Spring 2010



Overview

Advanced topics in computer vision with a focus on the use of machine learning techniques and applications in graphics and human-computer interface. Topics include image representations, texture models, structure-from-motion algorithms, Bayesian techniques, object and scene recognition, tracking, shape modeling, and image databases. Applications may include face recognition, multimodal interaction, interactive systems, cinematic special effects, and photorealistic rendering. Covers topics complementary to 6.801/6.866; these subjects may be taken in sequence.

General information

Lecture: Mondays/Wednesdays 1:00-2:30pm

Room: 2-139 ([where is this?](#))

Instructor: [Antonio Torralba](#)

E-mail: trrlb@mt.d (fill the missing vowels)

Office: D432

T.A.: Joseph Lim

Material:

- Textbook: [new book by Rick Szeliski](#) (not published yet, but a draft is available online)
- Textbook: Computer vision: a modern approach, by Forsyth and Ponce. Prentice Hall, 2002.
- The class will make use of [MATLAB](#).

Grading:

- problem sets: 1/3
- 2 take-home exams: 1/3
- final project: 1/3

Announcements

First class, Wednesday Feb 3rd

Schedule

| Lecture | Date | Topic | Slides | Readings | Assignments | Additional material |
|---------|------|-------|--------|----------|-------------|---------------------|
|---------|------|-------|--------|----------|-------------|---------------------|

Readings

Untitled3

Computer vision: Algorithms and Applications

<http://research.microsoft.com/en-us/um/people/szeliski/Book/>

Computer Vision: Algorithms and Applications

(c) [Richard Szeliski](#), Microsoft Research

Welcome to the repository for drafts of my computer vision textbook.

This book is largely based on the computer vision courses that I have co-taught at the University of Washington ([2008](#), [2005](#), [2001](#)) and Stanford (2003) with [Steve Seitz](#) and [David Fleet](#).

While I am working on the book, I would *love* to have people "test-drive" it in their computer vision courses (or their research) and [send me feedback](#).

The PDFs should be enabled for commenting directly in your viewer. Also, hyper-links to sections, equations, and references are enabled. To get back to where you were, use Alt-Left-Arrow in Acrobat.

This Web site is also a placeholder for the site that will accompany my computer vision textbook once it is published. Once I get further along with the project, I hope to publish supplemental course material here, such as figures and images from the book, slides sets, pointers to software, and a bibliography.

Latest draft

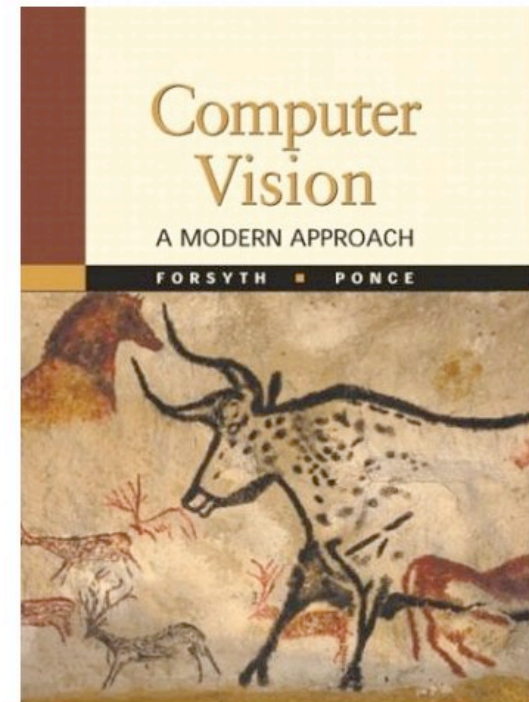
[February 2, 2009](#)

Older drafts

[January 20, 2009](#)

[January 12, 2009](#)

Done



Course requirements

- Two take-home exams (given out Monday, due back Wednesday)
- Five problem sets with lab exercises in Matlab (one to two weeks per problem set)
- Final project
(no final exam)

Grading

- Problem sets are graded check, check-plus, check-minus. (Outstanding solutions get extra credit.)
- Final grade:
 - 5 problem sets: 1/3
 - 2 take-home exams: 1/3
 - final project: 1/3

Collaboration policy

Problem sets may be discussed, but all written work and coding must be done individually. Please note on your problem sets who you discussed the homework problems with.

Take-home exams

Take-home exams may not be discussed.

Individuals found submitting duplicate or substantially similar materials due to inappropriate collaboration may get an F in this class and other sanctions.

Final project

The final project may be

- An original implementation of a new or published idea
- A detailed empirical evaluation of an existing implementation of one or more methods
- A paper comparing three or more papers not covered in class, or surveying recent literature in a particular area
- Something related to your research.

A project proposal not longer than two pages must be submitted by April 1st. I can provide ideas or suggestions for projects.

Prerequisites

- Familiarity with linear algebra
- Familiarity with probability
- Covers topics complementary to 6.801/866 and these subjects may be taken in sequence.
Prerequisites: 6.041 or 6.042; 18.06

Other classes

6.801/6.866 Machine Vision

6.869 Advances in Computer Vision

6.870 Advanced Topics in Computer Vision

6.815/6.865 Digital and Computational Photography, Fredo & Bill

MAS 132/532 Camera Culture: Future of Imaging, Ramesh

6.344 Digital Image Processing, J. Lim

6.342 Wavelets, Approximation, and Compression, V. K. Goyal

What is vision?

- What does it mean, to see? “to know what is where by looking”.
- How to discover from images what is present in the world, where things are, what actions are taking place.

The importance of images

Some images are more important than others

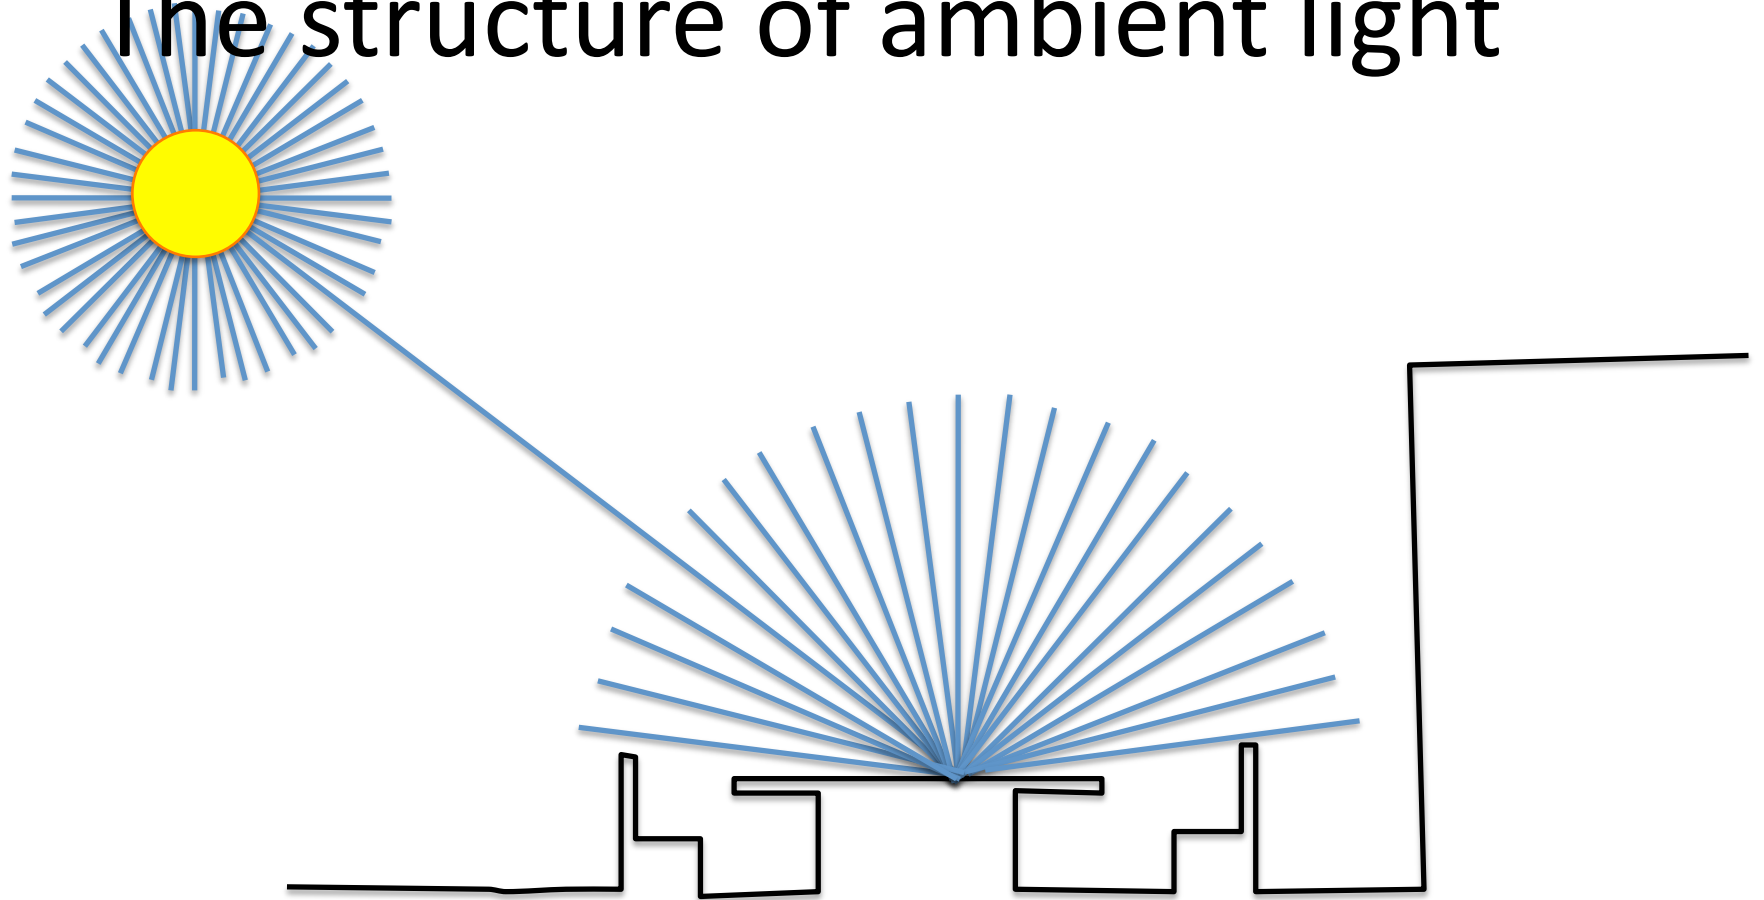


"Dora Maar au Chat"
Pablo Picasso, 1941

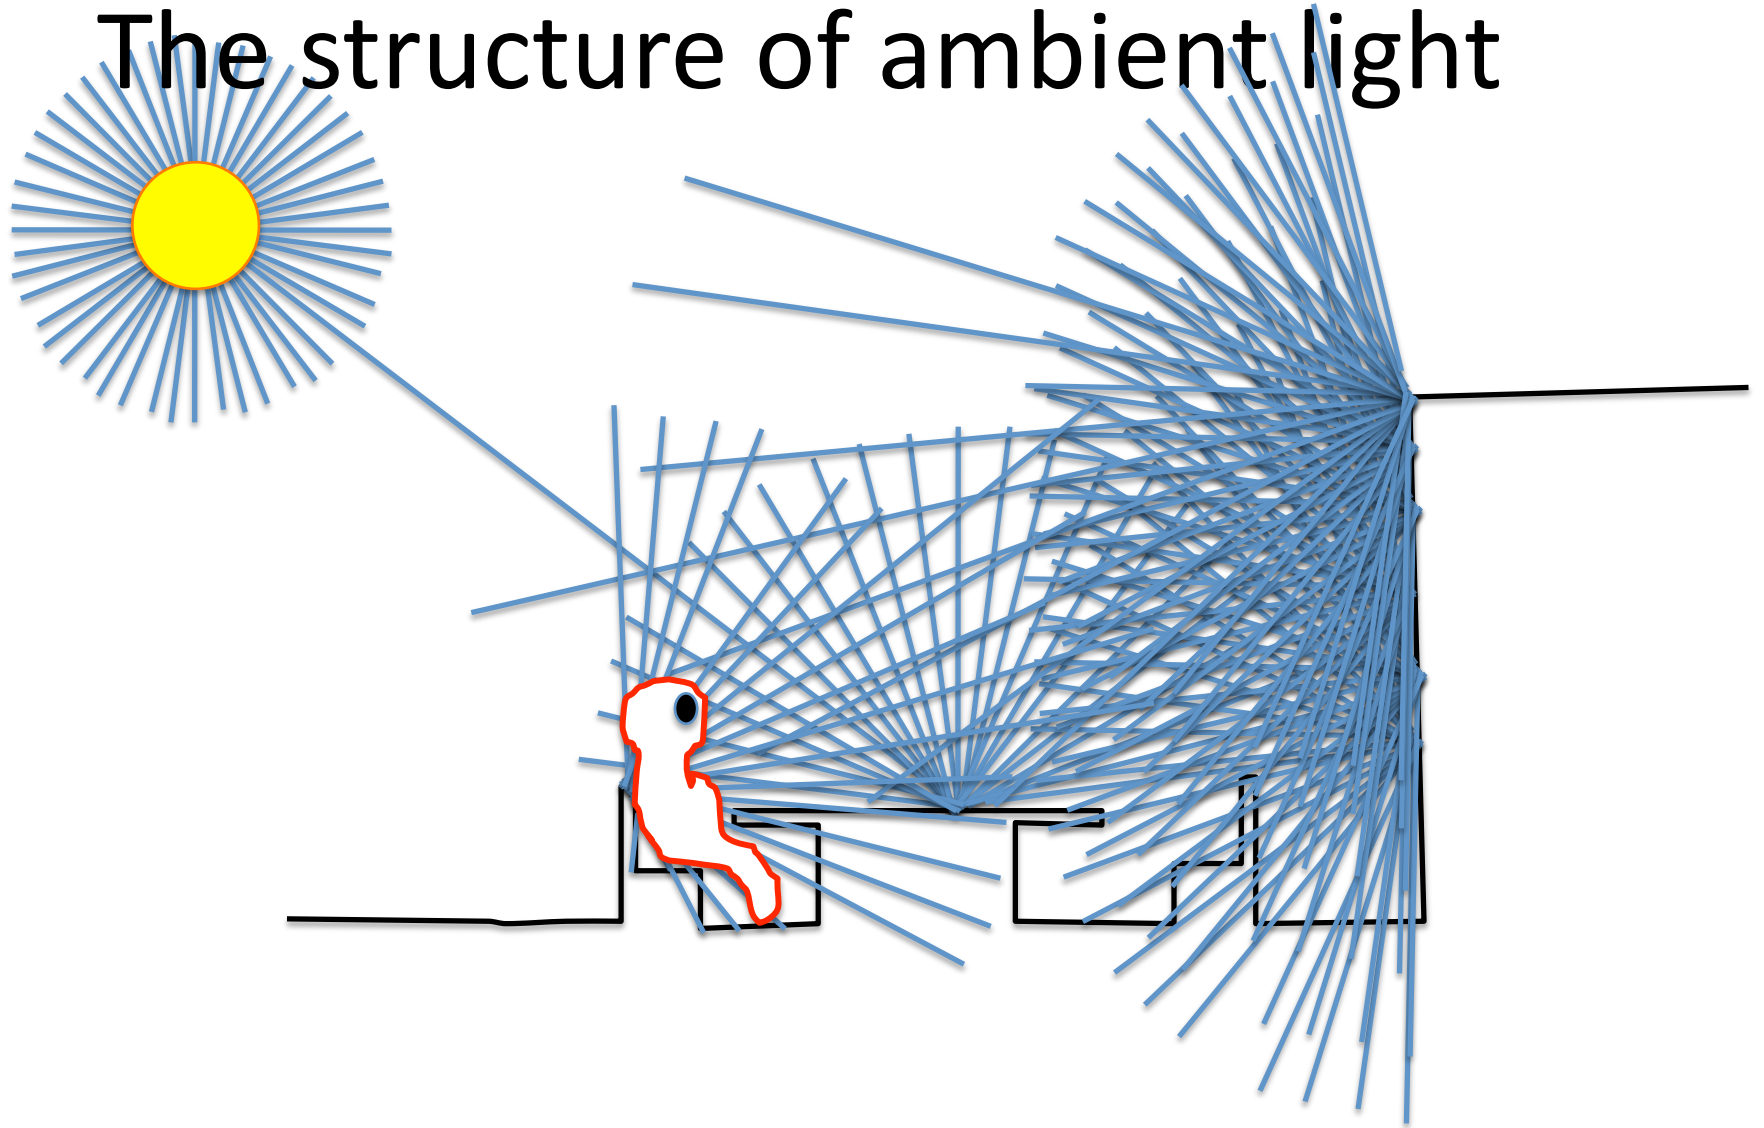
100 million \$

Why is vision hard?

The structure of ambient light

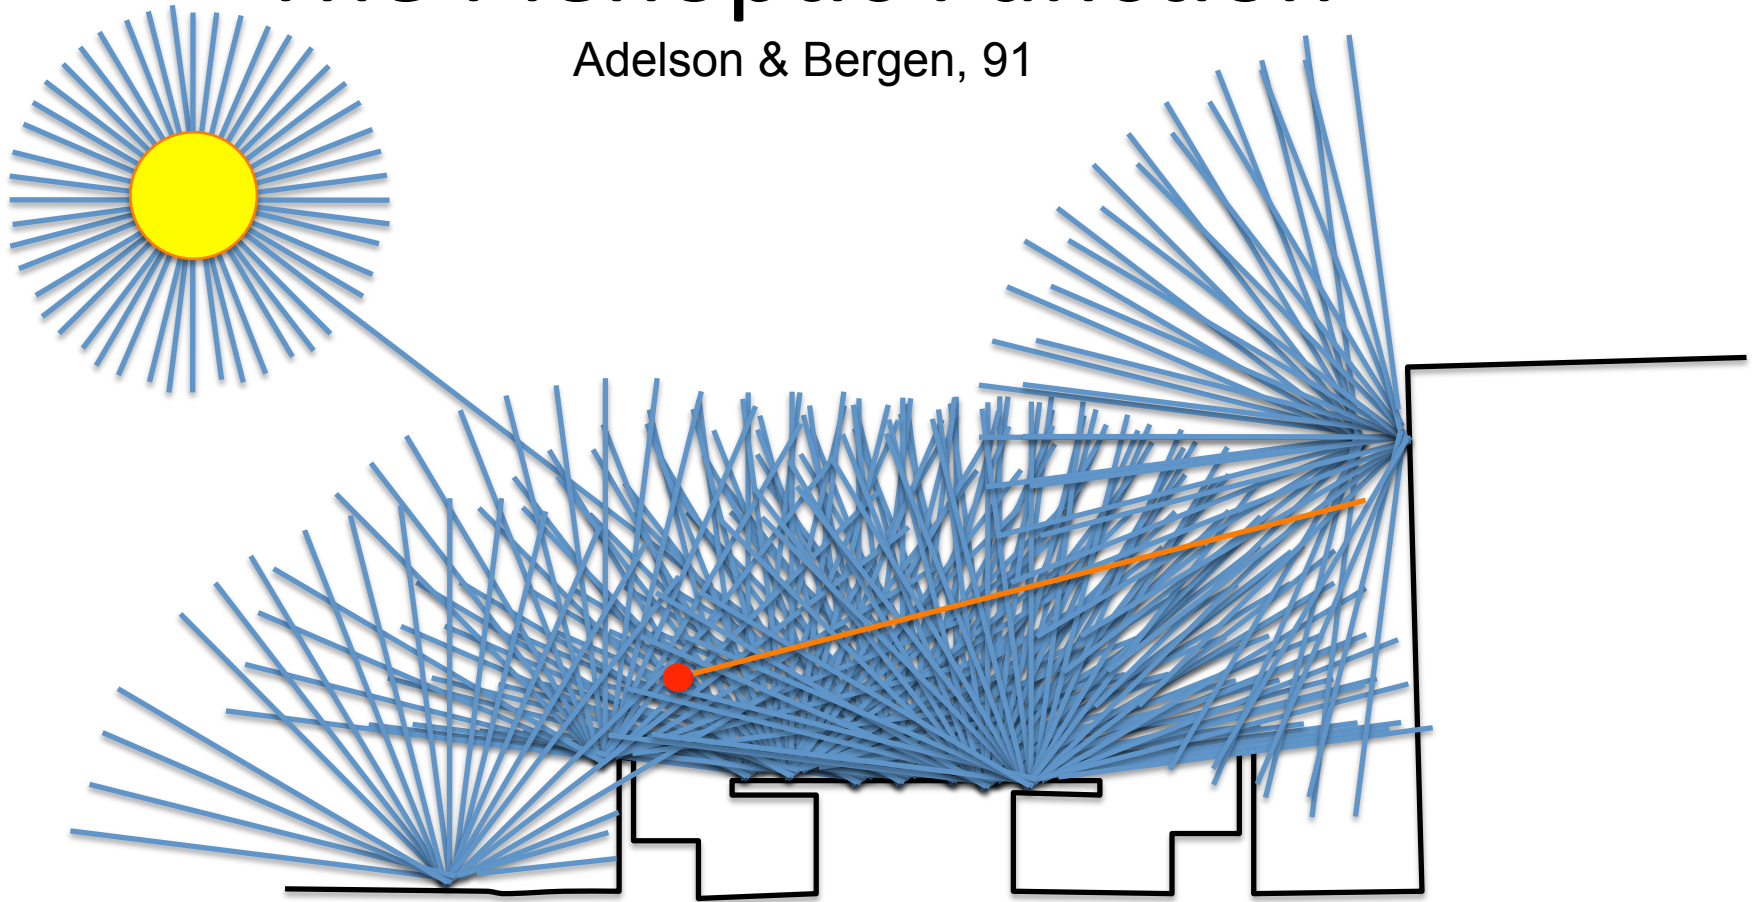


The structure of ambient light



The Plenoptic Function

Adelson & Bergen, 91



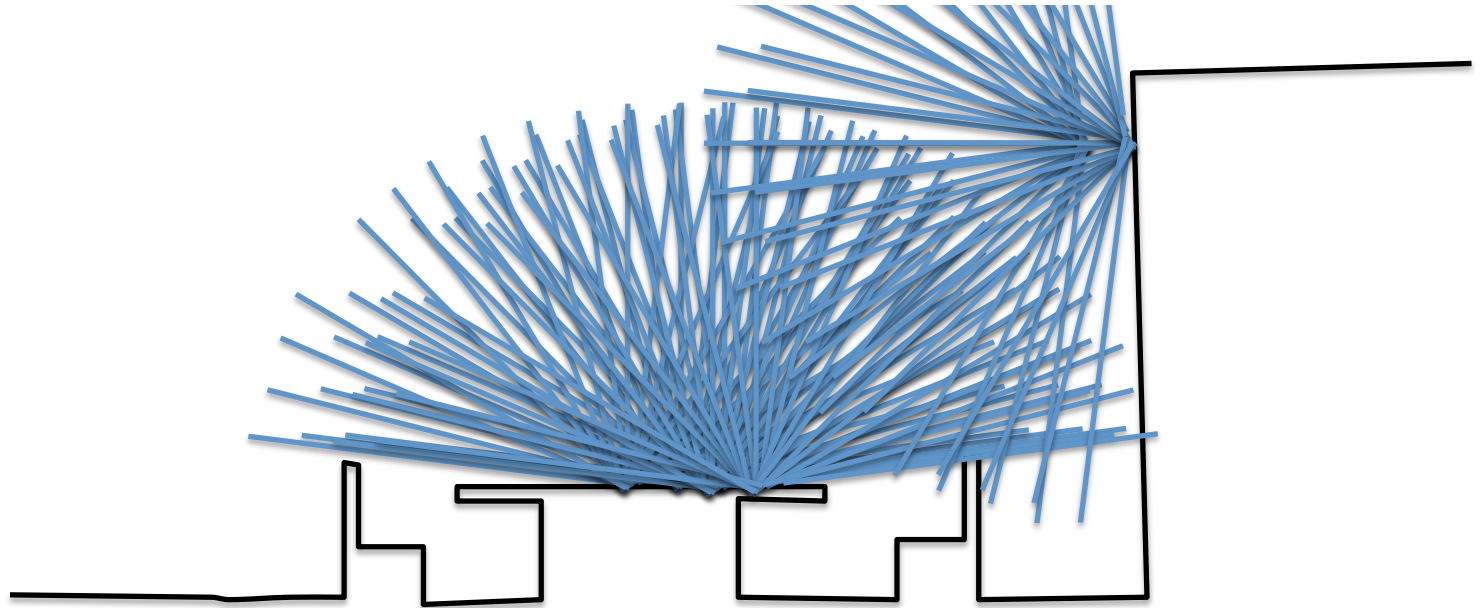
The intensity P can be parameterized as:

$$P(\theta, \phi, \lambda, t, X, Y, Z)$$

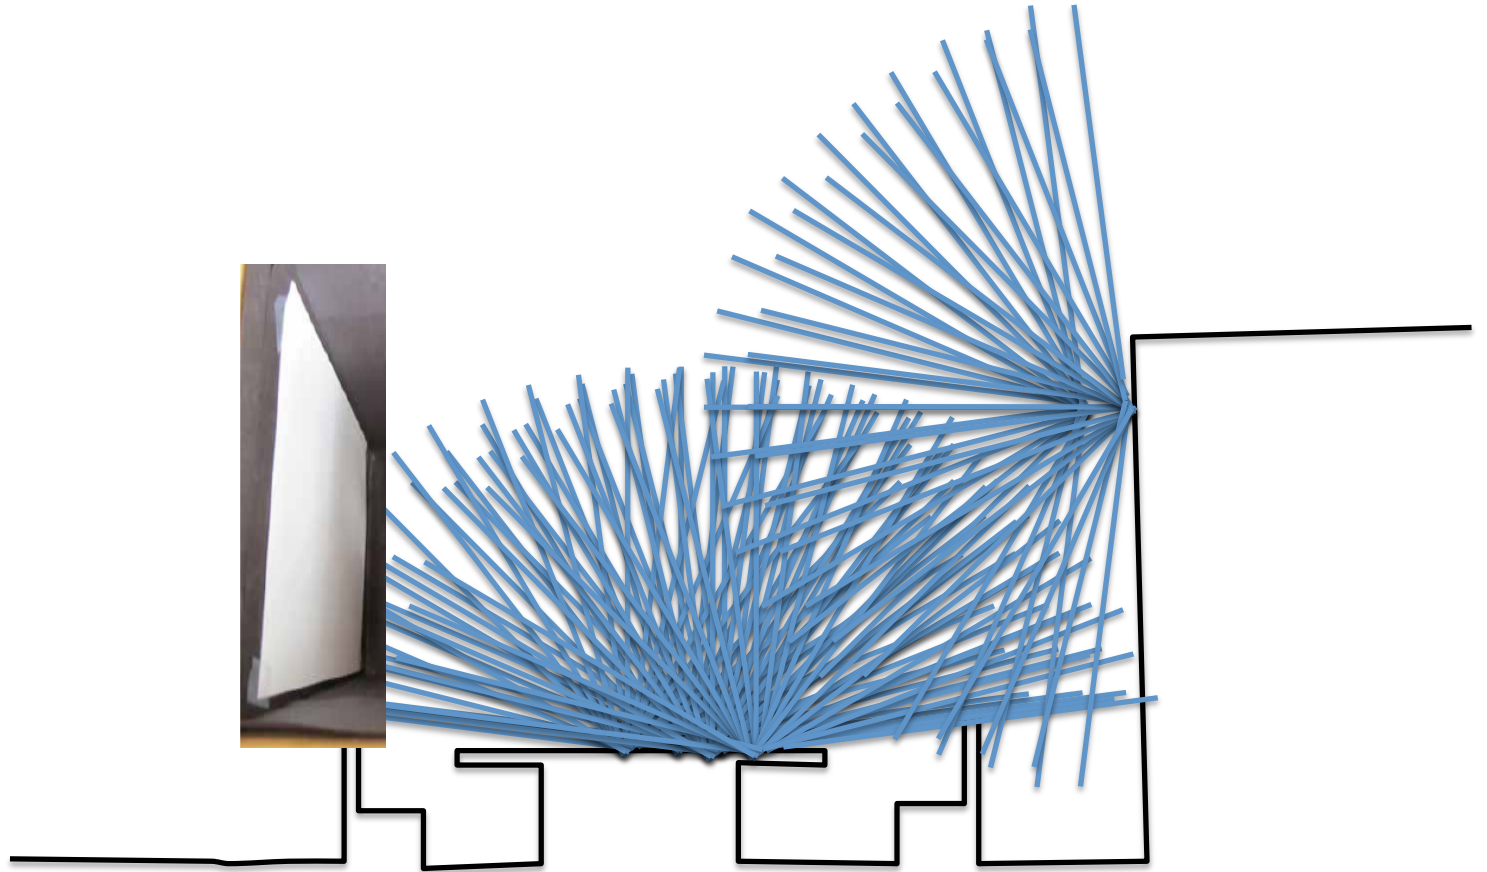
“The complete set of all convergence points constitutes the permanent possibilities of vision.” Gibson

Measuring the Plenoptic function

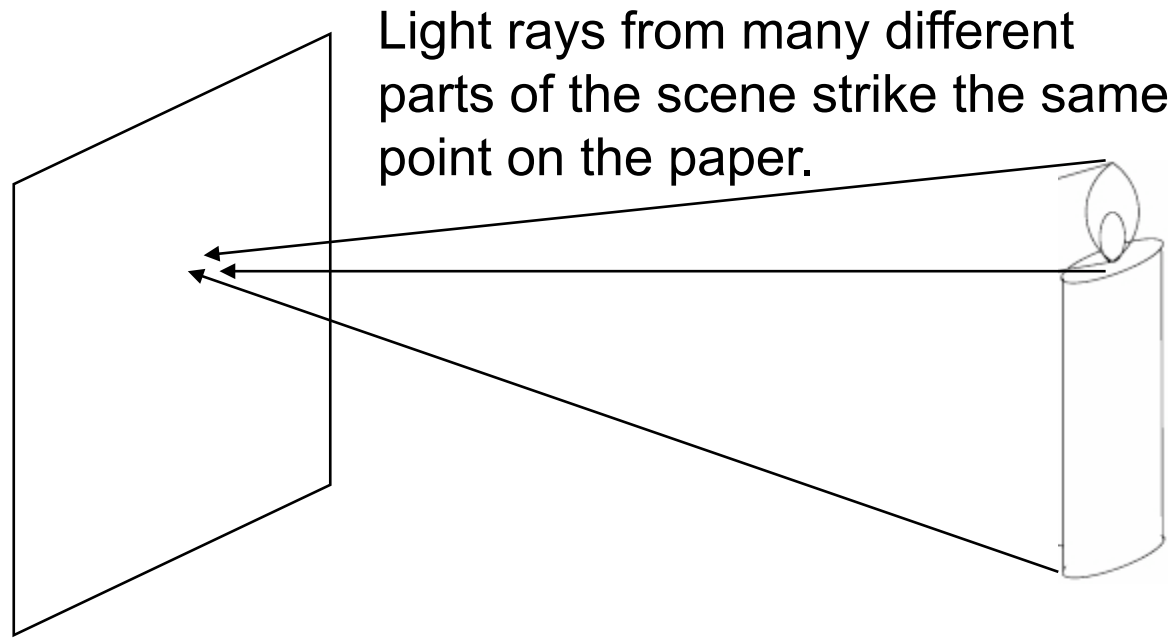
“The significance of the plenoptic function is this: The world is made of 3D objects, but these objects do not communicate their properties directly to an observer. Rather, the objects fill the space around them with the pattern of light rays that constitutes the plenoptic function, and the observer takes samples from this function.” Adelson, & Bergen 91.



Measuring the Plenoptic function

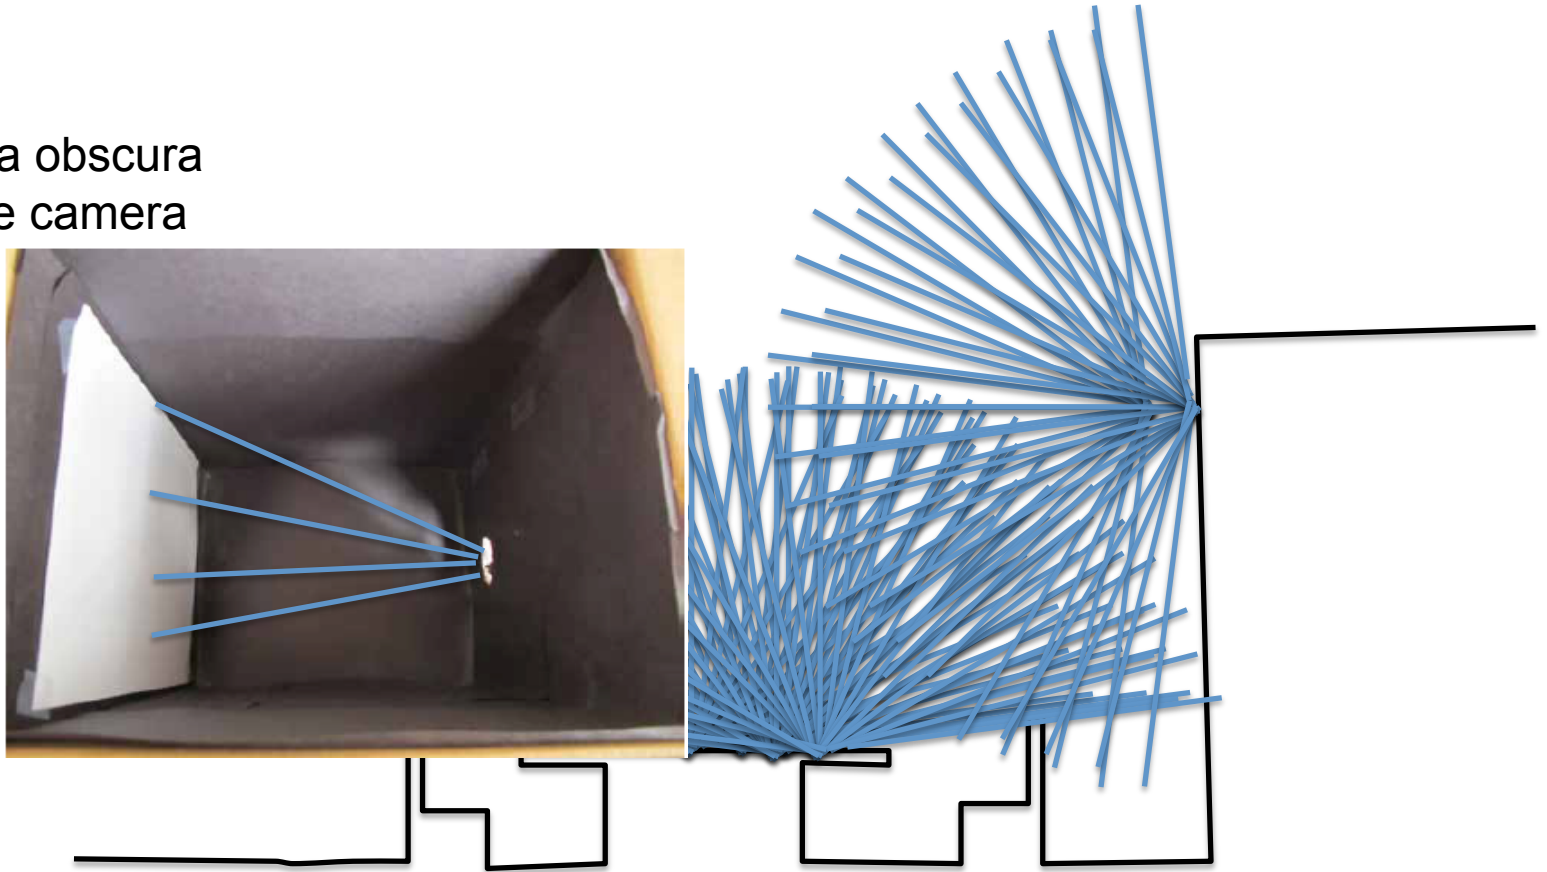


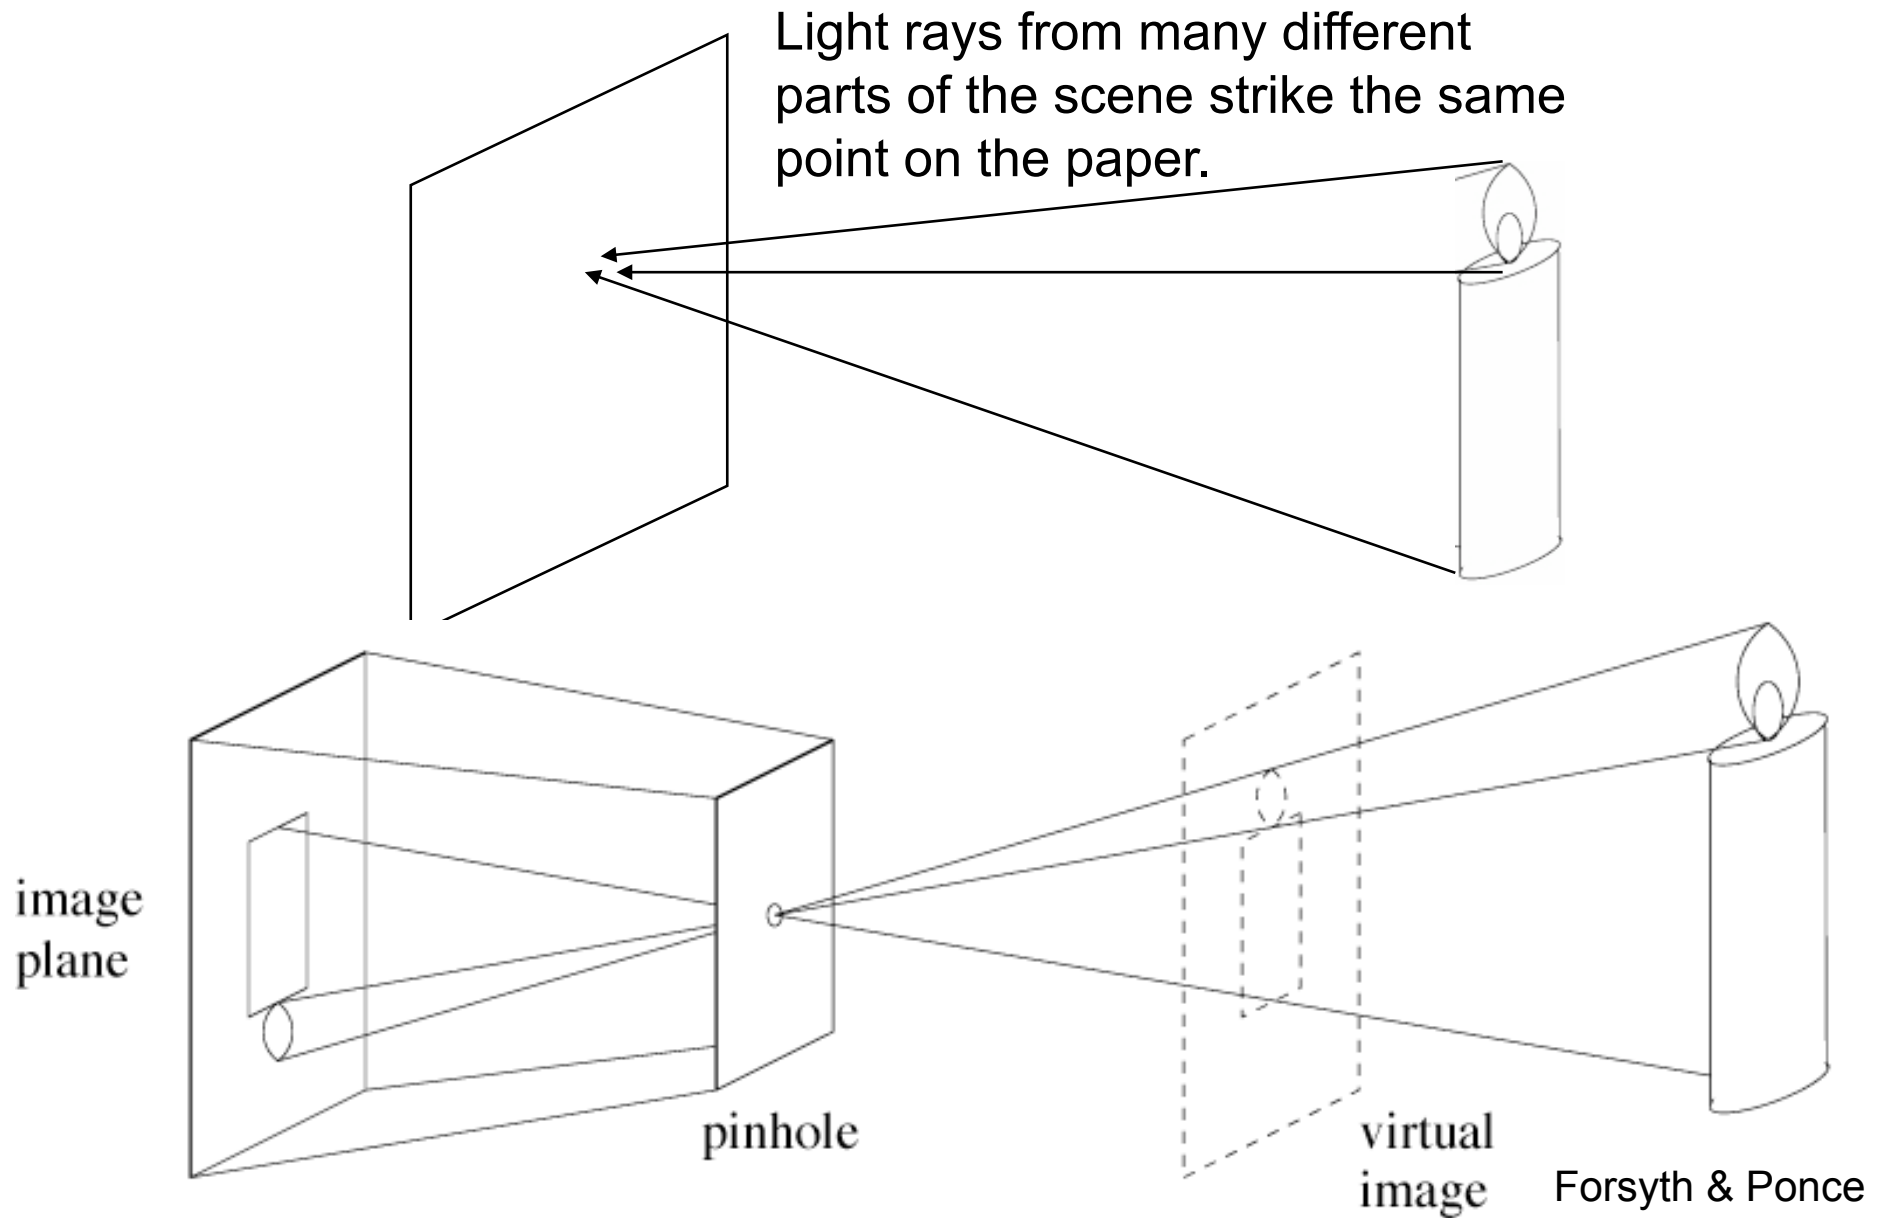
Why is there no picture appearing on the paper?



Measuring the Plenoptic function

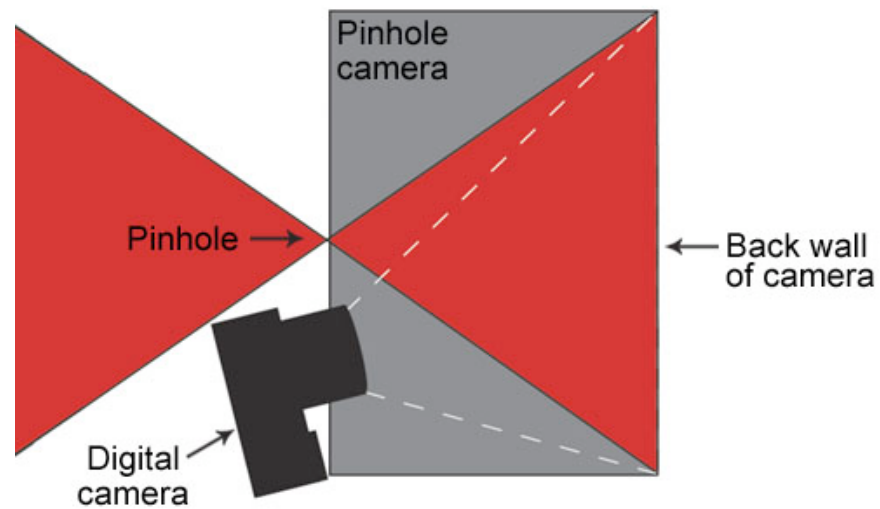
The camera obscura
The pinhole camera





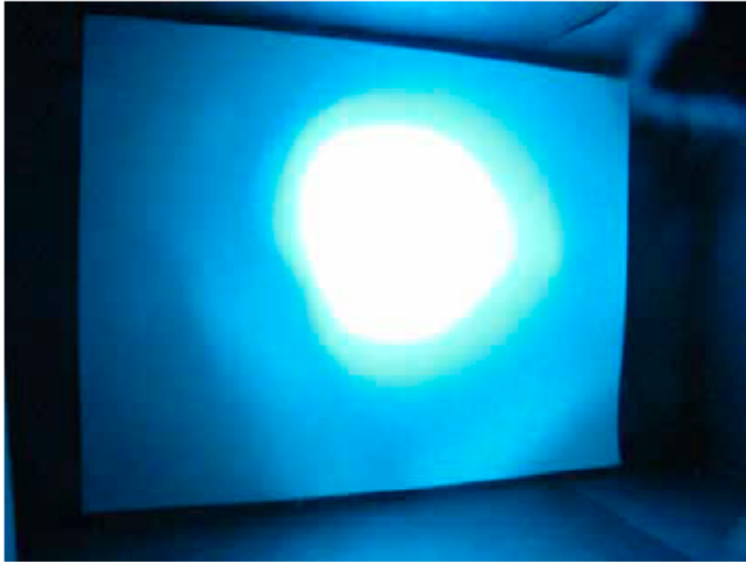
The pinhole camera only allows rays from one point in the scene to strike each point of the paper.

Problem Set 1

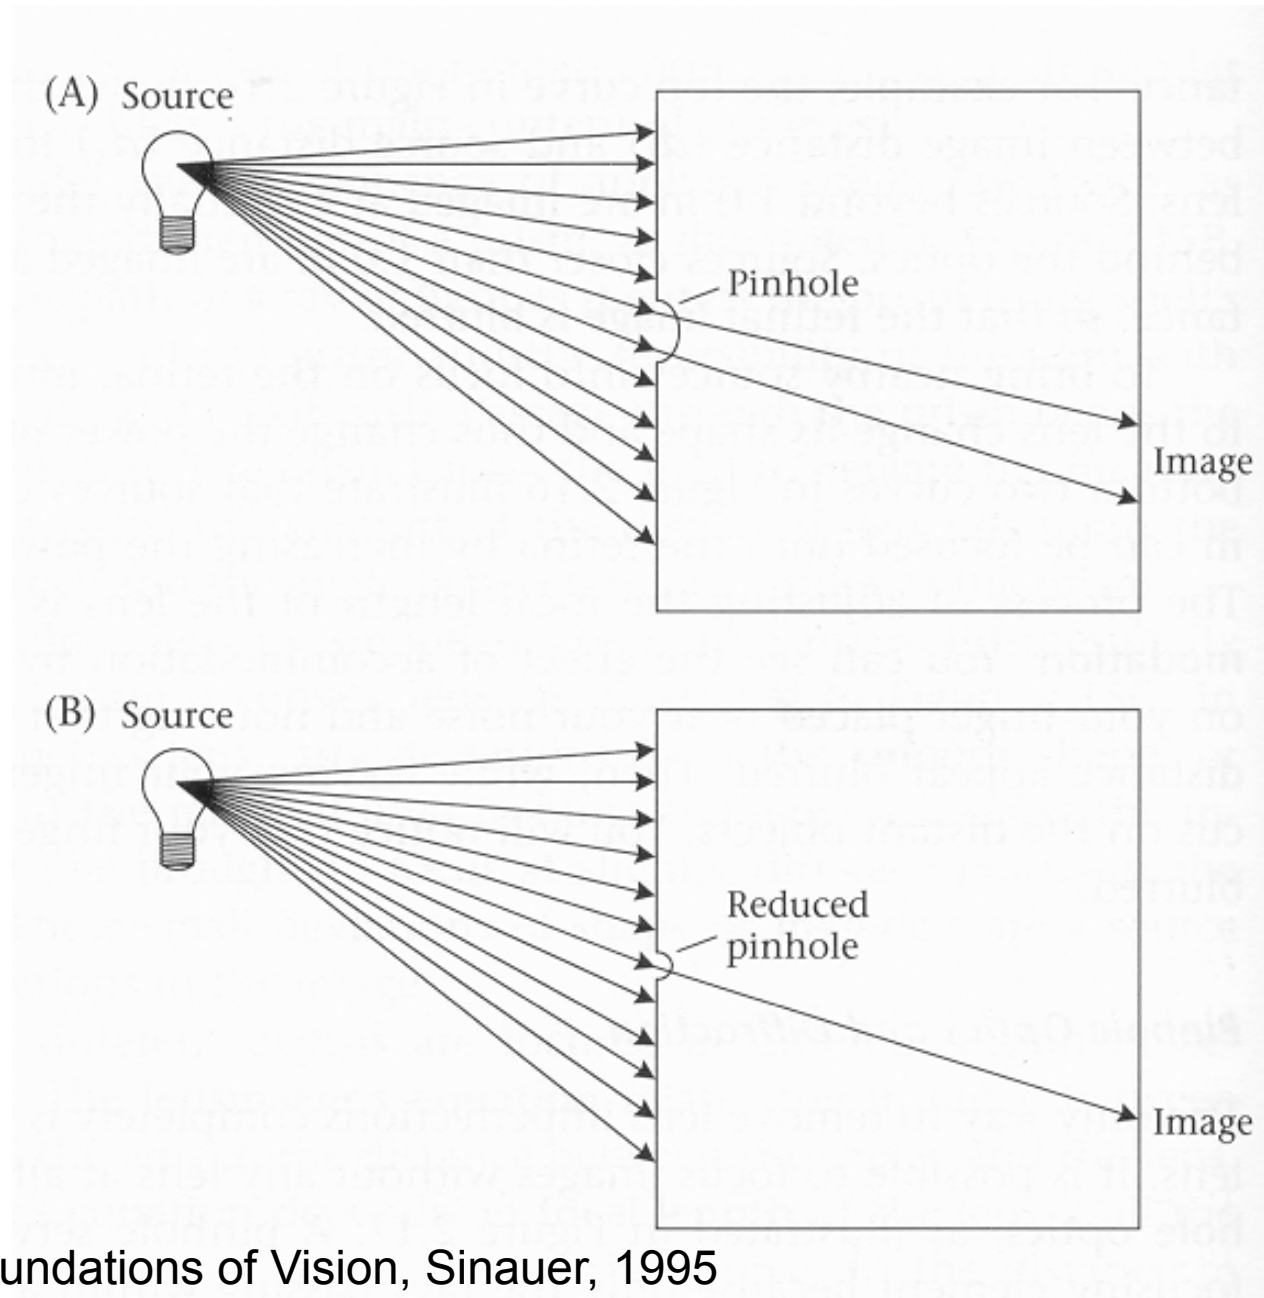


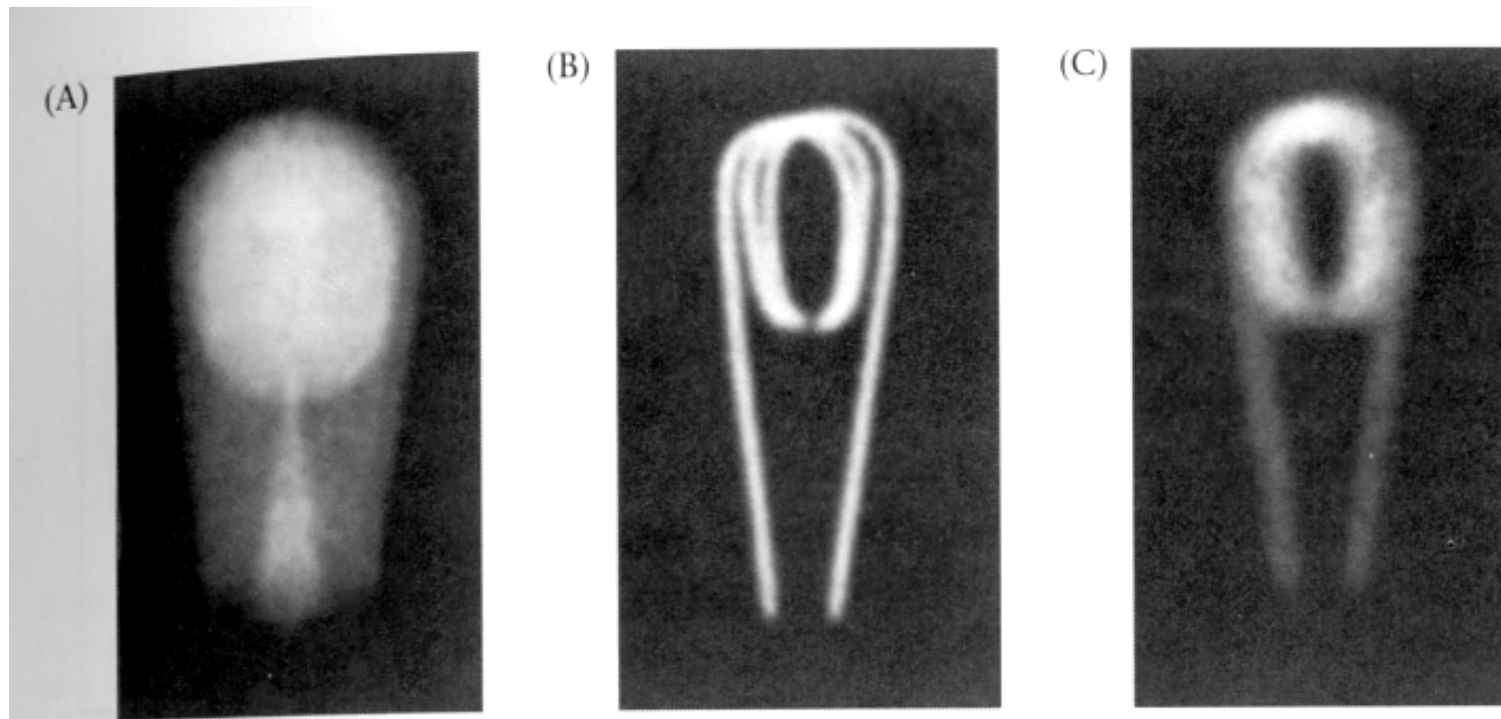
http://www.foundphotography.com/PhotoThoughts/archives/2005/04/pinhole_camera_2.html

Problem Set 1



Effect of pinhole size





2.18 DIFFRACTION LIMITS THE QUALITY OF PINHOLE OPTICS. These three images of a bulb filament were made using pinholes with decreasing size. (A) When the pinhole is relatively large, the image rays are not properly converged, and the image is blurred. (B) Reducing the size of the pinhole improves the focus. (C) Reducing the size of the pinhole further worsens the focus, due to diffraction. From Ruechardt, 1958.

Animal Eyes

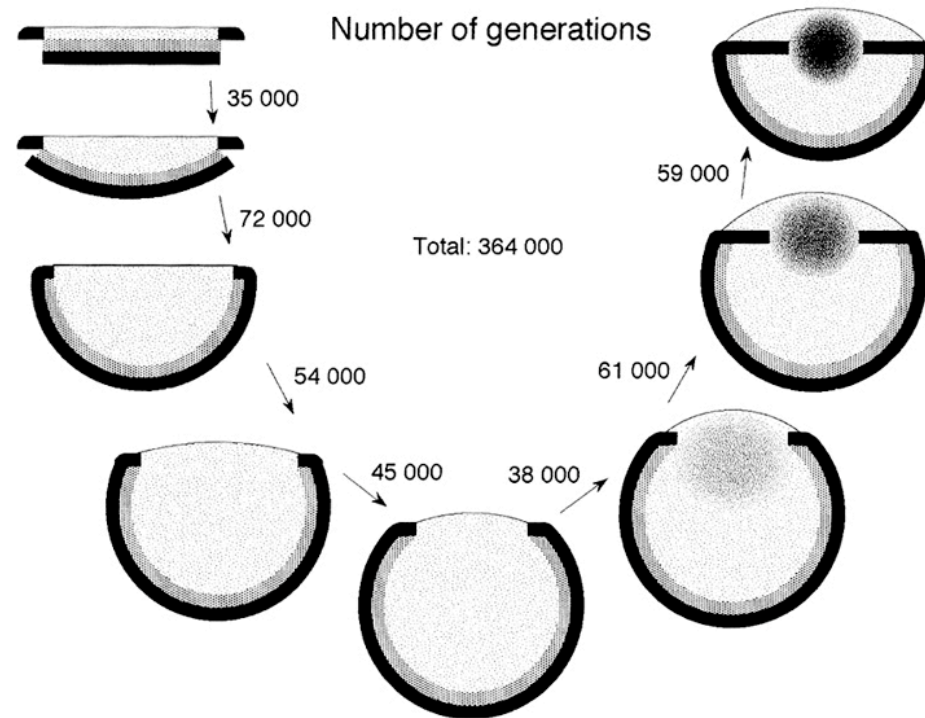
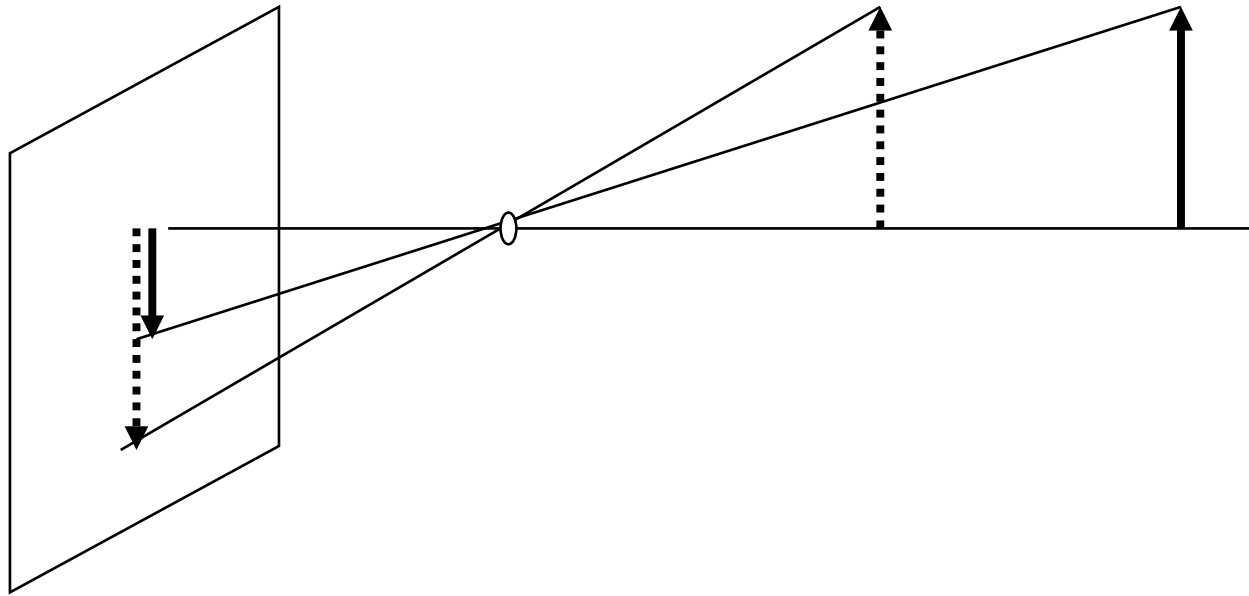


Fig. 1.6 A patch of light sensitive epithelium can be gradually turned into a perfectly focussed camera-type eye if there is a continuous selection for improved spatial vision. A theoretical model based on conservative assumptions about selection pressure and the amount of variation in natural populations suggest that the whole sequence can be accomplished amazingly fast, in less than 400 000 generations. The number of generations is also given between each of the consecutive intermediates that are drawn in the figure. The starting point is a flat piece of epithelium with an outer protective layer, an intermediate layer of receptor cells, and a bottom layer of pigment cells. The first half of the sequence is the formation of a pigment cup eye. When this principle cannot be improved any further, a lens gradually evolves. Modified from Nilsson and Pelger (1994).

Animal Eyes. Land & Nilsson. Oxford Univ. Press

Measuring distance

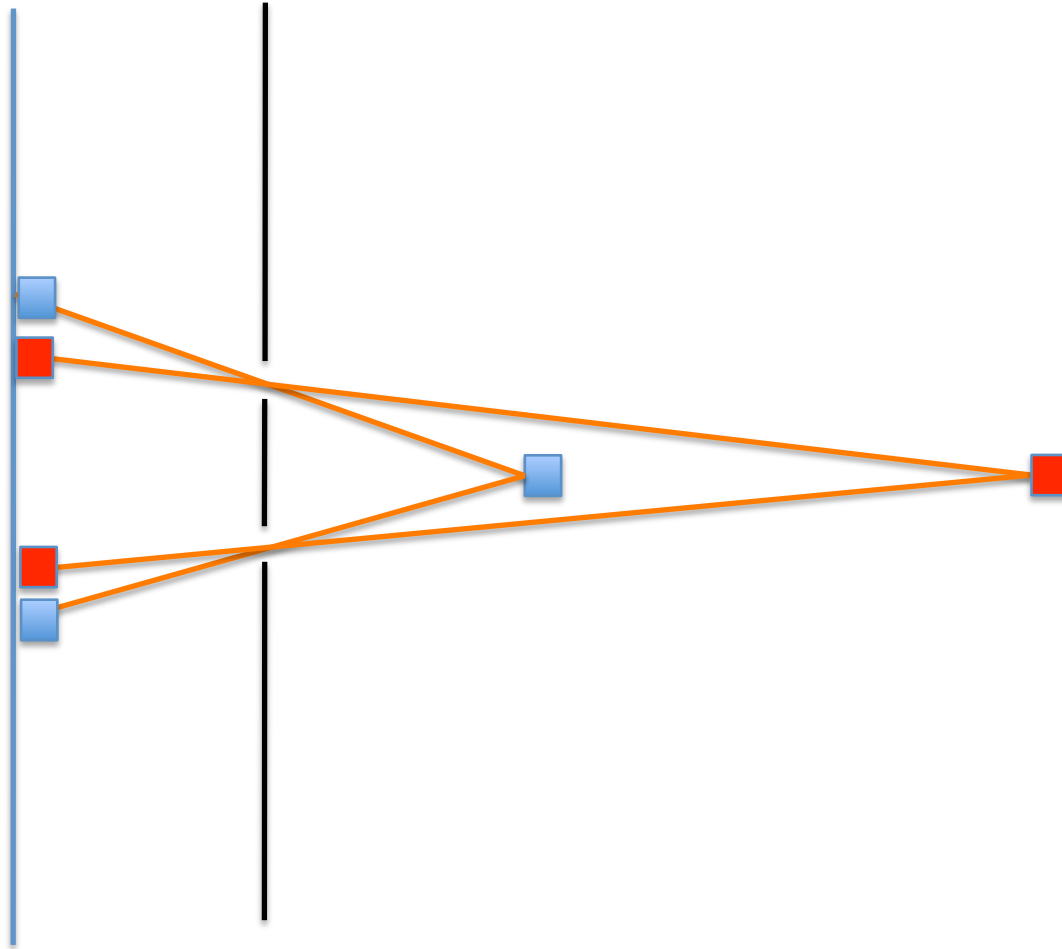


- Object size decreases with distance to the pinhole
- There, given a single projection, if we know the size of the object we can know how far it is.
- But for objects of unknown size, the 3D information seems to be lost.

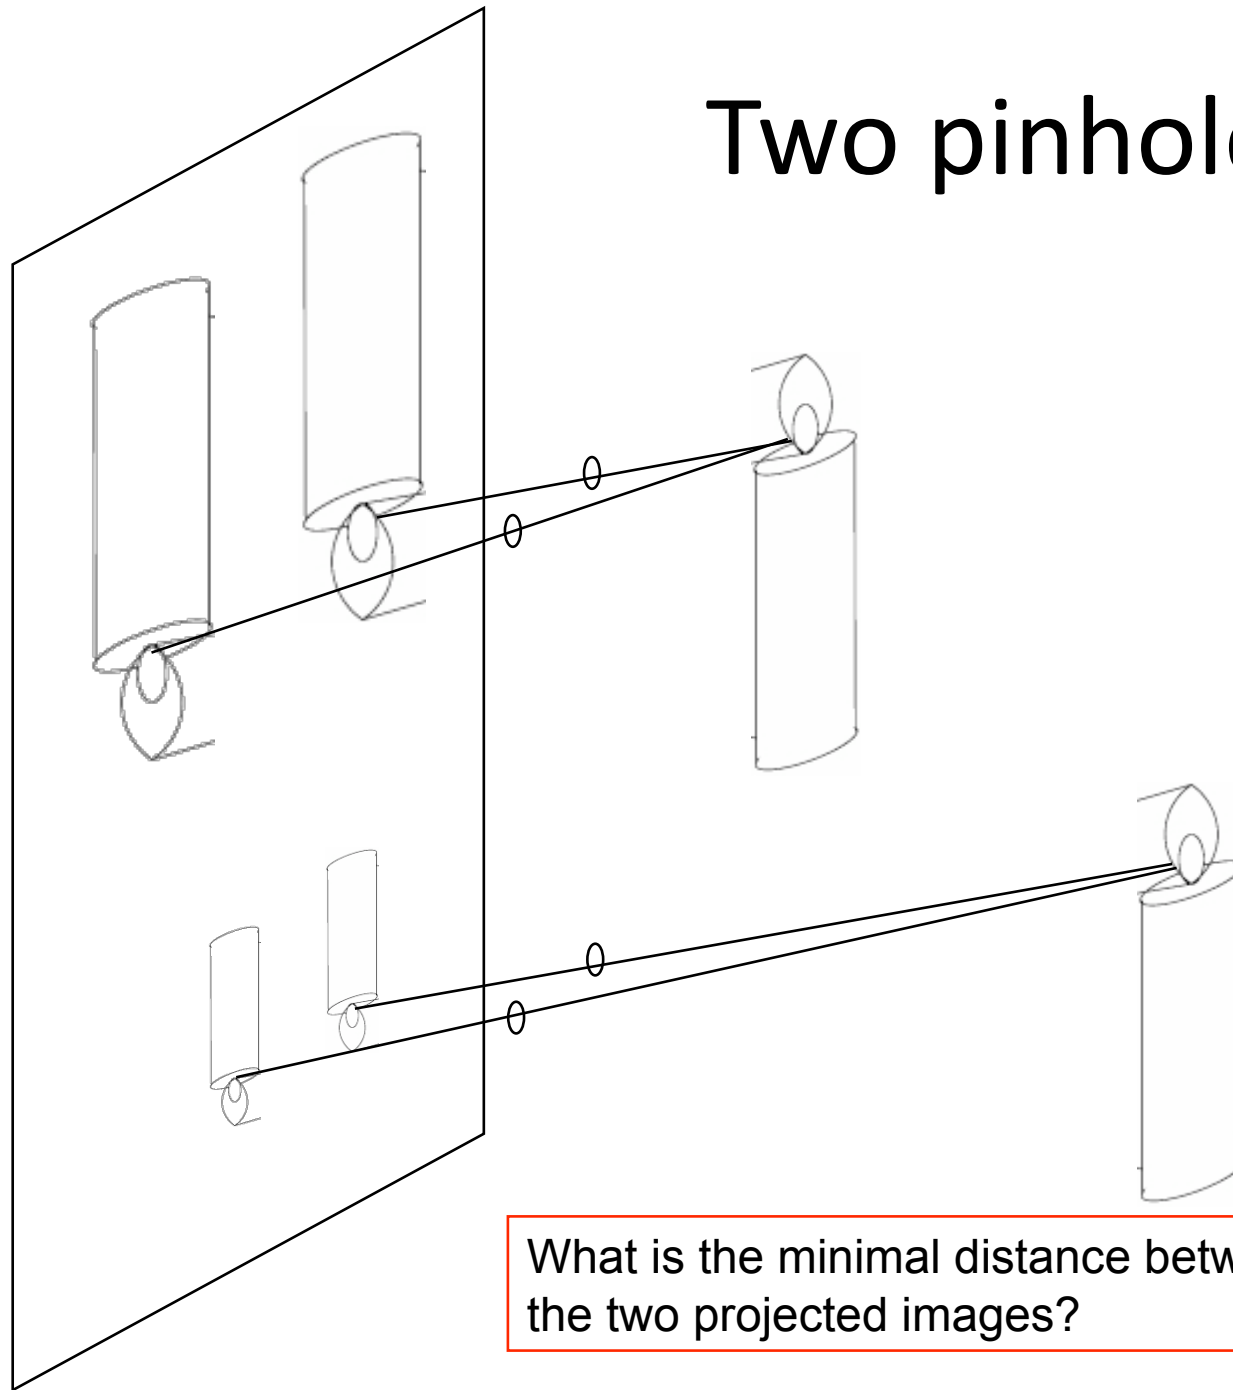
Playing with pinholes



Two pinholes

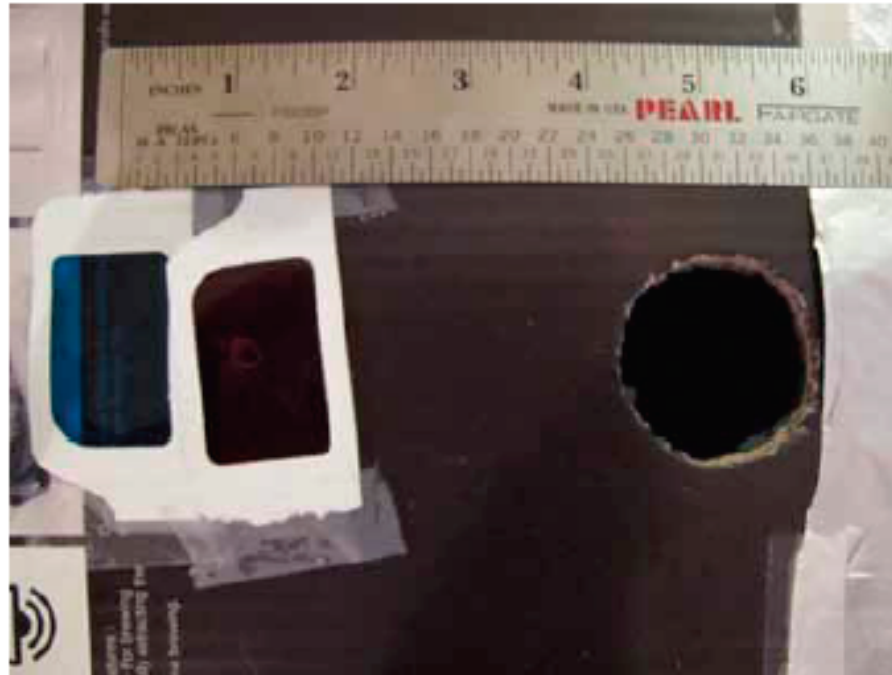


Two pinholes

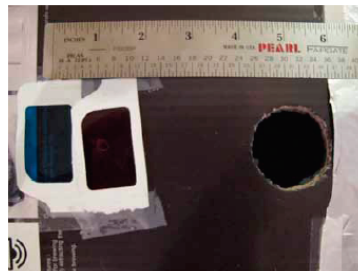
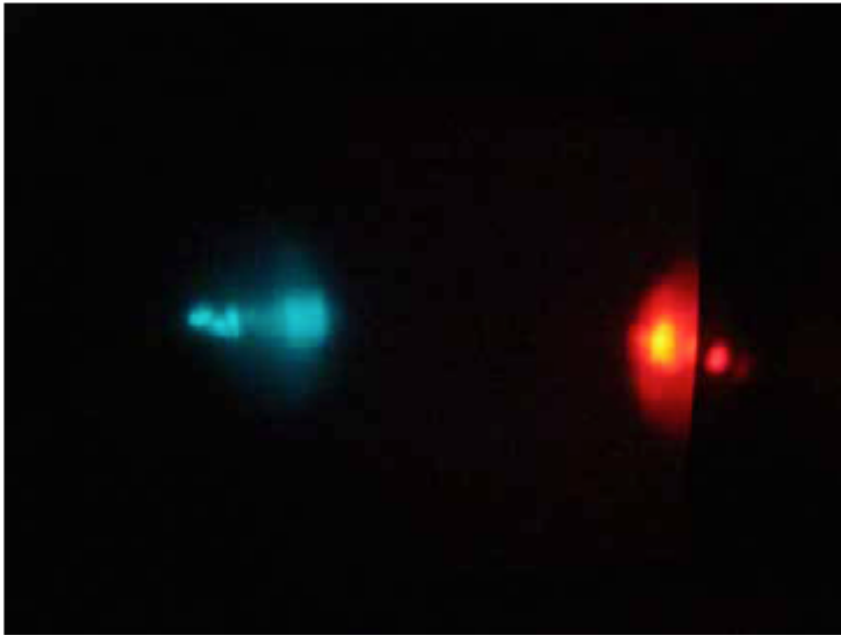


What is the minimal distance between the two projected images?

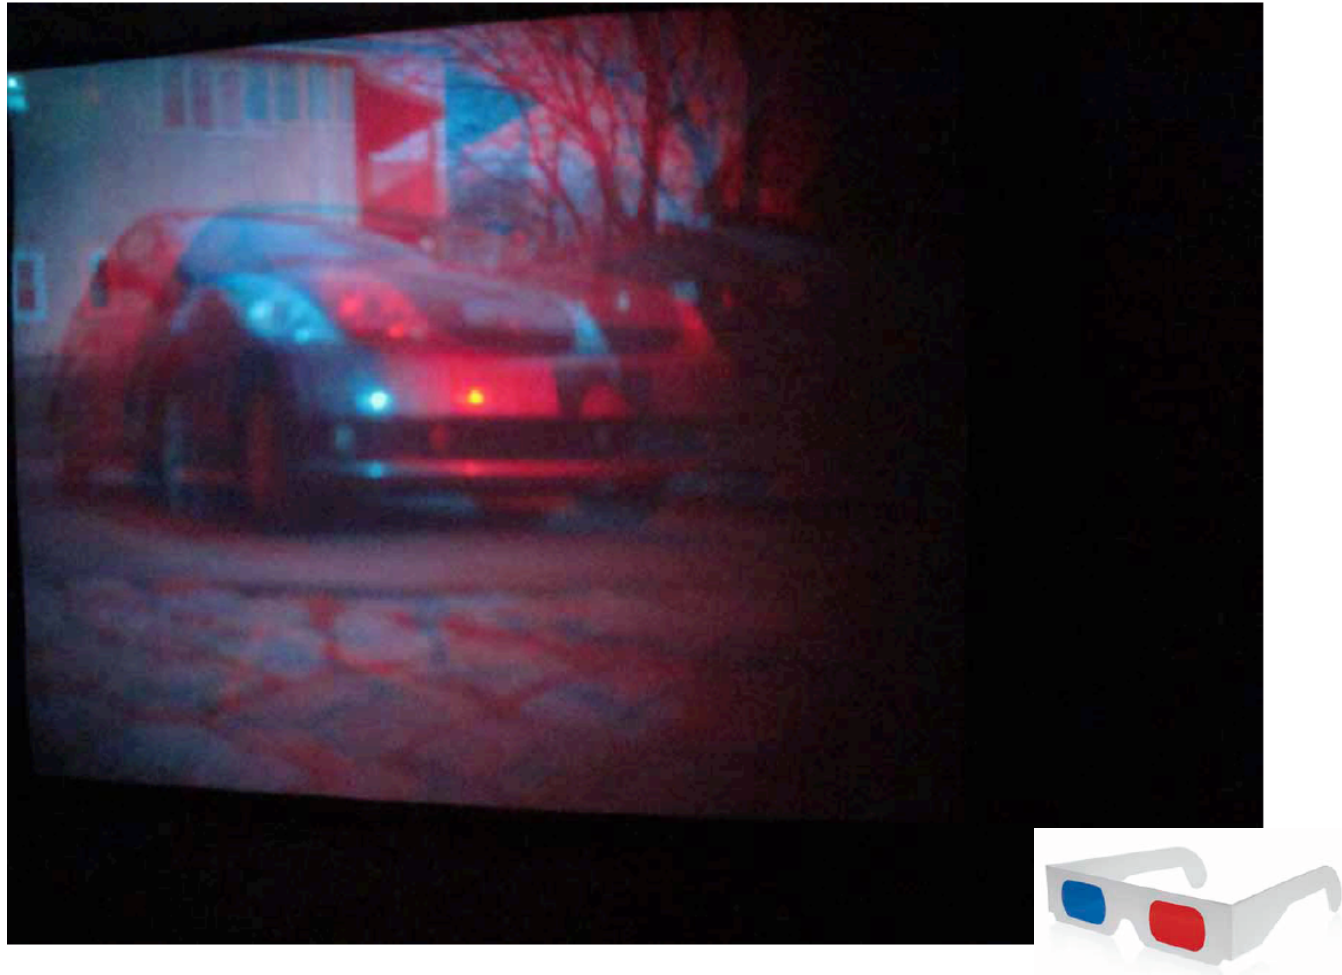
Anaglyph pinhole camera



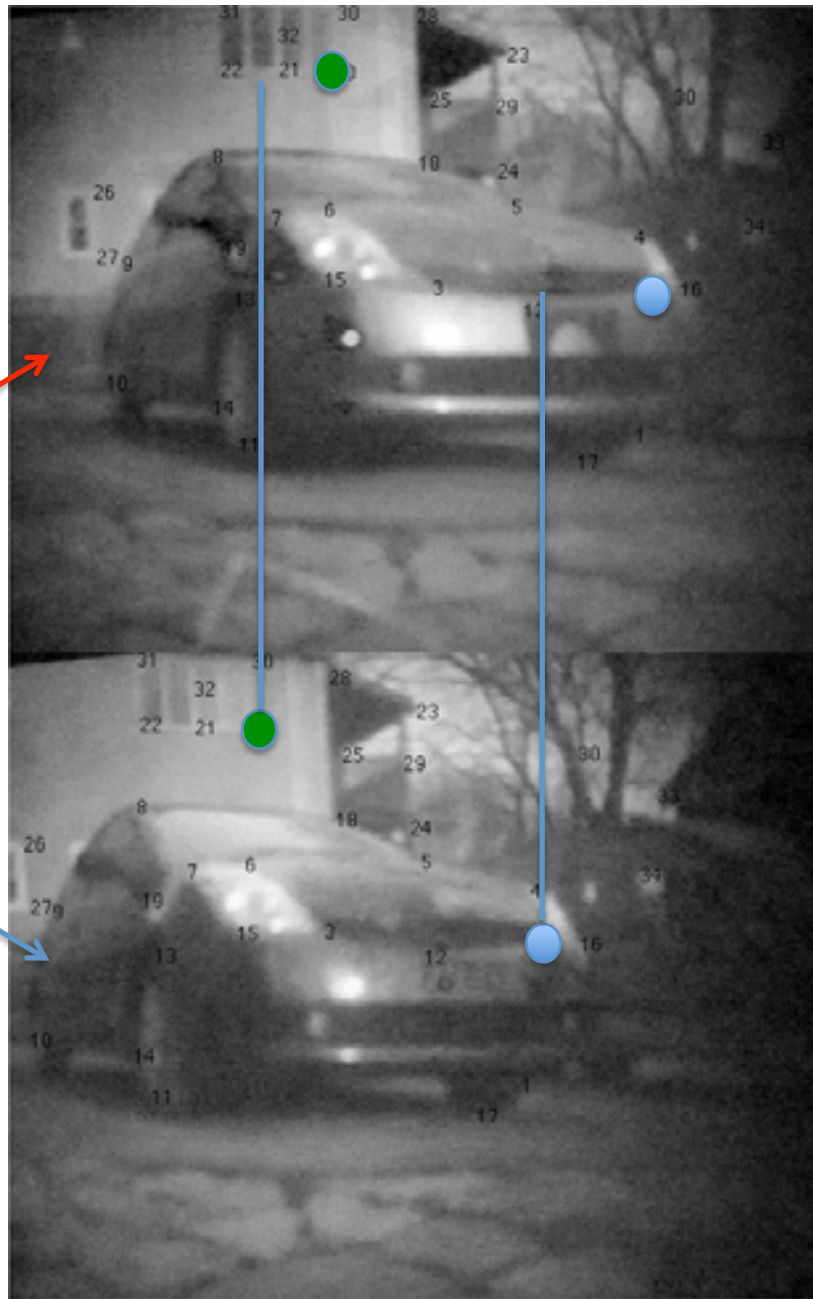
Anaglyph pinhole camera



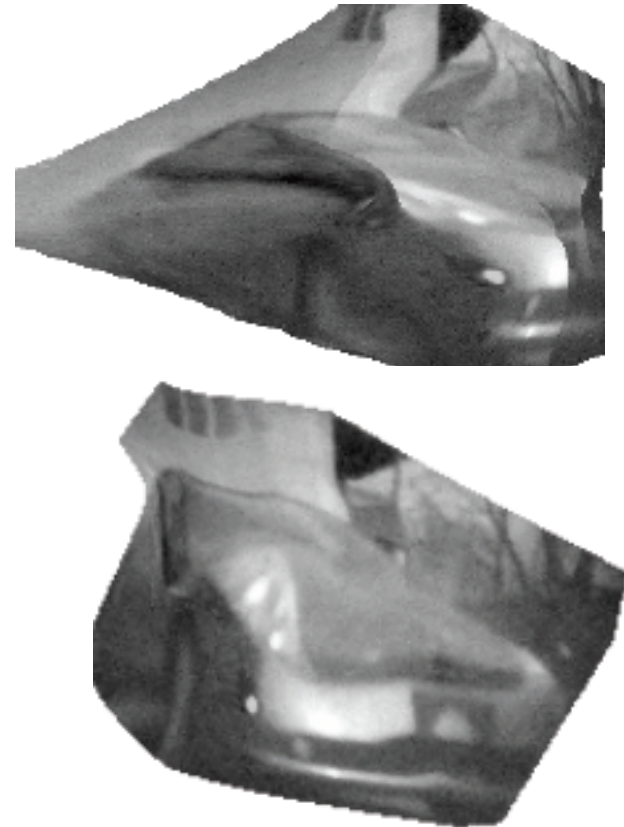
Anaglyph pinhole camera



Anaglyph



Synthesis of new views



Problem set 1

- Build the device
- Take some pictures and put them in the report
- Take anaglyph images
- Work out the geometry
- Recover depth for some points in the image

Why is vision hard?

Some things have strong variations
in appearance



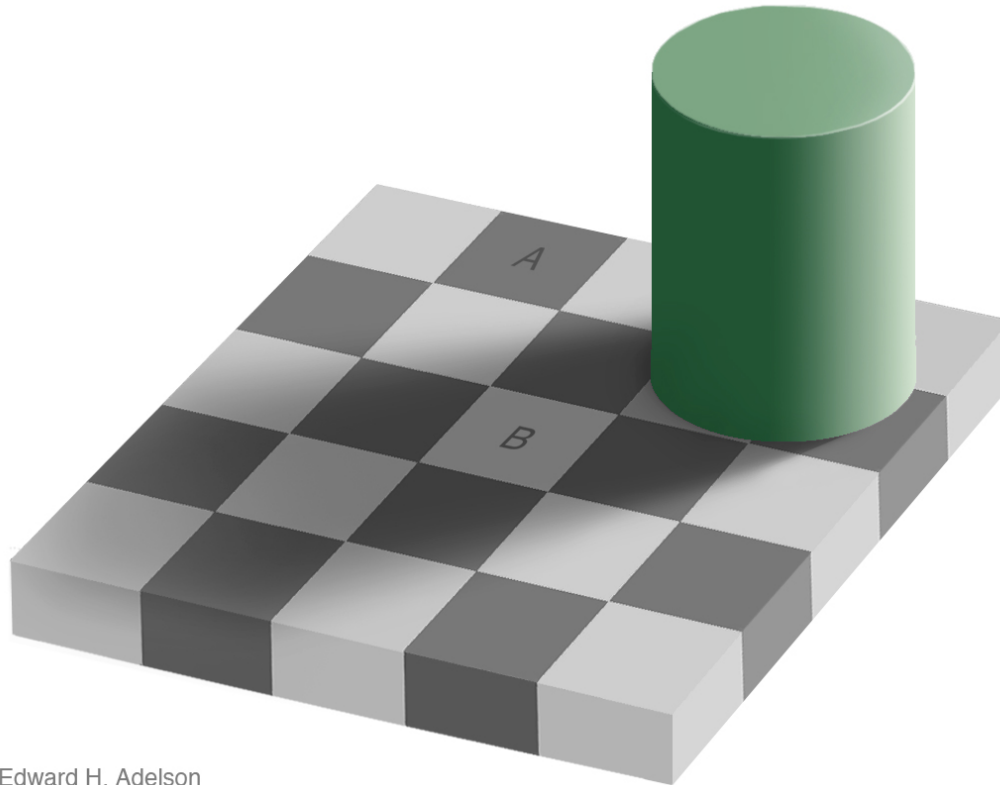
Some things know that you have eyes



Brady, M. J., & Kersten, D. (2003). Bootstrapped learning of novel objects. *J Vis*, 3(6), 413-422

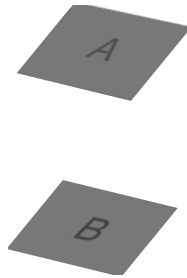
Why is vision hard?

Measuring light vs. measuring scene properties



Edward H. Adelson

Measuring light vs. measuring scene properties

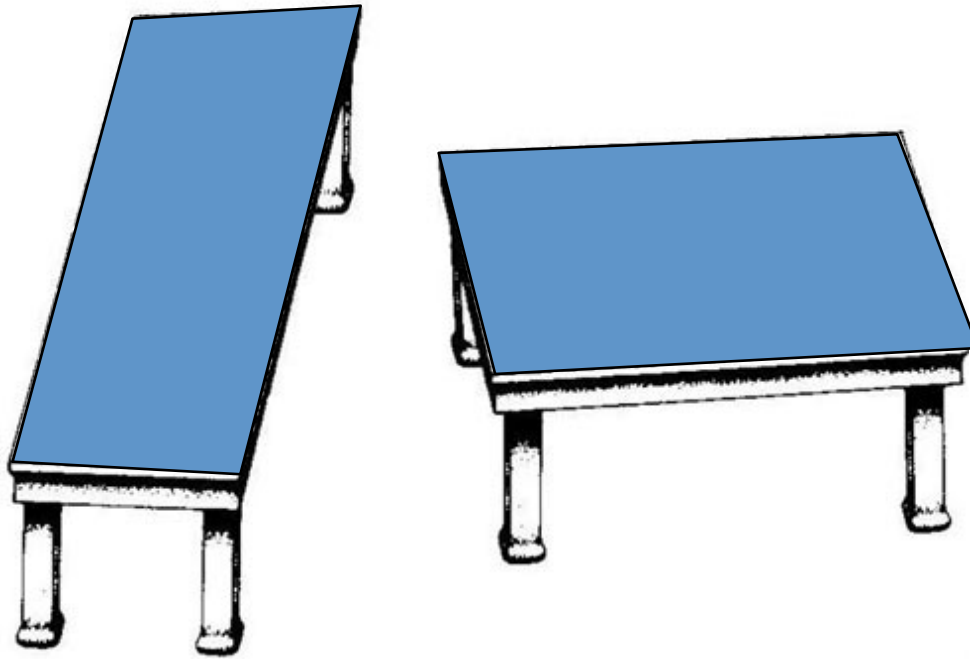


Measuring light vs. measuring scene properties



We perceive two squares, one on top of each other.

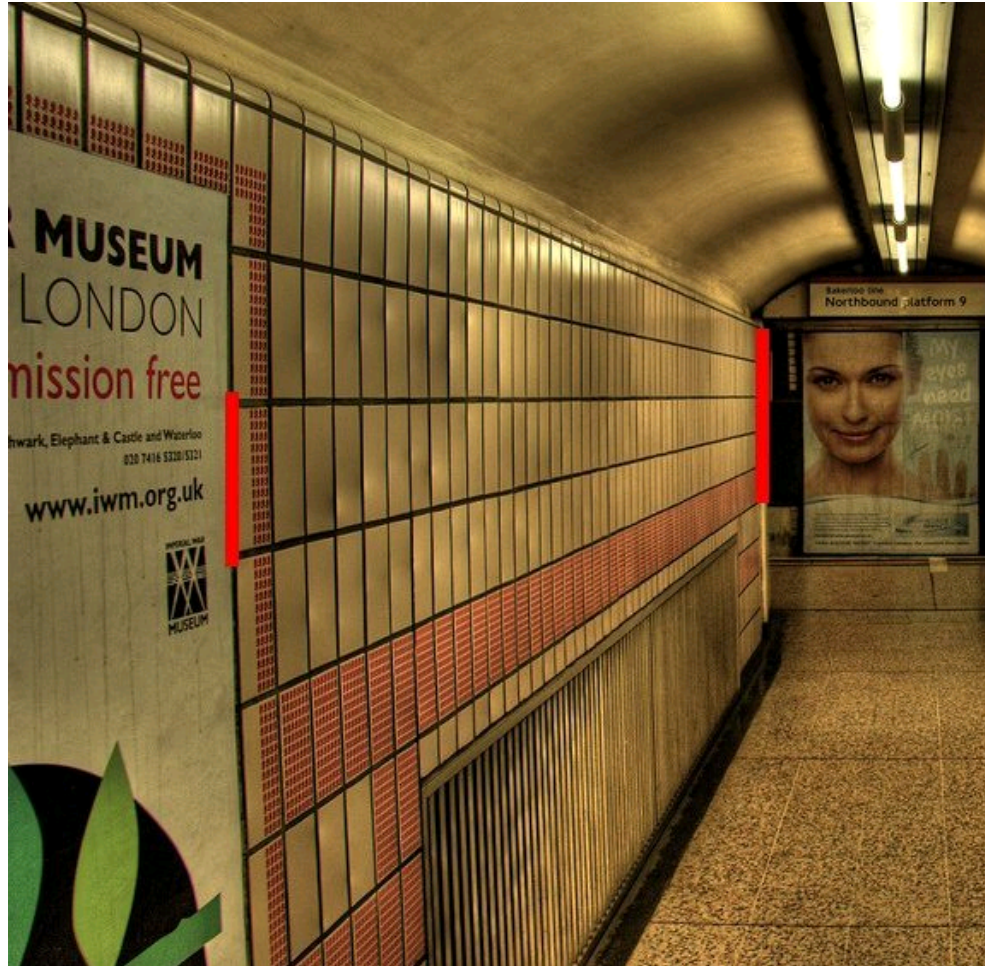
Measuring light vs. measuring scene properties



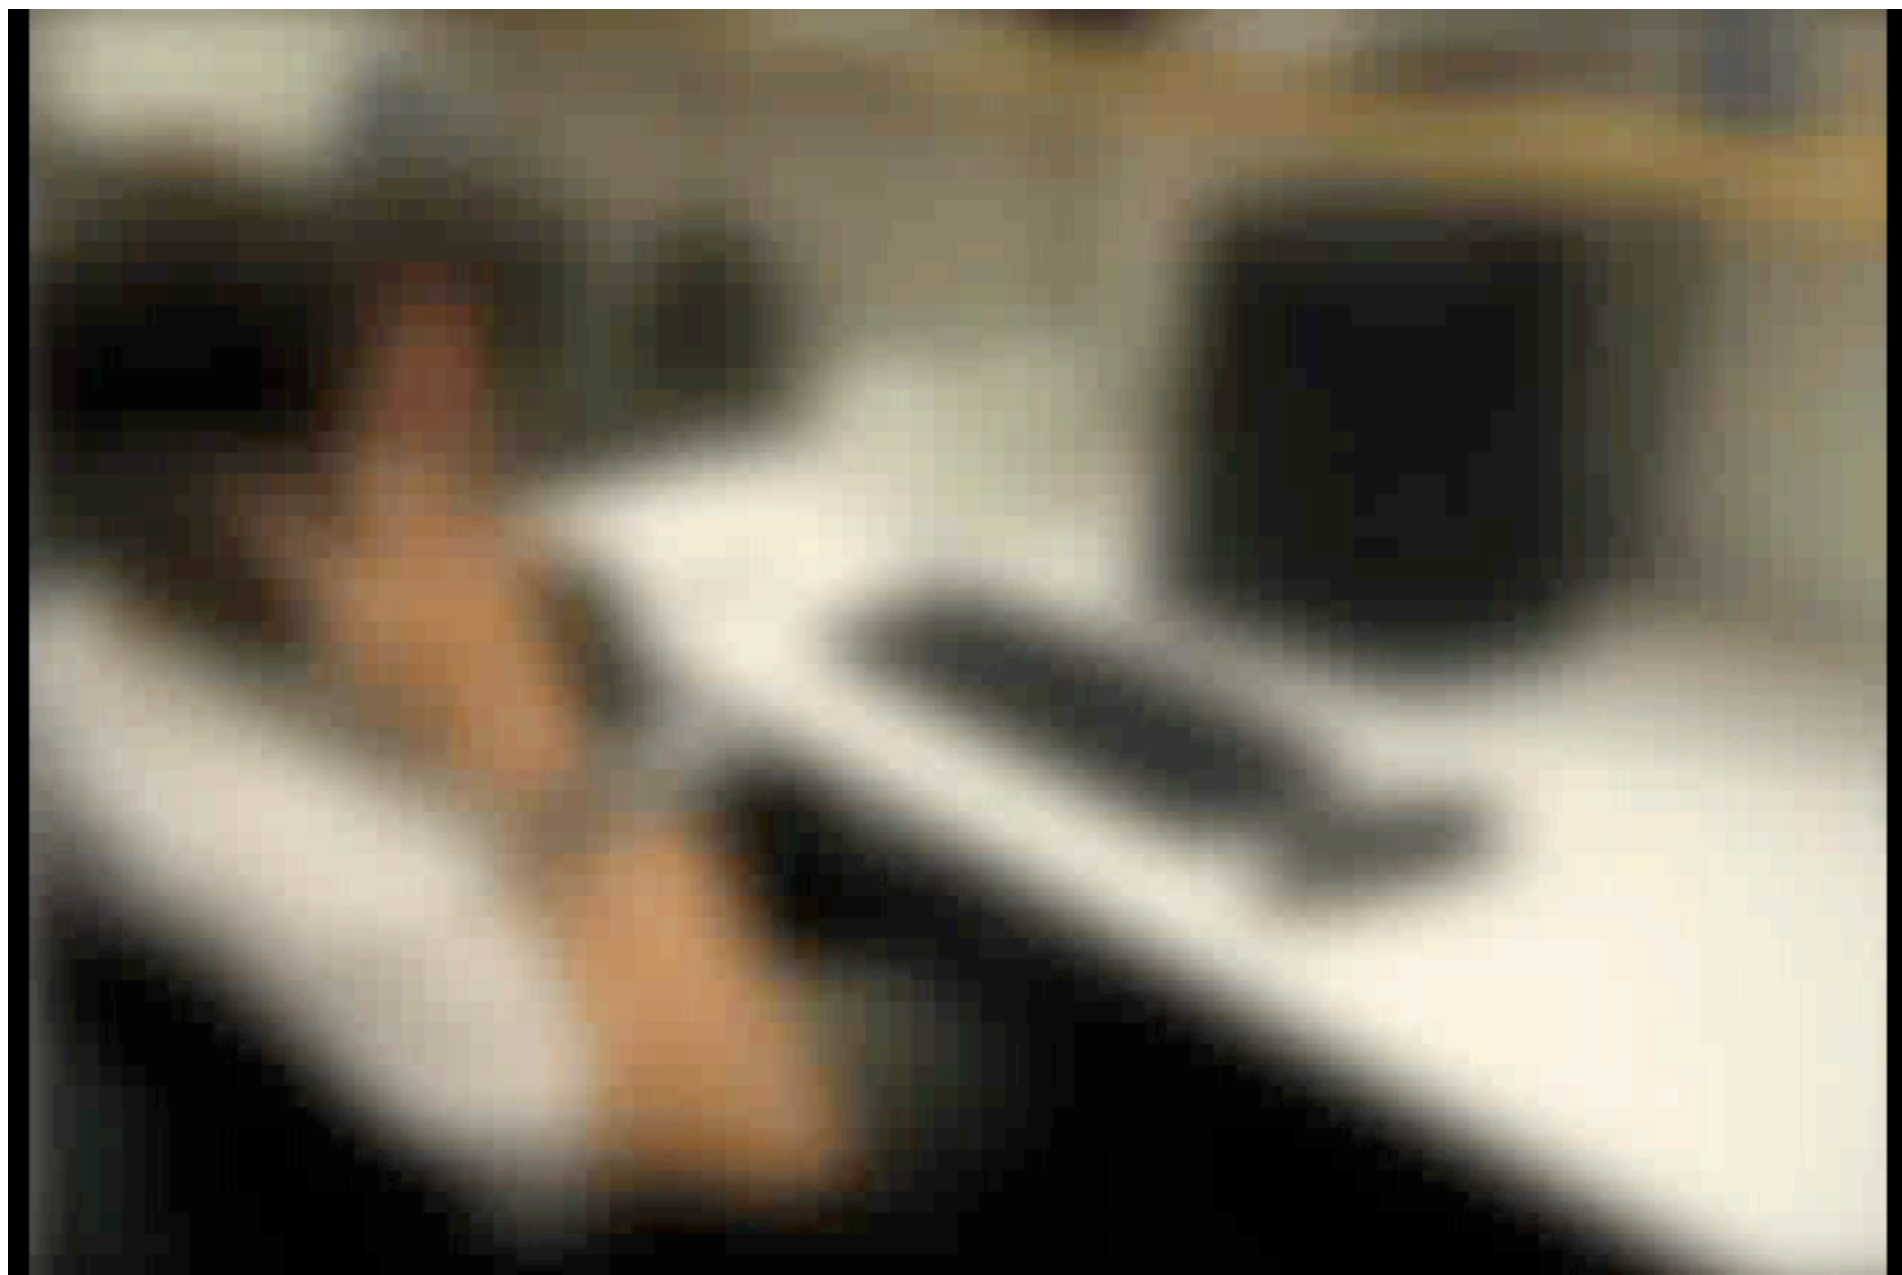
by Roger Shepard ("Turning the Tables")

Depth processing is automatic, and we can not shut it down...

Measuring light vs. measuring scene properties

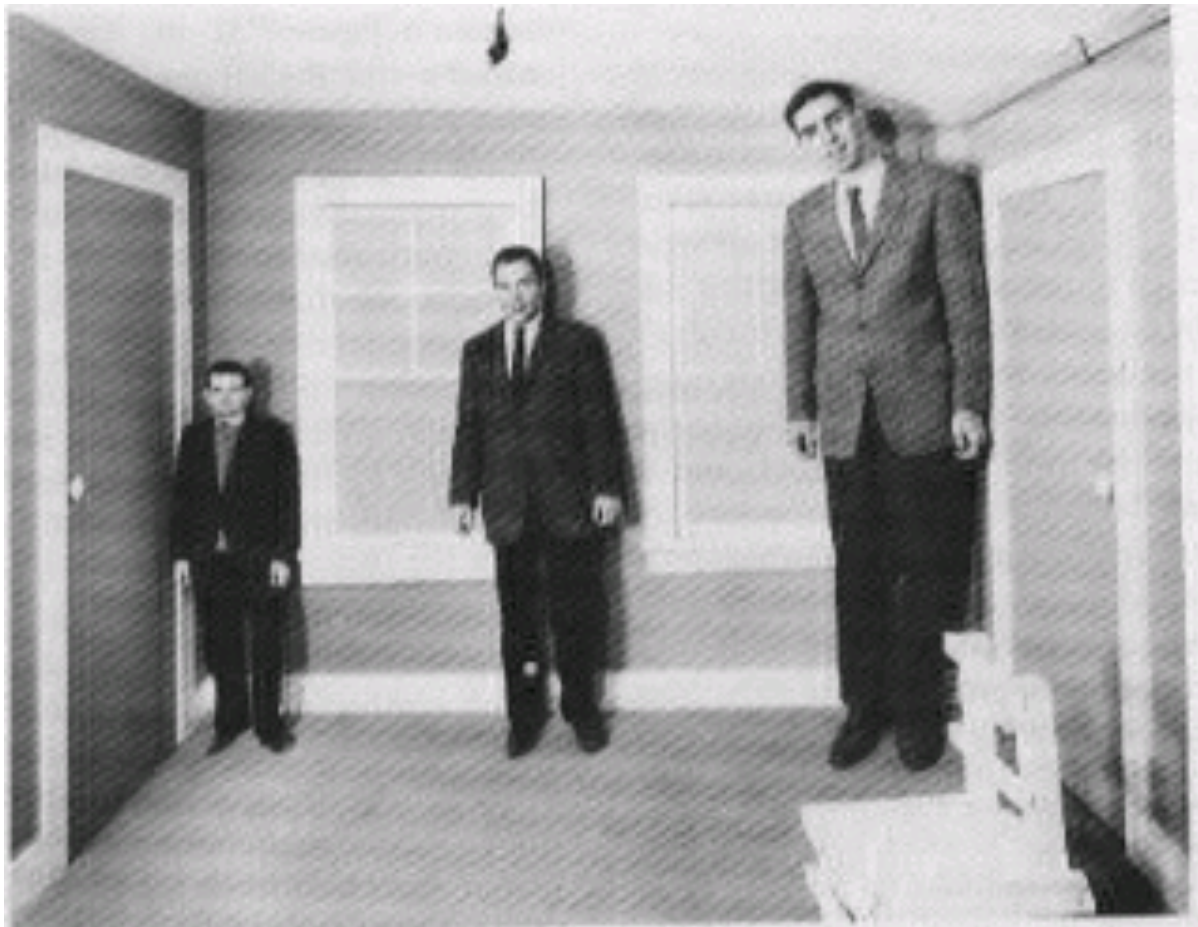


(c) 2006 Walt Anthony





Assumptions can be wrong



Ames room



By Aude Oliva

Vision has to solve an ill-posed problem

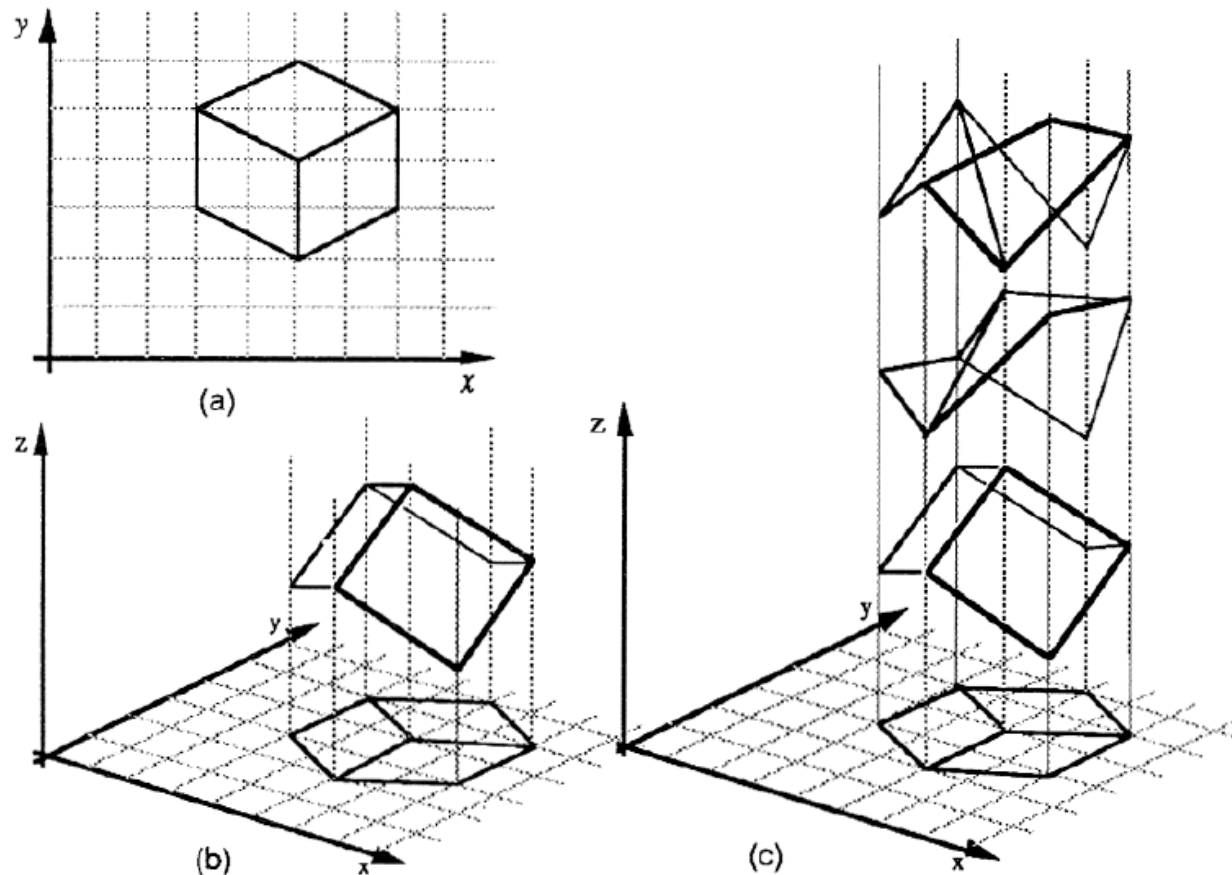
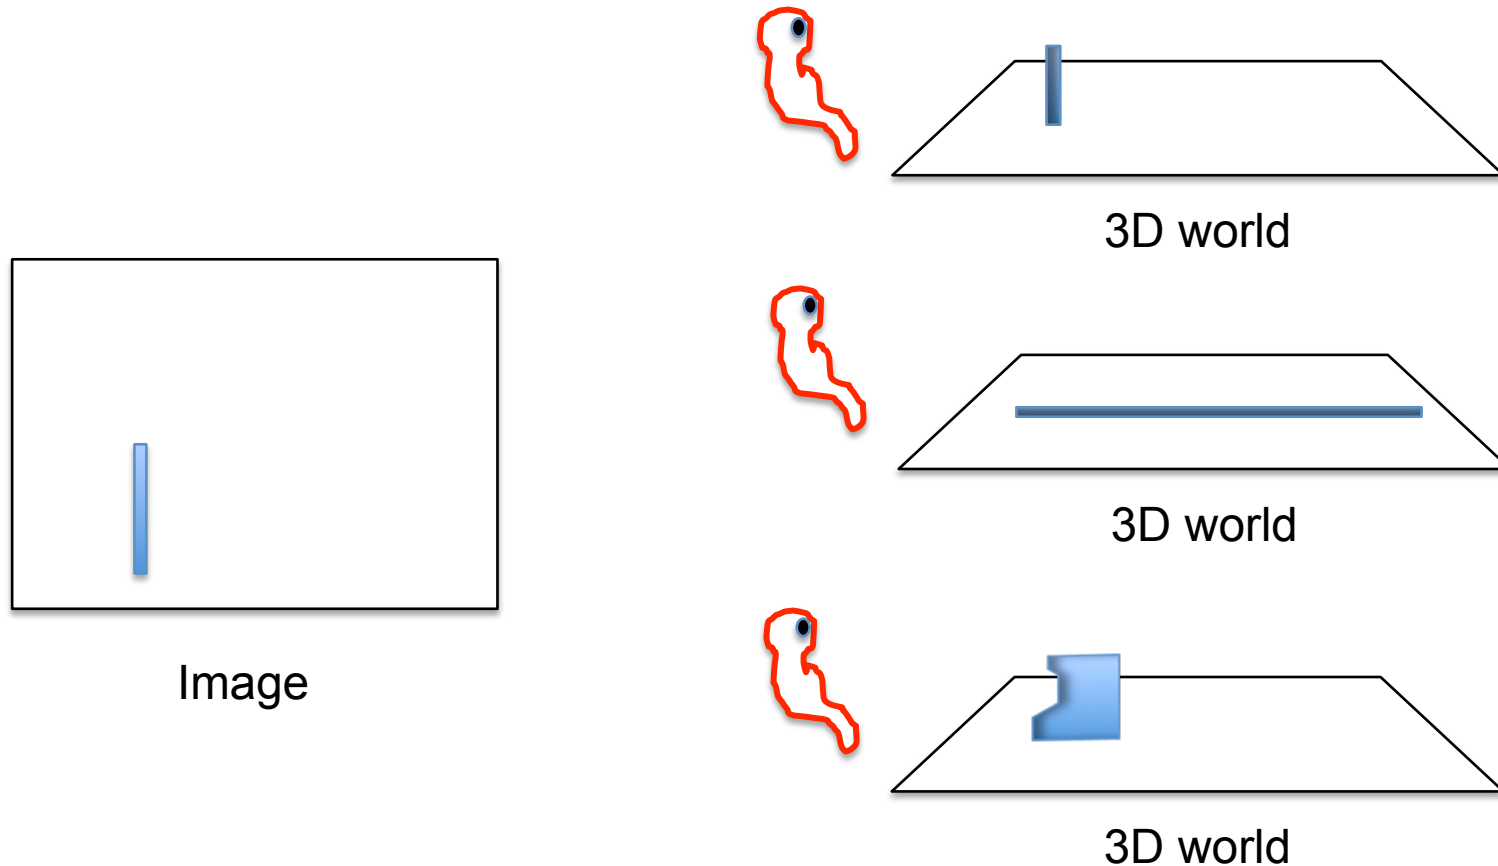


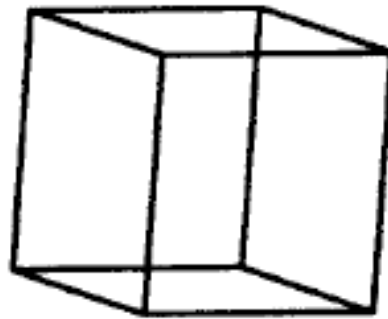
Figure 1. (a) A line drawing provides information only about the x, y coordinates of points lying along the object contours. (b) The human visual system is usually able to reconstruct an object in three dimensions given only a single 2D projection (c) Any planar line-drawing is geometrically consistent with infinitely many 3D structures.

Generic view assumption



Generic view assumption: the observer should not assume that he has a special position in the world... The most generic interpretation is to see a vertical line as a vertical line in 3D.

A simple idea to recover 3D shapes from line drawings



Task: Given the set of image coordinates of the vertices (x_i, y_i) , recover the world coordinates (X_i, Y_i, Z_i) .

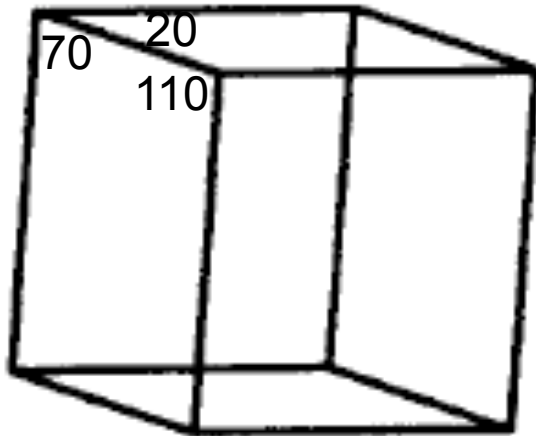
We will assume orthographic projection so that $X_i = x_i$ and $Y_i = y_i$

Then, the problem is: recover the missing depth Z_i

A simple idea to recover 3D shapes from line drawings

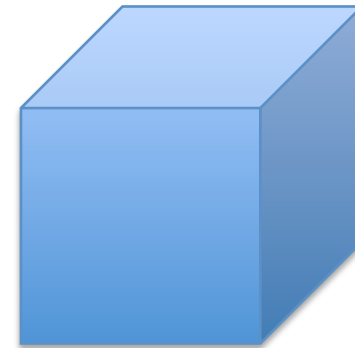
Heuristic: find Z_i that minimize the standard deviation of angles in the 3D object

Possible solution 1



$Z_i = 0$ for all i

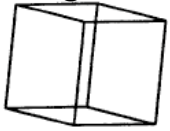
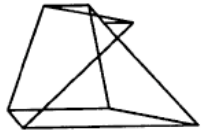
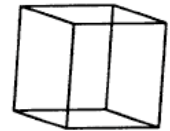
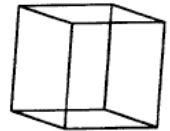
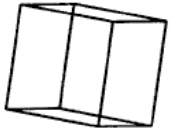
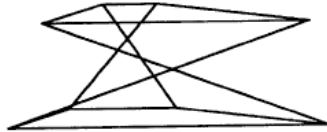
Possible solution 2



All 90 degree angles

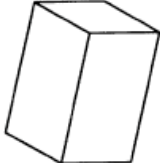
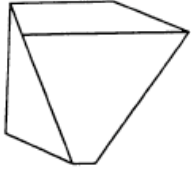
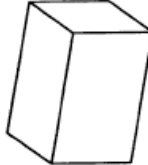
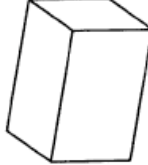
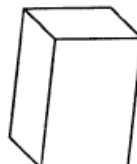

$$SDA = \sum (\alpha_{ij} - m)^2 \quad \text{with} \quad \alpha_{ij} = \cos^{-1}(u_i^T u_j)$$







A simple idea to recover 3D shapes from line drawings

| | | | | |
|-----------------------------|--|--|-----------------------------|--|
| Object rotated - 10 degrees | A | Reconstruction -10 degrees  | B |  |
| | C | INPUT  | D |  |
| | E | Reconstruction +10 degrees  | F |  |
| | Object created by CONSTRUCT program | | Arbitrary comparison object | |







Marill, AI-Memo-1136, 1989


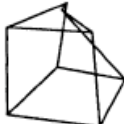
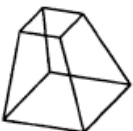

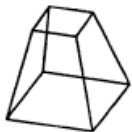

A simple idea to recover 3D shapes from line drawings

| | | | | |
|-----------------------------|-------------------------------------|---|-----------------------------|---|
| Object rotated - 10 degrees | A |  | B |  |
| | C |  | D |  |
| | E |  | F |  |
| | Object created by CONSTRUCT program | | Arbitrary comparison object | |

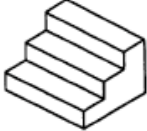

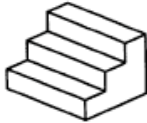
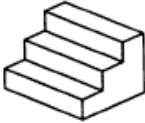
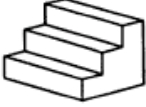

| | | | | |
|-----------------------------|-------------------------------------|---|-----------------------------|---|
| Object rotated - 10 degrees | A |  | B |  |
| | C |  | D |  |
| | E |  | F |  |
| | Object created by CONSTRUCT program | | Arbitrary comparison object | |

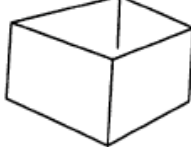
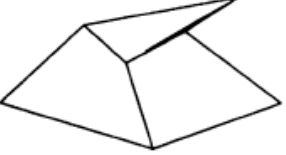
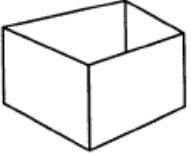
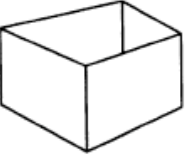
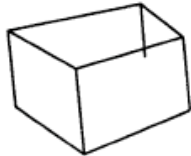
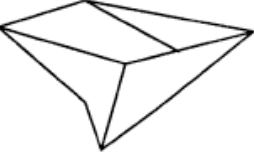
A simple idea to recover 3D shapes from line drawings

| | | |
|--|---|---|
| Object rotated - 10 degrees | A | B |
| |  |  |
| | C | D |
| Object rotated 0 degrees Input line-drawing |  |  |
| Object rotated +10 degrees | E | F |
| |  |  |
| | Object created by CONSTRUCT program | Arbitrary comparison object |

| | | |
|--|---|---|
| Object rotated - 10 degrees | A | B |
| |  |  |
| | C | D |
| Object rotated 0 degrees Input line-drawing |  |  |
| Object rotated +10 degrees | E | F |
| |  |  |
| | Object created by CONSTRUCT program | Arbitrary comparison object |

A simple idea to recover 3D shapes from line drawings

| | | | | |
|-----------------------------|-------------------------------------|---|-----------------------------|---|
| Object rotated - 10 degrees | A |  | B |  |
| | C |  | D |  |
| | E |  | F |  |
| | Object created by CONSTRUCT program | | Arbitrary comparison object | |

| | | | | |
|-----------------------------|-------------------------------------|---|-----------------------------|---|
| Object rotated - 10 degrees | A |  | B |  |
| | C |  | D |  |
| | E |  | F |  |
| | Object created by CONSTRUCT program | | Arbitrary comparison object | |

Where is now Computer Vision?

Application of statistical image model, and variational Bayesian inference: removing motion blur

Original



Variational Bayes



Fergus et al, 2006

Close-up

Original



Naïve Sharpening



Variational Bayes



Texture synthesis

Input

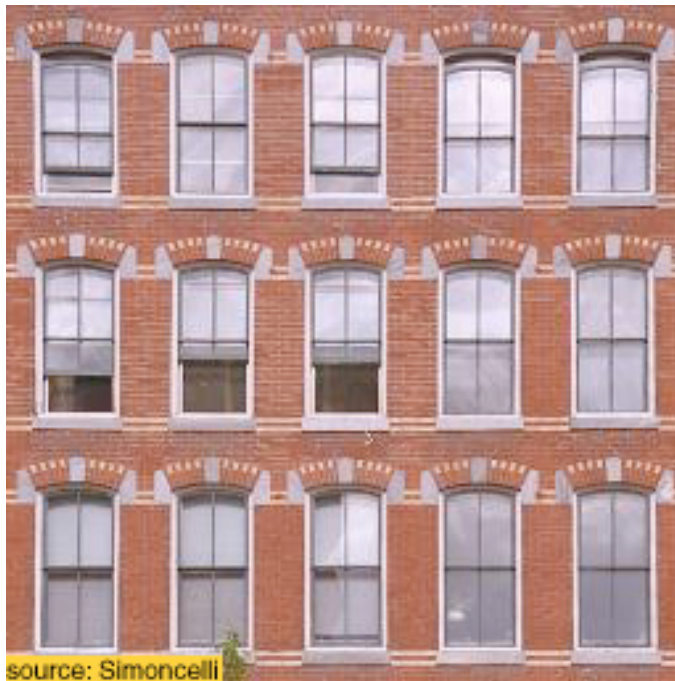


Output: new instances of the same "kind" of texture.



Texture synthesis

Input

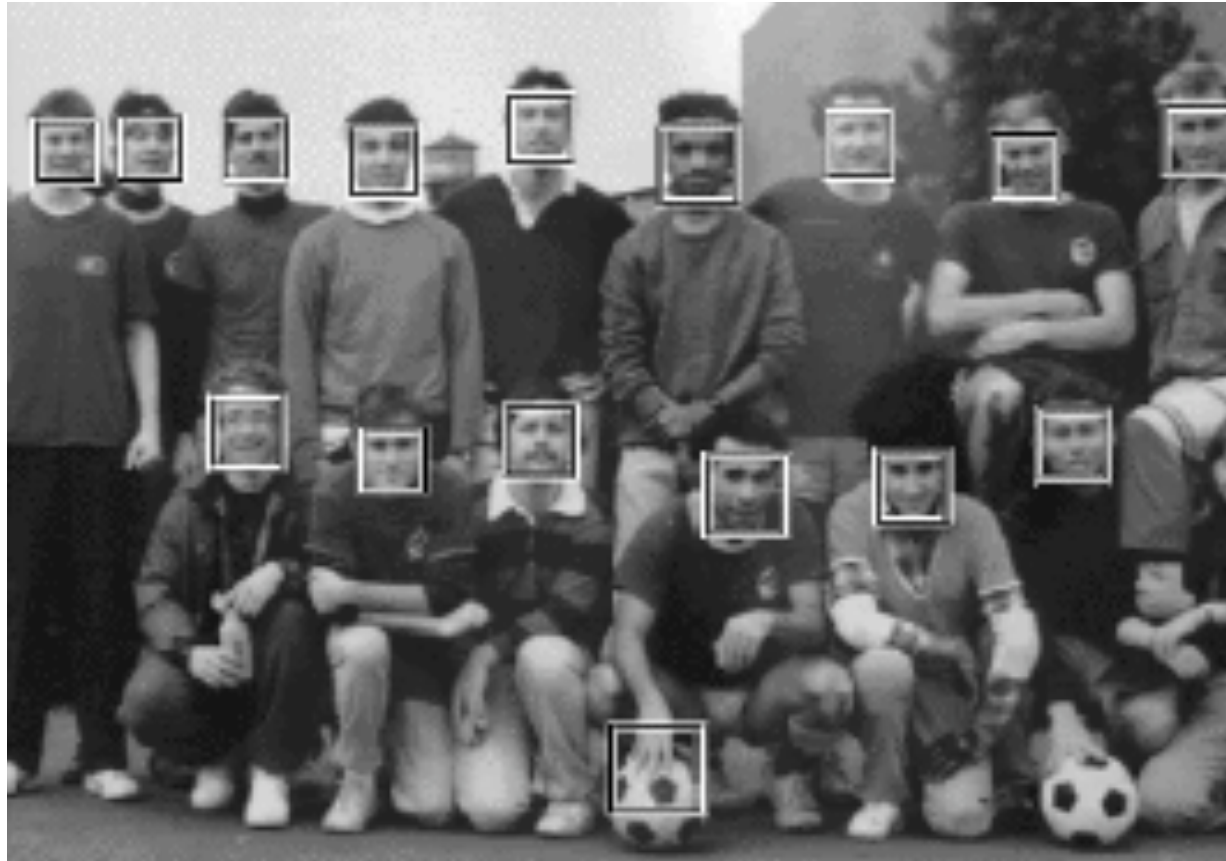


Output: new instances of the same "kind" of texture.



Portilla & Simoncelli, 1999

2D frontal face detection



Amazing how far they have gotten with so little...

Face detection

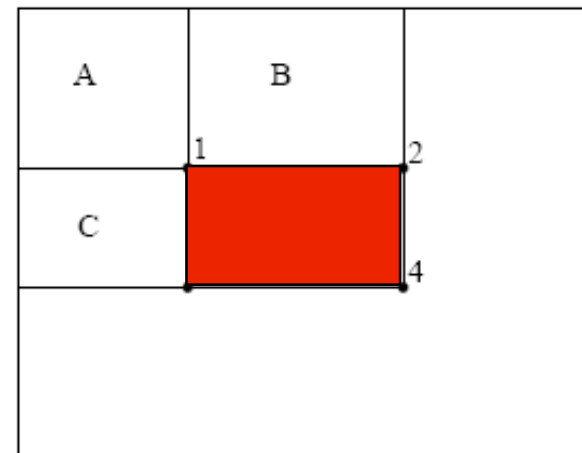
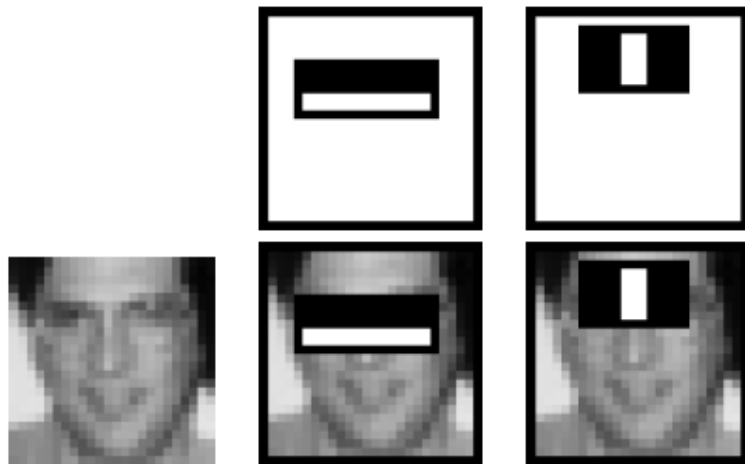


[Face priority AE] When a bright part of the face is too bright

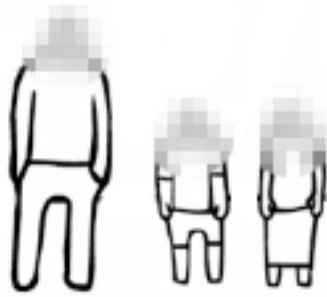
Face detection

Haar filters and integral image

Viola and Jones, ICCV 2001



The average intensity in the block is computed with four sums independently of the block size.



Google street view

[Web](#) [Images](#) [Maps](#) [News](#) [Shopping](#) [Mail](#) [more](#) ▼

[philipp.lenssen@gmail.com](#) | [My Profile](#) | [Help](#) | [Web History](#) | [My Account](#) | [Sign out](#)



e.g., "10 market st, san francisco" or "hotels near lax"

1504 broadway, nyc

Search Maps

[Show search options](#)

[Search the map](#)

[Find businesses](#)

[Get directions](#)

Search Results

My Maps

[Print](#) [Send](#) [Link to this page](#)

1504 Broadway
New York, NY 10036

[Make this my default location](#)

[Explore this area »](#)

Photos



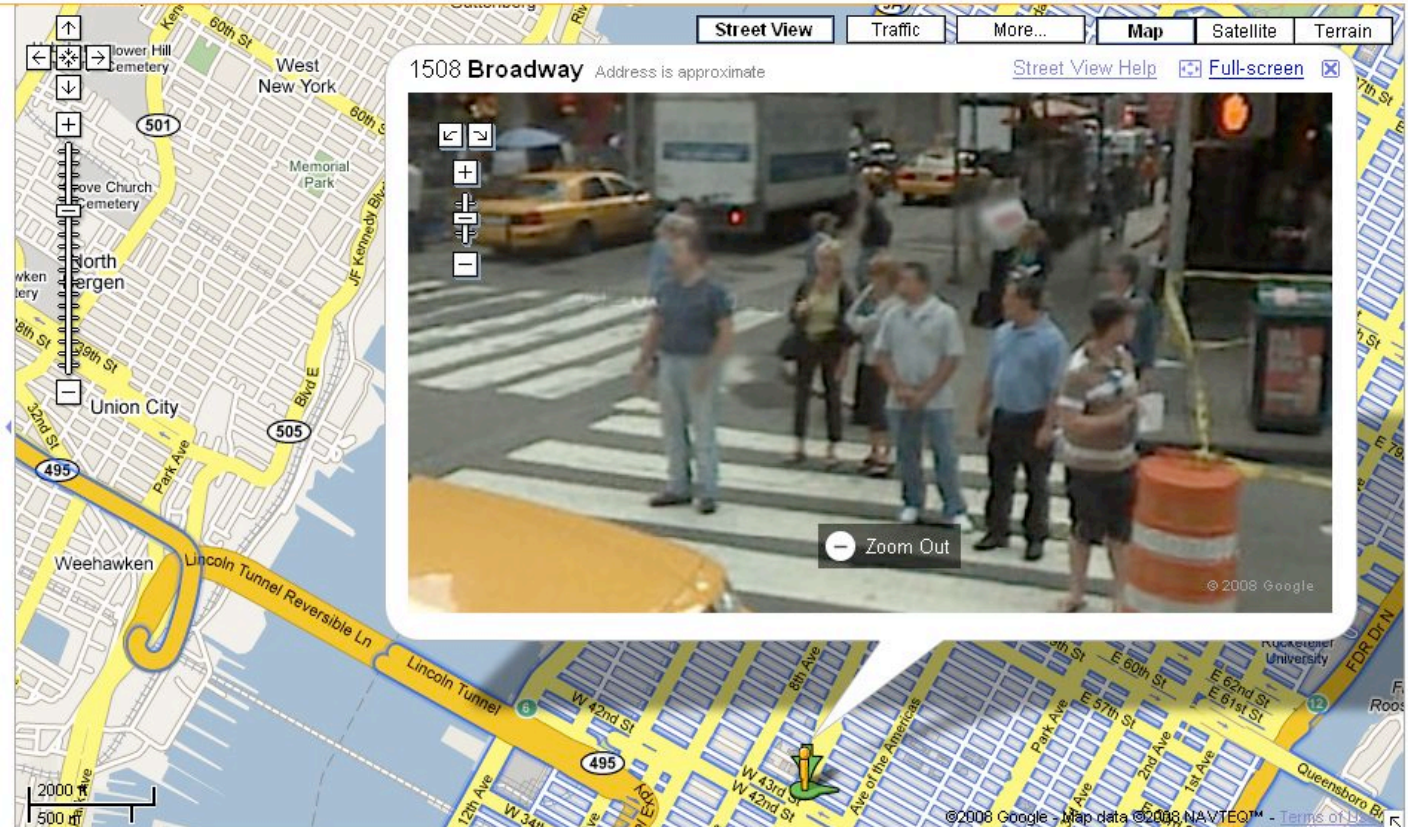
User-Created Maps

[Cheap Eats in NYC](#)

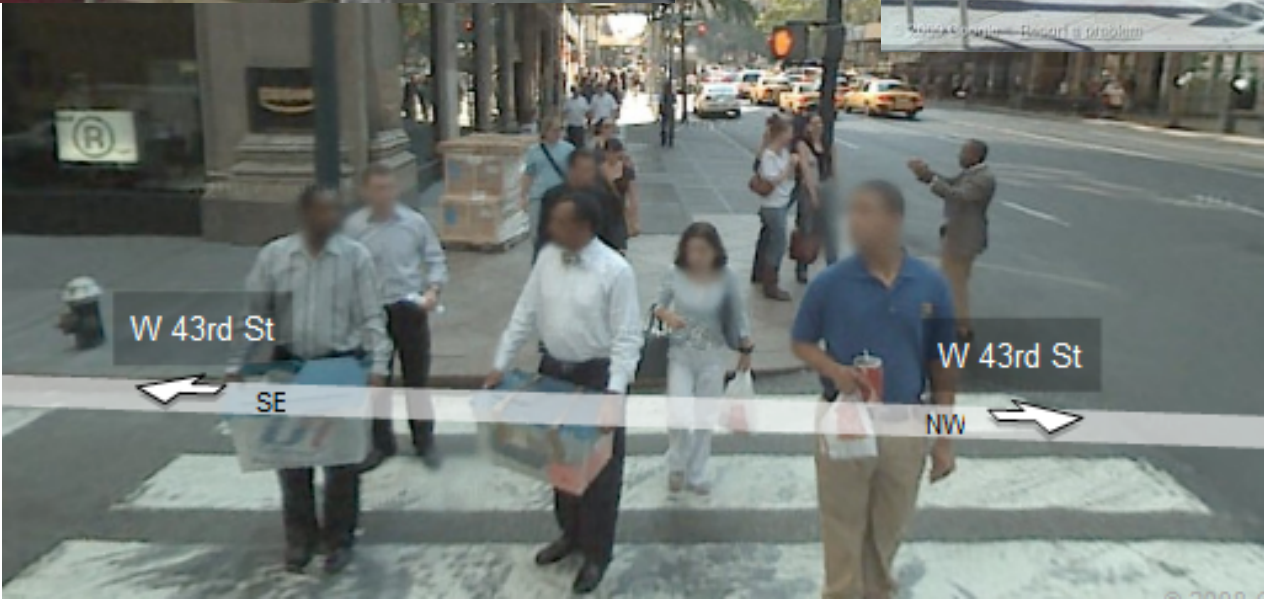
[North American New Humanist Forum](#)

[Barrel to Bottles](#)

[More photos, videos, and user-created maps »](#)

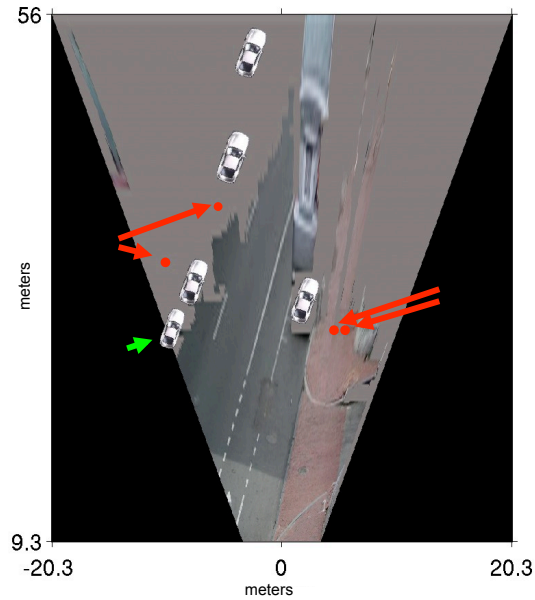


Google street view

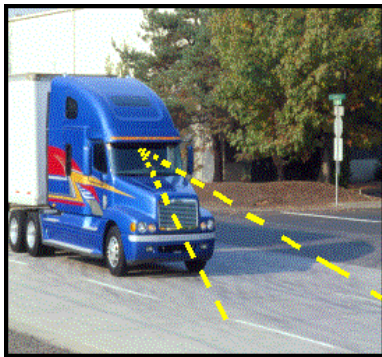


Assisted driving

Pedestrian and car detection

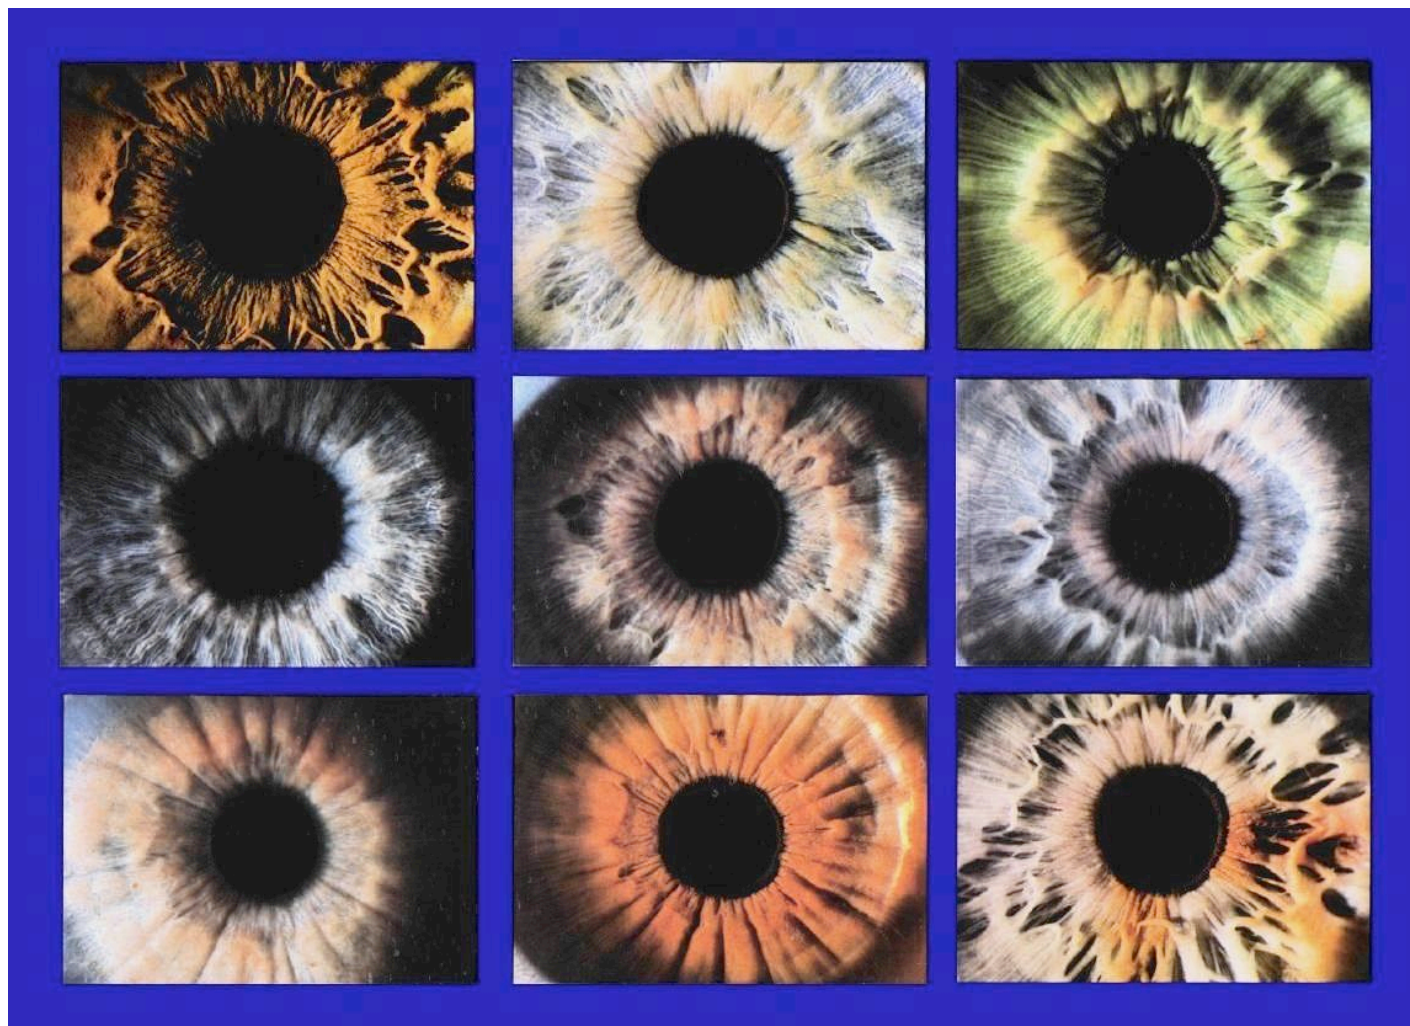
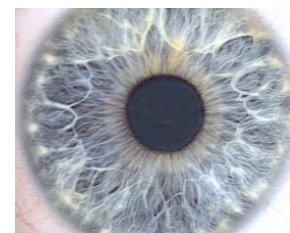


Lane detection



- Collision warning systems with adaptive cruise control,
- Lane departure warning systems,
- Rear object detection systems,

Iris recognition



<http://www.cl.cam.ac.uk/~jgd1000/iriscollage.jpg>

JOHN DAUGMAN

Iris recognition

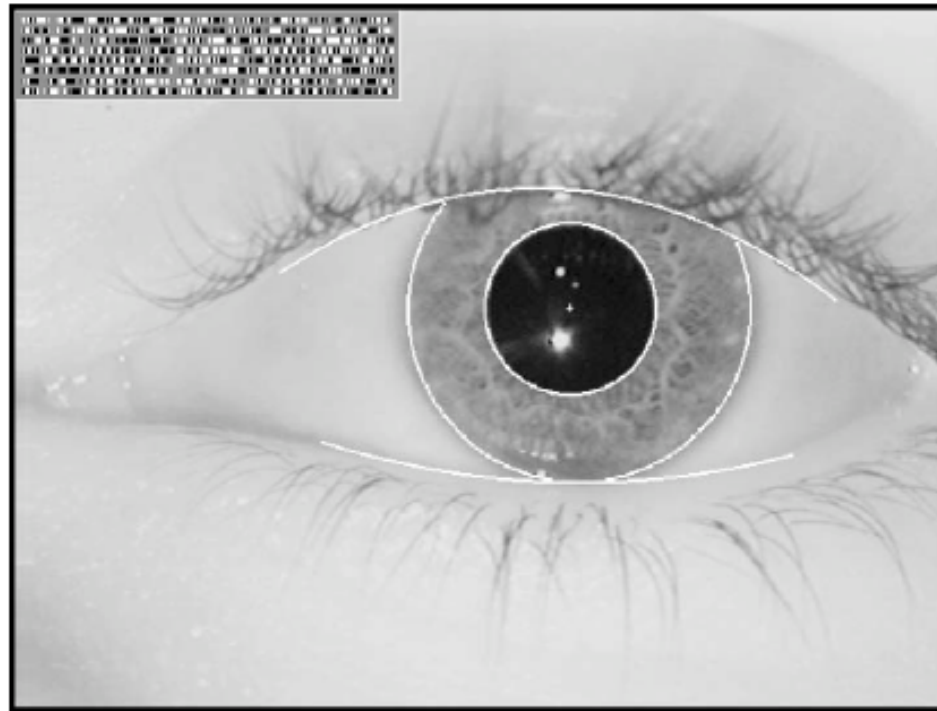


Fig. 1. Example of an iris pattern, imaged monochromatically at a distance of about 35 cm. The outline overlay shows results of the iris and pupil localization and eyelid detection steps. The bit stream in the top left is the result of demodulation with complex-valued two-dimensional (2-D) Gabor wavelets to encode the phase sequence of the iris pattern.

Genetically identical eyes also have different iris

JOHN DAUGMAN, 93

Iris recognition immigration system



<http://www.ind.homeoffice.gov.uk/managingborders/technology/iris/>

SIFT



Figure 12: The training images for two objects are shown on the left. These can be recognized in a cluttered image with extensive occlusion, shown in the middle. The results of recognition are shown on the right. A parallelogram is drawn around each recognized object showing the boundaries of the original training image under the affine transformation solved for during recognition. Smaller squares indicate the keypoints that were used for recognition.

SIFT vector formation

- Thresholded image gradients are sampled over 16x16 array of locations in scale space
- Create array of orientation histograms
- 8 orientations x 4x4 histogram array = 128 dimensions

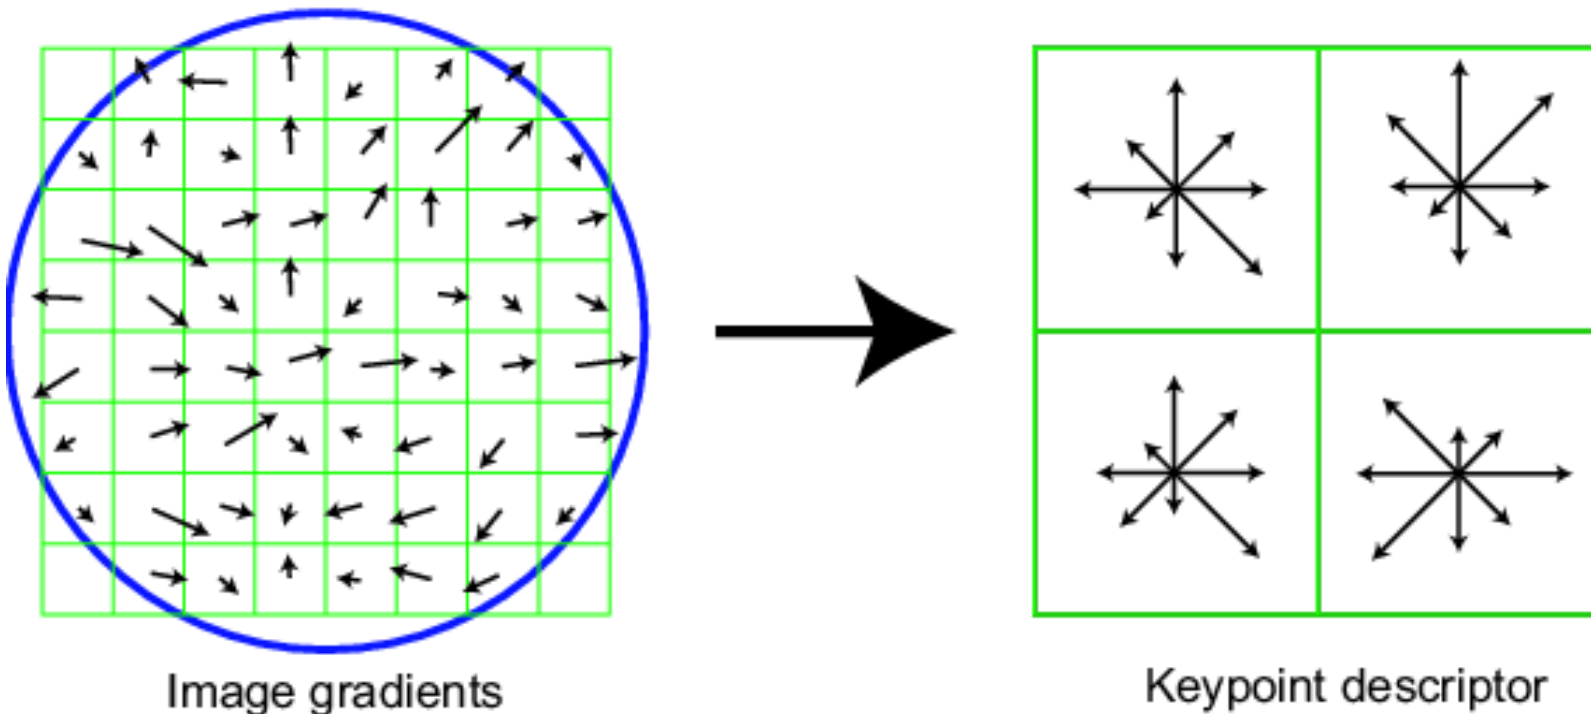
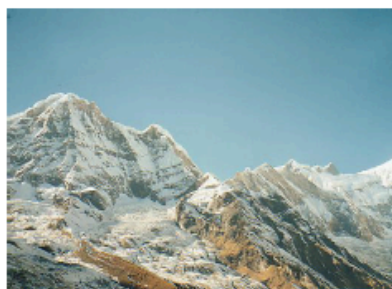


Image stitching



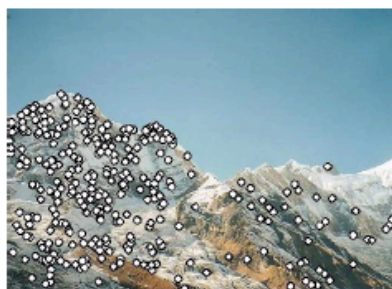
(a) Image 1



(b) Image 2



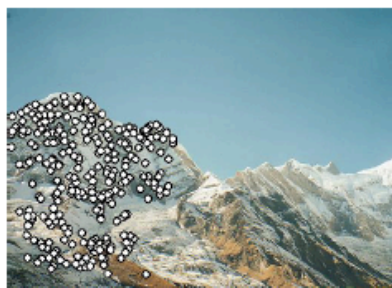
(c) SIFT matches 1



(d) SIFT matches 2



(e) RANSAC inliers 1



(f) RANSAC inliers 2

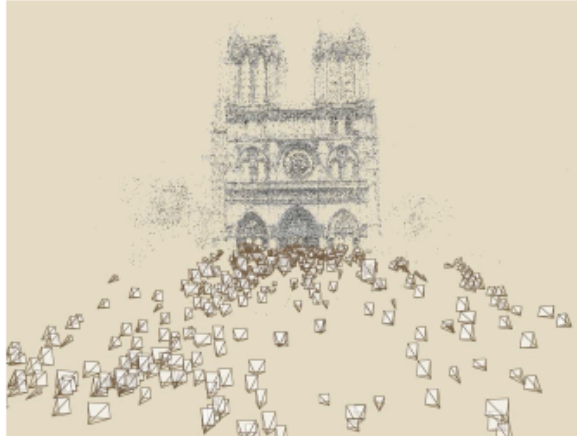


(g) Images aligned according to a homography

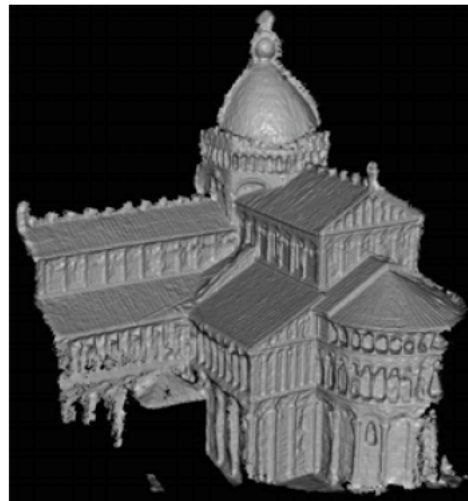
Brown, Lowe, 2007

Photo tourism

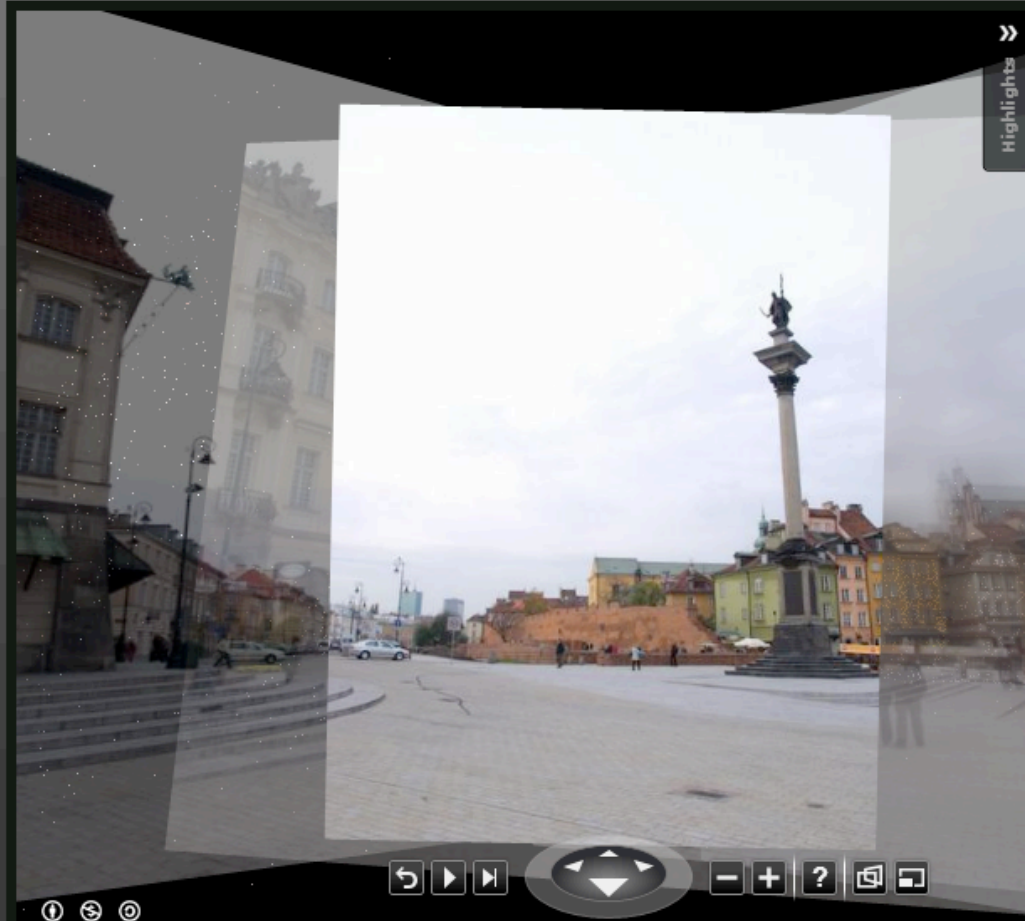
PhotoSynth



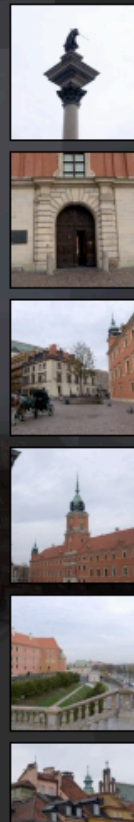
Snaveley et al. 2006



(Goesele et al. 2007).



Highlights



Royal Castle, Warsaw, Poland

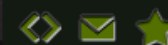
Mrozon 10/27/2008

300 Photos - 100% Synth

2226 Views

poland "royal castle" warsaw

See this synth on Bing Maps. (?)



Report Abuse



Figure 1: Top row: Two frames showing the same scene from very different camera viewpoints (from the film 'Run Lola Run'). Middle row: frames with detected affine invariant regions superimposed. 'Maximally Stable' (MS) regions are in yellow. 'Shape Adapted' (SA) regions are in cyan. Bottom row: Final matched regions after indexing and spatial consensus. Note that the correspondences define the scene overlap between the two frames.

eo

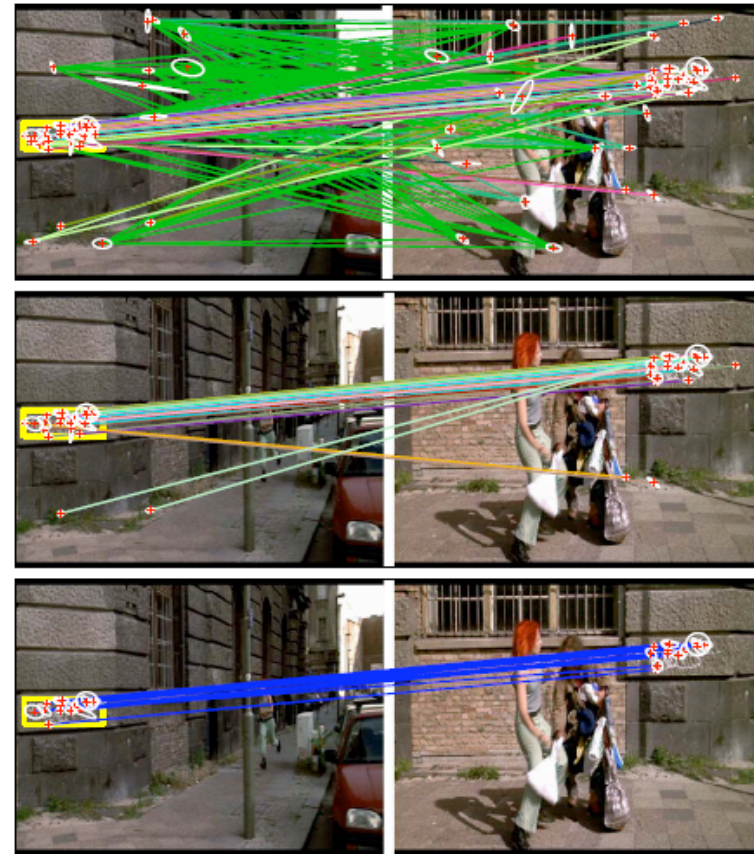


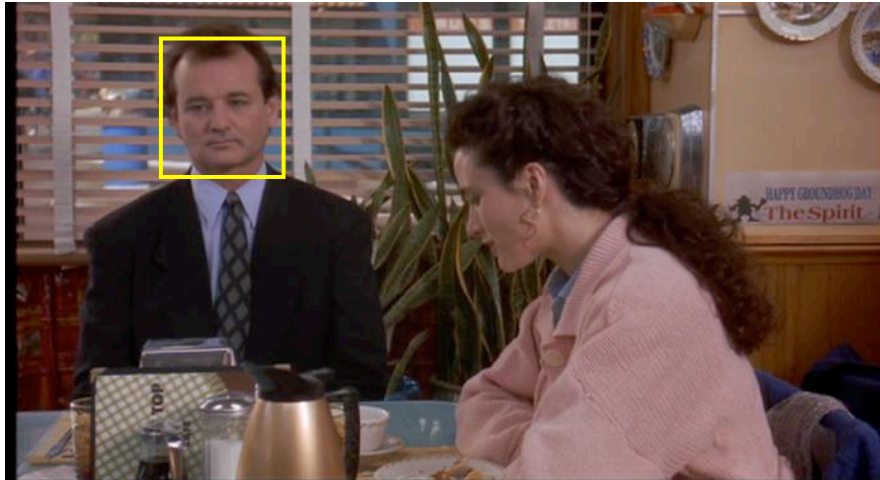
Figure 6: Matching stages. Top row: (left) Query region and (right) its close-up. Second row: Original word matches. Third row: matches after using stop-list, Last row: Final set of matches after filtering on spatial consistency.

Sivic, J. and Zisserman, A. Video Google: A Text Retrieval Approach to Object Matching in Videos
 Proceedings of the International Conference on Computer Vision (2003)

Visually defined search

Given an object specified by its image, retrieve all shots containing the object:

- must handle viewpoint change etc
- must be efficient at run time



objects



places



people

Object search in video: why is it hard?

- an object's imaged appearance varies ...

- scale changes



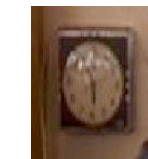
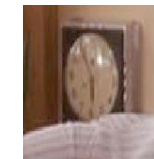
- lighting changes



- viewpoint changes



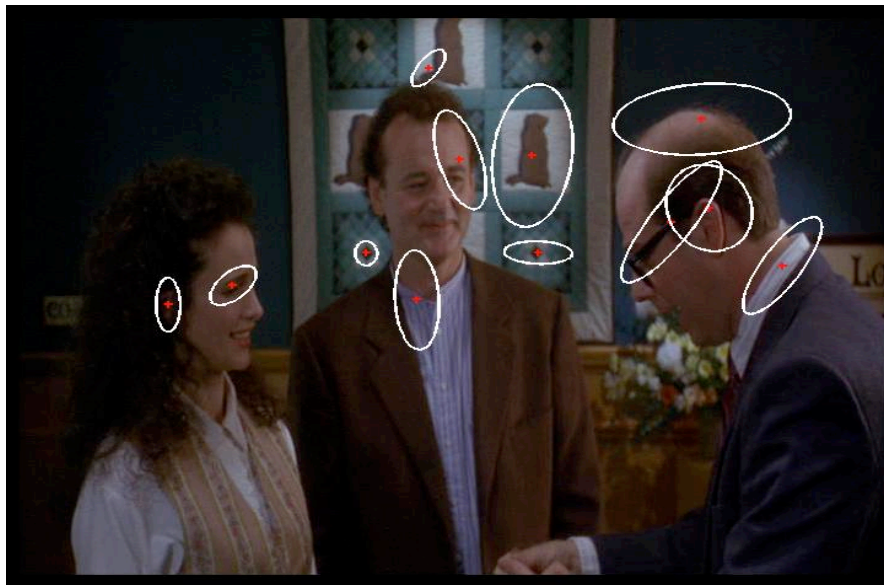
- partial occlusion



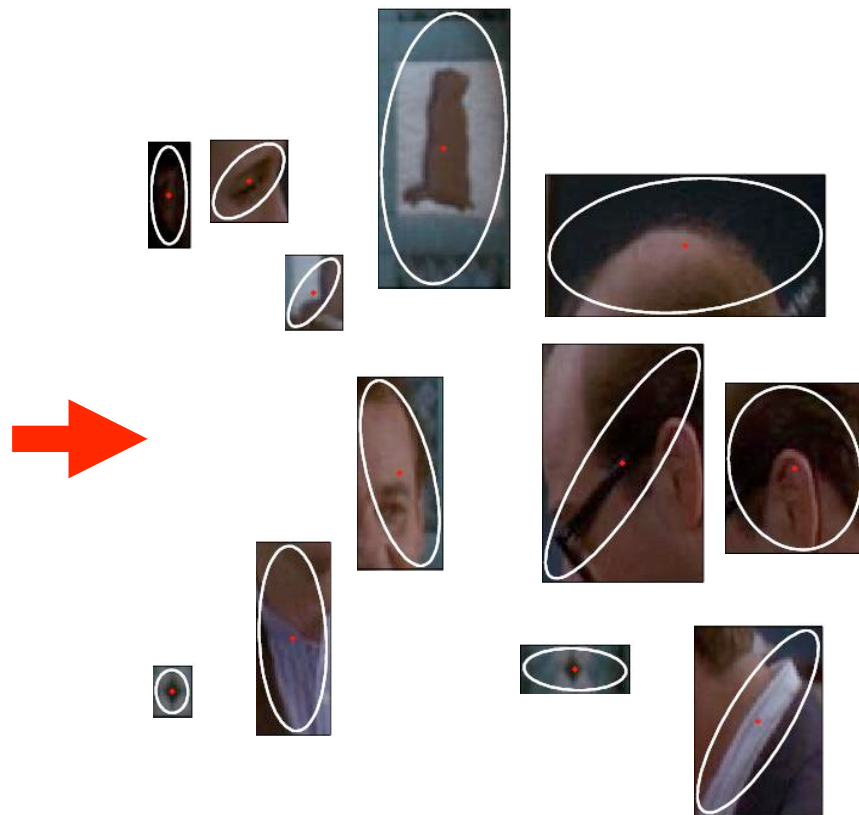
- sheer amount of data

- feature length movie ~ 100,000 -150,000 frames

Visual description – visual words

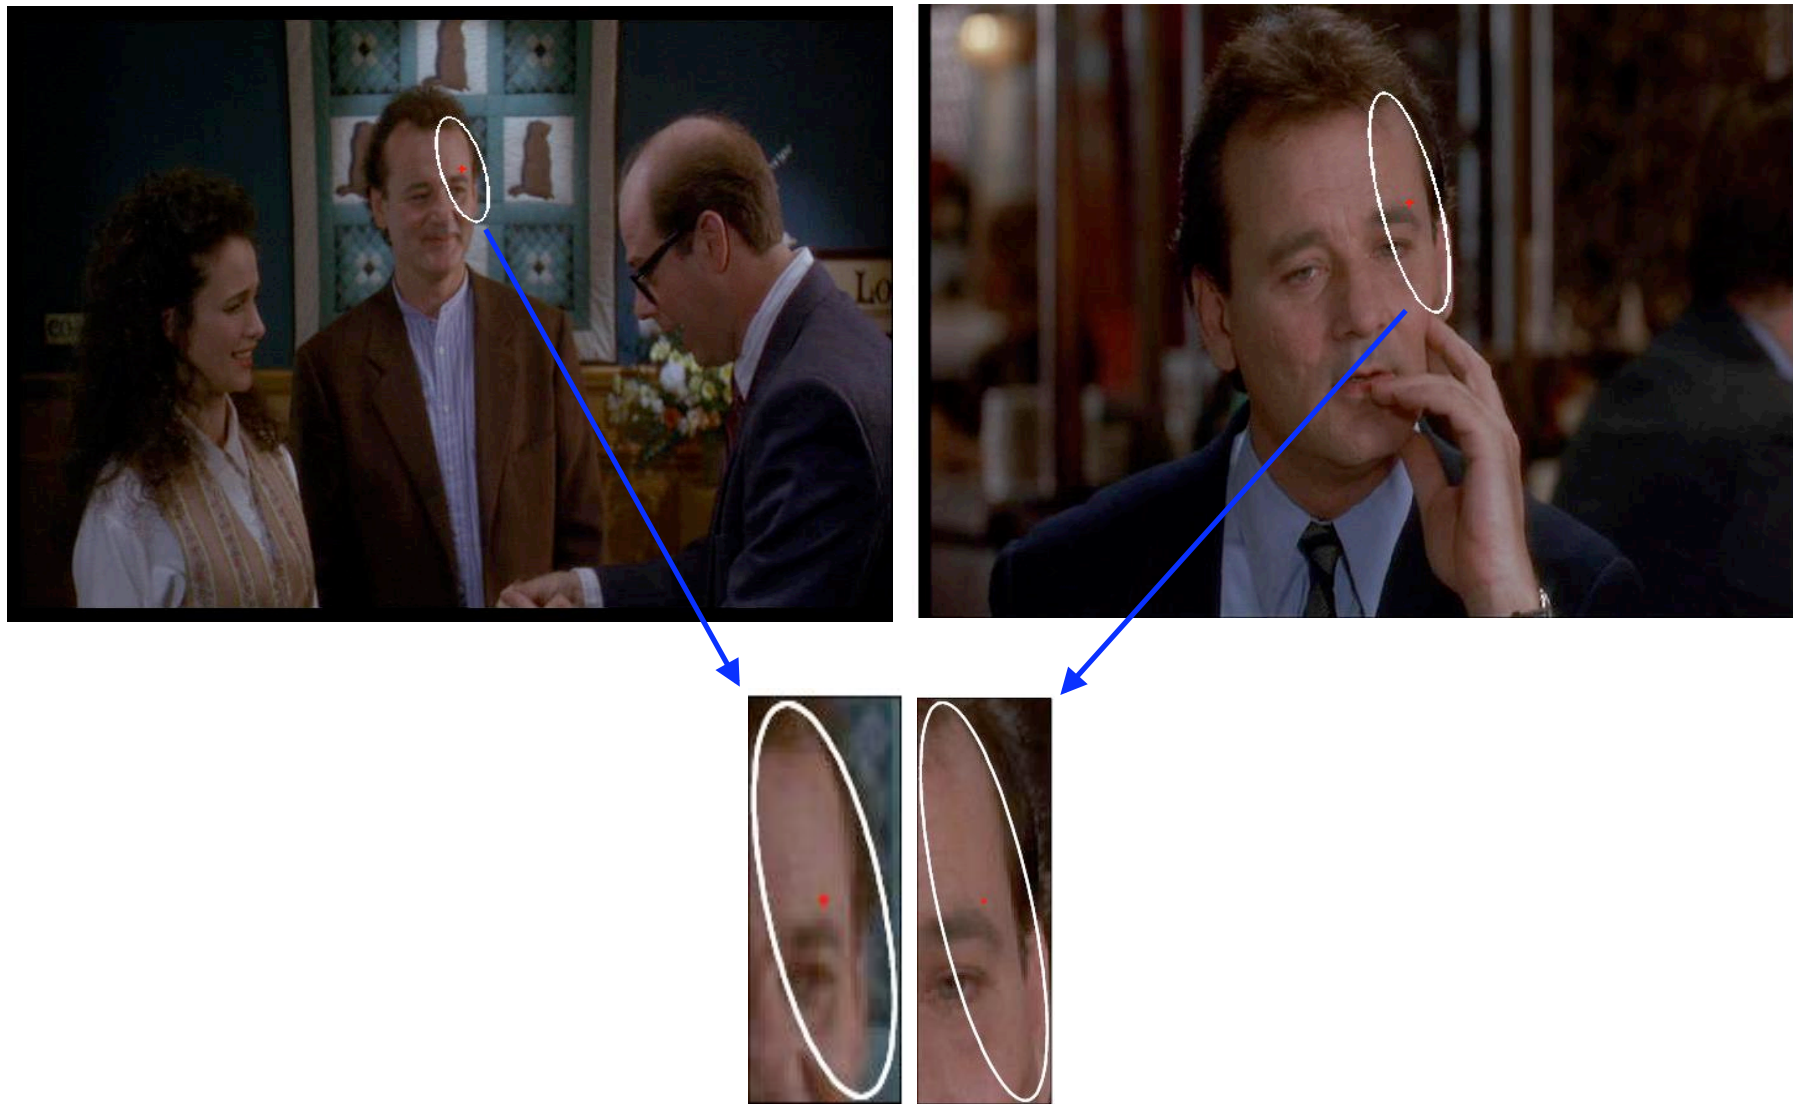


Image



visual nouns

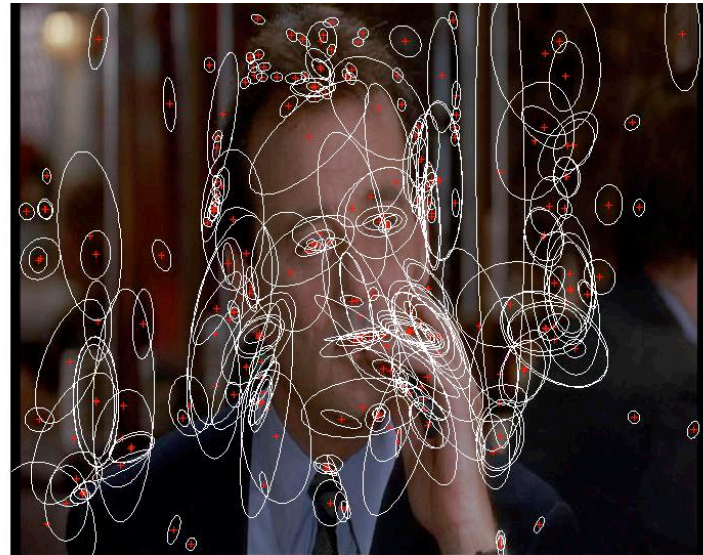
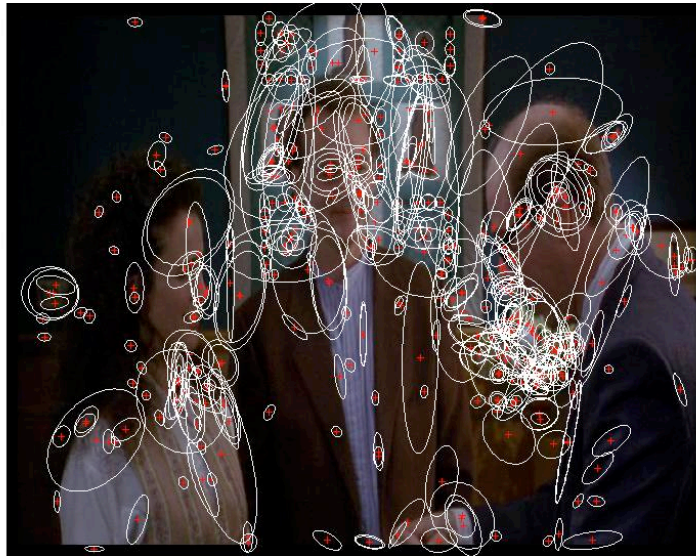
Visual vocabulary unaffected by scale and viewpoint



The same visual word

Slide by Josef Sivic

Image representation using visual words



Use efficient google like search on visual words

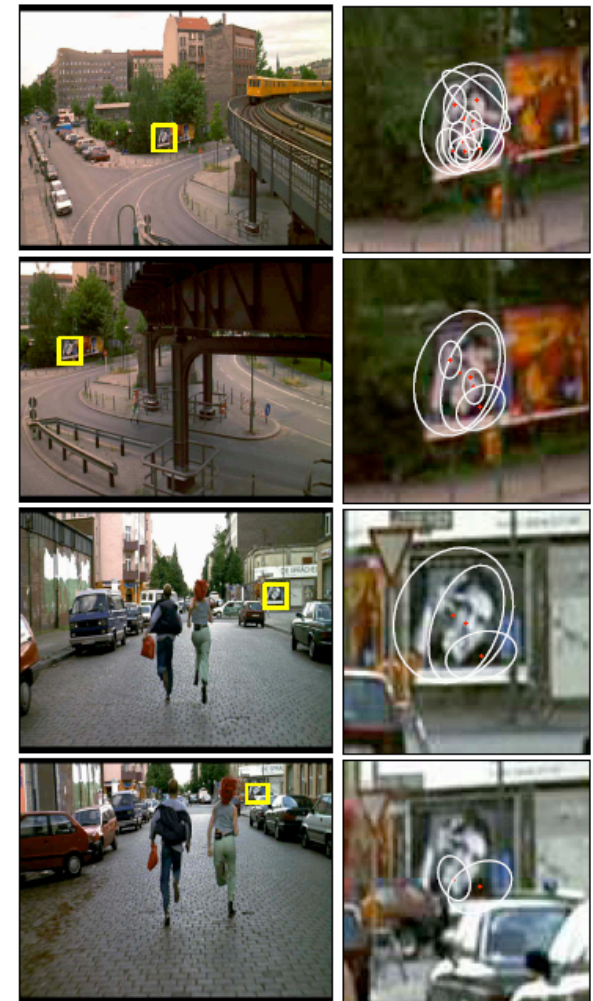
Video Google

Efficient search: In a classical file structure all words are stored in the document they appear in. An inverted file structure has an entry (hit list) for each word where all occurrences of the word in all documents are stored. In our case the inverted file has an entry for each visual word, which stores all the matches, i.e. occurrences of the same word in all frames. The document vector is very sparse and use of an inverted file makes the retrieval very fast. Querying a database of 4k frames takes about 0.1 second with a Matlab implementation on a 2GHz pentium.

Query:



Retrieved frames:



Example : Groundhog Day



retrieved shots

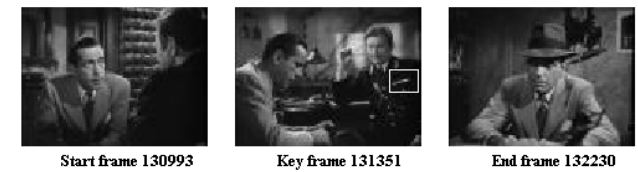
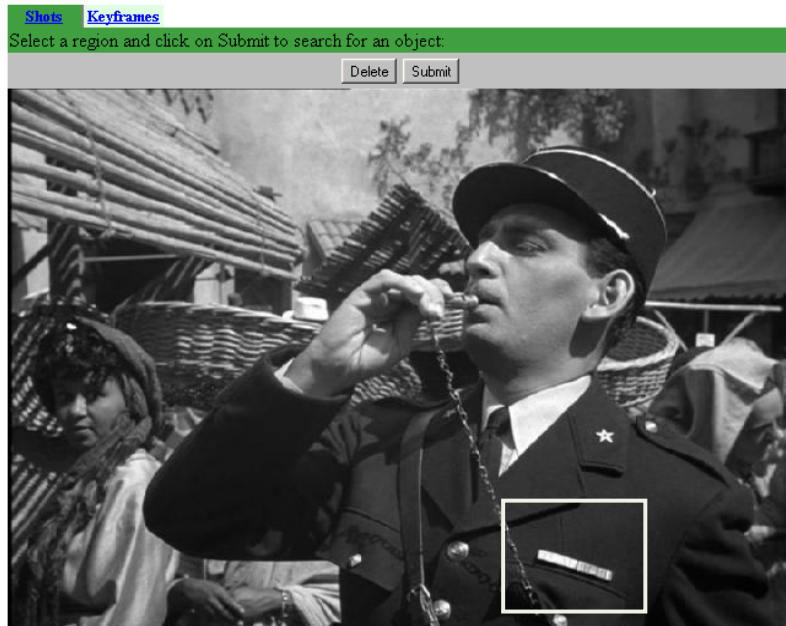


Video Google, Sivic & Zisserman, ICCV 2003

Slide by Josef Sivic

Example: Casablanca

retrieved shots



Slide by Josip Sivic

Improving online search

flickr

webshots

Ask Images
Cydral
Image & Site Search

picsearch

Google
Image Search

altavista













Query:
STREET

Google [vdeo](#) [images](#) [video](#) [news](#) [maps](#) [more »](#)

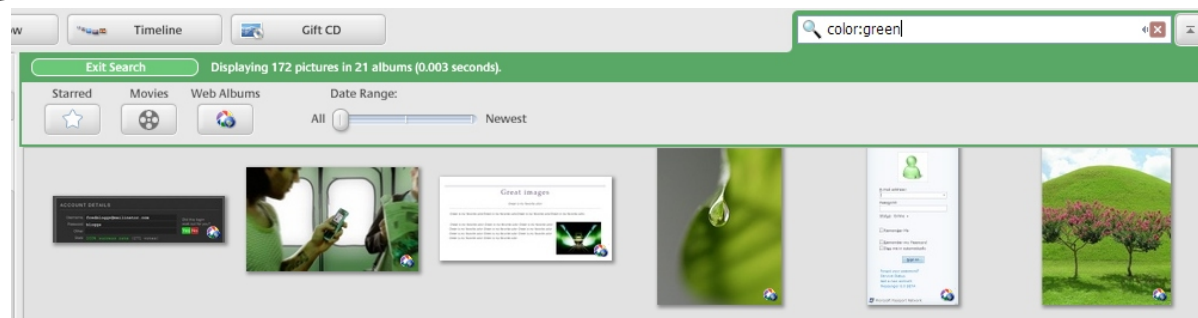
street [Advanced Image Search](#) [Preferences](#)

Moderate SafeSearch is on

Images Showing: All image sizes Results 19 - 36 of about 44,200,000 for street [definition] (0.04 seconds)

| | | | | | |
|---|--|--|---|--|---|
|  <p>Street sweeper 345 x 362 - 17k - jpg www.town.telluride.co.us</p> |  <p>Street Maintenance 407 x 402 - 18k - jpg www.town.telluride.co.us</p> |  <p>Main Street Station 360 x 392 - 30k - jpg www.rmaonline.org</p> |  <p>SHPO Wayne Donaldson at Main Street ... 410 x 314 - 41k - jpg ohp.parks.ca.gov</p> |  <p>Lombard Street, worlds crookedest See ... 500 x 387 - 59k - jpg www.inetours.com</p> |  <p>Street Bike (BS70-4A) Details 360 x 360 - 38k - jpg bashan.en.alibaba.com</p> |
|  <p>Street Lamps 360 x 360 - 18k - jpg syi.en.alibaba.com [More from img.alibaba.com]</p> |  <p>Washington D.C. Laminated Street Map 500 x 500 - 114k - jpg www.dcgiftshop.com</p> |  <p>street-riders-ss-3.jpg 550 x 309 - 53k - jpg www.pspworld.com</p> |  <p>Visually Street Riders is not nearly ... 550 x 309 - 52k - jpg www.pspworld.com</p> |  <p>STREET space ring Postcards To Space ... 1000 x 563 - 87k - jpg www.postcardstospace.com</p> |  <p>17 Fleet Street 492 x 681 - 74k - jpg www.pepysdiary.com</p> |

Organizing photo collections



bing™

car



CAR

Car Manufacturers

Car Models

Car Prices

Compare Cars

Car Reviews

Reference Articles
on car

Images of car

RELATED SEARCHES

Police Car

Car Crash

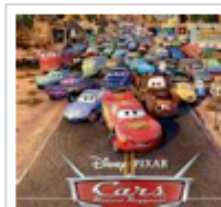
Cartoon Car

Smart Car

Red Car

Concept Car

IMAGES 16-35 of 36,400,000 results SafeSearch moderate ([Change](#))



Mikes Car

400 x 300 · 76 kB · jpeg

www.mikesguardianeaglefoundation.org

[Find more sizes](#)

[Show similar images](#)



bing™

car|



CAR

[Car Manufacturers](#)

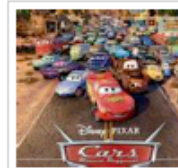
[Car Models](#)

[Car Prices](#)

[Compare Cars](#)

[Car Reviews](#)

IMAGES 16-35 of 36,400,000 results SafeSearch moderate ([Change](#))



Mikes Car

400 x 300 · 76 kB · jpeg

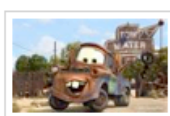
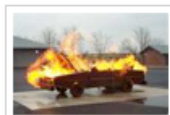
www.mikesguardianeaglefoundation.o

[Find more sizes](#)

[Show similar images](#)



IMAGES 1-20 of 240 results SafeSearch moderate ([Change](#))



bing™

car|



CAR

[Car Manufacturers](#)

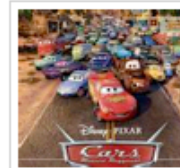
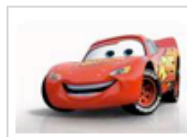
[Car Models](#)

[Car Prices](#)

[Compare Cars](#)

[Car Reviews](#)

IMAGES 16-35 of 36,400,000 results SafeSearch moderate ([Change](#))



Mikes Car

400 x 300 · 76 kB · jpeg

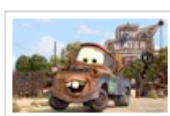
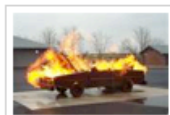
www.mikesguardianeaglefoundation.o

[Find more sizes](#)

[Show similar images](#)



IMAGES 1-20 of 240 results SafeSearch moderate ([Change](#))



Visual Dictionary

Teaching computers to recognize objects

[Download dataset](#)

[Download poster](#)

[Publication](#)

Search word



Visual Dictionary

Teaching computers to recognize objects

[Download dataset](#)

[Download poster](#)

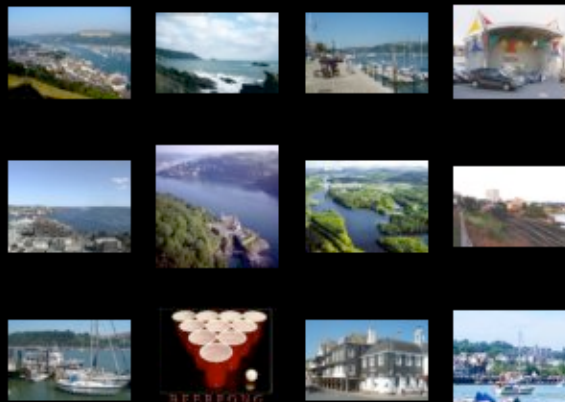
[Publication](#)

Search word

12618) Dartmouth

[Hide definition](#)

Click on the image if you think it is correct (a green frame will appear), and twice if it is wrong. For images that you are not sure leave the black frame around the image.



Click to submit selection

Next ▶

Definition (from Wordnet)

Dartmouth College, Dartmouth -- (a college in New Hampshire)

Wikipedia: [open wikipedia page](#)

Average Image



Visual Dictionary

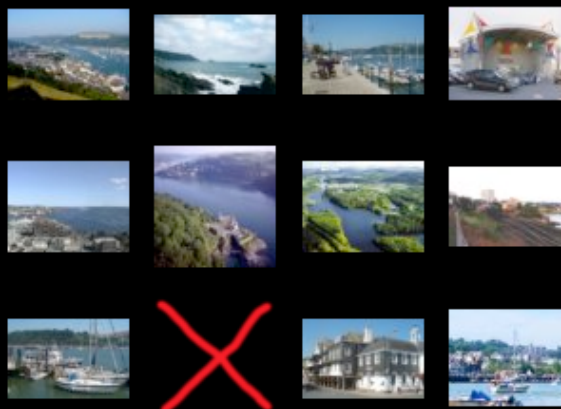
Teaching computers to recognize objects

[Download dataset](#)[Download poster](#)[Publication](#)

12618) Dartmouth

[Hide definition](#)

Click on the image if you think it is correct (a green frame will appear), and twice if it is wrong. For images that you are not sure leave the black frame around the image.



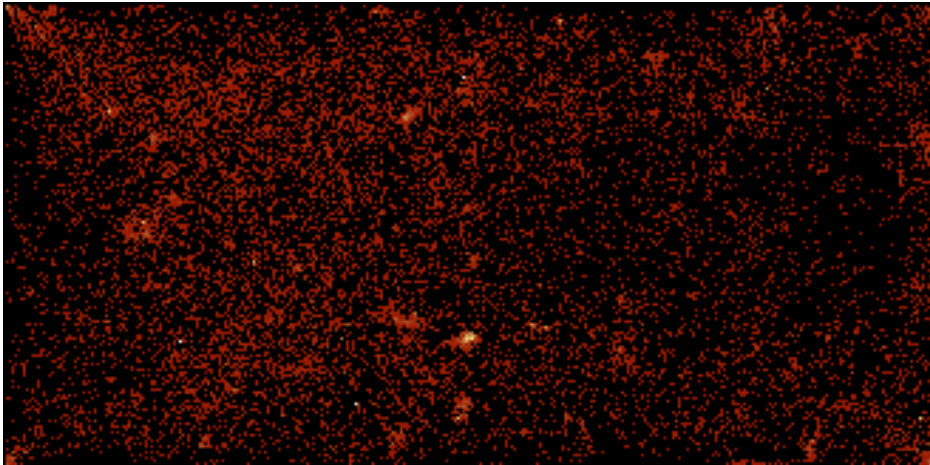
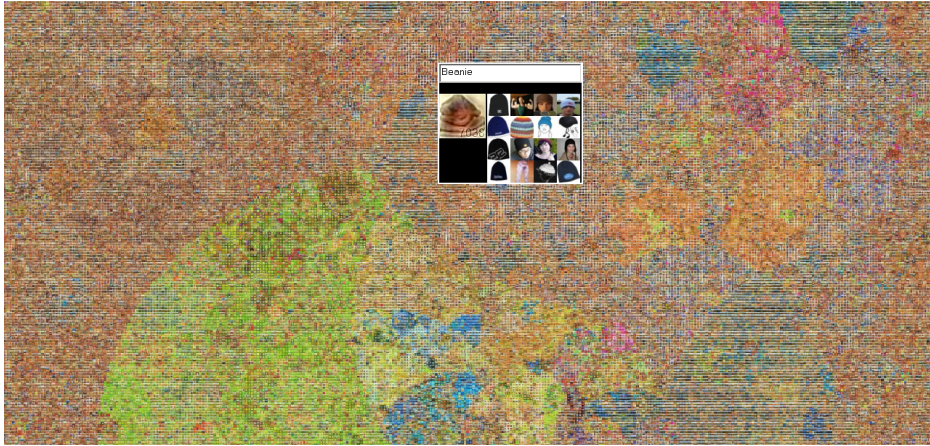
Definition (from Wordnet)

Dartmouth College, Dartmouth -- (a college in New Hampshire)

Wikipedia: [open wikipedia page](#)

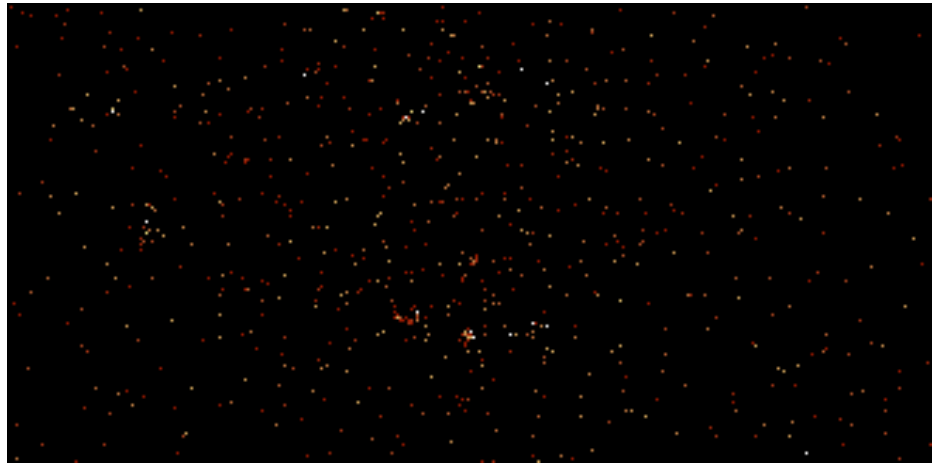
Average Image



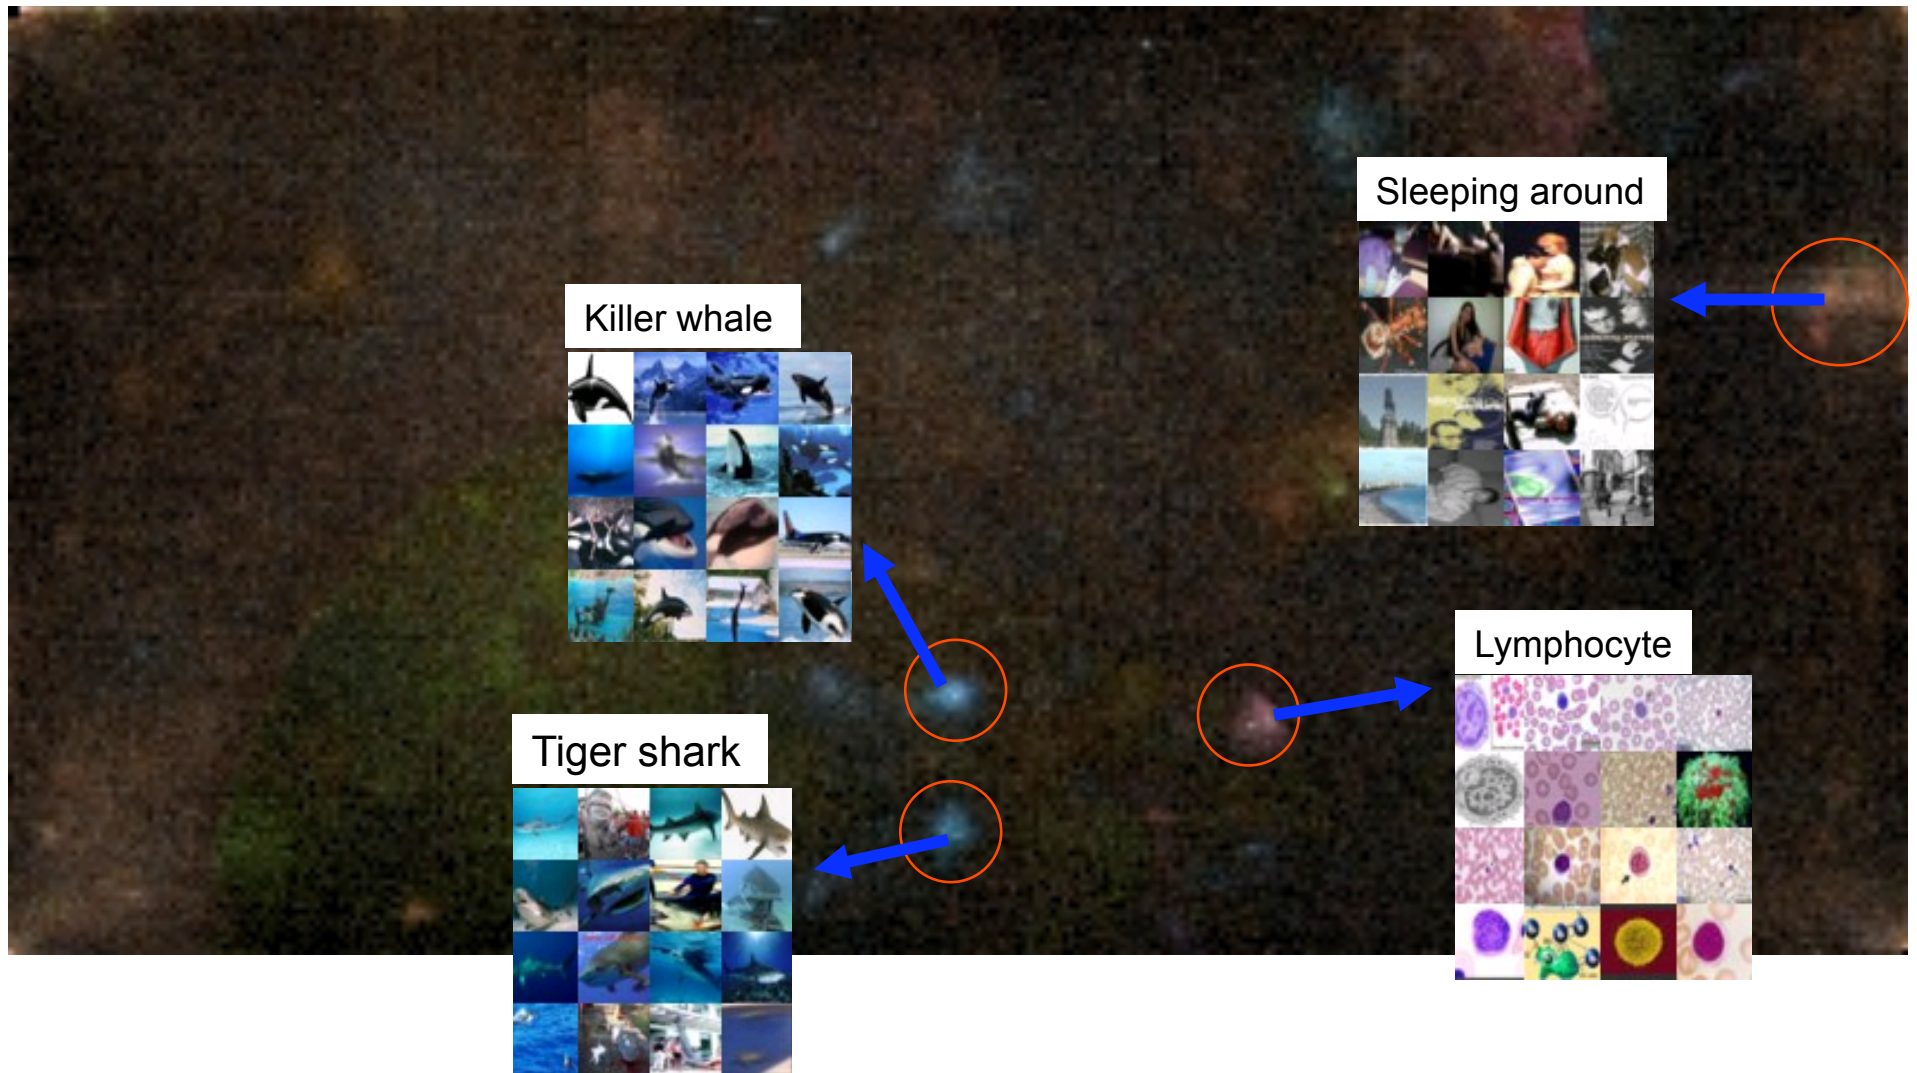


Views

Labels



Web users



Problem set 1

- Build the device
- Take some pictures and put them in the report
- Take anaglyph images
- Work out the geometry
- Recover depth for some points in the image