CSE 455: Computer Vision
Winter 2007

Instructor: Professor Linda Shapiro (shapiro@cs)
Additional Instructor: Dr. Matthew Brown
(brown@microsoft.com)

TAs: Masa Kobashi (mkbsh@cs)
Peter Davis (pediddle@cs)

Text: Shapiro and Stockman, Computer Vision
(chapters available from class web page)

Evaluation: 70% programming projects, 30% exams

Topics

- Basics: images, binary operations, filtering, edge operators
- Color, texture, segmentation
- Interest operators: detectors and descriptors
- Use of interest operators: object recognition, stitching, tracking
- Content-based image retrieval
- 2D object recognition
- Motion
- 3D: sensing, camera calibration, reconstruction, recognition

- What IS computer vision?

  the analysis of digital images by a computer

- Where do images come from?

Applications

- Medical Imaging

CT image of a patient's abdomen

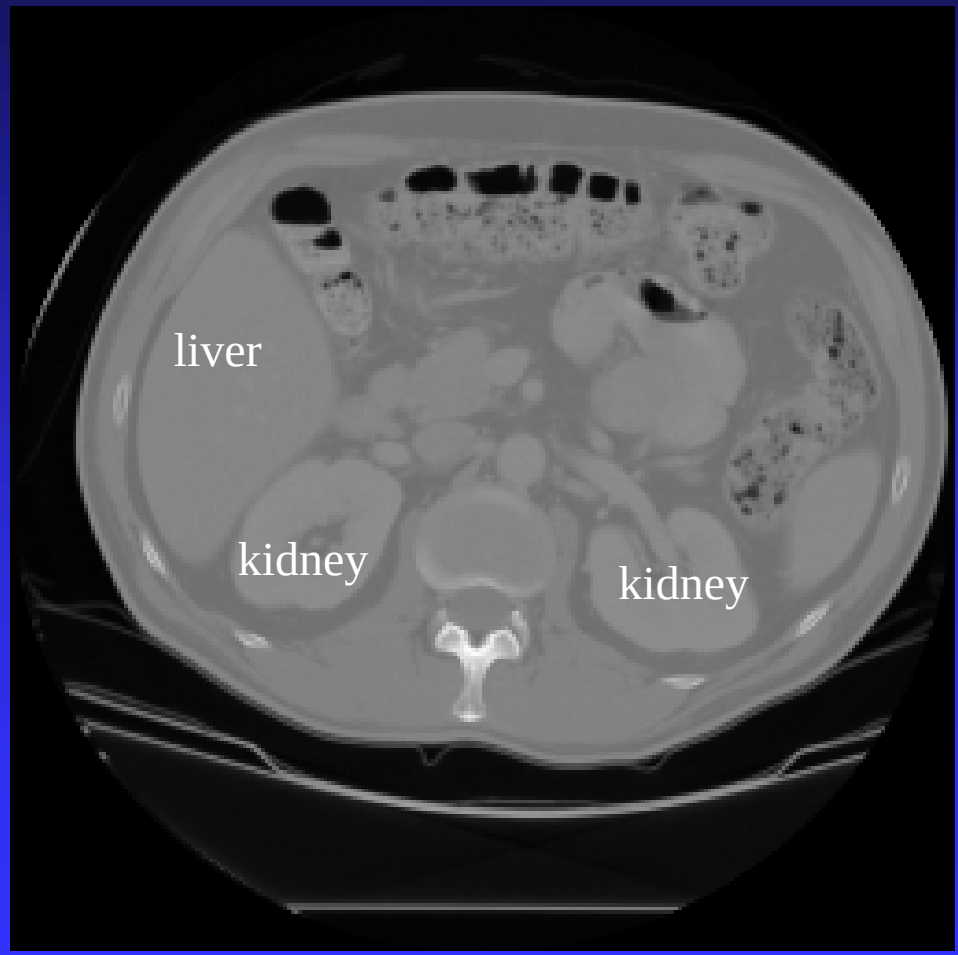

Visible Man Slice Through Lung

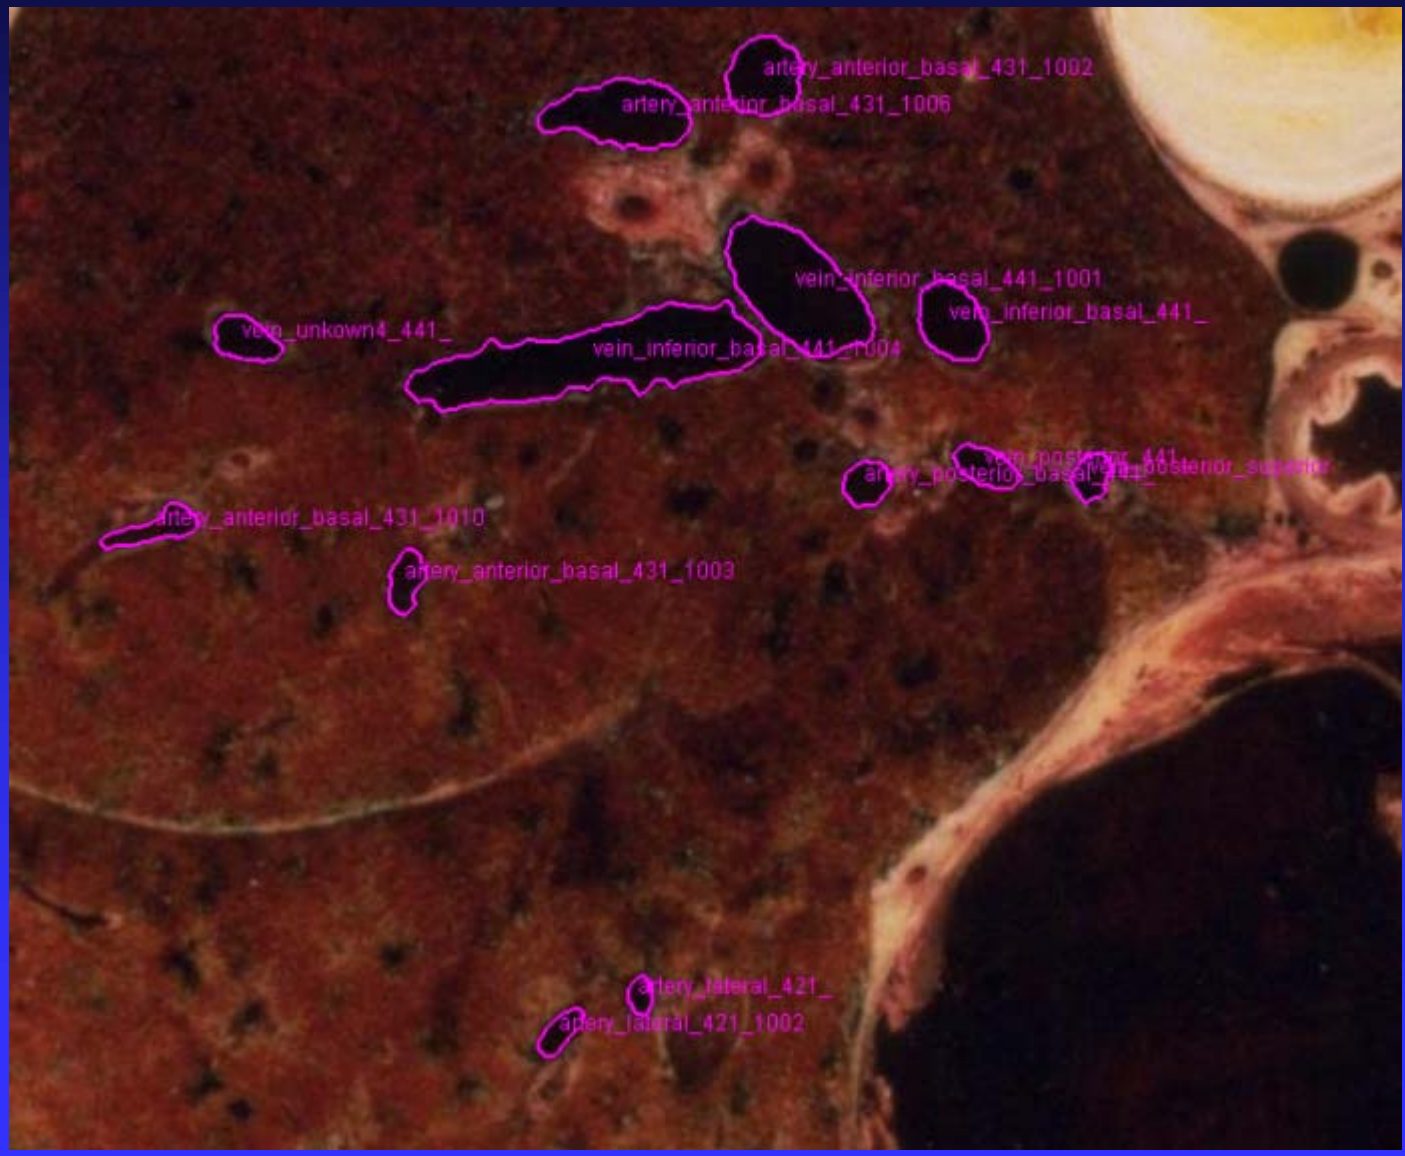

3D Reconstruction of the Blood Vessel Tree

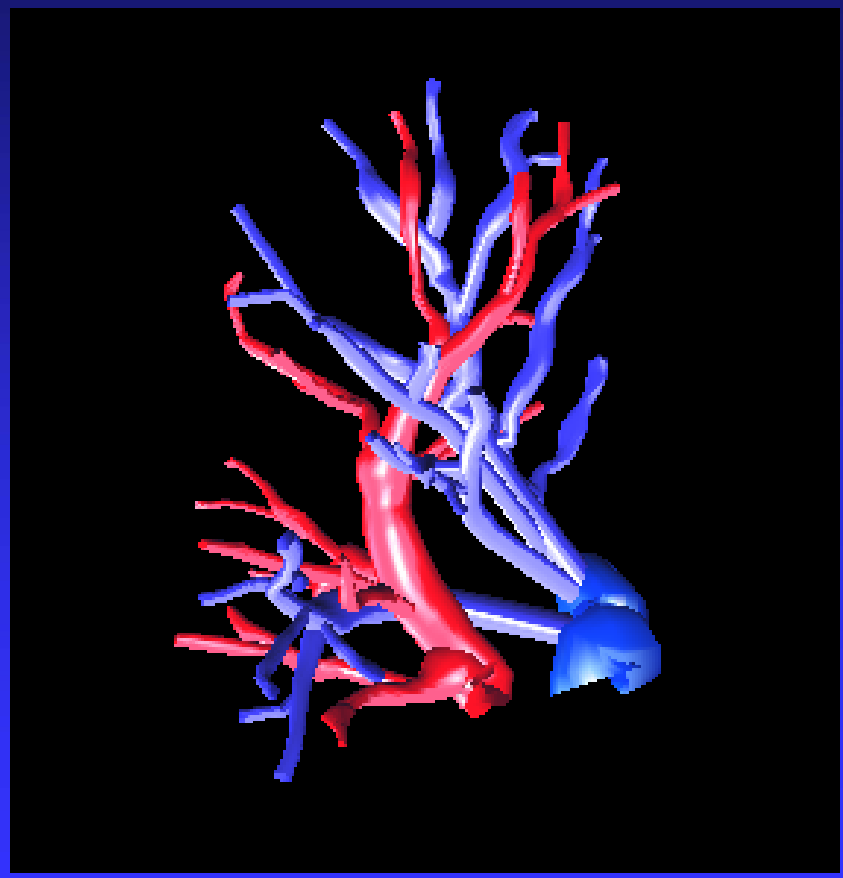

Symbolic Shape Descriptors for Classifying Craniosynostosis

sagittal synostosis

metopic synostosis

Robotics

- 2D Gray-tone or Color Images

"Mars" rover

- 3D Range Images

What am I?

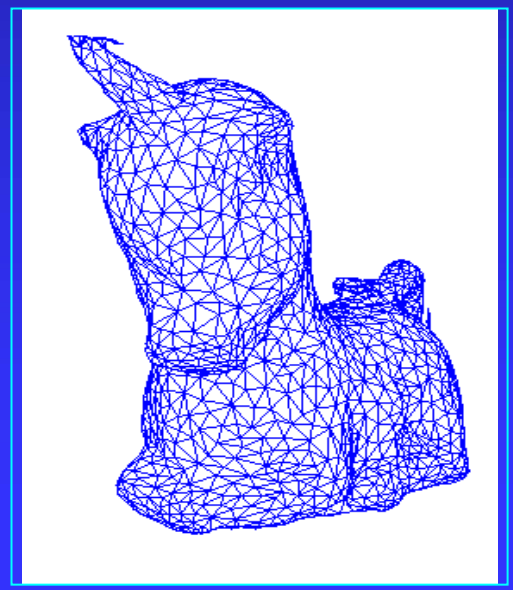

Recognition of 3D Object Classes from Range Data

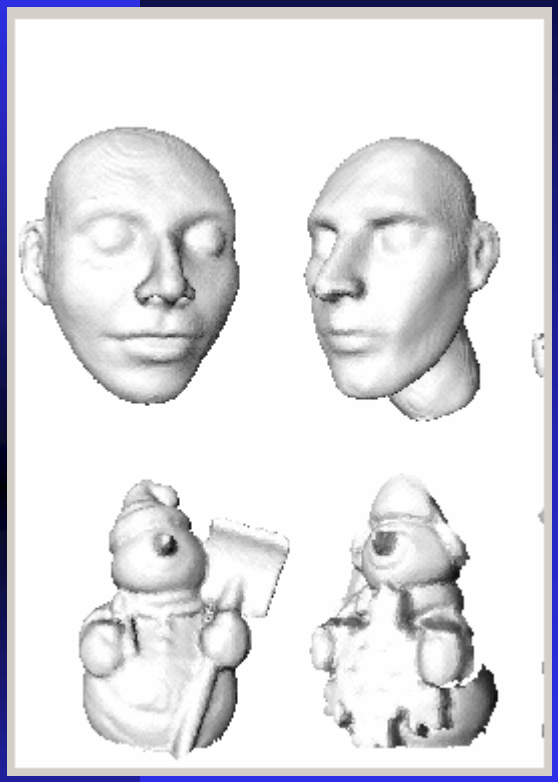

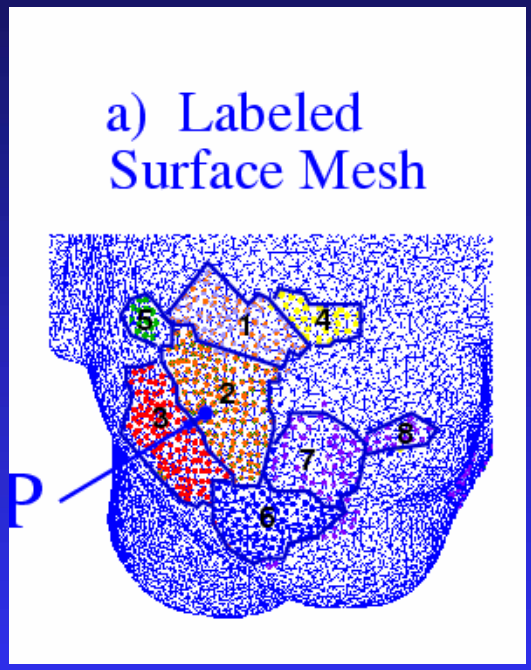

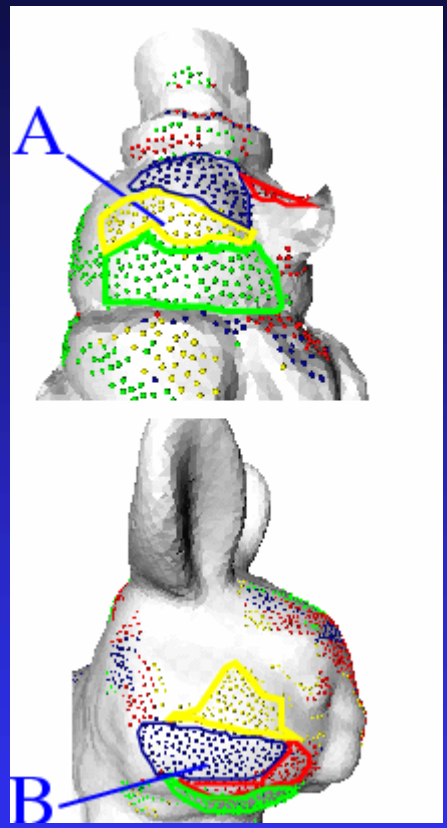

Image Databases: Content-Based Retrieval

Images from my Ground-Truth collection.



What categories of image databases exist today?

Similarity Retrieval of Brain Data

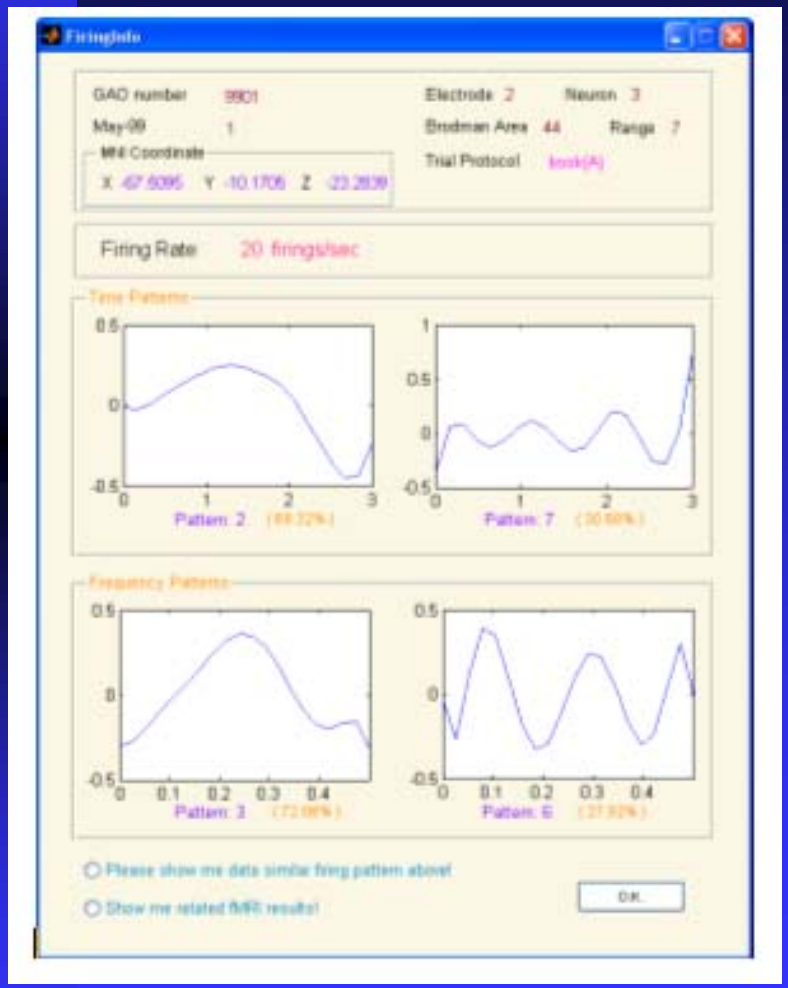

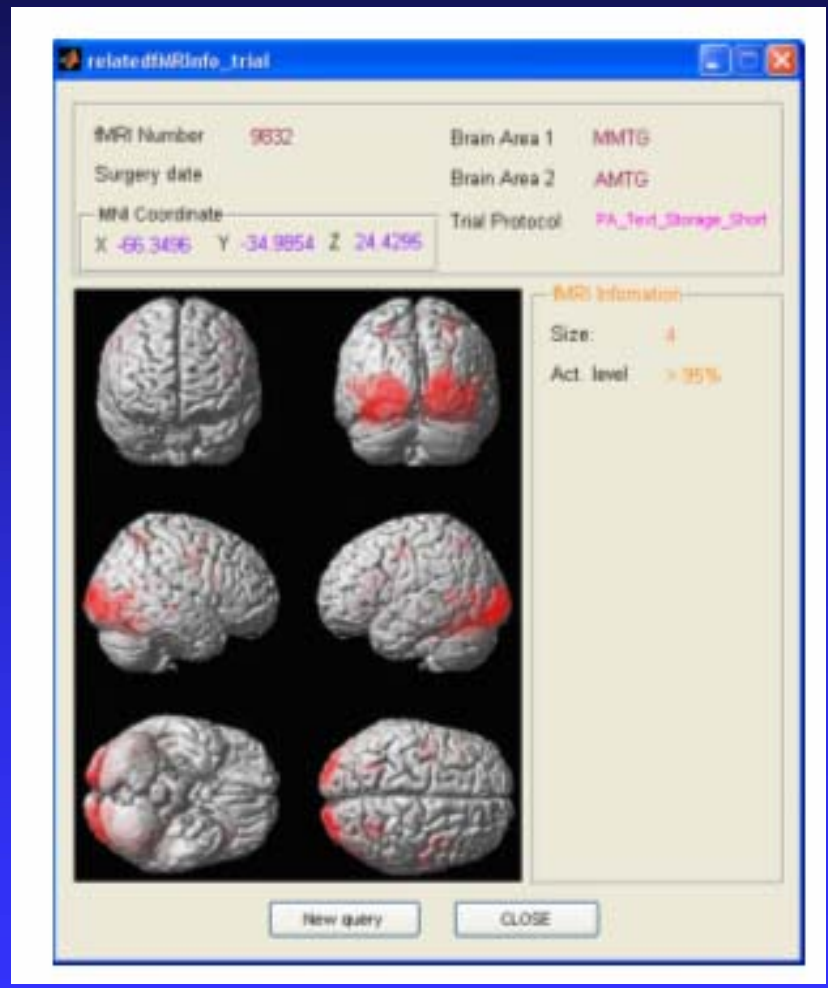

CBIR of Mouse Eye Images for Genetic Studies

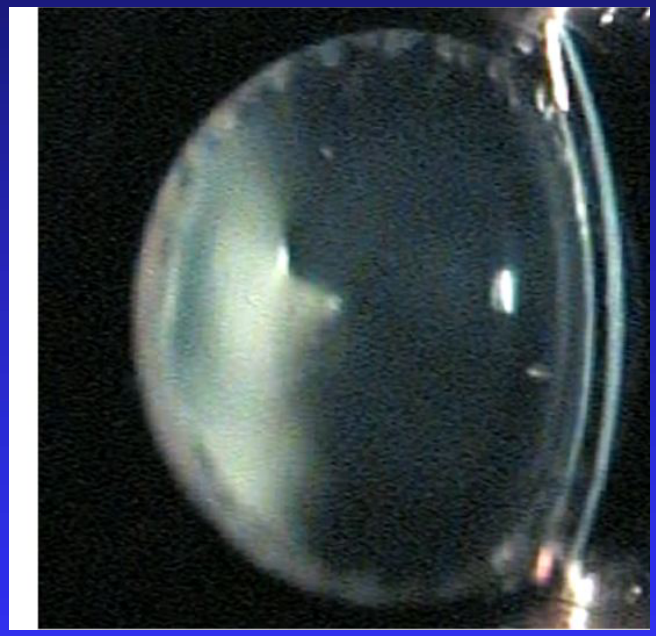

Abstract Regions for Object Recognition

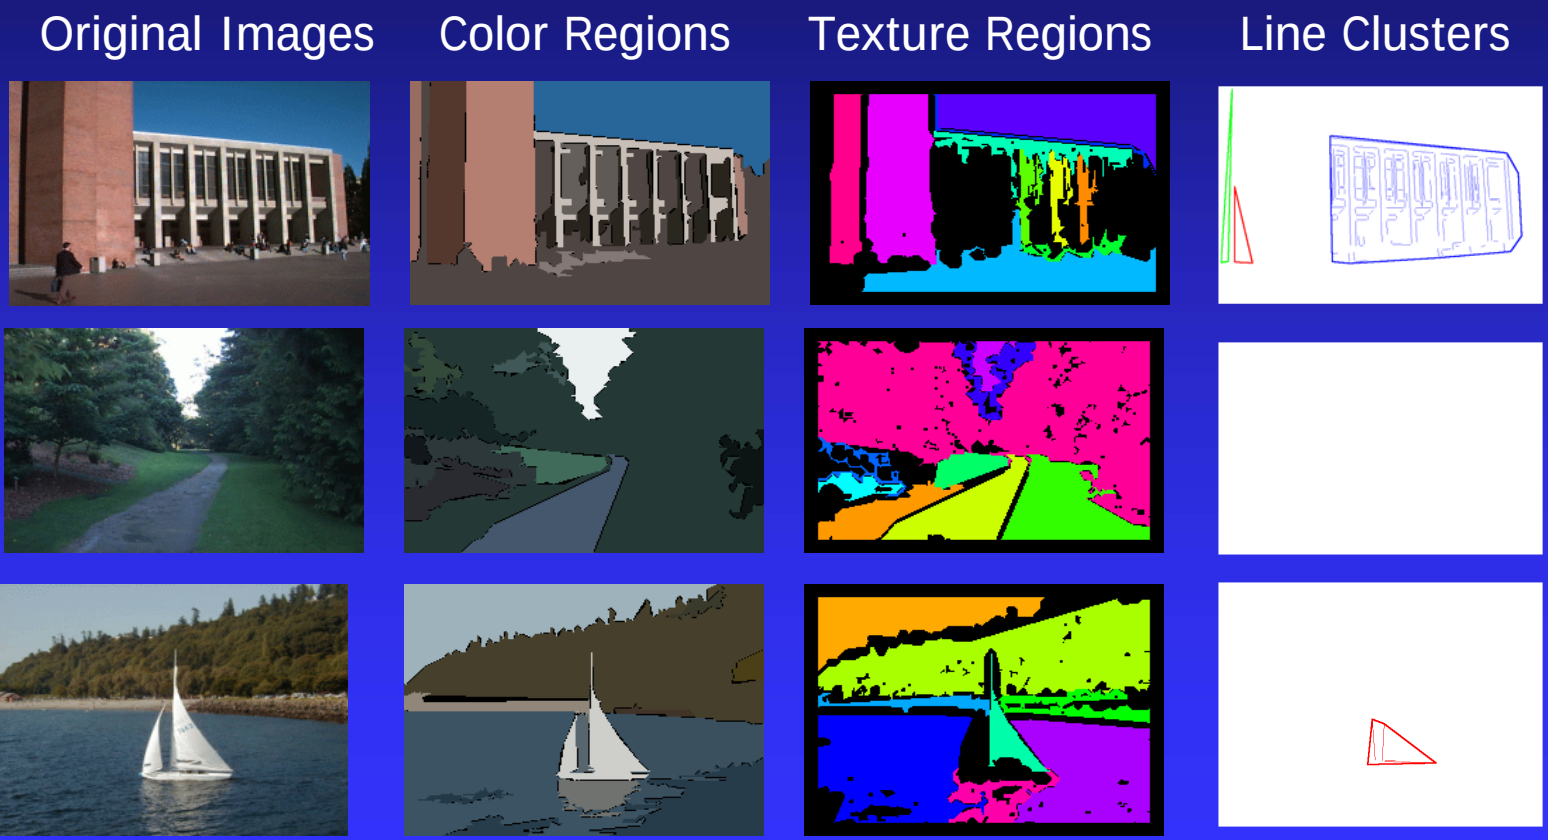

Insect Identification for Ecology Studies

Calineuria (Cal)

Doroneuria (Dor)

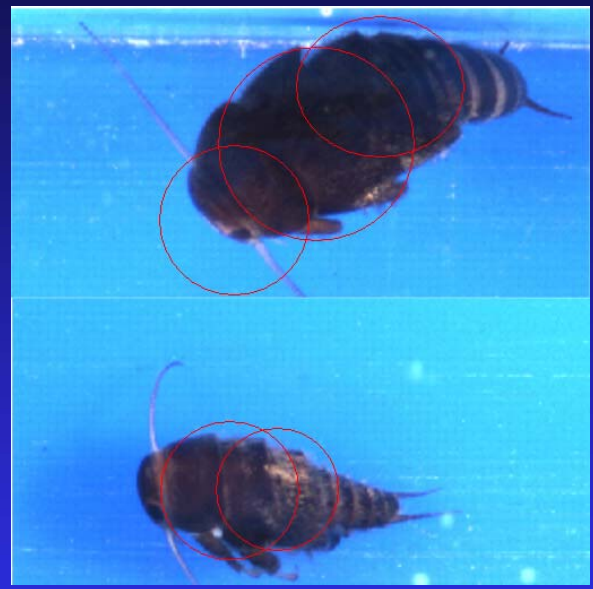

Yoraperla (Yor)

Document Analysis

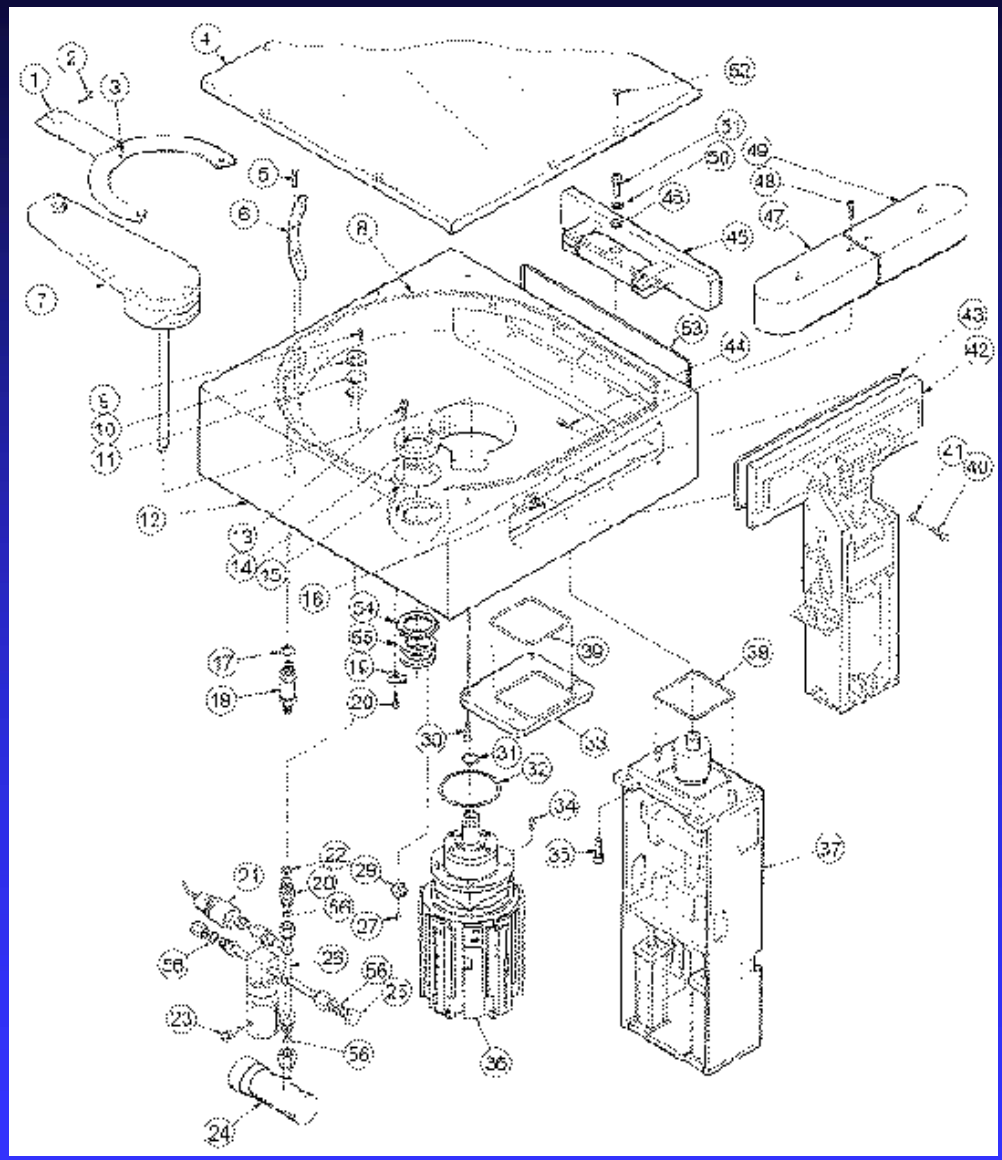

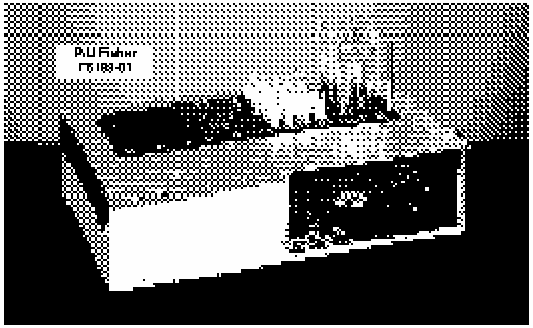

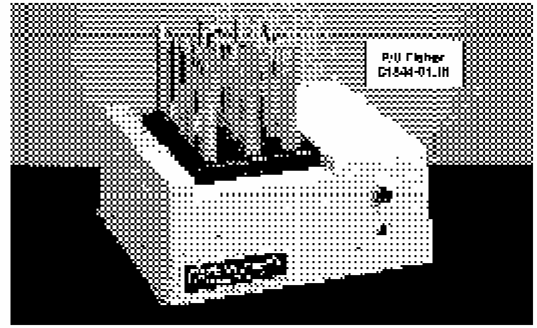

Model 145 Isotemp* Dry Bath Incubator

- *Holds 1 to 4 heating blocks with choice of 11 well sizes*
- *Maintains every sample to within ±0.1° C of temperature*

Unique sample wells are shaped so that aluminum blocks deliver same amount of heat to all parts of the sample tube. No temperature gradient—neither too hot at the bottom nor too cold at the top—that may invalidate tests in baths with drilled cylindrical wells. Sample tubes rest on insulating plugs to prevent localized heating. A low watt-density heater is mounted on a thick aluminum heat distribution plate in the floor of the bath. Plate is ⅜" thick (9.5 mm). Dry bath minimizes cleanup problems because tubes stay dry.

Ambient to 125° C (257° F) with ±0.1° C control. Dial temperature control graduated from 25° to 125° C. Ideal for enzyme reactions, inactivation of sera, Rh studies, cross-matching and thrombin determinations. Dimensions: 8 L x 15 W x 4" H (23 x 38 x 10 cm). With line cord and plug. **Heating blocks sold separately** (see lower right).

Electrical Requirements	Cat. No.	Each
120V 50/60 Hz, 300W (CSA approved)	11-715-100	[illegible]
240V 50/60 Hz, 600W	11-715-101	395.23

Average deviation from nominal 37° C
Intermediate Model

Incu-Block* Partial Immersion Thermometers

For all standard heating blocks and water baths. Critical temperatures (25°, 30°, 37°, 56° C) are marked with arrows. Available with shatter-proof, contamination-proof Teflon* coating. Total length: 175 mm. Immersion: 35 mm.

Range, °C	Div., °C	Teflon-Coated	Cat. No.	Each
25-57	0.5	No	14-982	45.84
25-57	0.5	Yes	14-983	45.15

More Thermometers

For more thermometers, including digital types,

see page 952

Model 147 Isotemp* Dry Bath

- *Holds single heating block with choice of 11 well sizes*

Similar to Model 145, but with ⅜" thick (22.2 mm) plate. Ideal for labs with smaller volumes of enzyme and coagulation assays, Rh studies, and dry incubations. Screwdriver-adjusted temperature control between ambient and 95° C (203° F). Observe thermometer placed in unused sample tube to set; adjust control through hole in front panel. Maintains set temperature with consistency and uniformity of ±0.5° C.

Supplied with strong nylon case, thermostatically controlled heater, pilot indicator lamp, line cord and plug, and instructions. Dimensions: 8 L x 6¾ W x 3" H (15 x 17 x 8 cm). CSA approved. **Heating blocks sold separately** (see below).

Electrical Requirements	Cat. No.	Each
120V 50/60 Hz, 120W	11-715-102	223.50

Interchangeable Heating Blocks for Isotemp* Dry Baths

For Models 145 and 147 Dry Baths. Composed of black anodized aluminum alloy. Chemical resistant. Dimensions: 4 L x 3 W x 2" H (10 x 7.5 x 5 cm).

The 11-715-123 block provides a safe, dry bath alternative for warming the Spatula of tissue for paraffin sections. Avoids hazardous use of burners around flammable histological solvents.

The 11-715-120 block is specifically designed to hold twenty 9.5 mm Roche Diagnostics Placentex* pregnancy test tubes. This special shallow well block is similar to the other blocks with 10 mm holes, but sample wells are only ⅜" deep (1.0 cm) to meet test requirements. Wells in all other blocks are 1¾" deep (4.4 cm).

Tube Size, mm	Wells/Block	Cat. No.	Each
6	35	11-715-105	71.[illegible]
10	20	11-715-107	71.[illegible]
9.5	20 (shallow)	11-715-120	71.[illegible]
12	12	11-715-109	71.[illegible]
12.5	12	11-715-121	71.[illegible]
13	12	11-715-111	71.[illegible]
15	12	11-715-113	71.[illegible]
16	8	11-715-123	71.[illegible]
18	12	11-715-115	71.[illegible]
20	6	11-715-117	71.[illegible]
25	5	11-715-119	71.[illegible]

[illegible]
[illegible]

Surveillance: Object and Event Recognition in Aerial Videos

Original Video Frame

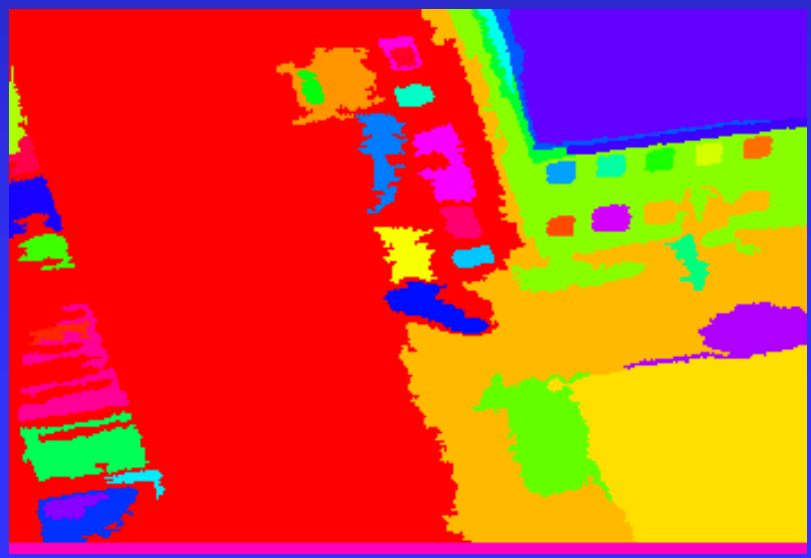

Color Regions

Structure Regions

Video Analysis

What are the objects? What are the events?

3D Scanning

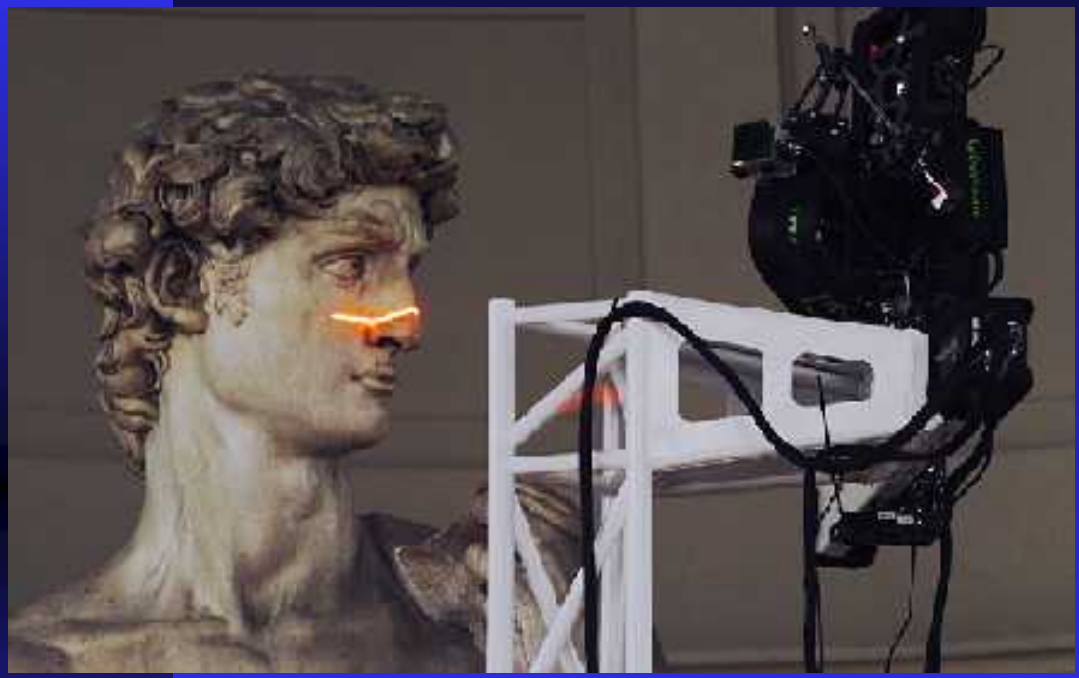

Scanning Michelangelo's *"The David"*

- The Digital Michelangelo Project
  - http://graphics.stanford.edu/projects/mich/
- UW Prof. Brian Curless, collaborator
- 2 BILLION polygons, accuracy to .29mm

The Digital Michelangelo Project, Levoy et al.

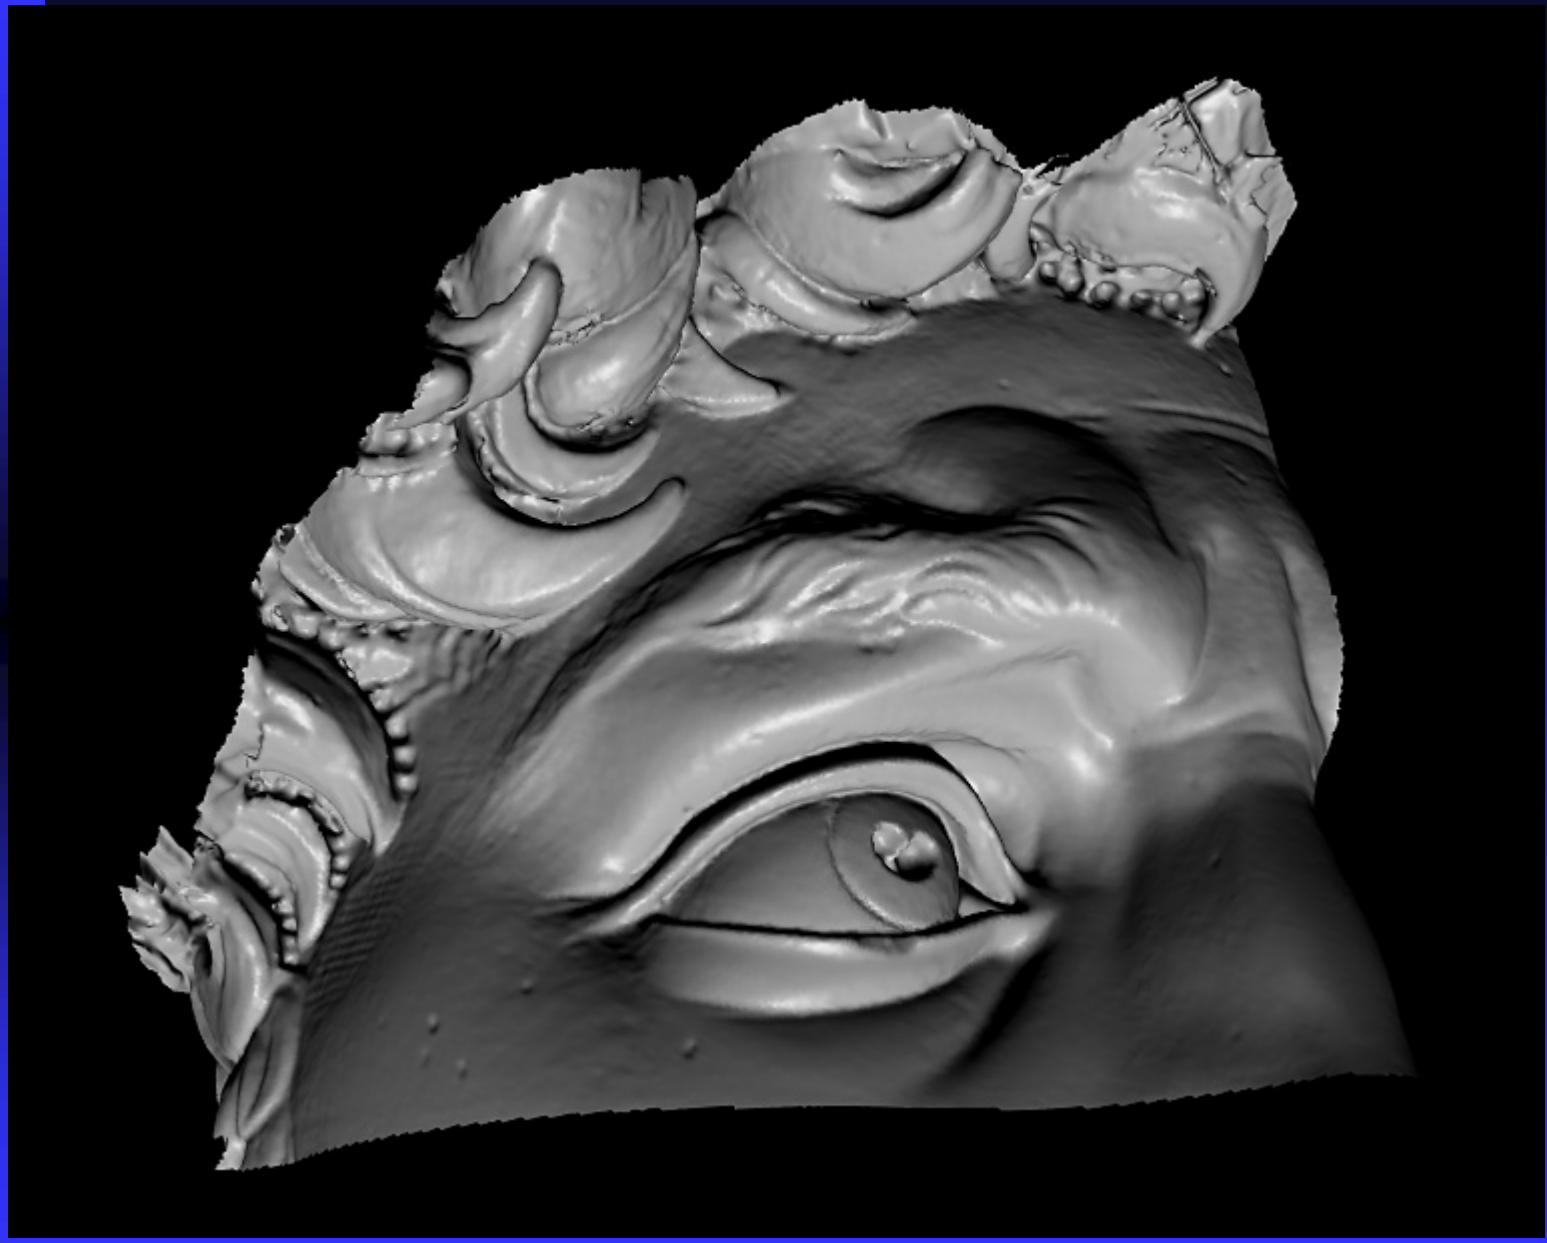

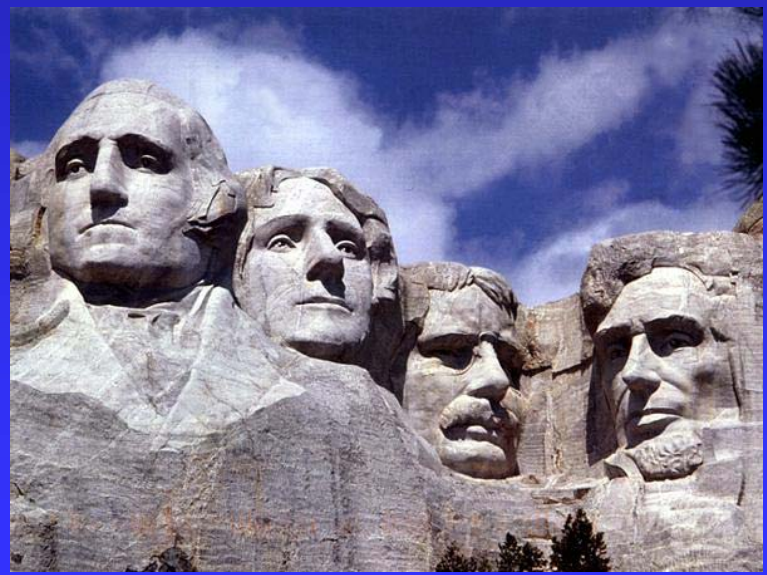

Motion Capture, Games

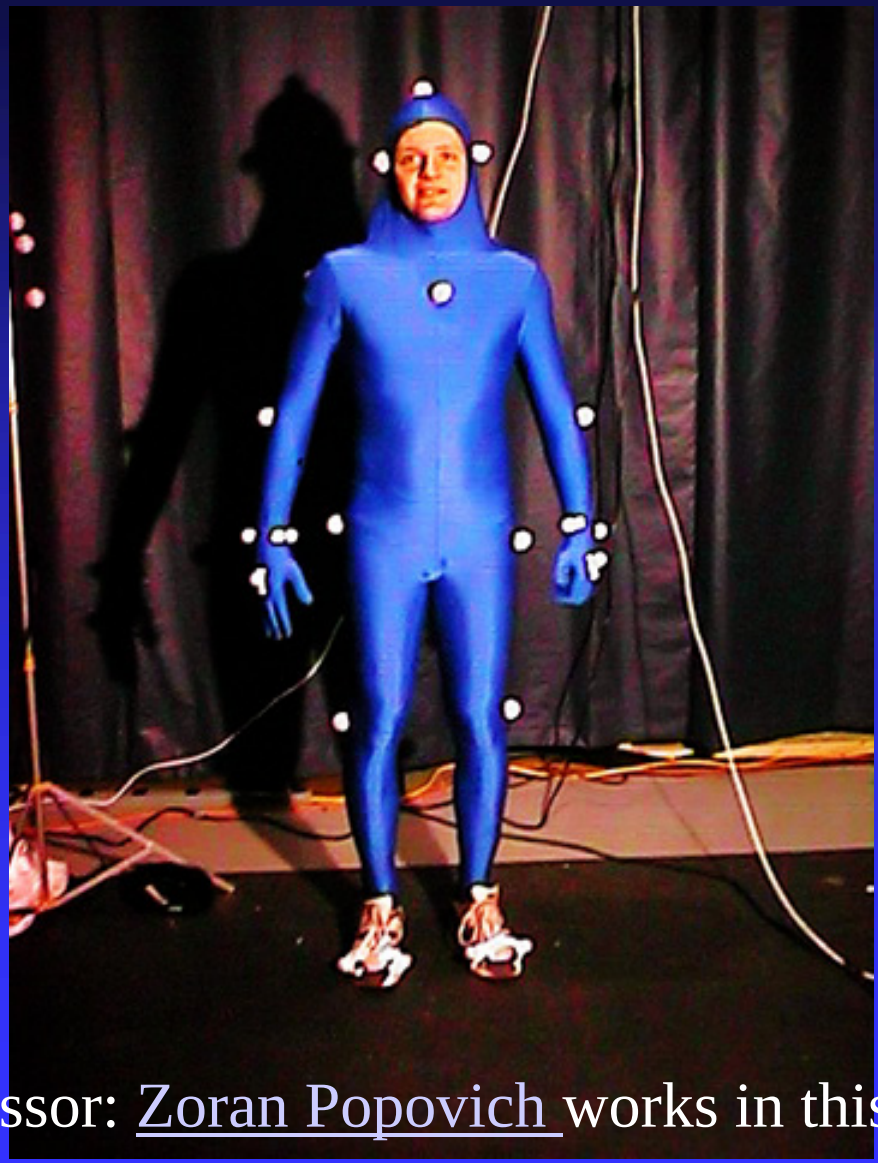

UW Professor: Zoran Popovich works in this area.

Effects

Andy Serkis, Gollum, Lord of the Rings

Imaging

Digital Image Terminology:

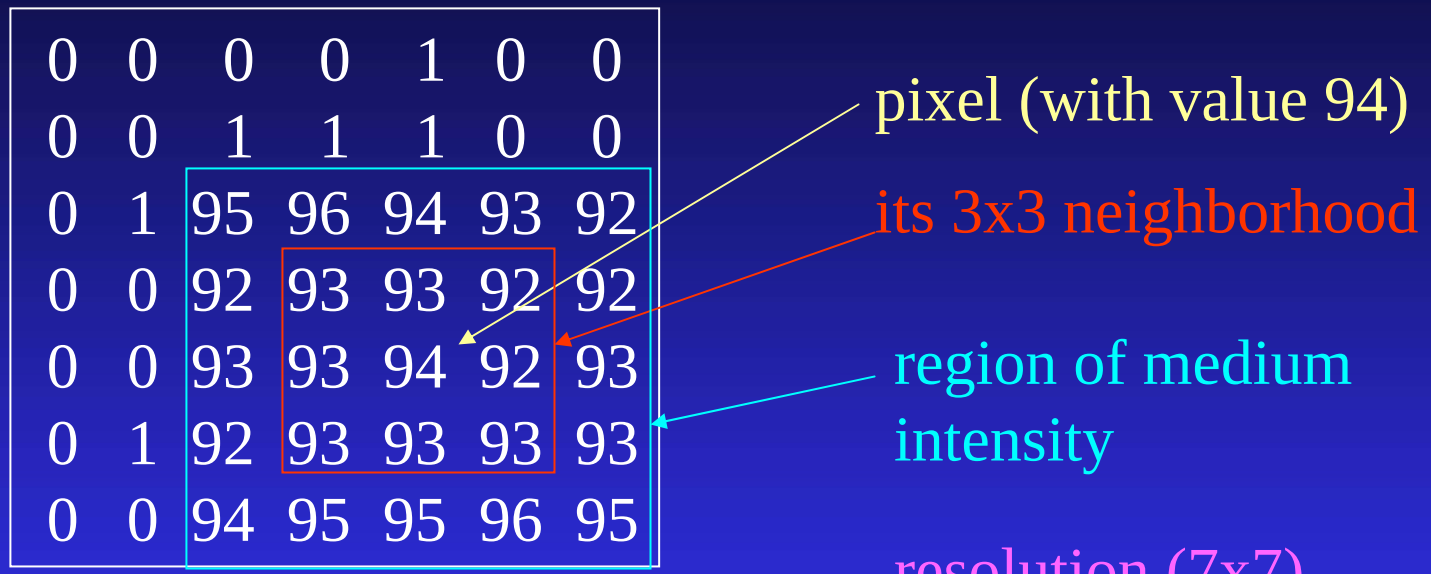

- binary image
- gray-scale (or gray-tone) image
- color image
- multi-spectral image
- range image
- labeled image

Goals of Image and Video Analysis

- Segment an image into useful regions
- Perform measurements on certain areas

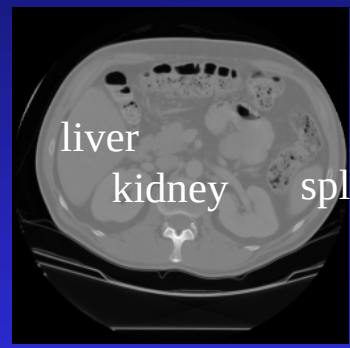

- Determine what object(s) are in the scene
- Calculate the precise location(s) of objects
- Visually inspect a manufactured object

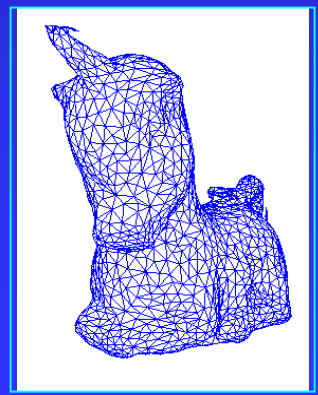

- Construct a 3D model of the imaged object
- Find “interesting” events in a video

•The Three Stages of Computer Vision

- low-level

  image → image

- mid-level

  image → features

- high-level

  features → analysis

Low-Level

sharpening

blurring

Low-Level

original image

Canny

edge image

Mid-Level

edge image

ORT

data structure

circular arcs and line segments

Mid-level

original color image

K-means clustering

(followed by connected component analysis)

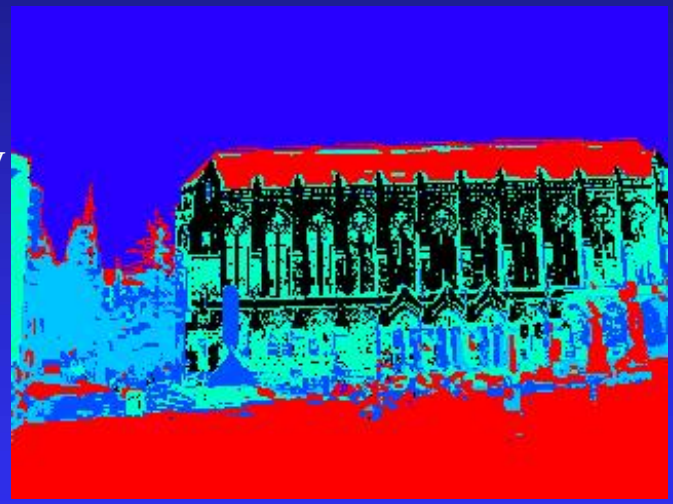

regions of homogeneous color

data structure

Low- to High-Level

low-level

edge image

mid-level

consistent line clusters

high-level

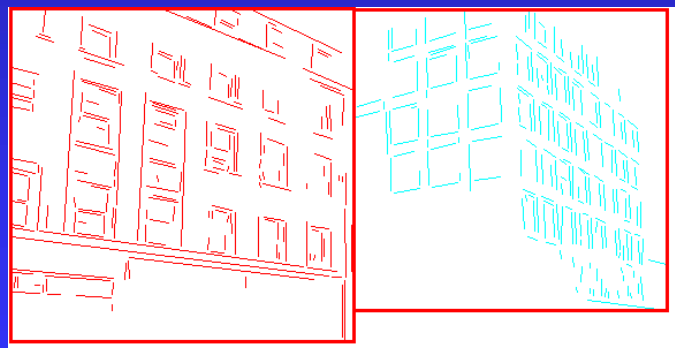

Building Recognition

Imaging and Image Representation

- **Sensing Process**
- **Typical Sensing Devices**
- **Problems with Digital Images**
- **Image Formats**
- **Relationship of 3D Scenes to 2D Images**
- **Other Types of Sensors**

Images: 2D projections of 3D

- The 3D world has color, texture, surfaces, volumes, light sources, objects, motion, …
- A 2D image is a projection of a scene from a specific viewpoint.

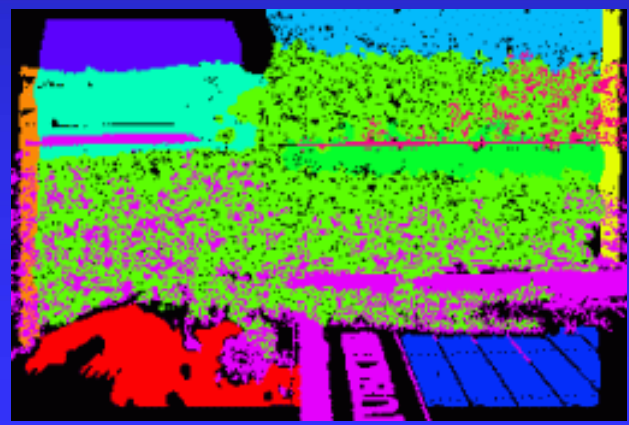

Images as Functions

- **A gray-tone image is a function:**

  g(x,y) = val or f(row, col) = val

- **A color image is just three functions or a vector-valued function:**

  f(row,col) =(r(row,col), g(row,col), b(row,col))

Image vs Matrix

Digital images (or just “images”) are typically stored in a matrix.

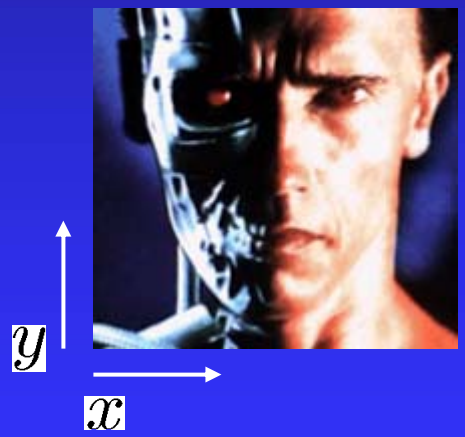

j

i

62	79	23	119	120	105	4	0
10	10	9	62	12	78	34	0
10	58	197	46	46	0	0	48
176	135	5	188	191	68	0	49
2	1	1	29	26	37	0	77
0	89	144	147	187	102	62	208
255	252	0	166	123	62	0	31
166	63	127	17	1	0	99	30

There are many different file formats.

Gray-tone Image as 3D Function

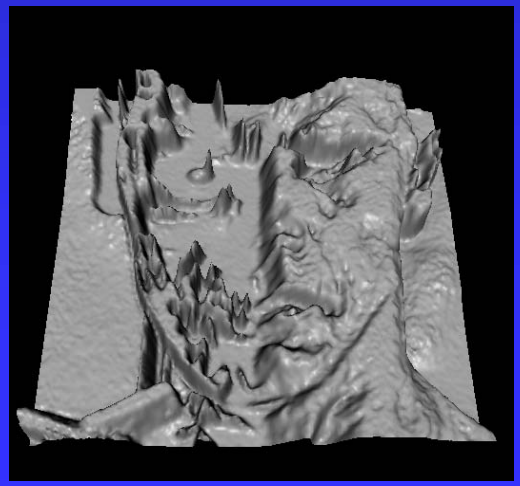

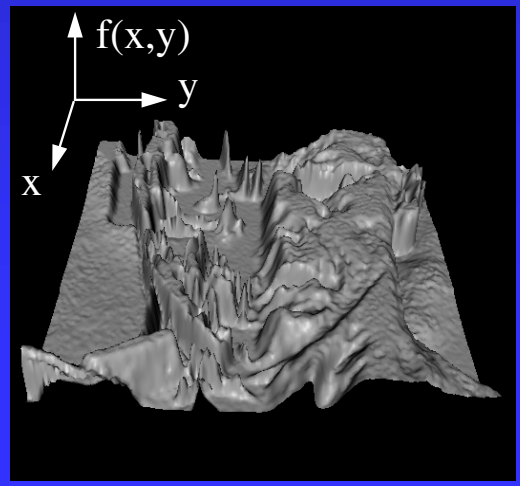

Imaging Process

- Light reaches surfaces in 3D
- Surfaces reflect
- Sensor element receives light energy
- Intensity counts
- Angles count
- Material counts

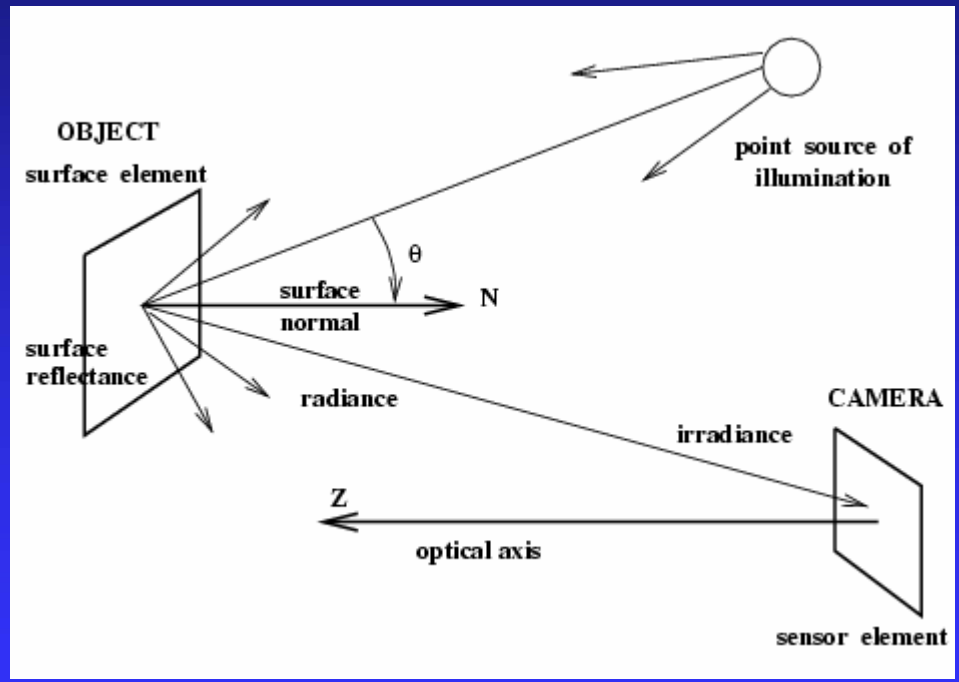

What are radiance and irradiance?

Radiometry and Computer Vision*

- **Radiometry** is a branch of physics that deals with the measurement of the flow and transfer of radiant energy.
- **Radiance** is the power of light that is emitted from a unit surface area into some spatial angle; the corresponding photometric term is **brightness**.
- **Irradiance** is the amount of energy that an image-capturing device gets per unit of an efficient sensitive area of the camera. Quantizing it gives image gray tones.

•From Sonka, Hlavac, and Boyle, *Image Processing, Analysis, and Machine Vision*, ITP, 1999.

CCD type camera:
Commonly used in industrial applications

- Array of small fixed elements
- Can read faster than TV rates
- Can add refracting elements to get color in 2x2 neighborhoods
- 8-bit intensity common

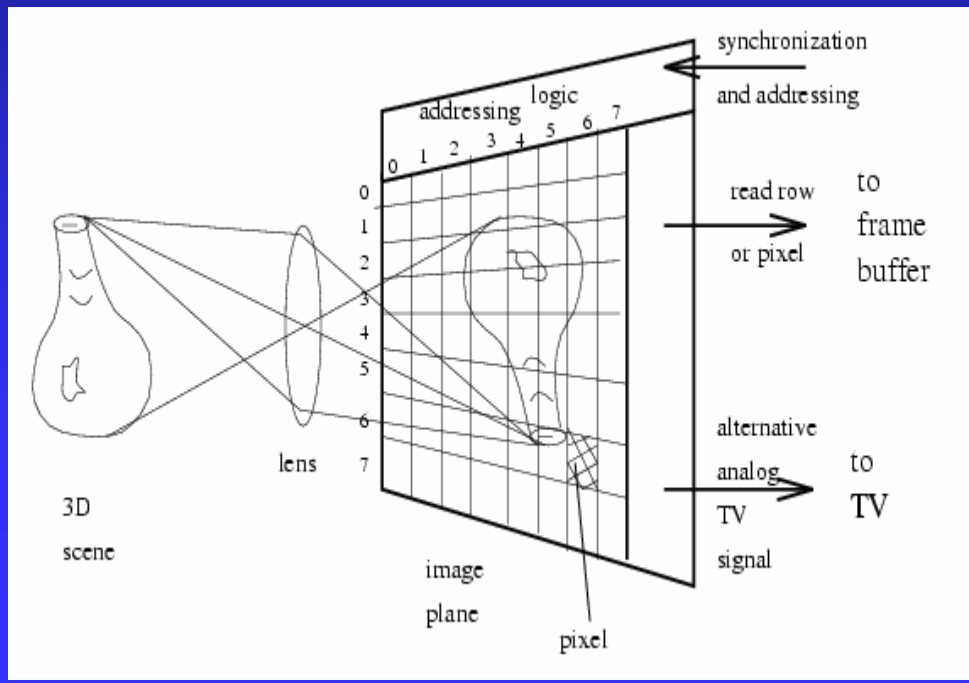

Blooming Problem with Arrays

- Difficult to insulate adjacent sensing elements.
- Charge often leaks from hot cells to neighbors, making bright regions larger.

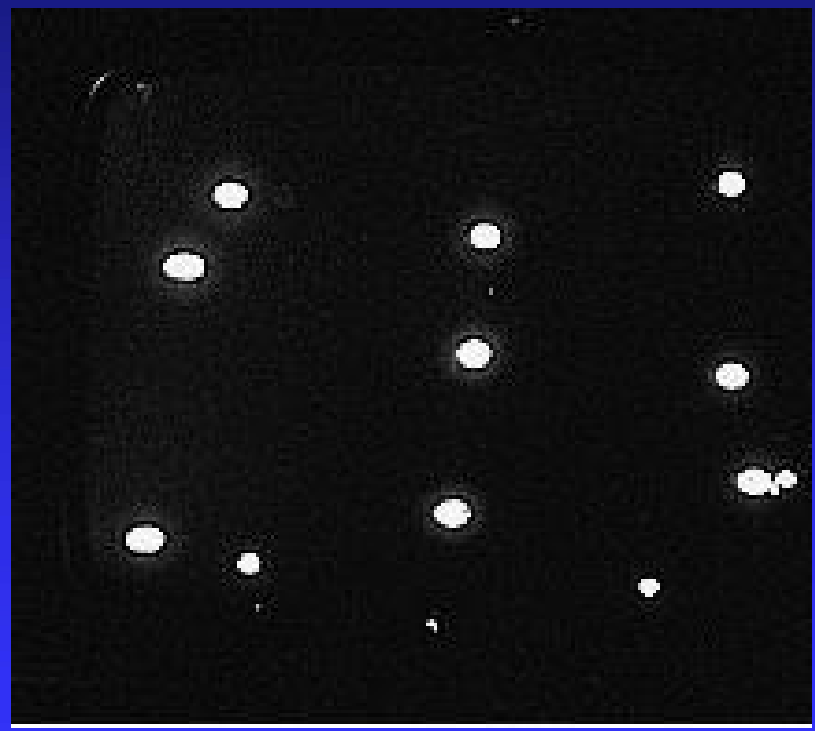

8-bit intensity can be clipped

- Dark grid intersections at left were actually brightest of scene.
- In A/D conversion the bright values were clipped to lower values.

Lens distortion distorts image

- “Barrel distortion” of rectangular grid is common for cheap lenses ($50)
- Precision lenses can cost $1000 or more.
- Zoom lenses often show severe distortion.

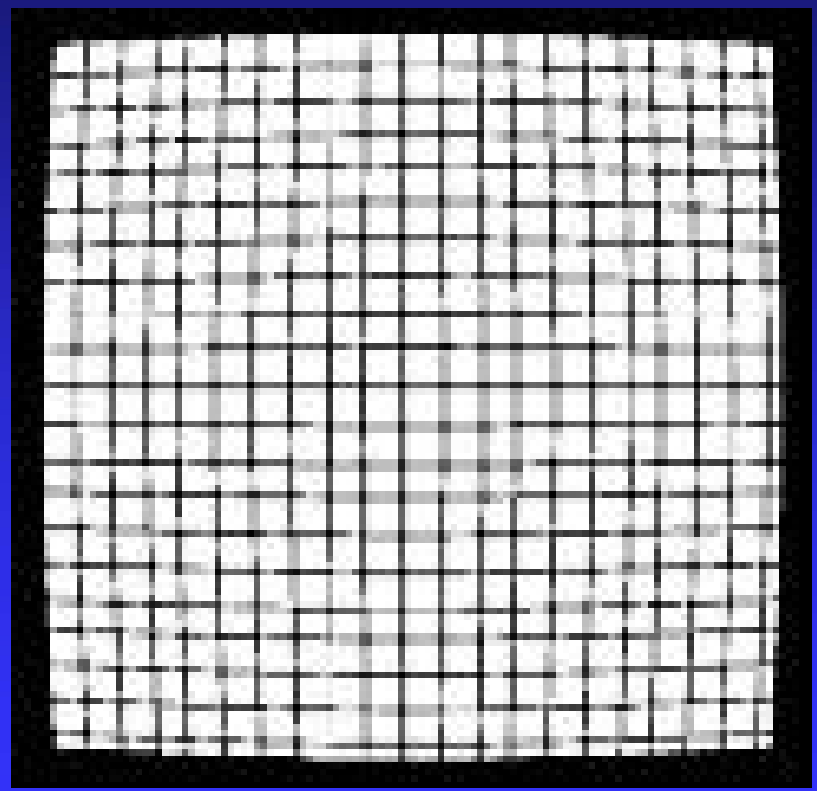

Resolution

- **resolution:** precision of the sensor
- **nominal resolution:** size of a single pixel in scene coordinates (ie. meters, mm)
- **common use of resolution:** num_rows X num_cols (ie. 515 x 480)
- **subpixel resolution:** measurement that goes into fractions of nominal resolution
- **field of view (FOV):** size of the scene a sensor can sense

Resolution Examples

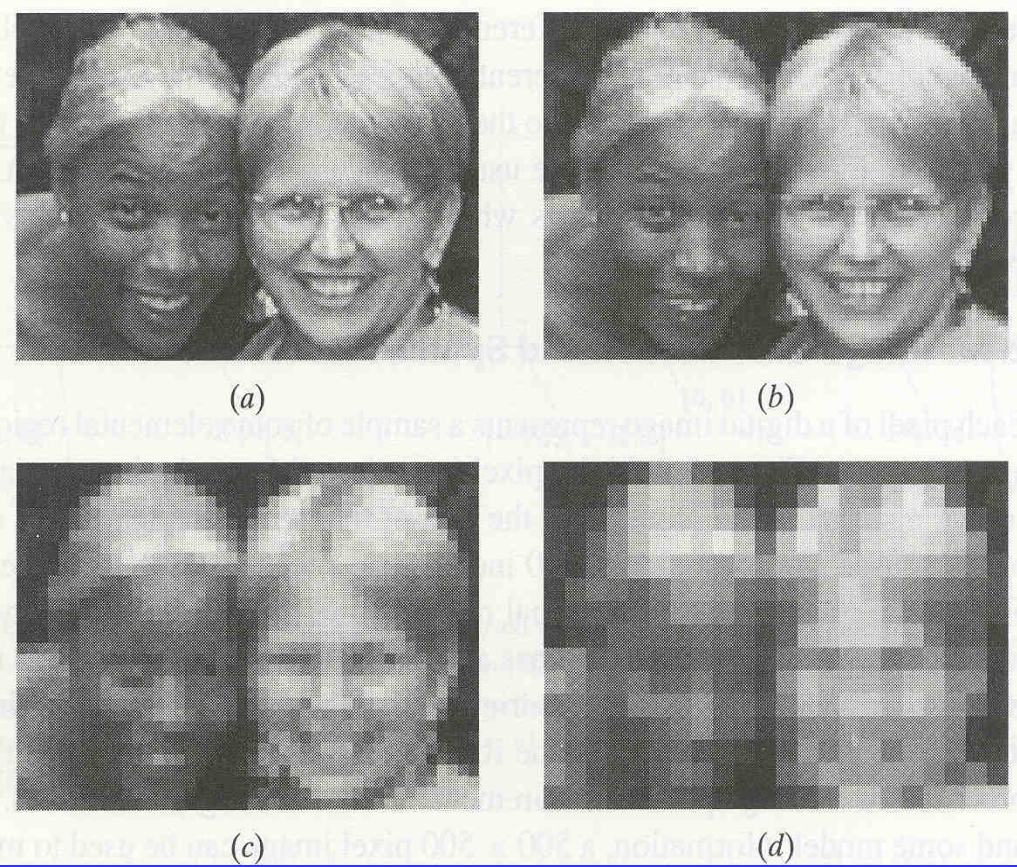

- Resolution decreases by one half in cases at left
- Human faces can be recognized at 64 x 64 pixels per face

Image Formats

- Portable gray map (PGM) older form
- GIF was early commercial version
- JPEG (JPG) is modern version
- Many others exist: **header plus data**
- Do they handle color?
- Do they provide for compression?
- Are there good packages that use them or at least convert between them?

PGM image with ASCII info.

- P2 means ASCII gray
- Comments
- W=16; H=8
- 192 is max intensity
- Can be made with editor
- Large images are usually not stored as ASCII

```
P2
# sample small picture 8 rows of 16 columns, max grey value of 192
# making an image of the word "Hi".
  16 8   192

 64 64  64  64  64  64  64  64  64 64 64  64  64 64 64 64
 64 64 128 128  64  64  64 128 128 64 64 192 192 64 64 64
 64 64 128 128  64  64  64 128 128 64 64 192 192 64 64 64
 64 64 128 128 128 128 128 128 128 64 64  64  64 64 64 64
 64 64 128 128 128 128 128 128 128 64 64 128 128 64 64 64
 64 64 128 128  64  64  64 128 128 64 64 128 128 64 64 64
 64 64 128 128  64  64  64 128 128 64 64 128 128 64 64 64
 64 64  64  64  64  64  64  64  64 64 64  64  64 64 64 64
```

PBM/PGM/PPM Codes

- **P1: ascii binary (PBM)**
- **P2: ascii grayscale (PGM)**
- **P3: ascii color (PPM)**
- **P4: byte binary (PBM)**
- **P5: byte grayscale (PGM)**
- **P6: byte color (PPM)**

JPG current popular form

- Public standard
- Allows for image compression; often 10:1 or 30:1 are easily possible
- 8x8 intensity regions are fit with basis of cosines
- Error in cosine fit coded as well
- Parameters then compressed with Huffman coding
- Common for most digital cameras

From 3D Scenes to 2D Images

- **Object**
- **World**
- **Camera**
- **Real Image**
- **Pixel Image**

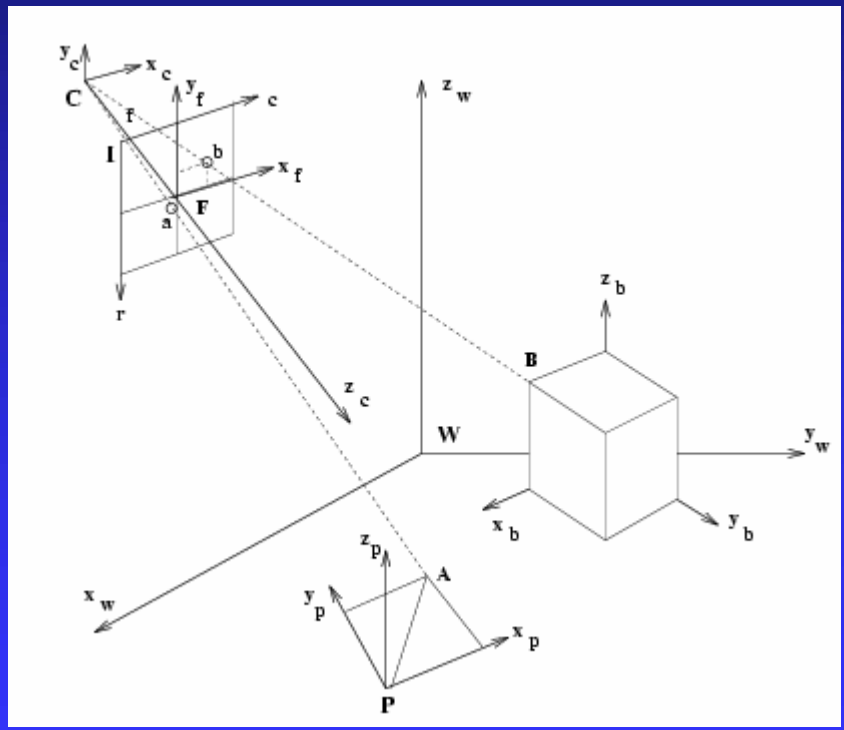

3D Sensors

- Laser range finders
- CT, MRI, and ultrasound machines
- Sonar sensors
- Tactile sensors (pressure arrays)
- Structured light sensors
- Stereo

- MRA (angiograph) showing blood flow.

Where do we go next?

So we've got an image, say a single gray-tone image.

What can we do with it?

The simplest types of analysis is binary image analysis.

Convert the gray-tone image to a binary image (0s and 1s) and perform analysis on the binary image, with possible reference back to the original gray tones in a region.