

Administrivia

If you don't have a CS account, get one! Also run “apply”, ASAP.
(Access to lab in BSE 328 will require these two steps)

Honors: See www.honors.arizona.edu/contracting.html

Course web page is now up: <http://www.cs.arizona.edu/classes/cs477/spring04>
(Linked from instructor's home page (<http://kobus.ca>))

First assignment as well as first two lectures are available on web page.

Lectures and assignments will require either connecting from a UA machine, OR a login id (“me”) and password (“vision4fun”)

TA: Scott Morris is awaiting your questions (smorris @ cs.arizona.edu)
Office 710D Office hours: TBA

Relation of course to book

Look for section numbers on slides.

My current plan is not that different from the syllabus on page xxi of the preface

Those specializing in vision should consider reading some of what we don't cover as well.

Light (review)

Light energy reaching a camera sensor has a distribution over wavelength, λ .
(*Recall from physics that wavelength is inversely related to photon energy)

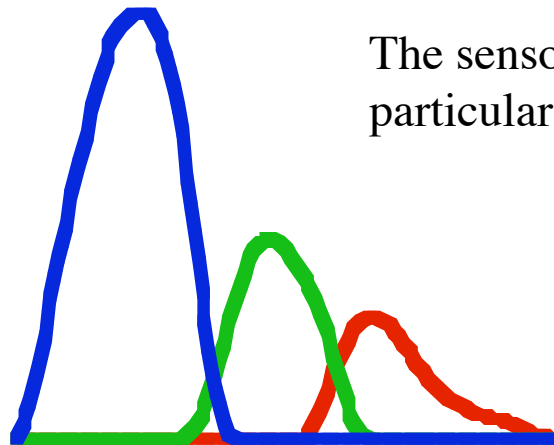


- * Things marked with red stars are optional comments in the current context. If cover them in more detail later on, then they might not be optional anymore.

Sensors

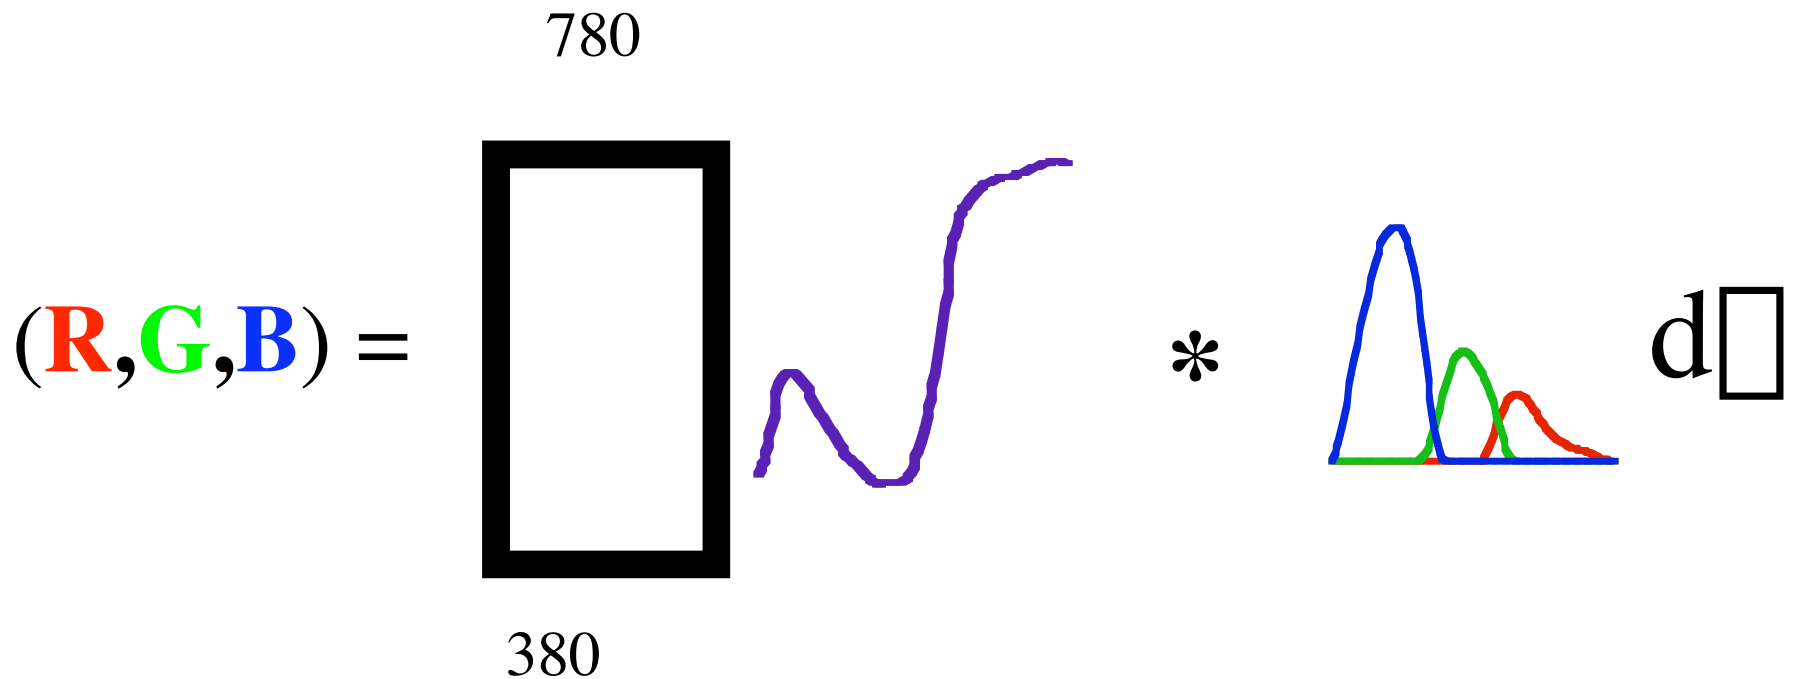
Sensors (including those in your eyes) have a varied sensitivity over wavelength

Different variations lead to different kinds of sensor responses (“colors” in a naïve sense)



The sensor sensitivity spectra for a particular color camera (they vary a lot).

Image Formation (Spectral)

$$(\mathbf{R}, \mathbf{G}, \mathbf{B}) = \int_{380}^{780} \mathbf{I}(\lambda) \mathbf{S}(\lambda) d\lambda$$


The diagram illustrates the spectral image formation equation. The equation is $(\mathbf{R}, \mathbf{G}, \mathbf{B}) = \int_{380}^{780} \mathbf{I}(\lambda) \mathbf{S}(\lambda) d\lambda$. The Spectral Response Matrix $\mathbf{S}(\lambda)$ is a 3x1 vector with rows for Red (R), Green (G), and Blue (B) channels, each with a corresponding colored curve. The Spectral Power Distribution $\mathbf{I}(\lambda)$ is a single curve showing the intensity of light across the visible spectrum (380-780 nm). The integration is performed over the wavelength λ from 380 to 780 nm.

More formally,

The response of an image capture system to a light signal $L(\lambda)$ associated with a given pixels is modeled by

$$I^{(k)} = F^{(k)}(I^{(k)}) = \int L(\lambda) R^{(k)}(\lambda) d\lambda$$

where $R^{(k)}(\lambda)$ is the sensor response function for the k^{th} channel, $I^{(k)}$ is the k^{th} channel response, and I is the k^{th} channel response linearized by the wavelength independent function $F^{(k)}$.

In this formulation, $R^{(k)}(\lambda)$ absorbs the contributions due to the aperture, focal length, sensor position in the focal plane. $F^{(k)}$ absorbs typical non-linearities such as gamma.

Discrete Version


Often we represent functions by vectors (or matrices for 2D functions). For example, a spectra might be represented by 101 samples in the range of 380 to 780 nm in steps of 4nm.

Then $L(\lambda)$ becomes the vector \mathbf{L} , $R^{(k)}(\lambda)$ becomes the vector \mathbf{R}^k , and the response (ignoring linearity issues) is given by a dot product:

$$\lambda^{(k)} = \mathbf{L} \cdot \mathbf{R}^{(k)}$$

Sensor/light interaction example

$\mathbf{R}=(0,0,1,3,7,4,2,0,0,0)$




$\mathbf{L}=(0,8,6,3,1,2,7,4,3,0)$


Multiply lined up
pairs of numbers
and then sum up

Sensor/light interaction example

$$\mathbf{R}=(0,0,1,3,7,4,2,0,0,0)$$



$$\mathbf{L}=(0,8,6,3,1,2,7,4,3,0)$$



$$\mathbf{L} \cdot \mathbf{R} =$$

$$(0*0, 0*8, 0*6, 1*3, 3*1, 7*2, 4*7, 2*4, 0*3, 0*0) \\ = (0, 0, 6, 9, 7, 8, 14, 0, 0, 0)$$

$$\mathbf{L} \cdot \mathbf{R} = 0 + 0 + 6 + 9 + 7 + 8 + 14 \\ = 44$$

Image Formation (Spectral)

Note that (other than $F^{(k)}$) image formation is linear.

Formally this means if:

$$L_1(\square) \square > \square_1^{(k)} \quad \text{and} \quad L_2(\square) \square > \square_2^{(k)}$$

Then:

$$aL_1(\square) + bL_2(\square) \square > a\square_1^{(k)} + b\square_2^{(k)}$$

Image Formation (Spectral)

$F^{(k)}$ is often ignored (assumed to be the identity), but this is not a safe assumption, especially when color or radiometric measurements matter. To compensate for the non-linearity of CRT display monitors, camera manufactures will “gamma” correct the signal, typically raising the signal (assuming it is between 0 and 1) to the $1/2.2$ power.

Note that in such an image, a number twice as large does not mean that the light had twice the power!


To linearize RGB's from such a signal we compute:

$$p = F(v) = 255 * (v/255)^{2.2}$$


Image Formation (Geometric)

Pinhole cameras

- Abstract camera model--box with a small hole in it
- Pinhole cameras work for deriving algorithms--a real camera needs a lens




Distant objects are smaller



Parallel lines meet*

Common to draw film plane
in front of the focal point
(virtual image)




*Exceptions?

Vanishing points

- Each set of parallel lines (=direction) meets at a different point
 - The *vanishing point* for this direction
- Sets of parallel lines on the same plane lead to *collinear* vanishing points.
 - The line is called the *horizon* for that plane
 - Standard horizon is the horizon of the ground plane.
- One way to spot fake images
 - scale and perspective don't work
 - vanishing points behave badly
 - supermarket tabloids are a great source
 - see also last question on Assignment 1

Example: The figure below is claimed to provide a perspective view of two identical cubes, with faces a and a' , and faces b and b' being parallel. Provide two (different kinds of) reasons why this could not be a real perspective drawing of the geometry described, marking any needed explanatory lines on the figure.




Representing Transformations

- Need mathematical representation for mapping points from the world to an image (and later, from an image taken by one camera to another).
- Represent linear transformations by matrices
- To transform a point, represented by a vector, multiply the vector by the appropriate matrix.
- To transform lines, transform endpoints
- To transform polygons, transform vertices

Transformation examples in 2D

- Scale (stretch) by a factor of k




$$M = \begin{vmatrix} k & 0 \\ 0 & k \end{vmatrix}$$

($k = 2$ in the example)

Transformation examples in 2D

- Scale by a factor of (S_x, S_y)




$$M = \begin{vmatrix} S_x & 0 \\ 0 & S_y \end{vmatrix}$$

(Above, $S_x = 1/2$, $S_y = 1$)

Transformation examples in 2D

- Rotate around origin by θ (Orthogonal)




$$M = \begin{vmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{vmatrix}$$

(Above, $\theta=90^\circ$)

Transformation examples in 2D

- Flip over y axis

(Orthogonal)




$$M = \begin{vmatrix} -1 & 0 \\ 0 & 1 \end{vmatrix}$$

Flip over x axis is ?

2D Transformations

- Translation ($\mathbf{P}_{\text{new}} = \mathbf{P} + \mathbf{T}$)



$\mathbf{M} = ?$

Homogenous Coordinates

- Represent 2D points by 3D vectors
- $(x,y) \rightarrow (x,y,1)$
- Now a multitude of 3D points (x,y,W) represent the same 2D point, $(x/W, y/W, 1)$
- Represent 2D transforms with 3 by 3 matrices
- Can now represent translations by matrix multiplications

2D Scale in H.C.

$$\mathbf{M} = \begin{vmatrix} S_x & 0 & 0 \\ 0 & S_y & 0 \\ 0 & 0 & 1 \end{vmatrix}$$

2D Rotation in H.C.

$$M = \begin{vmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{vmatrix}$$

2D Translation in H.C.

- $\mathbf{P}_{\text{new}} = \mathbf{P} + \mathbf{T}$
- $(x', y') = (x, y) + (t_x, t_y)$

$$\mathbf{M} = \begin{vmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{vmatrix}$$

Transformations in 3D

- Homogeneous coordinates now have four components - traditionally, (x, y, z, w)
 - ordinary to homogeneous: $(x, y, z) \rightarrow (x, y, z, 1)$
 - homogeneous to ordinary: $(x, y, z, w) \rightarrow (x/w, y/w, z/w)$
- Again, translation can be expressed as a multiplication.

Transformation examples in 3D

- Translation:

$$\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} tx \\ ty \\ tz \\ 1 \end{bmatrix}$$

- Anisotropic scaling:

$$\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \begin{bmatrix} sx & 0 & 0 \\ 0 & sy & 0 \\ 0 & 0 & sz \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

Rotations in 3D

- 3 degrees of freedom
- $\text{Det}(\mathbf{R})=1$
- Orthogonal
- Many representations are possible.
- Our representation: rotate about coordinate axes in sequence.
- Sequence of axes is arbitrary, but choice does affect the angles used (cannot use same angles with different order).
- Sign of rotation follows the Right Hand Rule--point thumb along axis in direction of increasing ordinate--then fingers curl in the direction of positive rotation).

Optional material **Not** covered in class

Rotations in 3D

- About x-axis

$$M = \begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Optional material **Not** covered in class

Rotations in 3D

- About y-axis

$$M = \begin{vmatrix} \cos\theta & 0 & -\sin\theta & 0 \\ 0 & 1 & 0 & 0 \\ \sin\theta & 0 & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Optional material **Not** covered in class

Rotations in 3D

- About z-axis

$$M = \begin{vmatrix} \cos\theta & -\sin\theta & 0 & 0 \\ \sin\theta & \cos\theta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Optional material **Not** covered in class

Rotations in 3D

- About X axis

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

- 90 degrees about X axis?

Optional material **Not** covered in class

Rotations in 3D

- About X axis

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

- 90 degrees about X axis

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Optional material **Not** covered in class

Rotations in 3D

- About Y axis

$$\begin{vmatrix} \cos\theta & 0 & \sin\theta & 0 \\ 0 & 1 & 0 & 0 \\ -\sin\theta & 0 & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

- 90 degrees about Y-axis?

Optional material **Not** covered in class

Rotations in 3D

- About Y axis

$$\begin{vmatrix} \cos\theta & 0 & -\sin\theta & 0 \\ 0 & 1 & 0 & 0 \\ \sin\theta & 0 & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

- 90 degrees about Y axis

$$\begin{vmatrix} 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Optional material **Not** covered in class

Rotations in 3D

- 90 degrees about X then Y

$$\begin{array}{c} \left| \begin{array}{cccc} 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{array} \right| \left| \begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{array} \right| = ? \\ \text{Y rot} \qquad \qquad \text{X rot} \end{array}$$

Optional material **Not** covered in class

Rotations in 3D

- 90 degrees about X then Y

$$\begin{array}{c} \left| \begin{array}{cccc} 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{array} \right| \left| \begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{array} \right| = \left| \begin{array}{cccc} 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{array} \right| \\ \text{Y rot} \qquad \qquad \text{X rot} \end{array}$$

Optional material **Not** covered in class

Rotations in 3D

- 90 degrees about X then Y

$$\begin{vmatrix} 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} \begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} = \begin{vmatrix} 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Y rot

X rot

- 90 degrees about Y then X

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} \begin{vmatrix} 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} = ?$$

X rot

Y rot

Optional material **Not** covered in class

Rotations in 3D

- 90 degrees about X then Y

$$\begin{vmatrix} 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} \begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} = \begin{vmatrix} 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Y rot

X rot

- 90 degrees about Y then X

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} \begin{vmatrix} 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} = \begin{vmatrix} 0 & 0 & -1 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$


X rot

Y rot

\neq


Optional material **Not** covered in class

Rotation about an arbitrary axis



Optional material **Not** covered in class


Rotation about an arbitrary axis



Strategy--rotate A to Z
axis, rotate about Z axis,
rotate Z back to A.

Optional material **Not** covered in class

Rotation about an arbitrary axis



Tricky part:


rotate A to Z
axis

Two steps.

- 1) Rotate about x to xz plane
- 2) Rotate about y to Z axis.

Optional material **Not** covered in class

Rotation about an arbitrary axis



Tricky part:
rotate A to Z
axis


Two steps.

- 1) Rotate about X to xz plane
- 2) Rotate about Y to Z axis.

As A rotates into the xz plane, its projection onto the YZ plane (red line) rotates through the same angle which is easily calculated.

Optional material **Not** covered in class

Rotation about an arbitrary axis



$$d = \sqrt{a_y^2 + a_z^2}$$


$$\sin \alpha_x = a_y / d$$

$$\cos \alpha_x = a_z / d$$

No need to compute angles,
just put sines and cosines into
rotation matrices

Optional material **Not** covered in class

Rotation about an arbitrary axis



Apply $R_x(\varphi_x)$ to A and renormalize to get A'

$R_y(\varphi_y)$ should be easy, but note that it is clockwise.

Optional material **Not** covered in class

Rotation about an arbitrary axis

Final form is

$$R_x(\alpha_x) R_y(\alpha_y) R_z(\alpha_z) R_y(\alpha_y) R_x(\alpha_x)$$

Projections

- Want to think about geometric image formation as a mathematical transformation taking points in the 3D world and mapping them into an image plane.
- Mathematical definition of a projection: $PP=P$
- (Doing it a second time has no effect).
- Generally rank deficient (non-invertable)--exception is $P=I$
- Transformation loses information (e.g., depth)
- Given a 2D image, there are many 3D worlds that could have lead to it.