Foundations of Computer Vision

Prof. Feng Liu

Winter 2022

http://www.cs.pdx.edu/~fliu/courses/cs410/

01/04/2022

Today

- Course overview
  - Computer vision
  - Admin. Info
- Computer Vision at PSU
- Image representation
- Color

https://electricautonomy.ca/2021/06/01/autonomous-driving-canada/

https://en.wikipedia.org/wiki/Amazon_Go

https://www.digitaltrends.com/computing/apple-faceid-scan-retinas-security/

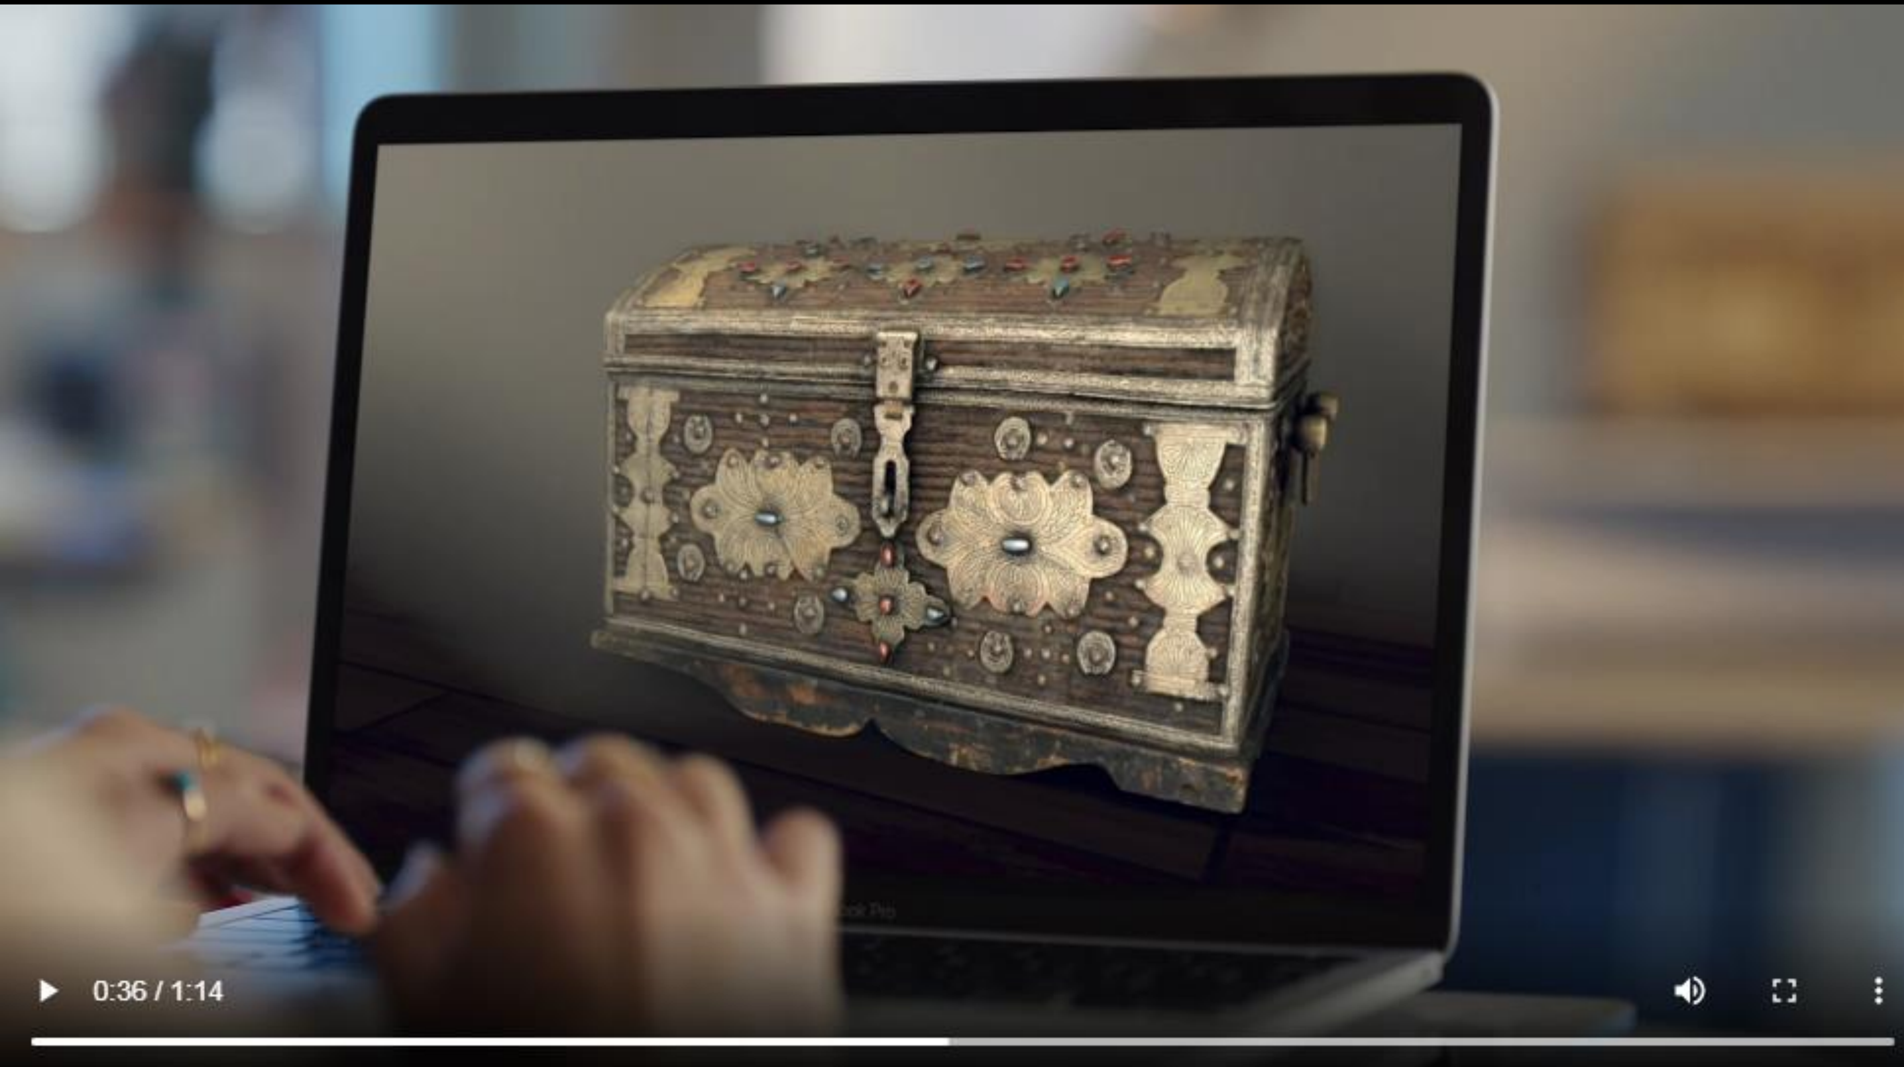

https://developer.apple.com/augmented-reality/object-capture/

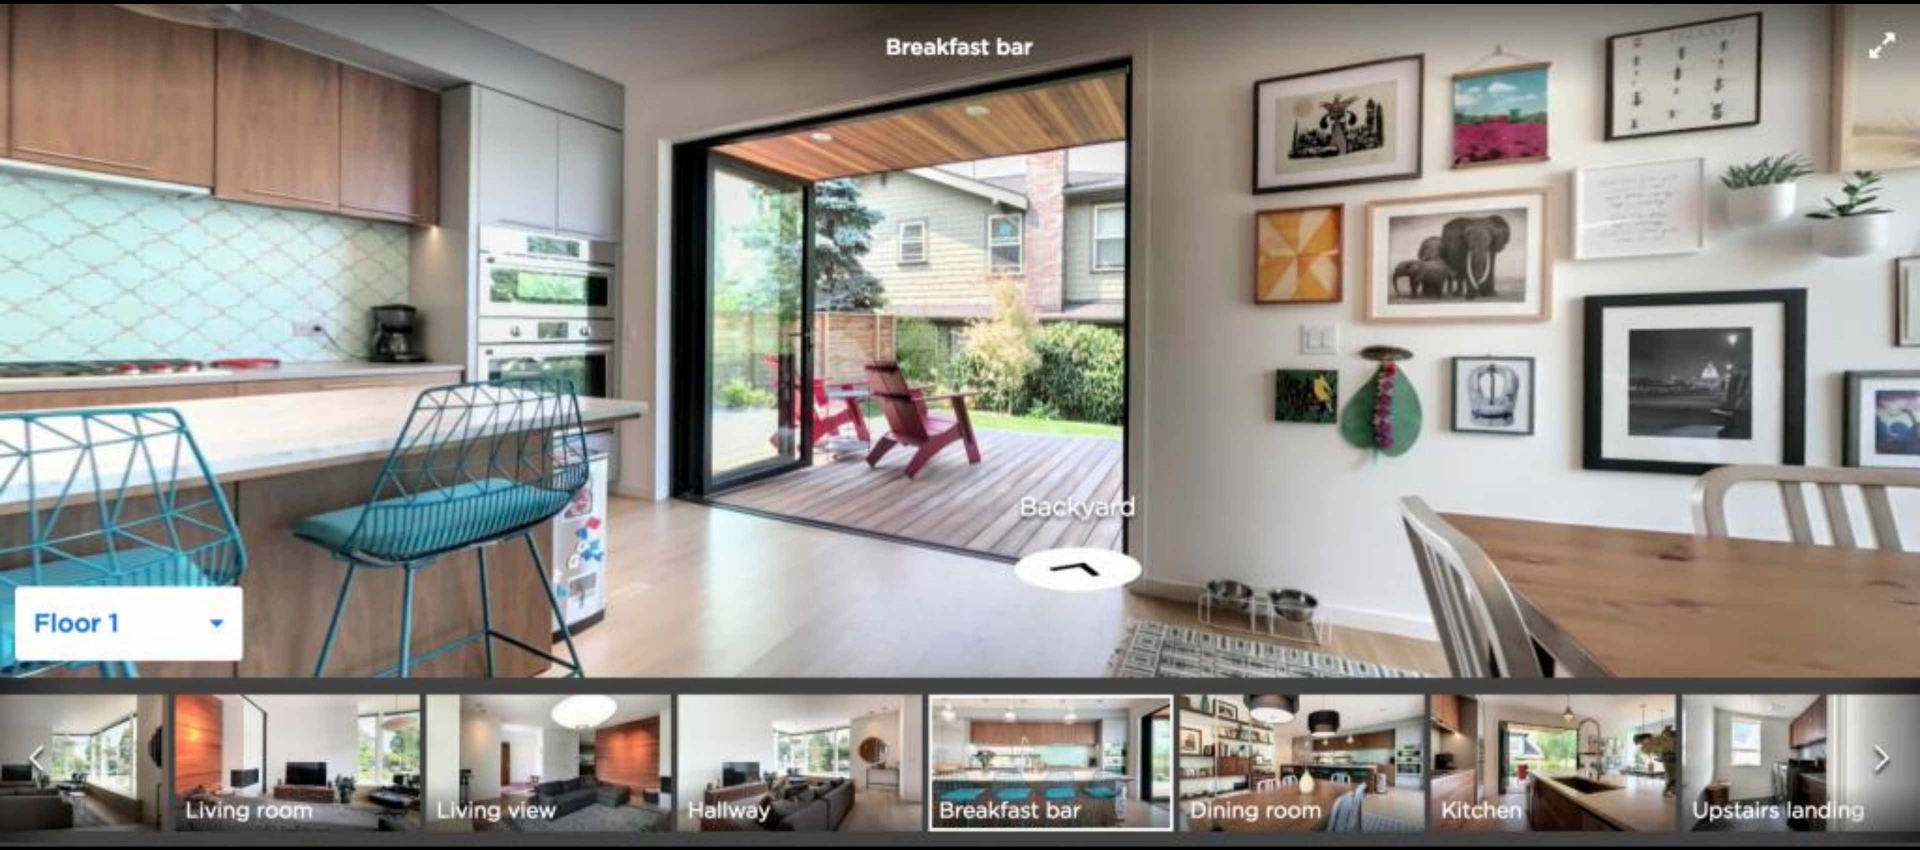

https://www.zillow.com/tech/behind-zillow-3d-home-backend-algorithms/

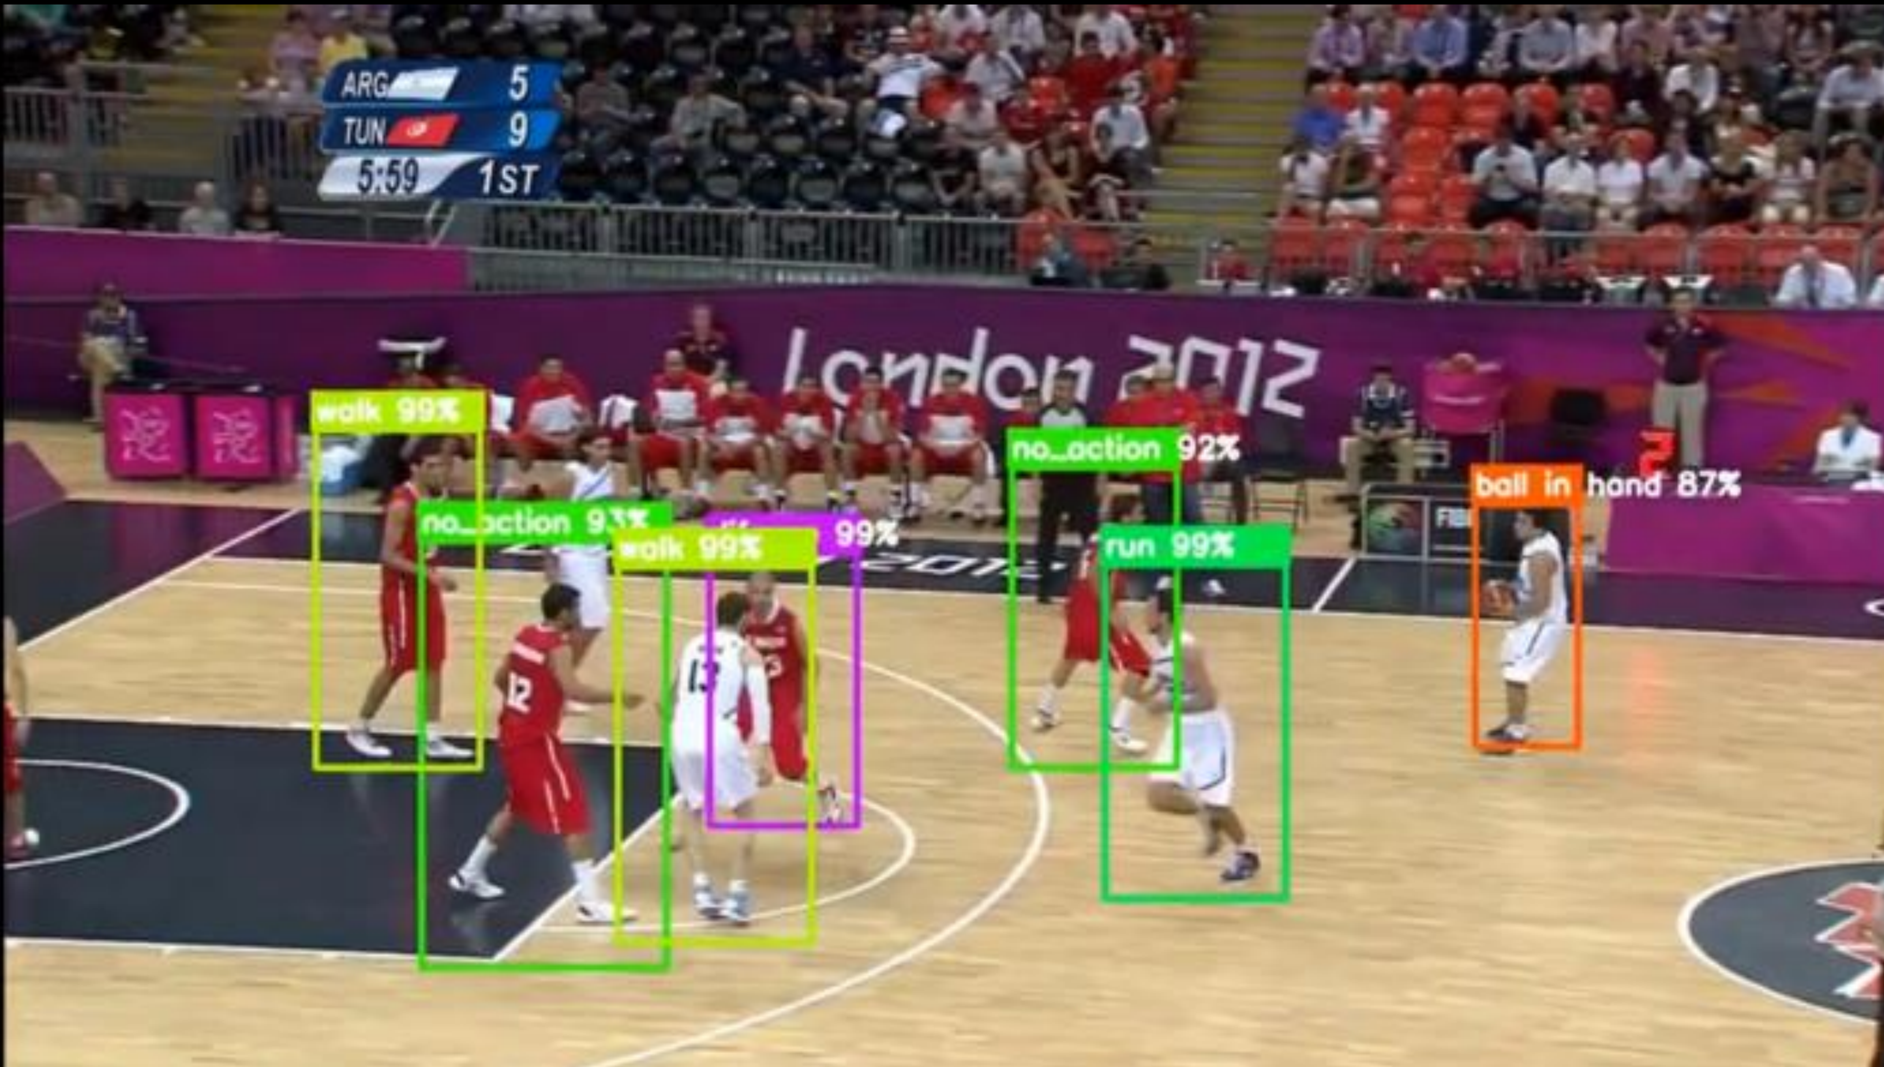

https://becominghuman.ai/machine-learning-applied-to-sports-60345867a5c1

Course objective

- A broad overview of computer vision
- Cover basic, common algorithms
  - Some advanced topics will be covered in CS 410/510 Computational Photography in Spring
- Exercise with computer vision toolkits
- Focus on the development of computer vision applications

People

- Lecturer: Prof. Feng Liu
  - Office Hours: by appointment
  - fliu@pdx.edu
- TA: Zhan Li
  - Location: https://pdx.zoom.us/j/84063267395
  - Office Hours: TR 2:30-3:30pm
  - lizhan@pdx.edu

Web and Computer Account

- Course website
  - http://www.cs.pdx.edu/~fliu/courses/cs410/
  - Class mailing list
    - Google Chat / Space
    - An invitation has been sent to your pdx.edu account
- Everyone needs a Computer Science department computer account to use CS lab machines
  - Get account at CAT
  - http://cat.pdx.edu

Textbooks & Readings

- No textbooks are required, but recommend
  - Computer Vision: Algorithms and Applications
    - By R. Szeliski
    - Available online, free
  - Learning OpenCV 3: Computer Vision in C++ with the OpenCV Library,
    - By Adrian Kaehler and Gary Bradski
- Papers recommended by the lecturer

Grading

- Homework: 65%
  - 3 assignments

- Project: 35%

Programming tools

- Python
  - Required for homework assignments
  - Recommended for the project
- C/C++ under Windows
  - OK for the project
  - Other graphics and vision libraries
- Others
  - OK for the project
  - As long as it works for you

OpenCV

- **Open** Source **C**omputer **V**ision library
  - http://opencv.org/
  - Perhaps the most popular toolkits for computer vision
  - Provides Python and C/C++ APIs for a wide range of vision algorithms
  - You can use the OpenCV library
    - unless prohibited for some homework assignments

Admin Questions?

Today

- Course overview
  - Computer vision
  - Admin. Info
- Computer Vision at PSU
- Image representation
- Color

Visual Computing @ PSU

- PSU Computer Graphics and Vision Lab
  - http://graphics.cs.pdx.edu/
- Projects
  - Video stabilization
  - Video frame interpolation
  - Shadow editing
  - Super resolution
  - Panorama synthesis
  - Video segmentation
  - Video interaction
  - Visual quality assessment
  - Stereoscopic imaging
  - VR

Stabilization

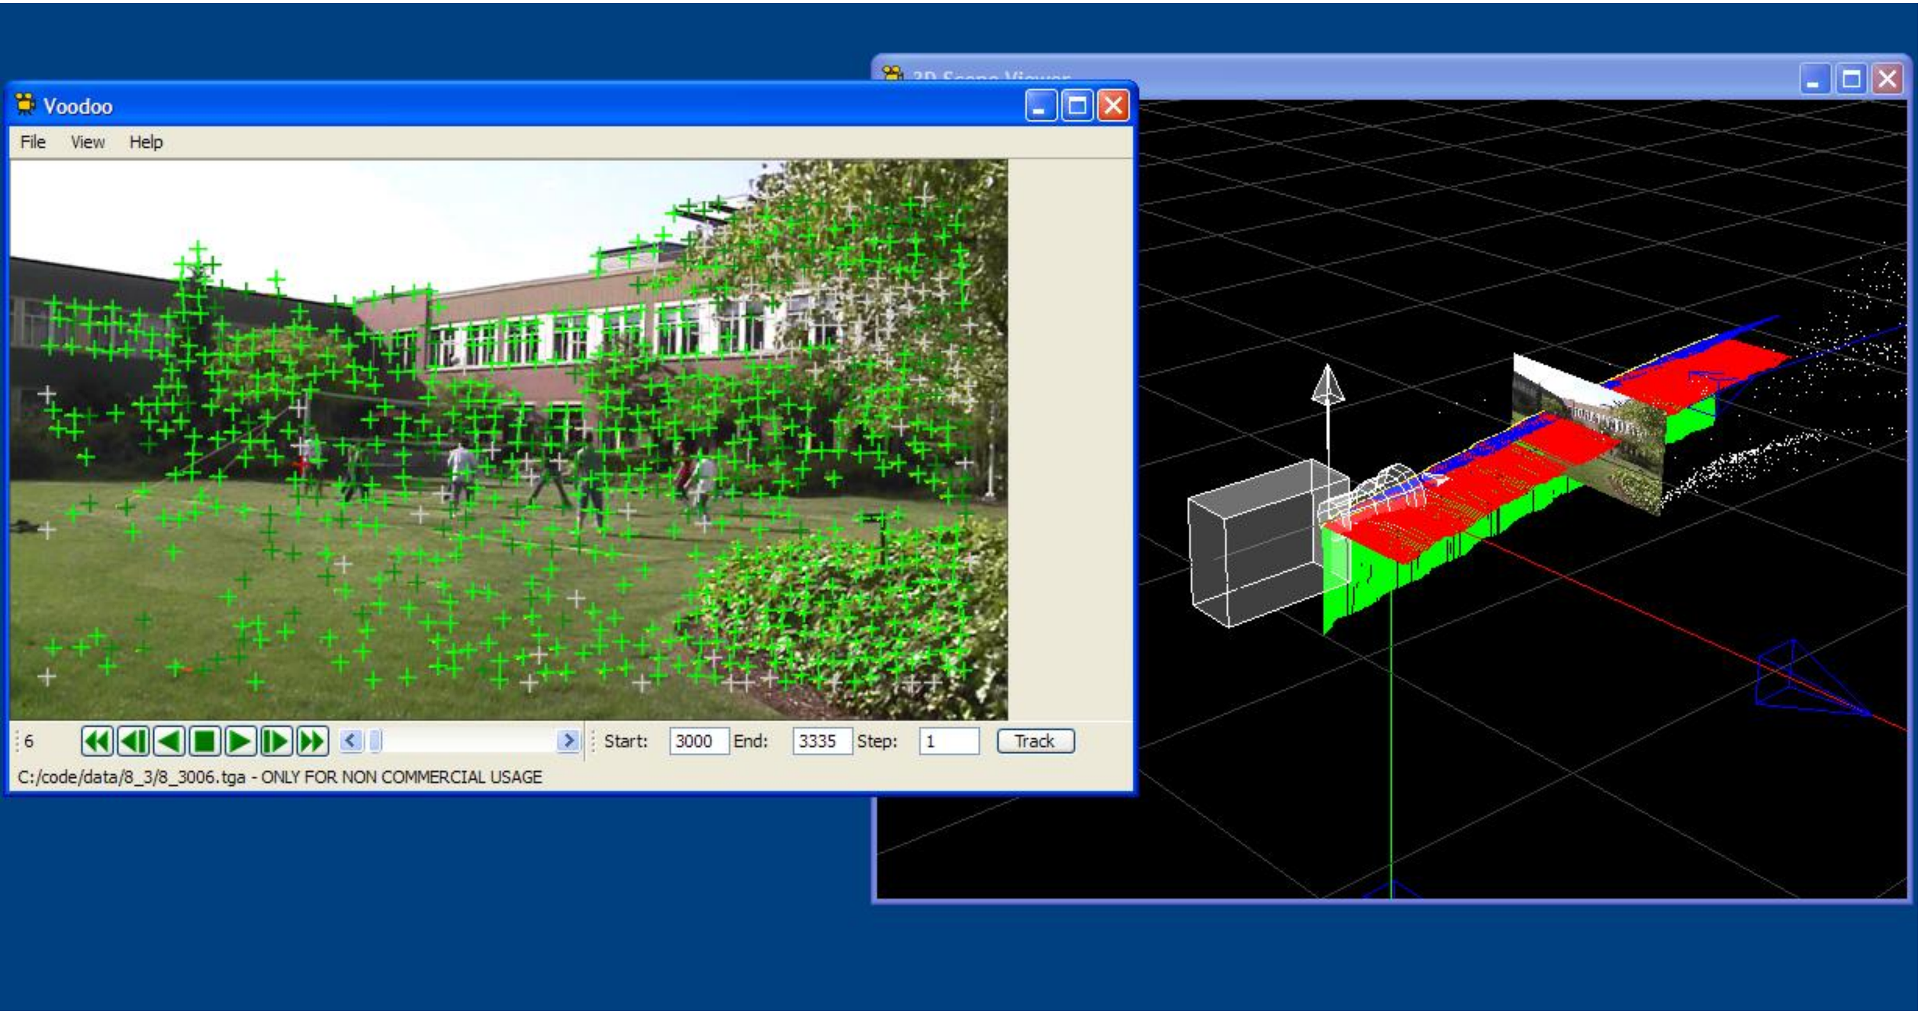

http://web.cecs.pdx.edu/~fliu/project/subspace_stabilization/

Neural Video Frame Interpolation

http://web.cecs.pdx.edu/~fliu/project/ctxsyn

Shadow Editing

Input

Output

Super-resolution

Low resolution High resolution

Panorama from Video

Segmentation

http://web.cecs.pdx.edu/~fliu/project/motionseg.htm

Visual saliency analysis

http://web.cecs.pdx.edu/~fliu/project/stereo-saliency

Visual quality assessment

http://web.cecs.pdx.edu/~fliu/project/thirdsrule.htm

Stereoscopic imaging

VR Video Editing

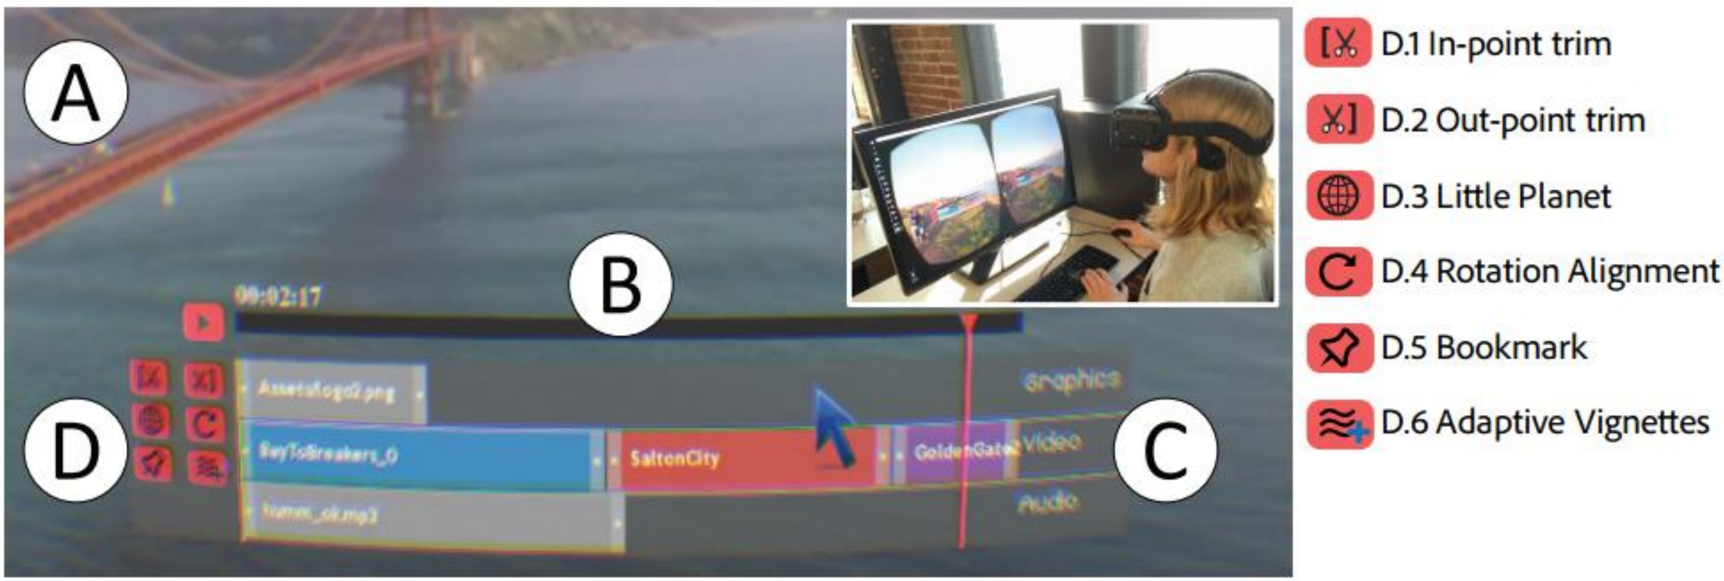

Today

- Course overview
  - Admin. Info
  - Computer vision
- Visual Computing at PSU
- Image representation
- Color

A Digital Image is a Matrix of Pixels

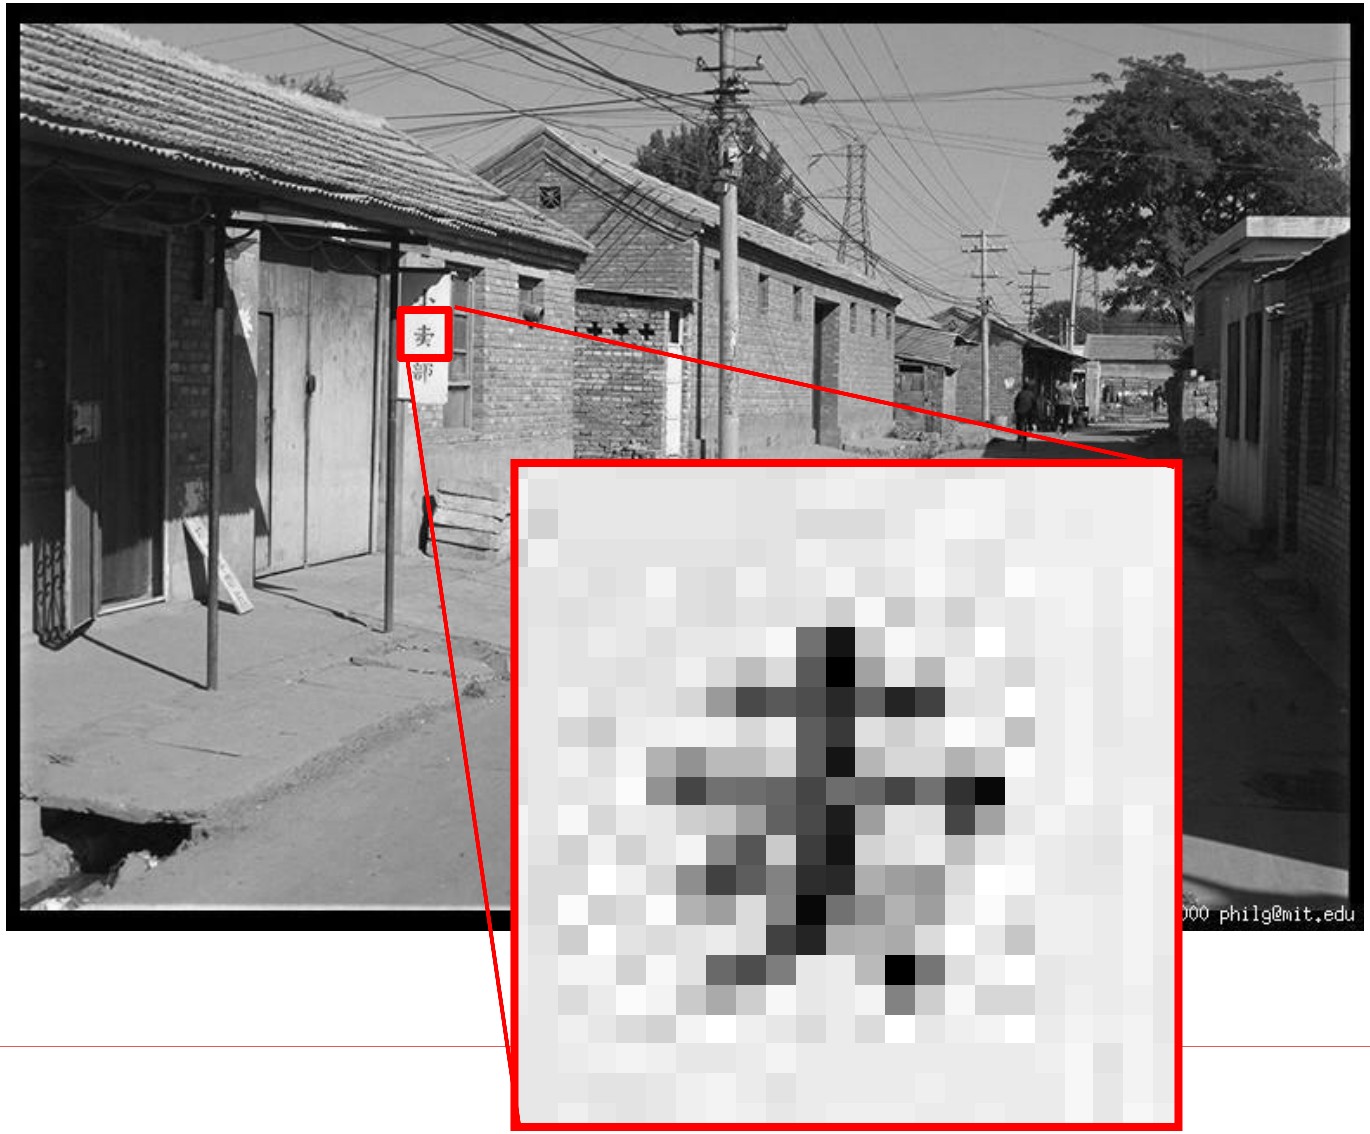

Slide credit: C. Dyer

Source: D. Hoiem

A Digital Image is a Matrix of Pixels

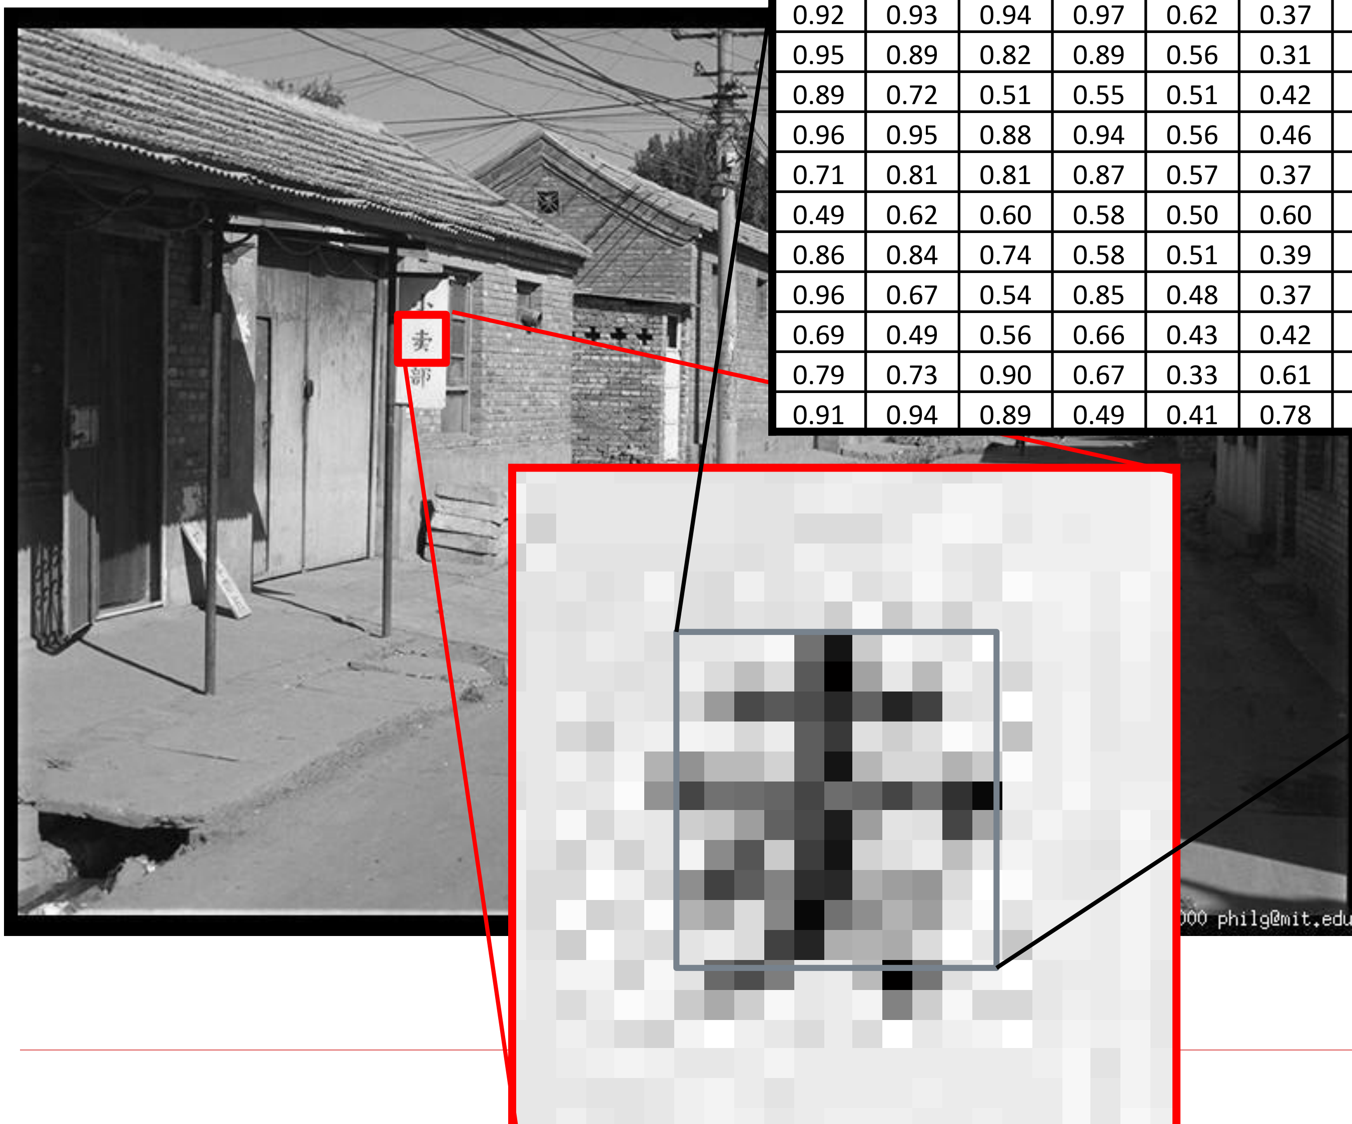

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Slide credit: C. Dyer

Source: D. Hoiem

Discretization Issues

- Can only store a finite number of pixels
  - Choose your target physical image size, choose your resolution (pixels per inch, or dots per inch, dpi), determine width/height in pixels necessary
  - Storage space goes up with square of resolution
    - 600dpi has 4× more pixels than 300dpi
- Can only store a finite range of intensity values
  - Typically referred to as *depth* - number of bits per pixel
    - Directly related to the number of colors available and typically little choice
    - Most common depth is 8, but also sometimes see 16 for grey
  - Also concerned with the minimum and maximum intensity – dynamic range
- What is enough resolution and enough depth?
  - Retina display?

Perception of Intensity

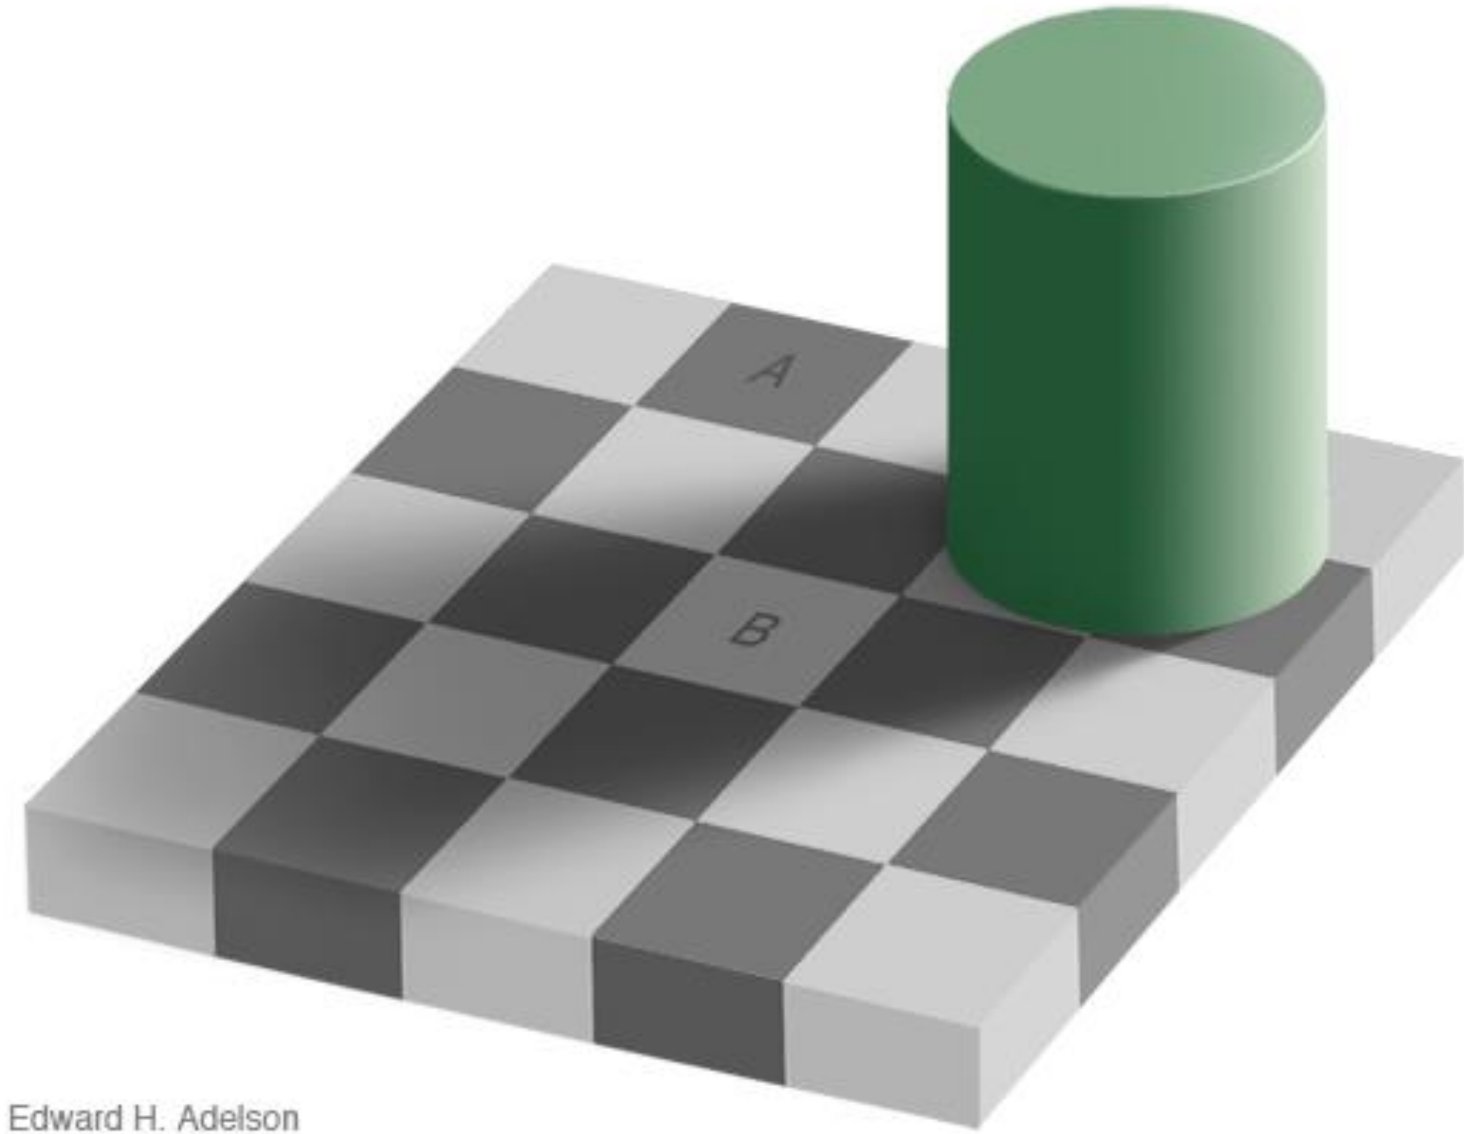

Slide credit: C. Dyer

Perception of Intensity

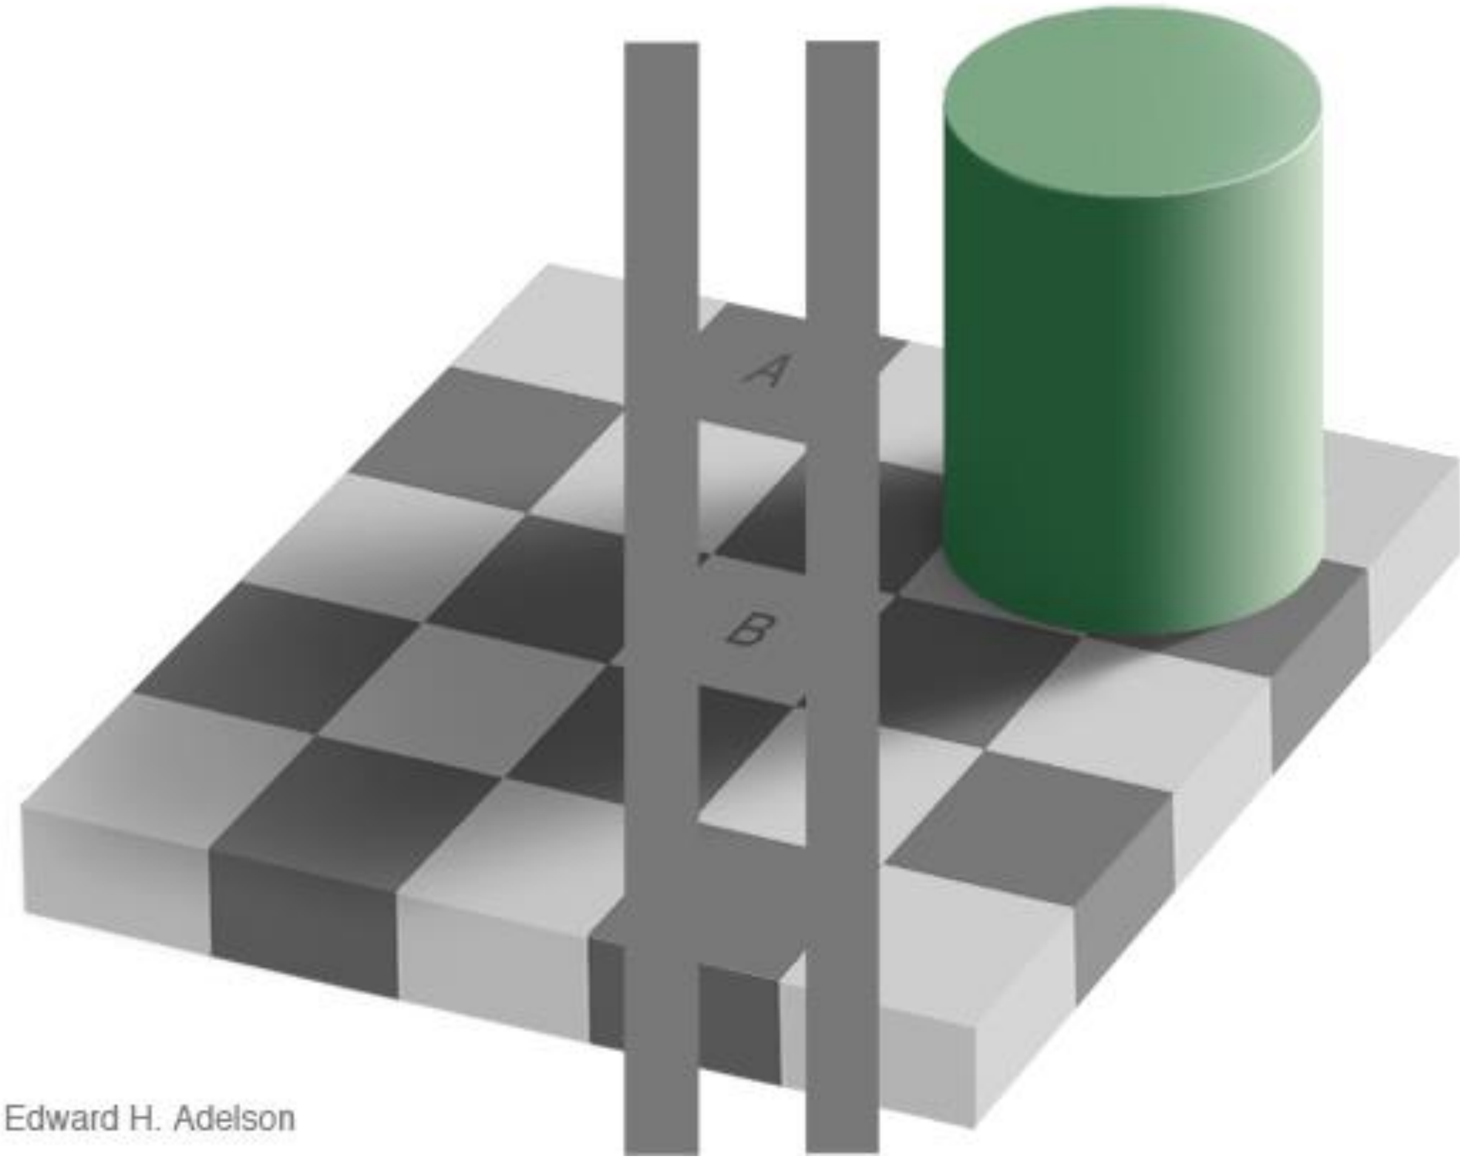

Slide credit: C. Dyer

Color Image

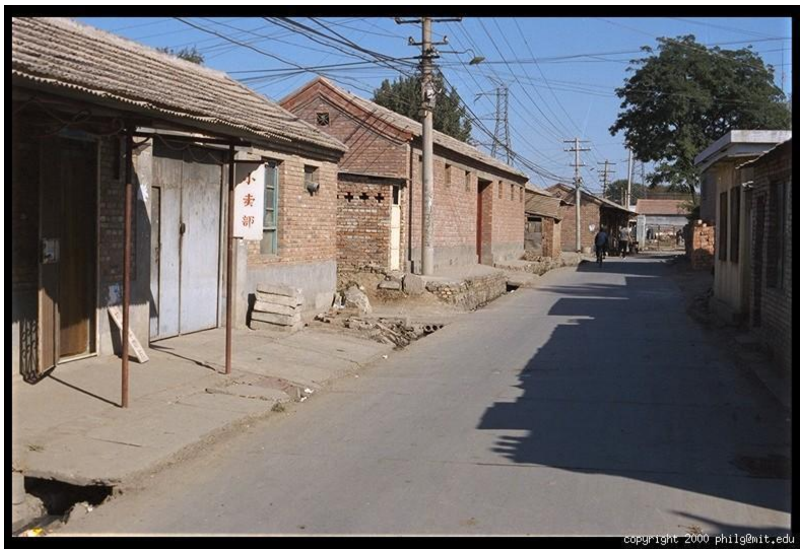

Slide credit: C. Dyer

Source: D. Hoiem

Light and Color

- The frequency of light determines its “color”
  - Wavelength is related
  - Energy also related
- We care about wavelengths in the visible spectrum: between the infra-red (700nm) and the ultra-violet (400nm)

Color and Wavelength

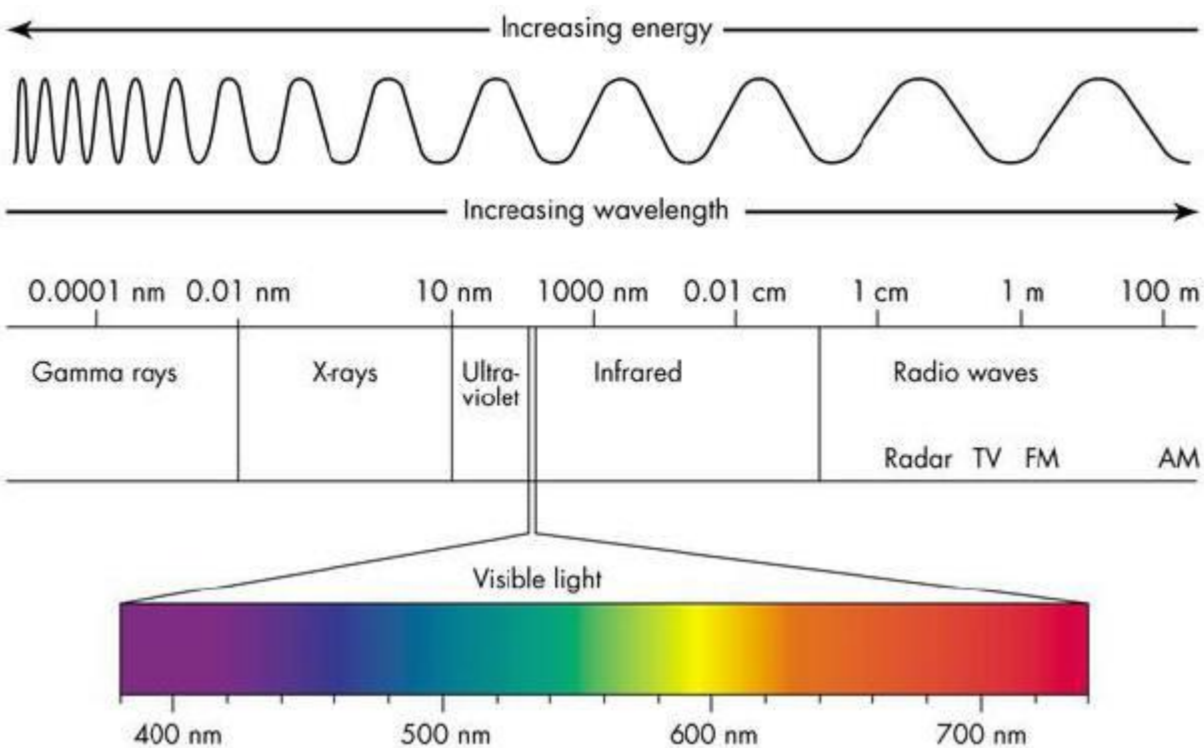

Light and Color

- The frequency of light determines its “color”
  - Wavelength is related
  - Energy also related
- We care about wavelengths in the visible spectrum: between the infra-red (700nm) and the ultra-violet (400nm)
- Describe incoming light by a *spectrum*
  - Intensity of light at each frequency
  - A graph of intensity vs. frequency

Normal Daylight

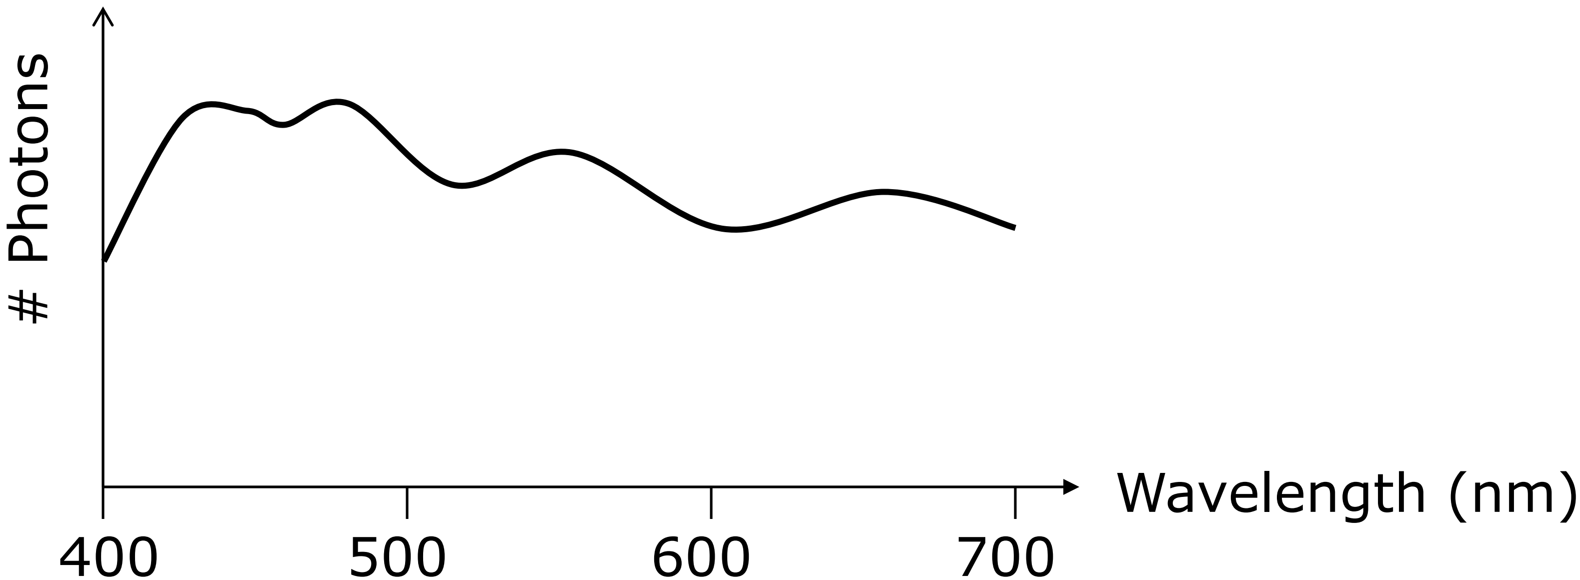

- Note the hump at short wavelengths - the sky is blue

Seeing in Color

- The eye contains *rods* and *cones*
  - Rods work at low light levels and do not see color
    - That is, their response depends only on how many photons, not their wavelength
  - Cones come in three types (experimentally and genetically proven), each responds in a different way to frequency distributions

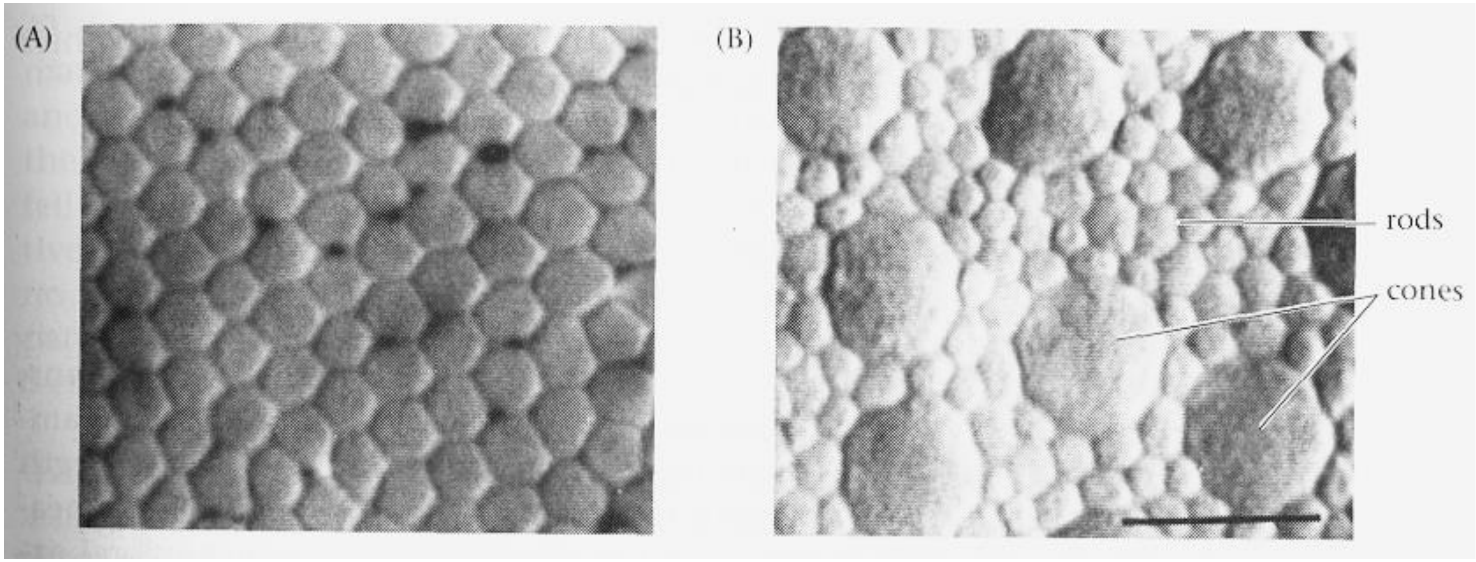

Color receptors

- Each cone type has a different sensitivity curve
  - Experimentally determined in a variety of ways
- For instance, the L-cone responds most strongly to red light
- "Response" in your eye means nerve cell firings
- How you interpret those firings is not so simple …

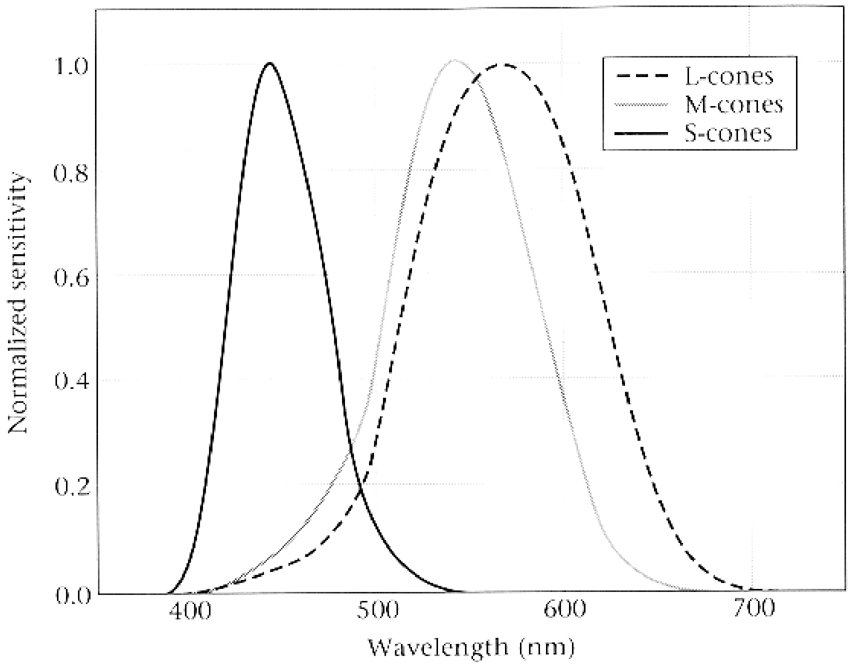

Color Perception

- How your brain interprets nerve impulses from your cones is an open area of study, and deeply mysterious
- Colors may be perceived differently:
  - Affected by other nearby colors
  - Affected by adaptation to previous views
  - Affected by “state of mind”
- Experiment:
  - Subject views a colored surface through a hole in a sheet, so that the color looks like a film in space
  - Investigator controls for nearby colors, and state of mind

The Same Color?

Image credit: https://www.mentalfloss.com/article/54448/5-color-illusions-and-why-they-work

Color Deficiency

- Some people are missing one type of receptor
  - Most common is red-green color blindness in men
  - Red and green receptor genes are carried on the X chromosome - most red-green color blind men have two red genes or two green genes
- Other color deficiencies
  - Anomalous trichromacy, Achromatopsia, Macular degeneration
  - Deficiency can be caused by the central nervous system, by optical problems in the eye, injury, or by absent receptors

Color Deficiency

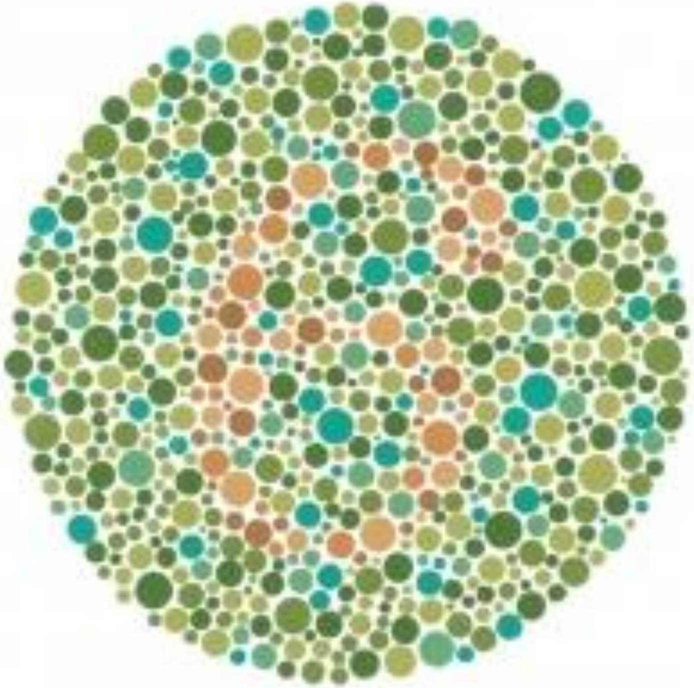

Trichromacy

- Experiments show that it is possible to match almost all colors using only three primary sources - *the principle of trichromacy*
  - Sometimes, have to add light to the *target*

- In practical terms, this means that if you show someone the right amount of each primary, they will perceive the right color

- This was how experimentalists knew there were 3 types of cones

Color Spaces

- The principle of trichromacy means that the colors displayable are all the linear combination of primaries
- Taking linear combinations of R, G and B defines the *RGB color space*
  - the range of perceptible colors generated by adding some part of R, G and B
- If R, G and B correspond to a monitor’s phosphors (monitor RGB), then the space is the range of colors displayable on the monitor

RGB Color Space

- Demo

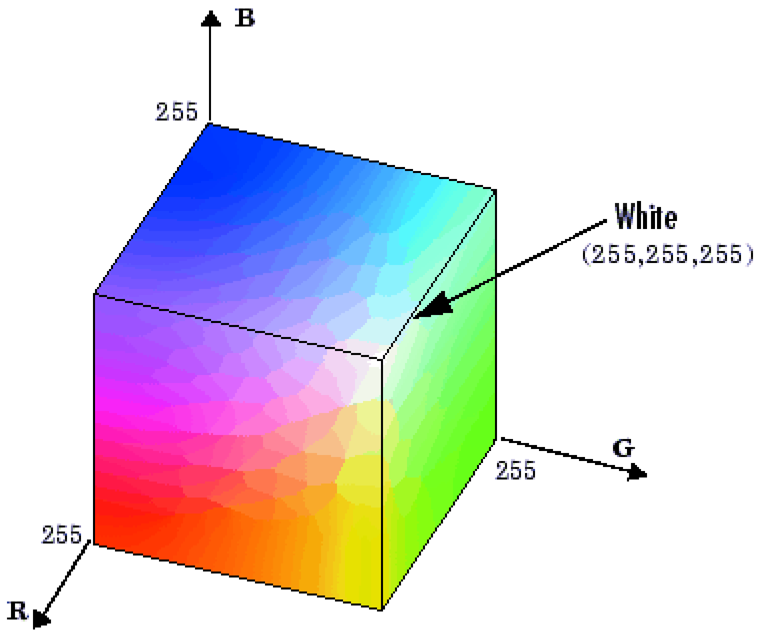

Problems with RGB

- Can only represent a small range of all the colors humans are capable of perceiving (particularly for monitor RGB)
- It isn’t easy for humans to say how much of RGB to use to make a given color
  - How much R, G and B is there in “brown”? (Answer: .64,.16, .16)
- Perceptually non-uniform

CIE XYZ Color Space

- Imaginary primaries
  - X, Y, Z
  - Y component intended to correspond to intensity
  - Cannot produce the primaries – need negative light!
- Defined in 1931 to describe the full space of perceptible colors
  - Revisions now used by color professionals
- Most frequently set $x=X/(X+Y+Z)$ and $y=Y/(X+Y+Z)$
  - x,y are coordinates on a constant brightness slice

Standard RGB $\boxed{\leftrightarrow}$ XYZ

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.4124 & 0.3576 & 0.1805 \\ 0.2126 & 0.7151 & 0.0721 \\ 0.0193 & 0.1192 & 0.9505 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

- Note that each matrix is the inverse of the other
- Recall, Y encodes brightness, so the matrix tells us how to go from RGB to grey

Accurate Color Reproduction

- Device-dependent RGB space
- High quality graphic design applications, and even some monitor software, offers accurate color reproduction
- A color calibration phase is required:
  - Fix the lighting conditions under which you will use the monitor
  - Fix the brightness and contrast on the monitor
  - Determine the monitor's γ
  - Using a standard color card, match colors on your monitor to colors on the card: This gives you the matrix to convert your monitor's RGB to XYZ
  - Together, this information allows you to accurately reproduce a color specified in XYZ format

More Linear Color Spaces

- Monitor RGB: primaries are monitor phosphor colors, primaries and color matching functions vary from monitor to monitor
- sRGB: A new color space designed for web graphics
- YIQ: mainly used in television
  - Y is (approximately) intensity, I, Q are chromatic properties
  - Linear color space; hence there is a matrix that transforms XYZ coords to YIQ coords, and another to take RGB to YIQ

HSV Color Space (Alvy Ray Smith, 1978)

- Hue: the color family: red, yellow, blue…
- Saturation: The purity of a color: white is totally unsaturated
- Value: The intensity of a color: white is intense, black isn't
- Space looks like a cone
  - Parts of the cone can be mapped to RGB space
- Not a linear space, so no linear transform to take RGB to HSV
  - But there is an algorithmic transform

HSV Color Space

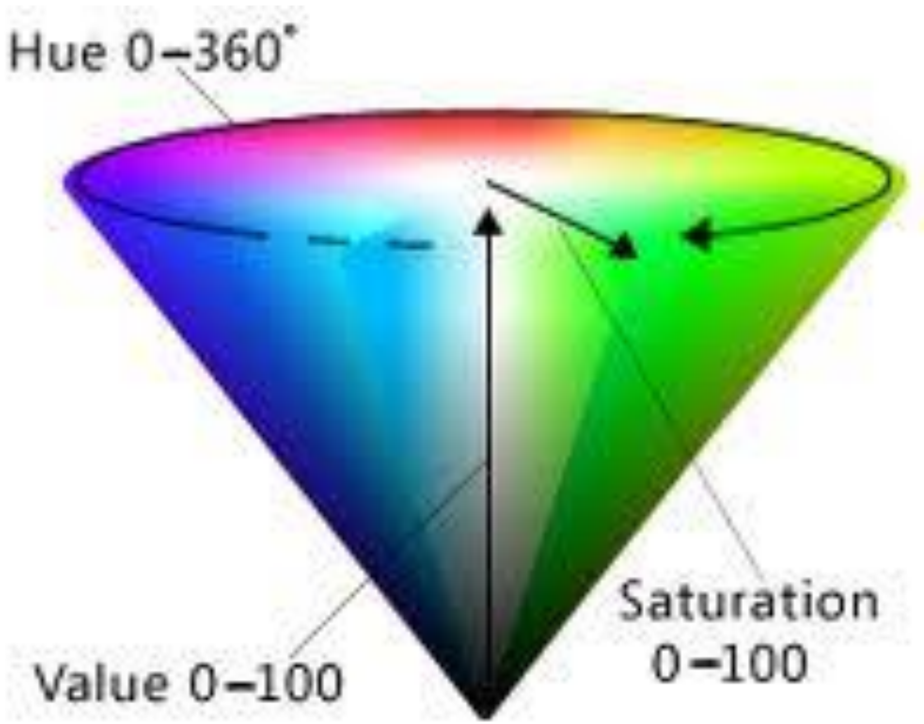

MacAdam Ellipses

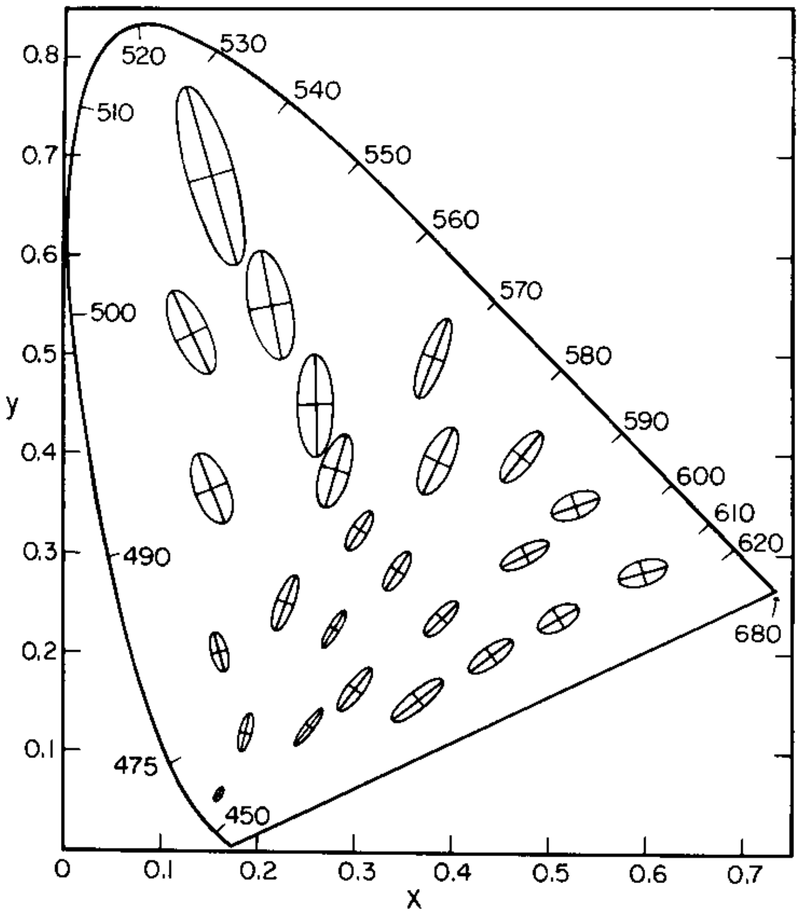

- Refer to the region which contains all colors which are indistinguishable
- Scaled by a factor of 10 and shown on CIE xy color space
- If you are shown two colors, one at the center of the ellipse and the other inside it, you cannot tell them apart
- Only a few ellipses are shown, but one can be defined for every point

CIE u’v’ Space

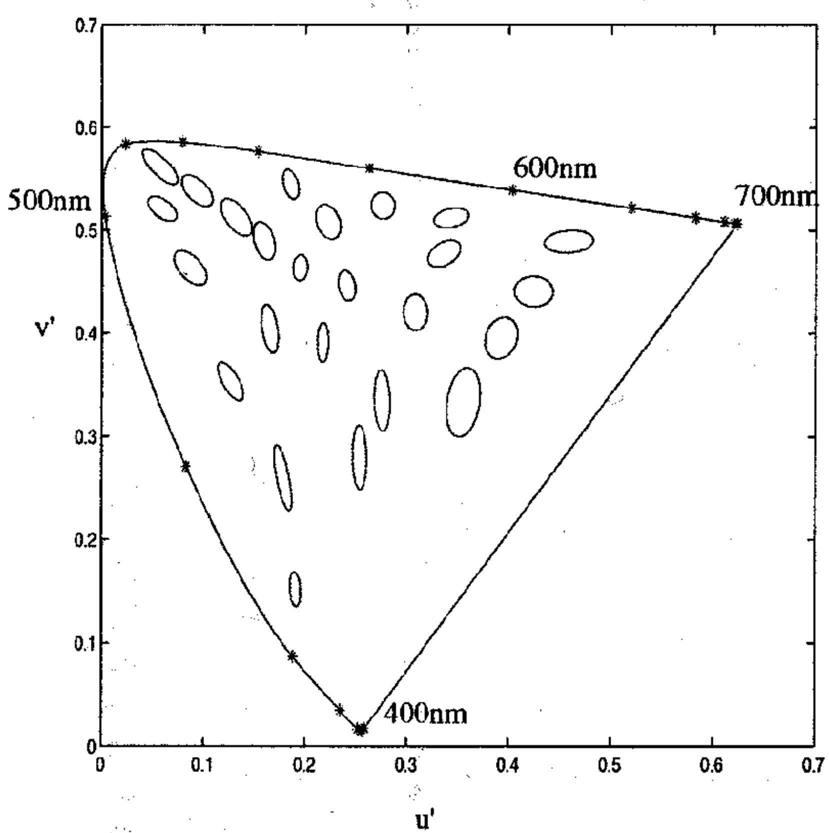

- CIE u’v’ is a non-linear color space where color differences are more uniform
- Note that now ellipses look more **like** circles
- The third coordinate is the original Z from XYZ

$$\begin{bmatrix} u' \\ v' \end{bmatrix} = \frac{1}{X + 15Y + 3Z} \begin{bmatrix} 4X \\ 9Y \end{bmatrix}$$

Next Time

- Filter