CS 376A
Digital Image Processing

01 / 25 / 2023

Instructor: Michael Eckmann

Today's Topics

- Introduction of myself
- Slight Review the syllabus
- Image processing vs. computer vision vs. computer graphics
- Examples of Image processing results
- Uses of image processing
- Digital Images defined
  - Color (RGB)
  - Grayscale
  - True black and white

Who is your instructor?

- I'm Mike Eckmann and I have taught at Skidmore since 2004. Before that I was at Lehigh University in PA.
- I studied Mathematics and Computer Engineering and Computer Science all at Lehigh University.
- I have research interests in computer vision (to be defined).
- I was employed as a programmer (systems analyst) for eight years.

Syllabus

- Office hours
- Readings
- Assignments
  - Hw's & programs
- Collaboration policy
- Grading schema
- Attendance
- Exams

meckmann.domainsskidmore.edu/2023Spring/cs376a

The most up-to-date syllabus will be found on the course web page.

Terms

- Image --- I will likely use the terms *image* and *digital image* interchangeably but digital images are not the only kinds of images that exist.
- Let's consider what we'll be talking about this semester in terms of digital image processing vs. some other image related computer science areas of study like computer vision and computer graphics.

Terms

- Roughly ...
- Digital Image Processing --- input and output are both images
- Computer Graphics --- input is a mathematical model description of a scene and “camera” and output is an image
- Computer Vision --- input is an image and output is some understanding of the image (possibly the inverse of the Computer Graphics)

Terms

- Some examples of Digital Image Processing problems:
  - enhance the contrast of an image
  - lighten, darken, or fix flat images
  - find the edges and create an edge image
  - blur or sharpen an image
  - creating an intensity (grayscale) image from a color image
  - segment an image into distinct parts based on how alike the pixels are within a segment vs. how different they are from other segments

Terms

- Some examples of Computer Vision problems:
  - face recognition from an image
  - track objects in video
  - classifying medical images into healthy vs. diseased

  - Note: these problems will often have image processing techniques as steps in their solutions

Terms

- An example of a Computer Graphics problem:
  - design a 3d model of a character and surrounding environment including lights and surface properties and generate images of the character in the environment from different viewpoints

Syllabus – course topics

- The course is on image processing techniques, specifically algorithms that take an image as input and produce another image as output after processing it in some way.
- Topics to be covered
  - computing intensity image of a color image
  - techniques to do edge detection
  - noise reduction techniques
  - blurring
  - learn how to compute histograms of color values
  - contrast enhancement via histogram equalization
  - segmentation

Syllabus – course topics

- Topics to be covered continued:
  - Morphological operators
  - Clustering via K-means
  - Discrete Fourier Transforms
  - JPEG compression algorithm
  - Geometric Transformations of the image plane
  - and more...

Syllabus – course goals

- Students will learn the details of the various digital image processing algorithms and will write programs in Java to implement many of these techniques.
- Students will learn the details of how they work instead of just using a program like Photoshop to perform the techniques.
- They will be encouraged to use their own digital images as inputs and to experiment with.

Examples of Digital Image Processing

- Let's look at some images and the results of some image processing techniques on those images.

Uses of Digital Image Processing

- Who needs digital image processing?
- Any thoughts?

Uses of Digital Image Processing

- Who needs digital image processing?
- Astronomical imaging processing
- Biological imaging
  - e.g.
  - http://www.ncbi.nlm.nih.gov/pubmed/23560739
- Chemical imaging
  - e.g.
  - http://ieeexplore.ieee.org/xpl/articleDetails.jsp?reload=true&arnumber=4351423
- Geological / Earth Science imaging
  - e.g. https://dr.library.brocku.ca/handle/10464/1891

What is a digital image?

- 2 dimensional grid of pixels
- Each pixel contains numeric data that describe what color the pixel is (or what brightness the pixels is for grayscale images.)
- Depending on whether it is a color or grayscale image, each pixel contains either 3 numbers (Red, Green and Blue) or 1 number (brightness).

What is a digital image?

- The smallest unit in an image is a pixel. Each pixel contains one color.

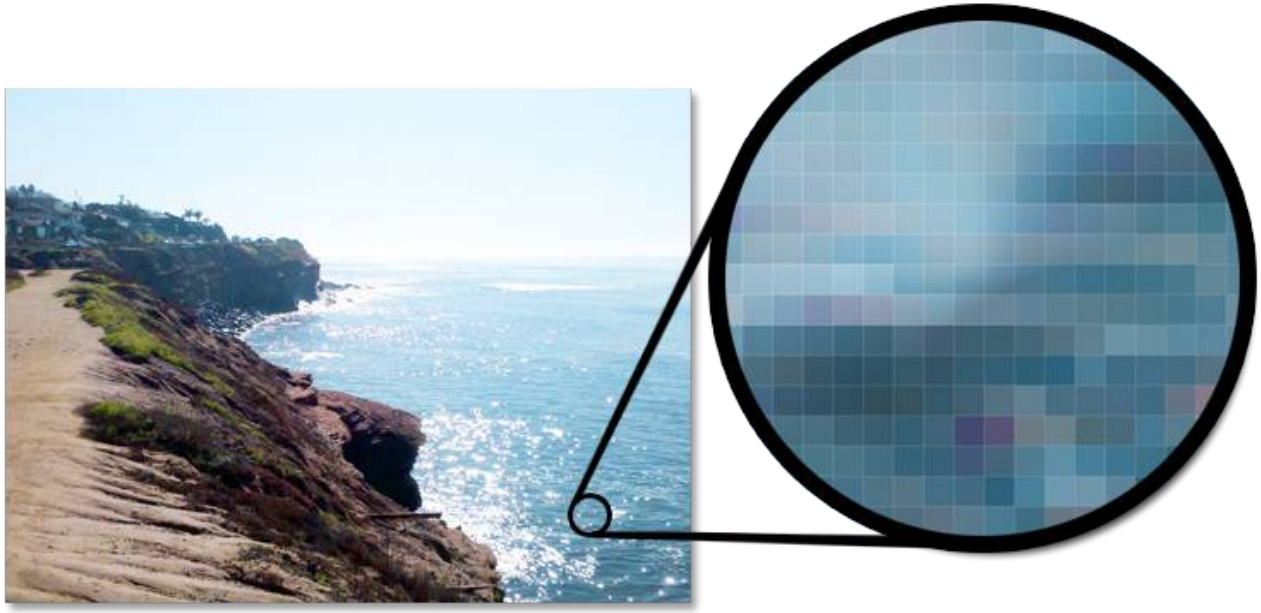

What is a digital image?

- Both images below are in public domain

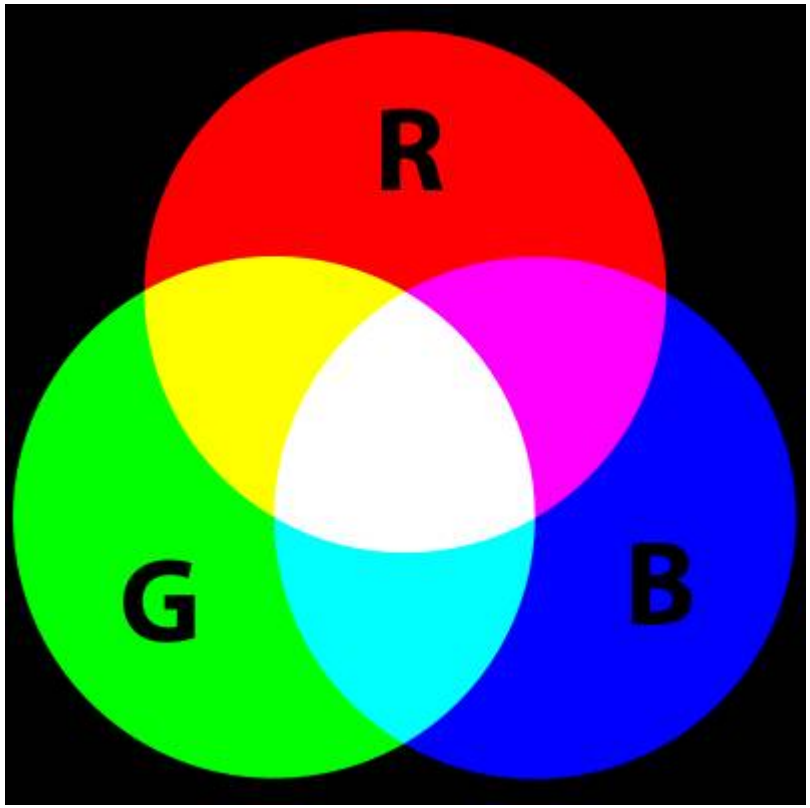

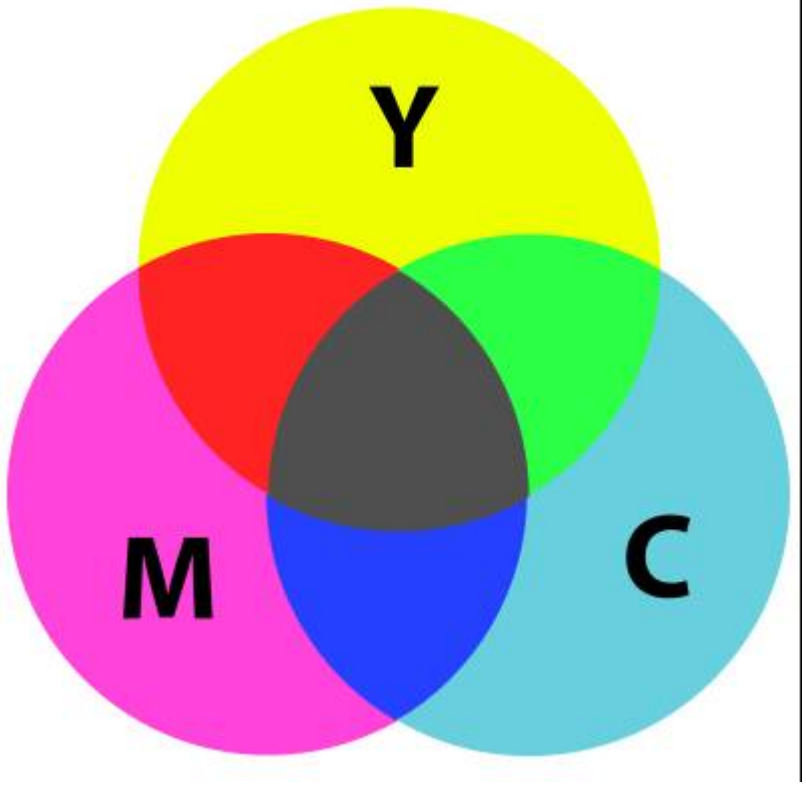

Color / light

(next 3 diagrams from Hearn & Baker, Computer Graphics with OpenGL).

- Physically, color is electromagnetic radiation in the visible light frequency range.

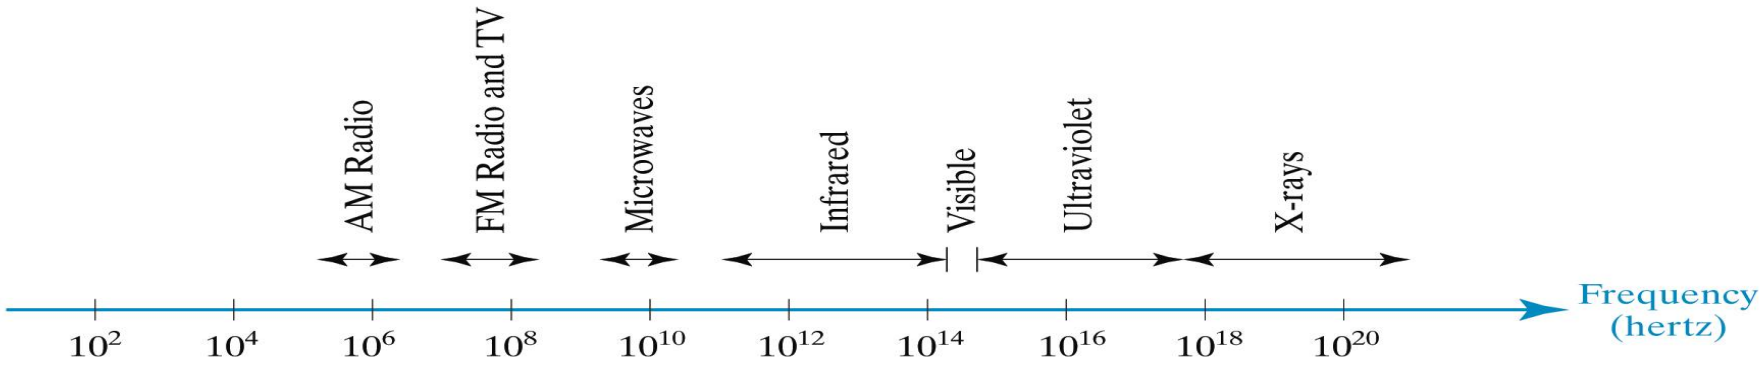

Figure 12-1
Electromagnetic spectrum.

Color / light

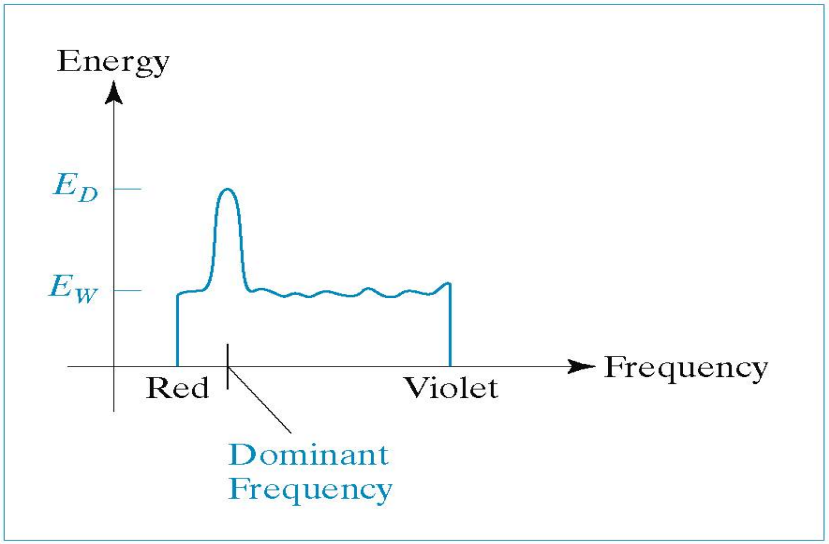

Figure 12-4

Energy distribution for a light source with a dominant frequency near the red end of the frequency range.

Color / light

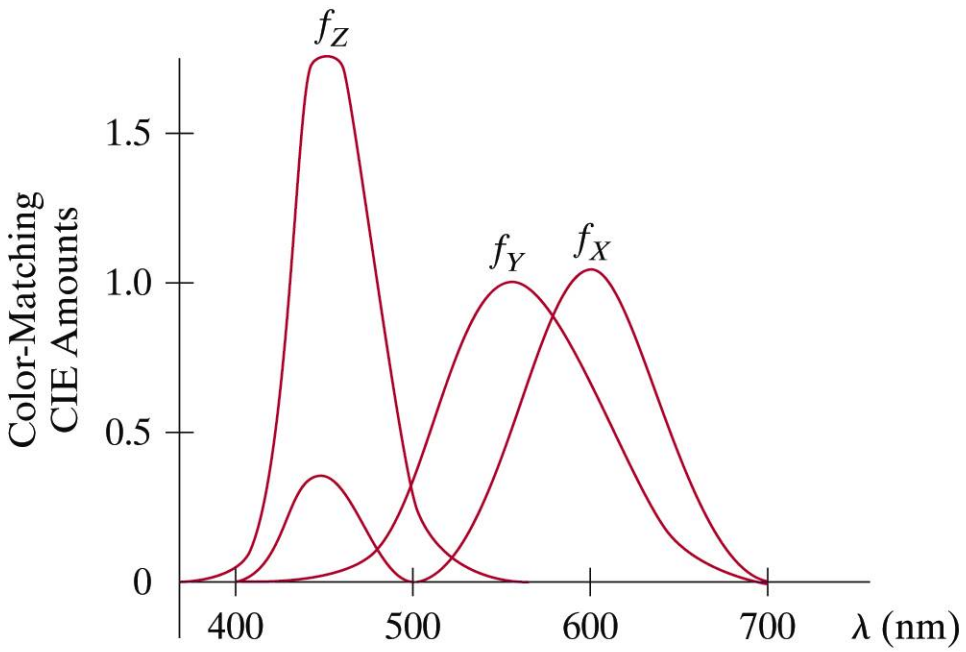

Figure 12-6

The three color-matching functions for the CIE primaries.

RGB

https://academo.org/demos/wavelength-to-colour-relationship/

- The link above has a way for us to see the RGB values for each pure wavelength color in the visible light spectrum.
- Note that these are only the pure single wavelength colors. Notice that there were no white, black or grays in that slider – recall those are when there are equal amounts of the red wavelength, blue wavelength and green wavelenght lights. Plenty of other examples too.

RGB

- Let's get a feel for what values of R, G and B channels describe what colors.
- Use the applet on this site:

https://yuilibrary.com/yui/docs/color/rgb-slider.html

Let's observe

- what colors we get when all 3 are equal (R=G=B)
- what color we get when two of the 3 are high, the other is low

RGB

- Often the size of each R, G and B value is 1 byte which has a range of 0-255.
- How many different colors can be made if each channel has 1 byte of space?

RGB

- Often the size of each R, G and B value is 1 byte which has a range of 0-255.
- How many different colors can be made if each channel has 1 byte of space?

- 256 * 256 * 256 = 16,777,216
- One of 16,777,216 different colors can be specified when 1 byte per color channel is used to store the color.

Byte

- A byte is 8 bits
- A bit holds one value (a 0 or a 1)
- Since there are 2 different values for a bit and there are 8 bits in a byte, there are 2^8=256 different values for a byte.
- Examples on the board.

Grayscale

- Grayscale images hold 1 number per pixel.
- Each pixel is usually 1 byte.
- So, one of 256 different brightness values can be stored per pixel.
- 0 = black
- 255 = white
- All numbers in between are varying levels of gray

True black and white (monochrome)

- In layman's terms a grayscale image is often described as “black and white”. This is not accurate.
- How much space would one need per pixel to store a true black and white image?

True black and white (monochrome)

- In layman's terms a grayscale image is often described as “black and white”. This is not accurate.
- How much space would one need per pixel to store a true black and white image?
- 1 bit (0=black, 1=white)
- Can anyone here come up with a technique to convert a grayscale image to true black and white?

Intensity from RGB

- Recall that white in RGB = 255, 255, 255
- Black is 0,0,0
- The higher the number for a color channel, the more it represents the amount of that color light.
- Any thoughts on how we can compute the intensity (a single number) from an RGB value?

Intensity from RGB

- Computing Intensity from RGB

- 0.333*R + 0.333*G + 0.333*B is a reasonable idea.
- 0.299 * R + 0.587 * G + 0.114 * B is a better choice to give more accurate human perceived intensity.
- Green light has more of an influence on the intensity that human eye perceives than Red does and Red has more influence than Blue.

Java classes for the course

- Let's review code I will provide that reads and writes jpeg images and allows access to the 2d array of pixels. Each pixel of which contains a red, green and blue value in the 0..255 range inclusive.