CS 331
Computer Vision

09 / 10 / 2026

Instructor: Michael Eckmann

Today's Topics

- Introduction to the course
  - topics etc.
- What is Computer Vision?
- Main topics today:
  - Digital images
  - Color models, RGB
- Etc.

Who is your instructor?

- I'm Mike Eckmann and I have been teaching at Skidmore since 2004. Before that I was at Lehigh University in PA.
- I studied Mathematics, Computer Engineering and Computer Science all at Lehigh University.
- I was employed as a programmer (systems analyst) for eight years.
- I have research interests in computer vision.
- I have taught this updated version of my Computer Vision course 4 times in the recent past.

Course Information

- For programming assignments in this course, we will be using python with many packages including OpenCV and tensorflow.
- Syllabus is posted to our class web page, where you can also find course notes etc.
- I will email a link to our course webpage later today.
- In class today you will fill out a questionnaire.

Course Information (cont.)

- Expect to have 4 or more programming assignments
  - one will make sure you understand how to use opencv and implement some image processing techniques
  - several will utilize machine learning

Course Information (cont.)

- In addition to the programming assignments, we may have a few homeworks dealing with some concepts that aren't covered in the programming assignments
- The second to last assignment will be an individual presentation of a paper

Reading

In Szeliski (https://szeliski.org/Book/)

section 1.1 What is computer vision?

pages 18-22 regarding the state of the computer vision field since the year 2000

section 1.5 A note on notation

section 1.6 Additional reading

Digital cameras, sampling, color models

section 2.3 The digital camera

2.3.1 Sampling and aliasing

2.3.2 Color

2.3.3 Compression --- skim it,

the details of compression can't be described in 2 pages, but it will give you an idea

Computer Vision

- The goal of Computer Vision, according to one text is to "make useful decisions about real physical objects and scenes from sensed images."
- That is, "understanding images."

Computer Vision

- Basic (oversimplified) observations on the relationship between Computer Vision, Computer Graphics and Image Processing.
  - Computer Graphics
    - Given models, generate images
  - Computer Vision
    - Given images, generate models
  - Image Processing
    - Given images, generate (other) images

Computer Vision

- Applications of computer vision are numerous – see chapter 1 in Szeliski
- e.g. OCR, defect detection, object recognition, medical imaging applications, self-driving cars, surveillance, face detection and recognition, etc.
- Are you aware of others?

Computer Vision

- Many computer vision tasks require image processing techniques to be done first.
- So, we will do some image processing early in the course.

Computer Vision Topics

- Image formation and representation
  - how do we acquire digital images and in what format are they stored
- Key Image processing techniques of correlation/convolution, histograms, etc.
- Binary Image Analysis
  - what information can we get out of a black and white image (as opposed to a color image or a greyscale image)?
  - what are typical operations on binary images and for what are they used?

Computer Vision Topics

- Image Filtering using convolution
  - how do we reduce different kinds of noise
  - detecting edges etc.
- Color models

Computer Vision Topics

- Machine Learning networks for computer vision
  - Details about convolutional neural networks and others for object classification
  - We will spend a large chunk of the middle of the course on the lead up to this
- Feature detection and matching

Computer Vision Topics

- Note: when you start reading the Szeliski book you will find out that there is a lot more to computer vision than what we can cover in this course. I decided on topics that we could cover well that give you experience with modern computer vision, while also teaching the necessary machine learning concepts.
- So, one big takeaway from this course will be the knowledge gained about neural networks and how they work in general.

Questionnaire

- Please fill out the questionnaire. We will wait until everyone is done before we continue.

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 pixel 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. This one color is represented as a combination of some amount of Red, Green and Blue light.

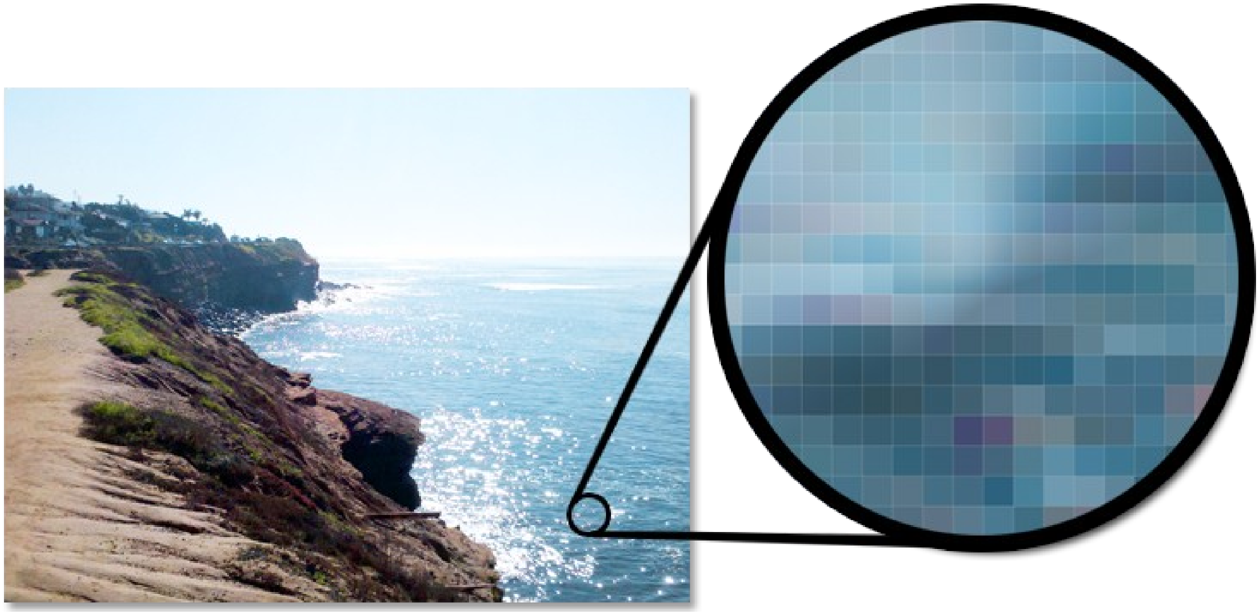

What is a digital image

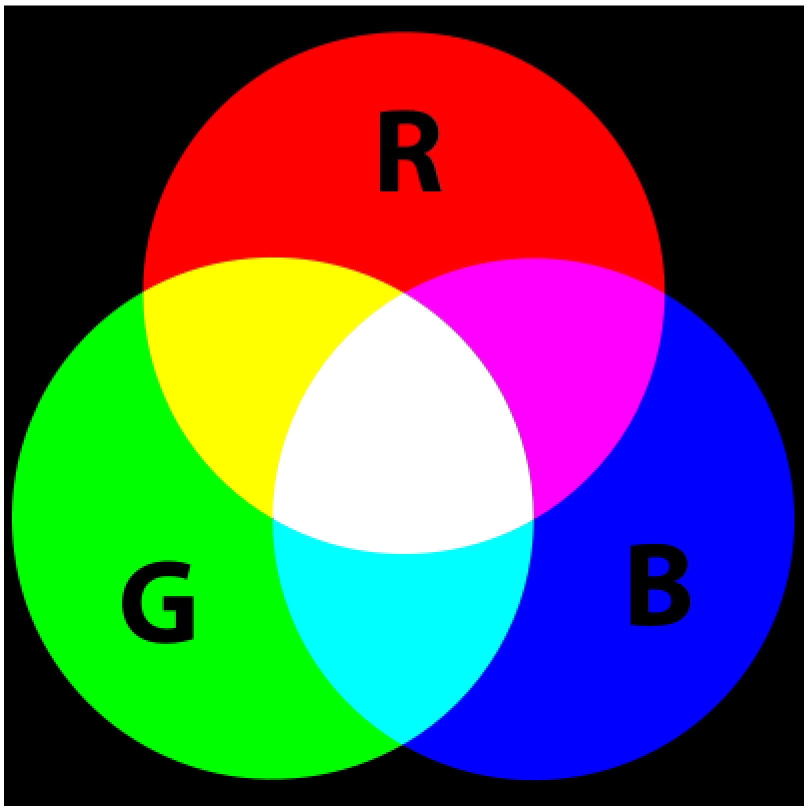

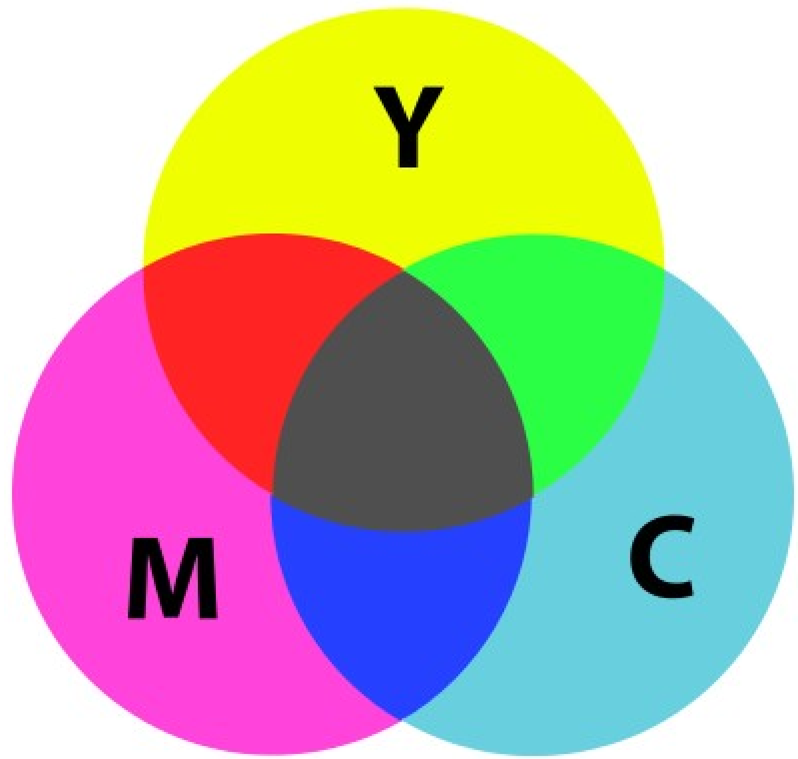

RGB Color Model

- **Color = r R + g G + b B**
- r, g and b range from 0.0 to 1.0
- notice:

  white's (r,g,b) = (1,1,1)

  black's (r,g,b) = (0,0,0)

  complementary colors are

  those that add up to (1,1,1)

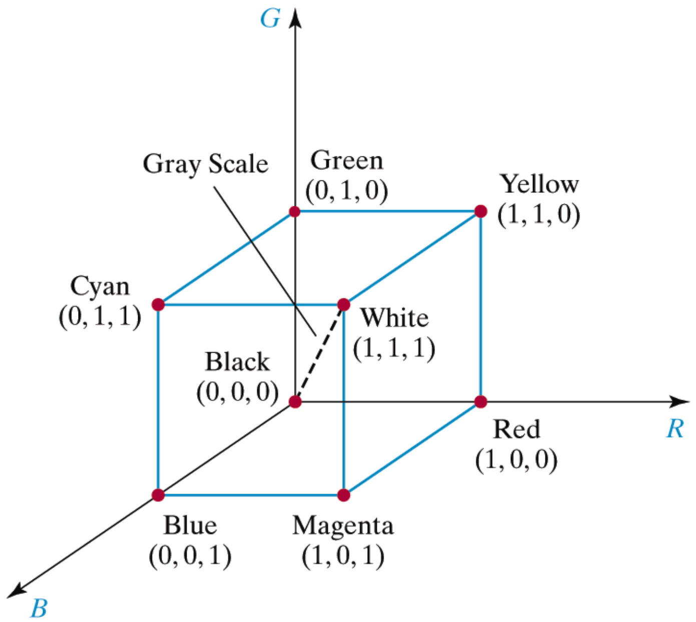

Figure 12-11

The RGB color model. Any color within the unit cube can be described as an additive combination of the three primary colors.

- image from “Computer Graphics” by Hearn and Baker

What is a digital image

- Let's examine a digital image using gimp (Gnu Image Manipulation Program) that has similar functionality to Photoshop.

- I'll zoom in and point out RGB values of pixels
- I'll convert to grayscale and show the brightness values of pixels

What is a digital image

What is a digital image

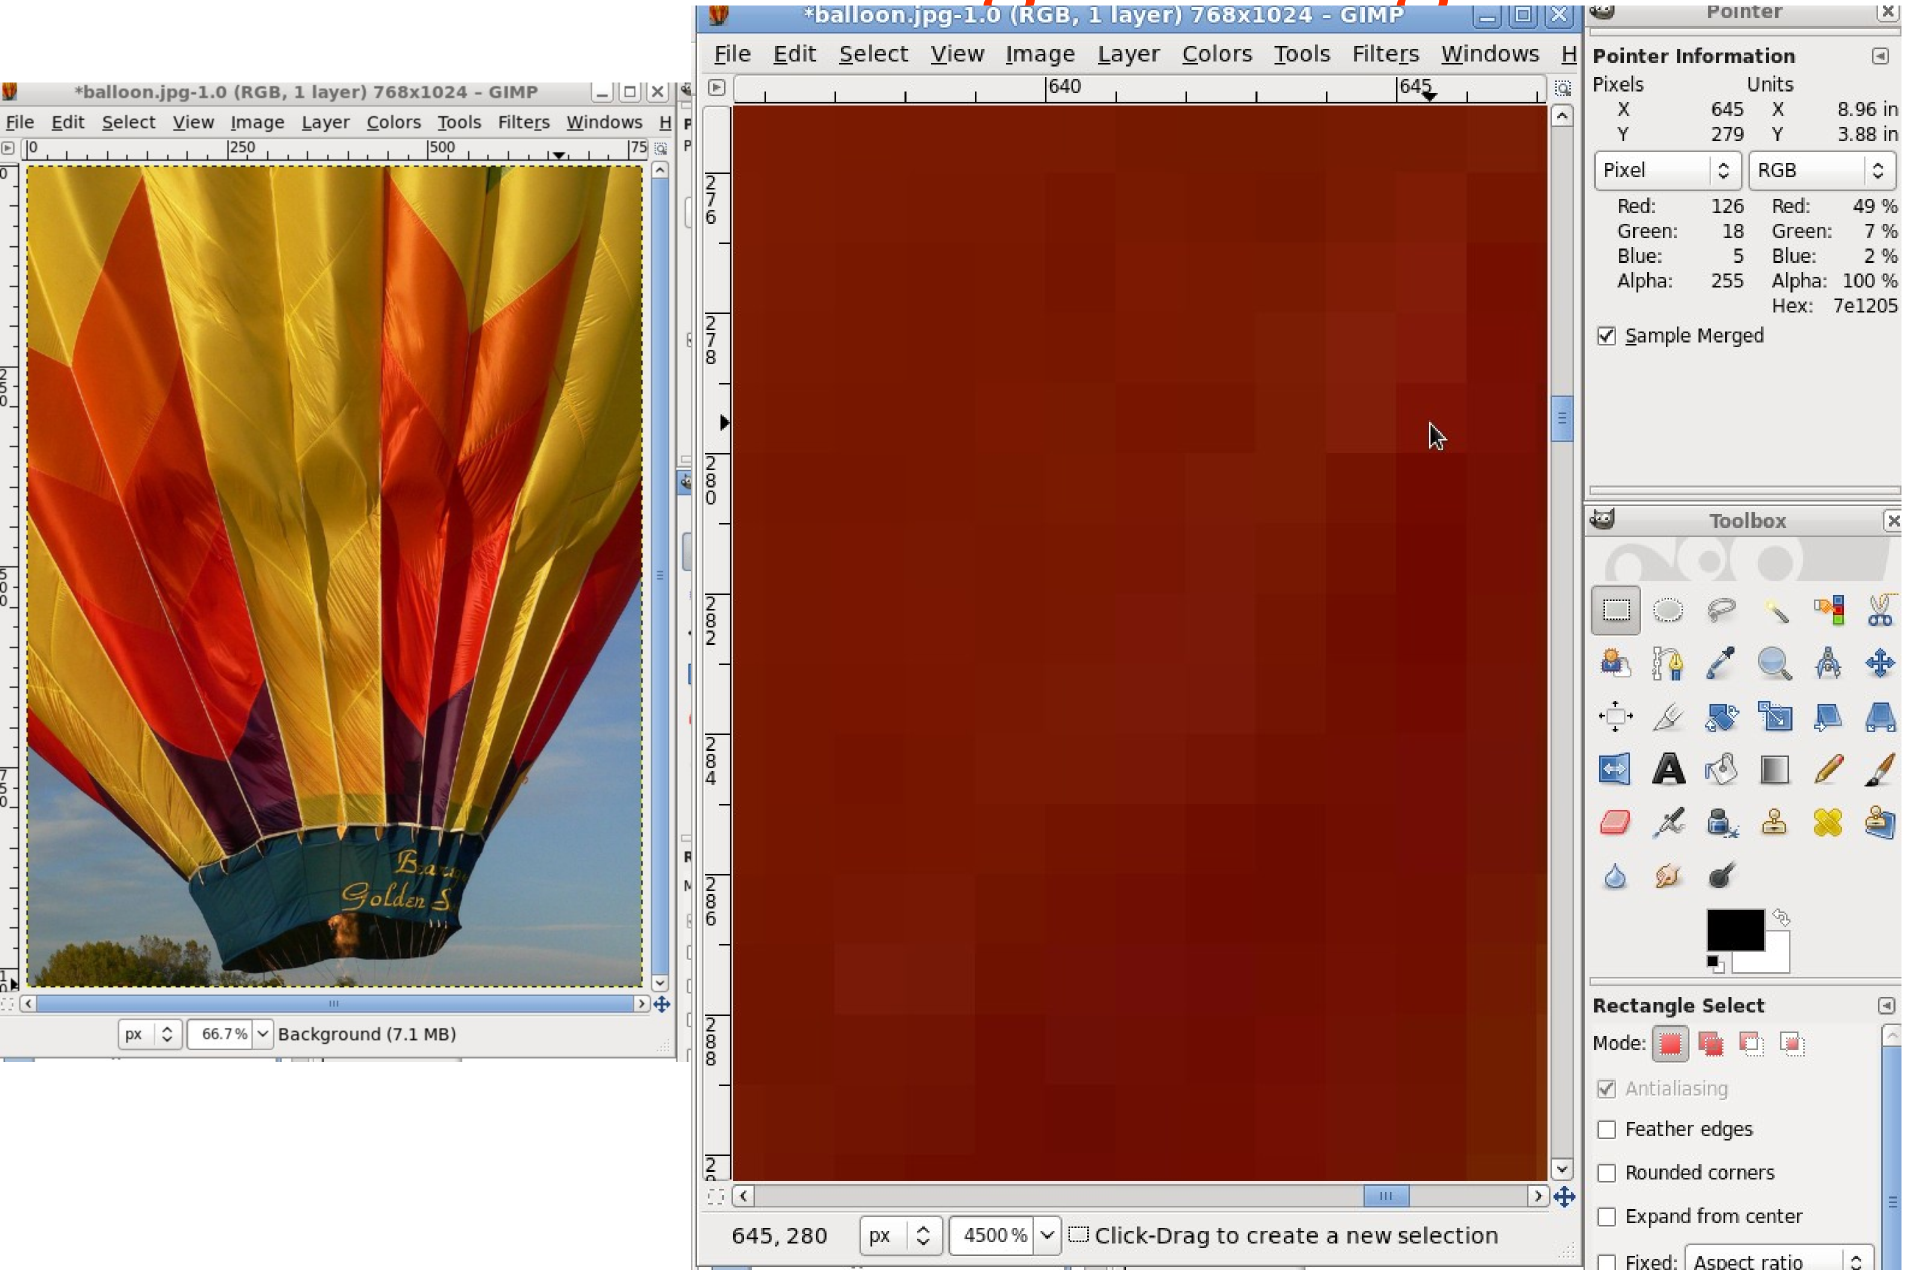

What is a digital image

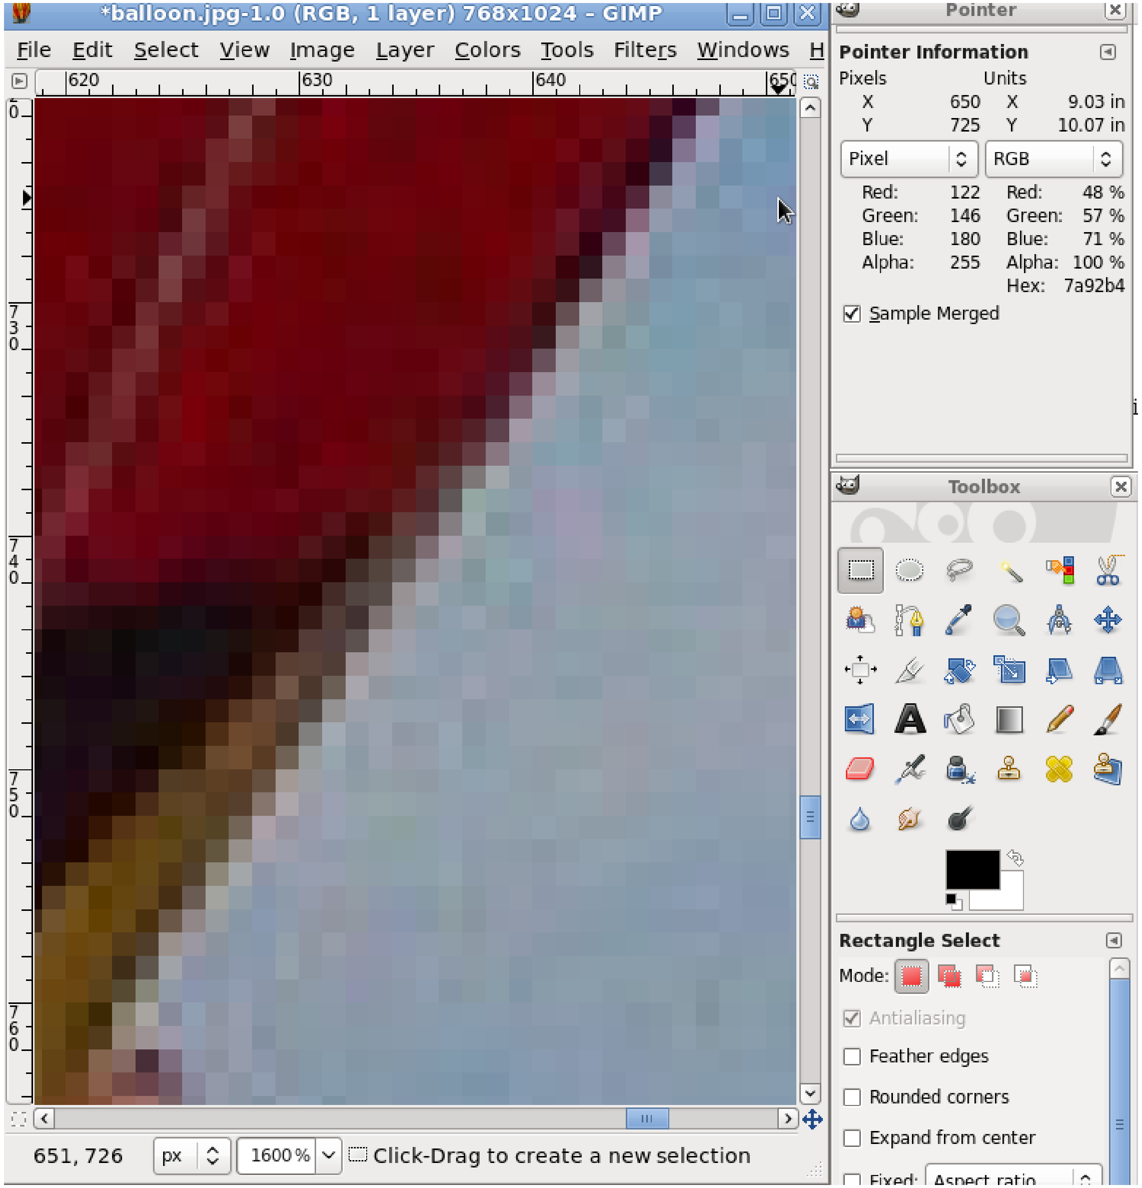

What is a digital image

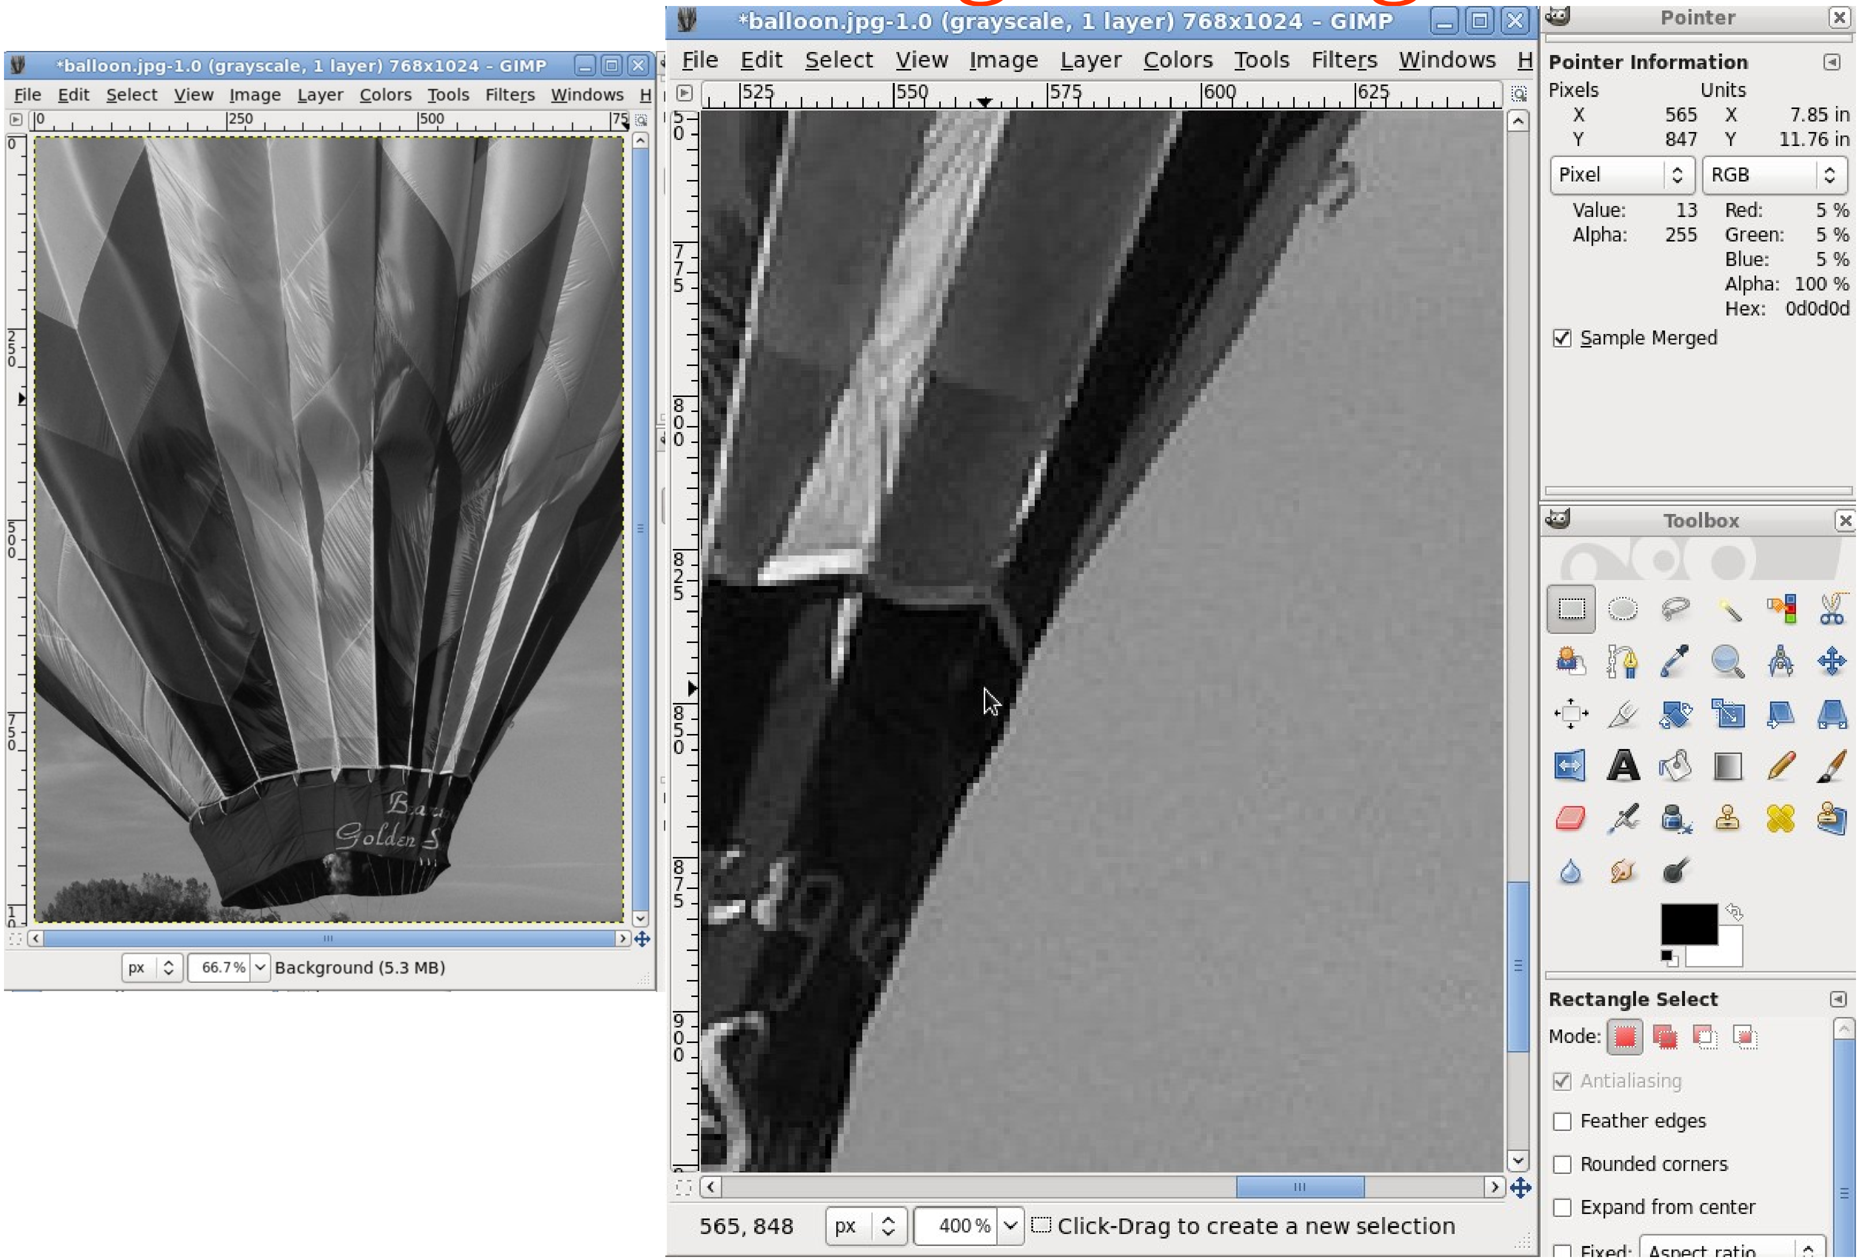

What is a digital image

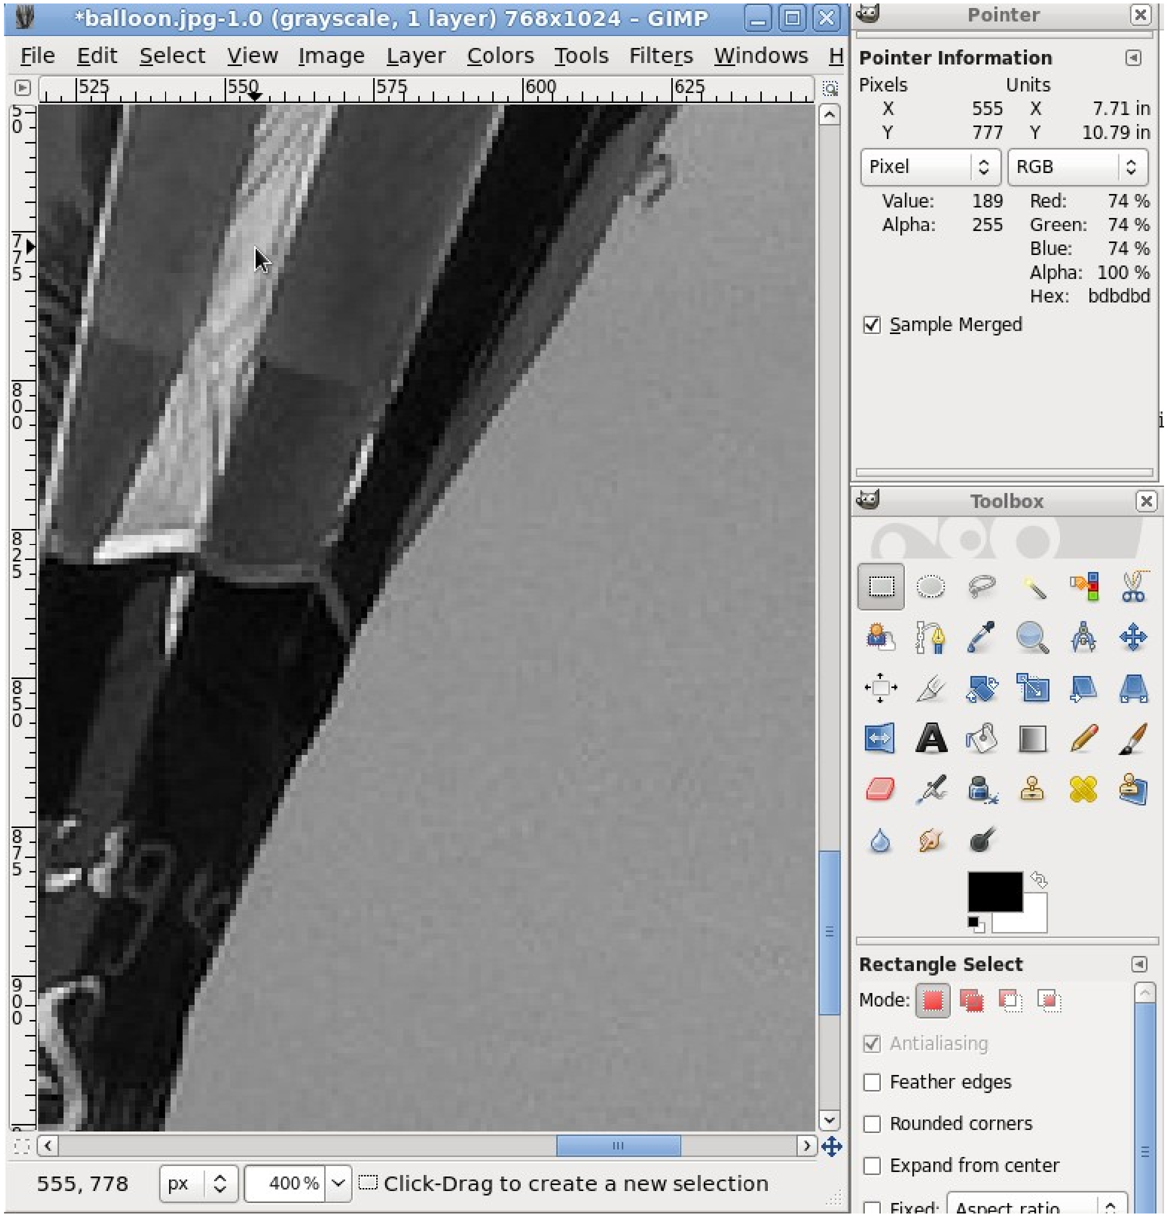

RGB

- You should get a feel for what values of R, G and B channels describe what colors.
- Let's use this site:

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

Observe what colors you get when all 3 are equal (R=G=B).

Let's keep two channels constant and slide the other channel among all values.

RGB

- Often the size of each of the R, G and B color channels 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 of the R, G and B color channels 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.

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)
- How might one convert a grayscale image to true black and white?