CS 200

Lecture 03
Pixel Graphics

Miscellaneous Notes

Abbreviations

dpi	=	dots per inch
ppi	=	pixels per inch
FAQ	=	frequently asked questions
lpi	=	lines per inch
RWS&HT	=	Real World Scanning and Halftones
NDWB	=	The Non-Designer's Web Book

Optional background reading

"Beyond the Mac is not a typewriter," by Robin Williams
(http://www.ratz.com/robin/realbio.html)

Administrivia

Remember . . .

UNDERSTANDING what you're doing, and why, is more important than how to do it

So as you listen today

- focus on WHAT we're doing, not how

This week's reading

- both supplements and *complements* this week's lecture
- from Learning Web Design

Topic
RGB colour (306-7)
Indexed Colour (648)
Bit Depth (696)
Monitor Resolution (658-9)
Resolution of Images (657-8)
File Formats (gif, jpeg, png) (645-646,653)
Anti-aliasing (677)
File size of images (657-660)
Saving images in Gimp (671,673-674)

- Please ask questions!

Assumptions

- You have used a camera.
- You may have done some photo editing.
- Terminology in this lecture may be brand new.

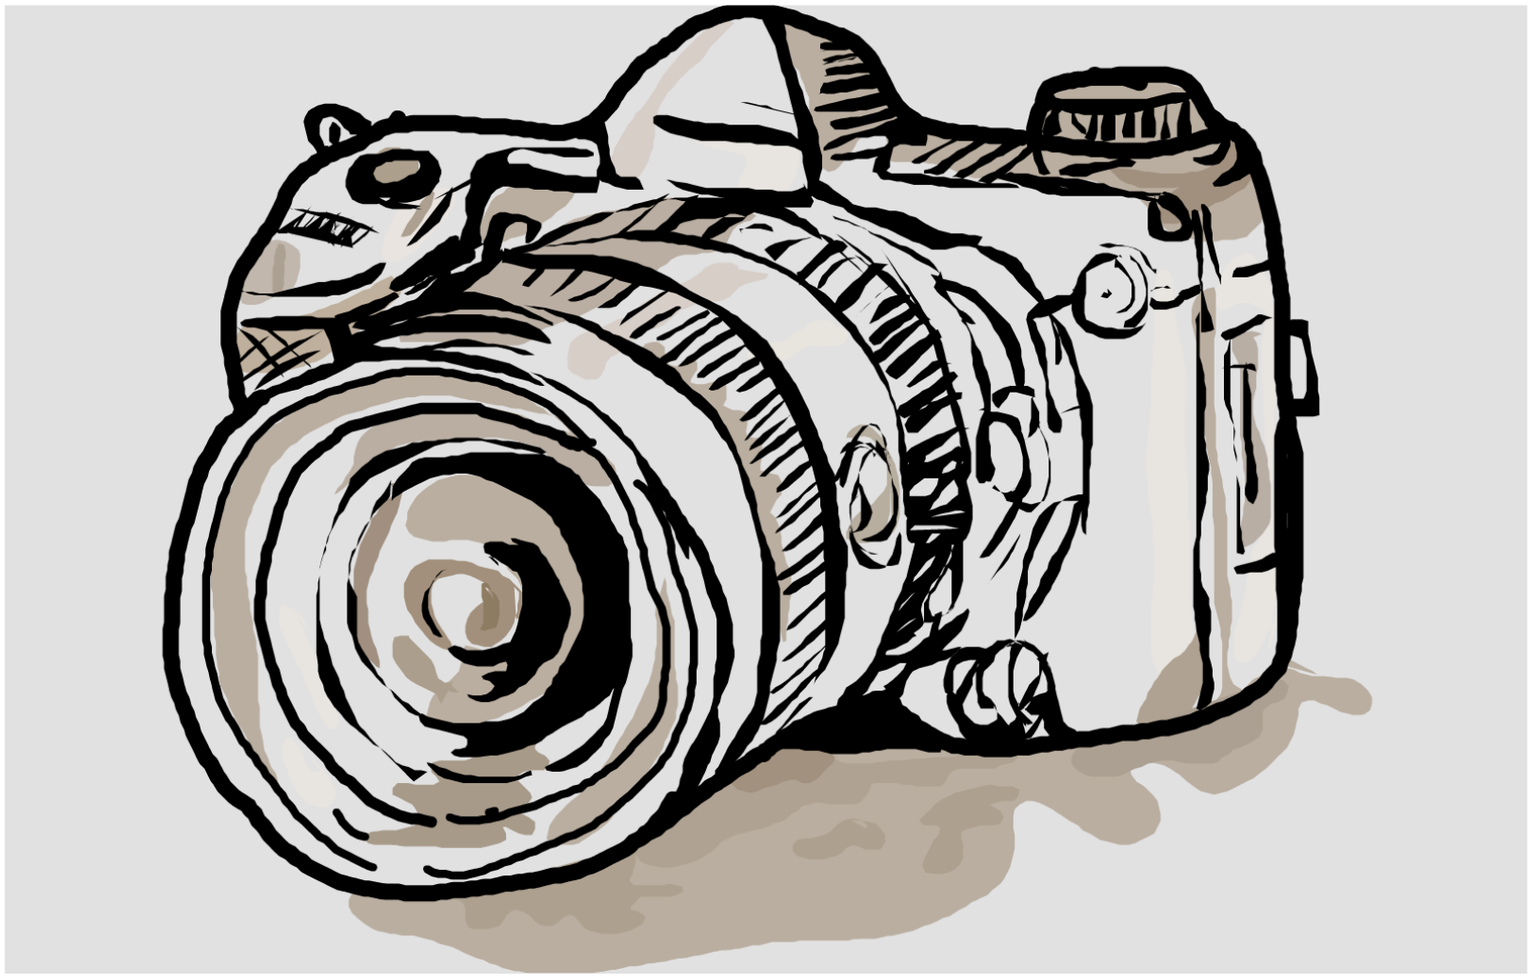

Things to Think About

- What are the data objects in a Pixel Graphics program?
- Is there more than one way to do any given task?
- What are the deficiencies of the interface?
- What are the efficiencies of the interface?

What *is* "Pixel Graphics?"

Everything you see on the screen is composed of "pixels"

- imagine the screen as a piece of graph paper
- draw an object by colouring in squares
  - "painting" with a "brush" (ie the "tracker")
- the squares are called "pixels"
- "pixel" is short for PICture ELement (from "pix" for picture)

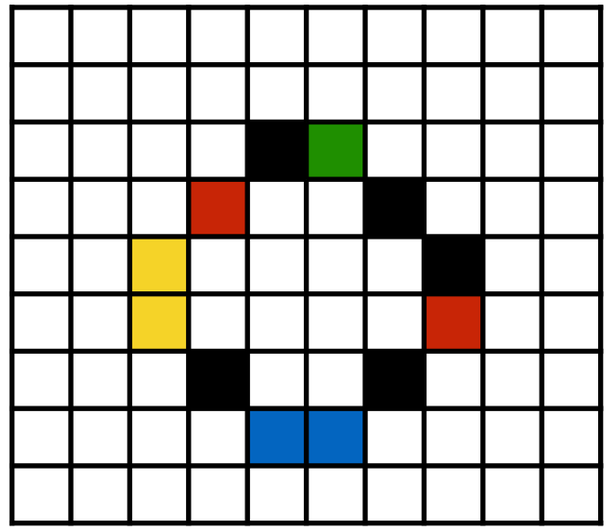

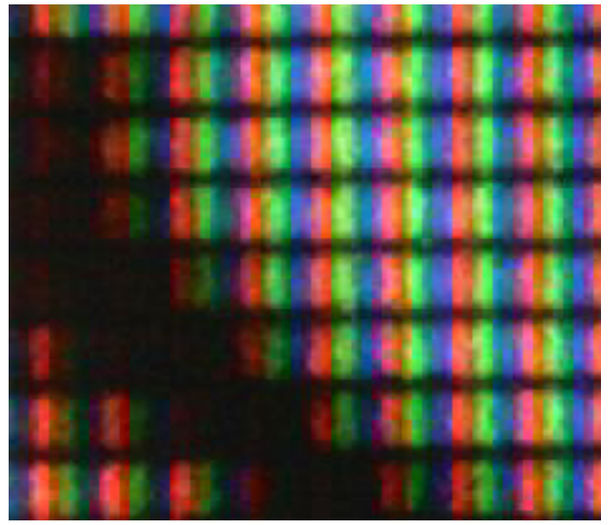

A piece of my Dell 2001FP LCD at 50x.

Each pixel (usually) has

- a red intensity (0...255)
- a green intensity (0...255)
- a blue intensity (0...255)

— vary these to get a variety of colours
Thus typically $0 \le R, G, B \le 255$

- 255 is the largest integer that fits in a byte (character) of memory

EVERYTHING displayed is represented by a pattern of pixels

For more information on how LCDs — and CRTs — actually work, see Wikipedia

Why does pixel graphics work?

RGB-based colour is "additive colour"

As distinct from the "subtractive (CMY-based) colour"
you may have seen in an art class

Additive colour is based on the red-, green-, & blue-sensitive "cones" in your eye

red, green and blue are said to be the "primary colours"

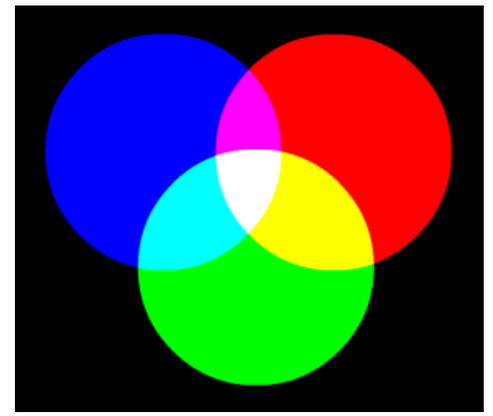

Additive Colour — eg an LCD

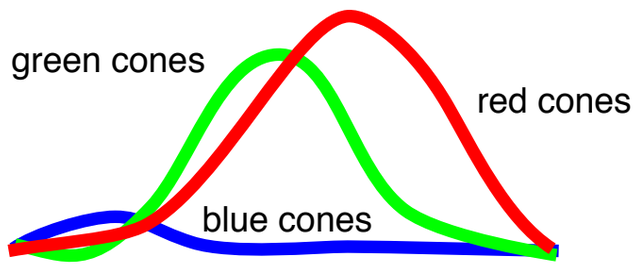

blue light Cone sensitivities across the visible spectrum. **red light**

"Subtractive Colour" — eg printing
(Adapted from www.edumedia.com.)

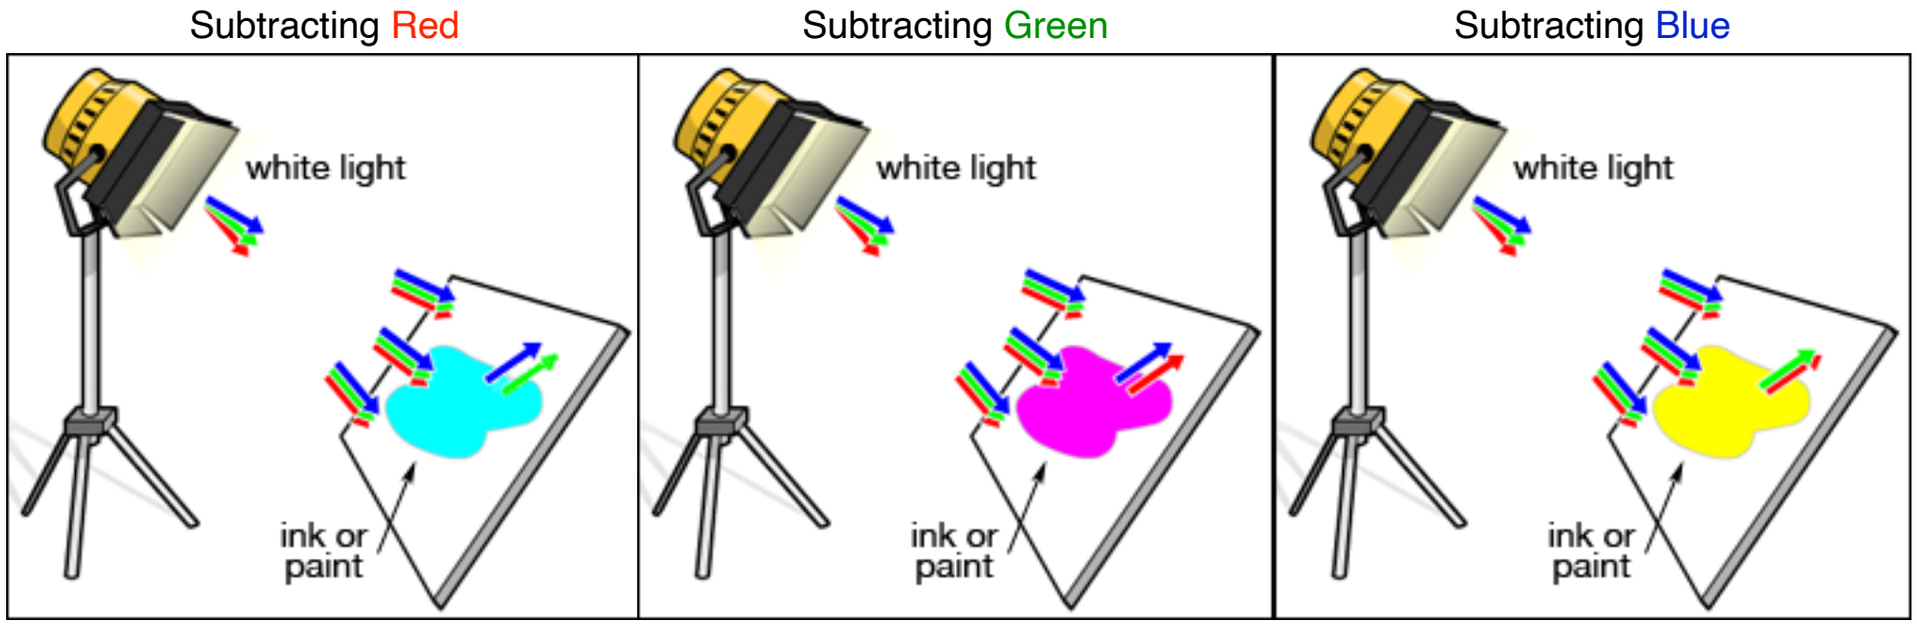

Some facts about additive colour

Some colour facts & terminology

- red + green = yellow ie R + G = Y
- green + blue = cyan ie G + B = C
- red + blue = magenta ie R + B = M
- for us, yellow, cyan and magenta are said to be "secondary colours"
- red + green + blue = white ie R + G + B = W

The "Colour Wheel" — a model for mixing colour

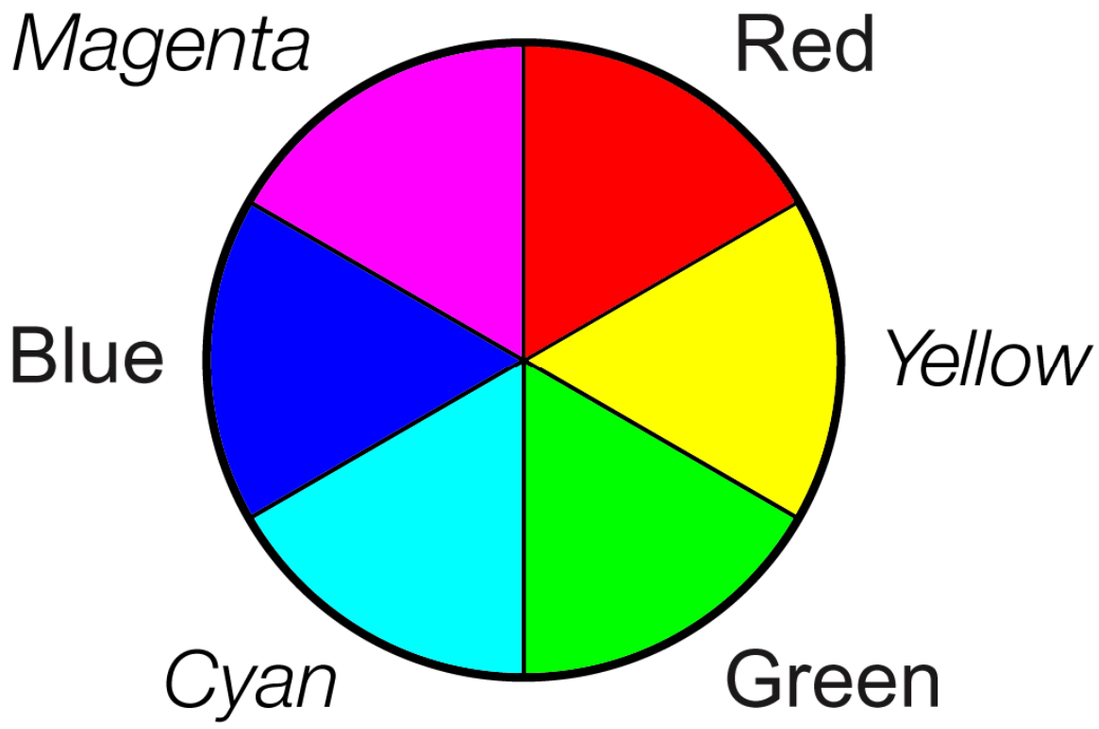

More facts about additive colour

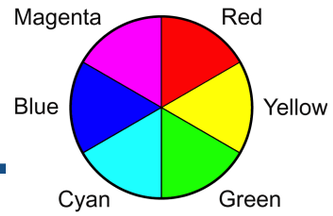

Therefore: white – blue = red + green = yellow (etc)

“nothing” is black, represented by K (because B is already taken for blue...)

“equal” amounts of red, green and blue combine to form grey

“hue” is what we usually think of as “colour”

as in reddish, greenish, etc

“saturation” refers to how “pure” or “vivid” the colour is

ie how much a colour departs from gray

“brightness,” “lightness,” & “luminance” refer to overall intensity

you can specify colour in terms of hue, saturation, & brightness

there’s an arithmetic transformation between the two

similarly for cyan, magenta and yellow, the (subtractive printing primaries)
or CMYK (printing with the four inks cyan, magenta, yellow and black gets better results)

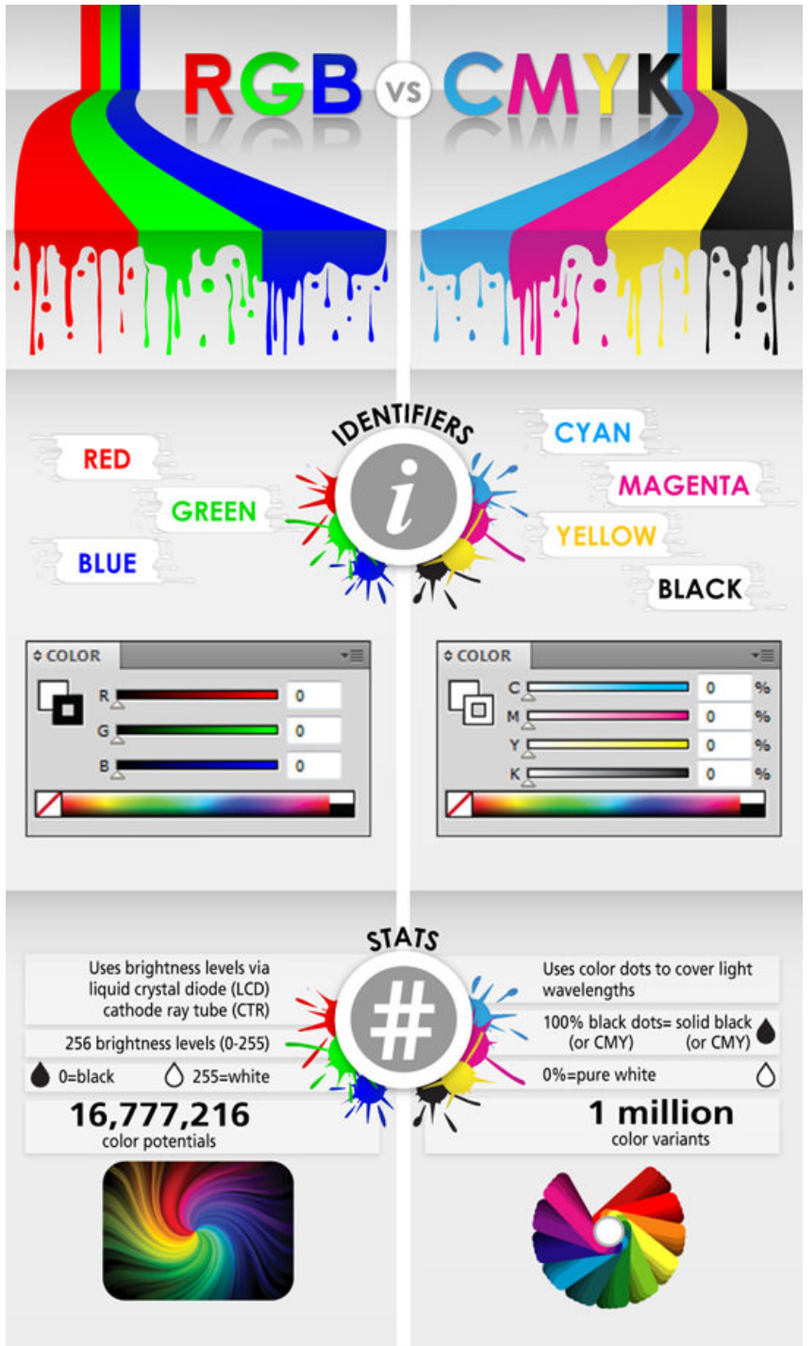
RGB VS CMYK
IDENTIFIERS
RED
GREEN
BLUE
CYAN
MAGENTA
YELLOW
BLACK
COLOR
R 0
G 0
B 0
COLOR
C 0 %
M 0 %
Y 0 %
K 0 %
STATS
Uses brightness levels via liquid crystal diode (LCD) cathode ray tube (CTR)
256 brightness levels (0-255)
0=black
255=white
16,777,216
color potentials
Uses color dots to cover light wavelengths
100% black dots= solid black (or CMY)
(or CMY)
0%=pure white
1 million
color variants

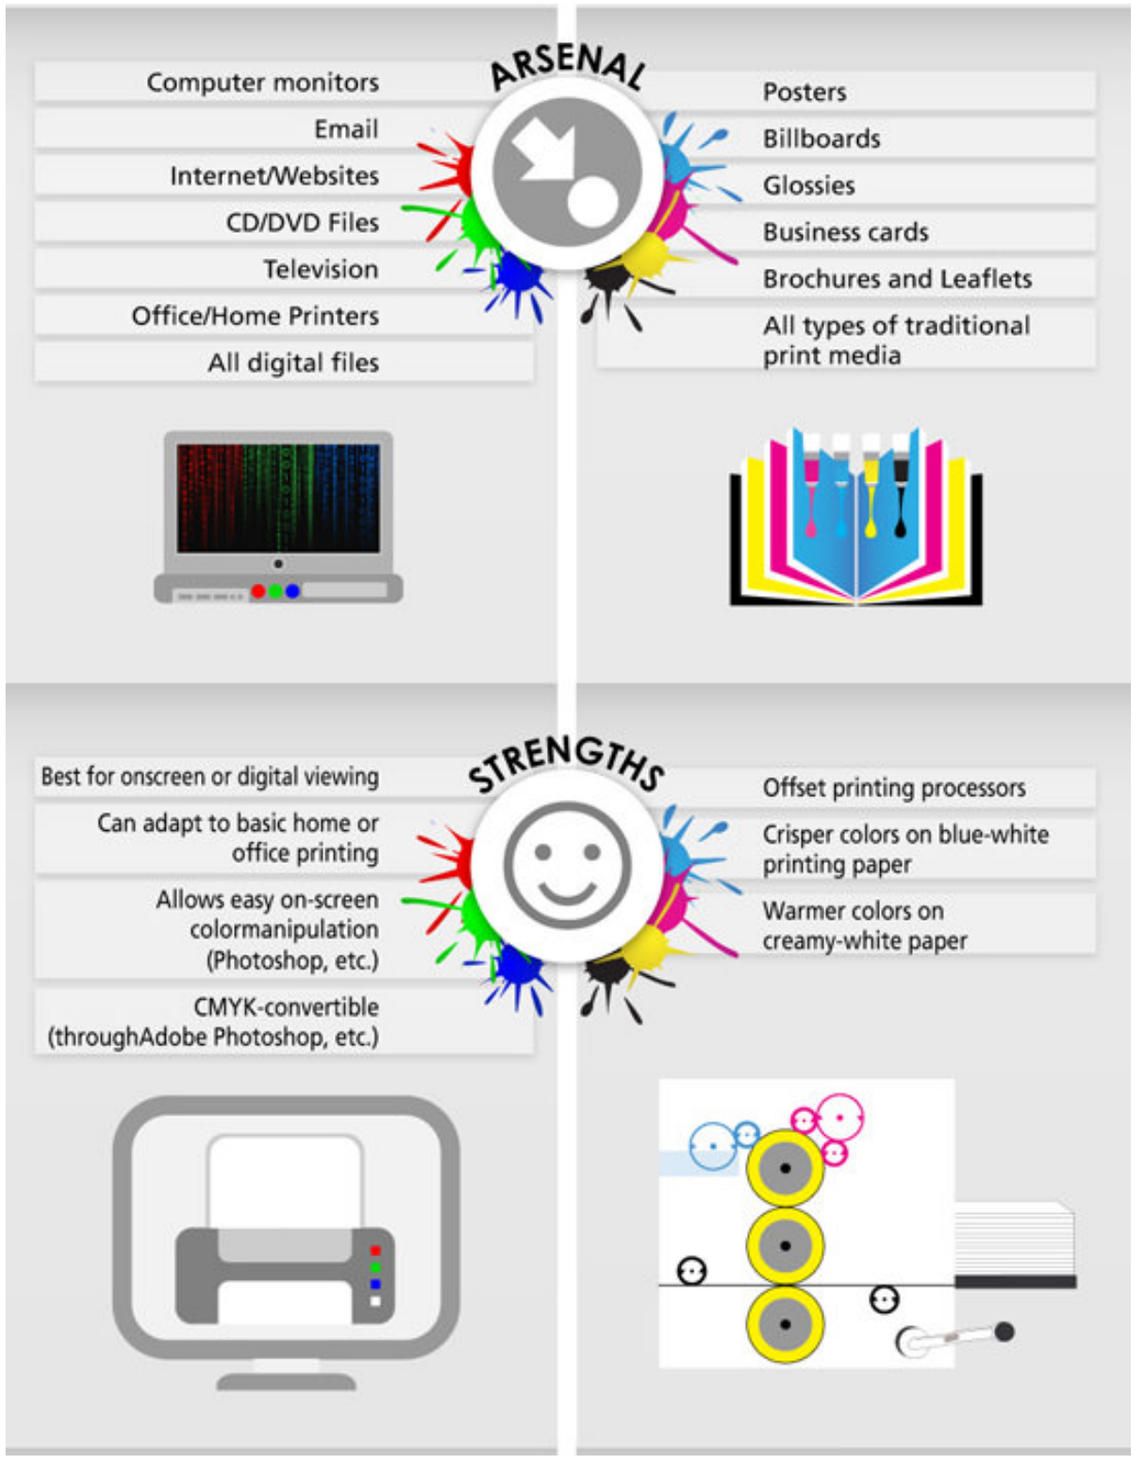
ARSENAL
Computer monitors
Email
Internet/Websites
CD/DVD Files
Television
Office/Home Printers
All digital files
Posters
Billboards
Glossies
Business cards
Brochures and Leaflets
All types of traditional print media
STRENGTHS
Best for onscreen or digital viewing
Can adapt to basic home or office printing
Allows easy on-screen colormanipulation (Photoshop, etc.)
CMYK-convertible (throughAdobe Photoshop, etc.)
Offset printing processors
Crisper colors on blue-white printing paper
Warmer colors on creamy-white paper

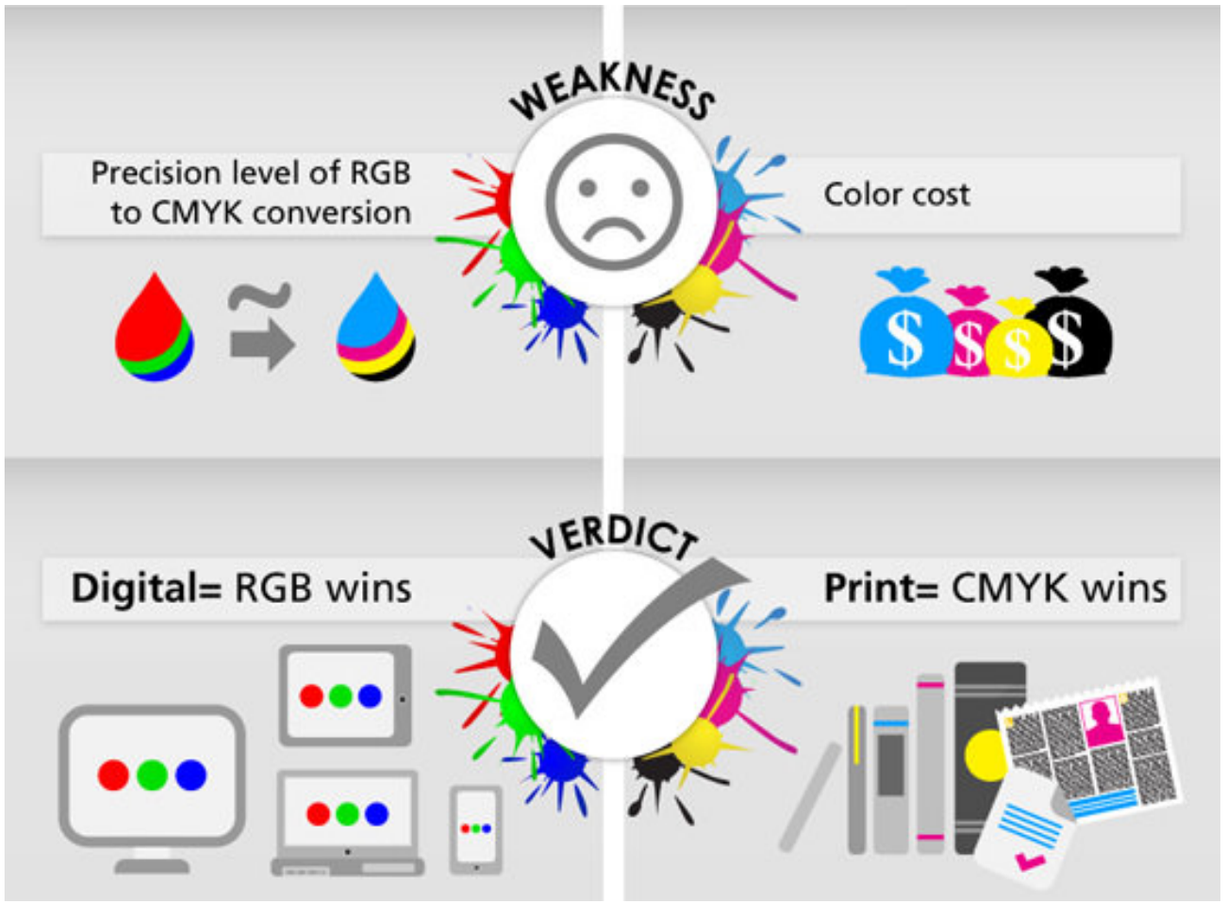
WEAKNESS
Precision level of RGB to CMYK conversion
Color cost
VERDICT
Digital= RGB wins
Print= CMYK wins

A Hue, Saturation & Brightness Model for Mixing Colour

In Gimp

- click on foreground or background colour in the Tools Palette
  - “paint colour” = “foreground colour”
  - “paper colour” = “background colour”

Desaturate by moving down, saturate by moving up

Darken by moving left, brighten by moving right

Click on the spectrum to change the hue

Or ... type R, G, B or H, S, B values into the text edit boxes

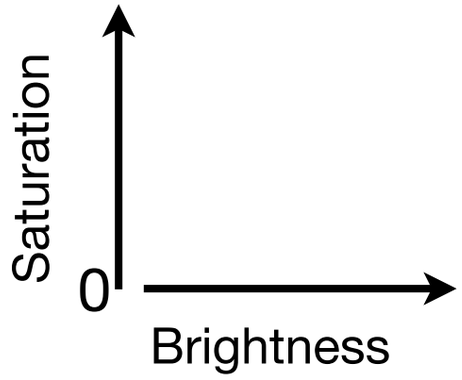

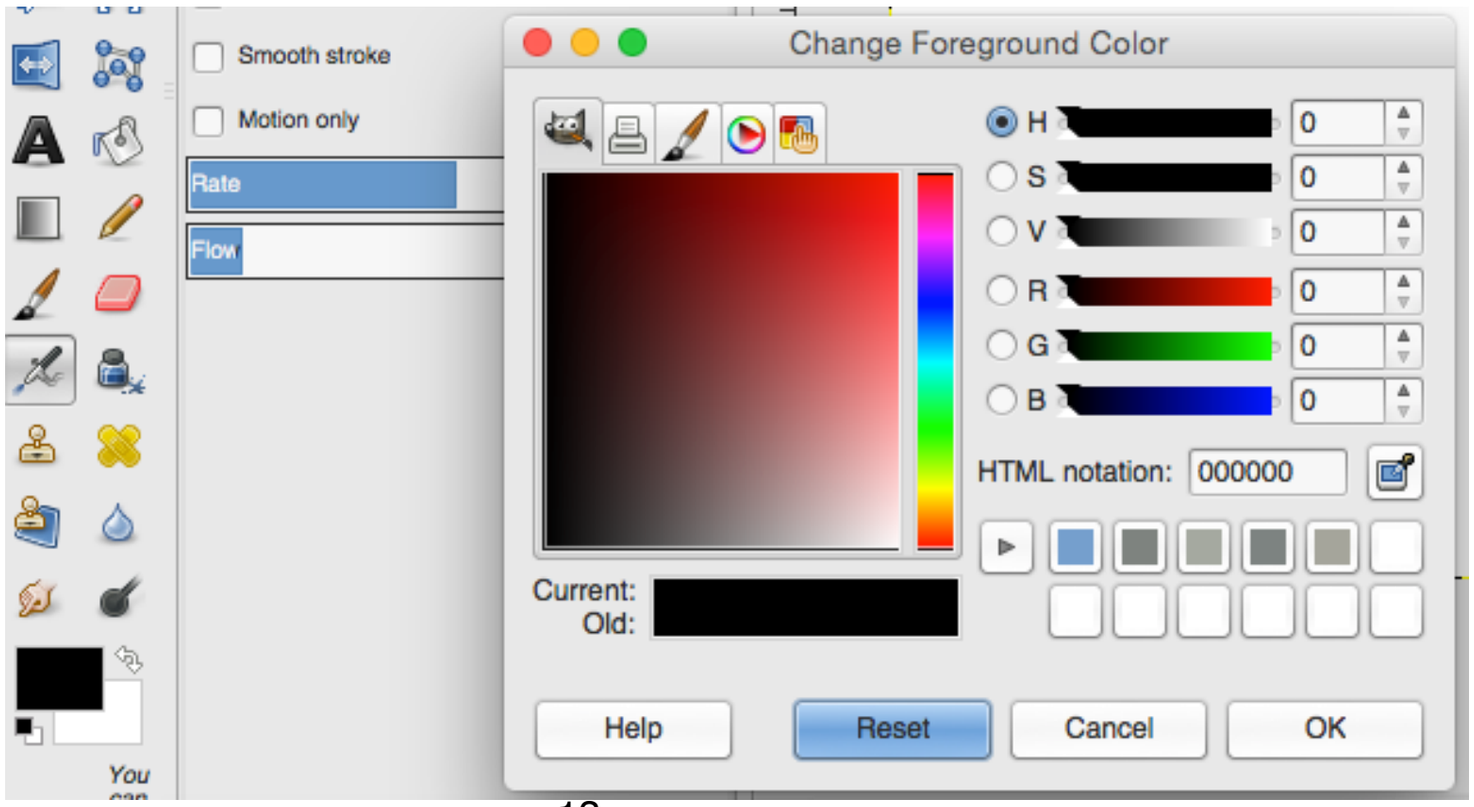

2^N Patterns for N Bits

1 bit
0
1

2 bits
0 0
0 1
1 0
1 1

3 bits
0 0 0
0 0 1
0 1 0
0 1 1
1 0 0
1 0 1
1 1 0
1 1 1

4 bits
0 0 0 0
0 0 0 1
0 0 1 0
0 0 1 1
0 1 0 0
0 1 0 1
0 1 1 0
0 1 1 1
1 0 0 0
1 0 0 1
1 0 1 0
1 0 1 1
1 1 0 0
1 1 0 1
1 1 1 0
1 1 1 1

5 bits
0 0 0 0 0
0 0 0 0 1
0 0 0 1 0
0 0 0 1 1
0 0 1 0 0
0 0 1 0 1
0 0 1 1 0
0 0 1 1 1
0 1 0 0 0
0 1 0 0 1
0 1 0 1 0
0 1 0 1 1
0 1 1 0 0
0 1 1 0 1
0 1 1 1 0
0 1 1 1 1
1 0 0 0 0
1 0 0 0 1
1 0 0 1 0
1 0 0 1 1
1 0 1 0 0
1 0 1 0 1
1 0 1 1 0
1 0 1 1 1
1 1 0 0 0
1 1 0 0 1
1 1 0 1 0
1 1 0 1 1
1 1 1 0 0
1 1 1 0 1
1 1 1 1 0
1 1 1 1 1

Typical Display “Resolutions” (“Addressability” Would Be Better)

“Spatial Resolution”

- h × v: eg 1024 × 768, 1152 × 870, 1280 × 1024, 1600 × 1200, 2048 × 1536
- h = horizontal, v = vertical
- pixels per inch (aka dots per inch, or dpi) is then determined by the screen’s actual size

Pixel Resolution

"Pixel Resolution" (bits per pixel)

- 24 bits per pixel — 8 for red, 8 for green, 8 for blue

  8 bit => 256 different values (intensities), so

  256 reds × 256 greens × 256 blues

  = 16,772,160 combinations — "millions" of colours in Apple-speak

  = "true colour" in Microsoft-speak

- 16 bits per pixel — 5 for red, 5 for green, 5 for blue

  5 bits => 32 different values (intensities), so

  32 reds × 32 greens × 32 blues

  = 32,768 combinations — "thousands" of colours in Apple-speak

  = "high colour" in Microsoft-speak

- 8 bits per pixel — is different: 256 "colour numbers" defined elsewhere in a "colour table" or "palette"

  each pixel consists not of an (R,G,B) triplet of intensities,
  but instead specifies an entry in the colour table

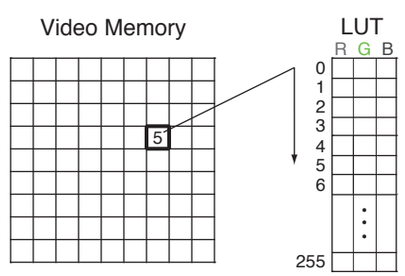

Changing Display Settings on a Mac

Lets you choose

how many pixel bits to use (“colour depth”)

among available spatial resolutions & “display rates”

The higher the resolution, the longer it takes to scroll

more bits to move around in display memory!

For Windows

use the Settings tab in the Display control panel

Multiple displays (both Mac & Windows now)

> Display

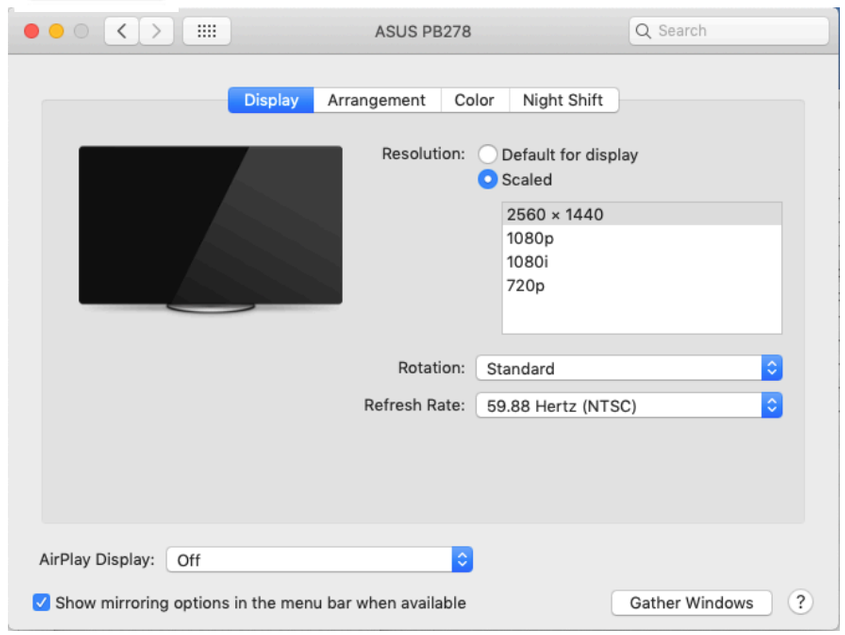

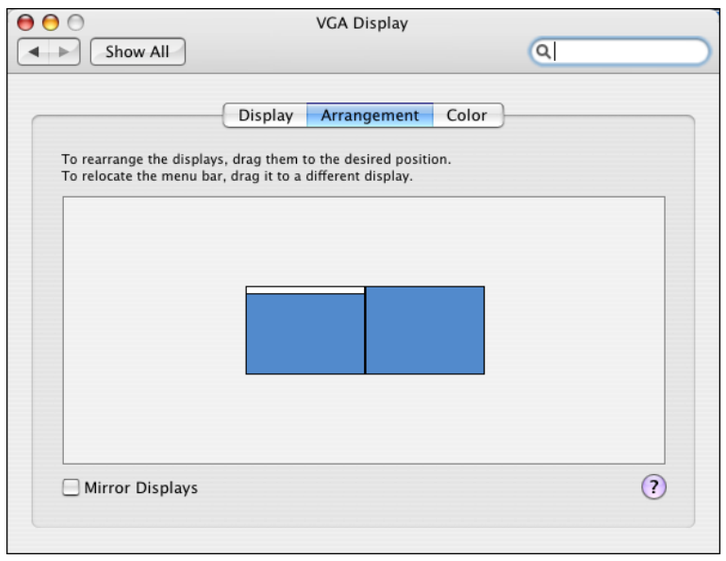

start > Settings > Control Panel >

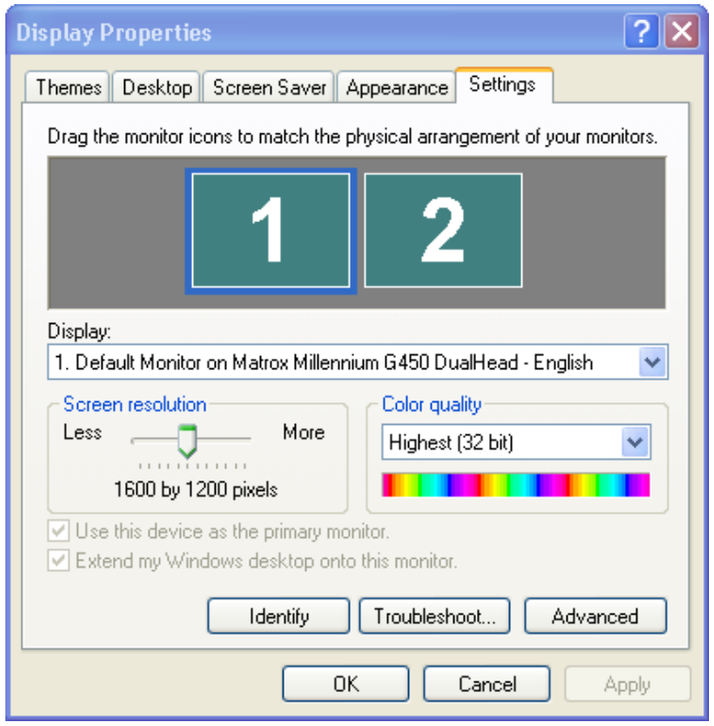

Pixel Graphics File Formats

Stored paintings/images are large
and take a long time to transmit over the internet

640 x 480 x 3 bytes = 0.92 Megabytes

roughly analog TV resolution

1280 x 1024 x 3 bytes = 3.75 Megabytes

1600 x 1200 x 3 bytes = 5.76 Megabytes

If you scan an 8.5” by 11” page in colour at 300 dpi ...

2550 x 3300 x 3 bytes = 22.275 Megabytes

So often you compress them

There are lots of different file formats for storing pixel images

of varying appropriateness for different kinds of images

some reflect different compression techniques

others reflect a programmer’s choice

Some Common Graphics File Formats (1)

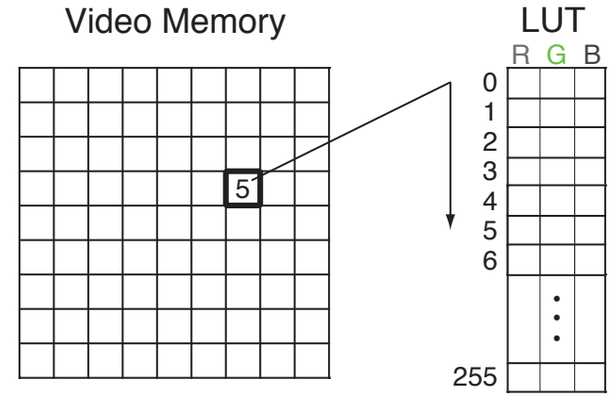

GIF (Graphics Interchange Format)

common on the internet (the “original” internet format)

“loss-less compression”

8-bit colour only, & supports simple animation

each “8-bit colour” actually refers to an entry in a separate 256-entry

“colour table” that defines each colour to be some particular 24-bit

(R,G,B) triple

was licensed by Unisys, and use required a fee
(paid by the software vendor; included in the purchase price of software)

— however, the patent expired in 2004

PNG (Portable Network Graphics)

invented in 1996 because of the GIF licensing fee

true-colour capable

loss-less compression

very flexible & cross-platform, but no animation (which GIF supports)

current and recent browsers support it; really old browsers didn’t (eg IE 3)

JPEG (Joint Photographic Experts Group, aka JPG)

true-colour

supports “lossy compression;” variations trade smaller file size for loss of detail

PNG vs JPEG Example (1) — Actual Size

The original "raw" data: 1,350 KB
becomes an 887 KB **png** file. (**1.5x**)

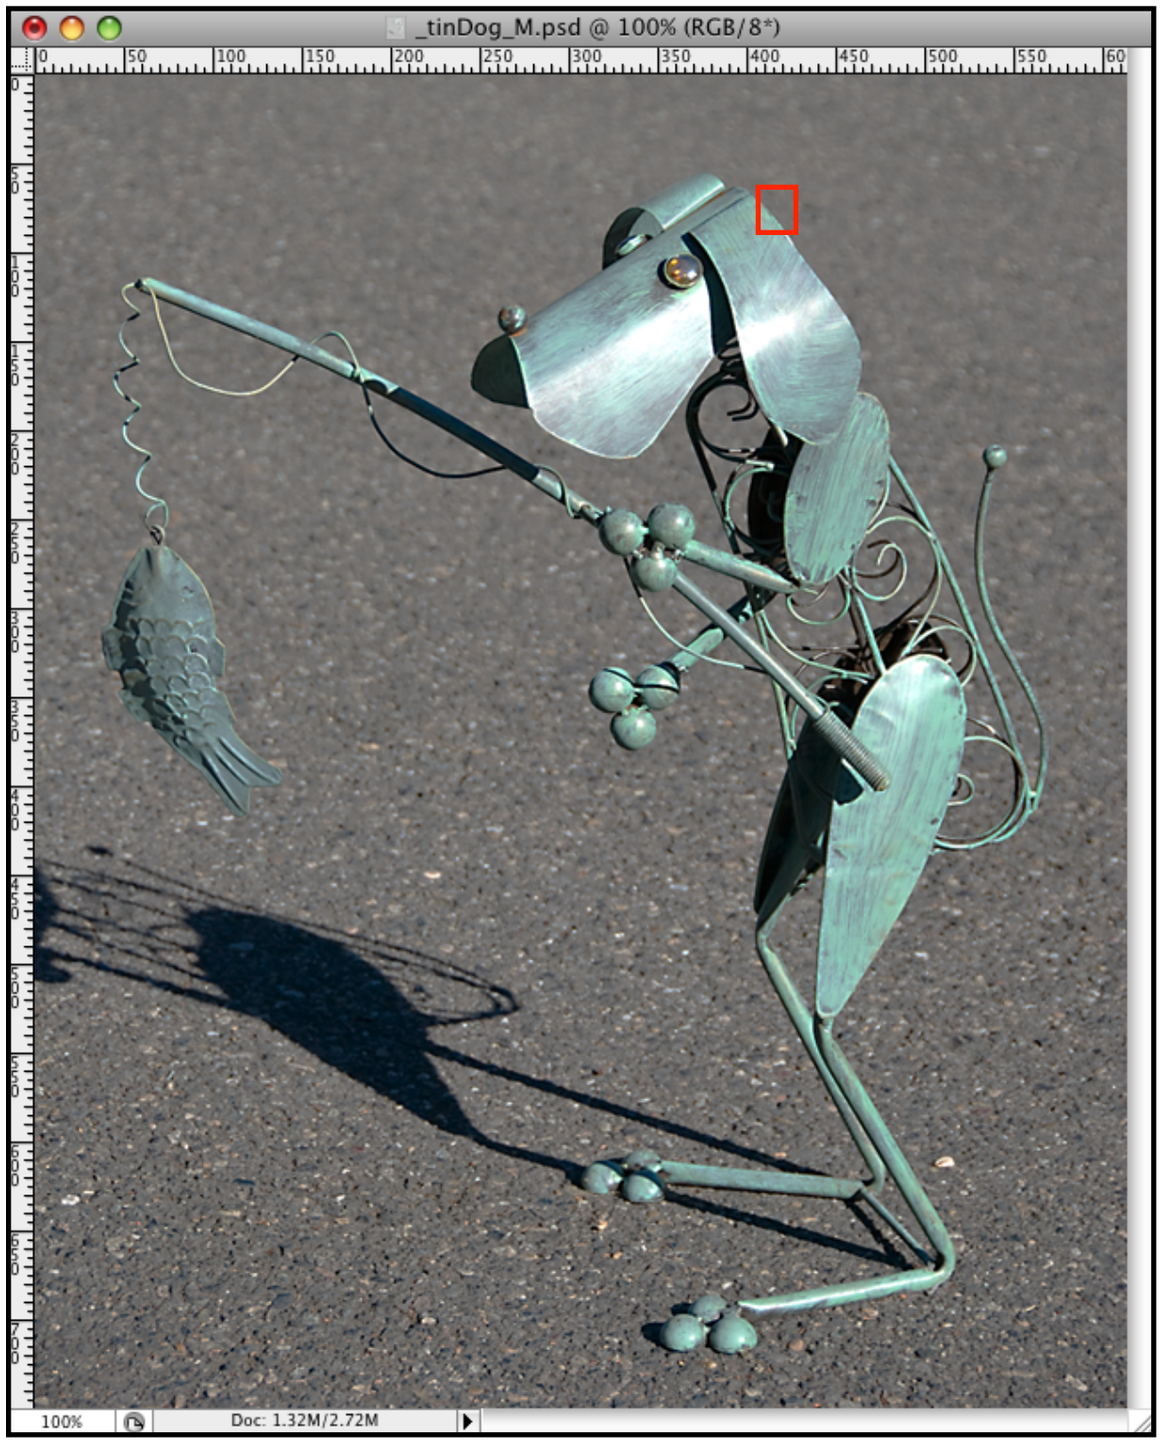

The same data as a "lowest quality"
56 KB **jpeg** file. (**24x**)

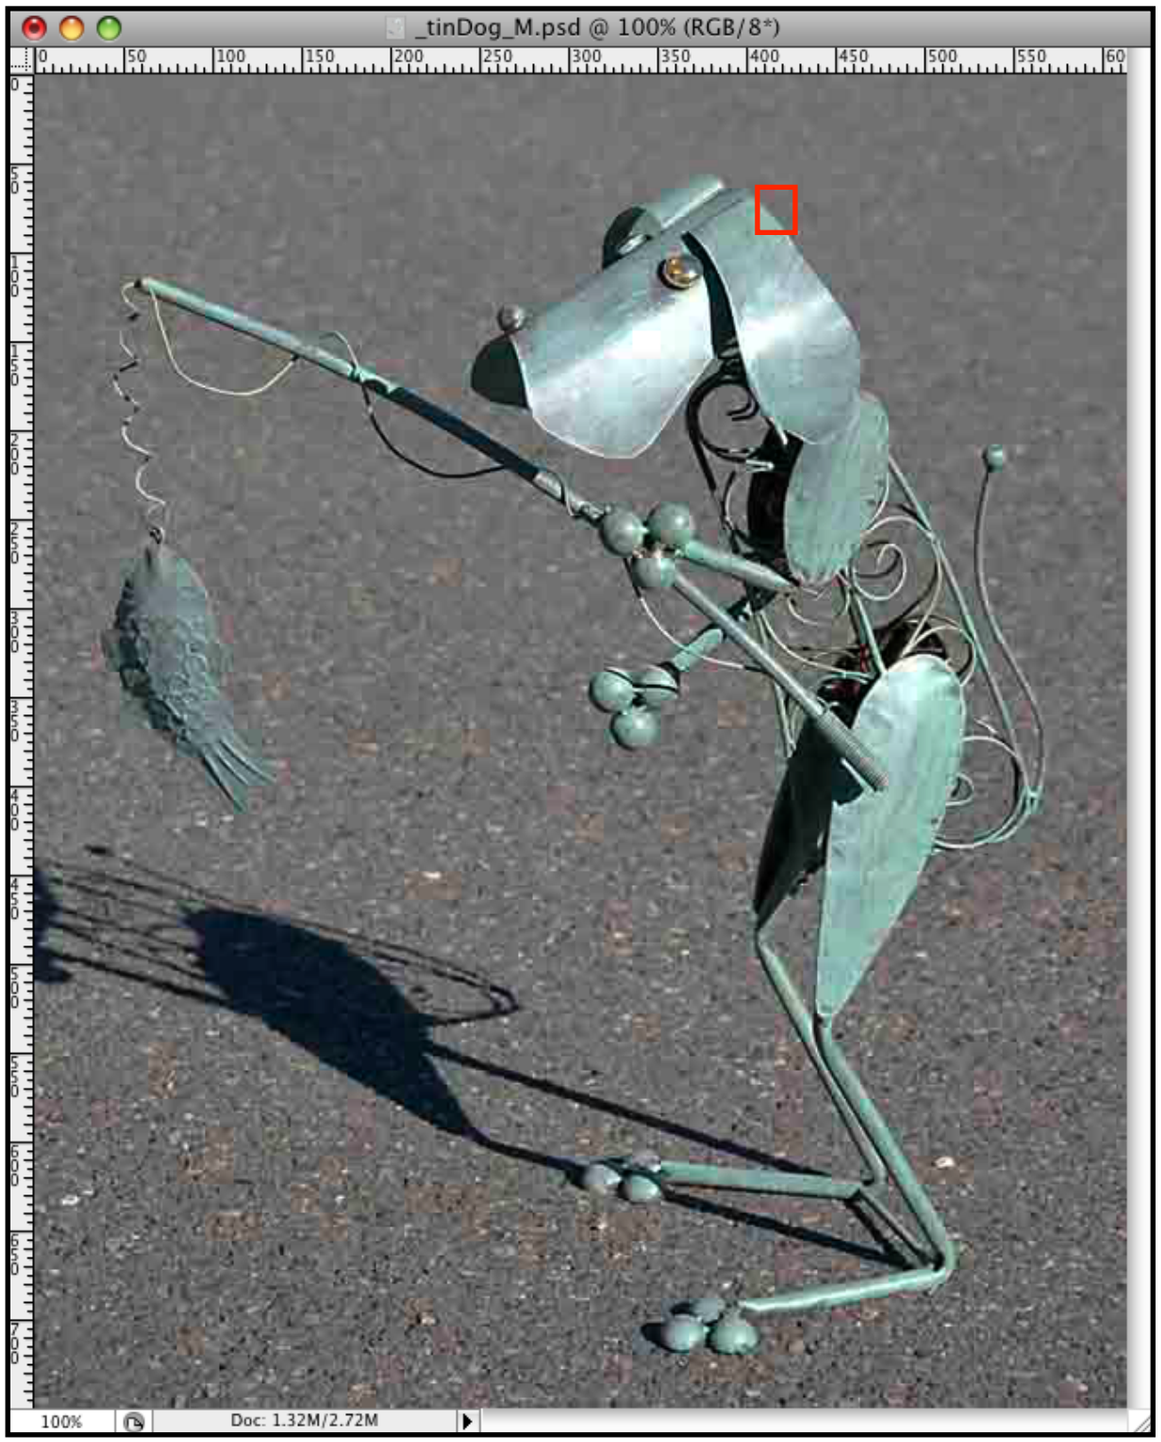

PNG vs JPEG Example (1) — The Selected Area Zoomed 2885 %

The highlighted raw pixels

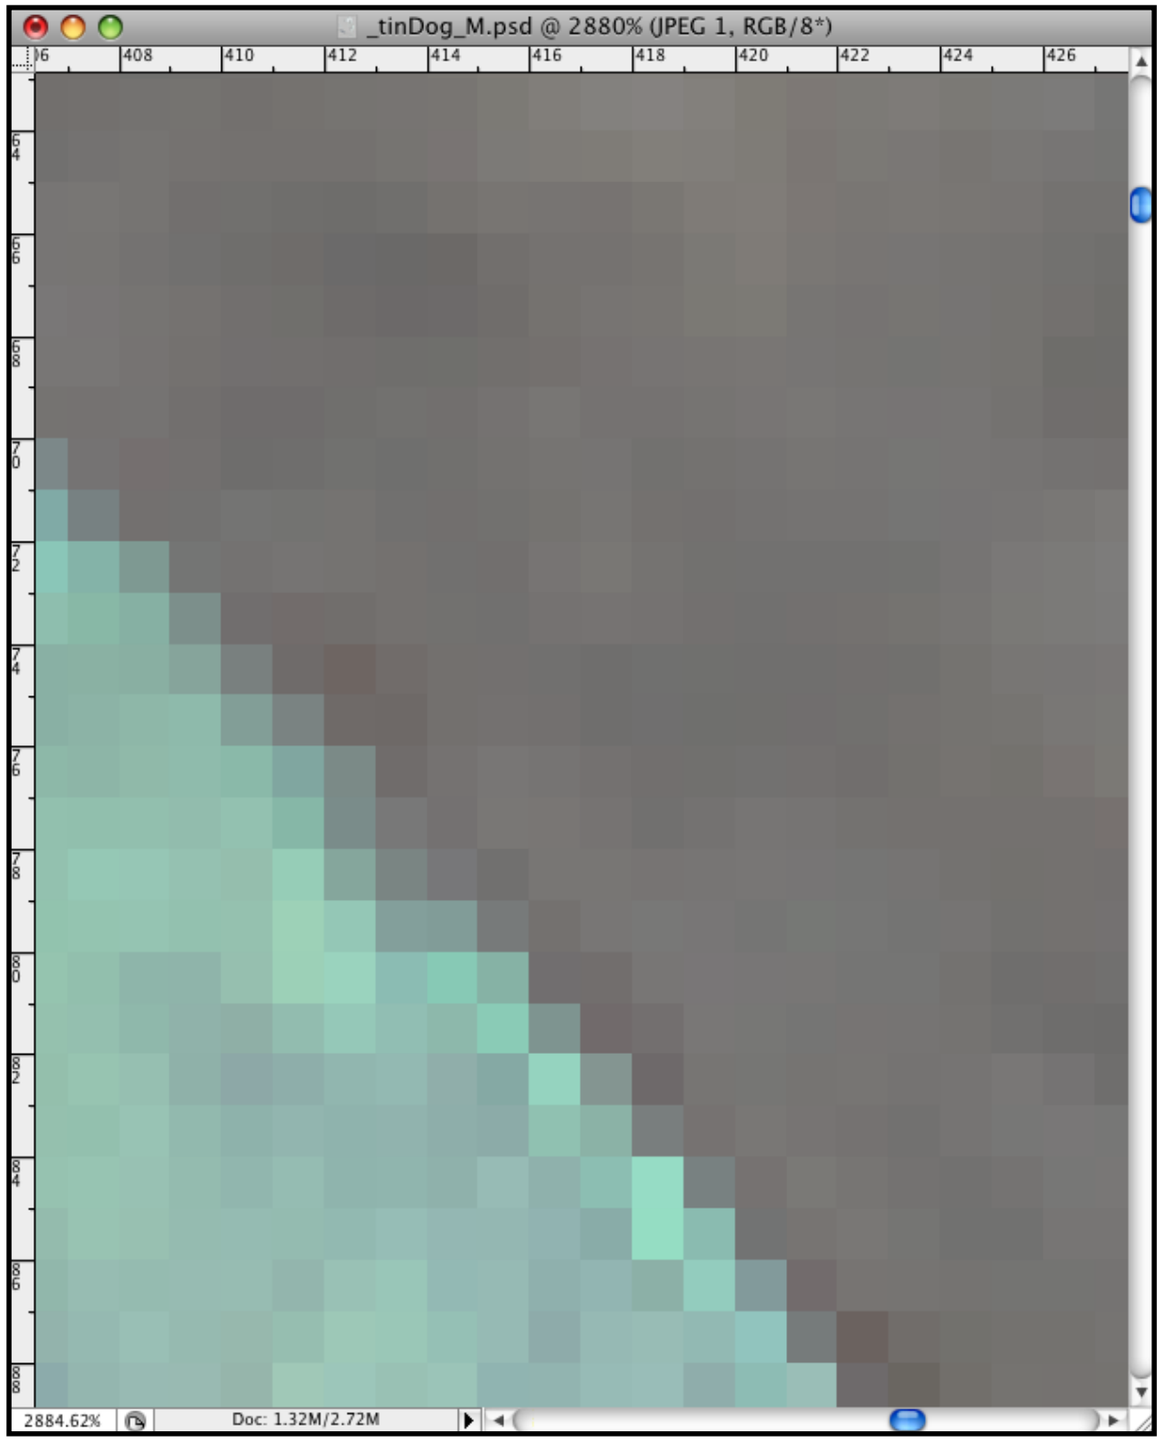

The highlighted jpeg-compressed pixels

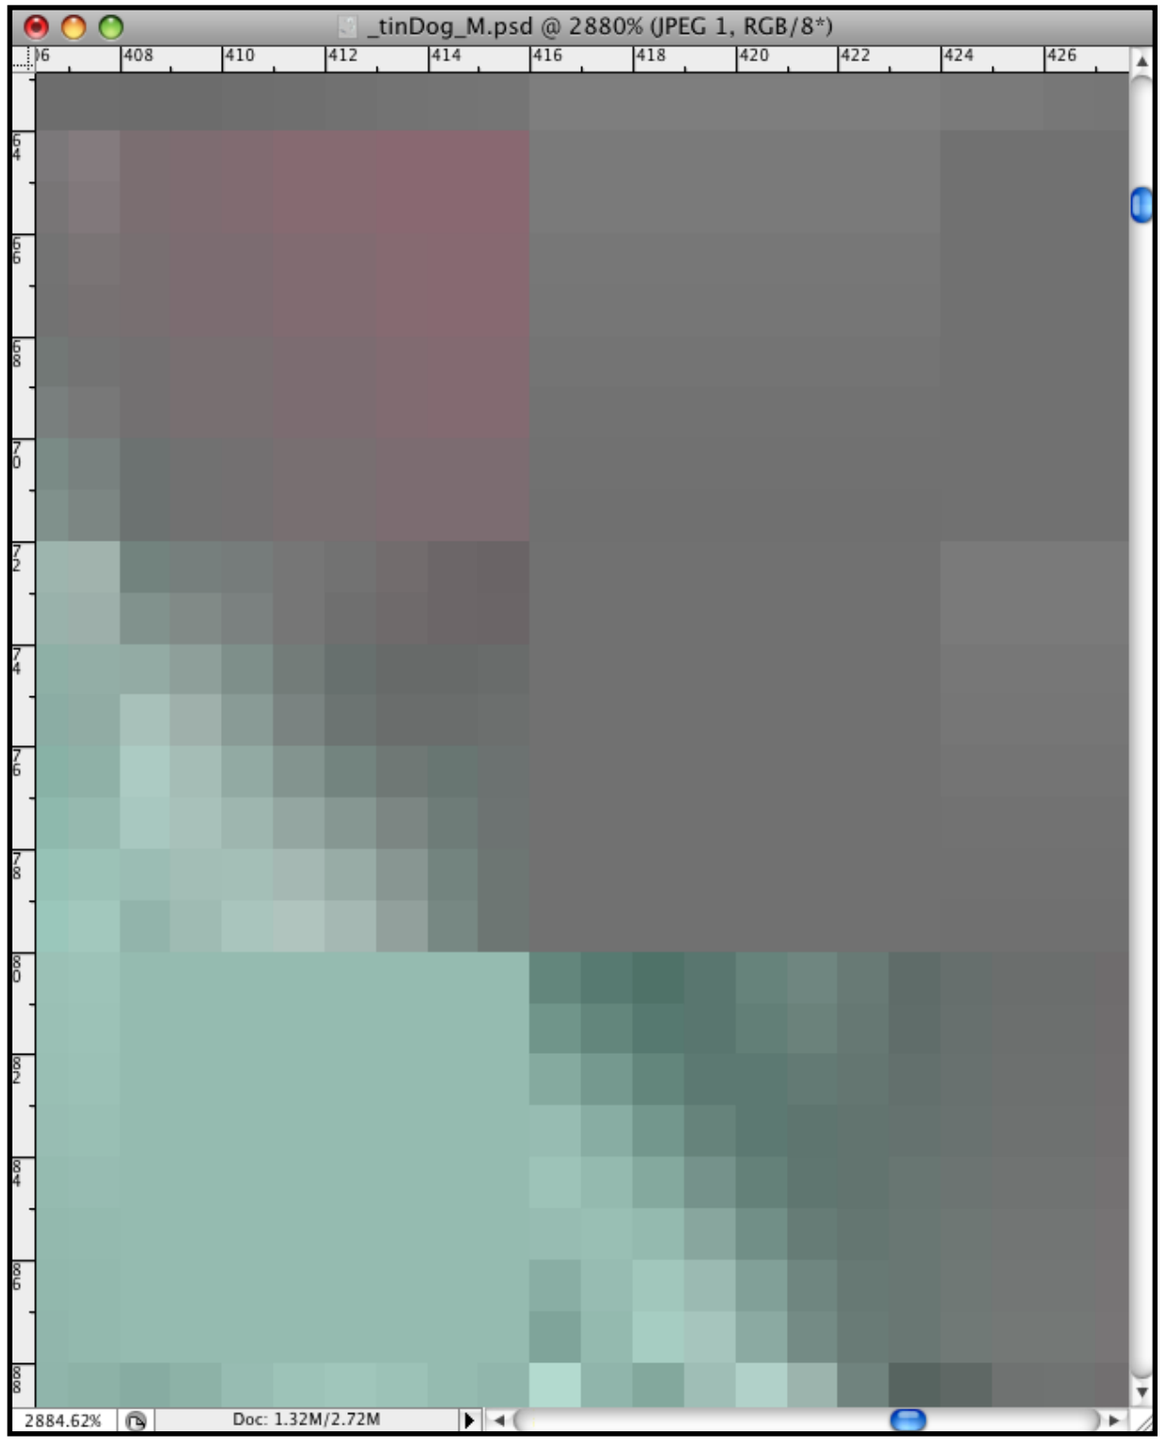

PNG vs JPEG Example (2)

The original – a perfectly sharp boundary.
219 KB of raw data;
728 bytes of **png**-compressed data. (309x)

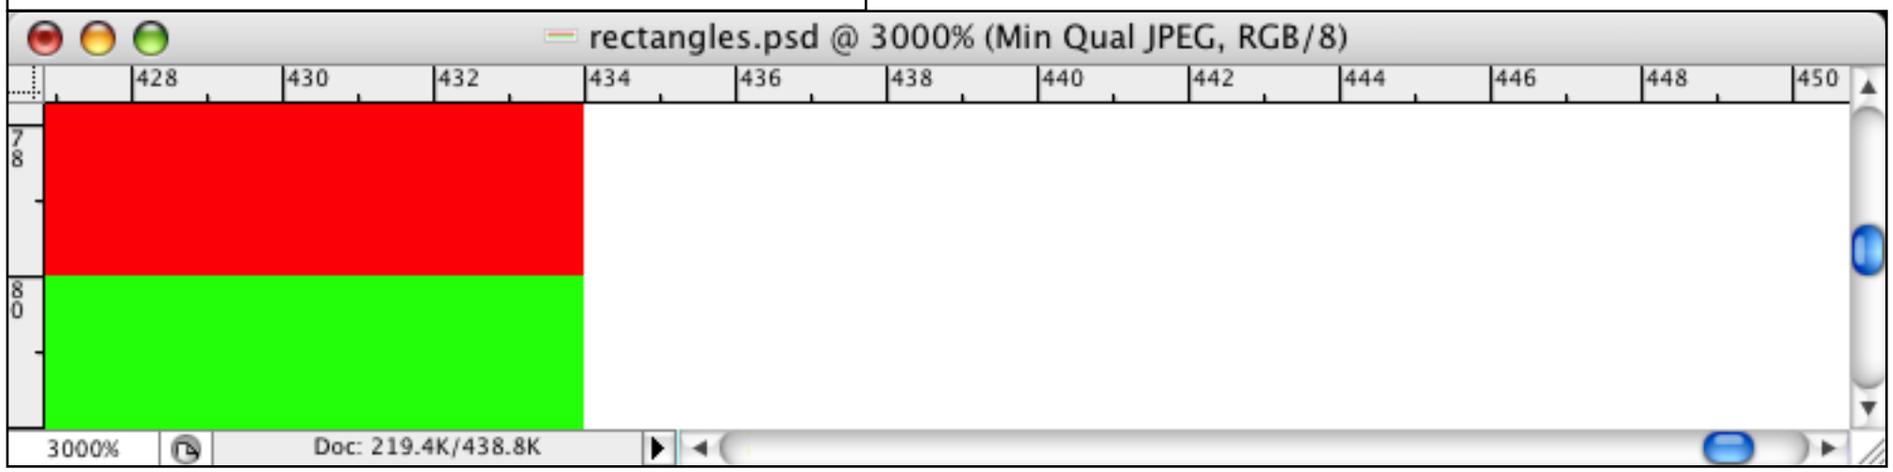

Minimum jpeg quality - border artifacts.
219 KB of raw data;
2,500 bytes of **jpeg**-compressed data. (89x)

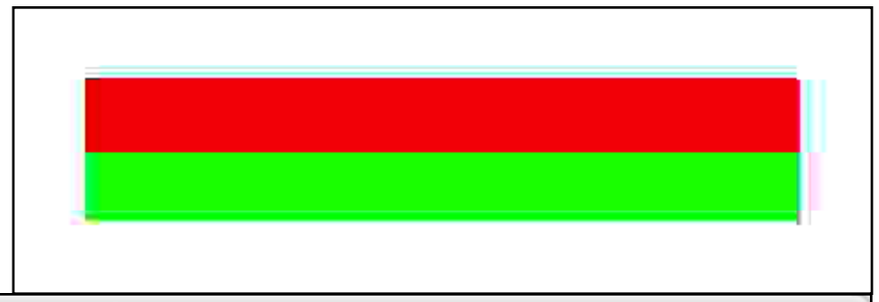

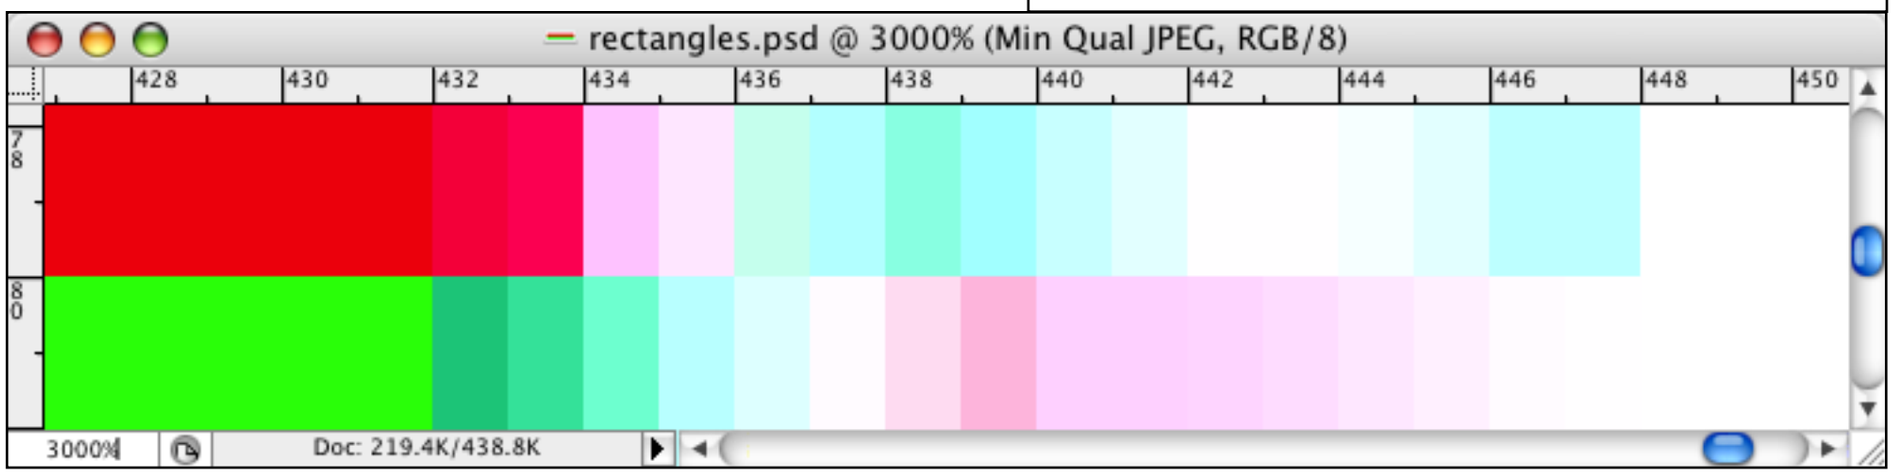

Some Common Graphics File Formats (2)

TIFF (Tagged Image File Format)

- both 8-bit & true-colour
- loss-less compression
- a cross-platform standard
- some browsers don't support

BMP (Windows Bitmap)

- both 8-bit & true-colour
- loss-less compression
- standard Windows format

"Raw" digital camera format

- just the RGB pixel intensities measured by camera
  (as distinct from JPEG, which is always an option)
  and w/o in-camera image processing (eg noise reduction)
- the file format produced is typically proprietary, however

Final words

- *"Different file format" means "the bit representation of the data is different."*
- *Renaming foo.png to foo.jpg doesn't change the file's format.*

Pixel Graphics with Gimp — Painting (1)

You “paint” by click-dragging the mouse

- which controls a “brush” (the tracker) (Paintbrush Tool)
- whose width, shape, colour, etc, you control

  (see “Paintbrush Tool” in online Help)

A hard, opaque brush

A soft (50%), opaque (50%) brush

A hard, semi-transparent brush

A soft (50%), semi-transparent (50%) brush

The tools palette

The brush tool

The foreground colour

Pixel Graphics with Gimp — Painting (2)

As you move the mouse

- pixels over which the brush passes are modified
- may be completely replaced by the brush colour
- may be combined with the brush colour

    — in Adobe-speak, according to the painting “mode”

    applies the “Paint” colour (aka the “foreground” colour)

    - erases to the “Paper” colour (aka the “background” colour)

    - (Bucket Fill Tool) pours the “Fill” colour (which is, in fact, the current foreground colour)

Both and are applied with a brush you can reshape

How do you select a Paint / Paper / Fill colour?

- apply the Poke pearl!

The options panel

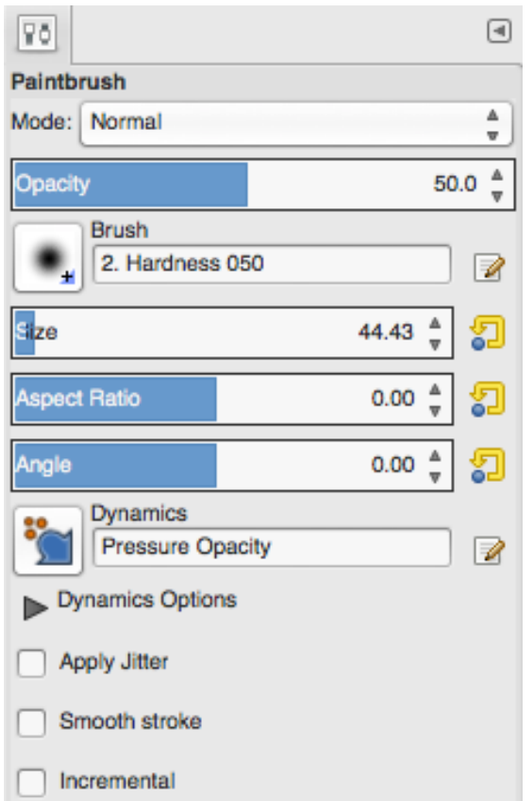

Pixel Graphics with Gimp — Painting (3)

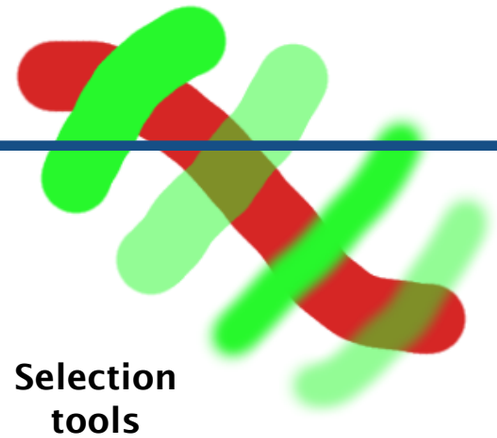

Brush “transparency” (the reverse of “opacity”)

the percentage of what’s underneath that shows through

“Wet” or “soft” edges

~ transparency is added to the edge of a brush stroke

0% for none (a “hard” or “sharp” edge)

Painting or filling a region: use the “marquee tool”

- then you can only paint within the “selected” region
- the remainder of the image is “masked”

“Feathering” a mask by some number of pixels causes smooth blending across the mask boundary; the “number of pixels” controls the width of the blending

To feather a selection, right click and then click Select > Feathering…

Selection tools

Rectangular Select Tool		*Ellipse Select Tool*
Free Select Tool		*Fuzzy Select Tool*
Select by Colour Tool		*Scissors Select Tool*

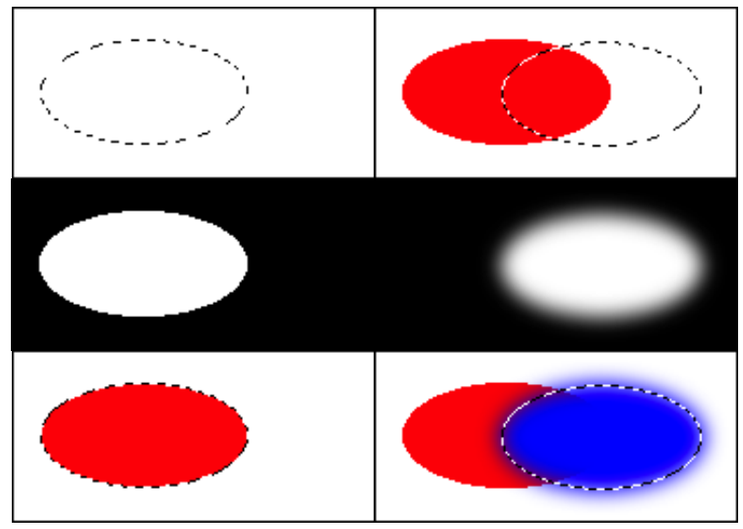

No feathering or anti-aliasing (a hard edge).

A feathered (5-pixel) fill.

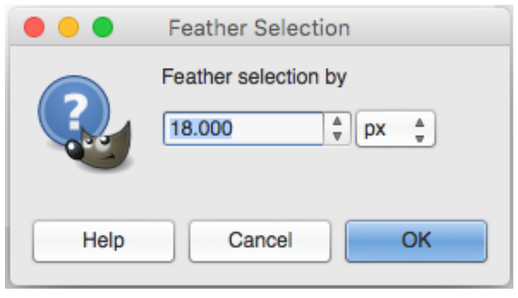

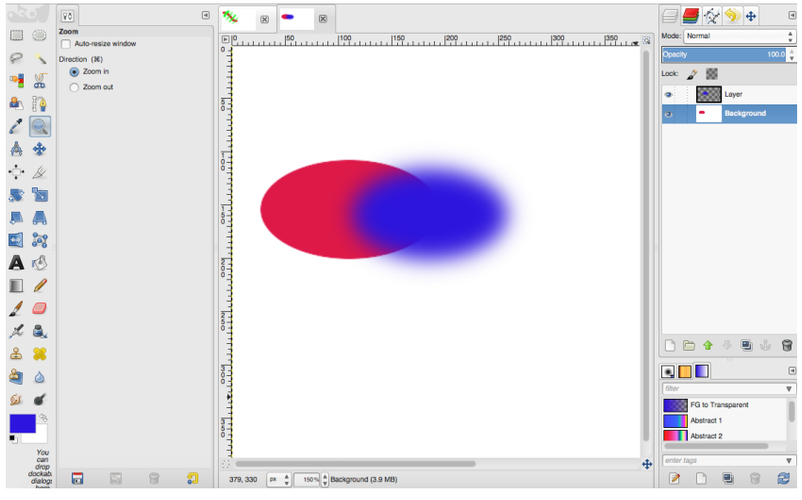

Retouching Scanned Images & Digital Photographs

Can be digitally manipulated in many interesting & useful ways

"Retouching" consists of such operations as manually

moderating highlights
(eg from flash reflection)

painting (or cloning)
over blemishes

— generally speaking, using a mouse to point to pixels that should be altered in some way, often based on the color of the pixels underneath the tracker

perhaps just changing the hue, the saturation, the brightness, etc (the painting "mode" again)

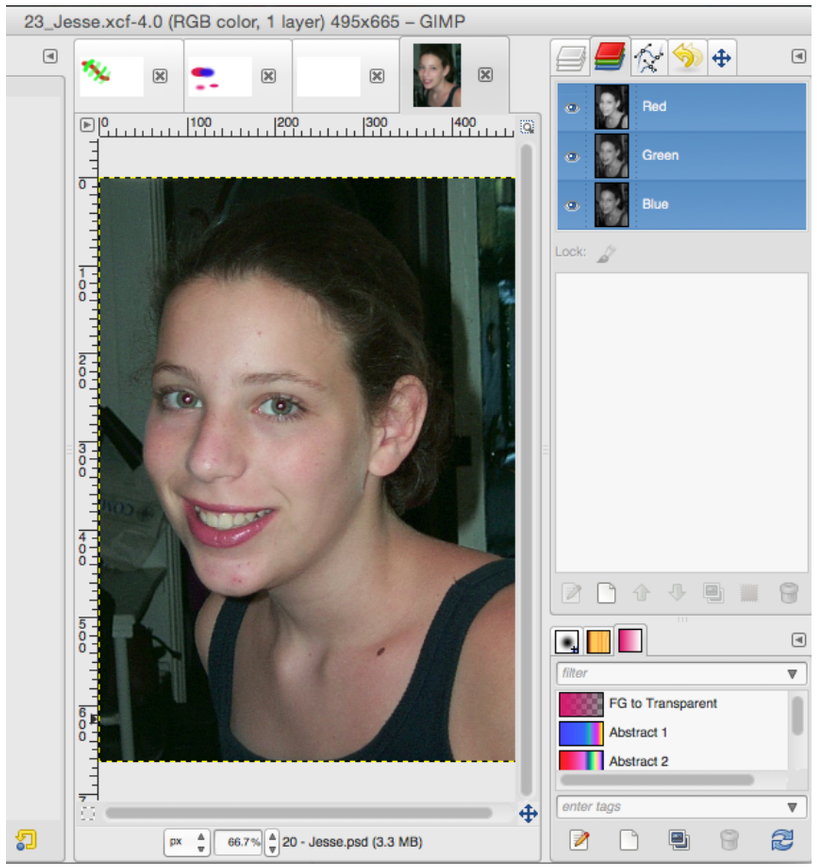

Demo!

Irrelevant aside: for lot's more interesting stuff about vision & the eye, see "The Eye—A Natural History," by Simon Ings, $15 at Chapters.
(Ings is an excellent science writer.)

Image Manipulation

Examples

- Level (Re)mapping
- Sharpening
- Colour Balance
- Brightness / Contrast
- Hue / Saturation / Lightness

These may be applied

- to the image as a whole
- to a selected portion of the image
  - There are a variety of tools and techniques for selecting portions of an image.
  - We will mostly use direct rectangular or elliptical selections
  - Other useful tools include the “fuzzy select tool” and the “scissors select tool”

For the images used to illustrate these, see

- Learn/ Week 3: Pixel Graphics: Files from Pixel Graphics Lecture
- NB: the manipulations we discuss don’t print accurately
(the “colour matching problem”)

Scanned Images & Digital Photographs (Channels in Color)

Are arrays of pixel intensities

- that come from a scanner or digital camera
- that MEASURES (R,G,B) pixel intensities on a rectangular grid, usually producing one byte (8 bits) for EACH of R, G and B at each grid position

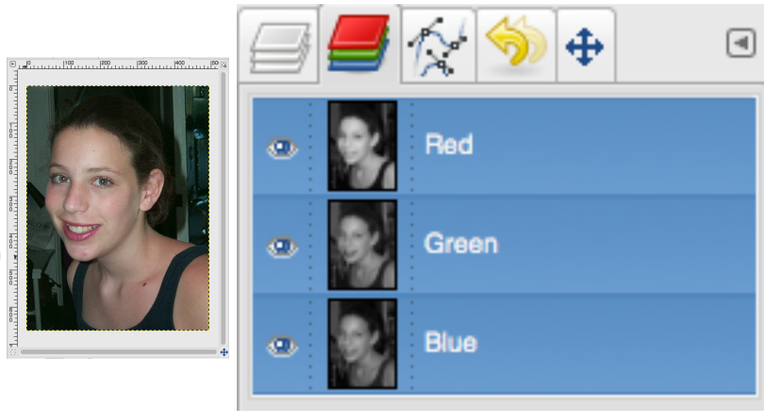

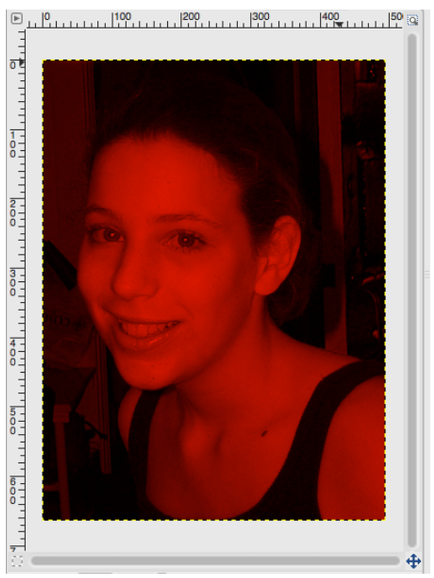

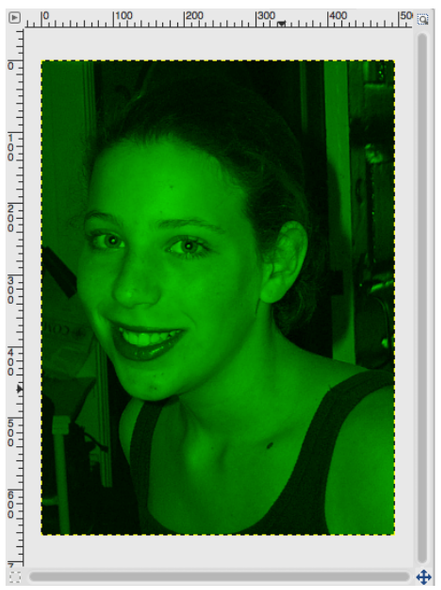

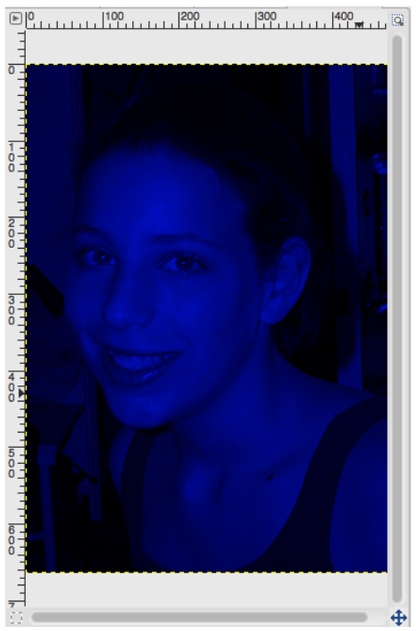

In the rightmost of these three images, which shows only the blue channel, the red and green intensity values at each pixel are temporarily set to zero while generating the display.

The red and green channels are handled similarly.

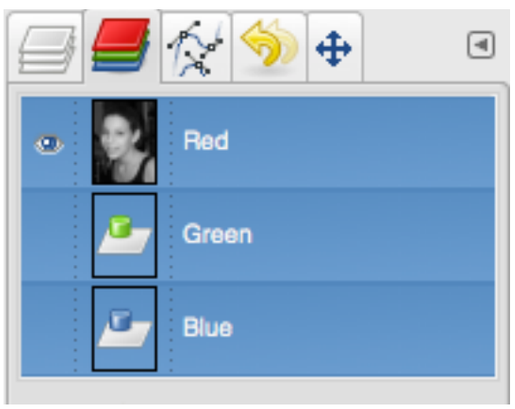

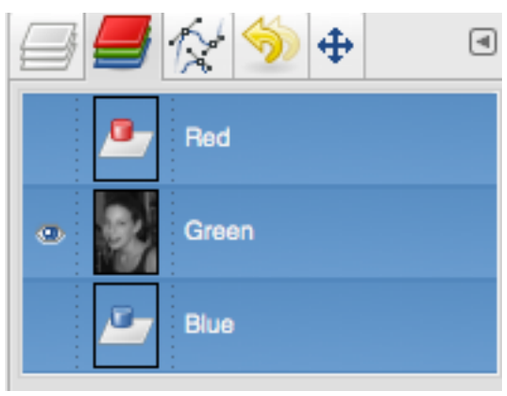

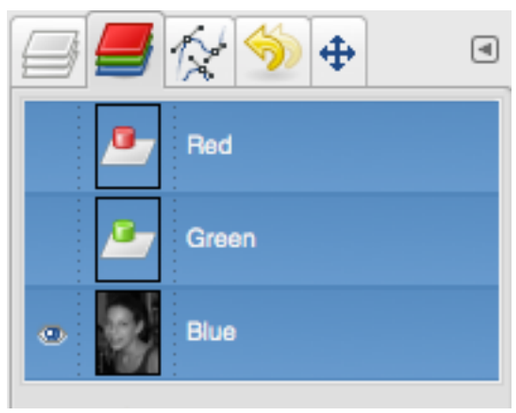

Scanned Images & Digital Photographs (Channels in Gray)

Are arrays of pixel intensities

that come from a scanner or digital camera
that MEASURES (R,G,B) pixel intensities on a rectangular grid
usually producing one byte (8 bits) for EACH of R, G and B
at each grid position

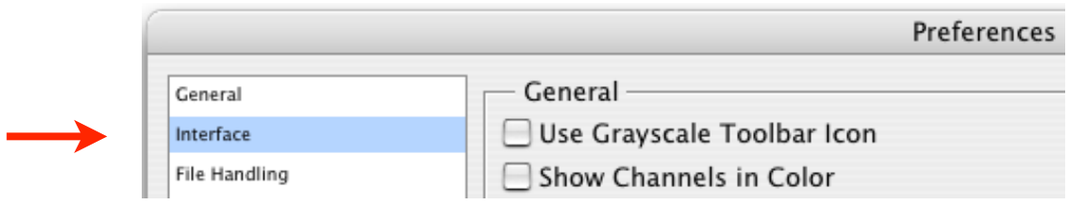

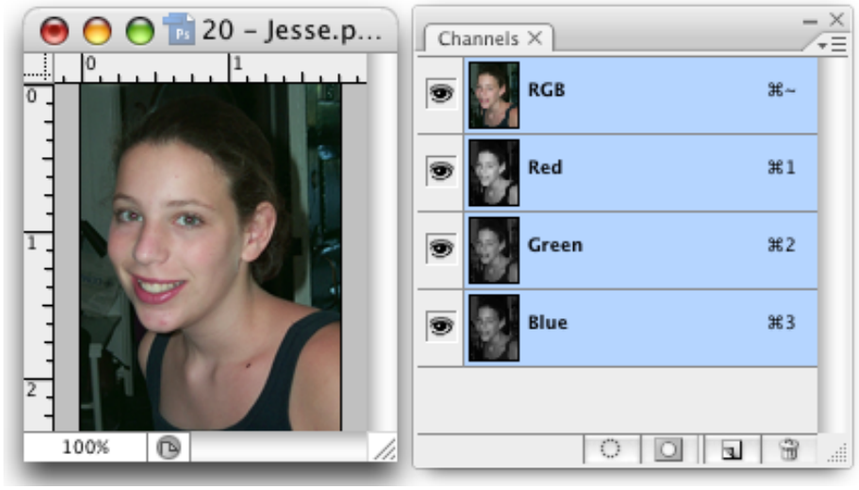

In the rightmost of these three images, which shows only the blue channel, the blue intensity value at each pixel temporarily replaces the red and green intensity values for that pixel while generating the display.

The red and green channels are handled similarly.

Selections (1)

Recall that a selection

is a "mask"

is an 8-bit grey-scale image in which

- white represents complete selection,
- black represents no selection, and
- gray represents partial selection

When you apply color ("paint"), the more fully selected a pixel is, the more paint is applied.

In Gimp you can

save a selection as a "channel,"
(*aka* an "alpha channel")

- restore a selection from channel

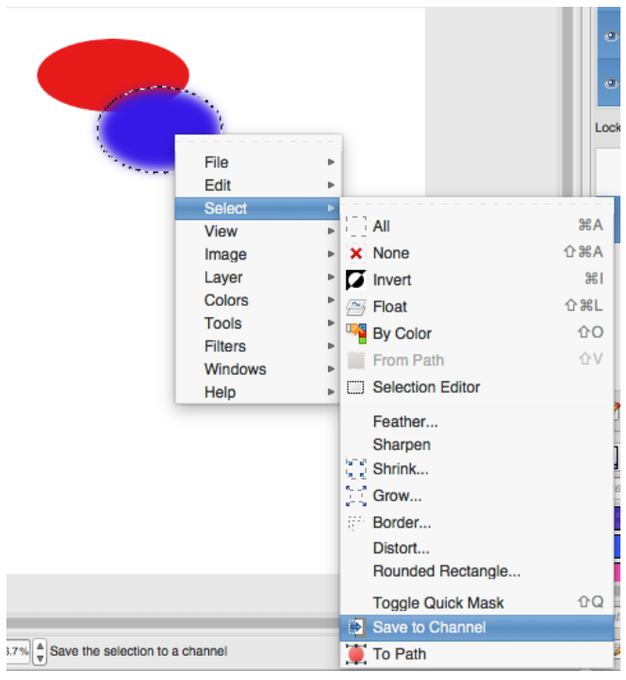

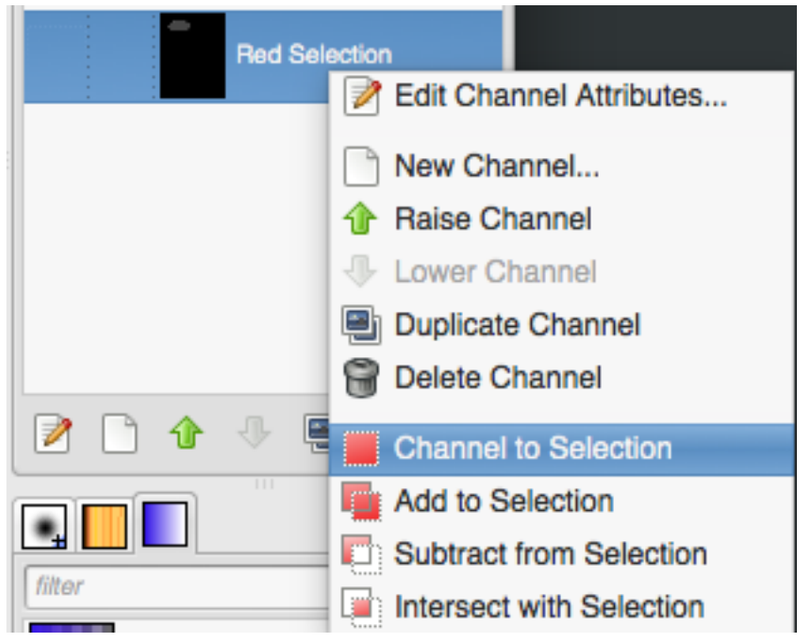

Selections (2)

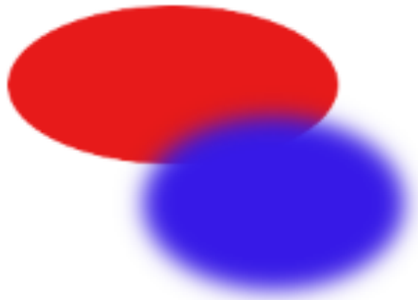

The complete image.

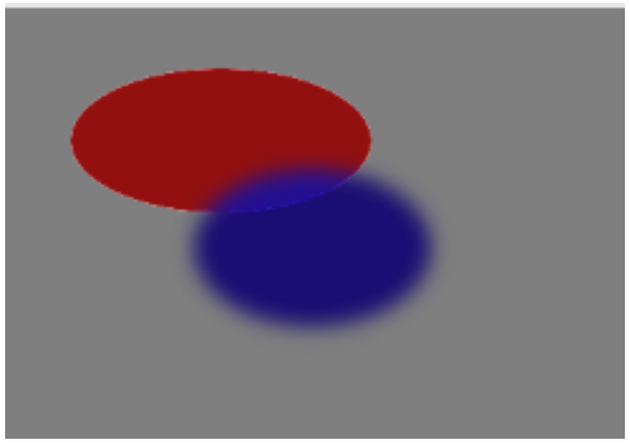

The red selection mask.

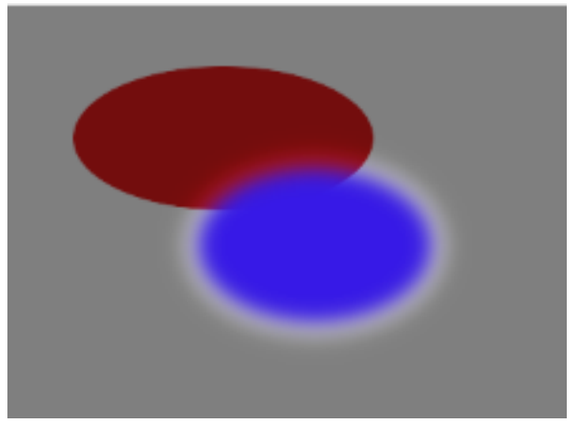

The blue selection mask. Note the gradual transition, which results from feathering.

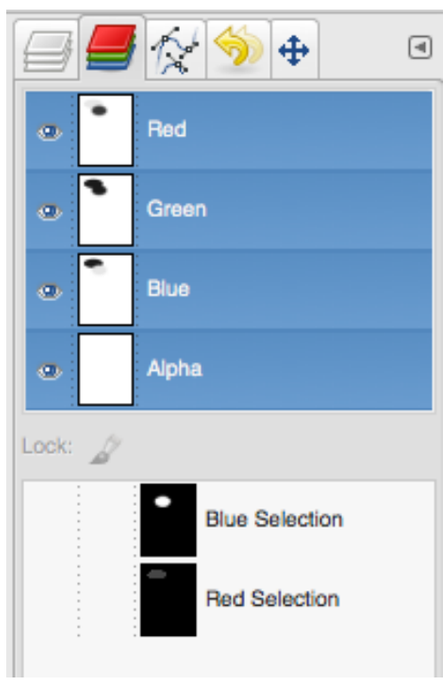

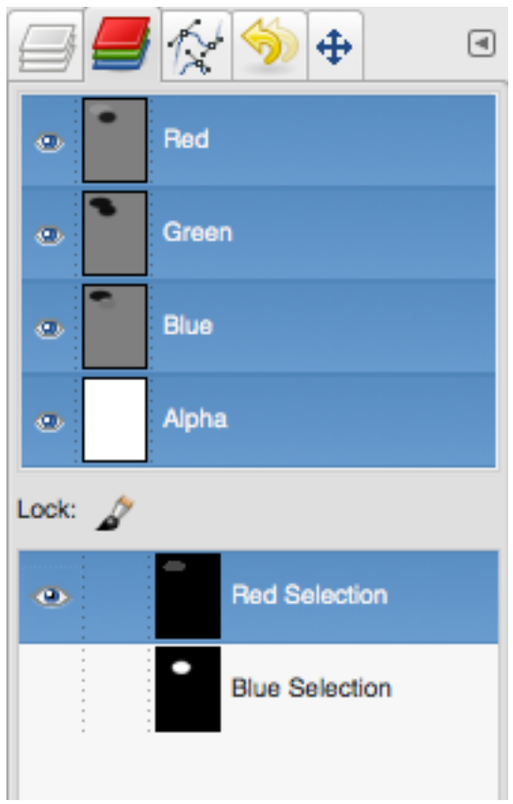

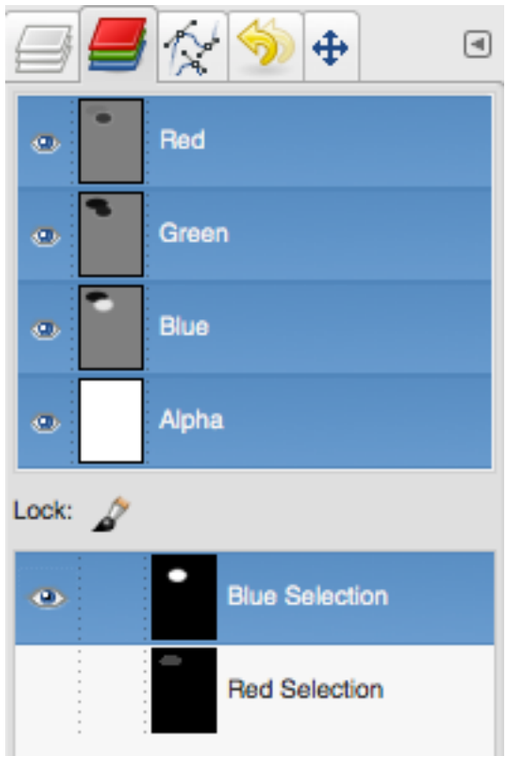

Pixel Histograms

Colors > Levels

The graph plots the number of pixels having each intensity value

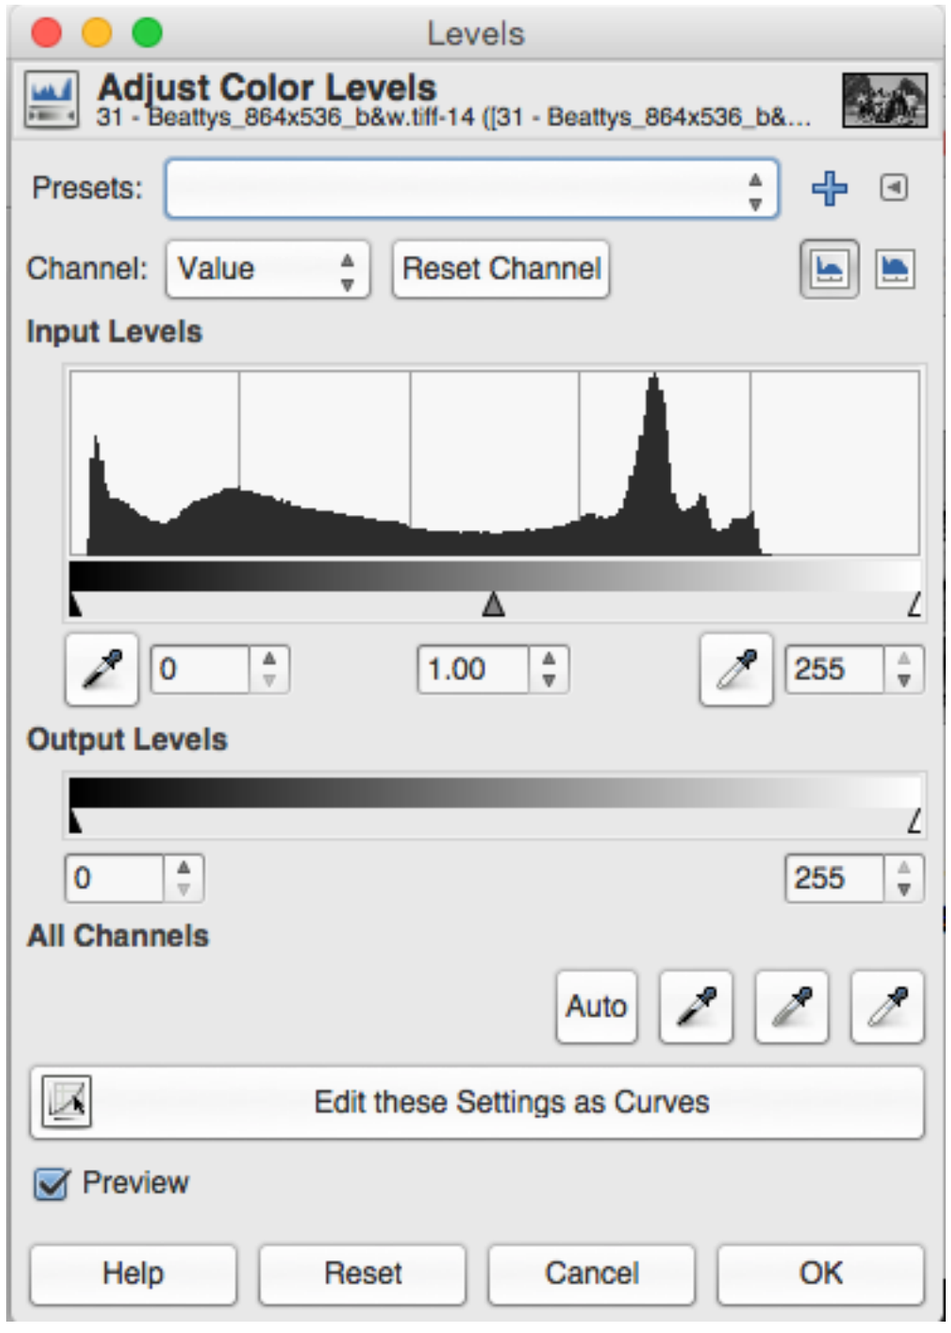

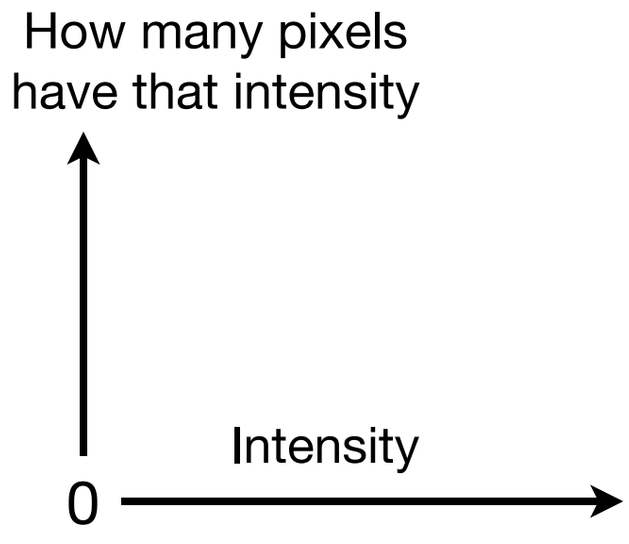

The Levels Dialog Box (1)

Colors > Levels…

Note the unused values at both ends.

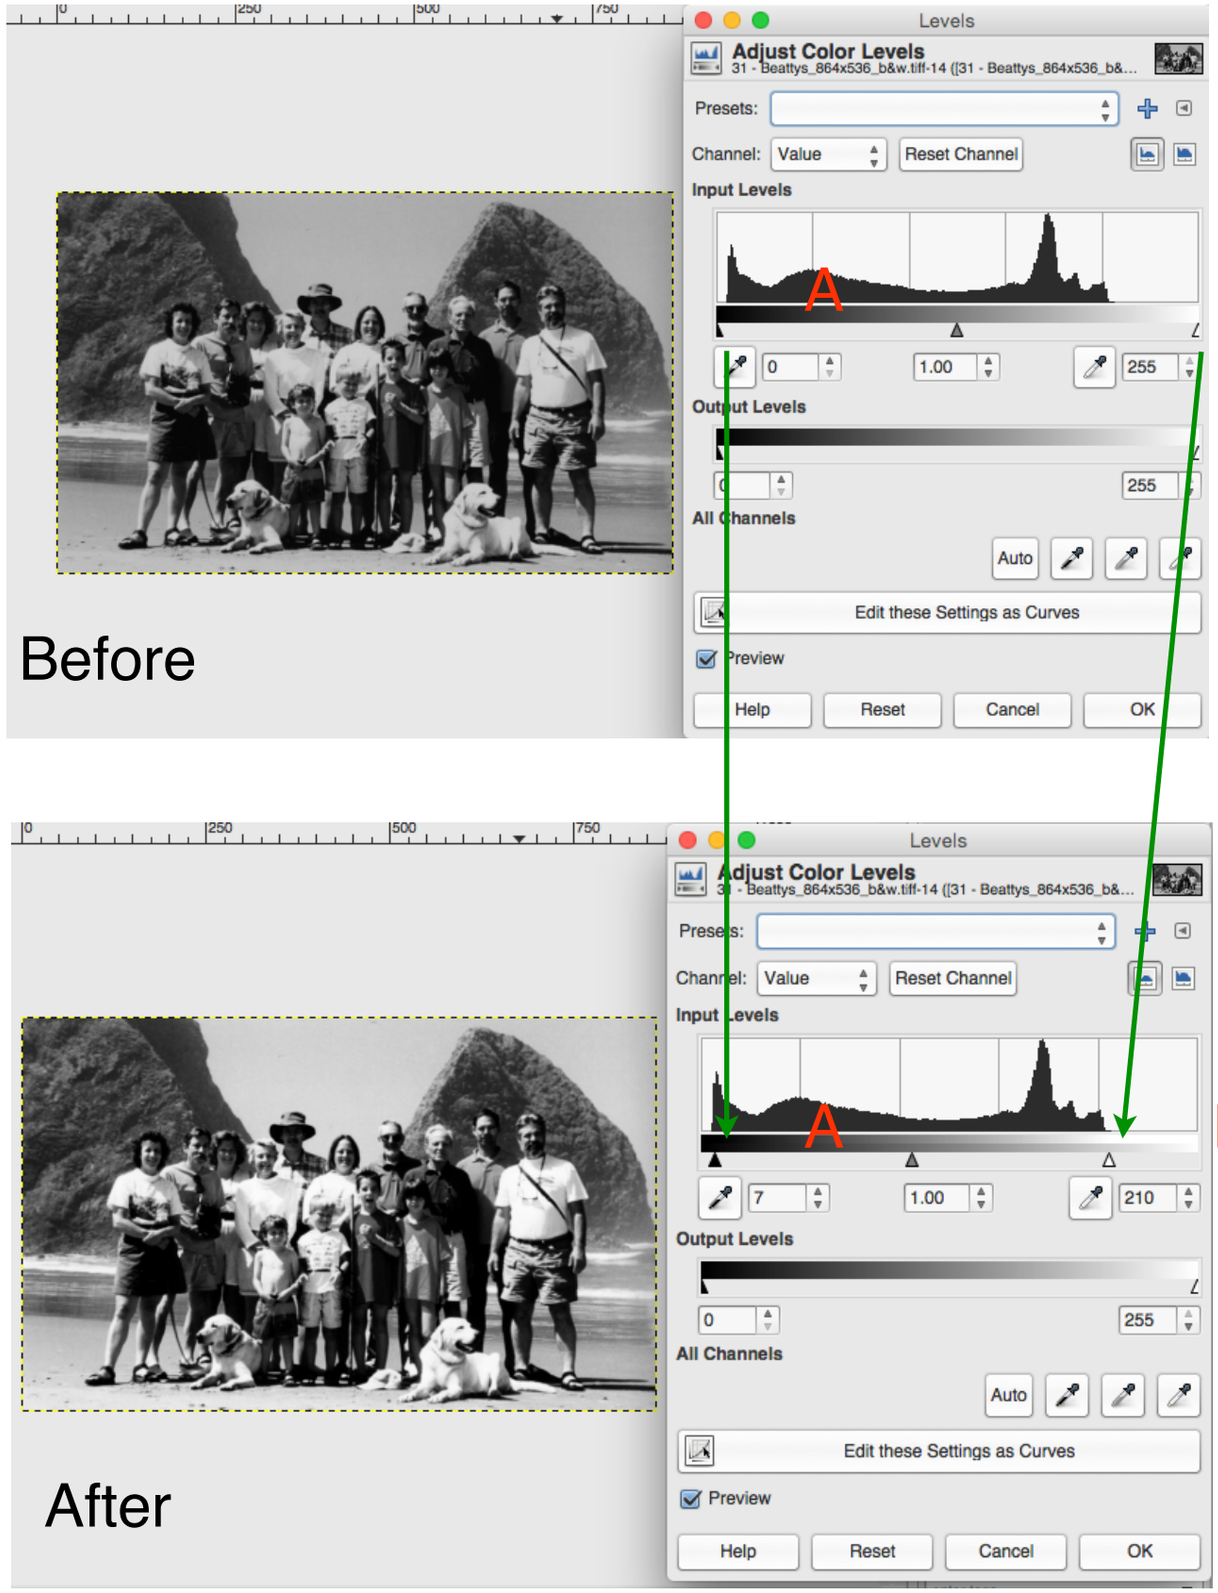

A has been moved right to pixel intensity 7.
B has been moved left to pixel intensity 210.

The Levels Dialog Box (2)

The effect of these particular changes to A and B is to use the full range of available intensities

pixels with intensity 0, 1, ...,7 take on intensity 0

- blacks become blacker

pixels with intensity 210, 211, ... 255 take on intensity 255

- whites become whiter

pixels with intermediate intensities shift accordingly

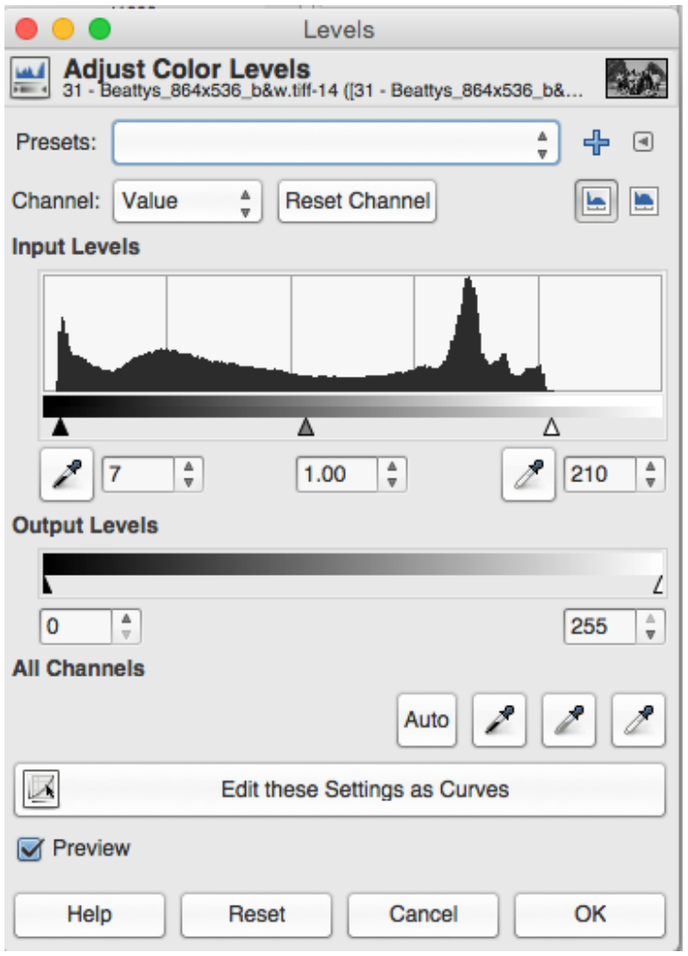

- the graph is “stretched uniformly”
- In other words, for the image as a whole, there’s ***more contrast*** and ***more dynamic range***

Before

After

Shifting Midtones

Move slider C…

left to lighten intermediate intensities or “midtones”

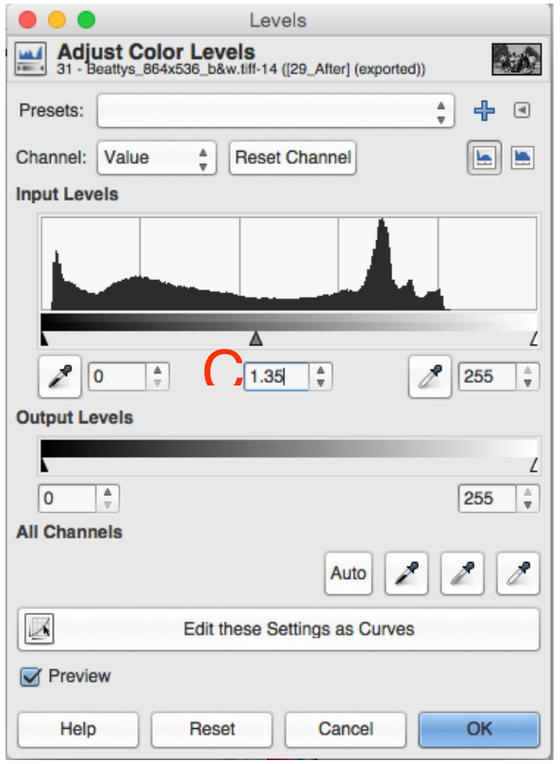

right to darken intermediate intensities

Lightened

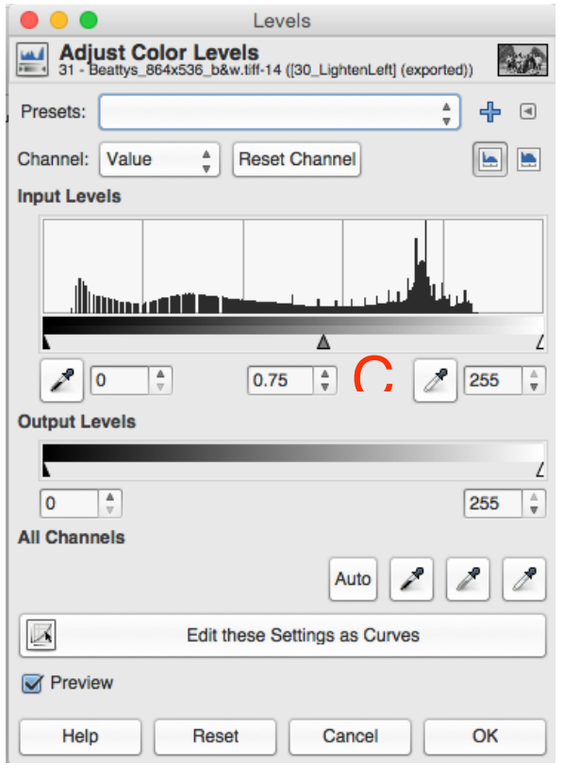

Sharpening

Filters > Enhance > Sharpen (Unsharp Mask)

each pixel value is replaced
by a weighted sum of its neighbours,
in such a way as to
sharpen the appearance of the image
by making edges more visible

Amount (%)

how much to sharpen (start with 30%)

Threshold

by how much adjacent pixels must differ

before sharpening occurs (start with 4)

an effort at edge detection

increase the threshold to focus on “real” edges

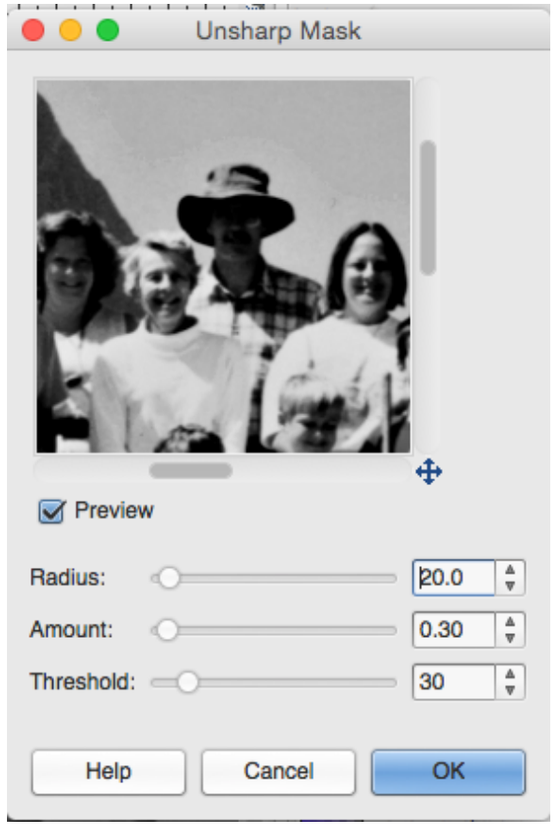

Radius

is the width in pixels of the effect on either side of an edge

Unsharp Masking—What's Happening

Before: RGB = (100,100,100) on the left & (200,200,200) on the right

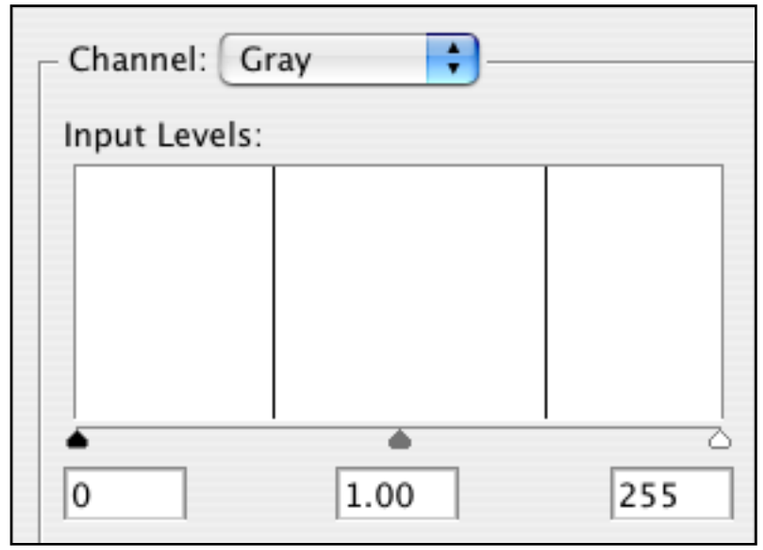

Pixel Histograms

After sharpening with amount = 100, radius = 2.0, threshold = 0

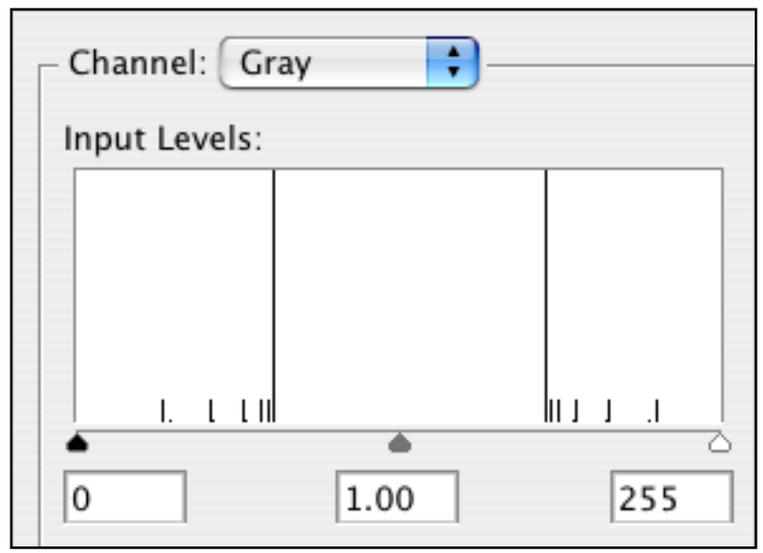

The contrast immediately to left and right of the boundary
has been increased—when not enlarged, the eye sees this as a sharper edge

What About Colour?

Apply Level and Sharpening adjustments

to the R, G and B channels simultaneously (RGB)

or individually

The Gimp “Curves” tool provides for very flexible adjustment of

levels

colour balance

contrast

Unfortunately it’s use is a bit beyond the scope of CS200.

For more, see:
http://www.cambridgeincolour.com/tutorials/photoshop-curves.html
(strongly recommended...)

Transformations specifically for colour:

Colour Balance

Hue / Saturation / Lightness

Brightness / Contrast

Colours > Curves...

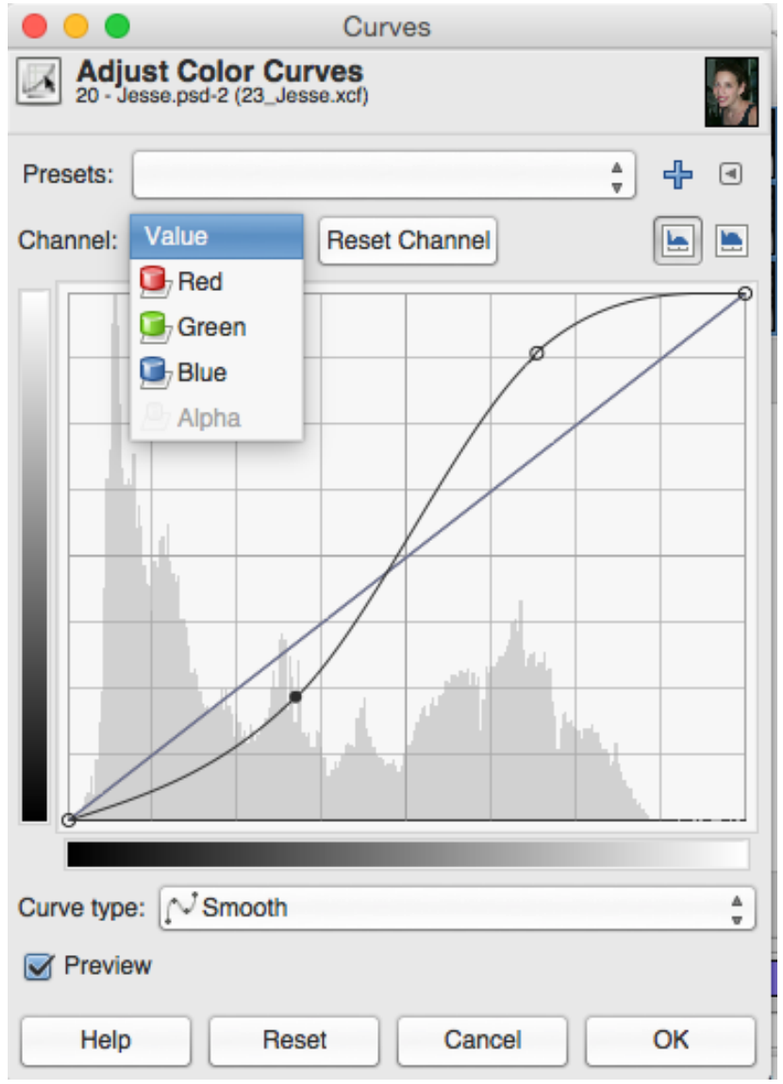

Hue, Saturation and Lightness (aka Brightness, Luminance)

Colours > Hue/Saturation...

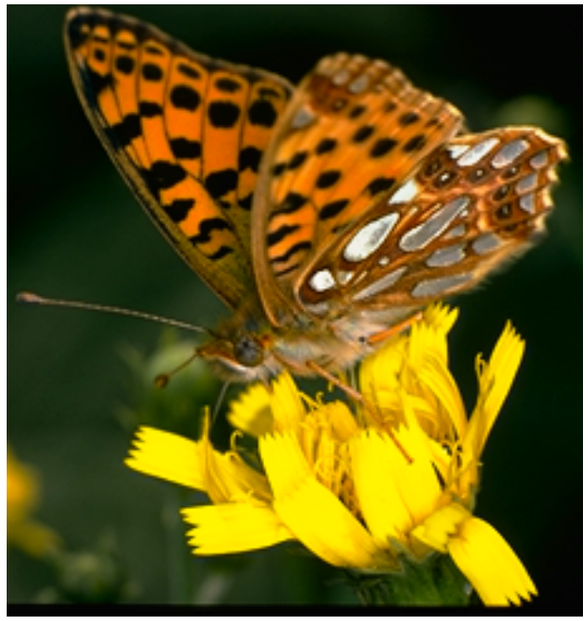

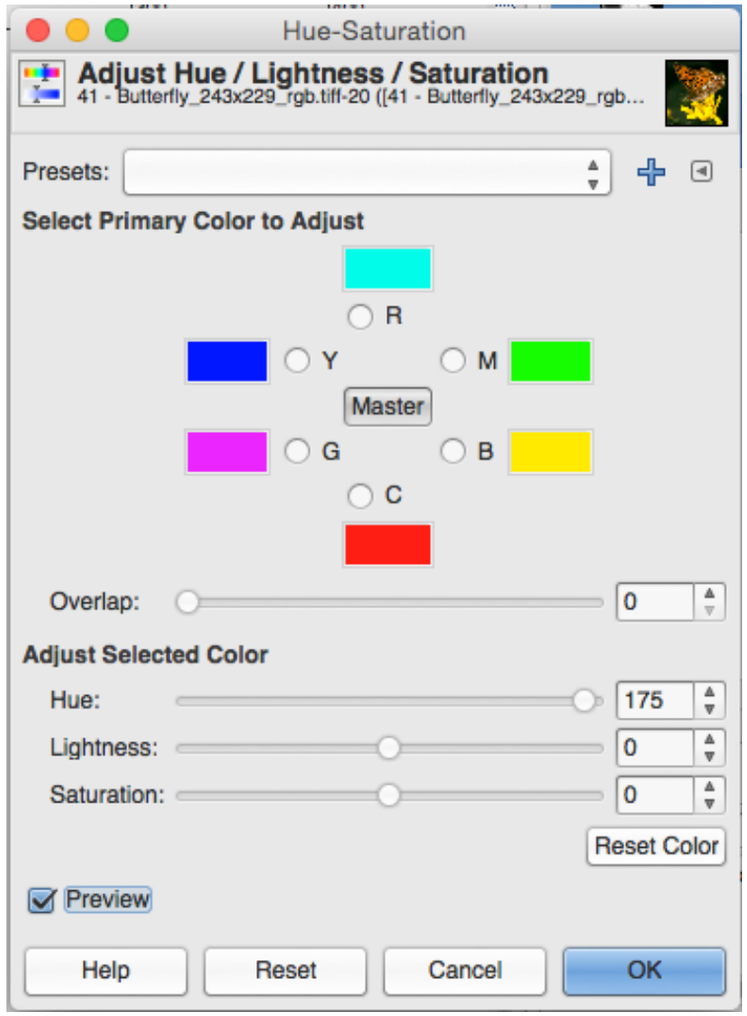

The Hue slider rotates colour around the colour wheel

Saturation

increasing saturation makes colours more vivid

fully desaturated colours are gray

Increasing lightness lightens the entire image

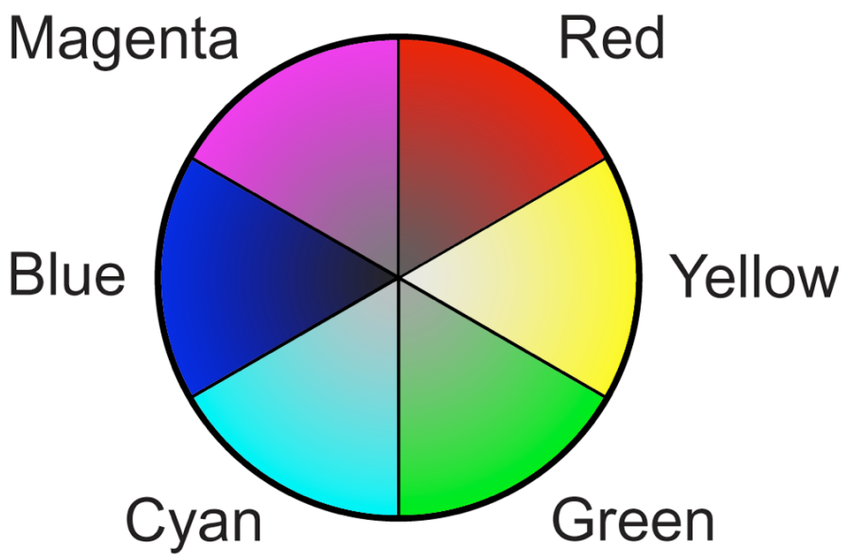

Brightness-Contrast

Colours > Brightness/Contrast...

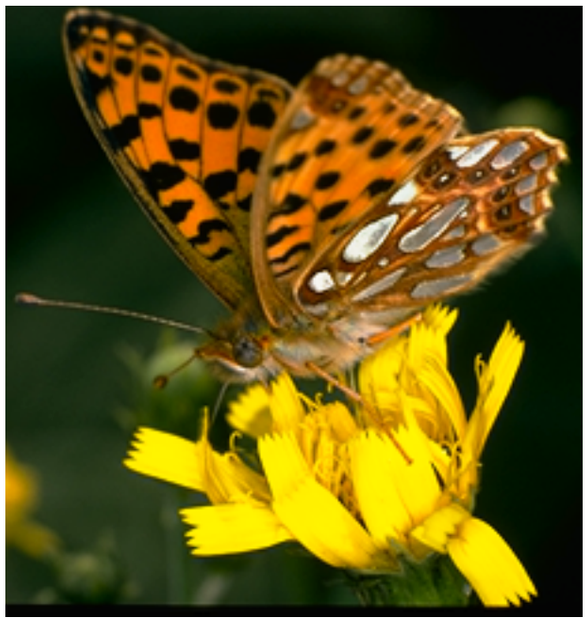

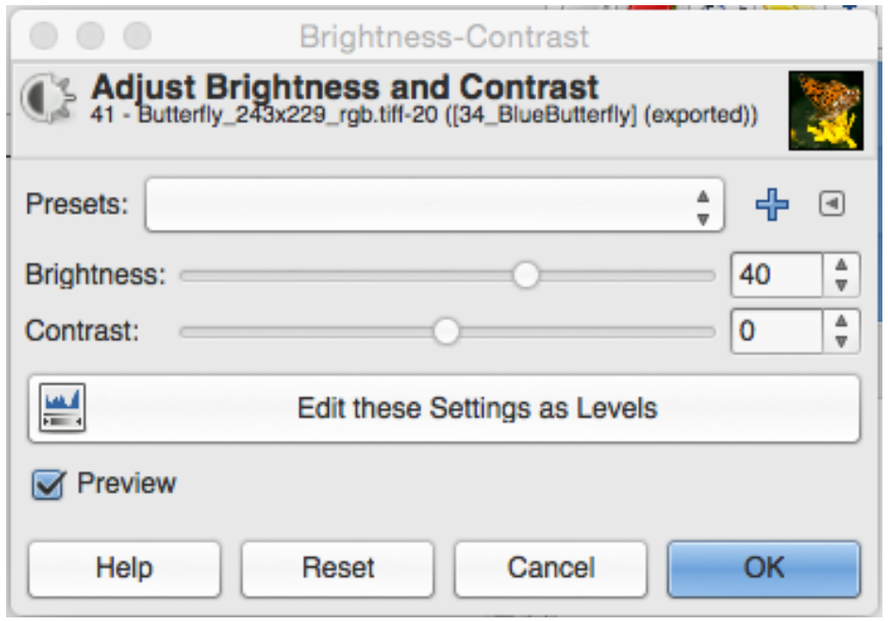

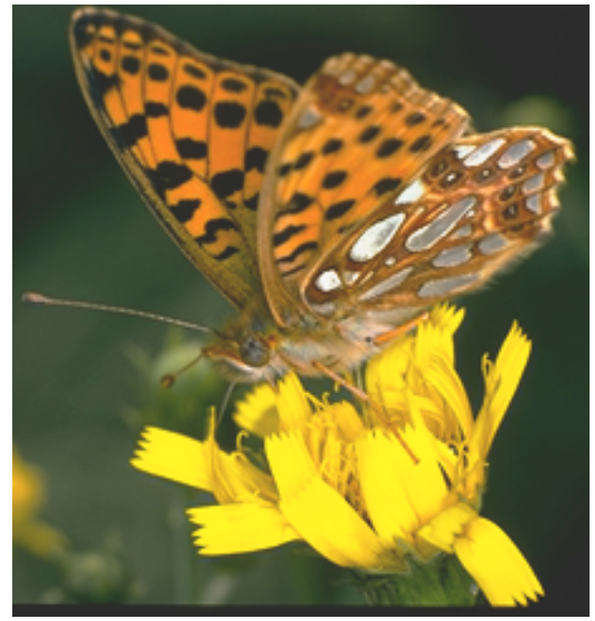

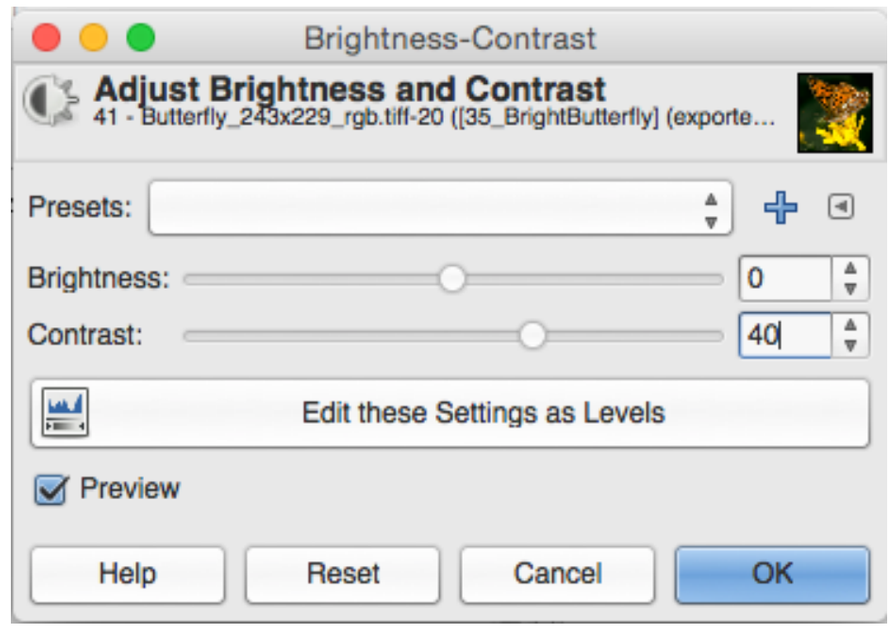

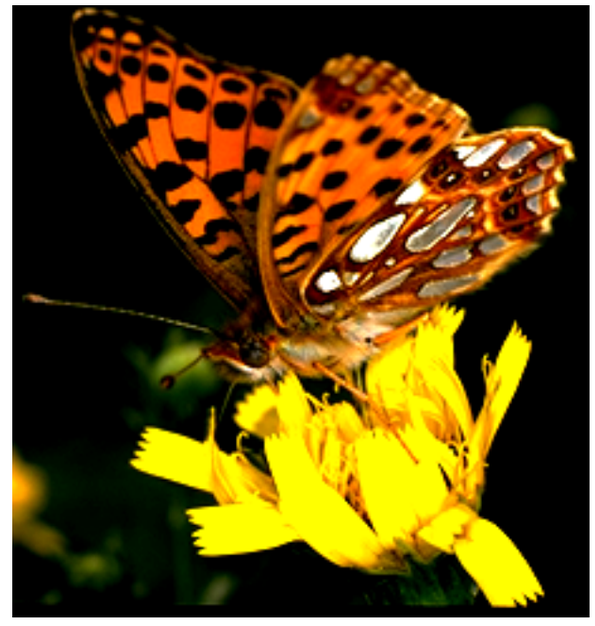

Applying Image Transformations to a Selected Region

Select a Marquee Tool

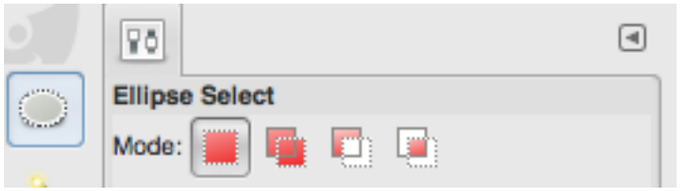

Use the tool to select a part of the image & apply an image transformation

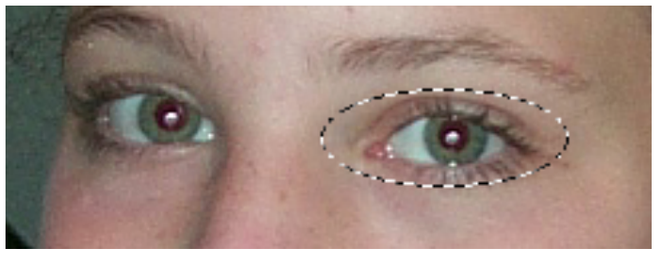

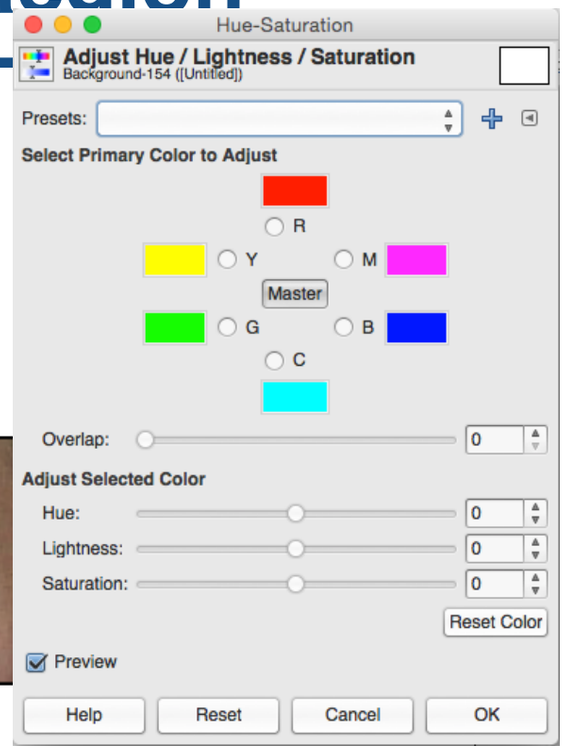

Feather the selection (remember feathering?) and apply the same image transformation

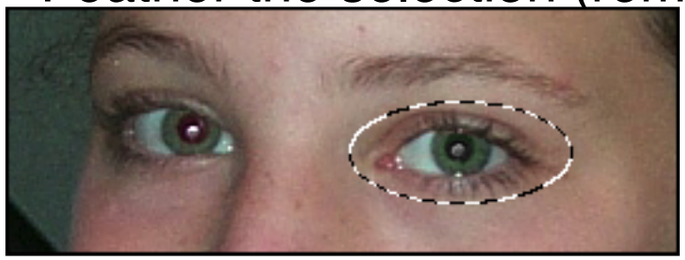

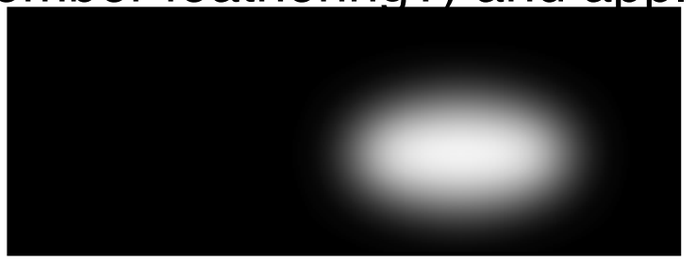

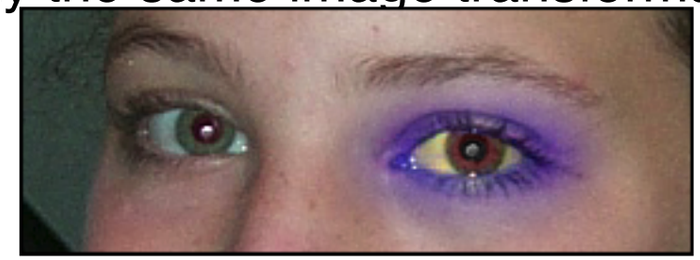

pixels inside a selection can be modified

pixels outside the selection are “masked” (remember masking?)

“Feathering” results in a gradual transition from full to no effect as you move outwards across the selection boundary

We’ve seen this before, when filling selections

Hint: Use Select > Feathering... to increase the feathering for an existing selection.

The Layers Palette (Another Data Model)

Each "layer" is a pixel grid (effectively, an individual & distinct painting)

stacking order determines visibility (painting order)
— the net effect is built up by painting the layers in order from bottom to top —
though objects may be partially transparent

each layer is ONLY a pixel grid (compare with the lecture on Geometric Graphics)

click-drag in the Layers palette to change a layer's stacking (ie pa

when you draw on the canvas, ink goes into (onto?) the selected

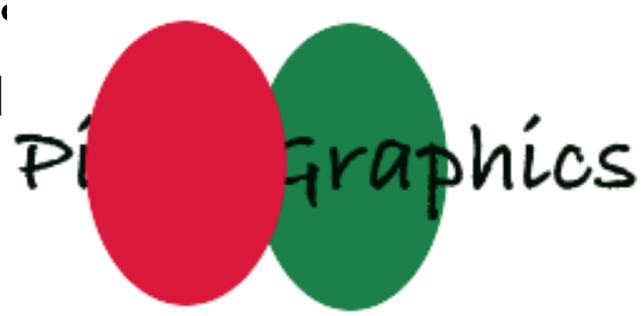

you can control whether a layer is visible

you can "lock" a layer so that it cannot be modified

click-drag the ink in an object to move the object in x and y

Gimp identifies the object to be moved by whose ink is frontmost

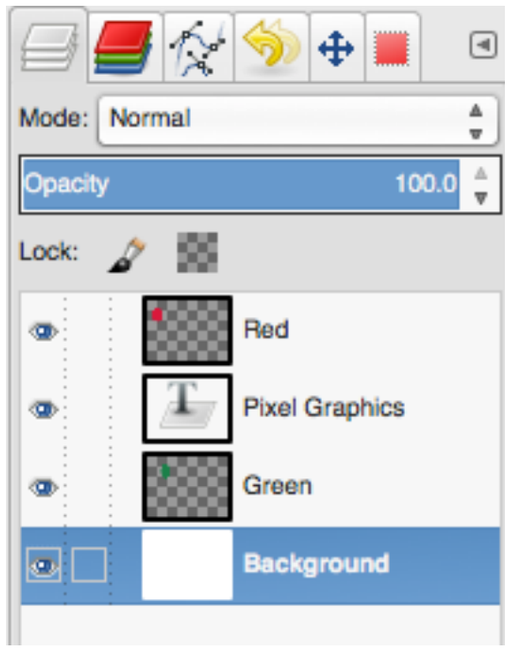

Using Layers to Control the Stacking Order (aka “Z-Depth”)

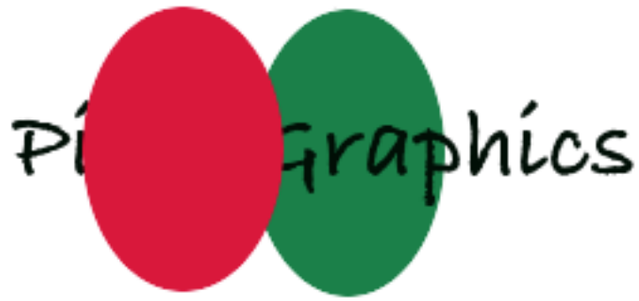

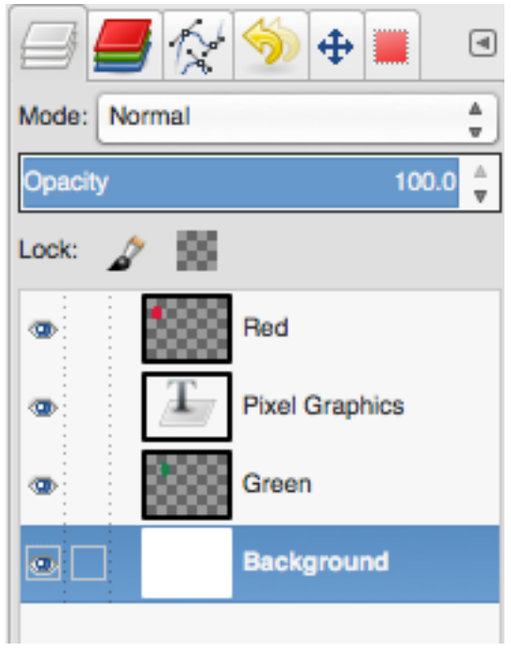

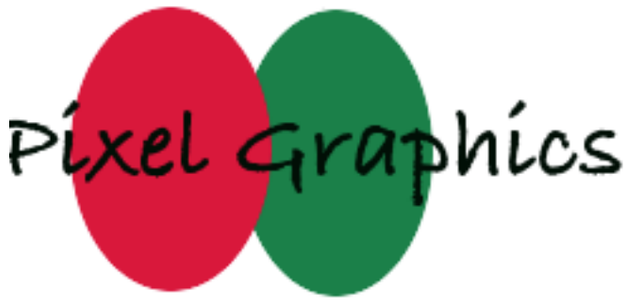

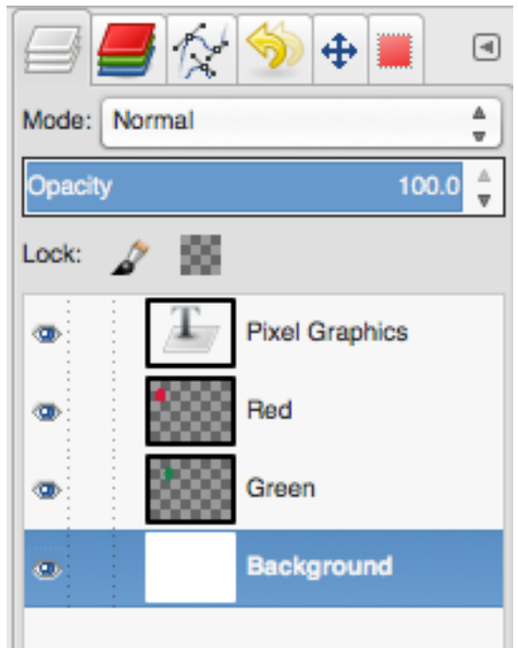

Adjustment Layers

An “adjustment layer” is a “layer” with an attached image transformation

the adjustment may affect the entire image, or only a selected (“masked”)

Warning

- You can apply image manipulations directly to an image via the Colour menu.
- It’s nearly always better to create an adjustment layer.
- Directly adjusting an image alters the image pixels themselves, whereas altering an adjustment layer changes only what’s displayed on the screen. (The adjustment layer is stored separately, and the transformation applied to the image whenever it is rendered.)

- Gimp does not have elegant adjustment layers like Photoshop does.
- Duplicate the previously modified layer to make an additional adjustment. This is described more fully in the video.

Layers, Layer Masks, and Selections

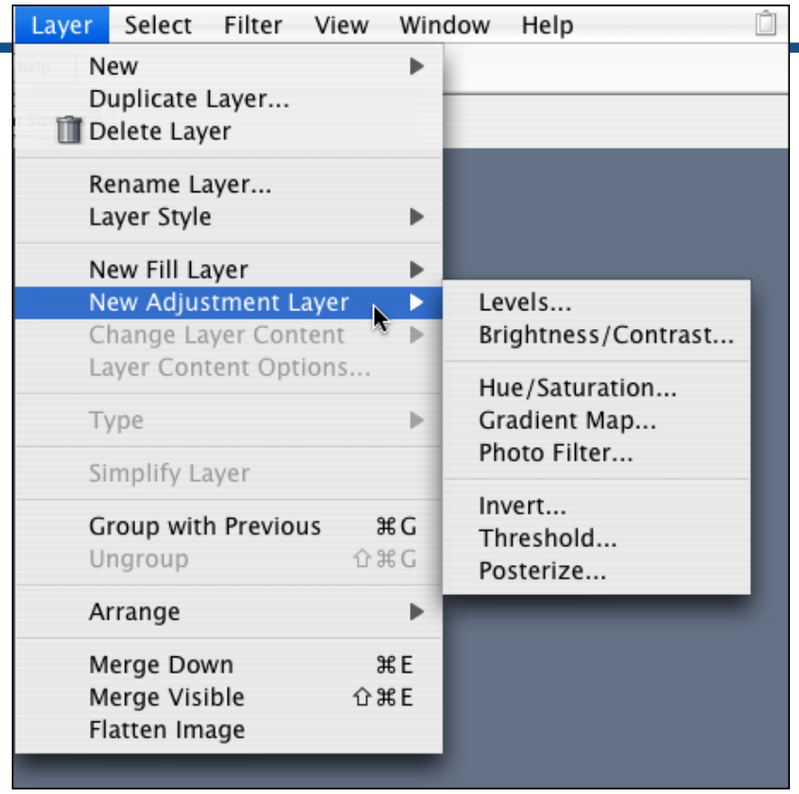

Selections can be converted to "layer masks"

- create a new adjustment layer while something's selected

A layer mask can be converted to a selection

- click on the layer in which you want the selection, then command-click on the layer mask

You can edit a layer mask with all of Photoshop's tools, just as you edit an image

- option-click on the layer mask

You can disable (turn off) a layer mask

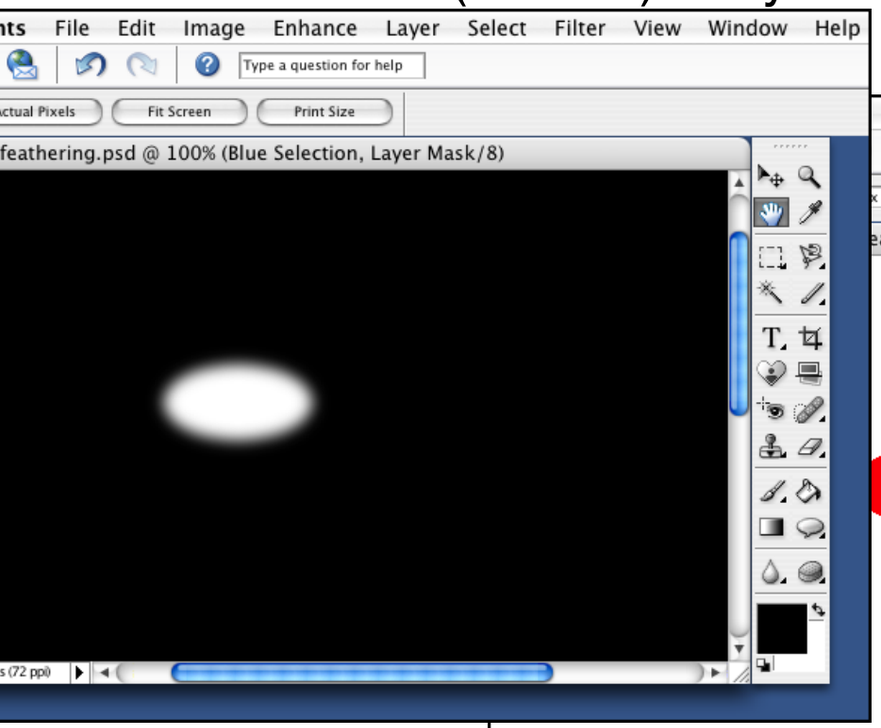

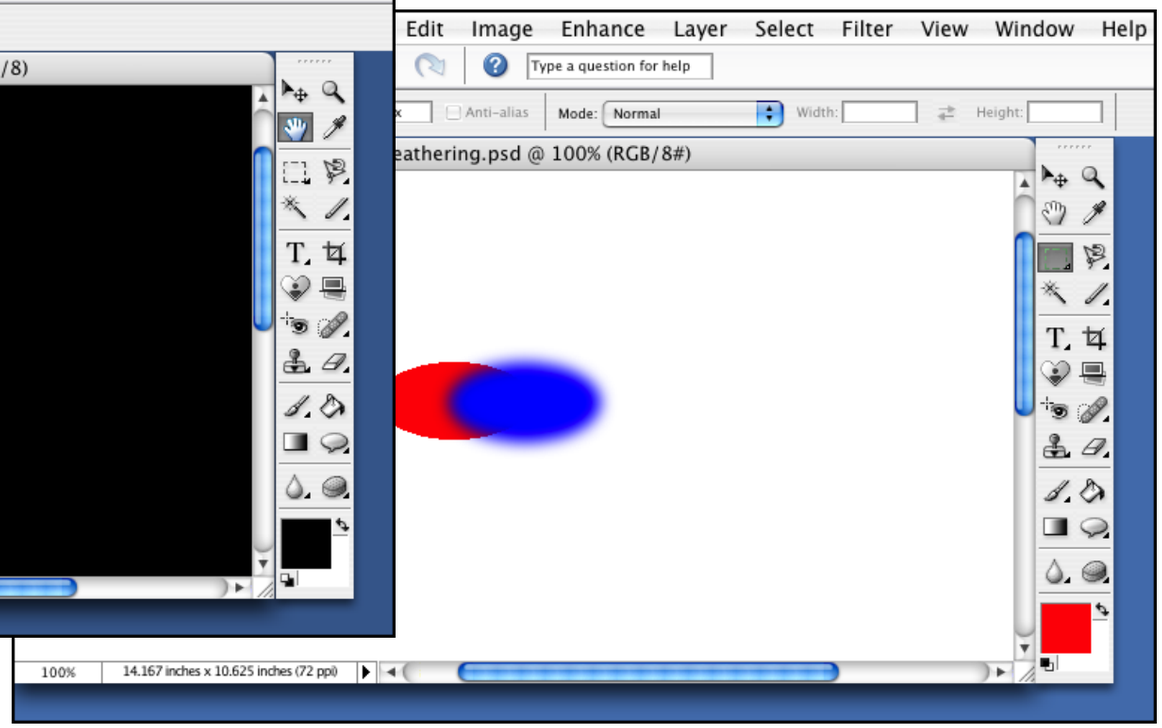

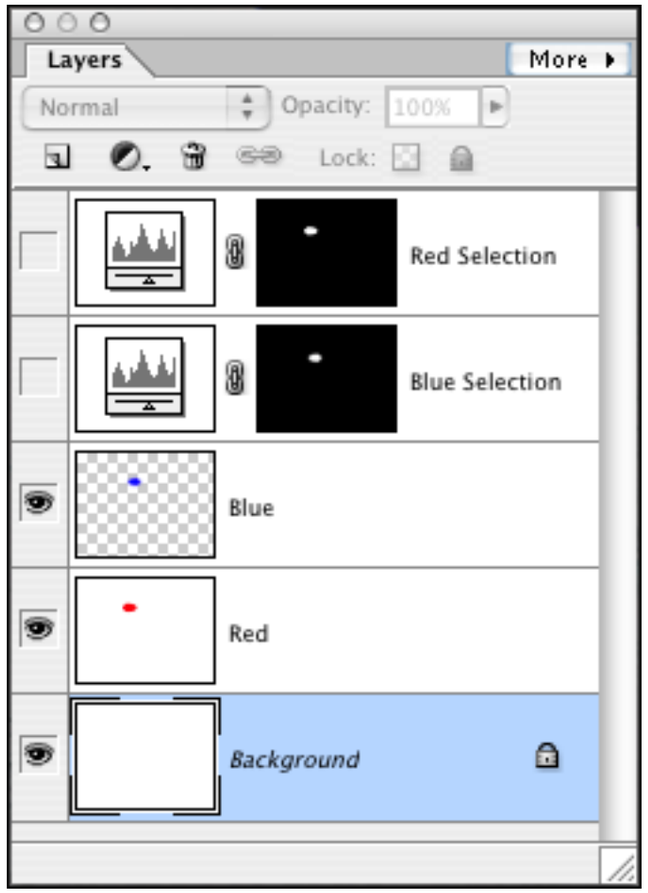

Final Words

Gimp is an *extremely* rich application

- rich both in features
- and in the variety of useful ways in which you can combine features

we've only scratched the surface / given you a skeleton to flesh out on your own

esp wrt making selections and masking

Places to go for more information

Real World Adobe Photoshop CS3 (© 2005)

by David Blatner, Conrad Chavez and Bruce Fraser
PeachPit Press, www.peachpit.com, ISBN 0-321-51868-3

Real World Scanning & Halftones, 3/e (© 2004)

by David Blatner, Conrad Chavez, Glenn Fleishman and Steve Roth
PeachPit Press, www.peachpit.com, ISBN 0-321-24132-0

The Non-Designer's Scan and Print Book (© 1999)

by Sandee Cohen and Robin Williams
Peachpit Press, www.peachpit.com, ISBN 0-201-35394-6

resources > Drop Shadows & Masks on the cws

(Like O'Reilly, Peachpit is a quality publisher whose books are generally recommended.)