

“Translation, not Manipulation”

Glen MacDonald
Digital Microscopy Center

Core for Communication Research
Virginia Merrill Bloedel Hearing Research Center
and
Neuroscience Core
Center on Human Development and Disability
University of Washington
Seattle, WA 98195-7923, USA

Is there an underlying principle to image processing?

Why do we need to do anything to the images we collect? After all, considerable time and money have been devoted to optics, detectors, sample preparation and attention to details. What could be wrong with our images?

Chances are, there is nothing inherently wrong with your images. The concept of image processing is best approached as an act of translation, regardless of whether the actions are for quantification or presentation. In fact, image processing is the last of several translations that occur between the illumination source and the final display. Incident light interacts with the sample and is translated into a diffraction pattern at the back focal plane of the objective. This diffraction pattern is translated into a real image, into analog electrical signals then to digital encodings and, eventually, translated through graphics processors, into electronic display systems or hard copy. Image processing is that last step where the image is translated into an output suitable for machine recognition or the human visual system. We take great care to minimize degradation of our images by the biases, noise, limits and losses imposed on the source image signal during this long series of optical and electronic translations. Why waste our investments by careless actions at the last step?

Unfortunately, image processing is often approached as arbitrary actions attempting to make an image look "better" (or more closely resemble a preferred reality). At best, this type of manipulation gives digital images an unsavory reputation and, at worst, may produce images subject to misinterpretation or that even become misleading.

Differences between the electronic detector and the human eye.

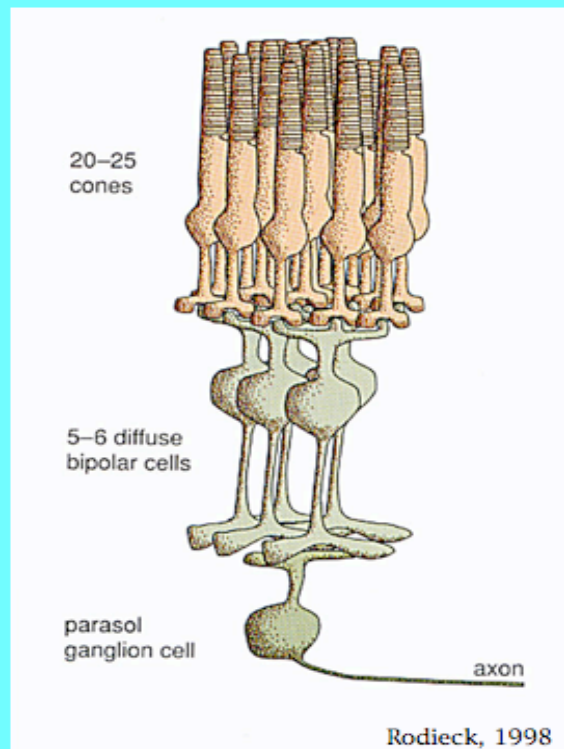
The intensity recorded by the detector for any point in the field depends upon the number of photons absorbed from that point. Detectors such as CCD arrays and photomultiplier tubes are usually operated in such a manner as to provide a linear response to changes in the number of absorbed photons, and such that the photons recorded for each sampled point reflects solely the signal from that point and not its neighbors. The range of intensities from a field recorded by CCD and PMT detectors are generally divided into scales ranging from 8-bits (256 intensity levels) to 16-bits (65,536 levels).

Unlike imaging detectors, our photoreceptors are arranged in overlapping functional groups of 4-5 cones innervating a diffuse bipolar cell. Visual fields from 5 to 6 of these bipolar cells feed into a single parasol ganglion cell. A diffuse bipolar cell fires maximally when one cone in its group is stimulated and the others are silent. Similarly the parasol ganglion cell is maximally active if only one of its bipolar cells fires. If intensities become uniform across the receptive field formed by the cluster of

cones, the ganglion cells reduce their firing rate. Of course, there are many levels of regulation within the retina that modify its response to the visual field. But, this very basic description is sufficient for discussion.

Center-Surround Principle

- Cones are functionally arranged in clusters
- Each cone's signal is integrated with others in its receptive field
- The output from a receptive field is proportional to the difference between the central cone and its surrounding neighbors
- We perceive *differences!*



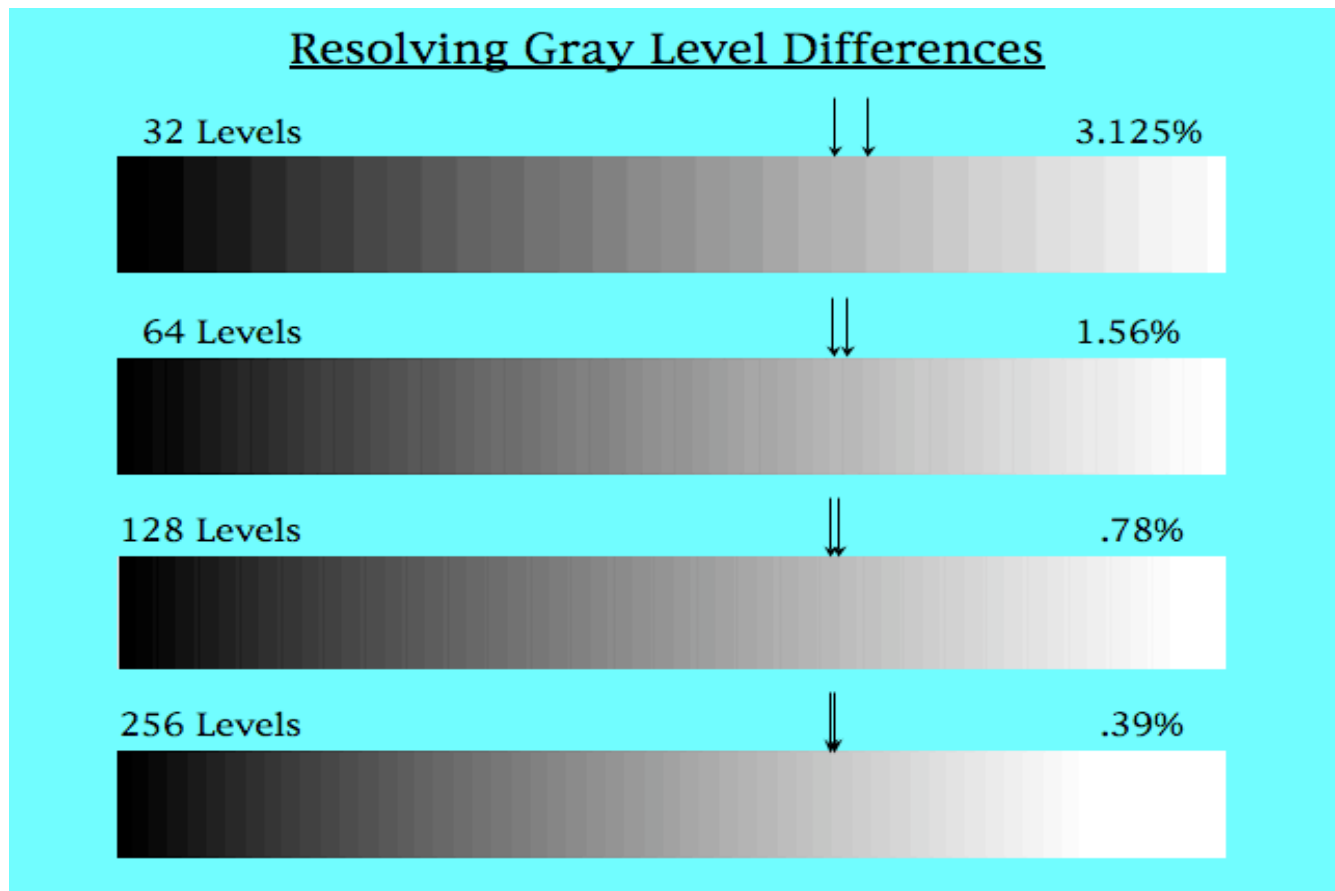
The result is that the output from the receptive field is proportional to differences between adjacent cones. This weights the visual response towards differences in the intensity gradients and boundaries in our visual field. Our visual system is well adapted to recognize patterns and movement, but, performs poorly as a spectrometer.

Again, processing digital images for visualization is not merely a matter of “manipulating” an image until it looks “good” or until it is consistent with your hypothesis. Image processing is a process of “translation” to convert it from the digital detection system into the human visual (detection) system.

Perception of intensity.

Although the human eye responds to a range of 10^9 intensity levels, this will only occur when viewing different fields. It can only distinguish about 100 intensity levels within any given field of view. Furthermore, the perceived intensities are relative to the setting, thus our perception is not absolute. Again, the eye is a poor spectrometer.

We require a degree of difference before we can perceive adjacent intensities. Consider a series of grayscale spectra at different bit-depths, all containing the range of intensities from white to black. In an 8-bit scale, with 256 intensity intervals, each step in the scale represents 0.39% of 256. Scales divided into 128 (7-bit), 64 (6-bit) and 32 (5-bit) shades present proportionally greater differences between adjacent intensities, 0.78%, 1.56%, and 3.12%, respectively.



The difference between each intensity step is readily observed up to the brightest and darkest pairs in the 5-bit scale (32 intensities). At 64 levels, the adjacent levels are readily observed, but the brightest and darkest are more difficult to discern from their neighbors. The 128 levels scale appears nearly uniform, but some banding is visible in the middle gray values with ambiguity at the extremes.

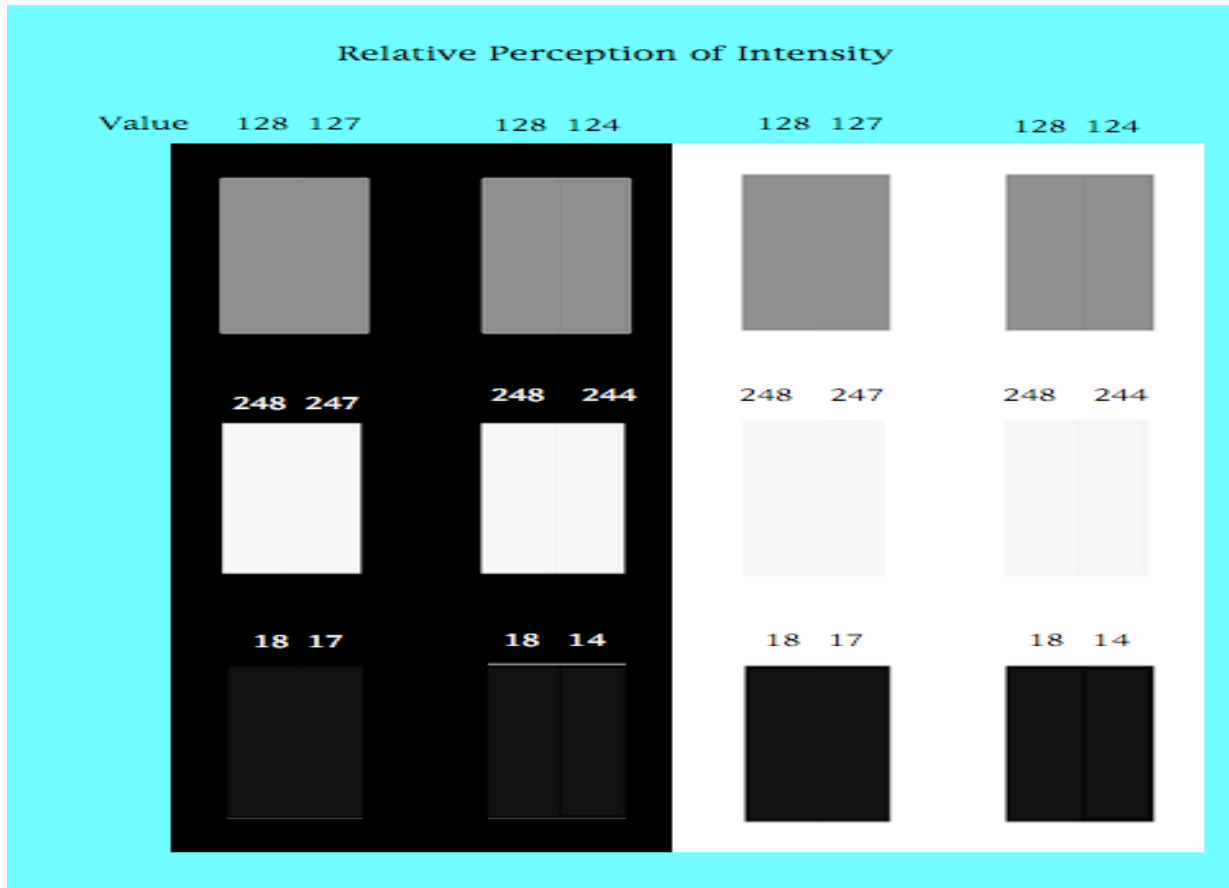
An 8-bit scale appears homogeneous, without any perceptible banding. EXCEPT, some banding does occur and will be moving. What is happening? Even though the intensity difference between adjacent intensity panels is less than we can resolve, the retina is perceiving intensity differences along the scale as best as it can. What happens once the difference is noted? If nothing changes, the lack of differential stimulation would make it appear as if no differences exist! We prevent that by “gaze shift”. Our eyes are continually moving and thus register new differences in the visual field. Even though adjacent values in an 8-bit scale are below the limits of discrimination, the differences are being detected wherever possible and are constantly updated by gaze shift. A caveat here is that printing this document may create some banding because the printer’s dynamic range may not reproduce all intensities.

What would happen if you stared at the 64 Levels scale without shifting your gaze?

The retina as a densitometer.

Perception of intensity is relative and depends upon context. Any region in an image will appear darker if it is radiating or reflecting fewer photons than its surroundings, and will appear brighter if the eye is registering more photons from that region than from its surroundings. This is an important consideration in presentation and interpretation of data. Consider the case of panels representing the

bright, mid-range and dark regions of intensity. They will appear relatively brighter or darker depending upon whether they are viewed on a white background or a black background.



Combining our ability to detect edges with relative perception of intensity differences can lead to illusions mistaken for structure. The next figure shows a grayscale spectrum with 32 intensity levels. Notice that each gray level panel appears to have a gradient with a lighter shade to the left edge of each panel that progresses to a lighter shade at the right edge of each panel. Furthermore, we will perceive a black line to the left of each intensity boundary and a white line paralleling the right side of each boundary.




However, if we measure the intensities of the panels across the panel (red line), an intensity profile plot demonstrates that each panel is filled with a uniform intensity level. If the plot does not show gradients or boundary lines, why do we see them?




Since our perception of intensity is relative, we perceive the intensity of each panel to be different as our vision approaches the left or right boundary and perceive intensity as relative to the adjacent panel. The difference between panels is perceived to be greatest at the boundaries.

Measured Intensities



Perceived Intensities



We can mimic our perception of the grayscale by performing an unsharp mask operation on the grayscale before plotting. In this operation, a copy of the image is blurred by a gaussian filter then the original image is subtracted from the blurred copy. The unsharp mask operation emphasizes boundaries, which makes it commonly used to increase image sharpness.

Grayscale after unsharp masking



Can we know whether a boundary really exists based upon purely visual perception?
If the unsharp mask filter emphasizes edges at arbitrary intensity gradients, can we really trust structure boundaries in the resulting image?

Persistence.

Persistence is the property of each field being retained by the visual system for a brief period of time, about 30 ms. A series of related views presented in rapid sequence will create a sensation of seamless motion due to the property of persistence.


Our perception of motion stems from the degree of difference between views, not velocity or frame rate. Slowly moving objects are easily missed, while something moving quickly, or suddenly changing position, readily attracts our attention.

Depth Perception.

There are several cues for depth, such as parallax (perspective of size and distortion), intensity differences (shadings), overlap of structures and expectations of shape or distribution. But, sequential presentation of images representing changes in viewing angle can also deliver the illusion of depth when the property of persistence allows the overlap of different viewing angles to recreate the stereoscopic information obtained by binocular vision. This, perhaps, is the only useful application of rotating maximum intensity projections.

Color Sensitivity.


Yet another facet of the relativistic nature of the human visual system is that we do not perceive all colors with equal quantum efficiency. The human retina is most sensitive to green wavelengths, and possesses about 20% less sensitivity in the red spectrum. The range of wavelengths described as “blue” is perceived with only 20% the efficiency of “green” wavelengths.




Even worse, there are no blue receptors in the fovea, the portion of the retina with highest visual spatial resolution. What effect might these differences of color perception and color acuity have upon our interpretation of images?

The detector channels in digital image capture devices are often displayed in pseudocolor during acquisition as red, green and blue in individual image panels on the instrument. Although esthetically pleasing, relying on pseudocolored images to make acquisition decisions or interpretations of images can lead to serious errors.

The figure below presents a cell labeled for DNA with DAPI, which fluoresces “blue”, labeled for actin by phalloid conjugated to Bodipy-FL, which fluoresces “green” and mitochondria are label “red” by Mitotracker Red. The channel for each label is displayed in grayscale, with the color merged image.




Now, let's view the mitochondria channel in different pseudocolor channels. The mitochondria image has been inserted into each color channel without any modification. What happened to the intensity levels in the "blue" channel? Blue wavelengths are perceived as being dim and of lower resolution than the other channels.



If you only view the sample image with false color during acquisition, might this affect decisions regarding acquisition conditions? How might you be misled to change the labeling?

What might you need to consider when designing an experiment? The choice of fluorophore emission may greatly affect the ability to present structures that are small or few in number. Mixing 3 channels into different false-color combinations has dramatic effects on the visibility of labeled components. This is one factor of contrast that requires no manipulation of pixel intensities.

Compare the ability to recognize individual mitochondria in these two images. The only difference is that the left image has the mitochondria displayed in the red channel and the right image displays the mitochondria in the blue channel.



Summary of the human visual system.


- Responds to intensity differences
- The differences response emphasizes edges
- Perceived intensity response is non-linear and relative
- Contrast requires a minimum difference between adjacent intensities
- Intra-scene dynamic range is about 100 intensity levels.
- Persistence provides sense of motion and 3D
- Differential color perception affects contrast and sensitivity
- Visual CNS filters for, and projects, patterns upon the signal

The role of the higher visual centers in our ability to recognize shapes and patterns is a discussion outside the scope of this course. However, it is useful to think of the higher visual centers as a programmable filter. If you are looking for a particular shape in the midst of many other shapes, you will find it quickly since your visual system has been “tuned” to recognize it. If you then decide to find a different pattern, the filter is re-set to find that pattern. You must be aware of this and take care not to interpret structure from what may turn out to be a random array of intensities.

Now, let’s introduce the basic terms and operations used in image processing...

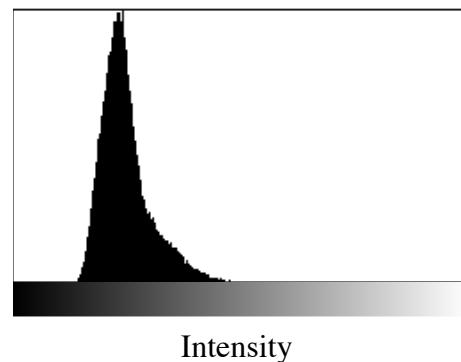
Histogram.

The histogram should be your best friend. It can display which levels contain the information in the image, suggest an approach to image processing and then display the effect of each processing step. It is simply a bar graph that displays the number of pixels possessing each intensity along the y-axis and the intensity range from 0 to maximum value along the x-axis. A small image is shown in the figure below, with its histogram. The graph shows that the intensities consist of dark gray to medium gray intensities. Most pixels are darker, which is typical of a fluorescence image, where background is generally dark and only a few pixels possess brighter intensities representing label. In this image, no pixels are truly black (0 intensity) or completely white (nearing maximum).




Histogram

Pixel Number



Histogram Stretch

A histogram stretch remaps the pixel intensities to re-distribute them across the full span of the available levels in the histogram. This accomplishes two functions; 1) adjacent intensities are given a greater degree of difference (contrast); 2) the bright and 'not so bright' intensities of signal are shifted to brighter intensities, making the labeled structures become more visible. The tools for histogram stretch have different names in various applications. In Photoshop, the 'Levels' dialog stretches the histogram, while ImageJ uses the "Adjust" tools and other applications may call it "renormalization". Two important points to note: 1) the relationships between gray levels are preserved; 2) It can be performed without losing levels or creating interpolated levels.






Contrast.

Contrast allows structures in an image to stand out and be noticed. It is the ratio of the brightest pixels in an image to the darkest pixels, or as the ratio between a region of interest and the image background. Note that the minimum signal is that signal above the noise floor. In other words, noise will reduce the contrast of an image by raising the noise floor to “drown” signal from label. Noise also reduces the number of available levels for depicting information.

$$C = \frac{I_{\max} - I_{\min}}{I_{\min}}$$

, where the minimum intensity is minimum signal above the noise floor.

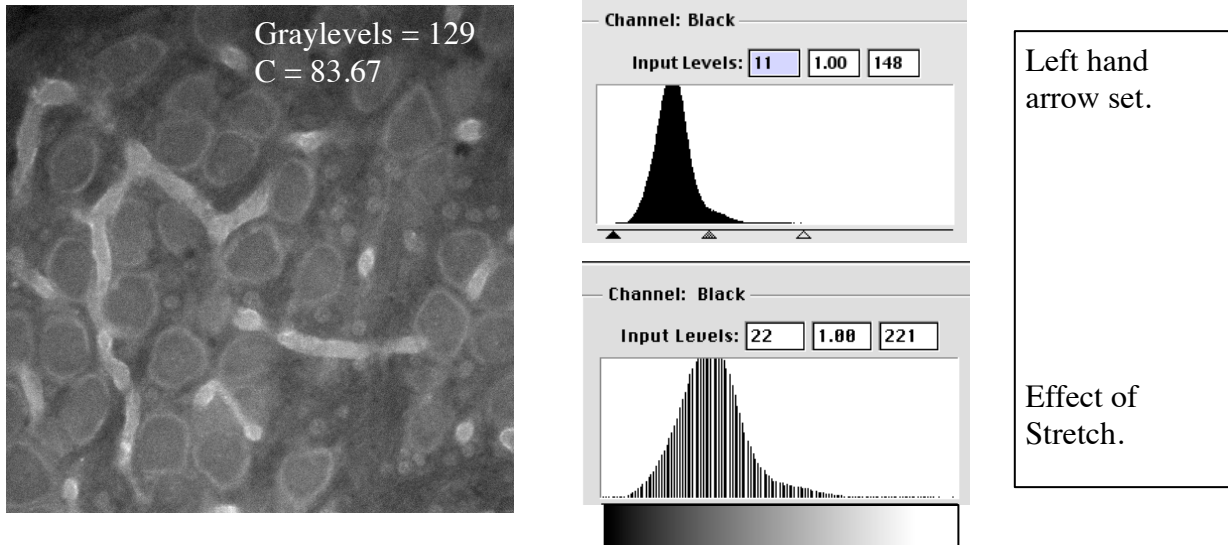
Here is the entire image. Note the overall lack of contrast and the small distribution of intensities in the initial histogram. The histogram is from the Photoshop Levels tool, with the contrast ratio, C, and the number of gray levels, GL. Notice that the histogram spans from nearly black to about the middle of the graph. The Input Levels shows that the image was coded for 0 to 255, and that the middle (gamma) is set to 1.



Original Levels Dialog

Right hand limit set.

The right hand and left hand arrowheads designate the limits of the histogram. If the arrowheads are moved to the ends of the curves, then those points of the curve will be mapped to the limits of the histogram range, as shown next.

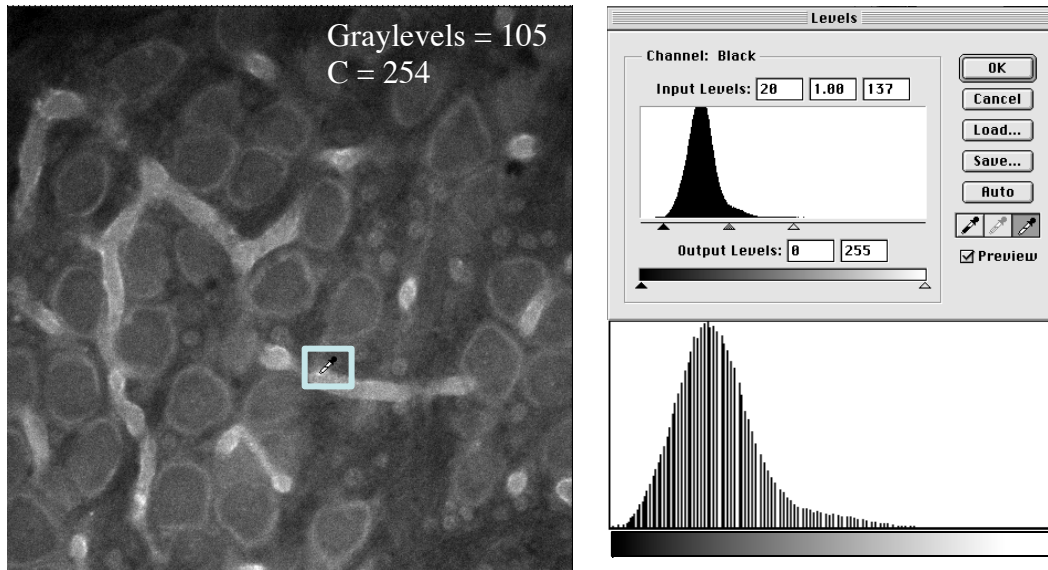


So, now we will increase the contrast by stretching the histogram. The initial contrast ratio was quite low. After stretching the bright levels we have neither lost any gray levels nor created any. But, the contrast ratio has greatly increased from 1.22 to nearly 84.

Stretching the left side of the histogram provides additional contrast and deepens the low intensity background. But still, we have not altered the number of gray levels.

Note the values in the Input Levels boxes. Any pixel with an intensity value less than the left hand value (controlled by the left hand triangle) will be set to zero. Similarly, pixels with intensities greater than the value of the right hand input level will be set to maximum and constitute saturation. Since multiple pixels are being set to the same value in either instance, the number of intensities are reduced.

We can semi-automatically stretch the histogram by using the 'Eyedropper' tools in the Photoshop dialog. These tools are placed over the point in the image you wish to designate as the darkest or brightest points (structures). The software will then perform the histogram stretch using the values from these tools. The black eyedropper, of course, sets the blackest levels and the white eyedropper sets the brightest point. In RGB images, the gray eyedropper tool becomes active, as we will see later.



In this example, the contrast was enhanced to the maximum for an 8-bit image, but that we have lost about 20% of our gray levels. Automatic adjustments of contrast tend to be overly aggressive and should be avoided.

Gamma.

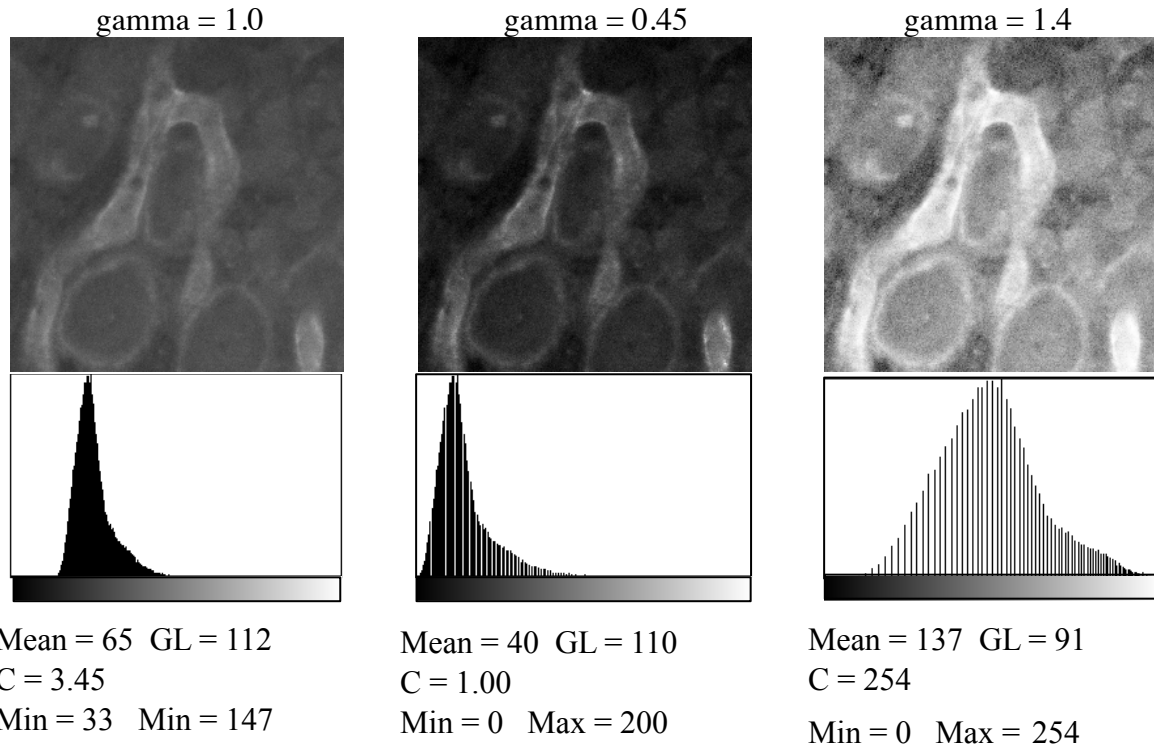
The relationship between the input signal and the output signal at any stage of the imaging process is described by an exponential function called “gamma”. Gamma is a tricky term, as it is applied at many stages of the image acquisition, display and printing processes as a by-product of limitations in our optics, electronics and display media. Gamma alters the middle ranges without saturating the pixels with the brightest or darkest intensities. **GAMMA CHANGES THE RELATIONSHIP BETWEEN BRIGHT AND DARK REGIONS!**

The degree of gamma is noted in the Photoshop Levels dialog as the middle value above the histogram in the Levels window, and is controlled by the middle triangle below the histogram.

Where gamma = 1, the input-output relationship is linear. Although we try to maintain a gamma = 1 through our imaging systems, the reality is that several different gamma functions may be applied to our signals, even choice of fluorescence filters may impose a gamma function. These factors may cancel each other, or they may accentuate each other. The act of calibrating a computer display usually involves changing its gamma. Printer inks, choice of paper and other printer settings and video projectors all have unique effects upon the gamma

Gamma < 1 means decreasing sensitivity with increasing intensity. Thus, the output signal from some tube cameras increased very quickly, once a threshold level was surpassed. CRT displays, on the other hand, require an exponentially increasing amount of voltage to increase their brightness as the input level increases. Note that gamma functions are employed as inverse functions, thus these curves may appear inverted in some examples.

Mindlessly changing the gamma, as well as stretching the histogram, can accentuate or depress background staining, autofluorescence, bleedthrough and other sources of unrecognized signals. The result is that you can create artificial results!

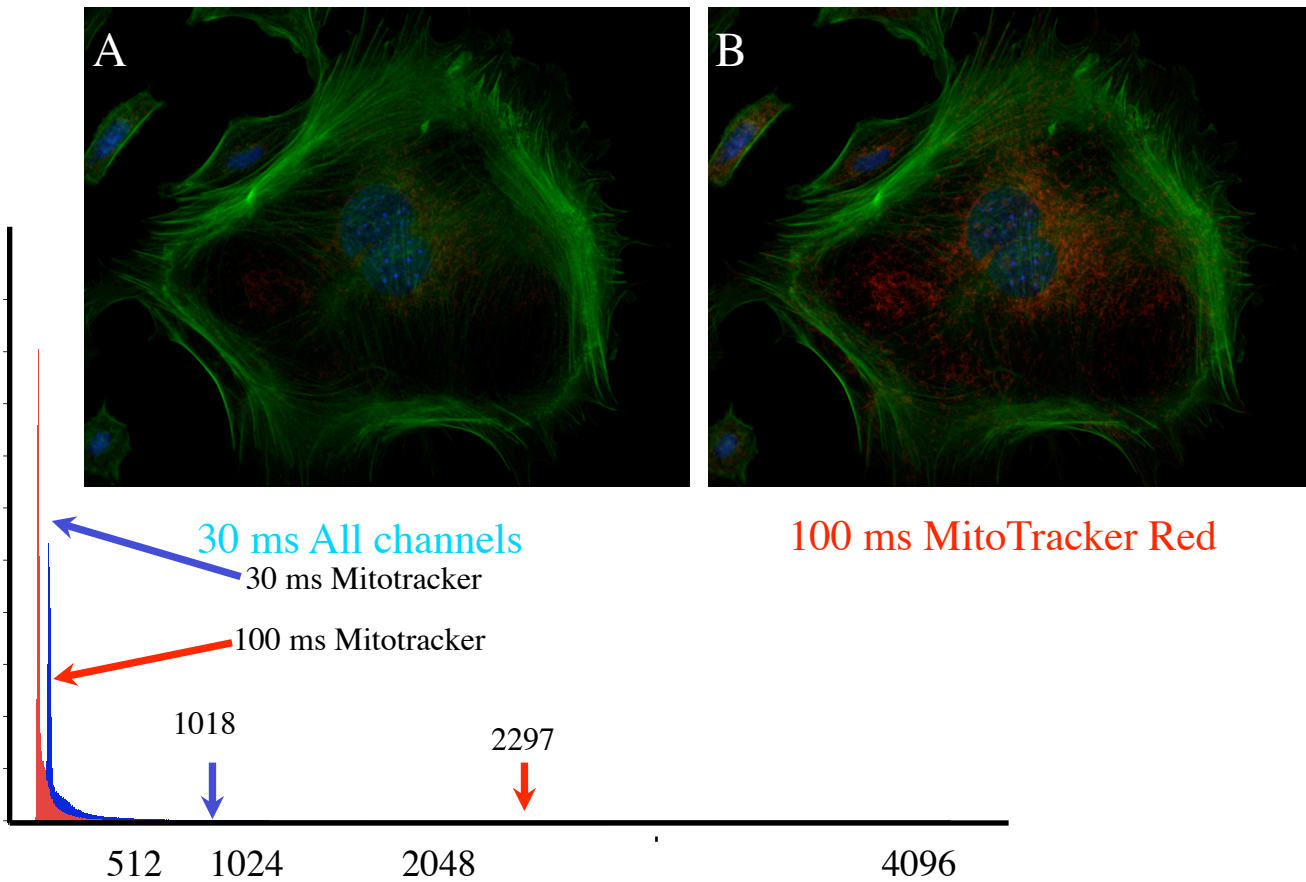


The left hand panel shows a portion of the actin labeled brain slice with no change in gamma. The middle and right panels show the effects of applying degrees of gamma in the histograms and in the images. The respective histograms are in the lower panels, with the statistics below. Applying gamma = .45 gives statistics of moderate increase in mean brightness and contrast by suppressing the darker intensities while leaving the brighter levels relatively unaffected. Notice how the darker intensities “pile up” while the brighter pixels at the left edge of the histogram have been slightly separated. A gamma > 1 dramatically increased the overall brightness by moving the entire histogram. Note that the bright pixels are “piled up” or compressed. Meanwhile, the background became more pronounced, the distribution of actin in the blood vessels is more visible, but with reduced dynamic range for the brightest values.


Since the RGB image is simply three images to which each is given a red, green or blue color table, we can apply all of the same manipulations to all channels simultaneously or to each channel individually.

Do we need to fill the histogram?


We rarely have enough photons to fully fill the histogram without using egregious amounts of excitation energy that probably causes the fluorophores to photobleach before they can produce enough signal to fill the histogram. Another issue is whether the intent of any image is to display all structures, or to indicate the relative intensities of different structures.



Panel A: Three channels, DAPI, MitoTracker Red and Phalloidin-Alexa 488, were recorded with equal excitation time of 30 ms. The mitochondria are barely visible, and probably would not be visible when printed.
Panel B: Excitation time for MitoTracker Red was increased brightness to 100 ms. This channel is now visible, but still more dim than the actin label. Notice that that red and blue arrows on the combined histograms show that a 3.3-fold increase in exposure resulted in only a 2-fold increase of maximum intensity.




30 ms, Histogram Stretch




100 ms, Histogram Stretch

However, if we stretch the histogram for MitoTracker Red for both the 30 ms exposure (C) and the 100 ms exposure (D), similar visibility is obtained. The shorter exposure is rendered acceptably bright to show up in print. However, is this more misleading than extending the exposure time in terms of demonstrating the actual intensity of labeling? Which creates a “better” image - adjusting the histogram or subjecting the sample to increased photobleaching?



30 ms, Histogram Stretch



30 ms, Histogram Stretch then
Gamma Adjust

Gamma affects relationships.


Notice in panel E that the mitochondrial intensities are greater in the perinuclear region even after histogram stretch is carried out. This could be due to more densely packed mitochondria, a thicker volume around the nuclei, reduced staining intensity in the periphery or simply that the periphery is out of focus. But, the relationships are maintained by the histogram stretch. In contrast, adjusting the gamma increases the weaker intensities without saturating the brightest organelles so that all mitochondria appear equally bright in panel F.

Black Level.


Any image should be acquired under conditions such that the darkest portion shows the noise floor of the instrument. The rise and fall of sea levels with the tide is a useful analogy. Low lying land disappears when the ocean level rises when the tide comes in and then reappears when the tide goes out. If the black level is set too high, low intensities are lost because they are below zero. The black level should be set such that no zero values are recorded in absence of a fluorescent signal, all pixel intensities are slightly above 0. This ensures that no information is “hidden”. The lowest intensities may represent non-specific label, cross-talk, bleedthrough or real and important structures!

This is a linear, additive change affecting all gray levels equally. Relationships are preserved, but values at one end of the histogram may be lost entirely if unduly adjusted.

Other names have been given to this operation, such as “offset” which describes the act of offsetting the signal relative to zero, to make it brighter or dimmer. “Pedestal” refers to the fact that the signal from a camera is lifted above some minimum value. “Brightness” is often found on CRT monitors (and televisions). If we graph the relationship between signal in on the x-axis and signal out on the y-axis, regardless of the gamma relationship, the entire curve is shifted to affect the output levels.



The signal output on a monitor, or hard copy, is equal to the input signal raised to the gamma plus the black level value. The middle graph we see that when the offset shifts the input signal downwards that it may set lower intensity pixels to zero. The right hand graph shows that the offset can raise the input signal so that the output contains no weak intensities and the brighter pixel information may be lost since they are all set to the maximum intensity.




$$I_{out} = I_{in}$$

$$I_{out} = I_{in} + \Delta I$$

$$I_{out} = I_{in} - \Delta I$$

Brightness and Contrast

The Brightness and Contrast levels available in most applications invoke histogram stretch (Contrast) and black level (Brightness). In this example the control is the Histogram Map from the applications NIH Image and Object Image. Contrast is manifested as the rate of change (slope) of the straight-line curve while Black Level is the offset of the curve from left to right. The location of the top and bottom ends of the graph line define the degree of the histogram subjected to the histogram stretch, analogous to the left hand and right hand triangles in the Photoshop Levels dialog.



Original

High Contrast

No Lost Levels

Caveats about your software.

- Know how your acquisition software is configured for histogram stretch and gamma during acquisition and display.
- Default software behavior may apply histogram stretch or gamma to the focus window only or also to the displayed captured image, and you never see the native intensities.
- Software might save images with unaltered intensities in some formats but not in others.
- Relative intensities of the channels may be preserved if the histogram stretch applies the darkest and brightest intensity from among all channels to each channel.

Summary.

- There is an intent to image processing, it is not just a bit of fiddling. The histogram is a very useful tool and should be consulted during acquisition as well as during processing. Unfortunately, it is not accessible in some software.
- Only a few basic tools are needed to cover most visualization needs. Properly applied, they preserve the number of intensity levels, can be easily described and are reproducible. The relationship between bright and dark intensities should be preserved, but may require alteration to compensate for the visualization medium, or to emphasize a particular region of interest or structure.
- Avoid actions that reduce the number of intensity levels in the image. Since all actions taken to an image has a cost of some small amount of information, you should keep your actions to a minimum. If testing an image with many different operations, apply the final actions to a fresh copy of the image once the proper actions have been identified.
- Keep track of your actions in image processing. An increasing number of journals are requiring that all actions applied to article figures be fully listed in the methods section. If you can't explain what you did to the image, why should anyone believe it? The standard for scientific inquiry is that an experiment can be replicated. This is equally applicable to processing the final images.
- Finally, explore a few sample images before committing to any particular image capture condition or image processing regimen for your entire experiment. Small changes at acquisition can have far greater effect than larger efforts in image processing.
- Labeling controls are essential to establish our acquisition parameters for assurance that we are recording signal rather than autofluorescence or bleedthrough. Stretching the histogram or altering the gamma without proper controls can accidentally accentuate background signal or autofluorescence in a manner that can be misleading.

We affect the appearance of our images in many ways without using any of the software tools described here. Even the choice of filters and excitation intensity can affect the gamma, apparent black level, and contrast. We also need to have an idea of intent – whether we are intending to demonstrate presence of labels or their relative intensity.