Image from NASA:
https://commons.wikimedia.org/wiki/File:Earth_Western_Hemisphere_transparent_background.png#filelinks

Remote Sensing

Contrast Stretching

AmericaView
Empowering Earth Observation Education
AmericaView.org - Est. 2003

WV VIEW

The goal of this module is to explore how we can improve the visualization and interpretability of image data by enhancing the contrast.

Module Topics

1. Contrast stretching and lookup tables
2. Radiometric resolution and DN values
3. Contrast stretching methods: standard deviation, Gaussian, linear, percent clip, and historgram equalization
4. Brightness, contrast, and gamma
5. Dynamic range adjustment

These are the topics covered in this module.

Contrast Stretching

- Mapping original DN values to new DN values to take full advantage of the dynamic range and enhance contrast between features

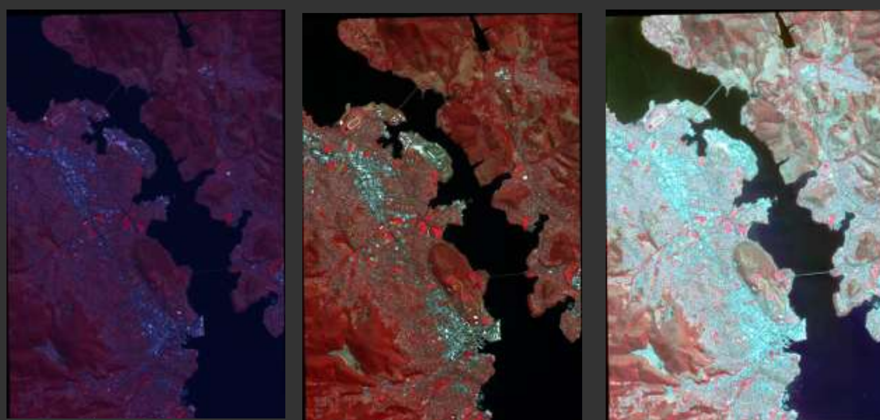

Contrast stretching is the process of mapping the original image DN or brightness values to new DN values to take full advantage of the dynamic range, or brightness levels, and enhance contrast between features. In this set of slides, we will discuss methods to accomplish this.

Pixels and DN Values

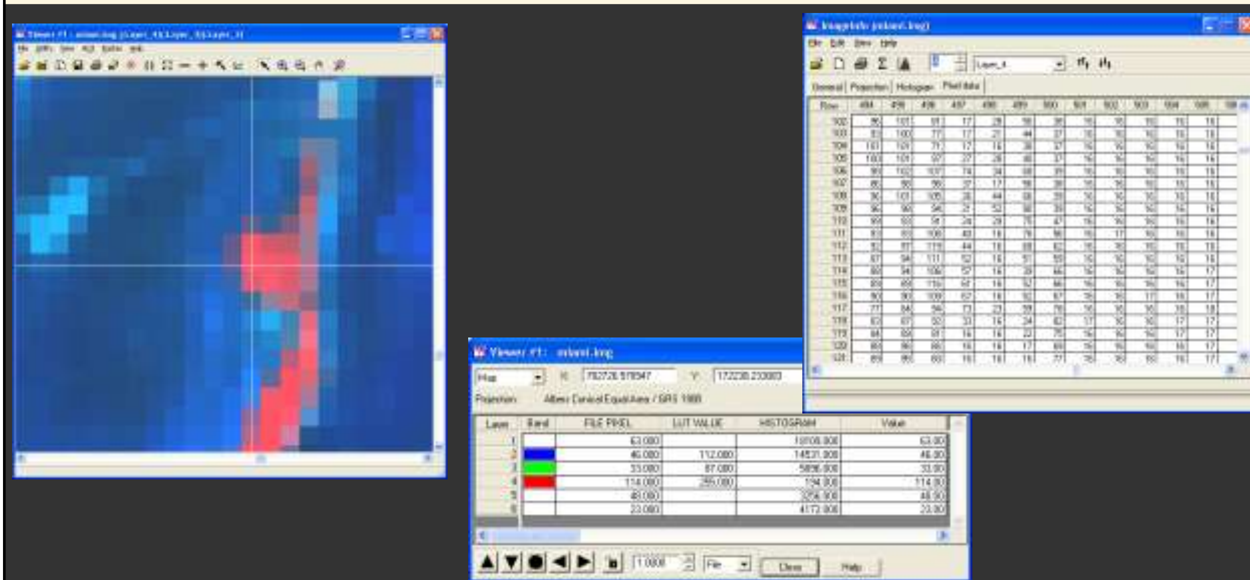

4

Remember that pixels store digital number, or DN, values that represent brightness. Each measured spectral range is stored as a separate channel or band. Essentially, image data are multi-dimensional arrays where the values represent brightness at different locations in different spectral ranges.

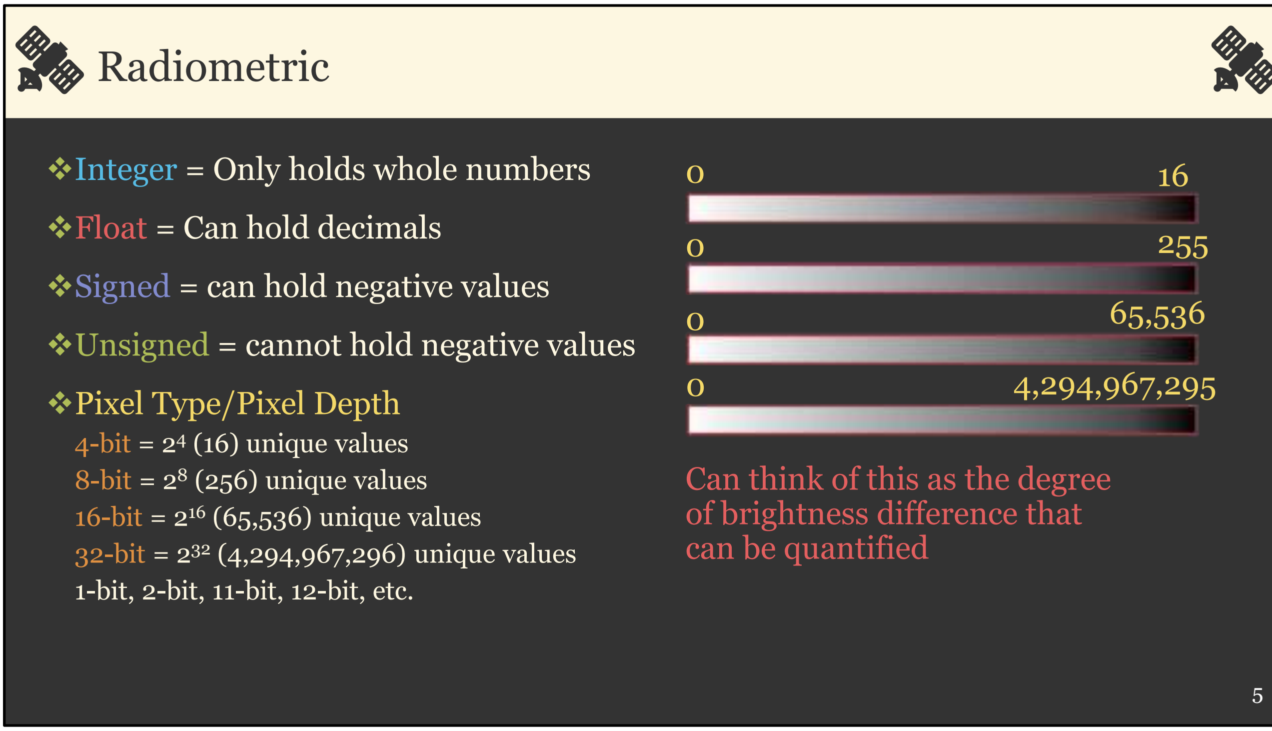

Radiometric resolution relates to what degree of differences in brightness can be distinguished. This relates to the pixel type or pixel depth of an image. For example, a 4-bit image would be able to differentiate 16, or 2 to the 4th power, discrete levels of brightness. In contrast, an 8-bit image could differentiate 256, or 2 to the 8th power, discrete levels of brightness. A 16-bit image could differentiate a large number of discrete brightness differences.

As an extreme case, a 1-bit image could only differentiate 2 to the 1st power, or 2, levels of brightness. So, the result would be black-and-white, or on-and-off.

You can think of this as the precision of the measurement. As an analogy, increasing radiometric resolution is like making length measurements to the nearest millimeter as opposed to the nearest meter. In other words, you will be able to differentiate smaller differences in length.

Many older sensors recorded 8-bit data. However, modern sensors are increasingly providing 11-bit or 16-bit data. Again, this allows for more precise differentiation of brightness or spectral reflectance. The trade-off is larger file size and greater data volumes.

Radiometric Resolution

- Gain = fineness of the divisions between successive brightness levels or sensor measurements (number of bits)
- Extent = range of brightness levels over which the sensor can differentiate change in radiance
- Signal-to-Noise Ratio = minimal radiance change that can be detected reliably

6

For radiometric resolution, grain relates to the fineness of the divisions between successive brightness levels or sensor measurements. In other words, it relates to the smallest difference in brightness that can be detected or quantified.

In contrast, extent relates to the range of brightness levels that can be detected. A sensor may not be able to detect a signal if it is too faint. If the signal is too strong, the measurement may become oversaturated. Think of the loss of detail in the dark shadows of an image or areas of glare or over exposure.

A related concept is the signal-to-noise ratio, which relates to how reliably we can separate signal, in this case differences in actual brightness, from random sensor noise.

When visualizing image data on a screen, we use the additive color space and the three additive primaries to display different channels or bands. For a simulated true color image, red = red, green = green, and blue = blue. For a false color image, we use the RGB color channels to map different bands (e.g., red = NIR, green = red, and blue = green).

Image Metadata

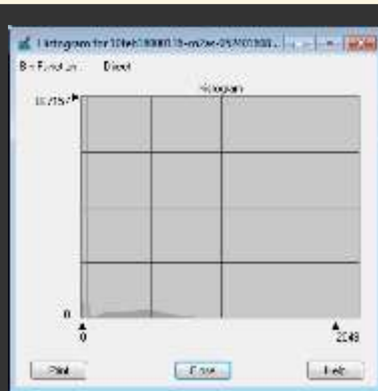

- Mean = average
- Median = central value
- Mode = most frequently occurring value

- Max – Min = Range

- Standard Deviation = measure of variability

8

Image data generally include metadata to highlight the distribution of DN values for each band. For example, it is common to report the mean or average DN value as a measure of central tendency. Alternatively, the median represents the most central DN value while the mode represents the most commonly occurring value. The distribution or dispersion of the data can be partially quantified using the range (maximum DN – minimum DN) and the standard deviation.

It is common for an image to not use the full radiometric range available and/or for the data values to be oddly distributed. For example, the provided histogram shows that many DN values in this specific band are near zero and most are low.

Lookup Table

- Mapping of file DN values to screen brightness levels

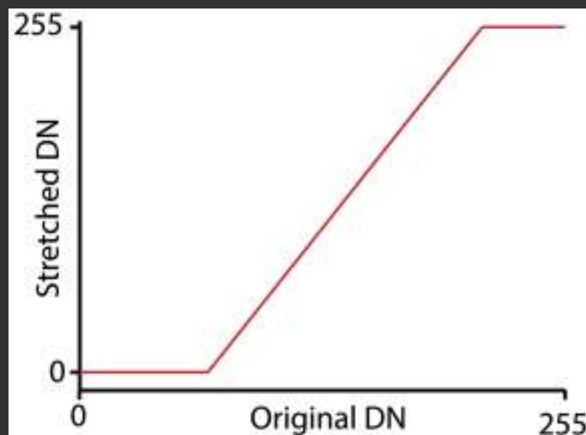

9

In order to map the raw, input DN values to brightness values to display on the screen, a lookup table, or LUT, is used. A lookup table describes how to map the input brightness to an output screen brightness. The provided graph explains this mapping process. The x-axis represents the original DN values while the y-axis represents the new or output values that will be displayed on the screen. The mapping is defined by the function, shown in red. In this example, low DN values will be mapped to 0, or be completely darkened to black. High values will be mapped to 255, or be completely brightened or mapped to white. Values in the middle will be stretched linearly to highlight brightness difference occurring in this range.

Imagine “Breakpoints” Tool

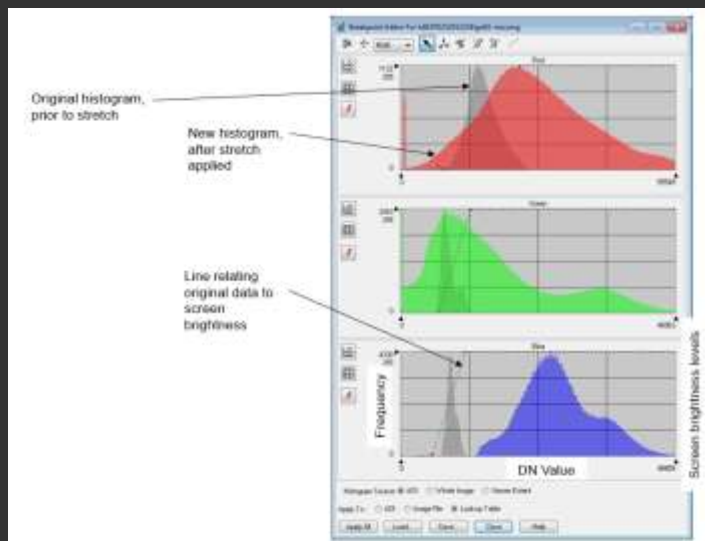

Example from Dr. Tim Warner

10

In Erdas Imagine, it is possible to define break points in the function manually and for each band separately. In the example, high and low thresholds are being set for each band to either darken (i.e., apply minimum screen brightness) or brighten (i.e., assign maximum screen brightness) the input DN values.

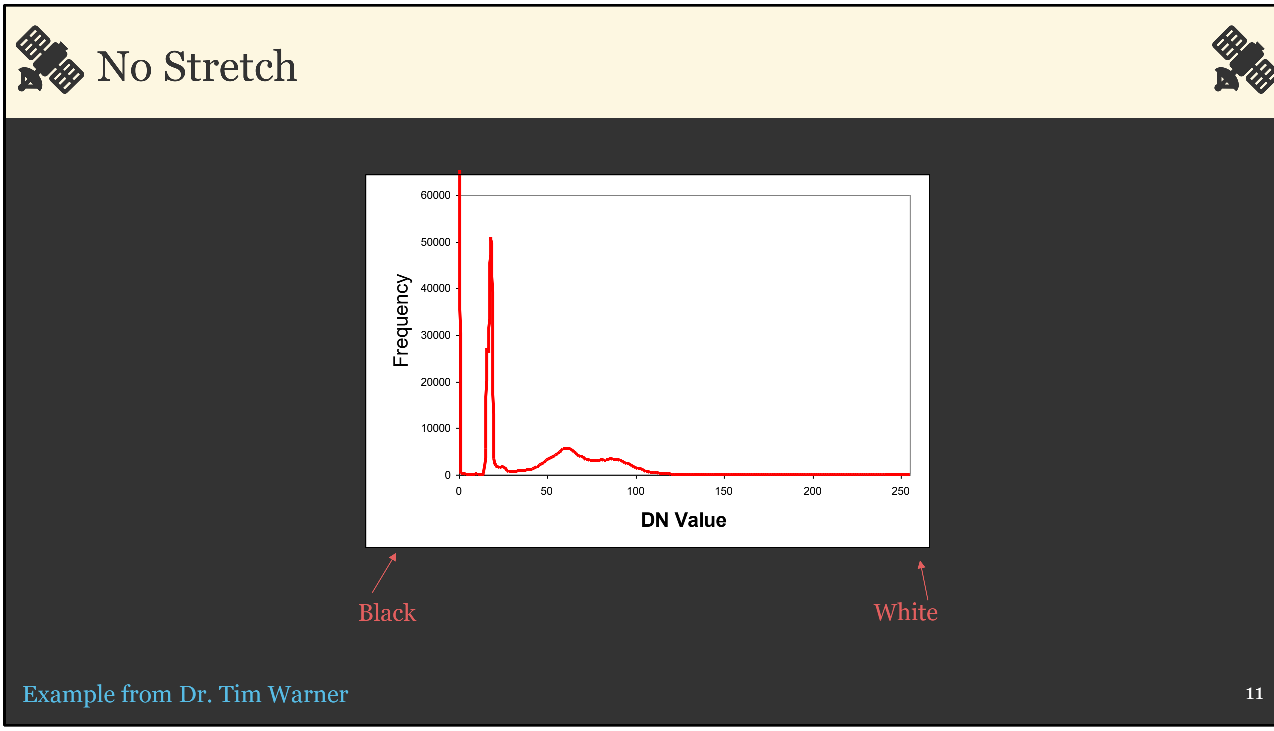

Histograms are used to graph or visualize the distribution of DN values per channel or band. The x-axis represents the DN values across the radiometric range (e.g., 0 to 255 for 8-bit data) while the y-axis represents the frequency of occurrence, or number of pixels having specific DN or brightness values. Peaks indicate common values while low points in the distribution represent rare values.

In the example, many cells have a low DN value assigned, with many near zero.

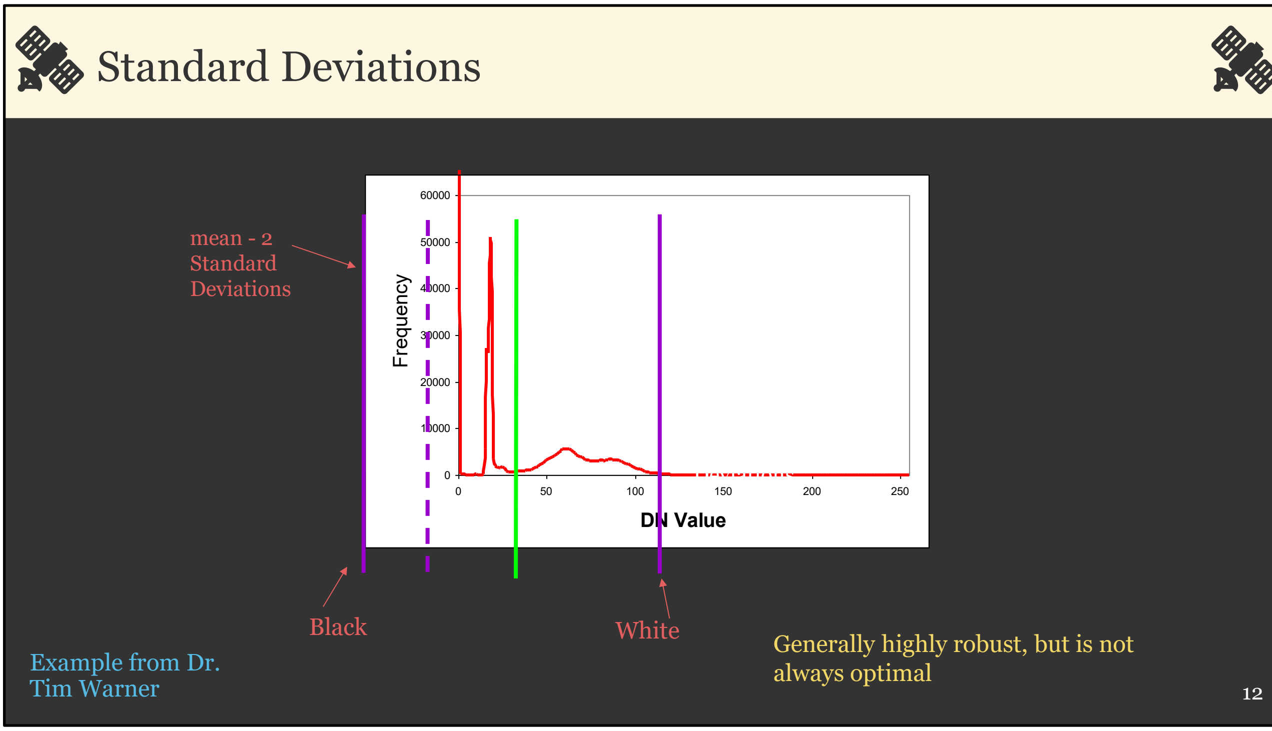

In order to more completely use the full radiometric range or screen brightness values, we can apply different data transformations to map from input DN to output DN. This slide provides an example of a 2 standard deviation stretch. Any value that is more than 2 standard deviations from the mean DN value will either be mapped to the lowest or highest output value. In other words, low values will map to black and high values will map to white. Intermediate values will be rescaled linearly. In the example, 2 standard deviations in the negative direction yields a negative number, so the range is readjusted to 0, or the lowest value.

Contrast Stretch and Enhancement Methods

- None = no stretching, input DN = output DN
- Standard Deviation Stretch = fully saturated or desaturated based on standard deviation from the mean
- Gaussian Stretch = center and scale data to a Gaussian distribution using mean and standard deviation
- Linear (Min/Max) Stretch = rescale linearly so that input min = output min and input max = output max
- Percent Clip Stretch = linear stretch, lower percentage will be desaturated, above a percentage will be saturated
- Histogram Equalization = spread out most frequent values and condense less frequent regions

13

Different contrast enhancement methods are available in different software packages. Here, I will describe some common methods.

First, none implies that no stretching is used: input DN values = output screen DN values. This is generally not preferred since the full screen brightness range will not be used.

Standard deviation was explained in the prior slide. Here, a distance from the mean DN value is defined based on standard deviation. Features that are higher or lower than these defined values will be displayed with either full brightness or zero brightness, respectively. All other values will be rescaled linearly to fill the full radiometric range.

A Gaussian stretch centers and rescales the data to approximate a Gaussian distribution using the DN value's mean and standard deviation.

Linear rescales the output so that the minimum input DN is scaled to 0 and the maximum input DN is scaled to the highest screen brightness (e.g., 255). This scaling is accomplished linearly.

Percent clip is similar to a linear stretch except that a defined upper and lower percent of the data are assigned either full or 0 brightness. All values in between are linearly stretched.

Lastly, histogram equalization alters the DN values such that frequent DN values are spread out to use more or a larger range of brightness values and rare DN values are condensed. This essentially highlights the difference between common brightness values while minimizing difference between less common values.

Example 1

Hobart, Tasmania

WorldView-2 data
Standard false color composite
Bands 7,5,3 as RGB

14

We will now explore these different enhancement techniques by visualizing the results obtained when applied to real datasets. Our first example is a WorldView-2 image of Hobart, Tasmania. We are specifically using a standard false color composite: red = NIR, green = red, blue = green.

Example 1

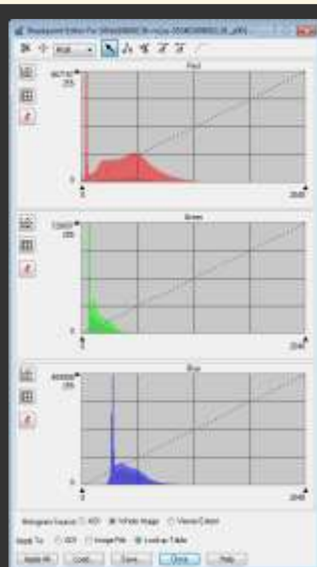

Linear “No Stretch”

15

First, no stretch is applied. As you can see, this results in a poor visualization since the full radiometric range is not being used. The image looks very “washed out” and dark. This should make sense since the histograms show that the DN values are clustered in the lower DN ranges for all displayed bands.

Example 1

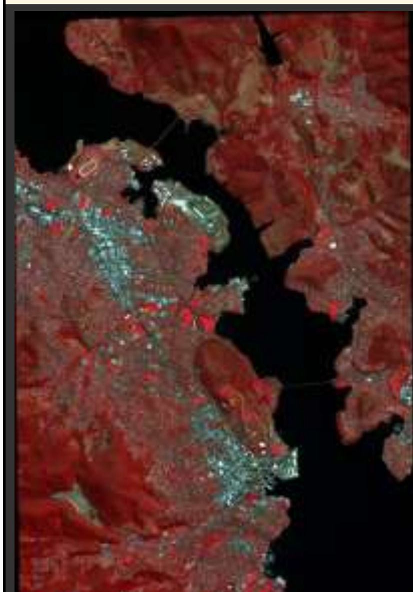

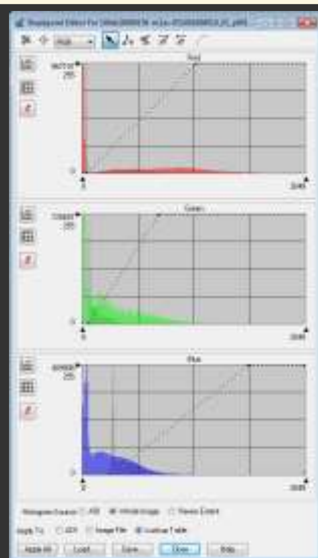

Saturate Upper 1%, Lower 2%

16

Here, I have used a percent-clip method in which the lower 1% of the DN values are converted to black or 0 and the upper 2% are converted to white or full brightness. The remaining values are rescaled linearly. This results in a much more pleasing and useful visualization of the data. In this example, it makes sense to provide a larger clip on the upper end since the values are concentrated in the lower DN value ranges. This will allow for more, brighter values to be used effectively.

Example 1

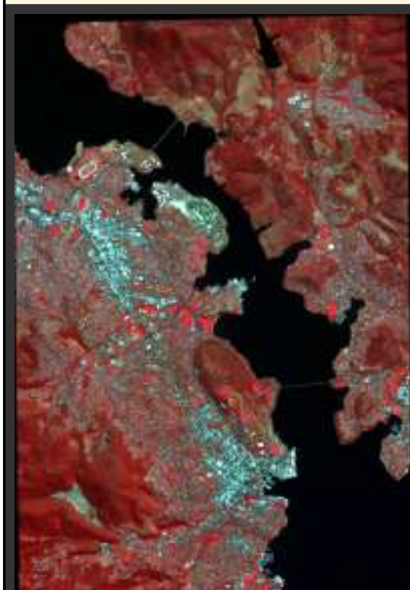

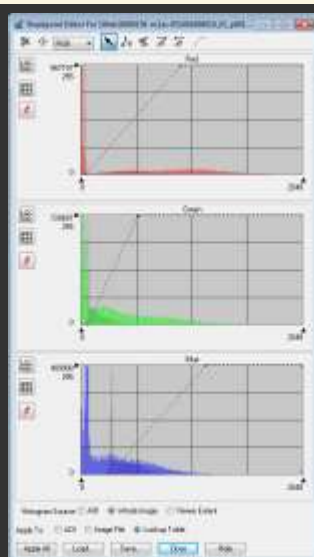

Saturate Upper 2%, Lower 2%

17

This image is similar to the last image; however, I have now expanded the lower clip to 2%.

Example 1

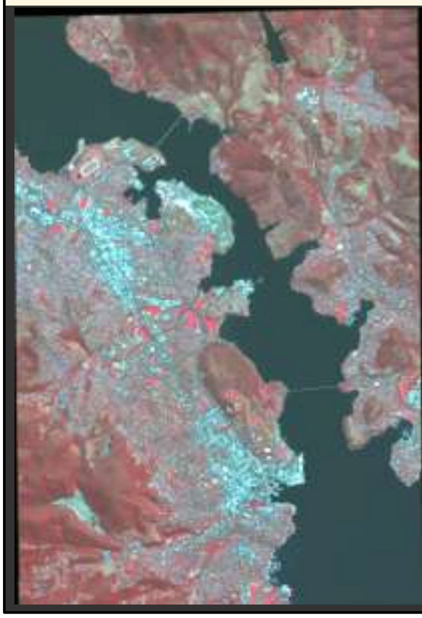

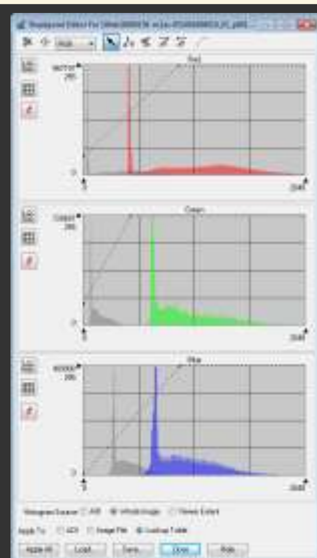

Two Standard Deviation Stretch

Here, a standard deviation stretch is applied where values lower than 2 standard deviations from the mean are mapped to 0 or black and values 2 standard deviations larger than the mean are mapped to full brightness. This stretch does improve the brightness in comparison to not using a stretch. However, the image looks "washed out" due to the odd distribution of the input data. In this case, more stretching is applied to higher DN values in comparison to low values.

Example 1

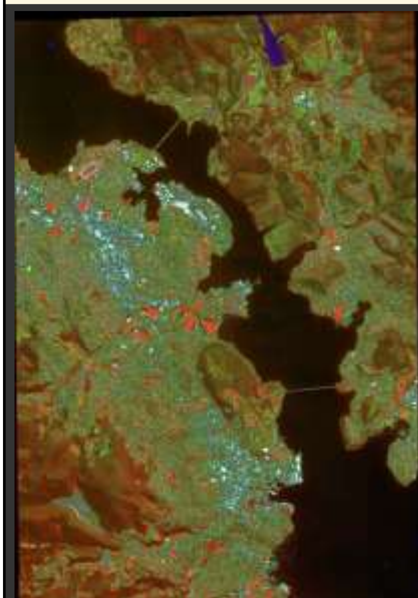

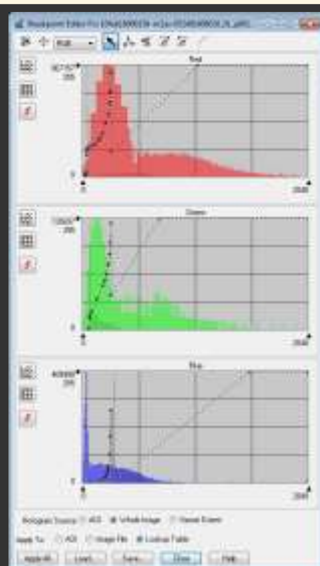

Gaussian

19

Here, I have applied a Gaussian stretch to make the data distribution approximate a Gaussian distribution. Note the odd function required to achieve this as defined by the curves plotted on the histograms.

Example 2

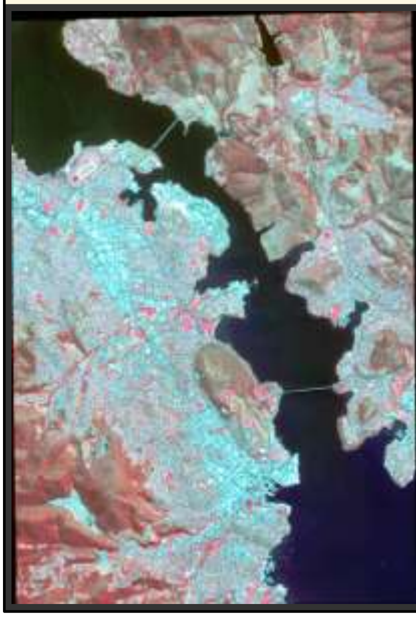

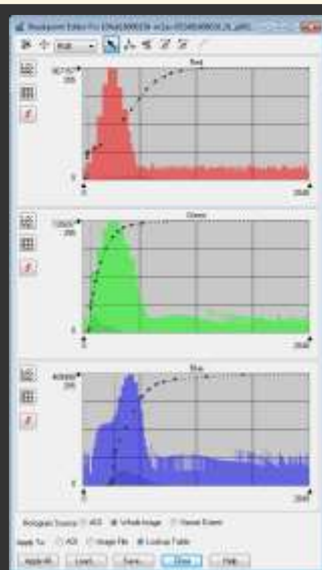

Histogram Equalization

20

Lastly, a histogram equalization is applied. Although this results in a strange visualization, it does highlight key differences in brightness. For example, changes in the spectral reflectance of the water across the scene is more discernable.

Example 2

- Miami, Florida
- Landsat 5Thematic Mapper

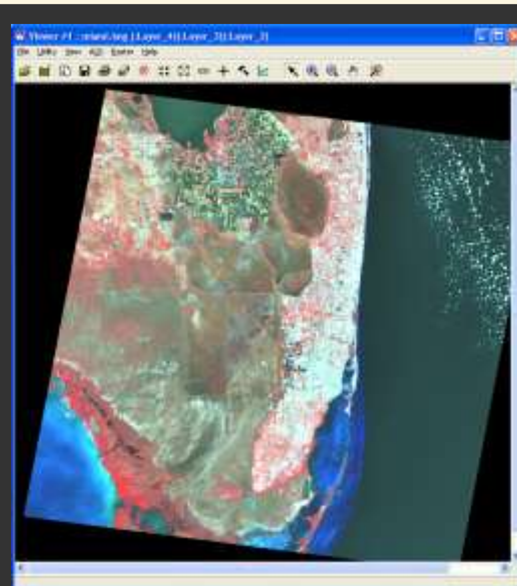

21

This second example is a Landsat 5 Thematic Mapper image of Miami, Florida. Again, we are using a standard false color composite (red = NIR, green = red, blue = green).

Example 2

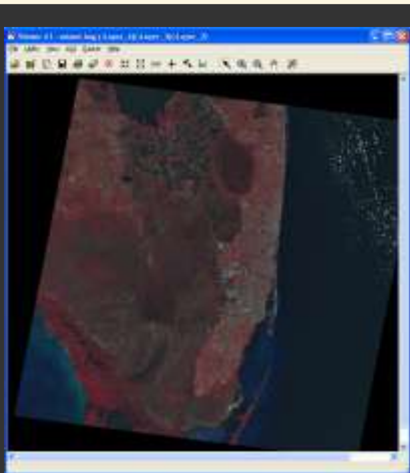

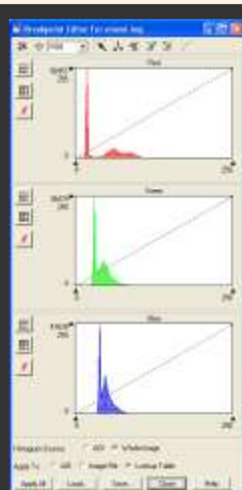

Linear “No Stretch”

22

Here, no stretch is applied. Again, the image is dark since the full radiometric range is not used and most of the DN values are low.

Example 2

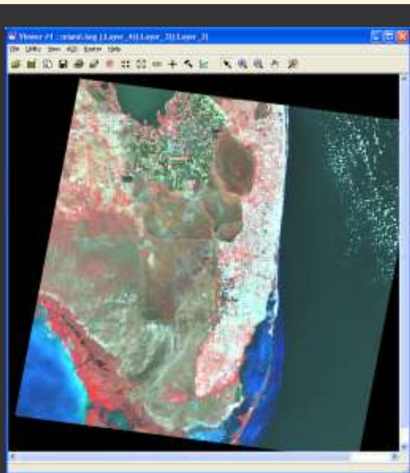

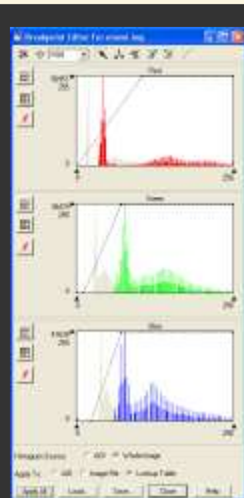

Two Standard Deviation Stretch

23

Here, I am using a 2 standard deviation stretch that brightens and highlights differences.

Example 2

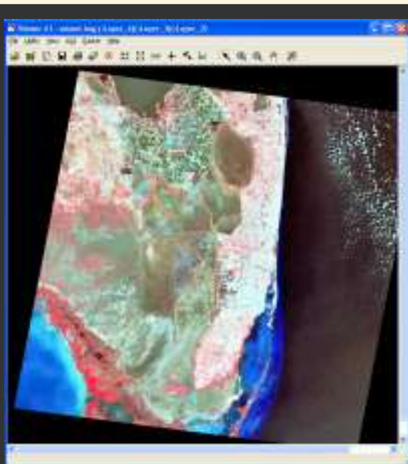

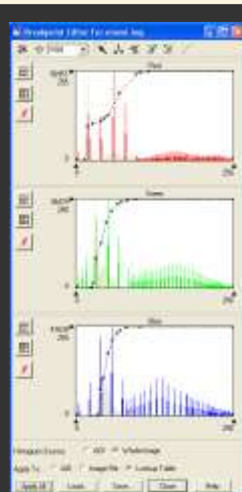

Histogram Equalization

24

Histogram equalization further highlights differences.

Example 2

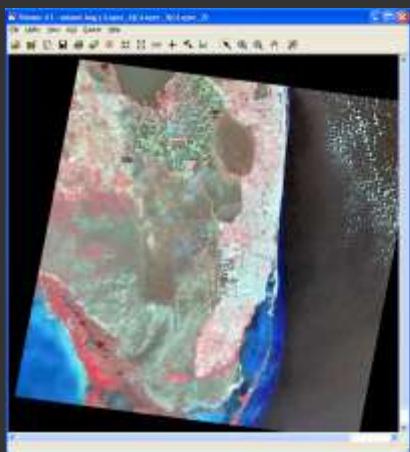

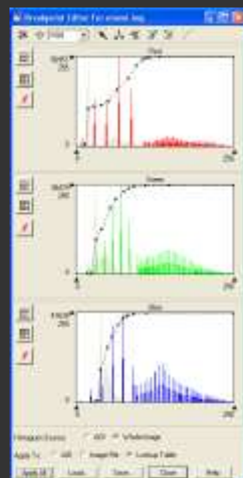

Gaussian

25

This example shows a Gaussian stretch.

Example 2

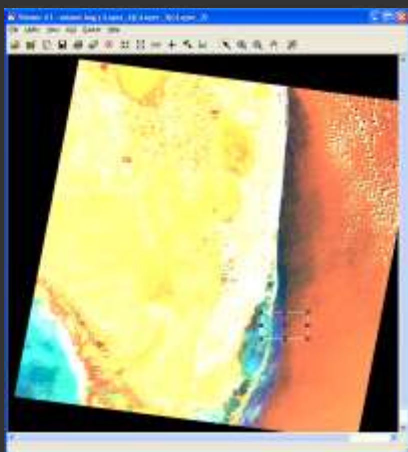

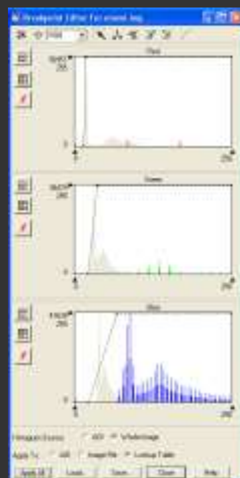

Two Standard Deviation Stretch, Just Over Water

26

Part of the reason the applied stretches look odd for this image is due to the odd distribution of the DN values, which is partially attributed to the nature of the image. This image is divided between land and water, which tend to have very different spectral reflectance.

If you are specifically interested in certain features, it is possible to apply stretches based on a subsample of the image containing the features of interest. In this example, a 2 standard deviation stretch has been defined based on a subset of water pixels. This highlights differences in water reflectance. Land pixels are poorly displayed. However, land may not be of interest, and, thus, this wouldn't be an issue.

Example 2

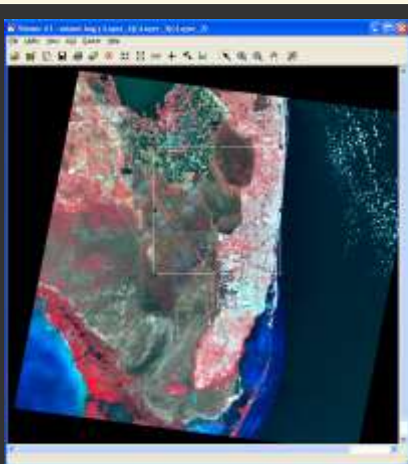

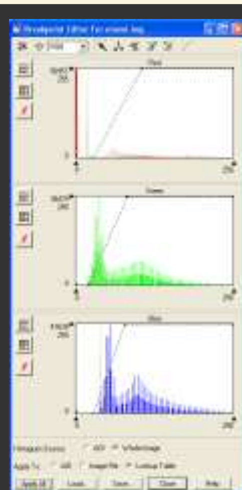

Two Standard Deviation Stretch, Over Land

27

Here, I have stretched the image using only a subset of land pixels, allowing for a greater differentiation of land features.

Example 2

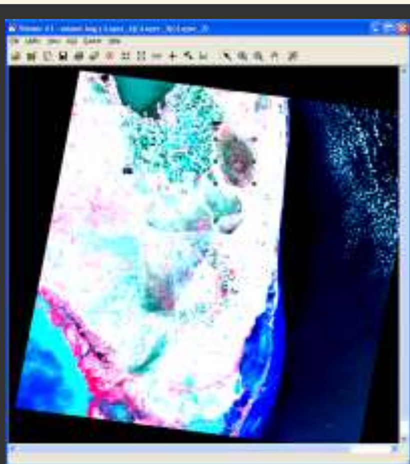

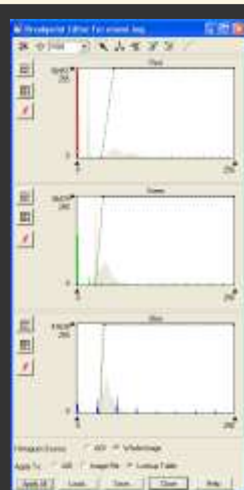

Two Standard Deviation Stretch, Over Swamp

I have now more specifically applied a 2 standard deviation stretch based on only pixels from a swamp. This results in a useful visualization of the swamp pixels specifically.

Custom Stretch

- Sometimes best to create custom stretch
  - Select statistics from image sub-set
  - Manually adjust
    - Upper and lower saturation %
    - # of standard deviations
    - Contrast and brightness
- Breakpoints tool to develop complex lookup table

29

It is possible to define custom stretches by manually adjusting the mappings.

Other Display Adjustments

- Brightness = increase or decrease overall brightness
- Contrast = difference between dark and light pixels
- Gamma = degree of contrast between the midlevel gray values (does not affect the black or white values, only the middle values)

30

It is also possible to apply some generic, global display adjustments. For example, ArcGIS Pro allows for adjustments of brightness, contrast, and gamma, which are defined on this slide.

Dynamic Range Adjustment (DRA)

- Stretch applied to subset of image
- Stretch applied to current view extent

31

Lastly, it is also possible to define a stretch based on DN values in a defined subset of the image or have the stretch update based on the current viewing extent. This last option is known as dynamic range adjustment, or DRA, and can be useful when you want to use the entire screen brightness range every time you change the viewing extent.

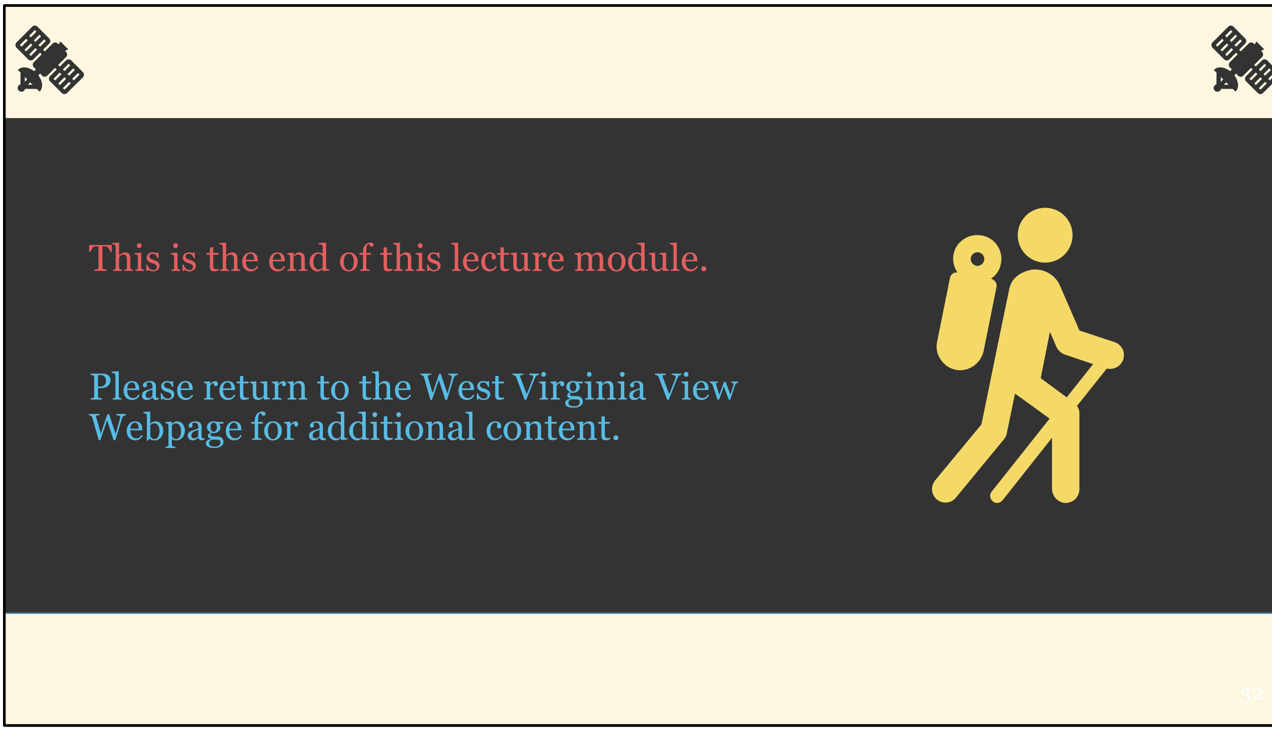

Thanks! Hope you found this useful.