Digital Cartography

Imagery and Terrain Data

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

WV VIEW

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

Imagery and digital terrain representations are commonly used to provide spatial context and visualize landscape conditions. So, it is important to appropriately symbolize these data for desired cartographic effect.

Module Topics

1. Image data review
2. False color composites
3. Contrast stretching
4. Brightness, contrast, and gamma
5. Dynamic range adjustment
6. Color ramps for elevation
7. Hillshades and terrain visualization
8. Hypsometrically tinted hillshades

2

These are the topics covered in this module.

Imagery

3

Types of Images

- Panchromatic = Single band for entire visible spectrum
- True Color = Red, Green, and Blue bands (RGB)
- Color Infrared = Red, Green, Blue, and Near Infrared (NIR)
- Multispectral = Multiple spectral bands collected
- Hyperspectral = Tens to hundreds of spectral bands

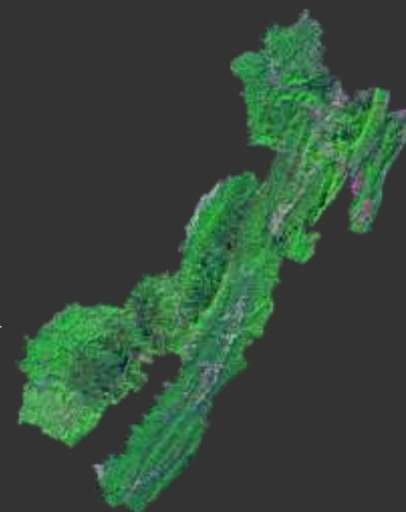

Data from Landsat 5 Thematic Mapper

4

Here is a brief review of image types.

Panchromatic images are black-and-white, single-band images that do not provide separate channels or bands. A wide range of wavelengths are sensed and recorded as a single gray-scale image. Although this is a bit limiting, panchromatic images are still used. In fact, panchromatic images are commonly collected along with multispectral data. The panchromatic image is generally collected at a higher spatial resolution and can be used to pansharpen multispectral data.

True color images are designed to mimic or simulate human vision. Your eyes contain red, green, and blue cones. Your visual system can then combine the primary colors to create a color view of the world. A true color camera works similarly. The red, green, and blue channels are collected separately and stored in separate bands. They are then combined to produce a color at each pixel. Most of the images we collect on a daily basis, such as pictures taken with a cell phone, are true color images.

Color infrared, or CIR, images collect visible bands and a near infrared (NIR) band. As previously discussed, near infrared wavelengths can be useful for differentiating and mapping vegetation and vegetation health. So, collecting this information can be of value. Older film-based methods were able to collect green, red, and near infrared data. Modern digital sensors can collect four bands of data: blue, green, red, and near infrared.

Multispectral images generally collect more than four bands. For example, the Landsat 5 Thematic Mapper sensor provides one blue, one green, one red, one near infrared, two shortwave infrared, and one thermal infrared band.

Hyperspectral images can provide tens or hundreds of bands as a data cube and more closely approximate the full spectral reflectance curve at a pixel location.

As opposed to a simulated true color composite, a false color composite can be used to symbolize data or bands outside of the visible spectrum.

For example, a standard false color composite uses the red channel to show near infrared, the green channel to show red, and the blue channel to show green.

In the image provided, red areas show up as red because they are reflecting a lot of near infrared radiation. This is an indicator of green, healthy, photosynthetically active vegetation.

So, using color channels to show different wavelength ranges can help us visualize the information made available in multispectral images.

Composites

- Red channel shows shortwave
- Green channel shows near infrared
- Blue channel shows green

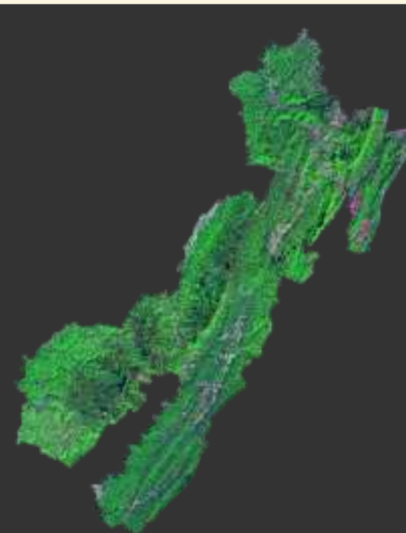

Data from Landsat 5 Thematic Mapper

6

Here is another example of a false color composite that is commonly used for Landsat data. The red channel shows a shortwave infrared band, the green channel shows near infrared, and the blue channel shows green. So, green areas appear green because they reflect a lot of near infrared, a signature of vegetation.

Red or pink areas would indicate areas that reflect a lot of shortwave infrared, such as bare soil.

Multiband Data in ArcGIS

- A stack of raster grids
- Each layer represents an image band

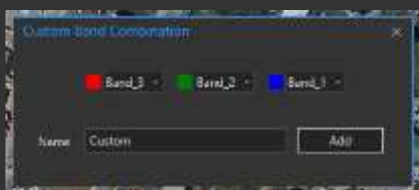

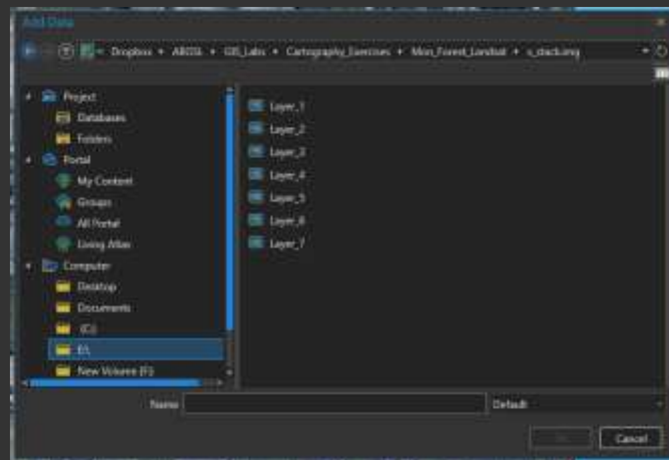

7

If you would like to experiment with changing image symbology in ArcGIS Pro, this can be accomplished using the Band Combinations options in the Appearance Tab.

You can also add individual layers or bands from a dataset.

When using image data in a map layout, it is important to consider which band combinations are most appropriate. For example, a general audience may find a false color composite confusing without additional explanation provided. However, professionals may prefer a false color composite that highlights features of interest.

Contrast Stretching

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

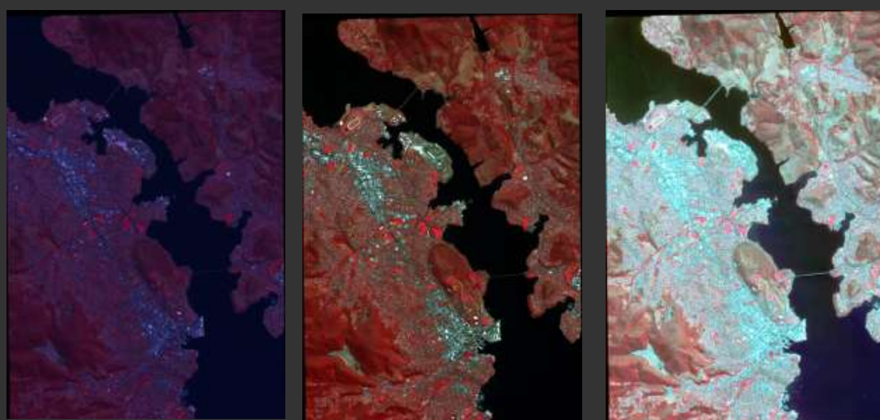

8

Contrast stretching is the process of mapping the original image digital number (DN) values to new DN values to take full advantage of the dynamic range and enhance contrast between features. In the next set of slides, we discuss methods to accomplish this.

Pixels and DN Values

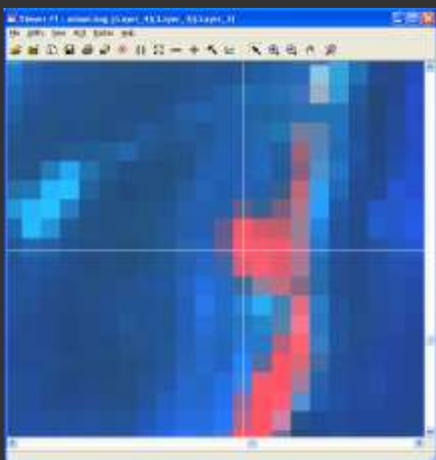

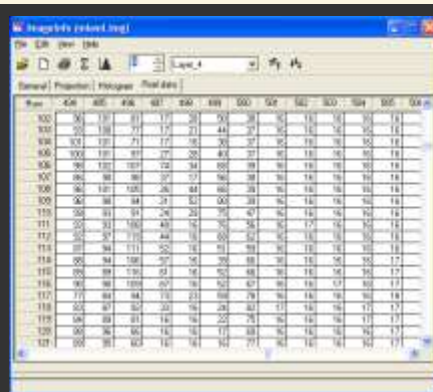

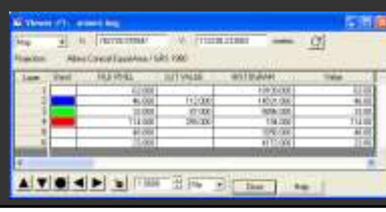

9

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.

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

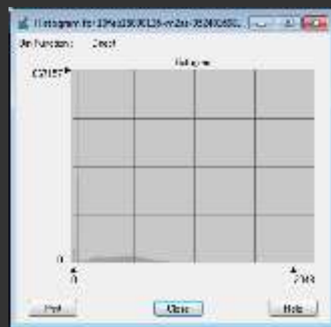

- Mean = average
- Median = central value
- Mode = most frequently occurring value
- Max – Min = Range
- Standard Deviation = measure of variability

11

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

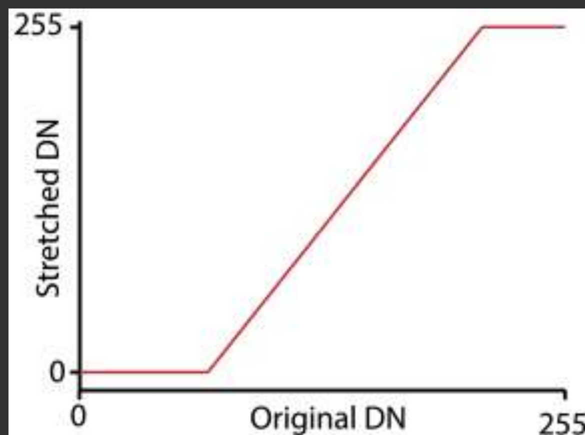

12

In order to map the raw, input DN values to brightness values to display on the screen, a lookup table 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.

No Stretch

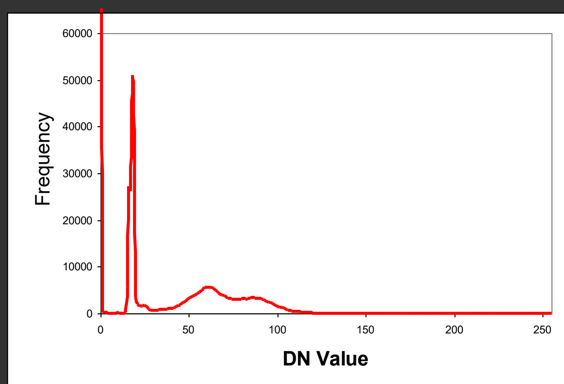

Black

White

Example from Dr. Tim Warner

13

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.

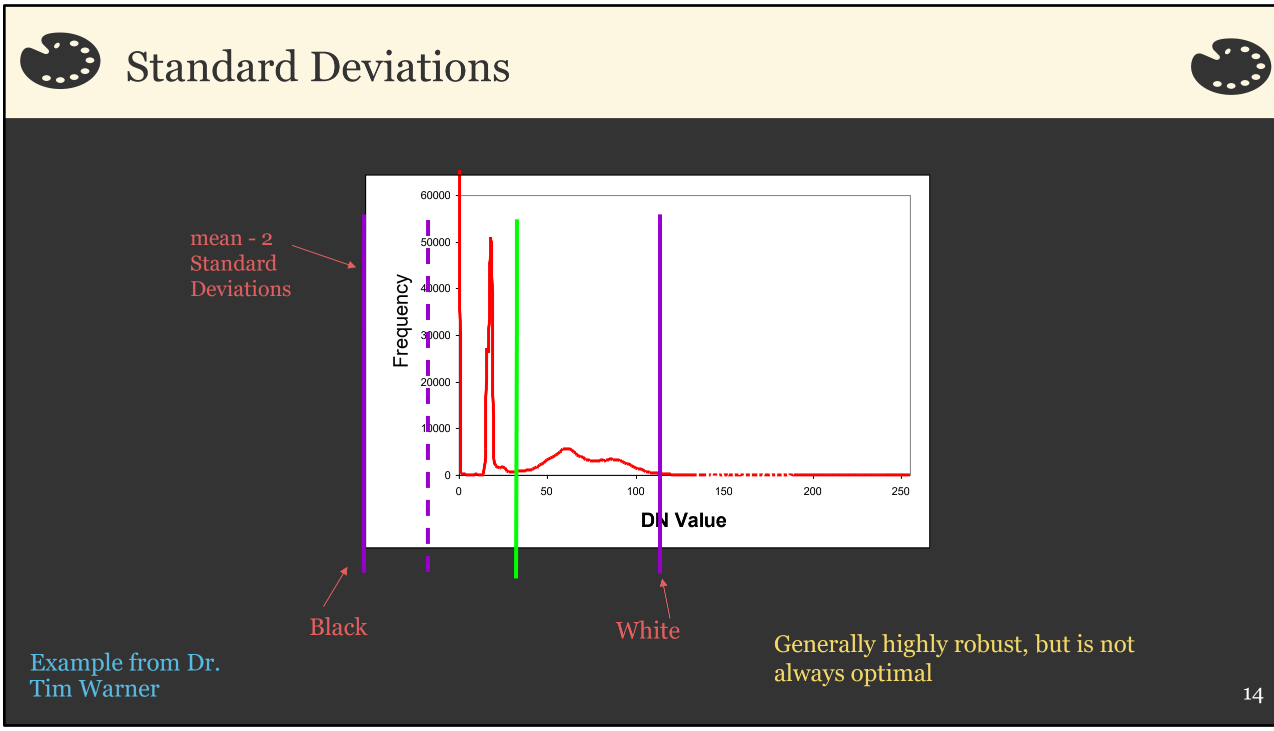

In order to more completely use the full radiometric range or screen brightness values, we can apply different data transformations to map from input DN to output DN. This slide provides an example of a 2 standard deviation stretch. Any value that is more than 2 standard deviations from the mean DN value will either be mapped to the lowest or highest output value: 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

15

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. Or, 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

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

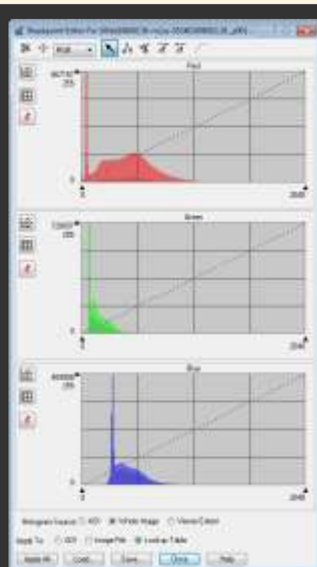

Linear “No Stretch”

17

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

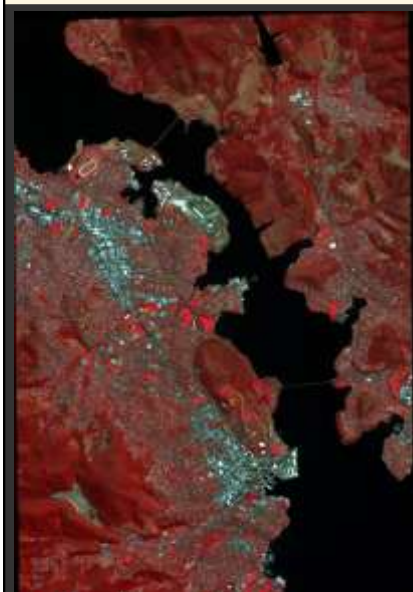

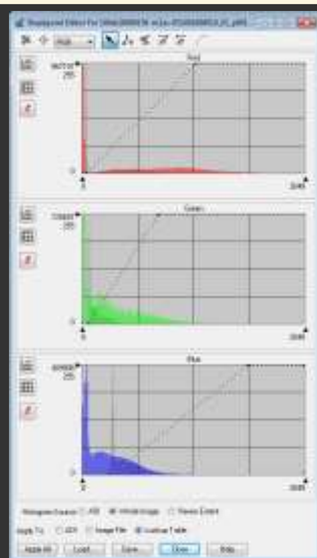

Saturate Upper 1%, Lower 2%

18

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

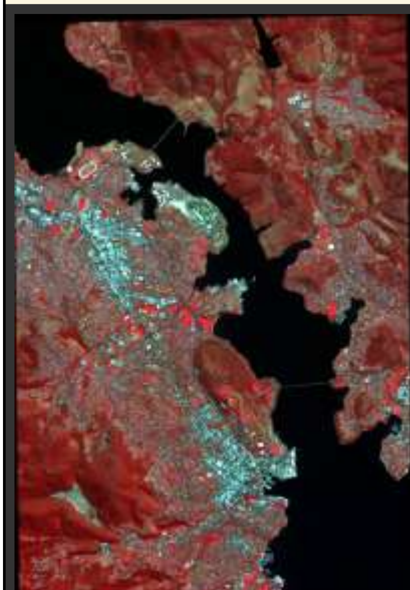

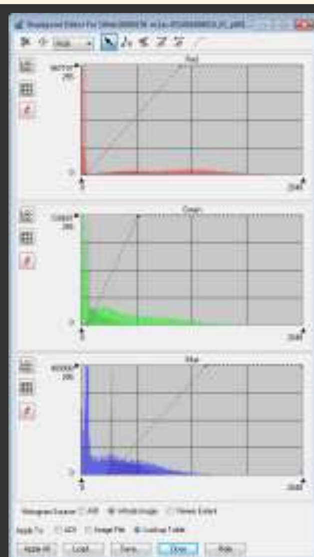

Saturate Upper 2%, Lower 2%

19

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

Example 1

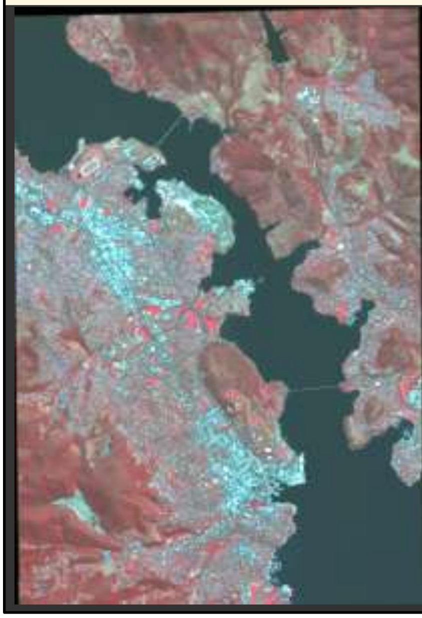

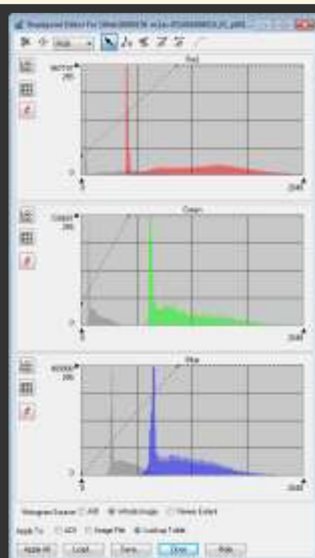

Two Standard Deviation Stretch

20

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

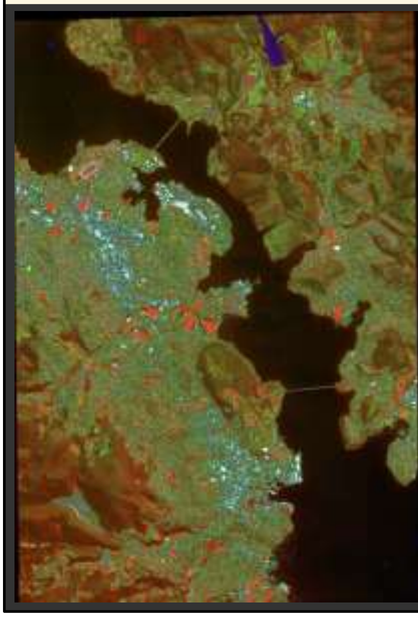

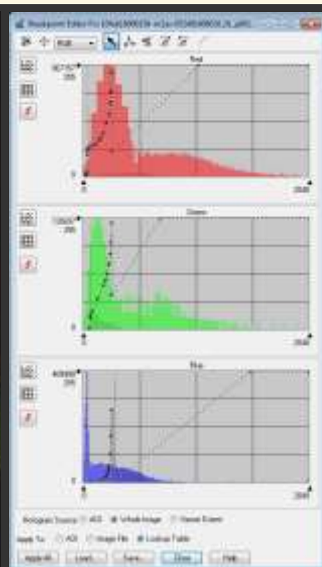

Gaussian

21

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 1

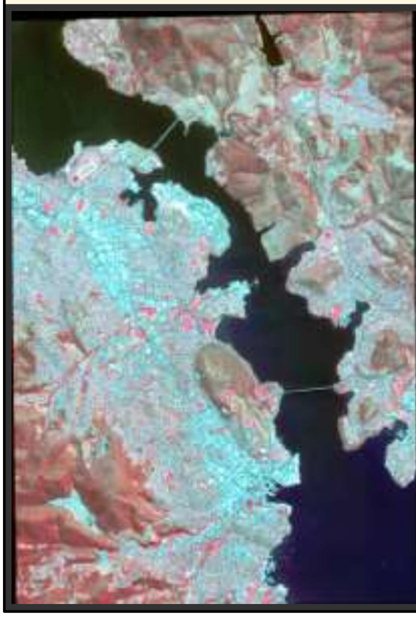

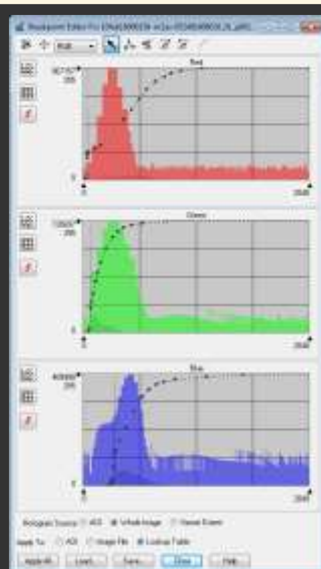

Histogram Equalization

22

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

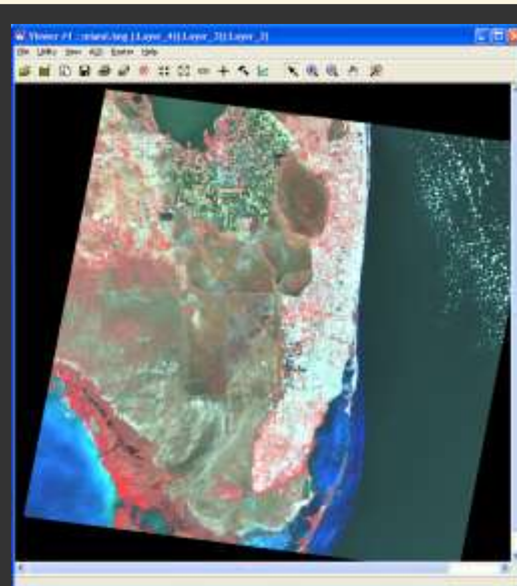

23

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

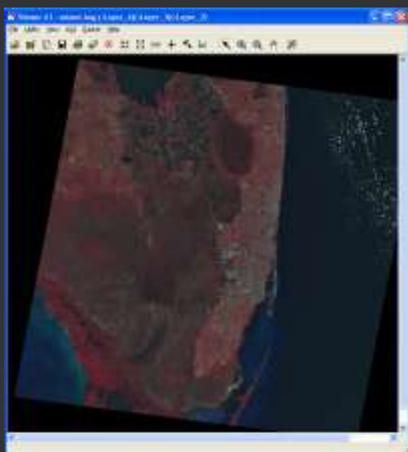

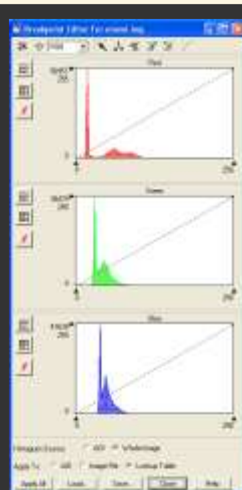

Linear “No Stretch”

24

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

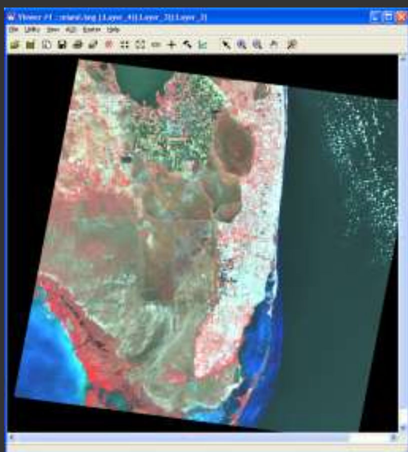

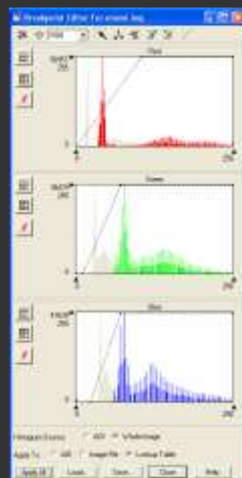

Two Standard Deviation Stretch

25

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

Example 2

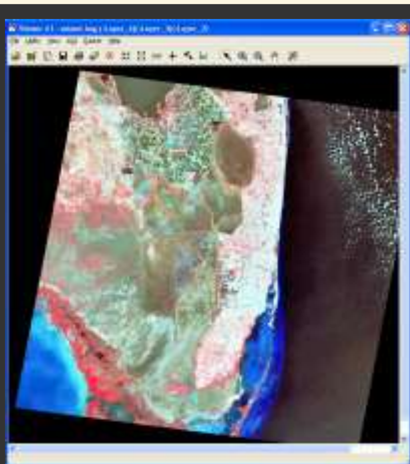

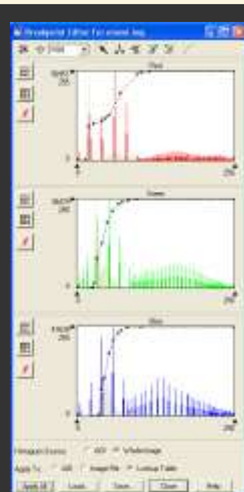

Histogram Equalization

26

Histogram equalization further highlights differences.

Example 2

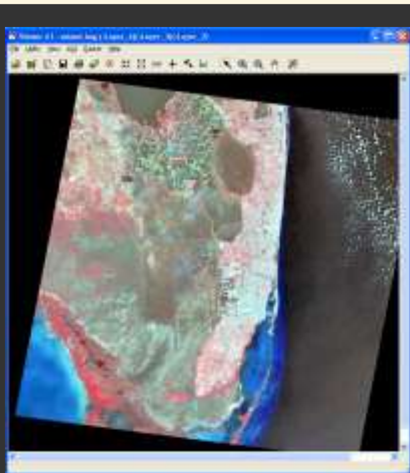

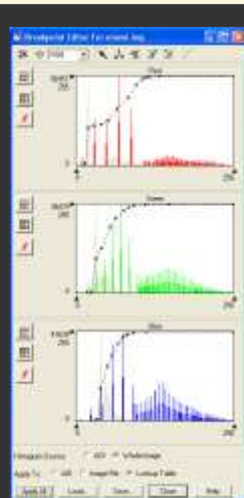

Gaussian

27

This example shows a Gaussian stretch.

Example 2

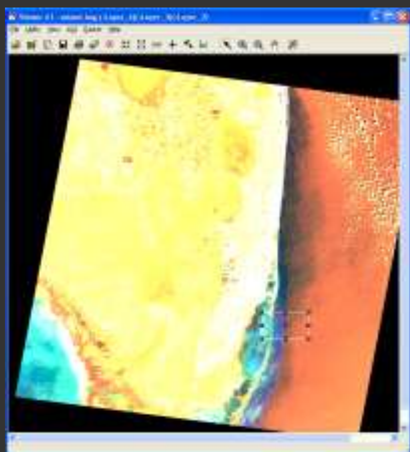

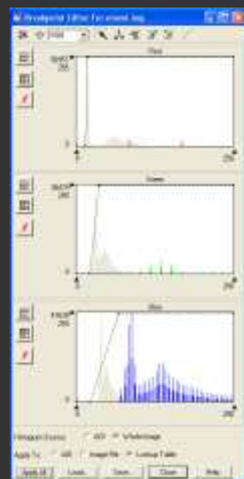

Two Standard Deviation Stretch, Just Over Water

28

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

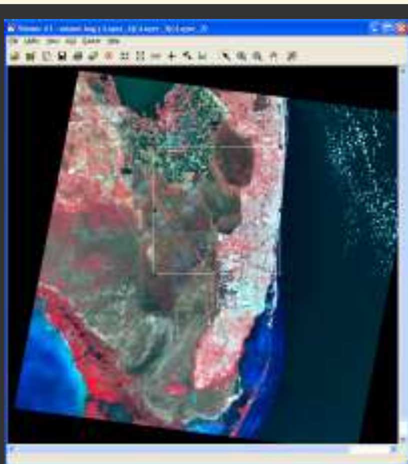

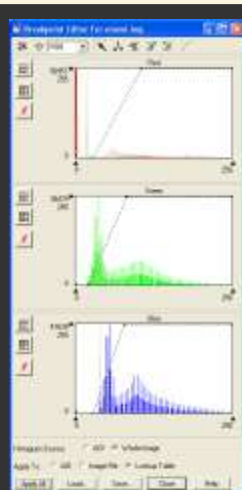

Two Standard Deviation Stretch, Over Land

29

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

Example 2

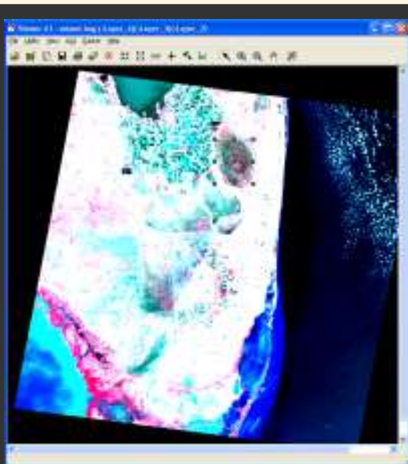

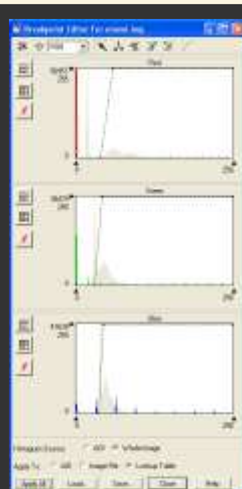

Two Standard Deviation Stretch, Over Swamp

30

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

31

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 colors
- Gamma = degree of contrast between the midlevel gray values (does not affect the black or white values, only the middle values)

32

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

33

Lastly, it is also possible 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.

Terrain Data

34

Symbolizing DEMs

- Can use light-to-dark sequence that also transitions through related hues to provide more contrast
- Can be binned or stretched
- Light = low elevation
- Dark = high elevation
- Make sure DEM is of adequate resolution for map
  - May need to find better data
  - May need to smooth or generalize high spatial resolution data

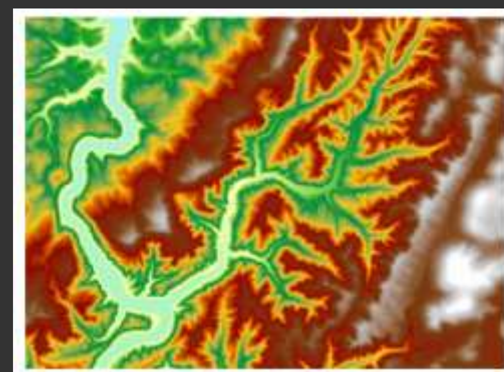

35

Raster-based digital elevation models (DEMs) or digital terrain models (DTMs) provide a representation of the terrain surface where each cell holds an elevation measurement.

When selecting a color scheme for elevation data, it is generally recommended to use a light-to-dark sequence that also transitions through related hues to enhance contrast. A color ramp can be applied using a classification, also known as binning, or mapped continuously as a stretch.

It is recommended that lighter shades be used to show the lower elevations while darker shades are used to show the higher elevations when using a single hue and when not using a black-to-white stretch. When using black-to-white, low values should map to black and high values should map to white.

You should also consider the spatial resolution of the DEM. A coarse DEM may not be appropriate for a large scale map. On the other hand, detailed DEM data, such as those derived from LiDAR, may be too detailed and may require some smoothing or generalization.

Contours

- ❖Raster → vector
- ❖Must define an interval
- ❖Make principal or index contours thicker
- ❖Make intermediate contours thinner
- ❖May require some smoothing

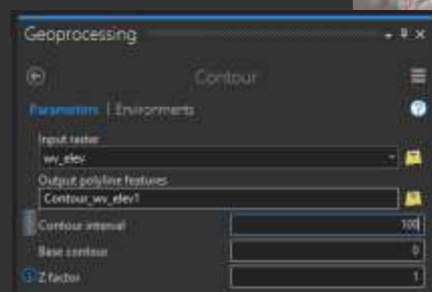

36

Elevation information can also be represented using contour lines. The principal or index contours are less frequent and are commonly labeled and use a thicker symbology. In contrast, intermediate contours are generally thinner and not labeled. The correct contour interval will depend on the scale of the map layout, the level of terrain detail desired, and the topographic conditions of the mapped landscape. For example, the contour interval may need to be smaller in areas with less topographic change.

As an example, USGS 1:24,000 scale topographic maps use a principal contour interval of 100 feet and an intermediate interval of 20 feet. So, at each 100 feet of change, a thicker, labeled contour line is provided. Between these principal contours, the intermediate contours are thinner and not labeled.

Line Generalization

- Remove unnecessary vertices
- Smooth curves
- For cartographic generalization
- Simplify Line Tool
- Smooth Line Tool

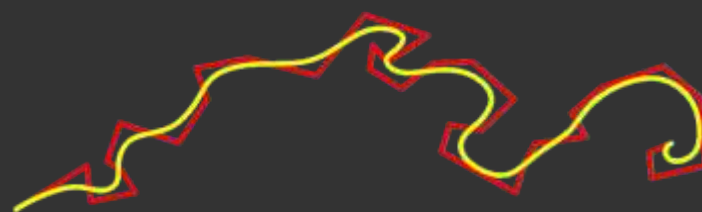

Available under Generalization Tools in the Cartography Toolbox

37

When contour lines are created from DEM data, especially if the data are very detailed, the lines may be overly complex. In order to improve the cartographic display, it may be desired to smooth the lines. The amount of smoothing to incorporate is case-specific and will likely require some experimentation.

Combining DEMs and Contours

❖Map color transitions with principal contours

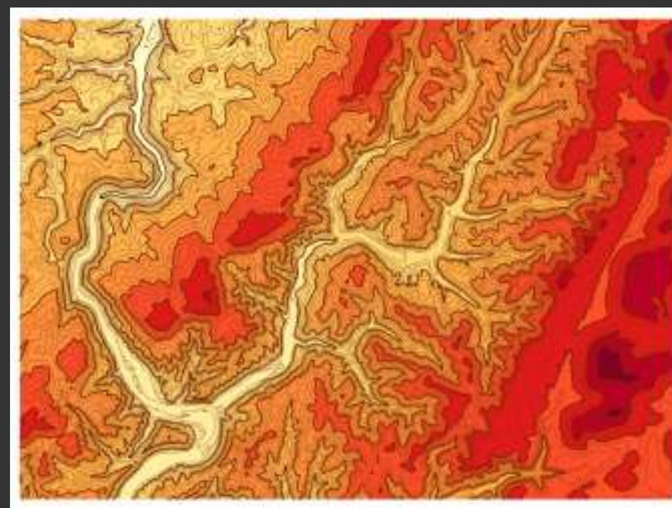

38

It is common to combine contour lines and hypsometric tinting from a DEM to visualize the terrain surface. When these representations are combined, it is generally recommended that the color breaks align with the principal contours so that they reinforce one another. This makes for a very nice visualization, as demonstrated in the provided example.

Hillshade

- A visual representation of a terrain using shadows
- Shadows are cast over the landscape to created a shaded relief surface
- Azimuth = compass bearing of Sun
- Altitude = height of Sun above horizon (as an angle)
- Color ramps must run light to dark

39

Hillshades offer a means to visualize digital terrain data. In a hillshade, a sun or illuminating source is placed in the sky to model illumination. The result is a shaded relief surface. Bright areas represent areas of light exposure whereas dark areas have less light exposure.

The position of the illuminating source is defined using a compass bearing or direction, this setting is known as azimuth, and the angle of the source above the horizon, or altitude. Using azimuth and altitude, you can position the illuminating source anywhere in the local sky.

Imagine a dome over your head. If you would like to locate an object on this dome, you would need to specify the compass direction or bearing in which to look, or azimuth, and how high up to look as an angle, or altitude.

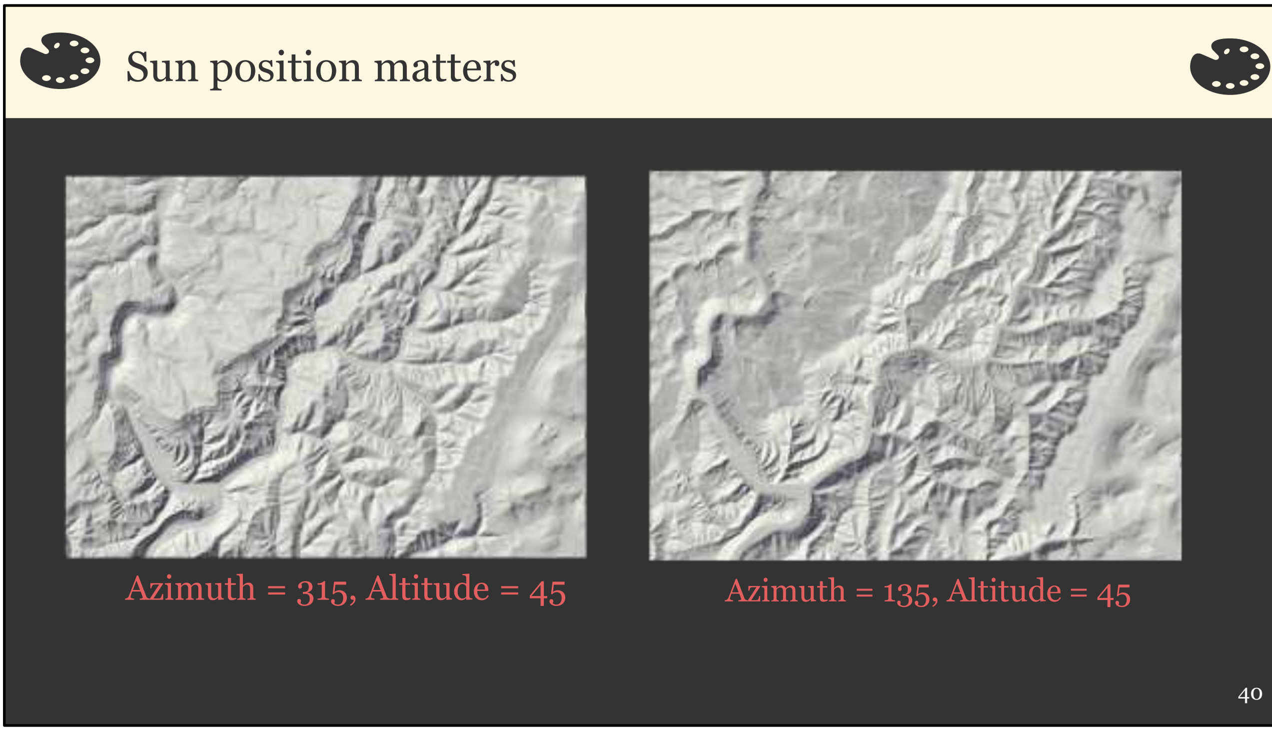

The location of the illuminating source impacts the effectiveness and look of the hillshade. In ArcGIS Pro, the default settings are an azimuth of 315 and an altitude of 45. So, the illuminating source would be in a northwest position and halfway between the horizon and zenith, or top of the sky.

Since we generally see features illuminated from above or by the Sun, we expect shadows to be cast downward. This is mimicked with a Sun position in the northwest azimuth. If we move the Sun to a southeast azimuth, then features will be illuminated from below. This will result in a strange optical illusion where the topography looks inverted: ridges will be interpreted as valleys and valleys as ridges.

So, we tend to avoid using southeast azimuth positions.

Sun position matters

Azimuth = 315, Altitude = 45

Azimuth = 315, Altitude = 10

41

We also tend to avoid using low altitudes as this can create very contrasted hillshade images with very dark and bright slopes. Placing the illumination source very close to the zenith can make hillshade images look very "washed out" or low contrast.

I have generally found the default settings to work well. However, you may want to experiment with different settings. Different settings may be optimal in different terrains.

Multidirectional Hillshades

- ❖Combine multiple Sun positions
- ❖North + Northwest + West + Southeast
- ❖Extra weight for northwest
- ❖Subtle detail on slopes that face different directions
- ❖Useful in subtle terrain

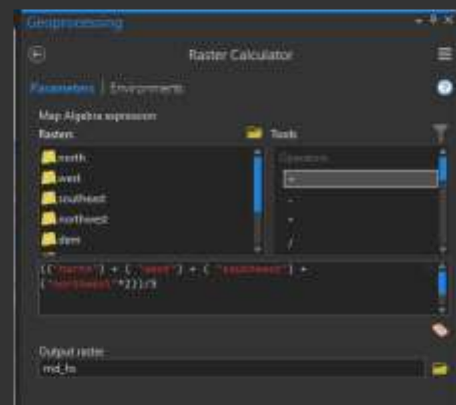

42

Another option is to calculate multiple hillshades using different illuminating positions and then combine them to create a multidirectional hillshade. In the example, I have created hillshades using north, west, southeast, and northwest positions. I then use Raster Calculator to average them. I double-weight the northwestern result so that it has the largest impact on the output.

Using multiple directions generally allows for a reduction in deep shadows or extremely bright cells in complex topography. If a hillshade using the default illuminating position or a single position is not adequate for your visualization, it is worth experimenting with multidirectional hillshades as an alternative.

Multidirectional Hillshades

- Combine multiple Sun positions
- North + Northwest + West + Southeast
- Extra weight for northwest
- Subtle detail on slopes that face different directions
- Useful in subtle terrain

43

This slide shows the result of the multidirectional hillshade calculation.

Slopeshade

- Display slope using a white-black stretch
- White = shallow slopes
- Black = steep slopes

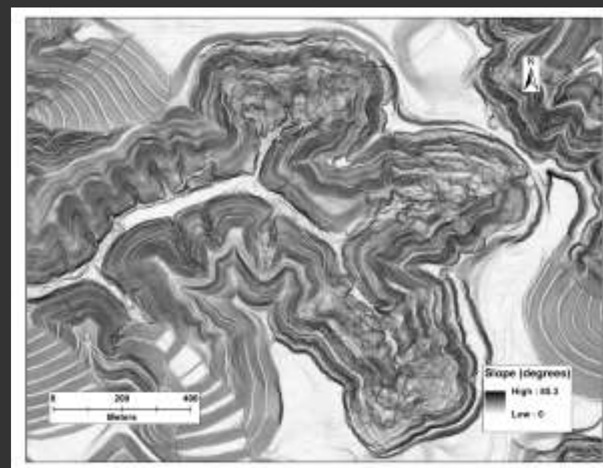

Map by Miles Reed

44

An alternative to the hillshade is the slopeshade. This is simply a topographic slope raster grid with a white-to-black color ramp applied. White shades indicate low slopes or flat areas while black indicates steep slopes. When quality, high spatial resolution DEM data are used, slopeshades can offer a very detailed terrain representation for interpretation. However, these visualizations are not commonly used in map layouts unless appropriate for the intended audience, such as geomorphologists.

Combining DEMs and Hillshades

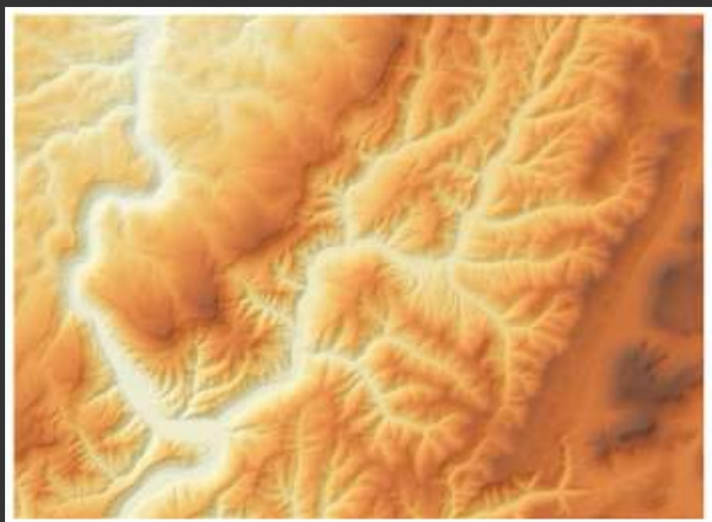

- Hillshade with 70% transparency
- DEM with 0% transparency

45

It is common to combine DEMs and hillshades to generate a textured and colorized terrain surface. These are often referred to as hypsometrically-tinted hillshades. Since both datasets are raster grids, this requires the use of transparency. In this example, the hillshade uses a black-to-white color ramp and has a 70% transparency applied. It is displayed above the DEM, which has an orange/red color ramp applied and no transparency. Determining the optimal amount of transparency will require some experimentation.

Combining DEMs and Hillshades

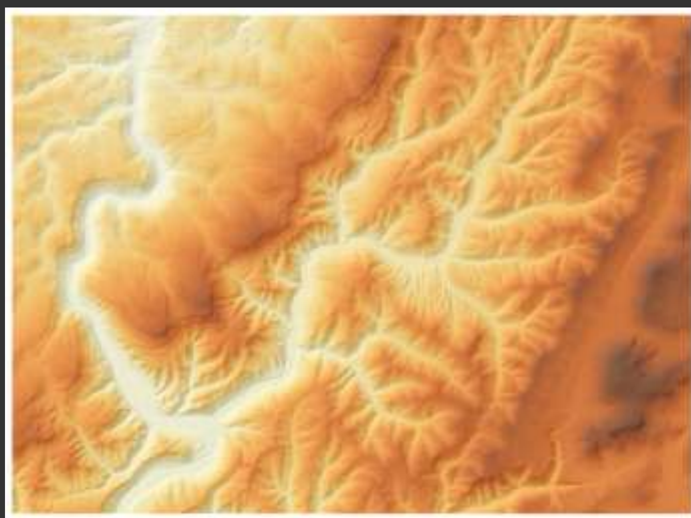

- DEM with 30% transparency
- Hillshade with 0% transparency

46

In this example, I have used the same color ramps but reversed the drawing order and applied a 30% transparency to the DEM data.

Combining DEMs, Hillshades, and Contours

47

In this example, I combine the DEM, hillside, and contour lines. The color binning applied to the DEM is designed to align with the principal contours, and transparency is used to add texture with the hillshade.

Surface Curvature

- Concavity/convexity of a surface
- Profile Curvature = curvature in direction of slope
- Planform Curvature = curvature perpendicular to direction of slope
- Can be used to highlight ridges vs. valleys
- Valley = dark
- Ridges = bright
- Black-to-white sequence

Negative = Concave
Positive = Convex
Near Zero = Flat

48

Surface curvature can also be applied to highlight ridges and valleys. Negative values indicate concavity, positive values indicate convexity, and zero indicates near flat terrain. So, applying a dark-to-bright color ramp will make valley bottoms appear dark and ridges appear bright.

Slope Position

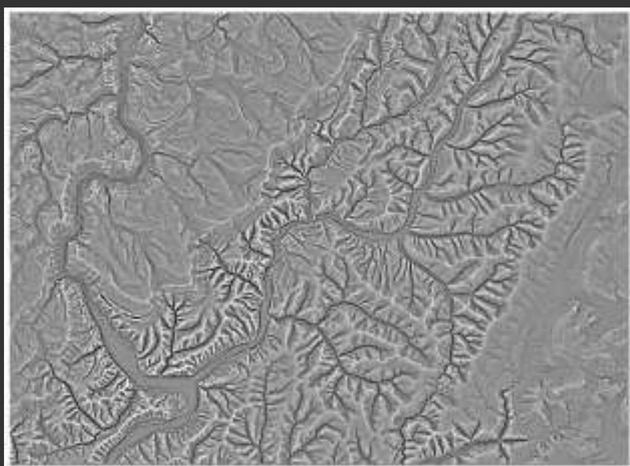

- Elevation at cells compared to elevation of surrounding cells
- Must calculate mean using neighborhood analysis
- “DEM” – “Mean”
- Ridge vs. valley

49

As an alternative to a surface curvature data layer, slope position can be used to obtain the same effect. Here, the mean elevation value in a moving window is subtracted from the center value. This relies on a moving window analysis, so the shape and size of the window can have a large impact on the output. Larger window sizes will generate more generalized patterns whereas smaller window sizes will capture more local variability. The appropriate window size is case-specific and will require some experimentation. In the output, high values indicate locations where the center cell is higher than the average of its neighbors. These are generally topographic high points, like ridges. In contrast, low values indicate areas where cells are lower than their neighbors, and would represent topographic low points or valleys. So, the result is similar to curvature, but calculated differently.

Combining DEMs, Curvature, and Hillshades

- Curvature with 70% transparency
- Hillshade with 40% transparency
- DEM with 0% transparency

50

This last example combines surface curvature, a hillshade, and a DEM. Transparency is applied to both the curvature surface and hillshade. The result highlights ridges and valleys due to the incorporation of curvature.

This is the end of this lecture module.

Please return to the West Virginia View Webpage for additional content.

Thanks! Hope you found this useful.