

Slide 1 – Title Slide

Hello and welcome to Week 11, part 4 of EGM310: Image decomposition and transformation. In this lesson, we will learn about the ways that we can transform our image data

Slide 2 – Image Analysis

In the last lesson, we learned about how we can visually interpret images in order to identify the objects in them. We've also spent some time talking about false-color composites, where we combine different bands to highlight different objects to make them easier to identify. But, RGB images, or even grayscale images, are not the only tools we have at our disposal. For instance, you might know that RGB is only one example of a color system – it is far from the only one. You may remember learning about the primary colors in school: Red, Yellow, and Blue. most printers use cyan, yellow, magenta, and black, or CYMK, to represent color. We also have something called Intensity, Hue, and Saturation – this is common for a lot of digital art and computer graphics, and it can be especially useful for remote sensing. In addition to transforming images into different color systems, we can also manipulate images to exploit spectral reflectance patterns, or to try to maximize the differences between different bands.

Slide 3 – Intensity, Hue, Saturation

Intensity, Hue, and Saturation or IHS is also known as hue, saturation, and lightness. Rather than choosing red, green, and blue components, it breaks a color value down as follows: the Hue component is the dominant wavelength, or what we would recognize as the color of the pixel, usually given as a value from 0 to 360 degrees – as you can see here, the color changes as we move around the color wheel here. Saturation refers to how much white is mixed with the color. As we move out from the center of the color wheel, we decrease the saturation. Finally, we have intensity, or lightness – this is how bright the pixel appears, ranging from 0 meaning totally dark, or black, and 1 or 100, representing totally light, or white.

Slide 4 – Example: IHS Fusion

One application that we can use for the IHS transformation is something called IHS Fusion. With this, we use high-resolution pan-chromatic data to 'sharpen' or increase the resolution of, the multispectral data. This is what was used to increase the apparent resolution of the image that you used for Practical 2, to take the 30 m Landsat multispectral data and resample it to 15 m. The way that we proceed is as follows: first, we re-sample the multispectral data to the same spatial resolution as the Panchromatic data. This doesn't actually increase the spatial resolution, it just changes the pixel size. Next, we transform the multispectral RGB image to IHS – we have the intensity component shown here. After that, we swap out the low-resolution intensity component for the high-resolution panchromatic image. Finally, we transform the swapped IHS image back to RGB – you can see here how the RGB image appears much sharper – we can see a lot more detail here in this image. This is because a lot of the

detail that we can see actually comes from the variations in brightness, while we can make do with less detail for the color.

Slide 5 – Band correlation

One thing that you might notice, looking at this graph, is that a lot of neighboring bands are correlated. That is, many surfaces, or objects, have similar reflectances in nearby wavelengths. This means that we often have a lot of redundant information in different bands. This can also make distinguishing different surfaces more difficult, as they will tend to look similar. For example, snow and clouds are quite similar in a number of wavelengths, distinguishing between different kinds of vegetation can be quite hard, as color differences can be quite subtle. To try to get around this, we have some techniques that we can use to help maximize the differences between bands, making it easier to extract information from an image.

Slide 6 – Principal Component Transformation

As an example, I have an image here showing two different landsat bands, the visible blue band on the top, and the visible green band on the bottom. The plot here shows a scatter plot of the pixel values for each band – the value in the green band on the x-axis, and the value in the blue band on the y-axis. And you can see that most of the values fall pretty close to this one-to-one line here. What this means is, most of the variation depends on how far along this line we are – we don't necessarily gain any extra information by having both bands. The principal component transformation looks to take advantage of this. We start by essentially finding the line that best fits the data – again, this is basically boiling down the most important information in the image, what we call the first principal component. After that, we can calculate the second principal component perpendicular to this line. And, we might want to shift the axes slightly so that the minimum value along the second principal component is 0. Because of the way that we have constructed this, these two bands are completely uncorrelated – we have maximized the differences between the two bands, so we also don't have any redundancy in the image. This is a bit of a toy example – we normally do this with all of the bands of a multispectral image, but this process is harder to visualize.

Slide 7 – Principal components

When we do a principal component decomposition, or transformation, on our original image, this is the result. For Landsat data, as we're using here, the first principal component, PC1, consists mostly of information from the near and shortwave infrared bands, while PC2 consists mostly of information from the visible bands – you can see this because water appears significantly brighter than the rest of the image. PC3 is mostly from the near-infrared – you can see that vegetation appears brighter than water, for example. As we go further down the principal components, we have less variation, less information – most of the variance, or the information, in the original 6-band image (3 visible bands, near infrared, 2 shortwave infrared bands) is contained within the first 3 or 4 bands – after that, we're adding less new information with each band.

Slide 8 – Principal Component Transformation

We can combine the first three principal components to create something called a “decorrelation stretch”, shown here on the right. The image on the left is the true-color composite. The decorrelation stretch is a way to increase the distribution of color within an image, while preserving the relative characteristics of the original image. And what we can see here is – most of the vegetation surfaces are a bright magenta color, indicating that their highest values are in PC1. Water is a bright green color, indicating that water surfaces have a significantly higher value in PC2 than in other bands. Most of the rest of the image appears to be shades of red/orange or green, with some small regions of dark blue. I’ll post larger versions of these images on blackboard – if you like, you can have a look at them and see if you can’t figure out what they might represent.

Slide 9 – Band Algebra

In addition to different kinds of transformations, we can also use bands algebraically – we can add them together, subtract, multiply, or divide them. These operations can all enhance the differences between different bands, helping us to identify surfaces, or convey other information, more easily than by using the reflectance values on their own. For example, we know that for healthy, chlorophyll-producing plants, the reflection in the near-infrared is significantly higher than in the visible red. But, you might also notice that for most other surfaces, the reflectance in the near-infrared and visible red is quite similar, or it even reflectance drops – for water or weathered basalt, it’s nearly the same; for concrete, it’s perhaps a slight increase; for snow, it’s a clear decrease. So, if we were to look at the difference between the near infrared and red bands, we might expect that vegetation will have strongly positive values, while most other surfaces will have slightly positive values, or even negative values. But, the details might change depending on the illumination in the scene. To help counteract this, we can do something called normalizing the data – dividing by the total reflectance in the two bands, which makes sure that the comparison is relative to how bright the scene is.

Slide 10 – Normalized Difference Vegetation Index (NDVI)

This leads us to an extremely common method for studying the health of vegetation, or to help identify vegetation, called the normalized difference vegetation index, or NDVI. The NDVI is calculated as follows: we subtract the red band from the near infrared band, then divide by the sum of the two bands. The NDVI then consists of values between -1 and 1. Values much greater than 0 typically indicate healthy, chlorophyll-producing vegetation, while values less than 0 typically indicate something else – either distressed vegetation, or some other surface like clouds, snow and ice; or soils. We can use the NDVI to help map healthy vegetation, or to help derive other information about vegetation, such as the area within a given pixel that represents plant leaves, or the chlorophyll concentration in leaves.

Slide 11 – Normalized Difference Water Index (NDWI)

Similar to the NDVI for vegetation, we can calculate a normalized difference index for water. The NDWI makes use of the fact that for water, the reflectance in visible green wavelengths is much higher

than in the near infrared. Thus, we calculate the NDWI by subtracting the near-infrared band from the green band, and dividing by their sum. The NDWI is used for, as you might have guessed, automatically mapping water bodies. We can also use it for flood detection and mapping in images – especially if we have an image from before a flood event, large changes in NDWI most likely indicate the presence of water. An alternative version of the NDWI, also called the NDWI, uses the near-infrared and shortwave infrared bands. This version is used to estimate the water content in plant leaves, as the reflectance in the shortwave infrared is more heavily affected by water content than in the near infrared.

Slide 12 – Summary

In this lesson, we have learned that by transforming image data, we can make more use of the information contained in the image.

For example, we can use it to sharpen multispectral images, by first transforming the image into intensity, hue, and saturation colorspace.

We can also use these techniques to improve the spectral differences between surface types, and remove redundant information from bands that are close to each other.

By using arithmetic operations, we can enhance differences in spectral reflectances between different surfaces, helping us to better identify them.

Slide 13 – Additional resources

Once again, you can read more about the concepts we've covered in this lesson in the textbooks, Chapter 7 of Lillesand, Kiefer & Chipman; and Chapter 11 of Campbell & Wynne. You can also read more about some of these topics in the Remote Sensing Tutorials from Natural Resources Canada. For more information about the HSL (or IHS) colorspace, I've linked to a video from Khan Academy; and I've included a link to a video about working with the NDVI in ArcMap. That's all for this lesson – I hope you found it interesting, and if you have any questions, please don't hesitate to e-mail me or post in the discussion forum on blackboard. Bye!