

## Slide 1 – Title Slide

Hello and welcome to Week 2, Part 2 of EGM703: Hyperspectral Image Analysis. In this lesson, we'll learn about what makes hyperspectral data different from most of the datasets we've worked with until now, along with some of the methods that we can use to analyse it.

## Slide 2 – Hyperspectral data

Remember that hyperspectral sensors, or instruments, are characterized by narrow, continuous bands – the resulting images can be used to create spectral signatures that we can use to directly identify surfaces or materials. Some of the applications that we can use these images to study are: surface mineralogy, water quality, crop or vegetation health, soil type, or even other properties. So, this should be easy, right?

## Slide 3 – Mo' data, mo' problems

The good thing about hyperspectral data or images is that they have lots of bands. This means that they provide more information about the physical and chemical composition of materials; as I've said a few times, it means that we can identify spectra directly; it also means that we have more information about atmospheric characteristics, since hyperspectral sensors normally acquire data at wavelengths that don't correspond to atmospheric windows, unlike most multispectral sensors. On the other hand, the bad thing about hyperspectral images is that they have lots of bands. This means that they have a comparatively high volume of data – as an example, the Landsat image shown here has only 11 bands, with a zipped size of about 1 GB. The EO-1 Hyperion scene shown in black and white and outlined here in blue, has a zipped size of about 250 MB, despite having the same spatial resolution as most of the Landsat bands, 30 m – this is because it has about 20 times more bands than the Landsat image. Like we've seen with multispectral data, there can be a lot of redundant information in adjacent bands, meaning that we have a high volume of data that doesn't necessarily give us more information. Many of these bands also have a lower signal to noise, or SNR, ratio – meaning that it can be more difficult to extract the information that we want, because there's more noise in the images. Because we're normally comparing these images to spectral signatures collected in the field or laboratory, atmospheric correction is also comparatively more important than it is for multispectral images. All of these things together mean that while hyperspectral remote sensing is an extremely powerful tool for many applications, it is not without significant challenges.

## Slide 4 – Dimensionality reduction

As mentioned on the previous slide, and as we covered in EGM702, the reflectance values in adjacent bands are often highly correlated, meaning that we end up with a lot of redundant information. This redundancy means that the “normal” classification methods we've seen, such as maximum likelihood, are significantly more difficult to apply to hyperspectral images directly. Because of the high number of bands, it's also computationally more expensive, since we need significantly more training samples

than we do with multispectral images. A general rule of thumb is that you need at least 10 times, and often 100 times, the number of bands worth of pixels per training sample – for an image with over 200 bands, this means we need at least 2000 to 20000 pixels per training sample. One approach that we can take is to reduce the dimensionality of data – we saw this in EGM702 with the principal component analysis, where we reduced an image with around 10 bands to 3-4 bands that contained most of the information in the image. For hyperspectral images, a common approach is using a similar method called the minimum noise fraction, or MNF. With this approach, we can take an image with over 200 bands, many of them redundant or noisy, and reduce it to around 10-20 bands that contain most of the information in the image.

## **Slide 5 – Spectral Angle Mapping**

We can represent each pixel in a hyperspectral image with  $n$  bands as an  $n$ -dimensional vector. The simple example here is for a two-dimensional vector, but we can generalize this approach to as many dimensions, or bands, as we like. In this approach, each pixel represents a test spectrum that we are trying to identify using a reference spectrum. We can use any number of reference spectra to compare to – we're only really limited here by the size of our reference library. We can calculate the angle,  $\alpha$ , between these two vectors  $t$  and  $r$  as the inverse cosine of the dot product of  $t$  and  $r$ , divided by the product of the magnitude of the two vectors. We then assign each of our pixels to a given reference spectrum, or "end member," based on the smallest angle between our test spectrum and our set of reference spectra. This approach allows us to directly compare lab or field-derived spectra and our atmospherically-corrected image. For more on this approach, I've included a link at the end of the lesson to this 1993 paper by Kruse and others that describes an application of spectral angle mapping.

## **Slide 6 – Spectral mixture analysis**

Remember that for most satellite images that we are going to be working with, each pixel represents multiple objects or surfaces that are smaller than the spatial resolution of our sensor. In the example here, our pixel contains grass, a few trees, some bushes, a sidewalk, and a house with a red tile roof. The "color" of the pixel is a linear combination of the "color," or reflectance, of each of these subpixel objects or surfaces. For example, if we were to take each of these colors and blend them together, we would probably end up with something that looks fairly green, like this. The same is true for the reflectance spectrum of the pixel – it's a linear combination of each of the individual spectra that we see here. In spectral mixture analysis, we use the measured spectra of each of these potential "end members" – concrete, grass, oak trees, roof tiles, agave bushes, etc. - to compute the percent coverage in each pixel of each of the different end members. Like with spectral angle mapping, this approach requires that we have identified our potential end members beforehand, so it helps to have some prior knowledge about the study area.

## **Slide 7 – Cross-correlogram spectral matching (CCSM)**

Another approach that attempts to identify the surfaces or materials present in each pixel is known as cross-correlogram spectral matching, or CCSM. With this method, we are actually computing the cross-

correlation, or similarity, between a reference spectrum and the spectrum for a given pixel. In the example here, we have three potential end member minerals, kaolinite, alunite, and buddingtonite, with the corresponding spectral signatures shown here on the left. With CCSM, we compute the cross-correlation between each spectrum directly, but we also shift the reference spectrum to higher and lower channels, or “match positions” – this gives us an idea of where the actual “peak” correlation is, how significant that peak is, and how skewed the correlation curve is – all of this helps protect somewhat against false matches. A higher correlation value, especially at match position 0, means that we have a better match between the test and reference spectrum – the example here, for kaolinite, shows that we have a high match (correlation = 1) for kaolinite, as we would expect – we can also see where the peaks for the other two end members are here. This approach finds the “best” match between our end members and the observed spectra – again, we’re only going to be limited here by the end member spectra that we’re using. These examples show a false-color image of the correlation, skewness, and significance for each pixel in the AVIRIS image and the given mineral, with Alunite on the left and kaolinite on the right. White pixels indicate a better match to the reference mineral – we can see that most of our Alunite is located more in the upper center and right portion of the image, while Kaolinite is located in more of a ring around these features here. The sides of these cubes indicate the correlation values for the top row and right-most column of the image at the different match positions, in order to show how the correlation changes with match position.

## Slide 8 – Continuum removal

One problem that we will often encounter is that a number of spectra are very similar. Another issue is that the albedo, which partly depends on the viewing angle and incident radiation, has a big influence on how deep the absorption features in a given spectral signature are. For example, in the image on the left here, we can see that the red absorption feature is quite small – only about 0.05-0.1. To help accentuate these absorption features and aid in identifying surfaces from their spectra, we can divide the reflectance values by the continuum value – basically, we take the smoothed-out shape of the curve, often the convex hull of the data – as a result, we end up with the curve on the right, where we’ve now emphasized the absorption features in the curve, which can make it easier to compare curves, especially if we have variable illumination. We can use this in conjunction with the methods we’ve just covered, such as spectral angle mapping or cross-correlogram spectral matching, where it can help improve our results.

## Slide 9 – Derivative analysis

In a similar vein, we have a technique known as derivative analysis, which is quite prevalent in lab-based spectroscopy. As mentioned on the previous slide, natural illumination can vary significantly, due to factors such as a different sun angle, cloud cover, or even topography. This changes the reflectance curves that we measure, and can make identification more challenging. Derivative analysis basically does what it says on the tin. By taking derivatives of the reflectance curve, we can more easily identify the location of different component features. The synthetic example here shows a combination of six different curves – in the combined curve, it’s difficult to identify where any one of these individual peaks is, meaning it will be more difficult to identify the individual components. However, once we

start taking derivatives of these curves, it's much easier to identify where these component peaks are – they show up as maxima, minima, or roots in the derivative curves. So, with derivative analysis, we can enhance absorption features, because the derivatives are typically insensitive to the variation in illumination. The real-world example here with reflectance spectra of soybean plants treated with manganese shows significant variation in reflectance depending on the illumination. But, once we start taking derivatives, we see that these differences are minimized. To help dampen the noise that is often present in the curves, we can also apply different smoothing techniques – the paper referenced here looked at a number of different techniques to aid in the derivative analysis.

## **Slide 10 – Spectral libraries**

Each of these different approaches that we've seen requires us to have a set of reference spectra to be able to compare our data to. One way that we can get reference spectra is by making our own lab or field measurements – of course, this assumes that we have access to the necessary equipment, which is expensive; we might also not be able to make it out to our field site very easily, which means that collecting our own data might be more difficult. Fortunately, there are a number of spectral libraries available. We've seen some of these already, such as the MODIS UCSB emissivity library or the ECOSTRESS spectral library, which combines a few older libraries, including the Johns Hopkins University spectral library, the Jet Propulsion Laboratory spectral library, and the USGS spectral library. These spectral libraries are a combination of both lab-derived, field-measured, and remotely-sensed spectra. They're a great resource for a lot of remote sensing applications, but especially for hyperspectral remote sensing. Additionally, some software packages include spectral libraries, making it relatively easy to include them in your workflow.

## **Slide 11 – Summary**

In this lesson, we've seen how hyperspectral images provide a lot of detailed information about reflectance spectra, which helps us to identify features in the images.

This is not without challenges, however, since the large data volume and low signal-to-noise ratio in many bands can make hyperspectral images difficult to work with.

Fortunately, there are a number of techniques we've covered so far that can help us mitigate these challenges and analyze hyperspectral images.

## **Slide 12 – Additional resources**

You can read more about the topics we've discussed here in the textbooks – Jensen, Chapter 11, Lillesand, Kiefer & Chipman, Chapter 7.21, or Campbell & Wynne, Chapter 15. I've included links to articles referenced in this lesson, including the Green et al. paper on minimum noise fraction reduction, the Kruse et al. paper that covers spectral angle mapping, the van der Meer and Bakker paper on cross-correlogram spectral matching, and the Tsai and Philpot paper on derivative analysis in hyperspectral remote sensing. That's all for this lesson – I hope you found it interesting, and you have any questions, please don't hesitate to e-mail me or post in the discussion forum on blackboard. Bye!