

Slide 1 – Title Slide

Hello and welcome to Week 12, part 1 of EGM310: Change Detection. In this lesson, we will learn about how we can use satellite images acquired at different times to observe and quantify changes.

Slide 2 – Week 12 Outline

In the rest of the lessons this week, we'll cover applications of remote sensing to different topics, including studies of water resources, surface motion, and archaeology.

Slide 3 – Mapping Change

Very often, we're not just interested in making a single observation – we want to be able to observe and quantify changes in landcover, or other variables. In this image, we see a normalized difference vegetation index for the Leeward Islands – remember that this indicates the overall health of vegetation: the higher the value, the healthier the vegetation. These are a group of islands situated where the Caribbean Sea meets the Atlantic Ocean, including the countries of Montserrat, Antigua and Barbuda, and St Kitts & Nevis. In September 2017, a Category 5 hurricane, Hurricane Irma, swept through the Caribbean. It was the strongest hurricane to ever hit the Leeward Islands, and it caused immense damage. The eyewall of the storm passed over Barbuda. 95% of the buildings on the island were damaged or destroyed, and the island was effectively stripped of its vegetation. You can see this in the post-Hurricane NDVI image – on a number of the islands, the NDVI values have dropped considerably. As Saint Kitts and Nevis was hit less directly, we see more or less the same picture from before to after.

Slide 4 – Change detection

Our goal with change detection is to analyze changes over time. for example, we might want to study the damage caused by a hurricane, or map the growth of urban areas, or other changes in landcover. We might want to understand how forest fires have changed over time, study landslides or other hazards, or any number of applications. We also want to have some kind of quantitative analysis – for example, the percent area change of glaciers, or the number of landslides. We may also want to understand the spatial distribution of these different things – for example, where do we see the largest changes in urban area? We usually also want some kind of accuracy assessment – to understand how accurate our results are, and gauge how reliable the changes that we observe are.

Slide 5 – Change detection considerations

When doing change detection, we have a number of considerations. The first is time – this can be the time period considered, the rate of change – how much time has to pass before we can expect to see any changes; or the time of year – is there a seasonal aspect to the change we want to see? For example, we wouldn't want to analyze vegetation changes by taking a winter image from one year and a summer

image from another year – that probably wouldn't give us a meaningful change. We also have to consider our remote sensing data – starting with the various resolutions we have discussed. Is the change we're trying to observe of a scale that we could actually see with our data? We wouldn't expect to see centimeter-scale changes with 30 m resolution imagery. We also have to think about the temporal resolution – for periodic changes, we need to have images frequently enough to actually capture change. And, depending on the radiometric resolution of our data, we may not be able to distinguish subtle changes in reflectance. Finally, we also want our images to be geometrically and radiometrically corrected – if they aren't, we can't be sure whether the changes we see are real, or due to differences between the images.

Slide 6 – Some change detection techniques

We have a number of different techniques at our disposal. Which technique we use will depend on the changes we are trying to see and the data that we have available. We'll start with visual analysis – this can be viewing two images side-by-side, for example. We often use visual analysis to help us determine which technique might be most appropriate, given the change and the data available. Of course, if we're doing any kind of manual mapping or classification, we might also use this on its own. The next category of technique we might consider is data transformation. This might be a principal component analysis – much like with image analysis or classification, this technique can help us reduce redundancy in the data and highlight variations or differences. Also like with image analysis, this can be difficult to interpret – it's not always clear what the principal components relate to. We can also create normalized difference indices, either of individual bands from two different points in time, or using band ratios, normalized difference indices, or other metrics. Band algebra, such as taking a band ratio or a difference, can also be a way to enhance differences between two different times. We could do a classification comparison, where we classify each image individually and compare the results. This may not always be the best option, however – classifications can be noisy, and so any differences may not reflect true changes. Change vector analysis helps us identify spectral changes based on the magnitude and direction of differences. Multi-temporal classification, unlike classification comparison, seeks to classify changes – for example, a change from bare soil to crops, or from forest to urban areas. Finally, with a sufficient number and density of observations, we can analyze the data as a time series – for example, here looking at a time series of glacier elevations.

Slide 7 – Dynamic visualisation

Dynamic visualization refers to somehow viewing the two images dynamically. For example, this could be using the swipe tool in our GIS software, flickering back and forth between the images, or with an animation as seen here. This shows a large landslide that occurred on 16 February 2014, on Mt La Perouse in Alaska. Dynamic visualization can be used to select the best technique given the data and the changes; it can also be used to perform the analysis, using the the visual interpretation techniques we discussed last week.

Slide 8 – Multitemporal false colour composites

Another type of visual analysis could be to form multi-temporal false color composites. With this, we composite bands from multiple dates. Areas without change should be shades of gray, black, or white, while changes should appear in color, depending on how we composite the image. In this example, I've set an image from March 2 to the red channel, and the green and blue channels are an image from February 7. In this case, red pixels correspond to shortening shadows – this is where the March 2 image is brighter than the February 7 image. Blue areas are the landslide, and white/gray pixels are areas with no change. We can also see some areas corresponding to glacier motion – this area down in the lower left corner of the image, for example. Here, crevasse motion leads to contrasting areas of red and blue.

Slide 9 – Image algebra

Much like single image analysis, we can use image algebra to enhance the differences between our two images. On the top here, we have a band ratio, where I've divided the February 7 image by the March 2 image. Pixel values close to 1 then correspond to no change, while pixel values below 1 correspond to areas where the March 2 image is brighter, and pixels values above 1 correspond to areas where the March 2 image is darker, and you can see the rock avalanche highlighted clearly. Below this, we have a normalized difference index, where I've taken the March 2 image minus the February 7 image divided by the sum of the 2. For this kind of change analysis, we usually have to pick a threshold value to classify where significant change has happened – we might pick a difference of ~ -0.15 for the normalized difference image, and maybe something like 2 for the band ratio. We can also do algebra with arithmetic images – as an example, a difference of differences can show us changes in multiple bands at once. Thinking back to the first slide of this lesson, a decrease in NDVI can show us where large changes in vegetation cover have occurred due to a hurricane, for example.

Slide 10 – Multi-temporal Classification

For a multi-temporal classification, we perform the classification using stacked images – all of the bands from time 1 and all of the bands from time 2 are combined in a single image. Unlike with single-image classification, the goal with multi-temporal classification is to classify change classes and non-change classes. We can illustrate this with a relatively simple example. Here, we have a scatter plot here for a given band, with the values at time 1 plotted on the x-axis, and the values at time 2 plotted on the y-axis. You can see that we have roughly 4 groups here that are clearly separated from each other. The groups that fall close to the 1-to-1 line are our non-change classes – they have roughly the same values in each image. And perhaps the low-value group represents grass – maybe this is the reflectance in the visible red band, so grass doesn't appear very bright. And we might also classify the high-value non-change group as snow – it has high values in both images. So now, we can see that we have a group of pixels that had high reflectance values at time 1, and low values at time 2 – so this could represent pixels that were covered with snow that melted in between time 1 and time 2, exposing the grass underneath. We also have the opposite – pixels that had low reflectance values at time 1, and high values at time 2 – these would probably represent pixels that are grass that was covered with snow in between time 1 and time 2. When we're doing this sort of classification, one thing we have to think

about is: how similar are our “change” classes to our “non-change” classes? This is going to depend on the image we have – images that only have visible bands might be difficult to distinguish, given how similar many surfaces appear in visible light. It will also depend on the type of change we’re trying to see, as well – if the changes are subtle, it might not be easy to distinguish them reliably, just like when we’re trying to classify single images. We can use principal component analysis to help with this. It will help reduce the redundancy in the image, as well as the dimensionality – we might be able to get by with only 3-4 bands from each time, rather than the full complement of 10-11 bands. We can also use PCA to do change detection – it will help maximize the variation that we see, but remember that this can be difficult to interpret.

Slide 11 – Change vector analysis

If we think about band values as forming vectors in space, we can analyse the differences between them. The example here shows the red brightness plotted along the x-axis, and the near-infrared brightness plotted along the y-axis. The longer vector represents the pixel values at time 1, and the shorter vector represents the pixel values at time 2. We can subtract the time 1 vector from the time 2 vector to get the change vector – in this case, the change in red brightness (x) is +12, and the change in nir brightness (y) is -39 – perhaps this represents a change in vegetation that has become stressed due to drought. We’ve shown an example in just 2 bands, but we can do this for as many bands as we like. This vector has an angle (-63.5 degrees, or about 300 degrees going anti-clockwise) and a magnitude (40.8). The magnitude represents how much change has occurred – how different the pixel is at time 2 when compared to time 1; while the angle can be interpreted to determine what kind of change has happened. We can threshold the magnitudes for significance – we don’t necessarily expect small changes to represent true changes, as they can be due to many different factors such as changes in scene brightness. We can also categorize the angles, similar to the cardinal directions on a compass. In this two-dimensional case, we might have 4 categories, or quadrants: where both components are positive, both negative, and then the two cases where one component is positive and one negative. This is an example for 2 bands, but we can do this for as many bands as we like – the more we add, the easier it can be to categorize and interpret the angles.

Slide 12 – Change vector analysis

An example here, taken from an article on remote-sensing dot org, shows the change vector magnitude and angle for two different images based on the red and near infrared bands, one from 1988 and the other from 2011. Here, you can see 3 different examples of change: the first represents no change – we had a forest pixel in 1988, and we still had forest in 2011. Because the pixel values are the same in 1988 and 2011, the angle and magnitude of change is of course 0. If we have deforestation, we might see a significant drop in near infrared reflectance between 1988 and 2011, and perhaps a slight positive change in red reflectance. The last change, from soil to water, perhaps indicative of flooding or the filling of a reservoir due to construction of a dam, is characterized by a drop in near-infrared reflectance, and a slight negative change in red reflectance.

Slide 13 – Summary

In this lesson, we have discussed how often, our goal is to analyse changes in time. This could be before and after a particular event, such as a landslide or a hurricane, or it could be more generally over time, using two or multiple observations. Visual analysis of change is often a first step toward determining which technique might work best for detecting the kinds of changes we are interested in studying, or it could be a way to analyze change on its own. Many of the techniques we've already covered, including principal component analysis and image algebra, can be adapted to analyze temporal changes as well. They can also be combined in a way to enhance the differences between images acquired at different points in time, similar to how we analyze individual images. Finally, the choice of technique is going to depend on our particular application and goal – some techniques are better suited to certain kinds of changes and available data.

Slide 14 – 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 16 of Campbell & Wynne. I've included links to two different websites – the first, Earth Observatory, has a number of articles and examples showing changes on the Earth's surface as observed by Landsat. The second website will enable you to make your own GIF using landsat images and Google Earth Engine – no coding required. Finally, there are also two videos showing examples of performing change detection using ArcMap and Landsat images. 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!