Slide 1 – Title Slide

Hello and welcome to Week 3, Part 4 of EGM703: Synthetic Aperture Radar. In this lesson, we'll introduce the principles of synthetic aperture radar, which is how we create radar images of the Earth's surface from space.

Slide 2 – Constructing a radar image

In the previous lesson, we showed how side-looking radars receive signals from objects at different slant ranges, which means we can effectively create one line of an image. Now we'll show how we construct the rest of the image, bit by bit. Here, we have our radar image, with the range and azimuth shown like this. We have our sensor flying at a height R and moving at a speed V in this direction, and it's creating a beam shown by the gray outline here. When the satellite gets the return from the ground here, we'll see a signal from the house here, because it's falling within the beam sent out from this location. And the width of that signal is going to be determined by our range resolution. As the satellite moves along, it will continue getting a return from the house, because the house is still within the beam, even though the satellite isn't yet at the house's azimuth location. And we can see how this will continue – the satellite will continue to measure a return from the house until the house is no longer in the satellite's beam. The number of lines in the image where we see a return from an object is related to the azimuth resolution of the sensor – as you might have guessed, this is effectively the beam width of our sensor on the ground.

Slide 3 – Azimuth resolution

So, how do we calculate the beam width? It turns out that the beam width, or the azimuth resolution, is related to the length of our antenna, L, and the distance to our target, R. It also depends on the wavelength of the signal being sent out. Using approximate values for the ERS satellite of 10 m for L, 850 km or so for R, and 5 cm for λ, we have a ground beam width of approximately 5 km. If we want to improve our azimuth resolution, we either need to decrease the wavelength, which isn't really something that we want to do, or we need to increase the antenna size. For example, using the same R and λ values, to get an azimuth resolution of about 10 m, we would need to have an antenna that is 5 km long. Hopefully, it's clear to you why this is not actually a workable solution to our problem – we need to think of something else.

Slide 4 – Putting the "synthetic" in SAR

Let's have another look at how we build our radar image in the azimuth direction. At this location here, our signal doesn't actually return a signal from the house – the house isn't within the beam. At this location, though, it is – this is the first point where we get a return from the house. And, we get another return here, and here, and so on, until we get to here, where the house is no longer in the beam. Remember also that at each point, the satellite is going to record signals that scatter from directly beneath it, but also from earlier locations. Because the platform is moving relative to the ground, these returns are going to be Doppler shifted – their frequencies will change, for the same reason that the sound of the ambulance siren changes as it drives toward or away from you. If the return is scattered from an area ahead of the platform, the frequency will be shifted higher, and if the return is scattered from an area behind the platform, the frequency will be shifted lower. So, if we process the returns based on the Doppler shift, then we can effectively treat all of these different returns as if they were received by one large antenna – we're creating a synthetic antenna that's much larger than our actual antenna. It turns out that if we do this, the azimuth resolution of our sensor is just half of the length of the antenna – it doesn't matter how far away the target is, our ability to distinguish between targets in the azimuth direction is only limited by the size of our antenna. As you might imagine, this is really great news – just like we saw with range resolution, by processing the signals in a clever way, we end up improving the resolution of our system without having to change the details of our system.

Slide 5 – Focus!

Unfortunately, when the signal measures these different returns, the resulting image looks like noise – we have many different overlapping returns that interfere with each other in a mostly random way. As with many things in life, if we want to try to make sense of it, we have to focus – in this case, we have two different steps. The first is one we've seen before: range compression, where we correlate, or convolve, the measured signal with the signal sent out over the range direction of the image. The second is very similar, known as azimuth compression. Here, we're correcting the Doppler shifts in the azimuth direction. Because we typically end up with a higher azimuth resolution than range resolution, the image will appear stretched out in the azimuth direction. We can also use a technique called multilooking, where we effectively average together a given number of azimuth pixels, in order to end up with an approximately square image.

Slide 6 – From raw images to multilook images

Here, I've used the python SAR processing tutorial from EO College, linked here and at the end of the presentation, which provides a nice demonstration of how this works. Starting from our raw, unfocused image, we first have to range-compress the image, resulting in something that looks like this. In this image, we can see a number of different pixel patterns that look like they've been smeared out in the azimuth direction here – think back to the first slide from this lesson, where the house showed up in mutiple different lines owing to the width of the beam on the ground. If we azimuth-compress the image, we end up with something that looks a bit more like we might expect – we can see that this is some kind of ground image, though it looks like it's been stretched out in the azimuth direction, owing to the different resolutions of the sensor. As a final step, we can average over every 5 pixels or so in the azimuth direction, to get an image that looks more like what we might be expecting, with pixels that are approximately square.

Slide 7 – Geometric distortion

Unfortunately, even after we've gone through these different steps, we're still going to have distortions in the image owing to the fact that our radar is side-looking.

The first kind of distortion that we'll discuss is known as foreshortening, and the diagram here shows how this happens – we have our satellite up here, a small mountain on the ground, the satellite beam in blue, and the slant range shown in red here. Remember that the wavefront of the signal has a circular pattern, at least in two dimensions. So point A is going to be seen at slant-range distance A', B at a slant range distance B', and C at a slant-range distance C'. And we can see here that the distance AB' in slant range is significantly shorter than the distance AB on the ground, because of the slope and the direction that it's facing – A and B are at similar distances to the sensor. We also tend to see bright pixel values associated with these slopes, because they're facing the sensor – they reflect a lot of energy directly back to the sensor.

The next kind of distortion that we'll talk about is layover, shown here. Layover is an extreme version of foreshortening, where A' is actually further away from the sensor than B', and so the top of the mountain effectively overlays the bottom of the mountain. This happens because the top of the mountain is actually closer to the sensor than the bottom. Like with foreshortening, we see bright pixel values associated with this type of distortion, for the same reason.

Finally, we have shadow, where the backslope is entirely hidden from the sensor by the top of the mountain. With this type of distortion, we have dark pixel values, because nothing returns to the sensor from the slope BC, or between points C and D.

Slide 8 – (Geometric) terrain correction

Once our image has been focused, it's still in radar geometry – that is to say, the pixel locations correspond to range and azimuth, rather than the location on the ground. If we want to compare the image to other remote sensing data, or use it in a GIS software, we have to georeference, or geocode, the image. When we download a SAR image, like for example the one we'll use in this week's practical, it has information about the satellite's location that we can use to do this, though we usually also want to have a DEM to help correct the different distortions we've discussed. This process works similarly to orthorectification of optical images – we're trying to take each pixel and map it to the correct location on the Earth's surface. We can see an example of the result of this here, courtesy of the Alaska Satellite Facility using data from the Japanese Aerospace Exploration Agency. The first image, the "before", shows what the image looks like when the pixels correspond to range/azimuth distance, rather than ground location. The second image shows how these distortions get corrected when we map them to the correct locations. We can see here that the slopes facing the sensor are bright – the satellite is flying in kind of a Southeast/Northwest direction here, looking toward the right. We can also see dark areas here in the northeast part of the image, where the steep canyon walls create shadows that are not seen by the sensor.

Slide 9 – (Radiometric) terrain correction

Remember that slopes that are facing the sensor are bright, because they efficiently reflect energy back toward the sensor. Slopes that are facing away from the sensor will appear dark, because they will reflect energy in a direction away from the sensor. If we want to study the actual properties of the ground using the backscatter, we often want to correct for this effect, a process known as radiometric terrain correction. We can see an example of the result of this here, courtesy of the Alaska Satellite Facility using data from the Japanese Aerospace Exploration Agency. The first image, the “before”, shows how slopes facing the sensor are bright, while slopes facing away from the sensor are dark. Again, the slopes facing the sensor are bright – the satellite is flying in kind of a Southeast/Northwest direction here, looking toward the right. In the radiometrically corrected image, we can see how many of these bright slopes appear darker, and some of the north-facing slopes appear brighter, but not all of them. The reason why some slopes still appear brighter than others is because of an actual difference in backscatter values – the south-facing slopes are primarily covered by deciduous trees such as birch, while the north-facing slopes are primarily covered by conifer species such as black spruce.

Slide 10 – Speckle

One thing that you will notice is that SAR images often look “noisy” – they have a salt-and-pepper, or speckle pattern, which we can see here in our example image from earlier. Although this looks like noise, it’s actually an interference pattern. Remember that the SAR sensor is recording both the amplitude and phase of the signal that is returned. Remember also that the phase of an individual pixel is the sum of all off the different subpixel scattering surfaces. This speckle pattern varies randomly over time, but if we were to average a number of scenes together, the measured backscatter will approximate the actual normalized radar cross-section of the pixel, assuming that the scenes remain coherent. Next week, we’ll see a few different examples of how we can use this speckle pattern to track the motion of surfaces over time.

Slide 11 – Summary

In this lesson, we’ve discussed how the azimuth resolution of a real-aperture radar image is restricted by the antenna size, which means that creating a real-aperture radar image from space is pretty much impossible.

The solution, as it often is, is to fake it until you make it! In this case, we simulate having a very large antenna, a “synthetic aperture”, which helps us get around this limitation.

With this, and the clever signal processing techniques we introduced in the third lesson, we can get high-resolution radar images from space.

Unfortunately, SAR images have a number of distortions that we’ll need to correct if we want to be able to use them – we’ll see how some of this works in this week’s practical on SAR image processing.

Slide 12 – Additional resources

You can read more about the topics we’ve discussed here in the textbooks – Lillesand, Kiefer & Chipman, Chapter 6.4, or Cambpell & Wynne, Chapter 7. I’ve included a link again to radartutorial.eu, which has a good overview of synthetic aperture radar. This 2013 article is a great tutorial on SAR, which covers the different steps and concepts in a bit more detail than we’ve gone into in the lecture. These two articles from NASA provide a good overview of SAR and different sensors, as well as a discussion of how we use SAR images to study other worlds, such as Saturn’s moon Titan. And finally, this video by Adrian Schubert is a funny but excellent explanation of how synthetic aperture radar works – I recommend giving it a watch if you have some time. 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!