Synthetic Aperture Radar Land Applications Tutorial

Part II Bilko Exercise

prepared by

Introduction

This SAR Land Applications Tutorial has three main components:

I. Background and Theory – an overview of the principles behind SAR remote sensing, data processing techniques, examples of land applications, and current and future sources of SAR data.

II. This Bilko Exercise – a computer practical using Bilko software with Enivsat-ASAR data, allowing students to apply the theoretical knowledge to the processing and interpretation of actual SAR data.

III. Answers and Examples – model answers to questions from both parts of the tutorial, including screenshots of what the students may expect to see when working with Bilko.

We recommend that you read through the Background and Theory before starting this practical. Chapter 1 and chapter 2, section 2.1 in particular provide necessary information for the successful completion of the exercise. You may also find it useful to refer to the theoretical part of the tutorial when you carry out the practical work.

Please note that the Bilko software is designed primarily for training purposes. Its strength lies in its transparency, which makes students consider the effects of different processing steps on the reliability and validity of remote sensing products. However, it lacks some of the advanced 'black-box' algorithms of more specialist software, and does not provide the facility for automated processing. For advanced processing of SAR data users should consider other software packages discussed in chapter 6 of the "Background and Theory".

Computer processing power and memory requirements

The ASAR images used in this lesson are large. Minimum recommended memory size for the successful completion of the exercise is

RAM:	512 MB
Virtual Memory:	1024-2048 MB

Note: The exercise has been run successfully on a machine with 256 MB RAM and a Virtual Memory setting of 768-1536 MB, but this required careful management of the number of open windows to avoid the system slowing down or freezing.

Using the module navigation tools

The exercise consists of several steps, each corresponding to a precessing step described in the section 2.1 of the theory section. While working with the exercise you may want to navigate around both the theory section and the exercise itself. The navigation tools at the bottom left of the pages to the left of the content lists are intended to help you move around between slides as easily as possible.

Bottom left navigation

Takes you to the main content slide; particularly useful for navigating around the theory section of the module

Takes you back to the slide you viewed previously, a useful way of returning to the content lists when you are trying to find a specific slide

Takes you to the previous slide in the series

Takes you to the next slide in the series

Navigation on contents pages

Takes you to the slide indicated by the adjacent entry in the contents list

Overview, aim and objectives

The Philippines have 2 crop season per year: June-October and December-April. For each crop season it is important to know the rice cultivated areas, which strongly depend on the water availability. The aim of this exercise is to map the rice fields at the different acquisition dates, and classify the area based on the emergence of rice from the rice paddies. Three ENVISAT ASAR AP (HH/HV) data (hence 6 images) have been acquired in an area, mainly covered by rice, in the Philippines on 13 November 2003, 18 December 2003, and 22 January 2004. Note that the ASAR AP data have already been focused and provided in SLC format.

At the end of the exercise you should be able to

- Import the ENVISAT ASAR AP data
- Convert the 1 look SLC data into 1 look Power image
- Extract the region of interest from the larger image
- Multi-look the extracted Power images
- Co-register the HH and HV images.
- Perform a speckle filter to each image
- Create and interpret a colour composite of the three HH images
- Create ratio images between images 1 and 2, and images 2 and 3 for both polarisations.
- Classify the area based on the 4 ratio images and interpret the results

Rice in the Philippines

Importance of rice

Agriculture, including forestry and fisheries, contributed approximately 22% to total GDP of the Philippines in 1990. Rice constituted about 16% of total crop production and 66% of total cereal production. Total harvested rice area was 3.5 million ha in 1980 and 3.4 million ha in 1991, 13% lower than the area under maize. The rice area of the country planted to traditional varieties (TVs) decreased from 23% in 1980 to 13% in 1988; while production from modern varieties (MVs) stood at 87% in 1991. Mean yields were 2.2 t/ha in 1980 (1.5 t/ha for TVs and 2.4 t/ha for MVs) and 2.6 t/ha in 1988 (1.9 t/ha for TVs and 2.8 t/ha for MVs). Rice production increased from 7.7 million t in 1980 to about 11.0 million t in 1995. The Philippines was self-sufficient in rice from 1980 to 1983 and in 1987, but imported 190,000 t of rice in 1984 and 593.000 t in 1990. A serious rice shortage in 1995 led to a 60% increase in rice prices, although production growth was satisfactory in 1993-95.

Rice environments

Of the approximately 3.4 million ha of ricelands, some 2.1 million ha (61%) are irrigated, 1.2 million ha (35%) are rain-fed lowland, and 0.07 million ha (2%) are upland. Much of the irrigated rice is grown on the central plain of Luzon, the country's rice-bowl. The rest comes mainly from coastal lowland areas and gently rolling erosional plains, e.g., in Mindanao and Iloilo. Rain-fed rice is found in the Cagayan Valley in northern Luzon, in Iloilo Province, and on the coastal plains of the Visayas region and Ilocos in northern Luzon. Upland rice is grown in both permanent and shifting cultivation systems scattered throughout the archipelago on rolling to steep lands.

Production constraints

Irrigated rice cultivation has increased in intensity to up to three crops per year in Central Luzon. Problems associated with such intensified systems include yield stagnation (and a possible yield decline) due to factors that are not yet fully understood, deterioration of irrigation systems (given low and decreasing investments in systems maintenance), and increased pests (insects, weeds) and diseases in some areas. Rain-fed lowland rice suffers from uncertain timing of arrival of rains, and drought and submergence - often on the same fields over the course of a single season or on different fields within a farm over the same season. Weeds, drought, diseases (blast), acidic soils, and soil erosion are major problems of upland rice in the Philippines.

Climate

The Philippines have a tropical monsoon climate, with a south-west monsoon from May to October, and a north-east monsoon from November to April. Temperatures vary little with the seasons, with average 24-hour temperatures falling mostly in the range 24-29 ^{o}C. Rainfall is highest during the south-west monsoon, and usually peaks in July-August, and is generally at it's lowest in Januarry to March, although there are considerable variations from year to year.

Rainfall and temperature statistics for selected sites near the region of the images you will be working with are shown on the right.
Source: www.worldclimate.com.

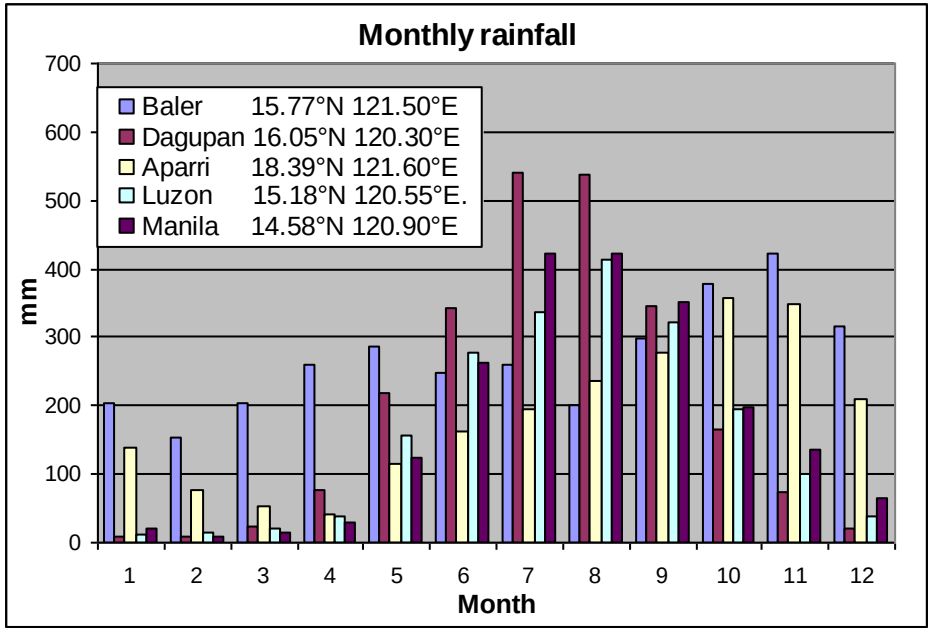

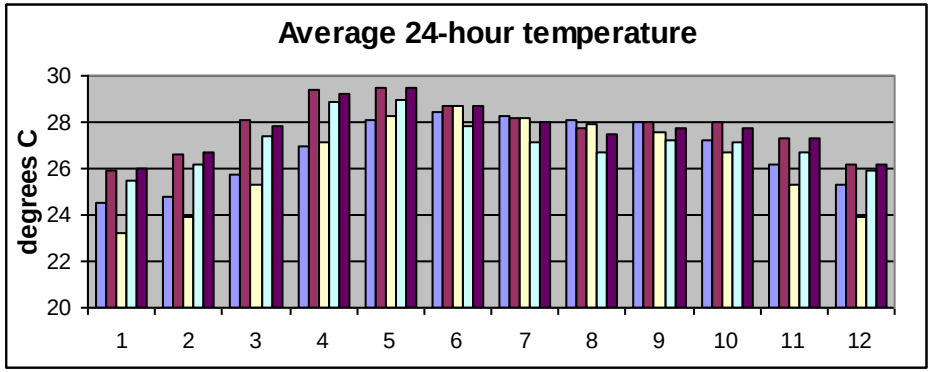

Exercise Content

1. Opening and examining Envisat ASAR data
2. Conversion from SLC data to Power 1-look image
3. Area selection and multi-looking of the Power image
4. Displaying and examining the image
5. Co-registration of the images
6. Speckle filtering
7. Creating a multi-temporal colour composite
8. Image Rationing
9. Classification based on the ratio images

1. Opening and examining Envisat ASAR data

The three images you will be working with are in the Envisat N1 format – a hierarchical data format containing metadata, ancillary data sets (georeferencing, viewing configuration etc.) and one or more measurement data sets (MDS). Opening N1 files in Bilko is simple. Using key stroke [CTRL+O] or clicking on the Open icon on the toolbar will bring up the Open dialog:

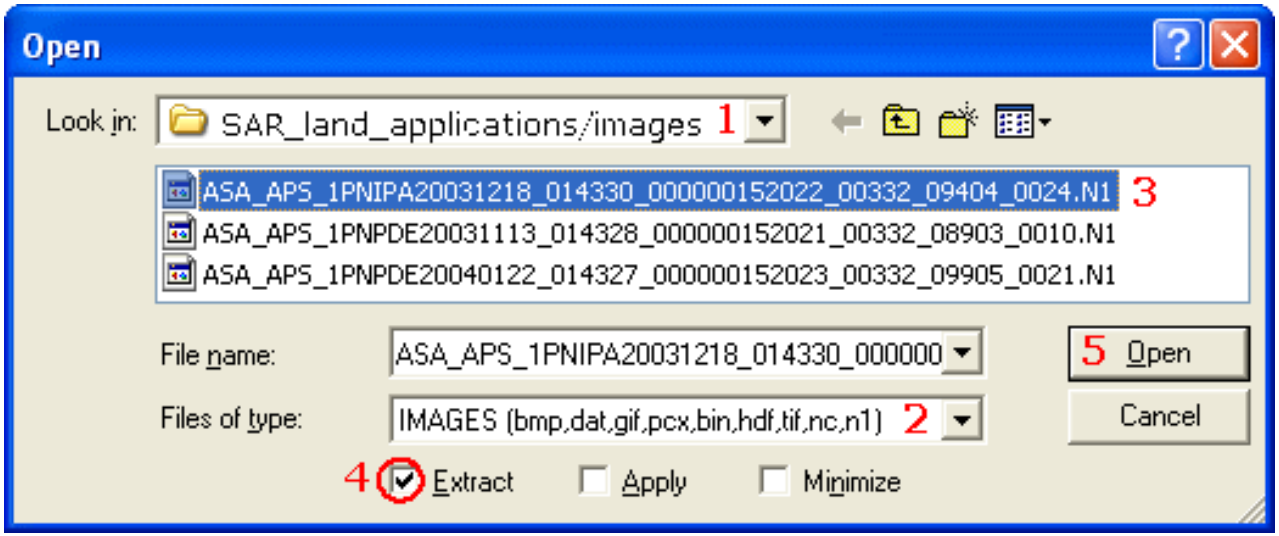

Set the folder option to look in the folder where the images are kept (1), use the arrow (2) to select IMAGES as the file type. select ASA_APS_1PNIPA20031218~.N1 from the list of images (3), make sure the Extract box is checked (4), and click Open (5).

1. Opening and examining Envisat ASAR data (cont.)

This will open the File structure window (below). The left frame lists the main data-sets in the N1 file. The right frame shows the content of the folder that has been selected in the left frame. The structure of N1 files are broadly similar for all Envisat data, with a header of metadata, ancillary data sets (ADS), and one or more measurement data sets (MDS).

Metadata contains information on orbit, time, etc. needed in order to interpret the data

Tie Point Grids contains ancillary data on geo-referencing and viewing configuration

Bands contains the measurement data sets (MDS).

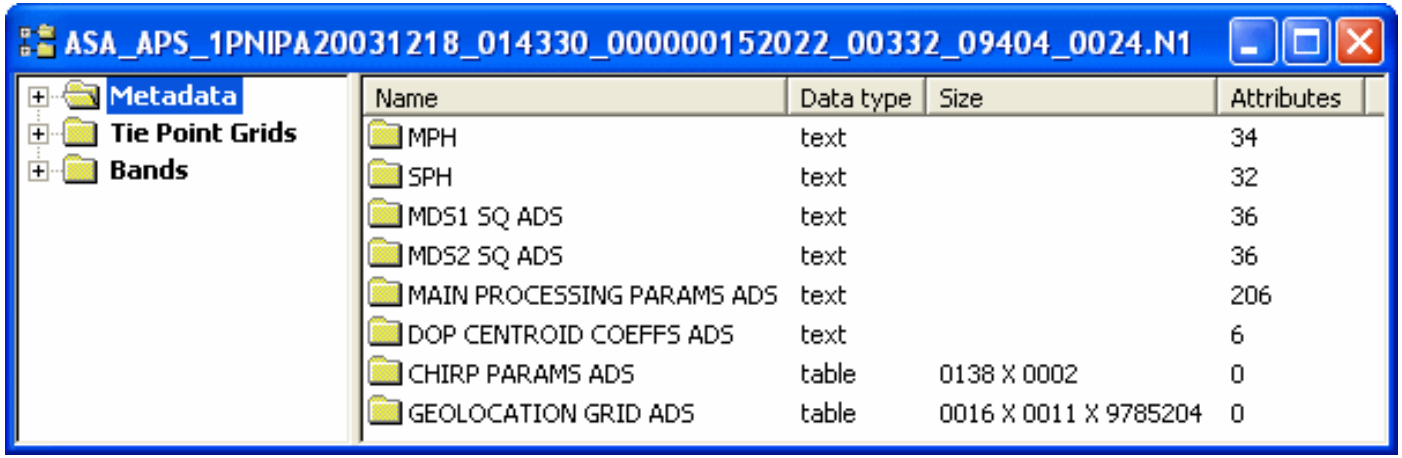

You can expand the folders by clicking on the + next to the folder in the left frame. To open a selected data set, double-click on it in the right frame.

1. Opening and examining Envisat ASAR data (cont.)

Activity / Question 1: Examining the file structure

Open the folders in turn and familiarize yourself with their content. Refer to the information in the background and theory part of the tutorial, section 1 What is SAR? To help you put the file information in context and answer the questions below.

a) What are the start and stop times for data acquisition?

b) What is the orbit number? Is the pass descending (N to S) or ascending (S to N)?

c) What is the pixel spacing in this image? Read about the SAR system in the theory section. How does the azimuth and range spacing relate to rows and columns in this image?

d) How many measurement data sets are there in this image? What are the polarisation of the transmitted (TX) and received (RX) signals for these data sets?

For more information about the structure of ASAR files see ESA's ASAR Product Handbook at http://earth.esa.int/envisat/dataproducts/asar/CNTR.htm

1. Opening and examining Envisat ASAR data (cont.)

Activity / Question 2: Examining the data properties

Expand the Bands folder and select MDS1 to reveal its content. In the right frame right-click on a folder and select Open Properties from the pop-up menu. This will reveal a description of the image. Refer to the Background and theory to help you answer the questions below. The following sections are particularly helpful in this:

1.2 SAR Specific Parameters
1.3 SAR Acquisition Modes and Techniques
1.6 SAR Data Statistics

a) The SPH folder in the Metadata describes the file as AP Mode SLC Image. What is meant by "AP mode SLC data"?

b) Which of the images in MDS1 contains the real part of the measurement? Which contains the imaginary part?

c) How would you calculate Power from the real and imaginary part of this measurement?

2. Conversion from SLC Data to 1-look Power Image

In reality each of the two data sets MSD1 and MSD2 of an AP Mode SLC image only contains the in-phase and quadrature channels. From these channels you can calculate a Power image using the answer to Question 2c.

However, because these ASAR images are so large, carrying out this calculation would require a large amount of computer memory and processing power. To avoid this problem Bilko calculates the Power image 'on the fly' from the real and imaginary part of the measurements. Therefore, all you need to do in order to open a Power image is to double-click on the p-folder in the right frame of the file structure window.

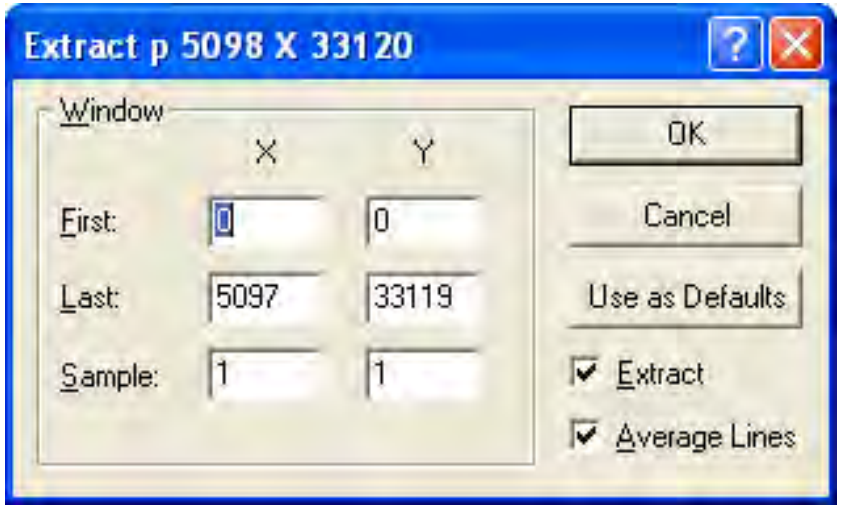

Having done this you should now see the Bilko Extract dialog (left). This allows you to select a sub-area of the original image. It also allows you to multi-look the Power image. The next slide takes you through this process.

3. Area selection and multi-looking of the Power image

Please refer to Background and Theory, section 2.1.2 Multi-looking for information relevant to this part of the exercise.

Selecting the area of interest

The Extract dialog allows you to display a selection of data from the original image. You select the area of interest by setting the start and end column (X) and start and end rows (Y) for the display window. For this image the area you want is columns 20 to 3319 and rows 4000 to 21999. Enter these Window parameters into the Extract dialog.

Setting the number of looks

The Extract dialog also contains an option for sub-sampling the image. A sampling of 1, 1 will select every column and every row, opening the image at the original resolution. However, in this case you want to load the image in such a way that the pixel size is approximately the same in both azimuth and range. This means averaging over a number of pixels in azimuth (lines) – a process known as multi-looking.

To do this you must:

- Make sure the Extract and Average Lines check boxes are both checked,
- Calculate the number of looks, and use the result to set the Sampling in the Y-dimension

3. Area selection and multi-looking of the Power image (cont.)

Activity / Question 3: Calculating the number of looks

Section 2.1.2 Mult-ilooking of the 'Background and Theory' provides and example of how to select and appropriate number of looks. You will also need to refer to the Metadata, and Tie Point Grid sections of the image file, in order to find the parameters you need to answer the questions below.

a) What parameters do you need to select an appropriate number of looks, and where would you find this information?

b) What is the ground resolution in range?

c) What is the equivalent pixel spacing in azimuth, and how many looks are required to obtain this azimuth resolution?

d) How would you set the Extract box to open the area of interest and convert the image to the desired azimuth resolution?

Use your answers to these questions to set the Window parameters in the Extract box, make sure the Extract and Average Lines boxes are checked, and OK.

4. Displaying and examining the image

Once you have accepted the Extract box settings, Bilko will start to calculate the Power Image. Because the file is very large, this may take some time. Use it to read Background and Theory, section 2.1.3 Co-registration, which you will need for the next step of the exercise. When the calculations have finished the Redisplay dialog will appear (see figure).

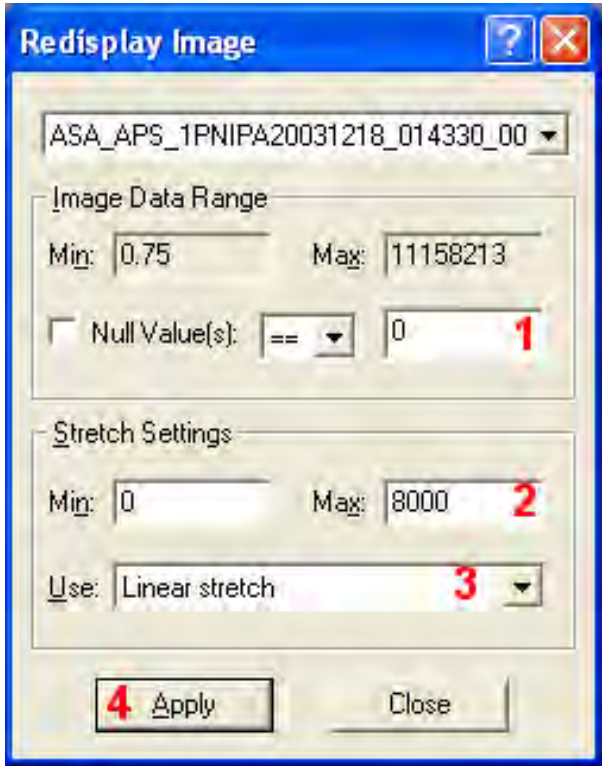

4.1 Setting the Redisplay options

There is no missing data in this image, so don't worry about setting the Null Values (1). Change the maximum used for the Linear stretch to 8000 (2), leave the stretch type as linear (3) and apply this choice (4). Save the image as hv_20031218_p.dat

The numerical format of the Power image is 32-bit float. This will be displayed as a grey-scale image, where each shade of grey corresponds to an 8-bit integer, in the range 0-255. 0 is black and 255 is white. The stretch settings tells Bilko how to map the 32-bit floating-point data to the 8-bit integers of the display. The default stretch option is a Linear Stretch, using the Max and Min values of the data. In this image there are a few Pixels with very high values, so the default would make the display look very dark. By changing the maximum setting for the stretch to 8000 you map this range to the available 0-254 display range. All pixels with values greater than 8000 will be displayed as white (255).

4. Displaying and examining the image (cont.)

4.2 Zoom

To see the whole image, right-click on the display, and select Preserve Shape. If you want to return to a higher resolution, right click again, set the % zoom factor, and click OK.

4.3 Using the cursor to examine the data

As you move the cursor around the image, its position and the pixel values will be displayed on the Status bar at the bottom of the image window.

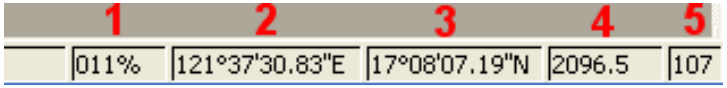

The Status Bar (above) shows the zoom factor (1), the cursor position: longitude (2) and latitude (3), the data value of the pixel (4) and the value used in the display (5).

4.4 Changing pixel position from [lon,lat] to [X,Y]

Because ASAR data has been geo-referenced, Bilko displays the position of a pixel using [lon,lat] as the coordinates. You can change this to display [x,y] coordinates, which give you the position in columns and rows relative to the top left corner of the image. Select View from the menu bar (left) and un-check Coords by clicking on it on the drop-down menu. You will need to use [x.y] coordinates for the next part of the exercise - coregistration.

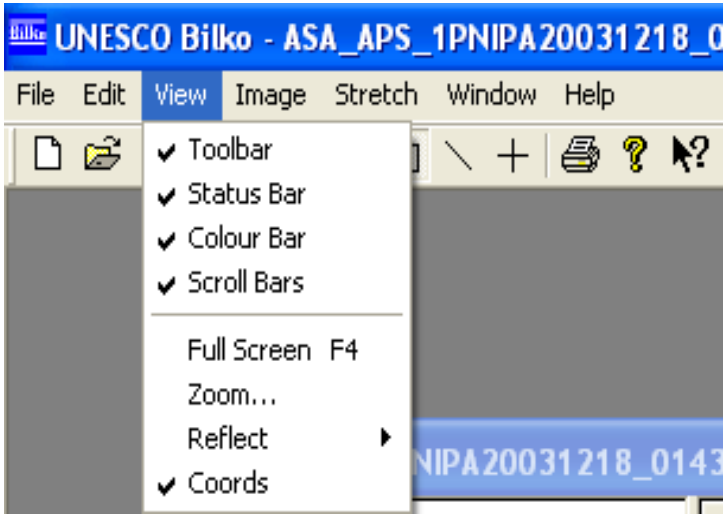

5. Co-registration of the images

Envisat has an orbit repeat time of 35 days. The three images are obtained at the same stage in the orbital cycle and should therefore cover the same area on the ground. However, in practice there are slight variations, particularly in the start and stop lines of the scene. You must therefore adjust the window of columns(x) and rows(y) you load in order to ensure that a particular point on the ground falls in the same [x,y] position in the display of all three images. A good way to do this is by finding constant features on the ground that are easy to identify in all three images, and use these as Ground Control Points (GCPs).

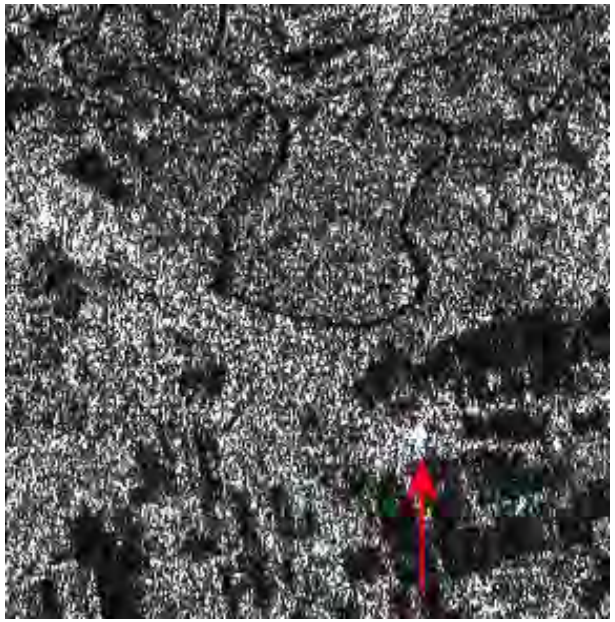

Prominent river or coastline features or highliy reflective man-made structures such as large buildings, road junctions and bridges may all provide good GCPs. By noting their [x,y] position in each of the three images, you can calculate the adjustment you must make to make sure they all show the same ground area.

The figure on the left shows one such feature, at an [x,y] position of [1540,695] in the image you have just opened.

5. Co-registration of the images (cont.)

Activity / Question 4: Identifying the position of suitable GCPs

The table below gives the position of three GCPs in the December image. Use the Go To dialog to find these features (key strokes: CTRL+G, or Edit>GoTo from the menu bar).

a) Find the position of these three GCPs in the other images, and fill in the table.

Open a small section of each of the other two images, each covering the same columns (20-3319) as in the image that is already open, and starting with the same line (4000), but ending at line 9999. (This reduces the memory requirements and processing time for the exercise).

b) Find the position of these three GCPs in the other images, and use it to fill in the table.

Image	20031113		20031218		20040122		Adjustment		Adjustment	
	x11	y11	x12	y12	x01	y01	x11-x12	y11-y12	x01-x12	y01-y12
Pixel position			2566	643						
			2493	914						
			1540	695						

c) Based on the differences in position of these GCPs, how would you fill in the Extract box to cover the same area on the ground in the November and January images as the area covered by the December image you have opened already?

5. Co-registration of the images (cont.)

Using the your answer to Question 4c, open the HV (MSD1) power image for November, using the same redisplay stretch as for the December image (Linear, max=8000). Save this as hv_20031113_p.dat.

To save time we have already created and saved the remaining images using the Bilko.dat image format. The rest of the exercise will use these prepared images together with the November and December HV power images you have created so far.

Activity / Question 5: Accuracy of the co-registration process

The table on the right is based on 36 GCPs. Use this information to answer the questions below.

Pixel Coordinates	x11-x12	y11-y12	x01-x12	y01-y12
Mean difference	-3.4	210.9	-14.0	126.3
Standard deviation (s)	1.1	1.1	1.2	1.2

a) Given 95% confidence limits (+/- 2s) how accurate will the co-registration of field boundaries be for these three images?

b) What effect could poor co-registration have on the calculation of image ratios in areas of highly variable land use?

c) Why would a simple shift of the columns and rows give insufficient coregistration of images that were overlapping but obtained at different stages of the orbital cycle?

6. Speckle filtering

You will probably have noticed the grainy appearance of the power images you have opened so far. This is due to speckle – explained in the Theory part of the tutorial, section 1.5 Speckle. Applying a filter to reduce speckle is a necessary step before carrying out image calculations or classification. Theory section 2.1.4 Speckle Filtering describes different types of filters used for this purpose. In this part of the exercise you will compare the results of applying different filters. To speed up processing and save time, you will use only a suba-area of the image. Before continuing, open the Windows menu and select Close All to close all open windows.

1. Open image hv_20031113_p.dat (saved earlier) and extract an area from [90, 3050] (top left, or first pixel) to [289, 3249] (bottom right or last pixel), choosing a linear redisplay stretch with min=0 and max=8000.
2. From the Image menu choose Filter, then Filter Options; make sure the Open New Window box is checked (right).
3. Select the whole image (keystroke: [CTRL+A] or Edit > Select All from the menu bar).
4. From the Image menu, choose Filter, then Mean.
5. Set the filter size to 5 rows and 5 columns, and press OK.
6. If the new image appears uniformly dark, right-click on the display to open the Redisplay dialog, and select the same stretch as in the original image.

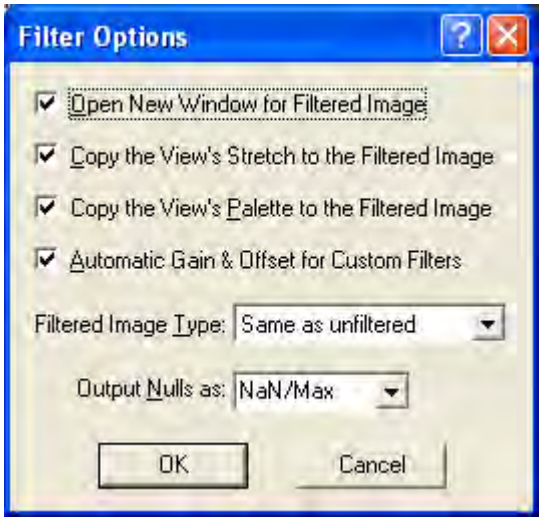

6. Speckle filtering (cont.)

Activity / Question 6: Comparing different speckle filters

When answering the questions below you may find it useful to keep section 2.1.4 Speckle Filtering to hand, as it describes all the filters you will be applying.

a) What is the main differences between the Mean filtered image and the original? What is the reason for this change?

b) Activate the original image by clicking on it, select all and apply a new filter, this time a Median filter. What are the main differences between this and the Mean filter you used earlier. How would you explain this?.

c) Activate the original image, select all and apply a Lee filter with a 5 x 5 window (remember to set the looks to 4). What are the main differences from the other filters? How would you explain this?

d) Finally activate the original image and apply a Frost filter with a 5 x 5 window (4 looks). How does this compare with the others?

e) Based on this exercise, which filter would you chose for filtering the 6 full size images?

7. Creating a multi-temporal colour composite

Colour composites are often used to display multi-band data, mapping three bands to red, green and blue to give the impression of a colour photograph. Less common, but also effective is the use of colour composite to show changes over time, by mapping images obtained at different times to the three colour channels. This is what you will do here.

Open the full size cross-polarisation (HV) images from November, December and January (hv_20031113_p.dat, hv_20031218_p.dat and hv_20040122_p.dat - to save time, the last of these has been prepared in advance for this exercise).

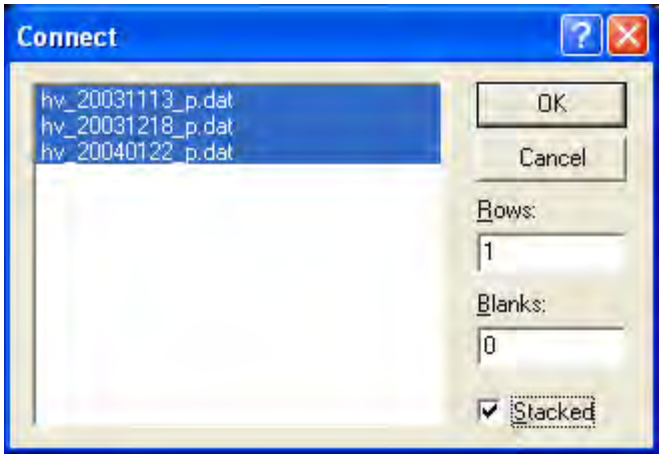

From the image menu choose Connect; this opens the connect dialog (right). Select all three images, check the Stacked checkbox and press OK. This opens a set containing the three images, one on top of the other. By using the tab key on the keyboard you can move through the stack image by image. Alternatively you can use the selector (below) to choose an image.

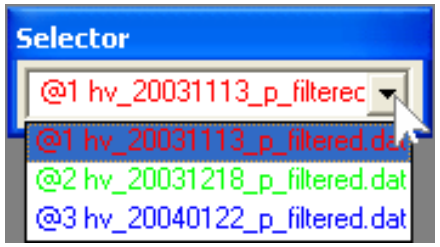

Note that the colour of the text in the selector corresponds to the colour tube to which this image will be matched in a colour composite. This is useful information that will help you interpret the composite image later on.

7. Creating a multi-temporal colour composite (cont.)

To create the colour composite, select Composite from the image menu. As you see this still contains speckle. You can filter all three images at the same time. In Filter Options uncheck Open New Window, and apply the filter given in the answer to Question 6e.

Filtering the three images may take some time. Use it to read the Background and Theory on Scattering Mechanisms (section 1.4) and Rice Mapping (section 3.9), which should help you answer the questions below.

You may find that having filtered the images once there is still some speckle left. Apply the Frost filter a second time, this time with a window size of 7 x 7. You may need to adjust the re-display stretch after filtering. A linear stretch of range 0-6000 should be a good choice.

Before continuing save the individual images of the composite as Bilko.dat files, adding _filtered to the file name, then save the composite as a Windows.bmp file, calling it nov_dec_jan_hv.bmp.

Note: The Bilko.dat format allows the original image data to be saved as 32-bit floats, thus retaining all the information for further analysis. When you open it again, you have to set the Redisplay stretch to be able to display the images in the same way as you did earlier.
In contrast the BMP format allows you to save the composite image just as you see it on the screen. This saves the display values (0-255) in a 24-bit (3 x 8-bit) image file, which can be used as an illustration in reports, but does not retain the original numerical data in a form that is useful for further analysis.

7. Creating a multi-temporal colour composite (cont.)

Activity / Question 7: Interpreting the composite

To help you answer the questions refer to the theory section 1.4- scattering mechanisms. It will also help you to know how red, green and blue combine to give other colours. The figures below show how full brightness (255) in all channels combines to give white, and the same, reduced brightness (128) in all channels give grey.

a) In the composite image, which month has been mapped to red, green and blue respectively?

b) The central area of the image contains two red areas that appear more uniform than the surrounding regions. Explain the difference in colour from the surrounding grey areas, and why this indicates that the red region is agricultural land used to grow rice.

c) Towards the lower edge of the composite there is a bright (white) area surrounded by grey. Using the information about scattering how would you interpret these colours?

d) At position [2565,640] a bridge crosses the river. The river appears dark, unlike the bridge and land on both sides. Use your knowledge of scattering to explain this.

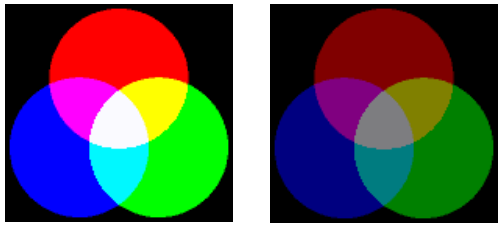

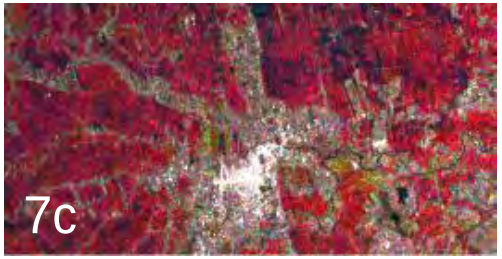

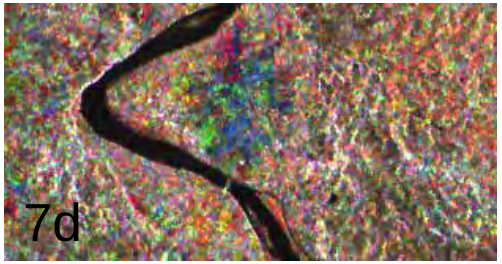

8. Image Rationing

The colour composite gave an indication of the extent to which different areas of the image change over time - a property linked to the use of land for agriculture. Classification can take advantage of this to sort pixels into land use groups based on temporal change patterns. It is often preceded by image-rationing, which helps to distinguish areas that change with time from areas that vary spatially but not temporally.

Bilko formulas allow you to carry out such image calculations. Open the file ratio.frm (to see this in the open dialog, change the file type from IMAGE to FORMULA). Bilko formulas usually consist of three parts:

- Comments, preceded by #, which explain what the formula or parts of the formula does.
- Constant declarations, which assigns a variable name to an image plane (@*) or a value.
- Statements – one or more mathematical or logical expression that contain the processing instructions to the software.

To apply a formula you must first connect images into to a set – usually a stack. The formula ratio.frm requires 3 images to be connected. @1 and @2 are images to be rationed; @3 is a blank image that will take the output of the calculations.

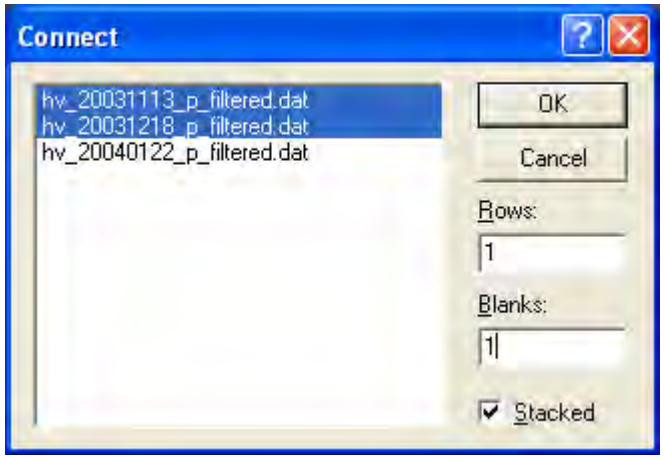

8. Image Rationing (cont.)

1. Connect the November and December images as indicated on the previous slide. Use the selector or tab key to display the blank image (@3).
2. To apply the formula, first select it (select it from the Windows menu or click on it to activate).
3. Select Options from the menu bar and make sure the output is set to float (or same as @1). With the formula still active, copy it ([CTRL+C]) and paste it onto the set (activate the set, then press [CTRL+P]), and watch how the blank image changes.
4. Activate the single blank image (called imageXX) and save it as ratio_nov_dec_hv.dat.
5. The image will appear uniformly dark, so you need to adjust the stretch. Right-click to open the Redisplay dialog and change the linear stretch range to 0-5.
6. Repeat 1-5 for the remaining 3 image pairs:
   - hv_20031218_p_filtered.dat / hv_20040122_p_filtered.dat
   - hh_20031113_p_filtered.dat / hv_20031218_p_filtered.dat
   - hv_20031218_p_filtered.dat / hv_20040122_p_filtered.dat

Question 8 – Interpreting the ratio data

Broadly speaking the images contain three different areas – bright, mid grey and dark. What do these shades represent?

9. Classification – Principal Component Analysis (PCA)

Close all open windows and re-open the ratio files, using the Extract box to remove 3 columns and 3 rows on each image edge as shown in the figure on the right.

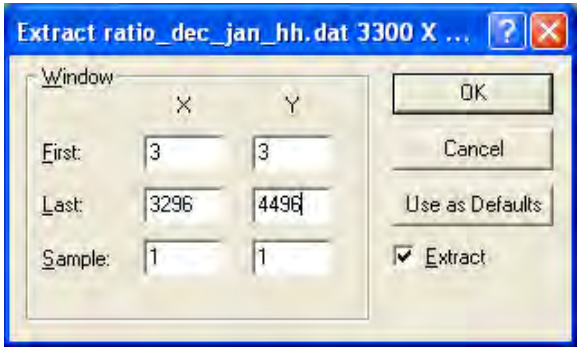

Connect these images into a stacked set.

From the Image menu, select PCA to open the PCA dialog and set this as illustrated below, and press OK.

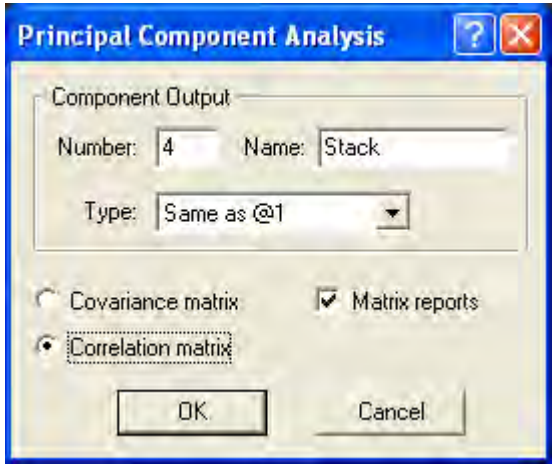

Question 9 – Preparing for PCA

a) Explain why it was necessary to remove the edge rows and columns from the ratio images. How many rows would have to be removed if the filter window used to remove speckle had been 15 x 15 pixels?

b) The Bilko help menu explains how to set the PCA dialog. Why was Correlation matrix a better option than Covariance matrix for the analysis of the ratio images?

When the PCA is completed, save the 4 new images as PC1, PC2 etc. chosing Bilko.dat as the file type. Redisplay the four images using a linear stretch of -5 to 5.

9. Classification – Principal Component Analysis (cont.)

Activity / Question 10: Interpreting the classification

To help you answer the questions below, read about Principal Component analysis in Background and Theory, section 2.1.11 Classification, and in the Bilko Help file.

a) The PCA created two tables, what do these show?

b) The loading represents the correlation between each PC and each of the original ratio images. You can use this to calculate how much of the variance is explained by each PC. Simply sum up the squares of the variable loadings for each PC and divide this by the number of variables (in this case 4). To do this, select the whole table [CTRL+A], copy the data (CTRL+C], and paste it to a spreadsheet, where you can carry out the necessary calculations. How many PCs do you need to account for more than 90% of the variability?

c) Compare PC3 to PC1 and PC2. Which seems most likely to be affected mainly by noise due to insufficient speckle filtering or co-registration errors?

d) Combine the first 3 PCs into a set (not stacked) and create a colour composite. How would you classify the land based on the results of this PCA?

e) How does this compare with rice mapping in this area during the following year 2004/20005 described in the example in Background and Theory, section 3.9 ?.

f) Unlike the first colour composite this one does not distinguish urban areas and permanent water bodies from other areas that do not change with time. Can you explain why?