

ILWIS USER's Guide Chapter 6

Adepted by Michiel Damen, ITC - September 2011

Introduction to Image Processing

ILWIS contains a set of image processing tools for enhancement and analysis of data from space borne or airborne platforms. In this chapter, the routine applications such as image enhancement, classification and geo-referencing are described. The image enhancement techniques, explained in this chapter, make it possible to modify the appearance of images for optimum visual interpretation. Spectral classification of remotely sensed data is related to computer-assisted interpretation of these data. Geo-referencing remote sensed data refers to geometric distortions, the relationship between an image and a coordinate system, the transformation function and resampling techniques.

Remotely sensed data

Remotely sensed data, such as satellite images, are measurements of reflected solar radiation, energy emitted by the earth itself or energy emitted by Radar systems that is reflected by the earth. An image consists of an array of pixels (*picture elements*) or gridcells, which are ordered, in rows and columns. Each pixel has a *Digital Number (DN)*, that represents the intensity of the received signal reflected or emitted by a given area of the earth surface. The size of the area belonging to a pixel is called the spatial resolution. The DN is produced in a sensor-system dependent range; the radiometric values. An image may consist of many layers or bands. Each band is created by the sensor that collects energy in specific wavelengths of the electro-magnetic spectrum.

Before you can start with the exercises, you should have installed ILWIS and copied all the exercise data to a folder on your own hard disk, for instance **D:\ILWIS\ExerciseData\Rs**.



- Double-click the ILWIS icon on the desktop to open ILWIS.
You can also go to the folder where the ILWIS program files are being stored and click:  **Ilwis30.exe**
- Use the Navigator to go to the folder where the ILWIS exercise data are being stored.

6.1 Visualization of single band images

For satellite images the *image domain* is used in ILWIS. Pixels in a satellite image usually have values ranging from 0-255 (byte value). The values of the pixels represent the reflectance of the surface object. The image domain is in fact a special case of a *value domain*. Raster maps using the image domain are stored using the '1 byte' per pixel storage format. A single band image can be visualized in terms of its gray shades, ranging from black (0) to white (255).

To compare satellite image bands before and after an operation, the images can be displayed in different windows, visible on the screen at the same time. The relationship between gray shades and pixel values can also be detected.

The pixel location in an image (rows and columns), can be linked to a georeference which in turn is linked to a coordinate system which can have a defined map projection. In this case, the coordinates of each pixel in the window are displayed if one points to it.

The **objectives of the exercises** in this section are:

- to understand the relationship between the digital numbers of satellite images and the display, and
- to be able to display several images, scroll through and zoom in/out on the images and retrieve the digital numbers of the displayed images.

Satellite or airborne digital image data is composed of a two-dimensional array of discrete *picture elements* or pixels. The intensity of each pixel corresponds to the average brightness, or radiance, measured electronically over the ground area corresponding to each pixel.

Remotely sensed data can be displayed by reading the file from disk line-by-line and writing the output to a monitor. Typically, the Digital Numbers (DNs) in a digital image are recorded over a numerical range of 0 to 255 (8 bits = $2^8 = 256$), although some images have other numerical ranges, such as, 0 to 63 (4 bits = $2^4 = 64$), or 0 to 1023 (10 bits = $2^{10} = 1024$). The display unit has usually a display depth of 8 bits. Hence, it can display images in 256 (gray) tones.

The following Landsat ETM bands of an area covering the village of Dedoplistskaro (Georgia) are used in this exercise:

Landsat Enhanced Thematic Emappr (ETM) multi spectral : band 1 (0.45-0.52 μm); band 2 (0.52-0.50 μm); band 3 (0.63-0.69 μm) band 4 (0.76-0.90 μm), band 5 (1.55-1.75 μm): and band 7 (2.08-2.35 μm).

Panchromatic band 8:

Path/Row: 169 / 31 Date of acquisition: 20 August 1999

Display of the satellite image

Go to the folder: **ILWIS_ExerciseData/Dedo_ETM_20Aug99**

Display the image ETM99 band 1 (the "blue" band) of the area..



- Double-click image **ETM99_band_1** in the Catalog. The Display Options - Raster Map dialog box is opened.
- In the Display Options - Raster Map dialog box, it is shown that the image **ETM99_band_1** has domain Image, has digital numbers ranging from 25 to 254 and will be displayed in gray tones (Representation: Gray).
- Click OK in the dialog box to accept these display settings. The image will be displayed in gray tones from black (lowest reflectance) to white (highest reflectance).

The image `ETM_band_1` is now displayed in a map window. The map window can be moved, like all windows, by activating it and dragging it to another position. The size of the map window can be changed in several ways.



- Reduce/enlarge the map window by dragging the borders of the window.
- Maximize the map window by clicking the Maximize button .
- Close the map window by opening the File menu in the map window and choose Exit, or click the Close button  in the map window.

Zoom in / out on a displayed satellite image

When you want to study details of the image, the zooming possibility can be used.



- Open **ETM_band_1** again, accept the defaults in the Display Options - Raster Map dialog box and click OK.
- Maximize the map window by clicking the Maximize button .
- Click the Zoom In button  in the toolbar to zoom in on a selected area.
- Move the cursor (now in the shape of a magnifying glass) to the first corner of the area that has to be displayed in detail and click on the left mouse button. Without releasing the button, drag the cursor a little bit to the second corner of the area of interest. Releasing the mouse button will result in an enlarged display of the selected area. You can also click in the map window to zoom in.
- Select the Zoom Out button  in the toolbar and click (repeatedly) in the map to zoom out.
- Click the Entire Map button  to show the entire map in the map window again.

By repeatedly zooming in or by zooming in on a very small area, the pixel structure of the image becomes clearly visible by the blocky appearance of the displayed image.



- Zoom in on the displayed image till you can clearly see the blocky pixel structure of the image. If necessary, change the size of the map window.
- Click the Normal button  in the toolbar to go back to the normal view.
- Press the left mouse button in the map window and the corresponding DN value of the pixel will appear. Move over some pixels with the cursor, selecting dark and bright toned pixels and note the changes in DN values.
- Close the map window.

Scrolling through a displayed satellite image

When part of an image is displayed in a map window, other parts of the image can be viewed by scrolling through the image.



- Display the image **ETM_band_4** and zoom in on an area.

- To roam through the image use the **Pan** button  or use the **Left/Right/Up/Down** scroll boxes in the horizontal or the vertical scroll bar.
- A fast scroll can also be achieved by dragging the scroll boxes in the scroll bars **Left/Right** or **Up/Down**, or by clicking in the scroll bar itself.
- Close the map window.

Displaying multiple images

It is often useful to display more than one image on the screen. This is done by opening a map window for each image to be displayed. Multiple map windows can be opened at the same time by selecting several maps and using the right mouse button menu to open them.

It is possible to display in one map window a raster image together with one or more point, segment or polygons maps but it is in ILWIS not possible to display two or more raster images in the same map window.

There are more ways of displaying images in map windows. Three bands of Landsat ETM, of the area, will be displayed using three different display windows.

- Double-click **ETM_band_1** in the **Catalog** and click **OK** in the **Display Options - Raster Map** dialog box. Move the map window to a corner of the screen.
- Double-click the **Show** item in the **Operation-tree**.
- In the **Open Object** dialog box select **ETM_band_3** and click **OK**. The **Display Options - Raster Map** dialog box is displayed.
- Accept the defaults and click **OK**. Move the window to a position close to the **ETM_band_1** window.
- Click the right mouse button on **ETM_band_4** in the **Catalog**. A context-sensitive menu appears.
- Select **Open** and accept the default display options, by clicking **OK** in the **Display Options - Raster Map** dialog box. Move the window to a position close to the **ETM_band_3** window.
- Drag-and-drop the polygon map **Dedoplistskaro** from the **Catalog** into the map window, which displays **ETM_band_1**. Select the option: **Boundaries Only** in the **Display Options-Polygon Map** dialog box. Select the **Boundary Color: Red**, make sure that the check box **Info** is selected and click **OK**
- If there is time: Do the same for the other ETM bands available.
- Study the individual windows and try to explain the gray intensities in terms of spectral response of water, agriculture and grasslands for the individual bands. Write down the relative spectral response as 'low', 'medium' or 'high', for the different land cover classes and spectral bands in Table 6.1.
- Close all map windows after you have finished the exercise.

Table 6.1: Relative spectral responses per band for different land cover classes. Fill in the table with classes: low, medium and high.

Band	Water	Agriculture	Grassland	Bare land	
1					
3					
4					

Digital numbers and pixels

The spectral responses of the earth's surface in specific wavelengths, recorded by the spectrometers on board of a satellite, are assigned to picture elements (pixels). Pixels with a strong spectral response have high digital numbers and vice versa. The spectral response of objects can change over the wavelengths recorded, as you have seen in the former exercise when comparing the three Landsat ETM bands with regard to for example water and rural areas. When a gray scale is used, pixels with a weak spectral response are dark toned (black) and pixels representing a strong spectral response are bright toned (white). The digital numbers are thus represented by intensities on a black to white scale.

The digital numbers themselves can be retrieved through a *pixel information window*. This window offers the possibility to inspect interactively and simultaneously, the pixel values in one or more images.

There are several ways to open and add maps to the pixel information window.

- Display map **ETM_band_1** in a map window.
- From the **File** menu in the map window, select **Open Pixel Information**. The pixel information window appears.
- Make sure that the pixel information window can be seen on the screen. Select **ETM_band_3** and **ETM_band_4** in the **Catalog** and drag these maps to the pixel information window: hold the left mouse button down; move the mouse pointer to the pixel information window; and release the left mouse button; The images are dropped in the pixel information window.
- Display **ETM_band_4**, click **OK** in the **Display Options - Raster Map** dialog box. Make sure that the map window does not overlap with the pixel information window.
- In the pixel information window the DNs of all three selected bands will be displayed. Zoom in if necessary.

You can also add maps one by one to the pixel information window (**File** menu **Add Map**).

Roam through image **ETM_Band_4** and write down in Table 6.2 some DN values for water, , Agricultural land, Grassland and Bare land.

Table 6.2: DN values of different land cover classes of selected spectral bands. Fill in the DN values yourself.

Band	Water	Agriculture	Grassland	Bare land	
1					
3					
4					

- Close the pixel information window.

Pixels and real world coordinates

When an image is created, either by a satellite or airborne scanner, the image is stored in row and column geometry in raster format. There is no relationship between the rows/columns and real world coordinates (UTM, geographic coordinates, or any other reference map projection) yet. In a process called geo-referencing, which is discussed in section 6.4, the relationship between row and column number and real world coordinates can be established.



- Double-click the map **ETM_band_4** in the Catalog. The Display Options - Raster Map dialog box appears.
- Click OK in the Display Options - Raster map dialog box and maximize the map window.
- Move the cursor to the lake area in the SE and note the Row/Col and XY, Lat/Long figures as given in the Status bar. Zoom in if necessary.
- Move to the hilly area in the NE corner of the image. Note the change in real world coordinates and the row/col numbers.

When a map has coordinates, distances can be measured.



- Click the Measure Distance button  on the toolbar of the map window, or choose the Measure Distance command from the Options menu.
- Move the mouse pointer to the northern edge of the lake (starting point) and click, hold the left mouse button down, and release the left mouse button at the southern edge of the lake (end point). The Distance message box appears.

The Distance message box will state:

From: the XY-coordinate of the point where you started measuring;

To: the XY-coordinate of the point where you ended measuring;

Distance on map: the distance in meters between starting point and end point calculated in a plane;

Azimuth on map: the angle in degrees between starting point and end point related to the grid North;

Ellipsoidal Distance: the distance between starting point and end point calculated over the ellipsoid;

Ellipsoidal Azimuth: the angle in degrees between starting point and end point related to the true North, i.e. direction related to the meridians (as visible in the graticule) of your projection.



- Click OK in the Distance message box.
- Close the map window after finishing the exercise.

Summary: Visualization of images

- Satellite or airborne digital images are composed of a two-dimensional array of discrete picture elements or pixels. The intensity of each pixel corresponds to the average brightness, or radiance, measured electronically over the ground area corresponding to each pixel.
- A single Landsat ETM image band can be visualized in terms of its gray shades, ranging from black (0) to white (255).
- Pixels with a weak spectral response are dark toned (black) and pixels representing a strong spectral response are bright toned (white). The digital numbers are thus represented by intensities from black to white.

- To compare bands and understand the relationship between the digital numbers of satellite images and the display, and to be able to display several images, you can scroll through and zoom in/out on the images and retrieve the DNs of the displayed images.
- In one map window, a raster image can be displayed together with point, segment or polygon maps. It is not possible in ILWIS to display two raster maps in one map window.
- An image is stored in row and column geometry in raster format. When you obtain an image there is no relationship between the rows/columns and real world coordinates (UTM, geographic coordinates, or any other reference map projection) yet. In a process called geo-referencing, the relationship between row and column number and real world coordinates can be established (see section 6.4).

6.2 Image enhancement

Image enhancement deals with the procedures of making a raw image better interpretable for a particular application. In this section, commonly used enhancement techniques are described which improve the visual impact of the raw remotely sensed data for the human eye.

Image enhancement techniques can be classified in many ways. **Contrast enhancement**, transforms the raw data using the statistics computed over the whole data set. Examples are: **linear contrast stretch**, **histogram equalized stretch** and **piece-wise contrast stretch**. Contrary to this, spatial or local enhancement only take local conditions into consideration and these can vary considerably over an image. Examples are *image smoothing and sharpening*.

6.2.1 Contrast enhancement

The objective of this section is to understand the concept of contrast enhancement and to be able to apply commonly used contrast enhancement techniques to improve the visual interpretation of an image.

The sensitivity of the on-board sensors of satellites, has been designed in such a way that they record a wide range of brightness characteristics, under a wide range of illumination conditions. Few individual scenes show a brightness range that fully utilizes the brightness range of the detectors. The goal of contrast enhancement is to improve the visual interpretability of an image, by increasing the apparent distinction between the features in the scene. Although the human mind is excellent in distinguishing and interpreting spatial features in an image, the eye is rather poor at discriminating the subtle differences in reflectance that characterize such features. By using contrast enhancement techniques these slight differences are amplified to make them readily observable.

Contrast stretch is also used to minimize the effect of haze. Scattered light that reaches the sensor directly from the atmosphere, without having interacted with objects at the earth surface, is called **haze**. Haze results in overall higher DN values and this additive effect results in a reduction of the contrast in an image. The haze effect is different for the spectral ranges recorded; highest in the blue, and lowest in the infra red range of the electromagnetic spectrum.

Techniques used for contrast enhancement are: the **linear stretching technique** and the **histogram equalization**. To enhance specific data ranges showing certain land cover types the **piece-wise linear contrast stretch** can be applied.

A computer monitor on which the satellite imagery is displayed is capable of displaying at least 256 gray levels (0 - 255). This corresponds with the resolution of most optical medium resolution satellite images, as their digital numbers also vary within the range of 0 to 255. To produce an image of optimal contrast, it is important to utilize the full brightness range (from black to white through a variety of gray tones) of the display medium.

The linear stretch (Figure 6.1) is the simplest contrast enhancement. A DN value in the low end of the original histogram is assigned to extreme black, and a value at the high end is assigned to extreme white. The remaining pixel values are distributed linearly between these extremes. One drawback of the linear stretch is that it assigns as many display levels to the rarely occurring DN values as it does to the frequently occurring values. A complete linear contrast stretch where (min, max) is stretched to (0, 255) produces in most cases a rather dull image. Even though all gray shades of the display are utilized, the bulk of the pixels is

displayed in mid gray. This is caused by the more or less normal distribution, with the minimum and maximum values in the tail of the distribution. For this reason it is common to cut off the tails of the distribution at the lower and upper range.

The histogram equalization technique (Figure 6.1) is a *non-linear stretch*. In this method, the DN values are redistributed on the basis of their frequency. More different gray tones are assigned to the frequently occurring DN values of the histogram.

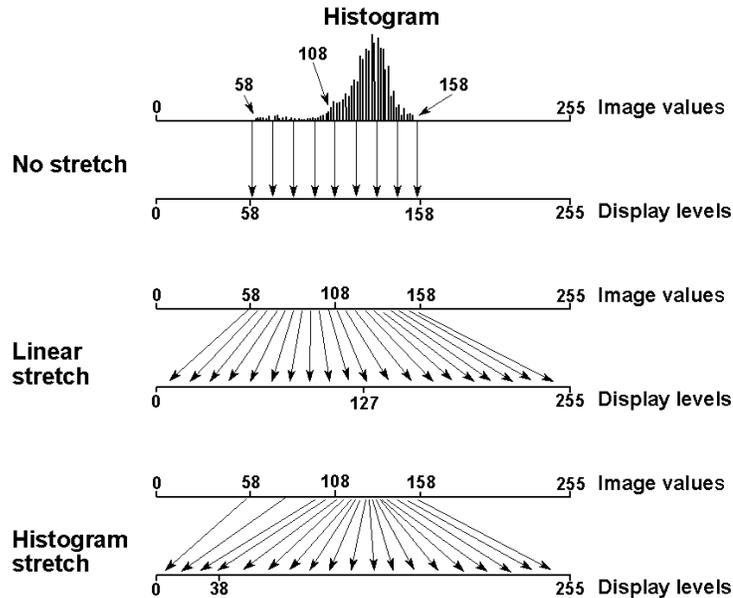


Figure 6.1: Principle of contrast enhancement.

Figure 6.1 shows the principle of contrast enhancement. Assume an output device capable of displaying 256 gray levels. The histogram shows digital values in the limited range of 58 to 158. If these image values were directly displayed, only a small portion of the full possible range of display levels would be used. Display levels 0 to 57 and 159 to 255 are not utilized. Using a *linear stretch*, the range of image values (58 to 158) would be expanded to the full range of display levels (0 to 255). In the case of linear stretch, the bulk of the data (between 108 and 158) are confined to half the output display levels. In a *histogram equalization stretch*, the image value range of 108 to 158 is now stretched over a large portion of the display levels (39 to 255). A smaller portion (0 to 38) is reserved for the less numerous image values of 58 to 108.

The effectiveness depends on the original distribution of the data values (e.g. no effect for a uniform distribution) and on the data values of the features of interest. In most of the cases it is an effective method for gray shade images.

Density slicing and *piece-wise linear contrast stretch*, two other types of image enhancement techniques, are treated in section 6.6.1 and 6.6.2 respectively.

The material for this exercise consists of a **Landsat ETM band 1** of the area..

Calculation of a histogram

Before stretching can be performed, a histogram of the image has to be calculated.

☞

- In the Operation-tree, double-click the Histogram operation. The Calculate Histogram dialog box is opened.
- Select the raster map **ETM_band_1** and click the Show button.
- The histogram window appears showing a graphical display of the histogram as well as a numerical display of the histogram in a number of columns:

N_{pix} (= number of pixels with a certain pixel value),
 N_{pixpct} (= the percentage of pixels compared to the total number of pixels),
 $Pctnotzero$ (= the percentage of pixels compared to the total number of pixels with non zero values),
 N_{pixcum} (= cumulative frequency distribution),
 $N_{pcumpct}$ (= cumulative frequency distribution in percentages) and, if pixel size is known,
 $Area$ (= area of pixels with a certain DN). In order to have a better view maximized the histogram window.

- Write down in Table 6.3 the DN values that belong to N_{pixpct} .

Table 6.3: DN values at specific cumulative percentages. Fill in the table from the histogram of ETM_band_1.

Cumulative Percentage	0	0.5	1	2
DN value				
Cumulative Percentage	100	99.5	99	98
DN value				

For an optimal stretching result, the histogram shape and extent has to be studied. This can be done using the numerical listing but a better picture can be obtained from a graphical display.

Use the graph and the histogram table to answer the following questions:

- 1) What is the total number of pixels of the image?
- 2) How many pixels have a digital number of 68?
- 3) How many pixels have a DN more than 93?
- 4) Which DN represents the average value of the image?

Linear stretching

After a histogram has been calculated for a certain image, the image can be stretched. A linear stretch is used here. Only the pixel values in the 1 to 99% interval will be used as input; pixel values below the 1% boundary and above the 99% boundary will not be taken into account.



- In the Main window, open the **Operations** menu and select **Image Processing, Stretch**. The **Stretch** dialog box is opened.
- Select **Tmb1** as **Raster Map**, accept **Linear stretching** as stretching method with a **Percentage** equal to **1.00** and type **ETM_band_1_stretch** as **Output Raster Map**.
- Accept all other defaults and click the **Show** button. The raster map **ETM_band_1_stretch** is calculated after which the **Display Options - Raster Map** dialog box appears.
- Click **OK** to display the map.

Compare original and stretched images

The original and the stretched image can be compared by displaying them in two map windows next to each other and by inspecting the image statistics in their histograms.



- Display the original image **ETM_band_1** and the stretched image **ETM_band_1_stretch**, each in a map window next to each other and assess visually the effect of linear stretching. Do not forget to

switch off the **Stretch** option in the **Display Options - Raster Map** dialog box of the original image!

- Display the histograms of **ETM_band_1** and **ETM_band_1_stretch** and study them.
- Close the map and histogram windows.

Histogram equalization

The histogram equalization stretching technique takes into account the frequency of the DN's. Just as in linear stretching, percentages for the lower and upper DN values can be specified, but also user-specified minimum and maximum values of the DN's themselves can be given instead.



- Open the **Stretch** dialog box by double-clicking **Stretch** in the **Operation-list**.
- Select **ETM_band_1** as input **Raster Map** and **Histogram Equalization** with 10 intervals as stretching method. Each interval should receive 10% of the data. Use the default **Percentage** to stretch from and type **ETM_band_1_equa1** for the **Output Raster Map**. Click the **Show** button to display the stretched map after the stretching operation. The **Display Options - Raster Map** dialog box appears.
- Display the stretched image, using a **Gray Representation**.
- Display the histogram of **ETM_band_1_equa1** and check the frequency per interval.
- Close the windows afterwards.

Summary: Image enhancement

- *Contrast enhancement*, transforms the raw data using the statistics computed over the whole data set.
- Techniques used for a contrast enhancement are: the *linear stretching technique* and *histogram equalization*.
- The linear stretch is the simplest contrast enhancement. A DN value in the low end of the original histogram is assigned to extreme black, and a value at the high end is assigned to extreme white.
- The histogram equalization technique is a non-linear stretch. In this method, the DN values are redistributed on the basis of their frequency. More different gray tones are assigned to the frequently occurring DN values of the histogram.

Remark: The subject image enhancement by filtering is not included in this part of the exercise.

6.3 Visualizing multi-band images

In this section, the objective is to understand the concept of color composites and to be able to create different color composites.

The spectral information stored in the separate bands can be integrated by combining them into a **color composite**. Many combinations of bands are possible. The spectral information is combined by displaying each individual band in one of the three primary colors: Red, Green and Blue.

A specific combination of bands used to create a color composite image is the so-called False Color Composite (FCC). In a FCC, the red color is assigned to the near-infrared band, the green color to the red visible band and the blue color to the green visible band.

The green vegetation will appear reddish, the water bluish and the (bare) soil in shades of brown and gray. A combination used very often for ETM imagery is the one that displays in red, green and blue the respective bands 5, 4 and 3. Other band combinations are also possible. Some combinations give a color output that resembles natural colors: water is displayed as blue, (bare) soil as red and vegetation as green. Hence this combination leads to a so-called *Pseudo Natural Color Composite*.

6.3.1 Color composites

Color composites are created and displayed on the screen, by combining the spectral values of three individual bands. Each band is displayed using one of the primary colors. A combination of pixels with high DN values for the individual bands results in a light color. Combining pixels with low DN values produces a dark color. Each point inside the cube produces a different color, depending on the specific contribution of red, green and blue it contains.

In ILWIS the relationship between the pixel values of multi-band images and colors, assigned to each pixel, is stored in the representation. A representation stores the values for *red*, *green*, and *blue*. The value for each color represents the relative intensity, ranging from 0 to 255. The three intensities together define the ultimate color, i.e. if the intensity of red = 255, green = 0 and blue = 0, the resulting color is bright red. Next to this, the domain assigned to color composites is the *picture* or the *color* domain, as the meaning of the original data has disappeared.

The pixel values from the three input images are used to set values for corresponding pixels in the composite. One input band gives the values for red (R), another the values for green (G) and the third for blue (B) (Figure 6.2).

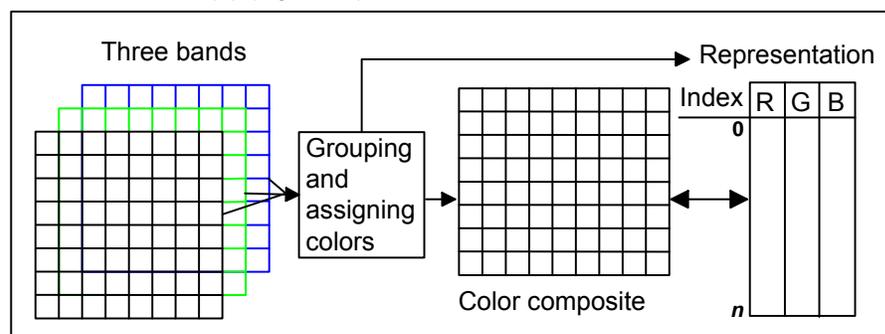


Figure 6.2: The link between a color composite, the source images and the representation. n is the number of colors.

In ILWIS, color composites can be created using different method: interactively by storing the screen results as a **Map View** or permanently by creating an output raster map with the **Color Composite** operation.

In this exercise only the permanent creation of a color composite is being explained.

Permanent Color Composite

The Color Composite operation is used to create permanent color composites. In this operation color composites can be created in various ways:

- Standard Linear Stretching;
- Standard Histogram Equalization;
- Dynamic;
- 24 Bit RGB Linear Stretching;
- 24 Bit RGB Histogram Equalization;
- 24 Bit HIS;

The different methods of creating a color composite are merely a matter of scaling the input values over the output colors. The exact methods by which this is done are described the ILWIS Help topic: “Color Composite: Algorithm”.



- In the Operation-tree, open the Image Processing item, and double-click the Color Composite operation. The Color Composite dialog box is opened.
- Select image **ETM_band_4** for the Red Band, **ETM_band_3** for the Green Band and **ETM_band_2** for the Blue Band.
- Type **ETM_CC_Bands_432** in the Output Raster Map text box, accept all other default and click Show. The permanent color composite is calculated and the Display Options - Raster Map dialog appears.
- Click OK to display the map and close it after you have seen the result.
- Try other image band combinations too.

Summary: Visualizing multi-band images

- The spectral information stored in the separate bands can be integrated by combining them into a *color composite*. Many combinations of bands are possible. The spectral information is combined by displaying each individual band in one of the three primary colors: red, green and blue.
- In a *False Color Composite (FCC)*, the red color is assigned to the near-infrared band, the green color to the red visible band and the blue color to the green visible band.
- In a *Pseudo Natural Color Composite*, the output resembles natural colors: water is displayed in blue, (bare) soil as red and vegetation as green.

6.4 Geometric corrections and image referencing

Remark: for this exercise the standard ILWIS User's Guide data for Chapter 6 are being used. This is an area in the Flevo polder of the Netherlands.

Data folder: **Flevoland**

6

Remote sensing data is affected by geometric distortions due to sensor geometry, scanner and platform instabilities, earth rotation, earth curvature, etc. Some of these distortions are corrected by the image supplier and others can be corrected referencing the images to existing maps.

Remotely sensed images in raw format contain no reference to the location of the data. In order to integrate these data with other data in a GIS, it is necessary to correct and adapt them geometrically, so that they have comparable resolution and projections as the other data sets.

The geometry of a satellite image can be 'distorted' with respect to a north-south oriented map:

- Heading of the satellite orbit at a given position on Earth (rotation).
- Change in resolution of the input image (scaling).
- Difference in position of the image and map (shift).
- Skew caused by earth rotation (shear).

The different distortions of the image geometry are not realized in certain sequence, but happen all together and, therefore, cannot be corrected stepwise. The correction of 'all distortions' at once, is executed by a transformation which combines all the separate corrections. The transformation most frequently used to correct satellite images, is a first order transformation also called *affine transformation*. This transformation can be given by the following polynomials:

$$X = a_0 + a_1 r_n + a_2 c_n$$

$$Y = b_0 + b_1 r_n + b_2 c_n$$

where:

r_n is the row number, c_n is the column number, X and Y are the map coordinates.

To define the transformation, it will be necessary to compute the coefficients of the polynomials (e.g. a_0 , a_1 , a_2 , b_0 , b_1 and b_2). For the computations, a number of points have to be selected that can be located accurately on the map (X , Y) and which are also identifiable in the image (`row`, `column`). The minimum number of points required for the computation of coefficients for an affine transform is three, but in practice you need more. By selecting more points than required, this additional data is used to get the optimal transformation with the smallest overall positional error in the selected points. These errors will appear because of poor positioning of the mouse pointer in an image and by inaccurate measurement of coordinates in a map. The overall accuracy of the transformation is indicated by the average of the errors in the reference points: The so-called *Root Mean Square Error (RMSE)* or *Sigma*.

If the accuracy of the transformation is acceptable, then the transformation is linked with the image and a reference can be made for each pixel to the given coordinate system, so the image is geo-referenced. After geo-referencing, the image still has its original geometry and the pixels have their initial position in the image, with respect to row and column indices.

In case the image should be combined with data in another coordinate system or georeference, then a transformation has to be applied. This results in a 'new' image where the pixels are stored in a new row/column geometry, which is related to the other georeference (containing information on the coordinates and pixel size). This new image is created by applying an interpolation method called *resampling*. The interpolation method is used to compute the radiometric values of the pixels, in the new image based on the DN values in the original image.

After this action, the new image is called *geo-coded* and it can be overlaid with data having the same coordinate system.

! In case of satellite imagery from optical systems, it is advised to use a linear transformation. Higher order transformations need much more computation time and in

many cases they will enlarge errors.

The reference points should be well distributed over the image to minimize the overall error. A good choice is a pattern where the points are along the borders of the image and a few in the center.

6.4.1 Geo-referencing using corner coordinates

When an image (raster map) is created, either by a satellite, airborne scanner or by an office scanner, the image is stored in row and column geometry in raster format. There is no relationship between the rows/columns and real world coordinates (UTM, geographic coordinates, or any other reference map projection). In a process called *geo-referencing* the relation between row and column numbers and real world coordinates are established.

In general, four approaches can be followed:

- *Georeference Corners*: Specifying the coordinates of the lower left (as x_{min} , y_{min}) and upper right corner (as x_{max} , y_{max}) of the raster image and the actual pixel size (see Figure 6.7). A georeference corners is always north-oriented and should be used when rasterizing point, segment, or polygon maps and usually also as the output georeference during resampling.
- *Georeference Tiepoints* (see section 6.4.2): Specifying reference points in an image so that specific row/column numbers obtain a correct X, Y coordinate. All other rows and columns then obtain an X, Y coordinate by an affine, second order or projective transformation as specified by the georeference tiepoints. A georeference tiepoints can be used to add coordinates to a satellite image or to a scanned photograph and when you do not have a DTM. This type of georeference can be used to resample (satellite) images to another georeference (e.g. to a georeference corners) or for screen digitizing.
- *Georeference Direct Linear*: A georeference direct linear can be used to add coordinates to a scanned photograph which was taken with a normal camera, and when you have an existing DTM to also correct for tilt and relief displacement (Direct Linear Transformation). With this type of georeference you can for instance add coordinates to small format aerial photographs without fiducial marks and for subsequent screen digitizing or to resample the photograph to another georeference (e.g. to a georeference corners).
- *Georeference Orthophoto*: A georeference orthophoto can be used to add coordinates to a scanned aerial photograph with fiducial marks, taken with a photogrammetric camera with known principal distance, and when you have an existing DTM to also correct for tilt and relief displacement (Differential rectification). You can use this type of georeference to add coordinates to professional near vertical aerial photographs and further monoplotting on the photograph and as the first step in creating an orthophoto (the second step is to resample it to a north-oriented georeference).
- *Georeference 3D*: A georeference 3D can be used to obtain a 3 dimensional picture of a raster map overlaid by vector maps. A DTM of the area is required. Georef 3D will be explained in section 10.5. (not in this exercise; see for this ILWIS UsersGuide)

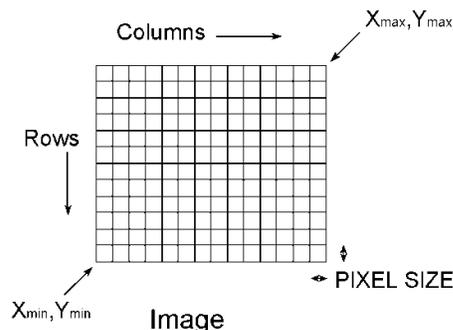


Figure 6.3: The principle of Georeference Corners. The coordinates are defined by the coordinates of the corners and the pixel size.

In this exercise we will create a north-oriented georeference corners to which all the maps will be resampled in a later stage.

- 
- Open the File menu in the Main window and select **Create, GeoReference**. The Create GeoReference dialog box is opened.
 - Enter **Flevo** for the GeoReference Name. Note that the option **GeoRef Corners** is selected.
 - Type: **GeoReference Corners for the Flevo Polder** in the text box **Description**.
 - Click the (yellow) **Create** button next to the **Coordinate system** list box. The **Create Coordinate System** dialog box is opened.

As already mentioned before, a *georeference* is linked to a *coordinate system*; the coordinate system contains the minimum and maximum coordinates of the study area, and optional projection parameters. It is always advised to create a coordinate system of type *Projection* for the study area in which you are working, even if you do not have projection information, instead of using the coordinate system *Unknown*. This is important if you want to transform later on other data that have different coordinate systems and/or projections.

- 
- Type: **Flevo** for the **Coordinate System Name**.
 - Type: **Coordinate system for the Flevo Polder** in the text box **Description**, select **CoordSystem Projection** and click **OK**. The **Coordinate System Projection** dialog box is opened. Select: **Define Coordinate Boundaries**
 - In the dialog box type the following coordinates for **Min X, Y: 180000** and **500000**, and for **Max X, Y: 210000** and **550000** and click **OK**. In a later stage we may add possible projection information.
 - You are back in the **Create Georeference** dialog box. Enter for the **Pixel size: 10 (meters)**.
 - Use the same coordinates for **Min X, Y** and **Max X, Y** as for **Coordinate System Flevo**, clear the check box **Center of Corner Pixels** and click **OK**. A georeference of 54500 lines and 3000 columns will be created.

The entered values for the X and Y coordinates, indicate the extreme corners of the lower left and upper right pixels and not the center of these pixels.

6.4.2 Geo-referencing a raster image using reference points

In many cases, you do not know the X, Y coordinates of the corners of the map or image. You only know the coordinates for some points in the image. In this case geo-referencing will be done by specifying *reference points (tiepoints)* that relate for distinct points their row/column number with the corresponding X, Y coordinate (Figure 6.4).

For a number of points that can be clearly identified, both in the image as on a topographic map, the coordinates are determined. Reference points have to be carefully selected. Corresponding locations of objects in the image and on a map, have to be carefully addressed and measured, especially when there is a large time difference between the production of the map and the acquisition date of the image. Examples of good reference points are highway intersections, fixed river crossings (bridges), railway crossings, runways/airports, etc.

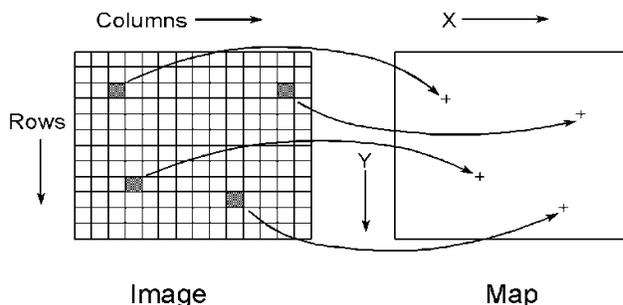


Figure 6.4: Geo-referencing using reference points or tiepoints.

The entered reference points are used to derive a polynomial transformation of the first, second or third order. Each transformation requires a minimum number of reference points (3 for affine, 6 for second order and 9 for third order polynomials). In case more points are selected, the residuals and the derived *Root Mean Square Error (RMSE)* or *Sigma*, may be used to evaluate the calculated equations.

After geo-referencing, the corresponding X, Y coordinate is returned each time a pixel is addressed in the display window. Note that the scanned image is not actually transformed. The original arrangement of the pixels is still intact and no resampling has taken place yet.

In the following exercises we will use raster map `Polder` which is a scanned topographic map of part of the Flevopolder near Vollenhove, the Netherlands. An analog map of the area is also available (see Figure 6.5).

- Display the scanned topographic map `Polder` and zoom in to see more detail. As you can see this map is a little bit rotated.
- Move the mouse pointer over the map and note that on the `Status bar` of the map window, the row and column location of the mouse pointer is shown and the message “No Coordinates”. This means that the image is not geo-referenced yet.

Study the coordinates on the analog topographic map. The georeference is going to be created using reference points. Selection of the *reference points* (in ILWIS also called *tiepoints* or *Ground Control Points (GCPs)*) is made on a reference map.

- In the map window, open the `File` menu and select `Create, GeoReference`. The `Create Georeference` dialog box is opened.
- Enter `Flevo2` for the `GeoReference Name`.
- Type `Georeference Tiepoints for the Flevo Polder` in the text box `Description`.
- Check if the option `Georef Tiepoints` is selected.
- Select `Coordinate System Flevo` (the one you created in the previous exercise) and click `OK`.

The `GeoReference TiePoints` editor (Figure 6.6) is opened. It consists of the map window, in which you will see the mouse pointer as a large square, and a table with the columns `X`, `Y`, `Row`, `Column`, `Active`, `DRow`, `Dcol`. For more information about the `Georeference Tiepoints` editor see section 6.4.3 or the ILWIS Help topic “Georeference Tiepoints editor: Functionality”.

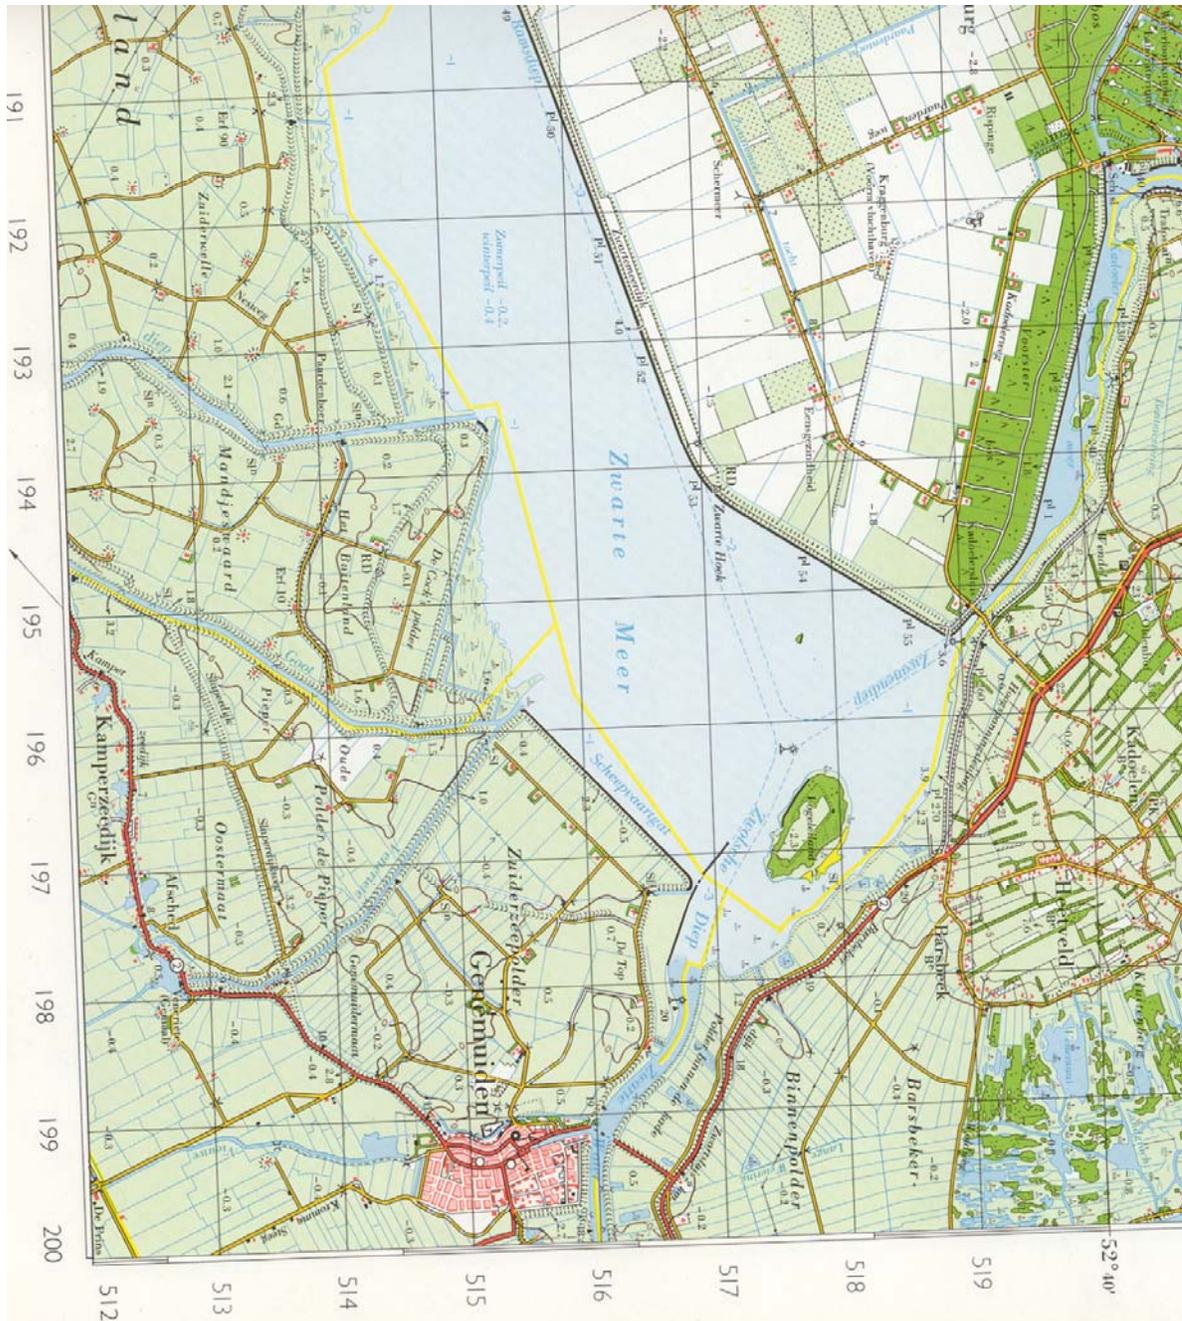


Figure 6.5: The topographic map of the Flevo Polder (The Netherlands). Note that the coordinates are in kilometers.

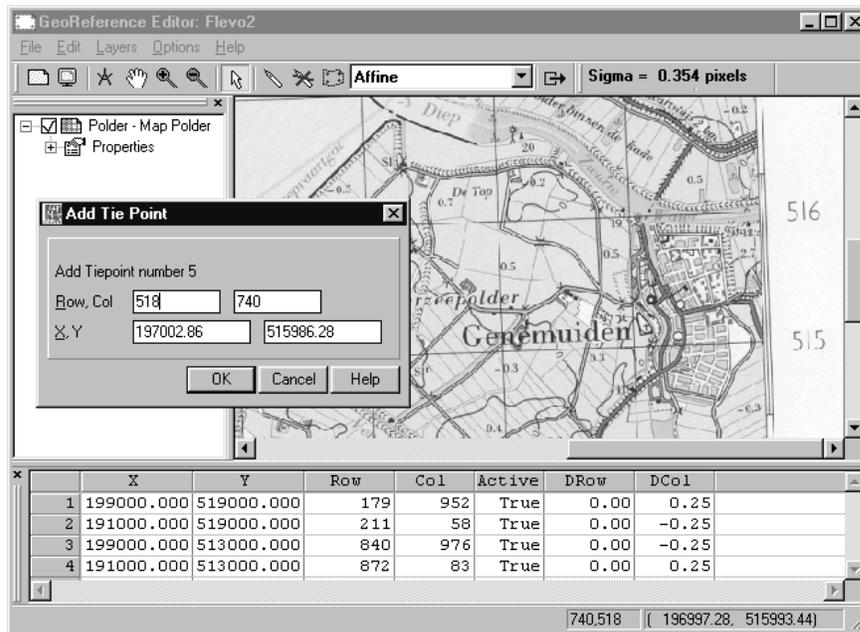


Figure 6.6: Georeference Tiepoints editor.

- Zoom in on the map to a point where horizontal and vertical gridlines intersect (for example P1: $x=191000$ and $y=514000$).
- Locate the mouse pointer exactly at the intersection and click. The Add Tie Point dialog box appears. The row and column number of the selected pixel are already filled out. Now enter the correct X, Y coordinates for this pixel ($x=191000$ and $y=514000$). Then click OK in the Add Tie Point dialog box. The first reference point (tiepoint) will appear in the tiepoints table.
- Repeat this procedure for at least 3 more points. You can use for instance the following points:
 - P2 ($x=199000$, $y=514000$);
 - P3 ($x=199000$, $y=519000$);
 - P4 ($x=191000$, $y=519000$).
- Accept the default **Affine Transformation** method and close the Georeference Tiepoints editor by pressing the Exit Editor button  in the toolbar when the Sigma value is less than 1 pixel
- You will return to the map window. Move the mouse pointer over the map and note that on the Status bar of the map window the row and column location of the mouse pointer as well as the X and Y coordinates are shown. The image is now geo-referenced.

6.4.3 Image-to-image registration

Accurate topographic maps are often not available to perform a standard geometric correction or georeferencing. In order to combine images of different dates and/or from different sensor systems, an *image-to-image registration* is executed. In an image-to-image registration, the distortions are not removed and no referencing to a coordinate system is needed. In image registration, one image is selected as reference and another image is matched to that one. The base image is often referred to as *master image* and the other one is called the *slave image*. The slave image is resampled to the master image. This requires the availability of objects identifiable on both images (Figure 6.7).

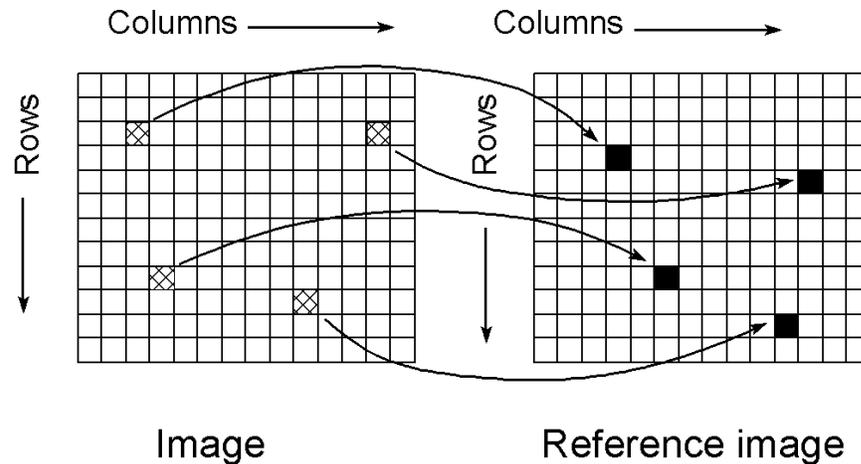


Figure 6.7: Image-to-image registration.

The registration operation does not imply correcting any geometric errors inherent in the images, since geometric errors are immaterial in many applications e.g. in change detection. For referencing an image to any geometric system, *tiepoints* are needed. These points should not be named ground control points or geo-reference points, since no relation is specified between the selected point and a location expressed in any map projection system. However, the master image has to be linked in an earlier process to a reference system. The slave image will inherit this reference system after registration. The definition of the transformation used for the registration and the resampling is further identical to the geo-referencing. The transformation links the row/column geometry of the slave to the row/column geometry of the master.

Now that the scanned topographic map **Polder** is georeferenced, we will use it (as master) to georeference the images **spotb1**, **spotb2** and **spotb3** (the slaves) of the same region.

☞

- Open satellite image **spotb3**.
- In the map window, open the File menu and select **Create, GeoReference**. The **Create GeoReference** dialog box is opened.
- Enter **Flevo3** for the **GeoReference Name**.
- Type: **GeoReference for the Flevo Polder** in the text box **Description**.
- Select the option **GeoRef Tiepoints**.
- Select the **Coordinate System Flevo** and click **OK**. The **Georeference Tiepoints** editor is appears.
- Position the window **spotb3** and **Flevo** next to each other so that you can see both clearly.

At this point you can start entering the reference points (tiepoints) by clicking pixels in the slave (to obtain Row/Column numbers) and then clicking at the same position in the master (to obtain X,Y coordinates).

The following procedures can be used to enter tiepoints in the **Georeference Tiepoints** editor (Figure 6.10):

- Firstly, click with the mouse pointer at a recognizable point in the background map without coordinates (the slave). The **Add Tie Point** dialog box appears. In this dialog box, the row and column values at the position of the click are shown. When already sufficient tiepoints are selected, the dialog box will come up with a suggestion for the XY coordinates (which should never be accepted without reflection).

Then:

- Click at the same position in a map of the same area which is already georeferenced (the master) and which is displayed in another map window (in this case the map `Polder`), or;
- Read the correct XY coordinates for this point from an analog paper map or a table, and type these XY coordinates in the dialog box.



- Choose a location in the slave image `Spotb3` for which you can find the coordinate pair in the master image `Polder`.
- Zoom in on the chosen location, and click the reference point in the slave image as accurately as possible. The **Add Tie Point** dialog box appears.
- Go to the map window with the master raster map and click the pixel, which corresponds to that position. In the **Add Tie Point** dialog box, the correct X and Y coordinates are filled out now.
- Click **OK** to close the **Add Tie Point** dialog box.
- Repeat this procedure for at least 5 other reference points.

When three control points are entered, a **Sigma** and residuals (`DRow` and `DCol`) are calculated. Columns `DRow` and `DCol` show the difference between calculated row and column values and actual row and column values in pixels. Very good control points have `DRow` and `DCol` values of less than 2 pixels.

The **Sigma** is calculated from the `DRow` and `DCol` values and the degrees of freedom, and gives a measure for the overall accountability or credibility of the active tiepoints. The overall sigma value indicates the accuracy of the transformation.

Based on the expected geometric distortions in the background image a transformation method can be selected. By default, an **Affine** transformation is used. You can choose another transformation method in the **Transformation** dialog box (**Edit** menu, or context-sensitive menu) as well as in the drop-down list box in the toolbar of the **GeoReference Tiepoints** editor.



For satellite images an affine transformation will usually do;
For a scanned photograph (without DTM), a projective transformation is recommended.

A tiepoint can be selected or excluded from the transformation computation by putting `True` or `False` respectively in the column `Active`. By excluding tiepoints, you can for instance evaluate the performance of the transformation; assuming that the excluded tiepoints are correctly chosen. Furthermore, in the **Georeference Tiepoints** editor, good tiepoints are shown in green, medium good tiepoints are shown in yellow, ‘bad’ tiepoints are shown in red, and passive tiepoints are shown in blue. You can set these colors in the **Customize Georeference Tiepoints** editor dialog box.

A tiepoint is considered well positioned if:

$$\sqrt{(DRow^2 + DCol^2)} \leq (1.2 * Sigma)$$

A tiepoint is considered to have a medium error if:

$$\sqrt{(1.2 * Sigma)} < \sqrt{(DRow^2 + DCol^2)} \leq (2 * Sigma)$$

A tiepoint is considered to have a large error if:

$$\sqrt{(DRow^2 + DCol^2)} > (2 * Sigma)$$

A tiepoint can be deleted by clicking its record number and by selecting **Edit**, **Delete Tiepoint**, or by pressing **<Delete>** on the keyboard. If confirmed, the tiepoint will be deleted.



- Inspect the DRow and DCol values and inspect the Sigma.
- Close the GeoReference Tiepoints editor when finished.
- You will return to the map window. See that the image Spotb3 now has coordinates. Move the mouse pointer around in the image and determine for some easy identifiable pixels the accuracy of the transformation. Use the zoom in option if necessary.

The georeference F1evo3 created for Spotb3 can also be used for other bands (Spotb1 and Spotb2) since all bands refer to the same area and have the same number of rows and columns.



- In the Catalog click Spotb1 with the right mouse button and select Properties. In the Properties of Raster Map "Spotb1" sheet, change the GeoReference of the image to F1evo3. Click: Apply, OK.
- Do the same for Spotb2 and close all windows when all three SPOT bands are georeferenced

6.4.4 Geo-coding a raster image by resampling

After geo-referencing an image, the image will have coordinates for each pixel, but its geometry is not corrected for geometric distortions and not adapted to a master map or image. To create a distortion free adapted image, the transformation that is defined during geo-referencing (see previous exercise) is executed. This process, geo-coding, results in a new image in which the pixels are arranged in the geometry of the master image or map, and the resolution is equal to the resolution of the master image or chosen in case the master is a topographic map. The radiometric values or pixel values of the new image are found by resampling the original image using a chosen interpolation method (Figure 6.8).

The value for a pixel in the output image is found at a certain position in the input image. This position is computed using the inverse of the transformation that is defined in the geo-referencing process.

Whether one pixel value in the input image, or a number of pixel values in that image are used to compute the output pixel value, depends on the choice of the interpolation method.

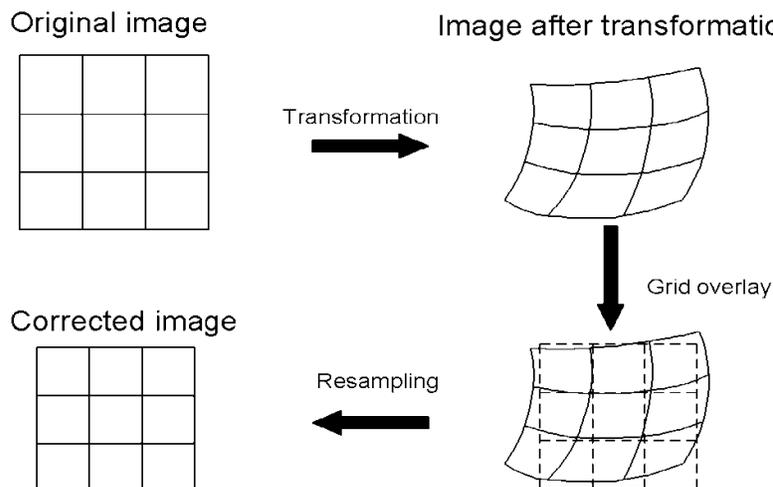
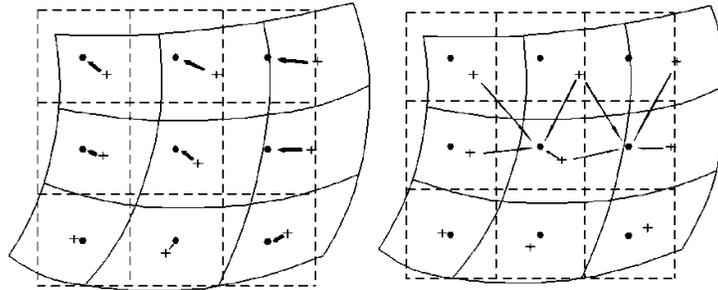


Figure 6.8: The process of geo-coding.

There are three different interpolation methods:

- **Nearest Neighbour.** In the nearest neighbor method, the value for a pixel in the output image is determined by the value of the nearest pixel in the input image (Figure 6.13 A).
- **Bilinear.** The bilinear interpolation technique is based on a distance dependent weighted average, of the values of the four nearest pixels in the input image (Figure 6.9 B).
- **Bicubic.** The cubic or bicubic convolution uses the sixteen surrounding pixels in the input image. This method is also called cubic spline interpolation.



A.

B.

Figure 6.9: Interpolation methods. A: Nearest Neighbour, B: Bilinear interpolation.

For more detailed information about the different interpolation methods see the ILWIS Help topics “Resample: Functionality” and “Resample: Algorithm”.

The images used in the previous exercises are the scanned topographic map `Polder` and the three Spot bands `Spotb1`, `Spotb2` and `Spotb3`. Up to now, they all use a georeference tiepoints. Now we will resample all maps to a north-oriented georeference. First the map `Polder` will be resampled to the north-oriented georeference corners `Flevo`; then you will resample the Spot images to the same georeference corners `Flevo`.



- Display the scanned topographic map `Polder` and indicate with a drawing how the map should be rotated.
- Close the map window.
- Double-click the **Resample** operation in the Operation-list in order to display the **Resample Map** dialog box.
- Enter `Polder` for the **Raster Map** and select the **Resampling Method** `Nearest Neighbour`.
- Type `Polder_resampled` for the **Output Raster Map**, select **GeoReference** `Flevo` and click **Show**. The resampling starts and after the resampling process the **Display Options - Raster Map** dialog box is opened.
- Click **OK** to display the map.
- Repeat this procedure for `Spotb1`, `Spotb2` and `Spotb3` successively. Call the output maps for instance `Spotb1_resampled`, `Spotb2_resampled` and `Spotb3_resampled`.
- Close the map windows when you finished the exercise.



What will happen to output pixel values if you use a bilinear or a bicubic interpolation during resampling?

All the maps now have the same georeference `Flevo` and it is possible now to combine the maps with other digitized information of the area (that is rasterized on georeference `Flevo` as well).

Additional

Adding projection information to a coordinate system **(additional exercise)**

Coordinate system `Flevo` as created in section 6.4.1 and used by georeference `Flevo`, did not yet have projection information. However, we know that the topographic map of the Flevopolder area uses the standard Dutch RD projection (Rijksdriehoeksmeeting). We will now add this information to coordinate system `Flevo`.



- In the **Catalog**, double-click coordinate system `Flevo`.
- In the appearing dialog box **Coordinate System Projection “Flevo”**, click the **Projection** button. The **Select Projection** dialog box is opened and shows a list of projections.
- Select **Dutch RD** and click **OK**.
- For the **Dutch RD** projection, datum and ellipsoid are already fixed, the information appears at the bottom in the dialog box. Click **OK** in the **Coordinate System Projection “Flevo”** dialog box.
- Display the map `Polder` again and note that the **Status bar** now also displays geographical coordinates in latitudes and longitudes.
- Close the map window afterwards.

Summary: Geometric corrections and image referencing

- Remote sensing data is affected by geometric distortions due to sensor geometry, scanner and platform instabilities, earth rotation, earth curvature, etc.
- Remotely sensed images in raw format contain no reference to the location of the data. In order to integrate these data with other data in a GIS, it is necessary to correct and adapt them geometrically, so that they have comparable resolution and projections as the other data sets.
- The geometric correction is executed by a transformation. The transformation most frequently used to correct satellite images, is a first order transformation also called *affine transformation*.
- To define the transformation, it will be necessary to compute the coefficients of the polynomials. For the computations, a number of points have to be selected that can be located accurately on the map (`X`, `Y`) and which are also identifiable in the image (`row`, `column`).
- The overall accuracy of the transformation is indicated by the average of the errors in the reference points: The so-called *Root Mean Square Error (RMSE)* or *Sigma*.
- After geo-referencing, the image still has its original geometry and the pixels have their initial position in the image, with respect to row and column indices.
- In case the image should be combined with data in another coordinate system or with another georeference, a transformation has to be applied. This results in a “new” image where the pixels are stored in a new line/column geometry, which is related to the other georeference. This new image is created by means of *resampling*, by applying an interpolation method. The interpolation method is used to compute the radiometric values of the pixels, in the new image based on the DN values in the original image.
- After this action, the new image is called *geo-coded* and it can be overlaid with data having the same coordinate system.

6.5 Multi-band operations

To enhance or extract features from satellite images, which cannot be clearly detected in a single band, you can use the spectral information of the object recorded in multiple bands. These images may be separate spectral bands from a single multi-spectral data set, or they may be individual bands from data sets that have been recorded at different dates or using different sensors.

The operations of addition, subtraction, multiplication and division, are performed on two or more co-registered images of the same geographical area. This section deals with multi-band operations. The following operations will be treated:

- *Vegetation indices*, some of them are more complex than ratio's only.

6.5.1 Normalized Difference Vegetation Index

Ratio images are often useful for discriminating subtle differences in spectral variations, in a scene that is masked by brightness variations. Different band ratios are possible, given the number of spectral bands of the satellite image. The utility of any given spectral ratio depends upon the particular reflectance characteristics of the features involved and the application at hand. For example a near-infrared / red ratio image might be useful for differentiating between areas of stressed and non-stressed vegetation.

Various mathematical combinations of satellite bands, have been found to be sensitive indicators of the presence and condition of green vegetation. These band combinations are thus referred to as vegetation indices. Two such indices are the simple *Vegetation Index (VI)* and the *Normalized Difference Vegetation Index (NDVI)*. Both are based on the reflectance properties of vegetated areas as compared to clouds, water and snow on the one hand, and rocks and bare soil on the other. Vegetated areas have a relatively high reflection in the near-infrared and a low reflection in the visible range of the spectrum. Clouds, water and snow have larger visual than near-infrared reflectance. Rock and bare soil have similar reflectance in both spectral regions. The effect of calculating VI or the NDVI is demonstrated in Table 6.4

Table 6.4: Reflectance versus ratio values.

	TM Band 3	TM Band 4	VI	NDVI
Green vegetation	21	142	121	0.74
Water	21	12	-9	-0.27
Bare soil	123	125	2	0.01

It is clearly shown that the discrimination between the 3 land cover types is greatly enhanced by the creation of a vegetation index. Green vegetation yields high values for the index. In contrast, water yield negative values and bare soil gives indices near zero. The NDVI, as a normalized index, is preferred over the VI because the NDVI is also compensating for changes in illumination conditions, surface slopes and aspect.

To create a NDVI map, Landsat ETM imagery of Georgia is used. The following bands are used: ETM_band_3 and ETM_band_4.

The creation of the NDVI is done with the **map calculator**.



- Double-click the Map Calculation operation in the Operation-list. The Map Calculation dialog box is opened.
- In the Expression text box type the following map calculation formula:

$$\frac{(\text{Dedo_ETM99_band_4} - \text{Dedo_ETM99_band_3})}{(\text{Dedo_ETM99_band_4} + \text{Dedo_ETM99_band_3})}$$
- Enter **Dedo_NDVI** for Output Raster map, select system Domain Value, change the Value Range to -1 and +1, and the Precision to 0.01.
- Click Show. The map NDVI is calculated and the Display Options - Raster Map dialog box is opened.
- In the Display Options - Raster Map dialog box select Pseudo as Representation. In this representation the colors range from blue (for the low DNs) through green (for the medium DNs) to red (for the high DNs). Click OK. The map **Dedo_NDVI** is displayed.
- In the map window open the pixel information window and add the maps **Dedo_ETM99_band_3** and **Dedo_ETM99_band_4** to it.
- Move through the image with the mouse pointer and study the DN values of the input maps and how they are combined in the NDVI map.
- Close all windows after finishing the exercise.



You can also use the NDVI function in MapCalc:
NDVImap=NDVI (Band_3, Band_4). Or use NDVI in the Operations List

Summary: Multi-band operations

- To enhance or extract features from satellite images, which cannot be clearly detected in a single band, you can use the spectral information of the object recorded in multiple bands.
- The operations of addition, subtraction, multiplication and division are performed on two or more co-registered images of the same geographical area.
- *Brightness variations*: a ratioed image of the scene effectively compensates for the brightness variation, caused by the differences in topography and emphasized by the color content of the data.
- *Normalized Difference Vegetation Index (NDVI)*: Ratio images are often useful for discriminating subtle differences in spectral variations, in a scene that is masked by brightness variations.
The utility of any given spectral ratio, depends upon the particular reflectance characteristics of the features involved and the application at hand.
- Various mathematical combinations of satellite bands, have been found to be sensitive indicators of the presence and condition of green vegetation. These band combinations are thus referred to as vegetation indices. Two such indices are the simple *vegetation index (VI)* and the *normalized difference vegetation index (NDVI)*.
- The NDVI, as a normalized index, is preferred over the VI because the NDVI is also compensating for changes in illumination conditions, surface slopes and aspect.

6.6 Image classification

Measured reflection values in an image depend on the local characteristics of the earth surface; in other words there is a relationship between land cover and measured reflection values. In order to extract information from the image data, this relationship must be found. The process to find the relationship is called *classification*. Classification can be done using a single band, in a process called *density slicing*, or using many bands (*multi-spectral classification*).

6.6.1 Density slicing

In theory, it is possible to base a classification on a single spectral band of a remote sensing image, by means of *single band classification* or density slicing.

Density slicing is a technique, whereby the DNs distributed along the horizontal axis of an image histogram, are divided into a series of user-specified intervals or slices. The number of slices and the boundaries between the slices depend on the different land covers in the area. All the DNs falling within a given interval in the input image are then displayed using a single class name in the output map.

The image used in the exercise is a SPOT-XS band 3 of an area in the Flevopolder near Vollenhove, The Netherlands.

Firstly, the ranges of the DN values representing the different land cover types have to be determined. Secondly, a domain group with the slice boundaries, names and codes has to be created.



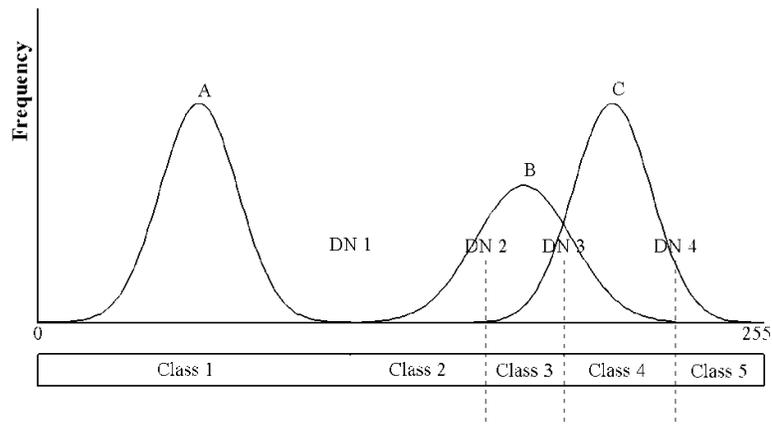
- Display the image Spotb3.
- Move through the image and write down the values of water and land. Fill in the data in Table 6.5. Zoom in if necessary. Select also some mixed pixels along the land/water boundary.

Table 6.5: DN values per cover class of Spot Band 3.

Cover class	DN values of SPOT-XS Band 3
Land	
Water	

By a visual interpretation of the spectral reflectance of the cover classes in band 3, differences are observed with regard to their gray tones. Density slicing will only give reasonable results, if the DN values of the cover classes are not overlapping each other, as indicated in Figure 6.10.

In Figure 6.10, a clear distinction is possible between cover type A and cover types B/C (slice with minimum number at position DN 1), as the distribution of the pixels over the digital range is different from cover type B and C. If the distributions of cover types B and C are studied, an overlap can be observed. To discriminate between B and C several options are possible, but all will result in an improper slice assignment. To classify cover class B, the slice DN 1 to DN 2 will include the lower DN values of B and the higher DN values of B are excluded. To include all the pixels belonging to B, the slice assignment is from DN 1 to DN 4. This slice assignment is including even more pixels that belong to cover class C. When trying to classify C, the same problem occurs.



Digital range splitted in a number of slices

Figure 6.10: Distribution of ground cover classes over the digital range.

Interactive slicing

Before a slicing operation is performed it is useful to firstly apply a method which shows the map, as if it was classified, by manipulating the representation.

- Display map `spotb3` with default `Gray Representation`.
- Reopen the `Display Options - Raster Map` dialog box and click the `Create` button next to the `Representation` list box. The `Create Representation` dialog box is opened.
- Type for `Representation Name`: `spotb3` and click `OK`. The `Representation Value` editor is opened.
- Open the `Edit` menu and select `Insert Limit`. The `Insert Limit` dialog box is opened. Enter the limit you have written down in Table 6.5 to differentiate land from water, and select the color `Blue`.
- Insert also the other limit shown in the table and select the color `Green` for it.
- Open the list box, which appears after you click the word `Stretch` between the two limits. Select `Upper`. Do the same for the other.
- Click the `Redraw` button  in the map window to see the map with the new representation.
- When satisfied, close the `Representation Value` editor.

The representation can thus be edited, and the result shown on the screen (by using the `Redraw` button in the map window). This allows you to interactively select the best boundaries for the classification.

Remark: Try to do the same for the lake area in Landsat ETM: `Dedo_etm99_band5`

Slicing operation

Now you will do the actual classification, using the `Slicing` operation.



- Double-click the Slicing operation in the Operation-list. The Map Slicing dialog box is opened.
- Select **spotb3** in the Raster Map list box.
- Type *Slices* in the Output Raster Map text box.
- Type *Classified SPOT* image in the Description text box.
- Click the **Create Domain** button. The Create Domain dialog box is opened.
- Type *Slices* in the Domain Name text box.
- Make sure that the **Group** check box is selected and accept all other defaults by clicking the OK button. The Domain Group editor is opened.

For a group domain, you can enter an upper boundary, a name and a code for each class/group. This domain will be used to slice or classify the raster map *spotb3*. The upper boundaries and group names that you have entered in Table 6.14 will be used in this exercise.



- Click the  button in the toolbar of the Domain Group editor. The Add Domain Item dialog box appears.
- Type the value from Table 6.5 in the Upper Bound text box.
- Type *Water* in the Name text box. The use of a Code and a Description is optional. It will not be used now.
- Click OK.
- Press the **Insert** key of the keyboard to enter the next upper boundary and name.

When you create a domain group, a representation is also created. It is possible to edit the colors of the output groups/classes from the Domain Group editor, or by opening the representation in the Catalog. To edit the representation:



- Press the Open Representation button  in the toolbar. The Representation Class editor appears, showing the groups/classes that are present in the domain with different colors.
- Double-click a color box in the editor and select a pre-defined color, or create your own color.
- Exit the Representation Class editor when you are satisfied with the colors.
- Exit the Domain Group editor.
- Click **Show** in the Slicing dialog box. The Display Options - Raster Map dialog box is opened.
- Accept the defaults by clicking OK.
- In the map window open the pixel information window and compare the values of the original map with the names in the classified map.
- Close the windows after finishing the exercise.

If you are not satisfied with the results, perform the same procedure again with a new set of boundaries or by editing the existing slices.

6.6.2 Piece-wise linear stretching

In the same way as described in the exercise on interactive slicing, you can also do a piece-wise linear stretching of images by creating a new representation. While for interactive slicing, we used ‘upper’ boundary values (Figure 6.11 left) for piece-wise linear stretch, we will use the ‘stretch’ option between limits (Figure 6.11 right).

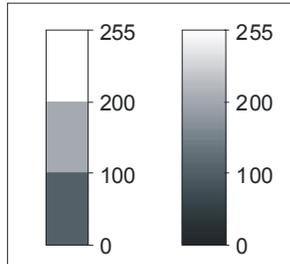


Figure 6.11: Upper boundary method (left) and stretched method (right).

The *piece-wise linear contrast stretch* is very similar to the linear contrast stretch, but the linear interpolation of the output values is applied between user-defined DN values. This method is useful to enhance a certain cover type, for example water.

- In the Catalog, double-click `spotb3`. The Display Options - Raster Map dialog box appears.
- In the Display Options - Raster Map dialog box click the (yellow) Create Representation button. The Create Representation dialog box is opened.
- In the Create Representation dialog box type `Piecewise` for the Representation Name accept the default `Image Domain` and click OK. The Representation Value editor is opened.
- Create a representation with the following limits and colors:
 - 255 (`LawnGreen`)
 - 68 (`ForestGreen`)
 - 60 (`MidnightBlue`)
 - 0 (`Cyan`)
 Make sure that in between two limits, the coloring method `Stretch` is selected and increase the `Stretch Steps` to 10 in `Edit`.
- Exit the Representation Value editor to return to the Display Options - Raster Map dialog box.
- Click OK to display the map.
- Compare the output map with map `Slices` created in the previous exercise and close both map windows when you are finished.

6.6.3 Multi-spectral image classification

Multi-spectral image classification is used to extract thematic information from satellite images in a semi-automatic way. Different methods for image classification exist; some of them are based on the theory about probabilities. Looking at a certain image pixel in M bands simultaneously, M values are observed at the same time. Using multi-spectral SPOT images, where $M=3$, three reflection values per pixel are given. For instance, (34, 25, 117) in one pixel, in another (34, 24, 119) and in a third (11, 77, 51). The values found for 1 pixel in several bands are called *feature vectors*. It can be recognized that the first two sets of values are quite similar and that the third is different from the other two. The first two probably belong to the same (land cover) class and the third one belongs to another class.

In classification jargon it is common to call the three bands “*features*”. The term features instead of bands is used because it is very usual to apply transformations to the image, prior to

classification. They are called “*feature transformations*” and their results are called “*derived features*”. Examples are: Principal Components, HSI transformations, etc.

In one pixel, the values in the (three) features can be regarded as components of a 3-dimensional vector, the feature vector. Such a vector can be plotted in a 2- or 3-dimensional space called a feature space. Pixels which belong to the same (land cover) class and which have similar characteristics, end up near to each other in a feature space, regardless of how far they are from each other in the terrain and in the image. All pixels belonging to a certain class will (hopefully) form a cluster in the feature space. Moreover, it is hoped that other pixels, belonging to other classes, fall outside this cluster (but in other clusters, belonging to those other classes).

A large number of classification methods exist. To make some order, the first distinction is between *unsupervised* and *supervised classification*. For satellite image applications, the second is generally considered more important.

In order to make the classifier work with thematic (instead of spectral) classes, some “knowledge” about the relationship between classes and feature vectors must be given.

Theoretically, this could be done from a database in which the relationships between (thematic) classes and feature vectors are stored. It is tempting to assume that in the past, enough images of each kind of sensor have been analyzed, as to know the spectral characteristics of all relevant classes. This would mean, for example, that a pixel with feature vector (44, 32, 81) in a multi-spectral SPOT image always means grass, whereas (12, 56, 49) is always a forest pixel.

Unfortunately, the observed feature vectors in a particular image are influenced by a large amount of other factors than land cover, such as atmospheric conditions, sun angle (as function of latitude, time of day, date and terrain relief), soil types, soil humidity, growing stage of the vegetation, wind (affecting orientation of leaves), etc. The problems we meet when trying to take all these influences into account vary from ‘quite easy’ to ‘practically impossible’; at least, vast amounts of additional data (DEMs, soil maps, etc.) would be required.

Sampling

Therefore, supervised classification methods are much more widely used. The process is divided into two phases: a training phase, where the user ‘trains’ the computer, by assigning for a limited number of pixels to what classes they belong in this particular image, followed by the decision making phase, where the computer assigns a class label to all (other) image pixels, by looking for each pixel to which of the trained classes this pixel is most similar.

During the training phase, the classes to be used are defined. About each class some “*ground truth*” is needed: a number of places in the image area that are known to belong to that class. This knowledge must have been acquired beforehand, for instance as a result of fieldwork, or from an existing map (assuming that in some areas the class membership has not changed since the map was produced). If the ground truth is available, training samples (small areas or individual pixels) are indicated in the image and the corresponding class names are entered.

A *sample set* has to be created in which the relevant data regarding input bands (map list), cover classes (domain codes) and background image for selecting the training areas is stored. The map presented below is giving the ground truth information, as well as the cover classes to be used. The images used in this exercise are SPOT images from the area of the Flevo Polder, The Netherlands:

SPOT-XS band 1: spotb1

SPOT-XS band 2: spotb2

SPOT-XS band 3: spotb3

As the selection of the training areas is done on a color composite, it has to be created beforehand.



- Create a map list, which contains the SPOT bands and display this as a color composite. Use **spotb3** for Red, **spotb2** for Green and **spotb1** for Blue.

You may save the color composite as a map view. In the following steps a sample set will be created. This will be used for the actual sampling procedure.



- In the map window, choose **Create, Sample Set** from the File menu.
- In the Sampling dialog box enter **spot_classes** for the Sample Set Name.
- Create a domain for the classes to be sampled by selecting the **Create Domain** button. The **Create Domain** dialog box is opened.
- Enter **spot_classes** for the Domain Name and click OK. The **Domain Class** editor appears.
- In the **Domain Class** editor, click the **Add Item** button. The **Add Domain Item** dialog box is displayed.
- Enter the **Name: Forest**, and the **Code: f**. Click OK. Repeat this for all classes as given in Table 6.6. New classes can always be added later on.

Table 6.6: The land-cover classes in the sample set for a supervised image classification.

Class name	Code
water	w
forest	f
grass land	gl
crop1	c1
crop2	c2
crop3	c3



- After having entered all classes, open the **Representation Class** editor by pressing the **Open Representation** button in the toolbar of the **Domain Class** editor. Select each class by double-clicking it and choose a color from the color list or define your own color.
- Exit the **Representation Class** editor and exit the **Domain Class** editor to return to the **Create Sample Set** dialog box.
- The **MapList** created before is already filled out.
- Click **OK** in the **Create Sample Set** dialog box.

The **Sample Set** editor will be started. Two windows are displayed: A map window showing the false color composite and a window showing the sample set statistics. During the sampling procedure, consult also the topographic map of the Flevoland region, shown in Figure 6.5.



- Zoom in on an area with water and press the **Normal** button (icon with hand & finder up) in the toolbar to return to the normal mode. To select pixels press the left mouse button and drag the cursor. When the mouse button is released, the statistics of the current selection of pixels are shown in the **Sample Statistics** window, **Current Selection** (Figure 6.12).

- Click the right mouse button to reveal the context-sensitive menu and select **Edit**. Select the appropriate class from the class list (`w: water`). Once this selection is added to the class, the total class statistics are shown in the upper part of the table.
- One can always compare the statistics of a sample with the statistics of a class, by selecting the class in the class list.
- Repeat the sampling procedure for a number of water samples and then continue with sampling the other land cover classes.
- If a land cover class has not yet been defined and, therefore, does not appear in the class list, a new class can be created by selecting `<new>`.
- Add a new class (**Name: Urban, Code: u**). Take samples of this class.

Band	Mean	StDev	Nr	Pred	Total
1:	60.4	1.0	230	60	532
2:	37.9	1.0	204	37	532
3:	21.3	0.7	257	21	532

Current Selection:					
Band	Mean	StDev	Nr	Pred	Total
1:	60.4	0.6	116	60	195
2:	37.4	0.5	114	37	195
3:	21.7	0.5	139	22	195

Figure 6.12: The sample set statistics window.

! To select multiple training pixels you can drag rectangles or you can hold the **Ctrl**-key.

The **Sample Statistics** window contains the code and name of the selected class as well as the number of bands. The following statistics are shown for a selected class (see Figure 6.12):

- Mean value of pixels (**Mean**).
- Standard deviation of pixels (**StDev**).
- The number of pixels having the predominant value (**Nr**).
- The predominant pixel value (**Pred**).
- The total number of the selected pixels (**Total**).

Displaying Feature Spaces

The classes can be plotted in distinct colors in the *feature space*, which enables a judgement of whether the classes can really be spectrally distinguished and whether each class corresponds to only one spectral cluster (if not, it is advisable to create sub-classes, which must be joined after the classification is finished). After training is done, a feature space with training samples may look like Figure 6.13.

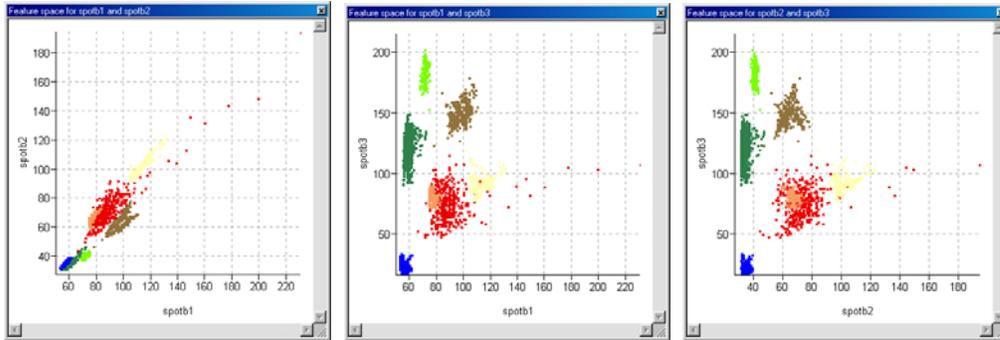


Figure 6.13: Three examples of ILWIS Feature Space windows.

The number of training samples should be between 30 and some hundreds of samples per class, depending on the number of features and on the decision making that is going to be applied. The feature space is a graph in which DN values of one band are plotted against the values of another. To view a feature space:

- Click the **Feature Space** button in the toolbar of the **Sample Set** editor. The **Feature Space** dialog box is opened.
- Select two bands for which a feature space should be created. Select `Spotb1` for the horizontal axis and `Spotb2` for the vertical axis. Click **OK**. The feature space of `Spotb1` against `Spotb2` is displayed.
- Try also the other band combinations for the feature space display.

You will probably see some overlapping classes in the feature space.

- After you have finished the sampling, close all windows and return to the **ILWIS Main** window.

6.6.4 Classification

It is the task of the decision-making algorithm to make a partitioning of the feature space, according to our training samples. For every possible feature vector in the feature space, the program must decide to which of the sets of training pixels this feature vector is most similar. After that, the program makes an output map where each image pixel is assigned a class label, according to the feature space partitioning. Some algorithms are able to decide that feature vectors in certain parts of the feature space are not similar to any of the trained classes. They assign to those image pixels the class label “Unknown”. In case the area indeed contains classes that were not included in the training phase, the result “unknown” is probably more realistic than to make a “wild” guess.

To find the relationship between classes and feature vectors is not as trivial as it may seem. Therefore, various decision-making algorithms are being used; they are different in the way they partition the feature space. Four of them are:

- *The Box classifier* is the simplest classification method: In 2-D space, rectangles are created around the training feature vector for each class; in 3-D they are actually boxes (blocks). The position and sizes of the boxes can be ‘exactly around’ the feature vectors (Min-Max method), or according to the mean vector (this will be at the center of a box) and the standard deviations of the feature vector, calculated separately per feature (this determines the size of the box in that dimension). In both cases, the user is allowed to change the sizes by entering a “multiplication factor”. In parts of the feature space where boxes overlap, it is usual to give priority to the smallest box. Feature vectors in the image that fall outside all boxes will be “unknown”.

- *The Minimum Distance-to-mean classifier*, first calculates for each class the mean vector of the training feature vectors. Then, the feature space is partitioned by giving to each feature vector the class label of the nearest mean vector, according to Euclidean metric. Usually it is possible to specify a maximum distance threshold: if the nearest mean is still further away than that threshold, it is assumed that none of the classes is similar enough and the result will be “unknown”.
- *Minimum Mahalanobis Distance classifier*. For the spectral values of a pixel to be classified, the distances towards the class means are calculated as Mahalanobis distance. The Mahalanobis distance depends on the distances towards class means and the variance-covariance matrix of each class. The class name with the shortest Mahalanobis distance is assigned, if this distance is smaller than the user-defined threshold value. Else, the undefined value is assigned.
- *Gaussian Maximum Likelihood classifiers* assume that the feature vectors of each class are (statistically) distributed according to a “multivariate normal probability density function”. The training samples are used to estimate the parameters of the distributions. The boundaries between the different partitions in the feature space are placed where the decision changes from one class to another. They are called decision boundaries.

In this exercise, different types of land-cover in the area will be detected by implementing a supervised classification, using the sample set prepared in the previous exercise on sampling. The output will be a map with domain type class representing land-cover units in the area.

The four standard techniques for classifying an image in ILWIS are:

- 1) The Box classifier;
- 2) The Minimum Distance classifier;
- 3) The Mahalanobis Distance classifier, and
- 4) The Maximum Likelihood classifier.

Each classifier uses a sample set, only the procedure to classify an image is different.

Classification methods

The box classifier uses class ranges determined by the DN values observed in the training set. These intervals result in rectangular spaces or boxes; hence, the name box classifier.

Having sampled a representative amount of pixels per class, the mean DN values of the class are used, together with the standard deviation of the sample and a multiplication factor, to determine the extent of the box. The boundaries of the box are defined by the product of the standard deviation and the multiplication factor.

A pixel not falling within a box, is not classified. If a pixel lies within two or more overlapping boxes, the pixel is classified according to the smallest box.

The bands used in this exercise are from a SPOT image of the Flevo Polder, The Netherlands:

SPOT-XS band 1: `spotb1`

SPOT-XS band 2: `spotb2`

SPOT-XS band 3: `spotb3`

When performing an automated box classification, the different land cover classes in the area will be detected by implementing a programmed classification rule, using a specified sample set. The output map will be a map with classes representing the land cover units in the area. The sample set, to be used, has been created in the part about sampling.



- In the **Catalog**, click the sample set **spot_classes** with the right mouse button and select **Classify**. The **Classify** dialog box appears.
- Select: **Box Classifier** as **Classification Method**. Accept the default **Multiplication Factor**. Type **spot_Box** for the **Output Raster Map** name and click **Show**. The classification will take place and the **Display Options - Raster Map** dialog box appears.

- Click OK to display the result.
- Repeat the procedure using a larger **Multiplication Factor** (for instance: 10).
- Visually compare the results with the classified map using the default multiplication factor.
- Perform other classifications using the **Minimum Distance**, **Minimum Mahalanobis Distance** and **Maximum Likelihood** method. Classify once without using a **Threshold Distance** and once with a **Threshold Distance** of around 15. Name the output images for instance **Spot_MinDist**, **Spot_MinMahaDist** and **Spot_MaxLikehd**.
- Compare the results with the box classified maps.
- Close all map windows when you have finished the exercise.

Evaluate the reclassified results

The final result is compared with the original values in the bands 1, 2 and 3 of the SPOT image. For the complex and more homogenous land covers the results can be compared with the original bands.



- Display the topographic map **Polder**. Make sure that it is georeferenced. Check this also for the SPOT bands.
- Open the pixel information window and add the maps **Spotb1**, **Spotb2**, **Spotb3** and **Spot_Box**.
- Evaluate for pixels the different cover classes defined and the result of your classification.
- Close the map window and the pixel information window.

To get a better idea of the overall accuracy of the classification, use has to be made of the test set which contains additional ground truth data, which have not been used to train the classifier. Crossing the test set with the classified image and creation of a *confusion matrix*, is an established method to assess the accuracy of a classification.



See the ILWIS Help topic “How to calculate a confusion matrix” for more information about creating a confusion matrix.

It is not recommended to use the same sample map for both the classification and the accuracy assessment, because this will produce figures that are too optimistic.

Post classification operations

The final classification may not meet the expected results. This can be due to spectral overlap in the sample set, or because the classifier has not been trained properly. If this is the case the classification process should be repeated, incorporating for example more spectral classes. Defining sub-classes, when a bi-modal distribution or large standard deviation/variance exists in a class, will improve the classification results. After the classification process these classes can be merged again. If the obtained results are still not according to expectations, incorporation of ancillary non-spectral information might be considered, for example elevation information might be useful to distinguish certain forest types.

In a classified image small areas occur, consisting of one or a few pixels, to which another class label has been assigned, compared to the larger homogeneous classified surrounding areas. Individual non-classified pixels may also occur throughout the classified image. If there is, for example, a need to integrate the results with other data sets, a large number of small

individual units will occur which can not be properly represented on a map. To circumvent these problems the classified image may be generalized in order to remove these small individual areas. Spatial filters are a mean to achieve this objective, e.g. individual pixels may be assigned to the majority of the surrounding pixels.

6.6.5 Unsupervised classification (clustering)

One way to perform a classification is to plot all pixels (all feature vectors) of the image in a feature space, and then to analyze the feature space and to group the feature vectors into clusters. The name of this process is *unsupervised classification*. In this process there is no knowledge about “thematic” land cover class names, such as town, road, potatoes, etc. All it can do, is to find out that there appears to be (for example) 16 different “things” in the image and give them numbers (1 to 16). Each “thing” is called a spectral class. The result can be a raster map, in which each pixel has a class (from 1 to 16), according to the cluster to which the image feature vector of the corresponding pixel belongs.

After the process is finished, it is up to the user to find the relationship between spectral and thematic classes. It is very well possible, that it is discovered that one thematic class is split into several spectral ones, or, worse, that several thematic classes ended up in the same cluster.

Various unsupervised classification (clustering) algorithms exist. Usually, they are not completely automatic; the user must specify some parameters such as the number of clusters (approximately) you want to obtain, the maximum cluster size (in the feature space), the minimum distance (also in the feature space), that is allowed between different clusters, etc. The process “builds” clusters as it is scanning through the image. Typically, when a cluster becomes larger than the maximum size, it is split into two clusters; on the other hand, when two clusters get nearer to each other than the minimum distance, they are merged into one.



- Double-click the **Cluster** operation in the **Operation-list**. The **Clustering** dialog box is opened.
- Select **Spotb1**, **Spotb2** and **Spotb3**. Use the same number of classes as when performing the supervised classification.
- Type **Unsupervised** for the **Output Raster Map** name and click **Show**. The supervised map is calculated and after the calculation the **Display Options - Raster Map** dialog box is opened.
- Accept the defaults in the **Display Options - Raster Map** dialog box by clicking **OK**. The map is displayed.
- Display one of the maps produced with the supervised classification method and compare the results.
- Close the map windows when you have finished the exercise.

Summary: Image classification

- In order to extract information from the satellite images, the relationship between pixel values and land cover types must be found. The process to find the relationship is called *classification*.
- *Density slicing* is a technique, whereby the DNs distributed along the horizontal axis of an image histogram, are divided into a series of user-specified intervals or slices. Density slicing will only give reasonable results, if the DN values of the cover classes are not overlapping each other.
- *Multi spectral image classification* is used to extract thematic information from satellite images in a semi-automatic way. The values of the various bands found for 1 pixel in several bands are called *feature vectors*. Such a vector can be plotted in a 3-dimensional space, called *feature space*. Pixels belonging to the same (land cover) class and having similar characteristics, end up near to each other in the feature space. In order to make the classifier work with thematic (instead of spectral) classes, some “knowledge” about the relationship between classes and feature vectors must be given.

- The classification process is divided into two phases: a *training phase*, where the user “trains” the computer, by assigning for a limited number of pixels to what classes they belong in this particular image, followed by the *decision making phase*, where the computer assigns a class label to all (other) image pixels, by looking for each pixel to which of the trained classes this pixel is most similar.
- About each class some “*ground truth*” is needed: a number of places in the image area that are known to belong to that class. It is the task of the decision- making algorithm to make a partitioning of the feature space, according to our training samples. For every possible feature vector in the feature space, the program must decide to which of the sets of training pixels this feature vector is most similar. After that, the program makes an output map where each image pixel is assigned a class label, according to the feature space partitioning.
- Various decision making algorithms are used: *box classifier*, *minimum distance to mean classifier*, *minimum Mahalanobis distance classifier* and *Gaussian maximum likelihood classifier*.
- *Unsupervised classification*. One way to perform a classification is to plot all pixels (all feature vectors) of the image in a feature space, and then to analyze the feature space and to group the feature vectors into clusters. In this process there is no knowledge about “thematic” land cover class names. The user must specify some parameters such as the number of clusters (approximately) you want to obtain, the maximum cluster size (in the feature space), the minimum distance (also in the feature space), that is allowed between different clusters, etc.