CHAPTER 6

Image Processing

ILWIS for Windows 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 start up ILWIS and change to the subdirectory c:\ilwis21\data\usrguide\chap06, where the data files for this chapter are stored.

☞

- Double-click the ILWIS program icon in the ILWIS program group.
- Change the working drive and the working directory until you are in the directory c:\ilwis21\data\usrguide\chap06.

6.1 Visualization of single band images

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

To compare bands or to compare 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 coordinate system belonging to a defined map projection. In this case, the coordinates of each pixel in the window are displayed if one points to it.

The objective of these exercises is 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 DNs 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 DNs constituting 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 also usually a display depth of 8 bits. Hence, it can display images in 256 (gray) tones.

The following Landsat TM bands, of an area near Cochabamba (Bolivia), are used in this exercise:

TM band 1	: Tmb1	TM band 4	: Tmb4
TM band 3	: Tmb3	Landuse map	: Landuse

Display a satellite image

Display the Tmb1 (blue) image of the area near Cochabamba.

☞

- Double-click Show in the Operation-list to open the Show Map dialog box.
- Select the image Tmb1 by double-clicking it. The Display Options-Raster Map dialog box is opened.

> • In the Display Options-Raster Map dialog box, it is shown that the image Tmb1 has a domain image with digital numbers ranging from 25 to 162 and that it will be displayed in gray tones (representation: gray).
>
> • Click OK to accept the selected image. The image will be displayed in gray tones from black (lowest reflectance) to white (highest reflectance).

The image Tmb1 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 the map window by clicking the Smaller button ▧.
>
> • Enlarge the map window by clicking the Larger button ▧.
>
> • Maximize the map window by clicking the Maximize button ▴.
>
> • Close the map window by selecting File from the map window menu bar and click Exit, or double-click on the upper left Control menu button of the map window.

Zoom in / out on a displayed satellite image

When you open a map or image, ILWIS displays it entirely in a map window. The size of the map or image is automatically reduced so that it will fit on the screen. When you want to study details of the image, the zooming possibility can be used.

> ☞
>
> • Double-click Tmb1, accept the defaults and click OK.
>
> • Maximize the map window by clicking the Maximize button ▴. Note that this button looks different in Windows 95.
>
> • Click the Zoom in button 🔍 in the button bar 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 to the second corner of the area of interest. Releasing the mouse button will result in an enlarged display of the selected area.
>
> • Click (repeatedly) the Zoom out button 🔍 to zoom out.

> • Click the Entire map button 🖼 to show the entire map in the map window.

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.
> • 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 Tmb4 and **Zoom** in on a selected area.
> • Use the Left/Right arrows of the horizontal scroll bar and the Up/Down arrows of the vertical scroll bar to roam through the image.
> • A fast scroll can be achieved by dragging the small square boxes in the scroll bars Left/Right or Up/Down, or click in the scroll bar itself (not on the arrow button).
> • You can also scroll through an image, by clicking in the map window and moving the mouse pointer outside of the map window, while keeping the mouse button pressed.
> • Close the Map window.

Displaying multiple images

Often it is 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. Although it is possible to display in one map window a raster image together with point, segment or polygons maps, it is, however, 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 TM, of the area near Cochabamba, will be displayed using three different
display windows.

☞

- Double-click `Tmb1` in the Catalog and select OK in the Display
  Options-Raster Map dialog box. Move the window to a corner of
  the screen.

- Open the File menu in the Main window and choose Open. In the
  Show Map dialog box select `Tmb3` and click on OK. The Display
  Options-Raster Map dialog box is displayed. Accept defaults and
  click OK. Move the window to a position close to the `Tmb1` window.

- Click the right mouse button while the mouse pointer is over `Tmb4`
  in the Catalog. A context-sensitive menu appears to the right side of
  the mouse pointer. Select Open and accept the default display
  options, by clicking OK in the Display Options dialog box. Move
  the window to a position close to the `Tmb3` window.

- Drag-and-drop the polygon map `Landuse` in the map window,
  displaying `Tmb1`. Select the option Boundaries Only in the Display
  Options dialog box. Select the color Red, make sure that the check
  box Info is selected and click OK. Do the same for the other two map
  windows. When you click in any of the map windows you will see the
  land use type.

- Study the individual windows and try to explain the gray intensities
  in terms of spectral response of *water, urban areas, agriculture* and
  grass lands for the individual bands. Note down the relative spectral
  response as *low, medium* or *high*, for the different land cover classes
  and spectral bands in table 6.1.

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	Urban area	Agriculture	Grass land
1				
3				
4				

☞

- Close the map windows.

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 receive high digital numbers and vise versa. The spectral responses of objects can change over the wavelengths recorded, as you have seen in the former exercise when comparing the three Landsat TM bands with regard to e.g. 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 of a tone.

The digital numbers representing the tone intensity of an image 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.

☞

- From the File menu in the Main window, select Open Pixel Information. The Pixel Information window appears. Open the File menu in the pixel information window and choose Add Map. Select Tmb1 and click on OK.

- Make sure that the pixel information window can be seen on the screen. Select Tmb3 in the Catalog and drag it to the pixel information window, by holding the left mouse button down. Move the mouse pointer to the pixel information window, and release the left mouse button: The image is dropped in the pixel information window.

- Drag-and-drop Tmb4 in the pixel information window also.

- Click the right mouse button while the cursor is over Tmb4 in the Catalog. A context-sensitive menu appears to the right side of the mouse pointer. Select Open and click on OK. Make sure that the map window does not overlap with the pixel information window.

- In the pixel information window the DN's of all three selected bands will be displayed. Zoom in if necessary.

Roam through the image and note down some DN values for water, urban area, irrigated and scrub land in table 6.2. For ground truth, the land use map of Cochabamba (map Landuse) can be used.

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

Band	Water	Urban area	Agriculture	Shrubs	Grassland
1					
3					
4					

☞

- Close all map windows and the pixel information window after finishing the exercise.

Pixels and real world 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, which is discussed in more detail later, the relationship between row and column number and real world coordinates are established.

☞

- Click the map Tmb4 with the right mouse button and select Open from the context-sensitive menu. Click OK in the Display Options dialog box. Maximize the map window.

- Move the cursor to the lake area near the city of Cochabamba and note the Row/Col and XY, Lat/Long locations in the status line. Zoom in if necessary.

- Move to the mountain area in the NW corner of the image. Note the change in real world coordinates and the Row/Col numbers.

The last button defined in the button bar of the map window is a tool to conduct measurements.

- Click the Measure distance button ⬚ and move the mouse pointer to the northern edge of the lake. Press the left mouse button and move the mouse pointer to the southern edge. When the mouse button is released, the length and direction are given.

Through linkage of the image with real world coordinates, e.g. distance, measurements can be performed.

> ☞
>
> - Close the map windows and the pixel information 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 band image 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 of a tone.

– 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, scroll through and zoom in/out on the images and retrieve the DNs of the displayed images.

– It is possible to display in one map window a raster image together with point, segment or polygons maps.

– 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 relationship between row and column number and real world coordinates are established.

6.2 Image enhancement

Image enhancement deals with the procedures of making a raw image more 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 on the human eye.

Image enhancement techniques can be classified in many ways. Contrast enhancement, also called global 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 takes into consideration local conditions only 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, have 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. 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* or *path radiance*. 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 a contrast enhancement are: the *linear stretching technique* and *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 256 gray levels (0 - 255). This corresponds with the resolution of most 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 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. However, linear contrast stretch, putting (min, max) at (0,255) in most cases still produces a rather dull image. Even though all gray shades of the display are utilized, the bulk of the pixels are displayed in mid gray. This is caused by the more or less normal distribution, within 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 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.
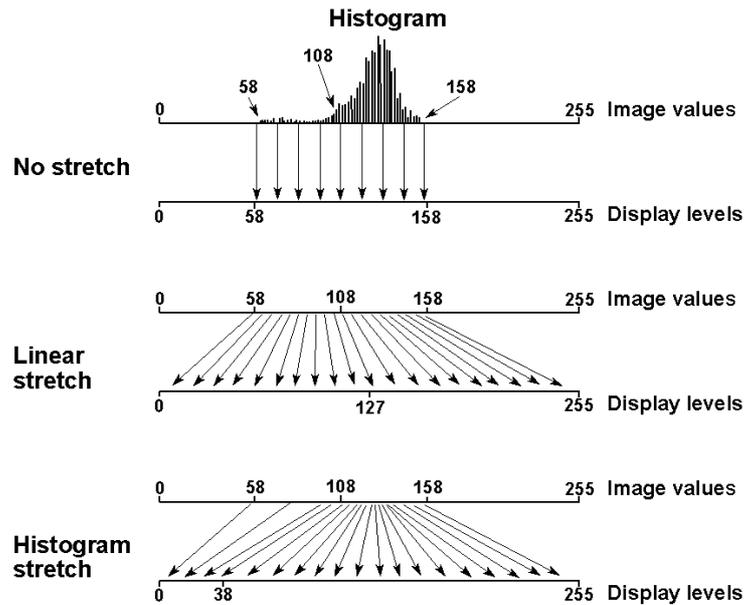
Figure 6.1: Principle of contrast enhancement

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 109 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.

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 only a certain cover type, for example water. The data values for this feature are in the range of 5 to 18, and in order to be able to discriminate as much as possible, it is wise to use all available gray levels for this feature only. In this way detailed differences within the feature of interest appear, where as the remaining features are assigned to a single gray tone.

The material for this exercise consists of a Landsat TM band 1 of the area near Cochabamba valley, Bolivia: Tmb1.

Calculation of a histogram

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

☞

- Double-click the Histogram icon in the Operation-list. The Calculate Histogram dialog box is opened.

- Select the raster map Tmb1. Select the option Show to show the histogram immediately after its calculation. Click OK.

- The histogram is shown in a table window. A numerical display of the histogram is shown in a number of columns:
  Npix (= number of pixels with a certain pixel value),
  Npixpct (= 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),
  Npixcum (=cumulative frequency distribution),
  Npcumpct (=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 the table window can be maximized.

- Write down in table 6.3 the DN values that belong to the corresponding cumulative percentages.

Table 6.3: DN values at specific cumulative percentages. Fill in the table from the histogram of Tmb1

Cumulative Percentage	0	0.5	1	2	5	10
DN value						
Cumulative Percentage	90	95	98	99	99.5	100
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.

☞

- In order to have a graphical display of the histogram, open the Options menu in the table window, and select Show Graph. The Graph dialog box is opened. Select Image value for the X-Axis, and Npixpct for the Y-Axis. The Edit Graph dialog box is opened.

- Accept the default display options (Type: Line and Color: Red) and a graphical display of the histogram is shown.

- Use the calculated histogram to answer the questions stated below.

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 predominant value (most frequently occurring) of the image?

☞

- In order to get an overview of the most important image statistics, make the table window which show the histogram values active and from the Edit menu choose Properties. Check if the DN values in table 6.3 are correct. Check also your answer to question 4.

- Close the dialog box, the graph window and the table window.

Linear stretching

After a histogram has been calculated for a certain image, the image can be stretched. A linear stretch is used here. The pixel values of the interval between the 1 and 99% boundaries, are used as input interval to expand the image values.

☞
- In the Main window, open the Operations menu and select Image Processing, Stretch. The Stretch dialog box is opened.
- Select Tmb1 as Raster Map, Linear stretching as stretching method with a Percentage equal to 1. Type Tmb1str as Output Raster Map name and select the domain Image.
- Select the check box Show and click OK. The raster map Tmb1str is calculated. The Display Options dialog box is opened. Click OK.

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 comparing the image statistics using their histograms.

☞
- Display the original (Tmb1) and the stretched image (Tmb1str), 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 dialog box of the original image!
- Display the histograms of Tmb1 and Tmb1s graphically and study them.
- Write down in table 6.4 the DN values belonging to the stretched image using Edit, Histogram properties. Indicate the major changes compared to the data written down in table 6.4.
- Close the map and table windows.

Table 6.4: DN values at specific cumulative percentages. Fill in the table from the histogram of Ttmb1str

Cumulative Percentage	0	0.5	1	2	5	10
DN value						
Cumulative Percentage	90	95	98	99	99.5	100
DN value						

Determine different linear stretch functions

In figure 6.2, four different linear contrast stretch functions are given using different lower and upper percentages to be ignored. These four functions were derived using:

1) No stretching

2) Stretching between minimum and maximum values

3) 1 and 99% as lower and upper percentages

4) 5 and 95% as lower and upper percentages

❢   Indicate for each line in figure 6.2 which linear stretch function is shown.
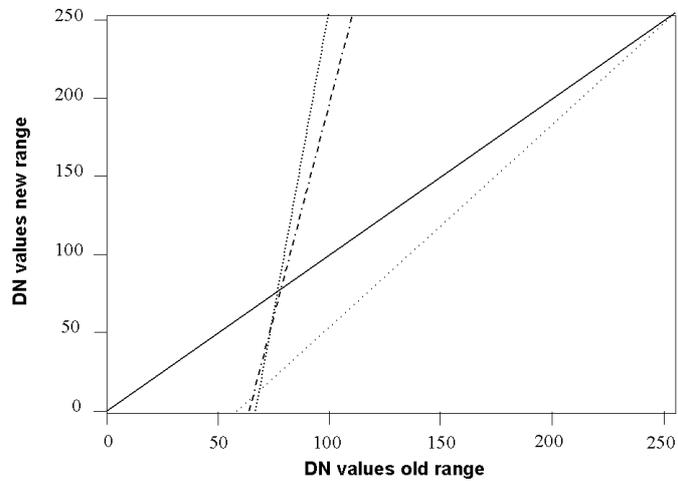
Figure 6.2: Different linear stretch functions

Histogram equalization

The histogram equalization stretching technique takes into account the frequency of the DNs. 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 DNs themselves can be given instead.

☞

- Open the Stretch dialog box by selecting Stretch in the Operations list.
- Select Tmb1 as Input Raster Map and Histogram Equalization as stretching method. Use 5 intervals, each interval should receive 20% of the data. To stretch from, use default percentage. Type Tmb1eq for the Output Raster Map. Select the check box Show to display the stretched map after the stretching operation.
- Display the stretched image, using a gray representation.
- Display the histogram of the previously created output map as a numerical listing and check the frequency per interval.

6.2.2 Spatial enhancement

The objective is to understand the concept of spatial enhancement using filters and to be able to apply different types of filters.

Spatial enhancement procedures result in modification of an image pixel value, based on the pixel values in its immediate vicinity (local enhancement). Filters are commonly used to (1) correct and restore images affected by system malfunctioning, (2) enhance the images for visual interpretation and (3) to extract features. Like all image enhancement procedures, the objective is to create new images from the original image data, in order to increase the amount of information that can be visually interpreted.

Spatial frequency filters, also often simply called spatial filters, may emphasize or suppress image data of various spatial frequencies. Spatial frequency refers to the roughness of the variations in DN values occurring in an image. In high spatial frequency areas, the DN values may change abruptly over a relatively small number of pixels (e.g. across roads, field boundaries, shore lines). Smooth image areas are characterized by a low spatial frequency, where DN values only change gradually over a large number of pixels (e.g. large homogeneous agricultural fields, water bodies). Low pass filters are designed to emphasize low frequency features and to suppress the high frequency component of an image. High pass filters do just the reverse.

Low pass filters. Applying a low pass filter has the effect of filtering out the high and medium frequencies and the result is an image which has a smooth appearance. Hence, this procedure is sometimes called image smoothing and the low pass filter is called a smoothing filter. It is easy to smooth an image. The basic problem is to do this without losing interesting features. For this reason much emphasis in smoothing is on edge-preserving smoothing.

High pass filters. Sometimes abrupt changes from an area of uniform DN's to an area with other DN's can be observed. This is represented by a steep gradient in DN values. Boundaries of this kind are known as edges. They occupy only a small area and are thus high-frequency features. High pass filters are designed to emphasize high-frequencies and to suppress low-frequencies. Applying a high pass filter has the effect of enhancing edges. Hence, the high pass filter is also called an edge enhancement filter.

Two classes of high-pass filters can be distinguished: *gradient (or directional) filters* and *Laplacian (or non-directional) filters*. Gradient filters are directional filters and are used to enhance specific linear trends. They are designed in such a way that edges running in a certain direction (e.g. horizontal, vertical or diagonal) are enhanced. In their simplest form, they look at the difference between the DN of a pixel to its neighbor and they can be seen as the result of taking the first difference (hence the term Gradient). Laplacian filters are non-directional filters because they enhance linear features having almost any direction in an image. They do not look at the gradient itself, but at the changes in gradient. In their simplest form, they can be seen as the result of taking the second difference.

A filter usually consists of a 3 x 3 array (sometimes called kernel) of coefficients or weighting factors. It is also possible to use a 5 x 5, a 7 x 7 or even a larger odd numbered array. The filter can be considered as a window that moves across an image and considers all DN values falling within the window. Each pixel value is multiplied by the corresponding coefficient in the filter. The 9 resulting values are summed and the resulting value replaces the original value of the central pixel. This operation is called *convolution*. Figure 6.3 illustrates the convolution of an image using a 3 x 3 kernel.

The material used for this exercise consist of band 4 of a Landsat TM Cochabamba, Bolivia: `Tmb4`.

Low pass filters

To remove the high frequency components in an image, a standard low pass smoothing filter is used. In this activity a 3x3 average filter is going to be used.

☞

- Open the Filtering dialog box by selecting the Filter operation in the Operations-list.

- Select for the Input Raster Map: `Tmb4`, use a linear filter type and the standard low pass filter `Avg3*3`. Enter `Tmavg` as Output Raster Map name and select domain Image. Select the check box Show. Click OK.

- Display the unfiltered and the filtered image next to each other in two map windows and make a visual comparison.

Create and apply a user defined low pass filter

Besides having pre-defined filters, it is possible to define, create and apply your own designed filter. A 3x3 low pass filter of the weighted mean type will be created and applied to the image.

☞

- Open the Create Filter dialog box by selecting New Filter in the Operations-list.

- Enter the name `Weigmean` for the filter to be created and clear the Real Matrix check box.

- Confirm the input parameters. The Edit Filter dialog box will be displayed.

- An empty filter will be displayed and the values for the filter can be entered as shown in table 6.5.

Applied filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

OR

1/9 x

1	1	1
1	1	1
1	1	1

Input image

Output image

25	27	32	20	21
23	27	19	25	22
27	26	24	26	22
23	25	26	30	27
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

= 1/9 (1x 25 + 1 x 27 + 1 x 32 + 1 x 28 +
1 x 27 + 1 x 19 + 1 x 27 + 1 x26 + 1 x 24)= 26

	26		

25	27	32	20	21
23	27	19	25	22
27	26	24	26	22
23	25	26	30	27
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

= 1/9 (1x 27 + 1 x 32 + 1 x 20 + 1 x 27 +
1 x 19 + 1 x 25 + 1 x 26 + 1 x 24 + 1 x 26) = 25

	26	25	

25	27	32	20	21
23	27	19	25	22
27	26	24	26	22
23	25	26	30	27
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

= 1/9 (1x 32 + 1 x 20 + 1 x 21 + 1 x 19 +
1 x 25 + 1 x 22 + 1 x 24 + 1 x 26 + 1 x 22)= 23

	26	25	23

25	27	32	20	21
23	27	19	25	22
27	26	24	26	22
23	25	26	30	27
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

= 1/9 (1x 23 + 1 x 27 + 1 x 19 + 1 x 27 +
1 x 26 + 1 x 24 + 1 x 23 + 1 x 25 + 1 x 26)= 24

	26	25	23
	24		

Final filtered image

	26	25	23
	24	25	27
	24	27	27

Figure 6.3: Convolution, using a 3*3 filter with all coefficients equal to 1/9

Table 6.5: Values for the low pass weighted mean filter

1	2	1
2	4	2
1	2	1

☞

- Accept the default Gain, and confirm the input parameters and close, if necessary, the open dialog boxes.

- What is the function of the gain value?

- Filter the image Tmb4 using the previously defined smoothing filter. Select the created filter Weigmean from the filter list. The domain for the output map is Image. Select the check box Show, and type Weigmean as Output map name.

- Compare this filtered image with the image filtered using the standard smoothing filter.

High pass filters

Below an example is given (figure 6.4) of a 3*3 laplace filter. The product (sum) of the applied filter is zero. A gain factor to correct the effect of the filter coefficients is therefore not necessary. When these type of filters are applied, the resulting values can be both negative and positive. It is for this reason that ILWIS treats the output images, storing the results of these computations in a value domain.

To enhance the high frequency components in an image, a standard high pass filter is used. The applied filter is an 3x3 edge enhancement filter with a central value of 16, eight surrounding values of -1 and a gain of 0.125.

☞

- Open the Filtering dialog box by selecting the Filter operation in the Operations-list.

- Select for the Input Raster Map Tmb4 and select the linear filter Edgesenh. Enter the name Edge for the Output Raster Map name. The domain should be Value (negative values are possible!). Accept the default value range and precision. Select the check box Show. Click OK. The Display Options dialog box is opened.

- Use a Gray representation and accept the other defaults for display.

- Display the unfiltered and the filtered image next to each other in two map windows and make a visual comparison.

Applied filter

0	-1	0
-1	4	-1
0	-1	0

Input image　　　　　　　　　　　　　　　　　　　　**Output image**

25	117	120	121	121
23	27	119	119	122
27	26	24	120	122
23	25	26	30	121
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

= (-1 x 117 + -1 x 23 + 4 x 27 + -1 x 119 + -1 x 26) = -177

-177			

25	117	120	121	121
23	27	119	119	122
27	26	24	120	122
23	25	26	30	121
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

=(-1 x 120 + -1 x 27 + 4 x 119 + -1 x 119 + -1 x 24) = 186

-177	186		

25	117	120	121	121
23	27	119	119	122
27	26	24	120	122
23	25	26	30	121
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

=(-1x 121 +- 1 x 119 + 4 x 119 + -1 x 122 + -1 x 120)= -6

-177	186	-6	

25	117	120	121	121
23	27	119	119	122
27	26	24	120	122
23	25	26	30	121
19	22	28	32	28

OUTPUT = MULTIPLY AND SUM

=(-1x 27 + -1 x 27 + 4 x 26 + -1 x 24+ -1 x 25) = 1

-177	186	-6	
1			

Final filtered image

-177	186	-6	
1	-195	185	
3	-3	-179	

Figure 6.4: Convolution, using a 3*3 Laplacian filter

Create and apply a user defined Laplace filter

A 3x3 high pass filter of the Laplacian type (the Laplace Plus filter) will be created and applied to the image.

☞

- Open the Create Filter dialog box by selecting New Filter in the Operations-list.
- Enter the Filter Name Lapplus and the size 3x3. Deselect the option Real Matrix.
- Confirm the input parameters and click OK. The Edit Filter dialog box will be displayed.
- An empty filter will be displayed and the values for the filter can be entered as given in the table below.

Table 6.6: Values for the Laplace Plus filter

0	-1	0
-1	5	-1
0	-1	0

☞

- Accept the default Gain and close, if necessary, the open dialog boxes.

What is the function of the gain value here?

☞

- Filter the image Tmb4 using the previously defined high pass filter. Use the default value range and step size. The domain is value.
- Filter the image Tmb4 using a standard Laplace filter. Use the domain Value and the default value range and precision.
- Compare this filtered image, displayed using a gray representation for the output map, with the image filtered using a standard Laplace high pass filter.

Directional filters

A directional filter is used to enhance specific linear trends. They are designed in such a way that edges running in a certain direction are enhanced. To enhance lineaments running North-South, an X-gradient filter can be used.

> ☞
>
> - Create a directional filter as given in the table below using the same procedures as given in the former exercise.

Table 6.7: Values for the directional filter

1	0	-1
1	0	-1
1	0	-1

> ☞
>
> - Filter the image `Tmb4` using the previously created directional filter.
> - Create the same type of directional filter, but now a filter to enhance the lineaments in a East-West direction and filter also the image `Tmb4`.
> - Display the original image `Tmb4` and the two filtered images.

Explain the difference between these images?

Summary: Image enhancement

– **Contrast enhancement**, also called global 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*. To enhance specific data ranges showing certain land cover types the piece-wise linear contrast stretch can be applied.

– 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.

– **Spatial Enhancement** procedures result in modification of an image pixel value, based on the pixel values in its immediate vicinity (local enhancement).

– Low pass filters are designed to emphasize low frequency features and to suppress the high frequency component of an image. High pass filters do just the reverse.

- Two classes of high-pass filters can be distinguished: *gradient (or directional) filters* and *Laplacian (or non-directional) filters*. Gradient filters are directional filters and are used to enhance specific linear trends.

- Laplacian filters are non-directional filters because they enhance linear features having almost any direction in an image.

- A filter usually consists of a 3 x 3 array (sometimes called kernel) of coefficients or weighting factors.

- Each pixel value is multiplied by the corresponding coefficient in the filter. The 9 resulting values are summed and the resulting value replaces the original value of the central pixel. This operation is called convolution.

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. For SPOT multi-spectral imagery, the bands 1, 2 and 3 are displayed respectively in blue, green and red. A combination used very often for TM 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*.

Bands of different images (from different imaging systems or different dates), or layers created by band rationing or principal component analysis, can also be combined using the color composite technique. An example could be the multi-temporal combination of vegetation indices for different dates, or the combination of 2 SPOT XS bands with a SPOT PAN band, (giving in one color image the spectral information of the XS bands combined with the higher spatial resolution of the panchromatic band).

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. In figure 6.5 the color cube is represented and the primary additive (red, green and blue) and subtractive colors (yellow, magenta, cyan) are given. 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 2.0 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, ranges 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* domain, as the relationship with the original data has disappeared.
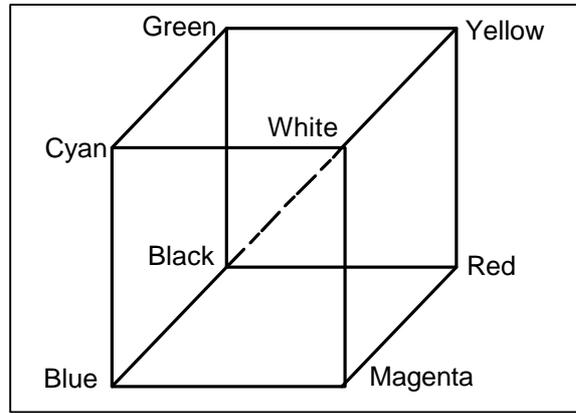
Figure 6.5: Color cube

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, another the values for Green and the third for Blue (figure 6.6). The number of colors used can be defined. A maximum of about 200 colors are available because 56 colors are reserved for Windows and ILWIS itself.
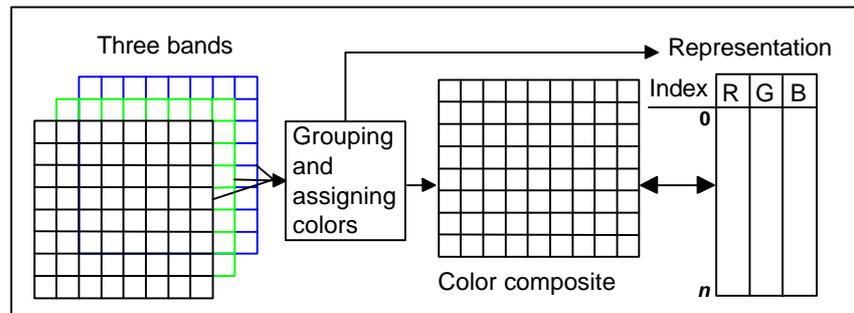
Figure 6.6: The link between a color composite, the source images and the representation. *n* is the number of colors (maximum 200)

Using ILWIS, color composites can be created using different methods: *24 bit*, *standard* and *dynamic*, and for the first two methods different stretching functions can be applied. Below, the dynamic and standard options are further elaborated.

By default ILWIS calculates dynamic color composites by combining 3 images using the *Heckbert* algorithm. This algorithm is based on the differences in the values of pixels in the three input images. All pixels of the three images are plotted in the color cube. The cube is then divided into two parts of equal volume. The

amount of pixels in each of the two parts is counted. After that, the box that contains the largest amount of pixels is again split into two. Again, the number of pixels is determined for all volumes, and the one containing the largest number of pixels is split. This process continues until the specified number of boxes (colors) is reached. The result of this procedure is that dense clusters of pixels (areas within the cube with relatively large number of pixels) are assigned more colors than the less dense clusters. This enhances the contrast within large, relatively homogenous areas. Finally, colors are determined according to the center of each box within the color cube. If one is interested in the image as a whole, i.e. a composite in which the input range is spread equally over an output range, the Dynamic option yields the best composite. On the other hand, if specific intervals of the input bands are of importance, the Standard option which applies a linear stretching function, can be used.

To create a standard color composite, all pixels of the three images are plotted in the color cube, using the chosen color assignment for the red, green and blue bands. The scaling of the axes of the cube is determined using the range of values of the input bands, which are linearly divided into six classes having equal length. Since this is done using three bands, the number of possible combinations is $6\times6\times6 = 216$, which is the maximum number of smaller cubes. Each of these cubes is assigned a color according to its position in the color cube. All pixel values that fall within one of these cubes will be given the color assigned to the center of that cube. The color assignment is stored in the color composite Representation.

False and pseudo natural color composites

In the exercises below, both a false color composite and a pseudo natural color composite, will be created using Landsat TM bands. For the creation of the false color composite, three TM bands have to be selected. Enter the corresponding TM bands for the spectral ranges indicated in table 6.8.

Table 6.8: Spectral ranges for selected TM bands

Spectral range	TM band number
Near infrared	
Visible Red	
Visible Green	

To create a Pseudo Natural Color composite using three TM bands, which bands have to be selected ? Note down the color assignment in table 6.9.

Table 6.9: Color assignments

Color	TM band
Red	
Green	
Blue	

In this exercise, a color composite image is created using TM band 4, 3 and 2. The red is assigned to the near infra red band, green to the red and blue to the green visible bands. The created color composite image should give a better visual impression of the imaged surface compared to the use of a single band image.

☞

- Double-click the ColorComp operation in the Operations-list.

- Accept the defaults for the Method and Colors.

- Select the images Tmb4 for Red band, Tmb3 for Green band and Tmb2 for Blue band.

- Click the check box Show.

- Type Tmcc432 in the text box Output Raster Map and click OK.

- Click OK in the Display - Options dialog box.

- Repeat the above described procedures and create a pseudo natural color composite, according to the color assignment given in table 6.9 and call the output image Pseudo.

- Display both color composites. Add the polygon map Landuse to the map windows (display only the boundaries in the black).

Explain the color differences between the two displayed images?

☞

- Close all open map windows when you have finished the exercises.

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.