
GEOG0027 Environmental Remote Sensing Documentation

Release 1.0.2019

P Lewis

Jan 04, 2023

CONTENTS

1	Course Tutors	3
2	Contents	7
2.1	Introduction	7
2.1.1	Course notes	7
2.1.2	Reading List	9
2.1.3	Use of these notes	9
2.2	Image Display and Enhancement using ENVI 5.5	10
2.2.1	Aims	12
2.2.2	Data	12
2.2.3	Introduction	13
2.2.4	Exploration	18
2.2.5	Spectral Features	18
2.2.6	Spectral Profiles	23
2.2.7	Enhancement	25
2.2.8	Comparing Images	27
2.2.9	More Exploring	31
2.2.10	Summary	32
2.3	Download data	33
2.3.1	Downloading and visualisation tools	34
2.3.2	Select Landsat Scenes	34
2.3.3	Download ordered data	35
2.3.4	Unpacking Downloaded data	35
2.3.5	Fixing Landsat 7 SLC issues	35
2.3.6	Download Scenes	39
2.4	Spatial filtering using ENVI 5.5	39
2.4.1	Aims	39
2.4.2	Data	40
2.4.3	Introduction	40
2.4.4	Convolution filtering	41
2.4.5	Convolution	46
2.4.6	4. Extras (optional)	51
2.4.7	Summary	52
2.5	Classification using ENVI 5.5	53
2.5.1	Aims	53
2.5.2	Data	53
2.5.3	Introduction	54
2.6	Classification using ENVI 5.5	57
2.6.1	Aims	57
2.6.2	Introduction	57

2.6.3	Examination of the data	57
2.6.4	Defining spectral classes	58
2.6.5	Image Classification	61
2.6.6	Accuracy Assessment	66
2.6.7	Further Work	68
2.6.8	Summary	68
2.7	Coursework	68
2.7.1	Coursenotes	68
3	Search	69



COURSE TUTORS

Prof. P. Lewis
Prof. M. Disney
Dr Qingling Wu
Department of Geography
University College London



CONTENTS

2.1 Introduction

2.1.1 Course notes

The material in this document is the main lecturing and course material for GEOG0027.

Educational Aims and Objectives of the Course

To enable the students to:

- Understand the nature of remote sensing data and how they are acquired
- Understand different types of remote sensing instruments and their missions
- Understand basic image representation and processing
- Understand how Earth Observation data can be combined with other sources of data and data techniques (e.g. GIS)
- Understand how EO data can be used in environmental science (particularly via classification and monitoring)
- Develop practical skills in these areas, which may be useful in planning of dissertations
- Develop links with the second year course on Geographic Information Systems Science and with other courses as appropriate (e.g. hydrology, environmental systems)
- Lay the foundations for the third year course on Earth Observation

Course workload and assessment

Expected Course Load

Component	Hours
Lectures	8
Private Reading	80
Supervised Laboratory Work (Computing)	24
Independent Laboratory Work (Computing)	20
Required Written Work	10
TOTAL	142

Usual range 100-150 for 1/2 course unit

Timetable 2018-19

Week/Day	Thursday 11:00-12:00	Thursday 12:00-13:00	Friday 16:00-17:00
Week 1	.	.	11/1/19 LECTURE 1
.	.	.	Introduction to course
Week 2	17/1/19	17/1/19	18/1/19 LECTURE 2
.	COMPUTING 1: <i>Image Display</i>	COMPUTING 1: <i>Image Display</i>	Image Display and Enhancement
Week 3	24/1/19	24/1/19	25/1/19 LECTURE 3
.	<i>DOWNLOAD</i>	COMPUTING 1: <i>Image Display</i>	Spatial Information
Week 4	31/1/19	31/1/19	01/2/19 LECTURE 4
.	COMPUTING 2: <i>Spatial Filtering</i>	COMPUTING 2: <i>Spatial Filtering</i>	Image Classification
Week 5	04/2/16	04/2/16	05/2/19 LECTURE 5
.	COMPUTING 3a: <i>Classification Intro</i>	COMPUTING 3: <i>Classification</i>	Spectral Information
Week 6	READING WEEK	READING WEEK	READING WEEK
.	.	.	.
Week 7	21/2/19	21/2/19	22/2/19 LECTURE 6
.	<i>Classification</i>	COMPUTING 3: <i>Classification</i>	Environmental Modelling I
Week 8	28/2/19	28/2/19	1/3/19 LECTURE 6
.	COMPUTING 4: <i>Project</i>	COMPUTING 4: <i>Project</i>	Environmental Modelling II
Week 9	07/3/19	07/3/19	08/3/19
.	COMPUTING 4: <i>Project</i>	COMPUTING 4: <i>Project</i>	COMPUTING 4: <i>Project</i>
Week 10	14/3/19	14/3/19	15/3/19 <i>Project</i>
.	COMPUTING 4: <i>Project</i>	COMPUTING 4: <i>Project</i>	Discussion
Week 11	21/3/19	21/3/19	No lecture
.	COMPUTING 4: <i>Project</i>	COMPUTING 4: <i>Project</i>	.

Lectures in Pearson G07_

Computing in Pearson Building, UNIX Computer lab, Pearson 110a_

Assessment

100% Assessed Practical (3500 words) - submission date standard 2nd year submission date i.e. Fri 22th March 2019 (12 noon).

N.B.

- Penalties for late submission and over length WILL be applied
- Different arrangements for JYA/Socrates (make sure you inform the lecturers if this affects you)

2.1.2 Reading List

- Jensen, John R. (2006) Remote Sensing of the Environment: an Earth Resources Perspective, Hall and Prentice, New Jersey, 2nd ed.
- Jensen, John R. (1995, 2004) Introductory Digital Image Processing: A Remote Sensing Perspective (Prentice Hall Series in Geographic Information Science)
- Jones, H. G and Vaughan, R. A. (2010) Remote Sensing of Vegetation, OUP, Oxford.
- Lillesand, T., Kiefer, R. and Chipman, J. (2004) Remote Sensing and Image Interpretation. John

2.1.3 Use of these notes

Use online notes

The simplest way to use these notes is to [read the documents online](#).

Use github notes

The main store for these notes is on [github](#) and you should always be able to access (or download) the notes from there.

Use binder

Click on the binder link

This lets you *run* the notebooks, rather than just viewing them.

Use Python locally

To use the notes as a notebook (assuming you have `git` [<http://git-scm.com>](http://git-scm.com) and `python` on your computer):

1. Copy all of the notes to your local computer (**if for the first time**)

```
mkdir -p ~/DATA/working
cd ~/DATA/working

git clone https://github.com/profLewis/geog2021.git

cd geog2021
```

2. Copy all of the notes to your local computer (**if for an update**)

```
cd ~/DATA/working/geog2021

git reset --hard HEAD

git pull
```

3. Run the notebook

```
ipython notebook SpatialFiltering.ipynb
```



2.2 Image Display and Enhancement using ENVI 5.5

Prof. P. Lewis

Prof. M. Disney

Dr H. Heorton

Department of Geography

University College London



2.2.1 Aims

After completing this practical, you should have a reasonable idea of a range of basic functions within an image processing software, *envi*. You should also have started to make some sense of how the theory you are learning in the lectures can be used to manipulate image data.

2.2.2 Data

The datasets you need for this practical should be available to you on moodle (<https://moodle.ucl.ac.uk/mod/folder/view.php?id=2700004>):

- ETM-190600
- ETM-190600.HDR
- TM-280589
- TM-280589.HDR

and you will need to download them into your personal N: Drive, e.g. N:\GEOG0027\data.

The data you will be using are:

- six wavebands of a **Landsat TM** image over Greater London, imaged on May 28th 1989. The data were obtained from the **GLCF** which maintains a large database of (freely available) Landsat and other imagery. The data are at an original pixel spacing of 28.5 m, but have been resampled to a 25 m grid here. The data are in a Transverse Mercator projection with OSGB 1936 datum (Great Britain Ordinance Survey).
- six wavebands (nominally the same wavelengths) of a Landsat ETM image with 25 m spatial resolution, covering the same spatial extent. These data were obtained on June 19th 2000. The data were obtained from Landmap which contains a database available to Universities and other users through an Athens login (done via the institution you are at).

The wavebands are:

1	2	3	4	5	6
blue	green	red	nir	swir 1	swir 2
450-520 nm	520-600 nm	630-690 nm	760-900 nm	1550-1750 nm	2080-2350 nm

The extent of the imagery is (Lat/Lon):

51°43'9.23" North, 0°36'18.37" West

to

51°16'29.32" North, 0°27'24.60" East

2.2.3 Introduction

***In this section, we will learn how to do a one time setup of the ENVI preferences and to load a false colour composite image*.**


The purpose of this practical is for you to gain experience in image display and enhancement of EO data. The session will be normally run as two hour supervised practical. You may not complete all tasks in detail in that time, so once you get the hang of how to use the tool, move on to the next section and return later to think more about the remote sensing.

For this term we are anticipating interruptions and absences though we are aiming to do all teaching face-to-face. We have set up the classes so that both options are available all the time by using UCL Desktop. The face-to-face teaching is in the departmental UNIX lab, and from these machines we access both MS Teams and UCL Desktop. ENVI will run straight away from UCL Desktop when you've navigated to it. There are several versions - use ENVI 5.5.3 (64 bit). **DO NOT USE ENVI Classic!**


There are alternate ways to work for those who wish: 1. The UNIX workstations also have the ENVI software installed, and you may find them better to use for image display. However when considering accessing saved work from multiple locations UCL Desktop is the easiest option for this course. 2. For personal machines the ENVI software is available from the [UCL Software Database](#), with installation instructions [for mac](#), [for windows](#). Note that running ENVI on your own machine will require you to first run the license server.

These notes are not intended to be a comprehensive introduction to ENVI functionality. They serve as an introduction to remote sensing and image processing, presuming no prior knowledge. Further information on envi tools and functions can be found in the [on-line help pages](#).

Starting ENVI on UCL Desktop reveals an empty display.

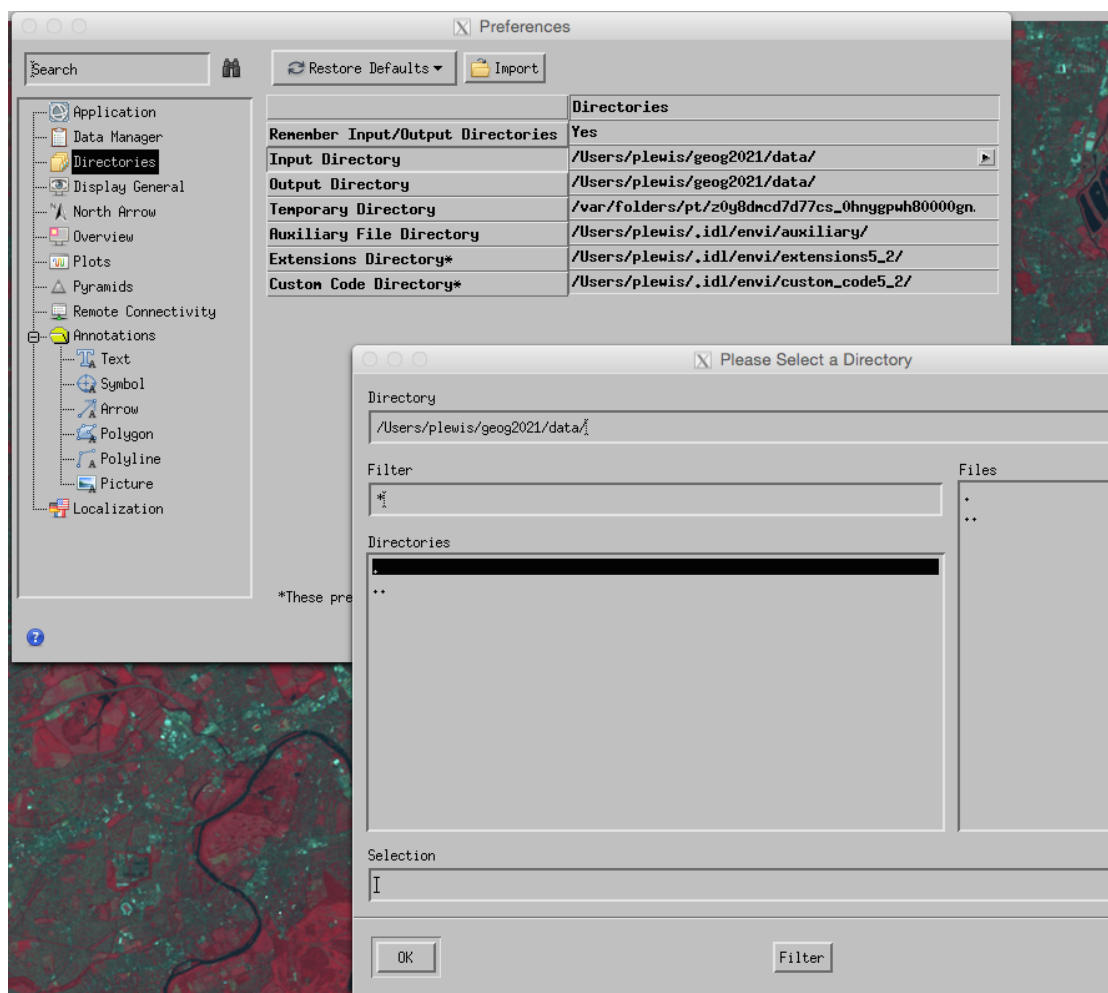


For personal machines or on a UNIX machine you get a similar display.

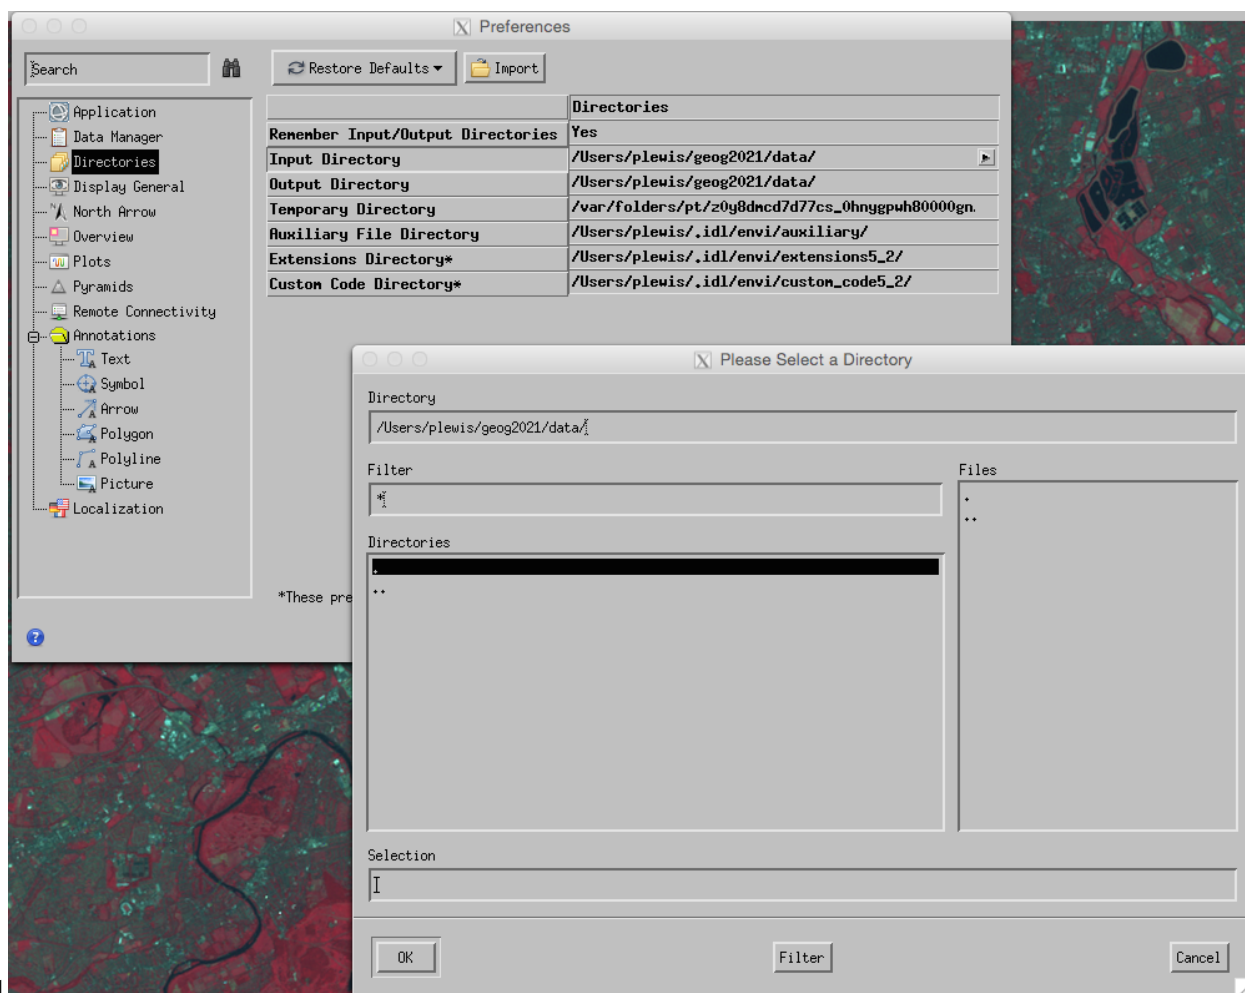


The first time you use the tool, you may find it convenient to update the default data location.

You can access this via the **File->Preferences** menu. Then select **Directories** from the left hand column. Select **Input Directory**, select the little arrow, and find the directory we made and put the data in earlier.



On UCL Desktop: this window:



for UNIX/Personal

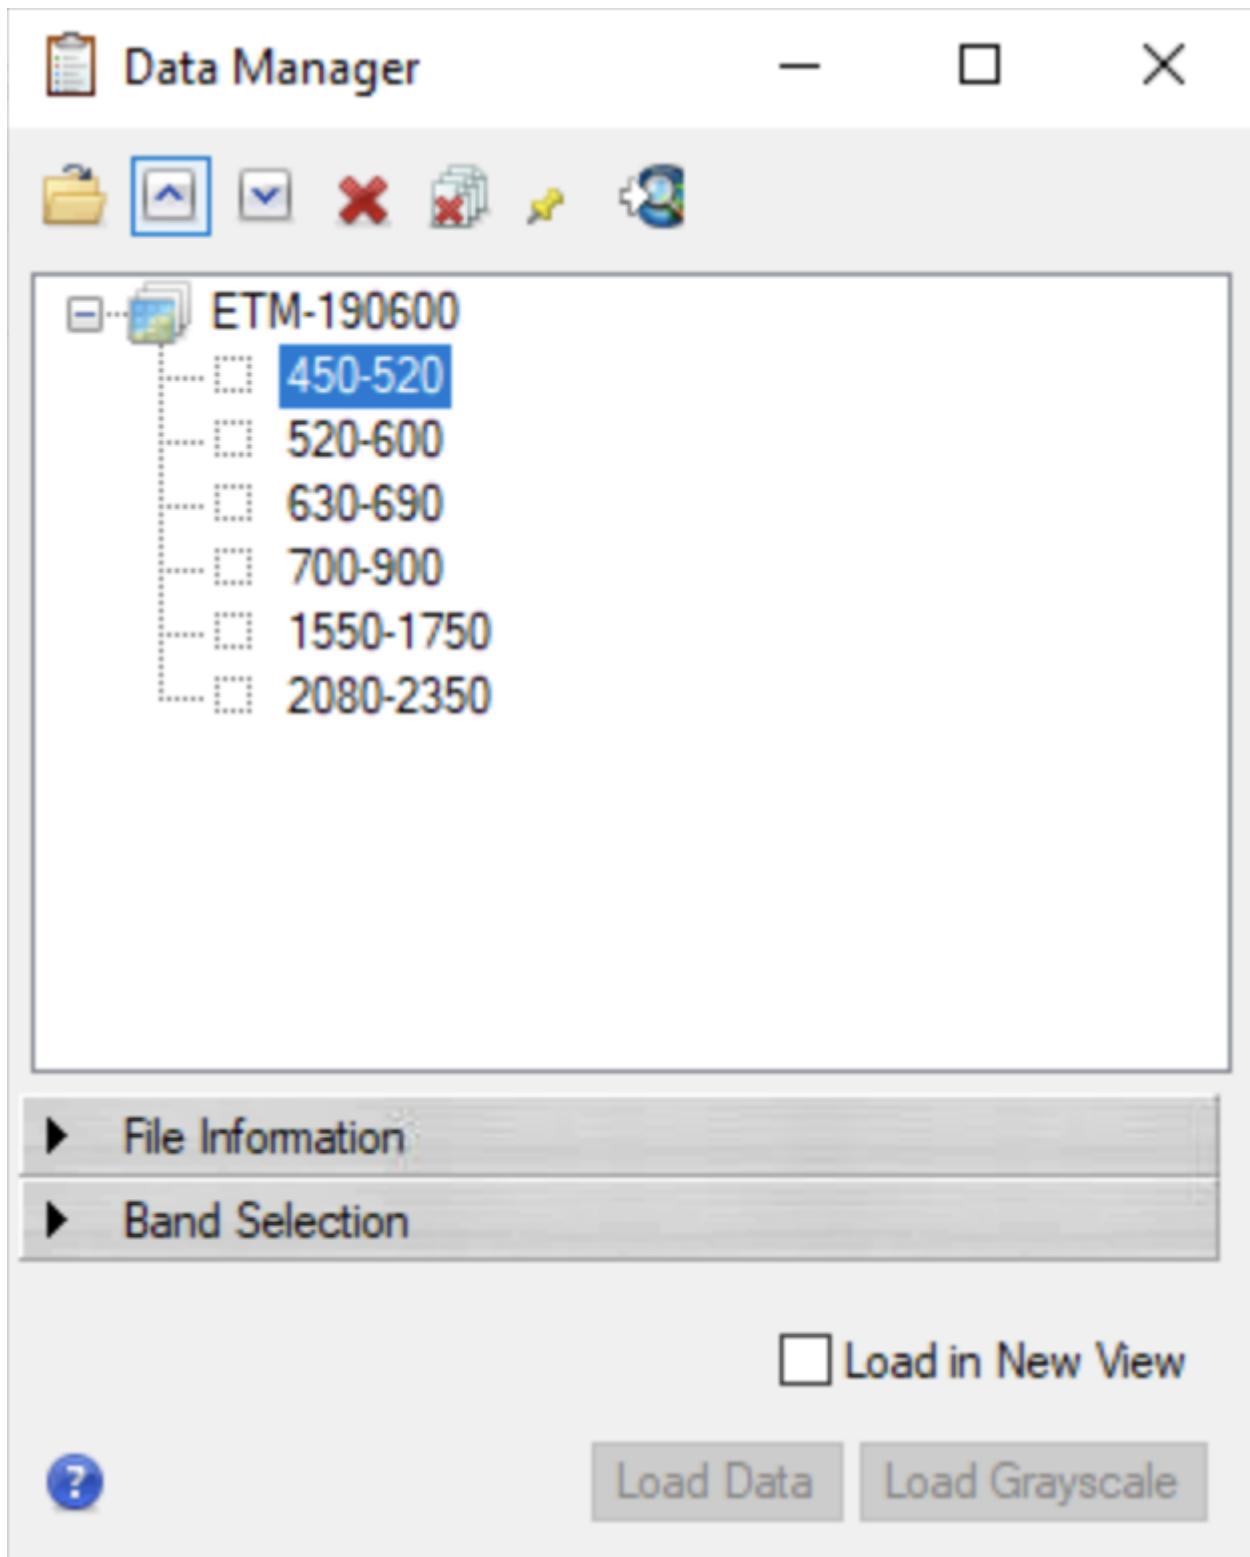
Typically, we will want to change the Input Directory and Output Directory to where you are currently working, e.g. N:\GEOG0027\data on UCL Desktop Anywhere (where we save the data earlier, when navigating this will be under This PC), or /Users/xxxx/GEOG0027/data on a Mac, or /home/ucfxxxx/DATA/GEOG0027/data or similar on the UCL Unix system.

Make a note of where your image data are, so that you can find them the next time!

For versions below 5.5.3, you have to quit and restart ENVI for these changes to take effect.

You should now be able to open an image dataset: select File->Data Manager from the menu.

Click on the file tab, and choose an image file to open e.g. ETM-190600. Do not select the HDR files as these are header files. You may need to first use the open icon (brown folder, top left) to find the file.



From which you can select one or more 'bands' to open and assign to the red, green and blue colour guns.

2.2.4 Exploration

***In this section, we use various tools for exploring the image*.**

Once you have loaded an image, explore how to pan and zoom around the image

Whilst doing this, examine some of the features you see in the image:

- The crosshairs can be used (and/or cursor value, looks like a “dropped pin”) to explore the geographical coordinates and the digital number (DN) in the three wavebands you are examining.
 - How do the DNs compare for lighter/darker regions of the image?
 - how could you check that the geographic coordinates seem sensible?
 - try to relate the visual cues you see in the data to the DN values being shown
 - use these ideas to interpret ‘colour’ as displayed in the image e.g. what does the ‘red’ colour *mean* and why is it so?

Find out how to Change RGB bands . . . for the image you are viewing and use this to display e.g. a real colour composite.


- Try some different band combinations and make sure you can interpret the colour you see
- What happens if you put the same band onto R G and B? why?

Make sure you make notes of what you have done and how you have done it so you can recall this later.

2.2.5 Spectral Features

***In this section, we will learn how to consider spectral features and perform contrast enhancement*.**

Clearly, we can ‘recognise’ much information in the image from spatial context.



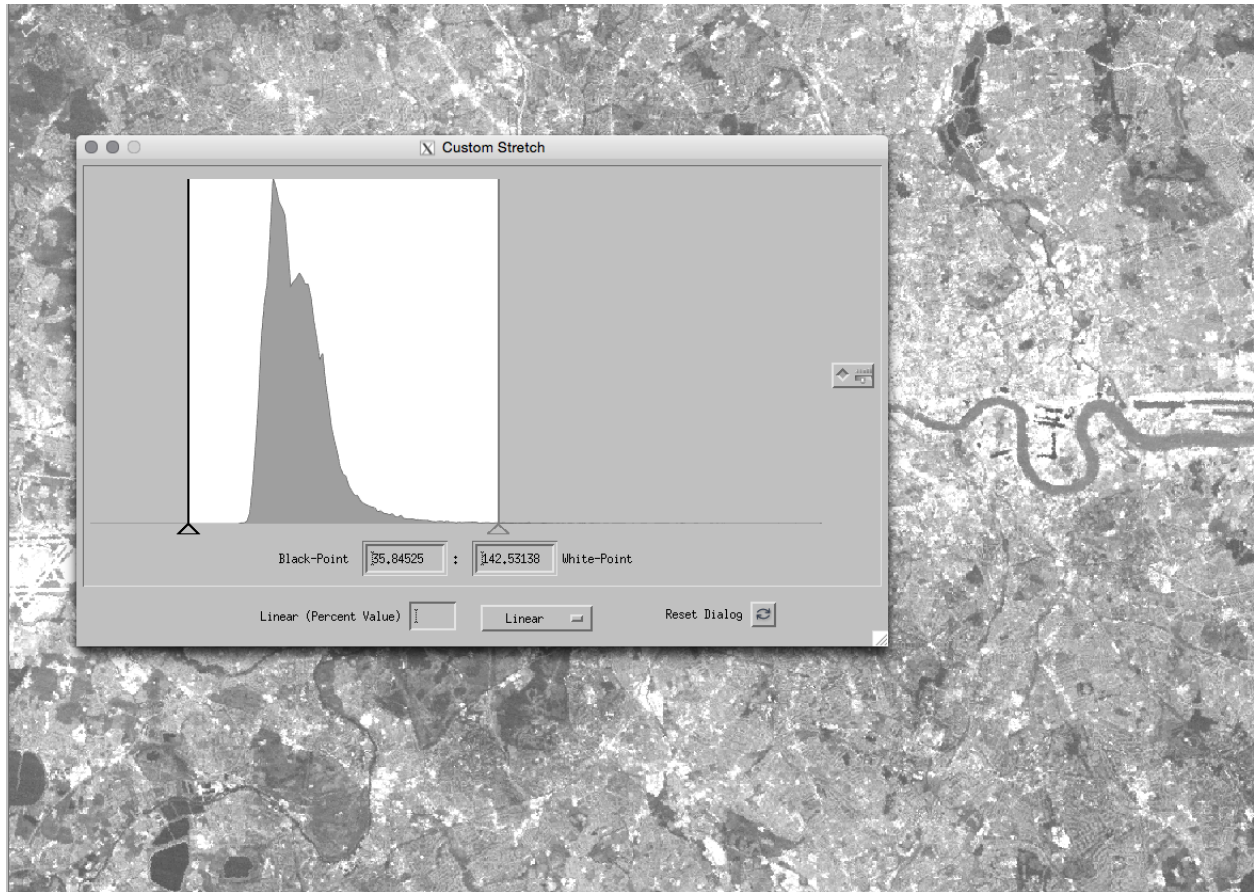
But there are other ways we can visualise the data to aid our interpretation.

For example, if we select the Histogram stretch button:



then we obtain a histogram view of the data, within which we can apply a contrast stretch by ‘moving’ the upper and lower thresholds.

For greyscale:



or three band colour:



This can be useful both for performing a contrast stretch

experiment with this so that you understand what the linear contrast enhancement is doing to the image DN

You may find it useful to save some histograms and images and put these in your notes. For UCL Desktop the easiest way to do quick saving is using screenshots. Find the On-Screen Keyboard and Paint apps. Use the combo of *Windows-Icon+Shift+S*, to start a screenshot on UCL Desktop, then go to Paint and *Paste*. The screenshot should appear in paint for you to save. Make sure the On-Screen Keyboard doesn't obstruct your image. You can also paste into a word doc if you want to take notes.


It is also of value in visualising the features of the histogram. The histogram shows the frequency of occurrence of the different DNs in the image. For example in the image above, we can describe the green waveband histogram (shown in green) as having two clear peaks and a long positive tail. The two peaks are most likely indicative of different land cover classes. The fact that there is a long positive tail may show a third 'bright' class present.

explore the histograms shown by some of the wavebands you have access to here

As you do this, think about the relationship between the informations shown in the image and that in the histogram.

- How could you try to use the histogram to perform a classification on the image into different land cover classes?
- What issues might you come across?
- Do some bands show better separability than others? if so, which (and why?)
- A histogram is a useful way of summarising the information in an image. In fact, we will often use even simpler descriptors to describe a frequency distribution such as this, one example being the *mean* and *standard deviation*. Would it be a good idea to describe these histograms in this way? If not, why not?

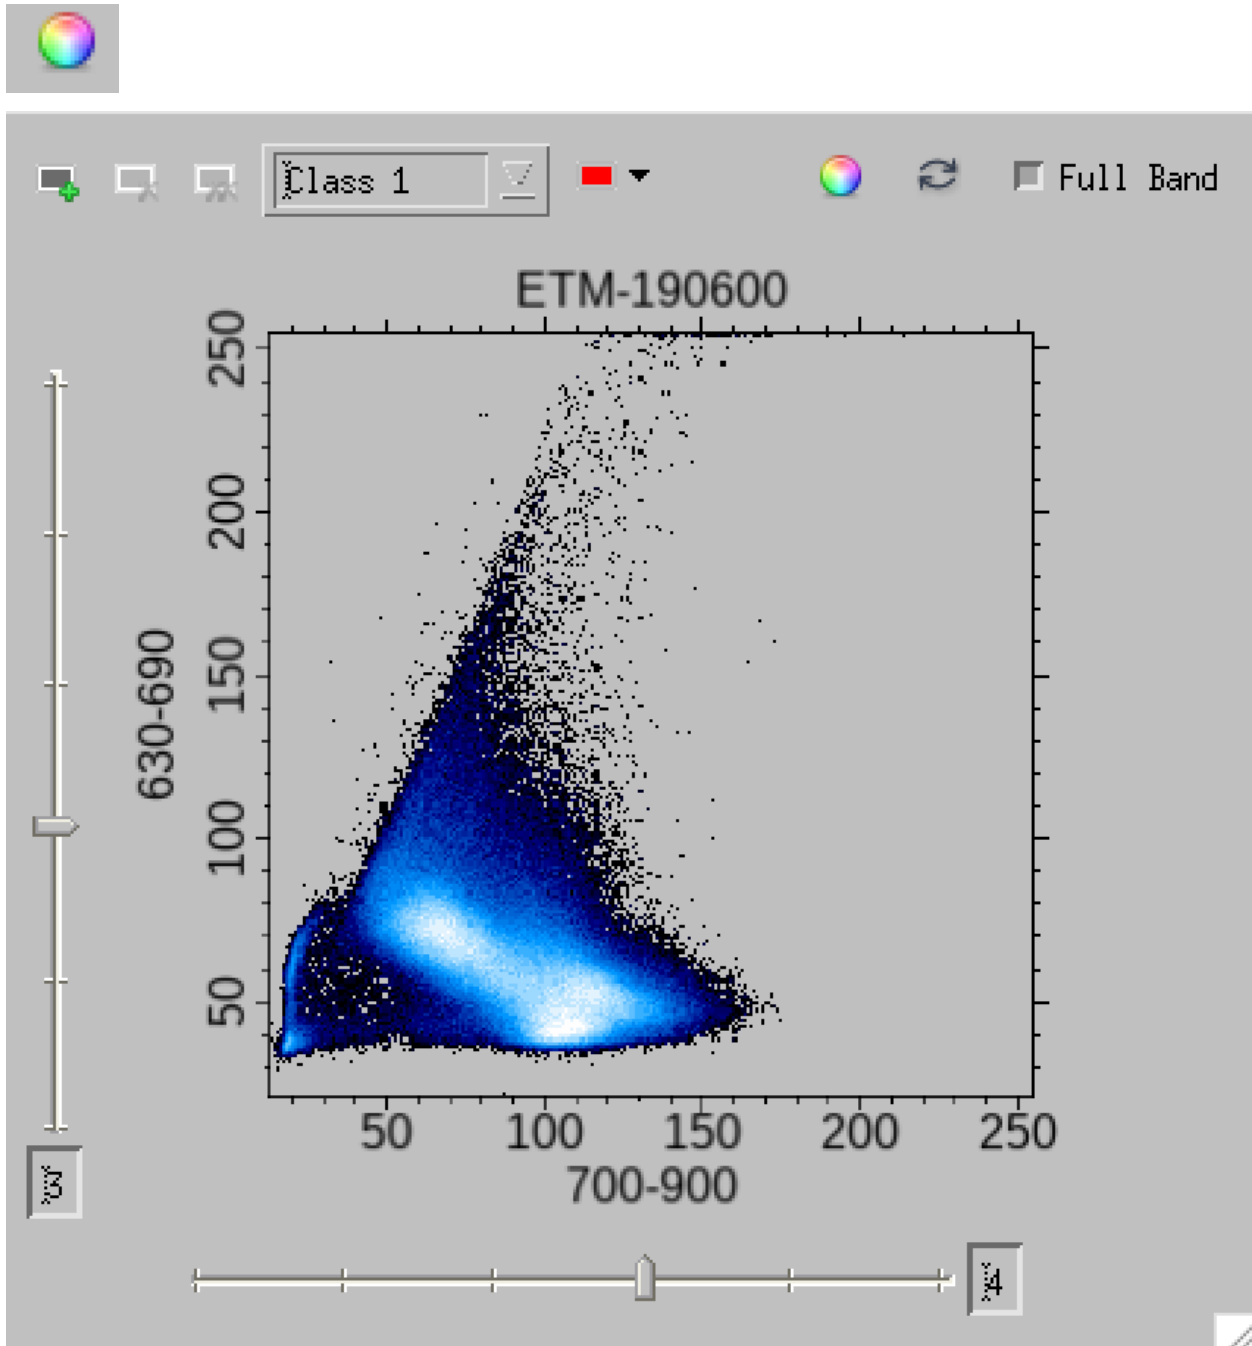
One way of extending our view of such summary data is to use a scatter plot (often called a 'feature space' plot).




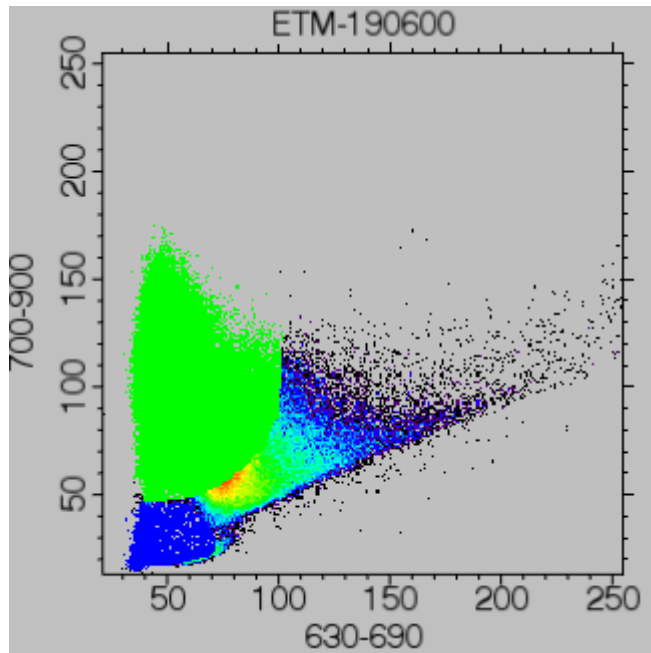
File->Select New Band

locate the Scatter Plot Tool button, and display a scatter plot with the Red waveband (630-690 nm) on the x axis and the Near InferRed (700-900 nm) on the y axis

Toggle the density slice button to show the scatter plot in colour (density slice)



You should now be able to define ('draw') regions onto the scatter plot, and the pixels within this specified region are then displayed in the defined colour on the image:



You should spend some time exploring this for different wavebands, as this concept forms the basis of many remote sensing algorithms (particularly those for classification)

As you do this, think about how the scatter plot and histogram information are related and also think about how you might get the computer to *describe* the regions you have drawn in the feature space plot.

Think also about how using only one or two wavebands may be a limiting factor in classification (i.e. there is often information in other bands that may allow separability).

Try using Options -> Mean All on the scatter plot tool. This shows you the class means across all bands. This should allow you to better think about how the classes you have drawn might be separable in different wavebands.

2.2.6 Spectral Profiles

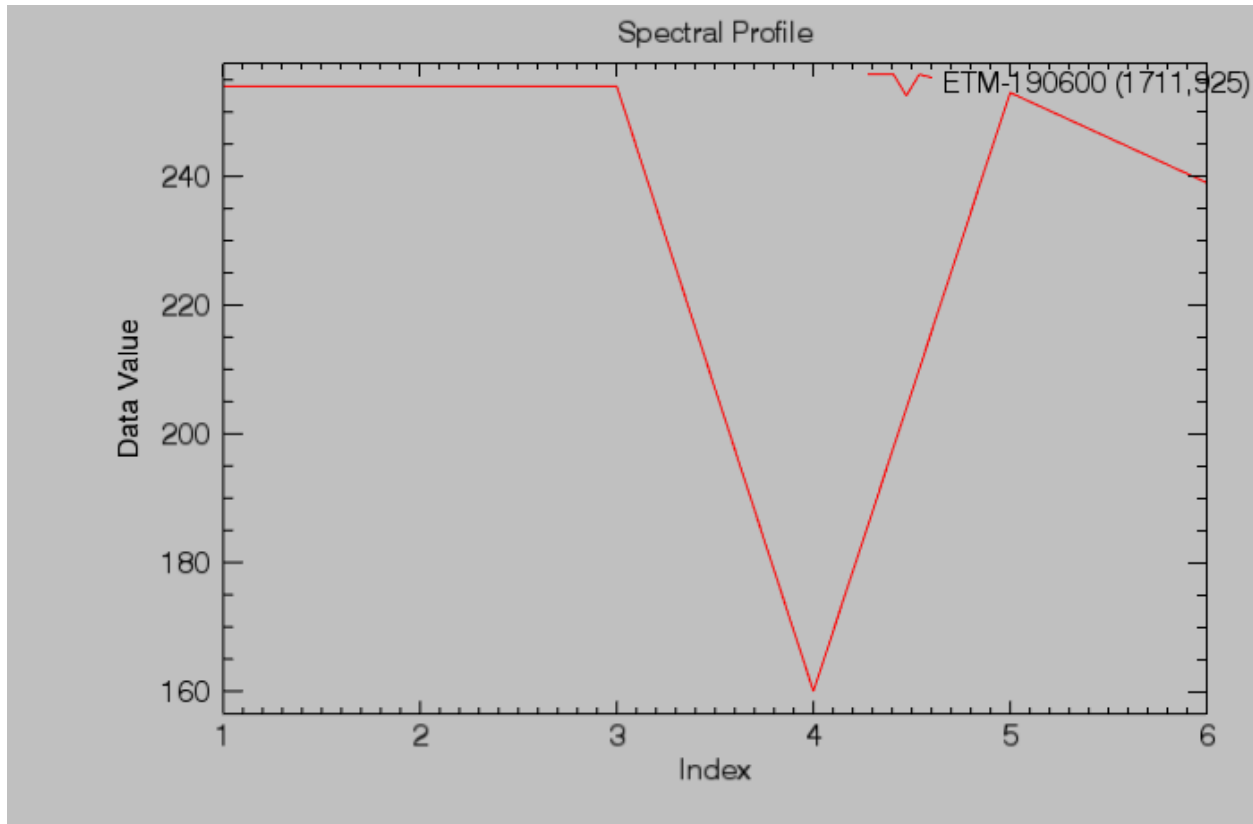
***In this section, we will learn how to do plot spectral profiles*.**

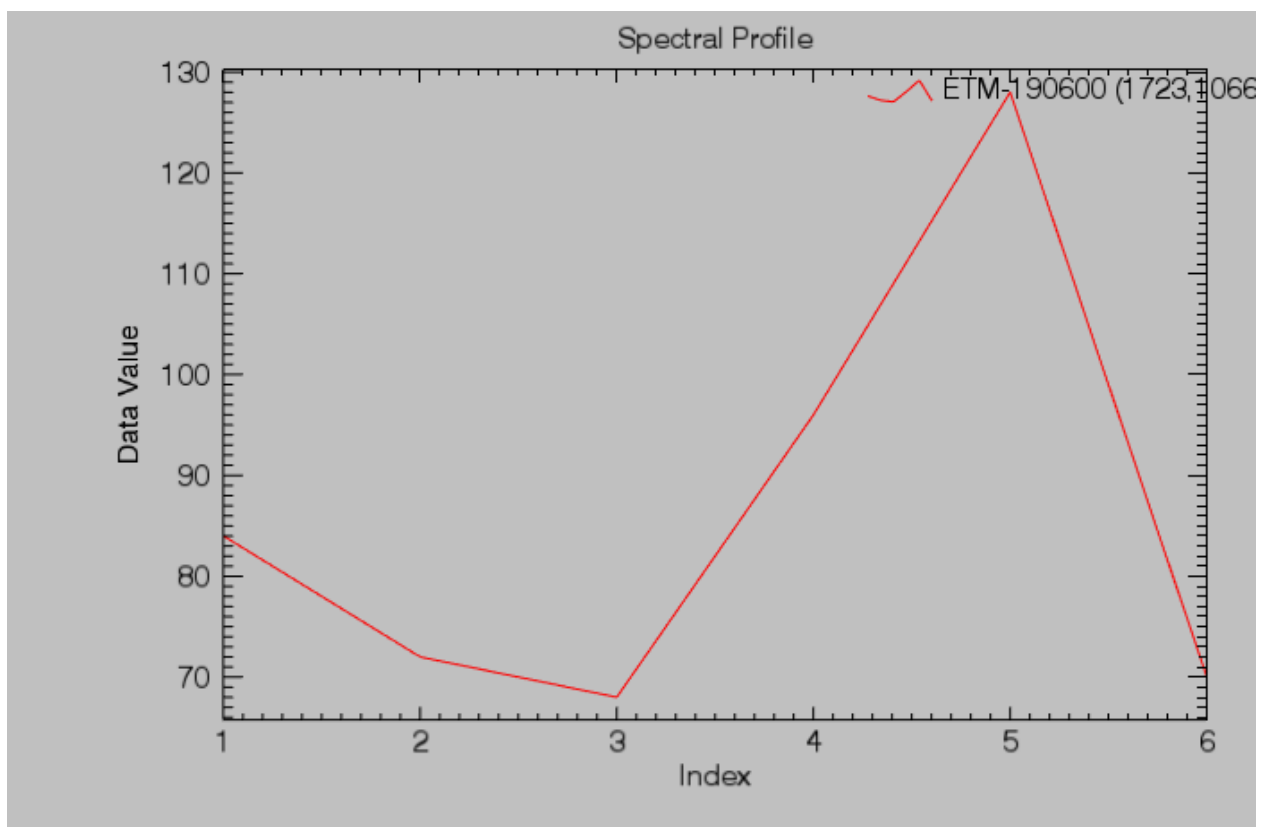
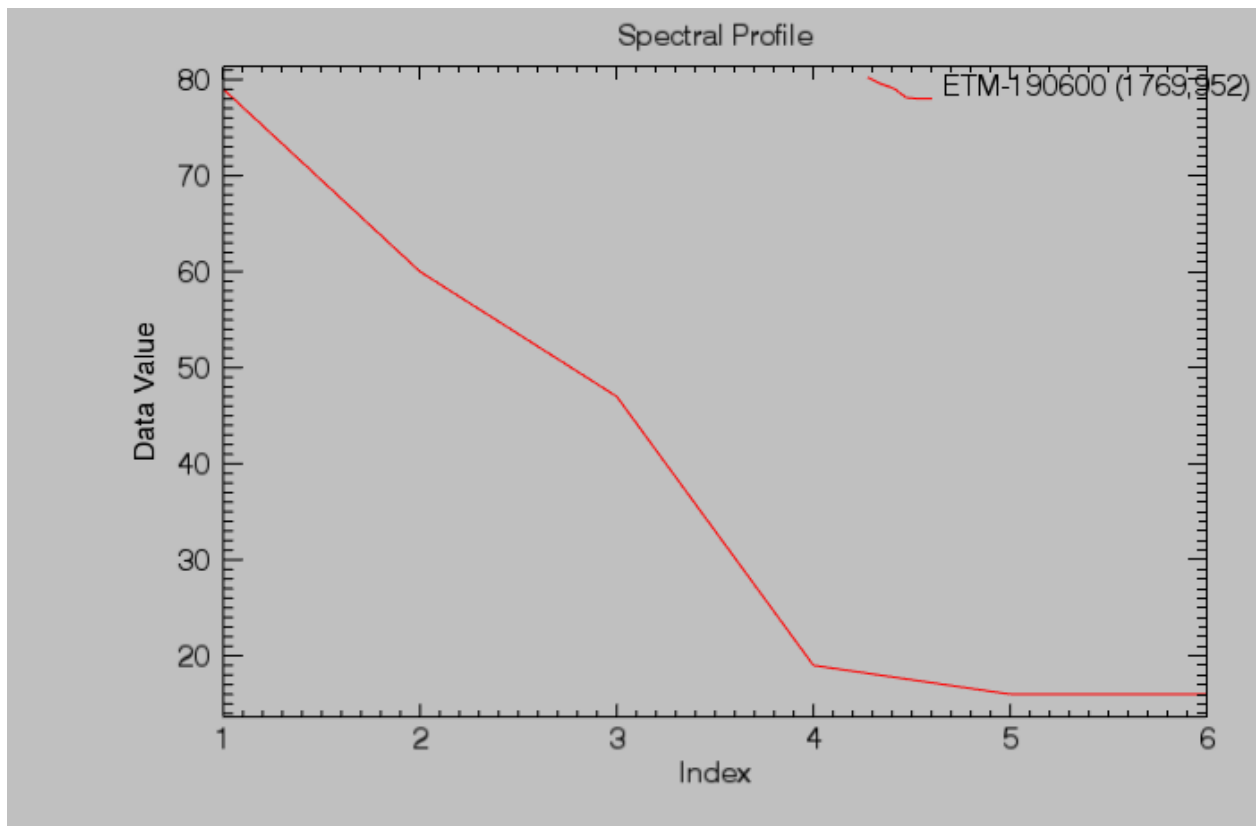
Note that you may need to close the scatterplot tool before doing this.

Although we can display (up to) three bands of information in a colour composite, we often wish to know more information.

The various Display->Profiles tools can help in this regard as we can look at the DN in all wavebands (for a particular pixel).

The following show some example spectral profiles:





See if you can navigate your way to these locations and/or work out what 'material' is shown by the spectral

profile. if you want to save these plots, make sure to give them a descriptive name that will make sense next time you find them.

2.2.7 Enhancement

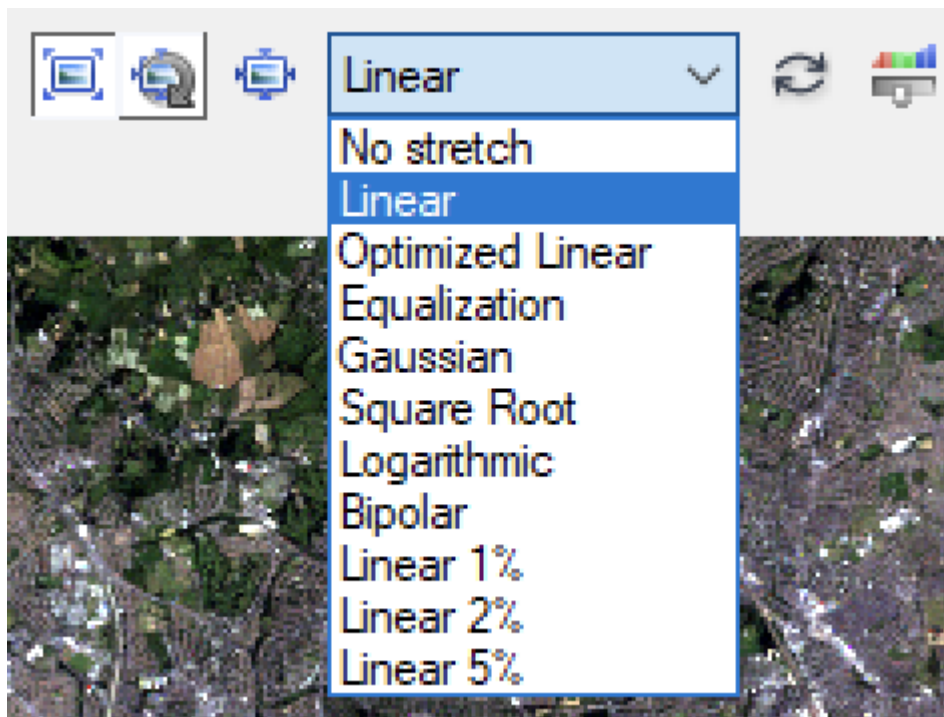
***In this section, we will learn how to perform a variety of enhancement tasks*.**

We have seen that ENVI has simple image display enhancement capabilities such as 'brightness' and 'contrast' variation via control sliders. Open the Histogram stretch window to visualise the histograms again.

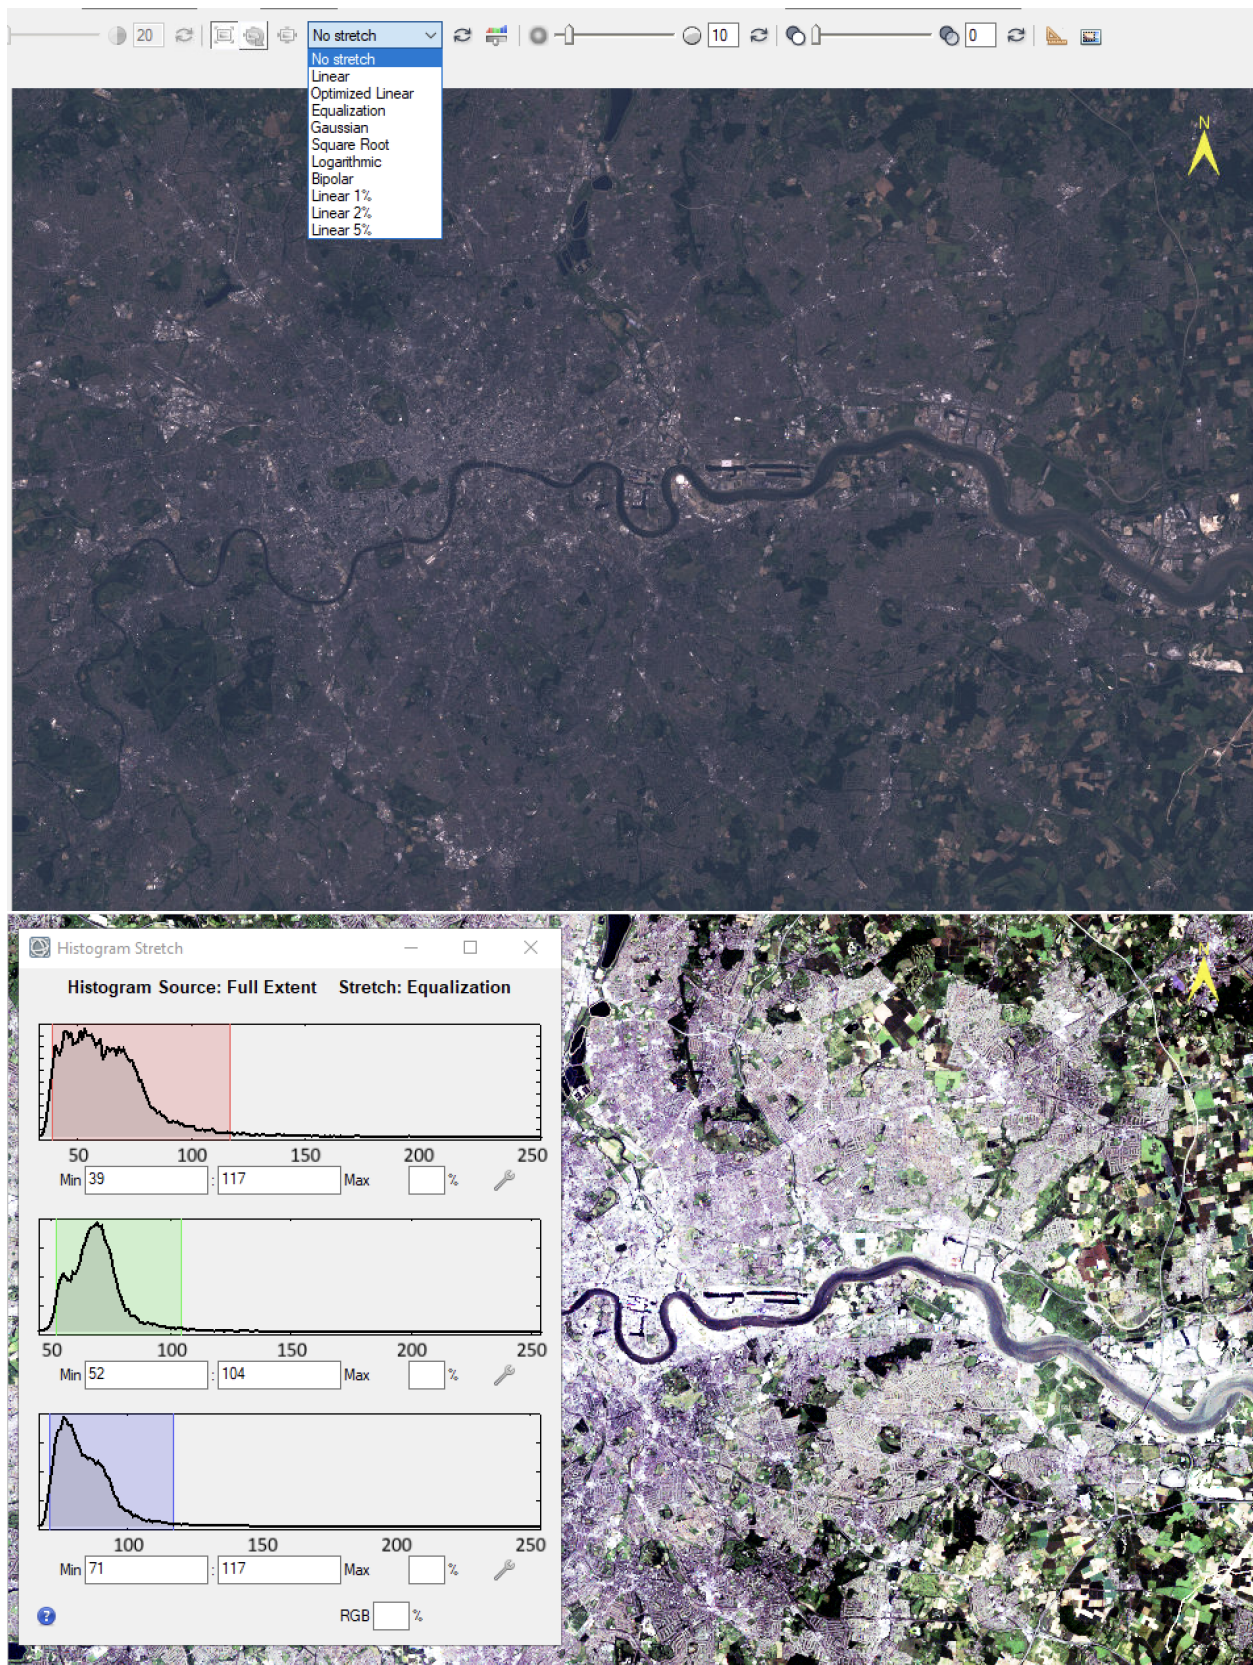
These are often useful first pass enhancement tools to enable you to more readily visualise features in the image.

Several 'automated' or semi-automated approaches also exist that may provide a useful enhancement under certain conditions.

These are available via a menu (Linear being the default) as:



For example, histogram equalisation is generally a useful visual enhancement method:



Explore the impact of various enhancement operations

You may find it easiest to do this using a 'greyscale' representation (so there is just a single waveband to think about). Think about how the different enhancements relate to the information in the histogram.

In particular, explore how you can use image enhancement to more readily differentiate different DN levels in a particular area of the image (e.g. the river)

2.2.8 Comparing Images

***In this section, we will examine some tools to use when comparing images*.**

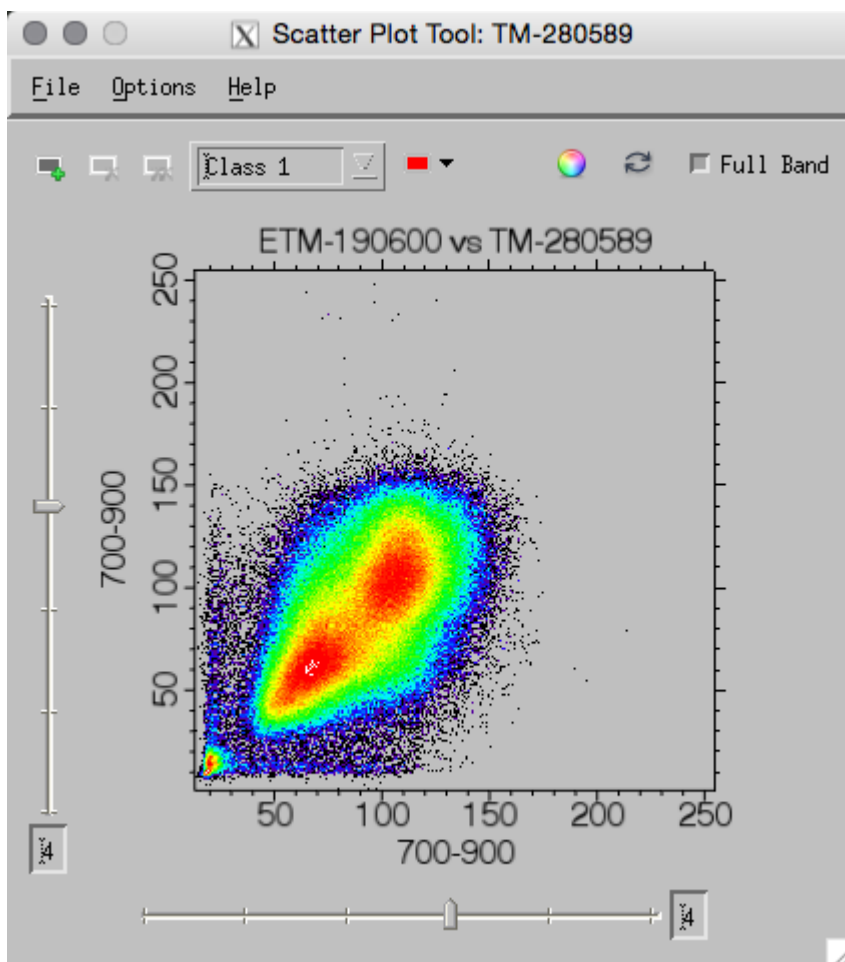
Often, we wish to compare features from more than one image (e.g. Landsat images taken over an interval of time).

There are several tools for exploring this in ENVI.


First, remove the current loaded images from ENVI (Note down how you do this!), then load two dataset ETM-190600 and TM-280589.

One immediately obvious thing you can do is simply toggle each dataset 'on and 'off once it is loaded (using the dataset check box).

You can use tools we have seen above, such as the scatterplot to visualise the scatter between the two datasets (File->Select New Band):



You can also very usefully 'draw' on this scatterplot to identify interesting features:



Transparency



Another option is to use the transparency slider to give different ‘weights’ of the two datasets. This might be useful, for instance to spot features that have changed between the two image dates:







Swipe and Flicker


Sometimes it is useful to ‘flicker’ between the various images or ‘swipe’ one over the other to spot features of (change) interest.

These can be done using the appropriate buttons on the command bar or under the Display option:

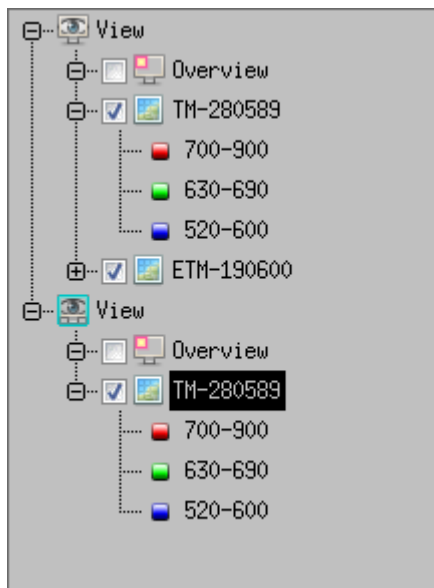


Views

One other option is to change the number of ‘views’ on the screen display:



This will create new ‘views’:



And you can drag and drop loaded datasets to the particular view you want.

Use these various comparison tools to highlight areas of difference between the two images

2.2.9 More Exploring

Some other interesting options to explore:

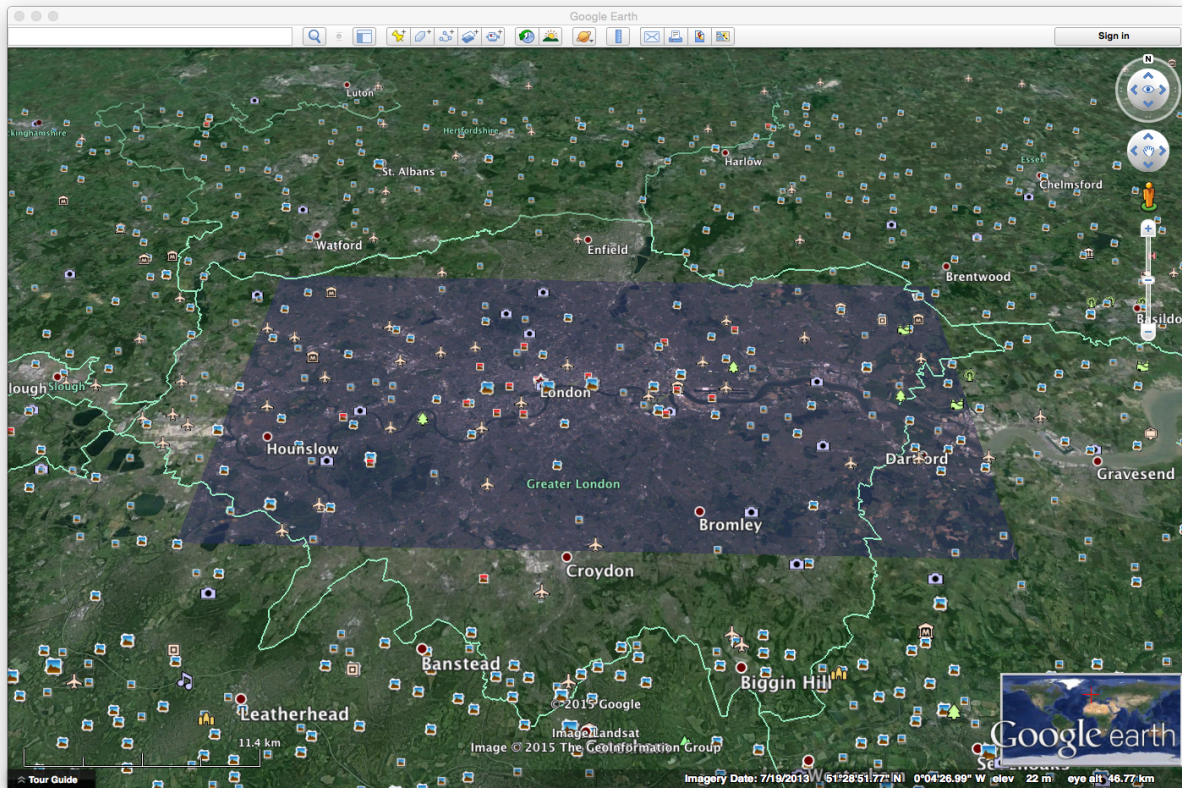
Annotation

On the control bar, you should see options for annotating the image. Experiment with annotating the image for a few features and work out how to save a 'picture' of what you have done (e.g. File -> Chip view to-> File ...):



Google Earth

If you have tools such as [Google Earth](#) set up on the computer you are using, try File->Chip view to->Google Earth



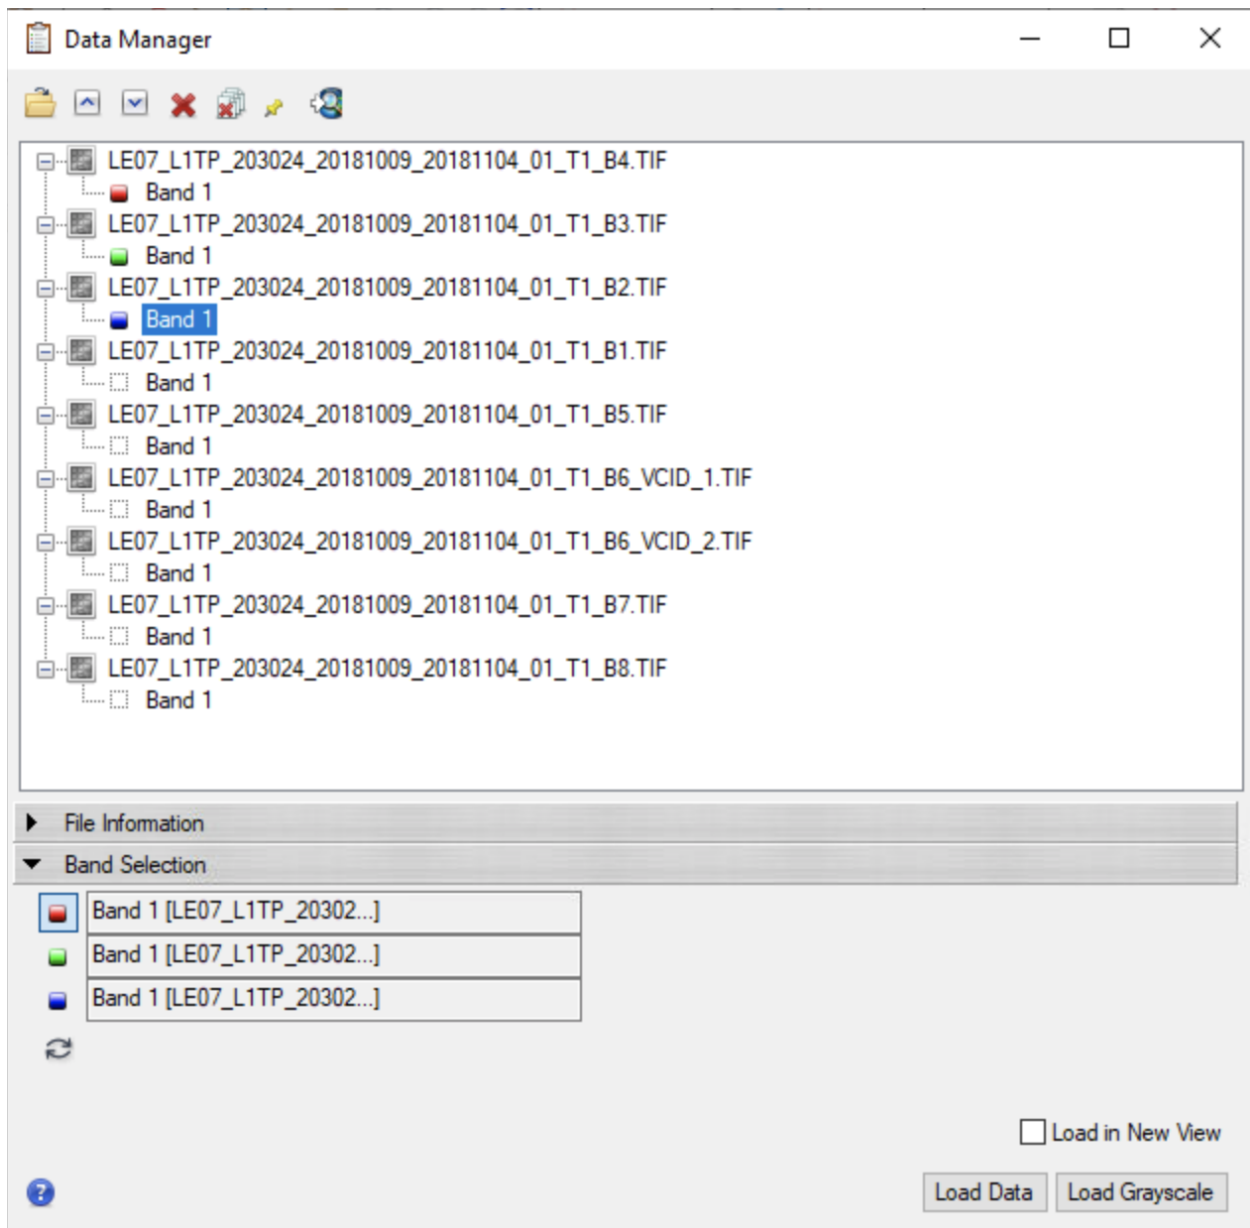
2.2.10 Summary

The main aim of this practical is to get you used to using the image processing software tool *envi*.

In this practical, we have loaded two Landsat images of London and used various tools within *envi* to visualise the data and their information content.

In particular, we have learned about:

- Image Display
 - False Colour Composites
 - Real Colour Composites
 - Greyscale display
 - Spectral profiles
- Feature Space
 - Histograms
 - Scatterplots
- Enhancement
 - Linear contrast enhancement
 - Histogram equalisation



2.3 Download data

In this section, you will learn how to order Landsat data from the USGS, including how to search only for the area we are interested in.

You will attempt to find two clear scenes (cloud-free or nearly so) for your area of interest.

You will then learn how to download the datasets.

Finally, we present one *manual* and one *automated* method for scaling the data, subsetting and producing a mask.

This is a skill you will need to complete the assessed coursework.

Once you have this skill, you can start collecting the datasets for the practical.

You do not need to use the ‘automated’ subsetting method: you can do everything manually if you wish.

2.3.1 Downloading and visualisation tools

The first task is to download the data that you will need.

There are several tools available to you for browsing and ordering Landsat data. See the [USGS](#) for a full description of the data products and where to order them.

You will most often be interested in **surface reflectance** products. These can be ordered through the [USGS Earth Resources Observation and Science \(EROS\) Center Science Processing Architecture \(ESPA\) On Demand Interface](#).

The first time you visit the ordering service, you will need to [register for an account](#). Make sure you remember the username and password you used for registration!

Once at the [USGS Earth Resources Observation and Science \(EROS\) Center Science Processing Architecture \(ESPA\) On Demand Interface](#) site, make yourself aware of the information on [surface reflectance](#).

2.3.2 Select Landsat Scenes

To select a particular landsat scene, go to the [Earth Explorer](#) site.

Enter the place you are interested in, in the [Search Criteria](#) (e.g. try: London, England) or select an area on the graphic using the [Use Map polygon](#) option.

Under [Data Sets](#), look in [Landsat](#) -> [Collection 1 Level 1](#) and check all of the boxes that have surface reflectance data. This will be for different sensors in the Landsat series (LS4-5, LS-7, LS-8).

You *may* wish to set a cloud cover limitation (e.g. 10%) under [Additional Criteria](#).

Click on 'Results' and look at a few of the images. In particular, you should look at *which* `PATH` and `ROW` https://landsat.usgs.gov/landsat_acq is appropriate for the area you want data for. This could, for example be

Path 201
Row 24

Now, you can go back to the [Search Criteria](#) and enter just the path/row that you want, and/or just go through and add the files you want to download to the basket. This is the bar you have: [Footprint](#), [Overlay](#), [Compare](#), [Metadata](#), [Download](#), [Bulk Download](#), [Order Scene](#), [Remove](#)



The first two are good for looking at, checking the cloud cover, and image cover. Choose either [Download](#), or [Bulk Download](#) to get the full image. Select the GeoTIFF, the largest file.

For bulk download you can view the basket, and submit the order. For this practical I suggest the [DOWNLOAD](#) option, getting immediate access to an image.

You should first try the Bulk Download with a small number of datasets, but it is quite straightforward to order many scenes if you need them.

If the order has worked, you will receive an email and also a notification when the data are ready for download. The order can be checked here <https://dds.cr.usgs.gov/queue/orders/>.

Once you get confirmation that the order is ready (this can, at times, take days but is usually less), you can then download the dataset.

2.3.3 Download ordered data

You can simply download the data by following the links in the email you receive.

Alternatively, there are various tools you can use for bulk download.

One example is `espa-bulk-downloader <>`_`, which is a python script. You use this, e.g.: <https://code.usgs.gov/espa/bulk-downloader>

```
download_espa_order.py -e p.lewis@ucl.ac.uk -d ~/Downloads -u p.lewis@ucl.ac.uk -p_
↳MYPASSWORD -o ALL
```

This is very handy if you wish to download a large amount of data to a remote network.

You can also use the USGS download app <https://dds.cr.usgs.gov/bulk/>, though this is unavailable on UCL Desktop.

2.3.4 Unpacking Downloaded data

The images from USGS come initially in a compressed `tar.gz` file. This can be unpacked using either a command line or terminal on a unix or Mac machine. For UCL Desktop use the Command Prompt app.

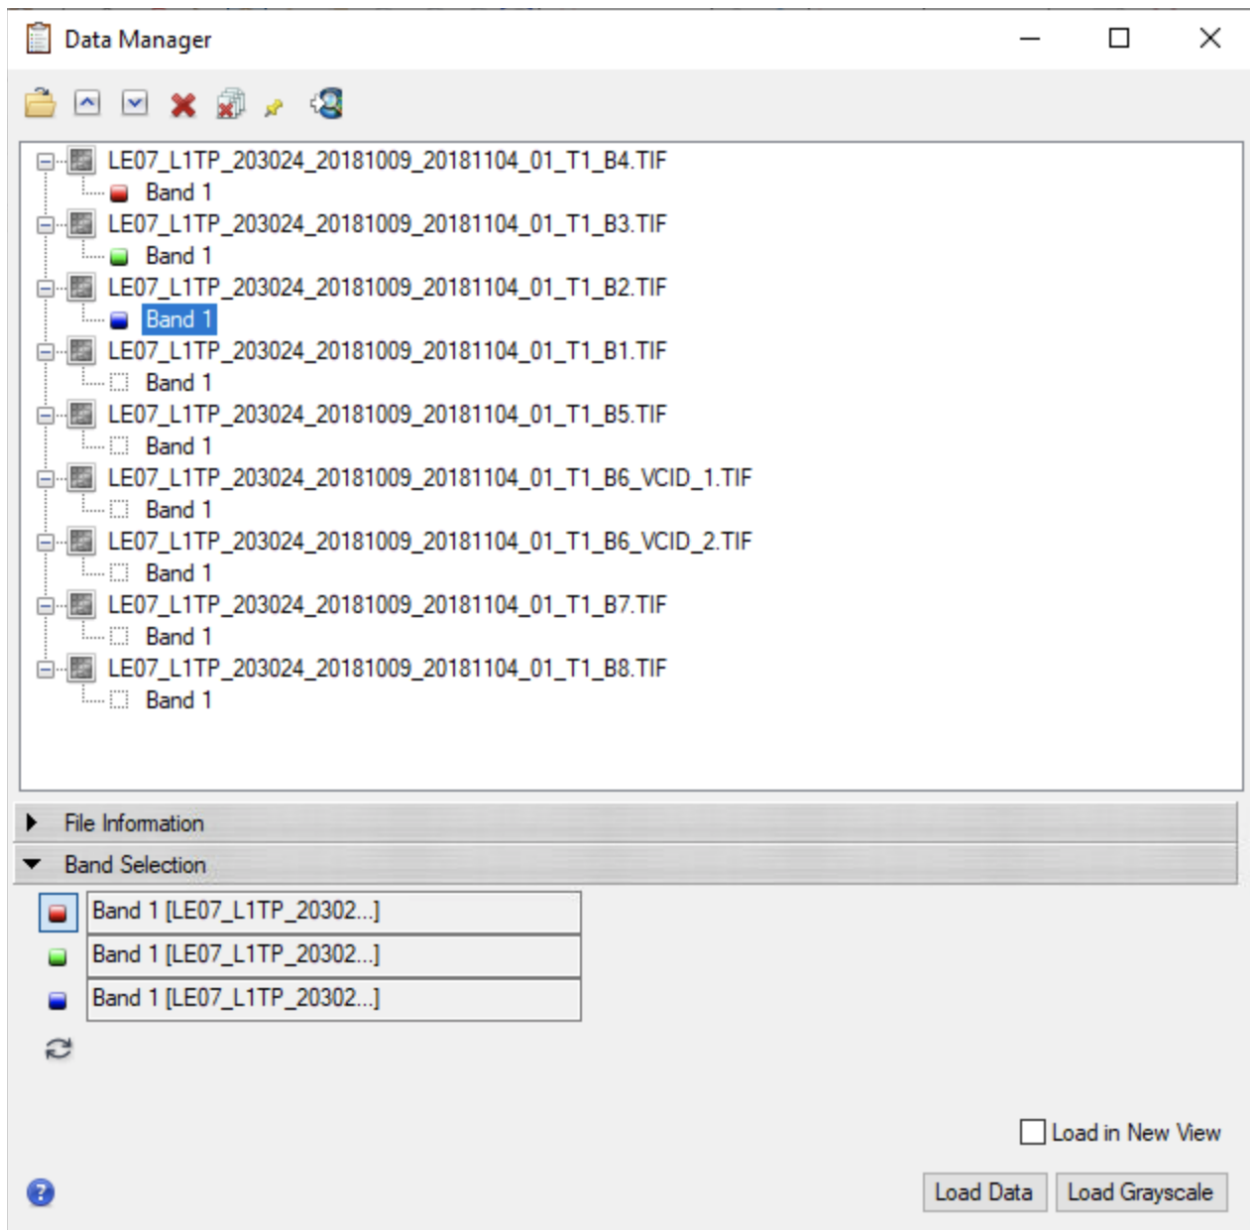
We will navigate to the downloaded file using typed commands. There exists a typed command for everything you may wish to do. First to get to the correct file system `N:` then get to the data folder `cd GEOG0027/data` (or wherever you downloaded the data) List all the files in this directory `dir` The downloaded file will be in this list. Then extract all the images using `tar -zxvf DOWNLOADED_FILE.tar.gz` (changing the `tar.gz` file name to the file you downloaded). You will see tiff files for each band.

2.3.5 Fixing Landsat 7 SLC issues

If you make use of Landsat 7 data (after [slc failure in 2003](#) you may wish to perform a ‘gap filling’ on the dataset prior to use.

There are many such algorithms. One you can use quite easily is mentioned [here](#).

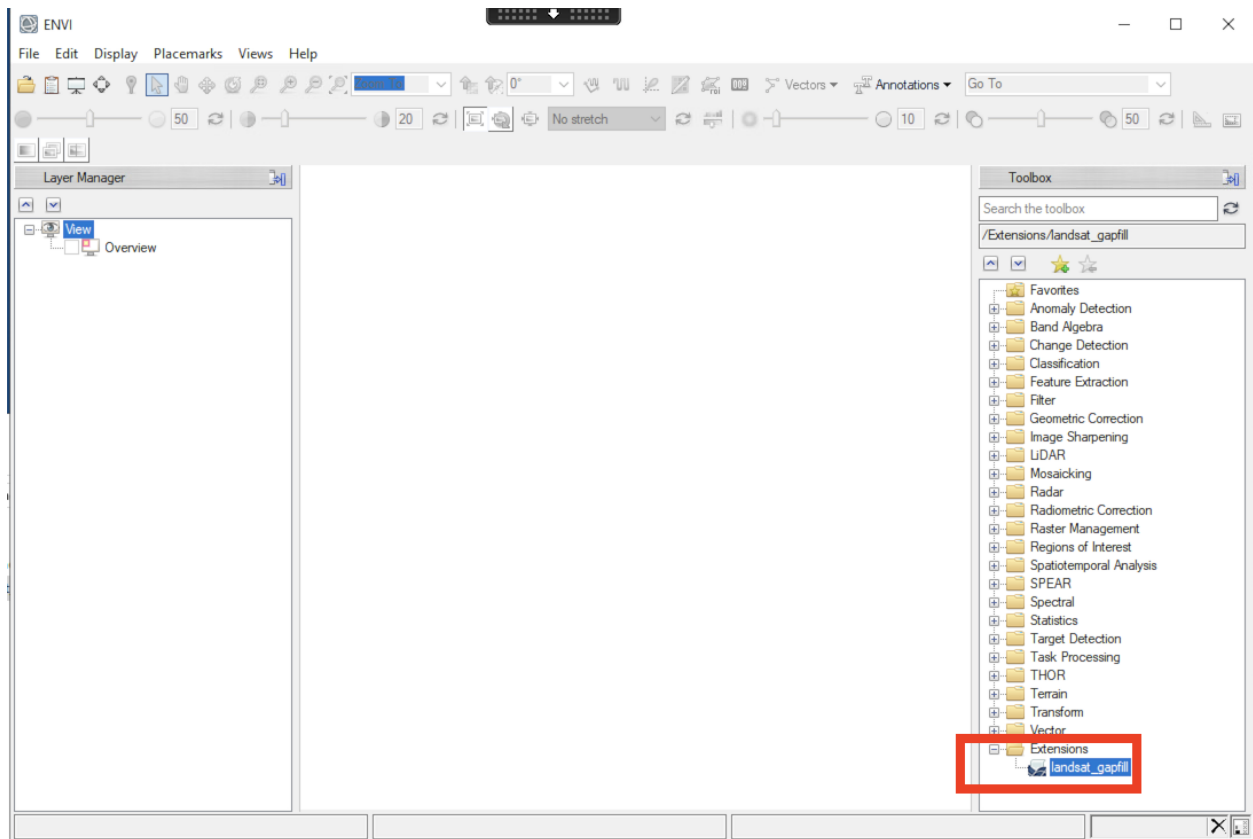
To use this, download the file [Landsat Gapfill IDL Model](#), either directly to, or copied to your `envi/extensions` directory (change folder permission if needed). From [Desktop@UCL](#), your `envi` extension folder should be `N:/.idl/envi/extensions5_5` or check your directory settings as shown below:



On a unix machine, this will involve e.g.:

```
cp ~/Downloads/landsat_gapfill1.sav ~/.idl/envi/extensions5_5
```

Then, you should find the model available to you the next time you run `envi`, under the `extensions` menu. To view the toolbox you need to press `F3`. You may need to use the on-screen keyboard on UCL Desktop to do this

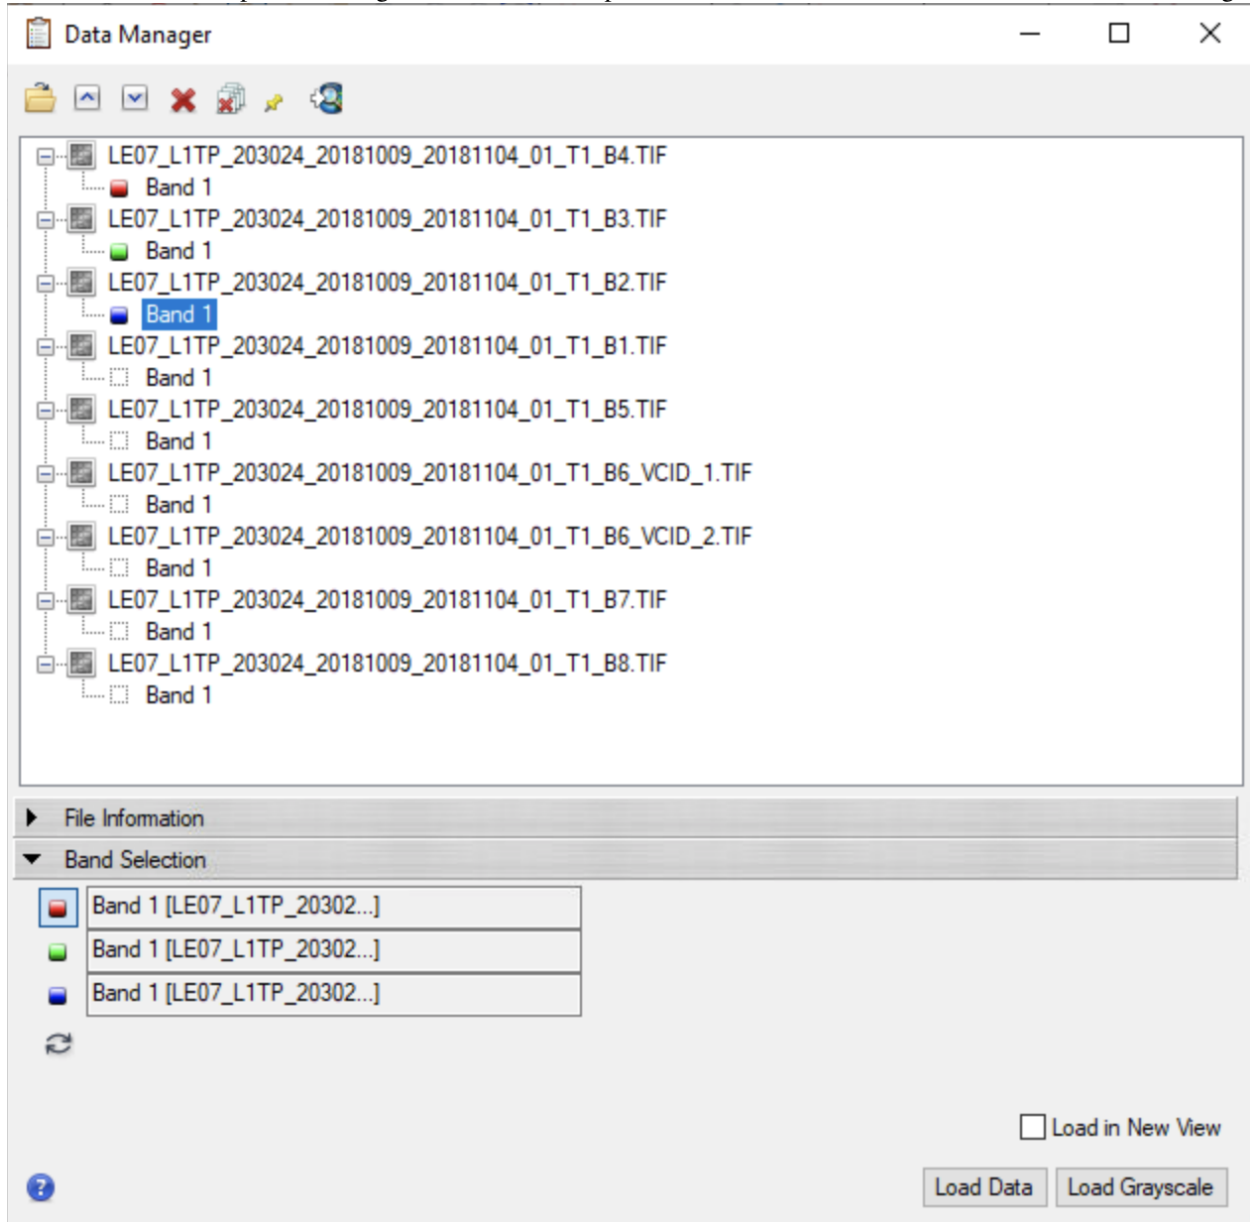


You may need to run this file separately for each waveband.

Once the extension is run it will wish to save the corrected image. You can rewrite the original image, or created a fixed version. As you can see we are creating a lot of data, so you may wish to make more folders to keep it all organised and tidy.

Viewing Multiband Files

To view a composite image from the separate .tiff files we use the Data Manager



To work out which file represents which wavelength read the documentation here <https://www.usgs.gov/faqs/what-are-band-designations-landsat-satellites>

Can you recreate a true colour image, or one of the diagnostic images discussed in this weeks lecture? Remember the skills presented in last weeks practical.

2.3.6 Download Scenes

You will be sent an email from the USGS.

You must keep this email to show evidence that you have ordered the data.

The email will give you a web address that will look something like <http://espa.cr.usgs.gov/ordering/order-status/p.lewis@ucl.ac.uk-01112017-113611-005>.

If you have recent or active orders, this web page should show you the status for each file (wait until it says ‘complete’).

Then, you can download the file (using the Download link).

You must keep a record of the download links to show evidence that you have ordered the data, e.g.:

<http://espa.cr.usgs.gov/orders/f.bloggs@nasa.gov-0101502242983/LE71220442000306-SC20150224104838.tar.gz>

Save all of the download links in a file to include in your report.

The filename will be something like:

LE71220441999255-SC20150217105511.tar.gz

There is no need to ‘unzip’ or ‘untar’ the file: just download it for the moment.

Make sure you remember which directory you have downloaded the data to.

You should probably download the datasets to your N: Drive (or DATA disk in the Unix lab), or you might run out of quota, or at least move the files to the N: Drive (or DATA disk) soon after.

As the files are quite large, it may take some minutes to download each scene.

[]:



2.4 Spatial filtering using ENVI 5.5

2.4.1 Aims

After completing this practical, you should be able to answer the questions: Which type of filter should I use for a given filtering application? What impact will the size and shape of the filter have on the output? You should have some understanding of the process (and issues) of spatial filtering of EO data using ENVI.

2.4.2 Data

The datasets you need for this practical are available from (skip this step if you've already downloaded them for the *Image Display* practical):

- ETM-190600
- ETM-190600.HDR
- TM-280589
- TM-280589.HDR

You should download these data and put them in a directory (folder) that you will remember!

The data you will be using are:

- six wavebands of a Landsat TM image over Greater London, imaged on May 28th 1989. The data were obtained from the GLCF which maintains a large database of (freely available) Landsat and other imagery. The data are at an original pixel spacing of 28.5 m, but have been resampled to a 25 m grid here. The data are in a Transverse Mercator projection with OSGB 1936 datum.
- six wavebands (nominally the same wavelengths) of a Landsat ETM image with 25 m spatial resolution, covering the same spatial extent. These data were obtained on June 19th 2000. The data were obtained from Landmap which contains a database available to Universities and other users through an Athens login (done via the institution you are at).

The wavebands are:

1	2	3	4	5	6
blue 450-520 nm	green 520-600 nm	red 630-690 nm	nir 760-900 nm	swir1 1550-1750 nm	swir 2 2080-2350 nm

The extent of the imagery is (Lat/Lon):

51°43'9.23" North, 0°36'18.37" West

to

51°16'29.32" North, 0°27'24.60" East

2.4.3 Introduction

***In this section, we load the image data we wish to explore*.**

The purpose of this practical is for you to build on practical 1 and learn about the process of spatial (convolution) filtering.


Note that convolution is a mathematical operation involving the modification of one function by another to produce a third (output) function. In spatial filtering this implies the operation of a filter (one function) on an input image (another function) to produce a filtered image (the output).

The session will be normally requiring at least two hours of effort. You may not complete all tasks in detail during the live session, so once you get the hang of how to use the tools, move on to the next section and return later to think more about the remote sensing.

There is a good material for this in text books (e.g. Jensen, Curran etc.) and some of this is online e.g. much of the Jensen material.

First, obtain and then load the TM and ETM images of London that we used in a previous practical.

View the ETM image as a FCC.

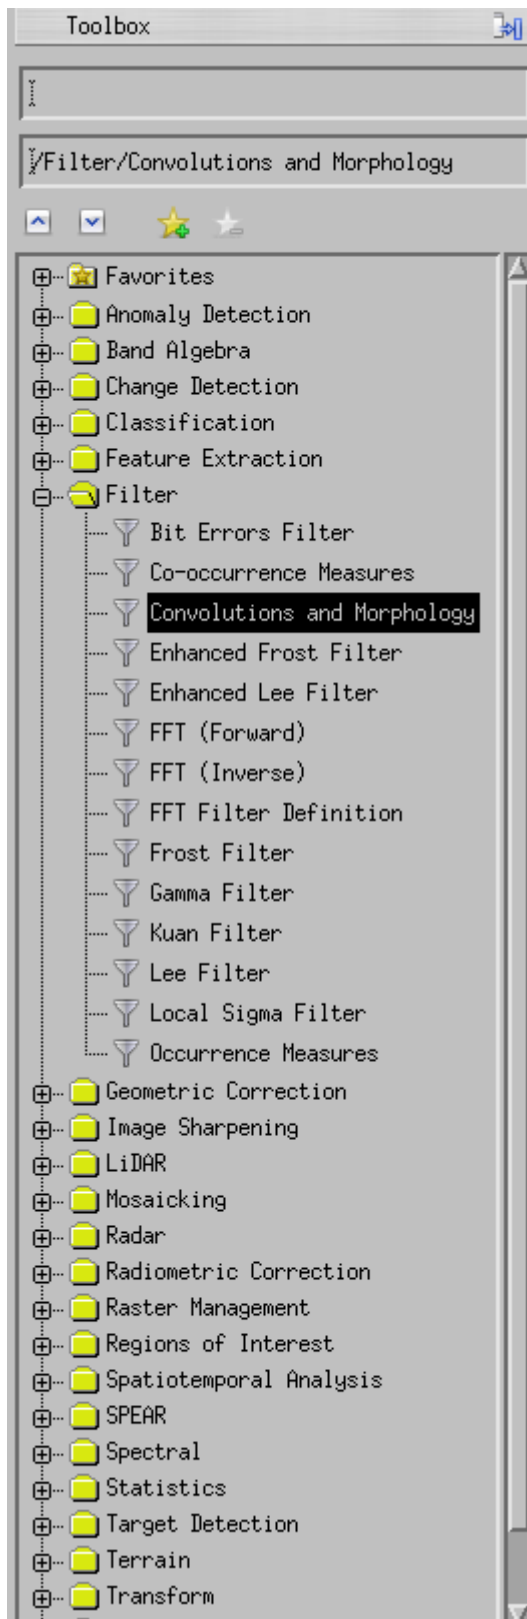


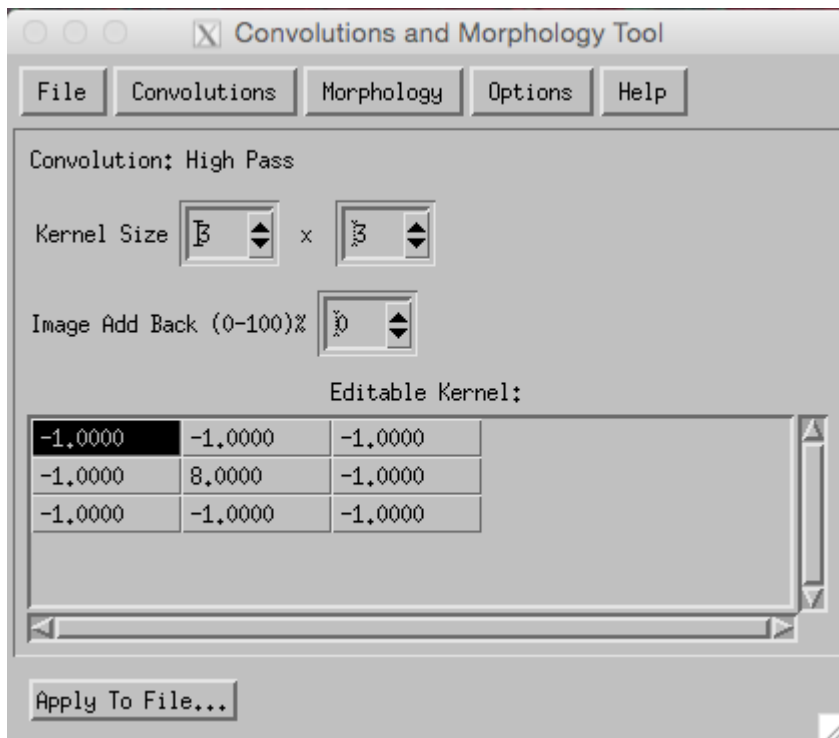
2.4.4 Convolution filtering

***In this section, we use various tools for image convolution*.**

A description of the various options for convolution and morphology are as [envi help pages](#). You should have a quick read over this if you are not familiar with the types of filter we will be using.

These operations are available via the Toolbox menu:






Apply a high pass filter (the one shown above) to the dataset

Now Display this and the original dataset using Views->Two Vertical Views with the convolution result in the left.


You should then Views->Link Views to 'geo-link' the two views (creating a new link, clicking on the left view as the 'anchor' and then the right as the 'reference'), so that moving around or zooming in one view is performed the same for both views. Make sure to move around the views to match each other before linking them (one simple way is to 'Zoom to full extent' for both views).



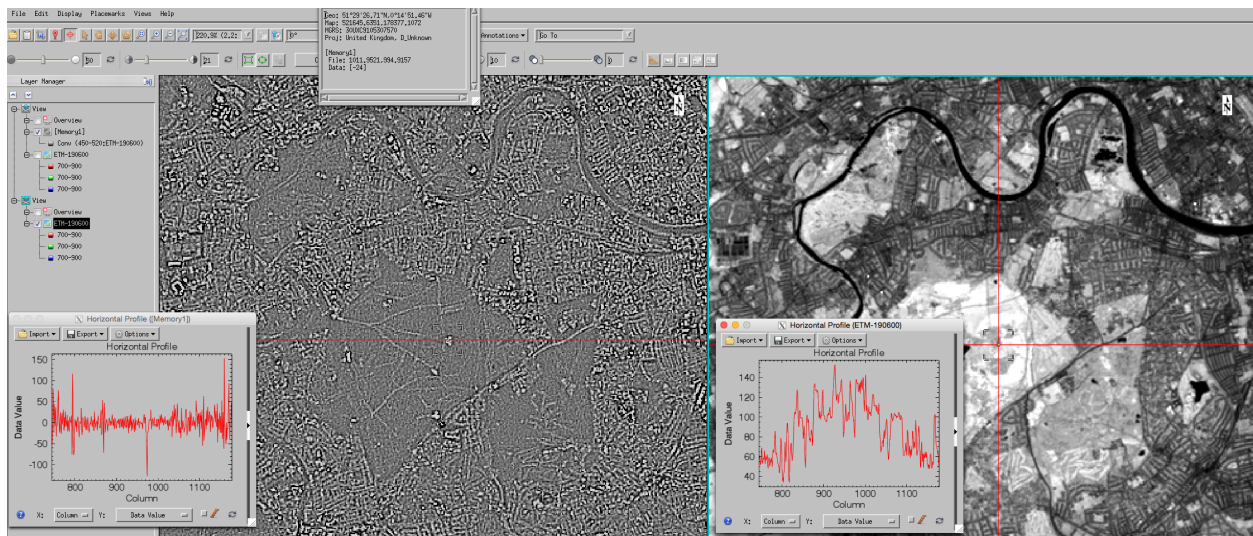
Make sure you make notes of what you have done and how you have done it so you can recall this later.

Now, find an interesting part of the image, for example Richmond Park, and try to understand how the filter is operating over the features you see.

You might find it easiest to use a greyscale image for this rather than false or real colour.



You may find it useful to examine horizontal or vertical profiles of the original satellite data:



You may need to zoom in to the image to see the detail of the transect, but what you should be interested in examining is the relationship between features in the original dataset and the high-pass dataset.

For instance, looking at the NIR band, for a transect going over Richmond Park we see higher DN over the park than surrounding (urban) areas (**why?**) in the original data. In the corresponding high pass filtered dataset (below), we see variation in the filtered result around zero, with some 'larger' features (spikes).

Consider carefully the numerical values in the convolution kernel:


—	—	—
-1	-1	-1
-1	8	-1
-1	-1	-1

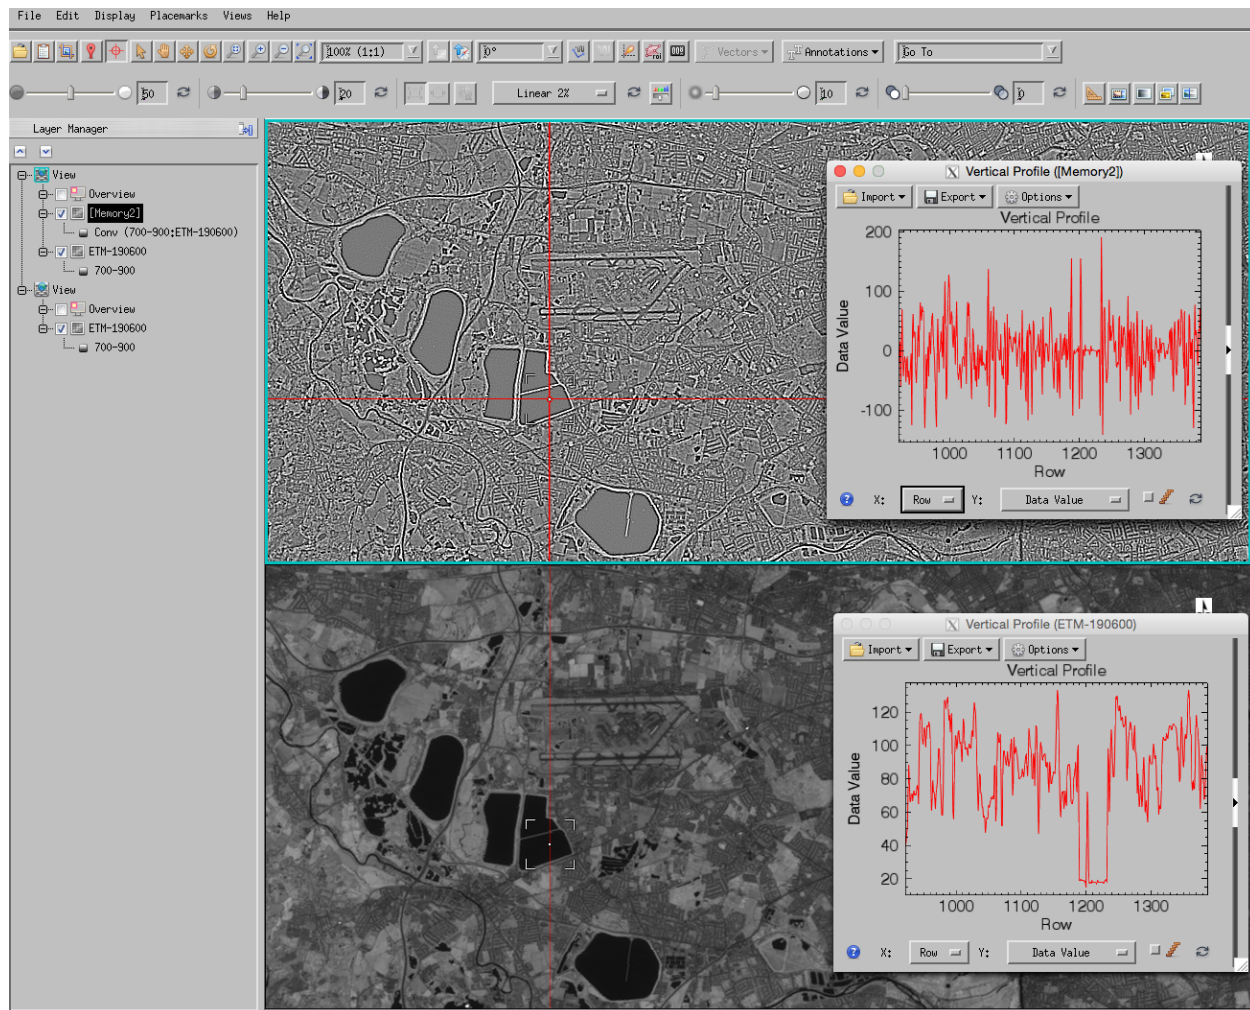
What does it indicate when there is a positive (or negative) spike in the filtered transect?

Why is there a number 8 in the centre of the filter?

what does the filter produce if the image data are constant? ('flat')

what does the filter produce if the image data show a step change? (see e.g. resevoirs)





2.4.5 Convolution

To think some more about this, it can be instructive to consider filtering in one dimension (rather than the two dimensions of an image). To aid this, we can consider some prototype ‘shapes’ and look at how they respond under the filtering operation.

Typical example would be step and ramp features.

```
[1]: %matplotlib inline
```

```
[2]: run python/proto.py
```

```
<IPython.core.display.Javascript object>
```

```
<IPython.core.display.HTML object>
```

```
<IPython.core.display.Javascript object>
```

```
<IPython.core.display.HTML object>
```

Two basic operations we can perform are looking at the *local average* (e.g. mean) or *local difference* with a convolution operator.

For a local mean over extent 3, we would use

—	—	—
0.333	0.333	0.333


What do the numbers add up to? why is that so

what would a filter of extent 7 look like?

We would expect the local mean to provide some ‘smoothing’ of the function:

```
[7]: run python/protoLow.py
```

```
filter [0.14285714285714285, 0.14285714285714285, 0.14285714285714285, 0.
↪ 14285714285714285, 0.14285714285714285, 0.14285714285714285, 0.14285714285714285]
```



We can see that this is the case in the illustration (filter width 7) the impact (of smoothing) on the ‘clean’ signals is mostly negligible (which is an effect we would want), and for the noisy datasets, the noise is reduced.

We can also note some ‘edge effects’ (e.g. the step at x = 100):

why do these occur?

does the extent of the edge effects depend on the extent of the filter? If so, why

what might be done to mitigate such effects?

A local mean (low pass) filter then, such as that shown above reduces the high frequency components of the signal and retains the low frequency components (so we can ‘get rid of’ (reduce)) random noise in the signal.

A local mean is effectively a local *integration* of the signal.

The ‘opposite’ of this is differentiation, i.e. the difference between one pixel value and the next.

The simplest filter for this is:

—	—
1	-1

which we can call a first order difference filter.

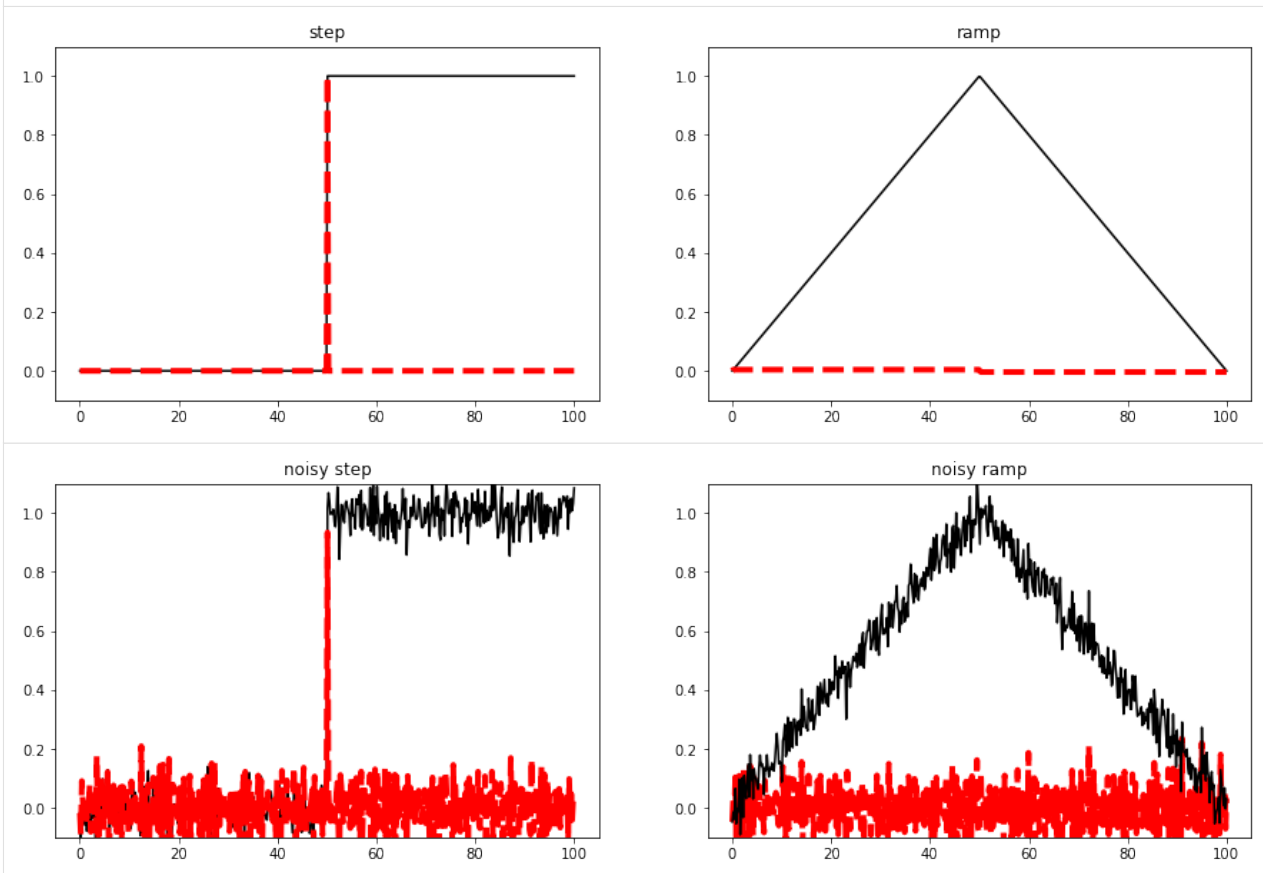
what do the values in the filter add up to in this case?

why is this so?

The impact of this is illustrated below.

```
[1]: run python/protoHil.py
```

```
filter [1.0, -1.0]
```



Think through *why* it has the impact that we see.

Of what use might such a filter be?

If we convolve the first order differential filter:

—	—
1	-1

we get a *second order differential filter*

—	—	—
1	-2	1

Whereas the first filter represents the ‘rate of change’ (e.g. slope) of the signal, the second order filter gives the rate of change of the rate of change (rate of change of slope).

what do the values in the filter add up to in this case?

why is this so?

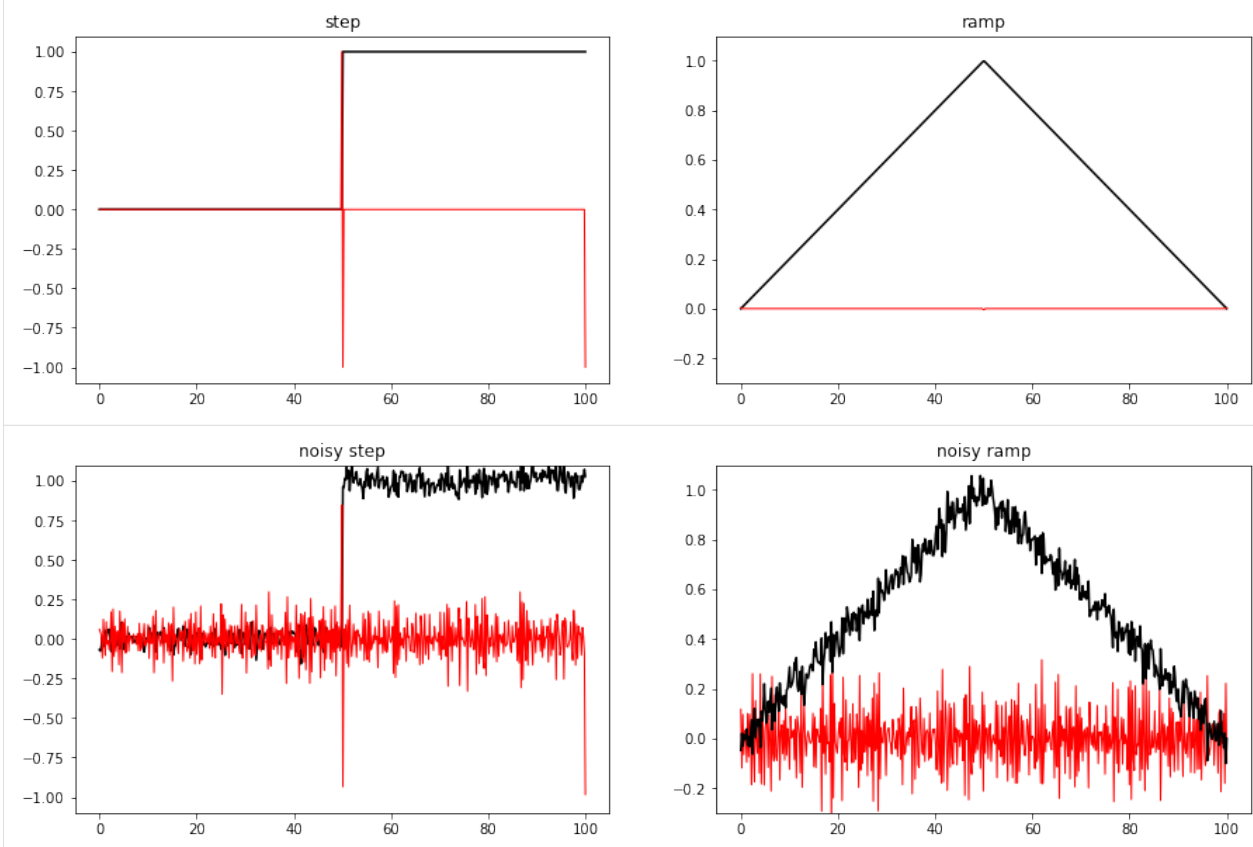
To illustrate this:

```
[9]: filt = [1.0,-1.0]
print (np.convolve(filt,filt))

[ 1. -2.  1.]
```

```
[10]: run python/protoHi2.py

filter [1.0, -2.0, 1.0]
```




At first sight, this might look quite similar to the result of the first order differential.

However, if we look more closely:

```
[11]: run python/protoHi2zoom.py
```

```
filter [1.0, -2.0, 1.0]
```



we see e.g. that the ‘edge’ of the step edge is mapped to a ‘zero crossing’ in this case, rather than a local maximum as was the case for the first order filter.

The peak of the ramp shows a small negative value in the second order differential data: ***why is this so?***

What would the result look like if we were to replace the step and ramp functions by one minus step and one minus ramp?

As a final example, let us consider what happens if we *smooth* (low pass) and then *differentiate* (high pass):

```
[12]: # show the result of convolving one filter with another filter
filt1 = [1.0, -1.0]
w = 3
filt2 = [1./w]*w
print ('differential',str(['%4.2f%i' for i in filt1]).replace("'",''))
print ('smoother    ',str(['%4.2f%i' for i in filt2]).replace("'",''))
conv = np.convolve(filt1,filt2,'full')
print ('combined    ',str(['%4.2f%i' for i in conv]).replace("'",''))

differential [1.00, -1.00]
smoother     [0.33, 0.33, 0.33]
combined     [0.33, 0.00, 0.00, -0.33]
```

Applying two filters to an image/signal has the same effect as convolving the two filters together and then applying this.


We see that convolving a filter of extent 2 with one of extent 3 gives a new filter of extent 4. ***Why is this so?***

To demonstrate, in the plots below, we first convolve the step and ramp with filter 1 (***what type of filter is this?***)

(blue line), then convolve this with filter 2 (***what type of filter is this?***) (green line). Applying the combined filter to the data, we get the dashed red line (which is the same as the green line).

```
[13]: run python/protoHiLow.py
```

```
filter 1:      [1.00, -2.00, 1.00]
filter 2:      [0.14, 0.14, 0.14, 0.14, 0.14, 0.14, 0.14]
combined filter: [0.14, -0.14, 0.00, 0.00, 0.00, 0.00, 0.00, -0.14, 0.14]
```



Thinking through these graphs should re-inforce what you have learned about spatial convolution in the lectures. You should be able to describe what the impact of different forms of filter on such prototype functions would be. You should then be able to translate this knowledge back to examining the image data.

Go back to the ``envi`` session and examine the impact of first and second order differential and smoothing filters on the image data

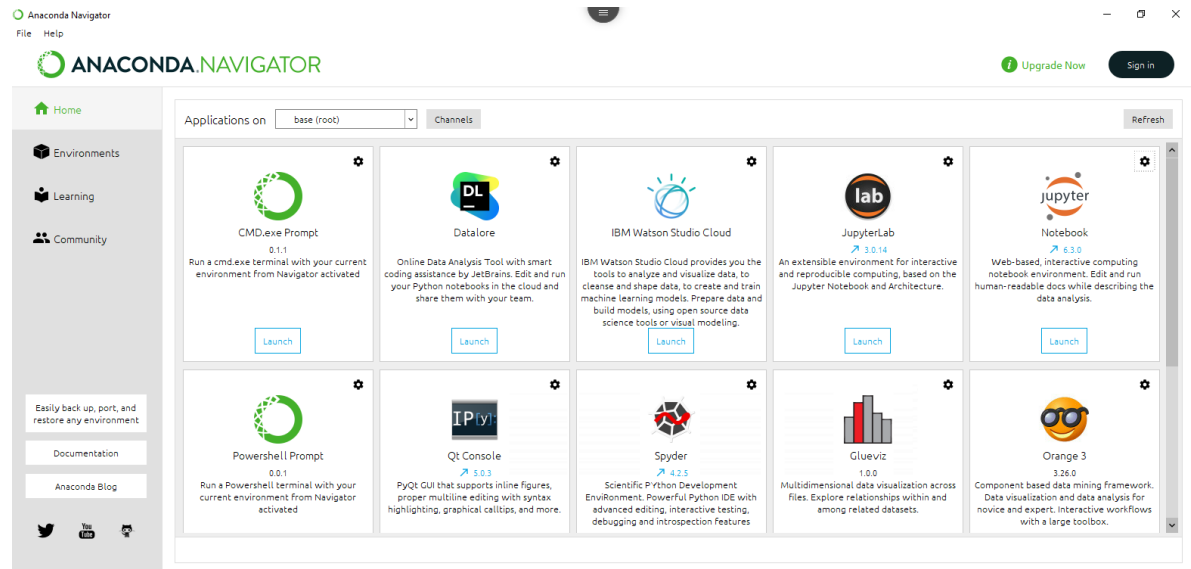
Relate what you see (e.g. in looking in detail at transects) to the impacts on the prototype shapes

Of what use might all of this be?

2.4.6 4. Extras (optional)

If you run these notes as an iPython notebook, then this shows an interactive tool to see the impact of filtering. **Additionally, if you open our practical notebooks (you may have noticed the xxx.ipynb file extension) with Jupyter, it also allows you to make changes to the notes, and further down the line, to modify or even write your own code (e.g. for the coursework)!**

- 1) If you don't have **iPython Jupyter** installed on your PC, please use Anaconda on UCL Desktop. Apps -> Anaconda3 (64bit) -> Anaconda Navigator Ignore requests to update.
- 2) Download the practical notes from here <https://github.com/UCL-EO/GEOG0027/>, using the green button say Code, under which you'll find Download zip. Then extract the .zip file.
- 3) Launch a Jupyter Notebook (Top right). This will appear as a webpage. Use the file navigator to find the extracted



notes.

4) Find this file `SpatialFiltering.ipynb`, and open. It will display it in a web browser.

5) If needed, delete the zip file(s) to free up space.

From here, you can use the following interactive tool to explore the trade-offs you need to consider when filtering: e.g. a wider smoothing filter will suppress a higher level of noise (i.e. do more smoothing) than a narrow filter. However, with a wider filter, you are likely to suffer from greater 'edge effects'.

[43]: `run python/interact.py`

```
-1 (498,) (500,) (499,)
```

```
<IPython.core.display.Javascript object>
```

```
<IPython.core.display.HTML object>
```

2.4.7 Summary

The main aim of this practical is to reinforce your understanding of convolution operations using the image processing software tool `envi`.

In this practical, we have loaded Landsat images of London and examined the application of high and low pass filters.



2.5 Classification using ENVI 5.5

2.5.1 Aims

After completing this practical, you should be able to analyse one or more image datasets using classification methods. This includes learning to identify land cover classes in a dataset and consider class separability (using histograms, scatterplots and other tools), and applying and assessing a classification product using Envi.

2.5.2 Data

The datasets you need for this practical are available from the [Classification data \(Rondônia\)](#) folder on moodle or you can download them individually:

- [ETM-110801](#)
- [ETM-110801.HDR](#)
- [TM-250792](#)
- [TM-250792.HDR](#)
- [SRTM-2002](#)
- [SRTM-2002.HDR](#)

You should download these data and put them in a directory (folder) that you will remember!

The data you will be using are:

- six wavebands of a Landsat TM image over Rondonia, Brazil, imaged on 25th July 1992. The data are at an original pixel spacing of 28.5 m.
- six wavebands (nominally the same wavelengths) of a Landsat ETM image with the same spatial resolution, covering the same spatial extent. These data were obtained on 11th August 2001.
- Digital Elevation model (DEM) data, obtained by RADAR interferometry from data on the SRTM ([Shuttle Radar Topography Mission](#)), are also available for the site. The data have been resampled to the same resolution and area as the TM/ETM data above.

The wavebands are:

1	2	3	4	5	6
blue	green	red	nir	swir 1	swir 2
450-520 nm	520-600 nm	630-690 nm	760-900 nm	1550-1750 nm	2080-2350 nm

The extent of the imagery is (Lat/Lon):

$$11^{\circ}1'31.29''S, 62^{\circ}58'27.57''W \rightarrow 11^{\circ}57'4.75''S, 62^{\circ}1'55.96''W$$

The full SRTM data can be loaded into [google earth](#), if you have access to this.

Although you have the data 'pre-packaged' for this practical, you can download your own datasets using the USGS [Glovis](#) tool:

USGS Home
Contact USGS
Search USGS

USGS Global Visualization Viewer System Notices: 1 (New), Critical

CRITICAL MESSAGE: Webpages and Services Unavailable Thursday, February 5, 2015 6:00pm - 10:00pm CT

Collection Resolution Map Layers Tools File Help

Downloadable

WRS-2 Path / Row: 231 68 Go
Lat / Long: -11.6 -62.4 Go
Max Cloud: 100% [Up] [Down]
Scene Information:
ID: LE72310682014227CUB00
CC: 0% Date: 2014/8/15
Qty: 9 Product: ETM+ L1T
Aug 2014 Go
Prev Scene Next Scene
Landsat 4 - Present List
Add Delete Send to Cart
1000m No Limits Set

USGS

Quick Start Guide	User Guide	What's New!
Browser Requirements	Download Source Code	About Browse Images

We can of course explore the area in Google Maps, which we may find useful for exploring the classification.


2.5.3 Introduction

***In this section, we load the image data we wish to explore*.**

Importantly, we have the ability to map these changes from the archive of satellite data, particularly data from the Landsat series of satellites. An excellent introduction to visualising environmental change from Landsat data is given by [Jeffrey Kluger](#).

Using data such as these, we can 'track' the changes in land cover over time.

For example, data produced by Google and Dr. Matthew Hansen at the University of Maryland which shows [global maps of forest change \(2000-2019\)](#) using Landsat data (see [Science article](#)), with red showing loss in 2019 through to yellow for the year 2000 (using pseudocolour), or with 2019 highlighted in aqua blue if you are viewing the default 'Forest Loss Year (2019 Highlight) product.



The purpose of this practical is for you to perform and test a land cover classification over this area, using data from two dates (1992 and 2001). The visualisations above show that there has been significant change *since* 2001 (and before 1992).

We will be doing this using separate classifications of the two image dates, but you should be thinking throughout about whether this is an appropriate method, and what else you might consider (especially if you had a long time series of data such as those shown in the animations).


We will be doing a *supervised* classification here.

The steps you will undertake are:

- Examine the data and explore the spectral characteristics
- Define a series of Regions of Interest (ROIs) describing the classes you wish to extract
- Perform the classification
- Test the result

First, obtain and then load the TM and ETM images of Rondonia noted above, along with the SRTM DEM file.

View the ETM image as a FCC.



You may need to edit the image file to associate the DEM data correctly. To do this, look under ‘Raster Management’ in the Toolbox, and edit the ENVI header (for the Landsat data). You should then edit the header attributes to associate the DEM with the image data.

[]:



2.6 Classification using ENVI 5.5

2.6.1 Aims

After completing this practical, you should be able to analyse one or more image datasets using classification methods. This includes learning to identify land cover classes in a dataset and consider class separability (using histograms, scatterplots and other tools), and applying and assessing a classification product using Envi.

2.6.2 Introduction

The datasets you need for this practical are available from the **`Classification data (Rondônia) folder on moodle <>`** or you can download them individually:

<https://moodle.ucl.ac.uk/mod/folder/view.php?id=2749790> * ETM-110801 * ETM-110801.HDR * TM-250792 * TM-250792.HDR

With accompanying elevation data * SRTM-2002 * SRTM-2002.HDR

You should download these data and put them in a directory (folder) that you will remember!

See [Classification Introduction](#) for more details on the context and datasets.

2.6.3 Examination of the data

Load up the two images and examine the data. Try to identify the various classes you might like to obtain for this exercise decide how you can identify them. Examine feature space plots (scatter plots) to help you decide what may be feasible (and what may not). **You may decide that transformations of the data (e.g. band ratios or Principal Components) might aid your ability (and the computer's ability) to discriminate between classes**, but you should simply explore the data to start with.

Some examples of the various classes you might consider (shown on a standard False Colour Composite (FCC) image):

Class	Notes	Example
Urban	May also include other ‘built’ structures such as roads. You should be able to recognise these from their spatial structure, even at this resolution	
Forest	This should be easy to spot, but there are sometime clear ‘shading’ effects (as in this example) that might complicate classification	
Rock	Rocks are quite easily identifiable in the FCC images. You would generally expect them to be static between the two dates.	
River	There are rivers and other water bodies in the scene, which you will be able to recognise by their shape. They will be difficult to use as training sites as they are quite narrow at this resolution.	
Farmland	You will see a broad patchwork of areas that have been cleared of forest and used to graze cattle or raise crops. The areas are quite easy to spot in the FCC images, but might represent a broad spectral class because of the various physical cover types involved	
Other	You may spot some areas that have rather different spectral properties to most of the other areas. One example is shown here of field-shaped areas (green and purple areas) that might be inferred to be farmland, but are clearly different spectrally to other areas of farmland. We cannot really determine what these areas are from the information available, so you might require an ‘other’ class to cope with such eventualities.	
Cloud	The images may contain a small amount of cloud or smoke/haze, an example of which is shown here. They are quite easy to recognise visually in the FCC, but may be difficult to classify unless they are quite thick. If there are any thick clouds, you may see cloud shadows on the ground as well.	

You may make use of Google Maps to explore detail of the areas, e.g., if you zoom in to the ‘rock’ area, you will find it is actually more complex than just ‘bare rock’:

When deciding which classes may be appropriate to use, you should make use of your understanding of histograms and scatterplots, and use these to help explore the image information content.

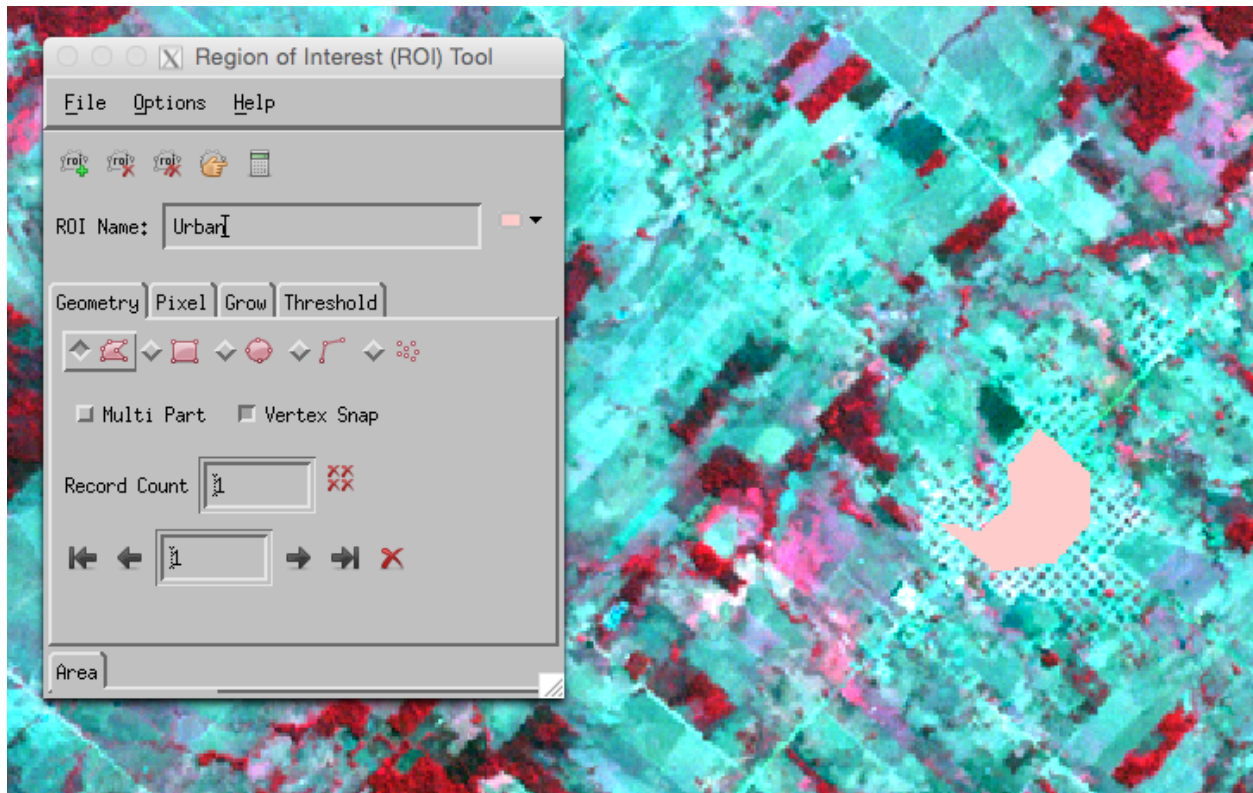
2.6.4 Defining spectral classes

In order to classify the image data you are required to define a set of “signatures” which represent each class. These are then used to “train” the classification algorithm.

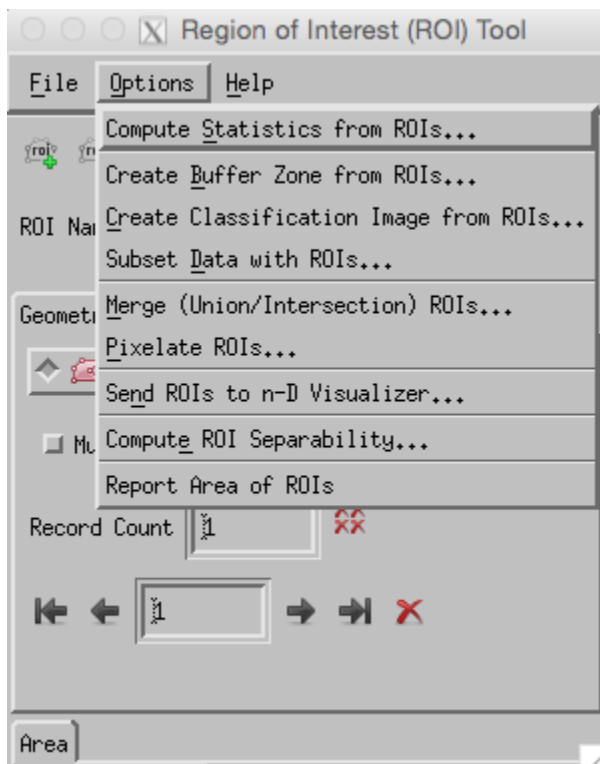
In envi, you need to define these classes via ROIs (Regions of Interest). Select the ROI tool:




and outline an ROI you want to define with the tool:



You may find the 'N-D visualiser' useful when doing this:



If you select only 2 bands to view, you will see information similar to the scatterplot (i.e. 2-dimensional).

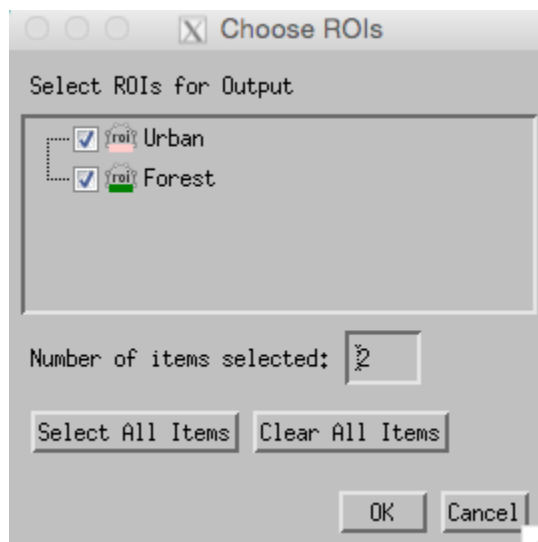


In such a view, you can readily ‘see’ how separable the classes might be.

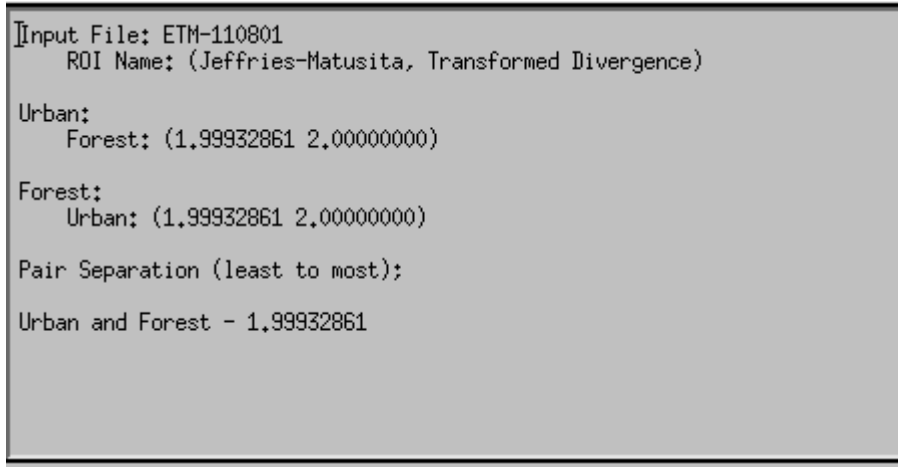
In higher dimensions, the visualiser ‘rotates’ the view so you can get different perspectives on the classes

Note that you will want to create an ROI for each class you are interested in, but that you can ‘merge’ (or delete) classes once you have created them.

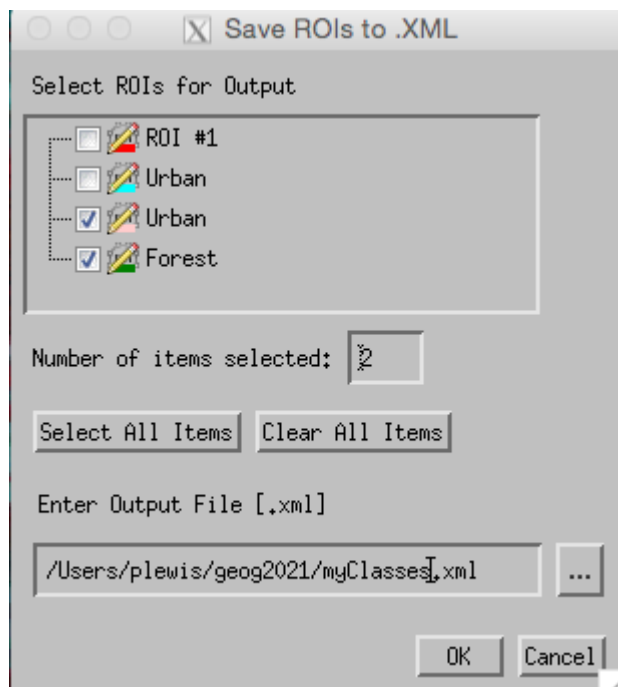
When you think you have a suitable set of ROIs, check the class separability:



This outputs Divergence metrics between the classes you have defined. These values range between 0 and 2.0. As a guide to interpretation, values greater than 1.9 indicate good separability of classes. If class separability is less than this, you might consider splitting the classes for the classification and recombining them post-classification (e.g. have two classes: forest1 and forest2).

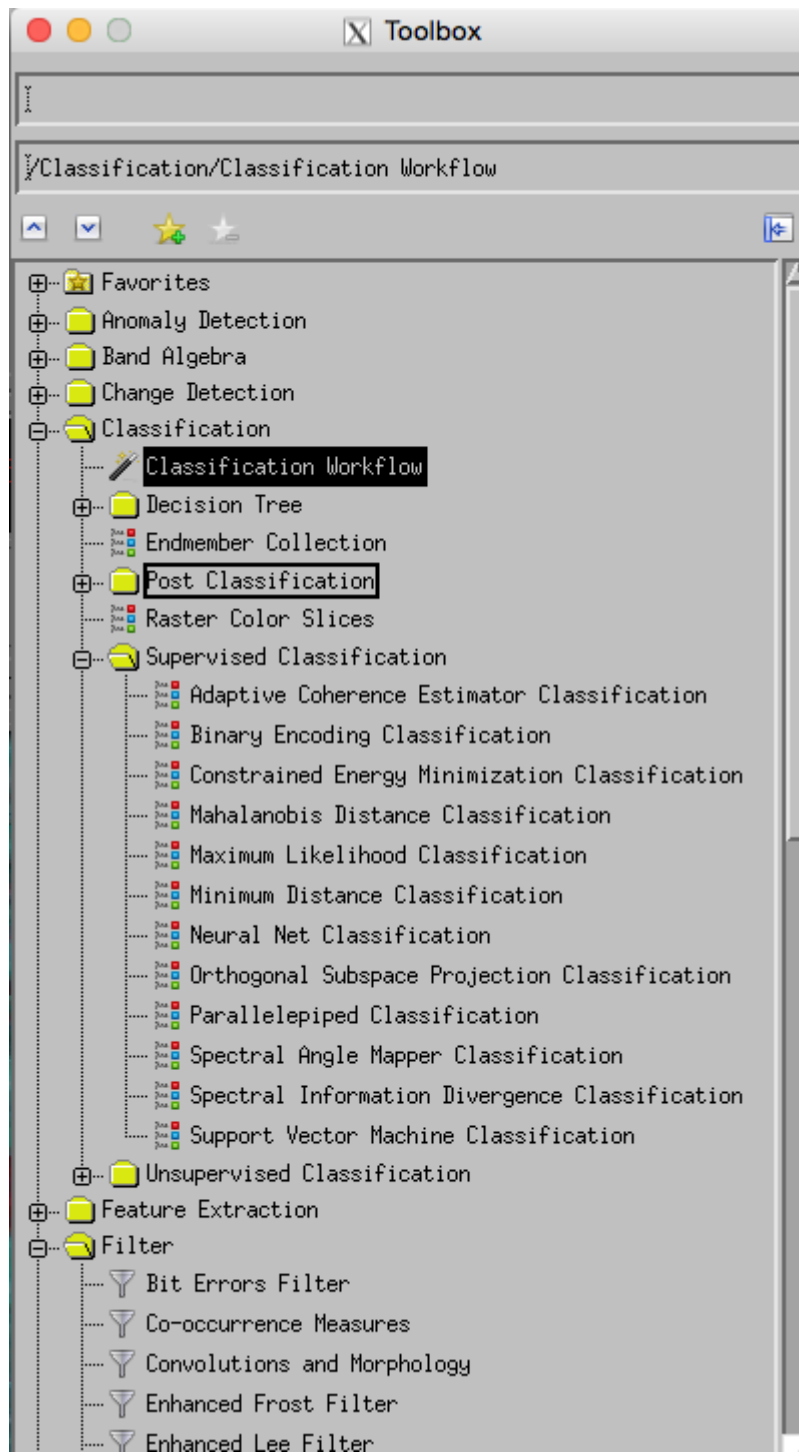


Then, make sure you save them (to xml format):



2.6.5 Image Classification

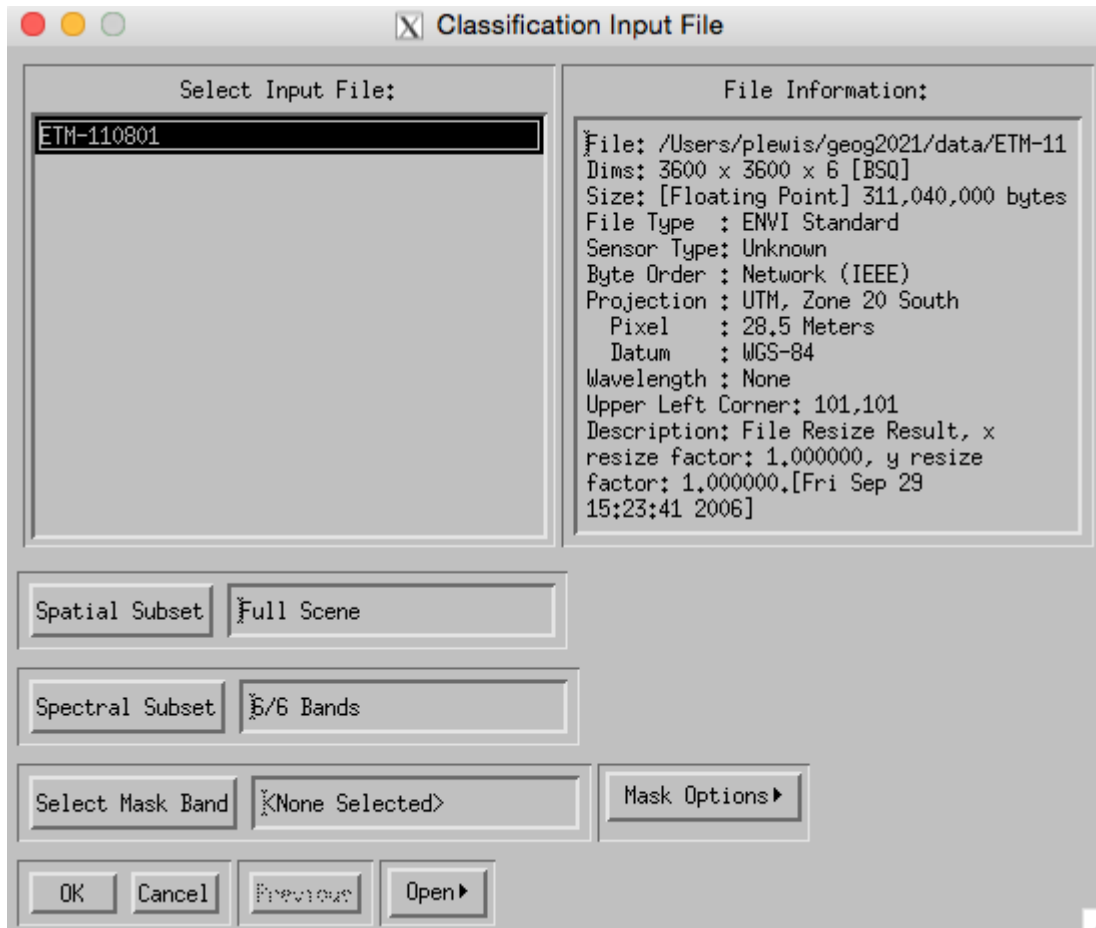
To perform a classification, first look at the options in the Toolbox:



As a first attempt, try the [Maximum Likelihood](#) classifier.

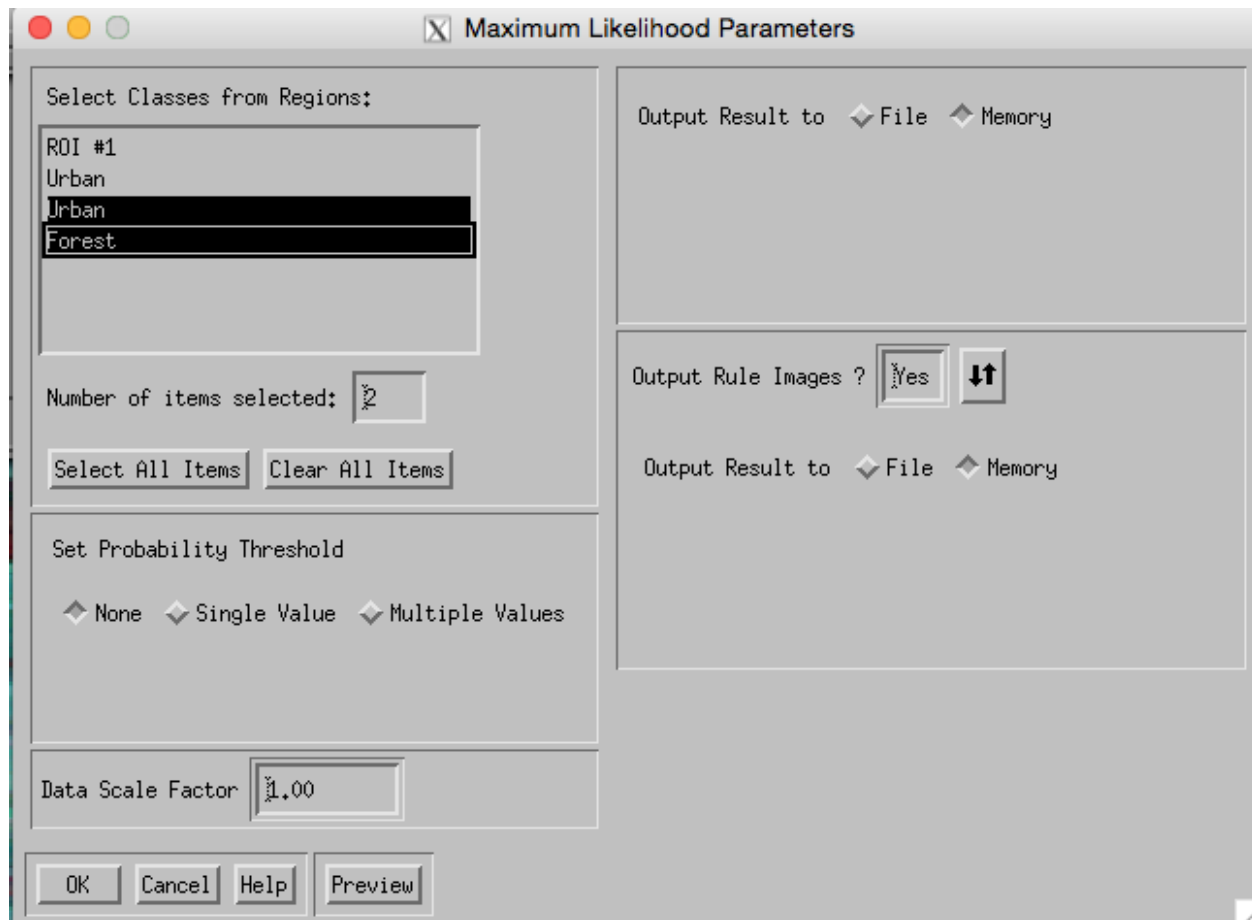
A [Tutorial](#) is available that will take you through some of the other options.

For the [Maximum Likelihood](#) classifier, select this item from the Toolbox:




and perform any subsetting or masking that you might require.

Then, select the Classes you want from the ROIs you have defined, along with making decisions about whether you want to save the result or not (if not, then just send it to ‘memory’, but it will not then be saved at the end of the session). If you do save the result, make sure you note down (in your notebook) what the file name was and what settings you used (e.g. which classes).




You should now have a classification result:



It is generally very instructive to visualise the ‘rule’ image associated with a result. This provides you with the reasoning the computer used to obtain the result it did.

For a method such as that used above, the training data are used to generate multivariate statistical distributions that we suppose to describe the full class. Each pixel then can be assigned a probability of class membership. The class which has the highest membership probability is usually assigned that class label.



What issues might occur if the probability of belonging to more than one class is very similar?

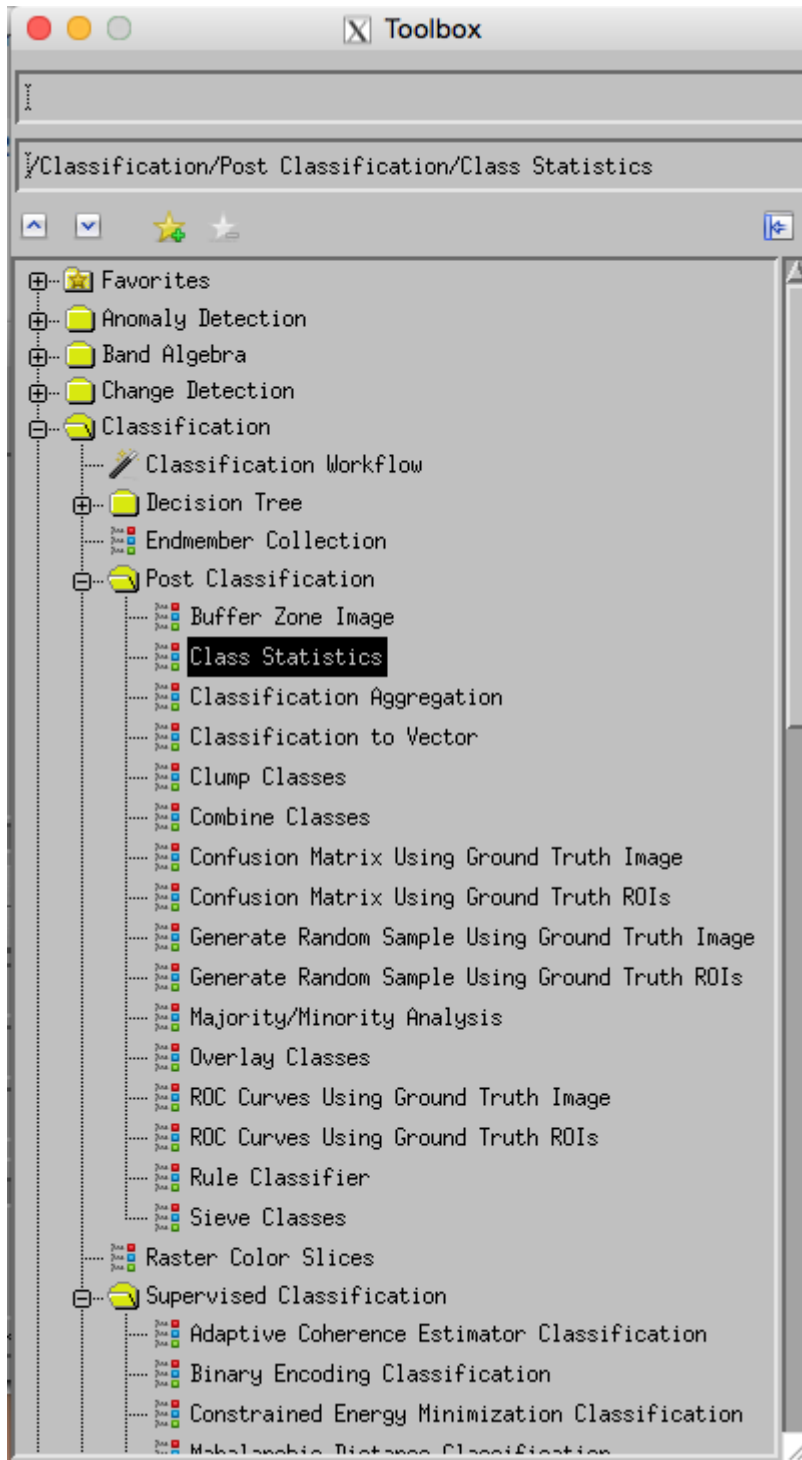
There appear to be topographic effects in the class probability images: why would that be so? and what might you do about it?

2.6.6 Accuracy Assessment

It is not very difficult to produce a classified map using earth observation data. You have now been through the process of supervised classification (using one method).

How can we tell how good this is though?

One thing you may wish to do is to examine the post-classification class statistics:



There are various other options that you may find useful to explore in the Post Classification section of the toolbox.

A vital part of the classification process though is an assessment of classification accuracy.

This is generally done as a [confusion matrix](#).

In setting this up, you need either to have a ground truth 'image', or a set of ROIs that can be used for ground truth.

You should first generate a new (independent) set of ROIs (or better still, use random samples) for your classes. If you

use random samples, you can check what you think the land cover class should be using Google Earth/Maps as above. Once you have your confusion matrix, make sure that you understand what it is telling you (and as far as possible, why that is so).

If the classification result seems poor, you can go back and edit your settings or class definitions and re-try, but try to keep the ROIs you use for checking independent of this process.

Make sure you understand the terms we use to describe the different accuracies shown in the confusion matrix, and also [what a kappa coefficient is](#).

2.6.7 Further Work

In this practical, you have gone through the process of performing an image classification and assessing its accuracy.

To finish the practical, you should classify *both* of the Landsat datasets you have, and calculate the change in forest area between the two dates. Since you have an accuracy assessment, it should be feasible for you to put an uncertainty on that estimate of change.

2.6.8 Summary

The main aim of this practical is to reinforce your understanding of the classification process and for you to gain practical experience at this.

It would be worthwhile exploring some of the options you have available (e.g. try some different classifiers).

Since there is quite a lot of ‘button clicking’ in this exercise, *make sure* that you understand what you are doing and why you are getting the result you do – there is very little value in the exercise otherwise!

If you have questions, ask the staff!

If you are very interested in change detection, you could explore the [change detection options in ENVI](#).

2.7 Coursework

2.7.1 Coursenotes

The material for the **associated coursework** is available via:

- [readthedocs](#)
- [github](#)
- [binder](#)
- [RStudio in binder](#)

SEARCH

- search