QGIS Training

1 Getting Started with QGIS

1.1 Software Installation

QGIS is a free GIS software. You can download the latest version here (select the supported operating system for your desktop/laptop): http://www.qgis.org/en/site/forusers/download.html

For Windows user, once you download the installer, please install the software as usual and follow the instructions.
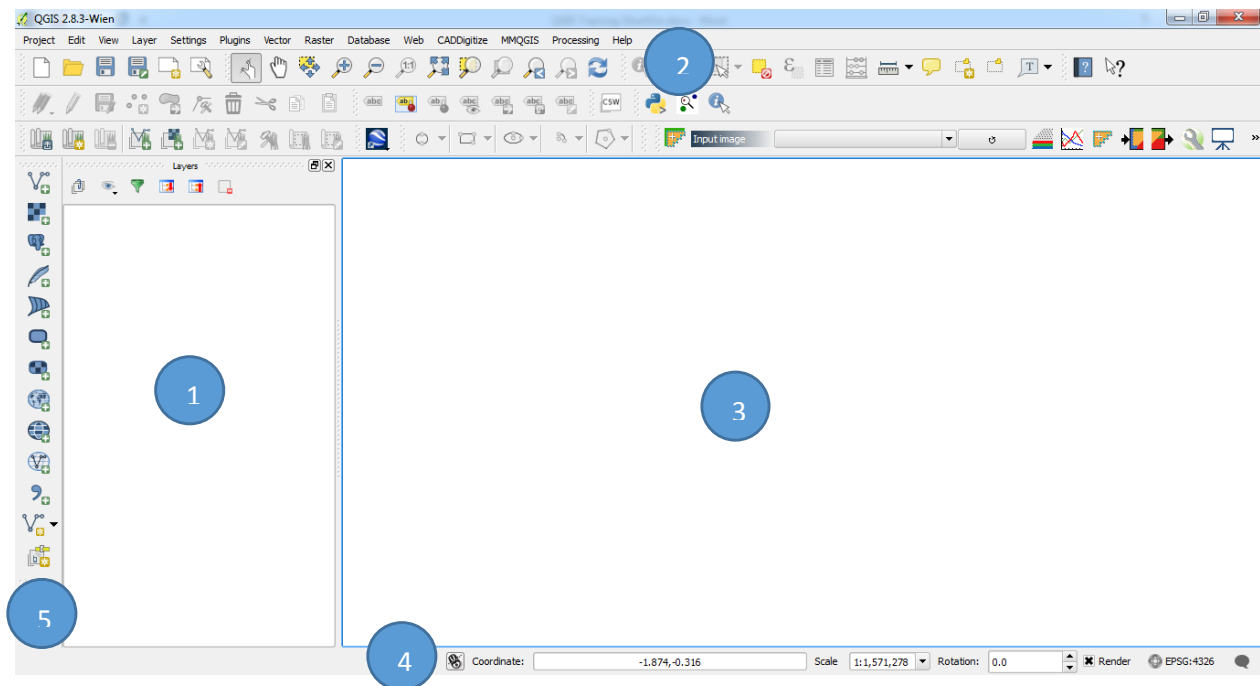
1.2 Preparing Exercise Data

The sample data provided with the Training Manual refers to the town of Swellendam and its surroundings. Swellendam is located about 2 hours' east of Cape Town in the Western Cape of South Africa. The dataset contains feature names in both English and Afrikaans. Anyone can use this dataset without difficulty, but you may prefer to use data from your own country or home town.

1.3 An Overview of the Interface

We will explore the QGIS user interface so that you are familiar with the menus, toolbars, map canvas and layers list that form the basic structure of the interface.

The goal for this lesson: To understand the basics of the QGIS user interface.

- Open QGIS. You will have a new, blank map.

The elements identified in the figure above are:

1. Layers List / Browser Panel
2. Toolbars
3. Map canvas
4. Status bar
5. Side Toolbar

The Layers List

In the Layers list, you can see a list, at any time, of all the layers available to you.

The Browser Panel

The QGIS Browser is a panel in QGIS that lets you easily navigate in your database. You can have access to common vector files (e.g. ESRI shapefile or MapInfo files), databases (e.g.PostGIS, Oracle, Spatialite or MSSQL Spatial) and WMS/WFS connections. You can also view your GRASS data.

Toolbars

Your most oft-used sets of tools can be turned into toolbars for basic access. For example, the File toolbar allows you to save, load, print, and start a new project. You can easily customize the interface to see only the tools you use most often, adding or removing toolbars as necessary via the View → Toolbars menu.

Even if they are not visible in a toolbar, all of your tools will remain accessible via the menus. For example, if you remove the File toolbar (which contains the Save button), you can still save your map by clicking on the File menu and then clicking on Save.

The Map Canvas

This is where the map itself is displayed.

The Status Bar

Shows you information about the current map. Also allows you to adjust the map scale and see the mouse cursor's coordinates on the map.

Try to identify the four elements listed above on your own screen, without referring to the diagram above. See if you can identify their names and functions. You will become more familiar with these elements as you use them in the coming days.
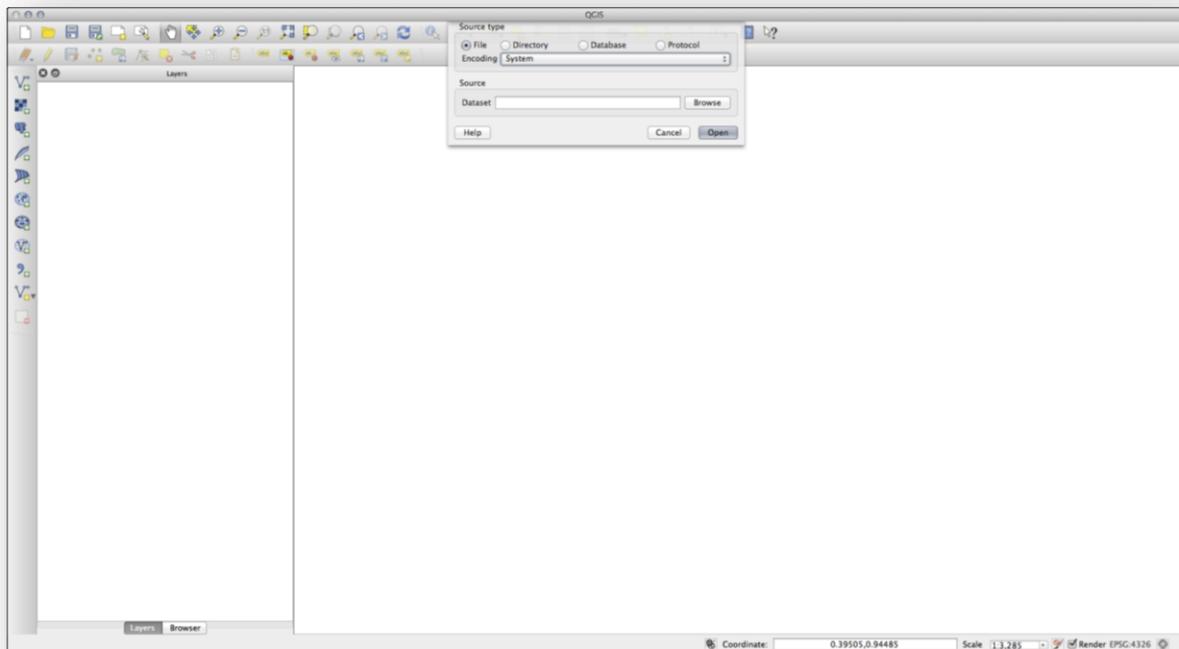
2 Working with Vector Data

Vector data is arguably the most common kind of data you will find in the daily use of GIS. It describes geographic data in terms of points that may be connected into lines and polygons. Every object in a vector dataset is called a feature, and is associated with data that describes that feature.

The goal for this lesson: To learn about the structure of vector data, and how to load vector datasets into a map.

2.1 Adding Your First Layer

We will start the application, and open a vector file to use for examples and exercises.

- Look for the Add Vector Layer button:

- Click on it to open the following dialog:

- Click on the Browse button and navigate to the file exercise_data/epsg4326/roads.shp (in your course directory). With this file selected, click Open. You will see the original dialog, but with the file path filled in. Click Open here as well. The data you specified will now load.

Congratulations! You now have a basic map. Now would be a good time to save your work.

- Click on the Save As button:

- Save the map under exercise_data/ and call it basic_map.qgs.

The standard vector file format used in QGIS is the ESRI shapefile. Support is provided by the OGR Simple Feature Library (http://www.gdal.org/ogr/).

A shapefile actually consists of several files. The following three are required:
1. .shp file containing the feature geometries
2. .dbf file containing the attributes in dBase format
3. .shx index file

Shapefiles also can include a file with a .prj suffix, which contains the projection information. While it is very useful to have a projection file, it is not mandatory. A shapefile dataset can contain additional files.

2.2   Viewing Layer Attributes

It's important to know that the data you will be working with does not only represent where objects are in space, but also tells you what those objects are.

From the previous exercise, you should have the roads layer loaded in your map. What you can see right now is merely the position of the roads.

To see all the data available to you, with the roads layer selected in the Layers panel:

- Click on this button:

It will show you a table with more data about the roads layer. This extra data is called attribute data. The lines that you can see on your map represent where the roads go; this is the spatial data.

These definitions are commonly used in GIS, so it's essential to remember them!

- You may now close the attribute table.

Vector data represents features in terms of points, lines and polygons on a coordinate plane. It is usually used to store discrete features, like roads and city blocks.

2.3   Loading Other Layers

QGIS uses the OGR library to read and write vector data formats, including ESRI shapefiles, MapInfo and Micro-Station file formats, AutoCAD DXF, PostGIS, SpatiaLite, Oracle Spatial and MSSQL Spatial databases, and many more. GRASS vector and PostgreSQL support is supplied by native QGIS data provider plugins. Vector data can also be loaded in read mode from zip and gzip archives into QGIS. As of the date of this document, 69 vector formats are supported by the OGR library (see OGR-SOFTWARE-SUITE in Literature and Web References). The complete list is available at http://www.gdal.org/ogr/ogr_formats.html.

2.3.1   Follow Along: Loading Vector Data from Shapefiles

The Shapefile is a specific file format that allows you to store GIS data in an associated group of files. Each layer consists of several files with the same name, but different file types. Shapefiles are easy to send back and forth, and most GIS software can read them.

Refer back to the previous exercise for instructions on how to add vector layers.

Load the data sets into your map following the same method:

- "places"
- "water"
- "rivers"
- "buildings"

2.3.2   Follow Along: Loading Vector Data from a Database

Databases allow you to store a large volume of associated data in one file. You may already be familiar with a database management system (DBMS) such as Microsoft Access. GIS applications can also make use of databases. GIS-specific DBMSes (such as PostGIS) have extra functions, because they need to handle spatial data.
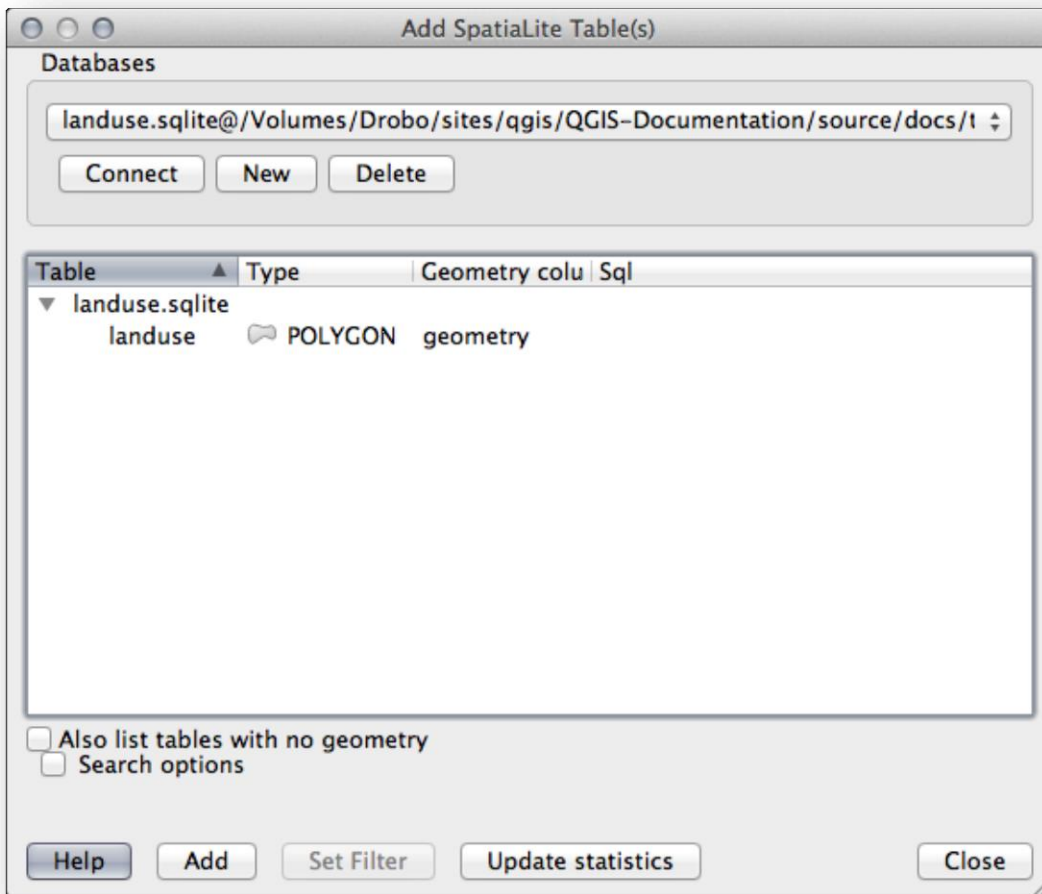
- Click on this icon:

(If you're sure you can't see it at all, check that the Manage Layers toolbar is enabled.)

It will give you a new dialog. In this dialog:

- Click the New button.
- In the same folder as the other data, you should find the file landuse.sqlite. Select it and click Open.

You will now see the first dialog again. Notice that the dropdown select above the three buttons now reads "land_use.db@...", followed by the path of the database file on your computer.

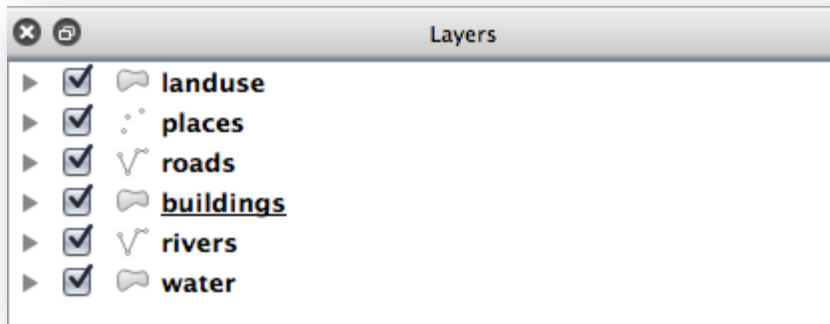- Click the Connect button. You should see this in the previously empty box:

- Click on the landuse layer to select it, then click Add

2.4 Reordering the Layers

The layers in your Layers list are drawn on the map in a certain order. The layer at the bottom of the list is drawn first, and the layer at the top is drawn last. By changing the order that they are shown on the list, you can change the order they are drawn in.

The order in which the layers have been loaded into the map is probably not logical at this stage. It's possible that the road layer is completely hidden because other layers are on top of it.

For example, this layer order...

... would result in roads and places being hidden as they run underneath urban areas.

To resolve this problem:

- Click and drag on a layer in the Layers list.
- Reorder them to look like this:

You'll see that the map now makes more sense visually, with roads and buildings appearing above the land use regions.

2.5   Creating Symbology for Vector Layer

The symbology of a layer is its visual appearance on the map. The basic strength of GIS over other ways of representing data with spatial aspects is that with GIS, you have a dynamic visual representation of the data you're working with.

Therefore, the visual appearance of the map (which depends on the symbology of the individual layers) is very important. The end user of the maps you produce will need to be able to easily see what the map represents. Equally as important, you need to be able to explore the data as you're working with it, and good symbology helps a lot.

In other words, having proper symbology is not a luxury or just nice to have. In fact, it's essential for you to use a GIS properly and produce maps and information that people will be able to use.

2.5.1   Follow Along: Changing Colors

To change a layer's symbology, open its Layer Properties. Let's begin by changing the color of the landuse layer.

- Right-click on the landuse layer in the Layers list.
- Select the menu item Properties in the menu that appears.

Note: By default, you can also access a layer's properties by double-clicking on the layer in the Layers list.

In the Properties window:

- Select the Style tab at the extreme left:

- Click the color select button next to the Color label.

A standard color dialog will appear.

- Choose a gray color and click OK.
- Click OK again in the Layer Properties window, and you will see the color change being applied to the layer.

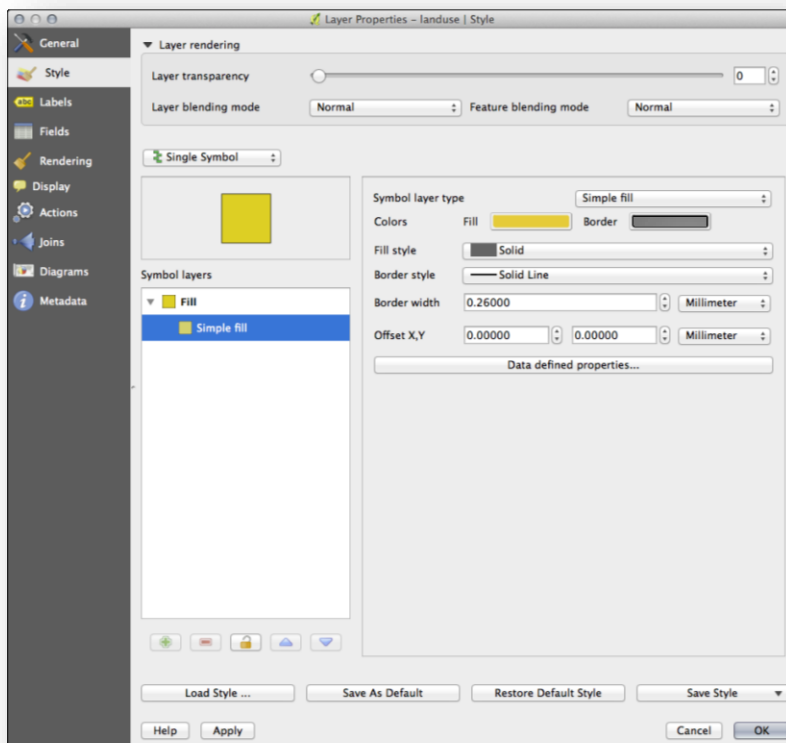**Try Yourself!**
Change the water layer to a light blue color.

2.5.2 Follow Along: Changing Symbol Structure

This is good stuff so far, but there's more to a layer's symbology than just its color. Next we want to eliminate the lines between the different land use areas so as to make the map less visually cluttered.

- Open the Layer Properties window for the landuse layer.

Under the Style tab, you will see the same kind of dialog as before. This time, however, you're doing more than just quickly changing the color.

- In the Symbol Layers panel, expand the Fill dropdown (if necessary) and select the Simple fill option:

- Click on the Border style dropdown. At the moment, it should be showing a short line and the words Solid Line.
- Change this to No Pen.
- Click OK.

Now the landuse layer won't have any lines between areas.

Try Yourself!
• Change the water layer's symbology again so that it is has a darker blue outline.
• Change the rivers layer's symbology to a sensible representation of waterways.

2.6   Classifying Vector Data

Classifying vector data allows you to assign different symbols to features (different objects in the same layer), depending on their attributes. This allows someone who uses the map to easily see the attributes of various features.

Up to now, none of the changes we have made to the map have been influenced by the objects that are being shown. In other words, all the land use areas look alike, and all the roads look alike. When looking at the map, the viewers don't know anything about the roads they are seeing; only that there is a road of a certain shape in a certain area.

But the whole strength of GIS is that all the objects that are visible on the map also have attributes. Maps in a GIS aren't just pictures. They represent not only objects in locations, but also information about those objects

2.6.1   Follow Along: Attribute data

The goal of this lesson: To explore the attribute data of an object and understand what the various data can be useful for.

Open the attribute table for the places layer (refer back to the section "Working with Vector Data" if necessary). Which field would be the most useful to represent in label form, and why?

In Conclusion
You now know how to use the attribute table to see what is actually in the data you're using. Any dataset will only be useful to you if it has the attributes that you care about. If you know which attributes you need, you can quickly decide if you're able to use a given dataset, or if you need to look for another one that has the required attribute data.

2.6.2   Follow Along: Using Labels

The goal for this lesson: To apply useful and good-looking labels to a layer.

Labels can be added to a map to show any information about an object. Any vector layer can have labels associated with it. These labels rely on the attribute data of a layer for their content.

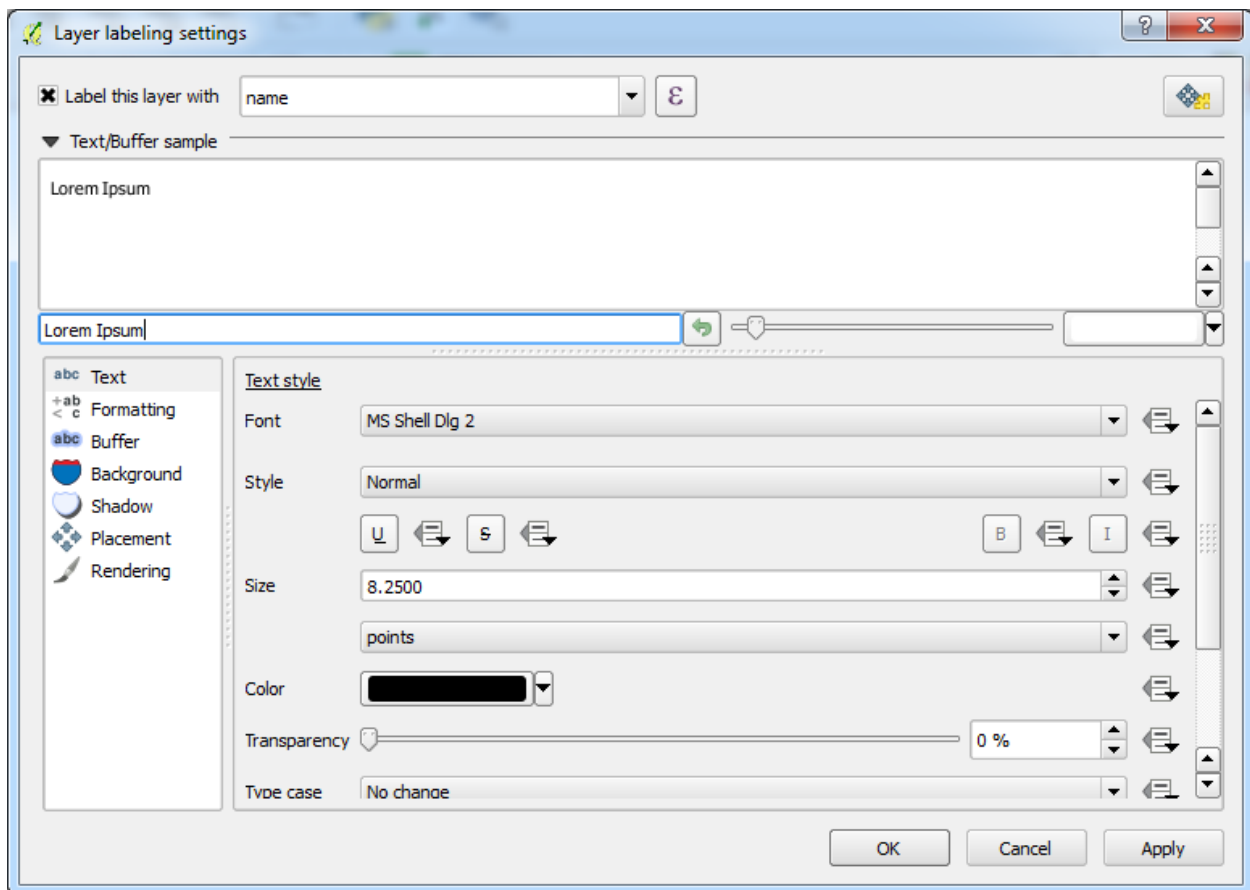Before being able to access the Label tool, you will need to ensure that it has been activated.

- Go to the menu item *View → Toolbars*.
- Ensure that the *Label* item has a check mark next to it. If it doesn't, click on the *Label* item, and it will be activated.
- Click on the *places* layer in the *Layers* list, so that it is highlighted.

- Click on the following toolbar button:

This gives you the *Layer labeling* settings dialog.
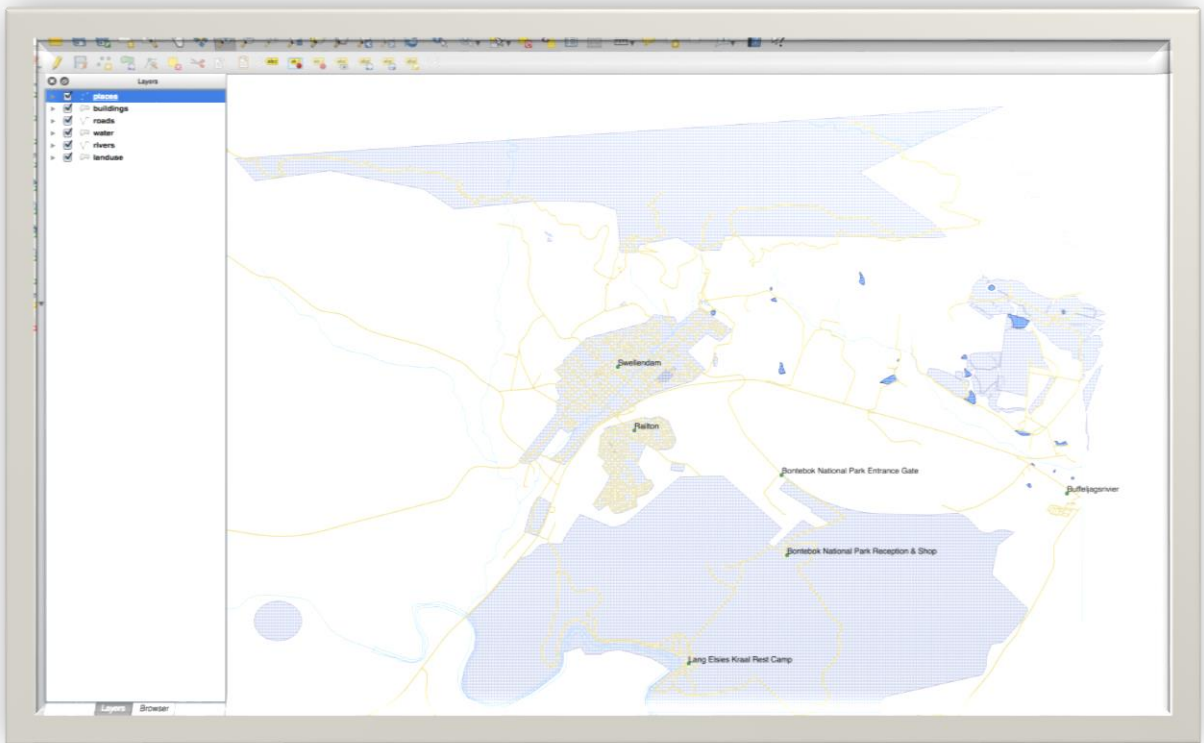
• Check the box next to *Label this layer with....*

You'll need to choose which field in the attributes will be used for the labels. In the previous lesson, you decided that the NAME field was the most suitable one for this purpose.

• Select *name* from the list:

- Click OK

The map should now have labels like this:

2.6.3 Follow Along: Changing Label Options

Depending on the styles you chose for your map in earlier lessons, you'll might find that the labels are not appropriately formatted and either overlap or are too far away from their point markers.

Let's look at the problem of the labels overlapping the points, but before we do that, let's take a look at the Buffer option.

- Open the *Label* tool again by clicking on its button as before.
- Select *Buffer* from the left-hand options list.
- Select the checkbox next to *Draw text buffer*, then choose options to match those shown here:

- Click *Apply.*

You'll see that this adds a colored buffer or border to the place labels, making them easier to pick out on the map.

Now we can address the positioning of the labels in relation to their point markers.

- In the Label tool dialog, go to the *Placement* tab.
- Change the value of *Distance* to 2mm and make sure that *Around point* is selected:

- Click *Apply*.

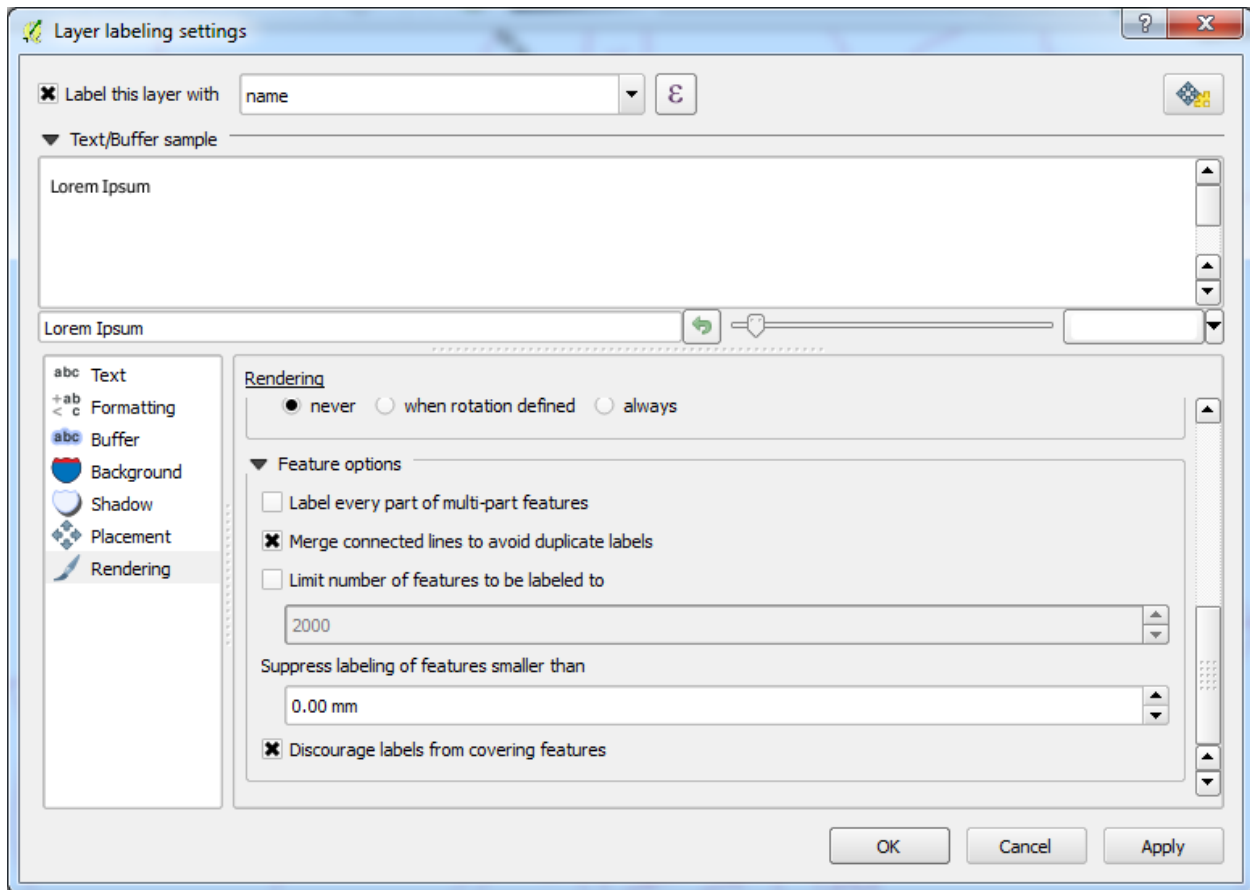You'll see that the labels are no longer overlapping their point markers.

2.6.4   Follow Along: Labeling Lines

Now that you know how labeling works, there's an additional problem. Points and polygons are easy to label, but what about lines?

We will now reformat the roads layer labels so that they are easy to understand.

- Hide the *Places* layer so that it doesn't distract you.
- Activate labels for the *streets* layer as before.

You'll see that some of the road names appear more than once and that's not always necessary. To prevent this from happening:

- In the *Label labelling settings* dialog, choose the *Rendering* option and select the *Merge connected lines to avoid duplicate labels*:

- Click *OK*

Another useful function is to prevent labels being drawn for features too short to be of notice.

- In the same *Rendering* panel, set the value of *Suppress labeling of features smaller than ...* to 5mm and note the results when you click *Apply*.

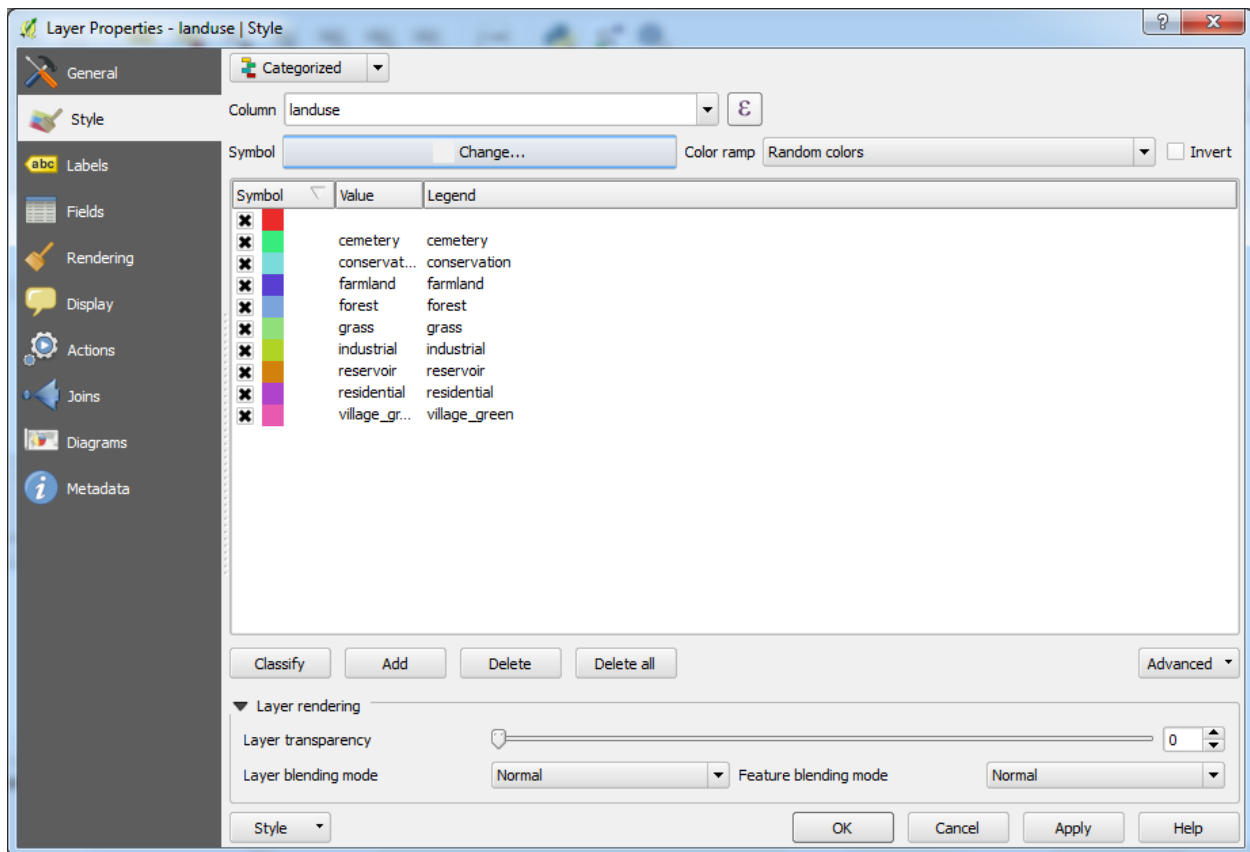Try out different *Placement* settings as well. Let's try the curved option.

- Select the *Curved* option in the *Placement* panel of the *Layer labeling settings* dialog.

2.6.5   Follow Along: Classifying Nominal Data

Labels are a good way to communicate information such as the names of individual places, but they can't be used for everything. Map's labeling sometime difficult to read and even overwhelming if there are numerous different names on the map.

The goal for this lesson: To learn how to classify vector data effectively.

- Open the *Layer Properties* dialog for the landuse layer.
- Go to the *Style* tab.
- Click on the dropdown that says *Single Symbol* and change it to *Categorized*.
- In the new panel, change the *Column* to *landuse*
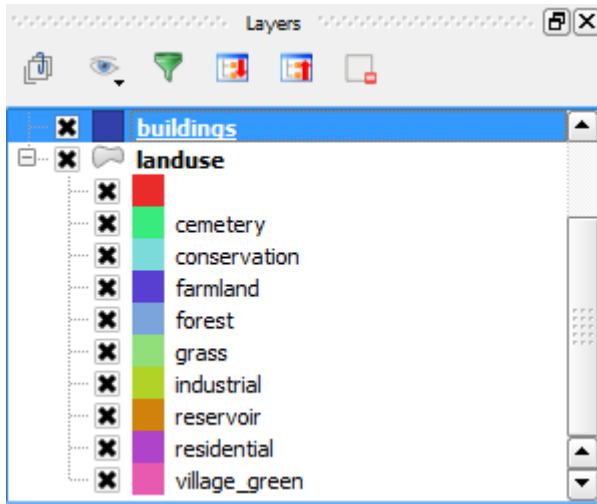- Click the button labeled *Classify*:

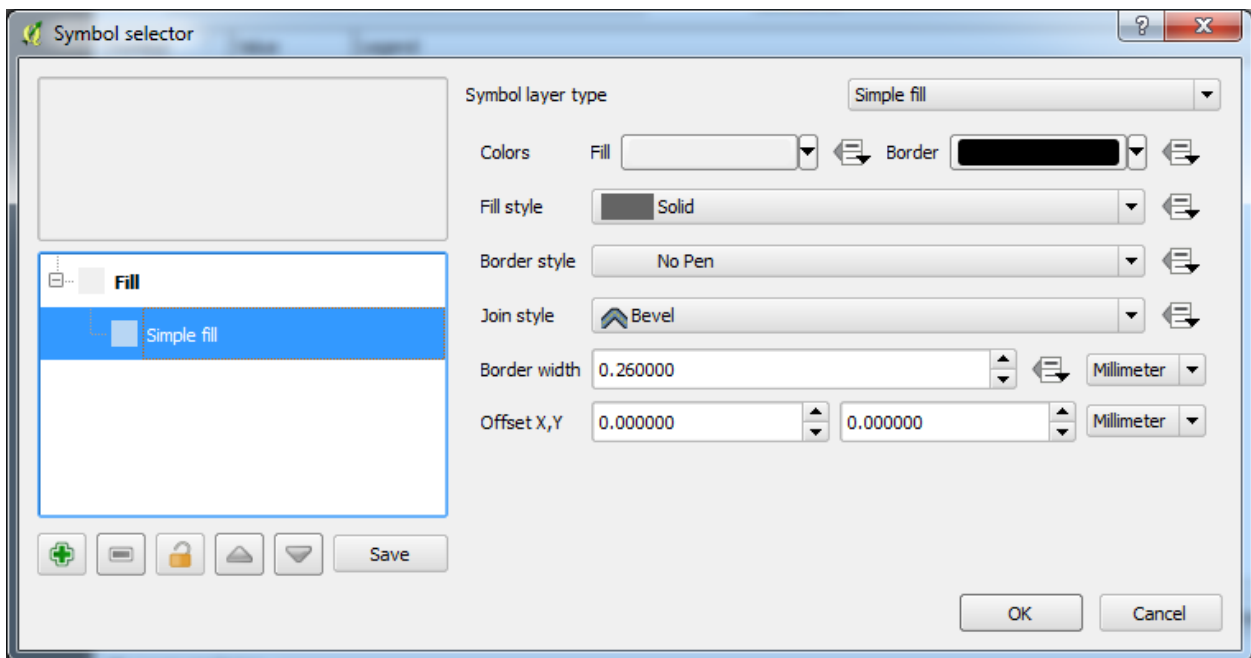- Click OK.

You'll see something like this:

- Click the arrow (or plus sign) next to landuse in the Layer list, you'll see the categories explained:

Now our landuse polygons are appropriately colored and are classified so that areas with the same land use are the same color. You may wish to remove the black border from the landuse layer:
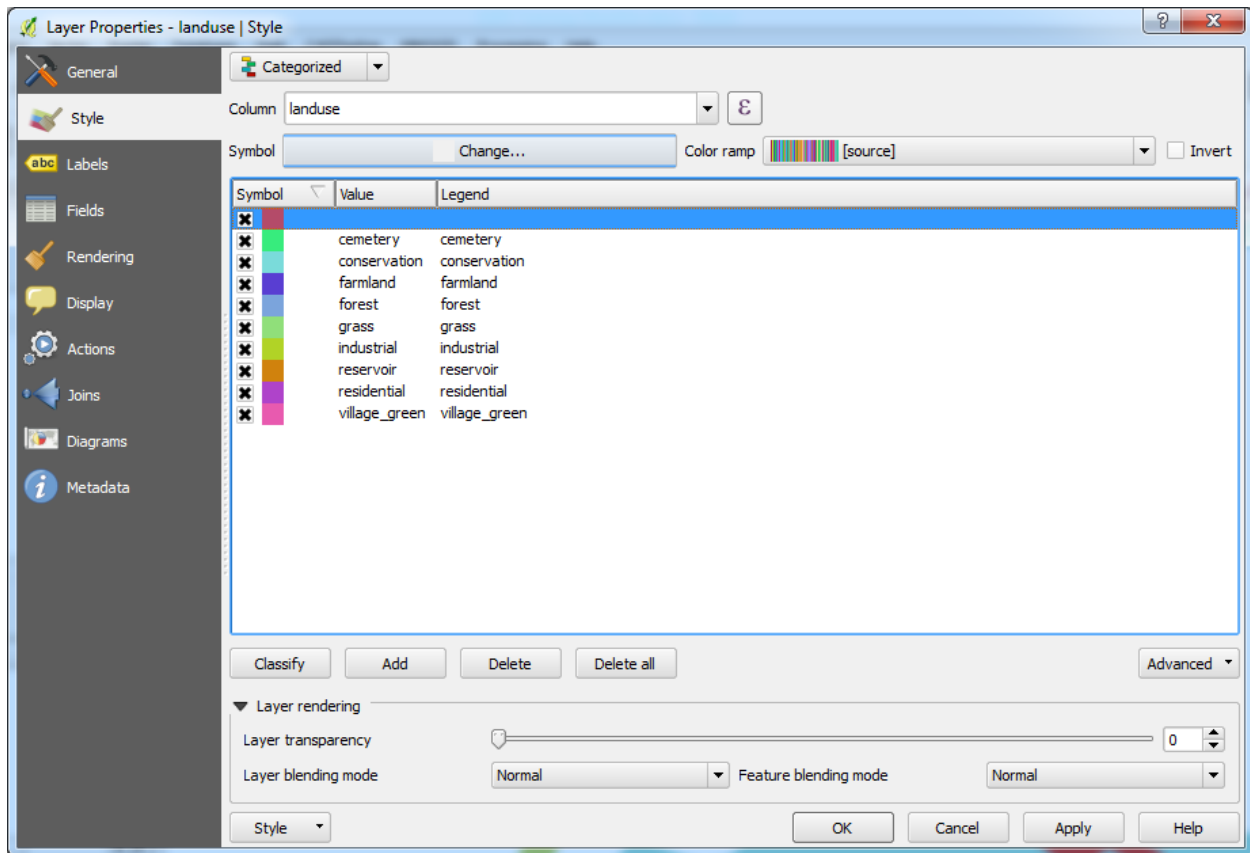
- Open *Layer Properties*, go to the *Style* tab and select *Symbol*.
- Change the symbol by changing the border to *Transparent* and click *OK*.

You'll see that the landuse polygon outlines have been removed, leaving just our new fill colours for each categorisation.

- If you wish to, you can change the fill color for each landuse area by double-clicking the relevant color block:

Notice that there is one category that's empty:

This empty category is used to color any objects which do not have a landuse value defined or which have a NULL value. It is important to keep this empty category so that areas with a NULL value are still represented on the map. You may like to change the color to more obviously represent a blank or NULL value.

2.7   Working with Table (Calculating Area)

Let's calculate the area of *buildings* layers.

- Open attribute table of *buildings* layer

- Enter edit mode by clicking this button:

- Add a new column with this button:
- Set up the dialog that appears, like this:

- Click OK.

The new field will be added (at the far right of the table; you may need to scroll horizontally to see it). However, at the moment it is not populated, it just has a lot of NULL values.

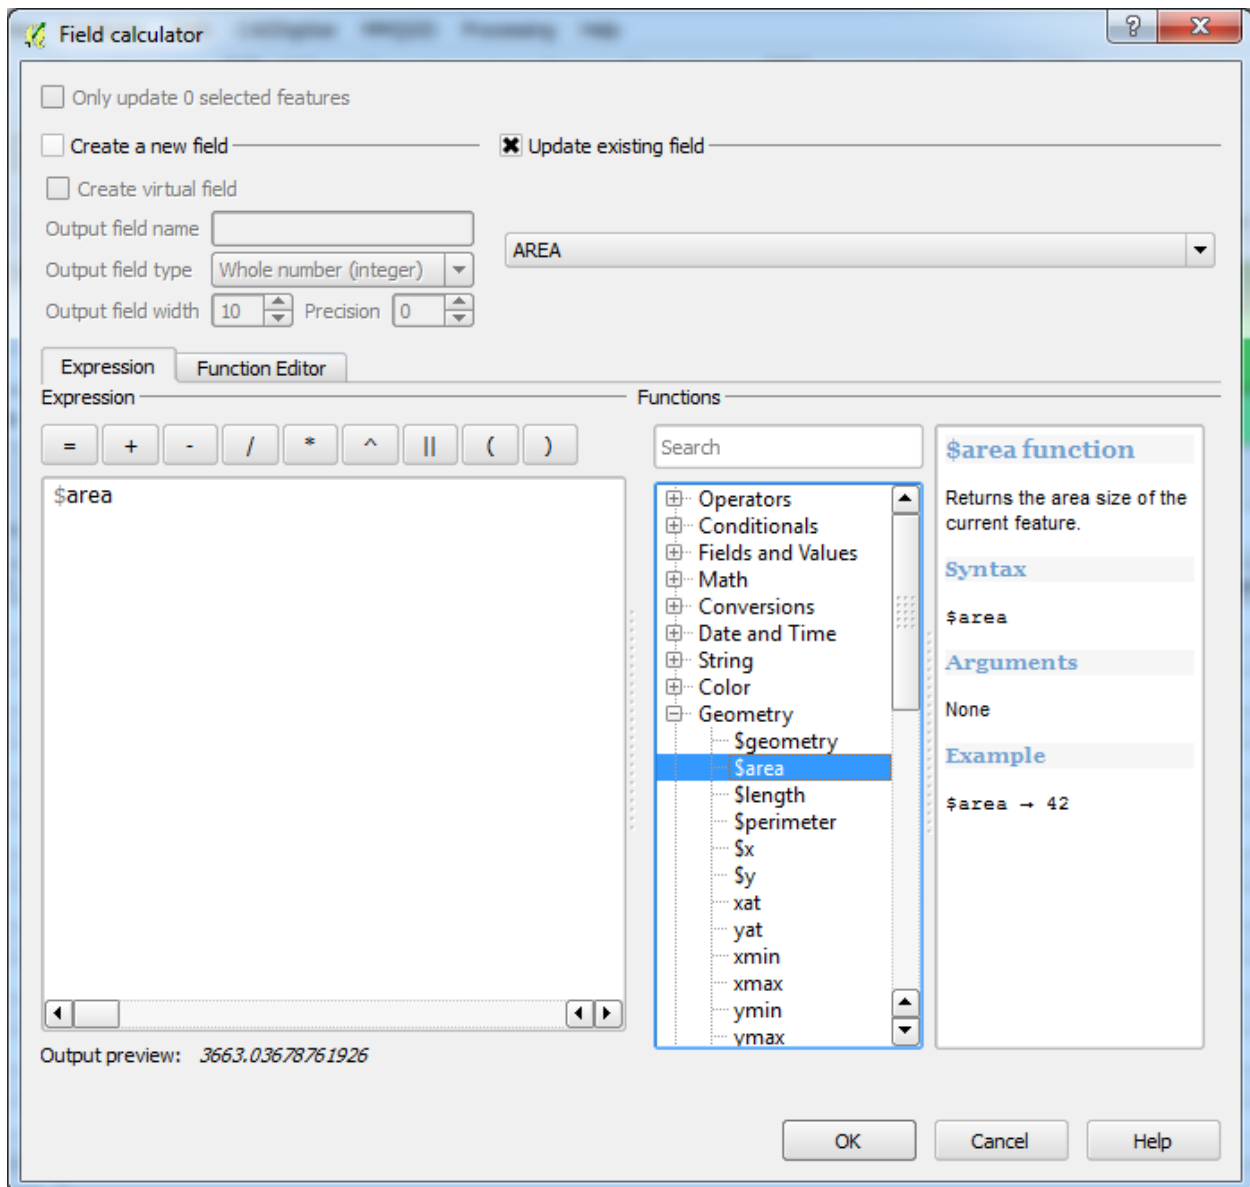To solve this problem, we'll need to calculate the areas.

- Open the field calculator:
- Click *Update existing field* and select *AREA*
- In the *Function List*, select *Geometry → $area*

Now your AREA field is populated with values (you may need to click the column header to refresh the data). Save the edits and click OK.

2.8   Reprojecting and Transforming Data

Let's talk about Coordinate Reference Systems (CRSs) again. We've touched on this briefly before, but haven't discussed what it means practically.

The goal for this lesson: To reproject and transform vector datasets.

2.8.1   Follow Along: Projections

The CRS that all the data as well as the map itself are in right now is called WGS84. This is a very common Geographic Coordinate System (GCS) for representing data. But there's a problem, as we will see.

- Then open the map of the world which you'll find under *exercise_data/world/continents.shp*.
- Zoom in to South Africa by using the *Zoom In* tool.
- Try setting a scale in the *Scale* field, which is in the *Status Bar* along the bottom of the screen. While over South Africa, set this value to *1:5000000* (one to five million).
- Pan around the map while keeping an eye on the *Scale* field.

Notice the scale changing? That's because you're moving away from the one point that you zoomed into at 1:5000000, which was at the center of your screen. All around that point, the scale is different.

To understand why, think about a globe of the Earth. It has lines running along it from North to South. These longitude lines are far apart at the equator, but they meet at the poles.

In a GCS, you're working on this sphere, but your screen is flat. When you try to represent the sphere on a flat surface, distortion occurs, similar to what would happen if you cut open a tennis ball and tried to flatten it out. What this means on a map is that the longitude lines stay equally far apart from each other, even at the poles (where they are supposed to meet). This means that, as you travel away from the equator on your map, the scale of the objects that you see gets larger and larger. What this means for us, practically, is that there is no constant scale on our map!

To solve this, let's use a Projected Coordinate System (PCS) instead. A PCS "projects" or converts the data in a way that makes allowance for the scale change and corrects it. Therefore, to keep the scale constant, we should reproject our data to use a PCS.

2.8.2 Follow Along: "On the Fly" Reprojection

QGIS allows you to reproject data "on the fly". What this means is that even if the data itself is in another CRS, QGIS can project it as if it were in a CRS of your choice.

- To enable "on the fly" projection, click on the *CRS Status* button in the *Status Bar* along the bottom of the QGIS window:
- In the dialog that appears, check the box next to Enable '*on the fly' CRS transformation*.
- Type the word *global* into the *Filter* field. One *CRS (NSIDC EASE-Grid Global)* should appear in the list below.
- Click on the *NSIDC EASE-Grid Global* to select it, then click *OK*.
- Notice how the shape of South Africa changes. All projections work by changing the apparent shapes of objects on Earth.
- Zoom in to a scale of 1:5000000 again, as before.
- Pan around the map.
- Notice how the scale stays the same!

"On the fly" reprojection is also used for combining datasets that are in different CRSs.

- Deactivate "on the fly" re-projection again:

  – Click on the *CRS Status* button again.

– Un-check the *Enable 'on the fly' CRS transformation* box.

– Clicking *OK*.

• In QGIS 2.0, the 'on the fly' reprojection is automatically activated when layers with different CRSs are loaded in the map. To understand what 'on the fly' reprojection does, deactivate this automatic setting:

– Go to *Settings → Options...*

– On the left panel of the dialog, select *CRS*.

– Un-check *Automatically enable 'on the fly' reprojection if layers have different CRS*.

– Click *OK*.

• Add another vector layer to your map which has the data for South Africa only. You'll find it as *exercise_data/world/RSA.shp*.

What do you notice?

The layer isn't visible! But that's easy to fix, right?

• Right-click on the RSA layer in the Layers list.
• Select Zoom to Layer Extent.

OK, so now we see South Africa... but where is the rest of the world?

It turns out that we can zoom between these two layers, but we can't ever see them at the same time. That's because their Coordinate Reference Systems are so different. The continents dataset is in degrees, but the RSA dataset is in meters. So, let's say that a given point in Cape Town in the RSA dataset is about 4 100 000 meters away from the equator. But in the continents dataset, that same point is about 33.9 degrees away from the equator.
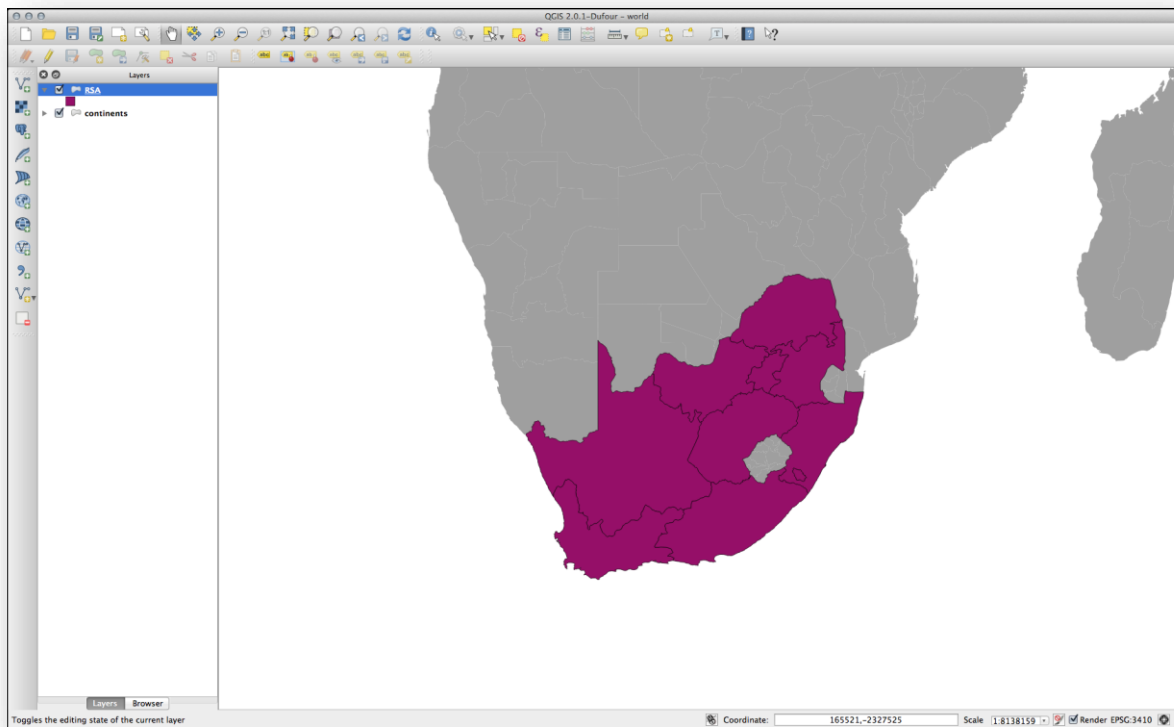
This is the same distance - but QGIS doesn't know that. You haven't told it to reproject the data. So as far as it's concerned, the version of South Africa that we see in the RSA dataset has Cape Town at the correct distance of 4 100 000 meters from the equator. But in the continents dataset, Cape Town is only 33.9 meters away from the equator! You can see why this is a problem.

QGIS doesn't know where Cape Town is supposed to be - that's what the data should be telling it. If the data tells QGIS that Cape Town is 34 meters away from the equator and that South Africa is only about 12 meters from north to south, then that is what QGIS will draw.

To correct this:

• Click on the CRS Status button again and switch Enable 'on the fly' CRS transformation on again as before.
• Zoom to the extents of the RSA dataset.

Now, because they're made to project in the same CRS, the two datasets fit perfectly:

When combining data from different sources, it's important to remember that they might not be in the same CRS. "On the fly" reprojection helps you to display them together.

Before you go on, you probably want to have the 'on the fly' reprojection to be automatically activated whenever you open datasets having different CRS:
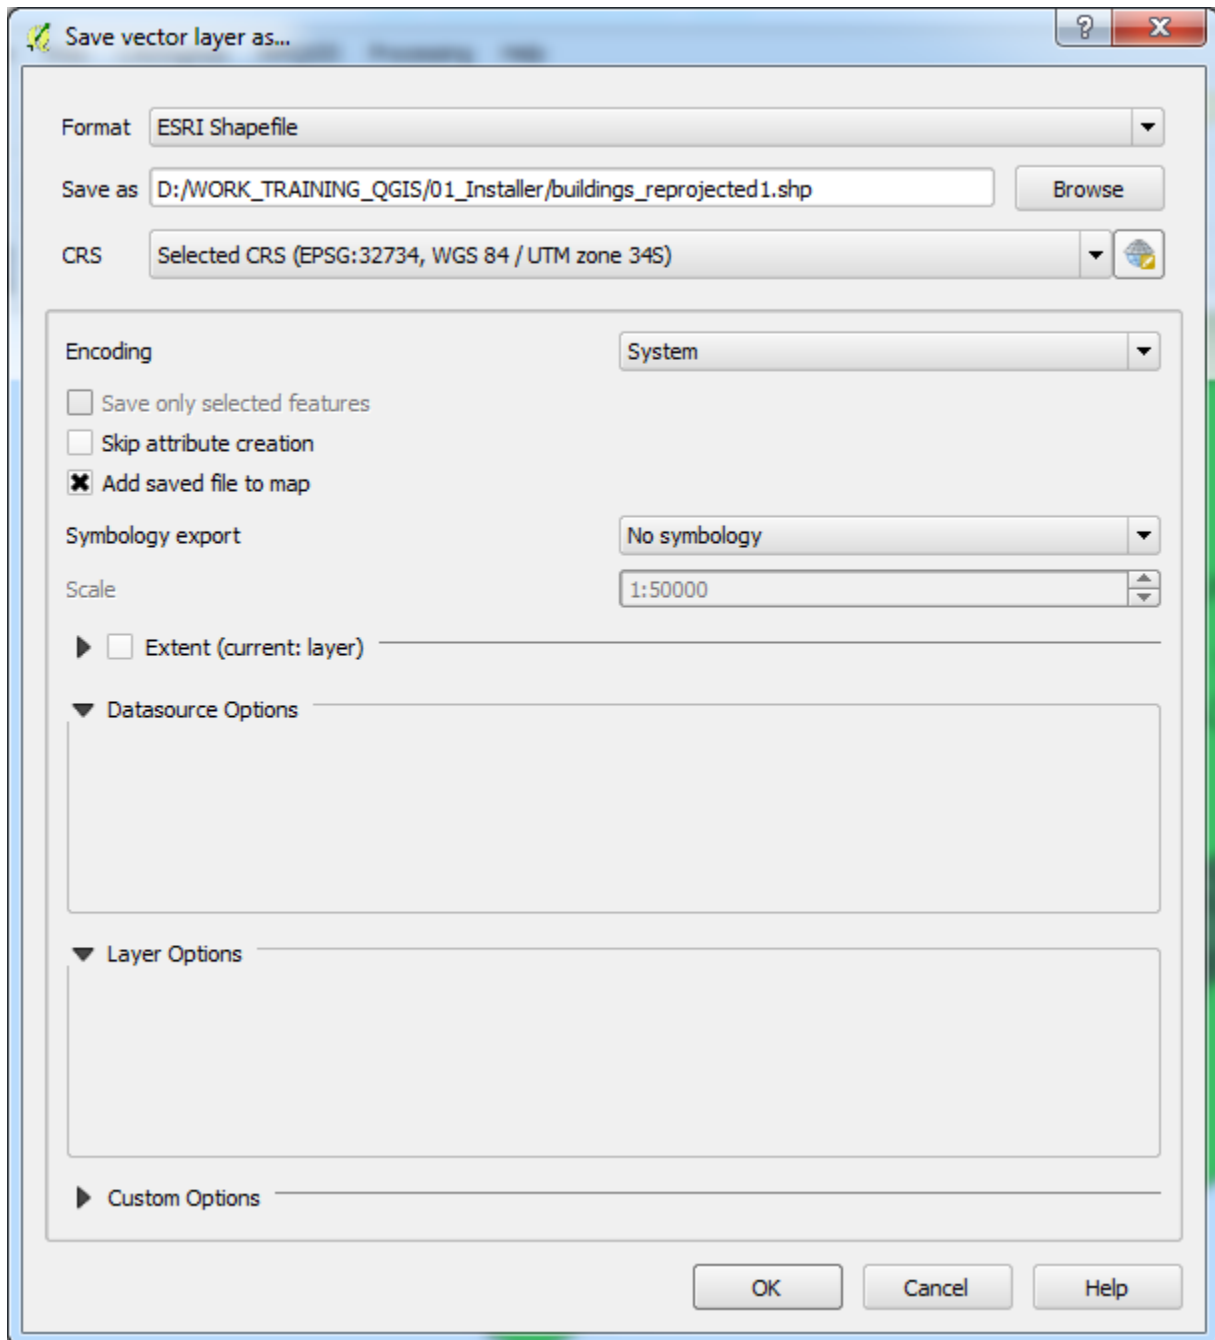
- Open again *Settings → Options...* and select *CRS*.
- Activate *Automatically enable 'on the fly' reprojection if layers have different CRS*.

2.8.3 Follow Along: Saving a Dataset to Another CRS

Let's change projection of buildings layer to another CRS.

- Right-click on the *buildings* layer in the *Layers* list.
- Select *Save As...* in the menu that appears. You will be shown the *Save vector layer as...* dialog.
- Click on the *Browse* button next to the *Save as field*.
- Navigate to *exercise_data/* and specify the name of the new layer as *buildings_reprojected.shp*.
- Leave the *Encoding* unchanged.
- Change the value of the *Layer CRS dropdown* to Selected CRS.
- Click the *Browse* button beneath the dropdown.
- The *CRS Selector* dialog will now appear.
- In its *Filter field*, search for *34S*.
- Choose *WGS 84 / UTM zone 34S* from the list.
- Leave the *Symbology* export unchanged.

The Save vector layer as... dialog now looks like this:

- Click *OK*.
- Start a new map and load the reprojected layer you just created.

3   Working with Raster Data

Raster data is quite different from vector data. Vector data has discrete features constructed out of vertices, and perhaps connected with lines and/or areas. Raster data, however, is like any image. Although it may portray various properties of objects in the real world, these objects don't exist as separate objects; rather, they are represented using pixels of various different color values.

The goal for this lesson: To learn how to work with raster data in the QGIS environment.

3.1   Follow Along: Loading Raster Data

- Click on the Load Raster Layer button:

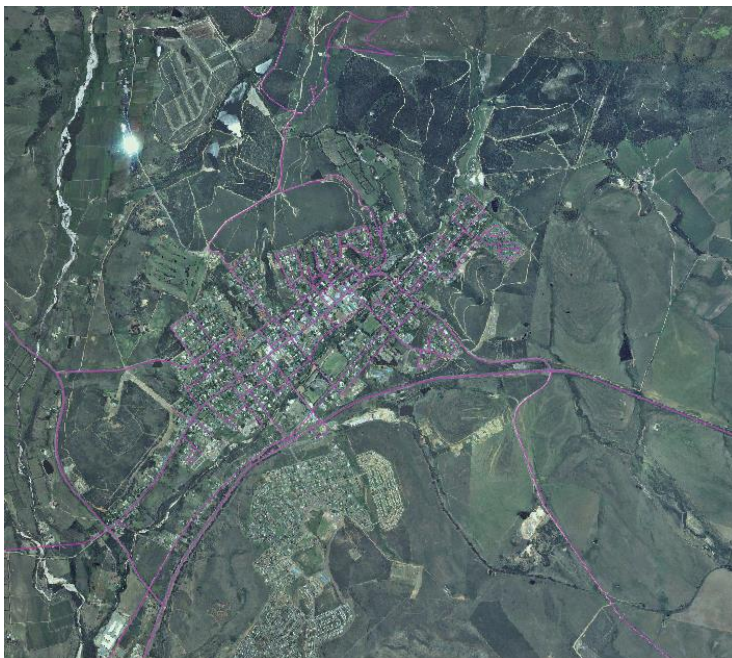The Load Raster Layer dialog will open. The data for this project is in exercise_data/raster.

- Either load them all in separately, or hold down ctrl and click on all four of them in turn, then open them at the same time.

The first thing you'll notice is that nothing seems to be happening in your map. Are the rasters not loading? Well, there they are in the Layers list, so obviously they did load. The problem is that they're not in the same projection.

Luckily, we've already seen what to do in this situation.

- Select Project –> Project Properties in the menu:
- Select CRS tab in the menu:
- Enable "on the fly" reprojection.
- Set it to the same projection as the rest of your data (WGS 84 / UTM zone 33S).
- Click OK.

The rasters should fit nicely:

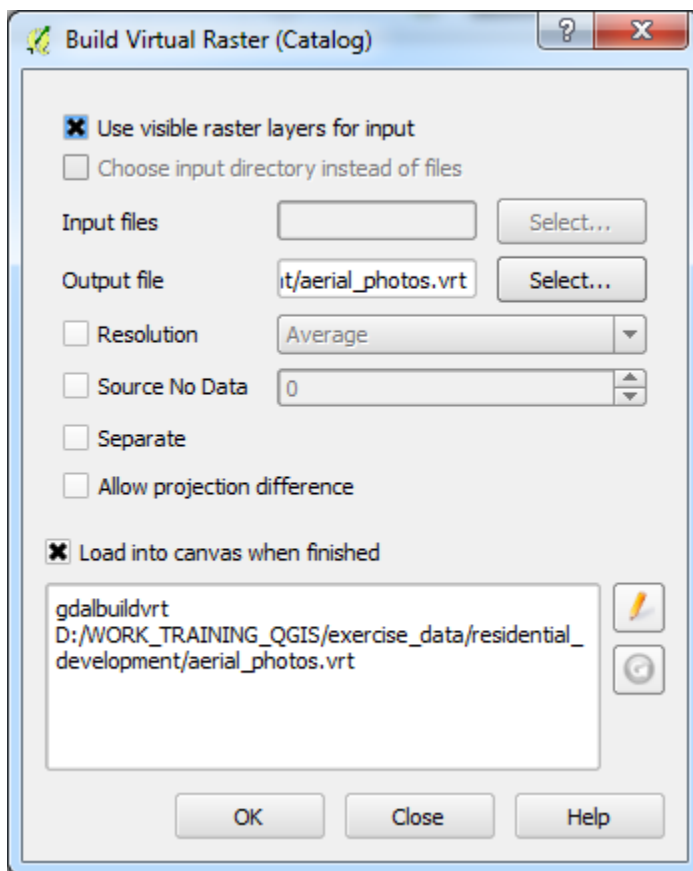## 3.2 Follow Along: Create a Virtual Raster

Now as you can see from this, your solution layer lies across all four photographs. What this means is that you're going to have to work with four rasters all the time. That's not ideal; it would be better to have one file for one (composite) image, right?

Luckily, QGIS allows you to do exactly this, and without needing to actually create a new raster file, which could take up a lot of space. Instead, you can create a Virtual Raster. This is also often called a Catalog, which explains its function. It's not really a new raster. Rather, it's a way to organize your existing rasters into one catalog: one file for easy access.

To make a catalog:

- Click on the menu item Raster → Miscellaneous →Build Virtual Raster (Catalog).
- In the dialog that appears, check the box next to Use visible raster layers for input.
- Enter the output location.
- Enter aerial_photos.vrt as the file name.
- Check the Load into canvas when finished button.

Notice the text field below. What this dialog is actually doing is that it's writing that text for you. It's a long command that QGIS is going to run.

- Click OK to run the command.

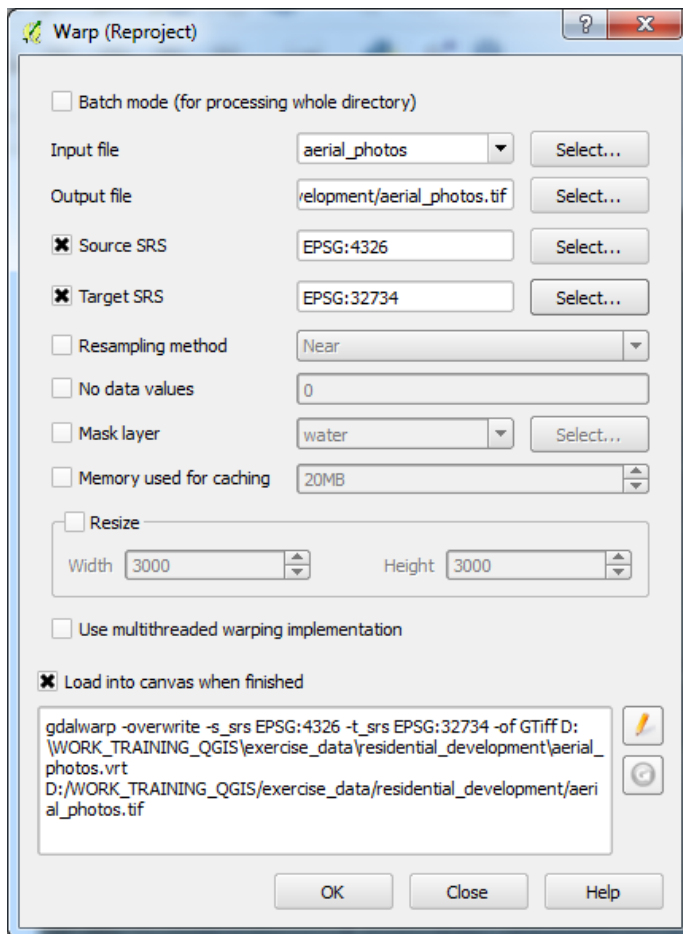It may take a while to complete. When it's done, it will tell you so with a message box.

- Click OK to chase the message away.
- Click Close on the Build Virtual Raster (Catalog) dialog. (Don't click OK again, otherwise it's going to start running that command again.)
- You can now remove the original four rasters from the Layers list.
- If necessary, click and drag the new aerial_photos raster catalog layer to the bottom of the Layers list so that the other activated layers become visible.

3.3   Follow Along: Transforming Raster Data

The above methods allow you to virtually merge datasets using a catalog, and to reproject them "on the fly". However, if you are setting up data that you'll be using for quite a while, it may be more efficient to create new rasters that are already merged and reprojected. This improves performance while using the rasters in a map, but it may take some time to set up initially.

- Click on the menu item *Raster → Projections → Warp (Reproject).*

Note that this tool features a handy batch option for reprojecting the contents of whole directories. You can also reproject virtual rasters (catalogs), as well as enabling a multithreaded processing mode.

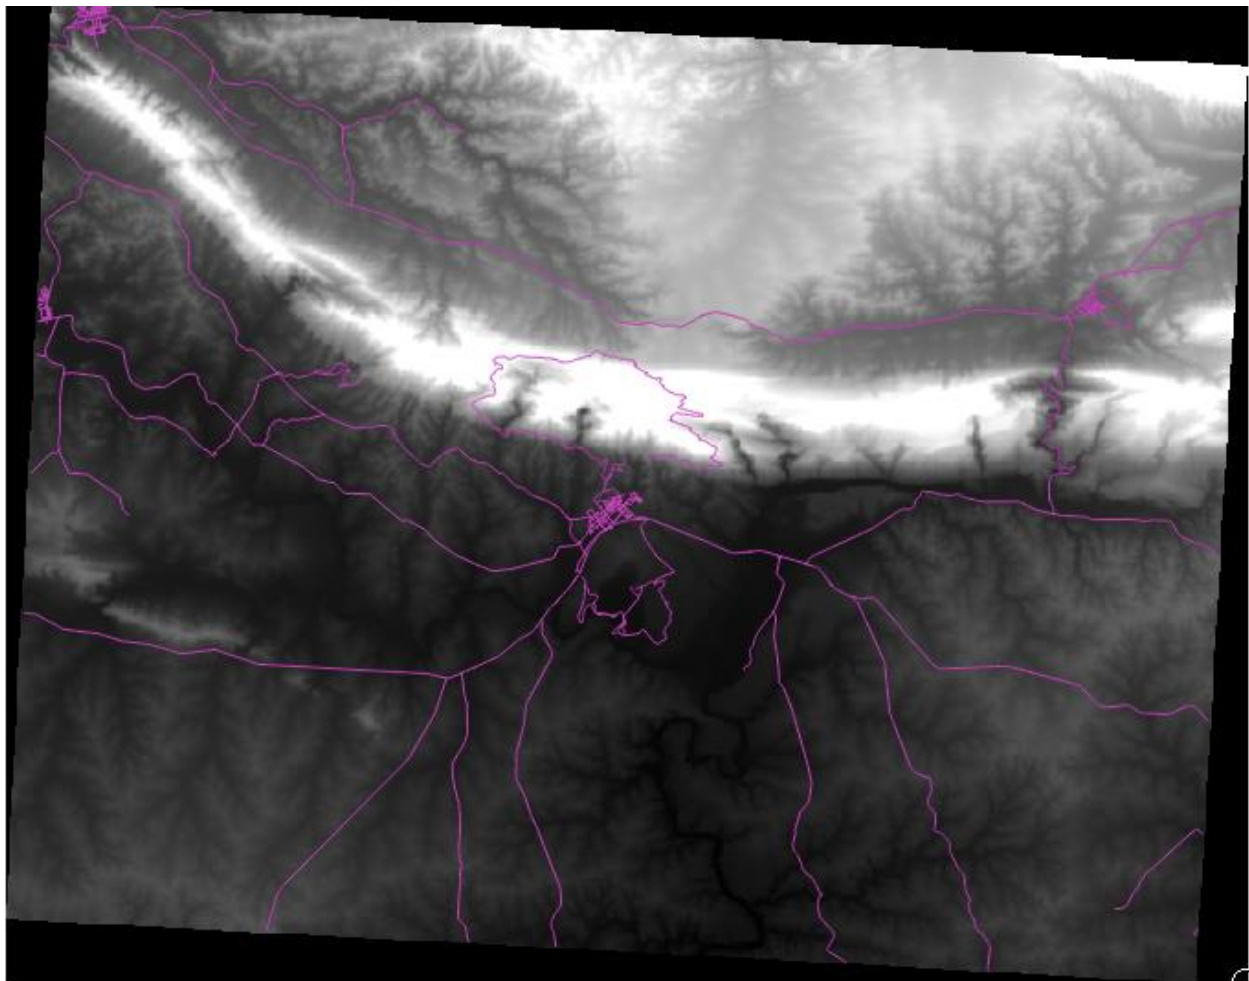## 3.4   Follow Along: Changing Raster Symbology

Not all raster data consists of aerial photographs. There are many other forms of raster data, and in many of those cases, it's essential to symbolize the data properly so that it becomes properly visible and useful.

The goal for this lesson: To change the symbology for a raster layer.

- Use the Add Raster Layer button to load the new raster dataset.
- Load the dataset srtm_41_19.tif, found under the directory exercise_data/raster/SRTM/.
- Once it appears in the Layers list, rename it to DEM.
- Zoom to the extent of this layer by right-clicking on it in the Layer List and selecting Zoom to Layer Extent.

This dataset is a Digital Elevation Model (DEM). It's a map of the elevation (altitude) of the terrain, allowing us to see where the mountains and valleys are, for example.

Once it's loaded, you'll notice that it's a basic stretched grayscale representation of the DEM. It's seen here with the vector layers on top:
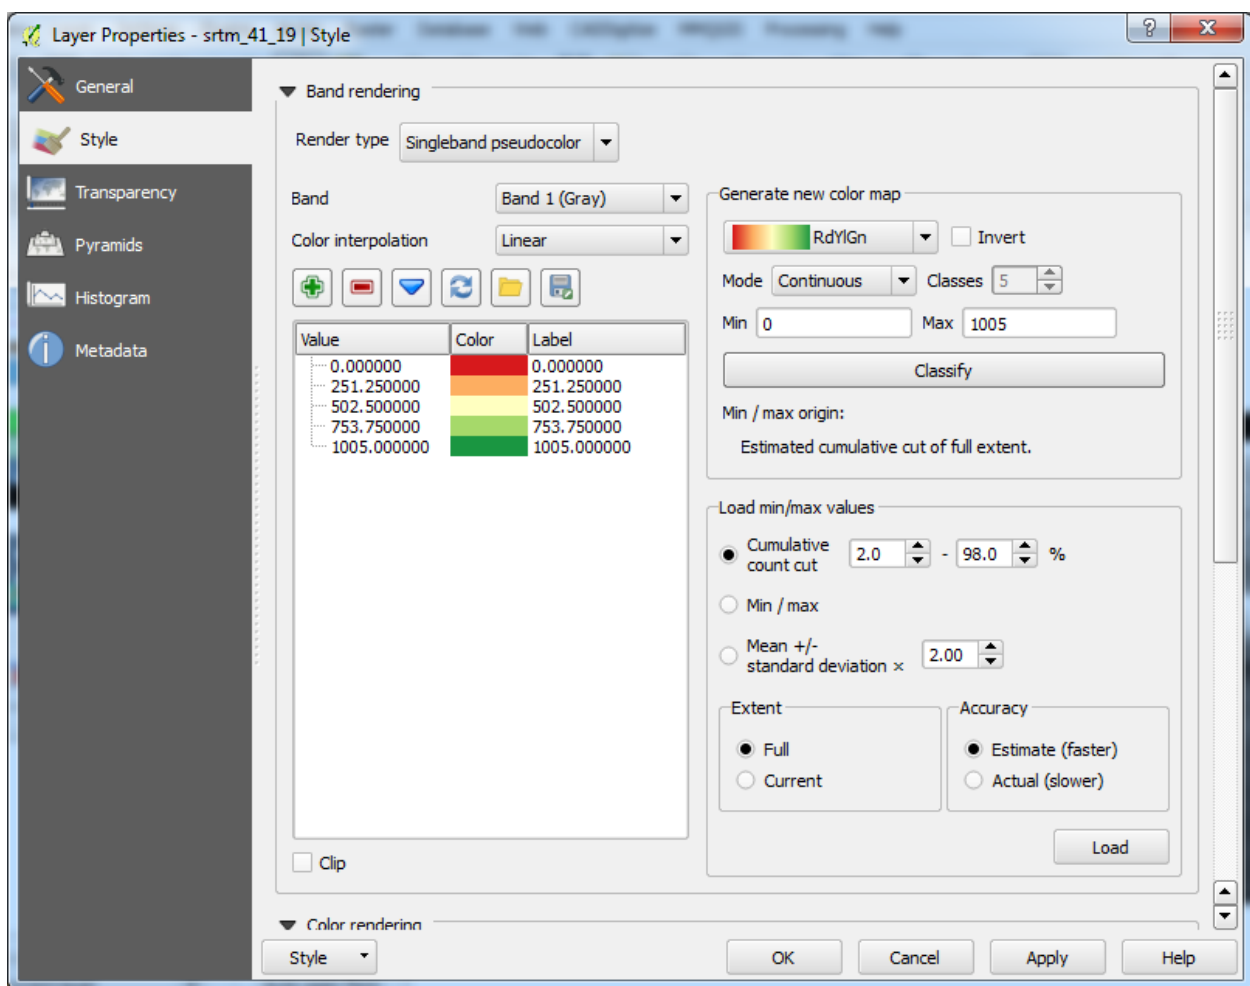
QGIS has automatically applied a stretch to the image for visualization purposes, and we will learn more about how this works as we continue.
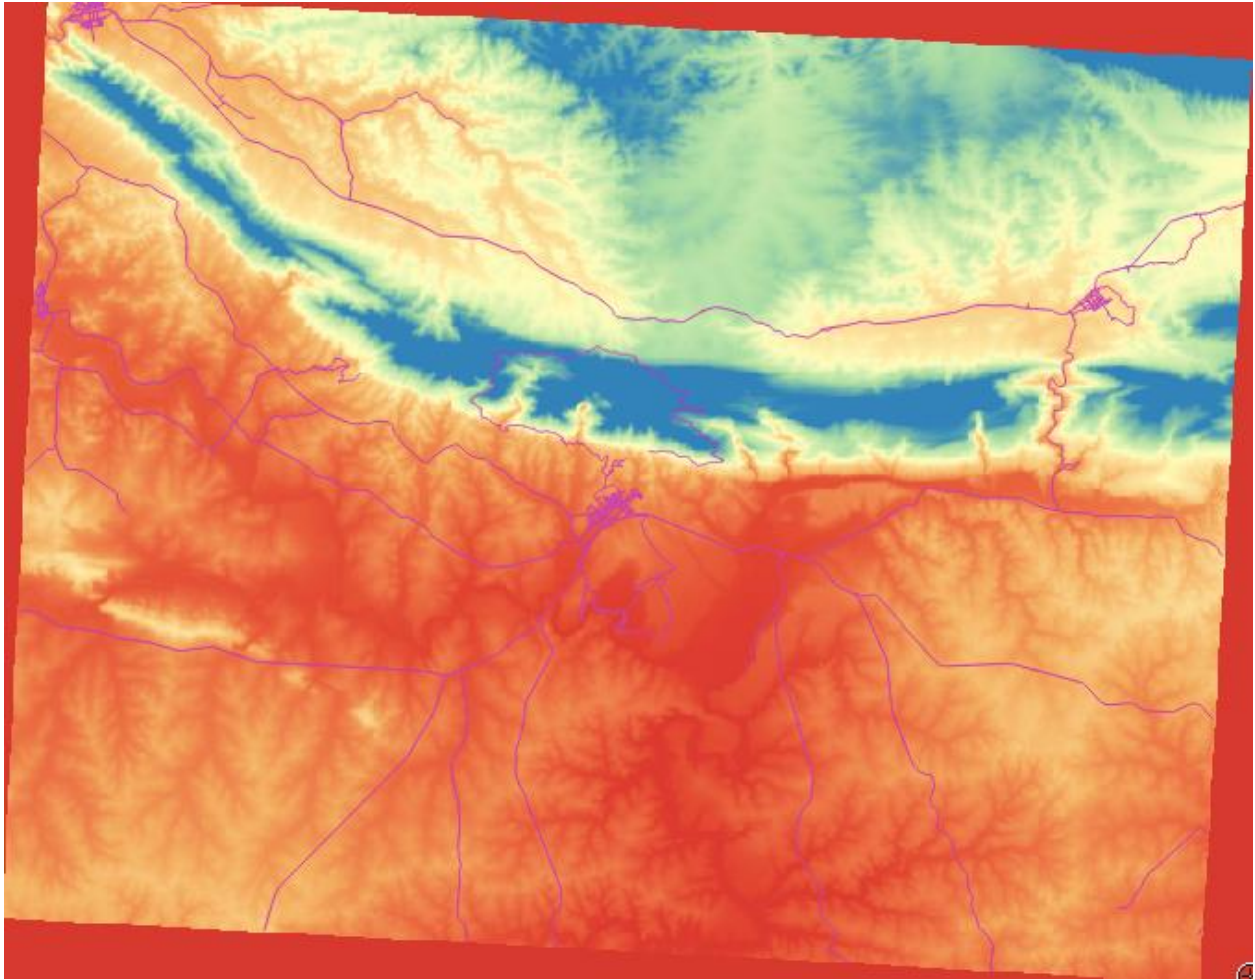
- Open the Layer Properties dialog for the SRTM layer by right-clicking on the layer in the Layer tree and selecting Properties option.
- Switch to the Style tab.

These are the current settings that QGIS applied for us by default. Its just one way to look at a DEM, so lets explore some others.

- Change the Render type to Singleband pseudocolor, and use the default options presented.
- Click the Classify button to generate a new color classification, and click OK to apply this classification to the DEM.

You'll see the raster looking like this:

4 Module: Online Resources

When considering data sources for a map, there is no need to be restricted to data which you have saved on the computer you're working on. There are online data sources which you can load data from as long as you are connected to the Internet.

In this module, we'll cover two kinds of web-based GIS services: Web Mapping Services (WMS) and Web Feature Services (WFS).

4.1 Follow Along: Loading a WMS Layer

A Web Mapping Service (WMS) is a service hosted on a remote server. Similar to a website, you can access it as long as you have a connection to the server. Using QGIS, you can load a WMS directly into your existing map.

From the lesson on plugins, you will remember that it's possible to load a new raster image from Google, for example. However, this is a once-off transaction: once you have downloaded the image, it doesn't change. A WMS is different in that it's a live service that will automatically refresh its view if you pan or zoom on the map.

The goal for this lesson: To use a WMS and understand its limitations.

For this exercise, you can either use the basic map you made at the start of the course, or just start a new map and load some existing layers into it. For this example, we used a new map and loaded the original places and landuse layers and adjusted the symbology:
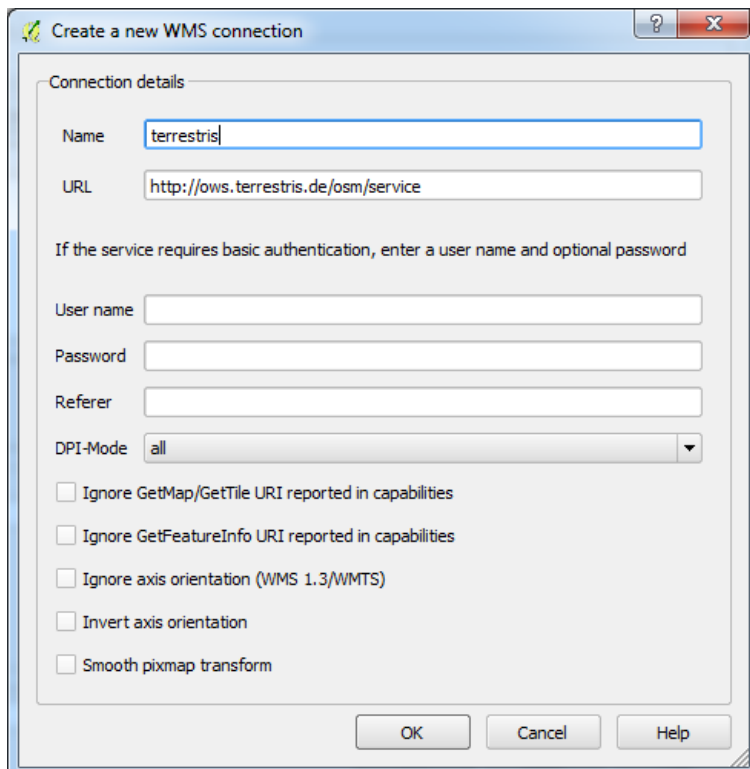
- Load these layers into a new map, or use your original map with only these layers visible.
- Before starting to add the WMS layer, first deactivate "on the fly" projection. This may cause the layers to no longer overlap properly, but don't worry: we'll fix that later.

- To add WMS layers, click on the Add WMS Layer button:

Remember how you connected to a SpatiaLite database at the beginning of the course. The landuse, places, and water layers are in that database. To use those layers, you first needed to connect to the database. Using a WMS is similar, with the exception that the layers are on a remote server.

- To create a new connection to a WMS, click on the New button.

You'll need a WMS address to continue. There are several free WMS servers available on the Internet. One of these is terrestris, which makes use of the OpenStreetMap dataset.

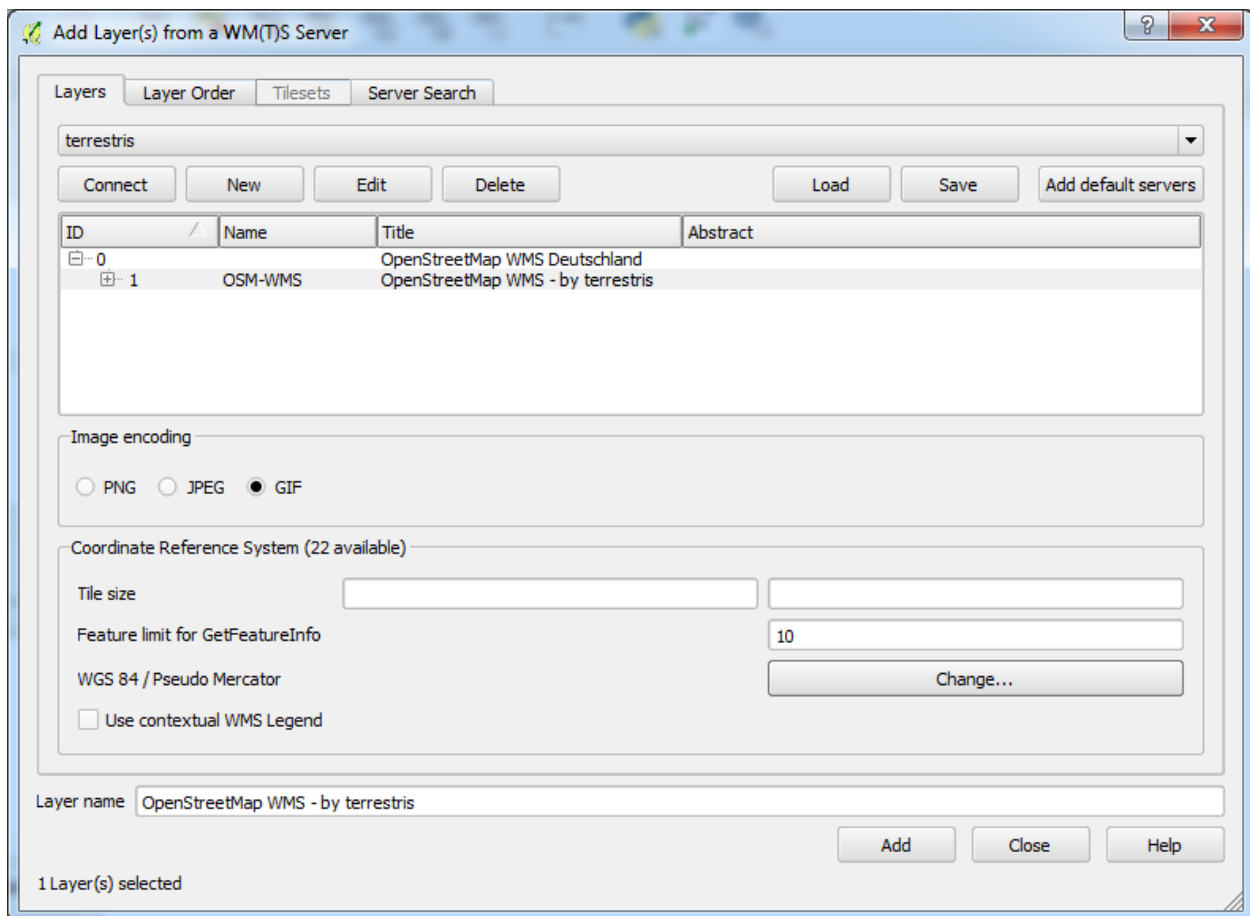- To make use of this WMS, set it up in your current dialog, like this:

- The value of the Name field should be terrestris.
- The value of the URL field should be http://ows.terrestris.de/osm/service.
- Click OK. You should see the new WMS server listed.
- Click Connect. In the list below, you should now see these new entries loaded.

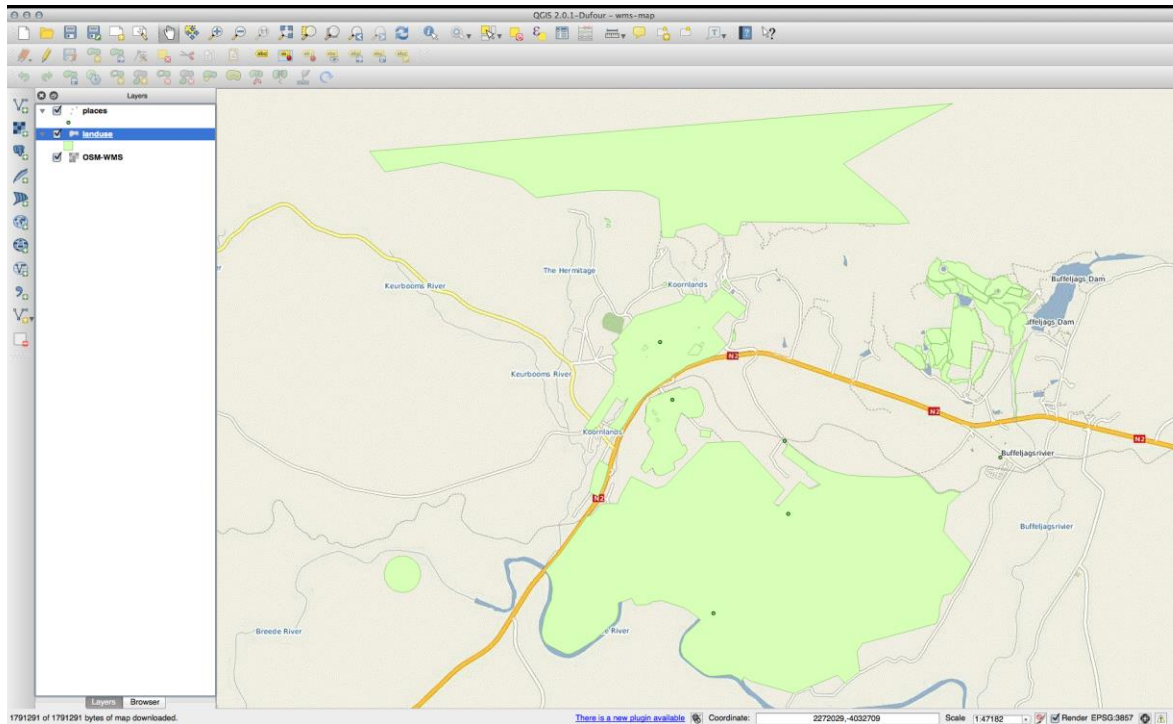These are all the layers hosted by this WMS server.

- Click once on the OSM-WMS layer. This will display its Coordinate Reference System

Since we're not using WGS 84 for our map, let's see all the CRSs we have to choose from.

- Click the Change button. You will see a standard Coordinate Reference System Selector dialog.
- We want a projected CRS, so let's choose WGS 84 / Pseudo Mercator.
- Click OK.
- Click Add and the new layer will appear in your map as OSM-WMS.

- In the Layers list, click and drag it to the bottom of the list.

The nature and limitations of WMS

By now you may have noticed that this WMS layer actually has many features in it. It has streets, rivers, nature reserves, and so on. What's more, even though it looks like it's made up of vectors, it seems to be a raster, but you can't change its symbology. Why is that?

This is how a WMS works: it's a map, similar to a normal map on paper, that you receive as an image. What usually happens is that you have vector layers, which QGIS renders as a map. But using a WMS, those vector layers are on the WMS server, which renders it as a map and sends that map to you as an image. QGIS can display this image, but can't change its symbology, because all that is handled on the server.

This has several advantages, because you don't need to worry about the symbology. It's already worked out, and should be nice to look at on any competently designed WMS.

On the other hand, you can't change the symbology if you don't like it, and if things change on the WMS server, then they'll change on your map as well. This is why you sometimes want to use a Web Feature Service (WFS) instead, which gives you vector layers separately, and not as part of a WMS-style map.

This will be covered in the next lesson.

Try Yourself!
- Hide the OSM-WSM layer in the Layers list.
- Add the "ZAF CGS 1M Bedrock Lithostratigraphy" WMS server at this URL: http://196.33.85.22/cgi-bin/ZAF_CGS_Bedrock_Geology/wms
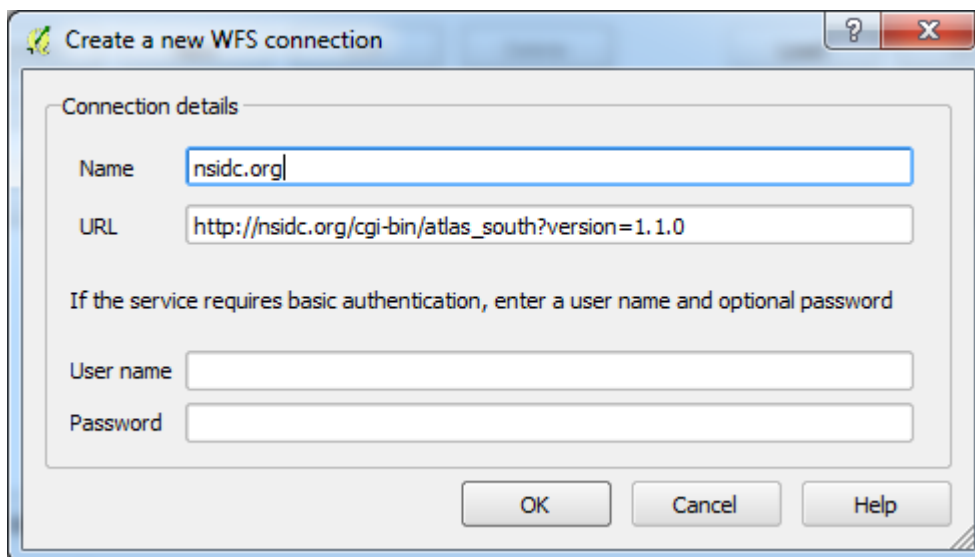
- Load the BEDROCKGEOLOGY layer into the map (use the Add WMS Layer button as before). Remember to check that it's in the same WGS 84 / World Mercator projection as the rest of your map!
- You might want to set its Encoding to JPEG and its Tile size option to 200 by 200, so that it loads faster:
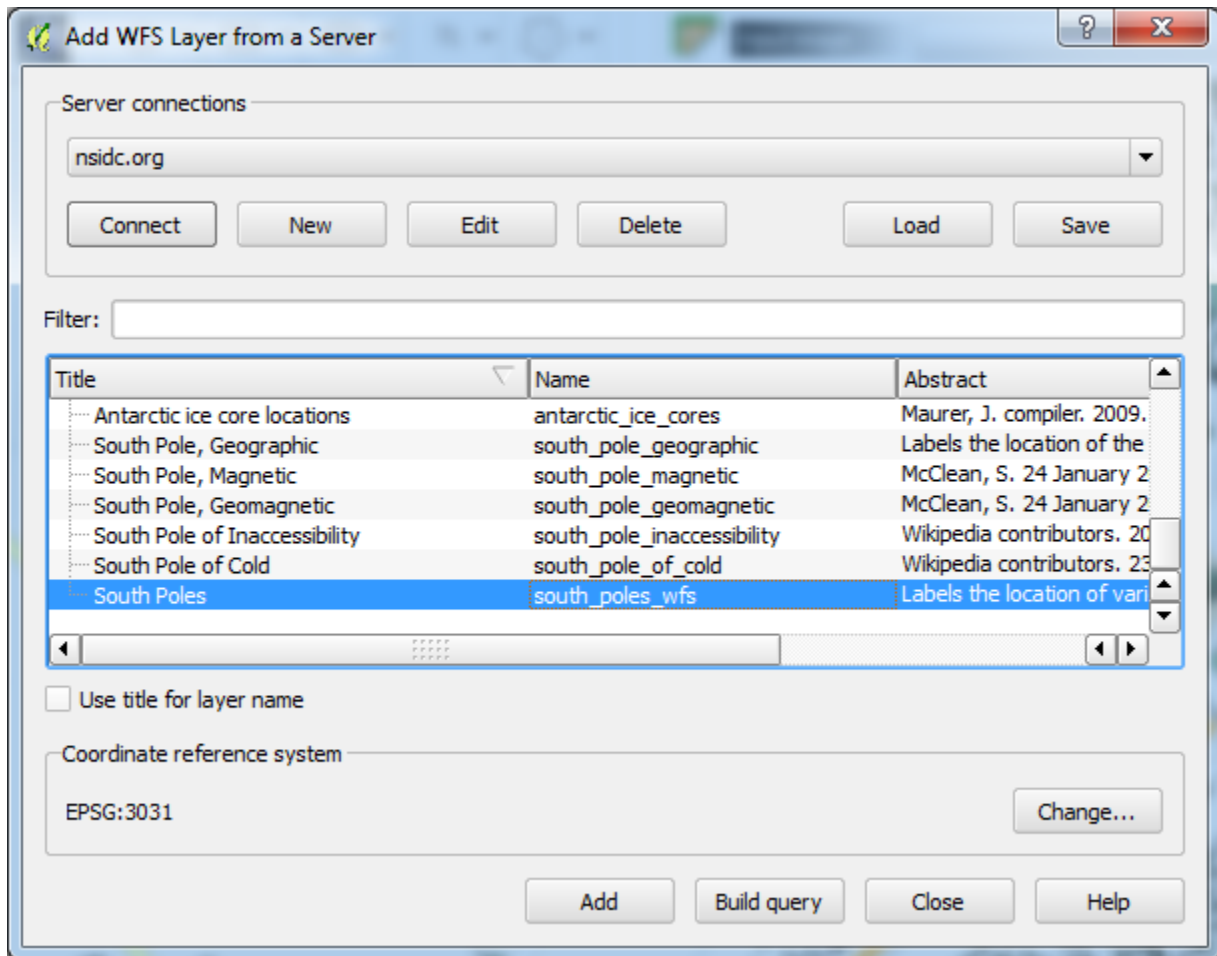
4.2 Follow Along: Loading a WFS Layer

A Web Feature Service (WFS) provides its users with GIS data in formats that can be loaded directly in QGIS. Unlike a WMS, which provides you only with a map which you can't edit, a WFS gives you access to the features themselves.

The goal for this lesson: To use a WFS and understand how it differs from a WMS.

- Start a new map. This is for demo purposes and won't be saved.
- Ensure that "on the fly" re-projection is switched off.

- Click the Add WFS Layer button:
- Click the New button.
- In the dialog that appears, enter the Name as nsidc.org and the URL as http://nsidc.org/cgi-bin/atlas_south?version=1.1.0.
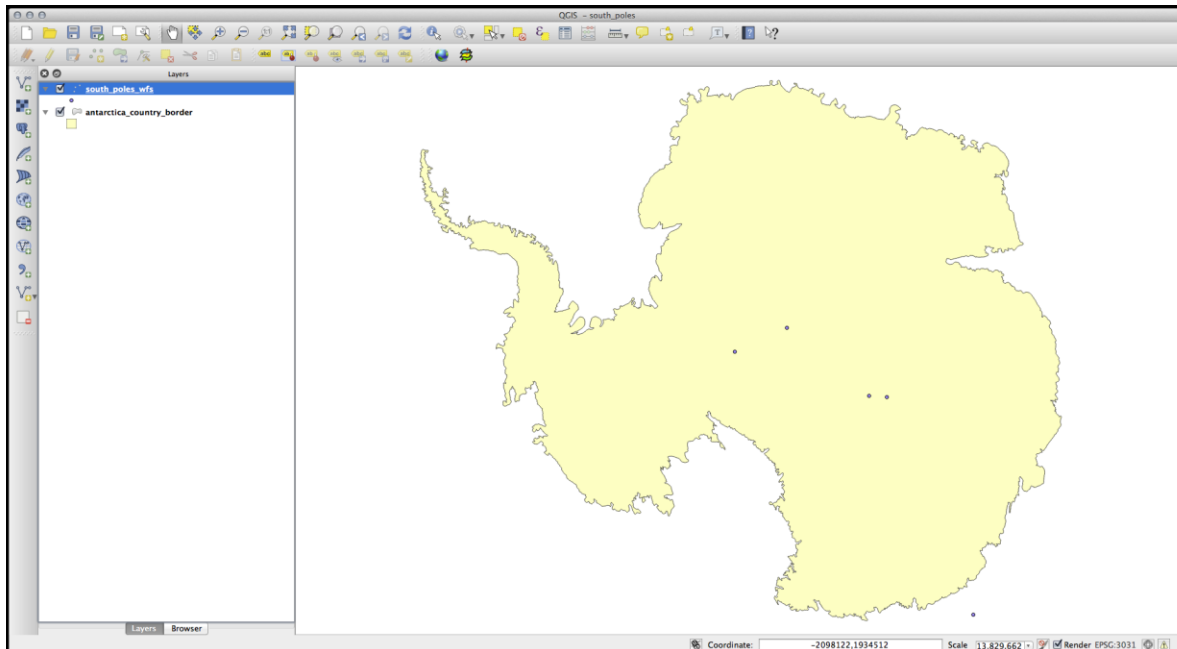

- Click OK, and the new connection will appear in your Server connections.
- Click the Connect. A list of the available layers will appear:
- Find the layer south_poles_wfs.
- Click on the layer to select it:

- Click Add.

It may take a while to load the layer. When it has loaded, it will appear in the map. Here it is over the outlines of Antarctica (available on the same server, and by the name of antarctica_country_border):
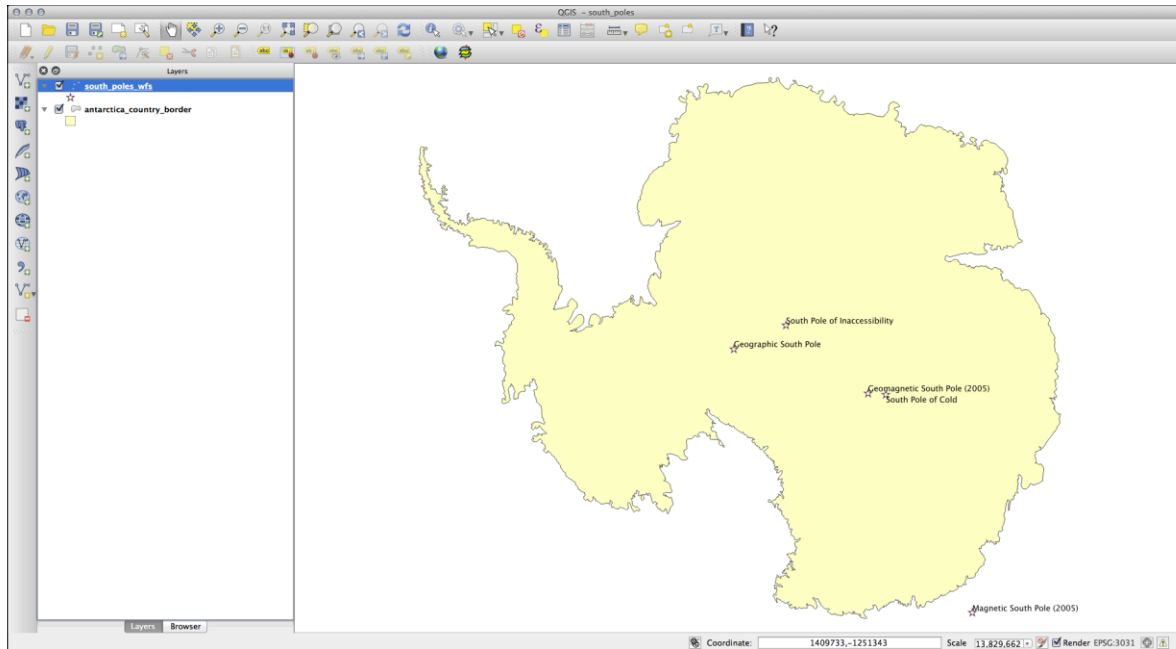
How is this different from having a WMS layer? That will become obvious when you see the layers' attributes.

- Open the south_poles_wfs layer's attribute table. You should see this:

Since the points have attributes, we are able to label them, as well as change their symbology. Here's an example:

- Add labels to your layer to take advantage of the attribute data in this layer.

Differences from WMS layers

A Web Feature Service returns the layer itself, not just a map rendered from it. This gives you direct access to the data, meaning that you can change its symbology and run analysis functions on it. However, this is at the cost of much more data being transmitted. This will be especially obvious if the layers you're loading have complicated shapes, a lot of attributes, or many features; or even if you're just loading a lot of layers. WFS layers typically take a very long time to load because of this.

4.3   Follow Along: Querying a WFS Layer

Although it is of possible to query a WFS layer after having loaded it, it's often more efficient to query it before you load it. That way, you're only requesting the features you want, meaning that you use far less bandwidth.
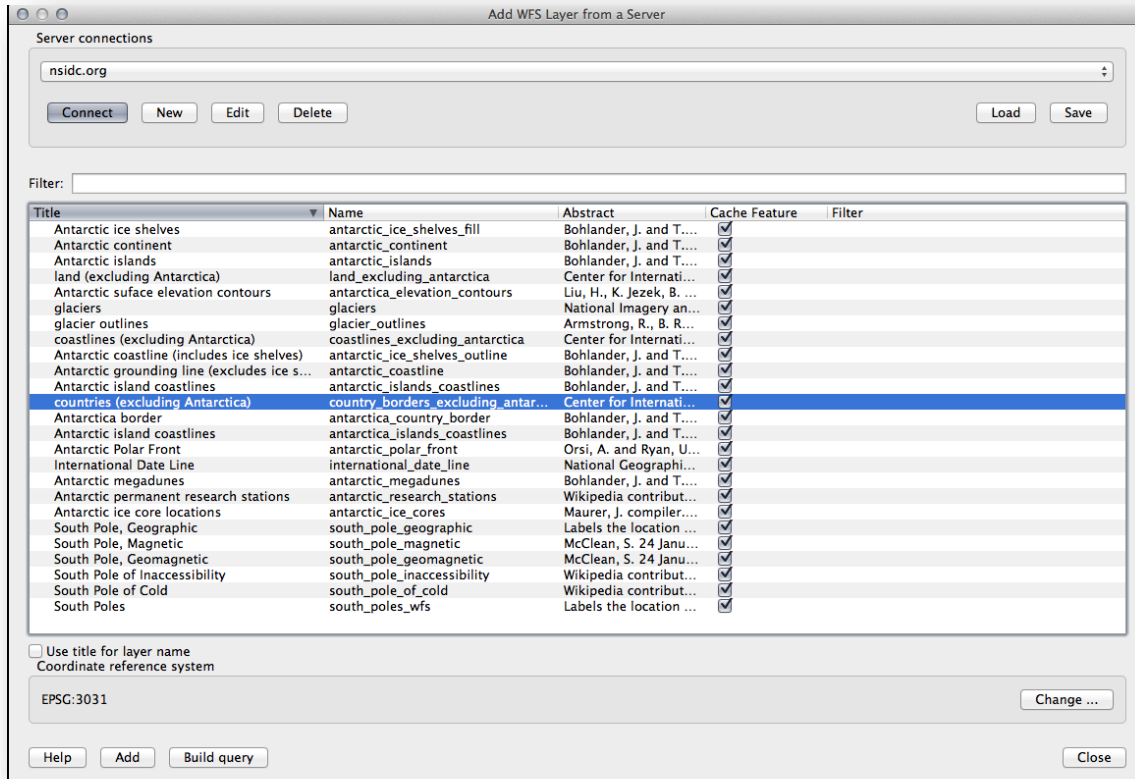
For example, on the WFS server we're currently using, there is a layer called countries (excluding Antarctica). Let's say that we want to know where South Africa is relative to the south_poles_wfs layer (and perhaps also the antarctica_country_border layer) that's already been loaded.

There are two ways to do this. You can load the whole countries ... layer, and then build a query as usual once it's loaded. However, transmitting the data for all the countries in the world and then only using the data for South Africa seems a bit wasteful of bandwidth. Depending on your connection, this dataset can take several minutes to load.
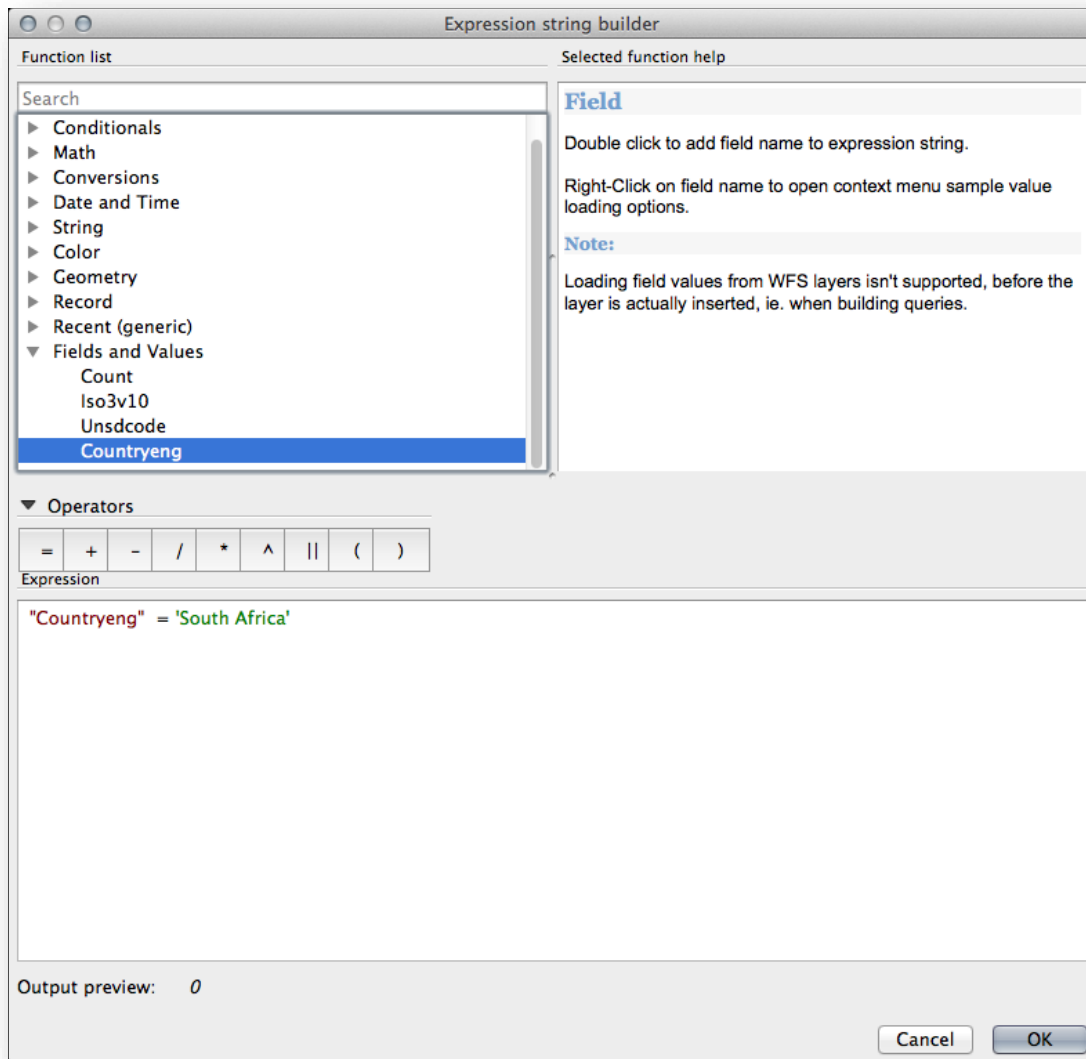
The alternative is to build the query as a filter before even loading the layer from the server.

- In the Add WFS Layer ... dialog, connect to the server we used before and you should see the list of available layers.
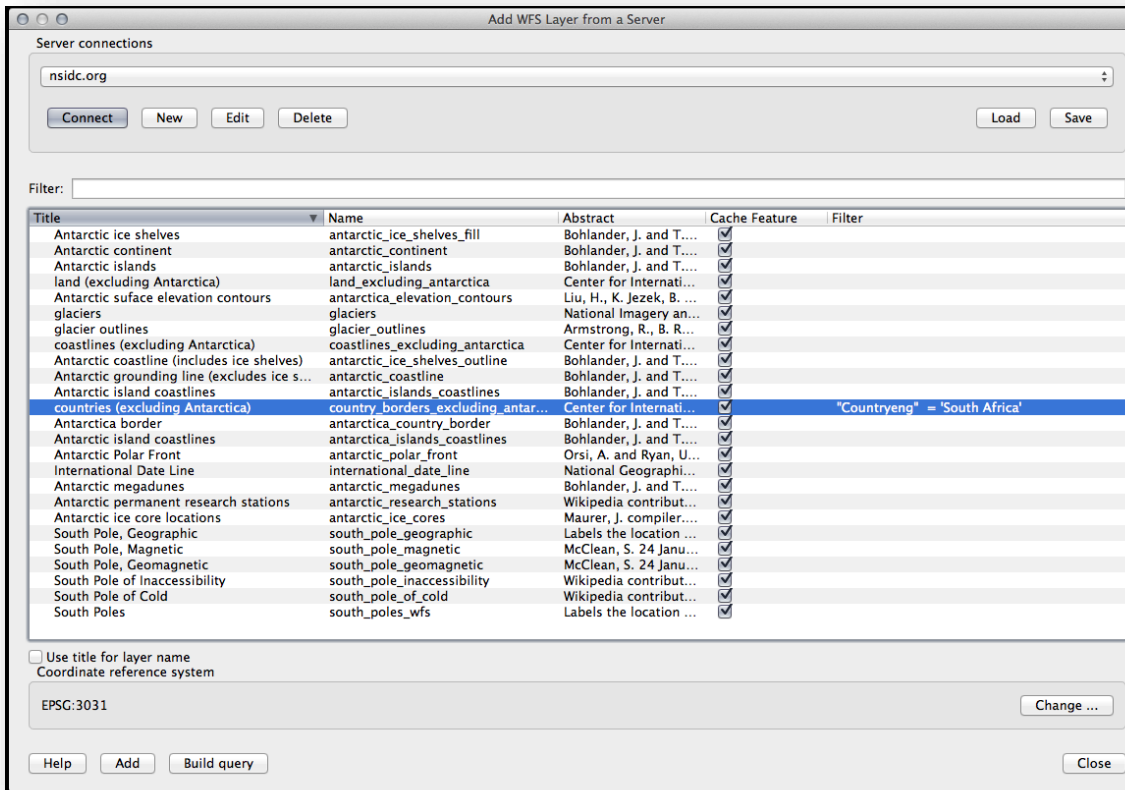
- Double-click next to the countries ... layer in the Filter field, or click Build query:
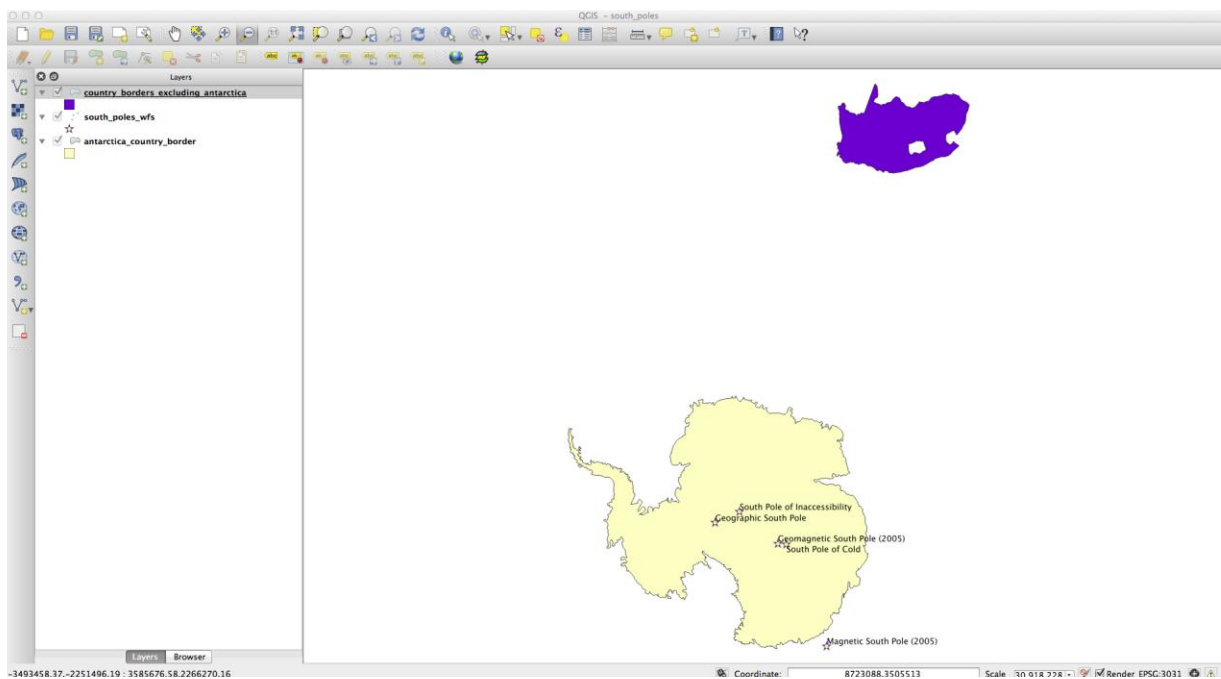

- In the dialog that appears, build the query "Countryeng" = 'South Africa':

- It will appear as the Filter value:

- Click Add with the countries layer selected as above. Only the country with the Countryeng value of South Africa will load from that layer:

You don't have to, but if you tried both methods, you'll notice that this is a lot faster than loading all the countries before filtering them!

Notes on WFS availability

It is rare to find a WFS hosting features you need, if your needs are very specific. The reason why Web Feature Services are relatively rare is because of the large amounts of data that must be transmitted to describe a whole feature. It is therefore not very cost-effective to host a WFS rather than a WMS, which sends only images.

The most common type of WFS you'll encounter will therefore probably be on a local network or even on your own computer, rather than on the Internet.

In Conclusion

WFS layers are preferable over WMS layers if you need direct access to the attributes and geometries of the layers. However, considering the amount of data that needs to be downloaded (which leads to speed problems and also a lack of easily available public WFS servers) it's not always possible to use a WFS instead of a WMS.