

Lab 1: Introduction to QGIS

What You'll Learn:

This is a basic introduction to starting QGIS, adding data, changing symbols, creating maps, and topology.

Data for this exercise are located in the L1 zip file, we assume it has been downloaded and expanded. Videos are in the QGIS 3.x Lab1 playlist of the Mapeez Youtube channel.

What You'll Produce: Three maps: one of lakes and roads, one of wetlands, and an inset map with raster and vector data.

Background: This is the first in a series of introductory exercises for QGIS, versions 3.x. Most were written under version 3.10 as the long-term releases. These are practical skills that complement the theory and practice of GIS described in the textbook "GIS Fundamentals: A First Text on Geographic Information Systems", by Paul Bolstad. You should read the first two chapters in that book before doing this exercise.

There are companion videos and playlists on the Youtube channel named Mapeez. They are organized by playlists. Playlist names correspond to the exercise, e.g., "Lab 1 QGIS3.x". Find and familiarize yourself with the Mapeez Youtube page. There are various playlists there, the videos referred to in this lab are in the QGIS3-Lab1 playlist. There are many other helpful Youtube channels for QGIS, for example, those by Klas Karlsson, Hans van der Kwast, and burdGIS.

We assume you have a functioning copy of QGIS running on your computer. If you don't, please see the video on the Mapeez channel on Youtube, named DownloadInstallQGIS3.x on loading QGIS, or visit the www.qgis.org website and follow the links to install.

We you have a copy of the needed data files, from an instructor or downloaded through links on the paulbolstad.net/gisbook.html or paulbolstad.org/gisbook.html web pages.

In the subsequent instructions, we will write "click" when we mean a left-click on the mouse button. We will explicitly identify those clicks of the mouse button that need to be right clicks.

Part 1: Starting QGIS, adding data and creating your first map

First, find the QGIS icon, shown to the right. The icon is often located



1) as a desktop or taskbar shortcut, or

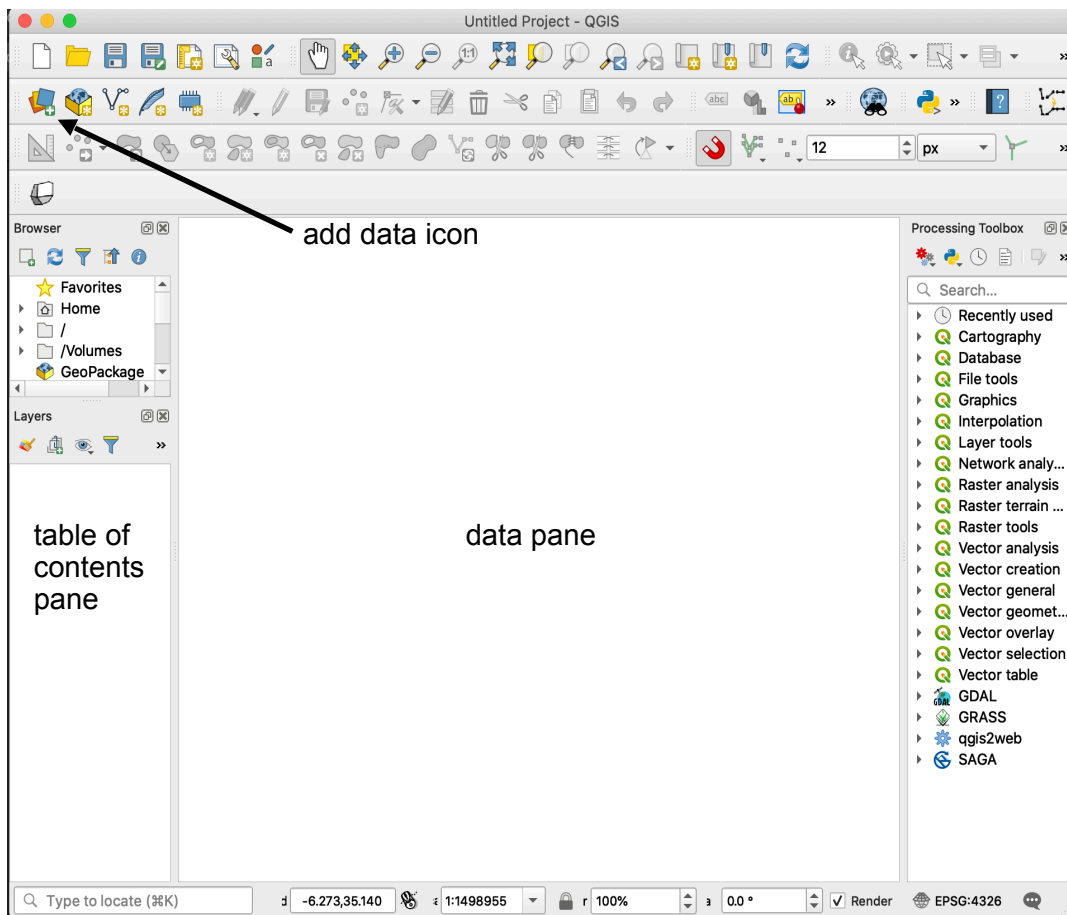
2) on a Microsoft Windows operating system (shortened to MSWin or Windows in future instructions) there may be a QGIS folder in the programs menu, opened by the button in the lower left corner of the screen.

(see the Youtube Mapeez site, QGIS3 Lab1 playlist, for the **video start QGIS 3.x**)

In Windows QGIS may often be found by left clicking on the Start button in the lower left of the screen and selecting Programs→OSgeo4W→QGIS. On a Mac, QGIS is usually found in the Applications Folder. The containing folder name and configuration on both operating systems depend on how the system was setup at install, so QGIS may be in different folders, and you may have to look a bit.

Double left click on the QGIS icon, and be patient while a start banner displays. You may first see a user tip, if so, dismiss it. The main QGIS window should appear, similar to that displayed below. Note there is **Table of Contents** pane (sometimes abbreviate TOC in this and later exercises) that QGIS calls a Layers pane, typically on the left side of the main window. This lists the layers you have currently loaded in your project. There is a Data pane, which displays your layers, as well as various tool bars.

One of the most used icons is for adding data layers, noted in the graphic below:

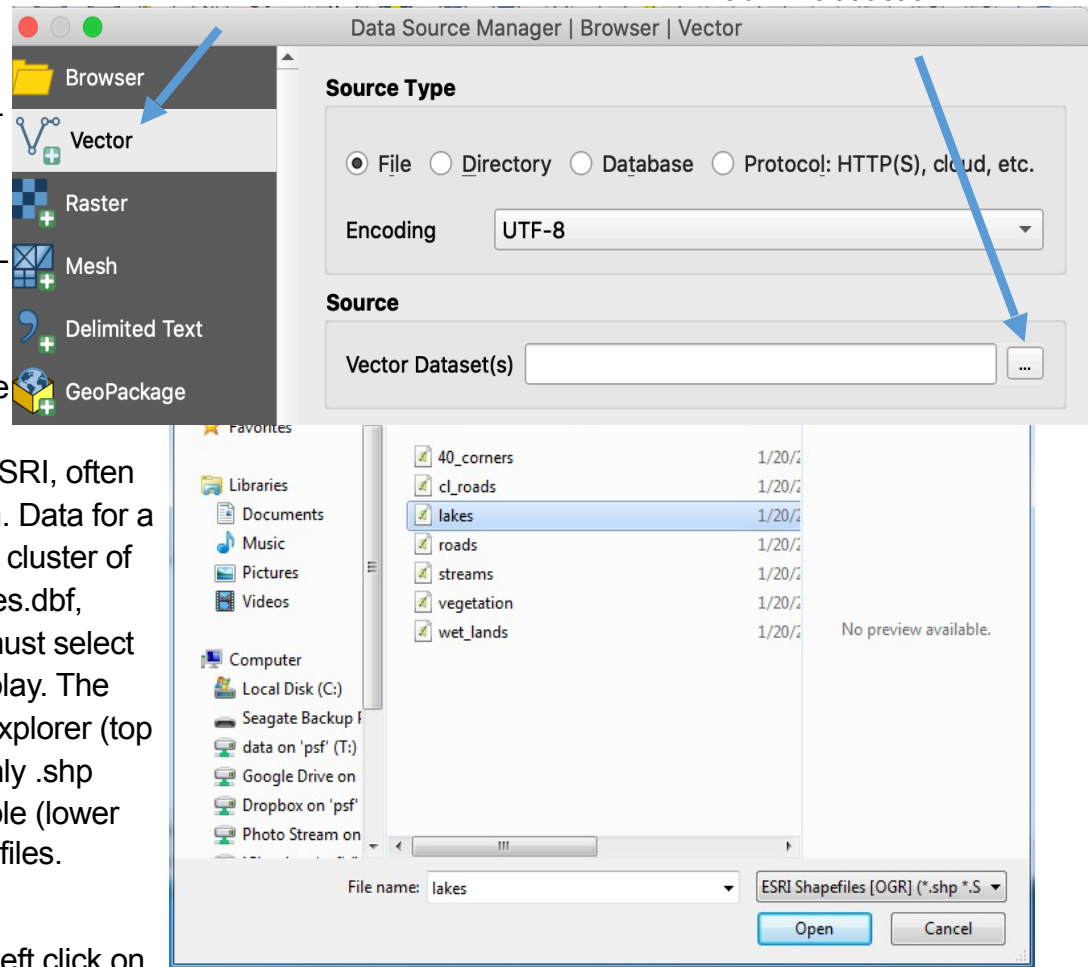


A left click on the add data icon will open a window, allowing you to identify a data type and file. Here we've selected a vector data type, and clicked on the browse button at the right end of the Source box:

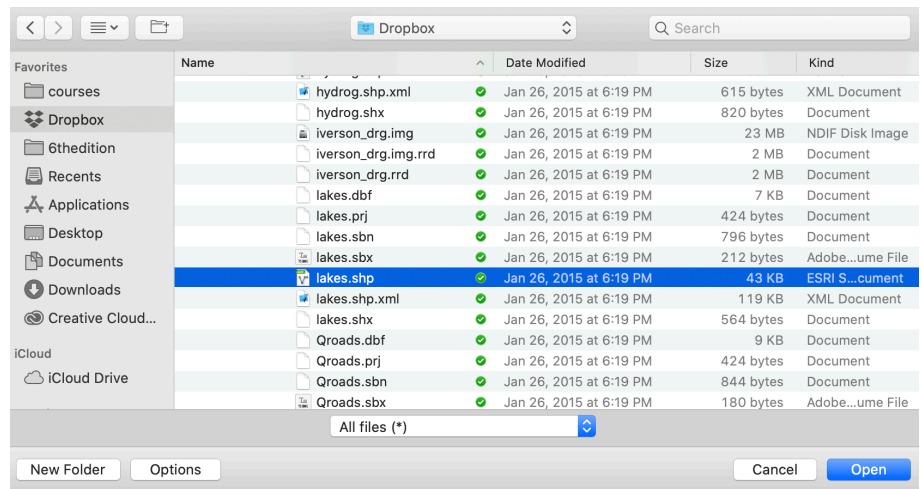

You should then see a navigation window similar to that shown at the right, although the general form will vary depending on your operating system. We show a version of the

QGIS 3.x Exercises
Windows directory that displays at right, top, and MacOS at right, below.

Note that for both operating systems, we have selected the lakes.shp file. This is a common file format, developed by a large software vendor, ESRI, often used for data distribution. Data for a shapefile are saved in a cluster of files, e.g., lakes.shp, lakes.dbf, lakes.sbn, etc., but we must select the lakes.shp file for display. The Windows 7 version file explorer (top right), is set to display only .shp files. The MacOS example (lower right) is set to display all files.



In both cases we single left click on the lakes.shp file to select it, then click the Open button. This should close the file display window and revert back to the Data Source selection window. Clicking on Add in the lower right to display the data, and click on Close to hide the window.



For efficiency, most future graphics will only show examples for one operating system (MacOS or Windows 7 or 10), and may switch back and forth. QGIS windows and menus are usually identical

across operating systems, and always similar enough that you should be able to interpret meaning. Major differences come in making calls to operating system functions, e.g., when selecting data files.

Repeat the add data process, using the roads shapefile.

You should have a data pane that looks something like the figure at right. Your colors will be different, we'll show you how to change symbology soon.



Panning and Zooming.

QGIS allows you to change the magnification and area that you view in your data pane.

There is a cluster of zoom buttons (see right). They are typically along the top of the main window, but because the toolbars are “dockable”, they may be elsewhere.




If you hold the cursor near each icon for a few seconds, descriptive text should appear, with the name/brief description of the display function.

If you hold the cursor near each icon for a few seconds, descriptive text should appear, with the name/brief description of the display function.

Left click on the zoom (plus/minus magnifying glass) and pan (hand) icons to change cursor function. Left clicking on the plus (+) magnifying glass changes it to a “zoom in” cursor, then click on the data pane will zoom in on a point. You can also left click and hold/drag to define a zoom area.

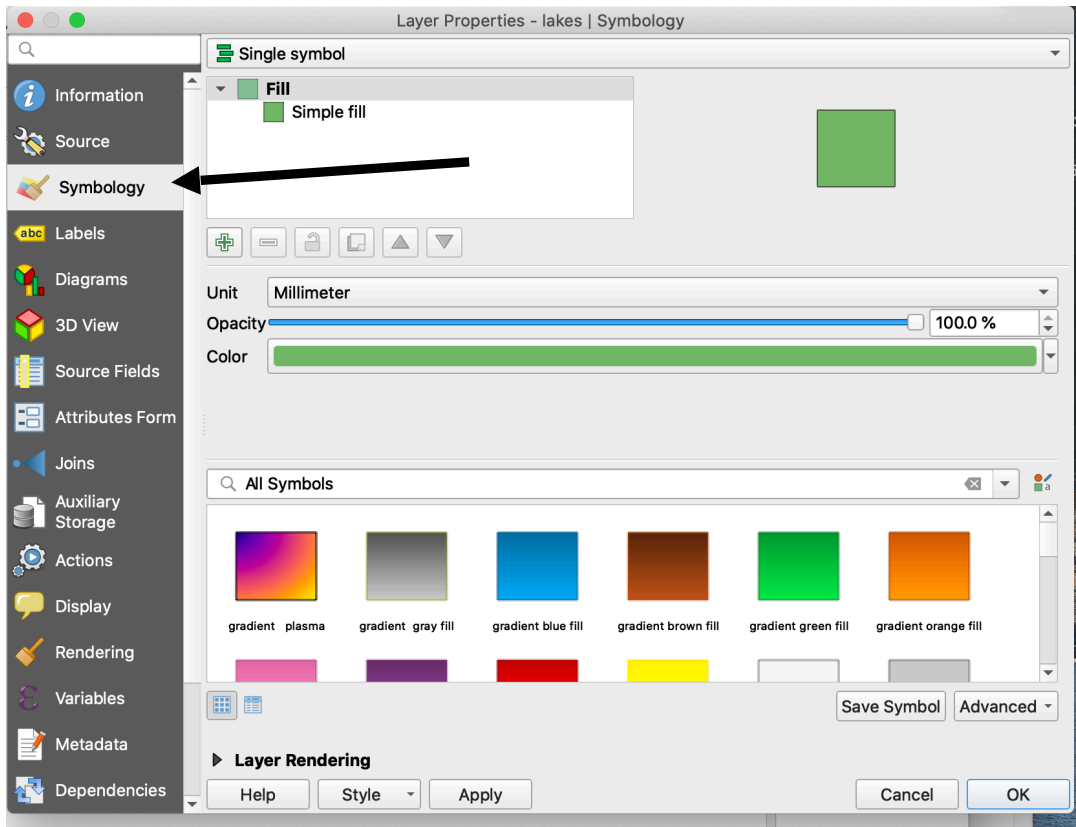
The pan button, a hand, allows you to click/drag to reposition views of displayed layers. There is also a three-arrow zoom button that zooms to the full “Extent” of your data. There is a magnifying glass with a yellow square behind that zooms to selected features (more on that later), and one with a gray square behind that zooms to a layer selected in the table of contents. There are also two buttons with carets that zooms back and forth among previous zoom levels.

Also note that some of the buttons can be hidden on a menu bars. Hidden buttons/icons are indicated by double right carets at the right end of a set of buttons. Clicking on the carets will show the missing buttons, e.g., for the zoom menu:



Layer Symbology

We can customize a layer's appearance. Left double-click on the name of the lakes data layer in the table of contents (TOC), or right click on the name and then scroll to near the bottom of the dropdown and select Properties. The properties window will appear, select on the Symbology option in the left-most column (see figure).



There is a **video** named [**QGIS 3.x Simple Symbology**](#) on the Youtube Mapeez channel.


Change the lake symbol to blue via the color and other controls on the right of this window.

Left click on the Apply and **OK**.

Repeat this process for the Roads layer, making the roads black and approximately double the original thickness, click OK, and note the changed symbology on the displayed layers

Creating a Map

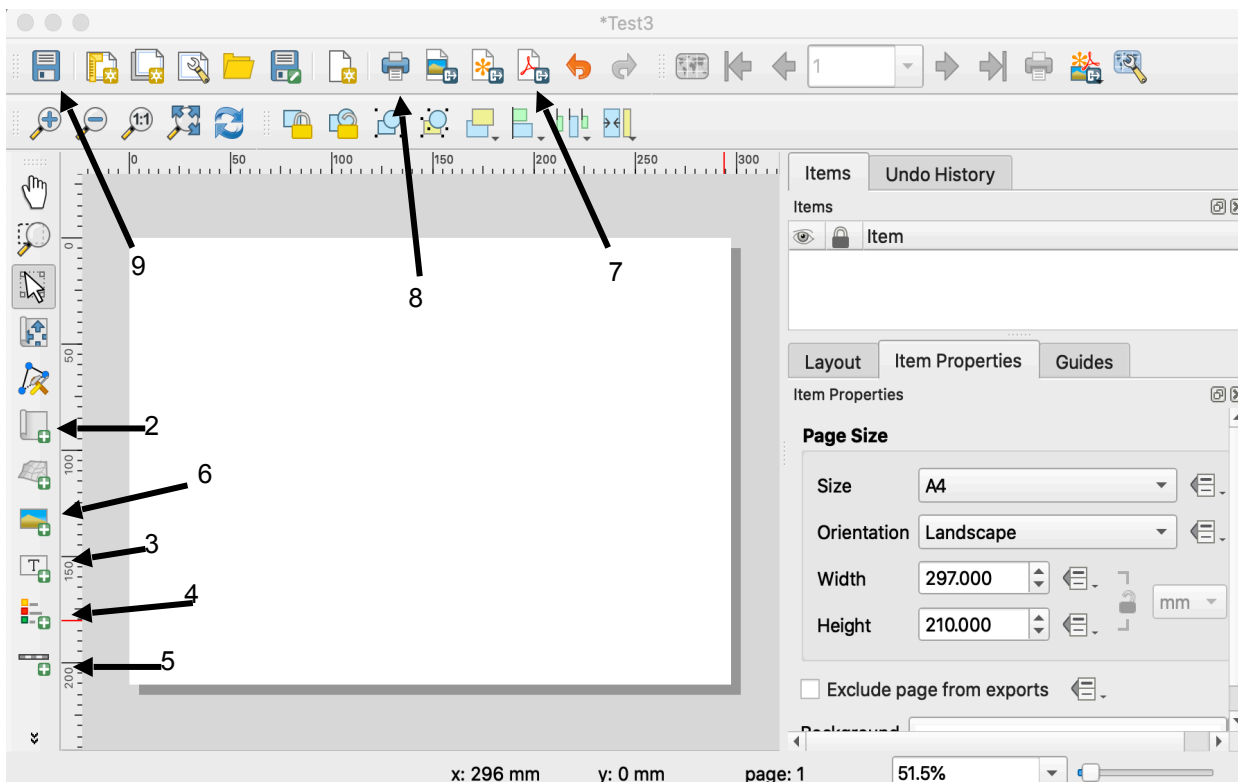
QGIS uses a print composer to create maps, typically showing data, legends, scalebars, and other features. The print composer is available through an icon in the main window (see right), from the Project/New Layout option from the main menu. See the ([video QGIS3x Basic Layout](#) on the Mapeez Youtube).



Clicking the wrench/page icon opens a layout manager (not shown), which asks you name a new Empty layout. Click the Create button and a window asks you to name a print layout, enter a name and click OK to display a composition panel (see below).

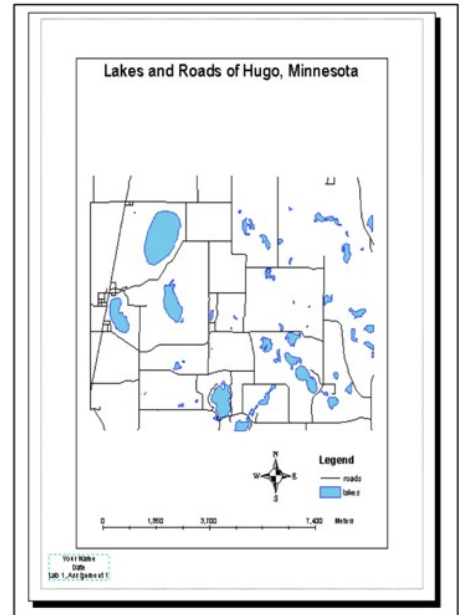
The typical sequence is to click to

- 1) select the paper size (right click on the page, then select page properties, and adjust the values for page size in the right panel), then click margin icons to
- 2) add a new map (select icon then click/drag a box to define an area),
- 3) add a title (click/drag, then modify item text and font properties in the right pane),
- 4) add a legend,
- 5) add a scalebar, and
- 6) add a north arrow, by clicking on add image, and then selecting and placing one of the images of an arrow (see the video referenced above),
- 7) export a pdf of your map, or 8) print it.



When you are done adding and adjusting map elements, you can export the map to a pdf (7, in the figure above), or print it (8), and save it (9).

Create a map that looks something like that in the figure above right and export it as a pdf.



Setting Relative Paths While Saving

QGIS project files have a .qgs extension. These project files do not save any real data. They only contain instructions on how to display data from various sources. Sources are most often files, and the file path to each data set is stored. This can present some problems when moving projects among computers, as the drive letters or directories may differ between systems. If you copy a project to another location, the path to data may no longer be valid, and the data may not display. We'll now show you how to avoid some of these problems.

First, create a new QGIS project (save your old project first), then left click on the New Map button, shown at right.

Add the roads.shp layer from the L1 directory.

Right click on the name roads in the table of contents window, and left click on Properties at the bottom of the menu.

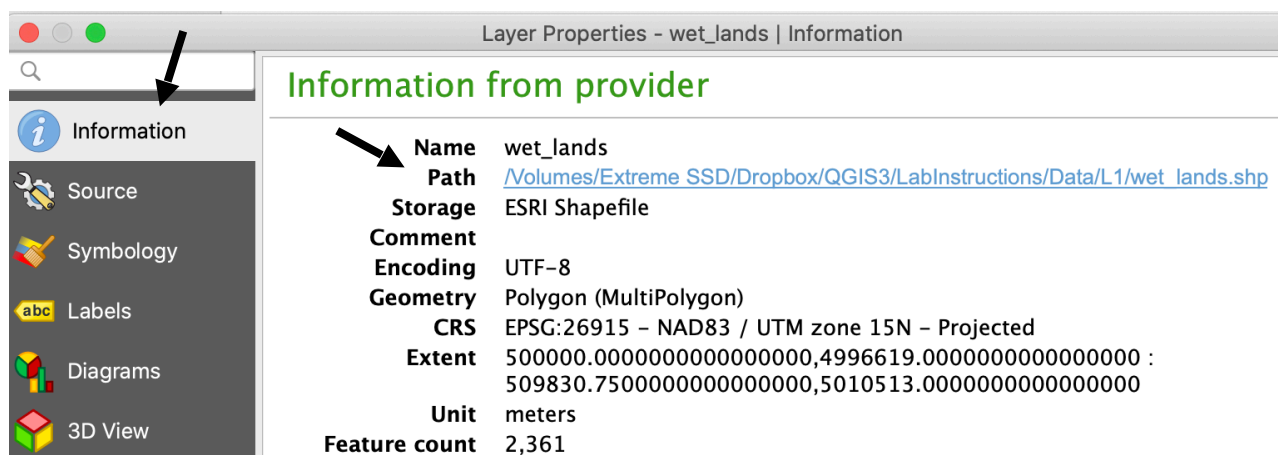
A new window will open, and click on the Information option at the top of the left-hand column (see arrow in figure below). Look at the layer path found listed second in the main window. Note that there is a path that may start with a drive letter, something like "E:\L1\roads.shp," or it may start with a directory, e.g., /Volumes/HD/QGIS3/Data/roads.shp, or an even longer path, as shown below.

The path is recorded for each data set within the QGIS project. This is where the project looks to assemble your data view each time you open any previously created and saved project. This path is the drive and sequence of subdirectories that lead to the displayed

data file. Your path will be different, depending on the directory you are using to store your data, but the important point is that if you are using a fixed or absolute path, your QGIS project will look along the full path for the data, and this can cause problems.

Unfortunately, a fixed path storage arrangement isn't very flexible, or portable. If you move your project files, including all data, to another computer, the drive letter or directory you save the data into will likely be different, for example, C:\ or D:\ instead of E:\. The path to the data will then be incorrect, and the project won't be able to find and display your data. In the current state, the map project is difficult to move between computers.

Perhaps worse, even if you don't move the data, but do something as simple as rename any directory in the path, the map project won't locate the data correctly. A fixed path isn't a problem if you always work on a computer with fixed drives, and you never change the subdirectories, but that often isn't the case.




Close the data layer information window.

There should be a ribbon across the top of your screen, or the window in which QGIS is open (it depends on operating system and configuration).

On a Mac using the iOS operating system it will look like this:



On Windows 10 it looks like this:

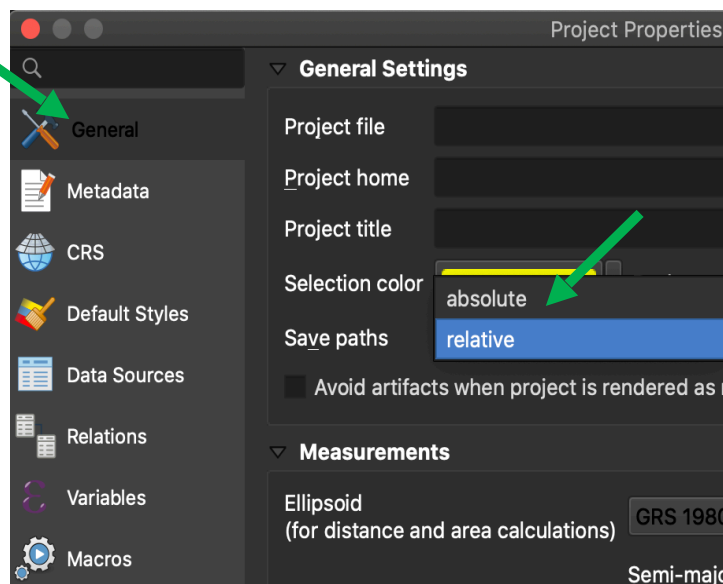


Left click on Project near the upper left of your screen/window, and then on Properties.... in the subsequent drop-down menu (figure at right). This will open a new menu, click on the General option (first in list in leftmost column), and then set save paths to relative (figure at right).

Relative paths only save as much of the path as needed to get from your data layer to your project.

If you store your project in the same directory as your data, or as near as possible to your data in the directory tree, a change in drive letter or base folder is less likely to result in wayward data.

Projects are more robust to folder name changes and more portable when you set the path as relative, and store the data and project files in the same directory/subdirectory.



Of course, sometimes it is not possible to save all the data and your project in the same directory. You may have large datasets organized in a directory structure for easy updates, with data on different drives, and so don't want to copy them to the same location as this would take up too much space, meaning a drive letter or complex path is inevitable for one of the data sets. Or you make work for a large organization, with data on linked, shared, or network drives, with different drive letters when you move between computers. However, when working on stand-alone projects, or with smallish data sets, it is best practice to have the project file (.qgs) and the data in the same directories.

Think of the .qgs file as the recipe, and the data as the ingredients. You need both of these to display a map. If whenever possible, you save the .qgs in the same directory as the data and you set relative path names, then you can easily move both the .qgs map project and the data the same time.

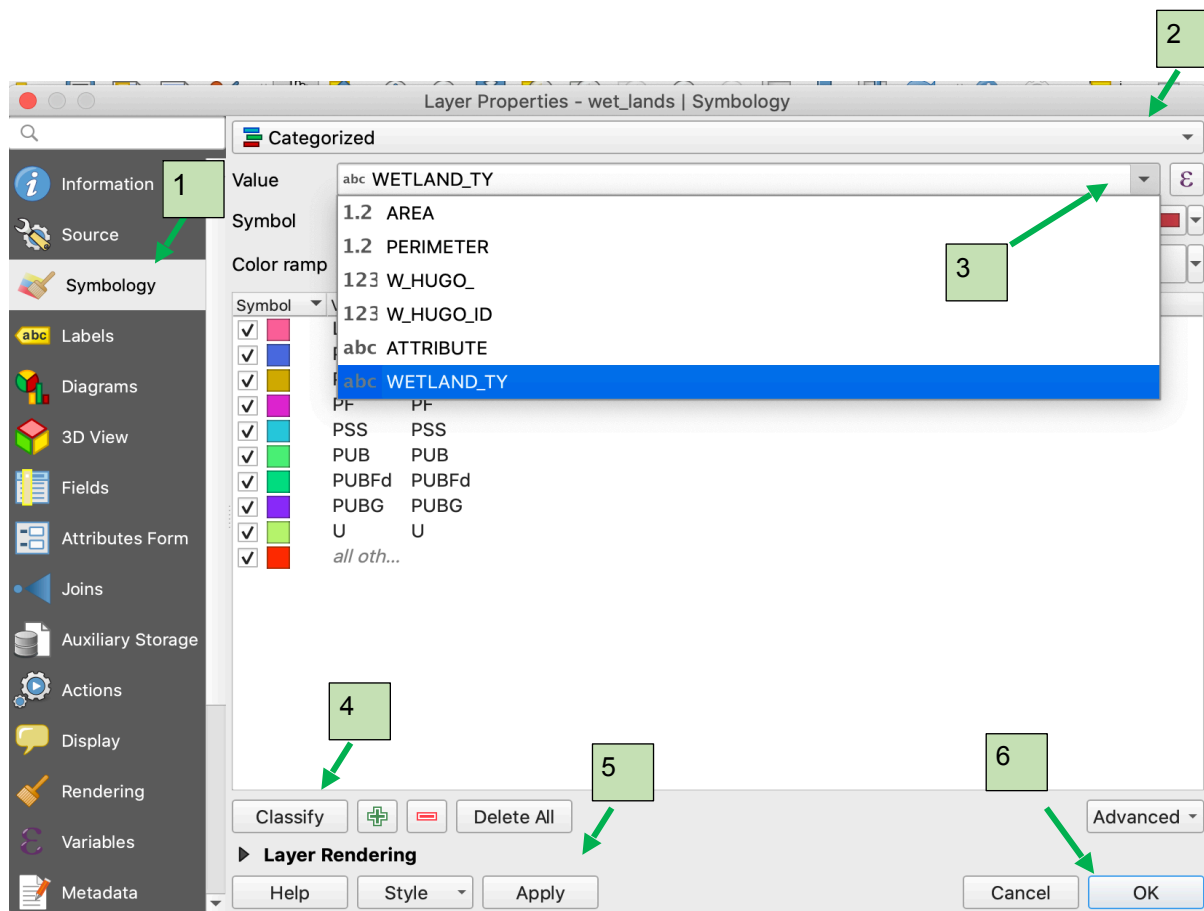
Manipulating Symbology

Remove the roads layer (right click on the name in the TOC, then left click on Remove Layer... in the dropdown menu), and add the layer "wet_land.shp" from your data directory. This layer shows polygons that depict the wetlands of the Hugo USGS quadrangle, in Minnesota.

After adding the data, left click on the name of the layer (*wet_lands*) and right click to select Properties. This opens a Layer Properties window, you saw earlier (see the graphic at right). Now we'll color the wetland polygons by different wetland categories (video [QGIS 3x Categorized Symbols](#), on Mapeez in Youtube).

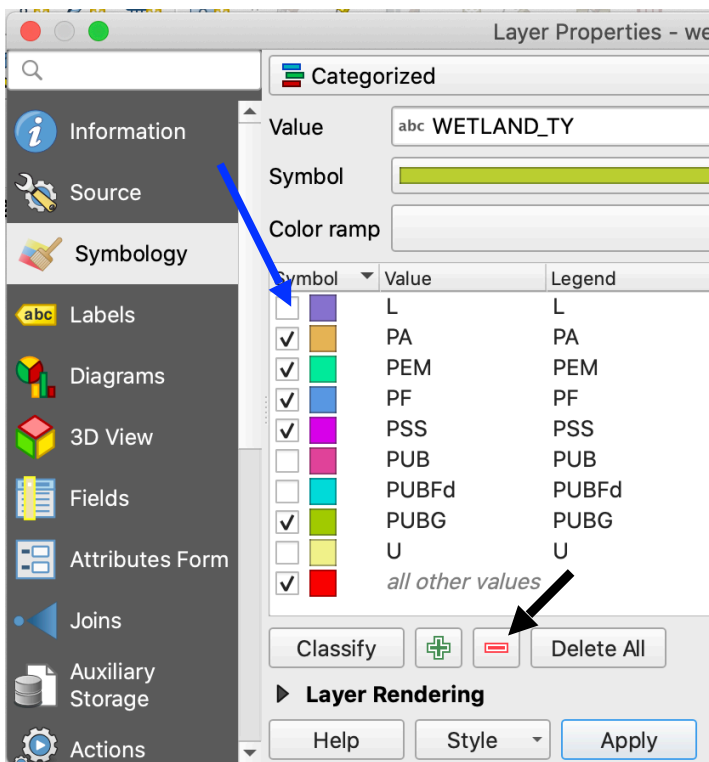
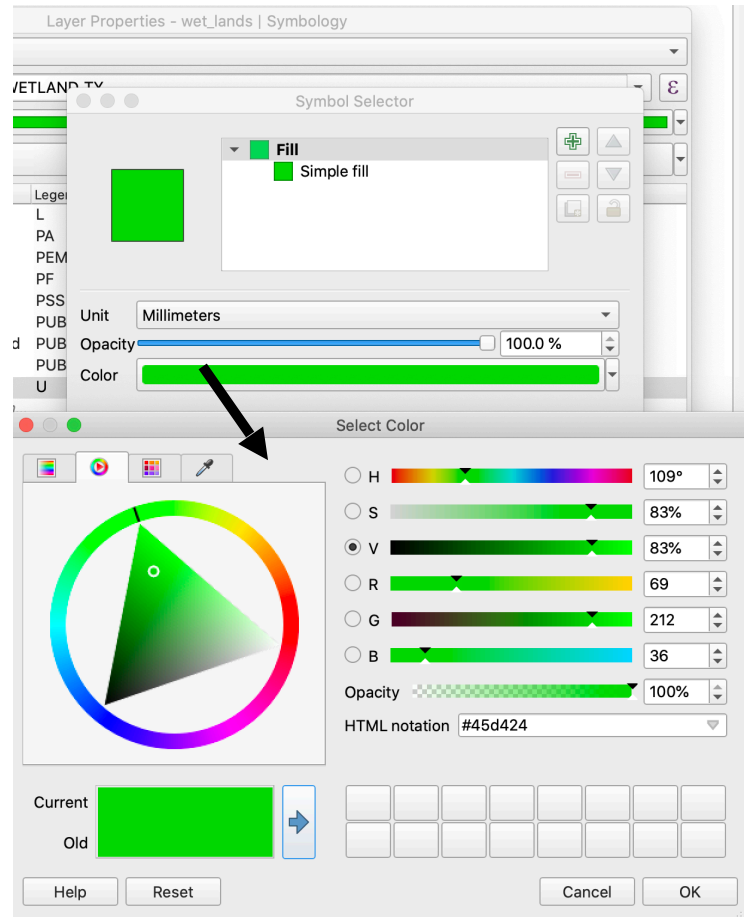
Start by

- 1) click on the Symbology option in the left-hand panel to activate,
- 2) select the color style, here Categories,
- 3) select the attribute column, in this case WETLAND_TYPE,
- 4) click the classify button to list the categories for the active column,
- 5) click the Apply button to color the layer per your settings, and
- 6) click the O.K. button if you're satisfied with the result



If you don't like the colors for a specific category, you can double click on the patch in the color ramp display (see figure above, red arrow) to set the individual color via the Symbol selector. For example, double clicking on the green patch for the "U" category symbol in the figure above would display a symbol selector menu something like the figure at right.

Double clicking on the fill color patch near the middle of the menu (arrow) displays a Select Color palette, through which you can select a different color. After specifying the color, clicking O.K. on the Select Color selector menu, then O.K. on the Symbol selector menu, then Apply/OK on the Layer Properties/Symbology menu to apply the change on your map.



You may hide categories from the display by unchecking their symbols in the Layer Properties-Symbology window (figure left).

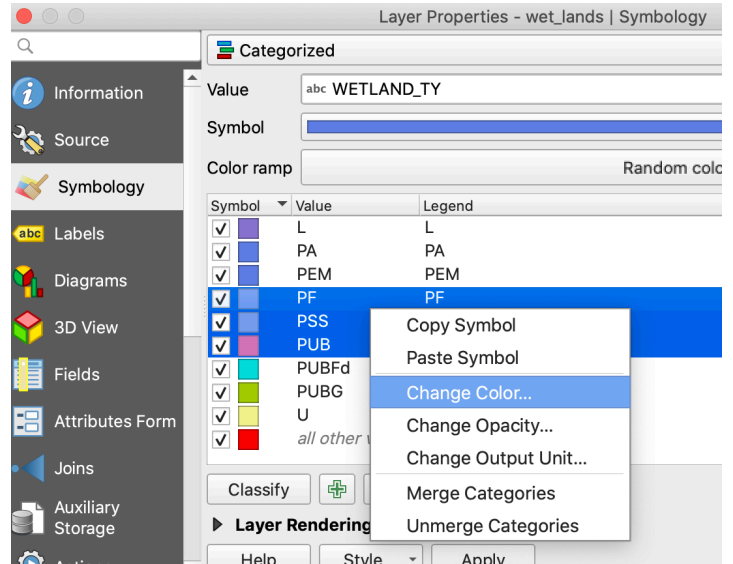
You may delete categories entirely by clicking on the value (e.g., PSS in the figure), and then clicking on the red minus (delete) key at the bottom of the patch window. Holding the shift key and clicking on a first, then a last value in a list will select a range, and control-select (PC) and CMD-select (Mac) will allow you to select multiple, distinct categories. You can batch delete selected items by hitting the minus button. These will then

not be displayed in your data pane.

You can also change multiple categories to the same color by selecting a group, and then right clicking over the selected group and selecting Change Color... from the drop-down menu.

This should open a color palette window as seen above.

After specifying the color and clicking OK, you should get the same color for all the selected categories.




Recolor the wetlands layer using the Properties - Symbology menu, basing the colors on the WETLAND_TY category. Since the U category is the background area, color this with some light color. The L category is lakes, make these a light blue.

Color each of the remaining categories with a distinct color for each of the different wetland types, so that your data look something like the panel on the right.

You don't need to use similar colors, except for the blue lakes, and light-colored uplands, as noted above.

Open the map composer, and add your layers map, a legend, title, scale bar, and north arrow. Note that you will see many images besides a north arrow using the icon referenced a few pages earlier, but select from the available north arrows.

Export the map to a PDF, and save/close this project.



Practice and More Tools

Open a new project, and add the following vector layers:

40_corners

Cl_roads (not the roads.shp file)

streams

Also add the raster layer named Iverson_drg. Do this by left clicking on the Open Data Source icon, then selecting the add raster icon, usually just below the add vector icon in the menu bar:




Clicking on it will open an add layer menu very similar to that for a vector layer. Make sure you set the Iverson_drg.img layer that contains the data, and not the other Iverson files with different file extensions (e.g., not Iverson_drg.rrd.img).

Setting the Layer Order

Notice that as you added layers, each new layer was placed on top of the preceding layers. When you added the Iverson raster, it is covered up the other layers.

This won't do. You can move the raster layer, or any layer, by left clicking on a layer in the table of contents pane, and holding and dragging it up or down. When you let go of the click, the layer is repositioned. Move the Iverson DRG to the bottom of the pile, and note that you can view the previously loaded vector layers. Reposition the points layer to be on top of all the other layers.


Sometimes the entire name of a layer doesn't appear in the table of contents pane (TOC) on the left side of the menu, or we want to resize the TOC. We can do this by "hovering" the cursor over the border line between the TOC and the data panes (see figure at right, the border is the gray band in the circled area).



As you hover over the bar, a resize cursor should appear, two vertical lines with horizontal arrows pointing outwards: ←||→

Left-clicking and holding when this cursor appears "grabs" the edge of the TOC, and lets you move it left or right to make it narrower/wider

Resymbolize the point (corners shapefile) and line features (roads and streams), and pan and resize the extent to make your data pane appear similar to the figure:



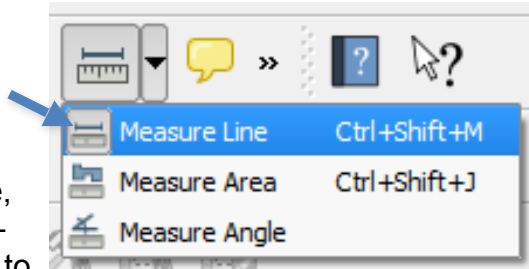
Don't draw the pink circles, those are added to help learn the measure tool, described below.

Tools used to change symbology for the line and point features are similar to the polygon symbol tools discussed above. With lines, the menu includes a line width option (set to 1 here for the roads, and 0.5 for the streams),

Note the gravel pit and lookout tower circled in pink in the image above. You may have to zoom in to see these features on the Iverson.drg map, the gravel pit is shown as a bare area with a faint X to the right of the label, and the tower is a transparent circle with a dot in the middle, to the left of the label. Notice that if you zoom in too far, the image appears pixelated, or blocky. If so, zoom out a bit, it will appear a bit better.

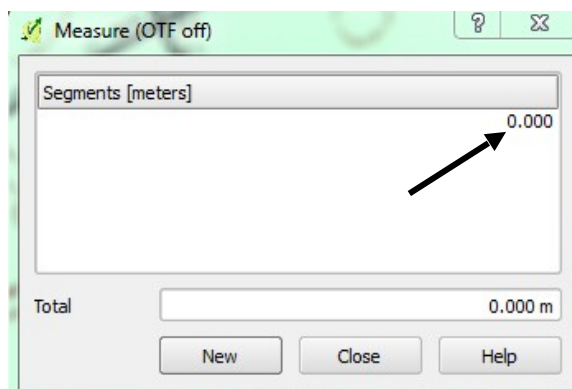
Pan and zoom until the map provides the highest definition without too much pixilation.

Use the measure tool, usually parked at the top middle-right of the main menu, to measure the distance between these two features. Left click on the triangle to the right of the tool, and from the drop-down menu select the Measure Line tool. Note that you can start this tool with a shortcut keyboard sequence, Ctrl+Shift+M on a MS Windows computer, and Command+Shift+M on a Mac. There are many shortcuts to learn, although in general we won't introduce them to not create too much detail.



With the line measurement tool selected, position the cursor over the gravel pit, then left click. You should see a measurement window:

Note that as you move the cursor to another location, the number that starts at 0.000 in the window will rise, reporting the distance for that segment, and the total in the bottom window will also increment upward.

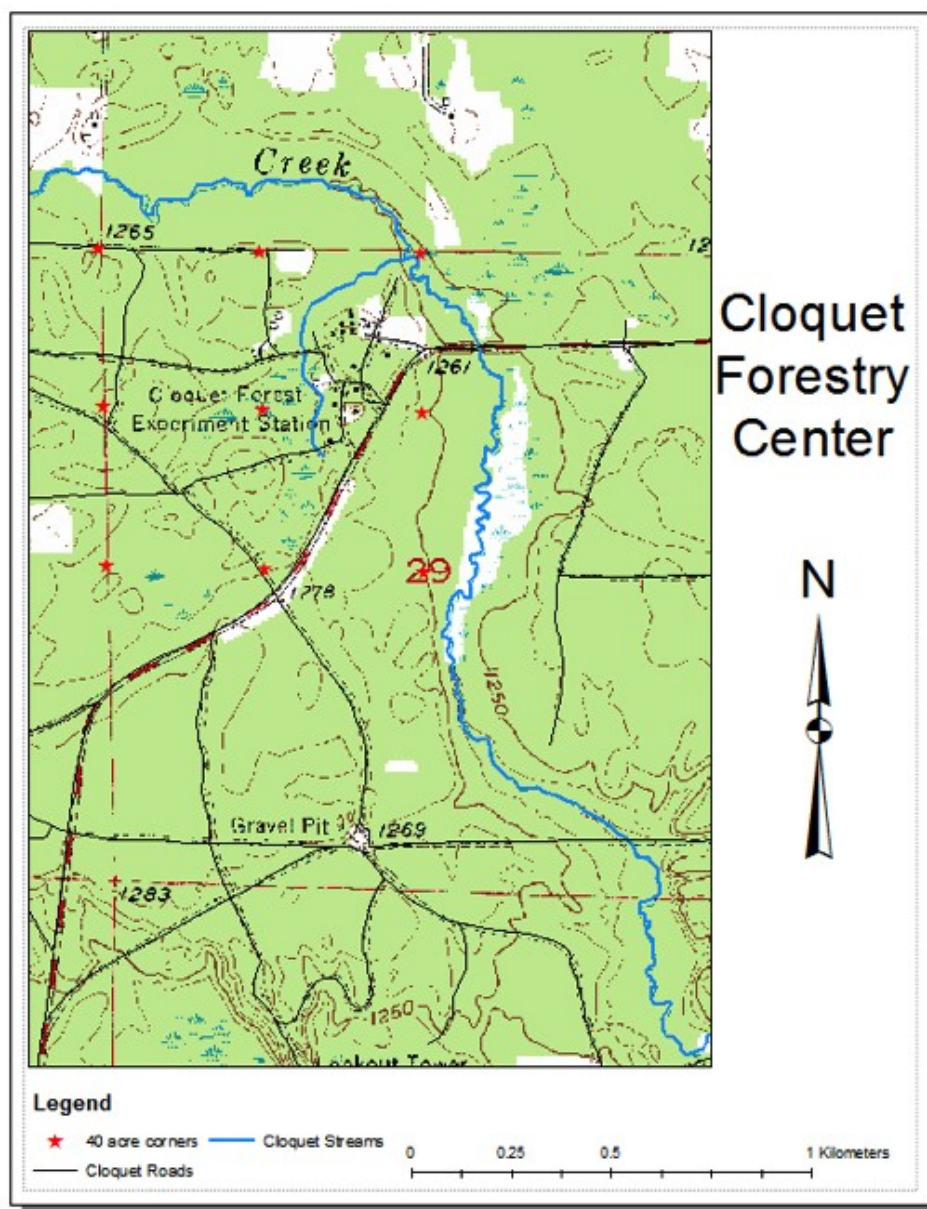


Move the cursor to the lookout tower, and click. Note that another number appears below the first, it should start at zero, but may be a small number if you move the cursor slightly.

Note this length, the distance between the two features. The units are those of the layer, meters in this example. We can set them to project units, but we'll save that until after you've learned about map projections.

If you had a multi-segment distance to measure, you would then move and click along the segments, noting the length of each segment, and the increasing total sum in the bottom window. You could start another measurement via the New button in the bottom. Close the window.

Create a map with a title, north arrow, legend, title, and description, and export a PDF as before.



Multiple maps on One Page

Sometimes you wish to show two maps on the same page, e.g., a larger overview map and an inset map. We do this in QGIS by:

- 1) loading all the data we wish to display in to a QGIS project,
- 2) setting the symbology for the first map we wish to display
- 3) opening a new map composition and setting the window, scale bar, and any other map elements you wish
- 4) “locking” the display elements for the map in the Item Properties (remember, tabs in the rightmost pane of the print layout), under layers, lock layers, and lock styles for layers check boxes.
- 5) going back to the QGIS Data/Map window and changing the zoom, pan, and symbology to that of the second map you wish to display
- 6) returning to the existing map composition, and adding a new map. This new map can have the same data as the first map for an inset, or you can display entirely different data, simply turning a layer on or off in the table of contents.

See the **video “Simple Inset Map”** in the QGIS 3.x Lab1 playlist on the Mapeez YouTube channel.


Note that the changes in the second map in your layout may not appear on the print layout until after you make it active, and perhaps zoom or pan at least a tiny bit.

As an exercise, start by creating a new QGIS project.

Load the Iverson DRG, and then the vegetation.shp data from the L1 data directory.

Apply the above steps to create a two-panel inset map similar to the map shown, with the inset map in one single color, and the large map a categories map colored by the values of cover_type.

Watch the video on inset maps if you have difficulties. Make sure the larger map is active when you add a scale bar.



Note that the inset and main map are usually at different scales, as here, they will have dif-

ferent scales scale bars. You only need to add a scale bar here for the larger map, and it is best to set the final zoom for that panel and then add the scale bar for the panel before you lock it. An added scale bar always applies to the currently active map.

Topologies

QGIS supports topologies in a rudimentary way. As described in the readings, topologies are the spatial relationships between objects, such as “overlaps,” or “contained within.” When we check for or enforce topologies, we’re trying to ensure correct spatial relationship, e.g., we might want to make sure all our digitized power poles fall within our block or parcel polygons, and not on our street polygons, or that all our river networks connect, or don’t have overshoots or dangling arcs. Many spatial operations assume or require correct topology, and will give wrong answers if there are topological errors.


Start by loading the basins.shp shapefile. It contains the watershed outline for the Red Cedar and Lower Chippewa rivers in western Wisconsin ([video](#), [Intro2Topology](#)).

Next, find and click on the topology checker icon, it typically is placed on the main menu frame:




This should open a pane to the right of your data window, depending on how you have your system configured. You may not initially be able to find the topology checker icon, though, and so you may have to enable the topology checker before you use it.

The topology checker is a “plug-in,” a contributed extension that represents one of the strengths of open source projects, because many developers can contribute to extending QGIS. It is a core plug-in, so can’t be un-installed. You can check to see if it is activated by clicking on the Plugins label in the main QGIS bar at the top of the screen or primary window (Mac example, right), then entering “topology” in the search bar. This should display the Topology Checker plugin, and you should check it if it is not checked. Then go back to the base QGIS window, and scour the icons to find the topology checker.




Once you find the topology checker (icon at right), click on it to open the Topology Checker Panel. To create rules, click on the manage rules button (the wrench, arrow in figure below).




This should open a rule-creation menu, as shown at right:

You create rules through dropdowns in the three buttons along the top, and then click on the Add Rule button to build your list. If you make a mistake or wish to change a rule, you can use the Delete Rule button and then add new rules to the list.




For example, if I click on the left-most button where it initially lists “No layer,” I’ll get a drop-down menu of the currently displayed layers. I can then select basin, and two buttons will be displayed immediately to the right of the basin button. Clicking on the middle one will give a dropdown list of rule choices. I can select “must not overlap” to ensure the polygons don’t overlap each other in a layer. If I then click the Add Rule button, it will place the rule in the pane below the buttons, and should look something like the figure to the right.



I can add several rules to be checked at once, depending on the layers I’ve displayed and the rules I want to enforce. I find it clearer to only check the topology on one layer at a time, rather than jumble up the display with rules and errors from multiple layers. It depends in part on how many errors you expect per layer.


You then check for errors by clicking on either the Validate button, which checks for errors in the entire data set (arrow at right).



You don’t need to fix the errors now, we’re just introducing the tools here. We’ll focus on them in a later lab.

An error should show up in the main topology checker pane, similar to that shown at the right. Clicking on the error will highlight the “offending” polygon, although in this case, both polygons overlap each other, so the choice it highlights is rather arbitrary. You’ll notice a border overlap in the lower left portion of the green polygon below.

Error	Layer	Feature ID
0 overlaps	basin	0




Typically, you would note where the errors are, and depending on their type, edit them manually or apply some sort of function to remove the errors. We’ll cover editing in a later lab, for now just note the error identified by topology.

Remove the basins data layer, delete the rule from the rule settings (click on a rule, then the Delete Rule button in the Topology Rule Settings window), then close the window.

Add the rivers.shp file to your canvas. Open a new Topology Checker Panel, and add a rule for the rivers data layers, specifying that it must not have dangles.

Save and then Validate the rule for the entire extent. You should see a list of errors appear in Topology checker. Red dots appear in the data pane, if you zoom to full extent you should see something like the figure to the right. Each dot represents a dangle. Some of these dangles are acceptable, e.g., where a river starts, but there should be no dangles where two rivers join. How many of gap errors where rivers join can you find?



Typically, we would edit to fix the topological errors, something described in a later exercise.

That’s all for this exercise, we’ll be using most of these data display and map symbolization skills in the rest of the exercises in this series.