Mapping Census Data

Before Beginning your Map

Before you start to build your map(s), you need to decide three things:

- Which general Census topic are you going to map?
- At what level of geography will you map your variable(s)?
- What geographical area do you want to map?

Which Census Topic?

The Census is arranged into seven broad topics:

- Aboriginal
- Age and Sex
- Ethnic Origin and Minorities
- Income and Earnings
- Labour, Occupation, and Education
- Language, Immigration, and Citizenship
- Marital Families and Households

Decide which of these broad topics relates best to the variables you want to map.

What Level of Geography?

Census variables are aggregated into nine geographic levels. In order of increasing size, they are:

- Dissemination Area
- Census Tract
- Designated Place
- Urban Area
- Census Subdivision
- Census Metropolitan Area
- Census Division
- Federal Electoral District
- Forward Sortation Area

Click on each geographic level for an explanation from Statistics Canada. Decide which geographic level will help you best present the variables you want to map.

What Geographical Area?

Finally, decide the geographical area you want to map – are you interested only in the Region of Waterloo? Southern Ontario? All of Ontario? All of Canada?

Building Your Map

To begin building your map, open ArcMap.

Add to the ArcMap layout the Census dataset which fits your chosen topic and geographic level. See Map Library staff for assistance in obtaining the correct dataset(s).

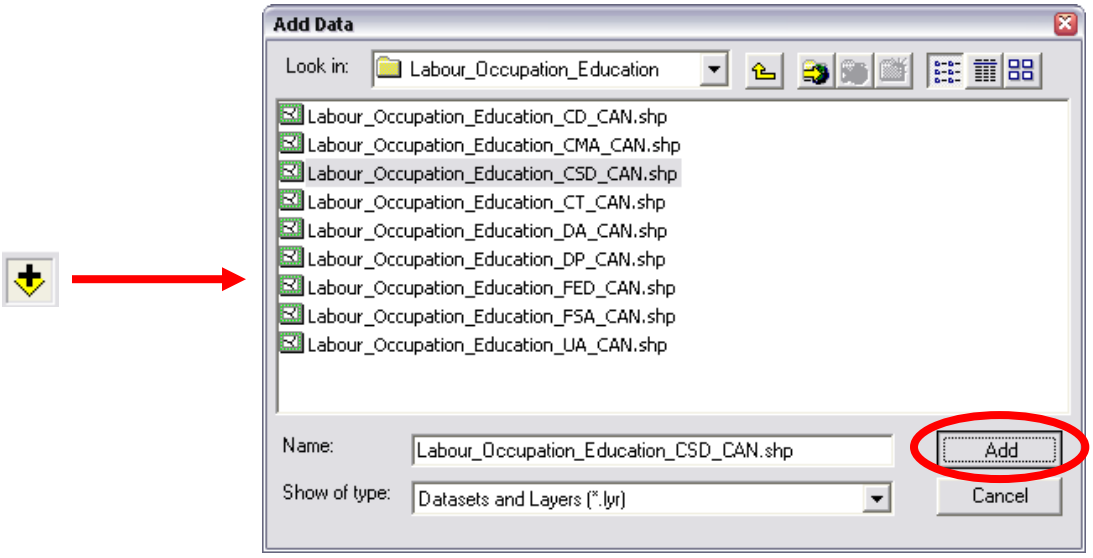

Open the attribute table of the dataset: right-click on the layer in the Table of Contents at the left of your screen, and click 'Open Attribute Table':

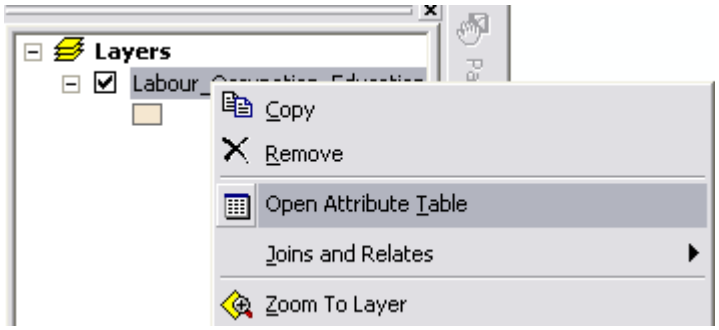

Examine the fields in the attribute table, and determine which field contains the Census variable you want to map. Make a note of ***which single attribute*** you wish to map.

Note: In many cases, the field names in the Shapefiles can be difficult to decipher, as they are frequently truncated, abbreviated, or otherwise shortened. Refer to the appropriate *Beyond 20/20* table to determine what each field name represents (See the following page for more information).

Using Beyond 20/20 to Decipher Field Names

Using Windows Explorer, browse to the location of the Beyond 20/20 file that matches your chosen Census topic and level of geography. Open the file.

At the same time, in ArcMap, open the Attribute Table of the Census dataset you chose previously, **or** open the Layer Properties window by double-clicking the layer in the Table of Contents, then switching to the 'Fields' tab.

In Beyond 20/20, scroll down until the fields in Beyond 20/20 begin to match up with the fields in ArcMap. This may take some guesswork, and you may need to skip over the first few fields in ArcMap.

Once the fields match up, you can use Beyond 20/20 as a reference to determine the meaning of each field:

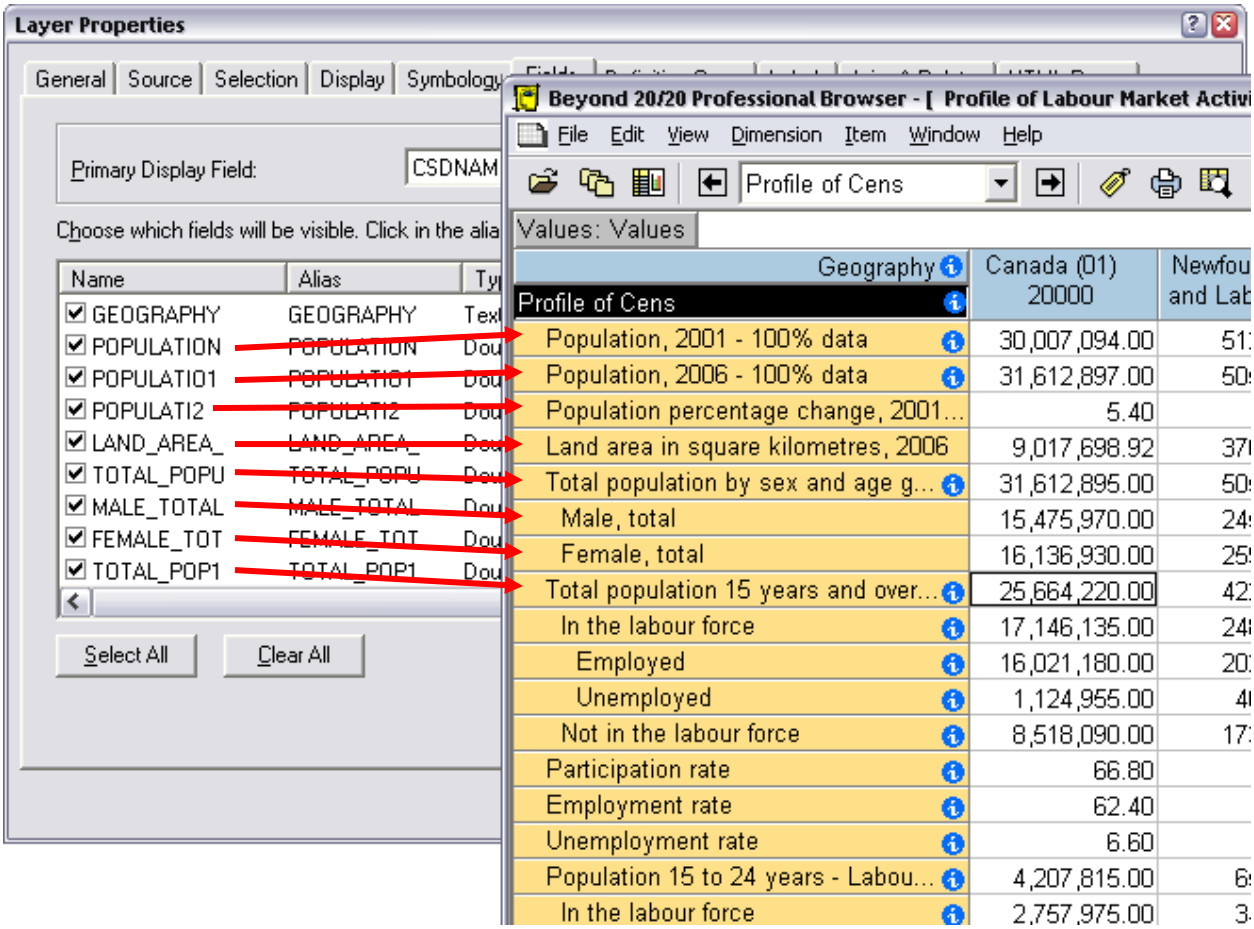

Symbolising the Census Data: Single-Variable Mapping

First, A Review

At this point, you have:

- Chosen the Census topic and geographic level you wish to map,
- Determined which geographic area you want to map,
- Added the appropriate Census dataset to ArcMap,
- Reviewed and deciphered the field names, and
- Decided which single attribute you wish to map.

And now, Symbolising the Data

Open the Layer Properties window by double-clicking the layer in the Table of Contents *or* by right-clicking it and selecting 'Properties.' Switch to the 'Symbology' tab.

At this point, the Census data is symbolised using an arbitrary colour chosen by ArcMap:

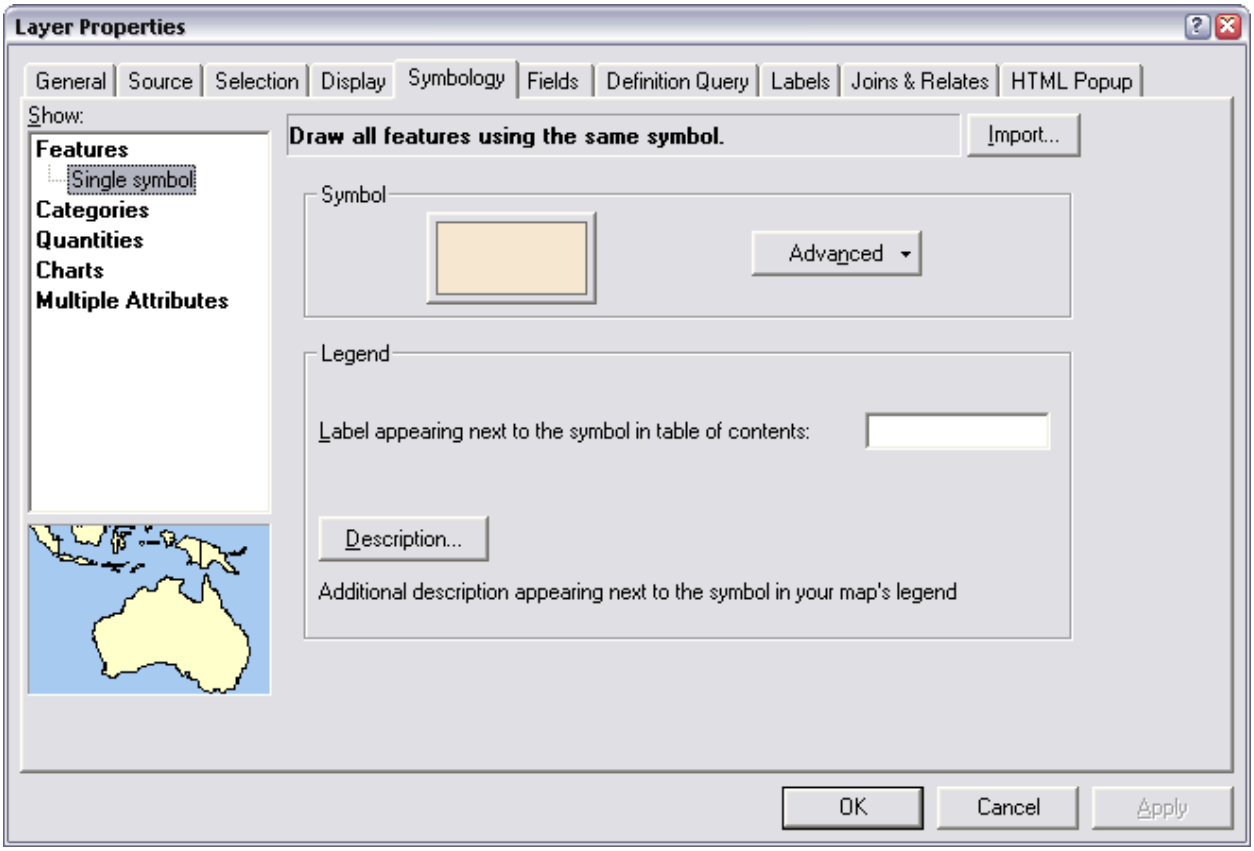

To make a meaningful map, we'll need to change the colours so that they illustrate the Census variable appropriately.

At the left of the Layer Properties window, click '**Quantities**', and then ensure 'Graduated colors' is selected. The Graduated Colors symbology style allows you to define a number of classes according to the Census variable, and then assign a different colour or shade to each class.

From the 'Value' dropdown box, select the field you want to map. The example below will ultimately result in a map illustrating levels of unemployment among males aged 15 to 25, by Census Subdivision.

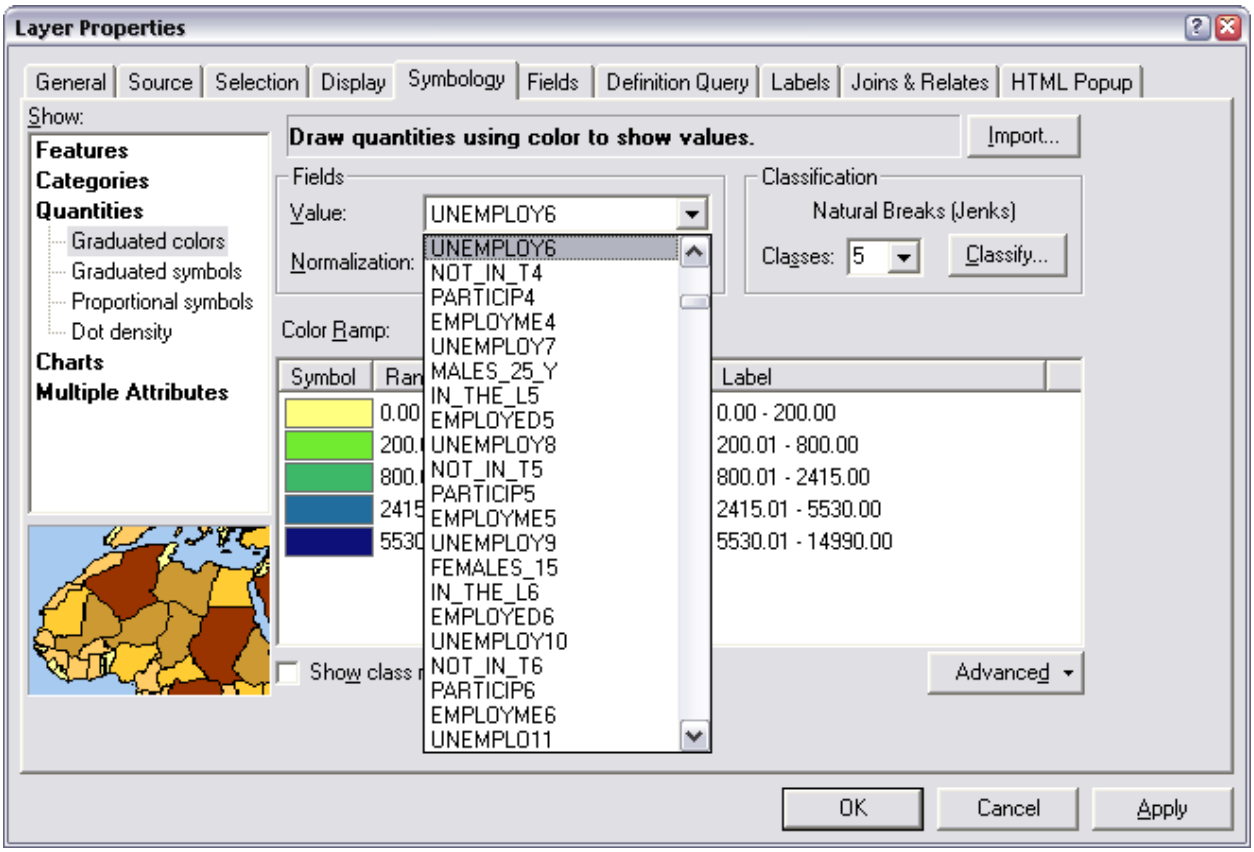

Take a quick look at the Class Ranges which result. Are they useful?

Chances are, they aren't. Is it useful to produce a map which shows that Census Subdivision 'X' has 200 unemployed males, while Census Subdivision 'Y' has 14,990 unemployed males?

It would be far more useful to create a map which shows that 14% of males in Census Subdivision 'X' are unemployed while 27% in Census Subdivision 'Y' are unemployed.

The problem, then, is that many Census variables are expressed in absolute terms – the number of persons in a given geographic area who fit a particular criterion – but, for comparative purposes, we need to display those variables in terms relative to the population of each geographic area.

This is where the 'Normalization' dropdown box comes into play.

Normalisation is the process of dividing one numeric attribute value by another in order to minimise differences in values based on the size of areas or the number of features in each area. For example, dividing the number of persons aged 18-30 in a given area, by the total population in that area, yields the percentage of 18-to-30-year-olds in that area. Similarly, dividing the total population in a given area by the geographic size of that area yields a density figure. (Adapted from the ESRI Desktop Help file).

When should you normalise your data? If your data is in raw counts that are expressed in large numbers.

When should you *not* normalise your data? When the data is already normalised.

When normalising your data, take special care that you are using the correct population or area value as a normalisation value. Using the wrong population value will result in an erroneous percentage or density value.

Consider the example below. If one were to create a map of the 'Unemployed' variable displayed below, it would be nonsensical to normalise that value by the total population, or by the total number of females, in the geographic area. It is similarly nonsensical to normalise that value by the 'Males 15 to 24 years – Labour force activity' variable – since that variable also includes those who do not or cannot work (for example, males under the age of 16).

Instead, the 'Unemployed' value should be normalised by 'In the labour force' to yield a logical unemployment rate among males 15 to 24 years of age.

Males 15 to 24 years - Labour force activity	2,145,565.00	33,210.00	17,085.00	1
In the labour force	1,407,685.00	17,085.00	9,555.00	
Employed	1,219,815.00	11,985.00	7,205.00	
Unemployed	187,870.00	5,100.00	2,345.00	
Not in the labour force	737,880.00	16,125.00	7,535.00	1

In ArcMap's Layer Properties window, select the appropriate field from the Normalisation dropdown box.

Examine the new class ranges. Are they appropriate and adequate for your purposes?

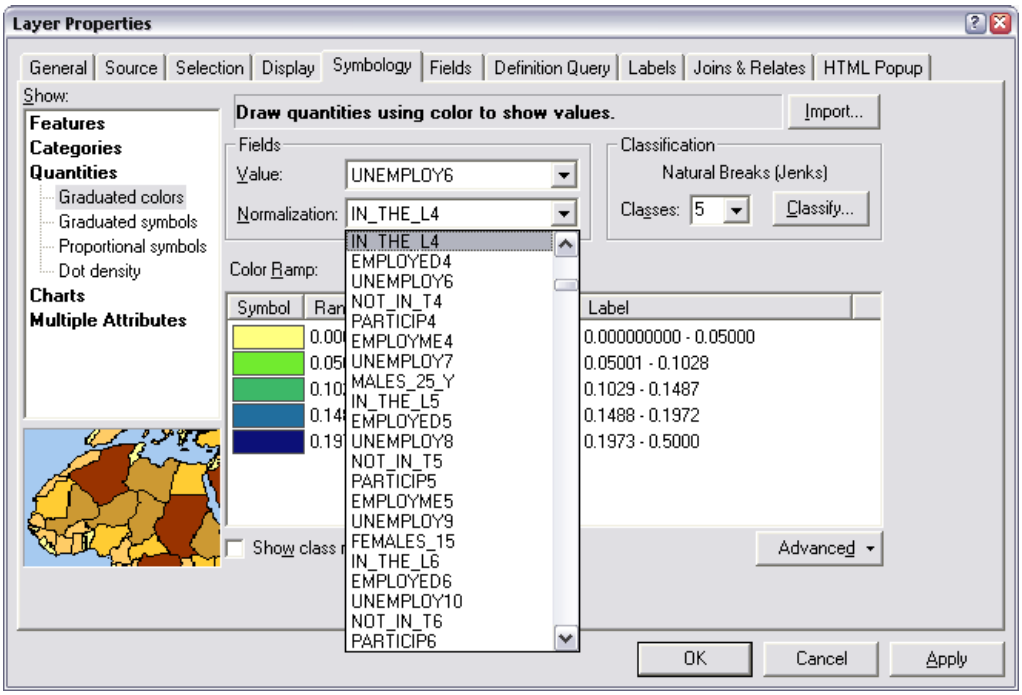

Often, one or more of the following problems will arise when normalising data:

- The breaks between classes tend to be arbitrary, which can make maps difficult to interpret;
- Unpopulated areas – that is, those with a normalisation value of '0' – are not drawn; and
- Statistics Canada randomly rounds values either up or down to either 5 or 10 for confidentiality purposes. This means that, theoretically, in an area with 12 males between the ages of 15 and 25, 7 of whom are employed and 5 unemployed, Statistics Canada could report that 10 males reside in that area, 10 of whom are employed and 10 unemployed. This rounding presents problems when trying to normalise and map areas with low population figures, and may result in normalisation values of greater than 100%.

Each of these problems will be dealt with below.

Making Class Breaks Logical

By default, ArcMap symbolises quantitative data using 'Natural Breaks,' thereby placing class breaks between adjacent feature pairs with relatively large differences in data value. In most cases, the 'Natural Breaks' classification method produces arbitrary class breaks which are not useful or logical.

To make your map more readable, consider using one of the following five classification methods:

Equal Interval

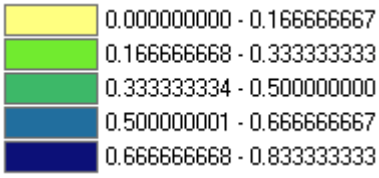

The class breaks for Equal Interval classifications are determined by dividing the difference between the maximum and minimum values by the number of classes. In the example above, the interval used is 16.67% or 0.1666(...). Equal Interval classifications may result in an imbalance in the number of features in each class, and the classes can be difficult to interpret.

Defined Interval

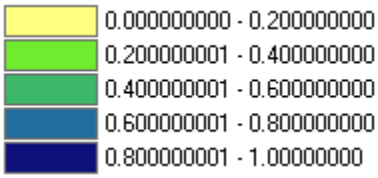

The class breaks for Defined Interval classifications are set by the user. In the example above, I set the class break at 20%, or 0.20. Defined Interval classifications may result in an imbalance in the number of features in each class, but the classes are very easy to interpret. This is a very common means of classifying quantitative data.

Quantiles

0.000000000	Lower Quintile
0.000000000 - 0.078711986	Second Quintile
0.078711987 - 0.120000000	Third Quintile
0.120000001 - 0.200000000	Fourth Quintile
0.200000001 - 0.833333333	Upper Quintile

The class breaks for Quantile classifications are based not on data values but on the percentage of features in each class. The Quantile classification is dependent on the number of classes – 4 classes produces quartiles, 5 classes quintiles, and so on. Quantiles are one of the best ways of symbolising linearly distributed quantitative data on maps, and classes are very easy to interpret. Care should be taken not to use too many classes, lest readability suffer. If the data is normally distributed, use Standard Deviations (below) instead.

Standard Deviations

0.000000000 - 0.022702839	< -0.50 Std. Dev.
0.022702840 - 0.146860806	-0.50 - 0.50 Std. Dev.
0.146860807 - 0.271018773	0.50 - 1.5 Std. Dev.
0.271018774 - 0.395176739	1.5 - 2.5 Std. Dev.
0.395176740 - 0.833333333	> 2.5 Std. Dev.

The class breaks for Standard Deviation classifications are based on, of course, the standard deviation or divisions thereof of the selected attribute. Standard Deviations are the best way to symbolise normally distributed quantitative data on maps. Classes are easy to interpret for those with some knowledge of statistics but may be inappropriate for other audiences.

Classifying Your Data Appropriately

Experiment with the various classification schemes until you find one that works for you and for the data you are using. To change the classification scheme, click the 'Classify' button while in the Symbology tab of the Layer Properties window:

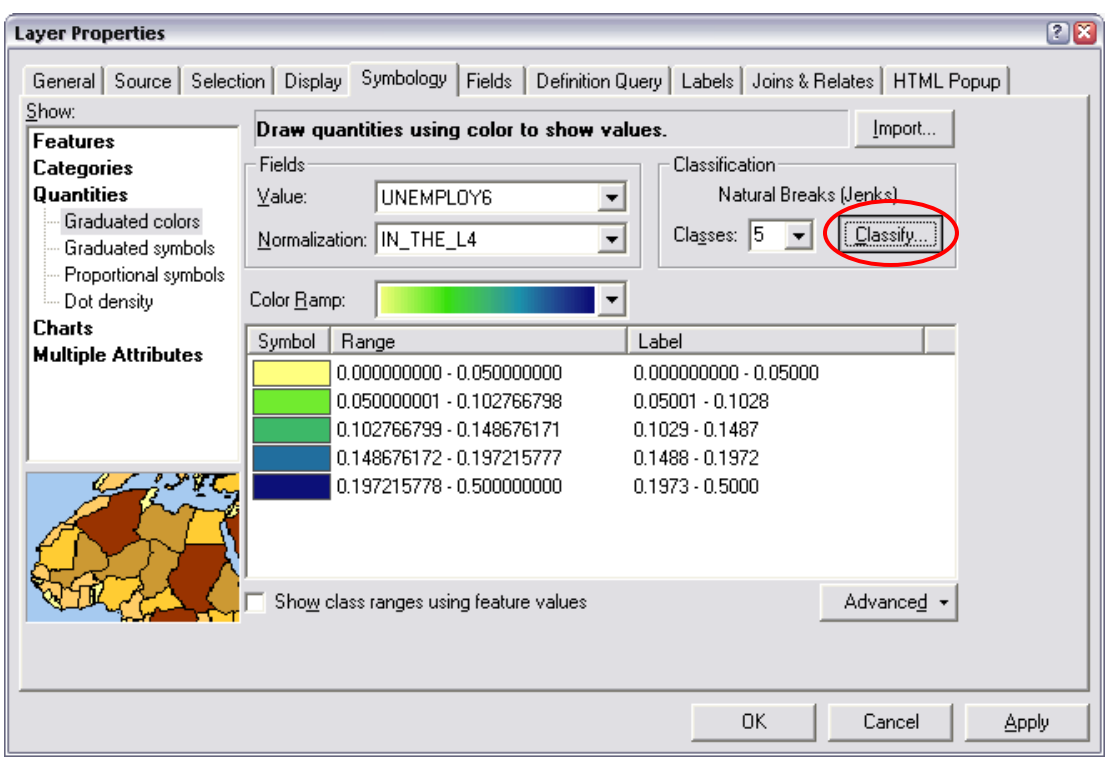

Choose the desired Method from the dropdown box in the Classification window.

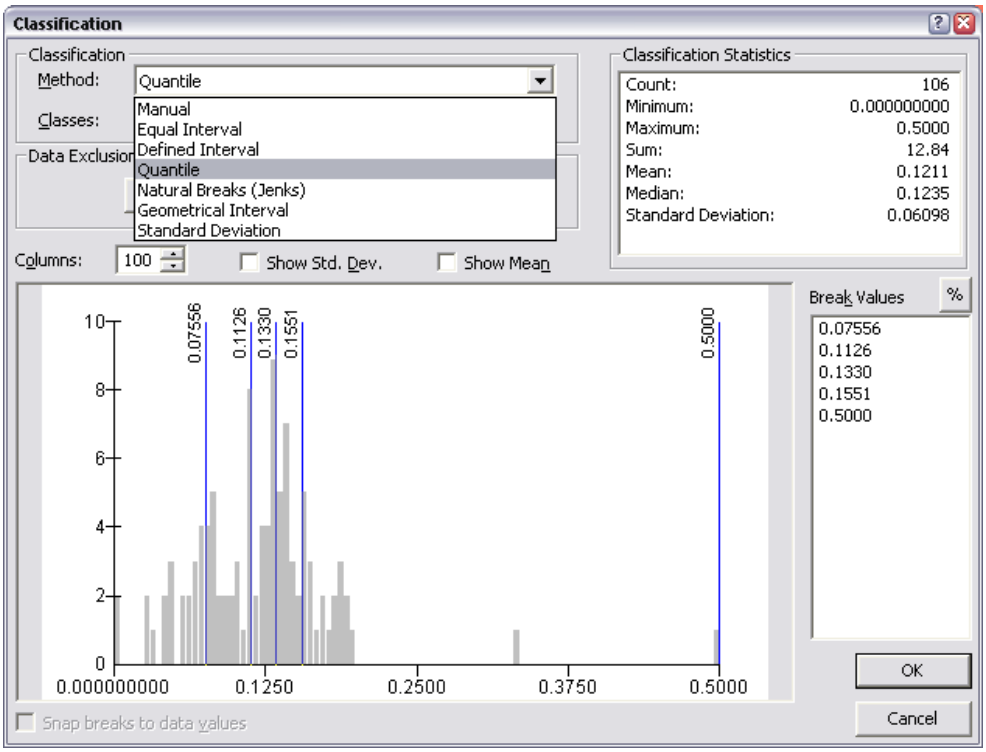

If necessary, change the number of classes (note: the number of classes is set automatically when using Defined Interval and Standard Deviation classification methods):

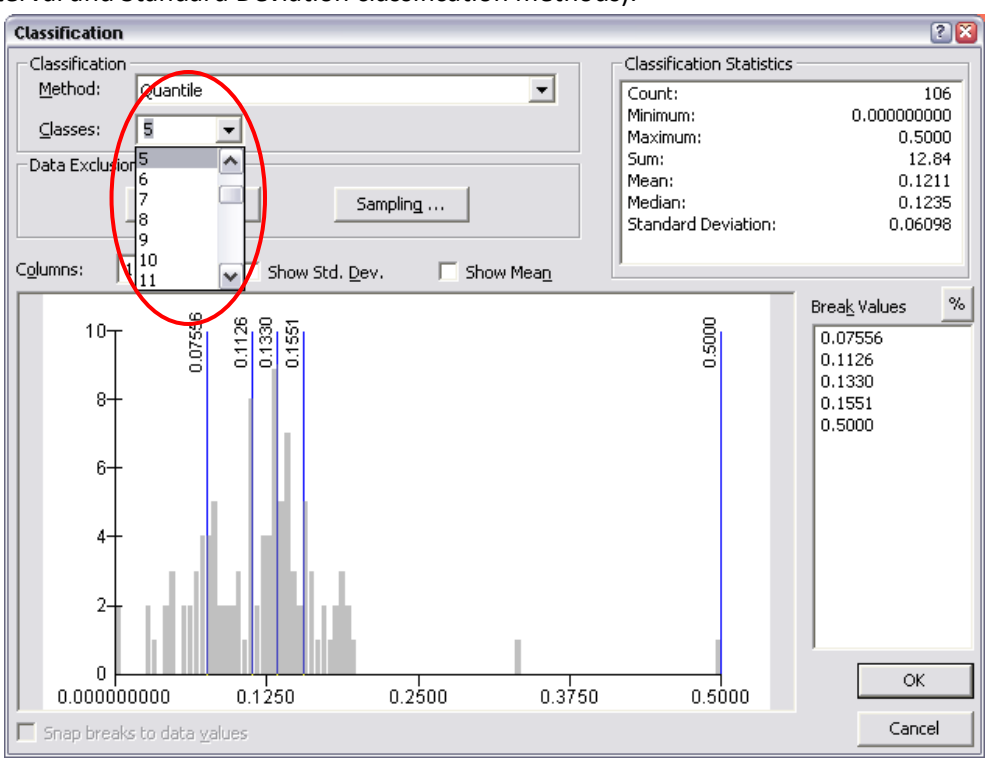

If you are using the Defined Interval classification method, set the interval:

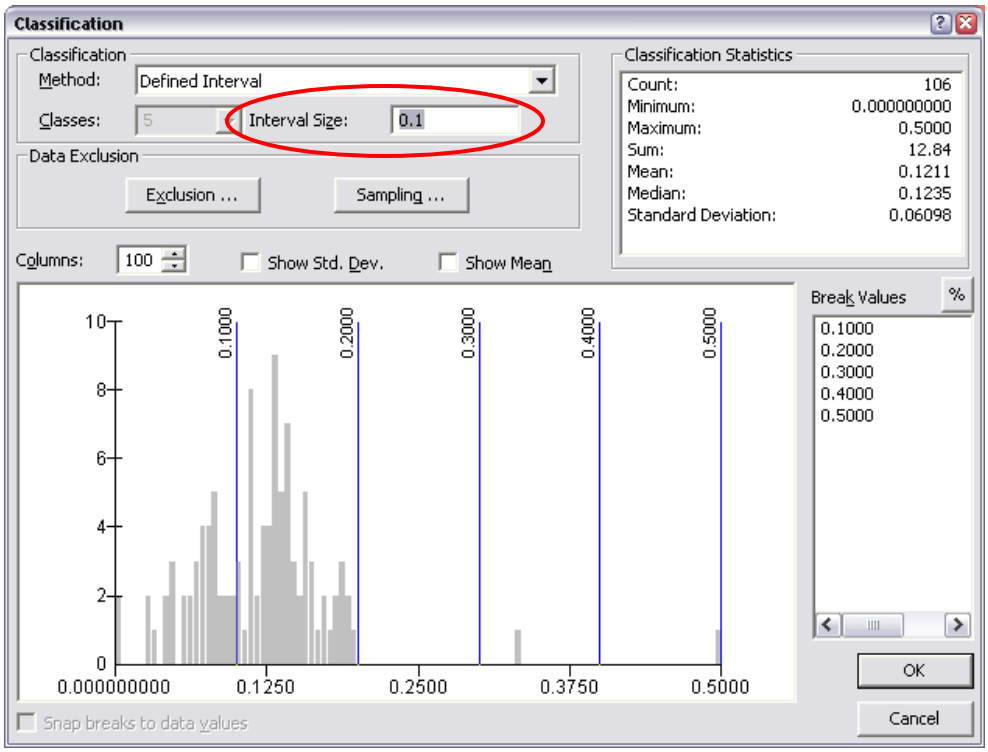

If you are using the Standard Deviation method, set the division of standard deviations to use:

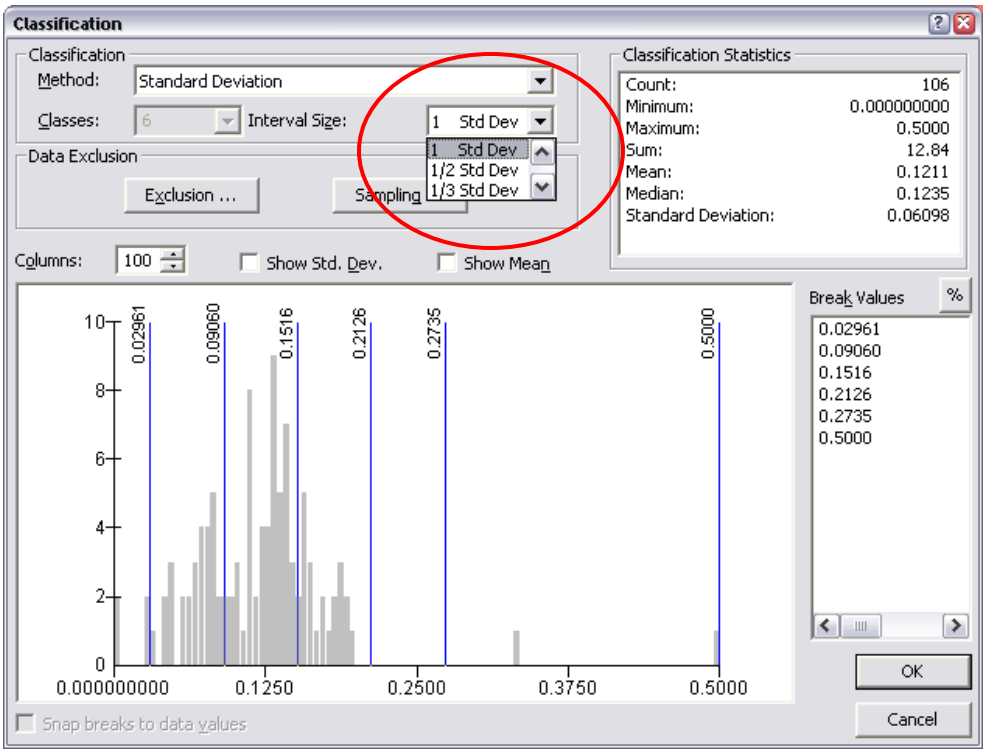

Setting Exclusions to Avoid Problems Created by Rounding

As noted earlier, Statistics Canada randomly rounds its statistics either up or down to the nearest 5 or 10. This can create problems whereby a given value is larger than the normalisation value, thus yielding a normalised percentage of over 100%. To avoid this problem, open the Classification dialogue box and click the 'Exclusion' button:

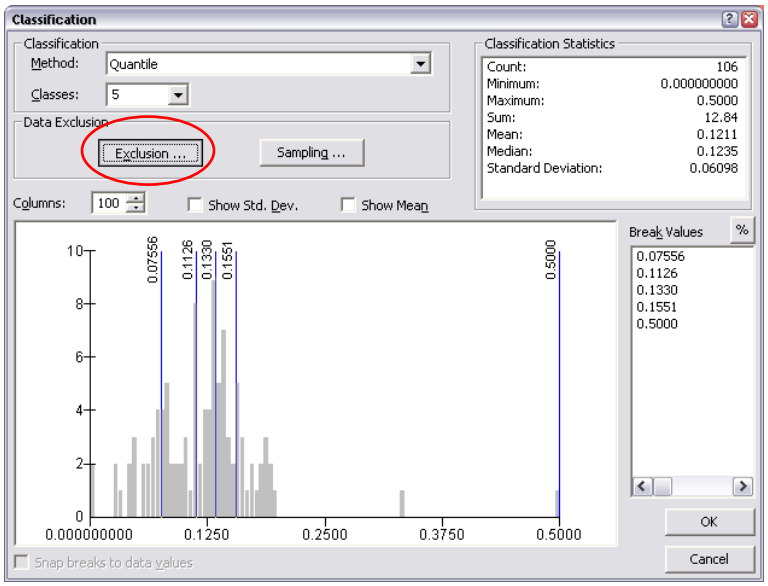

In the Data Exclusion Properties window which opens, scroll through the field list until you find the field you used to fill the 'Value' dropdown box earlier. Double-click it, then enter the '>' sign. Next, double-click the field you used to fill the 'Normalization' dropdown box. Click 'OK' to exit the Data Exclusion Properties window, and 'OK' again to exit the Classification window.

Any feature which satisfies the expression criteria – where the value is greater than the normalisation figure – will *not* be symbolised.

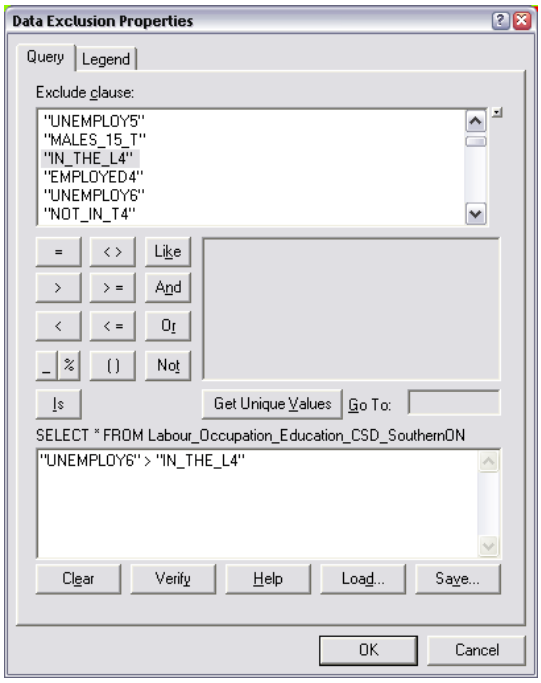

Changing Colours

Once you have selected an appropriate classification scheme, experiment with different colour ramps until you find one that best fits the number of classes and the type of classification scheme you are using.

To change the colour ramp, click the dropdown box denoted below and scroll through the options:

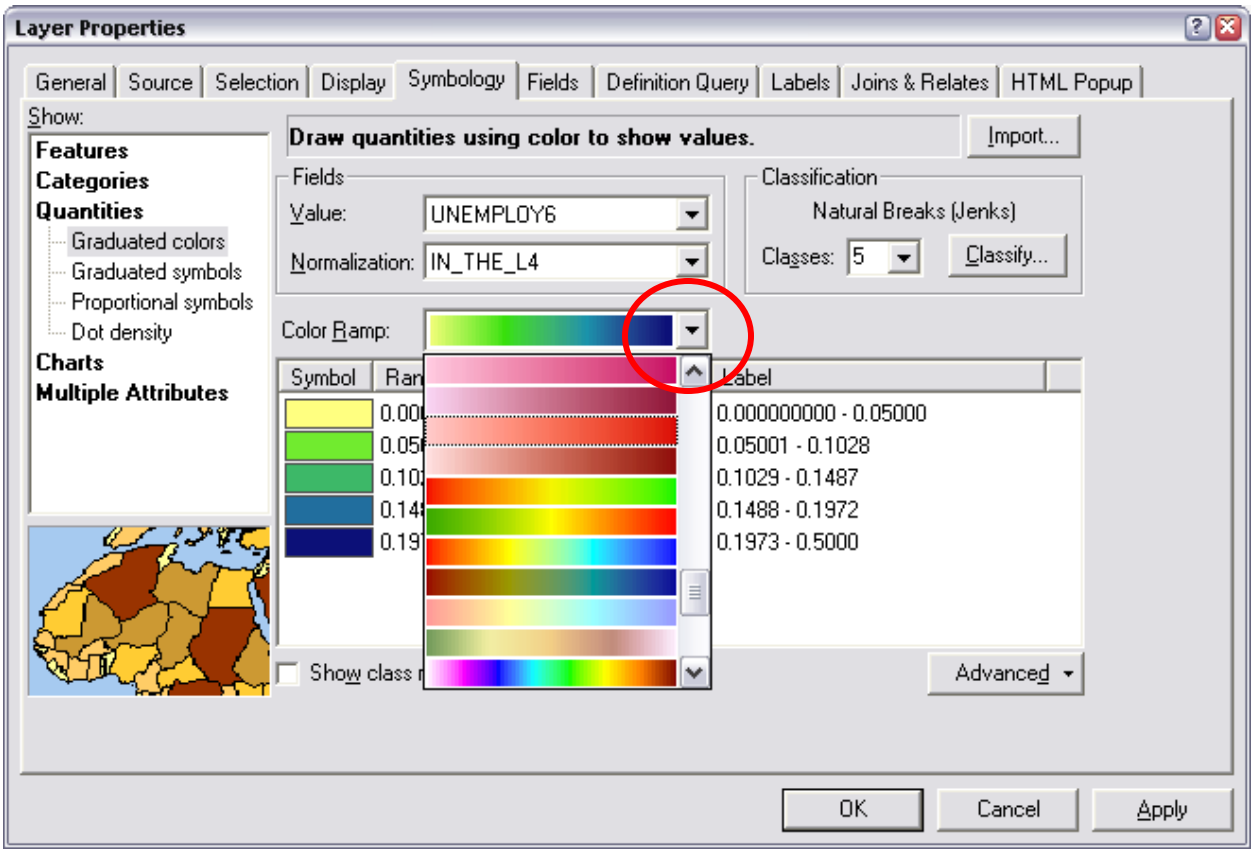

Some tips on selecting good colour ramps:

- For *Sequential* classification schemes such as Equal Interval and Defined Interval, use a colour ramp that goes from a light shade, for low values, to a dark shade, for high values, of the same colour, as used in the example above.
- For *Diverging* classification schemes such as Quantiles and Standard Deviations, use a colour ramp that grades between two contrasting colours with a neutral colour at the mean or median. Sample diverging colour ramps are shown below:

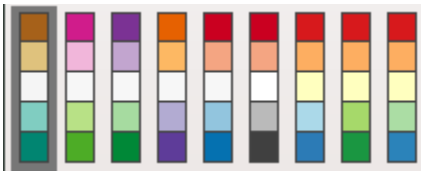

Try experimenting with various colour schemes using ColorBrewer2, an online tool created to aid cartographers in selecting colours for maps: http://colorbrewer2.org/

You may also wish to change the border colour, width, and style of each feature. Rather than doing so individually for each class, simply right-click on a class range and click 'Properties for All Symbols...'

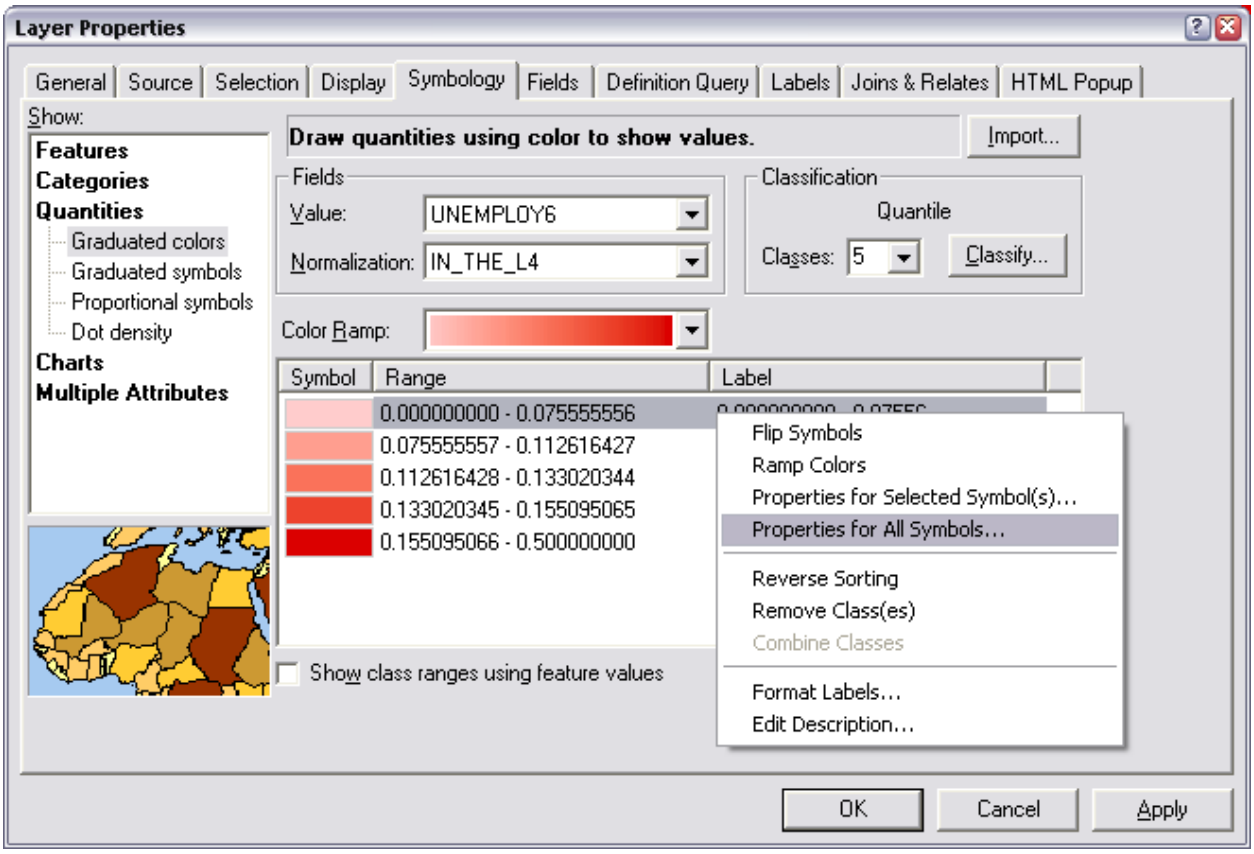

In the Symbol Selector window, change the Outline Width and Colour settings as desired, then click 'OK.'

Don't change the Fill Colour in the symbol selector – if you do, you'll have to re-apply your colour ramp!

Changing Legend Labels

Note: This should be done only when classification schemes and colours have been finalised. Any change you make to classifications will reset legend labels to their default value.

By default, the labels for each class are simply the class ranges:

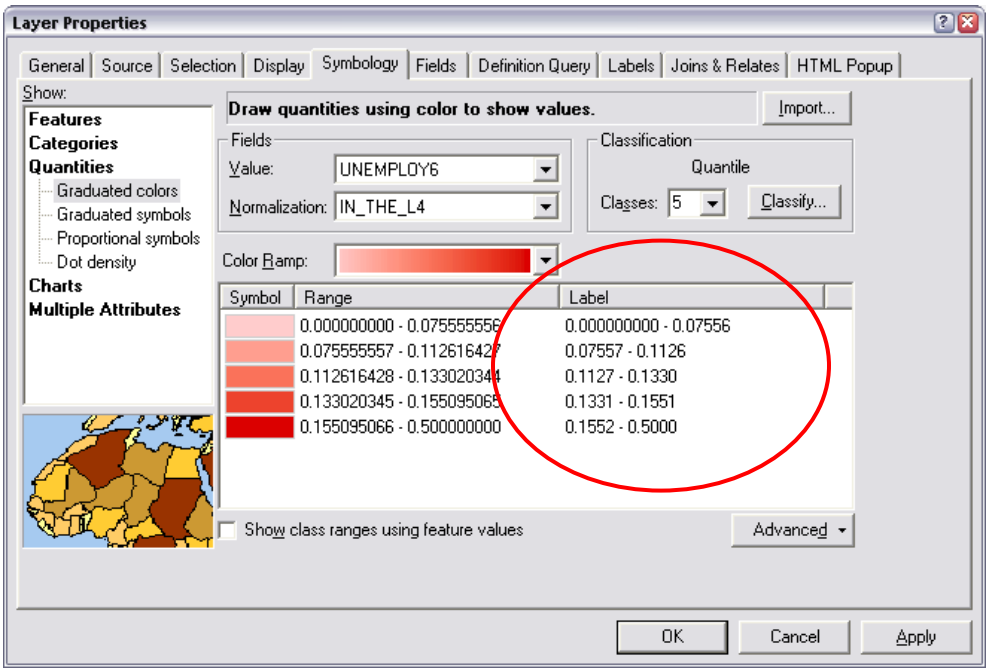

For readability, you may wish to change these labels. To do so, click once on the label text so the text is highlighted, and simply type in a more understandable label:

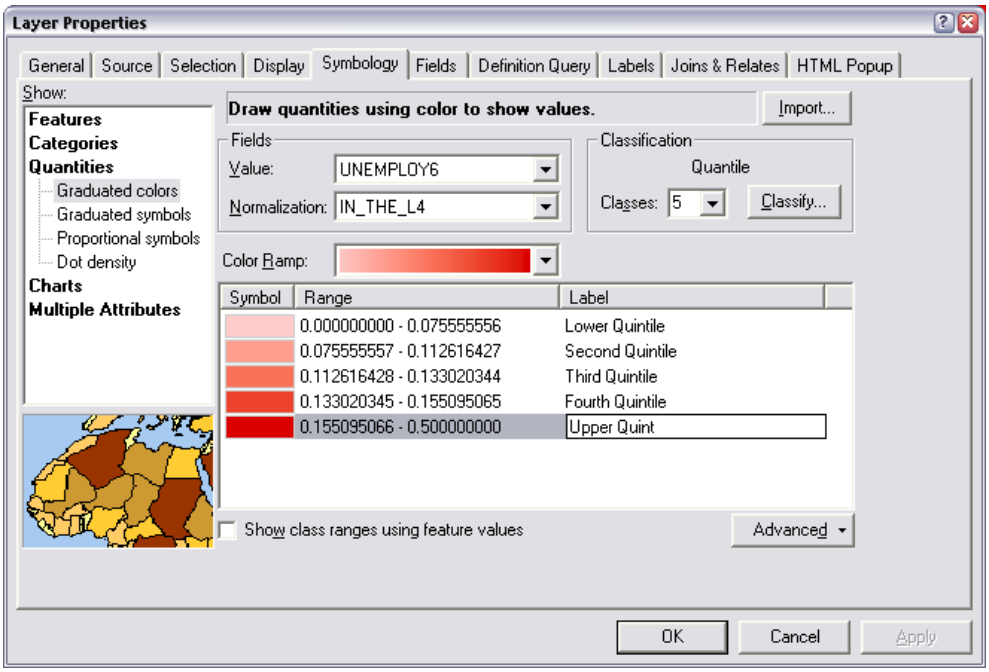

Verifying the Symbology

Once you have finished setting your symbology, click 'OK' to close the Layer Properties window and return to your map. Have a close look at the map, and make any changes to the symbology that might be necessary to make the map clear, understandable, and informative.

Zoom in to the geographical area of interest – be it Southern Ontario, all of Ontario, or all of Canada, and again ensure the symbology is appropriate.

In some cases, such as that presented below, you may find that the symbology needs to be adjusted when you are zoomed in to a small area. In this case, a Defined Interval classification of 20% was used:

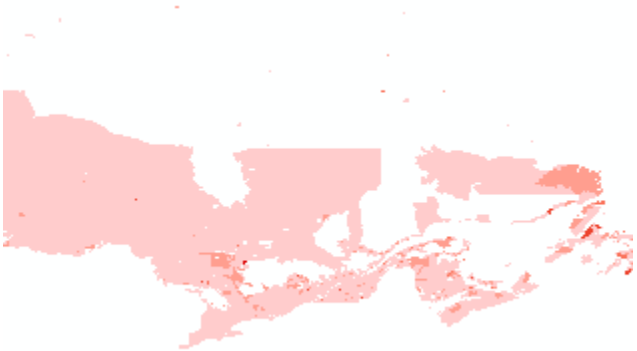

Evidently, most of Canada, and all of the Greater Golden Horseshoe, falls within the 'Under 20%' range. The solution here is to use a different classification scheme or reduce the class interval to 15% or 10%.

Keep in mind that **class definitions are based on all values in Canada**. If your map shows only a portion of Canada, it is likely that one or more classes will not occur in this smaller area. To rectify this, you may wish to create a new dataset containing only those features appearing on your map, and then classify that dataset.

To create a new dataset containing only those features appearing on your map, zoom in to your area of interest. Then, right-click on the original dataset in the Table of Contents, navigate to the 'Data' submenu, and click 'Export Data...':

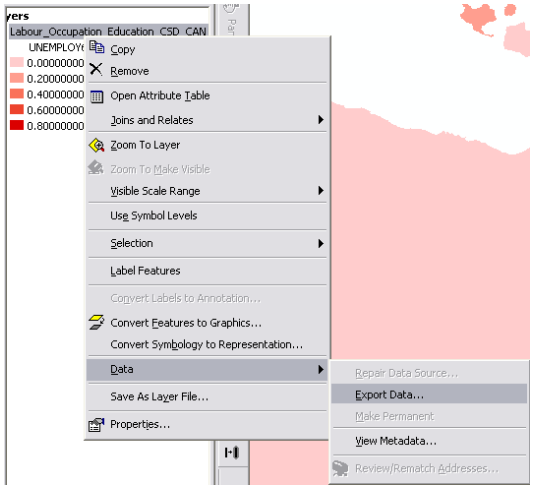

In the 'Export Data' dialogue box, choose 'All features in View Extent' from the Export dropdown box, enter or browse to a suitable output Shapefile name, and click 'OK.'

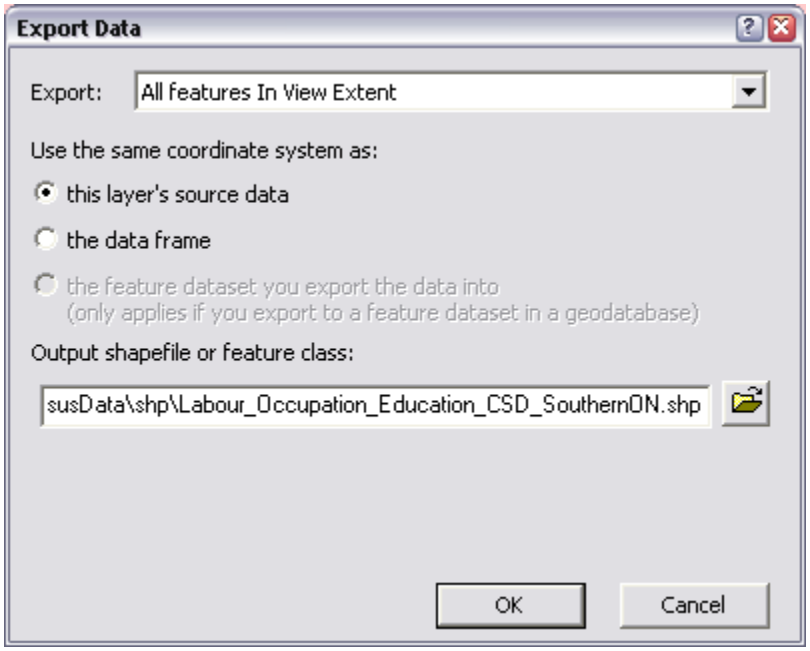

Click 'Yes' on the alert box which will open to add the new dataset to your map:

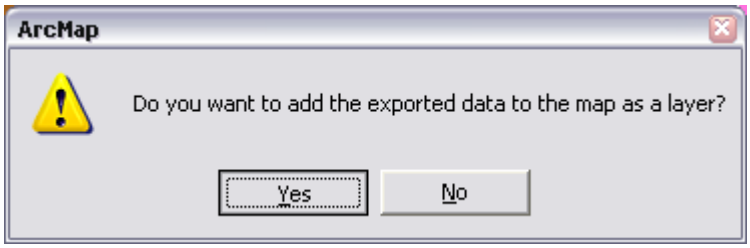

Now, symbolise the newly added dataset in the same way as before. The class definitions, and therefore the map colouring, will be significantly different, but likely much more readable.

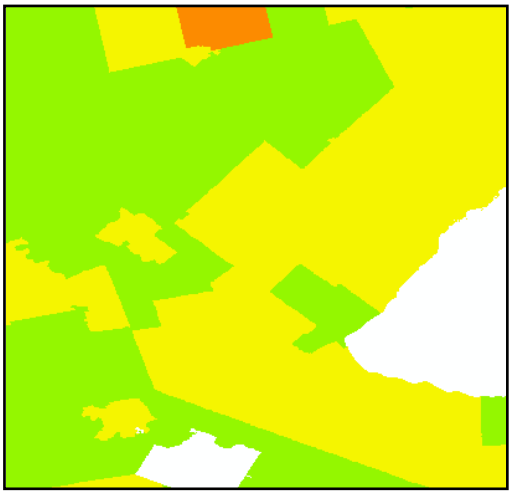

Original

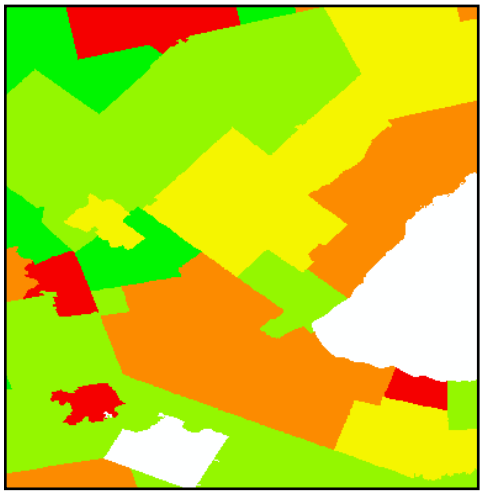

Exported and Reclassified

Handling 'No Data' Features

In some cases, some features will not be symbolised for various reasons – they may have been excluded from symbolisation, or there may have been insufficient data to normalise the variable (for instance, when the normalisation value is '0').

These 'no data' features will appear as holes on your map. To show that there are features there, simply add the Census dataset with the same geographic level as you are using on your map, and symbolise this new dataset as a single, unique colour – perhaps a light grey. Move the new layer such that it appears below the symbolised Census layer in the Table of Contents.

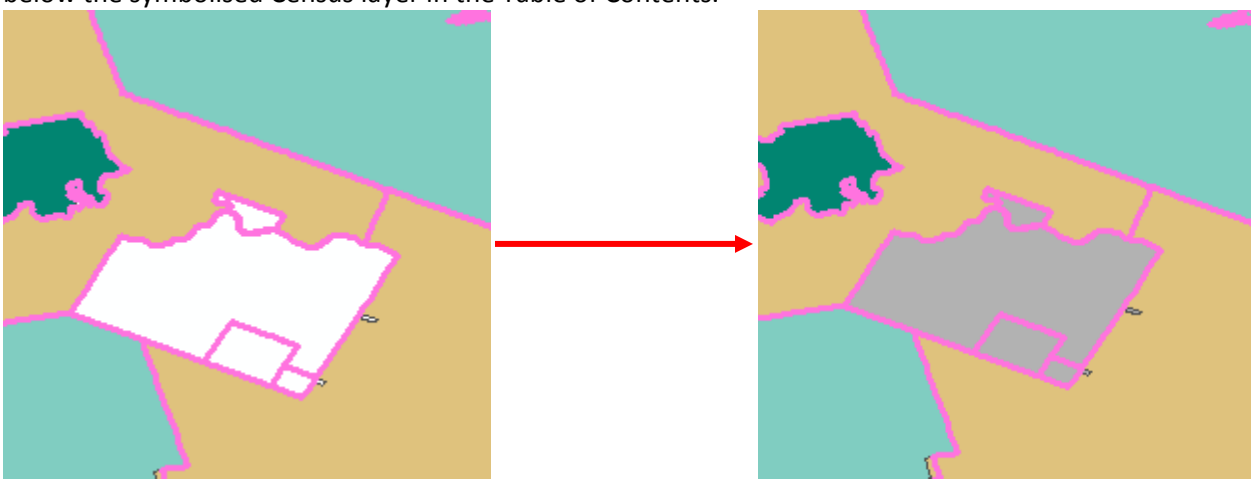

Labelling Census Features and Statistics

You may find it necessary to label Census geographies and their associated statistics to make your map more informative. Adding such labels can be complicated, but is often well worth it.

To add labels, open the Properties window by right-clicking on the Census dataset in the Table of Contents and clicking 'Properties,' or by double-clicking the Census dataset. Switch to the 'Labels' tab.

Check the 'Label features in this layer' box and click 'Apply.' Are these labels adequate?

Change the label font in the Properties window, if necessary, and add a mask to ensure the text is legible. To add a mask, click the 'Symbol...' button in the Properties window, then 'Properties...' in the Symbol Selector window. Switch to the 'Mask' tab in the Editor window, and select the 'Halo' radio button. Click 'OK' twice to return to the Properties window, then click 'Apply' to see the mask.

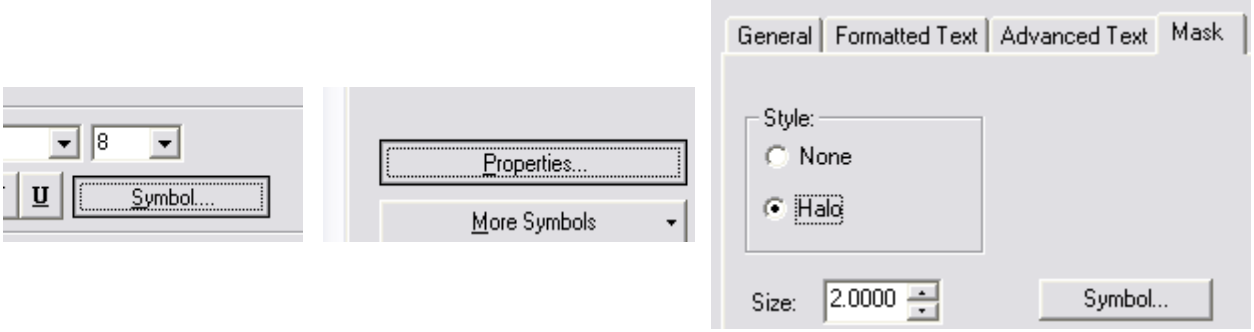

You may wish to include other information in your labels – for instance, a label which displays the Feature Name, the Feature Type, and then, on a new line, the value of the census variable.

To add a complex label, click the 'Expression' button in the Properties window.

In the Expression window, you can add fields to the label by double-clicking them in the field listings. You can also add plain text – which remains constant for every label – by enclosing the text in double-quotes. Include spaces in your text string. Combine fields and text components with ampersands (&)

For example, for a label which reads "Waterloo Census Division", "Hamilton Census Divsion", etc, the expression would read:

[CDNAME] & " Census Division"

Use the operator & vbnewline & to begin a new line in your label:

[CDNAME] & vbnewline & "Census Division"

will produce a label reading:

Waterloo
Census Division

For a complex label which displays the Feature Name, Feature Type, and the first four digits of the normalised statistic expressed as a percentage, the expression used is as follows:

[CSDNAME] & " C.S.D:" & vbnewline & left(([UNEMPLOY6] / [IN_THE_L4] * 100), 4) & "%"

and will produce a label reading:

Waterloo C.S.D:
13.3%

For additional help, sample label expressions, and suggestions, click the 'Help' button below the expression textbox:

Symbolising the Census Data: Multiple-Variable Mapping

Before beginning this section of the tutorial, complete or read through the section above on Single-Variable Mapping. This section assumes you are now familiar with setting symbology and using the various classification schemes.

Consider the examples below, and decide which technique will allow you to best symbolise your data. You may wish to experiment with each technique until you are satisfied with the result.

1. **Proportional symbols**

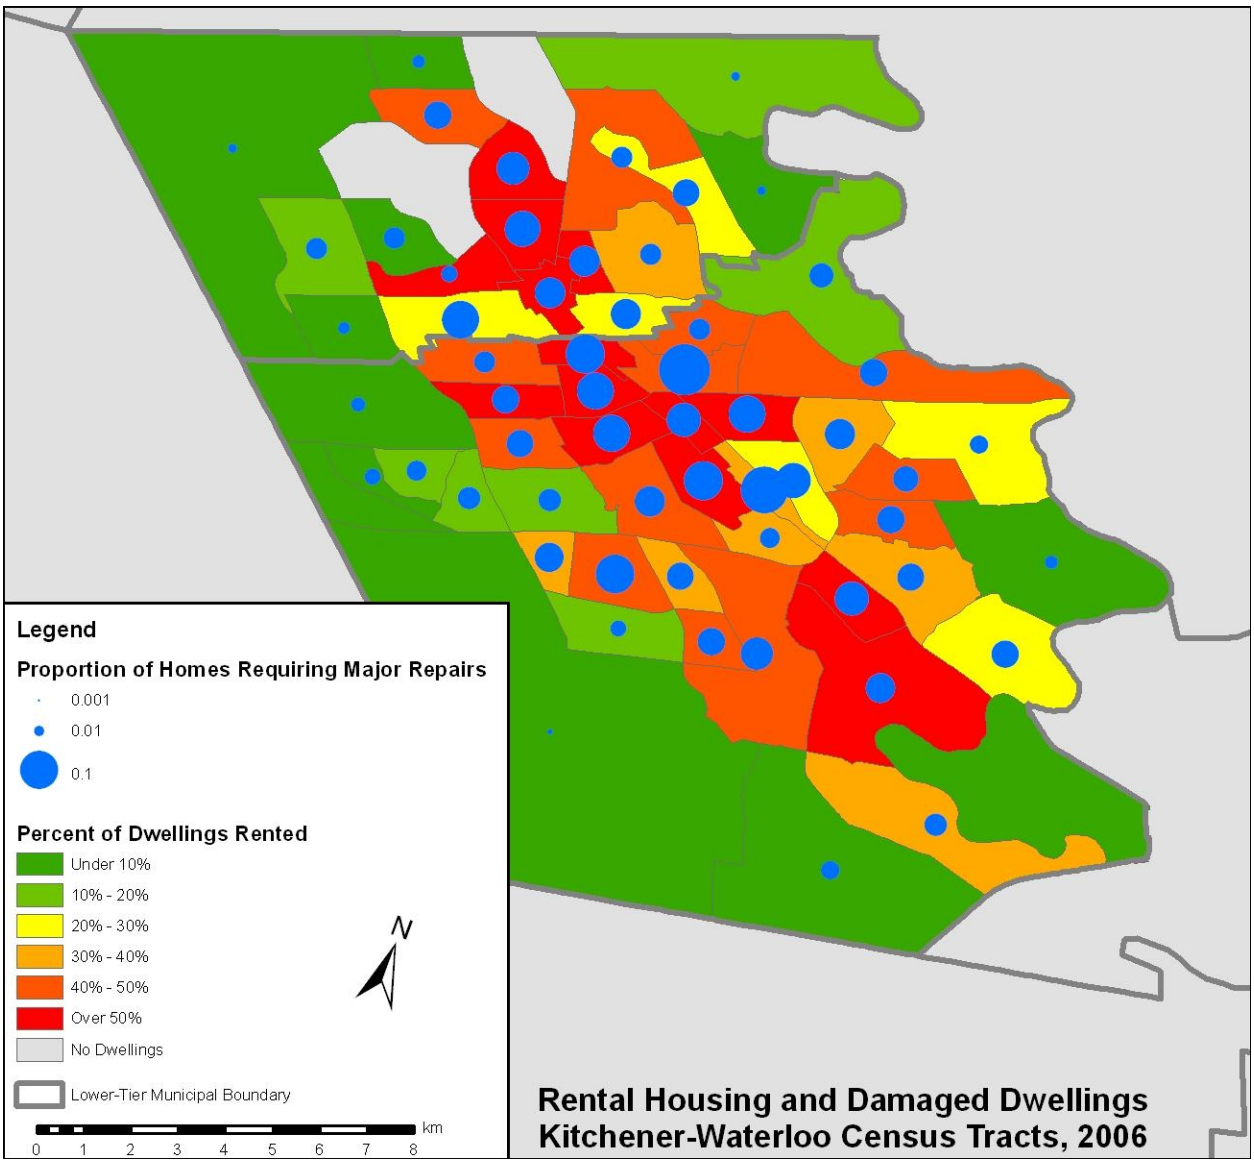

2. Dot density

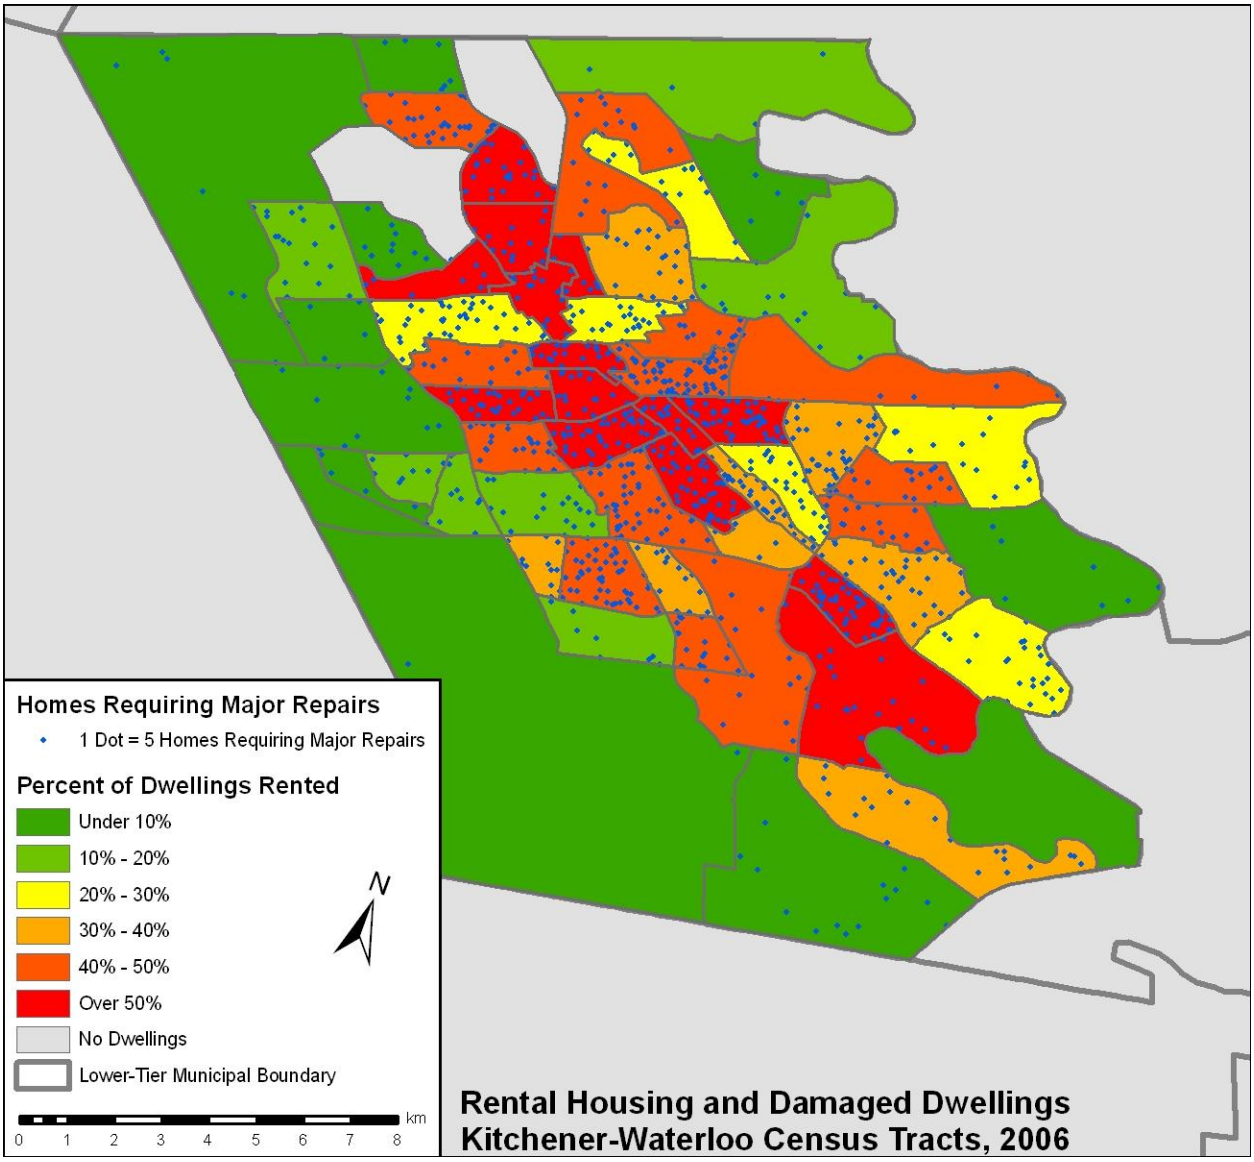

3. Charts: Pie Charts

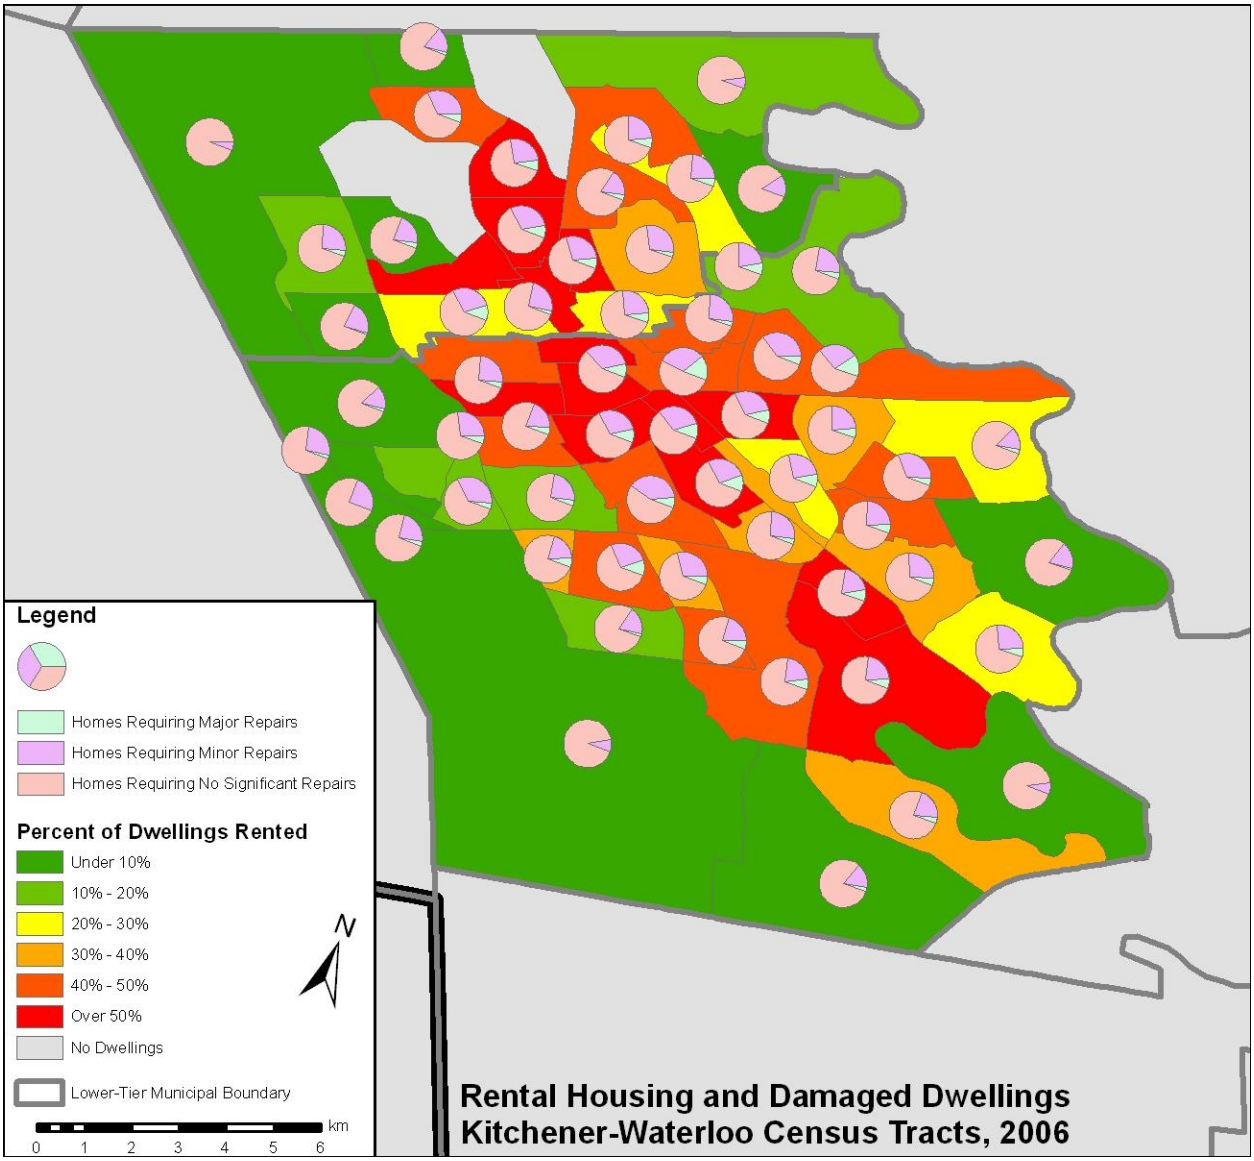

4. **Charts: Stacked Column Charts**

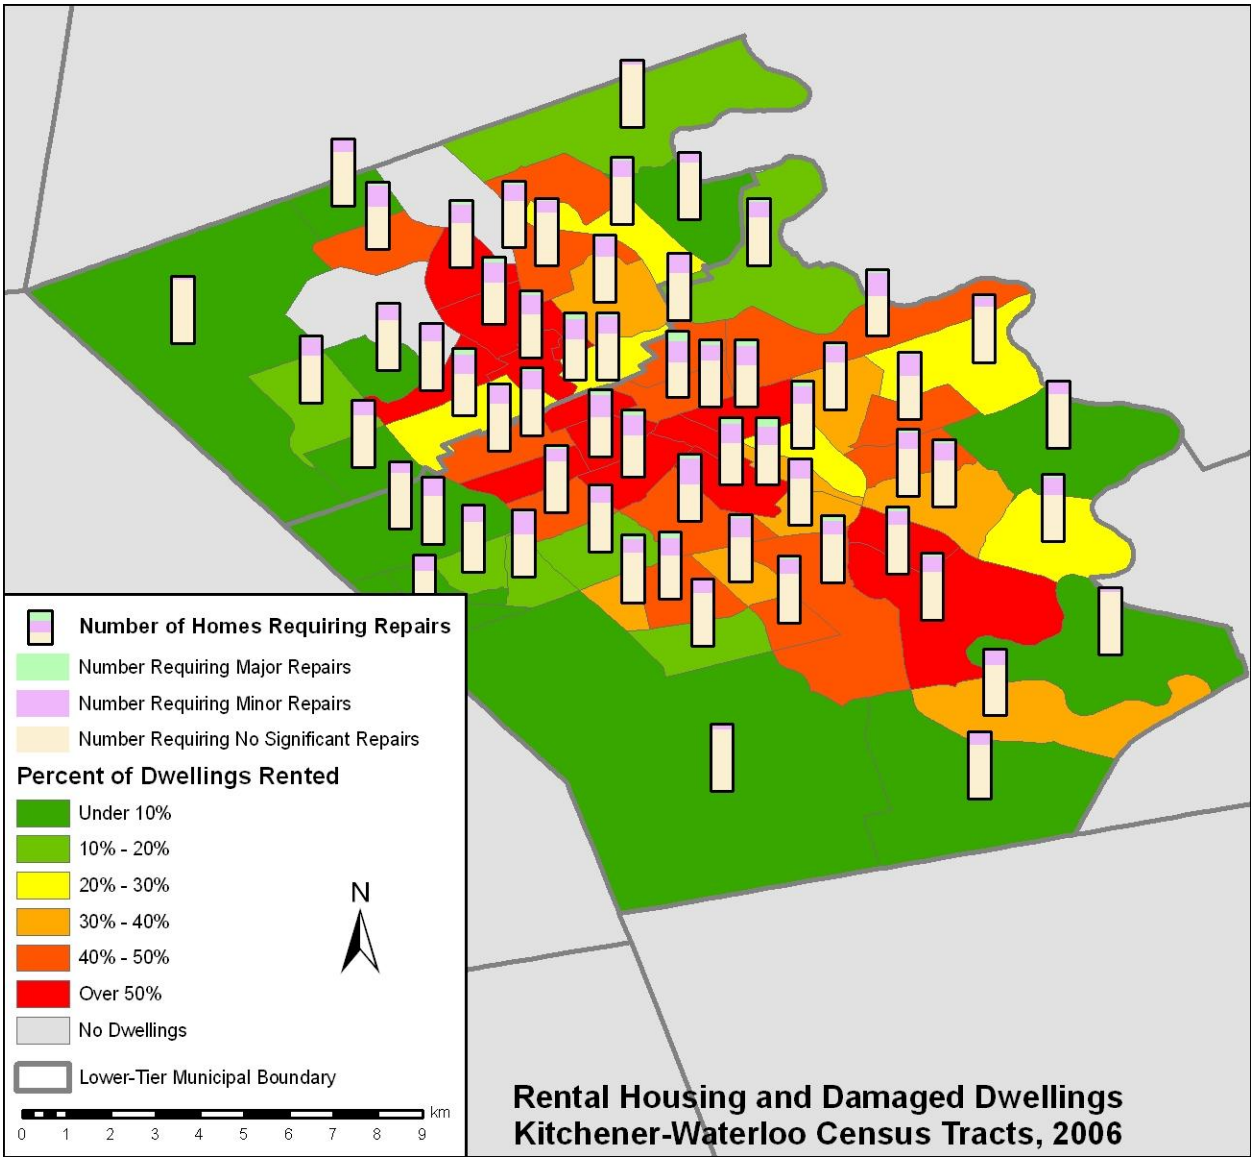

Symbolising Data by Proportional Symbols and Graduated Colours

Symbolising data by proportional symbols will produce a map displaying a single symbol for each geographic area. The larger the symbol, the higher the data value for that area.

Proportional symbols work well when overlaid atop a map symbolised by graduated colours, and therefore allow you to map two separate Census variables.

To begin, follow the directions above to create a map symbolised by graduated colours. Once this map is complete, add to your map the Census dataset which contains the variables you want to symbolise by proportional symbols. This can be the same dataset already on your map.

For example, I have created a map displaying the proportion of homes in each Census Tract in the Cities of Kitchener and Waterloo which are rented, symbolised by graduated colour. I also want to show the proportion of homes in need of major repairs in each Census Tract in the Cities of Kitchener and Waterloo, the variable which I will symbolise by proportional symbols. I need to re-add the Census dataset to the map, either by using the 'Add Data' button ( ) or by copying the dataset:

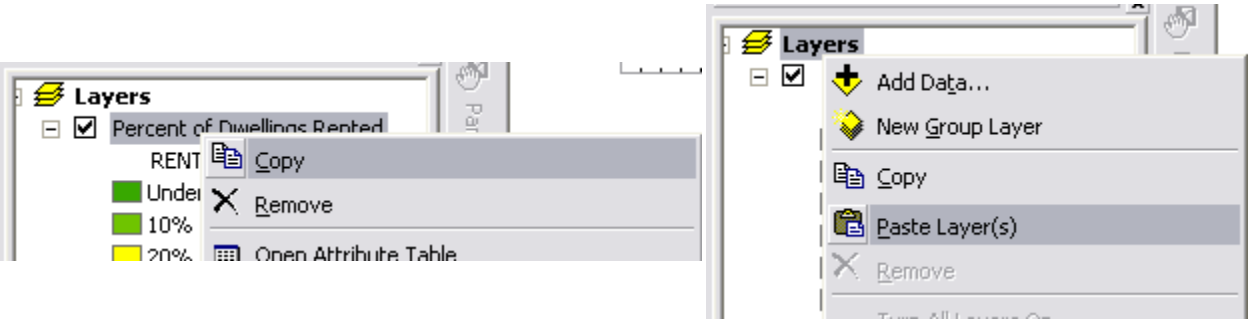

At this point, you should have *two* Census datasets showing in your Table of Contents in ArcMap. These can be the same *or* different datasets. Arrange these datasets such that the data symbolised by graduated colours appears below the new layer:

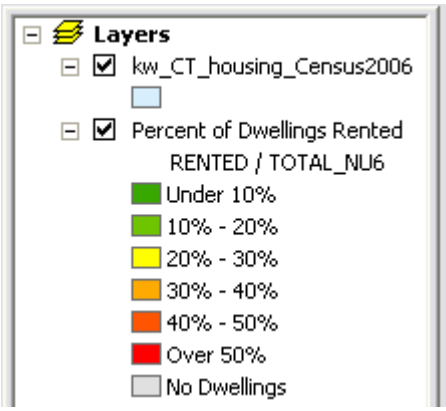

Open the ‘Layer Properties’ window of the new layer, and switch to the ‘Symbology’ tab. From the list at left, choose ‘Quantities’, then select ‘Proportional Symbols.’

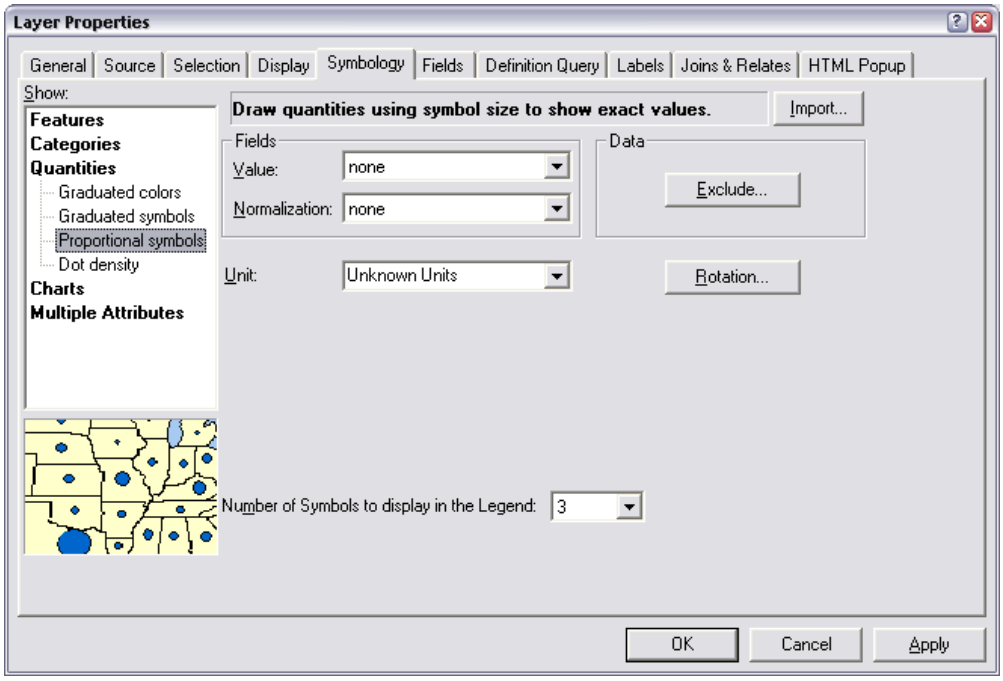

From the ‘Value’ dropdown box, select the field containing the values you wish to display as proportional symbols. If necessary, choose the field by which those values will be normalised from the ‘Normalization’ dropdown box.

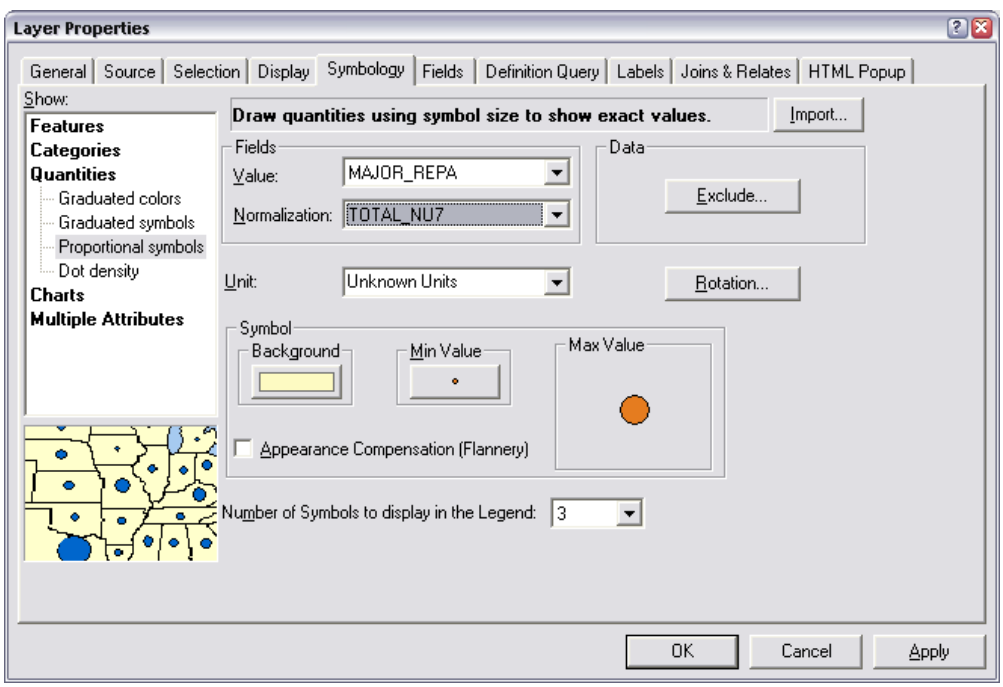

Choose the symbol and colour you wish to use to show data values by clicking the button below 'Min Value':

In the window which opens, choose an appropriate symbol from the list, set an appropriate colour, and set the size. The size setting in this window will be the size of the smallest symbol on the map: the larger this size is, the larger will be the size of the largest symbol on the map. You will likely need to experiment with different sizes until you find one which works.

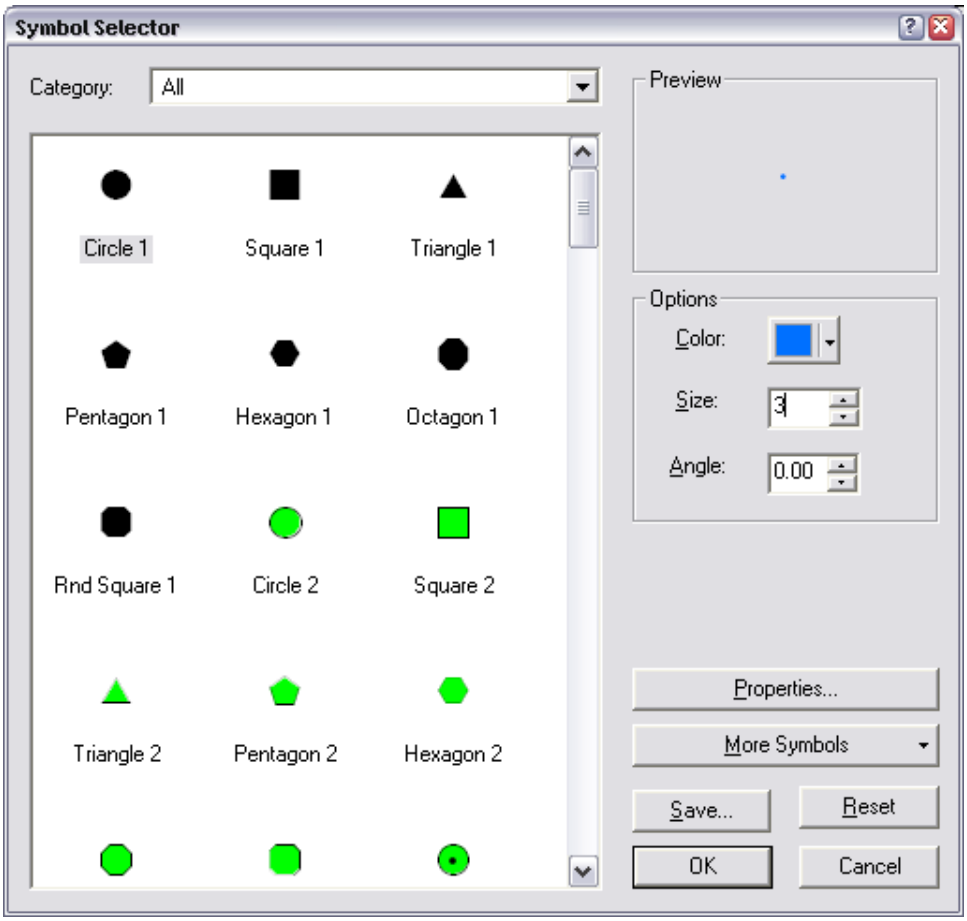

Click 'OK' to return to the Layer Properties dialogue box.

Click 'Apply' to see the proportional symbols on your map. If necessary, change the size of the symbols.

You will notice that you can't see the Graduated Colours on your map. This is because, by default, ArcMap applies an opaque background colour to the new Proportional Symbols layer. To turn this background off, click the button underneath 'Background' in the Layer Properties window, illustrated at right:

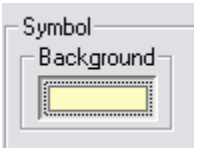

In the Symbol Selector window, set the Fill Color to ‘No Color’, the Outline Width to ‘0’, and the Outline Colour to ‘No Colour’:

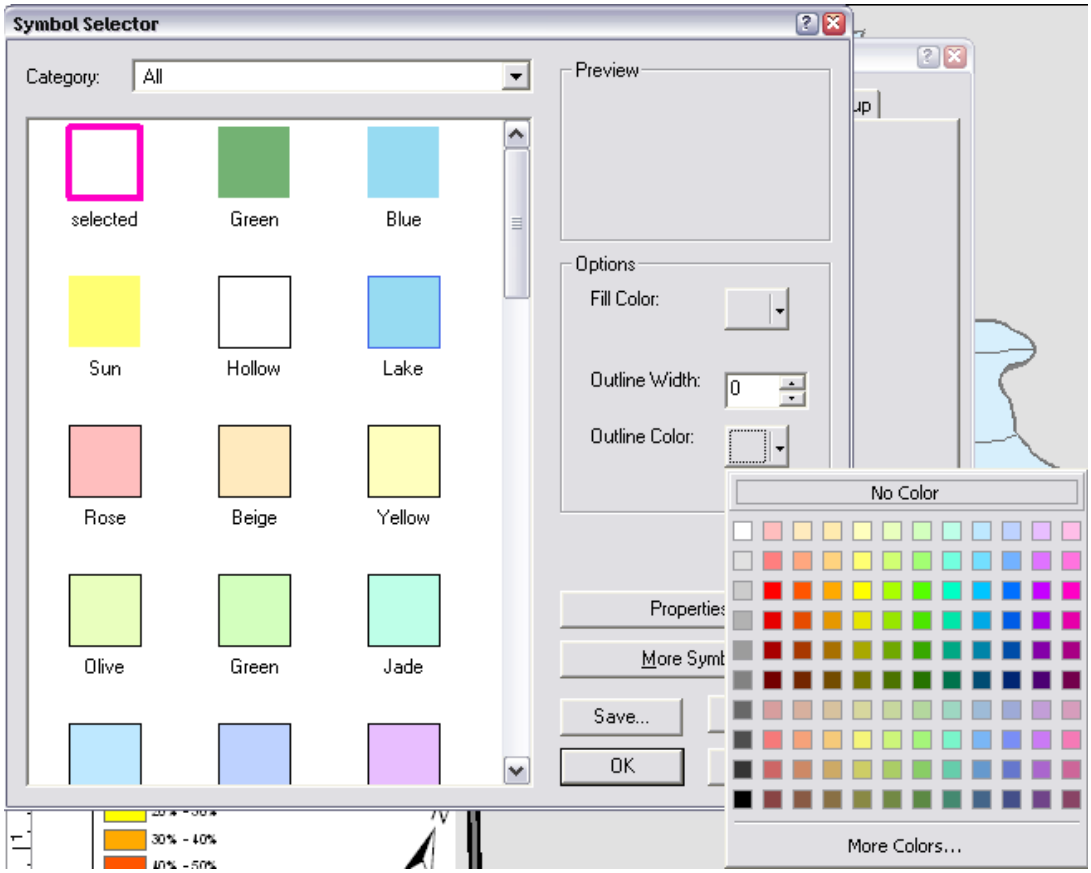

Click ‘OK’ to return to the Layer Properties window, and click ‘Apply’ to see the changes on your map.

Continue to experiment with symbol sizes and colours until you are happy with your map.

You may wish to use ‘Flannery Appearance Compensation’ which adjusts the size of larger symbols upwards to account for the fact that many map readers tend to underestimate the size of circular symbols. Try turning Appearance Compensation on and off using the checkbox shown below:

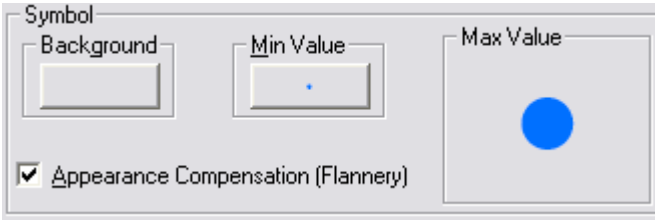

Once you are happy with your map, continue to the final section of this tutorial, “Completing the Map.”

Symbolising Data using Dot Density and Graduated Colours

Symbolising data using dot densities will produce a map displaying a number of dots, each of the same size, within each geographic area. The more dots in a given area, the higher the data value for that area.

Dot densities work well when overlaid atop a map symbolised by graduated colours, and therefore allow you to map two separate Census variables.

To begin, follow the directions above to create a map symbolised by graduated colours. Once this map is complete, add to your map the Census dataset which contains the variables you want to symbolise by dot densities. This can be the same dataset already on your map.

For example, I have created a map displaying the proportion of homes in each Census Tract in the Cities of Kitchener and Waterloo which are rented, symbolised by graduated colour. I also want to show the proportion of homes in need of major repairs in each Census Tract in the Cities of Kitchener and Waterloo, the variable which I will symbolise by dot densities. I need to re-add the Census dataset to the map, either by using the 'Add Data' button ( ) or by copying the dataset:

At this point, you should have *two* Census datasets showing in your Table of Contents in ArcMap. These can be the same *or* different datasets. Arrange these datasets such that the data symbolised by graduated colours appears below the new layer:

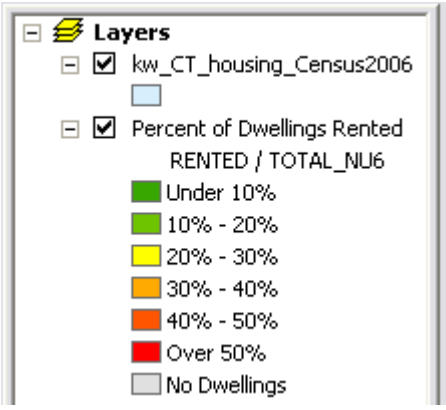

Open the ‘Layer Properties’ window of the new layer, and switch to the ‘Symbology’ tab. From the list at left, choose ‘Quantities’, then select ‘Dot density.’

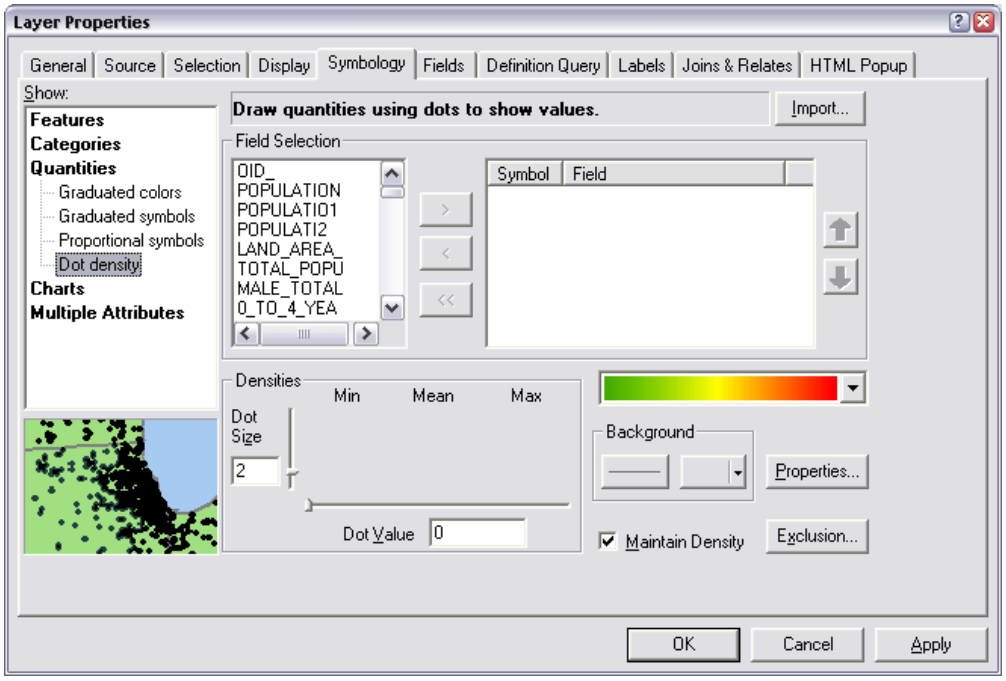

From the list of fields at the left of the window, choose the field containing the values you wish to display as dot densities. Click the [>] button to add that field to the display:

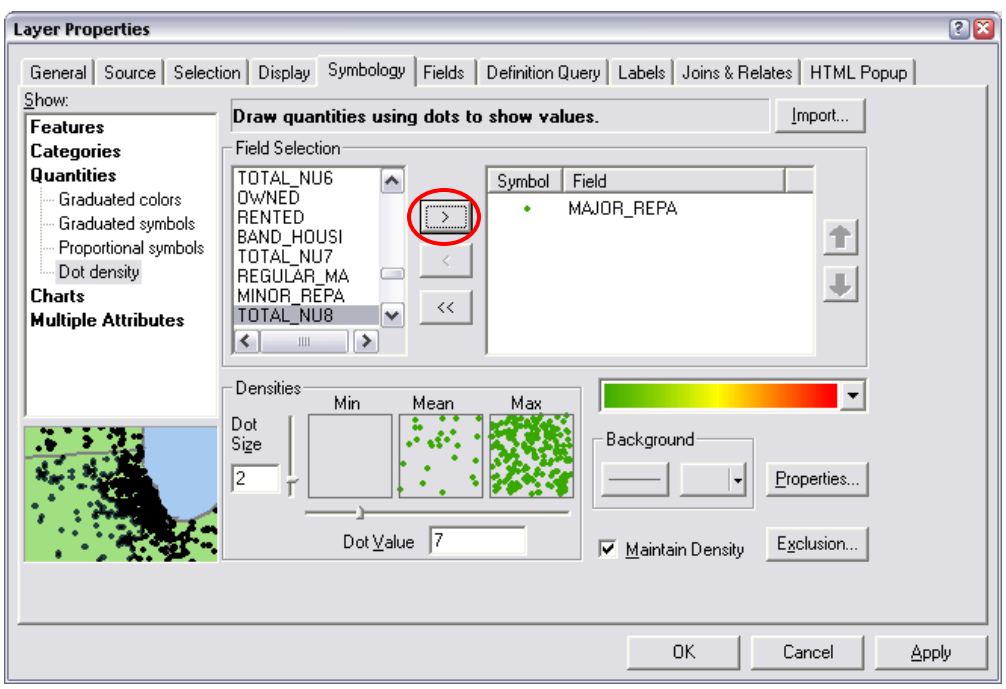

Choose the symbol and colour you wish to use by double-clicking the icon next to your chosen field:

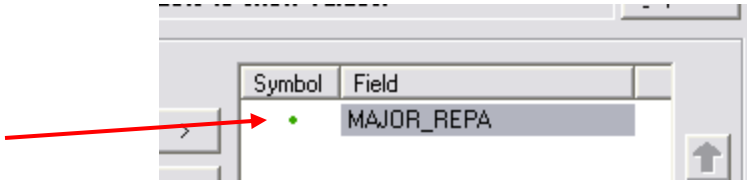

In the window which opens, choose an appropriate symbol from the list, set an appropriate colour, and set the symbol size:

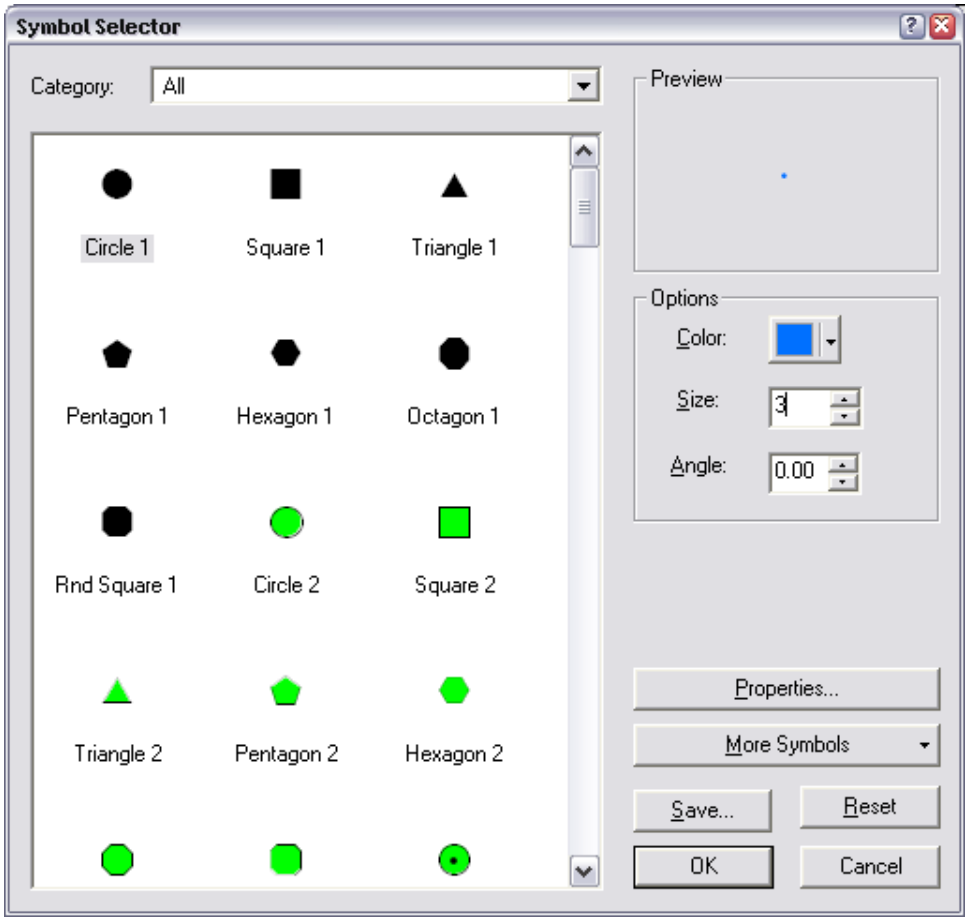

Click 'OK' to return to the Layer Properties dialogue box.

You will now need to adjust the Dot Value and Dot Size using the sliders illustrated at right. Experiment with different sizes and values until you find a combination which works well with your data. This may take some time. Adjust one value at a time, then click 'Apply' to see the changes on your map.

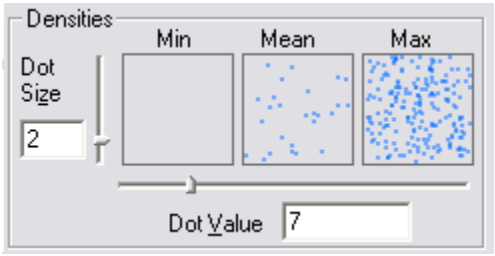

Once you are happy with your map, continue to the final section of this tutorial, "Completing the Map."

Symbolising Data using Pie Charts and Graduated Colours

Symbolising data using pie charts will produce a map displaying a pie chart for each geographic area showing the relative composition of two or more variables for each area. Pie charts can be difficult to see at small scales, and thus work best when showing a small geographical area.

To begin, follow the directions above to create a map symbolised by graduated colours. Once this map is complete, add to your map the Census dataset which contains the variables you want to symbolise by pie charts. This can be the same dataset already on your map.

For example, I have created a map displaying the proportion of homes in each Census Tract in the Cities of Kitchener and Waterloo which are rented, symbolised by graduated colour. I also want to show the proportion of homes in need of major repairs in each Census Tract in the Cities of Kitchener and Waterloo, the variable which I will symbolise by pie charts. I need to re-add the Census dataset to the map, either by using the 'Add Data' button ( ) or by copying the dataset:

At this point, you should have *two* Census datasets showing in your Table of Contents in ArcMap. These can be the same *or* different datasets. Arrange these datasets such that the data symbolised by graduated colours appears below the new layer:

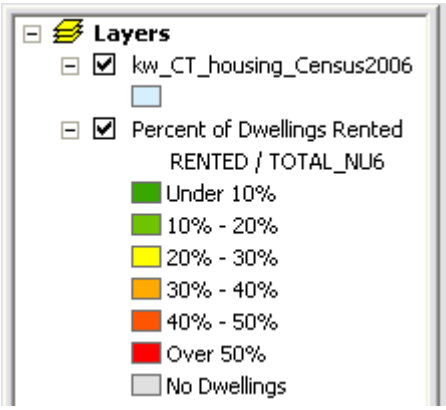

Open the ‘Layer Properties’ window of the new layer, and switch to the ‘Symbology’ tab. From the list at left, choose ‘Charts’, then select ‘Pie.’

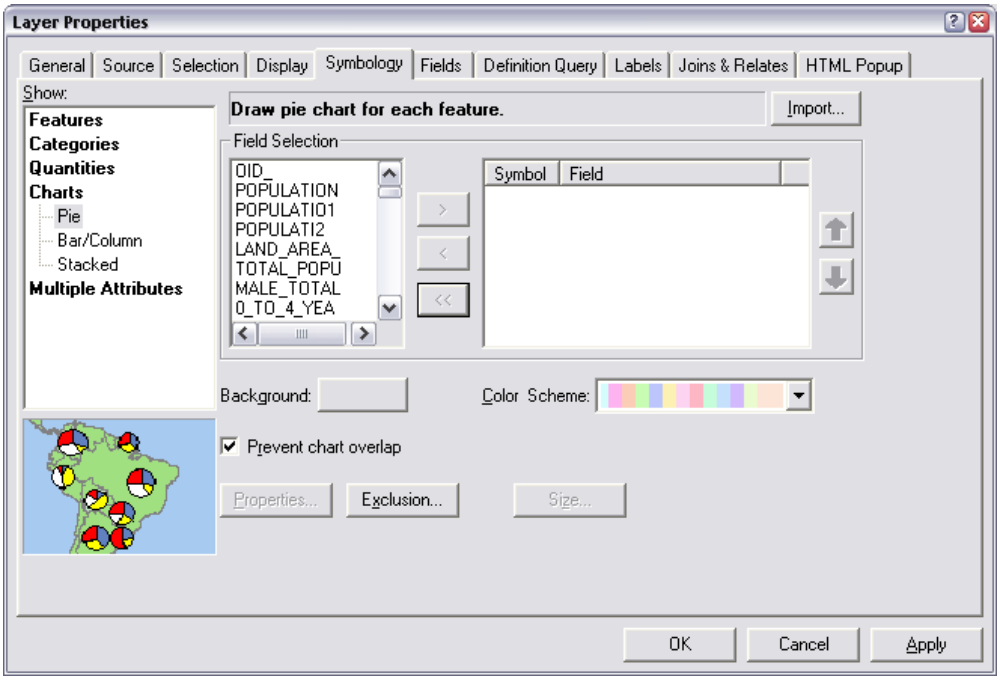

From the list of fields at the left of the window, choose the fields containing the values you wish to display in the chart. Try not to use more than three or four fields to ensure your charts are legible. Click the [>] button to add each field to the display:

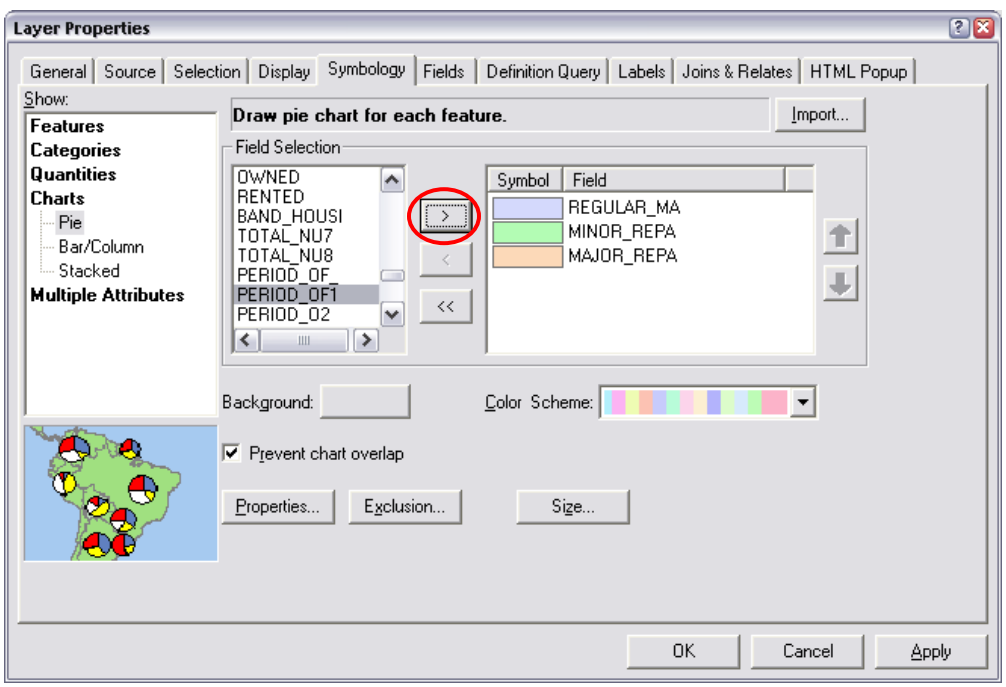

Choose the symbol and colour you wish to use for each field by double-clicking the icon next to the field:

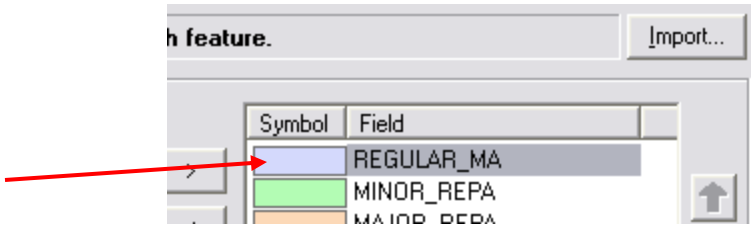

In the window which opens, choose an appropriate fill and outline colour. Set the outline width to a very small number – no larger than 0.5 – to ensure the outline does not detract from the chart's legibility.

Click 'OK' to return to the Layer Properties dialogue box. Set the colours for the other fields as well.

Click the 'Properties' button to open the Chart Symbol Editor.

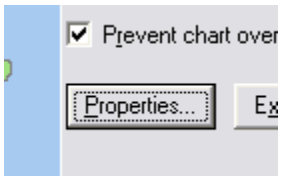

Here, experiment with the different settings until you are happy with how your chart looks.

Once you are done, click 'OK' to return to the Layer Properties window. Click 'Apply' to see the new charts on your map.

You may need to adjust the size of your charts. To do so, click the 'Size button':

Experiment with the 'size' value (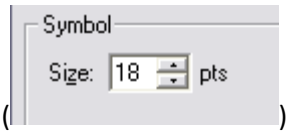

) until the charts are an appropriate size.

Once you are happy with your map, continue to the final section of this tutorial, "Completing the Map."

Symbolising Data using Stacked Bar/Column Charts and Graduated Colours

Symbolising data using stacked charts will produce a map displaying a chart for each geographic area showing the relative composition of two or more variables for each area. Stacked charts can be difficult to see at small scales, and thus work best when showing a small geographical area.

To begin, follow the directions above to create a map symbolised by graduated colours. Once this map is complete, add to your map the Census dataset which contains the variables you want to symbolise by a stacked chart. This can be the same dataset already on your map.

For example, I have created a map displaying the proportion of homes in each Census Tract in the Cities of Kitchener and Waterloo which are rented, symbolised by graduated colour. I also want to show the proportion of homes in need of major repairs in each Census Tract in the Cities of Kitchener and Waterloo, the variable which I will symbolise by a stacked chart. I need to re-add the Census dataset to the map, either by using the 'Add Data' button ( ) or by copying the dataset:

At this point, you should have *two* Census datasets showing in your Table of Contents in ArcMap. These can be the same *or* different datasets. Arrange these datasets such that the data symbolised by graduated colours appears below the new layer:

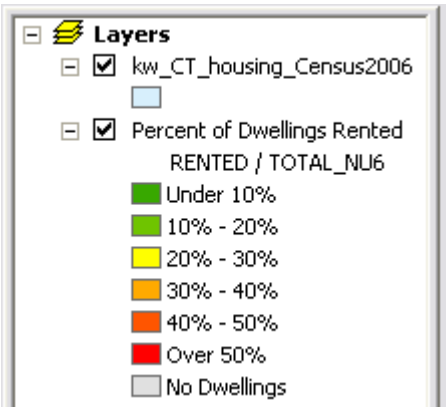

Open the ‘Layer Properties’ window of the new layer, and switch to the ‘Symbology’ tab. From the list at left, choose ‘Charts’, then select ‘Stacked.’

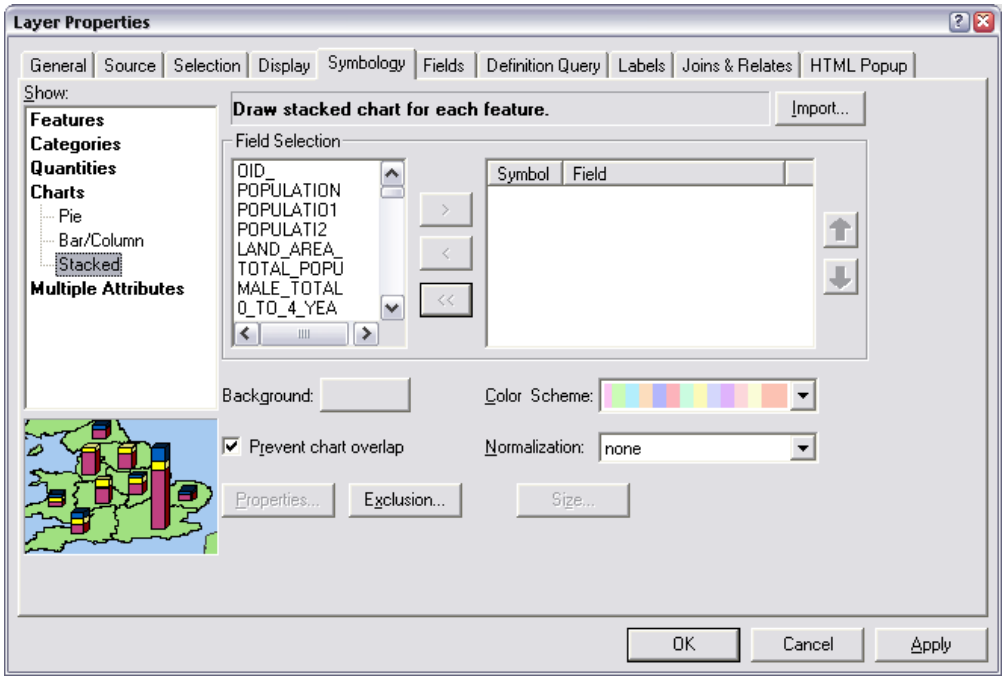

From the list of fields at the left of the window, choose the fields containing the values you wish to display in the chart. Try not to use more than three or four fields to ensure your charts are legible. Click the [>] button to add each field to the display:

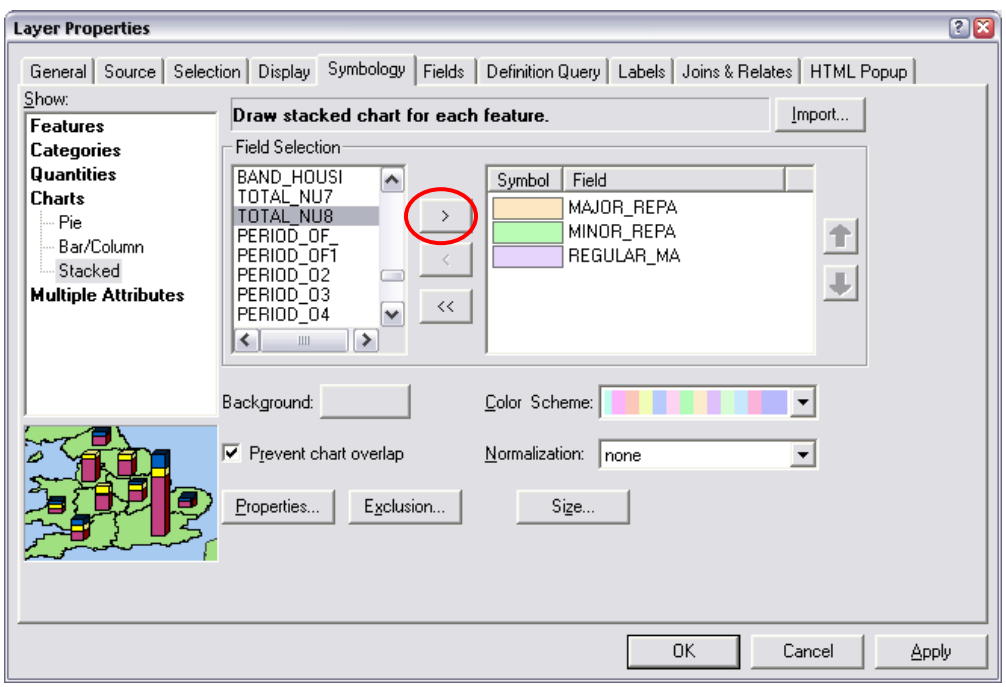

If necessary, select the field by which your data will be normalised from the 'Normalization' drop-down box: 

Choose the symbol and colour you wish to use for each field by double-clicking the icon next to the field:

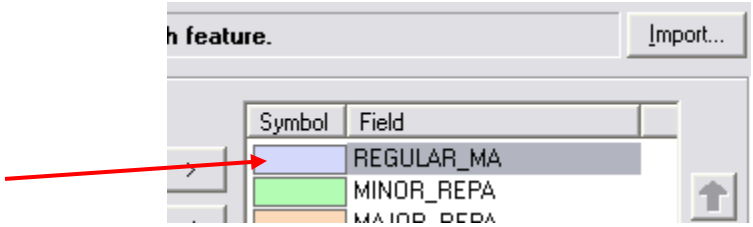

In the window which opens, choose an appropriate fill and outline colour. Set the outline width to a very small number – no larger than 0.5 – to ensure the outline does not detract from the chart's legibility.

Click 'OK' to return to the Layer Properties dialogue box. Set the colours for the other fields as well.

Click the 'Properties' button to open the Chart Symbol Editor.

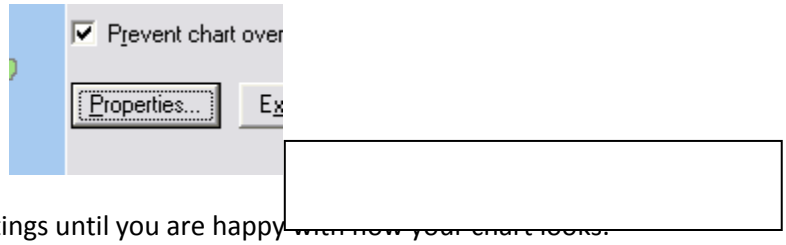

Here, experiment with the different settings until you are happy with how your chart looks.

Once you are done, click 'OK' to return to the Layer Properties window. Click 'Apply' to see the new charts on your map.

You may need to adjust the size of your charts. To do so, click the 'Size button':

Experiment with the ‘size’ value (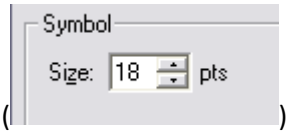

) until the charts are an appropriate size.

Note that adjusting the size in the ‘Size’ window will also alter the Bar Width in the Chart Symbol Editor window. After adjusting the size, you may need to click the ‘Properties’ button and increase or decrease the Bar Width value.

Once you are happy with your map, continue to the final section of this tutorial, “Completing the Map.”

Completing your Map

At this point, your Census data should be symbolised as desired – but you may wish to add additional layers to show context: perhaps municipal boundaries, water bodies, provincial and national boundaries, and/or roads. Map Library staff can help you acquire the necessary data.

Add the necessary data to your map and symbolise and label it appropriately.

Once you have done so, switch to 'Layout View' by clicking the button illustrated below:

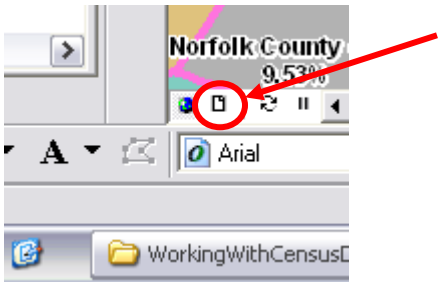

You should now see a rectangle, representing a sheet of paper, with your map centered within.

Change the orientation of the page if necessary from portrait to landscape or vice versa by navigating to the 'File' menu and clicking 'Page and Print Setup...' In the dialogue box which opens, change the Paper Size and Orientation as necessary:

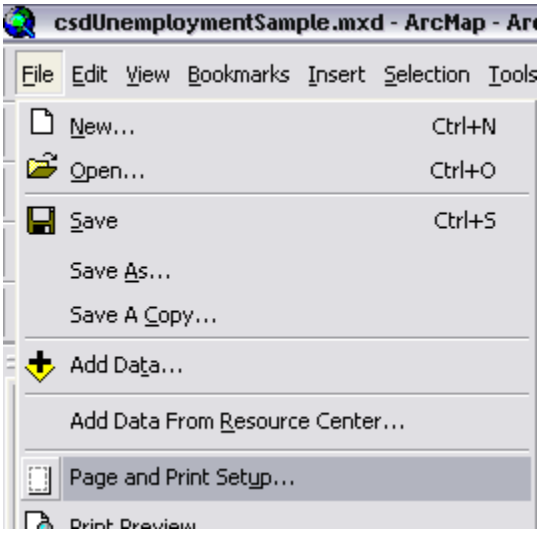

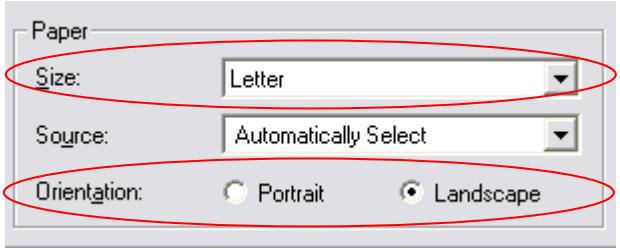

Once you are happy with the page setup, resize the map to fit the page. Click once on the map, then use the control handles on the corners and the edges to resize the map frame.

Zoom in or out or pan around your map until it is centred and zoomed appropriately in its frame. Use the appropriate tools to do so – the tools illustrated below on the **left** allow you to zoom in and out in your map; the tools on the **right** allow you to zoom in and out on the page itself:

Zoom in/out **in the map**

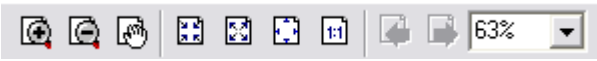

Zoom in/out **on the page**

Once you are happy with the size, position, and zoom level of your map, you will need to add a legend, north arrow, and scale bar to complete your map.

Add a Legend. Navigate to the 'Insert' menu, and click 'Legend.' In the dialogue box which opens, choose which layers to display in the legend, and reorder the layers as appropriate:

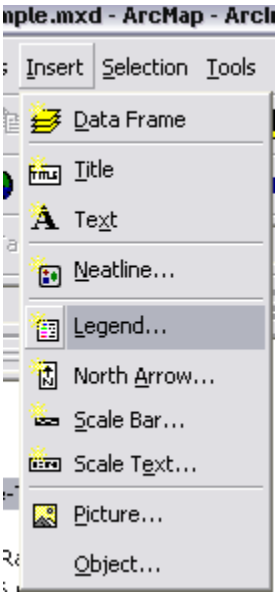

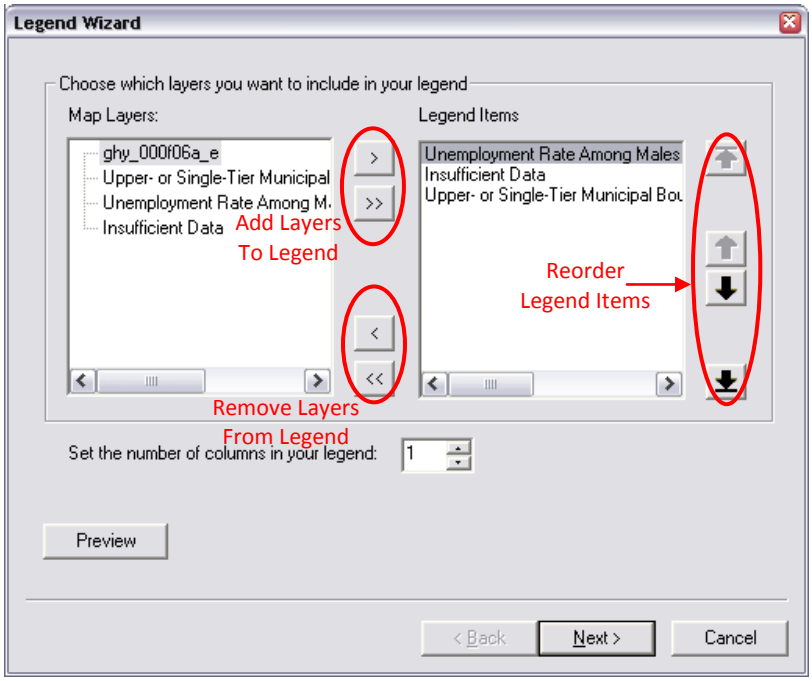

In most cases, it is not necessary to include the 'obvious' layers, such as water or roads, on the legend when creating thematic maps. Use your judgment – will your audience be able to understand the meaning of such layers if they are excluded from the legend?

Once you are satisfied with the legend items, click 'Next >' to format the title of the legend.

By default, the legend title is "Legend". If desired, change the legend title and the font and/or justification of the title.

In many cases, a legend title is unnecessary and redundant. Should you decide not to show a legend title, simply leave the settings in this window as-is – you can disable the title at a later stage.

Click 'Next' to move on.

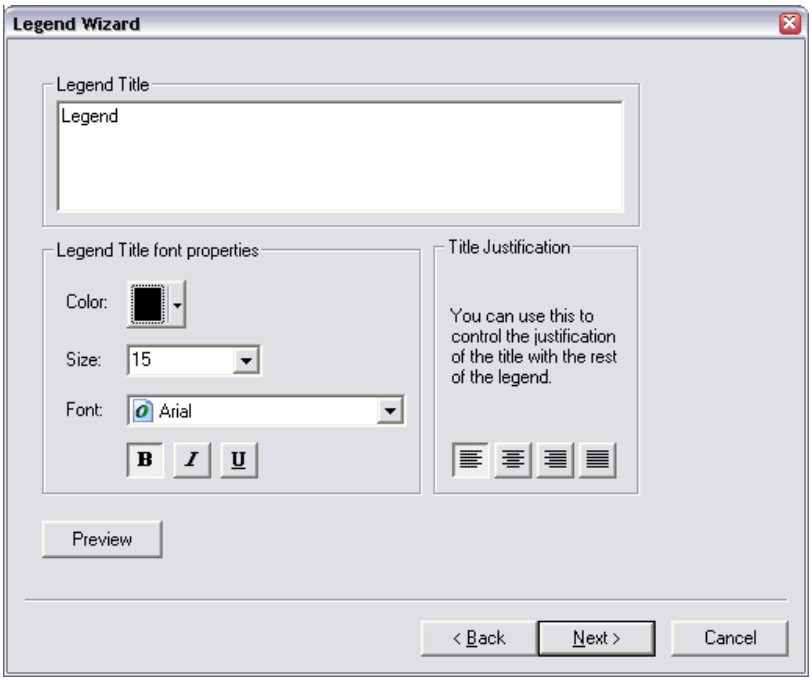

If desired, at this stage, you can add a border, background, and/or drop-shadow to your legend. Simply use the drop-down boxes to apply these settings:

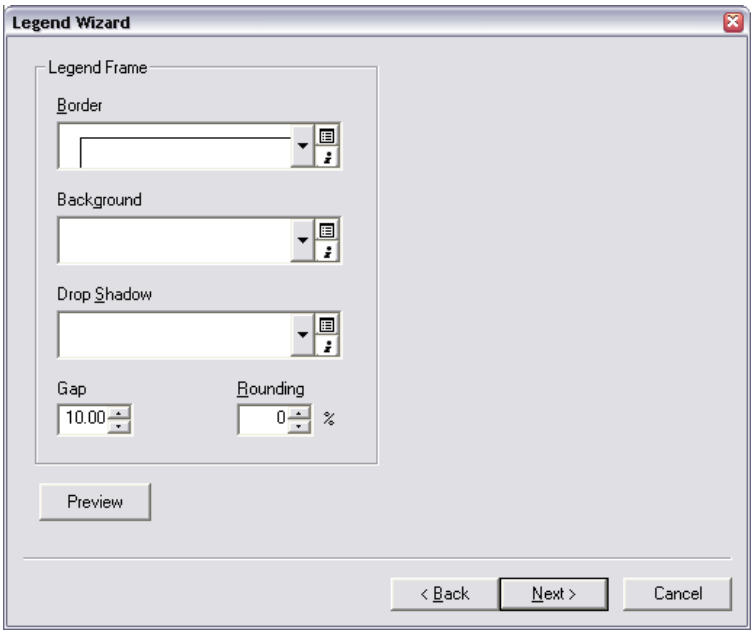

You may wish to increase the 'Gap' – leaving it set to 0 will result in your legend border appearing right next to the legend text, with no white space.

Click 'Next' to continue.

If you like, you can override the default settings for symbol patches in the legend. Normally, line features are displayed as straight horizontal lines, while polygon features are shown as rectangles. In most cases, the default settings are adequate and appropriate.

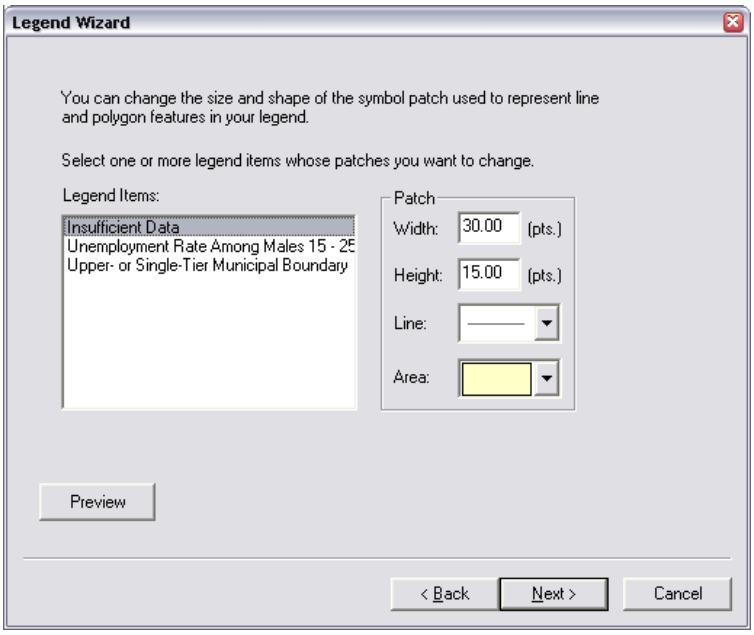

Click ‘Next’ to set the spacing between legend elements:

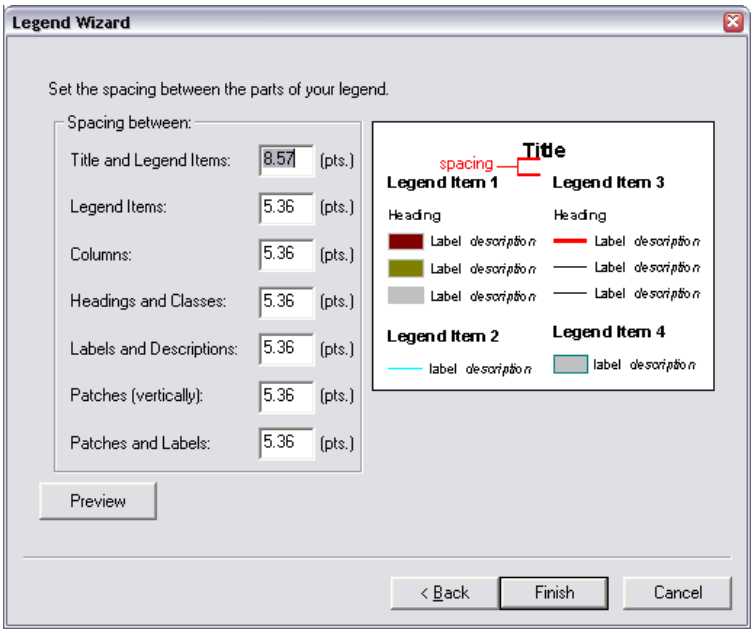

Usually, the default settings are fine. If necessary, adjust these values and then click ‘Finish.’

Preview your legend. There will probably be some problems with it:

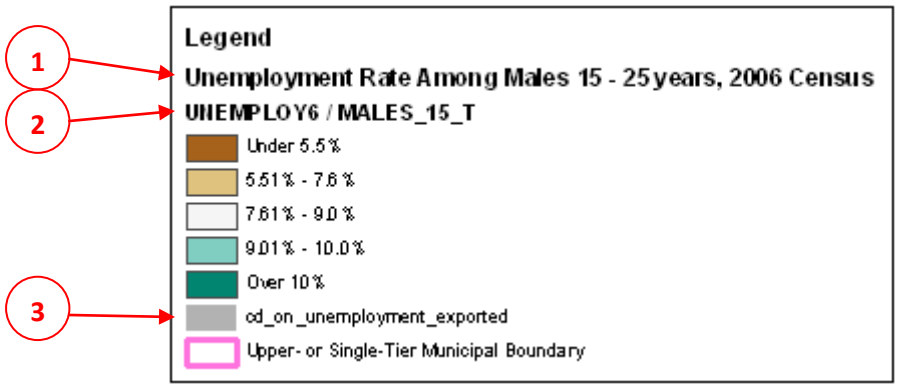

1: This layer heading will be redundant when I add a title for the map itself;
2: This heading is meaningless and will also be redundant when I add the map title;
3: The layer name displayed here does not explain the meaning of this layer.

Each of these problems is easily rectified.

First, open the Legend Properties window by double-clicking the legend or by right-clicking it and selecting “Properties...”

To remove the legend title: Uncheck the 'Show' box under the 'Title' textbox: 

To display a Map Title in place of the legend title: Ensure the 'Show' box is checked, and enter a meaningful, descriptive title in place of the legend title. Note that you can enter line breaks:

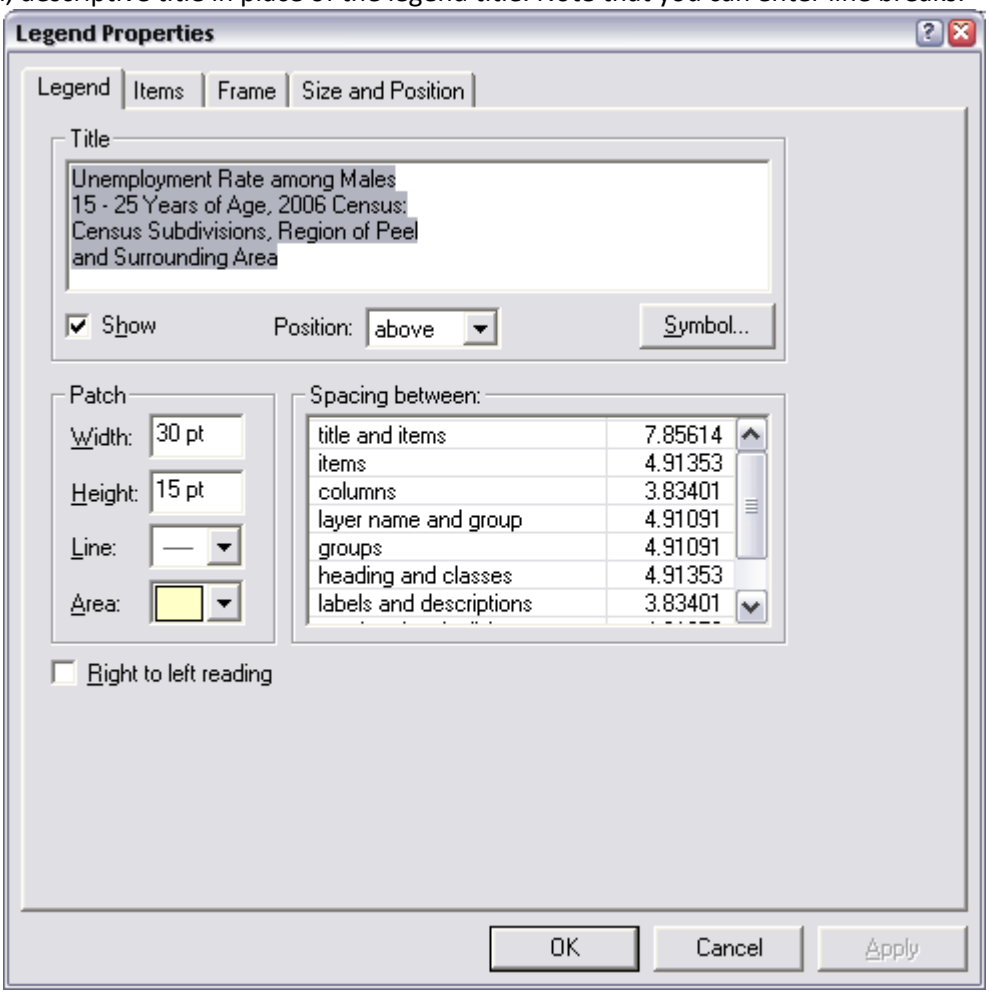

To reorder legend items: In the 'Items' tab, use the arrows to reorder items: 

To remove unnecessary layer headings (Problems '1' and '2'): Select the problematic layer in the 'Items' tab and click 'Style', then click 'Properties':

In the 'General' tab, uncheck 'Show Layer Name' and 'Show Heading', then click 'OK' twice to return to the Legend Properties window:

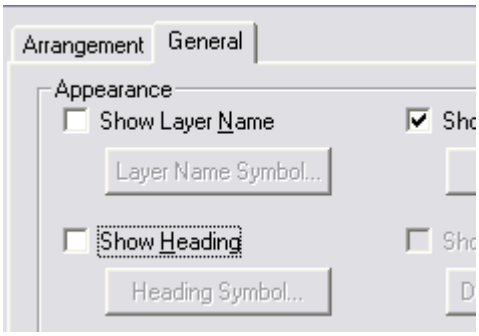

To change font settings for all labels: In the 'Items' tab of the Legend Properties window, use the dropdown box to select 'Apply to the label', and then click 'Symbol.' Change the font properties as desired, then click 'OK.'

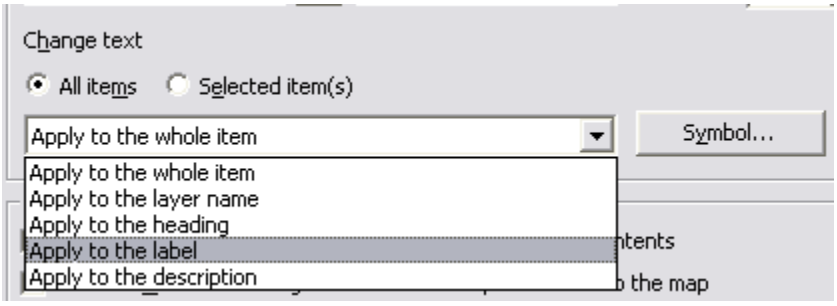

To rename legend items (Problem '3'): Exit the Legend Properties window. In the Table of Contents, rename the problematic layer(s) one at a time by selecting it, clicking once, editing the text, and pressing 'Enter'. The text in the legend will update automatically.

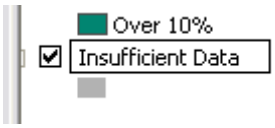

At this point, your legend should be complete:

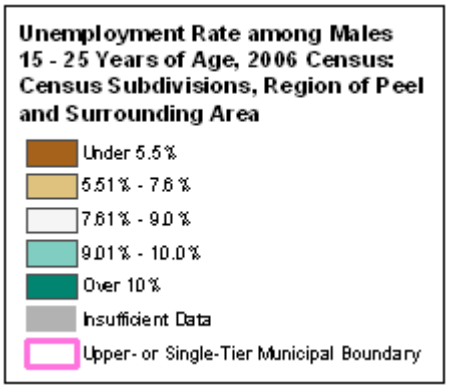

Add a north arrow to the page by navigating to the 'Insert' menu and clicking 'North Arrow...' Select an arrow from the North Arrow Selector window, and click 'OK'. Move and/or resize the new north arrow as appropriate.

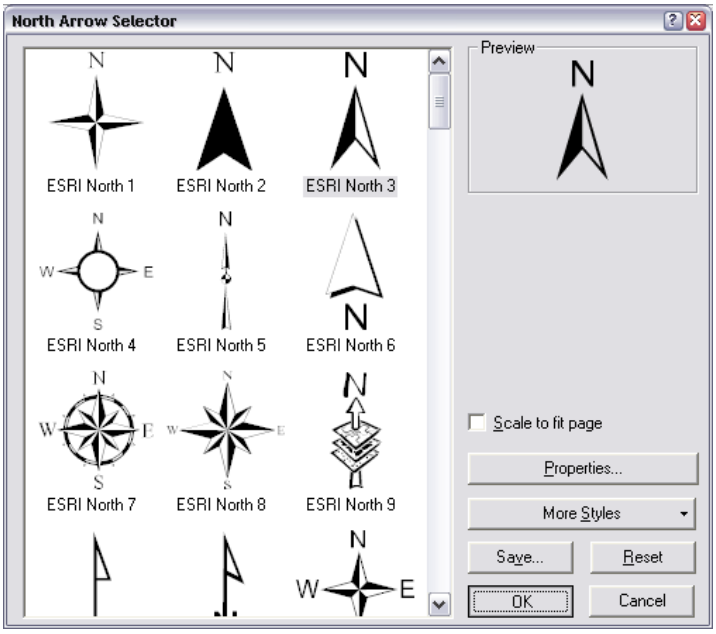

Add a scale bar to the page by navigating to the 'Insert' menu and clicking 'Scale Bar...'
Choose a scale bar style, and click 'OK.'

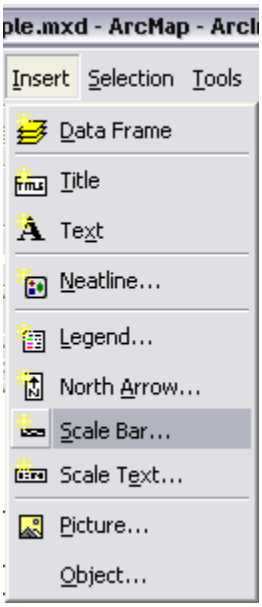

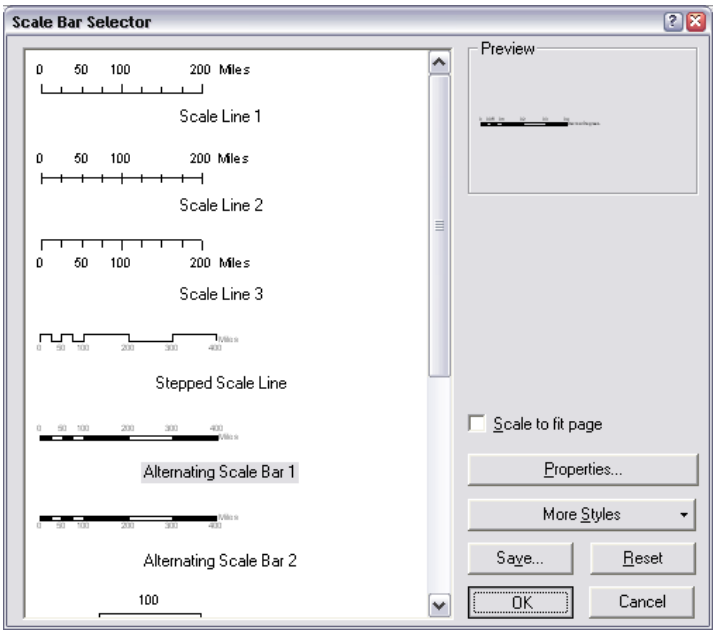

You may need to alter the settings for the scale bar. For instance, the scale bar below displays units in Decimal Degrees, a distance measurement which most people don't understand:

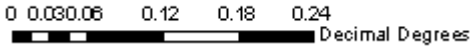

Open the Scale Bar Properties window by double-clicking the scale bar or by right-clicking it and selecting 'Properties...'

If necessary, change the Division Units:

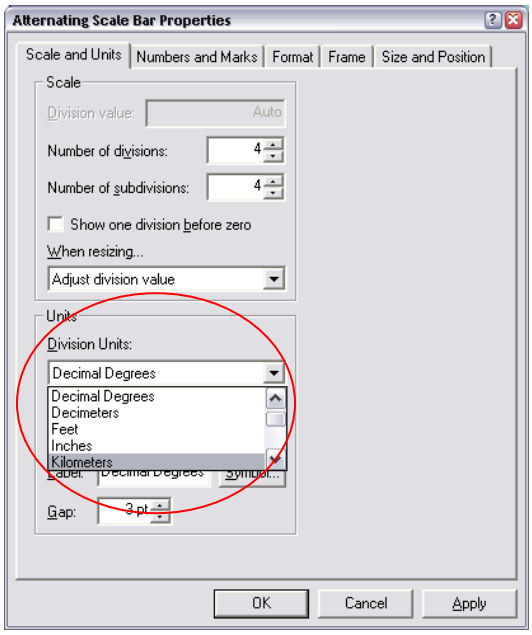

Click ‘Apply’ to see the changes:

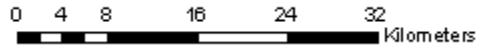

You may now wish to change the value of divisions – in the example above, it may be more appropriate to display divisions of 10km rather than 8km.

Change the ‘When Resizing...’ dropdown box to ‘Adjust number of divisions.’ Then, change the ‘Division Value’ to the desired value:

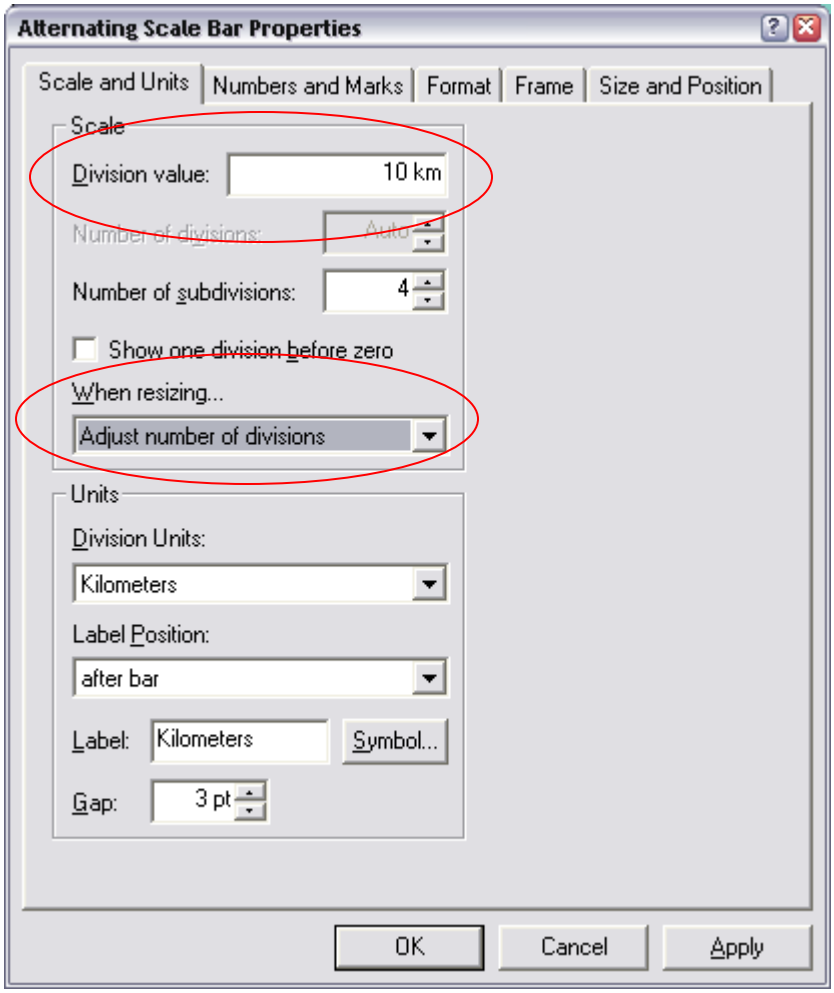

Click ‘OK’, then move and/or resize the scale bar as appropriate.