Green in the City

Atlanta, Georgia, USA

Teacher's Manual

Colophon

Imagery

Landsat, USGS and NASA

Sentinel-2, Copernicus program, ESA

TABLE OF CONTENTS

GENERAL INFORMATION

Module summary

In this module, students will examine the presence of urban green in the city of Atlanta. They will create a vegetation map of the city using satellite images and investigate the amount of green in the city as well as how it is spatially distributed. They will then compare the distribution of urban green with temperature data. Finally, students will conduct their own research and compare the distribution of urban green and temperature with various socio-economic data. This workshop is both challenging and instructive, as students will not only have the opportunity to work with satellite images in a comprehensive GIS software environment but also to formulate and answer research questions based on their interests and in consultation with the teacher.

Module structure

The workshop is divided into three parts. Part 1 covers the basics of remote sensing and familiarizes students with mapping vegetation from remotely sensed imagery. In this part, students will learn how to create a vegetation map for a city. The introductory part is comprised of four blocks, each starting with a video followed by several assignments. Each video prepares students for the exercises in which they will apply what they have learned.

Part 2 focuses on how to work with vector data in a GIS environment, map spatial distribution patterns, and analyze spatial relationships between different types of data, such as vegetation and temperature. Parts 1 and 2 of the workshop can be completed independently using the supplied worksheets. The worksheets serve as a guide and provide instructions on how to perform specific tasks using the QGIS software.

The table below provides suggested target times for the completion of each block of assignments.

		Independent / class	Time
PART 1			
	Block 1	Independent/class	45 min
	Block 2	Independent/class	30 min
	Block 3	Independent/class	30 min
	Block 4	Independent/class	30 min
PART 2	Block 5	Independent/class	60 min

Part 3 focuses on examining the relationship between vegetation in the city, temperature, and the spatial distribution of several socio-economic variables. This part of the module has a more open structure, and the intention is for students to work with the provided data, formulate research questions or hypotheses, and further investigate them based on the available data. The teacher plays a supportive and coordinating role in this process.

Group division

It is recommended that the module be carried out in teams. Students can work in pairs at a computer, watching the theoretical videos and completing the assignments together. In a strong class, students may work individually, but most students prefer to work with a partner so they can discuss the material. The results obtained can be discussed in class.

Material to be provided by the students:

- Writing utensils
- Ruler
- Calculator

Additional material to be provided by the students and/or the teacher/school:

- Access to computer(s) prepared for the workshop (see technical manual)
- Software to view the videos (format .mp4).
- A bundle of worksheets per student or per team
- One headset or earpiece set per student

If students are working in pairs at one computer, one audio splitter per PC is needed to split the single outgoing audio signal from the PC into two channels for the two headsets or earphone sets. Audio splitters, or jack splitters, can be found online and do not cost much.

THE WORKSHOP EXPLAINED STEP BY STEP

PART 1 – Make a vegetation map of Atlanta

In this part of the module, students will create a vegetation map for the city of Atlanta. They will do this by watching four videos and completing the corresponding assignments.

Block 1

THEORY

This video, which is approximately 5 minutes long, introduces students to the principles of remote sensing or Earth observation. They learn about how satellite images are captured, the concept behind measuring reflectance in different spectral bands, and how measurements within these bands are stored in multispectral images comprised of pixels.

EXERCISE 1.1

(Calculator, ruler)

In this exercise, students will familiarize themselves with a multispectral satellite image that is made up of different spectral bands, which are loaded as a raster dataset.

1) The students will use the *pan map*, *zoom in*, *zoom out*, and *zoom full* tools to observe different parts of the image and identify different features that are represented by image pixels.

2) Using the "*Identify features*" tool, the students will click on specific pixels and observe in the "identify results" window how pixel values differ in each spectral band. It is important to ensure that the students check the ***Atlanta_spectral_bands*** layer in the layers panel to be able to see its pixel values.

3) For three types of land cover (water, vegetation, and buildings), the students will look up the spectral values for three pixels in each of the four multispectral bands. They will use the "*Identify features*" tool to do this. Make sure that students only check the ***Atlanta_spectral_bands*** layer and uncheck the other layers.

   When you click on a pixel, the *Identify Results* window is displayed, with four values for Band 1, Band 2, Band 3, and Band 4. The value for Band 3 corresponds to the spectral value for the designated pixel within the blue band. Band 2, 1 and 4 correspond to the green, the red and the near infrared (NIR) band respectively. You should always write down the four pixel values in a row; you will therefore have three rows per land cover type (one row per pixel).

Students may find it difficult to locate water on the satellite image. This can be due to the fact that it is not always clearly visible in every city. In some cases, there may be a channel that runs through the city, and students may sample pixels within the channel. While this is not incorrect, it is better to choose water pixels in ponds or lakes. It is explained later in the manual (Exercise 1.3) why this is a better choice. Students can use Google Earth to get a better view of land cover within the city and identify locations belonging to each land cover type, especially water surfaces.

4) You have three measurement values per band for each type of land cover. To obtain the average spectral value, you need to calculate the mean of these three values using a calculator.

5) You can make a graph based on the calculated average spectral values for each land cover type. If the points are then connected for each type of land cover, you obtain a curve referred to as the spectral signature of the land cover type. Each type of land cover will reflect electromagnetic radiation to a greater or lesser extent, depending on the specific physical properties of the matter that you find within a type of land cover. Each type of material has its own characteristic spectral signature. The video in Block 3 delves further into this topic.

Tip: You may choose to conduct this exercise in class by preparing the same blank chart on the board or in an Excel spreadsheet. Divide the class into three groups, with each group performing the exercise for one of the land cover types. Once the students have completed the task, they can write their four mean values on the board or enter them into an Excel spreadsheet. Then, the graph can be drawn on the board or automatically generated in the Excel spreadsheet you have prepared. This approach can save time.

EXERCISE 1.2

This exercise encourages students to reflect on what they learned in the video. Understanding the meaning of low and high pixel values prepares students for the questions in EXERCISE 1.3.

The value of a pixel always represents a measurement of electromagnetic radiation. When a satellite image is captured, it measures the amount of reflected electromagnetic radiation from the portion of the Earth's surface covered by a pixel. If that part of the Earth's surface reflects the incoming electromagnetic radiation from the sun strongly, the corresponding pixel will have a high value. Conversely, if there is low reflection, the pixel will have a low value.

Note: This applies to each band. A band corresponds to a specific range of wavelengths in the electromagnetic spectrum. Within each band, the reflected radiation is measured, resulting in a raster of values. For instance, a low pixel value in the ***Atlanta_NIR*** layer means that the Earth's surface within that one pixel reflects near-infrared light to a lesser extent.

EXERCISE 1.3

Water is a strong absorber of electromagnetic radiation, which can be deduced from its relatively low pixel values, indicating low reflection. When a piece of the Earth's surface reflects less energy, it absorbs more of it.

Why should students avoid selecting pixels from a channel in **EXERCISE 1.1**? The pixel values in the channel are generally higher than average values for water. Drawing a curve based on these values would result in a graph that is skewed towards higher values, giving the impression that water is more reflective than it actually is. This is because a channel often contains pollutants that alter the spectral behavior of the water, causing it to reflect more electromagnetic radiation. As a result, pixels in a channel will often have higher values than those in other bodies of water and including them in the analysis would distort the spectral signature.

Vegetation reflects less energy in the blue and red parts of the spectrum compared to the green part because the energy in the blue and red parts of the spectrum is utilized for photosynthesis. In the near-infrared part of the spectrum, the spectral value is significantly higher due to the interaction of the infrared radiation with the internal leaf structure. However, it is not necessary to provide this information to the students at the moment. It will be explained in greater detail in Video 3.

Buildings are characterized by a spectral curve that shows relatively little variation in the visible and near-infrared parts of the spectrum. Unlike other land cover types, the spectral values in the four bands do not differ that much from each other.

Block 2

THEORY

In this 6-minute video, students will learn how to create a true-color image by combining the available multispectral bands. This is clarified using the RGB color coding system. The video will also demonstrate how to produce a false-color image, more specifically, a so-called 432 color image.

EXERCISE 2.1

In this exercise, students will open the ***Atlanta_spectral_bands*** raster layer that combines the four spectral bands of the image. They will see how three of the four spectral bands are used to produce a true-color image, similar to the one shown in the video.

The reflectance values of each image pixel in the blue, green and red part of the spectrum are stored in the layers ***Atlanta_Blue*** *(Band 3: B1_blue)*, ***Atlanta_Green*** *(Band 2: B2_green),* and ***Atlanta_Red*** *(Band 1: B3_red) respectively*. By linking these files to the corresponding colors on the computer screen ("Blue band", "Green band", "Red band") the intensity of each of the three colors in the image will be defined by the spectral value of each image pixel in the corresponding band.

The students should know that "Red band", "Green band", and "Blue band" on the Symbology window refer to the colors that are used to visualize the spectral reflectance in each of the three bands. Band 1, Band 2, and Band 3 correspond to the layers containing the spectral values for each pixel in the blue, green and red part of the electromagnetic spectrum.

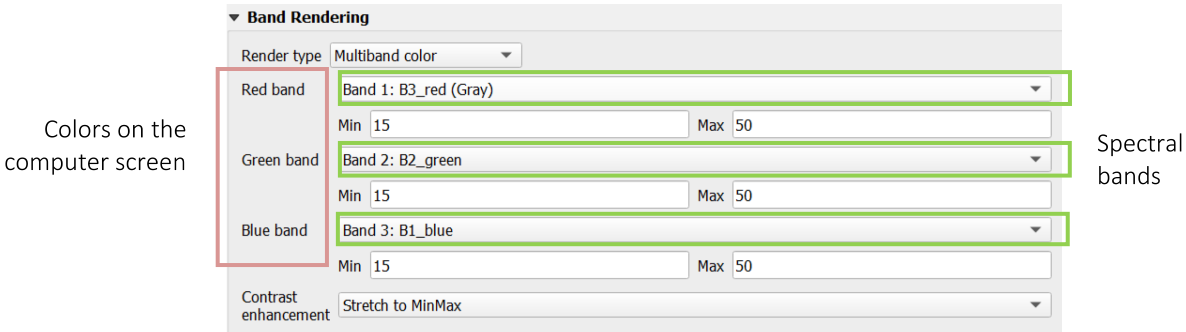

The next step is for the students to examine the true-color image once again and determine the color of the three land cover classes present in the image. Then, using the spectral curves they generated in EXERCISE 1.1, they must provide an explanation for the color of each land cover type.

Water reflects little energy in the visible parts of the spectrum, particularly in the blue, green, and red bands. As a result, water appears very dark (almost black) in the true-color image.

Healthy vegetation reflects the green wavelengths more than the blue and red wavelengths in the visible spectrum, resulting in higher reflectance values in the green band. For each pixel covered by vegetation, a color will therefore be obtained in which the green color component predominates over the blue and red color components. As a result, vegetation looks green.

Buildings appear grayish in the image because their spectral curve shows little variation in the visible part of the spectrum. When the blue, green, and red color components that form a true-color image have equal intensities, the resulting color is gray. The video also explains this concept.

Note: The interaction between matter and radiation is not always constant. For instance, a volume of water will absorb more red (and green) light as the depth of the water increases. This absorption occurs due to the interaction of electromagnetic radiation with water molecules. Due to selective absorption, a greater proportion of blue light is reflected from the surface of water, resulting in its apparent blue color. If one goes even deeper into the volume of water, all the radiation is absorbed.

EXERCISE 2.2

A false-color image is a type of multispectral image produced using at least one non-visible wavelength, such as near-infrared. Unlike true-color images, false-color images can visualize wavelengths that are not visible to the human eye, and this can enhance the interpretability of the data. In scientific research that analyzes vegetation using satellite images, false-color images are often used because they allow vegetation to be more easily distinguished.

There are many different band combinations used for producing a false-color image, and each combination is used to highlight different features. In this exercise, we choose NIR for the red band, Red for the green band, and Green for the blue band, which is a widely used combination to study plants.

In EXERCISE 2.1 students learned how a true-color image is created and how to interpret the spectral characteristics of different land cover types to explain their color in the image. In this exercise, they need to link the available spectral bands to the three main colors of the screen in a different way to create the requested false-color image. If they cannot recall how to do this, they can watch the video again. The students have to link *"Band 4: B4_NIR"* to the intensity of the red color component on their screen, "*Band 1: B3_red (Gray)*" to the green color component, and *"Band 2: B2_green"* to the blue color component, as shown below.

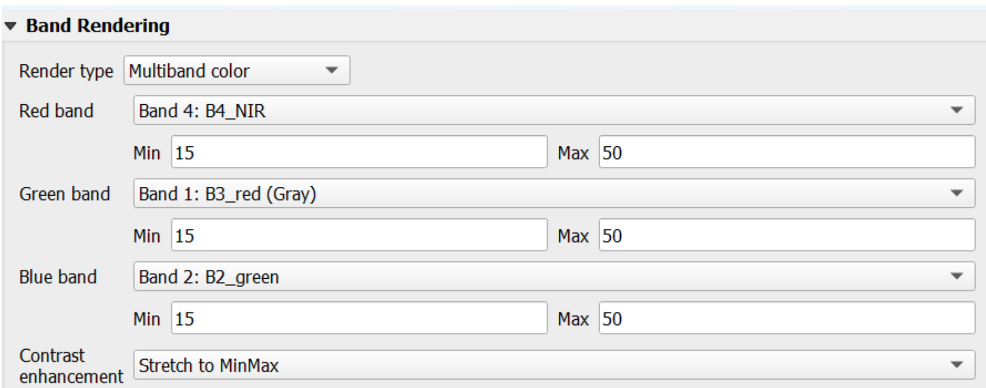

It may be useful to remind students again that "Red band", "Green band", and "Blue band" on the Symbology window refer to the colors that are used to visualize the spectral reflectance values of each pixel in the image. Band 1, Band 2, and Band 4 correspond to the layers containing the spectral values for each pixel and need to be linked to the proper color by the students.

EXERCISE 2.3

You can observe that the false-color image effectively highlights the presence of vegetation by showing green vegetation as a bright red color. The question that arises is why vegetation appears so intensely red in the false-color image. Despite the two hints provided (the band assigned to the red color component and the graph from **EXERCISE 1.1**), this question remains difficult for some students.

The explanation is as follows: In the false-color image, you assign the NIR band of the multispectral image (Band 4) to the red color component of the display. Thus, the pixel values of the near-infrared image determine the intensity of the red color component in the multiband image. In the graph from **EXERCISE 1.1**, it can be observed that vegetation strongly reflects near-infrared light. Therefore, for pixels corresponding to vegetation, a high measurement value is recorded in this part of the spectrum. That high reading is assigned to the red color component.

An RGB color is a composite color, and the green and blue color components also have their say. These color components are determined by the values of *Band 1 (Red) and Band 2 (Green),* respectively. However, relatively speaking, vegetation reflects near-infrared light much more strongly than red and green light. In other words, for each vegetation pixel, an RGB color will be obtained in which the red color component predominates the other two color components. Hence, the rather bright red color for vegetation. The image below shows the Atlanta Hemphill Water Treatment Facility in Atlanta City. The dark color corresponds to the lakes, while the red color shows the vegetation.

False-color image

This explanation may still be unclear for some students. It is important to understand that the students should not be distracted by the colors in the names of the spectral images. Instead, they should focus on the fact that they assign spectral values of particular bands (reflectances) to the different color components of the false color image.

Using a culinary analogy can make the concept of values or measurements that form colors more relatable. Let's say you're making soup and you have three ingredients: red bell pepper, leek, and carrot. These three ingredients represent the three color components of a hypothetical red-green-orange color coding system. We will assume that each ingredient colors the soup equally. If the soup consists of 80% carrots, 10% red bell pepper, and 10% leeks, the mixture will have a bright orange color because the carrots are given a much higher value in relative terms. This is a simple yet effective example.

To test whether the students have truly grasped the concept, you can ask the following question: How can I create a false-color image where vegetation is shown in blue? The answer is to assign Band 4 (NIR) to the blue color component. Since vegetation reflects a high amount of near-infrared light, assigning the corresponding spectral image to the blue color component will turn the vegetation in the image blue. For the red and green color component, the options are Band 1 (Blue), Band 2 (Green), and Band 3 (Red).

Give it a try in QGIS.

Block 3

THEORY

In this approximately 5-minute video, we will explore the unique spectral signature of vegetation and how it can be used to compute a vegetation index. The most widely used vegetation index is the NDVI (Normalized Difference Vegetation Index). The video will demonstrate how to compute this index and explain the insights that can be drawn from an NDVI map.

EXERCISE 3.1

In this exercise, students will create an NDVI map using the *Raster Calculator* tool. The student worksheet provides instructions on where to find the tool in the interface and how to use it. The result is a new raster layer that is loaded into the layers panel. The NDVI map of Atlanta is shown in the image below.

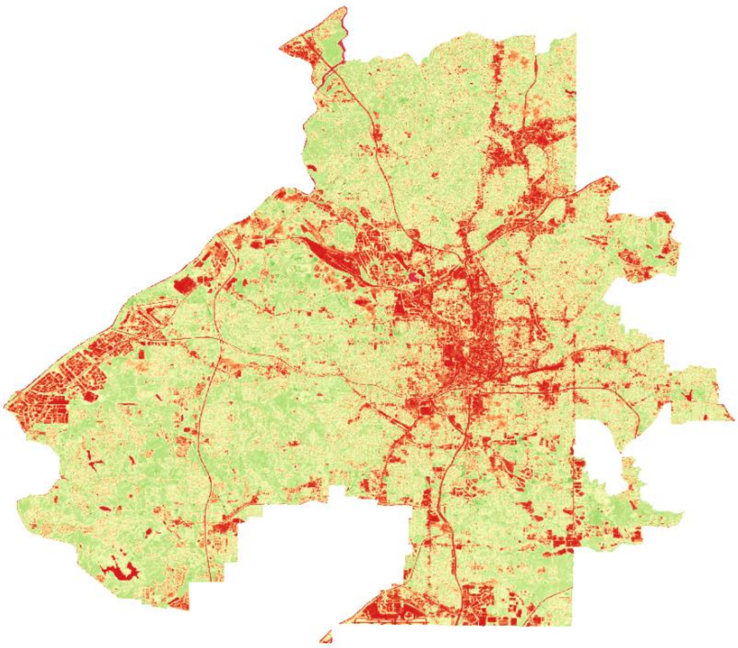

The following mistake is often made by students:

> The expression to calculate the NDVI value is provided in the worksheets: (NIR-RED)/(NIR+RED). Please avoid typing the formula as it is written on the worksheets into the white field of the *Raster calculator expression*, as it will result in an error.

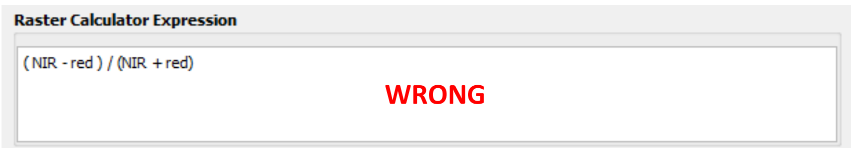

Students must keep in mind that they will be using raster layers as calculation elements, rather than typing the formula itself. For example, in the correct formula below, the calculation element ***Atlanta_NIR@1*** found under **Raster Bands** refers to the raster layer that contains the NIR values for each pixel.

Raster Calculator Expression

```
( "Atlanta_NIR@1" - "Atlanta_Red@1" )  /  ( "Atlanta_NIR@1" + "Atlanta_Red@1" )
```

CORRECT

Another mistake is not using parentheses in the expression. This can lead to erroneous NDVI values. If the NDVI map is calculated and added to the layers panel, you should get values between -1 and +1. You can check this in the legend of the map that you find in the layers panel.

EXERCISE 3.2

To answer this question, it is essential for students to have applied the color scheme that shades pixels from red to yellow to green, making it easier to read the map. Vegetation typically has the highest NDVI values, while streets, buildings, and water have lower NDVI values. The purpose of this exercise is for students to interpret the map and connect high NDVI values with vegetation. This understanding will aid them in comprehending the content of the subsequent video.

Block 4

THEORY

This approximately 4-minute video demonstrates a simple way to create a vegetation map using an NDVI map. First, the students are given an explanation of what a histogram is and how to construct one. To ensure the method is easy to understand, the flowchart for calculating the vegetation map is shown twice.

EXERCISE 4.1

In this exercise, students will create a vegetation map for Atlanta based on the histogram of the NDVI map. Double clicking on the ***Atlanta_ndvi*** layer opens the properties window for this layer. A click on the *Histogram* tab displays the histogram of all pixels in the NDVI map. If the histogram is not there yet, click the *Compute Histogram* button at the bottom of the screen. The idea is then to draw an imaginary vertical line between the two peaks of the graph, where the histogram reaches its lowest value between the two peaks. Where this vertical line intersects the X-axis (pixel value) is the threshold value. Based on this value, you can say that all pixels with a value less than the threshold value, or everything to the left of the vertical line, corresponds to what is not vegetation, and all pixels with a value greater than the threshold value, or everything to the right of the vertical line, corresponds to vegetation. So, we make a simple distinction between vegetation and non-vegetation based on a chosen threshold value.

While finding the NDVI value that corresponds to the dip in the histogram is a simple way to separate vegetated from non-vegetated areas, it is not necessarily the most optimal threshold value. Very often, the optimal threshold value must be defined by trial-and-error. In the case of Atlanta, the dip in the curve occurs around an NDVI value of 0.18. If we use 0.18 as a threshold and compare the vegetation map that we obtain with the true or false color image for Atlanta and zoom in on both layers, we will notice that many of the houses located in the suburban areas around the city center will be covered by vegetation in our vegetation map. By increasing the NDVI threshold value to 0.3, as shown in the figure below, the vegetation map obtained corresponds much better with the pattern of vegetation and built-up areas which we observe in the true and false color images. In other words, an NDVI threshold of 0.3 is more appropriate to create the vegetation map for Atlanta.

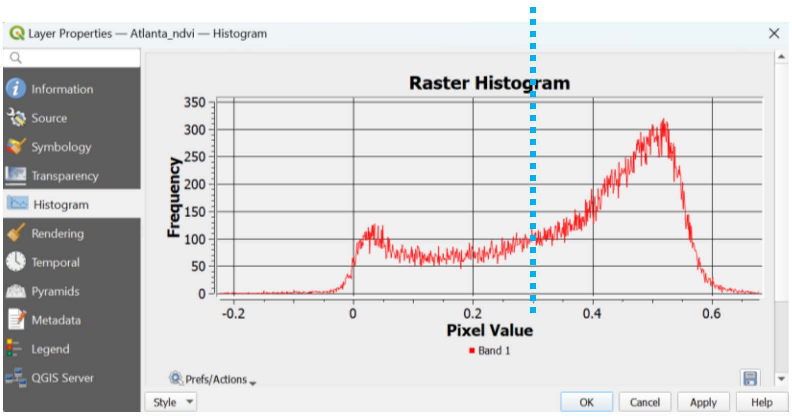

Note: The blue dotted line in the example above isn't part of the QGIS window screenshot; it's an illustrative line you can envision by holding a pen or pencil against the screen.

The calculation needed to apply the threshold is done with the *Raster Calculations* tool, which was also used for the creation of the NDVI map. The approach for producing the vegetation map is analogous to that of the NDVI map. The image below illustrates the vegetation map obtained for Atlanta.

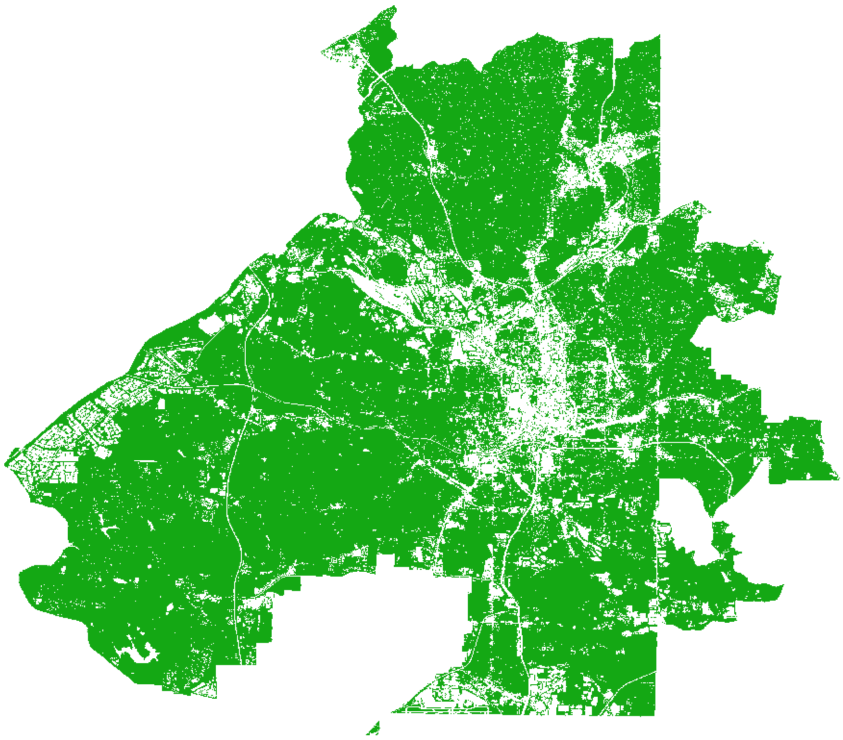

Since the procedure is almost the same as for the NDVI map, the possible errors that students can make are also analogous to the errors detailed in EXERCISE 3.1. Again, it is important to point out that raster maps are used as calculation elements. In the formula, the calculation element ***Atlanta_ndvi@1*** found under Raster Bands refers to the raster layer that contains the NDVI values for each pixel.

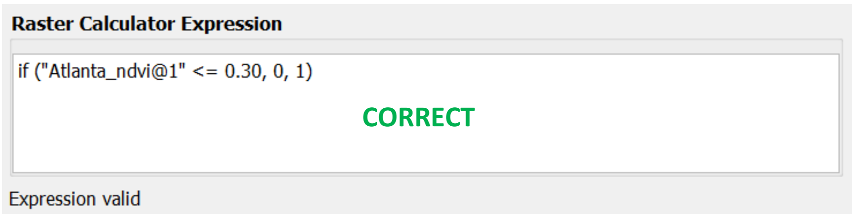

Note: The vegetation maps obtained by different students (or groups of students) may differ. As explained above, the map is based on a threshold value that the students determine. A lower threshold value will result in the identification of more vegetation pixels (see formula). A higher threshold value will lead to fewer pixels being labelled as vegetation.

EXERCISE 4.2

This simple exercise asks students to write down what they see on the map — do they notice any specific patterns? Do they recognize some of the forests or parks in the area?

PART 2 – Comparing green cover with temperature

This part focuses on learning how to work with vector data, map spatial distribution patterns, and analyze spatial relationships between different types of data, such as vegetation and temperature data.

Block 5

EXERCISE 5.1

In this exercise, students are introduced to the ***Atlanta*** vector dataset, which contains polygons representing all census tracts located within the city. They learn that each vector layer in a GIS has an attribute table that describes the characteristics of each polygon. Once the table attached to a vector file is open, you can click on any row in the table to select a polygon. You can also select multiple polygons at once by holding down the Ctrl key while selecting different rows in the table. Remember that to select a polygon, you have to click on the row number of the polygon that is on the far left of the table. The entire row will then be colored blue. If you look at the map, you will see that the selected polygons are also colored yellow on the map.

The selection principle also works in reverse. If you select a polygon on the map, which can be done via the *Select Features* tool in the search bar above the map, the polygon will also be highlighted in the table. From there, you can access all the information about the polygon. You will notice that there are several methods to select polygons on the map (within a circle, etc.).

As long as a polygon is selected, it remains marked in blue in the table and marked in yellow on the map. If you want to remove the selection, you must undo it. Both in the table and above the map canvas, you have a button in the search bar to undo the selection (*Deselect*).

Note: What often happens is that students forget to deselect polygons and afterwards do not know how to get rid of the yellow color of a selected polygon in a later phase of the exercise, when they start making maps. So, deselect.

Zonal statistics

The *Zonal Statistics* tool allows for the combination of a vector layer (in this case, the tracts) and a raster layer (in this case, the vegetation map derived from the remote sensing image), and the calculation of a statistic within each polygon of the vector layer. This calculation is based on the value of the overlapping pixels in the raster layer. Since the raster layer consists of pixels that have a value of 0 (no vegetation) or 1 (vegetation), calculating the average value of all pixels within each polygon will indicate the proportion of the pixels within the polygon that are covered with vegetation. The result of the calculation is added as an extra column named ***veg_mean*** to the attribute table of the vector map.

Note: Make sure students choose only the mean as a statistic. If they do not do this, other statistics are also calculated that needlessly add extra columns to the attribute table. This in itself is not a problem, but it may be confusing for the students.

EXERCISE 5.2

In this short exercise, the students will open the attribute table of the ***Atlanta_veg_mean*** layer to identify the tracts with the highest and lowest proportion of vegetation in the city. To do this, they have to look for the ***veg_mean*** column in the attribute table, which is added at the very end. Because the table contains many tracts, students should not go through the entire table tract by tract to find the one with the highest and lowest value. Instead, they can sort the tracts based on the values in one column by simply clicking on the name of the column at the top of the table. You can sort the column from low to high values, or vice versa. By repeatedly clicking on the name, the tracts are ordered alternately from lowest to highest and highest to lowest value. The numbers of the tracts that they must also note can be found in the first column, *Tract*.

EXERCISE 5.3

This is an important exercise, as the students will learn how to make a thematic map based on a chosen column in the attribute table (in this example, the proportion of vegetation, ***veg_mean***). The census tracts are grouped into classes (from low to high value tracts), and each class is displayed in a specific color.

Duplicate layer

Once the extra column with the proportion of green per polygon has been added to the attribute table of the ***Atlanta_veg_mean*** vector layer, we will make a copy of this active layer (right-click the layer in the layers panel and click *Duplicate Layer*). The layers panel now includes the new ***Atlanta_veg_mean_copy*** map layer. We will then rename this new map layer ***Atlanta_tracts_vegetation*** by right-clicking on it, selecting Rename Layer, and changing the name in the layers panel. This makes it possible to map the proportion of vegetation within each tract in this new map layer, without making any changes to the original map in the ***Atlanta_veg_mean*** vector layer.

Create map

To make a map, students have to make three choices:

a. Choosing a suitable color scheme: if we want to map data where the values vary from low to high, we usually choose a single-color scheme and let the color vary from light (low values) to dark (high values). We should also make sure that the color choice is logical, depending on what we want to map. For mapping vegetation, using a color scheme from light to dark green is a good choice.

b. Choosing a suitable number of classes: cartographers often choose a number of classes that varies from 5 to a maximum of 7. The fewer classes we choose, the better we will be able to perceive a spatial pattern; the more classes we choose, the more difficult this becomes. An odd number of classes creates a middle class or neutral class, neither high nor low, which is also generally preferred for a more balanced and easier to interpret map image.

c. Choosing a method to classify the spatial units (in this case, tracts): cartographers make use of different methods to subdivide data in classes, but when it comes to comparing patterns in different maps (which is the purpose of this module), the *quantile method* is usually chosen. The boundaries between the successive classes are defined in such a way that the same number of spatial units end up in each class.

The image below shows the proportion of vegetation per tract for Atlanta, as well as the corresponding legend that you can find in the layers panel under the name of the map layer after creating the map.

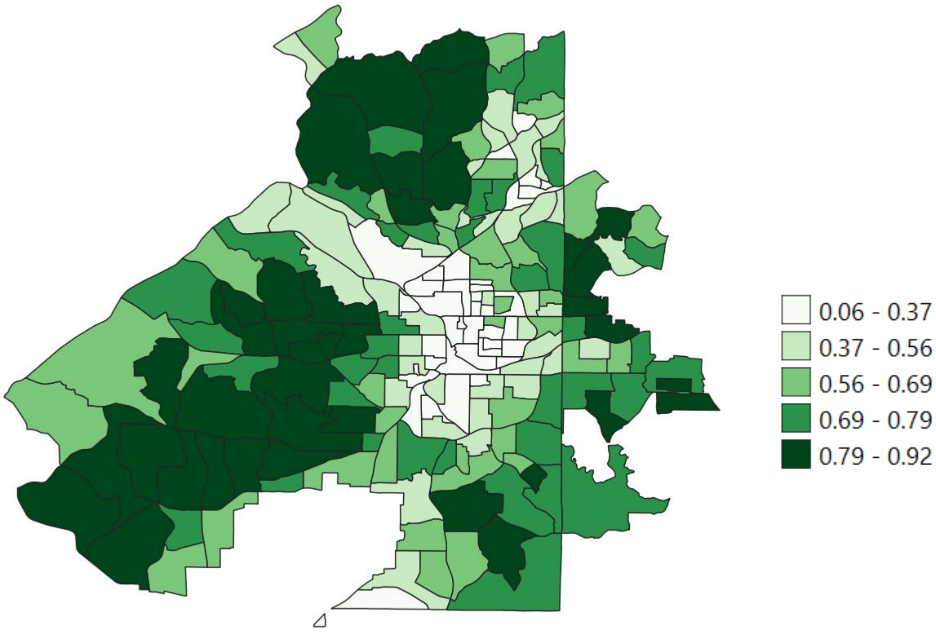

EXERCISE 5.4-5.5-5.6

The temperature values in the ***Atlanta_temperature*** raster layer are presented in degrees Celsius (°C), so we will convert the values to degrees Fahrenheit (°F). Make sure that students use the *Raster Calculator tool* for the conversion and save the new file as ***Atlanta_temp_fahrenheit*** in the Atlanta folder.

Then EXERCISES 5.1-5.2-3.3 are repeated, this time starting from the ***Atlanta_temp_fahrenheit*** raster layer instead of the ***Atlanta_vegetation*** raster layer to obtain a map of the temperature distribution at tract level. The method is completely identical. The result is a new ***Atlanta_tracts_temperature*** map.

Again, when creating the map, 5 classes are chosen, and the quantile method is used to divide the tracts into classes. However, this time the color scheme is a progression from light red to dark red.

The image below shows the average temperature per tract in Atlanta, as well as the corresponding legend that you can find in the layers panel under the name of the map layer after creating the map.

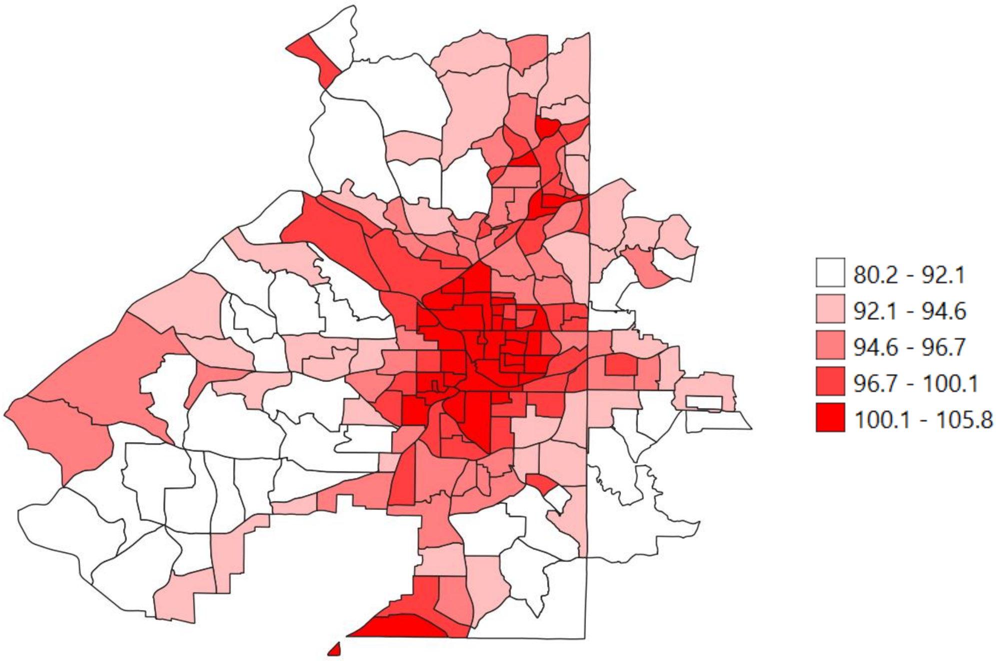

EXERCISE 5.7

Students will now compare the ***Atlanta_tracts_vegetation*** and ***Atlanta_tracts_temperature*** maps they previously created to explore the relationship between the proportion of vegetation and average temperature in the census tracts. A comparison of both maps clearly shows that where there is more vegetation, the average temperature is also significantly lower. This points to the cooling effect of vegetation in the city, which is one of the important ecosystem services provided by urban green.

To make the comparison of both maps easier, make sure that all map layers are unchecked except for the two you want to view. By clicking on and off the "top" layer in the layers panel, you can alternate between the two maps for easy comparison.

Before doing EXERCISE 5.7 in class, you may want to ask the students how they think vegetation and temperature in the city relate to each other. A comparison of the maps may or may not confirm their judgment or hypothesis. Or you can let them discover it themselves.

EXERCISE 5.8

This exercise introduces the use of a scatterplot to analyze the relationship between two variables (proportion green – mean temperature). The scatterplot can be saved as an .html file and viewed later in your browser. You can also view the plot directly in QGIS (Toolbox window on the right, under *View Results*, double-click the plot).

For Atlanta, the scatterplot looks like this:

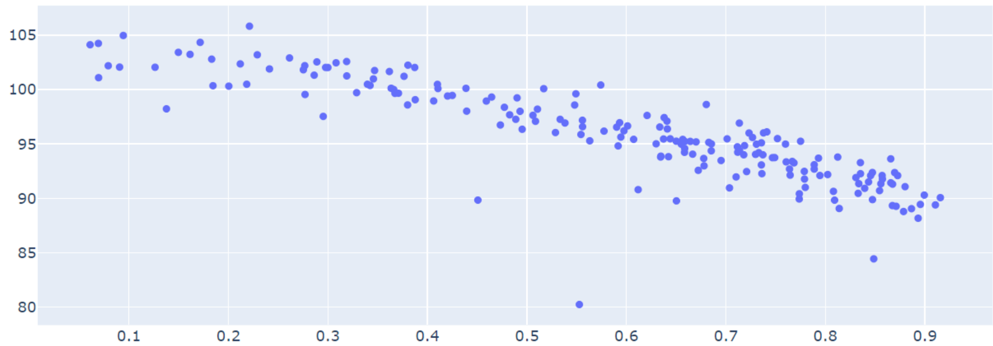

The plot clearly shows the inverse relationship between the proportion of green (x-axis) and the average temperature in each tract (y-axis).

EXERCISE 5.9

This exercise returns to the first question students were asked at the beginning of the workshop about the estimated percentage of urban green in the City of Atlanta. In this exercise, the students answer this question with QGIS. Make sure the students use the ***Atlanta_boundary*** polygon layer as the input layer when applying the *Zonal Statistics* tool. You can briefly discuss how well the students estimated the proportion of green in Atlanta, when they answered the question at the start of the workshop.

Unless you choose another method, QGIS shows the results of the majority of the analyses you can do with the software as a temporary layer in the layers panel. If you do not save the layer to disk, the temporary layer will be lost permanently when you exit QGIS. Temporary layers can be saved to a file with a different name by clicking on the icon on the right side of the layer's name in the layers panel.

PART 3 – Investigate the relationship between urban green, temperature and socio-economic data

This last part of the module has a more open structure. The intention is that the students themselves get to work with the data that is provided to them, formulate research questions or hypotheses, and further investigate these based on the available data. The data enables them to explore relationships between urban green, temperature, and a number of socio-economic variables.

FORMULATE RESEARCH QUESTIONS

A number of socio-economic data are available in the attribute table of the ***Atlanta_veg_temp_mean*** vector layer that the students produced during the previous exercises. The socio-economic data relate to the characteristics of the population living in the census tracts (density, age distribution, income, level of education, etc.) and to the characteristics of the dwellings in each tract (dwelling type, age). The aim is that, based on the available data, the students define one or more research questions themselves, formulate hypotheses, and then check whether their hypotheses are confirmed by the available data. Some examples of research questions are found below.

Example 1

Research question: Is there a spatial relationship between urban green abundance and income?
Hypothesis: greener neighborhoods attract higher incomes as a result of higher housing and rental prices; lower income families have less access to green residential areas.

To test the hypothesis, students can make a map of the median gross annual income per tract based on the skills they learned in Part 2 of the module and compare the resulting map with the map showing the proportion of green per tract. They can also produce a scatterplot to graph the relationship between the two variables.

Vegetation not only has a physically regulating role within the city (air quality, urban climate, water management, biodiversity), but also plays an important role in the general well-being of the population: people may appreciate a green environment to relax and escape from busy city life. Current scientific research examines the importance of these so-called ecosystem services provided by urban green. Socially just urban planning should therefore start from the principle that every inhabitant of the city is entitled to the benefits of green in their immediate living environment.

If a clear relationship is established between the presence of green space and income, a class discussion can be linked to this with regard to concepts such as environmental injustice (unequal distribution of environmental quality that mainly affects poorer populations) and green gentrification (the fact that providing more green space in the city often does not solve the problem but leads to an increase in house and rental prices that forces poorer populations to leave the neighborhood in the long term).

Example 2

Research question: Is the population living in parts of the city where the average temperature is noticeably higher during warmer periods also more vulnerable to the consequences of global warming?
Hypothesis: The share of older homes, which are generally less well insulated, is higher in densely built parts of the city where the average temperature is also higher.

To investigate the hypothesis, students can use the available data on the age of homes and create a map showing the proportion of homes per census tract built before 1960. In this case, the age of the homes serves as a proxy, indicating the presence of poorly insulated homes.

Examining the degree of insulation in homes is just one way to assess vulnerability to higher temperatures during heat waves. Other variables that can indicate vulnerability include the proportion of elderly or young people in the population. These age groups are most at risk of experiencing health complications during periods of persistently high temperatures. The dataset contains several variables related to the proportion of young and old within the population that could be used in such an exercise (with different age thresholds). Income and educational attainment are also indicators that can be linked to higher vulnerability.

Tip: When charting income, you may notice that some tracts are not shown on the map. This is because income data for these tracts is not publicly available. This does not necessarily mean that there are no people living in these tracts or that there is no data available for these tracts. However, some tracts have very few inhabitants. To respect privacy, the U.S. Census Bureau applies restrictions to disclose income data for tracts with few inhabitants. These tracts are therefore not colored, and if they are on the edge of the map, their boundary is not shown either. To clearly distinguish tracts without data from those with available data, you can follow these steps:

a. Make sure that your new map layer is positioned "above" the original ***Atlanta_veg_temp_mean*** vector layer in the layers panel. You do this by dragging the new map above the ***Atlanta_veg_temp_mean*** vector layer in the layers panel.
b. Adjust the symbology of the ***Atlanta_veg_temp_mean*** layer via the *Symbology* tab. Choose the Single symbol and gray fill as color.
c. Check both the new map layer and the ***Atlanta_veg_temp_mean*** vector layer so that both are visualized. The tracts without income data will appear gray on the map, as they are not colored in the new layer. This allows you to easily distinguish the tracts without data from the other tracts.

The image below shows an example of an income map for Atlanta. Tracts for which no data is available are grayed out by coloring all tracts in the ***Atlanta_veg_temp_mean*** map layer placed below it in gray.

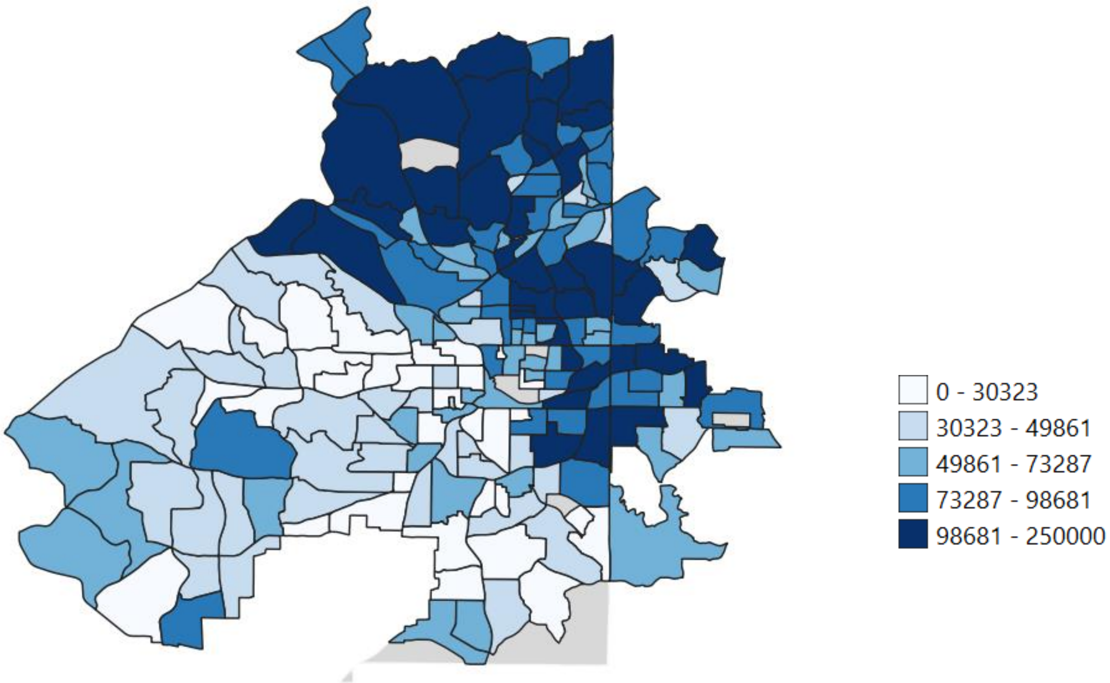

Note

The socio-economic data used in this module comes from the most recent Census of 2020.

Income statistics are publicly available, and anyone can download them. However, some tracts have very few residents. The U.S. Census Bureau has confidentiality requirements to protect the privacy of individuals and households. Based on these requirements, if a census tract has a population lower than a certain threshold, some demographic data, including the income levels of the population, are not released.

Finally, this...

The presence of certain variables in the attribute table does not necessarily imply that these variables all have a clear relationship with urban green provision and temperature. The goal of the research is to investigate relationships and validate or reject hypotheses.

FINISHING AND EXPORTING A MAP

The worksheets also contain instructions for creating a polished map and saving it as a .png file. The image file can then be used in PowerPoint to make a poster.

The layout of a map must always contain the following elements: map, title, legend, scale, and north arrow. The secret to a beautiful map design is:

a. Make optimal use of the available space on the layout page.
b. Nicely align all elements on the map with respect to each other and with respect to an imaginary frame that encompasses the entire map.
c. Provide sufficient free space around each element (i.e., do not place elements too close together).
d. Do not display the different elements too large in proportion to the map itself.

Encourage your students to produce a polished map design based on the above basic rules.

The image below gives an example of what such a design could look like for a map showing the median income per tract in Atlanta. Note that in this layout, the legend title and the gray legend item for tracts with missing data have been adjusted as described in the worksheets (via *Item Properties* > *Items for Legend*).

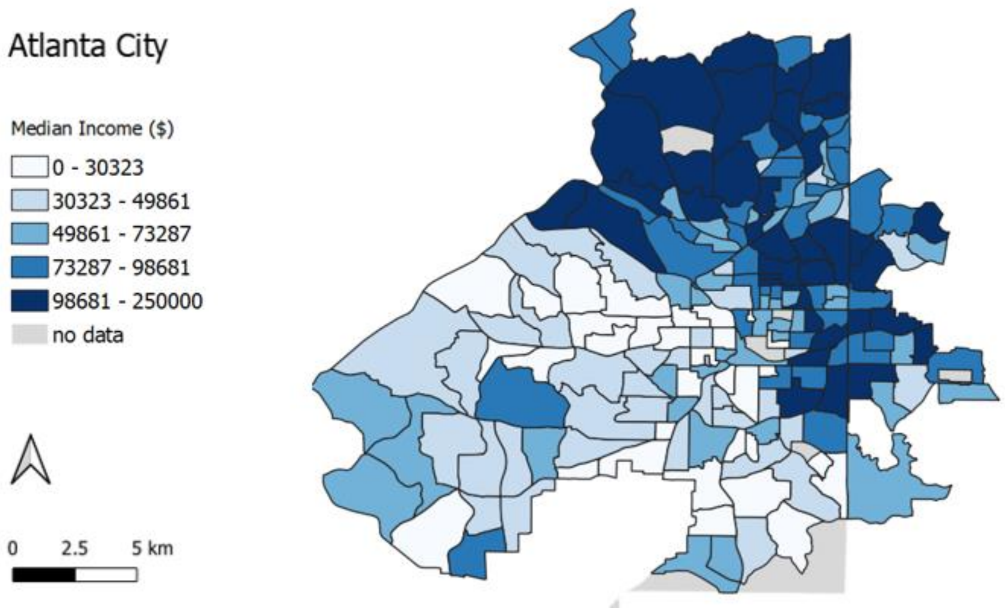

<u>Tip</u>: Students may find it confusing that in order to modify a map element, they must first click on the map element (for example, the legend) in their layout window. Only then will they be able to see the items specific to that map element in the items window on the right. This may seem logical, but students often forget this.

EXPORTING AN ATTRIBUTE TABLE TO EXCEL

A vector map's attribute table can be exported in a format that can be read directly into Excel. Working with the data in Excel offers students many more possibilities for creating various interesting charts from the available data. While QGIS is mainly suitable for creating maps, Excel provides additional tools for data analysis and visualization.

For example, after exporting the attribute table to Excel, students could:

- create a column or bar chart of the tracts with 0-20% green, 21-40% green,… (or a similar diagram for average temperature).

- calculate the average temperature across all tracts with 0-20% green, 21-40% green, and so on, and display the results in a chart.

There are numerous possibilities that you can explore with your students, which can be used to create compelling research posters that combine map and diagram visualizations.

AND FURTHER...?

You could also go a step further and, together with your students, look for additional data related to the Green in the City theme, in addition to the data provided in this module. You can use the QuickOSM or QGIS OpenStreetMap Plugin to access OpenStreetMap data. These plugins allow you to search for and download GIS data for Atlanta from OpenStreetMap, which is a collaborative mapping project that contains user-generated data for cities around the world. You can also search for GIS data on specialized websites such as the Georgia Geospatial Data Clearinghouse, which provides access to a variety of GIS data for the state of Georgia, including data for Atlanta. Once you have downloaded the data from these websites, you can load it into QGIS.

Working with additional datasets may require some extra skill in handling data within QGIS as this data will often cover a larger area than just the city boundaries. But that can also be a fun challenge. Your GIS knowledge and that of your students can only benefit from it.