GEO/EVS 425/525 Unit 4
Operations on Raster Images

In this exercise, you practice several different kinds of actions on raster images. These will involve three different classes of actions: actions on individual pixels, actions on local neighborhoods, and actions on regional neighborhoods.

Operations on Individual Pixels

Changing the ability of an image to communicate using the raster attribute editor: First, let us try the simplest sort of action on individual pixels improving the visibility of image by using the raster attribute editor. Copy the image LANDLC.IMG to your X: drive, and then open it. This is a land-use class model. The colors which appear in the image are designed and to be reasonably close to those which would have been found in the original thematic mapper image from which this land-use class model was derived. Open the raster attribute editor. You will notice that there are 15 land-use classes. Because the colors are what they are, they do not stand out legibly. By clicking on the color swatches in the raster attribute editor, change the colors for each land-use class to colors that better express the actual use of the land.

In doing this simple action, you have not changed any of the data in the image. However, you have made a major change in the ability of the image to express information. It is much easier to distinguish land uses now than it was previously. Save this image. You will need it later.

Generating new data using known relationships: Copy the image, MTWHITNEY.IMG into your X: drive, and load it into the viewer. *Be sure that you open it as a pseudocolor image*. This is a topographic image, commonly known as a Digital Elevation Model, or DEM, of part of the area to the west of Death Valley, California, on the east, across the Sierra Nevada and into the Central Valley to the west. Use ImageInfo to determine as much as you can about the image. It includes areas both above and below sea level, including the highest and lowest points in the lower 48 states. Use the Inquire Cursor function to get an idea of the variation in elevations from place to place.

As you know, the temperature of the land surface is a function of the elevation. We can approximate the relationship between the temperature at any point in time and elevation with the general equation:

AirTemp = GaugeTemp - (LapseRate * (ActualElevation - GaugeElevation))

The adiabatic lapse rate describes the theoretical relationship between elevation and temperature. It varies between 0.0036 °C / m for moist air and 0.01°C / m for dry air. The average value is 0.0066°C / m. For our purposes, we will use a higher value than this, 0.009 °C / m, closer to the dry lapse rate, since this is a fairly dry area. We shall also assume that the reference gauge is at an elevation of 150 m above sea level – approximately the elevation of Fresno. If we ask, therefore, what the temperature in this area when the sea-level temperature is 30°C, the temperature of each point in this area can be estimated by the specific formula:

AirTemp = 30 - (0.009 * (Actual Elevation - Gauge Elevation))

Since the attribute shown in the image is elevation in meters, we can easily use this equation to add a “temperature” attribute using the attribute editor.

Invoke the Raster Attribute editor, and click on the Column Properties icon. Choose to add a new column, and make its name editable. Give it a suitable name (e.g. Temperature), and fill the columns using the formula, 30 - (0.009 * ($”Value” - 150)). Be sure that you understand why this formula will calculate the temperature at each elevation throughout the region when you click on “apply.” Remember to check the “Apply on OK” box, so that the formula really does apply!

Now find the areas of likely snow cover when the Fresno temperature is 30°C (86°F). To do this, use the hidden row functions to set a search criterion, “Temperature < 0", where ‘temperature’ is whatever you named the column you added in the Raster Attribute editor. Change the color of the areas selected by this criterion to white. Now find the areas in which the temperature is greater than 88°F (31°C). Change the colors of these areas to red. Is it possible for you to develop a legend (or something) that will explain the meaning of the red and white colors? **Print this image for your portfolio.**

What you have done here is to create a new attribute in which each primary pixel attribute (i.e. elevation) is used as the base from which to calculate the temperature. You could, if you chose, add a series of attributes (e.g. Temp10, Temp20, Temp30, etc.) to calculate the temperature of the area where the sea-level temperature is 10°, 20°, 30°, etc. To be sure, the calculation is not precise, since there are differences in weather from the Fresno to the Sierra Nevada and Death Valley area, and the actual lapse rate *does* depend on both weather patterns and the relative humidity of the air, but the basic logic is quite sound.

Note also that you have used the elevation value for each pixel in the image to calculate the corresponding temperature. The lapse rate formula works on a place-by-place basis, and the temperature value of one place does not depend on the elevation of adjacent places.

Overlaying layers from different maps: So far, our operations have acted only on single images. Let's do an analysis in which we need two images. Using the familiar images of Gainesville, GA, determine which areas of deciduous forest have Pacelot soils.

Open two viewers and load LNSOILS.IMG and LNLANDC.IMG. Now Click on Interpreter-GIS Analysis-Overlay from the main Imagine Control Panel. You are prompted for two input images and an output image. Use the two images in your viewers as your inputs, and give the output a suitable name. Click on the button below the LNSOILS.IMG window labeled "Setup Recode." Select the soils which are Pacelot soils (i.e. $"Class_Names" contains Pacelot). Note that the Setup Recode window looks almost identical to the Raster Attribute window. There is one difference, however. There is a window labeled "New Value" and a button labeled "Change Selected Rows." Make sure there is a '1' in the New Value window, and press the Change Selected Rows button. Now invoke the hidden functions popup and choose "Invert Selection." Put '0' into the New Value window and press the Change Selected Rows button. You have just created a raster in which the areas with Pacelot soils have the value of '1' and all others have the value of '0'. This is a Boolean image that answers the query, "Identify all areas with Pacelot soils." Now click on the Setup Recode button below the LNLANDC.IMG window and change the areas of deciduous forests to '1' and the other areas to '0'. Again, this is a Boolean image that answers the query, "Identify all areas of deciduous forest."

Now look at the other things required in this dialog. You can choose the union or intersection of the two images. If you choose the union, your output will include all of the area in *both* of the input images, including areas that do not overlap. If you choose the intersection, your output will include only the area of overlap in the two images. Since your two images are of precisely the same area, the two are equivalent. More important, look at the "Choose Value to Dominate" window. If you choose the "Maximum Value", the output will code all areas in which *either* of the input images had a value of '1' as '1'. If you choose the "Minimum Value", the output will code only areas in which *both* of the input images had a value of '1' as '1'. Since you are interested in the areas that meet both conditions, you will choose the "Minimum Value" option. Click on OK, and look at the result. Compare it to the soils and land use images.

To make the comparison easiest, use the Raster Attribute editor to make an image in which all areas of Pacelot soils and deciduous forests are colored appropriate – and different – colors in their particular images, and the backgrounds are transparent (i.e. they have an opacity of 0). Take the image you have just created and change the opacity of the background to 0. Make sure the color of the areas with both deciduous forest and Pacelot soils is different from that of either color you adopted in LNSOILS.IMG and LNLANDC.IMG. Now open a viewer with LANIER.IMG as the base image, your "adjusted" LNSOILS.IMG and LNLANDC.IMG images in the middle, and your newly created image on top. **Print this image from the viewer and label the thematic areas by hand for your portfolio.**

Operations on Immediate Neighborhoods

Many operations with respect to a specific pixel depend not only on the attribute value of that pixel but also of the attribute values of the 8 adjacent neighbors. You have already run into one instance in which the neighbors were important – when you were doing the clumping operation in Unit 3.

Slope and Aspect: Consider a terrain model such as MTWHITNEY.IMG. The primary attribute for each pixel is the elevation of the point to which it refers. But we can compare the elevation of each pixel with each of its neighbors and calculate both a slope and an aspect (i.e. direction in which the slope faces).

Imagine, for example, that you are doing an ecological study of the central Sierra Nevada, and that you wish to determine which areas are likely to be very dry. In order to do this, you need to identify the areas with a southern exposure, so that they are heated most continuously by the sun, and sufficient slope that what rainwater falls on the surface will run off quickly. Click on Interpreter-Topographic Analysis-Slope and Interpreter-Topographic Analysis-Aspect. Use MTWHITNEY.IMG for the input image in both dialogs. Give the output images suitable (and different) names. Since this only has one layer, and it is a DEM, use the default for this window. When you have generated your images, open two viewers, one containing each of the two images you have just created – again as pseudocolor images. Using the Raster Attribute editor, change the colors so that all slopes greater than 35° (70%) are highlighted, and slopes less than 45° are white and have an opacity of 0. Likewise, change the colors so that all slopes with aspects greater than 140° and less than 220° are highlighted, and those with aspects other than this range are white and have an opacity of 0. Link the viewers, and see which of the steep slopes face south. These are the driest areas of your study area. Given the topography, which is some of the most rugged in the world, is there an easier way to pick out the really dry areas as quickly?

Operations on Extended Neighborhoods

Many operations beginning with a single pixel or a single object act not with regard to the 8 neighboring pixels, but rather over the entire image. The most common of these is those operations involving distance.

Simple Distance from an Object: Load LNHYDRO.IMG into a viewer. Use the Raster Attribute editor to see what this image of the hydrography in the Gainesville, GA area involves. The image contains several sensitive areas. Imagine, for the moment, that you wish to locate a secure landfill in the Gainesville area, and that it is important that it be as far as possible from any hydrologically sensitive area. How far is it possible that you can locate this landfill, and how far from the nearest sensitive area is it?

Look at the ImageInfo information about LNHYDRO.IMG. What are its height and width in pixels? Now click on Interpreter-GIS Analysis-Search. The dialog that opens accesses the distance algorithms in Imagine. Your input image is LNHYDRO.IMG. Give your output image a suitable name. Click on the "Setup Recode" button. Drag the mouse down the sensitive classes (perhaps classes 1-9, but you might choose another set) to select them, and change their value to '1'. Use the hidden functions to change the other classes to '0'. This is significant. Pixels with a class value of '1' are considered "origins" of a distance search; those with a class value of '0' are considered pixels to which a distance must be ascribed. In essence, the search algorithm will calculate the distance of any unallocated pixel (i.e. one of class value '0') to the *nearest* pixel with a class value of '1'. Be sure that the "Distance to Search (in pixels)" window contains a large enough number that every pixel in the image will be searched. Then press OK.

When the process has run, open the newly created image (in pseudocolor). Use Inquire Cursor to check on some of the values. What do the pixel values indicate? Where is the point in the image furthest from a hydrologically sensitive area? To help you find it, you might open the Raster Attribute editor and change the color of the largest distances in the table.

Buffer Zones: Gainesville, GA is an old city. Much of its expansion predated the environmental consciousness that we now take for granted. It would be interesting to see what sorts of land uses were within close proximity to the sensitive hydrologic areas – i.e. Lake Lanier, streams, and the "A" zones. Fairly arbitrarily, let us define 240 m as a buffer zone of interest that we would like to set around these sensitive zones. What land uses are there?

Again click on Interpreter-GIS Analysis-Search. Use LNHYDRO.IMG as your input image, and give your output image a suitable name. Set up your search exactly as you did in the previous instance *except* limit your search to the 240 m buffer zone size. To do this, insert the correct number of pixels in the "Distance to Search (in pixels)" window. What is the correct number? Remember that you are inserting the number of *pixels*, not meters. What is the pixel size for this image? If your buffer zone is 240 m, divide 240 by the pixel size in meters. This gives you the number of pixels required to give 240 meters.

When you run the model, you will get an image in which the images used as the "origins" of the distance model have the value '0' (i.e. they are 0 distance from the reference object). Points within the buffer zone have a value referring to the number of pixels they are from the reference object. Points outside the buffer zone have a value of 1 more than the number of pixels you entered in the "Distance to Search (in pixels)" window. Using the Raster Attribute editor, change the color of the background and highest classes (i.e. the areas *outside* the buffer zone) to white, and set the opacity from all classes *within* the buffer zone to 0. Then load the modified LNLANDC.IMG you created in the first portion of this exercise into the same viewer and use Arrange Layers to put the buffer zone image on top. This shows clearly which land uses lie within the buffer zone. **Print this image using Map Composer for your portfolio.**

Questions to Consider

1. Why was it important to open most of the images in this exercise as pseudocolor images? They came up in gray scale anyhow. Couldn't you just as easily have opened them as grayscale images and not had to worry about pseudocolor?

2. Can you think of a way to use the imaging tools within Imagine to visualize what this area looks like? Do they work? Why or why not?

3. Can you think of a way in which you can create a single image to show the portions of the central Sierra Nevada face south and have more than 70% slope?

4. Can you think of a way in which you can use the statistical tools in Imagine to calculate the area of each land-use class within the buffer zone?

Portfolio

Three images should be printed for your portfolio:

1. The image of the Mount Whitney/Death Valley area showing areas of snow cover and hot temperature.
2. The image showing areas around Gainesville, GA, in which deciduous forests occur on Pacelot soils.
3. The image showing land uses within the buffer zone of 240 meters of streams, Lake Lanier, and other hydrologically vulnerable areas.