SFU SIMON FRASER UNIVERSITY
THINKING OF THE WORLD

Course: Introduction to Remote Sensing
Course code: 253Q
Department: Geography

Term: Fall 2012
Instructor: Anders Knudby, aknudby@sfu.ca
TA: Marco Jorge, mjorge@sfu.ca

# Lab 07

**Goal :** Learn how to do assess the accuracy of classifications in ENVI.

**Terminology:**

**Regions of Interest (ROI):** Regions of Interest is ENVI's term for polygons that you can draw on an image, and use for a variety of purposes. In this lab we will use ROIs as training sites for our classification.

**Classes:** Classes (in this context) are categorical units that individual pixels are assigned to by a classifier, on the basis of the pixels' band values. Classes typically represent surface types, like "water", "forest", "desert" etc.

**Validation:** In remote sensing, the term validation is typically used to mean "assess the quality of". For a classifier, validation is the exercise of determining its accuracy. "Validation points" for that reason, are the individual points that are used to conduct the validation.

**Georeferenced:** This refers to anything that has a set of geographic coordinates attached to it. Validation points for classifications typically consist of digital photos with geographic coordinates, typically obtained from a GPS that recorded where the photo was taken. Modern cameras, including those on smartphones, can often automate the georeferencing process when a GPS and camera is included in the same device.

**"True class"**: This refers to the actual surface type of an area. It is opposed to what could be called the "*classified class*" - what the classifier says the surface type is.

1) ENVI 5.0 seems to me rather useless when it comes to classification, so we'll stick to ENVI Classic for this lab (Open with All Programs – GIS_RS – ENVI 5.0 – 32-bit – ENVI Classic (32-bit)).

2) Start this lab by opening your files from last week. You don't need to open the rules images, just the original Landsat data and the results of your Maximum Likelihood and Minimum Distance classifications.

a. Remember to open the Landsat image using the "File Open External File – Landsat – Fast" command.

b. You can open the others with "File – Open Image File".

**Q1: Classes are typically defined on the basis of one of two things – "Land Cover" or "Land Use". Land Cover refers only to what is physically present on the Earth's surface. Land Use, on the other hand, refers to the use humans (society) make of the land. Sometimes two different types of Land Cover have the same Land Use (for example both "deciduous trees" and "coniferous trees" can have the "production forest" land use). And sometimes two different kinds of Land Use can have the same Land Cover (for example "water" can be used either as "natural water body" or "wastewater pond"). Sometimes the difference between "Land Cover" and "Land Use" is obvious, but sometimes it is subtle. Consider the classes you have used in your classification, and give one example of a Land Cover class, and one example of a Land Use class. If you have only one kind of class, provide a hypothetical example of the other kind. 10 marks**

3) Ok, you are ready to assess the accuracy of your classification. Normally you would have some independent validation data for this, like georeferenced photos for example, to validate the classification accuracy. For the purposes of this exercise we don't have that, so we need to generate some data that can approximate real validation data. We will create such data to assess the accuracy of the Maximum Likelihood Classification first.

a. Use the tool in Classification – Post Classification – Generate Random Sample - Using Ground Truth Image.

i. Select the result of your Maximum Likelihood Classification and click "OK". This tells ENVI that you want to select random points from each class in the classification result.

ii. A new window opens up, showing you the different classes in your classification. All classes should have a ticked box to their left, so you just need to click "OK", because you want to generate a random sample for each class.

iii. Another new window opens up. What you choose here depends on how keen you are to know exactly how accurate your classification is. For the purposes of this lab we will not be super-keen, but we will do enough to understand what's going on.

iv. Under Sampling Type, select "Equalized Random". This option will generate the same number of points for each class. Each point will be randomly located within its proper class (as mapped by the Maximum Likelihood Classifier).

v. Under Sample Size, select "10". If you were doing this for real you should have (much) more than 10, but this is just an exercise to show you how it works.

vi. Under Output To select "Multiple ROIs", and select a filename (such as equalrandom.roi) for the new ROI file that will be created.

4) You have just created 10 random points (pixels) for each class, located somewhere inside your image. What we need to do now is find each of these points, and then rely on our visual interpretation of the Landsat data to tell ENVI what the "true class" is for each pixel. We use visual interpretation of the Landsat image because we don't have any field data on what the actual surface is at these points, but such field data would obviously be preferable.

a. Create a true-colour composite with the Landsat data. In ENVI Classic that's done this way:

i. In the window called "Available Bands List, make sure the radio button is set to "RGB Color", not "Gray Scale".

ii. Then click once (and only once) Landsat band 3 (it will be called "Meta (Band 3) (0.6614)).

iii. Then click once on band 2, and once on band 1, and then click the "Load RGB" button.

b. In the "Image" window of the true-colour composite, select "Overlay – Regions of Interest...". The ROI Tool opens up.

c. If there are any ROIs showing in the ROI Tool, delete them all (this ensures you will get things right if you follow the instructions!)

d. Go to “File – Restore ROIs”, select the .roi file you just created.

e. Delete the “Region #1” ROI that was in the ROI Tool when you opened it. You should now have one ROIs for each class listed in your ROI Tool.

5) What you need to do now is create an additional ROI for each class. For each class, you do that by first clicking New Region in the ROI Tool window. Then change the ROI name by clicking “Region #1” and renaming the ROI to something like “validation xx”, where xx is the name of the class. Don’t forget to press enter to save the new name you input for your ROI. You may have to look back to remember what each colour represents in your original ROI file.

6) You should now have 2 ROIs for each class, where the first ROI has 10 Pixels in it, and the second ROI has 0 pixels in it. Here’s an example of what it would look like with 5 classes:

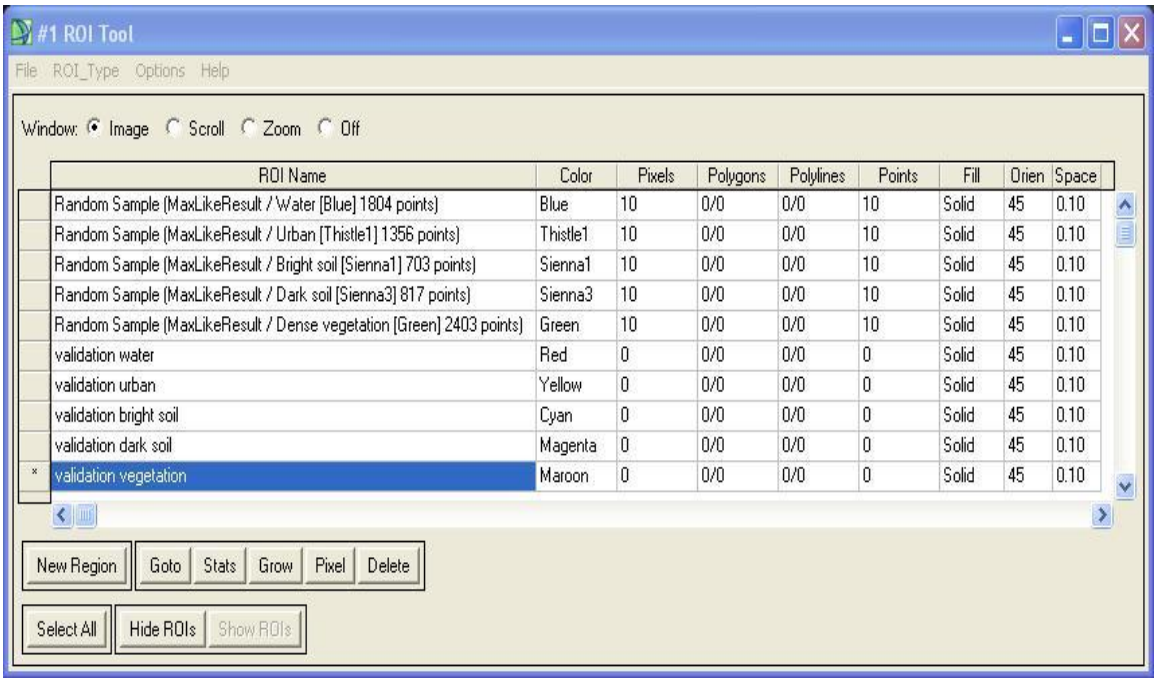


**Q2: Submit your own ROI Tool at this point as a screenshot. 10 marks**

7) Now we'll determine the true class for each of the randomly selected pixels for the first class.

a. **Make sure the ROI Tool's "Window" radio button is set to "Zoom".**

b. **Also, make sure in the menu bar in the ROI Tool window "ROI_Type" is set to "Point".**

c. This is where you can imagine that you have some independent data that tells you what the real surface type is at the location for each point, like a georeferenced photo or notes from someone who has visited the area. Of course you don't have that kind of data for this exercise, so instead you need to use visual interpretation of the Landsat data to determine what the "true class" is for each of the randomly sample pixels. You can use the true-colour display to do that, and you can also create new false-colour displays and link them all to help with the visual interpretation.

d. Select the top ROI (one with 10 pixels in it). The little star on the left indicates which ROI is selected. Then click on the "Goto" button in the ROI Tool. ENVI will navigate to the first point in this ROI.

e. You now need to determine what the surface type of that pixel is.

i. **Do not click anywhere in the Zoom window until you have determined the surface type of the pixel in question!** That would add a pixel to the ROI.

ii. The pixel itself will be coloured according to the ROI's colour, but sometimes you can determine its surface type from the neighbouring pixels (for example, if all the surrounding pixels are water, the pixel in question is probably also water). If you need to, you can click on the "Hide ROIs" button in the ROI tool, to see the pixel itself. To get the view of the ROI back you need to use "Show ROIs".

iii. Use different "Enhance" settings to get different views of the pixel and the area around it. That will help you form a better opinion of what the surface type of the pixel is.

iv. You will probably also benefit from zooming in a bit more in the "Zoom" window, but be careful to only click on the little "+" sign in the "Zoom" window, not anywhere else. If you click anywhere else in the "Zoom" window you will add another point to the ROI (you don't want to do that).

v. You can also open new displays with false-colour composites, and link them to the true-colour composite. Sometimes you can see different things in false-colour composites.

f. When you are pretty confident what the actual surface type of the pixel is, select the appropriate "validation xx" ROI in the ROI Tool (remember, the little star on the left indicates which ROI is selected) and then click once on the coloured pixel in the Zoom window of the true-colour composite.

g. You get two indications that you have successfully added the pixel to the ROI.

i. The pixel itself will be displayed with the colour of the validation xx ROI.

ii. In the ROI Tool, the "Pixels" value for the validation xx ROI will have changed from 0 to 1.

h. Now select the top ROI again, and click "Goto" again. ENVI will navigate to the next pixel in the top ROI, and you can repeat the procedure of figuring out what the true class of that pixel is, selecting the correct "validation xx" ROI, and clicking one in the Zoom window to assign the pixel to that ROI.

i. Repeat this to cover all 10 points in the top ROI. When the "Goto" button takes you back to the first pixel of the ROI, then you know it has cycled through all ten points. You can also check this by adding up the number of pixels in all the "validation xx" ROIs – they should add up to 10.

8) Repeat this procedure for all your "Random Sample" ROIs, so you end up with the same total number of pixels in the "validation xx" ROIs as you had in the "Random Sample" ROIs (10 times your number of classes). You will then have determined the true class for 10 randomly located points in each of class of the classification. This is the basis for assessing the accuracy of the classification. Make sure to save the ROIs (all of them) in a new file.

**Four notes:**

1) The ROI tool is one of ENVI's least user-friendly tools. If you add the wrong pixel to an ROI, or add the right pixel but to the wrong ROI, the only way to delete that pixel from the ROI is to right-click in the Zoom window, and select "Remove All Points From ROI". That's not so bad for the first point, but if you make a mistake on your 10th point for a given class, your only option is to remote all the points from that ROI and start over.

2) A bug in some versions of ENVI was that you sometimes had to zoom in or out in the Zoom window in order to be able to add more pixels to the ROI. I can't guarantee that this has been perfectly worked out in the version you are working with, so you may need to click the little "+" or "-" signs in the Zoom window before you can assign a new pixel to an ROI.

3) Because of this potential bug, and the general user-unfriendliness of this part of ENVI, do yourself a favour and save the ROIs often as you make them. You do this from the ROI Tool by clicking "File – Save ROIs", then "Select All Items", choose filename, and click "OK".

4) Two suggestions to keep track of your updated ROI files: 1) you could just overwrite one file with each update 2) save the file name with the time and date you created it. This will hopefully help you keep track of your ROI files.

**Q3: Submit your own ROI Tool at this point as a screenshot. 10 marks**

9) Now we'll build the "Error Matrix", which ENVI calls "Confusion Matrix".

a. From the main ENVI Menu Bar, select "Classification – Post Classification – Confusion Matrix – Using Ground Truth ROIs".

b. In the new window, select the result of the Maximum Likelihood Classification, and click "OK". A new window pops up.

c. Now you need to pair up the ROIs you just created (the ones called validation xx) with the classes in the Maximum Likelihood Classification. On the top part of the new window, you see two lists. The one on the left has your ROIs, and the one on the right has your classes.

d. Select a class in the right list, and select the corresponding "validation ROI" in the left list. For example, if you had a class called "Water", select "Water" in the right list, and "validation Water" in the left list. **Do not use the "Random Sample..." ROIs.**

e. With the pair selected, click "Add Combination". The combination should now be displayed in the list called "Matched Classes".

f. Repeat this for each of your classes. If you make a mistake, you can delete a combination from the "Matched Classes" list by clicking on it.

g. When you have selected all the correct combinations, click "OK". A new window pops up.

h. In the new window, you specify how you want the "confusion Matrix" to be reported. You can choose to use "Pixels" and/or "Percent" - make sure both are checked. Also make sure the "Report Accuracy Assessment" radio button is set to "Yes". Click "OK".

10) The confusion matrix is now displayed in a new window. The formatting is typically quite bad if you have more than a few classes, but you should still be able to read the important information. Here is how it is structured:

a. The first line tells you which image these results pertain to.

b. The second and third lines tell you the "Overall Accuracy" and the "Kappa Coefficient".

c. Then comes a Confusion Matrix like the one we looked at in the lecture, except that the rows and columns are switched, so that the classification result is displayed in rows and the reference data (the ROIs you just created) are displayed in columns. This is where the formatting is poor, so I've included an example on the next page (notes on the following page):

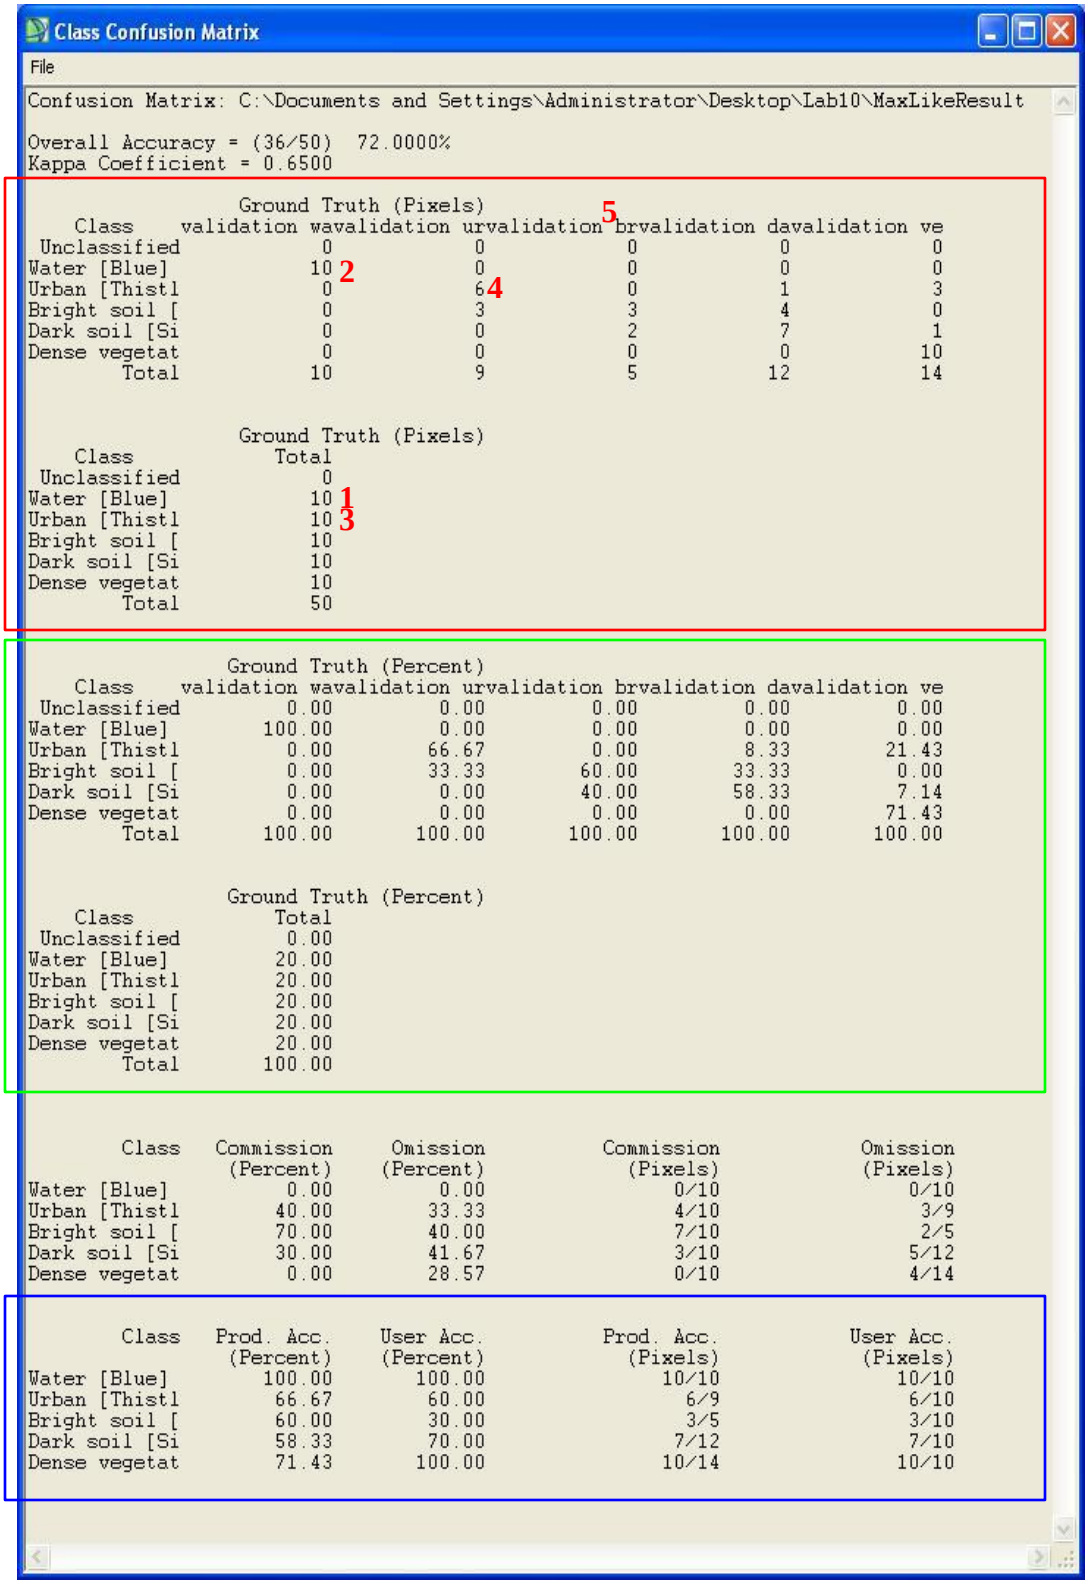
Class Confusion Matrix
File
Confusion Matrix: C:\Documents and Settings\Administrator\Desktop\Lab10\MaxLikeResult
Overall Accuracy = (36/50) 72.0000%
Kappa Coefficient = 0.6500
Ground Truth (Pixels)
Class validation wavalidation urvalidation brvalidation davalidation ve
Unclassified 0 0 0 0 0
Water [Blue] 10 0 0 0 0
Urban [Thistl 0 6 0 1 3
Bright soil [ 0 3 3 4 0
Dark soil [Si 0 0 2 7 1
Dense vegetat 0 0 0 0 10
Total 10 9 5 12 14
Ground Truth (Pixels)
Class Total
Unclassified 0
Water [Blue] 10
Urban [Thistl 10
Bright soil [ 10
Dark soil [Si 10
Dense vegetat 10
Total 50
Ground Truth (Percent)
Class validation wavalidation urvalidation brvalidation davalidation ve
Unclassified 0.00 0.00 0.00 0.00 0.00
Water [Blue] 100.00 0.00 0.00 0.00 0.00
Urban [Thistl 0.00 66.67 0.00 8.33 21.43
Bright soil [ 0.00 33.33 60.00 33.33 0.00
Dark soil [Si 0.00 0.00 40.00 58.33 7.14
Dense vegetat 0.00 0.00 0.00 0.00 71.43
Total 100.00 100.00 100.00 100.00 100.00
Ground Truth (Percent)
Class Total
Unclassified 0.00
Water [Blue] 20.00
Urban [Thistl 20.00
Bright soil [ 20.00
Dark soil [Si 20.00
Dense vegetat 20.00
Total 100.00
Class Commission (Percent) Omission (Percent) Commission (Pixels) Omission (Pixels)
Water [Blue] 0.00 0.00 0/10 0/10
Urban [Thistl 40.00 33.33 4/10 3/9
Bright soil [ 70.00 40.00 7/10 2/5
Dark soil [Si 30.00 41.67 3/10 5/12
Dense vegetat 0.00 28.57 0/10 4/14
Class Prod. Acc. (Percent) User Acc. (Percent) Prod. Acc. (Pixels) User Acc. (Pixels)
Water [Blue] 100.00 100.00 10/10 10/10
Urban [Thistl 66.67 60.00 6/9 6/10
Bright soil [ 60.00 30.00 3/5 3/10
Dark soil [Si 58.33 70.00 7/12 7/10
Dense vegetat 71.43 100.00 10/14 10/10
1
2
3
4
5

11) You can see in the figure that the pixel-based confusion matrix (outlined by the red rectangle) has been split so that the last column has been put on a new line.

a. What you read from this part of the figure is that of the 10 pixels checked for the "Water" class (point **1**), 10 were actually water (point **2**). However, of the 10 pixels checked for the "Urban" class (point **3**), only 6 were actually urban (point **4**).

b. Note that ROI names are truncated in the column headings (e.g. point **5** is next to the heading for "bright soil"). But they are in the same order (left to right) as in the columns (top to bottom).

c. You can interpret the remainder of the pixel-based confusion matrix based on the lecture notes, **keeping in mind that in the lecture notes the classification results are shown in columns, whereas in ENVI the reference data are shown in columns.**

12) Below the pixel-based confusion matrix is another confusion matrix (outlined by the green rectangle in the figure), where the number of pixels in each columns have been converted into the percentage of the total number of pixels in that column. If you have a very different number of validation points for each class, this is typically a preferred way to read the confusion matrix. We had 10 validation points for each class, so in our case the two version of the confusion matrix look quite similar.

13) At the bottom (outlined by the blue rectangle in the figure) you have the producer and user accuracies listed, first as percentages and then based on the actual numbers of pixels.

**Q4: Submit a screenshot of the Class Confusion Matrix window, similar to the example in the figure above. 20 marks**

**Q5: Which class(es) in your classification have the highest combined user and producer accuracies (add the two values to get a combined score)? If two classes have equal accuracies list both. Explain in detail why you think this/these classes have such high accuracies. You may want to have a look at some of your answers from last week as you answer this question. 20 marks.**

**Q6: Which class(es) in your classification have the lowest combined user and producer accuracies? If two classes have equal accuracies list both. Explain why you think this/these classes have such low accuracies. Again, you may want to have a look at some of your answers in Lab 9 as you answer this question. 20 marks.**

14) Although we won't do it in this lab, a quick analysis of the confusion matrix can often be used to improve the classification. For example, taking the class with lowest combined user and producer accuracies (e.g. "Bright Soil" in the example above), you can find out which other classes it is typically confused with. You can then go and refine the training areas for each of these classes to get better separability, or you can go even further and redefine the classes, maybe adding some new ones or merging others.

**Q7: In the confusion matrix, find the row that tells you what the actual surface type was for the 10 validation pixels selected in the class that had the lowest combined user and producer accuracy. For example, in the example above you can see that pixels classified as "Bright Soil" were actually either "Urban" (3 pixels), "Bright Soil" (3 pixels), or "Dark Soil" (4 pixels). Provide the same values from your own confusion matrix. 10 marks.**

**Note:** The process you have gone through in this lab is much faster if you already have your validation points ready in a suitable format, either as an ROI file or some other format that can be converted to ROI easily, for example ENVI can also import shapefiles from ArcGIS and use them instead of the ROIs we worked with in this lab. ArcGIS can actually also do both classification and accuracy assessment, so if you have validation points in a format ArcGIS can read that might be faster (and more user-friendly). However, the intention with this lab is that you learn not only how to do the accuracy assessment, but also learn how it works; the slow (and painful) progress of looking at each point helps do that.