Exercise #5D

Image Classification

(Decision Rules and Classification)

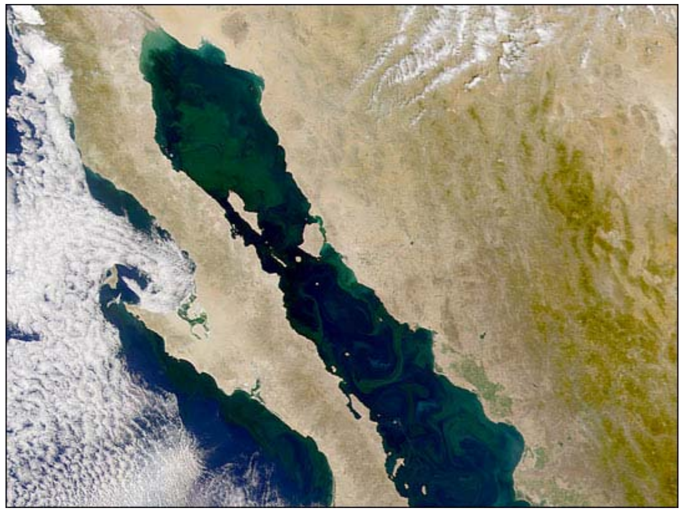

Objective

- **Choose how pixels will be allocated to classes**
- **Learn how to evaluate the classification**

Once signatures have been defined and saved in a signature file, whether they were defined from the image you are trying to classify, or from another image, you can use them to perform a classification. Each pixel in your image will then be placed in a class, placed in more than one class (overlap) or placed in no classes (unclassified). In order to decide which class to place any particular pixel in, we need to establish a set of decision rules.

If we use parametric rules (based on statistics) then all pixels are placed in classes. None are left unclassified. None are placed in more than one class. The parametric rules place the pixels in classes by either — finding the signature which is closest spectrally to the spectrum of each pixel (minimum distance rule); finding the signature that the pixel "most likely" belongs to (maximum likelihood rule); or finding the signature which minimizes the Mahalanobis distance (includes an evaluation of the covariances of the signatures). So all pixels are assigned to a single class.

1. Classify your image

Open your image to be classified (<yourfile.img>) in Viewer1. Open the Signature Editor with the signature file to be used in the classification. In the menu bar choose Classify -> Supervised. Enter a name for the output file (classified image), like <yourfile_class1> (for a maximum likelihood classification using the signatures you defined from the AOIs). Make sure the folder is the right one (i.e. yours). Select the distance file and give it a similar name (<yourfile_dist>), and the same folder. Choose <None> for the non-parametric decision rule. Choose maximum likelihood for the parametric rule. Press OK. The image will be classified. (But it may not be displayed – you might have to open it in another Viewer.)

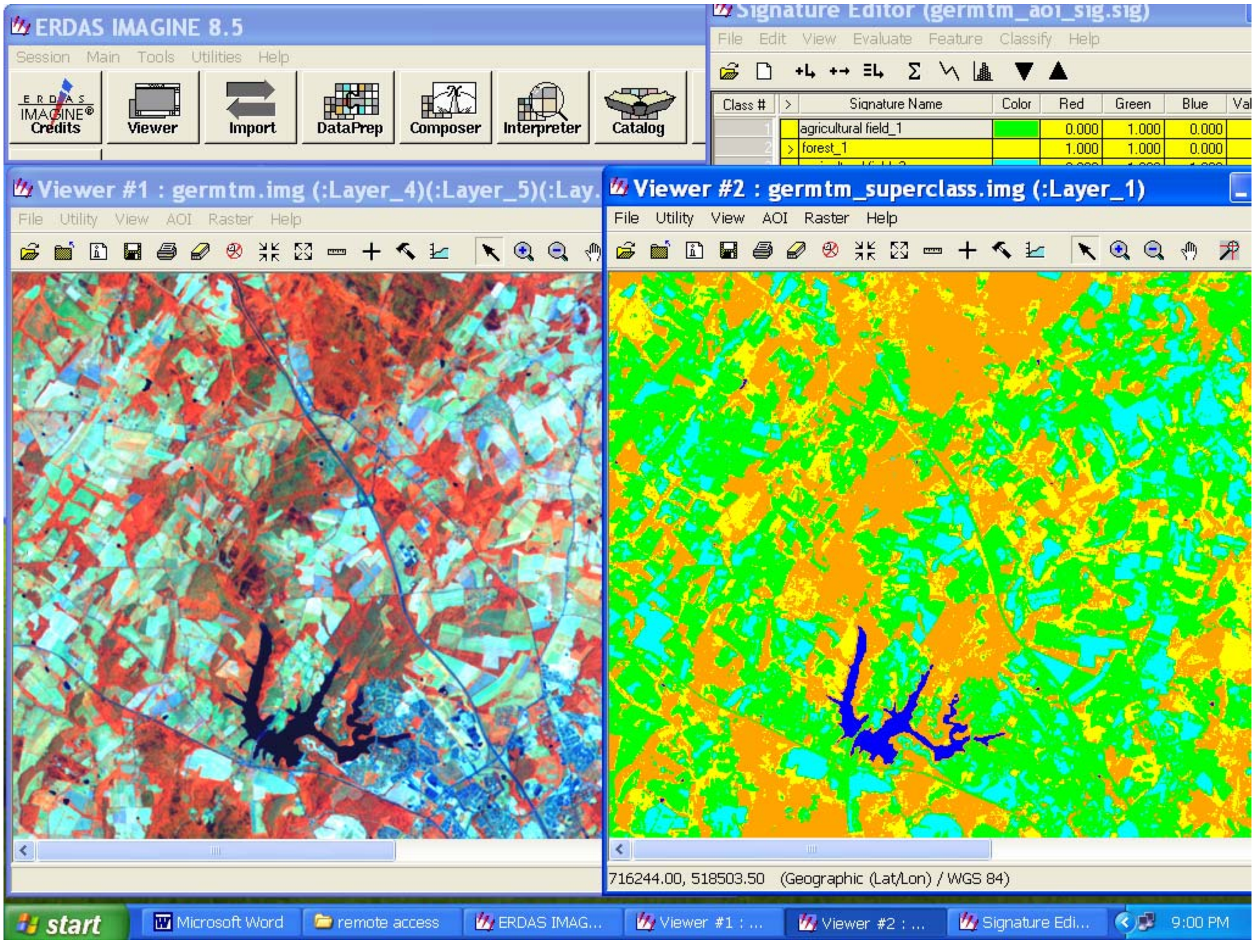

The original image is on the left. The classified one on the right.
Now classify your image using a different rule (e.g. minimum distance), and save it under another name (e.g. <yourfile_class2.img>).

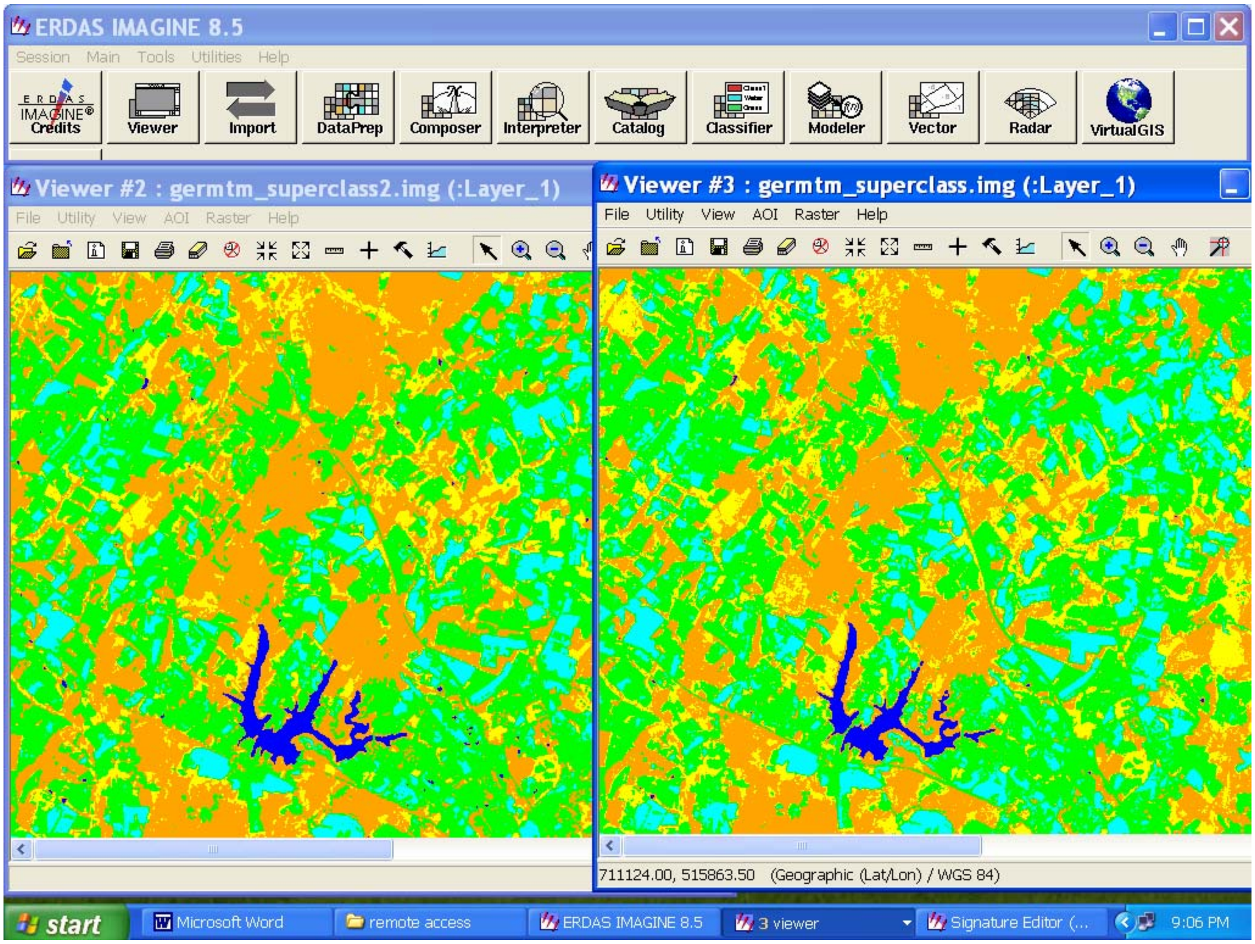

A comparison of these two classifications (minimum distance is on the left, maximum likelihood on the right) will show that in general they are very similar. There are some differences though.

1 Explain carefully why there are differences even when the original image is the same and the signatures used to classify the image are the same?

2 How would you know which was best?

Note: If you use a non-parametric rule to classify your image, some pixels may be left unclassified and some may fall into more than one class. That is because the decision rule is not based on a statistics like “minimum” or “maximum”, but the pixel would actually have to fit into the definition of one and only one signature class. If, for example, the ellipses of two signatures overlap in feature space and a pixel falls into this overlapping region it could be assigned to either class. This is where an overlap rule can be applied to decided which class to place the pixel in. Likewise, since the signatures are unlikely to cover every possible spectral combination there may be pixels in the image which fall into the gaps (in feature space) between the ellipses of the signatures. These pixels would then be unclassified unless we chose a parametric rule to classify them.

2. Evaluate your classification

There are many other ways to evaluate the accuracy of the classification. One is by viewing and analyzing the distance file that was created during the classification. The pixels in the distance file have a one-to-one correspondence with those in the original image. The brightness of the pixels in the distance file is indicative of their spectral distance from the signature to which they were assigned. Bright pixels are spectrally far from the class, and are more likely to be misclassified. Dark pixels are spectrally close to their assigned class. View the distance file for one of your classifications and use it to identify pixels which may be misclassified. (Note: You can't directly compare the distance values generated by a minimum distance classification with those generated by some other algorithm, like maximum likelihood, because the spectral distances are defined by different mathematical formulae.)

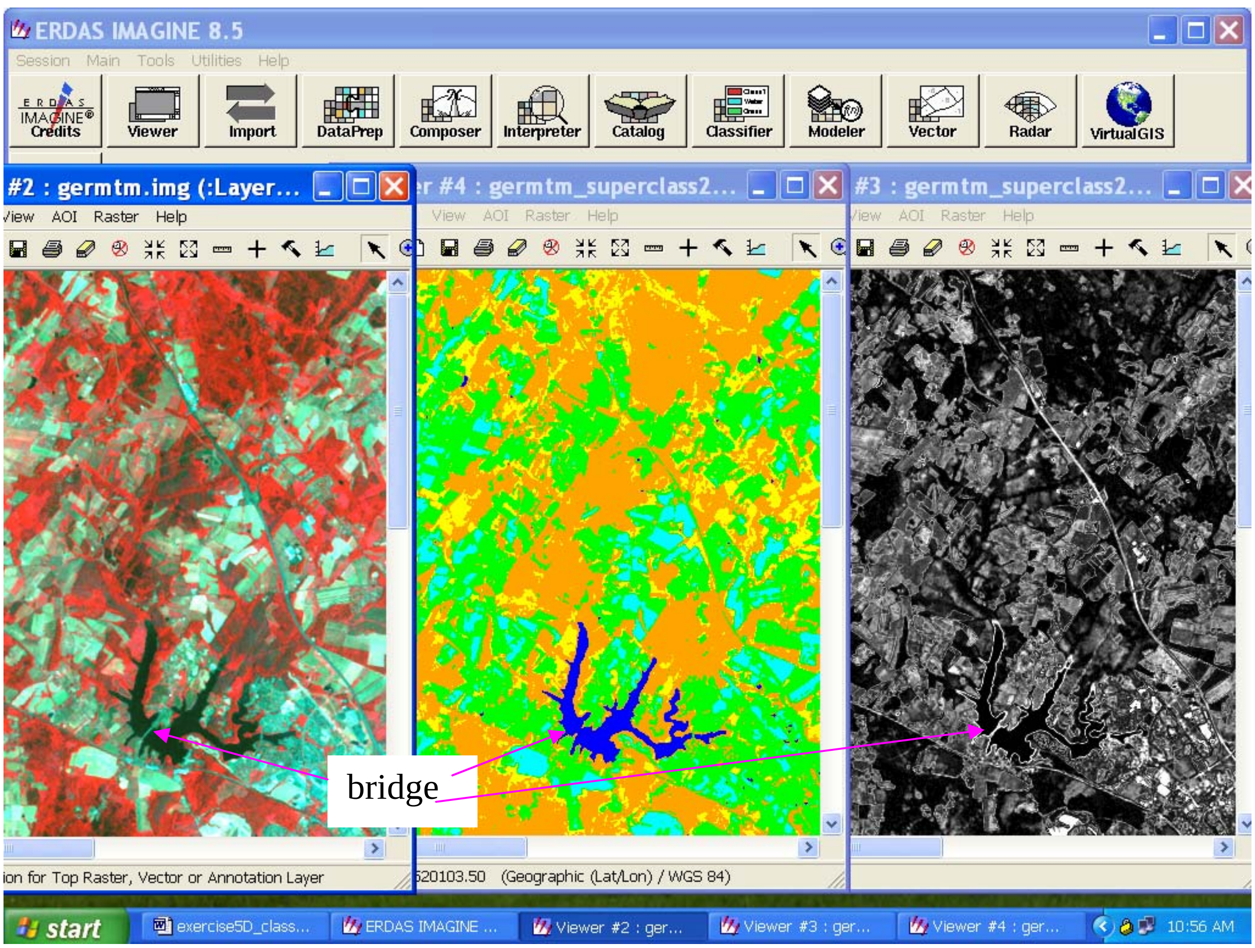

In the images above the left panel shows the unclassified image, the middle panel shows a classification done using minimum distance, and the right panel shows the distance file for this classification. You can see that the computer had some trouble distinguishing the bridge from the water (compare 1st and 2nd panels), and this shows up brightly in the distance file.

3. Identify some bright areas in your distance file. Do they correspond to areas which were classified differently by the two different decision rules that you used?

These are areas where the analyzer may choose to do "ground truthing" or refine the signature definitions.

3. Thresholding

In order to avoid misclassification in an image, sometimes analysts choose to define a threshold based on the distance value. This means that pixels whose distance value exceed this limit will be left unclassified. This avoids the possibility of biasing statistics with bad data. The easiest way to choose a threshold value is to look at the histogram of brightness values in the distance image.

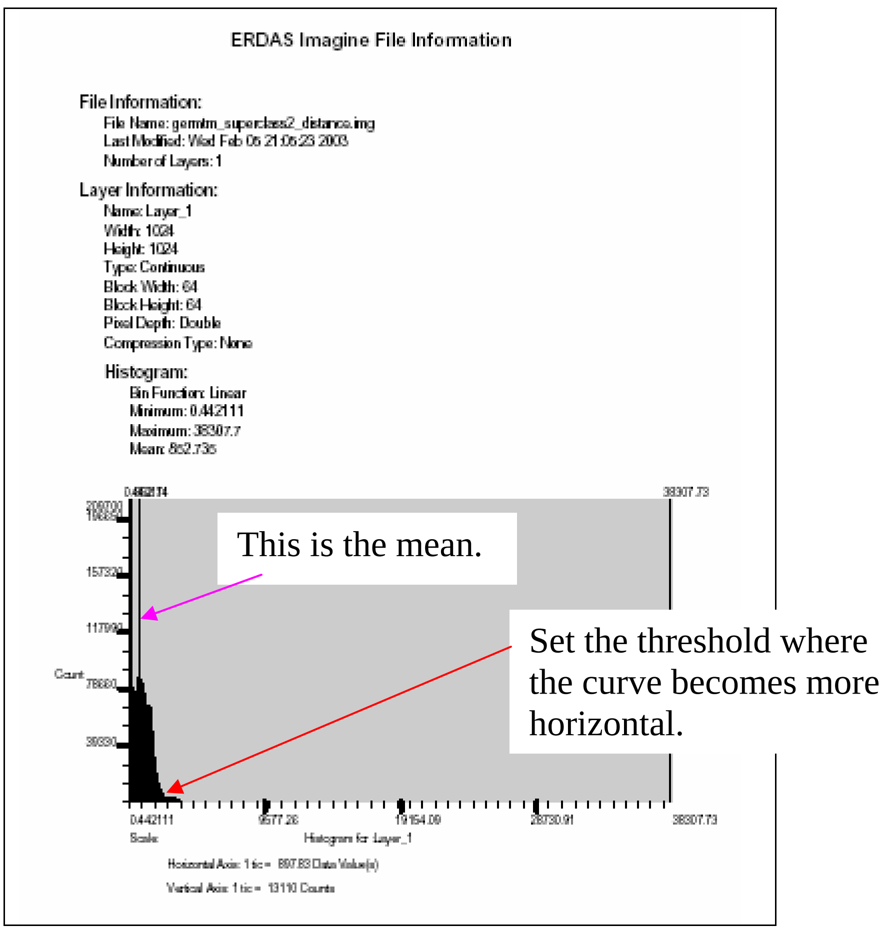

The report above was generated from the Histogram -> Print option. View the histogram of data values for your distance file. It should be a smooth chi-square shape like the one shown here.

If you want to experiment with thresholding, the Threshold option can be found under the Classifier icon in the Imagine header bar.

Last Modified: 7 January 2009
Helen Cox