An Urban Classification of Khartoum, Sudan

Project Summary and Goal: The primary goal of this project was to delineate the urban extent of Khartoum, Sudan from a Landsat ETM+ image captured in 2006. In the process five cover classes were created: urban, desert, agriculture, fallow, and water. Both unsupervised and supervised classifications were performed; in regards to the unsupervised classification different parameters and algorithms were examined and compared for their efficacy in establishing the urban class cover. The secondary goal of this project was to become familiar with the ENVI 5.0 interface and the tools required to complete a classification workflow using this software.

Processes Used: Discussion and Interpretation

The individual bands 1,2,3,4,5, and 7 were downloaded from the Global Land Cover Facility website and opened in ENVI 5.0. (Thermal bands 61 and 62 were discarded for the purposes of this project.) Several composite layers were output using the layer stacking tool: a 742 band image and a 432 band image were used to assess the terrain and visually identify the class cover types. A variety of stretch types were easily utilized in the ENVI software and applied to these false color images to create a crisper image of ground data and to enhance the identification of the various cover types. Additionally, an NDVI image was also created to assist in distinguishing agricultural areas from fallow fields. Given the geography of Khartoum it was expected that nearly all vegetation would either be natural growth along the Nile or agriculture.

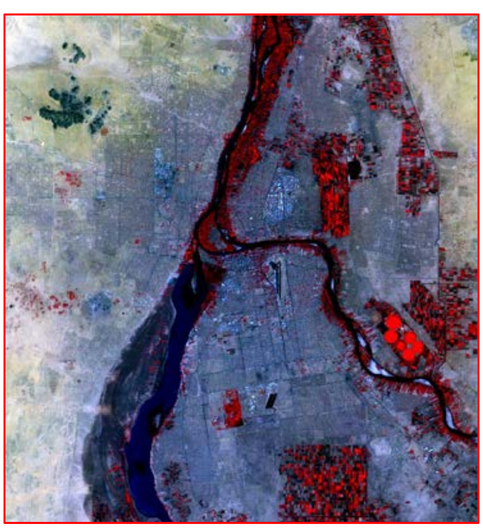

Bands: 432 Linear 5% Stretch

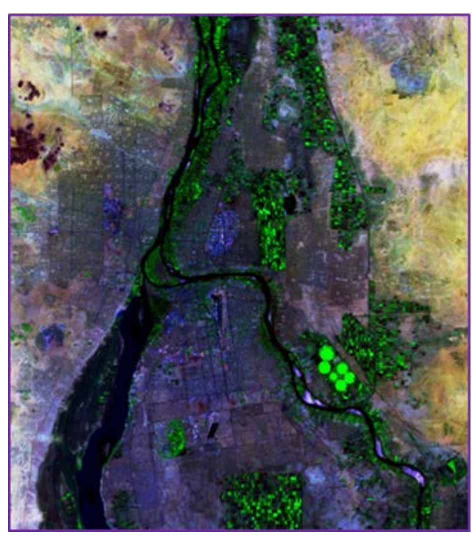

Bands 742: Equalization Stretch

In preparation of the classification process a composite image of all the bands was created. First an unsupervised classification was performed and assessed. Cover types were then paired with the class number through visual analysis of the output. Using the unsupervised results alongside the NDVI and the false color images, training samples were drawn. A total of 33 training samples were drawn: 9 for Agriculture, 9 for Desert, 5 for Fallow, 7 for Urban, and 3 for water. These shapefiles were then converted into an roi. file and used to perform a supervised classification.

Two supervised classification algorithms were used: maximum likelihood and Mahalanobis. A comparison of the results indicated that both methods resulted in confusion: the only difference being the particular classes which were confused. Below is an output comparison of the urban class as captured by these two methods in ENVI.

Maximum Likelihood

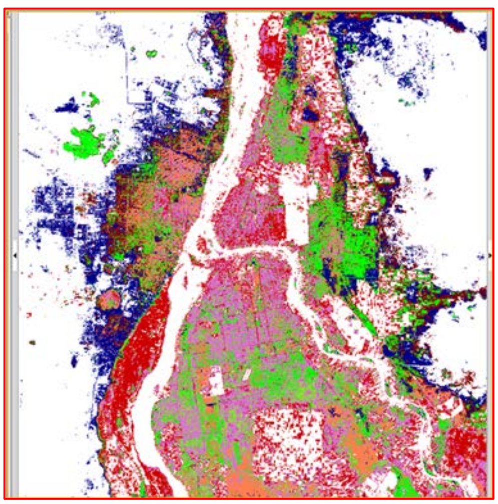

Urban Class 7 is depicted in red. It carries the western White Nile Flood plain very noticeably along with much of Khartoum Bahri (Northern Khartoum).

Mahalanobis

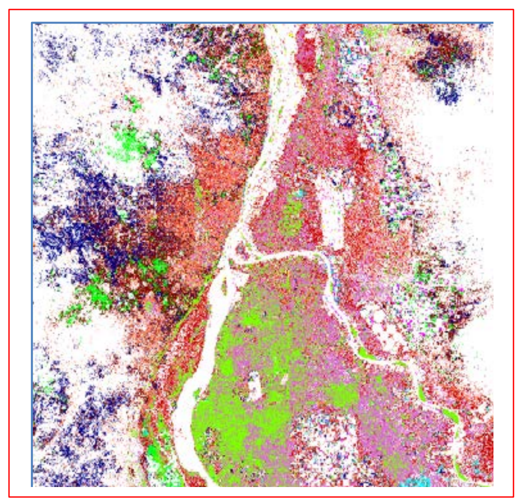

Urban Class 3 is depicted in blue. It carries some of western Omdurman along with much of the western desert.

With the maximum likelihood classification, Urban Class 7 was identified with agricultural and fallow lands, most noticeably on the western flood plain of the White Nile. The Mahalanobis classification output indicated confusion of Urban Class 3 with the surrounding desert. There was no apparent advantage of one classification method or the other, it being merely a case of trading one ambiguity for another. A decision was made to scrap the Mahalanobis classification and to proceed with the maximum likelihood classification, but to impose a probability parameter to increase accuracy. A threshold parameter of .9 or 90% was entered in the input and the maximum classification was performed again. The results were significant only for the urban class and desert class largely due to the fact that confused areas were unclassified in the output as a result of thresholding. The unclassified area included pixels from every major cover type.

Output of Urban Classes using Maximum Likelihood with a probability parameter of 90%.

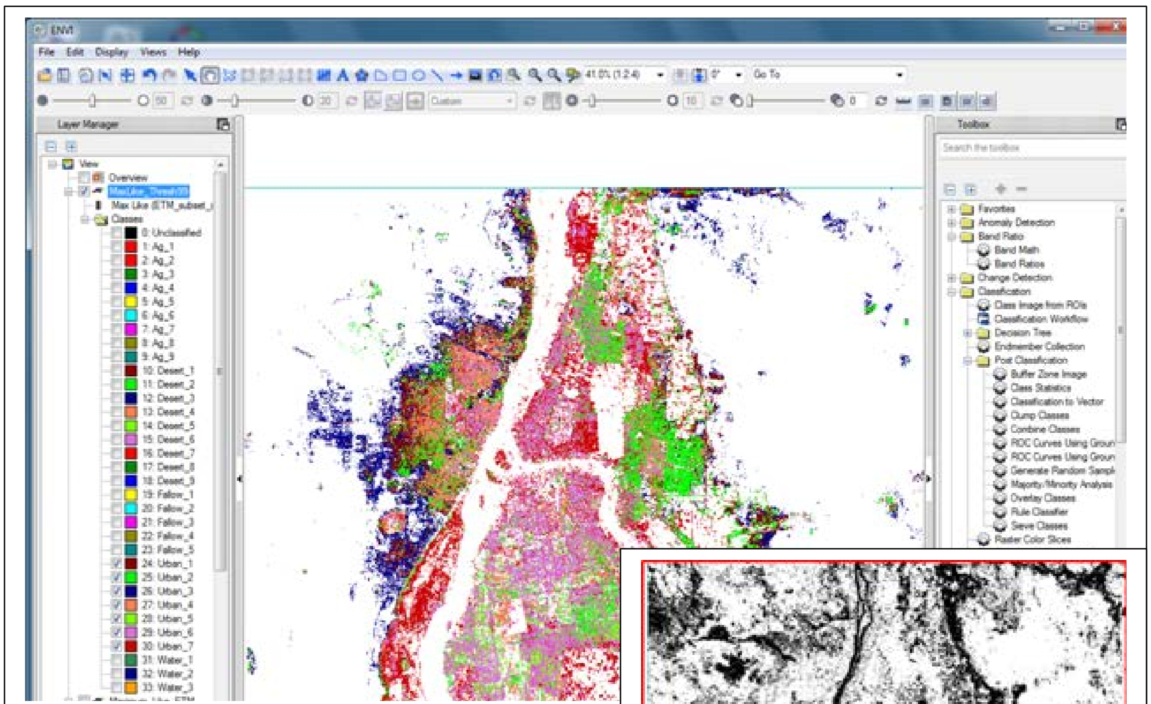

The unclassified output

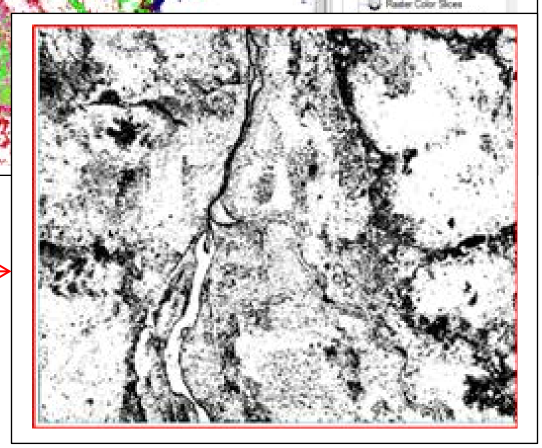

The result indicated that urban class confusion persisted along the western White Nile flood plain. Since thresholding alone did not prove sufficient in clarifying these confused classes, a secondary approach of subsetting out the problem class and further refining training samples was then employed. Urban Class 7, which had captured a large swath of fallow and agricultural land along the western White Nile in the first maximum likelihood classification, was exported as a vector file. Fifteen more training samples were drawn within its boundary, while paying close attention to the spectral profile of each sample, ensuring their integrity and uniqueness.

The Spectral Profile of Urban Class 9

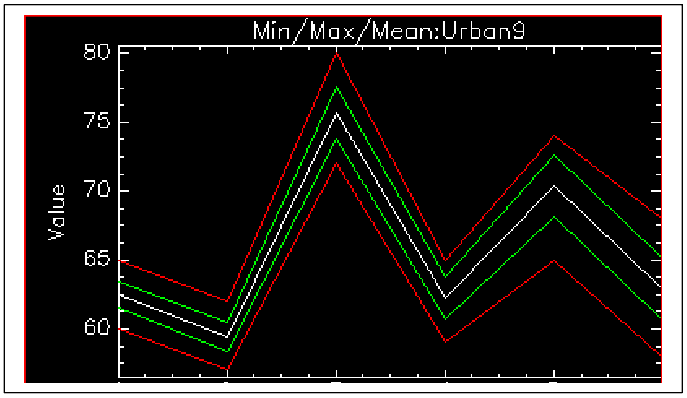

Using the Class 7 vector file for the input mask parameter along with thresholding, a supervised classification was run so that only areas within the confused class were reclassified: the goal being to test the efficacy of these training samples on the problem area alone.

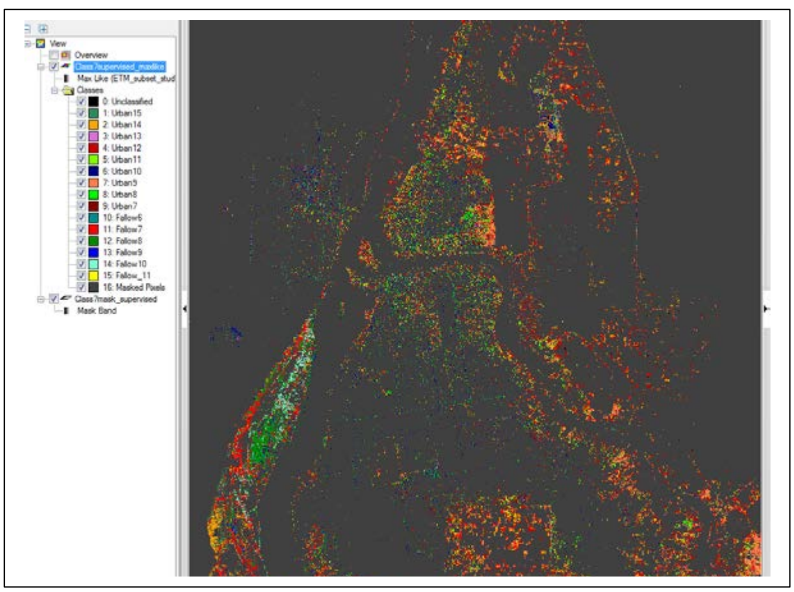

Second set of training samples used to run a supervised classification on the Urban Class 7 vector file.

After examining the output it was determined that these new training samples were quite effective in capturing the fallow lands along the White Nile. Additionally, upon inspection one class was changed from urban to agriculture.

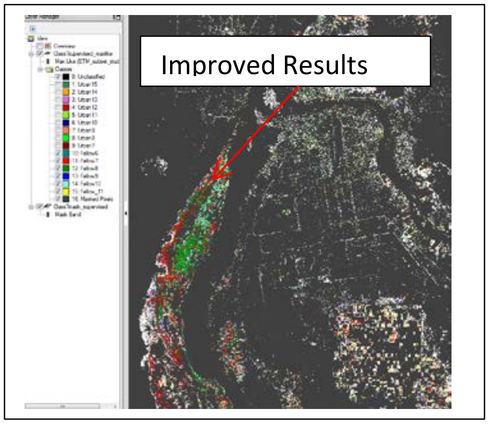

Fallow Classes

Finally, these training samples were merged with the first training samples (the original Class 7 removed and replaced), and converted into a roi. file consisting of 48 subclasses. A maximum likelihood classification with thresholding at 90% was run on the entire image. There were still a significant number of pixels from the surrounding desert that remained unclassified. This was a consequence of not having drawn more training samples of the desert. The final output results were then merged into class types using the Rule Classifier tool. It appears that the unclassified class morphed into the urban class when using the image classifier so that the urban boundary of Khartoum appears to extend farther into the desert than it actually does.

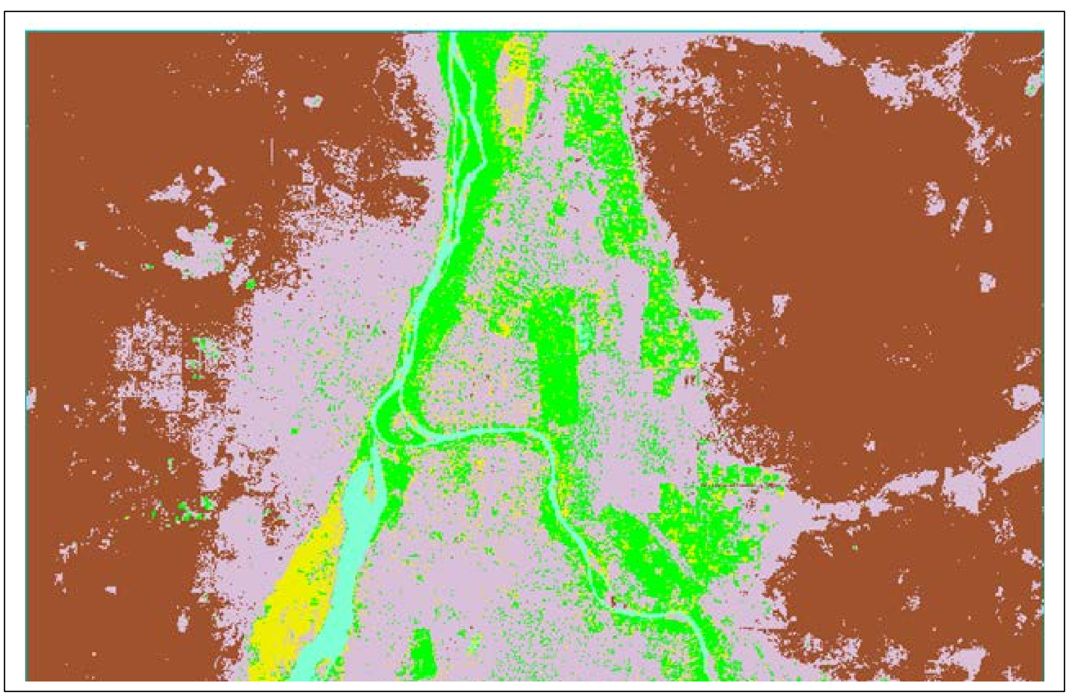

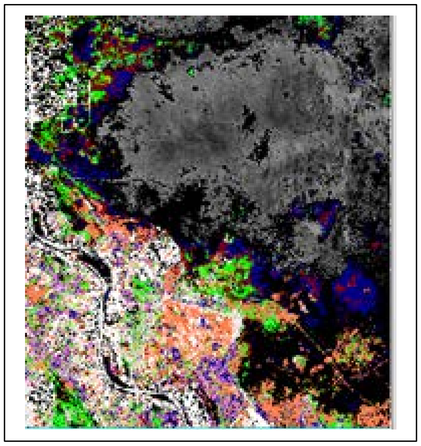 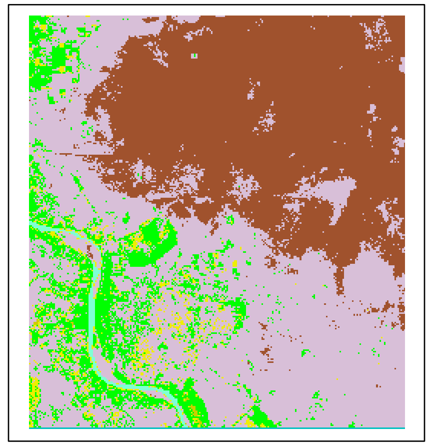

The rule classifier tool converted the unclassified desert area, shown in black on the left, to urban, shown as violet in the final output image on the right.

Outcomes: Conclusions and Inferences

My ability to draw more effective training samples and thereby increase the accuracy of the final output classification was hampered by both time constraints and geographical distance. A further set of training samples taken of the surrounding desert on the flanks of Khartoum would have provided a more precise urban/desert boundary for my final classification image. Additionally, acquiring ground truth would be necessary to resolve any class confusion which could not be resolved by spectral signatures alone. I suspect that the use of mud and clay in building many homes within Khartoum accounted for the confusion of urban classes with fallow cover. Additionally, I didn't really consider what the spectral profile of a rural area might be. Perhaps some of the fallow and agricultural lands contained rural dwellings, almost certainly constructed of mud and clay. Such hypotheses could only be tested by acquiring ground truth.

The most difficult, unresolved challenge of this project was posed by the fact that many fallow areas shared a nearly identical spectral profile with some urban areas. Given this fact no amount of polygon drawing would help to distinguish the two. This also presents a major challenge to the efficacy of an automated, spectral-based approach, which could be scaled up and used on a larger area of Sudan. Perhaps the only way to resolve such confusion would be by implementing a decision tree or some type of "localized" algorithm that is not solely based on spectral signatures alone.

What did you learn from conducting this project?

I learned how spectral signatures respond to a vastly different geography in ways that can be expected – e.g. band reflectance for water – to ways that were unexpected – band reflectance for the urban areas of Khartoum fell within a min-max range 60 – 70 (DNs for all bands), resulting in minimal covariance. Most of Khartoum appeared drab and dark whether looking at individual bands or false color image as a result of this low reflectivity. This provided an interesting contrast to individual band images and composite images of urban areas in North America which were previewed in textbook and used in our lab coursework.

Spectral classification is a laborious and tedious process which often presents a limited number of methods to resolve conflicts and the confusion of classes. After completing this project I think it unlikely that the best training samples would be able to capture class distinctions on the ground with 90% accuracy or greater given that the spectral differences between classes is often too subtle to be rendered. Hyper-spectral analysis may prove a fruitful alternative, but not a viable one, since such data is not readily available in the public domain.