

INCREASING THE DETAIL OF LAND USE CLASSIFICATION: THE IOWA 2002 LAND COVER PRODUCT

R. Peter Kollasch, Remote Sensing Analyst
Iowa Geological Survey
Iowa Department of Natural Resources
109 Trowbridge Hall
Iowa City, IA 52242
pkollasch@igsb.uiowa.edu

ABSTRACT

When producing the Iowa Land Cover 2002 product from multitemporal Landsat imagery, we tried a number of experimental approaches, which proved highly successful in producing a more detailed, and perhaps more accurate, land use classification product. The six factors which made the difference were 1) using path-oriented imagery, rather than georeferenced, to reduce the destructive effects of resampling; 2) utilizing 30 meter imagery at a resolution of 15 meters, which improves the positional accuracy of each pixel, and allows the intersection with the second image to form a higher accuracy end product; 3) taking great care to achieve sub-pixel accuracy in the scene to scene georeferencing, which allows for more accurate stacking of pixels; 4) using the “Class Grouping Tool” suite of tools in ERDAS Imagine to give the interpreter more power to manually refine the labelling of an unsupervised classification; 5) utilizing the “Fuzzy Recode” tool to resolve some of the spectral class confusion, where applicable, and 6) Designing the class structure to fit what can be accomplished from the imagery.

Three of these six elements are aimed at preserving as much of the spatial and spectral content as possible, while utilizing multitemporal imagery. Two of the other three involved the use of specialized tools for interpreting unsupervised classifications. Together these six elements allowed us to increase our number of land use classes from 8 to 17, and to produce a product which, at 15 meter resolution, appears more detailed than the 30 meter imagery from which it was derived.

INTRODUCTION

The Department of Natural Resources (DNR) of the State of Iowa has an ongoing program of developing current and historical land-use analyses in support of decision-making processes. In 2004 we were faced with the choice of trying to resolve some issues with the most recent previous (year 2000) land cover product, or of creating an entirely new product. The availability of nearly cloud-free multitemporal Landsat imagery over the entire state in 2002 compelled the decision to create an entirely new land use product for the year 2002.

Starting on a new project also enabled us to try a number of experimental approaches, which were motivated by a variety of factors. One important factor was that, in 2002, we had acquired statewide 1 meter resolution color infrared (CIR) digital orthophotography. However, attempts to perform image classification on this high-resolution CIR imagery had proved difficult, even using advanced tools such as eCognition, and this difficulty was attributed to the lack of spectral content of the imagery, as compared with the Landsat imagery that had been used in previous land-use classifications. Conversely, comparing the Landsat imagery with this high-resolution imagery had highlighted how crude the spatial resolution of Landsat imagery appeared by comparison, and motivated us to think about ways to mediate differences between these two important imagery sources, so that they might be used together more effectively. This effort is ongoing.

A variety of techniques were considered, with the goal of preserving as much of the internal pixel geometry of the imagery as possible, in order to construct a more detailed land-use classification. Ultimately, six separate approaches distinguished this effort from previous efforts done at the Iowa DNR and other projects in the author’s experience. The six factors were:

1. Developing the classification from path-oriented imagery, rather than North-up georeferenced imagery, to reduce the destructive effects of resampling;
2. Utilizing 30 meter imagery at a resolution of 15 meters, which reduces the absolute positional error of each pixel, enables the use of the pan band in the classification, and allows the intersection with a second 30 meter image to form a higher accuracy end product;
3. Taking extreme care to achieve sub-pixel positional accuracy in the scene to scene georeferencing, which allows for more accurate stacking of pixels;
4. Using the “Class Grouping Tool” suite of tools in ERDAS Imagine to give the interpreter more power to interactively refine the labeling of an unsupervised classification;
5. Utilizing the “Fuzzy Recode” tool to resolve some of the spectral class confusion, where applicable;
6. Designing the class structure to fit what can be accomplished from the imagery, rather than starting out with an abstract list of target classes.

The first three of these six elements are aimed at preserving as much of the spatial and spectral content as possible, while utilizing multitemporal imagery. This paper will focus on these first three elements. The remaining three are innovative approaches in working with unsupervised classifications, which involved the use of specialized tools for interactively labeling unsupervised classes, and for selectively resolving spectral confusion using each pixel’s neighborhood, and information derived from the labeling process. The following sections describe the rationale for these approaches, the processes used, and the results. Persons interested in any of these approaches which are not covered in sufficient detail here are encouraged to contact the author.

DESTRUCTIVE EFFECTS OF NEAREST NEIGHBOR RESAMPLING

The Landsat and SPOT satellites, and other major moderate resolution earth-imaging satellites, generally follow sun-synchronous orbits. These orbital tracks are designed so that, at each orbital pass, which average about 90 minutes apart, the satellite passes over an area of the earth at approximately the same local (sun) time, usually about 10:30 AM. Thus if the satellite passed over Nova Scotia at 10:30 AM local time, one orbit later it might pass over Iowa at 10:30 AM local time. The next orbital pass might take it over Nevada at 10:30 AM local time, etc. To accomplish this type of orbital timing, the orbit is at an angle to the meridian, in the case of Landsat and SPOT, travelling from Northeast to Southwest at approximately a 13 degree angle at the latitude of Iowa.

When imagery collected in this manner are georeferenced to North up, pixel resampling is used to reorient the Northeast Southwest pixel track to a North South orientation. Generally, when pixel resampling is performed, at least three modes of resampling are available, namely *Nearest Neighbor*, *Bilinear Interpolation*, and *Bicubic Spline*.

Both *Bilinear Interpolation* and *Bicubic Spline* derive the new value for a target pixel from at least four neighboring input pixels, resulting in some degree of spectral mixing. These two approaches sacrifice spectral content to preserve geometric relationships. Preserving spectral content is important in producing accurate image classifications. We have chosen to use *Nearest Neighbor* resampling because it has the advantage that each target pixel takes its value from just one source pixel, thus preserving whatever spectral content exists in the “raw” image. However, it does this at the expense of compromising the geometric relationship between neighboring pixels in serious ways. The following sections describes some of these effects.

When nearest neighbor resampling is done, at some scale, panels of pixels are preserved exactly as they occurred in the original unresampled image. The size of these panels is determined by the angle of rotation between the original (path-oriented) image and the target orientation (typically North up). If the angle of rotation is very small, these panels will be large, and as the angle of rotation gets larger, the size of these unchanged panels will shrink. At the latitude of Iowa, the size of these panels is approximately 5x5 pixels. Because the panels are irregularly shaped, some pixels are dropped by the resampling process, while others are duplicated. (Figure 1) Pixels near the center of one of these panels will be located very near to their “true” location, while pixels at the edge of a panel may be located as much as one-half of the pixel size distant from their true location (Figure 2). Between these panels of preserved pixel relationships, there are “shear” lines, separating each panel from its neighbor panels. Pixels on opposite sides of a shear line are certain to be displaced from their true location in a direction nearly opposite that of the pixels on the other side of the shear line, creating an offset of one full pixel across the line.

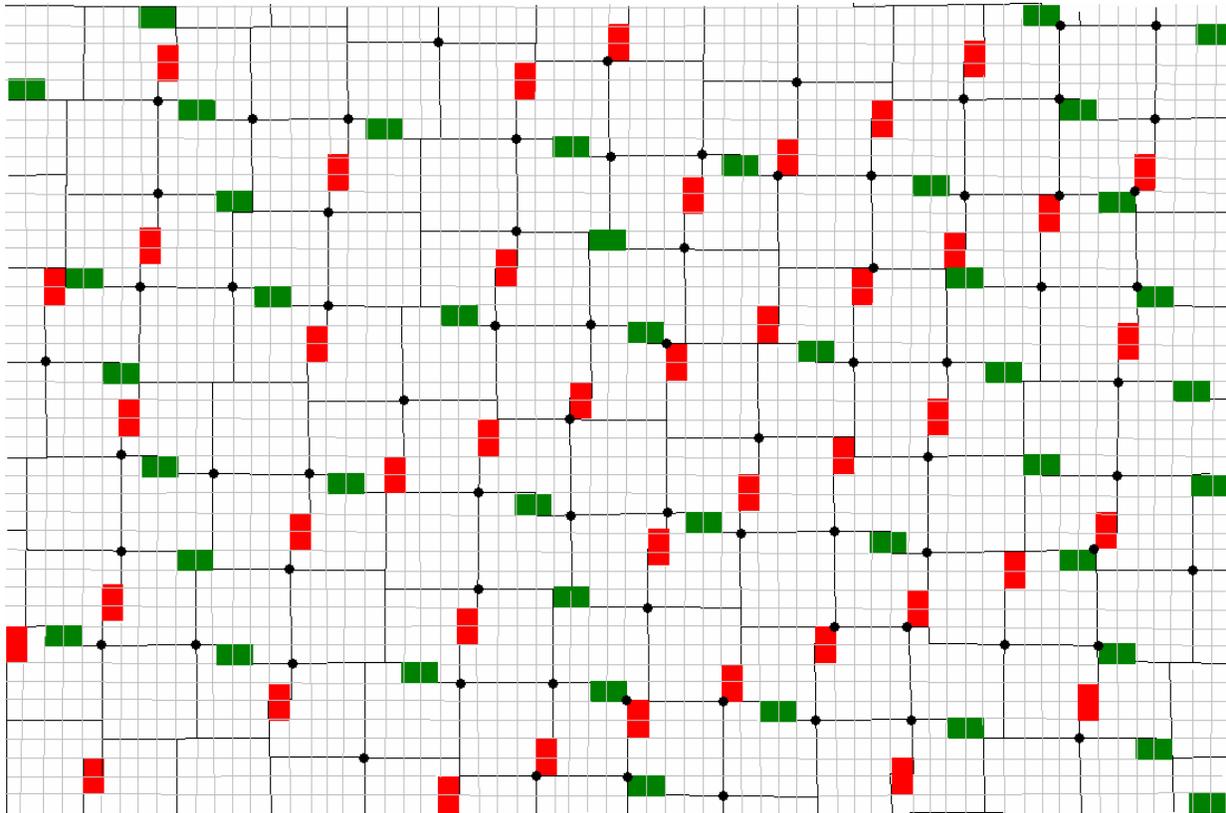


Figure 1. Analysis of actual resampling results in the target space. Panels of unchanged pixels are bounded by black “shear” lines. Red blocks show areas where a pixel was duplicated by vertical shear, and green blocks by horizontal shear. Black dots on intersections identify locations where pixels were dropped in the resampling process.

The most adverse effect of this occurs when working with multitemporal data. The way that imagery is normally combined to produce a multitemporal classification, is to produce a layer stack, where multiple images are brought into the same geographic reference system, then the layers of each image are merged, in a process called “stacking”, into a single image. The resulting layer stack is then used in much the same way as a single image would be to develop a classification.

Most users, when stacking multitemporal images, will use the reference system that will be utilized in the final map product, which in the case of the Iowa Land Use Map is UTM, Zone 15, using the NAD83 datum. This north-up reprojection involves a rotation of approximately 13 degrees, which results in an average panel size of 5x5 pixels, as described above and shown in Figures 1 and 2. With shear lines spaced at roughly a 4 to 5 pixel interval, more than half of all pixels are at the edge of a panel, where rotational displacement is maximized. The size of this displacement averages nearly one half pixel near these edges, with the direction of displacement distributed around the circle (Figure 2).

When layer stacking two multitemporal images that have been georeferenced to a north-up orientation, both images will have been independently sheared into panels in the manner described. When the images are stacked together into a single image, there is no way to control how the two independent sets of panels will intersect with each other. Since over half of the pixels in each panel have an average displacement around one half pixel, and since these displacements are oriented in all directions, many pixels will be stacked with others whose displacement is large enough in a reverse direction to create a net full pixel offset between their “real” locations. It is estimated that this will occur in more than 10 percent of all pixels, perhaps in a much higher percentage.

The preceding discussion shows that North-up georeferencing of multitemporal imagery can and will result, a significant percentage of the time, in stacking pixels from one image with pixels from the other image which would correspond to neighboring pixels in the original image. When the neighboring pixel represents a different land-use class, this mis-stacking of pixels can clearly affect the spectral quality of the pixel stack to induce a

misclassification. It is the suspicion of this author that this effect is at least partially responsible for the frequent occurrence of nearly random misclassified pixels in many multitemporal land-use classifications.

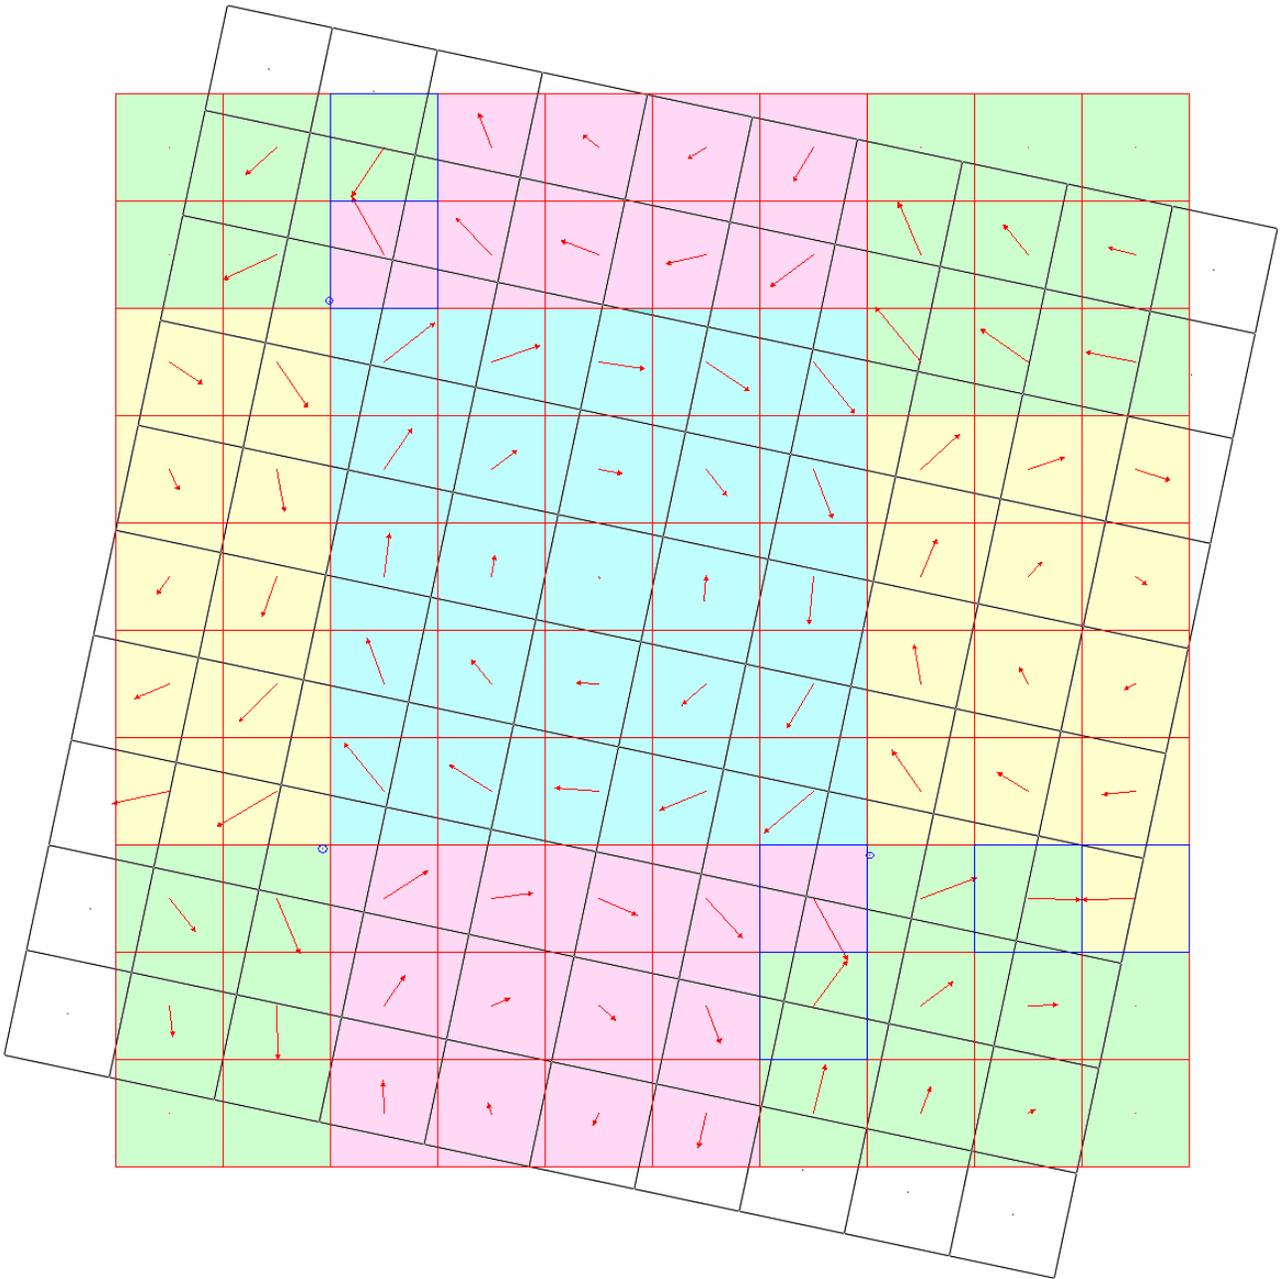


Figure 2. Illustrates the rotational displacement within panels of unchanged pixels, represented by the pastel regions. Arrows point from the center of the resampled pixel (red outlines) to their “real” location in the unresampled path-oriented image (black outlines), showing the rotational displacement. Note that across the “shear” lines separating panels, displacements are in opposite directions. Small blue circles identify pixels that were omitted from the resampled output. Duplicated pixels are outlined in blue.

RESAMPLING SOLUTIONS

It is clear from the preceding discussion that the destructive resampling effects described are being caused by the need to rotate the imagery from path-oriented to North-up. In previous work experiences I had encountered users who, in order to preserve spectral content, insisted on using only path-oriented imagery that had been resampled by nearest neighbor techniques. This idea influenced our decision to do the classification using only path-oriented imagery.

All of our previous land-use analysis projects had found that using multitemporal imagery produced a much better classification than any single image date could. We normally used one spring date and one summer date to improve class separability. Working with more than one date of imagery in path-oriented format introduced some new issues. What reference system would be used to bring the two images together for layer stacking, and at what resampling cost?

Resampling of at least one scene is essential to bring the two scenes together. We chose to minimize the effects of resampling by georeferencing one scene to match the path-oriented delivery format of the other. This scene-to-scene georeference would be from one path-oriented scene to another, of the same WRS path and row. Because both scenes are from the same nominal path, the angle of rotation between them is very small, nearly negligible. This means that the panels within which pixel relationships are preserved will be quite large, and the shear lines themselves will be widely spaced.

CHOICE OF RESOLUTION

We were fortunate that for nearly every Landsat image pair in our study, one member of the pair was a Landsat 7 scene, which therefore had a 15 meter panchromatic band available. We chose to use these 15 meter pan bands as the common reference for the layer stack operation. Using the 15 meter pan band as the basis for layer stacking led to the question, could this higher resolution pan band be used to improve the spectral resolution of the classification? We made the decision to try it, but only partially because of the contribution the pan data would make to the classification. We were also intrigued by the possibility that intersecting two 30 meter scenes at 15 meter resolution might have the effect, if carefully executed, of creating a higher resolution classification than the original data itself. This was possible only in the context of preserving as much of the internal pixel geometry as possible, which was accomplished by the use of path-oriented imagery, as described.

In order for the scenes to be intersected in a way that creates a higher resolution result, extreme care must be taken to assure the accuracy of the control points. Sub-pixel spatial accuracy must be achieved on nearly every control point collected. Experience has shown that it is possible to collect control points on 30 meter scenes with a precision that appears to be in the range of 5 to 15 meters. The approach used was to collect control points as accurately as possible, then resample the imagery, and, displaying the resampled imagery with the target image, use the Blend or Swipe viewer tools to compare the co-registration of the two images. With the Blend tool, and with extreme care, even a very small shift in the image registration can be detected. This is checked by viewing and blending both the target and the resampled images together, at full resolution, in every region of the scene. If even a slight shift is detected, the location of the problem is recorded, and the control points in that region of the scene are revisited and refined. This is repeated until acceptable sub-pixel co-registration has been accomplished. It is not always possible to achieve the nearly perfect results required to derive 15 meter results from 30 meter data with sub pixel precision, but it is essential that one come as close as possible.

Most of our control points are rural road intersections, which are very common in Iowa. The key to achieving sub pixel spatial accuracy lies in developing an accurate sense of where within a pixel the precise center of a road intersection lies, by studying the pattern of the surrounding pixels. Control points, then, are sometimes taken in the center of the pixel, sometimes on an edge, and sometimes in between. Every time a control point is collected on either image, this precise sense of location comes into play.

PROCESSING THE DATA

The choices made in producing this land use classification made many of the processing steps more difficult to perform than otherwise. The description of the georeferencing precision required, as described in the last paragraph, should by itself make that clear. We found that working with path-oriented imagery made it impossible to take

advantage of many of the features provided in the software. The lack of a projection-based reference system was the biggest handicap. The datasets were also larger and more cumbersome, at 15 meter resolution.

The Landsat data were purchased in multiscene format, with strips of three individual scenes in a single file, where possible. One multiscene file in each WRS path, usually a Landsat 7 scene, was chosen to use as the “reference” scene for layer stacking. The 30 meter bands were resolution-doubled, and stacked with the 15 meter pan band, if available. This scene became the target scene for control point collection for the scene-to-scene georeference. Other scenes from the same path were georeferenced with extreme care to match this reference scene, and were resampled to 15 meter resolution to overlay on this scene.

The resolution doubling step seemed simple, but was not. No tool is readily available to combine the 30 meter multispectral data with its 15 meter pan band. Although the two share a common projection, they are at different resolutions, which must be brought together. To accomplish this we carefully tweaked the image headers, and created a model in ERDAS Imagine Modeler that performed an implicit resampling that doubled the resolution. The 30 meter bands are then stacked with the 15 meter band, maintaining proper co-registration.

Both dates were reviewed for clouds, and a cloud mask for each was developed. These cloud masks were combined, and used to mask out areas of clouds, as well as areas of non overlap at the edges, from both scenes, so that all remaining areas would have useable data in every band. Principal Components Analysis was run on the masked individual scenes to reduce data redundancy, and the best three or four bands from each scene were selected to be layer stacked together. These masked, clipped, stacked principal components became the input to an ISODATA unsupervised classification, using 240 classes, which became the primary classified file. A second 1000 class unsupervised classification was also developed, but was used only when it was needed to help distinguish coniferous forest from deciduous.

The classifications were interactively reviewed using the Class Grouping Tool suite of ERDAS Imagine, and the classes were labeled. This tool offers a variety of ancillary tools, notably the Dendrogram tool, which proved very useful for producing initial class groupings, that were further refined interactively. The classification scheme was developed as we went, and some classes, notably the various grassland classes, such as Grazed Grassland, Ungrazed Grassland, and CRP (planted grassland), were introduced relatively late in the process, when we determined experimentally that they could be meaningfully separated.

The Fuzzy Recode tool was utilized to resolve some spectral confusion. This is a post classification procedure which can reduce the pixilated character of the classification. Fuzzy Recode utilizes the pixel neighborhood along with settings made in the Class Grouping Tool, to selectively resolve some cases of class confusion. Issues with its use prevented using it for classes with narrow spatial extents. However, it proved very useful for resolving spectral confusion between spectrally confused classes with broader extents, especially the Corn and Soybean classes.

Next, a manual edit was performed to resolve certain cases of class confusion which could not be done automatically. Two general cases to note are the Barren class, which generally represented quarries. These areas had been lumped by the classifier with the Commercial/Industrial urban class. Similarly, some bottomland forest classes were combined by the classifier with the Coniferous Forest classes. These were resolved by manual editing of the classification.

Because multiple files were involved, the classifications were iteratively reviewed and refined to minimize edge effects. The separate pieces were georeferenced to UTM Zone 15, and mosaicked together, using cutlines to maximize the area unaffected by clouds. Major areas still affected by clouds were filled in with classifications developed from a single image date. The full state file was clipped to a simplified boundary that provides at least a 10 mile buffer around the state for modeling purposes. This file can be downloaded from the Iowa DNR NRGIS Library at <http://www.igsb.uiowa.edu/nrgislibx/> under Statewide Data/Land Description/LC2002. Alternatively, it may be viewed online at http://csbweb.igsb.uiowa.edu/imgate/maps/watershed_atlas.asp.

THE CLASS GROUPING TOOL SUITE

Use of the Class Grouping Tool suite of tools was already described in the Processing section. However, its importance as a factor in producing the Iowa 2002 Land Cover product warrants further discussion. Far more time was spent using this tool to perform the class labeling process for this classification than was spent on any other step. The Class Grouping Tool is a suite of tools, available as part of the ERDAS Imagine 8.7 distribution, which facilitates labeling of classes in an unsupervised classification. This tool drives a standard Imagine viewer, and automatically performs many functions, such as highlighting, which were tedious when labeling classes with the

Raster Attribute Editor. One could write extensively about the many, varied ways of using this tool, but instead, a few important techniques, which made a significant difference, will be highlighted.

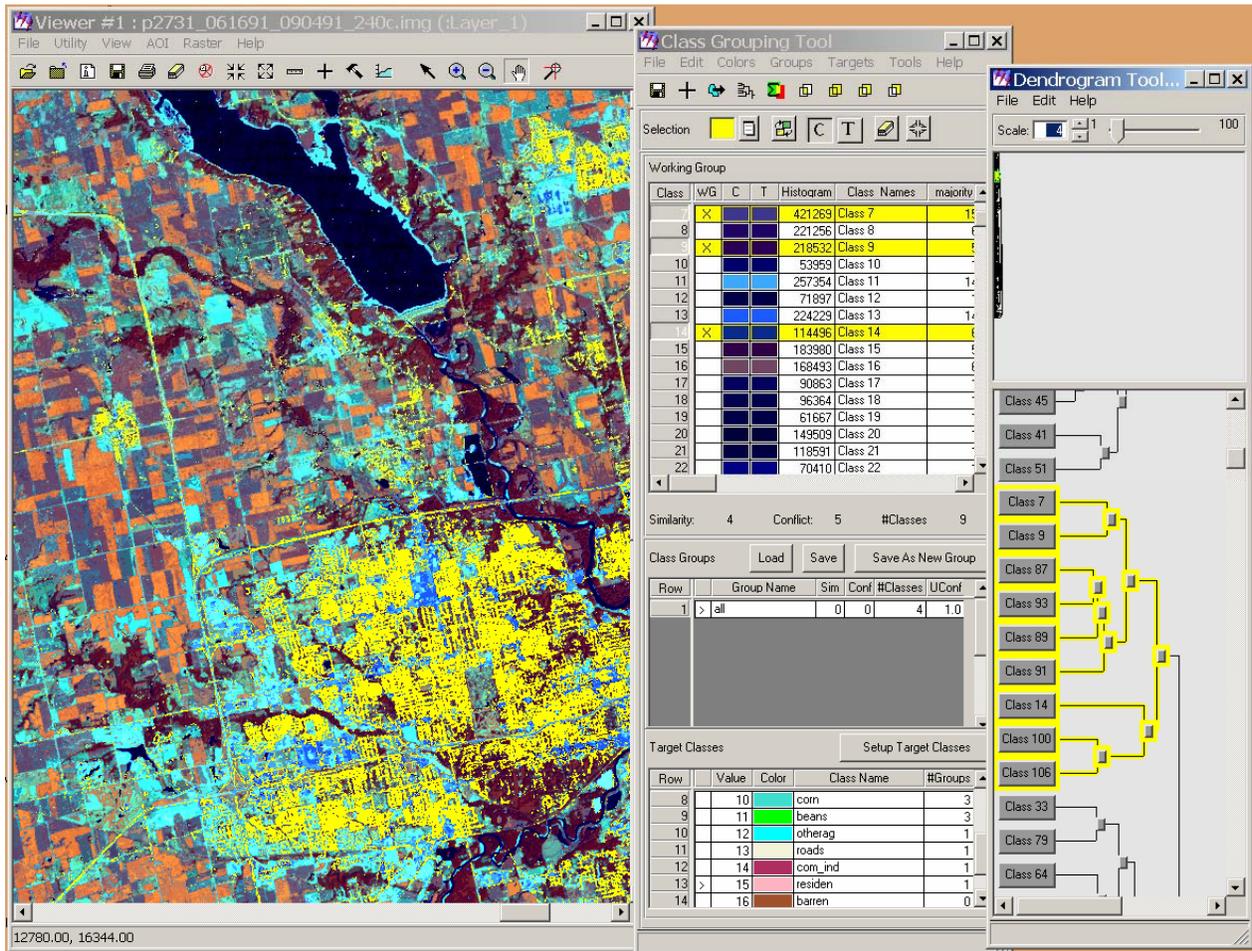


Figure 3. Illustrates the use of the Dendrogram Tool, with the Class Grouping Tool, to create a group of spectrally similar classes which represent urban residential areas. By clicking on any node in the Dendrogram Tool, all subordinate classes in the similarity hierarchy are brought into the current working group, and automatically highlighted in yellow in the viewer. In other words, once set up, the display above was produced by a single click on a hierarchy node.

One ancillary tool that is part of the Class Grouping Tool suite is the Dendrogram Tool. This tool is fully integrated with the Class Grouping Tool to create groups of classes (or clusters) based on their spectral proximity. The signature file that was built during the ISODATA process, and from which the classification was created, becomes the source of spectral information about each class. The Dendrogram Tool creates a hierarchical representation of the spectral similarity of the classes, in which the two classes most spectrally similar to each other are joined by a node. Then the next two most similar classes are joined by a node, and so on, until all classes are joined into a single hierarchy. Special *Agglomeration Methods* tell it how to define similarity to subtrees already joined by nodes. There is also a choice of *Distance Measures* available. In the author's experience, choosing the *Distance Measure* "Mean Scaled" and the *Agglomeration Method* "Average Linkage" seems to give the best behavior. What this means is that groups defined by a subtree in the hierarchy were more likely to highlight contiguous areas of similar character in the display. Figure 3 illustrates the use of the Dendrogram to highlight urban residential areas in the viewer.

Often when working with unsupervised classifications, a small number of classes cover most of the area of one final class, while a large number of spectrally similar classes fill in the holes. Traditional approaches to labeling unsupervised classifications leave the interpreter struggling to find all of the minor classes that complete a single

final class. With the Dendrogram Tool, one can start by finding all of the individual classes which may compose a final class. Then one can use other features of the Class Grouping Tool to refine the grouping by testing individual classes to see if they indeed belong in the larger group, and removing them if they do not. This process was used extensively in producing the 2002 Land Cover product, and helped to accomplish more solid fill-in of many areas.

Traditional approaches to labeling unsupervised classifications often allow each class to be assigned to only one final class. The Class Grouping Tool allows individual classes to become a member of any number of groups, belonging to different end classes, which adds flexibility to the process of grouping. Some important features of the tool are designed to aid the interpreter to find and resolve what it calls “conflict”, a situation where a class is assigned to more than one final class. In cases where true spectral “confusion” exists in the image, it may be appropriate to assign classes that are confused to more than one end class. This kind of situation may be partially resolved by the Fuzzy Recode tool, another member of the suite. The ideal, however, in a clean classification, is that each class will be assigned to one, and only one, end class.

When any group is *Loaded*, and becomes the current working group, conflict with other groups or end classes are automatically highlighted. The interpreter immediately knows whether more work needs to be done if all conflict is to be eliminated, and the highlights tell him where the conflict lies. The toolbar of the Class Grouping Tool has four Boolean tools, which are useful for resolving conflict. The *Intersect (And)* tool can quickly reveal which classes the working group has in common with another conflicting group. These classes can be reviewed, and a decision made which end class the conflicting classes should stay in. The *Exclusive Or* tool, followed by a *Save* operation, provides the quickest way of removing these classes from the end class they don’t belong in. Sometimes, when you have confidence that all of the classes in a group belong to that end class alone, the *Subtract* tool can be used to subtract that group from any other group with conflict, which will remove the conflict. The *Union (Or)* tool can be used to merge two groups. A good understanding of Boolean logic is useful when using these tools. Together, these Boolean tools provide the interpreter with great power to build groups of classes, compare them with each other, resolve conflicts between them, and assign them to end classes, with constant visual feedback provided by the automatic highlight capabilities of the Class Grouping Tool. These tools were an important asset in building the 2002 Land Cover product, and definitely contributed to making it a better product.

RESULTS

Together, use of these procedures has allowed us to increase our number of land use classes from 8 to 17, and to produce a product which, at 15 meter resolution, appears more detailed than the 30 meter imagery from which it was derived. No previous Iowa classification has discriminated more than a single grassland class, while this classification has four grass classes (including Alfalfa/Hay). Anecdotal evidence from a few users has attested to the utility of the CRP (planted grass) class. Likewise, there are three forest classes, where the earlier product had only one, and three urban classes versus one.

We understand well that increasing the class resolution has the potential of diluting class accuracy. However, it was our judgment that utility to the user, and not numerical accuracy, was more important. Comparison of the 2002 Land Cover (Figure 4) with the 2000 Land Cover (Figure 5) quickly reveals the higher class resolution of the 2002 product. It is easy to see that the new classification better distinguishes urban cover from rural, and gives a strong sense of seeing more detail. Because no formal accuracy assessment has been performed, no objective statement can be made about the accuracy of the product. However, we feel that the experimental approaches used in creating this product have been highly successful in producing a classification with higher spatial and class resolution, that better serves the needs of natural resource managers in the Iowa DNR, and of the people of the State of Iowa.

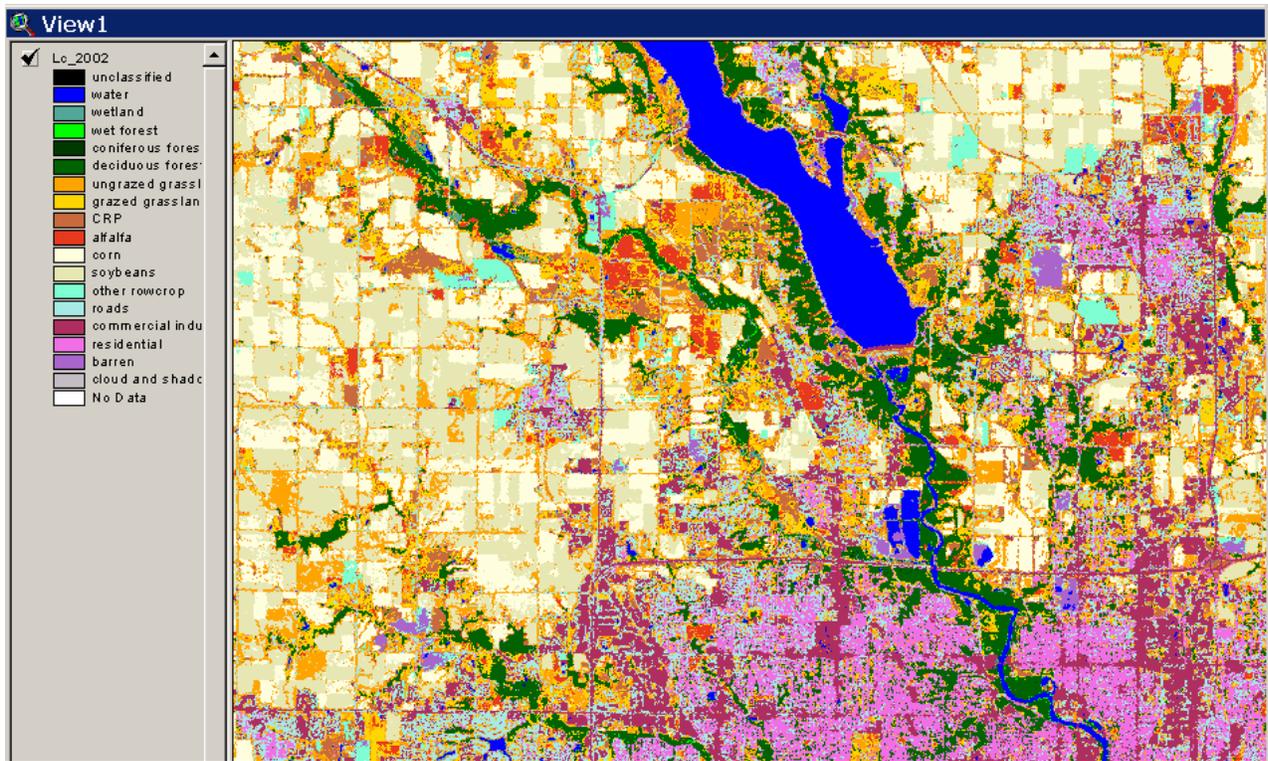


Figure 4. 2002 Land Cover and Legend. Saylorville Reservoir is visible north, and the city of Des Moines south.

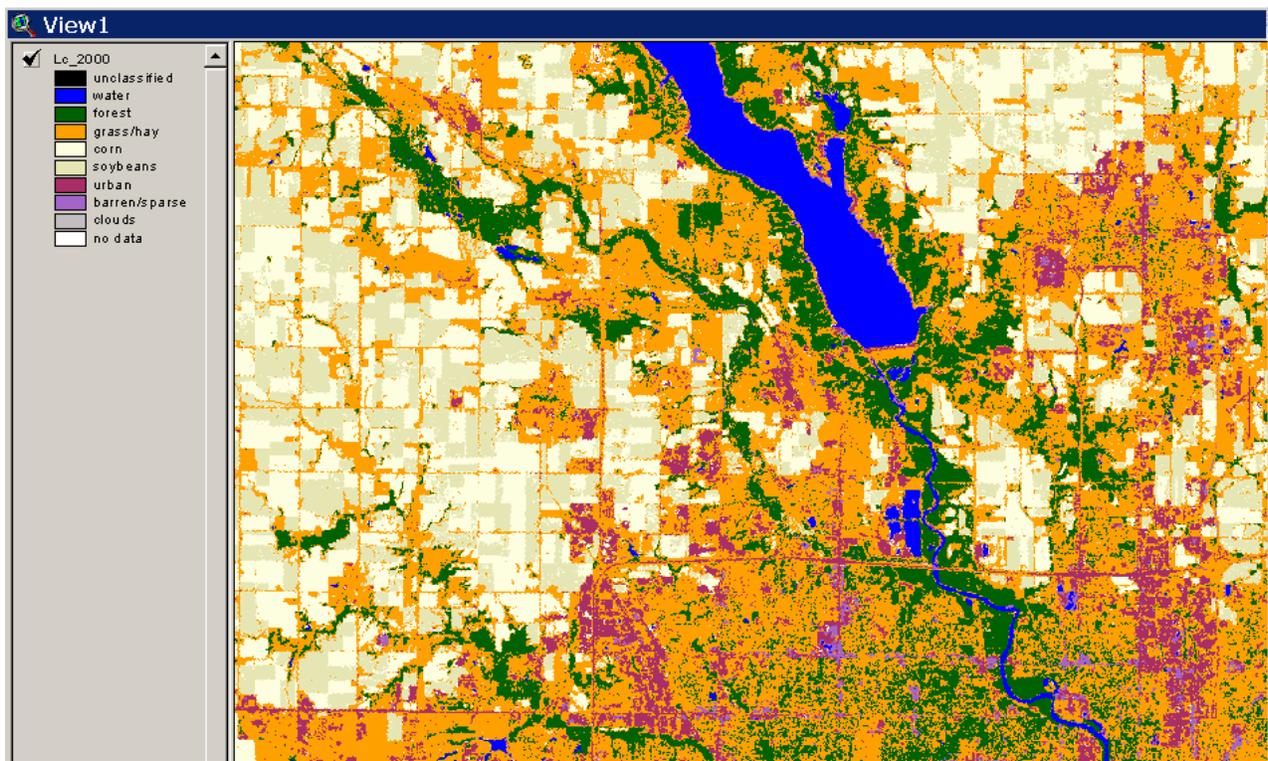


Figure 5. 2000 Land Cover and Legend, for comparison.