

High Resolution Lake Edge Extraction from Colour Orthophotography

Barry G. Pierce and York Law

*EBA Engineering Consultants Ltd.
1066 West Hastings St.,
Vancouver, B.C.*

Abstract

Permitting and development in the Canadian Arctic is made more difficult due to the low availability of quality baseline information. An overland route of approximately 160km is currently being developed for the southern portion of the Tibbitt to Contwoyto winter road, which has suffered from early thaws in recent years. Engineering designs were produced with the assistance of high resolution Lidar data; however, hard breaklines were required at water's edge to produce high quality elevation models. These lake positions were previously unknown with sufficient accuracy.

Multipath scatter of the raw Lidar returns over water prevented easy water boundary extraction from the Lidar data itself. Fortunately, the Lidar data also came with a series of 453 colour orthophotos. Water in these images were of variable siltiness, brightness and shallowness. An Earth Resource Mapper (ERMapper) algorithm was developed to select the water/land interface by choosing thresholds based on three derived measures, roughly described as Greyness, Blueness, and Smoothness. Ratios of grayness and blueness helped reveal transitions between water and land. Minimum smoothness criteria further helped reject strong edge transitions over land. Greyness was further used to remove surrounding black representing no-data.

Sections of the image sequence were captured as ERMapper mosaic algorithms. This allowed several images to be sequentially processed as one large virtual image. The image processing steps could in theory be run on one enormous virtual image mosaic containing all the individual scenes. It was experimentally found that vector conversion from ERMapper's .erv format to shapefiles seemed to be a memory-bound step, and limited the number of mosaic images in practice to about 20 at one time. Segmentation was performed on each mosaic section and then converted to vectors. The image segmentation cleaning was performed in ESRI's Arc/Info environment on vector polylines converted to coverages. Segment refinement was performed by rejecting line segments less than 50m in length, corresponding to erroneously selected land and rough water areas. The resulting cleaned vectors were then manually re-assembled to produce lake edges.

The algorithm performed well over a range of water surfaces. Thresholds conservative enough to returned unbroken water edges also falsely selected small terrain

May 2, 2008

features, choppy water and seams of adjacent ortho tiles as waters edge. The pre-cleaning steps made selection of the appropriate segments much easier. Occasionally, segments had to be drawn by hand where silt at water’s edge confounded the segmentation process. Final elevation models were produced from Lidar with hard breaklines at the extracted waters edge. These models were successfully employed for road engineering design and ecosystem and habitat mapping.

Acknowledgements

Many thanks to Andrew Deepwell for countless hours reassembling the cleaned vector output of this algorithm.

1 Introduction

For nearly thirty years, a joint venture of diamond mining interests in the Northwest Territories have maintained a winter access road penetrating far to the north of Yellowknife. The winter road is about 700km long, and consists of frozen lake crossings with connecting overland “portages” (Fig 1). It’s surface is rebuilt each new season, and it carries convoys of slow-moving long-haul trucks for a few months per year. The short duration of operation represents the only economical opportunity for heavy freight to and from the joint venture mine sites. Understandably, the early failure of the southern portion of the winter road in the 2006 season was met with anxiety, and a sense of urgency. In addition to the problem of extracting heavy equipment stranded far from civilization in a peat-swamp environment, there was the huge extra cost of flying in needed resources. With the warm-

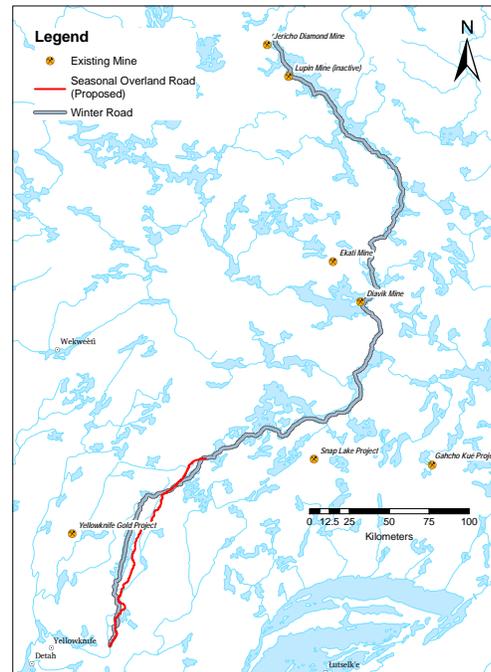


Figure 1. Proposed Seasonal Overland Route and existing winter road

ing trend in arctic climate, the joint venture looked for alternatives. EBA was contracted to design an overland route to replace the bottom, roughly 160km thaw prone portion of the route.

EBA’s transportation engineers acquired Lidar data over an approximate 500m swath for 165km of the proposed route. Point density was one metre with a 20cm vertical resolution. The Engineers wanted to develop detailed terrain models as part of the design-build process. In order for these models to be accurate, however, hard breaklines were

required to delineate water features. Reference vector data for this area of the Arctic was too coarse to be useable for this purpose. Further, Lidar over water exhibited a lower return density, but it was clear that multipath and delayed reflections resulted in false “terrain” elevations. Delineation of water from the raw Lidar was thus untenable. Fortunately, the Lidar data was bundled with 453 reference colour orthophotos. Applied methods to extract the water/land interface from these orthophotos was developed to extract the hard breaklines required.

One obvious method for extracting wet areas would be to look at near-infrared sensor returns. The “red edge” transition between red and near-infrared regions (cf. Gitelson & Merzlyak, 1986) provides a good demarcation between vegetated surfaces and open water. Strong absorption of near-infrared over water leads to characteristically low (dark) returns in these images. Unfortunately, the orthophotos available did not capture into the near-infrared region of the spectrum. Attempting to capitalize on the information available, a process was followed analogous to the development of the Tasseled Cap transform. This transform is best-known from the Landsat Thematic Mapper version (Crist & Cicone, 1984); however, the original paper by Kauth & Thomas (1976), although difficult to obtain, is a masterpiece. The requirements of the work presented here were very applied, but Kauth and Thomas’ paper suggests a future framework for making rigorous the methods developed here.

Ultimately, a simple image threshold would be desired which separates water from non-water. Gonzalez & Woods (1993) discuss the implications of thresholding exhaustively. Of particular note is their assertion that multiple thresholds derived for multi-modal histograms are less accurate than single thresholds. In the current work, segmentation based on the visual wavelength $\{Red, Green, Blue\}$ triplet could not be adequately performed with a simple bimodal threshold separation value. Water areas in the images exhibit a large range of siltiness. In many instances, algal mat features were visible below or on the water surface, as were underwater rock shelves. There was also a significant variability in image brightness, especially in the northern portion of the route. Gonzalez & Woods (1993) note that uneven illumination results in a broadening of the histogram, making it less modal and thus less able to be precisely segmented. Further, water’s edge becomes difficult to place with precision, even for a human operator, due to the prevalence of frost jack (large chunks of oblong rock that have split away from bedrock due to frost expansion) and deadfall. To meet the engineers’ requirements, the procedure developed would have to operate unmodified over a large range of image characteristics in a mostly automated manner.

2 Methods

Using the work of Kauth & Thomas (1976) as inspiration, band combinations of $\{Red, Green, Blue\}$ were explored to determine how well they could extract lake edges. From visual inspection, it was apparent that there was an immediate distinction between blue open water and greyish silted or algal bloom-laden water. Another apparent difference between water and land was the highly fragmented texture of the landscape at 20cm resolution. Both sparse tree and shrub canopy and frost jack broken rock led to a very different texture over land than over water (Figure 2). Three initial measures were developed to attempt to capture these distinctions:

$$\frac{\min(red, green, blue)}{\max(red, green, blue)} \quad (1)$$

$$\min(red, green)/blue \quad (2)$$

$$(red + green + blue)/3/255 \quad (3)$$

where Equation 1 measures “grey-ness”; Equation 2 measures “blueness”; and Equation 3 measures “smoothness”. Equations 1 and 2 were post-smoothed with a 3×3 median filter, while equ.3 had a two-pass 3×3 median filter applied. The filtering helped make surrounding values more similar without obscuring transitions between regional characteristics. These measures are suggestive of the properties their names connote, but were developed in a working applied environment where expediency was required, and thus empiri-

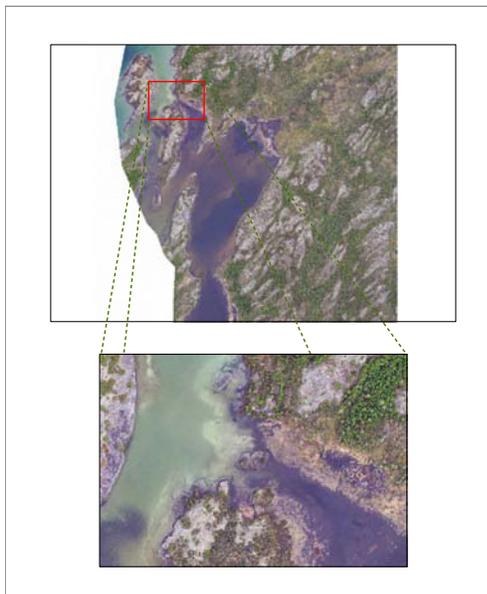


Figure 2. Image variability over water and land (route km 80)

cally rather than theoretically based. For example, values of Equation 2 actually decrease with increasing “blueness”. These issues were left to be cleaned up in a subsequent peer-review submission, where a more rigorous theoretical presentation can be made.

Boundary conditions were empirically derived from the three measures above to segment the water/land interface. Ratios of the derived measures were also tested for the ability to create a clean boundary segmentation. An Earth Resource Mapper (ERMapper) base algorithm was produced using subsequent images in the series as data sources. The base algorithm defined the measures above. It was found that ERMapper’s mosaic function could be used to first produce a mosaic of several images in sequence, and the mosaic could then be used as a data source to the base algorithm. A segmenta-

tion algorithm was then produced using the base algorithm as a data source. This last algorithm applied thresholds to the derived measures and ratios of those measures to produce the final segmentation logic. ERMapper’s raster to vector procedure was run, which resulted in both polyline and polygon segments. In general, the polyline layers encompassed the majority of the water/land boundary; unfortunately, significant portions of the boundary were occasionally encoded into the polygon layer as closed regions. This required treatment of the polygon layers as well, which contained very many small polygon features over land, which needed to be filtered out. Choosing more aggressive threshold values reduced data volume, but was unworkable, because significant lengths of true boundary features would then also drop out.

Treatment of the extracted vectors was done in Environmental Systems Research Institute’s (ESRI) Arc/Info software. The choice of Arc/Info over ArcGIS Desktop was necessitated by the sheer volume of vector data to be processed. Arc/Info’s coverage geospatial data model is very compact, and the design of the older Arc/Info system makes few assumptions about memory availability. ArcGIS Desktop seemed to attempt to load the entire vector layer into memory, and would fail for large datasets. EBA has since upgraded to the 9.2 release of ESRI’s ArcGIS platform. With the new File Geodatabase, or ArcSDE relational database system, the processing done in Arc/Info may be possible with the desktop soft-

ware.

ERMMapper *.erv* files were exported to ArcGIS *shapefiles* from ERMapper. These shapefiles were imported as *coverages* with double precision using the Arc/Info *shapearc* command. Polygon topology was not desired, so only the arc (polyline) portion of the result was used for polygon source data. Arcs from both source polygons and polylines were merged together in batch using the Arc/Info *append* command. Line (arc) topology was created with the *build* command. The output features were edited using the Arc/Edit *select* command, to select and delete arcs with lengths of less than 50m, thus removing the vast majority of spurious features from the original polygon data. This preprocessed output was exported to AutoCAD, where a technician manually selected and assembled the relevant pieces.

3 Results

3.1 Segmentation Algorithm Output

The workflow developed is illustrated in Figure 3 for the image tile from kilometre 102 of the route. The base image (also in Figure 4) is colour mapped as follows:

$$\{Smoothness, Greyness, Blueness\} \\ \implies \\ \{Red, Green, Blue\}$$

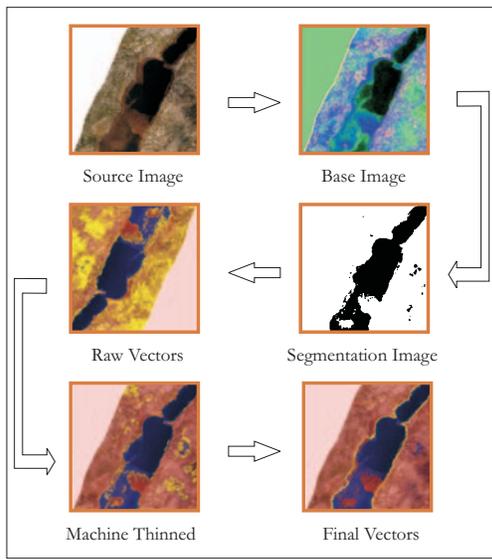


Figure 3. Workflow of steps in lake edge extraction

It is immediately apparent that the “blueness” measure is inverted, since land and silt show strong response; however, it is also clear that water is very well separated from land by this measure. In this particular image, water “smoothness” falls in the very bottom portion of the measure histogram (not shown). Land has a strong low value peak, which makes an effective segmentation threshold in this image. “Greyness” produces a less dramatic segmentation. A prominent confounding feature in this image is a prevalent underwater silt bed, which shows an erosion channel. The images are from summer, when meltwater has filled the lake basins. During early spring, features such as the runoff channel that are underwater in this image would be exposed. Transect positions of measure responses are shown in Figure 4. Profile values for these transects are shown in Figure 5.

Transect #1 follows a path from clear, open water, through underwater silt. It crosses the underwater meltwater channel and terminates near what appears to be the runoff channel of a small stream (bottom-left of image). The “smoothness” measure shows no interpretable pattern over open water, even given these underwater features; however “greyness” and “blueness” respond significantly to these features. “greyness” and “blueness” are identical over clear water, but there is a strong divergence of response over the underwater silt shelf, as shown between lines 1A and 1B on the figure. 1B is also the position where the transect crosses the underwater runoff channel. There is lesser divergence between the measures through silted water until 1C, where a small pocket of clear water exists. “Greyness” and “blueness” over clear water appear to be nearly mirror-images. Since the “blueness” measure is inverted (described above), it is likely that these two measures have a strongly correlated response.

Transect #2 is intended to show the transition between overland and water features. 2A is the position of the actual shoreline at the time of photography. There is siltation or a possible underwater shelf between 2A and 2B. 2A is a strong transition point in the “smoothness” measure. Values between 2A and the opposite shoreline at 2D are much smoother than over land. The graph illustrates that the “smoothness” measure is inverted, with lower values representing smoother textures. The picture of response in the “greyness” and

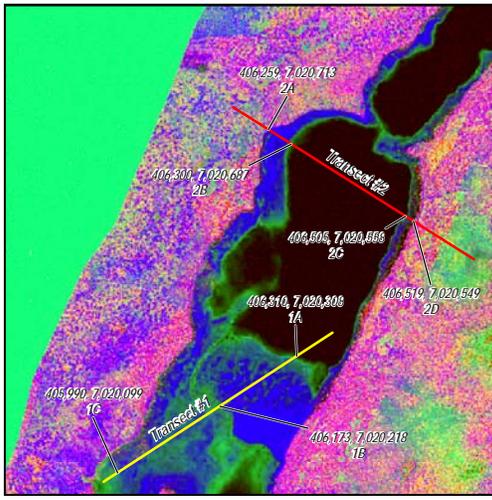


Figure 4. Image of derived measures and transect lines along lake at km 102

“blueness” measures is filled in by this transect. There is mirroring of the measures over land before 2A, but this effect *increases* over the slit shelf (greater than over land). Over open water, the measures are identical, as in Transect #1. The small separation in measures between 2C and 2D appears to be silt along the far shore, but the reason for the large diversion after 2D is not clear.

“Blueness” and “greyness” appear to have strong response over water, and also to underwater features. A ratio of these measures was developed to try to exploit their differing responses:

$$\frac{\min(\text{greyness}, \text{blueness})}{\max(\text{greyness}, \text{blueness})} \quad (4)$$

Output from Equation 4, the “grey/blue ratio”, is shown in Figure 6. Output saturates over clear water with a value of 1. For numerator and denominator of Equation 4 to be equal, both the “greyness” and “blueness” measures must be equal. This appears to

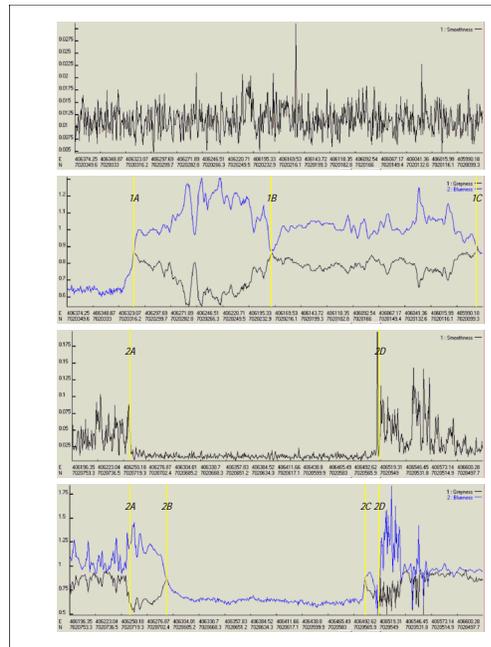


Figure 5. Transect plots for lake at km 102

be the case over clear water, as discussed for Figure 5. In Transect #1, we see this response up to 1A, where the underwater silt shelf occurs (Figure 7). There is a spike at 1B where a small patch of clear water shows through, and again at 1C where the underwater runoff channel is crossed. A final peak value occurs at 1D, at the location of the supposed creek outlet. Transect #2 shows that silt at water’s edge has a characteristically different response than land in this ratio, with $2A \rightarrow 2B$ and $2C \rightarrow 2D$ producing much lower values of the measure than over land or water.

Threshold values were empirically derived to produce good segmentation over the entire sequence of images. The values chosen are shown in Table 1. The limit on the “greyness” measure is simply to remove white null values from consideration. The imagery was strictly clipped by the

3.2 Segmentation Performance

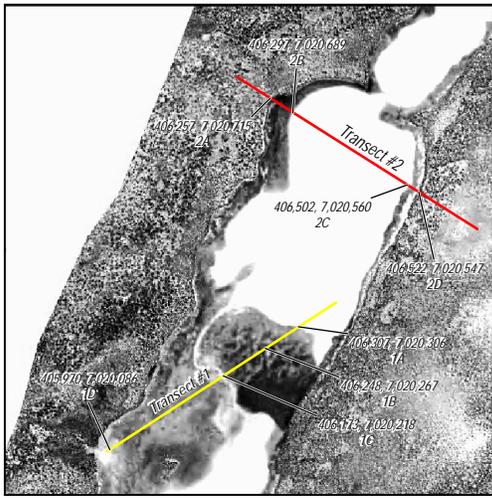


Figure 6. Image of “greyness” and “blueness” $\{min, max\}$ ratio measure–lake at km 102

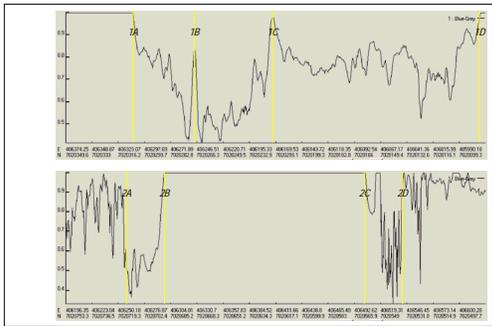


Figure 7. “greyness” and “blueness” $\{min, max\}$ ratio measure–lake at km 102

Table 1
Thresholds used for water/land interface image segmentation

$smoothness < 0.03$ $grey/blue\ ratio > 0.7$ $greyness < 0.99$
--

supplier to the limits of the Lidar returns, leaving areas outside the Lidar extents white.

An example of machine–derived and final, human–cleaned vectors are shown in Figure 8 (the semi-transparent fill showing the final water polygon is turned off in the bottom panel of these figures to allow the range of water colour to be seen). This lake, at km 80 of the route, contained a range of siltation. The segmentation performed very well over water. A large majority of extraneous arcs over land were removed automatically by machine thinning. The final operator–screened output uses the algorithm defined features; aside from assembling arcs, no additional intervention was required. Another lake at km 31 (Figure 9) was an anomaly. Lake water covers a bog/mud landscape, leading to a very atypical water colour. The algorithm created two potential lake edges, bounding a mud shelf that is in places above water and in others submerged. The inland portion which abuts treed landscape was selected in the final version. Staff who were onsite confirm that this lake exhibited rapid changes in water level, and that the shoreline was a viscous sludge of bog and mud (Figure 10). Nonetheless, the algorithm’s selection of a shoreline is plausible. A final example of a lake at km 117 (Figure 11) shows the impact of surface ripples and reflected cloud. The edge contrast between these features and calm, blue water led them to be delineated as water’s edge. This lake had a submerged shelf along the shoreline. Portions of the algorithm output select the lake–side extent of

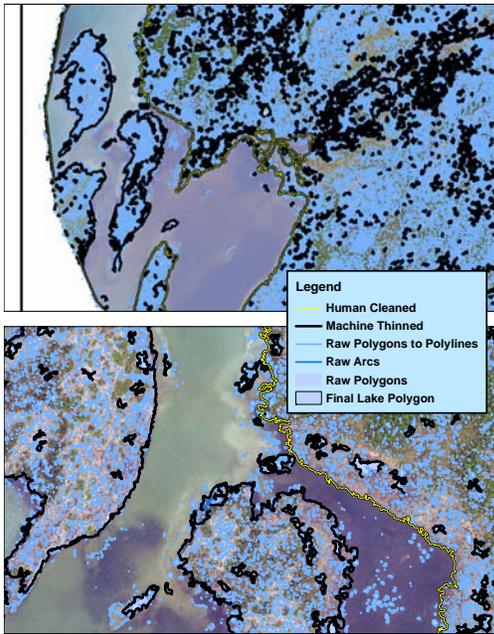


Figure 8. Algorithm and final cleaned output for lake at km 80

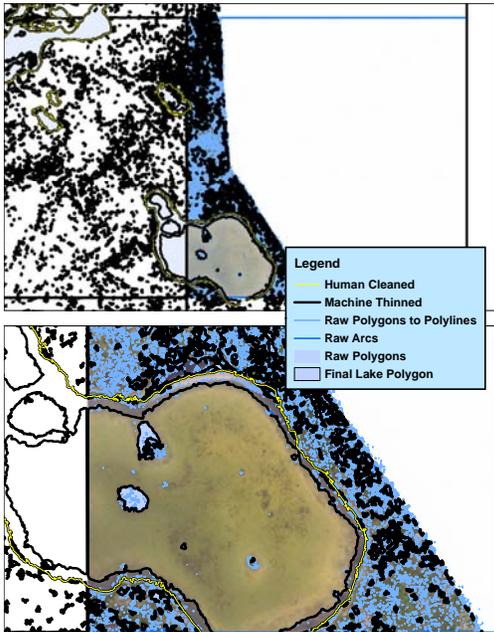


Figure 9. Algorithm and final cleaned output for lake at km 31

this shelf, but the shore-side portion is so fragmented that almost the entire shoreline had to be manually digitized for this section of the route.



Figure 10. Ground image of lake at km 31

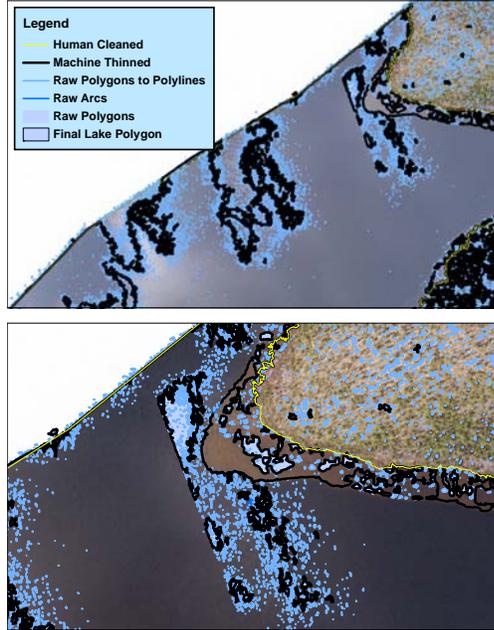


Figure 11. Algorithm and final cleaned output for lake at km 117

4 Discussion and Conclusions

In practical terms, the output from the algorithm discussed was successfully used to assemble complete lake edge coverage along the entire proposed overland route. These final lake edges were successfully employed as hard breaklines for surface modelling. The amount of human-operator time required to clean and

assemble the output was much less than would be required to manually digitize the lake boundaries unassisted. Pre-cleaning of short segment lengths helped to remove many extraneous features over land and water. The algorithm handled varying degrees of siltation or underwater features such as submerged bog, rendering plausible lake edges even under unusual conditions. Reflected cloud and surface ripples on water caused erroneous edge placement, and fragmented the shoreline so severely as to require entire sections to be digitized manually. Fortunately, these image conditions occurred rarely. The “blueness” and “smoothness” measures are actually inverted: larger values representing *less* of the property being measured. The scale of all the measures except the “grey/blue ratio” are unbounded on the upper extent, with values of the “smoothness” measure being extremely small. Future work would involve standardizing the direction and upper bounds of all measures, so that values increase for increasing presence of the property being measured, and so that a value of 1 represents total saturation.

References

- Crist, E.P., Cicone, R.C. (1984). A Physically-Based Transformation of Thematic Mapper Data—The TM Tasseled Cap. *IEEE Transactions on Geoscience and Remote Sensing*, **GE-22(3)**, 256-263
- Gitelson, A.A., Merzlyak, M.N. (1986). Signature Analysis of Leaf Reflectance Spectra: Algorithm Development for Remote Sensing of Chlorophyll. *Journal of Plant Physiology*, **148**, 494-500
- Gonzalez, R.C., Woods, R.E. (1993). *Digital Image Processing*. New York: Addison-Wesley.
- Kauth, R.J., Thomas, G.S. (1976). The Tasseled Cap—A Graphic Description of the Spectral-Temporal Development of Agricultural Crops as Seen by Landsat. In Proceedings of the Symposium on Machine Processing of Remotely Sensed Data. West Lafayette, Indiana, **4B41-4B51**