Modeling Buildings for Mission Rehearsals

Stuart Blundell
Visual Learning Systems, Inc.
PO Box 8226
Missoula, MT 59807
sblundell@vls-inc.com

Chris Finlayson
Visual Learning Systems, Inc.
PO Box 8226
Missoula, MT 59807
cfinlayson@vls-inc.com

David W. Opitz
The University of Montana
Computer Science Department
Missoula, MT 59812
opitz@cs.umt.edu

Abstract

The modeling of building footprints, heights, and roof shapes is critical to U.S. Defense Department mission rehearsal since buildings dominate and characterize urban landscapes. The purpose of this paper is to present a new system for creating, from elevation data, 3D models of buildings and trees. The system, built as an extension to Feature Analyst, simplifies building and tree shapes and automatically incorporates these models into ArcScene for the purpose of mission-specific rehearsal. The system also contains clean-up tools to assist the user in remedying common mistakes made during automated extraction.

Introduction

The ability to quickly and accurately update geospatial databases is crucial for most military mission rehearsal operations. Recent evolvement of some toolkits, such as Feature Analyst®, has greatly increased the proficiency of extracting features from panchromatic and multi-spectral imagery; however, effectively delineating certain features (such as buildings) is still proving to be difficult. The modeling of buildings and trees is critical to the U.S. Defense Department since these features dominate and characterize urban landscapes. Effective mission rehearsal for troops requires simulation units to have up-to-date detailed information of the landscape.

LIDAR data provides analysts with high-resolution elevation data to interpret and model information about the world. Such data is conducive to automated detection of buildings because, while rooftops come in too many colors and textures, buildings have consistent sharp drops in elevation. This paper presents a new LIDAR toolkit that exists as an extension to

Feature Analyst. The goal of the Feature Analyst LIDAR toolkit is to create a set of algorithms for extracting useful features from high-resolution elevation data. This paper concentrates on the extraction of bare earth, buildings, and trees.

Background

The objective of image classification is categorizing images into feature classes or themes. The thematic maps are loaded into a GIS database and support decision-makers for a wide variety of applications such as land-use planning, disaster and emergency services, and military applications. Image classifiers can also be used to extract from imagery some types of specific objects or targets, such as land-cover types. The real world task of feature extraction is difficult because: (a) two objects may look entirely different but comprise the same feature type and (b) spectral information alone is not enough to extract all features. It is not surprising that heads-up digitization, or hand classification, was for a long time the most common approach to feature extraction using imagery. It is often the most accurate and available known approach for many GIS analysts. In this process, a trained specialist manually traces the outline of a feature by clicking points on the screen. For many applications this laborious process is unfeasible for the following reasons:

- The process is slow and expensive.
- There is a lack of trained GIS analysts available to perform the task.
- The resulting quality of the extraction diminishes as the analyst becomes tired.

In addition to manual classification and traditional image processing techniques, a recent alternate approach is to model the change detection process using statistical and machine learning techniques (Maloof et al. 1998; McKeown 1996). The idea behind this approach is, with a sample of changed features from the image, a learning algorithm automatically develops a model that correlates known data (such as pixel values from images, terrain data, vector overlays, and grids) with targeted features. The learned model then automatically classifies and extracts the remaining features in the imagery set. Traditional supervised change-detection classification methods have predominately used basic statistical methods, such as Maximum Likelihood. These methods require priori statistical assumptions and often fail on small-feature capture because of their inability to: (a) classify disjunctive concepts, (b) take into account spatial context, and (c) remove clutter.

Feature Analyst is a new machine-learning, feature-extraction system that overcomes these obstacles. Feature Analyst provides: (a) a simple interface and workflow process, (b) the ability to take into account spatial context (contextual classification), (c) the ability to mitigate clutter, (d) the ability to learn disjunctive concepts, and (e) the ability to incorporate ancillary data sources such as elevation models. Feature Analyst is available as an extension to ArcGIS and ERDAS IMAGINE. Feature Analyst works by having an analyst provide sample target features in an image set. It then learns from these examples and automatically detects the remaining features.

When interpreting the contents of imagery, there are only a few attributes accessible to human interpreters. For any single set of imagery these are: Shape, Size, Color, Texture, Pattern, Shadow, and Association (Opitz, 2002). Traditional image processing techniques incorporate only color (spectral signature) and perhaps texture or pattern into an involved expert workflow process; Feature Analyst incorporates all these attributes as well as ancillary data sources.

Feature Analyst LIDAR Toolkit

The Feature Analyst LIDAR toolkit exists as an extension to Feature Analyst and currently contains functionality for automated and semi-automated extraction of various features from LIDAR. In this paper, we concentrate on the extractions of building footprints, bare-earth, and trees.

Bare Earth

Bare earth is the modeling of terrain with all elevated features (mostly buildings and trees) removed. Bare earth is critical for effective modeling and displaying of the building and tree volumes on an accurate real-world terrain. In order to extract bare earth, edges are detected from the last return elevation data, and areas of sharp change in elevation (i.e. edges of buildings, trees) are picked out and removed from the original image. When extracting bare earth with the Feature Analyst LIDAR toolkit, the user simply picks the final return (mandatory), and if available, the first return and the system effectively models bare earth. An analyst is given control to specify the appearance of the removed buildings and trees on the bare earth. In some cases, an analyst will prefer to have building footprints to be slightly elevated “landing pads” and in other cases, the analyst will prefer to have the visible ground seamlessly merge with the remaining area of the removed features. The detail of the visible earth is maintained, and the user has the option to automatically remove shrubs and small man-made objects such as cars.

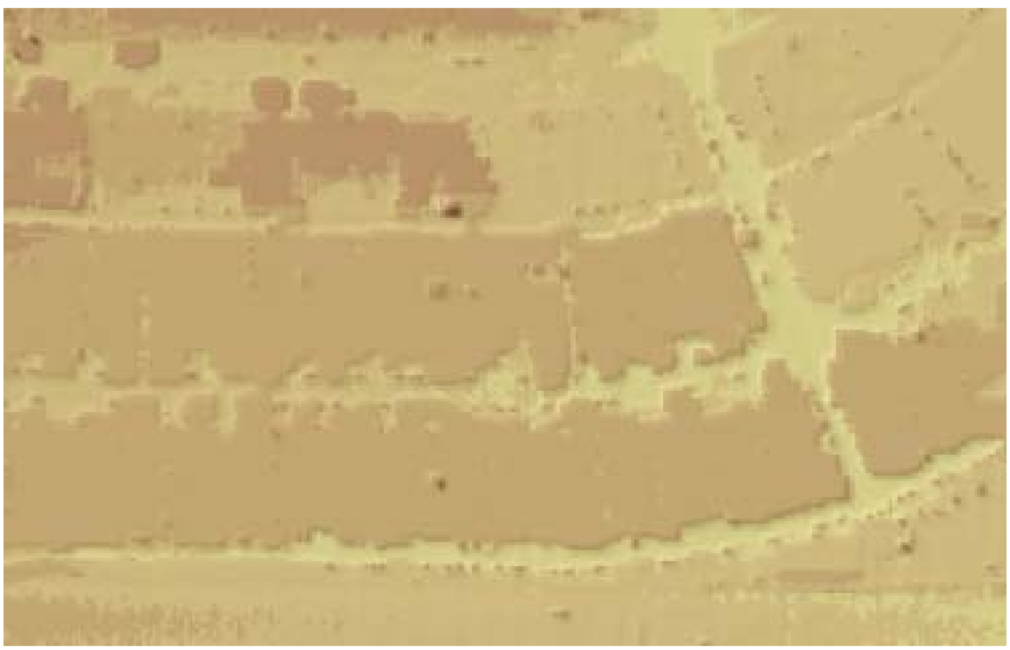

Figure 1: Bare earth representation of a heavily treed residential neighborhood.

Building Footprints

The Feature Analyst LIDAR toolkit automatically extracts building footprints by first finding bare earth. The bare earth is then subtracted from the original elevation data, at which point the software categorizes the remaining data into buildings, trees, and other features (shrubs, cars, etc.) based on shape and curvature. The manner in which this is done will be discussed in more detail in the section on tree extraction. We tested the accuracy of the automated building extraction algorithm of the toolkit and found that out of 184 buildings in a 1-meter resolution LIDAR data collected over a suburban/urban area; the automated extraction process had 8 false negatives (4.3%) and 18 false positives (9.8%). This significantly saved analyst time on the extraction of such buildings.

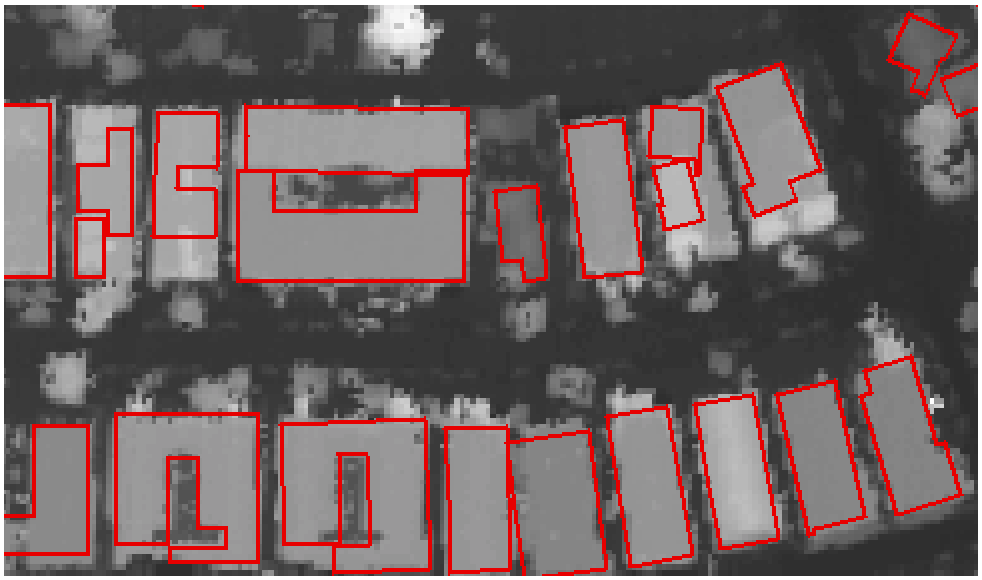

Figure 2: Unedited extracted building footprints overlaid on last return LIDAR

Clean-up tools

Most of the problems with extracting buildings are: 1) two buildings located so close together they are extracted as one, 2) buildings with “holes” where there should be none, 3) parts of buildings missed, 4) extra pieces added to a building, and 5) adding buildings that were missed entirely. In order to address these problems, the Feature Analyst LIDAR toolkit has numerous cleanup tools. Many GIS platforms, such as ArcGIS, support editing polygons, but fall short when it comes to addressing the problem of editing large numbers of polygons. The problem is that one must enter into distinctly different edit modes to complete these tasks. This may not seem a problem *prima facie*, but when working with large datasets, it can be time-consuming. For instance, a 3% error rate over a region with 5000 buildings means having to edit 150 buildings manually. Feature Analyst keyboard shortcuts allowing a user to enter different edit modes, greatly reduces this editing time. Figure 3 through Figure 8 below illustrate the shortcuts.

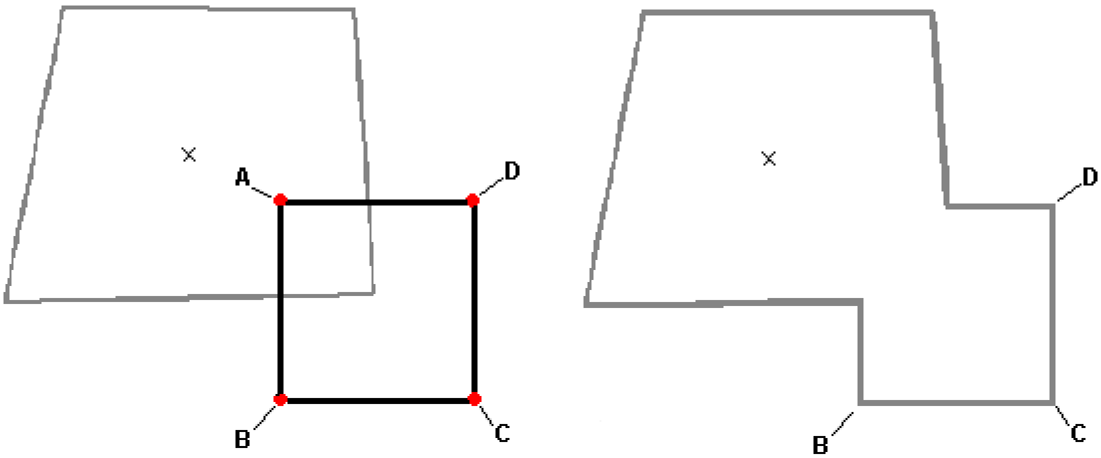

Figure 3: Extending a polygon feature. Clicking on points A through D extends the polygon as shown.

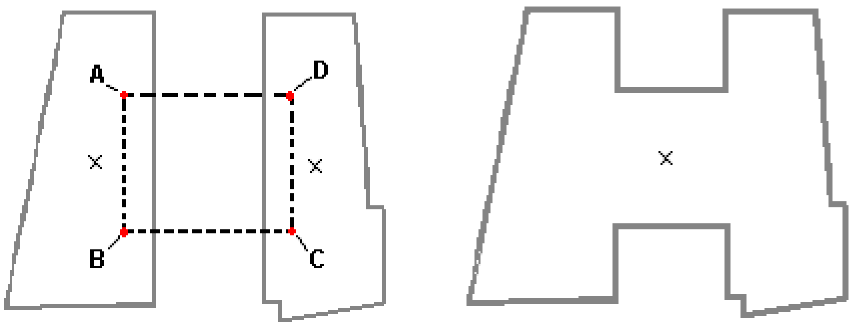

Figure 4: Merging polygon features. Clicking points A through D merges the two polygons automatically.

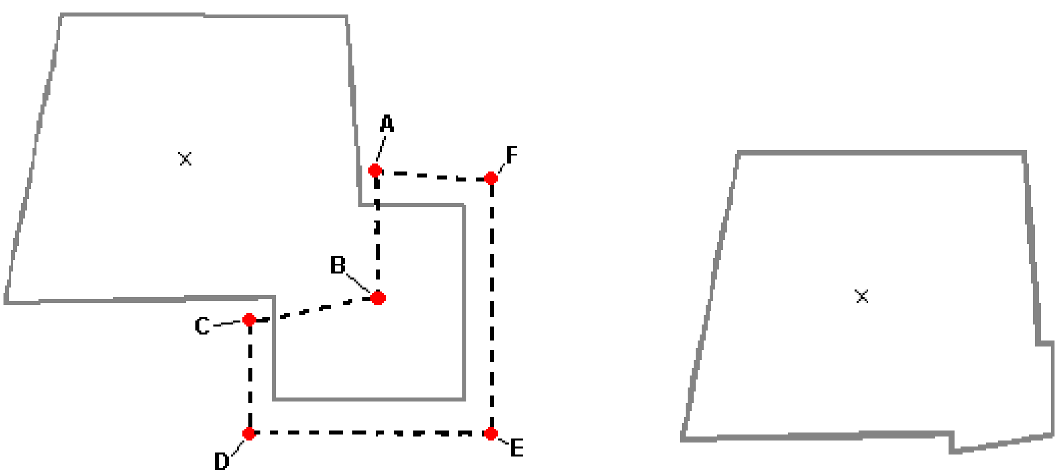

Figure 5: Trimming feature. While holding the shift key down, clicking on points A through E trims the polygon as shown.

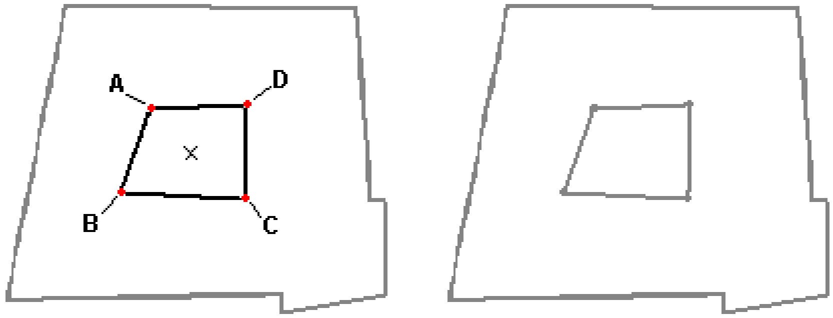

Figure 6: Cutting a hole in the middle. Clicking on points A through D while holding the shift key down cuts a hole in the middle of the polygon as shown.

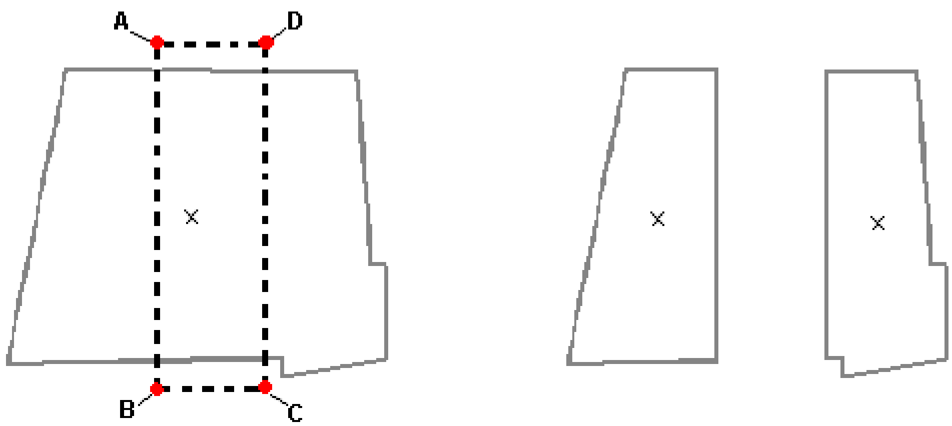

Figure 7: Splitting a polygon. Clicking on points A through D while holding the shift key down splits the polygon as shown.

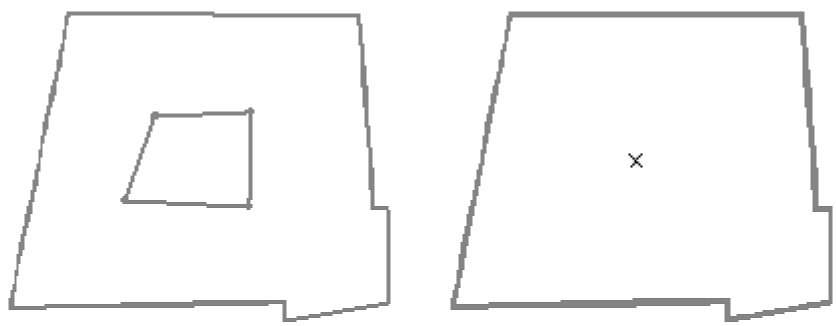

Figure 8: Closing hole in polygon. Holding the ctrl-key down and double-clicking clicking on a hole in a polygon will close the hole.

Trees

Trees are detected similar to buildings. First, the bare earth is determined, and subtracted from the original image. The remaining features (buildings, trees, other) are determined on their three-dimensional curvature. In most cases, building roofs are flat, or flat and sloped (like the four sides of a hipped roof). Trees and foliage are more dome-shaped with some irregularities within them; rarely would one find perfect geodesic domes. The observation that 3-D curvature is a good visual separation between trees and buildings is leveraged in our proprietary algorithm to separate trees from buildings.

Both the building extraction process and the tree exaction process require the LIDAR data and one numeric parameter representing a cut-off threshold. The cutoff threshold parameter is the 3-D curvature of the shape. The higher the value, the more curved the shape must be in order to be returned. Therefore, when extracting trees, if one has a really high value, then it is requiring that the features be very curvy and so only curved objects, which would be perhaps a few types of foliage, would be returned. One would expect few objects to be returned. If the value is really low, then objects with little curvature will be returned. Therefore, for building extraction the value is generally lower than for tree extraction.

An example of the extracted trees in a small region is shown in Figure 9.

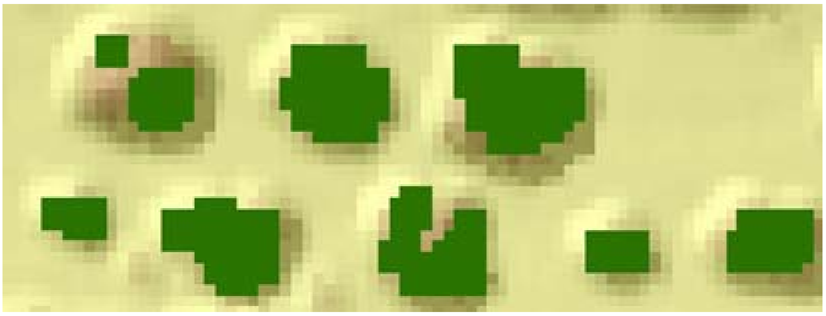

Figure 9: Extracted trees overlaid on first return LIDAR

Visualization

3-D visualization is fairly straightforward with the resultant building footprints. One may load the results in ArcScene. When looking at the image, some visual artifacts are obvious. This is due in part to the 1-meter resolution of the LIDAR data, as well as the fact that the building footprints don't match up exactly with the edge of the building. Some objects found on the top of buildings, such as air conditioning units and other structures often appear in the visualization. There structures can easily be removed by averaging the data of the data that lied under the buildings. To take account of potential spikes in the data, either because of poor LIDAR resolution or structures on top of the building, one can average only the data that lies within one or two standard deviations of all the points. An example of the latter is shown in Figure 10.

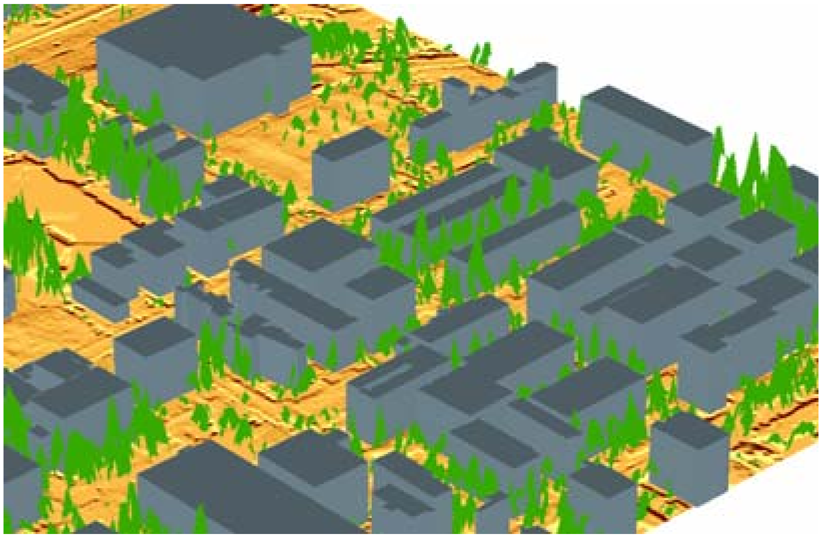

Figure 10: Visualizing extracted trees and buildings in ArcScene

In addition to the 3-D visualization, a standard hill-shade relief tool is provided to view the elevation data as a relief.

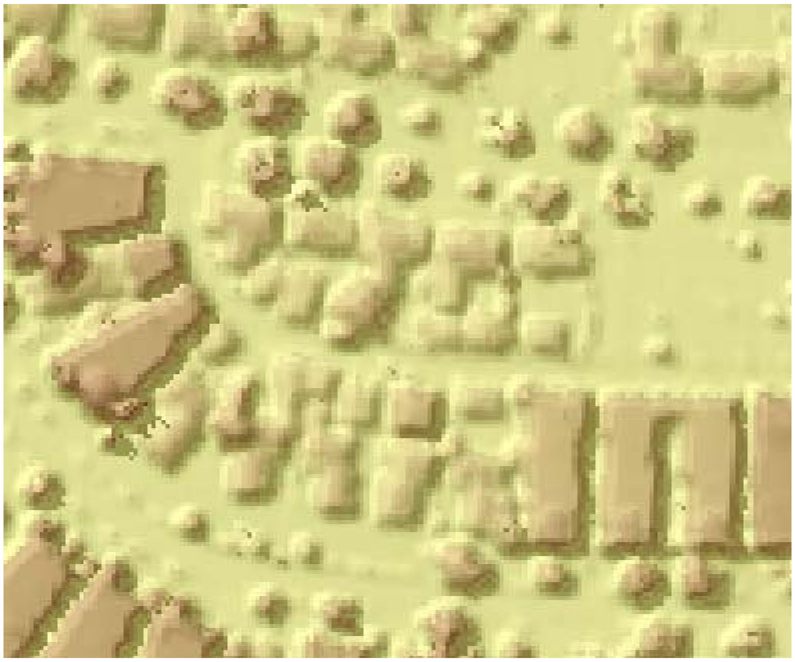

Conclusions and Future Work

The Feature Analyst LIDAR toolkit promises to change the way analysts collect data. Buildings are an example of features that are extremely hard to extract automatically from spectral imagery, but appear to be highly accurate from LIDAR. The solutions appears to be a fusion of spectral information with elevation information to allow nearly all features (such as drainage, vegetation, buildings, and roads) to be automatically extracted with the new Feature Analyst

LIDAR toolkit’s capabilities. Future work includes tools to facilitate clean up of the final results. These tools are described below.

Include Area in Bare Earth Model: Polygons are digitized over areas the user wants to include in the bare earth. For example, if the user digitizes a small area that corresponds to a building, the entire building will show up the revised model.

Remove Area from Bare Earth Model: Polygons are digitized over areas the user wants to remove from the bare earth model. For example, if user digitizes an area corresponding to a building that still remains, the building will be removed in the revised model.

Filtering (smoothing, sharpening, etc.): The user digitizes and selects a filter to run over a region to achieve more visually accurate results.

Remove Shrubs, Cars, Etc.: Polygons are digitized over areas the user wishes to remove artifacts, and the areas are smoothed.

Flatten to Specified Elevation: The user digitizes a polygon over a region and is prompted if he wants to compute the average elevation of the polygon. He is then prompted with the average value of the bare earth and the average value of the DEM for that polygon and asked to enter an elevation value to assign to each of the pixels in the polygon.

Region Grow Mode: A region grow tool can assist the user in many digitizing steps for example, when digitizing areas to filter, areas to remove artifacts, or areas where buildings are missed.

References

Maloof et al. 1998. "Learning to Detect Rooftops in Aerial Images." Image Understanding Workshop 835-845. Monterrey, CA.

McKeown, D. 1996. "Top Ten Lessons Learned in Automated Cartography," Technical Report CMU-CS-96-110, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA.

Opitz, D. 2002. "The Use of Spatial Context in Image Understanding." Ninth Biennial Remote Sensing Applications Conference.