

Visual Analysis and Semantic Exploration of Error Aware Urban Change Detection

Paper 1104

Abstract

Many previous approaches to detecting urban change from LIDAR point clouds do not consider error and uncertainty and simplify their calculations by interpolating into a regular grid structure, producing a 2D image as output. We present a method of LIDAR change detection that accounts for error and uncertainty, maintains accuracy, and extracts relevant changes as individual 3D models. We then utilize these models, alongside existing GIS data, within an interactive application that allows the chronological exploration of the changes to an urban environment. A three-tiered level-of-detail system maintains a scale-appropriate, legible visual representation across the entire range of view scales, from individual changes such as buildings and trees, to groups of changes such as new residential developments, deforestation, and construction sites, and finally to larger regions such as neighborhoods and districts of a city that are emerging or undergoing revitalization. Tools are provided to assist the visual analysis by urban planners and historians through semantic categorization and filtering of the changes presented.

Categories and Subject Descriptors (according to ACM CCS):

1. Introduction

In the mission statement for the USGS Center for LIDAR Information Coordination and Knowledge (CLICK) it is acknowledged that while "there has been increasing demand for research utilizing all information generated from LIDAR remote sensing data and not just [the resulting] bare earth [Digital Elevation Maps (DEM)]," "research on using the entire point cloud of this remote sensing data for scientific applications has been slowed by a steep learning curve on research and understanding involving utilizing the entire point cloud." [USG07] Indeed, when examining previous approaches to the detection of urban change through LIDAR data, we can see that most researchers choose to reduce the complexity and detail of the point clouds into raster based approximations for purposes of speed and ease of computation. It is evident that utilizing the unsimplified point cloud has great advantages both in retaining maximum quality in the resulting models and the ability to accurately determine the uncertainty of the data, leading to knowledge of the confidence one can have in those models and any calculations performed on them.

The goal of a change detection algorithm should not be to simply calculate the differences between two samplings, but instead to determine, with a measure of confidence, if the differences between the two are due to actual changes in the physical environment or are the result of a combination of various errors and uncertainties. As such, we start by requiring data quality measures for our input data. These values allow us to determine how much deviation between points in the datasets could be due simply to collection errors and therefore would not represent a change in the actual physical world being sampled.

Our system presents these change models alongside existing Geographic Information System (GIS) data, within an interactive application that allows the chronological exploration of the changes to an urban environment. The system has multiple coordinated views including a main 3D view, a heat map view and a zoomed-in inspection view. For the main 3D view a three-tiered level-of-detail system maintains a scale-appropriate, legible visual representation across the entire range of view scales, from individual changes such as buildings and trees, to groups and regions of changes.

Tools are provided to assist the visual analysis of changes and their patterns through semantic categorization and filtering of the extracted change models. Heat map signatures capture the relationship between the footprint areas and heights of change models. Semantic filters can be applied to suppress or highlight typical classes of physical objects such as trees, residential houses and road re-gradings.

2. Previous Approaches

A number of researchers have devised methods of detecting changes in urban environments between LIDAR scans taken at different times. Vu et al. [VMY04] focuses on updating GIS records and assessing damage from earthquakes, Murakami et al [HMI99] sees application in the enforcement of real estate taxes.

A common strategy these researchers take when attempting detection of changes between LIDAR scans is the transformation of the LIDAR point clouds into a grid using nearest-neighbor interpolation. As described in [VMY04] this tactic is chosen because it is effective at decreasing the amount of time needed to perform calculations. A grid resolution is determined based on the density of the LIDAR samples from the pair of scans, and the values for each grid location are generally determined by nearest-neighbor interpolation to preserve the sharp height changes at building edges.

Murakami subtracts the grid values of one scan from that of the other to produce a grid of values representing height changes; image processing morphology operations are used to filter out noise, and the cleaned-up difference image is merged with an orthoimage for manual interpretation [HMI99]. In [VMY04] Vu et al determine the statistical distribution of the intensity values of the difference image and define those values beyond a certain number of standard deviations from the mean to be possible new construction, and likewise those below a certain number of standard deviations from the mean to be possible demolition. Morphological opening and reconstruction is again used to filter out undesirable results. Output from these approaches are most commonly raster images or DEM files showing those buildings determined to be newly constructed or demolished/damaged.

We argue that gridding the LIDAR data for ease and speed of calculation not only introduces error to the data, but fails to properly accommodate scans of different resolutions and overlapping swaths, which can provide areas of higher resolution. The collection hardware used for each scan may have significant differences in accuracy which should be preserved. A likely failing of these techniques is their indifference towards the uncertainties of the sample collection methods themselves, although in fairness these are not likely to effect the detection of changes as large as entire buildings. To our knowledge there has been no interactive visual analysis of raw LIDAR point cloud data similar to what we present in this paper.

Finally, our analysis tools combines two 3D views with a heat map. Heat maps have been used successfully in visualizations tools in multiple application domains [AvH04] and are included with several commercial tools [Spo05]. In genomics, heat maps are used for the visualization of massive gene arrays [JSS02] and provide a compact overview of the data as well as a drill-down capability for detailed information. Similar to our system, Chang et al [CWK*07] combine an interactive heat map with a 3D visualization.

3. Implementation

The motivation for this work is to visualize the changes due to construction and development between yearly LIDAR scans of Mecklenburg County, North Carolina, which contains the city of Charlotte and roughly 370,000 buildings. The annual scans register quite well with each other. However, the density of the sample points leaves something to be desired, with roughly 3-4 meters between points. The only data used in this analysis is the raw LIDAR returns, in the form of x,y,z points, which have not undergone filtering to remove buildings or trees. This method is attractive because it does not require the use or availability of supplementary data such as classification or intensity values for the LIDAR returns or aerial/satellite imagery.

3.1. Pre-Processing

Before processing, values for the horizontal and positional accuracy of the LIDAR system used for collection must be specified. These values allow us to account for the uncertainty of sample points taken by different hardware systems, as would likely be encountered when comparing datasets collected across a span of years. By knowing the accuracy/error characteristics of each dataset, it can be determined how much deviation between points in the datasets could be due simply to collection errors and therefore would not represent a change in the actual physical world being sampled.

Firstly, irregular triangular meshes are generated from the LIDAR points using 2D Delaunay triangulation (we utilize the CGAL library [Yvi07]). The resulting meshes serve two purposes. During change detection points are projected into the mesh of the opposing scan which allows us to quickly determine surrounding points for comparison. The meshes, as well as simplified (using [GH97]) versions of them, are also stored for use as a reference surface models in the user interface.

3.2. Error Aware Change Detection

The goal of the change detection process is to identify and mark all the points in the new scan that are most likely to be changes in the actual physical world and not merely changes in measurement (collection error).

This is done by iteratively evaluating each point in the new scan in the following manner: First the point in the new scan is projected in 2D onto the triangulation of the old scan to determine which face it falls within. The distance along that face, from the projected point to the nearest vertex of the face, minus any positional errors, is calculated. This distance is then used to determine a height allowance based on the possible geologic variation that could be present between the known sample points. The height of the point above the interpolated face, minus any vertical measurement errors, is then compared against the height allowance. If the point exceeds the bounds of what could be accounted for with known errors and geologic variation, then it is considered to be a legitimate change point and is marked as such.

See [BCWR07] for a more complete description of the errors involved in generating terrain models from airborne LIDAR points, and the techniques we have used here to account for these errors. Specifically, the technique used here to classify single points is equivalent to the technique used there to determine the occupancy of voxels, with the exception that we have incorporated additional positional and elevation errors for the incoming point.

3.3. Filtering and Model Creation

The resulting set of marked points that have been classified as significant changes will generally contain many "noisy" objects caused by collection/data errors, low-resolution scan returns sporadically striking tall thin objects such as broadcast antennas, street signs, telephone poles, etc.

Vu et al [VMY04], Murakami et al [HMI99], and others have utilized morphological opening and reconstruction to filter out small noisy changes from their results. This was suitable for their purposes as their difference results were in the form of a 2D grid with each pixel marked as being the same or having changed. These image processing techniques both reduce the accuracy of the changes reported and are not applicable to our non-gridded point cloud results. We integrate our noise filtering process within the process of creating models of the detected change points.

We begin by selecting a marked point and examining its neighboring vertices in the triangulation to determine if any have also been marked as changed. Here we can implement the simplest form of noise reduction by disregarding any marked vertices with no marked neighbors. If the point has marked neighbors we can begin to build a model for this change. We add the incident faces of the original point and recursively visit the marked neighbors, adding their incident faces to the current change model until we can no longer reach any unvisited vertices. We can then store the extracted model and proceed to select the next marked point that has yet to be visited and continue this process until no unvisited marked points remain.

The output of this process is a collection of 3D models

representing the changes between the two scans. We can now perform yet another type of filtering by computing the sum of the 2D area of the triangles in each model. To conserve storage, change models below a specified area threshold can be discarded at this step. This method is suited for dismissing small legitimate changes, such as trucks on a highway, when the objective is to find more significant developments, such as new construction. Change models also retain a count of the marked change points they contained, allowing for filtering by that criteria as well, which is advantageous in some cases (such as removing small trees).

4. Presentation and Interaction

Our interactive application allows for a rich visual exploration of the output data as opposed to more traditional output formats (images or digital elevation maps) for this type of result. We provide the basic behaviors and capabilities of common existing GIS / 3D terrain applications (such as Google Earth or ArcGIS) as well as visualizations that allow the user to analyze the results in a number of different views. These include a main 3D view of the terrain with highlighted changes, a sub-window showing a zoomed in, inspection view of any user selected change objects, and a semantic filtering sub-window utilizing a heat map. In the heat map, the x and y axis correspond to the changed models' footprint area and height, and the intensities of the cells indicate the relative number of changed models matching the cells' area and height ranges. In addition to the semantic filtering, through the use of manual area and change-point-count thresholds, the user can permanently filter out spurious changes they do not care about.

Existing available GIS data, such as streets, highways, known building footprints, political boundaries, etc. provide ease of navigation, meta-data for change models (such as estimated street address) and show the potential for integration into a larger GIS environment. We draw the vector data directly onto the reference surface models and other 3D structures using the technique presented in [SK07].

Changes are presented to the user in the intuitive color scheme of green representing new construction and growth, and red representing demolitions, excavations, tree-clearing, etc. The user can control the temporal aspect of the visualization by switching between discrete years, where the terrain reflects only that year's scan with only the appropriate new construction or demolition that has occurred displayed overtop. A hybrid option is provided that presents data from multiple years; new constructions are shown overtop of the terrain from the previous scan in which they did not exist, and demolitions are shown overtop terrain from the latter scan in which the result can be seen. In this way, a change is always presented against its counterpart for reference. See Figure 4 for an example view of hybrid mode on an dense urban area.

Due to the large amounts of data and the inherently




Figure 4: Shown here is a new residential development at increasingly further camera distances. The individual models gradually begin to glow and are ultimately replaced by semi-transparent 'splats' which maintain near-constant screen-size regardless of distance.

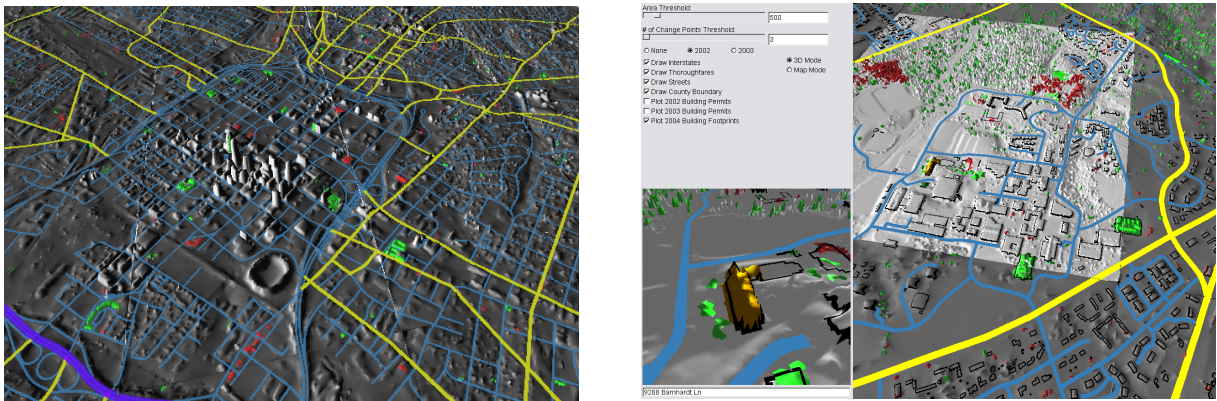


Figure 1: A view of the downtown area in hybrid mode. Change models of new construction are rendered in green and change models of buildings that have been demolished are rendered in red.

Figure 2: Changes to the university campus between 2002 and 2003. Green models represent new construction while red models represent demolition or tree clearing. 2004 tax database footprints are drawn in black. The selected building was currently under construction at the time, only filling half the future 2004 footprint.

broad range of scales at which the dataset can be examined (from individual houses at street level to county or statewide views,) it is necessary to implement some form of level-of-detail. It is here that we can exploit the level-of-detail concept to not only control the amount of data we render, but to change what it is we are rendering in order to appropriately abstract the data into a form that makes sense, and is useful to the user at that zoom-level.

4.1. Levels of Detail/Abstraction

When the user is zoomed in to a view that is significantly close to the terrain, the system shows the unaltered change models at full detail, as this allows for immediate inspection and interpretation of the change detection results. However, as the user zooms out, smaller change models quickly become little more than a pixel or disappear altogether.

We counteract this problem with the introduction of a second level-of detail/abstraction. As individual models recede away from the camera, they gradually begin to glow and are ultimately replaced by a semi-transparent 'splat' which is scaled to maintain near-constant screen-size regardless of distance. These splats seamlessly fade in and take over while the individual model detail level fades out.

The splats do not simply allow us to see individual changes/buildings beyond the point at which their models would disappear. They also provide an amalgamating behavior in which collections of individual changes cooperate to form larger, more significant glyphs in the visualization. This is shown in Figure 4 where in a series of increasingly distant views, a new residential housing development seamlessly transforms from a few rows of individual buildings

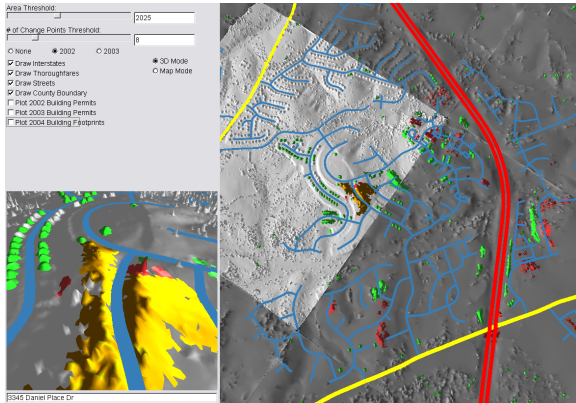


Figure 3: Examining a new residential development under construction. The selected (orange) volume shown in the inset view is an area where the earth has been moved and regraded to allow the level construction of a new street. We can also see the build up of earthen walls (green) on either side of the large interstate highway (red) in the lower-right region, presumably as sonic barriers against traffic noise.

into a single 'blob' representing the entire development as a whole. This blob is visible long after the individual models would have been too small to render as single pixels, aggregating and preserving the development.

The third and final level of abstraction is displayed when the user has zoomed out to a distance where depicting individual changes and even amalgamated groups in splat form no longer makes sense due to issues of overlap and extreme clutter. At this level, the user is presented instead with urban legibility regions. These regions are based on the level-of-detail clusters generated by [CBZ*06] which delineate sections of the city based on aspects of urban legibility, such as paths, districts, nodes, and other perceptual qualities. Originating from a LOD solution to city-viewing, these clusters are naturally useful to display city-wide data in clustered form at appropriate detail for a wide range of distances, as demonstrated with census data in [CWK*07]. This ability is clearly shown in Figure 7 where a significantly wide range of zoom levels result in continuously legible images of the concentration of changes present. The shading of these regions can be determined by several criteria, the simplest being the comparative ratios of local to global change footprint areas, number of change points, or number of individual changes.

4.2. Heat Map and Semantic Filtering

To allow a higher-level exploration of the results of the change detection we provide a semantic filtering interface utilizing a heat map. The heat map presents the distribution of the physical dimensions of the change models. The X-axis corresponds to 2D projected area (footprint size),

while the Y-axis corresponds to the maximum height of the change model above the surrounding terrain. The axes are divided into bins, forming a matrix of cells. Change models are enumerated into these cells, and the intensity of each cell indicates the relative number of change models meeting that cell's size ranges. A logarithmic scale is the default mapping for intensity, as it allows better observation of the less populous cells away from the main concentration areas, however linear and exponential scales are also provided in case changes follow uncommon distributions, or inspection deeper into the main concentration is desired.

Different types of physical objects often have a particular signature in the heat map. For instance, individual trees cut down will produce change models that are relatively tall with very small footprint areas, while grading of dirt at a construction site will tend to produce change models that are relatively flat (low heights), but cover very large areas. We can define these characteristics with sets of simple conditionals and bounding functions (relating growth in height to growth in area or vice versa) in order to delineate continuous regions in the heat map. These regions can then be associated with a particular group of physical changes.

We provide predefined functions that create regions in the heat map for trees, residential structures, commercial structures, and grading of earth. When the user selects one of these filter categories, the associated cells in the heat map are stippled with a unique color to indicate they are actively filtering the changes presented in the other 3D views. See Figure 5 for an illustration of the shapes these regions take on in the heat map.

The user can also define their own custom filters by selecting regions of cells in the heat map. This is advantageous when looking for changes that are not as simply defined as the provided categories, or to modify existing filters to better suit a particular dataset. Filters are stackable to permit more complex filtering.

An example of the usefulness of the custom semantic filtering is shown in Figure 8. Here the user opts to view only destruction changes, and creates a custom filter on the heat map to show only changes above a certain area and within a certain height range. The resulting 3D view allows the user to easily extract the red change models showing both the deforestation (left and center) due to clear cutting for construction, and the volumes of rock removed at the granite quarry (upper right).

Because the heat map provides a visual representation of the statistical distribution of the changes, it can also be used for a more general understanding of the types of changes across the entire dataset or in a sub-region. As shown in Figure 6 we can compare the heat map for the edge region of the city, which is under transformation from rural to suburban, to the more established urban and suburban areas closer to the city center. We see the wider distribution of changes, especially in the height dimension, for the city center region.

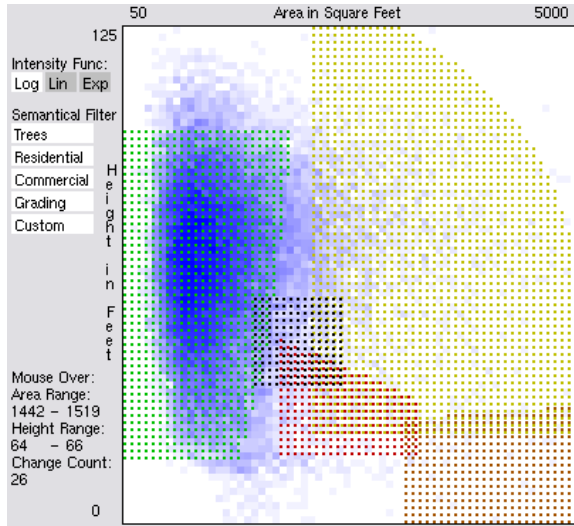


Figure 5: Shown here are the regions selected in the heat map generated by various semantic filters. The "trees" filter produces the green region, "residential" - orange, "commercial" - yellow, "grading" - brown, and a user defined region in black.

A possible explanation for this is that in the more established regions, structures tend to be larger and taller, and vegetation is more closely controlled.

5. Applications

Besides the stated application of analyzing the growth of a urban region based on detected new construction, there are other uses of this application that are being developed. The most promising is the detection of unreported or inaccurately reported construction. By comparing the changes extracted from the LIDAR scans through our method with existing real estate tax databases, we can look for buildings that do not exist in tax records, or those that are inaccurately recorded. Here, at the region level, instead of clustering changes to show overall trends, we would instead seek to highlight outliers and mismatches. We are seeking to collaborate with such enforcement agencies to evaluate the benefits of such a system over current methods of real estate tax enforcement and assessment, such as the use of aerial photography and in situ inspections.

6. Limitations

We did not rely on (nor did we possess a complete set of) any additional information besides the point locations of raw LIDAR returns. However, additional sources of information such as LIDAR return intensities or aerial imagery could be used to help classify the detected changes based on the

reflectance characteristics of their surfaces (concrete, water, vegetation, etc). When our system was presented to a group of GIS users, this was one of the most requested modifications.

While realistic rendering of the terrain is not of concern in this application, the method used for storage, retrieval, and display of the reference surface models is still quite primitive. This manifests itself in the seam artifacts present (roughly every square mile) between each reference surface model, a result of triangulations of points on each side of an arbitrary division line, with no triangles formed across said line. Clearly there is the need for the integration of a more advanced terrain engine to elegantly manage the storage and display of the reference models as a single entity. Hopefully this issue will be addressed by integration of the successful aspects of the visualization into an existing 3D terrain or GIS package.

There is much potential for the addition of algorithms such as automated target classification and pattern matching to further the abilities of the semantic filtering functions. This would be especially helpful in real GIS systems where the end users may know exactly what changes they want to ignore or expose, but are incapable of defining the specific parameters themselves.

7. Conclusions

In line with the mission statement of the USGS Center for LIDAR Information Coordination and Knowledge, we have developed a system that compares LIDAR data from different times and visualizes the differences in an intuitive and legible manner. This system, combined with building footprints from tax records and road information, provides a powerful tool for visualizing and identifying constructions and demolitions of buildings, grading of terrain, as well as increases and decreases in forestation.

Instead of using traditional methods for creating regularly gridded DEMs, our system maintains a high degree of accuracy by utilizing LIDAR point collections of any resolution and consistency. More importantly, our emphasis on preserving the accuracy in the LIDAR data is carried throughout the preprocessing and analysis processes in which any error introduced by LIDAR scanning is taken into account.

The presentation of the resulting LIDAR data uses a three-tiered LOD technique that is view dependent. This technique ensures that constructions or demolitions of buildings are clearly visible from all view distances and angles and allows the user to see overviews as well as detailed information at all times.

Lastly, by providing a heat map visualization of the distribution of changes and enabling semantic filtering, we allow the user to better explore the 3D results, and understand the patterns contained within, in ways beyond the capabilities of traditional 2D image results.




Figure 6: On the left is the distribution of changes for a region on the edge of the city, under transition from rural to suburban. In the center is the distribution of changes for a region in the center of the city, which is an established urban and suburban area. On the right is the difference of the two, showing the wider distribution of changes, especially in the height dimension, in the center city region.




Figure 7: Shown here are the urban legibility regions at increasingly distant camera locations.

References

- [AvH04] ABELLO J., VAN HAM F.: Matrix zoom: A visual interface to semi-external graphs. *infovis 00* (2004), 183–190.
- [BCWR07] BUTKIEWICZ T., CHANG R., WARTELL Z.,

RIBARSKY W.: Analyzing sampled terrain volumetrically with regard to error and geologic variation. Erbacher R. F., Roberts J. C., Grohn M. T., Borner K., (Eds.), vol. 6495, SPIE, p. 64950O.

- [CBZ*06] CHANG R., BUTKIEWICZ T., ZIEMKIEWICZ

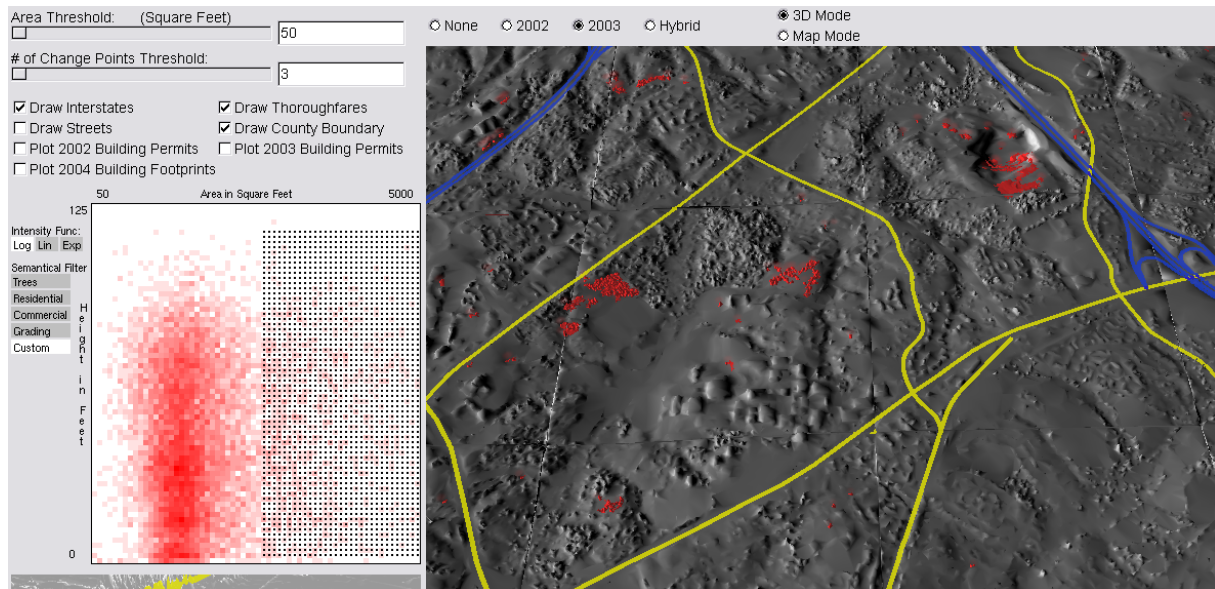


Figure 8: By opting to view only destruction changes, and creating a custom filter on the heat map to show only changes above a certain area and within a certain height range, we can easily see in the 3D view, the red change models showing both the deforestation (left and center) due to clear cutting for construction, and the volumes of rock removed at the granite quarry (upper right).

- C., WARTELL Z., POLLARD N., RIBARSKY W.: Hierarchical simplification of city models to maintain urban legibility. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Sketches* (New York, NY, USA, 2006), ACM Press, p. 130.
- [CWK*07] CHANG R., WESSEL G., KOSARA R., SAUDA E., RIBARSKY W.: Legible cities: Focus-dependent multi-resolution visualization of urban relationships. In *IEEE Transactions on Visualization and Computer Graphics (TVCG) InfoVis* (2007), pp. 1169–75.
- [GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1997), ACM Press/Addison-Wesley Publishing Co., pp. 209–216.
- [HMI99] HIROSHI MURAKAMI KATSUTO NAKAGAWA H. H. T. S., IWANAMI E.: Change detection of buildings using an airborne laser scanner. vol. 54 Issues 2-3, pp. 148–152.
- [JSS02] JINWOOK SEO; SHNEIDERMAN B.: Interactively exploring hierarchical clustering results [gene identification]. *Computer* 35, 7 (Jul 2002), 80–86.
- [SK07] SCHNEIDER M., KLEIN R.: Efficient and accurate rendering of vector data on virtual landscapes. *Journal of WSCG* 15, 1-3 (January 2007).
- [Spo05] SPOTFIRE: Decision site for functional genomics. <http://www.spotfire.com/>, 2005.
- [USG07] USGS.: The united states geological survey center for lidar information coordination and knowledge. <http://lidar.cr.usgs.gov/> (2007).
- [VMY04] VU T. T., MATSUOKA M., YAMAZAKI F.: Lidar-based change detection of buildings in dense urban areas. vol. 5, pp. 3413–3416.
- [Yvi07] YVINEC M.: 2d triangulations. In *CGAL User and Reference Manual*, Board C. E., (Ed.), 3.3 ed. 2007.