CHANGE DETECTION FOR TOPOGRAPHIC MAPPING USING THREE-DIMENSIONAL DATA STRUCTURES

D. M. Barber[a], D. Holland[b], J. P. Mills[c]

[a]DSTL, Porton Down, Salisbury, SP4 OJQ, UK – dbarber@dstl.gov.uk
[b]Ordnance Survey, Research Labs, Southampton, SO16 4GU, UK - david.holland@ordnancesurvey.co.uk
[c]Newcastle University, School of Civil Engineering and Geosciences, Newcastle upon Tyne, NE1 7RU, UK - j.p.mills@ncl.ac.uk

WG IV/9 – Mapping from high resolution data

KEYWORDS: Change Detection, Point Clouds, Lidar, Photogrammetry, Topographic Mapping

ABSTRACT:

Identifying significant changes to our urban areas is a prerequisite for accurate topographic mapping. This paper presents an approach based on octree data structures to identify change between two sets of point cloud data. The aim of the study was to establish if a method based on the comparison of point clouds could be used to detect simply where topographic change had occurred. As many of the changes that could occur are likely to result in a change of real world geometry (for example the construction or demolition of a building) the use of geometric data – rather than imagery alone as in other studies – is justified. Additionally, it means input data can be supplied directly from airborne lidar systems, or from aerial imagery using digital photogrammetric workstations which still remain the most commonly used apparatus for national mapping activities.Octrees are data structures that allow the partitioning of three-dimensional data into increasingly smaller units of space, using predefined criteria to control the level of subdivision (in this case a limit on the number of points in a node; and/or the total level of subdivision). Octrees have previously been used in applications where efficient searching and inspection of large volumes of three-dimensional data is required, such as in the rendering of computer graphics. By defining these structures, large data volumes of non-connected data (such as point clouds) can be efficiently managed and quickly compared with similar datasets collected at different epochs. In the study presented here, one approach compares entire octrees for differences between their structures, while a second approach compares individual data points to data contained in a reference octree.Two UK test areas form the basis of the study. Area 1 is the site of Heathrow Airport's new Terminal Five which has seen significant development over the last five years. Area 2 is an urban/peri-urban area of Bournemouth consisting of both commercial and residential properties. In both cases, multi epoch data was provided by Ordnance Survey allowing point cloud data to be generated from imagery collected by an Intergraph DMC digital airborne sensor. High resolution photogrammetric processing was undertaken using BAE's Socet Set and point cloud pre-processing using Terrasolid's TerraScan suite. While it was possible to recognize significant and pre-identified areas of change using the methodology, a large number of false identifications were also observed, making it difficult to interpret the results without prior knowledge. The lack of success can partially be attributed to the quality of the input data. Slight variations in the point clouds, perhaps arising from poor image correlation during surface extraction, led to subtle variations in the structure of the octree and thus in the changes identified.

1. INTRODUCTION

1.1 Opportunities for change detection

With the rapid expansion of our urban areas and the increasing desire to redevelop inner city brown field sites, ensuring that changes are quickly identified, and then recorded in national mapping databases, is increasingly important. At present within Ordnance Survey, Great Britain's National Mapping Agency, map revision is undertaken at two levels (Ordnance Survey, 2006). These comprise:

- a process of continuous revision, mainly directed at urban developments and aiming to capture new features within six months of their being completed;
- a process of cyclic revision for changes in the natural environment which is undertaken at intervals of between 2 and 10 years, depending on the nature of the area and the expected amount of change (Holland, 2008).

Typically, continuous revision is achieved through links to local planning processes and local intelligence. A surveyor may need to visit the site to undertake survey work. Cyclic revision, however, is most economically achieved through airborne imagery collected at fixed intervals, with this imagery being compared manually against the existing database and significant changes being specified for revision.

The introduction of digital aerial cameras has seen a streamlining of data capture flow lines resulting in two possible outcomes relevant to map revision. The first is economic: the cost of capturing and processing aerial imagery could potentially be reduced. It is beyond the topic of this paper to discuss this aspect, other than to recognise that this could result in a second outcome. That is the option to capture imagery more often and/or at a larger scale. Clearly this could improve the chances of small, yet significant changes being identified more quickly than before.

Also of relevance to map revision is the introduction of airborne and, for urban areas, ground based lidar systems (Barber et al., 2008). Over the past ten years such systems have seen rapid development and they are now commonplace in generating digital surface and urban city models. They potentially offer datasets with a much denser set of observations, and without outlier measurements relating to miss-correlation between image points, especially along depth discontinuities such as building edges. In applications seeking to determine changes between two or more epochs of data over an urban or peri-urban area, where the demolition or construction of buildings/structures is of interest, lidar data, in contrast with aerial photography, is potentially able to provide a point cloud dataset requiring significantly less manual editing than that produced by conventional photogrammetric processes.

However, while the efficiency of data capture might be improved, it is also necessary to improve the methodology relating to change detection in order to cope with this potential for an increase of data. Techniques operating on the point clouds themselves also have the advantage of being able to be applied to point clouds from any source, making such an approach more versatile. Efficient change detection methodologies, coupled with increasingly frequent data capture could lead to a merging of continuous and cyclic revision processes. Even without an increase in the quantity or scale of data it would be of interest to provide robust routines that can identify areas of interest automatically.

1.2 Point cloud based change detection

As high resolution three-dimensional samples of the real world, point clouds have potential for change detection at a variety of scales. Change detection on the point clouds themselves would help to increase the speed of the change detection process and improve the management of the datasets (both key parameters in determining the cost efficiency of the change detection process). There are, however, a number of obstacles to the demonstration of a robust change detection process based solely on a three-dimensional point cloud:

- ‘best guess’ point correspondence must be determined, even though no direct point to point correspondence can be assumed;
- known topology cannot be assumed as point clouds may originate from multiple sources;
- occlusions during data collection may need to be considered;
- differences between two point clouds do not imply an information change (e.g. the addition of a new building);
- large datasets require memory-efficient data handling routines.

The development of change detection procedures that operate on point cloud datasets would allow national mapping agencies to take advantage of improved data flow lines. As a result they would be able to provide change detection at a higher temporal and spatial scale, or simply be able to improve the economics of map revision.

This paper will outline work to establish a flow line to identify change between two point cloud datasets, based on the comparison of octrees, i.e. three-dimensional data structures representing the point cloud. By basing any comparison on these octrees, it is hoped to limit the amount of pre-processing required (for example the building of a digital surface model). The paper will start with a brief introduction to octrees. It will then present results of using the methodology on a variety of datasets. This will include datasets derived from digital aerial photography. Finally it will provide a discussion on the performance of the system and outline some conclusions.

2. METHODOLOGY

2.1 Octrees

An octree (Figure 1) is a recursive and regular subdivision of three-dimensional space originally used to optimise computer graphics rendering (Botsch at al., 2002). The bounding box for a three-dimensional geometry is divided into eight equal cubes, each of which are in turn divided into eight more cubes until a specified level of subdivision is achieved, or the cubes have less than a predefined number of data points within their bounds. At this point leaf nodes are formed. An octree representation of a point cloud provides an efficient represent of the spatial distribution of the 3D points. The tree structure assumes no previous knowledge of topology but allows efficient searching of the point cloud to determine possible correspondence.

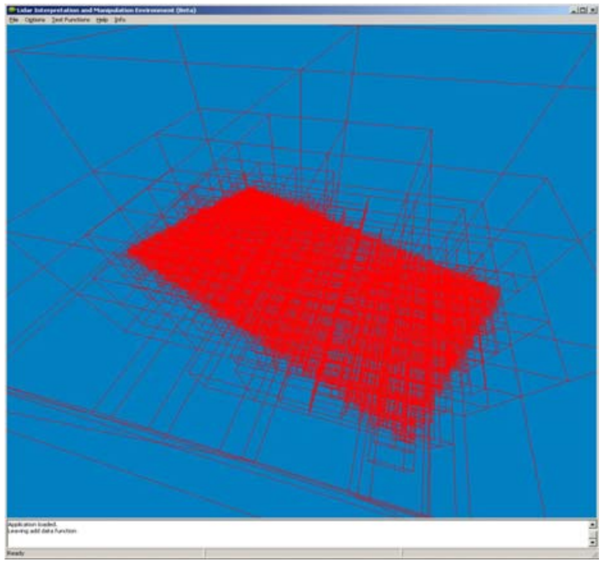

Figure 1 An octree generated from data used in the study

The comparison of point clouds against 3D CAD models is currently undertaken using point-to-mesh and mesh-to-mesh processes, often using algorithms based on the iterative closest point solution outlined by Besl and McKay (1992). This requires an existing reference model to be generated before change detection can take place. Girardeau-Montanut et al. (2005) outline a point-to-point comparison framework using octrees to optimise the performance of the change detection process, removing the need for a reference model; however this process is based on point clouds collected from similar positions and orientations and does not provide a general approach to the problem. Nevertheless, octrees clearly have the potential for facilitating the change detection directly between point clouds.

2.2 Node to node comparison

In this study the following methodology has been applied. Two octrees are generated based on the two point clouds which are to be compared. Each of these octrees share a common central node (described in the same geographic coordinate system) and the same controlling parameters, namely the maximum number of points allowed in each node and the maximum number of subdivisions allowed in the tree. As this study only deals with point clouds which are defined in a common coordinate system (for example the mapping system in use) one octree can be compared with the other. If the point cloud datasets were identical, the structure of each octree will be the same. Thus differences in the structure of each octree can be termed as change in the intervening time between epochs. Clearly, the cause of this change must be taken into account as such change could be related to the accuracy of datasets (for example errors in geo-referencing) or errors in the collection of the data (for example miss-correlation of images or multi-path errors in the return of a lidar signal). Also, each dataset must have a common ground sampling resolution for a comparison of this kind to be valid. For the purposes of this study this approach has been termed a node to node comparison.

2.3 Point to node comparison

The second approach in which the octree is used to aid the comparison between point datasets is its use to improve searching within a reference dataset. In this study this has been termed a point to node comparison, where a test point cloud searches for its closest corresponding data point in a pre-defined reference octree. Points with a corresponding point a large distance away could be considered change between epochs. This approach makes use of the spatial indexing of the octree to improve the speed at which this searching can be completed.

2.4 Advantages of using octrees

A number of advantages are anticipated from the use of octree data structures. Firstly, the approach is generic, in that it relies on point clouds from any source (for example ground based lidar, airborne lidar or airborne photogrammetry) providing, for a node to node comparison, the data resolution is the same. Secondly, once the process is established, it could be largely completed automatically limiting user interaction. Finally, the speed of the change detection can be optimised by limiting the searching times required. Note that for the purposes of this proof of concept study the efficiency of the routines has not been a priority, and thus the speed of the process has not been evaluated.

2.5 Test data

Two test areas were selected for use in the study. The first was around the area of Bournemouth on the south coast of Great Britain. The second was located around the ongoing development of Terminal Five at Heathrow Airport, London. Datasets from 2005 and 2006 were made available by Ordnance Survey. This included image datasets collected by the Ordnance Survey's Intergraph DMC. These datasets, delivered for this study as a pre-orientated photogrammetric dataset, were used to generate point cloud datasets using BAE systems Socet Set photogrammetric workstation. Gridded digital surface models were collected with a ground sample resolution of 1m using the standard Socet Set surface extraction and the NGATE surface extraction module. If such data was being used in an actual change detection flow line, it would be preferable not to have to manually edit the collected data – clearly in a change detection process it is assumed that the majority of data has already been collected and, when not required, the generation of new data is clearly uneconomic. Thus, in this study the surface data was used without any further editing.

Before octrees were generated, however, the extracted point cloud datasets were classified into ground, above ground and buildings using TerraSolid's TerraScan lidar processing system. A common classification routine was applied to both epochs to identify a ground surface, vegetation (or features above the ground) and buildings. This was applied to point clouds extracted using the standard and NGATE strategies.

Figure **2** and Figure 3 show the results of the classification on the Bournemouth test area.

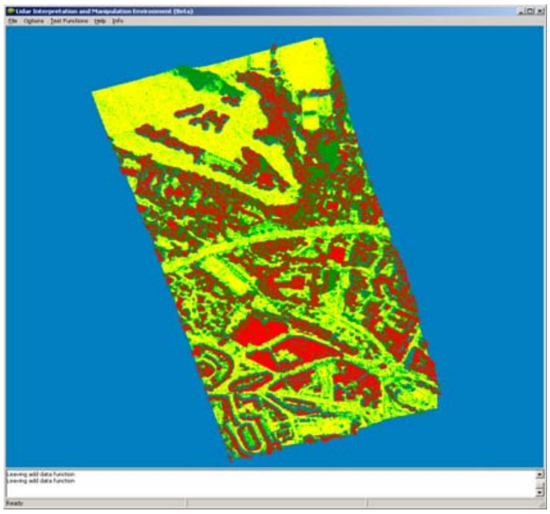

Figure 2 Bournemouth (2005) DSM collected by the NGATE Socet Set extraction module after classification with TerraScan.

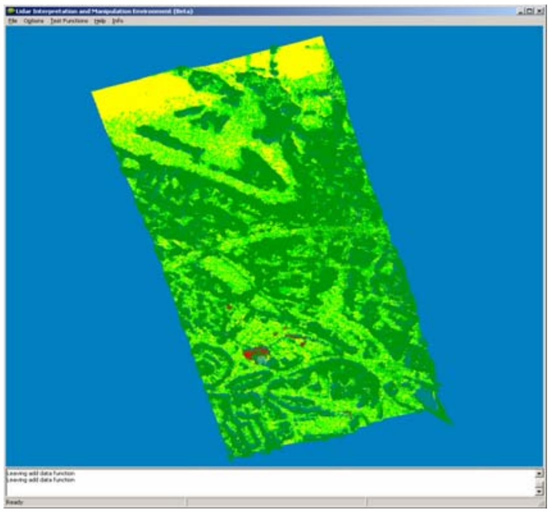

Figure 3 Bournemouth (2005) DSM collected by the standard Socet Set extraction module after classification with TerraScan.

It was noticeable that data collected using the NGATE system was more successfully classified at this stage (for the Bournemouth and Heathrow datasets), with buildings in particular being more clearly identifiable, where not obscured by vegetation. This indicated improved performance of the surface extraction around building edges compared to the standard extraction routine. In the case of the Bournemouth test area the standard DSM (Figure 3) contained a number of outlier errors, which resulted in a poorly classified ground surface, and ultimately a failure to identify any structures.

Following classification an octree was defined for each point cloud setting the maximum level of subdivision to 20 nodes and a maximum of 150 data points per node. These octrees were then used in node to node and point to node comparisons. In order to try and limit the number of changes identified due to changes in vegetation cover, an additional condition was placed on the comparison routine: only those points classified as buildings/structures in the TerraScan pre-processing should be considered in the comparison (except in the case of the Bournemouth dataset extracted using the standard Socet Set extraction module).

3. RESULTS

3.1.1 Bournemouth (node to node comparison)

The point clouds generated from the imagery over the city of Bournemouth from 2005 were compared with their corresponding dataset from 2006 using a node to node comparison. Results are given in Figure 4 and Figure 5.

Comparison of the NGATE datasets (Figure 4) highlights three major areas of change. Review of these areas against the image datasets reveals these changes to be: the redevelopment of an existing structure (1); the demolition of a building (2) and the construction of a new building (3). Other areas appear to be false identifications relating to errors in image correlation (at building edges) or vegetation growth that has been miss-classified at the point cloud classification stage.

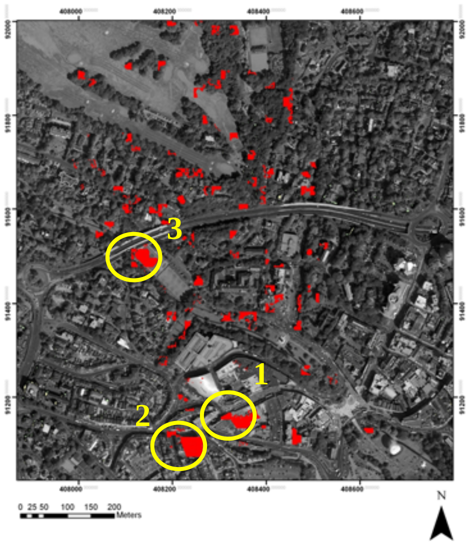

Figure 4 Results of the comparison between the 2005 and 2006 Bournemouth dataset from the NGATE module.

In the case of the standard Socet Set DSM (Figure 5) a greater level of variation between the datasets is evident. However, in most cases it is not conclusive whether this is due to differences in vegetation, surface extraction or due to actual real world changes. In a few cases however, it is possible to identify significant (large clusters of change) that do relate to real world changes (marked as 2, 3 and 4 in Figure 5): namely the demolition of a building (2), the construction of a new building (3) and a hot air balloon used as a viewing platform in 2005 and not present in 2006 (4).

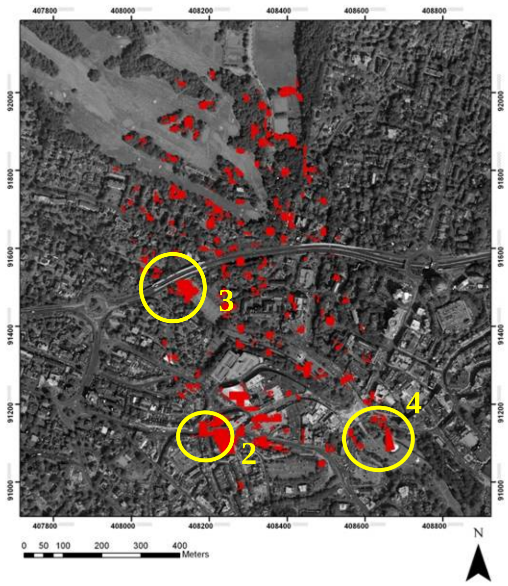

Figure 5 Results of the comparison between the 2005 and 2006 Bournemouth dataset from the standard Socet Set DSM generation.

Figure 6 gives enlarged examples of the real world changes found in the Bournemouth area and identified as site 1 in Figure **4**, and site 2 found in both Figure **4** and Figure 5.

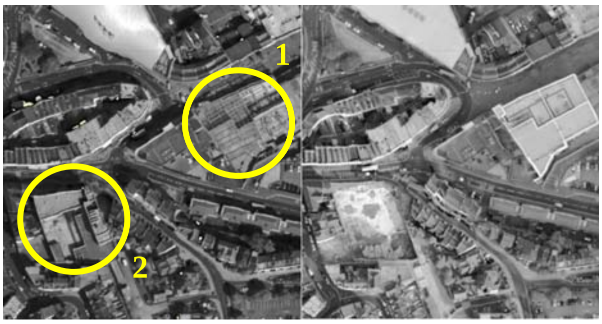

Figure 6 Example real world changes in Bournemouth dataset (left 2005, right 2006).

3.1.2 Heathrow (node to node comparison)

The point clouds generated from the imagery over the area close to Heathrow's Terminal Five from 2005 were compared with their corresponding dataset from the 2006 data using a node to node comparison. The results are given in Figure 7 and Figure 8.

The significant real world changes in this test area are: the construction of a building (1); the demolition of a building (2); the erection of a new motorway gantry (3). Some of the changes are shown in Figure 9 and Figure 10. For both datasets the results do not mirror the success of the Bournemouth test site. In the dataset produced from the Socet Set NGATE module none of the three major changes are identified. In the standard DSM only the addition of the motorway gantry is successfully indentified.

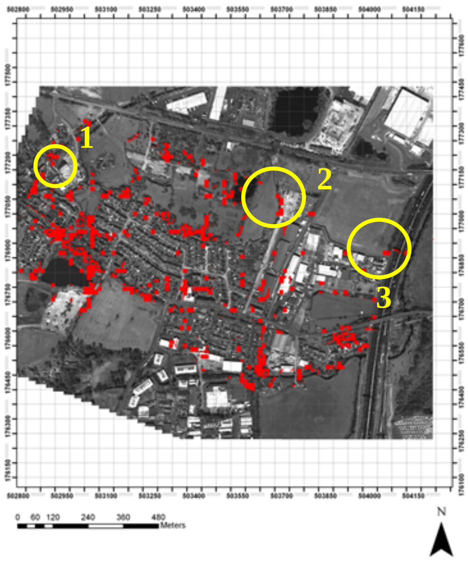

Figure 7 Results of the comparison between the 2005 and 2006 Heathrow dataset from the NGATE module.

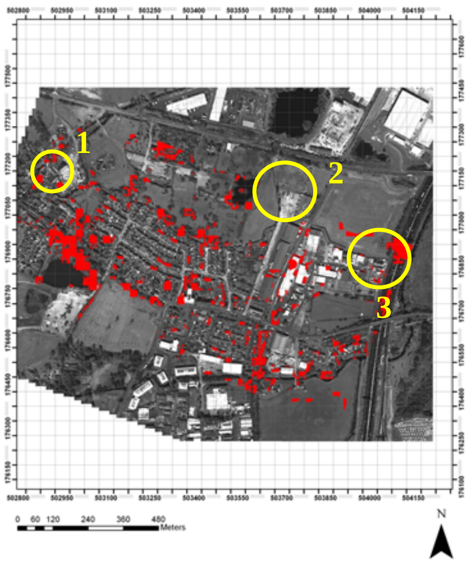

Figure 8 Results of the comparison between the 2005 and 2006 Heathrow dataset from the standard Socet Set DSM generation.

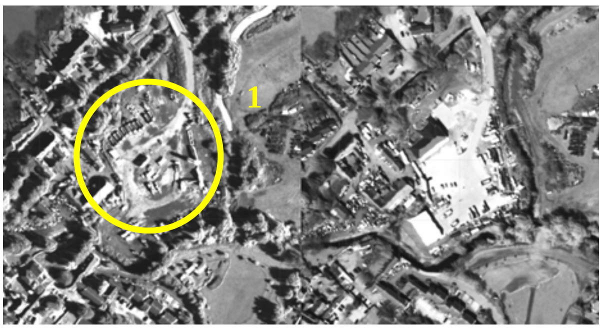

Figure 9 Example real world changes in Heathrow dataset (left 2005, right 2006).

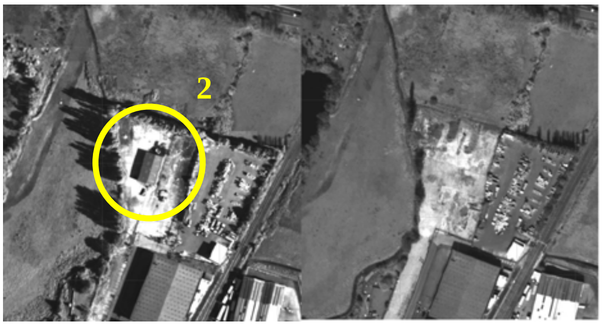

Figure 10 Example real world changes in Heathrow dataset (left 2005, right 2006).

3.2. Bournemouth/Heathrow (point to node comparison)

A second set of test results were run using the point to node routine described previously, comparing 2005 NGATE point cloud data against an octree derived from 2006 point data, and vice versa. The closest point was located and the elevation difference between the two points was computed. A threshold for the elevation difference was defined manually, consulting the data and determining the level at which an actual difference in elevation value indicated a possible real world change. Given the nature of the change being considered (demolition/construction of new buildings at least 5m in height) this was set at a threshold of 4.5m. Elevation differences greater than this value are displayed in Figure 11. Note: whereas in the node to node comparison octrees were generated with a maximum point per node value of 150 points, the octrees used in the point to node comparison limited the point per node value at 50, thereby deriving octrees of a higher resolution.

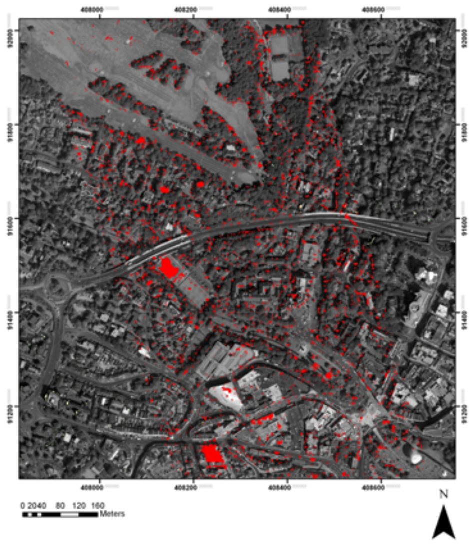

Figure 11 A point to node comparison highlighting change between 2005 and 2006 Bournemouth datasets.

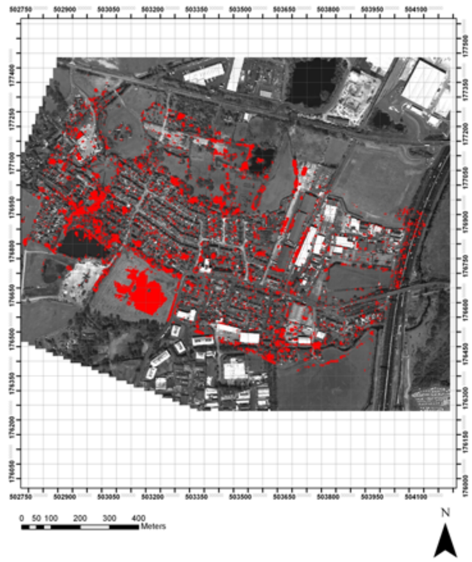

Figure 12 A point to node comparison highlighting change between 2005 and 2006 Bournemouth datasets.

4. DISCUSSION

The use of point cloud data derived from aerial photogrammetry has met with varying success in both of the test sites presented. A number of spurious points are visible in the results of each comparison. Mainly along significant edges in each scene, such as buildings and vegetation. Slight variations in the point clouds, perhaps arising from poor image correlation during surface extraction, led to subtle variations in the structure of the octree and thus in the changes identified. Especially along edges where image correlation may be poor. The quality of the input data is of prime importance when attempting to identify small changes. While it would be preferable to not have to edit digital surface models derived from imagery, in order to improve the success of the results editing of the surface models is required.

The proposed approach is perhaps better suited to datasets that are collected from a different perspective and with a larger number of points per m^2 with respect to the level of change anticipated. For example, this approach would be of interest in comparing datasets from ground based platforms, such as data collected by the StreetMapper mobile mapping system (Barber et al, 2008).

In this study, the approach has been to take a generic approach to the problem and identify differences between two sets of point cloud data, so a user could investigate any of the changes identified. If a user was only interested in mapping revision a more appropriate approach might be to compare newly collected data with the mapping database directly. However, by taking a generic approach to the problem the proposed solution may be of interest to a wider range of users. For example determining change between datasets where detailed existing mapping does not exist such as in aerial reconnaissance or in a situation where near real-time change detection is required.

This study has presented results from an urban/peri-urban area. The limiting factor in applying this type of algorithm operationally would be the frequency with which data could be collected. If, as for the continuous revision process, change needs to be recorded within six months of it occurring, data would need to be collected more frequently than this. While, as discussed earlier, this may be possible with increasingly more efficient data capture flow lines, at present this is unrealistic, so such an approach would only be of use for the cyclic revision process that requires data less frequently.

Further work includes the filtering of results so as to more clearly differentiate between significant changes, insignificant changes (such as vegetation growth) and changes relating to errors in the data collection process. Also, further testing is required using point clouds collected by airborne lidar sensors.

5. SUMMARY AND CONCLUSIONS

The use of an octree data structure for the detection of change between datasets has shown some limited success. Areas of real world change have been identified from real datasets, but a significant amount of spurious results have been generated, possibly relating to the quality of the input datasets which were generated from aerial photography. Improved results are anticipated when lidar derived point clouds are used as input data. Moreover, it is anticipated that the use of octree data structures would facilitate the future integration of airborne and ground based datasets.

ACKNOWLEDGMENTS

This investigation was conducted by Newcastle University whilst Dr Barber was a researcher in the School of Civil Engineering and Geosciences. The research was funded and supported by Ordnance Survey. BAE Systems kindly provided a licence for the NGATE module for use in the project.

REFERENCES

Barber, D. M., Mills, J. P. and Smith-Voysey, S. 2008. Geometric validation of a ground-based mobile laser scanning system. ISPRS Journal of Photogrammetry and Remote Sensing. 63(1): 128-141.

Besl, J. P. and McKay, N. D., 1992. A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence, 14(2): 239-256.

Botsch, M. Wirantanya, A. and Kobbelt, L., 2002. Efficient high quality rendering of point sampled geometry. Proceedings of Eurographics Workshop on Rendering.

Girardeau-Montaut, D., Roux, M., Marc, R. and Thibault, G., 2005. Change detection on point cloud data acquired with a ground based laser scanner. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, 36(3/W19).

Holland, D. 2008. Automating topographic change detection at Ordnance Survey. Geomatics World. 16(3) 30 – 32.

Ordnance Survey, 2006. Revision policy for basic scale projects. Ordnance Survey Information Paper. 3 pages.