Improving the update of geospatial information databases from imagery using semi-automated user-guidance techniques

D. A. Holland and W. Tompkinson

Research & Innovation, Ordnance Survey,
Romsey Road, Southampton, SO16 4GU, UK.
Telephone: +44 (0) 23 8079 2808/2867
Email: David.Holland@ordnancesurvey.co.uk
William.Tompkinson@ordnancesurvey.co.uk

Abstract

A common goal of mapping agencies is to optimise the collection of spatial information, by capturing ever-richer sets of high-quality data in as short a time as possible. One step towards achieving this goal is the development of automatic change detection methods. This paper concentrates on the requirements of such a system, and the development of a prototype which uses temporally-separated remotely-sensed images to guide a human photo-interpreter in the process of updating a geospatial database. The process requires one or more change detection procedures, several of which are described. These include the matching of polygons created by image segmentation; the detection of shadows cast by new buildings; and the analysis of areas which are recorded as "empty" in the geospatial database.

1 Introduction

In an increasingly consumer-led market, the demand for geospatial data and information continues to grow rapidly. In order to meet this demand, national mapping agencies require more efficient methods of spatial data capture; better spatial data management techniques; and faster, more user-targeted methods of data dissemination. This paper concentrates on the first of these – the rapid and efficient capture of spatial data. Many mapping agencies, including Ordnance Survey in Great Britain, have used remotely sensed imagery as a data source for several decades, but the processes of capturing spatial information from imagery have remained very manually-intensive. One long-term goal of a mapping agency is to develop or procure a system that can detect change and capture feature data from imagery, with little or no human intervention. While this goal of full automation remains unrealised, there are ways in which automatic techniques can be used to streamline at least some elements of the data capture system.

Almost all the large-scale spatial data products supplied by mapping agencies are based on attributed vector datasets. These datasets are derived from a variety of sources, to meet specifications that are defined at a regional or national level. Most of these products are stored in geospatial databases composed of points, lines and polygons which represent both natural features and man-made structures in the landscape. Aerial photography is used extensively by mapping agencies as a source of raw data from which such national geospatial datasets are derived and maintained. In most cases, the process of capturing any changes in the underlying spatial dataset relies solely on the ability of a photogrammetrist to manually identify and accurately capture the features that have changed.

One way to make the data capture process more efficient would be to automate the *detection* of change. Ideally, this would be followed by an automatic extraction of all the features for which a change has been detected. Although this subject has been an active area of research for many years (see, for example, Baltsavias, Gruen and Van Gool, 2001) the transfer of results from a research to a production environment is still quite limited. In this paper, we describe a prototype change detection system for use in a data collection production flowline, together with some examples of methods which may be used to aid the automated change detection processes embedded within such a system.

2 A Prototype Change Detection System

2.1 Rationale

Change detection, as performed by a mapping agency, often involves the use of an image pair – one image showing the current state of the region to be mapped; the second image showing the state at a previous point in time - often the time of the last update. In addition, many mapping agencies will have a geospatial database, in which topographic features are stored. One frequently encountered problem is that the real requirement is to determine the differences between the features in the geospatial database and the new image; rather than the simpler requirement of determining the differences between the two images. The problem lies in the fact that one is a vector model, determined by a specification designed to show cultural information; while the other is a raster model, dependent on the physical characteristics of the sensor and the conditions extant at the time of capture. One way to overcome this problem is to map one of the data sources onto the characteristics of the other – i.e. the vector data must be rasterized; or the imagery must be segmented into polygons. Either approach leads to some loss of integrity in the data, and it is generally better to use a like-for-like comparison. The method described here is similar to an image-to-image comparison; but the system first performs an image segmentation on each image, then compares the two sets of segments. This automatic stage is followed by a machine-assisted manual capture process.

Many researchers in the field of change detection are trying both to identify change, and to capture the exact outlines of those features which have changed. This has proved to be more difficult than one might expect. One reason for this is that, despite 20 years of research in digital photogrammetry and computer vision, no universal edge detector exists that can "both identify and track edges with sufficient success" (Agouris et al, 2000, p.2). Despite the lack of a complete solution, it is possible to develop a system that provides a marked improvement on a purely manual data collection system. Such a system takes into account the fact that much time is potentially wasted by human photo-interpreters, as they scan two or more images and identify those areas which require a closer look. Once this process is complete, the photo-interpreter must revisit each area of potential change, and employ different criteria to determine exactly what has changed, and which features must be edited, deleted or created in the geospatial database.

2.2 System Design

The prototype system, described in the following sections, is the focus of ongoing research investigating the extent to which the software interface of a data capture production system can be iteratively modified to increase the efficiency of geospatial database update. The underlying premise of this research is that, if the efficiency of such a production system is to be increased, there needs to be a strong appreciation of the fact that the system needs to be considered in terms of *both* 'software' and 'user' subsystems. As a result, the research concentrates upon improving the information flow between the software and user. In this paper we concentrate on the visual cues provided to the user by the system. To enable these concepts of improved information flow to be tested, this prototype represents a software interface that can have components inserted, removed or changed, as and when new methods for enhancing such information are developed.

The first incarnation of the system (described in more detail in Tompkinson *et al*, 2003) relies on the availability of imagery captured at two instances in time (usually separated by several years). Once the sensor distortions, chromatic differences and the different viewing geometries have been accounted for, the two images are compared and any differences are highlighted. In similar change detection systems, the process stops at this point, and it is left to the operator to search for all the areas which the system has identified as areas of potential change, to study these areas, and to capture the data manually. Figure 1 illustrates how such a change detection system might be organised according to a generic tiered structure commonly used in system design (e.g. as described in Bennet *et al.*, 1999).

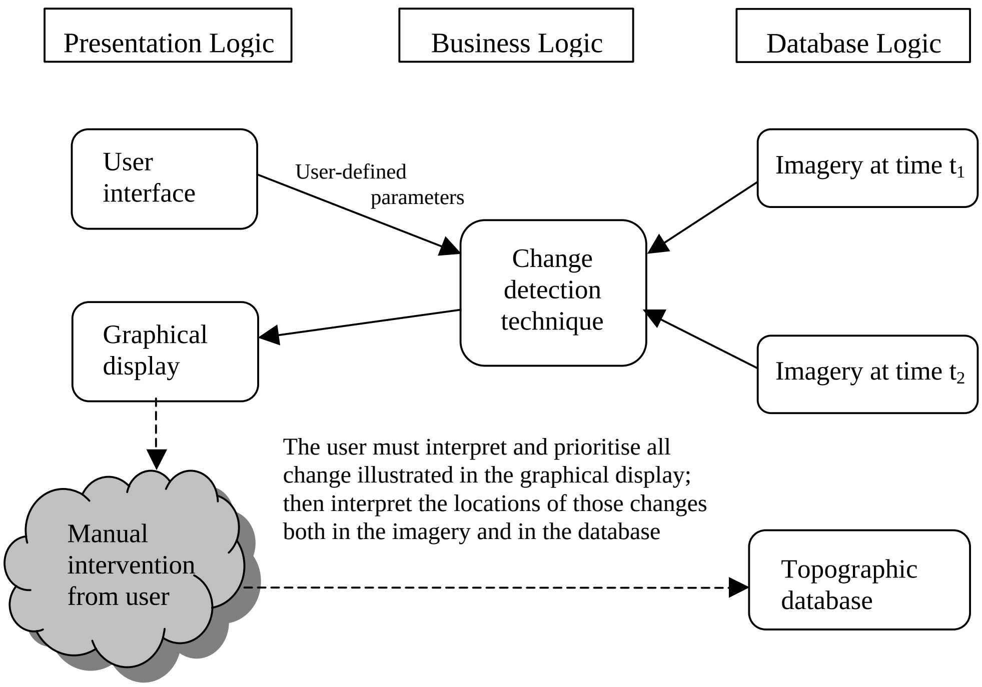

Figure 1: A tiered system design diagram illustrating a typical change detection system within a data capture software environment.

In the proposed system, the user is given more information and is automatically guided to the places where potential changes have been identified. In this system the following functions are executed:

- An automatic image-to-image change detection process is performed. The two images are first segmented into polygons, then these are compared, and matching polygons are identified. Differences between the matched polygons are used to indicate areas of potential change, using a technique similar to that described by Ware and Jones (1999). (Note: although image segmentation was used in the prototype, any change detection technique could be performed at this stage.)
- The system determines which of the areas of potential change are of greatest significance.
- Via a graphical user interface, the operator is guided to each area of change in turn. The user is presented with a view of the changed areas in both the images, together with a representation of the geospatial data currently held in the database.
- The user is also presented with a visual cue of the polygons which are most likely to have changed, as determined by the system.

By placing change detection within the context of the system in which it is applied, this research places an emphasis upon the application of computational methods to simulate the thought processes of a manual interpreter. In Figure 2, a diagram of the prototype system, it can be seen that emphasis is placed upon utilising cues of change detection to assess where change has taken place and presenting this information to the user. This information is then used to iteratively 'drive' the user to the location of the change. The user and the system work together to make the data collection process more efficient.

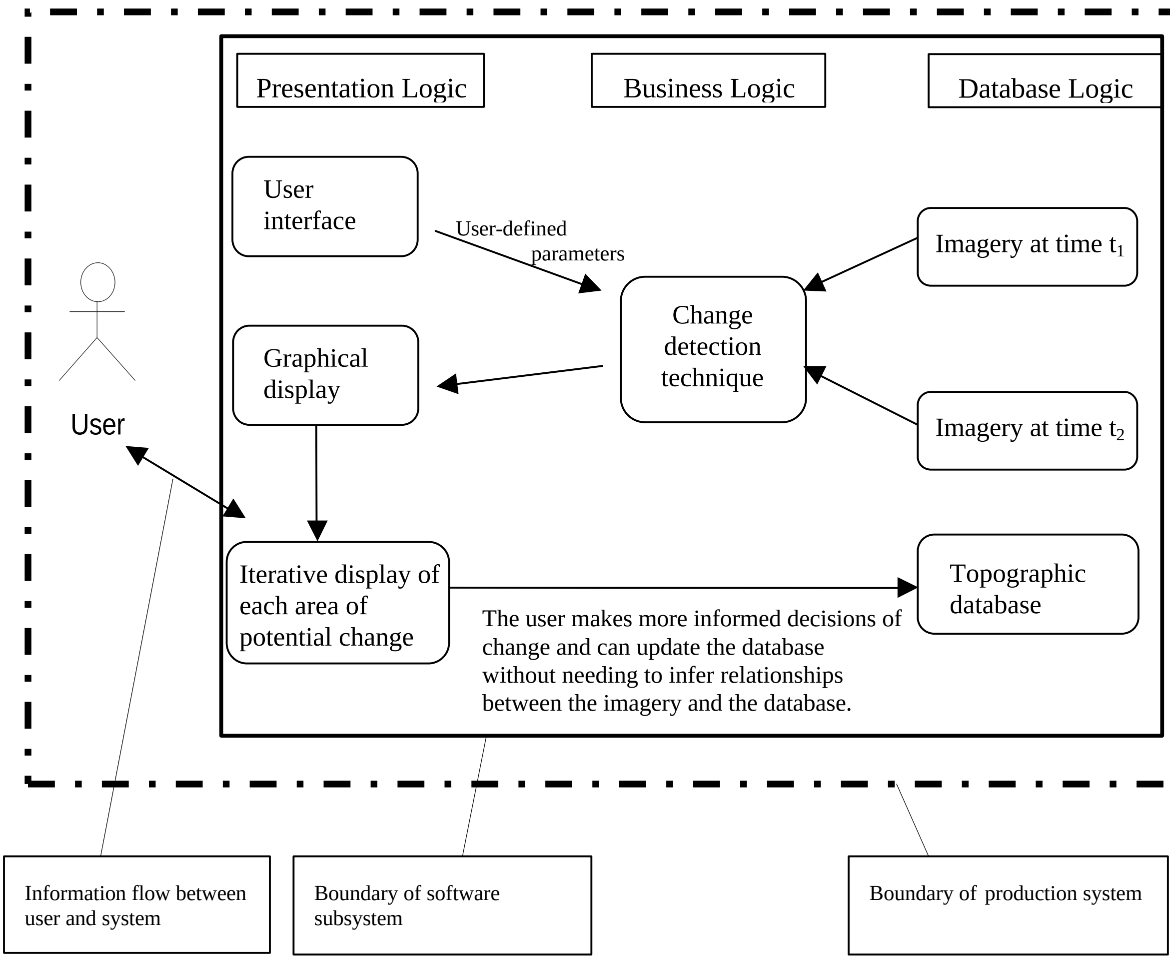

Figure 2: A tiered system design diagram illustrating the proposed change detection system within a data capture software environment.

2.3 Results of the prototype

The prototype system presently under development has been developed in Visual Basic using ESRI ArcObject™ technology, with a change detection technique implemented in C++. Since it produces vector segments from imagery in a consistent and efficient manner, the eCognition software package (from Definiens Imaging) is used to segment the images and generate the input vector data. The process flow in this system, from the input of the image pair, via the segmentation and the implementation of the change detection technique, through to driving the user to the location of change, is depicted in figure 3.

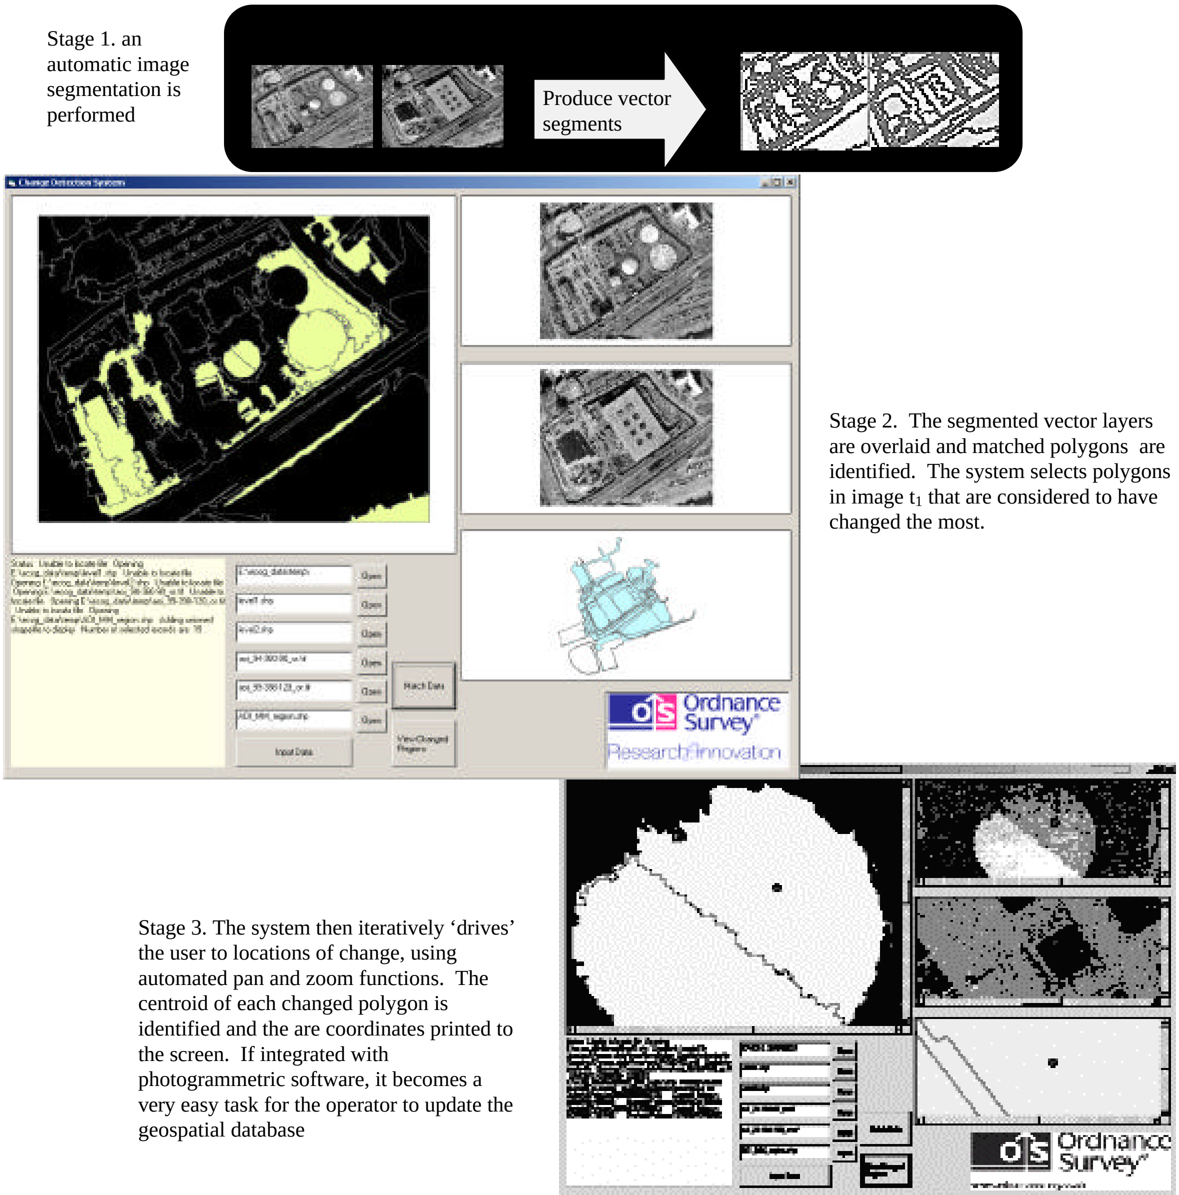

Figure 3: Process flowline through the prototype change detection system.

Initial tests have been performed using a small image subset, with dimensions of c.250m x c.210m. Initial demonstrations of this prototype, to those involved in the management of a large photogrammetric flowline, have been successful in conveying the concept of a system that enables more efficient interpretation of change locations through the automation of pan and zoom functions. It is envisaged that such an application has the potential to improve the speed of data capture from image blocks that cover ever larger areas of land. This becomes of greater importance if a mapping organization were to capture features from high resolution satellite imagery rather than aerial photography, since such imagery covers large areas without the need to produce aerial image mosaics.

Experimentation using larger (4km x 1km) areas of aerial orthophotographs is currently underway. It is intended that such trials will enable a better assessment of whether simply implementing the methods described in this system really do increase efficiency. These trials have indicated that limitations in the software platform and data structures used mean that the prototype can be slow to function.

Research is also underway that addresses the problems of edge effects, which currently present difficulty when processing sets of images taken from different orientations. These problems notably occur in the region between the convex hull and the envelope of a line of photographs, or in a region where photography is only available for one set on images. In such areas, polygons are often misclassified as having changed.

3 Other cues of change detection

3.1 Types of change detection techniques

The change detection technique in the prototype system described above uses a shape comparison between image segments. While this has produced good results in the area tested, this system could be improved by the incorporation of other change detection techniques. In a manual change detection process, the photogrammetrist will be skilled in identifying changes in the landscape from remotely sensed imagery. During manual comparisons between temporally-separated image pairs, or between an image and map data, changes may be identified in several ways. In one method, an operator may use indicators such as alterations in the shape of landscape features; in another, the operator may concentrate of the detection of artefacts within large homogeneous topographic objects. A new building development may be recognised by a change in texture over a large area of land; whilst a new wall may be identified through indirect evidence such as its shadow. If existing topographic data is available, this too can be used to supply extra information, either to the user or to the system, to help identify significant areas of change. Below we discuss some other techniques which may be used to indicate the incidence of change, and could be incorporated into future versions of our change detection system.

3.2 Using shadows to detect change

3.2.1 Shadows as indicators of topographic features

Shadows can be a good indicator of the presence of buildings and other objects, and they are usually easier to detect than the often-complex multi-textured objects which cast them (Irvin and McKeown, 1989). Until recently, the presence of shadows in satellite imagery was of very little interest, since the length of a shadow cast by a typical object (e.g. a suburban house) would often be smaller than the size of an image pixel. In contrast, the latest satellites from Space Imaging (IKONOS) and Digital Globe (QuickBird) provide imagery in which shadows are quite easy to detect. The QuickBird imagery used in our research (supplied by Eurimage) has pixels with a ground sample distance of 64cm. Even in the late morning, when the image of the study area was acquired, the shadows cast by a small house comprise 40-50 pixels (Figure 4).

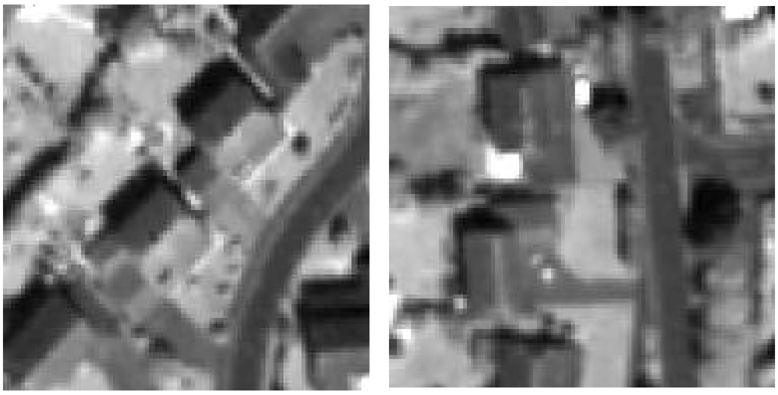

Figure 4: Examples of buildings and their shadows, visible in QuickBird satellite imagery (supplied by Eurimage).

3.2.2 Detecting shadows

One way of finding shadows is to simply threshold the image at a certain grey level, so that any pixels brighter than a given value are removed from the image, leaving only the areas of potential shadows. The most suitable threshold to use for this purpose depends on the characteristics of the image (e.g. what proportion of the image is in shadow and how much contrast is present in the image). After experimenting with various values, we chose a threshold calculated as the mean minus one standard deviation (these being the statistics of the pixel values in the image). For the test area of Christchurch in Dorset (Figure 5a), this results in the thresholded image shown in Figure 5b.

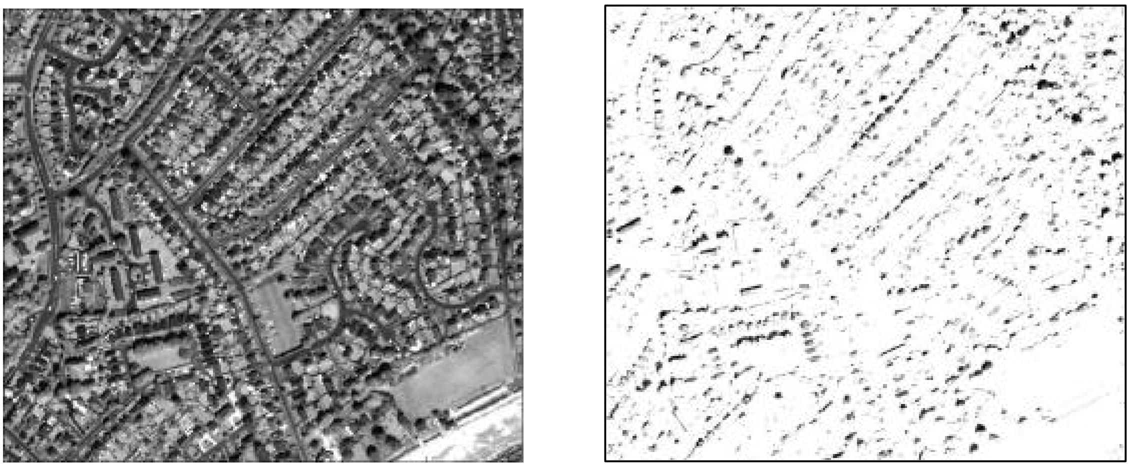

(a) (b)

Figure 5. Image of the study area (a) and the same image after thresholding (b)

After thresholding, the strongest shadows in the image were identified by passing a shadow detection filter over the image. The filter used was an ellipse, with its minor axis aligned with the direction of the sun at the time of image capture (23° East of South, as recorded in the image metadata) and its values weighted as the inverse distance from the central pixel, modified by the angular distance from the sun angle. The filter values were positive on the sunward side of the ellipse and negative on the leeward side. A diagram of the filter is shown in Figure 6, with the positive values in the filter represented by white to mid grey and the negative values represented by black to mid grey.

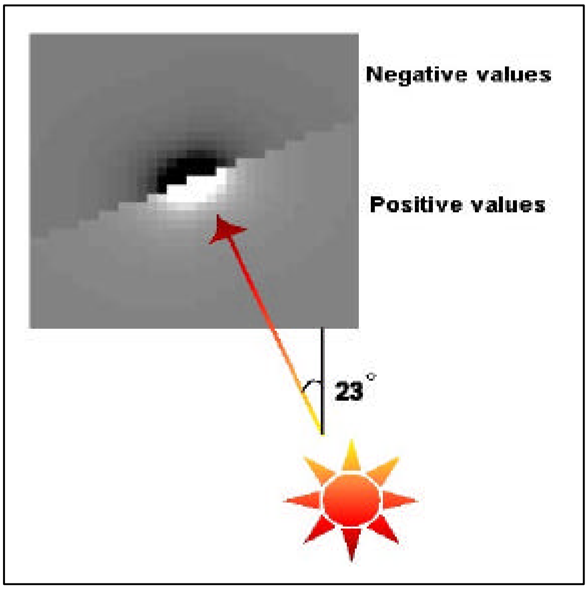

Figure 6. The shadow filter, with a sun azimuth of 23°.

The filtering process highlights the parts of the original image where the boundaries between light areas and dark areas are aligned with the predicted shadow angle (Figure 7a). These correspond to features such as buildings, walls, bushes and trees.

3.2.3 Detecting shadows cast by "new" features

To detect the features which are not in the topographic database, a set of "false shadows" was generated from the building footprints already held in the database. The buildings were given a nominal height of 6m (the average eave-height of a standard house in the study area), and we estimated the shadows cast by these buildings at the time and date on which the image was taken, using the sun azimuth and elevation angles recorded in the image metadata. Figure 7b shows the buildings and their predicted shadows.

The next step was to remove the areas covered by the buildings and their estimated shadows from the filtered image. Many very small areas of shadow were present in the resulting image, which were filtered out using simple dilation and erosion morphological operators. The final filtered image is shown in Figure 8.

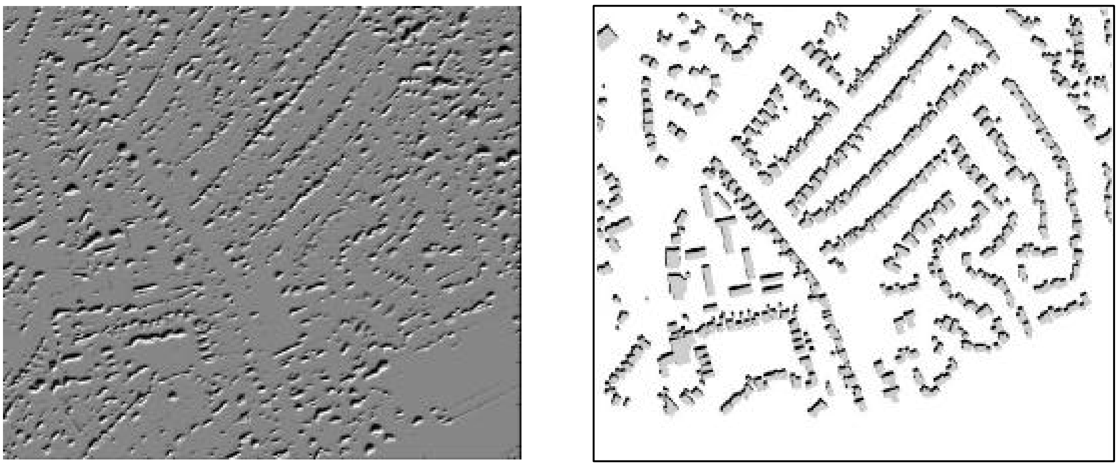

(a) (b)

Figure 7.

(a) The image after thresholding and filtering
(b) The buildings in the database, together with their predicted shadows

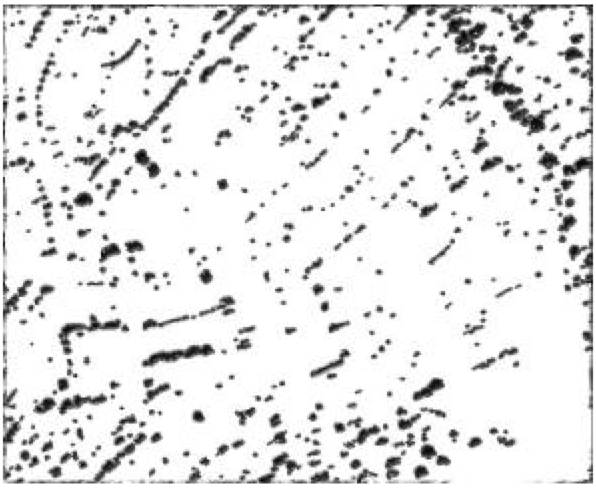

Figure 8. The filtered image after building+shadow removal, followed by dilation and erosion.

3.2.4 Results of the shadow detection

The processes described above leave an image of "unidentified shadows" (Figure 8) that could not be predicted from the existing spatial data. In an ideal world, these would indicate the presence of new features, which could then be captured. In reality most of the detected shadows were cast by bushes and trees, even in the urban area under consideration. Most of these "new" features, although real enough, are not part of a mapping agency's typical data specification. Since individual trees and bushes are not present in the topographic database, neither these features nor their shadows can be predicted.

In the Southern part of the study area, the topographic data was not used, so the final image in this area shows shadows cast by both buildings and by vegetation. A close-up view of part of this area is shown in Figure 9. From this example, it is clear that those shadows which are cast by buildings are very difficult to separate from those cast by trees. Without extra information, it is unlikely that an automatic technique would have much success in identifying new buildings from the shadow data. However, further analysis of the shadows of potential features could provide enough evidence to guide a semi-automatic process, such as the one described earlier in this paper.

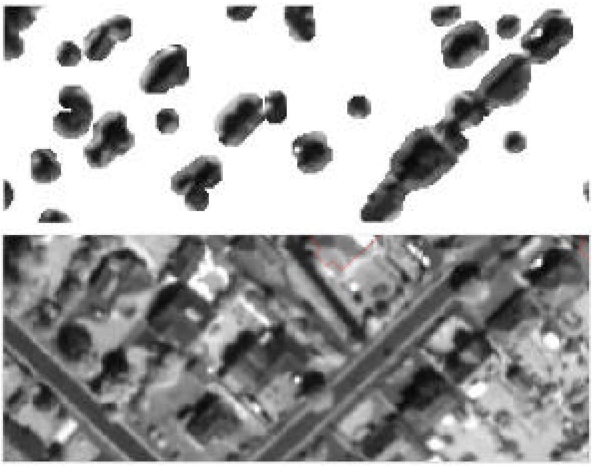

Figure 9. A comparison of the predicted shadows and the original image, showing shadows in the image cast by both trees and buildings.

3.3 Detecting change by examining "empty space"

A second indicator of change is currently under investigation. This research focuses upon utilising texture measurements to compare imagery with topographic data. This concentrates on areas in the image for which there are no existing features in the topographic database (i.e. the polygon is recorded as "general surface land"). In these areas, since there are no significant topographic features, the presence of significant image-texture differences will provide a further cue to the presence of changes to the features on the ground. However, as with the shadow detection process, initial results of this research show that trees in these areas severely restrict the effectiveness of the procedure. While one might expect that areas of "general surface land" would exhibit low image texture, we have found that image texture can often be very complex in these areas, usually due to the presence of trees and bushes. Paradoxically, it is often the case that features such as buildings in the image will exhibit less texture than the "general surface land" features. Buildings will show little texture difference within their large, flat areas of roofs (and also within vertical walls if the sensor is pointing off-nadir), while there will be large texture differences at the boundary between these features and their surroundings.

Further research will examine these areas more thoroughly, and will look at other cues which could be used to indicate change, such as the detection of straight lines and corners within areas recorded in the topographic database as featureless.

4 Summary

This paper presents a prototype change detection system that could be used to make the capture of spatial data from imagery a more efficient process. The prototype has been tested on sample images of a small area and has received favourable comments from personnel involved in the photogrammetric data capture area of a national mapping agency. In the initial prototype, image segmentation and polygon-matching techniques were used to determine the locations of change. Related research work has investigated the use of other cues of change, including shadow detection and the analysis of regions in the image which are “empty” in the geospatial database. These techniques could be used in later incarnations of the change detection system.

Current research concentrates on enhancing the prototype to take larger images; to overcome edge effects in the data processing; and to optimise the system to make it more useable in a production environment.

5 Conclusions

It has been shown that a technique that includes the user as an integral part of the change detection process can make the data capture process more efficient.

The use of shadows to detect change between features in a geospatial database and features in an image has met with limited success. While new buildings *can* be detected from their shadows, the presence of shadows from trees and other vegetation makes it almost impossible to determine whether or not the identified shadows indicate changes to the features required in the database. Further work is needed to determine whether other cues, such as shadows with straight edges, could be used to discriminate between building features and vegetation.

Initial research into the use of image texture to indicate the presence of features in areas recorded in a database as “empty space” has shown that many such areas are in fact rich in image texture, again mainly due to the presence of vegetation. Further research is needed in this area to determine whether the technique can be modified to produce more reliable indicators of change.

6 Acknowledgements

The authors wish to thank Eurimage for providing the QuickBird imagery used in this research.

7 References

AGOURIS, P., MOUNTRAKIS, G. AND STEFANIDIS, A., 2000, “Automated spatiotemporal change detection in digital aerial imagery”, *SPIE Proceedings Vol. 4054, Aerosense 2000 , Orlando, FL.*, pp. 2-12.

BALTSAVIAS, E.P., GRUEN, A., and VAN GOOL, L.(editors), 2001, Automatic extraction of Man-made Objects from Aerial and Space Images (III).

BENNETT, S., MCROBB, S., and FARMER, R., 1999, *Object-oriented systems analysis and design using UML* (Maidenhead: McGraw-Hill)

IRVIN, R. B., and MCKEOWN, D. M., 1989, Methods of Exploiting the Relationship Between Buildings and their Shadows in Aerial Imagery, *IEEE Trans. Sys. Man. Cyb.,* vol. 19, no. 6.

TOMPKINSON, W., SEAMAN, E., HOLLAND, D. and GREENWOOD J., 2003, An Automated Change Detection System: Making Change Detection Applicable To A Production Environment, *Proceedings of Multitemp 2003, Ispra.*

WARE, M. J., and JONES, C. B., 1999, "Error-Constrained Change Detection", Proceedings *of GeoComputation 1999*, *Fredericksburg*, Available at http://www.geovista.psu.edu/sites/geocomp99/Gc99/064/abs99-064.htm.