

DETECTING CHANGES TO TOPOGRAPHIC FEATURES USING HIGH RESOLUTION IMAGERY

D. A. Holland, C. Sanchez-Hernandez, C. Gladstone^a

^aResearch, Ordnance Survey, Romsey Road, Southampton, UK –
(David.Holland, Carolina.Hernandez, Catherine.Gladstone)^a@ordnancesurvey.co.uk

Commission VI, WG VI/4

KEY WORDS: Change Detection, Digital Mapping Camera, Cartography

ABSTRACT:

Detecting changes to topographic features is one of the major tasks of a national mapping agency. At Ordnance Survey, Britain's national mapping agency, the process of change detection has traditionally been a largely manual, labour-intensive task. One of the goals of the Research department is to develop automatic and semi-automatic change detection processes which could be developed into production systems. This paper describes the research undertaken, concentrating on the detection of changes to the built environment, specifically new and demolished buildings. The research has progressed from a set of potential methodologies; through research trials of different software packages and different change detection methods; to the adoption of one of the methods in a full production trial.

1. INTRODUCTION

1.1 Background

The detection of change is one of the most important aspects of the work of a mapping agency. In the case of Ordnance Survey, Britain's national mapping agency, this is still largely a manual process, relying on the observations of photogrammetrists and field surveyors and the notifications provided by external bodies. In the field, a network of approximately 300 surveyors is engaged in the process of updating the national topographic database, used in the production of the flagship OS MasterMap® large scale topographic data product. Part of the field surveyor's task is to note any changes to the natural or the built environment, in order to mark these changes for subsequent data capture. In addition to this, local authority planning departments provide information on planning applications which may affect buildings and other urban features. Major house builders inform Ordnance Survey of new developments, while a further source of intelligence on changes in the landscape is provided by a commercial survey organisation, which supplies information on changes which will affect Ordnance Survey's products.

The final source of changes is a group of image interpreters in the Photogrammetric Survey department, whose job is to scrutinise orthorectified digital aerial photography and mark up any changes. These changes are subsequently captured in stereo using digital photogrammetric workstations. In spite of all the other sources of information, this last phase uncovers many changes which have, up to that point, gone un-noticed. Manually scanning through the images is a laborious task, and one which requires concentration and a disciplined methodology. In an investigation carried out by the Ordnance Survey Research last year it was found that when scanning an image from top left to bottom right, constantly zooming in and out, it can be difficult to keep track of which areas of the image have already been looked at. In some cases, a tendency was found for the image interpreters to gradually migrate to one side

of the image, hence missing out one corner of the image and therefore miss any changes in that corner.

1.2 Automating change detection

In order to help the data collection process, automation of the workflow could be achieved at various points. Fully automatic feature capture has been a long-term goal of several ISPRS working groups, but it is unlikely to be completely realised in the near future. Semi-automation of the capture of topographic objects is a feasible option – especially in areas where many features are repeated (such as housing estates where every house is almost identical to its neighbour). In Great Britain, however, this is very seldom the case. The continuous nature of topographic data revision usually means that changes are very local in nature, often involving only one or two features within an area of many un-changed features.

Automation of the change detection process is more feasible, especially if the system is only required to *detect* that a change has occurred, without having to *identify* exactly what that change is. In recent years there has been continued improvement to photogrammetric and image processing software and the hardware on which they run. This has enabled systems to become *more* automated, either by aiding the interactive data capture process or by performing a pre-processing step on images before the human operator becomes involved. The nature of the workflow at Ordnance Survey lends itself to the latter process, in which an automated process would identify potential changes, then an operator would confirm or reject each change. It was envisaged that such a process would save a considerable amount of time in the data collection workflow.

There has been much research in the last two decades on techniques to automatically detect changes between images taken on different dates. Such techniques can be of value, but often they highlight many changes which are of no interest – such as minor changes to vegetation, or changes due to the

movement of traffic, shipping containers, and other transient features. The effectiveness of the technique also depends, to a greater or lesser extent, on the nature of the data that is to be extracted. For example, a process that can identify changes to road networks in a rural area may not work as well in an urban setting, and may be entirely unsuitable for the detection of changes to buildings. The nature of the input data also plays a major role in the process. The resolution (ground sample distance) of the images will often dictate which techniques are most appropriate to use, while the presence of an infra-red component in an image can be of great importance when separating vegetation from the built environment.

One task of the data collection area of a mapping agency is to detect changes between the features held in a topographic database, and the features present in an image or set of images. In an ideal situation, an automatic technique is required which compares the existing topographic data with a single up-to-date image and identifies the differences between them. It would be very useful if the process could also filter out any changes which are unlikely to be of interest (e.g. traffic) and present to the photogrammetrist only those changes which are required by the data capture specification. This paper presents various methods of automatic change detection investigated by Ordnance Survey during the last year. These methods are largely based on image classification, followed by feature comparison techniques. Per-pixel and per-object classification methods were tested, using both off-the-shelf systems and techniques developed in-house.

2. CHANGE DETECTION METHODOLOGY

2.1 Which changes are important?

Of the many different features present in the national topographic database, changes to the built environment proved to be the most critical. The construction and demolition of buildings are both important to many users of spatial data and, as uncovered by the manual change detection processes, are often overlooked by third-party change intelligence sources. It was therefore decided that our research should concentrate on the detection of new buildings and demolitions. Once this decision had been made, the exact method of change detection had to be determined. Rather than taking one single approach, several different methods were tried and compared, to find the one showing the most promise for future implementation into the production system.

2.2 Source data

Since the target of this research is the implementation of a change detection process within the photogrammetric data collection production system, it was important to use as inputs only those types of data which would be readily available to that system. For several years, Ordnance Survey has used an Intergraph Z/I Imaging Digital Mapping Camera (DMC) as its primary image data source. The inputs to this research were therefore constrained to the data which can be extracted from this DMC imagery. Both the panchromatic and 4-band 12-bit multispectral images from the DMC were used at some point in the research. Rather than relying solely on the spectral aspects of the data, the projects also used a digital surface model, created fully automatically from the overlapping panchromatic imagery. This allowed us to discriminate more easily between man-made and natural objects in the scene.

Two test sites were chosen, both of which have undergone many changes recently: one near the Heathrow Airport Terminal 5 junction of the M25 London orbital motorway, the other in the urban centre of Bournemouth, a city on the south coast of England. The images were collected during the 2005 and 2006 flying seasons. The two sites gave us the opportunity to test the algorithms in different physical environments, to indicate whether the process is likely to be transferable for use in different parts of the country.

2.3 Change detection via image classification

Each of the change detection methods investigated involved an image classification process. Both per-pixel and per-object classifications were undertaken, using various methodologies and several different software packages. Initially, the classification techniques were applied to the images to identify buildings, roads, trees, other vegetation, water bodies and roads. These were then filtered to identify the changes to buildings, which were the main focus of the research.

2.4 Per pixel classification

One of the main challenges in the detection of urban change is the spectral heterogeneity of the urban land cover (Small 2001). The many different surface types and objects present within an urban scene can often generate different spectral responses for essentially a single land cover. In order to deal with this problem, two relatively new per-pixel image classification techniques, which had the potential to discriminate between objects in an urban scene, were applied to the images. The first of these was the Support Vector Data Description (SVDD), a one-class classifier developed by Tax and Duin (1999). The SVDD is based upon the principles of Support Vector Machines. Being a one class classifier, it works on the basis that only target class data are used in the training stage. The target class refers to the class of interest and it is assumed that it is sampled well and that enough training data is available. However, in the testing and validating stage the classifier will encounter outlier data that was not present in the training stage. The classifier must therefore have the capacity to distinguish if the data in a testing set belongs to the target class or it is unknown and as such belongs to the outlier class (Tax and Duin, 1999). In order to give the maximum information about each class and to allow comparison between different images taken at different times, the variables used included band ratios and texture for each of the bands, NDVI (Normalised Difference Vegetation Index) and a DSM (digital surface model) for the area. The training set was composed of a mixture of different pixels from different rooftop materials. The testing set was composed of pixels, 50% belonging to the target class and 50% belonging to all the other classes present in the image. All the pixels were chosen randomly from across the image. The pixels that formed the testing set were totally independent from those used in the training sets to avoid any biases in the confidence level of accuracy.

The second per-pixel classifier was a decision tree method, specifically the CART decision tree software developed by Salford Systems. Decision trees have long been popular in machine learning, statistics and other disciplines for solving classification problems. Decision trees are very flexible and can handle non-linear relationships between features and classes (Friedl and Botley, 1997). CART uses several different approaches to the splitting of the decision tree, including the Gini Index, entropy, and class probability. CART uses an over-

grow/prune-back strategy in which the decision tree is allowed to grow to a high level of complexity, and is then pruned back to a more manageable level. This approach helps prevent the classifier getting stuck at a local optimum point too early in the classification process.

2.5 Object-based classification

An object-based classification was adopted by the third method, using Definiens software. This object-based method offers some advantages over the traditional pixel-based approach, in that, in addition to the spectral data, it can incorporate shape, texture and local context into the classification process (Benz et al, 2004). Also, by segmenting the image into homogeneous objects prior to classification, this approach helps to reduce the spectral variability within each class.

The first stage of the classification process was the segmentation of the image into homogeneous objects. Definiens employs a region-merging segmentation method, starting with individual pixels and merging adjacent regions until a user-defined threshold of heterogeneity is reached (Benz et al, 2004). After the initial segmentation, the objects were classified using a hierarchy that divides the objects into increasingly refined categories, using user-defined membership functions. The membership functions are determined by observation, i.e. by manual examination of the characteristics of the objects of interest in the image. Although this is a slow process, it should only be required once, in the expectation that the objects of interest in other images will have similar characteristics. Earlier splits were simpler and more reliable (e.g. vegetated objects were split from non-vegetated objects using only the Normalized Difference Vegetation Index, NDVI, with very high reliability). Subsequent classes became progressively more difficult to separate and the reliability of the classification decreased. Spectral properties, shape and texture features were all used in the identification of buildings. Shadow was also used to constrain the buildings class to those objects within a defined distance from a shadow object. This methodology is discussed further in Sanchez Hernandez et al (2007).

3. CLASSIFICATION RESULTS

3.1 Classification accuracy

The first test of the classification methods used a 300 m by 270 m subset of the Heathrow dataset. The aim of the test was to determine the accuracy of each classification, in order to assess its usefulness in a subsequent change detection step. A training set of 600 pixels was selected, with 50% of the pixels in the buildings class and 50% in the non-building class. A further set of 200 randomly selected pixels formed the test set. The overall accuracies of the test pixels are shown in Table 1. The first set shows the accuracy achieved using only the spectral components of the image, while the second set shows the accuracy achieved when the digital surface model was also included in the source data. As can be seen, the DSM data increased the accuracy of the classification in all cases. This was as expected, since the buildings can be spectrally very similar to the surrounding man-made roads and other surfaces; while the difference in height distinguishes them immediately from the ground surface.

Classification	SVDD	Decision Tree	Object-Based
Spectral only	70.0%	73.0%	76.0%
Spectral + DSM	85.0%	90.5%	91.0%

Table 1: Overall accuracy of the classified Heathrow subset

For the object-based classification, a slightly different method was applied when the DSM data were included. First, the image was classified by a user-defined sequence of processes within Definiens. The main features used in this classification were shape, slope, and height in context to neighbouring features. The spectral information used in the process was of a lesser importance and was limited to the normalised difference vegetation index (NDVI) and the brightness. These were chosen since other spectral features are likely to differ significantly between images captured in different areas, at different dates and in different lighting conditions. Localised differences in height between building features and the surrounding land are less likely to differ between images and these should prove more reliable. It is thought that this will allow the process to be generalised and will therefore increase the potential for transferring the rule-set to different images.

The results on a larger test area were similar to those shown in Table 1. As with any classification process a certain amount of misclassification is inevitable. In the per-pixel classifications, the main reasons for misclassification were the presence of objects on the surface which could not be distinguished from the surrounding buildings. These objects included large vehicles and shipping containers, which are similar in height and area to small buildings. A further cause of misclassification was the use of the DSM data as absolute height values (rather than heights relative to the surrounding pixels). Using the absolute heights works well in a flat area, but in areas of undulating terrain there will be problems between rooftops and man-made surfaces at the top of slopes.

In the object-based classification, vehicles and containers were also often misclassified as buildings. This classification, however, could distinguish successfully between rooftops and manmade surfaces, by using the local slopes around each of the objects. Some misclassifications still remained, due to very low-rise buildings such as sheds and garages, which were not recognised as buildings. This occurred because these objects failed to meet the height and area thresholds within the rule-set. This was expected to occur because the rule-set was devised to meet the Ordnance Survey specification for significant changes ("Category A" changes). Category A includes newly built residential buildings and demolitions, but small, non-residential buildings are not within this specification. A Definiens rule-set could be constructed to detect these small buildings, but this would inevitably lead to misclassifications of objects of a similar size and height, such as vehicles. In the end, a trade-off must be made between the proportion of building features correctly identified as buildings (true positives) and the proportion of non-building features that are misclassified (false positives).

The results indicated that the decision tree and the object-based methods were the most reliable, so these two methods were further developed, by introducing post-classification change detection processes which are described in the following section.

4. POST-CLASSIFICATION CHANGE DETECTION

4.1 Identifying changes

In order to test the accuracy of the classifications in a change-detection process, the classified buildings were compared with the building objects in the Ordnance Survey topographic database (used to generate the OS MasterMap® product). In the test area, all the significant building objects which were either greater than 50 m² in area, or had a postal address (and hence were residential buildings) were determined from the topographic database. In total, this revealed 965 buildings to test, of which 34 had been demolished. In addition to these, there were 17 significant new buildings on the image, making a total of 51 Category A changes.

4.2 Detecting demolitions and new buildings

Both the decision tree and object based classifications were tested using the same method of post-classification change detection. Demolitions and new buildings were considered independently.

Demolitions were identified by intersecting the areas classified as buildings with the known OS MasterMap® building polygons. For each building polygon, if at least 50% of its area was classified as a building, that building polygon was considered to be verified by the classification. If less than 50% of its area was classified as a building, then it was considered to be a change (i.e. the building was considered to have been demolished).

To identify new buildings, the first step was to mask out all the regions in the test site where constructions would be unlikely to have occurred. These consisted of all roads, rail or water bodies present in the OS MasterMap® data. All existing buildings in the data were also masked out, together with a 3 m buffer around each building, to help eliminate any remnant-objects produced by misalignment between the image and the topographic data or by the draping effect of the DSM. The remaining area, consisting of vegetation, farmland and man-made surfaces, was then searched for any objects classified as buildings. Objects smaller than a given size threshold were filtered out, to leave a set of potential new buildings.

4.3 Results of post classification change detection

Table 2 shows the results for the decision tree classification, and Table 3 shows the results for the object-based classification. It can be seen in both cases that, of the 51 Category A changes on the image, 49 were successfully identified, with only one actual change not flagged as a change (false negative) each for demolitions and for new buildings. These errors were caused by a single feature - a residential building that had been demolished and rebuilt with a similar footprint. Such rebuilds are inevitably difficult to detect when the footprint in the map database of the recently demolished building is similar to the footprint in the image of the building constructed in its place.

	Demolitions	New	Total
Actual changes	34	17	51
Objects classified as changes	288	161	449
Actual changes correctly classified (True Positives)	33	16**	49
Non-changed objects classified as change (False Positives)	255	150	405
Actual changes not classified as changes (False Negatives)	1	1	2
% Classified as changes that were actual changes	11%	10%**	11%
% of actual changes classified as changes	97%	94%	96%

Table 2. Results of change detection from decision tree classification for the 2 km² Heathrow test site.

In the decision tree results (Table 2) there are a large number of false positives, in which objects which haven't actually changed are falsely identified as changes. These errors are caused by a variety of factors. One factor was that the change detection works on classified features, rather than individual pixels. In order to do this, the groups of pixels classified by the decision tree had to be grouped into contiguous areas and converted to vector form. The process of grouping and vectorising inevitably leads to a slight degradation in the quality of the results. A second factor is the misclassification of an area of active construction as a building. In the construction site, earthworks for a new road were of a similar spectral nature to the buildings, and were elevated above the average ground surface height, making them similar in height to the building objects nearby. Other false alarms were caused by the presence of large vehicles and shipping containers – features which were much in evidence in this area of active construction work.

	Demol	New	Total
Actual Changes	34	17	51
Objects classified as changes	79	96	175
Actual changes correctly classified (True Positives)	33	16***	49
Non-changed objects classified as change (False Positives)	46	84	130
Actual changes not classified as changes (False Negatives)	1	1	2
% Classified as changes that were actual changes	42%	17%	28%
% of actual changes classified as changes	97%	94%	96%

Table 3. Results of change detection from object-based classification for the 2 km² Heathrow test site.

** Of the 161 objects flagged as new builds, 11 contained actual changes, but this included 16 individual new buildings as many were close together and so were merged in the classification

*** Of the 88 objects flagged as new builds, 12 contained actual changes, but this included 16 individual new buildings as many were close together and so were merged in the classification

In the object-based classification, the true positives were the same as for the other method. The difference in the object-based method is in the number of false positives. Although there are still many features incorrectly flagged as changes, there are significantly fewer than in the decision tree method. Of the 965 buildings present in the topographic database, 886 (92%) were detected by the method, leaving 79 predicted to be demolished. Of these, 46 were false alarms, caused mainly by low-rise buildings such as industrial sheds and small residential properties that were in the topographic database but were below the height threshold required to be classified as buildings. For new buildings, 16 of the 17 actual changes were identified, but 84 other objects were flagged as new buildings but turned out to be false alarms. These errors were mainly from objects such as caravans, lorries and shipping containers being mis-identified as new buildings and from garages and that fall outside the Category A specification.

5. CHANGE DETECTION PRODUCTION TRIAL

The results of the change detection led us to choose the object-based method for further development. To test the method in a more realistic environment, a prototype production system was developed, using a combination of software already used within the production area and software required for the classification and change detection process. It was decided that the process would be tested on a live production job and directly compared with the manual process of change detection currently employed.

Two test sites in Sunderland in NE England were chosen for the production trial, one area (site A) of 23 km², the other (site B) of 25 km². The DMC imagery was processed to provide the inputs required by the object based classifier. These inputs consisted of an orthorectified 4-band image mosaic, a DSM and a slope map derived from the DSM. The newly-acquired NGATE module of SOCET-SET was used to produce the DSM fully automatically (i.e. without any seamline editing or other manual processing). The same rule-set used on the first test dataset was also used in the trial, to test whether the rules could successfully be applied to different areas.

The object-based change detection method was applied to the Sunderland image data and OS MasterMap topographic data. This resulted in a set of polygons representing potential demolitions and new buildings. These were presented to the image interpreters using a similar user-interface to the one they would normally use for manual change detection. The interface was modified slightly to automatically direct the user to each potential change in turn, and zoom in to the image at that point. The user then compared the image to the topographic data in the area of potential change and clicked a button to either accept or reject the change. Once a button was clicked, the result was recorded and the user was immediately presented with the next potential change on the list.

5.1 Results of the trial

In site A, 142 potential changes were detected. Of these, 35 were accepted as real Category A changes (representing 25% completeness) and just one real change was missed. In site B, 427 potential changes were detected, of which 77 were accepted as genuine (18%). There were 14 real changes that were missed in site B, which are discussed below. In terms of time taken to identify the changes in the two sites, it was found that the

application of the automatic change detection process reduces the overall time by 50% (compared with the manual process). This is a significant improvement on the manual process, and would be improved further by some small changes to the user interface, which were requested by the operators.

5.2 Changes that were missed

The results show that 14 of the changed features were not detected by the automatic process, even though the initial test had very good completeness statistics. The nature of the omissions in site B were:

- 5 demolitions
- 4 new buildings
- 5 minor modifications to school buildings

Of the 5 demolitions, 3 were small buildings which did not have an address in the database, and therefore were assumed to be insignificant, non-residential buildings. In the manual change detection process, these were recorded as changes, even though they did not meet the exact criteria to be considered as Category A change. This is a situation in which the human operator will err on the side of caution, while the automated process simply filters these features out. A second reason for the omissions also involves an interpretation of the specification, this time relating to minor changes to school buildings. The automated process ignored minor changes to existing buildings, while the human operator treated school buildings as a special case and marked up any changes in such features. This could be overcome by identifying schools in the topographic data and treating them as special cases in the automatic process. A third reason for the discrepancies was the time lapse between the imagery and the topographic data. In several cases, there were genuine changes between the map data and the image, but the map data was more up-to-date, having been updated via field survey since the imagery had been flown. The human operator can readily verify this, while the automated process cannot.

5.3 Opinions of the operator

When asked for his opinion of the automated process, the operator considered that it would halve the time taken to collect the change data, and that it was "a very useful tool for change intelligence". To determine whether this evaluation of the process is widely accepted, further tests of the system are planned.

6. CONCLUSIONS AND FURTHER WORK

An automated tool for detecting changes to buildings using an object-based classifier has proved to be sufficiently successful in a research environment to be taken up in a production trial. In this trial, the tool was well received and was shown to significantly reduce the amount of time taken to identify the changes over a 48 km² area. At the time of writing, a large production trial is planned, which will use recently-captured imagery in an area which has changed significantly since the previous update to the topographic data. This will allow us to test the method in another area, and to determine its effectiveness and efficiency in an area with a large number of changes.

REFERENCES

- Benz, U.C., Hofmann, P., Willhauck, G., Lingenfelder, I., Heynen, M., 2004. Multi-resolution, object-oriented fuzzy analysis of remote sensing data for GIS-ready information. *Photogrammetry & Remote Sensing* 58, pp. 239-258
- Sanchez Hernandez, C., Gladstone, C., Holland, D., 2007. Classification of urban features from Intergraph's Z/I Imaging DMC high resolution images for integration into a change detection flowline within Ordnance Survey. Proceedings of the ISPRS Joint Conference URBAN 2007/URS 2007, Paris, France.
- Small C., 2001. Scaling Properties of Urban Reflectance Spectra, *AVIRIS Earth Science and Applications Workshop*, 27 Feb–2Mar, Pasadena, CA. <ftp://popo.jpl.nasa.gov/pub/docs/workshops/aviris.proceedings.html> (accessed April 2008)
- Tax D. M. J., Duin R. P. W., 1999. Data domain description using support vectors. Proceedings of the European Symposium on Artificial Neural Networks, Bruges, Belgium, April, 21-23, pp. 251–256