AUTOMATIC REGISTRATION OF IMAGES WITH MAPS USING POLYGONAL FEATURES

I J Dowman, A Morgado and V Vohra
University College London

Commission III, Working Group III/2

KEY WORDS: Automation, Matching, Orientation, Registration

ABSTRACT

The automation of the full orientation procedure for images from aircraft or satellites has progressed significantly during the past 4 years. In particular inner and relative orientation have matured into automatic production processes and aerial triangulation, with the exception of a general method of identifying ground control points, is well on the way to the same stage. This paper reports on the development of tools and on an overall strategy which permits automated registration of an image or 3D model to a map or a digital data set using large polygonal features, stored in raster or vector format. No manual identification of discrete points is required and the methods takes into account any changes which might have taken place between the original information being compiled and the current imagery being obtained. The method can also be used for registering images to images and allows a wide range of image types to be used, ranging from aerial photographs to synthetic aperture radar (SAR) data.

1. INTRODUCTION

The problem of absolute orientation of aerial photography can be tackled in a number of ways. Automated aerial triangulation is a promising method in which photographs in a block can be joined using a large number of conjugate points which are automatically determined. A relatively small number of ground control points then have to be identified manually, although Gülch (1995) has shown that certain types of ground control points can be identified automatically. If the ground points are premarked then the problem is tractable. The orientation of single pairs or of satellite images is more difficult because a relatively larger number of ground control points are required. For high accuracy work at large scales, high level feature extraction is required and is the topic of a large amount of research. At smaller scales, and particularly when map revision is the main application, then an alternative approach is possible in which the image is registered directly to the existing map or data base, referred to as the reference data. A necessary condition of such an approach is that the matching of common features must be done in two dimensions. In some cases, for example using satellite images over flat ground, this might be sufficient, but is other cases the third dimension must be introduced after matching. This approach offers a great deal of flexibility and many of the advantages seen in the matching strategies used for automatic digital elevation computation. For example the determination of many redundant matched points and the ability to tune the system for different types of image.

The system described here is designed as a flexible generic system which will allow a range of image types to be registered with a reference data set. The system will be fully automatic but will offer the user a number of tools to employ to optimise the solution for the particular data being used and to validate the end result.

The method is based on the selection of conjugate polygons which can be uniquely identified in the image and the reference data. A number of well established methods are available for segmentation and edge detection for use in defining homogenous areas on an image. Polygons are more easily identified uniquely in vector data than are points or lines, hence polygons are highly suited to this task. Recognition of the possible distortions due to tilt and relief must be made but these are generally easily modelled.

The paper first outlines the strategy to be used and then describes a number of algorithms which have been developed and tested. Three examples are given in the paper. New work which carries out the full absolute orientation process on a pair of aerial photographs using a 1:10 000 map is described in some detail. Previously published results showing the matching of woodland from attributed layers of a 1:50 000 map and a Landsat Thematic Mapper image, and matching of large buildings on a 1:10 000 raster image and a 1:10 000 map are summarised.

2. A STRATEGY FOR IMAGE TO MAP REGISTRATION

2.1 Preparation of the reference data

The strategy is shown in figure 1. Each stage will be discussed in outline but reference should be made to the papers cited for further detail. The method involves the separate preparation of the image and the reference data. The preparation of the reference data will differ according to the source. If the data is available in digital form, with attributes attached, than the process of extracting polygons of a given type is straightforward. A paper map by itself provides a greater challenge and some manual operations are still necessary to identify particular features. If the original layers used in the printing stage are available then this will enable features such as woods or water bodies to be easily isolated. If such layers are not available then manual identification may be necessary. One approach to this is described by Vohra and Dowman (1996) in which polygons have been identified semi-automatically and then only those exceeding a specified area are retained.

The degree of processing on the image depends on the method to be used. If general polygons are to be used then edge extraction and segmentation are needed. Generally homogeneous areas are looked for and a multispectral classification may assist if features such a woodland or water are used.

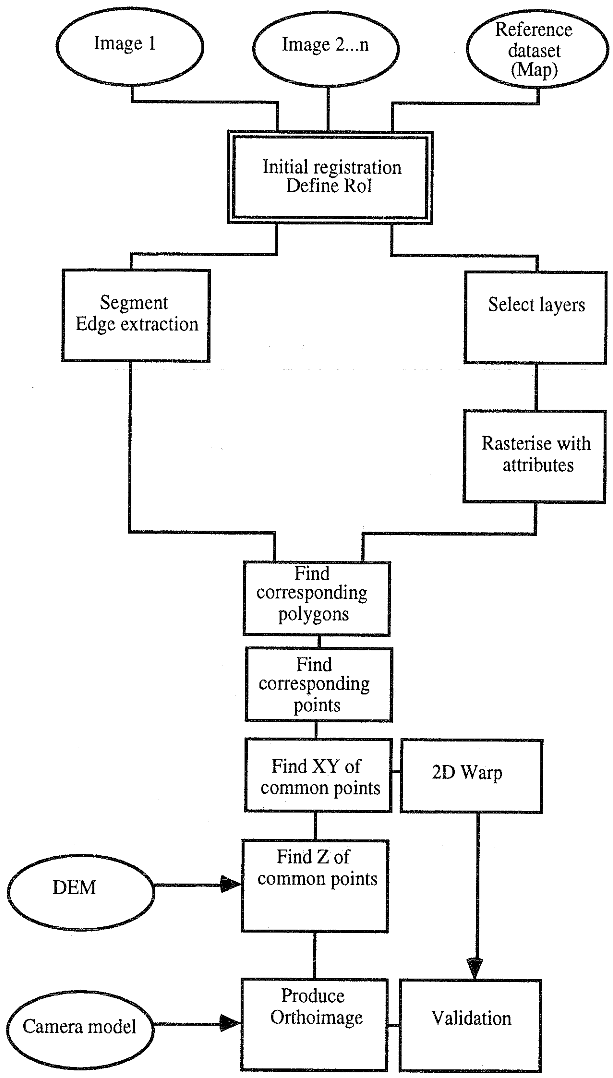

Figure 1. A strategy for image to map registration.

The general stages used in extracting polygons from the map are smoothing, segmentation, edge enhancement, edge thinning and removal of small polygons. It must be remembered that for the process of registration only a small proportion of features are needed and the rejection of some polygons is not a problem.

2.2 Matching the polygons

Polygon matching is required at two levels. First corresponding polygons must be established and then the edges of these must be matched on a point by point basis.

To carry out the first stage an initial, approximate correspondence must be established. A number of techniques are available. Alternatives which are available include the manual selection of 2 to 4 control points or an automatic identification of position and orientation by matching large polygons, possibly embedded into an image pyramid. The first two methods have been used but it is thought that the automatic method may not be robust in all circumstances using only a single layer.

The basic techniques of matching polygons is adapted from Abbasi-Dezfouli and Freeman, (1994). Polygons are characterised by a number of parameters such as shape and area. Shape is defined by dimensions parallel to defined axes and orientation and also by the chain code method described by Abbasi-Dezfouli and Freeman. The initial translation and azimuth must be fixed by first defining a few polygons which have good matches based on a first pass through the selected points. An iterative approach then allows corresponding polygons to be identified. A large number of polygons are not necessary but it is important that they are distributed in a suitable pattern over the image. The method is described in detail in section 3.

Once established he corresponding polygons must then be exactly matched in order to extract conjugate points. A method of dynamic programming developed at UCL is one method of doing this, (Newton et al, 1994). The perimeter of the feature is followed and a best fit obtained. The matching of the edges uses the method first developed by Maitre and Wu (1989). Costs are determined by a number of measures relating the predicted edge pixel position projected into the map and the edge pixel under consideration. The difference in gradient direction between the map boundary pixel and the edge pixel under consideration are used as costs. The method also allows the detection of changes between the two polygons which may represent true change or an error in detection, in either case such points will not be selected as conjugate. The technique makes allowance for the fact that the image may be distorted due to terrain effects or geometric effects from the camera or sensor.

Other techniques are under investigation such as recognition of corners.

2.3 Transformation

Once the corresponding points are selected then the required transformation can be computed. If absolute orientation of a pair of aerial photographs are under consideration then three dimensional reference data must be available. Model co-ordinates can be found using the co-ordinates of one image, the relative orientation elements and an approximate surface model to predict the area on the second photograph within which the conjugate point lies. The exact position can be found by stereo matching. If finite element modelling was used to find the relative orientation then the output from this could be used to determine the conjugate co-ordinates. If only a single image is to be matched then a two dimensional transformation will be used or a space resection carried out.

3. REGISTRATION OF FIELD BOUNDARIES ON AERIAL PHOTOGRAPHY TO 1:10 000 MAP DATA

The first example is close to a fully automated system in that polygons are selected in an entirely automatic manner from the image and from the map and the method of polygon matching described by Abbasi-Dezfouli and Freeman, (1994) is used. Furthermore the two dimensional conjugate points derived from the image are given a height value so that a full three dimensional transformation can be carried out. The method is described by Morgado (1996).

An aerial photograph of the Isle of Wight was scanned using the Sharp JX-600 scanner at 600dpi resulting in a ~26Mb image. The scale of the photography is 1:11 600 and the pixel size of the resulting digital image is 0.50m. The

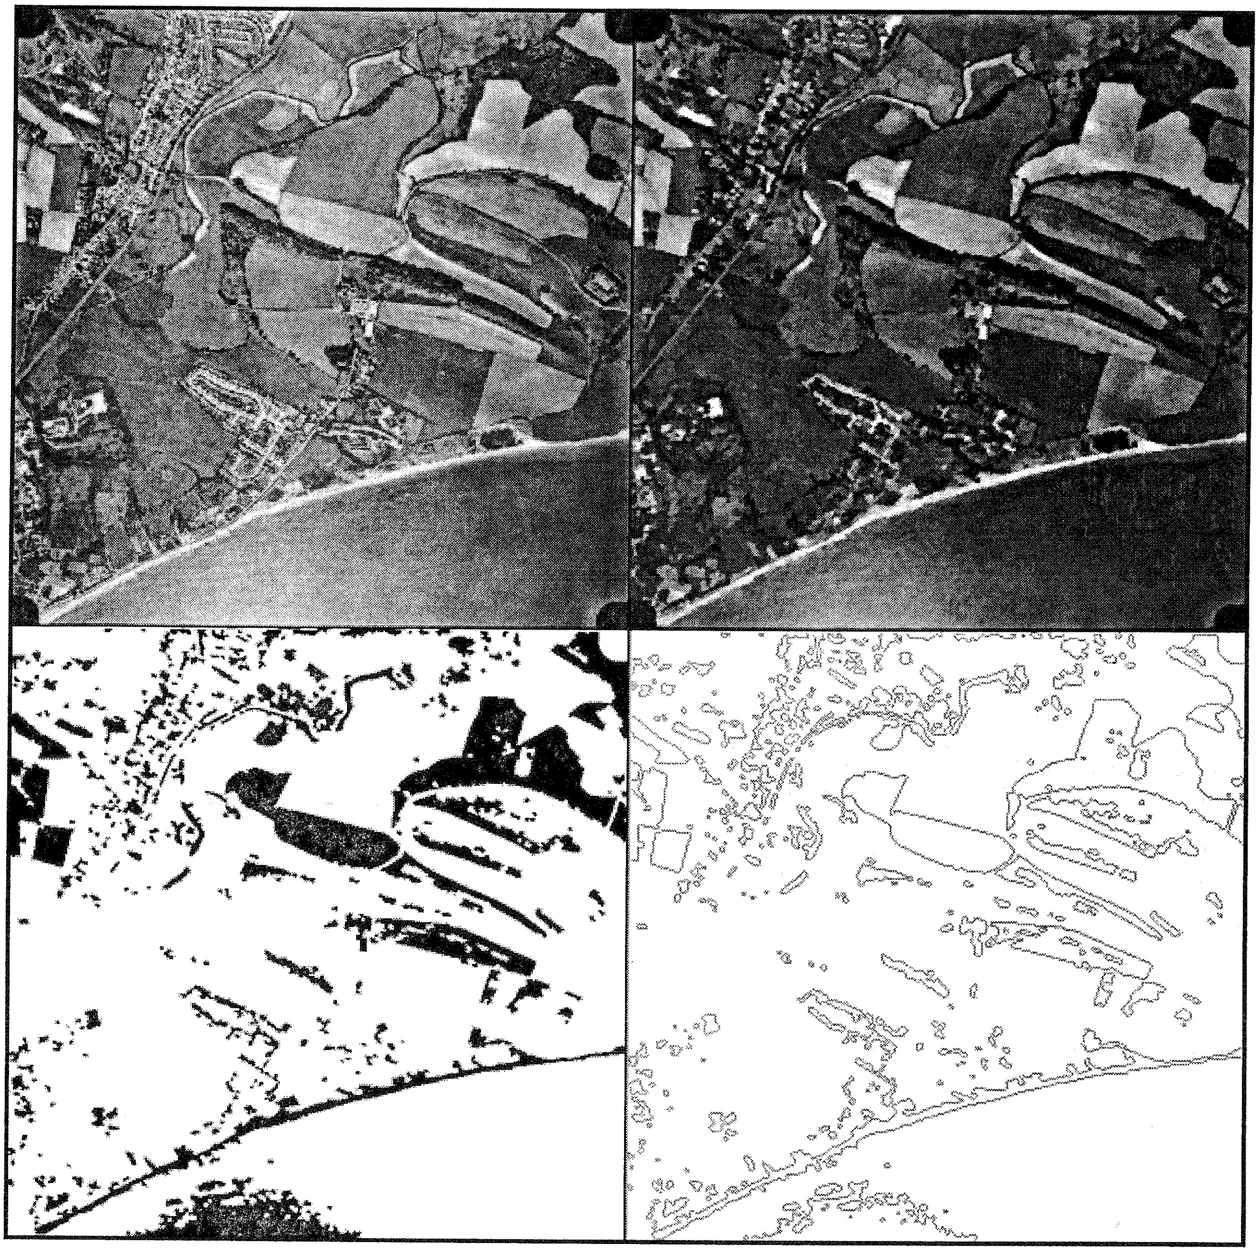

Figure 2 . The Isle of Wight original image (top left), the smoothed image (top right), the thresholded image (bottom left) and the edges detected in the image (bottom right)

photograph was taken in 1985. The 1:10 000 Ordnance Survey map of the same area represented on the photograph was used to digitise the field boundaries. This map was surveyed at 1:2 500 scale between 1969-72 and revised for significant changes in 1973, the publication date is 1975. The vector file obtained was processed as described in section 2.1. The maximum resolution on a map is 0.2mm, therefore a precision greater than 2m can not be expected from a measurement taken from a 1:10 000 map. This means that even if a resolution smaller than 2m per pixel is used, this will not increase the accuracy of the measurements made on the digital map. Taking this into account and at the same time trying to minimise the amount of data to be processed, both the map and the image were scaled to a pixel size of 3.92m, thus resulting in a ~660x660 pixel image. In this case the highest point on the photograph has a height of 100m, even if this point was at the extremity of the photograph where the relief distortion is maximum, this distortion would be 0.6mm, which corresponds to 1.4 pixels on the digital image. Since the image was reduced 8 times, this distortion is irrelevant. The digital photograph was prepared using the procedure described in section 2.1: it was smoothed, thresholded, the edge detector was applied and finally the smallest polygons were removed from the image. The results of each step of this process are shown in Figure 2.

Figure 3 shows both the image and the map prepared for the matching process, the polygons are numbered according to the order used to read them. A record of the frequencies of the directions between neighbouring pixels, the frequencies of the change in direction segments, the area of the polygon, the perimeter of its boundary, the width and height of the bounding rectangle and the coordinates of the centre of gravity of each polygon were derived. Figure 4 shows an example of the information recorded for one of the polygons represented in the Isle of Wight image.

Polygon number = 1

perimeter = 135
area = 227

Code = 61 3 1 13 39 13 3 1
Differences = 93 23 1 0 0 0 0 17

Width of bounding rectangle = 65
Height of bounding rectangle = 9

Centre of gravity (X) =377.8
Centre of gravity (Y) = 3.6

Figure 4. Information extracted from a polygon represented in the Isle of Wight image

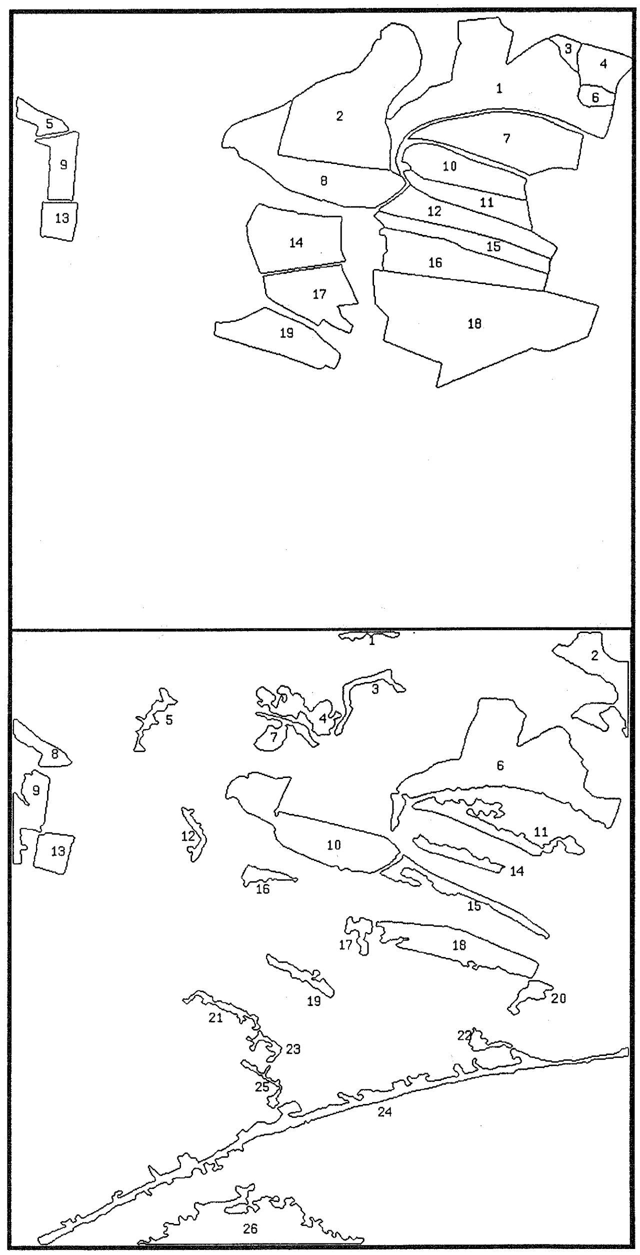

Figure 3 . The Isle of Wight map and the image prepared for matching

A program was written to compare the information relating to each of the polygons extracted from the map and from the image. For each polygon in the map differences were computed between each of its attribute values and the corresponding attribute values for all polygons in the image. After adding all the differences, the minimum value was found and the pair of polygons that correspond to that minimum is considered matched. As mentioned by Abbasi-Dezfouli and Freeman the chain code approach followed in this study is very often a reliable way to distinguish between different shapes. However the method by itself is not infallible and the other parameters used to characterise a polygon should provide the information to compensate for a possible uncertainty derived by the frequencies of the chain code direction segments and change in direction, if these frequencies were used alone.

It was observed that different weights should be given to the differences between the parameters that characterise the polygon shape. More precisely, it had to be considered that the difference in area of the polygons to be matched had to be scaled to at least to half of its value since it was adversely affecting the matching results. This problem can be easily understood by considering that the area of the polygon includes all the pixels inside the boundary of the polygon, depending on the size of the polygon a slight difference in shape can contribute to a high difference in the area. If not reduced the difference in area between the polygons would be the main characteristic to be considered when adding all the differences between the different parameters used to characterise the polygons. Taking this into account the matching was performed with no problems. For each polygon in the map one polygon in the image was found whose characteristics were considered to be the most similar. The resulting value from adding all the differences between the parameters of each pair of polygons was then divided by the perimeter of the polygon to compensate for the fact that more differences are expected to occur on polygons with a longer perimeter. It was observed that to eliminate the bad map to image correspondences the matching value for each match did not provide a very robust way of identifying the bad matches. To solve this problem, it was also necessary to consider the relative geometry of the patches to detect the wrongly matched patches. Two approaches were considered to perform the selection of the good matches. The first method consists of computing geometric relations between the map polygons, such as distance and angles, and then checking that the corresponding image polygons maintain similar geometric characteristics among them. The second technique starts by considering that the best matching value should belong to a correct match. Based on this, a transformation is computed between the best matched polygons using the centre of gravity, and the furthest and the closest point from the centre of gravity of both map and image corresponding polygons. The parameters found by this transformation are then applied to transform each map polygon centre of gravity into the corresponding image polygon centre of gravity. The transformed centre of gravity is then compared with the centre of gravity of the image polygon that is considered to match the map polygon under consideration. If these two values coincide then the matching is accepted, if they are different then this means that the original correspondence is wrong and that map and image polygons do not match. Both techniques were implemented. It was observed that the second approach can cause problems because the map to image transformation initially used focuses on a restricted area of the image where the polygon that starts the process is located, this may be quite different from the transformation to be used on polygons distributed in the different areas of image. Therefore the first procedure to filter the bad matches was adopted since it applies extra geometric constrains that assure a correct result. Table 1 shows the matching values, the values highlighted represent the good matches.

Once the common polygons are identified a refinement of the matching and a consequent identification of conjugate points can be performed using the map to image matching technique based on the dynamic programming algorithm described in section 2.2. Now that the conjugate polygons have been found using the map to image matching algorithm based on shape, the knowledge of the centre of gravity of these polygons together with the width and height of the bounding rectangle can be used to section the original image into areas only containing the polygon to be further processed. The resolution of the image and the map for this matching is 2m per pixel which corresponds to the minimum resolution to be used with a 1:10 000 map. Using the method outlined in section 2.2, the initial estimation of the transformation between the map and the image can be computed.

Table 1 Isle of Wight matching values and corresponding matched polygons

Matching Value	Map Polygon	Image Polygon
0.629	**1**	**6**
1.641	2	6
1.476	3	17
1.826	4	13
0.810	**5**	**8**
1.370	6	20
0.839	7	10
0.660	**8**	**10**
1.491	9	9
1.644	10	18
1.780	11	4
0.623	12	18
0.555	**13**	**13**
2.778	14	18
1.008	15	18
1.284	16	18
1.846	17	18
2.498	18	6
1.463	19	18

Figure 5 shows one of the map polygons with the respective image polygon used. These two polygons were independently processed using the techniques described in section 2.2 in order to prepare them to be matched. The points matched by the dynamic programming are also shown.

4. REGISTRATION OF LANDSAT THEMATIC MAPPER DATA USING FOREST AREAS

The basic method is described in Newton et al (1994). The objective of this work was to detect change in forest and the strategy adopted was to manually identify corresponding forest areas and then to use the dynamic programming routine to match the edges of the polygons and to detect any changes between the map and the image. A major part of this work was to identify and classify the forest using multispectral data. Landsat Thematic Mapper data was used and after classification the forest area was segmented using a disc filling routine and the edges extracted. The corresponding forest is found from the green layer in 1:50 000 vector data.

The boundary matching of Thematic Mapper data was tested by Holmes (1994) by selecting control points along the boundary and determining the parameters of an affine transformation. The root mean square error on 4 points was 0.46 pixels in plan and inspection showed a distribution of residuals which was consistent with an affine transformation but which could be corrected with a plane projective transformation. The assessment also showed that the method would not be entirely reliable when complex edges were involved. It was found that the relief differences over the TM image, although quite large would not affect the accuracy of the boundary determination.

In this part of the work only the boundary matching is automatic and the output is a set of two dimensional conjugate points.

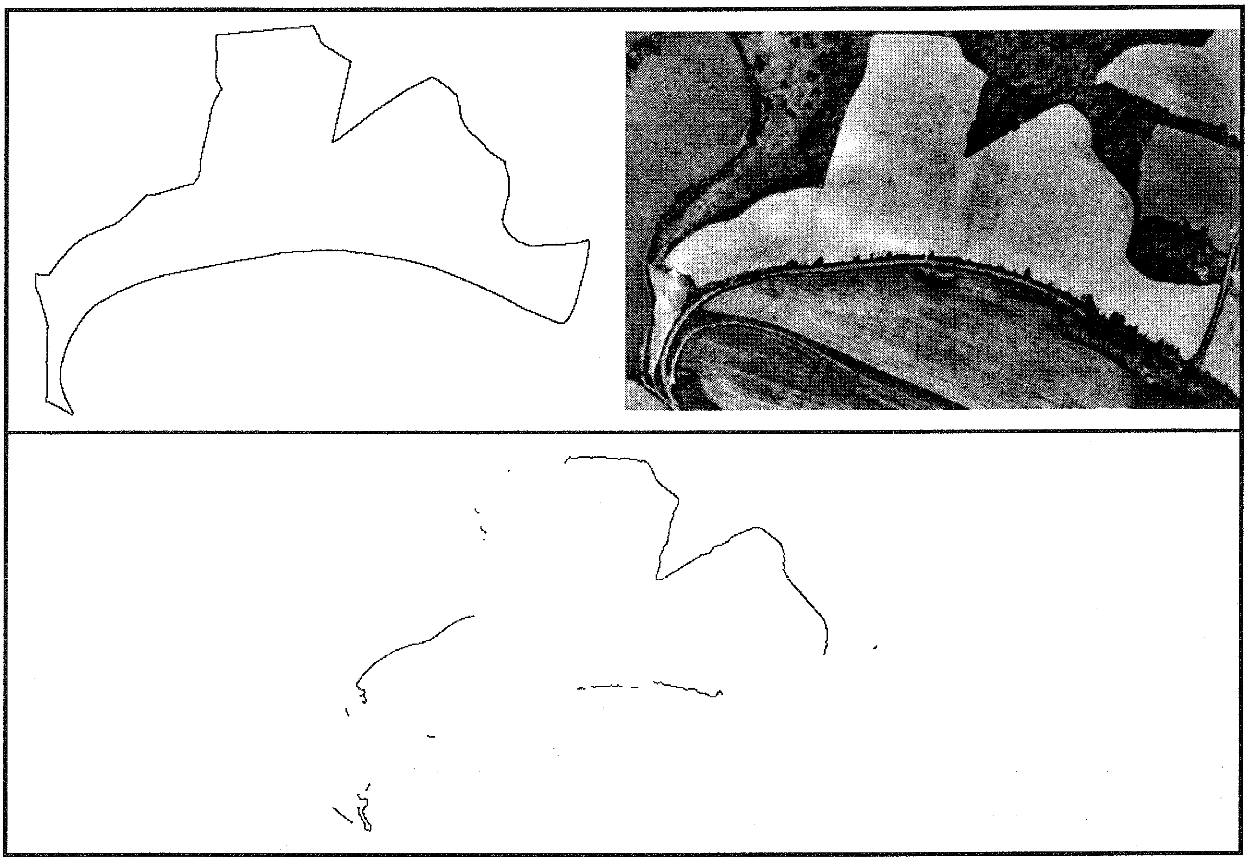

Figure 5. One of the corresponding polygons in the map and the image and the matched points.

5. REGISTRATION OF BUILDINGS ON DD5 DATA TO 1:10 000 MAP

A further test was carried out on the extraction of large buildings from satellite data with a 2.5m pixel size. In this work the buildings were registered with a 1:10 000 scale raster map. The work is described in detail by Vohra and Dowman (1996).

The preparation of the map involves considerably more effort than in the two examples discussed above, as a lot of unwanted information has to be removed. Buildings are first identified manually and then assigned a unique DN value. Objects below a given size are removed and the edges of the remaining objects are extracted. In order to carry out the polygon patching the edge strength and direction are needed and these are also extracted at this stage. The image is prepared in the way described in 2.1 above.

The matching of the polygons produced 923 matched points and root mean square residual after transformation with an affine transformation of 2.0 pixel. Inspection of the residuals clearly showed the effect of perspective distortion of the buildings and a arbitrary figure of 0.5 pixel was adopted as the acceptable limit. This left 64 points with a rmse of 0.3 pixel.

Once again the area is flat and only a two dimensional set of points are determined. At this scale the perspective view of the buildings has a significant effect and the residuals on the matched points show systematic displacements which are related to the camera look angle and the orientation of the buildings. Future work will look for methods whereby the perspective effects can be identified and removed using a suitable model of the distortion. Shadows also have an effect which must be removed.

6. CONCLUSIONS

The three examples described illustrate a number of important component of the automated registration system. There is however still a considerable amount of work to be done to remove manual operations, refine the algorithms and introduce new algorithms to allow other data to be used. The results show that the core algorithm, the polygon edge matching, is robust and flexible. It allows matching to take place despite distortion and differences in the actual polygons. The results allow an analysis to be made which indicates that errors due to relief and perspective can be modelled and corrected for. They also show that a large number of conjugate points can be extracted and provided that the strategy used ensures that there is a good distribution of points, sufficient points should be obtained to allow three dimensional absolute orientation of two dimensional registration.

7. FUTURE DEVELOPMENTS

The method described will be developed in an European Union 4th Framework project involving University College London, The Technical University of Stockholm, University of Stuttgart, University of Oporto, Swedish Space Corporation and Earth Observation Sciences Ltd. This work will concentrate on refining the algorithms already developed, automated the areas at present done manually and developing new algorithms to handle other types of data such as radar.

REFERENCES

Abbasi-Dezfouli M and Freeman T, 1994. Patch matching in stereo-images bases on shape. International Archives of Photogrammetry and Remote Sensing 30(3):1-8

Gülch E, 1995. Automatic Control Point Measurement. Photogrammetric Week '95. Wichmann, Heidelberg. Pp185-196.

Holmes A, 1994. The accuracies and application of forest change detection algorithms for Ordnance Survey. MSc report, University of London. 64 pages.

Maitre H, Wu Y, 1989. A Dynamic Programming Algorithm for Elastic Registration of Distorted Pictures Based on Autoregressive Model, IEEE Transactions on Acoustics, Speech, and Signal Processing, 37(2).

Morgado A, 1996. Automated Procedures for Orientation of Digital Images. PhD Thesis, Univesrity of London. 252 pages.

Newton, W Gurney C, Sloggett D, Dowman I, 1994. An approach to the automated identification of forests and forest change in remotely sensed images. International Archives of Photogrammetry and Remote Sensing, 30(3):607-614.

Vohra V and Dowman I, 1996. Automated extraction of large buildings from high resolution satellite images for regsitartion with a map. International Archives of Photogrammetry and Remote Sensing, 31(3)