Building verification from geometrical and photometric cues

C. Beumier
SIC Dept., Royal Military Academy, 30 Avenue de la Renaissance, 1000 Brussels, Belgium

ABSTRACT

Damage assessment, change detection or geographical database update are traditionally performed by experts looking for objects in images, a task which is costly, time consuming and error prone. Automatic solutions for building verification are particularly welcome but suffer from illumination and perspective changes. On the other hand, semi-automatic procedures intend to speed up image analysis while limiting human intervention to doubtful cases. We present a semi-automatic approach to assess the presence of buildings in airborne images from geometrical and photometric cues. For each polygon of the vector database representing a building, a score is assigned, combining geometrical and photometric cues. Geometrical cues relate to the proximity, parallelism and coverage of linear edge segments detected in the image while photometric factor measures shadow evidence based on intensity levels in the vicinity of the polygon. The human operator interacts with this automatic scoring by setting a threshold to highlight buildings poorly assessed by image geometrical and photometric features. After image inspection, the operator may decide to mark the polygon as changed or to update the database, depending on the application.

Keywords: Remote sensing, building verification, geographical database, linear segment matching, shadow

1. INTRODUCTION

Damage assessment, change detection or geographical database update are traditionally performed by experts looking for objects in images, a task which is costly, time consuming and error prone. Automatic solutions for building verification are particularly welcome but suffer from the wide variety of object appearance in images due to illumination and perspective changes. Typical difficulties concern the large change in image intensities, occlusions and shadow areas. On the other hand, semi-automatic procedures benefit from the rapidity of automatic image analysis while limiting human intervention to doubtful cases.

A very good example of semi-automatic development is described in Flamanc et Al [1] and concerns the platform developed by MATIS (laboratory of Institut Géographique National, France) to combine automatic and user assisted tools for building reconstruction (3D city models) using aerial and 2D ground maps. In the field of built-up area verification, the Signal and Image Centre of the Belgian Military Academy has developed an automatic system to estimate the urbanization changes from SPOT5 images relatively to the Belgian National Geographic vector database (ETATS project, Lacroix [2]). This development does not address change detection of individual buildings.

In this paper we present a semi-automatic approach to assess the presence of buildings from geometrical and photometric cues detected in an aerial image. For each polygon of the vector database representing a building, a score related to the confidence in the building presence is assigned. This score combines a geometrical part associated to the proximity, parallelism and coverage of linear edge segments detected in the image with a photometric factor measuring shadow evidence based on intensity levels in the vicinity of the polygon.

The approach is semi-automatic in the sense that the human operator interacts with the automatic polygon scoring by setting a threshold to highlight buildings poorly assessed by image geometrical and photometric features. After image inspection, the operator may decide to mark the polygon as changed or to update the database, depending on the application.

The rest of the paper is organized as follows. Section 2 describes the general approach. Section 3 details how database polygons are scored from geometrical and photometric image features. Results are presented in section 4. Section 5 concludes the paper and presents perspectives for future work in building verification.

2. METHODOLOGY

The overview of the approach is presented in Figure 1. In our application, an airborne image is available with a vector database consisting of polygons including buildings and roads. After some vector filtering based on geometry to eliminate roads, the objective of the image processing part is to automatically score remaining polygons thanks to image features like linear segments and shadow cues. Finally, the operator selects a threshold, possibly thanks to the score histogram, to highlight buildings poorly assessed by image features. The separation into data, image processing and visual inspection is further described in the next paragraphs and detailed in the next sections.

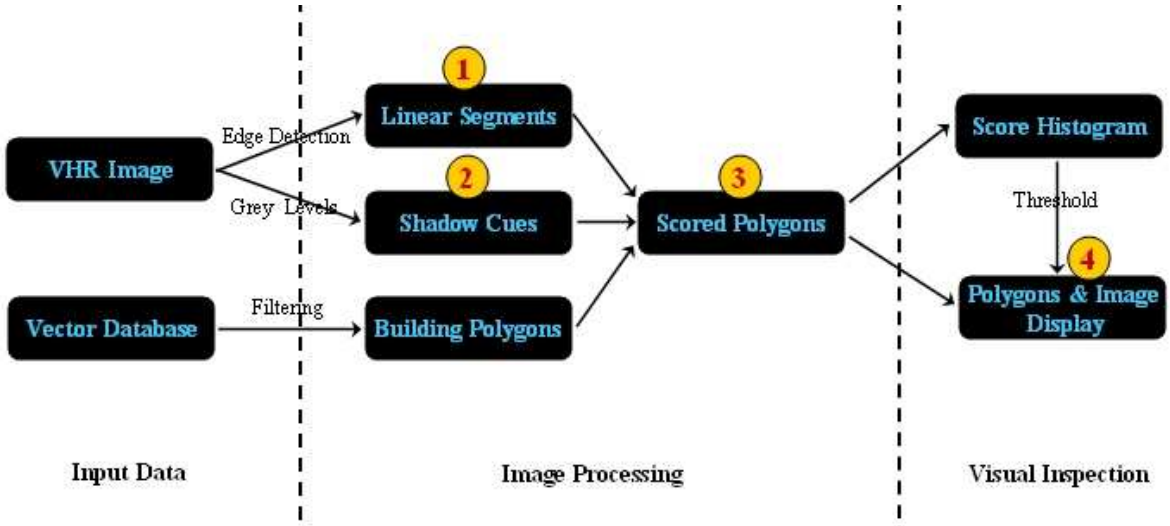

Figure 1: Approach overview

The VHR image is captured from an airplane and scanned at a corresponding ground resolution of 0.3m. As Belgian National Geographic Institute is still using black and white images from scanned camera pictures, we could not make photometric analysis based on colour or vegetation (thanks to near infrared) and concentrated our research on the useful shadow feature. In the near future, we hope to benefit from the better quality (less noise) of multi-spectral digital cameras. The fact that the vector database that we received was not separated into layers necessitated filtering to extract polygons related to buildings as explained in section 4.1. This complication was easy to handle and led to the presence of interesting 'impostor' tests.

In aerial images, buildings exhibit distinctive linear features like their borders, roof ridges or shadow profiles. Building elevation normally leads to surrounding dark areas due to shadows, providing strong additional support to building presence. Our approach consists in combining scores obtained from analyzing detected linear segments and grey levels in the vicinity of polygon edges. Each polygon edge is scored based on the proximity, parallelism, mutual coverage of image segments and polygon edges. Segments inside polygons provide additional score based on parallelism or distance to a polygon vertex to account for roof ridges. The shadow score of a polygon is based on the outside grey level for polygon edges compliant with the shadow direction. This direction and the shadow grey range are estimated for the entire image from the grey level histogram of pixels close to polygons. Shadow and linear edge scores are fused into a combined polygon score by a weighted sum. Details (points '1', '2' and '3' in Figure 1) are given in the next section.

The human operator interacts with the automatic scoring through the selection of the score threshold, so that building polygons unsupported by image features are superimposed in red on the image. The operator may then inspect the image for building verification and depending on the application, update the geographical database or mark the area as changed or damaged. In the case of a construction not in the database, linear segments not assigned to any building may draw the operator's attention and can help him to delineate or check delineation of the new building. Point '4' of Figure 1 is detailed in section 4.

Combining multiple cues for building detection in aerial images is not new. In Müller et Al. [3], geometric (size and form), photometric (mean and most frequent hue) and structural (neighbourhood, shadow) features are integrated. In the survey of Shufelt [4], the author compares four systems for building extraction in aerial images with many different parameters and test images. All of them support shadow modelling. We already described in Beumier [5] the verification of buildings from geometrical cues. We show here the advantage to include shadow information. Later, we intend to integrate multi-spectral cues like colour and vegetation index, combine confidence scores obtained from several images and include height information derived from stereo pairs.

3. SCORING POLYGONS

As outlined in the previous section, building polygons of the vector database are scored according to linear segments detected in the image and surrounding grey levels in the shadow direction. We present in three subsections the specific scoring from supporting linear edge segments and from surrounding dark grey-levels and their combination into a unique confidence score.

3.1 Scoring polygons from linear edge segments

Building polygons of the vector database are scored according to the mutual coverage, the proximity and the parallelism of supporting linear edge segments detected in the image. We summarize in this subsection the way linear edge segments are detected in the image and how these are used to score the polygons as detailed in Beumier [5].

3.1.1 Linear edge segment detection

The first image processing step consists in detecting image linear segments since they mostly relate to human structures and in particular buildings. Pixels of sufficient gradient magnitude are linked into connected segments based on their gradient orientation. Short and curved segments are rejected according to the length and straightness measured by the moment of inertia (refer to Beumier [7]).

The parameters involved in this step are the minimum gradient magnitude to start a segment, the minimum segment length and minimum segment straightness. They were all set to common sense values and have little influence on the results. The image quality is much more important than the tuning of those parameters.

3.1.2 Polygon scoring from edge segments

All the detected image segments are positioned relatively to each polygon edge thanks to four measures: the signed distances d1, d2 (< 0 inside the polygon) and the indices of segment ends along the polygon edge i1, i2 (see Figure 2a). From these values, we define the measures:

a) MC, the mutual coverage, equal to i2-i1, the clipped projection of the segment on the polygon edge;

b) D, the average signed distance of the segment to the polygon (d1+d2)/2;

c) P, factor of parallelism, equal to abs(d2-d1)/MC;

d) R, the confidence value for the segment/edge pair, equal to (50*MC)/edge_length. The factor 50 (compared to 99) allows for the contribution of several segments overlapping the same polygon edge.

According to the positioning of a segment relatively to a polygon edge, which is measured by the values P and D, three cases may occur as depicted in Figure 2b:

A) the segment is close to the polygon and parallel to one edge (SegA);

B) the segment is inside the polygon and parallel to one edge, likely to correspond to a roof ridge (SegB);

C) the segment is inside the polygon with one end close to a vertex, likely to be a diagonal roof ridge (SegC).

Each case increases the edge score by the value R and the final score is saturated to 99 (maximal confidence).

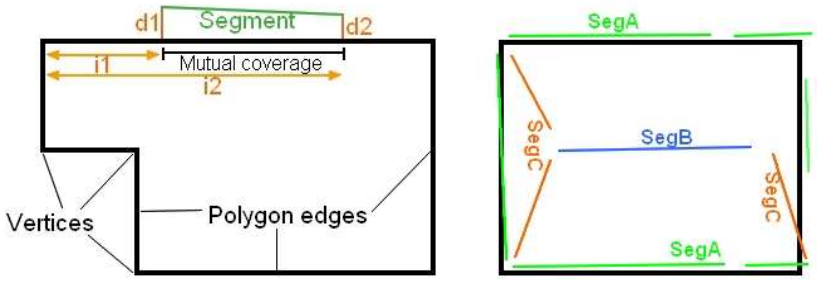

Figure 2: a) Definitions for the building polygon and b) cases of segment matching

The polygon geometrical confidence score is the average of individual polygon edge scores:

Polygon geometrical score = ∑ edge_score[i] / num_edges

In the current implementation, the thresholds related to parallelism P and distance D (used to separate the three cases) are quite flexible to allow for registration imprecision or image variability. Distance tolerance is typically 5 pixels (1.5m). The angular tolerance related to P is typically 0.05 (3°) but may be larger for images with gable roofs and large parallax leading to segments not parallel to the roof polygon as delineated by the operator when constructing the database.

To normalize the distribution of scores in the [0..99] range, we apply the following correction to the scores:

Normalized polygon geometrical score = 50.0 + (Polygon geometrical score – mean)*12.0 / std_dev

where 'mean' and 'std_dev' are respectively the mean and standard deviation of the scores. This formula transforms the distribution to have a mean value of 50 and a standard deviation of 12.0. This better fills the 0..99 range. The value 12.0 was roughly derived from the case of a normal distribution, for which the confidence interval for 3 times the standard deviation (36) contains 99.73 % of the population.

3.2 Scoring polygons from shadow

Adding shadow to the geometric analysis is particularly well suited since it captures efficiently a side effect of the elevation property of buildings and possibly discriminates them from other man-made structures like roads or parking lots. The visibility required by the optical sensors imposes that images are acquired in cloud free condition during the day. These conditions are sufficient to ensure that shadow can be exploited although its form is dependent on the sun elevation and azimuth. In the extreme case, the sun is at the zenith and no shadow is visible.

The use of shadow to support building extraction is not new. Two references have been given in section 2 but interested readers can also refer to Jin [6] where shadow has been detected to support building extraction in satellite images. In our context of building verification, the shadow outline is not necessary and we preferred to check shadow presence from the darkness of grey levels in an appropriate direction close to the known building positions.

3.2.1 Detecting shadow orientation

To assess the presence of shadow, we need to know the shadow direction. A first method consists in deriving the sun azimuth (and possibly elevation) from the image acquisition date and geographical coordinates generally accompanying the image.

A second approach, applicable in the case of building verification, considers image grey levels in the vicinity of building polygons to derive the average grey levels in each orientation. Although grey levels around buildings are not uniform all over the image, the values associated with shadow areas are much darker and reveals a clear directionality in angle dependent histograms.

For all polygons fully contained in the image limits, each polygon edge is attributed to one of the 16 orientations evenly separating the 360° range. For each point of the polygon edge, the histogram associated to the edge orientation is accumulated with the grey value of the minimum found perpendicularly to and outside of the polygon, up to a distance of 5 pixels. The histogram of each orientation is then replaced by its median value and a vector of 16 values is returned. This highlights the directionality of grey levels around the polygons as induced by shadow. In Table 1, each direction covers 22.5°, starting with the horizontal 'East' direction and continuing clockwise. For instance, direction '12' points to the north in the image.

Table 1: Median of minima values surrounding polygon edges classified in 16 directions (each covering 22.5°).

Direction	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Intensity	94	91	100	106	102	96	77	98	92	38	34	36	32	31	34	55

This vector allows first to assess that shadow is visible in the image. About half its elements should clearly be lower than the other elements, with a minimum opposite to the sun direction. Secondly, this vector helps to establish in which direction shadow will have to be looked for around polygons (orientation 9..15 (north) in the example). Finally, it allows for the definition of a shadow threshold 'shad_thres' to score the shadow confidence. By taking the average of all values, we obtained the value 70 as shadow threshold for this example.

The approach that has been adopted supposes that a sufficient proportion of polygons match buildings that are effectively present in the image. Some latitude in localization of the polygons is tolerated thanks to the minimum search perpendicularly to the polygon edges. In difficult situations for which no clear orientation is detected for shadow, the sun direction can be obtained automatically from image metadata like the coordinates and time. If not available, an acceptable solution for a semi-automatic system consists in the interactive specification of the shadow orientation from the operator through a single parameter.

3.2.2 Polygon scoring from shadows

Scoring polygons for shadow evidence is achieved by measuring the average grey level along segments parallel to the polygon edges (Figure 3).

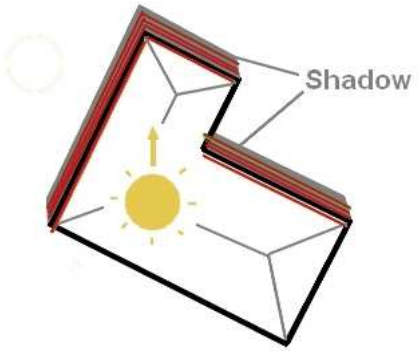

Figure 3: Shadow scoring by averaging grey levels along segments parallel to compliant polygon edges.

Each polygon edge compliant with the shadow direction (bins 9 to 15 in Table 1) contributes to the shadow score by the difference between 'shad_thres' (estimated in 3.2.1) and 'min_average_grey', the minimum value of average grey levels along five segments parallel to the polygon edge at distances ranging from -1 to 3 pixels. The polygon shadow score is the average of the scores of compliant polygon edges ('num' is typically 2 or 3):

$$\text{Polygon shadow score} = \sum (\text{shad\_thres} - \text{min\_average\_grey}) / \text{num}$$

An adaptation of the scores to have a mean of 50 and a standard deviation of 12.0 is also applied for shadow:

$$\text{Normalized polygon shadow score} = 50.0 + (\text{Polygon shadow score} - \text{mean}) * 12.0 / \text{std\_dev}$$

3.3 Combining scores

Although edge and shadow scores can be used separately by the graphical interface to highlight unsupported buildings, we also integrated both confidences into a unified score to allow for automatic evaluation. As presented in section 4, some erroneous polygons are present in the database so that impostor tests are available and ROC curve or equivalent statistics can be derived.

Since we already normalized the score distributions (mean 50 and standard deviation 12), we simply average both scores to derive the unified score:

Polygon score = (Normalized polygon geometrical score + Normalized polygon shadow score)/2.

4. RESULTS

After describing the data on which the approach has been applied and associated pre-processing, two types of results are presented. The first one describes the way the graphical interface is used to highlight unsupported polygons. The second result presents numerical figures in terms of false acceptance and false rejection as a few polygons present in the database do not correspond to any building in the image and serves as impostor tests.

Figure 4: Small part of a test image with superimposed vectors containing one phantom building and one blind alley. Notice the numerous small outbuildings.

4.1 Data and pre-processing

The approach was applied to two aerial images (ground resolution: 0.3m) of the same region used by the National Geographical Institute of Belgium for stereo restitution. A part of the images covering a rural area was used (4000x3000, about 1km^2) to limit computation time while including a large number of buildings (more than 650). Originating from scanned photographs, these noisy images were filtered with a 3x3 uniform low-pass filter. Buildings are mainly isolated houses and small outbuildings.

The area is covered by a polygon database (scale 1:10000) that mixes buildings and the road network. Separate layers could have been obtained but the blind alleys, part of the road network, are impostors interesting for evaluation. Database elements which are not closed polygons and that contain less than four vertices were filtered out.

4.2 Highlighting low scores

In this first experiment, polygons were superimposed on the image with a colour dependent on their geometrical score (shadow not used). Polygons with a score higher than a threshold are drawn in green while others are drawn in red. The operator modifies the threshold, possibly using the cumulated histogram of scores, to highlight a few percents of buildings in red. Then he or she analyses the underlying image to verify the presence of the building and updates the database or marks the building as modified, depending on the application.

Figure 5: Building verification from edge segments: polygons of verified buildings in green and unverified in red; blind alleys circled by light blue ellipses and phantom building polygons by white circles.

The use of low threshold values to highlight buildings weakly supported by edges revealed the presence of phantom buildings and blind alleys (Figure 5). Phantom buildings are polygons which do not correspond to real buildings. They may be planned or destroyed constructions or reveal that the database is more recent than the image. Here, blind alleys refer to squares ending a road ("cul-de-sac") and are present in the database that we received although they normally belong to the road network. Setting the threshold to the largest score of the six phantom buildings, five blind alleys out of seven received image support (due to road edges), and less than 4% of the buildings were falsely rejected, many of them being small (less than 30 m^2).

Introducing shadow improves the situation of blind alleys and phantom buildings since they are flat or inexistent constructions. Figures are given in the next section.

4.3 False acceptance and false rejection

Thanks to the blind alleys and phantom buildings present in the region under test, we can present a table of false acceptance and false rejection tests, similar to results produced for recognition systems (in particular, the ROC curves). In the next table, the number of falsely accepted blind alleys, falsely accepted phantom buildings and falsely rejected buildings are reported against the score value based on edge.

Table 2: Number of accepted blind alleys and phantom buildings and rejected genuine buildings versus geometrical score.

Scores	**15**	**18**	**21**	**24**	**27**	**30**	**33**	**36**	**39**	**42**	**45**	**50**	**60**	**70**
# Blind Alleys	7	7	7	7	5	5	4	1	0	0	0	0	0	0
# Phantoms	6	6	4	1	1	0	0	0	0	0	0	0	0	0
# Buildings	0	0	1	4	7	24	42	58	111	166	236	352	536	611

Setting the threshold to 30 to reject all polygons related to phantom buildings, 24 (3.7%) out of 657 buildings were erroneously rejected, many of them being small buildings (less than 30 m^2). Geometrical scores are expected to have much more potential thanks to multi-spectral images (colour or vegetation index) which were not available to us but which allow to better discriminate the roofs from their neighbourhood.

Considering the scores related to shadow, we see that blind alleys are easily discarded (low scores). With a threshold of 16, no impostor is supported, while only two small buildings are not verified. The shadow cue appeared to be very discriminative.

Table 3: Number of accepted blind alleys and phantom buildings and rejected genuine buildings versus the shadow score.

Scores	**5**	**10**	**15**	**20**	**25**	**30**	**35**	**40**	**45**	**50**	**55**	**60**	**70**	**80**
# Blind Alleys	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# Phantoms	4	2	1	0	0	0	0	0	0	0	0	0	0	0
# Buildings	2	2	2	2	9	22	40	98	190	277	402	562	656	657

The combination of edge and shadow cues (Table 4) better separates the impostor (blind alleys, phantom buildings) and genuine building score distributions. Mention however that the number of impostor tests is limited (13 polygons, 2%).

Table 4: Number of accepted blind alleys and phantom buildings and rejected genuine buildings versus the combined score.

Scores	**12**	**15**	**18**	**21**	**24**	**27**	**30**	**35**	**40**	**45**	**50**	**60**	**70**	**80**
# Blind Alleys	5	2	1	1	0	0	0	0	0	0	0	0	0	0
# Phantoms	5	3	1	0	0	0	0	0	0	0	0	0	0	0
# Buildings	0	0	0	0	0	4	10	21	61	152	313	593	656	657

Processing time is about 15 s on a 1.3 GHz laptop, with 50% for linear segment extraction, 40% for segment edge scoring and 10% for shadow scoring. The results obtained for both images of the same region gave similar results.

5. CONCLUSIONS

We have presented and detailed in this paper the successful integration of shadow and geometrical cues to highlight building polygons insufficiently supported by image data. Shadow verification requires simple and little computations while offering most of the discriminative power. Geometrical support, implemented by linear segment detection, is currently limited by the roof visibility in intensity images although multi-spectral images have large potential for roof segmentation thanks to colour and vegetation index analysis. Possible applications of the proposed semi-automatic solution are geographical information update, change detection and damage assessment of buildings for which a vector database exists.

This work represents one of the first steps to realize man-made structure verification where additional information will be gathered (road network, colour, height from stereo disparity) from several aerial or satellite images (spectral images, stereo pairs). We also plan to add to the approach the possibility to make building hypothesis from detected cues so as to detect buildings and propose candidates not yet in the database.

ACKNOWLEDGEMENTS

This study has been supported by the Belgian Ministry of Defence. Image and vector data has been gracefully obtained from IGN (Institut Géographique National) of Belgium.

REFERENCES

1. C. Flamanc, G. Maillet and H. Jibrini, "3D City Models: an operational approach using aerial images and cadastral maps", *ISPRS Archives, Vol. XXXIV, Part 3/W8*, Munich 17-19 Sept 2003, 53-58.
2. V. Lacroix, M. Idrissa, A. Hincq, H. Bruynseels and O. Swartenbroekx, "SPOT5 pour la détection d'urbanisation", *Revue Française de Photogrammétrie et de Télédetection*, Numéro 178, 2006.
3. S. Müller and D. Zaum, "Robust Building Detection in Aerial Images", CMRT05, IAPRS, Vol. XXXVI, Part 3/W24, Vienna, Austria, Aug 29-30 2005, pp 143-148.
4. J. Shufelt, "Performance Evaluation and Analysis of Monocular Building Extraction From Aerial Imagery", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 21, NO. 4, Apr 1999, pp 311-326.
5. C. Beumier, "Building verification from geometrical features", 27th *Earsel Symposium - Geoinformation in Europe*, Bolzano Italy, 4-7 June 2007.
6. X. Jin and C. Davis, "Automated Building Extraction from High-Resolution Satellite Imagery in Urban Areas Using Structural, Contextual and Spectral Information", *in EURASIP Journal on Applied Signal Processing,* 2005:14, pp 2196-2206.
7. C. Beumier, "Straight-line Detection Using Moment of Inertia", *IEEE International Conference on Industrial Technology 2006 (ICIT2006)*, Mumbai, India, Dec 15-17, 2006.