Methods for exploiting the relationship between buildings and their shadows in aerial imagery

R. Bruce Irvin
David M. McKeown

December 25, 1988

CMU-CS-88-200 2

Digital Mapping Laboratory
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

This research was primarily sponsored by the U.S. Army Engineering Topographic Laboratories under Contract DACA72-87-C-0001 and partially supported by the Defense Advanced Research Projects Agency, DoD, through DARPA order 4976, and monitored by the Air Force Avionics Laboratory Under Contract F33615-87-C-1499. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Engineering Topographic Laboratories, or the Defense Advanced Research Projects Agency, or of the United States Government.

Table of Contents

List of Figures

Abstract

This paper describes computational techniques for utilizing the relationship between shadows and man-made structures to aid in the automatic extraction of man-made structures from aerial imagery. Four methods are described that perform the prediction of structure shape, grouping of related structures, verification of individual structures, and structure height estimation. In each method the relationship between structure and cast shadows is exploited in a unique fashion. Key issues involve the accurate localization of the structure/shadow boundary and the shadow edge, and attribution of shadow segments to structure hypotheses. We present several examples that show how each method is used within the task of building detection, delineation, and height estimation.

1. Introduction

The extraction of man-made structures in complex urban scenes often elude traditional image processing techniques. Such structures often violate many assumptions such as uniform pixel intensity, strong edges, and predictable shapes found in traditional image segmentation systems. One important property of man-made structures, particularly buildings, is that they normally have height and, when illuminated by the sun, they cast shadows. This paper shows how the relationship between man-made structures and cast shadows is easily exploited to aid and improve extraction of the man-made structures from the aerial imagery. While shadow detection and analysis is difficult, their extraction from aerial imagery is often easier than the analysis of the buildings that cast them. Shadow textures and colors are not affected by the texture and color of the structures that cast them. While shadow intensity can vary greatly depending on the color of the surface onto which they are cast, shadows are usually among the darkest regions in images and their extraction is often amenable to simple image processing techniques. Finally, as we will see, shadows provide a great deal of information concerning the structure of man-made object without the necessity of developing an explicit model of the structure. Given the great variability of building shapes, heights, and structure, it is difficult to imagine how an explicit model-based approach to building interpretation would function without first having some scene analysis cues to constrain search while matching models. This work focuses on the generation and refinement of such scene analysis cues.

1.1. Previous Work

The use of shadows to perform height determination of objects in aerial imagery has a long tradition in photogrammetry and manual photo interpretation techniques [1]. The primary assumptions are that the objects being measured are vertical, that the shadows are cast from the top of the object and not the sides, and that the shadows fall on open, level ground. The use of shadow mensuration and its inverse, the determination of the sun elevation angle using objects of known height, are techniques that are taught to all novice photo interpreters.

The computer vision community has also recognized that shadows are a valuable source of information in the interpretation of monocular aerial imagery. Nagao and Matsuyama [12] used shadow regions detected in color infrared aerial imagery by simple thresholding to guide the extraction of adjacent regions, presumed to be buildings. Lowe and Binford [5] used the correspondence between shadow edges and geometric edges in an airplane model to recover a coarse 3-dimensional model of the airplane. Shafer [15] introduced a general theory of the relationships between shadows and surfaces for line drawings using line labeling and a gradient space representation for surfaces. However, no

Figure 1-1: DC37405: Suburban Housing Scene

examples of shadow extraction in real imagery are presented and assumptions concerning knowledge of the shape of the surface onto which the shadow is being cast is unrealistic for remotely sensed imagery. Huertas and Nevatia's work in building analysis used structural analysis, by the generation of lines and corners in the image. It then verified the local organization using shadow constraints [4]. It was the first system to utilize shadows to interpret "object" corners and sides in complex aerial imagery explicitly accounting for the shadow geometry.

Our work extends that of Huertas and Nevatia, which used shadows to verify the consistency of structure hypotheses as they were being generated, in several ways. First, we decouple the process of hypothesis generation from that of verification using shadow information. This is important because there may be multiple techniques for building hypothesis generation [10] based upon edge/line analysis [3], region analysis [7], or stereo analysis [8]. Each technique will generate a different set of hypotheses with inherent strengths and weaknesses. As we have seen, the use of multiple methods [11] in aerial image analysis greatly improved accuracy and robustness of the scene analysis system. Rather than embedding shadow analysis techniques within each method, we believe that an independent shadow verification capability is required.

Second, we define two new sources of information based upon shadow analysis, structure prediction and structure aggregation. Structure prediction involves the automatic detection of shadow regions and the hypothesization of structures that could generate such shadows. Structure aggregation groups together hypotheses that can be explained by the same shadow region. This is in recognition of the fact that scene analysis techniques often fragment buildings into multiple structure hypotheses.

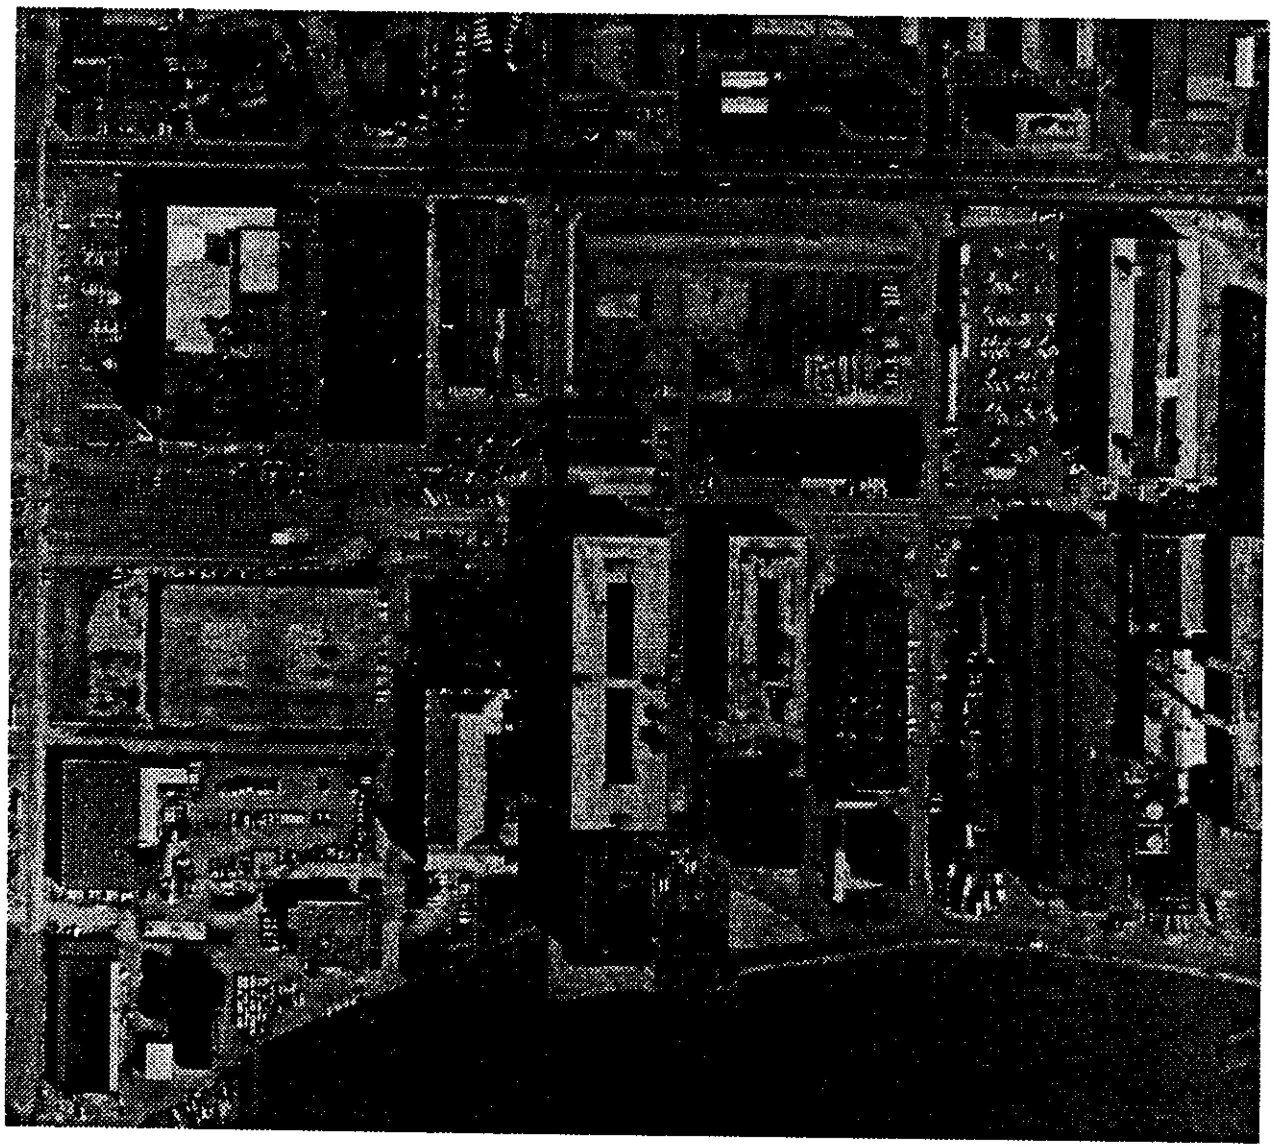

Figure 1-2: DC38008: Industrial Area Scene

In the remainder of this paper we describe a set of techniques to automatically extract and analyze shadow information in aerial imagery. In Section 2 we describe SHADE and GROUPER, two programs that focus processing attention on specific areas of an image that are likely to contain man-made structures. The two main tasks in this domain are prediction of structure shape and the aggregation of structure hypotheses. In Sections 3 and 4 we describe SHAVE (SHAdow VErification), a program that predicts and delineates shadows for building hypotheses and assigns confidence values for each hypothesis. SHAVE also measures heights of buildings using their mean shadow lengths.

For each program we present examples showing steps of execution and completed results. We have chosen high resolution images that illustrate many of the problems common to building extraction and shadow analysis. Figures 1-1 and 1-2 are two such scenes that we have used for all of the examples in this paper. Figure 1-1 is a suburban housing area and Figure 1-2 is an industrial area in Washington D.C. We have found that use of scenes with a variety of buildings, rather than isolated buildings, greatly increases the task complexity. Our methods do not correctly delineate or verify every building in every scene. However, these programs are not meant to exist as a stand-alone building extraction system, rather they are components of a cooperative building extraction system. Our cooperative methods paradigm assumes that no single method can perform robust analysis across a variety of imagery. Shadow analysis is one of several methods which add information to the overall interpretation.

2. Shadow extraction

This Section describes two programs, SHADE and GROUPER, that extract shadows from aerial imagery to focus attention on specific areas that are likely to contain man-made structures. Each program begins its processing by extracting possible shadow regions. Shadow region extraction is accomplished using a simple set of image processing techniques including image smoothing, thresholding, and connected region extraction. Initially the intensity image is smoothed using an edge preserving smoothing algorithm. Thresholding is performed using a single image-wide intensity threshold. After the image is thresholded, the four-connected regions are extracted and labelled as separate "dark" regions. These "dark" regions are a collection of true shadow regions as well as other "dark" surfaces/objects in the image. At this point, very small regions are discarded.

The image-wide threshold is provided as a result of monocular analysis for building hypothesis by the BABE system [3] while it evaluates plausible structure hypotheses based upon line/corner/structure generation. BABE examines those areas immediately adjacent to the structures to determine an appropriate shadow intensity pixel sample. The actual values produced by BABE are the mean and standard deviation of of those pixels immediately adjacent to buildings. SHADE and GROUPER use the mean plus one standard deviation as the shadow intensity threshold.

BABE also produces an estimate of the sun direction based upon a statistical analysis of its shadow/structure boundaries. However, this information and the sun angular elevation can also retrieved on a per image basis from the CONCEPTMAP database [6, 9]. The sun angular elevation angle can be determined based upon knowing the time of day that the image was acquired, the latitude of the imagery, and the sun's declination corrected to Greenwich Mean Time [1]. The sun angular elevation is required to calculate the actual structure elevation based upon shadow length.

2.1. Shadow analysis to predict the shape of man-made structures

The goal of SHADE is to predict the location and shapes of man-made structures given their shadows. To do this SHADE examines the sun-side of each "dark" region, determines if it is indeed a shadow of a building, then using the shape of the shadow, predicts the shape of the corresponding building.

Quite often it is the case that shadow extraction produces regions that have shapes characteristic of

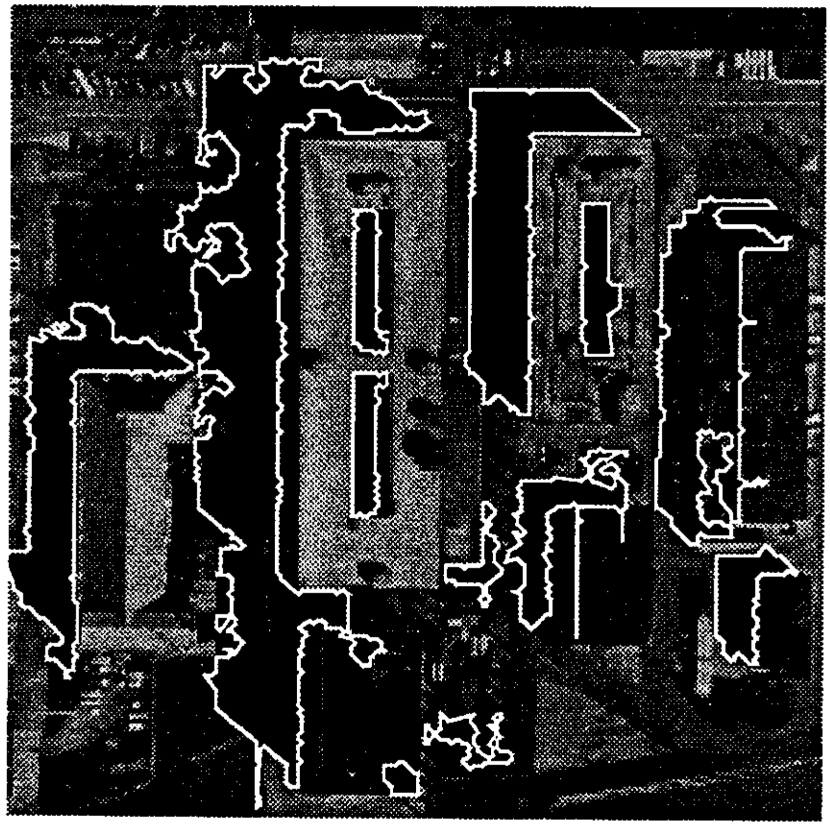

Figure 2-1: Shadow regions found by SHADE

shadows cast by man-made structures. In particular, the characteristic "L" shaped shadows cast by rectangular buildings give strong clues to the shapes of their corresponding buildings. SHADE exploits the relationship between man-made structures and cast shadows by identifying nearly right angle corners in the shadow region boundary. In the following section we describe how SHADE forms structure hypotheses as logical extensions to the delineated shadow regions.

2.1.1. How to predict structure shape using shadows

Figure 2-1 shows the result of shadow extraction on a subarea of Figure 1-2 containing several complex buildings and their shadows. SHADE examines each edge of each smoothed shadow region and tests whether or not it is on the sun-side of the shadow. This test involves back projecting the edge's midpoint toward the sun along the sun direction vector a nominal distance, normally one pixel. If the projected point falls outside of the original shadow region, then it is labelled as a sun-side edge. Otherwise, it is labelled as a non-sun-side edge. Because there may still be noise in the shadow edges or complex structural behavior along the building/shadow edges, SHADE uses a sequence finder [2] to locate the imperfect sequences of sun-side edges in each shadow region.

Figure 2-2 shows the result of the determination of those portions of the shadow region that represent the shadow building boundary. After localizing the shadow/building edges, SHADE searches for instances of the characteristic "L" shapes. To do this it breaks these edges into nearly straight-line segments [2] and passes them through a corner finder. The corner finder returns a list of all the corners in the edge data that satisfy nominal range and angle constraints. Because the only corners of interest are those whose concave sides are oriented toward the sun, SHADE rejects all corners whose bisecting vectors do

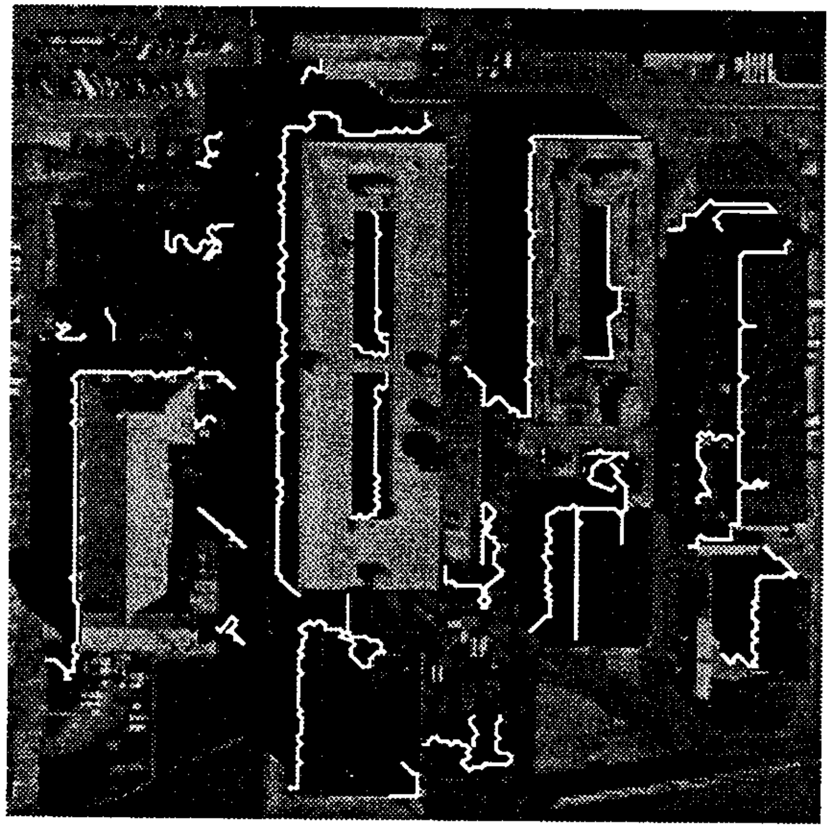

Figure 2-2: Shadow/Building Edges

not fall within the same quadrant as the sun. Figure 2-3 shows the result of generalizing the region boundaries into segments bounded by points of high local curvature, often corresponding to building corners.

The final processing step of SHADE simply extends the building/shadow corners into parallelograms and rejects very small regions. These parallelograms are hypotheses for occurrences of man-made structures. Because buildings have height they often occlude parts of their shadows in perspective projections. As a result the actual delineation of the buildings may be displaced. Figure 2-4 is the completed SHADE analysis. It is evident that while further processing is necessary to precisely delineate the buildings and interprete the various 3-dimensional roof structures, SHADE provides an important model-independent detection and delineation.

2.1.2. Results of hypothesis prediction on a complete image

Figure 2-5 shows the results of running SHADE on a complete industrial area scene. Of the 34 hypotheses returned by SHADE, 15 are complete detections, 13 are partial detections, 5 are intersections, and only one is a false alarm. However, 10 buildings are missed. Partial detections are hypotheses that are correctly aligned with the actual building and contain at least 50% of the building area. Intersections are partial detections with less than 50% overlap with the building. At least 80 percent of the hypothesis area must be within the actual building area in order for it to be scored as a complete, partial, or intersection detection.

Most of the misses and partial detections can be attributed to SHADE's relatively simple structure model.

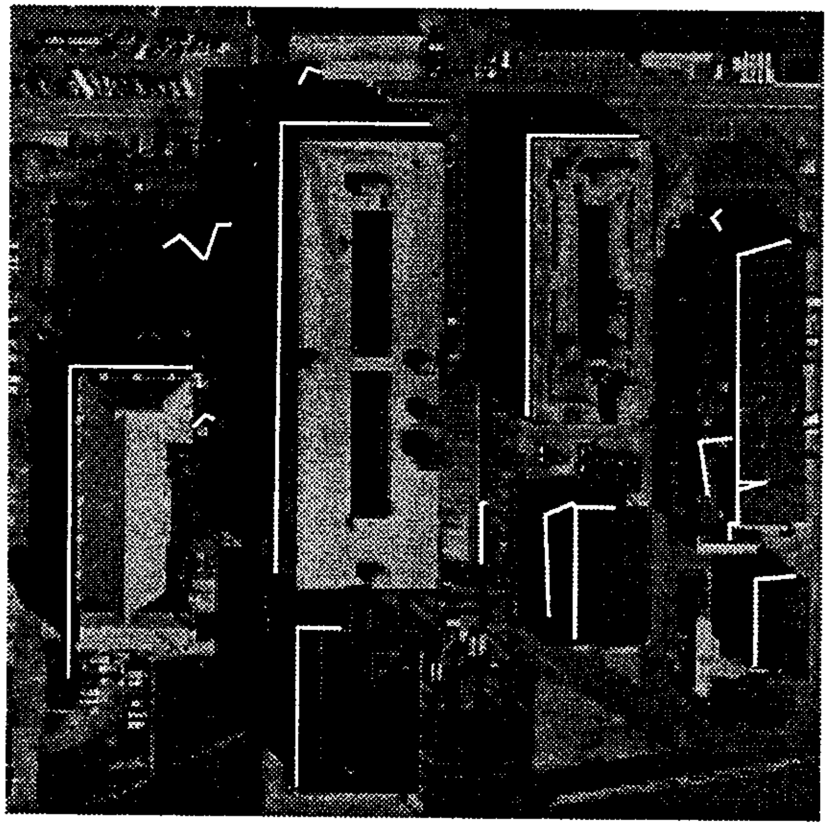

Figure 2-3: Generalization of Shadow/Building Edges

For instance, SHADE assumes that a building's shadow may be extracted as a single connected region and does not attempt to use separate shadow fragments to predict single structures. Therefore, buildings such as the one in the bottom center of Figure 2-4 are not extracted properly. Also, complicated shadow/building edges that do not perfectly fit the "L" pattern are often misinterpreted. The large buildings in the upper left and upper right of Figure 2-5 have large corners within their shadow edges that cause SHADE to terminate the shadow/building edges prematurely.

SHADE would no doubt benefit from a structure model that includes arbitrary shapes. For example, collections of circular storage tanks could not be predicted using the current method of corner analysis. However, the prime rationale for SHADE is to *suggest* building locations to higher-level processes. Within this context it is not clear whether additional analysis should be performed by SHADE or whether the appropriate strategy should be to rely on alternative sources of information. In the following Section we see how shadow analysis can be used to aggregate or group hypotheses produced using other hypothesis generation techniques.

2.2. Shadow analysis to perform structure grouping

It is often the case that building/shadow edges do not predict complete shapes of buildings as is assumed in SHADE. This is certainly true when a building is oriented parallel to the sun direction angle or when shadows are occluded by other elements of the scene. However, even in these cases the building/shadow boundary still contains important information. In this section we describe a technique to focus processing attention on the areas adjacent to the sun-side of the shadow in order to perform

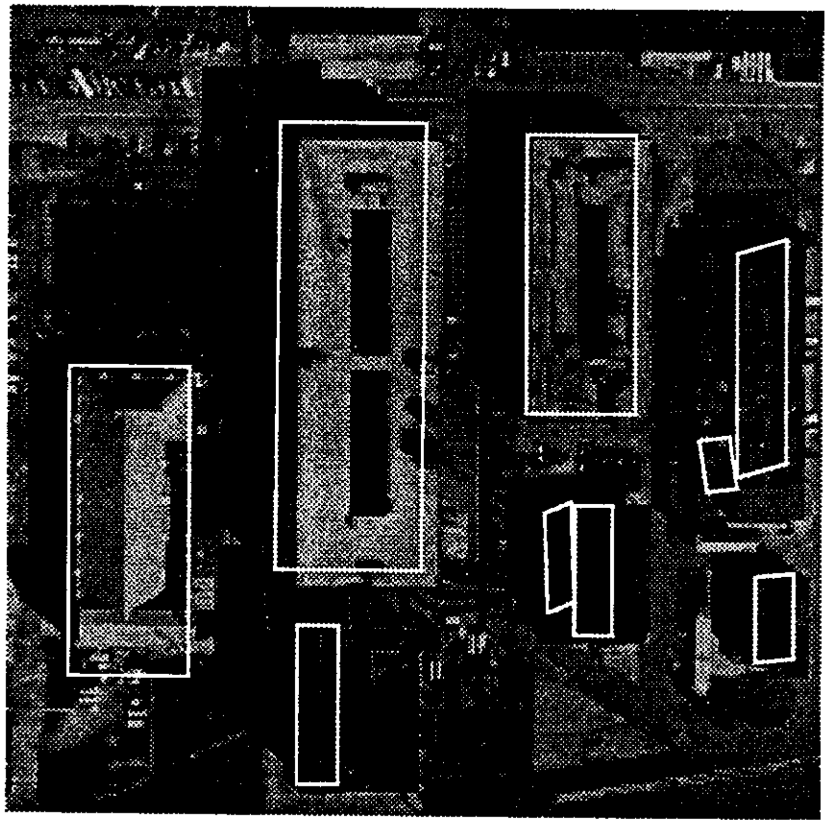

Figure 2-4: Building Hypothesis Results

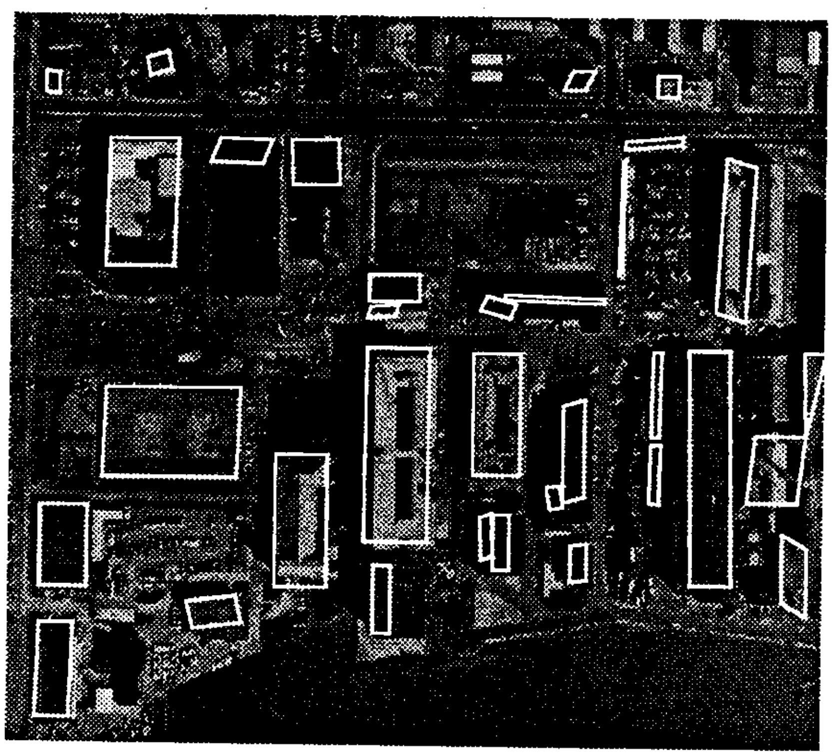

Figure 2-5: Complete Building Hypothesis Generation for DC38008

grouping of structure hypotheses. Such grouping is necessary to overcome segmentation fragmentation of complex buildings. Although the essential characteristics of the structures are preserved when fragmented, it is difficult to piece the fragments together to form complete structures. However, by relating several fragments to a shared shadow, they may be fused into one complete structure.

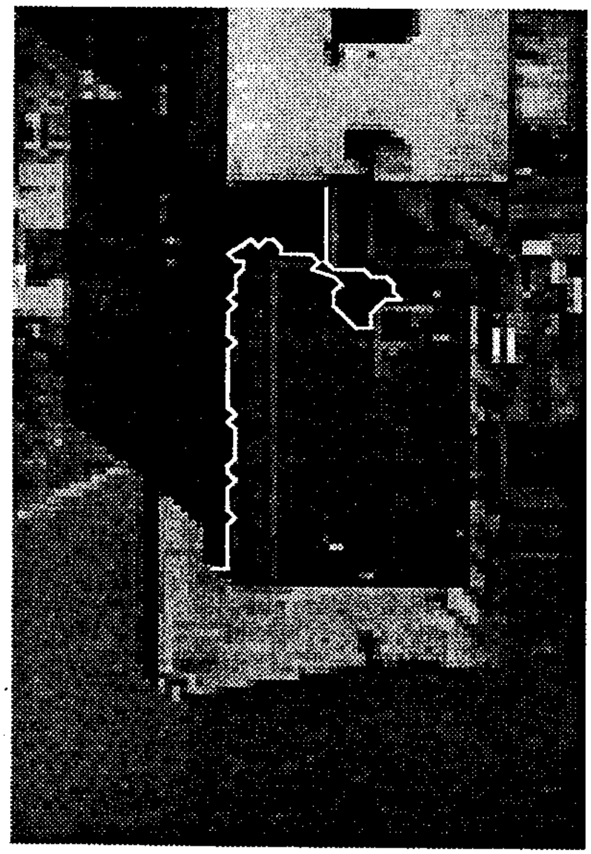

Figure 2-6: A Shadow/Building boundary produced by SHADE

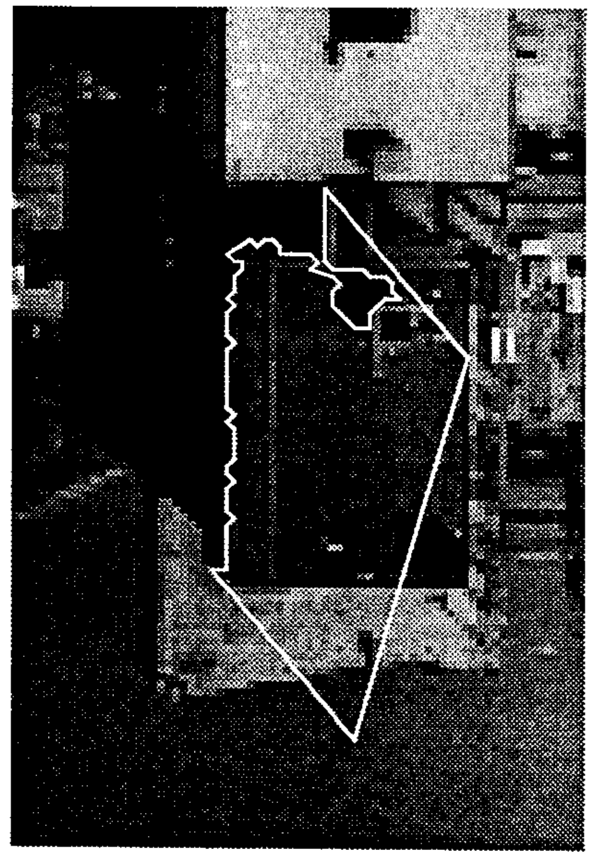

Figure 2-7: Area-of-Interest generated by GROUPER

GROUPER exploits the relationship between man-made structures and cast shadows by declaring "regions of interest" that lie on the sun-side of shadow regions. For each such region, GROUPER searches through lists of building hypotheses provided by other scene analysis techniques and groups the hypotheses according to whether or not they overlap the same region of interest and therefore share the same shadow region.

2.2.1. How to group structure hypotheses

GROUPER begins its processing with the sun-side shadow/building edges produced by SHADE as previously illustrated in Figure 2-2. It does not calculate generalized shadow/edge boundaries and perform verification based upon corner orientation and angle, as in SHADE, but rather assume that all of the dark regions will generate one shadow/building edge. GROUPER back projects the endpoints of the shadow/building edge a distance proportional to the shadow length along the sun direction vector. These projected points are joined to close the region of interest. Figure 2-6 shows a single shadow/building edge segment. Figure 2-7 shows the result of back projecting the shadow/building edge endpoints toward the sun along the sun angle. The region-of-interest constructed in Figure 2-7 is then intersected with hypothesized building regions generated by BABE. Figure 2-8 shows the set of building hypotheses within the general area of the region of interest. Of these possible regions, only those in Figure 2-9 are actually found to overlap the region-of-interest and form a group of hypotheses.

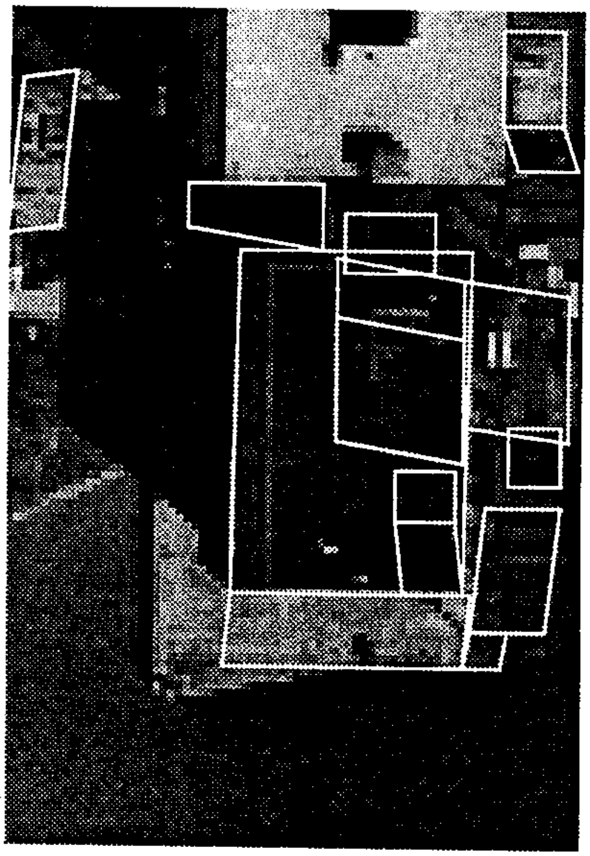

Figure 2-8: Building Hypotheses produced by BABE in the Area-of-Interest

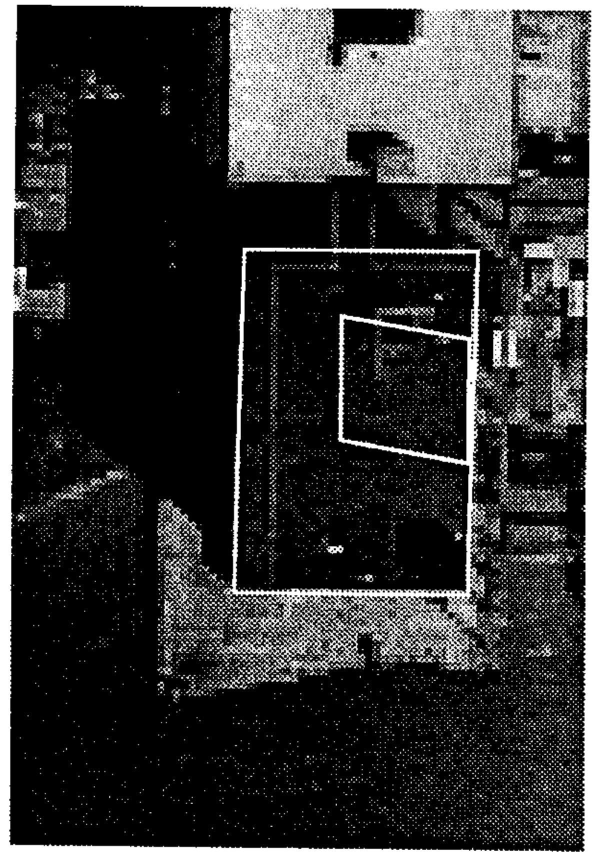

Figure 2-9: Building Hypotheses consistent with GROUPER Area-of-Interest

Overlap is determined by the percentage of the hypothesis region contained within the region-of-interest. We currently use 75% overlap as a selection cutoff. The rationale is that we are interested in grouping together hypotheses within close physical proximity that also appear to fall within an area that shadow analysis predicts should contain all or part of a building. These groups can then be processed using a more detailed and expensive evaluation procedure. As we see in Figure 2-9 only two of the fifteen BABE hypotheses are grouped with respect to the shadow/building boundary in Figure 2-6.

2.2.2. Results of hypothesis grouping on a complete image

The ability to analyze and group building hypotheses generated by BABE allows us to perform a parallel analysis and evaluation of these hypotheses. That is, we can greatly reduce the number of hypothesized buildings by requiring that hypotheses be supported by an area-of-interest generated by a shadow/building boundary. Figure 2-10 shows all of the 503 hypothesized structures generated provided by BABE before it performs its own hypothesis verification step. Figure 2-11 shows all of the 54 shadow/building boundaries produced by SHADE in image DC38008.

Figure 2-12 shows the results of running GROUPER using the BABE and SHADE results shown in Figures 2-10 and 2-11 on the complete image DC38008. Of the 503 original BABE hypotheses 132 are grouped into 37 resultant aggregates. Note that most of the erroneous hypotheses evident in Figure 2-10 are not

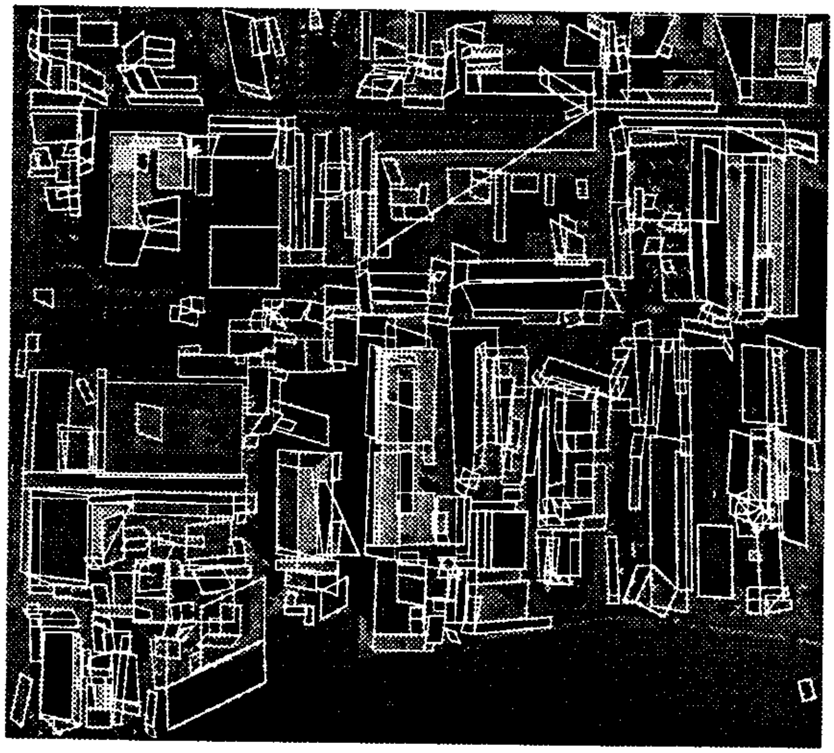

Figure 2-10: All Building Hypotheses Generated By BABE in DC38008

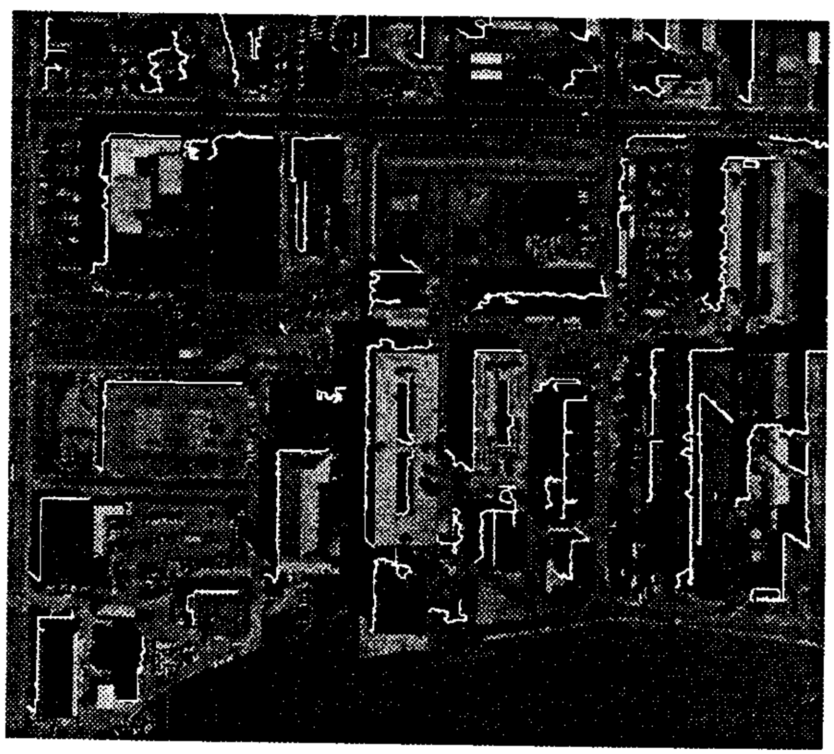

Figure 2-11: All Shadow/Building Boundaries Generated By SHADE in DC38008

aggregated into any building group. However, it is also the case that a number of hypotheses that appear to cover actual buildings are also discarded.

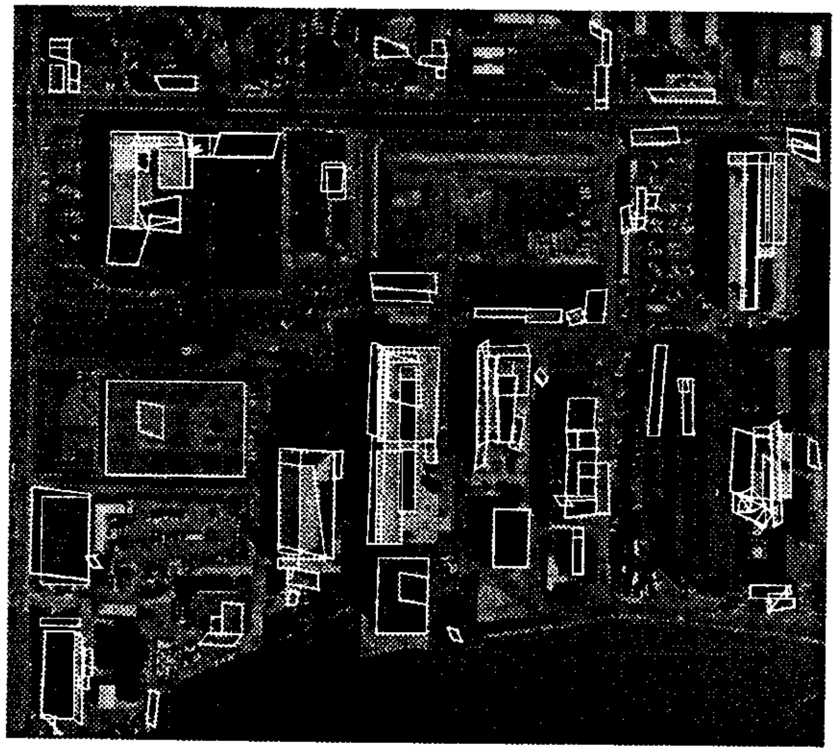

Figure 2-12: All Building Groupings Generated in DC38008

3. Shadow analysis to perform structure verification

Hypothesis verification is an important component for any scene analysis system. One way to verify regions that are hypothesized to be buildings is to verify that these hypotheses cast shadows in the image. Whereas man-made structures such as road intersections or parking lots as well as accidental alignments between objects may exhibit rectangular structure similar to that of a building, we do not expect these to cast shadows consistent with the object shape. Shadow verification is of particular importance when hypotheses are generated by monocular image analysis because it ameliorates problems due to lack of 3-dimensional information. However, shadow verification is also useful to evaluate and refine the boundaries of hypotheses generated by stereo analysis since the position of object boundaries are often displaced by the stereo matching process. This section describes SHAVE, a program that rates hypothesized building regions based on the extent to which they cast shadows. SHAVE searches for "dark" pixels starting with the edges of the regions that face away from the sun. Again, "dark" pixels are those that meet the same image specific intensity range derived by BABE that was used in SHADE.

There are three stages used by SHAVE in evaluating a region based on its shadow. First, it determines which edges of the region should be adjacent to a shadow. Next it delineates the shadow region and finally it scores the region based on the quality of the shadow.

To identify the building/shadow edges of a hypothesis region SHAVE uses virtually the same method that is used by SHADE. For each midpoint of every edge in the region, it translates the midpoint a nominal distance along the sun direction vector. If the translated point falls outside the original region then the midpoint's edge is considered to be a shadow edge.

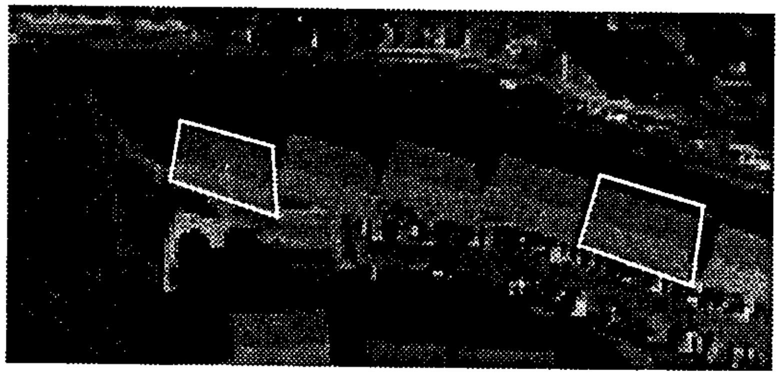

Figure 3-1: DC37405 Area-of-Interest With Two Building Hypotheses

Figure 3-2: Shadow Detection and Delineation For Two Building Hypotheses

After SHAVE generates a list of the building/shadow edges, it delineates the shadow region. Walking from the building/shadow edge along the sun direction vector, SHAVE classifies each pixel as "dark" or "bright" with respect to the estimated shadow intensity. A sequence finder [2] is used to find the imperfect sequence of "dark" pixels that correspond to the shadow region. The far end of that sequence is call the shadow terminator. This walk is performed for every point along each building/shadow edge. The length of the shadow at each point is recorded as the number of steps taken during the walk from the building/shadow edge to the shadow terminator.

After the shadow region is delineated, SHAVE scores the hypothesized building region and each of its shadow edges as a function of the mean (μ) and standard deviation (σ) of the component shadow lengths.

$$\mathrm{Score}(\sigma,\mu) = 0.0 \;\ldots\ldots\; \{0.0 <= \mu < N\}$$
$$\mathrm{Score}(\sigma,\mu) = 1.0 - \sigma/\mu \;\ldots\ldots\; \{\mu >= N\}$$

The term (σ/μ) represents the coefficient of variance for the shadow length. This term should be very small for shadows of nearly constant length and larger for shadows that exhibit a lot of variability. Shadows with a mean length less than 2 pixels (N=2) are ignored. Since our purpose is to perform a *relative* ranking of hypotheses based upon shadow verification we are not concerned with an absolute quality measure. The ranking score is calculated for each building/shadow edge separately because it is often the case that one or more of the candidate edges do not have detectable shadows. This is the case when a building is adjacent to another building as in the rightmost building hypothesis in Figure 3-1. Note that due to the structure of the townhouses, only the leftmost building has shadows completely consistent with both of its building/shadow edges.

By computing scores for each shadow/building edge we also get additional information to predict when shadows may be occluded. In Figure 3-2 we see the shadow boundaries generated for each edge of the hypothesized buildings. Note that the rightmost building has a minimal shadow attributed to it on one of the two sides that SHAVE predicts should have a significant shadow. Normally, if this were an isolated building, we would expect a shadow similar to the two exhibited by the leftmost building. While such information is not utilized by SHAVE it is recorded and is available to an analysis component that could recognize these scores as an important anomaly and attempt to find a reason for the lack of shadow along the vertical edge (ie. the other building).

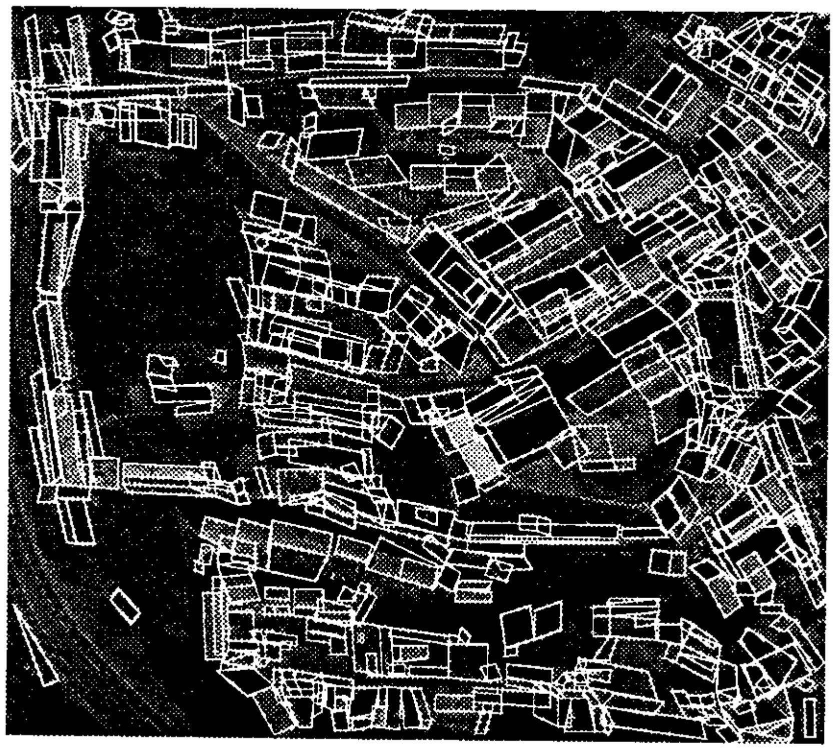

Figure 3-3: All Building Hypotheses For DC37405

3.1. Results of shadow verification on DC37405

Figure 3-3 shows the utility of SHAVE in processing large sets of hypotheses. This scene includes 605 building hypotheses generated by BABE for the complete image DC37405. As we have described SHAVE rank ordered these hypotheses based upon their consistency with casting a shadow given the sun direction. Figure 3-4 shows the best 15% of the 605 original regions, or the top 94 regions. These regions correlate quite well with actual buildings. Thus SHAVE can be used to greatly reduce the search through the hypotheses generated by BABE by examining hypotheses with the highest likelihoods first.

Figure 3-4: Building Hypotheses Verified By SHAVE

4. Shadow analysis to derive height estimation

For applications such as cartography, land use studies, urban planning, and flight simulation it is often necessary to have the height of man-made structures in addition to their geo-location. Digital stereo matching using various correlation techniques has been used to automate the generation of digital elevation models primarily in areas not characterized as urban or suburban. There are a variety of reasons why automated stereo matching techniques are not reliable in urban areas. Difficulties in matching are due to scene complexity, especially depth discontinuities due to buildings and other manmade objects, repetitive patterns and textures, and scene occlusions. Besides these problems in automatic measurement of building height, in many situations only monocular views of the scene may be available, or the imaging conditions may preclude good stereo geometry even if there is overlapping coverage. Therefore, alternatives to direct methods involving stereo mensuration are of interest. In this section we describe a simple procedure to estimate building heights using information provided by cast shadows.

SHAVE derives heights for man-made structures by applying a well known trigonometric relationship between the sun inclination angle and the values measured for shadow lengths [1, 14]. After SHAVE has measured the mean (μ) shadow length of a building in image space (pixels), it projects the distance into ground space (meters) and applys the following function to determine the building's height.

$$\text{Height}(L,\Phi) = L * \tan(\Phi)$$

L -- the length of the shadow

Φ -- the sun inclination angle

As we described in Section 1.1, such techniques, whether manual or automated, assume that the building and its environment satisfy two conditions. First, the building has uniform height generally descriptive of flat roof on top of a rectangular solid. Peaked roofs and slanted roofs generally do not cause major problem to this model as long as the base-to-height of the peak or slant is small with respect to the overall height of the structure. A more significant problem are complex buildings that are composed of multi-leveled rectangular solids. Such buildings must be modeled by decomposing the structure into components of relatively constant height. The second assumption is that the building casts its shadow on a surface that is locally flat, ie., whose surface elevation does not change. While these two assumptions are rarely completely valid, they allow SHAVE to provide good *estimations* for building heights in situations where automated stereo analysis is likely to fail. Figures 4-1 and 4-2 show the results of

Figure 4-1: Perspective View of DC38008 Using Building Heights From Shadow Estimation

generating 3D perspective views using the height information derived in SHAVE. These scenes are generated by image perspective transformation of the vertical images shown in Figures 1-2 and 1-1. The basic procedure is to use the original intensity image as a texture map for an underlying digital elevation model. In addition to a coarse terrain model, we use a ground truth file containing building locations. This ground truth file contains polygonal descriptions of the <latitude,longitude> location of the buildings. Building heights are derived from the output of SHAVE and are added to the scene database. The scene can then be rendered from an arbitrary 3-dimensional viewing position. In both cases, we have chosen a viewing position that views the original image from the bottom of the scene at an oblique angle with the scene center in the center of the original image.

5. Conclusions

This paper has described a set of techniques that perform shadow analysis on high resolution aerial imagery. Each exploits the relationship between man-made structures and cast shadows to perform structure shape prediction, structure grouping, structure verification, and height estimation. Detailed examples of the use of these programs were presented using aerial imagery of complex urban scenes. These modules provide important information about the 2-dimensional and 3-dimensional scene structure without having an explicit model of objects in the scene. Such techniques are important to indirectly

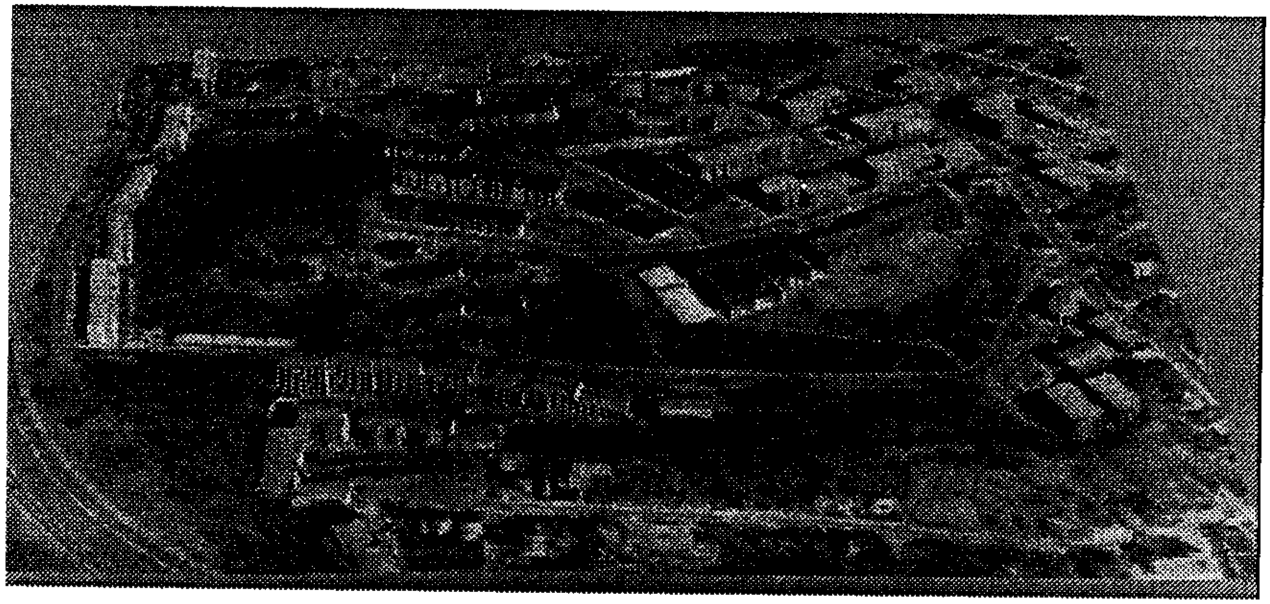

Figure 4-2: Perspective View of DC37405 Using Building Heights From Shadow Estimation

obtain height information when stereo image coverage is not available or when stereo matching is infeasible. These techniques can be used to generate new pieces of information in the scene interpretation or to rank order existing hypotheses.

There are several areas for future work. Scene registration for stereo matching requires the determination of conjugate points in left/right image stereo pairs in order to perform an initial local orientation. Since shadows fall on the same point on the ground for imagery taken at the same time of day, techniques for automated shadow corner detection may be usable to automate the selection of ground control points [13]. Along a similar line, the use of building/shadow edges may be useful to provide coarse matching for a feature-based stereo system. That is, in certain cases, it may be possible to directly match two sides of building using techniques similar to those described in SHADE. Finally, we are working on integrating these techniques into a high performance building extraction system that utilizes shadow analysis to aid in area-based region segmentation, edge/corner structure generation, correlation-based and feature-based stereo.

We believe that shadow analysis will play an important role in a variety of scene analysis techniques, just as it appears to be an important cue for human photo interpreters.

6. Acknowledgements

We would like to thank our colleagues in the Digital Mapping Laboratory for an interesting and congenial working environment. We would particularly like to acknowledge the contribution of Aviad Zlotnick and Yuan Hsieh for their work on the building hypothesis generator, and Lambert Wixson and Eric Thompson for their work on the 3D scene generation system.

7. Bibliography

[1] Avery, T. E.
Interpretation of Aerial Photographs.
Burgess Publishing Company, 1977.

[2] Aviad, Z.
Locating Corners in Noisy Curves by Delineating Imperfect Sequences.
Technical Report CMU-CS-88-199, Carnegie-Mellon University, December, 1988.

[3] Aviad, Z., McKeown, D. M., Hsieh, Y.
The Generation of Building Hypotheses From Monocular Views.
Technical Report, Carnegie-Mellon University, January, 1989.
to appear.

[4] Huertas, A. and Nevatia, R.
Detecting Buildings in Aerial Images.
Computer Vision, Graphics, and Image Processing 41:131-152, April, 1988.

[5] Lowe, D. G., and Binford, T. O.
The Interpretation of Geometric Structure From Image Boundaries.
In *Proceedings: DARPA Image Understanding Workshop*, pages 39-46. April, 1981.

[6] McKeown, D.M.
MAPS: The Organization of a Spatial Database System Using Imagery, Terrain, and Map Data.
In *Proceedings: DARPA Image Understanding Workshop*, pages 105-127. June, 1983.
Also available as Technical Report CMU-CS-83-136.

[7] McKeown, D.M., Denlinger, J.L.
Map-Guided Feature Extraction from Aerial Imagery.
In *Proceedings of Second IEEE Computer Society Workshop on Computer Vision: Representation and Control.* May, 1984.
Also available as Technical Report CMU-CS-84-117.

[8] McKeown, D.M., McVay, C.A., and Lucas, B. D.
Stereo Verification In Aerial Image Analysis.
Optical Engineering 25(3):333-346, March, 1986.
Also available as Technical Report CMU-CS-85-139.

[9] McKeown, D.M.
The Role of Artificial Intelligence in the Integration of Remotely Sensed Data with Geographic Information Systems.
IEEE Transactions on Geoscience and Remote Sensing GE-25(3):330-348, May, 1987.
Also available as Technical Report CMU-CS-86-174.

[10] McKeown, D. M.
Building knowledge-based systems for detecting man-made structures from remotely sensed imagery.
Philosophical Transactions of the Royal Society London A324:423-435, March, 1988.

[11] McKeown, D.M. and Denlinger, J. L.
Cooperative Methods for Road Tracking in Aerial Imagery.
In *Proceedings IEEE Computer Vision and Pattern Recognition Conference*, pages 662-672. June, 1988.

[12] Nagao, M., Matsuyama, T., and Ikeda, Y.
Region Extraction and Shape Analysis in Aerial Photographs.
Computer Graphics and Image Processing 10(3):195-223, July, 1979.

[13] Perlant, F. P., McKeown, D.M.
Scene Registration in Aerial Image Analysis.
In *SPIE Proceedings on Reconnaissance, Astronomy, Remote Sensing, and Photogrammetry.* January, 1989.

[14] *The Manual of Photogrammetry*
Fourth Edition edition, American Society of Photogrammetry, Falls Church, VA., 1980.

[15] Shafer, S. A.
Shadows and Silhouettes in Computer Vision.
Kluwer Academic Publishing, 1985.