Detecting Runways in Aerial Images[1]

A. Huertas, W. Cole and R. Nevatia
Institute for Robotics and Intelligent Systems
University of Southern California
Los Angeles, California 90089-0273

Abstract

We are pursuing the detection of runways in aerial images as part of a project to automatically map complex cultural areas such as a major commercial airport complex. This task is much more difficult that appears at first. We use a hypothesize and test paradigm. Hypotheses are formed by looking for instances of long rectangular shapes, possibly interrupted by other long rectangles. We use runway markings, mandated by standards for runway construction, to verify our hypotheses.

I. Introduction

Our aim is to develop general techniques for automated mapping and photointerpretation tasks. We have chosen major commercial airports as a test domain that has a variety of interesting characteristics.

Airports contain a variety of objects, such as the transportation network (runways, taxiways, and roads), a variety of building structures (hangars, terminals, storage warehouses), and a variety of mobile objects (automobiles, aircraft, humans). Further, the airport complexes are under continual change, usually due to expansion. The images themselves are rather complex due to the large number of objects present in them. The mapping of this domain, thus, offers a variety of challenging problems.

Our goal is to map all of the interesting objects in the scene and also to devise integrated descriptions that include the functional relationships of the objects in the scene. In this paper we concentrate on the mapping of runways (we are pursuing mapping of buildings in separate work [Huertas and Nevatia, 1987]). The runways and taxiways may appear to be modeled easily — namely as long, thin, rectangular strips of uniform brigthness. However, the real images are much more complex, as shown in figure 1, a portion (LOGAN:800 × 2200 resolution) of Logan International Airport in Boston, and in figure 2, a portion (JFK:1500 × 2600 resolution) of John F. Kennedy International Airport in New York. These help illustrate the following:

- *Object complexity:* Runways have a variety of markings. These are applied to the paved areas of runways and taxiways to identify clearly the functions of these areas, and to delimit the physical areas for safe operation and aid pilots. In many cases there are visible signs of heavy use, such as tire tread marks, oil spots, and exhaust fume smears. Also, runways have shoulders of various widths.
- *Object composition:* Runways may not be of uniform material. The landing surface and the shoulders may be of the same or different for different runways in the same airport. Runways may be extended using different materials. In certain geographical locations, the runway surfaces develop defects that need to be repaired; the repair work, usually in the form of patches is not necessarily homogeneous with the original surface material, and can have random shapes.
- *Object functionality:* Runway surfaces may be occluded by trucks and aircraft. Runways have access taxiways and service roads in a variety of positions with respect to the runway. Runways can intersect with other runways. Also, old runways or portions of them may be now used for other purposes.

One of the major causes of difficulties in detecting runways and other objects in real aerial scenes is that the low level segmentation rarely give complete and accurate results. In our work we have chosen to work primarily with the line *segments* computed from the intensity edges in the image. These lines may be fragmented, due in part to inadequacies in the line detection process, and in part due to actual structures in the image. In general, we assume that the images are of fairly good quality and of adequate resolution.

Our method uses the hypothesis formation and verification paradigm to detect runways. Our approach uses a generic model of the objects of interest derived from the following sources of knowledge:

- *Geometry* and *Shape*: We know that we are looking for instances of objects whose outlines represent a rectangular shape having a large aspect ratio of length to width. We know that runways have *ends* as opposed to nearby straight stretches of highways and roads.

[1]This research was supported by the Defense Advanced Research Projects Agency, monitored by the Air Force Wright Aeronautical Laboratories under contract F33615-84-K-1404

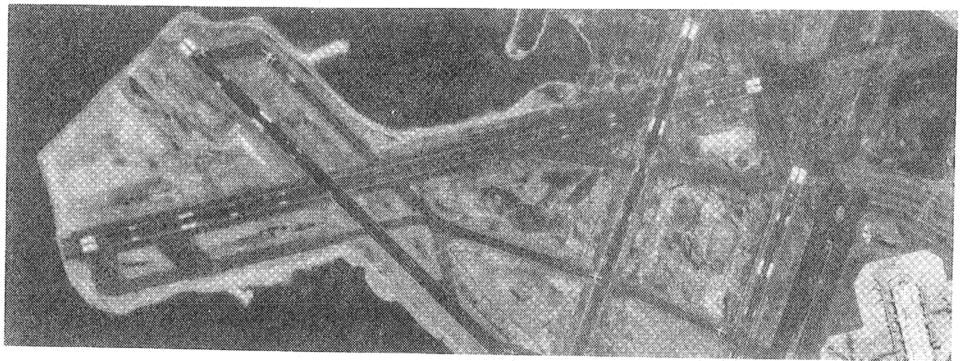

Figure 1: Logan International Airport image (LOGAN)

Figure 2: John F. Kennedy International Airport image

- *Specific Knowledge* of airport design: We know the features that make a visible long strip in the image an airport runway: the standard markings applied to the surfaces, according to FAA specifications. From airport engineering we also know the range of angles between runways, range of widths and so on.
- *Photometric Knowledge*: Intensity data may be of some help in verifying runway hypotheses when runway markings are non existent or not available due to lack of contrast or lack of resolution. Our current implementation does not make use of this knowledge but only uses the image resolution information.

In work reported here, our verification step consists only of finding the various markings we expect. We have not yet combined the different criteria to give an overall confidence value. This process should, ideally, take place in the context of the larger system that is also reasoning about other objects in the scene, such as the remainder of the transportation network, buildings and the mobile objects. Location of these objects will mutually affect the confidence levels of the descriptions of other objects. Thus, the system described here should be viewed as a module for the larger system to operate on.

The software architecture in our system consists of collections of functions that operate on linear features on the basis of constraints imposed by the object's geometry. Extensive work on rule based systems for aerial image analysis has been reported by McKeown at CMU (see for example [McKeown et al, 1987]). Their approach however is based on region features rather than linear features.

II. Description of the Method

A. Formation of Runway Hypothesis

1. Detection of Line Segments and Apars

We have chosen to work primarily with line *segments* extracted from the image. Geometric knowledge of the desired structures indicate that they should be characterized by parallel lines of opposite contrast. We call such pairs of lines "anti-parallel", and abbreviate them as *apars.* Apars form the basic unit of our further analysis.

We use the USC "LINEAR" line detection system [Nevatia and Babu, 1980] to obtain line segments and apars. Each linear segment is described by its length, orientation, contrast, and position of its end points. Additionally we also know if a segment connects to another segment at either end. Figure 3 shows the 8262 line segments computed from our LOGAN example. The center axis lines of 9,498 apars, shown in figure 4, were computed from the LOGAN segments by specifying the minimum (in our examples, 1 pixel) and maximum (60 pixels) distance between the anti-parallel pairs of segments. The range is derived from the known image resolution. The apars are described by their length, orientation, end points, width and color (brighter or darker than surround). We also know if apars are connected to other apars at either end.

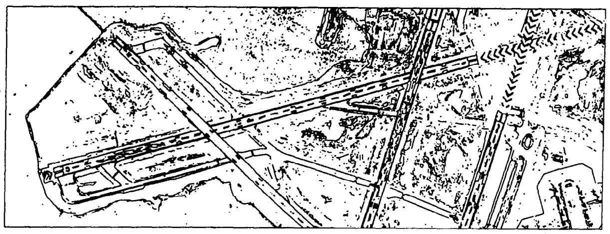

Figure 3: Line Segments from LOGAN image.

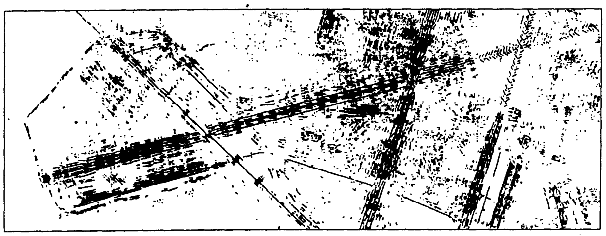

Figure 4: Anti-Parallels from Segments in LOGAN image

2. Reduction of Search Space

Each line segment may contribute to many apars, as is the case along runway features where there may be a large number of linear features parallel to the runway. This leads to a large search space that we reduce by implementing a focus of attention mechanism that facilitates the detection of "targets" in the presence of a large number of "distractors". We accomplish this by computing estimates of the directions and widths of potential runways. Using these estimates we extract from the set of apars, those in the selected directions and having a range of widths, and form sets of apars presumably representing fragments of runways.

First, we estimate the direction of the runways by computing a length-weighted histogram of the apar *orientations*. The histogram for the LOGAN apars is shown in figure 5. The three sharp peaks denote the dominant orientations of the linear features (including runways) in the image.

To estimate the runway widths we compute a length-weighted histogram of the apar *widths* including only those apars oriented in the estimated runway directions. This histogram (not shown) typically shows three width groups: a group of wide apars including runway and shoulder fragments, a middle group including taxiways and service roads, and in some cases, narrow shoulders, and a group of thin apars including the surface markings.

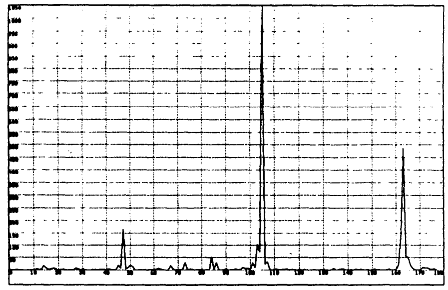

Figure 5: Length-Weighted Histogram of Apar Orientations

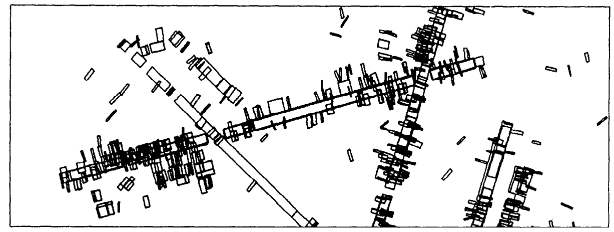

Figure 6: Apars representing initial set of Runway Fragments

We extract form the set of apars those in the selected directions and belonging to the width group for possible runways. We construct one set of runway fragments for each orientation peak, allowing for a tolerance of 5° on both sides of the peaks. The three sets for the LOGAN example are combined and shown in figure 6. We show the apars as rectangles to depict their width. A comparison of the original set of 9,498 apars to the 518 shown in figure 4 gives, in this example, a 94% reduction in the search space.

3. Joining Apars on the Basis of Continuity

Apars are usually broken due to noise in the image and inadequacies in the low-level processes. However, some of the breaks are due to real structures in the image. Consider for example where taxiways join runways. One one of the boundaries of the runway is continuous while the other boundary is broken at the junctions. The runway portions on both sides of the junction form collinear apars having the same width. We join these apars allowing a 5° tolerance in collinearity and 5 pixels tolerance in width. The resulting longer apar must have an orientation that is compatible with the estimated direction of the runway within a small tolerance (5°).

In some cases, as in our LOGAN example, there is sufficient resolution and contrast in the image for the edge detector to be able to resolve both boundaries of the white side stripes that bound the landing surfaces of some runways. In these cases the outside boundaries of the side stripes result in apars that contain apars resulting from the inside boundaries of the same side stripes. We remove *properly contained* apars from the sets. Apars that overlap however are preserved. We also remove apars having an aspect ratio smaller than 1, as they are considered unreliable. The result of these processes is shown in figure 7.

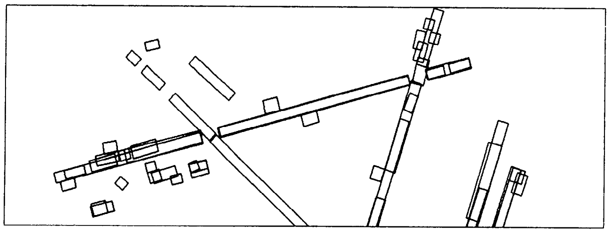

Figure 7: Apars joined on the basis of boundary continuity and filtered on containment and aspect ratio.

4. Joining Apars on Collinearity and Analysis of Gap Texture

Next we join collinear (within 5°) apars that have similar widths (within 5 pixels) on the basis of examining the gap between the fragments. Many runway apars may remain fragmented due to noise and occlusion. Consider for example where two runways cross or when there are aircraft on the runways.

In general, this process is quite liberal in the analysis of the information in the gaps, as long as the resulting apar has a direction consistent (within 5°) with the hypothesized runway direction. For instance, if the gap contains mostly segments that are oriented in the direction of the apars, we join them. If the gap contain mostly segments oriented at an angle consistent with the angles allowed between crossing runways then we join them. However, as in our JFK example, repair work, changes in surface material, signs of heavy use, oil spots and tire tread marks, can result in basically random arrangements of segments (texture) in the gaps. Thus, we also consider the lengths of the apar candidates and the size of the gap. A more precise way to support these decisions would include the use 3-D information to determine if the surface is smooth and flat. The result of this process for our LOGAN example is shown in figure 8.

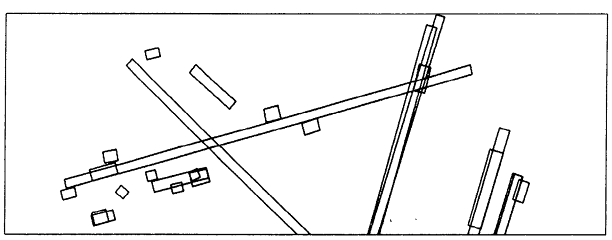

Figure 8: Apars joined on segment texture and gap analysis.

5. Final Runway Hypotheses

At the end of the joining process, short apars are removed from the sets if they have an aspect ratio smaller than 20:1. This will preserve those apars possibly representing partially visible runways. The resulting apars constitute the instances of the shapes found in the image that match our geometric model for airport runways. These are shown in figure 9.

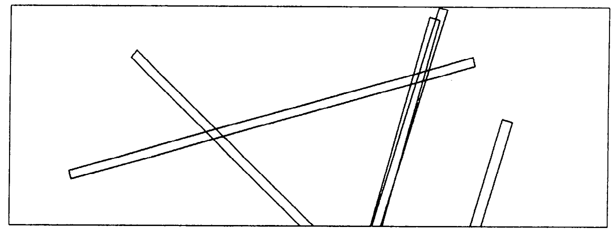

Figure 9: Runway Hypotheses.

B. Runway Verification

Hypotheses disambiguation and verification of runways is accomplished primarily by detection and identification of runway markings. We currently look for centerlines, side stripes, threshold marks, touchdown marks, distance marks, and blast pad marks. Most of these are shown in figure 10 (from [Ashford and Wright, 1984]). We have specific knowledge of their dimensions and position [Federal Aviation Administration, 1980].

We map these knowledge onto the image's coordinate system for the available image resolution. Fractions of pixel indicate lack of resolution and, instead of looking for, say two close markings, we look for one wider marking, equivalent to the fusion of the individual non-resolved markings.

The visibility of runway markings is primarily determined by the following factors:

- *Image Resolution*: Determines if the markings can be resolved.
- *Surface Material*: The contrast between markings and background depends on the underlying surface. White markings on a dark asphalt surface are quite visible. Concrete runways are brighter and perhaps make it more difficult to detect the markings. In some cases contrast depends also on the material in the runway shoulders.
- *Usage and Upkeep*: Tire tread marks, oil spots and exhaust fumes obscure the markings along and at the ends of runways. On the other hand, tire tread marks form quite visible and high contrasting dark regions in the center of concrete runways, and can be used for verification purposes. Our current technique relies on markings detected elsewhere to predict the presence of obscured markings.

The size and position of each runway hypothesis determines the window where we search for the markings. To find them we first look for thin bright apars in the window. If necessary we also look at the line segments. Figure 11 shows the markings found for our LOGAN example. The two overlapping (competing) hypotheses in figure 10 are disambiguated early because the incorrect hypothesis has only a few centerlines compared to those in the hypothesis that remains valid.

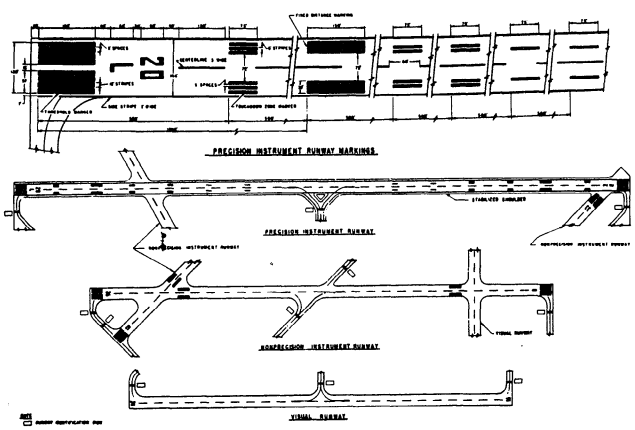

Figure 10: Standard Runway Markings

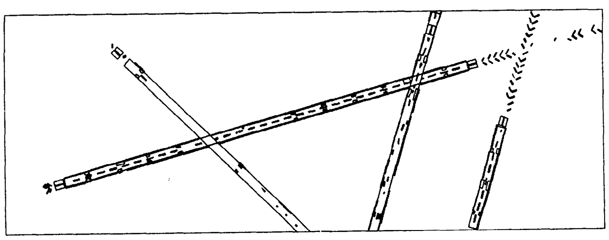
Figure 11: LOGAN Runways with Markings Detected.

1. Detection of Runway Centerlines

Centerlines are equally spaced along the landing surface of runways. To detect these we look in the middle of the hypothesized runway for bright apars having the desired dimensions. We do not enforce the separation constraint between centerlines to allow detection of broken or incomplete individual markings due to exhaust burns, tread marks, etc. We also look for individual segments (that do not form thin apars) down the middle of the runway.

2. Detection of Side Stripe Markings

The sides of the landing surface of runways are bound by side stripes. They result in thin bright apars. We look for these, at or near the boundaries of our runway hypotheses, and test that they are oriented parallel to the estimated runway direction. These thin apars are often broken mostly due to lack of contrast, and we do not attempt to join them. We however require that the fragments be collinear and that they have a consistent width.

3. Threshold Mark Detection

The threshold are probably the most important set of markings that can be used to verify a runway; they give pilots the position of the start and end of the runway. Often, these marks are partially worn away by exhaust fumes due to their position, so we expect our search to look for partial markings.

At the resolution in our examples the threshold marks appear as white rectangles separated by a dark zone. This results in two bright wide apars for each mark and a dark apar between them. In our search first look for the bright apars. These apars must be oriented in the direction of the runway (within a 5° tolerance). If we find only one of these apars, we hypothesize the position of the missing mark, and look in the line segment information for line segments to support our hypothesis. If no bright apar is found we look for the dark apar. It must meet the length and orientation constraints for the dark zone between the threshold marks. From its position and orientation we predict the position and orientation of the two threshold marks, and look for support evidence in the set of line segments.

4. Touchdown Mark Detection

Touchdown marks are located at a specific distance from the threshold marks, on each side of the runway. When present, at the resolution in our examples, they generate two bright apars and a bright apar between them. We look for these, and test them for consistent orientation.

5. Distance Marking Detection

Runways have a series of distance markings extending from the touchdown marks, equally spaced but of varying width. They generate specific bright and dark apars that we can look for. We look for the first (large) pair of distance marks first. For this we rely on the position of the threshold marks to predict their approximate position. Locating the small distance markings proceeds in a similar manner. We estimate their position from the large distance marks (if these are available, otherwise we use the position of the threshold marks) and do a search in the area for apars of the desired characteristics.

6. Blast Pad Mark Detection

Blast pad markings are optionally located at the ends of runways. They consist of pairs of white lines oriented at 45° angles with respect to the runways, and meet at the runway central axis. The separation between these pairs of lines varies thus, we detect them by looking for thin bright apars in the proper configuration.

III. More Results

The runways at LOGAN consist of dark asphalt, well maintained surfaces and markings, while JFK present a wide variety of problems. We therefore selected a portion of this airport as our second example. The level of complexity of most major commercial airports lies between our two examples.

In our JFK example, the partially visible apparent runways have no discernable markings on them. The complete runway running across the image shows increasing amounts of repair work, of a different material than that of the original surface. The darker material, however, makes some of the markings more visible. On the left side of the runway, the end of the runway becomes narrower as it turns into a taxiway. The accurate detection of the runway end thus depends on being able to locate the threshold markings. As shown below, we were able to locate them.

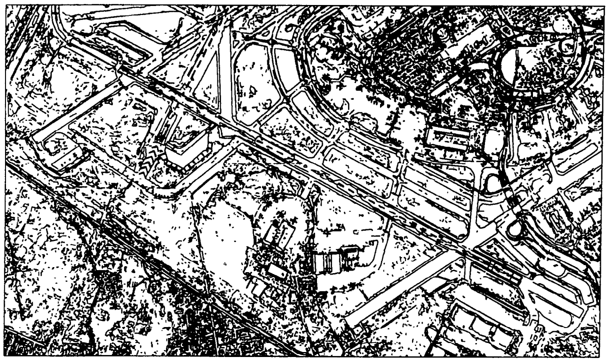

Figure 12: Line Segments from JFK Image.

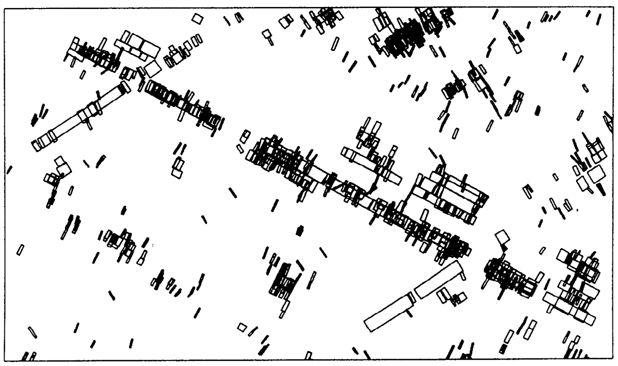

Figure 13: Initial Set of Runway fragments in JFK Image.

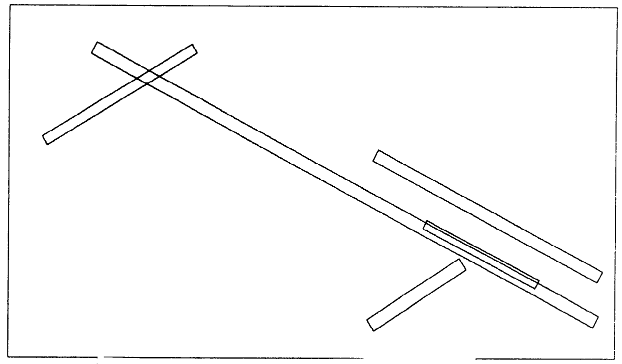

Figure 14: Runway Hypotheses.

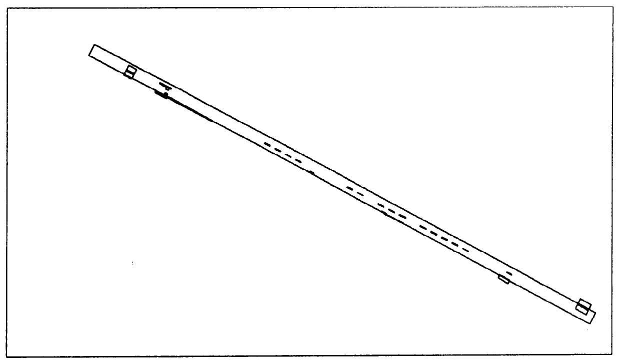

Figure 15: JFK Runway and Markings Detected.

The line segments computed from the JFK image are shown in figure 12. The reduced search space and apars representing the initial set of runway fragments is shown in figure 13. The runway hypotheses are shown in figure 14. Figure 15 shows the results of the verification process.

IV. Conclusion

We have described a technique, based on geometry and shape as the sources of knowledge suitable to form and test hypotheses representing instances of a known object shape, airport runways, using linear features.

We presented results on two very different airports to show the strength of the hypotheses formation process. Together with a sound search space reduction mechanism, and an object-specific feature verification technique, our method represents the state-of-the-art in runway detection. We have tested the technique on images of several major airports, varying in complexity between our two examples, with very encouraging results. In all our test the system parameters were the same.

Our basic technique can be easily extended to use the intensity image if necessary, feedback mechanisms, and analysis of non-standard markings. We point out that our hypotheses formation/verification technique can be useful for similar tasks, such as road detection and in general, transportation network detection.

We have not yet combined the different criteria to give confidence values. This process should, ideally, take place in the context of the larger system that is also reasoning about other objects in the scene, such as the remainder of the transportation network, buildings and the mobile objects. Location of these objects will mutually affect the confidence levels of the descriptions of other objects. Thus, the system described here should be viewed as a module for the larger system to operate on.

References

[Ashford and Wright, 1984] N. Ashford and P.H. Wright. *Airport Engineering, 2nd Ed.* Wiley and Sons, 1984.

[Federal Aviation Administration, 1980] FAA Advisory Circular. AC 150/5340-1E, November 4, 1980.

[Huertas and Nevatia, 1987] A. Huertas and R. Nevatia. Detecting Buildings in Aerial Images. To appear in *Computer Vision, Graphics and Image Processing.*

[McKeown et al, 1987] D.M. McKeown and W.A. Harvey. Automatic Knowledge Acquisition for Aerial Image Interpretation. In *Proceedings, Image Understanding Workshop*, Vol. 1, February 1987, pp. 205-226.

[Nevatia and Babu, 1980] R. Nevatia and R. Babu. Linear Feature Extraction and Description. In *Computer Vision, Graphics and Image Processing*, Vol. 13, 1980, pp. 257-269.