A System for Aircraft Recognition in Perspective Aerial Images

Subhodev Das[*], *Bir Bhanu*, *Xing Wu*, and *R. Neil Braithwaite*
College of Engineering
University of California
Riverside, CA 92521-0425

Abstract

Recognition of aircraft in complex, perspective aerial imagery has to be accomplished in presence of clutter, occlusion, shadow, and various forms of image degradation. This paper presents a system for aircraft recognition under real-world conditions that is based on the use of a hierarchical database of object models. The particular approach involves three key processes: (a) The *qualitative object recognition* process performs model-based symbolic feature extraction and generic object recognition; (b) The *refocused matching and evaluation* process refines the extracted features for more specific classification with input from (a); and (c) The *primitive feature extraction* process regulates the extracted features based on their saliency and interacts with (a) and (b). Experimental results showing the qualitative recognition of aircraft in perspective, aerial images are presented.

1 INTRODUCTION

Aircraft recognition is an important subproblem of photointerpretation (PI) which continues to be a major application domain of image understanding (IU) techniques for nearly two decades [2, 5]. Very few of the proposed aircraft recognition systems have actually addressed the concerns of real-world such as occlusion, shadow, cloud cover, haze, seasonal variations, clutter, and various other forms of image degradation. To illustrate the difficulty of the problem, we show typical aircraft images in Figure 1 where recognition has to be accomplished under low contrast and in presence of shadows. There exists no viable approach to date that can work satisfactorily on all these cases.

In this paper, we describe an end-to-end IU system for aircraft recognition under development which has been demonstrated to be effective in presence of shadows, clutter, and low image contrast. Our system uses a hierarchical representation of aircraft models consisting of generic aircraft, aircraft classes (e.g., jumbo aircraft), specific aircraft (e.g., Boeing 747), and aspects of a specific aircraft. Such representation is in terms of qualitative-to-quantitative descriptions that vary from advance concepts (e.g., aircraft wing) to primitive geometric entities (e.g., points, lines) and allow increasingly focused search of the precise models in the database. To account for image variabilities, our system exploits heterogeneous models such as those of camera/platform, sun, shadow to derive symbolic features in a robust manner. Finally, the system regulates the extracted primitive features based on their saliency, an ability that helps to distinguish relevant features from image clutter.

Section 2 describes the background and motivation behind the work reported in this paper. Section 3 describes the novel features of our aircraft recognition system. Section 4 presents the details of an algorithm that integrates feature refinement and object classification. Section 5 gives the details of implementation and the experimental results for qualitative recognition of aircraft using real-world data. Section 6 presents concluding remarks.

2 BACKGROUND

The different approaches to aircraft recognition that have been proposed so far can be broadly classified into the following categories: (a) *moment invariants*, which use moment invariant features of the aircraft silhouette and silhouette border, e.g., [4]; (b) *syntactic/semantic grammar*, which use linguistic pattern recognition techniques to analyze aircraft shapes represented by piecewise linear border approximations, e.g., [9]; (c) *Fourier descriptors*, in which the shape of the aircraft's closed contour is represented using a Fourier descriptor (FD), e.g., [10]; and (d) *model and knowledge-based*, in which an aircraft is represented in a hierarchical part-subpart fashion and the recognition process assembles image features in a forward- or backward-chaining fashion using the system's model and knowledge base, e.g., [3].

The "knowledge-free" techniques are inadequate for real scenes as they treat the object of interest in isolation from the rest of the image. Also, these approaches have very limited capability to handle clutter, shadow, occlusion, etc. Among the knowledge-based techniques, the ACRONYM system by Brooks [3] is closest to the system described in this paper. However, the primi-

[*]Current address: PEB, 47 Hulfish St, Princeton, NJ 08542.

0-8186-6410-X/94 $4.00 © 1994 IEEE

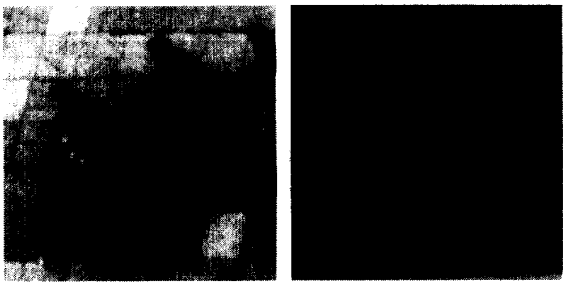

(a) (b)

Figure 1: Representative aircraft recognition scenarios: (a) shadow and (b) shadow and low contrast.

tive feature extraction process in ACRONYM is independent of the subsequent analysis. So errors in preprocessing of primitives persist throughout analysis. Also, ACRONYM does not possess any hypothesis verification step and the complexity of its constraint-based modeling is limited. All of these factors limit its application to images that are relatively simple photometrically, e.g., free of shadows and clutter.

Most model-based aircraft recognition approaches rely on some form of segmentation of the raw image primitives into meaningful groups. But good segmentation cannot be guaranteed for real-world, complex images. Instead, a more effective approach would be to decompose the overall recognition problem into one of successive identification of a more specific instance of an already recognized generic class. One of the key problems of poor recognition performance in real-world scenarios is the presence of shadows, as seen in Figure 1. However, shadows can be helpful to identify image features that are caused by raised structures such as aircraft. Occlusion, which in aerial imagery is primarily due to self-occlusion or close parking of aircraft, is another issue that a robust aircraft recognition system needs to address. To handle image clutter, the feature selection process must be able to distinguish between image features that are caused by the aircraft structure and those that constitute the clutter. Detection of perceptually salient contours in aerial images, which could very well be aircraft contours, should considerably simplify the segmentation and recognition tasks.

3 SYSTEM DESCRIPTION

The system for model-based aircraft recognition is schematically shown in Figure 2. It has four key features whose descriptions are given below.

The *hierarchical object model database* is useful for generic-to-specific object recognition. There are two important considerations in designing such a database: the choice of features to represent a particular object class and the matching process. In our system, the choice of

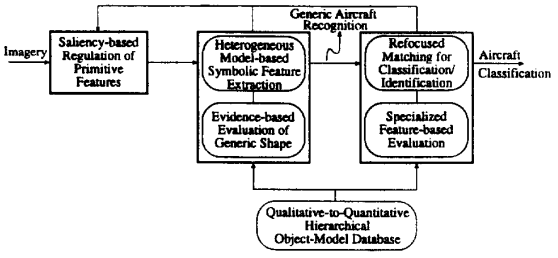

Figure 2: A schematic of the model-based aircraft recognition system.

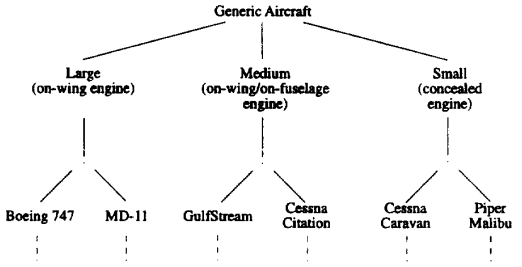

Figure 3: A partial hierarchy of generic-to-specific aircraft.

features is driven by their ability to discriminate among objects at the same level of the hierarchy. The selected features are ranked according to their relative importance in recognizing that particular object. The top-level of the hierarchy contains the shape (qualitative) attributes of a generic aircraft in terms of its structural subparts. The progressingly deeper levels embody more specific knowledge that becomes completely quantitative at the terminal nodes, i.e., location of geometric models. A partial hierarchy illustrating this particular database structure is shown in Figure 3. Using this database, reasoning about objects and its classes can be cascaded without requiring the presence of the same features (for matching) at all levels. The matching process can also search a lower level for distinguishing features should a categorization be not possible at a particular level because of the lack of suitable features. Thus, the flow of control during matching is bi-directional – between a generalized class and its more specialized subclasses.

The flow of low-level features (primitives) that are used to derive the symbolic features of aircraft models is *regulated* based on the saliency of these primitives. Perceptual saliency is an useful property which can distinguish an object from its background clutter in real scenes [8]. Our approach utilizes local measures of saliency based on the strengths of detected edge pixels, and lengths and local curvatures of edge segments. Regulation may also be based on the "specialized" nature of the features as required by the refocused matching process.

Symbolic features of particular aircraft models are *de-*

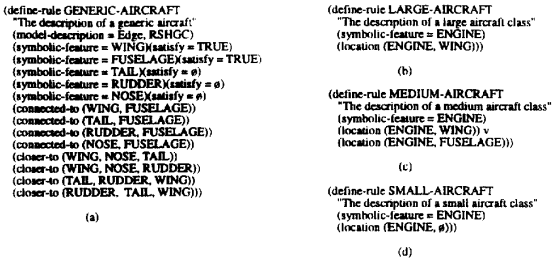

Figure 4: Simplified examples of advance concepts: (a) A generic aircraft, and the three aircraft classes (b) Large, (c) Medium (d) Small.

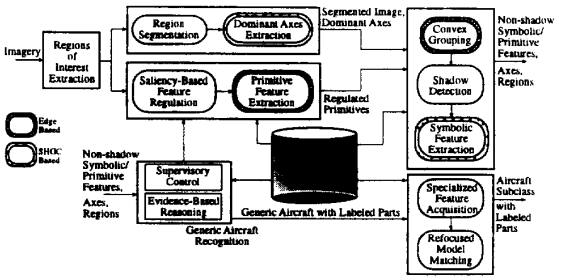

Figure 5: A flow diagram of the *qualitative* aircraft recognition algorithm.

rived by incorporating domain-specific knowledge into the geometrical and physical constraints of the mappings of the primitives. Domain knowledge is embedded in the production rules describing these symbolic features. The production rule definitions of a generic aircraft and its three subclasses are illustrated in Figure 4. To extract these features, the recognition search process is initiated at the top level of the database hierarchy and is looped through the production rules of the node being visited in a goal-decomposition fashion until a a rule is encountered whose conditions require mapping of primitives. The *evaluation* of the extracted features involves the verification of the global semantic shape components of the generic model. It utilizes heterogeneous models: edge/gray scale-based model of image segmentation, models of shadow casting process, and models of image acquisition. Also used are the dominant axes that characterize the shape of the generic aircraft class. During evaluation, feedback from the generic recognition module to the feature regulation module (see Figure 2) helps to acquire additional low-level features in the event of recognition failure or low recognition confidence.

The *refocused matching and evaluation* module is responsible for further classification of an object whose category has been determined by the generic recognition process. The key difference between this step and multiresolution approaches to recognition is that the former views the object model database at increasing resolution instead of the image. Data reduction is achieved by deriving symbolic features that are more "focused" or localized with respect to a particular level. Usually, the symbolic features of the identified generic model guide the search for image primitives required for deriving these focused symbolic features. In addition, the former type of symbolic features may be subjected to mensuration during refocused recognition.

4 ALGORITHM

The emphasis of this paper is on qualitative recognition of aircraft using the system described in the previous section. Initially, a solid modeling system is used to build B-spline CAD models of aircraft with *multiple representations*. In this work, we utilize two representations, polyhedral edge-based (approximations of B-spline CAD models) and straight homogeneous generalized cylinder (SHGC). Given an input 2D image and ancillary data about the imaging parameters and scene conditions, the algorithm for qualitative recognition of aircraft consists of the steps illustrated in Figure 5. First, the regions of interest (ROIs) likely to contain aircraft are identified in the high resolution input image by performing an intelligent search [7]. Each ROI is then analyzed following the steps described below.

4.1 Saliency-based Feature Regulation

Edge pixels are detected in an input ROI by applying multiple thresholds. This is motivated by the fact that *no* single threshold is suitable for all the different images that may be encountered in practice. Subsequently, these edge pixels are used to extract perceptually salient contours and other primitive features.

Our approach to identifying perceptually salient contours is based on finding long, smooth edge segments that are made up of high-magnitude edge pixels. This may be formulated as a problem of finding an edge segment of length N starting at a terminal pixel, corresponding to $s = 0$ (s being the segment parameter), and subject to the following optimization:

$$\max_{C \in C_0^N} \left[\int_C [w_1 \nabla(s) + w_2 \lambda(s)] ds - w_3 \int_C (d\theta/ds)^2 ds \right], \quad (1)$$

where $0 < w_1, w_2, w_3 < 1$. Here, C_0^N denotes the set of all contours, C, of length N beginning at $s = 0$. The variable $\nabla(s)$ is the magnitude of an edge pixel along the contour and denotes the *strength* component of the criterion function; $\lambda(s) = 1$, if $\nabla(s)$ is greater than a chosen threshold and $\lambda(s) = 0$, otherwise, and it represents the *length* component; $d\theta/ds$ denotes the local curvature at the selected pixel. where $\theta(s)$ is the slope along the contour, and it is a measure of local roughness. To reduce complexity. the above optimization model is decom-

posed into n sub-optimization steps, each of which finds an edge segment of length N_i such that $N = \sum_{i=1}^{n} N_i$. These sub-optimization steps differ on the threshold of edge magnitude (hence the use of n thresholds) which affects the $\lambda(s)$ term of (1).

At the first sub-optimization step, edge segment-following is initiated at a terminal pixel (one which has a single neighboring edge pixel) in the corresponding edge image. To continue edge-following, a neighbor of the last selected edge segment pixel is chosen that maximizes (1). To account for noise, our approach allows a gap length of up to two pixels in the edge segment. Once it is terminated within the current edge image, the process is continued (i.e., the next sub-optimization step) in the image obtained with the next lower threshold. The optimization process ends when the current edge image is the last of the edge-image set. Repetition of the optimization process for different starting pixels in the same initial edge image yields different contours having the same degree of saliency, while contours obtained with different initial edge images are said to have different saliency. The contours having the same saliency are handed over to the primitive feature extraction process simultaneously, starting with the top-level configuration, i.e., contours whose initial edge image is obtained with the highest threshold.

4.2 Primitive Feature and Dominant Axes Extraction

The extraction of primitive features from regulated edge contours utilizes edge-based representation of shapes. Since this representation is derived from polygonal approximations of B-spline CAD models, the primitive features comprise linear segments. In our system, we have implemented a line extraction algorithm similar to the one proposed by Lowe [6].

The dominant axes correspond to the longest axes of the SHGC representations of CAD models. The potential dominant axes of the generic aircraft shape are obtained by connecting the extremities of a labeled region within a segmented image. The segmentation process is based on a joint relaxation of edge- and region-based approaches maximizing edge and region-border coincidence [1]. To determine the extreme points, the smallest convex polygon surrounding the object region is found. It can be shown that the vertices of this polygon lie close to the local extrema of curvature points along the boundary of the labeled region. For multiple local extrema, nearby points are grouped into clusters and the cluster centers are chosen to represent the region extremities.

4.3 Non-shadow Feature Extraction

Assuming that any arbitrarily shaped shadow boundary can be locally represented by straight lines, our algorithm to detect potential shadow lines is based on the test of bimodality of the local histogram. Initially, region segmentation [1] is carried out within a window on either side of an extracted line and the largest region is retained. The most significant modes of the histograms of these two regions are then subjected to the bimodality test. If the separation between the modes is less than a threshold or the smaller of the two is greater than another threshold, the line is ignored. Otherwise, it is marked as a potential shadow line.

In order to separate the shadow lines from the shadow-making ones, the algorithm first obtains convex groups of lines based on proximity and collinearity. Each line from a selected group is subjected to a convexity test by pairing it with another line from the same group. Lines failing the test are removed and put in a new group. After all the initial groups have been considered, this process creates the *first* set of convex groups and isolated lines removed during the convexity test. The *second* pass considers whether an isolated line can be put in a convex group based on proximity, collinearity, and convexity.

For each marked shadow line in a convex group, a corresponding shadow-making line from another group is sought by searching in a direction towards (or away from) the projection of the illumination point, i.e., the sun. The latter is determined from the ancillary data about the camera-platform position/orientation and the sun position together with the imaging parameters. The matching score of a pair of shadow-shadow making lines is computed from the degree of overlap of the two lines in the predicted direction. All the candidate matches of a selected shadow line are arranged according to the matching scores and marked with the corresponding group identifier. This entire matching process is repeated for all other shadow lines in that particular group. The most promising matching group is determined from the group identifiers of the candidate matches. Each line in the selected group is assigned a unique match from the candidate group based on the matching scores and enforcing similarity of spatial ordering of the selected lines and their matches. If most of the lines in the selected group have been assigned unique matches, then the group as a whole is marked as a shadow group and the matching group is marked as a shadow-making (i.e., non-shadow) group.

4.4 Symbolic Feature Extraction

The symbolic features of the generic aircraft class include trapezoid-like shapes for wings, tails, and rudder, and wedge-like shape for the nose part. Convex groups of non-shadow lines are used to derive these symbolic features. To identify the trapezoid-like shapes, groups of three (partially closed contour) and four (fully closed contour) lines are considered. Pairwise intersections of lines are verified to occur near detected corners. Non-overlapping line pairs that are far apart are prevented to

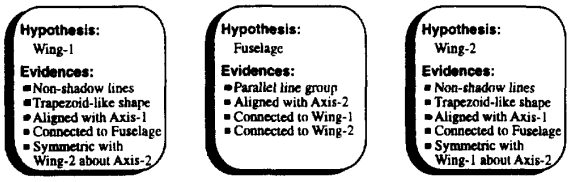

Figure 6: Interacting hypotheses for the key structural subparts of a generic aircraft. The evidences (positive) are listed in the order they are sought while verifying the corresponding hypotheses.

have a high collinearity value by enforcing the condition that the average separation between the lines of a pair is proportional to the smaller line length. To overcome the problem due to oversegmentation, i.e., fragmentation of long lines into smaller parts, candidate groups of lines are merged based on collinearity measures. If the total length of the lines in a group is smaller than a certain fraction, T_{peri}, of the perimeter of the trapezoid-like shape obtained by connecting these lines, then that group is discarded.

It is important to note that a range of threshold values is associated with each perceptual measure of parallelism, proximity, and collinearity. Initially, the threshold values – T_p, T_c, and T_ρ, respectively – are set to the maxima of the corresponding ranges. However, these values can be relaxed during symbolic feature extraction based on the flow of evidence when multiple mutually supporting hypotheses interact (described next).

4.5 Evidence-Based Reasoning

Once the symbolic features have been derived, these need to be matched to the generic aircraft model through an evidence accumulation process which determines the number of positive evidences in support of a hypothesis. (In this paper, we do not consider negative evidences for a hypothesized object.)

During the matching process, the supervisory control module (see Figure 5) checks each hypothesis to determine its combined support based on the evidences associated with it. The combined support may be low if the evidence body is incomplete. Some pieces of evidence, known as the *critical evidences*, contribute higher support than others and must be present to accept a hypothesis. Examples of such evidences (see Figure 6) are trapezoid-like shape, alignment with a dominant axis, and interconnection with the fuselage for the generic aircraft-wing hypothesis. During run time, the supervisory controller determines which critical evidences are missing. Typically, this situation is caused by insufficient data, either due to the screening of features at the regulator level or due to the constraints on the imaging process for that particular viewpoint. In the former situation, the controller interacts with the low-level feature regulator so that less salient features may now be available along with the existing ones. Constraints due to the viewpoint warrants the access of the qualitative database by the supervisory controller to obtain further knowledge about the alternate evidences for this hypothesis.

4.6 Refocused Matching

Further refinement of the detected aircraft shape is required to improve upon the extracted symbolic information. This usually involves completing the generic aircraft description by accounting for the missing elements of the symbolic features. This is followed by obtaining a skeleton of the refined shape which is composed of the axes of symmetry of the structural subparts. The skeleton can be directly used for mensuration purposes when performing quantitative matching. The final output are the identified symbolic parts of the generic aircraft.

The labeled symbolic parts are next used to direct image-based search for more localized model features. Availability of these features at progressingly lower levels of the database hierarchy allows more precise classification of the recognized generic aircraft. The refocused matching process may utilize the symbolic/primitive features that have not been utilized in the generic aircraft recognition step or may request new or less-salient primitive features. Currently, our algorithm handles only qualitative model features.

5 EXPERIMENTAL RESULTS

The aircraft recognition system described in this paper is implemented in C on Sun Sparcstation. A UNIX shell-level program controls the entire system. Ancillary data about camera-platform position/orientation, weather condition (sunny/cloudy/hazy), sun angle, and camera parameters are provided in an external file. The hierarchical model database has three levels: generic, intermediate, specific. The intermediate level consists of three categories based on engine location - large (on-wing engine), medium (on-wing/on-fuselage engine), small (concealed engine). The experimental results of qualitative aircraft recognition are presented using aerial photographs of an air-base. The examples are ordered according to increasing level of recognition complexity.

Example 1: An aerial image is shown in Figure 7(a). In this example, we present the results of analyzing the ROI of Figure 7(b). The top-level (most salient) configuration of perceptually salient contours is shown in Figure 7(c). The linear segments of this structure are displayed in Figure 8(a). The ROI is segmented into three sets – shadow (Figure 8(b)), object, and background (Figure 8(c)) – using recursive application of the segmentation algorithm. The dominant axes of the largest object region are shown in Figure 8(d). The potential

shadow lines based on the bimodality test of neighborhood histograms are displayed in Figure 8(e). To identify the non-shadow lines, the illumination point projection is first determined using the ancillary data: the coordinates of the camera-platform position are (527, 337, 560) m and that of the point of the intersection of the line-of-sight (LOS) with the ground are (-500, -290, -560) m; the roll of the camera about the LOS is 20°. The sun position is noted to be *in front* of the camera. The shadow lines that paired up with "non-shadow" lines are displayed in Figure 8(f). Next, convex groups of the non-shadow lines are used to extract trapezoid-like symbolic features of Figure 9(a).

Figure 9(b) shows the results of initial recognition during which only one trapezoid-like feature is available to support a wing concept (for the right wing) due to its alignment with the wing axis (axis-1). Additional evidence for this hypothesized wing is obtained from a successful search for elongated blobs representing engine. Next, a search region for the second wing is set up along axis-1 on the other side (away from the hypothesized wing) of the fuselage axis (axis-2). A line pair is located within this region that aligns with axis-1 and for which the presence of engine can also be verified. Thus, the symbolic feature and the line pairs are retained as candidates for the two wings as shown in Figure 9(b).

The supervisory control now interacts with the feature regulator to identify candidate lines for the missing parts of the fuselage in the image region specified by the two hypothesized wings and the hypothesized fuselage subpart. The "shadow" lines removed earlier are considered first, since these are perceptually more salient than any other "non-shadow" lines that may be obtained from the next set of less salient contours. Two such lines are found that made up for the missing bottom part of the fuselage. The final results are shown in Figure 9(c) when an aircraft is said to be have been detected. The symbolic parts are further refined as shown in Figure 9(d). Finally, the class of the generic aircraft is determined to be *large* based on engine locations as seen in Figure 9(e),

Example 2: The second aerial image is shown in Figure 10(a) of which Figure 10(b) constitutes an ROI. The segmentation result and the extracted dominant axes are shown in Figure 10(c). The illumination point projection is determined using the ancillary data: the coordinates of the camera-platform position are (-984, -115, 700) m, that of the point of the intersection of the LOS with the ground are (1150, 150, -700) m, and the roll of the camera about the LOS is $-21°$. The sun position is recorded as *behind* the camera.

Figure 11(a) shows the top-level structure. The line-fits to this structure and the next incremental salient structure are displayed in Figure 11(b)-(c). Here, the salient structures constitute only 26% of the lines obtained from the edge image corresponding to the lowest edge threshold. This example illustrates the difficulty posed by shadows as nearly equal number of lines belong to the actual aircraft contour (38% of the salient structure lines) and the shadows (35% of the salient structure lines). Convex groups are formed using the lines of Figure 11(b). The potential shadow lines identified among the lines of Figure 11(b) are shown in Figure 11(d). To resolve the non-shadow lines, the rightmost of the two nearly parallel axes of Figure 10(c) is determined to be due to shadow based on the illumination projection point location. The final shadow lines are displayed in Figure 11(e). The groups of non-shadow lines are next used to obtain the trapezoid-like symbolic features of Figure 11(f).

One of the symbolic features of Figure 11(f) that aligned with one of the dominant axes, axis-1, is hypothesized as the wing, wing-1 (say). This is shown in Figure 12(a). Further evidence of wing is obtained from a successful search for engine feature in the vicinity of wing-1. To obtain evidence of the other wing, wing-2 (say), a search region is set up as shown in Figure 12(a) based on the image location of wing-1 and the condition of symmetry of the wings about the fuselage axis. However, the non-shadow lines contained within this region fail to identify any symbolic feature that may support a wing hypothesis. Additional non-shadow lines derived from less salient contours of Figure 12(b) are combined with the existing non-shadow lines within the search region (Figure 12(c)). The initial (maximum) values of the various perceptual constraints fail to produce any meaningful perceptual grouping as the lines are few in number and are separated apart. However, since the hypotheses of wing-1 and wing-2 are mutually reinforcing, the constraints are relaxed in steps. Particularly, the lowering of the threshold of the proximity, T_p, of a line-pair intersection to a detected corner and T_{peri} (initial values of $T_p = 30.0$ and $T_{peri} = 0.9$ in steps of $\Delta T_p = 7.0$ and $\Delta T_{peri} = 0.1$) cause grouping of lines to occur. As a result, a trapezoid-like symbolic feature emerges (the left wing tip) that drives the subsequent steps of recognition shown in Figures 12(d)-(e). The processes of shadow identification and symbolic feature extraction have reduced the number of useful lines for aircraft recognition to only 6% of the total number of lines that otherwise would have to be considered. Finally, the refocused matching process determines the class of this aircraft as *large* as indicated in Figure 12(f).

6 CONCLUSIONS

In this paper, we have described a system for recognition of aircraft in complex, perspective aerial imagery and presented an algorithm to recognize and classify aircraft using qualitative features. Our approach is motivated by the difficulties posed by real-world scenarios, such as occlusion, shadow, cloud cover, haze, seasonal

variations, clutter and various forms of image degradation. The *main contributions* of this research are the extraction of perceptually salient primitive features and their use in a regulated fashion, use of heterogeneous geometric and physical models associated with image formation for feature extraction and subsequent recognition, and integration of high-level recognition processes with low-level feature extraction ones. Real-world data, highlighting the difficulties of aircraft recognition in practical situations, is used to demonstrate the effectiveness of our proposed approach. Elsewhere, we have utilized the qualitative results to search for more specific aircraft models. Our system needs improvement in the integration and the simultaneous hypothesis verification steps to be carried out to handle effectively the situations like the one illustrated in Figure 1(b).

References

[1] B. Bhanu and R. D. Holben. Model-based segmentation of FLIR images. *IEEE Trans. Aerospace Elec. Sys.*, 26(1):2–11, 1990.

[2] J. F. Bogdanowicz and A. Newman. Overview of the SCORPIUS program. In *Proc. DARPA Image Understanding Workshop*, pages 298–308, Palo Alto, CA, May 1989.

[3] R. A. Brooks. Symbolic reasoning among 3-dimensional models and 2-dimensional images. *Artificial Intell.*, 17:285–349, 1981.

[4] S. A. Dudani, K. J. Breeding, and R. B. McGhee. Aircraft identification by moment invariants. *IEEE Trans. on Computers*, C-26(1):39–46, 1977.

[5] J. Edwards, S. Gee, A. Newman, R. Onishi, A. Parks, M. Sleeth, and F. Vilnrotter. RADIUS: Research and development for image understanding systems phase 1. In *Proc. DARPA Image Understanding Workshop*, pages 177–184, San Diego, CA, Jan. 1992.

[6] D. G. Lowe. *Perceptual Organization and Visual Recognition*. Boston, MA: Kluwer, 1985.

[7] H. Nasr, B. Bhanu, and S. Lee. Refocused recognition of aerial photographs at multiple resolution. In *Proc. SPIE*, Orlando, FL, March 1989.

[8] A. Sha'ashua and S. Ullman. Structural saliency: The detection of globally salient structures using a locally connected network. In *Proc. IEEE Second Intl. Conf. Comp. Vision*, pages 321–327, Tarpon Springs, FL, Dec. 1988.

[9] G. Y. Tang and T. S. Huang. Using the creation machine to locate airplanes on aerial photos. *Pattern Recognition*, 12:431–442, 1980.

[10] T. P. Wallace, O. P. Mitchell, and K. Fukunaga. Three-dimensional shape analysis using local shape descriptors. *IEEE Trans. Patt. Anal. and Mach. Intell.*, PAMI-3:310–323, 1981.

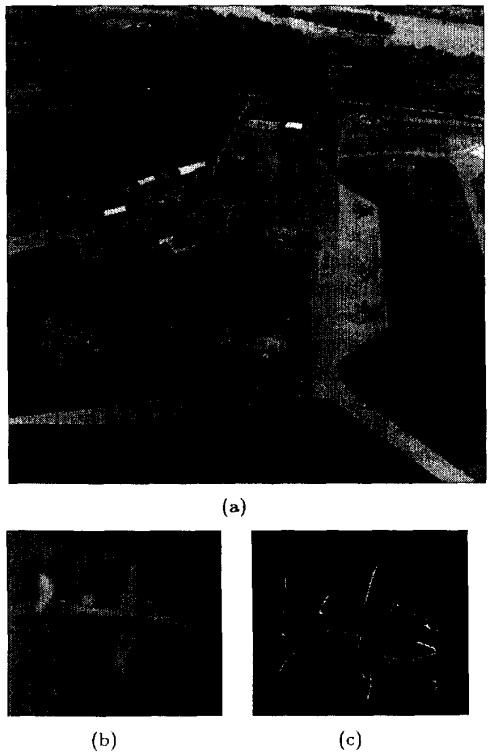

(a)

(b) (c)

Figure 7: Aerial view of an airfield: (a) Original image ($4K \times 4K$), (b) A ROI image (120×140) of the aircraft marked with a $\times$ in (a), (c) Most salient contour configuration.

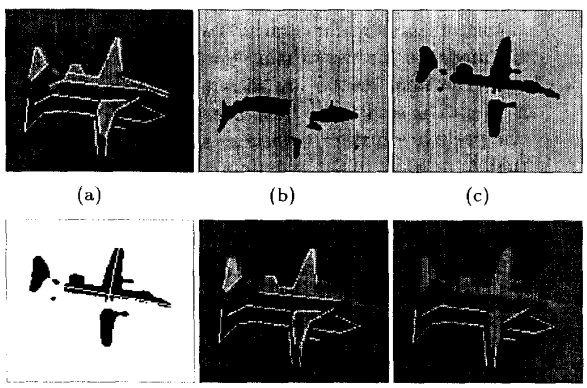

(a) (b) (c)

(d) (e) (f)

Figure 8: Results of feature extraction: (a) Straight lines fitted to the most salient contours, (b) Segmented regions where 0=shadow and 1=background, (c) Regions where 0=object and 1=background, (d) Dominant axes for the largest object region where axis-1=wing axis, axis-2=fuselage axis, (e) Potential shadow lines, (f) Resolved shadow lines. 0=black, 1=white.

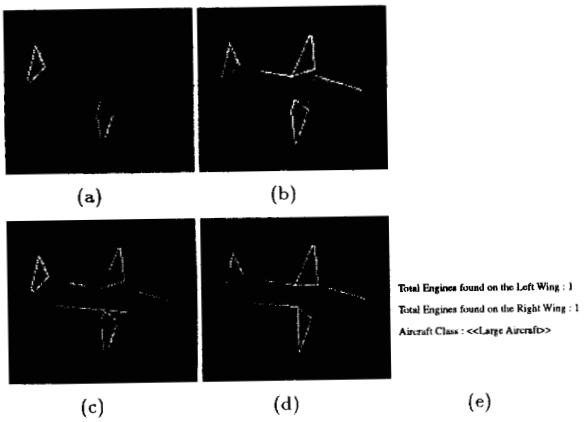

Figure 9: Results of qualitative object recognition: (a) Trapezoid-like shapes identified using non-shadow groups, (b) Symbolic features recognized with low confidence, (c) Improved recognition using regulated features, (d) Refined structural parts, (e) Class recognition.

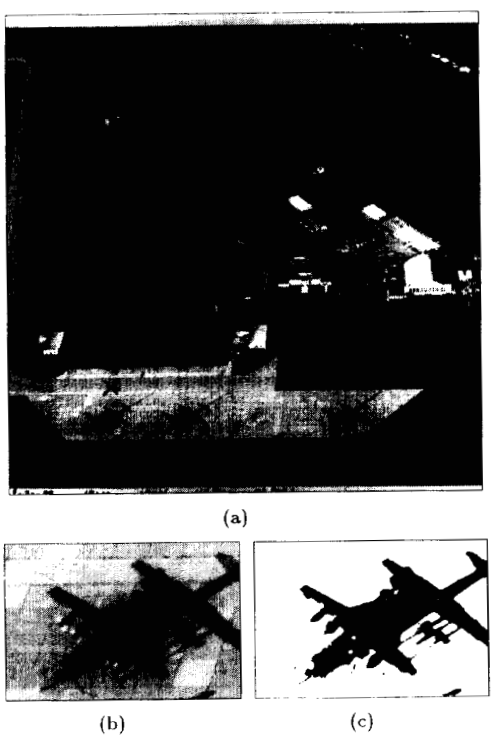

Figure 10: A second aerial image: (a) Original image ($4K \times 4K$), (b) A ROI image (300×450) of the aircraft marked with a $\times$ in (a), (c) Extracted dominant axes.

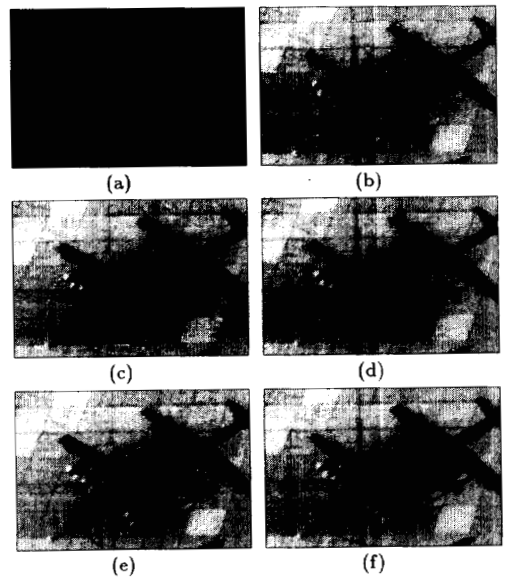

Figure 11: Results of feature extraction: (a) Detection of most salient structure, (b) Straight lines fit to the structure of (a), (c) Lines fit to the next incremental salient structure, (d) Potential shadow lines, (e) Resolved shadow lines, (f) Trapezoid-like shapes identified using non-shadow groups.

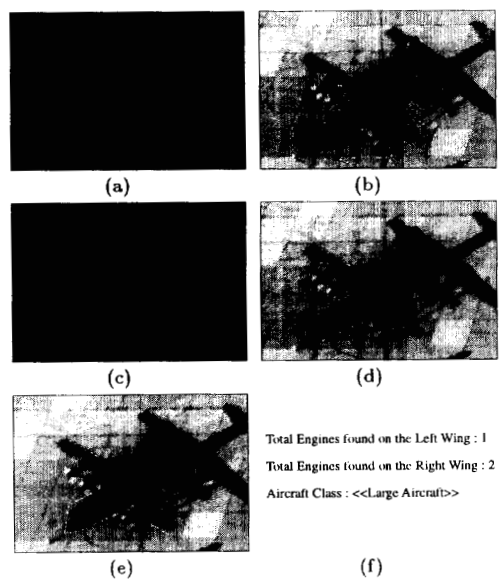

Figure 12: Results of qualitative object recognition: (a) A hypothesized wing and a search region for the second wing, (b) Current non-shadow lines together with those of Figure 11(c), (c) Additional non-shadow lines within the search region of (a), (d) Recognition after verification of the second wing, (e) Refined structural parts, (f) Class recognition.