A Vision-Based Automatic Safe Landing-Site Detection System

YU-FEI SHEN
Old Dominion University

ZIA-UR RAHMAN, Senior Member, IEEE
NASA Langley Research Center

DEAN KRUSIENSKI, Member, IEEE
JIANG LI, Member, IEEE
Old Dominion University

An automatic safe landing-site detection system is proposed for aircraft emergency landing based on visible information acquired by aircraft-mounted cameras. Emergency landing is an unplanned event in response to emergency situations. If, as is usually the case, there is no airstrip or airfield that can be reached by the unpowered aircraft, a crash landing or ditching has to be carried out. Identifying a safe landing-site is critical to the survival of passengers and crew. Conventionally, the pilot chooses the landing-site visually by looking at the terrain through the cockpit. The success of this vital decision greatly depends on external environmental factors that can impair human vision and on the pilot's flight experience, which can vary significantly among pilots. Therefore, we propose a robust, reliable, and efficient detection system that is expected to alleviate the negative impact of these factors. We focus on the detection mechanism of the proposed system and assume that image enhancement for increased visibility and image stitching for a larger field-of-view (FOV) have already been performed on the terrain images acquired by aircraft-mounted cameras. Specifically, we first propose a hierarchical elastic horizon detection algorithm to identify the ground in the image. Then, the terrain image is divided into nonoverlapping blocks, which are clustered according to a "roughness" measure. The adjacent smooth blocks are merged to form potential landing-sites, whose dimensions are measured with principal component analysis and geometric transformations. If the dimensions of a candidate region exceed the minimum requirement for safe landing, the potential landing-site is considered a safe candidate and is highlighted on the human machine interface. At the end the pilot makes the final decision by confirming one of the candidates, and also by considering other factors such as wind speed and wind direction, etc. Preliminary experimental results show the feasibility of the proposed system.

Manuscript received June 14, 2011; revised October 31, 2011 and January 31, 2012; released for publication March 27, 2012.

IEEE Log No. T-AES/49/1/944354.

Refereeing of this contribution was handled by M. Efe.

A significant portion of this work was supported under NASA Cooperative Agreement NNL07AA02A with ODU. Z. Rahman had several grants from NASA for this research effort. The research of D. Krusienski is supported by the National Science Foundation (NSF) and the National Institutes of Health (NIH).

Authors' addresses: Y-F. Shen, D. Krusienski, and J. Li, Department of Electrical and Computer Engineering, Old Dominion University, Norfolk, VA 23529, E-mail: (yshen002@odu.edu); Z. Rahman (passed away in Dec. 2010), Electromagnetics and Sensors Branch, NASA Langley Research Center, Hampton, VA 23681.

I. INTRODUCTION

The top-five leading factors of unplanned landing, which is also called emergency landing, are engine failure, running out of fuel, extremely bad weather, medical emergency, and aircraft hijack. Under the two most emergent situations, engine failure and running out of fuel, the aircraft may quickly lose flying power, and its maneuverability may be restricted to gliding. Once these happen a forced landing process has to be immediately carried out. If, as is usually the case, there is no airport, or even a runway, that can be reached by the unpowered aircraft, a crash landing or ditching is inevitable.

Finding a safe landing-site is vital to the survival of the passengers and the pilot. Conventionally the emergency landing-site is visually selected by the pilot by looking at the terrain through the cockpit. This is a required, fundamental skill acquired in the flight training program. However, many external environmental factors, i.e., fog, rain, illumination, etc., can significantly affect human vision so that the decision of choosing the optimal landing-site greatly depends on the pilot's flight experience—the most significant internal factor—which can vary a lot among different pilots. In addition the visual angle that the human eyes can simultaneously cover is limited: when the pilot looks to the left, what is on the right is missed and vice versa. Since time is of supreme importance in the scenario we are considering, the inability to simultaneously scan on both sides of the cockpit is a distinct disadvantage. Imaging sensors can alleviate this problem by creating panorama images that encompass the entire field-of-view (FOV) in front of the aircraft. In order to compensate for the natural inadequacies of human vision and also to alleviate the negative effects of both external and internal factors, a robust, reliable, and efficient process for safe landing-site detection is greatly desirable. Therefore, we present a vision-based, automatic safe landing-site detection system [1, 2].

Before introducing the design of the system, we first investigate appropriate criteria to assess the safeness of the landing-sites. Two geographic concepts, elevation and landform, are taken into consideration. The gradient of elevation generally determines the roughness of the terrain. Landform describes terrain covering, i.e., forest, grass, water, rock, buildings, etc. Smooth elevation gradient by itself is not sufficient to guarantee a safe landing-site since the associated landform could be hazardous to the landing procedure. In addition the landing-site must have sufficient length and width—which can vary with the type of airplane—to enable a safe emergency landing. In summary we evaluate the "safeness" of a potential landing-site by considering its surface roughness and its dimensions.

A landing-site is considered safe only if its surface is smooth and if its length and width are adequate. The proposed safe landing-site detection system is designed to automatically detect landing-sites that meet both of the requirements.

Many of the achievements of autonomous landing have been accomplished [3–10] by utilizing vision-based approaches to guide unmanned aerial vehicles (UAVs) or helicopters to known landing-sites. Landing marks, which often appear in high-contrast in the image so that can be easily detected, play an important role in these approaches by providing relative position information for state estimation and control. Nevertheless, for a landing strategy to be feasible in unknown environments, which is usually the case for emergency landings, the dependence on known landing marks is limiting, and, therefore, a flexible means of finding safe landing-sites is desired.

To date there are relatively few publications on automatic aircraft safe landing-site detection. As stated in [11]–[13], no automated forced landing research or automated forced landing system was available at their time of writing. In [11] Garcia-Pardo, et al. designed a two-step autonomous safe landing-site detection strategy. First, they applied a local contrast descriptor μ/σ, which is derived by normalizing the neighborhood of the to-be-tested pixel and then by calculating the mean μ and the standard deviation σ of its neighborhood, to assess the roughness of the ground under the assumption that the boundaries of hazards appear as high-contrast edges in the image, reflected by small values of μ/σ. A contrast threshold needs to be selected to differentiate smooth areas and boundaries, and the optimal contrast threshold is found to have a linear relationship with the ratio of mean and standard deviation of the whole image. Second, round landing-sites with a sufficient size are found in the smooth areas. The system was tested in an off-line fashion on 10 image sequences, which are captured by real flights over a synthesized environment, i.e., placing white boxes (obstacles) on grassy ground. The detection results are evaluated by a "failure rate" defined as the percentage of images in which the system fails to find any safe landing-site.

Fitzgerald, et al. also applied a two-step safe landing-site detection strategy [12]. First, Canny edge detector [14] is employed to describe the edges in the image. This detection is computationally more efficient than the local contrast descriptor mentioned above. Second, the safe landing-sites are found by scanning the smooth area with a set of rectangular masks that are predefined in different scales and rotation angles. Three problems associated with the second step are as follows. 1) It is inconvenient, or even impossible, to predefine a sufficiently large number of masks with all possible scales and angles. For example, if a potential safe landing-site has a shape which is not covered in the predefined mask set, the system is very likely to miss it. 2) Various aircrafts have different requirements for safe landing-sites in terms of the minimum length and width. Using a predefined set of masks limits the application of the system to different aircrafts. 3) It is computationally expensive to move all the masks over the smooth area. The computational cost is proportional to the number of masks so that the requirement of time conflicts with the requirement of detection accuracy.

Related research of spacecraft landing has been conducted by many groups in recent years. The NASA Jet Propulsion Laboratory (JPL) proposed a LIDAR-based hazard avoidance approach for safe landing on Mars [15]. They made use of elevation maps generated by scanning synthetic terrains with a simulated LIDAR model. Later, JPL introduced a fuzzy rule-based safety index to assess landing-sites [16, 17]. Furthermore, they brought multi-sensor images into their approach [18]. Based on a ballistic analysis, the JPL also proposed a method to estimate the reachable area for the spacecraft [19]. In addition to its application to landing on Mars, autonomous landing and hazard avoidance technologies (ALHAT) are also utilized for lunar landing [20–22] and UAV landing [12]. Therefore, the proposed system has a wide range of potential applications.

The contributions of the present paper consist of the following. 1) A delicate automatic safe landing-site detection mechanism is developed by seamlessly combining some existing image-processing and analysis techniques, including the block-wise roughness assessment, the classification of blocks based on their edge strength, the segmentation of candidate safe landing-sites, the dimension assessment of candidate landing-sites, and the visualization of detected safe landing-sites on the human-machine interface. 2) We propose a hierarchical elastic horizon detection algorithm to identify the ground in the aerial image so that the camera is relieved from the limitation of looking straight down to the ground. In the looking-forward mode, the system can detect safe landing-sites in front of the aircraft, thus providing more time to the pilot to prepare for landing, which is especially helpful for unpowered aircraft in emergency situations. 3) We improve the efficiency of the detection system by applying the Canny edge detector, instead of the local contrast descriptor [11], as a part of the roughness assessment algorithm and by utilizing the principle component analysis as the means to measure the dimension of smooth areas, instead of using a predefined set of masks [12]. 4) We develop a performance metric to comprehensively evaluate the detection results.

The remainder of the paper is organized as follows. In Section II we describe modules of the proposed system in details. Experimental results are shown in Section III. We discuss the necessity and the

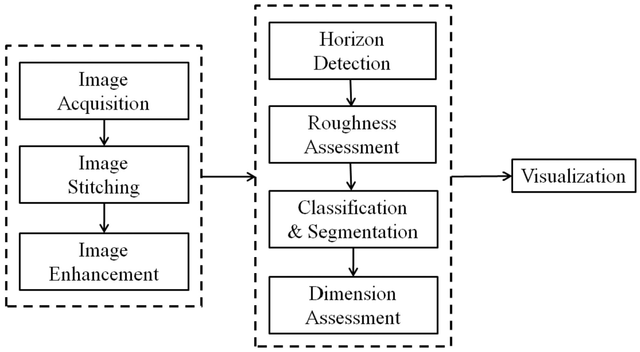

Fig. 1. Flow diagram of proposed automatic safe landing-site detection system.

feasibility of the vision-based system in Section IV, followed by conclusions in Section V.

II. METHODS

The safeness of an emergency landing-site is mainly determined by its surface roughness and dimensions. In general the roughness of the terrain can be measured by the gradient of elevation. If we have the elevation map of the terrain, the gradient information can be easily found, and the safeness can be accurately estimated. However, in this specific scenario, safe landing is not only determined by the elevation variation of the land but also threatened by hazards upon the ground, i.e., trees, rocks, vehicles, etc., which are usually not captured in elevation maps. Therefore, a vision-based information channel is necessary, which provides real-time imagery of the ground. Ideally, when the aircraft is flying in the upper air, it can be guided to an approximately smooth area according to the gradient information extracted from the elevation map. Then, the proposed computer-aided-detection (CAD) system leads the aircraft to a safe landing-site. In practice most aircrafts do not have either a database of elevation maps or a LIDAR sensor system. The imagery captured by aircraft-mounted cameras is the only available information source, so the proposed CAD system plays a crucial role in this scenario.

The proposed safe landing-site detection system consists of eight main modules as shown in Fig. 1. In the first module images are acquired by aircraft-mounted cameras. Each camera looks in a specific direction that covers a portion of the region in front of the airplane. Multi-spectrum sensors are preferred to obtain complementary information. In the second module the separate images that are acquired at the same time instant are registered and stitched together to form a larger panorama image that covers the full FOV in front of the airplane. In the third module, if the images are captured under poor illumination or weather conditions, we make use of the nonlinear retinex image-enhancement method [23, 24] to ameliorate the effect of environmental factors and to improve the contrast and sharpness of the images. The first three modules are necessary for getting high-quality images and directly affect the performance of the subsequent modules. However, we do not further discuss the preprocessing procedures in this paper since the emphasis of this paper is not on the development of the preprocessing procedures but on the safe landing-site detection algorithm.

A. Horizon Detection

Before assessing the roughness of the ground, the first problem that we need to solve is to identify where the ground is when the sky and the ground both appear in the image. Many efforts have been conducted in horizon detection. Williams and Howard proposed a horizon detection algorithm for a specific ground-based rover application of segmenting the foreground plane from distant mountains and the sky in glacial environments [25]. Due to the specialty of that application, the following two strong but reasonable assumptions were made in the algorithm. 1) It is assumed that the bottom third of the image is ground because the camera is mounted on a ground-based rover. 2) The ground is assumed to appear all white with very little variance because the rover is in glacial environments. Based on these two assumptions, the edge map, generated by applying a Canny edge detector to the original image, is examined column by column. An edge point in a column can be considered a point of the horizon when the pixels below it in that column appear all white with little variance. The problem of this algorithm is that the two assumptions often fail in other environments, so we do not further discuss it.

Dusha, et al. [26] applied the Hough transform to recognize straight lines from the binary edge map, which is also generated by using Canny edge detection, based on the assumption that the horizon is the strongest boundary in the image. That assumption does not always hold and can be easily disturbed by the appearance of other strong edges.

Ettinger, et al. proposed a horizon detection method in a greedy search manner [27] based on two assumptions: 1) the horizon is a straight line that partitions the image into two parts, namely sky and ground, and 2) little variance appears in either part, i.e., pixels of the sky part look like pixels of the sky part and not like pixels of the ground part, and vice versa. Thus, the detection of the horizon becomes to search for the optimal straight line such that the sum of the variances of both parts reaches the lowest value. The lines of all possible locations and angles are tested, and the optimal one that meets the above criteria is considered the horizon. Two concerns of that method are found: 1) it is computationally expensive due to the greedy search scheme, and 2) the second assumption fails when a part of the sea,

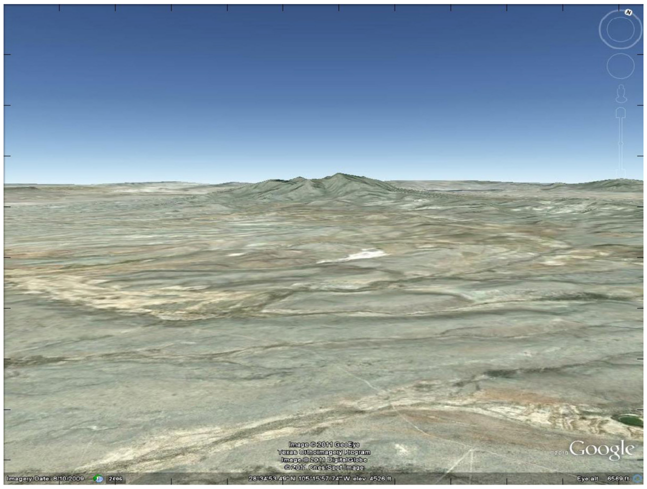

Fig. 2. Sample image.

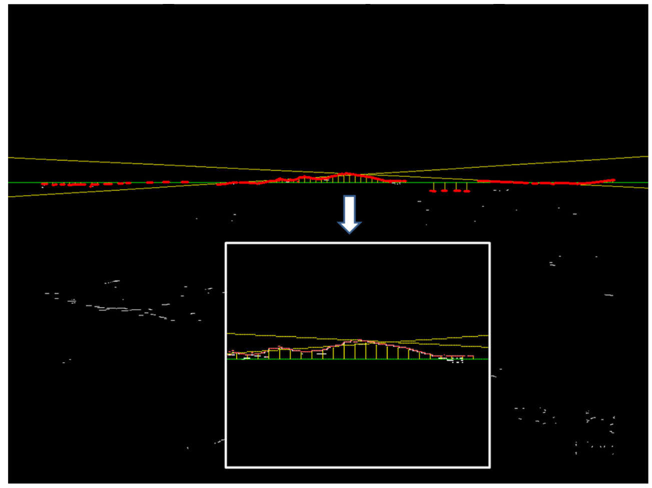

Fig. 3. Coarse adjustment.

rivers, or anything that has similar colors to the sky appear in the ground part of the image.

In the present paper a hierarchical elastic horizon detection algorithm is proposed to provide a robust and efficient way, which is detailed as follows. First of all the original image, as shown in Fig. 2, is blurred by a low-pass filter, i.e., a Gaussian low-pass filter with a large sigma value, so that all the fine edges are ignored and only the strongest bounds remain. Secondly, an edge detector is utilized to find the major bounds. In this paper we make use of the Canny edge detector [14] because it can provide edge-strength information in addition to edge-location information. Based on the edge strength of each pixel of the image, an edge-strength histogram is computed, and the top $p\%$ of the points are obtained as possible points that comprise the horizon. It is worthy to note that the edge-strength threshold p is adaptive for different images. On one hand the value of p should be a small number because, by experience, we know that the horizon is the strongest bound in the image most of the time. On the other hand it is not always the case, so the value of p should not be too small. In other words we would rather conservatively include some nonhorizon points in this step and exclude them in a later step than hastily lose some horizon points in this step. By experiments we set

$$p = \frac{1}{N_H} \times 500 \tag{1}$$

where N_H is the total number of columns in the image, which varies among images. In general images in this particular application usually consists of hundreds or thousands of columns so that the value of p can be guaranteed to take values that are less than or equal to 100. Based on the p value and on the corresponding edge-strength threshold, a binary map can be generated. Then, the standard Hough transform (SHT) [28] is applied to search probable lines in the binary map. Ideally, the horizon should be the highest peak of the voting result of the SHT so that, by taking the highest peak value, we can find the slope and the intercept of the horizon in the image. However, the horizon is not always the highest peak in the SHT result. Therefore, we cannot just simply take the highest peak as the horizon. Instead, we take into account the top N_L highest peaks by comparing the average edge strength within the dual-side narrow bands along the N_L lines. The line that has the highest average edge strength within the dual-side narrow bands is called the true peak of the Hough transformation. By experiments, we set N_L to 5. This strategy makes the detection of the horizon reliable and robust. In Fig. 3, the green line is the true peak of the Hough transformation along with two yellow lines, which are among the top N_L highest peaks but not the true peak. Since the horizon is not often a strictly straight line, the Hough transformation result may not perfectly match the real, curvy horizon. Therefore, it is necessary to adjust the line in a pixel-wise manner by searching the dual-side neighborhood of each pixel. The coarse adjustment, as shown in Fig. 3, is applied based on the edge map of the blurred image. The subfigure inside Fig. 3 provides a closer view at a local segment. The yellow arrows denote the directions of adjustment, and the red dots denote the position of each horizon-point after the coarse adjustment.

After coarsely finding the position of the horizon, we get back to the original image and do a fine adjustment to accurately locate the horizon, as shown in Fig. 4. For each pixel of the coarsely-found horizon rendered in green, we search its dual-side narrow band in a fine-edge map rendered in a gray scale, which is computed from the original image, and then slightly adjust its position to the pixel that has the largest edge strength of that neighborhood. Yellow arrows point to the directions of adjustment, and red points denote the position after the fine adjustment. Noisy points may emerge due to the discontinuity of the horizon in the

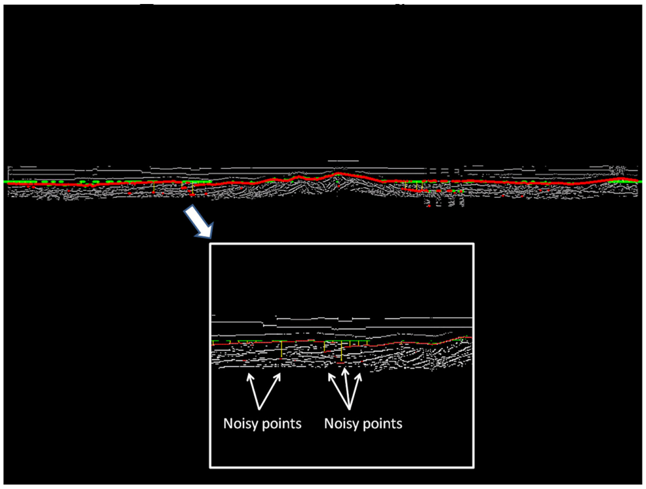

Fig. 4. Fine adjustment.

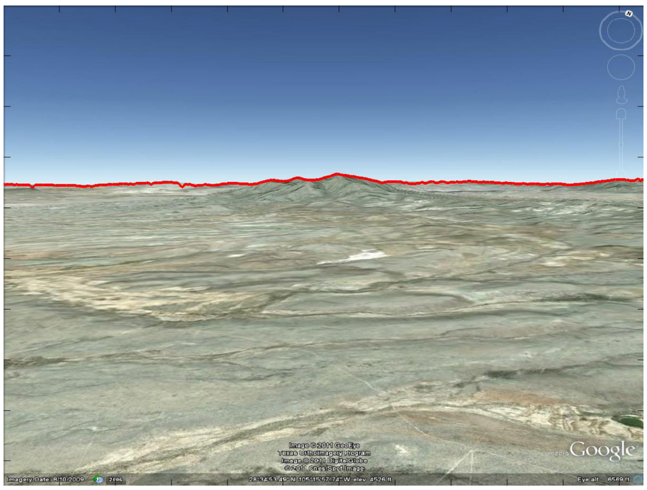

Fig. 5. Detected horizon.

fine-edge map. We remove these discontinuous points by interpolation based on their neighboring points on the left and right. In this paper we use the B-spline interpolation method [29, 30]. After removing all the noisy points, a smoothing technique is applied to local segments of the detected horizon to get the smooth final detection result as shown in Fig. 5. Then, the roughness assessment can be applied to the ground part in the image.

B. Roughness Assessment

The roughness of the ground and the presence of hazards are often reflected as boundaries and as a high-variance of pixel intensity values in visible images. If high-resolution elevation maps are not available, it is plausible to assume that identifying rough areas or hazardous objects on the ground is equivalent to the process of edge detection in visible images. The Canny edge detector [14] is an efficient tool for computing the sharpness of edges, which is, from smoothest to sharpest, quantified to the range from 0 to 255. This method is applied at

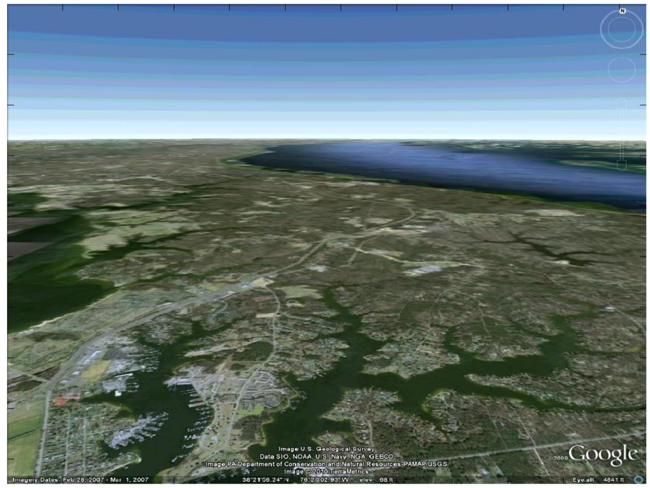

Fig. 6. Sample image.

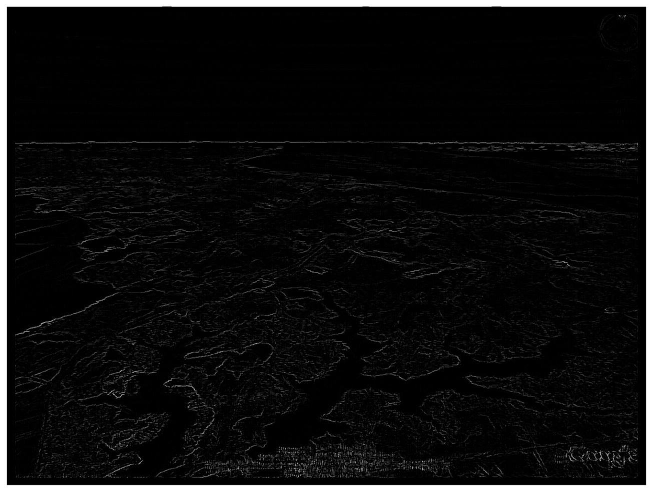

Fig. 7. Edges found with Canny detector.

the beginning of the roughness assessment module. Figure 6 shows a sample image provided by Google Earth®, and its edge detection result using the Canny detector is shown in Fig. 7. Brighter pixels represent sharper edges, and vice versa.

Different edge patterns appear among diverse regions in terms of the edge strength within a certain range. To characterize the difference the edge map is first divided into nonoverlapping blocks. We define the cumulative hazard strength (CHS) of each block as follows

$$\mathrm{CHS}_B = \sum_{p \in B} H(\mathrm{ES}_p) \tag{2}$$

$$H(\mathrm{ES}_p) = \begin{cases} 1 & \mathrm{ES}_p > T \\ 0 & \mathrm{ES}_p \le T \end{cases} \tag{3}$$

where ES_p is the edge strength of each pixel p in block B, and $H()$ is the hazard-indicator function. If ES_p is greater than the prespecified safeness threshold T, then the pixel p is considered hazardous, and the CHS of block B, CHS_B, is incremented by 1. In contrast, if ES_p is no greater than T, then the pixel p

is considered safe, and CHS_B remains the same. Thus, blocks of smooth areas have a zero or low CHS value, but blocks of rough areas have a high CHS value. The block size (BS) in the unit of pixels is adaptively determined based on the height of the camera h_c in the unit of ft

$$\mathrm{BS} = \begin{cases} 20 \times 20 & h_c \leq 10000 \text{ ft} \\ 15 \times 15 & 10000 \text{ ft} < h_c \leq 20000 \text{ ft} \\ 10 \times 10 & h_c > 20000 \text{ ft} \end{cases} . \quad (4)$$

For example, if the aircraft is flying at a higher elevation, the image covers a relatively larger area on the ground, and the realistic size of each pixel is relatively larger is compared with the image captured at a lower height. As a result, to keep the consistency of the realistic area of each block to some extent, the BS is set as a smaller number when the image is captured at a higher height, and vice versa. In addition the prespecified safeness threshold T is related to the requirement of acceptable smoothness. A lower value of T means a stricter requirement for smoothness because the edge strength of more pixels will be beyond the safeness threshold, and they will be considered hazardous. It is more reasonable to utilize a unified strict safeness threshold rather than an adaptive safeness threshold according to the change of h_c, because loosing the requirement for smoothness as the h_c increases will bring risk to the landing process. It is empirically determined that T is not sensitive to the final detection results if its value is picked in the range from 15 to 30 since the edge strength of the hazards' boundaries is usually much higher than 30. In this paper, T is selected to be 20 as a relatively stricter requirement for smoothness.

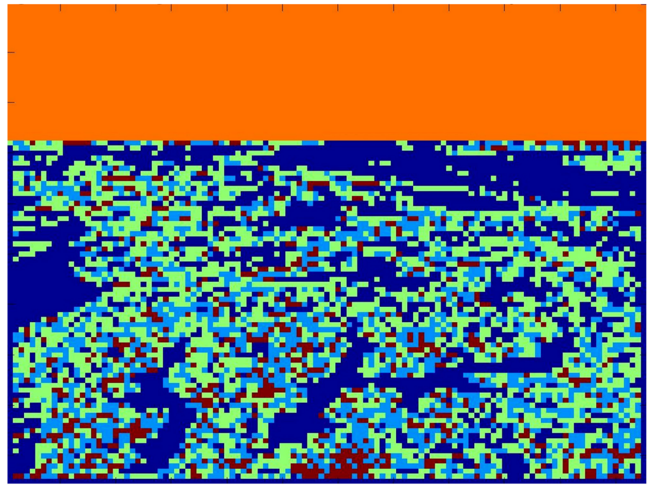

Fig. 8. Clustering result based on CHS.

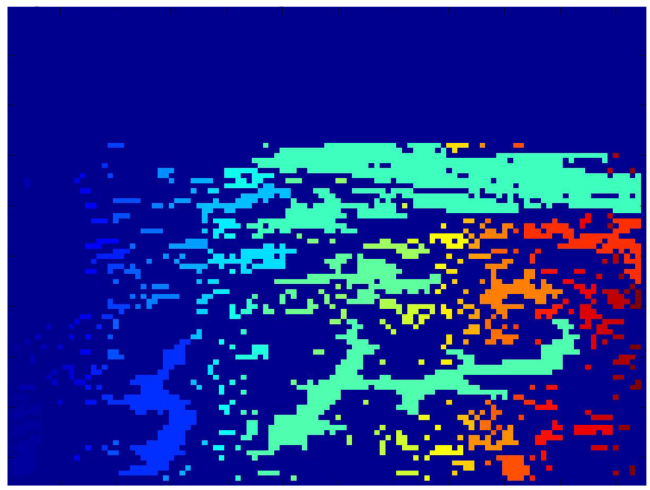

Fig. 9. Multi-region growing result.

C. Classification and Segmentation

The classification module utilizes the K-mean clustering method [31, 32] to classify the CHS of each block into a number of clusters. For example, if the number of clusters is specified as seven, the clusters can be interpreted as "very rough," "rough," "moderate rough," "median," "moderate smooth," "smooth," and "very smooth." The number of clusters is, first, set to seven by default and then, automatically, reduced in the clustering procedure. That is, if any cluster loses all of its members, that cluster will be removed [31, 32]. Figure 8 shows the clustering results of the sample image shown in Fig. 6. In this case four clusters are obtained: dark blue renders the smoothest areas, red renders the roughest areas, and green and light blue represent the areas in between. Based on the clustering result, the adjacent "smoothest" blocks are merged to form larger, smooth areas by using the region-growing method [33]. The result of connected areas is shown in Fig. 9, where each area is labeled with a unique color. For the concern of efficiency, isolated tiny spots and narrow branches of merged areas can be removed by applying the morphological operation of image erosion [33] without assessing their dimensions since they are obviously undersized.

D. Dimension Assessment

After the above steps potential landing-sites are obtained as shown in Fig. 10. In this module we measure their realistic dimensions and determine which are qualified to be candidate landing-sites. The realistic dimensionality of each potential landing-site is measured by converting its size from the image coordinate system to the realistic world coordinate system. In flight dynamics changing the orientation of the aircraft to any direction can be decomposed to three kinds of rotations: yawing, rolling, and pitching, which are, respectively, to rotate the aircraft along the vertical axis, the longitudinal axis, and the lateral axis. Given those three rotation angles, this procedure can be described by the intrinsic or extrinsic matrices composition [34, 35] with

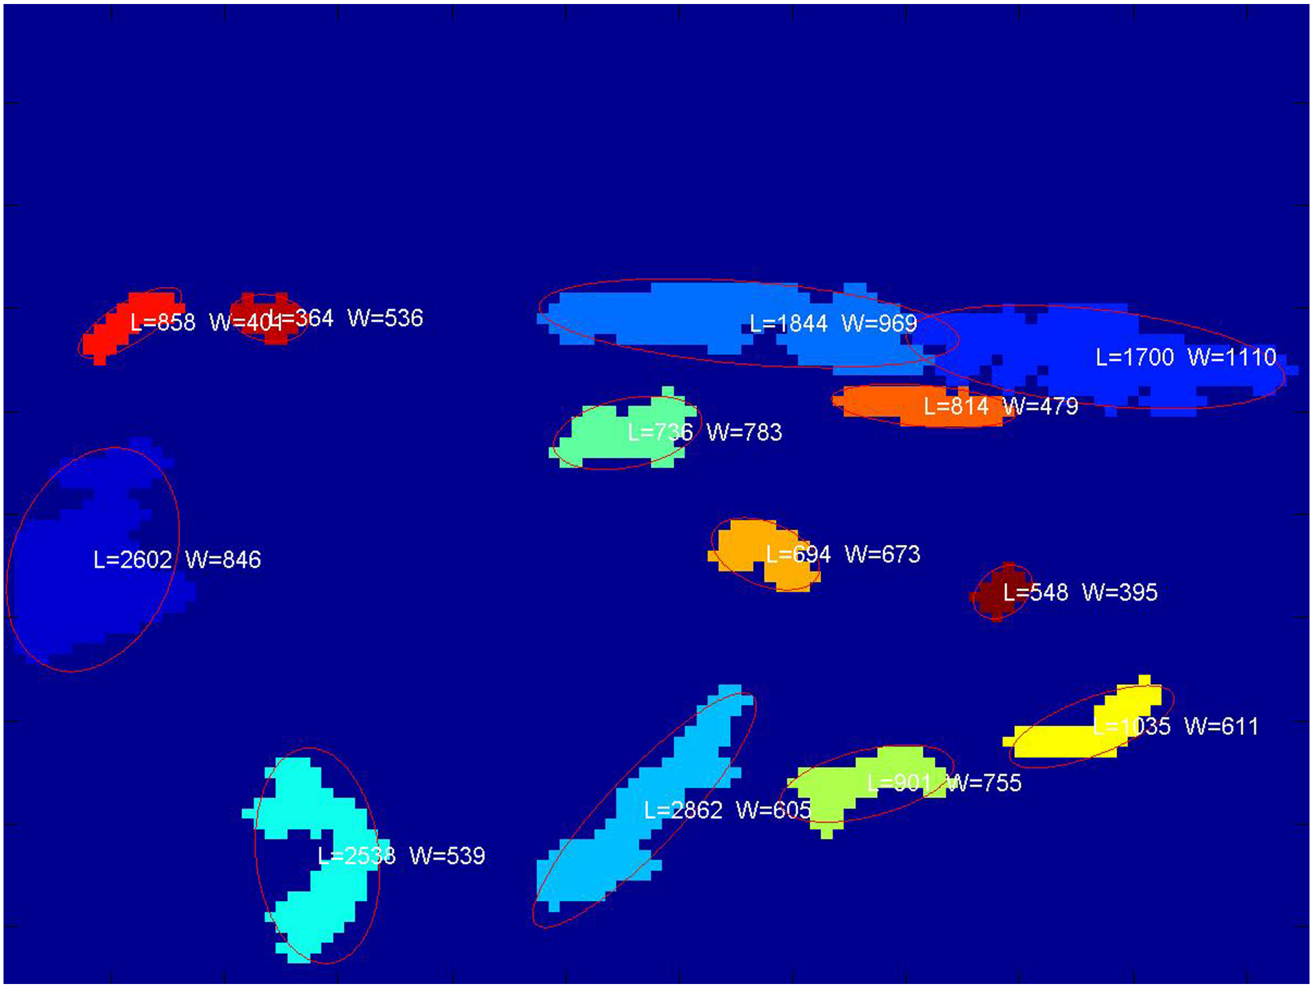

Fig. 10. Potential landing-sites identified and realistic dimensions measured.

which one can map the world coordinate system to the aircraft coordinate system, and vice versa. In other words two arbitrary points in an aerial image can be mapped to the world coordinate system so that the realistic distance between the two points on the ground is measurable if the three rotation angles are known. In practice most aircrafts have the necessary equipment to record the three angles so that they can be synchronically stored with real time aerial images. We use images provided by Google Earth® in this pilot study. Because of the lack of the information of the three rotation angles, we simplify the imaging process with only pitching but no yawing and rolling. Therefore, the imaging model in the vertical direction of the image coordinate system can be described as shown in Fig. 12. Then, the realistic size of each pixel along the vertical direction of the image can be computed as follows

$$d_0 = h_c \cdot \tan\left(\alpha - \frac{\mathrm{FOV}_V}{2}\right)$$
$$d_i = h_c \cdot \tan\left(\alpha - \frac{\mathrm{FOV}_V}{2} + \sum_{j=1}^{i}\theta_j\right) - \sum_{k=0}^{i-1} d_k \tag{5}$$

where h_c is the height of camera, α is the pitching angle, FOV_V is the FOV along the vertical direction of the image, N_V is the total number of pixels along the vertical direction, d_0 is the distance between the vertical line and the first pixel, and θ_i and d_i $(i = 1,2,\ldots,N_V)$ are, respectively, the angle and the realistic distance corresponding to pixel p_i along the vertical direction. For large N_V $\theta_i (i = 1,2,\ldots,N_V)$ can be considered to have the same approximate value θ so that (5) can be simplified as

$$d_i = h_c \cdot \tan\left(\alpha - \frac{\mathrm{FOV}_V}{2} + i\theta\right) - \sum_{k=0}^{i-1} d_k$$
$$\theta = \frac{\mathrm{FOV}_V}{N_V}. \tag{6}$$

In addition, since it is assumed that there is no yawing or rolling rotation, the realistic size of the pixels along the horizontal direction of the image is the same

$$d_H = \frac{2h_c}{N_H}\tan\left(\frac{\mathrm{FOV}_H}{2}\right) \tag{7}$$

where FOV_H is the FOV along the horizontal direction of the image and where N_H is the total number of pixels along the horizontal direction.

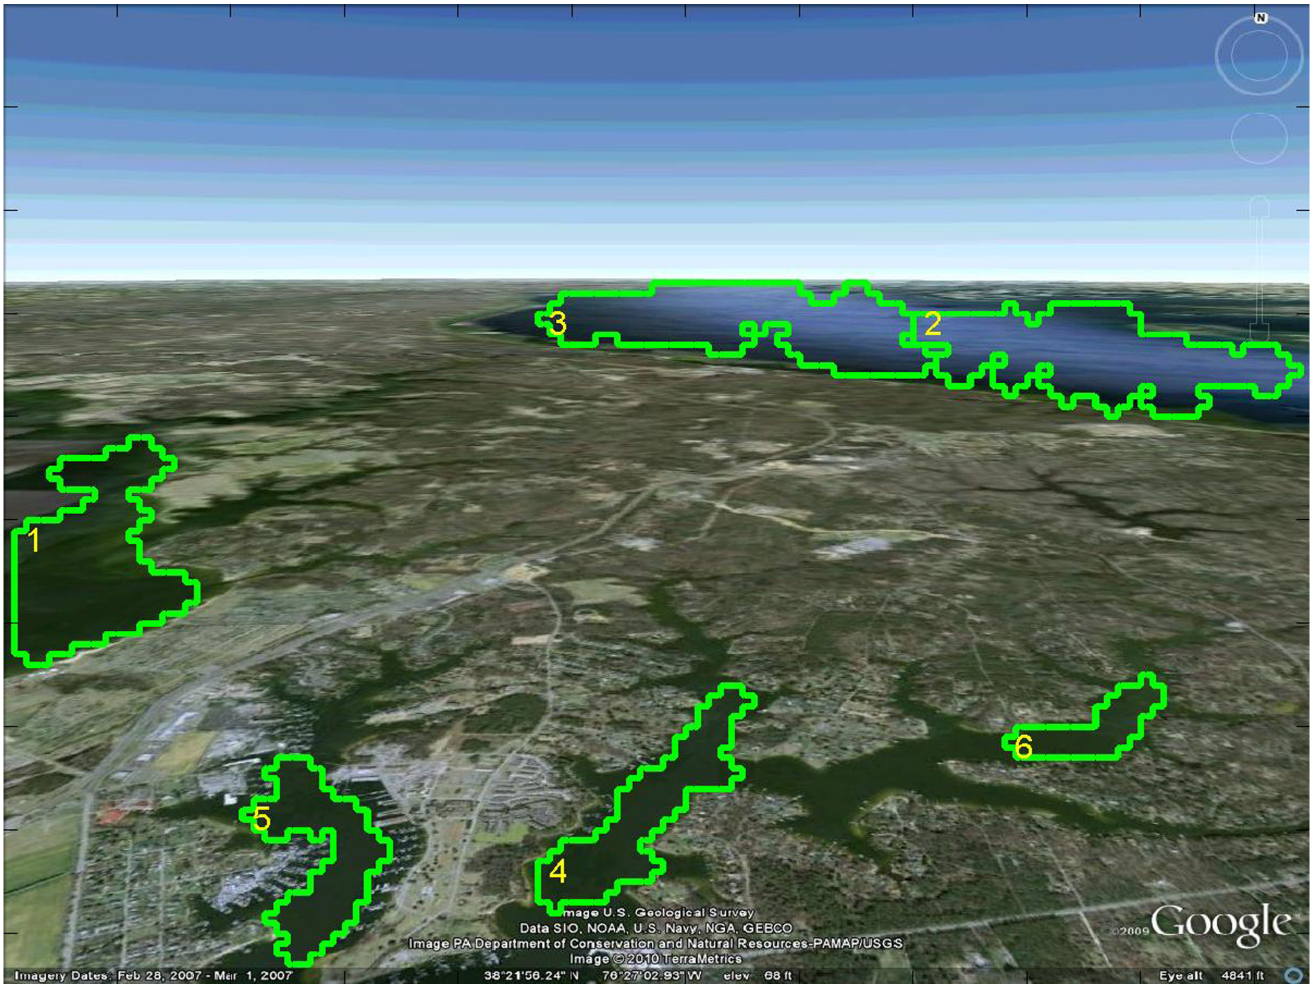

Fig. 11. Landing-sites sorted in descending order based on areas for pilot to evaluate candidates.

The dimensions of each potential landing-site are estimated by measuring the major axis and minor axis of its best fit ellipse, which are obtained using the principle component analysis method [36]. Once the major axis and the minor axis are found, the realistic length L and the width W in the unit of feet can be gained as

$$L = \sqrt{(d_H(x_{a2} - x_{a1}))^2 + \left(\sum_{k=y_{a1}}^{y_{a2}} d_k\right)^2}$$
$$W = \sqrt{(d_H(x_{b2} - x_{b1}))^2 + \left(\sum_{k=y_{b1}}^{y_{b2}} d_k\right)^2} \tag{8}$$

where (x_{a1}, y_{a1}), (x_{a2}, y_{a2}) are coordinates of the two end-points of the major axis and (x_{b1}, y_{b1}), (x_{b2}, y_{b2}) are coordinates of the two end-points of the minor axis, in the image coordinate system. Figure 10 shows the length and the width of each potential landing-site in the unit of ft. Small areas with insufficient length or width are ruled out, and only large areas with sufficient length and width can be considered safe emergency landing-sites.

E. Visualization

The visualization module is designed to highlight, at most, the five largest safe landing-site candidates on the human-machine interface for the pilot's final decision, though the system may detect more than five safe landing-sites. If the system provides the pilot with all the possible choices, he may get confused when seeing too many recommended areas on the screen, and the time cost of making a decision is very critical under the emergency situation. Therefore, only up to five largest candidate landing-sites are visualized on the human-machine interface and labeled with preference indices. The landing-sites are sorted in a descending order based on their areas as shown in Fig. 11, so that the pilot can efficiently evaluate the recommended candidates in a rational order. The pilot will make his final decision by choosing one emergency landing-site from the recommended candidates and by taking into account other factors as well, i.e., wind direction, wind speed, maneuvering ability, etc. In general larger areas are preferable when compared with smaller ones.

F. Performance Metric

To quantitatively evaluate the results generated by the proposed system, we ask two veteran professional

TABLE I
Four Possible Results of the Detection of Landing-Sites

#	Categories	Ground-Truth Results	CAD System Results
1	True Positive	The area is **safe** for emergency landing and labeled as a candidate.	Hazards are **not found** in the area small capsand its dimensions **equal or exceed** the requirement.
2	False Positive	The area is **not safe** for landing.	Hazards are **not found** in that area small capsand its dimensions **equal or exceed** the requirement.
3	True Negative	The area is **not safe** for landing.	Hazards are **found** in that area small capsor its dimensions are **below** the requirement.
4	False Negative	The area is **safe** for emergency landing and labeled as a candidate.	Hazards are **found** in that area small capsor its dimensions are **below** the requirement.

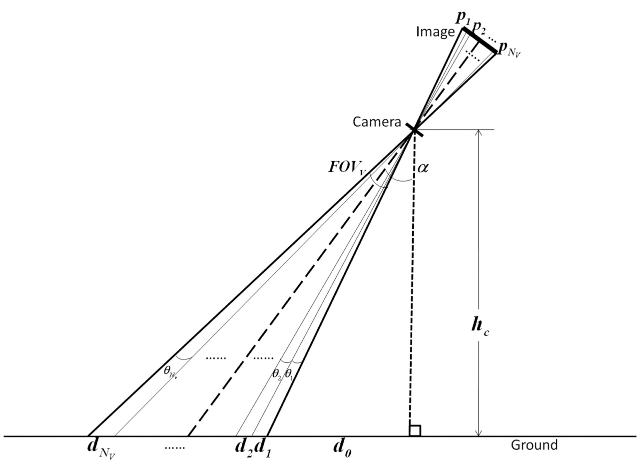

Fig. 12. Simplified imaging model.

pilots to manually pick all the possible landing-sites in the original images. Their judgment is mainly based on the apparent smoothness of the areas shown in the images. Next, the realistic dimensions of these manually-selected areas are measured by using the same dimension assessment module mentioned above. This step is necessary because it is hard to accurately estimate the length and the width of candidate landing-sites in the images captured at different heights by just looking at them. As shown in the left column of Fig. 15, if the dimensions of a selected area meet the minimum requirement, it is labeled in green as a safe landing-site. Otherwise, it is labeled in red as an unsafe landing-site. After manual selection all the selected regions are sorted in a descending order according to area. To fully evaluate the performance of the proposed CAD system, the manually-selected and labeled regions are utilized as the ground-truth, with which we compare the complete detection results produced by the proposed CAD system, that is, the results before the visualization module, including the five largest safe landing-sites shown on the human-machine interface and the remaining smaller ones omitted in the visualization module.

If an area is found as a candidate landing-site by the proposed CAD system, the detection result of this area is considered positive. Furthermore, if this detection result is consistent with the ground-truth, it is called a true positive (TP) detection. Otherwise, it is a false positive (FP) detection. Similarly, if an area is not selected as a candidate landing site by the proposed CAD system, the detection result of this area is considered negative. In addition, if this detection result is consistent with the ground-truth, it is called a true negative (TN) detection. Otherwise, it is a false negative (FN) detection. Table I lists the interpretations of the four exclusive and exhaustive situations. TP and TN are desired correct diagnoses. FP and FN are wrong diagnoses and have to be eliminated. It is worth noting that an FP is the worst situation since it can mislead the aircraft to a dangerous place.

A scoring mechanism is proposed to quantitatively evaluate the performance of the proposed CAD system

$$\begin{aligned} S &= S_0 + S_B - S_P \\ S_0 &= \sum_{j=1}^{N} p_j, \qquad S_B = \sum_{i=1}^{M} b_i T_i \\ S_P &= \sum_{j=1}^{N} p_j F_j. \end{aligned} \tag{9}$$

For each test image the score S consists of three parts: bonuses S_B for TP detections, penalties S_P for FP detections, and the base score S_0. In the ground-truth of each test image, there are M manually-selected areas. According to the priority index, different bonus weights b_i $(i = 1,\ldots,M)$ are given to these M areas in a descending order of size. If the CAD system successfully detects the ith largest safe landing-site in the ground-truth, then bonus b_i is earned. T_i is the flag that indicates if the ith largest area in the ground-truth was successfully detected by the system. In addition to the reward mechanism, a punishment mechanism is also used. The complete detection results produced by the system before the visualization module contain N candidates, where N may be greater than five and different from M. Penalty p_j $(j = 1,\ldots,N)$ is imposed if the jth recommended candidate landing-site is an FP detection. F_j is the flag indicating if the jth detected area is an FP detection. The initiative of using bonus and penalties with different weights is to emphasize the priority of each safe landing-site in the ground-truth. In general the larger the dimensions of the landing-site are, the easier and safer the forced landing process is. Therefore, it is reasonable to give higher bonuses for TP detection of larger safe landing-sites. Also, since an FP detection result labeled as a higher priority has more negative effect than an FP detection result labeled as a lower priority, it is reasonable to impose higher penalties to the former. In this paper we set b_i, respectively, with 30, 25, 20, 15, 10 for $1 \leq i \leq 5$ and 5 for $i \geq 6$. In the same way we set p_j, respectively, 30, 25, 20, 15, 10 for $1 \leq j \leq 5$ and 5 for $j \geq 6$. The base score S_0 is the potential maximum of penalties that a set of detection results can get. It is used to guarantee that S is nonnegative, even if in the worst case all of the N detection results are FP, $S = 0$ is the lowest score that the CAD system can get. Since N can significantly vary among images, the scores need to be normalized to a unified range. The normalized score is computed as follows,

$$\hat{S} = 100\left(\frac{S}{S_{FM}}\right), \qquad S_{FM} = S_0 + \sum_{i=1}^{M} b_i \tag{10}$$

where $\hat{S}$ is the normalized score that is obtained by normalizing S using S_{FM}. S_{FM} is the possible full score for each experiment when all detection results are consistent with the ground-truth. After normalization, scores range from 0 to 100: a score of 100 shows a perfect match, while lower scores show decreasing matches between the ground-truth and the CAD results. Equation (10) can also be interpreted as follows

$$\begin{gathered} \hat{S} = 100\left(\frac{S' - S_{\min}}{S_{\max} - S_{\min}}\right) = 100\left(\frac{S' + S_0}{S_{FM}}\right), \\ S_{\min} = -\sum_{j=1}^{N} p_j, \qquad S_{\max} = \sum_{i=1}^{M} b_i \\ S' = S_B - S_P, \qquad S_{FM} = S_{\max} - S_{\min} \end{gathered} \tag{11}$$

where S' is the actual performance of the N detection results, $S_{\min}$ and $S_{\max}$ are, respectively, the lower and upper bound of S'. $S' - S_{\min}$ is equivalent to $S' + S_0$ shown in (9), both of which transfer S' from the interval $[S_{\min}, S_{\max}]$ to the interval $[0, S_{FM}]$ so that S is guaranteed to be nonnegative. Then, $\hat{S}$ is obtained by normalizing S from the interval $[0, S_{FM}]$ to the unified interval $[0, 100]$.

Equation (10) fails under two scenarios. 1) For images captured above rough terrains, there may be no safe landing-sites in the ground-truth ($M = 0$, $S_{\max} = 0$), so the best corresponding detection result should be no recommendations ($N = 0$, $S_B = 0$, $S_0 = 0$, $S_P = 0$). For this special case (10) is ill defined since its denominator is 0. 2) When all the detection results are FP ($S_B = 0$, $S_0 = S_P$), (10) fails to differentiate results with different numbers of FPs since the numerators of those situations are all zero. For example, suppose there is one safe landing-site in the ground-truth ($M = 1$, $S_{\max} = 30$), the detection result of zero TP and zero FP ($S_B = 0$, $S_0 = 0$, $S_P = 0$) should be better than the result of zero TP and two FPs ($S_B = 0$, $S_0 = 55$, $S_P = 55$). However, based on (10), $\hat{S}$ of both the above is 0, which means it fails to differentiate between these two scenarios.

To solve these two problems, a correction is made to (10) by adding a small augment α to the numerator and denominator,

$$\hat{S} = 100\left(\frac{S + \alpha}{S_{FM} + \alpha}\right), \qquad S_{FM} = S_0 + \sum_{i=1}^{M} b_i. \tag{12}$$

α should be a relatively small number so that it has little effect on the ratio of the non-zero numerator and the denominator when there is at least one TP in the detection results ($S_B > 0$). In this paper α is 1. By using (12) both problems are solved: in scenario 1, $\hat{S}$ gets the expected score 100; in scenario 2, the normalized scores of the two scenarios are, respectively, 3.2 and 1.2. Thus, the two scenarios get different scores. Though both are relatively low, "no recommendations" is better than false recommendations in this particular application.

III. EXPERIMENTAL RESULTS

The validation of the proposed system consists of three parts in this pilot study.

A. Experiment 1

We tested the reliability of the proposed hierarchical horizon detection algorithm on 108 sample images provided by Google Earth®. The horizon in the 108 sample images appears in various angles, and those images are captured at different elevations, which range from 1000 ft to 30000 ft and over different types of terrains. We also tested on the same set of images with the Hough transform method [26], the greedy search method [27], and the simplified version of the proposed method in which the finer adjustment step is omitted. To quantitatively compare the accuracy of the above horizon detection methods, we develop a measurement, average maximum bias (AMB) defined as follows,

$$\text{AMB} = \frac{1}{H_M}\sum_{h=1}^{H_M} MB_h, \qquad H_M = L_H \times 0.1 \tag{13}$$

where L_H is the total pixel number of the detected horizon, H_M is one-tenth of L_H, and MB_h ($h = 1, 2, \ldots, H_M$) is the hth maximum bias from the detected horizon to the position of the true horizon

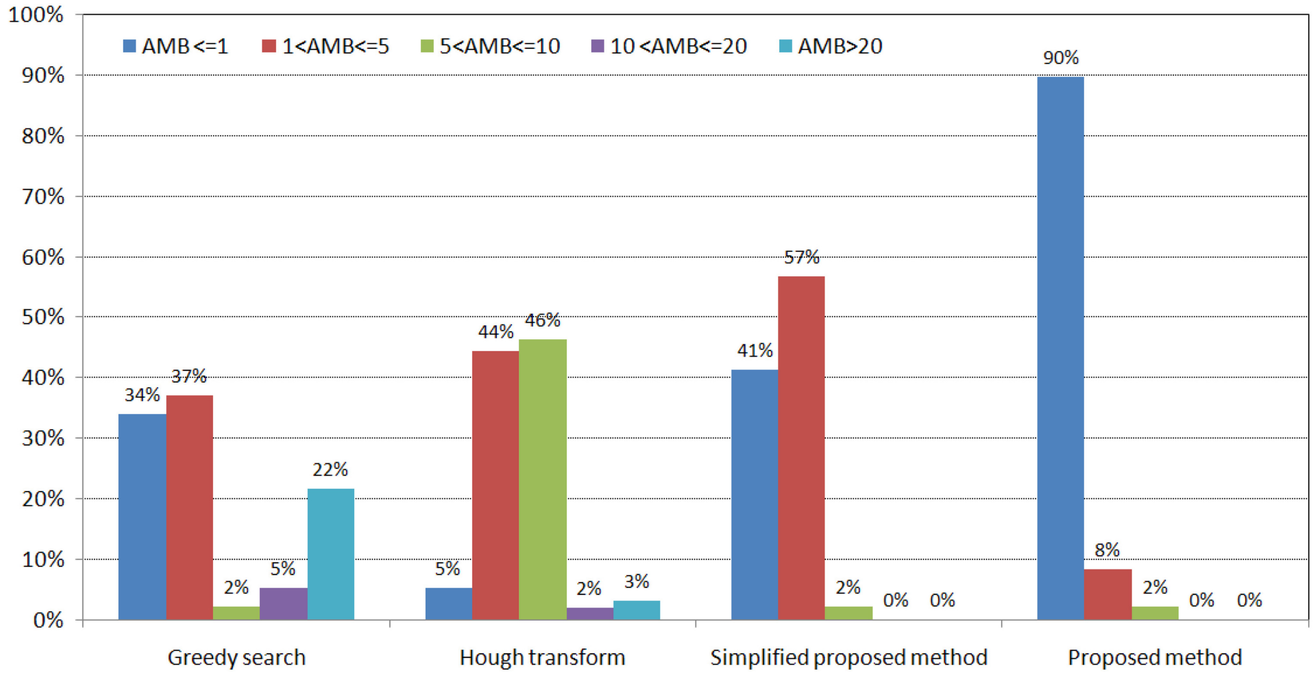

Fig. 13. Horizon detection results by using greedy search method [27], Hough transform [26], simplified version of proposed method, and proposed method.

in the unit of pixel. The reason that we are only interested in the top 10% of the maximum bias is because the horizon is often long. If we compute the average bias for the whole detected horizon, some significant bias may be hidden by other well-aligned parts of the detected horizon when the averaging operation is taken. Therefore, evaluating the most biased segment of the detected horizon can tell us the true performance of the method. In other words, if the most biased segment of the detection result can be considered satisfactory, the rest of the detected horizon can be guaranteed to be better than the most biased segment. Figure 13 shows the horizon detection results by using four methods, respectively, the greedy search method [27], the Hough transform method [26], the simplified version of the proposed method, and the proposed method. Five AMB intervals are used, which are, respectively, perfect detection (AMB ≤ 1), good detection ($1 <$ AMB ≤ 5), acceptable detection, ($5 <$ AMB ≤ 10), biased detection ($10 <$ AMB ≤ 20), and false detection (AMB > 20). The percentage shown above the bar of each AMB interval reflects the ratio of the number of images, for which the AMB value falls into that AMB interval, to the total number of tested images. For the greedy search method, it achieves perfect and good detection (34% perfect detection and 37% good detection), which means that the method successfully finds the true horizon in 71% of the images, though it fails in 22% of the images when a part of sea, rivers, or anything that has similar colors to the sky appears in the ground part of the image. The reason that it cannot achieve perfect detection is because it assumes that the horizon is a straight line so that it cannot fit the curvature of the true horizon, though the general position of the true horizon is detected. For the Hough transform method, due to its intrinsic mechanism, it is easily corrupted by other nonhorizon edges shown in the image. This explains why it only achieves 5% perfect detection but a high rate of good and acceptable detection (44% and 46%). In addition, when the nonhorizon edges are even stronger than the horizon, the method fails as shown in 3% of false detection. The proposed method achieves 90% perfect detection and 8% good detection. Some results generated by the proposed method are shown in Fig. 14, in which the red line represents the detected horizon. To show the importance of its hierarchical detection strategy of coarse detection and fine adjustment on the final detection accuracy, we also tested a simplified version of the proposed method, in which the fine adjustment is omitted. By doing so the perfect detection rate significantly decreases to 41%, and the good detection rate jumps to 57%. It is also worth noting that both the proposed method and its simplified version produced no false detections on the test images. This robustness and reliability is achieved by the examination of the top N_L lines ($N_L = 5$ in this paper) generated by the Hough transform, instead of only by examination of the top one line, as in [26].

B. Experiment 2

The reliability and accuracy of the proposed system was validated on independent static sample images with projection angles of 0° and 60°. A total of 169 images captured at 1,000–30,000 ft, and 100

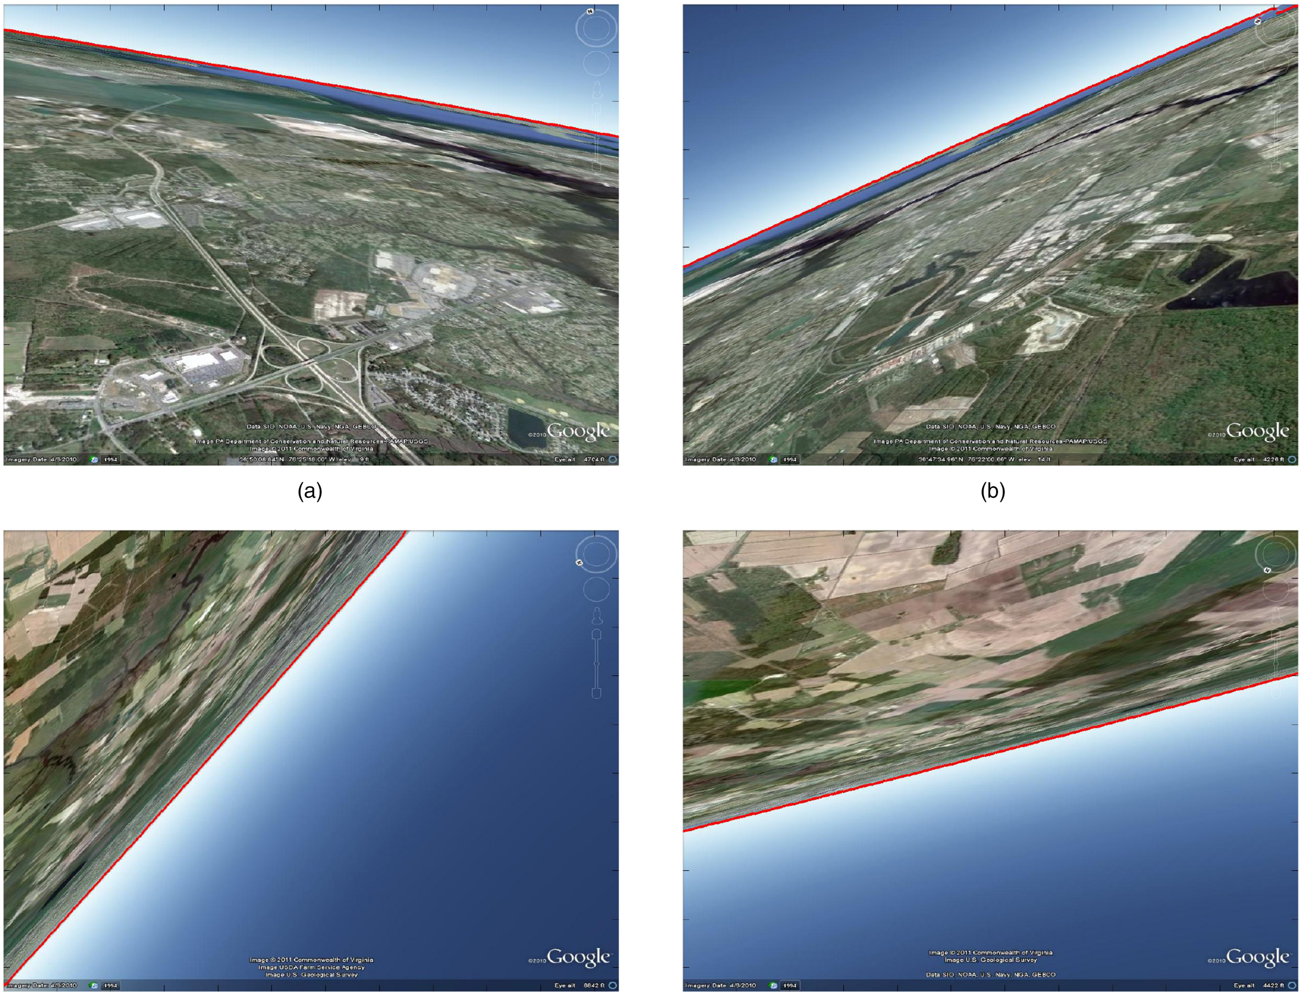

(a) (b) (c) (d)

Fig. 14. Results produced by proposed horizon detection algorithm.

images captured at 5,000 ft were used for testing the CAD algorithm. The left column of Fig. 15 shows five samples of the labeled manual selection, and the right column shows the corresponding results produced by the proposed CAD system. Two performance metrics are applied to evaluate the detection results. 1) A group-wise TP detection rate (GTPR) is defined as,

$$\mathrm{GTPR} = \frac{\sum_{i=1}^{N_{img}} DTP_i}{\sum_{i=1}^{N_{img}} MTP_i} \times 100\% \tag{14}$$

where DTP_i is the number of TP detections in the ith image, MTP_i is the number of safe landing-sites in the ith image provided by the ground-truth, and N_{img} is the total number of images in the testing set. The GTPRs of the 1,000–30,000 ft testing set and the 5,000 ft testing set are, respectively, 81.2% and 87.1%. 2) By using the performance metric presented in Section II-F, normalized scores of detection results of the two testing sets are obtained, and the distribution of the normalized scores is shown in Table II. 69.2% of the experimental results of the first set and 74.0% of the experimental results of the second set completely match the ground-truth ($\hat{S} = 100$), and 80.6% of the first set and 86.0% of the second set generally match the ground-truth ($\hat{S} \geq 60$), which demonstrates the feasibility of the proposed CAD system. We did notice that there were 18.3% failures in the first set and 14.0% in the second set ($\hat{S} < 5$). Two major causes are found. 1) Fewer details can be seen from higher elevations since the realistic area covered by each pixel in the image becomes larger when camera height increases, so some areas appear to be smooth in the image, while in reality they are not. This explains why the results of the 5,000 ft set is better than that of the 1,000–30,000 ft set. This image resolution problem is, of course, a characteristic of the image-capture device. 2) Artificial or pseudo boundaries may cause the system to miss safe landing-sites. For example, an area may be flat, but changes in its soil color or soil texture appear as sharp edges that can confuse the CAD system. This can be attributed to the assumption that elevation

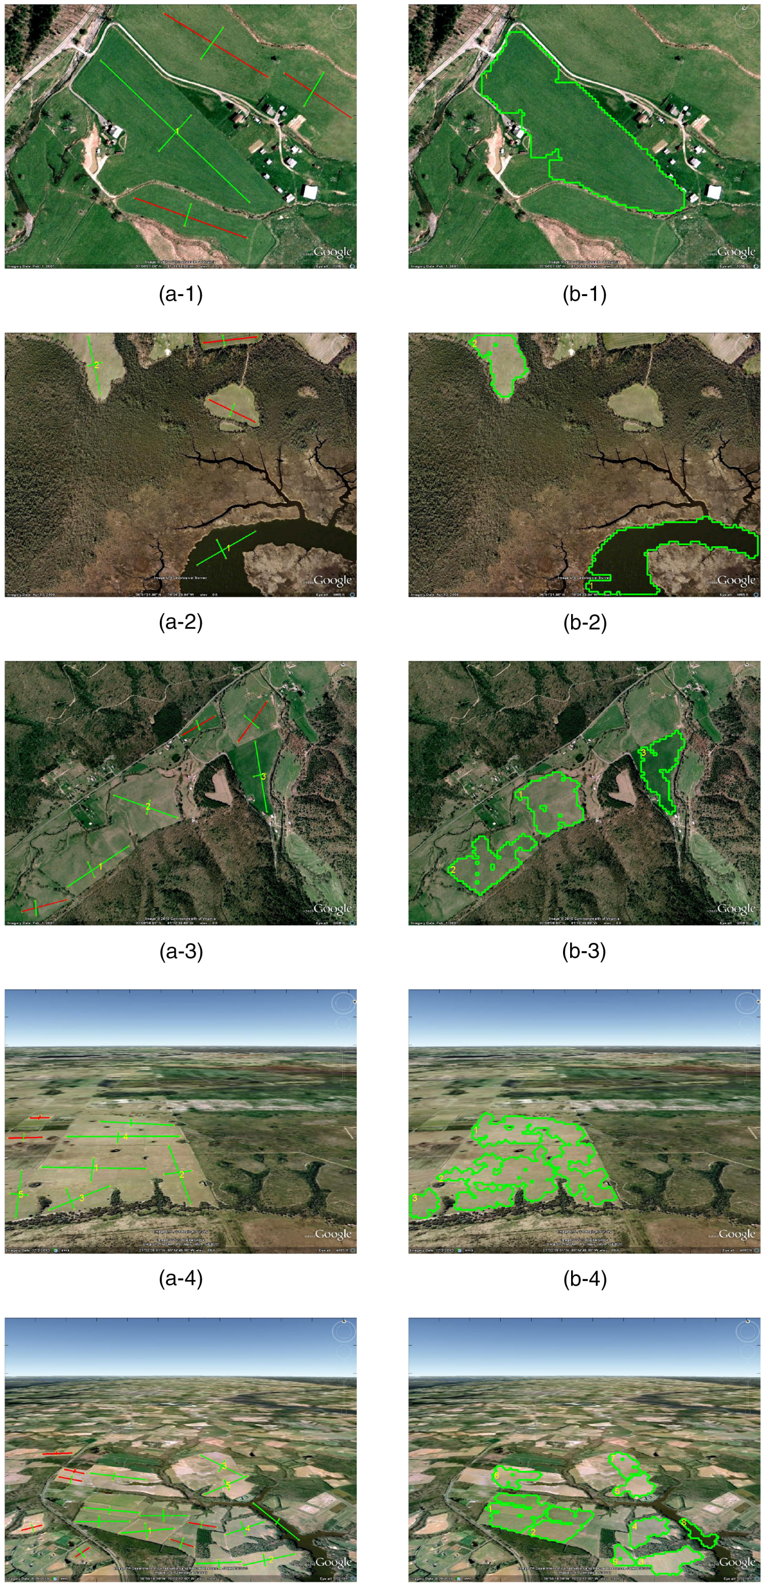

Fig. 15. Comparison between manual selection and automatic detection. (a) Manually-selected landing-sites. (b) Recommended landing-sites detected by proposed system.

changes can be mapped by edges in visible images. In the absence of elevation information (most small, general aviation aircraft do not have an elevation database on board), imagery captured by aircraft cameras is the only source for computer algorithms to evaluate the surface roughness. The proposed CAD system can play an important role under this situation.

TABLE II
Distribution of Normalized Scores of Experimental Results

Normalized Score	0	1	2	3	4	5–45	50	55	60	65	70	75	80	85	90	95	100
Percentage (%), 1 k–30 k ft	0	4.1	8.9	0	5.3	0	1.2	0	3.0	0	0.6	0.6	1.8	3.6	0.6	1.2	69.2
Percentage (%), 5 k ft	0	0	9.0	0	5.0	0	0	0	4.0	0	1.0	2.0	5.0	0	0	0	74.0

C. Experiment 3

The consistency of detection results is important because the proposed system eventually will be utilized as a real-time system to detect safe landing-sites as the aircraft moves. By applying the proposed system to 10 sequences of images, each of which has 10 to 15 images, the consistency of detection system is validated. The motion of the aircraft is reflected by the relative position of the sequential images. For example, if the aircraft moves along the direction parallel to the ground, the corresponding sequence of images covers a band along the trace of the aircraft. The distance between centers of every two consecutive images is determined by the flying speed. In other words two consecutive images captured by an aircraft with lower speed have more overlap than those captured by one with higher speed. One sample sequence is shown in Fig. 16(a), in which the distance between two consecutive images is about 2.5 s of latitude and 0 s of longitude, that is, approximately 253 ft or 77 m. In addition we also tested 10 sequences of images captured over a same spot but at different heights along the direction perpendicular to the ground. Figure 16(b) shows a sequence of images along the vertical direction. A pair-wise consistency rate (CR) of detection results between two adjacent images is defined as

$$\mathrm{CR} = \frac{CN}{CM} \times 100\% \tag{15}$$

where CM is the number of common safe landing-sites in the ground-truth between two adjacent images and where CN is the number of common TP detection results between two adjacent images. The average consistency rate (ACR) is defined as

$$\mathrm{ACR} = \frac{1}{F-1}\sum_{f=1}^{F-1} CR_f \tag{16}$$

where f is the index of images, F is the total number of images in one sequence, and CR_f is the consistency rate between the fth image and the $(f+1)$th image. The average ACR of 20 tested sequences is 84.1%.

IV. DISCUSSION

A vision-based real-time information source is indispensable in the application of seeking safe emergency landing-sites, although there is existing advanced equipment that contains a database that can indicate the locations of plain areas which are suitable for emergency landing, because the system should possess the capability to identify transitory hazards or moving objects on a real-time basis. The relatively low update frequency of the database mentioned above often cannot satisfy this expectation. Ideally, the proposed system can be combined with such a database to work together. For example, if the database can provide the location of a potential safe area and if this area is within the reachable radius of the aircraft, the pilot can first follow the direction prompted from the database. As a result the aircraft is expected to go in a generally correct direction. Then, by using the proposed system, safe landing-sites can be found after ruling out hazardous sites in that generally safe area. This hierarchical methodology is also applicable for exploratory landing on the Moon, Mars, or other planets, in order to eliminate damage to the spacecraft during the landing procedure.

Image quality is directly related to the reliability of the detection results. Given that the camera is at the same height, few details can be seen in a low-resolution image because each pixel of the image covers a large area on the ground. Similarly, for a given camera, fewer details can be seen from the higher elevations, so some areas appear to be smooth in the image, while in reality they are not. The resolution is, of course, a characteristic of the image-capture device, and a high-quality imaging device is always desirable.

In this paper the candidate landing-sites recommended to the pilot are sorted in descending order only according to their areas. We have not taken into account the factor of maneuverability. The assumption that we use in this paper is that a larger and wider landing-site can provide relatively more room for emergency landing, which is, indeed, a positive factor for the emergency landing process. However, in reality, access to the largest candidate landing-site might not be the safest and easiest option for the unpowered aircraft in terms of maneuverability. Therefore, if we can develop a method to estimate the degree of safeness, reliability, and difficulty of landing at each candidate site by evaluating the recommended candidate landing-sites with other factors which are not reflected in the image, i.e., the controllability of the aircraft, wind direction, wind speed, and so on, the automatic detection system can provide a comprehensive index

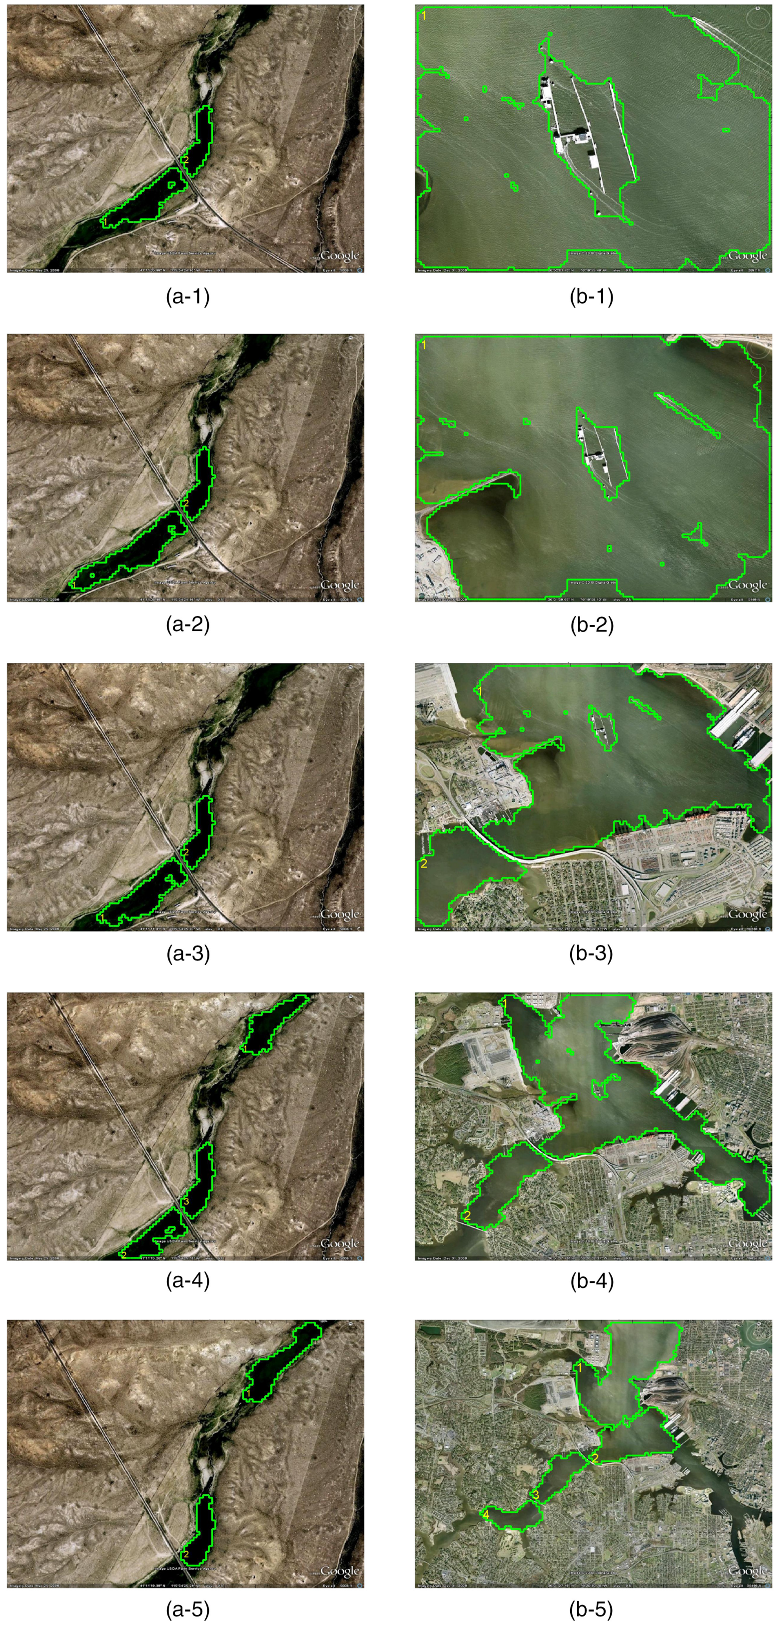

Fig. 16. Detection results of sequential images. (a) Aircraft moves along direction parallel to ground. (b) Images captured over same spot at different elevations.

of the priority of each candidate landing-site, which will be a plus to the proposed system, although that is out of the scope of this paper. Currently, we leave this task to the pilot who will make the final decision by evaluating the recommended candidate landing-sites in an all-inclusive manner.

V. CONCLUSIONS

This paper presents an automatic safe landing-site detection system for robust, reliable, and efficient emergency landing. The proposed system makes up for the limitations of human eyes, assists the pilot to find safe landing-sites, and more importantly, saves time for the pilot to devote to other necessary actions under emergency conditions. The promising results show the feasibility of the vision-based system. In the next step the proposed system will be further developed to better meet practical demands and applications.

ACKNOWLEDGMENT

The authors wish to thank the NASA Aviation Safety Program for the funding which made this work possible. The authors deeply cherish the memory of Dr. Rahman and his outstanding leadership and contribution in this study.

REFERENCES

[1] Shen, Y-F. and Rahman, Z.
An automatic computer-aided detection system for aircraft emergency landing.
Proceedings of the AIAA Infotech at Aerospace Conference, St. Louis, MO, vol. 1, 2011, pp. 779–788.

[2] Shen, Y-F., et al.
Automatic detection of aircraft emergency landing sites.
Proceedings of SPIE, Defense, Security, and Sensing, vol. 8056, Orlando, FL, Apr. 25–29, 2011.

[3] Dickmanns, E. D. and Schell, F-R.
Autonomous landing of airplanes by dynamic machine vision.
Proceedings of the IEEE Workshop on Applications of Computer Vision, Palm Springs, CA, Nov. 30–Dec. 2, 1992, pp. 172–179.

[4] Petruszka, A. and Stentz, A.
Stereo vision automatic landing of VTOL UAVs.
Proceedings of the 23rd Annual Association for Unmanned Vehicle Systems International Symposium and Exhibition (AUVSI '96), Orlando, FL, 1996, pp. 245–263.

[5] Kaminer, I., Pascoal, A., and Kang, W.
Integrated vision/inertial navigation system design using nonlinear filtering.
Proceedings of the American Control Conference, vol. 3, 1999, pp. 1910–1914.

[6] Shakernia, O., et al.
Vision guided landing of an unmanned air vehicle.
Proceedings of the 38th IEEE Conference on Decision and Control, vol. 4, Phoenix, AZ, 1999, pp. 4143–4148.

[7] Shakernia, O., et al.
Landing an unmanned air vehicle: Vision based motion estimation and nonlinear control.
Asian Journal of Control, **1**, 3 (Sept. 1999), 128–145.

[8] Sharp, C. S., Shakernia, O., and Sastry, S. S.
A vision system for landing an unmanned aerial vehicle.
Proceedings of the IEEE International Conference on Robotics and Automation, Seoul, Korea, 2001, pp. 1720–1728.

[9] Werner, S., et al.
A vision-based multi-sensor machine perception system for autonomous aircraft landing approach.
Proceedings of the SPIE, vol. 2736, Orlando, FL, 1996, pp. 54–63.

[10] Yang, Z. F. and Tsai, W. H.
Using parallel line information for vision-based landmark location estimation and an application to automatic helicopter landing.
Robotics and Computer-Integrated Manufacturing, **14**, 4 (1998), 297–306.

[11] Garcia-Padro, P., Sukhatme, G., and Montgomery, J.
Towards vision-based safe landing for an autonomous helicopter.
Robotics and Autonomous Systems, **38**, 1 (Jan. 2002), 19–29.

[12] Fitzgerald, D. L., Walker, R. A., and Campbell, D. A.
A vision based forced landing site selection system for an autonomous UAV.
Proceedings of the International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), Melbourne, Australia, 2005.

[13] Fitzgerald, D. L.
Landing site selection for UAV forced landings using machine vision.
Ph.D. thesis, School of Engineering Systems, Queensland University of Technology, Brisbane, Australia, 2007.

[14] Canny, J.
A computational approach to edge detection.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-8**, 6 (1986), 679–698.

[15] Johnson, A., et al.
Lidar-based hazard avoidance for safe landing on Mars.
Proceedings of the 11th Annual AAS/AIAA Space Flight Mechanics Meeting, Santa Barbara, CA, 2001, pp. 323–337.

[16] Howard, A. and Seraji, H.
A fuzzy rule-based safety index for landing site risk assessment.
Proceedings of the 5th Biannual World Automation Congress, vol. 14, Orlando, FL, June 2002, pp. 579–584.

[17] Howard, A.
A novel information fusion methodology for intelligent terrain analysis.
Proceedings of the IEEE International Conference on Fuzzy Systems, Honolulu, HI, 2002, pp. 1472–1475.

[18] Howard, A. and Seraji, H.
Multi-sensor terrain classification for safe spacecraft landing.
IEEE Transactions on Aerospace and Electronic Systems, **40**, 4 (2004), 1122–1131.

[19] Ploen, S. R., Seraji, H., and Kinney, C. E.
Determination of spacecraft landing footprint for safe planetary landing.
IEEE Transactions on Aerospace and Electronic Systems, **45**, 1 (2009), 3–16.

[20] Paschall, S., et al.
A self contained method for safe & precise lunar landing.
Proceedings of the IEEE Aerospace Conference, Big Sky, MT, 2008, pp. 1–12.

[21] Cappellari, Jr., J. O.
Where on the moon? An Apollo systems engineering problem.
Bell System Technical Journal, **51** (1972), 961–1126.

[22] Johnson Space Center, Solar System Exploration Division, Lunar and Mars Exploration Program Office
Development of a site selection strategy for a lunar outpost—Simulating the selection process.
A Site Selection Strategy for a Lunar Outpost: Science and Operational Parameters Workshop Report, Houston, TX, Aug. 13–14, 1990, pp. 1–80.

[23] Rahman, Z., Jobson, D. J., and Woodell, G. A.
Retinex processing for automatic image enhancement.
Journal of Electronic Imaging, **13**, 1 (2004), 100–110.
[Online], available: http://link.aip.org/link/?JEI/13/100/1.

[24] Jobson, D. J., Rahman, Z., and Woodell, G. A.
A multi-scale Retinex for bridging the gap between color images and the human observation of scenes.
IEEE Transactions on Image Processing (special issue on color processing), **6**, 7 (1997), 965–976.

[25] Williams, S. and Howard, A. M.
Horizon line estimation in glacial environments using multiple visual cues.
Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Shanghai, China, May 9–13, 2011, pp. 5887–5892.

[26] Dusha, D., Boles, W., and Walker, R.
Attitude estimation for a fixed-wing aircraft using horizon detection and optical flow.
Proceedings of *the 9th Biennial Conference of the Australian Pattern Recognition Society on Digital Object Digital Image Computing Techniques and Applications*, Glenelg, Australia, Dec. 3–5, 2007, pp. 485–492.

[27] Ettinger, S., et al.
Towards flight autonomy: Vision-based horizon detection for micro air vehicles.
Proceedings of the 15th Florida Conference on Recent Advances in Robotics, Miami, FL, May 23–24, 2002.

[28] Duda, R. O. and Hart, P. E.
Use of the Hough transformation to detect lines and curves in pictures.
Communications of the ACM, **15**, 1 (Jan. 1972), 11–15.

[29] Rueckert, D., et al.
Nonrigid registration using free-form deformations: Application to breast MR images.
IEEE Transactions on Medical Imaging, **18**, 8 (Aug. 1999), 712–721.

[30] Lee, S., Wolberg, G., and Shin, S. Y.
Scattered data interpolation with multilevel B-splines.
IEEE Transactions on Visualization and Computer Graphics, **3**, 3 (1997), 228–244.

[31] Tou, J. T., Gonzalez, P., and Rafael C.
Pattern Recognition Principles.
Reading, MA: Addison-Wesley, 1974.

[32] Duda, R. O. and Hart, P. E.
Pattern Classification and Scene Analysis.
Hoboken, NJ: Wiley, 1973.

[33] Gonzalez, R. C. and Woods, R. E.
Digital Image Processing (3rd ed.).
Upper Saddle River, NJ: Prentice-Hall, 2008.

[34] Goldstein, H.
Classical Mechanics (2nd ed.).
Reading, MA: Addison-Wesley, 1980.

[35] Landau, L. D. and Lifshitz, E. M.
Mechanics (3rd ed.).
Oxford, UK: Butterworth-Heinmann, 1976.

[36] Lay, D.
Linear Algebra and Its Applications.
Reading, MA: Addison-Wesley, 2000.

Yu-Fei Shen received his B.S. degree in electronic engineering from Fudan University, Shanghai, China, in 2006 and his M.S. and Ph.D. degrees in electrical and computer engineering from Old Dominion University, Norfolk, VA, in 2009 and 2012, respectively.

His primary research focus is on the development of computer-aided detection/diagnosis systems. His research interests include digital signal/image processing and analysis, computer vision, machine learning, and pattern recognition.

Dr. Shen is a member of the AIAA and SPIE.

Zia-ur Rahman (M'89—SM'07) received his B.S. degree in physics from Ripon College, Ripon, WI, in 1984 and his M.S. and Ph.D. degrees in electrical engineering from the University of Virginia, Charlottesville, in 1986 and 1989, respectively.

He joined the Department of Electrical and Computer Engineering at Old Dominion University as an associate professor in 2006. Before that he was a research associate professor with the Department of Applied Science at the College of William & Mary.

His primary area of specialization was digital signal processing, with an emphasis on digital image processing. His research focused on the application of digital image enhancement algorithms to increase situational awareness of a pilot in landing an aircraft in poor visibility conditions. He had several grants from NASA for this research effort. He was a member of SPIE and OSA. He passed away in 2012.

Dean J. Krusienski (M'01) received B.S., M.S., and Ph.D. degrees in electrical engineering from The Pennsylvania State University, University Park, PA, in 1999, 2001, and 2004, respectively. He completed his postdoctoral research at the New York State Department of Health's Wadsworth Center Brain-Computer Interface (BCI) Laboratory in Albany, NY.

He is an associate professor in the Electrical and Computer Engineering Department at Old Dominion University. His primary research focus is on the application of advanced signal processing and pattern recognition techniques to brain-computer interfaces, which allow individuals with severe neuromuscular disabilities to communicate and interact with their environments using their brainwaves. His research interests include decoding and translation of neural signals, digital signal and image processing, machine learning, evolutionary algorithms, artificial neural networks, and biomedical and musical applications.

Jiang Li (M'05) received his B.S. degree in electrical engineering from Shanghai Jiaotong University, China, in 1992, his M.S. degree in automation from Tsinghua University, China, in 2000, and his Ph.D. degree in electrical engineering from the University of Texas at Arlington, TX, in 2004.

Dr. Li worked as a postdoctoral fellow in the Department of Radiology, National Institutes of Health, from 2004 to 2006. He joined Old Dominion University as an assistant professor in 2007. He is also affiliated with Virginia Modeling, Analysis, and Simulation Center. His research interests include machine learning, computer-aided medical diagnosis systems, medical signal/image processing, and modeling and simulation.