ARMY RESEARCH LABORATORY

Autonomous Identification of Locations for Unmanned Ground Vehicle (UGV) Cover and Concealment With the Use of Video Images

by William F. Oberle

ARL-TR-3517   **May 2005**

NOTICES

Disclaimers

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5066

Autonomous Identification of Locations for Unmanned Ground Vehicle (UGV) Cover and Concealment With the Use of Video Images

William F. Oberle
Weapons and Materials Research Directorate, ARL

		Form Approved
REPORT DOCUMENTATION PAGE		*Form Approved* *OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)* May 2005	2. REPORT TYPE Final	3. DATES COVERED (From - To) October 2004 to April 2005
4. TITLE AND SUBTITLE Autonomous Identification of Locations for Unmanned Ground Vehicle (UGV) Cover and Concealment With the Use of Video Images		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) William F. Oberle (ARL)		5d. PROJECT NUMBER 622618AH03
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory Weapons and Materials Research Directorate Aberdeen Proving Ground, MD 21005-5066		8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-3517
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Autonomously identifying appropriate locations that provide cover and concealment for an unmanned ground vehicle (UGV) is a difficult task. This report expands an earlier algorithm for accomplishing this task, based on video images to address the difficulty of handling thin or narrow objects in the UGV environment. The vertical line approach used in the earlier algorithm is expanded to include an analysis of local regions to assess color homogeneity and size criteria. Careful selection of seed pixels attempts to ensure that they represent a potential thin or narrow object and not the background. The expanded algorithm provided very good results, based on the video sequences analyzed.

15. SUBJECT TERMS
Canny; edge detection; Hough transform; unmanned ground vehicle; video

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON William F. Oberle
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified	SAR	35	19b. TELEPHONE NUMBER *(Include area code)* 410-278-4362

Contents

List of Figures

List of Tables

Acknowledgments

INTENTIONALLY LEFT BLANK

1. Introduction

Integral to the successful completion of many tactical Army missions is the ability to use natural and man-made features within the battlefield environment for cover and concealment from the enemy. Unfortunately, for unmanned ground vehicles (UGVs) operating in an autonomous[1] or semi-autonomous[2] mobility mode, the identification and evaluation of potential locations for cover and concealment is a difficult behavior to model and program. The exact degree of the difficulty is influenced by the sophistication and resolution of the UGV's on-board sensors.

In an earlier work (Oberle, 2004), an algorithm for the autonomous identification of locations was presented, which might offer acceptable cover and concealment with the use of video images. The algorithm was predicated on the assumption that most acceptable locations would be man-made structures characterized by vertical lines (e.g., buildings). In that algorithm, a set of candidate vertical lines was determined via the Canny edge detector (Canny, 1986; Faugeras, 1993; Trucci and Verri, 1998) and the Hough transform (Klaus and Horn, 1986; Faugeras, 1993; Jahne, 1997; Trucci and Verri, 1998). We reduced this set by imposing a continuous length criterion, with the final choice being the vertical line terminating lowest in the image. The choice of this line is used since that location should be closest to the camera. This line is referred to as the dominant vertical line. Additional details are available in Oberle (2004). Results of applying the algorithm to a sequence of video images are provided in figure 1. As can be observed in figure 1, the algorithm is reasonably effective at identifying appropriate locations. However, as noted in the earlier work, the algorithm needs to be improved to eliminate locations associated with the vertical edge of a thin or narrow object. This is illustrated in frame[3] 5 where the dominant vertical line is associated with the polyvinyl chloride (PVC) pipe protruding from the ground. The objective of this report is to detail efforts to expand this earlier algorithm developed in Oberle (2004) to address the problem of thin or narrow objects.

The remainder of this report is organized as follows. A description of the methodology selected to address the problem is presented in section 2. Application of the expanded algorithm to several video sequences and an assessment of the results are provided in section 3. Conclusions and limitations of the algorithm are discussed in section 4.

---

[1]No human direction.

[2]Intermittent human direction.

[3]In figures 1 and 9, the frames are numbered from left to right, starting with the top row. For example, the frames in the second row are frames 4, 5, and 6 from left to right.
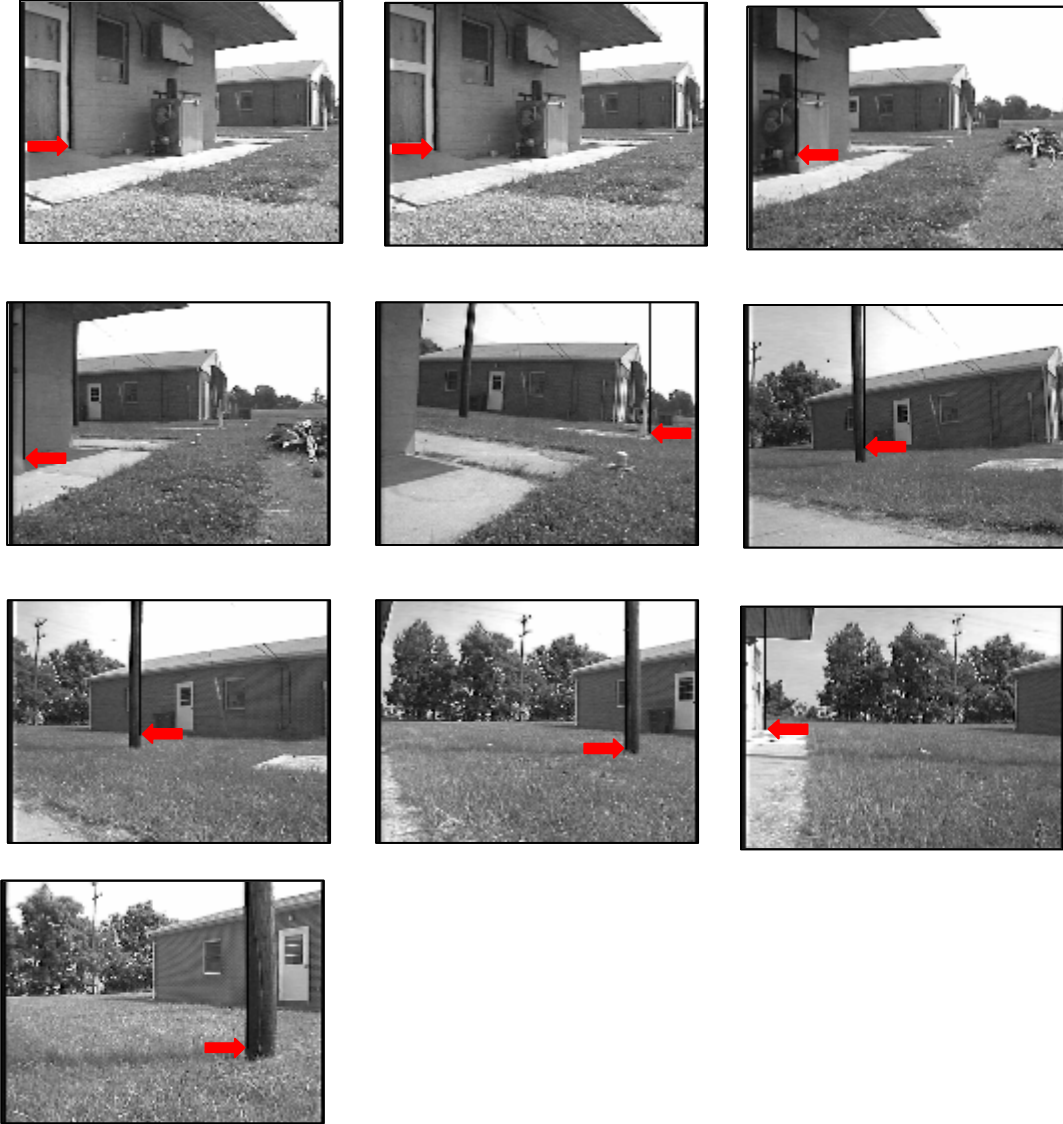
Figure 1. Grey scale of video sequence 1 (VS1) with dominant vertical line inserted. (Bottom of dominant vertical line indicated by red arrow [Oberle, 2004].)

2. Methodology

In the detection scheme used in the original algorithm, edges are defined by pixel locations at or around which the grey scale intensity gradient is large. This implies that in the red, green, and blue (RGB) color space, the image colors to the left and right of a vertical edge will be different, as illustrated in figure 2. Although figure 2 shows the regions to the left and right of the vertical line as homogeneous in color, this is not necessarily the case. For example, in figure 1, frame 10, the dominant vertical line is the left edge of the telephone pole. The region to the right of the edge is homogeneous in color (i.e., the telephone pole); however, to the left of the edge are three

distinct regions: the grass at the bottom, the brick building wall in the middle, and the sky at the top. Thus, a vertical line can represent the boundary between a number (possibly more than two) of different color regions. However, in this work, only the regions immediately adjacent to the left and right of the lowest point on the vertical line are considered, and therefore, any reference to region(s) will be referring to one or both of these two regions. The reader should also note that the regions may be small, e.g., when they concern thin or narrow objects.

Region 1   **Region 2**

Vertical Edge

Figure 2. Schematic of vertical edge and
different color regions to left
and right.

Several assumptions (in addition to those of the earlier work) need to be made in order to expand the vertical line approach of the earlier work to address the presence of thin or narrow objects. These assumptions are now listed:

1. Region size in pixels can be used to estimate the physical size of the region, or conversely, given a desired physical size (i.e., area sufficiently large enough to conceal the UGV), the region size in pixels can be determined.

2. A homogeneously colored region corresponds to a contiguous physical surface suitable for cover and concealment. Unfortunately, grass and sky regions are almost always the most homogeneously colored regions in a video image. Since these regions do not generally offer acceptable cover and concealment, they need to be eliminated from consideration. The following approaches will be used in an attempt to minimize/eliminate consideration of such regions.

a. In the RGB color space, if the value of the G component is larger than the R and B components, the pixel is most likely grass. Note that this assumption eliminates other vegetation such as tree leaves that could indicate tree lines.

b. Regions terminating lowest in the image will almost always avoid sky regions.

3. Because of noise in the data acquisition, differences in illumination, etc., homogeneously colored regions do not consist of pixels all with exactly the same RGB values but consist of pixels satisfying specified similarity criteria.

Under these assumptions, the original algorithm is expanded by the following three tasks in order to meet the objectives of this work:

1. Determine the region dimensions in pixels, given desired physical dimensions.

2. Apply similarity criteria for the vertical lines identified in the image, to determine if the left, right, or both associated regions are sufficiently homogeneous in color to represent a continuous surface.

3. Select from all acceptable regions a single region for cover and concealment.

2.1   Region Dimensions

Physical dimensions cannot be determined from a single camera image, as illustrated in figure 3. Points A and B in the image plane are the image of any point along rays r_1 and r_2, respectively. Thus, line segments A_1B_1 and A_2B_2 both map to the same line segment, AB, in the image. In order to convert between physical and pixel measurements, the physical distance to the object must be known. Fortunately, most UGVs are equipped with some method for estimating distance, e.g., ladar[4] or stereo cameras. In fact, if the distance is measured directly as with ladar and the resolution is sufficiently high, then range images could be used in place of video images in the algorithm being developed with appropriate modification. Vertical depth discontinuities would replace vertical lines, and homogeneously colored regions would be replaced by regions of equal depth. For the purposes of this work, the distance to the object is considered user input as a surrogate to being supplied by an on-board UGV sensor.

Assuming that the distance to an object of interest is known, a minimal region size can be estimated with the knowledge of the camera vertical and horizontal field of view (FOV) and the image resolution in pixels. Let H_F be the camera horizontal FOV (radians), V_F the camera vertical FOV (radians), H_P the number of pixels per scan line, and V_P the number of scan lines, i.e., the image resolution in H_P pixels wide by V_P pixels high. Now consider figure 4; this shows the geometry of the camera and an object of interest represented by the line segment AB. This view can be considered either a top view, so that α is the subtending angle in the horizontal plane

---

[4]An acronym for laser detection and ranging, ladar uses light for detection of speed, altitude, direction and range; it is often called laser radar. See the photonics dictionary – web site: http://www.photonics.com/dictionay/.

of the object or a side view with α the subtending angle in the vertical plane of the object. If d is the distance to the object, the length of the line segment AB can be approximated by the length of the arc AB (in fact, this is always an overestimation, which in this case is acceptable since a minimal region size is the objective) and is given by the standard geometric equation for the arc length of a sector of a circle,

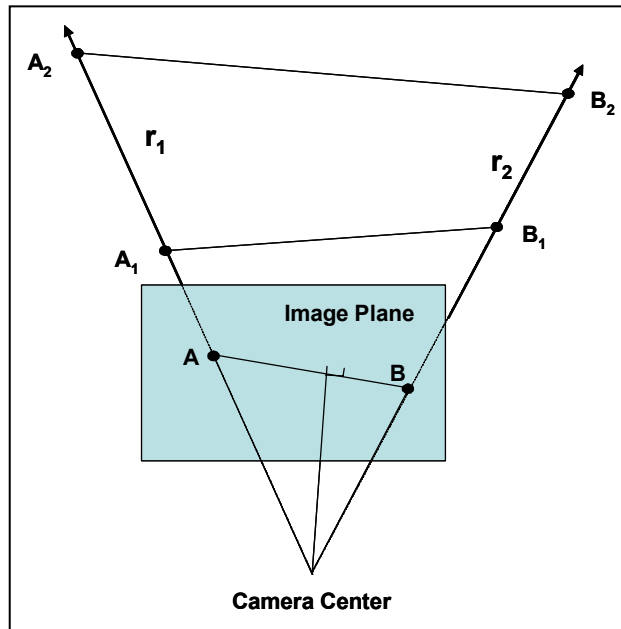$$\overline{AB} \approx d\,\alpha\,. \tag{1}$$

Figure 3. Schematic of camera image which illustrates
that physical dimensions cannot be
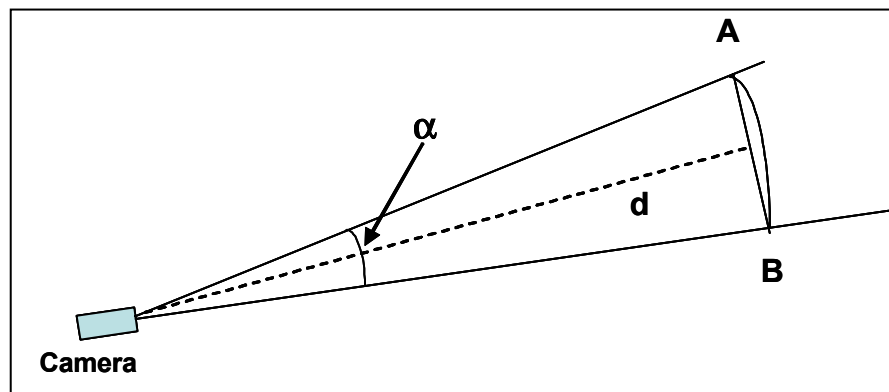determined from a single video image.

Figure 4. Geometry of camera and object of interest (line AB).

Now the length of the line segment AB is the desired height (width) of the region and is a known value. Thus, the angle α can be estimated. Next, consider the ratio H_P/H_F (V_P/V_F), which gives the camera's pixels per radian in the horizontal (vertical) direction. Therefore, the number of

image pixels in the horizontal (vertical) direction corresponding to a desired physical size is given by

$$Pixels_{Horizontal} = \frac{\overline{AB}}{d} * \frac{H_P}{H_F} \qquad and \qquad Pixels_{Vertical} = \frac{\overline{AB}}{d} * \frac{V_P}{V_F}. \qquad (2)$$

When implemented in the expanded algorithm, the size of the region in pixels will be a square region of length equal to the maximum of the horizontal and vertical pixel sizes calculated with equation 2.

2.2   Similarity Criteria

The second task to be completed by the expanded algorithm involves the selection and application of similarity criteria to determine the suitability of regions associated with identified vertical lines. Although this task appears similar to image segmentation, the task is simpler than segmentation. Segmentation attempts to divide the image into regions with like characteristics (e.g., color). In this problem, the regions of interest are well defined in location (adjacent to vertical lines) and size.

This suggests a straightforward approach to completing the task. Specifically, one should decide if a region of interest consists of at least a user-defined threshold percentage of similar pixels. The similarity criteria are based on color. Two pixels will be considered similar if the absolute difference of their values in each of the three color components in RGB color space is less than a user-specified value. As stated earlier, to eliminate grass, all pixels in which the value of the G component of the RGB signal is greater than the value of the R and B components will automatically be classified as dissimilar. Note that other criteria can be used, especially if additional color information about the scene is available.

Now the question of how to determine the set of similar pixels must be addressed. First, note that similar as defined is not transitive, i.e., if pixel A is similar to pixel B, and pixel B is similar to pixel C, then pixel A is not necessarily similar to pixel C. For example, if pixels are similar if their value is within 10 units, and if the red component of pixel A is 142, of pixel B 150, and of pixel C 159, then in the red component pixel A is similar to pixel B, and pixel B is similar to pixel C, but pixel A is not similar to pixel C. Thus, determining similarity is relative to a specific pixel or values for the RGB component values termed a "seed". Different approaches to selecting a seed pixel or RGB values have been used in various approaches to image segmentation. These approaches include the use of a randomly selected pixel in the region, the use of the region's centroid pixel, or the use of the average RGB values of all the pixels in the region. However, all these approaches leave open the possibility that for a thin or narrow object, the seed will be more representative of the background and thus may not eliminate a thin or narrow object, as shown in figure 5. In the figure, the seed for the left region is randomly chosen while in the right region, the seed is the centroid. To minimize the possibility of this situation, the seed pixel should be selected at the vertical edge. However, because of acquisition noise and

edge "bleeding," no single pixel on or near the edge could be representative of the region. As a compromise, three seed pixels will be chosen for each region. Similar pixels will be determined relative to each seed pixel. This will produce three sets of similar pixels. If any one of the three sets is large enough to satisfy the user-supplied percentage threshold, the region will be considered to contain a sufficient number of similar pixels. The three seed pixels are chosen close to the edge to increase the probability that they will be located within the boundaries of a thin or narrow object. Specifically, the pixels L1, L2, and L3 are chosen for the left region, and pixels R1, R2, and R3 are chosen for the right region, as shown in figure 6.
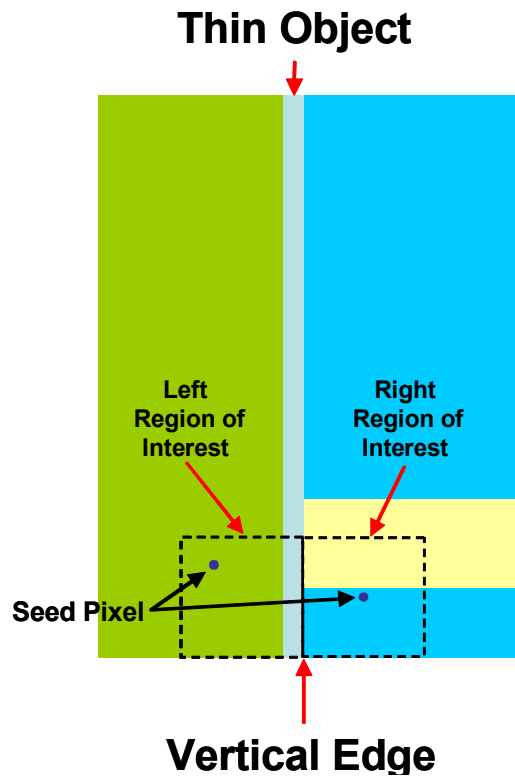

Figure 5. Example illustrating that the seed pixel should be selected close to the vertical edge to ensure that the similarity test is representative of the thin object and not the background.

With these similarity criteria, it can be decided whether all the regions associated with all the identified vertical lines in the image can be classified as possible cover and concealment locations. For each vertical line, there are four possible outcomes for the regions, as shown in table 1.
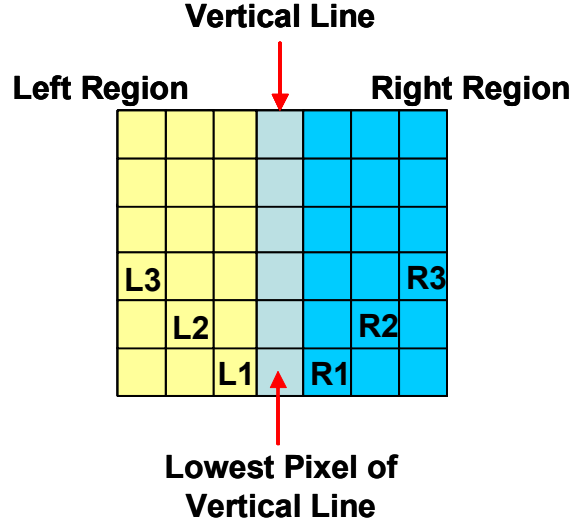
Figure 6. Location of seed pixels in left and right
regions. (Each square represents a
single pixel.)

Table 1. Possible outcomes for vertical line regions.

Outcome	Left Region Acceptable	Right Region Acceptable
I	Yes	Yes
II	Yes	No
III	No	Yes
IV	No	No

Outcomes II, III, and IV are clear. However, outcome I requires additional analysis. If the left
and right regions are statistically similar in the sense that both regions are samples from the same
underlying distribution, then the regions should be merged. If the regions are not statistically
similar, the region with the highest percentage of similar pixels is considered the most
appropriate region. To determine if the two regions are statistically similar, a chi-squared
statistic is used. Each component of the RGB signal is tested separately in the determination of
whether the regions are statistically similar via the following procedure.

1. For both regions and each component of the RGB signal, a histogram is constructed.
    This produces three pairs (R, G, and B components) of binned data.

2. For each of the three pairs of binned data, the chi-squared statistic is computed by the
    formula

$$\chi^2 = \sum_i \frac{\left(\sqrt{S/R}\,R_i - \sqrt{R/S}\,S_i\right)^2}{R_i + S_i} , \tag{3}$$

in which

$$R \equiv \sum_i R_i \qquad S \equiv \sum_i S_i$$

8

R_i is the number of events in bin i for the first data set, and S_i is the number of events in the same bin for the second data set (Press, Teukolsky, Vetterling, and Flannery, 2002).

3. For each of the chi-squared results in step 2, the incomplete gamma function is used to determine the probability that the left and right regions are statistically similar in each color component. A user-supplied threshold value is used to decide if the regions are statistically similar in each color component. The smaller the probability, the less likely that the color component is statistically similar (Press et al., 2002).

4. If all three color components are determined to be statistically similar in step 3, the left and right regions are merged. Otherwise, the regions are considered dissimilar.

2.3 Selection of Single Region

The third task for the expanded algorithm is the selection of a single region for cover and concealment from regions determined to be acceptable. The criterion for completing this task is the same as in the original algorithm for selecting the most appropriate vertical line, namely, the region situated lowest in the image. For the reasons discussed in the earlier report, this region should be closest to the vehicle. In addition, as stated in the assumptions, this choice should eliminate regions dominated by the sky.

---

3. Application of Expanded Algorithm to Video Sequences

The region selection criteria are based on two factors: a sufficient number of similar pixels and the lowest acceptable region in the image. Thus, instead of analyzing the regions for each vertical line, we can reduce computation time by first ordering the set of vertical lines in descending order, i.e., lowest line in the image first, and then stopping at the first acceptable region. This is the approach selected for the expanded algorithm. A flow diagram for the region-selecting algorithm is provided in figure 7.

The results of applying the expanded algorithm to video sequence 1 are shown in figure 8 (a) through (j). The image to the left of figure 8 (a) through (j) is the original color image from the video sequence (Oberle, 2004); the corresponding grey scale conversion of this image together with the dominant vertical line is shown in figure 1. The image on the right of figure 8 (a) through (j) provides the results via the expanded algorithm. A green line is used to indicate the vertical line whose region(s) meet the selection criteria discussed in the previous section. If a single green square is in the image, then only one of the regions (outcome II or III, table 1) associated with the vertical line met the criteria. Green and red squares indicate that both regions (outcome I, table 1) met the criteria, but the regions were not statistically similar. The green square indicates the region with the higher percentage of similar pixels. If the regions were

9

statistically similar, there would be two green squares. In the original algorithm, the cover and concealment location (the waypoint for the UGV) is the bottom end point of the vertical line. For the expanded algorithm, the way point for the UGV is the centroid of the region. In addition, the expanded algorithm should be run continuously in order to refine the UGV way points. The rationale behind this choice is discussed next.
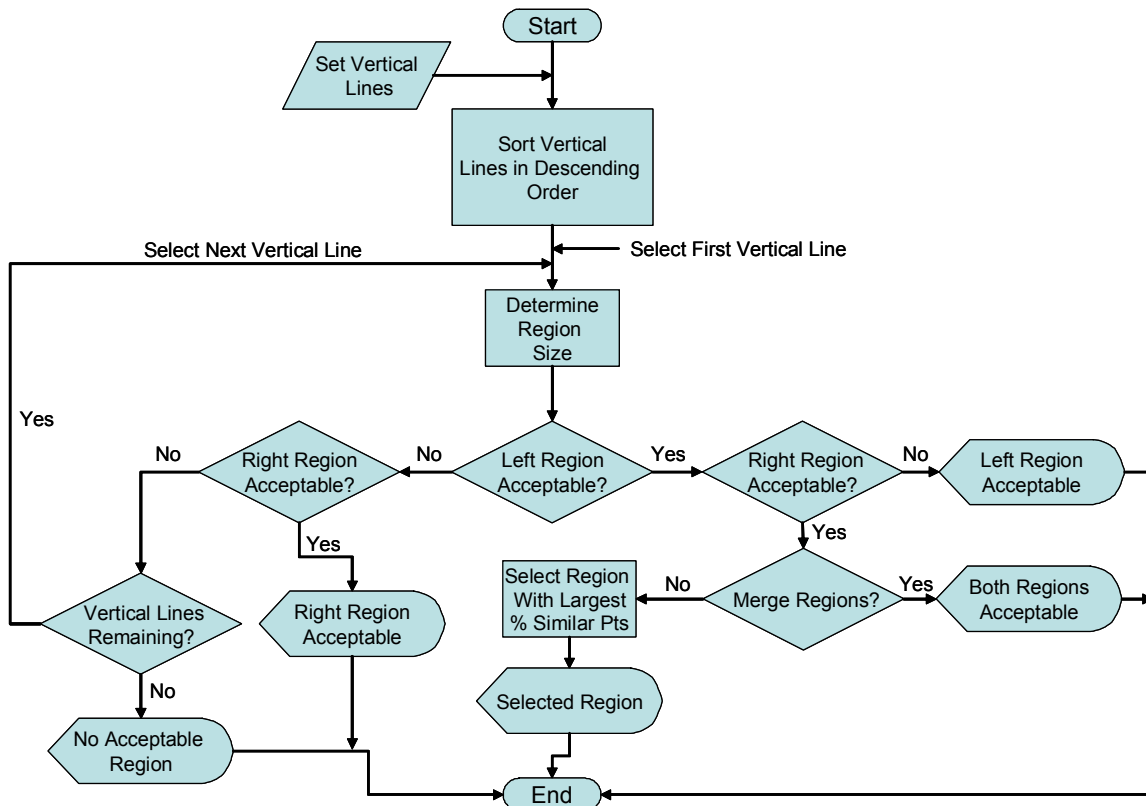
Figure 7. Flow diagram for region-selecting algorithm.

The prime objective of this work is to develop an algorithm that eliminates the identification of locations associated with thin or narrow objects. Frame 5 of figure 1 was used earlier to illustrate the problem. As can be seen in figure 8 (e), the expanded algorithm no longer identifies the vertical line associated with the PVC pipe but instead selects a region associated with the telephone pole. Note that the vertical line selected by the expanded algorithm terminates higher in the image than the vertical line associated with the PVC pipe. Thus, consistent with logic of the algorithm, the regions associated with the PCV pipe's vertical line are tested and rejected as acceptable. The region to the left contained seed pixels associated with the white of PCV pipe, a color not representative of the region, while the region to the right is mainly grass.

10

(a) VS1 – frame 1

(b) VS1 – frame 2

(c) VS1 – frame 3

11

(d) VS1 – frame 4

(e) VS1 – frame 5

(f) VS1 – frame 6

(g) VS1 – frame 7

(h) VS1 – frame 8

(i) VS1 – frame 9

13

(j) VS1 – frame 10

Figure 8. Results of region selection for VS1. (Images on left are the original color images [Oberle, 2004]; those on the right are the results for the region selection algorithm.)

Figures 8 (h) and (j) illustrate that a second objective of the expanded algorithm is also achieved, namely, grass regions are eliminated. A comparison of figure 1, frames 8 and 10 and figure 8 (h) and (j), shows that the vertical line selected is in the same general location, i.e., along the side of the telephone pole. However, in the original algorithm, figure 1, the vertical line terminated at the base of the pole with regions to the right and left being grass. The expanded algorithm eliminated these lines and selected lines slightly higher in the image where the regions would not include the grass. These figures also illustrate the rationale for the selection of the centroid of the region and not the end point of the vertical line as the UGV way point location. If the UGV headed toward the end point of the vertical line, the location would be the telephone pole, an object that is too narrow for the specified region size. Heading to the centroid of the region sends the UGV toward the better cover offered by the brick wall. Intuitively, it seems strange to head toward the brick wall when the vertical line is associated with the telephone pole. A human would recognize that the wall continues behind the telephone pole, and heading toward the "middle" of the wall would hinder the UGV's lines of sight. However, from the algorithm's perspective, the vertical line is the boundary for both the foreground and background object. What the algorithm is determining is whether a region to the left and/or right of the boundary is sufficiently large in size and homogeneous in color. The expanded algorithm is concerned with the regions, not the vertical line. If robust depth information is available, that information could be used to identify the foreground object. Now consider what would happen if the analysis were restricted to only those regions associated with foreground objects. If the region(s) did not meet the acceptance criteria, the location would be correctly rejected as a thin or narrow object and background objects that could simply be interior portions of a larger continuous object would not be considered (as in figure 8 (h) and (j)). This would cause the algorithm to consider vertical lines higher in the image. In the frames considered, this could be the left edge of the building, which is the most desirable location. Fortunately, if the expanded algorithm is run continuously, the same effect can be achieved. Consider figure 8 (j); as the UGV heads toward the centroid of

14

the indicated region, it will eventually pass the telephone pole. At that point, the dominant line will no longer be associated with the telephone pole but will most like be the left edge of the building. Furthermore, as the UGV gets closer to the wall, the distance will decrease, causing the region size to decrease. Assuming that the left edge of the building remains the dominant vertical line, this will cause the centroid of the region and thus, the UGV way point, to move toward the edge of the building.

In summary for video sequence 1, if the UGV is directed toward the centroid of the selected region, adequate cover and concealment would be provided by the location and the difficulty with thin or narrow objects of the original algorithm is avoided. In addition, both grass and sky regions are eliminated.

In the original report (Oberle, 2004), a second video sequence (VS2) was also used. In VS2, the distances to possible cover and concealment locations are considerably greater than in VS1. Results for the original algorithm with VS2 are shown in figure 9 and for the expanded algorithm in figure 10.
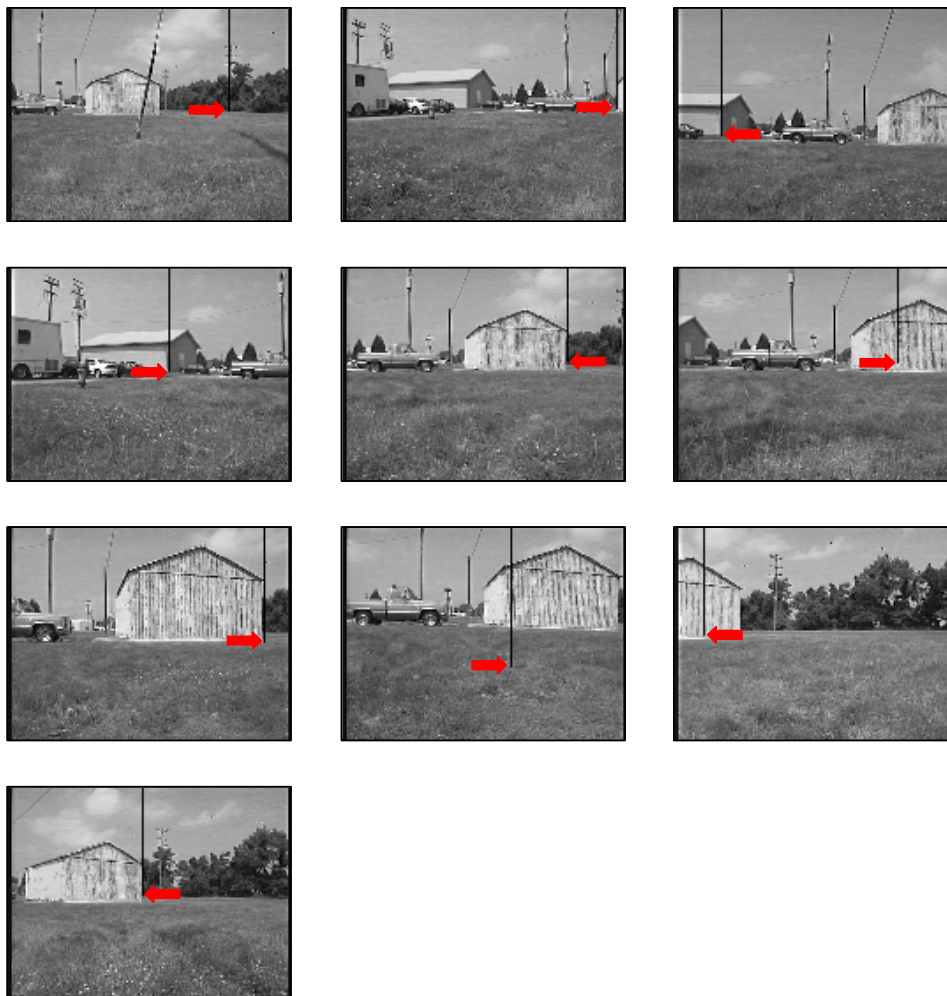

Figure 9. Grey scale of video sequence 2 (VS2) with dominant vertical line inserted.
(Bottom of dominant vertical line indicated by red arrow [Oberle, 2004].)
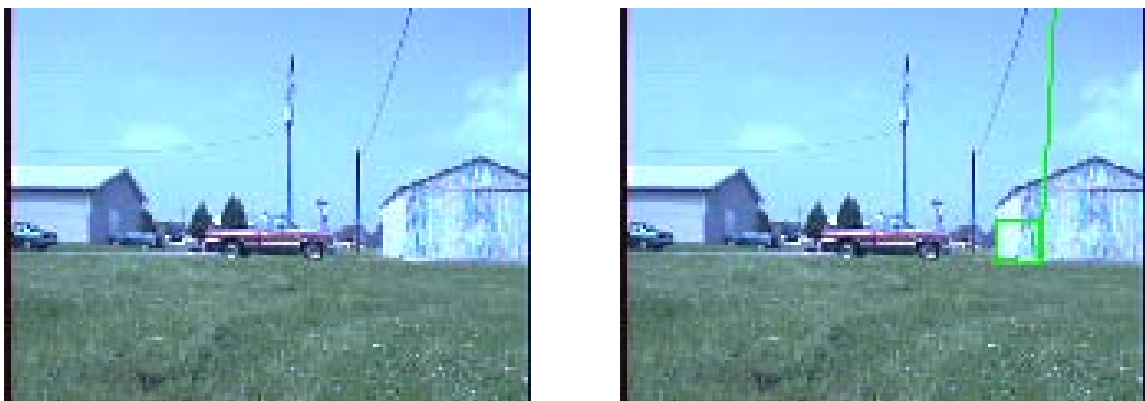
As can be seen in figure 10, the expanded algorithm performed well in identifying acceptable regions for cover and concealment for this video sequence.

(a) VS2 – frame 1

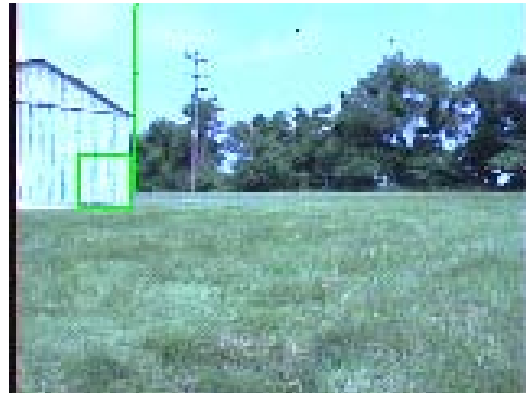**(b) VS2 – frame 2**

(c) VS2 – frame 3

(d) VS2 − frame 4

(e) VS2 − frame 5

(f) VS2 − frame 6

17

(g) VS2 – frame 7

(h) VS2 – frame 8

(i) VS2 – frame 9

18

(j) VS2 – frame 10

Figure 10. Results of region selection for VS2. (Images on left are the original color images [Oberle, 2004]; those on the right are the results for the region selection algorithm.)

4. Conclusions and Limitations

For tactical military missions, the use of cover and concealment is often central to the success of the mission and survivability of the forces. Identification of acceptable locations for cover and concealment is a relatively simple task for a human. However, automating the process so that a UGV can successfully perform this task is a significant challenge that is undergoing active research by the robotic community. The objective of this report is to expand an earlier algorithm (Oberle, 2004) for finding cover and concealment locations, based on video images to address the difficulty of handling thin or narrow objects in the UGV environment. To assess color homogeneity and size criteria, the vertical line approach used in the earlier algorithm is expanded to include an analysis of regions adjacent to the vertical lines. By careful selection of the seed pixels to ensure that they represent a potential thin or narrow object and not the background, the expanded algorithm provided very good results, based on the video sequences analyzed. Results for two such video sequences are provided in the report.

Although the expanded algorithm produced very good results, several limitations and an area for improvement are noted. First, the algorithm requires an estimate of the distance from the UGV to the location or object being analyzed. The distance estimate is required to size the image regions used in the analysis. If the image regions are too small, the physical location may also be too small to provide adequate cover and concealment; if too large, acceptable physical locations may be rejected because the image region may include extraneous pixels. Since a single video image cannot provide this estimate, the algorithm must rely on other UGV sensor input. A second limitation is that the algorithm is not applicable to general terrain. The algorithm is predicated on object boundaries being vertical lines. This limits the algorithm to considering

mostly man-made objects in the analysis. General terrain would require the use of pattern recognition techniques such as neural networks or pattern matching. Finally, the expanded algorithm does not make any real value judgment between regions identified as meeting the acceptance criteria. It simply selects the region lowest in the image. Although this has proved to be an effective solution to identifying a location for cover and concealment, it may not always identify the best location. Extensions of the algorithm to expand its capabilities in this area are recommended.

5. References

Canny, J. A Computational Approach to Edge Detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **November 1986**, *8* (6).

Faugeras, O. *Three-Dimensional Computer Vision: A Geometric Viewpoint*, The MIT Press: Cambridge, Massachusetts, 1993.

Jahne, B. *Digital Image Processing Concepts, Algorithms, and Scientific Applications*, Springer-Verlag: Berlin, Heidelberg, New York, 1997.

Klaus, B.; Horn, P. *Robot Vision*, The MIT Electrical and Engineering and Computer Science Series, The MIT Press, McGraw-Hill Book Company: New York, 1986.

Oberle, W. F. *Autonomous Identification of Vertical Lines in Video Images to Assist in UGV Cover and Concealment*; ARL-TR-3365; U.S. Army Research Laboratory: Aberdeen Proving Ground, MD, November 2004.

Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; Flannery, B. P. *Numerical Recipes in C++, The Art of Scientific Computing,* Second Edition, Cambridge University Press: New York, 2002.

Trucco, E.; Verri, A. *Introductory Techniques for 3-D Computer Vision*, Prentice Hall: Upper Saddle River, NJ, 1998.

22