

VBASR: The Vision System

Vision-Based Autonomous Security Robot

Kevin Farney, Joel Schipper

Bradley University
1501 West Bradley Avenue
Peoria, IL 61625
309-677-2260

kevin.farney@gmail.com, jschipper@bradley.edu

ABSTRACT

The goal of this project is to develop a computer vision system that enables a robot to navigate the hallways of Bradley University's engineering building using a generic webcam as the only sensor. OpenCV2.0 programmed in C++ is the primary programming tool used to develop the vision system.

Three algorithms were developed to identify the center of the hallway and guide the robot in the correct direction. The first two algorithms use a generic filter (normal, median, or Gaussian) followed by edge detection and then corner detection on the edge-detected image. The first algorithm identifies the strongest vertical lines on an image. Averaging the horizontal coordinates of the vertical lines indicates the location of the center of the hallway relative to the robot. The second algorithm utilizes the trapezoidal shape of the hallway formed where the floor meets the walls, as seen from the perspective of the robot. The y-coordinates of the trapezoid's lower corners are then compared to estimate robot orientation with respect to the walls. The third algorithm uses color to segment the floor from the rest of the features in the image (walls, ceiling, and obstacles). Once again, the trapezoidal shape appears and the center of the hallway is determined based on the location of the highest y-valued pixels identified as floor pixels.

Test data indicates that none of these algorithms are singularly sufficient; however, combining their results they can identify the direction a robot must turn to remain in the center of the hallway with 96% accuracy. Furthermore, leveraging the results of multiple algorithms produces more robust navigation, where one algorithm covers over the shortcomings of another. The vision system architecture is designed to execute algorithms in parallel. Such a structure enables the addition and removal of algorithms without adversely affecting the system as a whole. Further algorithms may be developed and easily added to improve navigation. Additionally, the system may intelligently ignore results from algorithms that are recognized as inappropriate for certain situations.

KEYWORDS

Machine Vision, Image Processing, Mobile Robot Navigation, Autonomous Vehicles

1. INTRODUCTION

VBASR (Vision-Based Autonomous Security Robot) is designed to patrol the second floor hallway of the engineering building during the after-hours of a regular school day. Essentially, VBASR is a mobile, intelligent security camera able to locate and navigate to specific rooms and photograph any intruders it encounters. VBASR: The Vision System is a proof-of-concept for the design and development of the machine vision system necessary to implement VBASR. The primary goal for this project is a robust, vision-based navigation system. Similar work using vision and security systems was performed by Sage, Kingsly and Young, particularly on detecting motion [1].

A great resource for familiarization of the computer vision terms found throughout this paper can be found in the text, [Computer Vision](#) [2].



Figure 1. iRobot Create and Accessories [3]

VBASR is primarily a machine vision project. Therefore a chassis was selected that requires little modification. Figure 1 shows the iRobot Create chassis selected as the robot platform. A simple webcam mounted in the cargo bay and an onboard computer are the only additional hardware necessary for VBASR. Microsoft Robotics Developers Studio (MRDS) is used to control the iRobot Create and OpenCV2.0 computer vision libraries were utilized to implement the vision algorithms. VBASR is programmed in C++. Since VBASR is currently a proof-of-concept, no mountable onboard computer was selected. Ideally, a single-board computer running MRDS in a Windows environment should be implemented. For development, a generic HP laptop running the Windows 7 OS was used.

1.1 The Problem Description

Given an image of the hallway, such as the one shown in Figure 2, how would a robot choose which direction to turn to travel down the center of the hallway? Humans can solve this problem using vision intuitively without a concentrated effort. VBASR must make similar decisions using similar information.



Figure 2. Example image of hallway

One obvious difficulty with the image in Figure 2 is the lack of depth perception. VBASR is designed to use a single webcam, making depth perception on a single image extremely difficult. Other sensors could be used to measure distance from the walls or ends of the hallway, but is it possible to navigate only using machine vision? Three different algorithms were developed to navigate in such a manner. The compilation of those three algorithms constitutes VBASR: The Vision System.

Each algorithm independently determines the direction VBASR should navigate, selecting from one of seven general directions: Hard Left, Left, Slight Left, Straight, Slight Right, Right and Hard Right. Each direction is given an integer representation: 0=Hard Left through 6=Hard Right. A resolver function then considers information from each individual algorithm and determines the final direction of travel.

A library of three hundred hallway images was used to test the accuracy of the VBASR's vision system. These three hundred images represent a cleaned data set where markedly similar images were removed to prevent redundancy. Skewing the data with such redundancy could lead to an overly optimistic or pessimistic evaluation of the system. Given that VBASR is successfully navigating its environment, VBASR should rarely encounter situations where it is in close proximity to a wall. Consequently, images taken within one foot of the wall were also discarded. The library was preprocessed to assign an ideal direction to each image, which was encoded in the filename. Every image was examined and the desired direction was assigned by human observation. The data was normalized by calculating the success rate for each direction and then averaging the percentages. Therefore the number of pictures in each category does not need to be equal.

Each algorithm in the vision system was tested on all three hundred images in the library. The results calculated by the algorithms were then compared to the desired navigation results to evaluate the success rate of each individual algorithm. The algorithm's final decision was considered successful if it was within one step of the correct direction. As shown in Figure 3, the

ideal direction chosen is Slight Right. However, if VBASR decides to go Straight or Right in this case, it will still be travelling in a proper direction.



Figure 3. Slight Right example

2. LINES ALGORITHM

2.1 Theory

The first idea attempted was to find the strongest vertical lines of the image. Main vertical lines in a hallway include: windows, doors, pictures, etc. All of these are found on the walls. Thus if the wall locations can be determined on either side, then the average of the wall locations should be the approximate center of the hallway, as shown in Figure 4. (Note that the lines shown on Figure 4 were added with an image editing program and are not mathematically accurate. The lines were added solely to demonstrate the theory of the lines algorithm.) Red lines represent 'strong' vertical lines and the maroon line represents the average of the x-values of the red lines.



Figure 4. Lines algorithm theory

2.2 Preprocessing

To find the strongest vertical lines of the image, a line (edge) detection algorithm is required. A generic filter must be used on the image as a prerequisite for the line detection algorithm. Without the use of a blurring filter shown in Figure 5, the edge detection algorithm will detect many artifacts that are undesirable, as seen in Figure 6. Example blurring filters are the normal blur, median blur or Gaussian blur. Each of these filters have similar effects to the one shown in Figure 5. The differences between the filters are simply the mathematical methods utilized to achieve the desired end result: a blurred image. For example, the normal blur

shown in Figure 5 uses a normalized box filter which simply sums the pixels over the given neighborhood. The best overall VBASR algorithm utilizes a median blur which simply returns the median of the neighborhood of the given pixel. Figure 7 shows the desired result of using edge detection on a filtered image. In Figure 7 there are very few artifacts and the image is considerably clearer than Figure 6. The Canny algorithm is used for all the edge detection required by VBASR.



Figure 5. Example of a Normal Blur (compare to Figure 2)



Figure 6. Edge detection on an un-blurred image



Figure 7. Edge detection on a blurred image

One major problem with the image in Figure 7 is that computer still has no simple way of identifying the strongest vertical lines. Therefore, corner detection is used for the final stage of the preprocessing. Corner detection was performed on the edge-detected image enabling the program to obtain data points on the

lines. It was discovered that corners all fall on the lines, so this information was used to find the strongest vertical lines. As shown in Figure 8, the corner detection algorithm marks all of the corners with small white circles and outputs the x and y coordinates of each corner found into an array.



Figure 8. Corner Detection on image in Figure 7

2.3 Processing

The next step is to find the strongest vertical lines using the information shown in Figure 8. First, the image from the webcam is split into sixteen different bins. These bins allow a histogram-like transformation by counting the number of corners found within vertical bins. The end result is the summed number of corners found in each bin. Lastly, the number of corners in each bin is compared to a constant value, and if the number of corners is greater than that threshold, then that bin is considered a strong vertical line.

The image shown in Figure 9 is the same image used in all the other figures of this section. The thin white lines delineate each bin, the black lines represent the strongest vertical lines (as determined by the corners in Figure 8), and the thick white line represents the x-value average of the strongest vertical lines.



Figure 9. Lines algorithm processing example

To determine the direction found by the lines algorithm, the average of the strong vertical lines is found and compared to seven equally distributed direction bins (hard left, left, etc.). If the thick white line in Figure 9 were located on the left edge of the image, then it would evaluate to Hard Left. Likewise, if it was in the center of the image it would evaluate to Straight. The

particular image shown in Figure 9 evaluates to Slight Left as shown in Figure 10.



Figure 10. Direction bins displayed

2.4 Results and Shortcomings

In practice, the lines algorithm is successful on 76% of the images in the test set. The algorithm had the lowest success rate on images requiring the action of Hard Right, where a success rate of only 26% was achieved.

One shortcoming with this method appears when vertical lines fall directly on the separation line for a bin. When this happens, the corners found on that line may be split in between two bins and the line may be ignored completely. Notice the leftmost vertical line in Figure 11. This line falls directly on a bin line marked in Figure 12. The corners are split between the bins and the strong vertical line is ignored. (Note that in Figure 12 the bins are marked with thin black lines whereas the strong vertical lines are marked in thicker white lines.)



Figure 11. Lines shortcoming example



Figure 12. Lines shortcoming example

A second shortcoming occurs when VBASR is oriented directly at a wall (i.e. the image does not contain the center of the hallway at all). In these cases, the algorithm generally finds only one or two strong vertical lines. Depending on where these few lines are found it may determine a wildly inaccurate direction. Most often, the lines algorithm fails and the resolver function ignores the lines algorithm when deciding the final direction for VBASR. If the lines algorithm detects one or zero strong vertical lines, then the algorithm fails.

Incidentally, the lines algorithm did not work as originally intended. In practice, most corners for an image are found in the center of the hallway, not along the walls. As a result, the bins near the center of the hallway all are marked as strong vertical lines, which enables the lines algorithm to perform well regardless of this unexpected outcome.

3. CORNERS ALGORITHM

3.1 Theory

After observing several images of the hallway it was observed that the floor of most images forms a trapezoidal shape, as outlined in Figure 13. The trapezoid is created by the intersection between the floor and the walls. The corners algorithm was developed to take advantage of the observation. If one leg of the trapezoid is higher on the image than the other, then VBASR is facing the longer, lower leg's wall and needs to adjust in the opposite direction

In Figure 13, the legs of the trapezoid are marked in orange and the top is blue. In practice, the edge detection algorithm actually finds the edge of the colored tile rather than the corner of the floor and wall. Even though this is technically incorrect it works just as well.




Figure 13. Corners algorithm theory

3.2 Preprocessing

The preprocessing for the corners algorithm is similar to that of the lines algorithm. First, a blur is used on the image to eliminate artifacts from the Canny line-detection algorithm. After the Canny algorithm, corner detection is performed on the line-detected image. In Figure 14, the lower left and right hand sections of the image are boxed-off to aid viewers in understanding how the corners algorithm processes. These boxes denote the thresholds where the algorithm searches for the legs of the trapezoid.




Figure 14. Preprocessing for the corners algorithm

Frequently, the lines detected accurately define the trapezoid legs, but the corner-detection algorithm fails to find a corner (denoted by white circles) on the legs of the trapezoid. Note that the left-hand box in Figure 14 is an example of such a chase. Therefore, to aid the corners algorithm, two vertical lines are drawn on the image frame on either end of the image. The extra vertical lines on help the algorithm locate corners on those legs, as shown in Figure 15. (Interestingly, no corner was found at the intersection of the lines in the right-hand box.)




Figure 15. Extra vertical lines (compare to Figure 14)

3.3 Processing

Each of the corners found within the boxed-off sections of the image are generally on the legs of the trapezoid (i.e. along the border where the floor meets the wall). The x and y-value of each corner within the boxes are averaged to minimize the effect of outliers. Finally, the y-values of the final averages are compared and the leg with the higher y-value indicates the direction VBASR

should turn. The distance between these final averages also indicates the strength of the turn.

Figure 16 shows an example of the complete corners algorithm. The target marks show the final averages of the corners located in each box. When the target marks' y-values are compared, Figure 16 evaluates to Slight Left.



Figure 16. Corners algorithm example

3.4 Results and Shortcomings

The corners algorithm was successful on only 51% of the test images. Even though the percentage is low, the corners algorithm still improves the overall performance of the system because the corners algorithm sometimes finds the correct direction for images on which the other two algorithms fail.

Surprisingly, the corners algorithm fails the most for Slight Right and Slight Left images. Because the corners algorithm averages all the corners found within the boxes it is more likely to find large differences rather than smaller ones. As such, the corners algorithm performs better for larger turns and thus complements the lines algorithm well, which tends to fail on the Hard Left and Hard Right turns.

An obvious shortcoming of this algorithm is that not all of the corners found within the boxed-off regions of the image are directly on the trapezoidal legs. As shown in Figure 17, extra corners pull the target mark off of the desired position and effect the decision made by the corners algorithm.




Figure 17. Corners algorithm shortcoming

The corners algorithm fails altogether if no corners are found within either of the boxed-off regions. If the corners algorithm fails, then the resolver function ignores the corners algorithm when calculating the final direction for VBASR.

4. COLORS ALGORITHM

4.1 Theory

The third algorithm implemented takes advantage of the color difference between the floor and the walls. In most buildings, the floor color is distinguishable from the wall color. Bradley University's engineering building is no exception, as seen in Figure 2. If the floor can be identified and marked, then the image becomes a binary image of "floor" and "not-floor."

4.2 Preprocessing

An OpenCV library command called "flood fill" is used for the colors algorithm. A single pixel is picked as the seed point and then the neighborhood of that pixel is evaluated. If the neighboring pixels are similar enough to the seed point then all of the similar neighboring pixels are set to a predefined value, such as the red shown in Figure 18. The command continues to expand outward until there are no more similar pixels found. Flood fill only evaluates outward from the seed pixel and does not evaluate the entire image. As a result, the ceiling is not painted red even though it is a very similar color to the floor. The seed pixel for all Figures 18-24 are shown as blue circles.




Figure 18. Example of the Flood Fill command

4.3 Processing

After a binary image is achieved the resulting image is scanned from the top down. The first row with more than twenty red pixels is selected, and the x-values for those red pixels are averaged. The result is considered the center of the hallway, as shown in Figure 19, where the thick pink line indicates the decision line. Finally, the direction is determined by comparing the location of the decision line with the seven direction bins, in the same manner as the lines algorithm discussed above. This particular example evaluates to Straight, as shown in Figure 20.




Figure 19. Example of the colors algorithm




Figure 20. Colors algorithm with direction bins

If the center of the hallway is not in frame, then the highest row of red pixels still indicates the correct direction to navigate. In Figure 21, a hard right is required and the colors algorithm finds the correct direction successfully.



Figure 21. Colors algorithm for off-center images

4.4 Results and Shortcomings

Easily the best of the three algorithms, the colors algorithm has a success rate of 95%. This algorithm has no particular category of images for which it performs poorly. Unfortunately, many shortcomings still exist for this algorithm.

The first shortcoming is that the seed-point cannot be adjusted once it is set. It is possible for the seed-point to fall on the wall instead of the floor. If this occurs then the flood fill command will paint the walls red, which is clearly undesirable. Due to the

mounting configuration of the camera on the robot platform and the lack of inclines on the hallway floors, the horizon line for the camera should never change. If the camera is placed such that the horizon line is approximately halfway up the image, any red pixels found above the horizon line indicate that something other than the floor has been marked red. In this case, the colors algorithm fails and is ignored when the resolver processes the final direction for VBASR.

The second shortcoming is that tiles of different colors can confuse the algorithm. The orange tiles in Figure 22 are rightly identified as not-white, however, they are still part of the floor. Likewise, if the seed-point falls on an orange tile (similar to the first shortcoming), only the orange tile is filled while the rest of the floor is ignored as “not-orange,” see Figure 23.




Figure 22. Discolorations in the floor




Figure 23. Seed-point location causing algorithm failure

Lastly, reflections on the floor are another challenge. The flood fill command may not recognize the reflections if they contrast too much from the neighboring pixels. Notice that in Figure 24 many reflections cause gaps in the red floor.



Figure 24. Colors algorithm and reflection

To work around these shortcomings several different seed-points are used. Each pixel value at each seed point is identified and if that point is either white or orange then it is evaluated (orange to catch the random orange tiles). Figure 25 shows the benefit of adding extra seed points. If a seed point has been evaluated it is also filled in black for simple troubleshooting.



Figure 25. Many seed points

5. RESOLVER

After all three algorithms independently determine a direction to navigate, they are resolved into a single direction for the entire system. The resolver ignores algorithms when it detects a failure (e.g. the color algorithm is ignored if it paints the walls and ceiling red). Since the resolver evaluates each algorithm in parallel, the system architecture is such that algorithms can be added and removed without compromising the integrity of the system as a whole. Programming with such a parallel architecture enables the effortless addition or removal of an algorithm as necessary. As such, failed algorithms can easily be ignored in the following computations.

Each direction from each algorithm is given a numerical value (0 = Hard Left through 6 = Hard Right). The numerical values are averaged to determine the final direction value. Normally, integer division truncates, which bias the system towards the left. The averaging algorithm eliminates the left-bias. To address this issue, VBASR’s averaging algorithm is biased towards the center, favoring outcomes closer to driving straight. If the average is greater than three (straight) then it is rounded down. If the final average is less than three it is rounded up. Therefore, the resolver biases the VBASR to travel straight rather than biasing in one direction or the other. Averaging the results of several algorithms

and favoring the center also buffers the robot from erratic directional swings in its navigation decisions.

6. RESULTS

The table shown in Table 1 details the results for the best vision algorithm. The success rating for each algorithm is broken into each type of picture. Certainly the best algorithm is the colors algorithm although if either of the other two algorithms is eliminated then the total percentage lowers.

Overall the vision system is a great success and has exceeded expectations.

Table 1. Final Vision System Results (%)

	Lines	Corners	Colors	Resolved
Hard Left	53.3	26.7	93.3	100.0
Left	93.1	72.4	93.1	98.3
Slight Left	97.1	34.3	91.4	94.3
Straight	100.0	53.6	96.4	98.2
Slight Right	97.6	46.3	92.7	100.0
Right	61.2	81.6	96.9	98.0
Hard Right	26.3	42.1	100.0	84.2
Totals	75.5	51.0	94.8	96.1

7. FUTURE WORK

The discussion in this paper focuses on navigation of a hallway where obstacles are located only along the walls. Although generally not observed in the halls of the engineering building, it is possible that obstacles may be placed in the center of the hallway, obstructing the path of the robot. To handle these situations, obstacle avoidance will be implemented utilizing an algorithm similar to the colors algorithm. Creating a binary image of “floor” and “not-floor” enables simple detection of obstacles

because the obstacles should generally be a different color than the floor. The orange tiles in Figure 16 that are not marked red give a general feel for how the algorithm could identify obstacles. These orange tiles also pose a challenge as they should be identified as “floor” rather than “not-floor.” Using the more robust colors algorithm shown in Figure 25 the floor can be identified, and, thus, the obstacles are isolated.

In addition, VBASR should be able to react faster than humans. To do so, VBASR must be able to process an image and start reacting within 190ms[4]. Currently, VBASR processes about ten images per second which meets the requirement.

8. ACKNOWLEDGMENTS

We thank Bradley University’s Electrical and Computer Engineering Department and Northrup Grumman for sponsoring VBASR: The Vision System.

9. REFERENCES

- [1] Sage, Kingsly, and Stewart Young. Security Applications of Computer Vision. Tech. IEEE AES Systems Magazine, 1999. IEEEEXPLORE. Web. 20 Oct. 2009. <<http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=756080&isnumber=16393>>
- [2] Shapiro, Linda G., Linda G. Shapiro, and George Stockman. Computer Vision. Upper Saddle River: Prentice Hall, 2001. Print
- [3] "iRobot Create Premium Development Package." IRobot. IRobot Corporation, 2009. Web. 10 Nov. 2009. <<http://store.irobot.com/product/index.jsp?productId=2591901&cp=2591511&parentPage=family>>.
- [4] Kosinski, Robert J. "Literature Review on Reaction Time." Clemson University, Aug. 2009. Web. 10 Nov. 2009. <<http://biae.clemson.edu/bpc/bp/Lab/110/reaction.htm>>.