Vision-Based Obstacle Detection and Avoidance for the CWRU Cutter Autonomous Lawnmower

Alexander Schepelmann, Henry H. Snow, Bradley E. Hughes, Frank L. Merat, Roger D. Quinn
Case Western Reserve University, Cleveland, Ohio, 44106, axs287@case.edu

James M. Green
MTD Products, Inc., Valley City, Ohio, 44280, jim.green@mtdproducts.com

Abstract **–** This paper describes the vision-based obstacle detection system of the CWRU Cutter, an autonomous lawnmower developed for the annual "Institute of Navigation (ION) Autonomous Lawnmower Competition." Unlike LIDAR sensors commonly found on autonomous vehicles, computer vision systems can provide similar information at drastically reduced prices. Though significantly more cost-effective than LIDAR, these systems have inherent problems due to changing lighting conditions and shadows. This paper investigates the use of image hue and intensity to create a robust, real-time vision-based obstacle detection system for use during the ION competition. Data abstraction methods used to process incoming images for easy combination of information from multiple sensors are also discussed. Using this system, CWRU Cutter correctly identified obstacles in 89% of frames containing fence, 78% of frames containing flowerbeds, and 84% of frames containing boundary lines.

INTRODUCTION

Autonomous robots are increasingly becoming part of our daily lives. Such robots are designed to perform specific household chores like vacuuming the floor or ironing clothes [1][2]. Autonomous lawnmowers have recently gained popularity, yet, currently available consumer versions do not possess the ability to mow parallel lines or other more complicated patterns and cannot sense obstacles from a distance [3][4]. The CWRU Cutter, an autonomous lawnmower developed at Case Western Reserve University, accomplishes both.

CWRU Cutter was designed for the "Institute of Navigation (ION): Autonomous Lawnmower Competition." The competition emulates a backyard environment where commercial autonomous mowers could operate and has a variety of common obstacles like fence, flowerbed lining, and a mobile obstacle (a stuffed dog mounted on an RC car) inside the contest course, which is bounded by a series of white lines [5].

CWRU Cutter uses a variety of sensors to intelligently cut grass during the competition. Navigational information is determined using a differential GPS, Christa IMU, and wheel encoders, which are combined using a Kalman filter. Previous versions of the CWRU Cutter used a SICK LMS291

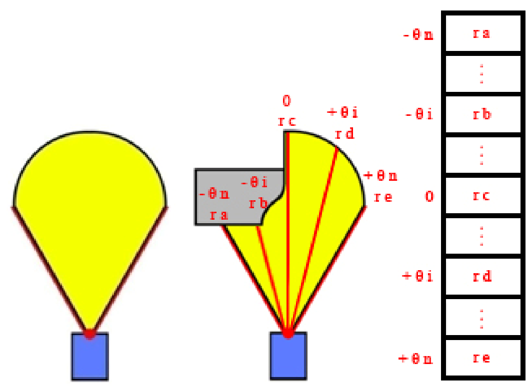

Figure 1 - CWRU Cutter polar freespace representation and 1D range array. Left & Center: CWRU Cutter (blue) and LIDAR field of vision (yellow) between –θn to +θn (outer red lines) centered at r=0 located on a point on mower (red dot). Left: LIDAR field of view when an obstacle is not present. Center: LIDAR field of view when an obstacle (grey) is present. Right: LIDAR returns a range at every observable θ sequentially placed in a 1-D array, to which a polar function is fit. This gives a polar representation of the reachable space in front of the mower. As the mower travels, these ranges are shifted, creating a 360º representation of the polar freespace around the robot.

LIDAR unit for obstacle detection, which provided ranges to the nearest obstacle at 1º degree angular increments [6]. This data was used for reactive obstacle avoidance and navigation during mower operation. The data output by the LIDAR is naturally represented in a polar coordinate system, creating range images of the free-space around the mower. This is referred to as the "polar freespace" (Figure 1).

However, LIDAR's price makes its inclusion in consumer versions of the CWRU Cutter prohibitively expensive. Therefore, alternative methods of obstacle detection are being investigated.

A computer vision-based obstacle detection method is being developed as one cost-effective alternative for use in the competition environment. Though significantly cheaper, computer vision systems have inherent problems with consistent identification of an object due to changing lighting conditions and shadows [7]. This has kept vision-based obstacle identification methods from being widely adopted in other outdoor consumer products.

Unlike LIDAR data, camera images are typically represented in rectangular coordinates. Since all previous CWRU Cutter software and control routines utilized the polar representation of freespace output by the LIDAR and it was desired to continue using the same software, measurements from other sensors must therefore be abstracted and represented in polar coordinates to allow for easy integration with the existing software. Abstracting data in this manner also allows for easy combination of information from multiple sensors.

This paper discusses the methods used to create a robust, real-time vision-based obstacle detection system for use during the ION competition. The use of the hue color plane to detect competition obstacles is discussed first. Next, data abstraction methods used to process incoming images are laid out. A brief explanation of the robot's polar freespace observer, to which the abstracted data is passed, follows. White line observation methods, which utilize both hue and intensity to keep the robot inside of the course, are also explained. The paper ends with a discussion of metrics used to measure the success of CWRU Cutter's obstacle detection in a competition environment using these methods, results, and a discussion of how additional processing methods can be used to create an even more robust obstacle detection system.

HARDWARE & SOFTWARE SPECIFICATIONS

Images are captured at 640x480 resolution by a DFK 21AF04 firewire camera at a maximum 30 frames/sec. Incoming images are processed by a 1.83 GHz Mac Mini running Windows XP with 4GB of memory.

Due to the processing limitations of the Mac Mini, algorithm complexity was a primary concern. Though lighting problems can be addressed by computationally intensive processing methods, the CWRU Cutter's onboard computer limits the complexity of the algorithms that can be used, since the robot needs to capture and analyze incoming images in real-time. To safely track obstacles while the mower is operating, the target speed for algorithm operation was set at a minimum of 10 frames/sec.

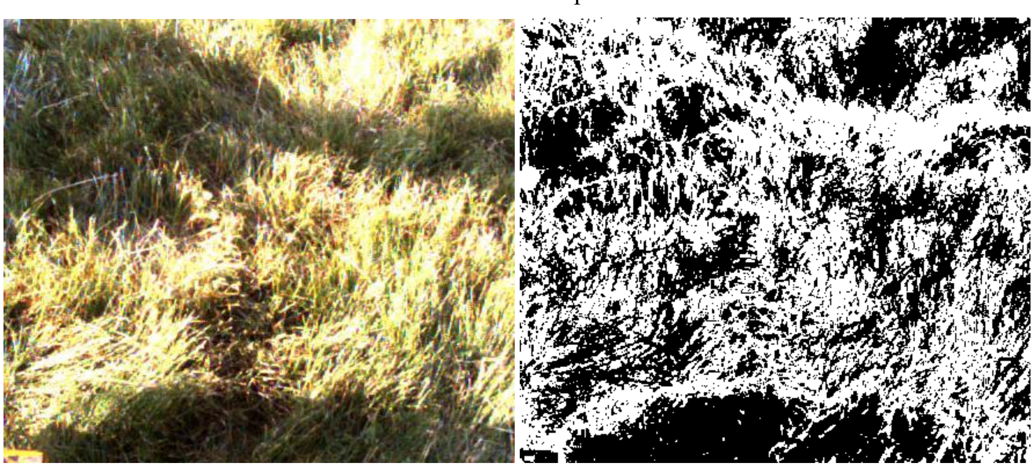

Figure 2 – Incoming camera image (L) and resulting image after a binary threshold in the red color plane within one standard deviation of the mean image color. Though the entire image contains grass, the color threshold is unable to successfully quantify the entire contents due to shadows and illumination overexposure.

CWRU Cutter is programmed using National Instrument's LabVIEW 8.6 programming language. LabVIEW offers several advantages, including easy sensor interfacing and integration as well as the ability to operate in real-time. Additionally, LabVIEW's "Image Acquisition" (IMAQ) Vision package contains many optimized functions to allow for the creation of a functional robotic system capable of operating in real-time. As such, CWRU Cutter uses available IMAQ vision functions where possible.

PROCESSING METHODS

Obstacle Detection

Observed RGB color values of an image change under varying lighting conditions [8]. This makes it difficult to recognize an object that is partially occluded by shadows (Figure 2). Additionally, the exposure of consecutive camera frames can vary drastically. This is especially problematic on an autonomous robot that may be travelling between illuminated and shaded areas in an outdoor environment. Because of these problems, it is impossible to use simple RGB color thresholding to consistently recognize an object such as a green lawn in an image.

Unlike the RGB planes, the hue plane has been found to be relatively insensitive to changing lighting conditions and shadows. Instead of indicating how much red, green, and blue are present in a color, hue is an indication of *how much* of a certain color is present at a pixel. Therefore, though grass may be lighter or darker in some regions of the image, all grass containing pixels are classified under similar hue values. This creates fewer false positives that can be more easily filtered out through post processing (Figure 3).

Using hue plane is especially well suited for obstacle identification in the ION competition. Since all obstacles are either white (i.e. the picket fence) or black (the flowerbed lining and stuffed dog) they do not fall within the hue range that corresponds to grass and are therefore easily identified within an image.

Image processing for obstacle detection is accomplished on the CWRU Cutter in the following way. Before mowing, a sample image of the competition field is taken and the mean hue value of grass in the image is calculated for reference. Since the hue value of grass in the mowing area is confined within a narrow band, this only needs to be done once per mowing location. During mower operation, the input image is transformed from pixel to real-world coordinates. We refer to this as "calibration," which was determined prior to mowing. Next, the hue plane is extracted from the image using the "IMAQ Threshold" function and a threshold is applied within ± one standard deviation of the mean hue value of the grass. The threshold extracts the matching pixels to a binary image with the same real-world calibration as the image, creating a binary representation of what is passable terrain (represented by a 1) and potential obstacles (represented by a 0). Binary image dilation and hole-filling operations are performed on the array to remove small false positives in the image. Large neighborhoods of obstacle locations remain in the array and correspond to obstacle locations in the current image. The resulting array is then converted into a range image in polar coordinates, which

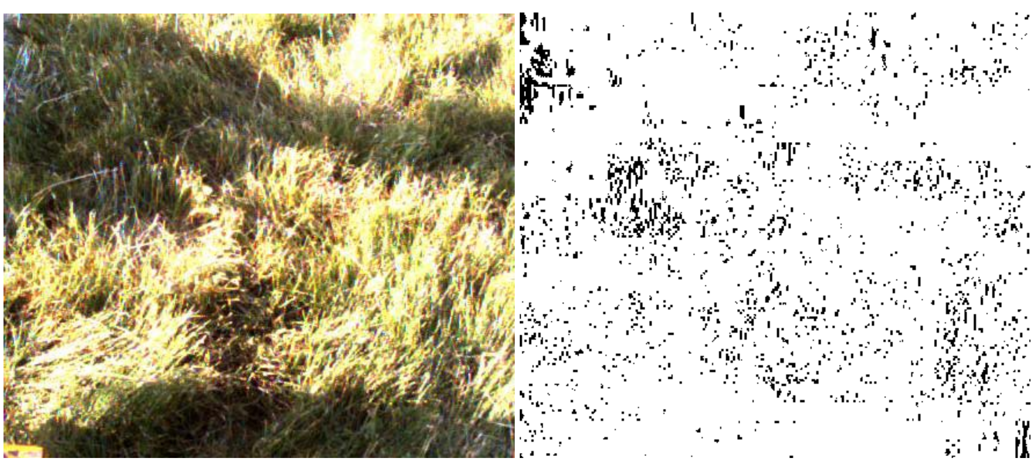

Figure 3 – Incoming camera image (L) and resulting image after a binary threshold in the hue plane. Though heavy amounts of shadow exist in the image, most of the pixels in the image have a similar hue value, which appears as one continuous object in a binary array. Small false positives are later filled with image dilation.

integrates directly with CWRU Cutter's "Polar Freespace Observer."

The range image conversion occurs in the following way. First, the binary array is converted from rectangular real-world coordinates to a polar representation using LabVIEW's "Rectangular-to-Polar Array" function. This creates a binary range-image of the mowable terrain and obstacle locations in the current camera frame, where the (0,0) location of the camera is located in the middle element of the bottom row of the array. Corresponding (r,θ) pairs are then sorted using LabVIEW's "Sort 1D Array" method, and the shortest range for each integer value of theta in the camera's field of view is inserted into a 1-D Array shifted so that the ranges are relative to the (0,0) location of the robot body. Ranges are inserted into the array sequentially, which creates a data structure identical to what is output by the LIDAR, where the first element corresponds to the range to the nearest obstacle at $\theta=1°$, the second element corresponds to the range to the nearest obstacle at $\theta=2°$ degrees, etc. If an obstacle is not observed at a θ value, a range of 4 meters is entered into the corresponding 1-D array element, a distance equivalent to the maximum range of the SICK LMS291, indicating that no obstacle is present at that angle. This 1-D range array is referred to as a "pseudo-LIDAR" scan. To visualize the pseudo-LIDAR scan, a polar function is fit to the 1-D array and plotted in LabVIEW's front-panel. This allows the user to monitor what obstacles CWRU Cutter was or was not able to see (Figure 4). This range image processing occurs at a minimum of 10Hz during mower operation. The resulting 1-D array is passed to the freespace observer.

The freespace observer accomplishes two things. First, it combines the new pseudo-LIDAR observations with previously observed freespace to create an accurate estimate of traversable terrain around the robot. Second, it shifts observed ranges in front of the robot to behind the robot in accordance with how the robot moves through the environment. This shifted freespace is simply the previous freespace estimation, shifted according to the change in position and orientation of the robot's center. This creates a 360º representation of the currently traversable space around the robot, which is passed to an onboard NI compactRIO (cRIO) responsible for navigation.

White Line Detection

The ION competition field requires CWRU Cutter to observe white lines to keep the robot in bounds. White lines are observed in the hue and saturation planes. Calibration is applied to the incoming image. Next, a threshold is applied in both planes to extract the white line in the image. This is used to create two binary images of the observed line. A line is fit through the observed points in each plot using a "Random Sample Consensus" (RANSAC) algorithm, which returns a linear fit of the observed points and its standard deviation. RANSAC was utilized because it has been very successful for line observations on other autonomous robots [9]. The robot's displacement is calculated from this line. If the robot's distance to the line falls below a minimum value, the robot turns away from the line to stay in bounds.

Both hue and saturation are used to ensure robustness. If the standard deviation of the new line

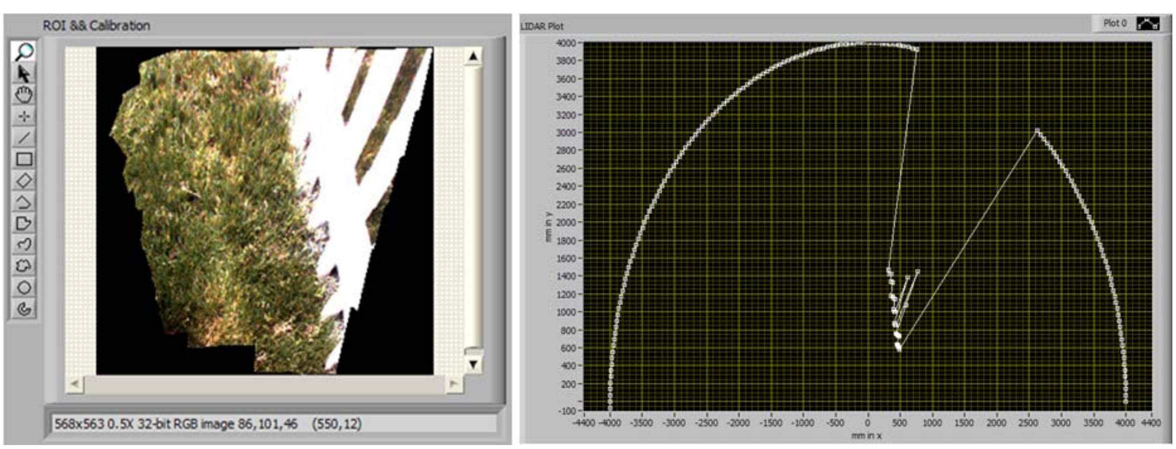

Figure 4 – Incoming calibrated camera image of the white competition fence (L) and resulting pseudo-LIDAR scan. In this example, since the camera is only able to observe ranges for θ values between 90-55°, ranges at non-observable θ values are filled with 4000mm.

in hue or saturation is too large relative to previous line fits, the linear fit of the line in the other is used for the estimate, and subsequent measurements rely on that. In this way, the robot switches between hue and saturation observations to ensure that the line is accurately observed over time.

Testing Procedure

CWRU Cutter's onboard camera recorded images while the robot was driven past the competition fence, flowerbed lining, and white lines via remote control in a field which mirrored the competition environment. During recording, the robot approached the obstacles from multiple angles to collect representative data of how the robot could approach an obstacle during the competition. As in the competition, only one type of obstacle was present in each frame.

The recorded data was then played back in real-time executing on CWRU Cutter's onboard computers to ensure that real-time processing was maintained. From the recorded data, representative frames were sampled of when the mower was near an obstacle and analyzed individually. In total, 205 frames near fence, 244 frames near flowerbed lining, and 188 frames near white lines were analyzed.

A frame was marked as "incorrect" if the plotted pseudo-LIDAR scan identified false positives within the camera's field of view in frames that did not contain obstacles or failed to identify an obstacle in a frame. Identification success rate was determined by dividing the number of correctly identified frames by the total number of frames for each obstacle type.

CONCLUSION

CWRU Cutter successfully processed incoming images at a speed of 10Hz or greater throughout the duration of the simulation. CWRU Cutter correctly identified 89% of representative fence frames, 78% of flowerbed frames, and 84% of white line frames. Optical overexposure of grass areas and insufficient obstacle width within a frame were the two primary contributing factors to incorrect frames.

If an area of grass was strongly illuminated, the region appeared to be white and the color information of that region was lost. This caused the algorithms to incorrectly perceive an obstacle in that location. If the overexposed area was large enough, it was not fully filtered out by the dilation operation, which caused the remnants to be marked as obstacle locations. It is important to note that subsequent frames did not contain identical false positives, since the perceived color of a region changed as the robot's position relative that that region changed.

False positives were also introduced if the observed obstacle width in a frame was not sufficiently large, since small identified regions were filtered out by the subsequent dilation operation. This caused the robot to fail to identify an obstacle in a frame. Again, this was remedied in subsequent frames as the mower's position and orientation relative to the obstacle changed.

During mower operation, false positive and obstacle identification failures in individual, non-subsequent frames would not be a problem, since pseudo-LIDAR data would be combined with previous observations. If an obstacle suddenly appeared in a frame where no obstacle was present in previous observations, the confidence of the observation would drop and would not be heavily weighted in the freespace around the robot. Similarly, if an obstacle suddenly disappeared from a location previously observed to contain an obstacle multiple times, the robot's observed freespace would remain largely unaffected.

Future research will involve improving camera-based obstacle recognition techniques using directional visual texture to supplement obstacle observations. Visual texture has been shown to be robust in an outdoor environment, and, due to its computational efficiency, is potentially well suited for use on the CWRU Cutter.

REFERENCES

[1] http://store.irobot.com/product/index.jsp?productId=2804958

[2] http://www.kitchengadgetry.com/kitchen/guides/aid-74/fagor-driron

[3] R.W. Hicks II, E.L Hall. "A survey of robot lawn mowers," *Proc. of SPIE.* San Diego, USA, vol. 4719, pp. 262-269, 2000.

[4] http://www.friendlyrobotics.com/faq/?kb=9

[5] "3.3 Field Description," in *The Sixth Annual Robotic Lawnmower Competition Rulebook*. Dayton, Ohio: Institute of Navigation, 2009. [Online]. Available: http://www.ion.org/satdiv/alc/rules2009.pdf. [Accessed: July 24, 2009].

[6] *CWRU Cutter 2 Technical Report*. Cleveland, Ohio: Case Western Reserve University, 2009. [Online]. Available: http://www.ion.org/satdiv/alc/reports2009/CWRU_Cutter_2_Technical_Report.pdf. [Accessed: July 24, 2009].

[7] E. Salvador, A. Cavallaro, T. Ebrahimi. "Spatio-temporal shadow segmentation and tracking." *Proc. of SPIE. Image and Video Communications and Processing*. Santa Clara, USA, vol. 5022, 2003, pp. 389–400.

[8] J. G. Liu, J. M. Moore "Hue image RGB colour composition. A simple technique to suppress shadow and enhance spectral signature." *International Journal of Remote Sensing*, vol. 11, iss. 8, 1990, pp. 1521 – 1530.

[9] S. McMichael. "Lane Detection for DEXTER, an Autonomous Robot, in the Urban Challenge." M.S. Thesis, Case Western Reserve University, Cleveland, OH, USA, 2008.