Is the Sun Too Bright in Queensland?
An Approach to Robust Outdoor Colour Beacon Detection

Ashley Tews, Jonathan Robert, Jonathan Roberts, and Kane Usher

CSIRO Autonomous Systems Laboratory
ICT Centre
P.O. Box 883, Kenmore 4069
Queensland, Australia
Email: Ashley.Tews,Jonathan.Robert,Jonathan.Roberts,Kane.Usher@csiro.au

Abstract

Using cameras onboard a robot for detecting a coloured stationary target outdoors is a difficult task. Apart from the complexity of separating the target from the background scenery over different ranges, there are also the inconsistencies with direct and reflected illumination from the sun, clouds, moving and stationary objects. They can vary both the illumination on the target and its colour as perceived by the camera. In this paper, we analyse the effect of environment conditions, range to target, camera settings and image processing on the reported colours of various targets. The analysis indicates the colour space and camera configuration that provide the most consistent colour values over varying environment conditions and ranges. This information is used to develop a detection system that provides range and bearing to detected targets. The system is evaluated over various lighting conditions from bright sunlight, shadows and overcast days and demonstrates robust performance. The accuracy of the system is compared against a laser beacon detector with preliminary results indicating it to be a valuable asset for long-range coloured target detection.

1 Introduction

Image sensors are relatively inexpensive and provide a rich source of information about the environment and objects within their vision. However image data typically requires extensive processing to produce the same data as other, more dedicated sensors such as the distance information supplied from scanning laser rangefinders. The focus of this paper is towards replacing laser beacons with coloured beacons for robot localisation. A coloured beacon is considered to consist of one or more coloured targets in close proximity. From this perspective, the project is split into two phases - how to maintain colour consistency of the beacons with outdoor lighting conditions, and developing the beacon detection algorithm.

In the first phase, there are many environmental and system influences that affect how a coloured object is perceived. The illumination conditions of the environment change throughout the day due to variations in position of the sun, the influence of cloud cover and the reflections of light from objects and infrastructure surrounding the object (e.g. buildings, bright concrete surfaces). Each of these impacts on the colours of an object perceived by the camera, and therefore the image produced, even if the camera is automatically compensating for the varying conditions. The size of the object in the image also determines how much influence the surrounding scenery has on the auto-adjustments of the camera. That is, the further the target object is from the camera, the smaller its area and the camera adjustments will be influenced by more of the scenery than the target. A major goal is to minimise these effects to achieve relatively consistent colour readings.

There is a distinction between our research and research into colour constancy. Our research is more pragmatic in the sense that it is not important to maintain colour constancy over the objects in the image, only for the selected targets' colours. As demonstrated by [Austin and Barnes, 2003], the changes in a target's colour values over varying lighting conditions can be non-linear which is more evident in some colours than others for the colour space chosen. With the constraints of the colours of the targets available, we focus primarily on how to maintain the consistency of their colours. We initially analyse the performance of several colour spaces with varying sunny, cloudy and shadowy illumination conditions, camera parameters and beacon ranges in an outdoor environment. The outcomes highlight the most consistent colour space, beacon colours and camera parameters. The resulting configuration is used in the beacon detection system. For uniqueness against the environment, beacons consist of two coloured targets located adjacent to each other. The detection system uses a simple histogram thresholding approach across the image to detect beacons. The range and bearing information is returned from the size and location of the beacon in the image. Experiments are conducted under various illumination conditions and with various backgrounds to validate our approach.

The remainder of this paper is outlined as follows. Section 2 overviews research into colour consistency and visual tracking in outdoor environments. Section 3 provides the analysis of the environment and system influences on coloured targets. Section 4 provides the details of the beacon detection system initialised with the settings highlighted in the previous section. The results of locating coloured beacons in varying lighting conditions is included. The conclusions from this research are presented in Section 5 including future work.

2 Related Work

The search for solutions to maintaining an object's colour despite varying illumination conditions extends across many research disciplines including image processing, computer graphics and robotics. It also extends across illumination sources from artificial lighting, to natural indoor and outdoor lighting. Our research focus is on robotic applications in outdoor environments with completely natural lighting. The related work overviewed in this section is a sample of the research associated with this focus.

[Austin and Barnes, 2003] analysed variations in colour over time and varying indoor lighting conditions. They used a HSV model and determined that the change in colour over varying lighting conditions was mostly non-linear. The non-linearity was worse for natural lighting conditions in their experiments.

The general problem of maintaining colour constancy under varying illumination is difficult due to non-linearities in colour shift. Thresholding techniques can be used but increase in complexity as the number of colours in the image increases. [Gevers and Stokman, 2005] state their technique of using variable kernel density estimators to model the colour histograms is more robust than thresholding techniques on traditional colour spaces such as normalised RGB. Contrasting to this approach, [Stokman and Gevers, 2005] apply an online weighting technique to discriminate between colours in images from an initial training set using manual segmentation. These techniques have been demonstrated under controlled lighting with good results.

In outdoor environments where there is little control on daytime lighting, maintaining consistent colours is more difficult. In the area of general illumination compensation, [Todt and Torras, 2001] propose a method using colour ratios for colour constancy in outdoor images. The goal is for images to be as consistent as possible across natural illumination intensities so downstream processing can extract salient landmarks.

In an effort to determine the effects of sunlight and skylight on colour, a CIE daylight model was developed by [Judd *et al.*, 1964]. The model was an accumulation of data samples taken from various constrained areas of the sky at different times of the day in different countries. From the model, the effect of illumination can be used in determining the apparent colour of an object. However, [Buluswar, 2002] determines that the CIE colour space is inaccurate for outdoor image processing due to the inaccuracies with modelling the illumination at the source (i.e. sun and sky) rather than at the destination where reflections also contribute. Instead, the author proposes a daylight model based on parameterising natural lighting features taken from images of particular objects over the course of varying lighting conditions (cloud cover, sun angle, etc).

More closely related to our research, [Browning and Veloso, 2005] demonstrate a fast adaptive coloured object tracking approach for an outdoor RoboCup application. The algorithm uses thresholded YUV histograms to classify pixels into various colour classes. Regions of each colour class are further analysed to adjust the thresholds for the current conditions. The technique adapts to small changes in illumination under the assumption that each colour will change by the same amount. Subsequent steps use template matching and geometrical constraints to classify regions as known objects.

It is clear that there are many techniques for researchers to try to achieve consistent colours over varying illumination conditions for their applications. However, many involve significant image processing which may be unpractical for embedded systems. We seek a simplistic approach, partly for processing efficiency, and the need to be robust to natural lighting conditions in a complex environment, varying distances and viewing angles. The remainder of this paper outlines our initial analysis of system parameters and environment conditions, leading to a target detection mechanism capable of detecting targets over distances considered large for our application.

3 System and Environment Variable Analysis

There are many parameters that affect the colour of an object as perceived by a camera in outdoor environments including illumination sources, scenery colour and reflectivity, camera hardware, camera settings, shadows, object surface properties, object shape and how the image output from the camera is processed. Of these, the main parameters that can be controlled (without extensive engineering of the environment) are the choice of camera, camera settings, object surface covering and image processing steps. Generally, the first stage of image processing is to convert the image output from the camera to a moderated feature space - usually a pixel conversion to a more application-dependent colour space. This selection acts as feedback to the hardware parameters. An effective system results from balancing the hardware parameters with the colour space, and selecting a colour space that produces stable target colours with changes in the environment.

This section provides an analysis of the effects of the above-mentioned parameters on colour consistency of various targets over different ranges. The goal is to find the best values for the controllable parameters that provides the smallest change and deviation of colour value for the non-controllable parameters. The following parameters

Figure 1: *The experimental platform consisting of several cameras, GPS, IMU and laser rangefinder. At the rear of the vehicle are the vision processing and control computers.*

are analysed in the remainder of this section:

1. target colour
2. colour spaces
3. camera parameters
4. cameras
5. environment conditions

The test vehicle used is CSIRO's Autonomous Tractor, (Figure 1), which is a ride-on mower which has been retrofitted with an array of actuators, sensors, and computer systems enabling the implementation and testing of sensors, and algorithms for control and navigation. Shown at the front of the tractor are two sideways-facing Unibrain firewire cameras with fisheye lenses, a forward-facing stereo camera head and a Marlin F145C2 camera configured as an omni-cam. The stereo head was replaced with a Unibrain camera with its standard lens which was used for the majority of the testing in this section. The test environment is the worksite at CSIRO's QCAT which is a mostly static environment containing large sheds, equipment and barrels.

3.1 Targets

The targets (Figure 2) used for the analysis in this section were common sheets of A2 (420 mm by 594 mm) craft paper obtained from a stationery store. Several different coloured targets were used including red, green, light blue, and yellow with each having consistent colour and specular reflection characteristics.

Our standard method adopted for gathering data from the targets was to locate them at one end of the environment take images from the camera every 3 m to the opposite end of the environment, approximately 50 m away. With the standard image size set to 640 by 480 pixels, the resulting target size in the images ranges from approximately 150 by 115 pixels to 7 by 7 pixels.

Figure 2: *Yellow, green, blue and red targets taped onto a portable office partition at 9 m from the camera.*

3.2 Robustness of Different Colour Spaces Outdoors

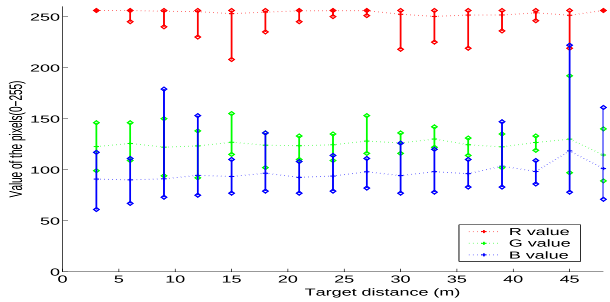

(a) Red.

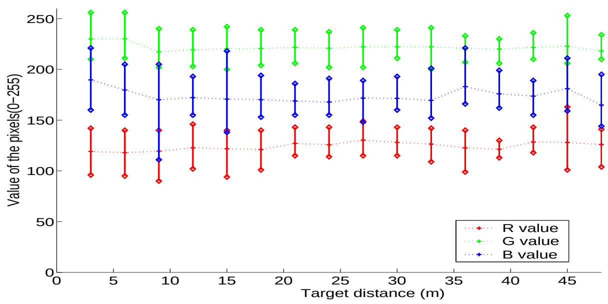

(b) Green.

Figure 3: *RGB values for the targets.*

The colour spaces tested were HSV, YUV, RGB and normalised RGB (r = R/R+G+B, g = G/R+G+B). The camera was set to auto-exposure and auto-white balance, and the image resolution set to 24 bit RGB. The images were taken on a sunny day (common in Queensland). Figures 3 to 6 show the mean, maximum and minimum values of the histograms of the red and green targets.

In evaluating the stability of the target colours over various ranges, minimal variation in the average colour value and histogram spread highlight a high performing colour space. The evaluation indicates that normalised RGB space outperforms the others for the system configuration and environment conditions. The YUV space has also shown high performance, although the colour variation and histogram spread over distance are more non-

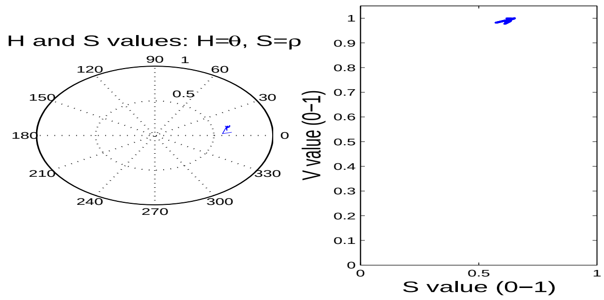

(a) Red.

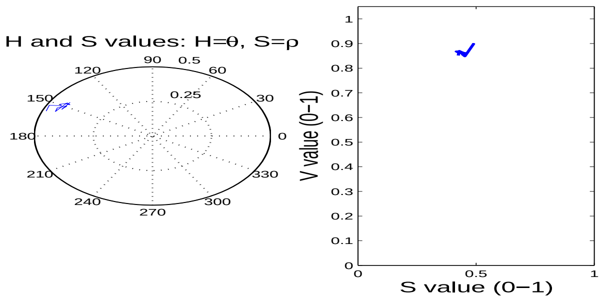

(b) Green.

Figure 4: *HSV values for the targets. The left figures represents H versus S and the right represents V versus S.*

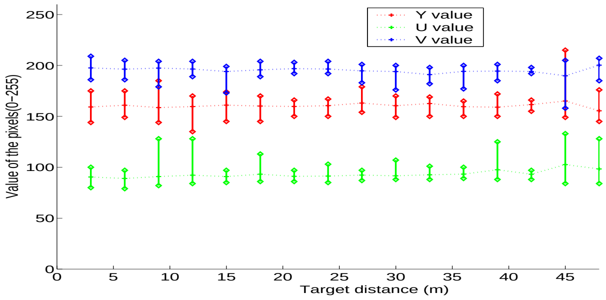

(a) Red.

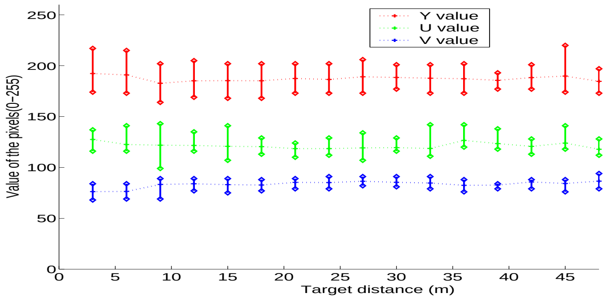

(b) Green.

Figure 5: *YUV values for the targets.*

linear than the normalised RGB space. The HSV space generally presents a non-linear shift in the target colours. The results also indicate that range to target has little influence over the target's reported colour.

3.3 Camera Settings

Using the normalised RGB values of the preceding analysis as a baseline, the effects of using manually adjusted white balance or exposure are analysed in this section to determine the utility of the camera's automatic functions for stabilising the colours in the images over varying distances. By changing the distance from the camera to the target, the reflectance and colours of surrounding objects can influence the perceived colour of the targets. The auto functions should be able to override these effects if the target area is made the subject of the image. In practice, this is generally not the case and the camera's exposure and white balance control will adapt to the entire image. Assuming the environment is static and illumination is constant as was the case in this analysis, manually setting the exposure and white balance should produce consistent target colours over distance. An example of the results are shown in Figures 7 and 8.

Generally, the manual camera settings shown in the figures demonstrate similar average colour values to the automatic settings shown in Figure 6. Manual exposure produced slightly higher variations in colour values.

The choice of settings for manual white balance and exposure are subjective since they are based on the user's analysis of the live image which will be different between calibration sessions and users. However, it indicates that for the colours tested, the automatic settings of the Unibrain camera provide equal or higher colour stability than the preset parameters. Had this not been the case, the camera auto-functions would need to be overridden by exposure control or white balance programs external to the camera.

3.4 Colour Space Analysis with the Marlin Camera

Aligned with the goal of determining the most appropriate parameters for outdoor coloured-beacon detection, alternative hardware needs to be evaluated. The Unibrain firewire camera was used for the analysis in the preceding sections. In this section, a Marlin camera fitted with an 8 mm lens is used. Images are obtained with all camera settings set to automatic except exposure which is not a feature of the camera. The images are taken at 3 m intervals up to approximately 50 m and are processed through the various colour spaces similar to the analysis in Section 3.2. Example results for light blue and green targets are shown in Figures 9 to 12.

The most noteable differences between the images from the Marlin and the Unibrain are the lower reported illumination of the image (evident in the Y values in YUV space), and the colour values in each colour space. The RGB, HSV and YUV values generally vary significantly from their Unibrain equivalents for the given targets. The normalised RGB space is more similar and stable indicating it is less sensitive to the camera hardware tested.

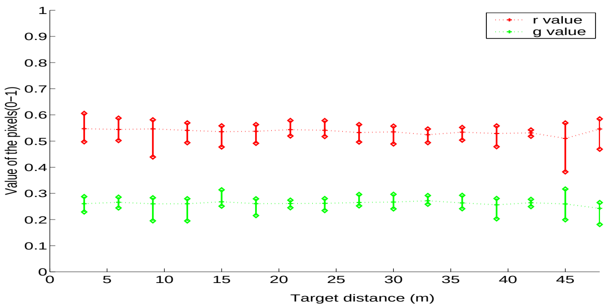

(a) Red.

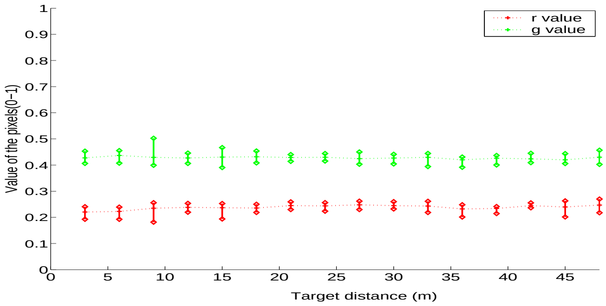

(b) Green.

Figure 6: *Normalised RGB values for the targets.*

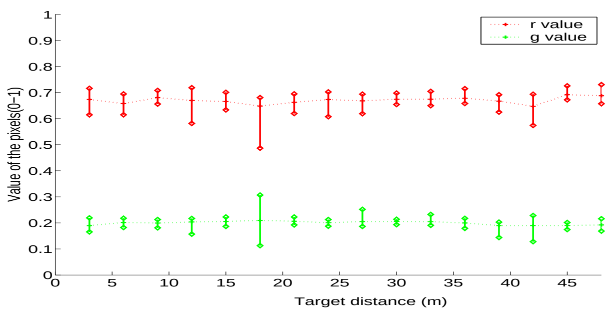

(a) Red.

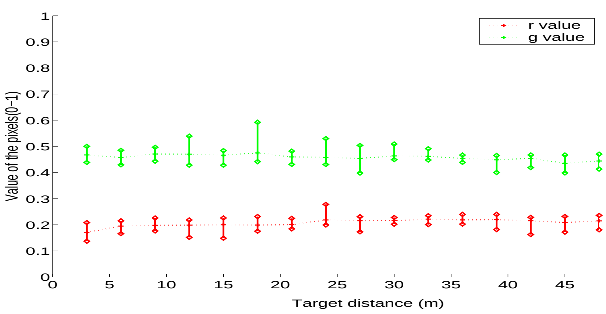

(b) Green.

Figure 7: *Normalised RGB values for the targets with manual exposure.*

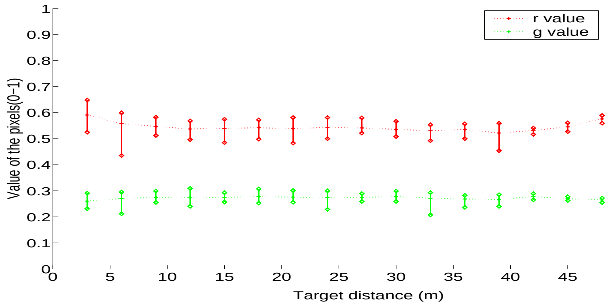

(a) Red.

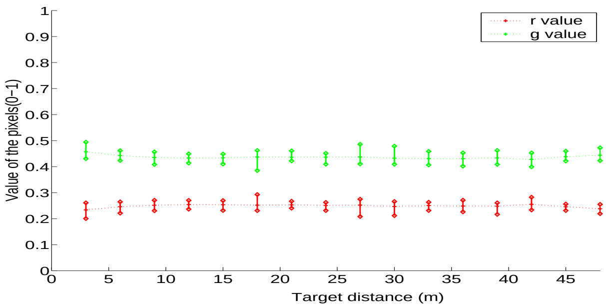

(b) Green.

Figure 8: *Normalised RGB values for the targets with manual white balance.*

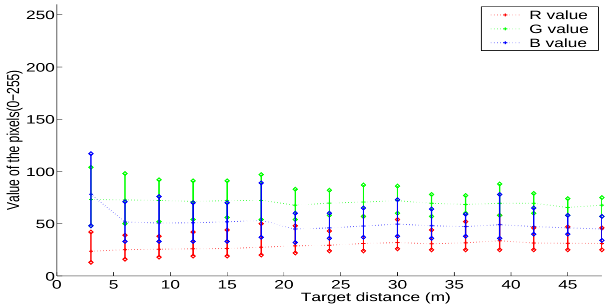

(a) Green.

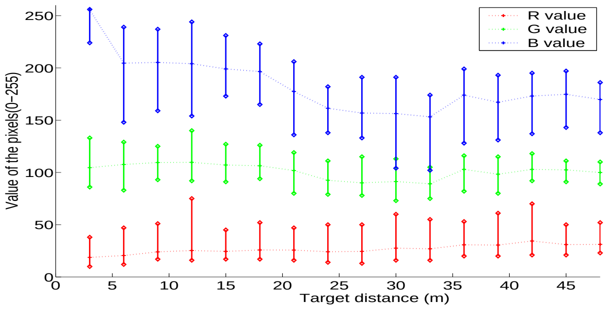

(b) Blue.

Figure 9: *RGB values for the targets from the Marlin.*

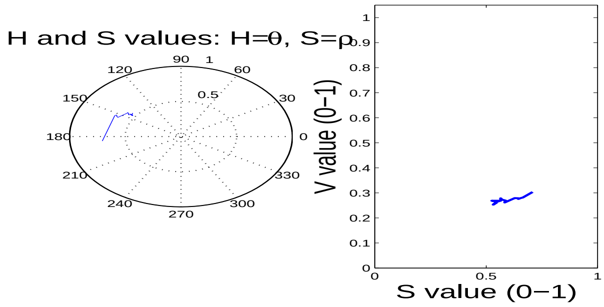

(a) Green.

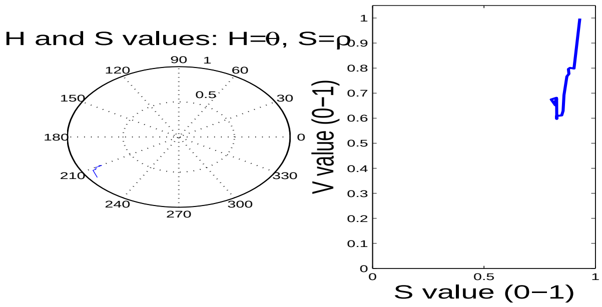

(b) Blue.

Figure 10: *HSV values for the targets from the Marlin.*

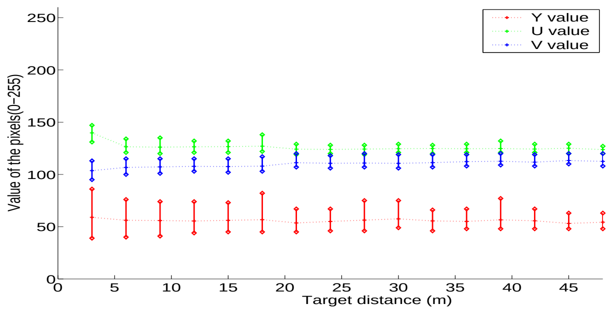

(a) Green.

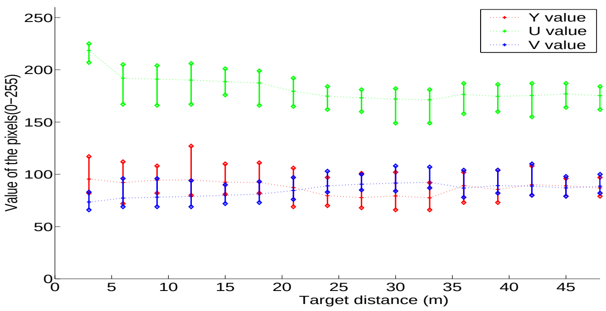

(b) Blue.

Figure 11: *YUV values for the targets from the Marlin.*

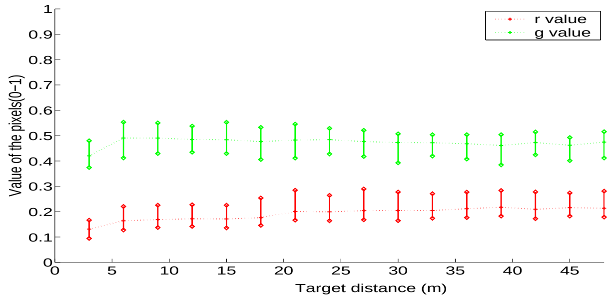

(a) Green.

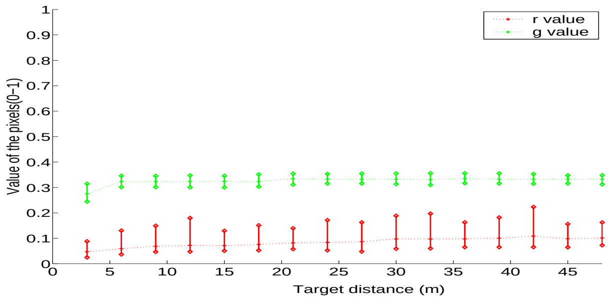

(b) Blue.

Figure 12: *Normalised RGB values for the targets from the Marlin.*

3.5 Colour Shift Between Sunny and Cloudy Days

The purpose of the analysis in this section is to determine the consistency of the RGB, YUV and normalised RGB colour spaces over days of collected data. Figures 13 to 15 shows the averages and standard deviations of selected targets over three cloudy days and two sunny days. It is clear that the normalised RGB space provides the most consistency between days of the same type and days of different types.

3.6 Conclusions from the Analyses

The main objective of the analyses in this section was to determine the best camera configuration and colour space that provide consistent colour under varying distances and illumination conditions outdoors, as objectively as possible. This was determined with two different cameras, on sunny and cloudy days with a limited set of coloured targets. More conclusive evaluation would occur using a larger range of target colours, cameras, and camera settings. However, the analyses highlighted the required camera parameters, and stability of colours in the target environment.

The normalised RGB space was the highest and most consistent performer over the tests, indicating it is less sensitive to environment changes than the other colour spaces. [Khan and Reinhard, 2005] found a similar result in their analysis of 11 colour spaces for their discrimination abilities to determine shadow and reflectance edges on different coloured targets in outdoor images. One of the conclusions of their analysis was that normalised RGB space had the

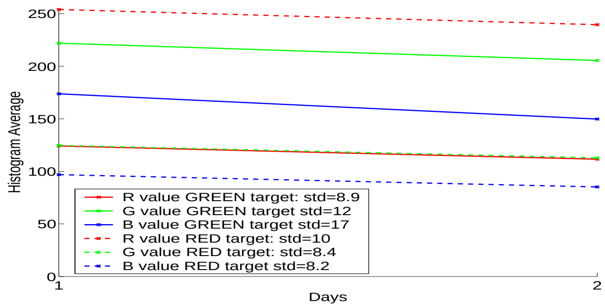

(a) Sunny.

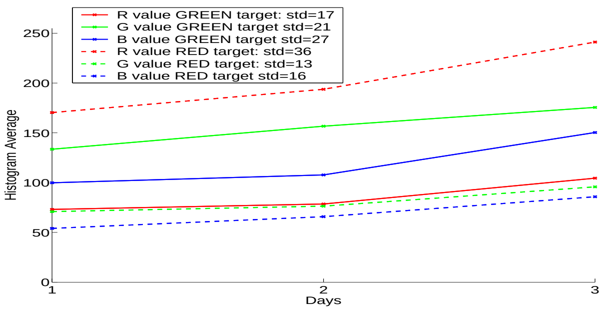

(b) Cloudy.

Figure 13: *The mean of RGB values between days.*

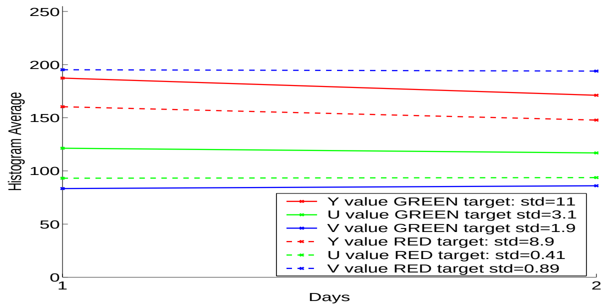

(a) Sunny.

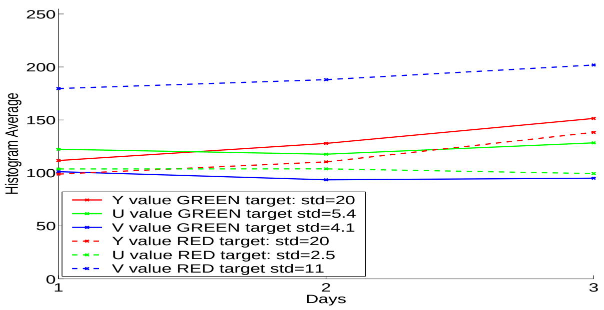

(b) Cloudy.

Figure 14: *The mean of YUV values between days.*

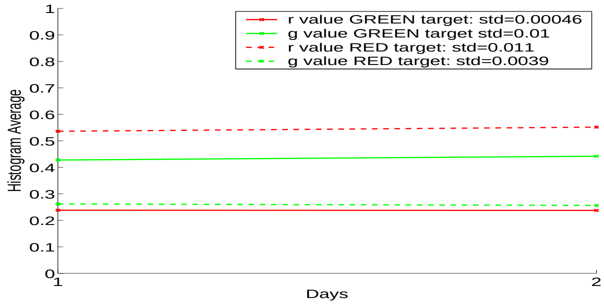

(a) Sunny.

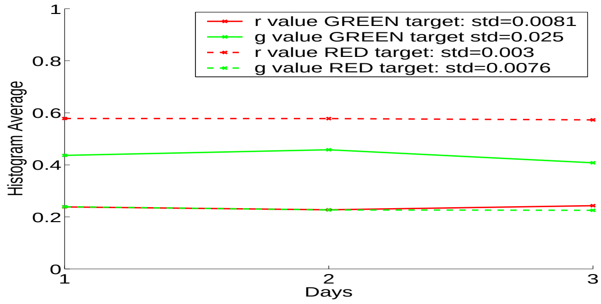

(b) Cloudy.

Figure 15: *The mean of normalised RGB values between days.*

lowest discrimination ability. In normalised RGB space, we have found that:

- the target colours are relatively constant over varying distances,

- the target colours are constant over similar days (e.g. cloudy),

- between sunny and cloudy days, the colour shift for each target is minimal,

- the Unibrain camera's automatic exposure and white balance control are acceptable for producing consistent colours over varying environment conditions, and

- between the cameras tested, the normalised RGB space produces similar results.

4 Detecting Coloured Beacons

The first step in coloured beacon localisation is to determine the underlying system that will provide consistent beacon colours over large distances. The second step is to develop a beacon detection system that provides range and bearing information. An intermediate step is to determine target colours that provide unique signatures in the environment. The third step is to use the detection system information to triangulate the robot's location. The previous section provides the first step with this section providing the intermediate and second steps.

For the intermediate step of selecting unique colours, the features and objects in the environment must also be considered since they can produce false beacons. This occurs in our worksite from coloured barrels, equipment,

Figure 16: *A coloured beacon consisting of dark blue and pink targets.*

buildings and the sky. To reduce false detections, we use two different coloured targets positioned vertically to form a beacon as shown in Figure 16. Each coloured target tested previously was a candidate and by viewing the targets against various backgrounds in normalised RGB space, dark blue and pink proved to be the most unique in sunny, cloudy and shadowy conditions. Their profiles are shown in Figure 17 and Figure 18 respectively.

The detection algorithm is based on finding target-coloured blobs in each normalised RGB image and evaluating their proximity to each other. Initially, the target colours are seeded online by the user selecting a point on each target's image. A six pixel by six pixel window is sampled around the point to determine the average pixel values as the target's seed. Upper and lower thresholds are initially set from the analysis results (Section 3.2) and fine tuned empirically. Once these values are set, they are saved and used by the detection algorithm. The algorithm (Algorithm 1) is shown below and uses the notation of target 1 and target 2 for the targets making up the beacon.

The algorithm makes use of various thresholds to allow for aliasing in the image (edge thresholds) and geometric constraints (aspect ratio and coverage). The values for these parameters once set, have not had to be changed.

4.1 Evaluation of the Detection Algorithm

The purpose of the detection system is to provide range and bearing information to detected beacons. Our tests were conducted in a large concreted area measuring approximately 50 m by 30 m surrounded by buildings, barrels and equipment. Beacons were placed approximately halfway along the north, south, east and west borders. The robot conducted traverses from directly in front of each beacon, reversing in a straight line to the opposite side of the area. Traverses have been made on different days at different times of the day under varying sunny, cloudy and shadowy conditions and produce similar results. For ground truthing the range and bearing calculations of the method, retro-reflective beacons were placed under the visual beacons. The onboard SICK LMS was used to provide the range and bearing to the laser beacon which is considered to be highly accurate. However, it can only detect the

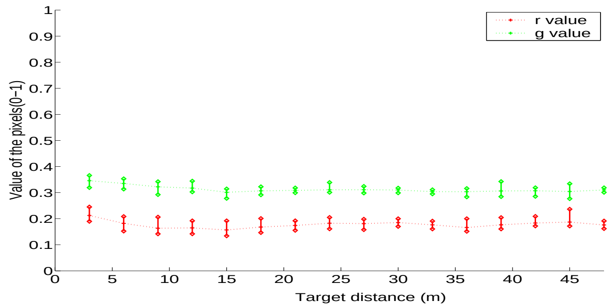

(a) Sunny.

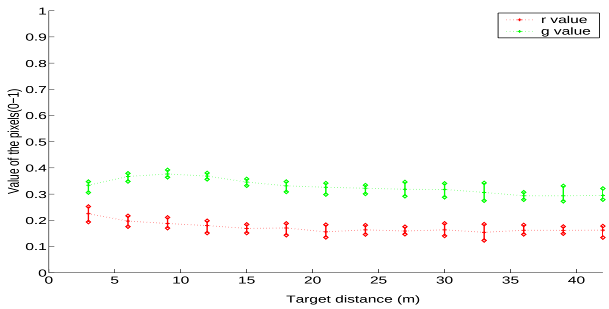

(b) Cloudy.

Figure 17: *Normalised RGB colour profiles for the dark blue target in sunny and cloudy conditions.*

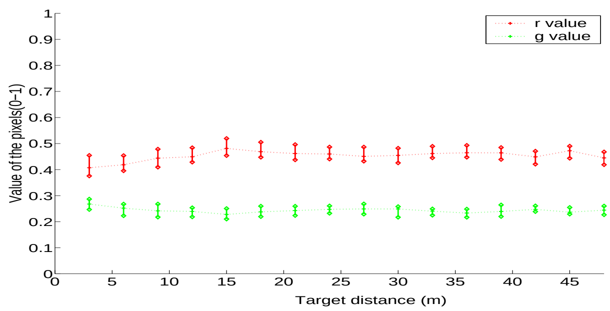

(a) Sunny.

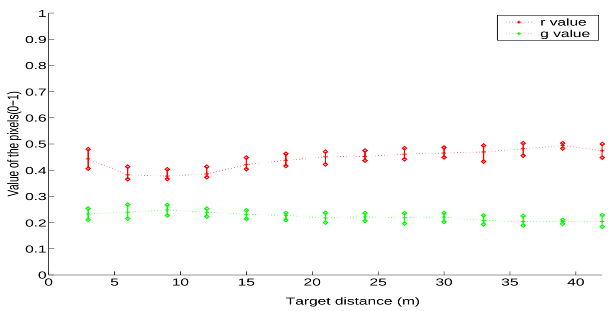

(b) Cloudy.

Figure 18: *Normalised RGB colour profiles for the pink target in sunny and cloudy conditions.*

Algorithm 1 The beacon detection algorithm.

```
using the seed and threshold values, classify every
  pixel as either target 1, target 2 or neither
create blobs from the target 1 and target 2 pixels
compute the dimensions and centroid of each blob
compute the coverage (width/height) and aspect ratio
  (number of pixel in blobs/(width+height)) of each blob
for for each blob of target 1 do
  for each blob of target 2 do
    if each blob's coverage < a coverage threshold
      and each blob's aspect ratio < an aspect ratio
      threshold then
      if the pixel distance between the bottom edge
        of target 1 and the top edge of target 2 < a hor-
        izontal edge threshold then
        if the pixel distance between the leftmost
          edges or rightmost edges < a vertical edge
          threshold then
          /* the blobs constitute the beacon */
          calculate the range from beacon height
          calculate bearing from the vertical centre
            of the beacon and the centre of the image
        end if
      end if
    end if
  end for
end for
```

laser beacon up to 25 m on the intensity channel. Figures 19 and 20 shows the results of two typical trials.

We have found that the range error tends to dip negative or provides a zero-sum average spread after 12 m and is a relatively consistent phenomenon in our tests. Analysing this trend is a subject of future work. Since the range reported by the system is dependant on the height of the detected beacon, as it increases, the physical area covered by a pixel increases. Therefore, variations in the size of the detected beacon at the same range will result in a noisy distance reading. This is apparent by the variations in the range measurements when comparing them to the reported beacon size. However, the bearing error reported is within two degrees over the tested range.

The results in Figures 19 and 20 are indicative of straight approaches to the target. In this case, the beacon is centred in the image and the effect of lens distortion is minimal. As the beacon moves towards the edge of the image, its size and shape will change non-linearly since the edges of the lens are typically more curved than the centre. To determine the effects of this distortion, four traverses were conducted to place the target towards the edge of the image. Each traverse ran parallel to the longitudinal axis of the target, at approximately 2.5 m lateral intervals. At larger lateral distances, the beacon would move towards the edge of the image as the range decreased. Results in Figure 21 show the range error, bearing error and bearing

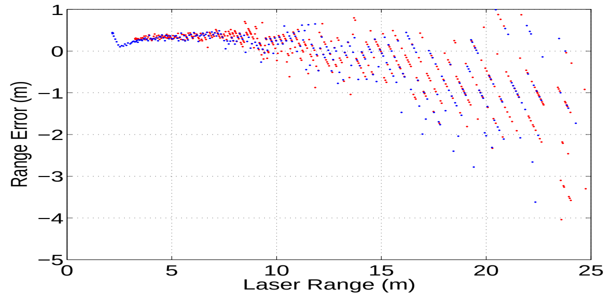

(a) Range errors.

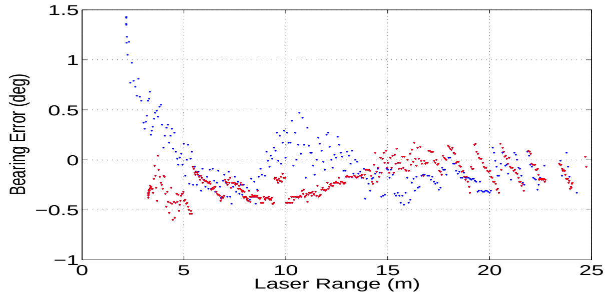

(b) Bearing errors.

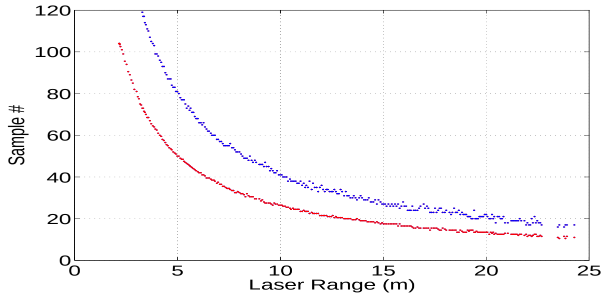

(c) Beacon size.

Figure 19: *Subfigures a and b compare the range and bearing against those reported by the laser for the laser beacon. Subfigure c shows the size of the beacon (height is blue, width is red) in pixels versus distance. Note the correlation between the height variation and range error.*

when compared with the laser beacon readings. Consistent with approaching the beacon head on, the range error shows a negative trend at increasing distances from approximately 12 m. Contrasting to this, the bearing error decreases with distance but increases with bearing angle as a function of the lens distortion. Also apparent is the 21 degree angular limit of the lens in Figure 21c.

5 Conclusions and Future Work

The main conclusion to this research is that the normalised RGB colour space provides better colour consistency for our application under varying natural outdoor lighting and target ranges than HSV, YUV and RGB. The colours of the beacons need to be determined to be unique for the environment. In practice, this is difficult since en-

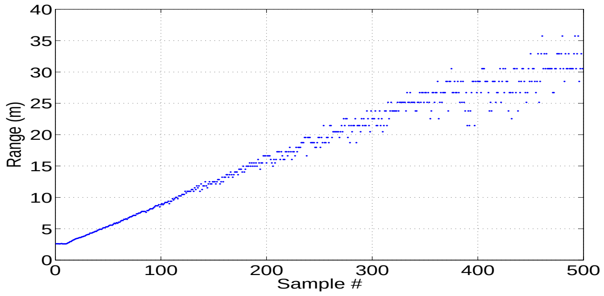

Figure 20: *This figure shows the maximum range reported by the detection algorithm.*

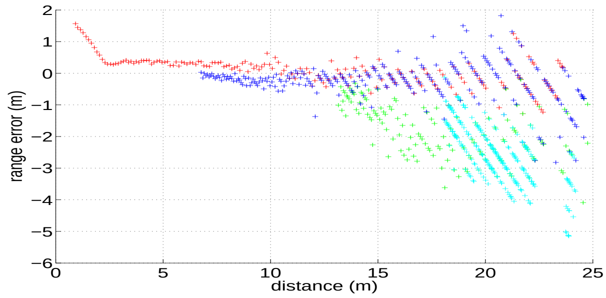

(a) Range Errors.

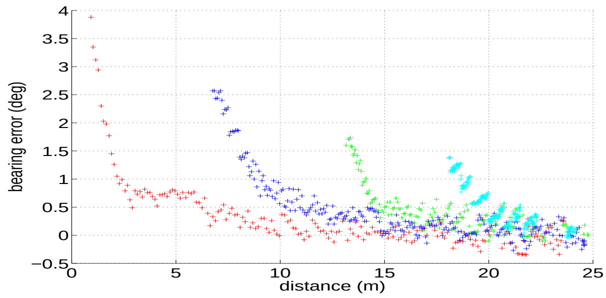

(b) Bearing Errors.

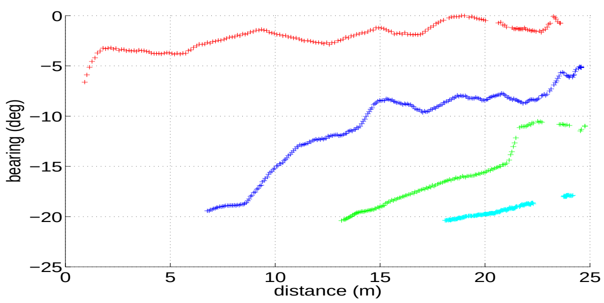

(c) Bearings.

Figure 21: *Results comparing the visual beacon detector to the laser beacon detector. Traverses were undertaken at 0m (red), 2.5m (dark blue), 5m (green) and 7m (light blue) laterally parallel to the longitudinal axis of the beacon. Subfigure 'a' shows the range differences, 'b' shows the bearing differences and 'c' shows the bearings taken from the laser rangefinder.*

vironment changes may produce these colours. To overcome this limitation, a combination of colours can be used with constraints on their proximity and geometry. We have found this technique to be robust in uniquely identifying beacons, in conjunction with thresholding seed values for each target's colour. Over a week of daily testing at different times and illumination (bright morning sun, large building shadows late in the afternoon, and cloudy), we have not adjusted the parameters and are able to detect beacons beyond 35 m which is greater than our laser-based beacon system can manage.

Although beacon detection has proven robust and consistent, the range and bearing calculations are noisy. This part of the system is still at the preliminary stage and we have yet to evaluate the approach more thoroughly to determine its limitations and potential accuracy. Subpixel resampling will be investigated to reduce the noise. For the trend of the range error, initially a more throrough analysis will be conducted to determine if this is a systematic manifestation of our approach. If not, it can be modelled to remove the errors in the beacon detection calculations. Once these errors have been dealt with, the system will be a valuable utility for colour beacon localisation for mobile robots.

References

[Austin and Barnes, 2003] David Austin and Nick Barnes. Red is the new black - or is it? In *Australasian Conference on Robotics and Automation*, Brisbane, Australia, 2003.

[Browning and Veloso, 2005] Brett Browning and Manuela Veloso. Real-time, adaptive color-based robot vision. In *International Conference on Robotics and Automation*, page in press, Barcelona, Spain, 2005.

[Buluswar, 2002] Shashi Buluswar. Color-based models for outdoor machine vision. In *PhD Thesis*, University of Massachusetts, Amherst, 2002.

[Gevers and Stokman, 2005] Theo Gevers and Harro Stokman. Robust histogram construction from color invariants for object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(1):113–118, 2005.

[Judd *et al.*, 1964] D. Judd, D. MacAdam, and G. Wyszecki. Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America*, 54(8):1031–1040, 1964.

[Khan and Reinhard, 2005] Erum Khan and Erik Reinhard. Evaluation of color spaces for edge classification in outdoor scenes. In *International Conference on Image Processing*, 2005.

[Stokman and Gevers, 2005] Harro Stokman and Theo Gevers. Selection and fusion of color models for feature detection. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, 2005.

[Todt and Torras, 2001] E. Todt and C. Torras. Color constancy for landmark detection in outdoor environments. In *4th European Workshop on Advanced Mobile Robots (Eurobot'01)*, pages 75–82, 2001.