A Novel and Practical Approach towards Color Constancy for Mobile Robots using Overlapping Color Space Signatures

Christopher Stanton and Mary-Anne Williams

Innovation and Technology Research Laboratory
Faculty of Information Technology
University of Technology Sydney
{cstanton, Mary-Anne}@it.uts.edu.au

Abstract. Color constancy is the ability to correctly perceive an object's color regardless of illumination. Within the controlled, color-coded environments in which many robots operate (such as RoboCup), engineers have been able to avoid the color constancy problem by using straightforward mappings of pixel values to symbolic colors. However, for robots to perform color vision tasks under natural light the color constancy problem must be addressed. We have developed a color vision system which allows for the color space signatures of different symbolic colors to overlap. This raises the question: if a specific pixel value can be mapped to multiple symbolic colors, how does the robot determine which color is the "correct" one? Context plays an important role. We adopt a knowledge driven approach which allows the robot to reason about uncertain color values. The system is fully implemented on a Sony AIBO.

1 Introduction

Within the color-coded world of RoboCup[1], most teams use color image segmentation techniques to assist with the identification of relevant objects, such as the ball, goals, landmarks and other robots. Generally, the vision systems developed for RoboCup take advantage of the engineered RoboCup environment – where the lighting is bright, evenly dispersed and constant, and the colors of important objects highly distinct – by providing a one-to-one mapping between pixel values and symbolic colors, i.e. for any given raw pixel value there is a maximum of one corresponding symbolic color class.

Such systems ignore the reality that even within the controlled RoboCup environment the colors of important objects do indeed overlap within the color space. For example, many teams within the legged-league are familiar with the problem that when in shadow or dim light, the orange of the soccer ball can appear to be the same color as the red uniforms worn by one team of robots. Such color misclassifications can have dire consequences for system performance, as witnessed by many a robot

[1] http://www.robocup.org

chasing the red uniform worn by a fellow robot, guided by the misconstrued belief that the red robot is in fact the orange soccer ball.

Color constancy is the ability to perceive an object's color correctly regardless of illumination [1]. To help overcome this problem, we have developed a color vision system which allows for a pixel's raw data value to be mapped to a set of many possible symbolic color classes. However, introducing such a relationship raises the question: if a specific pixel value can belong to multiple symbolic colors, how does the robot determine which color is the "correct" one? In this paper we detail our approach which allows the robot to reason about uncertain color values. The system is completely implemented on a Sony AIBO[2], and an initial version was used with great success at RoboCup 2004.

2 Color Constancy and Mobile Robotics

The appearance of a surface's color is dependent upon the complex interaction of a number of factors, including the reflectance properties of the surface, the camera, and the lighting conditions. A change in any of these factors can affect an object's apparent color. Illumination is rarely constant under natural light. Even when lighting is relatively constant, the viewing geometry for mobile robots is not. If we consider the legged-league of RoboCup, in which teams of Sony AIBOs play soccer, the viewing geometry is consistently shifting as the robot's camera is located in, and moves with, the robot's head. For example, the robot's own head often casts shadows over the ball.

Mobile robots must also compensate for imperfect and noisy sensors. For example, the camera on the AIBO ERS7 has a "fisheye" effect which produces a blue discoloration around the edge of the camera. The camera also discriminates dark colors poorly, making it difficult to distinguish between colors such as the dark grey skin of an AIBO ERS210, the black pants of a referee, field green in shadow, or the blue team's uniform. Also, the fact that a mobile robot is indeed mobile can affect camera performance. In robotic soccer robots frequently collide, and for legged robots there is an element of "bounce" when the robots walk. Motions such as these can cause color distortion and blur within the images captured by the robot's camera. Finally, any solution must be capable of operating in real-time within the limited computational resources provided by the robot's hardware.

3 Prior Work

There is an enormous body of literature regarding computational color constancy. However, the vast majority of this research focuses on static images, database image retrieval, and off-board processing. The general aim of computational color constancy is to determine the effect of the unknown illuminant(s) in a scene, and then to correct

---

[2] http://www.sony.net/Products/aibo/

the image by either mapping to an illumination invariant representation, or by correcting for the illuminant.

In terms of color constancy applied to mobile robots there is a much smaller body of knowledge. Many approaches use color insensitive algorithms to assist with object or color recognition, so that once an object or color is recognized, the robot can survey the pixel values within the image, and then use these values to update color tables dynamically, e.g. [2], [3], [4], [5]. Another method is to use image statistics, either from a single image or a series of images, to determine a global scene illuminant - the rationale being if lighting conditions can be accurately classified, then an appropriate color mapping between raw pixel values and symbolic colors can be selected by the robot, e.g. [6] [7]. Lastly, an alternative approach is to pre-process an image to improve the separation of colors within the color space so that symbolic classes are more tightly clustered around a central point, e.g. [8].

The area of specific concern in this paper is determining symbolic color class membership in robotic color vision tasks when the symbolic colors have substantial overlap within a color space, even when the lighting conditions are relatively constant. Within RoboCup most teams avoid the problem by adding controls to their color calibration process which govern overlap and outliers for symbolic colors within the color space (e.g. [9], [10]). Mayer *et al*. [11] reported that when playing middle-size league soccer under natural light they experienced substantial overlap between white and other symbolic colors. Their unsatisfactory solution was to simply give “priority” to colors other than white. A common problem within the legged-league is that when looking down at the orange ball it can appear red, and one team [12] tried to compensate for this by building two color tables – one for when the robot is looking down, and one for all other situations. However, they report mixed success, as orange tends to merge with red and having two color tables did not solve their problems of color misclassification. In the most related approach to our work, [13] report on initial attempts to identify overlapping color signatures within the color space. They describe pixel values for which no overlap exists as “core” colors, and pixel values for which overlap exists as “maybe” colors.

4 Our Approach

Rather than focusing on building a color constancy system that can overcome drastic lighting changes through mathematical calculations of the scene illuminant, we have focused our efforts on developing a vision system which can provide an expressive representation for reasoning about the uncertainty of colors. We are motivated by our longer term aim of allowing the robots to reason about the color of pixels and objects using their knowledge about the environment, such as lighting conditions, camera, and prior experiences.

The first step of our approach to the color constancy problem is to identify the pixel values of different symbolic colors that overlap in color space, and instead of removing or ignoring these particular pixel values, we provide the robot with the complete set of possible candidate colors for any given pixel value. Importantly, this reduces the search space for classifying pixel values whose color signatures overlap.

Secondly, we created a symbolic color class called "dark noise" to capture dark areas within the image in which much color overlap occurs, such as shadow. Next, the color classification algorithm assigns each pixel a value which indicates the set of possible colors for that pixel. For pixel values with more than one possible color, color classification relies upon local image area statistics, the pose of the robot, and other heuristic based knowledge.

4.1 Image Sampling and Training

We use a color labeling process in which a human trainer labels regions within the image that correspond to the objects of interest, such as the ball, field, robot uniforms, landmarks, and so forth. A custom built software system, using a relational database, stores every unique raw pixel value "*p*" that the user selects for every symbolic color "*c*". Thus, given a set of symbolic colors, e.g. C = {*white, green, pink, orange …*}, it is possible for any pixel value to be a member of an arbitrarily assigned subset of C, depending upon the pixel to symbolic color relationships identified by the human trainer. We call this subset of C the *candidate colors* for a pixel value. While many pixel values will share the same symbolic color relationships, and hence candidate colors, e.g. $p_1 \equiv p_2$, invariably a large proportion of pixel values will have different symbolic color relationships. For example, $p_1 \in$ {*green, robot blue, beacon blue*}, $p_2 \in$ {*orange, red*}, $p_3 \in$ {*orange*}, and so forth. In accordance with the terminology used in [18], we call pixel values for which there is only one candidate color "core-colors" (e.g. $p \in$ {*orange*}), and pixel values for which there are multiple candidate colors "maybe-colors" (e.g. $p \in$ {*orange, red*}). In other words, a core color is a pixel value for which there exists no overlap within the color space – they have only ever been assigned to one symbolic color - while a "maybe-color" is a pixel value which has been assigned to two or more symbolic colors.

At any point during the training process, the user can generate three artifacts that are required by the robot's vision system:

1. A structured file containing the complete set of unique candidate color combinations, with each combination possessing a unique index for the purposes of identification.
2. A color lookup table, which for every possible raw pixel value, provides an index to the corresponding set of candidate colors.
3. A file containing the mean (prototypical) value for each symbolic color in terms of raw pixel value.

4.2 Color Labeling and Image Segmentation

Our color calibration system has provided our robots with a more detailed level of color perception. In previous systems a particular pixel value was either unknown, or it belonged to a specific symbolic color. Now, a pixel value can be either unknown, belong to one specific symbolic color, or belong to a specific subset of the entire spectrum of symbolic colors. Thus, for many pixels within the image we are forced to make a new decision: which color is the correct one? To answer this question we trialed a variety of simple and efficient computational techniques, all of which can operate in real-time on both an ERS7, as well as the older and more computationally challenging ERS210.

The algorithm which provided best results for varying lighting conditions within our research laboratory was surprisingly simple. The algorithm takes advantage of the distinction between core colors and maybe-colors, by treating core colors as influential local area predictors for maybe-colors. For example, if a maybe color pixel could be either red or orange, but is surrounded by more orange core colors than red core colors, then it will be assigned the color orange. By only considering candidate colors, and not the complete set of colors, we are able to reduce the search space, increase the speed of the algorithm, and provide surprisingly natural results. In the absence of candidate core colors within a local area of the image (which can occur in images in which there are large concentrations of maybe-colors), or when there is an equal abundance of different neighboring candidate colors (e.g. 4 red and 4 orange neighboring pixels), Manhattan distance metric is used to find the closest candidate color. However, our aim is not to present, or find, the most sophisticated algorithm for correctly color segmenting an image, but rather to demonstrate how knowledge of relationships between pixel values and overlapping symbolic color signatures is a powerful alternative for overcoming color constancy issues. Code containing the implementation of this algorithm can be obtained from [14].

4.3 Results

Fig. 2 displays a raw image taken from an ERS7, together with images indicating the maybe colors within the image and the final segmented image.

Fig. 1. An image from an AIBO ERS7 *(left)*, the overlapping colors in the corresponding image are represented in purple (*centre*), and the processed image in which overlapping colors are assigned to symbolic colors (*right)*. In the raw image there is a blue discoloration around the edge of the image, and there is little contrast or separation between the robot's blue uniform, shadows on the field, and the darker colors of each robot. The blue uniform consists almost entirely of maybe-colors.

A consistently surprising feature of processed images was the ability to accurately classify the regions of the image which corresponded to shadows on the field, and in some cases also on the robot. While such features are currently not used by our object recognition routines, models of color constancy which involve some level of scene understanding will require robots to detect such features. Fig. 3 displays an image in which two robots almost collide, and the proximity of the two robots causes a decrease in the illu mination within the image.

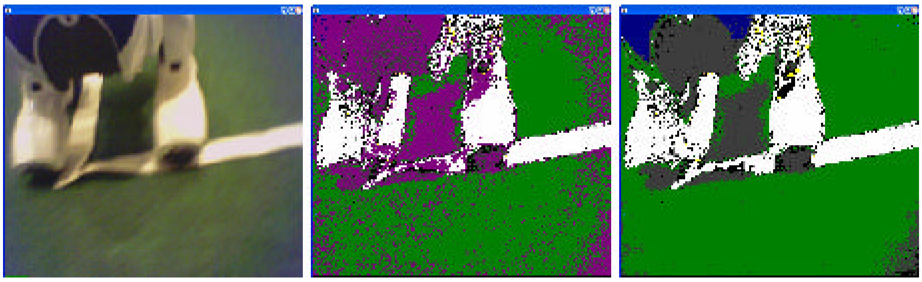

Fig. 2. Raw image from an AIBO ERS7 *(left)*, the overlapping colors in the corresponding image are represented in purple (*centre*), and the processed image in which overlapping colors are assigned to symbolic colors (*right)*. The vision system is able to correctly segment the blue of the robot's uniform, and the area of shadow underneath the robot.

To demonstrate our results we have displayed images that we feel are indicative of the vision system's general performance. It is interesting to note that evaluating the performance of a color constancy system empirically is challenging. For any color constancy system there exists no automated method for recording the number or percentage of pixels within each image that are classified "correctly". Such notions of correctness or ground truth must be specified manually by a human tester. In much of the robotic research relating to color constancy systems are evaluated through behavioral performance tests. However, the performance of behaviors is also related to the performance of higher level routines (e.g. object recognition). Evaluating the performance of perception systems is an area of increasing significance for robotics [15].

One method adopted to evaluate the system's robustness was to vary the lighting conditions, and to also change the camera settings of the system. We were able to create color calibration tables that could function over a range of camera settings and lighting conditions. Fig. 4 displays an image in which the camera shutter speed was set at "fast", but the calibration tables used were created when using the "medium" shutter speed (effectively decreasing the brightness within the image).

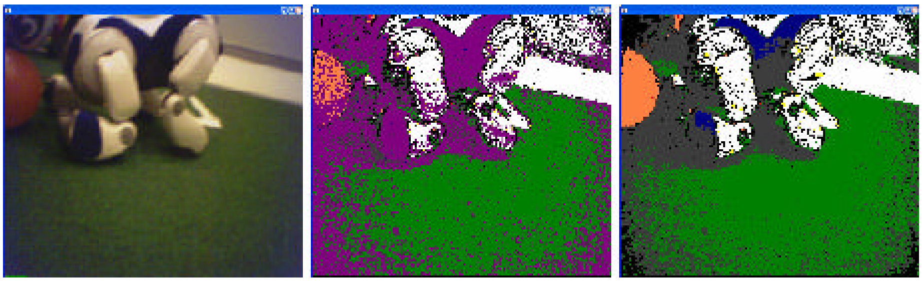

Fig. 3. Raw image from an AIBO ERS7 taken at fast shutter speed *(left).* The color tables used to segment the image were created at medium shutter speed. Due to the darker conditions an increased amount of overlap colors were present in the image (*centre*). Large parts of the ball overlap with robot red. The processed image effectively segments the ball and most of the robot's uniform (*right)*.

5 Discussion

We have implemented a novel approach to deal with color constancy. Rather than avoiding or removing overlapping color space signatures, we have developed a system which uses the relationships between pixel values and overlapping symbolic color signatures to segment color images.

Our approach offered several immediate benefits. Color calibration can be undertaken more quickly, as the calibration method encourages the human trainer to identify all possible pixel values for each color of interest, rather than avoiding those that may cause misclassification (e.g. those that occur in shadow or on the borders of different objects within the image). Image segmentation has improved due to richer and more expressive color tables. Lastly, object recognition has also improved, due to not only image segmentation performance, but because object recognition routines can reason about the different levels of color uncertainty indicated by core colors, maybe colors, and unknown colors. For example, object recognition routines can exploit simple statistics, such as the percentage of maybe-colors within a blob, to reason about the likelihood of false identification of an object.

Future research will involve developing mechanisms for automatically generating the rules for determining the color membership of overlapping pixel values. When a human trainer labels the colors of pixels within an image, a wealth of contextual knowledge and scene understanding affects our interpretation of a pixel's color. A longer term aim is to investigate color training mechanisms that can embed this knowledge within the robot. For example, the human trainer compensates for the blue discoloration around the edge of the ERS7's image without conscious effort. Thus we are recording features such as the pixel's location in the image which allow us to calculate probabilistic rules for color membership which consider constant distortions of the camera.

References

1. Horn, B. Robot Vision, MIT Press, Cambridge, MA, 1987.
2. Cameron, D., & Barnes, N.: Knowledge-based autonomous dynamic color calibration. In: Robocup 2003, Padua, Italy (2003)
3. Jungel, M.: Using Layered Color Precision for a Self-Calibrating Vision System. In Daniele Nardi, Martin Riedmiller, and Claude Sammut, editors, RoboCup-2004: Robot Soccer World Cup VIII, Springer Verlag, Berlin, 2005.
4. Schulz, D., & Fox, D.: Bayesian Color Estimation for Adaptive Vision-based Robot Localization. IROS-04.
5. Gönner, C., Rous, M., and Kraiss K.: Real-Time Adaptive Colour Segmentation for the RoboCup Middle Size League. In Daniele Nardi, Martin Riedmiller, and Claude Sammut, editors, RoboCup-2004: Robot Soccer World Cup VIII, Springer Verlag, Berlin, 2005.
6. Lenser, S. and Veloso, M.: Automatic detection and response to environmental change. In: Proceedings of the International Conference of Robotics and Automation, May 2003.
7. Sridharan, M. and Stone, P.:Towards Illumination Invariance in the Legged League. In Daniele Nardi, Martin Riedmiller, and Claude Sammut, editors, RoboCup-2004: Robot Soccer World Cup VIII, Springer Verlag, Berlin, 2005.
8. Mayer, G., Utz, H., Kraetzschmar, G.:Towards Autonomous Vision Self-Calibration for Soccer Robots. In Proceedings of the 2002 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems, 2002.
9. German Legged Team Report, 2004. Available at: http://www.tzi.de/4legged/bin/view/Website/Teams2004
10. University of Texas Legged Team Report. 2004. Available at: http://www.tzi.de/4legged/bin/view/Website/Teams2004
11. Mayer, G., Utz,H., Kraetzschma, G.: Playing Robot Soccer under Natural Light: A Case Study. 238-249. In RoboCup 2003.
12. University of Newcastle Legged Team Report. 2004. Available at: http://www.tzi.de/4legged/bin/view/Website/Teams2004
13. University of New South Wales Team Report. 2003. Available at: http://www.cse.unsw.edu.au/~robocup/reports.phtml
14. University of Technology Sydney Team Report and Code Release, 2004. Available at: http:// unleashed.it.uts.edu.au
15. Gardenfors, P., Karol, A., McCarthy, J., Stanton, C, and Williams, M. A Framework for Evaluating Groundedness of Representations in Agents: From Brains in Bottles to Mobile Robots. Workshop on Agents in real-time and dynamic environments, IJCAI 2005.