

Color recognition in outdoor images

Shashi D. Buluswar
Department of Computer Science
University of Massachusetts
Amherst, MA 01003, U.S.A.

Bruce A. Draper
Department of Computer Science
Colorado State University
Ft. Collins, CO 80523, U.S.A.

Abstract

The color associated with an object in machine vision images is not constant; under varying illuminating and viewing conditions (such as in outdoor images), the perceived color of an object can vary significantly, thus making color-based recognition difficult. Existing methods in color-based recognition have been applied mostly to indoor and/or constrained imagery, but not to realistic outdoor data.

This work analyzes the variation of object color in outdoor images with respect to existing models of daylight illumination and surface reflectance. Two approaches for color recognition are then proposed: the first develops context-based models of daylight illumination and hybrid surface reflectance, and predicts the color of objects based on scene context. The second method shows that object color can be nonparametrically "learned" through classification methods such as Neural Networks and Multivariate Decision Trees. The methods have been successfully tested in domains such as road/highway scenes, off-road navigation and military target detection.

1 Introduction

The color (or rather, the *apparent* color) of an object varies with changes in illuminant color, illumination geometry (i.e., angle of incidence), viewing geometry (angle of reflectance), and miscellaneous sensor parameters [11]. In outdoor images, the color of the illuminant (i.e., daylight) varies with the time-of-day, cloud cover and other atmospheric conditions [12]; the illuminant and viewing geometry vary with changes in object and camera position and orientation. In addition, certain sensor response parameters [15], shadows and inter-reflectances [10], all of which can be difficult to model in outdoor scenes, may also affect the apparent color of objects. Consequently, at different times of the day, under different weather conditions, and at various positions and orientations of the object and camera, the apparent color of an object can be different.

Figure 1 (top) shows the variation in the apparent color of two simple matte surfaces (white and green) under different lighting and viewing conditions from about 50 images; the variation for each surface is represented by a cluster in *RGB* space. The overall variation for each surface is about 250% of the distance between the centroids of the two clusters. In other words, the variation in the apparent color of a single




Figure 1: Samples of objects in outdoor images, along with the variation of *RGB* color over 100 images. Top: two matte surfaces; Middle: camouflaged military vehicle; Bottom: shiny traffic sign.

surface can be greater than the difference (in terms of color-space Cartesian distance) between two distinct colors (white and green, in this case). The variation in the apparent color of more realistic objects, such as a road surface and a camouflaged military vehicle (also in figure 1), can be greater.

Much of the work in computational color recognition under varying illumination has been in the area of color constancy, the goal of which is to match object colors under varying illumination without knowing the spectral composition of the incident light or surface reflectance. An illuminant-invariant measure of surface reflectance is recovered by first determining the properties of the illuminant. While there have been many interesting advances in color constancy [7, 8, 16, 21], their applicability to unconstrained outdoor images has not yet been established. Indeed, as Forsyth [8] states, "*Experimental results for [color constancy] algorithms running on real images are not easily found in the literature... Some work exists on the processes which can contribute to real world lightness constancy, but very little progress has been made in this area.*" In the interests of brevity, this paper does not discuss the previous approaches in detail; Funt [9] and Forsyth [8] discuss the existing approaches, along with their

methodology, strengths and weaknesses.

This paper analyzes the variation of the color of objects with respect to existing models of daylight [12] and surface reflectance [17, 14, 21], and proposes two approaches to color recognition based on the analysis. The first approach develops a surface reflectance model for hybrid surfaces and a context-based model of daylight illumination; thereafter, given information about the scene context (sun-angle, sky conditions and surface orientation), the apparent color of the surface is predicted. The second approach shows that the *RGB* distributions representing the apparent color of objects under varying conditions (see figure 1) assume shapes which can be approximated using nonparametric classification schemes such as Neural Networks [19] and Multivariate Decision Trees [2]; these classification methods effectively "learn" an object's color from training samples. The methods have been tested, to various degrees, in real outdoor images from a number of domains, such as road scenes, off-road navigation and military target detection.

There are two color spaces that are used: the classification (second) approach uses the standard *RGB* space; the model-based (second) approach uses the normalized *rgb* color space, where $r = R/(R+G+B)$, $g = G/(R+B+G)$, and $b = B/(R+G+B)$.

2 Color shift in outdoor images: Causes and analysis

The primary causes for the variation in the apparent color of a given object in outdoor images are changes in the illuminant (daylight) color, illumination geometry, viewing geometry and imaging parameters. The effect of illumination geometry and viewing geometry is determined by the surface reflectance properties. Although some models exist for daylight [12], surface reflectance [11, 14, 17, 21] and imaging parameters (from camera specifications), these models have not been applied to realistic outdoor machine imagery.

2.1 Illumination

The variation in the color of daylight is caused by changes in the sun-angle, cloud cover, and other weather conditions. The CIE daylight model [12] shows the extent of the variation in color space as a parabola, defined by the following equation in the CIE chromaticity space (*RGB* is a linear transform of the CIE chromaticity space).

$$y = 2.87x - 3.0x^2 - 0.275, \quad (1)$$

where $0.25 \leq x \leq 0.38$. In *RGB* space, the parabola stretches out into a paraboloid surface [3]. In the normalized *rgb* space, the function remains a parabola.

Although the CIE model has empirically been validated through radiometric measurements, the model has not been used in the context of machine vision to predict the color of daylight in order to determine the color of objects under daylight; Ohta [16], for instance, uses the CIE model as a constraint on artificial illumination.

2.2 Illumination geometry, surface reflection and viewing geometry

Illumination geometry affects the composition of the light incident upon the surface. Daylight has two components, sunlight and (ambient) skylight, and the surface orientation determines how much light from each source is incident on the surface. Surface reflectance (specifically, the relative strengths of the Lambertian and specular components), based on the combined geometry of illumination and viewing, affects the composition and amount of light reflected by the surface onto the camera.

A number of reflectance models exist for realistic surfaces (which have components of both Lambertian and specular reflectance), most notably the Phong shading model [17], Shafer's Dichromatic model [21], and Nayar's hybrid reflection model based on photometric sampling [14].

The Dichromatic reflection model (originally proposed by Shafer [21] and subsequently used and extended by Klinker [13]) models the net reflection of a surface as the linear combination of the specular and Lambertian reflection components.

$$L(\lambda, i, e, g) = m_s(i, e, g)c_s(\lambda) + m_b(i, e, g)c_b(\lambda) \quad (2)$$

where $L(\lambda, i, e, g)$ is the intensity of light at wavelength λ , angle of incidence i , angle of reflection e and phase angle g (angle between direction of incident light and viewing direction); $m_s(i, e, g)$ is the geometric scale factor (determined by the illumination and viewing geometry) of $c_s(\lambda)$, the spectral power distribution of the specular component of the reflected light, and $m_b(i, e, g)$ and $c_b(\lambda)$ are the same quantities for the Lambertian reflection component. Specularities in the image are used to determine the weights for each component. Under this the apparent color of a surface in *RGB*, lies somewhere along a \perp -shaped (inverted T) distribution, where the base of the \perp represents the Lambertian component, and the vertical height represents the specular component.

The Phong shading model [17] approximates the falloff of the brightness of specular reflection as $\cos^n(\alpha)$, where α is the difference between the viewing angle and the angle of maximal specular reflectance. The value of n is determined empirically for a given surface, and varies from 1 (for matte surfaces) to 200 (for highly specular surfaces). At $\alpha = 0$, the brightness is the maximum (i.e., 1), and falls off as the surface is rotated, to the minimum (i.e., 0) at -90° and 90° . The Phong model has been very widely applied by the computer graphics community as an effective method of achieving shading effects in rendering grey-scale and color images.

Nayar [14] describes the brightness of surface reflectance as a linear combination of the Lambertian and specular components (a concept similar to [21]):

$$I = IL + IS, \quad (3)$$

where I is the total intensity at a given point in the surface, and IL and IS the intensities of the specular and Lambertian components. $IL = A\cos(\theta_s - \theta_n)$

(Lambert's law), where A is the constant representing the weight of the Lambertian component, and θ_s and θ_n are the directions of the illumination source and the surface normal. I_S , modeled by the delta function [14], is $B\delta(\theta_s - 2\theta_n)$, where B is the weight of the specular component. Hence, $I = A\cos(\theta_s - \theta_n) + B\delta(\theta_s - 2\theta_n)$. The model is adapted for an extended light source to determine the weights of the reflection components by photometric sampling, a method by which brightness samples are obtained for multiple angles of illumination and viewing.

Each of these models has proven successful in constrained domains, and offers an insightful understanding of reflection. However, their applicability to color recognition in outdoor imagery has not been established. For instance, the Nayar model has been used for shape recovery in grey-scale images, but not in color imagery. The Phong Shading model is used for graphical rendering, but has not been widely used to recover the reflectance of surfaces. The Dichromatic model has been used to recover reflection components in constrained, indoor color images, but not under outdoor illumination (Klinker [13], for instance, uses plastic cups under indoor lighting).

2.3 Shadows, inter-reflections, and imaging parameters

Inter-reflections and shadows can cause a variation in color by altering the color of the light incident upon the surface [10]. Inter-reflections, for instance, cause light reflected off other surfaces in the scene to be incident upon the surface being examined. Shadowing can cause the elimination of incident sunlight (if the surface is self-shadowed), and further inter-reflection (if the surface is shadowed by another surface).

Color calibration studies have shown that a number of imaging parameters can cause a color shift. For instance, wavelength-dependent displacement of light rays by the camera lens, called chromatic aberration, may lead to color mixing and blurring [1]. Nonlinear camera response and digitization errors can skew the ratio of the values in the three color bands (red, green and blue), and the dynamic range of brightness in outdoor scenes accentuates the possibility of blooming and clipping [15]. Unfortunately, the precise affect of these factors on realistic outdoor color images is not known.

As suggested above, color images of outdoor scenes are complicated by a host of difficult-to-model phenomena; as a result, computational color recognition has been difficult and largely unsolved.

3 Color prediction using context-based models

The model-based approach to color recognition uses context-based models of illumination and reflectance to first predict the color of the incident light from an approximate description of the illuminating conditions (sun-angle, cloud cover, etc.) and then predict the apparent color of the surface from the illuminant color and surface orientation.

A context-based illumination model is constructed, indexed by sun-angle (vertical rise), cloud cover (per-




Figure 2: The normalized Dichromatic reflectance model. The apparent color of the surface is a point on the line between the points representing the illuminant and the Lambertian component of the surface.

centage) and sun-visibility (0-1); for each set of conditions, the model predicts two illuminant colors in rgb : one when the surface is facing the sun (which is the same at all angles of orientation as long as the surface faces the sun) and the other when it faces away (which, again, is the same for all angles when the surface faces away from the sun). Table 1 lists some sample entries from the complete table [4].

Sun \angle	Cloud %	Sun vis.	rgb_{sun}		
			r	g	b
0-5	0-20	1	.436	.360	.204
			.288	.302	.410
6-15	0-20	1	.394	.360	.246
			.288	.302	.410

Table 1: Context-based model of daylight.

A context-based reflectance model, known as the Normalized Photometric Function (NPF) is developed based on the three models described above. First, the Dichromatic model is modified for the normalized rgb space (figure 2), in order to avoid the effect of unpredictable intensity variations in outdoor images [4]. In the normalized color space, the apparent color of a surface under a given illuminant lies on a straight line between two points, one of which represents the color of the illuminant, and the other, the color of the Lambertian component of the surface. The greater the strength of the specular component, the closer the apparent color will be to the color of the illuminant (as the illumination angle approaches the viewing angle). Hence, by using photometric sampling in normalized color, a continuous function can be plotted in two dimensions, where the horizontal axis represents the relative viewing angle (i.e., the combination of illumination and viewing angles, which is 0° when the two angles are equal; the range along this axis is -90° to 90°), and the vertical axis represents the relative (on a scale of 0-1) distance in normalized color space between the points representing the colors of the illuminant and the Lambertian component (Figure 3). In this formulation, Lambertian reflectance forms a straight line at vertical position 0 for all relative viewing angles, and a highly specular surface forms a symmetric function that is at 0 for almost all relative viewing angles, but abruptly jumps to 1 as the angle approaches 0° . The relative viewing angle at which specular reflectance makes the abrupt jump depends on the sharpness of




Figure 3: Normalized Photometric Functions (left to right): hypothesized for narrow illuminant, for a matte surface and for a shiny surface under sunlight.

the illuminant — for a pure point-source light, the spike is very narrow at 0° , and for the sun (which can be modeled as a thin Gaussian [20]) the spike starts ranges from about -15° to 15° . Figure 3 (middle and right) also shows the NPF's for the matte paper and shiny traffic sign shown in figure 1.

Specular peaks are caused only when the illuminant is a point-source or sharp (e.g., like the sun); in other words, when a surface is illuminated by an extended source (such as ambient skylight), then no specular reflection is caused, regardless of the surface reflectance. The shiny street sign, for instance, has no specular peak under skylight, and the corresponding NPF is flat, at a vertical position 0. To compensate for this effect, a scale function is required for the illuminant, which flattens the NPF appropriately; for the context-based illuminant model shown in table 1, the sun visibility factor serves this purpose — the factor is 1 when the sun is completely visible, and 0 if the sun is covered by clouds or if the surface is facing away from the sun. This study has found that four samples of a surface are sufficient for constructing the NPF for the surface, one when the surface is facing the sun at a relative viewing angle of 0° and the three more at increments of 15° [4].

The NPF model can be used to predict the color of a surface given the sun-angle, sky conditions and the relative surface orientation. The rgb value of the illuminant is first determined from the context-based illumination table. The next step is to determine the color of the Lambertian component of the surface; this is done using the 9-parameter color-coefficient model of Lambertian reflectance [4]:

$$\begin{bmatrix} r \\ g \\ b \end{bmatrix} = \begin{bmatrix} W_{rr} & W_{rg} & W_{rb} \\ W_{gr} & W_{gg} & W_{gb} \\ W_{br} & W_{bg} & W_{bb} \end{bmatrix} \begin{bmatrix} r_i \\ g_i \\ b_i \end{bmatrix} \quad (4)$$

where $[r, g, b]$ is the apparent color of the Lambertian surface reflecting an illuminant of color $[r_i, g_i, b_i]$. The 3×3 matrix contains the color coefficients; W_{rr} , W_{rg} and W_{rb} are the coefficients (albedo) of the red component of the surface under the red, green and blue bands of the illuminant, and so on. The coefficient matrix for the Lambertian component of a surface can be calculated from three samples of the surface under different (non-specular) lighting conditions [3].

Once the rgb value of the Lambertian component has been calculated from the illuminant color and the

coefficient matrix for the surface, the relative orientation can be used (with the NPF of the surface) to determine where along the line (from figure 2) the apparent color of the surface will lie. This approach was tested on 90 images of various objects in outdoor images, such as the matte surface and traffic sign from the previous example, concrete and road slabs; the illuminating conditions ranged from clear and sunny to overcast. On the average, about 89% of the target pixels were correctly predicted; after grouping the pixels predicted as potentially belonging to the target surface, 85 out of 90 instances of the target surface were detected in the images, with 39 false positives. The false positives were mostly from “neutral” colored surfaces like the concrete slab [3] (see figure 5, first and second columns).

4 Color recognition by nonparametric classification

Scene context may not always be available, as in the case of pre-existing images databases; the second approach presented in this paper shows color recognition can be achieved in the absence of models or context. The models described in the previous section suggest the following hypothesis about the overall RGB of apparent surface color in outdoor images: under daylight, the RGB distribution representing a matte surface (or a surface with a dominant Lambertian component, such as the top and middle examples in figure 1) will lie in a single continuous region, and the distribution for a specular surface (or a surface with a dominant specular component, such as the one shown at the bottom of figure 1) will be divided into two clusters, one near the Lambertian component, and the other near the color of the illuminant. If we take into account the effect of shadows, inter-reflections and the sensor-induced distortions described above, we can conclude that the RGB distribution representing a surface with unknown reflectance under arbitrary conditions can be complicated.

The nonparametric approach to color recognition assumes that no models of illumination, reflectance or any of the other processes are available; rather, the color of an object is represented by the RGB distribution of the object under various conditions. To recognize the color of the object, the distribution (i.e, the boundary of the region in RGB) is “learned” from samples of the object (positive training samples) and the background (negative training samples) under various conditions (this is known as the training phase). Thereafter, to recognize the color of an object, image pixels are classified with respect to the RGB boundaries associated with the object (the classification phase).

A number of nonparametric classifiers have been used in various problems in machine vision, such as minimum distance classifiers for color-based road-detection [6] and neural networks for road-following [18]; however, none of these methods was intended for the purpose of generic color recognition.

In choosing a classifier, the primary issue to consider is the ability of the technique to generalize the




Figure 4: Recursive discriminants of an MDT, separating the '+'s (positive samples) from the '-'s (negative samples).

function so as to adequately represent the distribution in color space without being too loose (resulting in the inclusion of samples from a different class), or too tight (resulting in the exclusion of samples belonging to the class). Based on this criterion, Neural Networks (with backpropagation) [19] and Multivariate Decision Trees [2] were chosen. Of these two techniques, NN's are well-known, while MDT's are relatively new; consequently, a brief description of MDT's follows.

4.1 Multivariate Decision Trees (MDT's)

Multivariate Decision Trees [2] create piecewise-linear approximations of regions in feature space by recursively dividing feature space with hyper-planes (figure 4).

MDT's recursively subdivide the feature space by linear threshold units (LTU's), which are binary tests represented by linear combinations of feature values and associated weights. Each division attempts to separate, in a set of known instances (the training set), target instances from non-targets. If the two resulting subsets are linearly separable, a single LTU will separate them and the multivariate decision tree consists of the single node. If not, the LTU linearly divides the feature space so as to separate the instances to the extent possible, and the MDT creates and trains new LTU's on the two divisions of the instances. The result is a tree of LTU's recursively dividing the feature space so as to perform a piecewise linear approximation of the region in color-space consisting of the positive samples. The terminal nodes in the tree correspond to inseparable sets, which are labeled as individual classes. Thus, each node in a decision tree is either a decision or a class. Each LTU is fitted using the Recursive Least Squares (RLS) algorithm, which minimizes the mean squared error between the estimated \bar{y}_i and true y_i values ($\sum(y_i - \bar{y}_i)^2$) of the selected features over a number of training instances. RLS incrementally updates the weight vector W according to

$$W_k = W_{k-1} - K_k(X_k^T W_{k-1} - y_k) \quad (5)$$

where W_k is the weight vector for the instance k , of size n , W_{k-1} is the weight vector for instance $k-1$, X_k is the instance vector; X_k^T is X_k transposed, and y_k is the class of the instance. $K_k = P_k X_k$, where P_k is the $n \times n$ covariance matrix for instance k , reflecting the uncertainty in the weights, and

$$P_k = P_{k-1} - P_{k-1} X_k [1 + X_k^T P_{k-1} X_k]^{-1} X_k^T P_{k-1} \quad (6)$$




Figure 5: Left to right: Sample images of three objects (on easel, circled where not obvious) — matte paper, traffic sign (with specular reflection) and concrete slab; results from model-based prediction, MDT's and NN's.

The weights are initialized randomly, and the matrix consists of 0 values everywhere except along the diagonal, which is set to a very large value. If at any level, the LTU results in a non-negative value, the corresponding set of pixels is labeled as belonging to the object (target), otherwise, it is labeled negative (non-target).

4.2 Comparison: MDT's and NN's

Comparative tests were conducted on about 100 images of surfaces such as matte paper, traffic signs, and concrete slabs. The number of training images varied from a minimum of 10 up to a maximum of 50; the tests used the images in the set of 100 that were not used for training (cross-validation). The difference in performance between the MDT and the NN was negligible; in other words, the result of pixel classification using the two systems were virtually identical. Figure 5 shows the representative classification results using decision trees and neural nets. The percentage of target pixels correctly classified by the MDT was about 97% for the matte paper, 94% for the traffic sign, and 96% for the concrete slab; the corresponding numbers for the NN were 97%, 96%, and 94%. The percentage of false-positives (i.e., percentage of pixels classified as targets that were false alarms) for the MDT were 1% for the matte paper, 5% for the traffic sign, and 12% for the concrete slab; the corresponding values for the NN were 1%, 4%, and 18%.

Although there appears to be no significant difference between the two classification techniques, there are a few issues in the context of color-based classification that favor decision trees. To begin with, the MDT hyperplanes in RGB space are much more intuitive for analysis than are the logistic functions of the hidden layer of the neural network. In addition, the notion of topology and contiguity with respect to distributions in RGB is preserved with the MDT approach; with a neural net, contiguous regions in RGB



Figure 6: Representative results for MDT-based classification for lane-markers and obstacles: Original image (left), classification for lane-markers (middle), classification for road vs. non-road (right).

may no longer remain contiguous in the hidden layer space. In theory, this could pose a problem if training samples from one part of a contiguous distribution are not provided.

4.3 Applications: highways, off-road, and military scenarios

Implementations of MDT-based classification have been tested in several domains, such as automated highway systems, off-road obstacle-detection, military target detection, wildlife detection from aerial imagery and skin finding [5]. Results from the highway and target detection tests are discussed below. The following tests were conducted using cross-validation, where half (or fewer) of the images were used in training, and the others for testing.

MDT-based classification is currently being used in the National Automated Highway System (AHS) project for detecting lanes and obstacles in highway scenes for autonomous vehicles. Figure 6 shows a sample image from a highway scene; the goal in this application is to find the lane-markers and obstacles (i.e., non-road pixels). On 10 sequences of 100 images each (i.e., a total of 1000 images) of highway scenes, about 83% of the lane-marker pixels were correctly classified, and about 64% of the obstacle pixels were correctly classified. The false positive classification percentage was less than 2% for lane markers and about 14% for obstacles. By grouping pixels in the binary images, 100% of the lane-markers, and 1480 out of 1497 obstacles were detected. Obstacles included vehicles on the highway in the current and adjacent lanes, as well as miscellaneous objects cluttering the highway. The obstacles that were not detected were portions of black rubber tires that were almost the color of the highway tarmac.

Like highway obstacle-detection systems, many off-road systems use depth-from-stereo for generating an obstacle map. A system based on depth estimation is expected to mark all objects above the ground plane as obstacles. In off-road navigation tests in Colorado, most of the objects marked as obstacles by such a system were actually yucca bushes, which the vehicle can drive over without risk of damage. In such a scenario, MDT-based classification can be used to detect yucca bushes and eliminate them from the obstacle map. Out of 45 test images, approximately 84% of the yucca pixels were correctly detected, with about 29% of the pixels identified as yucca being misclassified (almost all of these pixels were from the grass). Out of




Figure 7: Results from MDT-based classification for yucca bushes. Left to right: original image (actual rocks marked with circles), simulated depth-based obstacle map, classification for yucca, final obstacle map.




Figure 8: Results from MDT-based classification for camouflaged target detection - original color images (left, targets marked with circles), binary classification (right).

a total of 212 yucca bushes, 176 were successfully detected. (The many false positives at the pixel level from grassy regions did not affect the performance of the system because they were not in the initial obstacle map.) Figure 7 shows results from one image with a simulated obstacle map; the yucca bushes (pixels) are detected, and those pixels are removed from the obstacle map, leaving the only rocks in the final obstacle map.

The most challenging and comprehensive application of the MDT-based classification method has been in domain of camouflaged military target detection using autonomous vehicles. This task is particularly difficult because the goal of camouflage patterns and colors on military targets is precisely to blend the targets into the background vegetation. However, it is not always possible to get a perfect match between the background color and camouflage because the color of vegetation is not constant. Consequently, the hyperplanes of the MDT can make fine distinctions between target color and the background. The MDT-based system was tested on tests sanctioned by the US Dept. of Defense. Figure 8 shows the results from two color images from the Ft. Carson images. Targets are extracted from the binary classification image by clustering target pixels and applying region-level heuristics such as the expected vehicle size and aspect ratio. Out of 96 images, 176 out of 211 targets were detected, along with 180 false positives.

5 Discussion and conclusions

The nonparametric approach has proven successful in several applications, and has been implemented in several real-time systems. The model-based approach, on the other hand is more recent, and is yet

to be tested as thoroughly; however, this approach offers more insight into the physics of the formation of apparent color. Model-based prediction is deterministic (i.e., a surface of a given reflectance under a given set of conditions will appear in a predictable fashion), whereas nonparametric classification is more ambiguous (because the *RGB* regions for objects of two different colors may overlap, there is a possibility of false positives, as the discussed results show). This paper has been compressed to fit the format of the proceedings; longer descriptions of the two methods can be found at <http://vis-www.cs.umass.edu/~buluswar> [4, 5].

References

- [1] T.E. Boult and G. Wolberg, "Correcting Chromatic Aberrations Using Image Warping", *DARPA Image Understanding Workshop*, 1992.
- [2] C.E. Brodley and P.E. Utgoff, "Multivariate decision trees", *Machine Learning*, 1995.
- [3] S. Buluswar, *Trichromatic model of Daylight Variation*, University of Massachusetts Computer Science Department, technical report, UM-CS-1995-012.
- [4] S. Buluswar, *Color prediction in outdoor images from context-based models*, University of Massachusetts Computer Science Department, technical report, UM-CS-1997-030.
- [5] S. Buluswar, *Color recognition in outdoor scenes by nonparametric learning*, University of Massachusetts Computer Science Department, technical report, UM-CS-1997-029.
- [6] J. Crisman and C. Thorpe, "Color Vision for Road Following", *Vision and Navigation: The Carnegie Mellon NAVLAB*, Kluwer, 1990.
- [7] G.D. Finlayson, "Color Constancy in Diagonal Chromaticity Space", *Proceedings of the Fifth International Conference on Computer Vision*, 1995.
- [8] D. Forsyth, "A Novel Approach for Color Constancy", *International Journal of Computer Vision*, 5:5-36, 1990.
- [9] B.V. Funt, G.D. Finlayson, "The State of Computational Color Constancy", *Proceedings of the First Pan-Chromatic Conference*, Inter-Society Color Council, 1995.
- [10] R. Gershon, A. Jepson and J. Tsotsos, *The Effects of Ambient Illumination on the Structure of Shadows in Chromatic Images*. RBCV-TR-86-9, Dept. of Computer Science, University of Toronto, 1986.
- [11] B.K.P. Horn, *Robot Vision*, MIT Press, Cambridge, MA, 1987.
- [12] 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.
- [13] G.J. Klinker, S.A. Shafer and T. Kanade, "Color image analysis with an intrinsic reflection model", *Proceedings of the International Conference on Computer Vision*, 1988.
- [14] S.K. Nayar, K. Ikeuchi and T. Kanade, "Determining Shape and Reflectance of Hybrid Surfaces by Photometric Sampling", *IEEE Transactions on Robotics and Automation*, 6:418-431, 1990.
- [15] C. Novak, S. Shafer and R. Wilson, "Obtaining Accurate Color Images for Machine Vision Research", *Proceedings of the SPIE*, v 1250, 1990.
- [16] Y. Ohta and Y. Hayashi, "Recovery of Illuminant and Surface Colors from Images Based on the CIE Daylight", *Proceedings of the Third European Conference on Computer Vision*, 1994.
- [17] B.T. Phong, "Illumination for Computer Generated Images", *Communications of the ACM*, 18:311-317.
- [18] D.A. Pomerleau, *Neural Network Perception for Mobile Robot Guidance*, Kluwer Academic Publishers, Boston, 1993.
- [19] D.E. Rumelhart, G.E. Hinton and J.L. McClelland, "A general framework for parallel distributed processing", *Parallel Distributed Processing: Explorations in the microstructures of cognition*, Bradford Books/ MIT Press, Cambridge, MA, 1986.
- [20] Y. Sato and K. Ikeuchi, "Reflectance analysis under solar illumination", *Proceedings of the IEEE Workshop for Physics-based Modeling in Computer Vision*, 1995.
- [21] S.A. Shafer, "Using Color to Separate Reflection Components", *Color Research Application*, 10:210-218, 1985.