Color Camera Characterization with an Application to Detection under Daylight

Yves Bérubé Lauzière[†§], Denis Gingras[§], Frank P. Ferrie[†]

[†]Centre for Intelligent Machines, and Dept. of Computer and Electrical Engineering
McGill University, 3480 University, Montréal, Québec, H3A 2A7 Canada

[§]INO, 369 Franquet, Ste-Foy, Québec, G1P 4N8 Canada

Abstract

A technique for measuring the spectral sensitivity curves of color cameras is described. Details on the difficulties that might be encountered in such measurements are discussed. These are mostly related to the non-ideal behavior of cameras, and the properties of the quasi-monochromatic light that must be presented to the camera. The usefulness of these curves is illustrated on a specific example: that of detecting colored objects of known spectral reflectances under daylight illumination conditions.

Keywords: color camera, characterization, spectral sensitivity, physics-based vision, color detection, daylight

1 Introduction

In this paper, a technique for measuring the spectral sensitivity curves for each channel of a color camera will be presented. These curves are useful in applications requiring precise color measurements, such as in some quality control and detection tasks. The aim is to discuss the basic issues involved in doing such measurements, and to present a procedure for performing them.

Camera spectral sensitivity characterization for color research has been discussed in recent times by other authors. Based on work by Sharma and Trussel [1], Barnard and Funt [2] have developed a method in which the camera views the color samples of a chart (Macbeth color checker) illuminated by several different spectra. For each spectrum and color sample, the spectral content of the light reflected back by the sample (the stimulus) and seen by the camera is measured using a spectroradiometer. At the same time, the camera R, G, B responses are stored. The relationship between these responses and the spectral content of the stimulus is linear and given by the camera spectral sensitivities. The authors present an optimization approach for inverting this relationship, thereby obtaining the spectral sensitivities. A similar method has been described in [3] in order to map the color responses of a camera to the CIE $L^*u^*v^*$ color coordinates.

Closer in spirit to the present work is that of Vora *et al.* [4], where the camera is stimulated by monochromatic light obtained from a monochromator. The difference between the approach of Vora *et al.* and the present, is that here the monochromatic stimulus is directly sent onto the sensor, whereas in [4], the monochromatic stimulus is sent onto a white standard, with the camera looking at that standard.

The content of the paper is as follows. In the first section, some important facts about color cameras will be discussed. This is adapted from work by Novak *et al.* The motivation is that real cameras deviate from idealized camera behavior, and this must be taken into account while measuring the spectral sensitivity curves of a given camera. More generally, these facts should be known to anyone making use of a color camera for computer vision purposes, since they have consequences on the quality of color images, and also influence the efficiency of algorithms that rely on ideal camera behavior. Of major interest in this respect is the linearity of the response of a camera to the amount of light that impinges on it. A method for characterizing this linearity is explained and illustrated as well. In the second section, a measurement technique for the spectral sensitivity curves is discussed in detail with emphasis on important issues related to the properties of the monochromatic light required. Finally, in the third section, the usefulness of having the spectral sensitivity curves of a given camera is illustrated in a specific example[1].

2 Facts about Color Cameras

The problems that can occur in real color images caused by non-ideal behavior of color cameras have been thoroughly

[1]Note: The color images in this paper are available by anonymous ftp at ftp.cim.mcgill.ca from the directory pub/people/berube/VI99.

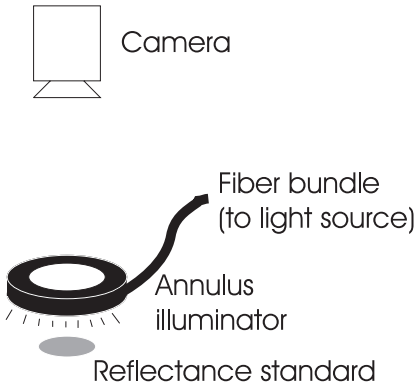

Figure 1: Schematic of the set-up for characterizing the linearity of the camera. The annular light source is a fiber optic light source.

investigated by Novak, Shafer, and Willson [5]. This will be reviewed in some detail, as it is of importance in the sequel.

2.1 Linearity

Of major importance in applications requiring accurate color measurements is to insure that the camera has a linear response. In the negative, appropriate corrections must be made, because the equations modelling color formation in cameras rely on this linearity. Explicitly, these equations have the form

$$C_i = \int E(\lambda) s_i(\lambda) d\lambda, \quad i = R, G, B, \tag{1}$$

with C_i being the response in channel i, E the spectral distribution of irradiance reaching the sensor, and s_i the spectral sensitivity of the combination imaging-optics/color-filter/sensor/framegrabber in channel i. Typically in RGB cameras, the spectral sensitivities are distributed so that one channel detects long visible wavelengths, one medium wavelengths, and one short wavelengths, with some degree of overlap. These are commonly called the red, green, and blue channels respectively. When use of those equations is made, it is thus assumed that as more light enters the camera, the resulting camera measurement will grow proportionately at a given pixel. However, many cameras do not behave in such a linear way. The non-linearities occuring in cameras are described by,

$$C_{i,meas} = C_i^{1/\gamma_i}, \tag{2}$$

with C_i being the value that would be expected from Eq. (1), and $C_{i,meas}$ the value that is actually output [5]. To correct for this (gamma correction), $C_{i,meas}$ has to be raised to the power γ_i, thus requiring γ_i to be known.

The response of a camera with respect to the entering light intensity can be measured with gray calibration standards consisting of several gray samples made from materials of nearly constant reflectances over the visible spectrum.

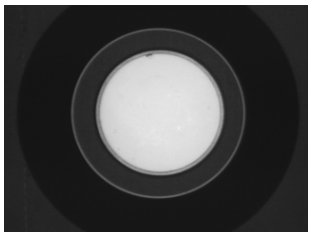

Figure 2: Image of the reflectance standard (white disk at the center) having an 80% reflectance as seen by the red channel of the camera. The darker annulus around the standard is the light source seen from the back, see Fig. 1.

Ideally, when no non-linearity is present, values measured by the camera viewing these gray patches under evenly distributed illumination should lie on a straight line when plotted against percentage reflectance. If it is not the case, then γ_i can be estimated by using a log-log plot. An experimental set-up for performing such measurements is sketched in Fig. 1. In such measurements, great care must be used in evenly illuminating the calibration standards. It must also be ensured that no light other than that reflected by the calibration standard reaches the camera.

Results of an experiment performed to check the linearity of a Panasonic GP-US502 3-CCD micro head color video camera will now be described (this camera is used throughout the following). Reflectance calibration standards from Labsphere Inc. with nominal percent reflectances of 2%, 10%, 20%, 40%, 60%, 80%, and 99% were employed. With the standards evenly illuminated, the camera responses are easily related to the percent reflectance. The light reaching the sensor is first expressed as the light from the source which is reflected by the standard

$$E(\lambda) = R(\lambda) L(\lambda), \tag{3}$$

with $L(\lambda)$ being the light from the source and $R(\lambda)$ the spectral reflectance of the standard. Combining this with Eqs. 1 and 2 gives

$$C_{i,meas} = \left(\int R(\lambda) L(\lambda) s_i(\lambda) d\lambda \right)^{1/\gamma_i}, \quad i = R, G, B. \tag{4}$$

These last equations are the so-called color formation equations with gamma correction included. Taking into account that the spectral reflectance of a standard is constant over the visible spectrum, it can be pulled out of the integral, giving rise to

$$C_{i,meas} = a_i \cdot (\%reflectance)^{1/\gamma_i}, \tag{5}$$

with $a_i = \left(\int L(\lambda) s_i(\lambda) d\lambda \right)^{1/\gamma_i}$. Obviously, in such an experiment, the light source must be set to a constant value throughout the measurements on the different reflectance

standards. Taking the logarithm on both sides results in

$$\log C_{i,meas} = \frac{1}{\gamma_i} \log(\%reflectance) + \log a_i. \quad (6)$$

The slopes and the ordinates at the origin for the different channels were estimated by least-square fitting a line to the data in the log-log plot. For a given reflectance standard, the data point has been obtained by averaging the values of the pixels over the region where the standard is located in the image, see Fig. 2. The following values were obtained: $\gamma_R = 2.0869$, $\gamma_G = 2.0575$, and $\gamma_B = 2.0267$. Fig. 3 illustrates the data points and the fitted lines.

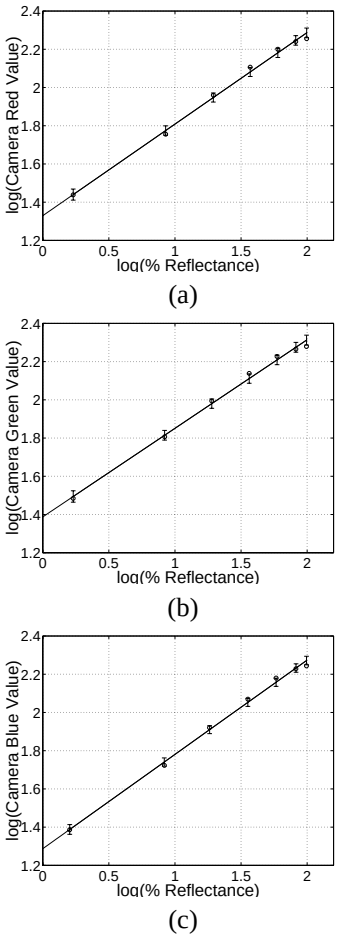

Figure 3: Log-log plots for obtaining the gamma parameters of the camera for the (a) red, (b) green, and (c) blue channels respectively. The lines are least-square fits to the data points, *c.f.* Eq. (6), with error bars showing the estimated errors of the fit at the data points. The data points are within the error bounds of the fit.

2.2 Sensitivity to Non-Visible Wavelengths

Another problem that may occur is that the camera is sensitive to wavelengths out of the visible range. In the ultraviolet (UV), this is not a problem because silicon is negligibly sensitive. Moreover, glass, out of which most lenses are made, is not transparent to such wavelengths. However, silicon is sensitive to infrared (IR) radiation. The peak sensitivity of silicon detectors is typically reached at around 750 nm. Thus, depending on the color filters used in a given color camera, it may be necessary to use an IR blocking filter. In the case of the camera used here, the responses of the channels to IR radiation were found to be negligible, see Fig. 7 which shows spectral sensitivity curves of the camera (how these curves are measured is the subject of next section).

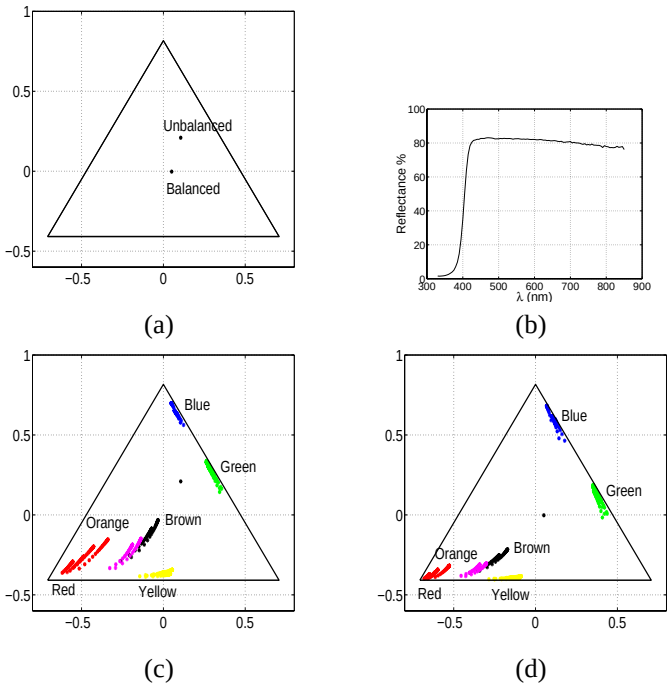

Figure 4: Color balancing versus no balancing. When balancing is performed, whites should be close to the center. In the normalized color coordinates used here, a perfect white, (*i.e.* with a perfectly flat spectral reflectance), should ideally be centered at $(0, 0)$. In (a), a white having the spectral reflectance plotted in (b) is shown with D_{65} (daylight with correlated color temperature of $6500^{\circ}K$) as the illuminant and with a camera having the balanced and unbalanced (factory settings) spectral sensitivities illustrated in Fig. 7. In (c) and (d), the effect of color balancing is illustrated; (c) shows the color distribution of a few colored samples illuminated by different phases of daylight when no balancing is done, the colors are shifted towards the blue corner of the color triangle. In (d) balancing has been performed. Note: in (c) and (d), the distributions for orange and brown are shown in magenta and black respectively.

2.3 Color Imbalance

A further problem is color imbalance in which colors are shifted towards certain portions of color space. The cause can be that the camera is more sensitive to certain ranges

of wavelengths, and/or the light source is emitting more strongly in certain ranges of wavelengths. In both cases, the solution is to change the sensitivity of the camera to the different portions of the spectrum. This is usually done by adjusting the gains of the color channels of the camera using the knobs on the camera unit. In doing so, however, it should be kept in mind that while the signal is boosted, the noise is also amplified, resulting in no improvement in the signal-to-noise ratio (SNR)[2]. This may be important in certain applications, but with the quality of cameras available today it should not be of concern in usual applications. Fig. 7 (b) of the next section shows the spectral sensitivities of the camera for the three color bands, when the factory settings are used. The maximum sensitivity of the blue band is significantly higher than that of the other two. In those settings, the manufacturer has probably compensated for the fact that the incoming light is generally weaker in the blue end of the spectrum for usual indoor artificial light sources. However, depending on the illuminant, those settings might not be appropriate since whites may not be well centered in color space, and that colors tend to be displaced towards blue. Fig. 4 shows an instance of this for daylight, which may have a strong component in the blue end of the spectrum depending on the conditions, see Fig.9 (a). Such drifting due to imbalance can lead to difficulties when colors must be discriminated, because colors then tend to get more closely packed in certain portions of color space.

2.4 Clipping, Blooming, and Chromatic Aberration

In [5] other problems are mentioned: clipping, blooming, and chromatic aberration. Clipping is related to the finite dynamic range of real cameras, and the problem can be avoided by adjusting the aperture, so that the camera does not saturate. Other than that, not much can be done. Blooming, as well, can be prevented in this manner. Concerning chromatic aberration, in common applications it is generally not critical, and can be brought to an acceptable level by resorting to good quality optics. When critical, the method of [5], which consists of correcting chromatic aberration by mechanically adjusting the lens geometry and camera position seems somewhat outdated, and can definitely not be used in real-time applications. The only solution is to invest in high quality compensated optics.

[2]See [5] and [6] for an interesting but perhaps somewhat outdated approach of dealing with color imbalance, since a monochrome camera with a color filter wheel is used. Moreover, this method is impractical in view of real-time applications

3 Spectral Sensitivity Curves

In this section, a method for determining the spectral sensitivity curves of a color camera is presented.

3.1 Method

In the present approach, the spectral sensitivity curves are measured by presenting monochromatic light stimuli to the camera at evenly spaced wavelengths in the visible spectrum. The wavelengths considered were from 400 to $750\,nm$ in steps of $10\,nm$. Monochromatic stimuli are obtained with a monochromator, an instrument that disperses an incident flux of radiation into its spectrum, and from which any narrow band of wavelengths centered around a given wavelength can be isolated with a slit aperture. Details about monochromators are given in [7]. The width of the band of wavelengths centered at a given wavelength is a measure of the monochromaticity of the light exiting the monochromator. It depends on the focal length L of the monochromator, the width w of the entrance and exit slits[3], and the pitch p of the diffraction grating (in $lines/mm$), in the following manner,

$$\Delta\lambda = \frac{w}{p \cdot L}. \tag{7}$$

Because of this finite width, it will thus be more accurate to talk in terms of quasi-monochromatic signals in the sequel. For the monochromator used here, the values are $L = 500\,mm$, $p = 1200\ lines/mm$. In the measurements to be presented later, the width of the slits was set to $3\,mm$, resulting in a bandwidth of $5\,nm$ (in [4], the authors report that their stimuli had a bandwidth of $8\,nm$). This is not narrow for an optical bandwidth, and much better can be achieved (on the order of tenths of nanometers), but this is sufficient for the purpose of characterizing a camera for most computer vision applications. Moreover, there is a trade-off between monochromaticity and light intensity. Ideally one would like to have a bandwidth $\Delta\lambda$ as small as possible, for greater purity, but then less light is available at the exit of the monochromator, because the slits are narrower. In the present case, enough light intensity was required, and thus the widths had to be made large enough.

To properly image the light emerging from the exit slit onto the detector, and to measure in which amount light reaches it, an optical set-up must be used. This is schematically illustrated in Fig. 5.

For each narrow band of wavelengths in the visible spectrum, an image of the slit is grabbed (see Fig. 6 (a) for an example of a stimulus centered at 700 nm), and the amount of power in the stimulus reaching the camera is measured using an optical power meter. This last measurement is necessary

[3]Those widths should always be the same, because the optics within the monochromator image the entrance slit onto the exit slit.

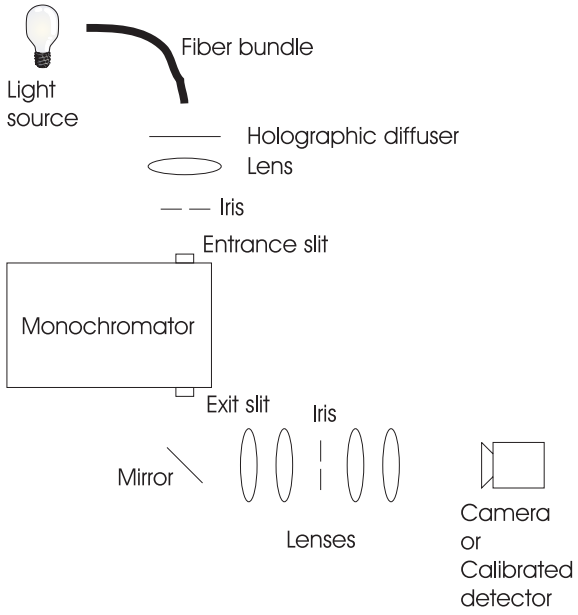

Figure 5: Schematic of the set-up for measuring the spectral sensitivities of the camera using the substitution method. Light from the source is coupled through an incoherent fibre bundle and a holographic diffuser in order that the light be as uniform as possible. In the exit arm of the set-up, lenses are used in combination to compensate for chromatic aberration. Had the direct method been used, a beam splitter would have been placed between the last lens and the camera, and the calibrated optical power detector would be located in the other path of light generated by the beam splitter.

in order to refer the R, G, B responses of the camera to the amount of power contained in the stimulus which is not the same for all wavelengths. The measurement of the power can be done while the image of the slit is being grabbed, in which case a beam splitter must be used to send part of the beam to a calibrated optical power detector, and the other part to the camera. This is an example of a so-called direct method. It has the drawback that not all the light exiting the monochromator is sent to the camera. Since this amount is already small for a reasonable light source (200 Watts), the full dynamic range of the camera cannot be covered in the measurements, leading to SNRs that can definitely be improved. To alleviate this, one must use more poweful light sources, which require substantial cooling, and tend to wear out rapidly. Moreover, such light sources tend to have large temporal fluctuations, making the measurements less accurate. The fact that only small amounts of light are actually available at the exit of the monochromator is also due to the need of having a very uniform patch of light at its output as explained below. If no care is used, the problem encountered is that the filament of the lamp is imaged onto the camera, making the patch seen by the camera highly non-uniform. One way to avoid this is to ensure that the light fed to the monochromator is uniform. This can be done with the use of an incoherent illumination fibre bundle and a diffuser for coupling the light from the source to the monochromator. The price paid is a reduced amount of source light that can be converted into quasi-monochromatic light.

Instead of a direct method, substitution can be used. Here the quasi-monochromatic stimuli at equally spaced wavelengths are sent sequentially onto the camera, and an image of each of them is acquired. The camera is then replaced by a calibrated optical detector, and the same sequence of stimuli is presented to the detector to measure the power contained in the stimuli. This requires the source to be stable in time since the images of the stimuli and the measurements of their respective power are not made a the same time. This is easily achieved for not too intense sources powered with a stabilized current source. The advantage of this method is that all the light available at the output of the monochromator is sent onto the sensor that must be characterized. The disadvantage is that the source may have fluctuated between the image acquisition and the measurement of the power contained into the patch. Averaging several measurement sequences can be used to minimize such fluctuations. In what follows, the substitution method with averaging is the one that has been employed.

Once the image of the stimulus at a given wavelength has been captured, and the power measured, the data must be processed to extract the sensitivities at that wavelength. Since the stimulus covers a region of pixels, a spatial average is first performed. The region over which this average is computed is chosen to be the same *for all* wavelengths. In this manner, it is always the same pixels that are considered. It will thus be the average sensitivities over these pixels that will be obtained in the end. To determine this region, three binary masks are generated by thresholding the blue, green and red images of stimuli centered at 480, 565, and $700\,nm$ respectively, wavelengths at which each channel is the most sensitive for the camera considered here (for other types of cameras these wavelengths may be different, and must be determined). The masks are ANDed to get the final mask. Refer to Fig. 6 for images of the stimuli and masks. For the measurements to be meaningful, great care must be taken in making the intensity of the patch of quasi-monochromatic light as uniform as possible. How to do this was discussed above (incoherent fiber bundle and holographic diffuser in Fig. 5). This prevents the responses of the pixels within the stimulus from being too scattered. Looking at the masks for the red, green, and blue channels in Fig. 6 (g), (h), and (i), it is seen that not all pixels in the stimulus are present. Some were eliminated because they had values too far from those of the rest of the pixels. This further reduces the scattering of the responses, and is easily determined by using a histogram.

Let $C_{avg,i}(\lambda_0)$ denote the spatial average in channel i at wavelength λ_0. In terms of spectral integration, this is equal

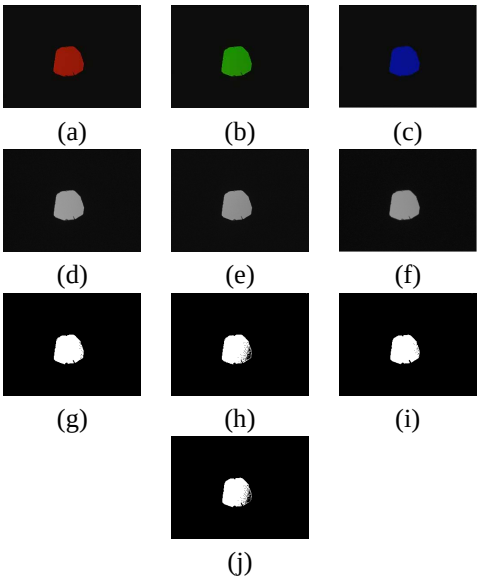

Figure 6: Illustrating the image of the exit slit onto the detector and the binary masks generated to measure the spectral sensitivity curves of the camera. (a), (b), and (c) show quasi-monochromatic stimuli centered at 700 nm, 565 nm, and 480 nm respectively. (d), (e), and (f) are respectively: the image of the red channel of (a), the image of the green channel of (b), and the image of the blue channel of (c). These images are thresholded to generate the binary masks in (g), (h), and (i). The three masks are ANDed to generate the final mask in (j).

to

$$C_{avg,i}(\lambda_0) = \left(\int L_{\lambda_0}(\lambda) s_i(\lambda) d\lambda \right)^{1/\gamma_i}, \qquad (8)$$

where $L_{\lambda_0}(\lambda)$ represents the quasi-monochromatic stimulus centered at λ_0. In first approximation it can be set equal to a delta function with a multiplicative factor proportional to the measured power contained in the stimulus,

$$L_{\lambda_0}(\lambda) = b \;\; power(\lambda_0) \;\; \delta(\lambda - \lambda_0). \qquad (9)$$

The delta function approximation gets better as the width of the quasi-monochromatic signal gets narrower. Combining Eqs. 8 and 9 gives

$$b\, s_i(\lambda_0) = \frac{C_{avg,i}^{\gamma_i}}{power(\lambda_0)}. \qquad (10)$$

The spectral sensitivities are thus obtained within a constant factor which is the same for the three channels and for all wavelengths. In other words relative values are obtained.

3.2 Results

Using the method just described, the spectral sensitivity curves of the camera were obtained. Fig. 7 depicts the results for the camera equipped with an objective. It is important to characterize the camera with the objective that will

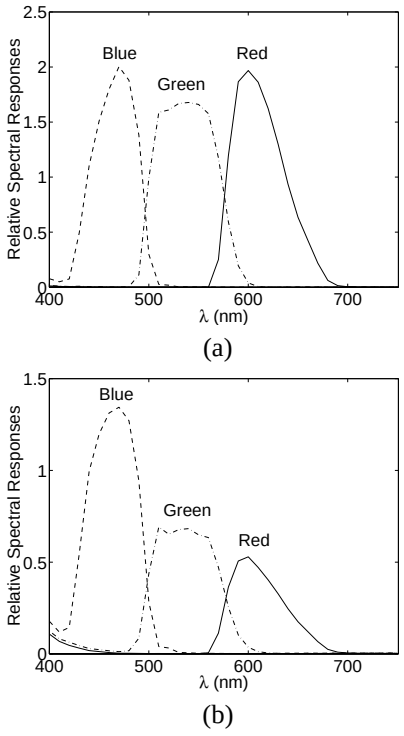

Figure 7: Spectral sensitivity curves for the camera characterized with an objective. These were obtained by scanning the spectrum at 10 nm intervals between 400 nm and 700 nm. In (a) color balancing has been performed, whereas this is not the case in (b) where the curves are those of the factory settings.

be used in the application, as well as for the channel gains at which it will be set. These are important reasons why one should do the measurements himself, because manufacturers cannot provide the spectral sensitivity curves for every camera objective on the market, and for everyone's settings. All measurements should be made in complete darkness, and with the automatic gain control (AGC) of the camera disabled, in order for the camera not to adjust the gains itself. Of course, when the camera is actually used, the AGC must also be off for the same reasons. In reality, what is so characterized is the combined effect of sensor, filters, gain adjustments, optics and frame grabber. For computer vision applications, this is what is necessary.

4 Application

A physics-based approach for detecting objects under daylight illumination will be briefly presented. For more details, the reader is referred to [8]. This is to illustrate the use of the spectral sensitivities of a camera in a specific example where the illumination is not controlled, a case not often considered in physics-based vision. The challenge is to be able to use color as a useful means of identification, even

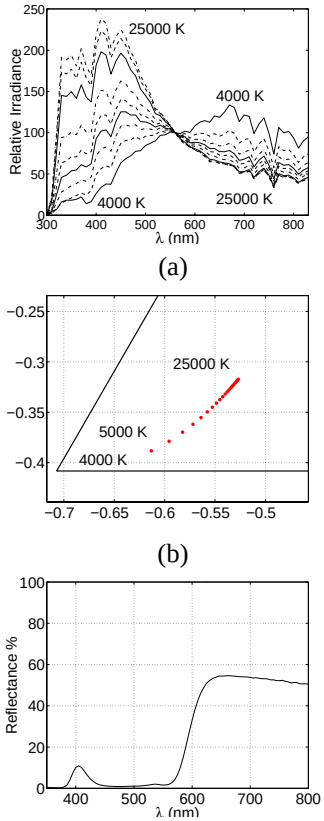

Figure 8: (a) Spectral distributions of daylight with correlated color temperatures ranging from 4000 $^{\circ}K$ to 25000 $^{\circ}K$ in steps of 1000 $^{\circ}K$ obtained from the model of Judd *et al.* [9]. (b) Trace of predicted color points using Eq. 11 for a material with red appearance having the spectral reflectance shown in (c).

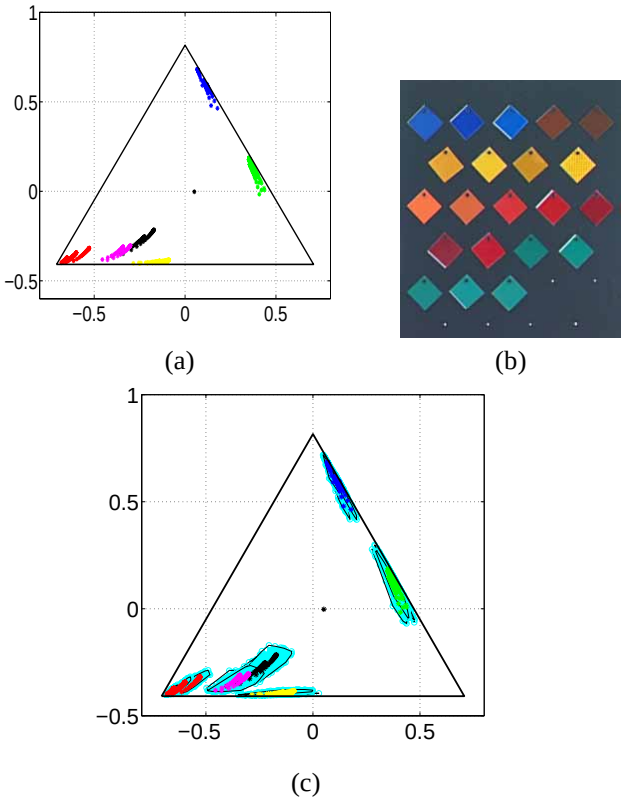

Figure 9: (a) Traces of predicted colors under daylight for the various samples on the chart shown in (b). (c) Detection regions used for pixel classification.

though the response of the camera to a given colored object changes due to variations in the illumination.

More precisely, the approach is based on the color formation equations, which require the spectral sensitivities to be known, and uses a semi-empirical model of the spectral distributions of daylight due to Judd *et al.* [9] to *predict* the color responses of a given camera to an object of known spectral reflectance. In this context, the color formation equations, with gamma correction included, are expressed by

$$C_{i,meas}(T_c) = \left(\int R(\lambda) S_D(\lambda, T_c) s_i(\lambda) d\lambda \right)^{1/\gamma_i} \quad i = R, G, B, \tag{11}$$

where $R(\lambda)$ is the spectral reflectance of the object of interest (which is measured with a spectrophotometer), and $S_D(\lambda, T_c)$ is the spectral distribution of a phase of daylight with correlated color temperature T_c, see Fig. 8 (a). The spectral sensitivities $s_i(\lambda)$ are those of Fig. 7 (a). When plotted in normalized color space, the predictions (for a given material) corresponding to the different phases of daylight give a trace of points. This is shown in Fig. 8 (b) for a red sample having the spectral reflectance plotted in (c). In Fig. 9 (a), such points are shown for the samples of the chart shown in Fig. 9 (b). Around each such trace, a region is determined by introducing uncertainty into the model via a Monte Carlo simulation in order to take factors into account which cannot be embodied by the model. This is shown in Fig. 9 (c). Overlapping regions associated with different samples, but having the same color are merged to form a detection region associated to that color. Pixels in an image can then be labelled according to whether or not their values fall in such regions. Neighboring pixels having the same label are then grouped. In this manner objects can be detected. Note that orange and brown pixels must be further processed due to the overlap between their corresponding regions which cannot be avoided [8].

The method has been tested under various daylight illumination conditions with very satisfying results, a few of which are shown in Fig. 10, see [8] for more. The chart used for the tests is that of Fig. 9 (b). In the results presented in Fig. 10, pixels detected as being blue are highlighted in vivid blue, those detected as brown are depicted in light brown,

(a)

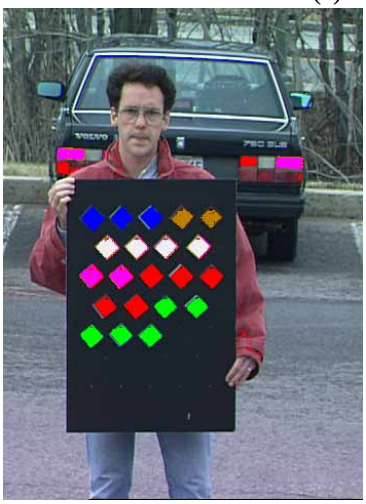

(b)

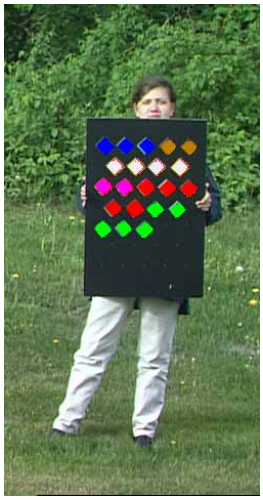

(c)

Figure 10: Detection examples for a few different illumination conditions. The following conditions prevailed: (a) captured during winter under sunny conditions, (b) spring, heavy overcast sky, and (c) summer, light overcast.

yellow pixels appear in white, orange pixels in magenta, red pixels in vivid red, and green ones in vivid green [4]. It happens that other objects than those on the chart are detected, because the values of the pixels they contain fall in the detection regions, see Fig. 10 (b). It is not the purpose of the method to eliminate them, this will have to be done by other means, *e.g.* shape analysis.

[4]The colors used for highlighting and described here are those seen electronically on the computer screen. When printed on paper, these may be altered depending on the color printer used and its settings. For instance, with our color printer vivid blue on the screen appeared as dark blue on paper, and magenta as pink. Concerning pixels detected as orange and yellow, these are highlighted in magenta and white respectively, for better visibility on the screen as well as on paper.

5 Conclusion

A method for measuring the spectral sensitivity curves of color cameras was discussed. It was presented for a 3-CCD camera, but can be applied to 1-CCD chip color cameras as well. Practical aspects were discussed, with some emphasis on fundamental issues related to non-ideal camera behavior, and the properties of the quasi-monochromatic light that should be applied to the camera. An example of the use of those curves was given, but other applications requiring precise color measurements and using physics-based methods could have been considered as well.

Acknowledgements The authors would like to thank the reviewers for helpful comments on the initial version of the manuscript. Special thanks to Mario Marois, Jérôme Comallonga, and Pascal Parrein for their help in the experiments. Y.B.L. acknowledges financial support from the Natural Sciences and Engineering Council of Canada, and from McGill University through its McGill Major Fellowships Program.

References

[1] G. Sharma and H. Trussel, "Characterization of Scanner Sensitivity," *IS&T and SID's Color Imaging Conference: Transforms and Transportability of Color*, pp. 103–107, 1993.

[2] K. Barnard and B. Funt, "Camera Calibration for Colour Vision Research," *Electronic Imaging IV, SPIE Proceedins*, vol. 3644, 1999.

[3] N. Strachan, P. Nesvadba, and A. Allen, "Calibration of a Video Camera Digitising System in the CIE $L^*u^*v^*$ Colour Space," *Pattern Recognition Letters*, vol. 11, pp. 771–777, 1990.

[4] P. Vora, J. Farrel, J. Tietz, and D. Brainard, "Digital Color Cameras-2-Spectral Response." Available from http://color.psych.ucsb.edu/hyperspectral/, 1997.

[5] C. L. Novak, S. A. Shafer, and R. G. Willson, "Obtaining Accurate Color Images for Machine Vision Research," *Perceiving, Measuring, and Using Color, SPIE Proceedings*, vol. 1250, pp. 54–68, 1990.

[6] G. J. Klinker, *A Physical Approach to Color Image Understanding*. Wellesley, Massachusetts: A. K. Peters, 1993.

[7] G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*. New York: John Wiley and Sons, 2 ed., 1982.

[8] Y. B. Lauzière, D. Gingras, and F. P. Ferrie, "Physics-Based Color Detection under Daylight," *submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence*, August 1998.

[9] D. B. Judd, D. L. MacAdam, and G. Wyszecki, "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature," *Journal of the Optical Society of America*, vol. 54, pp. 1031–1040, August 1964.