Kobus Barnard*

Computer Division
University of California, Berkeley
Berkeley, CA 94720-1776

Lindsay Martin, Brian Funt, and Adam Coath

Department of Computing Science
Simon Fraser University
8888 University Drive
Burnaby, BC, Canada, V5A 1S6

A data set for color research

We present an extensive data set for color research which has been made available on-line (http://www.cs.sfu.ca/~colour/data). The data is especially germane to research into computational color constancy, but we have aimed to make the data as general as possible, and we anticipate a wide range of benefit to research into computational color science and computer vision. Since data is only useful in context, we provide the details of the collection process, including the camera characterization, and the data used to determine that characterization. The most significant part of the data is 743 images of scenes taken under a carefully chosen set of 11 different illuminants. The data set also has several standardized sets of spectra for synthetic data experiments, including some data for fluorescent surfaces.

INTRODUCTION

In this paper we present an extensive data set for color research which has been made available on-line (http://www.cs.sfu.ca/~colour/data). The data was collected over a period of several years as part of our ongoing investigation into computational color constancy[1,2]. In addition to providing the

---

*Work completed while the first author was at Simon Fraser University.

data, we have taken care to identify the correspondence between data sets and published results. We believe that this information is important for efficient collective and collaborative progress in computational color science. Although the data is especially germane to research into computational color constancy, we have aimed to make the data as general as possible, and we anticipate a wide range of benefit to research into computational color science and computer vision. We organize the data into the 5 different components. They are:

1. The data used to characterize our camera.

2. The camera sensor responses, 4 sets of illuminant spectra, and reflectance spectra.

3. Images of scenes taken under a carefully chosen set of illuminants split into four groups

    a) Images with minimal specularities

    b) Images with dielectric specularities

    c) Images with metallic specularities

    d) Images with fluorescent surfaces.

4. Images of single objects in various poses in front of a black background.

5. Reflectance data for the fluorescent surfaces which occur in some of the images.

We will now describe each component in more detail.

CAMERA CHARACTERIZATION DATA

We include the data we used to characterize our Sony DXC-930 3CCD camera[3]. There are two reasons why it is important to include such data. First, camera characterization is an on-going research endeavor, and the data we provide will allow others to compare our characterization results, and improve upon them. Second, although we provide the characterization for convenience, other researchers should not be bound to use this characterization. Specifically, if they feel that color constancy can be improved by better characterization, then they should be able to pursue this line of research while still using the corresponding image data.

In the context of computational colour vision, we define camera characterization as determining the relationship between input energy spectra and camera responses. Normally this is broken down into a linearization function and a sensitivity function for each channel. Our data for calibration is a

collection of corresponding input energy spectra (radiance) and camera responses. We measured the radiance spectra using a PhotoResearch spectroradiometer. Our target was a Macbeth ColorChecker® which has 24 different colored patches which we illuminated with a number of illuminant/filter combinations. The black patch was not used because it did not reflect enough light with the darkest illuminants to allow reliable spectroradiometer measurements.

The main criterion for camera characterization was that the camera and spectroradiometer measured the same signal. Furthermore, we required the camera data to be from the center of the image. Therefore we mounted the color checker horizontally on an XY table which moved it under computer control. The camera and the spectroradiometer were mounted on the same tripod, with their common height being controlled with the tripod head height adjustment mechanism. We set the two optical axes to be parallel. The tripod was raised and lowered between capturing camera data and spectroradiometer data. Thus we captured an entire chart worth of camera data before capturing an entire chart worth of spectra.

We took additional steps to obtain clean data. As indicated above, it is important that the camera and the spectroradiometer are exposed to the same signal. To minimize the effect of misalignment, we made the illumination as uniform as possible. To reduce the effect of flare, the target was imaged through a hole in a black piece of cardboard, exposing the region of interest, but as little else as was practical. We extracted a 30 by 30 window from the image which corresponded as closely as possible to the area used by the spectroradiometer. The 8 bit RGB values of the pixels in this window were averaged. Finally, the camera measurements were averaged over 50 frames to further reduce the effect of photon shot noise, and the spectroradiometer measurements were averaged over 20 capture cycles.

We used 26 illuminant/filter combinations for each of the 23 target patches for a total of 598 measurements. Some of the pixels have one or more channel values of 255. These data values are likely clipped and should not be used. In fact, in our characterization procedure we normally exclude values which are greater than 240. Nonetheless, we include all data, since most pixels have at least one non clipped value. For example, even if the red channel is clipped, the green or

blue channel may contain useful information. We also include the camera response to no light, measured by averaging pixel values over a large number of frames taken with the lens cap on.

SPECTRAL DATA

The second component of the data set is measured spectral data. These data sets have been used for a number of published results[2,4-7], and should be useful for comparison and further work. All the spectra consist of 101 measurements from 380 nm to 780 nm in steps of 4 nm, which is the form used by the PhotoResearch® PR-650 spectrophotometer. We include the camera sensor functions (shown in Figure 1), 4 sets of illuminant spectra, and a set of reflectance spectra.

The 1995 surface reflectance are compiled from several sources. The surfaces include the 24 Macbeth ColorChecker patches (measured by ourselves), 1269 Munsell chips, 120 Dupont paint chips[8], 170 natural objects[8], the 350 surfaces in Krinov data set[9], and 57 additional surfaces measured by ourselves.
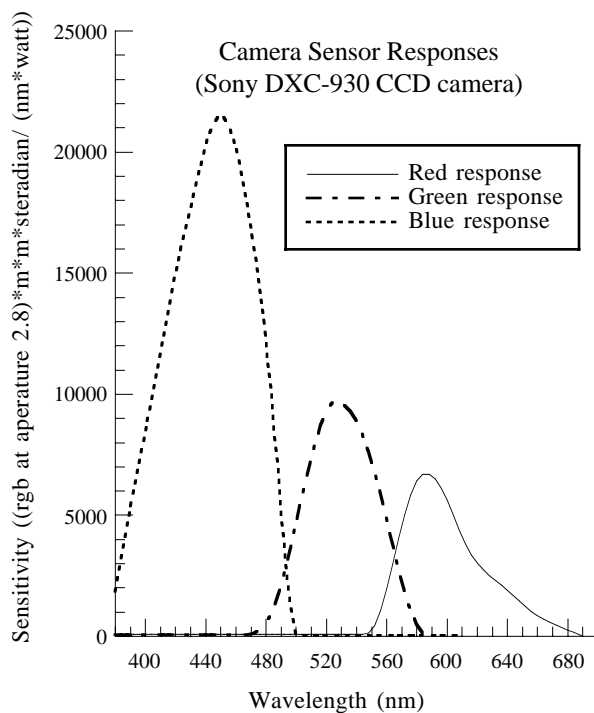

FIG 1. The camera sensitivities for the Sony DXC-930 used in this study.

We provide 4 sets of illuminant spectra. The first is for a set of 11 sources used for the image data. These were selected to span the range of chromaticities of common natural and man made illuminants as best as possible, while bearing in mind the other considerations of stability over time and physical suitability. These sources include three fluorescent lights (Sylvania warm white, Sylvania cool white, and Philips Ultralume), four different 12 volt incandescent lights, and those four used in conjunction with a blue filter (Roscolux 3202). The spectrum of one of the incandescent sources (Sylvania 50MR16Q) is very similar to a regular incandescent lamp. The other three have spectra which is similar to daylight of three different color temperatures (Solux 3500K, Solux 4100K, Solux 4700K). When used in conjunction with the blue filter, these bulbs provide a reasonable coverage of the range of outdoor illumination. The chromaticities (relative to our camera) of all 11 illuminants are shown in Figure 2(a).

The second illuminant set is based on 81 spectra measured in and around the Simon Fraser University campus, at various times of the day, and in a variety of weather conditions. Unusual lighting, such as that beside neon advertising lights, was excluded. However, care was taken to include some reflected light, provided that it was not too extreme. This set of illuminants was augmented with measurements of 27 sources, including the 11 above. The chromaticities of this illuminant set are shown in Figure 2(b).

From this second set we created two more illuminant sets, the first of which we have used
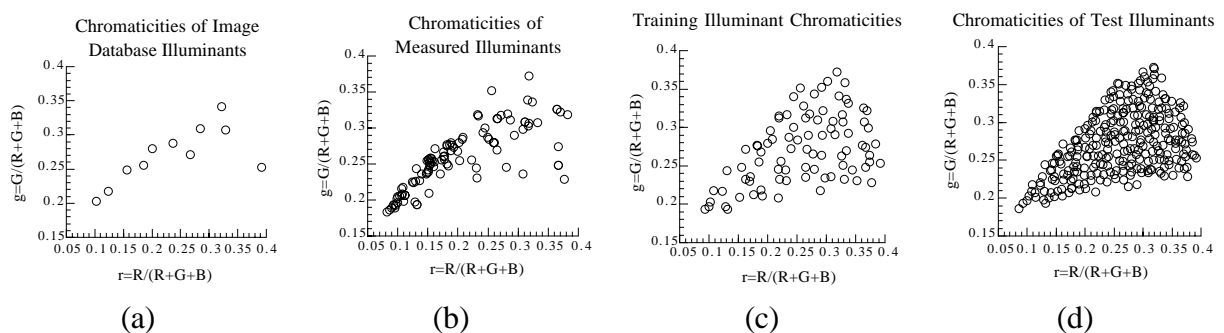

FIG 2.   The chromaticity distributions of the four illuminants sets of this data. The chromaticities are specific to the camera used, but are based on spectra which can be used in conjunction with any camera model or standard color space.

extensively for algorithm training, and the second for algorithm testing. To create the illuminant set used for training, we divided (r,g) space into cells 0.02 units wide, and placed the 11 sources described above into the appropriate cells. We then added spectra from the second measured set, provided that their chromaticity bins were not yet occupied. Finally, to obtain the desired density of coverage, we used random linear combinations of spectra from the two sets. The fourth illuminant set was produced using the same procedure, but the space was filled 4 times more densely. The chromaticities of these illuminant sets are shown in Figures 2(c) and 2(d).

Using linear combinations of spectra is valid because illumination is often the blending of light from two or more sources. In addition, to the extent that illumination changes can be modeled by a diagonal transform[7,10-14] (global scaling of the three channels), these constructed illumination spectra will behave like those from physical sources with the same chromaticities.

IMAGE DATA (SCENES)

The most significant component of this data is undoubtedly the image data. There are two groups of images—images of scenes, discussed in this section, and images of objects, discussed in the next section. The images of scenes are divided into four categories. These are a set of images with minimal specularities (21 scenes), images with non-negligible dielectric specularities (9 scenes), images with metallic specularities (14 scenes), and image of scenes with at least one fluorescent surface (6 scenes). Each scene was imaged under the 11 sources described above. Some images were omitted due to deficiencies in the calibration data. This left 223 valid images in the first set, 98 in the second, 149 in the third, and 59 in the fourth. The different images of a given scene are not absolutely guaranteed to be registered, as some scene shift between illuminants could have occurred, but for the most part, the registration is very close. If registration is required, then it is easily checked. This data has been used for a number of published results[5,6,15,16].

The experimental routine to capture the images was as follows. First we constructed a new scene. We then placed a reference white standard in the center of the scene, perpendicular to the direction of the illuminant. The position of the illuminant was set so that the number of clipped

pixels was small (usually zero). This meant that if the image had bright specularities, then it was purposely under-exposed. We then took a picture of the scene with the reference white in the center, and captured the spectrum of the light reflected from the reference white. Finally, we removed the reference white, and took 50 successive pictures which were averaged to obtain the final image. We then repeated the process for the remaining 10 illuminants, and then we moved onto the next scene.

The images with the reference white were used to determine the (R,G,B) of the illuminant for the corresponding image. We extracted the central 30 by 30 pixel window of each of these images, and used the average (R,G,B) over these windows as the estimate of the illuminant (R,G,B). Both the final images and the target (R,G,B) values were mapped into a more linear space, and received other corrections discussed more fully below. This method provided a good estimate of the chromaticity of the illuminant, but the error in the illuminant (R,G,B) magnitude for any given picture can be quite high–easily 10%, because of the difficulties in keeping the white reflectance standard perpendicular to the light source. Furthermore, the fluorescent sources were spatially extended, and here we simply attempted to find the orientation which maximized the brightness of the reflectance standard.

Because of the frame averaging described above, the images have large dynamic range, and the image pixels have more than the usual 8 bit precision which was preserved by averaging with floating point arithmetic and storing the results in a floating point format. Several pre-processing steps were taken to improve the data. First, we removed some fixed pattern noise determined by averaging a large number of frames taken with the lens cap on. The pattern is essentially vertical striping at a spatial frequency of 4 pixels with a magnitude of roughly 1. This noise is negligible in non-frame averaged images, but is quite noticeable when dark regions of the extended dynamic range images are brightened, thus justifying its removal.

The next step in image preparation was to linearize the image to correct for a minor non-linearity in our camera[3]. This step also removes the sizeable camera offset (response to no light). The resulting images are such that pixel intensity is essentially proportional to scene radiance.

We then corrected for a spatially varying chromaticity shift due to the camera optics. The correction was based on an image of a uniform surface illuminated at a distance by a single source. This calibration image was then spatially smoothed and used to compute a correction factor for each pixel for each channel. The factors for a pixel, i, are given by:

$$f_{red}(i) = \left(\frac{R_{Middle}}{R(i)}\right)\left(\frac{S(i)}{S_{Middle}}\right) \quad f_{green}(i) = \left(\frac{G_{Middle}}{G(i)}\right)\left(\frac{S(i)}{S_{Middle}}\right) \quad f_{blue}(i) = \left(\frac{B_{Middle}}{B(i)}\right)\left(\frac{S(i)}{S_{Middle}}\right)$$

where

$$S_{Middle} = R_{Middle} + G_{Middle} + B_{Middle} \quad \text{and} \quad S(i) = R(i) + G(i) + B(i)$$

and where $(R_{Middle}, G_{Middle}, B_{Middle})$ is the (R,G,B) in the middle of the calibration image and $(R(i), G(i), B(i))$ is the (R,G,B) of pixel i in the calibration image. We use the middle of the image as the reference because the camera characterization was computed for the middle of the field of view.

Because of the preprocessing and extended dynamic range, it is possible to scale the images by a factor of 10 without incurring too much noise. Thus the images can be re-scaled to emulate capture with an automatic aperture. For example, if a scene has significant specularities, then they would normally be clipped by an automatic aperture mechanism. Such camera behavior can be emulated by scaling the image, and clipping the result at 255. Thus our approach allows the study of higher dynamic range images, which will become more readily available and are worthy of consideration in many applications, but does not rule out the emulation of more standard camera behavior.

In order to allow researchers to experiment with the extra data depth we provide the images as 16 bit tiff images, as well as 8 bit tiff images. However, the 8 bit images contain most of the information in the 16 bit images, and are likely adequate for most applications. We emphasize that the 8 bits in these images are very clean compared to most 8 bit images.

IMAGE DATA (OBJECTS)

The images of objects were captured using the same methodology described above for the scenes. There are two main differences. First, the images in this second set are of single objects on a black

background. Second, the object was rotated between each illuminant change. This set of images is therefore designed along the paradigm of image indexing experiments under illumination change[2,17,18].

FLUORESCENT SURFACE CHARACTERIZATION

We also provide data which characterizes the fluorescent surfaces which occur in some of the images. To describe the reflectance from fluorescent surfaces one normally uses a matrix which specifies the complete spectrum emitted as the result of being exposed to each wavelength sample. Such data is difficult to obtain, so we took a different approach. Since the dimensionality of common illuminants is limited, one can characterize fluorescent surface reflectance adequately by simply specifying the response to a number of representative illuminants[15].

For each surface thought to be fluorescent, we measured the light reflected back while the surface was under a number of different illuminants. If the surface was not fluorescent, then the ratio between the output and the input (the reflectance spectrum) would be constant (ignoring noise and regions of no energy). However, for fluorescent surfaces, this ratio is a function of the illuminant. Therefore our data for each fluorescent surface consists of a number of input/output energy spectrum pairs. The format of the data is explained on the web site.

CONCLUSIONS

We have described a data set for computational color and vision research. The most significant component of the data is a large number of images of scenes captured under 11 carefully selected illuminants. The data set is unique in scope and completeness, containing the required information to follow the development of color constancy algorithms, starting from simple synthetic data, going to images which have more and more difficulties associated with them, and ending with object recognition applications. Thus we look forward to hearing about both anticipated and novel uses of the data.

ACKNOWLEDGMENTS

REFERENCES

1. K. Barnard, Computational colour constancy: taking theory into practice, Simon Fraser University School of Computing Science, M.Sc. thesis, available from ftp://fas.sfu.ca/pub/cs/theses/1995/KobusBarnardMSc.ps.gz (1995).
2. K. Barnard, Practical colour constancy, Simon Fraser University School of Computing Science, Ph.D. thesis, available from ftp://fas.sfu.ca/pub/cs/theses/1999/KobusBarnardPhD.ps.gz (1999).
3. K. Barnard and B. Funt, Camera characterization for color research, *Color Res. and App.* (in press).
4. K. Barnard and B. Funt, Experiments in Sensor Sharpening for Color Constancy, in *Proceedings of the IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications*, The Society for Imaging Science and Technology, Springfield, Va., 43-46 (1998).
5. K. Barnard, Improvements to Gamut Mapping Colour Constancy Algorithms, in *Proceedings of the 6th European Conference on Computer Vision*, Springer, 390-402 (2000).
6. K. Barnard, L. Martin, and B. Funt, Colour by correlation in a three dimensional colour space, in *Proceedings of the 6th European Conference on Computer Vision*, Springer, 375-389 (2000).
7. K. Barnard, F. Ciurea, and B. Funt, Sensor Sharpening for Computational Color Constancy, *J. of the Opt. Soc. of Am. A*, **18**, 2728-2743 (2001).
8. M. J. Vrhel, R. Gershon, and L. S. Iwan, Measurement and Analysis of Object Reflectance Spectra, *Color Res. Appl.*, **19**, 4-9 (1994).
9. E. L. Krinov, *Spectral Reflectance Properties of Natural Formations*, National Research Council of Canada, Ottawa, 1947.
10. G. D. Finlayson, Coefficient Color Constancy, Simon Fraser University, School of Computing, Ph.D. thesis, (1995).
11. G. D. Finlayson, M. S. Drew, and B. V. Funt, Spectral Sharpening: Sensor Transformations for Improved Color Constancy, *J. Opt. Soc. Am. A*, **11**, 1553-1563 (1994).
12. G. D. Finlayson, M. S. Drew, and B. V. Funt, Color Constancy: Generalized Diagonal Transforms Suffice, *J. Opt. Soc. Am. A*, **11**, 3011-3020 (1994).
13. G. West and M. H. Brill, Necessary and sufficient conditions for von Kries chromatic adaptation to give colour constancy, *J. Math. Biol.*, **15**, 249-258 (1982).
14. J. A. Worthey, Limitations of color constancy, *J. Opt. Soc. Am.*, **2**, 1014-1026 (1985).
15. K. Barnard, Color constancy with fluorescent surfaces, in *Proceedings of the IS&T/SID Seventh Color Imaging Conference: Color Science, Systems and Applications*, The Society for Imaging Science and Technology, Springfield, Va., 257-261 (1999).
16. K. Barnard and B. Funt, Color constancy for specular and non-specular surfaces, in *Proceedings of the IS&T/SID Seventh Color Imaging Conference: Color Science, Systems and Applications*, The Society for Imaging Science and Technology, Springfield, Va., 114–119 (1999).
17. B. V. Funt and G. D. Finlayson, Color Constant Color Indexing, *IEEE Pattern Anal. Mach. Intell.*, **17**, 522-529 (1995).
18. B. Funt, K. Barnard, and L. Martin, Is Colour Constancy Good Enough?, in *Proceedings of the 5th European Conference on Computer Vision*, Springer, I:445-459 (1998).