

Illuminating Illumination

Jeffrey M. DiCarlo, Feng Xiao and Brian A. Wandell
Stanford University
Stanford, California

Abstract

We introduce an active imaging method to measure scene illumination. The system implementation is divided into four steps. First, the system acquires two images: one is an ordinary image of the scene under ambient light and the other is a corresponding image in which light from the camera flash is added to the scene. Second, the image pair is analyzed to obtain an image that represents the scene as if it had been illuminated by the flash alone. Third, the flash-only image is used to estimate object reflectance functions. Fourth, using the estimated reflectance functions, the ambient illumination spectral power distribution is estimated. We present results that evaluate the method's stability with respect to changes in the mean reflectance function of the scene. Finally, we discuss limitations of the current implementation and alternative implementations.

Introduction

Images acquired by either film or digital cameras pass through several image-processing stages. The goal of the processing is to render the acquired sensor values into an image reproduction that is both pleasing and faithful to the appearance of the original scene. The performance of several of these processing steps can be substantially improved when the physical properties of objects or illuminants (e.g., distance, shape, motion, reflectance, spectral power distribution, etc) in the scene are known.

Color balancing, one of these processing steps, refers to the transformation of the image data to correct for differences in color appearance when the image is acquired under one illuminant but rendered under a second, different illuminant. For example, say an image of a scene is captured indoors under a tungsten ambient illumination. If rendered without color balancing, the image will have a yellowish appearance (color cast) when viewed under natural outdoor ambient illumination. When properly color balanced, the image appears more similar to the original. Knowledge of the ambient illumination is helpful for color balance processing.

We introduce a novel active imaging method (AIM) that estimates the scene ambient illumination from image data. The method is active because it involves emitting a light into a scene in order to measure a scene characteristic. The illuminant estimation calculations can be performed using a conventional (three-channel) digital imaging system.

The estimated ambient illuminant derived by the method can guide color-balancing transformations. The method can be used in any application, such as light metering, which requires an estimation of the ambient illumination.

Background

The great majority of illuminant estimation algorithms are designed to work in a passive mode: the estimation is based on light recorded passively by film or a digital image sensor. These methods range from calculations based on simple image statistics to more sophisticated calculations that incorporate knowledge about the likely distribution of surfaces and illuminants.

The gray-world algorithm is a simple and widely known passive mode transformation (see Hunt¹ for a description). The method assumes that the average surface reflectance of objects in a scene corresponds to a gray surface; thus, the algorithm uses the average color of an image to set the parameters for color balancing. Buchsbaum offered a clear physical interpretation and mathematical foundation for the method.² If the average reflectance function in the scene is a constant function, then the average of the image data is a measure of the ambient illumination.

Color-by-correlation is a recent passive mode illuminant estimation algorithm that incorporates information about the likely distribution of image information under different lighting conditions.³ The method tests which of the possible illuminants is the most likely one given the image data (see also ⁴⁻⁹). Color-by-correlation basically compares the image chromaticity gamut with the chromaticity gamut derived for each of several illuminants.

Passive mode algorithms incorporate assumptions about the physical properties of the illuminants and objects. These assumptions usually take the form of a linear model for surfaces or illuminants and some assumptions about the likely content of the scene. Passive mode algorithms fail when the assumptions made about the image or its contents are incorrect. An important source of error can be traced to assumptions about the distribution of surface reflectances in the scene. For example: is the mean reflectance gray? One way to improve the performance of these algorithms is to reduce the likelihood of such errors.

Active imaging methods (AIMs) differ from passive algorithms in that they emit a signal into a scene. In principle, AIMs can improve on passive methods because

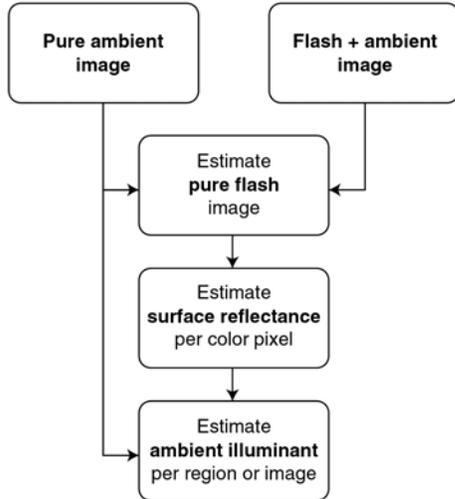


Figure 1. A flowchart of the active imaging method (AIM) for illuminant estimation.

the emitted signal can measure properties of the scene and thus provide a better basis for the assumptions about the image. AIMs are used in various types of camera applications, such as sonar range finders used for auto-focusing.¹⁰ As far as we are aware, AIMs have not been applied to ambient illumination estimation and color balancing.

Computational Methods

Figure 1 provides an overview of the active imaging method we use for illuminant estimation. The method assumes that the camera is calibrated and the spectral power distribution of the camera flash is known. The major steps of the algorithm are outlined below:

1. *Image acquisition.* A sequence of two images of the scene is acquired. The first image is acquired under only the ambient illuminant, and the second is acquired with the camera flash added to the ambient illuminant.
2. *Flash estimation.* Data from the two images are combined to estimate an image that would be measured using only the flash and no ambient illumination. This step provides an image of the scene under a known illuminant.
3. *Surface estimation.* Using conventional estimation methods, surface reflectance estimates are obtained at a variety of image locations from the illuminant-known image.
4. *Ambient illuminant classification.* Finally using the surface reflectance estimates and the pure ambient image, we use conventional methods to determine the most likely ambient illumination.

The details of these algorithm steps are described in the following sections.

Image acquisition

The AIM illuminant estimation procedure begins with a small adjustment to the conventional image acquisition procedure: two pictures instead of one are acquired. The first picture is an image representing the scene under the ambient illuminant. The second picture is an image representing the scene under both the flash illuminant and the ambient illuminant.

Pure flash image estimation

By proper combination of the two acquired images, we create an image in which the ambient illumination is effectively “turned off.”^{11,12} This pure flash image is computed by subtracting the irradiance at the sensor in the ambient image from the sensor irradiance in the ambient plus flash image. Scaling adjustments must be made for exposure time differences, aperture differences and the camera transduction function to convert image digital values to sensor irradiance before subtracting. For example, suppose that both images of the sequence were acquired using a linear camera with the same aperture setting, but different exposure times. To correctly estimate the pure flash image, each image must be scaled by its own exposure time prior to the subtraction. This estimated pure flash image is important because it is a representation of the scene under a known illuminant.

Surface reflectance estimation

Given an image of the scene under a known illuminant (the pure flash image), many methods can be used to estimate object surface reflectances.¹³⁻¹⁵ We used a linear estimation procedure based upon a linear model approximation of the surface reflectances. Suppose \mathbf{B} represents the basis functions of a surface linear model, \mathbf{e}_f represents the flash spectral power distribution, $\text{diag}(\mathbf{x})$ represents placing the vector \mathbf{x} along the diagonal of an identity matrix and \mathbf{R} represents the spectral responsivities of the camera sensors, then the transformation from pure flash image responses, \mathbf{r}_f , to surface reflectance functions, $\hat{\mathbf{s}}$, is:

$$\hat{\mathbf{s}} = \mathbf{B} (\mathbf{R}^T \text{diag}(\mathbf{e}_f) \mathbf{B})^{-1} \mathbf{r}_f \quad (1)$$

Three caveats should be considered when using linear estimation. First, we do not want to include pixels whose values are outside of the camera compliance range. Hence, color pixels that have a response less than 5 (noise) or greater than 250 (saturation) for any sensor in any of the three images (pure ambient, pure flash or combination) should be discarded before estimating the reflectances. Second, the dimensionality of \mathbf{B} should be chosen to match the number of independent sensors; including more or less dimensions can degrade the estimate. Finally, we recommend when building \mathbf{B} that the example surface

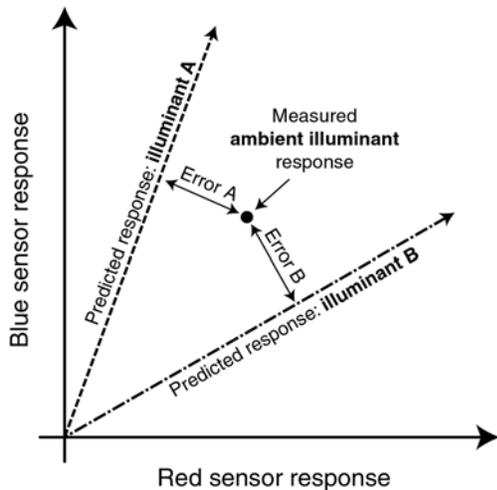


Figure 2. Error computation for two different test illuminants. The dot denotes the measured sensor response due to the ambient light. The vectors denote the predicted sensor response for two different illuminants accounting for the unknown scale factor. The error for each illuminant is the shortest distance from the measured sensor response to the illuminant line.

reflectance functions are normalized to unit length. This operation discards absolute level information. As we will discuss in the next section, the AIM method cannot estimate the absolute level of the illuminant, therefore there is no need to represent scale information in the surface model.

Ambient illuminant classification

Next, we use the estimated surface reflectances and the pure ambient image to classify the ambient illuminant. The classification process estimates the ambient illuminant by comparing the sensor values measured in the pure ambient image with the expected sensor values assuming each of many possible ambient illuminants. The ambient illuminant with the lowest error is chosen.

In general, the expected sensor values under any ambient illuminant, \mathbf{e}_a , can be predicted from the spectral responsivities of the camera sensors, \mathbf{R} , and the estimated surface reflectances, $\hat{\mathbf{s}}$. We can predict the sensor response, $\check{\mathbf{r}}_a$, from the equation:

$$\check{\mathbf{r}}_a = \mathbf{R}^T \text{diag}(\hat{\mathbf{s}}) \mathbf{e}_a \quad (2)$$

The illuminant incident at any position in the image depends upon the scene geometry; we do not have this information so we can only estimate the illuminant (and surface) up to an unknown scale factor. Therefore, we compute the error between the measured responses and the predicted using an error measure that allows for the unknown illumination scale factor (Figure 2). Because the illumination intensity is unknown, the sensor response predicted for a particular illuminant will fall somewhere along a line through the origin. For each illuminant, we

compute the error by measuring the distance between the observed pixel values and this line. We then cumulate these errors across image pixels and select the illuminant with the smallest error.

Experimental Methods

In this section, we describe the specific equipment, test images and linear models that we used to build, test and implement a complete AIM estimation system.

Device Calibration

The experimental AIM imaging device is shown in Figure 3a. It consists of a QImaging Retiga 1300 color camera, a Vivitar 283 flash and a HP Omnibook 4150 portable computer running Matlab 6.0. The portable computer controls both the camera and the flash and performs all the image processing.

We calibrated the imaging device using an Oriel 74000 monochromator. The sensors have a linear transduction function; the spectral responsivity of each of the color sensors is shown in Figure 3b. Lastly, we measured the spectral power distribution of the flash using a SpectraScan PR650 (Figure 3c).

Test Images

We used a Gretag MacBeth Spectralight III light booth to acquire the majority of our images. We collected images under the “A” illuminant (A) setting, the cool white fluorescent (CWF) setting, the daylight (D65) setting and the horizon (HOR) setting. For each illuminant, we acquired nine images of the same scene but with different colored backgrounds. The backgrounds of each image were changed to perturb the image means and image gamuts as much as possible. The goal was to create a set of test images in which it was difficult to make assumptions about the surfaces in the scene.

Figure 4a shows an example image. Figure 4b and 4c show the sensor gamuts of all the images under the four illuminants as well as the image means for each scene. Figure 4b plots the gamut and means in the sensor RB plane, and Figure 4c plots these values in the rb chromaticity plane.

Surface and Illuminant Models

We built a linear model of possible surface reflectance functions using the patches of the Macbeth color chart. The first three principal components of the normalized patch reflectance functions were used in the model.

Our set of common illuminants consisted of blackbody and fluorescent illuminants. We generated blackbody illuminants with temperatures from 2K to 8K. A total of one hundred different blackbody illuminants were generated, spaced evenly in mireds. For the fluorescent illuminants, we used three standard fluorescent illuminants: F2, F7 and F11.¹ Overall, the algorithm had a total of 103 different illuminants to choose from when classifying the true ambient illuminant in a scene.

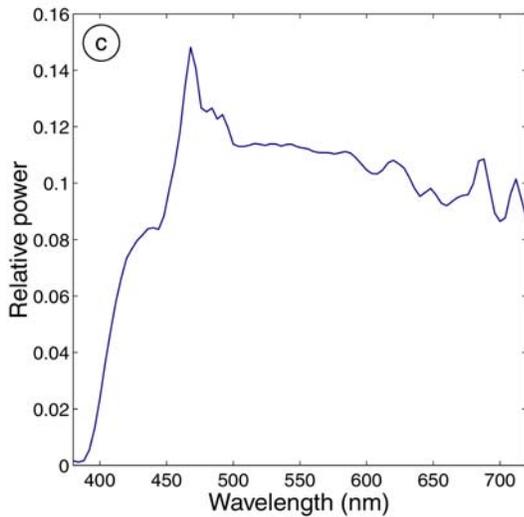
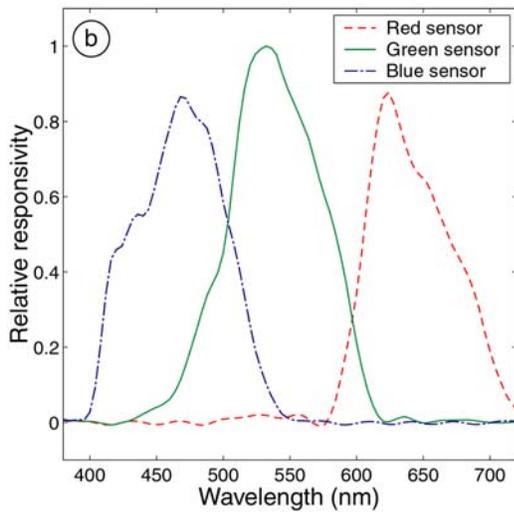


Figure 3. The experimental imaging system. (a) The system includes a color camera, a flash that provides the active illumination, and a portable computer to control the camera and flash. (b) The measured spectral responsivities of the camera sensors. (c) The spectral power distribution of the flash.

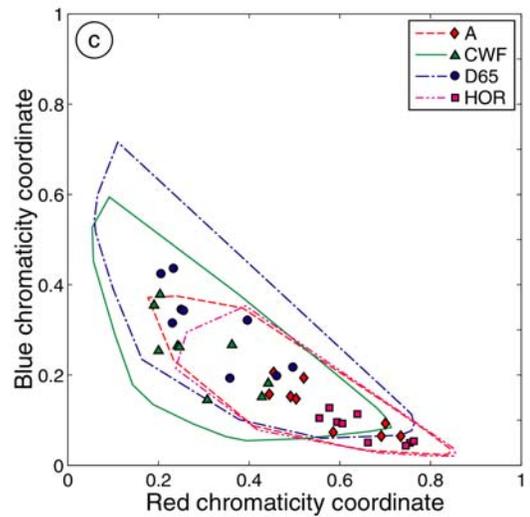
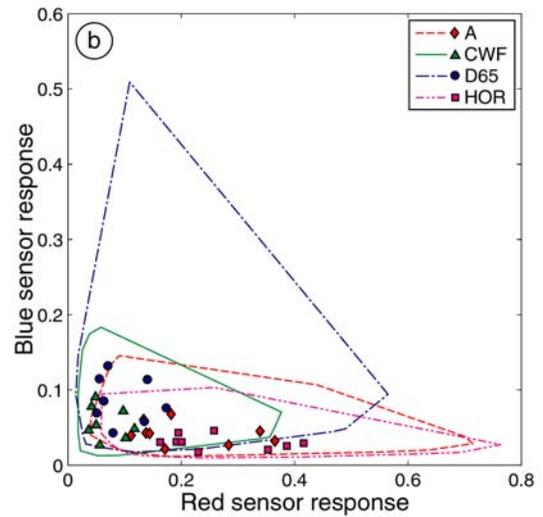


Figure 4. Test images. (a) An example test image with an orange background under the CWF illuminant. (b) The sensor gamuts and image means under all four illuminants plotted in the RB sensor plane. (c) The same gamuts and means plotted in the rb chromaticity plane.

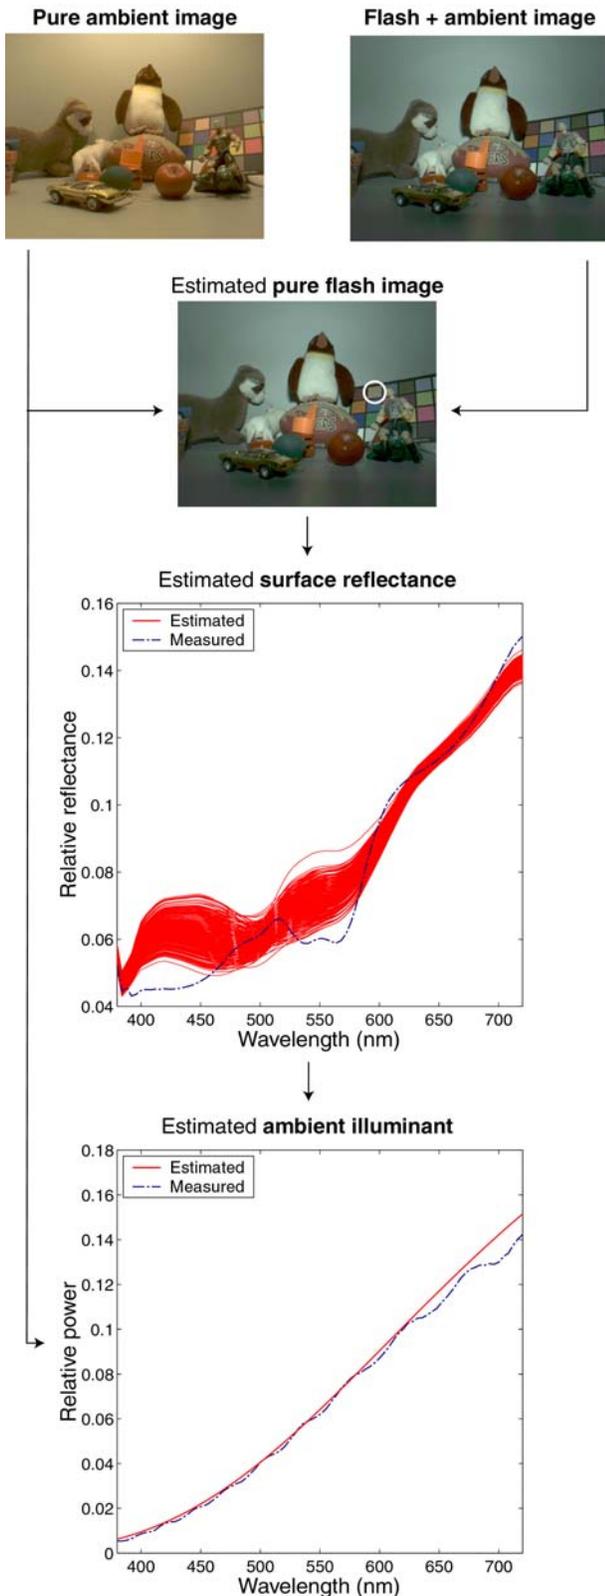


Figure 5. An example AIM illuminant estimation. See the text for details.

Results

An Example

Figure 5 illustrates the AIM illuminant estimation process. The two images at the top were acquired under pure ambient illumination and a combination of ambient and flash illumination. After accounting for their exposure times, the images were subtracted to produce the pure flash image estimate.

A region of the pure flash image containing the brown patch of the Macbeth color chart (circled in white) was selected to illustrate the surface reflectance estimation. This region contains a number of different pixels; the estimated reflectance functions for these pixels are plotted along with the measured reflectance function. The variance among the estimates indicates the size of the measurement error and limits of the Lambertian reflectance model. Differences between the estimates and the true reflectance also occur because the estimate must fall within the linear model of the surface reflectance functions.

The final panel shows the measured illuminant and the illuminant chosen from among the 103 in the classification set. Even though the surface reflectance estimates were not precise, they were adequate to yield a reliable and accurate classification of the ambient illuminant. In the complete process, the surface reflectance estimation and illuminant classification steps are repeated for pixels distributed across the entire image, improving the accuracy further.

Single Scene, Multiple Illuminants

Figure 6 shows the results of processing one scene acquired under the four ambient illuminants available in the Gretag light booth: Illuminant A, CWF, D65 and HOR. The four curves show the classification error as we test with the collection of classification choices. Blackbody illuminant errors are on the left side of the graph and the three fluorescent illuminants are on the right. The circles denote the minimum error solution, that is, the illuminant that the classification algorithm picked. The method correctly classifies these four illuminants. The steepness of the error curves near the local minima shows that the method makes a sharp distinction between the different classification choices.

Multiple Scenes, Multiple Illuminants

In the previous section, we showed that a single scene could be classified according to illuminant type. A more difficult test is to evaluate whether the illuminant estimate is robust with respect to variations in the scene composition. Hence, we repeated the calculations of the previous section using test images (described in the test images section) with very different surface reflectance functions. Again, we used the four different ambient illuminants in the Gretag light booth.

The results for the nine different scenes are shown in Figure 7. The four panels show images under illuminant A, CWF, D65 and HOR. The algorithm correctly classified the illuminants well, though the scene background and

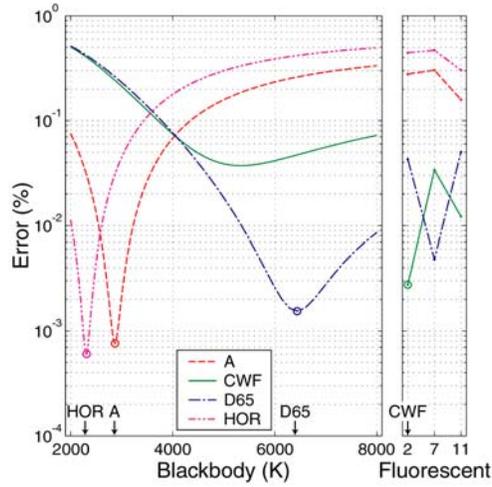


Figure 6. Illuminant classification errors based on the entire image for one scene under the four test illuminants.

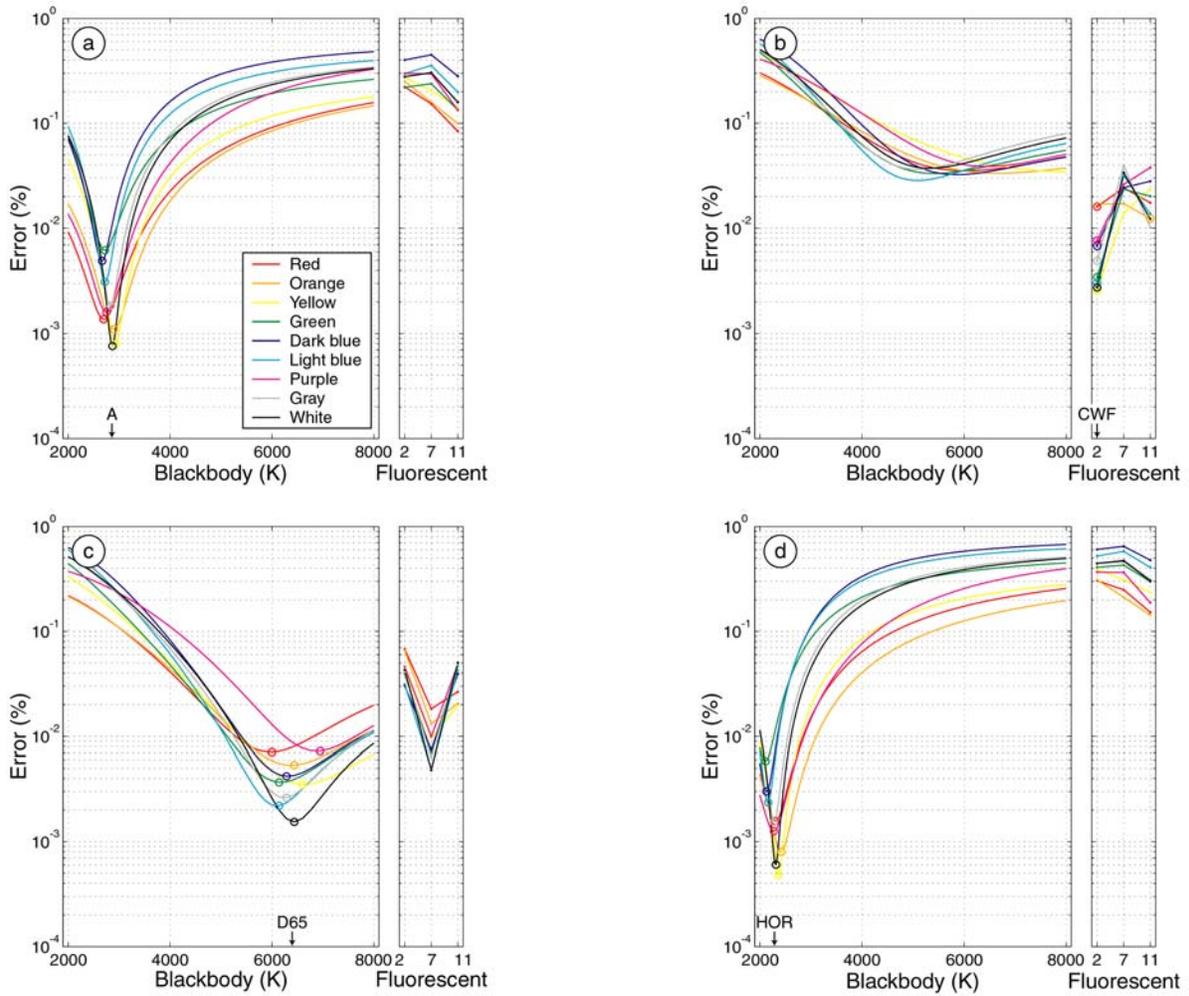


Figure 7. AIM illuminant estimation comparing images that differ substantially in their mean surface reflectance. Each panel is analogous to Figure 6 but shown for multiple surfaces acquired under a specific illuminant: (a) A, (b) CWF, (c) D65, and (d) HOR.

illuminant type varied across the test scenes. For the illuminant A scenes, the average deviation of the blackbody temperature was 90K. For the D65 and HOR scenes, the average deviation was 271K and 98K respectively. For the fluorescent scenes, only the orange background scene was misclassified; the algorithm chose F11 instead of F2. All other fluorescent test scenes were classified correctly.

The graphs in Figure 7 show that the size of the classification error depends on both the illuminant and the surfaces. Further, notice that the classification error was not consistent across the scenes. The estimation accuracy will depend on the spectral sensitivity of the camera sensors, the surface reflectance model, and the spectral power distribution of the ambient light. As the illuminant varies, this combination of factors permits the method to estimate some reflectance functions more accurately than others. Hence, the classification error for different scenes depends on multiple factors.

Spatially Varying Illuminant Classification

The AIM illuminant estimation can be applied to relatively small regions of the image. The quality of the result depends only on the quality of the local surface reflectance estimate. Passive methods, such as gray-world or color-by-correlation, can also be applied locally to the image. However, the core assumptions of the methods (local reflectance is gray; the pixel distribution spans the illuminant gamut) are very unlikely to be true over a local image region. Hence, the AIM illuminant method may have an advantage for measuring spatial-variation in the illuminant.

Figure 8 shows a scene with two illuminants: a fluorescent illuminant and a blackbody illuminant (under the desk lamp). Illuminant estimates were obtained in individual image blocks using AIM. The image blocks with white centers were classified as a blackbody illuminant and the blocks with black centers were classified as a fluorescent illuminant. The blocks with neither a black nor white center had indeterminate estimates because of a lack of light in either the pure ambient image or the pure flash image or both (e.g., a black surface). The AIM illuminant estimates are locally accurate. We have not yet determined the gray-world or color-by-correlation estimates over these local regions.

Limitations

The AIM illuminant estimates are accurate over a range of imaging conditions but not under all conditions. Some of the limitations are inherent to the computational methods, while other limitations are imposed by the specific hardware implementation. In this section, we discuss the limitations that we have observed, and we describe some alternative implementations that might overcome these limitations.

We have implemented the system using conventional camera hardware: a digital camera and a flash. By using this equipment, we could prototype and evaluate a system rapidly. The choice of a flash as an illuminator has an



Figure 8. Spatially varying ambient illumination. The scene is a mixture of fluorescent and blackbody illuminants. The blocks denote the type of illuminant that the algorithm determined in each region of the image. White blocks denote blackbody, black blocks denote fluorescent, and the absence of a block denotes a region where there was not adequate light to make a determination.

important limitation: the method can only be used to estimate the illuminant in those portions of the scene that reflect light from the flash back to the camera. Distant portions of the scene do not return light from the flash; they do not permit an estimate of the ambient illumination. Hence, for this implementation the imaging volume must be on the order of the size of a room.

It is possible to extend the operating range of the method from small spaces to larger spaces by building a system with a more concentrated illumination source. This represents a design tradeoff: the illuminant estimate is obtained from a smaller portion of the scene, but the imaging volume can be increased.

A second limitation of the current implementation concerns the linear model of surface reflectance functions. When used with a single flash and a conventional color camera, the surface model can only be three-dimensional. This is smaller than most estimates of the dimensionality of natural surfaces.¹⁶⁻¹⁸ This causes errors in the surface reflectance estimates that become errors in the illuminant classification.

This limitation can be overcome by acquiring a third image using a second active illuminant. If the spectral power distribution of the second active illuminant differs from the first, one can estimate additional surface reflectance dimensions and improve the accuracy of the method. Such instrumentation might be appropriate, say, for a light meter that could be used in photographic applications, such as film making, in which precise illuminant control is needed.

Finally, we note that the range of illuminants and surfaces we have used to demonstrate the system is modest. We continue to acquire new test images, but at present we

can only say with confidence that the method has worked very well on the images we have acquired in indoor conditions within fairly small (room size) spaces.

Summary

We have described a novel active imaging method for estimating the ambient illumination in a scene. The method introduces light into the scene, and the reflected light is used to estimate object surface reflectances and the ambient illuminant. Initial experiments, using a conventional camera and flash system, yield accurate estimates of the ambient illuminant across a range of test surfaces and conventional illuminants. The preliminary performance is comparable to that found using much more complex image processing methods⁹. Beyond simplicity, the method has two additional advantages: it measures the scene and thereby avoids making strong assumptions about the image contents, and it provides a space-varying estimate of the illuminant.

Acknowledgments

We thank Ben Olding, Ulrich Barnhoefer, Peter Catrysse, Julie DiCarlo and Abbas El Gamal for their help, and Pixim Inc. for use of the light booth. This work was supported by the Programmable Digital Camera project at Stanford whose founding members are Agilent, Canon, Kodak, and Hewlett-Packard. Jeffrey DiCarlo is a Kodak Fellow.

References

1. R. W. G. Hunt, *The Reproduction of Colour*, Fountain Press, Tolworth, England (1995).
2. G. Buchsbaum, A Spatial Processor Model for Object Color Perception, *J Franklin Institute*, 310, 1 (1980).
3. G. D. Finlayson, P. M. Hubel, and S. Hordley, Color by correlation, in *1997 5th Color Imaging Conference: Color Science, Systems, and Applications*, Soc Imaging Sci Technol Springfield VA USA, Scottsdale, AZ, USA, pp. 6 (1997).
4. M. H. Brill, A device performing illuminant-invariant assessment of chromatic relations, *J. Theor. Biol.*, 71, 473 (1978).
5. M. H. Brill, and G. West, Contributions to the theory of invariance of color under the condition of varying illumination, *J. Math. Biology*, 11, 337 (1981).
6. D. A. Forsyth, A Novel Algorithm for Color Constancy, *International Journal of Computer Vision*, 5, 5 (1990).
7. K. Barnard, L. Martin, and B. Funt, Colour by correlation in a three dimensional colour space, in *6th European Conference on Computer Vision*, Springer-Verlag, pp. 275 (2000).
8. D. H. Brainard, and W. T. Freeman, Bayesian color constancy, *J Opt Soc Am A*, 14, 1393 (1997).
9. S. Tominaga, S. Ebuisi, and B. A. Wandell, Scene illumination classification: brighter is better, *Journal of the Optical Society of America, A*, 18, 55 (2001).
10. G. Lighthart, and F. C. A. Groen, A Comparison of Different Autofocus Algorithms, in *Proc. of IEEE Int. Conf. on Pattern Recognition*, pp. 597 (1982).
11. S. Ando, and A. Kimachi, Time-Domain Correlation Image Sensor: First CMOS Realization of Demodulator Pixels Array, in *Proc. '99 IEEE CCD/AIS Workshop*, Karuizawa, pp. 33 (1999).
12. F. Xiao, J. M. DiCarlo, P. B. Catrysse, and B. A. Wandell, Image analysis using modulated light sources, in *Proceedings of the SPIE Electronic Imaging 2001 conference*, Vol. 4306, San Jose, CA (2001).
13. B. K. P. Horn, Exact reproduction of colored images, *Computer Vision, Graphics and Image Processing*, 26, 135 (1984).
14. B. A. Wandell, The synthesis and analysis of color images, *IEEE PAMI, PAMI-9*, 2 (1987).
15. J. M. DiCarlo, and B. A. Wandell, Illuminant Estimation: Beyond the Bases, in *Eighth Color Imaging Conference: Color Science, Systems, and Applications*, IS&T, Scottsdale, AZ, pp. 91 (2000).
16. L. T. Maloney, Evaluation of linear models of surface spectral reflectance with small numbers of parameters, *J. Opt. Soc. Am. A*, 3, 1673 (1986).
17. J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, Characteristic spectra of Munsell colors, *J. Opt. Soc. Am.*, 6, 318 (1989).
18. S. Westland, J. Shaw, and H. Owens, Colour statistics of natural and man-made surfaces, *Sensor Review*, v.20.50 (2000).