4-Sensor Camera Calibration for Image Representation Invariant to Shading, Shadows, Lighting, and Specularities

Graham D. Finlayson
School of Information Systems
The University of East Anglia
Norwich,
England NR4 7TJ
graham@sys.uea.ac.uk

Mark S. Drew
School of Computing Science
Simon Fraser University
Vancouver, British Columbia
Canada V5A 1S6
mark@cs.sfu.ca

Abstract

Most lighting can be accurately modeled using a simplified Planckian function. If we form logarithms of color ratios of camera sensor values, then in a Lambertian plus specular two-lobe model of reflection the temperature-dependent term is separate and is seen as a straight line: i.e., changing lighting amounts to changing each pixel value in a straight line, for a given camera. Here we use a 4-sensor camera. In this case, forming color ratios reduces the dimensionality to 3. Applying logarithms and projecting onto the plane in the 3D color space orthogonal to the light-change direction results in an image representation that is invariant to illumination change. For a given camera, the position of the specular point in the 2D plane is always the same, independent of the lighting. Thus a camera calibration produces illumination invariance at a single pixel. In the plane, matte surfaces reduce to points and specularities are almost straight lines. Extending each pixel value back to the matte position, postulated to be the maximum radius from the fixed specular point, at any angle in the 2D plane, removes specularity. Thus images are independent of shading (by forming ratios), independent of shadows (by making them independent of illumination temperature) and independent of specularities. The method is examined by forming 4D images from hyperspectral images, using real camera sensors, with encouraging results.

1. Introduction

The association of specularity detection with 4-band images has a long history (cf. [3]). Here we wish to accomplish a great deal more than suppression of specularities,

however, but again using a 4-band image acquired using a 4-sensor imaging device. We wish to recover an image representation which is, firstly, based on ratios of color bands, a type of chromaticity, so that image components which are derived from a Lambertian model are independent of shading and illuminant direction. As recognized by Lee [7], in a dichromatic model with neutral interface model for specular color since color equals the matte color plus a contribution deriving from specularities given by the illuminant color, chromaticity plots consist of one point for the matte component of each surface, with a line indicating specular content. All such lines point to the specular color.

Though the structure of specularities is thus simple to model in chromaticity space, finding the specular color and removing specularities from images is hard to do [7]. One way round this problem is to assume a known white illuminant [6]. In this case, the specular color is always found in the same place and as such it is a simple matter to derive a function of chromaticity that is independent of lighting geometry and specularity.

Of course, we would prefer not to constrain the light color. This is in fact what we achieve in this paper.

Here we show how, in a 4-sensor camera system, it is possible to make the specular color *independent of lighting*. To do so, we adopt a near-Planckian model for lighting, which describes many lights extremely well. Then taking logs of color ratios, the lighting color itself appears as a simple additive term, times the inverse temperature. If we start with color 4-vectors, and then divide by the geometric mean of the four values, we have independent color values given by 3-vectors. Thus the lighting direction part that depends on temperature is itself a 3-vector that depends only on the camera used.

By projecting orthogonal to the light direction we produce an image representation independent of the light. Thus we end up with a camera calibration that removes illumina-

tion dependence and also drives all images into the same lighting, with the same specular point.

Previously, Finlayson et al. [4] utilized RGB log color ratios to remove dependence on illuminant color for matte images: there, starting with 3-vector color, forming ratios reduces the vector values to 2-vectors. Then in the plane, changes in Planckian lighting temperature generate approximate lines in 2-space. The single direction orthogonal to this set of lines gives a grayscale image, independent of shading and of illumination temperature, and hence resistant to shadowing. This was shown to be highly useful for image indexing.

Here, we start with 4 dimensions. Forming logs of ratios reduces the dimensionality to 3, and projecting into the plane orthogonal to the lighting direction (fixed, for a particular camera), gives generalized color 2-vectors. Here we consider specularities as well. In the 2-space, specularities are approximately linear streaks pointing to the specular point. Therefore we suggest the simple expedient of replacing each 2D color by the maximum 2-vector position at its direction from the specular point: pixels that include matte plus specular contributions are driven out to the matte point. Although the generalized color images achieved this way are in fact still two-dimensional, we can rotate back into the original 3D color space of generalized colors (logs of color ratios) and produce RGB color images that are tied to the original generalized color directions. Resulting images are approximately independent of shading, shadows, lighting, and specularities.

Note here that we acquire all four bands at once, and do not use four separate photometric images (cf. [3, 8]). The use of the logarithm of color ratios is not new: Berwick and Lee used a 2D color space of logs of band ratios, taking 3D color into 2D chromaticities [2]. Under the assumption that a diagonal transform models illumination change, in 2-space a lighting change is modeled by a translation. In a sense, the present work extends that idea but *explicitly* calculates what a particular camera would see, under lighting change. Using higher dimensionality for sensors, we can approximately eliminate lighting change directly, rather than searching for it.

Gevers et al. classify pixels into highlight and matte but making the assumption that illumination is equi-energy white [6]. However, illumination which is “white” may appear to a camera anything but white — e.g., tungsten lighting is quite red.

Since we do not in fact have available a 4-sensor camera, we develop images to test the ideas presented here by synthesizing 4-band images from hyperspectral images, employing real camera sensors. However, we point out that 4-sensor cameras do exist: some video cameras have Cyan, Magenta Yellow and Green sensors. So assuming these sensitivities are linearly independent, the methods presented

here are directly applicable.

Here we use the measurements presented by Brelstaff and colleagues [10]. These images are subject to all the noise that would accompany captured images, along with additional problems due to image registration and lens corrections. Results are quite good: we find that shading, shadows, and specularities are greatly attenuated.

2. Theory and Proof in Principle

2.1. Planckian Lighting and Two-Lobe Reflection Model

Suppose we consider a fairly narrowband camera, with four sensors. We actually do not have a 4-sensor camera, so instead we use the measured sensors for a real camera and replicate and shift one of the sensor curves. Here we use a Sony DXC-930 3-CCD camera. Fig. 1 shows the sensor responses for this camera; however since we require a 4-sensor camera, we augment the sensor set by replicating the Blue sensor response, multiplied by 0.75 and shifted to the right by 70 nm. Below, we investigate how the Planckian assumption impacts models of image formation using a 4-sensor delta-function sensitivity camera. However, it is evident that the real sensors are far from idealized delta functions: each is sensitive to a wavelength interval over 100nm in extent. Nevertheless, as we shall see, they behave sufficiently like narrowband sensors for our theory to work.

To be able to fully understand results of our assumption, let us start with a synthesized image. For a simple test image consisting of three shaded spheres, with surface colors equal to patches 1, 4, and 9 (dark skin, moderate olive green, moderate red), of the Macbeth ColorChecker [9], under illuminant D65 the matte image in a Lambertian model is as in Fig. 2(a), where we have normalized sensor curves so that they sum to unity — the curves are color balanced for equi-energy white light. Here, the three components corresponding to RGB are shown, but the image is actually 4-dimensional (for full-size images see [1]).

Now consider adding a very strong specular spike for each surface. Using a simple Phong specular model with roughness 0.05 and multiplicative factor 1.0, one sees Fig. 2(b). For visualization, consider a histogram-equalized version, Fig. 2(c). To see how strongly the specular spike affects pixel values, Fig. 2(d) shows the red channel response, for the first sphere.

Now let us consider image formation for a general reflectance function comprising Lambertian and specular parts. For the Lambertian component: suppose there are $i = 1..L$ lights, each with SPD $E^i(\lambda)$ given by a good approximation of a Planckian source [15], with

$$E^i(\lambda) = I^i c_1 \lambda^{-5} e^{-c_2/(\lambda T)} \quad (1)$$

in directions $\mathbf{a}^i$ with intensity I^i . If the surface projecting to retinal point x has spectral surface reflectance $S(\lambda)$ then, for a 4-channel delta-function narrowband sensor camera with spike sensor sensitivities $q_k = q_k(\lambda_k)$, $k = 1..4$, the 4-vector response would be

$$\begin{aligned}\rho_k &= \sum_{i=1}^L c_1 \mathbf{a}^i \cdot \mathbf{n} S(\lambda_k) I^i \lambda^{-5} e^{-c_2/(\lambda_k T)} q_k \\ &= \left[\sum_{i=1}^L (c_1 I^i \mathbf{a}^i) \right] \cdot \mathbf{n} S(\lambda_k) \lambda^{-5} e^{-c_2/(\lambda_k T)} q_k \\ &\equiv \tilde{\mathbf{a}} \cdot \mathbf{n} S(\lambda_k) \lambda^{-5} e^{-c_2/(\lambda_k T)} q_k \quad .\end{aligned}\quad (2)$$

For the specular component: Suppose the surface's reflectance map falls in the m -lobed class [13]. This class consists of the constant, plus $j = 1..(m-1)$ terms of the form $\Phi^j(\mathbf{p}_j^T \mathbf{n})$ with Φ^j a non-negative, monotonically decreasing function. Further, the vector $\mathbf{p}_j$ must lie in the *principal plane* formed by the light direction and the vector $\mathbf{n}_S$ half-way between the illuminant direction and the viewer. E.g., the Torrance-Sparrow model [14] can be written in terms of a function with $\mathbf{p} \equiv \mathbf{n}_S$. Clearly the Lambertian model also falls into the m -lobed class, with $\mathbf{p}_j \equiv \mathbf{a}$, the light source direction.

For color images, consider a simple dichromatic model [11], whereby reflectance is comprised of a diffuse or body term plus a surface (specular) term. The specular term is effectively 1-lobe if all the lights have the same SPD: in general,

$$\rho^{\text{Specular}} = \sum_{i=1}^L \mathbf{b}_S^i \Phi(\mathbf{n}_S^i{}^T \mathbf{n}) \quad , \quad (3)$$

where $\mathbf{b}_S^i$ is the color of the specularity for the i^{th} light. In the Phong specular model [5],

$$\Phi(\mathbf{n}_S^i{}^T \mathbf{n}) = (\mathbf{n}_S^i{}^T \mathbf{n})^N \quad . \quad (4)$$

But in a neutral interface model [7], the color of the specular term is

$$\mathbf{b}_S^i \equiv \text{color of the light} \quad (5)$$

Hence for Lambertian plus Specular reflectance, we arrive at a 2-lobe model:

$$\rho_k = \left[\tilde{\mathbf{a}} \cdot \mathbf{n} S(\lambda_k) + \sum_{i=1}^L c_1 I^i \Phi(\mathbf{n}_S^i{}^T \mathbf{n}) \right] \lambda^{-5} e^{-c_2/(\lambda_k T)} q_k \quad . \quad (6)$$

For each pixel, the second term in the brackets is a constant, β , say, that depends only on geometry and not on the light color. Therefore we have

$$\rho_k = [\tilde{\mathbf{a}} \cdot \mathbf{n} S(\lambda_k) + \beta] \lambda^{-5} e^{-c_2/(\lambda_k T)} q_k \quad (7)$$

with possibly several specular spikes on any surface.

If we define

$$\alpha = \beta/(\tilde{\mathbf{a}} \cdot \mathbf{n}) \quad , \quad (8)$$

then our final expression is quite simple:

$$\rho_k = (\tilde{\mathbf{a}} \cdot \mathbf{n}) [S(\lambda_k) + \alpha] \lambda^{-5} e^{-c_2/(\lambda_k T)} q_k, \quad k = 1..4 \quad (9)$$

2.2. Log Color Space

For generality, we consider p -channel sensor values; the above equation holds, independent of the number of camera channels.

Now consider band-ratio chromaticities r_k defined as the ratio of the first $p-1$ sensor values ρ_k , $k = 1..(p-1)$, divided by the last sensor value ρ_p . Here, we use $p = 4$, so our chromaticities are 3-vectors:

$$r_k = \frac{\rho_k}{\rho_p} \quad , \quad k = 1..(p-1) \quad . \quad (10)$$

Since we divide, the shading term $(\tilde{\mathbf{a}} \cdot \mathbf{n})$ is removed. Taking logarithms, we can concisely state the result by defining several terms:

$$\begin{aligned}e_k &= -c_2/\lambda_k \quad , \quad s_k = S(\lambda_k) \quad , \\ u_k &= \lambda_k^{-5} q_k \quad , \quad w_k = \log(u_k/u_p) \quad .\end{aligned}\quad (11)$$

Then we have

$$\log r_k = \log \left(\frac{s_k + \alpha}{s_p + \alpha} \right) + w_k + (e_k - e_p) \frac{1}{T} \quad , \quad k = 1..(p-1) \quad . \quad (12)$$

The meaning of the above expression is that, for every surface patch, with geometric parameter α , the $(p-1)$ -vector is formed as a constant vector plus a vector $(e_k - e_p)$ times the inverse temperature: as lighting changes (i.e., as temperature T changes), pixel values are constrained to a straight line in log color space, as in Fig. 3(a).

2.3. Broadband Sensors

For broadband sensors the p -sensor color must be formed by integration, since our sensors are no longer delta functions:

$$\rho_k = (\tilde{\mathbf{a}} \cdot \mathbf{n}) \int [S(\lambda) + \alpha] \lambda^{-5} e^{-c_2/(\lambda T)} q_k(\lambda) d\lambda \quad (13)$$

with p sensor functions $q_k(\lambda)$, $k = 1..p$.

While we maintain the calculation of ratio values as in eq. (10) and calculate log values $\log r_k$ as above as well, in principle we now replace the definitions of parts of the log ratios by:

$$\begin{aligned}\sigma_k &= \int q_k(\lambda) d\lambda \quad , \\ e_k &= 1/\sigma_k \int -c_2/\lambda q_k(\lambda) d\lambda \quad , \quad s_k = 1/\sigma_k \int S(\lambda) q_k(\lambda) d\lambda \quad , \\ u_k &= \int \lambda_k^{-5} q_k(\lambda) d\lambda \quad , \quad \text{and } w_k = \log(u_k/u_p) \quad .\end{aligned}\quad (14)$$

2.4. Geometric Mean Log Color Space

Division by a single color channel can be replaced by division by the geometric mean of all p channels, with increase in robustness to noise. Suppose we define the mean ρ_M by

$$\rho_M = \sqrt[p]{\prod_{k=1}^p \rho_k} \quad . \quad (15)$$

Then, with $r_k = \rho_k / \rho_M$, from eq.(9) we have

$$\log r_k = \log \left(\frac{s_k + \alpha}{s_M + \alpha} \right) + w_k + (e_k - e_M) \frac{1}{T} \quad (16)$$

for $k = 1..(p-1)$, where we set

$$\begin{aligned} e_M &= -c_2/p \sum_{j=1}^p \lambda_j, \quad u_M = \{\Pi_{j=1}^p \lambda_j^{-5} q_j\}^{1/p}, \\ w_k &= \log(u_k/u_M) \quad \text{and we define} \\ s_M + \alpha &\equiv \{\Pi_{j=1}^p (s_j + \alpha)\}^{1/p}. \end{aligned} \quad (17)$$

We shall use this version in the following.

2.5. Recovery of Log Light Vector

To find the vector $(e_k - e_M)$, we image matte Lambertian color patches. Here we use the 18 non-gray patches of the Macbeth ColorChecker [9]. We form $\log r_k$ values using temperatures T from 5500°K to 10500°K.

According to eq. (16) (with no specular contribution), for each surface we should see a set of points in 3-space that falls on a straight line along $(e_k - e_M)$. Thus for each surface, if we then subtract the mean in each channel of $\log r_k$, we see a set of nearly coincident lines through the origin, as in Fig 4 (plotted as 2D projections).

Therefore, as pointed out in [4] in a 2D setting, we can find vector $(e_k - e_M)$ by forming the covariance matrix of $\log r_k$ values and calculating eigenvectors. The first eigenvector is exactly the desired approximation of $(e_k - e_M)$.

3. Recovery of Image with no Shading or Specularities

3.1. Removal of Illumination Dependence

As a generalization of Finlayson and Hordley [4], we can remove the effects of lighting temperature, and hence of shadows, by projecting in $(p-1)$ -space onto the plane orthogonal to the vector $(e_k - e_M)$. Similarly, all the other vectors in eq.(12) are $(p-1)$ -vectors. If $p = 4$, then, all vectors live in a 3-space.

Suppose the projector onto the vector $(e_k - e_M)$ itself is

$$P_e = (e_k - e_M)(e_k - e_M)^T / \|e_k - e_M\|^2. \quad (18)$$

This is a $(p-1) \times (p-1)$ matrix. Then the projector onto the subspace orthogonal to $(e_k - e_M)$ is [12]

$$P_e^\perp = I - P_e \quad (19)$$

Such a projector has only $(p-2)$ nonzero eigenvalues; for our $p = 4$ sensor value space, we are projecting onto a 2-space.

We can rotate the coordinate frame into the plane by forming the eigenvector decomposition of the projector $P_e^\perp$:

$$P_e^\perp = (U_e^\perp)^T \text{diag}(1, 1, 0) U_e^\perp. \quad (20)$$

Then $U_e^\perp$ is a rotation matrix taking us into the plane; we make use of only the 2×3 portion of $U_e^\perp$, in order to discard components in the $(e_k - e_M)$ direction. This is the same as using the projector $P_e^\perp$: the projector equals transforming into the plane coordinate frame using the 2×3 matrix $U_e^\perp$ and then transforming back to the original coordinate frame using the 3×2 matrix $(U_e^\perp)^T$.

This is just what is wanted because as much as possible, we would like to continue associating the ratios s_k/s_M with the *original* directions of these quantities, so that when we come to reconstitute images, pixels' colors are as close as possible in meaning to the actual colors we started with. Rotating back with a 3×2 matrix $(U_e^\perp)^T$ will accomplish this: although we still only have values that derive from the plane, and so are not truly 3D, we can view the pixel values as 3D colors nonetheless. Therefore we choose some set of basis vectors in the plane, and project into the plane — thus removing dependence on illumination. Then we rotate back into the original first $p-1$ generalized color axes, as in Fig. 3(b).

Fig. 5(a) shows how projected log ratios for Fig. 2(b) appear in the 2D space resulting from projecting into the plane orthogonal to vector $(e_k - e_M)$. Now all similarly colored *matte* points project onto the same point in 2-space. However, independent of temperature T , *specularities* are spread out in 2-space according to

$$U_e^\perp \log \left(\frac{s_k + \alpha}{s_M + \alpha} \right) + U_e^\perp w_k. \quad (21)$$

The first term above is the log of a straight line. However, most of the points on the line will concentrate around the matte point: for pixels with very little specular content, the projected 2-vector is at the matte point, shown in Fig. 5(a) for each object as 'M'.

For highly specular points, the projected 2-vector is near the specular log color, shown as 'S' in Fig. 5(a). The specular color is the projection of the 3-vector $w_k + (e_k - e_M)/T$, so the projected specular color is $U_e^\perp w_k$. Note that *the projected specular point is the same no matter what the lighting is*, provided the illumination is well-described by a Planckian model. Matte colors are given by the projection of 3-vectors $\log(s_k/s_M)$.

Most points that make up the observable curve in the log of the straight line $(s_k + \alpha)/(s_M + \alpha)$ are due to pixels in the highly specular part of the image. For these, the curve asymptotically approaches values

$$\log[(s_k + \alpha)/(s_M + \alpha)] \rightarrow 1/\alpha (s_k - s_M) \quad (22)$$

as $(1/\alpha)$ becomes small, so that near the specular point the curve approaches a straight line even in log space.

As we shall see later, the specular curves tend to be fairly straight even quite far from the specular point. This may not be altogether surprising if we examine (22) a little closer. Value s_M is the geometric mean of the p -sensor responses. If for a moment we let $p=3$ and so consider the conventional RGB camera case we see that $s_M = 0$ when $R=0$, or $G=0$, or $B=0$, and hence when R and $G=0$, or R and $B=0$, or G and $B=0$. It is easy to show that for these cases the specular line is a line in log and non-log space. The specular curves for other surfaces are constrained to lie between the straight specular lines for the basic colors and this effectively bounds the maximum curvature of the lines.

For our 4-sensor analysis, there are 14 basic colors (i.e., $2^4 - 2$) which have provably straight specular lines in log and non-log space. Because all other specular curves lie between pairs of basic specular lines the extent of the possible specular curvature, the possible departure from linearity, is very limited indeed.

3.2. Recovery of Matte Image Representation

We have calibrated for the projector $P_e^\perp$ simply using matte patches (as in Finlayson and Hordley [4]). As well, one can obtain the projection of the w_k vector easily by (1) imaging a light source, and then (2) projecting onto 2-space the result of forming the 3-vector log chromaticity. However, even without subtracting $U_e^\perp w_k$, but instead just projecting log chromaticities $\log r_k$, one can recover an image representation free from specularities. For our simple image, clearly the right course is to attempt to get as close to point marked ‘M’ as possible, by lengthening vectors in the direction from the specular point out to a maximum radius from that point.

If we did so for a maximum overall radius, we would arrive at approximately three spots on the maximum-radius circle. This way, we end up with a 2-vector representation: the two components are of course not independent, since the invariant is in fact the angle formed by each of the lines in Fig. 5(a). Thus a *grayscale* specularity-free image can easily be derived as the image made from angles formed by vectors from the specular point to log color ratio values. This would be in keeping with the grayscale invariant images produced in [4]. However since we do have 3D color correlates here, we instead choose to maintain a 3D space.

A better approach than using a single circle would be to find the maximum radius for every small increment of angle. For our simple image, this means finding three circles, as in Fig. 5(b).

For a more complex image, the idea of using radii simply amounts to finding the maximum radius at each angle for all image points in the 2D log color space, and then replacing

each radius at that angle with the maximum. This means representing every pixel at the same angle from the specular center by the same log chromaticity — all points with the same color end up at the same radius.

Since we are still in a log color space, it makes sense to exponentiate the $P_e^\perp \log r_k$ values, in order to return to color units. The resulting image is as in Fig. 5(c) — the $(U_e^\perp)^T$ rotation takes us back to 3D colors. Using the maximum-radius projected log ratio vector to characterize every pixel of the same ratio, for our simple exemplar image we use three circles, as in Fig. 5(b). Fig. 5(c) shows the recovered generalized color image, almost without shading or specularities. The worst sphere is the middle one, with red channel shown in perspective in Fig. 5(d). Specularities have been reduced to a remarkable degree.

4. Real Images

4.1. Use of Hyperspectral Images

Since we are using a theoretical camera, derived from the Sony DXC-930 camera, we can devise real images by making use of hyperspectral images acquired using multiple narrowband filters, and synthesize a real, but composed image that approximately comprises what the camera in Fig. 1 would record from a real scene.

Here we use images derived from 31-band hyperspectral image captures due to Brelstaff et al. [10]. Fig. 6(d) shows a color image as it would appear for a camera balanced for the actual illumination, which is tungsten lighting (very similar to a Planckian with $T=3000$). In contrast, Figs. 6(a,b) show the image as it would be acquired by our 4-sensor camera, in the first three bands and the fourth band, respectively. The lighting is very red, but the camera is balanced for equilibrium white light, so the acquired image is very red.

The image has a Kodak Graycard at one corner: this approximately reflects 18% radiance. Fig. 6(c) displays the radiance spectrum reflected from the Graycard: a Planckian with $T=3000$ is also shown as a solid curve.

Fig. 7(a) shows the pixels for the 4-band image, in terms of their log-ratios projected onto the 2D color space. The figure shows the recovered specular point via crosshairs; the dashed crosshairs show the position of the projected components of the Kodak Graycard.

In terms of a 1D angle image, these projected values, as vectors from the specular point, show angles as in Fig. 7(b). Here pixels with radii less than the 10th quantile away from the specular point have indeterminate angle and have been re-mapped to the specular point. In comparison, Fig. 7(c) shows the histogram-equalized version of Fig. 6(b). Clearly, the 1D angle image which attempts to remove shading, shadows, and specularities does a superior job in regards to image understanding.

Finally, Fig. 7(d) shows the exponentiated color image when each generalized color is forced to the maximum value, for the particular angle that the generalized color makes with the specular point. This image has successfully almost removed the effects of shading, shadows, and specularities, and as well is invariant to changes in lighting.

Pixels that are very close to being specular in color must necessarily be mapped to the specular point. E.g., in the original 4D image the graycard has the chromaticity of the illuminant. Figs. 8(a,b) show in detail the specularity on a leaf, in the four channels. In Fig.8(c), we plot color values for pixels in the subimage, in blue circles, along with the colors observed from the graycard, shown as red points. We observe that there are pixels from the leaf that have colors indistinguishable from the specular, and these are necessarily mapped to the specular generalized color rather than to the matte generalized color.

Nevertheless the result is very successful: shadows and specularities have been greatly attenuated, and we are left with a representation of image matte values independent of these artifacts, as well as of shading. Such an image is precisely what is needed for database search. Fig. 9 shows further results, all about as good as the one studied in detail above.

5. Conclusions

Images without shading, shadows, or specularities are very useful in database search (cf. the excellent results achieved in [4] using grayscale images similar to the grayscale angle images we produce here). As well, generalized color representations of images may be useful for invariant characterization of the intrinsic data in the image. Shading information, specularities and shadows are recoverable, so it should be possible to reconstruct images of faces, say, under different lights than used in image acquisition.

Here we decided to drive every pixel back towards the matte point by replacing 2D generalized colors in log space by the generalized color at the maximum radius for a given angle from the specular point. Clearly other schemes may be possible. We did find that the algorithm seems to be fairly independent of the size of angular wedge chosen in which to find a maximum radius. We usually used wedges of size 4 degrees. Results show that the algorithm operates as a kind of automatic segmentation routine: colors are correctly generated for parts of flowers etc. hiding in the background that the eye can easily miss.

Of course, the present analysis can be simply extended to higher dimensional sensor sets. The truly invariant part of the present 4D→1D analysis is the angle in the 2D plane; or, in other words, a circle of generalized colors in that plane. If we started with 5D sensors instead, then the set of

generalized colors recovered would be essentially 2D, i.e., they would live on a sphere. And starting with 6D sensors, we would recover a fully 3D invariant image. The basic contribution of the paper is that, without needing to know the sensor curves themselves, we can calibrate a higher-dimensional camera to eliminate lighting artifacts.

References

- [1] Full-size images may be viewed at <http://www.cs.sfu.ca/~mark/lccv01>.
- [2] D. Berwick, S.W., and Lee. A chromaticity space for specularity-, illumination color- and illumination pose-invariant 3-d object recognition. In *ICCV98*, pages 165–170, 1998.
- [3] E. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surface using four-source photometry. *Comp. Graphics and Image Proc.*, 18:308–328, 1982.
- [4] G. Finlayson and S. Hordley. Color constancy at a pixel. *J. Opt. Soc. Am. A*, 2001. To appear. Also, UK Patent application no. 0000682.5. Under review, British Patent Office.
- [5] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley, 2nd edition, 1990.
- [6] T. Gevers and H. M. G. Stokman. Classifying color transitions into shadow-geometry, illumination, highlight or material edges. In *ICIP 2000*, pages 1521–1524. IEEE, 2000.
- [7] H. Lee. Method for computing the scene-illuminant chromaticity from specular highlights. *J. Opt. Soc. Am. A*, 3:1694–1699, 1986.
- [8] S. Lin, S.W., and Lee. Estimation of diffuse and specular appearance. In *ICCV99*, pages 855–860, 1999.
- [9] C. McCamy, H. Marcus, and J. Davidson. A color-rendition chart. *J. App. Photog. Eng.*, 2:95–99, 1976.
- [10] C. Parraga, G. Brelstaff, T. Troscianko, and I. Moorhead. Color and luminance information in natural scenes. *J. Opt. Soc. Am. A*, 15:563–569, 1998.
- [11] S. Shafer. Using color to separate reflection components. *Color Res. Appl.*, 10:210–218, 1985.
- [12] G. Strang. *Linear Algebra and its Applications*. Harcourt, Brace, Jovanovich, 3rd edition, 1988.
- [13] H. Tagare and R. deFigueiredo. A framework for the construction of general reflectance maps for machine vision. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 57:265–282, 1993.
- [14] K. Torrance and E. Sparrow. Theory of off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.*, 57:1105–1114, 1967.
- [15] G. Wyszecki and W. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley, New York, 2nd edition, 1982.

Figure 1. Four-sensor broadband camera.

(a)

(b)

(c)

Figure 2. (a): Matte image, under (first three of a set of four) broadband sensors. (b): Image including specularities, for first three sensors. (c): Image with specularity, histogram equalized. (d): Lambertian plus specular.

(b)

Figure 3. (a): Pixels, including those with specular content, are formed by adding a vector in the $(e_k - e_M)$ direction. (b): Points are made independent of lighting by projection onto the plane orthogonal to $(e_k - e_M)$; the original generalized color directions are kept the same by rotating back into the original log color space directions.

Figure 4. Calibration using Lambertian patches: mean-subtracted $\log r_k$ lines are nearly coincident (here 2D projections are plotted).

(a)

(b)

Figure 5. (a): Projected 2D log color ratios. (b): Using the maximum log ratio radius in each direction in log space.

(c)

(d)

Figure 5. continued. (c): Recovered generalized color image without shading, specularities, and independent of lighting. (d): Recovered Red generalized color image, for second sphere.

(a)

(b)

(c)

(d)

Figure 6. (a): Synthetic image derived from hyperspectral images: RGB = first three bands. (b): Fourth band. (c): Radiance from Kodak Graycard: solid line shows Planckian illumination at $T=3000$. (d): Image as it would be acquired by a color-balanced RGB camera.

Figure 7. (a): Log ratio pixel values, projected into 2D color space. (b): Angle image for log ratios, calculated from specular point. (c): For comparison, the histogram-equalized version of Fig. 6(b). (d): Generalized color image, with each generalized color driven to the maximum at each angle from the specular point.

Figure 8. (a,b): Sub-image for a specularity, in the four channels. (c): Pixels values (for two channels): Specularity in blue circles, Graycard in red points— some specularity pixels are identical to graycard points.

(a)

(b)

(c)

(d)

(e)

(f)

Figure 9. RGB images and corresponding invariant images derived from 4-sensor data.

(g)

(h)

(i)

(j)

(k)

(l)

Figure 9. continued. RGB images and corresponding invariant images derived from 4-sensor data. (Full-size images may be viewed at [1].)