Eurographics Symposium on Rendering (2004)
A. Keller, H. W. Jensen (Editors)

Estimating source spectra and spectral albedos from RGB data for rerendering

J. J. Koenderink

Department of Physics and Astronomy, Universiteit Utrecht, The Netherlands

Abstract

I consider the problem of estimating material properties (the spectral albedo) on the basis of "object colors" (at worst only RGB data say). I show how to obtain a priori likely estimates for the white point, the spectral composition of the source, and the spectral albedos of the objects in a scene. I also show how to construct the general solutions. These general solutions are so broad as to render them practically useless. There are good reasons to disregard the larger part of the solution space, because very general considerations suggest that the specific solutions constructed with the methods discussed here are very likely to yield sensible and useful results in practice. From a principled perspective it is desirable to be able to construct the full solution space though. Since the results are in the scene, rather than the image domain, they are suitable for rerendering purposes.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three–Dimensional Graphics and Realism

1. Nature of the problem

The potential image data at a fixed vantage point are fully specified by the spectral radiance as a function of direction and resolution. Sampling the radiance yields "pixels", where each pixel has the value of a (continuous) spectral photon number flux density. Sampling the spectrum with a CCD chip yields a vector, referred to as the "color" or the "RGB" value. The color is a 3D projection of an ∞D Hilbert space of spectra. The spectral albedo is a material property. It specifies the fraction of photons scattered and the fraction absorbed. I assume (an extreme simplification of reality) that the scattering is Lambertian. The colors depend multiplicatively on the spectral density of the source, the spectral albedo of the materials and on the nature of the projection to RGB space.

Suppose one samples any number of surface colors. These colors will be due to the spectral albedos, the spectral density of the source, and effects of shading. I assume that multiple scattering (which introduces nonlinear mixture of the beams) can be neglected. Can one estimate the source and the albedos? The importance of the problem is obvious: Suppose one shuffles the objects and changes the source, how can one retrace the objects by their material properties, their "spectral signatures"? For rerendering purposes one needs estimates of the spectral reflectances, how does one obtain such estimates given only an image?

The problem can be subdivided into a number of conceptually simpler problems. For instance, instead of asking for the *spectrum* of the source one may ask for its *color*. The problem can also be stratified in terms of the objectives, does one need *some* solution (*any* solution or perhaps the *most probable* solution in some sense) or *all* solutions?

In this paper I consider specific solutions of practical value as well as the complete solution. The latter is (of course) of conceptual interest, but of no practical value whatsoever. The real problem is how to select useful specific solutions from the (enormous) set of all solutions. The present method achieves this in a very direct way.

2. Estimation of the white point

Given a scene, what is the color of the source? In the context of the present setting this problem may be specified more precisely: what would be the color of a spectrally nonselective and maximally light ("white") surface irradiated normally by the source?

This problem is quite tractable. Any beam is a convex sum (coefficients—continuous of course—in the range zero to one) of the monochromatic beams that make up the beam remitted by a white surface. Thus all object colors must lie inside the region of intersection of the *spectrum cone* with vertex at the black point (origin) and the inverted spectrum cone with vertex at the white point in RGB space [Bou47,Sch20]. Here is an algorithm to estimate the white point (see figure 1):

— *Start* by finding the convex hull of the colors and discard all colors in its interior. Although this step is not necessary, it greatly improves efficiency.

— The white point must lie in the intersection of the half-spaces defined by planes parallel to tangent planes of the inverted spectrum cone, moved to each color, an ∞ volume. Pick the solution nearest the origin.

— This is a linear programming problem: *Done*.

In practice one approximates the cone with about a hundred tangent planes. Then the solution is found fast. Of course this is only one *particular* solution. *All* solutions lie (very nearly) inside the spectrum cone with this solution as vertex.

In the absence of prior knowledge the intensity of the source will be uniformly distributed on a logarithmic scale (no scale singled out). Then (by Bayes' theorem) the specific solution is the most probable one. In simulations I find that the chromaticity of the white point is recovered quite accurately from samples of ten or more random surface colors. In one experiment (I assume the CIE observer throughout in this paper) I based estimates upon sets of 16 samples drawn at random from the NCSU data base [NCSU]. The intensity is typically estimated up to fifty percent too low. This is indeed to be expected when all colors are darkish as they usually are (the "average" reflectance is like an 18% gray chart). The presence of only a single light color saves the solution. The *chromaticity* of the white point is recovered quite accurately, even when the intensity is too low (figure 2).

This method is free of unrealistic assumptions ("gray world assumption" say) and is very robust against sampling bias because it doesn't rely on averaging processes. In fact, 3 maximally bright colors of different hue in the sample guarentee a good estimate.

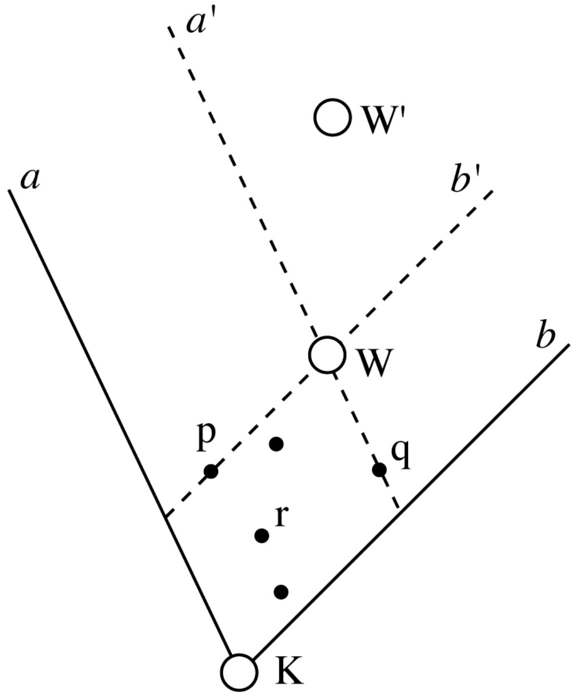

Figure 1: *A 2D schematic diagram that illustrates the estimation of the white point. Here* ab *is the spectrum cone surface with vertex at the black point K, whereas* a'b' *is the surface of the inverted spectrum cone at the estimated white point W. Notice that* a'b' *is just the inverted cone* ab. *The estimate W is the point nearest K such that the intersection of the two cones contains the gamut of colors (black dots). Notice that the colors p, q are relevant for the estimate, whereas colors such as r are not. The point W' is another possible solution.*

3. Estimation of the spectrum of the source

Suppose one knows the white point, what can one say concerning the spectrum of the source? The general solution is immediate. The colors are projections from the ∞D space of beams to 3D RGB space [Bou47]. Thus the pseudoinverse of the projection yields a *specific* solution and adding the kernel of the projection yields *all* solutions. Linear methods are involved, thus solutions are not necessarily *physically possible* though. For real beams the spectral density has to be nonnegative throughout. One has to prune the general solution. The required real solutions can be constructed with linear programming.

Although the general solution might appear to be the answer, it is actually of little use. Most of the solutions are freaks even when physically possible. In order to arrive at practical results one has to mind prior probabilities. Unfortunately, one can hardly put a prior on the source spectrum unless one has specific ecological knowledge.

A practical approach is to constrain source spectra via a truncated development into basis functions of graded complexity, for example, by linear combinations of a flat spectrum, a linear trend, and a spectral "curvature". Such sums fit both daylight and artificial light situations well, but avoid freak solutions. One easily finds a particular solution of this type by precomputing the colors of the basis functions. A basis transformation converts the white color to the coefficients of the combination. (See figures 3 and 4.) It appears to be next to impossible to improve on this fast and easy method. All existing methods boil down to this in the final analysis, although one typically uses principal components (PC's) of all sources in the universe [MW86]. Although this

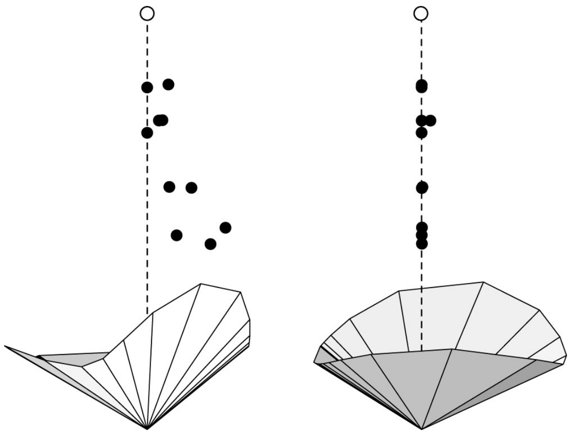

Figure 2: *Ten estimates of the white point from random samples of 16 items each from the NCSU data base, irradiated with average daylight (CIE D65). The true white point is the open circle, estimations the closed circles. On the left a projection on the yellow–blue plane, on the right on the red–green plane. Estimated intensities are systematically too low, chromaticities are quite accurate (the deviations have to be judged against the spread of the spectrum cone, the dashed line is the white chromaticity). The deviations are mainly in the yellow–blue dimension. When the estimated intensity is not too low these deviations are only slight.*

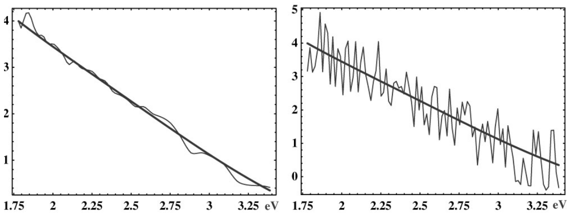

Figure 3: *On the left the average daylight (CIE D65) spectrum (thin line) and the spectrum computed from the white point (thick line). On the right another solution, obtained by adding a random black.*

seems a good idea it really isn't (although fairly harmless since the PC's are likely to resemble the simple basis functions closely). It is more important to obtain *as smooth a solution as possible* though, for any nonrealistic features will propagate to all further estimations.

4. Estimation of the spectral albedo

After estimation of the white point and the source spectrum one might attempt to estimate the spectral albedo of individual materials on the basis of their color. These estimates will also be very ambiguous, that is to say, one has to reckon with ∞D solution spaces. However, for very *vivid* colors the solution is very constrained, moreover, of the ∞D space of

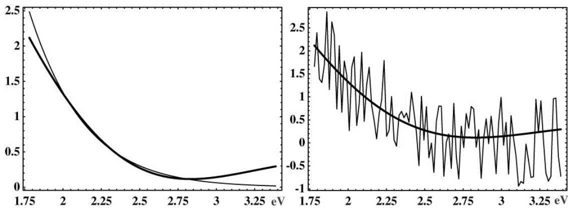

Figure 4: *On the left the Tungsten light (black body at* 3000°K*) spectrum (thin line) and the spectrum computed from the white point (thick line). On the right another solution, obtained by adding a random black, in this example physically impossible.*

solutions only a small subspace will represent "reasonable" solutions. Of course the latter distinction presumes a prior on the spectral albedo. I will assume the spectral photon number flux density of the source as "given". This will often be an estimate itself, and one has to compound the various ambiguities.

When the white point is fixed all object colors lie in the intersection of the spectrum cone at the black point and the inverted spectrum cone at the white point. This is a finite volume. It is limited through the finite photon number flux of the source and the nature of the spectral characteristics, especially their overlap, of the high, medium and low photon energy sensitive photoelements of the CCD device (these determine the shape of the spectrum cone). Not all points inside this double cone represent possible object colors though. For instance, a point on the spectrum cone would have to be due to a monochromatic beam. For an object color this entails that the source would contain a *monochromatic* component of *finite* power (like a laser beam). *No object color can be a spectrum color if the source has a continuous spectrum.* When the spectrum of the source is known one can calculate the volume occupied by all object colors [Bou47, Sch20]. It is a spindle shaped volume, its surface smooth throughout except for conical singularities at the white and black points. At these points the surface is tangent to the double cone, at all other places it lies strictly inside it. Like the double cone itself this "color solid" has central symmetry (see figure 5). For an average daylight spectrum and the CIE color matching functions one calculates that the volume of the color solid is about two thirds of the volume of the double cone.

Schrödinger [Sch20] showed that the spectral albedos for colors on the surface of the color solid (these are the brightest colors of a given chromaticity, known as "optimal colors") are *uniquely determined*. The spectral albedos of the optimal colors are of the following nature:
— the spectral albedo is either zero or one;
— there exist no more than two transitions throughout the spectrum.
There exist two generic types, namely the materials that scat-

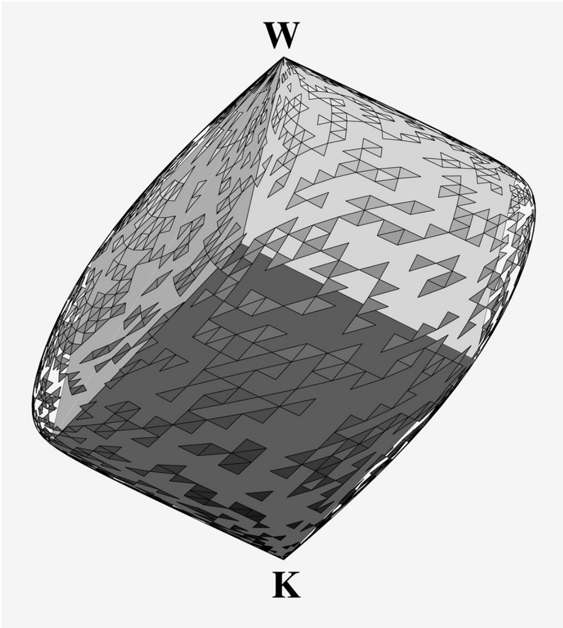

Figure 5: *The color solid for average daylight and human vision is rendered transparantly such that the inscribed parallelepiped of maximum volume can be seen. This "RGB–crate" exhausts much of the volume of the color solid. Since colors near the boundary of the color solid occur rarely in nature (they are close to the Schrödinger "optimal colors"), almost all scene colors fall within the crate. That is why a rough approximation of the color solid by an RGB–crate often works surprisingly well.*

ter only medium energy photons and absorb photons from both ends of the spectrum, and the materials that absorb photons of medium energy and scatter those from the ends of the spectrum. Roughly, the former look greenish, the latter purplish. These generic types occur in simply connected regions of the surface of the color solid. These regions meet at two singular curves that connect the black and white point and spiral over the surface of the color solid. Colors on these curves are the so called "boundary colors" (discovered by Goethe [Goe82]). One type scatters low energy photons and absorbs high energy ones, the other type scatters high energy photons and absorbs low energy ones. Roughly, the former look reddish, the latter bluish.

Although the optimal colors are the *brightest* for a given chromaticity (chromaticity is simply color modulo intensity, *i.e.*, RGB data considered as homogeneous coordinates of "chromaticity space"), they are not necessarily *vivid* colors. For instance, those near the black point look blackish and those near the white point whitish. Both are therefore almost "achromatic" in the intuitive sense of "colorless" (but watch out: those near the black point have very high saturations!). The most vivid colors occur on a curve that is the locus of optimal colors at maximum distance from the achromatic (white–black) axis. This curve is obtained when one circumscribes the color solid with a cylinder with generators parallel to the achromatic axis. The cylinder touches the solid along a curve, the curve of "full colors". The full colors are maximally vivid in the following sense: When one represents a color by the sum of a spectrum color (monochromatic beam) and an achromatic color (attenuated source), the amount of spectrum color is a maximum for the full colors [Hel1896, Gra1853]. The full colors scatter exactly "half of the spectrum". What is meant by this is that the transitions occur at photon energies such that the colors corresponding to the monochromatic beams of those energies are coplanar with the achromatic axis (one also says they are "complementary"). The full colors were discovered and their nature explained by Ostwald [Ost17, Ost36]. So how does all this help one to estimate the spectral albedo of any given color? Well, when the color happens to be an optimal color (or nearly so) one is done. The spectral albedo is uniquely determined, it has to be of the Schrödinger type. When the color lies in the interior of the color solid one can do the following:
— a *specific* solution can be constructed from a linear combination of an optimal color, black and white. All these colors have uniquely determined spectra, thus we indeed obtain a spectral albedo. The result will be piecewise constant with no more than two transitions between two levels and be physically realistic (albedo between zero and one);
— one obtains *all* solutions when one adds arbitrary "black material", taking care that the result stays in the physically realistic range. "Black materials" look black when irradiated by the source, thus the RGB values of a black material's spectral albedo times the source spectrum are $\{0,0,0\}$. Such blacks have physically unrealistic albedos (negative for some photon energies).

In the absence of prior knowledge it is prudent to select solutions that display minimal features. Since the particular solutions toggle between only two levels with at most two transitions, they are indeed desirable solutions. One may cut down the choice even further by selecting solutions that are binary mixtures of an optimal color with the even mixture of white and black (let's agree to call that "median gray"). The median gray is the symmetry center of the color solid, so we treat all colors similarly. (Notice that the median gray looks much lighter than the "medium gray" as used by photographers. Medium gray is 18% white, 82% black.) When one selects such a solution one stays as close as possible to the median gray (50% scattering regardless of incident photon energy) which may be considered the "default" assumption in the absence of any prior knowledge. (See figure 6.)

At this point the present method departs significantly from other, existing methods. One typically relies on PC's of albedos of all materials in the universe. I don't claim such knowledge. Moreover, the method is nonlinear and *the solutions*

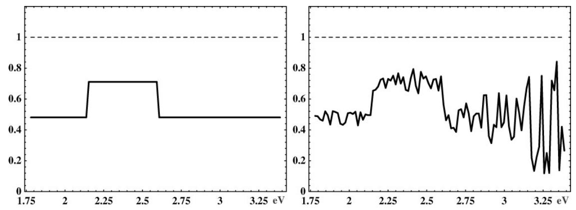

Figure 6: *Typical estimated albedo (left) and another solution obtained by adding a random black (right).*

are by no means linear combinations of any triple of fiducials. Therefore these albedos can actually represent *all object colors*, which is not the case for *any* PC based method (these cannot represent the optimal colors for instance). This is not often a problem in practice because very vivid colors are rare, but it is a conceptual flaw of linear methods in general.

5. Nature of the ambiguities

Although it is easy enough to construct quite reasonable particular solutions, this involves some rather specific assumptions:
— the intensity of the source is uniform on a logarithmic scale;
— the source spectrum can be reasonably well approximated with a second order curve (*e.g.*, in the photon energy, photon number flux density representation, it doesn't matter much though);
— spectral albedos are generally smooth and much like median gray.
Any of these assumptions might easily be replaced with more realistic ones. For instance, one might know beforehand to be confronted with beach scenes during daytime, one might know that the sources will be fluorescent tubes, one might know that the materials are all vegetables, and so forth.

In a great many cases one will lack such information though. Even worse, I have assumed knowledge of the spectral characteristics of the camera (the shape of the spectrum cone, nature of the projection). When one is handed *any* image one doesn't have such information, even when one uses the camera oneself one may lack spectral calibration. In such cases one has little option but to replace the spectrum cones with the first octant of RGB space (nonnegative coordinate values). Although only a coarse guess, this will usually work out well enough. The reason is that for color matching functions close to the CIE curves the volume of the color solid exhausts a large fraction of the circumscribed color crate of minimum volume (the nearest "RGB cube").

The remaining ambiguities are staggering. Suppose, one samples the spectrum at 100 photon energies. Then the kernel of the projection is a 97D space! When one considers the nature of this space it becomes clear that most of the "black spectra" are freaks though. They sport dozens of wriggles throughout the spectrum. However, such freaky spectra are by no means ruled out by the physics. Very complicated spectral albedos exist (*e.g.*, didymium glass powder [Kor69]), the same holds for source spectral densities (*e.g.*, the neon line spectrum), but they are rare under daily life conditions. General solutions are useless because *almost all* instances are freaks. Remarkably, one may safely take bets on many *specific* estimates. Bright yellow objects will absorb high energy photons, a grayish object in daylight will look grayish under a light bulb, and so forth. Such bets are quite safe for very generic reasons:
— *instrumental*: The spectral characteristics of the low, medium and high photon energy sensitive elements of the CCD chip are broad and smooth. High frequency wriggles average out;
— *ecological A*: Most source spectra and most spectral albedos are smooth and simple. Here "smooth" has to be taken *cum grano salis*, minor wriggles don't count (see below);
— *ecological B*: High frequency wriggles in the source spectra are unlikely to *match* those in the spectral albedos. One may safely assume the sources and materials to be uncorrelated. Since object colors involve averages over the products of source spectra and spectral albedos one expects the influence of high frequency wriggles (when present) to cancel out.

These reasons are sufficiently generic and certain in daily life circumstances that the estimation of material properties on the basis of their colors is not altogether impossible or even ludicrous as is sometimes suggested. The use of more informed priors is not likely to bring worthwhile gains except in very constrained circumstances.

6. Practical problems

I consider the general reasons discussed above compelling. Thus one may indeed construct useful estimates of material properties (the spectral albedo) from RGB data. I expect little gain from extensive ecological research to establish prior distributions for the various entities involved in the estimation as is commonly assumed to be necessary. One doesn't need *any* research to put trust in the fact that source spectra and material properties will on the average be uncorrelated or that the spectral characteristics of the apparatus introduce sufficient smoothing to discard most of the "black" that make up the remaining ambiguities. There are much more serious problems to address if one wants to make these methods work in practice.

Important problems have to do with the validity of the prior assumptions: Lambertian surfaces, single source, no interreflections, and so forth.

Most scenes are not dominated by a single source, even in cases where this might seem to be the case. Consider a portrait, taken on the beach under full sunlight. The face has an irradiated and a shadow side. The former is irradiated by the sun, the latter cannot be. It is irradiated by photons that have been scattered from other parts of the scene. It cannot be expected that this "ambient" irradiation has the spectrum of sunlight. Such simple considerations apply to almost any scene. This means that some prior segmentation into compartments of roughly uniform "radiant atmosphere" is necessary. Such segmentation is likely to depend at least partly on an interpretation of the scene.

Most surfaces are not Lambertian, some are somewhat close (at least for restricted ranges of directions of incidence and viewing), some completely different. One thing one has particularly to watch for is *bright specularities*. The estimation of the white point is essentially based upon outliers. This is *good* because a few outliers of different color suffice to determine the white point quite precisely, the actual density of pixels in RGB space being irrelevant. It is *bad* because even a single specularity will throw off the estimate dramatically. Again, this calls for a prior segmentation, best based on an interpretation of the scene. When a surface is not Lambertian, but has some weird BRDF [DGN*97], this needn't be a problem, for instance, when the direction of incidence and viewing are fixed one wouldn't even notice. This is roughly how the human observer handles such materials. They look one color seen one way, another color seen another way (think of Peacock feathers). One "stratifies" the material properties according to circumstances. Again, this asks for an interpretation of the scene.

In order to estimate material properties with some degree of success then, the present methods have to be embedded in a more extensive method that involves segmentation and interpretation of the scene. It is not likely that this can be avoided. Even human observers often change their judgments of material properties as they change their interpretation of the scene as a whole. This is especially evident in the interpretation of "difficult" realistic paintings.

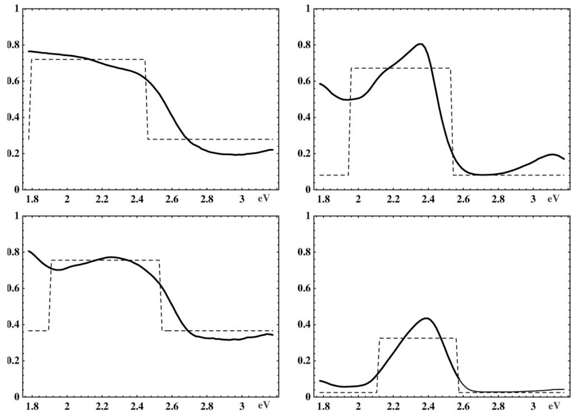

Figure 7: *Some examples of albedos, estimated from the color under daylight ilumination (D65), source spectrum assumed known. Samples are taken from the NCSU data base. Estimated albedo dashed, actual albedo drawn curve.*

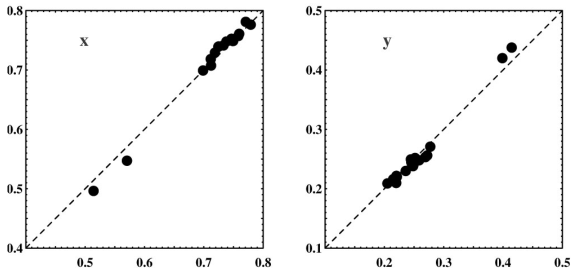

Figure 8: *The shift of the X (left) and Y (right) chromaticity coordinates (CIE XY–diagram) for a number of materials from the NCSU data base. On the horizontal axis the shift due to a change from average daylight (D65) to artifical illumination (light bulb,* i.e.*, black body at* 3000°K*). On the vertical axis the shift predicted from an estimate of the shift to artificial light obtained from the color under average daylight. The spectra of the sources assumed known.*

7. The "color constancy problem"

One way to put "the color constancy problem [MW86, [DZI93a,b, DZI94]" is: Given an image, rerender it in the "true colors". Here the "true colors" can be understood as the colors objects would have under standard conditions.

Clearly the estimation of spectral albedos yields one possible approach towards the color constancy problem. The spectral albedos having been estimated, one may proceed to render the scene under any assumed (in this case the standard) source.

In figures 7 and 8 I show an example. Random samples from the NCSU data base were irradiated with average daylight (CIE D65) and their colors computed. From these colors the spectral albedos of the materials were estimated (figure 7). The source spectrum was assumed known. Notice that the estimates are quite reasonable nonlinear approximations of the actual spectral albedos. The colors of the materials under another source (black body at 3000°K, roughly an ordinary light bulb) were found for the estimated and for the true albedos. In figure 8 I show scatter plots of the differences between the daylight colors and real or estimated artificial light colors in the X and Y directions in the CIE XY–diagram. Clearly the predictions are very good indeed. In this case we have practically perfect color constancy. The distances in the XY–diagram between the predicted and actual artificial light colors were at least twenty times less than the distances between the daylight and artificial light colors in all cases. This experiment illustrates both the results of albedo estimations and the efficiency of the present methods in classical color constancy.

Of course the estimation of source spectra and spectral albedos of materials opens possibilities for many applications such as automatic color correction for photographs taken under non–standard conditions.

Apart from such applications, the present method has the conceptual advantage that it allows the explicit construction and formal overview of all possible solutions. This will be useful in cases that deviate strongly from generic daily life conditions.

8. References

[Bou47] BOUMA, P. J.: *Physical Aspects of Colour.* N.V. Philips Gloeilampenfabrieken, 1947.

[DGN*97] DANA, K. J., GINNIKEN, B. VAN, NAYAR, S. K., KOENDERINK, J. J.: Reflectance and texture of real–world surfaces. *Proc. IEEEE Conf. on Comp.Vis. and Patt.Rec. (CVPR))* (1997).

[DZI93a] D'ZMURA, M., IVERSON, G.: Color constancy. I. Basic theory of two–stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A10* (1993), pp. 2148–2165.

[DZI93b] D'ZMURA, M., IVERSON, G.: Color constancy. II. Results for two-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A10* (1993), pp. 2166-2180.

[DZI94] D'ZMURA, M., IVERSON, G.: Color constancy. III. General linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A11* (1994), pp. 2389–2400.

[Goe82] GOETHE, J. W. VON: *Theory of Colours.* The M.I.T. Press, Transl. Ch. L. Eastlake, Introduction by D. B. Judd, 1982.

[Gra1853] GRASSMANN, H.: Zur Theorie der Farbenmischung. *Annalen der Physik 89* (1853), pp. 69–84.

[Hel1896] HELMHOLTZ, H. L. F VON: *Handbuch der physiologischen Optik.* (2nd Ed.), Voss, 1896.

[Kor69] KORTÜM, G.: *Reflectance Spectroscopy: Principles, Methods, Applications.* Springer, 1969.

[MW86] MALONEY, L. T., WANDELL, B. A.: Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America A3* (1986), pp. 29–33.

[Ost17] OSTWALD, W.: Das absolute System der Farben. *Zeitschrift für physikalische Chemie 92* (1917), pp. 222–226.

[Ost36] OSTWALD, W.: *Er und Ich.* Theodor Martins Textilverlag, 1936.

[Sch20] SCHRÖDINGER, E.: Theorie der Pigmente von grösster Leuchtkraft. *Annalen der Physik 62* (1920), pp. 603–622.

[NCSU] VRHEL, M.: The NCSU data base of 170 objects contributed by Mike Vrhel is to be found at fttp:///ftp.eos.ncsu.edu/pub/spectra/.