Application of Spectral Estimation Techniques to the Improvement of a 3D-color Digitizing Camera

Réjean Baribeau
National Research Council Canada
Ottawa, Canada

Abstract

This paper reports on the colorimetric improvement of a multispectral 3D digitizer through scanning at optimized wavelengths. These wavelengths were first established theoretically based on the criteria of minimal CIEDE2000 color difference over the set of reflectance curves from the full OSA-UCS catalog. A PCA-based and a spline-based spectral estimation method were considered, and sets of three, four and five optimal sampling wavelengths were derived for each method. This provided a basis for the selection of HeCd, ArKr, HeNe and DPSS commercial laser lines for which the colorimetric performance was predicted. This was then tested in the lab, where colour rendition charts were scanned with the camera at seven wavelengths, after which the charts were computer-rendered on a CRT display. Both the theoretical prediction and the experimental observation indicate that four well-chosen wavelengths are adequate for proper rendition of the charts.

Introduction

This study discusses the improvement of the colorimetric performance of a multispectral 3D laser scanner through proper selection of the scanning wavelengths. The scanner is designed to capture the shape and color of three-dimensional objects, such as artist paintings and archeological artifacts,[1] which can then be computer-rendered with high realism. The camera works by projecting a laser spot on the object, which is then imaged on a CCD looking at the scene from a slightly different direction from that of the incoming beam. Spatial coordinates of surface elements are obtained by triangulation. An auto-synchronized scanning mechanism allows the full coverage of the scene. In its commonly used configuration, superimposed red (633nm), green (532nm), and blue (442nm) laser beams are used for the projection, and the CCD image is split in its three components with a prism. The amplitudes of the three peaks are converted into reflectance values through a calculation that takes into account the particular geometry of illumination and detection at each surface element.[2]

The reflectance values obtained for the surface elements represent true intrinsic properties of the objects and are useful for monitoring the physical state of an artifact, for example before and after some restoration work. In the case of virtual reality applications, however, it is desirable to attribute some perceptual color values to the surface elements. The recommended practices[3] for measuring the color of reflecting objects require that the complete reflectance spectra be obtained across the visible range. This is then weighted by spectral power distributions of standard illuminants and by color matching functions of standard colorimetric observers to obtain tristimulus values that correlate with the visual perception of color. Until now, the laser scanner captured only three points on the reflectance curve and earlier studies[4] indicated that for typical reflecting surfaces, a CIELAB error of 3 to 5 units was to be expected. Even worse was that for some of the colors the error could be in excess of 10 units, enough for example to turn a red into an orange-red. Improvement of the color accuracy requires the incorporation of additional laser lines.

In a study published earlier,[4] we determined optimal sets of three, four, five and six sampling wavelengths that minimized the mean ΔE_{94} error over the 554 samples of the OSA-UCS catalog. The approach was to approximate the full spectra across the visible range using interpolation/extrapolation from the sparse data and then calculate the CIE colors for a standard illuminant. PCA-based and spline-base estimation methods were considered. ΔE_{94} color differences between true and estimated spectra were predicted. The optimal wavelengths were obtained through unconstrained optimization of the average ΔE_{94} error over the collection. This paper starts by revisiting this previous study in light of the newly proposed CIEDE2000 color difference formula[5] for which optimal sets of three, four and five sampling wavelengths are derived. Since these optimal wavelengths are not readily available from commercial lasers, we next investigate which combination of commercial laser lines to use. We start by predicting the average ΔE_{00} over the OSA-UCS collection for various combinations of wavelengths available from HeNe,ArK,HeCd and diode-pumped solid state lasers. We follow this theoretical analysis with its experimental counterpart for which we moved into the lab and coupled the 3D camera with seven different laser lines and digitized some color charts. Renderings of the charts on a CRT display, for various combinations of three, four and five wavelengths, are judged visually to provide a qualitative feeling of the benefit of using the extra wavelengths.

There is abundant literature on the subject of spectral estimation by imaging methods. Much of it relates to 2D imaging and is relevant for those 3D systems that rely on a 2D imaging subsystem for the capture of color. The NRC laser scanner is quite unique for its ability to capture both the 3D and the color images from a single optical system, requiring the use of laser lines. This offers unique advantages but poses unique optical engineering challenges as well.

Optimal Wavelengths

Optimal sampling wavelengths are determined assuming some spectral estimation method and some comprehensive set of color surfaces for which the spectral reflectance across the visible is already known. A color difference formula is also needed to determine the difference between the true and estimated color. The method for determining, say, three optimal sampling wavelengths (λ_1, λ_2, λ_3), starts with the following mean error computation function:

MeanError(λ_1, λ_2, λ_3)

1. for each color sample of the set:
    1.1. Read the reflectance R(λ_1), R(λ_2), R(λ_3) from its true spectral reflectance curve $R_{true}(\lambda)$.
    1.2. Estimate the spectral reflectance $R_{estimated}(\lambda)$ over the visible range.
    1.3. Calculate CIE L* a* b* values from the estimated spectrum as well as from the true spectrum.
    1.4. Apply a color difference formula to the true and estimated L*a*b* to obtain ΔE.
2. Return the average ΔE over all the color samples.

The set of wavelengths that minimize the MeanError() function is then sought with the Nelder-Mead simplex method.[6]

The above scheme works for a given set of color samples, a given spectral estimation method, and a given color difference formula. In Ref. [4], we determined optimal sets of three, four, five and six sampling wavelengths that minimized the mean ΔE_{94} error over the 554 samples of the OSA-UCS catalog, which is a comprehensive color atlas that covers fairly evenly a large fraction of the full color space. PCA-based and spline-base estimation methods were considered. Since that time, the CIE has proposed an improved color difference formula,[5] which incorporates parametric corrections for hue, chroma and lightness dependencies as well as for interaction between chroma and hue. The formula furthermore assumes the 10º observer while our early calculations with the CIE94 formula were for the 2º observer. These calculations were repeated with the new formula. Results appear in Table 1. The optimal wavelengths for the CIEDE2000/10º observer are found to stay within a few nanometers to what was found with the CIE94/2º observer. The change from 2º to 10º observers has more influence than the parametric corrections, as was verified by trying the 2º observer in the CIEDE2000 formula. We also confirm our previous findings that: 1- the average error is cut by about half going from three to four optimal wavelengths, and by another half going from four to five, and 2- that the PCA-based method performs slightly better than the spline-based method when the optimal wavelengths are used. Optimal wavelengths are listed in Table 1.

Table 1. Optimal wavelengths for the OAS-UCS set.

Method	[Optimal wavelengths], mean ΔE_{00}
PCA	[443.4 534.0 608.5], 1.7
spline	[456.3 533.4 607.3], 2.1
PCA	[446.4 512.4 562.5 612.7], 0.61
spline	[453.2 528.0 577.3 627.4], 1.1
PCA	[442.0 496.2 536.7 581.6 618.1], 0.34
spline	[443.9 511.8 557.6 597.8 633.8], 0.55

Perhaps one day the laser technology will allow the tuning of any of wavelengths in Table 1 while meeting the practical requirements for the 3D camera. Until then, one must resort to continuous-wave gas lasers and diode pumped solid state lasers and use Table 1 to help decide among what is available commercially. We test this idea in the next section.

Near-Optimal Commercial Laser Lines

The NRC 3D camera requires TEM00 continuous wave lasers with power outputs roughly 10 mW or higher. Since HeCd (442nm), DPSS green (532nm) and HeNe (633nm) were already available in the lab, we considered adding a ArKr tunable laser to the system, bringing a choice among 488nm, 514nm, 568nm and 647nm for additional wavelengths. Using Table 1 as a guide, we tested various combinations of wavelengths that fell close to the optimal ones by predicting the mean error for the OSA-UCS catalog. Table 2 presents these theoretical predictions.

As can be seen, the spline method and the PCA method become pretty much equivalent for the sets of non-optimal wavelengths considered. There appears to be a substantial cost for deviating from the optimal wavelengths. Most notable is the doubling of the predicted mean error when the non-optimal 633 red is used instead of the optimal 608 for scanning at three wavelengths. We believe the reason behind this is that many of the yellow-orange-red OAS-UCS samples owe their color due to a rather sharp transition between low spectral reflectance values and high spectral reflectance values occurring somewhere in the λ>550nm region. Hue is very sensitive to the exact location of that transition. Spectral estimation methods can only approximate this transition with some sort of ramp between sampled green and sampled red. If the sampled red is too far to the right of the λ axis, all the yellows, oranges and reds start looking orange-red since they are all approximated with the same ramp. The same idea applies to the blue-green shades but seems less sensitive.

When tied to using the HeNe 633 nm red, it appears very beneficial to incorporate a fourth sampling wavelength in the yellow region in order to catch the yellow-orange-red transition. The combination [442 532 568 633] is seen to perform by a factor of two better than [442 532 633] alone. It appears also that there is much less improvement to be gained from adding a fifth wavelength, either 488 or 514 to catch the blue-green transition. Finally, introduction of a sixth and a seventh wavelength does not bring significant improvement to the average color difference.

This exercise demonstrates the importance of being close to the optimal wavelengths, as the performance from scanning at all seven available wavelengths is still poorer than that predicted for four optimal wavelengths.

Table 2. Predicted performance for some combinations of commercial wavelengths.

Method	[wavelengths], mean ΔE_{00}
spline	[442 532 633], 3.0
PCA	[442 532 633], 3.1
spline	[442 532 568 633], 1.3
spline	[442 514 568 633], 1.4
PCA	[442 514 568 633], 1.5
PCA	[442 488 532 568 633], 0.88
spline	[442 488 532 568 633], 0.94
spline	[442 514 532 568 633], 1.1
spline	[442 488 514 532 568 633 647], 0.92
PCA	[442 488 514 532 568 633 647], 0.93

Experimentation

Our next move was to get a practical sense of these predictions in the lab. It is possible to interface with the camera as many lasers as one wishes and capture the 3D-spectral images in multiple scans. Each laser line is then coupled to the camera through a single-mode optical fiber and used to capture a 3D + intensity image of the object. Images captured with all the laser lines can then be merged into a single 3D multispectral image for which the first three channels contain the (x,y,z) spatial information, and the rest of the channels contain the spectral reflectance at the sensing wavelengths. Ideally however, one would like to use just three or four lasers and benefit from the camera's optics and peak detection electronics that allow the simultaneous detection of up to four peaks. 3D multispectral acquisition in one single scan then becomes possible.

Many sources of uncertainty other than spectral estimation errors affect the color measurement results. These can be of systematic nature, (e.g. calibration errors, dark signal offsets, etc.) or purely random. An example of the latter is speckle noise that affects both the amplitude and centroid location of the peaks detected on the CCD.[7] Since these errors may dominate the spectral estimation errors, it is worth experimenting before investing in extra lasers and modified procedures.

For testing the performance of the color rendition system, a Macbeth Color Checker™ (we used the small version of it) and a Macbeth Color Checker DC™ were scanned at all seven wavelengths available to us. A first scan was with the HeCd, DPSS green and HeNe lasers simultaneously coupled to the camera. This was followed by four additional scans with the extra wavelengths from the ArKr laser coupled one line at a time. CRT renderings of the charts were then produced for various combination of wavelengths, allowing visual judgment.

The following steps are involved in the experimental workflow:

1. An intensity calibration model is first established prior to scanning the objects. This is extracted from series of scans of a nearly-Lambertian white reference target placed at various depths inside the volume. The spatial dependence of the intensity response of the camera over its whole 3D volume of view is then determined. This is repeated for each scanning wavelength.
2. After a given object is scanned, the calibration model is used to recover reflectance values at each surface element (surfel). A compensation for shading effects, which makes use of the 3D information, is embedded in this step as well as in step 1.[2]
3. For each surface element, an associated full visible spectrum is estimated based on the reflectance values for the chosen set of sampling wavelengths. The spline interpolation method was favored for this test.
4. CIE XYZ tristimulus values for the D_{65} standard illuminant and the 2º standard observer are then derived from the estimated spectrum.
5. CIE XYZ tristimulus are next converted into sRGB values. (Note that the 2º observer is imposed by the sRGB standard while CIEDE2000 applies to the 10º observer; this is nevertheless consistent since there is no relationship between the standard used for encoding and the standard used for color difference evaluation.)
6. The sRGB part of the 3D color image is finally displayed on a CRT operated close to the sRGB standard conditions.

We note that our testing made use of flat 2D charts placed perpendicular to the camera's and viewer's lines of sight. The 3D information served only at the reflectance recovery stage. For more complicated 3D shapes, synthetic shadows would be applied on the virtual objects for realistic renderings. This shadowing must be performed in XYZ color space prior to conversion to sRGB.

For the comparison, the CRT was placed next to a Macbeth viewing booth simulating daylight and holding the real charts. The visual comparison confirms that the combination [442 532 633] fails to deliver the pure reds, which are rendered orange-red instead. Other “warm” colors suffer as well. This problem appears to be solved by the addition of the 568 nm wavelength. The large majority of the colors of the charts then come out fine, at least as far as can be judged with this kind of test. Although some visible improvement could be observed with the addition of a fifth wavelengths at 488 nm, this improvement was found to be very subtle and probably not worth the extra cost.

Conclusion

3D laser scanning at a number of discrete wavelength has the ability to measure surface reflectance values in perfect registration with the spatial coordinates of surface elements. When the reflectance data is used to calculate the perceptual color values of the surface elements, the choice of the laser lines used for scanning affects the achieved color accuracy. In virtual reality applications, there is no need to be more accurate than what the graphics engine can deliver and what the visual system can detect. In consequence one can limit the number of wavelengths to something practical.

We started this study by deriving optimal sets of three, four and five sampling wavelengths through simulations that made use of available color reflectance data sets and of the recently proposed CIEDE2000 color difference formula. The new sets of wavelengths were found to be very similar to those obtained with the CIE94 formula. The two formulae appear pretty much equivalent for this application, and the choice of the 10º observer over the 2º observer seems to have more influence than the parametric corrections.

The experimental counterpart of this study was then to optimize our camera though proper selection of commercially available laser lines. We started by predicting the mean ΔE_{00} errors over the OSA-UCS collection for various combinations of HeCd, ArKr, HeNe and DPSS laser lines. The best combinations of three, four and five wavelengths were found to be [442 532 633], [442 532 568 633] and [442 488 532 568 633]. Performance degrades substantially using these commercial lines instead of the unconstrained optimal ones. Sensing in the green-yellow to orange-red region of the spectrum was found critical for proper rendition of the pure reds. Our theoretical predictions were confirmed in the lab by coupling the camera to seven different lasers lines and capturing 3D-multispectral images of color charts.

For the laser lines we tested, a four-wavelength combination appears necessary in order to capture the pure reds. Incorporation of additional wavelengths is found to bring only marginal improvement. This holds true for the color charts that were used, which are believed to be representative of a large fraction of the color surfaces encountered in real life.

References

1. R. Baribeau, L. Cournoyer, G. Godin and M. Rioux, Colour Three-Dimensional Modeling of Museum Objects, Imaging the Past, British Museum Occasional Paper no. 114, pp. 199-209, Tony Higgins et al. editors.
2. R. Baribeau, M. Rioux et G. Godin, Color Reflectance Modelling Using a Polychromatic Laser Range Sensor, IEEE Trans. Pattern Anal. Machine Int., 14(2):263-269, 1992.
3. R. S. Berns, Billmeyer and Saltzman’s Principles of Color Technology, 3rd edition, John Wiley, New York, 2000.
4. R. Baribeau, Spectral Estimation from Laser Scanner Data for Accurate Color Rendering of Objects, SPIE 4421 (2001).
5. M.R. Luo, G. Cui and B. Rigg, The Development of the CIE 2000 Colour-Difference Formulae: CIEDE2000, Color Res. Appl., 26(5):340-350, 2001.
6. J.A. Nelder and R. Mead. Computer Journal, Vol. 7, pp. 308-313, 1965.
7. R. Baribeau and M. Rioux. Influence of speckle on laser range finders. App. Opt., 30(20):2873-2878, 1991.

Biography

Réjean Baribeau received his B.Sc. (physics) in 1977, his M.Sc. (physics) in 1979, and his Ph.D. (physics) in 1992, all from Université Laval. From 1988 to 1997 he worked with NRC’s Institute for Information Technology on the museum application of 3D laser scanning technology. He joined NRC’s Institute for National Measurement Standards in 1997 where he works in the field of color science applied to color imaging technology.