Answering Hunt's Web Shopping Challenge: Spectral Color Management for a *Virtual Swatch*

Mitchell R. Rosen, Lawrence A. Taplin, Francisco H. Imai, Roy S. Berns, and Noboru Ohta
Munsell Color Science Laboratory, Rochester Institute of Technology Rochester, New York

Abstract

In his keynote address to the Eighth Color Imaging Conference, Dr. Robert Hunt investigated the limitations of current display technology with an emphasis on how well it can be expected to represent product color on the World Wide Web. His conclusion was that for purchases where color was important, "the best safeguard is to ask for a swatch before placing an order." While his analysis was compelling, we considered the topic a suitable launching point for discussion of a potential use for emerging reproduction technologies. Designed to maintain spectral integrity, some of these new approaches are now worthy of public demonstration. For discussion, a scenario is described where image files contain spectral data and are available for download from on-line vending sites. On a user's computer the downloaded files may be spectrally color managed so that when printed on a local multi-ink printer, the result would be a *virtual swatch*. If successful, such a print would reflect a good approximation of the original sample spectra. Although some of the limitations described by Hunt would remain and some new ones would be introduced, if spectral color management were fast, cheap, widely available and accurate, there would be many situations in which Web consumers would be well prepared to make color-critical purchasing decisions without the use of factory-produced swatches. Contained in these proceedings are examples of spectral reproductions using a 6-ink printer juxtaposed with original fabric samples. The examples are crude, but point toward the potential of spectral reproduction. Discussion will follow examining why these renderings are both successful and unsuccessful.

Preamble

At last year's Color Imaging Conference, Dr. Robert Hunt presented a keynote address[1,2] entitled "How to Shop on the Web Without Seeing Red." Through a careful set of well documented arguments, Dr. Hunt concluded that when color was important, reliable decisions could not be made over the Web and needed to rely on off-line processes such as postal delivery of a swatch. During the presentation, one author of this paper turned to another and whispered "print them spectrally." The idea of the *virtual swatch* has grown into this paper.

In a recent editorial[3] in the Journal of the Society for Information Display, the editor described Dr. Hunt's findings as defining a "major challenge that should involve display hardware, display systems, and imaging software designers, technologists and developers." We agree. We hope to inspire the community to consider the advantages to Web commerce of a successful implementation of spectral color management. It is our belief that seeing red while shopping on the Web, as well as seeing green, yellow or chartreuse, with approximate spectral accuracy, is technically feasible and could be implemented within hardware not very different from that currently available and affordable.

Introduction

Virtual Swatch Examples

Each of Figures 1 and 2 contain three parts. In the center are actual fabric swatches. On the left is a standard CMYK ICC color managed reproduction of a sample taken from the same bolt of fabric. To the right of the swatches are spectrally color managed 6-ink CMYKOG reproductions. Details of the workflows are reported later in this paper.

The reproductions contain many flaws exemplifying some inherent problems for both reproduction techniques and point toward areas where the spectral approach could be improved. As is particularly well demonstrated in Figure 1, out-of-gamut colors will continue to be troublesome for spectral reproduction as they have been in the colorimetric domain.

Each page was individually printed on a 6-ink CMYKOG ink-jet printer so the examples were at the mercy of slight hue shifts when ink was changed and to jet-clogs. The ICC reproductions only used 4 inks, CMYK, so they were less vulnerable to jet-out than the 6 ink, CMYKOG, spectral reproductions. Horizontal banding can thus be found one-third more frequently in the spectral examples. On the other hand, vertical banding in the spectral examples is not a printer artifact but is a consequence of the image processing. An explanation for these latter effects is described below.

Examine the figures under changing light conditions. See if the color appearance of the spectral examples follow more closely the color changes in the fabric swatches.

Hunt: Monitor Fails Color-Critical Decisions

In 1970, Hunt enumerated the six objectives for color reproduction.[4] Recalling these objectives in his Web shopping talk,[1] Hunt declared that *equivalent* and *corresponding* were the best fits to the needs of on-line color shoppers. *Equivalent color reproduction* attempts to match both the appearance of a product and the luminance level under which it would be typically viewed. For situations where luminance matching was unrealizable, *corresponding color reproduction* that reproduces the appearance of the original as viewed under reproduction viewing conditions, was described as being most appropriate for the Web.

Hunt made it clear that "there are many goods and services offered on the web where color is sufficiently unimportant for web marketing to be very successful."[5] He also explained that a monitor calibration approach can be very helpful. "There is no doubt that standardising the display set-up can remove an important source of variability in web marketing activities."[1] Unfortunately for the Web merchant, it was shown that when color is important, the shortcomings that remain are unaddressed by such tools or any other universal remedy. Persistent problems include: surface colors with appearance that change as illumination changes, camera instrumental metamerism, observer metamerism, surround effects, limited color gamut and limited resolution.

Hunt stated that although *spectral color reproduction* would be a very good objective for the paper catalog industry, spectral matching was not possible for Web shopping "because the spectral power distribution of displays from the web are usually different from those of most goods." This paper asks whether we could not outfit the end-user with a system that could strive to produce quality on par with a high-end paper catalog.

The Swatch, the Web and the Virtual Swatch

The swatch, of course, is generally free from the criticisms that disqualify the display for color-critical commerce. If the original product changes appearance under different illumination, so does the swatch. With a swatch, concerns about camera metamerism, observer metamerism, display gamut and display resolution are no longer relevant since the swatch has all the same spectral and spatial properties of the original – it is, in fact, a piece of the original.

While the swatch has these desirable aspects, it also has a number of drawbacks relative to the Web. *Instantaneous gratification* is the battle-cry of the Web shopper. No-cost retail sales is the dream of the Web merchant. Ordering and delivering a swatch adds time, uncertainty and cost to the deal. The turn-around time for such a sale would be measured in days or weeks with customer attrition rate, no doubt, quickly asymptoting toward 100%.

What if the Web consumer could be provided with a machine which could, in minutes, produce any number of swatches of any desired size, up to a reasonable, say, 8-and-half by 11 inches? The major impediment to color shopping on the Web would have been removed. Users would be able to experiment with ideas, change their mind, try again, and hop from site to site in search of just the right colors. Transactions could take place at the speed of Web commerce. Extra cost to the Web user would be associated with consumables that feed the swatch-making machine, but lower merchandise costs and more reliable purchases would be the pay-off. Web masters would have to stock their sites with recipes for the swatch machines and hope that compatibility did not become a problem. The advantage to the Web sites would be faster sales, more satisfied customers, fewer returns and no need to produce or pay postage on samples.

A swatch-making machine is, of course, a fantasy, but the next best thing is already or will soon be in many consumers' hands: the high resolution multi-ink printer. Today, 6-color, 720 DPI printers cost less than $300 with prices, quality and speed-per-page being adjusted constantly. For sufficient spectral dexterity, 6-inks may prove to be insufficient with perhaps 8 or 10 inks being the general optimum. Regardless, it is clear that with today's hardware, or some small variation from it, combined with the right inks, proper characterization and appropriate software, the *virtual swatch* can be supported.

Spectral Printing

Spectral color reproduction is a worthy aim for catalog merchants. This is because color decisions improve when the original and reproduction are spectrally identical. The color match between original and a spectral copy transcend changes in illumination, individual observer variations and even instrumental metamerism. Admittedly, there are limitations to the value of a spectral reproduction since it does not promise to maintain other physical aspects of the original such as surface structure and gloss. Also, due to the limited number of inks used in any realistic printer, there is going to be some level of spectral error between most originals and their reproductions. In spite of those shortcomings, high resolution, approximately accurate spectral copies can be very useful in making color-critical decisions.

Five- and six-ink printers have been shown to be capable of making approximate spectral reproductions.[6-9] Printer characterization approaches described previously in the literature depend upon physics-based models and their variations. The characterization method discussed in this paper will be purely empirical based on a lookup table built from sampling regularly spaced combinations of predetermined ink levels. Linear interpolation is used to estimate spectra. Regardless of how a mapping is determined, as long as it can predict accurately the spectra resulting from printing combinations of printer digits, the inversion of the mapping can be used to render spectral

Fig 1. Brite Lites: comparison of reproduction techniques. ICC approach on left- scanner to CMYK using standard profiling tools. Original fabric in center. End-to-end spectral approach on right. Blue background is out-of-gamut spectrally and colorimetrically.

Fig 2. Assorted Bulbs: comparison of reproduction techniques. ICC approach on left. Original fabric in center. End-to-end spectral approach on right. Vertical features in spectral image are artifacts of raster processing.

Note for PDF version of this paper: Figures 1 and 2 were specially prepared for hardcopy book publication. Original fabric samples were glued to the page in the center column. A six color printer was used to produce the reproductions found in the left and right columns as described in the above text. A limited number of inserts are available while supply lasts. Please send request and a self-addressed, stamped envelope large enough to hold an 8.5 by 11 inch sheet of paper to:

Mitchell Rosen
Munsell Color Science Laboratory
RIT
54 Lomb Memorial Drive
Rochester, NY 14623

reproductions. Inversion techniques vary according to the nature of inherent or forward mappings. For characterizations based on directly invertible functions, inverse relationships can be solved mathematically. For those based on complex models not easily inverted or lookup tables, optimization techniques should be employed to determine the mapping from spectra to printer digits.

Approach

Multi-channel Capture of the Fabric

A pair of fabric bolts were acquired. A pattern called *Brite Lites* is illustrated in Figure 1 and *Assorted Bulbs* in Figure 2. Brite Lites has large areas that are out-of-gamut for the output printer, particularly the blue background. Assorted Bulbs is substantially in-gamut.

Images of the fabric samples were captured on a Photometrics Quantix camera with a Kodak KAF6303E monochrome sensor, 3072 by 2048 at a 9µm pitch. See Figure 3. The camera can deliver 12-bit images at 5MHz or 1MHz. When using the 5MHz speed, a ghosting of the image, possibly associated with weak electron surface traps in the silicon substrate, was observed for very short exposure times, so 1MHz was used for this study. A CRI tunable filter[10] was mounted between the lens and the camera body. This solid state filter, based on Lyot-Ohman principles,[11] selects spectral bandpasses under computer control. Our filter can be physically configured into either of two modes, a nominal 40nm bandpass with medium contrast or a nominal 10nm bandpass with high contrast. For these experiments, the wider 40nm was chosen. A Rodenstock Rodagon 105mm lens at f5.6 completed the optical setup of the camera.

Figure 3. Photometrics Quantix camera (foreground) imaging Brite Lites fabric (background). Lens and CRI tunable filter (not visible) mounted in front of camera. Also out of view for this illustration are Bogen quartz lamps to the left and right.

Exposure times for each filter setting were determined using a target with large areas of highly reflective white. Exposure times were designed so that several hundred 12-bit digits of "headroom" were left for the white areas.

Images of the fabric samples and for a gray card were captured with 3 repeats at each wavelength-setting using the predetermined exposure times. The gray card was repositioned three times and a full series of 33 bandpass exposures were made for each position. Approximately 240 pixels were captured per inch of fabric.

In IDL, the captured repeats were averaged to reduce random noise. Over the 3 gray card positional images, a median filter was applied to each x,y coordinate. This reduced the affect of any random scratches or defects on the gray card. This median gray collage was used to perform flatfielding for the averaged fabric images. Each averaged fabric wavelength-setting image was divided, pixel-by-pixel by the value in the median gray collage at that same pixel and then multiplied by 100, thus correcting for non-uniform illumination and varying pixel sensitivities. Floating point tiff files of these flatfielded and normalized images were saved.

Spectral Estimation of the Fabric Images

The reflectances for a variety of imaged colors were spectrally measured on a GretagMacbeth Spectrolino reporting reflectance factor from 380nm to 730nm in 10nm increments. The colors included 6 from each fabric. In Figure 4, the measurements from the Brite Lites Fabric are illustrated. Figure 5 has the measurements from the Assorted Bulbs.

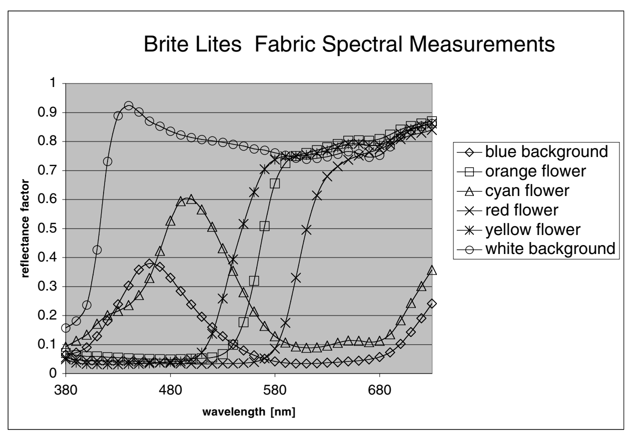

Figure 4. Spectral measurments for Brite Lites fabric.

For deriving spectral estimates from the flatfielded and normalized digits a number of techniques were attempted. The first approach was very direct and was based upon measurements of the various system components and knowledge of the exposure time for each wavelength-setting. A matrix, *C*, was derived from these first principles relating how the system would produce digits given input spectra:

$$D_{nxm} = C_{pxm} \cdot S_{nxp} \quad (1)$$

where D is an nxm matrix of m flatfielded and normalized digital counts for n samples, S is an nxp matrix of p spectral values for n samples and C is a pxm correction matrix which relates spectra to digits. While the correction matrix, C, worked reasonably well relating spectra to digits, its pseudo-inverse, $[C^T C]^{-1} C^T$, produced wholly unacceptable results when relating digits back to spectra. This may be due to the extreme overlap of spectral transmittances between filter-settings leaving the inverted relationship with extreme sensitivity to slight amounts of noise and imprecision.

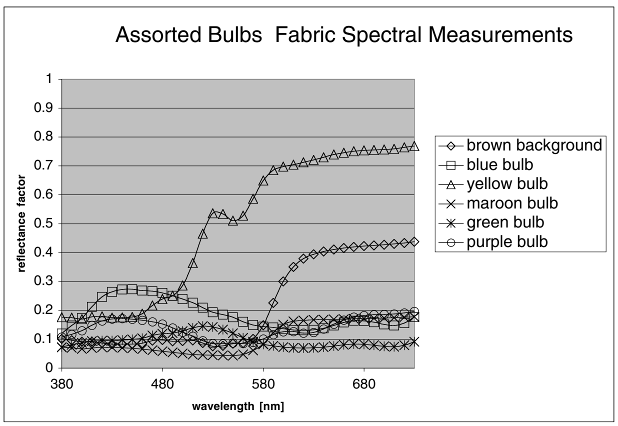

Figure 5. Spectral measurements for Assorted Bulbs fabric.

A second approach was attempted where B, the matrix which directly relates digits to spectra was derived through optimization methods. The Powell routine12 from the Numerical Recipes library13 was used for this purpose. B is defined as giving the lowest error in satisfying the following equation where i spans all tested spectra-digit combinations:

$$S_i = B \cdot D_i \quad (2)$$

S_i is the spectrum of sample i, and D_i are the flatfielded and normalized digital counts for sample i. Although the results here were far more usable, they predicted spectra that were very jagged. The matrix tended to be unstable for untested digits.

Methods such as principal component analysis have been shown to work well for spectral estimation. For this study, the six dominant colors from each fabric, as illustrated in Figures 4 and 5, were chosen as spectral bases. The Powell routine was used to process the entire image. An optimization was applied, finding the weighting of digits associated with each basis spectra that would best represent that pixel's flatfielded and normalized digits. The derived coefficients were then applied to the basis spectra to build the spectral estimate.

Spectralizer[14] a spectral image analysis tool, introduced at last year's Color Imaging Conference, allowed easy visual assessment of the quality of the spectral reconstruction. [Spectralizer requires an IDL environment. It may be downloaded free-of-charge from www.cis.rit.edu/mcsl/online/Spectral/Programs.] Figures 6 and 7 are Spectralizer screen shots showing the estimated spectra at two different points in the Brite Lites spectral image.

An artifact associated with the optimization approach used here is that the derived coefficients are strongly dependent on the seeds given to the processing engine. Because the image was processed in a raster style from left to right, top to bottom, local minima were often propagated in this same pattern. See Figure 8.

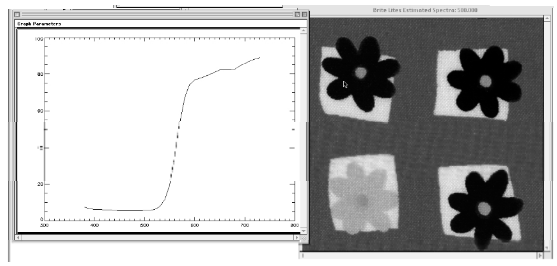

Figure 6. Spectralizer probe of 36-band spectral estimate of Brite Lites fabric. Image on right represents the 500nm band of the estimate. Cursor can be seen over upper left flower (orange). Graph shows spectral estimate for selected pixel.

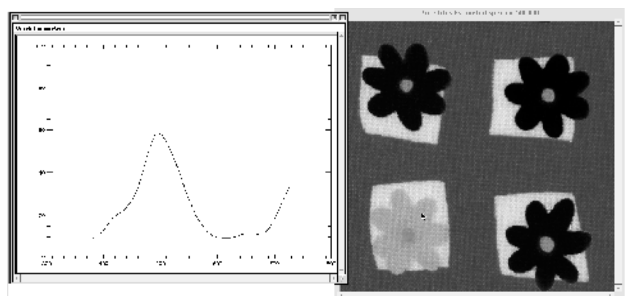

Figure 7. Cursor is now over lower left flower (cyan).

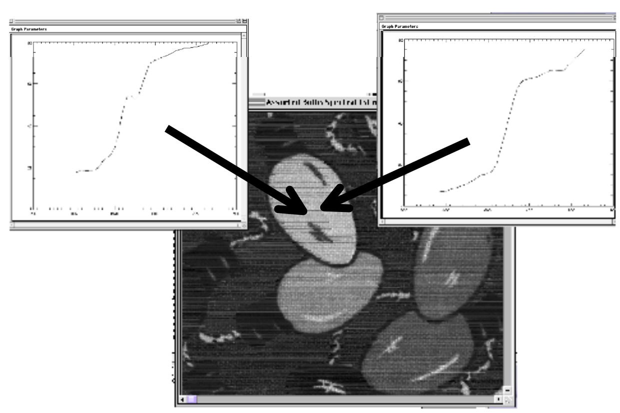

Figure 8. Image represents the 500nm band of the spectral estimate for the Assorted Bulbs reproduction. The two graphs are spectral estimates derived for pixels one line apart in the yellow bulb. Differences are propagated in linear fashion.

Spectral Characterization of the Printer

A six-ink 720 DPI Epson 1200 ink-jet printer was specially retrofitted to have continuous ink feeding of four standard process inks plus an orange and a green ink. A printer driver was obtained which allows direct specification of the individual CMYKOG planes. The printer is shown in Figure 9. 8-and-a-half by 11 inch Hammermill Jet Print Ultra Matte Radiant white paper with 94 Brightness was used as the printing medium.

Figure 9. Epson 1200 retrofitted for continuous 6-ink feeding

Ramps of each color were printed to determine the relationship between individual ink digits and percent dot coverage. Using this derived relationship, a full factorial sampling of all ink combinations was designed. Interspersed on the characterization targets were the midpoints of the hypercubes described by the factorial design.

For quality control on each page, five copies of 16 different patches were placed, one set in each corner and one set in the page center. Also, individual color ramps for the six inks were printed on six of the pages. Each page held 768 1/3 inch square patches in 24 columns of 32 rows each. In the final design, 32 pages were printed, for a total of 24,576 patches. Test prints were made before and after each page to ensure that no jets were clogged.

A GretagMacbeth Spectrolino/SpectroScan was used to measure the reflectance spectra of the samples. Reflectance was sampled at 36 wavelengths spanning 380nm to 730nm in 10nm increments. A lookup table indexed by percent dot coverage was built based on these measurements. It used linear interpolation to predict the spectra printed by a set of requested percent dot coverages.

Accuracy was determined by comparing measured hypercube midpoint spectra to LUT-interpolated spectra. See Figure 10.

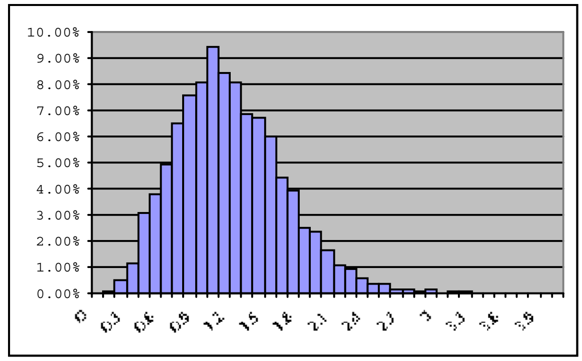

Figure 10. ΔE^*_{ab} histogram of difference between interpolated and measured spectra for the 4096 hypercube midpoints.

Processing the Spectral Images

Standard optimization techniques proved useful for inverting the relationship between percent dot coverage and printed spectra. The data at each pixel of the spectral estimate image was processed through the printer characterization table, determining the best combination of percent dots from the six inks that would produce the closest spectral match. The percent dot images were then passed through 1 dimensional lookup tables that converted the values to byte digits for printing. As described in the previous discussion on the making of the spectral estimates and illustrated in Figure 8, again, the seeding of the process had a significant impact on the results. For example, Figure 11 shows the 6 bands used to print the spectral image found in Figure 1. Although the banding looks severe in Figure 11, the spectral copy in Figure 1 does not show off the severity because of the tremendous spectral and colorimetric redundancy in the printer 6-ink space. Figure 2 shows that reliance on such redundancy can fail, as the banding is visible there.

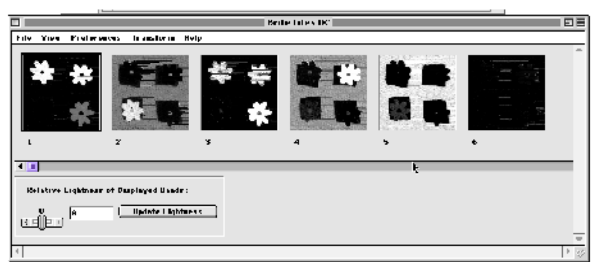

Figure 11. Digit bands in OGYMCK order for Brite Lites Virtual Swatch as displayed in Spectralizer. Banding appears severe here, but is not visible in spectral reproduction of Figure 1. (Note Figures 1 and 2 rotated and reflected relative to this illustration, so printed artifacts would run vertically.)

The ICC Representation

GretagMacbeth Eye-one Match was used to produce ICC profiles for an Epson 640u scanner and the same Epson 1200 printer described above but in this case using only the CMYK inks. In PhotoShop, the Image Mode "Profile to Profile" interface was used to convert an RGB scan of the fabric to the CMYK of the printer. The images were post-processed to reduce sharpness so that they would match the approximately 240 DPI of the spectral image. Reproductions are found on the left-hand side of Figures 1 and 2.

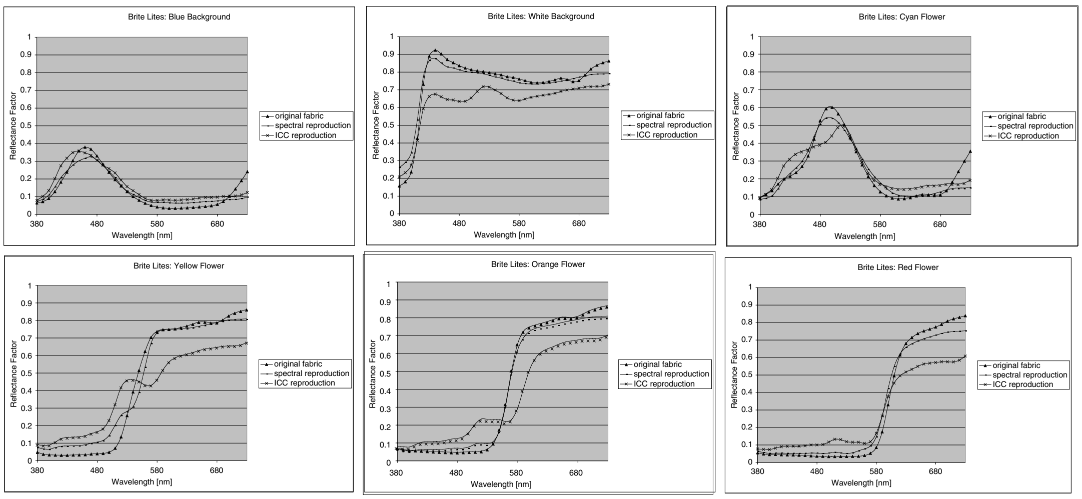

Figure 12. Spectral measurements of spectral and ICC reproductions compared with original fabric measurements.

Results

Success of a virtual swatch demonstration will be measured in how well the spectral reproduction matches the original under all viewing environments. While spectral reproductions in Figures 1 and 2 are not perfect, they do represent a good step in the right direction. Figure 12 shows a comparison between original and reproduced spectra for the Brite Lites fabric. In all cases, as expected, the spectral reproductions are far closer to the original spectra than the ICC reproductions.

Conclusions

The concept of the Virtual Swatch has been explained. Its potential usefulness to the Web consumer described. An early attempt at spectral reproduction was undertaken as the heart of this paper and actual samples included in this publication. Many areas for improvement have been identified.

Acknowledgements

We were very privileged this summer to have a pair of high school interns in the laboratory, Yoshi Fujita and David Swan. They were great sports and fun to have around. In between experiments and instrument calibration work, they found the time to cut over 1700 fabric samples for this publication. The gluing crew for Figures 1 and 2 at various times included Yoshi Fujita, Yongda Chen, Pano Spiliotis, Lawrence Taplin, Mitchell Rosen and Alma, Peter and Marissa Balonon-Rosen (wife, son and daughter of the primary author). Fuji Xerox partially supported this work. Special thanks to RWG Hunt for his initial inspiration, for his insightful comments and criticisms of an early draft and for his kind encouragement.

References

1. R. Hunt, How to Shop on the Web Without Seeing Red, *Proc. 8th IS&T/SID Color Imaging Conf.*, 2-7 (2000).
2. R. Hunt, How to Shop on the Web Without Seeing Red, *Journal of the SID*, **9**, 79-86 (2001).
3. A. Lakatos, Forward, *Journal of the SID*, **9**, 77 (2001).
4. R. Hunt, Objectives in colour reproduction, *Journal of Photographic Science*, **18**, 205-215 (1970).
5. R. Hunt, *personal correspondence*, July, 4 (2001).
6. D. Tzeng, R. Berns, Spectral-Based Ink Selection for Multiple-Ink. Printing I. Colorant Estimation of Original Object, *Proc. 6th IS&T/SID Color Imaging Conf.*, 106, (1998).
7. D. Tzeng, Spectral-Based Color Separation Algorithm Development for Multiple-Ink Color Reproduction, Ph.D. dissertation, R.I.T., Rochester, N.Y. (1999).
8. M. Rosen, X. Jiang, Lippmann2000: A spectral image database under construction, *Proc. Intl. Symp. on Multispectral Imaging and Color Repro. for Digital Archives*, 117-122 (1999).
9. F. Imai, M. Rosen, R. Berns, D. Tzeng, Spectral Reproduction from scene to Hardcopy I: Input and Output, *Proceedings of SPIE, Sensors and Cameras Systems for Scientific, Industrial, and Digital Photography Applications II*, **4306**, 346-357 (2001).

10. R. Slawson, Z. Ninkov, E. Horch, Hyperspectral Imaging: Wide-area spectrophotometry using a liquid-crystal tunable filter, *Publications of the Astronomical Society of the Pacific*, **111**, 621-626 (1999).
11. B. Lyot, US Patent 2,718,170 (1955).
12. W. Press, S. Teukolsky, W. Vetterling, B. Flannery, Numerical Recipes in C, The Art of Scientific Computing, Second Edition, Cambridge Univ. Press, NY (1992).
13. Numerical Recipes, The Art of Scientific Computing, Code CDROM, v 2.08 (1997).
14. M. Rosen, M. Fairchild, G. Johnson, D. Wyble, Color Management within a Spectral Image Visualization Tool, *Proc. 8th IS&T/SID Color Imaging Conf.*, 75-80 (2000).

Biography

Mitchell R. Rosen is a Senior Color Scientist at the Munsell Color Science Laboratory at RIT. He is also an Imaging Science Ph.D. candidate in the Center for Imaging Science. He joined the staff of the Munsell Laboratory in 1998 after 10 years with Polaroid as a member of its Image Science Laboratory. He received his B. S. degree in computer science from Tufts University and his M. S. in Imaging Science from RIT. Prior to his master's work, he spent several years as staff programmer in the Visible Language Workshop, Media Laboratory, Massachusetts Institute of Technology. During his years at Polaroid he worked in image quality evaluation, imaging system specification, algorithm development and the evolution of industry-wide efforts to standardize desktop color management. His work in the Munsell Laboratory is focused upon development of spectral imaging infrastructure. He is a member of IS&T.