

Spectral Reproduction from Scene to Hardcopy II: Image Processing

Mitchell Rosen, Francisco Imai, Xiao-Yun (Willie) Jiang, Noboru Ohta

Munsell Color Science Laboratory, RIT

ABSTRACT

Traditional image processing techniques used for 3- and 4- band images are not suited to the many-band character of spectral images. A sparse multi-dimensional lookup table with inter-node interpolation is a typical image processing technique used for applying either a known model or an empirically derived mapping to an image. Such an approach for spectral images becomes problematic because input dimensionality of lookup tables is proportional to the number of source image bands and the size of lookup tables is exponentially related to the number of input dimensions. While an RGB or CMY source image would require a 3-dimensional lookup table, a 31-band spectral image would need a 31-dimensional lookup table. A 31-dimensional lookup table would be absurdly large. A novel approach to spectral image processing is explored. This approach combines a low-cost spectral analysis followed by application of one from a set of low-dimensional lookup tables. The method is computationally feasible and does not make excessive demands on disk space or run-time memory.

1. BACKGROUND

Although many research centers around the world are looking at various aspects of the spectral imaging workflow,¹⁻⁵ there are very few publications concerned with spectral hardcopy output. There are notable exceptions to this observation including several papers from the Munsell Color Science Laboratory.⁶⁻⁸ Spectral hardcopy remains an elusive holdout for several reasons: 1) it is not necessarily obvious how one would utilize typical spectrally broad printing inks for spectral output and, even if one could use these inks, natural intuition leads to the assumption that a plethora of inks would be needed to make a reasonable spectral match to an arbitrary input spectrum; 2) very few spectral models exist for hardcopy devices describing the relationship between requested output digit and the delivered spectrum; and, 3) even in the case where a model is given, if it is assumed that the model is computationally complex, then the problem remains as to how to apply the inverse of the model to a large spectral image within a reasonable time. These first two concerns have already been addressed in previous publications. It has been shown that for certain images, four process inks plus two or three readily available inks can be used to produce highly spectrally accurate output⁶. Even for cases where best obtainable spectral accuracy is still low, reduction of metamerism is a worthy goal for many applications. Spectral models do exist for certain printing technologies⁹ and have been utilized in the above referenced studies^{6,7}. Application of the inverse of a model to a complex and large spectral image is where the effort for this study was applied.

In any color reproduction system, there are generally three major components. There is the image capture or synthesis stage, then the image processing stage and finally an image output stage, sometimes referred to as image display or image rendering. Traditionally, color capture devices have been engineered to take into account the spectral integration which is an important aspect of human color perception. Thus, typical color capture devices have three channels each having spectrally wide sensitivity, much like the human. For image output, a computer monitor is similar to the typical image capture device in that it relies on three spectrally wide channels. Typical printers have four channels where a black is added to a set of three spectrally wide subtractive primaries. The black is primarily used to increase the color gamut for dark colors. A standard image processing regime is designed to manipulate the three incoming channels in anticipation of how the output stage renders colors. For accurate color reproduction, the image processing stage must be aware of the characteristics of both the image capture device and the image rendering device. Figure 1a describes a general color reproduction system. The International Color Consortium (ICC) defines the industry standard method for informing color processing about the stimulus and response character of the two ends of the color reproduction chain. The input device ICC profile describes the relationship between input digit and colorimetry under the D50 illuminant. Colorimetry, based on the human color matching functions, is a three-dimensional space. For most color capture devices the mapping between its three channels and the three dimensions of colorimetry is overall one-to-one. Where the actual mapping is many-to-one or one-to-many, this is due to differences between the instrumental metamerism of the input device and the metameric characteristics of the "standard observer." The ICC profile for the output device describes the relationship between D50 colorimetry and output digit. For most printing devices, there is much redundancy between colorimetry and output digits. This is due to the presence of the fourth colorant. Makers of ICC profiles must deal with these mapping issues and provide color processing with single

mappings in each instance. In a typical ICC compliant image processing stage of a color reproduction system, a pair of profiles will be used, one referring to the image capture device and the other referring to the image rendering device, as shown in Figure 1b. Optimizations will take place to ensure that the image processing proceeds in as efficient a manner as possible since it is typical that each pixel in a captured image would undergo some level of manipulation. The fact that each ICC profile is tied to the common space of colorimetry makes the concatenation of a series of transforms described in the profiles a fairly straight forward and computationally efficient matter. Figure 1c shows a typical ICC image processing block diagram which includes concatenation. After applying these efficiencies to the transformation parameters, it is likely that at its most complex, the image processing stage would consist of applying in parallel to each input channel a one-dimensional lookup followed by a three-dimensional lookup and possibly a final one-dimensional lookup applied to each output channel. Computationally and with respect to memory requirements the image processing chain described here is very feasible and highly efficient.




Figure 1a: General color reproduction block diagram.




Figure 1b: Logical ICC image processing block diagram.




Figure 1c: Typical ICC image processing with efficiencies. Transform box may include 1-D and 3-D lookups and matrices.

When the input and output devices in a color reproduction system possess a sufficient number of channels, spectral characterizations rather than colorimetric ones are possible. The image processing stage for such a system would necessarily be more complex than the ICC approach described above. In this scenario, a request could be to match input and output reflectance (or transmittance) or to match input and output radiance under user selected illuminants. These spectral matching capabilities would require far more computational power than that needed for the ICC colorimetric approach. In order to discuss efficient approaches to this problem, it is first necessary to outline the type of demands to which such a system might be subject.

2. SPECTRAL COLOR MANAGEMENT

The discussion of how spectral color management strategies might proceed has only recently begun within the community.¹⁰⁻¹² Once the spectra of an incoming image is known, either in terms of reflectance or radiance, those data could be harnessed to generate a number of useful outcomes. A reproduction which matches the reflectance spectra of an original object would be highly desirable. Theoretically, such a reproduction would preserve color matches over the range of all illuminants and for all observers. Reflectance-based rendering would be less sensitive to small characterization and calibration errors and to rendering noise. Radiance matching could take current viewing conditions into consideration when reconstructing an object's appearance as it would be under other viewing conditions. Again, sensitivity to noise and error would be reduced. It should be noted, however, that certain appearance attributes such as surface characteristics, surround conditions, self-radiating versus reflective copy, viewing distance, etc. are no more accounted for within a simple spectral color management system than they would be within a traditional ICC system, but such concerns are outside the purview of this paper. The reflectance and radiance matching tasks referred to above are only the most obvious spectral-based extensions of a color

reproduction infrastructure. New capabilities not previously appreciated will also present themselves. Some examples include: an increased ability to analyze original scenes improving the recognition and reproduction of certain memory colors such as skin tones, grass, sky, etc.; higher levels of within-gamut reproduction quality through the use of multi-channel input devices for which instrumental metamerism, previously unavoidable and the scourge of today's color reproduction world, would be rare; the introduction of fluorescence as a positive force in color reproduction, unmasked through the acquisition of multiple spectral images taken under different light sources; and, improvement in making spot-color and specialty ink choices.

The spectral color reproduction system will have the same basic stages as discussed before and shown in Figure 1a: input, image processing and output. Characterization of input and output devices will need to be carried out, but this time with respect to spectra. It would be very convenient if Figure 1b which describes the ICC logical image processing work flow could be replaced by a similar diagram which simply replaces the word 'colorimetry' with 'spectra' and the word 'colorimetric' with 'spectral.' Unfortunately, it does not turn out to be so simple. By increasing the generality of the system, it now needs to obtain added functionality. Within a spectral color management system the following requirements would be necessary:

Input profiles could describe any of the following transforms:

- a) Input digits to reflectance
- b) Input digits to radiance
- c) Input digits to colorimetry (not spectral, but good for backwards compatibility)

Color processing would need to be able to make the following conversions:

- d) reflectance to radiance by multiplying by illuminant
- e) radiance to reflectance by dividing by illuminant
- f) reflectance to colorimetry by multiplying by illuminant, multiplying by color matching functions and then integrating
- g) radiance to colorimetry by multiplying by color matching functions and then integrating

Output profile could describe any of the following transforms:

- h) reflectance to output digits
- i) radiance to output digits
- j) colorimetry to output digits

These basic operations can be applied in series to carry out complex tasks. For example, consider the problem of choosing the right wallpaper for one's tungsten lit living room. It is well known that the cool white fluorescents in the store can be misleading. Any thinking customer could take out a multispectral camera, capture a picture of the wallpaper in the store and take a second shot of the store's lights. At home the two images would be downloaded to the computer. Using the camera's spectral profile, each image would be converted to radiance (functionality 'b', above). Using the radiance image of the store light source, the wallpaper radiance image could then be transformed to reflectance (functionality 'e', above – this is not *true* reflectance. This pseudo-reflectance will be discussed below). A subsequent picture of the living room light allows for calculation of what the wallpaper would have looked like at home (functionality 'd', above). Finally, to view the color image on the home monitor, the radiance image needs to be converted to XYZ (functionality 'g', above) and then the ICC monitor profile would be used to transform to monitor RGB (functionality 'j', above). Figure 2 illustrates a cartoon of the user action in this example. Figure 3 shows the series of actions taken by the spectral color management system.




Figure 2: Spectral color management would provide a way to impose a new light source on the image of a captured object.




Figure 3: Block diagram illustrating the logical color management system actions which would accompany the wallpaper example described above. The wallpaper radiance image has the store light source divided away and then is converted to a new radiance image with the living room light source multiplied into it. The new image is displayed colorimetricly on a monitor.

Dividing the radiance image by the store light source would not return, strictly, a reflectance image. The resultant image would still contain the effect of uneven illumination and gloss. Also, it would be impossible without referring to a known reflector in the scene to assign an absolute reflectance level to anywhere in the image. All values would be relative. After the division, a first pixel with twice the value of a second pixel in the same band would certainly have reflected twice the photons toward the camera as the second pixel did. It would be impossible, though, to determine if this was because the surface imaged by the first pixel had twice the reflectance of the surface associated with the second pixel or whether the first surface was illuminated with twice the light as the second surface, or some combination of differences in illumination level and

reflectance. Specular highlights would be impossible to distinguish from highly reflective areas of the scene. If specular highlights were interpreted as lambertian reflectors, then the overall absolute reflectances in the scene would be under-predicted. This pseudo-reflectance should not be dismissed as being without value, though. By using such an image, as was done in the previous example, a resultant image can simulate the appearance of replacing the original store light bulbs with the living room light bulbs along with the same levels of uneven illumination.

For a second example, the reflectance reproduction of a painting is described. This example contains a far less complicated set of transformations than the previous, wallpaper, example. Here, a multi-channel scanner would be chosen to capture the input image. The input profile relates scanner digits to reflectance. For an output device, a multi-ink printer would be the desired rendering engine. An output profile relates reflectance to printer digits. The reflectance image is manipulated by the output transformation which yields printer digits. The digits are printed and the reflectance match is consummated. Figure 4 is a cartoon representing this example and Figure 5 shows the color management steps.




Figure 4: Spectral color management would provide a way to match the reflectance of a captured object.




Figure 5: Block diagram illustrating the logical color management system actions which would accompany the painting reproduction example described above. The painting reflectance image is converted to printer digits and then printed.

3. THE PROBLEM: SPECTRA TO DIGIT TRANSFORMS

The painting reflectance reproduction outlined in Figure 5 seems like a comparatively simple task when compared to the wallpaper matching example, found in Figure 3. Although it has far fewer transformations than the previous example, it is deceptive in its apparent simplicity. The major question being explored by this paper concerns how an efficient transform from spectra to output digit might be designed. In Figure 5, the box labeled “Apply Output Profile Transform to Printer Digits” is central to this question. It is there that a transform which accepts spectra must determine which digits are to be fed to the output device. A spectral sampling frequency might be every 10nm from 400nm to 700nm. This would mean that every pixel for an image with such a sampling would have 31 values associated with it, in essence a 31-dimensional image. To make the problem easier, one might sample the spectra every 20nm yielding a 16-dimensional image or every 40nm for an 8- or 9-dimensional image. Further downsampling is possible, but the loss of spectral detail will eventually return the image to the wide-band domain with no advantage over current RGB devices. Traditional color management tools available to carry out transforms in highly efficient manners include matrix multiplies, one-dimensional lookups and multi-dimensional lookups. Applying one-dimensional lookups to the individual bands of a multi-channel image would not be a costly enterprise. Even performing a matrix multiply would not be particularly expensive for modern computers. It is the application of a multi-dimensional lookup table to an 8-, 16- or 31-dimensional image which would be overwhelming due to the memory requirements of such a huge table. Due to the non-linear characteristics of printer physics, it is inevitable that a transformation between spectra and printer digits will eventually require a computationally expensive step. Within the realm of color management when such transforms need to be applied in an efficient manner, multi-dimensional lookups are the transforms of choice.

4. METHOD

The method described here can easily be customized to particular needs where tradeoffs can be biased among desired precision, processing time requirements, available computational power and memory constraints. A statement of the aim of the method is as follows: reduce the dimensionality of the input to no more than the dimensionality of the output. In other words, if there are input spectra described by N sample points per pixel and the image is being processed for output to an M-ink per pixel printer, the task begins as an N to M problem but reduces to an M to M problem. The savings can be tremendous. If transforming from a 31 sample point per pixel input to a 6-ink printer, the problem reduces from a 31- to 6-dimensional transform to no greater than a 6- to 6-dimensional transform. In our own case, the model we have been exercising for a 6-ink printer⁹ assumes no more than 4 inks printed for any particular pixel. For such a case, the problem is further reduced to a 4- to 4-dimensional transform.

Reduction of input dimensionality from N to M is handled by carefully choosing a set of M spectral curves to be used for analyzing the source spectra. While processing a source pixel, weightings are derived for the M chosen curves such that the weighted sum best represents the original N point spectrum. It is these weightings which are used as lookup values for the final M to M lookup as shown in Figure 6.




Figure 6: Reduction of dimensionality, starting with a 31- to 4-dimensional problem and reducing it to a 4- to 4-dimensional problem. (a) 31 input bands. (b) Each pixel represents 31 samples of a reflectance spectrum. (c) Derive weightings for a set of fixed spectra such that the sum of the weighted spectra is an approximation of the original spectrum. (d) Weights are used as input to a 4- to 4-dimensional lookup table.

In a standard colorimetry-based image processing approach, the printer characterization starts with either the construction of a physical model of the printer or, more commonly, a brute force gamut-wide sampling takes place. This printer model can be used to report the colorimetry expected from rendering a set of printer digits. If a physical model were derived and it was mathematically invertible, then the inverse relationship would be easily found. For a sampling-based model or for mathematically non-invertible models, search methods are employed to efficiently guess which digits might yield a particular colorimetric value. Typically off-line, a three-dimensional lookup table indexed by regularly spaced colorimetric values will be built. Search methods are used to populate the lookup table with printer digits. Figure 7 demonstrates the building of such a lookup table.




Figure 7: Process for creating a printer characterization for a colorimetric-based color management system.

The spectral approach has a flow similar to the colorimetric characterization process. It is illustrated in Figure 8. A printer model is needed which reports reflectance given a printer digit. This might be a physically derived model or one based upon the measurement of a plethora of printed samples. As in the colorimetric approach, the inversion process would use




Figure 8: Process for creating a printer characterization for a spectral-based color management system.

optimization techniques to exercise the printer model until a grid request was satisfied. It is the production of the grids which deviates significantly from the colorimetric approach. If, alternatively, there were no deviations and the colorimetric model had been faithfully followed, a regular grid of spectral domain values would need to be processed. As already pointed out, this could be an 8-, 16- or 31-dimensional grid. Obviously such lookup tables would be absurdly large. Our approach, instead, is to choose a set of M spectral curves which will be referred to as “grid curves.” A regular M-dimensional grid of gain factors is constructed where each dimension is associated with a single grid curve. For each node in the grid, the gain factors are applied to the respective grid curves and the M gained spectra are summed. The summed spectra produce a unique spectrum and this spectrum becomes associated with the gain factors which lookup to that node. This unique spectrum is then inverted and the derived digits are stored at the node. Thus, if a spectrum were completely decomposed into the M grid curves, the gain factors for that spectrum could be looked up in the table and the appropriate printer digits would be found such that when those digits were fed to the printer, the reflectance spectrum would be rendered. This is shown in Figure 9.




Figure 9: Image processing with spectral lookup method. Each reflectance image pixel is decomposed into gain factors associated with the grid curves. The lookup table constructed in Figure 8 expects these values as input and reports printer digits.

5. CONCLUSIONS

A method for reducing the dimensionality demands of spectral color management has been introduced. Without such approaches, spectral color management would remain a very slow process or would make memory demands far exceeding today's capabilities. The method is fairly simple. It involves creating a low-dimensional lookup table or set of lookup tables which can be accessed with values derived through the analysis of image spectral curves. These analyses involve decomposing the spectral curves into a standard set of curves known as “grid curves.” Future discussions will include describing how to choose grid curves, decomposition approaches, and alternative spectral spaces within which such decomposition has greater physical meaning and thus lookup error is reduced.

6. ACKNOWLEDGEMENTS

The authors would like to express their gratitude to Fuji Xerox Corporation for its support of this work. We would also like to acknowledge the many members of the Munsell Color Science community who are working on research which either directly feeds this study or which conjures new questions helping to guide our approach. We recognize Mark Fairchild for encouraging the realization of a spectral image visualization tool, now grown into Spectralizer.¹² The process around designing and building that tool has helped to cement many of the issues discussed in this paper.

7. REFERENCES

1. Good examples of such papers may be found in Proc. of the International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives (1999) and papers from sessions 1, 2 and 3 in Proc. of SPIE 3963 (2000).
2. S. Tominaga, , “Spectral Imaging by a Multi-Channel Camera”, Proc. of SPIE 3648, pg. 38 (1999).
3. P. D. Burns and R. S. Berns, “Analysis of multispectral image capture”, Proc. 4th IS&T/SID Color Imaging Conference, 19 (1996).
4. F. H. Imai and R. S. Berns, “High-resolution Multi-Spectral Image Archives: A Hybrid Approach”, Proc. of The Sixth Color Imaging Conference: Color Science, Systems, and Applications, 224 (1998).

5. G. M. Johnson and M. D. Fairchild, "Full-Spectral Color Calculations in Realistic Image Synthesis," IEEE Computer Graphics and Applications, 19, 4 (1999).
6. D. Tzeng," Spectral-Based Color Separation Algorithm Development for Multiple-Ink Color Reproduction", Ph.D. dissertation, Rochester Institute of Technology (1999).
7. M. Rosen, X. Jiang, "Lippmann2000: A spectral image database under construction", Proc. of the International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives (1999).
8. F.H. Imai, M. Rosen, R. Berns, D. Tzeng, "Spectral Reproduction from Scene to Hardcopy II: Image Input and Output", Submitted to EI 2001.
9. K. Iino and R. Berns, "Building Color Management Modules Using Linear Optimization II. Prepress System for Offset Printing", Journal of Imaging Science and Technology, 42, 2 (1998).
10. P. Hung, "Color Reproduction Using Spectral Characterization", International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives, pp. 98-105 (1999).
11. B. Hill, "Color capture, color management and the problem of metamerism: does multispectral imaging offer the solution?", Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts V, R. Eschbach, G. G. Marcu, Editors, Proc. of SPIE 3963 Bellingham, WA, pp.2-14 (2000).
12. M. Rosen, M. Fairchild, G. Johnson, D. Wyble, "Color Management within a Spectral Image Visualization Tool", submitted to CIC 2000.