

Spectral Color Processing using an Interim Connection Space

Mitchell R. Rosen and Noboru Ohta
Munsell Color Science Laboratory, RIT
Rochester, New York

Abstract

The use of an Interim Connection Space (ICS) is proposed as a means for extending the concept of device independent color management to support spectral imaging. Color management, in its standard practice, relates color rendering capability of devices through a Profile Connection Space (PCS). The International Color Consortium (ICC) defines a set of encodings for PCS derivable from CIEXYZ. Multi-channel imaging systems where the goal is the reproduction of image spectra cannot use a colorimetrically-based PCS. The obvious solution is to use a new PCS based on spectral reflectance, transmittance or radiance. In the presence of a spectral PCS, though, a straightforward analog to the typical ICC color processing approach breaks down due to the likely high dimensionality of a the new PCS. To alleviate such problems the ICS is introduced as a standard dimensionally reduced image processing stage. Eigenvectors were rejected as candidate bases for ICS. Pseudoprimaries were introduced as potential ICS bases.

Background

The few reported approaches to spectral printer output of complex scenes have relied on highly computationally intensive model inversion on a pixel-by-pixel basis. Typical transformation approaches¹ require the use of a non-linear iterative optimization step. It is only due to the fantastic computational power of recent high-end computers that one would seriously consider, even for demonstration purposes, the inversion of a spectral printer model on every pixel of a complex image. For production of many images, such an approach is impractical.

Traditional colorimetrically-based color management gets around the computational bottleneck of inverting a printer model for every pixel by collapsing large numbers of pre-inverted answers into a large multi-dimensional lookup table. Assuming the sampling density of the lookup table is sufficient, standard interpolation techniques between the samples enables an efficient and accurate means for transforming the image.

It is tempting to look for ways to help spectral color management enjoy the benefits of off-line pre-inversion stored in a lookup table much in the same way that colorimetric color management does. For camera to printer

processing, the ICC approach starts with a digital image made up of camera digital counts and transforms the image to printer digital counts based on a mapping that logically passes through a Profile Connection Space (PCS) based on CIEXYZ. Spectral color processing could follow a similar route where the Profile Connection Space is based on spectral reflectance, transmittance or radiance.² A spectral camera to printer transformation chain is illustrated in Figure 1.

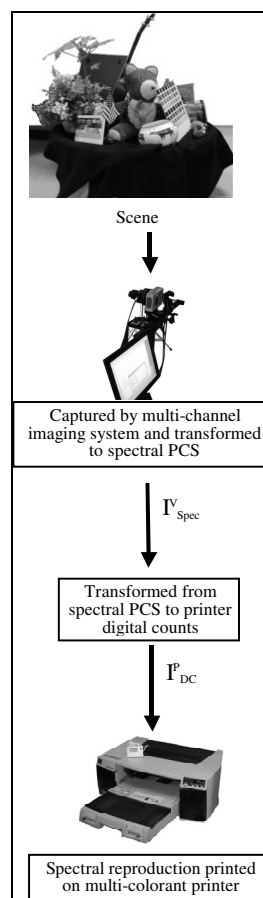


Figure 1. Spectral image processing workflow from scene to print.

Equation 1 shows the formula for calculating the size of a lookup table.

$$\text{Size}_{\text{bytes}} = L_i^{B_i} \times B_o \times b_o \quad (1)$$

where b_o is the number of bytes per output band; B_o is the number of output bands; B_i is the number of input bands; L_i is the number of levels per input band assuming that all input bands have the same number of levels; and, $\text{Size}_{\text{bytes}}$ is the calculated size in bytes of the lookup table.

Table I assumes that all output bands are one byte each ($b_o = 1$).

Table I: Size in Bytes of Lookup Tables

B_i Number Bands In	B_o Number Bands Out	L_i Levels per Band In	$\text{Size}_{\text{bytes}}$
3	6	17	30KB
6	6	17	145MB
9	6	17	700GB
31	6	17	8×10^{27} GB

Because colorimetric connection spaces are typically based on a derivation of CIEXYZ, they are 3-dimensional. Table I shows that 3-dimensional lookup tables tend to be a reasonable size. For the parameters shown, a 3-dimensional lookup table would be only 30KB. On the other hand, the connection space for spectral image processing would be spectral reflectance, spectral transmittance or spectral radiance. For any of these cases, the dimensionality would be very large, typically 31 samples per spectrum. From the table it is clear that a lookup table with 31 dimensions would be much too large to consider.

Interim Connection Space

As a starting place, it has been suggested that a low-dimensional *interim* connection space (ICS) be formed³. An ICS must have several important qualities. There must be a low-computational-cost transformation from spectra to ICS and from ICS back to spectra. It should be low-dimensional so that a reasonably sized lookup table may be created from it to printer digital counts. For the case of printing to a 6-ink printer, choosing a 6-dimensional ICS has merit since there are at most 6 degrees of spectral freedom for such a printer.

In image processing, the lowest computational-cost transformations are generally those that can be implemented through a series of 1-dimensional lookup tables and matrices. While there do exist other specialized fast software operations and also hardware implementations that can dramatically reorder relative speed of operations, for these discussions, computationally inexpensive processing chains for converting between spectral PCS and ICS should contain nothing more than a string of 1-dimensional LUTs and matrices.

Each dimension of ICS is associated with a basis function. In Reference 3, these are called “grid curves.” Mixing of grid curves can be performed within the linear domain or some other domain that only requires the use of our basic operations, 1-D LUTs and matrices, to get there.

The image processing goal illustrated in Figure 1 can now be expanded to include the use of the low-dimensional Interim Connection Space. Figure 2 shows the workflow where the transform from spectral PCS to printer digital counts is broken into two steps: transformation from PCS to ICS, and then transformation from ICS to printer digital counts. A reasonably sized multi-dimensional LUT will be introduced to implement that second transformation.

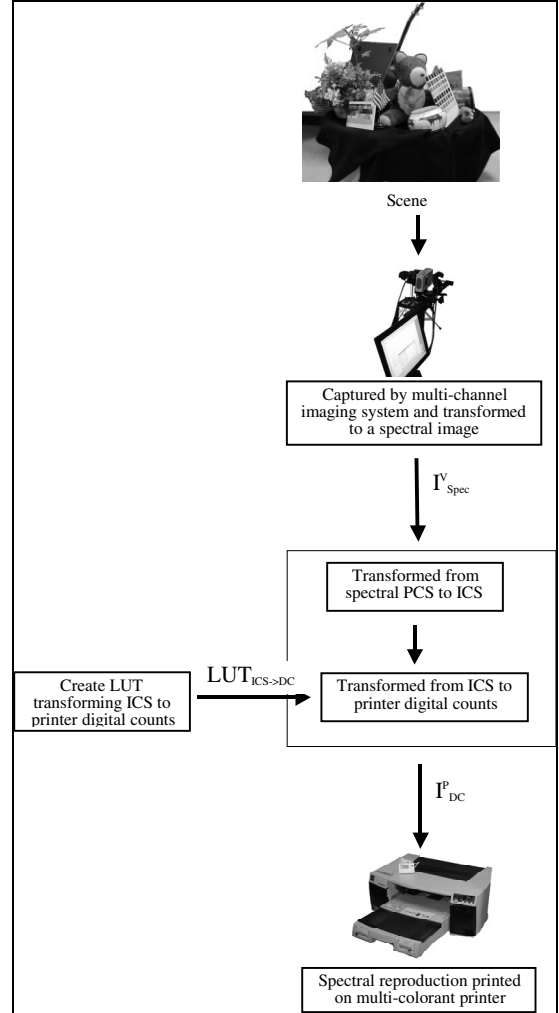


Figure 2. Image processing workflow from scene to print through low-dimensional ICS.

Once the capability illustrated in Figure 2 is implemented, further efficiencies for faster runtime transformation of images may be introduced. However, the Figure 2 workflow is a basic necessity even if additional concatenations take place in certain implementations of a spectral image processing system. Based on the byte quantities shown in Table I, it is desirable to limit the dimensionality of an ICS to 6 or fewer so that the size of $\text{LUT}_{\text{ICS} \rightarrow \text{DC}}$ may be maintained at a reasonable size.




Figure 3. Flowchart for building the LUT that transforms ICS to printer digital counts.

All of these steps take place off-line:

1. Spectrally characterize the printer.^{2,4}
2. Choose Interim Connection Space. The interim space should be relatively low-dimensional and have a computationally inexpensive relationship with spectra. While it is not necessary for the space to have physical meaning, the more physically justifiable, the more likely that a grid of *interim* values will have limited “wasted regions.” Such regions would be physically unrealizable spectra or spectra out-of-gamut to the printer.
3. Produce $\text{GRID}_{\text{interim}}$ of ICS by a regular sampling in each of the space’s dimensions.
4. Convert to $\text{GRID}_{\text{interim} \rightarrow \text{Spec}}$ by applying the inexpensive transformation to spectra.
5. Use standard LUT inversion routines to determine for each node point in the grid, the printer digital counts that would provide the associated spectrum. The relationship between the grid of *interim* connection space values and printer digital counts is a lookup table, $\text{LUT}_{\text{ICS-DC}}$.

Eigenvectors Rejected as ICS Candidate Bases

As mentioned, the goal of the ICS is to provide a spectrally derivable bottleneck where the dimensionality of the system is reduced to a reasonable size so that a lookup table may be invoked to provide the final transformation to printer digital counts. ICS must be readily transformable to spectra, hopefully able to be implemented through the use of

nothing more than a series of 1-dimensional lookup tables and matrices. 6 dimensions make good sense because a 6-ink printer has 6 degrees of spectral matching freedom and, according to Table I, a 6-dimensional LUT is at the limit of reasonable size.

One obvious candidate for the ICS is one based on a set of eigenvectors derived through Principal Components Analysis (PCA) of samples from either the scene’s or the printer’s spectral gamut. This is particularly appealing because a simple matrix describes the relationship between eigen values and spectra. This approach was attempted. In the end, the approach was proven to be unworkable, a number of interesting observations were collected.

8 spectral images taken by a multi-channel camera system were selected for this analysis. They are shown in Figure 4. The top 6 eigenvectors describing the relative reflectance of a color target captured in Image 1 were derived. The histograms of eigen weightings for each image show interesting trends. The eigenvalues for the full set of images tend to strongly cluster around a common set of means for all the vectors with the possible exception of the first eigenvector. There, a number of images are bi-modal. The summary histograms are found in Figures 6 to 11. Figure 5 displays the key for these figures.

The 8 spectral images were decomposed into weighting for the eigenvectors. At each pixel there were 6 eigenvalue weightings. Based on statistics describing the decomposition $\text{GRID}_{\text{interim}}$ was built as a sampling of eigenweights. Step 4 of Figure 3 followed, creating a $5 \times 5 \times 5 \times 5 \times 5 \times 5$ lookup table that converted from the eigenvector-based ICS to spectra.




Figure 4. 8 spectral images used to evaluate ICS choices.




Figure 5: Key for Figures 6 – 11.




Figure 6: Eigenvector 0 histograms for 8 images of Figure 4.




Figure 8: Eigenvector 2 histograms for 8 images of Figure 4.




Figure 7: Eigenvector 1 histograms for 8 images of Figure 4.




Figure 9: Eigenvector 3 histograms for 8 images of Figure 4.




Figure 10: Eigenvector 4 histograms for 8 images of Figure 4.




Figure 11: Eigenvector 5 histograms for 8 images of Figure 4.

There are many ways to evaluate the effectiveness of a lookup table. One simple way is to discern the quantity of wasted space in the lookup table. In this case, all spectra associated with grid points were evaluated to see if they had any reflectance factors outside the range of 0 to 1. Recall that the spectra were calculated by applying the weightings at a particular grid node to the eigenvectors.

In the full 6-dimensional grid, there were 5^6 grid points. Of this full set of 15,625 eigenvector calculated spectra, almost every spectrum was found to be invalid. There were, in the entire set, only two exceptions! This means that although there is quite a large spread of values in the histogram of eigenweightings for the images, the 6-dimensional spread must be very tight. Unfortunately, this ruled out eigen vector-based spaces as appropriate for use as ICS.

This finding is not surprising. The method for deriving Principal Components demands that each vector is orthogonal with all the others. A consequence of this basic aspect is that only one of the vectors can be all positive. Thus, when combining the vectors in an unconstrained manner as is the case in building a LUT, unrealizable spectra should be expected as common. Now, it is

understood, physically acceptable spectra have an extremely low probability of being produced.

Pseudo-Primaries as Bases for ICS

There are many arguments for designing an ICS so that it is particularly efficient at describing a printer's spectral gamut. Far less important is the need to describe spectra that are potentially captured by a camera but unprintable on the printer of interest. Of very little importance is the ability to describe spectra outside the limits of physical surface colors.

It was decided to build an ICS based on the reflectance spectra of the printer's inks. The true physics of how the printer inks combine to produce spectra was left aside for the ICS. A simple mixing model, as called for in Reference 3, was instituted, instead. The mixing model is a use of the Beer-Lambert assumption that ink spectra multiply in linear space and add in log space. Only two constraints were built into this pseudo color mixing space.

1. When one pseudo-ink is at unity concentration and all other pseudo-inks are at 0 concentration, the resultant spectrum is that of the measured single full ink from the printer.
2. When all 6 pseudo-inks are at unity concentration, the resultant spectrum is that of the measured combination of all six full inks from the printer.

All other combinations of the pseudo-inks relied on the Beer-Lambert assumption.

Constraints 1 and 2 were implemented through a derived spectral gain and offset. Equation 2 shows the computation chain for transforming from pseudo-concentrations to reflectance.

$$\begin{aligned} & \text{gain}(\lambda) \times \text{reflectance}_{\text{paper}}(\lambda) \\ & \times \prod_i (\text{reflectance}_{\text{fullinks},i}(\lambda)^{\text{concentration}_i}) \\ & + \text{offset}(\lambda) \end{aligned} \quad 2)$$

where i spans from 1 to 6; concentration_i is the pseudo-concentration for the i^{th} pseudo-ink, $\text{reflectance}_{\text{fullinks},i}(\lambda)$ is the measured spectral reflectance of the i^{th} printer ink at full area coverage with the measured reflectance of paper divided out; $\text{reflectance}_{\text{paper}}(\lambda)_r$ is the measured spectral reflectance of the paper medium on which the printer is printing; and, $\text{gain}(\lambda)$ and $\text{offset}(\lambda)$ are the spectral gains and offsets derived to maintain constraints 1 and 2, above.

Equation 2 is a non-linear from pseudo-concentration to reflectance. In order to put it in a format that would allow its implementation with 1-dimensional LUTs and matrices, Equation 3 is presented for calculating spectra from pseudo-concentrations.

$$\begin{aligned} & \exp[\log(\text{gain}(\lambda)) + \log(\text{reflectance}_{\text{paper}}(\lambda)) \\ & + \log(\text{reflectance}_{\text{paper}}(\lambda) \times \text{concentration}) \\ & + \text{offset}(\lambda) \end{aligned} \quad (3)$$

where *concentration* is an array of pseudo-concentrations; and, $\text{reflectance}_{\text{fullinks}}(\lambda)$ is an array of the measured spectral reflectances of the full coverage printer inks.

An implementation of Equation 3 using 1-dimensional lookup tables and matrices could look like Figure 5.




Figure 5. Implementable processing chain from pseudo-concentrations to reflectance.

For this ICS example, a set of grid nodes were chosen based on the following criteria:

1. First, a uniform sampling was taken between pseudo-concentrations of 0 and 1 in .25 increments.
2. Second the spectra of all printer individual inks printed with area coverages of 0%, 25%, 50%, 75% and 100% were converted to pseudo-concentrations. If any of the derived pseudo-concentrations were negative or greater than 1, that pseudo-concentration was added into the grid. Also, if the derived pseudo-concentration was more than .1 away from those already in the sampling, it was added to the grid. If the derived pseudo-concentration was very close to one of the original samplings, it replaced that pseudo-concentration in the grid.

This grid is different from those discussed previously because there are an uneven number of levels depending upon the dimension. A straightforward generalization of Equation 1 allowing for non-uniform LUT levels is shown in Equation 4.

$$\text{Size}_{\text{bytes}} = \prod_{n=1}^{B_i} L_{i,n} \times B_o \times b_o \quad (4)$$

where b_o is the number of bytes per output band; B_o is the number of bands out; $L_{i,n}$ is the number of levels in the n^{th} dimension; B_i is the number of bands in; and $\text{Size}_{\text{bytes}}$ is the size of the lookup table in bytes.

The grid used for this demonstration was $7 \times 7 \times 8 \times 8 \times 7 \times 7$. The size in bytes for a lookup table built for output to a 6-ink printer would be less than 1MB.

Image 1 from Figure 4 contains a target consisting of 231 Munsell Chips. The set of spectra from the Munsell Chips were converted to pseudo-concentrations, transformed to printer digits and then reconverted back to spectra via the printer's characterization lookup. The original and simulated printed sets of spectra were compared under equal energy illuminant. Mean ΔE^*_{ab} were under 5, but maximum exceeded 12.

An analysis of pseudo-concentrations for the images showed that a great many of the image pixels fell outside the pseudo-concentration ranges for the table. Thus the lookup table was too limited for those pixels. To address this situation, additional pseudo-concentrations were added to each dimension of the grid in order to envelop all image pseudo-concentrations. The new table was now $8 \times 8 \times 9 \times 9 \times 8 \times 8$. By adding two nodes per dimension, the size in bytes for the new lookup table, according to Equation 4, raised to just under 4MB. Mean ΔE^*_{ab} for the simulated print fell to under 4, and maximum fell to under 8.

Conclusions

The concept of the Interim Connection Space has been introduced. In a spectral color management scheme, the ICS will be in a processing step that reduces the dimensionality of a spectral PCS. Eigenvectors were ruled out as axes for ICS. Pseudo-primaries and the development of pseudo-concentrations as ICS values were explained.

References

1. L. Taplin and R. Berns, Spectral Color Reproduction Based on a Six-Color Inkjet Output System, *Proc. 9th CIC*, pp. 209-213 (2001).
2. M. Rosen, M. Fairchild, G. Johnson and D. Wyble, Color Management within a spectral image visualization tool, *Proc. 8th CIC*, pp.75-80 (2000).
3. M. Rosen, F. Imai, X. Jiang and N. Ohta, Spectral reproduction from scene to hardcopy II: Image Processing, *Proc. SPIE 4300*, pp. 33-41 (2001).
4. M. Rosen, L. Taplin, F. Imai, R. Berns and N. Ohta, Answering Hunt's Web Shopping Challenge: Spectral Color Management for a Virtual Swatch, *Proc. 9th CIC*, pp. 267-273 (2001).

Biography

Mitchell Rosen is a Senior Color Scientist with the Munsell Color Science Laboratory at RIT. His recent research has concentrated on data- and computational-efficiencies for use in spectral reproduction. He has a BS in Computer Science from Tufts University and a Ph.D. in Imaging Science from RIT. He is Color Imaging editor for the Journal of Imaging Science and Technology.