### Introduction

In the Academy Color Encoding System, an Input Transform processes non-color-rendered RGB image values from a digital camera system's capture of a scene lit by an assumed illumination source into white-balanced ACES RGB relative exposure values. The recommended procedure for creating a "best practices" Input Transform requires spectral sensitivity data for the camera. Spectral sensitivity data is rarely provided by camera manufacturers, so it is sometimes necessary to measure the sensitivities of the color filter array on the sensor.

In this experiment, a monochromator was used to expose a camera with narrow bands of light centered at regular wavelength increments across the visible spectrum. Data was collected and processed to derive spectral sensitivity data. The validity of the spectral sensitivity data was evaluated independently by testing linearity of sensor response as well as color accuracy between theoretical calculated data and empirical photography. Once confidence in the linearity and spectral sensitivity data was established, the procedure specified in "Academy P-2013-001: Recommended Procedures for the Creation and Use of Digital Camera System Input Device Transforms (IDTs)" was followed. Finally, the results of applying the IDTs for four different cameras were compared to the ACES Reference Input Capture Device (RICD) and to each other.

## **Data Collection**

The procedure was performed with four digital stills cameras:

- a Nikon D810 using a Nikon 24-70mm f/2.8 lens
- a Sony a7 using a Zeiss 24-70mm f/4.0 lens
- a Canon 5D Mark II using a Nikon 50mm f/1.4 lens
- a Canon 5D Mark III using a Canon 24-70mm f/2.8 lens

For each digital still camera, a total of 120 images were collected:

- monochromator "spot" images from 350 nm to 800 nm in 5 nm increments (91 images)
- a black frame with the lens cap on and camera covered (1 image)
- TE226 chart "normal" exposure (1 image)
- TE226 chart exposure wedge -3 stops to +3 stops in ½ stop increments (13 images)
- Macbeth chart "normal" exposure (1 image)
- Macbeth chart exposure wedge -3 stops to +3 stops in ½ stop increments (13 images)

The following spectral readings were collected:

- spectral power distribution of integrating sphere backlight (1 measurement)
- spectral radiance of TE226 chips plus int. sphere backlight (45 measurements)
- spectral radiance of Nikon 24-70mm lens plus int. sphere backlight (1 measurement)
- spectral radiance of Zeiss 24-70mm lens plus int. sphere backlight (1 measurement)
- spectral radiance of Nikon 50mm lens plus int. sphere backlight (1 measurement)
- spectral radiance of Canon 24-70mm lens plus int. sphere backlight (1 measurement)

In addition, the computer attached to the monochromator produced a text file of power readings (91 total) at each wavelength.

#### Measurement Procedure

For each digital still camera, raw images of the monochromator's integrating sphere were captured at 5nm increments from 350nm to 800nm (91 images). A black frame was also captured with the lens cap on and camera covered.

The Labview software that controls the monochromator also communicates with a power meter. The radiant flux at each wavelength was automatically measured and recorded into a text file for later use in normalizing the data.

After capture of the 91 frames, the raw images were processed to linear camera RGB by using the 'dcraw' command-line utility to black frame subtract, white balance, debayer, and linearize into 16-bit linear TIFF files.

Using a Matlab script, the resulting TIFF frames were opened and a large square region (>500 pixels square) was averaged from the center of the colored spot in each image to get the native exposures in red, green, and blue channels at each wavelength increment (Figure 1). The averaged values were assembled into a table where each row corresponds to the RGB values at a wavelength.



Figure 1. Example of the region of interest read from the monochromator "spot" image.

Because the illuminant in the monochromator is tungsten-based, radiant flux is not equal across the spectrum; there is much higher power toward the red end of the spectrum. The power data was extremely consistent among the measurement runs (Figure 2).



To compensate for any variations in power, the averaged RGB values from each wavelength snapshot were normalized by dividing each wavelength triplet by the radiant flux at that wavelength. When the resulting RGB values are plotted per wavelength, the spectral sensitivity curves emerge (Figure 3).



For completeness, the spectral transmittance for each lens was also calculated by dividing the integrating sphere backlight from the spectral radiance of the backlight as measured through the lens. There was noise in the shorter wavelengths where power was low, so the lens data was extrapolated to follow the trend of the data across that region. The extrapolation only affected wavelengths less than 400 nm which is a region with little weighting in any spectral calculations since camera sensitivity is also very low in those regions. Therefore, the impact of the smoothing by extrapolation is negligible.

It is recognized that lens transmittance is part of the capture system and modulates the spectral sensitivity data. However, the lens was present during capture of the monochromator images and also the test charts, so nothing is actually done with the lens transmittance in this experiment, since it is already factored into the spectral sensitivity data. If one wanted an IDT exact to the camera sensitivities only (i.e. no lens), then this data could theoretically be factored out of the spectral calculations.



## **Verification of Linearity**

Linearity can be tested by comparing values for a test target at regular exposure increments. If 'dcraw' is correctly undoing the camera Opto-Electronic Conversion Function (OECF), then a doubling in exposure (1 stop difference) should equal a doubling in camera values. Also, linearity should hold across a series of exposures.

To confirm this visually, RGB values from an exposure series of the TE226 chart were plotted as in Figure 5. The exposures form mostly straight lines on a log-log plot (with the exception of where clipping takes effect and for the "black" filter where SNR is low), and therefore indicate that linearity is achieved. Ratios of RGB values for exposures 1 stop apart were also compared and confirmed to follow a linear response with a pattern of either ½ or 2x the adjacent values.



#### **Verification of Spectral Sensitivities**

The normal exposure of the TE226 chart was used to verify the measured spectral sensitivities by comparing it to theoretical RGB exposure values calculated using the spectral radiance of the integrating sphere backlight as modulated by each Wratten filter and the measured spectral sensitivities. If spectral sensitivities were measured correctly, then the calculated RGB exposure values should match the actual RGB exposure values.

The captures of the TE226 chart were processed the same as the monochromator spot images, using the 'dcraw' command-line utility to black frame subtract, white balance, debayer, and linearize files into a 16-bit linear TIFFs.

The actual TE226 image was then opened in Matlab with square regions averaged from each Wratten filter chip to get RGB values as captured by the camera. These RGB values were then scaled using the center gray patch from the bottom row to set exposure and white balance.

The result for each camera can be seen in Figure 6. The top half shows the actual RGB values and the bottom half shows the theoretical RGB values. For display, the RGB values have had a gamma of ½ applied in order to brighten their appearance.

Note: Colors in Figure 6 appear muted because no adjustment has yet been made to the color primary encoding and no rendering has been applied.



Figure 6. TE226 charts for each camera. The top half shows the actual RGB values and the bottom half shows the theoretical RGB values.

For a different viewpoint, the values for the grayscale row at the bottom of the TE226 chart were plotted for each camera (Figure 7).



. (blue).

# **Creating the Input Transform**

Once confidence in the spectral sensitivity data has been established, optimizing a color transformation matrix to convert camera RGB to ACES is straightforward by following the procedure described in Academy P-2013-001. One can obtain matrices and scale factors for each camera for tungsten and daylight illumination. Note that there is an infinity of possible Input Transform matrices that can be generated, but this infinity is usually reduced to a minimum of two – one for "tungsten" and one for "daylight" balance.

The matrices have been written into CTL files and are available in the supplemental files.

## Verifying the Input Transform

There are two comparisons performed and tabulated. The first compares each camera to the RICD. The RICD is the theoretical "ACES camera", but is not representative of real cameras because it has spectral sensitivities that are distinctly different from real color filters used in physical capture devices. Therefore, the second and more meaningful comparison is to designate one of the actual cameras as the reference to which the other cameras are then compared. This comparison demonstrates the cameras' abilities to match each other and avoids the distraction of the RICD values potentially being quite different from results of actual exposures.

By comparing the Macbeth charts captured by each camera processed through the tungsten IDT matrix for each camera, the degree of match between cameras can be assessed. Because

all photography captured in this experiment used incandescent lighting, only the 3200K IDT is verified. The physical light source used for capture will never exactly match a theoretical design illuminant. Therefore, to make all things equal, a one-time balance using one of the gray chips on the Macbeth was used to create scale factors for each camera that would normalize exposure and white balance. These scale factors are derived once and then can be applied to all shots under those same lighting conditions and camera settings. For the ColorChecker chart, the third chip from the bottom right (Neutral 5) was balanced in ACES RGB space so that it would equal [0.18 0.18 0.18].

The ACES values were then processed through the RRT and Rec. 709 ODT in order to create display-referred imagery to compare. Output code values were converted to XYZ tristimulus values, the CIE L\*a\*b\*, and finally dE2000.



Canon 5D Mk II





Figure 8. Result of balanced and exposure-adjusted ACES images as rendered through the RRT and Rec 709 ODT.

Figure 9 shows the images tiled together using Photoshop such that each patch can be compared in a 2x2 layout. The ordering is the same as for all the other images in this document.



Figure 9. Tiled and overlaid version of the images from Figure 8. Nikon D810 (top left); Sony a7 (top right); Canon 5D Mk II (bottom left); Canon 5D Mk III (bottom right)





Figure 10. (left) Macbeth ColorChecker - Nikon D810 (top left); Sony a7 (top right); Canon 5D Mk II (bottom left); Canon 5D Mk III (bottom right); RICD (center)

(right) Macbeth ColorChecker - Nikon D810 (top left); Sony a7 (top right); Canon 5D Mk II (bottom left); Canon 5D Mk III (bottom right

	Compared to RICD					Compared to Nikon D810			
	D810	a7	5D Mk II	5D Mk III		a7	5D Mk II	5D Mk III	
Dark skin	2.680	2.455	3.293	2.887		0.326	0.941	0.836	
Light skin	7.585	7.576	6.226	5.091		0.431	1.664	2.890	
Blue sky	3.052	3.236	2.946	3.144		0.810	0.892	1.497	
Foliage	2.398	3.138	2.638	2.636		0.742	1.543	0.680	
Blue flower	6.286	5.470	6.711	5.916		1.163	0.834	1.401	
Bluish green	3.566	3.377	3.263	3.404		0.955	1.713	2.501	
Orange	3.175	3.057	3.641	3.151		0.411	0.492	0.768	
Purplish blue	2.059	2.258	2.644	3.164		0.360	0.603	1.253	
Moderate red	3.958	4.171	4.070	3.403		0.281	0.429	0.715	
Pruple	2.987	1.840	2.980	2.429		1.264	0.686	0.683	
Yellow green	3.613	4.095	2.840	3.504		0.753	1.455	0.914	
Orange yellow	0.991	1.148	1.313	1.435		0.220	1.103	0.469	
Blue	1.415	1.389	1.689	2.177		0.365	0.622	0.966	
Green	4.276	4.237	3.459	3.835		0.703	1.080	0.779	
Red	5.359	4.888	5.483	5.355		0.724	0.400	1.263	

Table 1. dE2000 values for each camera, as compared to RICD and as compared to Nikon D810

Yellow green	2.351	2.302	2.219	2.322	0.179	0.952	0.089
Magenta	4.348	4.264	4.611	4.567	0.585	0.341	0.595
Cyan	3.012	2.239	2.151	2.554	1.012	1.916	2.870
White	1.401	1.044	1.845	1.830	0.935	1.047	1.286
Neutral 8	2.406	2.119	2.790	3.059	0.870	0.736	2.226
Neutral 6.5	3.817	3.743	4.144	4.597	0.337	0.709	1.821
Neutral 5	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Neutral 3.5	0.590	0.612	1.409	0.868	0.359	1.552	0.807
Black	0.795	0.899	1.388	1.508	0.839	1.440	1.519
Mean	3.005	2.898	3.073	3.035	0.609	0.964	1.201

To assess performance on a larger variety of spectra, a similar comparison was done for the TE226 chart.





Figure 11. (left) TE226. Nikon D810 (top left); Sony a7 (top right); Canon 5D Mk II (bottom left); Canon 5D Mk III (bottom right); RICD (center)

(right) TE226. Nikon D810 (top left); Sony a7 (top right); Canon 5D Mk II (bottom left); Canon 5D Mk III (bottom right)

	Compared to ACES RICD					Compared to Nikon D810			
	D810	a7	5D Mk II	5D Mk III		a7	5D Mk II	5D Mk III	
R1C1	3.103	4.004	2.304	2.860		0.929	0.927	0.561	
R1C2	1.136	2.093	0.760	1.070		1.509	1.449	1.839	
R1C3	3.410	3.712	1.381	1.637		0.621	3.219	2.709	
R1C4	4.730	5.136	4.338	4.905		1.173	0.670	0.470	
R1C5	2.971	3.005	3.507	3.782		0.404	1.388	1.637	
R1C6	2.079	2.147	2.539	2.488		0.652	1.890	1.660	
R1C7	4.231	5.117	3.347	3.952		0.894	1.552	0.873	
R1C8	1.095	0.931	1.643	1.152		0.878	1.895	1.310	
R1C9	1.733	1.576	1.259	1.803		1.087	0.650	1.018	
R2C1	2.963	2.365	4.159	3.778		0.623	1.449	1.013	
R2C2	6.031	5.954	5.164	5.442		0.522	1.749	1.639	
R2C3	2.006	3.143	1.702	2.005		1.169	0.706	1.270	
R2C4	1.160	0.996	0.831	0.859		0.176	0.381	0.334	
R2C5	6.755	7.276	5.762	7.423		0.737	2.255	1.574	
R2C6	2.414	2.597	2.324	3.151		0.506	0.199	0.951	
R2C7	1.384	1.261	0.821	1.078		0.505	1.021	0.742	
R2C8	2.572	2.196	2.574	2.040		0.476	0.633	0.572	
R2C9	0.860	1.478	2.938	1.000		0.661	2.177	0.533	

Table 2. dE2000 values for each camera, as compared to RICD and as compared to Nikon D810

50.04								
R3C1	3.134	2.114	2.247	1.876	1	.180	1.552	1.597
R3C2	0.727	0.855	1.263	1.413	1	.259	0.813	1.149
R3C3	6.443	6.468	4.316	5.421	0	.080	2.620	2.347
R3C4	4.441	4.686	3.810	4.775	0	.329	0.791	0.804
R3C5	1.662	1.876	1.839	2.292	0	.442	1.596	1.751
R3C6	3.639	3.591	2.594	3.266	0	.349	2.191	1.482
R3C7	3.147	4.189	2.583	3.289	1	.090	0.586	0.598
R3C8	2.789	3.063	3.037	2.582	0	.604	1.461	0.379
R3C9	2.395	2.049	1.619	1.322	0	.773	2.606	1.710
R4C1	1.683	1.367	1.996	1.486	1	.016	0.459	0.504
R4C2	4.732	4.210	3.653	4.110	0	.733	1.898	1.337
R4C3	2.526	2.478	1.225	1.534	0	.553	2.434	2.441
R4C4	2.608	1.819	2.174	1.698	0	.943	0.504	1.078
R4C5	1.682	2.026	1.259	1.642	0	.534	0.907	0.750
R4C6	1.095	1.023	1.348	0.825	0	.462	1.798	1.161
R4C7	12.735	10.398	7.951	9.334	2	.635	5.537	3.941
R4C8	1.309	2.283	1.160	1.611	1	.079	0.409	0.496
R4C9	4.053	3.849	2.861	2.828	0	.276	1.714	1.880
R5C1	0.554	1.706	2.306	0.583	1	.586	2.679	1.018
R5C2	0.703	0.632	2.263	1.045	0	.420	2.522	1.233
R5C3	0.771	0.773	1.358	1.195	0	.332	1.744	1.375
R5C4	0.648	1.877	0.786	0.949	1	.353	1.167	1.249
R5C5	0.000	0.000	0.000	0.000	0	.000	0.000	0.000
R5C6	0.416	1.476	0.984	0.712	1	.193	0.647	0.318
	0 990	1 479	1.351	1.509	0	.859	0.554	0.654
R5C7	0.550	11175						
R5C7 R5C8	0.748	0.714	0.800	0.824	0	.204	0.284	0.191
R5C7 R5C8 R5C9	0.748	0.714 0.253	0.800 0.253	0.824 0.253	0	.204 .006	0.284	0.191 0.006

## Conclusions

With a one-time grading operation consisting of a scale in linear space, the IDTs match very closely. In Table 1 and Table 2, it can be seen that the mean DeltaE2000 is ~3 or less when cameras are compared to the RICD, and that value drops to ~1 or less when the cameras are compared to each other.

There is room for improvement in IDTs by working out how the "default" scale factor and white balance scalars for the IDT should produce an image at the expected levels for an image exposed properly according to the camera sensitivity rating, exposure time, and aperture. This requires the cameras to be properly rated and for meters to give exposure settings consistently, but should be attainable, or at least able to be improved from the current state.

However, if a one-time grade is derived for a gray card to force 18% gray to [0.18 0.18 0.18] in ACES space, then after that, all colors fall into place quite nicely and produce a match between cameras.

There are further options allowed by the procedures in Academy P-2013-001 that were not explored in this experiment. These include alternate training data sets, cone response matrices, choices in optimization spaces and regression algorithms, etc. Exploring these alternatives could potentially improve results but were beyond the scope of this initial exploration into the state of current IDT "best practices".