

Efficient illuminant correction in the Local, Linear, Learned (L^3) method

Francois G. Germain^a, Iretiayo A. Akinola^b, Qiyuan Tian^b, Steven Lansel^c, and Brian A. Wandell^{b,d}

^aCenter for Computer Research in Music and Acoustics,
Stanford University, Stanford, CA-94305, USA;

^bElectrical Engineering Department, Stanford University, Stanford, CA-94305, USA;

^cOlympus Corporation of the Americas, Sunnyvale, CA-94085, USA;

^dPsychology Department, Stanford University, Stanford, CA-94305, USA.

ABSTRACT

To speed the development of novel camera architectures we proposed a method, L^3 (Local, Linear and Learned), that automatically creates an optimized image processing pipeline. The L^3 method assigns each sensor pixel into one of 400 classes, and applies class-dependent local linear transforms that map the sensor data from a pixel and its neighbors into the target output (e.g., CIE XYZ rendered under a D65 illuminant). The transforms are pre-computed from training data and stored in a table used for image rendering. The training data are generated by camera simulation, consisting of sensor responses and CIE XYZ rendered outputs. The sensor and rendering illuminant can be chosen equal (same-illuminant table) or different (cross-illuminant table). In the original implementation, illuminant correction is achieved with cross-illuminant tables, and one table is required for each illuminant. We find, however, that a single same-illuminant table (D65) effectively converts sensor data for many different same-illuminant conditions. Hence, we propose to render the data by applying the same-illuminant D65 table to the sensor data, followed by a linear illuminant correction transform. The mean color reproduction error using the same-illuminant table is on the order of $4\Delta E$ units, which is only slightly larger than the cross-illuminant table error. This approach reduces table storage requirements significantly without substantially degrading color reproduction accuracy.

Keywords: Local Linear Learned (L^3), image processing pipeline, illuminant correction

1. INTRODUCTION

The L^3 method automatically generates an image processing pipeline (sensor correction, demosaicking, illuminant correction, noise reduction) for novel camera architectures.¹⁻⁴ The L^3 method comprises two parts:¹ L^3 processing pipeline and L^3 training module. The L^3 processing pipeline first classifies each sensor pixel into one of many possible classes based on a statistical analysis of the responses at the pixel and its neighborhood. Then, it applies the class-associated local linear transform to map the pixel to the target output space.

The L^3 training module pre-computes and stores the table of class-associated linear transforms that convert sensor data to the target CIE XYZ values. The transformed are learned from the training data using Wiener estimation. The training data consist of sensor responses and desired rendering data (i.e. CIE XYZ values for consumer photography). They are generated through simulation from a camera model that is implemented using the Image Systems Engineering Toolbox (ISET).^{5,6}

2. MOTIVATION: ILLUMINANT CORRECTION

Illuminant correction is an essential part of the image processing pipeline that is necessary to account for the perceptual phenomenon of *color constancy*.⁷ The visual system adapts to changes in the ambient light level and

Send correspondence to Francois G. Germain. E-mail: FG (fgermain@stanford.edu), IA (iakinola@stanford.edu), QT (qytian@stanford.edu), SL (steven.lansel@olympus.com), BW (wandell@stanford.edu)

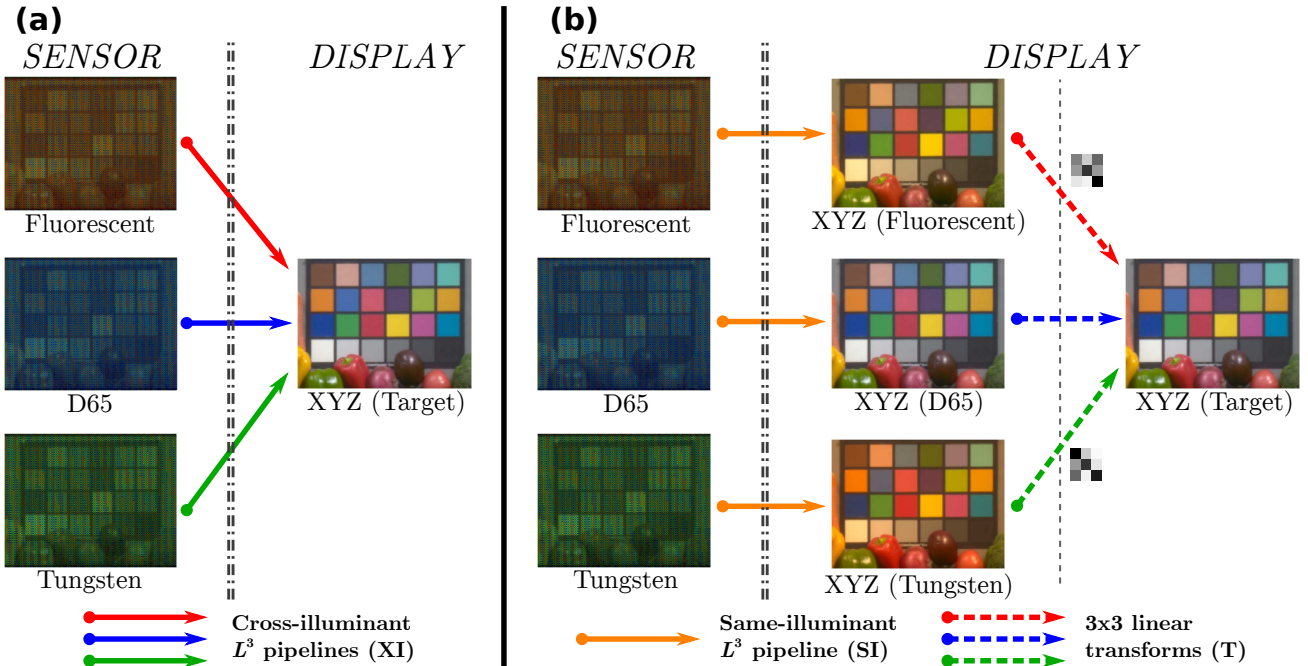


Figure 1. **Illuminant correction with L^3 image processing pipelines.** (a) Cross-illuminant (XI) illuminant-dependent L^3 processing. (b) Same-illuminant (SI) illuminant-independent L^3 processing followed by a linear illuminant correction transform (T).

color, and this adaptation has the general effect of preserving the color appearance of an object (e.g. a white shirt) across conditions (daylight to tungsten light). Because the human visual system adapts to the illuminant, to preserve color appearance the linear transforms in the L^3 method used for rendering must adapt, too.

In the original algorithm, we anticipated learning a table of linear transforms for each illumination condition (cross-illuminant, Figure 1(a)). These cross-illuminant tables would be stored and applied for rendering. This approach pays a significant cost in computation and storage. Here, we explore an approach that separates illuminant correction from the rest of the pipeline. We learn a same-illuminant table, that renders data acquired under one illuminant into XYZ values under the same illuminant. This accomplishes demosaicking, denoising and sensor correction, but it does not perform illuminant correction (Figure 1(b)). We then apply an illuminant correction transform to obtain the final rendering. This architecture requires training and storing only one table of L^3 transforms to be used for all illuminants and one illuminant correction matrix for each illuminant of interest.

3. METHODS

3.1 General simulation conditions

We used ISET^{5,6} to simulate digital camera processing. The ISET camera model includes the simulation of scene radiance, optics, sensor electronics and image processing pipeline (here the L^3 pipeline). We model the camera as an f/4 lens with diffraction limited optics and focal length of 3mm. The sensor uses a RGB/W color filter array (CFA),¹ with parameter values as in Figure 2.

The L^3 algorithm uses simulated training data to learn a table of linear transforms. Each pixel in the simulated sensor data is assigned to a distinct class based on (a) its color type (R, G, B, W), the neighborhood (b) saturation pattern (no saturation, W-saturation, W&G-saturation), (c) response voltage level (20 samples spanning the voltage swing), (d) spatial contrast (uniform or textured).¹ The transform for a given class is derived by optimizing a Wiener filter between the data in the pixel neighborhood in the sensor, and the corresponding ideal value in the target color space (XYZ).

Optics(diffraction-limited)	
Focal length(mm)	3
F-number	4
Aperture diameter (mm)	0.75
Sensor	
Rows/Columns	204 × 240
Size (mm)	0.4488 × 0.528
Dark noise nonuniformity (mV)	1.41
Photoreceptor nonuniformity (%)	2.218×10^{-3}
Analog gain	1
Analog offset (V)	0
Sensor pixel	
Width/height (μm)	2.2×2.2
Fill factor	0.45
Dark voltage (V/sec)	1×10^{-5}
Read noise (mV)	1.34
Conversion gain ($\mu\text{V}/e$)	200
Voltage swing (V)	1.8
Well capacity (electrons)	9000




Figure 2. **Optical and sensor parameter values (left) and Natural-100 color test chart (right)**

The rendering process identifies the class membership for each pixel in the sensor data and then applies the appropriate stored linear transform. The image processing pipeline, which performs demosaicking, denoising, sensor conversion comprises the application of these precomputed local, linear transforms.

3.2 L^3 tables

Learning the cross-illuminant tables. Two cross-illuminant (XI) tables of L^3 transforms are learned by simulating the sensor data under tungsten (XI_{Tun}) and fluorescent (XI_{Flu}) illumination and rendering to a target representation (XYZ) under a D65 illuminant.

Learning the same-illuminant tables. Same-illuminant (SI) tables of L^3 transforms are learned by simulating sensor data and the target display XYZ values using the same illuminant. We learned same-illuminant tables for D65 (SI_{D65}), tungsten (SI_{Tun}) and fluorescent (SI_{Flu}). To complete the rendering, we apply an illuminant correction transform from the acquisition illuminant to the final D65 rendering. The matrix T is derived by a linear regression between the XYZ values in the acquired and target (D65) representations. T is derived from first principles and is not dependent on the camera structure or L^3 .

In principle, any of the SI tables can be used. In fact, it is possible to learn a same-illuminant table using sensor data from an ensemble of illuminants (e.g., a mixture of D65 and tungsten scenes). The linear transforms then reflect the properties of the multiple illuminants used in the ensemble, rather than the requirements of a specific illuminant. For example, the response level of a blue pixel under tungsten illuminant is typically quite low. This is reflected in the images obtained from the tungsten-based same-illuminant table as the L^3 processing compensates by using linear transforms with more blurring on those pixels than the D65-based same-illuminant table. We learned tables for an ensemble comprising data from a mixture of D65 and Tungsten illuminants (SI_{Ens}) which results in a more robust table for both illuminants.

3.3 Color evaluation

To compare the XI and SI methods, we used data from hyperspectral scenes and simulated test charts.

Multispectral scenes. We used a small set of multispectral scenes including the Macbeth (Gretag) color chart, fruits, and people.⁸ These images compare the algorithms visually.




Figure 3. **Cross-illuminant and same-illuminant algorithms perform equally well in color rendering for a tungsten illuminant.** (a) The sensor image under tungsten, uncorrected. (b) The ideal image under a D65 illumination. (c) A cross-illuminant analysis for the same image but acquired under a tungsten light. The L^3 table was built from the tungsten-acquisition data and XYZ target values for a D65-rendered image. (d) A same-illuminant rendering. In this case the L^3 table was built using the same-illuminant D65 table, and the final rendering combined this table with an illuminant correction matrix T. The cross-illuminant and same-illuminant architectures produce similar images.

Natural-100 reflectance chart. To generate a systematic and quantitative assessment of the color reproduction error, we created a Natural-100 test chart (Figure 2). The chart includes 100 samples with natural object reflectances from five categories,⁹ and a 10-sample gray color strip. The reflectances are grouped by category (20 food, 20 clothes, 20 nature, 5 hair, 15 skin and 20 paint color samples). To select the category samples we calculated 7 principal components for the entire set of natural reflectances, represented each sample by its component weights, and then formed clusters of the samples. If N samples are needed from a category, we formed N clusters for that category and chose one sample from each cluster. By selecting one sample per cluster per category we ensure that the final chart has a variety of reflectances that span the sample pool.

We quantified color reproduction error using CIE2000 ΔE units.¹⁰ The error was the difference between the chart XYZ under a D65 illuminant and the data rendered using either the SI or XI method.

4. EXPERIMENTAL RESULTS

We compared the quality of the color rendering using two natural hyperspectral scenes (Figures 3 and 4). Each column shows a different representation of the scene. The leftmost columns show the scene as measured under a tungsten illuminant (Figure 3(a)) and a fluorescent illuminant (Figure 4(a)). The scene rendered under a D65 illuminant, calculated from the multispectral data, is shown in panel (b). The original data rendered using the XI tables is shown in panel (c), and rendered using the SI_{D65} table and linear transformation T ($SI_{D65} - T$) is shown in (d). The XI and $SI_{D65} - T$ renderings (panels c and d) are good approximations to the calculated D65 rendering (panel b).

We quantified the color rendering quality accuracy using the Natural-100 test chart (Table 1). The XI color reproduction errors are very similar to the SI - T color errors. Hence, the color accuracy of the efficient single-illuminant approach matches the cross-illuminant table for the tungsten illuminant.

Interestingly, we found that the SI_{D65} table trained with tungsten illumination can replace the SI_{Tun} table with only a small reduction in color reproduction accuracy ($1.5\Delta E$). Using a SI table derived from training on an ensemble (tungsten and D65) of illuminants reduces the color accuracy by even less ($0.4\Delta E$). As a result, we believe that training a SI table from data with an ensemble of illuminants may provide the best SI approach.

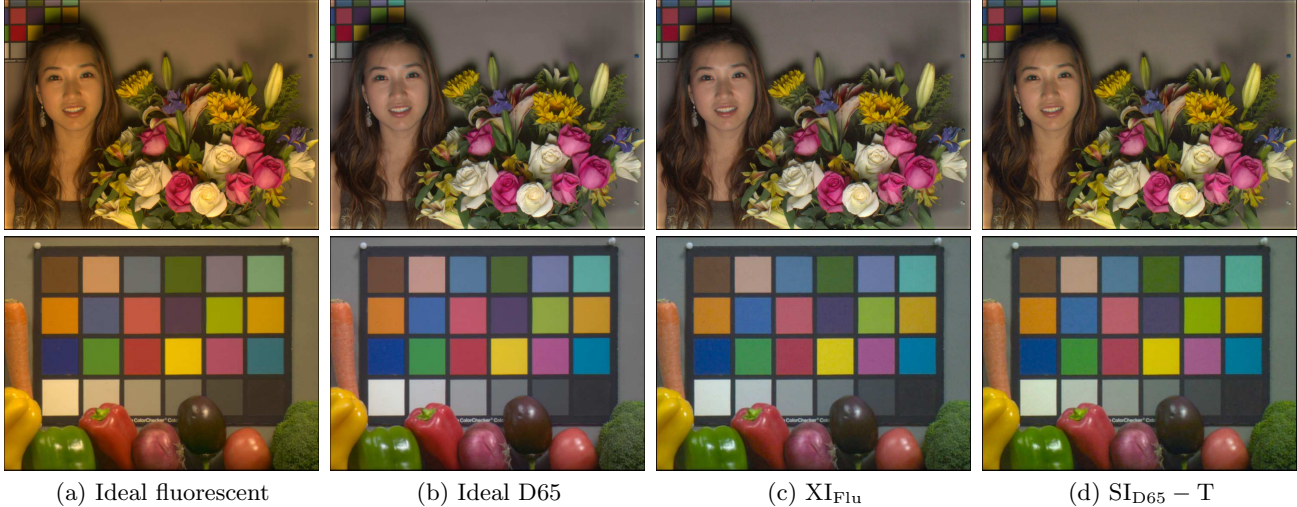


Figure 4. **Cross-illuminant and same-illuminant algorithms perform equally well in color rendering for a fluorescent illuminant.** (a) The sensor image under fluorescent, uncorrected. (b) The ideal image under a D65 illumination. (c) A cross-illuminant analysis for the same image but acquired under a fluorescent light. The L^3 table was built from the fluorescent-acquisition data and XYZ target values for a D65-rendered image. (d) A same-illuminant rendering. In this case the L^3 table was built using the same-illuminant D65 table, and the final rendering combined this table with an illuminant correction matrix T. The cross-illuminant and same-illuminant architectures produce similar images.

A more detailed account of the color reproduction error for the cross-illuminant tables (XI_{Tun}) and three different same-illuminant tables followed by the illuminant transform ($SI_{Tun} - T$, $SI_{D65} - T$ and $SI_{Ens} - T$) is presented in Figure 5 with the distribution of the ΔE values recorded for each reflectance on the Natural-100 chart and the scatter plots comparing the CIE XYZ as well as the $L^*a^*b^*$ color values of the true and reproduced colors.

5. CONCLUSIONS

We compared color reproduction accuracy following illuminant correction for two versions of the L^3 algorithm. In the SI version we store a single, same-illuminant table and a set of 3×3 illuminant correction transformations. In the XI version, we store a set of tables, one for each illumination. The SI method requires far less storage than the XI method and costs very little in terms of color reproduction accuracy. We present different options regarding the choice of the same-illuminant table. We suggest that the best option for a same-illuminant table might be derived by training the L^3 algorithm with data simulated using an ensemble of illuminants.

Color reproduction error (ΔE)	Average (Std. dev.)	90th percentile
$XI_{Tun} L^3$	2.5 (1.8)	4.9
$SI_{Tun} L^3$ with linear illuminant transform T ($SI_{Tun} - T$)	2.5 (2.0)	5.4
$SI_{D65} L^3$ with linear illuminant transform T ($SI_{D65} - T$)	4.0 (3.0)	8.3
$SI_{Ens} L^3$ with linear illuminant transform T ($SI_{Ens} - T$)	2.9 (2.2)	6.0

Table 1. **Color reproduction errors (ΔE) for the Natural-100 color chart.** The original images were simulated using a tungsten illuminant. The left column shows the training method used to derive the L^3 table. The mean reproduction error (middle) and the relatively large reproduction errors (right) are shown. The smallest reproduction error is from the cross-illuminant table. The errors for the different same-illuminant tables followed by a linear illuminant correction transform are only slightly larger.

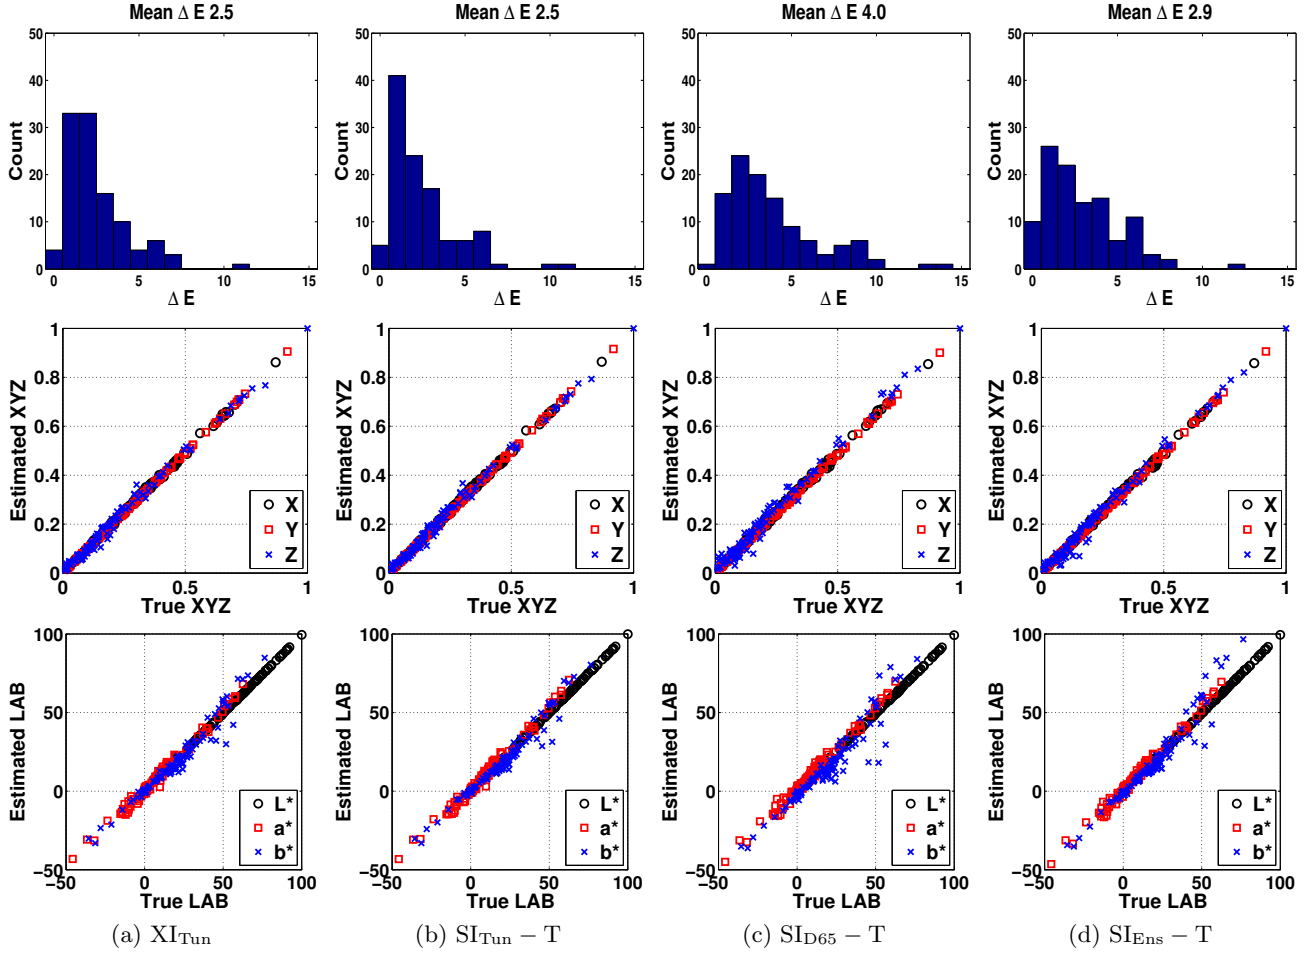


Figure 5. Color error distributions for a Natural-100 chart acquired under a tungsten illuminant and rendered using different L^3 tables. Histograms of the distribution of color reproduction errors in ΔE units (top), scatter plots between true and estimated XYZ coordinates (middle) and scatter plots between true and estimated $L^*a^*b^*$ coordinates (bottom) for the 110 chart reflectances for different L^3 tables: (a) the cross-illuminant table (XI_{Tun}), (b) the same-illuminant table trained under tungsten followed by the illuminant transform T ($SI_{Tun} - T$), (c) the same-illuminant table trained under D65 followed by the illuminant transform T ($SI_{D65} - T$) and (d) the same-illuminant table trained under an ensemble of illuminants followed by the illuminant transform T ($SI_{Ens} - T$).

REFERENCES

- [1] Tian, Q., Lansel, S., Farrell, J. E., and Wandell, B. A., “Automating the design of image processing pipelines for novel color filter arrays: Local, Linear, Learned (L^3) method,” in [*IS&T/SPIE Electronic Imaging*], 90230K–90230K, International Society for Optics and Photonics (2014).
- [2] Lansel, S. P. and Wandell, B. A., “Local linear learned image processing pipeline,” in [*Imaging Systems and Applications*], Optical Society of America (2011).
- [3] Lansel, S. P., *Local Linear Learned Method for Image and Reflectance Estimation*, PhD thesis, Stanford University (2011).
- [4] Lansel, S. P. and Wandell, B. A., “Learning of image processing pipeline for digital imaging devices,” (Dec. 7 2012).
- [5] Farrell, J. E., Xiao, F., Catrysse, P. B., and Wandell, B. A., “A simulation tool for evaluating digital camera image quality,” in [*Electronic Imaging 2004*], 124–131, International Society for Optics and Photonics (2003).

- [6] Farrell, J. E., Catrysse, P. B., and Wandell, B. A., “Digital camera simulation,” *Applied Optics* **51**(4), A80–A90 (2012).
- [7] Wandell, B. A., [*Foundations of Vision*], Sinauer Associates (1995).
- [8] <http://imageval.com/scene-database/>.
- [9] Vrhel, M. J., Gershon, R., and Iwan, L. S., “Measurement and analysis of object reflectance spectra,” *Color Research & Application* **19**(1), 4–9 (1994).
- [10] Sharma, G., Wu, W., and Dalal, E. N., “The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations,” *Color Research & Application* **30**(1), 21–30 (2005).