Simple Analytic Approximations to the CIE XYZ Color Matching Functions

Chris Wyman Peter-Pike Sloan Peter Shirley
NVIDIA

Abstract

We provide three analytical fits to the CIE $\bar{x}$, $\bar{y}$, and $\bar{z}$ color matching curves commonly used in predictive and spectral renderers as an intermediate between light spectra and RGB colors. Any of these fits can replace the standard tabulated CIE curves. Using tabulated curves can introduce typos, encourage crude simplifying approximations, or add opportunities to download curves from sources featuring inconsistent or incorrect data. Our analytic fits are simple to implement and verify. While fitting introduces error, our fits introduce less than the variance between the human-subject data aggregated into the CIE standard. Additionally, common rendering approximations, such as coarse spectral binning, introduce significantly more error. We provide simple, analytic fits in Equations 2 and 3, but even our more accurate fit in Equation 4 only requires ten lines of code.

1. Introduction

We have long understood that human vision is trichromatic. In graphics, the suitability of three-dimensional color spaces and display devices stems from this physiology. To understand this trichromatism, researchers in the 1920s measured human responses to lighting stimuli and empirically derived various spectral response curves spanning human trichromatic vision. The Commission Internationale de l'Éclairage (CIE) standardized these measured curves. The most common standard used in graphics is the CIE XYZ color space, consisting of spectral curves known as the *standard observer* or *color matching functions*; it is preferred over alternate CIE spaces for its nice numerical properties.

The curves $\bar{x}$, $\bar{y}$, and $\bar{z}$ of the CIE standard observer provide a standardized way of converting spectral radiance $L(\lambda)$ to a trichromatic color space:

$$\begin{aligned} X &= \int L(\lambda)\,\bar{x}(\lambda)\,d\lambda, \\ Y &= \int L(\lambda)\,\bar{y}(\lambda)\,d\lambda, \\ Z &= \int L(\lambda)\,\bar{z}(\lambda)\,d\lambda. \end{aligned} \tag{1}$$

Simple transforms can take these XYZ values to other common color spaces, such as RGB, sRGB, Yuv, or Lab.

In graphics, these curves are typically used only in physically-based and predictive renderers. For other applications, the additional accuracy of spectral lighting does not justify the cost of maintaining full spectra throughout all computations.

One part of this extra cost is maintaining a tabulated representation of the color matching functions, which are provided at a 1 nm sampling rate. Such large tabulated data sets limit usage in bandwidth- or memory-limited applications, encourage the use of crude approximations, and provide opportunities to download incorrect data or introduce hard-to-detect typos. We provide fits to the $\bar{x}$, $\bar{y}$, and $\bar{z}$ curves. This allows analytical evaluation at arbitrary wavelengths and reduces programmer cost. While these fits introduce some error, we show that this is less than the measurement error as well as the errors introduced by many commonly used approximations (e.g., coarse spectral binning).

1.1. Background

Our initial goal for this project was a quick analytical fit to the $\bar{x}$, $\bar{y}$, and $\bar{z}$ curves to simplify a new spectral-path tracer. This goal was justified by our collective experience in these curves' robustness; we have used 1 nm, 5 nm, 10 nm, and coarser samplings without apparent variation in image quality.

Prior to our evaluation in Section 3, we explored the color-science literature to identify which aspects of these curves needed preservation in an analytical fit. Surprisingly, we found few answers.

In fact, there are two different CIE-standard $\bar{x}\bar{y}\bar{z}$ functions, the 1931 and 1964 standard observers [Schanda 2007]. The measurements are valid for 2° and 10° fields of view, respectively. As color science often experiments with very small fields of view, the 1964 curves are often ignored, so the 1931 data is more widely available. For graphics, the 10°, 1964 curves may make more sense, as screens typically subtend well over a 2° field of view. We provide fits for both standards.

While apparently not widely known, a few analytic fits to $\bar{x}$, $\bar{y}$, and $\bar{z}$ have been proposed [Moon and Spencer 1943; Tannenbaum 1974]. These use poorly fitting or computationally expensive analytic functions, and all were developed without extensive computing support, which prevented extensive parameter space searches; all our fits exhibit significantly lower error.

Various aspects of the standard curves make them difficult to fit with simple curves. The literature (e.g., [Guild 1932]) is unclear if these are measurement issues, sampling artifacts, or the underlying physiology. Rather than overfit, we ensure our errors are smaller than measured variances in humans.

2. Fitting the XYZ Color Matching Curves

During development, our goals broadened from a quick hack to avoid typing tabulated data to include a more careful fit suitable for predictive renderers. Because of this, we tested numerous fits with varying properties and concluded that two types would be useful to the community: a simple, $\mathbb{C}^\infty$ analytic fit and a more accurate, $\mathbb{C}^1$ piecewise fit.

2.1. Simple, Single-Lobe Fit

To achieve our initial, simple fit, we plugged the CIE color matching functions* into an online curve fitting tool† that uses a Levenberg-Marquardt algorithm [Press et al. 2007] to fit to various functions. This provided us single-lobe analytic functions to fit $\bar{y}$ and $\bar{z}$, and a sum of two lobes to fit $\bar{x}$. Somewhat unsurprisingly, given Houston's [1930] early fitting for photometry, the best fits used Gaussian and log-normal distributions.

For the curves most commonly used in graphics, the CIE 1931 standard observer, the following $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$ approximate $\bar{x}$, $\bar{y}$, and $\bar{z}$:

$$\begin{aligned}
\tilde{x}_{31}(\lambda) &= 1.065 \exp\left(-\tfrac{1}{2}\left(\tfrac{\lambda-595.8}{33.33}\right)^2\right) + 0.366 \exp\left(-\tfrac{1}{2}\left(\tfrac{\lambda-446.8}{19.44}\right)^2\right), \\
\tilde{y}_{31}(\lambda) &= 1.014 \exp\left(-\tfrac{1}{2}\left(\tfrac{\ln\lambda-\ln 556.3}{0.075}\right)^2\right), \\
\tilde{z}_{31}(\lambda) &= 1.839 \exp\left(-\tfrac{1}{2}\left(\tfrac{\ln\lambda-\ln 449.8}{0.051}\right)^2\right).
\end{aligned} \tag{2}$$

These fit the CIE standard curves reasonably well; $\tilde{x}_{31}$ and $\tilde{y}_{31}$ have root mean square errors below 0.015 and maximum absolute errors below 0.046 (roughly 4%). However, the $\bar{z}$ curve is difficult to fit near the peak and base, so $\tilde{z}_{31}$ has almost three times that error.

For the CIE 1964 standard observer, we found the following $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$ fit:

$$\begin{aligned}
\tilde{x}_{64}(\lambda) &= 0.398 \exp\left(-1250\left[\ln\left(\tfrac{\lambda+570.1}{1014}\right)\right]^2\right) \\
&\quad + 1.132 \exp\left(-234\left[\ln\left(\tfrac{1338-\lambda}{743.5}\right)\right]^2\right), \\
\tilde{y}_{64}(\lambda) &= 1.011 \exp\left(-\tfrac{1}{2}\left(\tfrac{\lambda-556.1}{46.14}\right)^2\right), \\
\tilde{z}_{64}(\lambda) &= 2.060 \exp\left(-32\left[\ln\left(\tfrac{\lambda-265.8}{180.4}\right)\right]^2\right).
\end{aligned} \tag{3}$$

*Tabulated CIE XYZ standard observer data was obtained from the Rochester Institute of Technology Munsell Color Science Lab at http://www.cis.rit.edu/mcsl/online/cie.php.

†Curve fitting done using ZunZun, http://www.zunzun.com.

While more complex than our fits to the 1931 data, the fit in Equation 3 is better behaved; $\tilde{x}_{64}$, $\tilde{y}_{64}$, and $\tilde{z}_{64}$ all have root mean square errors below 0.016 and maximum absolute errors below 0.056 (roughly 3%).

2.2. Multi-Lobe, Piecewise Gaussian Fit

The single-lobe Gaussian fits presented above should suffice for many applications (see Section 3), except those where closely matching the CIE standard is vital. Unfortunately, all the curves have asymmetric lobes and $\bar{x}$ and $\bar{z}$ have strangely flat-topped lobes near 450 nm. None of our simple analytical curves fit these characteristics. To reduce the remaining errors, we explored Gaussian mixture models, polynomial-windowed Gaussian lobes, and piecewise continuous Gaussians. The online fitting tool we used did not handle these models, so we wrote code to perform these fits via the Simplex method [Press et al. 2007].

Our best fits use multiple piecewise continuous Gaussians for each curve. With three Gaussians for $\tilde{x}$ and two for $\tilde{y}$ and $\tilde{z}$, we achieved squared error rates below the within-observer variance in the experimental measurements used to form the CIE standards [Nimeroff et al. 1962]. While adding additional Gaussians could reduce numerical error, we would effectively be fitting noise in the data; more error is likely added elsewhere, e.g., during monitor and printer color calibration.

Our piecewise fit to the 1931 standard observer uses the following form:

$$\begin{aligned}
\tilde{x}_{31}(\lambda) &= \sum_{i=0}^{2} \alpha_{x_i} \exp\left(-\tfrac{1}{2}\left[\,(\lambda-\beta_{x_i})\,\mathbf{S}\,(\lambda-\beta_{x_i},\gamma_{x_i},\delta_{x_i})\,\right]^2\right), \\
\tilde{y}_{31}(\lambda) &= \sum_{i=0}^{1} \alpha_{y_i} \exp\left(-\tfrac{1}{2}\left[\,(\lambda-\beta_{y_i})\,\mathbf{S}\,(\lambda-\beta_{y_i},\gamma_{y_i},\delta_{y_i})\,\right]^2\right), \\
\tilde{z}_{31}(\lambda) &= \sum_{i=0}^{1} \alpha_{z_i} \exp\left(-\tfrac{1}{2}\left[\,(\lambda-\beta_{z_i})\,\mathbf{S}\,(\lambda-\beta_{z_i},\gamma_{z_i},\delta_{z_i})\,\right]^2\right),
\end{aligned} \tag{4}$$

where $\mathbf{S}(x,y,z)$ is the selector function represented in C code as (x<0)?y:z, or mathematically using a pair of Heaviside functions, $\mathbf{H}$:

$$\mathbf{S}(x,y,z) = y(1-\mathbf{H}(x)) + z\mathbf{H}(x).$$

	$\tilde{x}_0$	$\tilde{x}_1$	$\tilde{x}_2$	$\tilde{y}_0$	$\tilde{y}_1$	$\tilde{z}_0$	$\tilde{z}_1$
α	0.362	1.056	-0.065	0.821	0.286	1.217	0.681
β	442.0	599.8	501.1	568.8	530.9	437.0	459.0
γ	0.0624	0.0264	0.0490	0.0213	0.0613	0.0845	0.0385
δ	0.0374	0.0323	0.0382	0.0247	0.0322	0.0278	0.0725

Table 1. Fitting coefficients for our multi-lobe Gaussian fit in Equation (4).

```
// Inputs:  Wavelength in nanometers

float xFit_1931( float wave )
{
    float t1 = (wave-442.0f)*((wave<442.0f)?0.0624f:0.0374f);
    float t2 = (wave-599.8f)*((wave<599.8f)?0.0264f:0.0323f);
    float t3 = (wave-501.1f)*((wave<501.1f)?0.0490f:0.0382f);
    return 0.362f*expf(-0.5f*t1*t1) + 1.056f*expf(-0.5f*t2*t2)
                                    - 0.065f*expf(-0.5f*t3*t3);
}

float yFit_1931( float wave )
{
    float t1 = (wave-568.8f)*((wave<568.8f)?0.0213f:0.0247f);
    float t2 = (wave-530.9f)*((wave<530.9f)?0.0613f:0.0322f);
    return 0.821f*exp(-0.5f*t1*t1) + 0.286f*expf(-0.5f*t2*t2);
}

float zFit_1931( float wave )
{
    float t1 = (wave-437.0f)*((wave<437.0f)?0.0845f:0.0278f);
    float t2 = (wave-459.0f)*((wave<459.0f)?0.0385f:0.0725f);
    return 1.217f*exp(-0.5f*t1*t1) + 0.681f*expf(-0.5f*t2*t2);
}
```

Listing 1. Simple C code for our multi-lobe Gaussian fit from Equation 4.

The fitting parameters for α, β, γ, and δ are provided in Table 1, and simple C code to implement these functions is provided in Listing 1.

3. Evaluation and Discussion

This paper provides simple, easy-to-use fits as an alternative to large tabular forms of the standard color matching curves. To convince error-sensitive users of the validity of these substitutions, we evaluate these fits by asking these questions: What errors do our fits introduce relative to common simplifications and the standard curves? How do these fits compare to prior analytic fits? How do these fits affect output RGB images?

Common Simplifications. We explore two simplifications used in some spectral renderers, to compare their error with our analytic fits. In particular, we compare with coarse binning and linearly interpolating sparse $\bar{x}\bar{y}\bar{z}$ samples. Unless a path tracer stochastically assigns a wavelength to each photon, maintaining 1 nm sampled spectra along the path increases the computation at reflections in addition to increasing the storage needed for spectral measured BRDFs. Most renderers, thus, use some bin size larger than 1 nm; we compare against a 10-bin spectral representation (30 nm-wide

	Max Squared Error			Mean Squared Error		
	$\tilde{x}$	$\tilde{y}$	$\tilde{z}$	$\tilde{x}$	$\tilde{y}$	$\tilde{z}$
Moon and Spencer [1945]	9.5e-3	3.9e-3	5.9e-2	9.0e-4	4.1e-4	2.5e-3
1931 CIE data, 10 bins	7.1e-2	2.5e-1	5.8e-1	5.6e-3	4.1e-3	2.0e-2
1931 CIE data; 10 nm, interpolated	1.7e-4	5.2e-5	3.2e-3	1.4e-5	4.2e-6	1.3e-4
1931 Single-lobe fit (Sec. 2.1)	1.3e-3	2.1e-3	2.5e-2	2.2e-4	2.2e-4	1.6e-3
1931 Multi-lobe fit (Sec. 2.2)	2.0e-4	6.4e-5	4.9e-4	3.1e-5	7.1e-6	1.6e-5
Measured error, between subjects		N/A		1.2e-3	3.8e-4	3.7e-3
Measured error, within subject		N/A		3.8e-5	1.2e-5	1.1e-4
1964 Single-lobe fit (Sec. 2.1)	2.1e-3	7.2e-4	3.0e-3	1.9e-4	1.2e-4	2.4e-4
Tannenbaum [1974]	1.8e-1	7.7e-2	7.7e-2	3.0e-2	8.4e-3	5.0e-3

Table 2. Maximum and mean squared error for fits of the CIE standard observers, relative to the corresponding 1 nm sampled curves. Experimental errors come from Nimeroff et al. [1962]. Moon and Spencer [1945] fit the 1931 standard observer, Tannenbaum [1974] the 1964 observer.

bins). While less common, using a subset of the $\bar{x}\bar{y}\bar{z}$ tables and linearly interpolating reduces the necessary memory footprint.

Numerical Errors. Table 2 compares the errors of our fits, previously published analytic fits by Moon and Spencer [1945] and Tannenbaum [1974], and the common simplifications discussed above. Errors are relative to 1 nm sampled CIE standard observer curves. We also include in Table 2 the errors from the experimental measurements used to create the 1964 standard observer [Nimeroff et al. 1962]. Similar numbers apparently do not exist for the earlier standard, though Moon and Spencer [1943; 1945] suggest the errors may have been larger.

These error metrics show even our simple, single-lobe fits are closer to the CIE standard curves than prior analytic models; they are also simpler to compute. Our simple fit introduces an order of magnitude less error than using coarse binning, a common simplification in many spectral renderers. Additionally, our simple fits exhibit less error than the experimental variation between different human subjects.

For users desiring more accuracy, our multi-lobe Gaussian fit in Section 2.2 introduces less error than the experimental variations between repeated measurements on a single subject. The error is similar to that introduced by linearly interpolating between 10 nm samples on the curves. Given that the CIE curves themselves were interpolated to 1 nm [Judd 1931], closer fits may end up overfitting to noise, sampling artifacts, and prior interpolation schemes.

Rendering Comparisons. Figure 1 shows a worst-case image for visualizing differences in the $\bar{x}\bar{y}\bar{z}$ color matching curves, a rainbow where the spectrum in each pixel is non-zero at one wavelength. Despite significant variations in the curves, most are indistinguishable from the standard CIE observers (Figure 1(f,g)). Only the Tannen-

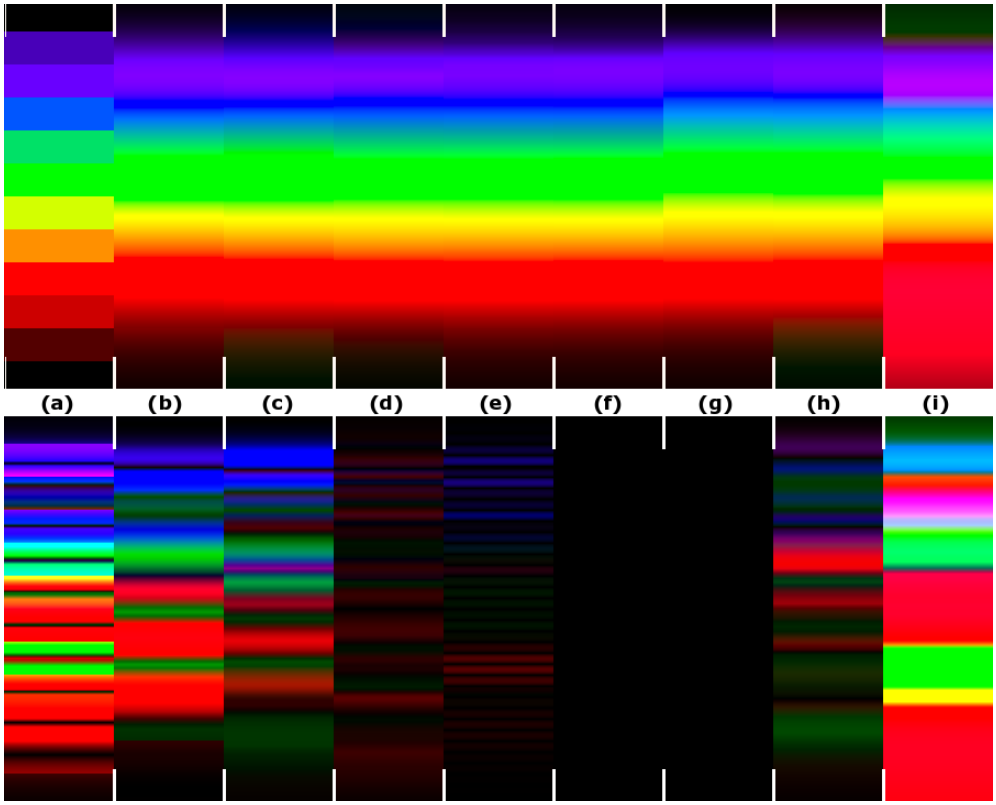

Figure 1. A rainbow under various $\bar{x}\bar{y}\bar{z}$ curves. The spectrum in each pixel is a Dirac function, varying from 375 nm to 725 nm. Top: The result using each curve; Bottom: 8x differences with the 1931 or 1964 CIE standards. The curves come from (a) ten 30 nm-bins of 1931 CIE data, (b) Moon and Spencer [1945], (c) our single-lobe fit ($\tilde{x}_{31}\tilde{y}_{31}\tilde{z}_{31}$) from Section 2.1, (d) our multi-lobe fit from Section 2.2, (e) 1931 CIE data sampled every 10 nm and linearly interpolated, (f) 1931 CIE standard observer, (g) 1964 CIE standard observer (h) our single-lobe fit ($\tilde{x}_{64}\tilde{y}_{64}\tilde{z}_{64}$) from Section 2.1, and (i) Tannenbaum [1974].

baum [1974] fit and binning approach give drastically different results. The Moon and Spencer [1945] and our single-lobe fit closely approximate the standard curve, but slightly blur the color boundaries. Our multi-lobe fit best matches the 1931 CIE standard, but does have a slight variation apparent in the purple band.

Figures 2 and 3 shows the MacBeth color checker using the $\bar{x}\bar{y}\bar{z}$ curves discussed in the paper. We also plot the curves and their absolute deviations from the standard observers. As in Figure 1, our multi-lobe Gaussian fits closely to the 1931 standard observer, with slightly higher error than interpolating a 10 nm sampled standard observer. Either our single-lobe fit or the Moon and Spencer [1945] fit should work for most renderers, though the Moon and Spencer fit is fairly expensive to compute and has precision issues when implemented poorly.

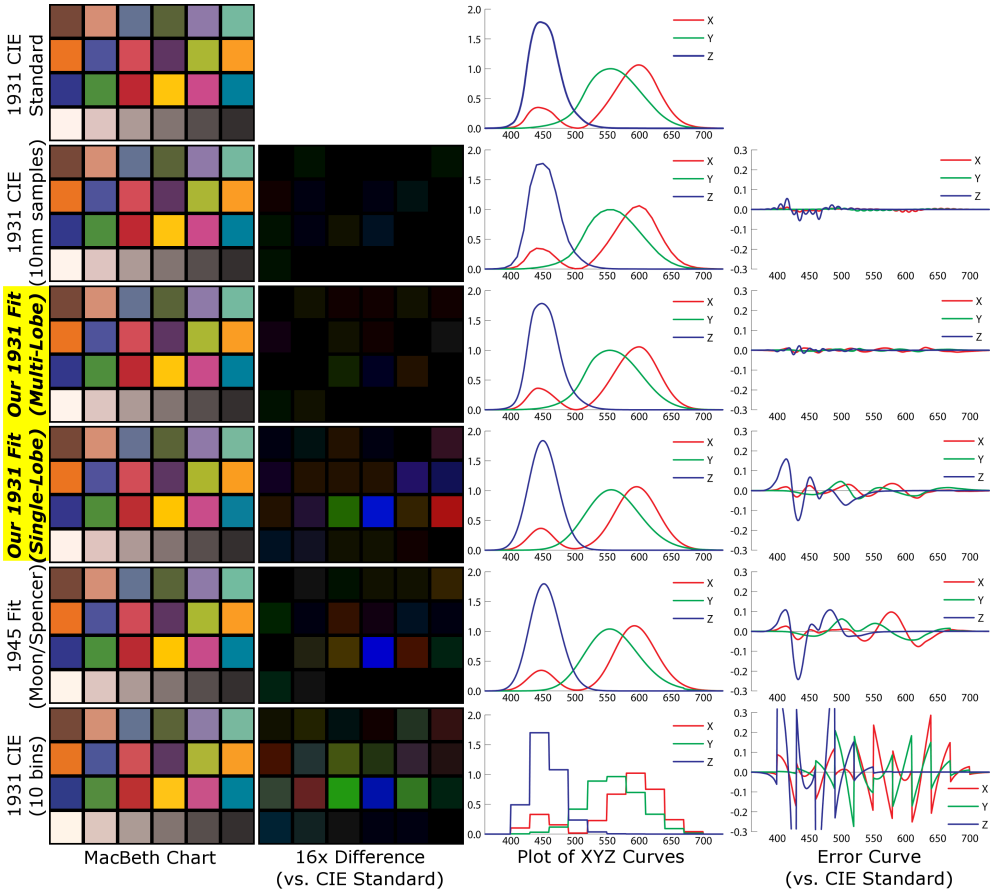

Figure 2. Comparing various $\bar{x}\bar{y}\bar{z}$ fits to the 1931 CIE standard observer. Top to bottom: 1931 CIE standard observer, CIE standard sampled at 10 nm and linearly interpolated, our multi-lobe fit from Section 2.2, our single-lobe fit from Section 2.1, the Moon and Spencer [1945] fit, and a coarse 10-bin fit of the 1931 CIE data,

Our single-lobe approach closely fits the 1964 CIE standard observer. Given the closeness of this fit and the reliance on the 1931 standard in the graphics community, we decided a better multi-lobe fit was unnecessary. Clearly, renderers should avoid using the Tannenbaum [1974] fit, as it significantly overshoots in the red region of the spectrum and undershoots in green and blue.

While our results use sRGB, rather than spectral or XYZ color space, they still provide insight into expected errors (after all most renderings get displayed as sRGB or CMYK). The MacBeth color checker is relatively unsaturated, so few colors in the accurate fits lie outside of gamut. In the rainbow comparison, the interesting errors occur at saturated color boundaries. These differences show up clearly in our difference images.

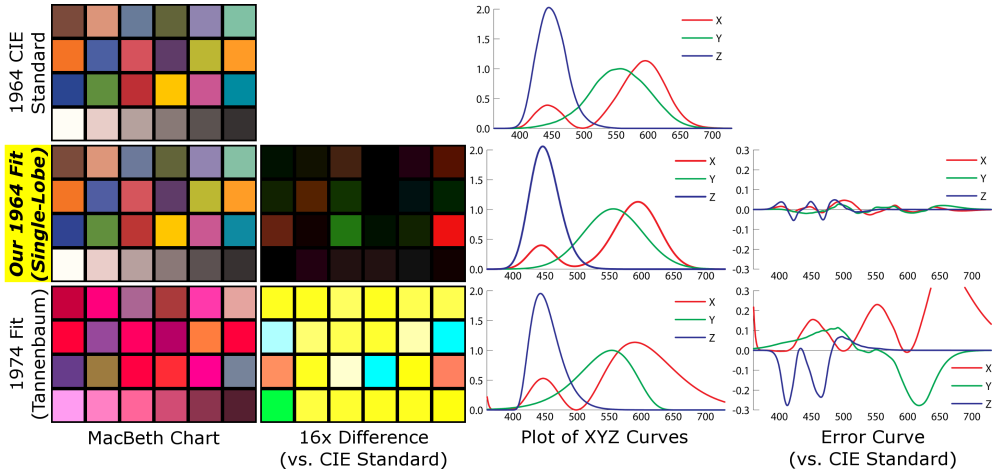

Figure 3. Comparing various $\bar{x}\bar{y}\bar{z}$ fits to the 1964 CIE standard observer. Top to bottom: 1964 CIE standard observer, our single-lobe fit, and Tannenbaum's [1974] fit.

Performance Considerations. By using analytic fits rather than tabulated samples, our approach avoids polluting memory caches with tabulated $\bar{x}\bar{y}\bar{z}$ samples. This is especially important in GPUs, DSPs, and other stream processors where thousands of pixels may be processed in parallel. To demonstrate this effect, we timed each $\bar{x}\bar{y}\bar{z}$ representation from Figure 1 with both CPU and GPU implementations (see Table 3). These simulations show best-case results for tabulated approaches, since we timed a tight loop over hundreds of thousands of $\bar{x}\bar{y}\bar{z}$ evaluations. No other operations accessed memory, so the tabulated $\bar{x}\bar{y}\bar{z}$ curves resided in CPU cache throughout our test.

	1931 CIE data						1964 CIE data		
	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)
CPU evaluation (msec / 100k)	2.4	31.1	20.8	17.5	6.8	5.7	5.3	19.0	18.0
GPU evaluation (μsec / 100k)	3.8	0.5	0.8	0.9	11.7	13.9	14.3	0.4	0.9

Table 3. Cost to evaluate $\bar{x}\bar{y}\bar{z}$ at 100,000 random wavelengths. Evaluations performed back-to-back, so tabulated values still reside in cache, especially for large CPU caches. (a) Ten 30 nm bins, (b) Moon [1945], (c) our single-lobe fit, (d) our multi-lobe fit, (e) 10 nm 1931 CIE data, (f) 1 nm 1931 CIE data, (g) 1 nm 1964 CIE data, (h) our single-lobe fit, and (i) Tannenbaum [1974].

4. Summary

We introduced three new fits of the CIE $\bar{x}\bar{y}\bar{z}$ color matching curves in Equations (2), (3), and (4). We believe any are suitable as a replacement for the standard $\bar{x}\bar{y}\bar{z}$ curves for rendering applications, and we showed they all exhibit less error than the experi-

mental measurements used to generate the standardized curves. For users desiring the most accurate representation possible, our multi-lobe fit in Equation (4) provides a fit comparable to interpolating 10 nm data samples yet requires just ten lines of code.

Acknowledgments

We would like to thank Dave Luebke, others at NVIDIA, and the anonymous reviewers for their comments, suggestions, and other feedback.

References

GUILD, J. 1932. The colorimetric properties of the spectrum. *Philisophical Transactions of the Royal Society of London A 230*, 149–187. 2

HOUSTON, R. A. 1930. Visibility of radiant energy equation. *Philisophical Magazine 10*, 416–432. 3

JUDD, D. 1931. Extension of the standard visibility function to intervals of 1 millimicron by third-difference osculatory interpolation. *Journal of the Optical Society of America 21*, 5, 267–275. 6

MOON, P., AND SPENCER, D. E. 1943. Analytic representation of standard response curves. *Journal of the Optical Society of America 33*, 2 (Feb), 89–103. 2, 6

MOON, P., AND SPENCER, D. E. 1945. Analytic representation of trichromatic data. *Journal of the Optical Society of America 35*, 6 (June), 399–427. 6, 7, 8, 9

NIMEROFF, I., ROSENBLATT, J., AND DANNEMILLER, M. 1962. Variability of spectral tristimulus values. *Journal of the Optical Society of America 52*, 6 (June), 685–691. 4, 6

PRESS, W., TEUKOLSKY, S., VETTERLING, W., AND FLANNERY, B. 2007. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, Cambridge, UK. 3, 4

SCHANDA, J. 2007. *Colorimetry: Understanding the CIE System*. Wiley, New York. 2

TANNENBAUM, P. 1974. Analytic approximations for the CIE 1964 ten-degree-field color-matching functions. *Journal of the Optical Society of America 64*, 1 (Jan), 89–91. 2, 6, 7, 8, 9

Index of Supplemental Materials

We provide supplemental material including C++ code for each of the $\bar{x}\bar{y}\bar{z}$ functions discussed in the paper, and a simple curve viewer to interactively compare the curves. The code is organized as follows:

curves/ - C++ implementation of XYZ curves used in the paper

data/ - headers including tabulated data for various CIE standard curves

iglu/ - headers and library dependencies for curve viewer

Release/ - precompiled curve viewer binaries for Windows 32

main.cpp - curve viewer source

Author Contact Information

Chris Wyman (chris.wyman@acm.org)
Peter-Pike Sloan (ppsloan@nvidia.com)
Peter Shirley (pshirley@nvidia.com)
NVIDIA
2150 South 1300 East, Suite 500
Salt Lake City, UT 84106-4315

Chris Wyman, Peter-Pike Sloan, and Peter Shirley, Simple Analytic Approximations to the CIE XYZ Color Matching Functions, *Journal of Computer Graphics Techniques (JCGT)*, vol. 2, no. 2, 1–11, 2013
http://jcgt.org/published/0002/02/01/

Received:	2012-11-01		
Recommended:	2013-01-14	Corresponding Editor:	Larry Gritz
Published:	2013-07-12	Editor-in-Chief:	Morgan McGuire