

Colour

2015-04-04

Creative Commons Attribution-ShareAlike 4.0 (CC BY-SA 4.0)

<http://creativecommons.org/licenses/by-sa/4.0/>

This is a human-readable summary of (and not a substitute for) the license.

You are free to:

- Share – copy and redistribute the material in any medium or format.
- Adapt – remix, transform, and build upon the material for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:

- Attribution – You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
- ShareAlike – If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

No additional restrictions – You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Almost all content is taken from Wikipedia with some light editing. Figures and plates were made with `gnuplot` and C programs using equations and data sets found online. The document was typeset in Latin Modern Sans using \LaTeX .

[15] <http://en.wikipedia.org/wiki/SRGB>.

[16] http://en.wikipedia.org/wiki/Portable_Network_Graphics.

[17] <http://www.libpng.org/pub/png/libpng.html>.

[18] http://en.wikipedia.org/wiki/ICC_profile.

[19] International Color Consortium.
http://www.color.org/XYZprofiles.xalterfileD50_XYZ.icc.

Further Reading

“Completely Painless Programmer’s Guide to XYZ, RGB, ICC, xyY, and TRCs” – Elle Stone

<http://ninedegreesbelow.com/photography/xyz-rgb.html>

References

- [1] <http://en.wikipedia.org/wiki/Light>.
- [2] http://en.wikipedia.org/wiki/Black-body_radiation.
- [3] http://en.wikipedia.org/wiki/Planck%27s_law.
- [4] http://en.wikipedia.org/wiki/Boltzmann_constant.
- [5] http://en.wikipedia.org/wiki/Planck_constant.
- [6] http://en.wikipedia.org/wiki/Speed_of_light.
- [7] http://en.wikipedia.org/wiki/Standard_illuminant.
- [8] International Commission on Illumination.
<http://cie.co.at/>.
- [9] <http://files.cie.co.at/204.xls>, sheet "daylight comp".
- [10] http://en.wikipedia.org/wiki/Cone_cell.
- [11] http://en.wikipedia.org/wiki/Color_vision.
- [12] Colour & Vision Research Laboratory
University College London.
<http://cvrl.ioo.ucl.ac.uk/>
section cone fundamentals, data set linss2_10e_5.csv,
Stockman and Sharpe (2000) 2-deg cone fundamentals
based on the Stiles and Burch 10-deg CMFs (adjusted to
2-deg).
- [13] http://en.wikipedia.org/wiki/CIE_1931_color_space.
- [14] <http://files.cie.co.at/204.xls>, sheet "1931 observer".

Contents

0	Introduction	4
1	Light	5
2	Daylight	6
3	Eyesight	9
4	Observer	10
5	Chromaticity	13
6	sRGB	14
7	PNG	16
8	Profiles	17
9	Example	17

0 Introduction

Colour is complicated. From electromagnetic radiation to conscious human perception – much physics, many biology, even psychology. Technology to reproduce colour is limited by physics because it's hard to generate arbitrary electromagnetic spectral power distributions, and so metameric tricks exploit biology to give the illusion thereof. Some illusions are more convincing than others, but mathematics can transform between them.

This booklet came to be while I was writing a ray-tracer (a program that simulates optical physics to generate images of virtual objects). Some optical behaviours like refraction (and to a lesser extent reflection) vary with the wavelength of light. Simulating any given wavelength of light gives an intensity image, but our human vision has chromaticity as well as brightness. So I needed a way to combine many different intensity images, each corresponding to a different wavelength of light, into a single colour image.

The CIE XYZ colour matching functions (also known as the standard observer) allow to reduce an intensity spectrum into a 3-channel image. This can then be converted to other colour spaces like sRGB, either “offline”, or via embedded ICC profiles in PNG image files and other interchange formats. Using embedded colour profiles allows colours outside the gamut of sRGB to be preserved until the final output device (whether that be an sRGB monitor, or a multi-ink printer, or laser beams pointed at the moon).

Listing 3: profile.s

```
.section .rodata
.global profile
profile:
.incbin "D50_XYZ.icc"
```

Listing 4: observer.h

```
const float observer_data[81][3] =
    { { 0.001368f, 0.000039f, 0.006450f }
      . . . . . // CIE 1931 standard observer
      , { 0.000042f, 0.000015f, 0.000000f }
5     };

void observer // linear interpolation of data
    ( float x, float *xx, float *yy, float *zz ) {
    float y = fminf(fmaxf(80.0f * x, 0.0f), 80.0f);
10    int i0 = floorf(y);
    int i1 = i0 + 1;
    float t1 = y - i0;
    float t0 = 1.0f - t1;
    *xx = observer_data[i0][0] * t0
15    + observer_data[i1][0] * t1;
    *yy = observer_data[i0][1] * t0
    + observer_data[i1][1] * t1;
    *zz = observer_data[i0][2] * t0
    + observer_data[i1][2] * t1;
20 };
```

Listing 5: Makefile

```
spectrum: spectrum.c profile.h profile.s \
    D50_XYZ.icc observer.h
5 gcc -std=c99 -Wall -pedantic -Wextra \
    -O3 -fopenmp -march=native \
    -o spectrum spectrum.c profile.s \
    'pkg-config --cflags --libs libpng'
```

```

    } else {
        xyz[j][3*i+0] = quantize(gx); // just grey
        xyz[j][3*i+1] = quantize(gy);
80     xyz[j][3*i+2] = quantize(gz);
    }
}
}
}
85
int main() {
    const float dpi = 600.0f;
    const int width = 2700, height = 300;
    int bytes // allocate
90     = height * sizeof(uint16_t *)
    + height * sizeof(uint16_t) * width * 3;
    uint16_t **xyz = calloc(bytes, 1);
    if (! xyz) { abort(); }
    for (int j = 0; j < height; ++j) {
95     xyz[j] =
        ((uint16_t *) (xyz+height)) + j*width*3;
    }
    render_image(xyz, width, height); // render
    FILE *out = fopen("spectrum.png", "wb"); // save
100    if (! out) { abort(); }
    bool ok = write_png
        ( out, (png_bytepp) xyz, width, height, dpi );
    if (! ok) { abort(); }
    fclose(out);
105    return 0; // success
}

```

Listing 2: profile.h

```

const char *profile_name = "D50_ XYZ";
const float profile_white[3] = // white point
    { 0.9642f, 1.0000f, 0.8249f };
const uint32_t profile_size = 528; // in bytes
5 extern const uint8_t profile[528]; // "profile.s"

```

1 Light

Light [1] is electromagnetic radiation within a certain portion of the electromagnetic spectrum. Visible light is usually defined as having a wavelength in the range of 400nm to 700nm – between the ultraviolet (with shorter wavelengths) and the infrared (with longer wavelengths).

The most common light sources are thermal: a body at a given temperature emits a characteristic spectrum of black-body radiation [2]. A simple thermal source is sunlight, the radiation emitted by the chromosphere of the Sun at around 6000 Kelvin peaks in the visible region of the electromagnetic spectrum.

Planck's law [3] describes the electromagnetic radiation emitted by a black body in thermal equilibrium at a definite temperature. Max Planck showed that the spectral radiance at wavelength λ of a body at absolute temperature T is given by:

$$B_{\lambda}(\lambda, T) = \frac{2hc^2}{\lambda^5 \left(\exp \frac{hc}{k_B T \lambda} - 1 \right)}$$


where $k_B = 1.3806488(13) \times 10^{-23} \text{ JK}^{-1}$ is the Boltzmann constant [4], $h = 6.62606957(29) \times 10^{-34} \text{ Js}$ is the Planck constant [5], and $c = 299792458 \text{ ms}^{-1}$ is the speed of light [6]. The distribution B_{λ} peaks at energy:

$$E = \left(5 + W \left(\frac{-5}{\exp 5} \right) \right) k_B T \approx 4.965 k_B T$$

where W is the Lambert W function. The wavelength of the peak is $\lambda_{\max} = \frac{hc}{E}$, resulting in the relation:

$$\lambda_{\max} T \approx 2.898 \times 10^{-3}$$

Figure 1: Planck's law for black body radiation.



2 Daylight

A standard illuminant [7] is a theoretical source of visible light with a spectral power distribution which is published. Standard illuminants provide a basis for comparing images or colours recorded under different lighting. The International Commission on Illumination [8] (abbreviated CIE) publishes all of the well-known standard illuminants.

Illuminants A, B, and C were introduced in 1931, with the intention of respectively representing average incandescent light, direct sunlight, and average daylight. Illuminants D represent phases of daylight, Illuminant E is the equal-energy illuminant,

```

png_set_IHDR                                     // image settings
  ( png, info, width, height, 16
40   , PNG_COLOR_TYPE_RGB, PNG_INTERLACE_ADAM7
    , PNG_COMPRESSION_TYPE_DEFAULT
    , PNG_FILTER_TYPE_DEFAULT );
uint32_t dpm = roundf(dpi * 1000.0f / 25.4f);
png_set_pHYs                                     // dots per meter
45   ( png, info, dpm, dpm, PNG_RESOLUTION_METER );
png_set_iCCP                                     // ICC profile
  ( png, info, (png_charp) profile_name, 0
    , (png_charp) profile, profile_size);
png_write_info(png, info);
50 png_write_image(png, rows);                   // pixel data
png_write_end(png, info);
png_destroy_write_struct(&png,&info); // cleanup
return true;
}
55
uint16_t quantize(float c) {
  return htons // PNG uses network byte order
              (fminf(fmaxf(65535.0f * c, 0.0f), 65535.0f));
}
60
void render_image
  ( uint16_t **xyz, int width, int height ) {
  float gx = 0.5f * profile_white[0]; // grey
  float gy = 0.5f * profile_white[1];
65  float gz = 0.5f * profile_white[2];
  #pragma omp parallel for schedule(static)
  for (int i = 0; i < width; ++i) {
    float x = (i + 0.5f) / width;
    float xx, yy, zz;
70    observer(x, &xx, &yy, &zz);
    for (int j = 0; j < height; ++j) {
      if (height <= 4 * j && 4 * j < 3 * height) {
        // darken and mix with grey
75        xyz[j][3*i+0] = quantize(xx * 0.25f + gx);
        xyz[j][3*i+1] = quantize(yy * 0.25f + gy);
        xyz[j][3*i+2] = quantize(zz * 0.25f + gz);
      }
    }
  }
}

```

Listing 1: spectrum.c

```


#include <math.h>
#include <stdbool.h>
#include <stdint.h>
#include <stdio.h>
5 #include <stdlib.h>
#include <arpa/inet.h>
#include <png.h>
#include "profile.h"
#include "observer.h"

10 void write_png_err // longjmp for error handling
    ( png_structp png, png_const_charp msg ) {
    fprintf(stderr, "error: \u005Cpng: \u005Cs\n", msg);
    jmp_buf *jmp = (jmp_buf *)
15     png_get_error_ptr(png);
    if (jmp) { longjmp(*jmp,1); } else { abort(); }
}

bool write_png // save a PNG image with XYZ data
20 ( FILE *out, png_bytepp rows
    , int width, int height, float dpi) {
    jmp_buf jmp; // initialize
    png_structp png = png_create_write_struct
        (PNG_LIBPNG_VER_STRING,&jmp,write_png_err,0);
25     if (! png) { return false; }
    png_info info = png_create_info_struct(png);
    if (! info) {
        png_destroy_write_struct(&png,0);
        return false;
30     }
    if (setjmp(jmp)) {
        png_destroy_write_struct(&png,&info);
        return false;
    }
35     png_init_io(png, out); // start writing
    png_set_compression_level
        ( png, Z_BEST_COMPRESSION );

```

Figure 2: Illuminants D, normalized to black-body levels.



while Illuminants F represent various fluorescent lamps.


The spectral power distribution S of Illuminants D can be approximated by using the mean S_0 and first two characteristic vectors S_1 and S_2 :

$$S(\lambda, T) = S_0(\lambda) + M_1(T)S_1(\lambda) + M_2(T)S_2(\lambda)$$

The vector S_1 corresponds to yellow-blue variation, accounting for changes in the correlated colour temperature due to presence or absence of clouds or direct sunlight. The vector S_2 corresponds to pink-green variation caused by the presence of water in the form of vapor and haze.

The coefficients M_1 and M_2 are derived from from the correlated

Figure 3: CIE Illuminant D daylight components [9].



colour temperature T via chromaticity coordinates (x, y) in the CIE 1931 colour space, setting $t = 10^3/T$:

$$x = \begin{cases} 0.244063 + 0.09911 t + 2.9678 t^2 - 4.6070 t^3 & 4000 \text{ K} \leq T \leq 7000 \text{ K} \\ 0.237040 + 0.24748 t + 1.9018 t^2 - 2.0064 t^3 & 7000 \text{ K} < T \leq 25000 \text{ K} \end{cases}$$

$$y = -3.000 x^2 + 2.870 x - 0.275$$

$$M = 0.0241 + 0.2562 x - 0.7341 y$$

$$M_1 = (-1.3515 - 1.7703 x + 5.9114 y)/M$$

$$M_2 = (0.0300 - 31.4424 x + 30.0717 y)/M$$

8 Profiles

In colour management, an ICC profile [18] is a set of data that characterizes a colour input or output device, or a colour space, according to standards promulgated by the International Color Consortium (ICC). Profiles describe the colour attributes of a particular device or viewing requirement by defining a mapping between the device source or target colour space and a profile connection space, which is either CIE LAB ($L^*a^*b^*$) or CIE XYZ. Mappings may be specified using tables, to which interpolation is applied, or through a series of parameters for transformations. A profile might define several mappings, according to rendering intent. These mappings allow a choice between closest possible colour matching, and remapping the entire colour range to allow for different gamuts.

The reference illuminant of the profile connection space is a 16-bit fractional approximation of D50; its white point is:

$$XYZ = (0.9642, 1.000, 0.8249)$$

There are a few ICC profiles for storing CIE XYZ data in image files [19]; `D50_XYZ.icc` is the simplest as no conversion is needed to the profile connection space.

9 Example

Here are some excerpts from a program that plots the spectrum of the CIE 1931 standard observer mixed with grey into a 16-bit PNG image file, adding an embedded ICC D50 XYZ colour profile, and setting the pixel dimensions to 600dpi.

7 PNG

Portable Network Graphics [16] (PNG) is a raster graphics file format that supports lossless data compression, and is the most used lossless image compression format on the Internet. A PNG file starts with an 8-byte signature, followed by a series of chunks.

A decoder must be able to interpret critical chunks to read and render a PNG file. “IHDR” must be the first chunk; it contains the image’s width, height, bit depth and colour type. “IDAT” chunks contains the actual image data, which is the output stream of the compression algorithm. Finally an “IEND” chunk marks the end of the image.

Other ancillary chunks that can be stored in PNG files include chunks for colour management (“cHRM” chromaticity, “gAMA” gamma, “sRGB” indicator, “iCCP” colour profile), image properties (“hIST” colour histogram, “pHYs” intended pixel dimensions, “sTER” stereoscopic indicator) and textual metadata information (“iTXt”, “tEXt”, “zTXt”).

PNG supports palette-based, grayscale, and colour images with three channels plus an optional alpha (opacity) channel. Each image can contain 8-bit or 16-bit values (integers from 0 to 255, or from 0 to 65535, respectively) in each channel.

PNG images can be interlaced, allowing a low-resolution image to be visible more quickly when transferring over a slow connection or when working with huge images.

libpng [17] is a library for decoding and encoding PNG files.

3 Eyesight


Cone cells [10], are one of the two types of photoreceptor cells that are in the retina of the eye which are responsible for colour vision. Cones are less sensitive to light than the rod cells in the retina (which support vision at low light levels), but allow the perception of colour. Humans normally have three kinds of cones, responding most to long (L) reddish wavelengths, medium (M) greenish wavelengths, and short (S) bluish wavelengths. The three types have peak wavelengths near 564–580nm, 534–545nm, and 420–440nm, respectively, depending on the individual.

The perception of colour [11] is achieved by a complex process that starts with the differential output of the cone cells in the retina and it will be finalized in the visual cortex and associative areas of the brain. Colour processing begins at a very early level in the visual system. Trichromacy arises at the level of the receptors, and opponent processes arise at the level of retinal ganglion cells and beyond.

The trichromatic theory, proposed in the 19th century by Thomas Young and Hermann von Helmholtz, states that the retina’s three types of cones are preferentially sensitive to blue, green, and red. Ewald Hering proposed the opponent process theory in 1872. It states that the visual system interprets colour in an antagonistic way: red vs. green, blue vs. yellow, black vs. white.

Nothing categorically distinguishes the visible spectrum of electromagnetic radiation from invisible portions of the broader spectrum. In this sense, colour is not a property of electromagnetic

Figure 4: Stockman & Sharpe 2-deg cone fundamentals [12].



radiation, but a feature of visual perception by an observer. Furthermore, there is an arbitrary mapping between wavelengths of light in the visual spectrum and human experiences of colour.

4 Observer

The CIE 1931 colour spaces [13] are a quantitative link between wavelengths of light and perceived colours. The mathematical relationships that define these colour spaces allow one to translate different physical responses to visible radiation in colour inks, illuminated displays, and recording devices such as digital

sRGB primary chromaticities:

	x	y	Y
red	0.6400	0.3300	0.2126
green	0.3000	0.6000	0.7152
blue	0.1500	0.0600	0.0722
white	0.3127	0.3290	1.0000 (D65)

CIE XYZ to sRGB (where C is each RGB channel):

$$\begin{pmatrix} R_{\text{lin}} \\ G_{\text{lin}} \\ B_{\text{lin}} \end{pmatrix} = \begin{pmatrix} 3.2406 & -1.5372 & -0.4986 \\ -0.9689 & 1.8758 & 0.0415 \\ 0.0557 & -0.2040 & 1.0570 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$$

$$C_{\text{srgb}} = \begin{cases} 12.92 C_{\text{lin}} & 0 \leq C_{\text{lin}} \leq 0.0031308 \\ 1.055 C_{\text{lin}}^{1/2.4} - 0.055 & 0.0031308 < C_{\text{lin}} \leq 1 \end{cases}$$

sRGB to CIE XYZ:

$$C_{\text{lin}} = \begin{cases} \frac{C_{\text{srgb}}}{12.92} & 0 \leq C_{\text{srgb}} \leq 0.04045 \\ \left(\frac{C_{\text{srgb}} + 0.055}{1.055} \right)^{2.4} & 0.04045 < C_{\text{srgb}} \leq 1 \end{cases}$$

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 0.4124 & 0.3576 & 0.1805 \\ 0.2126 & 0.7152 & 0.0722 \\ 0.0193 & 0.1192 & 0.9505 \end{pmatrix} \begin{pmatrix} R_{\text{lin}} \\ G_{\text{lin}} \\ B_{\text{lin}} \end{pmatrix}$$

6 sRGB

sRGB [15] is a standard RGB colour space created in 1996 for use on monitors, printers and the Internet. sRGB uses the ITU-R BT.709 red, green, and blue primaries (also used in studio monitors and HDTV), and a nonlinear transfer function that makes efficient use of the small integer values in an image file to display human-discernible light levels. sRGB is designed to match typical home and office viewing conditions, rather than the darker environment typically used for commercial colour matching. The triangle defined by the primaries is well inside the gamut of human vision, meaning many colours cannot be represented in sRGB.

Figure 6: sRGB gamut in CIE xyY colour space.





Figure 5: CIE 1931 standard colorimetric observer [14].



cameras into a universal human colour vision response.

The CIE XYZ colour space encompasses all colour sensations that an average person can experience. A set of colour-matching functions, like the spectral sensitivity LMS cone functions, associate physically produced light spectra with specific tristimulus values.

Two light sources with different spectra may appear to be the same colour; this effect is called metamerism. Light sources have the same apparent colour to an observer when they produce the same tristimulus values, independent of their precise spectra. Moreover, perception of colour depends on context. A

white page under blue or pink light will reflect mostly blue or pink light to the eye; the brain, however, compensates for the effect of lighting (based on the colour shift of surrounding objects) and interprets the page as white under both conditions, a phenomenon known as colour constancy.

Most wavelengths stimulate more than one type of cone cell, because the spectral sensitivity curves of the three types of cone cells overlap. Some tristimulus values are thus physically impossible. In a normal trichromatic additive colour space (e.g. RGB), pure spectral colours would imply negative values for at least one of the three primaries, since the chromaticity would be outside the colour triangle defined by the primary colours. To avoid negative RGB values, “imaginary” primary colours and corresponding colour-matching functions were formulated.

The CIE’s colour matching functions \bar{x} , \bar{y} and \bar{z} are the numerical description of the chromatic response of the observer. They can be thought of as the spectral sensitivity curves of three linear light detectors yielding the CIE tristimulus values X , Y and Z . Collectively, these three functions are known as the CIE standard observer. The tristimulus values for a colour with a spectral power distribution I are given in terms of the standard observer by:

$$X = \int_{380\text{nm}}^{780\text{nm}} I(\lambda) \bar{x}(\lambda) d\lambda$$

$$Y = \int_{380\text{nm}}^{780\text{nm}} I(\lambda) \bar{y}(\lambda) d\lambda$$

$$Z = \int_{380\text{nm}}^{780\text{nm}} I(\lambda) \bar{z}(\lambda) d\lambda$$

where λ is each wavelength of monochromatic light.

5 Chromaticity

The concept of colour can be divided into two parts: brightness and chromaticity. For example, grey is a less bright white, while the chromaticity of white and grey are the same. The Y parameter is a measure of the brightness or luminance of a colour. The chromaticity (x, y) is derived by a reversible normalization:

$$x = \frac{X}{X + Y + Z} \quad y = \frac{Y}{X + Y + Z}$$

$$X = \frac{Y}{y} x \quad Z = \frac{Y}{y} (1 - x - y)$$

Plate 1 shows the chromaticity diagram, a tool to specify how the human eye will experience light with a given spectrum. The chromaticities are unlikely to be reproduced correctly by any given output device, and the chromaticity observed while looking at an object also depends on the light source.

The gamut of all visible chromaticities on the CIE plot is the tongue-shaped or horseshoe-shaped figure shown in colour. The curved edge is called the spectral locus and corresponds to monochromatic light. The straight edge on the lower part is called the line of purples. These have no counterpart in monochromatic light. Less saturated colours appear in the interior of the figure with white at the center.

All the colours that lie in a straight line between any two points on the diagram can be formed by mixing those two colours. Distance on the chromaticity diagram does not correspond to the degree of difference between two colours. No number of real primaries can cover the gamut of human vision. Light with equal power in every 1 nm interval corresponds to the point $(\frac{1}{3}, \frac{1}{3})$.