CSE 291F: Computational Photography

Lecturer: Ben Ochoa

Notes by Owen Jow

1 Color

1.1 Terminology

- **Spectral power distribution (SPD)**: power (e.g. radiance) at each wavelength. Can represent discretely with 31 components, one for each 10 nm band in the 400-700 nm range.
- **Luminance**: radiant power weighted by predefined sensitivity function (*perceived lightness*).
- **Hue**: value described by dominant SPD wavelength (*pure color*).
- **Saturation**: degree of concentration of SPD at one wavelength (*colorfulness*).

1.2 Color Spaces

1.2.1 CIE XYZ

While we *can* represent a color with 31 components, this is inefficient for image coding. Due to the trichromatic nature of human vision (we have three types of cones), it is in fact sufficient to weight an SPD with only three spectral response functions, leading to a three-component representation.

Accordingly, in 1931, the CIE defined three "standard observer" spectral weighting functions – the CIE *color matching functions* – which are based on the famous 1920s color matching experiment and which allow us to distinguish in value the colors that we see as distinct in real life.

These color matching functions produce the three (linear) XYZ tristimulus values. Note that the Y value corresponds to luminance, while the X and Z values describe the actual color information. The CIE XYZ color space is special as a "common denominator" from which you can transform to any other color space; it is often used as an intermediate color space for this purpose.

Chromaticity To represent color without luminance, you can turn X, Y, Z values into x, y chromaticity values (and plot the resulting color representation in a 2D chromaticity diagram[1]).

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

[1] note: the outer curve of the chromaticity diagram consists of the chromaticity coordinates of *spectral colors* (colors arising from single wavelengths), while the interior consists of the chromaticity coordinates of non-spectral colors

There is also a $z = 1 - x - y$ chromaticity value, but it is redundant if x and y are given.

To convert from x, y to X and Z:

$$\begin{aligned} X &= \frac{Y}{y} x \\ Z &= \frac{Y}{y}(1 - x - y) \end{aligned}$$

[Assumption: Y is already known (a color is specified by chromaticity x, y and luminance Y).]

1.2.2 Additive RGB Systems

An additive RGB system can be specified by the chromaticities of its primaries and its white point. These chromaticities "anchor" the RGB system by telling us how to interpret RGB values in terms of fixed XYZ colors. We must be able to identify the primaries that we are additively mixing!

The primaries define the gamut (anything in the primaries' triangle on the chromaticity diagram). The white point, i.e. the chromaticity x, y arising from equal amounts of R, G, and B, characterizes how the chromaticity space is divided between the three primary points. It determines the "step" through the gamut for each unit of R, G, and B. Also, the triangle given by the white point and two of the primaries describes the sub-gamut of colors that you can make with just those two primaries. A common choice of white point is Illuminant D_{65}, which approximates daylight.

sRGB, Adobe RGB, and ProPhoto RGB are examples of additive RGB systems (/standards for encoding different gamuts). sRGB has a gamut containing about 35.9% of visible colors.

1.2.3 Gamuts: Converting Between Color Spaces

We can map between RGB color spaces with a 3×3 transformation matrix[2]. The basic idea is to map from the first RGB space (with some chromaticities) to XYZ, and then map from XYZ to the second RGB space (with some other chromaticities). To transform from an RGB color space to XYZ:

$$\begin{aligned} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} &= \mathbf{M}_{\text{RGB}\to\text{XYZ}} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \\ &= \begin{bmatrix} R_X & G_X & B_X \\ R_Y & G_Y & B_Y \\ R_Z & G_Z & B_Z \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \end{aligned}$$

where the transformation matrix $\mathbf{M}_{\text{RGB}\to\text{XYZ}}$ depends on the XYZ values of the primaries R, G, and B. It becomes simpler to solve for this matrix if we split it into two matrices:

$$\begin{aligned} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} &= \begin{bmatrix} R_x(R_X + R_Y + R_Z) & G_x(G_X + G_Y + G_Z) & B_x(B_X + B_Y + B_Z) \\ R_y(R_X + R_Y + R_Z) & G_y(G_X + G_Y + G_Z) & B_y(B_X + B_Y + B_Z) \\ R_z(R_X + R_Y + R_Z) & G_z(G_X + G_Y + G_Z) & B_z(B_X + B_Y + B_Z) \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \\ &= \begin{bmatrix} R_x & G_x & B_x \\ R_y & G_y & B_y \\ R_z & G_z & B_z \end{bmatrix} \begin{bmatrix} R_X + R_Y + R_Z & 0 & 0 \\ 0 & G_X + G_Y + G_Z & 0 \\ 0 & 0 & B_X + B_Y + B_Z \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \end{aligned}$$

[2] it's like a change of basis: how do you represent the same color in terms of different primaries and white points?

which follows from the idea that (e.g.) R_X, the X-value of the primary R, is equivalent to the x-value of R times the sum of the X, Y, and Z values of R (a definitional chromaticity conversion):

$$R_X = R_x(R_X + R_Y + R_Z) \quad \text{because} \quad X = x(X + Y + Z)$$

Note that the chromaticities of the primaries (i.e. the first matrix) are known; they are part of the color space description. So we only need to solve for the sums of the XYZ values of the primaries. To do so, we can substitute in other known values. Specifically, we know that RGB white is $(1, 1, 1)$ and that XYZ white is $(w_x/w_y, 1, w_z/w_y)$, where w_x, w_y, w_z are the chromaticities of the white point.

$$\begin{aligned}
\begin{bmatrix} w_x/w_y \\ 1 \\ w_z/w_y \end{bmatrix} &= \begin{bmatrix} R_x & G_x & B_x \\ R_y & G_y & B_y \\ R_z & G_z & B_z \end{bmatrix} \begin{bmatrix} R_X + R_Y + R_Z & 0 & 0 \\ 0 & G_X + G_Y + G_Z & 0 \\ 0 & 0 & B_X + B_Y + B_Z \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \\
&= \begin{bmatrix} R_x & G_x & B_x \\ R_y & G_y & B_y \\ R_z & G_z & B_z \end{bmatrix} \begin{bmatrix} R_X + R_Y + R_Z \\ G_X + G_Y + G_Z \\ B_X + B_Y + B_Z \end{bmatrix}
\end{aligned}$$

We can easily solve for $\begin{bmatrix} (R_X + R_Y + R_Z) & (G_X + G_Y + G_Z) & (B_X + B_Y + B_Z) \end{bmatrix}^\top$ and plug it back in to obtain the desired transformation matrix $\mathbf{M}_{\text{RGB}\to\text{XYZ}}$.

Therefore, to transform between arbitrary RGB color spaces with different gamuts, we can map from the first RGB color space to XYZ using $\mathbf{M}_{\text{RGB1}\to\text{XYZ}}$, and then map from XYZ to the second RGB color space using $(\mathbf{M}_{\text{RGB2}\to\text{XYZ}})^{-1}$. Note that everything should be linear at this stage.

1.2.4 Gamuts: Converting Between Color Spaces with Different White Points

However, if the color spaces have different white points (again, white points represent illuminants), we also need to perform chromatic adaptation. This involves converting from XYZ to LMS (the "long-medium-short" color space), then applying the adaptation matrix, then converting from LMS back to XYZ. The XYZ $\to$ LMS conversion is easy – just apply Bradford's transformation matrix:

$$\mathbf{M}_{\text{XYZ}\to\text{LMS}} = \begin{bmatrix} 0.8951 & 0.2664 & -0.1614 \\ -0.7502 & 1.7135 & 0.0367 \\ 0.0389 & -0.0685 & 1.0296 \end{bmatrix}$$

The adaptation matrix is what we use to convert between the different white points. To compute it, we first map each of the XYZ white points to LMS using Bradford's transformation matrix. Let the first white point in LMS be (L_1, M_1, S_1), and let the second white point in LMS be (L_2, M_2, S_2). Then the adaptation matrix is

$$\mathbf{M}_{\text{LMS1}\to\text{LMS2}} = \begin{bmatrix} L_2/L_1 & 0 & 0 \\ 0 & M_2/M_1 & 0 \\ 0 & 0 & S_2/S_1 \end{bmatrix}$$

Now we're ready to perform the full conversion.
If the color spaces have different white points, the full 3×3 transformation matrix is

$$(\mathbf{M}_{\text{RGB2}\to\text{XYZ}})^{-1} \cdot (\mathbf{M}_{\text{XYZ}\to\text{LMS}})^{-1} \cdot \mathbf{M}_{\text{LMS1}\to\text{LMS2}} \cdot \mathbf{M}_{\text{XYZ}\to\text{LMS}} \cdot \mathbf{M}_{\text{RGB1}\to\text{XYZ}}$$

1.3 Gamma Correction

Humans perceive light nonlinearly. Thus, for the sake of memory efficiency, it's better to encode images in a physically nonlinear (but *perceptually* linear) way! This ensures that the spacing between successive levels in an 8-bit (256-level) luminance image appears somewhat uniform to us. When such is the case, we are able to represent the full range of brightnesses more precisely.

1.3.1 Linear → Nonlinear

$$\mathbf{x}_{\text{nonlinear}} = \mathbf{x}_{\text{linear}}^{1/\gamma}$$

This will use more of the levels to represent very-dark values, and less of the levels to represent very-bright values (because humans are better at distinguishing dark values). Generally, $\gamma = 2.2$.

1.3.2 Nonlinear → Linear

$$\mathbf{x}_{\text{linear}} = \mathbf{x}_{\text{nonlinear}}^{\gamma}$$

1.3.3 Other Linear ↔ Nonlinear Mappings

Different standards use different gammas and/or transfer functions. See this page for other standards' mappings (pay attention to the "Color component transfer function" for each encoding).

As an example, the mapping from linear RGB[3] to (nonlinear) sRGB is

$$\mathbf{x}_{\text{nonlinear}} = \begin{cases} 12.92\ \mathbf{x}_{\text{linear}} & \text{for } \mathbf{x}_{\text{linear}} \leq 0.0031308 \\ 1.055\ \mathbf{x}_{\text{linear}}^{1/2.4} - 0.055 & \text{for } \mathbf{x}_{\text{linear}} > 0.0031308 \end{cases}$$

2 Image Restoration

2.1 Degradation Model

We assume that image degradation can be modeled as follows:

$$(\texttt{degraded image}) = (\texttt{degradation function}) * (\texttt{original image}) + (\texttt{noise image})$$

where $*$ denotes convolution.

Under this model, image restoration (i.e. reversing the degradation) comes down to

1. Removing the noise
2. Estimating the degradation function and performing a deconvolution

[3] with sRGB chromaticities

References

1. Ben Ochoa, Spring 2020 CSE 291F lectures.
2. Charles Poynton, A Guided Tour of Color Space.