

Color



15-463, 15-663, 15-862
Computational Photography
Fall 2021, Lecture 8

Course announcements

- Homework assignment 2 is out.
 - **Due Monday October 4th (extended deadline).**
 - **Make sure to download updated version (fixed .EXR issue and a typo in an equation).**
 - Requires camera *and* tripod.
 - Start early! *Substantially* larger programming and imaging components than in Homework 1.
 - Generous bonus component, up to 50% extra credit.
 - No really: start early!
 - Make sure to check the pinned FAQ thread on Piazza.
- Homework assignment 3 will be posted tonight.
- Please propose topics for reading group on Piazza: <https://piazza.com/class/ksm9uc16vsg4bf?cid=48>
 - Zero suggestions so far.

Overview of today's lecture

- Recap: color and human color perception.
- Retinal color space.
- Color matching.
- Linear color spaces.
- Chromaticity.
- Color calibration.
- Non-linear color spaces.
- Some notes about color reproduction.

Slide credits



Many of these slides were inspired or adapted from:

- Todd Zickler (Harvard).
- Fredo Durand (MIT).

Recap: color and human color perception


Color is an artifact of human perception

- “Color” is not an *objective* physical property of light (electromagnetic radiation).
- Instead, light is characterized by its wavelength.




What we call “color” is how we *subjectively* perceive a very small range of these wavelengths.

Light-material interaction



Light-material interaction




Illuminant Spectral Power Distribution (SPD)


- Most types of light “contain” more than one wavelengths.
- We can describe light based on the distribution of power over different wavelengths.



We call our sensation of all of these distributions “white”.




Light-material interaction




Spectral reflectance

- Most materials absorb and reflect light differently at different wavelengths.
- We can describe this as a ratio of reflected vs incident light over different wavelengths.



Light-material interaction




Human color vision


retinal color

$$\mathbf{c}(\ell(\lambda)) = (c_s, c_m, c_l)$$

$$c_s = \int k_s(\lambda) \ell(\lambda) d\lambda$$




LMS sensitivity functions



perceived color

object color

color names




Retinal vs perceived color



Retinal vs perceived color

Retinal vs
perceived color.




Retinal vs perceived color

- Our visual system tries to “adapt” to illuminant.
- We may interpret the same retinal color very differently.



Retinal vs perceived color




Human color vision

We will exclusively discuss retinal color in this course


retinal color

$$\mathbf{c}(\ell(\lambda)) = (c_s, c_m, c_l)$$

$$c_s = \int k_s(\lambda) \ell(\lambda) d\lambda$$



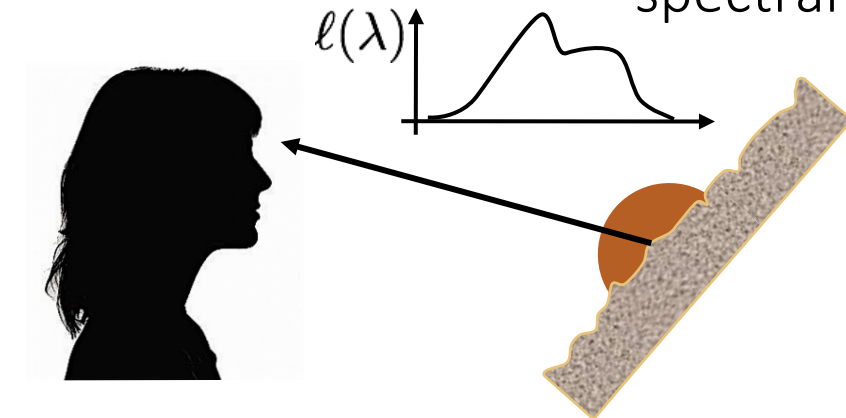
LMS sensitivity functions



perceived color

object color

color names



Retinal color space

Spectral Sensitivity Function (SSF)


- Any light sensor (digital or not) has different sensitivity to different wavelengths.
- This is described by the sensor's *spectral sensitivity function* $f(\lambda)$.
- When measuring light of some SPD $\Phi(\lambda)$ the sensor produces a *scalar* response:

$$\text{sensor response} \longrightarrow R = \int_{\lambda} \overset{\text{light SPD}}{\Phi(\lambda)} \overset{\text{sensor SSF}}{f(\lambda)} d\lambda$$

Weighted combination of light's SPD: light contributes more at wavelengths where the sensor has higher sensitivity.

Spectral Sensitivity Function of Human Eye

- The human eye is a collection of light sensors called cone cells.
- There are three types of cells with different spectral sensitivity functions.
- Human color perception is three-dimensional (*tristimulus color*).




cone distribution
for normal vision
(64% L, 32% M)

“short” $S = \int_{\lambda} \Phi(\lambda) S(\lambda) d\lambda$

“medium” $M = \int_{\lambda} \Phi(\lambda) M(\lambda) d\lambda$

“long” $L = \int_{\lambda} \Phi(\lambda) L(\lambda) d\lambda$




The retinal color space



$$\mathbf{c}(\ell_{\lambda_i}) = (c_s, c_m, c_l)$$



$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$



LMS sensitivity functions




“pure beam” (laser)

The retinal color space





$$\mathbf{c}(\ell_{\lambda_i}) = (c_s, c_m, c_l)$$



$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$



LMS sensitivity functions



"pure beam" (laser)

- "lasso curve"
- contained in positive octant
- parameterized by wavelength
- starts and ends at origin
- never comes close to M axis

← why?


← why?

The retinal color space





$$\mathbf{c}(\ell_{\lambda_i}) = (c_s, c_m, c_l)$$



$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$



LMS sensitivity functions



"pure beam" (laser)


if we also consider variations in the *strength* of the laser this "lasso" turns into (convex!) radial cone with a "horse-shoe shaped" radial cross-section

The retinal color space




$$\mathbf{c}(\ell_{\lambda_i}) = (c_s, c_m, c_l)$$




$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$



LMS sensitivity functions



colors of mixed beams are at the interior of the convex cone with boundary the surface produced by monochromatic lights




= convex combination of pure colors

The retinal color space





$$\mathbf{c}(\ell_{\lambda_i}) = (c_s, c_m, c_l)$$



$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$




LMS sensitivity functions



= convex combination of pure colors


- distinct mixed beams can produce the same retinal color
- these beams are called *metamers*

There is an infinity of metamers

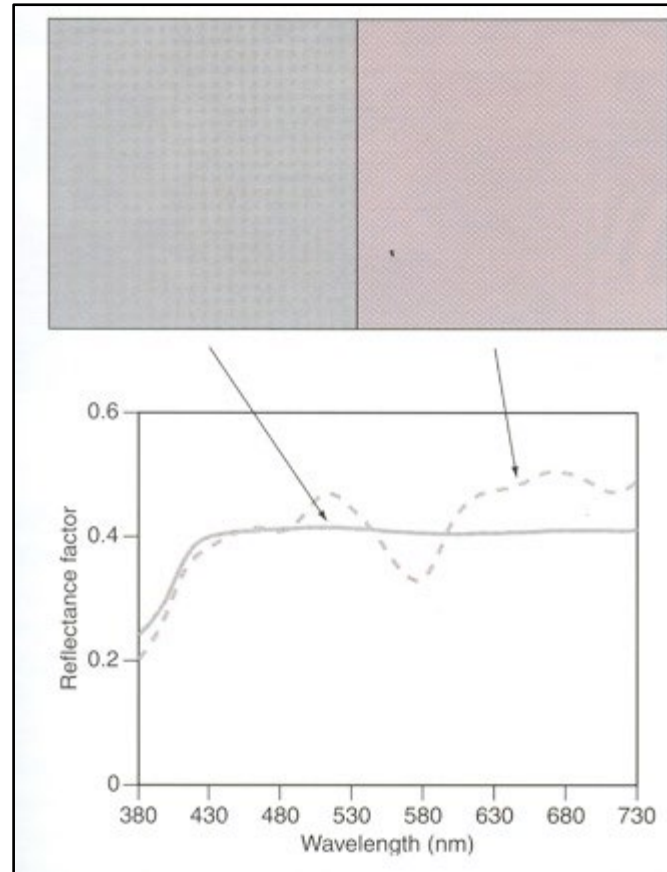


Ensemble of spectral reflectance curves corresponding to three chromatic-pigment recipes all matching a tan material when viewed by an average observer under daylight illumination. [Based on Berns (1988b).]


Example: illuminant metamerism



day light




scanned copy



hallogen light

Color matching

CIE color matching

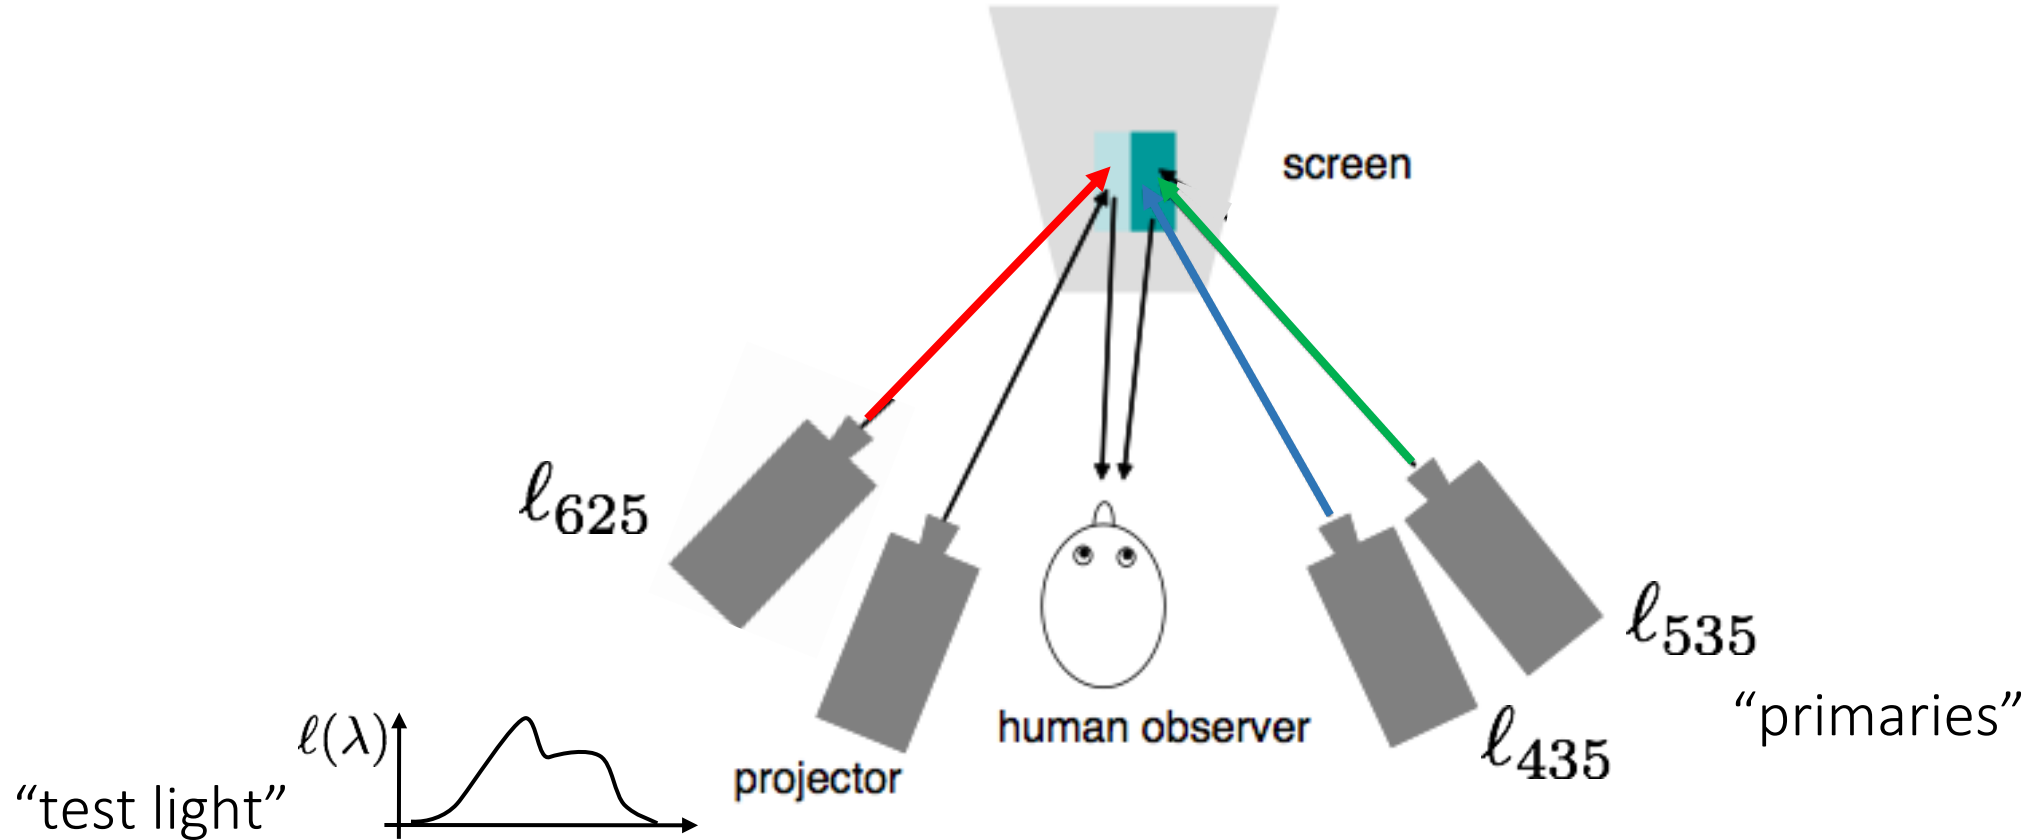


Adjust the strengths of the primaries until they re-produce the test color. Then:

$$\mathbf{c}(l(\lambda)) = \alpha \mathbf{c}(l_{435}) + \beta \mathbf{c}(l_{535}) + \gamma \mathbf{c}(l_{625})$$

← equality symbol means "has the same retinal color as" or "is metameric to"


CIE color matching






To match some test colors, you need to add some primary beam on the left (same as "subtracting light" from the right)

$$\begin{aligned} \mathbf{c}(l(\lambda)) + \gamma \mathbf{c}(l_{625}) &= \alpha \mathbf{c}(l_{435}) + \beta \mathbf{c}(l_{535}) \\ \longrightarrow \mathbf{c}(l(\lambda)) &= \alpha \mathbf{c}(l_{435}) + \beta \mathbf{c}(l_{535}) - \gamma \mathbf{c}(l_{625}) \end{aligned}$$

Color matching demo



CIE color matching






Repeat this matching experiments for pure test beams at wavelengths λ_i and keep track of the coefficients (negative or positive) required to reproduce each pure test beam.

$$\mathbf{c}(\lambda_i) = k_{435}(\lambda)\mathbf{c}(l_{435}) + k_{535}(\lambda)\mathbf{c}(l_{535}) + k_{625}(\lambda)\mathbf{c}(l_{625})$$

note the
negative values




CIE color matching



Repeat this matching experiments for pure test beams at wavelengths λ_i and keep track of the coefficients (negative or positive) required to reproduce each pure test beam.

$$\mathbf{c}(\lambda_i) = k_{435}(\lambda)\mathbf{c}(l_{435}) + k_{535}(\lambda)\mathbf{c}(l_{535}) + k_{625}(\lambda)\mathbf{c}(l_{625})$$


CIE color matching



What about "mixed beams"?

Two views of retinal color


$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$



LMS sensitivity functions

Analytic: Retinal color is produced by analyzing spectral power distributions using the color sensitivity functions.

$k_{435}(\lambda)$ $k_{545}(\lambda)$ $k_{625}(\lambda)$




Matching experiment matching functions

Synthetic: Retinal color is produced by synthesizing color primaries using the color matching functions.


What is each view of retinal color best suited for?

Two views of retinal color


$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$



LMS sensitivity functions



$k_{435}(\lambda)$ $k_{545}(\lambda)$ $k_{625}(\lambda)$



Matching experiment matching functions


Analytic: Retinal color is produced by analyzing spectral power distributions using the color sensitivity functions.

Synthetic: Retinal color is produced by synthesizing color primaries using the color matching functions.

How do they relate to each other?

Two views of retinal color


$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$



LMS sensitivity functions

Analytic: Retinal color is produced by analyzing spectral power distributions using the color sensitivity functions.

$k_{435}(\lambda)$ $k_{545}(\lambda)$ $k_{625}(\lambda)$



Matching experiment matching functions

Synthetic: Retinal color is produced by synthesizing color primaries using the color matching functions.

The two views are equivalent: Color matching functions are also color sensitivity functions. For each set of color sensitivity functions, there are corresponding color primaries.

Linear color spaces

Linear color spaces

1) Color matching experimental outcome:

$$\mathbf{c}(\lambda_i) = k_{435}(\lambda)\mathbf{c}(\ell_{435}) + k_{535}(\lambda)\mathbf{c}(\ell_{535}) + k_{625}(\lambda)\mathbf{c}(\ell_{625})$$

same in matrix form:

$$\begin{bmatrix} | \\ \mathbf{c}(\lambda_i) \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} k_{435} \\ k_{535} \\ k_{625} \end{bmatrix}$$



how is this matrix formed?

Linear color spaces

1) Color matching experimental outcome:

$$\mathbf{c}(\lambda_i) = k_{435}(\lambda)\mathbf{c}(\ell_{435}) + k_{535}(\lambda)\mathbf{c}(\ell_{535}) + k_{625}(\lambda)\mathbf{c}(\ell_{625})$$

same in matrix form:

$$\begin{bmatrix} | \\ \mathbf{c}(\lambda_i) \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} k_{435} \\ k_{535} \\ k_{625} \end{bmatrix}$$

2) Implication for arbitrary mixed beams:

$$\begin{bmatrix} | \\ \mathbf{c}(\ell(\lambda)) \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} \int k_{435}(\lambda)\ell(\lambda)d\lambda \\ \int k_{535}(\lambda)\ell(\lambda)d\lambda \\ \int k_{625}(\lambda)\ell(\lambda)d\lambda \end{bmatrix}$$



where do these terms come from?

Linear color spaces

1) Color matching experimental outcome:

$$\mathbf{c}(\lambda_i) = k_{435}(\lambda)\mathbf{c}(\ell_{435}) + k_{535}(\lambda)\mathbf{c}(\ell_{535}) + k_{625}(\lambda)\mathbf{c}(\ell_{625})$$

same in matrix form:

$$\begin{bmatrix} | \\ \mathbf{c}(\lambda_i) \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} k_{435} \\ k_{535} \\ k_{625} \end{bmatrix}$$

2) Implication for arbitrary mixed beams:

$$\begin{bmatrix} | \\ \mathbf{c}(\ell(\lambda)) \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} \int k_{435}(\lambda)\ell(\lambda)d\lambda \\ \int k_{535}(\lambda)\ell(\lambda)d\lambda \\ \int k_{625}(\lambda)\ell(\lambda)d\lambda \end{bmatrix}$$



what is this similar to?

Linear color spaces

1) Color matching experimental outcome:

$$\mathbf{c}(\lambda_i) = k_{435}(\lambda)\mathbf{c}(\ell_{435}) + k_{535}(\lambda)\mathbf{c}(\ell_{535}) + k_{625}(\lambda)\mathbf{c}(\ell_{625})$$

same in matrix form:

$$\begin{bmatrix} | \\ \mathbf{c}(\lambda_i) \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} k_{435} \\ k_{535} \\ k_{625} \end{bmatrix}$$

2) Implication for arbitrary mixed beams:

$$\begin{bmatrix} | \\ \mathbf{c}(\ell(\lambda)) \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} \int k_{435}(\lambda)\ell(\lambda)d\lambda \\ \int k_{535}(\lambda)\ell(\lambda)d\lambda \\ \int k_{625}(\lambda)\ell(\lambda)d\lambda \end{bmatrix}$$



representation of retinal
color in LMS space



change of basis matrix



representation of retinal
color in space of primaries

Linear color spaces

basis for retinal color \Leftrightarrow color matching functions \Leftrightarrow primary colors \Leftrightarrow color space

$$\begin{bmatrix} \mathbf{c}(\ell(\lambda)) \\ | \\ | \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}_1 & \mathbf{c}_2 & \mathbf{c}_3 \\ | & | & | \end{bmatrix} \begin{bmatrix} \int k_1(\lambda)\ell(\lambda)d\lambda \\ \int k_2(\lambda)\ell(\lambda)d\lambda \\ \int k_3(\lambda)\ell(\lambda)d\lambda \end{bmatrix}$$

$$\begin{bmatrix} \mathbf{c}(\ell_{435}) & | & \mathbf{c}(\ell_{545}) & | & \mathbf{c}(\ell_{625}) & | \\ | & | & | & | & | & | \end{bmatrix} \mathbf{M}^{-1} \quad \begin{bmatrix} k_1(\lambda) \\ k_2(\lambda) \\ k_3(\lambda) \end{bmatrix} = \mathbf{M} \begin{bmatrix} k_{435}(\lambda) \\ k_{545}(\lambda) \\ k_{625}(\lambda) \end{bmatrix}$$

$\mathbf{M}^{-1}\mathbf{M}$ can insert any invertible \mathbf{M}


$$\begin{bmatrix} \mathbf{c}(\ell(\lambda)) \\ | \\ | \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \mathbf{c}(\ell_{435}) & \mathbf{c}(\ell_{545}) & \mathbf{c}(\ell_{625}) \\ | & | & | \end{bmatrix} \begin{bmatrix} \int k_{435}(\lambda)\ell(\lambda)d\lambda \\ \int k_{535}(\lambda)\ell(\lambda)d\lambda \\ \int k_{625}(\lambda)\ell(\lambda)d\lambda \end{bmatrix}$$

representation of retinal
color in LMS space


change of basis matrix

representation of retinal
color in space of primaries

A few important color spaces



LMS color space



CIE RGB color space

not the “usual” RGB color space encountered in practice

Two views of retinal color

Analytic: Retinal color is three numbers formed by taking the dot product of a power spectral distribution with three color matching/sensitivity functions.

Synthetic: Retinal color is three numbers formed by assigning weights to three color primaries to match the perception of a power spectral distribution.


How would you make a color measurement device?

How would you make a color measurement device?

Do what the eye does:

- Select three spectral filters (i.e., three color matching functions.).
- Capture three measurements.

Can we use the CIE RGB color matching functions?



CIE RGB color space


How would you make a color measurement device?

Do what the eye does:

- Select three spectral filters (i.e., three color matching functions.).
- Capture three measurements.

Can we use the CIE RGB color matching functions?

Negative values are an issue (we can't "subtract" light at a sensor)




CIE RGB color space

How would you make a color measurement device?

Do what the eye does:

- Select three spectral filters (i.e., three color matching functions).
- Capture three measurements.

Can we use the LMS color matching functions?



LMS color space


How would you make a color measurement device?

Do what the eye does:

- Select three spectral filters (i.e., three color matching functions).
- Capture three measurements.

Can we use the LMS color matching functions?


- They weren't known when CIE was doing their color matching experiments.



LMS color space

The CIE XYZ color space

- Derived from CIE RGB by adding enough blue and green to make the red positive.
- Probably the most important *reference* (i.e., device independent) color space.




CIE XYZ color space

Remarkable and/or scary: 80+ years of CIE XYZ is all down to color matching experiments done with 12 “standard observers”.

The CIE XYZ color space


- Derived from CIE RGB by adding enough blue and green to make the red positive.
- Probably the most important *reference* (i.e., device independent) color space.




CIE XYZ color space

How would you convert a color image to grayscale?


A few important color spaces



LMS color space



CIE RGB color space



CIE XYZ color space

Two views of retinal color

Analytic: Retinal color is three numbers formed by taking the dot product of a power spectral distribution with three color matching/sensitivity functions.

Synthetic: Retinal color is three numbers formed by assigning weights to three color primaries to match the perception of a power spectral distribution.


How would you make a color reproduction device?

How would you make a color reproduction device?

Do what color matching does:

- Select three color primaries.
- Represent all colors as mixtures of these three primaries.

Can we use the XYZ color primaries?



CIE XYZ color space


How would you make a color reproduction device?

Do what color matching does:

- Select three color primaries.
- Represent all colors as mixtures of these three primaries.

Can we use the XYZ color primaries?


- No, because they are not “real” colors (they require an SPD with negative values).
- Same goes for LMS color primaries.



CIE XYZ color space

The Standard RGB (sRGB) color space

- Derived by Microsoft and HP in 1996, based on CRT displays used at the time.
- Similar but not equivalent to CIE RGB.



Note the negative values

sRGB color space

While it is called “standard”, when you grab an “RGB” image, it is highly likely it is in a different RGB color space...


The Standard RGB (sRGB) color space

- Derived by Microsoft and HP in 1996, based on CRT displays used at the time.
- Similar but not equivalent to CIE RGB.

There are really two kinds of sRGB color spaces: linear and non-linear.


- Non-linear sRGB images have the following tone reproduction curve applied to them.

$$C_{\text{non-linear}} = \begin{cases} 12.92 \cdot C_{\text{linear}}, & C_{\text{linear}} \leq 0.0031308 \\ (1 + 0.055) \cdot C_{\text{linear}}^{\frac{1}{2.4}} - 0.055, & C_{\text{linear}} \geq 0.0031308 \end{cases}$$




sRGB color space


A few important color spaces




LMS color space



CIE RGB color space




CIE XYZ color space




sRGB color space

A few important color spaces




LMS color space




CIE RGB color space

Is there a way to
“compare” all these color
spaces?




CIE XYZ color space



sRGB color space

Chromaticity

CIE xy (chromaticity)



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

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


$$(X, Y, Z) \longleftrightarrow (\underline{x, y}, Y)$$

chromaticity

luminance/brightness

Perspective projection of 3D retinal color space to two dimensions.

CIE xy (chromaticity)




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

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

$$(X, Y, Z) \longleftrightarrow (x, y, Y)$$


Note: These colors can be extremely misleading depending on the file origin and the display you are using

CIE xy (chromaticity)



What does the boundary of the chromaticity diagram correspond to?


Color gamuts



We can compare color spaces by looking at what parts of the chromaticity space they can reproduce with their primaries.

But why would a color space not be able to reproduce all of the chromaticity space?

Color gamuts




We can compare color spaces by looking at what parts of the chromaticity space they can reproduce with their primaries.

But why would a color space not be able to reproduce all of the chromaticity space?

- Many colors require negative weights to be reproduced, which are not realizable.


Color gamuts




sRGB color gamut:

- What are the three triangle corners?
- What is the interior of the triangle?
- What is the exterior of the triangle?

Color gamuts




Color gamuts




Gamuts of various common industrial RGB spaces

The problem with RGBs visualized in chromaticity space




Device 1 ———
 Device 2
 Device 3 - - -




RGB values have no meaning if the primaries between devices are not the same!

Color gamuts



- Can we create an RGB color space that reproduces the entire chromaticity diagram?
- What would be the pros and cons of such a color space?
- What devices would you use it for?

Chromaticity diagrams can be misleading



Different gamuts may compare very differently when seen in full 3D retinal color space.

Some take-home messages about color spaces

Analytic: Retinal color is three numbers formed by taking the dot product of a power spectral distribution with three color matching/sensitivity functions.

Synthetic: Retinal color is three numbers formed by assigning weights to three color primaries to match the perception of a power spectral distribution.


Fundamental problem: Analysis spectrum (camera, eyes) cannot be the same as synthesis one (display) - impossible to encode all possible colors without something becoming negative

- CIE XYZ only needs positive coordinates, but need primaries with negative light.
- RGB must use physical (non-negative) primaries, but needs negative coordinates for some colors.

Problem with current practice: Many different RGB color spaces used by different devices, without clarity of what exactly space a set of RGB color values are in.

- Huge problem for color reproduction from one device to another.

See for yourself



Images of the same scene captured using 3 different cameras with identical settings, supposedly in sRGB space.


Color calibration and affine transform estimation

Color calibration

Apply linear scaling and translation to RGB vectors in the image:

$$c' = M \cdot c + t$$

transformed RGB vector original RGB vector



What are the dimensions of each quantity in this equation?


Color calibration

Apply linear scaling and translation to RGB vectors in the image:

$$c' = M \cdot c + t$$

transformed RGB vector


original RGB vector

The diagram shows the equation $c' = M \cdot c + t$ centered on the page. Below the equation, there are two labels: "transformed RGB vector" on the left and "original RGB vector" on the right. An arrow points from the label "transformed RGB vector" to the variable c' in the equation. Another arrow points from the label "original RGB vector" to the variable c in the equation.

What are the dimensions of each quantity in this equation?

How do we decide what transformed vectors to map to?

Using (again) a colorchecker



Color patches manufactured to have pre-calibrated XYZ coordinates.

Calibration chart can be used for:

1. color calibration
2. radiometric calibration (i.e., response curve) using the bottom row

Using (again) a colorchecker

Color Checker Chart						
100	Dark Skin R=115 G= 82 B= 68	Light Skin R=194 G=150 B=130	Blue Sky R= 98 G=122 B=157	Foliage R= 87 G=108 B= 67	Blue Flower R=133 G=128 B=177	Bluish Green R=103 G=189 B=170
200						
300	Orange R=214 G=126 B= 44	Purple Red R= 80 G= 91 B=166	Moderate Red R=193 G= 90 B= 99	Purple R= 94 G= 60 B=108	Yellow Green R=157 G=188 B= 64	Orange Yellow R=224 G=163 B= 46
400						
500	Blue R= 56 G= 61 B=150	Green R= 70 G=148 B= 73	Red R=175 G= 54 B= 60	Yellow R=231 G=199 B= 31	Magenta R=187 G= 86 B=149	Cyan R= 8 G=133 B=161
600						
700	White R=243 G=243 B=242	Neutral 8 R=200 G=200 B=200	Neutral 65 R=160 G=160 B=160	Neutral 5 R=122 G=122 B=121	Neutral 35 R= 85 G= 85 B= 85	Black R= 52 G= 52 B= 52
800						
	X-Rite					
	200	400	600	800	1000	1200

Color patches manufactured to have pre-calibrated XYZ coordinates.

Can we use any colorchecker image for color calibration?

Calibration chart can be used for:

1. color calibration
2. radiometric calibration (i.e., response curve) using the bottom row

Using (again) a colorchecker

Color Checker Chart					
Dark Skin R=115 G= 82 B= 68	Light Skin R=194 G=150 B=130	Blue Sky R= 98 G=122 B=157	Foliage R= 87 G=108 B= 67	Blue Flower R=133 G=128 B=177	Bluish Green R=103 G=189 B=170
Orange R=214 G=126 B= 44	Purple Red R= 80 G= 91 B=166	Moderate Red R=193 G= 90 B= 99	Purple R= 94 G= 60 B=108	Yellow Green R=157 G=188 B= 64	Orange Yellow R=224 G=163 B= 46
Blue R= 56 G= 61 B=150	Green R= 70 G=148 B= 73	Red R=175 G= 54 B= 60	Yellow R=231 G=199 B= 31	Magenta R=187 G= 86 B=149	Cyan R= 8 G=133 B=161
White R=243 G=243 B=242	Neutral 8 R=200 G=200 B=200	Neutral 65 R=160 G=160 B=160	Neutral 5 R=122 G=122 B=121	Neutral 35 R= 85 G= 85 B= 85	Black R= 52 G= 52 B= 52

X-Rite

Color patches manufactured to have pre-calibrated XYZ coordinates.

Can we use any colorchecker image for color calibration?

- It needs to be a *linear* image!
- Do radiometric calibration first.

Calibration chart can be used for:


1. color calibration
2. radiometric calibration (i.e., response curve) using the bottom row

Color calibration

Apply linear scaling and translation to RGB vectors in the image:

$$c' = M \cdot c + t$$

transformed RGB vector original RGB vector



What are the dimensions of each quantity in this equation?

How do we decide what transformed vectors to map to?

How do we solve for matrix M and vector t?

Color calibration

Apply linear scaling and translation to RGB vectors in the image:

$$c' = [M \quad t] \begin{bmatrix} c \\ 1 \end{bmatrix}$$

Color calibration

Apply linear scaling and translation to RGB vectors in the image:

$$c' = \underbrace{[M \quad t]}_T \underbrace{\begin{bmatrix} c \\ 1 \end{bmatrix}}_C$$

Color calibration

Apply an affine transform to homogeneous RGB vectors in the image:

$$c' = T \cdot C$$

heterogeneous
transformed RGB vector

homogeneous
original RGB vector

The diagram shows the equation $c' = T \cdot C$ centered at the top. Below the equation, there are two arrows pointing towards the variables. The left arrow points from the text "heterogeneous transformed RGB vector" to the variable c' . The right arrow points from the text "homogeneous original RGB vector" to the variable C .

How do we solve for an affine transformation?

Determining the affine transform matrix

Write out linear equation for each color vector correspondence:

$$c' = T \cdot C \quad \text{or} \quad \begin{bmatrix} r' \\ g' \\ b' \end{bmatrix} = \begin{bmatrix} t_1 & t_2 & t_3 & t_4 \\ t_5 & t_6 & t_7 & t_8 \\ t_9 & t_{10} & t_{11} & t_{12} \end{bmatrix} \begin{bmatrix} r \\ g \\ b \\ 1 \end{bmatrix}$$

Determining the affine transform matrix

Rearrange into an equation involving a vectorized form of T:

$$\begin{bmatrix} r' \\ g' \\ b' \end{bmatrix} = \begin{bmatrix} r & g & b & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & r & g & b & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & r & g & b & 1 \end{bmatrix} \begin{bmatrix} t_1 \\ t_2 \\ t_3 \\ t_4 \\ t_5 \\ t_6 \\ t_7 \\ t_8 \\ t_9 \\ t_{10} \\ t_{11} \\ t_{12} \end{bmatrix}$$

Determining the affine transform matrix

Stack equations from multiple color vector correspondences:

$$\underbrace{\begin{bmatrix} r' \\ g' \\ b' \\ \vdots \\ r' \\ g' \\ b' \end{bmatrix}}_b = \underbrace{\begin{bmatrix} r & g & b & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & r & g & b & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & r & g & b & 1 \\ \vdots & & & & & & & & & & & \\ r & g & b & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & r & g & b & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & r & g & b & 1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} t_1 \\ t_2 \\ t_3 \\ t_4 \\ t_5 \\ t_6 \\ t_7 \\ t_8 \\ t_9 \\ t_{10} \\ t_{11} \\ t_{12} \end{bmatrix}}_x$$

Solving the linear system

Convert the system to a linear least-squares problem:

$$E_{LLS} = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$$

Expand the error:

$$E_{LLS} = \mathbf{x}^\top (\mathbf{A}^\top \mathbf{A}) \mathbf{x} - 2\mathbf{x}^\top (\mathbf{A}^\top \mathbf{b}) + \|\mathbf{b}\|^2$$

Minimize the error:

Set derivative to 0 $(\mathbf{A}^\top \mathbf{A})\mathbf{x} = \mathbf{A}^\top \mathbf{b}$

Solve for \mathbf{x} $\mathbf{x} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{b}$ ←

In Matlab:

$$\mathbf{x} = \mathbf{A} \setminus \mathbf{b}$$

Note: You almost never want to compute the inverse of a matrix.

An example



original

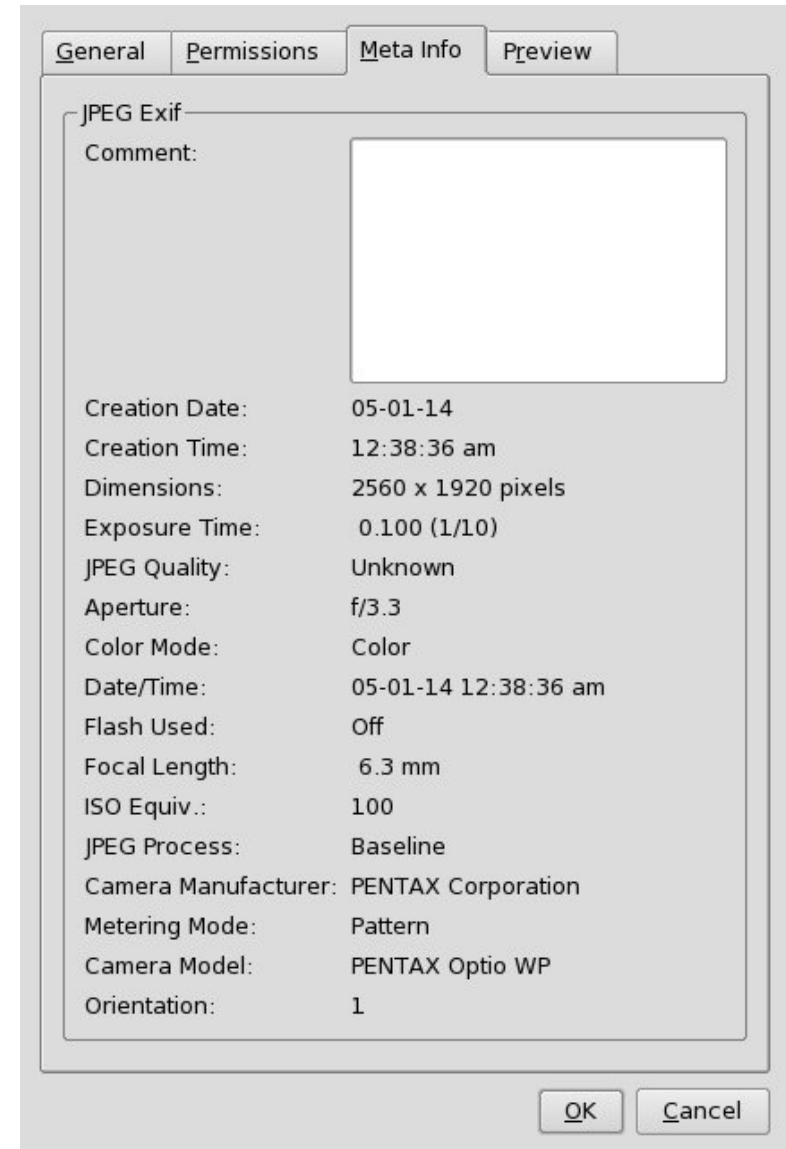


color-corrected

Quick note

If you cannot do calibration, take a look at the image's EXIF data (if available).

Often contains information about tone reproduction curve and color space.



Color profiling for displays



colorimeter: device calibrated to measure displayed radiance in some reference color space (usually CIE XYZ)

Exactly analogous procedure for figuring out the color space of a display.


Note: In displays, color *calibration* refers to changing the display's primaries so that colors are shown differently. This is a completely separate procedure from color profiling.

Note also the discrepancy in terminology between cameras and displays.


program displaying multiple color patches with known coordinates in the same color space as the colorimeter

Non-linear color spaces


A few important linear color spaces




LMS color space



e RGB color space




CIE XYZ color space



sRGB color space

What about non-linear color spaces?

CIE xy (chromaticity)



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

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

$$(X, Y, Z) \longleftrightarrow (\underline{x, y}, Y)$$

chromaticity

luminance/brightness


CIE xyY is a non-linear color space.

Uniform color spaces

Find map $F : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that perceptual distance can be well approximated using Euclidean distance:

$$d(\vec{c}, \vec{c}') \approx \|F(\vec{c}) - F(\vec{c}')\|_2$$

MacAdam ellipses





Areas in chromaticity space of imperceptible change:

- They are ellipses instead of circles.
- They change scale and direction in different parts of the chromaticity space.

MacAdam ellipses

Note: MacAdam ellipses are almost always shown at 10x scale for visualization. In reality, the areas of imperceptible difference are much smaller.



The Lab (aka L*ab, aka L*a*b*) color space

The L* component of *lightness* is defined as

$$L^* = 116f\left(\frac{Y}{Y_n}\right), \quad (2.105)$$

where Y_n is the luminance value for nominal white (Fairchild 2005) and

$$f(t) = \begin{cases} t^{1/3} & t > \delta^3 \\ t/(3\delta^2) + 2\delta/3 & \text{else,} \end{cases} \quad (2.106)$$

is a finite-slope approximation to the cube root with $\delta = 6/29$. The resulting 0...100 scale roughly measures equal amounts of lightness perceptibility.

In a similar fashion, the a* and b* components are defined as

$$a^* = 500 \left[f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right) \right] \quad \text{and} \quad b^* = 200 \left[f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right) \right], \quad (2.107)$$

where again, (X_n, Y_n, Z_n) is the measured white point. Figure 2.32i–k show the L*a*b* representation for a sample color image.

The Lab (aka L*ab, aka L*a*b*) color space

The L* component of *lightness* is defined as

$$L^* = 116f\left(\frac{Y}{Y_n}\right), \quad (2.105)$$

where Y_n is the luminance value for nominal white (Fairchild 2005) and

What is this?

$$f(t) = \begin{cases} t^{1/3} & t > \delta^3 \\ t/(3\delta^2) + 2\delta/3 & \text{else,} \end{cases} \quad (2.106)$$


is a finite-slope approximation to the cube root with $\delta = 6/29$. The resulting 0...100 scale roughly measures equal amounts of lightness perceptibility.

In a similar fashion, the a* and b* components are defined as

$$a^* = 500 \left[f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right) \right] \quad \text{and} \quad b^* = 200 \left[f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right) \right], \quad (2.107)$$

where again, (X_n, Y_n, Z_n) is the measured white point. Figure 2.32i–k show the L*a*b* representation for a sample color image.

Perceived vs measured brightness by human eye





Human-eye *response* (measured brightness) is linear.

However, human-eye *perception* (perceived brightness) is *non-linear*:

- More sensitive to dark tones.
- Approximately a Gamma function.


The Lab (aka $L^*a^*b^*$, aka $L^*a^*b^*$) color space



Hue, saturation, and value

Do not use color space HSV! Use LCh:

- L^* for “value”.
- $C = \sqrt{a^2 + b^2}$ for “saturation” (chroma).
- $h = \text{atan}(b / a)$ for “hue”.



How could you make an image like this from a color image?




How could you make an image like this from a color image?

Zero saturation

Control saturation with red-pass filter



Higher saturation




LCh

Easier to do color processing in ~~HSV~~ **LCh**

Some thoughts about color reproduction

The image processing pipeline

The sequence of image processing operations applied by the camera's image signal processor (ISP) to convert a RAW image into a "conventional" image.



Color reproduction notes

To properly reproduce the color of an image file, you need to?

Color reproduction notes

To properly reproduce the color of an image file, you need to convert it from the color space it was stored in, to a reference color space, and then to the color space of your display.

On the camera side:

- If the file is RAW, it *often* has EXIF tags with information about the RGB color space corresponding to the camera's color sensitivity functions.
- If the file is not RAW, you *may* be lucky and still find accurate information in the EXIF tags about what color space the image was converted in during processing.
- If there is no such information and you own the camera that shot the image, then you can do color calibration for the camera.
- If all of the above fails, assume sRGB.


On the display side:

- If you own a high-end display, it likely has accurate color profiles provided by the manufacturer.
- If not, you can use a spectrometer to do color profiling (not color calibration).
- Make sure your viewer does not automatically do color transformations.

Be careful to account for any gamma correction!

Amazing resource for color management and photography: <https://ninedegreesbelow.com/>

The METACOW spectral image database



Amazing dataset for color management and photography: https://www.rit.edu/cos/colorscience/rc_db_metacow.php

How do you convert an image to grayscale?

How do you convert an image to grayscale?

First, you need to answer two questions:

1) Is your image linear or non-linear?

- If the image is linear (RAW, HDR, or otherwise radiometrically calibrated), skip this step.
- If the image is nonlinear (PNG, JPEG, etc.), you must undo the tone reproduction curve.
 - i. If you can afford to do radiometric calibration, do that.
 - ii. If your image has EXIF tags, check there about the tone reproduction curve.
 - iii. If your image is tagged as non-linear sRGB, use the inverse of the sRGB tone reproduction curve.
 - iv. If none of the above, assume sRGB and do as in (iii).

2) What is the color space of your image?

- If it came from an original RAW file, read the color transform matrix from there (e.g., dcraw).
- If not, you need to figure out the color space.
 - i. If you can afford to do color calibration, use that.
 - ii. If your image has EXIF tags, check there about the color space.
 - iii. If your image is tagged as non-linear sRGB, use the color transform matrix for linear sRGB.
 - iv. If none of the above, assume sRGB and do (iii).

With this information in hand:

- Transform your image into the XYZ color space. (If it is in sRGB, you may need to do whitepoint adaptation!!)
- Extract the Y channel.
- If you want brightness instead of luminance, apply the Lab brightness non-linearity.

How do you convert an image to grayscale?

Why You Should Forget Luminance Conversion and Do Something Better

Rang M. H. Nguyen
National University of Singapore
nguyenho@comp.nus.edu.sg

Michael S. Brown
York University
mbrown@eecs.yorku.ca

Abstract

One of the most frequently applied low-level operations in computer vision is the conversion of an RGB camera image into its luminance representation. This is also one of the most incorrectly applied operations. Even our most trusted softwares, Matlab and OpenCV, do not perform luminance conversion correctly. In this paper, we examine the main factors that make proper RGB to luminance conversion difficult, in particular: 1) incorrect white-balance, 2) incorrect gamma/tonne-curve correction, and 3) incorrect equations. Our analysis shows errors up to 50% for various colors are not uncommon. As a result, we argue that for most computer vision problems there is no need to attempt luminance conversion; instead, there are better alternatives depending on the task.

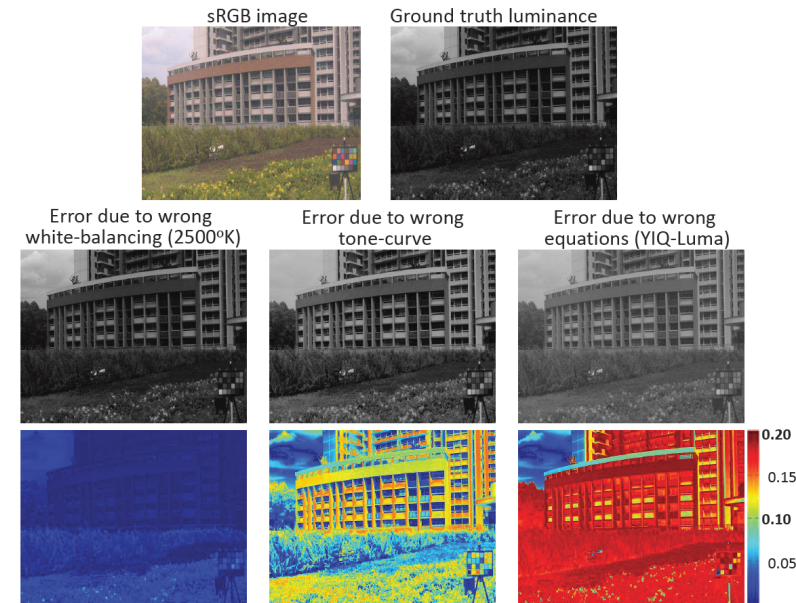


Figure 1. This figure shows examples of errors that arise due to improper luminance conversion. The ground truth luminance for this experiment is captured from a hyperspectral camera.

References

Basic reading:

- Szeliski textbook, Section 2.3.2, 3.1.2
- Michael Brown, “Understanding the In-Camera Image Processing Pipeline for Computer Vision,” CVPR 2016, Very detailed discussion of issues relating to color photography and management, slides available at: http://www.comp.nus.edu.sg/~brown/CVPR2016_Brown.html
- Gortler, “Foundations of 3D Computer Graphics,” MIT Press 2012.
Chapter 19 of this book has a great coverage of color spaces and the theory we discussed in class, it is available in PDF form from the CMU library.

Additional reading:

- Reinhard et al., “Color Imaging: Fundamentals and Applications,” A.K Peters/CRC Press 2008.
- Koenderink, “Color Imaging: Fundamentals and Applications,” MIT Press 2010.
- Fairchild, “Color Appearance Models,” Wiley 2013.
All of the above books are great references on color photography, reproduction, and management. The book by Reinhard et al. is my go-to reference on color.
- Nine Degrees Below, <https://ninedegreesbelow.com/>
Amazing resource for color photography, reproduction, and management.
- Bruce Lindbloom’s website, <http://brucelindbloom.com/>
An online page with a lot of information about color transforms, adaptation, and so on.
- MetaCow, https://www.rit.edu/cos/colorscience/rc_db_metacow.php
The best colorchecker dataset ever.
- Kim et al., “A New In-Camera Imaging Model for Color Computer Vision and Its Application,” PAMI 2012.
A detailed discussion of color processing in the image processing pipeline of modern cameras, and how to do color calibration for accurate color reproduction.
- Nguyen and Brown, “Why You Should Forget Luminance Conversion and Do Something Better,” CVPR 2017.
A detailed discussion of all the intricacies and challenges in trying to convert a color image to grayscale.