




1950
FUTURE VISION



11 FEBRUARY 2019
COMPUTER VISION

Proj 2

Looks like the evaluation function changed in converting to Python, and 80% on Notre Dame is more tricky to reach.

We will tweak the percentages.

Leaderboard / Gradescope is up.

Extra Credit

Please tell us which extra credit you attempted in its own section of your writeup.

I've amended the writeup.tex to make this explicit.

Alternative Textbook

Concise Computer Vision, Klette, 2014

Find & Borrow Articles, Journals, & Databases Research Help Hours, Locations, & Events Ask a Question Now Of



BROWN UNIVERSITY LIBRARY

[Search Home](#) [Catalog](#) [Reserves](#) [Feedback](#) [Bookmarks](#) [My History](#) [Classic Josia](#)

Concise computer vision an introduction into theory and algorithms


Reinhard Klette.
London : Springer, [2014] ©2014.


"Many textbooks on computer vision can be unwieldy and intimidating in their coverage of this extensive discipline. This textbook addresses the need for a concise overview of the fundamentals of this field. Concise Computer Vision provides an accessible general introduction to the essential topics in computer vision, highlighting the role of important algorithms and mathematical concepts. Classroom-tested programming exercises and review questions are also supplied at the end of each chapter. This concise and easy to read textbook/reference is ideal for an introductory course at third- or fourth-year level in an undergraduate computer science or engineering programme." --

Availability

Online Resources

- [Full text available from SpringerLINK ebooks - Computer Science \(2014\)](#)








Lenses

Let's design a camera


Idea 1: Put a sensor in front of an object
Do we get a reasonable image?




Let's design a camera

Idea 2: Add a barrier to block most rays

- Pinhole in barrier
- Only sense light from one direction.
 - Reduces blurring.
- In most cameras, this **aperture** can vary in size.




Pinhole camera model



f = Focal length
 c = Optical center of the camera

Figure from Forsyth

Projection: world coordinates \rightarrow image coordinates




$\mathbf{p} = \begin{bmatrix} U \\ V \end{bmatrix}$
 \mathbf{p} = distance from image center

$$U = -X * \frac{f}{Z} \quad V = -Y * \frac{f}{Z}$$

What is the effect if f and Z are equal?

Camera Obscura

Camera Obscura, Gemma Frisius, 1558



The first camera

- Known to Aristotle
- Depth of the room is the effective focal length


Home-made pinhole camera




Why so
blurry?

<http://www.debevec.org/Pinhole/>

Shrinking the aperture



2 mm



1 mm




0.6mm




0.35 mm

Less light gets through

Integrate over fewer angles



Shrinking the aperture




Less light gets through

Why not make the aperture as small as possible?


- Less light gets through
- Diffraction effects...

Shrinking the aperture - diffraction




Light diffracts as wavelength of aperture equals wavelength of light

The reason for lenses




Focus and Defocus



A lens focuses light onto the film

- There is a specific distance at which objects are “in focus”
 - other points project to a “circle of confusion” in the image
- Changing the shape of the lens changes this distance

Thin lenses



Thin lens equation:
$$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$$

$$\frac{1}{f} - \frac{1}{d_o} = \frac{1}{d_i}$$

Any object point satisfying this equation is in focus


What is the shape of the focus region?

How can we change the focus region?


Thin lens applet: <https://sites.google.com/site/marclevoylectures/applets/operation-of-a-thin-lens>

(by Andrew Adams, Nora Willett, Marc Levoy)

Beyond Pinholes: Real apertures



Bokeh:




[Rushif – Wikipedia]

Depth Of Field


Depth of Field



Depth of Field



Aperture controls Depth of Field



Changing the aperture size affects depth of field

- A smaller aperture increases the range in which the object is approximately in focus
- But small aperture reduces amount of light – need to increase exposure

Varying the aperture

Large aperture = small DOF



Small aperture = large DOF



Accidental Cameras



Accidental Pinhole and Pinspeck Cameras
Revealing the scene outside the picture.
Antonio Torralba, William T. Freeman

Accidental Cameras



a) Input (occluder present)



b) Reference (occluder absent)



c) Difference image (b-a)



d) Crop upside down





e) True view

DSLR – Digital Single Lens Reflex Camera




DSLR – Digital Single Lens Reflex Camera

“See what the main lens sees”




1. Objective (main) lens
2. Mirror
3. Shutter
4. Sensor
5. Mirror in raised position
6. Viewfinder focusing lens
7. Prism
8. Eye prescription lens

Shutters



Shutters



Shutters



Sensors: Rolling shutter vs. global shutter

Most modern cameras have purely digital shutters.



Sensor ISO


ISO = old film terminology
= sensitivity to light

ISO 200 is twice as sensitive as ISO 100.

Digital Photography:


ISO = 'gain' or amplification of sensor signal





Field of View (Zoom)

Field of View (Zoom)



17mm



28mm




50mm




85mm

From London and Upton

Field of View (Zoom) = Cropping



135mm



300mm




500mm



1000mm

From London and Upton

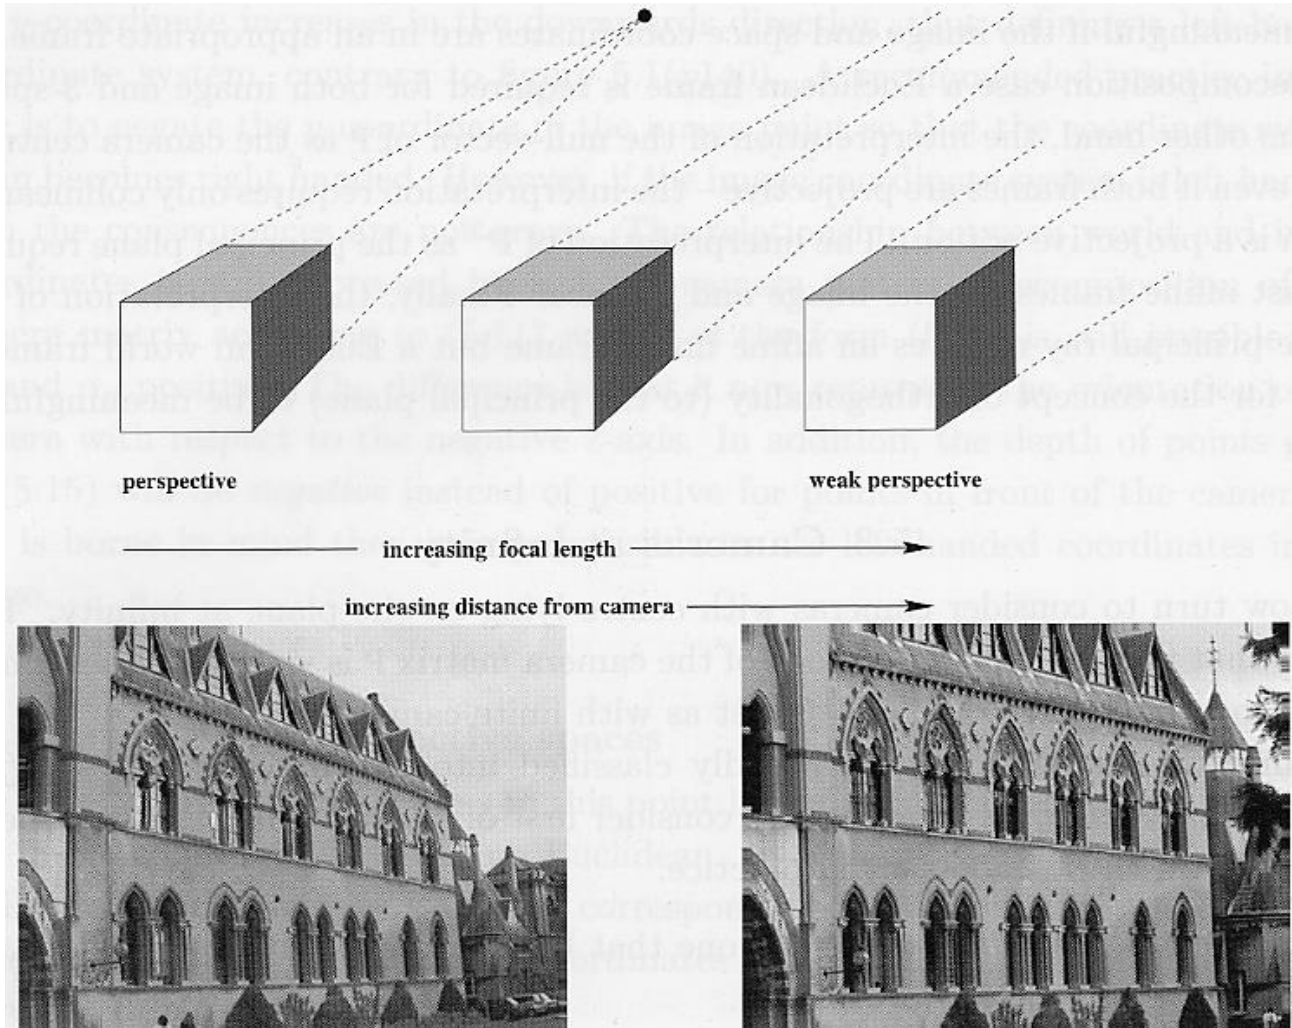
FOV depends of Focal Length



Size of field of view governed by size of the camera retina:

$$\varphi = \tan^{-1}\left(\frac{d}{2f}\right)$$

Smaller FOV = larger Focal Length



Field of View / Focal Length




Large FOV, small f
Camera close to car



Small FOV, large f
Camera far from the car

Fun with Focal Length (Jim Sherwood)



<http://www.hash.com/users/jsherwood/tutes/focal/Zoomin.mov>



Figure 5.1




Figure 5.2


Lens Flaws

Lens Flaws: Chromatic Aberration

Dispersion: wavelength-dependent refractive index

- (enables prism to spread white light beam into rainbow)

Modifies ray-bending and lens focal length: $f(\lambda)$



Color fringes near edges of image

Corrections: add 'doublet' lens of flint glass, etc.

Chromatic Aberration

Near Lens Center



Near Lens Outer Edge




Radial Distortion (e.g. 'barrel' and 'pin-cushion')


Straight lines curve around the image center




Radial Distortion



No distortion



Pin cushion



Barrel

Radial distortion of the image

- Caused by imperfect lenses
- Deviations are most noticeable for rays that pass through the edge of the lens



Corrected Barrel Distortion

Vignetting

Optical system occludes rays entering at obtuse angles.

Causes darkening at edges.


'Old mode' - but WHY?

Computer-aided lens design
(optimization) and
manufacturing made
removing (all) these flaws
much easier.




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- **Absorption**
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- **Diffuse Reflection**
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- **Specular Reflection**
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- **Transparency**
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- **Refraction**
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- **Fluorescence**
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- **Subsurface scattering**
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- **Phosphorescence**
- Interreflection



A photon's life choices


- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- **Interreflection**




Lambertian Reflectance

- In computer vision, surfaces are often assumed to be ideal diffuse reflectors with no dependence on viewing direction.

Grayscale intensity



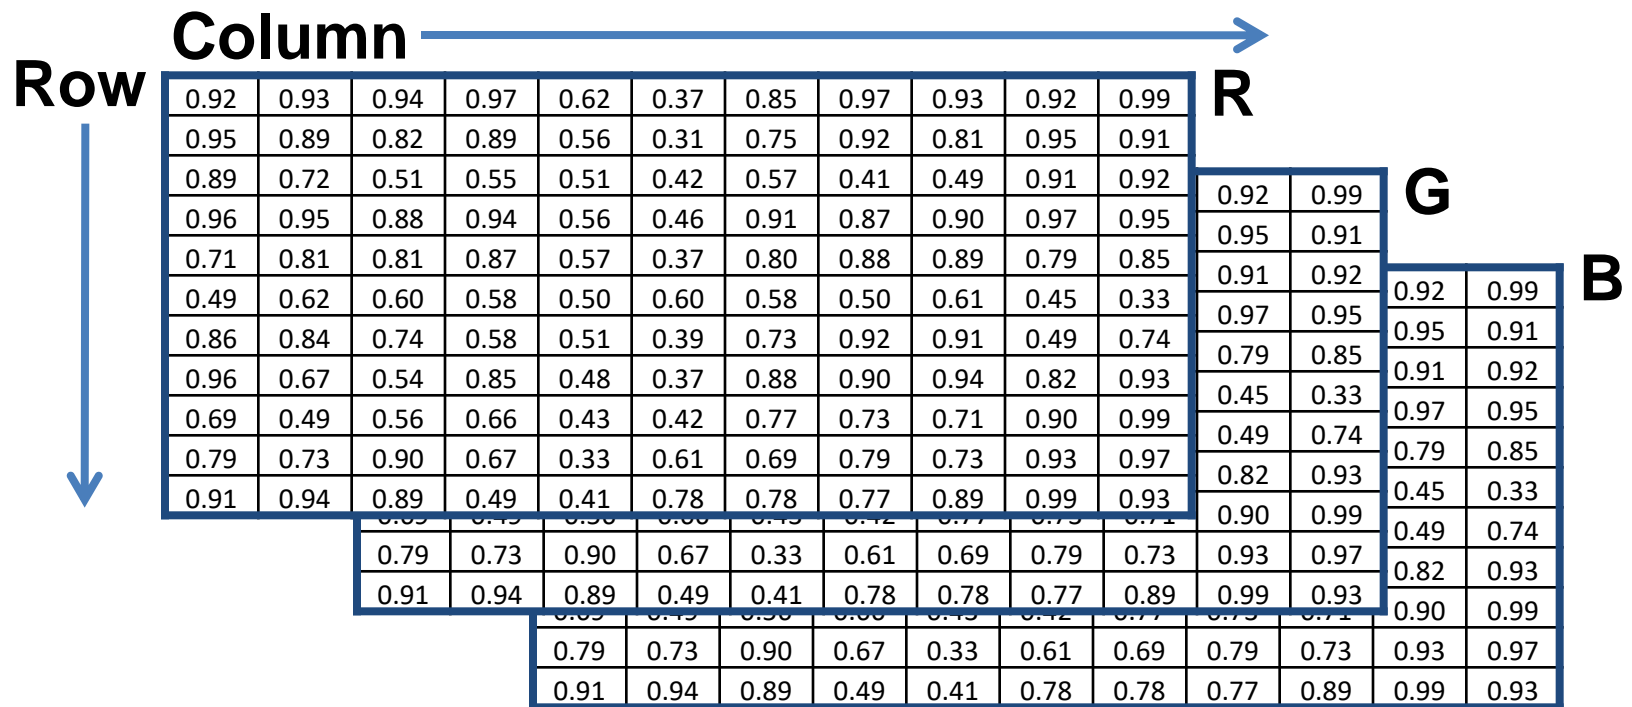
Color



Images in Python Numpy

N x M RGB image "im"


- `im[0,0,0]` = top-left pixel value in R-channel
- `im[x, y, b]` = x pixels to right, y pixels down in the bth channel
- `im[N-1, M-1, 3]` = bottom-right pixel in B-channel



But what is color?


ANATOMY

The Eye



- The human eye is a camera
 - **Iris** - colored annulus with radial muscles
 - **Pupil** - the hole (aperture) whose size is controlled by the iris
 - What's the sensor?
 - photoreceptor cells (rods and cones) in the **retina**

The Retina




Wait, the blood vessels are in front of the photoreceptors??


https://www.youtube.com/watch?v=L_W-IXqoxHA

What humans don't have: tapetum lucidum



Human eyes can reflect a tiny bit and blood in the retina makes this reflection red.





Tapetum lucidum exposed (cow eye)


Two types of light-sensitive receptors

Cones


cone-shaped
less sensitive
operate in high light
color vision

Rods

rod-shaped
highly sensitive
operate at night
gray-scale vision




Distribution of Rods and Cones




Night Sky: why are there more stars off-center?

Averted vision: http://en.wikipedia.org/wiki/Averted_vision

Rod / Cone sensitivity



Does the eye alias?




4x downsample
nearest neighbor

Spatially, apparently not.

The retina (sensor) has high resolution, but the optics (lens) of the eye cannot meet that resolution.

The image is blurred optically before being sampled (removes high-frequency content!)



Electromagnetic Spectrum



Human Luminance Sensitivity Function

Physiology of Color Vision


Three kinds of cones:



The Physics of Light

Any patch of light can be completely described physically by its spectrum: the number of photons (per time unit) at each wavelength 400 - 700 nm.


Photons
(per ms.)




The Physics of Light

Some examples of the spectra of light sources


A. Ruby Laser




B. Gallium Phosphide Crystal



C. Tungsten Lightbulb



D. Normal Daylight







The Physics of Light

Some examples of the reflectance spectra of surfaces



% Photons Reflected



400

700

400

700

400

700



400

700

Wavelength (nm)

Physiology of Color Vision

Three kinds of cones:




- Why are M and L cones so close?
- Why are there 3?

Impossible Colors


Can you make the cones respond in ways that typical light spectra never would?

http://en.wikipedia.org/wiki/Impossible_colors




CHIMERICAL COLOR DEMO TEMPLATES


Fatigue template
(stare at "x")






Target field
(glance at "x")






Approximate
Rendering



STYGIAN BLUE
(simultaneously deep
blue and black)




SELF-LUMINOUS RED
(simultaneously red and
brighter than white)



HYPERBOLIC ORANGE
(more than 100%
color saturation)

Tetrachromatism




- Most birds, and many other animals, have cones for ultraviolet light.
- Some humans seem to have four cones (12% of females).
- True tetrachromatism is rare; requires learning.

Bee vision




Metamers

Input




Frequency

Input




Frequency

Result



Result



by Jeff Beall, Adam Doppelt and John F. Hughes

(c) 1995 Brown University and the NSF Graphics and Visualization Center

What is color?

Why do we even care about
human vision in this class?


Why do we care about human vision?

- We don't, necessarily.
- But biological vision shows that it is possible to make important judgements from images.

Why do we care about human vision?

- We don't, necessarily.
- But biological vision shows that it is possible to make important judgements from images.
- It's a human world -> cameras imitate the frequency response of the human eye to try to see as we see.

Ornithopters



"Can machines fly like a bird?"

No, because airplanes don't flap.



"Can machines fly?"

Yes, but airplanes use a different mechanism.

"Can machines perceive?"

Is this question like the first, or like the second?


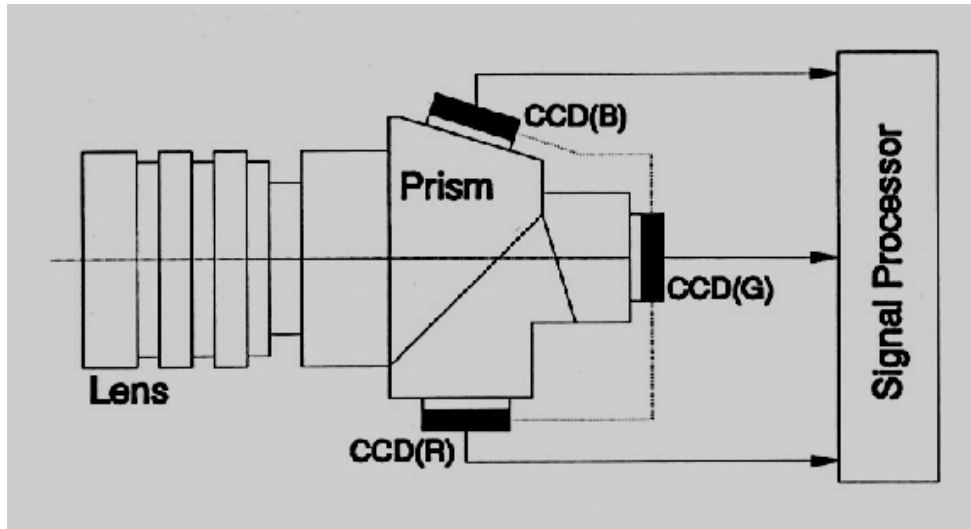

Cameras with Three Sensors



[Edmund Optics; Adam Wilt]

Color Sensing in Camera (RGB)

- 3-chip vs. 1-chip: quality vs. cost
- Why more green?





Bayer filter

Ruff Works

Why 3 colors?

<http://www.cooldictionary.com/words/Bayer-filter.wikipedia>


Cheaper/More Compact Color Sensing: Bayer Grid



Estimate RGB
at 'G' cells from
neighboring
values

Why more green?

Approximate human spectral sensitivity



←
Less than
~400nm to 10nm
= ultraviolet (UV)

Human visible portion
of electromagnetic
(EM) spectrum

→
Greater than
~700nm to 1mm
= infrared (IR)

RGB Camera Color Response



Canon 450D Quantum Efficiency



Display Color Response






PA248Q LCD Monitor Color Calibration Testing Report



Every ASUS PA248Q is equipped with pre-tuned sRGB and has undergone rigorous tests and calibration processes to ensure that color difference, ΔE , is less than 5, thus preventing color inaccuracy and inconsistency on screen.
ASUS advanced gray-scale tracking technology ensures smoother color gradation delivered by every ASUS PA248Q.

Serial No.	J3LMQ5158406
Test Equipment	Minolta Color Analyzer CA210



Note: The sRGB calibrations of every ASUS PA248Q are pre-tuned and tested under ASUS standardized procedures using calibrated facilities at the factory manufacturing line. This report is a certificate only for the newly manufactured ASUS PA248Q monitor unit. Test results may vary under different test procedures, equipment and patterns.





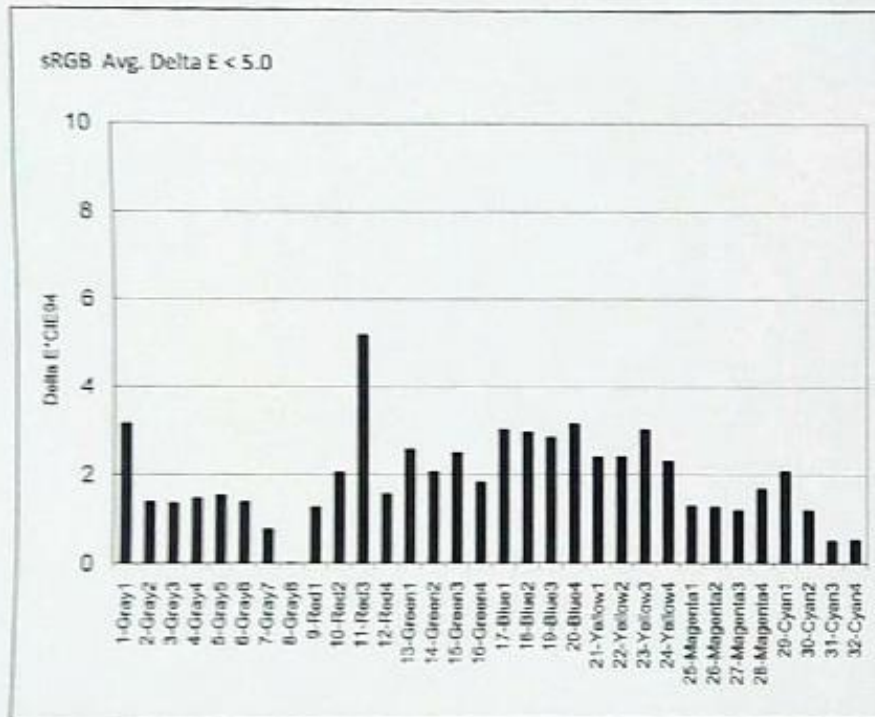
Display Color Response



PA248Q LCD Monitor Color Calibration Testing Report

Every ASUS PA248Q is equipped with pre-tuned sRGB and has undergone rigorous tests and calibration processes to ensure that color difference, ΔE , is less than 5, thus preventing color inaccuracy and inconsistency on screen.
ASUS advanced gray-scale tracking technology ensures smoother color gradation delivered by every ASUS PA248Q.

Serial No.	J4LMQS157085
Test Equipment	Minolta Color Analyzer CA210




Note: The sRGB calibrations of every ASUS PA248Q are pre-tuned and tested under ASUS standardized procedures using calibrated facilities at the factory manufacturing line. This report is a certificate only for the newly manufactured ASUS PA248Q monitor unit. Test results may vary under different test procedures, equipment and patterns.




Color spaces

How can we represent color?



Color spaces: RGB

Default color space



R = 1
(G=0,B=0)



G = 1
(R=0,B=0)



B = 1
(R=0,G=0)

Any color = $r \cdot R + g \cdot G + b \cdot B$


- Strongly correlated channels
- Non-perceptual

Got it. $C = r*R + g*G + b*B$

IS COLOR A VECTOR SPACE?
THINK-PAIR-SHARE

Color spaces: HSV

Intuitive color space



If you had to choose, would you rather go without:

- intensity ('value'), or
- hue + saturation ('chroma')?

Think-Pair-Share

Most information in intensity




Only color shown – constant intensity

Most information in intensity



Only intensity shown – constant color




Most information in intensity



Original image

Color spaces: HSV

Intuitive color space



H
(S=1,V=1)




S
(H=1,V=1)



V
(H=1,S=0)

Color spaces: YCbCr

Fast to compute, good for compression, used by TV



Y
(Cb=0.5,Cr=0.5)



Cb
(Y=0.5,Cr=0.5)



Cr
(Y=0.5,Cb=0.5)

Most JPEG images & videos subsample chroma




PSP Comp 3
2x2 Chroma subsampling
285K

Original
1,261K lossless
968K PNG

**IS COLOR PERCEPTION
A VECTOR SPACE?**

Color spaces: $L^*a^*b^*$

“Perceptually uniform”* color space



L
($a=0, b=0$)

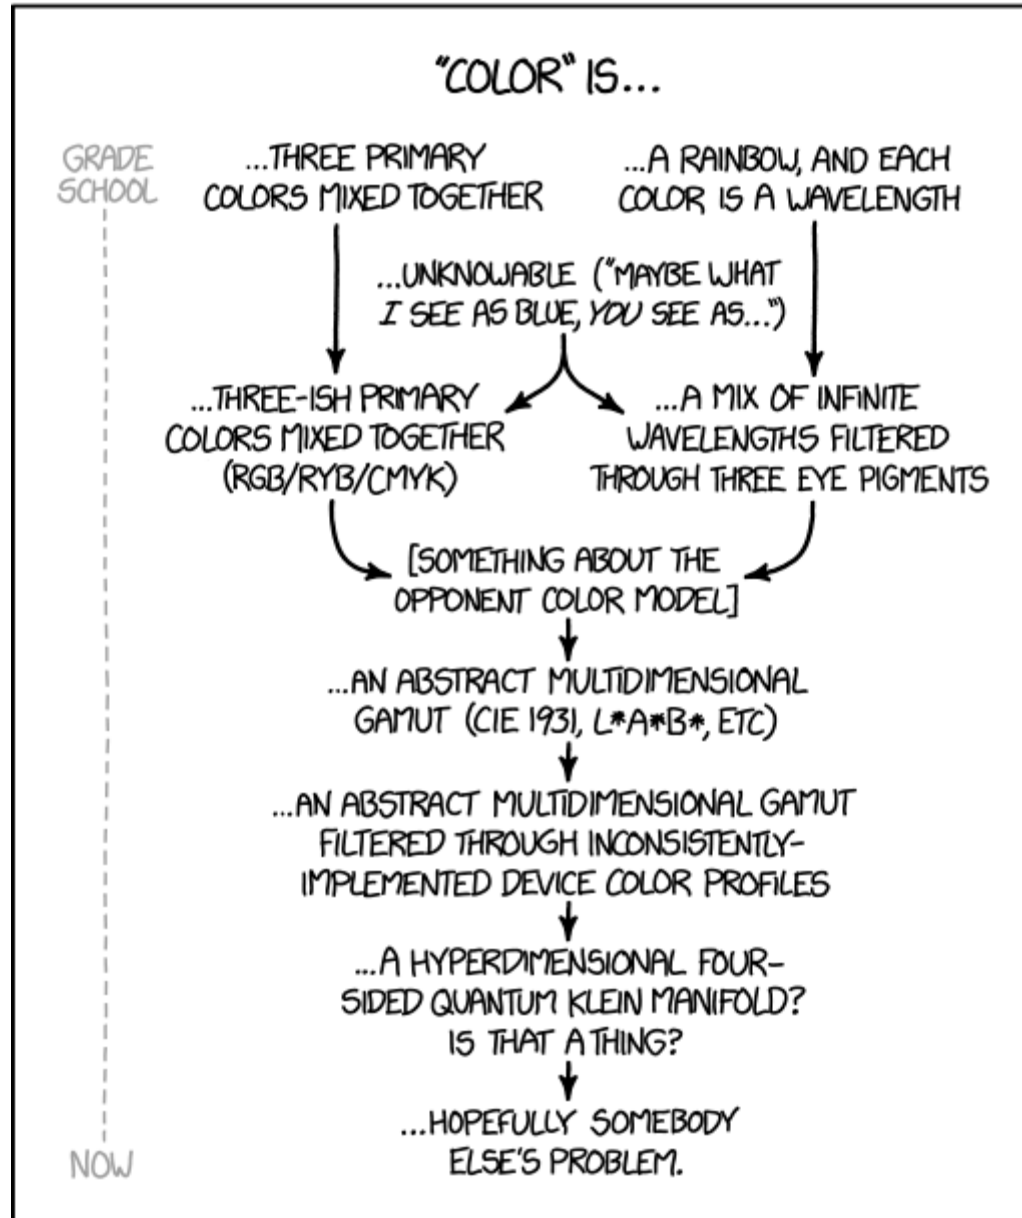


a
($L=65, b=0$)




b
($L=65, a=0$)

EVOLUTION OF MY UNDERSTANDING OF COLOR OVER TIME:



More references

- <https://www.colorsystem.com/>
- A description of many different color systems developed through history.
- Navigate from the right-hand links.
- Thanks to Alex Nibley!



57 37 17 -3 -23 -43 -63 -83 -103



Homestead

Key Largo

Key West




Hurricane Irma - VIIRS - I-Band 5 - 11µm
10 September 2017, 3:40 am EDT





Rainbow color map considered harmful

Borland and Taylor



“Intuitive” color space?



Wait a minute...

**WHY DOES COLOR LOOK LIKE IT
MAPS SMOOTHLY TO A CIRCLE?**

'Color' != position on EM spectrum


Our cells induce color perception by interpreting spectra.

Most mammals are dichromats:

- Lack 'L' cone; cannot distinguish green-red
- 1% of men (protanopia color blindness)

Trichromaticity *evolved*.

No implicit reason for effect of extra cone to be linear.



'Color' != position on EM spectrum

Many different ways to parameterize color.

Ask Prof. Thomas Serre for a qualified answer.

Or...

“When some primates started growing a *third* cone in their retinas, the old bipolar system remained, with the third cone adding a 2nd dimension of color encoding: red versus green. since color is now encoded in a 2d space, you find that you can draw a circle of colors in that space, which when you think about the fact that wavelength is 1d is really weird.”

- aggasalk, Reddit.

Held and Hein (1963)




FIG. 1. Apparatus for equating motion and consequent visual feedback for an actively moving (A) and a passively moved (P) S.