

Color



Prof. Dr. Elli Angelopoulou

Pattern Recognition Lab (Computer Science 5)

University of Erlangen-Nuremberg



What is Color?



What is Color?



- We use the same word to denote 2 different things.
- Consider for example the color of an object.
- We can use the word color to refer to the color we *perceive* an object has.
 - Subjective (influence by culture, individual differences on people's optic system)
 - Vague descriptors (e.g. fuschia, mauve, lilac... How does a mauve object differ from a lilac one?)
- We can use the word color to describe the part of the visible light that is not absorbed by the object.

Outline




- The importance of color (other presentation)
- Color in physics
- Trichromatic color
- Color perception
- Color capture
- Applications
 - Specular Highlights
 - Color Constancy

Physics of Color




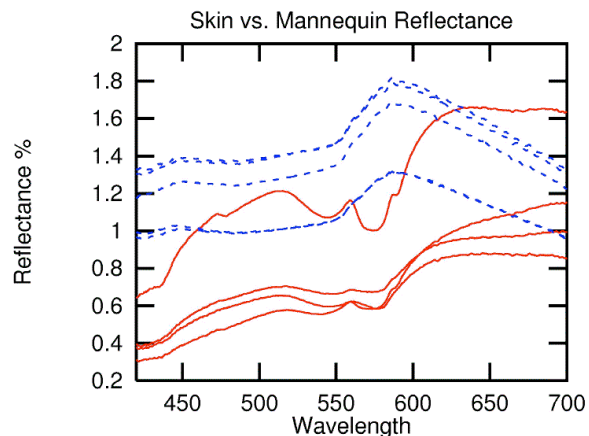
- Color is a property related to the wavelength of the EM spectrum. It is only applicable over the visible range.



Spectral Distribution




- Most objects emit/reflect light at a collection of wavelengths.
- At different wavelengths they emit/reflect a different “amount” of light.
- If we know the amount of emitted/reflected light for each wavelength, then we know the spectral distribution (or spectrum) of an object.
- Lasers by definition are single wavelength light sources (i.e. for HeNe lasers 632.8nm)





Sample Flower Spectra



Measurements by E. Koivisto.

Spectral Irradiance



- Irradiance: power of light falling on a surface patch

$$E = \frac{dP}{dA}$$

- Measured in W/m^2
- Spectral Irradiance: power of light falling on a surface patch per unit wavelength

$$E^\lambda = \frac{d^2P}{dAd\lambda}$$

- Measured in W/m^3

Spectral Radiance



- Radiance: power of light falling on a surface patch from a specific direction

$$L = \frac{d^2P}{d\omega dA \cos\vartheta}$$

- Measured in $\text{W}/\text{sr} \cdot \text{m}^2$
- Spectral Radiance: power of light falling on (emitted from) a surface patch from a specific direction for a unit wavelength

$$L^\lambda = \frac{d^3P}{d\omega dA \cos\vartheta d\lambda}$$

- Measured in $\text{W}/\text{sr} \cdot \text{m}^3$

Light Sources



- The biggest body of work in computer vision focuses on surfaces that reflect light.
- In order to observe such a surface, there must be light falling on the surface.
- The same object can produce different spectral (ir)radiances depending on the spectrum of the incident light.






The Color of Light Sources

- In order to understand/analyze the color of an object we need to know the color of the incident illumination.
- Different types of light sources produce different illumination spectra.
- One can talk about
 - Indoor Illumination
 - Outdoor Illumination
- Under indoor illumination we can have different types of light sources:
 - Incandescent lamps
 - Fluorescent lamps
 - Arc lamps
- There is an overlap between indoor and outdoor illumination. A large number of these lights are **black body radiators**.



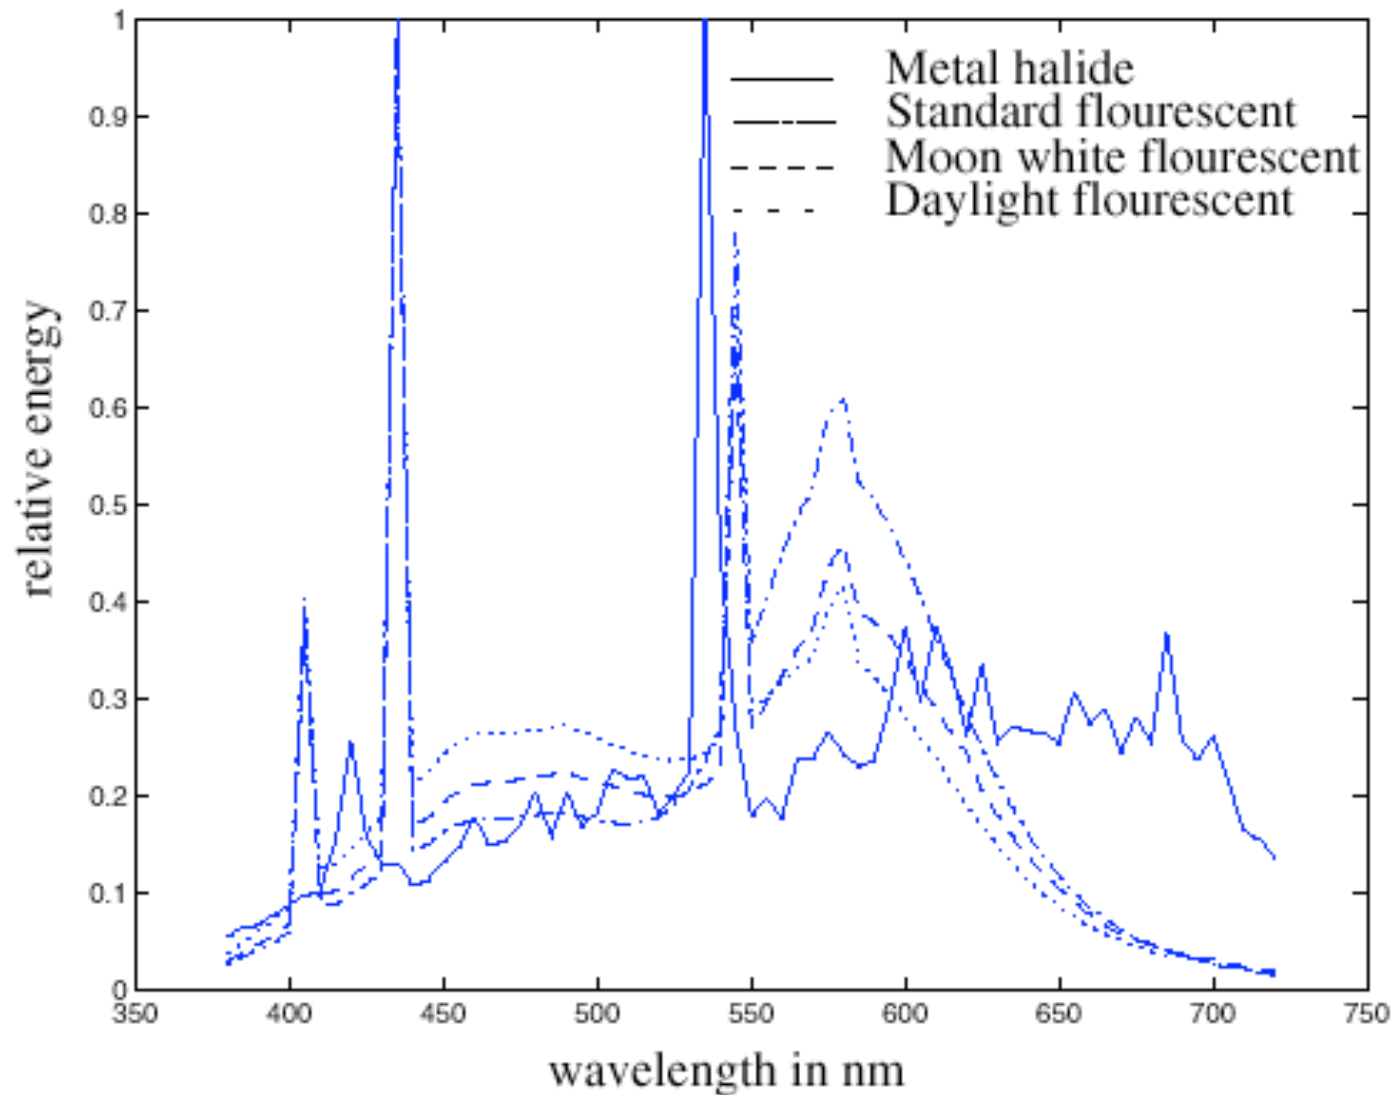
Illuminant Spectra



Measurements by J. Parkkinen and P. Silfsten.



Spectra of Fluorescent Light




Measurements by H. Sugiura.

Black Body Radiators



- Black body: A body (object) that reflects no light.
- If we heat a black body (e.g. the sun, the filament in an incandescent light bulb) it will start emitting EM radiation.




Spectrum of Black Body Radiators



- The spectrum of a black body radiator (BBR) depends only on the temperature of the black body.
- Planck's law closely approximates the spectrum of a BBR, that is why BBRs are often also called Planckian Illuminants.

$$I(\lambda, T) = \frac{2hc}{\lambda^5} \frac{1}{e^{hc/\lambda kT} - 1}$$

h: Planck's constant
 c: speed of light in vacuum
 k: Boltzmann's constant

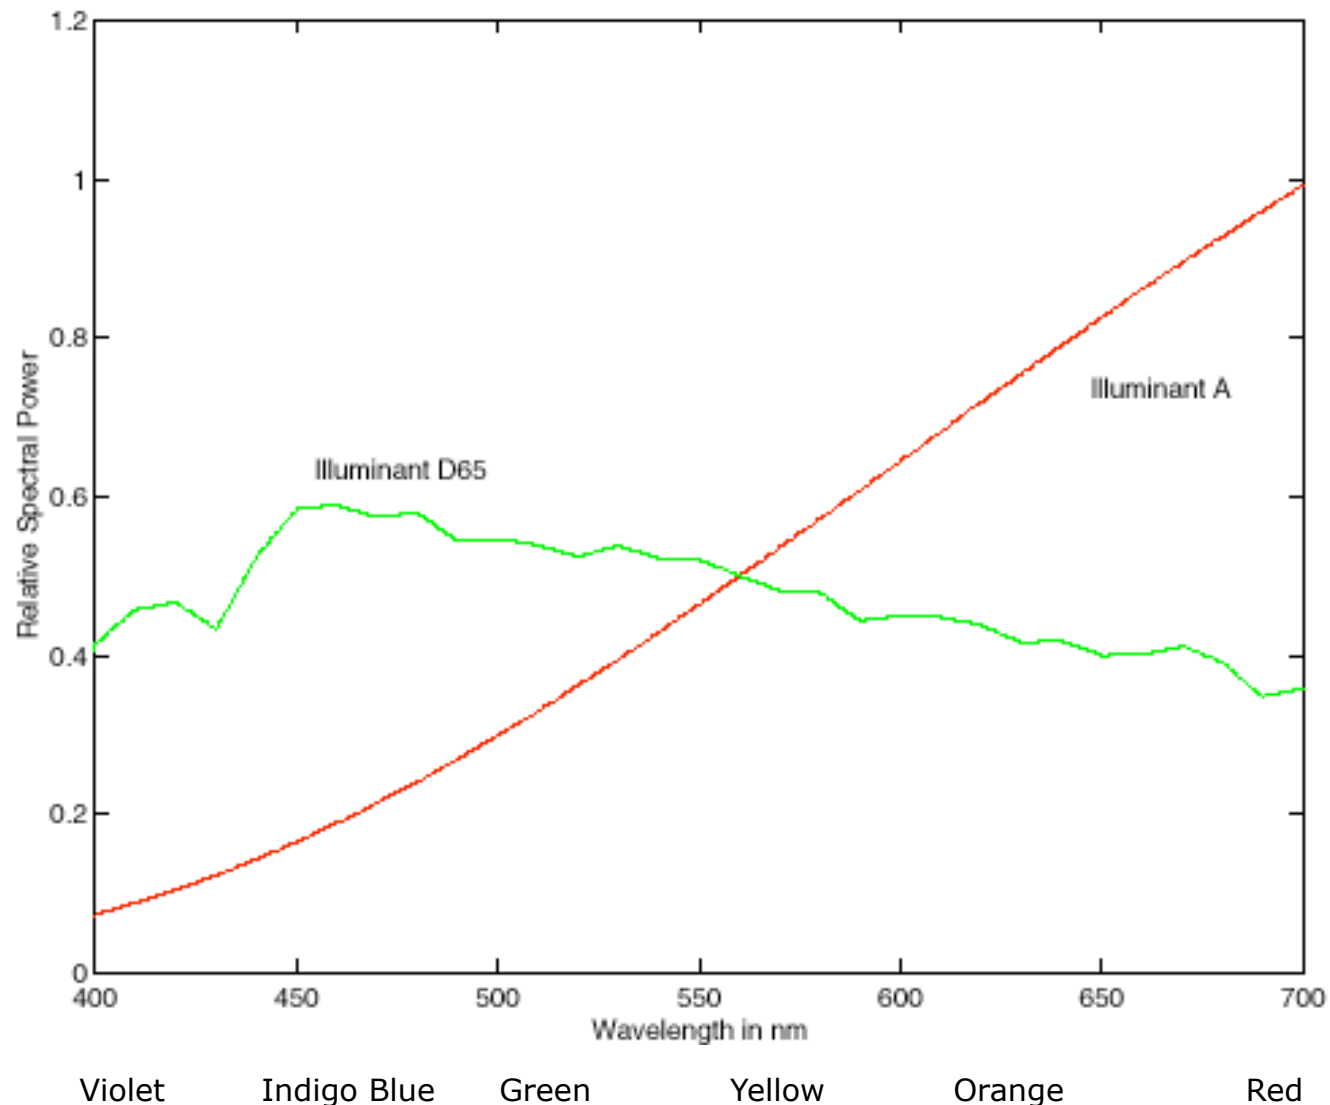


BBRs and Real Illuminants



- Incandescent lamps are black body radiators.
- The sun is a black body radiator. A BBR at 5780°K closely approximates the sun spectrum.
- Though the earth's atmosphere scatters short wavelengths more than the longer ones, we still treat outdoor light as a black-body radiator. For the sky we use a BBR of higher temperatures:
 - Overcast sky, BBR at 7000°K
 - Clear blue sky, BBR at $10,000^{\circ}\text{K}$
- Fluorescent lamps are *not* black body radiators.

Standard Illuminant Spectra



Relative spectral power of two standard illuminant models. D65 models sunlight, and illuminant A models incandescent lamps. Relative spectral power is plotted against wavelength in nm. The visible range is about 400nm to 700nm.



The Color of Surfaces

- A large variety of mechanisms affect the color of surfaces.
 - Reflection
 - Refraction
 - Diffraction
 - Scattering
- The physics community looks at light interaction at the *microscopic level* and has complex models that can describe a variety of light and surface interactions.
- The computer vision community bundles all these effects into *macroscopic models* which describe how the color will change as the geometry and the illumination changes.
- Example reflectance model for diffuse+specular reflection:

$$L(P, \vartheta_o, \phi_o, \lambda) = \rho_d(P, \lambda) \int_{\Omega} L(P, \vartheta_i, \phi_i, \lambda) \cos \vartheta_i(P) d\omega + \\ + \rho_s(P, \lambda) L(P, \vartheta_s, \phi_s, \lambda) \cos^n(\vartheta_s(P) - \vartheta_o(P))$$



Trichromacy

- In Computer Vision and Computer Graphics we could do all the color computations using a continuous color space (full spectrum).
- However, experiments have shown that for most people a combination of three basis colors, formally known as *three primary colors*, is sufficient to represent the entire color space that we can sense.
- So most algorithms in CV and CG operate in *trichromatic* space.
 - Space and time efficiency
 - Early work inspired by CG and sensors which have a human observer.



Trichromacy Experiment

- Observer faces a monitor with a black background. A test light is presented in one half. The observer is asked to adjust a mixture of lights in the other half until the 2 lights are perceived to have the same color.
- For most observers adjusting three primary colors was sufficient in order to match the two colored lights.
- In other words, any light T can be described as:

$$T = w_1 P_1 + w_2 P_2 + w_3 P_3$$

where $w_1, w_2, w_3 \geq 0$

More Trichromacy Results



$$T = w_1 P_1 + w_2 P_2 + w_3 P_3$$

- Almost any perceived color can be expressed as a linear combination of three primary colors.
- For the three colors to be primary colors they have to be:
 - independent
 - span the space of perceived color.
- Most observers select the same mixture of primaries (same w_i 's)

Grassmann's Laws



- In 1853 H.G. Grassmann developed a theory on color mixing that became known as Grassmann's Laws.

“If two simple but non-complementary spectral colors be mixed with each other, they give rise to the color sensation which may be represented by a color in the spectrum lying between both and mixed with a certain quantity of white.”

- Consider 2 colored lights T_a and T_b :

$$T_a = w_{a_1} P_1 + w_{a_2} P_2 + w_{a_3} P_3$$

$$T_b = w_{b_1} P_1 + w_{b_2} P_2 + w_{b_3} P_3$$

Grassmann's Laws (continued)



1. Mixing the lights = Mixing the weights (matches)

$$T_a + T_b = (w_{a_1} + w_{b_1})P_1 + (w_{a_2} + w_{b_2})P_2 + (w_{a_3} + w_{b_3})P_3$$

2. If two lights are matched by using the same weights (matches) then they must be the same.

$$w_{a_i} = w_{b_i}, \forall i \Rightarrow T_a = T_b$$

3. Matching is linear:

$$kT_b = kw_{b_1}P_1 + kw_{b_2}P_2 + kw_{b_3}P_3$$

Human Physiology



- For most people there are 4 distinct receptors, 3 cones and 1 rod, on the retina.
- Each cone is sensitive to a different part of the visible spectrum, roughly blue, red and green.
- Red cones "fire" when red light falls on them, etc.
- Rods "fire" in low light and fire independent of the wavelength (as long as it is in the visible range)







Fig. 1.1. A drawing of a section through the human eye with a schematic enlargement of the retina.

Cones



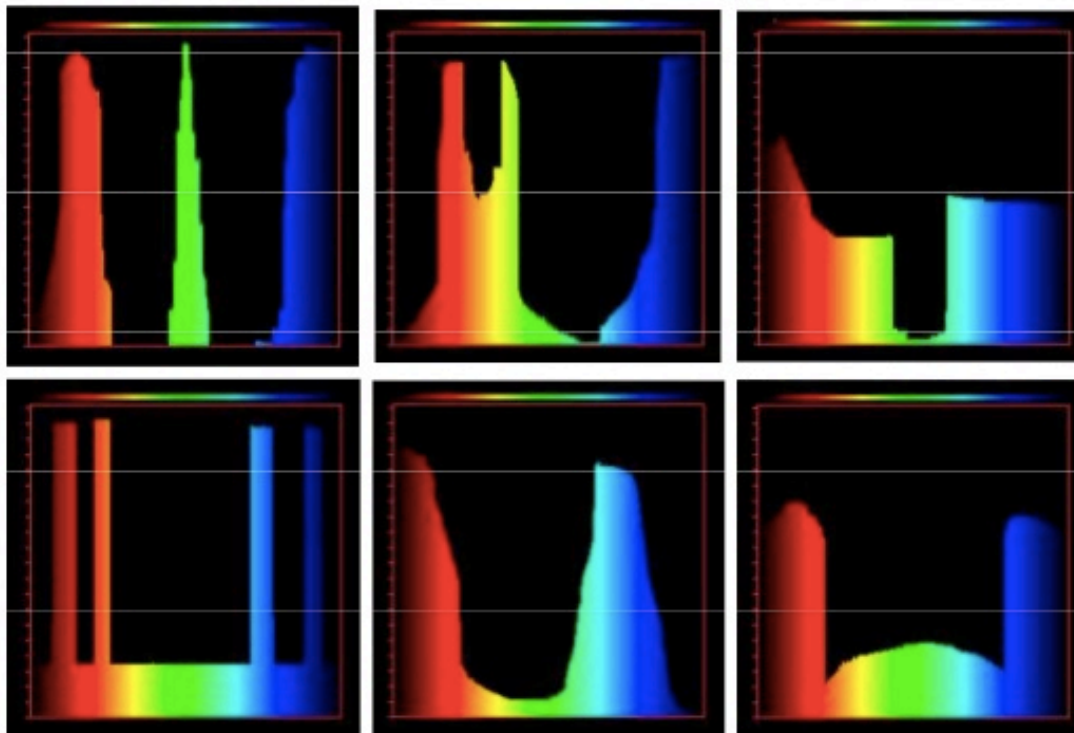
- The distribution of cones and rods varies in people.
- The response (sensitivity) to different wavelengths varies among people.
- Three types of cones:
 - L cones: respond most to light of long wavelengths, typically peak at red, near 564–580nm.
 - M cones: respond most to light of medium wavelengths, typically peak at green-yellow, near 534–545nm.
 - S cones: respond most to light of short wavelengths, typically peak at blue, near 420–440nm.




Metamerism



- The reduction of full spectral information down to 3 values, introduces a significant loss of information.
- Metamerism: the phenomenon where distinct spectra produce the same trichromatic response.



These six spectra look the same purple to normal color-vision people.*




* "We are ALL colorblind",
<http://www.visualmill.com/>

10

Standardized Trichromatic Color



- CIE Commission International d'Eclairage was established in order to develop color standards (similar to IEEE for electronic engineering).
- First goal was to standardize the three primaries.
- CIE XYZ primaries:
 - $x(\lambda)$: small peak at 440nm, large peak at 590nm, local minimum at 490nm, sensitive from 380 to 700nm.
 - $y(\lambda)$: single peak at 548nm, sensitive from 400 to 700nm.
 - $z(\lambda)$: single peak at 430nm, sensitive from 380 to 550nm.



Spectrum to CIE XYZ.




- The CIE xyz primaries are formally known as the *CIE color matching functions*.
- Given a spectrum $I(\lambda)$, one can compute the CIE X, Y, Z *tristimulus values* as follows:

$$X = \int_0^{\infty} I(\lambda)x(\lambda)d\lambda$$

$$Y = \int_0^{\infty} I(\lambda)y(\lambda)d\lambda$$

$$Z = \int_0^{\infty} I(\lambda)z(\lambda)d\lambda$$

RGB Color Matching Functions




RGB: primaries are monochromatic, energies are 645.2nm, 526.3nm, 444.4nm.


Color matching functions have negative parts -> some colors can be matched only subtractively.



CIE XYZ Color Space



CIE Color Space




A qualitative rendering of the CIE (x,y) space.

The horseshoe-like region represents visible colors.

Each slice, typically shows a constant brightness section of the color space.

There are sets of (x, y) coordinates that don't represent real colors, because the primaries are not real lights (so that the color matching functions could be positive everywhere).

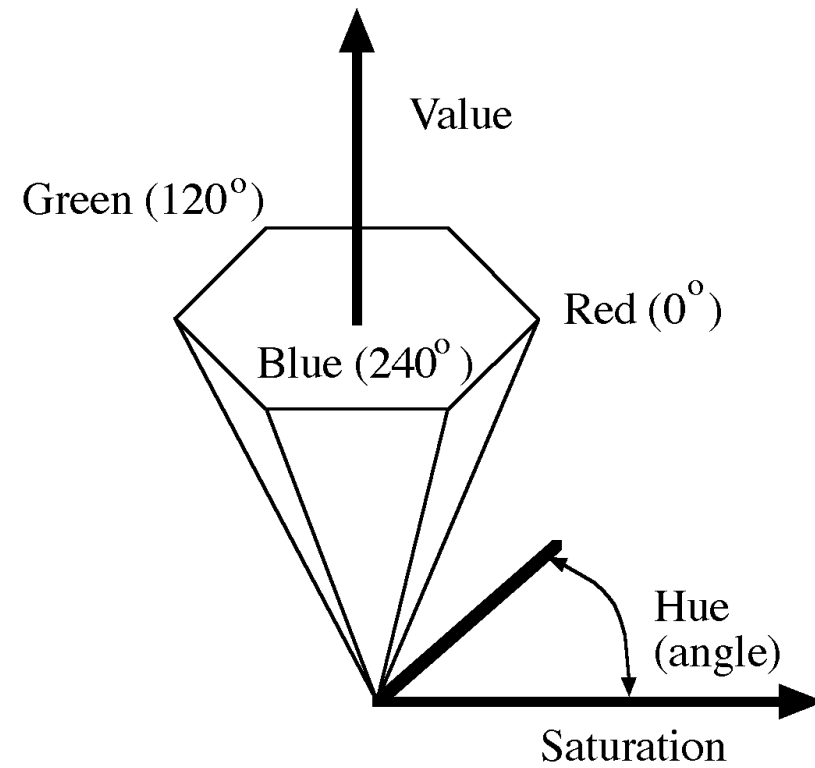

Illuminants in CIE Color Space



A plot of the CIE (x,y) space. The spectral locus (the colors of monochromatic lights) and the black-body locus (the colors of heated black-bodies) is shown, as well as the range of typical incandescent lighting.





HSV Hexcone







CIE XYZ vs. RGB cube





Color Matching in CIE x,y Space



At the center of the ellipse is the color of a test light. The size of the ellipse represents the scatter of lights that the human observers matched to the test color. The boundary of the ellipses corresponds to just noticeable differences.

The ellipses on the left have been magnified 10x for clarity. On the right they are plotted to scale.

The ellipses are known as MacAdam ellipses after their inventor.

Elli Angelopoulou


Color Matching in CIE u',v' Space



CIE u', v' is a projective transformation of the CIE x,y space.

We transform (x,y) points so that ellipses are more similar.

Distance metric in CIE u',v' space are more uniform.




Color Cameras- review



- Most color cameras give a triplet of color values per pixel (R,G,B).
- Either a separate chip is used per color, or a filter composed of a mosaic of smaller individual color filters is laid over the CCD chip.


Bayer filter

50% G, 25% R, 25% B




RGBE filter

equal distribution



3 CCD chip



Images courtesy of Wikipedia <http://en.wikipedia.org>

Image courtesy of Canon
<http://www.usa.canon.com/tro>

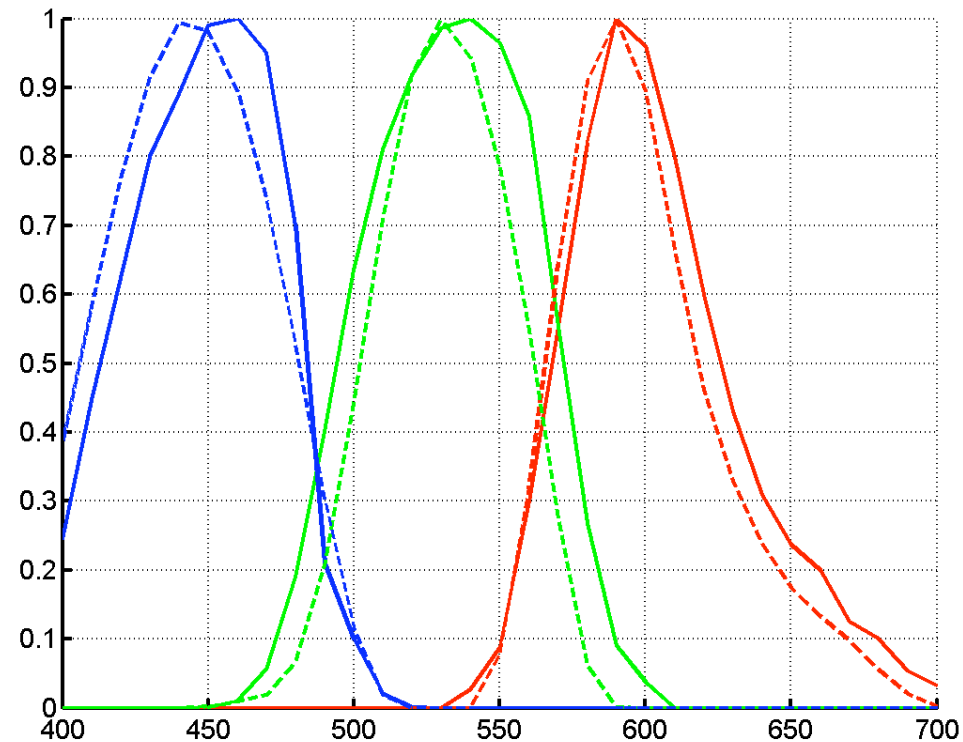


Incident Irradiance on Sensor

- Under orthographic projection, the amount of light (incident irradiance) arriving at a pixel p for a particular color filter k (e.g. $k=R,G,B$) depends on:
 1. The spectral irradiance $E(\lambda)$ of the light falling on the surface.
 2. The spectral reflectance $S(\lambda, \vartheta_i, \phi_i, \vartheta_r, \phi_r)$ of the surface patch projected on the pixel p (e.g. Lambertian, or Specular or a mixture).
 3. The spectral response $\sigma_k(\lambda)$ of the k^{th} color filter of the sensor.
- Thus, the value at pixel p for color filter k is:

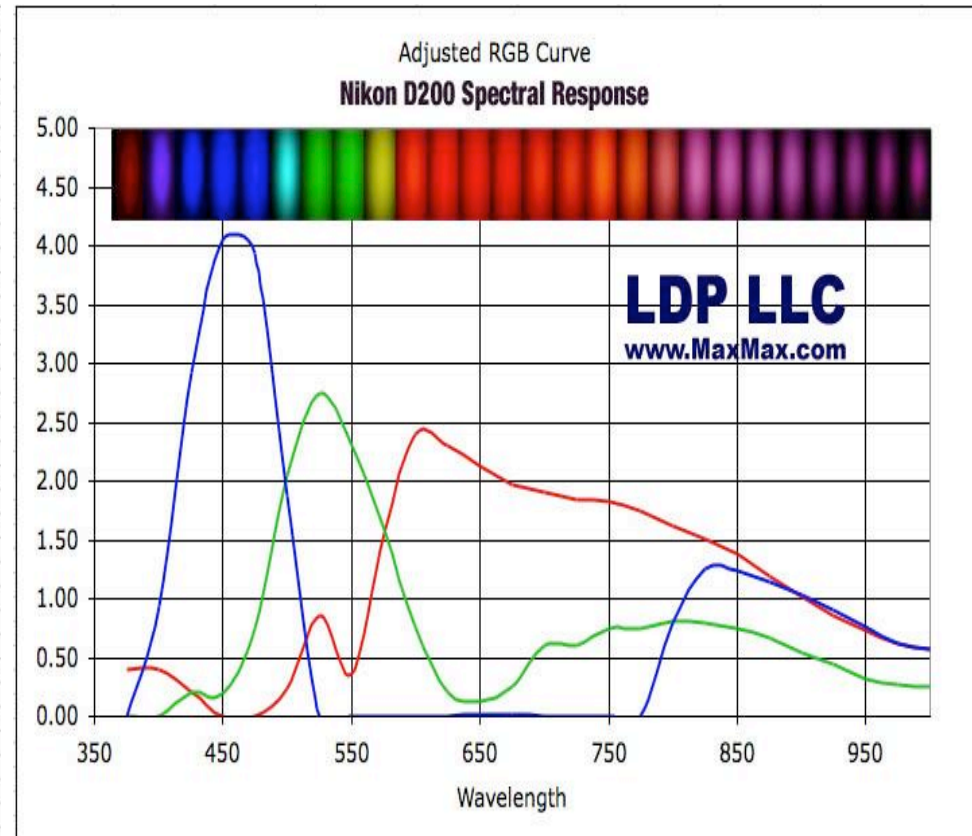

$$p_k = \int_{\lambda} \sigma_k(\lambda) \underbrace{S(\lambda, \vartheta_i, \phi_i, \vartheta_r, \phi_r)}_{I(\lambda)} E(\lambda) d\lambda$$

Camera Spectral Sensitivities





- solid line: 3CCD Sony DXC-755P (as published by Oulu University)
- dashed line: 3CCD Sony DXC-930 (as published by K. Barnard)

Camera Spectral Sensitivities 2



- Canon 40D and Nikon D200. Both Bayer pattern sensors, Canon 40D uses a CMOS chip, while Nikon D200 uses a CCD chip. (images courtesy of http://www.maxmax.com/spectral_response.htm)

Camera Spectral Sensitivities 3








- Nikon D700 vs. Nikon D200. Both Bayer pattern sensors, Nikon D700 uses a CMOS chip, while Nikon D200 uses a CCD chip. (images courtesy of http://www.maxmax.com/spectral_response.htm)

Foveon CMOS color sensor



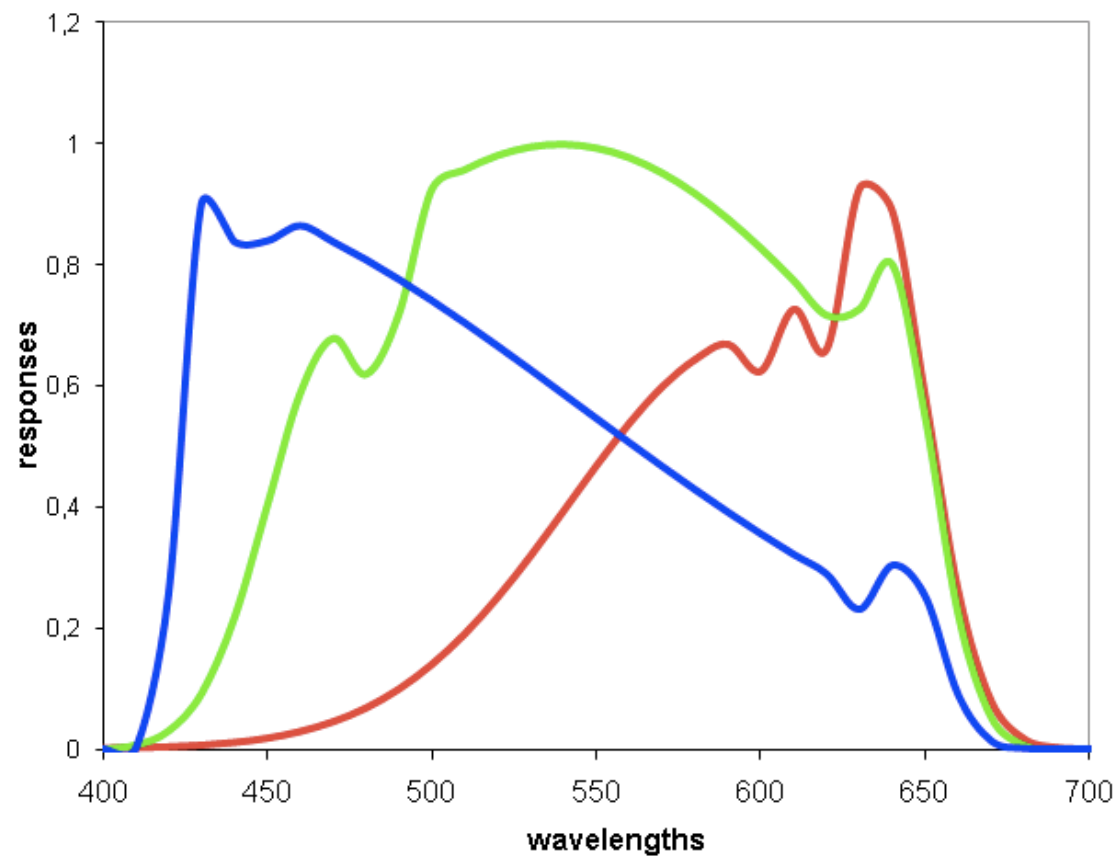
- Newest color technology. The Foveon X3 chip can currently be found on cameras by Sigma Corp.



Elli Angelopoulou

Color

Foveon Camera Spectral Sensitivities



- Spectral sensitivities of Sigma Foveon D10, which uses the foveon x3 chip. (Image courtesy of J. Vazquez i Corral, Universitat Autònoma de Barcelona.)

Simplified Model for Pixel Response



- Recall that the value at pixel p for color filter k is

$$p_k = \int_{\lambda} \sigma_k(\lambda) S(\lambda, \vartheta_i, \phi_i, \vartheta_r, \phi_r) E(\lambda) d\lambda$$

- The spectral response $\sigma_k(\lambda)$ is often unknown.
- The majority of the surfaces exhibit a mixture of specular and diffuse reflectance.
- Thus a popular model for the value at a camera pixel \mathbf{x} is:

$$p_k(\mathbf{x}) = g_d(\mathbf{x})d_k(\mathbf{x}) + g_s(\mathbf{x})s_k(\mathbf{x})$$

where

d_k is the *image* value for the k^{th} color filter of the *diffuse* reflection of an equivalent flat frontal surface viewed under the same light.

g_d is a geometric term that captures the variation in brightness caused by changes in the surface orientation.

s_k is the *image* value for the k^{th} color filter of the *specular* reflection of an equivalent flat frontal surface viewed under the same light.

g_s is a geometric term that captures the variation in the amount of energy that is specularly reflected.

Example Application: Specularity Detection



- We have many algorithms (stereo, tracking, shape recovery) that assume purely diffuse reflection.
- Other algorithms (estimation of light source position) assume specular reflection.
- Most real surfaces exhibit a mixture of diffuse and specular reflectance.





Specularity Detection- continued




- How should we treat images which include diffuse, specular and diffuse+specular pixels?
- Identify/Separate specularities.



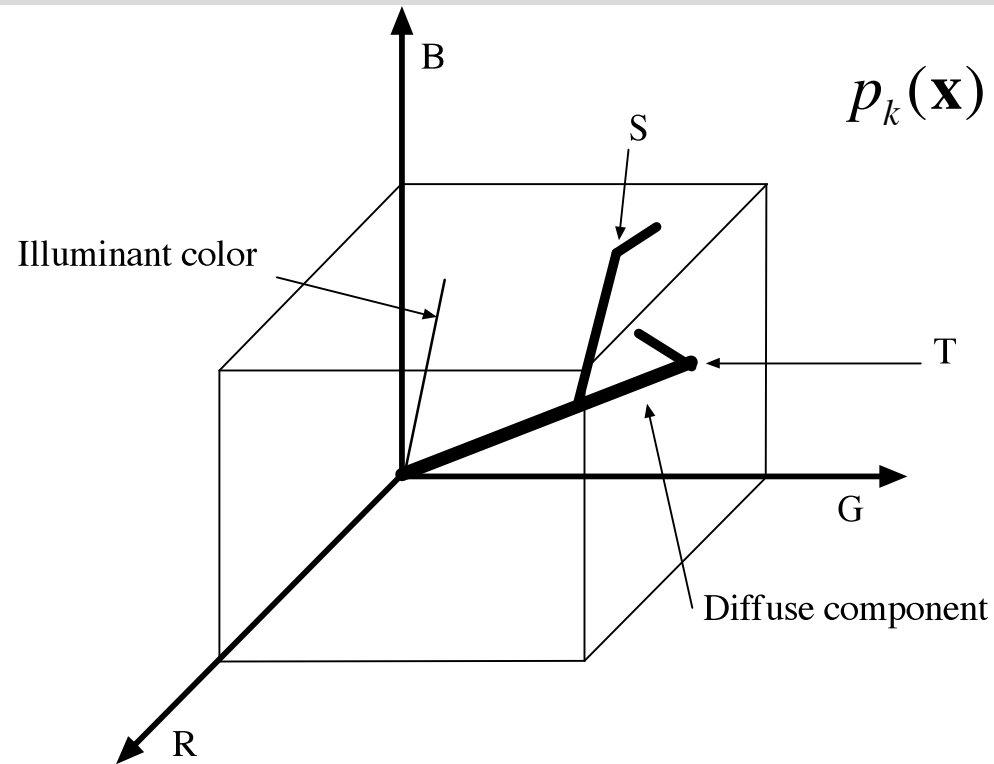


Specularity Detection - continued


- How should we treat images which include diffuse, specular and diffuse+specular pixels?
- Identify/Separate specularities.



Finding Specularities




$$p_k(\mathbf{x}) = g_d(\mathbf{x})d_k(\mathbf{x}) + g_s(\mathbf{x})s_k(\mathbf{x})$$




- Consider a picture of a single-colored object.
- If we plot each pixel in RGB space, we should obtain a distribution of RGB points (a gamut) that looks like the plot on the left.
- The term $g_d(\mathbf{x})d_k(\mathbf{x})$ produces a line T that should pass through the origin.

Finding Specularities




$$p_k(\mathbf{x}) = g_d(\mathbf{x})d_k(\mathbf{x}) + g_s(\mathbf{x})s_k(\mathbf{x})$$



- If there is a specularity, a 2nd line, S , is formed, which is caused by the term $g_s(\mathbf{x})s_k(\mathbf{x})$. It is parallel to the illuminant color.
- The specular line, S , is parallel to the illuminant color.
- It is an offshoot of the diffuse line, T , because in many pixels we have a combination of specular and diffuse.



Finding Specularities



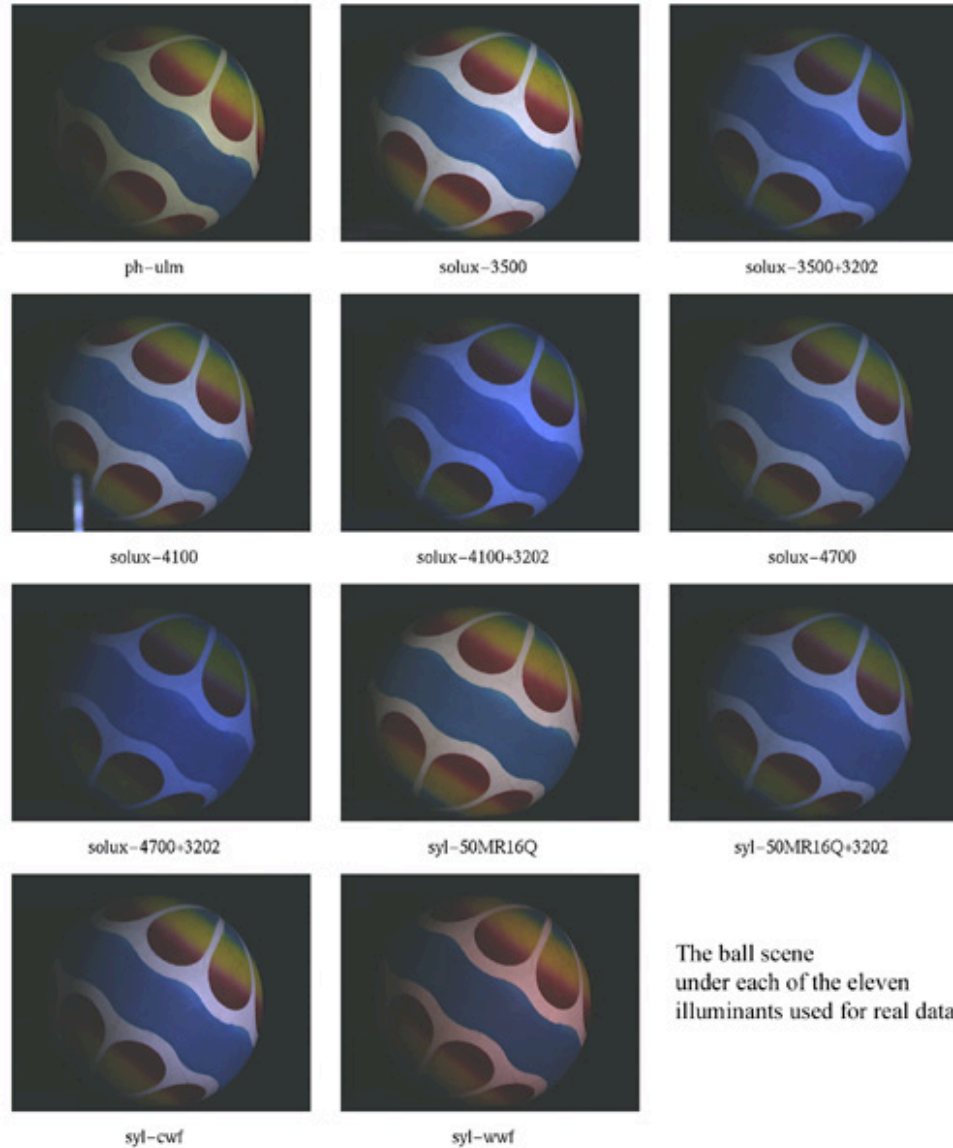
- All pixels on S are specular pixels.
- We want to identify and exclude such specular pixels if our algorithm assumes diffuse reflectance.

On Specularity Detection

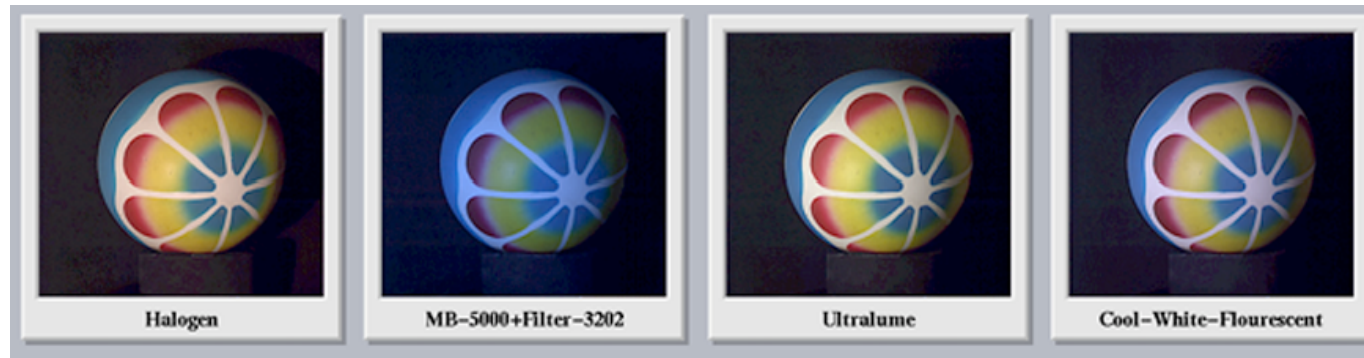


- There exist a large number of algorithms for specularity detection/elimination that analyze the distribution of pixels in color space and try to identify the lines S and T .
- Complexities:
 - Unknown illumination color
 - The model is too simple and thus does not capture all the complexities of light interaction. => We do not get nice straight lines.
 - Multi-colored objects
 - Perspective projection
 - Non-linear camera response (camera-gamma)

Example Application 2: Color Constancy




Color Constancy



- Color constancy is the term we use to describe the mechanism that humans have, which allows them to extract the spectral content (color) of a scene relatively independent of the spectral content (color) of the illumination of the scene.

Humans "See" the Color Difference




Camera response varies with the illuminant color.

This figure shows a **uniform reflectance** illuminated by five different lights, and the result plotted on CIE x,y



Humans "See" the Color Difference




The same object can generate different color responses due to different illuminations.

A **blue flower** illuminated by five different lights, and the result plotted on CIE x,y . Notice how it looks significantly more saturated under some lights.



Humans "See" the Color Difference



The same object can generate different color responses due to different illuminations.

A **green leaf** illuminated by five different lights, and the result plotted on CIE x,y. Notice how it looks significantly more saturated under some lights.

Color Constancy Algorithms



- The variation in captured color as the illumination changes affects recognition and tracking.
- Color constancy algorithms aim at cancelling out the effects of illumination variation.
- They typically involve:
 1. Estimation of illuminant color
 2. Mapping to a canonical (standardized) color
- Step 1 is sometimes omitted.
- It often involves, like specularities detection, analysis of pixel distributions in color space.
- Color constancy is related to white-balancing.

Grey World Algorithm



- One of the simplest methods for estimating the color of the incident illumination.
- Key assumption: The average color of a scene is grey.
- Any deviation from the expected average grey must be caused by the color of the illuminant.
- Let $\mathbf{p}^k = (p_R^k, p_G^k, p_B^k)^T$ be the observed color at pixel k .
- Then according to the grey-world assumption, the illuminant color $\mathbf{i} = (i_R, i_G, i_B)^T$ in an image I , $|I| = s$, can be estimated as follows:

$$i_c = \frac{1}{s} \frac{\sum_{k \in I} p_c^k}{\left\| \sum_{k \in I} p_c^k \right\|_2}$$

Grey Edge Algorithm



- Extension of the Grey World to spatial derivatives.
- Key assumption: The average color *difference* of a scene is grey.
- Differentiation increases the robustness against noise.
- Higher order derivatives mimic how the human eye handles color constancy.

$$i_c = \sqrt[m]{\sum_{k \in I} \left| \frac{\partial^n p_{c, \sigma_G}^k}{\partial \mathbf{u}^n} \right|^m}$$

where $p_{c, \sigma_G} = p_c \otimes G_{c, \sigma_G}$ is the color channel c smoothed with a Gaussian of standard deviation σ_G and $\mathbf{u} = (u_1, u_2)$ is the vector of the directions of differentiation.

- For $n = 0$ and $m = 1$ we obtain the Grey World (up to a scalar).

White-Balancing: Laboratory Images



Original image



White-balanced image using the ground truth



White-balanced image using the grey world estimates

White-Balancing: Real World Images



Original image



White-balanced image using the ground truth




White-balanced image using the grey world estimates


Ground Truth for Two Illuminants





=




+



Per-illuminant influence in the two-illuminant image



Electromagnetic Spectrum & Trichromacy





With traditional RGB cameras a continuum of EM radiation values is converted to 3 discrete values.

Beyond Tri-chromatic Sampling




- **Solution:** Capture a denser, finer spectral sampling of the visible light. Complex light reflectance interactions affect.
- Image a scene under a multitude of different **narrow** color filters.



- Example: 10 filters each 10nm wide (compared to the traditional 75nm wide RGB filters).

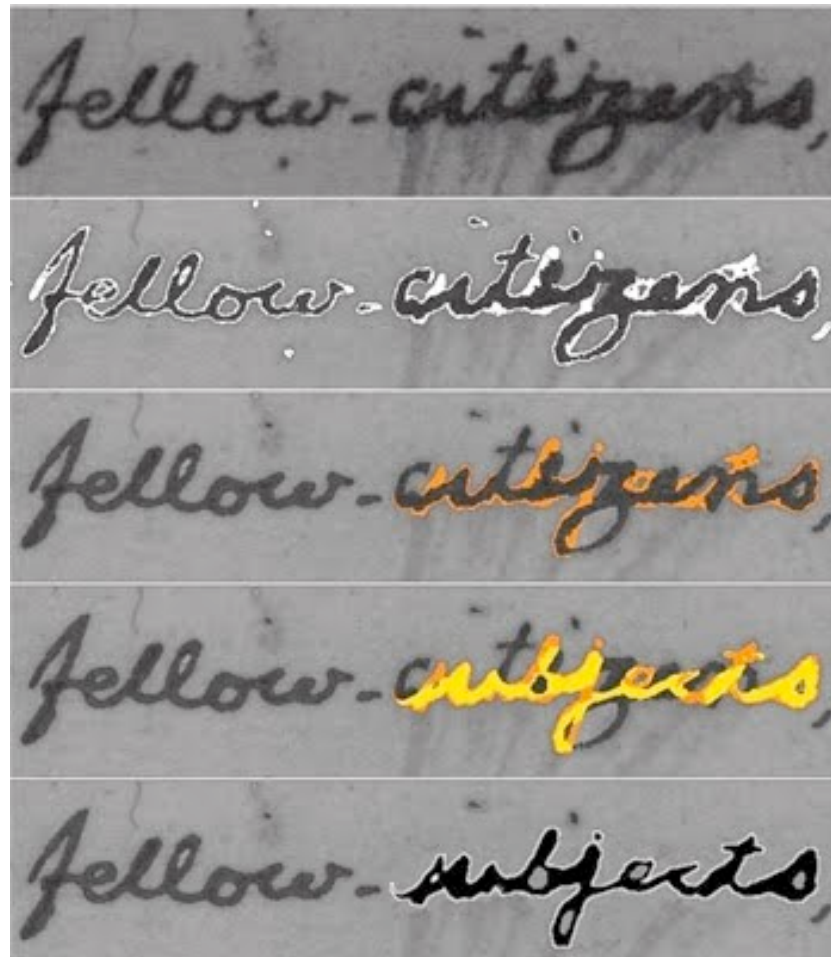


Illustration of a Multispectral "Image"





A Teaser: Declaration of Independence.



Multispectral analysis of Jefferson's rough draft of the Declaration of Independence.

Image Sources



1. The laser spectrum is courtesy of http://www.antonine-education.co.uk/physics_gcse/Unit_1/Topic_5/em_spectrum.jpg
2. The plot of the black body radiators is courtesy of http://www.ucolick.org/~bolte/AY4_00/week2/blackbodies.html
3. The maple leaves rendered under different colors are courtesy of http://static.creativecrash.com/tutorialimages/352/light_color_dihe_img_6.png
4. The retina drawing is courtesy of <http://webvision.umh.es/webvision/imageswv/Saqschem.jpeg>
5. The SML cones plot is courtesy of <http://www.handprint.com/HP/WCL/IMG/conesens3.gif>
6. The RGB cones plot is courtesy of <http://www.physicsclassroom.com/class/light/u12l2b2.gif>
7. The middle RGB cone plot is courtesy of Cvonline http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/OWENS/LECT14/cones.gif
8. The CIE XYZ curves and the CIE space are courtesy of <http://escience.anu.edu.au/lecture/cg/Color/Image/>
9. The 3D CIE XYZ color space is courtesy of C. Ulbricht and A. Willie <http://www.cg.tuwien.ac.at/research/publications/2006/ulbricht-2006apw/image-orig.jpg>
10. The colorful CIE XYZ slice is from http://www.knowledgerush.com/wiki_image/4/40/Cie_chromaticity_diagram_wavelength.png
11. The colorful MacAdam ellipses are courtesy of wikipedia,
12. The CIE plots are from the slides by D.A. Forsyth, University of California at Urbana-Champaign.
13. The malachite pictures were posted by panoguy at <http://www.pentaxforums.com/forums/pentax-q/219163-humorous-comparison-q-vs-sigma-dp2-merrill.html>
14. The Foveon colorful picture is titled "Buffalo Gap, Texas", copyright 2004 Barry Bird.
15. The Foveon leaves picture is titled "Wasatch Mtns., Utah", copyright 2004 Mike Morrison.
16. The example images of specular and diffuse reflection are courtesy of S. Mallick, <http://graphics.ucsd.edu/~somallick/research/suv/index.html>
17. The color constancy database is courtesy of K. Barnard http://www.cs.sfu.ca/~colour/data/objects_under_different_lights/index.html
18. The 4-ball color constancy example is courtesy of P. Gabbur and K. Barnard http://www2.engr.arizona.edu/~pgsangam/IBM_Almaden.ppt
19. Image of Declaration of Independence, <http://philosophyofscienceportal.blogspot.com/2010/07/thomas-jeffersons-faux-pas.html>