TIEA311 Tietokonegrafiikan perusteet
kevät 2018

("Principles of Computer Graphics" – Spring 2018)

TIEA311 Tietokonegrafiikan perusteet – kevät 2018 (“Principles of Computer Graphics” – Spring 2018)

Adapted from: *Wojciech Matusik*, and *Frédo Durand*: 6.837 Computer Graphics. Fall 2012. Massachusetts Institute of Technology: MIT OpenCourseWare, https://ocw.mit.edu/.

Frontpage of the local course version, held during Spring 2018 at the Faculty of Information technology, University of Jyväskylä:
`http://users.jyu.fi/~nieminen/tgp18/`

TIEA311 - Event horizon of this course

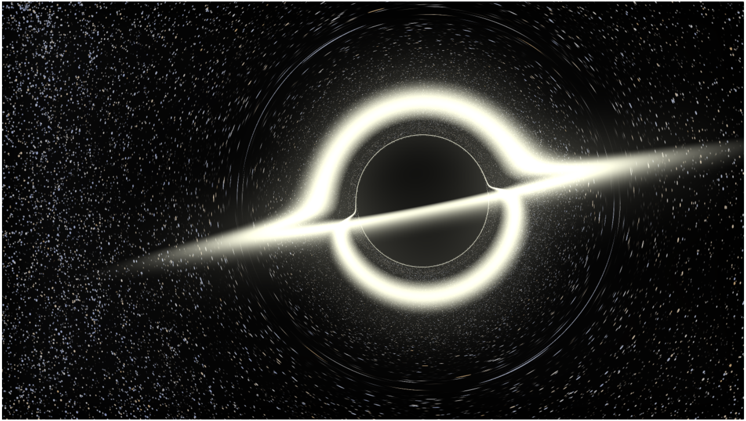

Image and rendering code by Sakari Kapanen. Used with permission. Source:

`https://github.com/flannelhead/blackstar/blob/master/example.png` Code repo:

`https://github.com/flannelhead/blackstar`

TIEA311 - Final week in Jyväskylä

Last week starts!

- Ray Casting (ideas on lecture; practicals in the final Assignment!)
- Ray Tracing issues $\rightarrow$ possible to continue as a "hobby project"; teachers of "TIEA306 Ohjelmointityö" may be contacted regarding credit for (any) hobby projects.
- When **learning**, don't think about deadlines. **Think about learning!**
- My deadlines are flexible, since this course is about graphics and programming and **not about meeting time-to-market**. Leave that to project courses and traineeship periods!

Plan for the final two lectures:

- First, a "debriefing" of Instanssi 2018.
- Color, Shading, texture mapping: Cover the principles up to Phong model (last week) and texture coordinates (today).
- Cherry-pick title slides from advanced stuff that we mostly defer to the follow-up course (starts next week) and/or future self-study.

Instanssi 2018 debrief

- Students from this and last year's courses shined and sparkled in the competitions!
- I wanted to make one new feature for the "courselogo.js" library. Did too.. Copy, paste, modify two lines or so... I have now have scaled geometries, but it's quite a hack :). (But these competitions are a fine way to practice time-to-market and resulting trade-offs:))
- Third place in demo competition achieved by a student of this course!! Congrats!

TIEA311 - Today in Jyväskylä

Super fast-forward!

Today we **rush** through the MIT OCW slides about **color**.

Notice that we'll end up with our "old friend": intensities of red, green, and blue (and "alpha" for transparency). But the following things are worth noticing:

- Color and the human visual processing system is a colorful research topic on its own
- Even as we use RGBA in real-time graphics, we need to know at least something of **why** we do that
- True "hardcore" photorealistic rendering needs more than just RGBA!
- Some of the things touched on the slides have quite interesting connections to our top research in Jyväskylä!!

Color

Wojciech Matusik MIT EECS

Many slides courtesy of Victor Ostromoukhov, Leonard McMillan, Bill Freeman, Fredo Durand

Does color puzzle you?

Answer

- It’s all linear algebra

Plan

- Spectra
- Cones and spectral response
- Color blindness and metamers
- Color matching
- Color spaces

Color

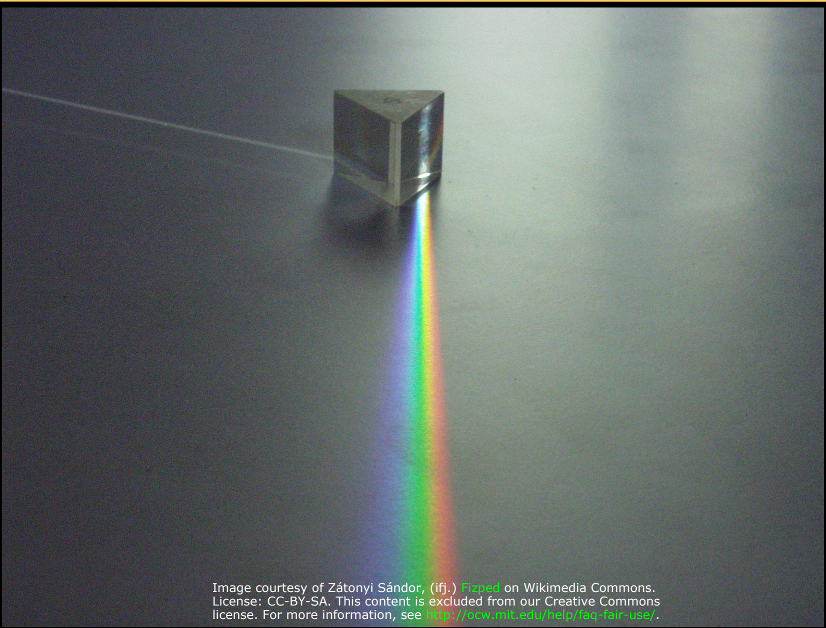

Spectrum

Light is a wave
Visible: between 450 and 700nm

Spectrum

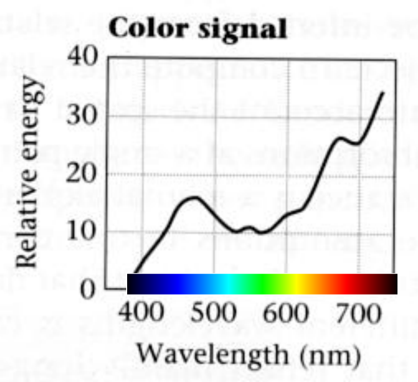

Light is characterized by its spectrum:
the amount of energy at each wavelength
This is a full distribution:
one value per wavelength
(infinite number of values)

Light-Matter Interaction

Where spectra come from:
- light source spectrum
- object reflectance (aka spectral albedo)
get multiplied wavelength by wavelength

There are different physical processes that explain this multiplication
e.g. absorption, interferences

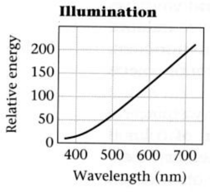

.*

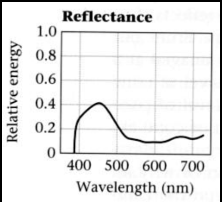

=

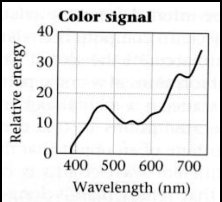

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Spectrum demo

- Diffraction grating:
  - shifts light as a function of wavelength
  - Allows you to see spectra
  - In particular, using a slit light source, we get a nice band showing the spectrum
- See the effect of filters
- See different light source spectra

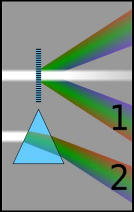

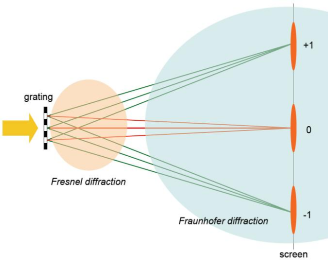

Questions?

So far, physical side of colors: **spectra**

an infinite number of values
(one per wavelength)

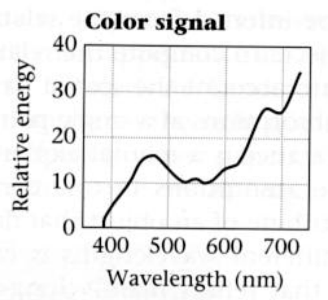

Plan

- Spectra
- **Cones and spectral response**
- Color blindness and metamers
- Color matching
- Color spaces

What is Color?

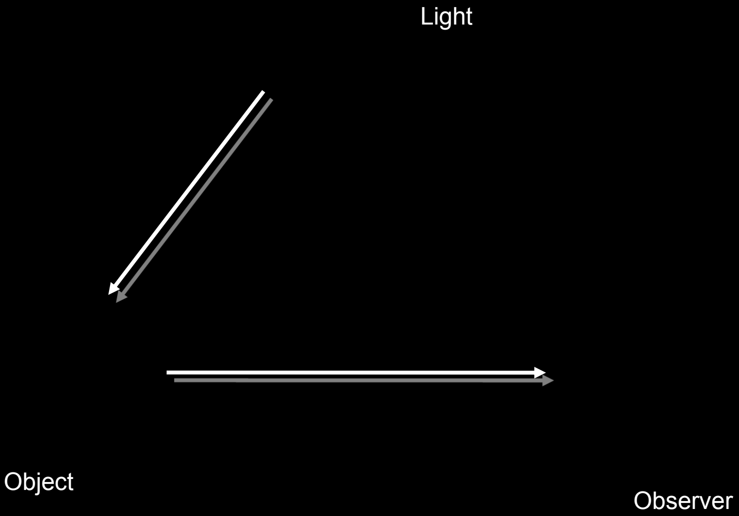

What is Color?

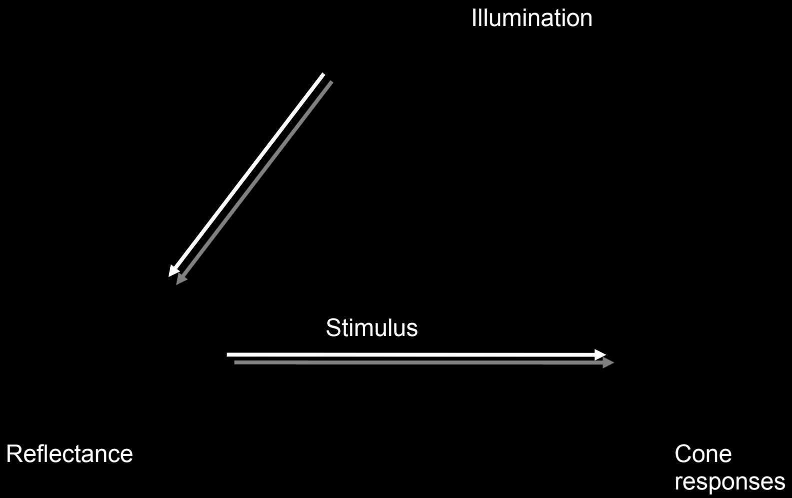

What is Color?

Light Illumination

Object Reflectance

Final stimulus

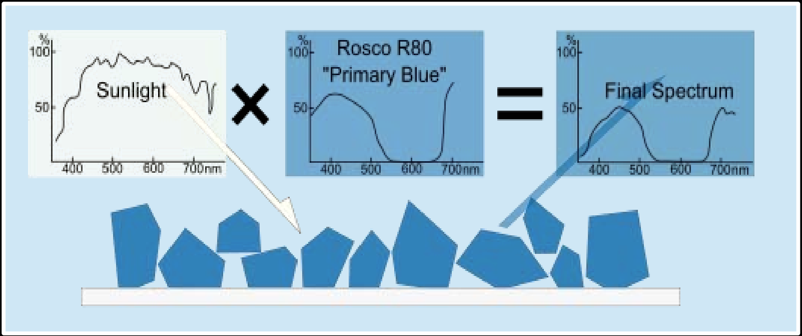

Then the cones in the eye interpret the stimulus

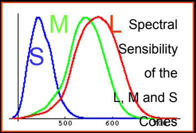

Cones

- We focus on low-level aspects of color
  - Cones and early processing in the retina
- We won't talk about rods (night vision)

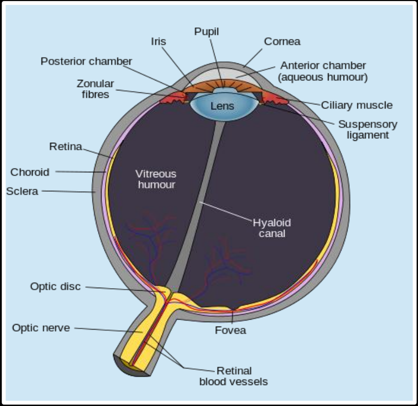

This image is in the public domain. Source: Wikimedia Commons.

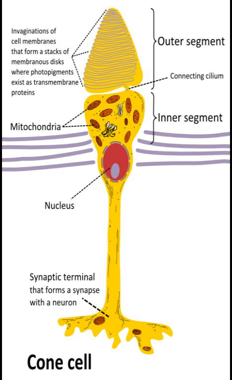

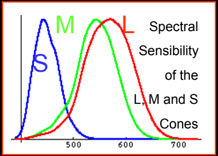

Summary (and time for questions)

- Spectrum: infinite number of values
  - can be multiplied
  - can be added
- Light spectrum multiplied by reflectance spectrum
  - spectrum depends on illuminant
- Human visual system is complicated

Cone spectral sensitivity

- Short, Medium and Long wavelength
- Response for a cone
  $= \int_\lambda \text{stimulus}(\lambda) * \text{response}(\lambda)\, d\lambda$

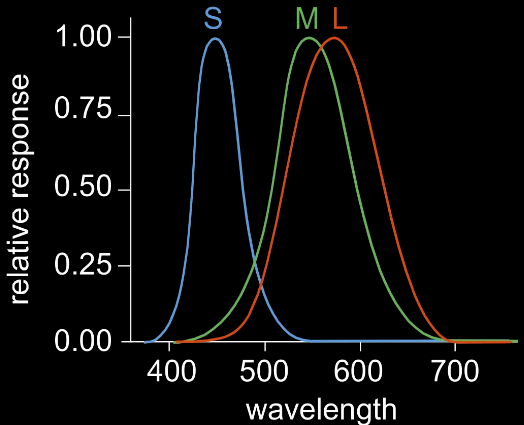

Cone response

Start from infinite number of values (one per wavelength)

End up with 3 values (one per cone type)

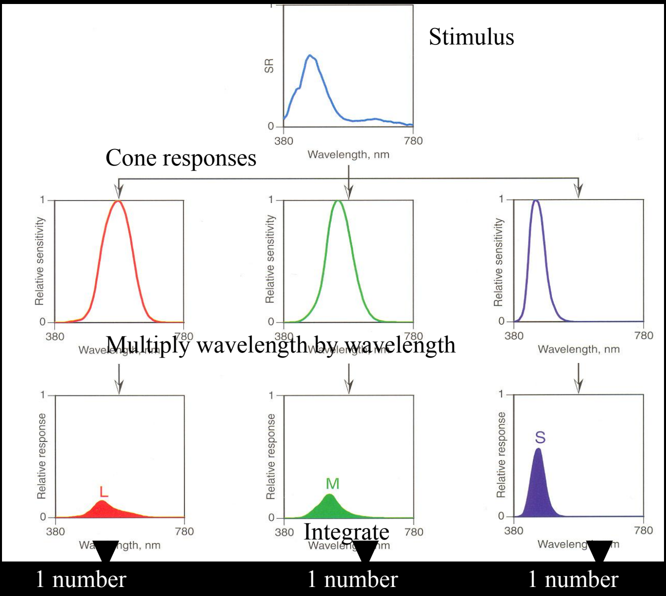

For matrix lovers

- Spectrum: big long vector size N where N=∞
- Cone response: 3xN matrix of individual responses

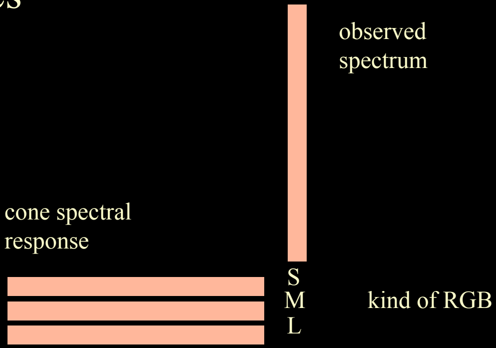

Big picture

- It's all linear!

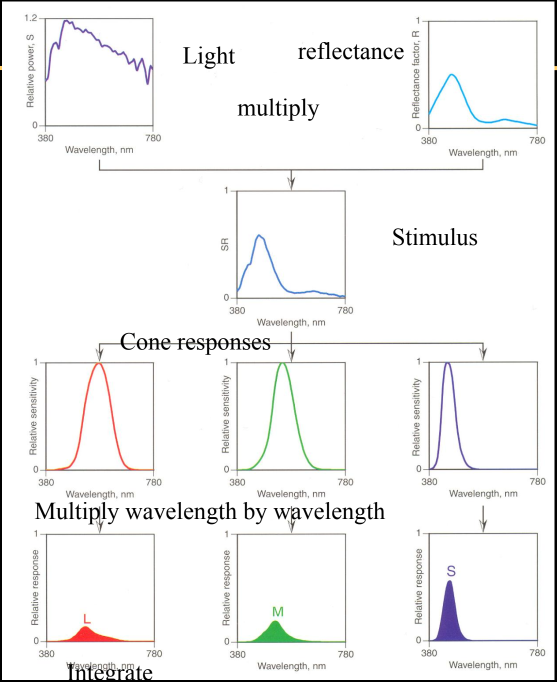

Big picture

- It's all linear!
  - multiply
  - add
- But
  - non-orthogonal basis
  - infinite dimension
  - light must be positive
- Depends on light source

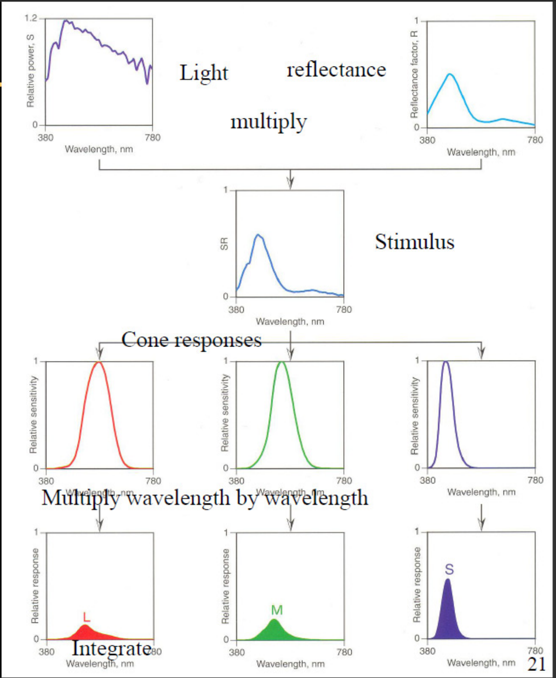

Questions?

A cone does not "see" colors

- Different wavelength, different intensity
- Same response

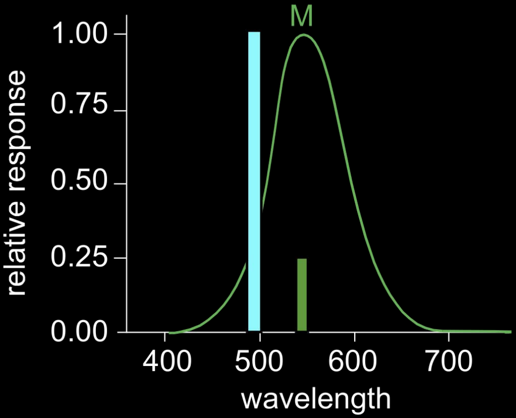

Response comparison

- Different wavelength, different intensity
- But different response for different cones

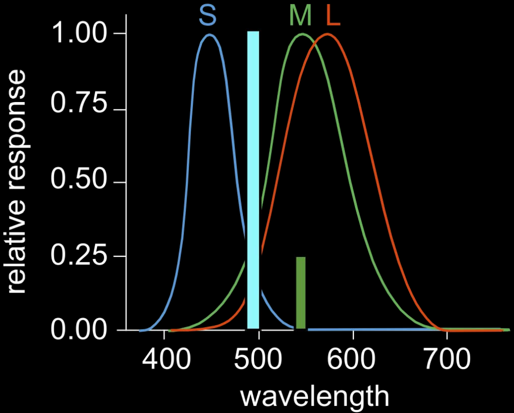

von Helmholtz 1859: Trichromatic theory

- Colors as relative responses (ratios)

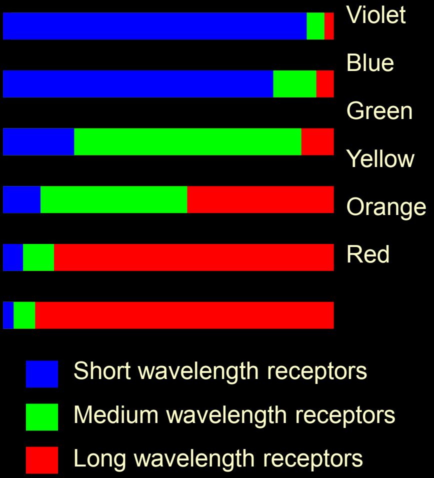

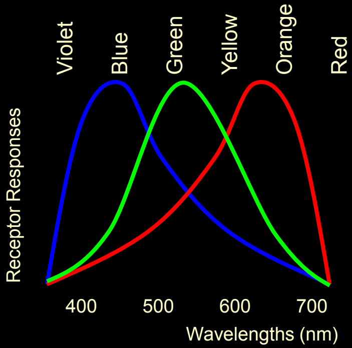

Questions?

Plan

- Spectra
- Cones and spectral response
- **Color blindness and metamers**
- Color matching
- Color spaces

Color blindness

- Classical case: 1 type of cone is missing (e.g. red)
- Makes it impossible to distinguish some spectra

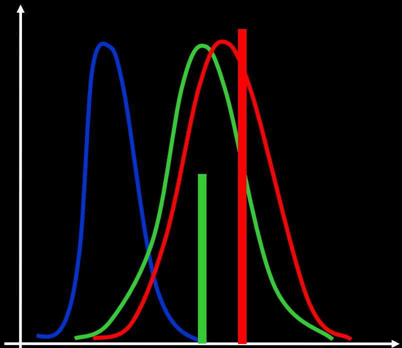

differentiated

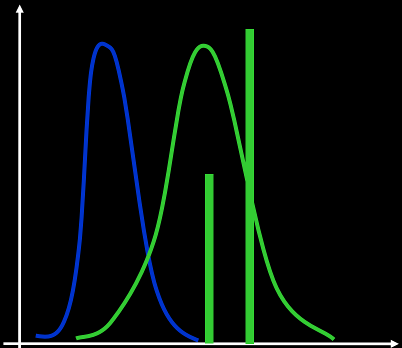

Same responses

Color blindness – more general

- Dalton
- 8% male, 0.6% female
- Genetic
- Dichromate (2% male)
  - One type of cone missing
  - L (protanope), M (deuteranope), S (tritanope)
- Anomalous trichromat
  - Shifted sensitivity

Color blindness test

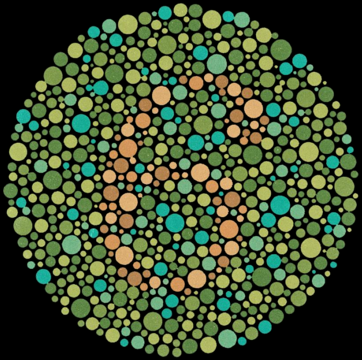

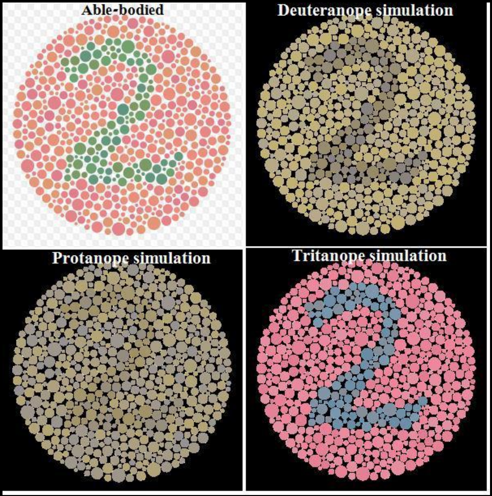

Color blindness test

- Maze in subtle intensity contrast
- Visible only to color blinds
- Color contrast overrides intensity otherwise

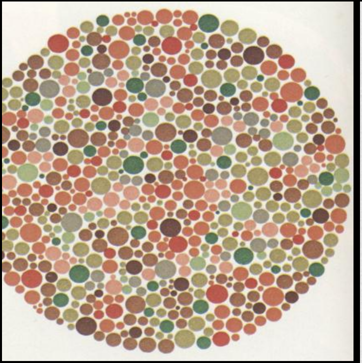

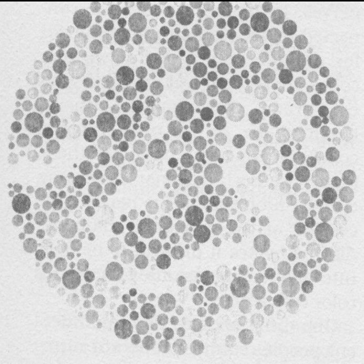

Questions?

- Links:
  - Vischeck shows you what an image looks like to someone who is colorblind.
  - http://www.vischeck.com/vischeck/
  - Daltonize, changes the red/green variation to brightness and blue/yellow variations.
  - http://www.vischeck.com/dalton
  - http://www.vischeck.com/daltonize/runDaltonize.php

Metamers

- We are all color blind!
- These two different spectra elicit the same cone responses
- Called metamers

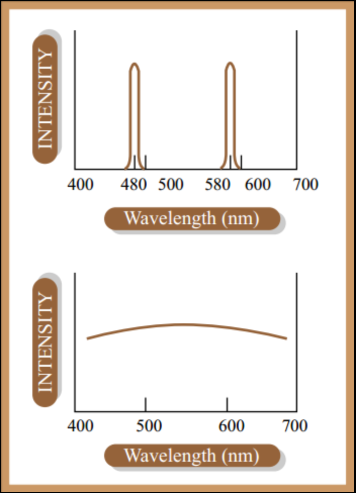

Good news: color reproduction

- 3 primaries are (to a first order) enough to reproduce all colors

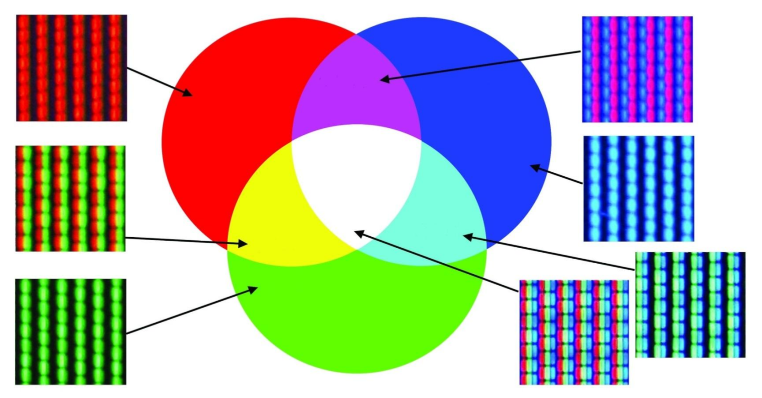

Recap

- Spectrum: infinite number of values
- projected according to cone spectral response
  => 3 values
- metamers: spectra that induce the same response
  (physically different but look the same)

- Questions?

Metamerism & light source

- Metamers under a given light source
- May not be metamers under a different lamp

Illuminant metamerism example

- Two grey patches in Billmeyer & Saltzman's book
  look the same under daylight
  but different under neon or halogen
  (& my camera agrees ;-)

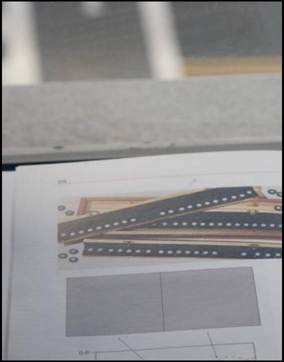

Daylight

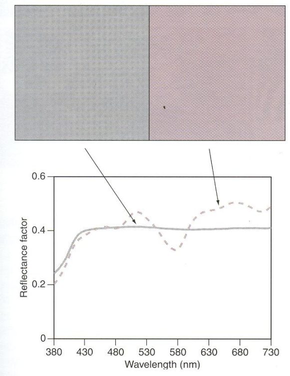

Scan (neon)

Hallogen

Bad consequence: cloth matching

- Clothes appear to match in store (e.g. under neon)
- Don’t match outdoor

Recap

- Spectrum is an infinity of numbers
- Projected to 3D cone-response space
  - for each cone, multiply per wavelength and integrate
  - a.k.a. dot product
- Metamerism: infinite-D points projected to the same 3D point
  (different spectrum, same perceived color)
  - affected by illuminant
  - enables color reproduction with only 3 primaries

Questions?

Analysis & Synthesis

- Now let's switch to technology
- We want to measure & reproduce color as seen by humans
- No need for full spectrum
- Only need to match up to metamerism

Analysis & Synthesis

- Focus on additive color synthesis
- We'll use 3 primaries (e.g. red green and blue) to match all colors

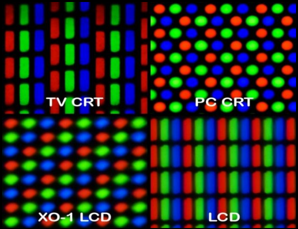

- What should those primaries be?
- How do we tell the amount of each primary needed to reproduce a given target color?

Warning

Tricky thing with spectra & color:

- Spectrum for the stimulus / synthesis
  - Light, monitor, reflectance

- Response curve for receptor /analysis
  - Cones, camera, scanner

They are usually not the same

There are good reasons for this

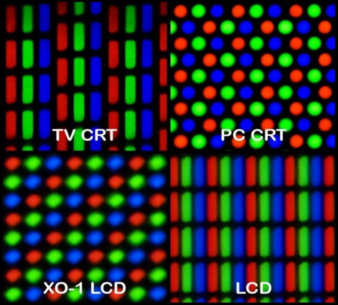

Additive Synthesis - wrong way

- Take a given stimulus and the corresponding responses s, m, l (here 0.5, 0, 0)

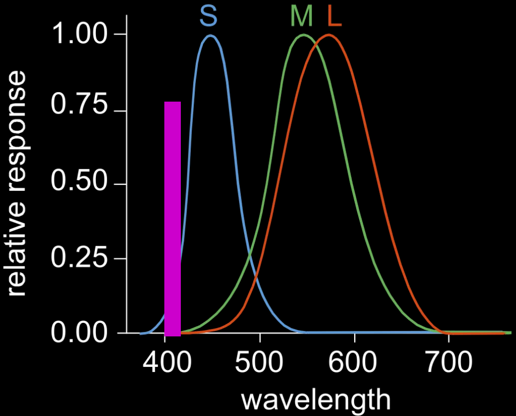

Additive Synthesis - wrong way

- Use it to scale the cone spectra (here 0.5 * S)
- You don't get the same cone response!
  (here 0.5, 0.1, 0.1)

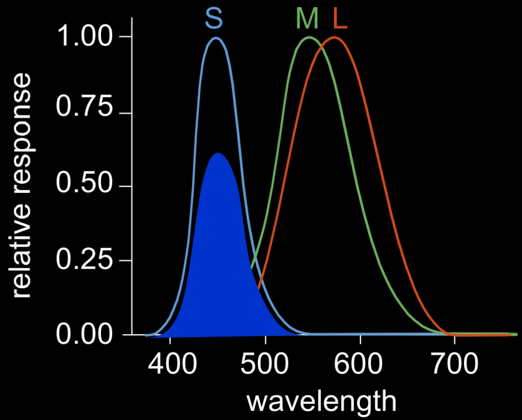

What's going on?

- The three cone responses are not orthogonal
- i.e. they overlap and "pollute" each other

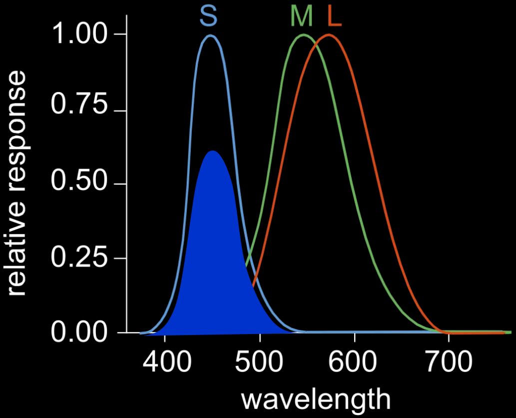

Fundamental problems

- Spectra are infinite-dimensional
- Only positive values are allowed
- Cones are non-orthogonal/overlap

Summary

- Physical color
  - Spectrum
  - multiplication of light & reflectance spectrum
- Perceptual color
  - Cone spectral response: 3 numbers
  - Metamers: different spectrum, same responses
    - Color matching, enables color reproduction with 3 primaries
- Fundamental difficulty
  - Spectra are infinite-dimensional (full function)
  - Projected to only 3 types of cones
  - Cone responses overlap / they are non-orthogonal
    - Means different primaries for analysis and synthesis
  - Negative numbers are not physical

Questions?

Standard color spaces

- We need a principled color space
- Many possible definition
  - Including cone response (LMS)
  - Unfortunately not really used, (because not known at the time)

- The good news is that color vision is linear and 3-dimensional, so any new color space based on color matching can be obtained using 3x3 matrix
  - But there are also non-linear color spaces (e.g. Hue Saturation Value, Lab)

Overview

- Most standard color space: CIE XYZ
- LMS and the various flavor of RGB are just linear transformations of the XYZ basis
  - 3x3 matrices

Why not measure cone sensitivity?

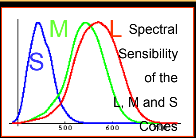

- Less directly measurable
  - electrode in photoreceptor?
  - not available when color spaces were defined
- Most directly available measurement:
  - notion of metamers & color matching
  - directly in terms of color reproduction:
    given an input color,
    how to reproduce it with 3 primary colors?
  - Commission Internationale de l'Eclairage (International Lighting Commission)
  - Circa 1920

CIE color matching

- Choose 3 synthesis primaries
- Seek to match any monochromatic light (400 to 700nm)
  - Record the 3 values for each wavelength
- By linearity, this tells us how to match any light

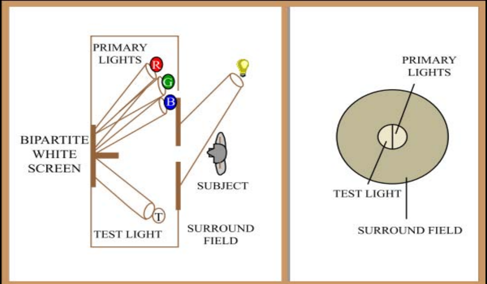

Image adapted from:
Wandell, Brian A. *Foundations of Vision: Behavior, Neuroscience and Computation*. Sinauer Associates, Inc. May 1995. ISBN: 0878938532.

CIE color matching

- Primaries (synthesis) at 435.8, 546.1 and 700nm
  - Chosen for robust reproduction, good separation in red-green
  - Don't worry, we'll be able to convert it to any other set of primaries (Linear algebra to the rescue!)
- Resulting 3 numbers for each input wavelength are called tristimulus values

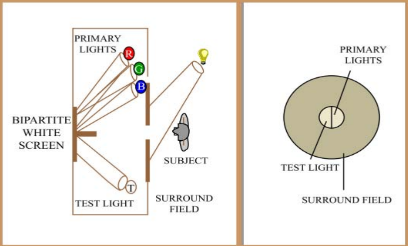

Image adapted from:
Wandell, Brian A. *Foundations of Vision: Behavior, Neuroscience and Computation.* Sinauer Associates, Inc. May 1995. ISBN: 0878938532.

Now, our interactive feature!

You are…

THE LAB RAT

Color Matching Problem

- Some colors cannot be produced using only positively weighted primaries
- Solution: add light on the other side!

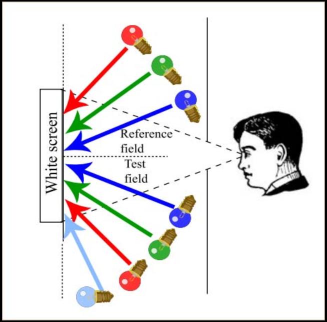

CIE color matching

- Meaning of these curves: a monochromatic wavelength λ can be reproduced with
  b(λ) amount of the 435.8nm primary,
  +g(λ) amount of the 546.1 primary,
  +r(λ) amount of the 700 nm primary

- This fully specifies the color perceived by a human

- Careful: this is not your usual rgb

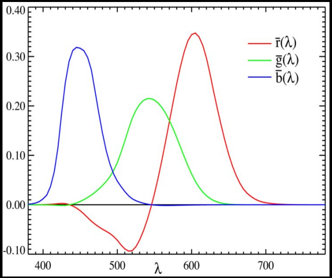

CIE color matching

- Meaning of these curves: a monochromatic wavelength λ can be reproduced with b(λ) amount of the 435.8nm primary, +g(λ) amount of the 546.1 primary, +r(λ) amount of the 700 nm primary

- This fully specifies the color perceived by a human
- However, note that one of the responses can be negative
  - Those colors cannot be reproduced by those 3 primaries.

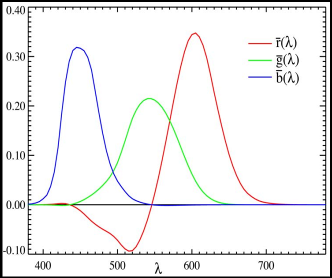

CIE color matching: what does it mean?

- If I have a given spectrum X
- I compute its response to the 3 matching curves (multiply and integrate)
- I use these 3 responses to scale my 3 primaries (435.8, 546.1 and 700nm)
- I get a metamer of X (perfect color reproduction)

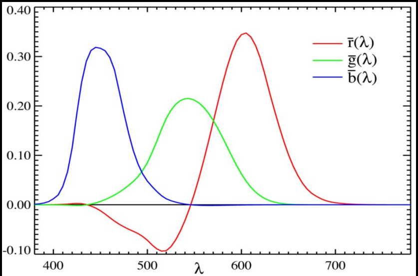

Relation to cone curves

- Project to the same subspace
  - b, g, and r are linear combinations of S, M and L
- Related by 3x3 matrix.
- Unfortunately unknown at that time. This would have made life a lot easier!

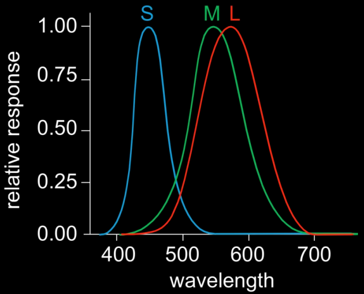

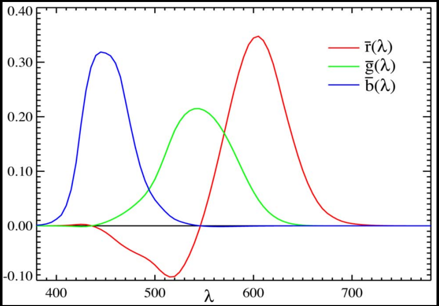

Recap

- Spectra : infinite dimensional
- Cones: 3 spectral responses
- Metamers: spectra that look the same (same projection onto cone responses)
- CIE measured color response:
  - chose 3 primaries
  - tristimulus curves to reproduce any wavelength

- Questions?

How to build a measurement device?

- Idea:
  - Start with light sensor sensitive to all wavelength
  - Use three filters with spectra b, r, g
  - measure 3 numbers
- This is pretty much what the eyes do!

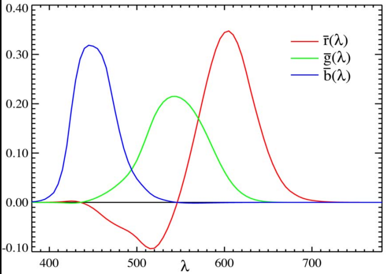

CIE's problem

- Idea:
  - Start with light sensor sensitive to all wavelength
  - Use three filters with spectra b, r, g
  - measure 3 numbers
- But for those primaries, we need negative spectra

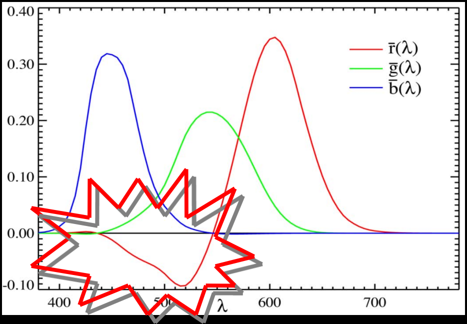

CIE's problem

- Obvious solution:
  use cone response!
  - but unknown at the time
- =>new set of tristimulus curves
  - linear combinations of b, g, r
  - pretty much add enough b and g until r is positive

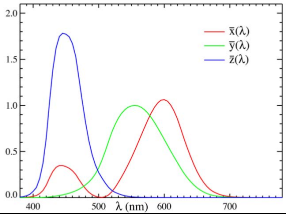

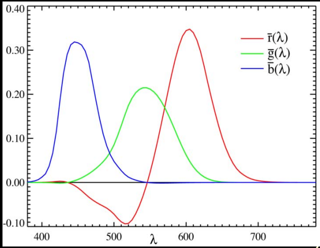

Chromaticity diagrams

- 3D space are tough to visualize
- Usually project to 2D for clarity
- Chromaticity diagram:
  - normalize against X + Y + Z:

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

CIE XYZ -recap

- THE standard for color specification
- Lots of legacy decision - I wish it were LMS
- Based on color matching
  - 3 monochromatic primaries
  - Subjects matched every wavelength
  - Tricks to avoid negative numbers
  - These 3 values "measure" or describe a perceived color.

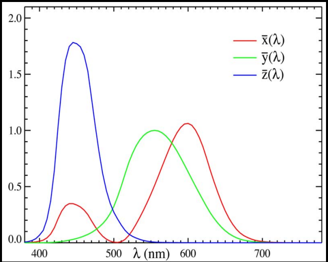

Questions?

Other primaries

- We want to use a new set of primaries
  - e.g. the spectra of R, G & B in a projector or monitor
- By linearity of color matching,
  can be obtained from XYZ by a 3x3 matrix

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} 3.24 & -1.54 & -0.50 \\ -0.97 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$$

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

one example RGB space

Other primaries

- We want to use a new set of primaries
  - e.g. the spectra of R, G & B in a projector or monitor
- By linearity of color matching,
  can be obtained from XYZ by a 3x3 matrix
- This matrix tells us how to match the 3 primary spectra from XYZ using the new 3 primaries

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} 3.24 & -1.54 & -0.50 \\ -0.97 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$$

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

one example RGB space

XYZ to RGB & back

- e.g.
  http://www.brucelindbloom.com/index.html?Eqn_RGB_XYZ_Matrix.html

- sRGB to XYZ

0.412424	0.212656	0.0193324
0.357579	0.715158	0.119193
0.180464	0.0721856	0.950444

XYZ to sRGB

3.24071	-0.969258	0.0556352
-1.53726	1.87599	-0.203996
0.498571	0.0415557	1.05707

- Adobe RGB to XYZ

0.576700	0.297361	0.0270328
0.185556	0.627355	0.0706879
0.188212	0.0752847	0.991248

XYZ to Adobe RGB

2.04148	-0.969258	0.0134455
-0.564977	1.87599	-0.118373
-0.344713	0.0415557	1.01527

- NTSC RGB to XYZ

0.606734	0.298839	0.000000
0.173564	0.586811	0.0661196
0.200112	0.114350	1.11491

XYZ to NTSC RGB

1.91049	-0.984310	0.0583744
-0.532592	1.99845	-0.118518
-0.288284	-0.0282980	0.898611

Color gamut

- Given 3 primaries
- The realizable chromaticities lay in the triangle in xy chromaticity diagram
- Because we can only add light, no negative light

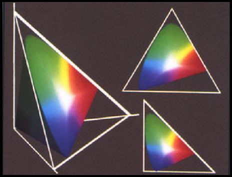

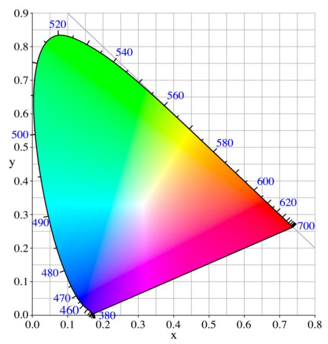

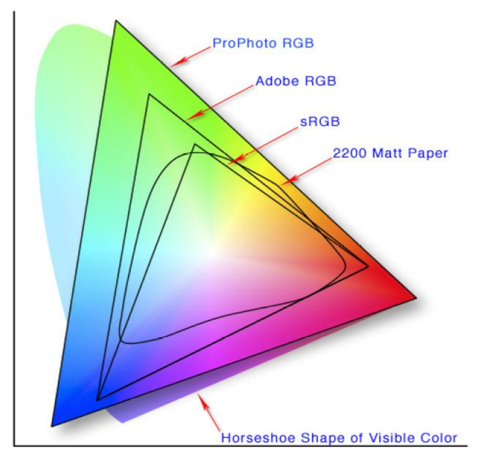
ProPhoto RGB
Adobe RGB
sRGB
2200 Matt Paper
Horseshoe Shape of Visible Color

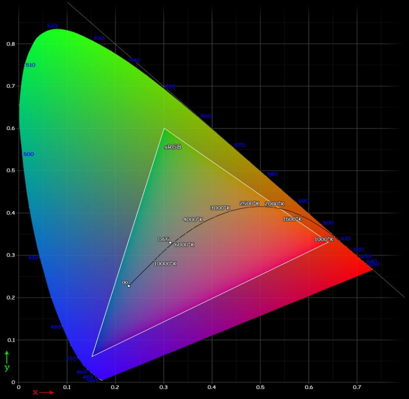
sRGB
3000°K
2500°K
2000°K
4000°K
1500°K
D65
6000°K
1000°K
10000°K
y
x

In summary

- It's all about linear algebra
  - Projection from infinite-dimensional spectrum to a 3D response
  - Then any space based on color matching and metamerism can be converted by 3x3 matrix
- Complicated because
  - Projection from infinite-dimensional space
  - Non-orthogonal basis (cone responses overlap)
  - No negative light
- XYZ is the most standard color space
- RGB has many flavors

Questions?

Gamma encoding overview

- Digital images are usually not encoded linearly
- Instead, the value **X**1/γ is stored

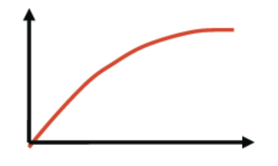

- Need to be decoded if we want linear values

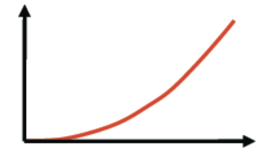

Color quantization gamma

- The human visual system is more sensitive to ratios
  - Is a grey twice as bright as another one?
- If we use linear encoding, we have tons of information between 128 and 255, but very little between 1 and 2!
- Ideal encoding?

  Log
- Problems with log?

  Gets crazy around zero

Solution: gamma

Color quantization gamma

- The human visual system is more sensitive to ratios
  - Is a grey twice as bright as another one?
- If we use linear encoding, we have tons of information between 128 and 255, but very little between 1 and 2!
- This is why a non-linear gamma remapping of about 2.0 is applied before encoding
- True also of analog imaging to optimize signal-noise ratio

Color quantization gamma

- The human visual system is more sensitive to ratios
  - Is a grey twice as bright as another one?
- If we use linear encoding, we have tons of information between 128 and 255, but very little between 1 and 2!
- This is why a non-linear gamma remapping of about 2.0 is applied before encoding
- True also of analog imaging to optimize signal-noise ratio

Gamma encoding

- From Greg Ward
- Only 6 bits for emphasis

Important Message

- Digital images are usually gamma encoded
  - Often $\gamma = 2.2$ (but 1.8 for Profoto RGB)
- To get linear values, you must decode
  - apply $x \Rightarrow x\gamma$

Questions?

Selected Bibliography

Vision Science by Stephen E. Palmer
MIT Press; ISBN: 0262161834
760 pages (May 7, 1999)

Billmeyer and Saltzman's Principles of Color Technology, 3rd Edition
by Roy S. Berns, Fred W. Billmeyer, Max Saltzman
Wiley-Interscience; ISBN: 047119459X
304 pages 3 edition (March 31, 2000)

Vision and Art : The Biology of Seeing by Margaret Livingstone, David H. Hubel Harry N
Abrams; ISBN: 0810904063
208 pages (May 2002)

The Reproduction of Color by R. W. G. Hunt
Fountain Press, 1995

Color Appearance Models by Mark Fairchild
Addison Wesley, 1998

Color for the Sciences, by Jan Koenderink
MIT Press 2010.

Questions?

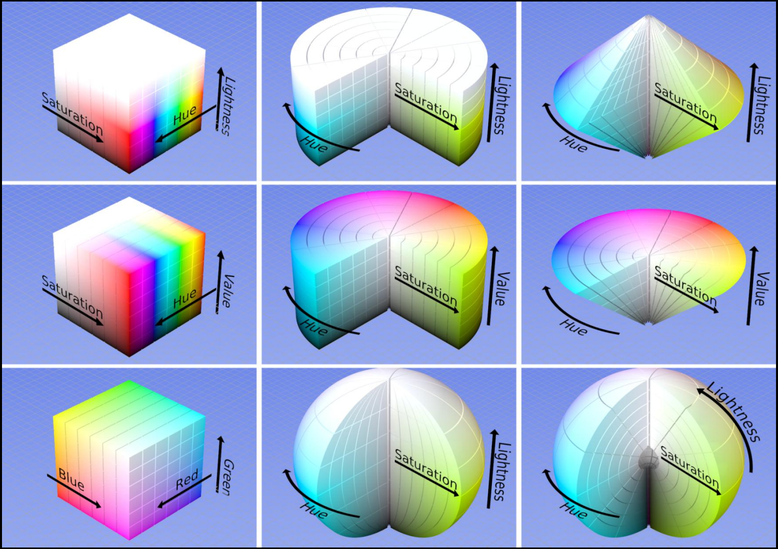
Saturation
Hue
Lightness
Saturation
Hue
Lightness
Saturation
Hue
Lightness
Saturation
Hue
Value
Saturation
Hue
Value
Saturation
Hue
Value
Blue
Red
Green
Saturation
Hue
Lightness
Lightness
Saturation
Hue

Shading & Material Appearance

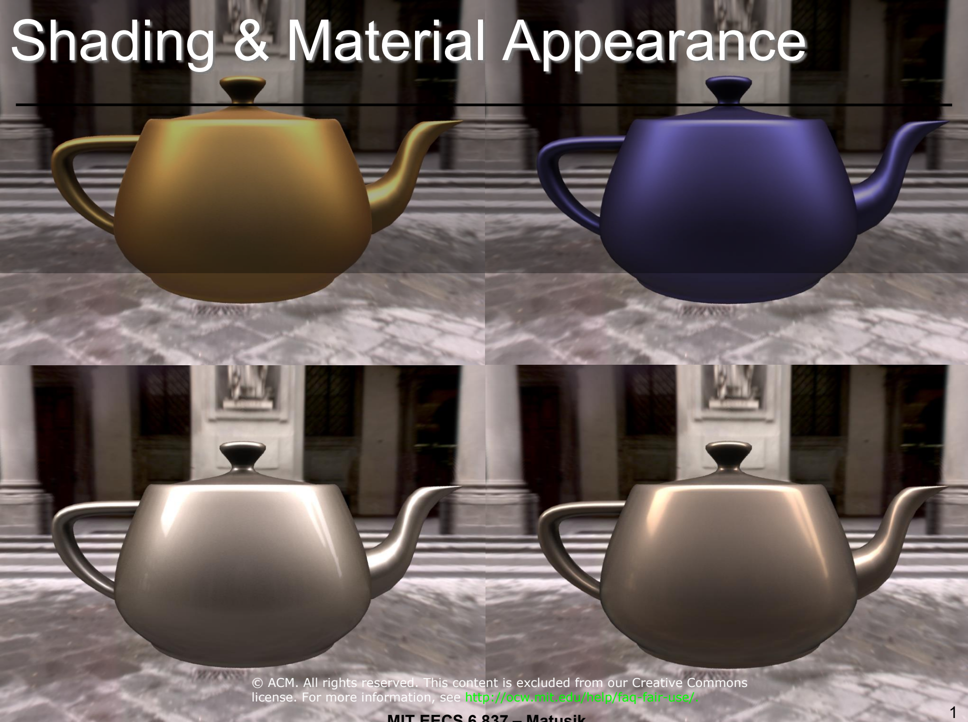

Lighting and Material Appearance

- Input for realistic rendering
  - Geometry, Lighting and Materials
- Material appearance
  - Intensity and shape of highlights
  - Glossiness
  - Color
  - Spatial variation, i.e., texture (next Tuesday)

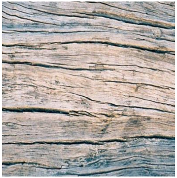

Slide Addy Ngan

Unit Issues - Radiometry

- We will not be too formal in this class
- Issues we will not really care about
  - Directional quantities vs. integrated over all directions
  - Differential terms: per solid angle, per area
  - Power? Intensity? Flux?

- Color
  - All math here is for a single wavelength only; we will perform computations for R, G, B separately
    - Do not panic, that just means
      we will perform every operation three times, that is all

Light Sources

- Today, we only consider point light sources
  - Thus we do not need to care about solid angles
- For multiple light sources, use linearity
  - We can add the solutions for two light sources
    - I(a+b) = I(a) + I(b)
  - We simply multiply the solution when we scale the light intensity
    - I(s a) = s I(a)

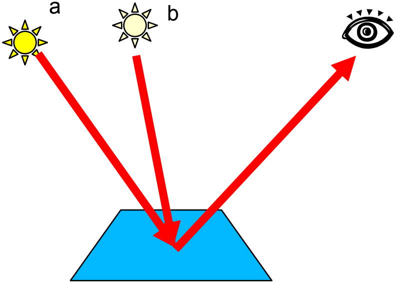

Yet again, linearity is our friend!

Intensity as Function of Distance

- $1/r^2$ fall-off for isotropic point lights
  - Why? An isotropic point light outputs constant power per solid angle ("into all directions")
  - Must have same power in all concentric spheres
    - Sphere's surface area grows with r^2 => energy obeys $1/r^2$

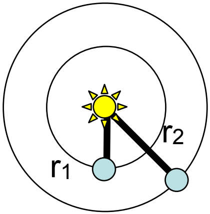

- … but in graphics we often cheat with or ignore this.
  - Why? Ideal point lights are kind of harsh
    - Intensity goes to infinity when you get close – not great!
  - In particular, $1/(ar^2+br+c)$ is popular

Incoming Irradiance

- The amount of light energy received by a surface depends on incoming angle
  - Bigger at normal incidence, even if distance is const.
    - Similar to winter/summer difference
- How exactly?
  - Cos θ law
  - Dot product with normal

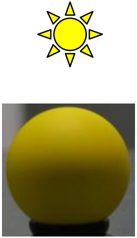

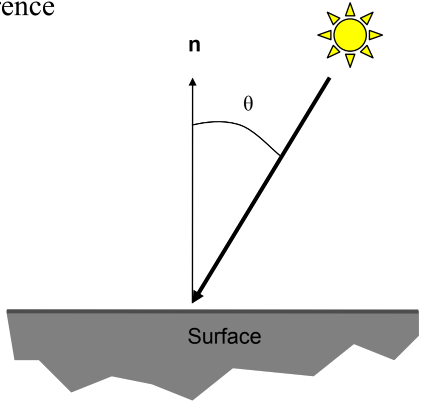

Incoming Irradiance for Pointlights

- Let's combine this with the $1/r^2$ fall-off:

$$I_{\text{in}} = I_{\text{light}} \cos\theta / r^2$$

  - I_{in} is the irradiance ("intensity") at surface point **x**
  - I_{light} is the "intensity" of the light
  - θ is the angle between light direction **l** and surface normal **n**
  - r is the distance between light and **x**.

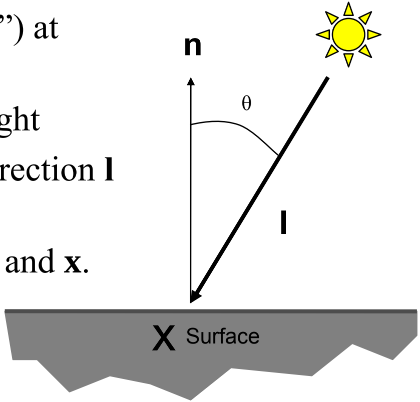

Directional Lights

- "Pointlights that are infinitely far"
  - No falloff, just one direction and one intensity

$$I_{\text{in}} = I_{\text{light}} \cos\theta$$

  - I_{in} is the irradiance at surface point x from the directional light
  - I_{light} is the "intensity" of the light
  - θ is the angle between light direction **l** and surface normal **n**
    - Only depends on **n**, not **x**!

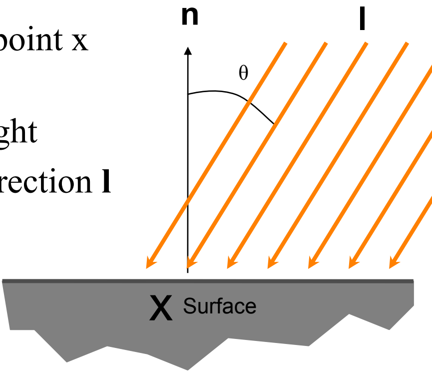

Spotlights

- Pointlights with non-uniform directional emission
- Usually symmetric about a central direction **d**, with angular falloff
  - Often two angles
    - "Hotspot" angle:
      No attenuation within the central cone
    - "Falloff" angle: Light attenuates from full intensity to zero intensity between the hotspot and falloff angles

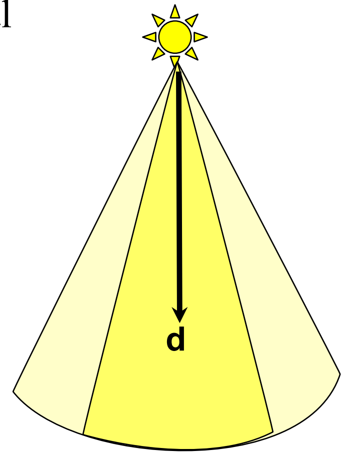

- Plus your favorite distance falloff curve

Spotlight Geometry

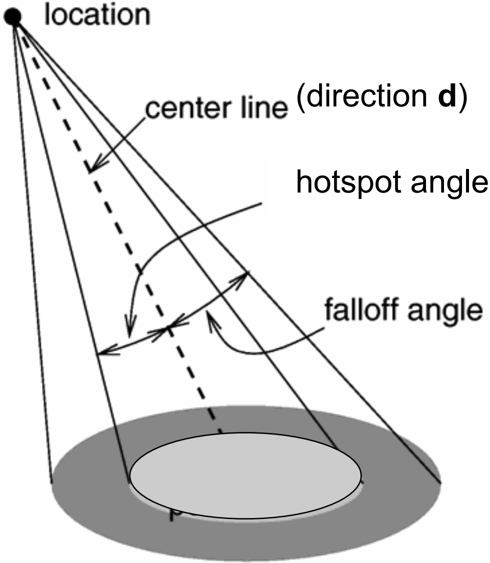

Adapted from
POVRAY documentation

Questions?

Quantifying Reflection – BRDF

- Bidirectional Reflectance Distribution Function
- Ratio of light coming from one direction that gets reflected in another direction
  - Pure reflection, assumes no light scatters into the material
- Focuses on angular aspects, not spatial variation of the material
- **How many dimensions?**

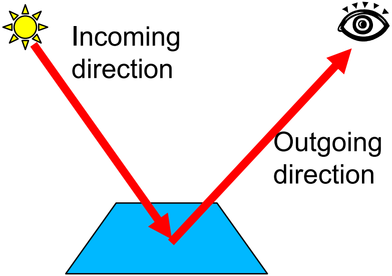

BRDF f_r

- Bidirectional Reflectance Distribution Function
  - 4D: 2 angles for each direction
  - BRDF = $f_r(\theta_i, \phi_i; \theta_o, \phi_o)$
  - Or just two unit vectors: BRDF = $f_r(\mathbf{l}, \mathbf{v})$
    - $\mathbf{l}$ = light direction
    - $\mathbf{v}$ = view direction

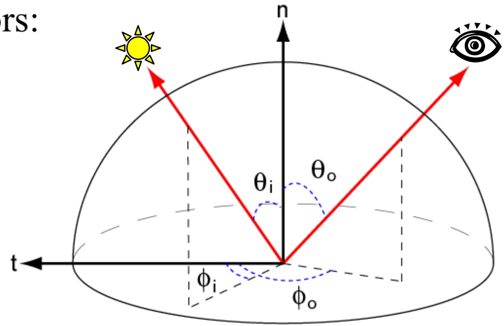

BRDF f_r

- Bidirectional Reflectance Distribution Function
  - 4D: 2 angles for each direction
  - BRDF = $f_r(\theta_i, \phi_i; \theta_o, \phi_o)$
  - Or just two unit vectors: BRDF = $f_r(\mathbf{l}, \mathbf{v})$
    - $\mathbf{l}$ = light direction
    - $\mathbf{v}$ = view direction
  - The BRDF is aligned with the surface; the vectors **l** and **v** must be in a local coordinate system

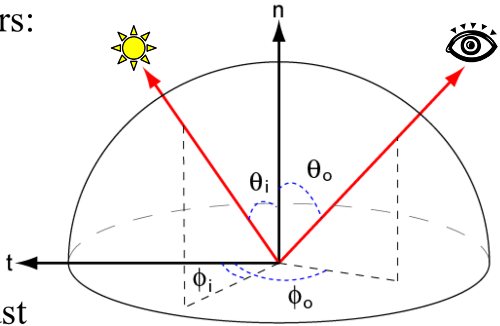

BRDF f_r

- Relates incident irradiance from every direction to outgoing light. How?

l = light direction (incoming)
v = view direction (outgoing)

$$I_{\text{out}}(\boldsymbol{v}) = I_{\text{in}}(\boldsymbol{l})\, f_r(\boldsymbol{v}, \boldsymbol{l})$$

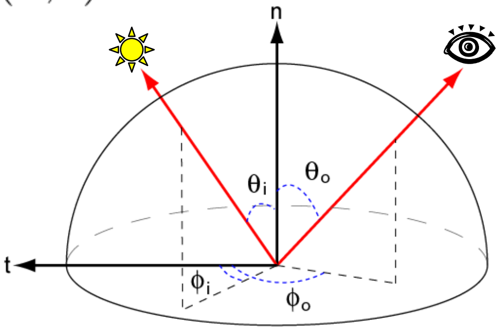

BRDF f_r

- Relates incident irradiance from every direction to outgoing light. How?

l = light direction (incoming)
v = view direction (outgoing)

$$I_{\text{out}}(\boldsymbol{v}) = I_{\text{in}}(\boldsymbol{l})\, f_r(\boldsymbol{v}, \boldsymbol{l})$$

- Let's combine with what we know already of pointlights:

$$I_{\text{out}}(\boldsymbol{v}) =$$

$$\frac{I_{\text{light}} \cos\theta_i}{r^2} f_r(\boldsymbol{v}, \boldsymbol{l})$$

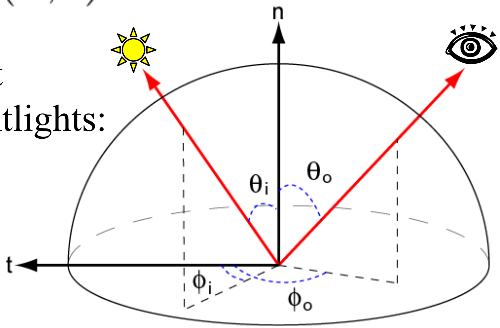

2D Slice at Constant Incidence

- For a fixed incoming direction, view dependence is a 2D spherical function
  - Here a moderate specular component

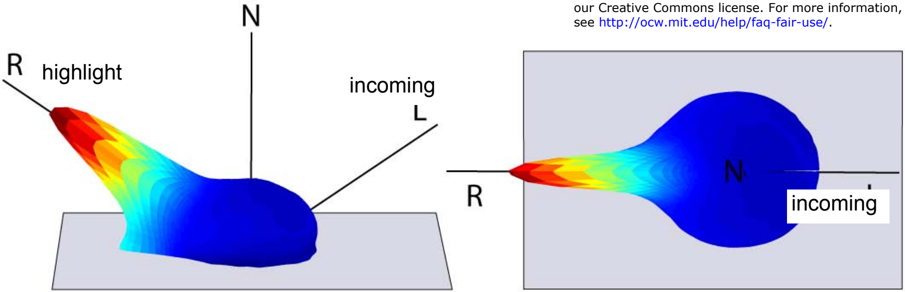

Example: Plot of “PVC” BRDF at 55° incidence

Isotropic vs. Anisotropic

- When keeping **l** and **v** fixed, if rotation of surface around the normal does not change the reflection, the material is called isotropic
- Surfaces with strongly oriented microgeometry elements are anisotropic
- Examples:
  - brushed metals,
  - hair, fur, cloth, velvet

Westin et.al 92

How do we obtain BRDFs?

- One possibility: Gonioreflectometer
  - 4 degrees of freedom

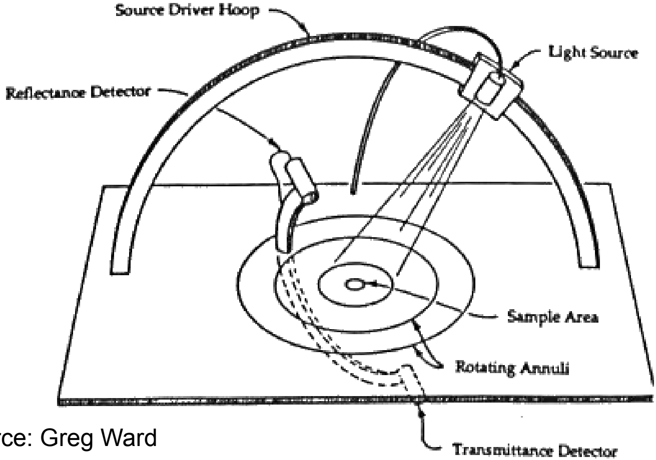

Source: Greg Ward

How Do We Obtain BRDFs?

- Another possibility: Take pictures of spheres coated with material, rotate light around a 1D arc
  - This gives 3DOF => isotropic materials only

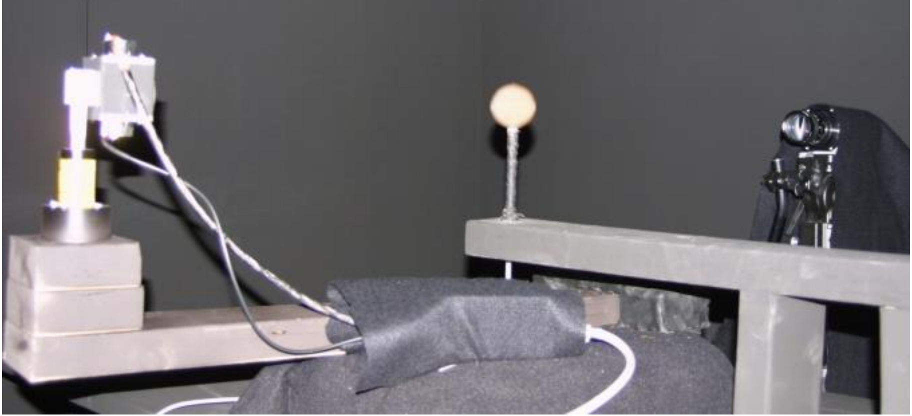

Matusik et al.

Parametric BRDFs

- BRDFs can be measured from real data
  - But tabulated 4D data is too cumbersome for most uses
- Therefore, parametric BRDF models represent the relationship between incident and outgoing light by some mathematical formula
  - The appearance can then be tuned by setting parameters
    - "Shininess", "anisotropy", etc.
  - Physically-based or Phenomenological
  - They can model with measured data (examples later)
- Popular models: Diffuse, Blinn-Phong, Cook-Torrance, Lafortune, Ward, Oren-Nayar, etc.

Questions?

Ideal Diffuse Reflectance

- Assume surface reflects equally in all directions.
- An ideal diffuse surface is, at the microscopic level, a very rough surface.
  - Example: chalk, clay, some paints

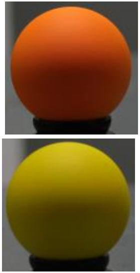

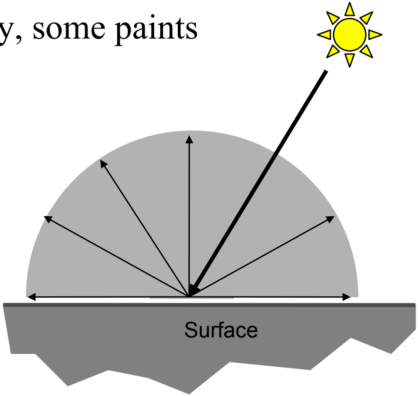

Ideal Diffuse Reflectance

- Ideal diffuse reflectors reflect light according to Lambert's cosine law
  - The reflected light varies with cosine even if distance to light source is kept constant

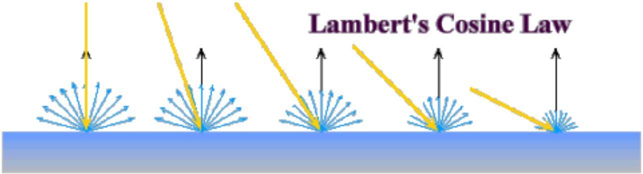

Ideal Diffuse Reflectance

- Ideal diffuse reflectors reflect light according to Lambert's cosine law
  - The reflected light varies with cosine even if distance to light source is kept constant

Remembering that incident irradiance depends on cosine, what is the BRDF of an ideally diffuse surface?

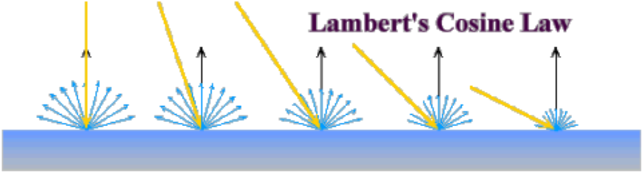

Ideal Diffuse Reflectance

- The ideal diffuse BRDF is a constant $f_r(\mathbf{l}, \mathbf{v})$ = const.
  - What constant ρ/π, where ρ is the *albedo*
    - Coefficient between 0 and 1 that says what fraction is reflected
  - Usually just called "diffuse color" k_d
  - You have already implemented this by taking dot products with the normal and multiplying by the "color"!

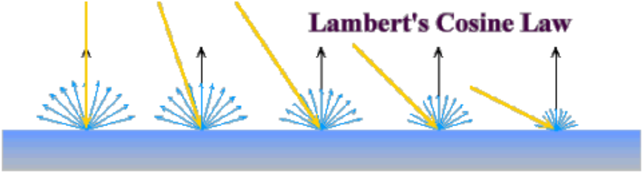

Ideal Diffuse Reflectance

- This is the simplest possible parametric BRDF
  - One parameter: k_d
    - (One for each RGB channel)

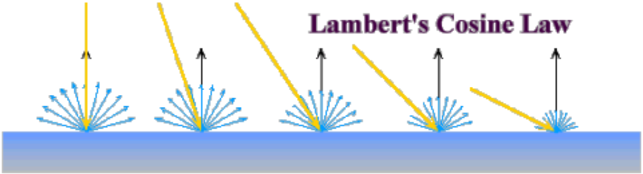

Ideal Diffuse Reflectance Math

- Single Point Light Source
  - k_d: diffuse coefficient (color)
  - **n**: Surface normal.
  - **l**: Light direction.
  - L_i: Light intensity
  - r: Distance to source
  - L_o: Shaded color

$$L_o = k_d \max(0, \boldsymbol{n} \cdot \boldsymbol{l}) \frac{L_i}{r^2}$$

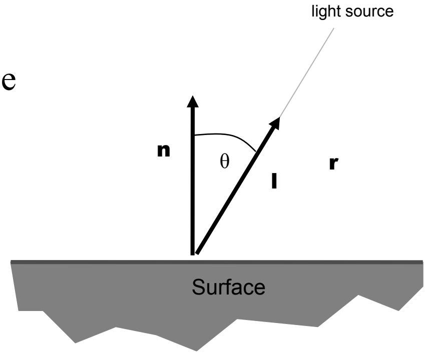

Ideal Diffuse Reflectance Math

- Single Point Light Source
  - k_d: diffuse coefficient (color)
  - **n**: Surface normal.
  - **l**: Light direction.
  - L_i: Light intensity
  - r: Distance to source
  - L_o: Shaded color

$$L_o = k_d \boxed{\max(0, \boldsymbol{n} \cdot \boldsymbol{l})} \frac{L_i}{r^2}$$

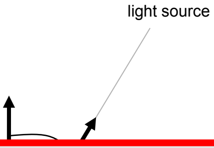

Do not forget to normalize your **n** and **l**!

We do not want light from below the surface! From now on we always assume (on this lecture) that **dot products are clamped to zero** and skip writing out the max().