EECS 487: Interactive Computer Graphics

Lecture 23:

- Intensity
- HDR and Tone Mapping

Intensity

Intensity/luminance: achromatic light, shades of gray, light with color (chromaticity) removed

Brightness: human perception of intensity

Steven's power law: human senses response to stimuli following a non-linear power law $S = Y^p$

Sense	Exponent (p)
brightness	0.33
smell	0.55
sound	0.60
taste	0.80
length	1.00
heaviness	1.45

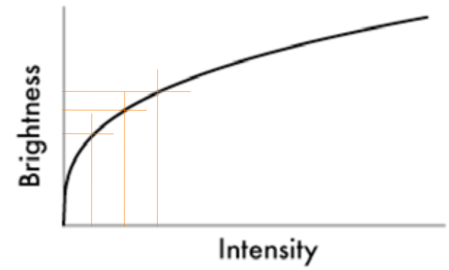

Hanrahan09

Luminance Ratio

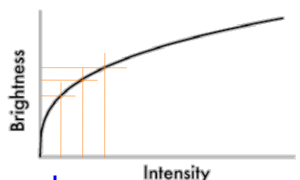

As with hearing, the human eye is more sensitive to ratio of luminance levels than to absolute values of luminance

- on a 50-100-150-watt light bulb, we see a higher jump in brightness from 50 to 100 watt (100% increase) than from 100 to 150 watt (50% increase)

Weber's Law of Just-Noticeable Difference: human eyes cannot perceive differences in brightness smaller than 1%:
$JND = \Delta Y/Y \approx 0.01$

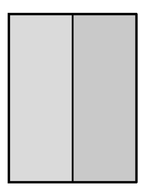

Hanrahan09

Luminance Encoding

Intensity (Y) can be represented as a real number, 0 (black) $\leq Y \leq$ 1(white)

- intermediate values are called grayscales

Given k bits to represent the intensity of each pixel we get $n = 2^k$ different intensity values per pixel

Shall we encode intensity values at linear step sizes?

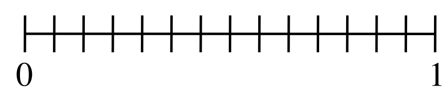

TP3

Luminance Encoding

Consequent to Steven's Law

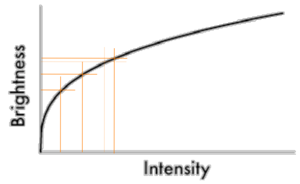

- the intensity values should be spaced out logarithmically
- keeping a constant luminance ratio (r)
- to achieve linear increases in brightness (perceived luminance)

$$r = (Y_{max} / Y_{min})^{1/(n-1)}$$

$$Y_j = r_j Y_{min} = (Y_{max} / Y_{min})^{j/(n-1)} Y_{min}$$

for $Y_{max} = 1$ (white), $Y_j = Y_{min}{}^{1-j/(n-1)}, 0 \leq j \leq n-1$

Levels of Grayscale

How many levels of grayscale (n) do we need?

Following Weber's JND Law,

$$(Y_{max} / Y_{min})^{1/(n-1)} \geq 1.01 \text{ and}$$

$$n = 1 + \log_{1.01}(1/Y_{min})$$

typical monitors have Y_{min} not 0 but $\approx 1/256$ to $1/300$

$$\Rightarrow n \approx 500$$

Display Transfer Function

Display devices also have non-linear relationship between the voltage input (V) and the light output (Y): $Y = a\,(V+\varepsilon)^{\gamma}$, ε is the black level (brightness) setting, a and γ are constants, $\gamma \in [1.5, 3.0]$

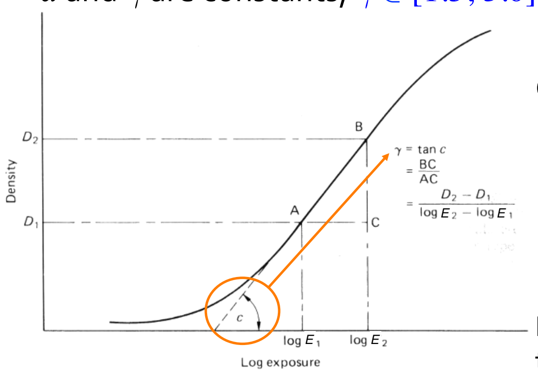

Gamma in terms of density and log exposure

Jacobson et al., Manual of Photography

Consequences:

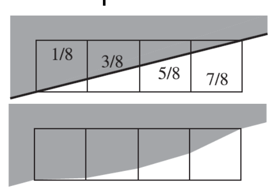

RTR3

Images appear darker than intended

Gamma Correction

Pre-adjust input values to achieve linear relationship between input and output values

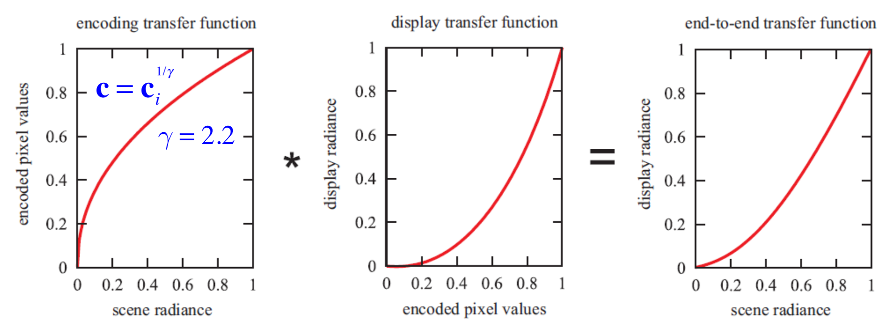

RTR3

Dynamic Ranges

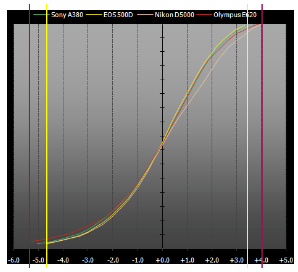

Dynamic range: the ratio between the brightest to the darkest point in a scene, a.k.a., the maximum contrast ratio, usually presented as $\log_{10}(Y_{max}/Y_{min})$

For a monitor, Y_{min} is about $1/256 - 1/300$

Display/capture media	Dynamic range
monitor	2.4 (256:1)
digital cameras	3.6 (4096:1)
photo print	2 (100:1)
photo slide	3 (1000:1)
b&w newsprint	1.7 (50:1)

In contrast, the human eye has a dynamic range of 100,000:1 adjustable to a wider range

High Dynamic Range

Natural illumination can have a dynamic range of 10^9:1, though typical scene has 10^5:1

The human eye can simultaneously perceive illumination with dynamic range of 10^5:1, and adapting more gradually to the full range of natural illumination

Luminance level of white surface under	in cd/m²
direct sunlight	10^5
indoor lighting	10^2
moonlight	10^{-1}
starlight	10^{-3}

The human eye can also do local adaptation to a relatively small solid angle in the visual field, about 1sr around the current fixation point

High Dynamic Range

Humans can't see most of the natural spectrum of light, but cameras can't capture and monitors can't display even most of what humans can see

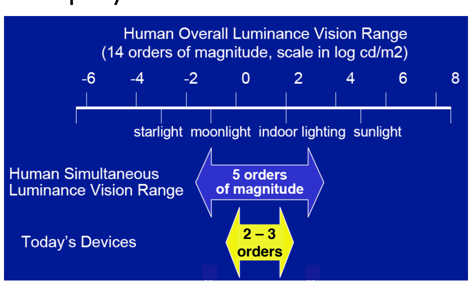

Media	Dynamic Range
nature	9 (10^9:1)
human eye	5 (10^5:1)
digital camera	3.6 (4×10^3:1)
display monitor	2.4 (2.6×10^2:1)

How can we capture and display high dynamic range images?

Hodgins

High Dynamic Range Imaging

A two-part process:

1. Compensate for the low dynamic range of sensors:
   a) generate image with (physically-based) global illumination rendering methods
   b) recover HDR from multiple-exposure low dynamic range images
2. Compensate for low dynamic range of display devices (tone reproduction/tone mapping):
   - fit a wide illumination range (10^4–10^6:1) to within a limited range (10^2–10^3:1)

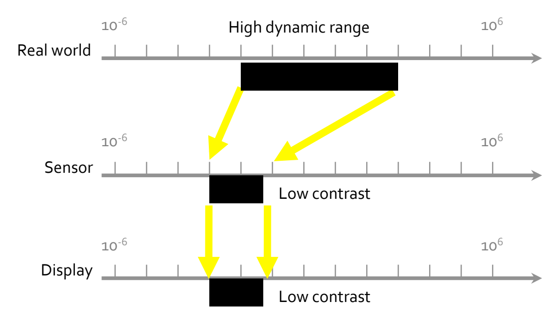

Freeman&Durando6

Camera Exposure

Exposure (X) $=$ Irradiance (E) $\times$ Time (T)

Irradiance (E):
- controlled by aperture

Exposure time (T):
- controlled by shutter speed

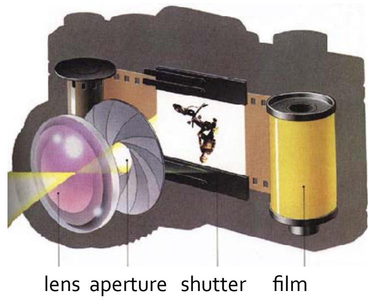

Shutter Speed

How long sensor is exposed to light

Doubling the open time doubles exposure (X) … until sensor saturates

Denoted in multiples of a second:
- $1/1000, 1/500, 1/250, 1/125, 1/60, 1/30, \ldots, 1, \ldots$

Aperture

Diameter of lens opening (controlled by diaphragm)
1 stop down halves X

Expressed as a fraction of focal length, in f-number
- f/2.0 on a 50 mm lens means aperture is 25 mm
- f/2.0 on a 100 mm lens means aperture is 50 mm
- typical f numbers:
  - f/1.8, f/2.0, f/2.8, f/4, f/5.6, f/8, f/11, f/16, f/22, f/32

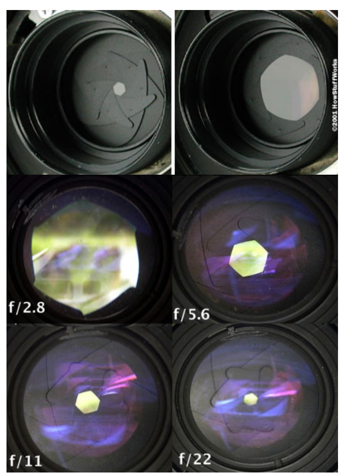

Law of Reciprocity

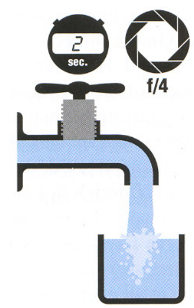

The same exposure is obtained with an exposure twice as long and an aperture opening half as big
- hence $\sqrt{2}$ progression of f stops vs. power of two progression of shutter speed

For a target exposure, given the combinations of aperture and shutter speed:

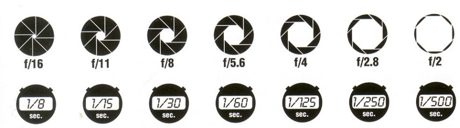

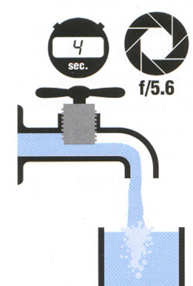

how do we choose?

London et al., Photography

High Dynamic Range Imaging

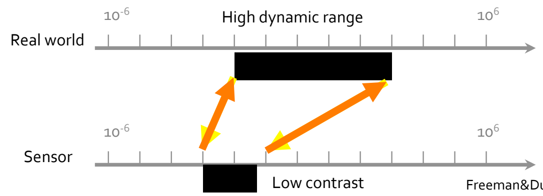

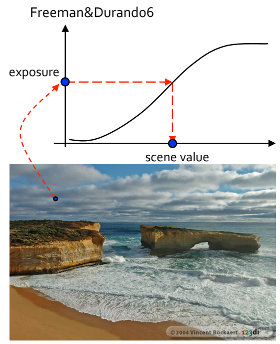

Two-part process:

1. How to capture HDR images with low dynamic range sensors (film or CCD)?
   - idea: if we know the response curve of the sensor, we can inverse map captured exposure to real world intensity!
2. How to display? tone mapping

If We Know The Response Curve…

Difficult because of discretization, clipped highlights and clipped shadows, and that response curves are non-linear and poorly documented (often times unpublished)

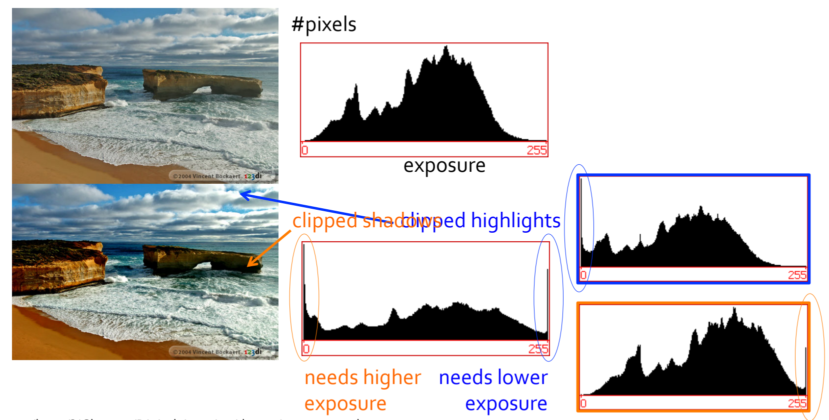

dpreview.com/learn/?/Glossary/Digital_Imaging/dynamic_range_01.htm

Multiple-Exposure HDR Imaging

Sequentially measure all segments of the range

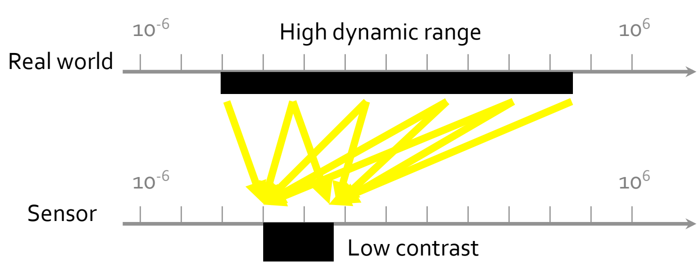

Freeman&Durando6

Multiple-Exposure HDR Imaging

Take P photos at different exposures

- for each photo, measure exposure at N spots that receive different luminance ($X_{ij}=E_iT_j$, measured exposure at pixel i in photo j)
- let the observed pixel value in image j at pixel i be Z_{ij}
- reconstruct non-linear response curve ($g(\bullet)$)
- recover actual scene irradiance at each pixel i (E_i)

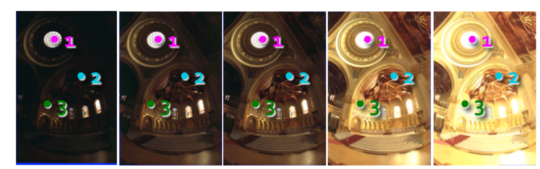

$T =$ ⅛ sec | $T =$ ½ sec | $T =$ 2 sec | $T =$ 8 sec | $T =$ 30 sec

$Z_{ij} = f\left(E_iT_j\right), f(\cdot)$ unknown but assumed monotonically increasing

$$f^{-1}\left(Z_{ij}\right) = E_iT_j$$

$$\log f^{-1}\left(Z_{ij}\right) = \log(E_iT_j)$$

Let $\log f^{-1}(\cdot)$ be $g(\cdot)$:

$$g\left(Z_{ij}\right) = \log(E_iT_j)$$

Freeman&Durando6

Response Curve

Irradiance (E_i) is unknown, fit to find a smooth curve

Assuming unit irradiance at midpoint pixel

After adjusting irradiances to obtain a smooth response

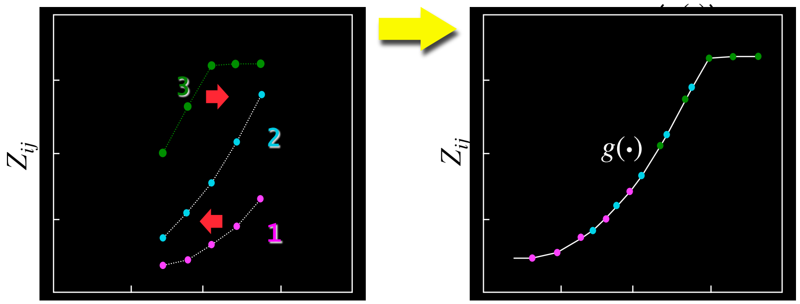

log Exposure ($\log(E_i T_j)$)

log Exposure ($\log(E_i T_j)$)

Debevec;Efros;Freeman&Durand06

Recovering Irradiance

From the fitted response curve $g(Z_{ij})$, compute
$g(Z_{ij}) = \log(E_i T_j) \Rightarrow \log(E_i) = g(Z_{ij}) - \log(T_j)$

Practical caveat if you actually want to try this:

- the resulting irradiances are relative valued
  - if absolute values are needed, match against a photograph of a calibrating luminaire of known radiance
- color photo would need a separate response curve ($g(\cdot)$) for each color
  - for correct color balance, may need to calibrate for each color

Smartphone Camera Apps have HDR functionality built-in

Tone Mapping

Goal: to fit HDR to within the limited dynamic range of display media

$$X_i = h(E_i)$$

Different objectives:

- reveal as much detail as possible
- match what one could realistically see
- "artistic" digital photography (similar to dodging and burning when developing B&W film)

What would be a good $h(\cdot)$?

Other Tone Mapping Operators

Gradient-domain operators

Frequency-based operators

- perceptually driven tone-reproduction

Spatial operators: global or local

- division
- sigmoids
- photographically motivated tone reproduction

Only the log, histogram, and sigmoids methods can be done in real-time

See Martin Čadík's website

http://cadik.posvete.cz/tmo/

HDR Encoding Formats

Encoding Format	Visible Gamut	bits/ pixel	Dynamic Range	Luminance Step
sRGB	✗	24	1.6 (1:0.025)	non-log
PixarLog	✗	33	3.8 ($25:4\times10^{-3}$)	0.4%
RGBE XYZE	✗ ✔	32	76 ($10^{38}:10^{-38}$)	1%
LogLuv24	✔	24	4.8 ($15.9:2.5\times10^{-4}$)	1.1%
LogLuv32	✔	32	38 ($10^{19}:10^{-20}$)	0.3%
ILM EXR	✔	48	10.7 ($6.5\times10^{4}:1.2\times10^{-6}$)	0.1%
scRGB	✔	48	3.5 ($7.5:2.3\times10^{-3}$)	non-log
scRGB-nl scYCC-nl	✔ ✔	36	3.2($6.2:3.9\times10^{-3}$)	non-log

Ward's *High Dynamic Range Image Encodings*
http://www.anyhere.com/gward/hdrenc/hdr_encodings.html

See http://www.extremetech.com/article2/0,2845,1170684,00.asp for more on Microsoft/HP's sRGB and Microsoft Vista's scRGB

32-bit LogLuv

32 bits/pixel

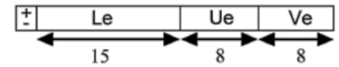

- 1 bit for the intensity sign (negative intensity allowed)
- 15 bits for the intensity value
- 16 bits for chromaticity

Take advantage of Steven's and Weber's (JND) Laws and encode intensity (L) at a logarithmic scale:

$$L_e = \lfloor c_1(\log_2 L + c_2) \rfloor,$$
$$L = 2^{[L_e/c_1 - c_2]}$$