Measuring Light

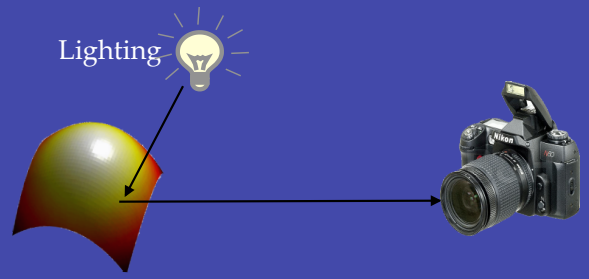
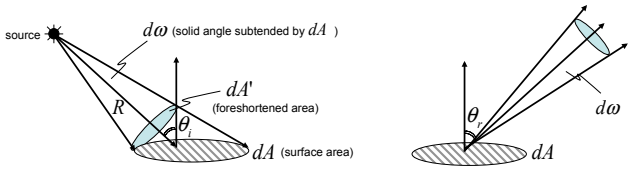


Image Quantization discretizes scene radiance at each pixel into a "brightness" or color value

Radiometric Terms



Surface Irradiance

$$E = \frac{d\Phi}{dA} \quad (\text{watts / m}^2)$$

- Light Flux (power) incident per unit surface area coming from a hemisphere of directions
- Does not depend on where the light is coming from

Surface Radiance

$$L = \frac{d^2\Phi}{(dA\cos\theta_r)\, d\omega} \quad (\text{watts / m}^2\text{-steradian})$$

- Flux emitted per unit foreshortened area per unit solid angle
- L depends on direction θ_r
- Surface can radiate into whole hemisphere.
- L depends on reflectance properties of surface.

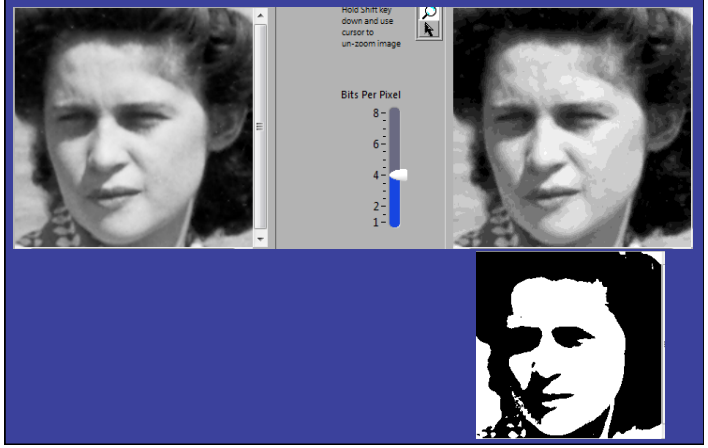
Photometric Terms

- **Photometry** is the measurement of light as detectable by the human eye
- Just like radiometry except weighted by the spectral response of the eye
- **Luminance** is the analog of **radiance**
  - Measured in lumens/m^2-steradian (= nit)
- **Illuminance** is the analog of **irradiance**
  - Measured in lumens/m^2 (= lux)

Digital Image Quantization

- A well-exposed photograph has a histogram that has values close to 0 near the minimum and maximum brightness values so as not to lose information
  - No saturated (over-exposed) regions
  - No dark (under-exposed) regions
- Brightness values representing smoothly changing radiance should not be noticeable
  - Too few gray levels leads to false contours in areas of the image where brightness changes slowly
- Most digital cameras represent brightness by 8 bits per pixel, or 8 bits per color (R,G,B)

False Contours



Dynamic Range

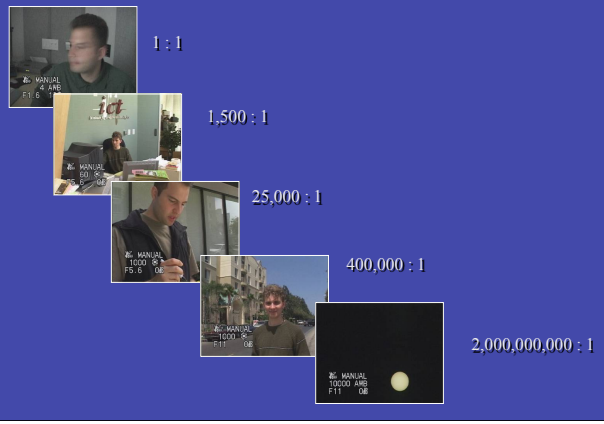
Dynamic range is a measure of the ***contrast*** in the scene, i.e., the ratio between the brightest area and darkest area



Real-World Scene Dynamic Range is Often Large

- Some luminance levels of real scenes
  - Starlight 10^{-3} cd/m^2
  - Moonlight 10^{-1}
  - Indoor lighting 10^{2}
  - Sunlight 10^{5}
- Contrast ratio often 10,000 : 1 in a scene

Real-World Scene Dynamic Range is Often Large



Dynamic Range of Various Display Devices

- LCD display 700 : 1
- Print film 128 : 1
- Color negative 256 : 1
- Positive slide 4,096 : 1
- Human eye 100 : 1 static dynamic range
- Human eye 1,000,000 : 1 dynamic range by adapting exposure geometrically and chemically to allow a sensitivity of $\sim 10^{-6}$ to 10^{6} cd/m^2

Dynamic Range of Image Sensors

- Ratio of maximum possible signal ("full well capacity") and total noise signal in the dark
- Common CCD sensors have dynamic range around 4,000 : 1, requiring 12 or 13 bits per pixel
- A/D converters have conversion uncertainty that reduces the usable dynamic range by ~1 bit
- High-end CCDs have larger dynamic range
- CMOS sensors have lower dynamic range

- Bottom line: **Image sensor dynamic range is not high enough to capture high dynamic range scenes**

Image Dynamic Range is Too Small

- Today's cameras: limited dynamic range (LDR)



High Exposure Image

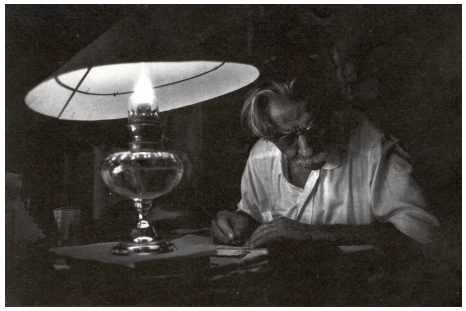


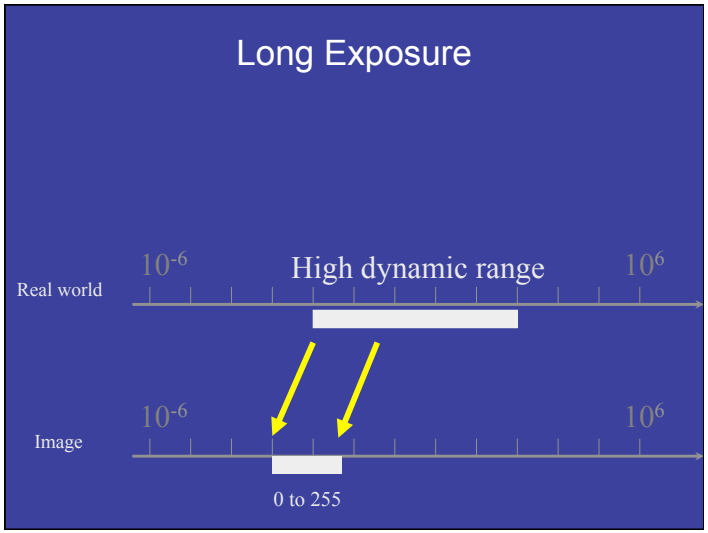
Low Exposure Image

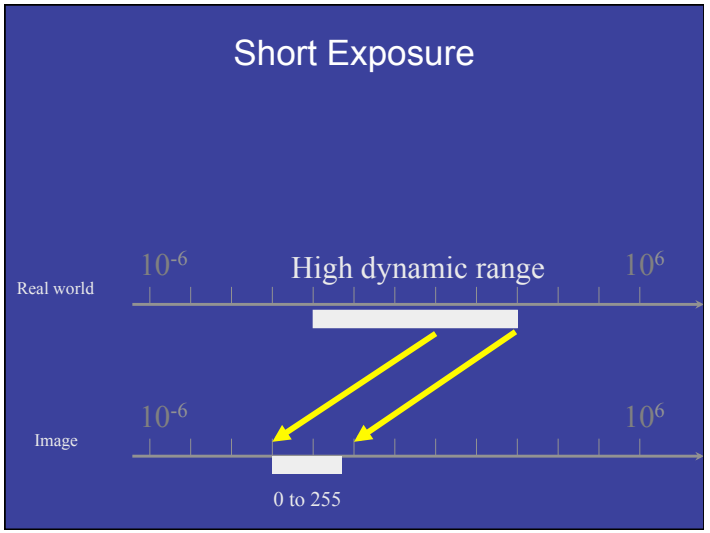
- We need 5-10 million values to store all brightnesses around us
- But, typical 8-bit cameras provide only 256 values

Method 1: Limit Dynamic Range

- W. Eugene Smith photo of Albert Schweitzer
- Overexpose bright areas
- Correctly expose dark areas
- 5 days to print

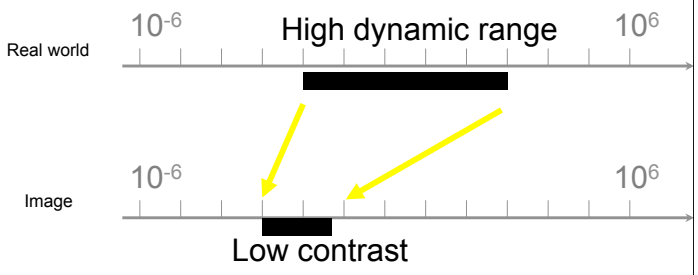






Method 2: Contrast Reduction

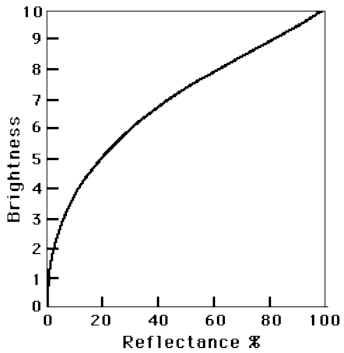
- Match limited contrast of the medium
- Preserve details



How Humans Deal with Dynamic Range

- We're sensitive to contrast (multiplicative)
  - A ratio of 2 : 1 is perceived as the same contrast as a ratio of 200 : 100
  - Illumination has a multiplicative effect
  - Use the log domain as much as possible
- Dynamic adaptation (very local in retina)
  - Pupil
  - Neural
  - Chemical
- Different sensitivity to different spatial frequencies

Perceived Brightness is Non-Linear



The perceived brightness of light diffusely reflected from a surface is not a linear function of the actual reflected light. Psychophysical experiments indicate that the perceived brightness increases approximately as the logarithm of the luminous flux. With a doubling of the flux leads to an increase of about 1.5 units on the arbitrary brightness scale at left.

Contrast Reduction: Fill-in Flash

- Use a flash to reduce contrast



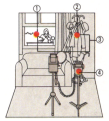
Exposure for outside



Exposure for inside



Average exposure



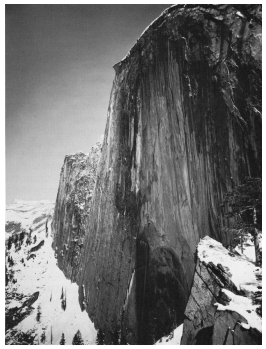


Using fill flash

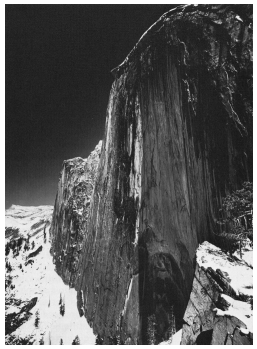
From Le Livre de la Photo Couleur (Larousse)

Filtering: Black and White

Red/orange/yellow filters darken the sky



No filter



With red filter

Source: Ansel Adams

Graduated Neutral Density Filtering

- Art Wolfe: “In the late evening light, I composed this image using a graduated neutral-density filter to bring the overall exposure into alignment, thus preserving the detail in the clouds in the sky and the reflections on the water.”

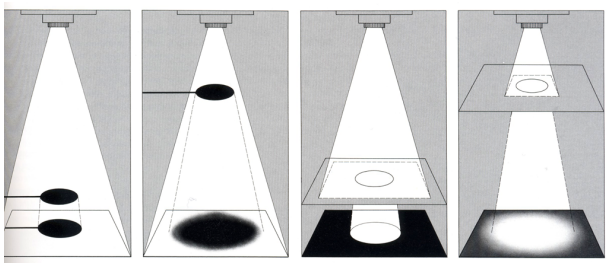


Portage Lake, Horizontal

http://www.artwolfe.com/

Dodging and Burning

- During print making process
- Hide part of the print during exposure
  - **Manually** select areas to increase or decrease exposure



From The Master Printing Course, Rudman

Dodging and Burning

- Must be done for every single print!



Straight print

After dodging and burning

High Dynamic Range Digital Imaging

Idea: Take multiple photos to cover the full dynamic range, then combine best parts from each photo



HDR Examples



John Adams



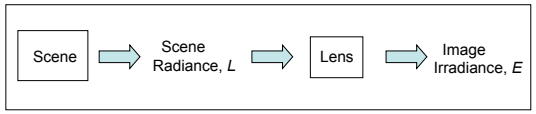
John Adams



Alberto Carrozzo

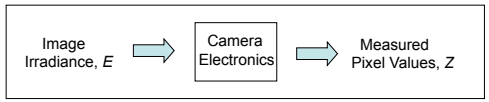
Relationship between Scene Radiance and Image Brightness

- Before light hits the image plane:



Assume known mapping

- After light hits the image plane:



Non-linear mapping

Can we go from measured pixel value, *Z*, to scene radiance, *L*?

In-Camera Digital Pipeline

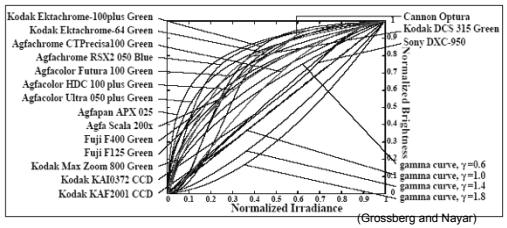
- Photosites transform photons into charge (electrons)
  - Sensor itself is linear
- Then goes through analog-to-digital (A/D) converter
  - Usually 12 or 14 bits/channel
- Stop here when shooting in RAW mode
- Then image processing and a “response curve” are applied
- Quantized and saved as 8-bit JPEG

Relation between Pixel Value, *Z*, and Image Irradiance, *E*

Image Irradiance *E* ⇒ Camera Electronics ⇒ Measured Pixel Values, *Z*

The **Camera Response Function** relates **image irradiance**, *E*, at the image plane to the measured **pixel intensity** value, *Z*

$$g : E \rightarrow Z$$



(Grossberg and Nayar)

***g* is monotonic and smooth for all cameras**

Camera is *Not* a Photometer

- Limited dynamic range
  - Use multiple exposures
- Unknown, nonlinear response
  - Difficult to convert pixel values directly to radiance
- Solution:
  - Recover *radiometric response curve* from multiple exposures, then reconstruct the *radiance map*

High Dynamic Range (HDR) Imaging

1. Capture multiple (usually 3 or 5) images with different exposure settings
2. Estimate the **Camera Response Function**
3. Estimate **radiance map**: for each pixel, combine the calibrated images (for example, by weighted averaging)
4. **Tone map** the image into a displayable range



(Mitsunaga)

Ways to Vary Exposure

- Shutter speed
- F-stop (aperture)
- Neutral Density (ND) filters
- ISO / ASA



Exposure $X = E\,\Delta t$ = *irradiance* × *time*
⇒ Halving *E* and doubling Δ*t* will *not* change exposure – ***reciprocity* property**



Exposure Value (EV)

- Combinations of shutter speeds and apertures that give the ***same exposure*** (= irradiance * time) (reciprocity property)

• 1/1000	@	f/1.4
• 1/500	@	f/2
• 1/250	@	f/2.8
• 1/125	@	f/4
• 1/60	@	f/5.6
• 1/30	@	f/8
• 1/15	@	f/11
• 1/8	@	f/16
• 1/4	@	f/22

• Each increment in speed or aperture is a "stop"

• Similarly, doubling the ISO value means it is twice as sensitive, so
1/125 @ ISO 400 =
1/250 @ ISO 800

How to Best Vary Exposure?

- Varying aperture changes the depth of field and can cause vignetting
- Varying ISO changes the "graininess"
- ***Varying shutter speed is best***

- Assume scene is static, camera is static, and lighting is static, so all images will be in register

Shutter Speed

- Ranges: Canon D30: 30 to 1/4,000 sec
  Sony VX2000: ¼ to 1/10,000 sec
- Pros:
  – Directly varies the exposure
  – Usually accurate and repeatable
- Issues:
  – Noise in very long and very short exposures

Step 1: Capture Images with Different Shutter Speeds (Varying Exposure)



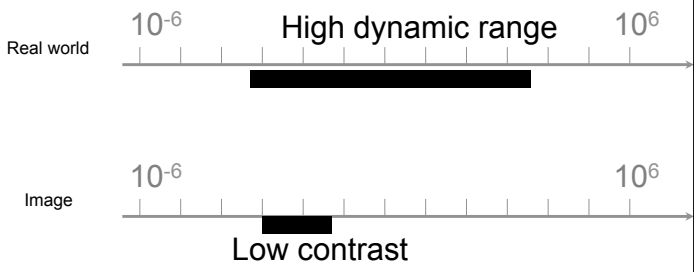
Assume scene is static, camera is static, and lighting is static, so all images are in register

Capturing a Set of Images

- Commonly, take 3 images at -2 EV, 0 EV, +2 EV, or 5 images at increments of 1 EV
- Many modern digital cameras have "**AEB mode**" meaning "**automatic exposure bracketing**," which automatically takes (usually) 3 photos at ±2 EV
- Expensive DSLR cameras can bracket 3, 5, 7, 9 images at various EV increments

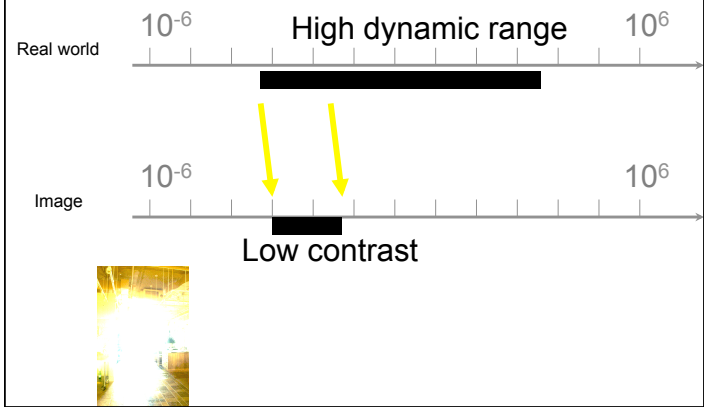
Multiple Exposure Photography

- Sequentially measure all segments of the range



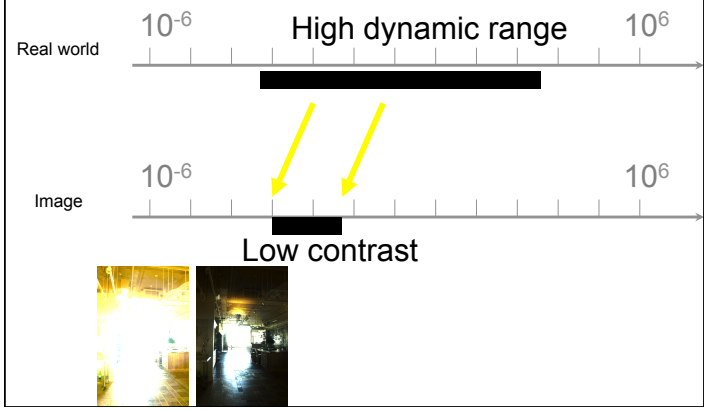
Multiple Exposure Photography

- Sequentially measure all segments of the range



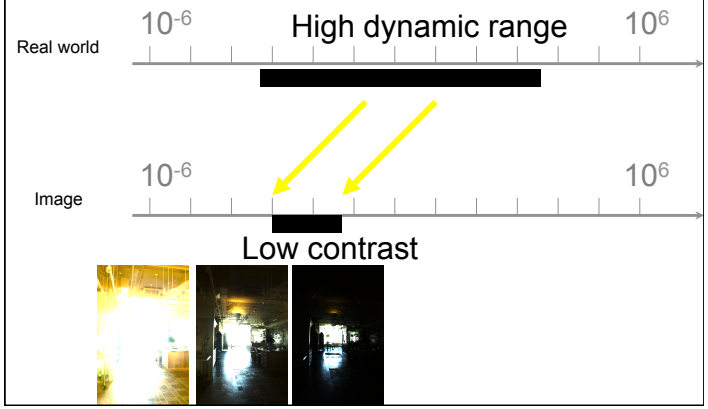
Multiple Exposure Photography

- Sequentially measure all segments of the range



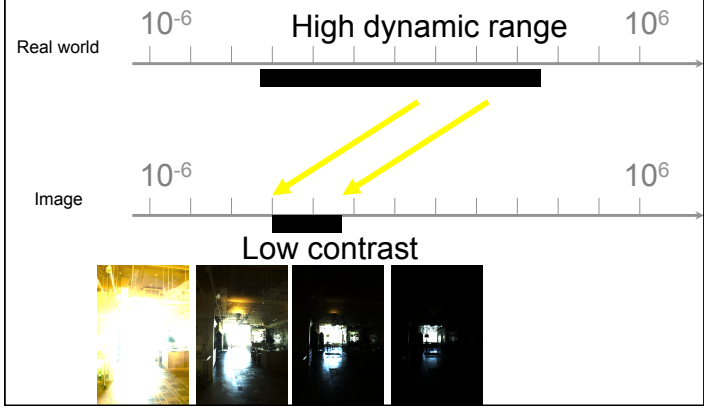
Multiple Exposure Photography

- Sequentially measure all segments of the range



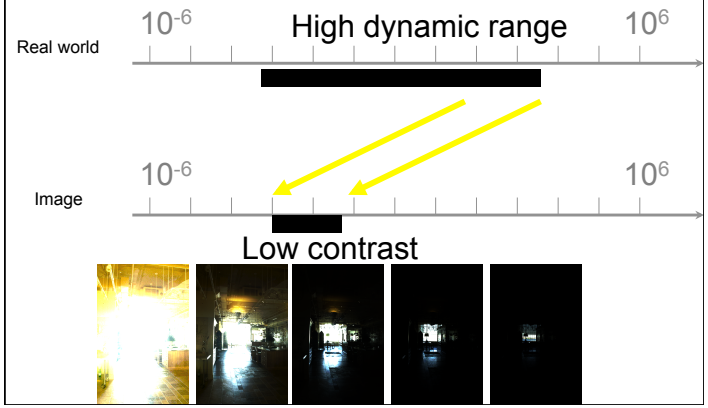
Multiple Exposure Photography

- Sequentially measure all segments of the range



Multiple Exposure Photography

- Sequentially measure all segments of the range



Step 2: Recover Camera Response Function

Registered images



Pixel value $Z = f(\text{Exposure } X)$

$X = \text{Irradiance } E \times \Delta t$

$\log X = \log E + \log \Delta t$

$g(Z) = \log f^{-1}(Z) = \log E + \log \Delta t$

Camera Response Function

- Let $g(Z)$ be the *discrete* inverse response function
- For each pixel i in image j, want:

$$\ln E_i + \ln \Delta t_j = g(Z_{ij})$$

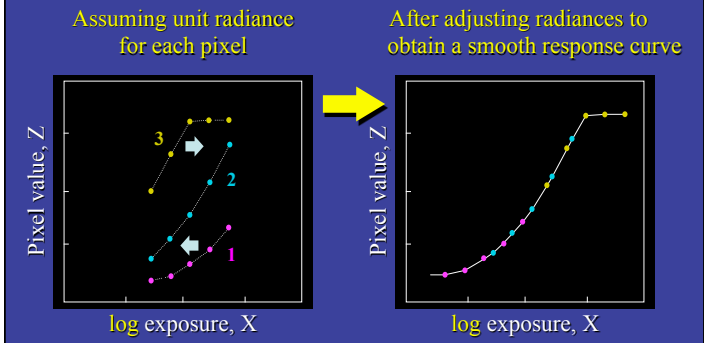
Known: Δt and Z
Unknown: g and E

- Solve the overdetermined linear system (using linear least squares):

$$\sum_{i=1}^{N} \sum_{j=1}^{P} \left[\ln E_i + \ln \Delta t_j - g(Z_{ij})\right]^2 + \lambda \sum_{z=Z_{min}}^{Z_{max}} g''(z)^2$$

fitting term — smoothness term

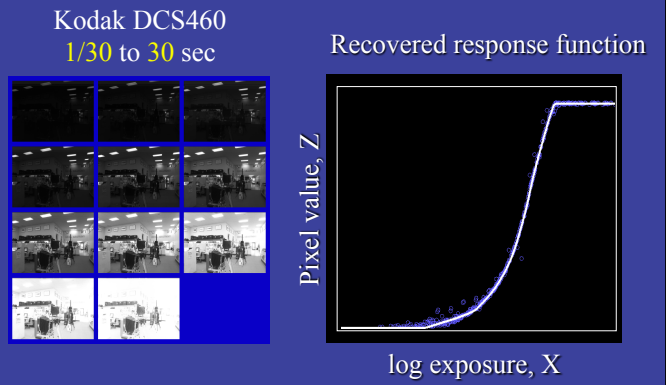
Camera Response Function



Camera Response Function Assumptions

- Minimal assumptions on g: monotonic and smooth
- Other methods make stronger assumptions on g:
  - Low-order polynomial (Mitsunaga and Nayar)
- Other methods simultaneously solve for Δt values

Results: Digital Camera

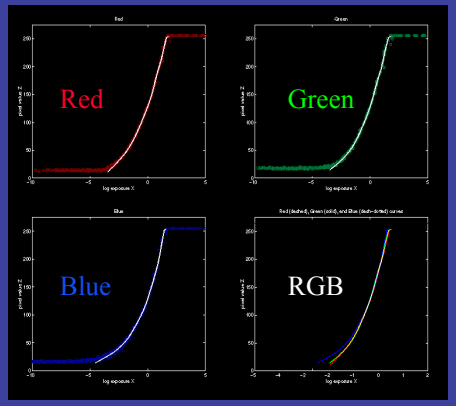


Results: Color Film

- Kodak Gold ASA 100, PhotoCD



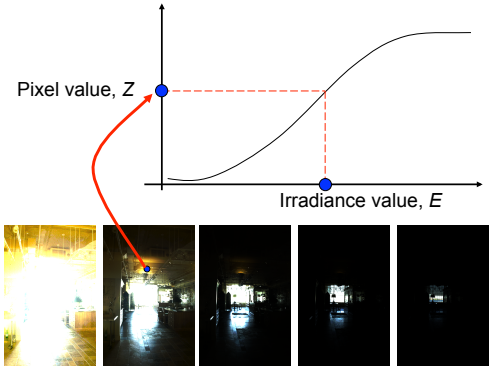
Recovered Response Functions



Recover 1 response function for each R, G, B channel

Step 3: Reconstruct the Radiance Map

(i.e., an image of radiance or irradiance values)



Relation between Scene Radiance and Image Irradiance

$$E = L\frac{\pi}{4}\left(\frac{d}{f}\right)^2 \cos^4\alpha$$

where E is image irradiance, L is scene radiance, d is lens diameter, f is focal length, and α is the angle from the optical axis to the ray to the pixel

Step 3: Reconstruct the Radiance Map

- Assuming a perfect response function and pixel values have no noise, just look up the value from the response function for a given pixel value:
  - For each pixel, *i*, in each image, *Z*, taken at shutter speed Δt_j
    - Map pixel value, Z_{ij}, to exposure value, X_i
    - Map exposure value to irradiance, $E_i = X_i / \Delta t_j$
    - Or, equivalently: $\ln E_i = g(Z_{ij}) - \ln \Delta t_j$
- Store each R, G, B value as a 32-bit float

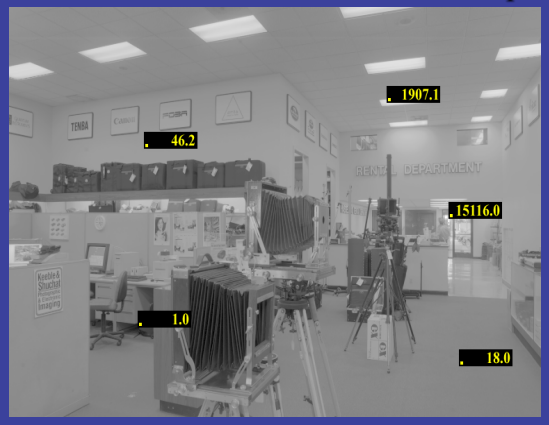
Reconstruct the Radiance Map

- Assuming noise, improve mapping by weighting pixels differently
  - E.g., weight pixels more that have high values and large gradients in the camera response function:

$$w(Z) = g(Z)/g'(Z)$$

$$\ln E_i = \frac{\sum_j w(Z_{ij}) g(Z_{ij}) - \ln \Delta t_j}{\sum_j w(Z_{ij})}$$

Reconstruct the Radiance Map



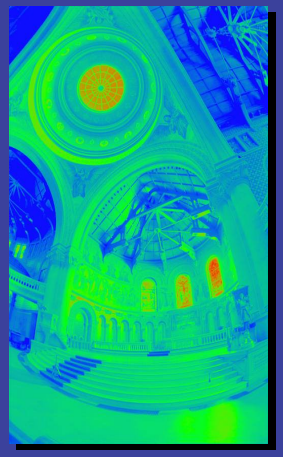
$E = g(Z)$

Log radiance values shown as intensity

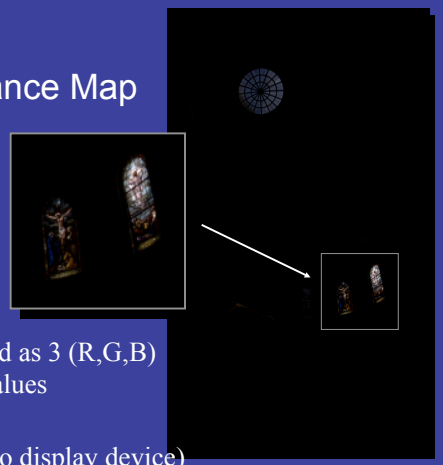
Radiance Map as a False-Color Image

Range of radiance values is nearly 250,000 : 1





The Radiance Map

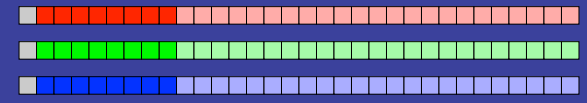


Pixel represented as 3 (R,G,B) floating point values

(linearly scaled to display device)

HDR Image Formats: Portable FloatMap (.pfm)

- 12 bytes / pixel, 4 for each channel



Text header similar to .ppm image format:

Floating Point TIFF similar

```
PF
768 512
1
<binary image data>
```

Radiance Format (.pic, .hdr)

4 bytes / pixel



(145, 215, 87, 149) =
(145, 215, 87) * 2^(149-128) =
(1190000, 1760000, 713000)

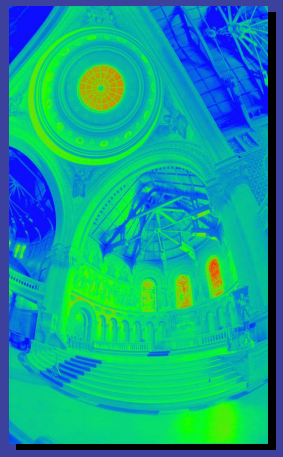
(145, 215, 87, 103) =
(145, 215, 87) * 2^(103-128) =
(0.00000432, 0.00000641, 0.00000259)

G. Ward, "Real Pixels," in *Graphics Gems IV*, J. Arvo, ed., 1994

Now What?

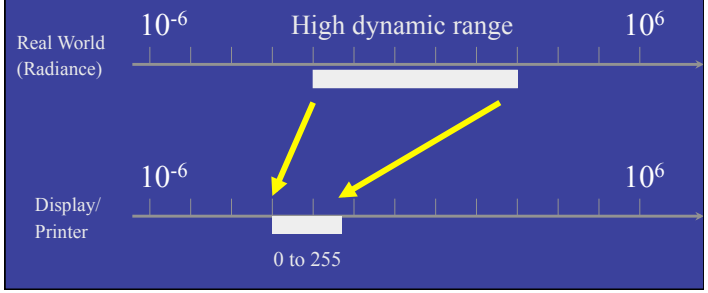
Display device doesn't usually have enough bits per pixel!





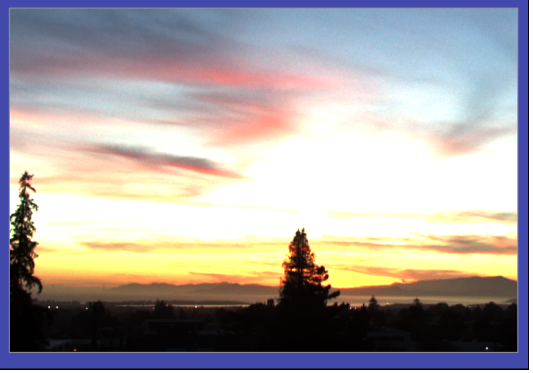
Step 4: Tone Mapping

- Goal: Approximate the appearance of an HDR image on an LDR display, i.e., **reduce the contrast but preserve the details**
- How can we do this?



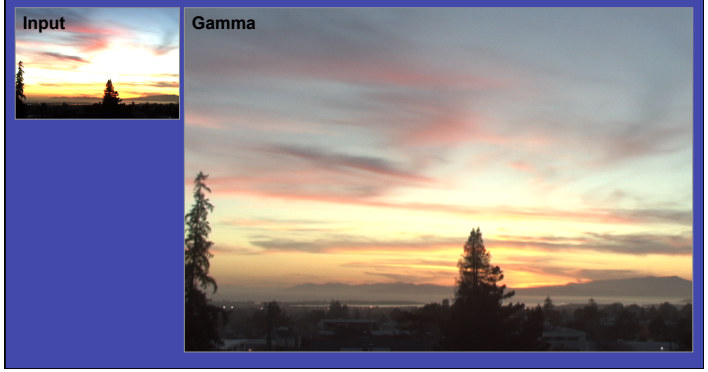
Tone Mapping

- Input: HDR image
  - floating point value for each color per pixel

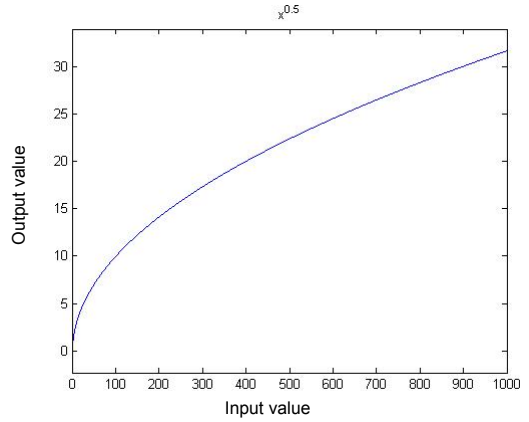


Global Operator: Gamma Compression

- $X \rightarrow X^{\gamma}$ for each color channel (e.g., $\gamma = 0.5$)
- Causes colors to wash-out (less saturated) because high values weakened more

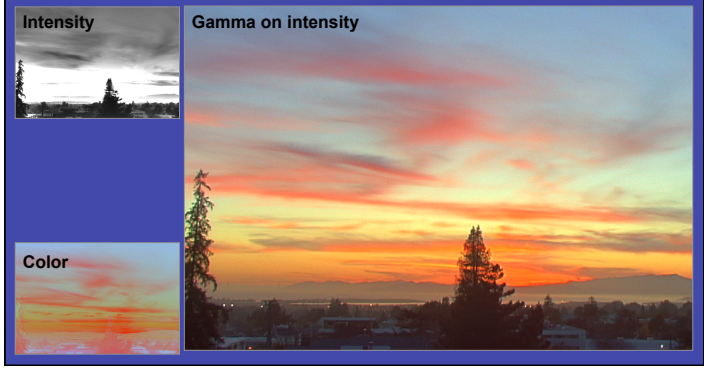


Gamma Compression



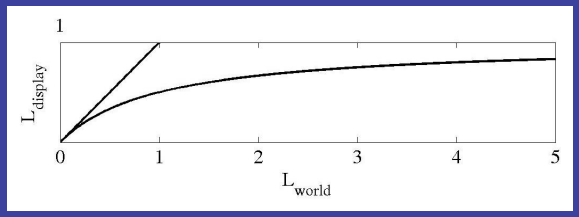
Gamma Compression on Intensity Only

- Convert RGB to L*a*b* (luminance plus 2 chrominance components)
- Gamma correction applied to L channel *only*
- Colors are OK, but details (intensity high-frequencies) are blurred



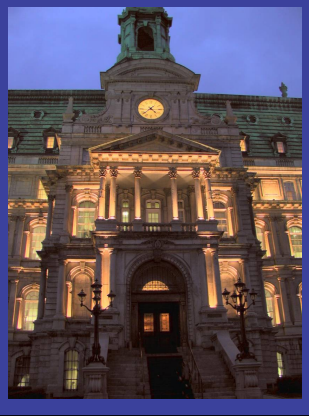
Another Global Operator (Reinhart *et al.*)

$$L_{display} = \frac{L_{world}}{1 + L_{world}}$$

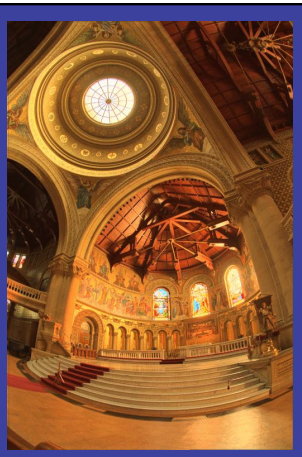


- Brings everything within range
- Leaves dark areas alone

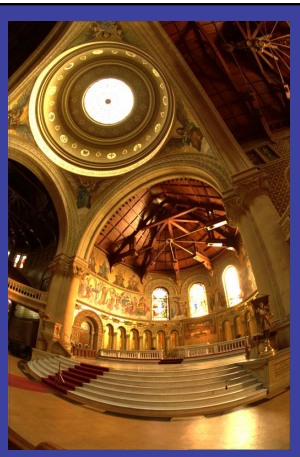
Reinhart Global Operator Results







Reinhart Operator



Darkest 0.1% scaled to display device

Problems with Global Operators

- Fails to preserve details in regions with widely varying exposures
- They don't adaptively attenuate or brighten different regions of the image
- What is needed is a kind of **local dodge and burn** process that compares each pixel to the average brightness in a region around the pixel in order to preserve the local contrast

Local Tone Mapping



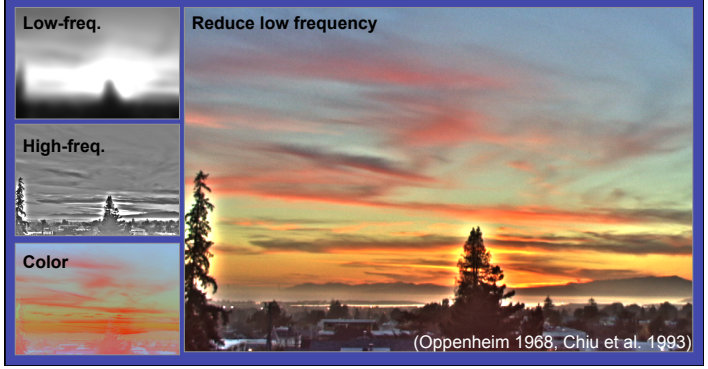
Global contrast reduction



Local tone mapping

Local Tone Mapping

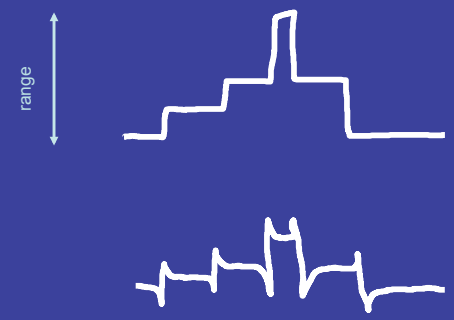
- Split Luminance into large-scale structure and small-scale texture
- Reduce contrast of large-scale features (low-frequencies) *only*
- Keep small-scale features (high frequencies)



Local Tone Mapping using Linear Filters

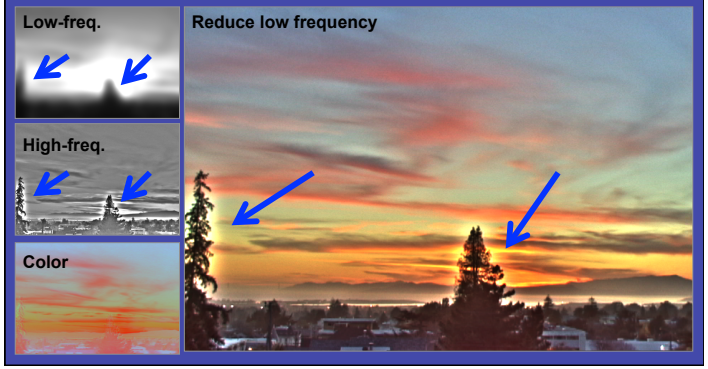
1. Separate luminance and chrominance channels
2. Compute log luminance image, $H = \log L$
3. Low-pass filter (i.e., blur) log luminance image, $HL = H * G$
4. Compute high-pass filtered image from log luminance image, $HH = H - HL$
5. Reduce contrast of low-pass filtered image, $HL' = s * HL$
6. Add low-pass and high-pass images, exponentiate, and then add back chrominance image

Compressing Dynamic Range



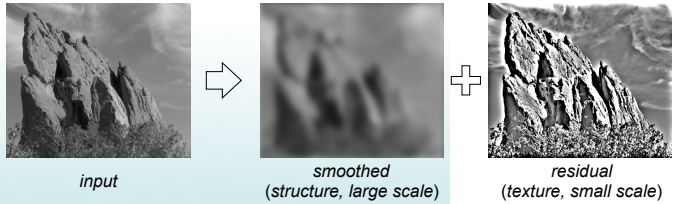
The Halo Problem

"Halos" occur at strong edges because they contain high frequencies



Gaussian Blur

Gaussian filtering is not edge preserving, which causes decomposition to contain "halos" in high-frequency regions



Gaussian Convolution

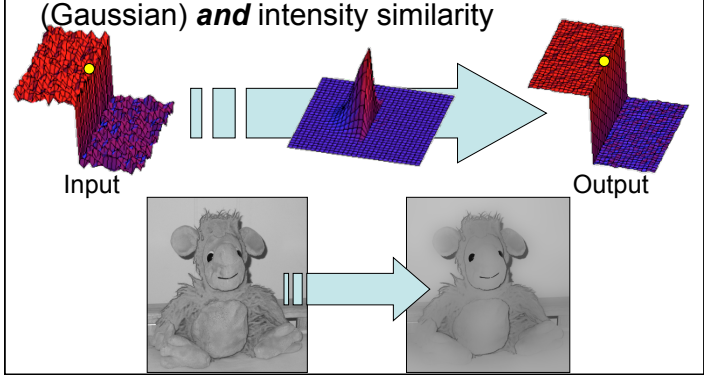
Bilateral Filter Method (Durand and Dorsey, 2002)

- Do *not* blur across edges
- Non-linear, edge-preserving filtering
- No visible halos



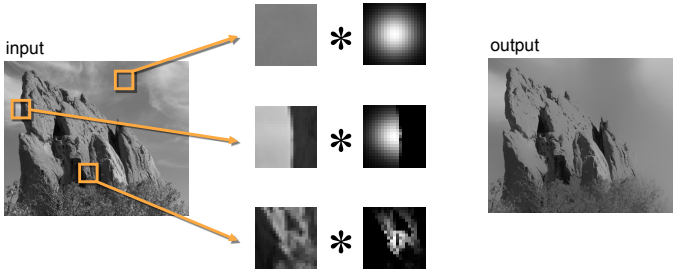
Coarse-Scale Layer

Bilateral filter – edge-preserving filter that weights neighboring pixels by their distance (Gaussian) ***and*** intensity similarity



Bilateral Filter: No Averaging across Edges

[Aurich 95, Smith 97, Tomasi 98]



The kernel shape depends on the image content

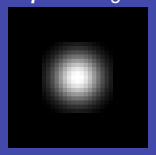
Bilateral Filter Definition: An Additional Edge Term

Idea: **weighted average of pixels around *p***

$$BF[I]_{\mathbf{p}} = \underbrace{\frac{1}{W_{\mathbf{p}}}}_{\text{new}} \sum_{\mathbf{q}\in S} \underbrace{G_{\sigma_s}(\|\mathbf{p}-\mathbf{q}\|)}_{\text{not new}} \underbrace{G_{\sigma_r}(|I_{\mathbf{p}}-I_{\mathbf{q}}|)}_{\text{new}} I_{\mathbf{q}}$$

normalization factor

space weight

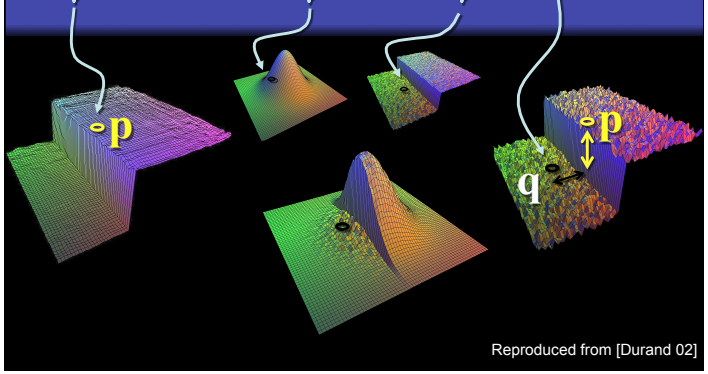


range weight



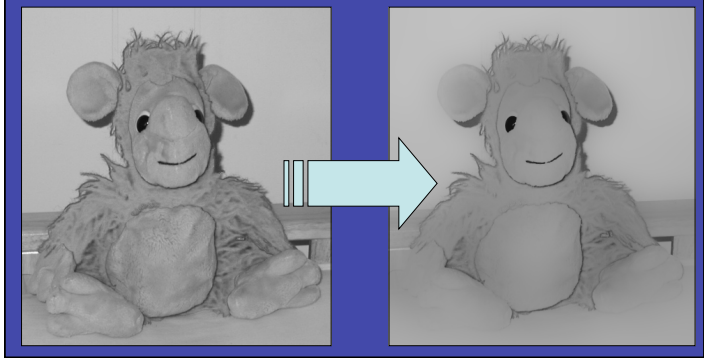
Bilateral Filter on a Height Field

$$BF[I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q}\in S} G_{\sigma_s}(\|\mathbf{p}-\mathbf{q}\|)\; G_{\sigma_r}(|I_{\mathbf{p}}-I_{\mathbf{q}}|)\; I_{\mathbf{q}}$$

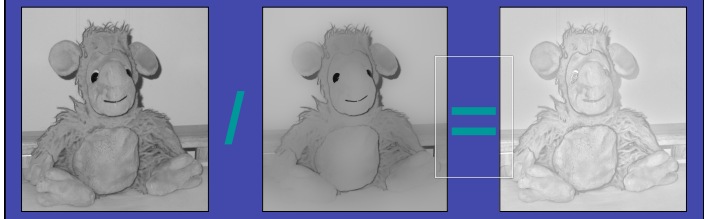


Reproduced from [Durand 02]

Coarse-Scale Layer of Bilateral Filter

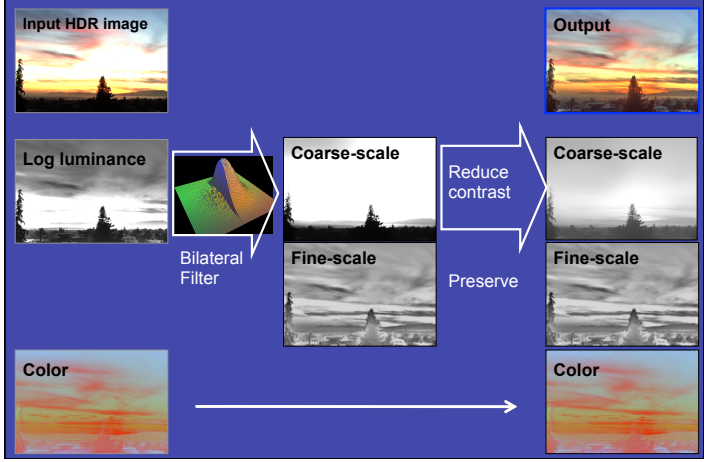


Fine-Scale Layer of Bilateral Filter



Recombination: Coarse-scale * Detail = Intensity

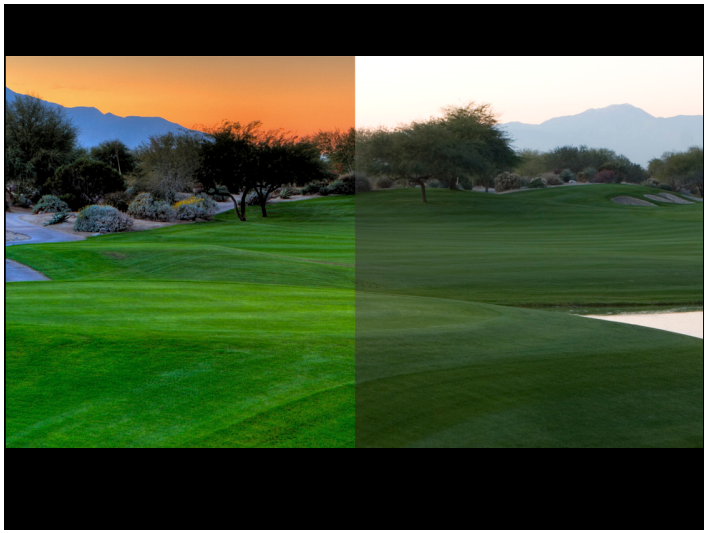
Bilateral Filter



Results









HDR Software

- Photoshop
- Photomatix
- HDRShop
- and many more

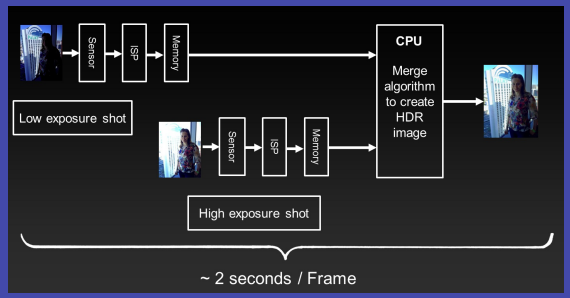
Photomatix
HDR Imaging for Photography

Photoshop

HDR SHOP

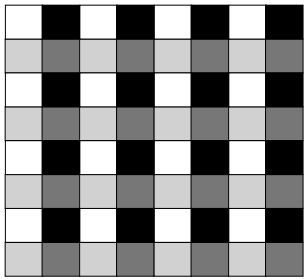
In-Camera HDR Hardware

- Conventional In-Camera HDR: Consecutive shots + merge



- 2 – 5 sec to capture, merge and create HDR images on smartphones ⇒ motion blur and ghosting artifacts
- Note: ISP = Image Signal Processor

HDR Image Sensors: Spatially Vary Exposure





Base Pattern of Neutral Density Filters

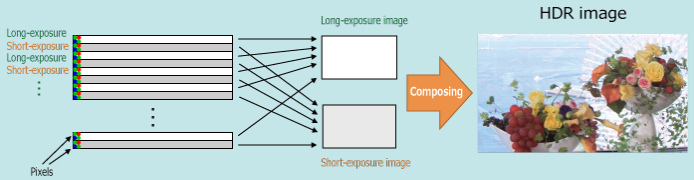
Spatially vary exposure in 2 x 2 block of pixels, then combine 4 pixels into 1

Nayar et al., 2000

HDR Image Sensors

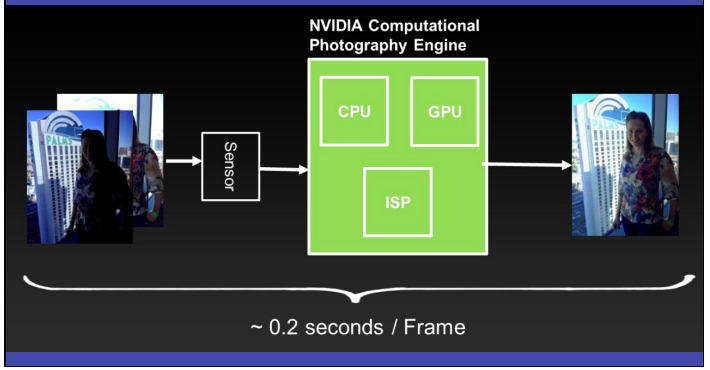
Toshiba Single-Frame HDR

– Captures alternate lines of an image with different exposure times and composes them into a single HDR image



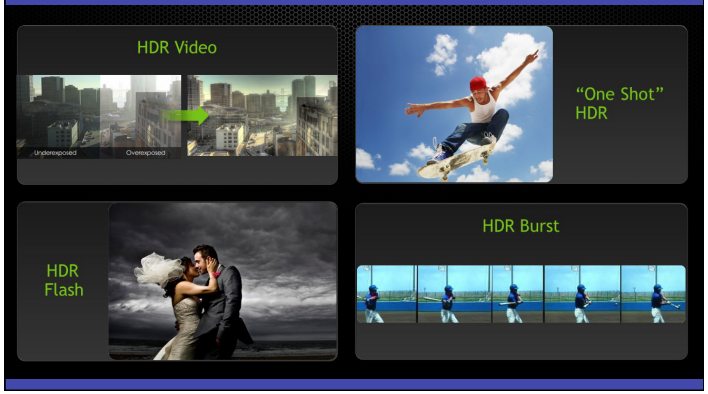
NVIDIA Chimera Computational Photography Architecture

"One Shot HDR" captures 2 images at "almost the same time"



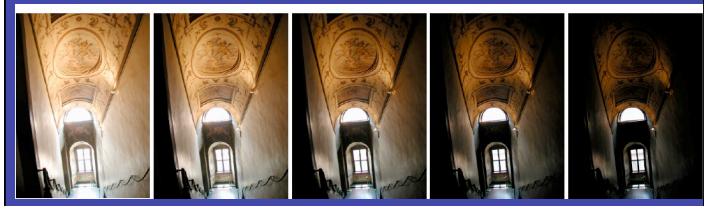
"Always On" HDR in Mobile Devices

Live preview, no motion blur, can use flash, 30 fps HDR video



Better HDR Processing

- Image registration (no need for tripod)
- Lens flare removal
- Ghost removal



Images by Greg Ward

HDR Video

Alternate frames with short and long exposures

– Kang et al., *Proc. SIGGRAPH* 2003



gure 1: High dynamic range video of a driving scene. *Top row: Input video with alternating short and long exposures. Bottom row: H*
namic range video (tonemapped).

HDR Video

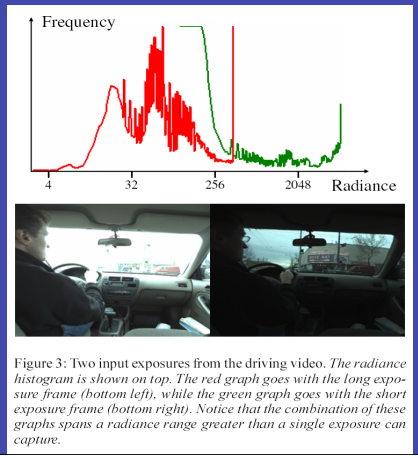


Figure 3: Two input exposures from the driving video. *The radiance histogram is shown on top. The red graph goes with the long exposure frame (bottom left), while the green graph goes with the short exposure frame (bottom right). Notice that the combination of these graphs spans a radiance range greater than a single exposure can capture.*