High-Dynamic-Range Imaging and Tonemapping

Kari Pulli
Senior Director
NVIDIA Research

Problems with setting the camera exposure level

- Under-exposed
  - Highlight details captured
  - Shadow details lost
- Over-exposed
  - Highlight details lost
  - Shadow details captured

Dynamic range

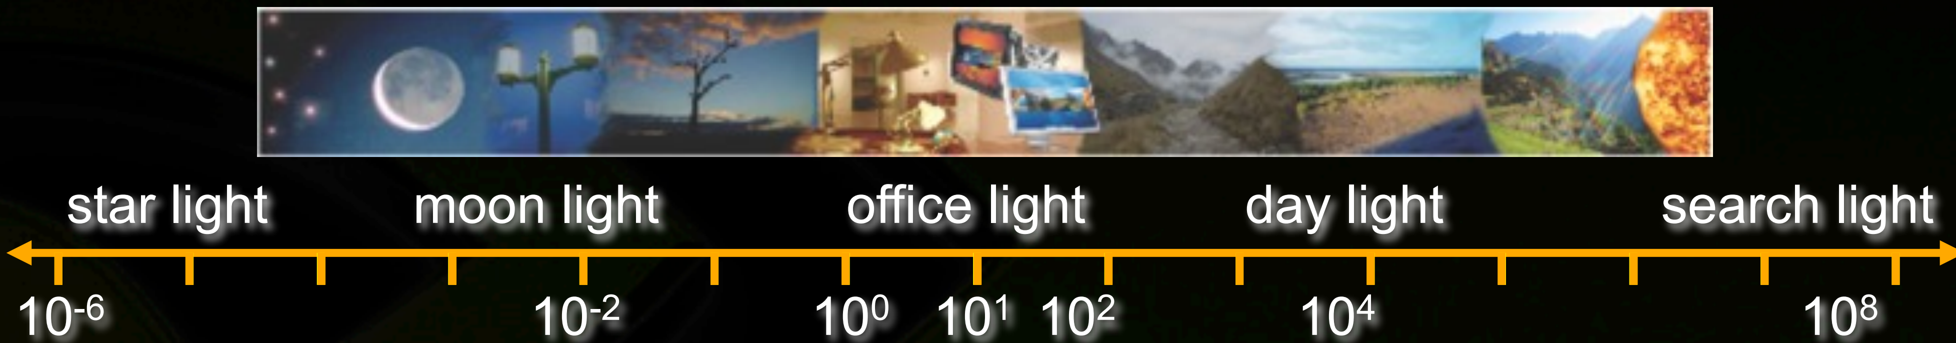

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2

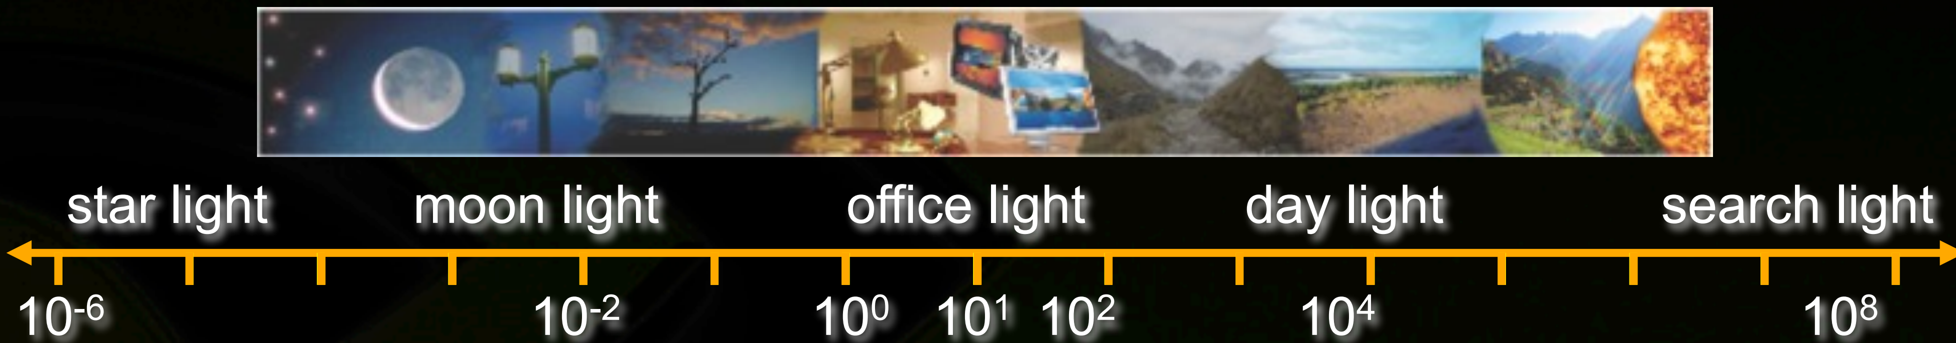

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2

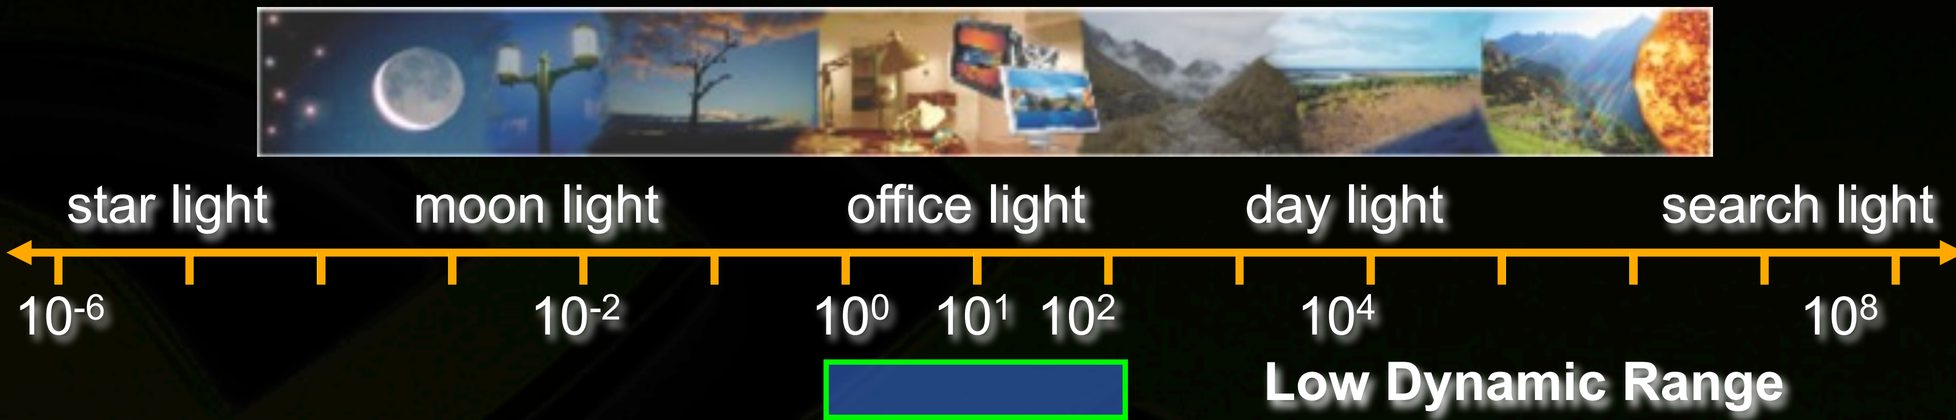

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2

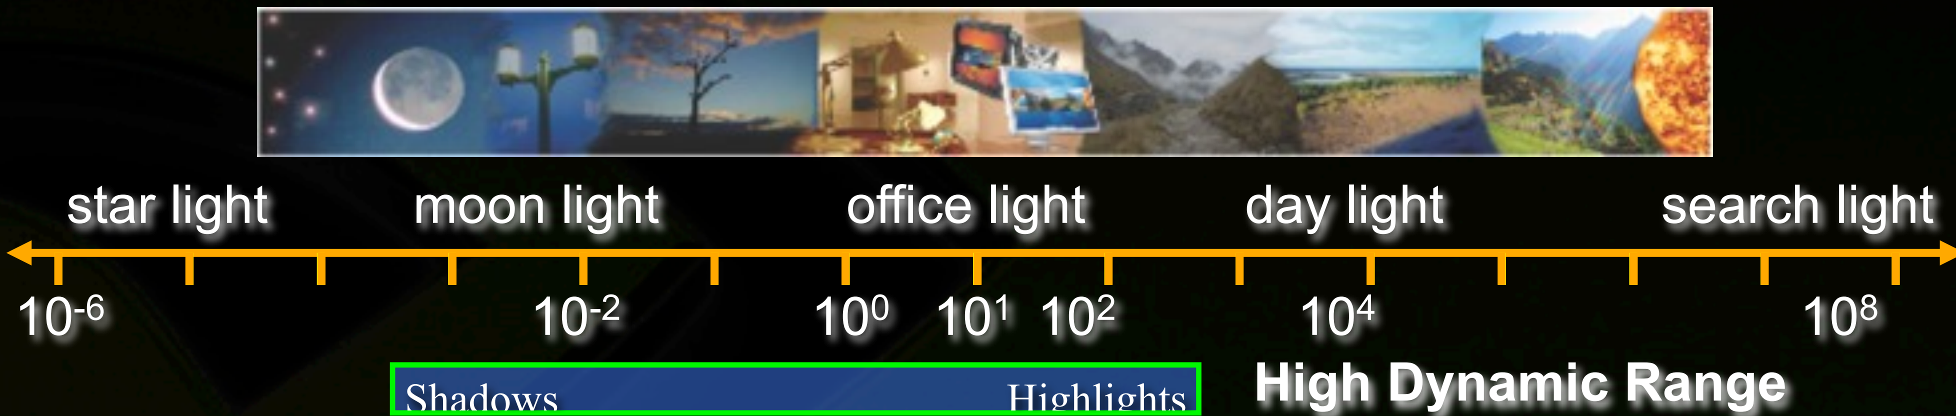

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2
- Sometimes 1 : 100,000 in a scene

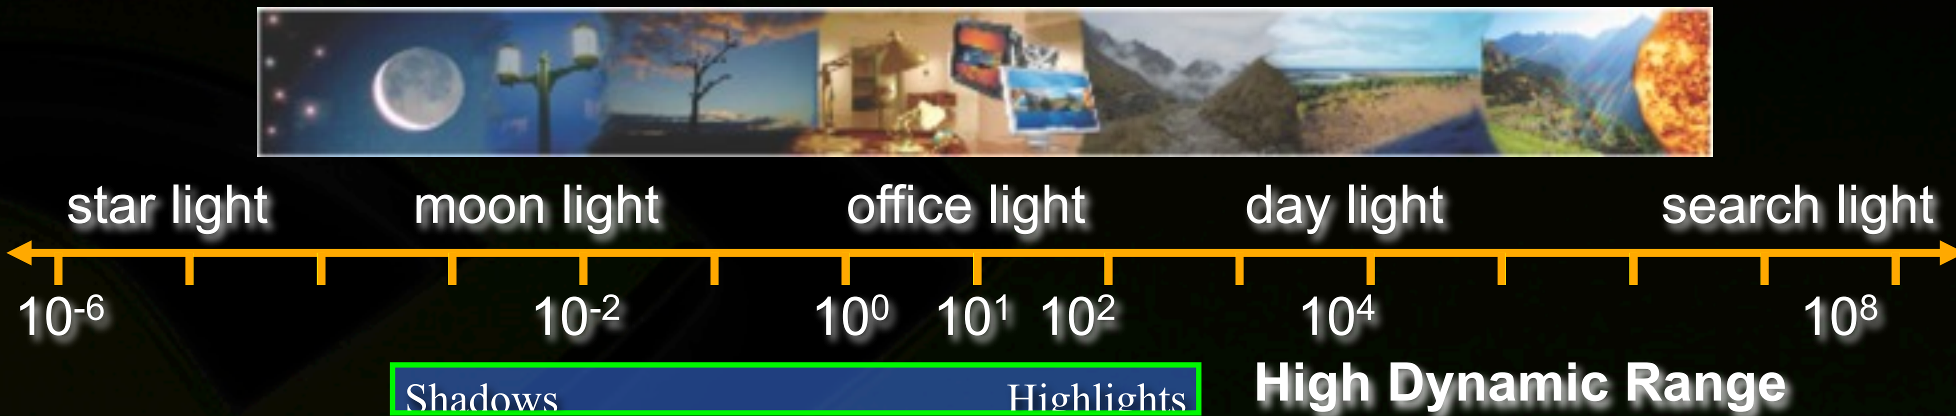

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2
- Sometimes 1 : 100,000 in a scene

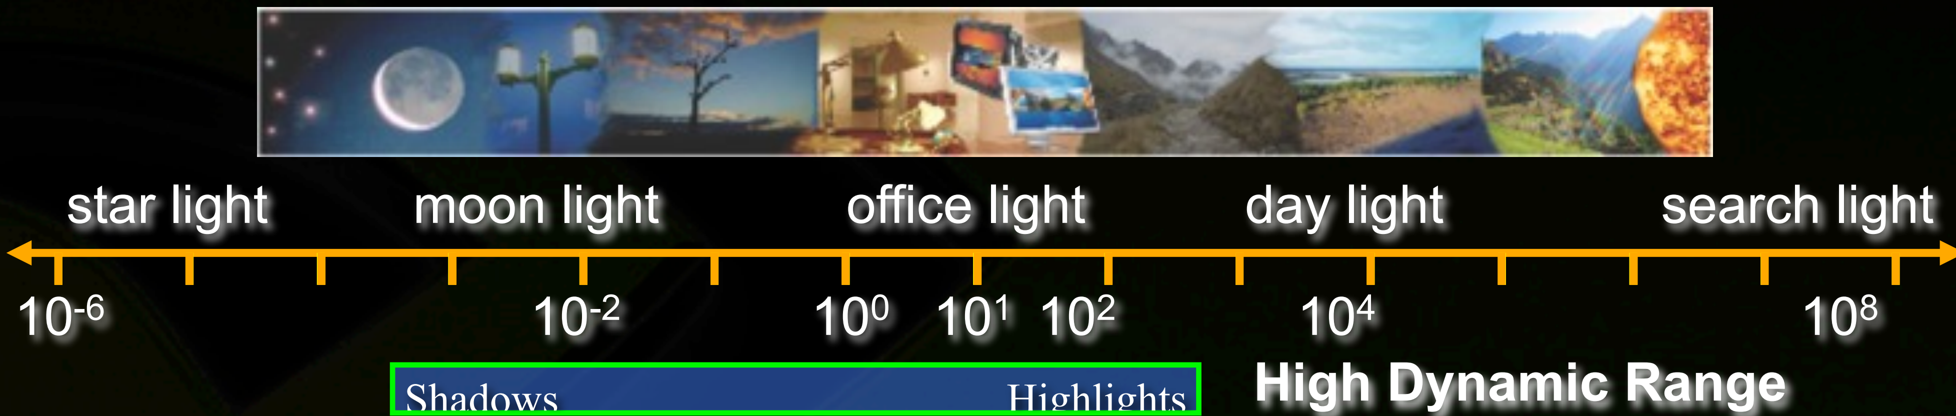

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2
- Sometimes 1 : 100,000 in a scene

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2
- Sometimes 1 : 100,000 in a scene

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2
- Sometimes 1 : 100,000 in a scene

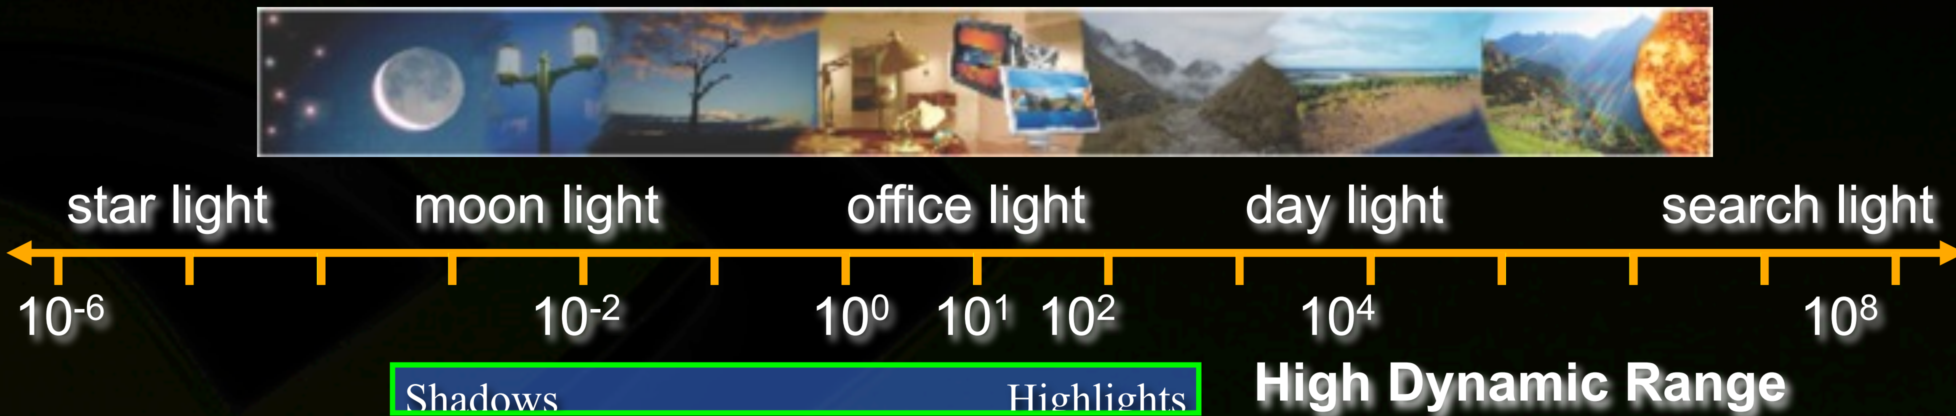

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2
- Sometimes 1 : 100,000 in a scene

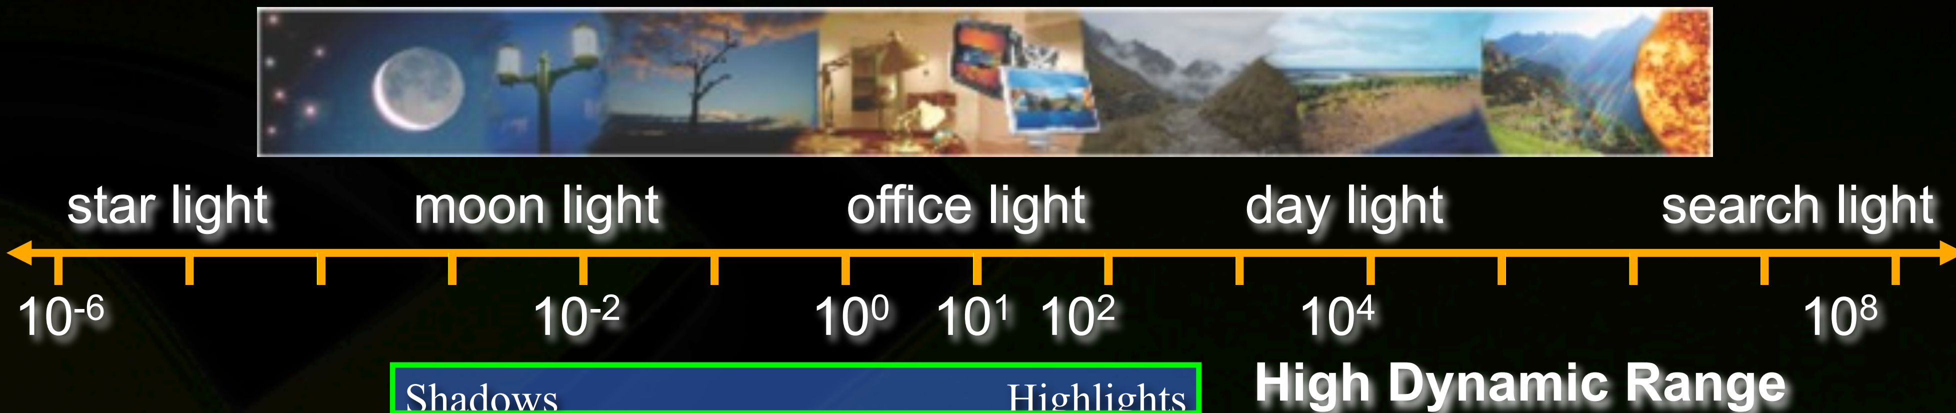

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m²
- Sometimes 1 : 100,000 in a scene

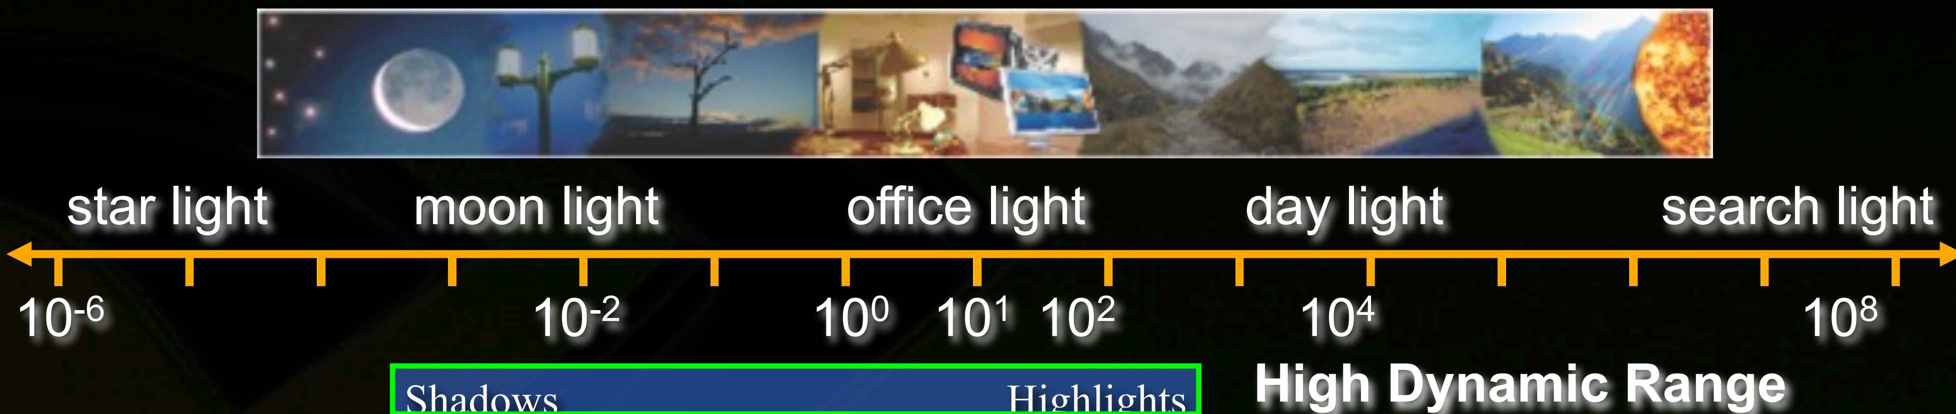

- Without adaptation eye can handle about 1 : 10000
  - Scotopic Mesopic Photopic
  - Even 1 : 1000 easily enough for scenes with non-specular reflectance

Dynamic range

- Eye can adapt from ~ 10^{-6} to 10^{8} cd/m^2
- Sometimes 1 : 100,000 in a scene

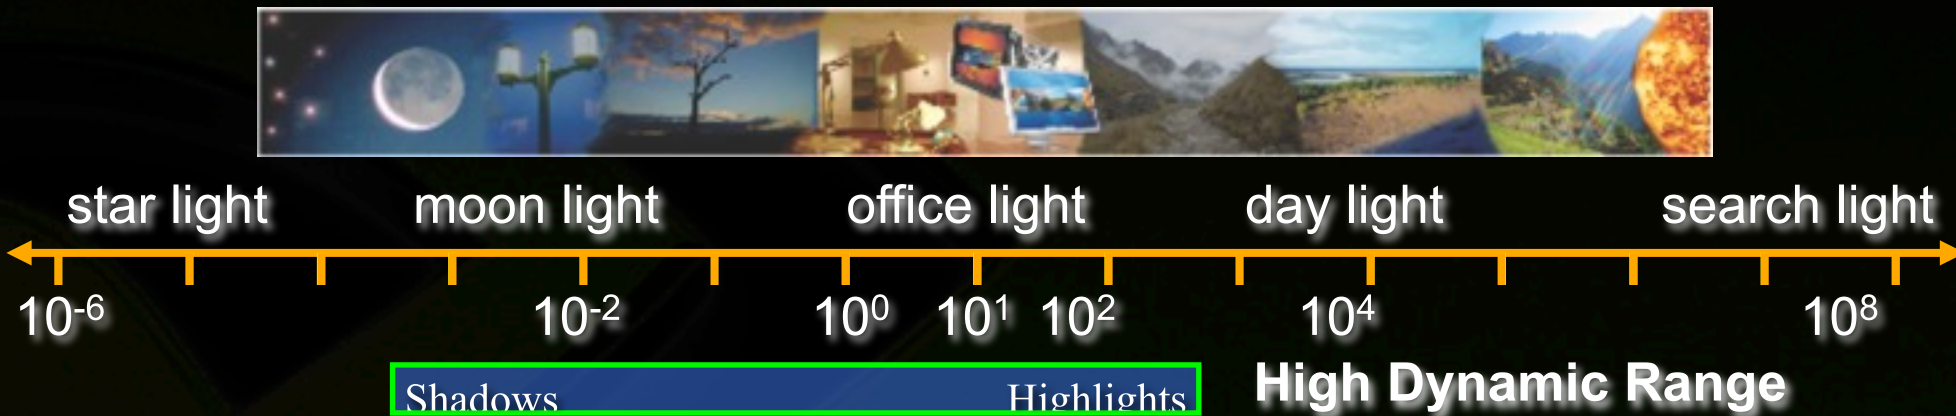

- Without adaptation eye can handle about 1 : 10000
  - Scotopic Mesopic Photopic
  - Even 1 : 1000 easily enough for scenes with non-specular reflectance
- Most displays can handle less than 1 : 100

Range of Typical Displays:

from ~1 to ~100 cd/m^2

Dynamic Range = Contrast Ratio

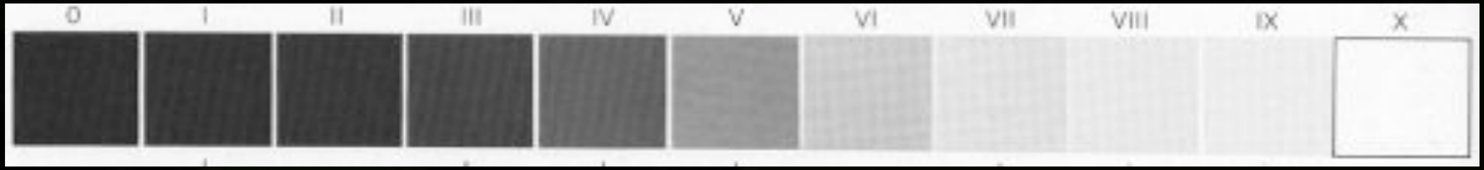

Dynamic Range = Contrast Ratio

- Depends on the overall brightness range and the smallest visible step
  - the dynamic range of the image below is 10:1

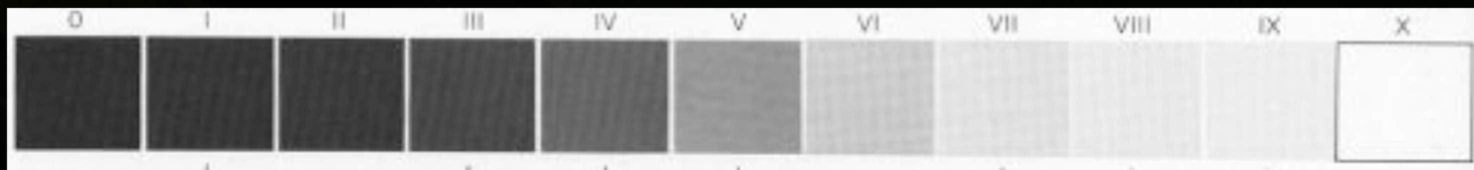

  - we could increase the dynamic range by either
    - using smaller steps
    - adding even brighter whites
    - adding even darker blacks

Dynamic Range = Contrast Ratio

- Depends on the overall brightness range and the smallest visible step
  - the dynamic range of the image below is 10:1

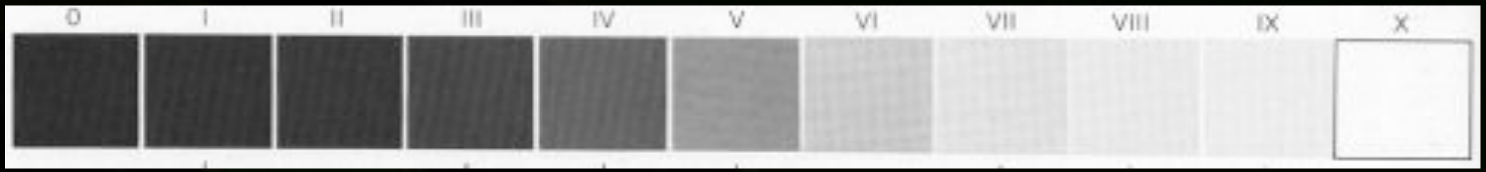

  - we could increase the dynamic range by either
    - using smaller steps
    - adding even brighter whites
    - adding even darker blacks
- Originally audio term
  - ratio of the max intensity over the base noise level

Dynamic range of print

- Typically 1:20 to 1:50
  - Black is ~50x darker than white
  - Max 1:500

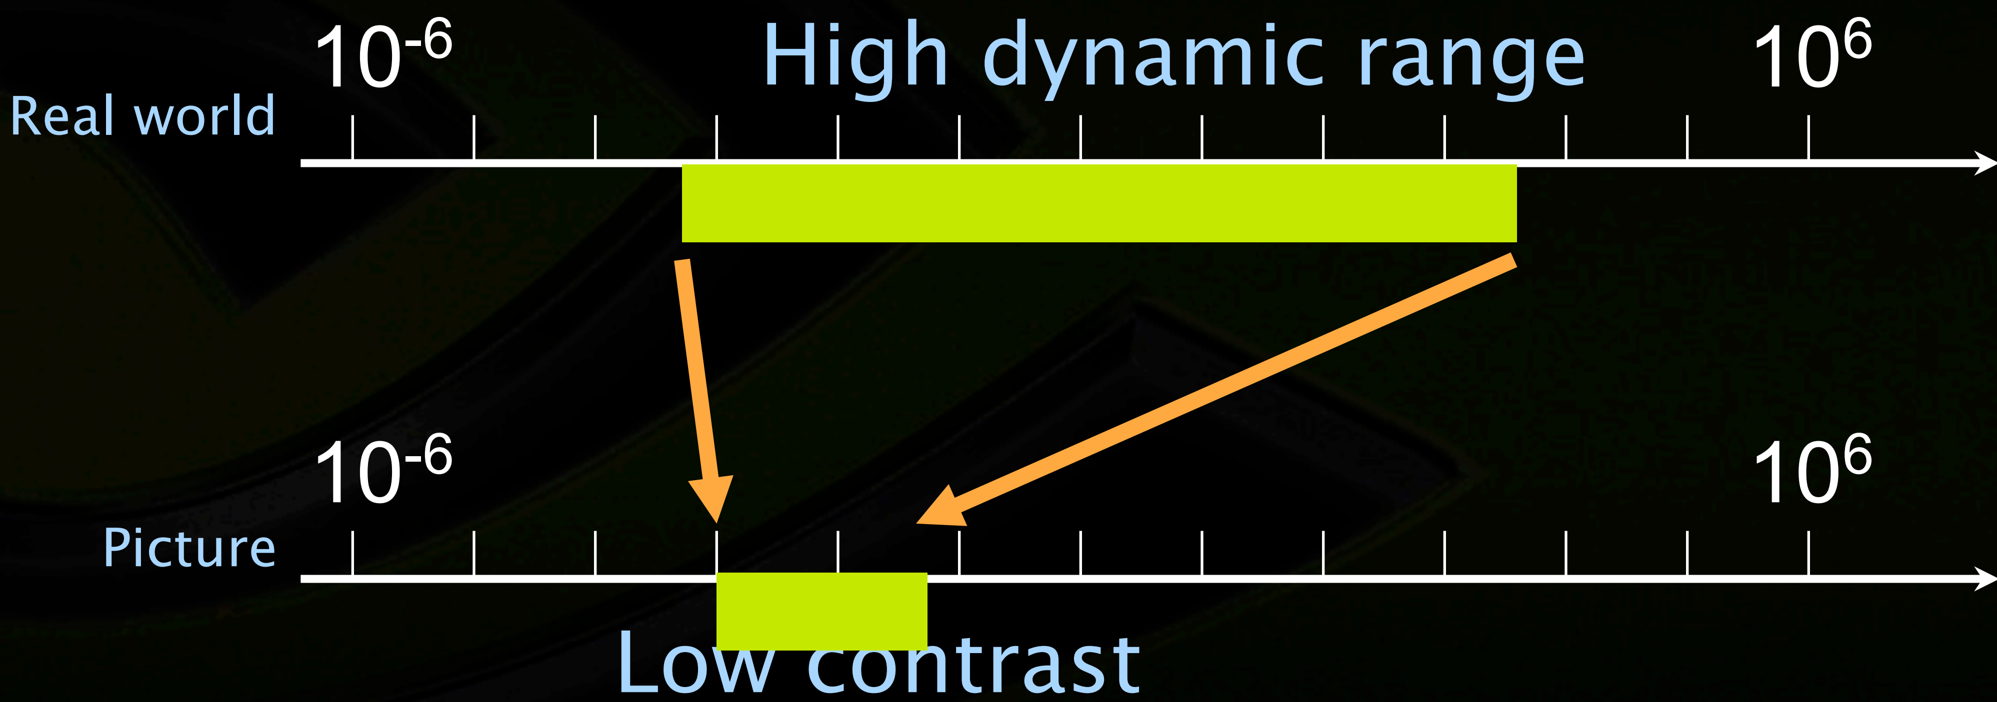

Why is it difficult ?

Why is it difficult ?

- Is it harder to obtain good blacks, or good whites?

Why is it difficult ?

- Is it harder to obtain good blacks, or good whites?

Why is it difficult ?

- Is it harder to obtain good blacks, or good whites?
- Black is harder. It’s hard to absorb all the light.
  - See the history of painting: good blacks appeared late

Why is it difficult ?

- Is it harder to obtain good blacks, or good whites?
- Black is harder. It’s hard to absorb all the light.
  - See the history of painting: good blacks appeared late

Why is it difficult ?

- Is it harder to obtain good blacks, or good whites?
- Black is harder. It’s hard to absorb all the light.
    - See the history of painting: good blacks appeared late
- We can achieve excellent white
    - Albedo >100%
    - How is this possible?
    - Use fluorescence
    - Most white materials (paper, paint, fabric) are fluorescent

Photo paper dynamic range

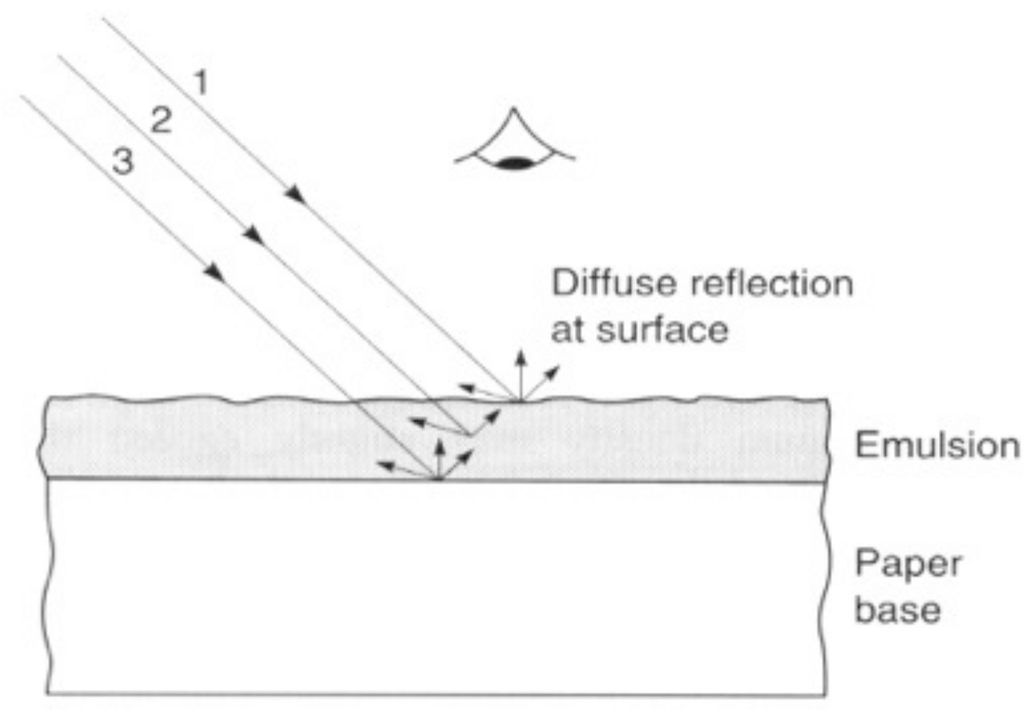

(a) Matt paper

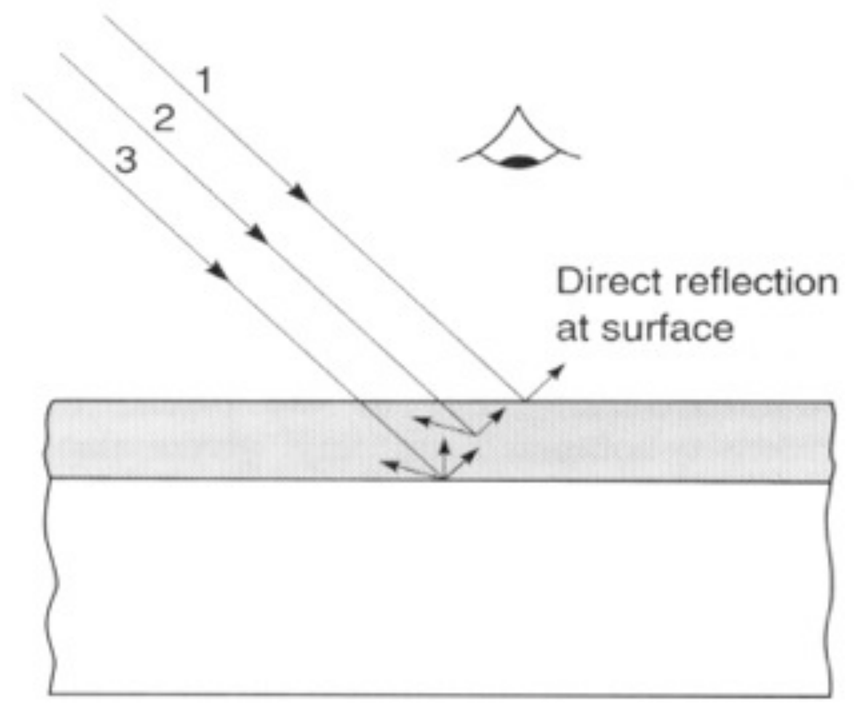

(b) Glossy paper

From The Manual of Photography, Jacobson et al.

Photo paper dynamic range

- Matte vs. glossy: which has the highest dynamic range?
  - Glossy because for some directions, it does not reflect light at all, while matte reflects equally in all directions

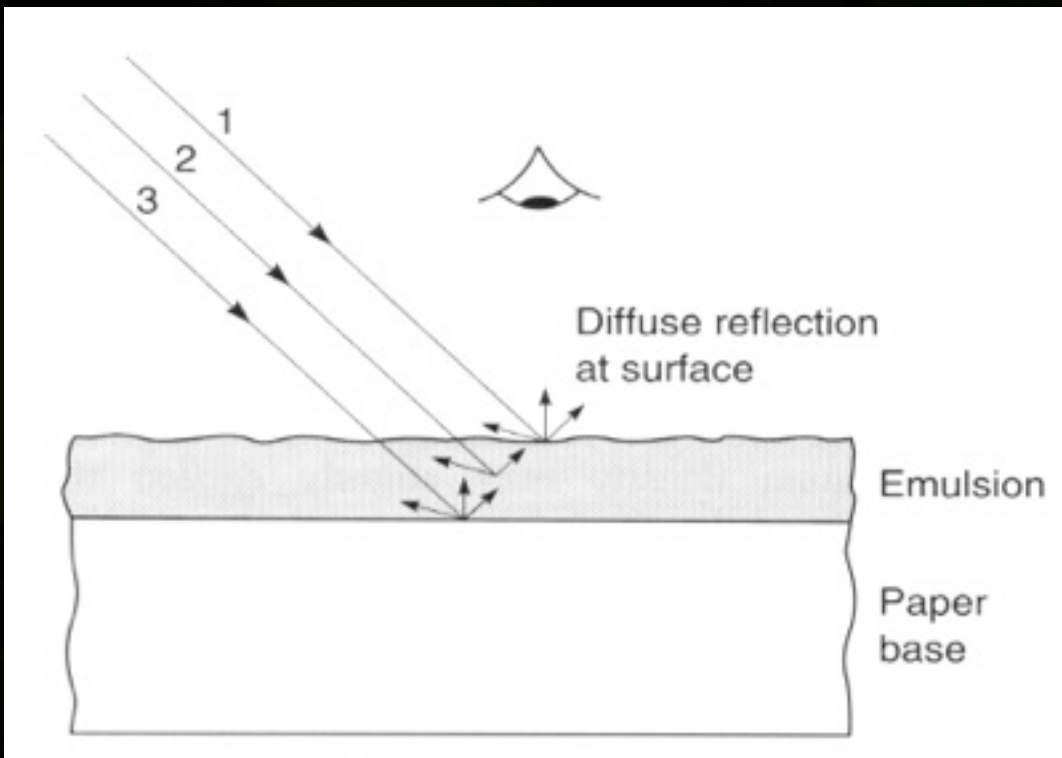

(a) Matt paper

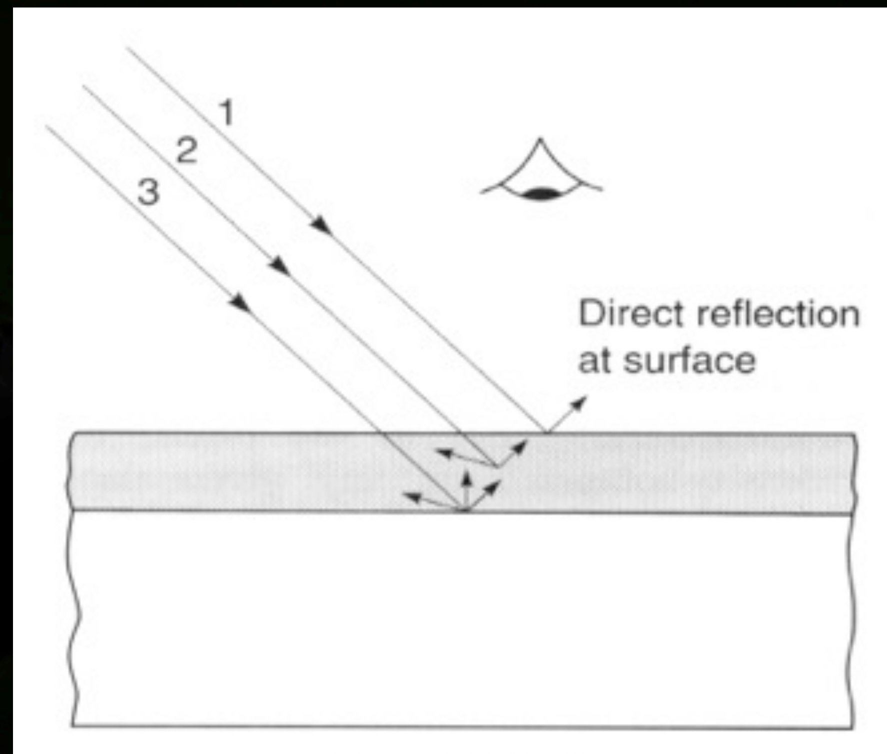

(b) Glossy paper

From The Manual of Photography, Jacobson et al.

Sunnybrook HDR display

- Bright back light at low resolution
  - modulate with the front panel
  - over 1 : 10,000 doable

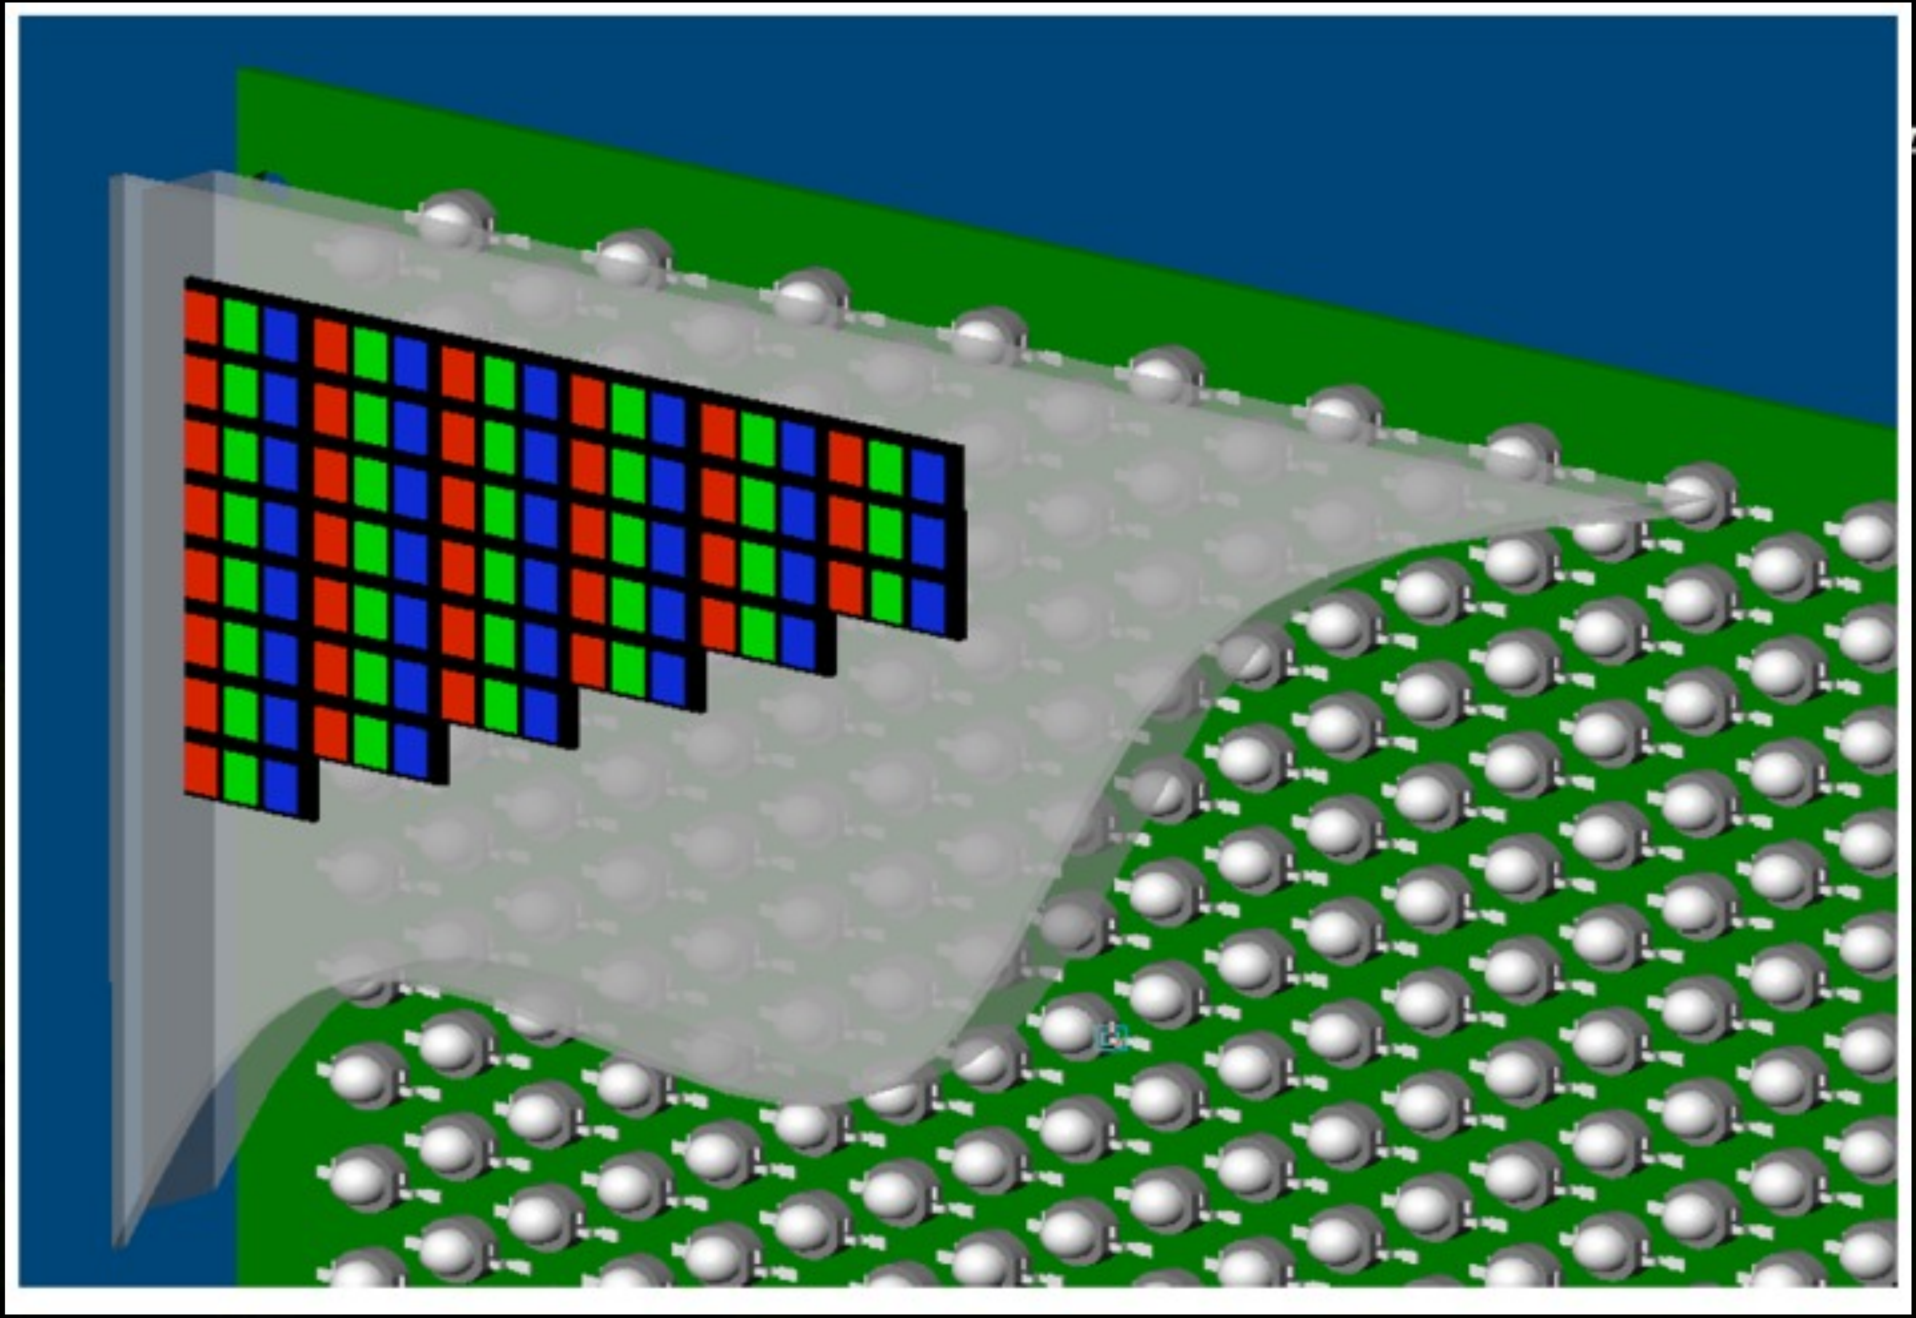

Monday, March 5, 12

How it works

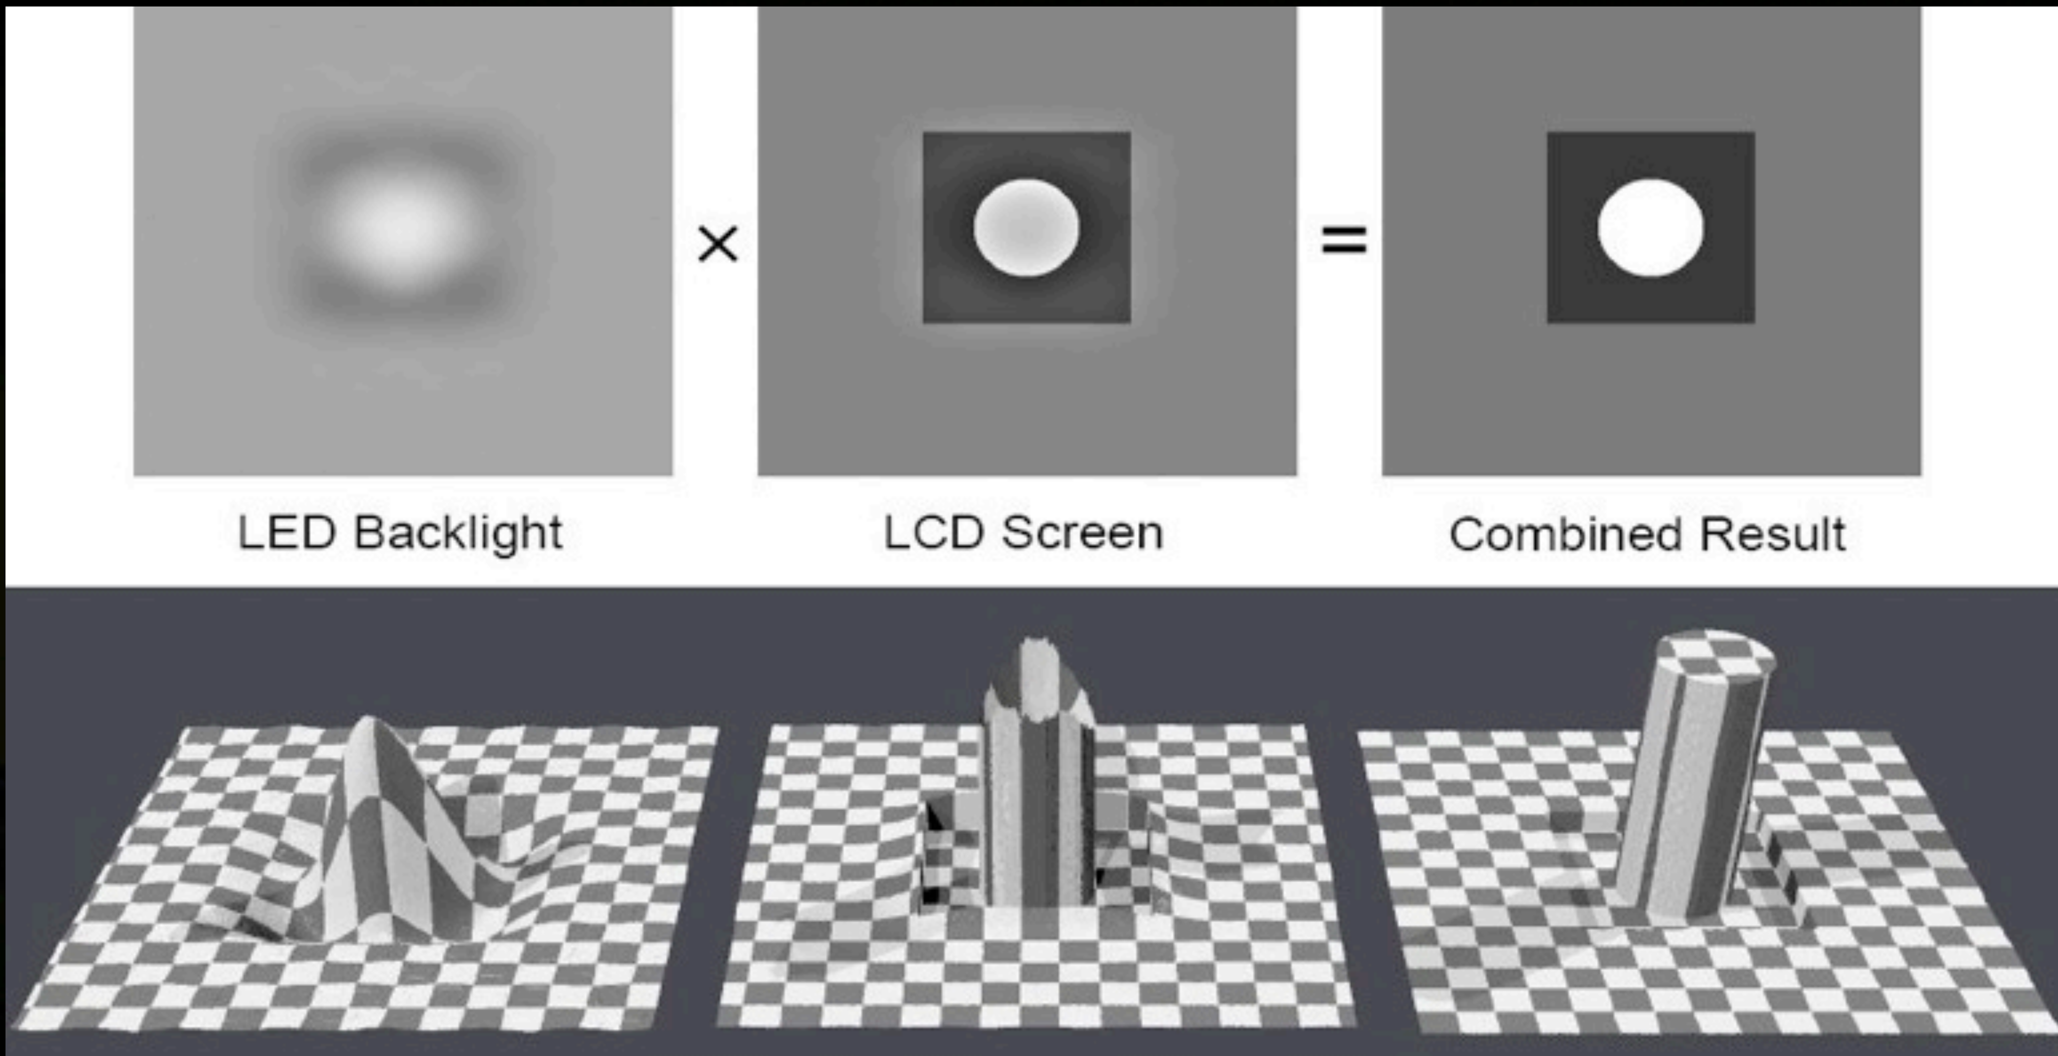

How humans deal with dynamic range

- We're sensitive to contrast (multiplicative)
  - A ratio of 1:2 is perceived as the same contrast as a ratio of 100 to 200
  - Use the log domain as much as possible
- Dynamic adaptation (very local in retina)
  - Pupil (not so important)
  - Neural & chemical
    - can adapt ~ 10^{10}
- Transmit the signal to brain
  - only 10^3 – 10^4
  - spatial contrast-based processing already in the eye

How humans deal with dynamic range

- We're sensitive to contrast (multiplicative)
  - A ratio of 1:2 is perceived as the same contrast as a ratio of 100 to 200
  - Use the log domain as much as possible
- Dynamic adaptation (very local in retina)
  - Pupil (not so important)
  - Neural & chemical
    - can adapt ~ 10^{10}
- Transmit the signal to brain
  - only 10^3 – 10^4
  - spatial contrast-based processing already in the eye

Dim Light

Pupil dilates
More light enters the eye

How humans deal with dynamic range

- We're sensitive to contrast (multiplicative)
  - A ratio of 1:2 is perceived as the same contrast as a ratio of 100 to 200
  - Use the log domain as much as possible
- Dynamic adaptation (very local in retina)
  - Pupil (not so important)
  - Neural & chemical
    - can adapt ~ 10^{10}
- Transmit the signal to brain
  - only 10^3 – 10^4
  - spatial contrast-based processing already in the eye

Dim Light

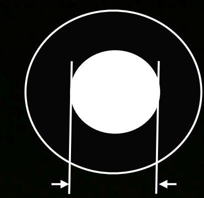

~6 mm
Pupil dilates
More light enters the eye

Bright Light

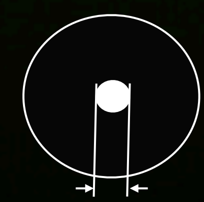

~1 mm
Pupil constricts
Less light enters the eye

How humans deal with dynamic range

- We're sensitive to contrast (multiplicative)
  - A ratio of 1:2 is perceived as the same contrast as a ratio of 100 to 200
  - Use the log domain as much as possible
- Dynamic adaptation (very local in retina)
  - Pupil (not so important)
  - Neural & chemical
    - can adapt ~ 10^{10}
- Transmit the signal to brain
  - only 10^3 – 10^4
  - spatial contrast-based processing already in the eye

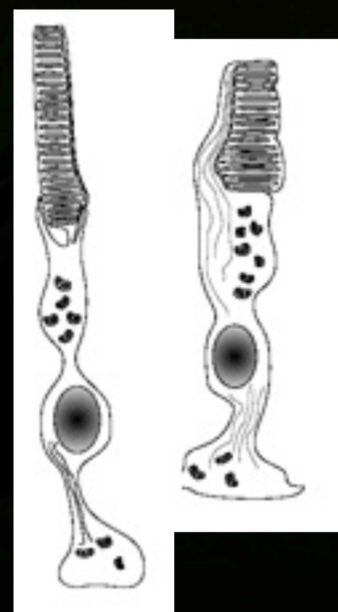

Dim Light

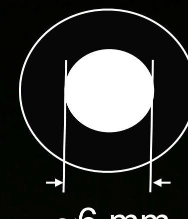

Pupil dilates
More light enters the eye

Area ratio
~16 : 1

Bright Light

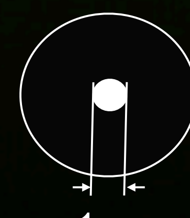

Pupil constricts
Less light enters the eye

How humans deal with dynamic range

- We're sensitive to contrast (multiplicative)
  - A ratio of 1:2 is perceived as the same contrast as a ratio of 100 to 200
  - Use the log domain as much as possible
- Dynamic adaptation (very local in retina)
  - Pupil (not so important)
  - Neural & chemical
    - can adapt ~ 10^{10}
- Transmit the signal to brain
  - only 10^3 – 10^4
  - spatial contrast-based processing already in the eye

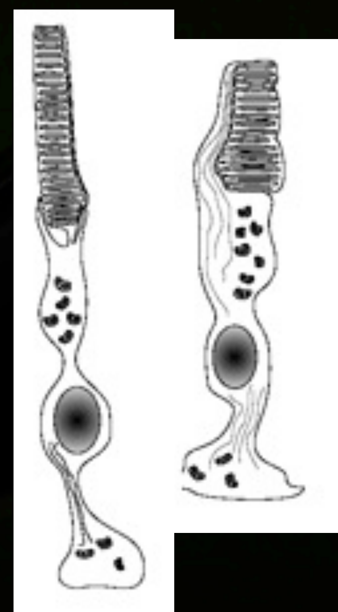

Dim Light

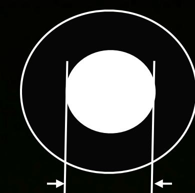

~6 mm
Pupil dilates
More light enters the eye

Area ratio
~16 : 1

Bright Light

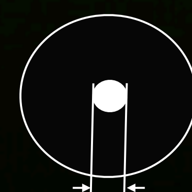

~1 mm
Pupil constricts
Less light enters the eye

Mesopic vision

Perceptually Based Tone Mapping for Low-Light Conditions

Adam G. Kirk   James F. O'Brien
University of California, Berkeley

- 3 cones + 1 rod map to 3 signals from eye to visual cortex

Mesopic vision

- 3 cones + 1 rod map to 3 signals from eye to visual cortex

Perceptually Based Tone Mapping for Low-Light Conditions

Adam G. Kirk James F. O'Brien
University of California, Berkeley

Figure 6: *Top: An HDR image of the Fremont Troll with no tone mapping. Bottom: The image has been tone mapped for low-light conditions.*

Mesopic vision

- 3 cones + 1 rod map to 3 signals from eye to visual cortex

Perceptually Based Tone Mapping for Low-Light Conditions

Adam G. Kirk James F. O'Brien
University of California, Berkeley

Figure 4: *Left: An HDR night scene with no tone mapping featuring UC Berkeley's Sather Tower. Right: The image has been tone mapped for low-light conditions. Artifacts on the clock face occur because the hands moved during image acquisition.*

...ure 6: *Top: An HDR image of the Fremont Troll with no tone mapping. Bottom: The image has been tone mapped for low-light conditions.*

Multiple exposure photography

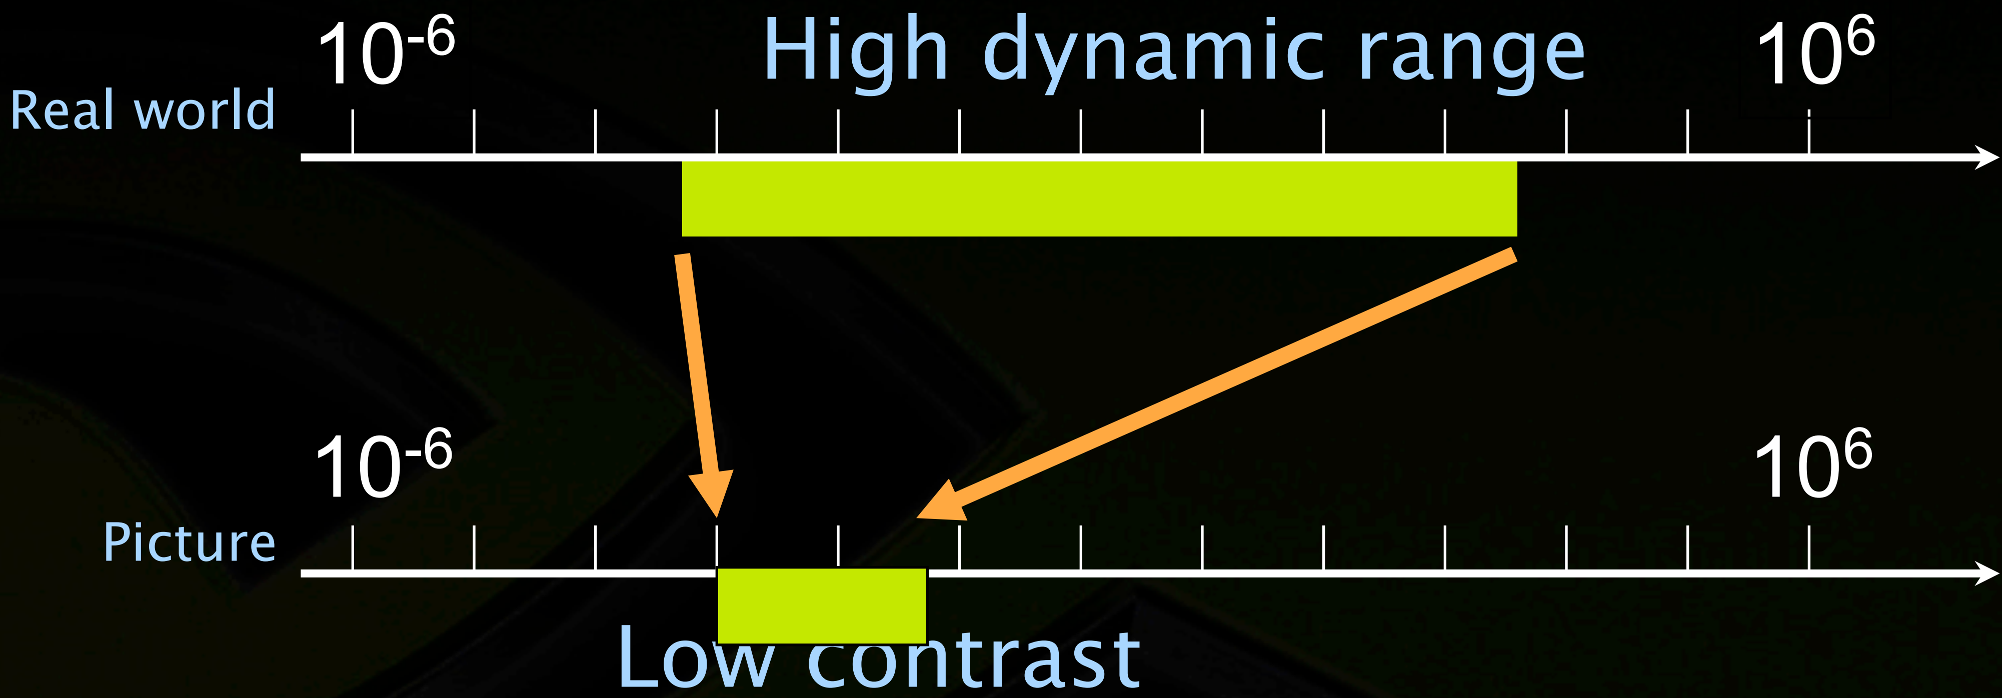

Multiple exposure photography

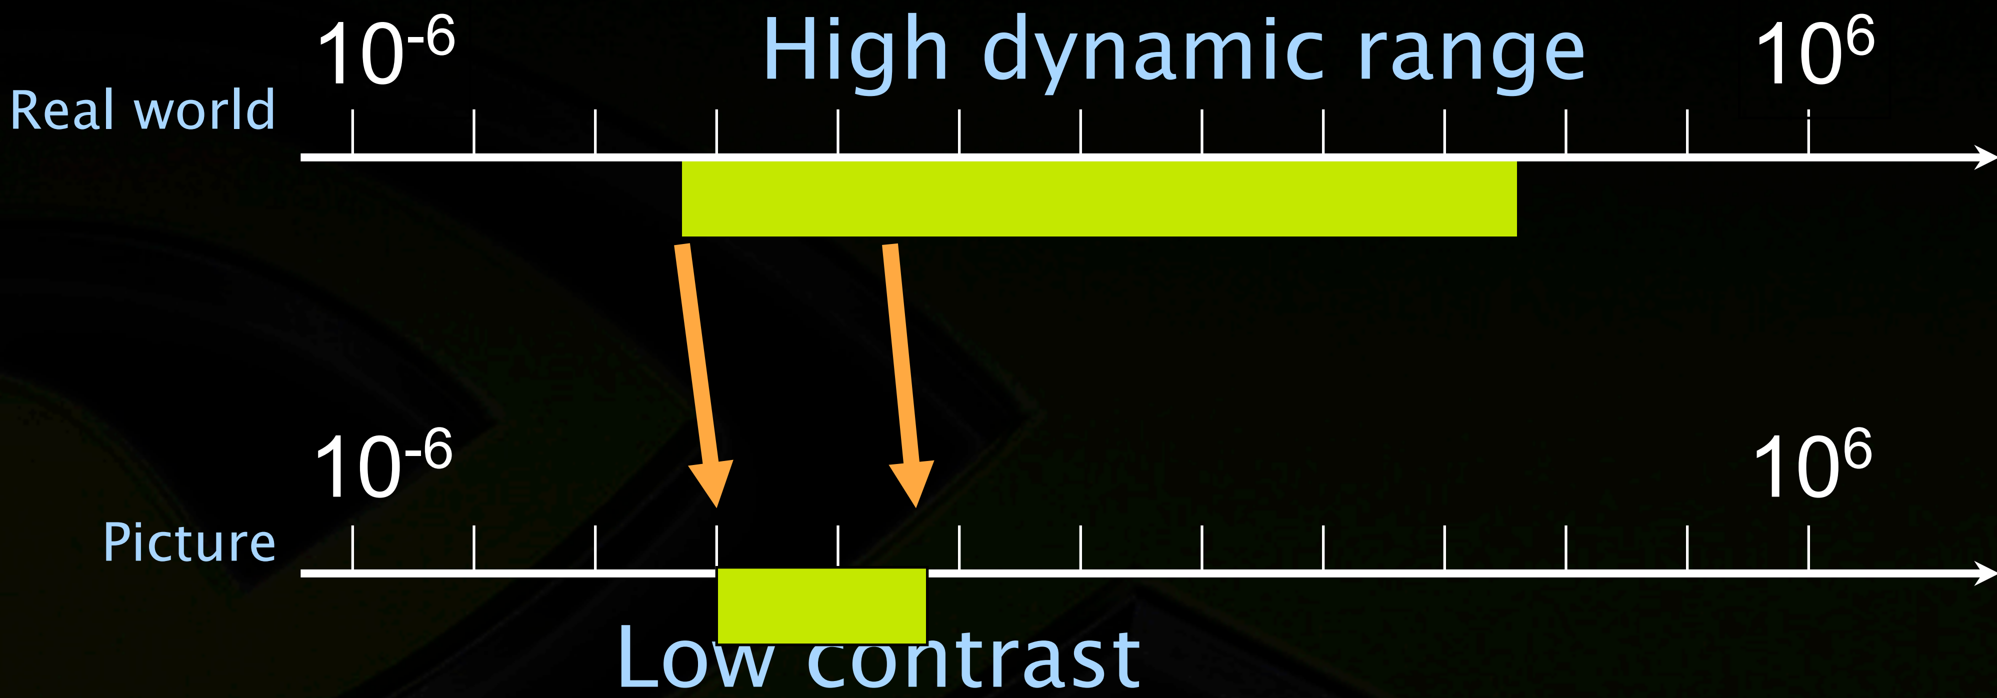

Multiple exposure photography

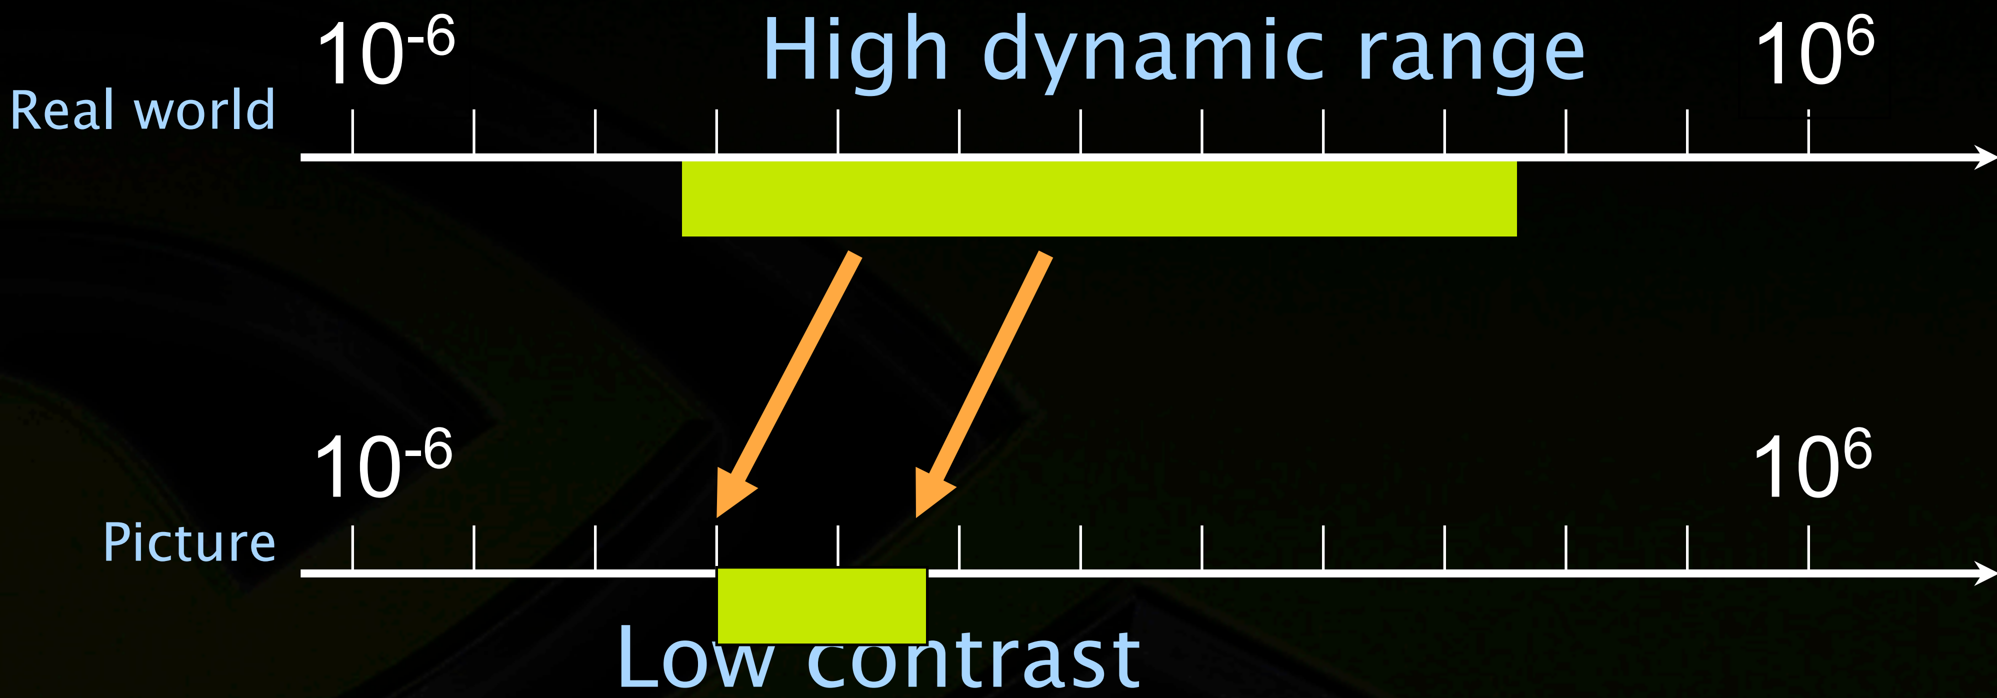

Multiple exposure photography

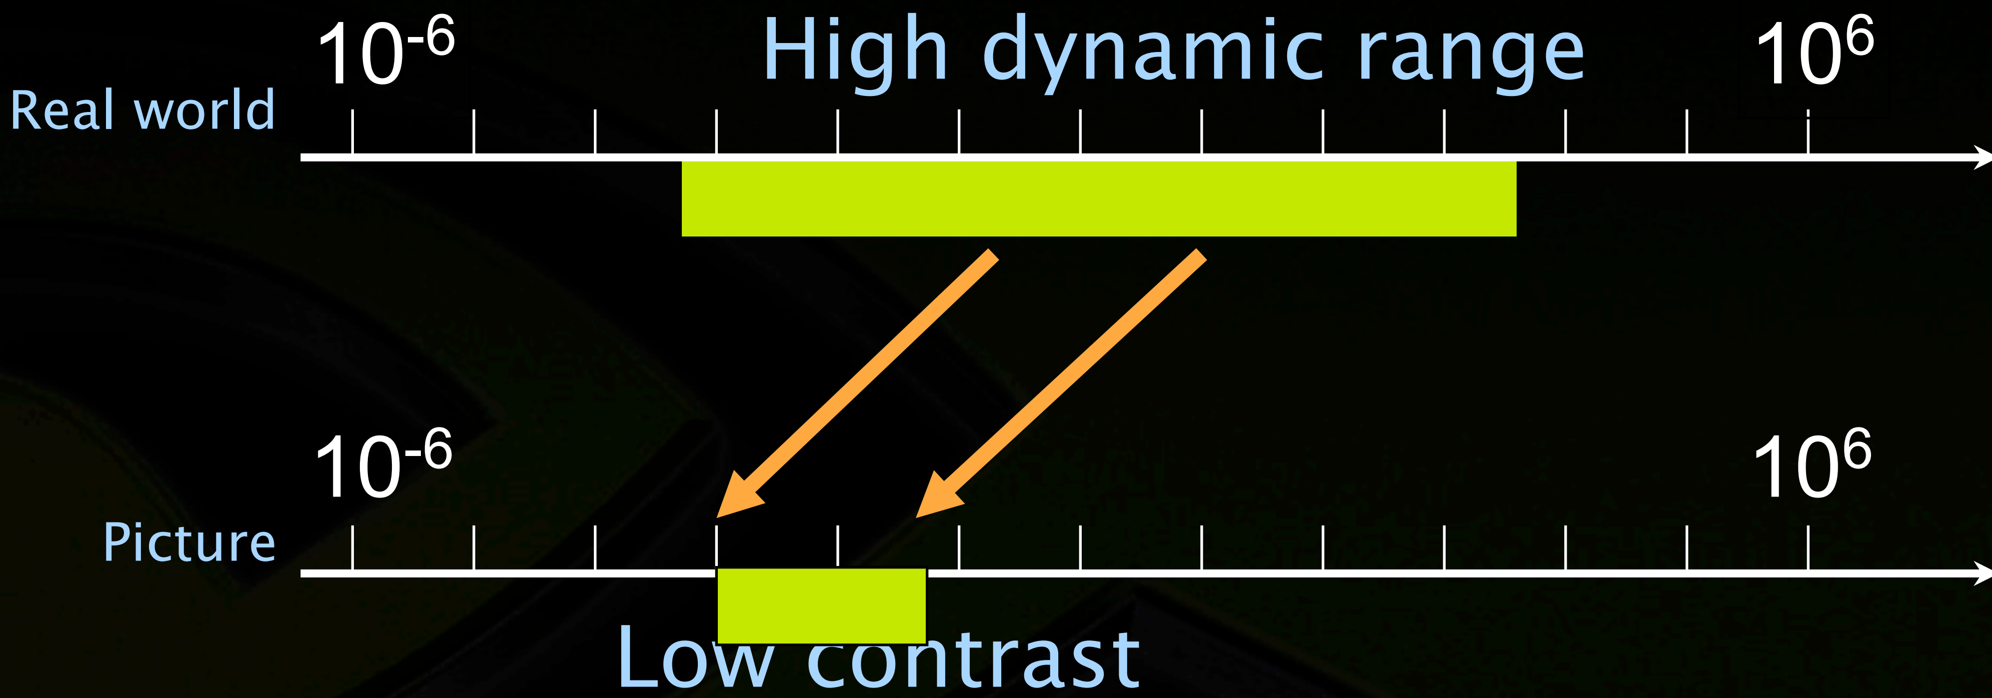

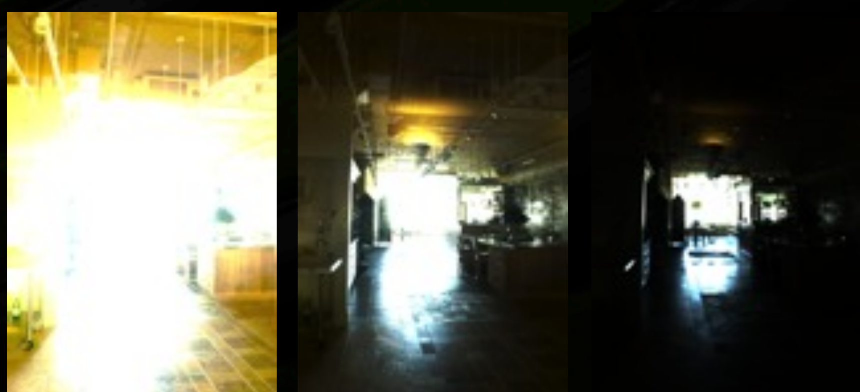

Multiple exposure photography

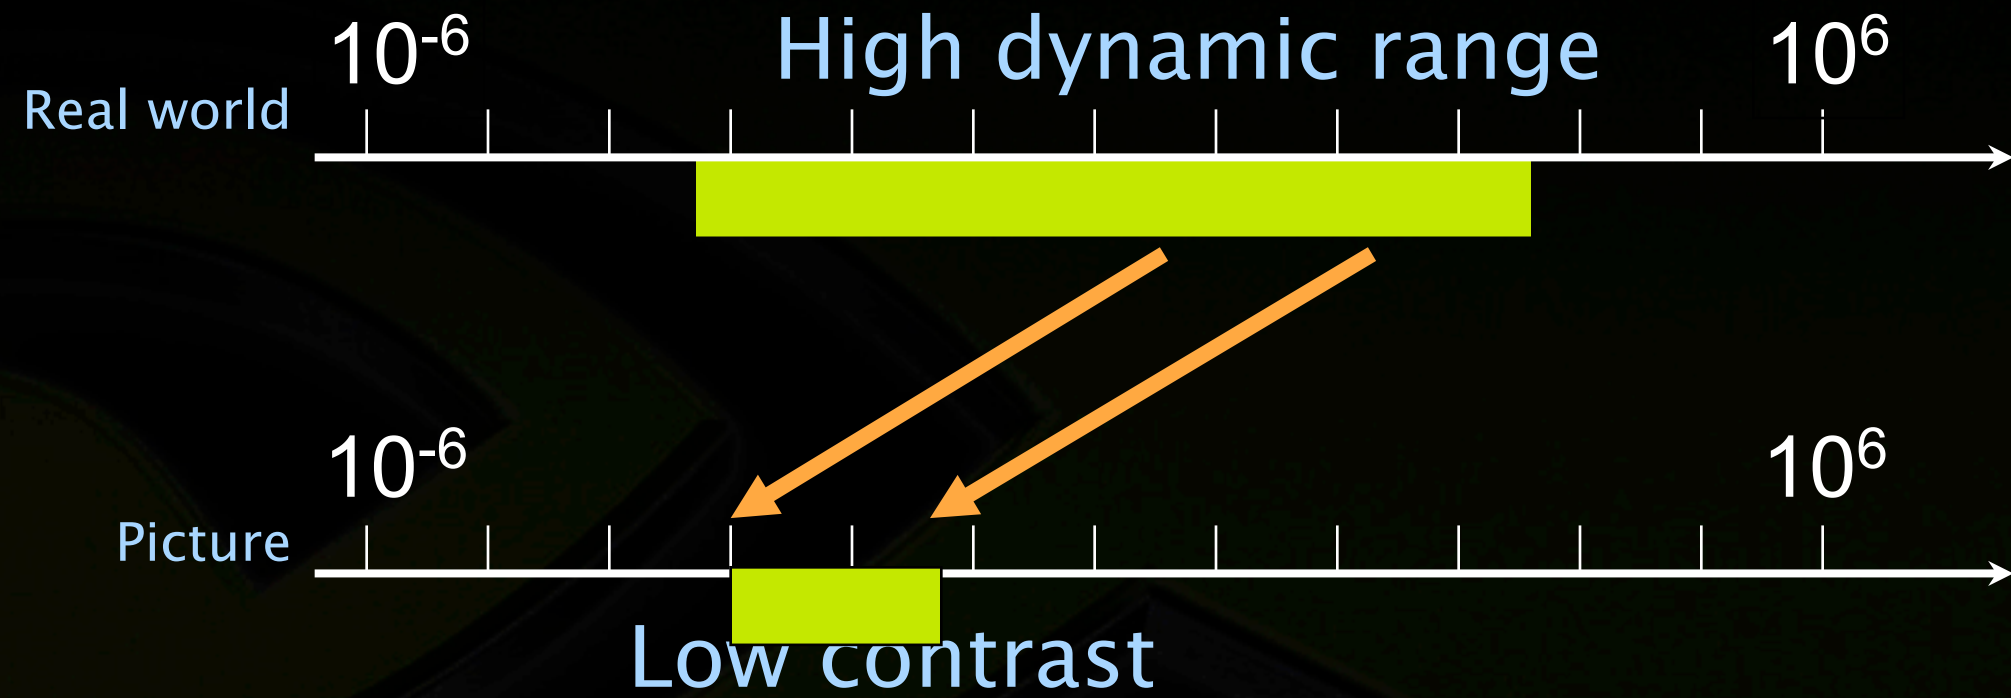

Multiple exposure photography

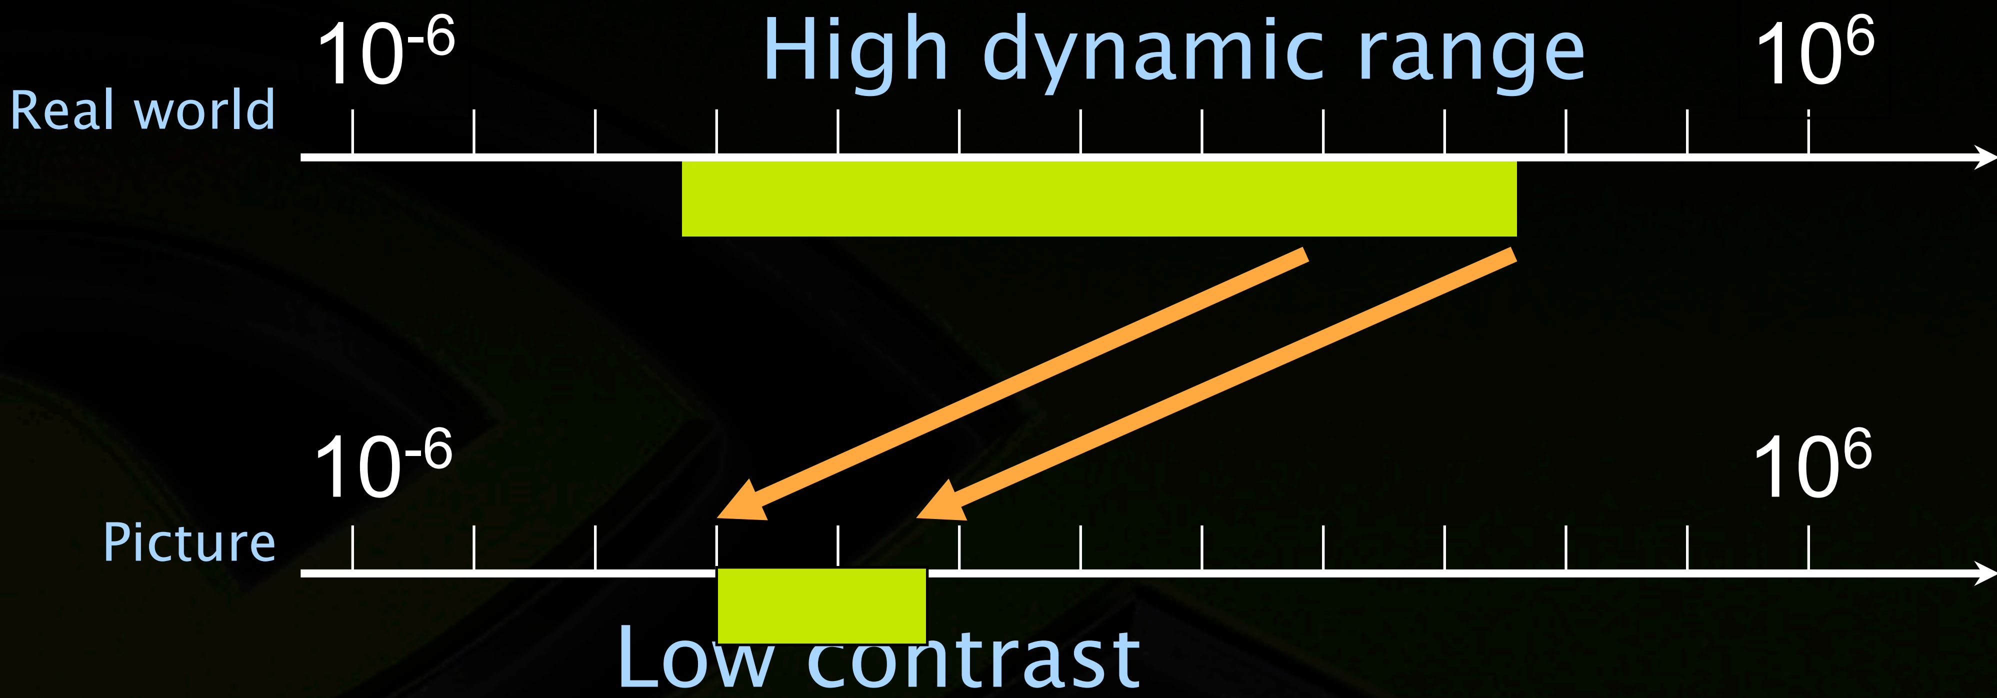

How can we vary exposure?

- Options:
  - Shutter speed
  - Aperture
  - ISO
  - Neutral density filter

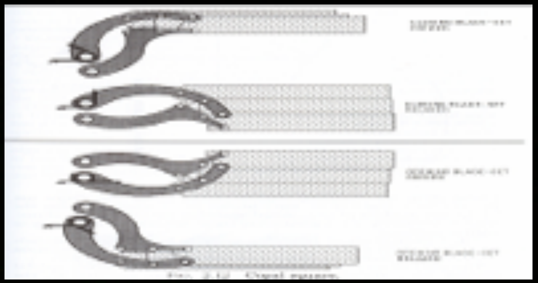

Slide inspired by Siggraph 2005 course on HDR

Tradeoffs

Tradeoffs

- Shutter speed
  - Range: ~30 sec to 1/4000sec (6 orders of magnitude)
  - Pros: reliable, linear
  - Cons: sometimes noise for long exposure

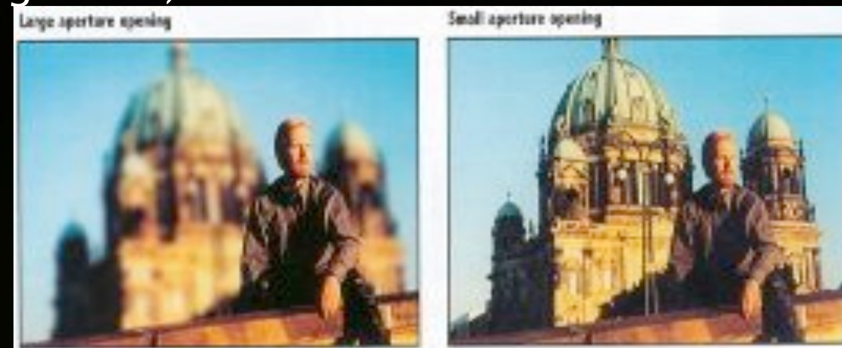

Tradeoffs

- Shutter speed
  - Range: ~30 sec to 1/4000sec (6 orders of magnitude)
  - Pros: reliable, linear
  - Cons: sometimes noise for long exposure
- Aperture
  - Range: ~f/1.4 to f/22 (2.5 orders of magnitud
  - Cons: changes depth of field
  - Useful when desperate

Tradeoffs

- Shutter speed
  - Range: ~30 sec to 1/4000sec (6 orders of magnitude)
  - Pros: reliable, linear
  - Cons: sometimes noise for long exposure
- Aperture
  - Range: ~f/1.4 to f/22 (2.5 orders of magnitud
  - Cons: changes depth of field
  - Useful when desperate
- ISO
  - Range: ~100 to 1600 (1.5 orders of magnitude)
  - Cons: noise
  - Useful when desperate

Tradeoffs

- Shutter speed
  - Range: ~30 sec to 1/4000sec (6 orders of magnitude)
  - Pros: reliable, linear
  - Cons: sometimes noise for long exposure
- Aperture
  - Range: ~f/1.4 to f/22 (2.5 orders of magnitude)
  - Cons: changes depth of field
  - Useful when desperate
- ISO
  - Range: ~100 to 1600 (1.5 orders of magnitude)
  - Cons: noise
  - Useful when desperate
- Neutral density filter
  - Range: up to 4 densities (4 orders of magnitude) & can be stacked
  - Cons: not perfectly neutral (color shift), not very precise, need to touch camera (shake)
  - Pros: works with strobe/flash, good complement when desperate

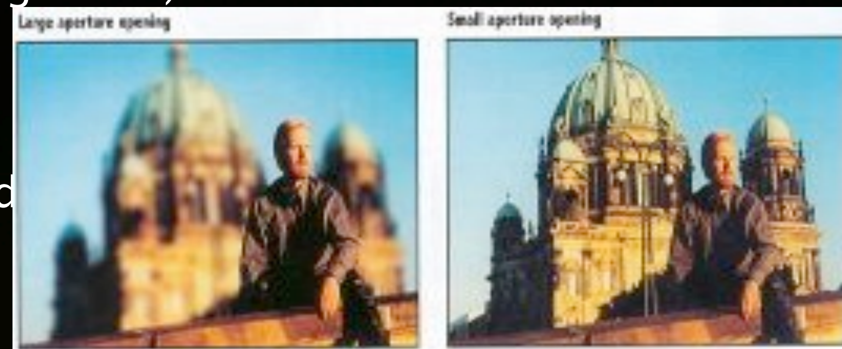

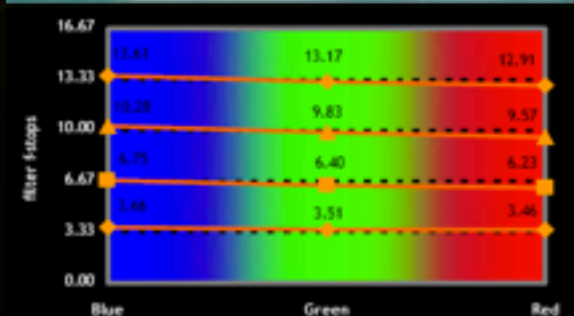

Early HDR photos: Gustave Le Gray (~1850)

- Take two shots
  - one for the sky – direct light
  - one for the rest – reflected light
  - cut and paste the negatives, and develop

Early HDR photos: Gustave Le Gray (~1850)

- Take two shots
  - one for the sky – direct light
  - one for the rest – reflected light
  - cut and paste the negatives, and develop

HP Robinson (1858) Fading Away, 5 negatives

Robust image registration of different exposures

- Exposure stack capture
  - camera may move
  - details look different
- Median-Threshold Bitmap (MTB)
  - use a black and white version of the image thresholded at the median
- Find the translation that minimizes difference
- Accelerate using pyramid

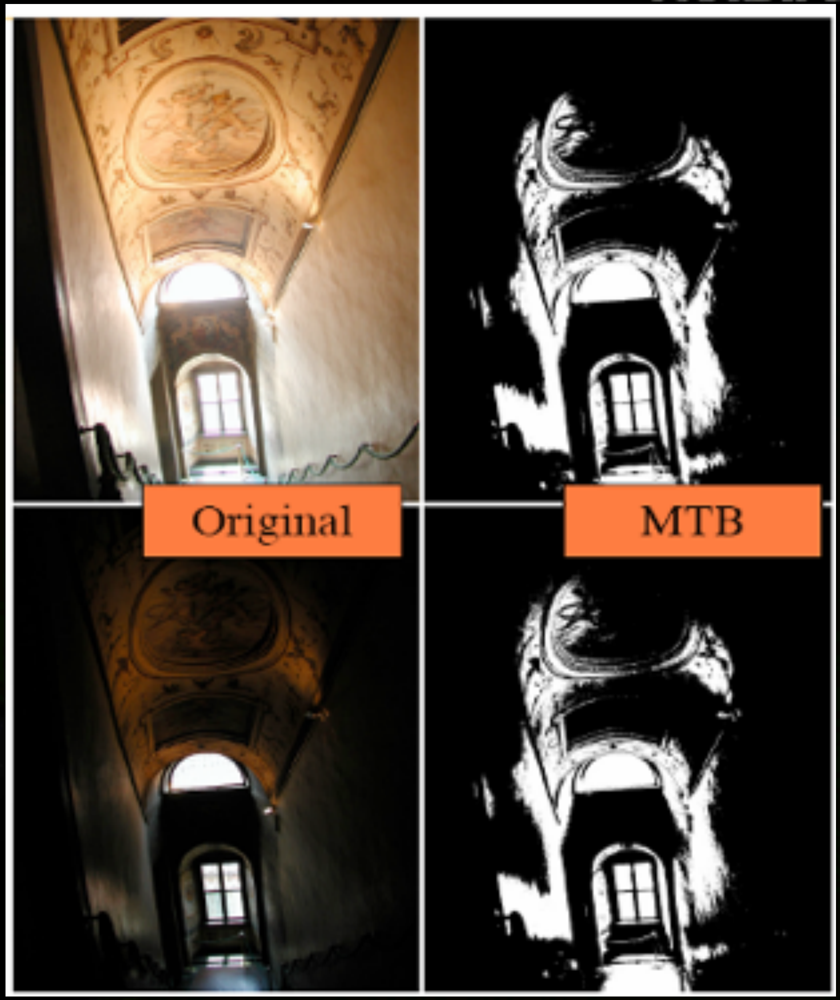

Alignment Results

SIGGRAPH2005

5 unaligned exposures | Close-up detail | MTB alignment

Time: About .2 second/exposure for 3 MPixel image

Feature-based registration

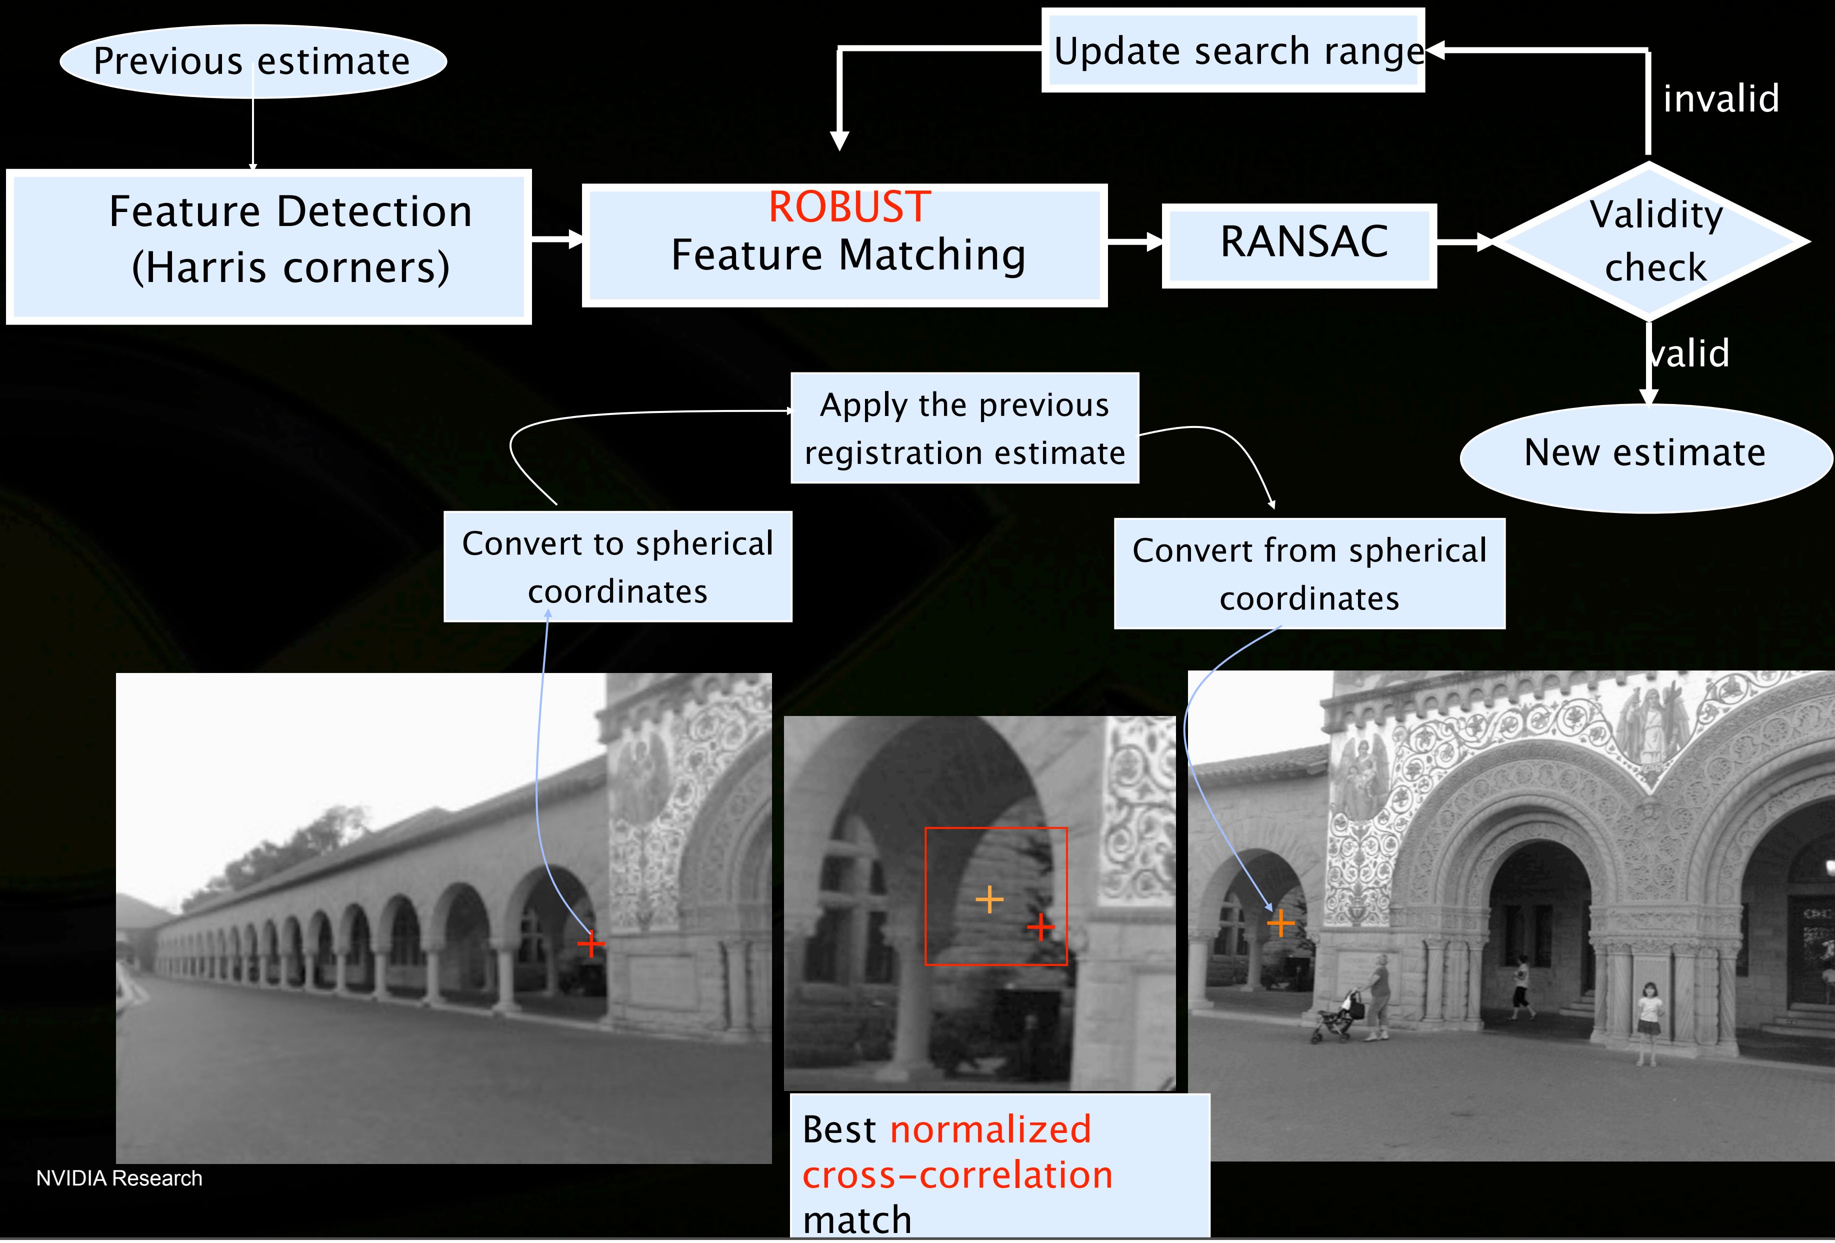

Don't forget lens distortion!

- Modern lenses are pretty good
  - often the distortion is quite small
- Homography only works for planar scenes
  - with rotations around entrance pupil homography is OK
- But if you want to align and accumulate images
  - you should account for all sources of warpage

Response curve calibration

- Debevec and Malik (1997)
  - select a small number of pixels from the images
    - using all pixels would make the matrix too large
  - optimize the response curve with a smoothness constraint
- Robertson et al. (1999)
  - optimize over all pixels in all images
    - no need to solve a linear equation with a large matrix
  - iterate
    - calculate HDR image using the response curve
    - find a better response curve using HDR image
  - refinement (2003): better weights

Calibration [Debevec & Malik 97]

$\Delta t =$ 1/64 sec

$\Delta t =$ 1/16 sec

$\Delta t =$ 1/4 sec

$\Delta t =$ 1 sec

$\Delta t =$ 4 sec

Calibration [Debevec & Malik 97]

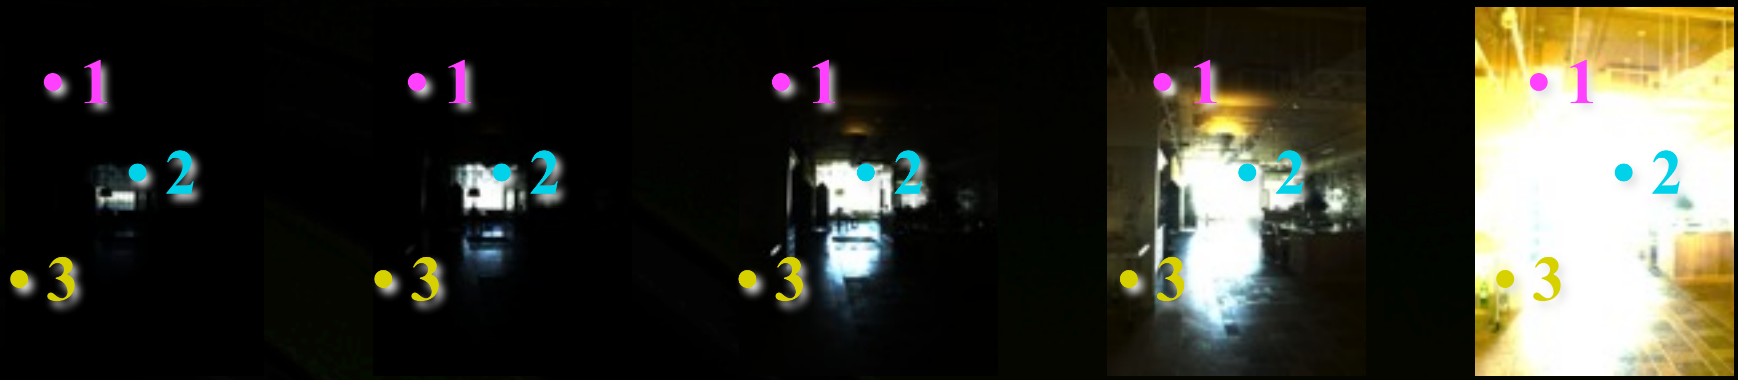

Δt = 1/64 sec | Δt = 1/16 sec | Δt = 1/4 sec | Δt = 1 sec | Δt = 4 sec

Pixel Value Z = f (Exposure)

Calibration [Debevec & Malik 97]

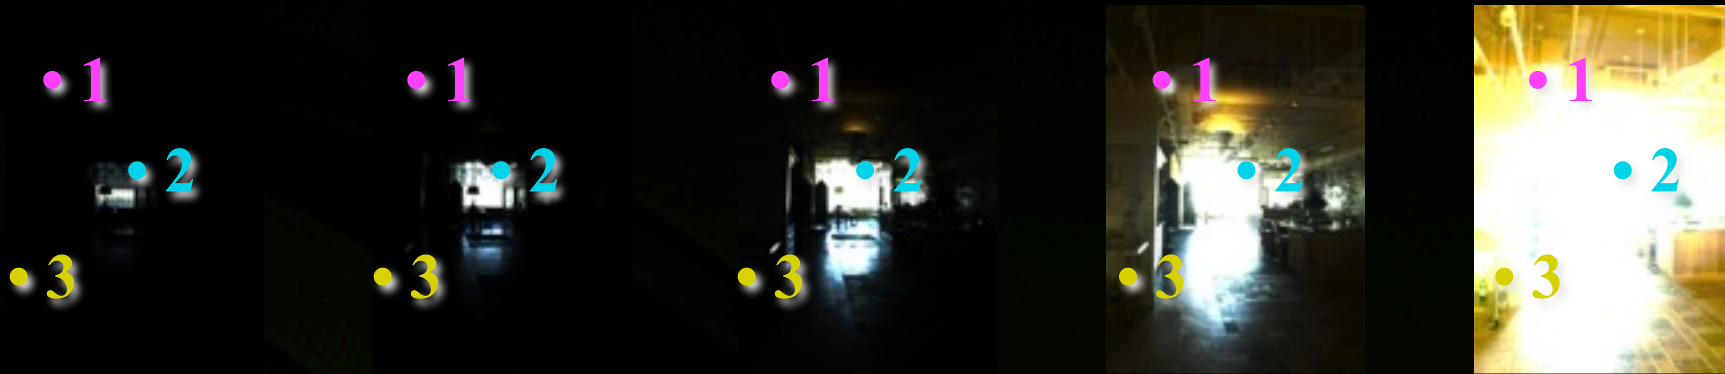

$\Delta t =$ 1/64 sec | $\Delta t =$ 1/16 sec | $\Delta t =$ 1/4 sec | $\Delta t =$ 1 sec | $\Delta t =$ 4 sec

Pixel Value $Z = f($ Exposure $)$

Exposure = Radiance $* \Delta t$

Calibration [Debevec & Malik 97]

$\Delta t =$ 1/64 sec	$\Delta t =$ 1/16 sec	$\Delta t =$ 1/4 sec	$\Delta t =$ 1 sec	$\Delta t =$ 4 sec

Pixel Value $Z = f($ Exposure $)$

Exposure = Radiance $* \Delta t$

$\ln$ Exposure $= \ln$ Radiance $+ \ln \Delta t$

Adjust exposure to find a smooth response curve

Assuming the same exposure for each pixel

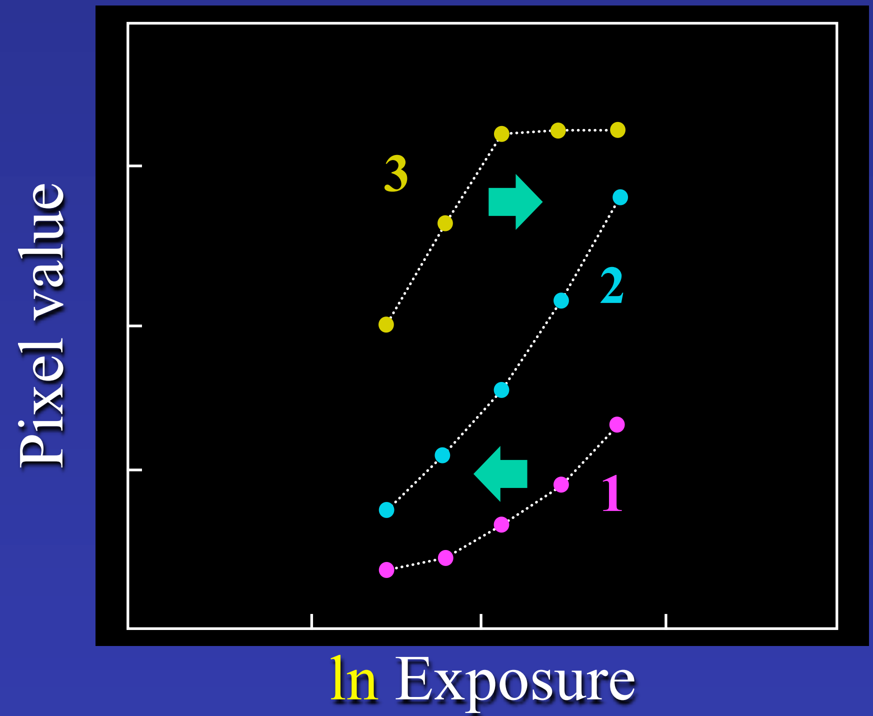

After adjusting radiances to obtain a smooth response curve

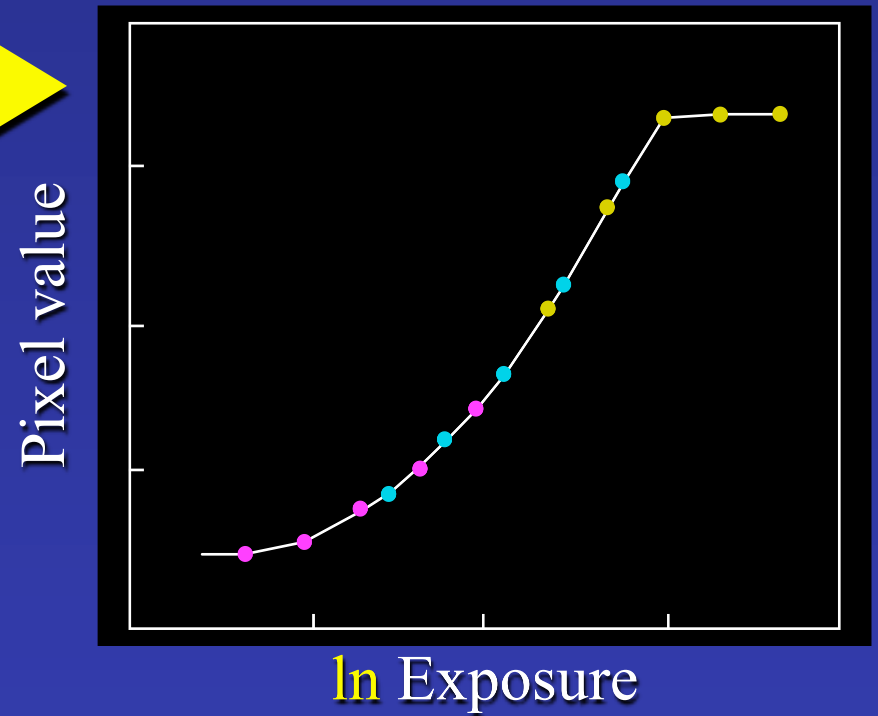

The Math

- Let f be the response function: $Z_{ij} = f(E_i \Delta t_j)$
- Let g be the logarithm of the inverse response function: $g(Z_{ij}) = \ln f^{-1}(Z_{ij}) = \ln E_i + \ln \Delta t_j$
- Solve the overdetermined linear system:
  - unknown E_i, $g(\,)$

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

Matlab code

```
%
% gsolve.m - Solve for imaging system response function
%
% Given a set of pixel values observed for several pixels
% in several images with different exposure times, this
% function returns the imaging system's response function
% g as well as the log film irradiance values for the
% observed pixels.
%
% Assumes:
%
% Zmin = 0
% Zmax = 255
%
% Arguments:
%
% Z(i,j) is the pixel values of pixel location number I
%        in image j
% B(j)   is the log delta t, or log shutter speed, for
%        image j l is lamdba, the constant that
%        determines the amount of smoothness
% w(z) is the weighting function value for pixel value z
%
% Returns:
%
% g(z) is the log exposure corresponding to pixel value z
```

```
function [g,lE] = gsolve(Z,B,l,w)
n = 256;
A = zeros(size(Z,1)*size(Z,2)+n+1,n+size(Z,1));
b = zeros(size(A,1),1);

%% Include the data-fitting equations
k = 1;
for i=1:size(Z,1)
  for j=1:size(Z,2)
    wij              = w(Z(i,j)+1);
    A(k,Z(i,j)+1)    = wij;
    A(k,n+i)         = -wij;
    b(k,1)           = wij * B(i,j);
    k = k+1;
  end
end

%% Fix the curve by setting its middle value to 0
A(k,129) = 1;
k = k+1;

%% Include the smoothness equations
for i=1:n-2
  A(k,i)   = l*w(i+1);
  A(k,i+1) = -2*l*w(i+1);
  A(k,i+2) = l*w(i+1);
  k=k+1;
end

%% Solve the system using SVD
x  = A\b;
g  = x(1:n);
lE = x(n+1:size(x,1));
```

Results: Digital Camera

Kodak DCS460
1/30 to 30 sec

Recovered response curve

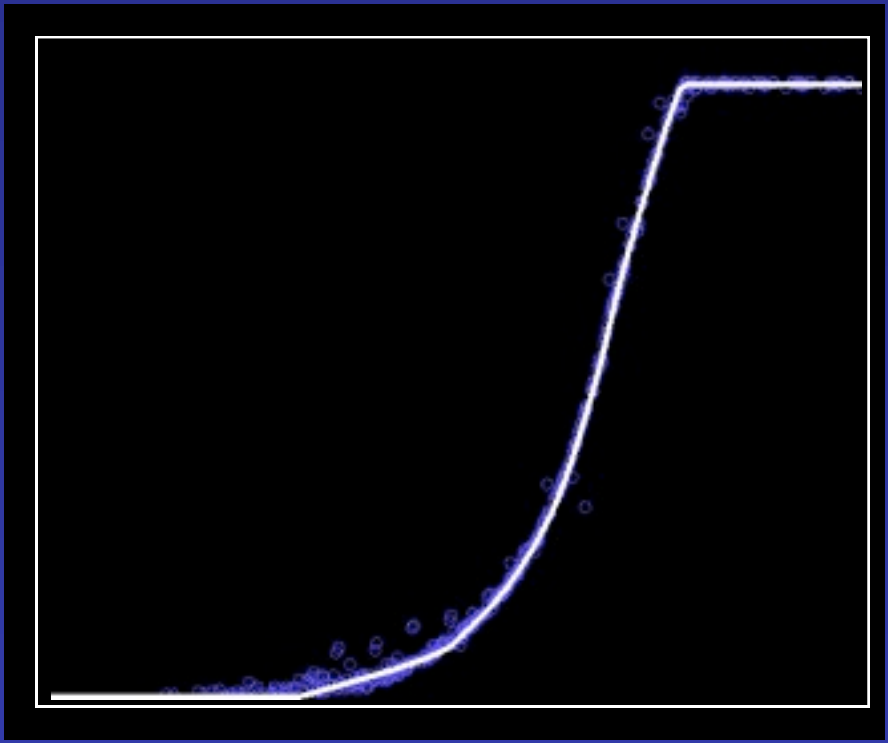

Pixel value

log Exposure

Reconstructed radiance map

Monday, March 5, 12

HDR representations

HDR representations

- Portable Float Map, Floating point TIFF
  - RGB as three floats
  - 3 * 32b = 96b

HDR representations

- Portable Float Map, Floating point TIFF
  - RGB as three floats
  - 3 * 32b = 96b
- Radiance
  - 8bit R, G, and B; 8bit shared exponent
    - $(R,G,B) / 255 * 2^{(E - 128)}$
  - 4 * 8b = 32b

HDR representations

- Portable Float Map, Floating point TIFF
  - RGB as three floats
  - 3 * 32b = 96b
- Radiance
  - 8bit R, G, and B; 8bit shared exponent
    - $(R,G,B) / 255 * 2^{(E - 128)}$
  - 4 * 8b = 32b
- OpenEXR
  - RGB each a half float, matches gfx HW
  - 3 * 16b = 48b

HDR representations

- Portable Float Map, Floating point TIFF
  - RGB as three floats
  - 3 * 32b = 96b
- Radiance
  - 8bit R, G, and B; 8bit shared exponent
    - $(R,G,B) / 255 * 2^{(E - 128)}$
  - 4 * 8b = 32b
- OpenEXR
  - RGB each a half float, matches gfx HW
  - 3 * 16b = 48b
- Many others that are better at compressing
  - but require conversions / LUT to operate on

Calibration [Robertson et al. 99]

- Input
  - series of **i** images with pixel values m = $\mathbf{y_{ij}}$ due to exposure $\mathbf{I_{ij}}$ = irradiance $\mathbf{x_j}$ * time $\mathbf{t_i}$

$$y_{ij} = f(I_{ij}) = f(x_j t_i)$$

  - weighting function $\mathbf{w_{ij} = w_{ij}(y_{ij})}$
    - assume central values most sensitive
- Task:
  - find pixel irradiance (luminance)
  - recover response curve
  - use EM (expectation maximization)

Calibration [Robertson et al. 99]

- Input
  - series of **i** images with pixel values m = $\mathbf{y_{ij}}$ due to exposure $\mathbf{I_{ij}}$ = irradiance $\mathbf{x_j}$ * time $\mathbf{t_i}$
  - weighting function $\mathbf{w_{ij} = w_{ij}(y_{ij})}$
    - assume central values most sensitive

$$y_{ij} = f(I_{ij}) = f(x_j t_i)$$

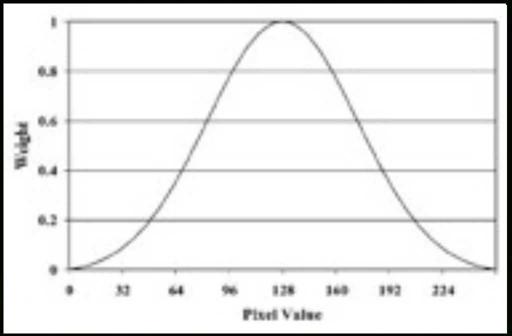

- Task:
  - find pixel irradiance (luminance)
  - recover response curve
  - use EM (expectation maximization)

$$x_j = f^{-1}(y_{ij}) / t_i$$

Calibration [Robertson et al. 99]

- Input
  - series of **i** images with pixel values m = y_{ij} due to exposure I_{ij} = irradiance x_j * time t_i
  - weighting function $w_{ij} = w_{ij}(y_{ij})$
    - assume central values most sensitive

$$y_{ij} = f(I_{ij}) = f(x_j t_i)$$

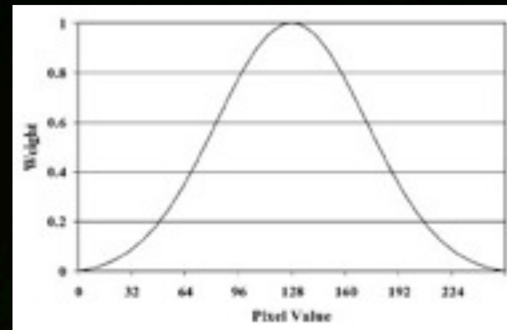

- Task:
  - find pixel irradiance (luminance)
  - recover response curve
  - use EM (expectation maximization)

$$x_j = f^{-1}(y_{ij}) / t_i$$

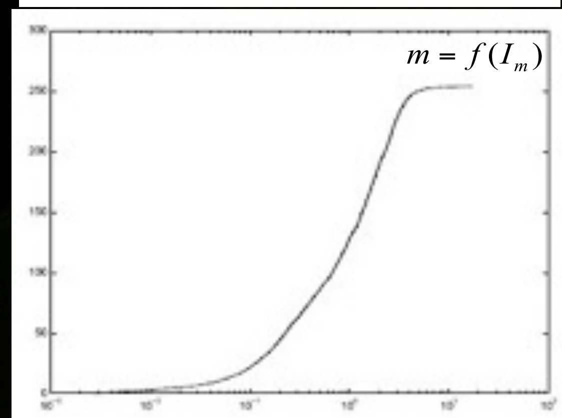

Calculate irradiances x_j given the images

- First iteration: unknown response curve
  - assume linear response with: $I_{128} = 1.0$
- Solve for irradiance at pixel j given
  - exposure times t_i
  - pixel values $m = I_{y_{ij}}$
  - weights w_m

$$\min_{x_j} \sum_{i,j} w_{ij} \left(I_{y_{ij}} - t_i x_j \right)^2 \quad \Rightarrow \quad x_j = \frac{\sum_i w_{ij} t_i I_{y_{ij}}}{\sum_i w_{ij} t_i^2}$$

Estimate response function for each m

$$I_m = \frac{1}{\mathrm{Card}(y_{ij} = m)}$$

Estimate response function for each m

- For all pixel values m = 0, ..., 254
  - ignore 255, that's saturated
  - collect every pixel with value m
  - calculate average I_m
    - note: weight is constant for a given m

$$I_m = \frac{1}{\mathrm{Card}(y_{ij} = m)}$$

Estimate response function for each m

- For all pixel values m = 0, ..., 254
  - ignore 255, that's saturated
  - collect every pixel with value m
  - calculate average I_m
    - note: weight is constant for a given m

$$I_m = \frac{1}{\mathrm{Card}(y_{ij} = m)}$$

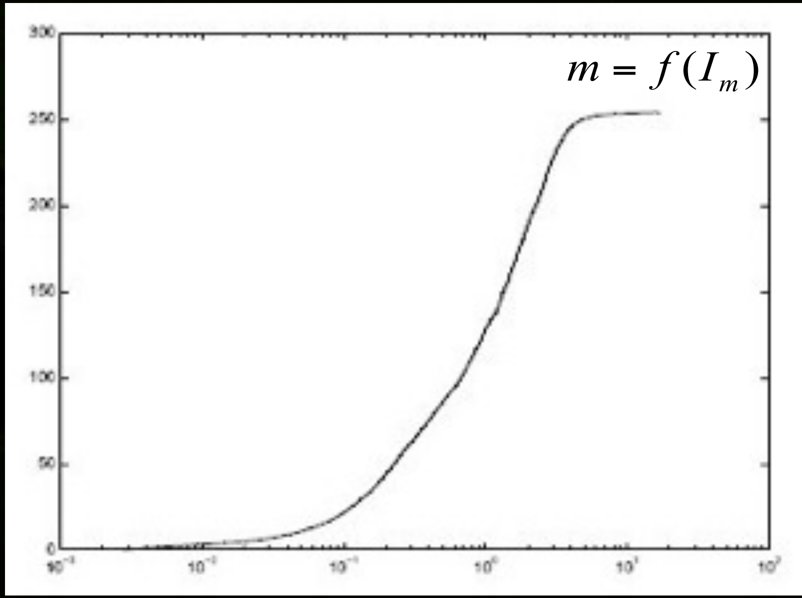

Estimate response function for each m

- For all pixel values m = 0, ..., 254
  - ignore 255, that's saturated
  - collect every pixel with value m
  - calculate average I_m
    - note: weight is constant for a given m
- Normalize
  - set $I_{128} = 1.0$

$$I_m = \frac{1}{\mathrm{Card}(y_{ij} = m)}$$

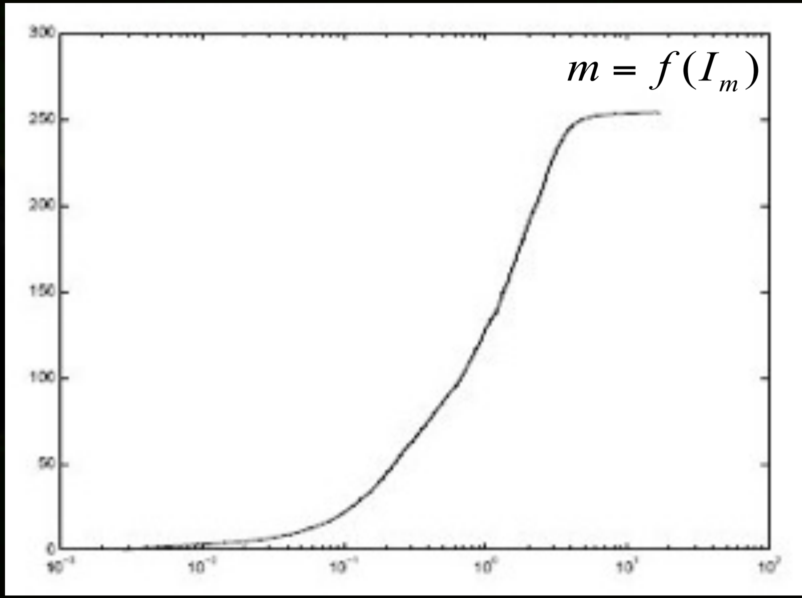

Estimate response function for each m

- For all pixel values m = 0, ..., 254
  - ignore 255, that's saturated
  - collect every pixel with value m
  - calculate average I_m
    - note: weight is constant for a given m

$$I_m = \frac{1}{\mathrm{Card}(y_{ij} = m)}$$

- Normalize
  - set $I_{128} = 1.0$

- Repeat (Expectation Maximization)
  - estimate irradiances
  - estimate response function

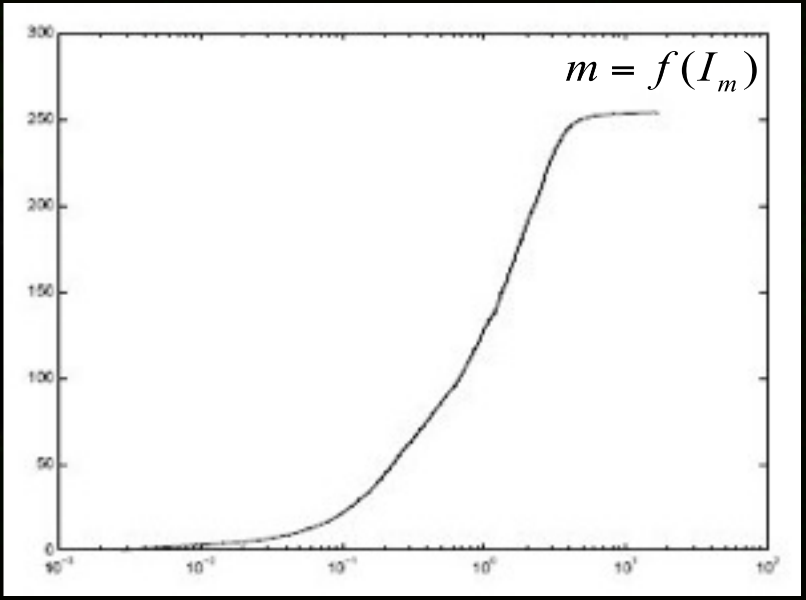

Calibration [Robertson et al. 03]

- Estimate also the weight
  - high derivative → high weight

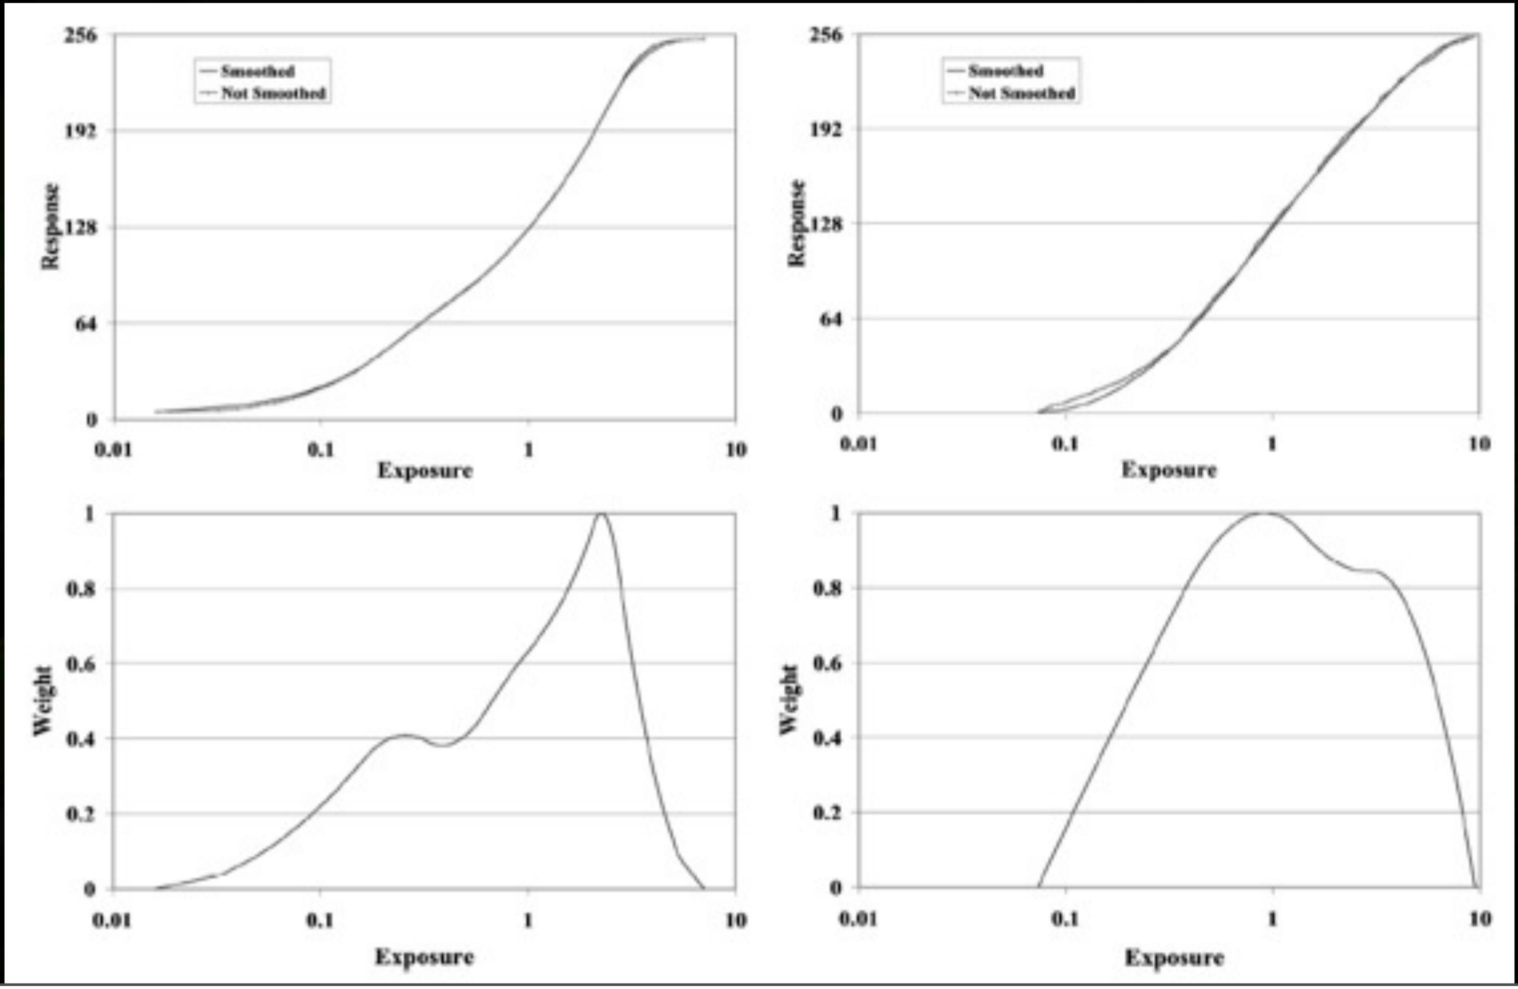

Tone mapping is not easy

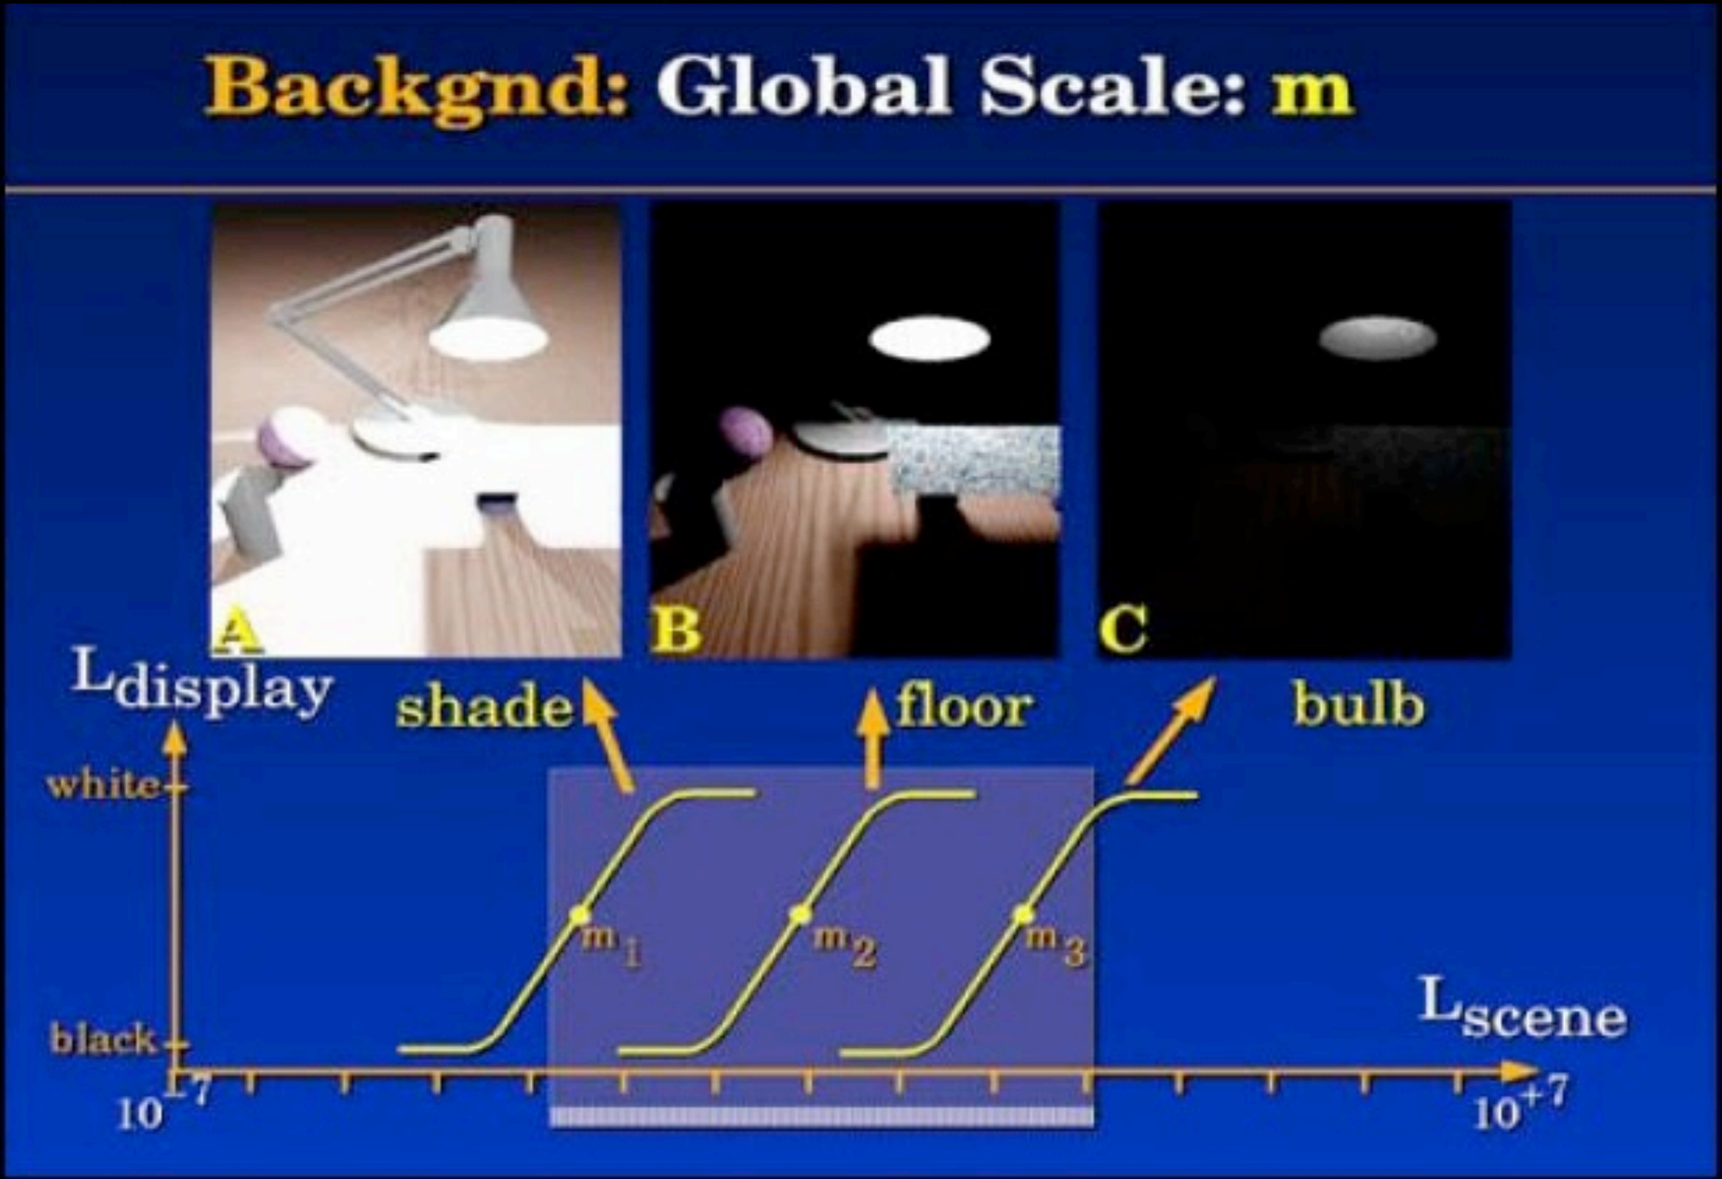

Monday, March 5, 12

Tone mapping is not easy

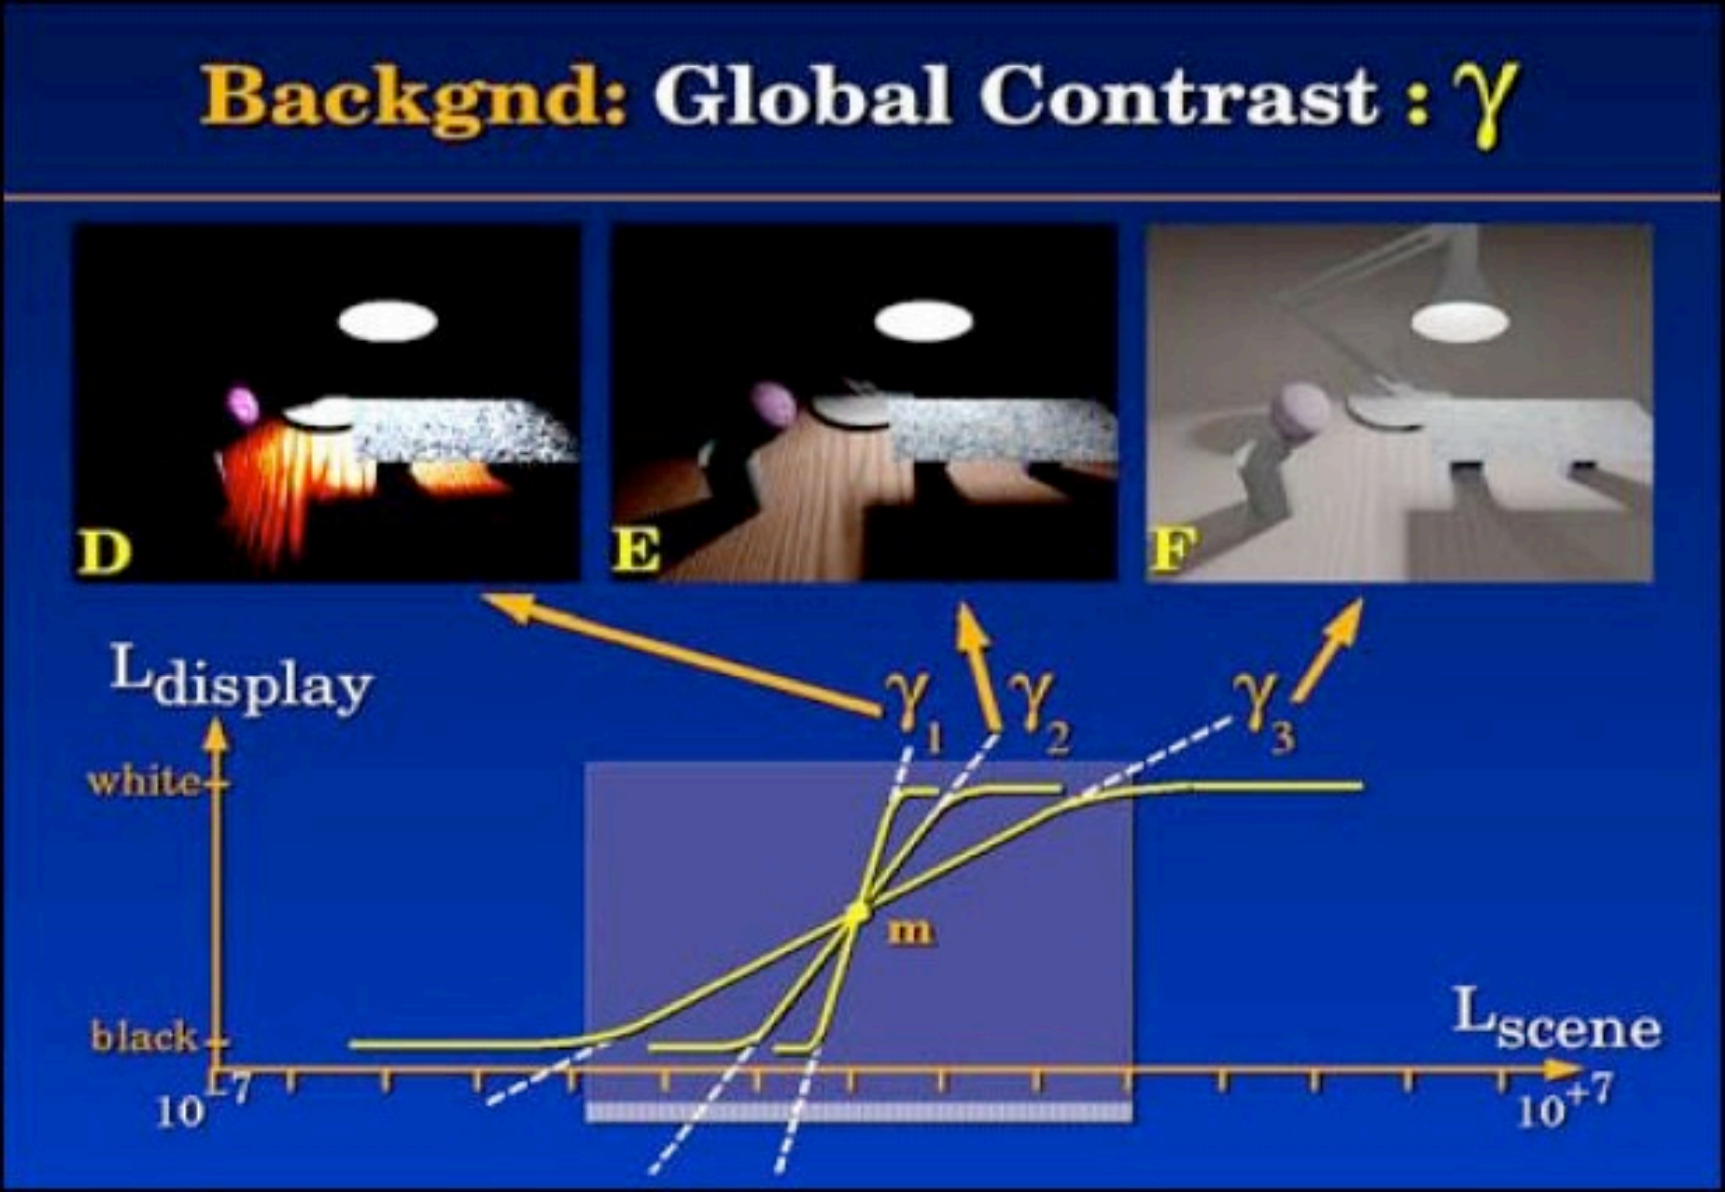

Monday, March 5, 12

Tone mapping is not new

- Painters needed to deal with HDR forever
  - dynamic range of the world is much higher than that of paints
  - change the contrasts to give an effect
- Photographers have done it for a long time
  - dynamic range of the film is much higher than that of paper
  - developing prints required manual tone mapping

Early painters couldn't handle HDR

- Go for local contrast, sacrifice global contrast

Paris Psalter, 10^{th} century

Go for global contrast

- Local contrast suffers
  - a flat painting

Simone Martini, c. 1328

Leonardo masters contrast: HDR painter

Leonardo invents Chiaroscuro

Madonna by Giotto

Madonna by Leonardo

Leonardo invents Chiaroscuro

Madonna by Giotto

Madonna by Leonardo

Caravaggio

Ansel Adams

Ansel Adams

- Design and plan the photo while you are taking it
  - know the medium: both the film, development, and paper
  - standard film & development for the masses using Kodak Brownie
    - global tone map curve, OK on the average
  - virtuosos like Adams
    - capture full dynamic range on the film
    - add spatially varying contrast during development

Dodging and burning

- Hide a part of the print during exposure
  - dodge → keep the bright color of the paper
- Let more light be exposed to a region
  - burn → creates a darker print
- Smooth circular motions & blurry mask avoid artifacts

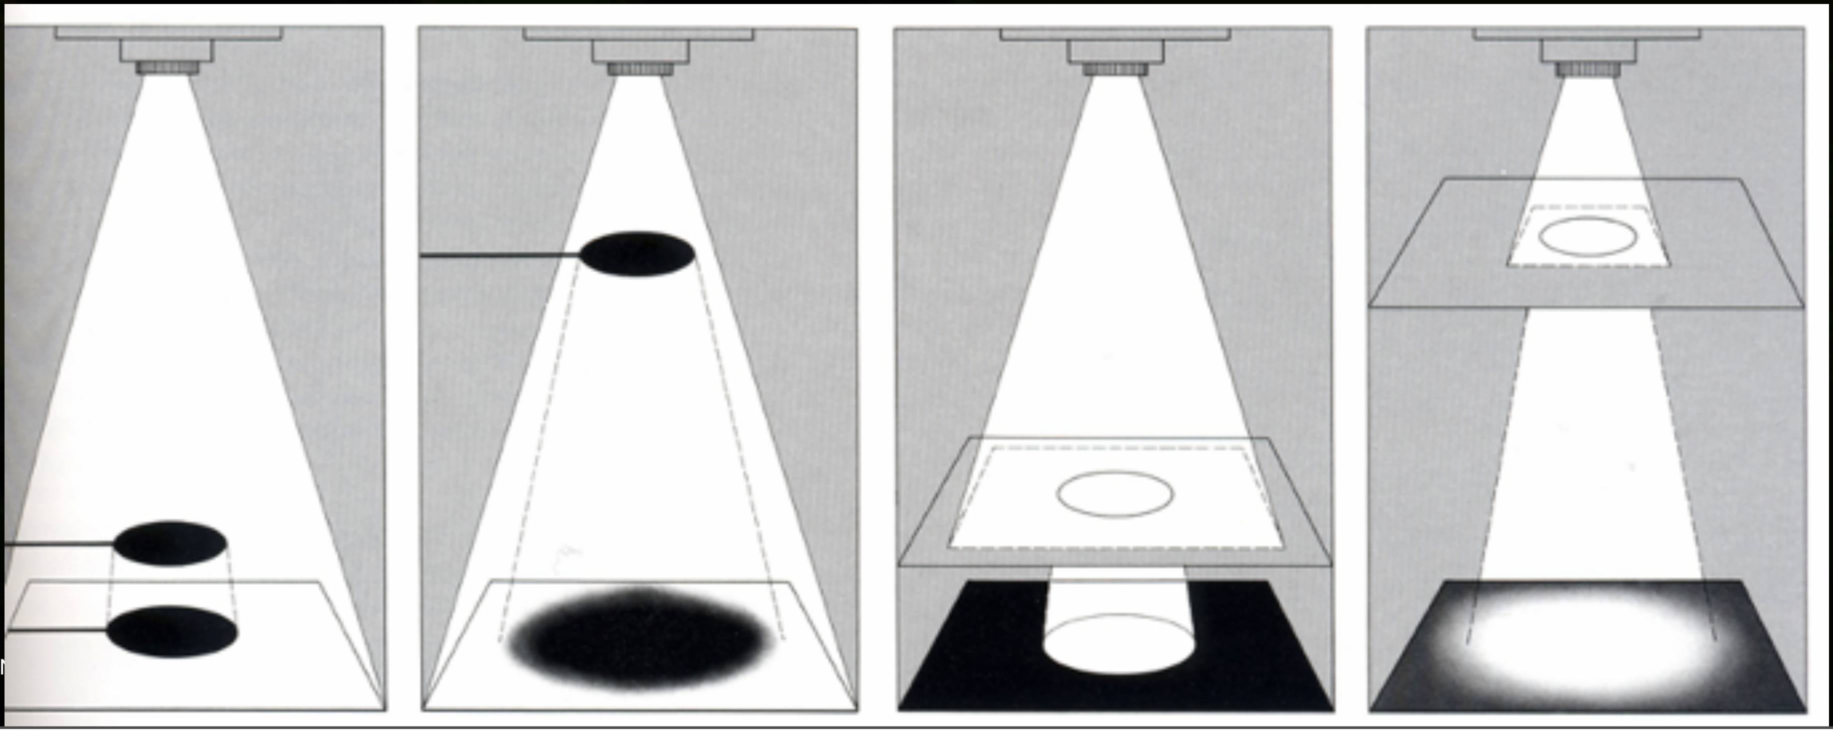

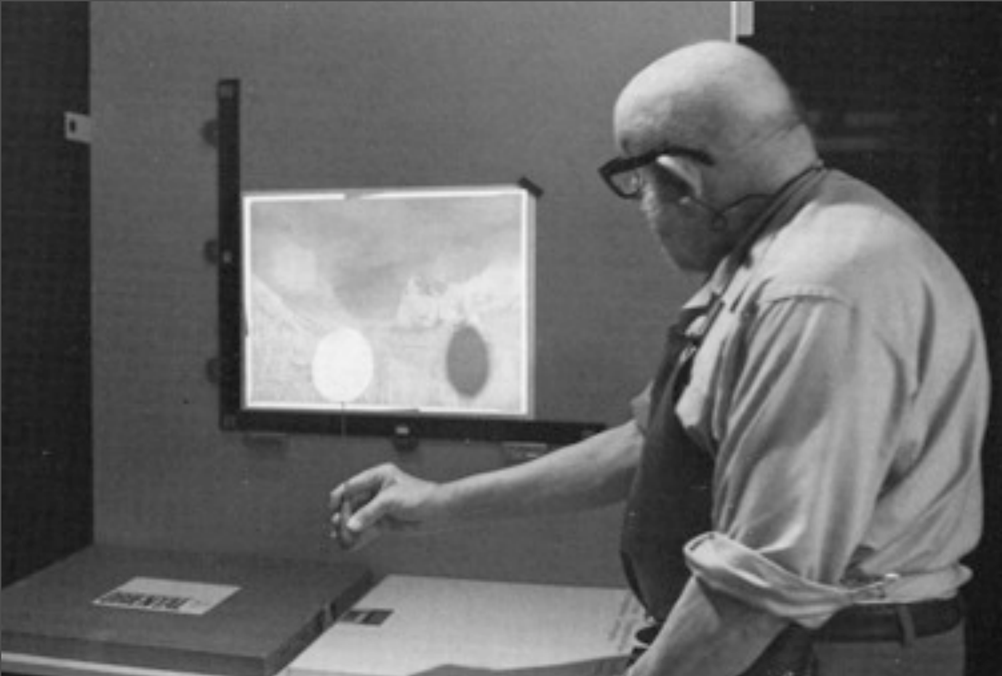

Dodging *holds back light during the basic printing exposure to lighten an area.*

Burning *adds light after the basic exposure to darken an area.*

Manual instructions – repeat for each print

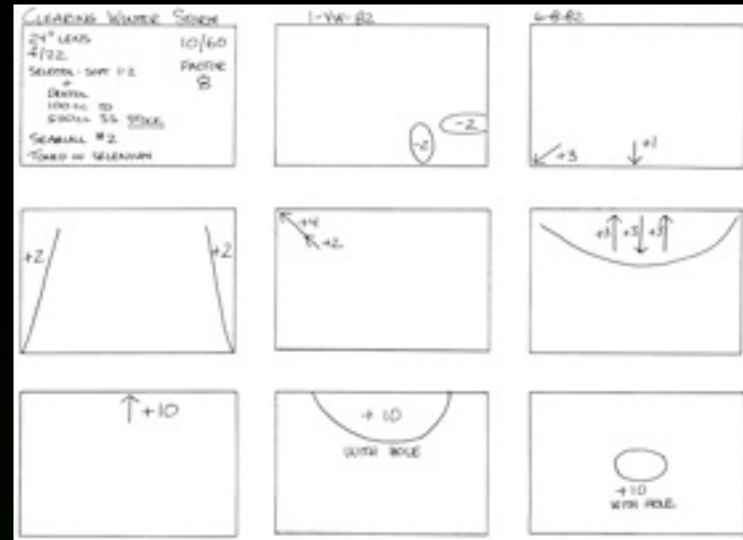

Straight print

After dodging & burning

Contrast reduction in the digital world

- Scene has 1:10,000 contrast, display has 1:100
- Simplest contrast reduction?

Naïve: Gamma compression

- $X \rightarrow X^{\gamma}$ (where γ=0.5 in our case)
- But… colors are washed-out

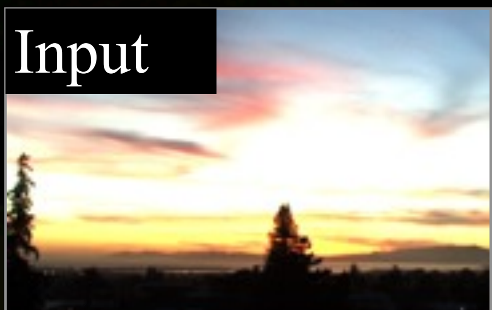

Gamma compression on intensity

- Colors are OK,
  but details (intensity high-frequency) are blurred

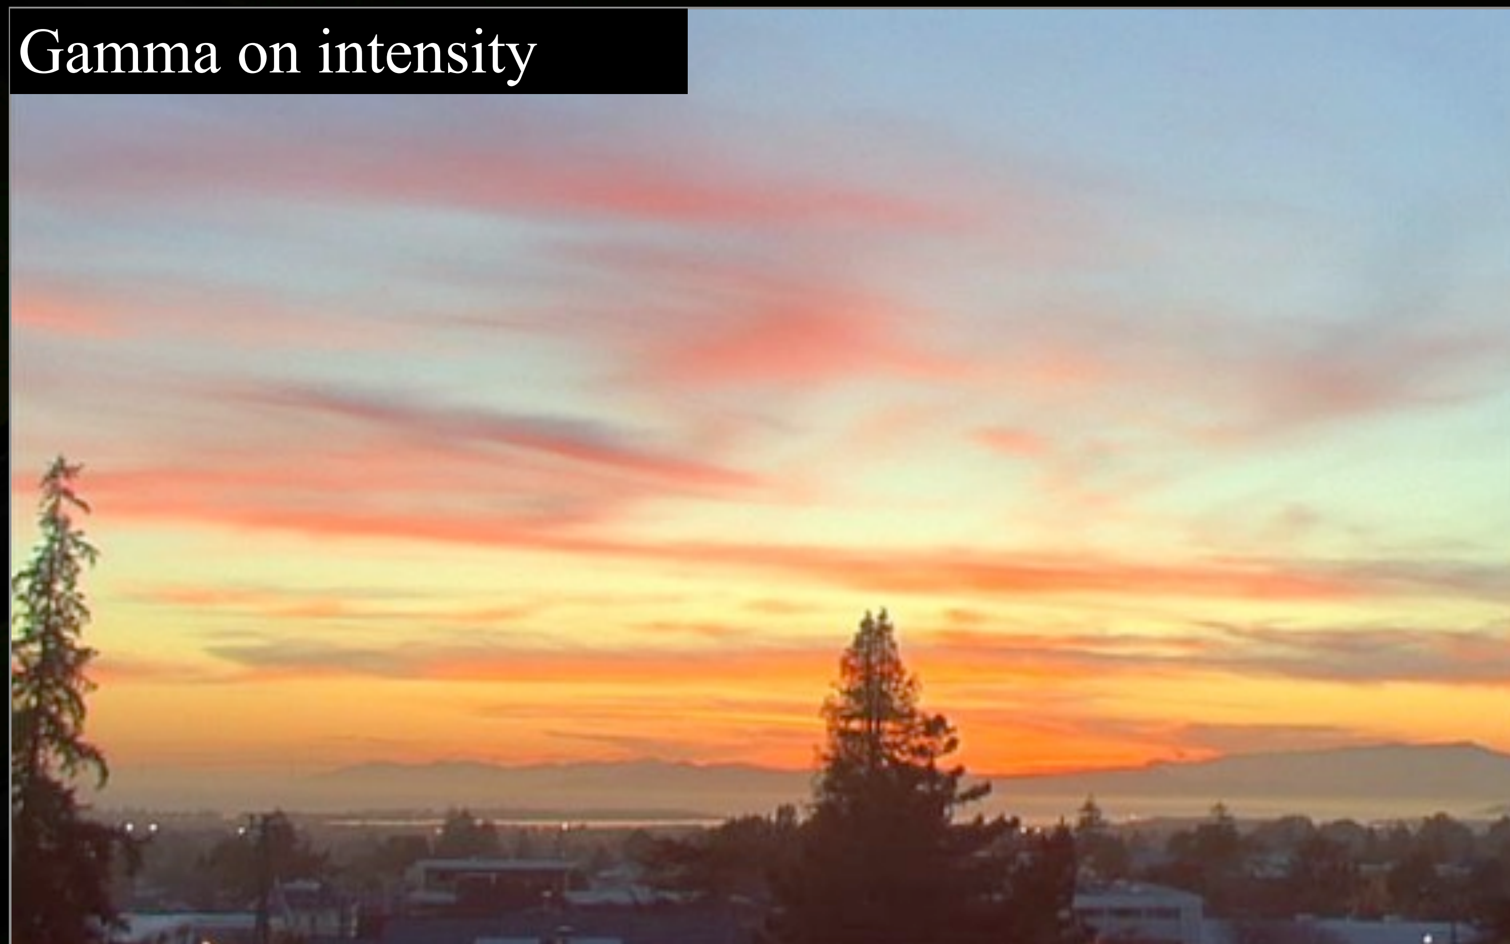

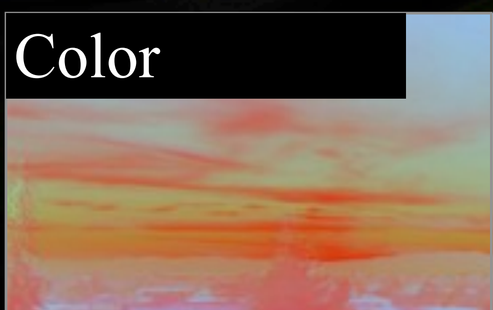

Let highlights saturate

- Darkest 0.1% scaled to display device

Simple global operator (Reinhard et al.)

http://www.cs.utah.edu/~reinhard/cdror
tonemap.pdf

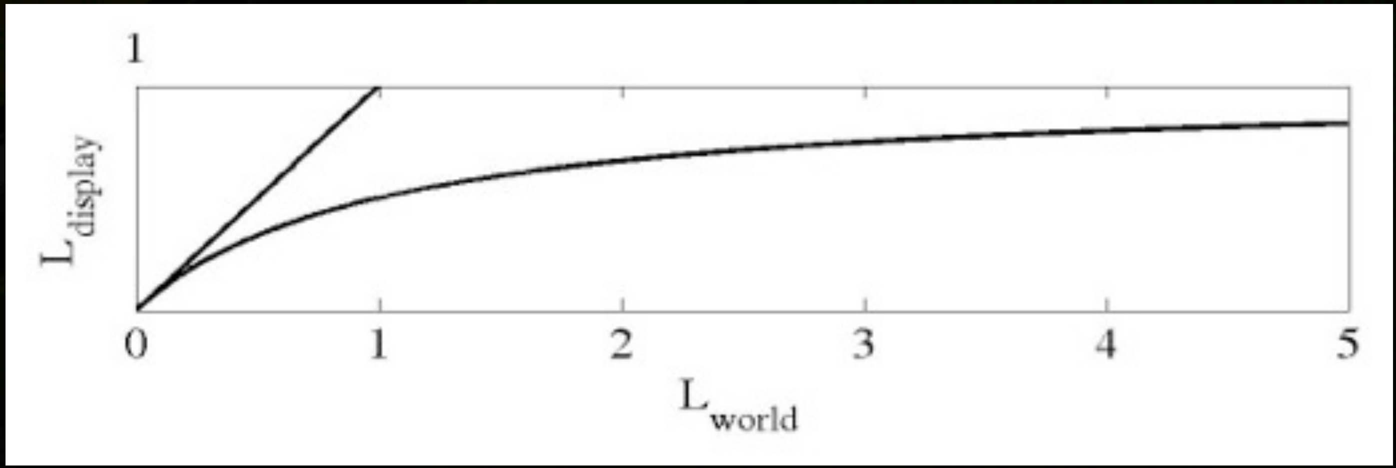

Simple global operator (Reinhard et al.)

- Compression curve needs to
  - bring everything within range
  - leave dark areas alone

http://www.cs.utah.edu/~reinhard/cdrom/tonemap.pdf

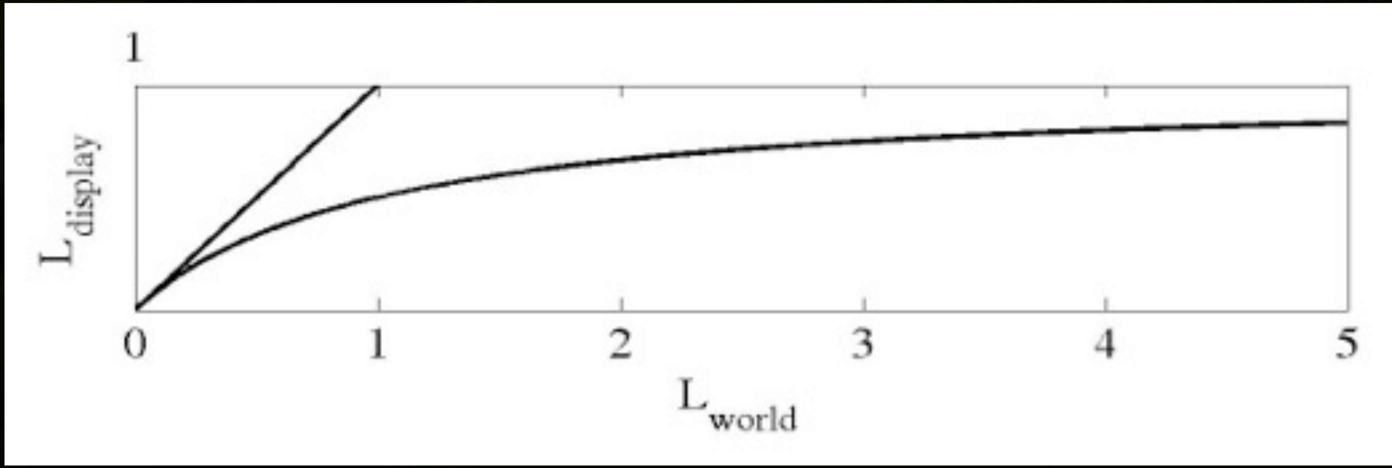

Simple global operator (Reinhard et al.)

- Compression curve needs to
  - bring everything within range
  - leave dark areas alone
- In other words
  - asymptote is 1
  - derivative at 0 is 1

http://www.cs.utah.edu/~reinhard/cdrom/tonemap.pdf

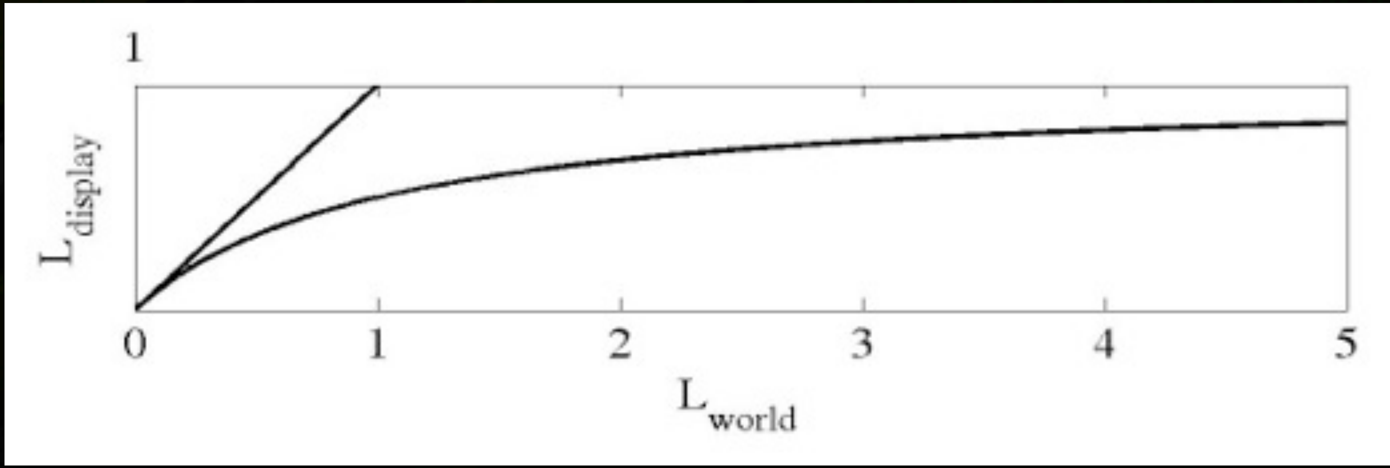

Simple global operator (Reinhard et al.)

- Compression curve needs to
  - bring everything within range
  - leave dark areas alone
- In other words
  - asymptote is 1
  - derivative at 0 is 1

http://www.cs.utah.edu/~reinhard/cdrom/tonemap.pdf

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

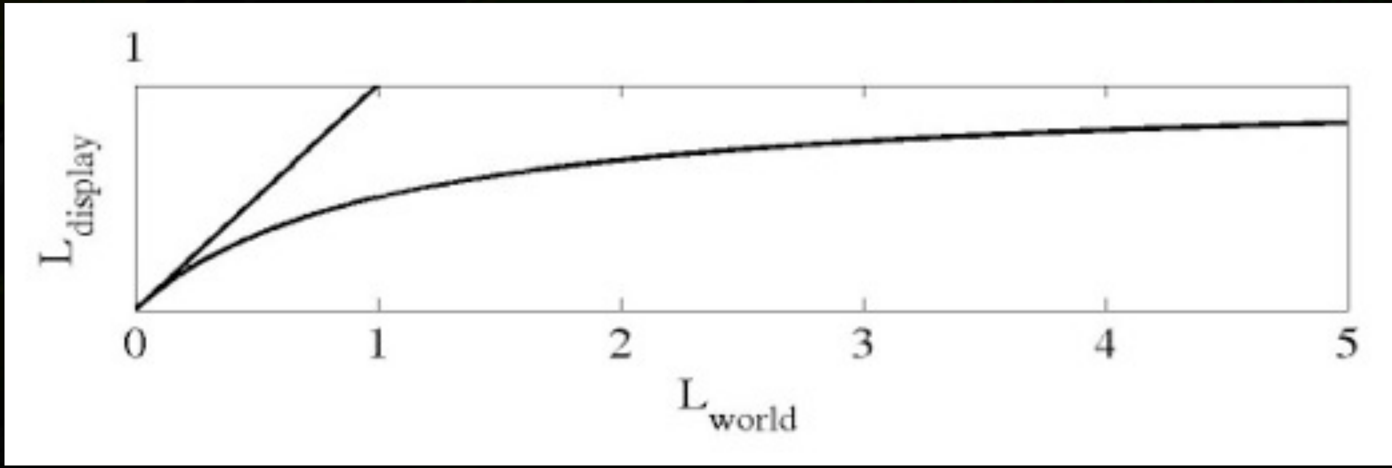

The same in log L
closer to brightness perception

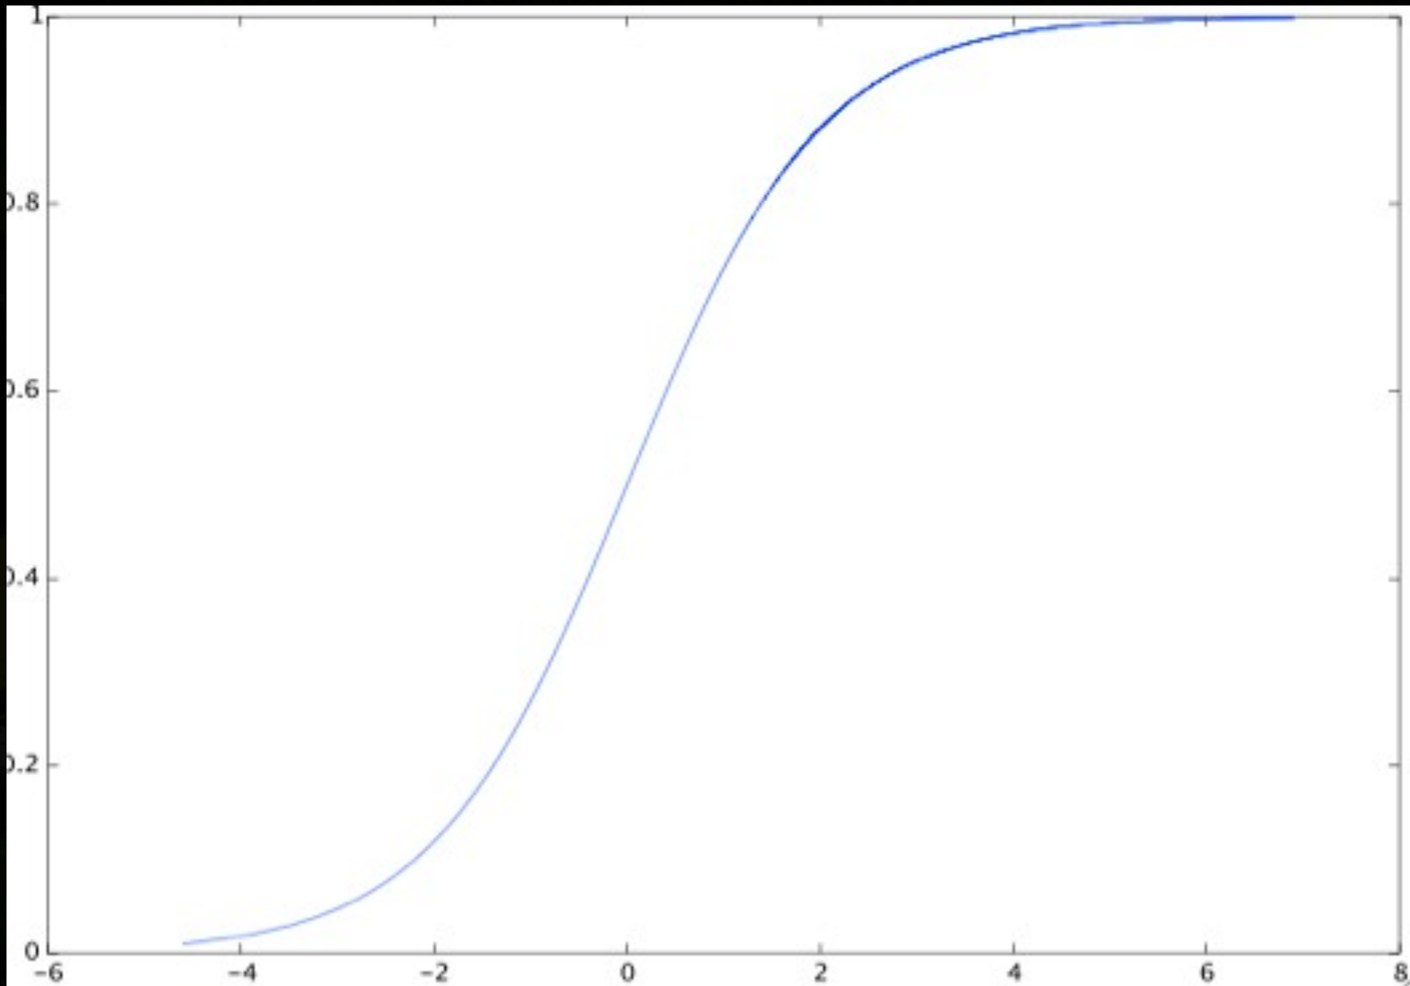

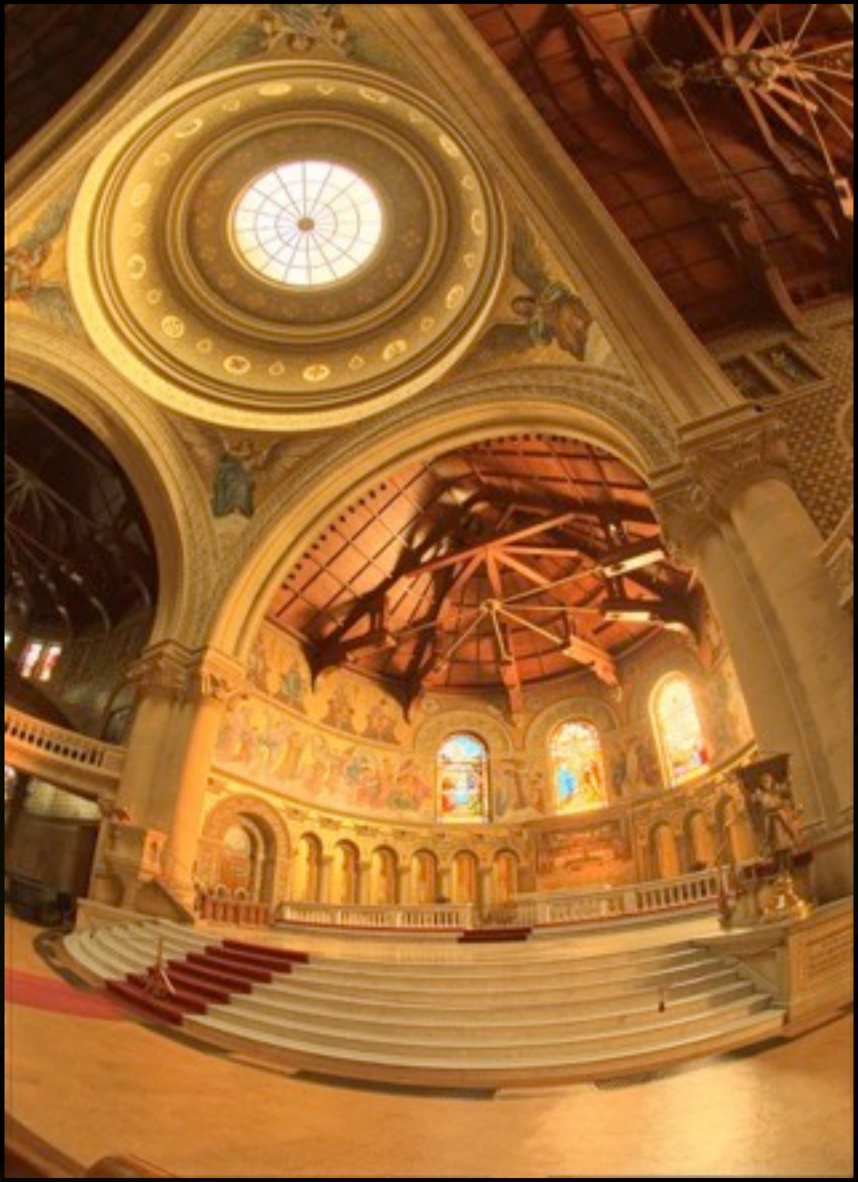

Reinhard operator

Darkest 0.1% scaled to display device


Histogram adjustment [Ward et al. 1997]

- Histogram equalization
  - well-known method to increase contrast
  - luminance is not evenly spread, spread it
- Basic approach
  - lump pixels with 1deg area together
  - calculate histogram in log(luminance) space
- Problem
  - doesn’t just compress contrast, but also expands it
- Solution
  - put a ceiling to contrast by trimming large bins
  - not equalization, but adjustment

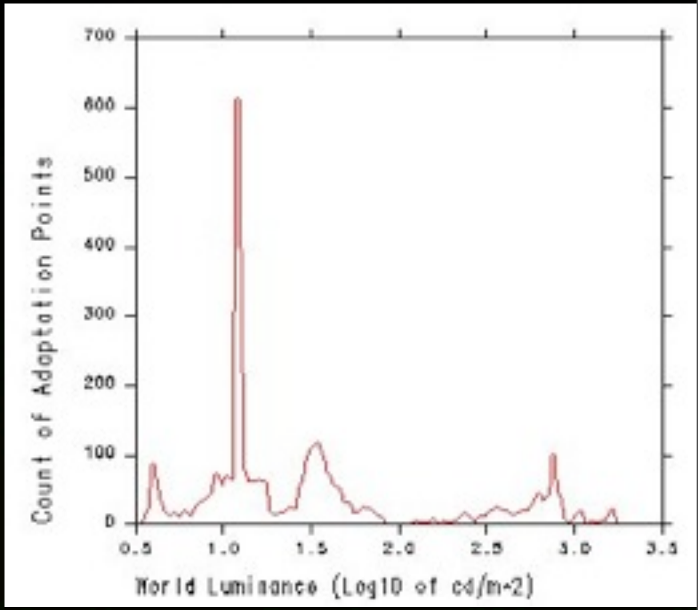

Equalization vs. adjustment

Linear Adjustment

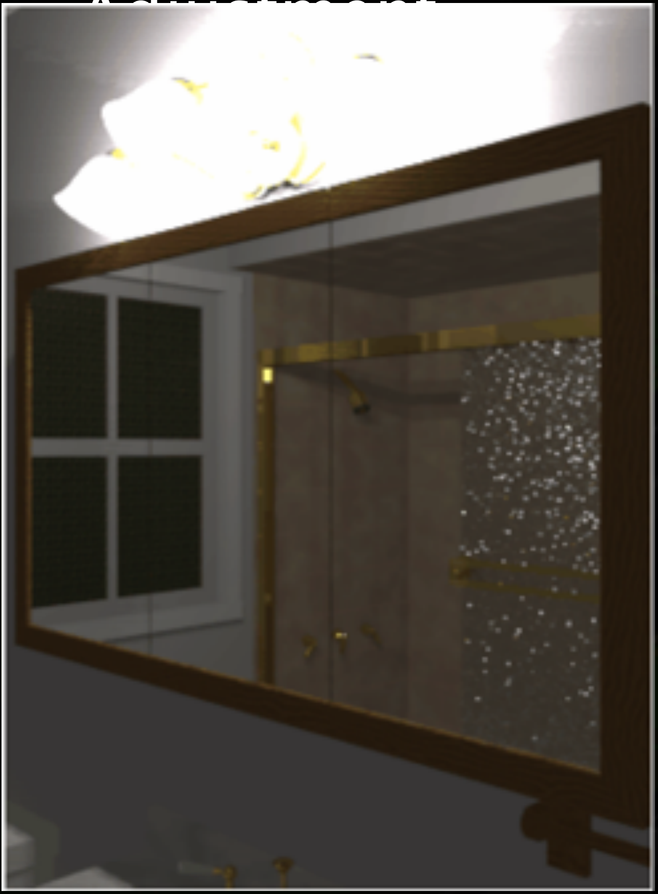

Equalization

Equalization vs. adjustment

Linear Adjustment

Equalization

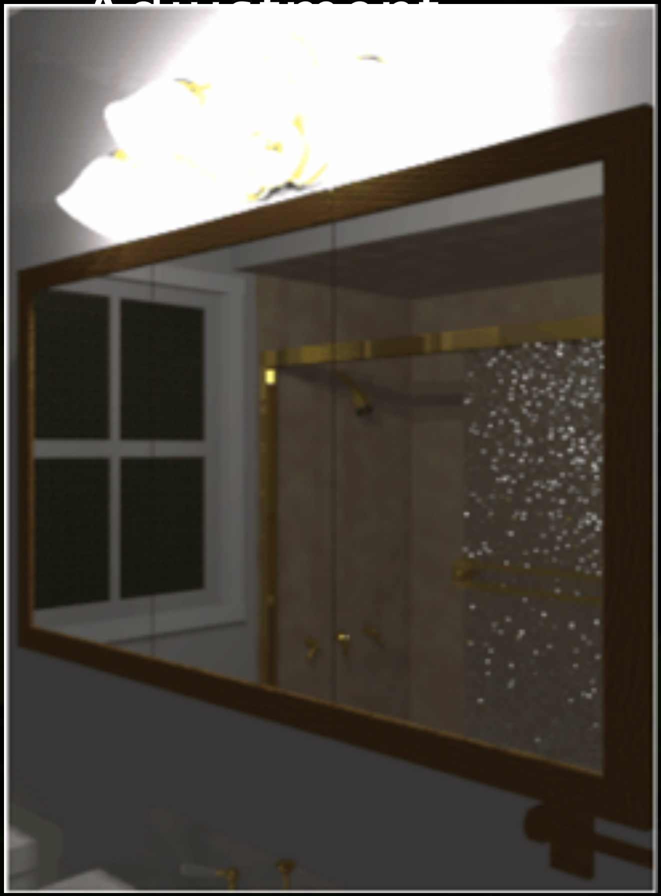

Equalization vs. adjustment

Linear Adjustment

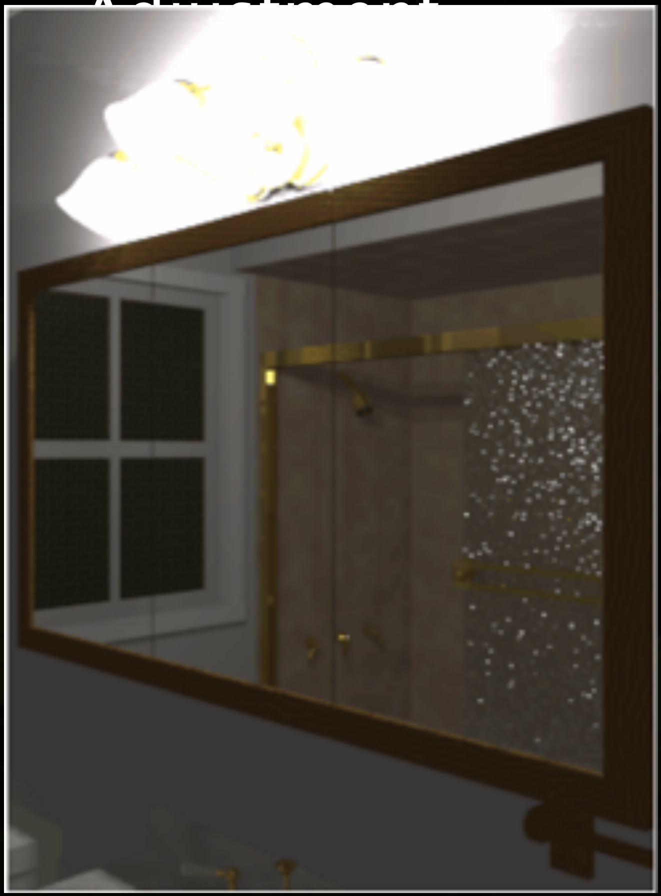

Equalization

Oppenheim 1968, Chiu et al. 1993

- Reduce contrast of low-frequencies
- Keep high frequencies

Contrast sensitivity function

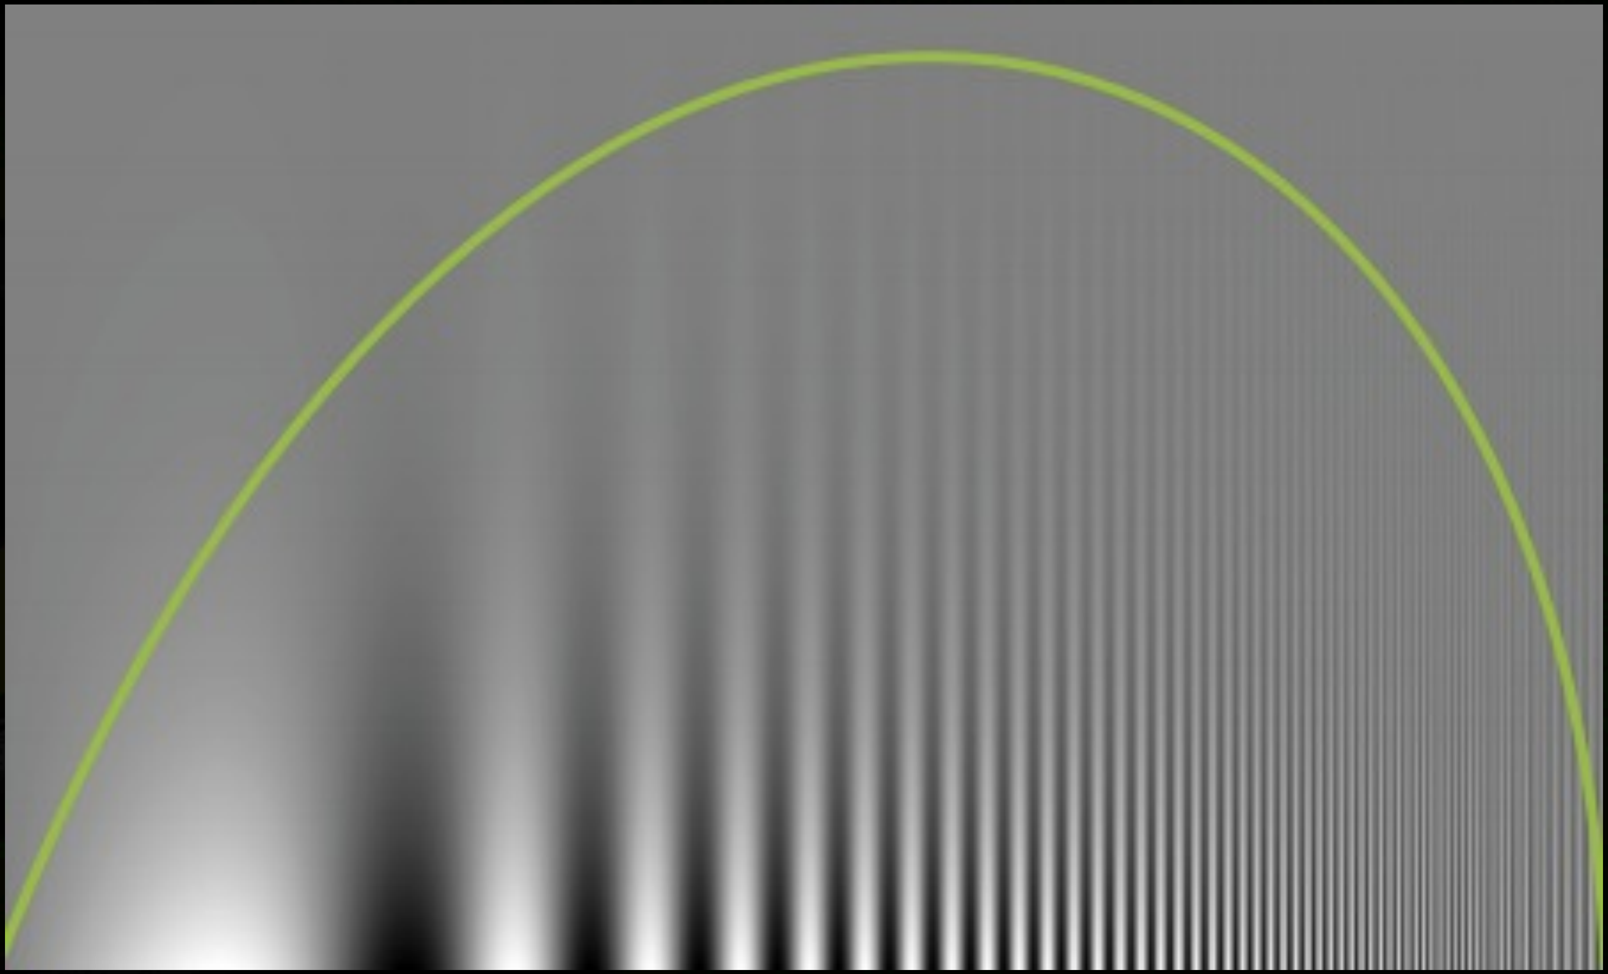

Contrast sensitivity function

- Low sensitivity
  - to low frequencies
- Higher sensitivity
  - medium to high frequencies
- Most methods to deal with dynamic range
  - reduce the contrast of low frequencies
  - but keep the color

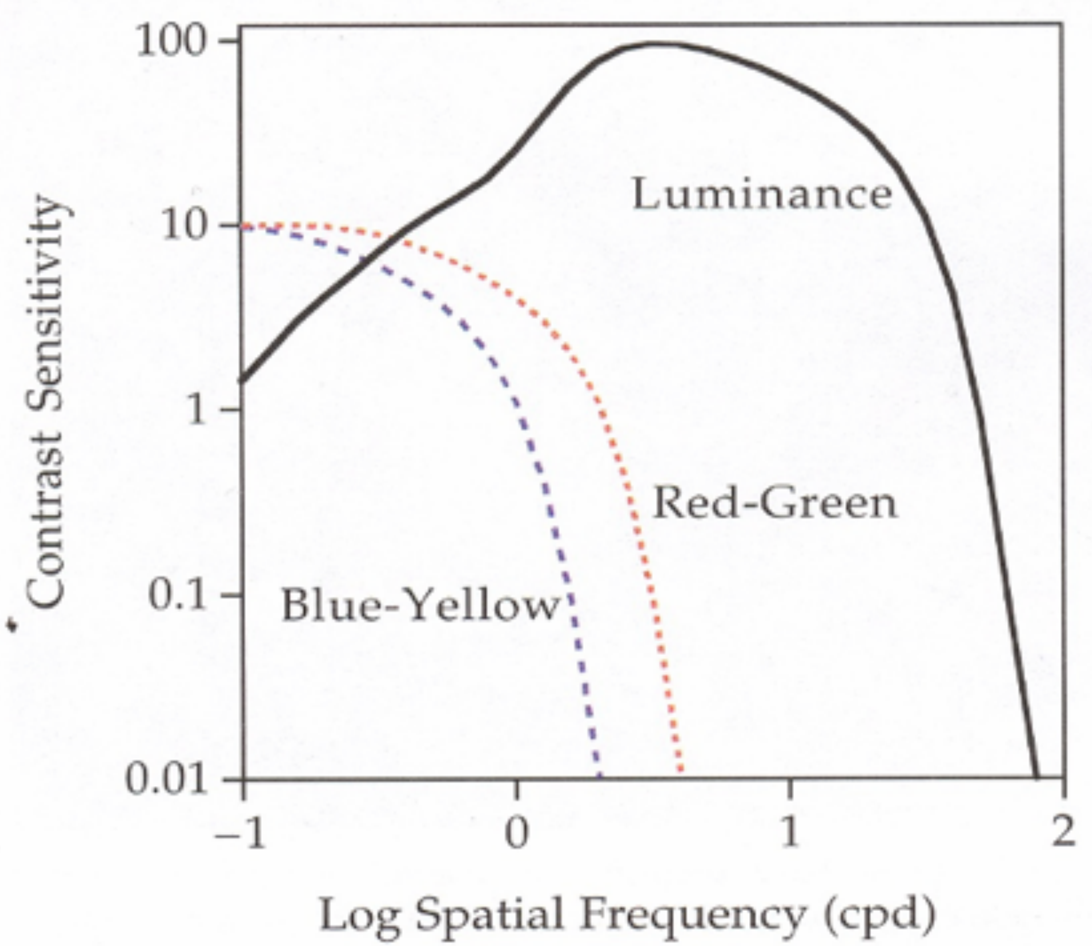

Figure 1-18. Spatial contrast sensitivity functions for luminance and chromatic contrast.

The halo nightmare

- For strong edges
- Because they contain high frequencies

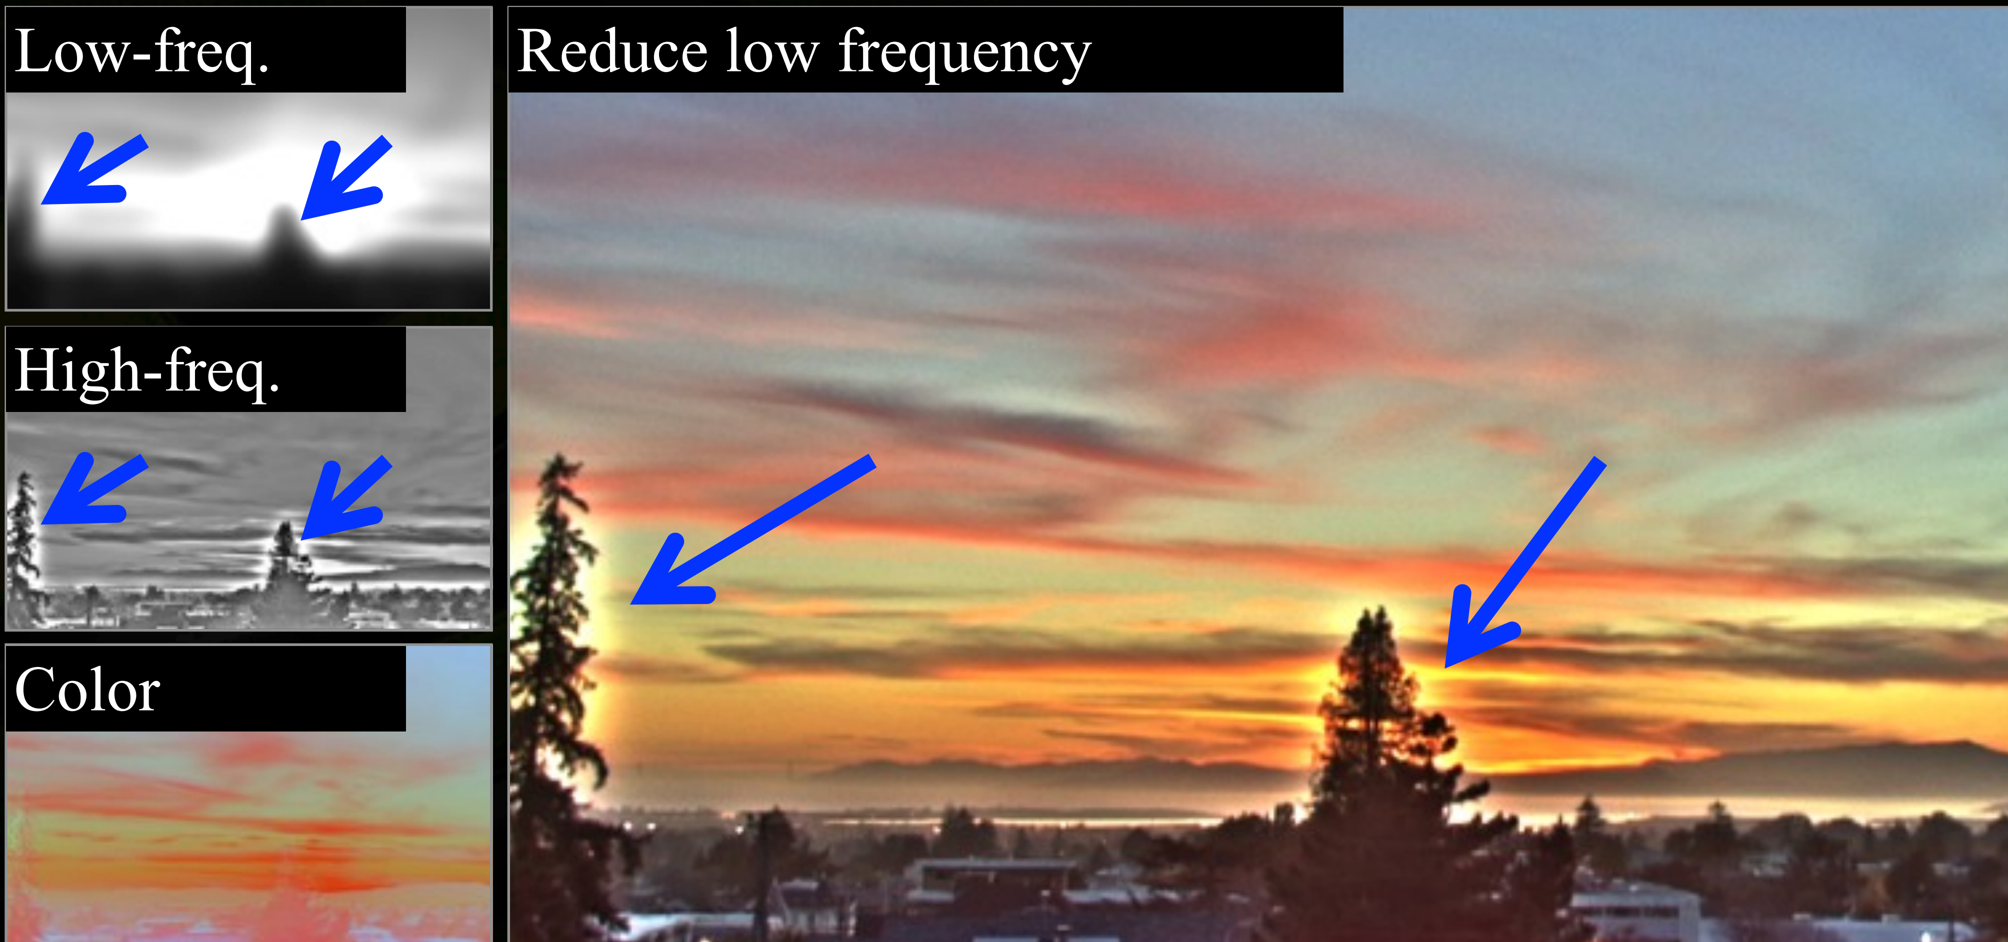

Durand & Dorsey 2002: Bilateral

- Use non-linear filtering
  - to better separate details
  - without blurring across edges

Durand & Dorsey 2002: Bilateral

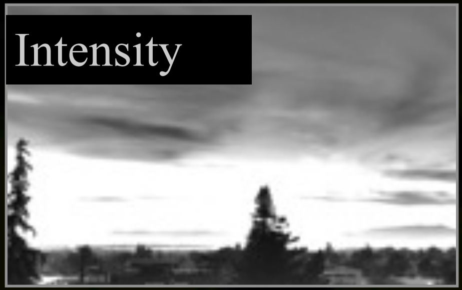

- Use non-linear filtering
  - to better separate details
  - without blurring across edges

Durand & Dorsey 2002: Bilateral

- Use non-linear filtering
  - to better separate details
  - without blurring across edges

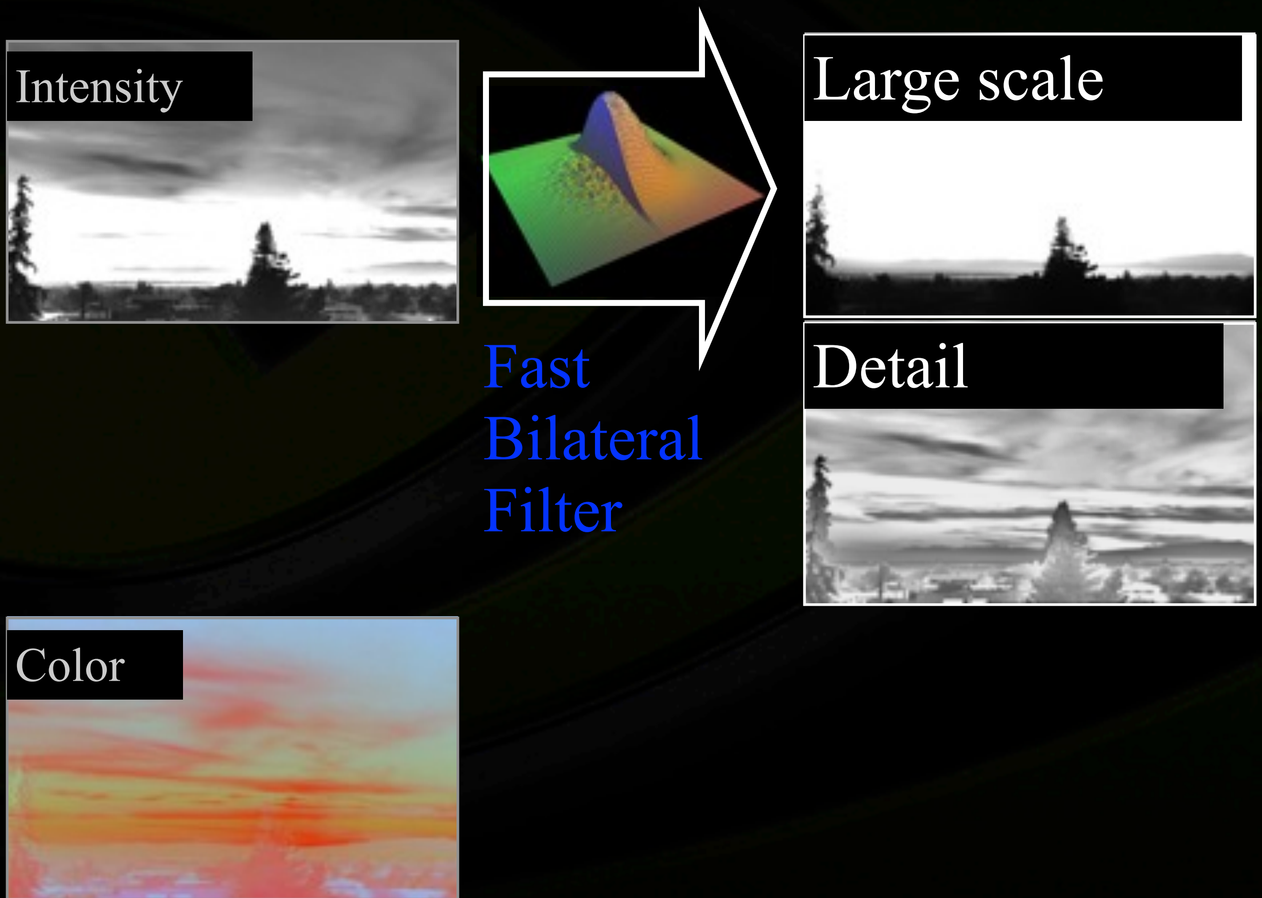

Durand & Dorsey 2002: Bilateral

- Use non-linear filtering
  - to better separate details
  - without blurring across edges

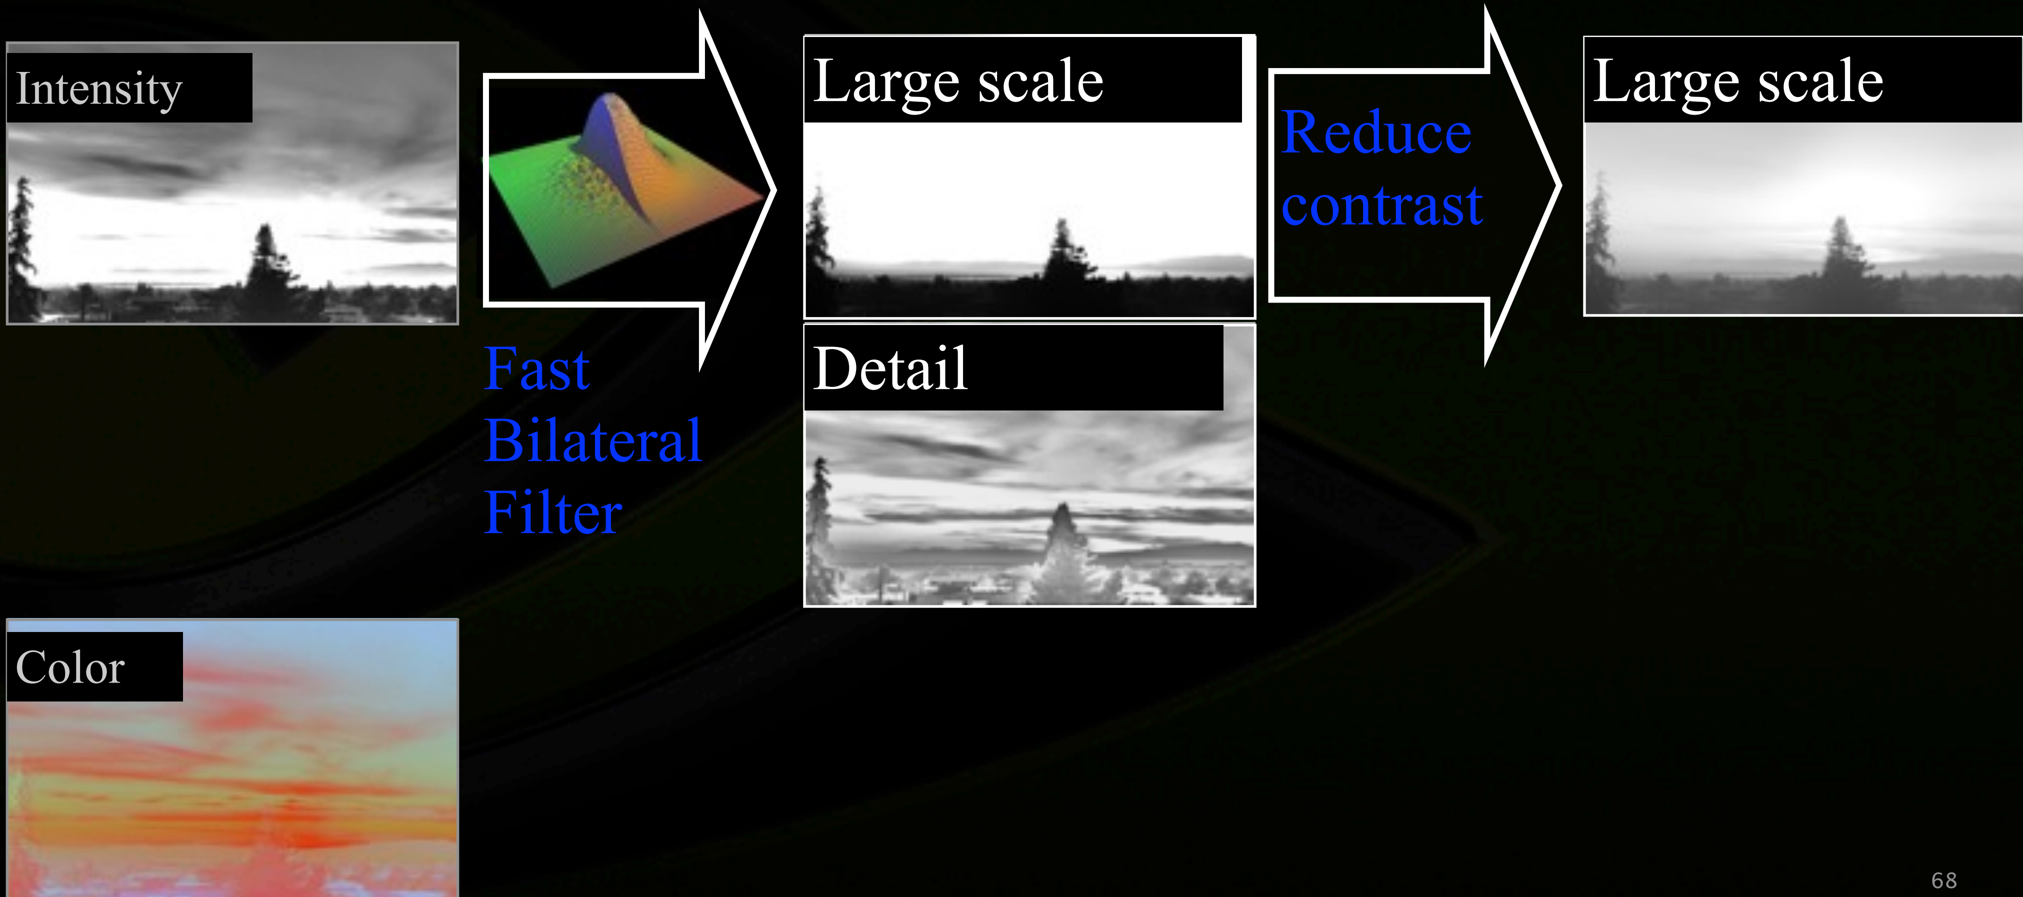

Durand & Dorsey 2002: Bilateral

- Use non-linear filtering
  - to better separate details
  - without blurring across edges

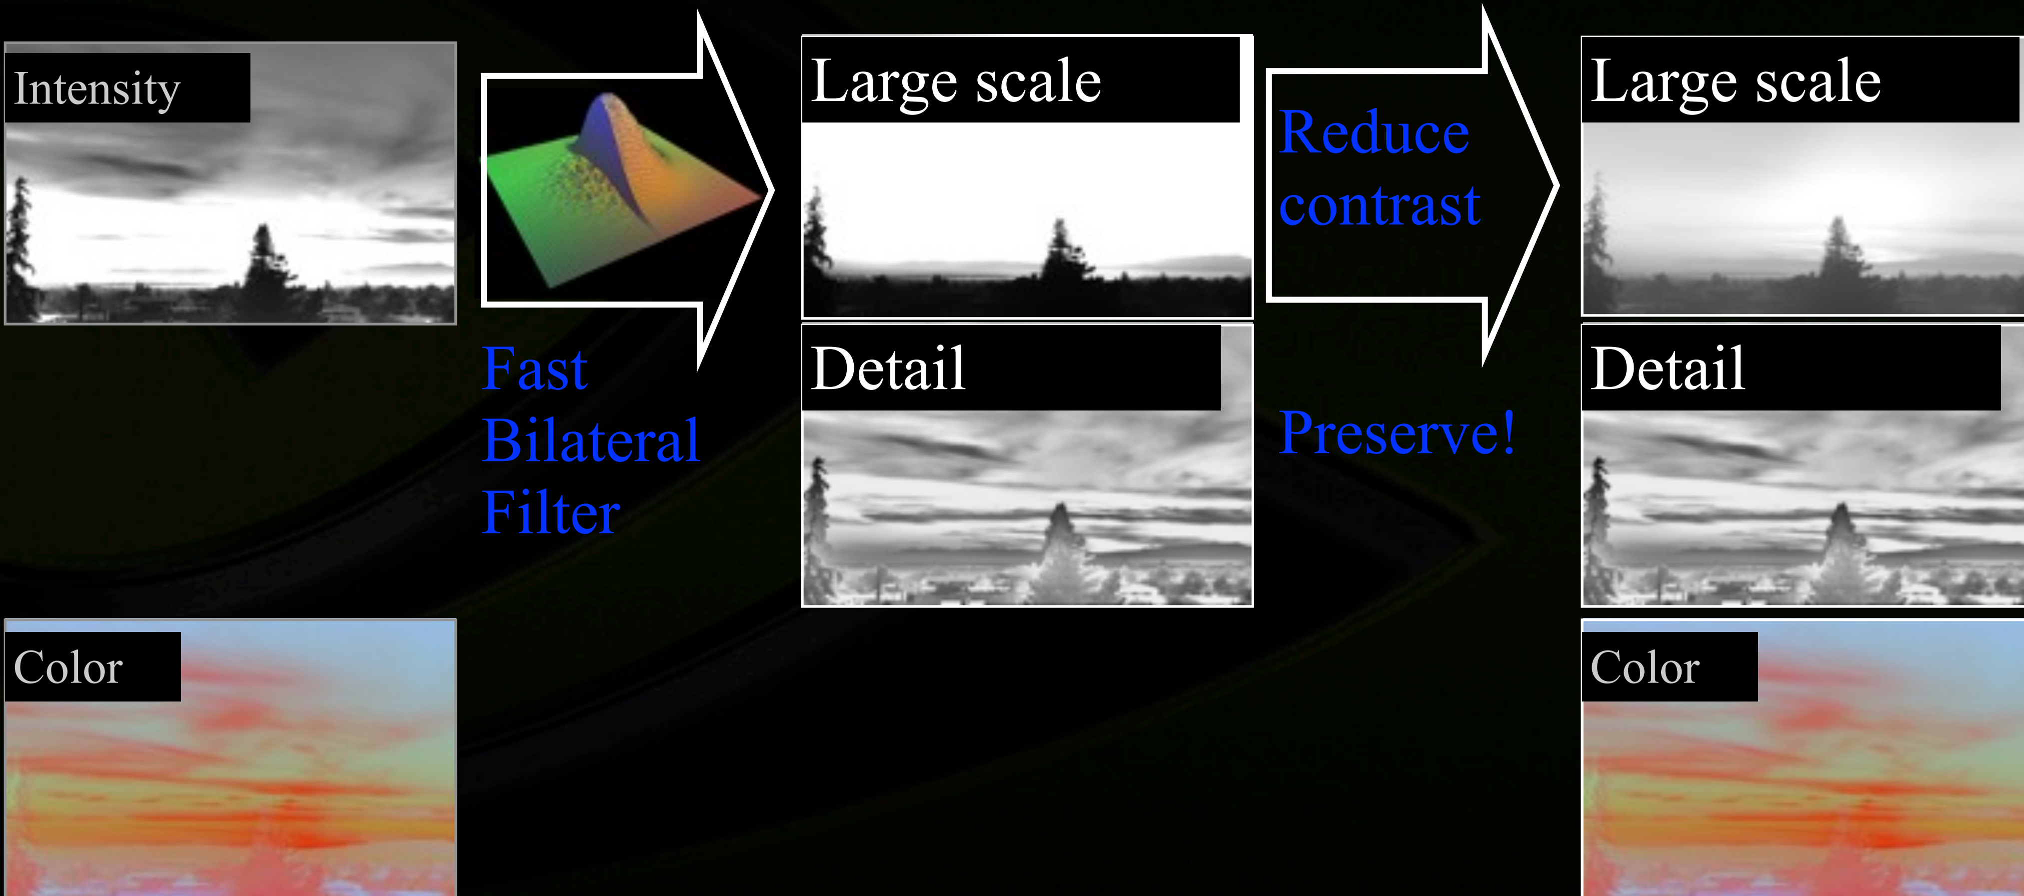

Durand & Dorsey 2002: Bilateral

- Use non-linear filtering
  - to better separate details
  - without blurring across edges

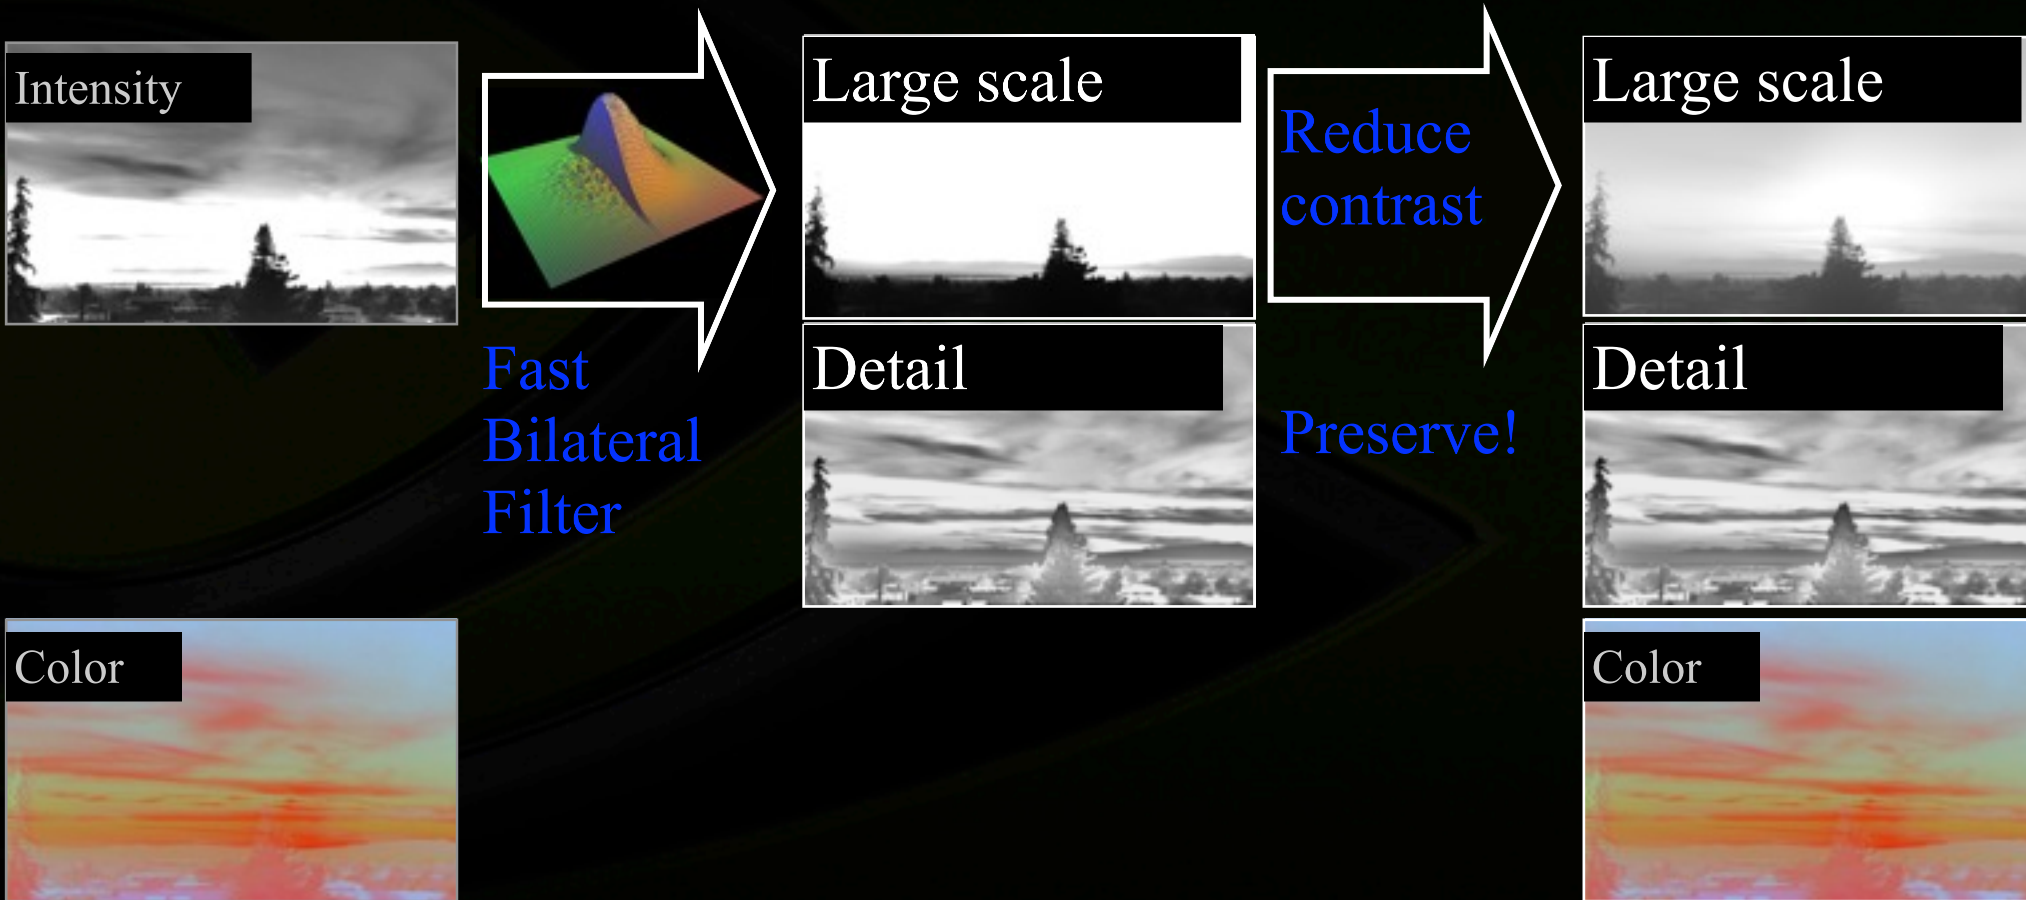

SIGGRAPH2005

Compressing and Companding High Dynamic Range Images with Subband Architectures

Yuanzhen Li, Lavanya Sharan,

Edward Adelson

Massachusetts Institute of Technology

Halo Artifacts

Halos are widely believed to be inherent with multiscale methods.

We fixed the halos

A subband/wavelet method that minimizes halos.

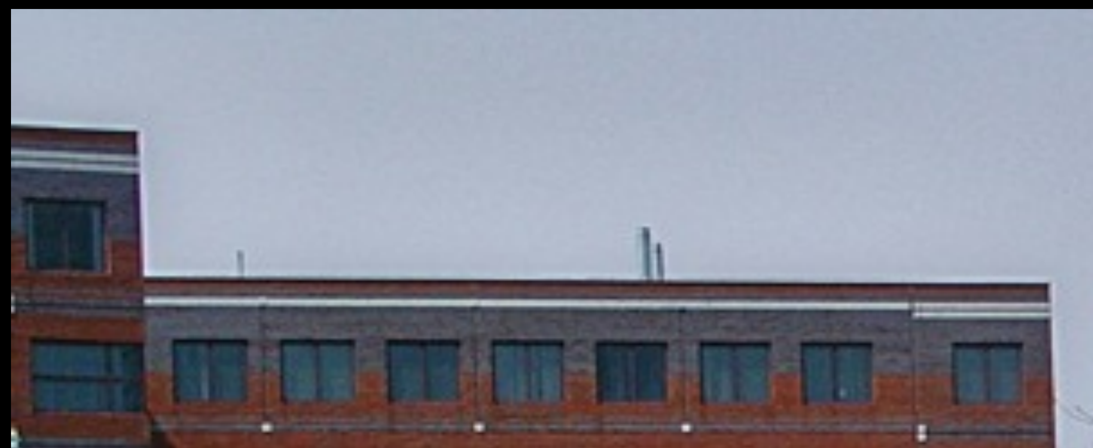

Range Compression

Method: Gamma or log on intensities
Problem: loss of detail

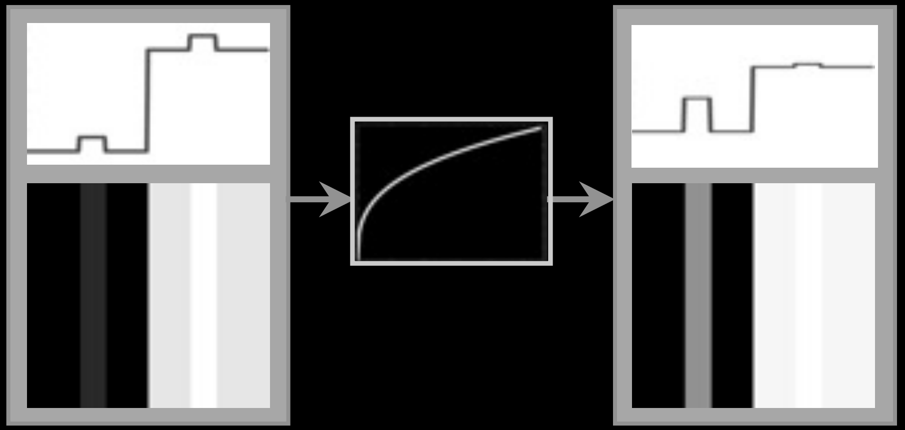

Range Compression

Method: Gamma or log on intensities
Problem: loss of detail

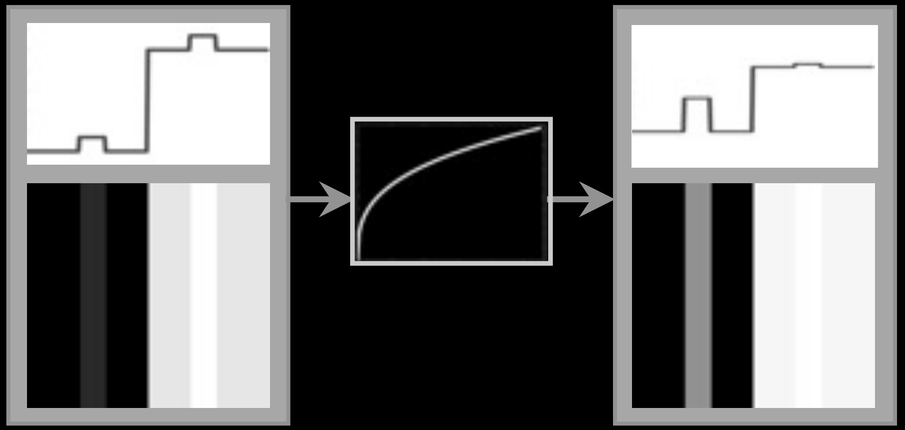

Solution: filtering
Problem: halos

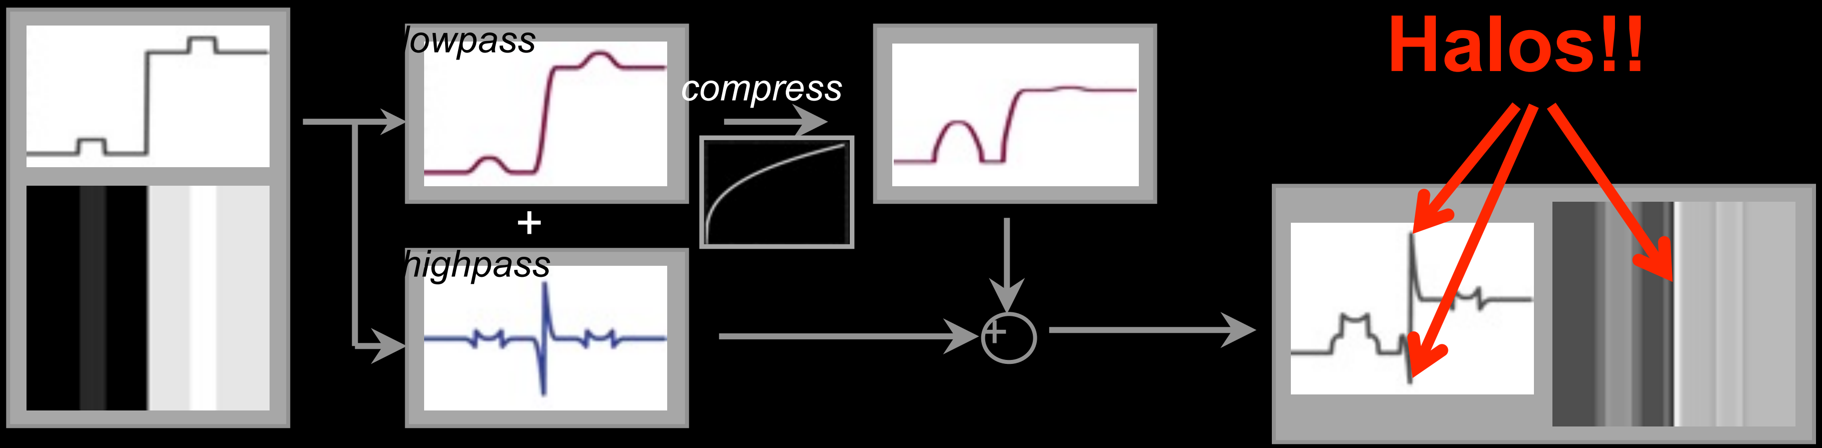

How did we fix the halos?

- Analysis-synthesis subband architecture, e.g., wavelets

- Smooth gain control on subbands

Analysis-Synthesis Architecture

Analysis filters

Synthesis filters

Original
signal

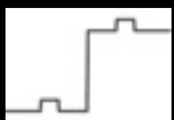

. . .

N

**Non-
linearity**

Reconstructed
signal

Point Nonlinearity on Subbands

Original subband

point nonlinearity limits range

Modified subband

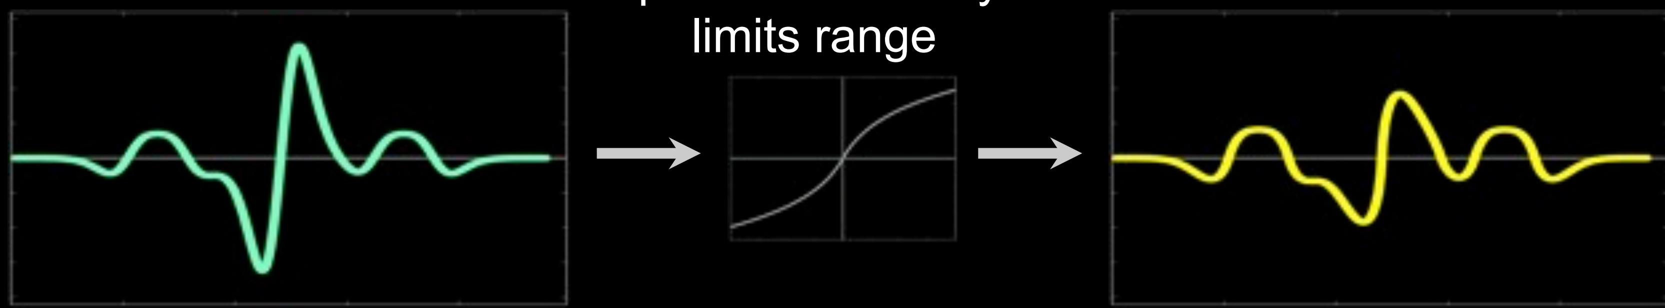

Point Nonlinearity on Subbands

Original subband

point nonlinearity limits range

Modified subband

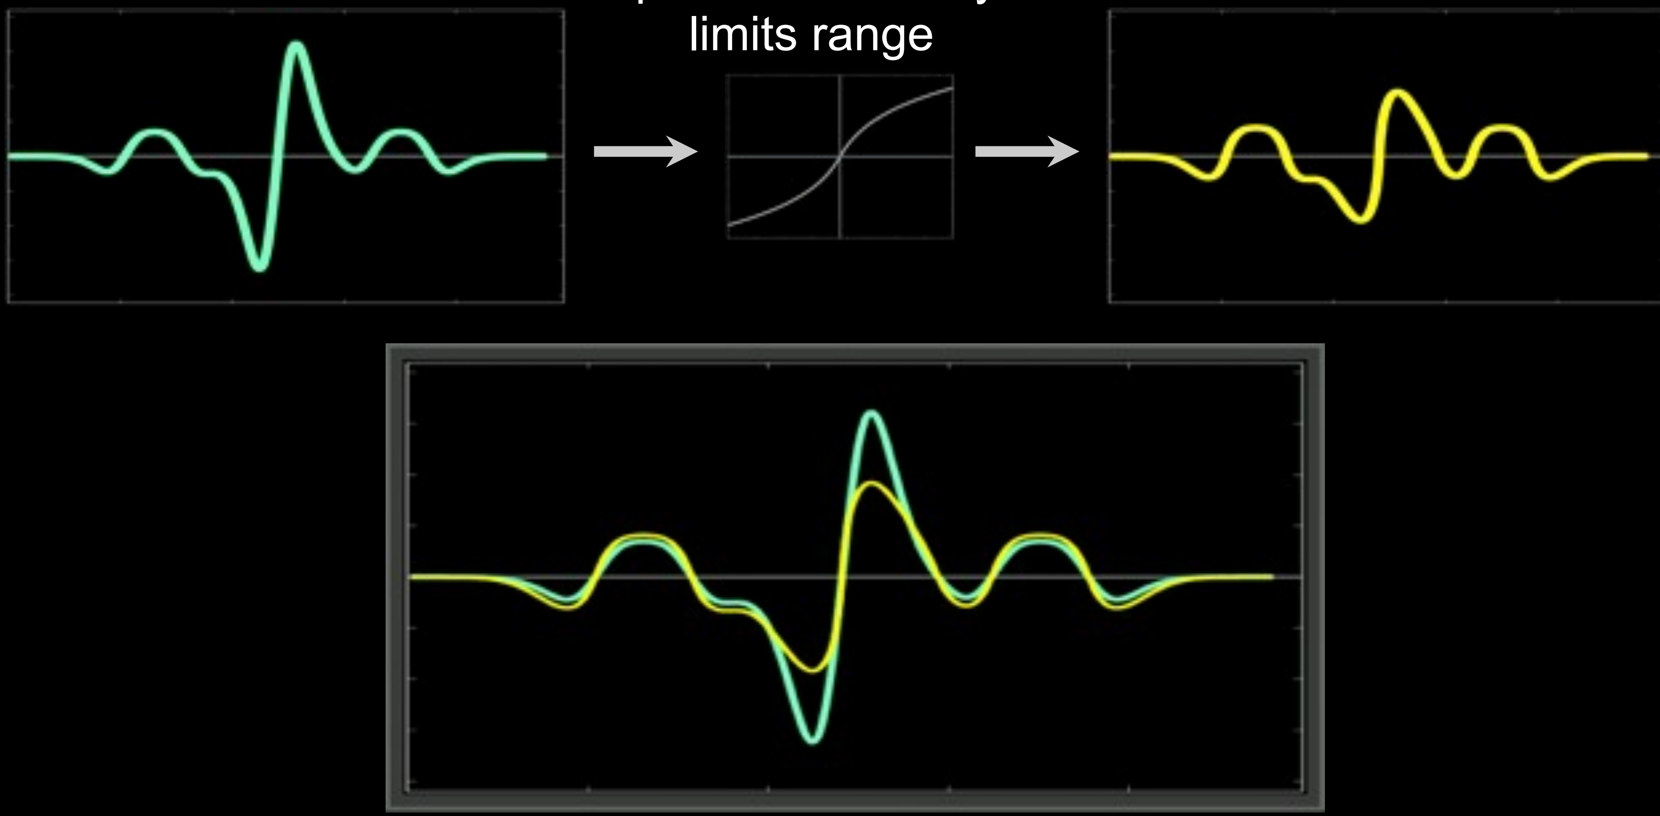

Problem: Nonlinear distortion

Point Nonlinearity on Subbands

Original subband

point nonlinearity limits range

Modified subband

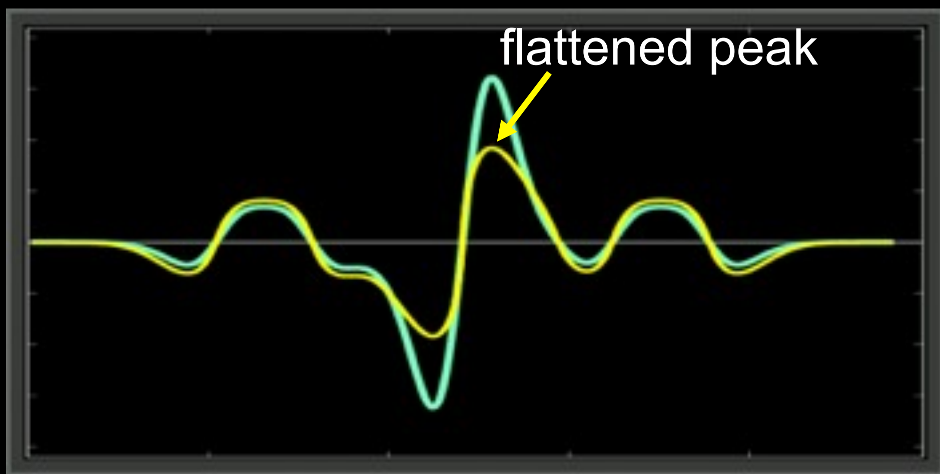

Problem: Nonlinear distortion

Smooth Gain Control

Smooth Gain Control

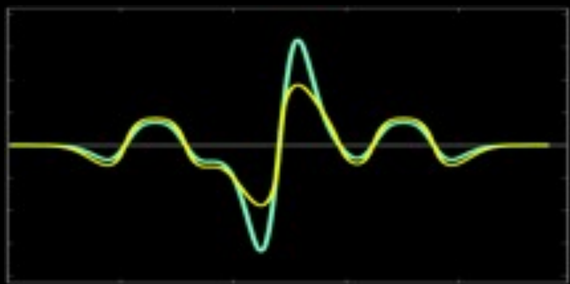

Smooth Gain Control

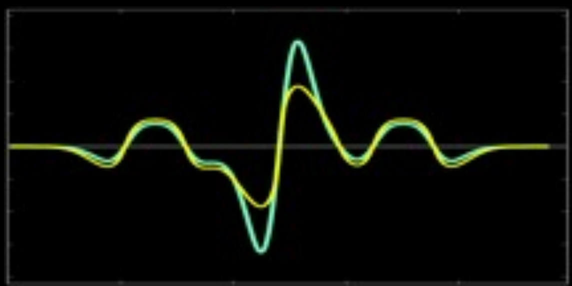

$gain(x) = b'(x) / b(x) =$

$=$

Smooth Gain Control

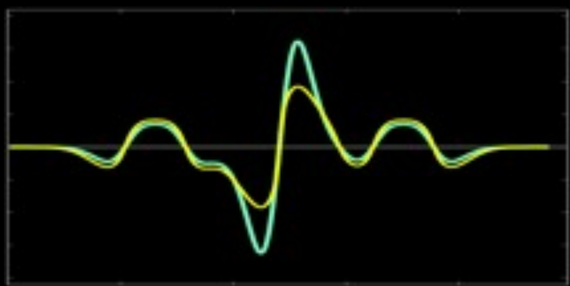

$$gain(x) = b'(x) \,/\, b(x) =$$

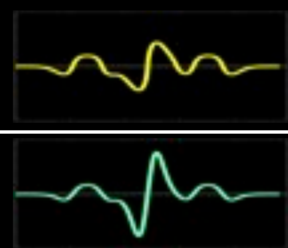

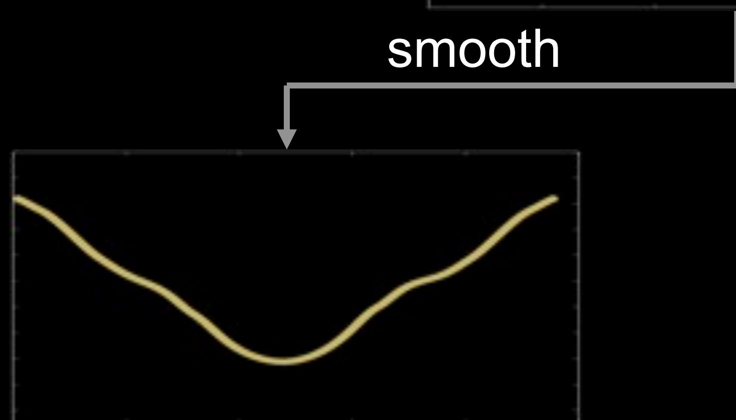

Smooth Gain Control

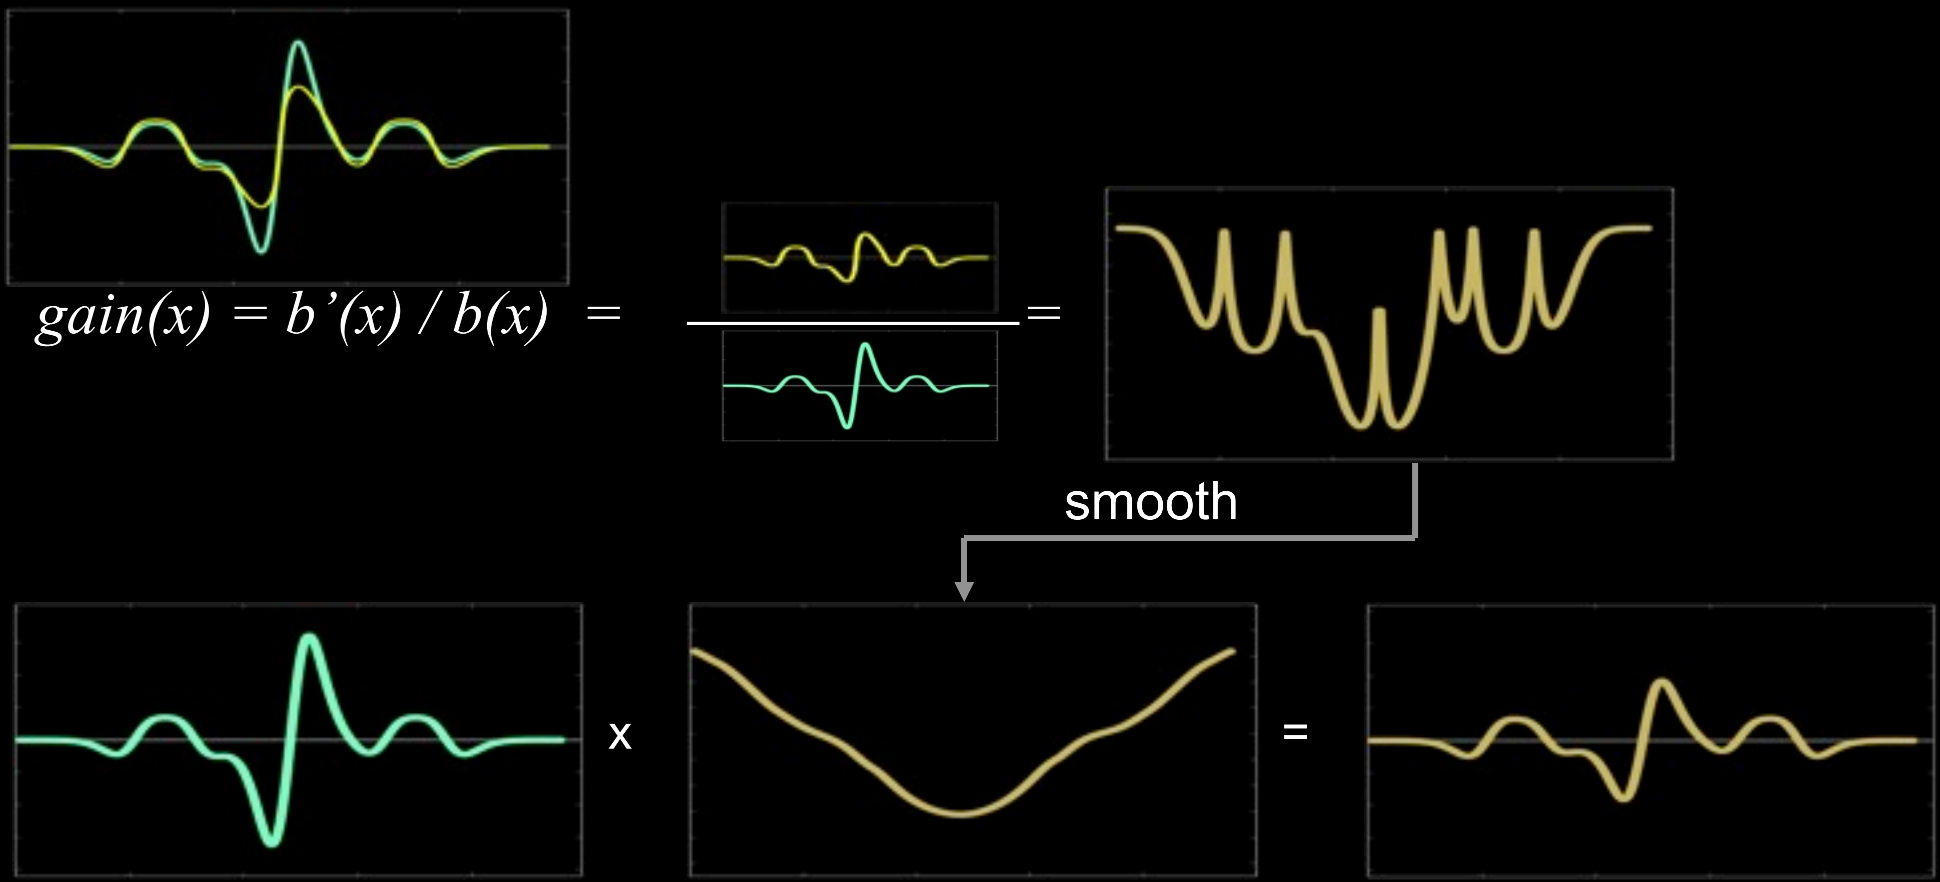

Smooth Gain Control Reduces Distortion

Point nonlinearity

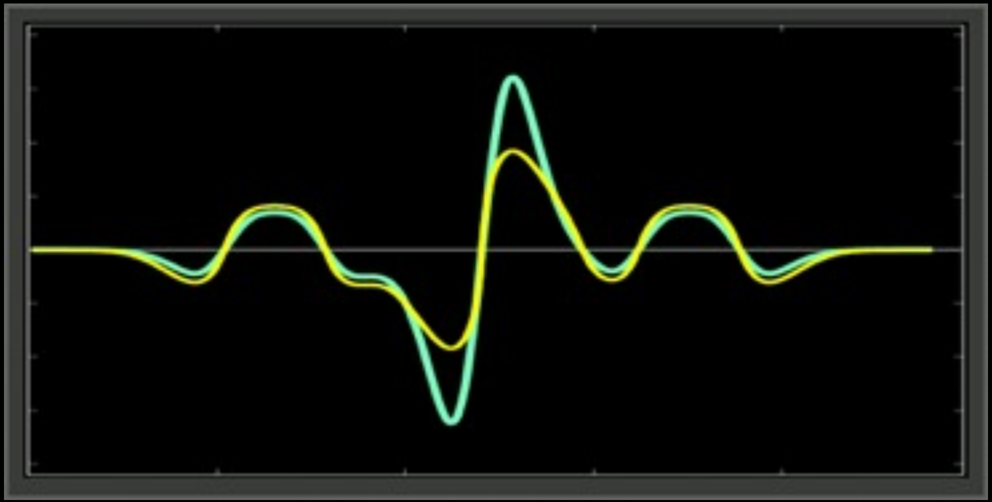

Distorted

Smooth gain control

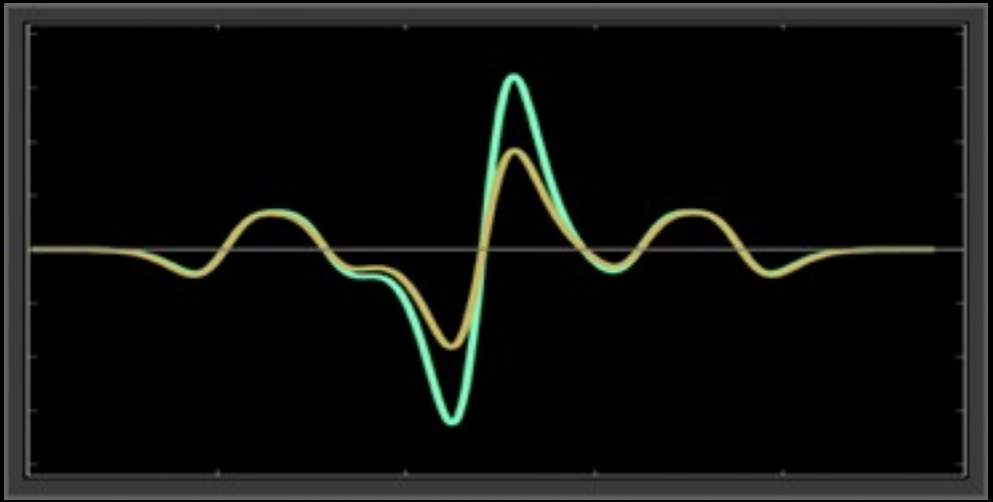

Distortion reduced

Smooth Gain Control on Subbands

Smooth Gain Control on Subbands

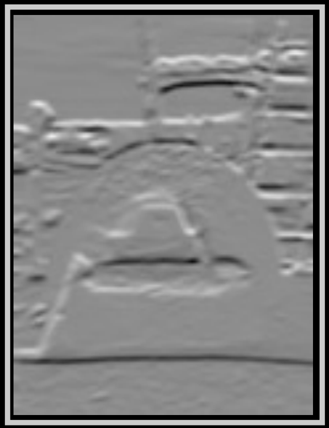

Smooth Gain Control on Subbands

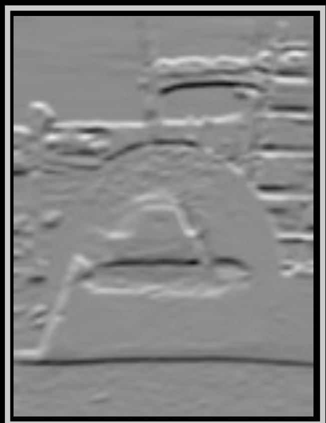

rectify (by taking abs)

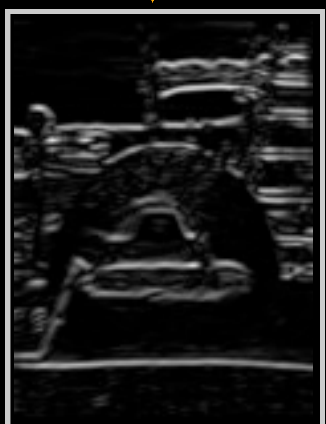

Smooth Gain Control on Subbands

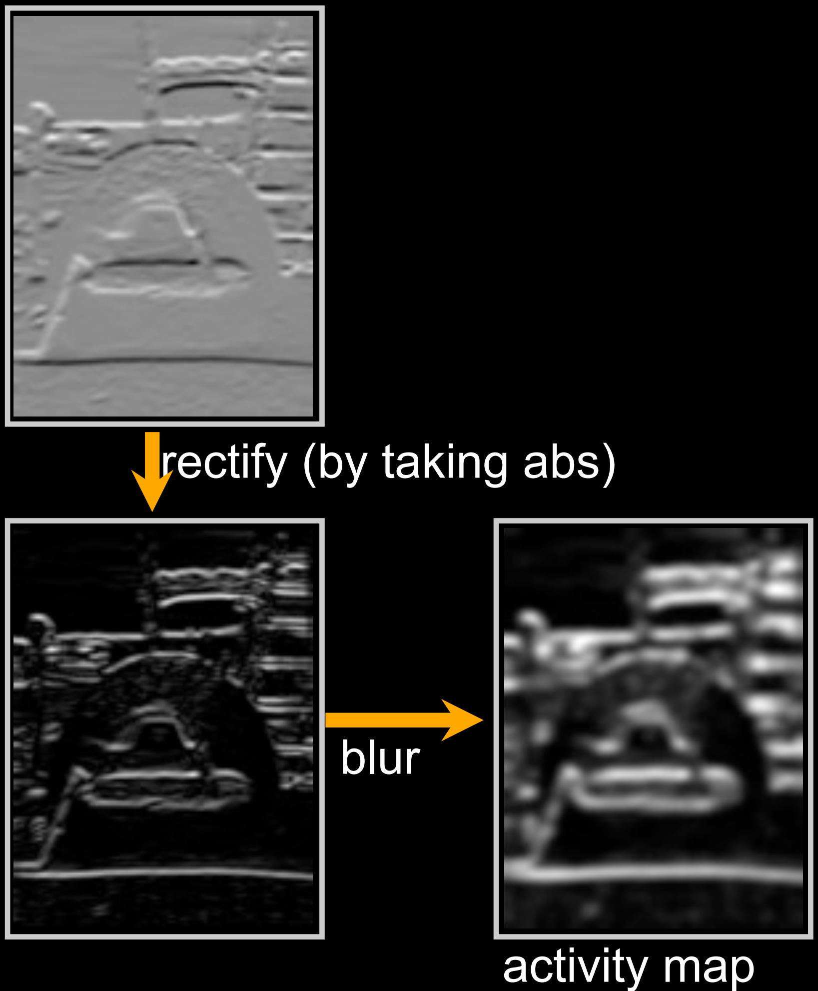

Smooth Gain Control on Subbands

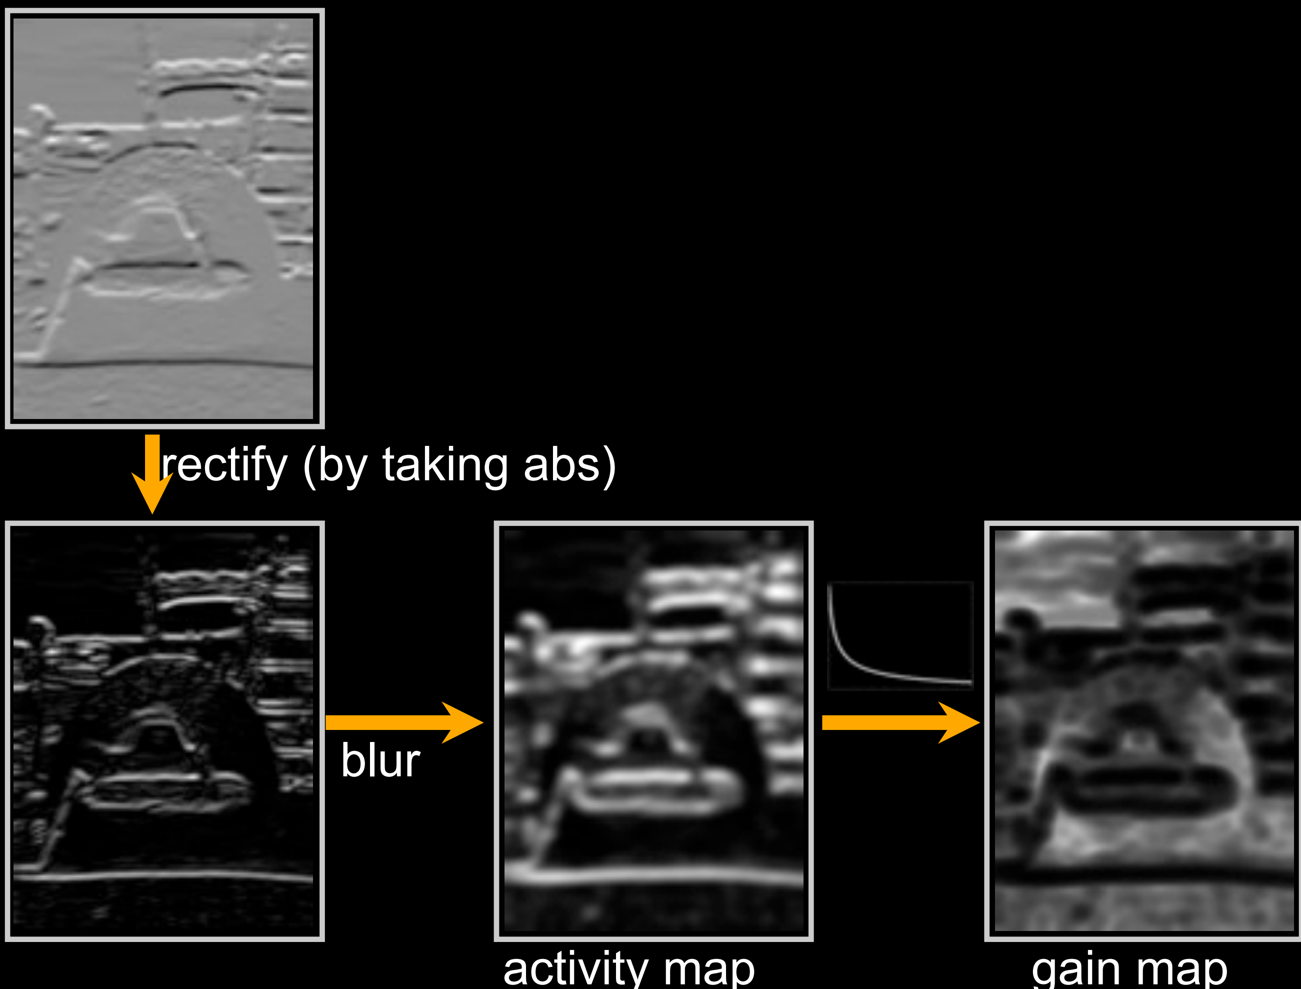

Smooth Gain Control on Subbands

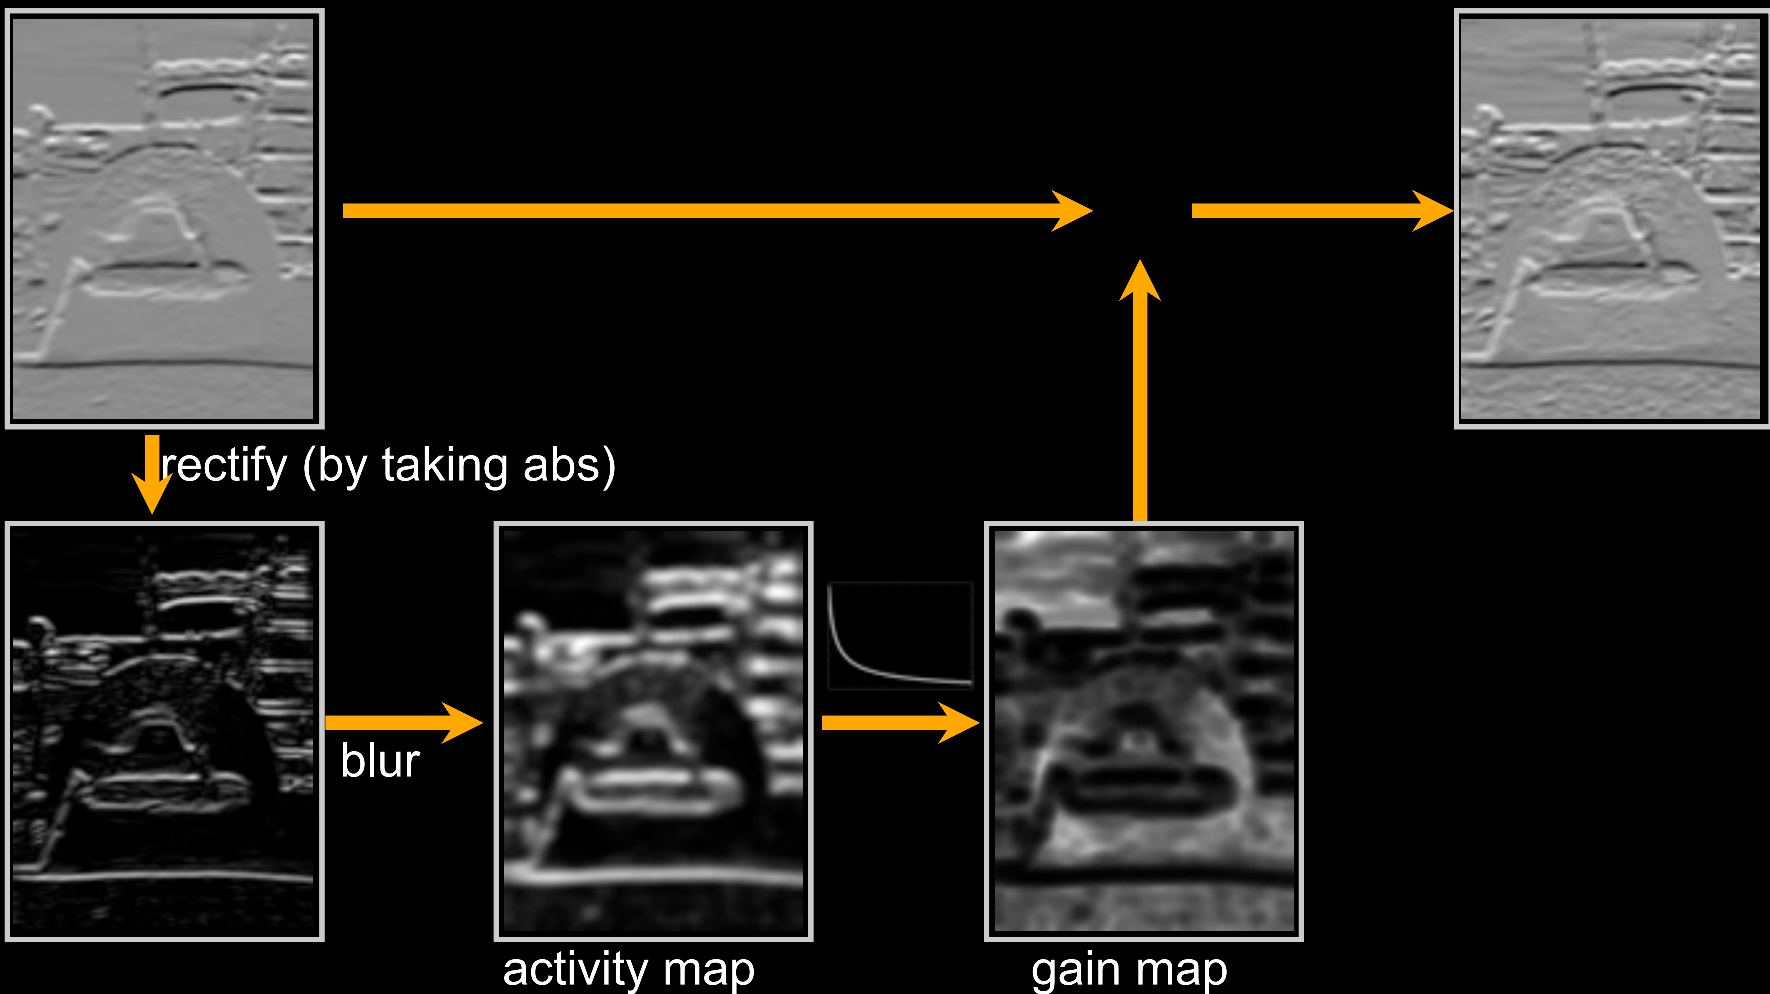

Ours

Durand & Dorsey 2002

Reduced contrast

- Tonemapping reduces dynamic range
  - usually also contrast is reduced
  - for example the skyline below becomes less visible

Countershading enhances contrast

- Make use of perceptual illusion
  - Craik – O'Brien – Cornsweet effect

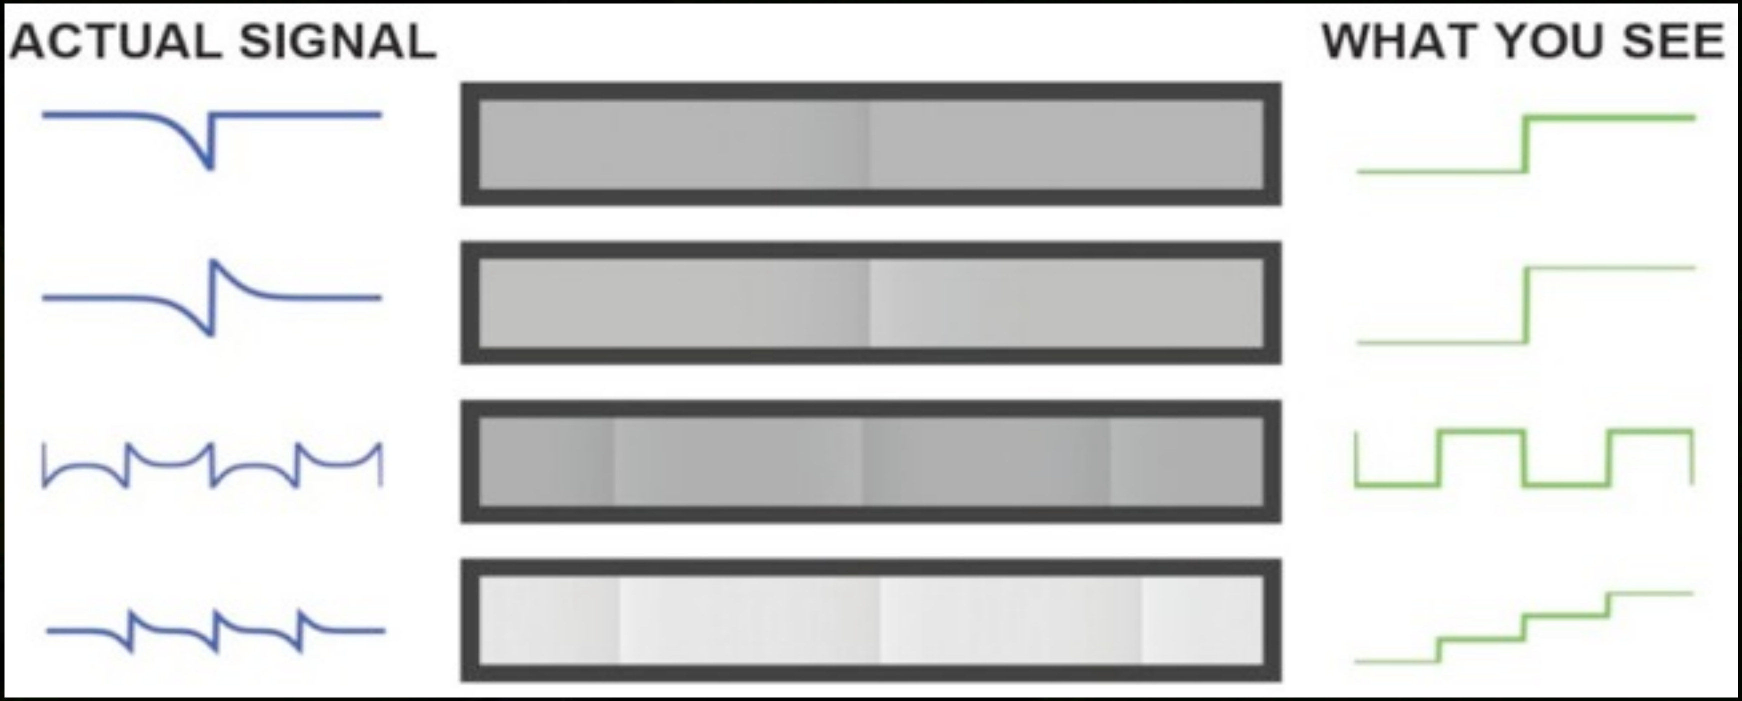

Unsharp masking

- Tonemapped HDR images lose contrast, add some back

Tonemapped LDR

Original HDR

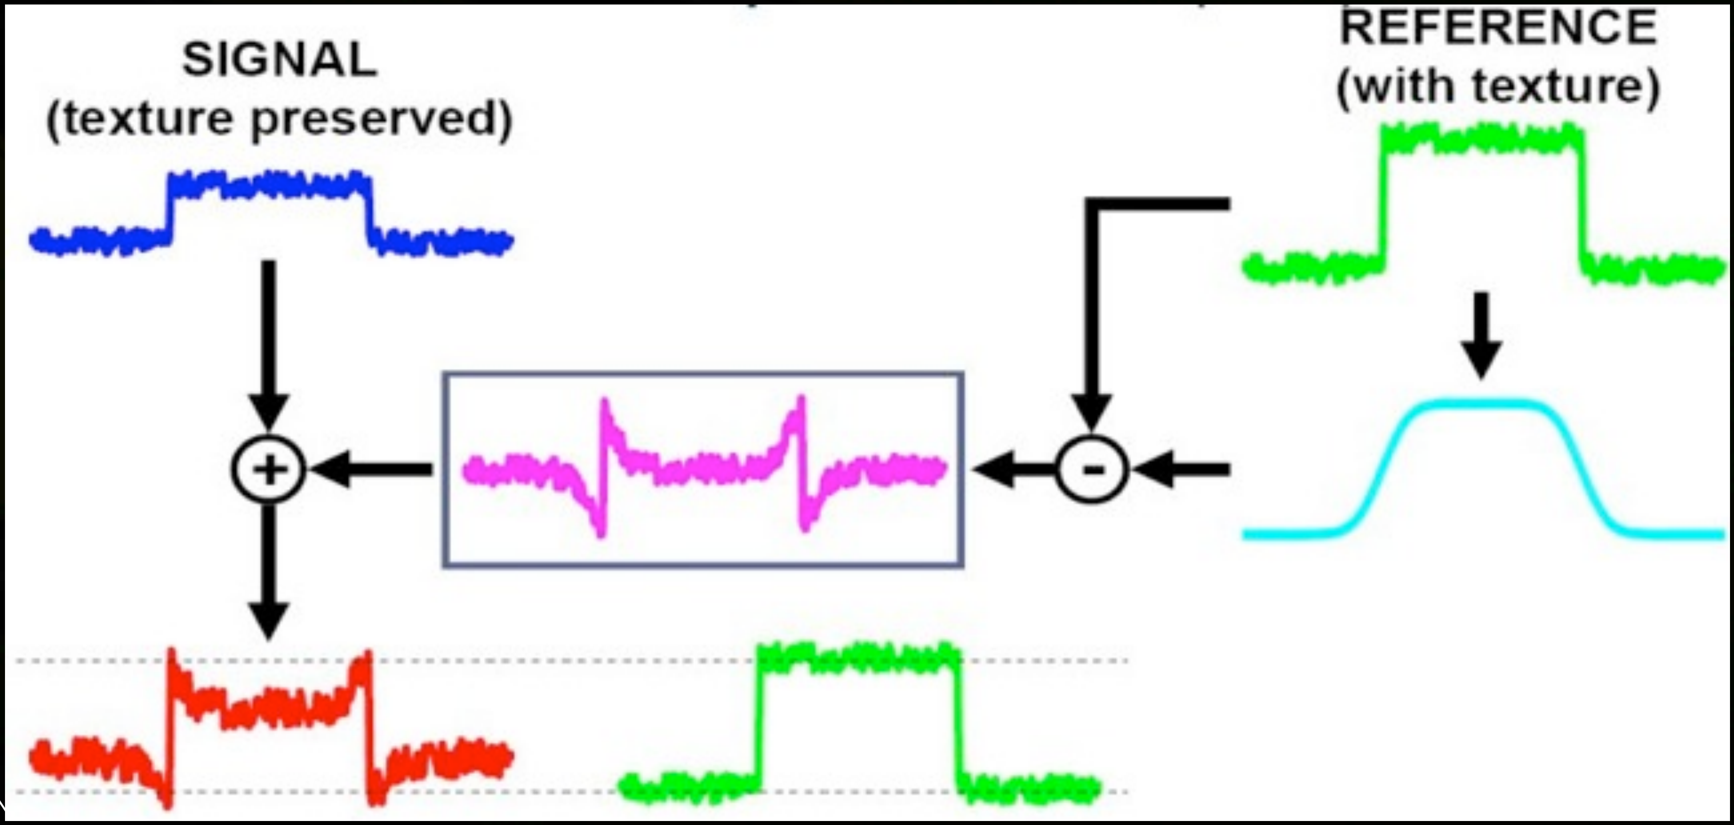

Monday, March 5, 12

Unsharp masking

- Too much contrast changes the visual appearance
- Global scale for contrast enhancements
  - features have different scales

Adaptive Countershading

- [Krawczyk et al. 07]
- Multiresolution local contrast metric adapts contrast to features
- Visual detection model avoids halos

Unsharp Masking, Countershading and Halos: Enhancements or Artifacts?

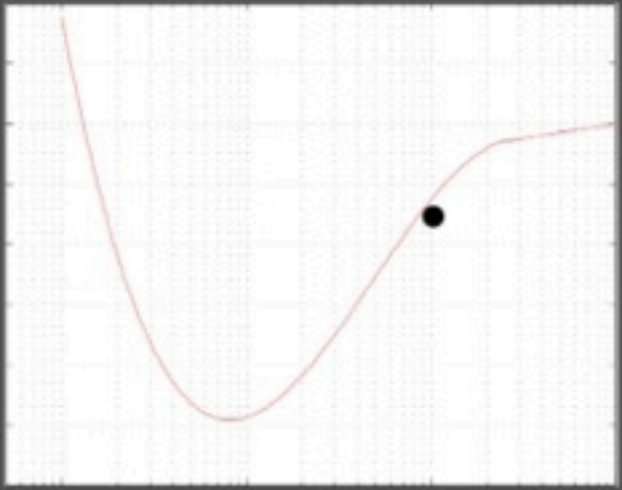

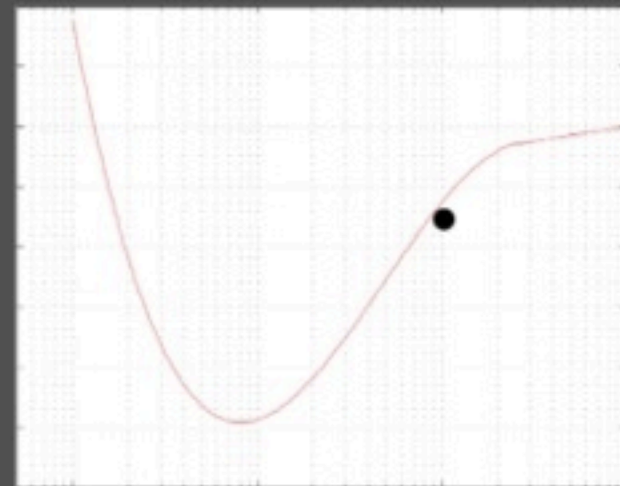


Adjust contrast
Adjust width

0:30 – 3:00; 3:35 – 4:10

0:30 – 3:00; 3:35 – 4:10

0:30 – 3:00; 3:35 – 4:10

Exposure Fusion: Simplified HDR

Mertens, Kautz, van Reeth PG 2007

- Choose the best pixel from one of the images
  - Use heuristics for a smooth selection, such as
    - Exposure
    - Color saturation
    - Contrast

LDR images

Weight maps

The Laplacian pyramid

Gaussian Pyramid

The Laplacian pyramid

Gaussian Pyramid

The Laplacian pyramid

Gaussian Pyramid

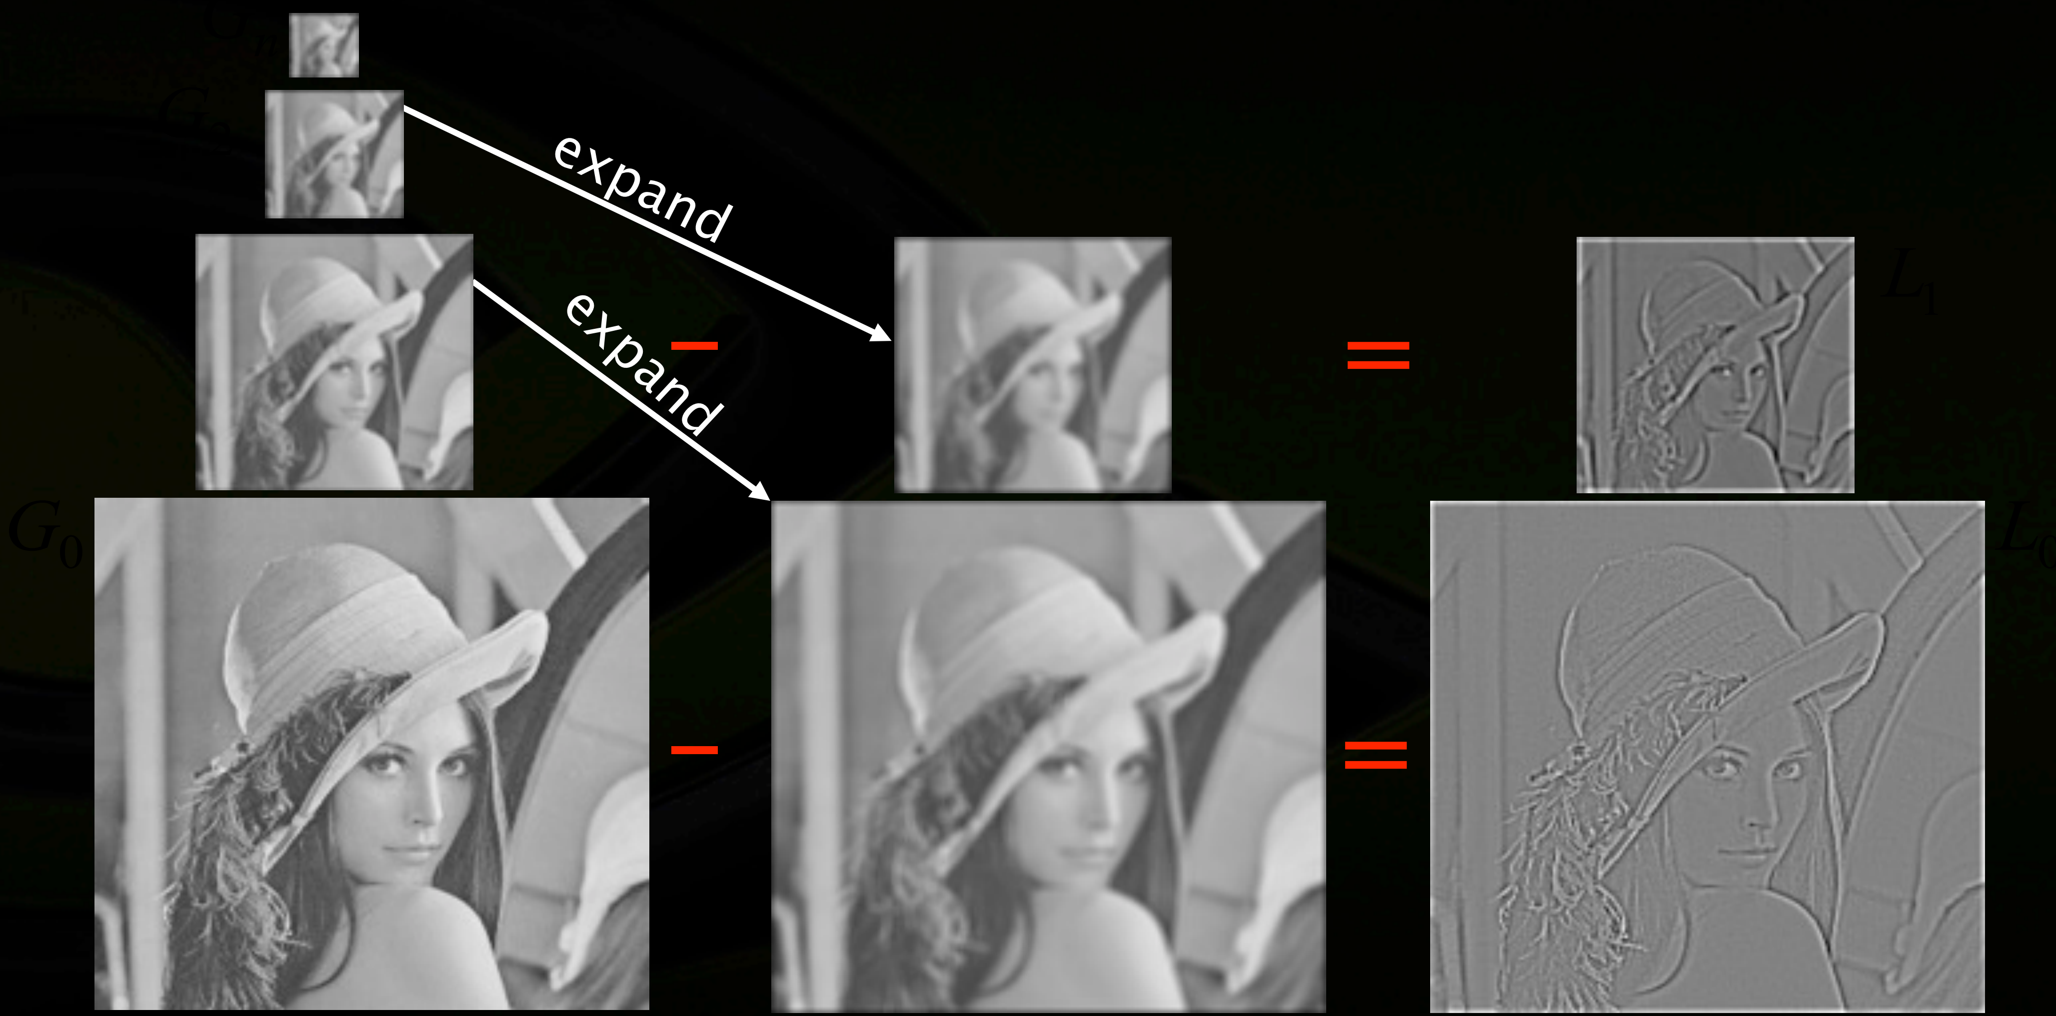

The Laplacian pyramid

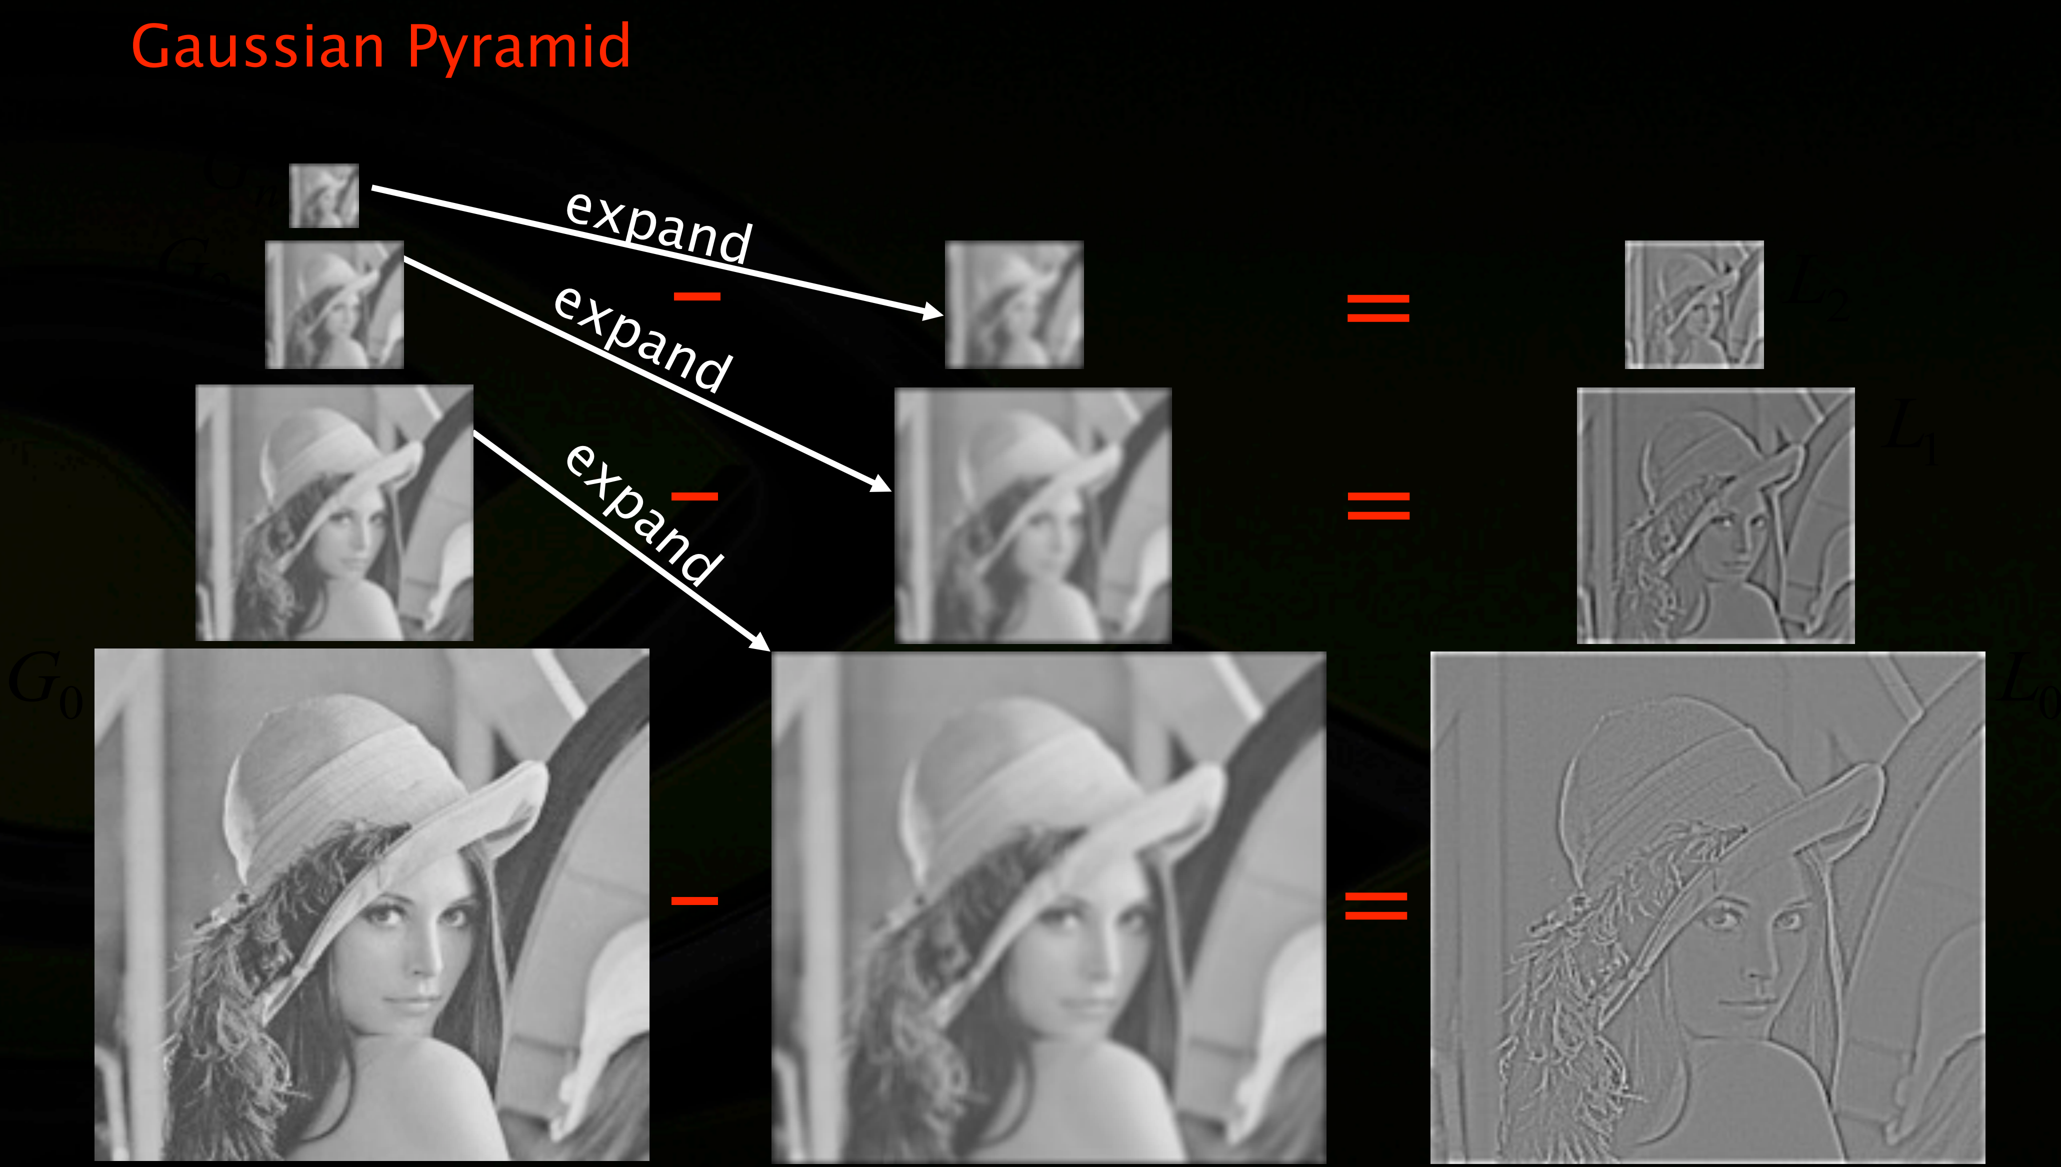

The Laplacian pyramid

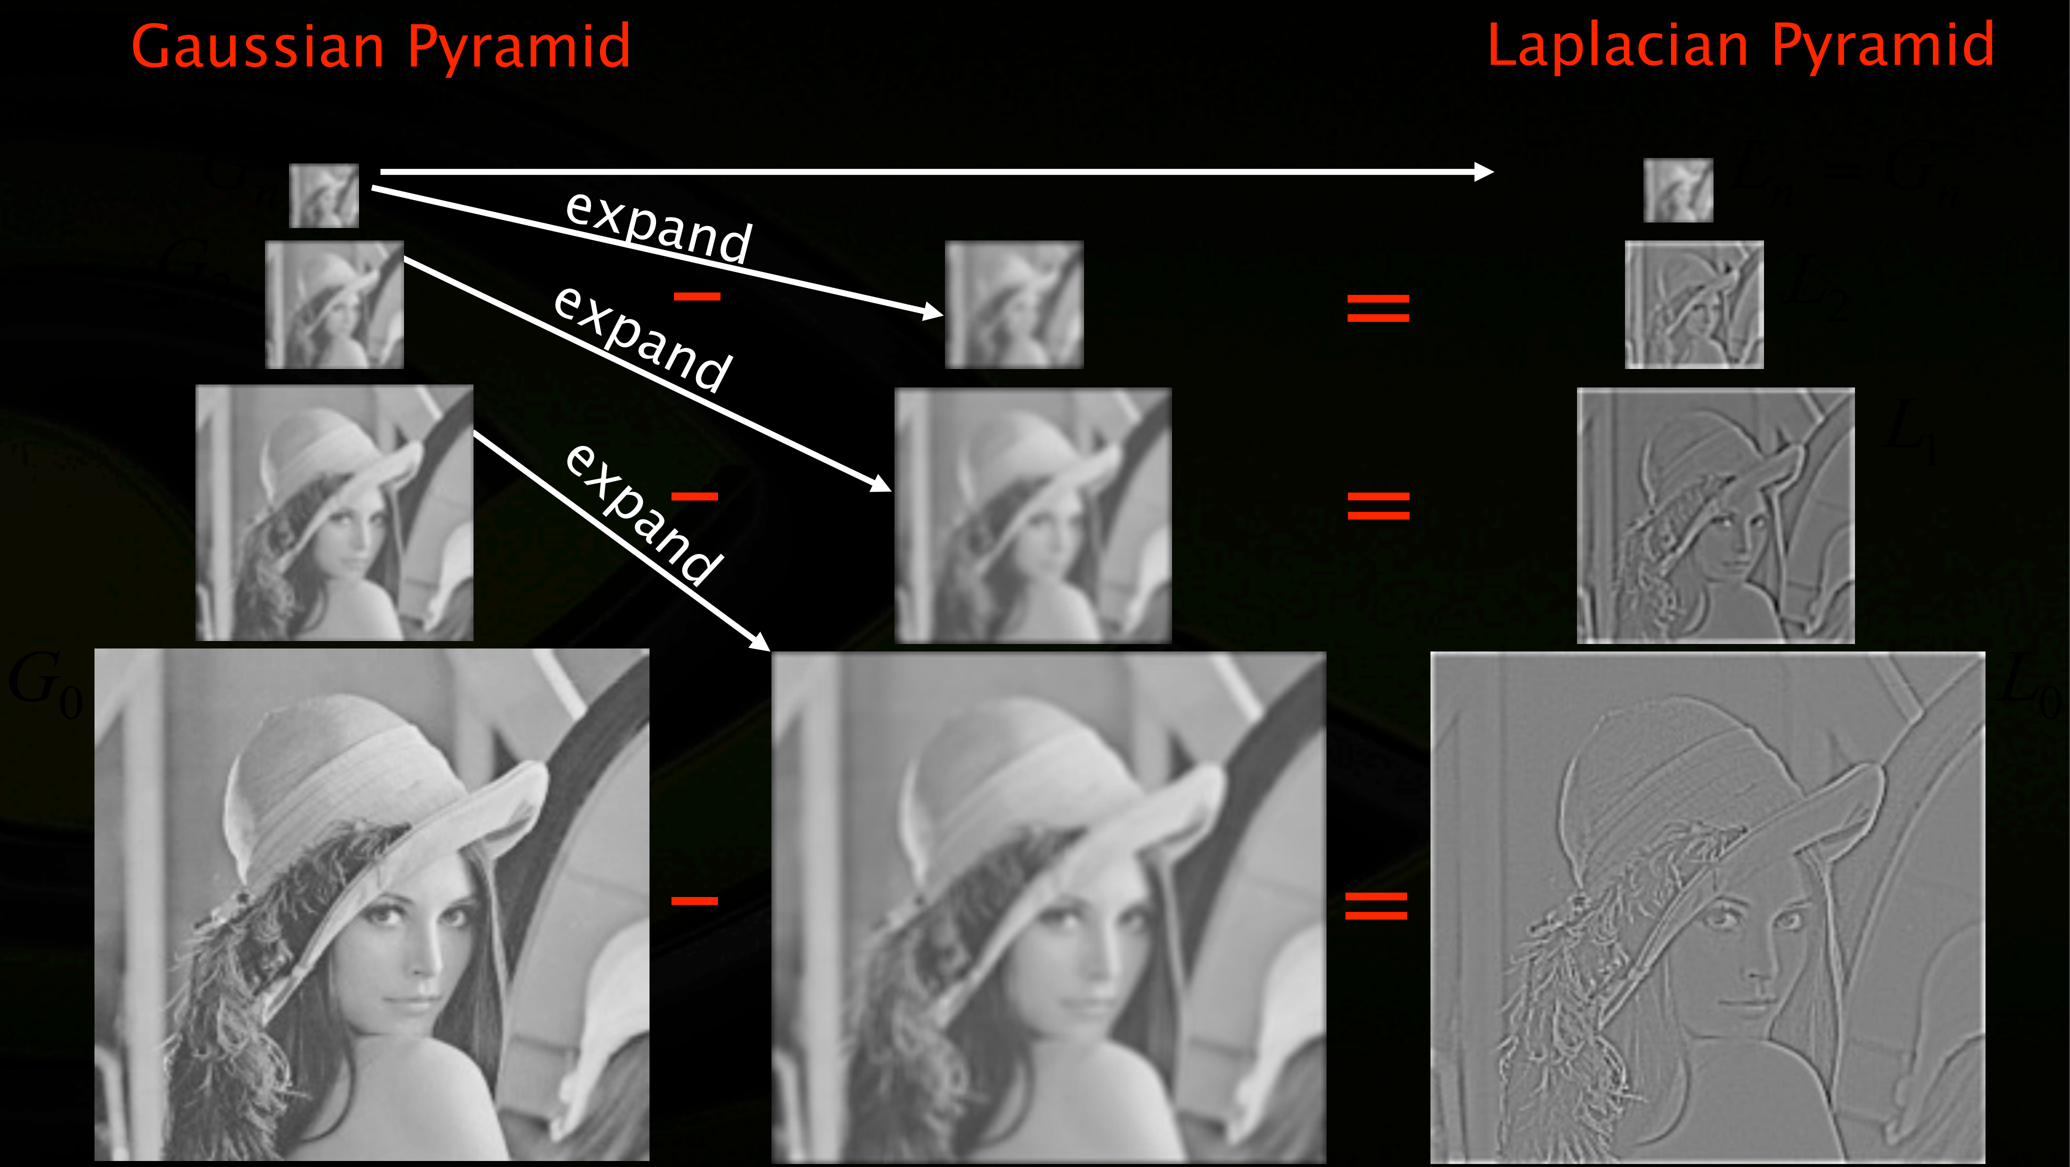

Multi-resolution fusion

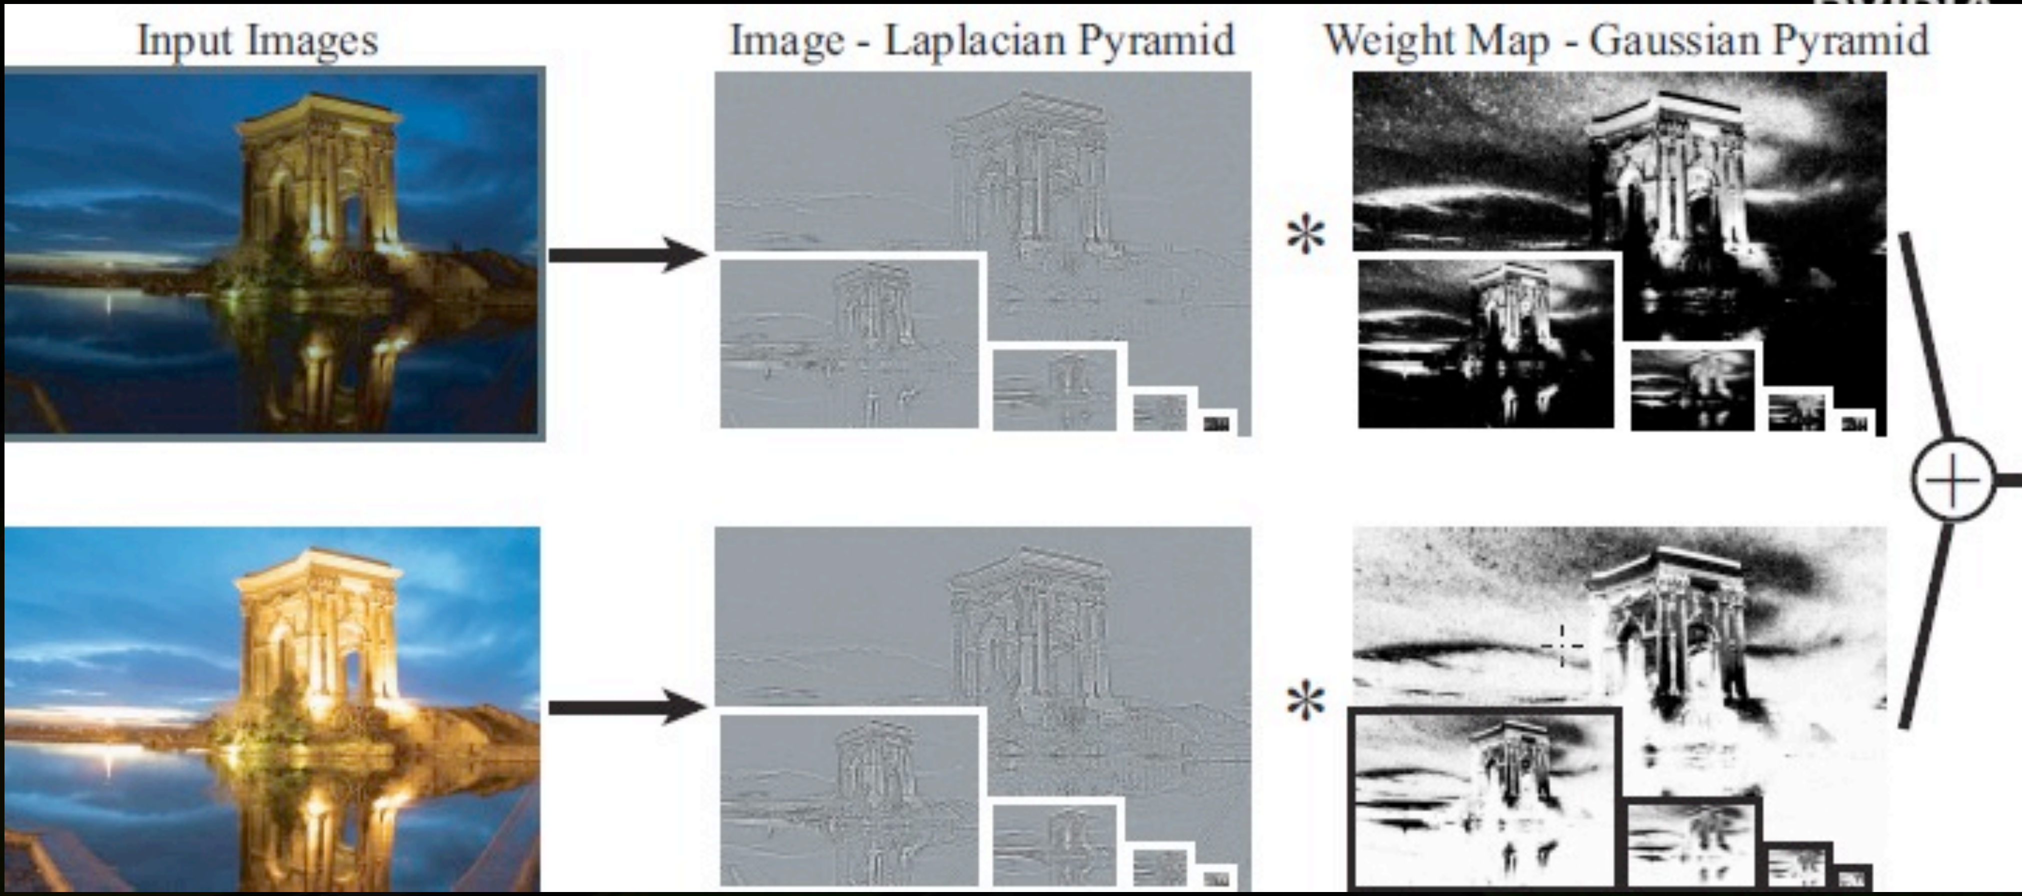

Multi-resolution fusion

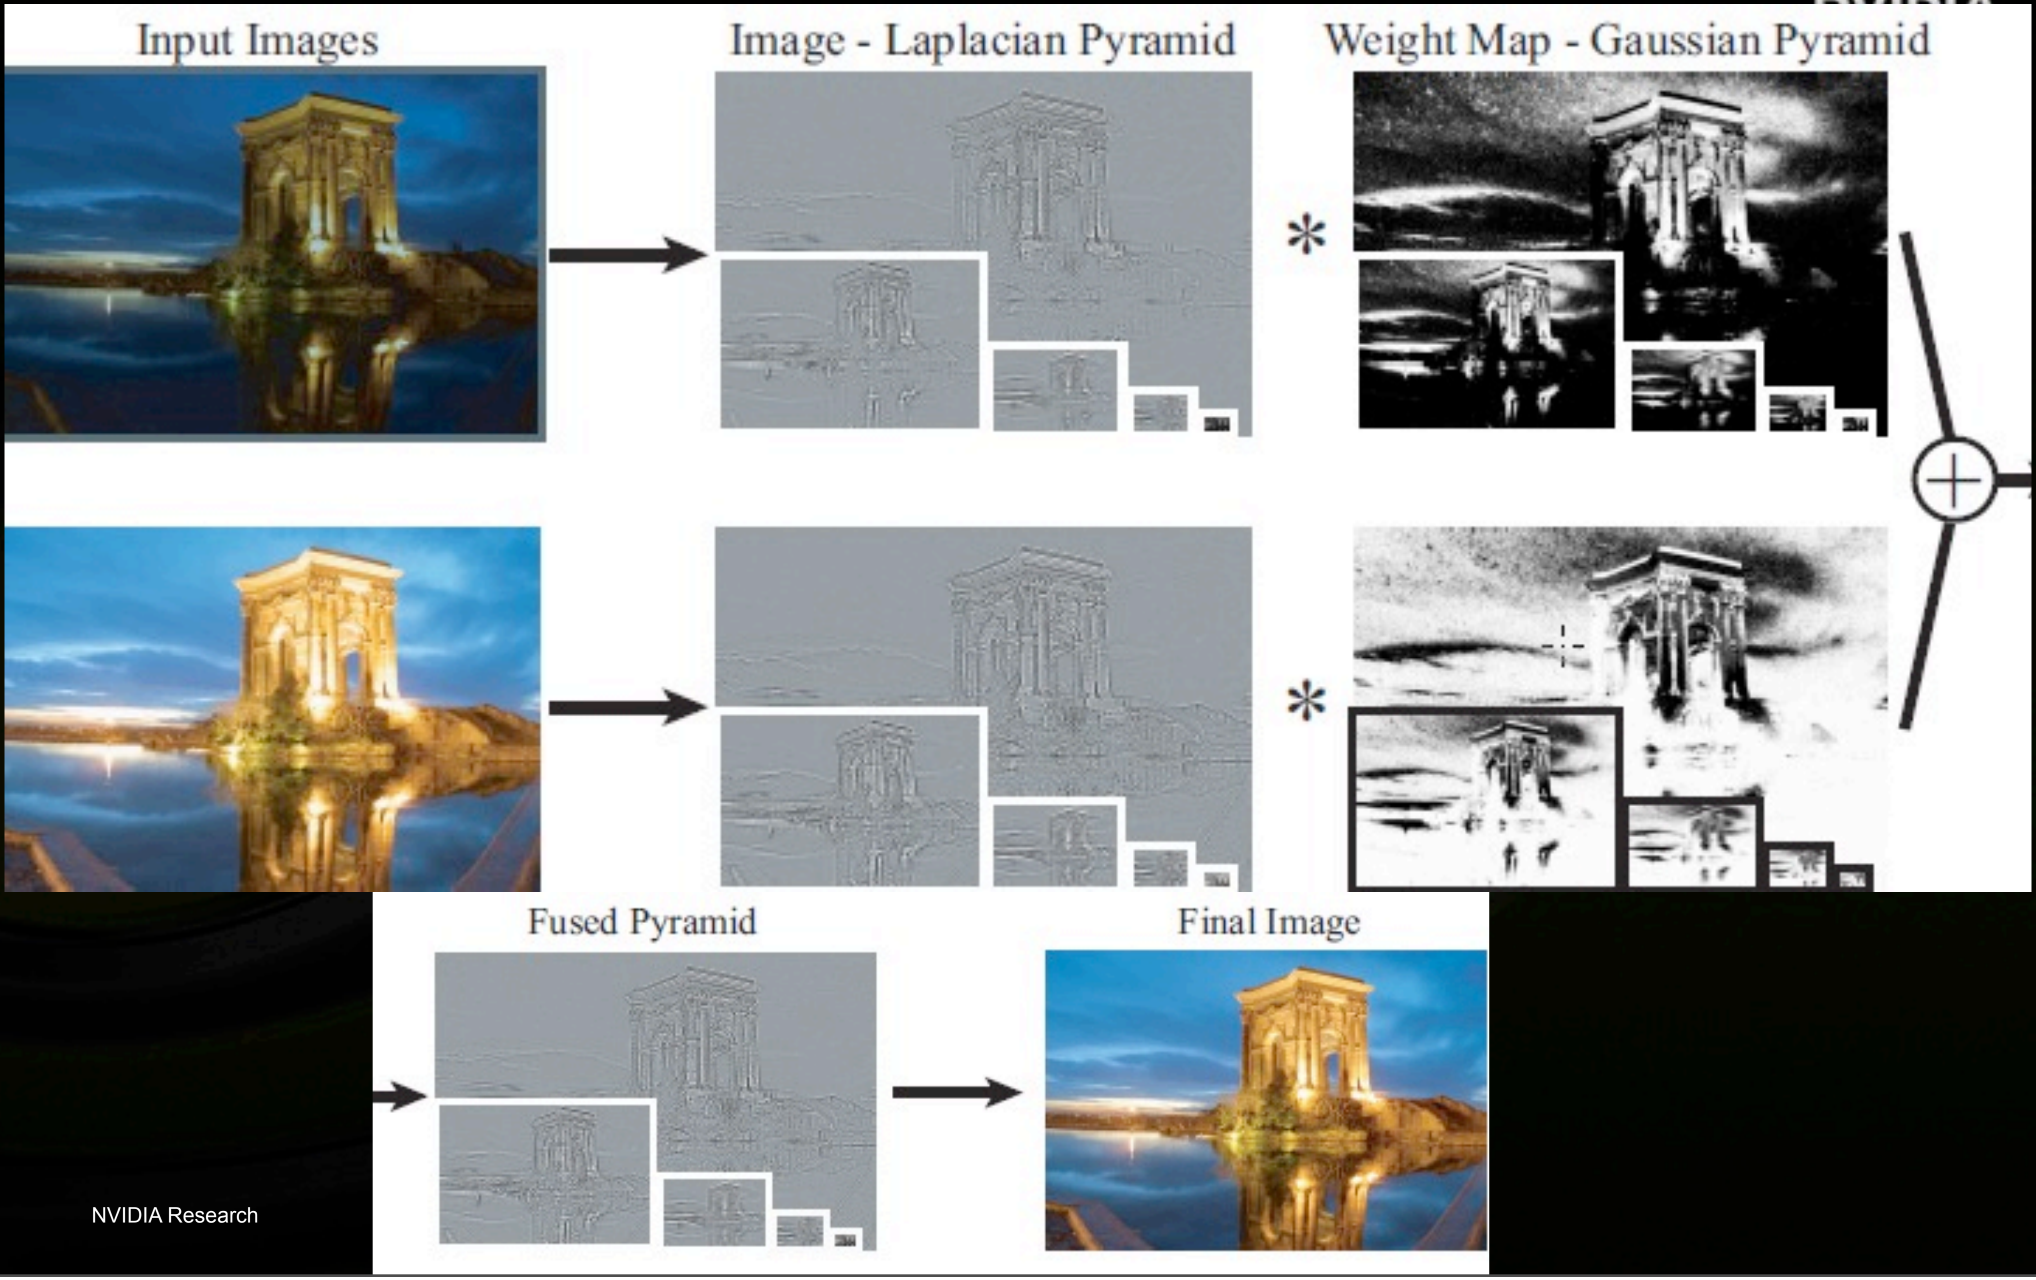

Metering

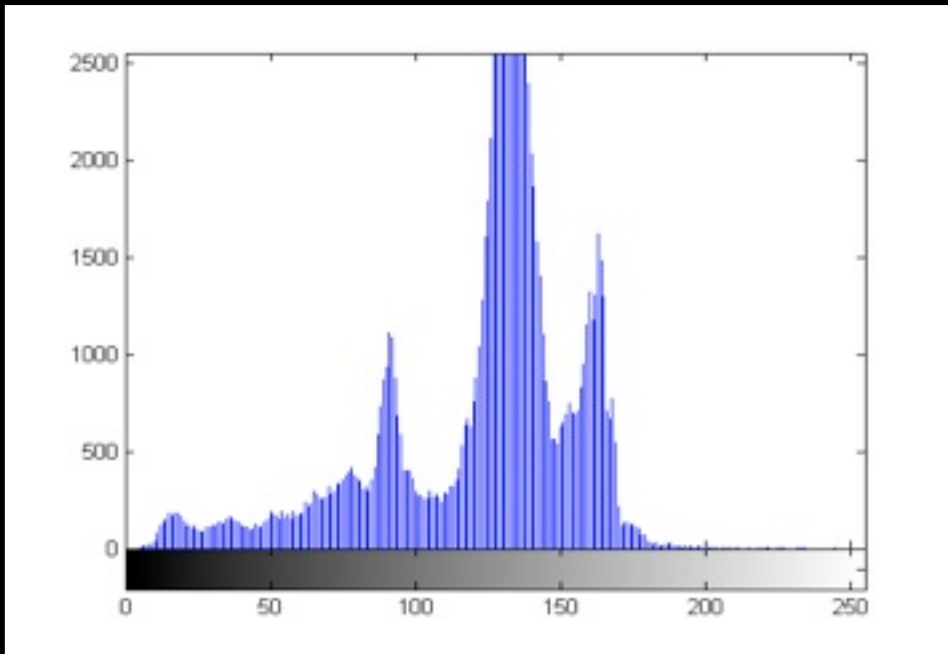

O. Gallo, M. Tico, N. Gelfand, K. Pulli
Metering for Exposure Stacks
Eurographics, Cagliari, Italy, May 13 - 18, 2012

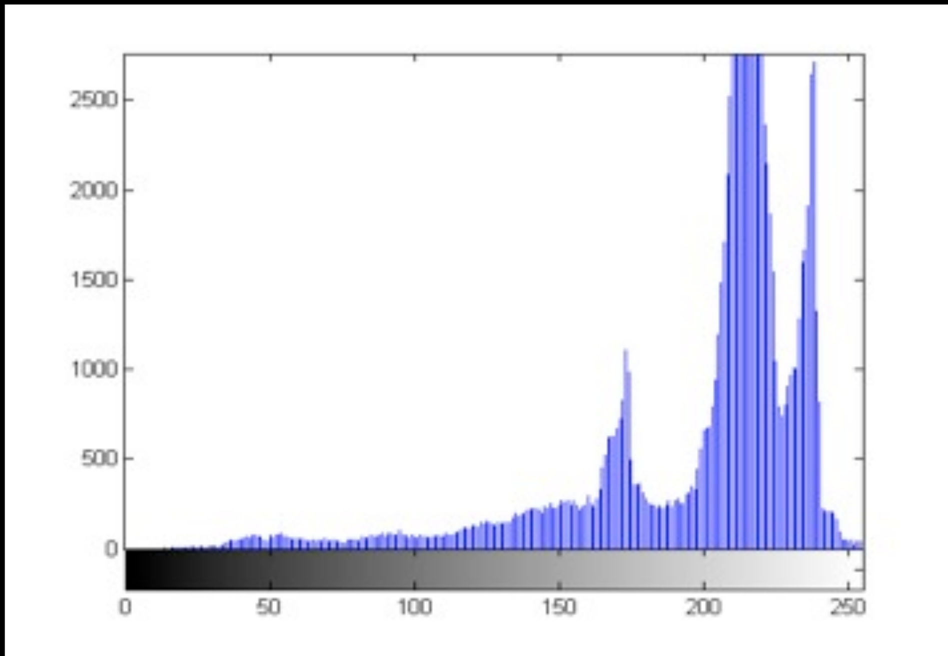
2500
2000
1500
1000
500
0
0
50
100
150
200
250


Baskin
Engineering
NOKIA

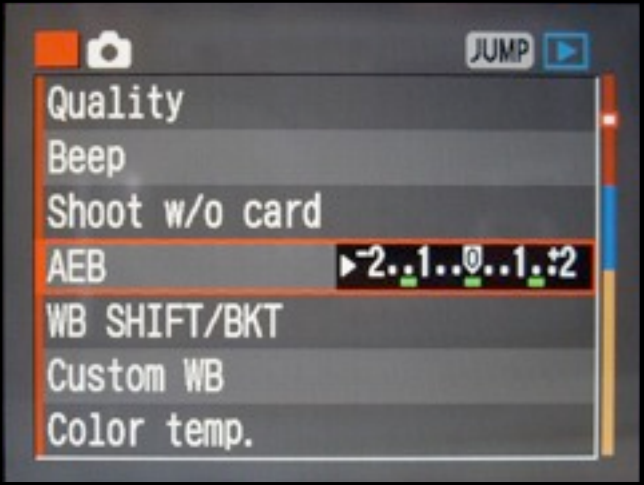
JUMP
Quality
Beep
Shoot w/o card
AEB
-2..1..0..1..+2
WB SHIFT/BKT
Custom WB
Color temp.

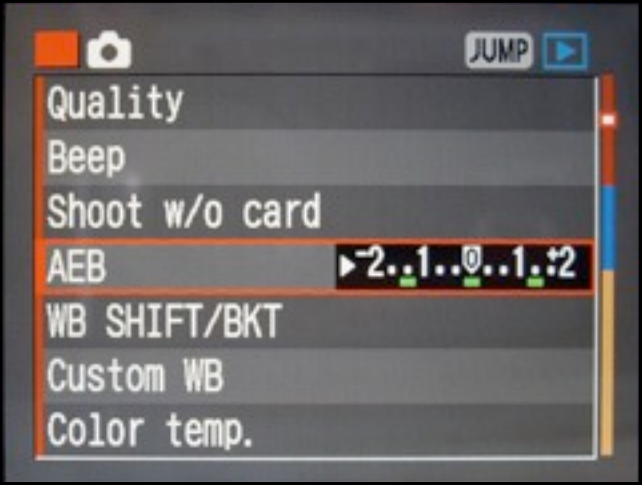

- “Minimal-Bracketing Sets for High-Dynamic-Range Image Capture”, *Barakat et al., Transaction on Image Processing* 2008
- “High Dynamic Range Imaging on Mobile Devices”, *Bilcu et al., ICECS 2008*
- *“Optimal Scheduling of Capture Times in a Multiple Capture Imaging System”, Chen and El Gamal, SPIE 2002*
- “Optimal HDR Reconstruction with Linear Digital Cameras”, Granados et al., CVPR 2010

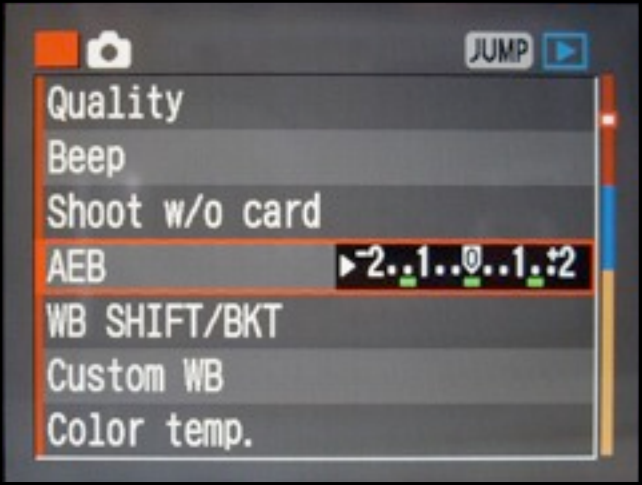

- “Minimal-Bracketing Sets for High-Dynamic-Range Image Capture”, *Barakat et al., Transaction on Image Processing* 2008
- “High Dynamic Range Imaging on Mobile Devices”, *Bilcu et al., ICECS 2008*

- *“Optimal Scheduling of Capture Times in a Multiple Capture Imaging System”, Chen and El Gamal, SPIE 2002*
- “Optimal HDR Reconstruction with Linear Digital Cameras”, Granados et al., CVPR 2010

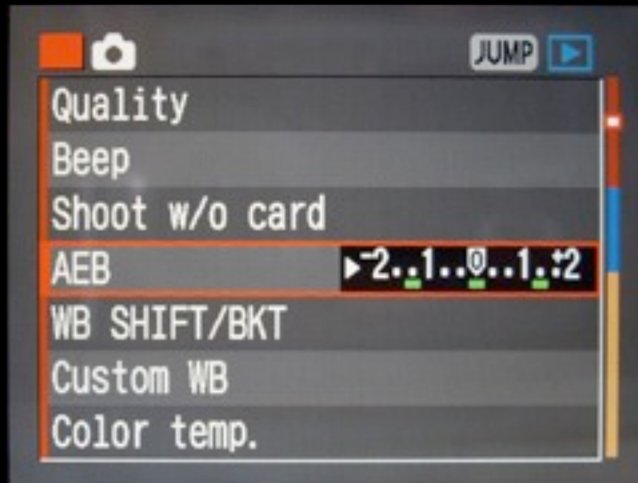

- “Minimal-Bracketing Sets for High-Dynamic-Range Image Capture”, *Barakat et al., Transaction on Image Processing* 2008
- “High Dynamic Range Imaging on Mobile Devices”, *Bilcu et al., ICECS 2008*
- ~~*“Optimal Scheduling of Capture Times in a Multiple Capture Imaging System”,*~~ *Chen and El Gamal, SPIE 2002*
- “Optimal HDR Reconstruction with Linear Digital Cameras”, Granados et al., CVPR 2010

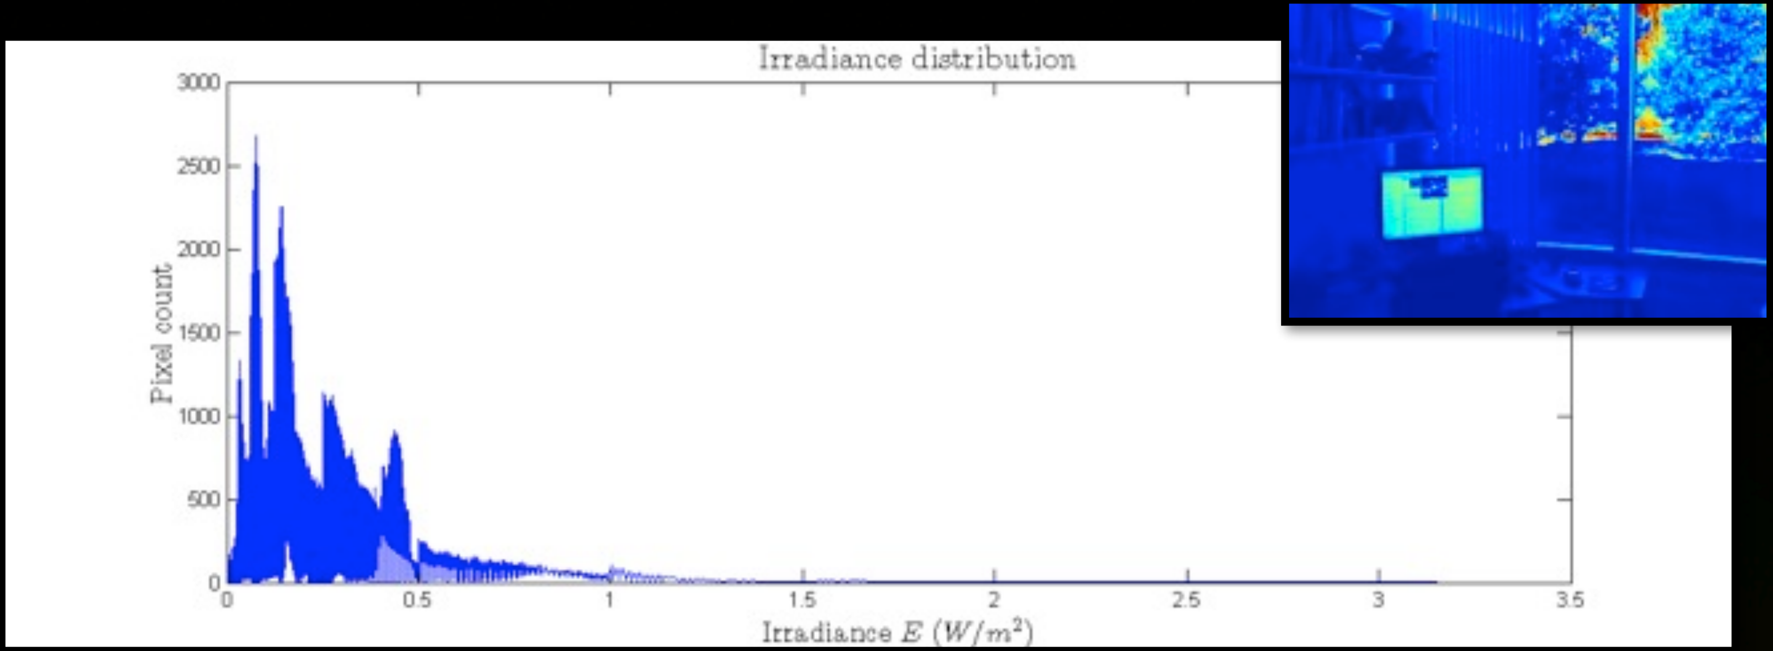

NVIDIA Research

Monday, March 5, 12

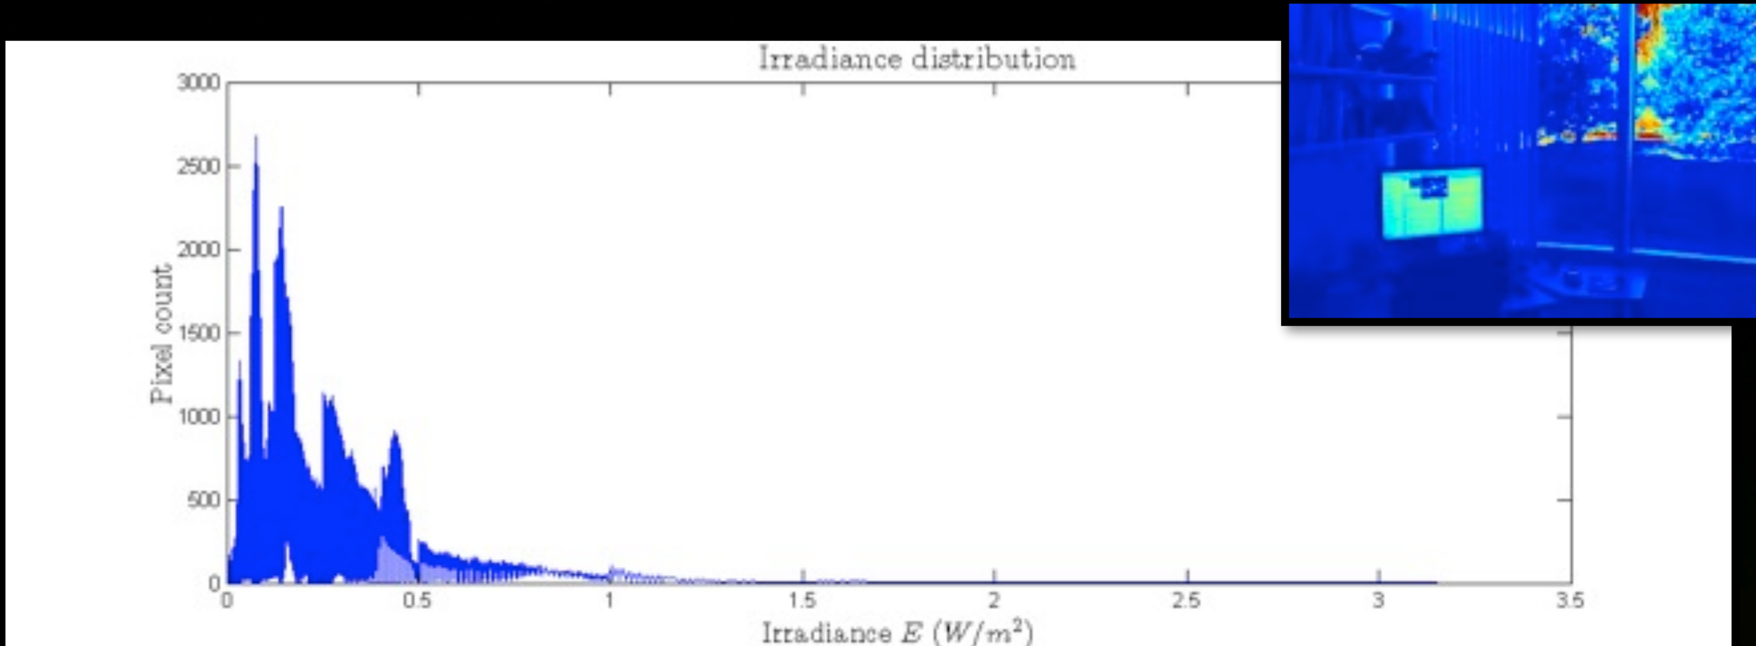
Irradiance distribution
Pixel count
3000
2500
2000
1500
1000
500
0
0
0.5
1
1.5
2
2.5
3
3.5
Irradiance E (W/m^2)

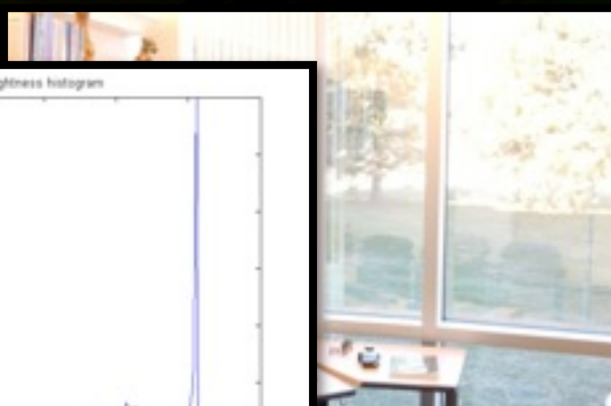

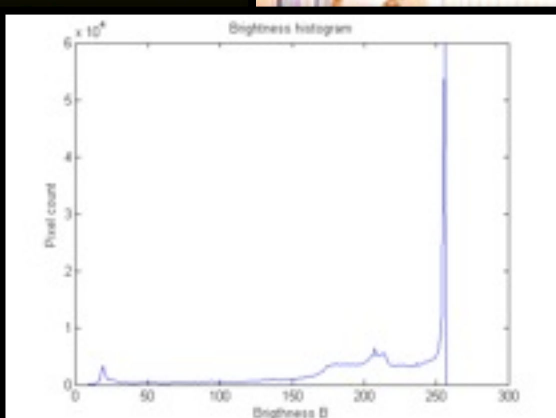
Brightness histogram
x 10^4
Pixel count
Brightness B
0
50
100
150
200
250
300

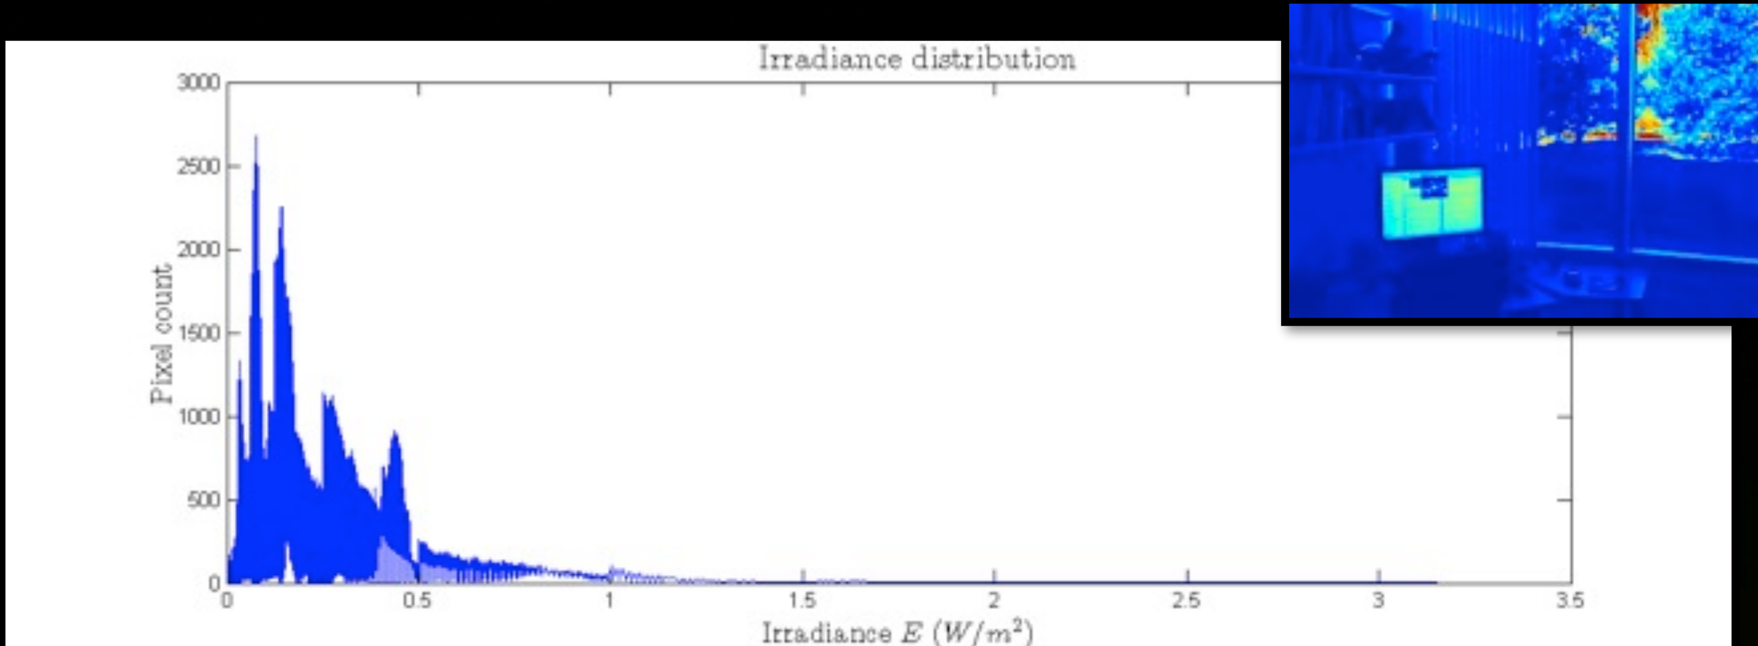
Irradiance distribution
Pixel count
Irradiance E (W/m²)
3000
2500
2000
1500
1000
500
0
0
0.5
1
1.5
2
2.5
3
3.5

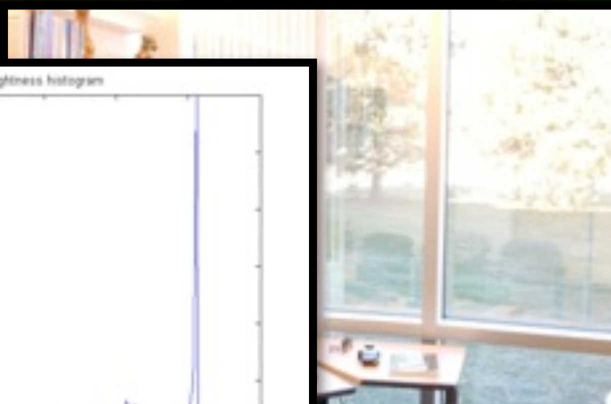

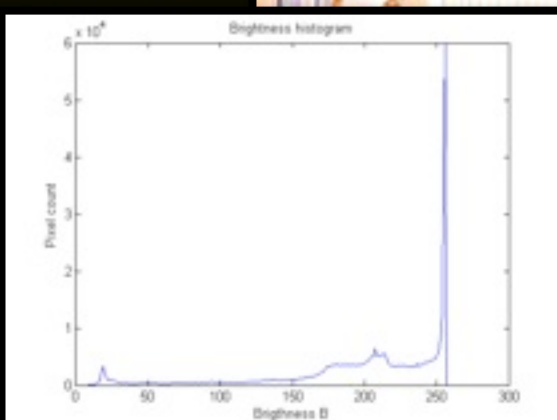
Brightness histogram
Pixel count
Brightness B

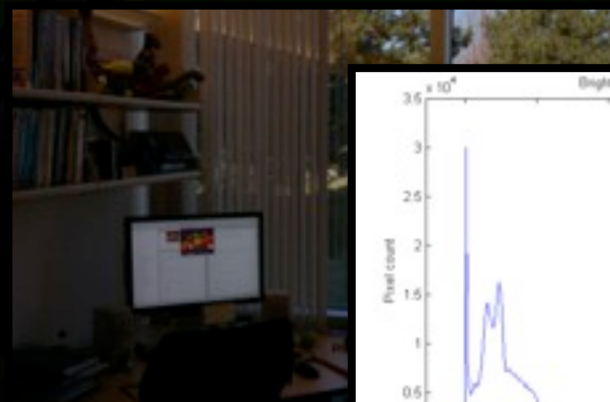

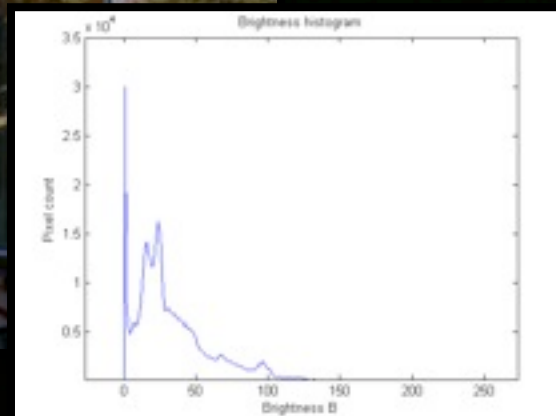
Brightness histogram
Pixel count
Brightness B

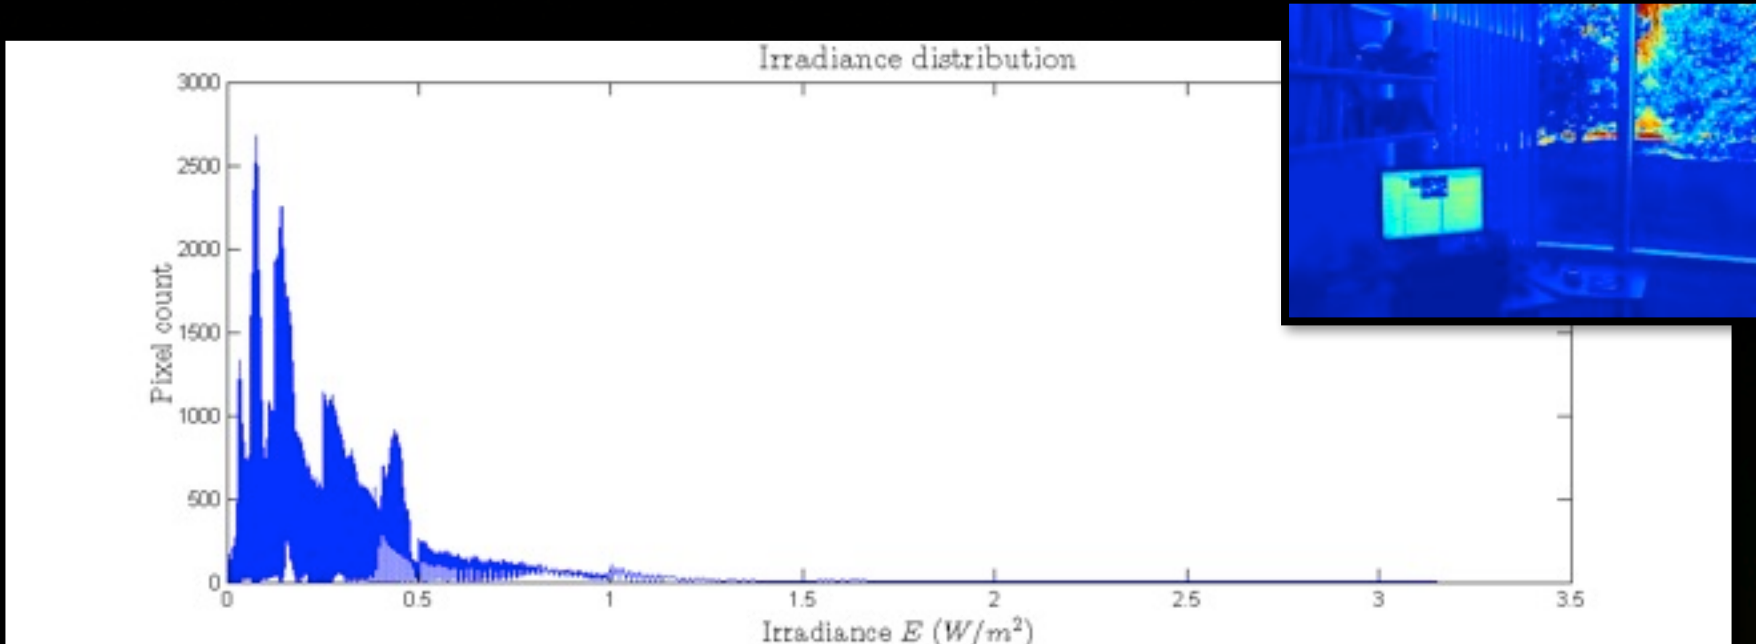
Irradiance distribution
Pixel count
3000
2500
2000
1500
1000
500
0
0
0.5
1
1.5
2
2.5
3
3.5
Irradiance E (W/m^2)

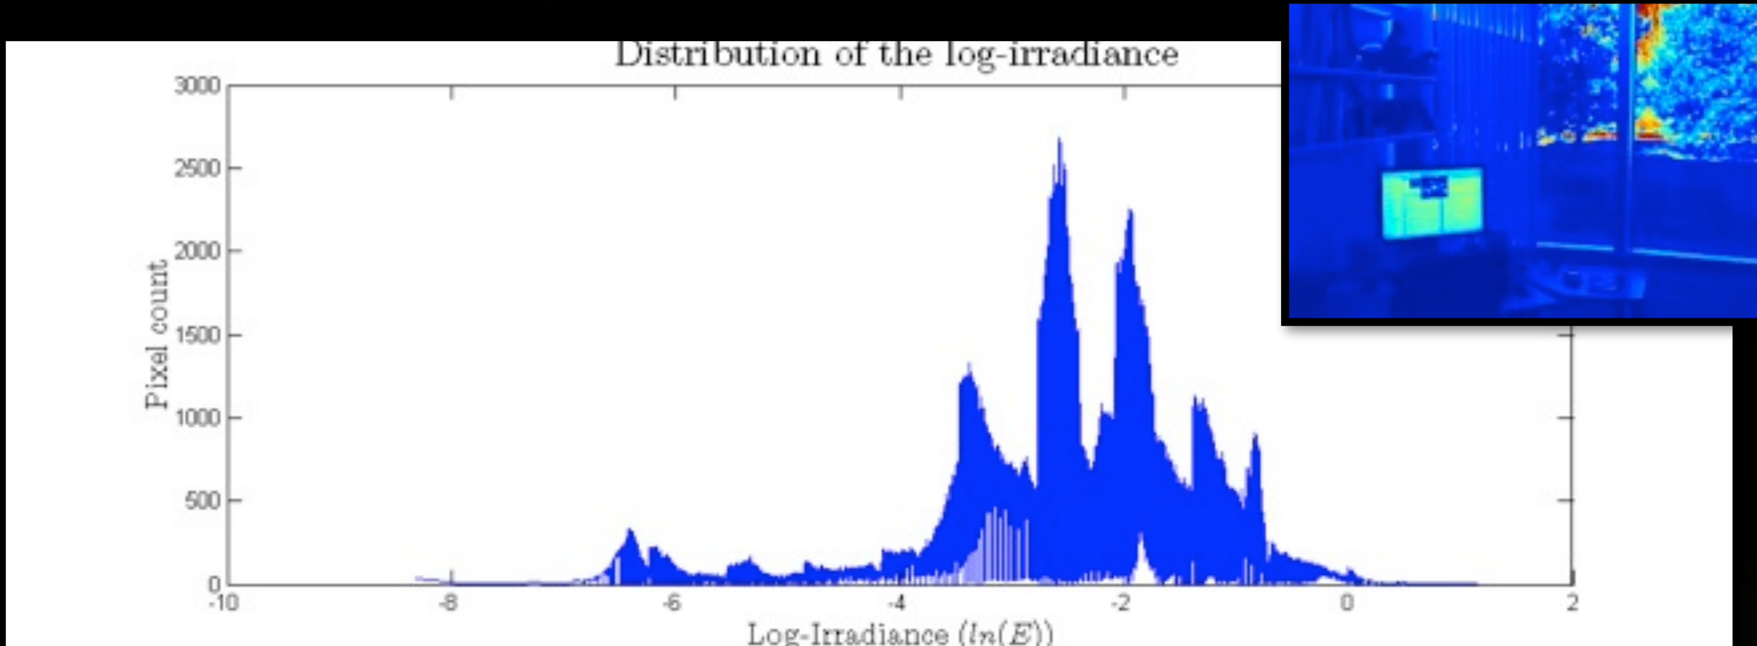

Distribution of the log-irradiance
Pixel count
3000
2500
2000
1500
1000
500
0
-10
-8
-6
-4
-2
0
2
Log-Irradiance ($ln(E)$)

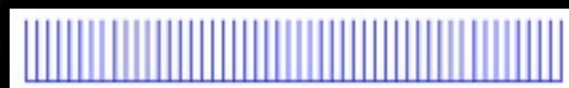

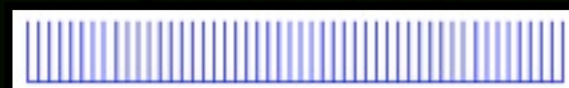


Pixel count
ln(E)

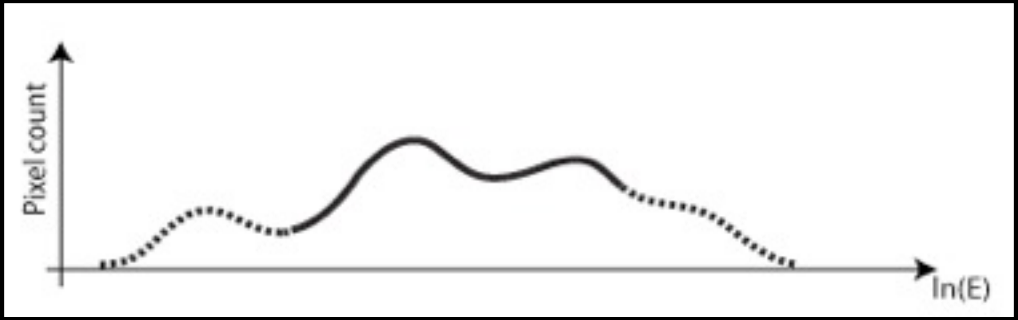

Δt_1

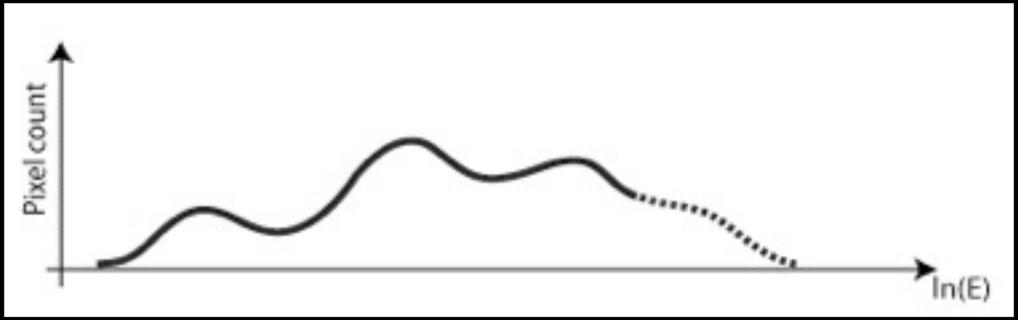

Δt_1

Δt_2

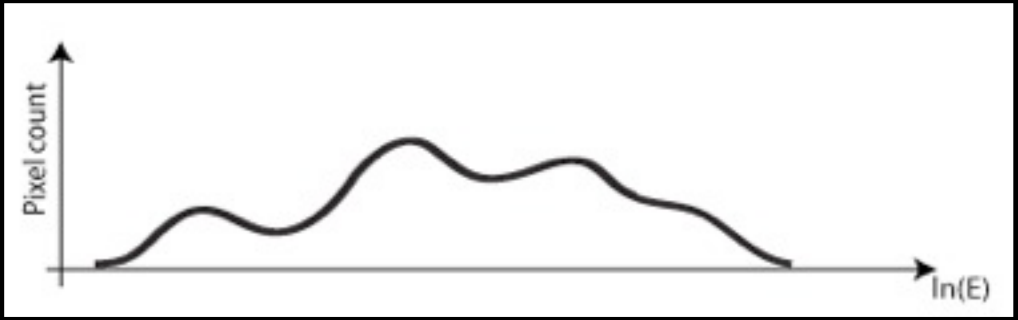
Pixel count
ln(E)

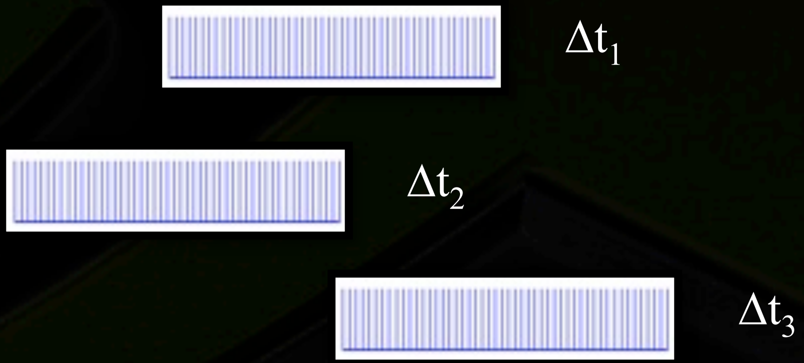
Δt_1
Δt_2
Δt_3

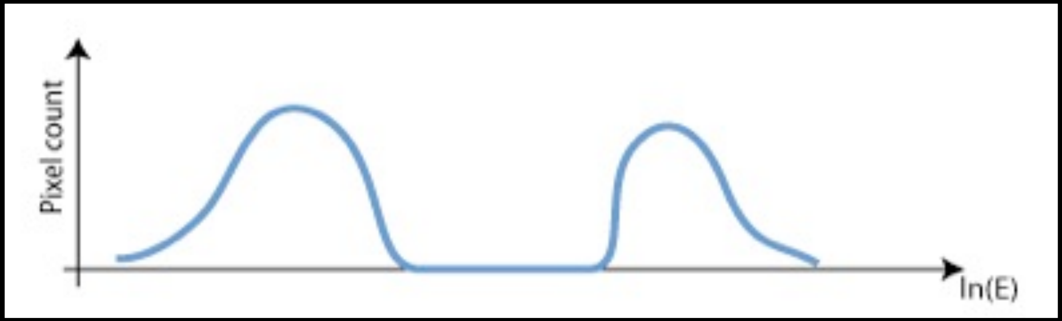
Pixel count
ln(E)

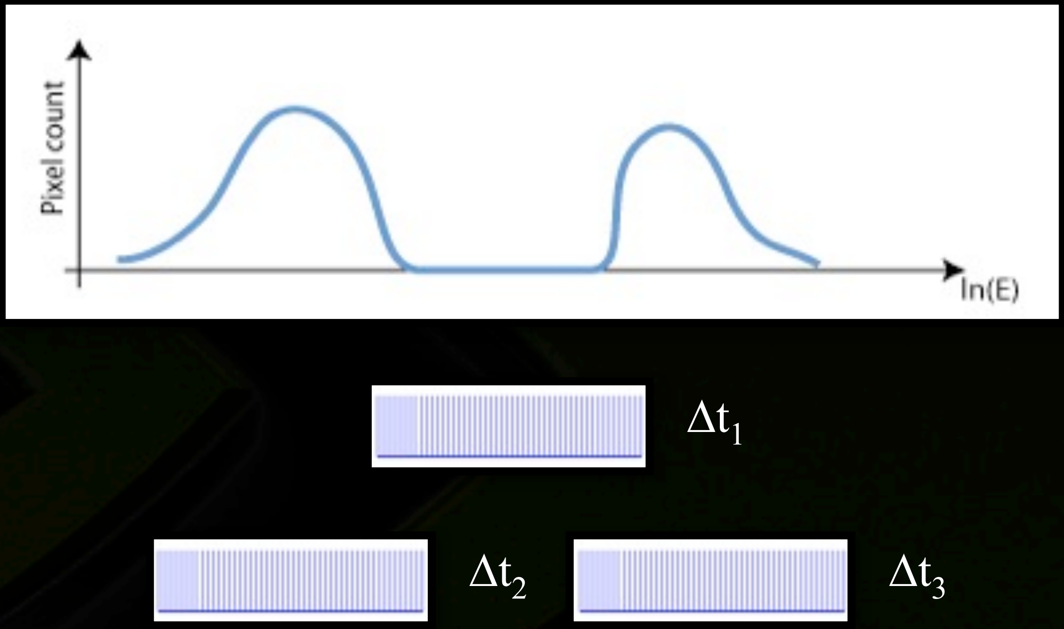
Pixel count
ln(E)
Δt_1
Δt_2
Δt_3

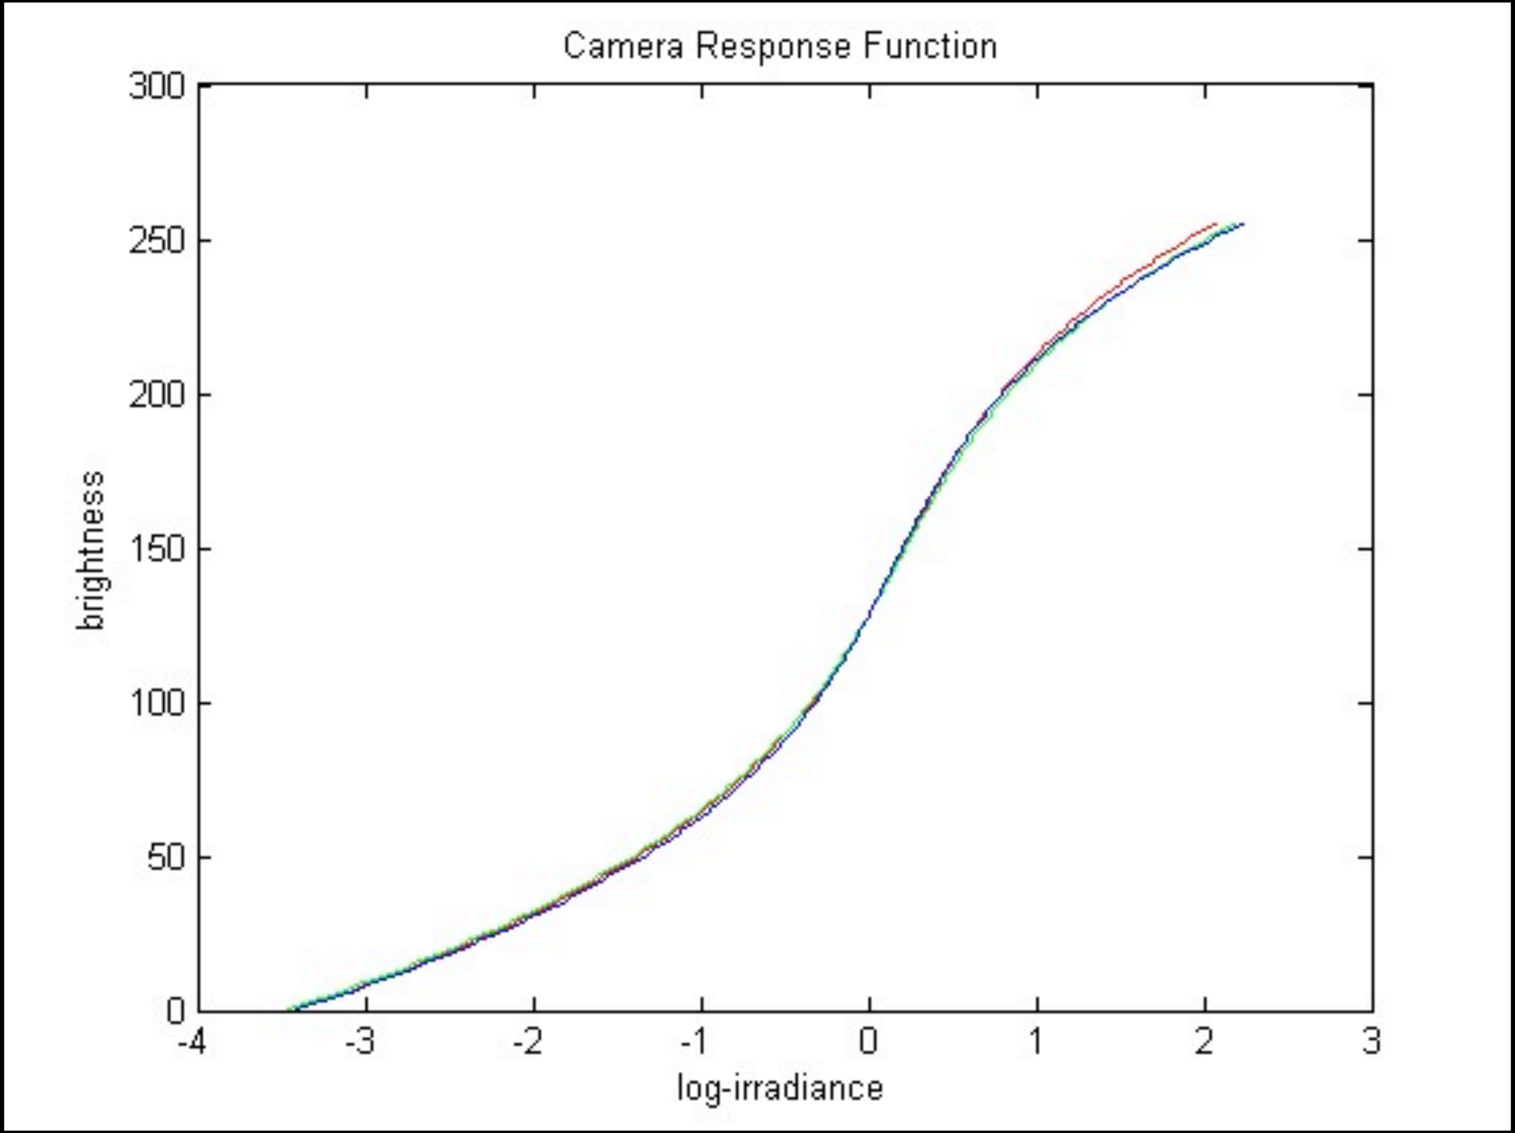
Camera Response Function
300
250
200
150
100
50
0
brightness
-4
-3
-2
-1
0
1
2
3
log-irradiance

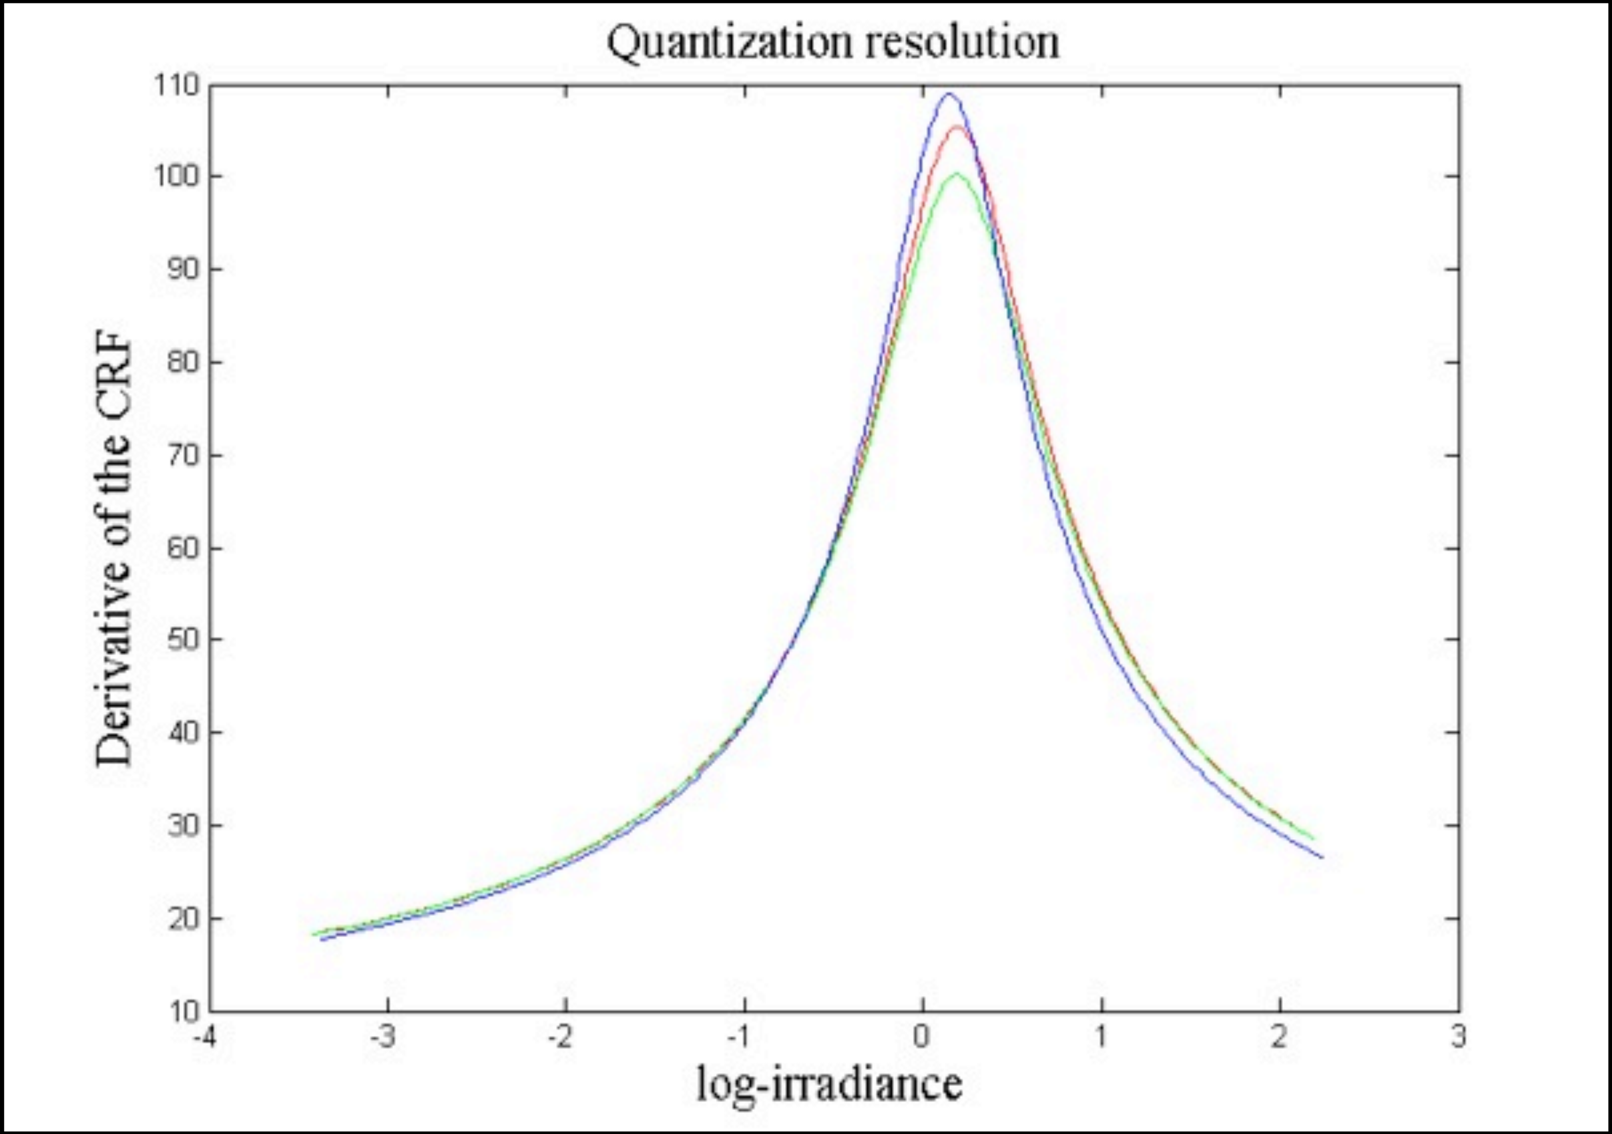

Quantization resolution
Derivative of the CRF
log-irradiance
110
100
90
80
70
60
50
40
30
20
10
-4
-3
-2
-1
0
1
2
3

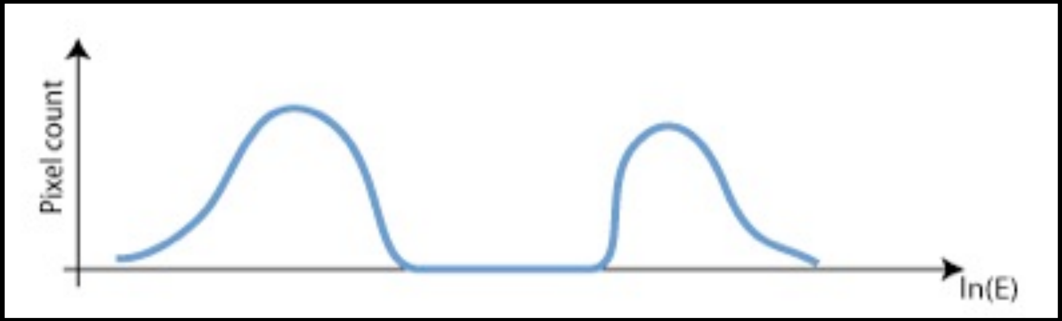
Pixel count
ln(E)

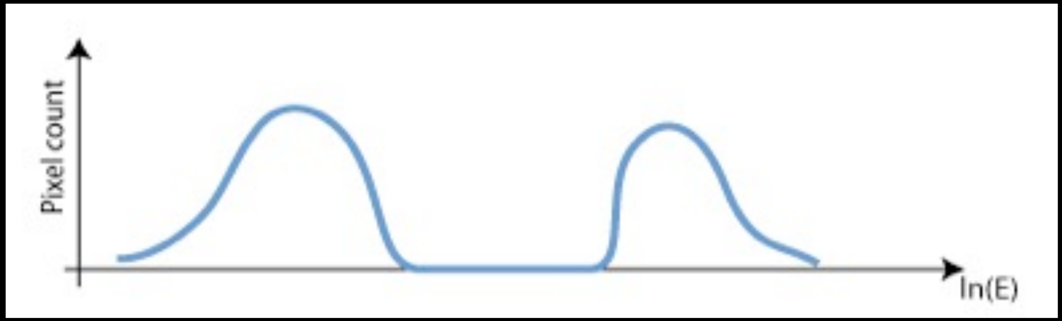

Pixel count
ln(E)

1
2

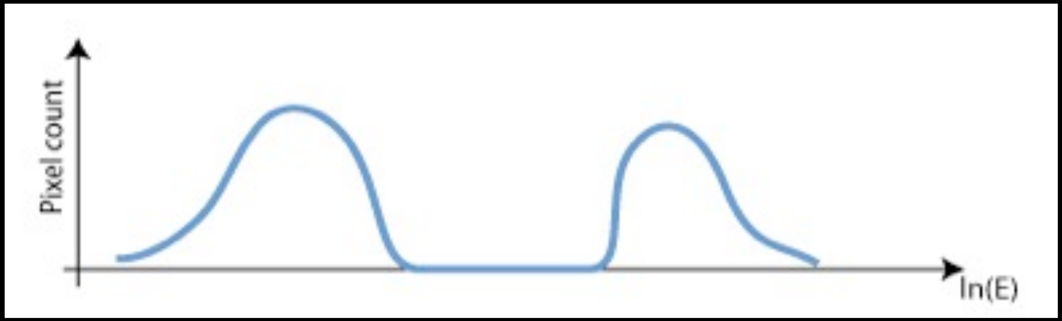

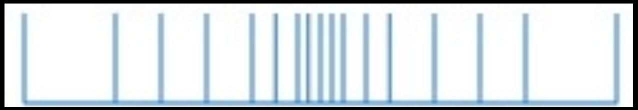

Δt_1

Δt_2

Histogram estimation

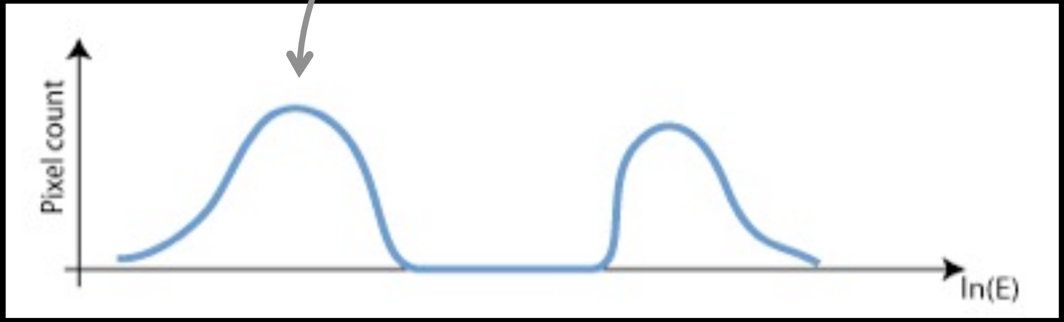

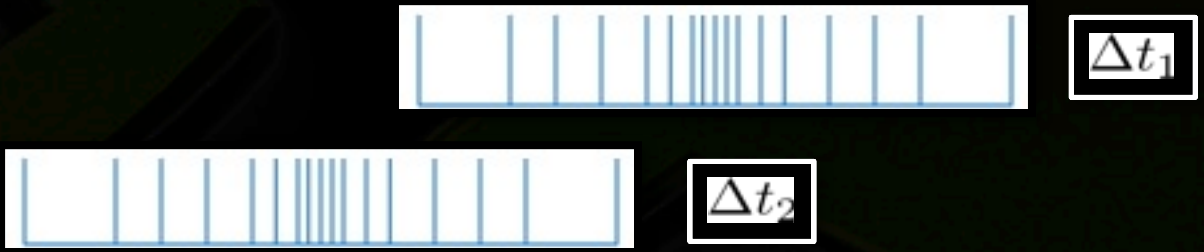

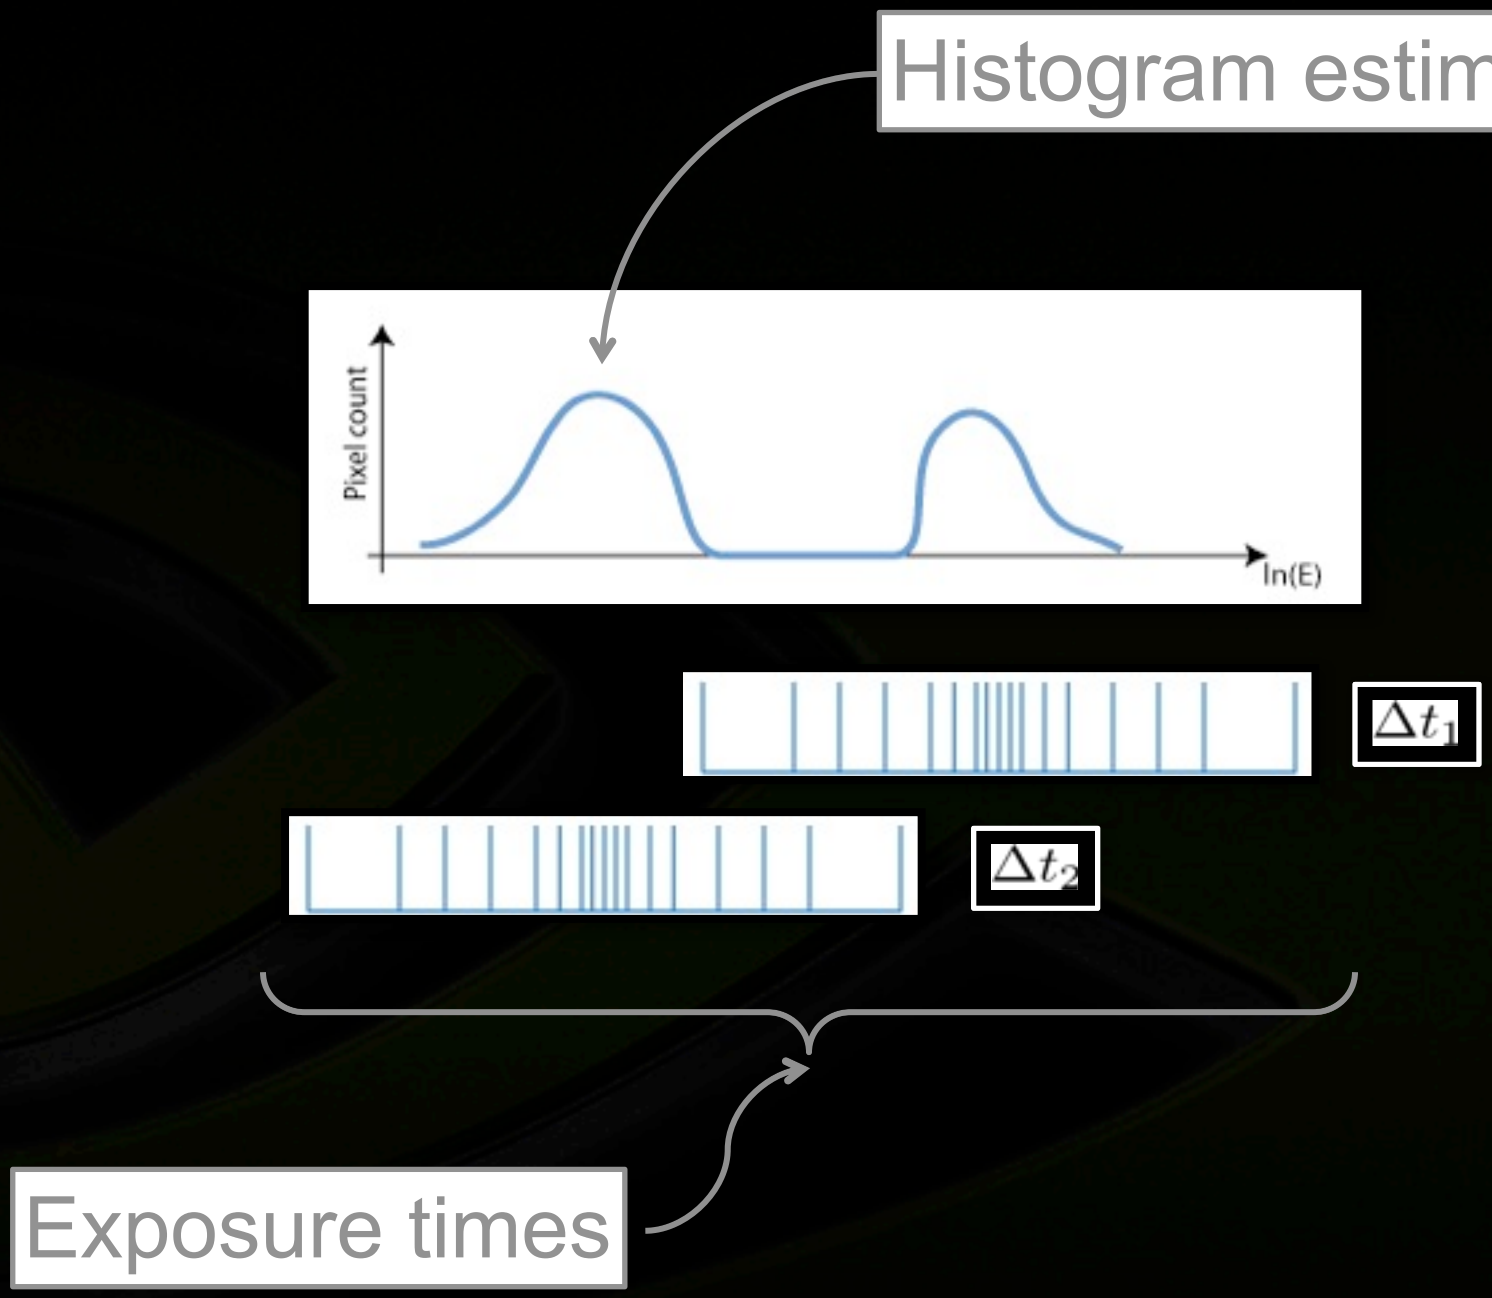
Histogram estimation
Pixel count
ln(E)
Δt_1
Δt_2
Exposure times

Histogram estimation

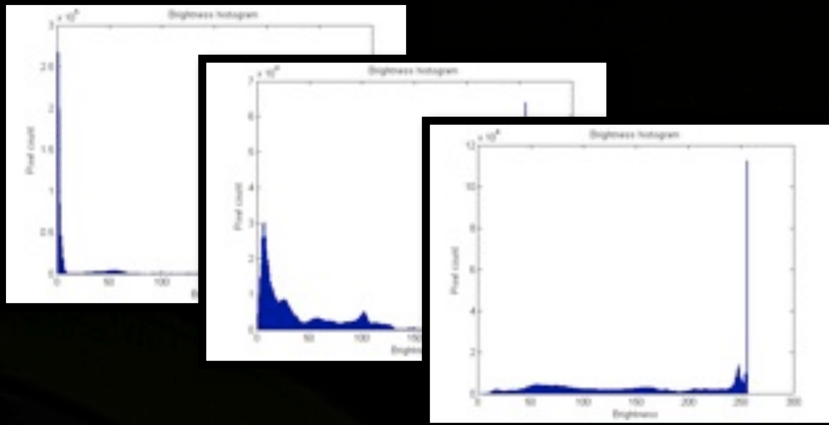

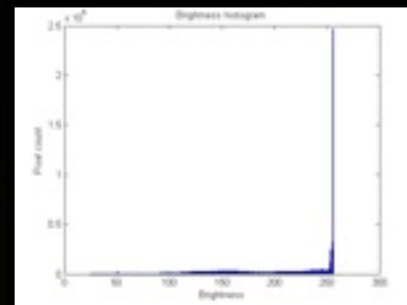

Histogram estimation

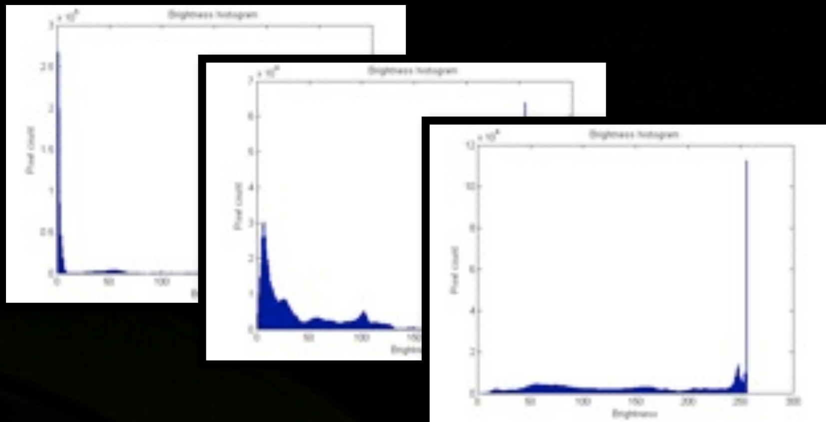

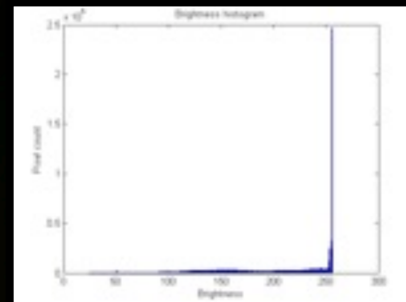

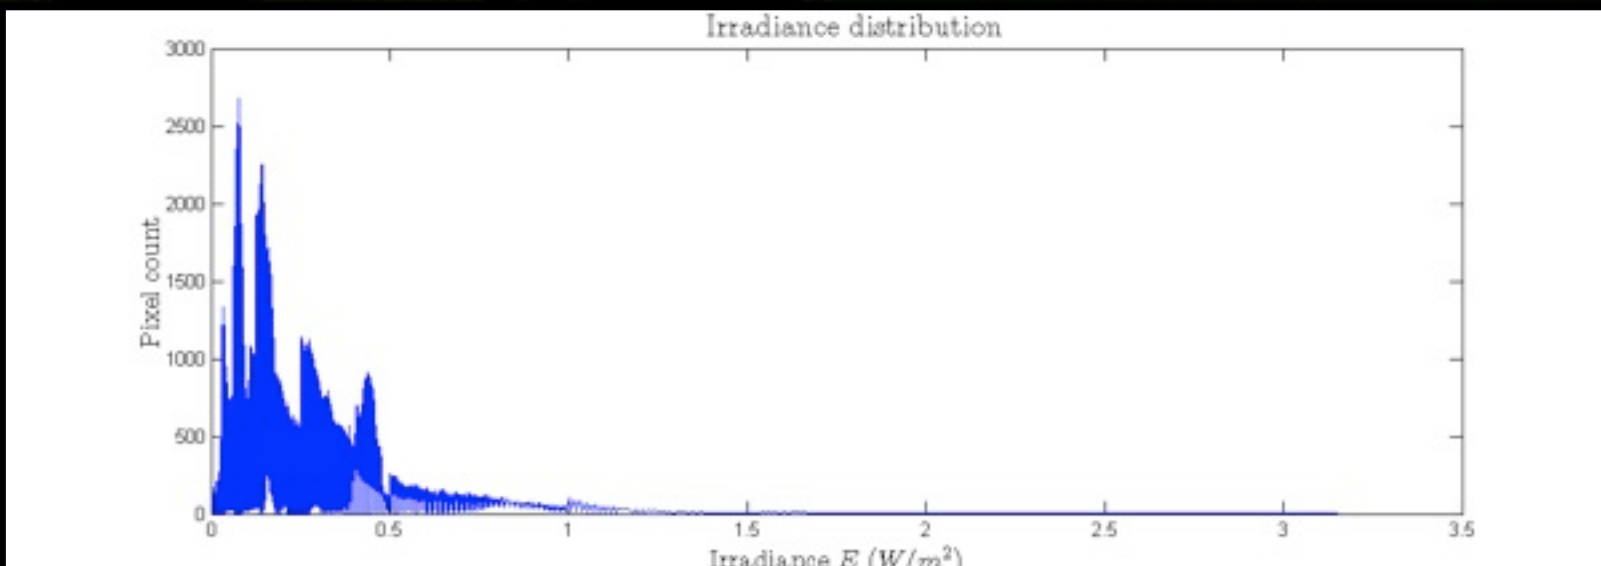

Cumulative Distribution Function (CDF)

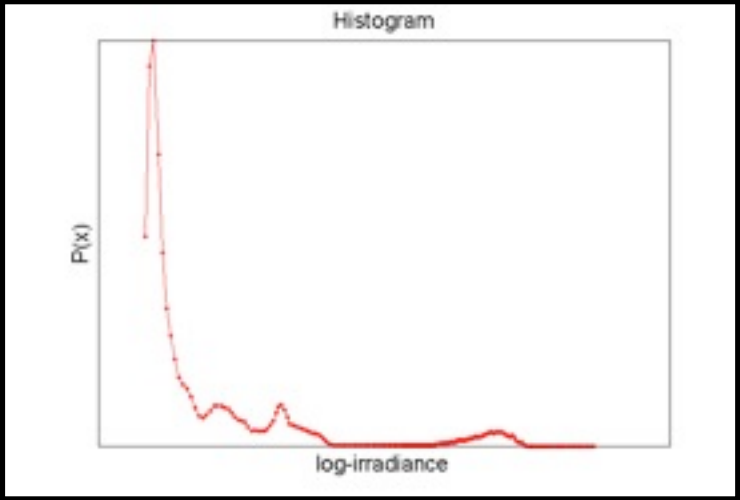

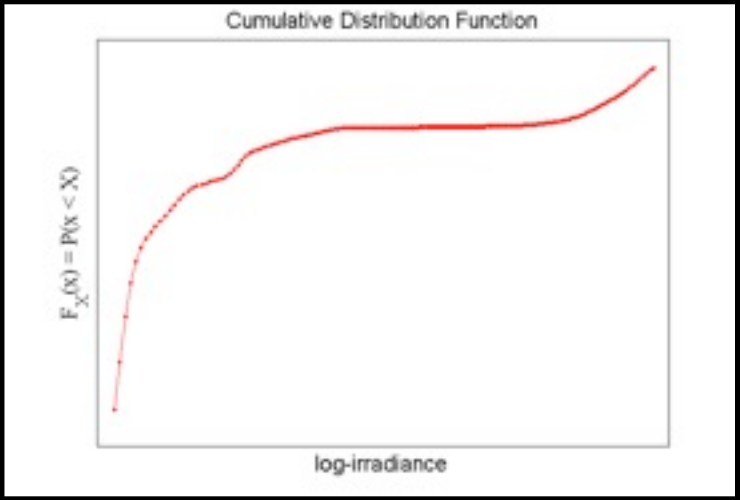

Histogram estimation

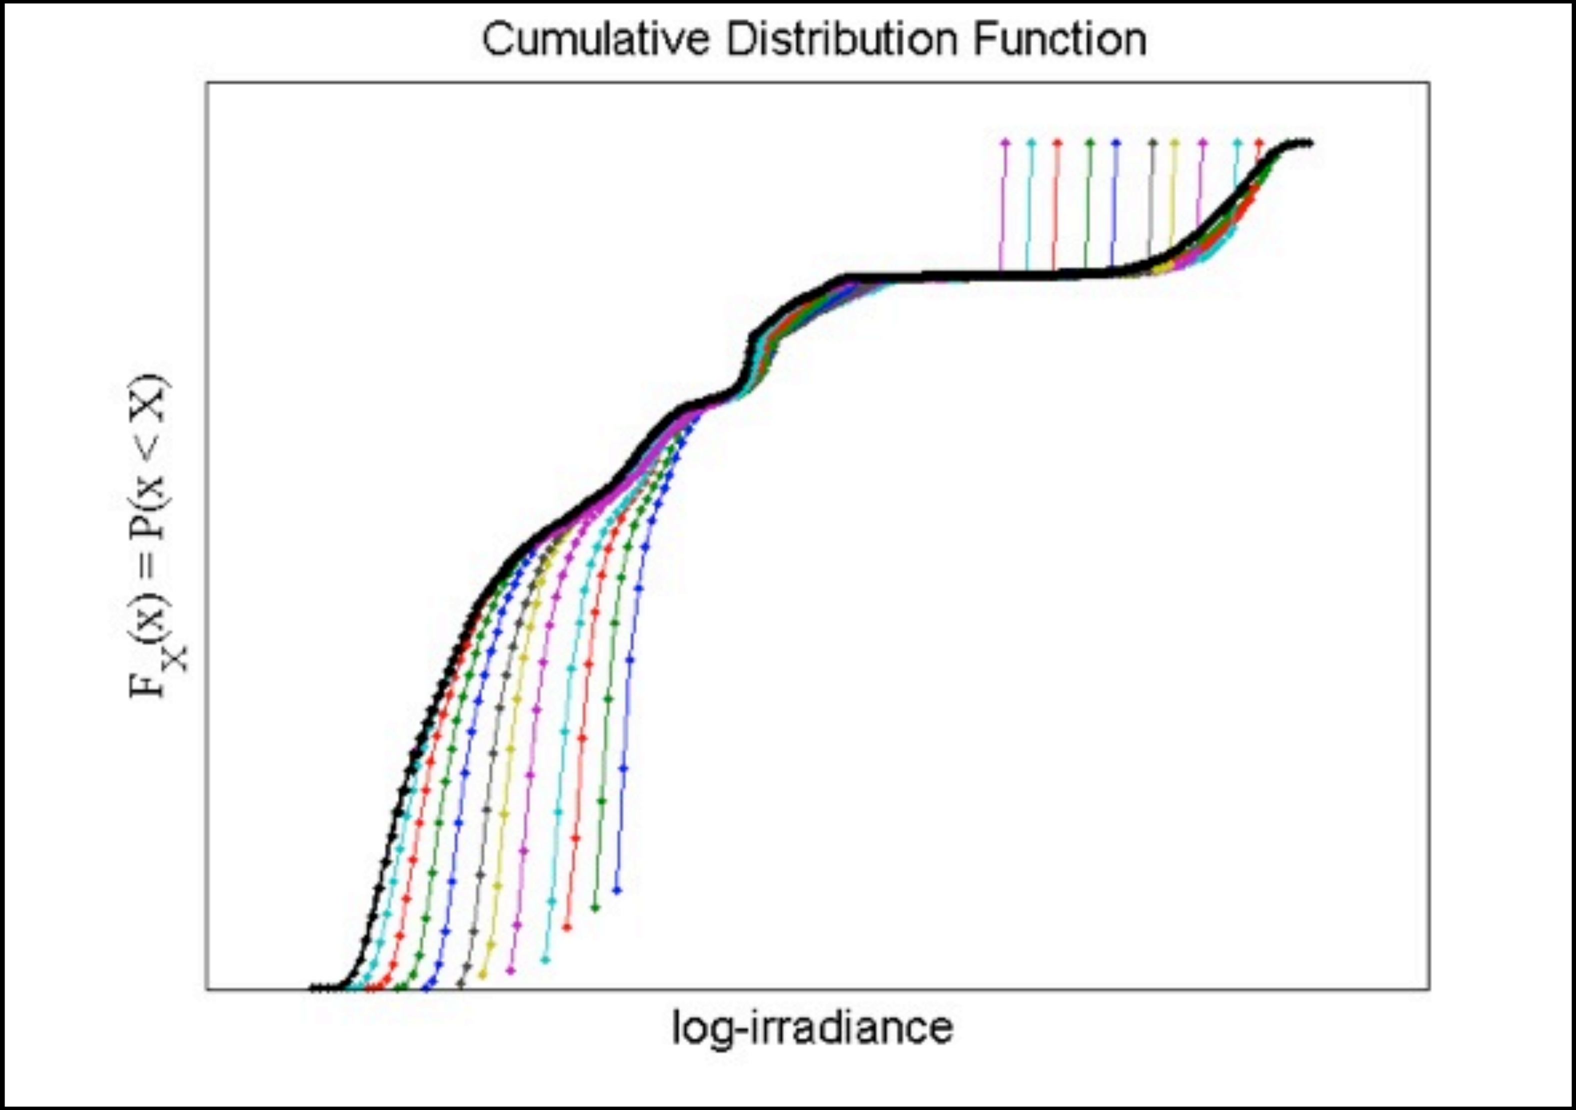

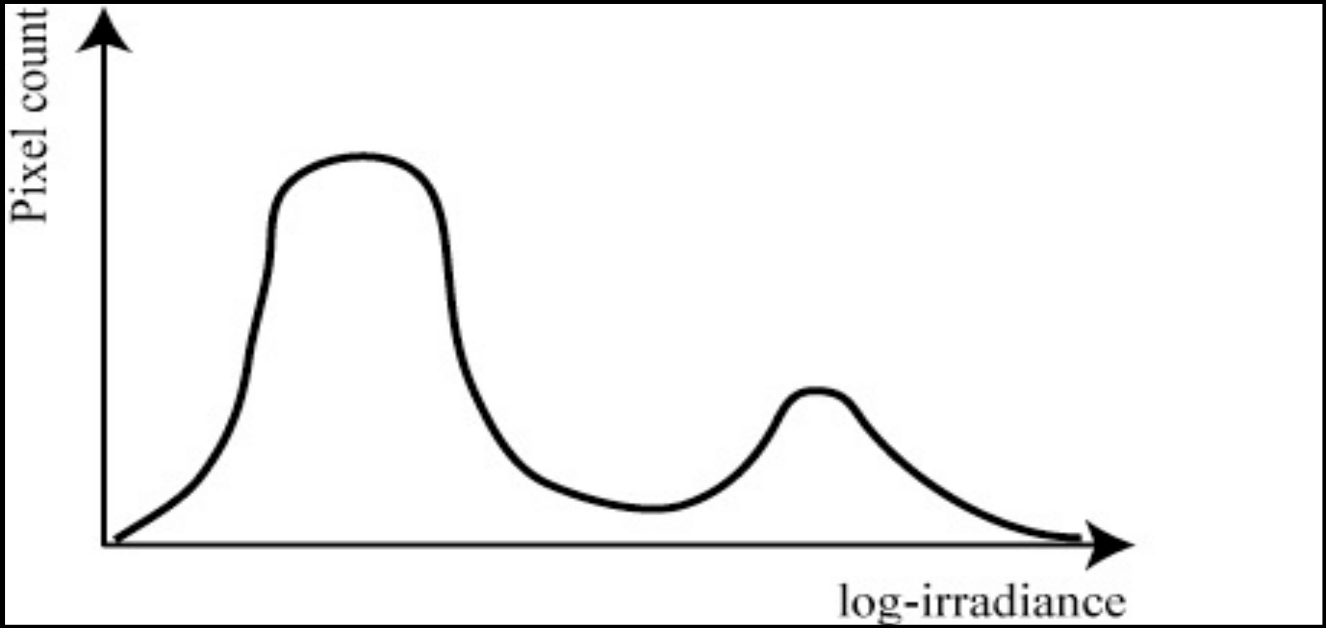
Pixel count
log-irradiance

$$t = \underset{t}{\operatorname{argmax}} \sum_i P_{\tilde{E}}(\tilde{E}_i) \frac{\left[f'(\tilde{E}_i + \tilde{t})\right]^2}{\left[a' e^{\tilde{E}_i + \tilde{t}} + b'\right] \left[f'(\tilde{E}_i + \tilde{t})\right]^2 + \frac{\Delta^2}{12}}$$

HDR histogram

$$t = \operatorname*{argmax}_t \sum_i P_{\tilde{E}}(\tilde{E}_i) \frac{\left[f'(\tilde{E}_i + \tilde{t})\right]^2}{\left[a' e^{\tilde{E}_i + \tilde{t}} + b'\right] \left[f'(\tilde{E}_i + \tilde{t})\right]^2 + \frac{\Delta^2}{12}}$$

Noise

CRF

Quantization

$$t = \underset{t}{\operatorname{argmax}} \sum_i P_{\tilde{E}}(\tilde{E}_i) \frac{\left[f'(\tilde{E}_i + \tilde{t})\right]^2}{\left[a' e^{\tilde{E}_i + \tilde{t}} + b'\right] \left[f'(\tilde{E}_i + \tilde{t})\right]^2 + \frac{\Delta^2}{12}}$$

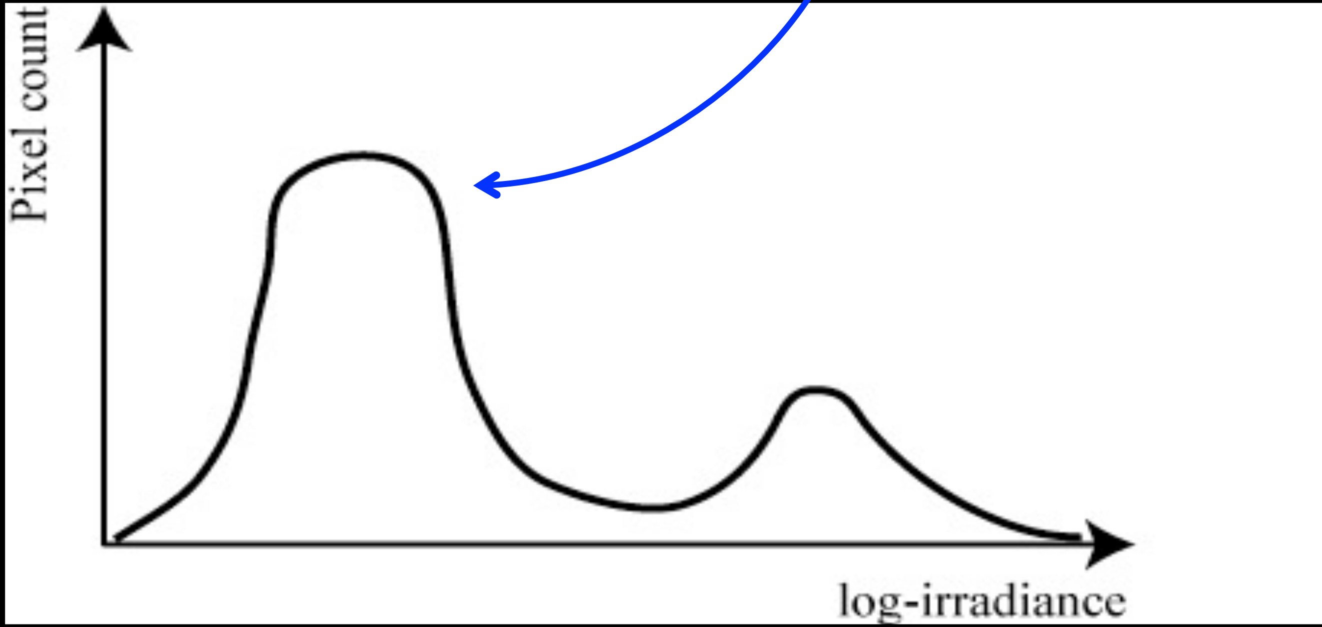

$$t = \underset{t}{\operatorname{argmax}} \sum_i P_{\tilde{E}}(\tilde{E}_i) \frac{\left[f'(\tilde{E}_i + \tilde{t})\right]^2}{\left[a' e^{\tilde{E}_i + \tilde{t}} + b'\right]\left[f'(\tilde{E}_i + \tilde{t})\right]^2 + \frac{\Delta^2}{12}}$$

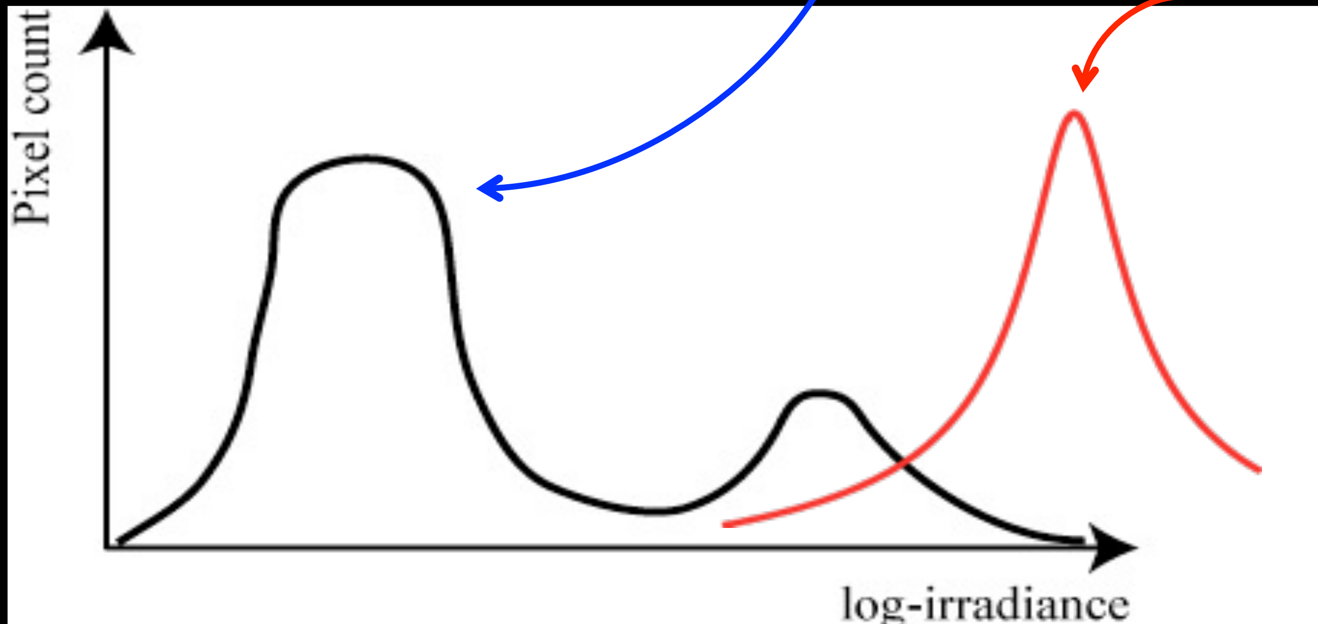

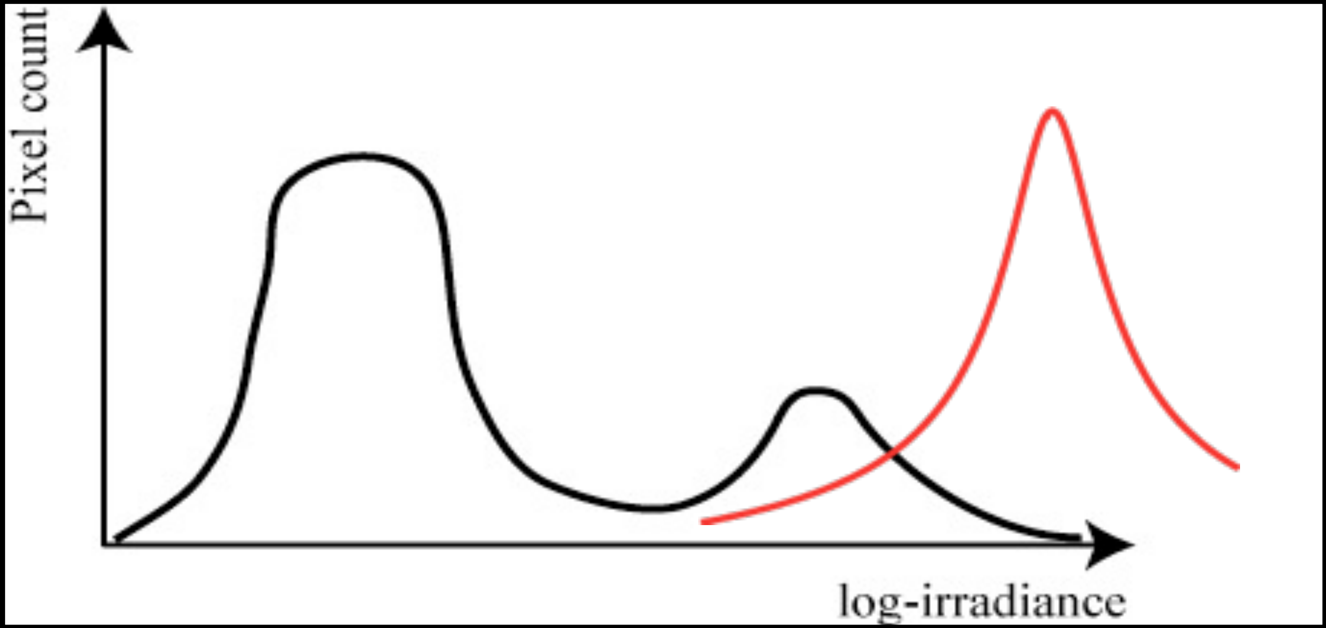
Pixel count
log-irradiance

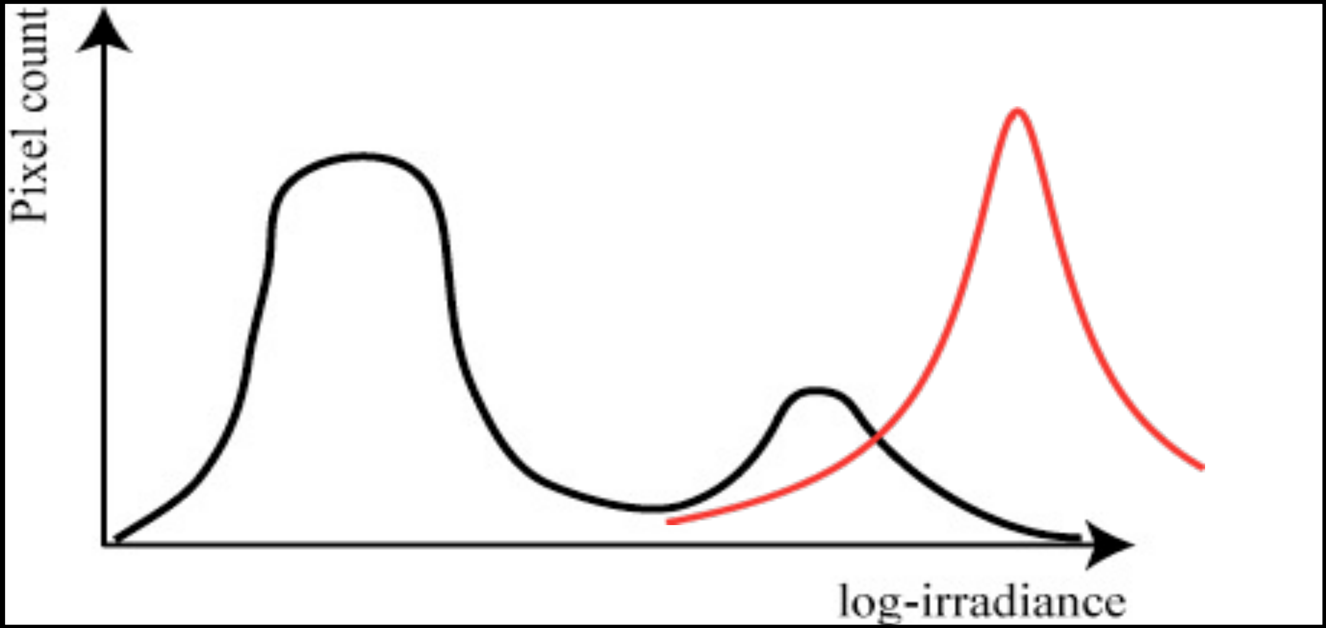
Pixel count
log-irradiance

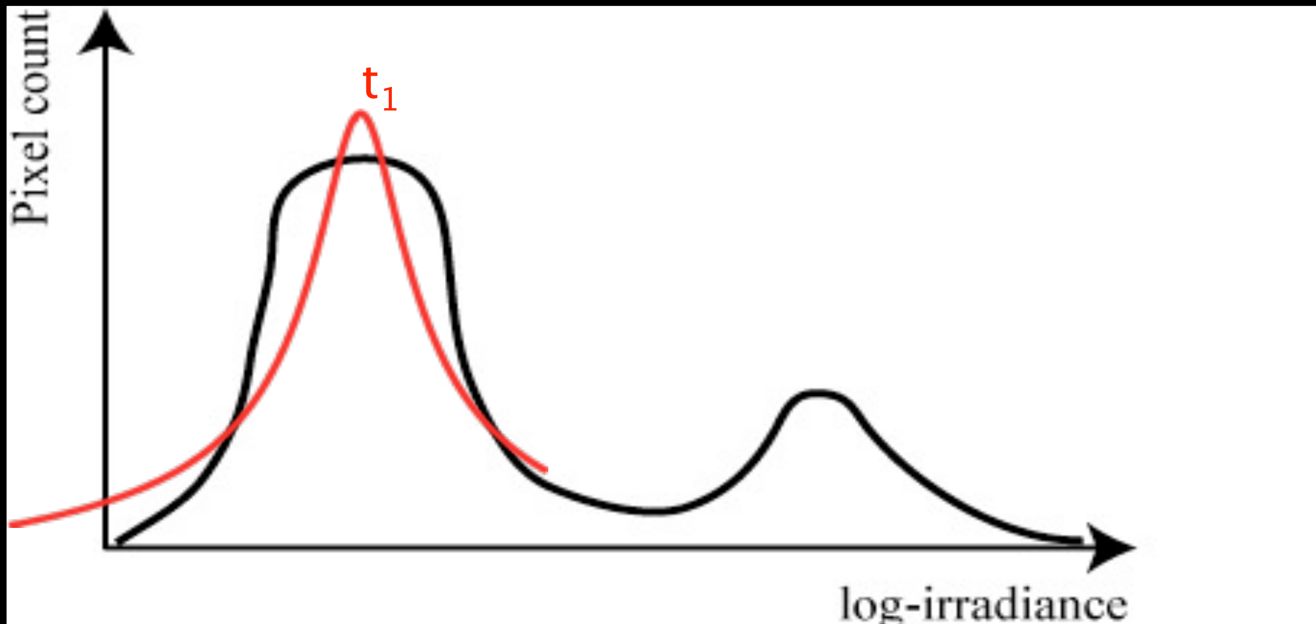
Pixel count
t_1
log-irradiance

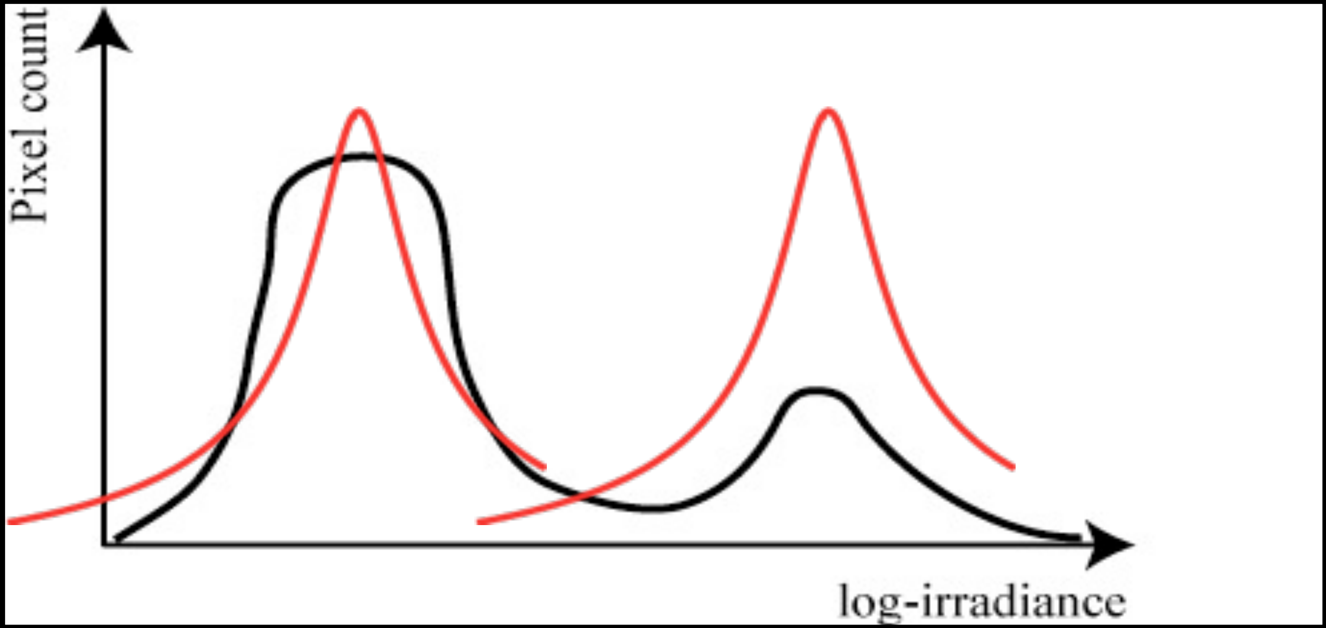
Pixel count
log-irradiance

NVIDIA

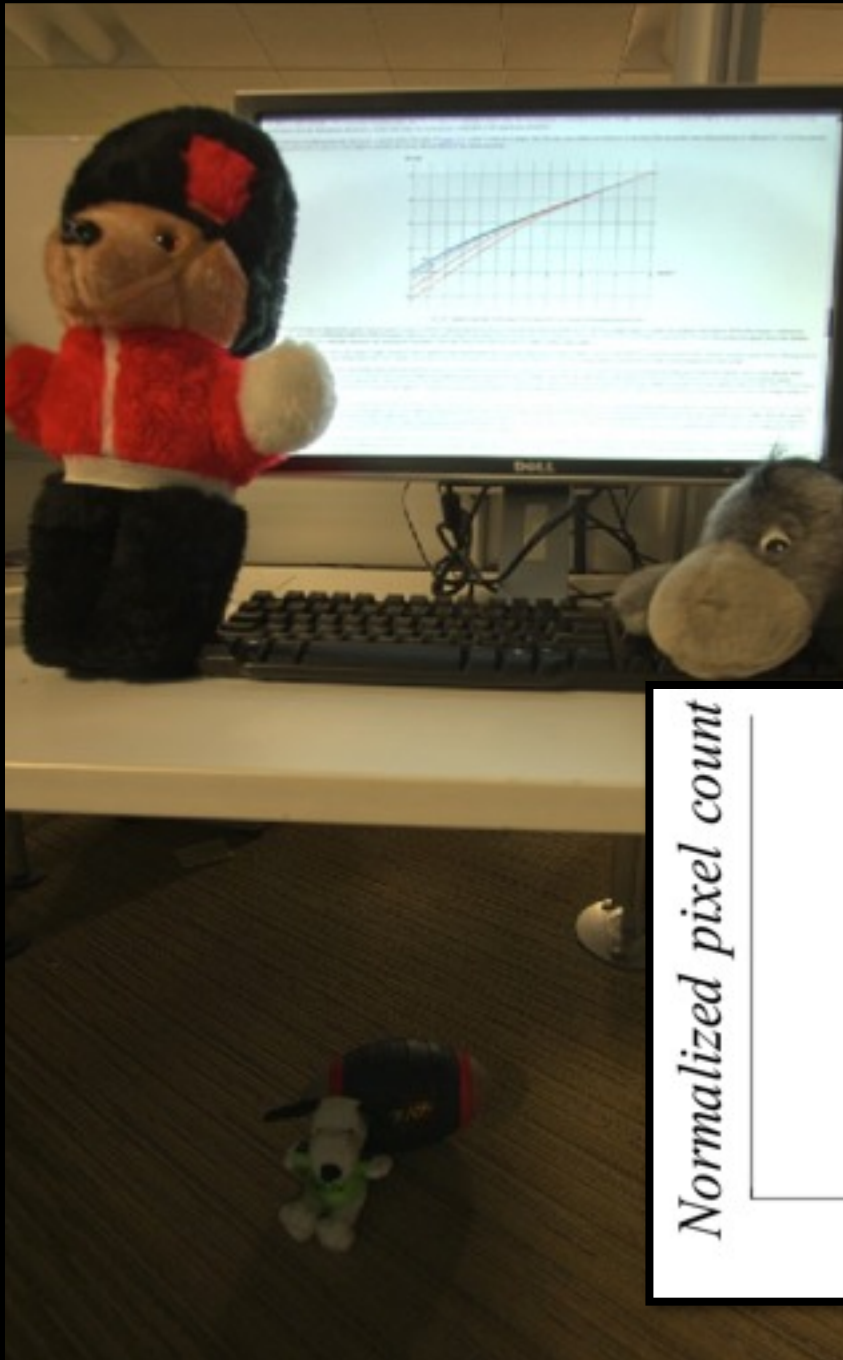

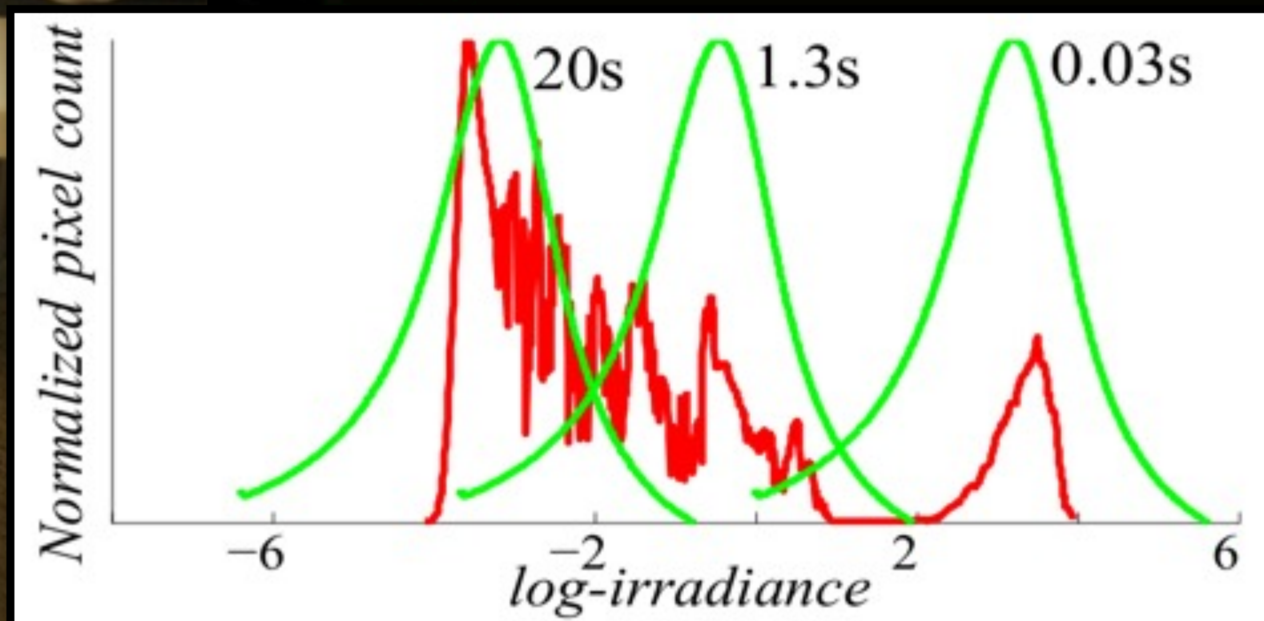

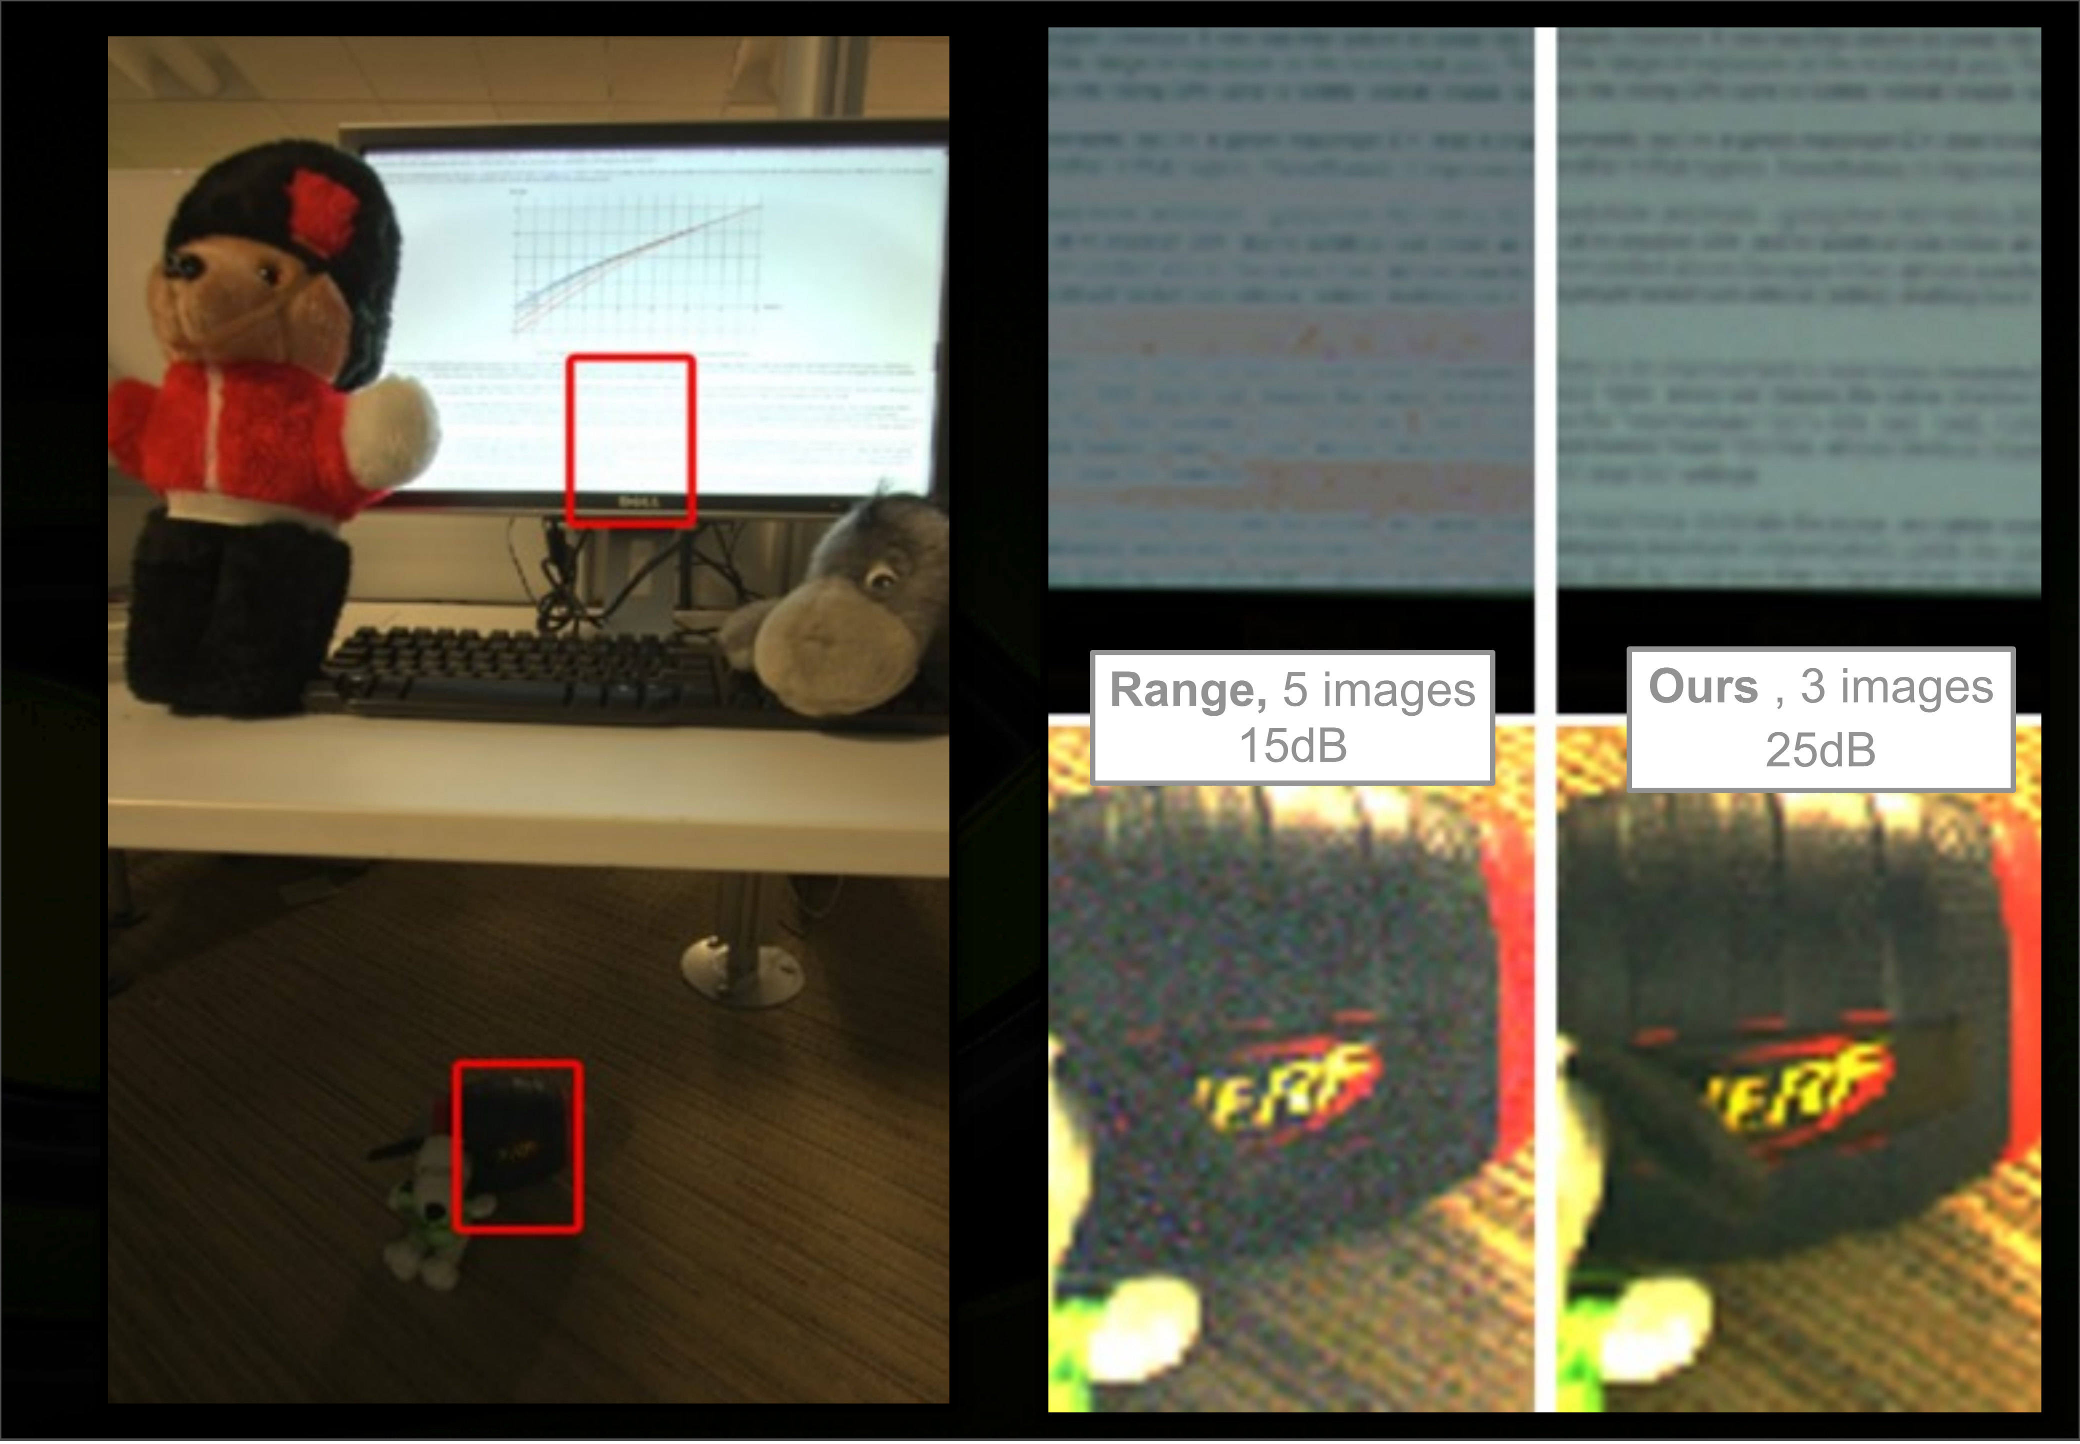
Range, 5 images
15dB
Ours , 3 images
25dB

O.Gallo, W-C Chen, N.Gelfand, M.Tico, K.Pulli
Artifact-free High Dynamic Range Imaging
IEEE International Conference on Computational Photography 2009

O.Gallo, W-C Chen, N.Gelfand, M.Tico, K.Pulli
Artifact-free High Dynamic Range Imaging
IEEE International Conference on Computational Photography 2009

O.Gallo, W-C Chen, N.Gelfand, M.Tico, K.Pulli
Artifact-free High Dynamic Range Imaging
IEEE International Conference on Computational Photography 2009

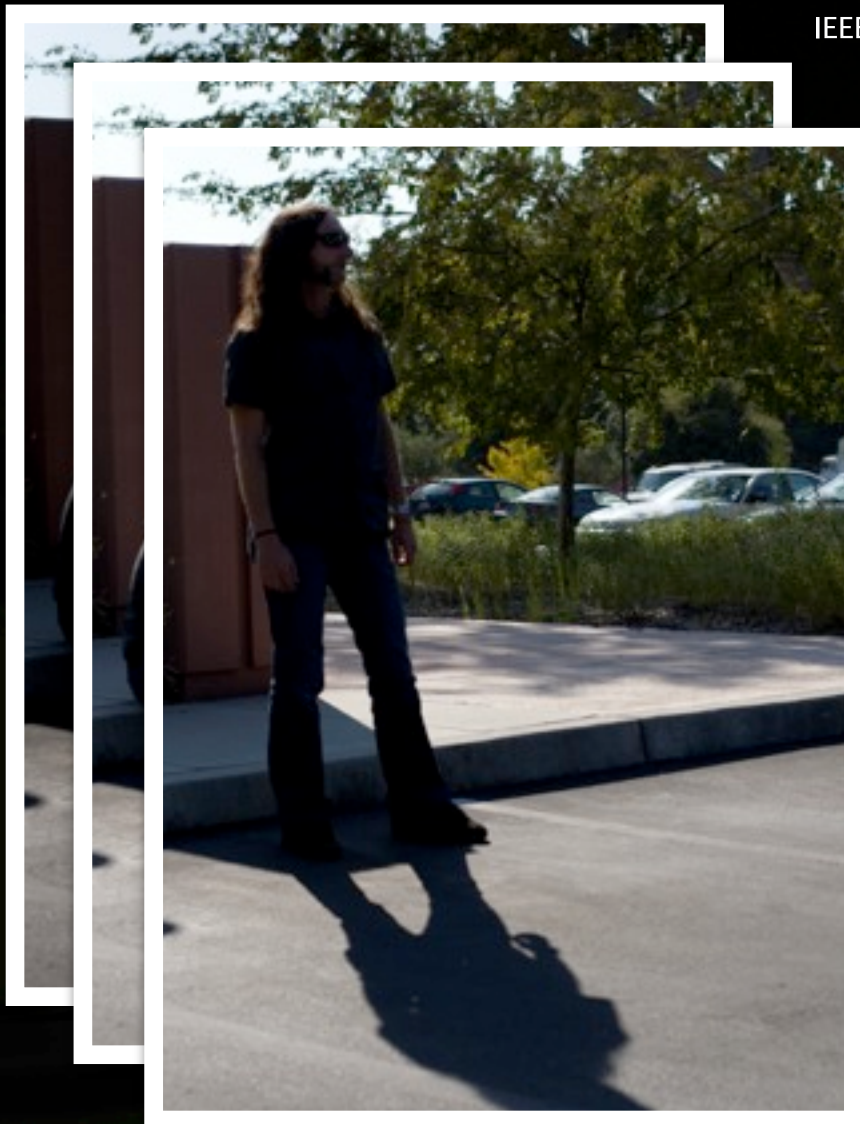

O.Gallo, W-C Chen, N.Gelfand, M.Tico, K.Pulli
Artifact-free High Dynamic Range Imaging
IEEE International Conference on Computational Photography 2009

O.Gallo, W-C Chen, N.Gelfand, M.Tico, K.Pulli
Artifact-free High Dynamic Range Imaging
IEEE International Conference on Computational Photography 2009

NVIDIA

O.Gallo, W-C Chen, N.Gelfand, M.Tico, K.Pulli
Artifact-free High Dynamic Range Imaging
IEEE International Conference on Computational Photography 2009

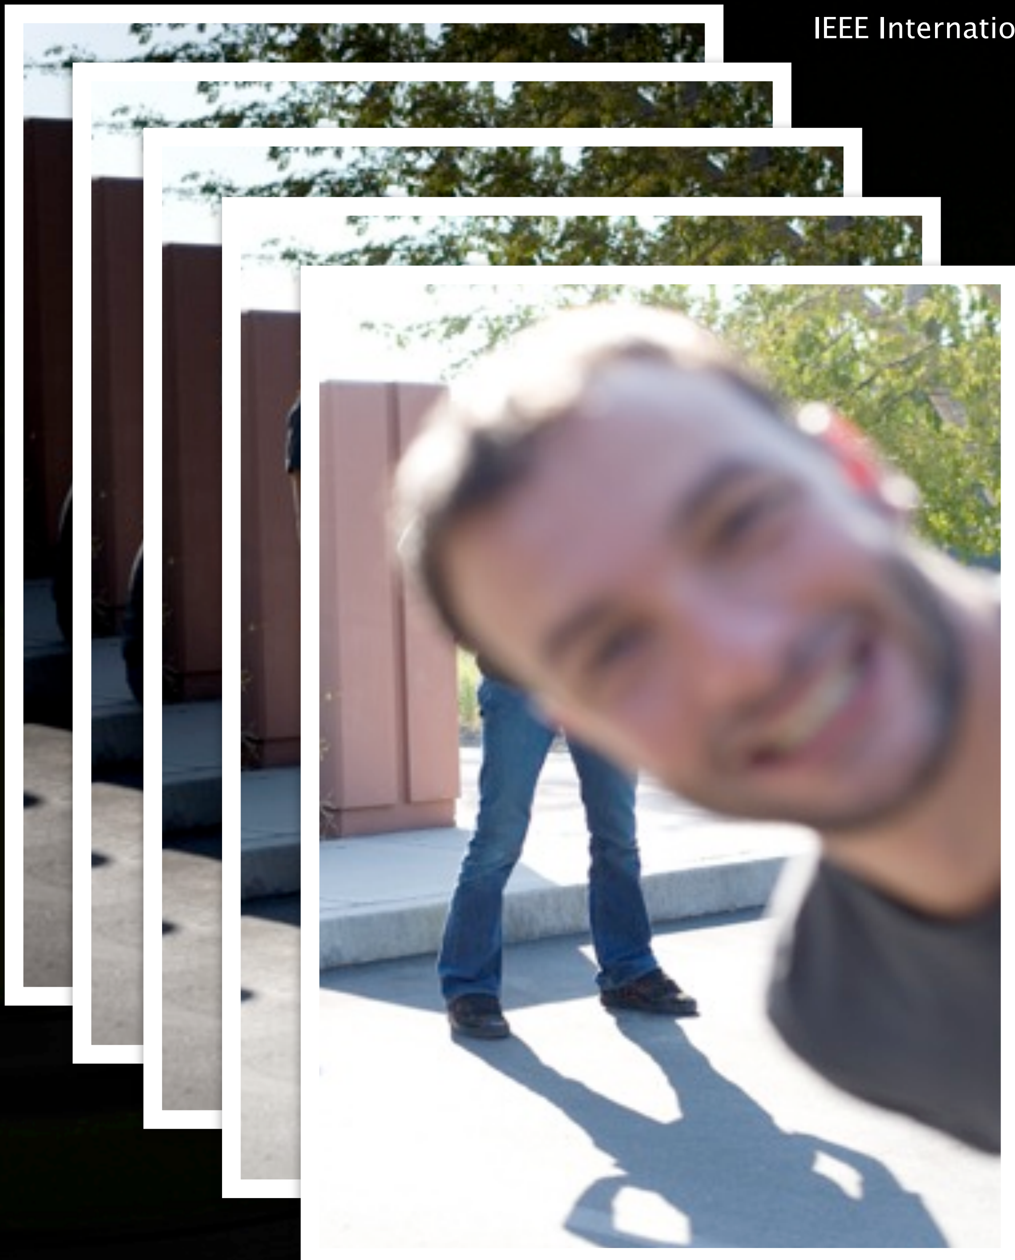

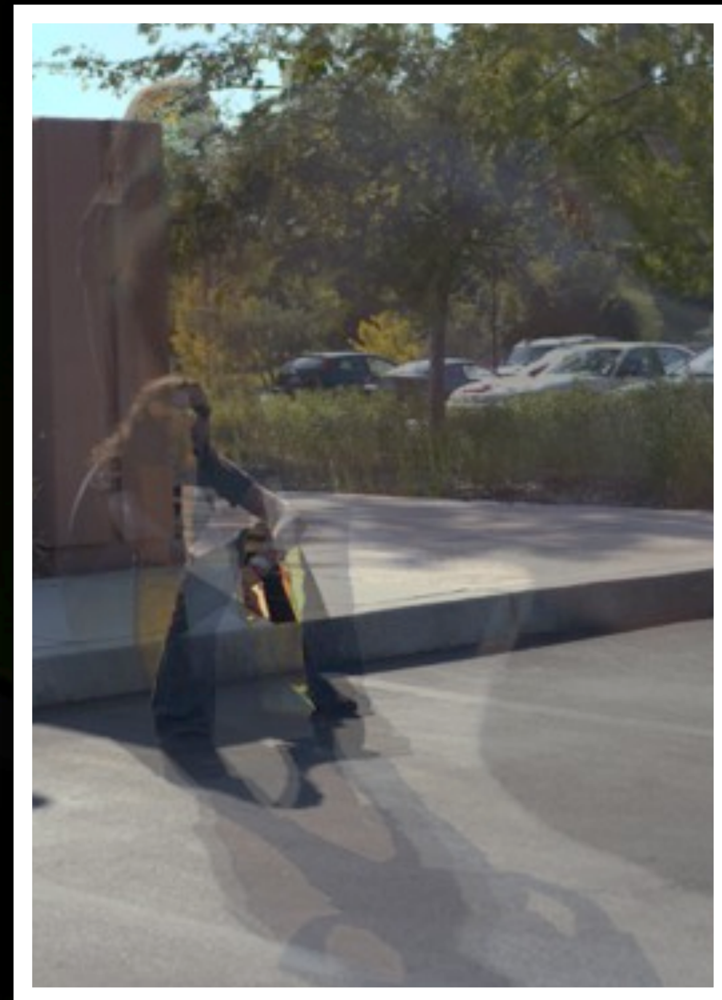

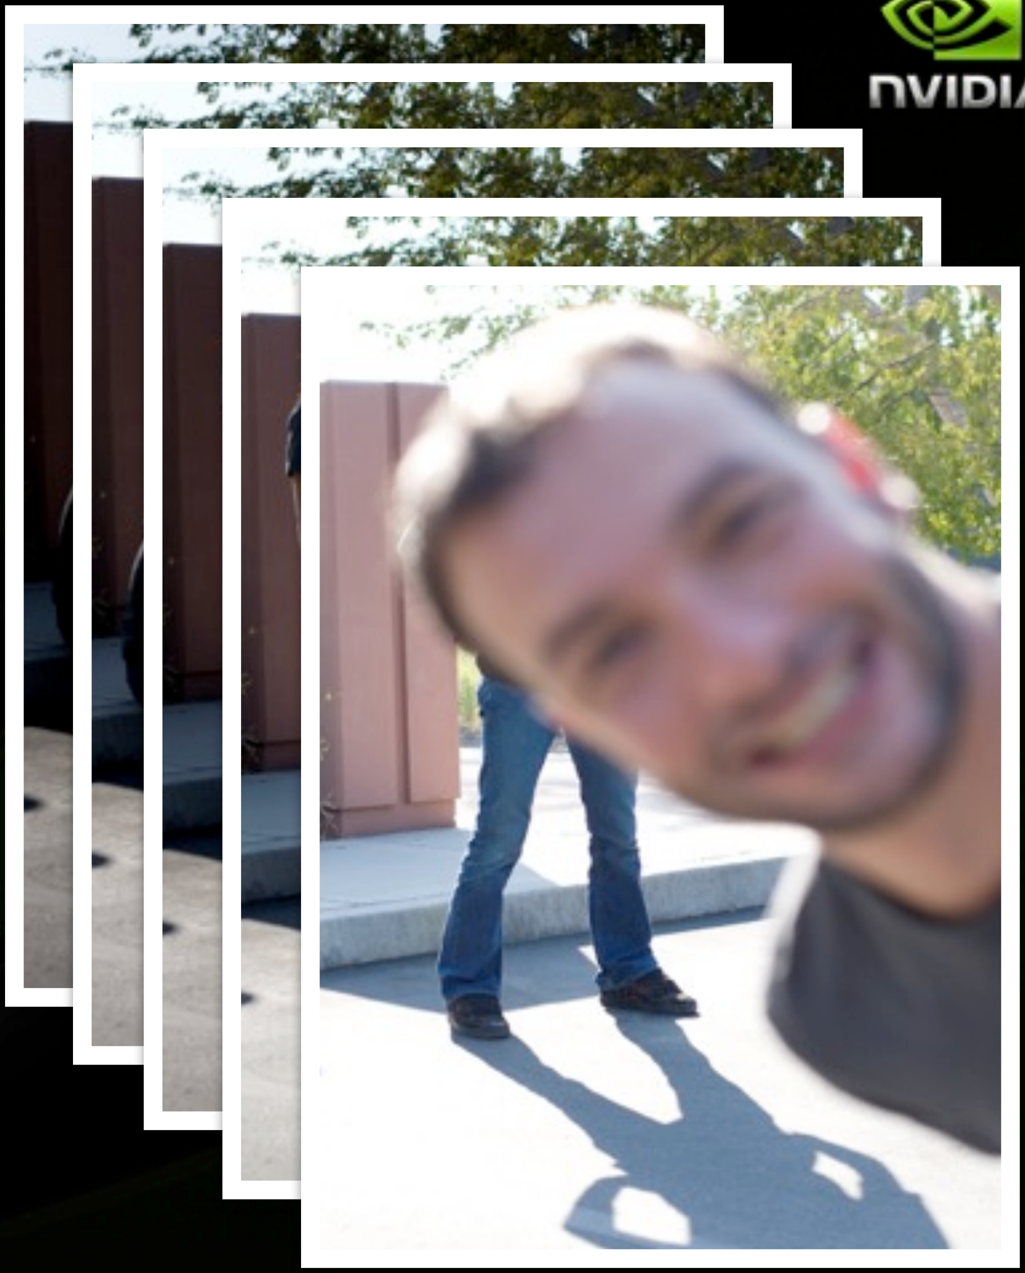

Reference Frame Selection

Consistency Detection

HDR Generation

Poisson Blending

NVIDIA Research

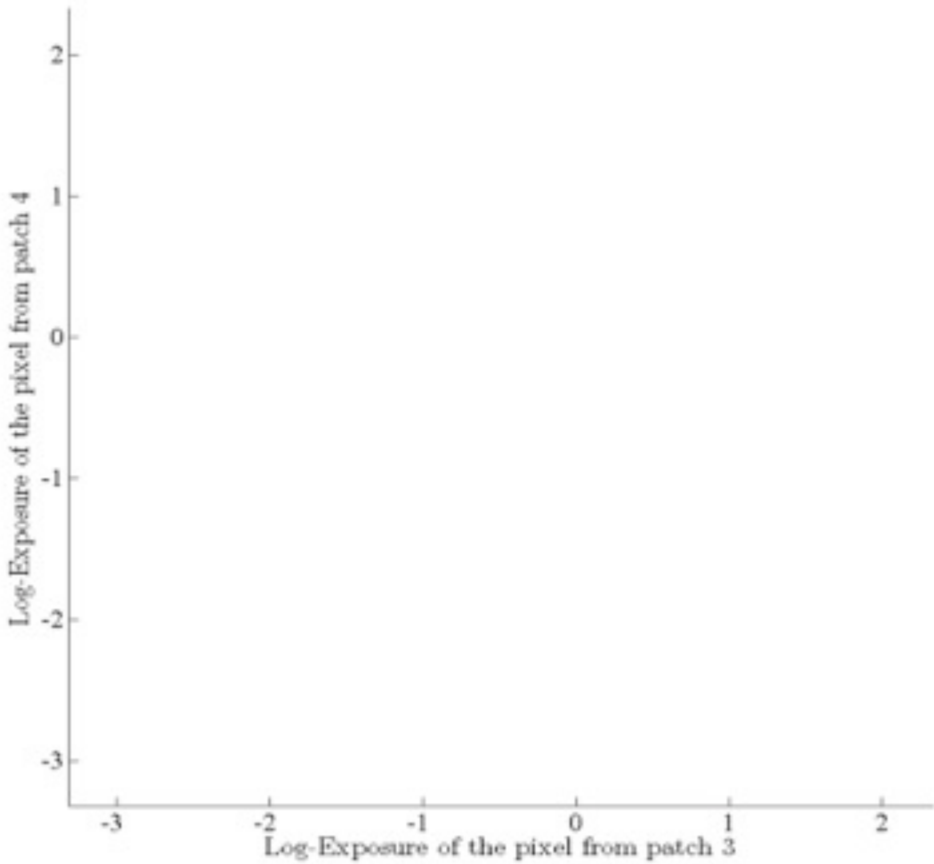

Log-Exposure of the pixel from patch 4
2
1
0
-1
-2
-3
-3
-2
-1
0
1
2
Log-Exposure of the pixel from patch 3

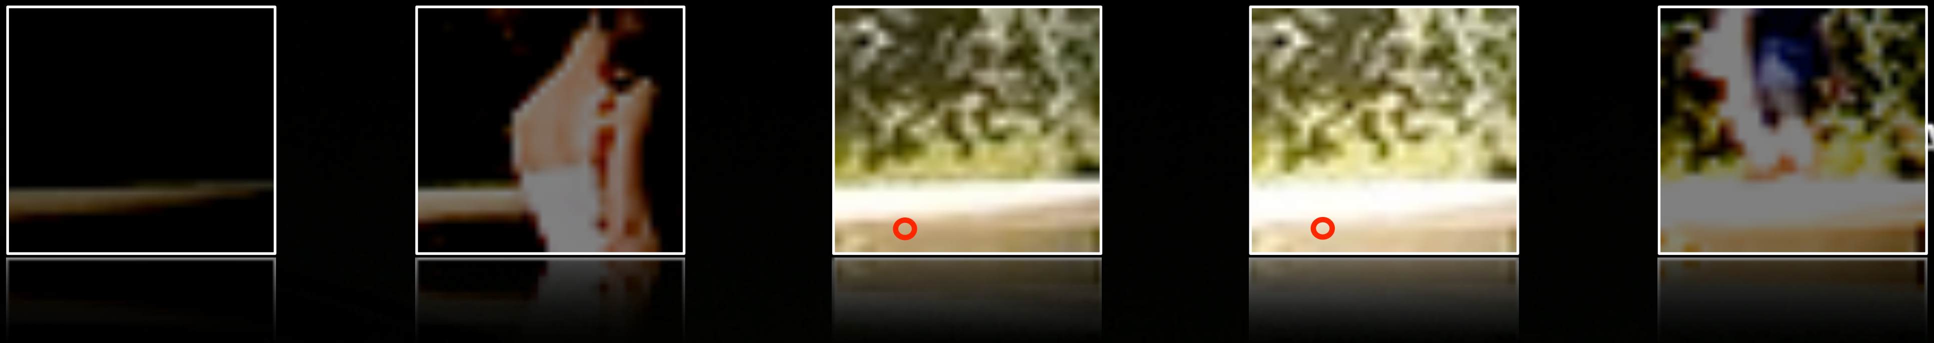

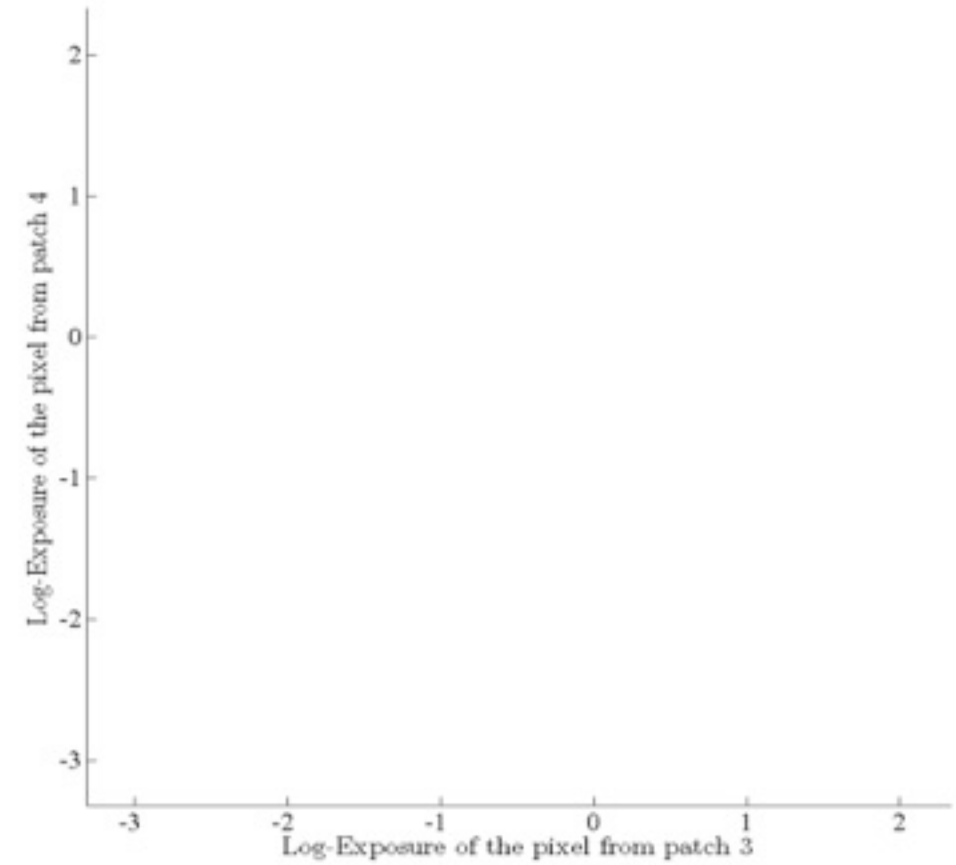
Log-Exposure of the pixel from patch 4
Log-Exposure of the pixel from patch 3
2
1
0
-1
-2
-3
-3
-2
-1
0
1
2

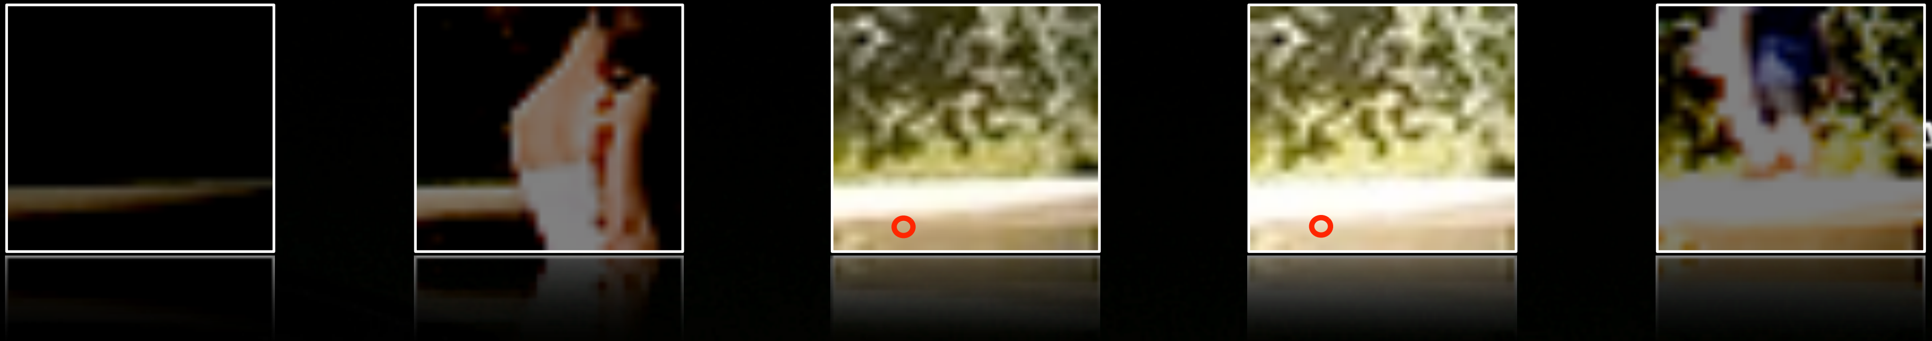

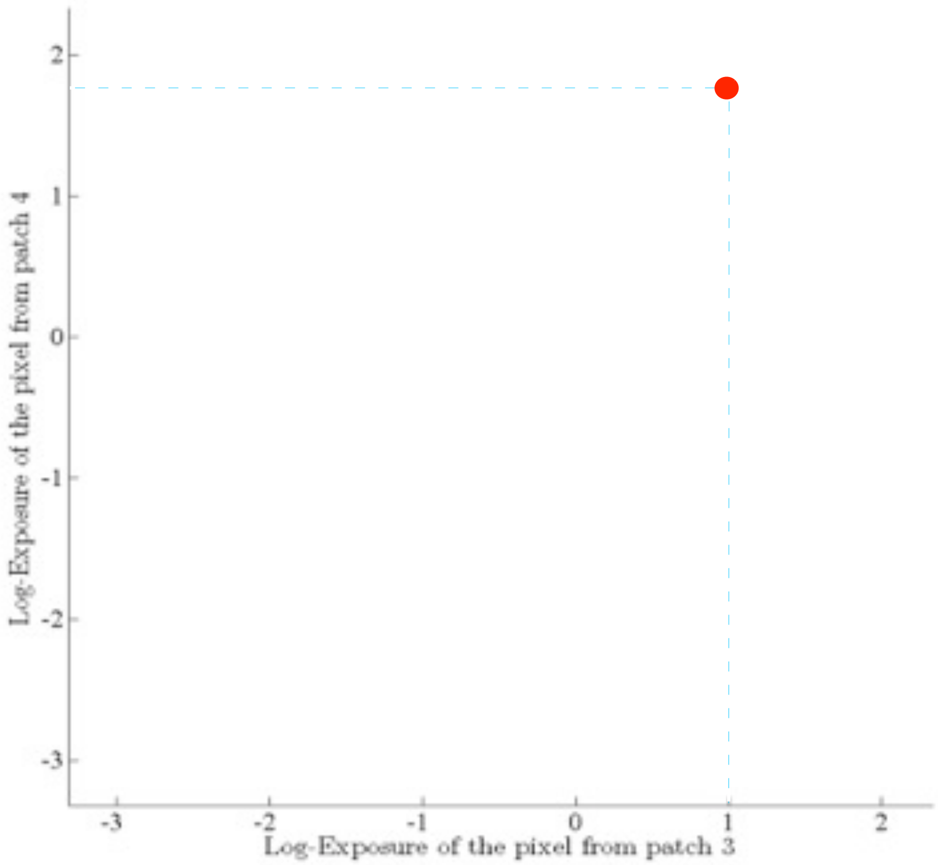

NVIDIA Research

Monday, March 5, 12

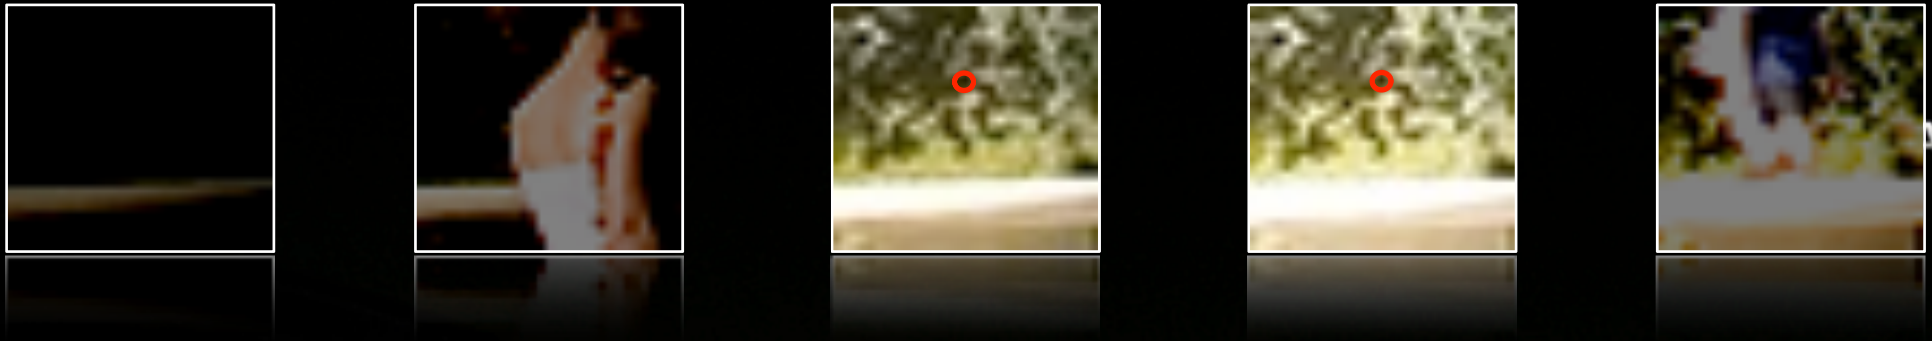

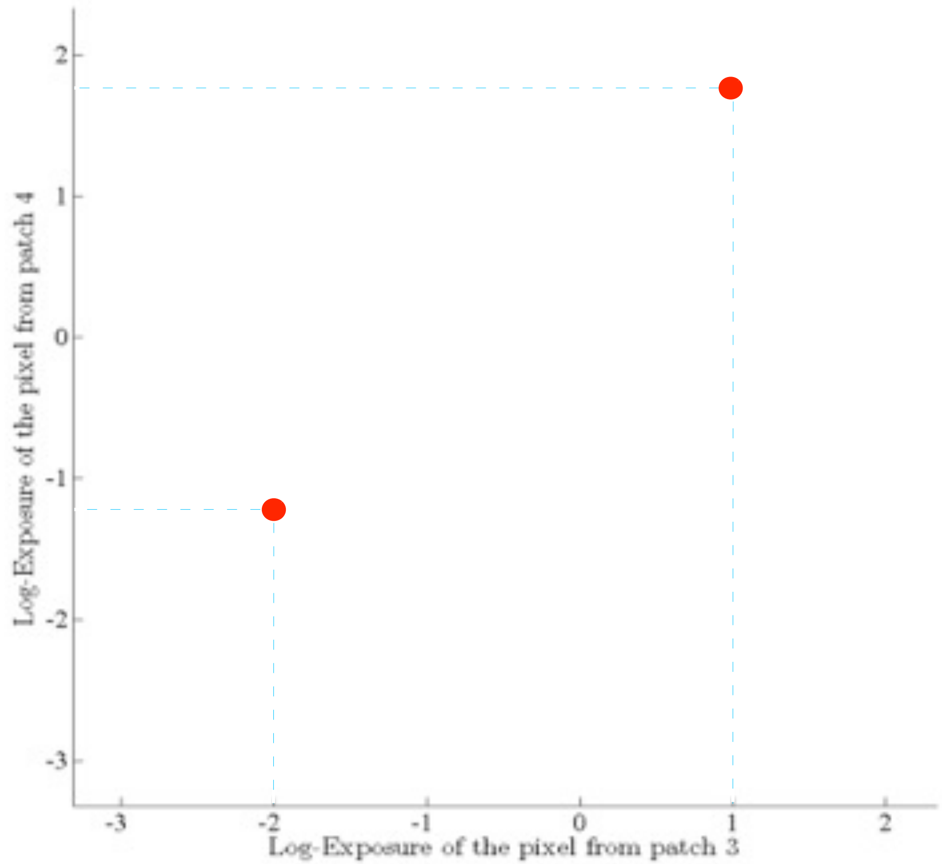
2
1
0
-1
-2
-3
Log-Exposure of the pixel from patch 4
-3
-2
-1
0
1
2
Log-Exposure of the pixel from patch 3

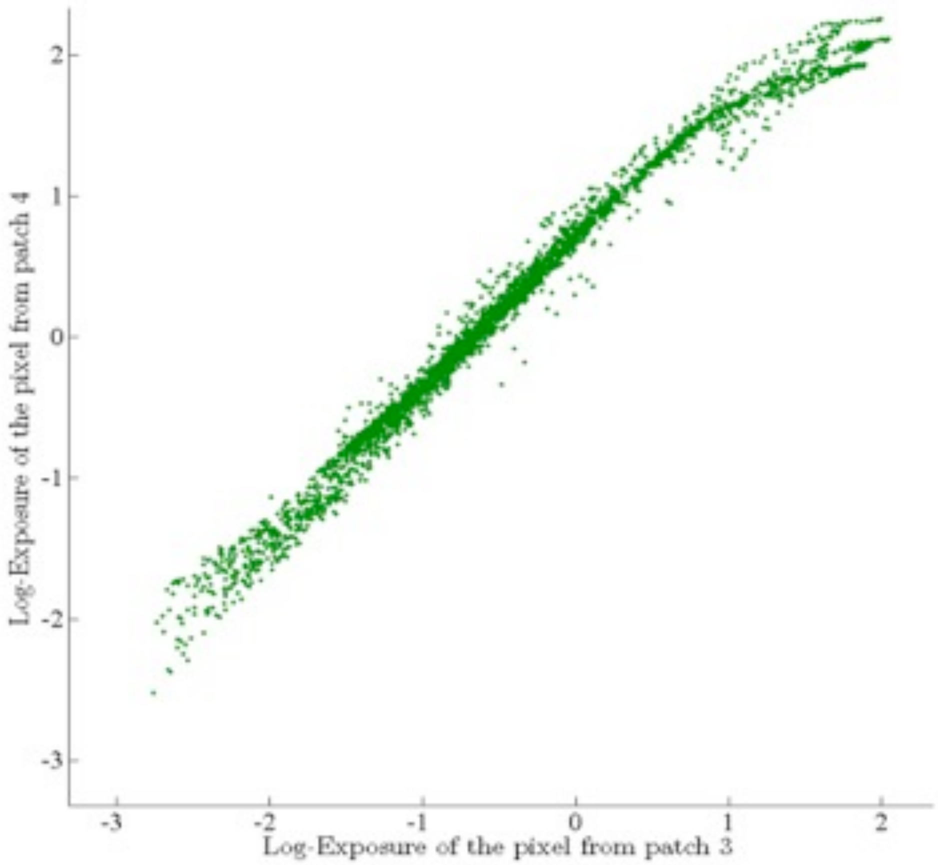
Log-Exposure of the pixel from patch 4
Log-Exposure of the pixel from patch 3
2
1
0
-1
-2
-3

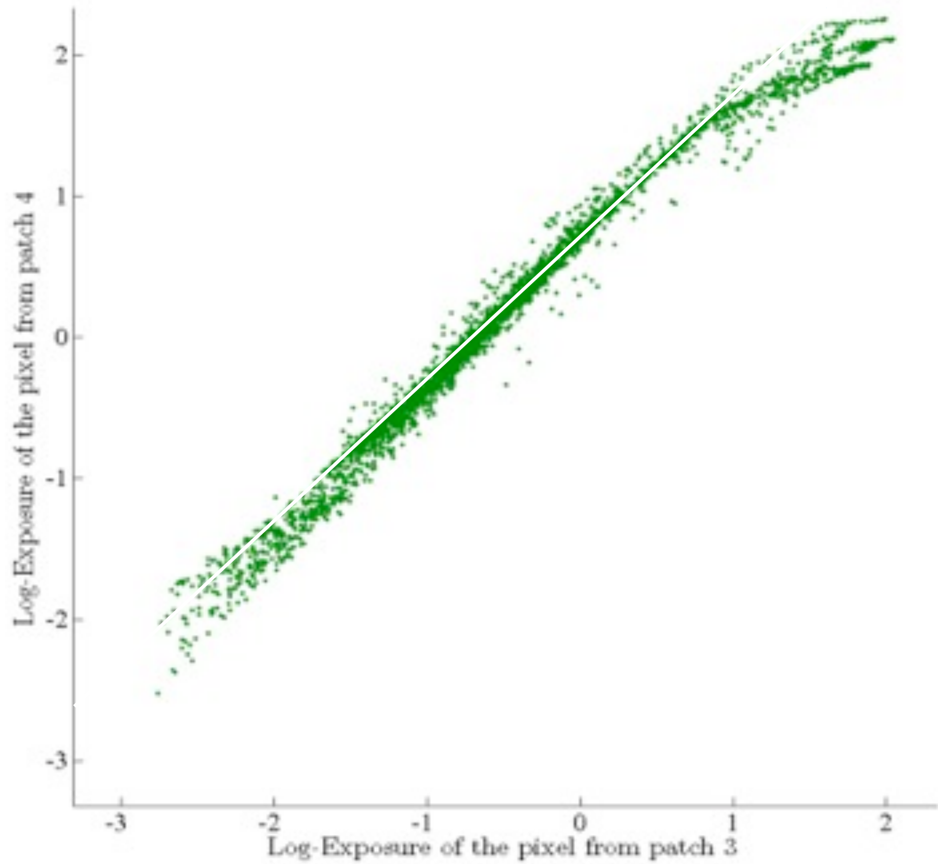
Log-Exposure of the pixel from patch 4
Log-Exposure of the pixel from patch 3
2
1
0
-1
-2
-3
-3
-2
-1
0
1
2

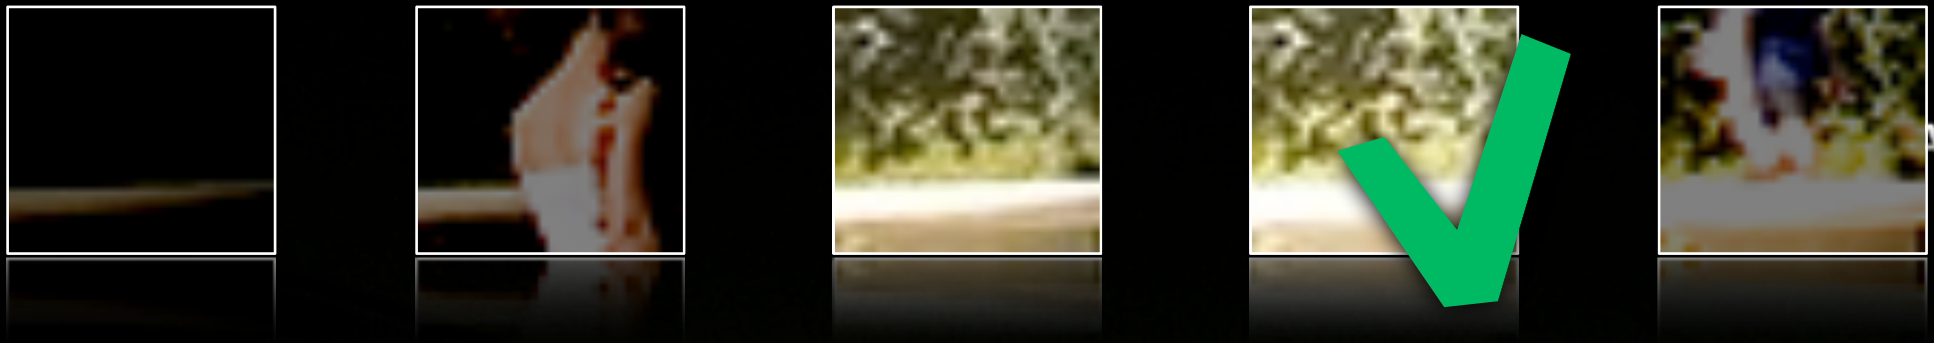

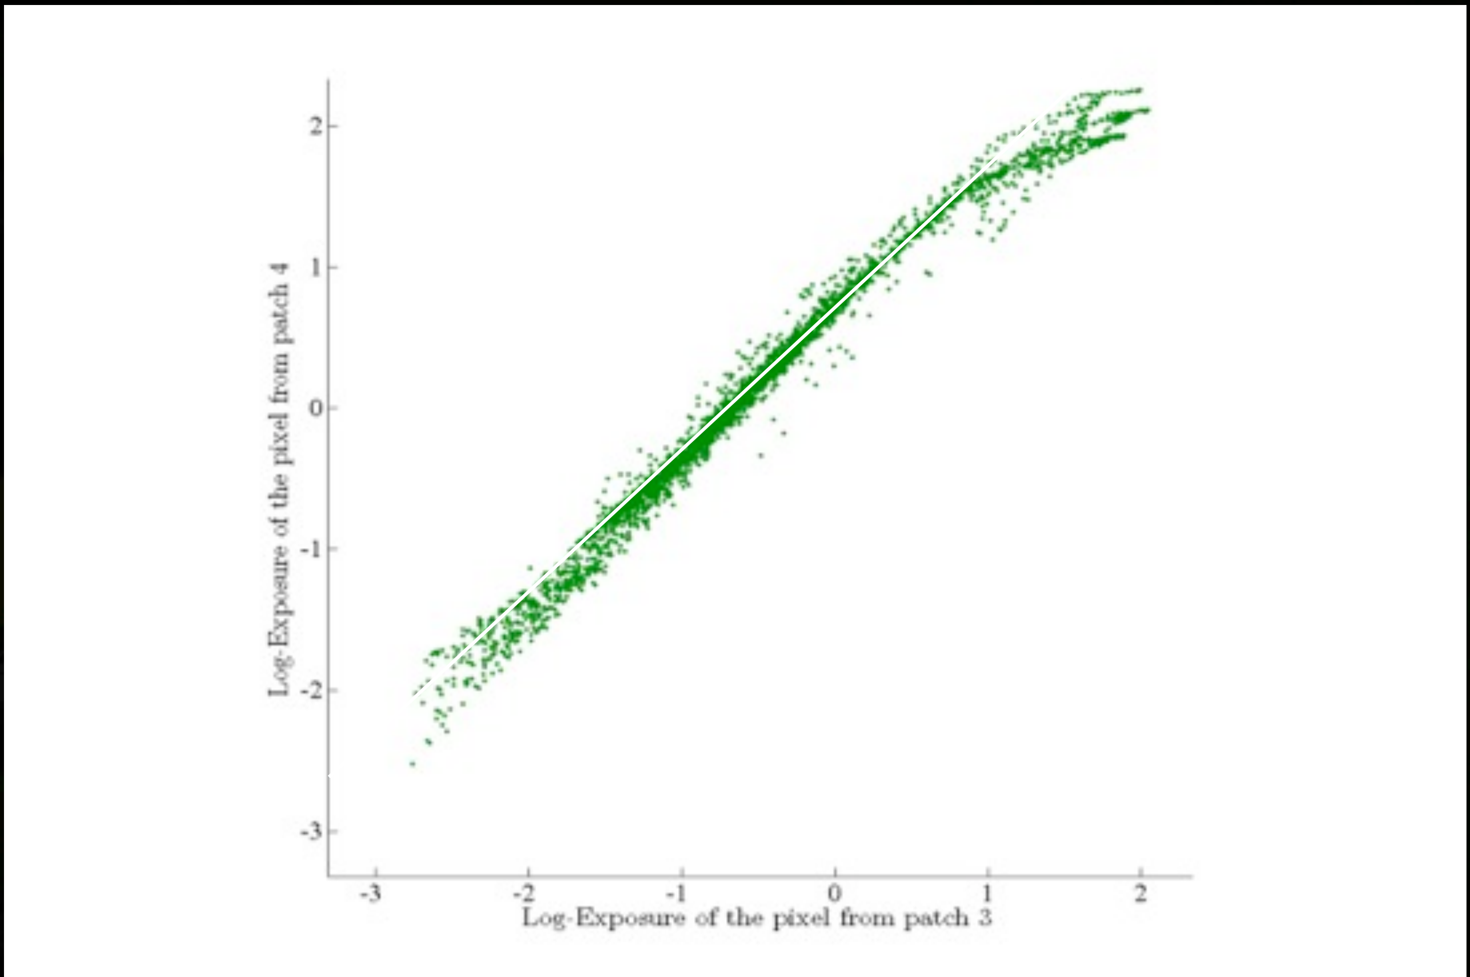

Log-Exposure of the pixel from patch 4
Log-Exposure of the pixel from patch 3
2
1
0
-1
-2
-3

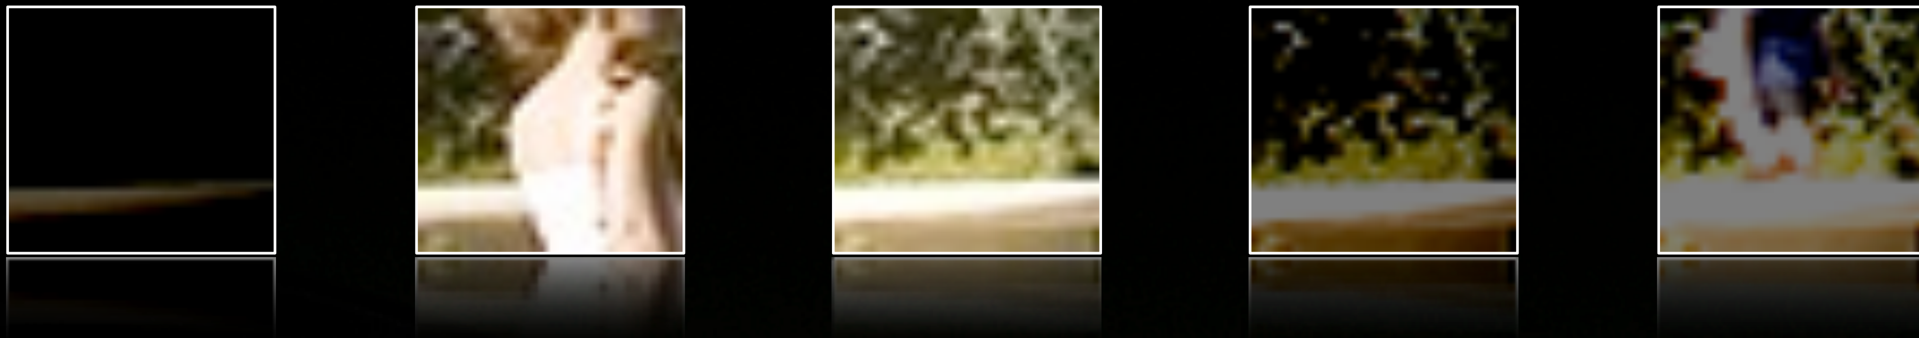





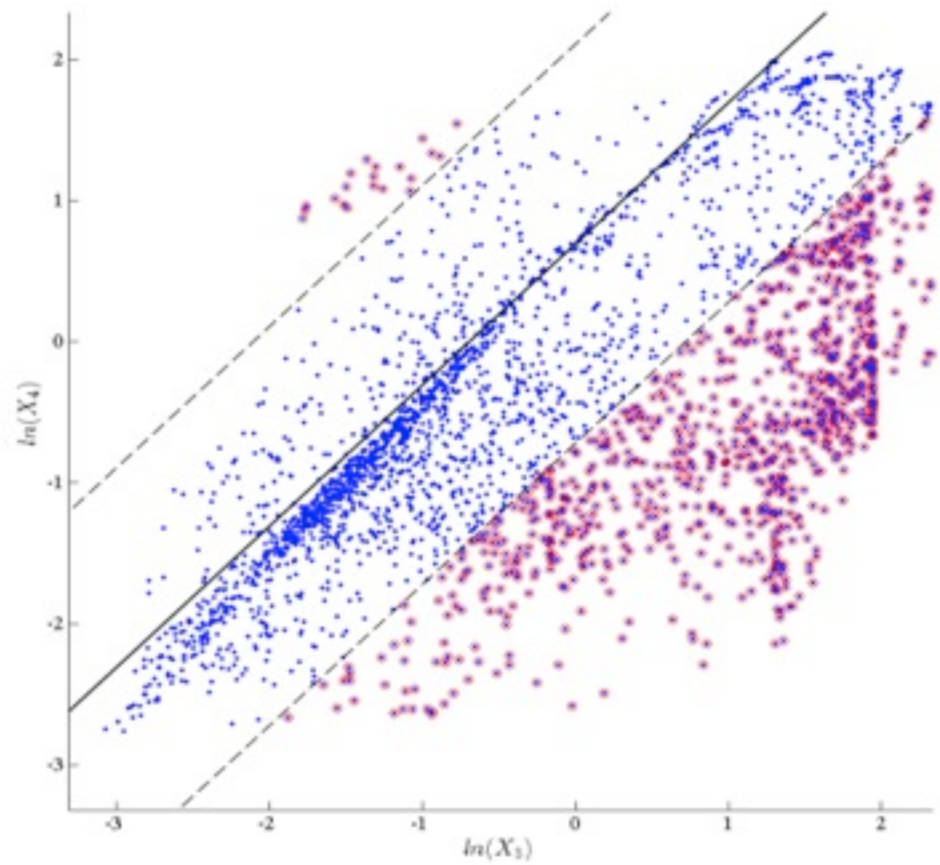

2
1
0
-1
-2
-3
$ln(X_4)$
-3
-2
-1
0
1
2
$ln(X_3)$

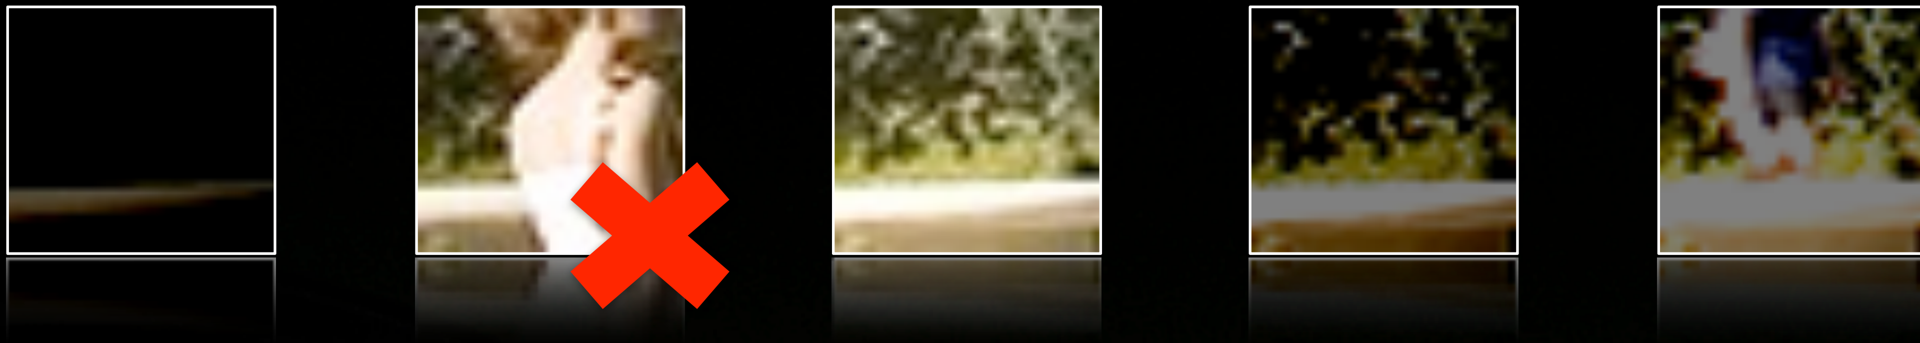

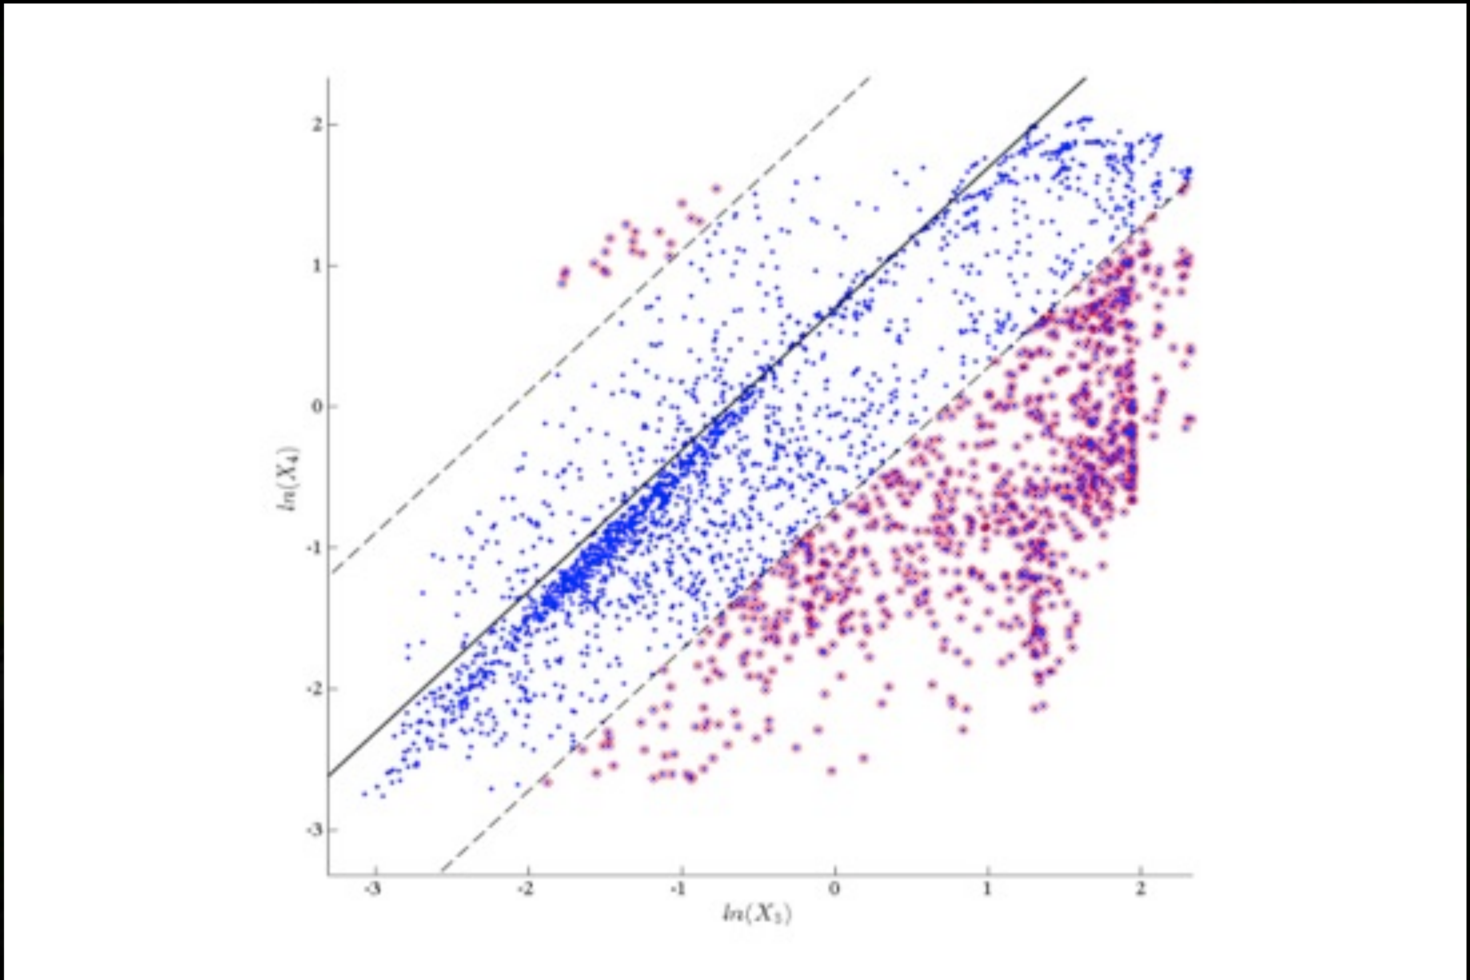

NVIDIA Research





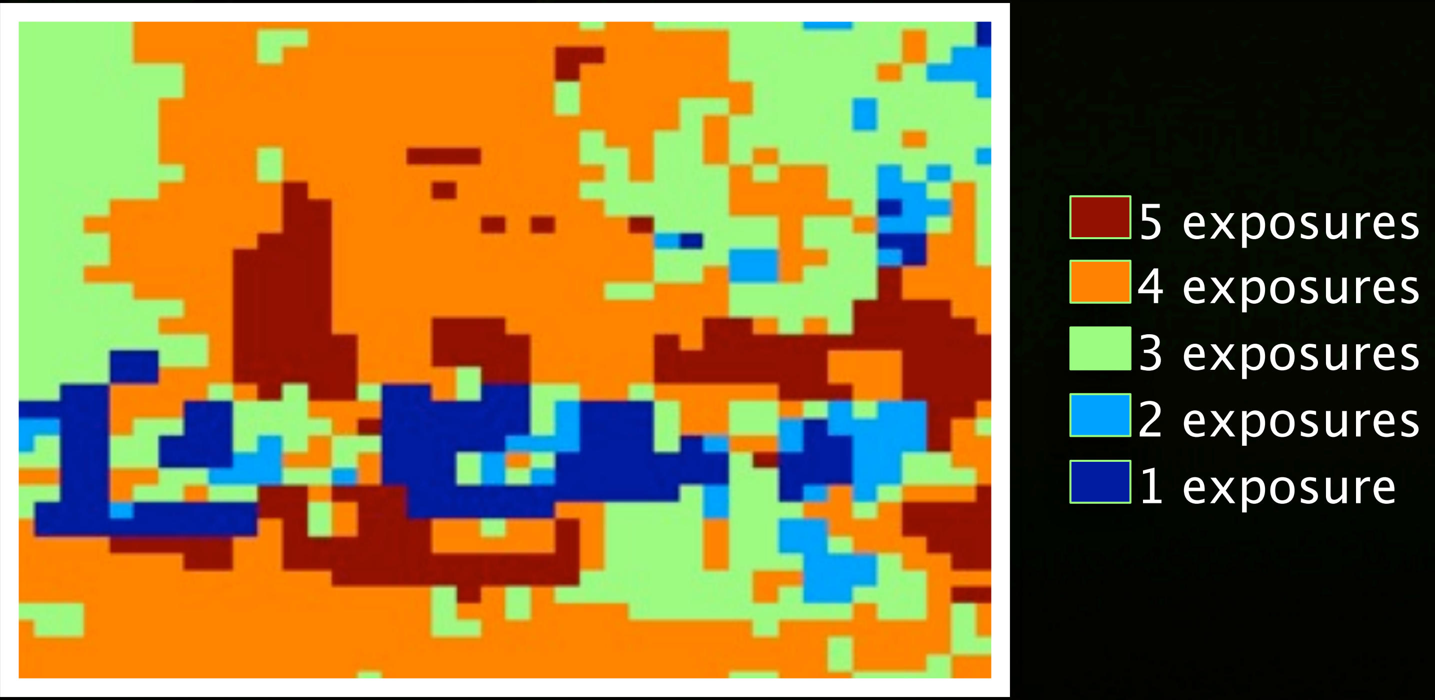
5 exposures
4 exposures
3 exposures
2 exposures
1 exposure

HDR video

- Kang et al. 2003
  - automatic exposure control
  - register neighboring frames (motion compensation)
  - tonemapping

Monday, March 5, 12

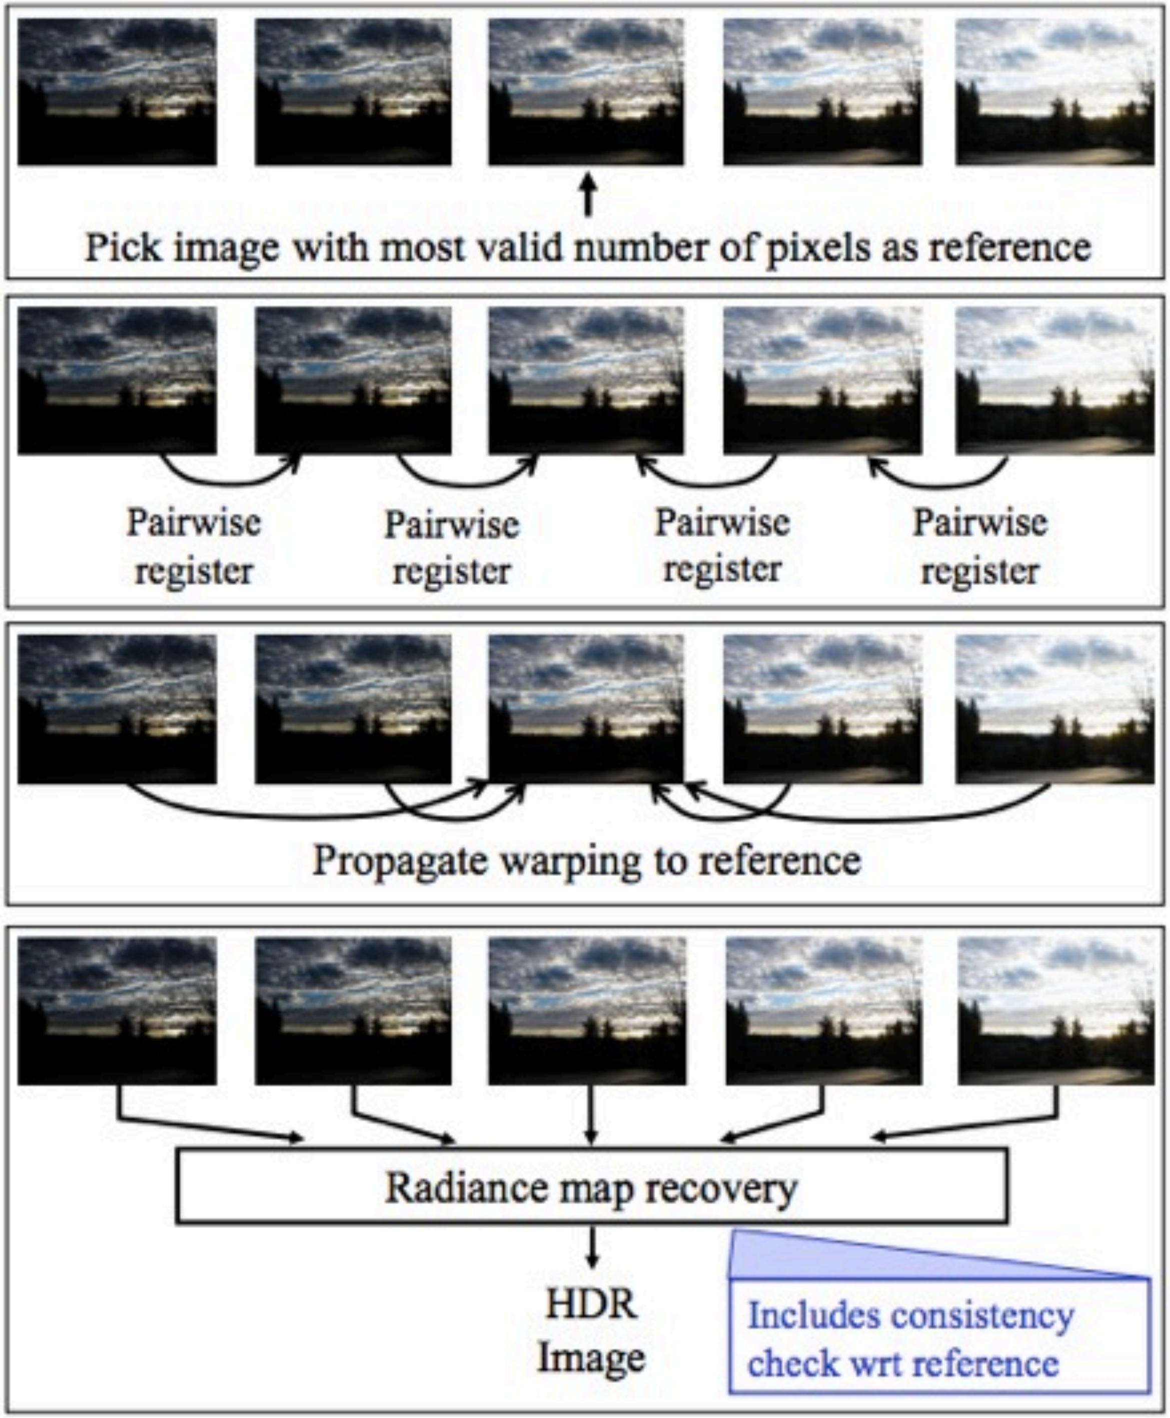

NVIDIA Research

A Versatile HDR Video Production System

ACM SIGGRAPH 2011

NVIDIA

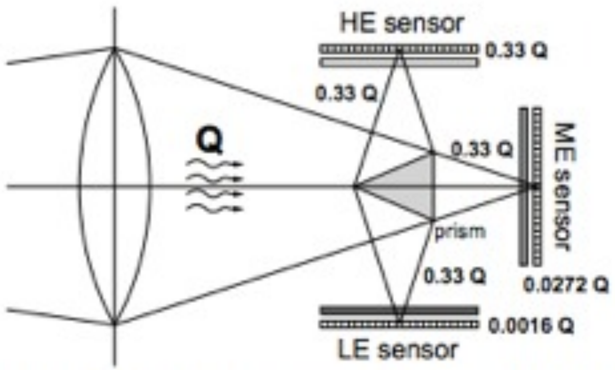

Figure 2: A traditional beamsplitting HDR optical system. *Here a beamsplitting prism breaks up the light into three parts, one for each sensor fitted with different filters. Designs that use absorptive filters like this one make inefficient use of light.*

A Versatile HDR Video Production System

ACM SIGGRAPH 2011

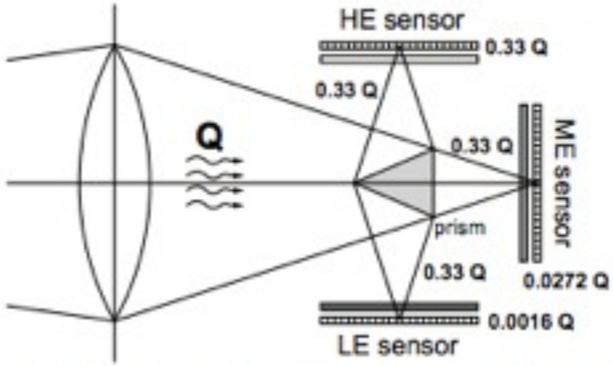

Figure 2: A traditional beamsplitting HDR optical system. *Here a beamsplitting prism breaks up the light into three parts, one for each sensor fitted with different filters. Designs that use absorptive filters like this one make inefficient use of light.*

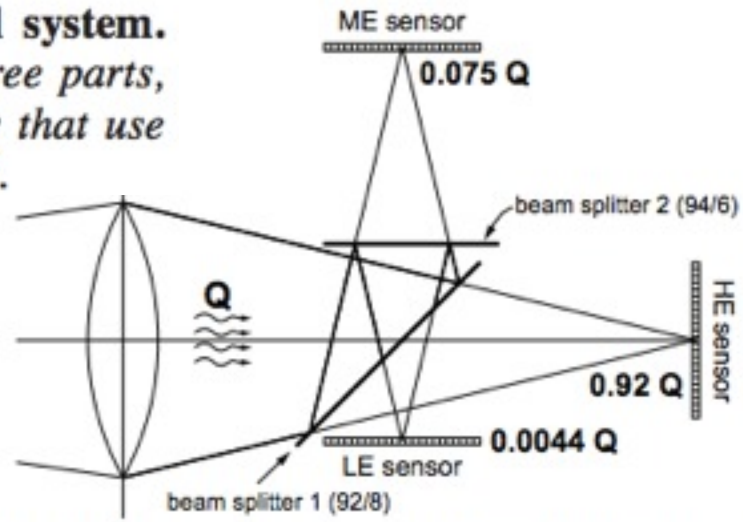

Figure 3: Illustration of our optical architecture. *We also use beamsplitters between the lens and sensors, but the key difference is that we re-use the optical path to improve our light efficiency. In the end, 99.96% of light entering the aperture arrives at the sensors. Light efficiency is important in all imaging applications.*

A Versatile HDR Video Production System

ACM SIGGRAPH 2011

NVIDIA Research

A Versatile HDR Video Production System

ACM SIGGRAPH 2011

NVIDIA Research

A Versatile HDR Video Production System

ACM SIGGRAPH 2011

NVIDIA Research

LDR image processing
= asking for trouble

- Physically accurate image processing requires floats
  - 8bit or 16bit ints are not enough
    - inherent quantization between operations
    - e.g., applying gamma to brighten or darken maps levels that were separate to the same levels, can't separate any more
  - saturation
    - at the high end
    - can't deal with really bright pixels (direct light sources)
  - non-linearity
    - for better encoding, but not for physical processing

Image processing example: motion blur

- Processing LDR gamma-corrected images (sRGB) yields artifacts

blurred LDR

blurred HDR

blurred real photo