High-Dynamic-Range Photography + Tone mapping

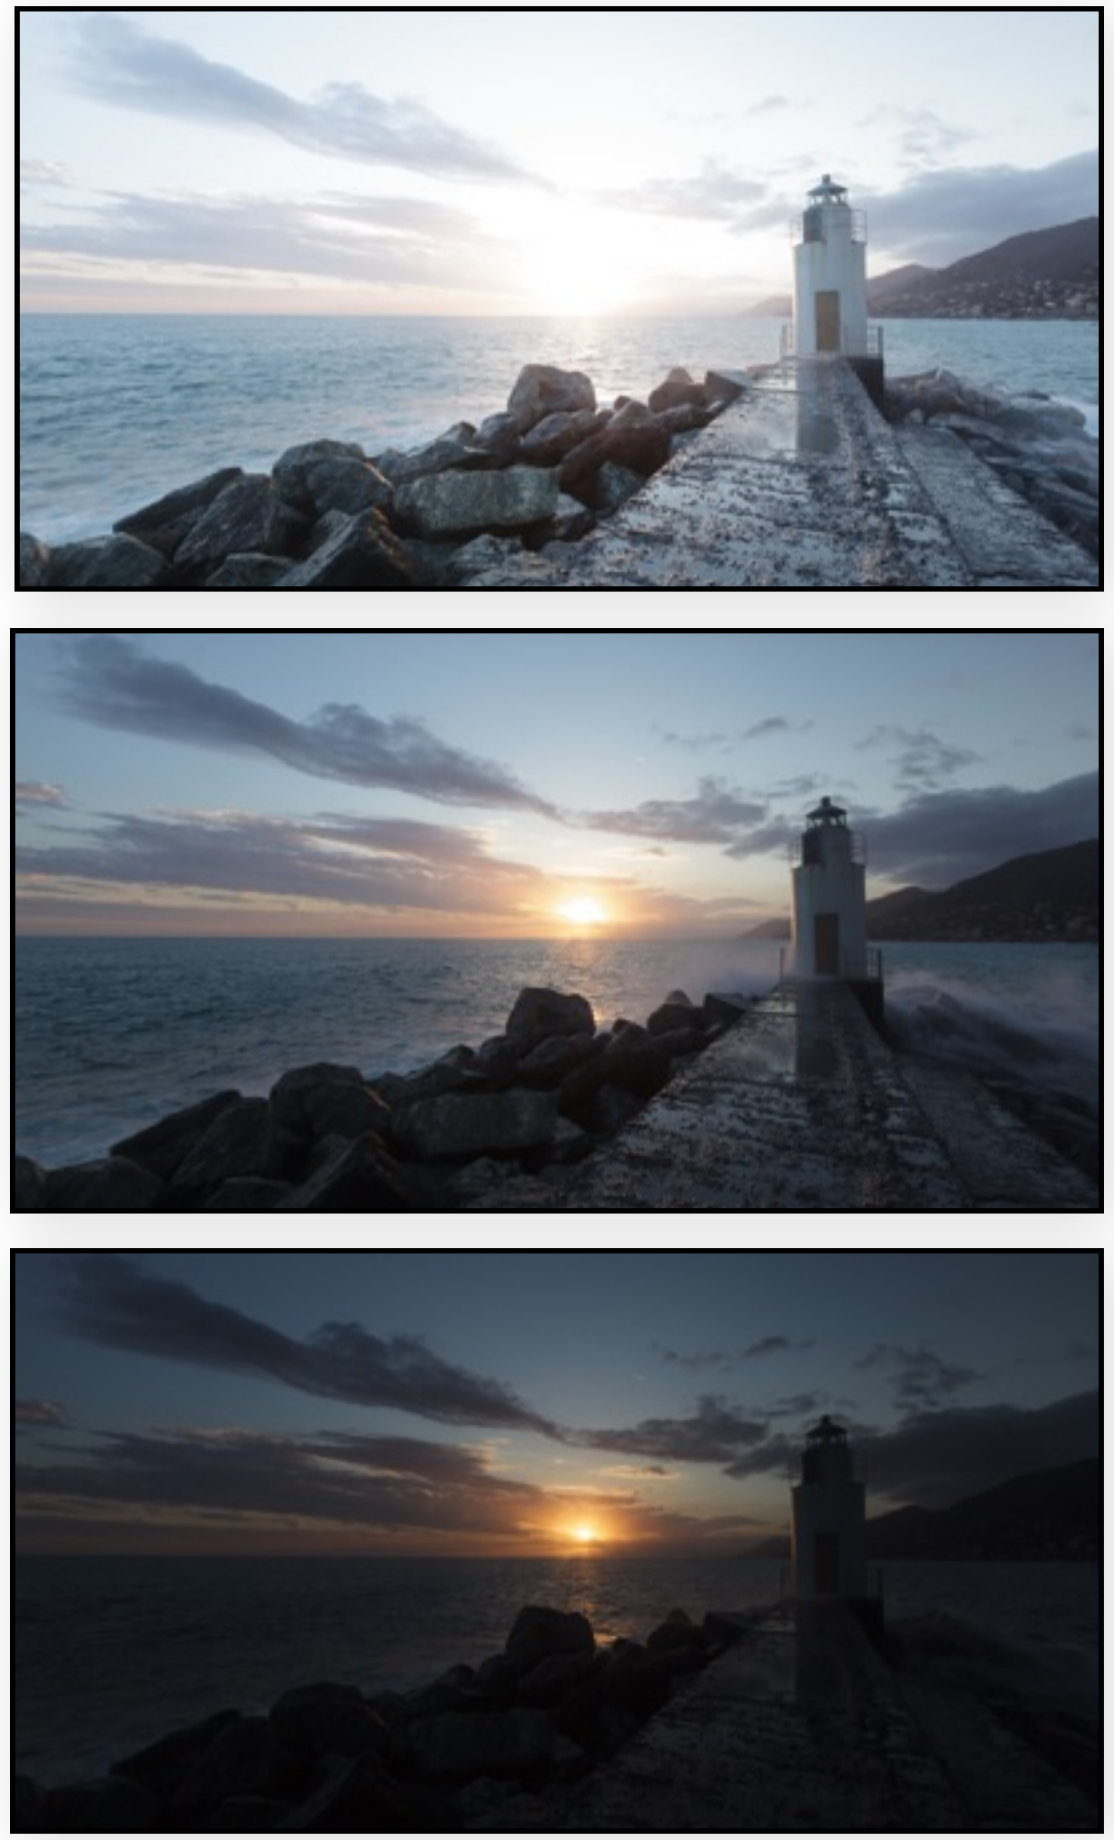

Dartmouth

Wojciech Jarosz
wojciech.k.jarosz@dartmouth.edu

CS 89.15/189.5, Fall 2015

Light, exposure and dynamic range

Exposure: how bright is the scene overall?

Dynamic range: contrast in the scene

- ratio of brightest to darkest intensity

The dynamic range challenge

Examples

Inside is too dark

Outside is too bright

Sun overexposed

Foreground too dark

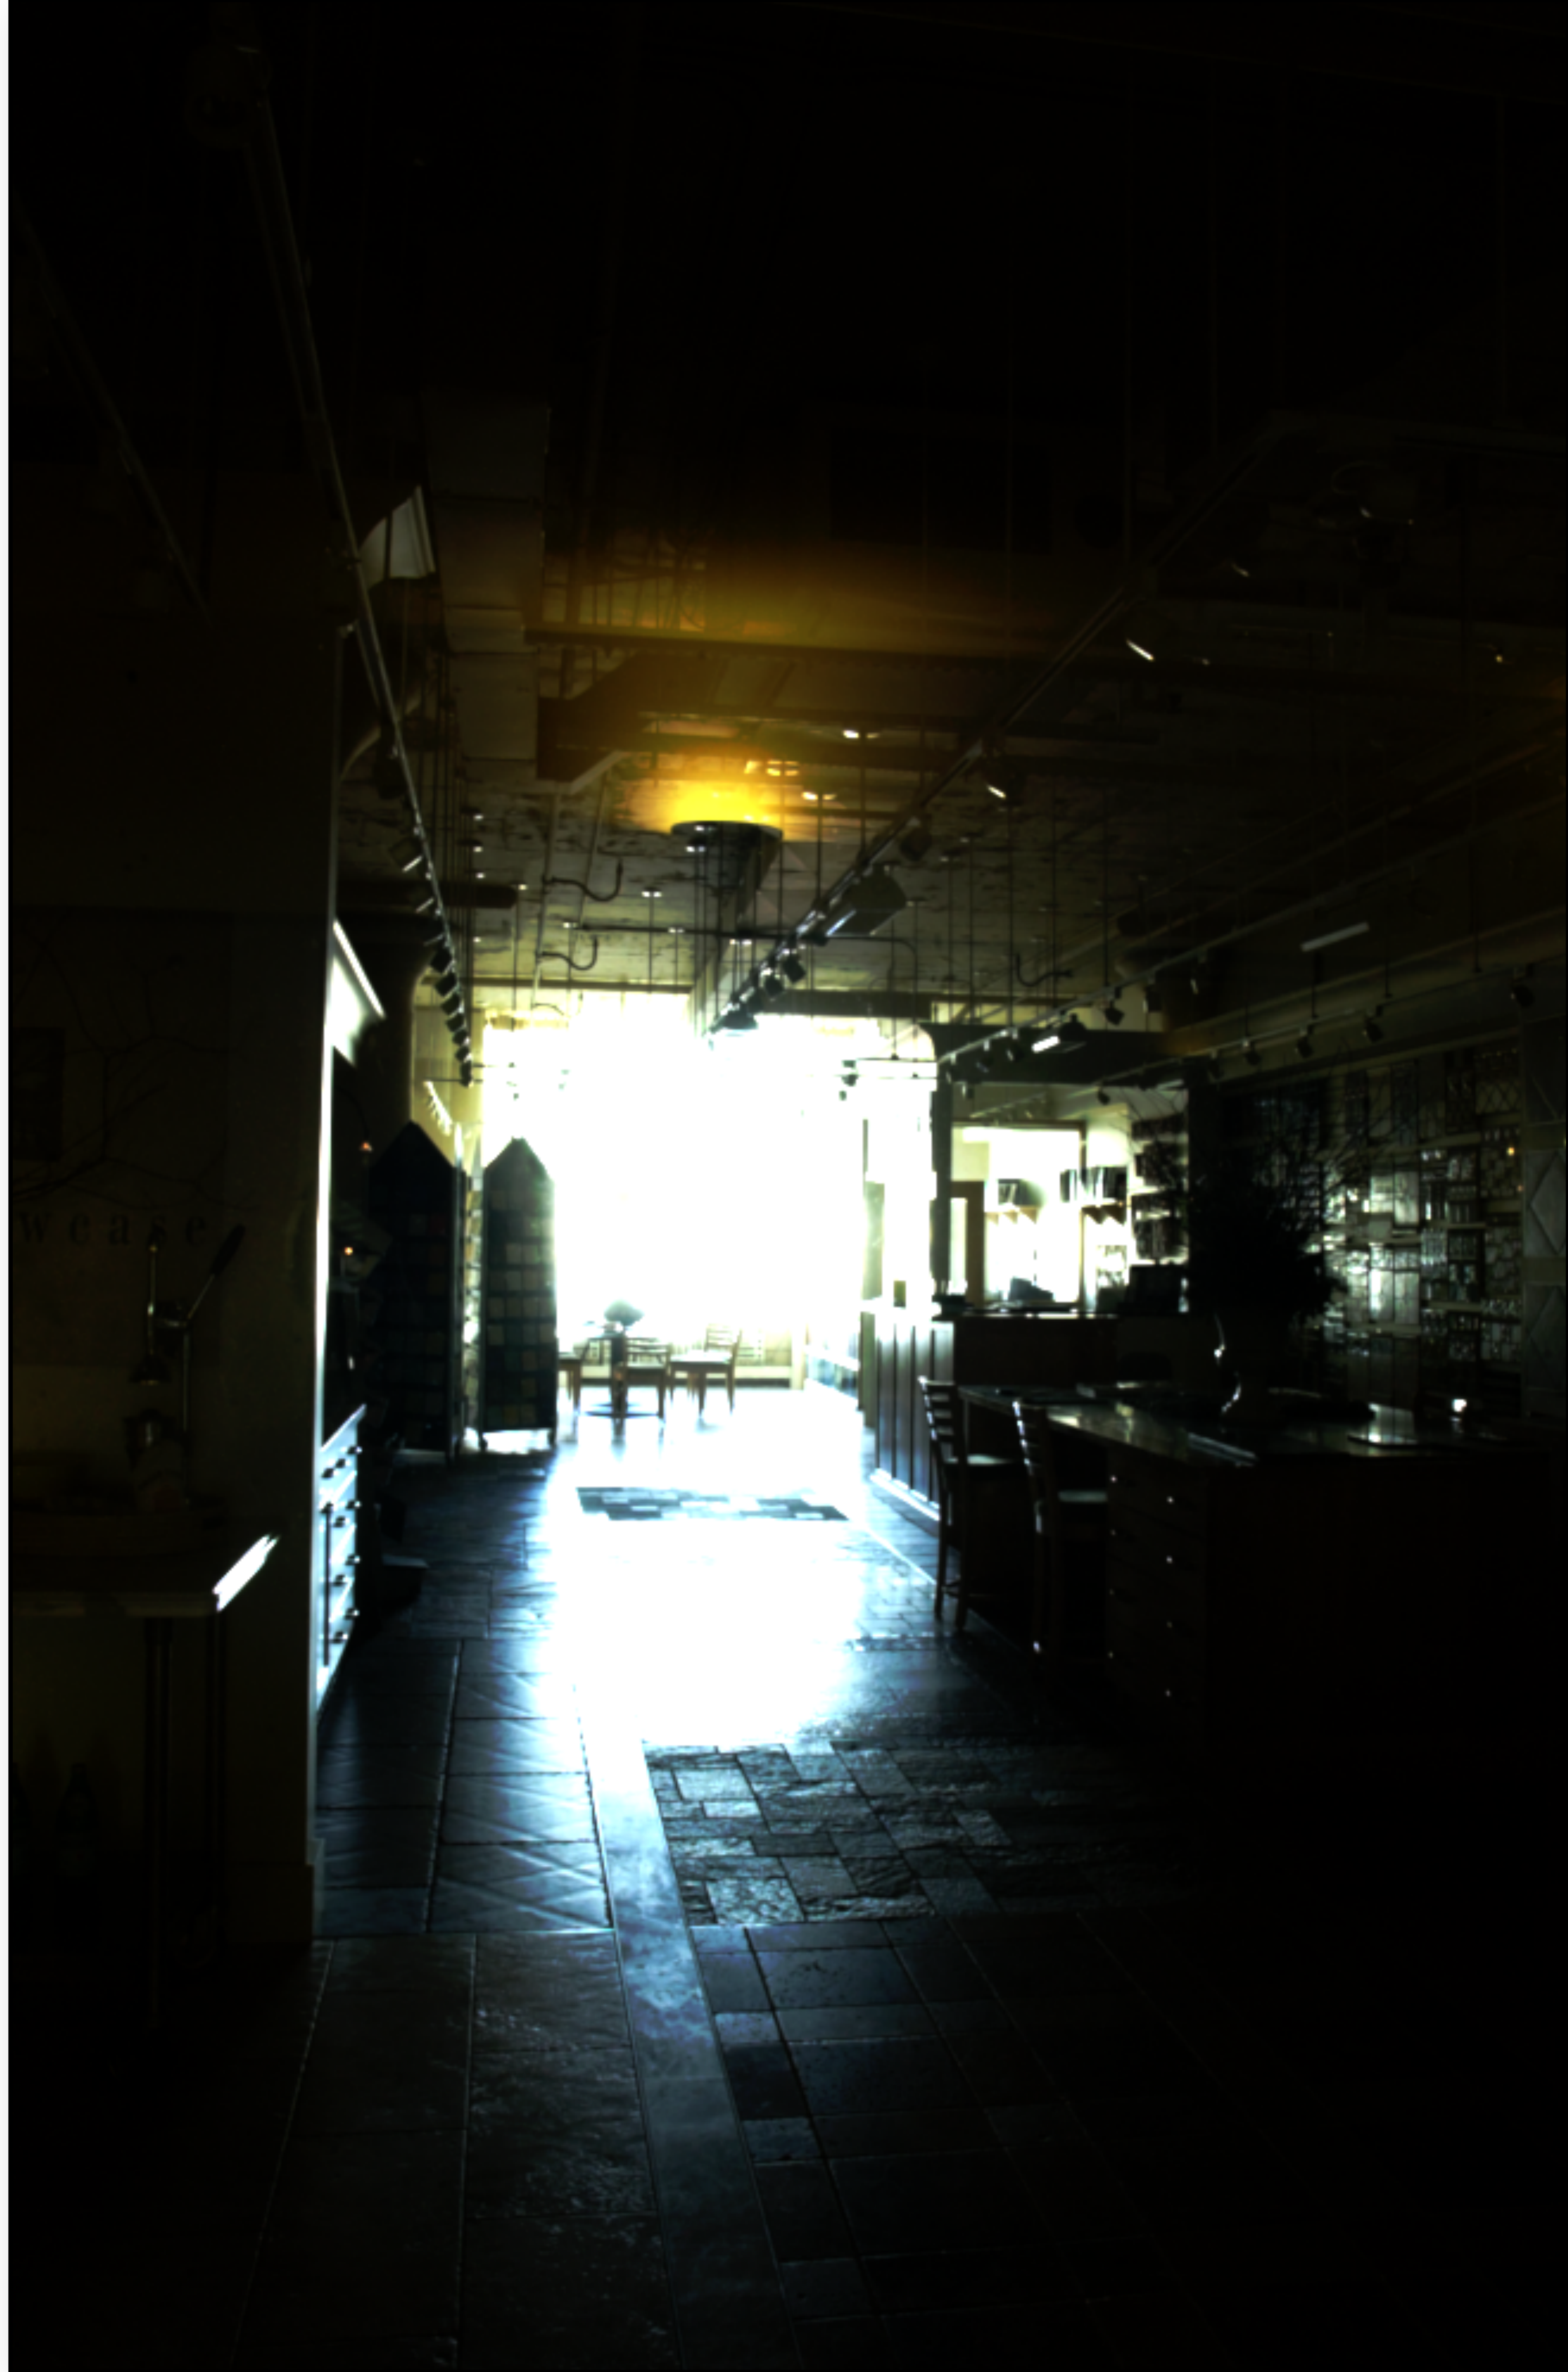

After a slide by Frédo Durand

Low Dynamic Range (LDR)

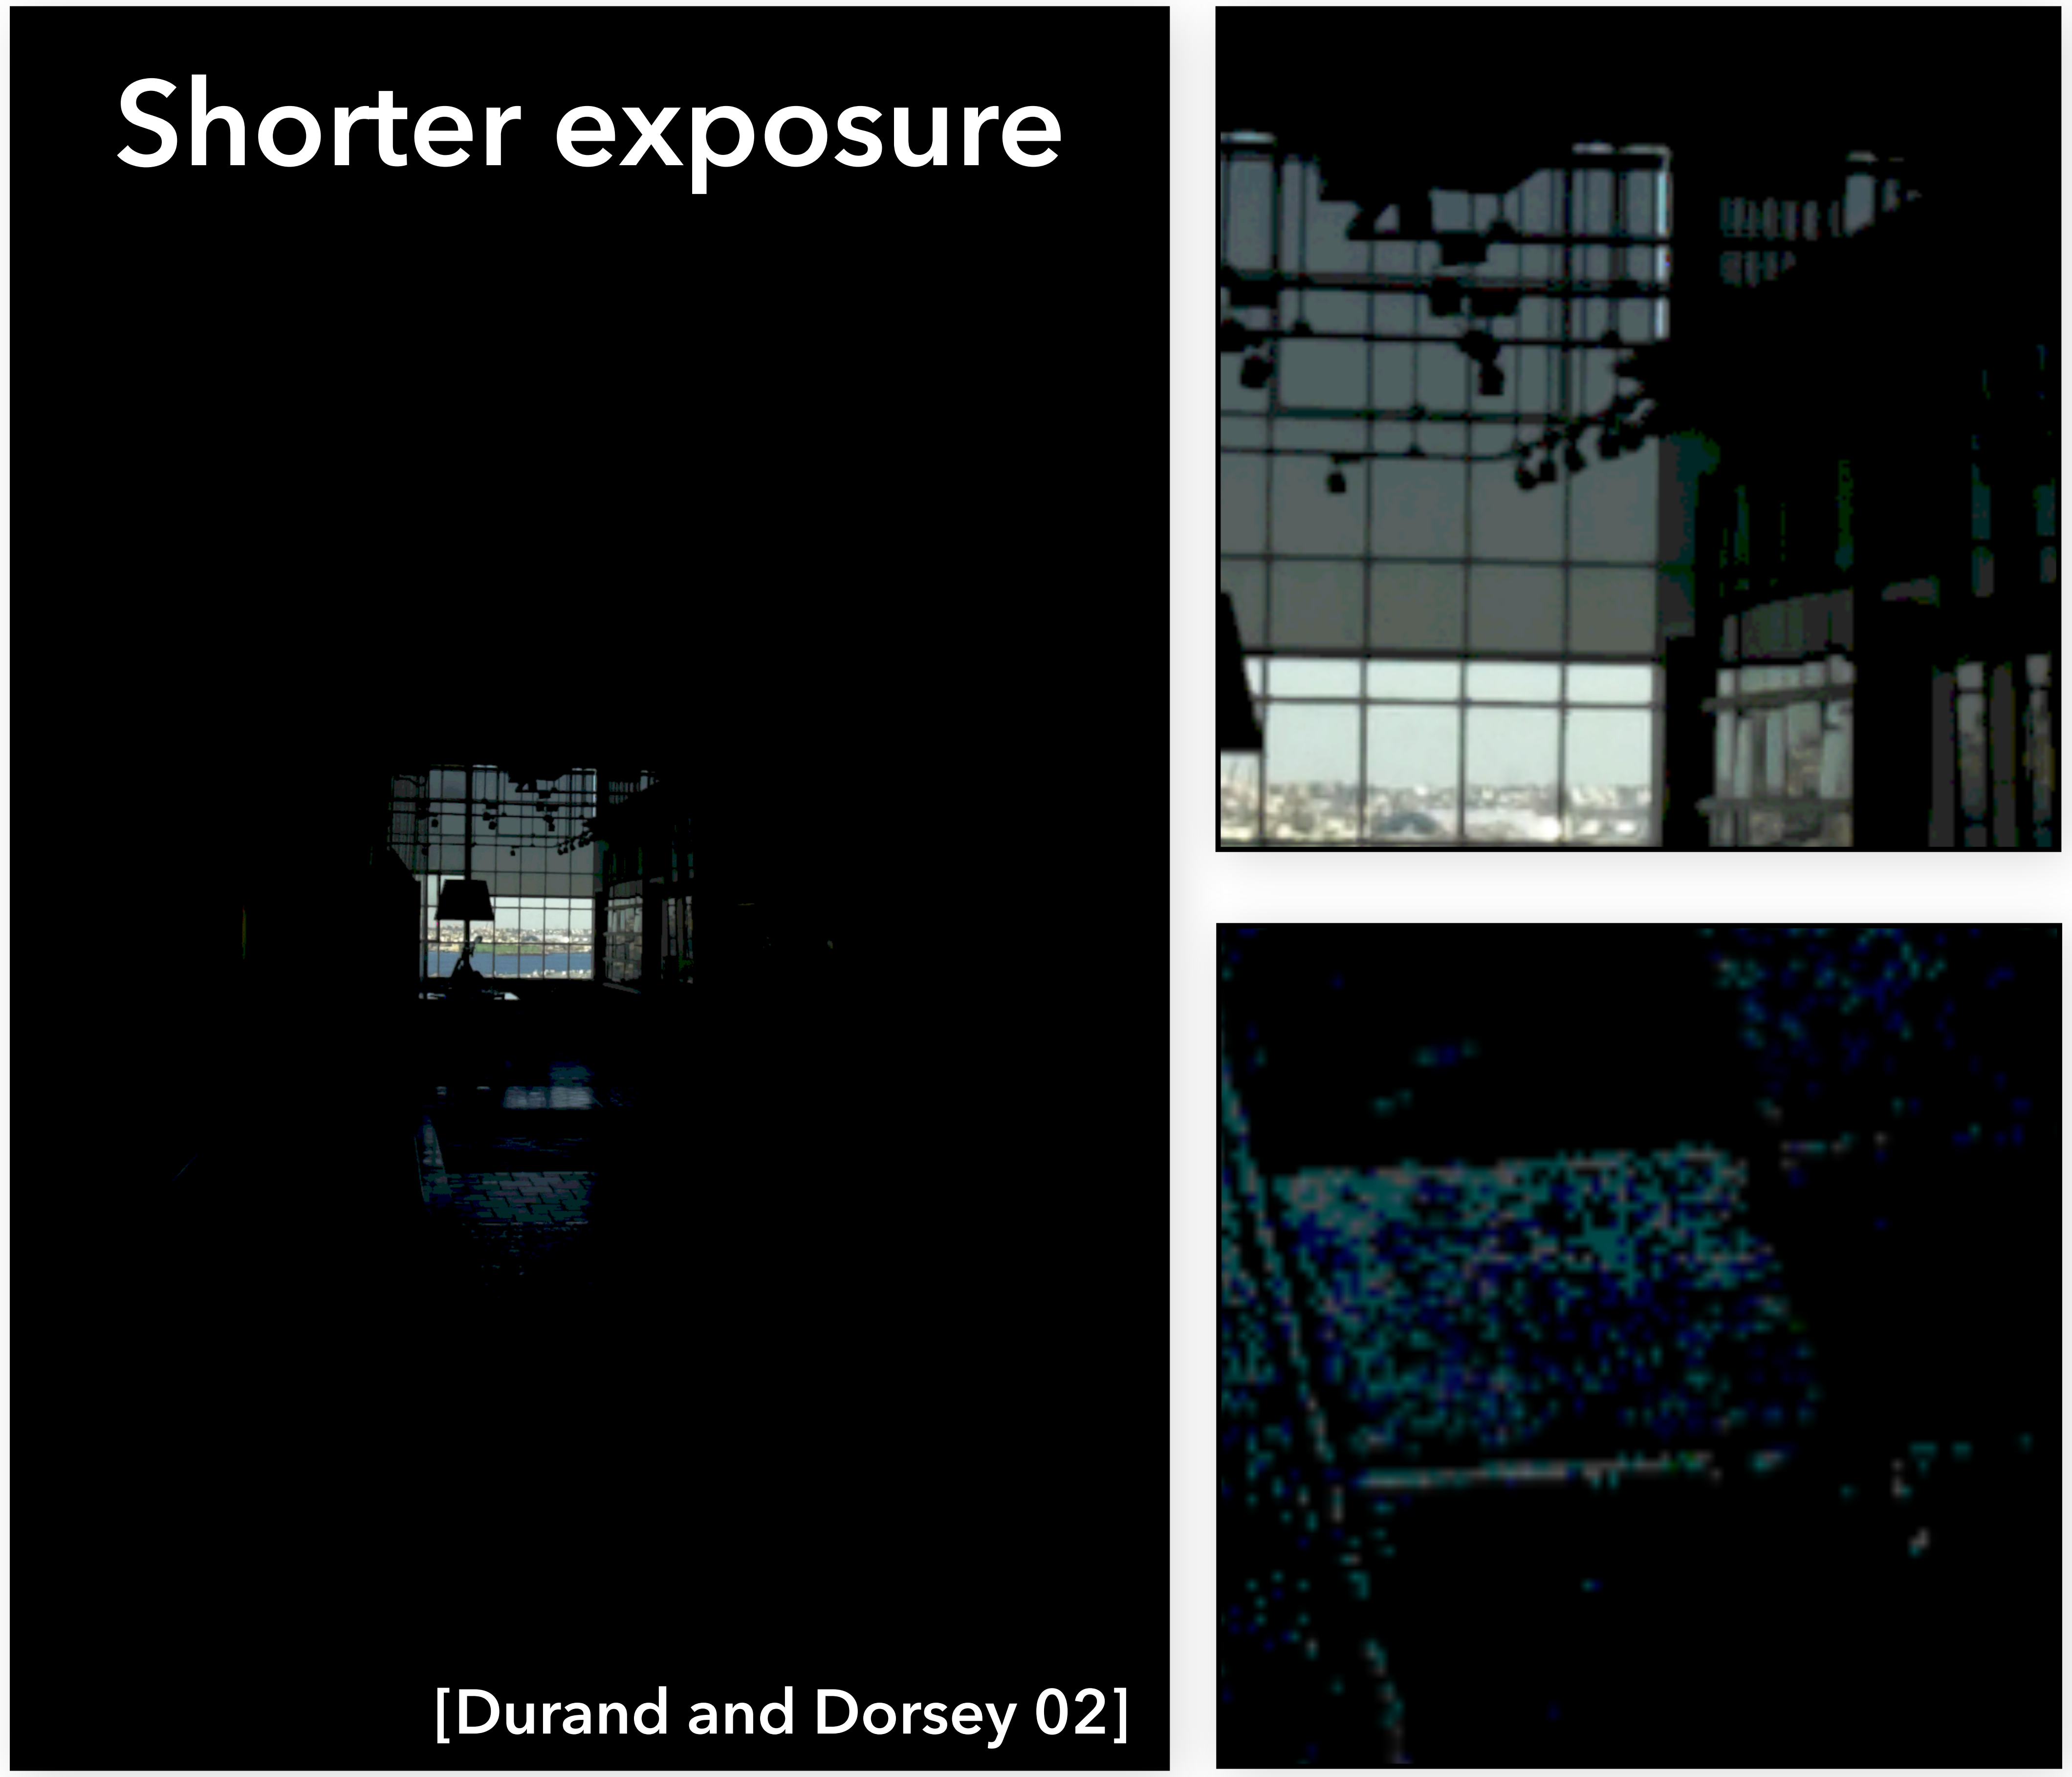

✓ detail in shadows
✗ clipped highlights

✓ detail in highlights
✗ noisy/clipped shadows

Real world dynamic range

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

Often 1 : 100,000 in a scene

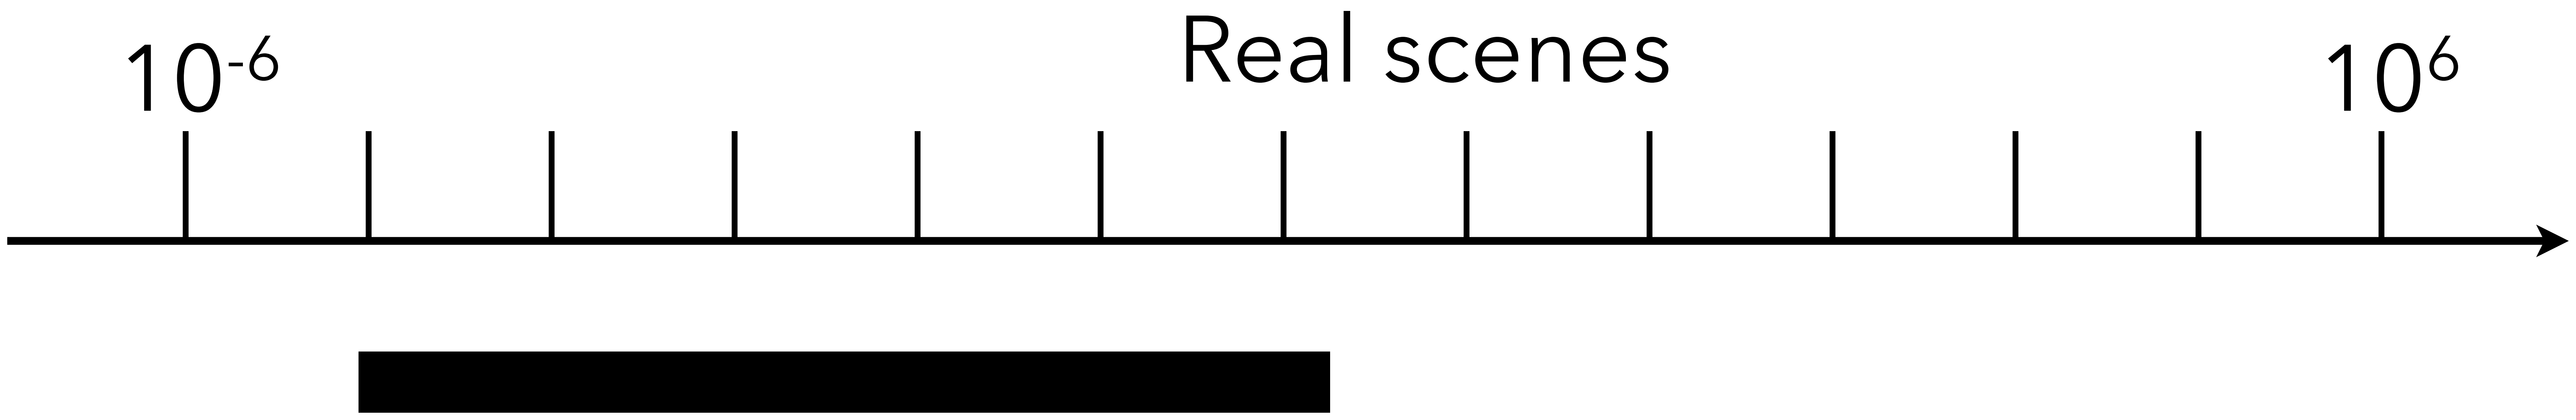

After a slide by Frédo Durand

The real world is high dynamic range.

Slide from Paul Debevec

Dynamic Range Examples

Example	Contrast ratio / Dynamic Range	Bits
Photographic print (higher for glossy paper)	1:10	4
Artist's paints	1:20	5
Slide film	1:200	8
Typical display / JPEG image	1:255	8
Negative film	1:500	9
High-end DSLR	1:8192	13
Human eye	1:1,048,576	20
Range of observable scenes	1:8,589,934,592	33

*approximate and debatable

Problem 1: Record the information

The range of illumination levels that we encounter is 10–12 orders of magnitude

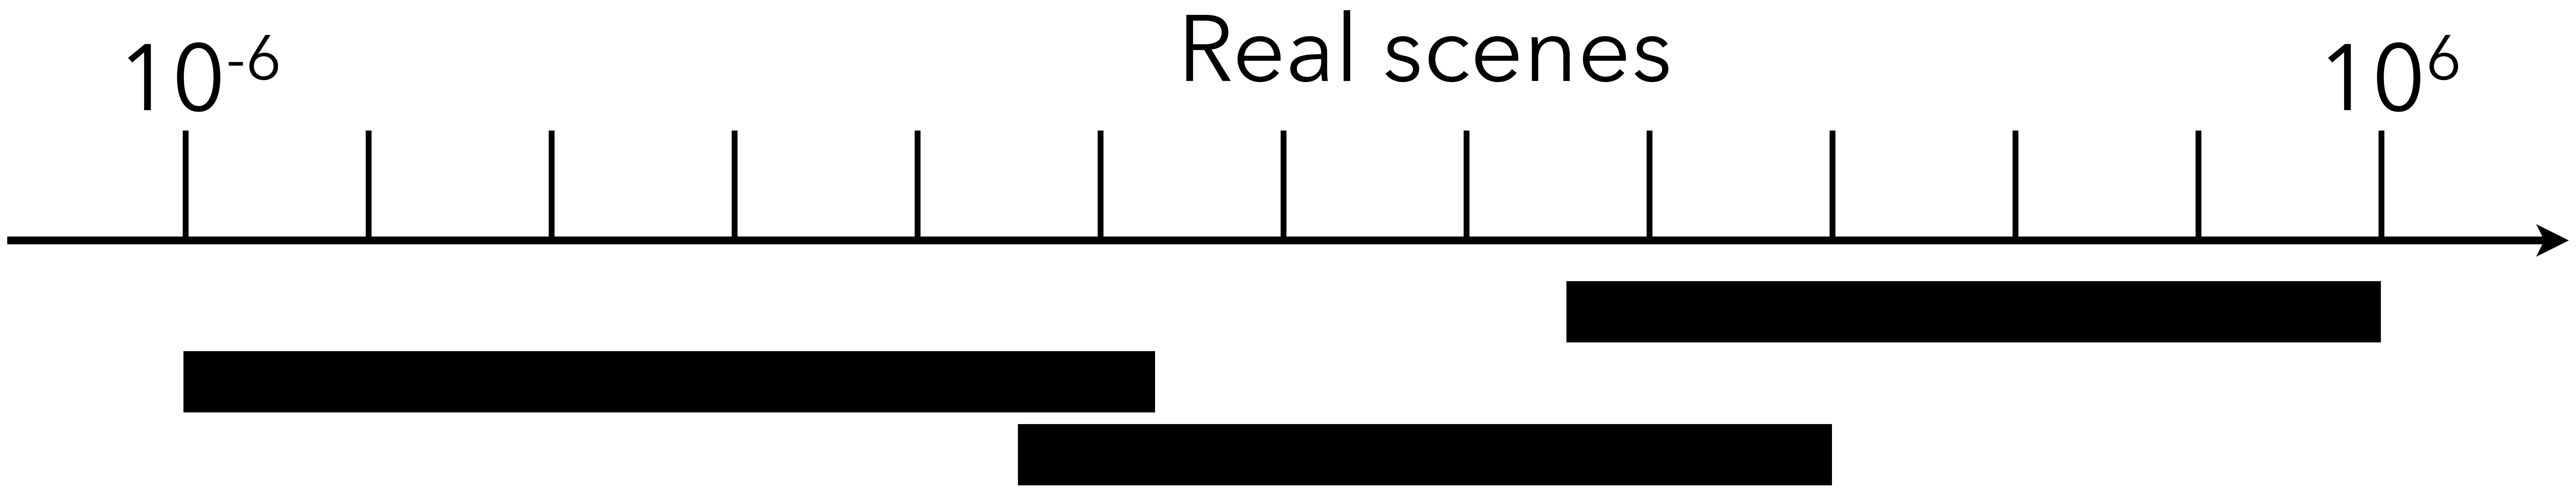

Film/sensors can record 2-3 orders of magnitude

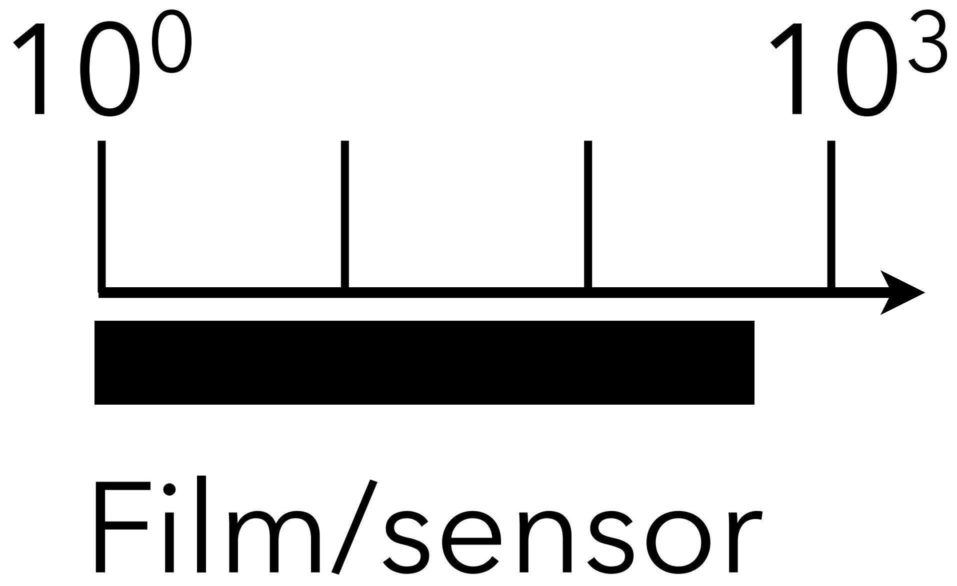

After a slide by Frédo Durand

Problem 2: Display the information

Match limited contrast of the medium while preserving details

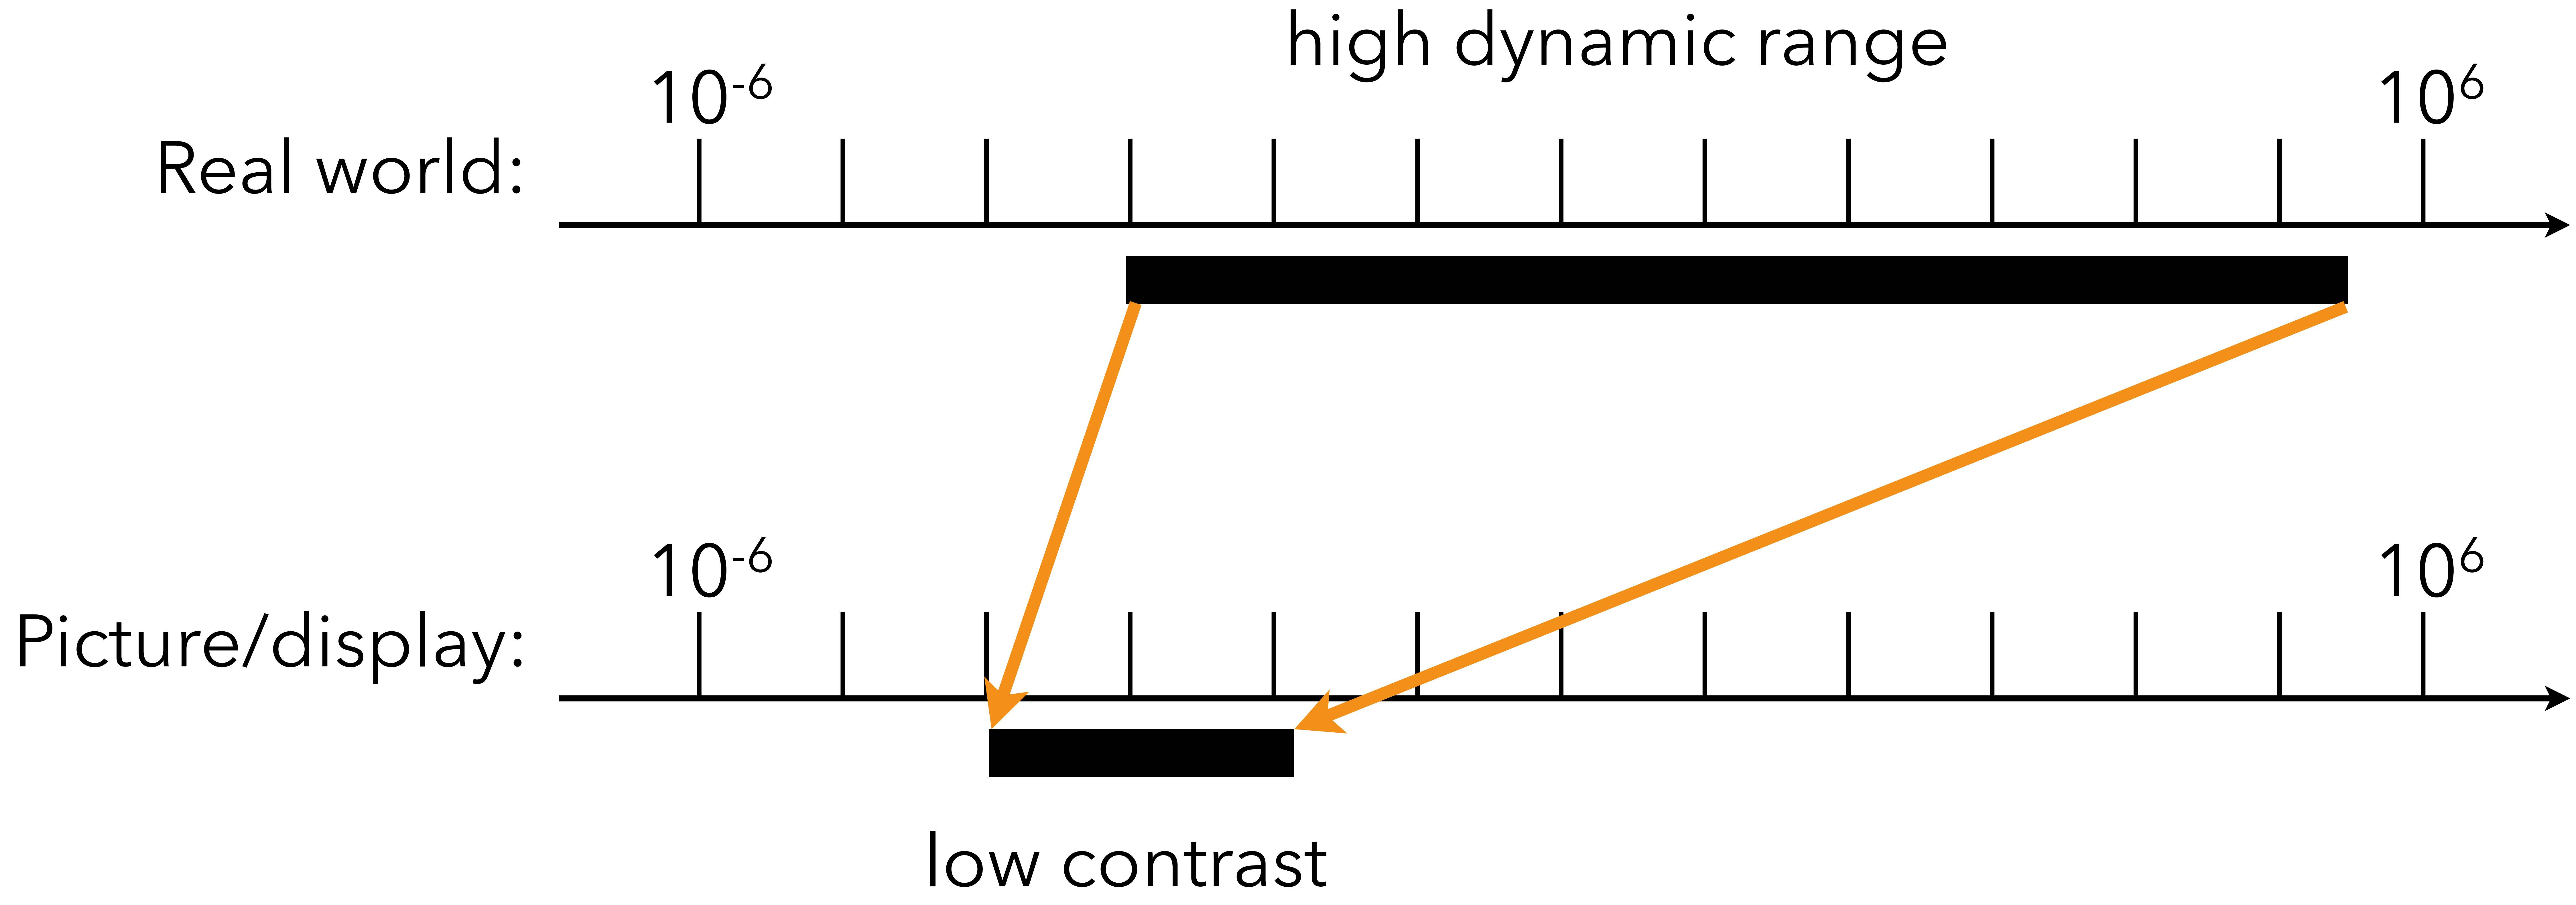

After a slide by Frédo Durand

Without HDR & tone mapping

[Wojciech Jarosz]

[Wojciech Jarosz]

HDR today

HDR Off

HDR On

Application: Motion blur

Simulated Motion

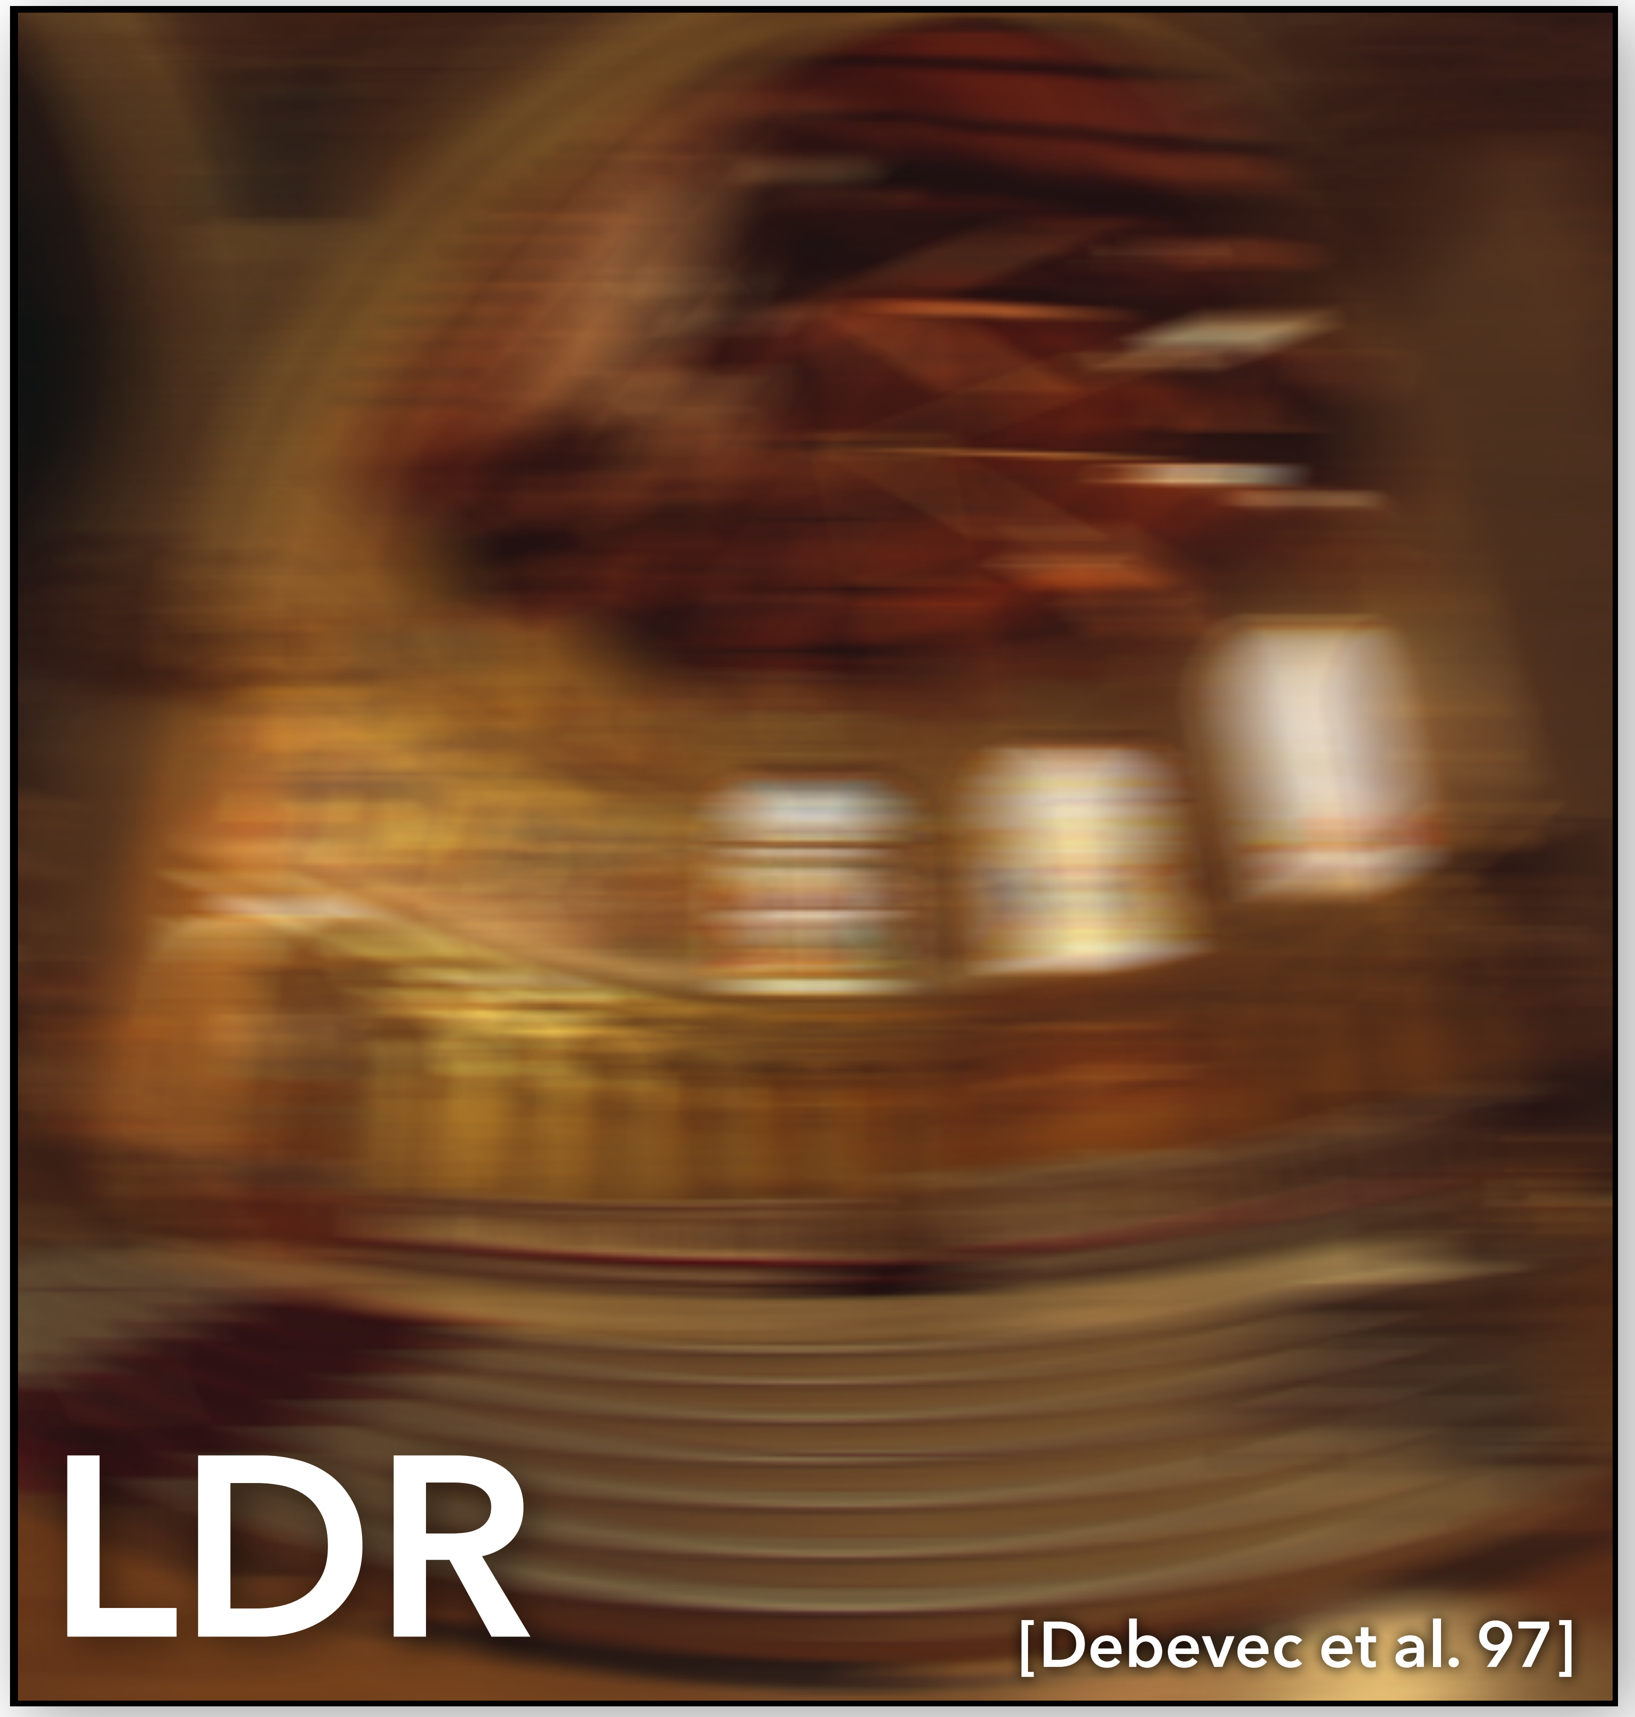

Simulated Motion Blur

Actual Motion blur

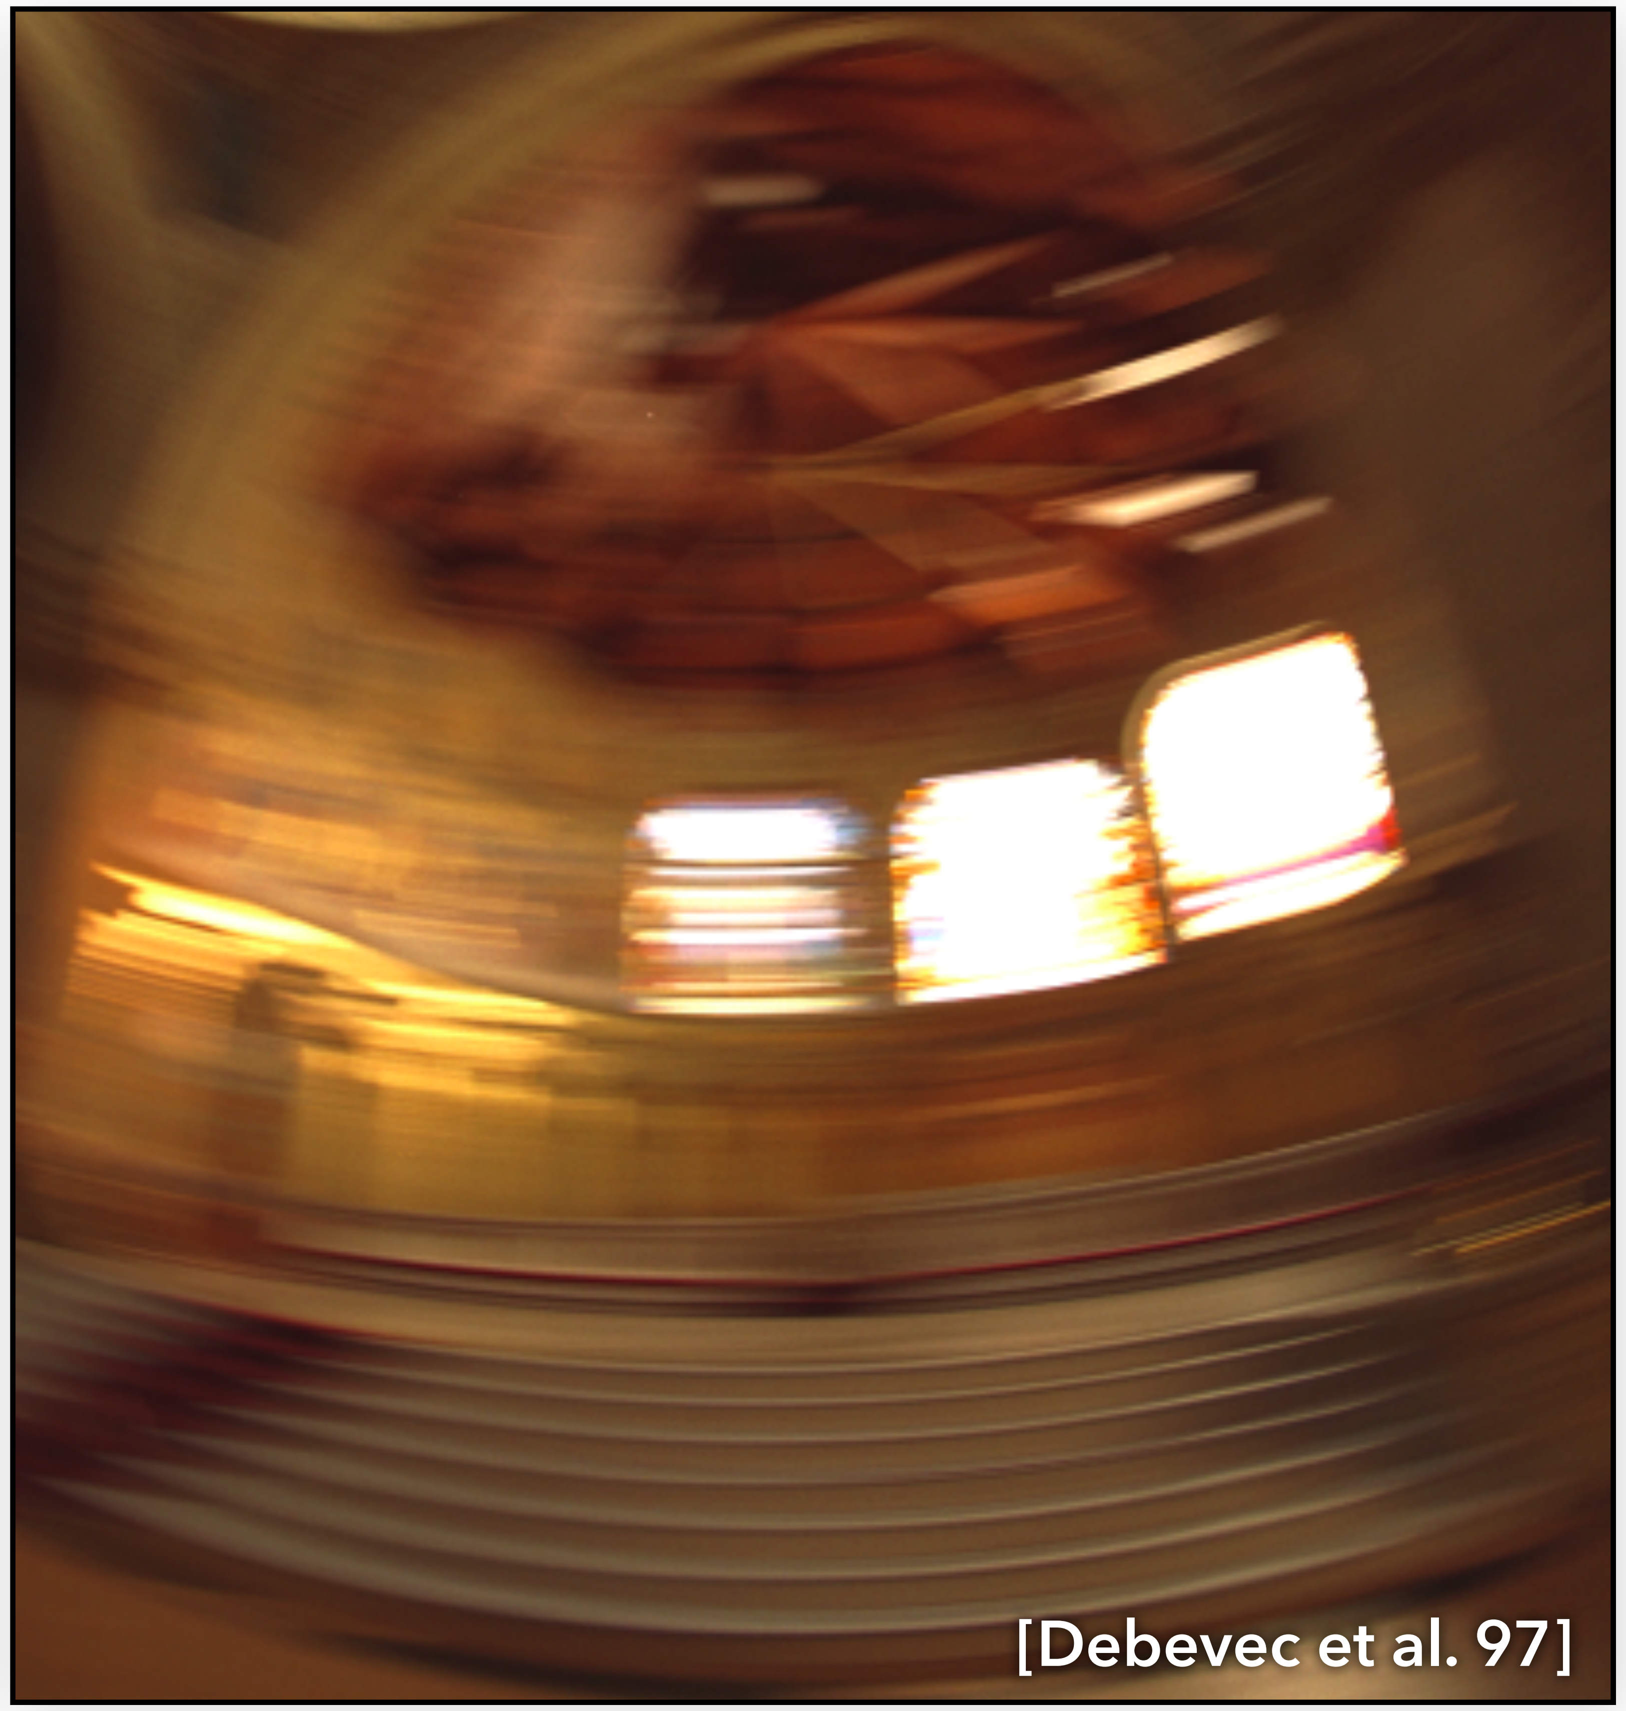

Application: Inserting Synthetic Objects

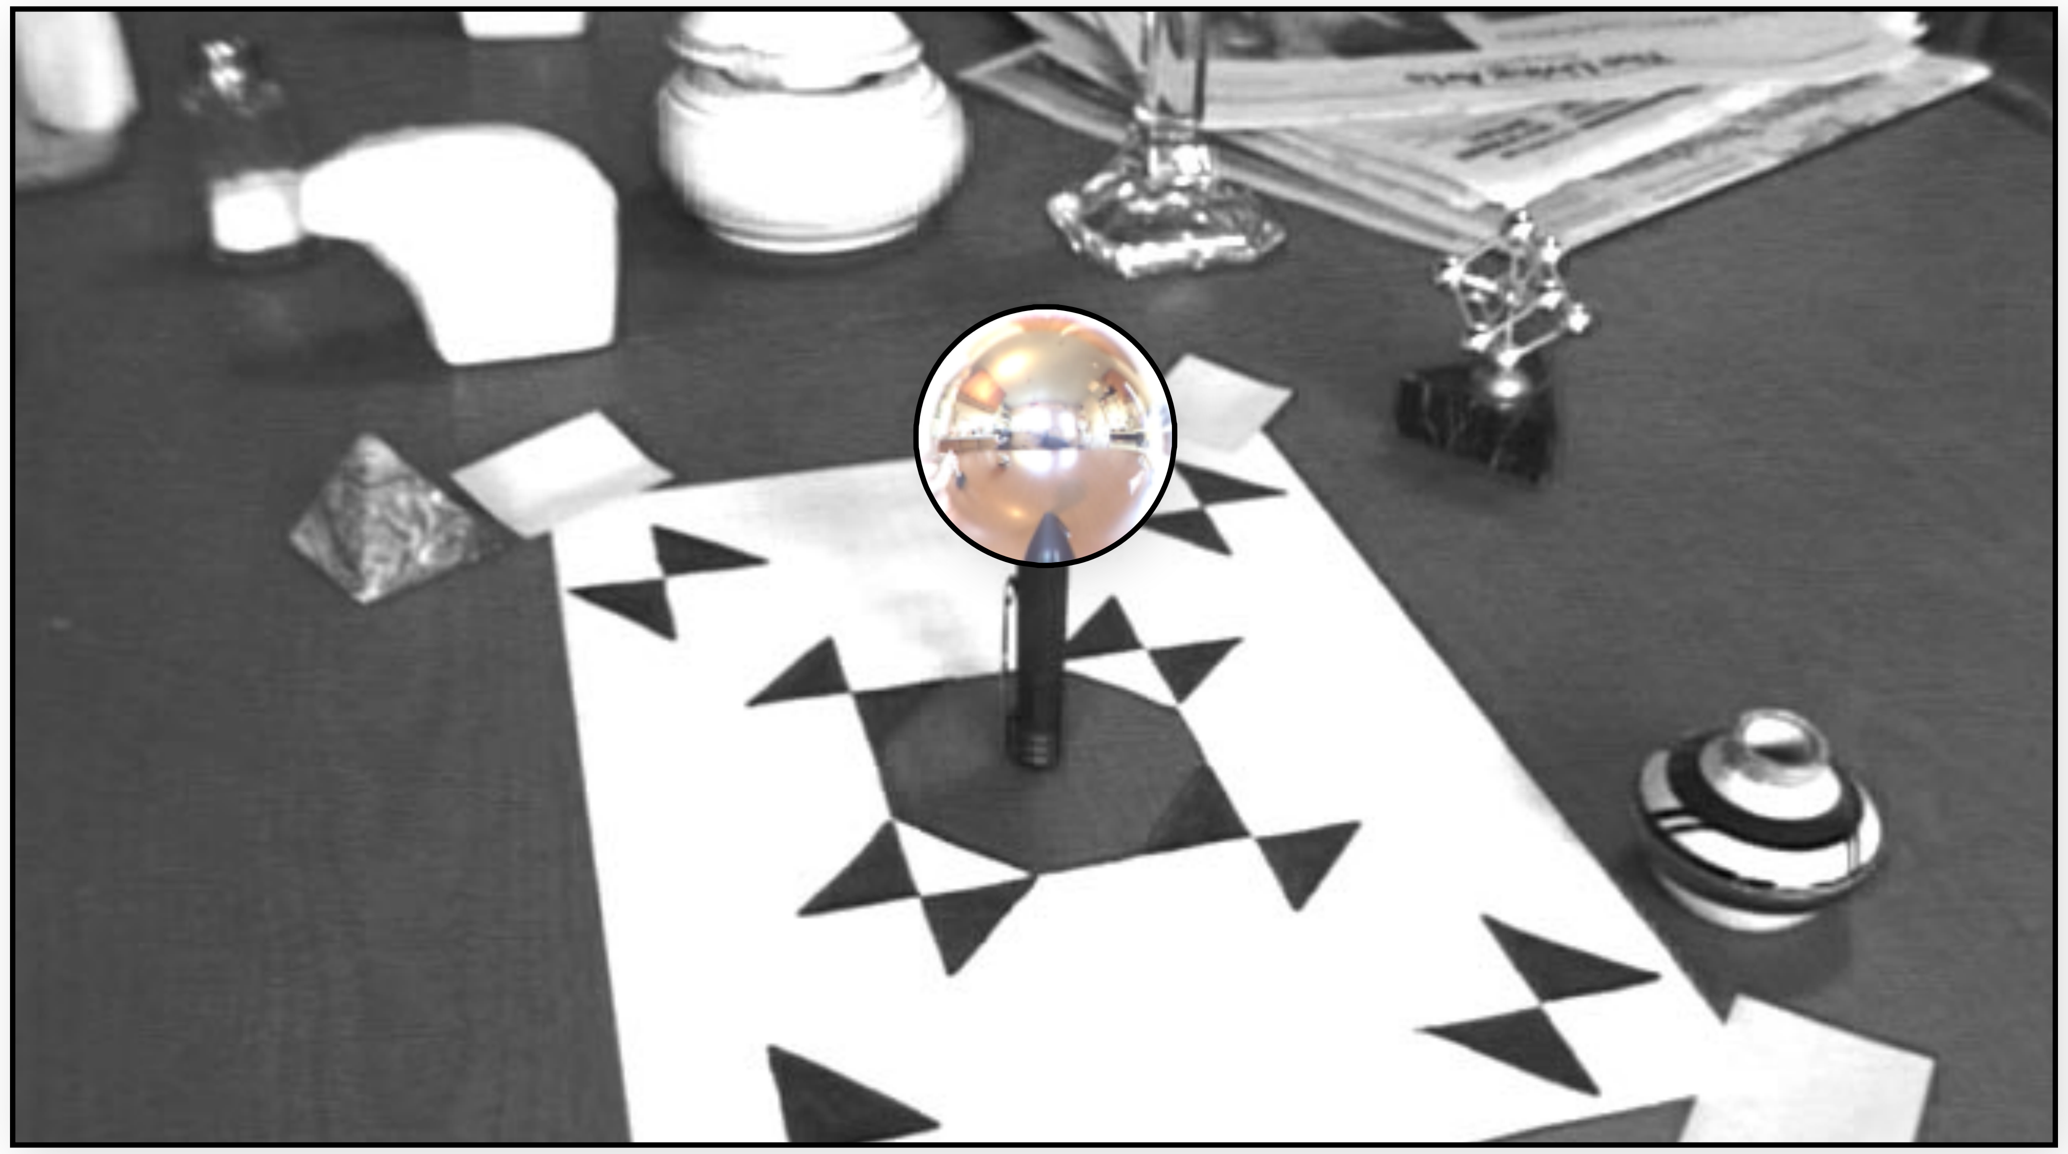

Application: Inserting Synthetic Objects

Image Based Environment Lighting

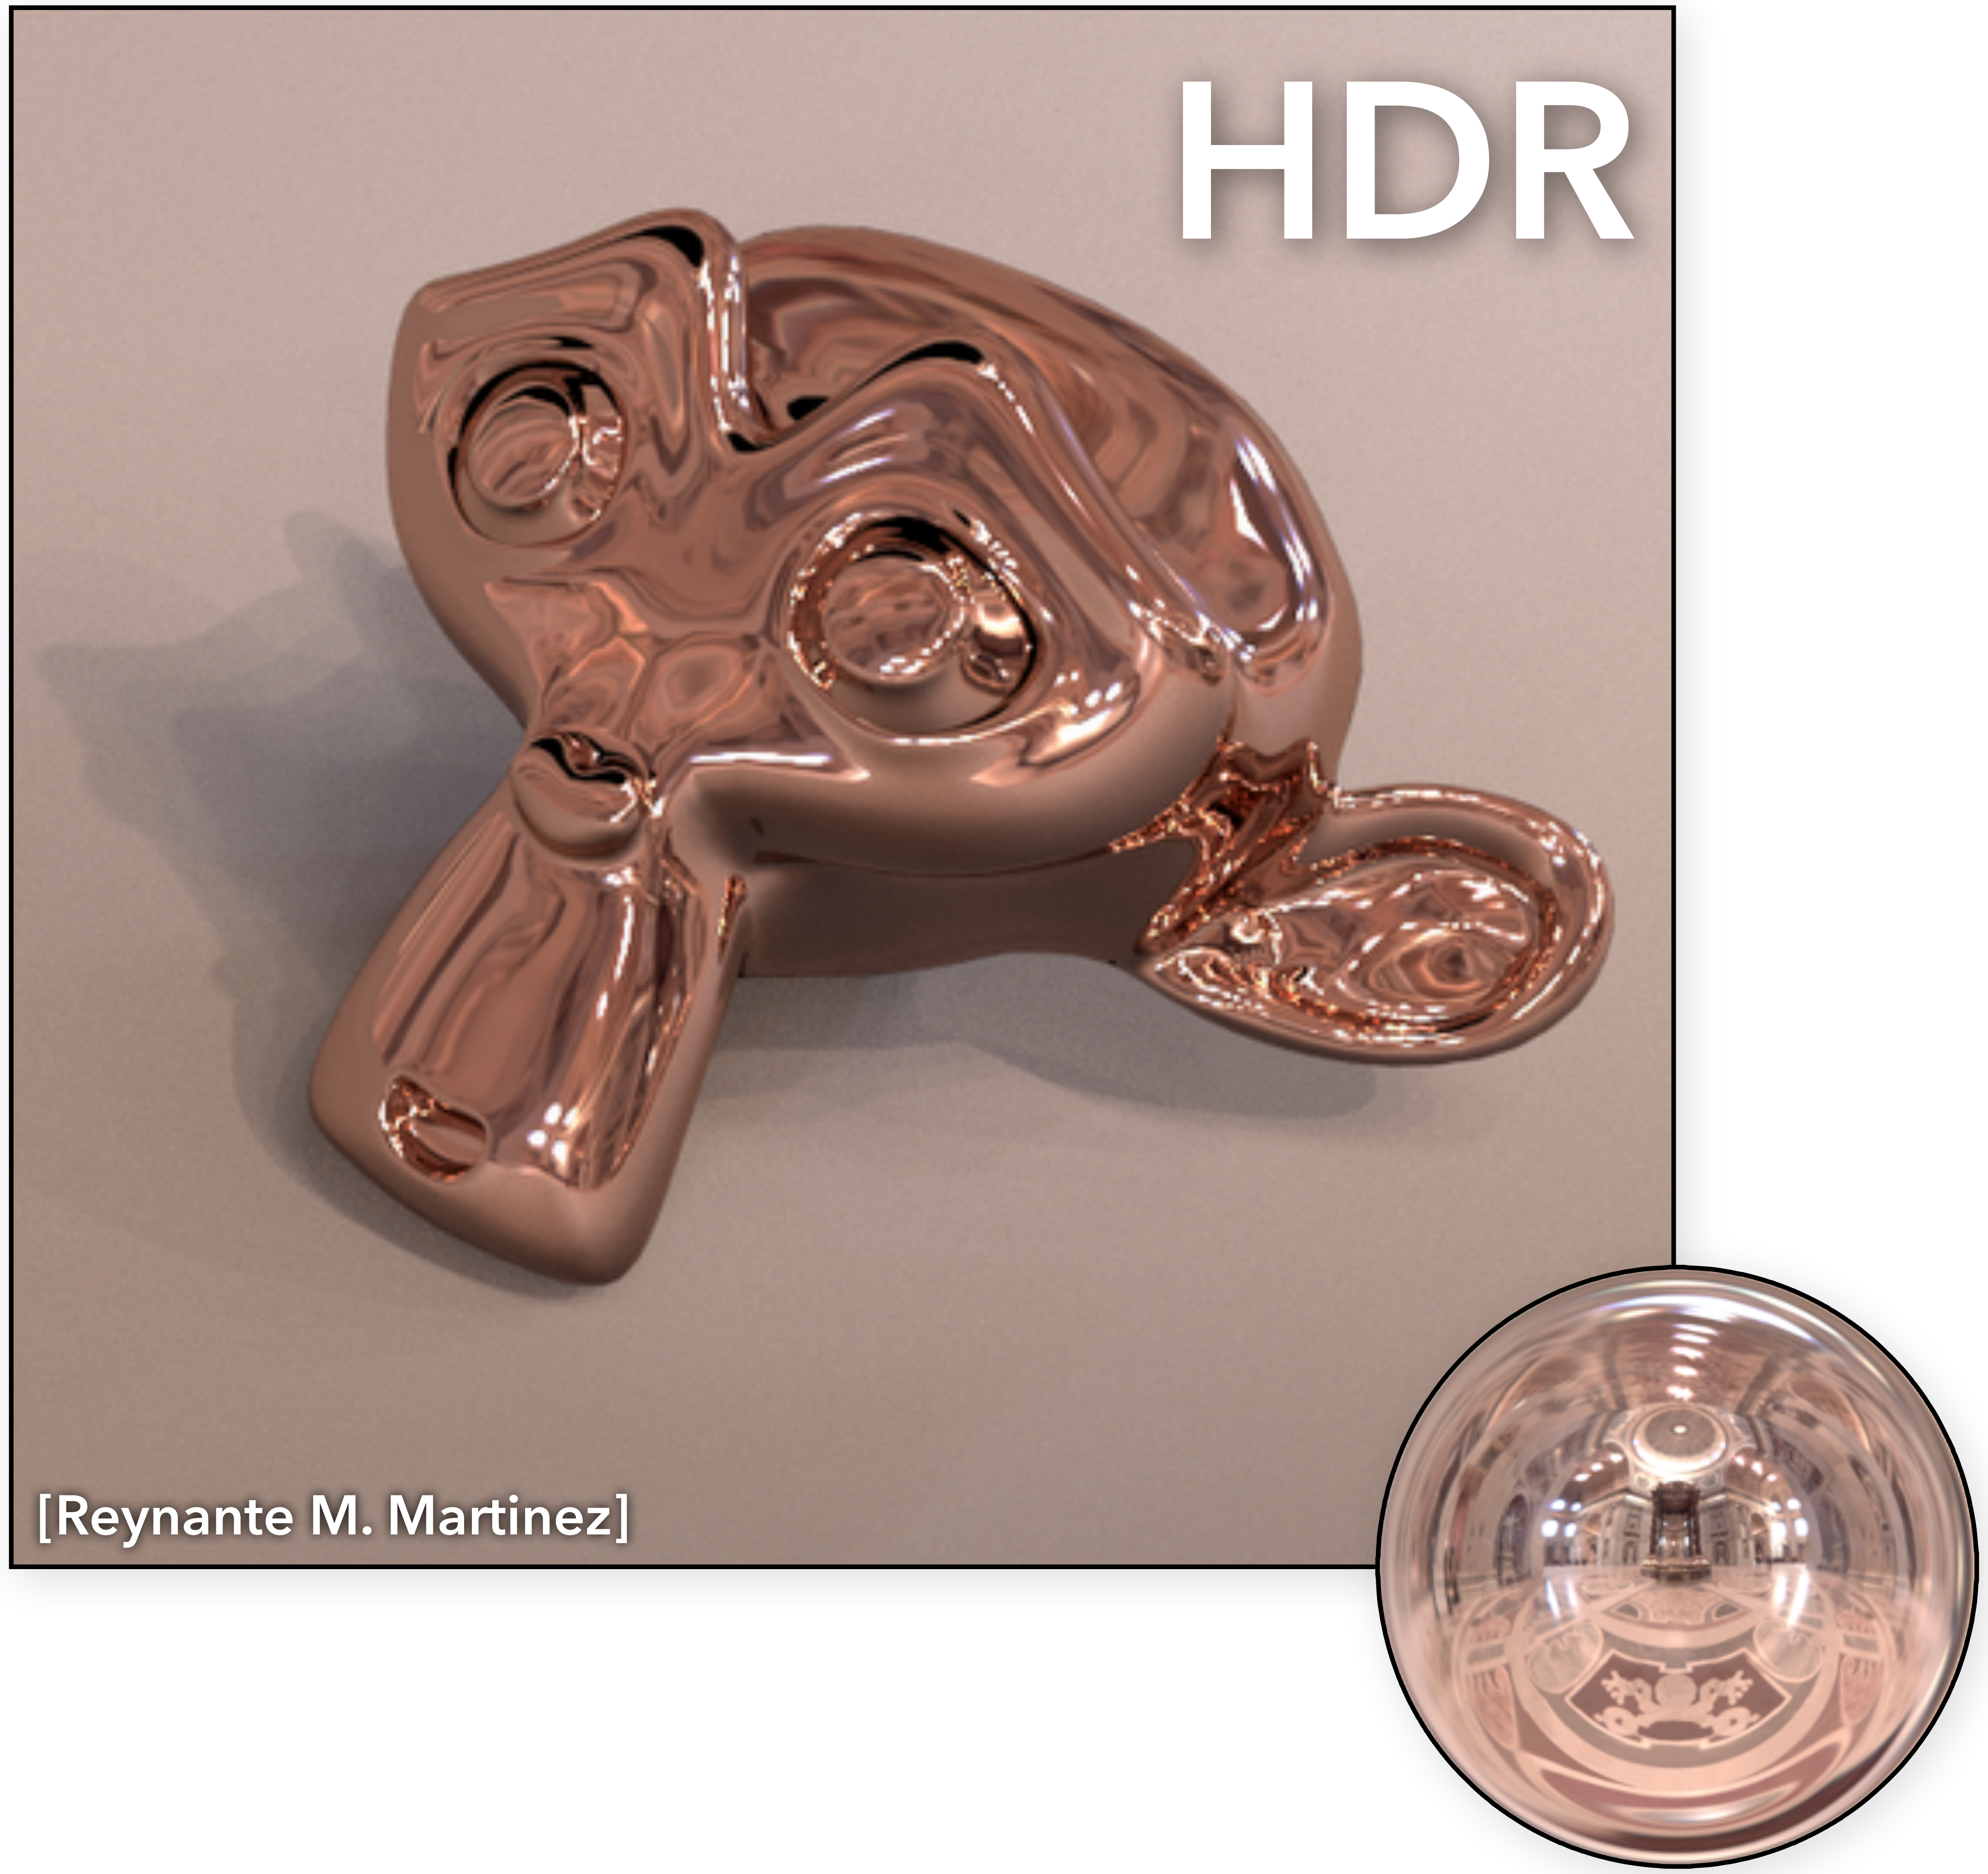

App: Measuring material properties

Marble Block

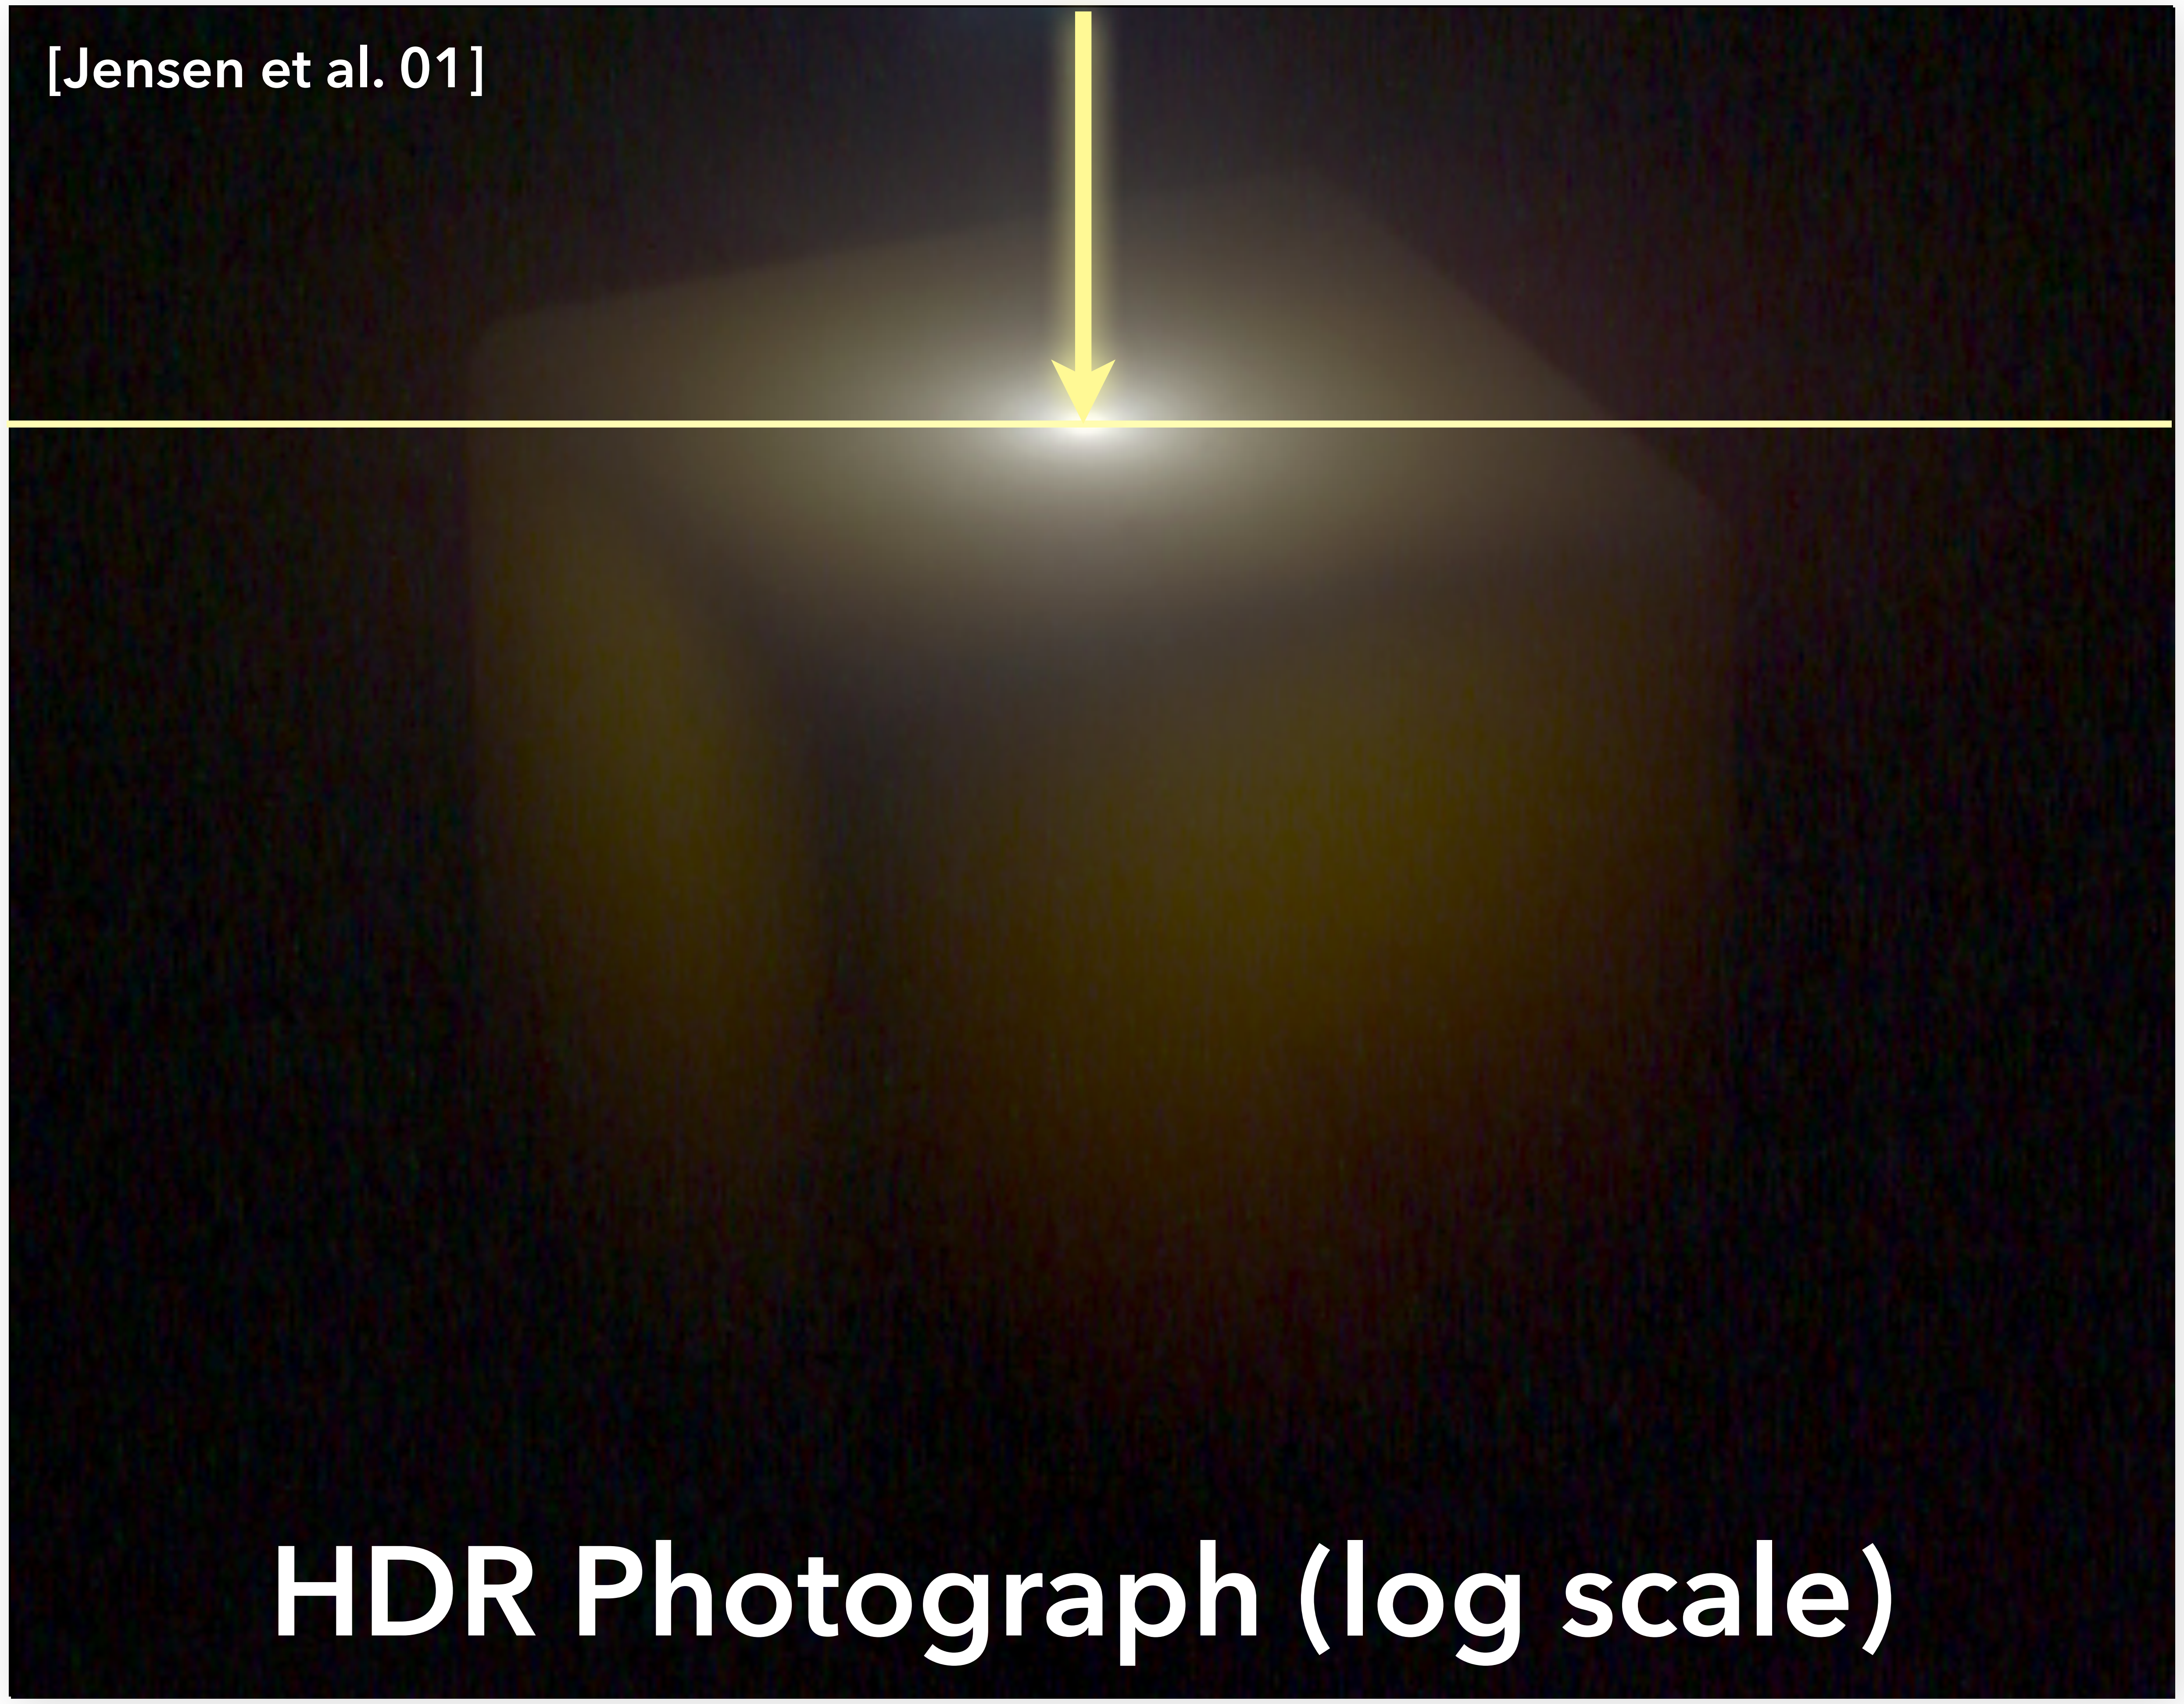

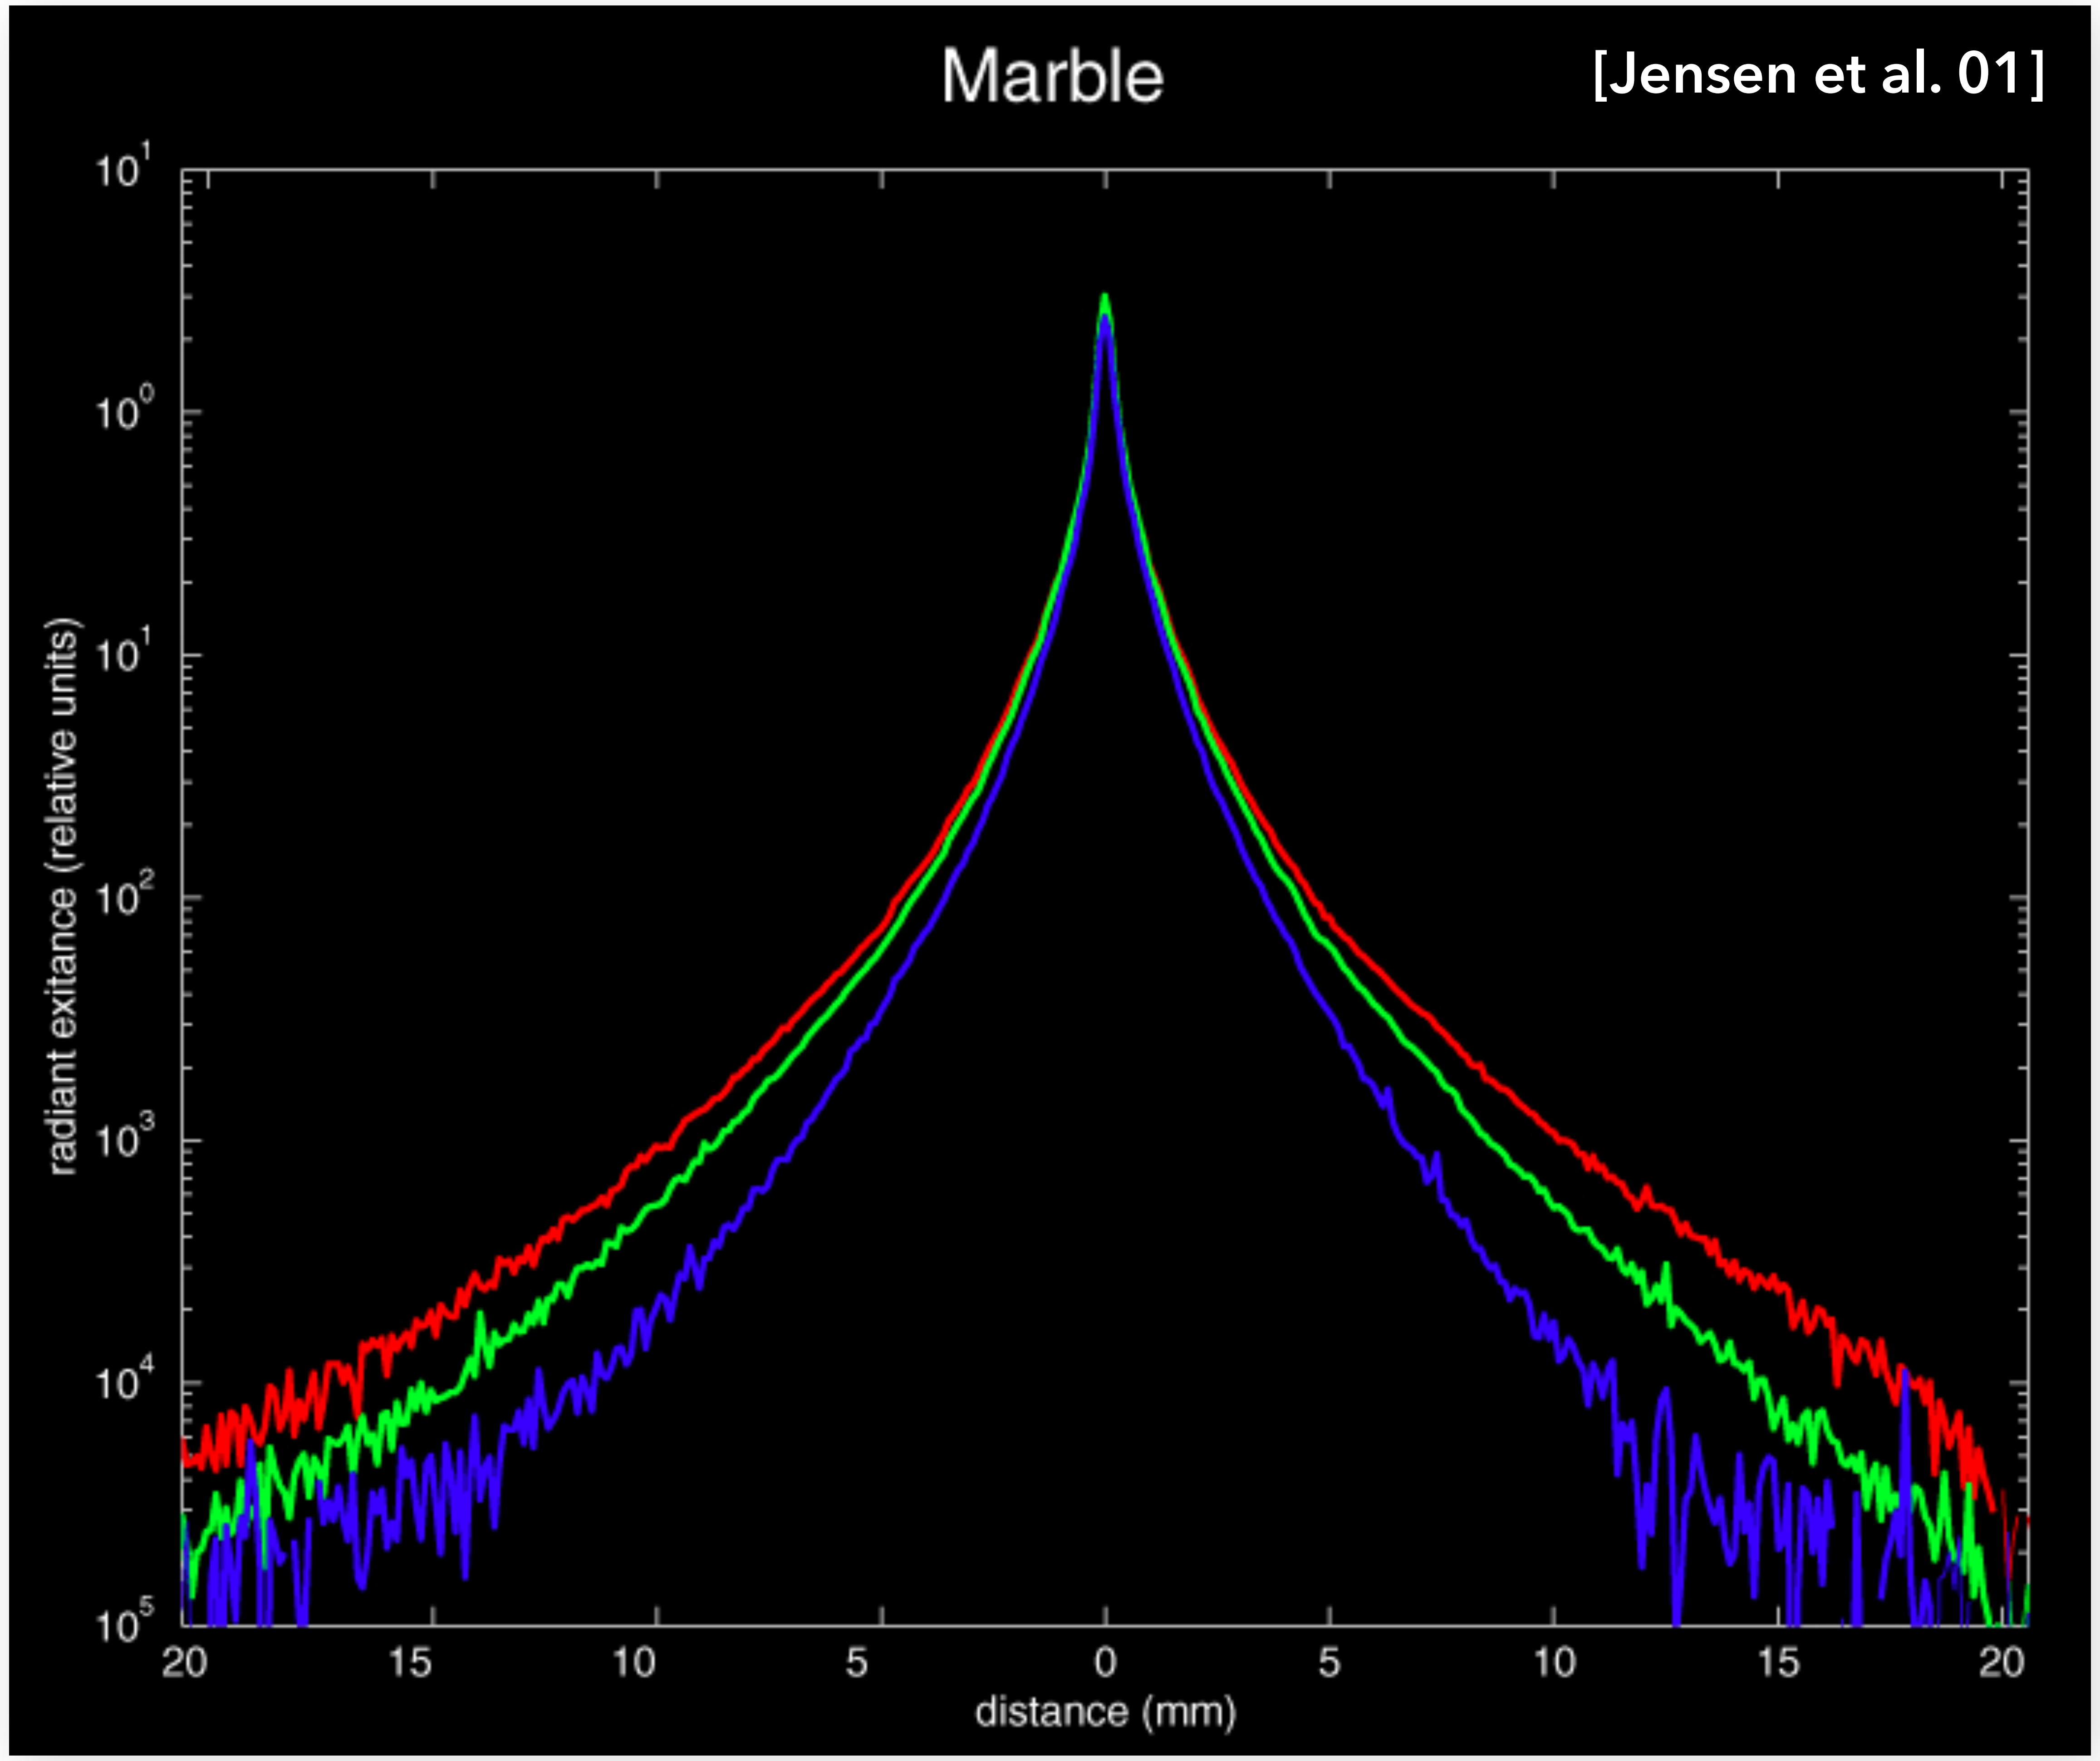

Application: Photography

“Sunset from Rigi Kaltbad”

[Wojciech Jarosz 2014]




"Sunset from Rigi Kaltbad"

[Wojciech Jarosz 2014]

Can be extreme

By Anthony Wong
http://abduzeedo.com/20-beautiful-hdr-pictures-part-3

After a slide by Frédo Durand

Not always cheesy

By Alexandre Buisse
http://luminous-landscape.com/essays/hdr-plea.shtml

After a slide by Frédo Durand

Today

Multiple-exposure High-Dynamic-Range imaging

Tone mapping the image for display/print

After a slide by Frédo Durand

At the end

Some practical tips when taking the photo

Representing HDR images

In general, need to represent extremely small and extremely large values

Typically some form of floating-point representation

Our FloatImage class already does this!

We'll talk about storing/encoding HDR images later

Capturing HDR images

Capturing HDR

If the scene doesn't have extreme dynamic range

- might not need "HDR" at all
- could just take one RAW image (> 8 bits)
- already higher dynamic range than display/JPEG
- make sure to use low ISO (and tripod)

“Abandoned Ship at Point Reyes”

[Wojciech Jarosz 2014]

“Abandoned Ship at Point Reyes”

[Wojciech Jarosz 2014]

“Abandoned Ship at Point Reyes” [Wojciech Jarosz 2014]

Capturing HDR

If the scene doesn't have extreme dynamic range

- might not need HDR at all
- could just take one RAW image (> 8 bits)

High dynamic range scene

- bracket exposures

Multiple exposure photography

Sequentially measure all segments of the range

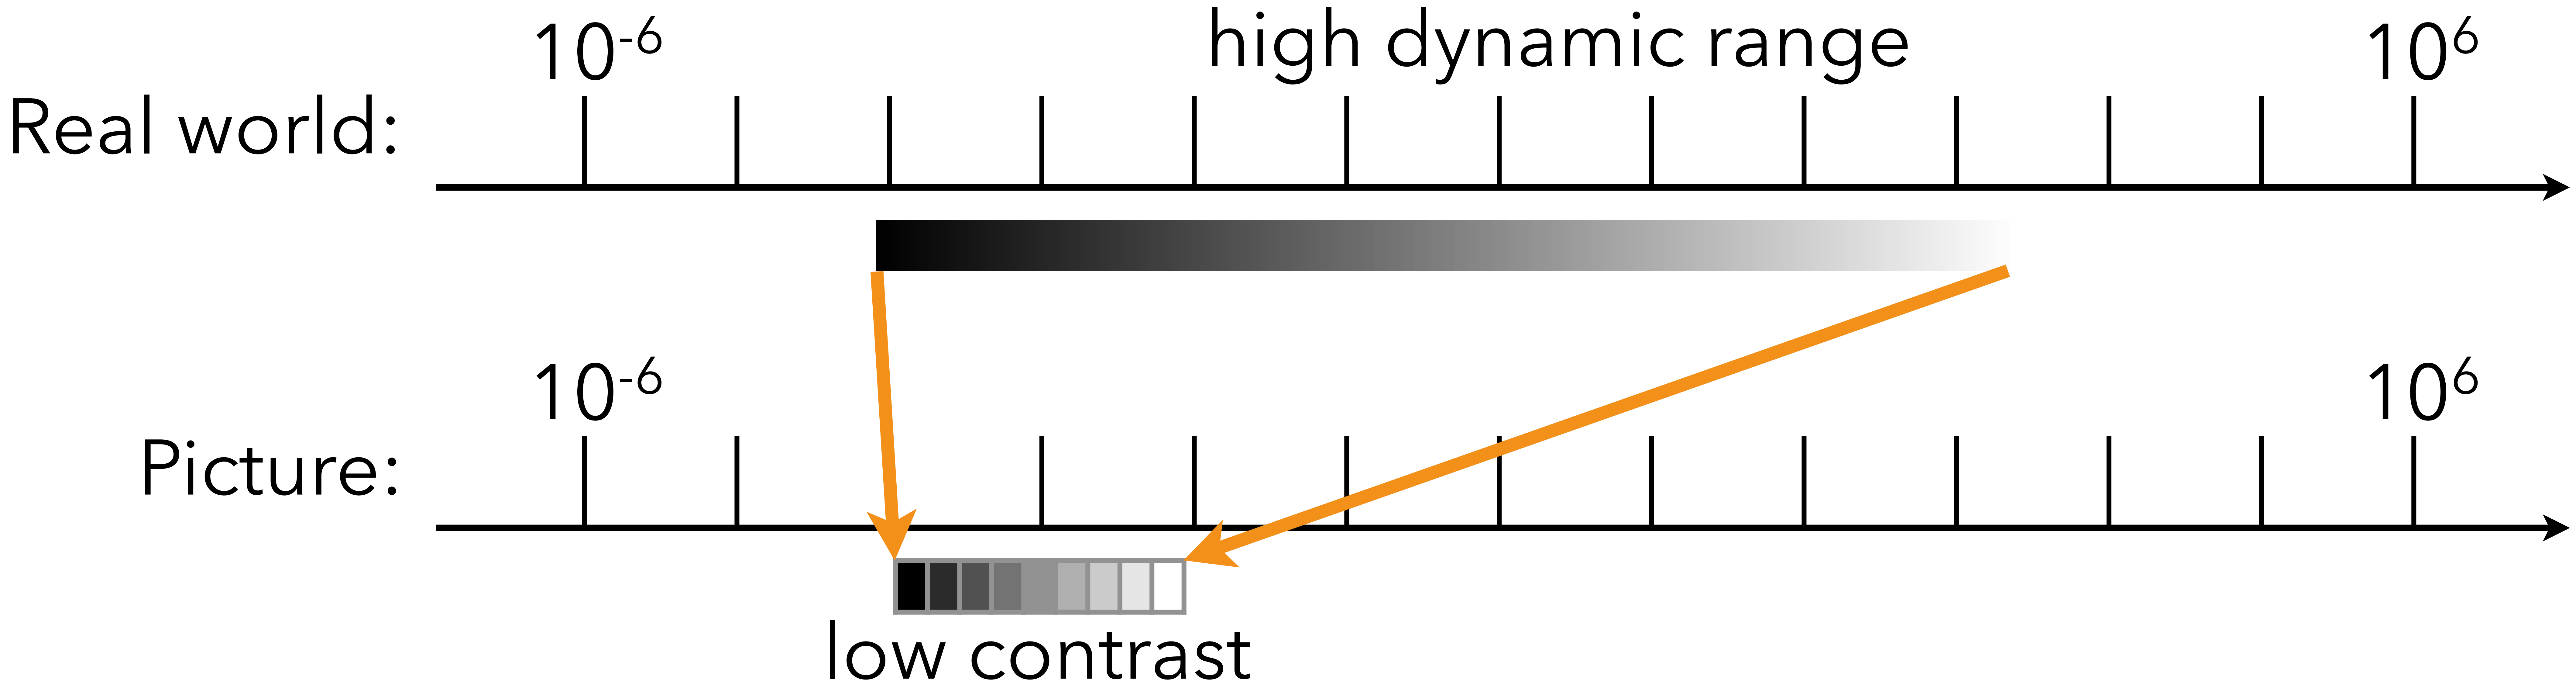

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

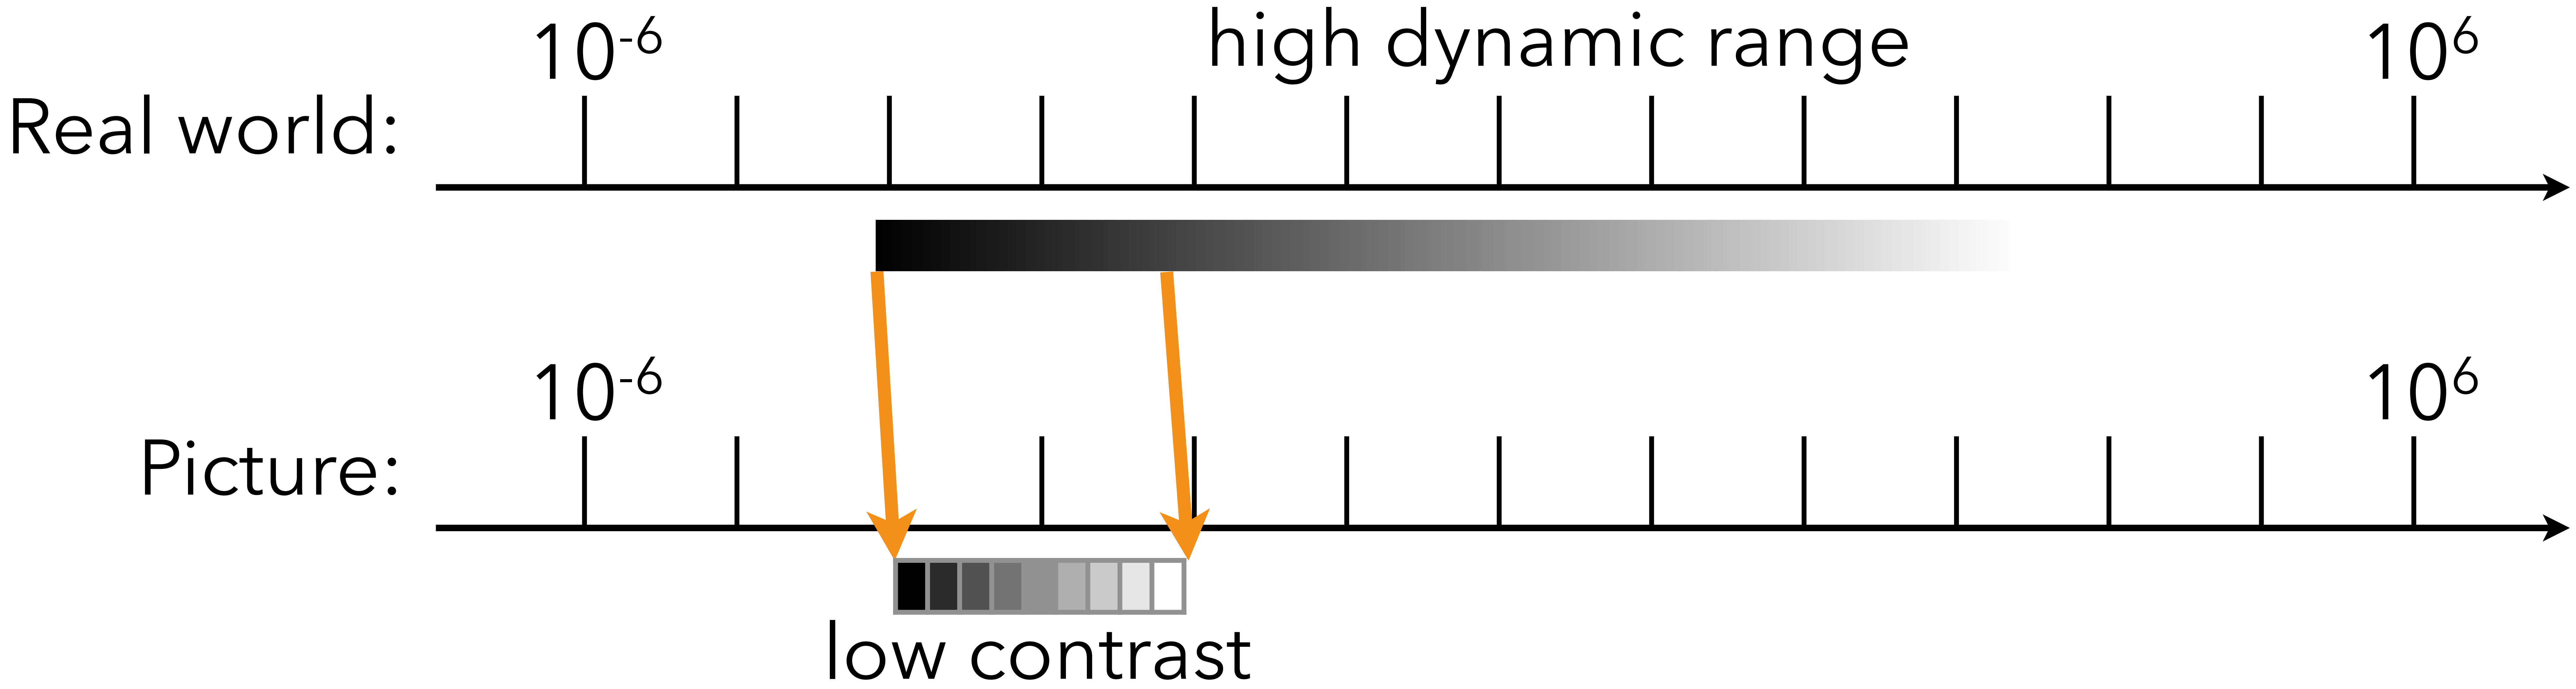

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

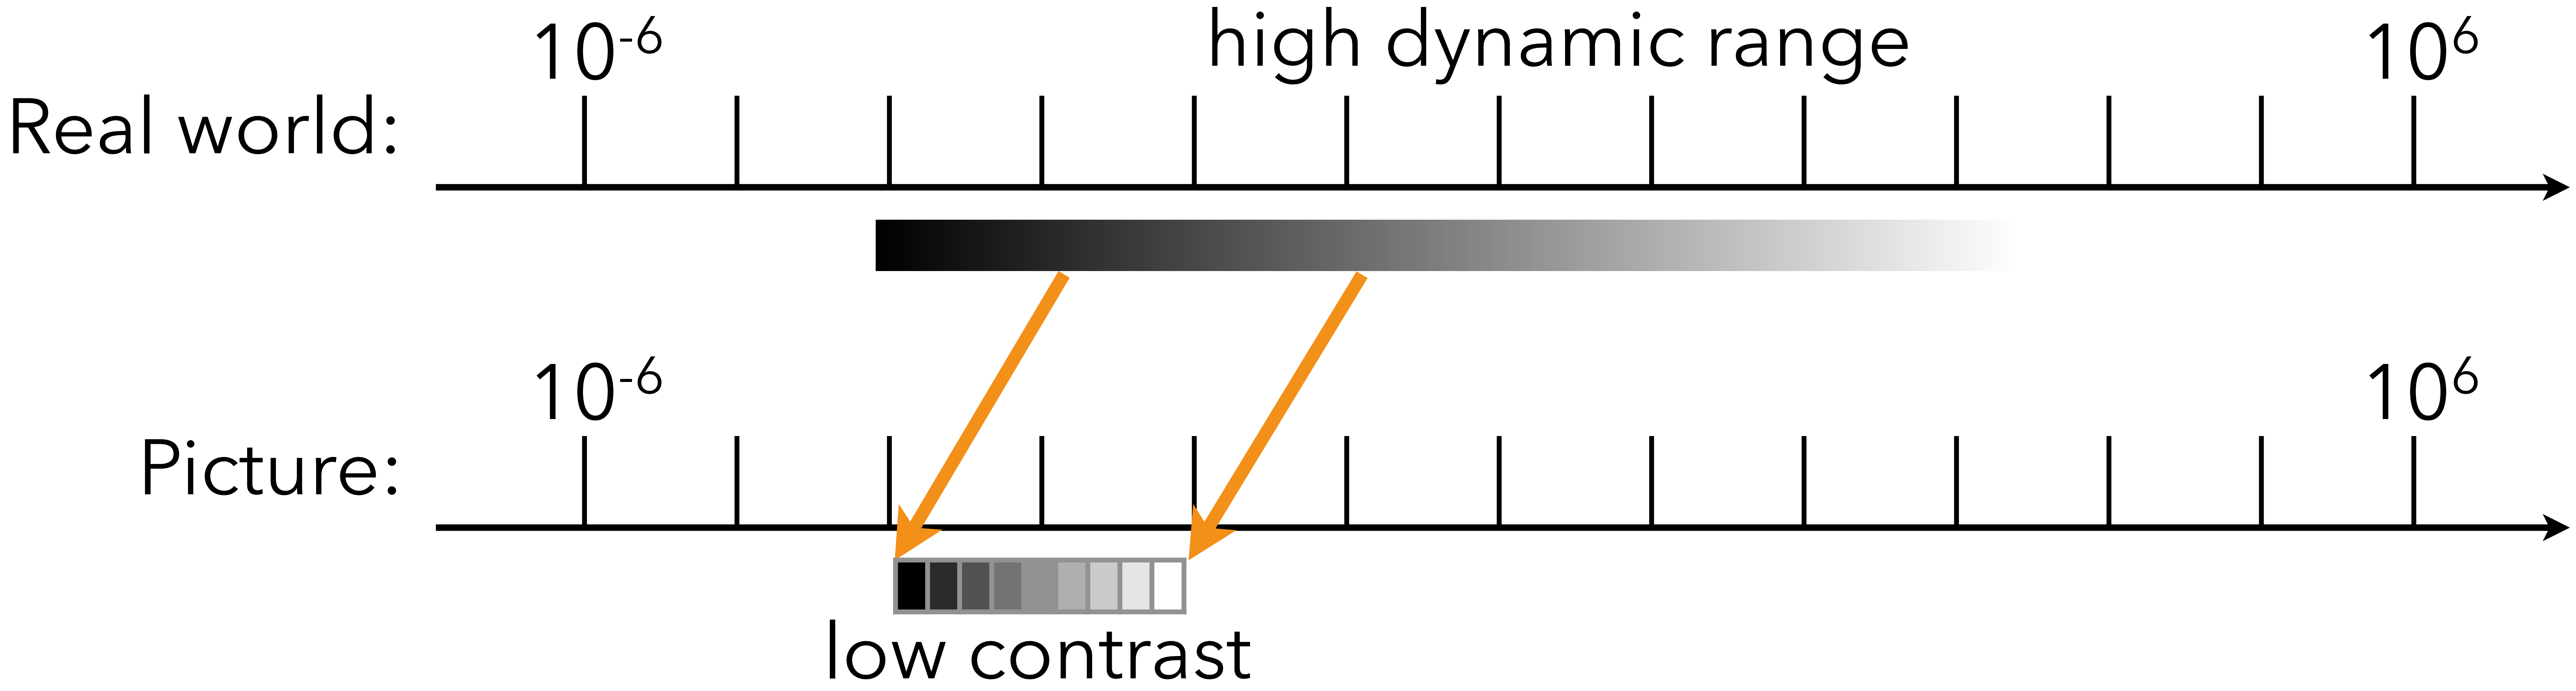

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

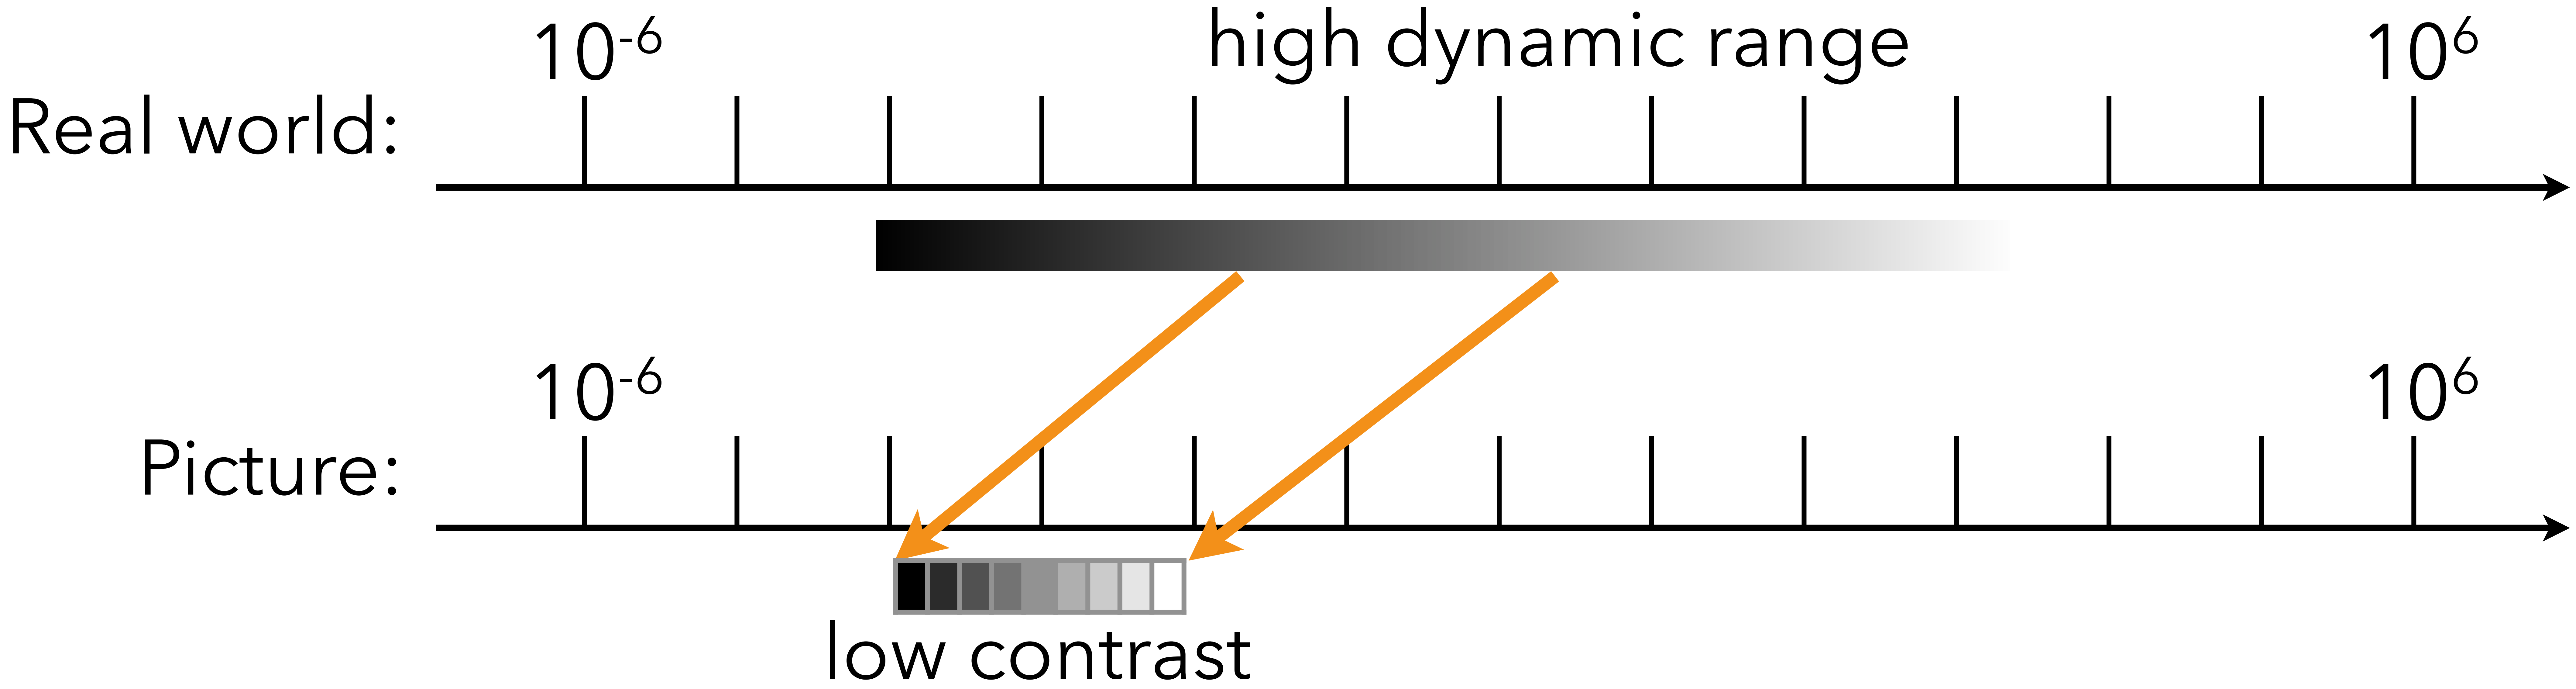

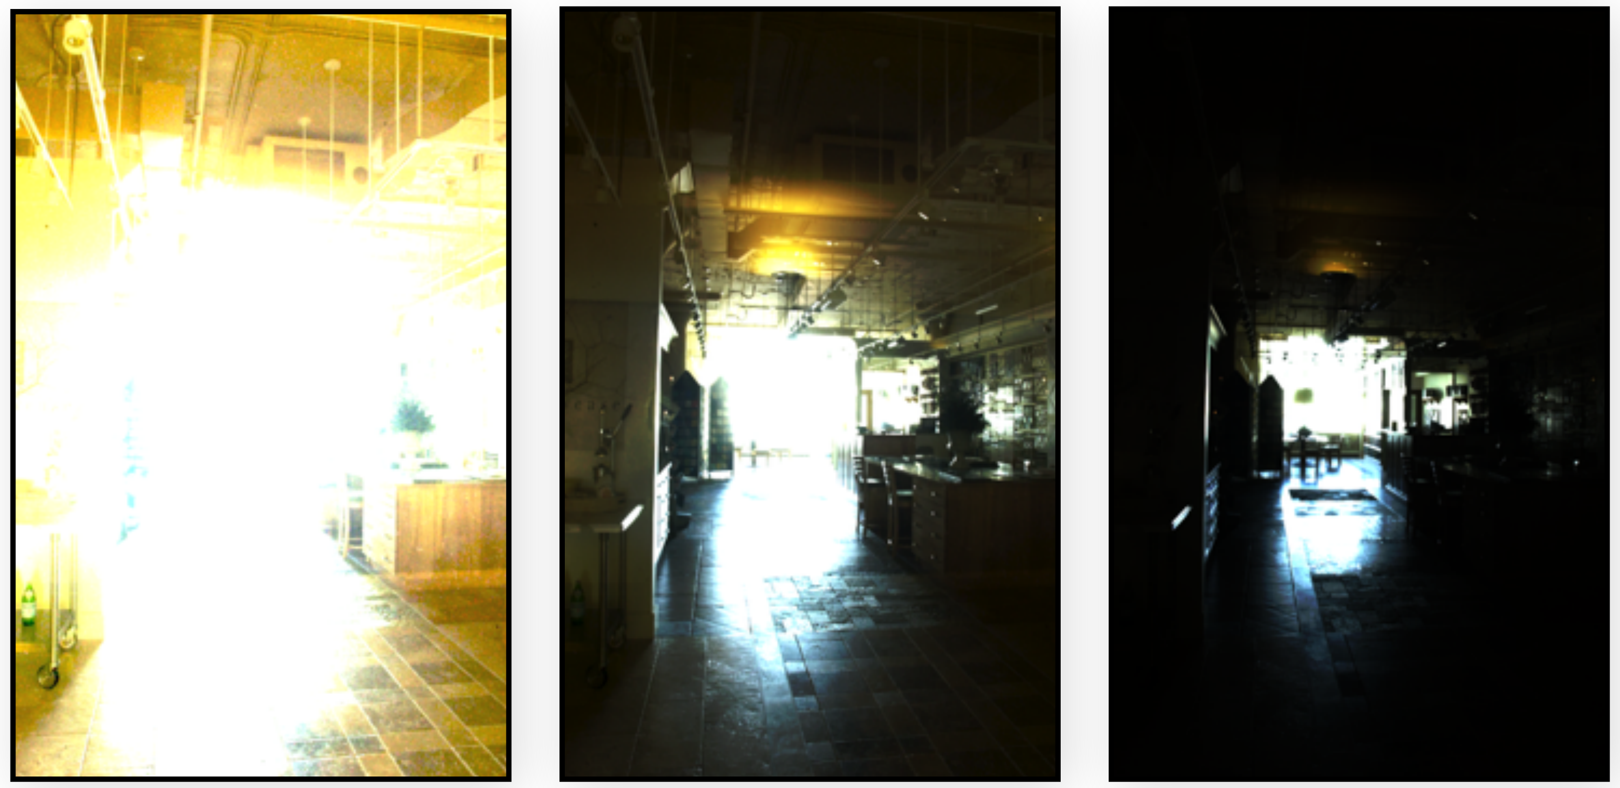

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

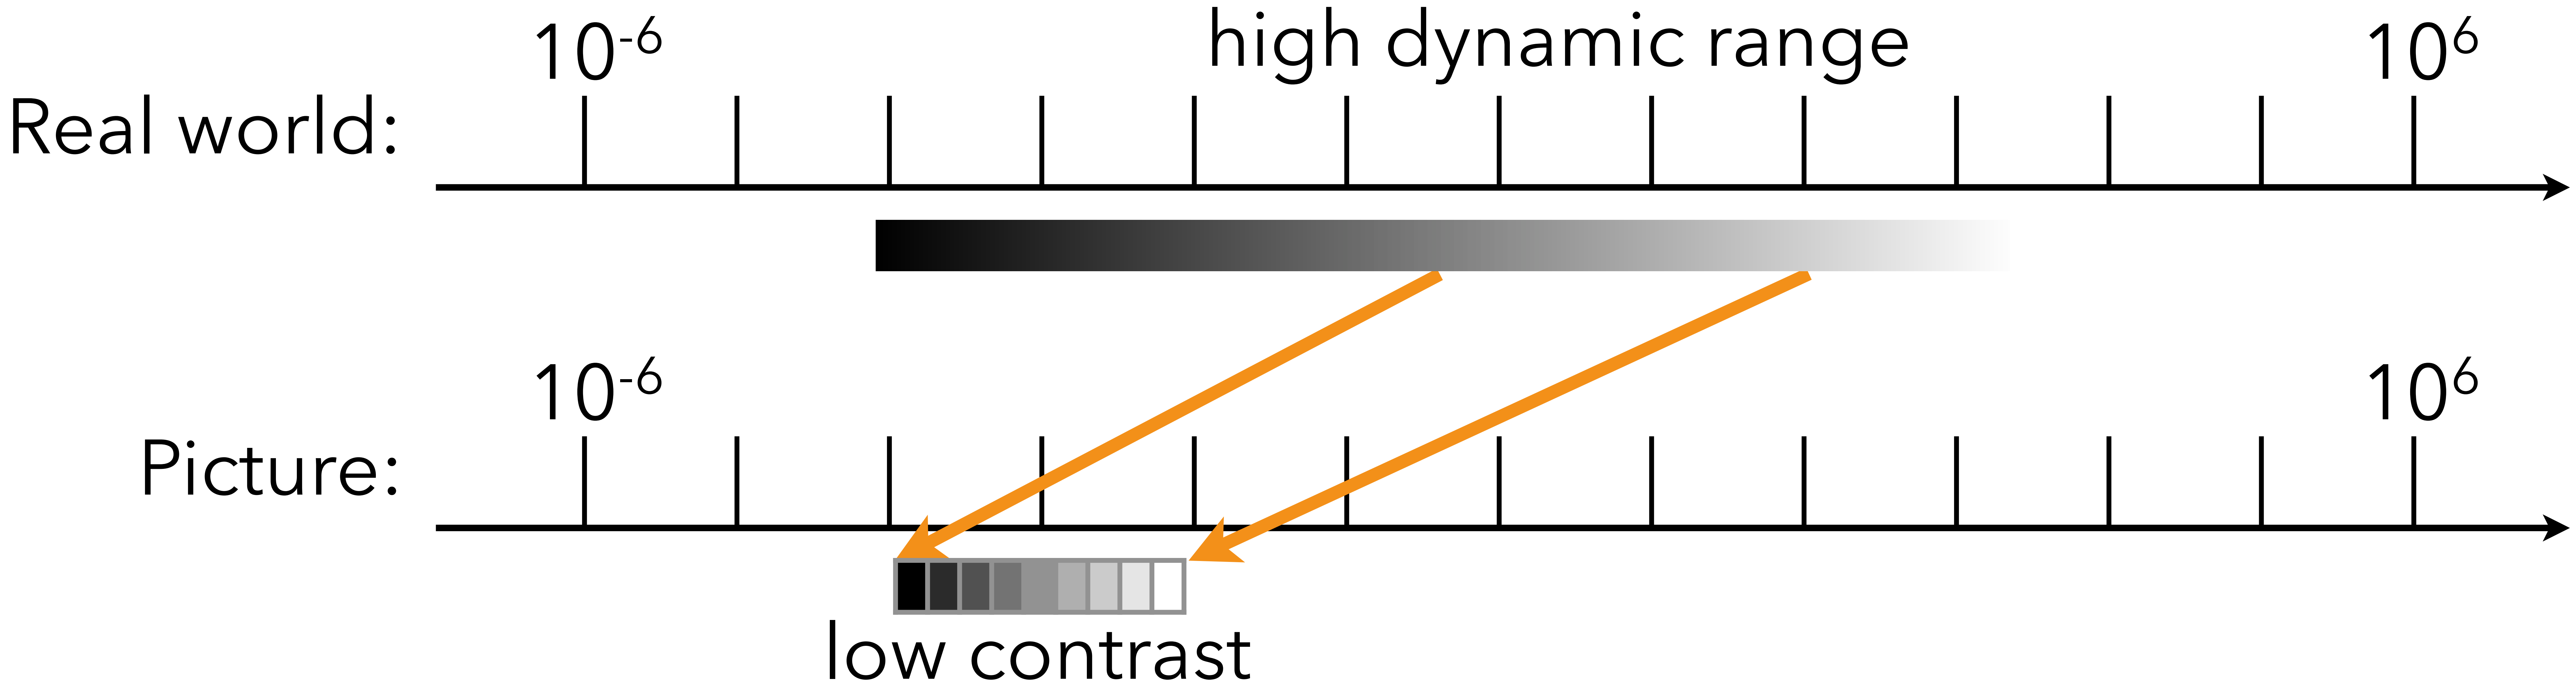

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

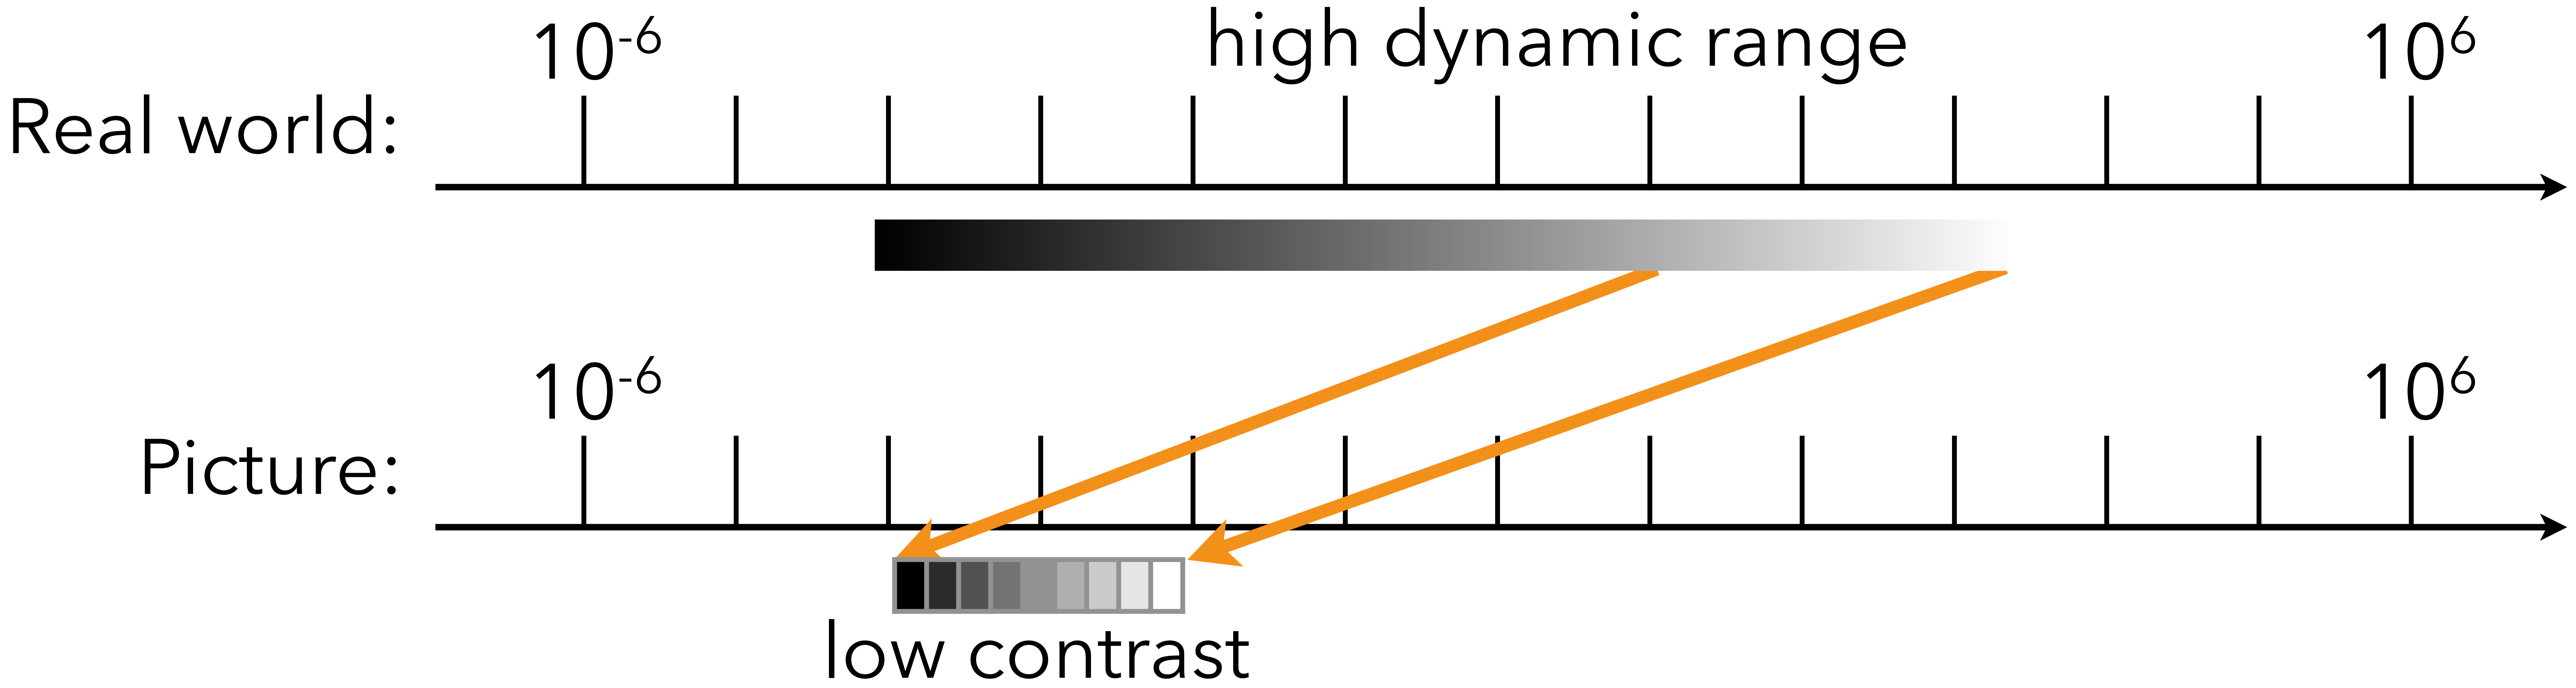

After a slide by Frédo Durand

“Camogli Lighthouse” [Wojciech Jarosz 2012]

“Camogli Lighthouse”

[Wojciech Jarosz 2012]

“Florence”

[Wojciech Jarosz 2011]

“Florence”

[Wojciech Jarosz 2011]

"Matterhorn and Riffelsee"

[Wojciech Jarosz 2010]

"Matterhorn and Riffelsee"

[Wojciech Jarosz 2010]

How do we vary exposure?

Options:

- Shutter speed
- Aperture
- ISO
- Neutral density filter

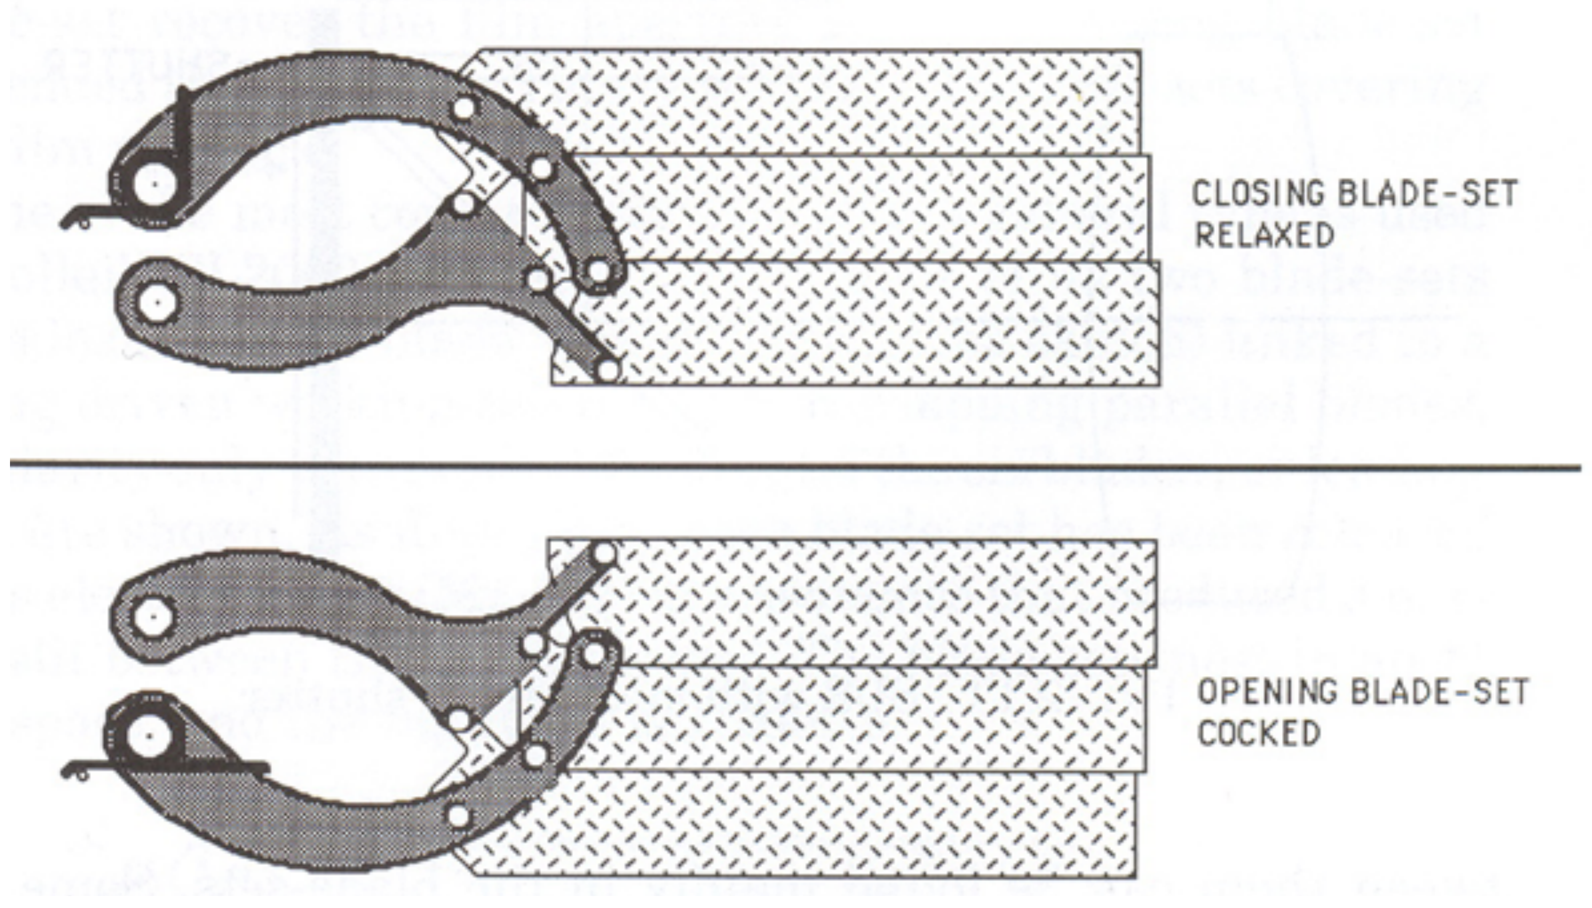

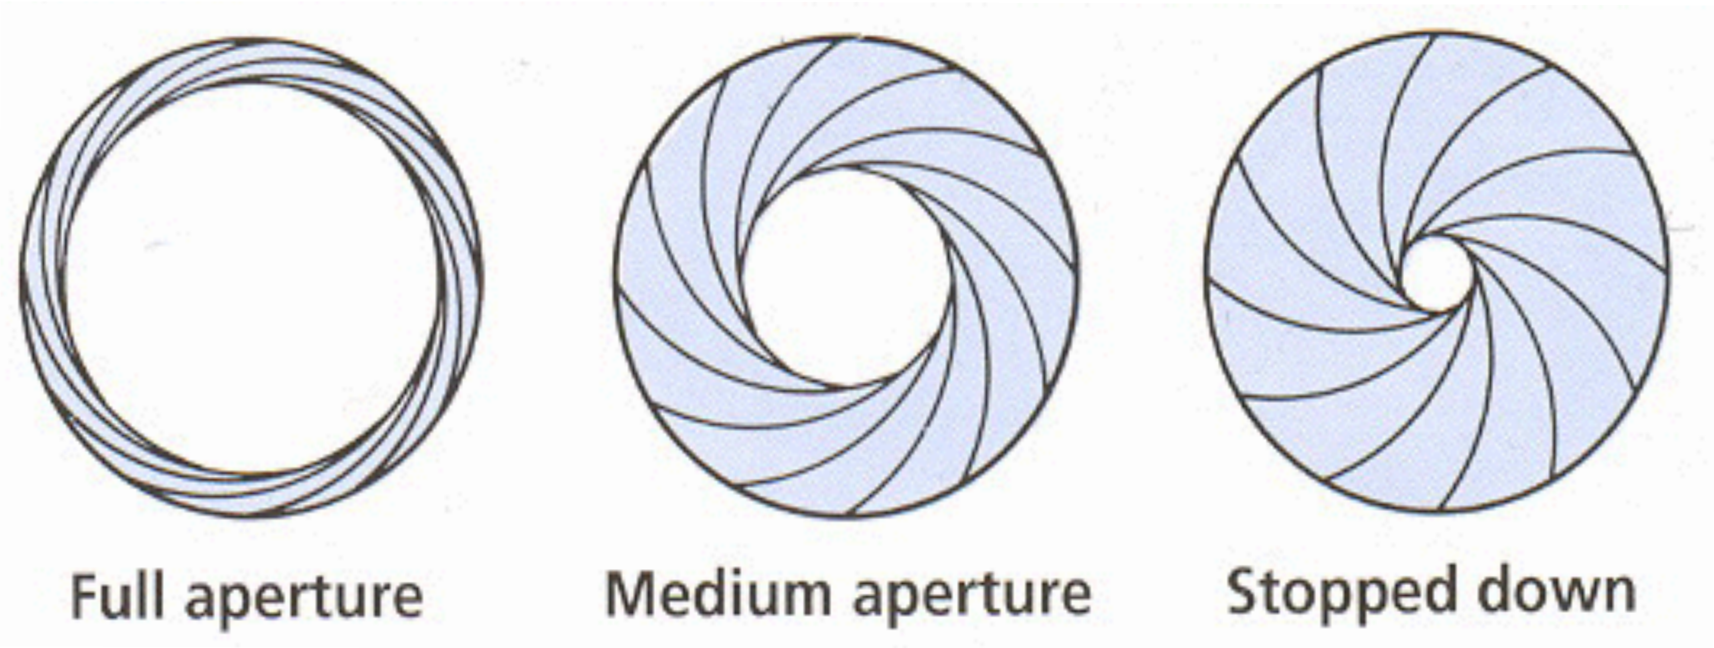

After a slide by Frédo Durand

Slide inspired by Siggraph 2005 course on HDR

Tradeoffs

After a slide by Frédo Durand

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

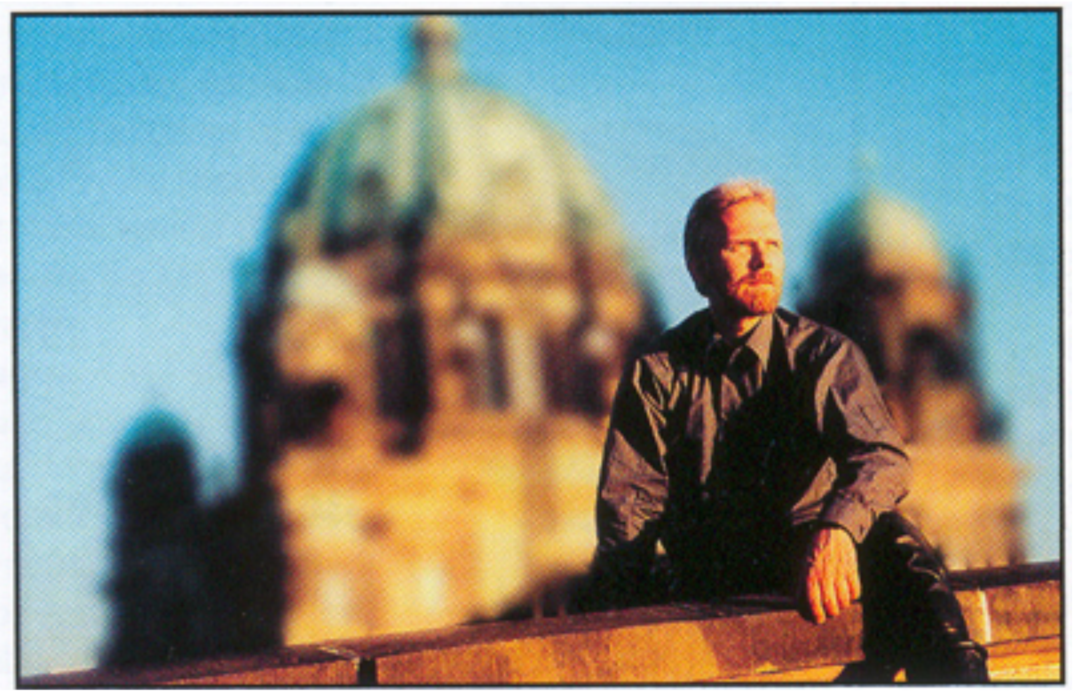

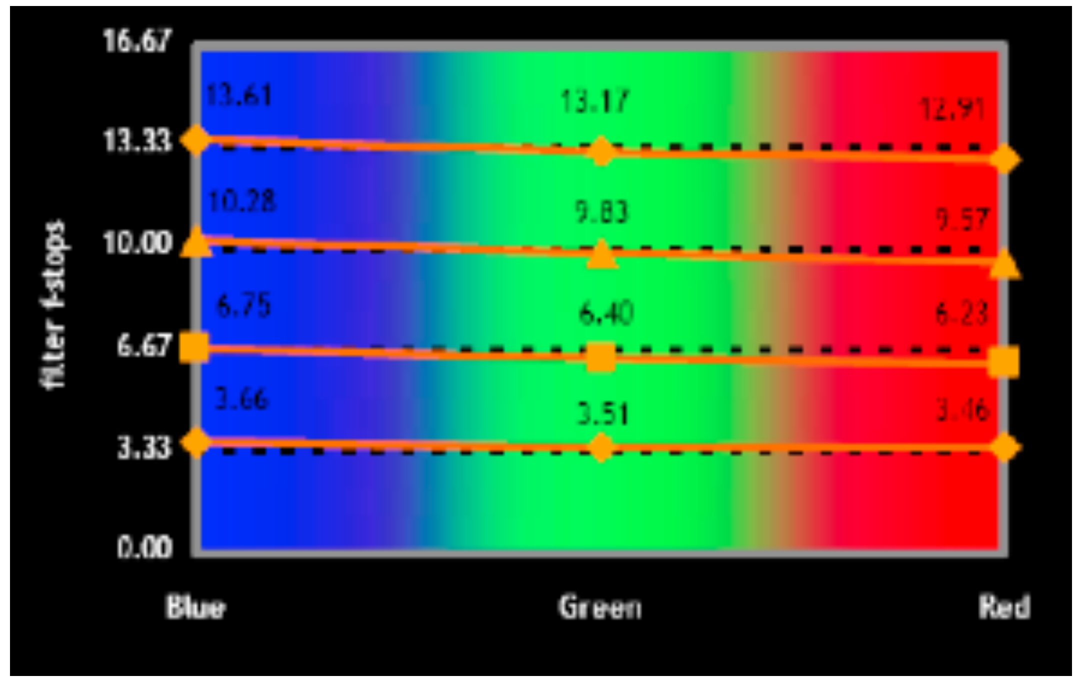

HDR merging: linear case

Problem statement

We have N images

- images are encoded linearly
- only exposure changes: no motion

We want 1 single HDR image

- encoded with FloatImage class
- one value per x, y, c
- values may be >1

After a slide by Frédo Durand

Getting linear images

http://www.mit.edu/~kimo/blog/linear.html

http://www.luminous-landscape.com/forum/index.php?topic=25064

http://www.guillermoluijk.com/tutorial/dcraw/index_en.htm

./dcrawx86 -v -H 0 -g 2.2 0 -o 1 DSC_*nef

After a slide by Frédo Durand

Image formation: photons to floats

Scene radiance L(x,y) reaches the sensor at a pixel x, y

Value of pixel?

- depends on shutter speed, aperture, ISO (multiplicative factor)
- clips if >1
- noise gets added

After a slide by Frédo Durand

Image formation: photons to floats

Scene radiance L(x,y) reaches the sensor at a pixel x, y

For each image i,

- radiance gets multiplied by exposure factor k_i (depends on shutter speed, aperture, ISO)
- noise n gets added
- values above 1 get clipped (depends on photosite well capacity)

$$I_i(x, y) = \text{clip}(k_i * L(x,y)+n)$$

After a slide by Frédo Durand

Dynamic range

$I_i(x, y) = \text{clip}(k_i * L(x,y) + n)$

In the highlights, we are limited by clipping

In the shadows, we are limited by noise

Brightened many times

After a slide by Frédo Durand

2-image example

Simple in principle:

- imageA = 1/30th second ("brighter" image)
- imageB = 1/120th second ("darker" image)
- imageHDR = average(4·imageB, remove-clipped(imageA))
- assumes images have been linearized

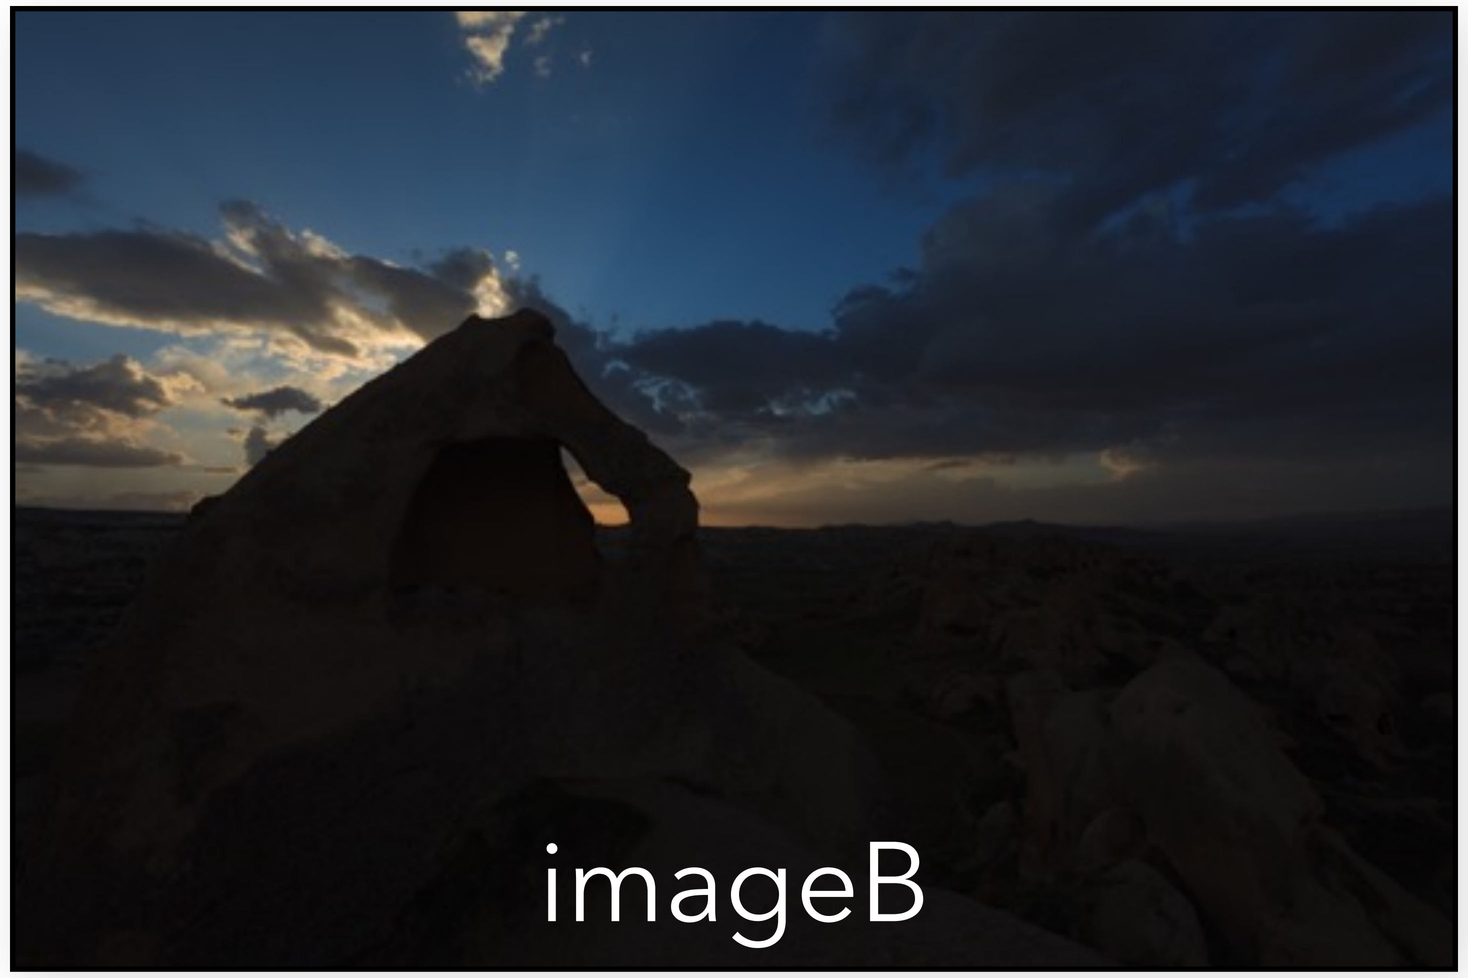

General HDR merging

$I_i(x, y) = \text{clip}(k_i * L(x,y) + n)$

For each pixel

- figure out which images are useful
- scale values appropriately (ideally according to k_i)
- average scaled values from useful images

After a slide by Frédo Durand

Which images are useful?

Eliminate clipped pixels

- e.g. >0.99

Eliminate pixels that are too dark / too noisy

- e.g. <0.002

After a slide by Frédo Durand

Eliminate bad pixels

In Assignment 5

We compute a weight map for each image

We use binary weights (0 or 1)

- but can be extended to full scalar (to better handle noise)

Weights can be different for different channels

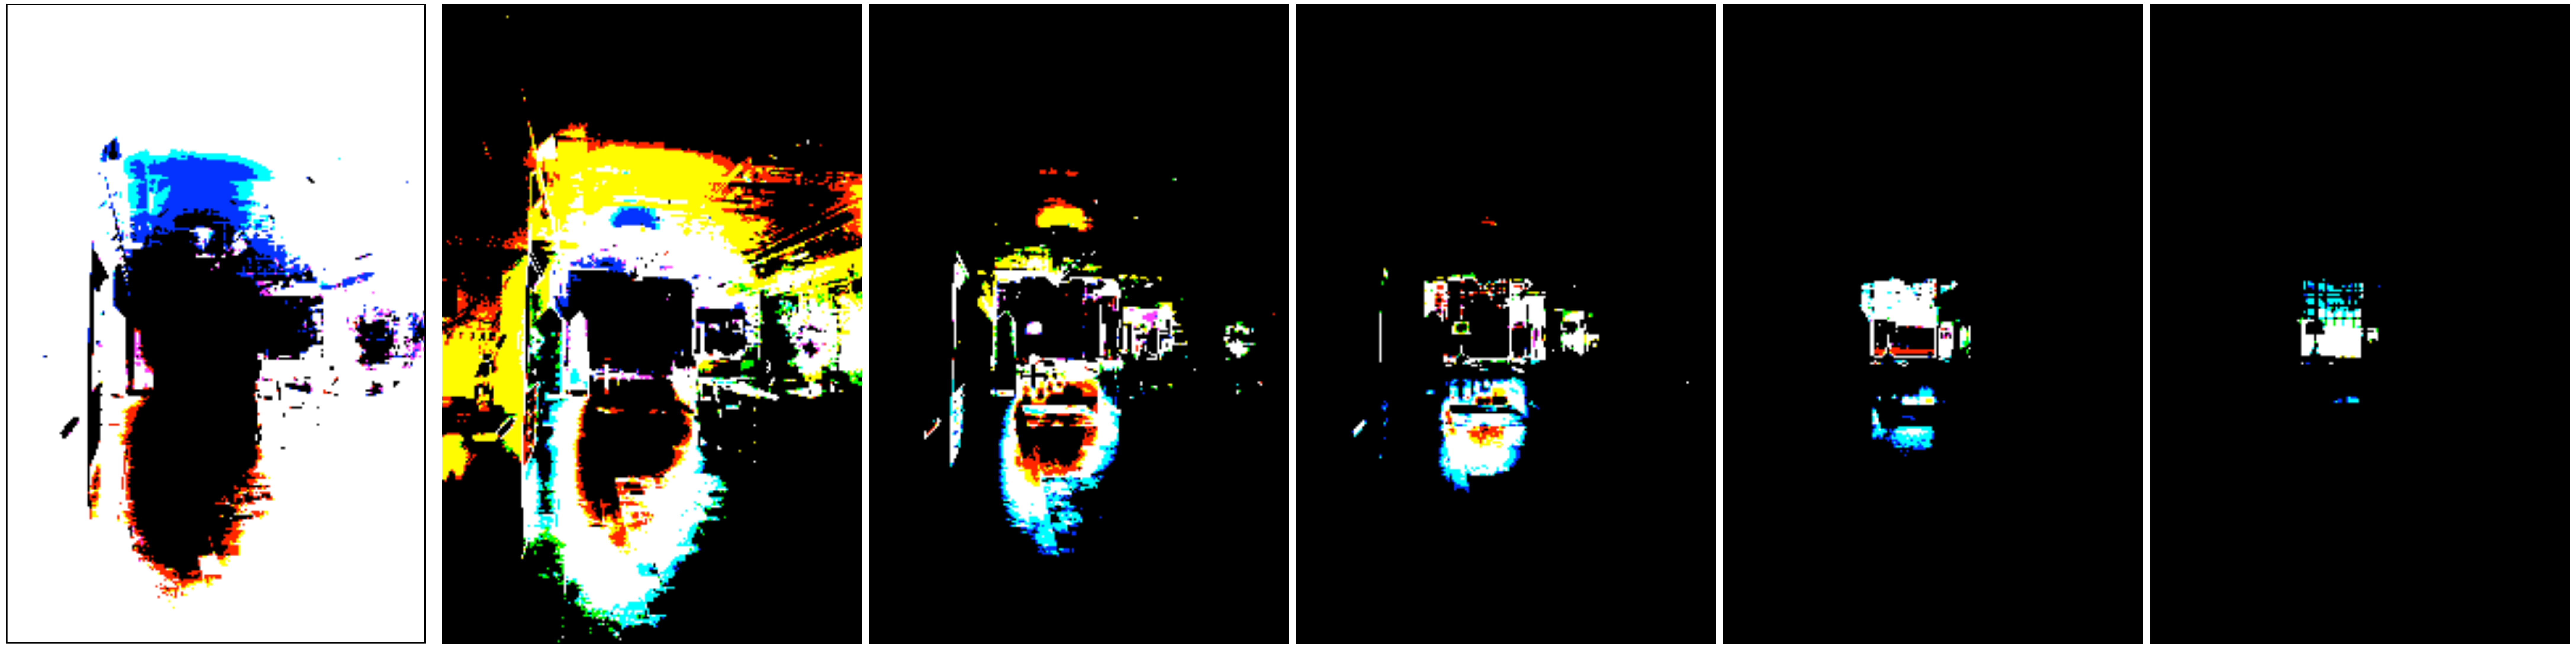

After a slide by Frédo Durand

Assembling HDR

Figure out scale factor between images

- from exposure data, or
- by looking at ratios $I_i(x,y)/I_j(x,y)$ (only when both are good)

Compute weight map w_i for each image

Reconstruct full image using weighted combination

$$out(x,y) = \frac{1}{\sum w_i(x,y)} \sum w_i(x,y) \frac{1}{k_i} I_i(x,y)$$

After a slide by Frédo Durand

Computing k_i in PA 5

$I_i(x, y) = clip(k_i * L(x,y)+n)$

Only up to global scale factor

Actually compute k_i/k_j for pairs of images

Focus on pixels where

- no clipping occurs & noise is negligible

$I_i(x, y) = k_i*L(x,y)$

get k_i/k_j by considering I_i/I_j

If linearity holds, should be the same for all pixels

Use median for extra robustness

After a slide by Frédo Durand

Computing k_i in PA 5

$I_i(x, y) = \text{clip}(k_i * L(x,y)+n)$

Only up to global scale factor, e.g. k_i/k_0

Actually compute k_i/k_j for pairs of images

$k_i/k_j = \text{median}(I_i(x,y)/I_j(x,y))$

- for pixels st. $w_i(x,y)>0$ AND $w_j(x,y)>0$

Then compute k_i/k_0 by chaining these ratios

After a slide by Frédo Durand

Special cases

Some pixels might be underexposed or overexposed in all images

Simple solution: don't eliminate dark pixels in the brightest image or bright pixels in the darkest one.

In the end: HDR image

Encoded with same FloatImage class

One single float value per x, y, c

numbers may be >1

Questions?

HDR combination papers

Steve Mann http://genesis.eecg.toronto.edu/wyckoff/index.html

Paul Debevec http://www.debevec.org/Research/HDR/

Mitsunaga, Nayar, Grossberg
http://www1.cs.columbia.edu/CAVE/projects/rad_cal/rad_cal.php

http://people.csail.mit.edu/hasinoff/hdrnoise/

Tone mapping

Today

Multiple-exposure High-Dynamic-Range imaging

Tone mapping the image for display/print

Problem 2: Display the information

Match limited contrast of the medium while preserving details

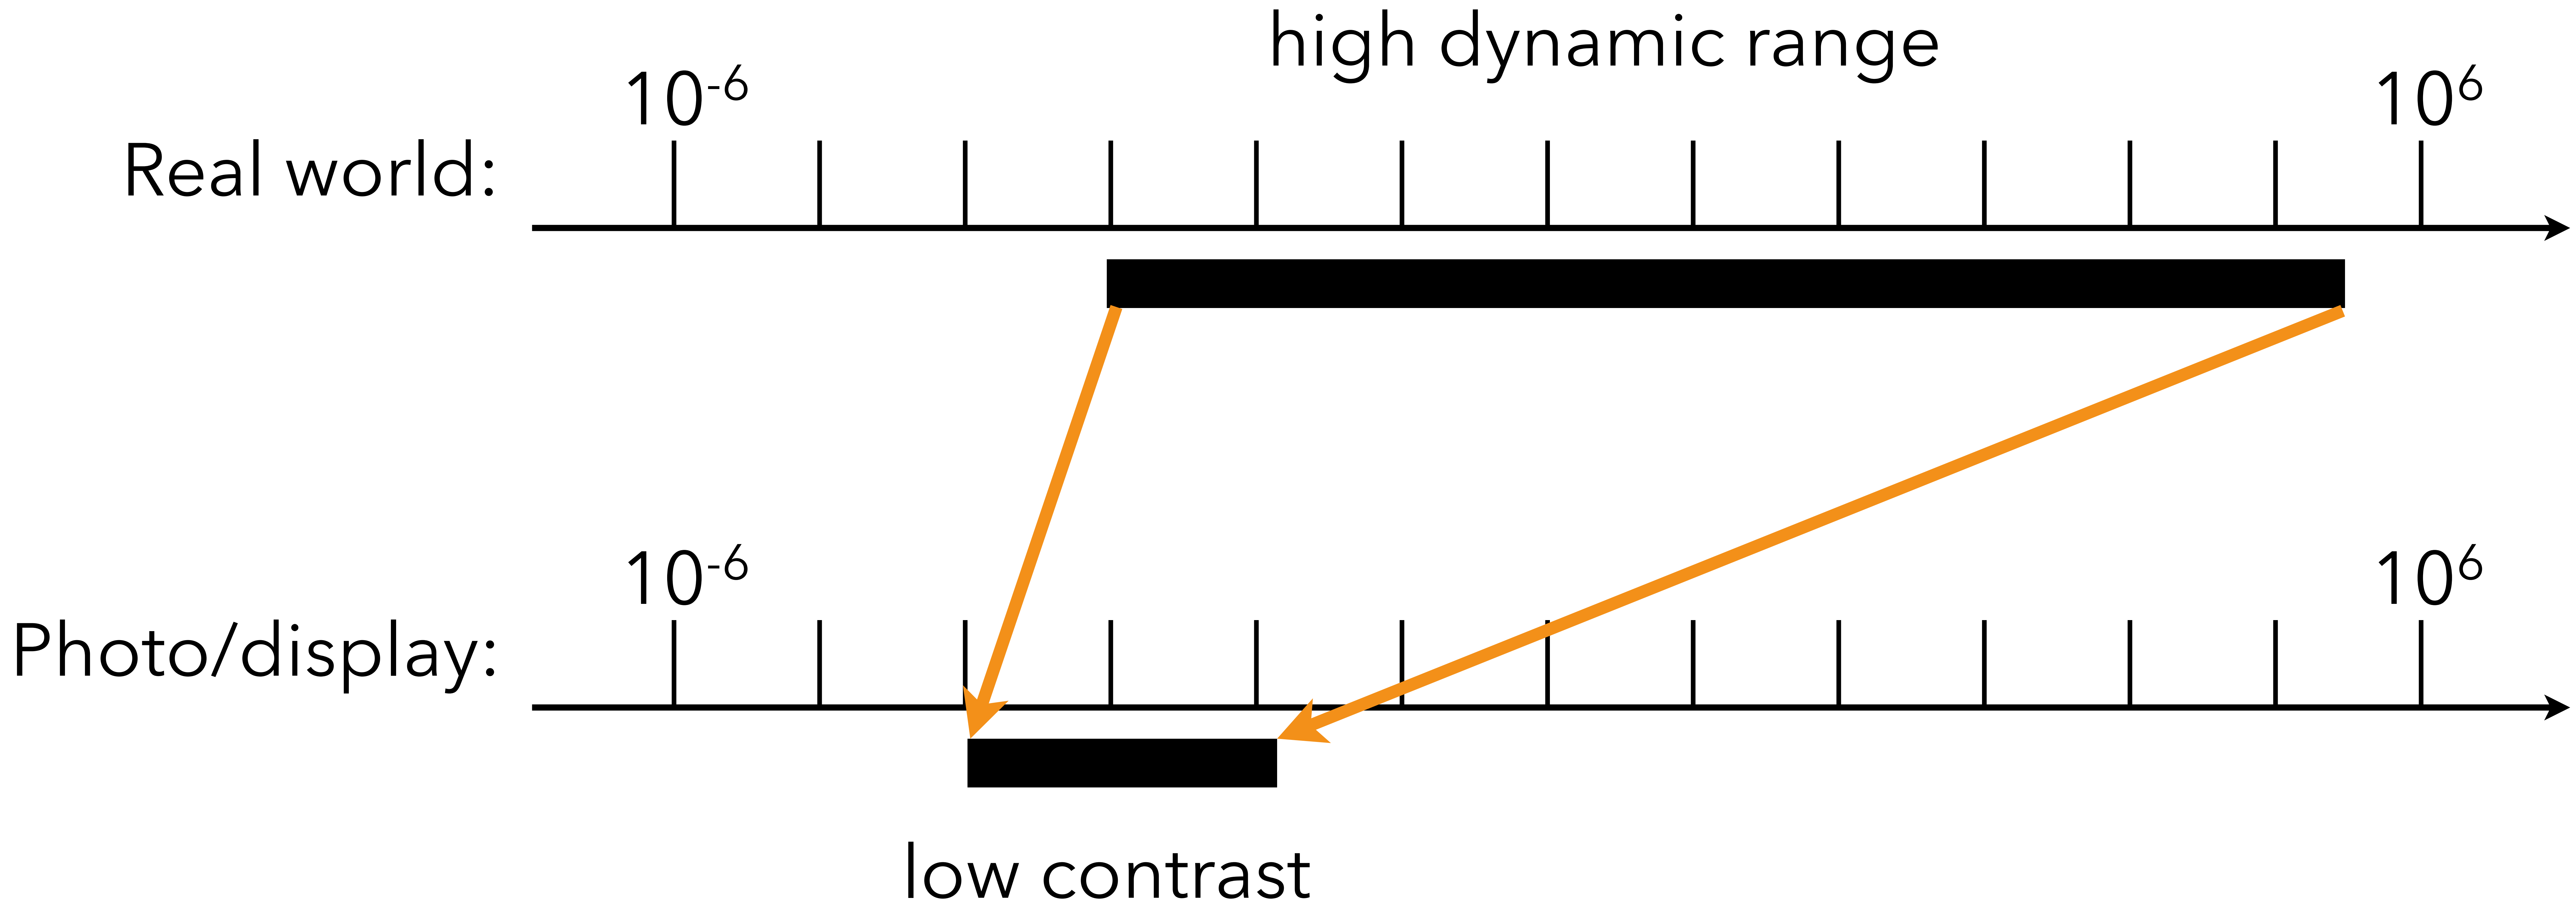

After a slide by Frédo Durand

Tone mapping

Called tone mapping operators

Two general categories:

- Global (spatially invariant)
- Local (spatially varying)

Tone mapping for very HDR scenes

Scene has >100,000:1 dynamic range, JPEG has 255:1

How can we compress the scene's dynamic range?

After a slide by Frédo Durand

Naïve technique?

Scene has >100,000:1 dynamic range, JPEG has 255:1

How can we compress the scene's dynamic range?

Scale linearly?

- If we scaled linearly from 100,000:1 to 255:1, everything but the sun would be black!

After a slide by Frédo Durand

Global tone mapping operators

Gamma compression, applied independently on R,G,B

- output = e · input$^{\gamma}$ (γ = 0.5 here)

Colors become washed-out.

- Why?

In addition to the gamma transform during RAW-to-JPEG conversion

After a slide by F. Durand and M. Levoy

Global tone mapping operators

Gamma compression on intensity only

Colors are OK, but details (high-frequency intensity) not

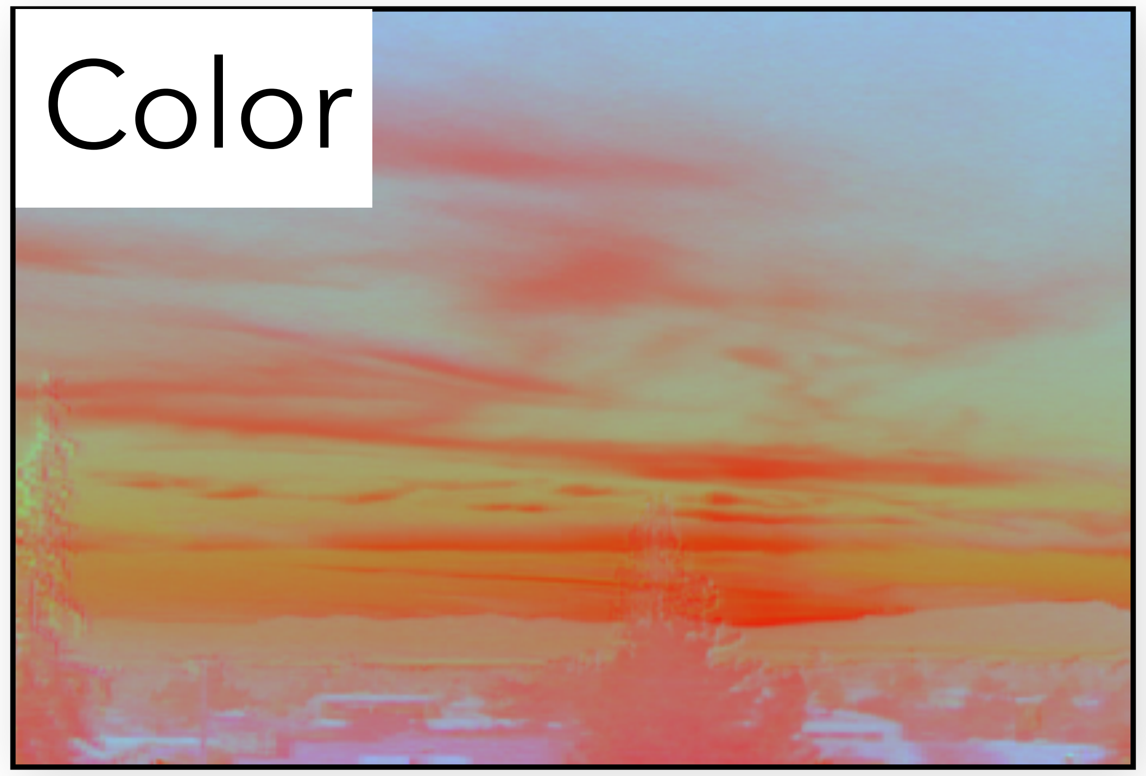

After a slide by Frédo Durand

The importance of local contrast

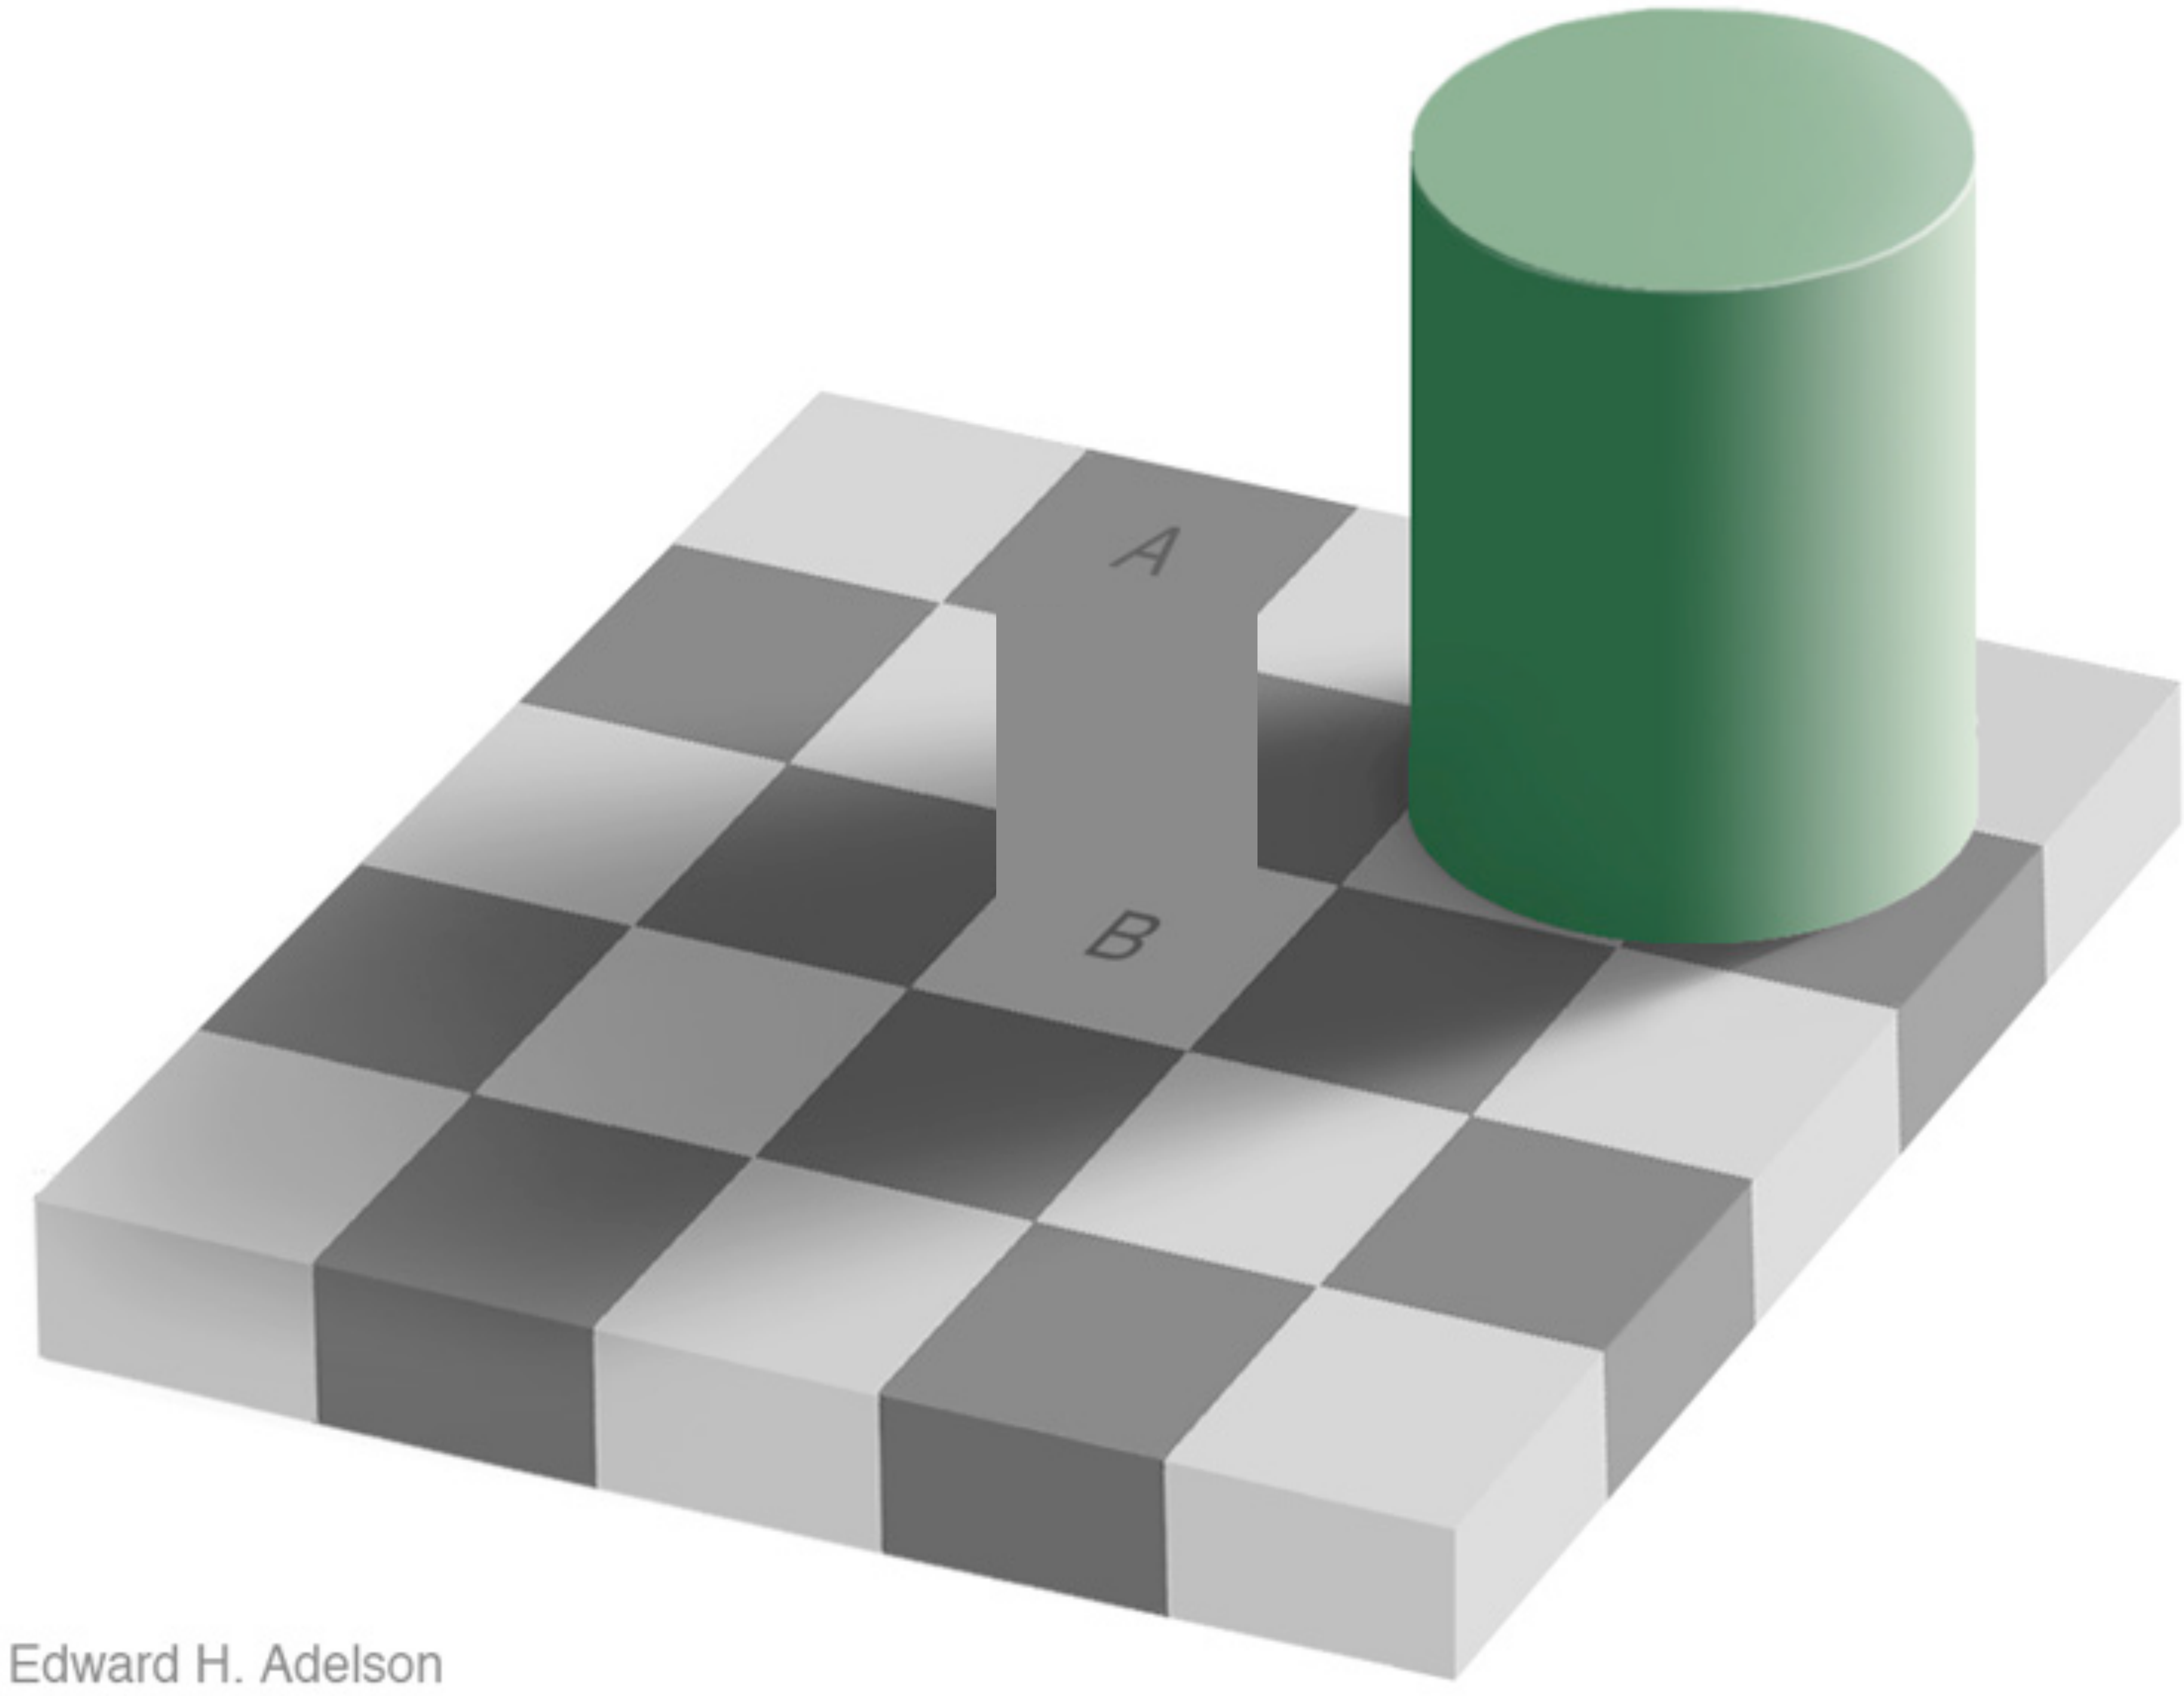

Purposes of tone mapping

Technical:

- fitting a wide range of values into a small space while preserving differences between values as much as possible

Artistic

- reproduce what the photographer/artist feels she saw
- stylize the look of a photo

Mach bands

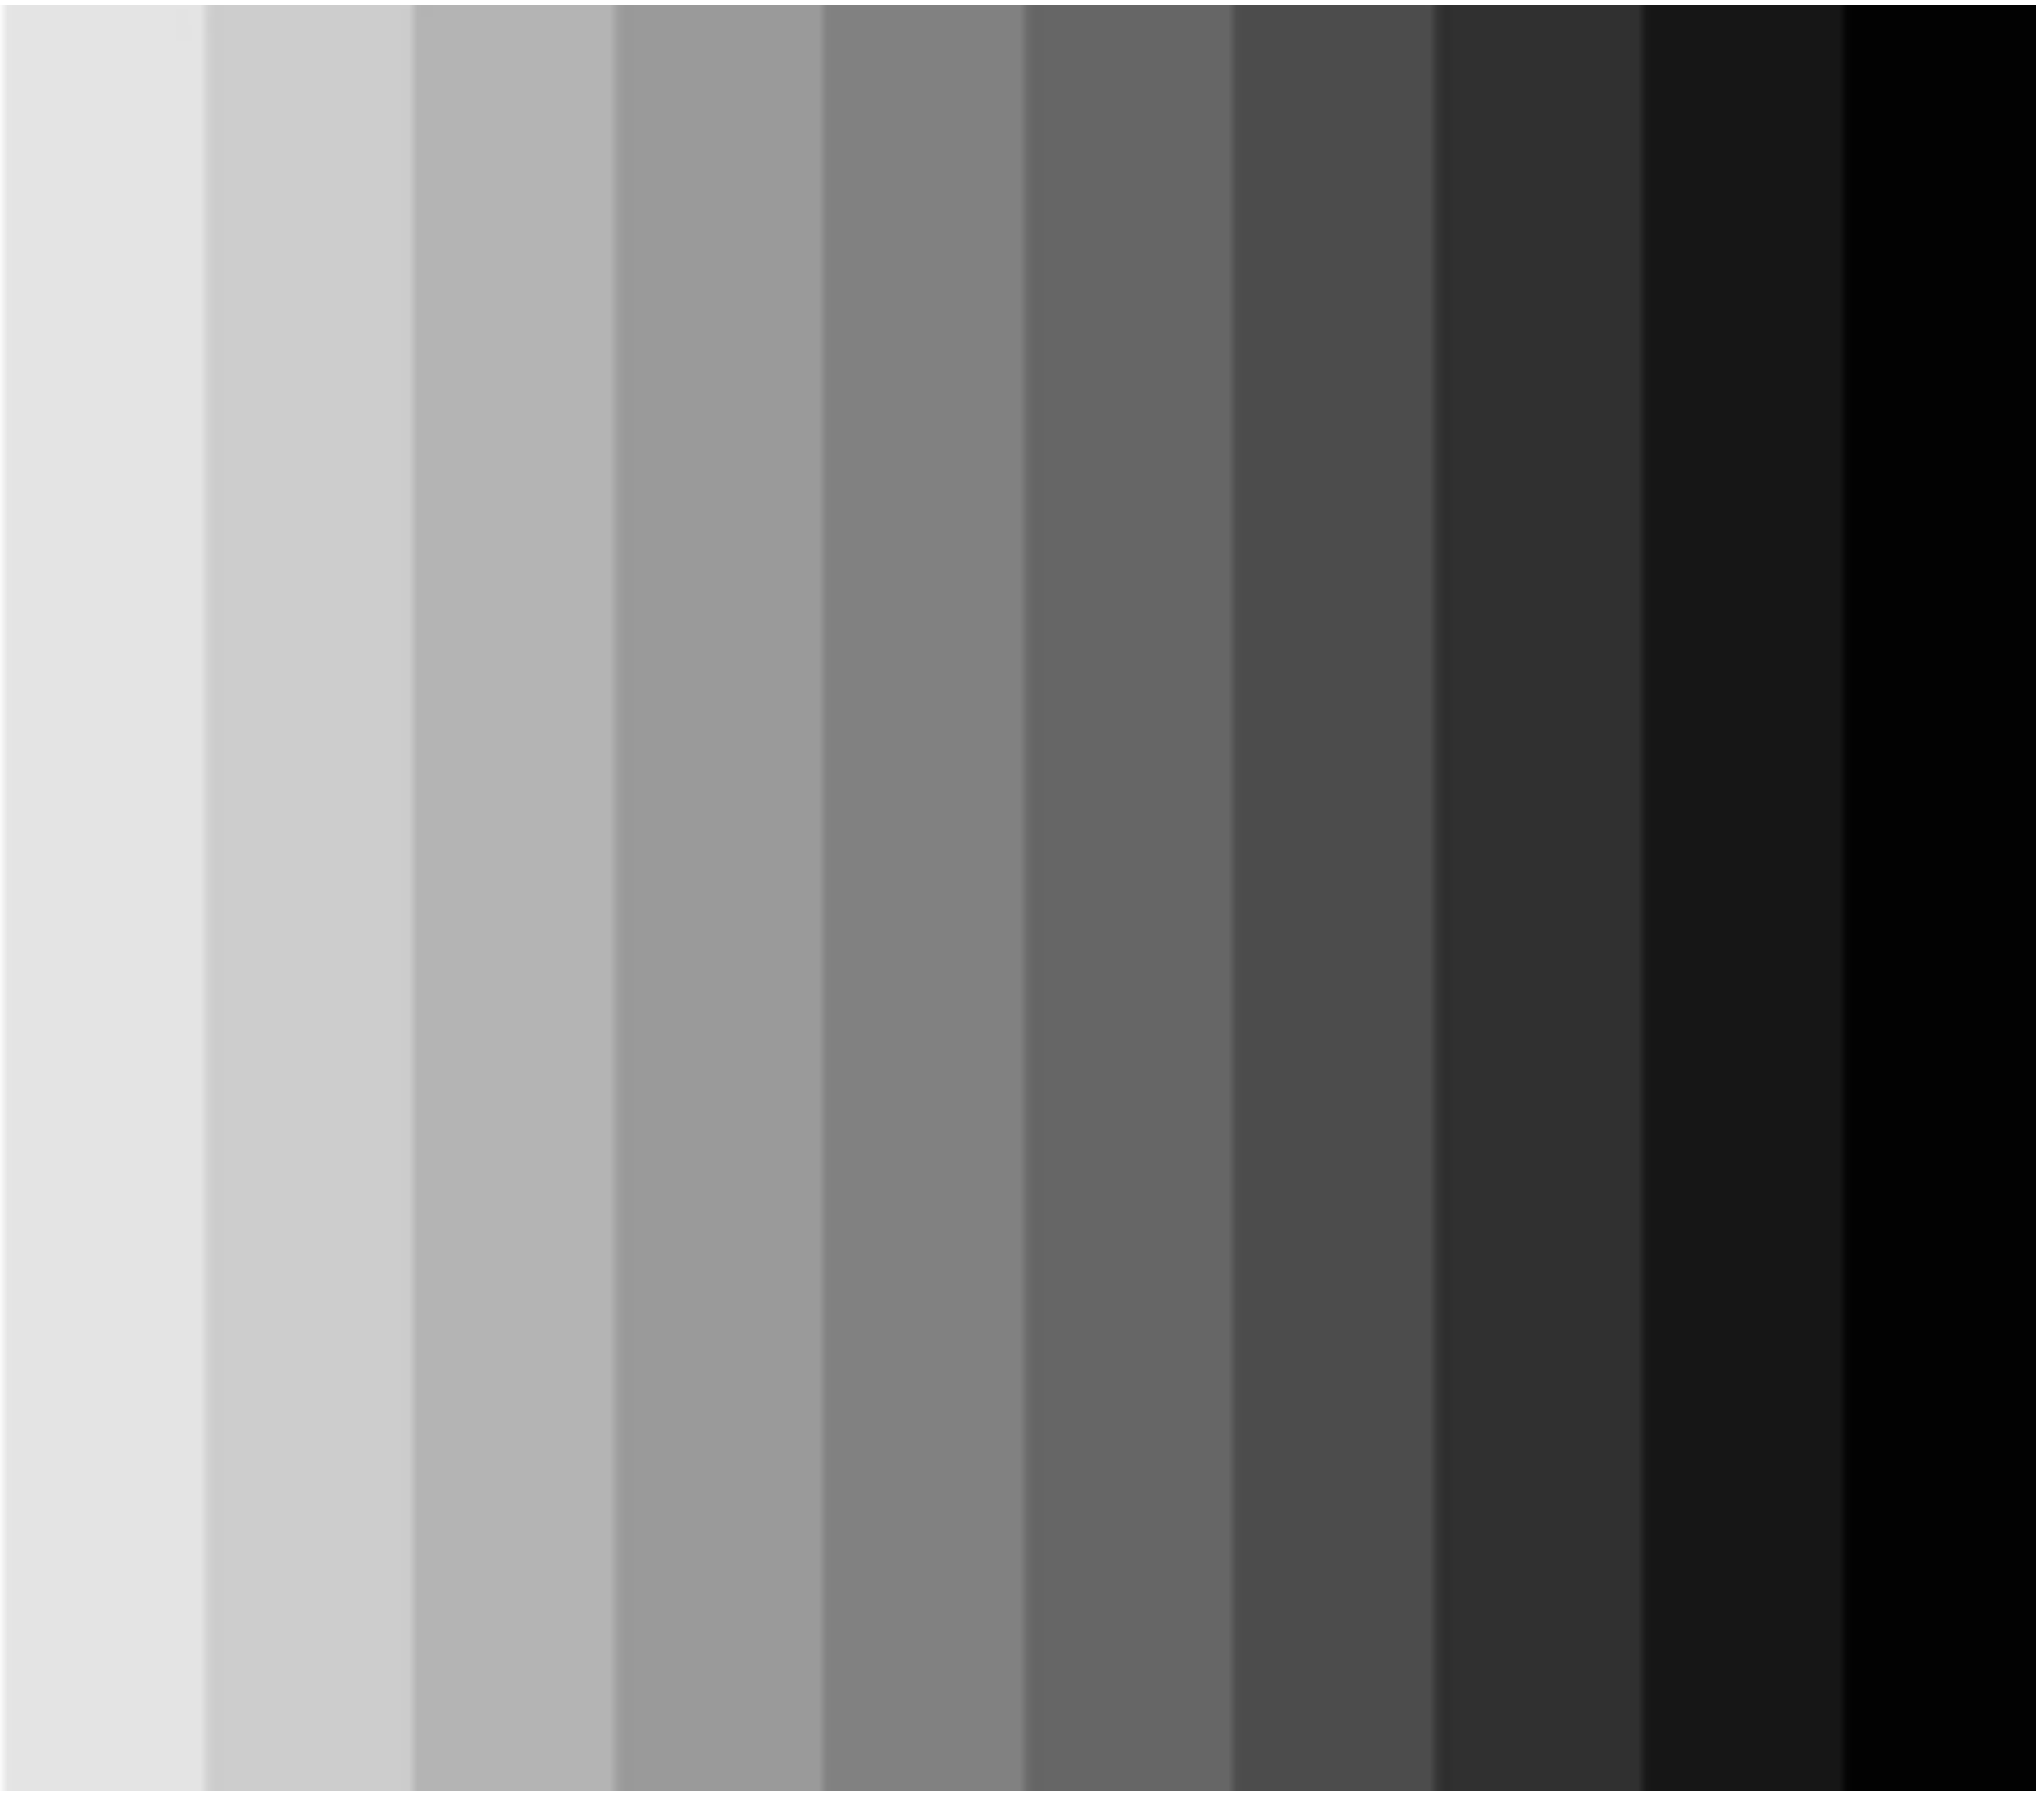

After a slide by Marc Levoy

La Grande Jatte, Georges Seurat, 1884

After a slide by Marc Levoy

Dodging & Burning

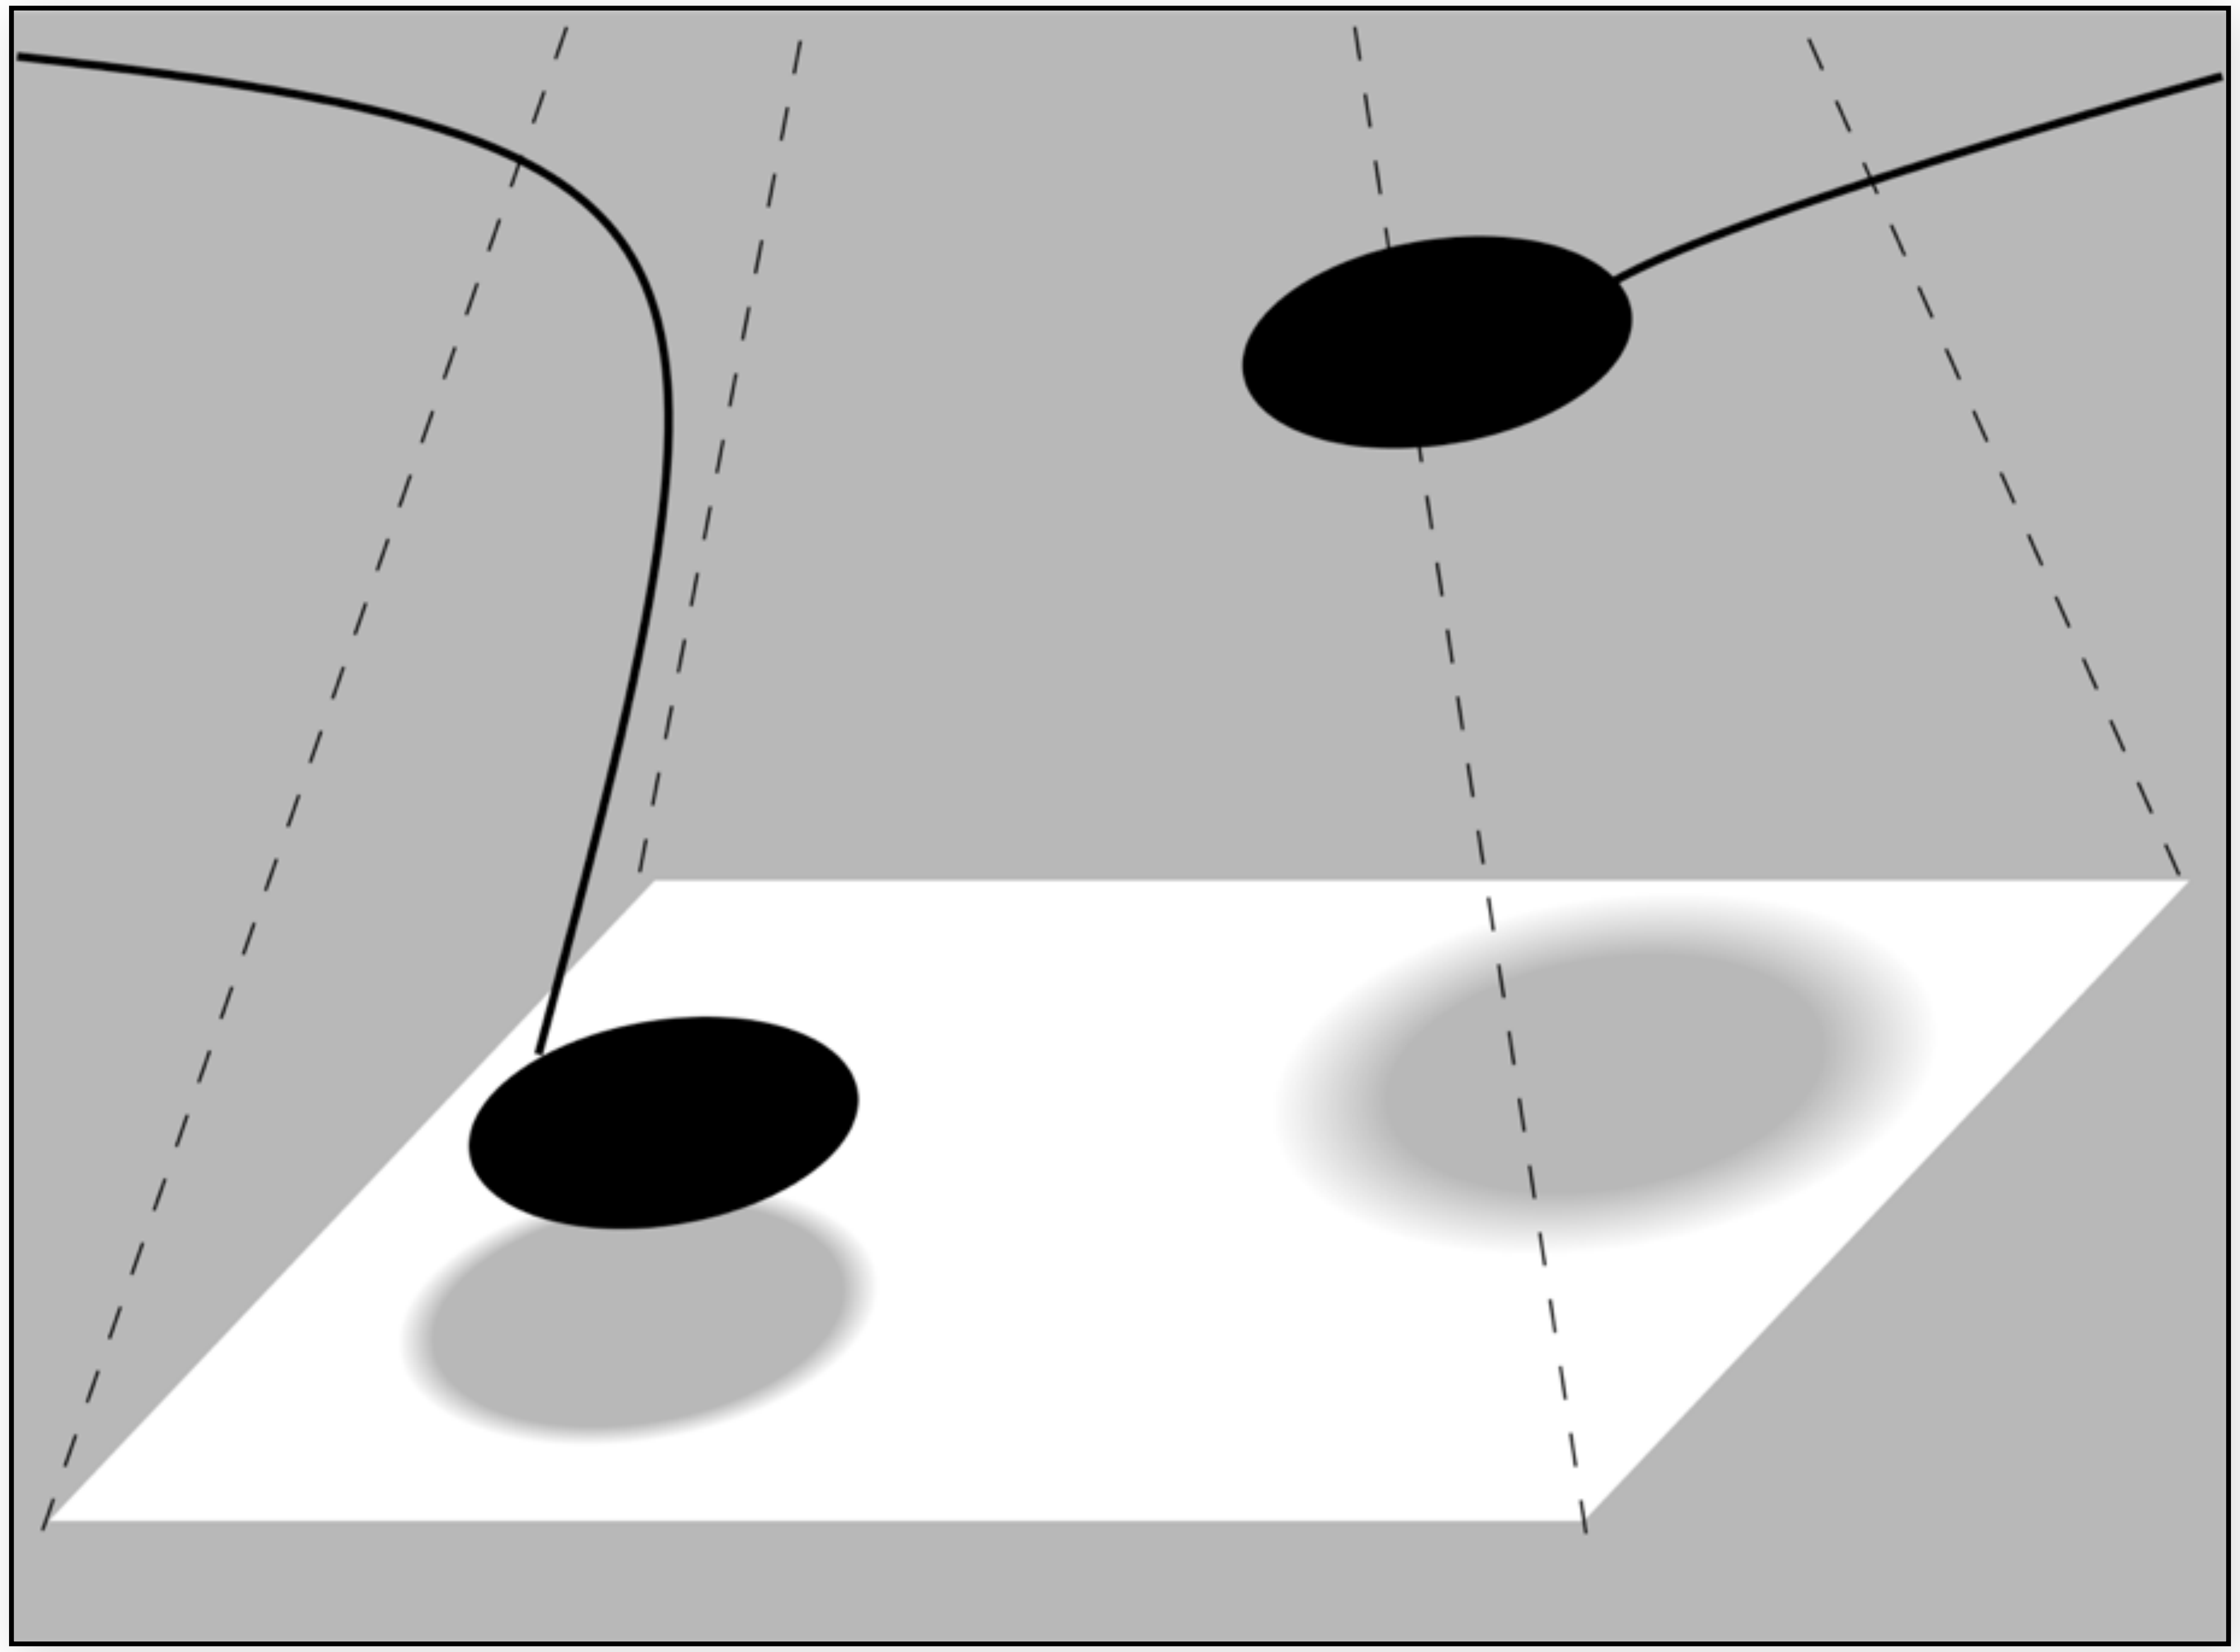

Dodging
(makes print lighter)

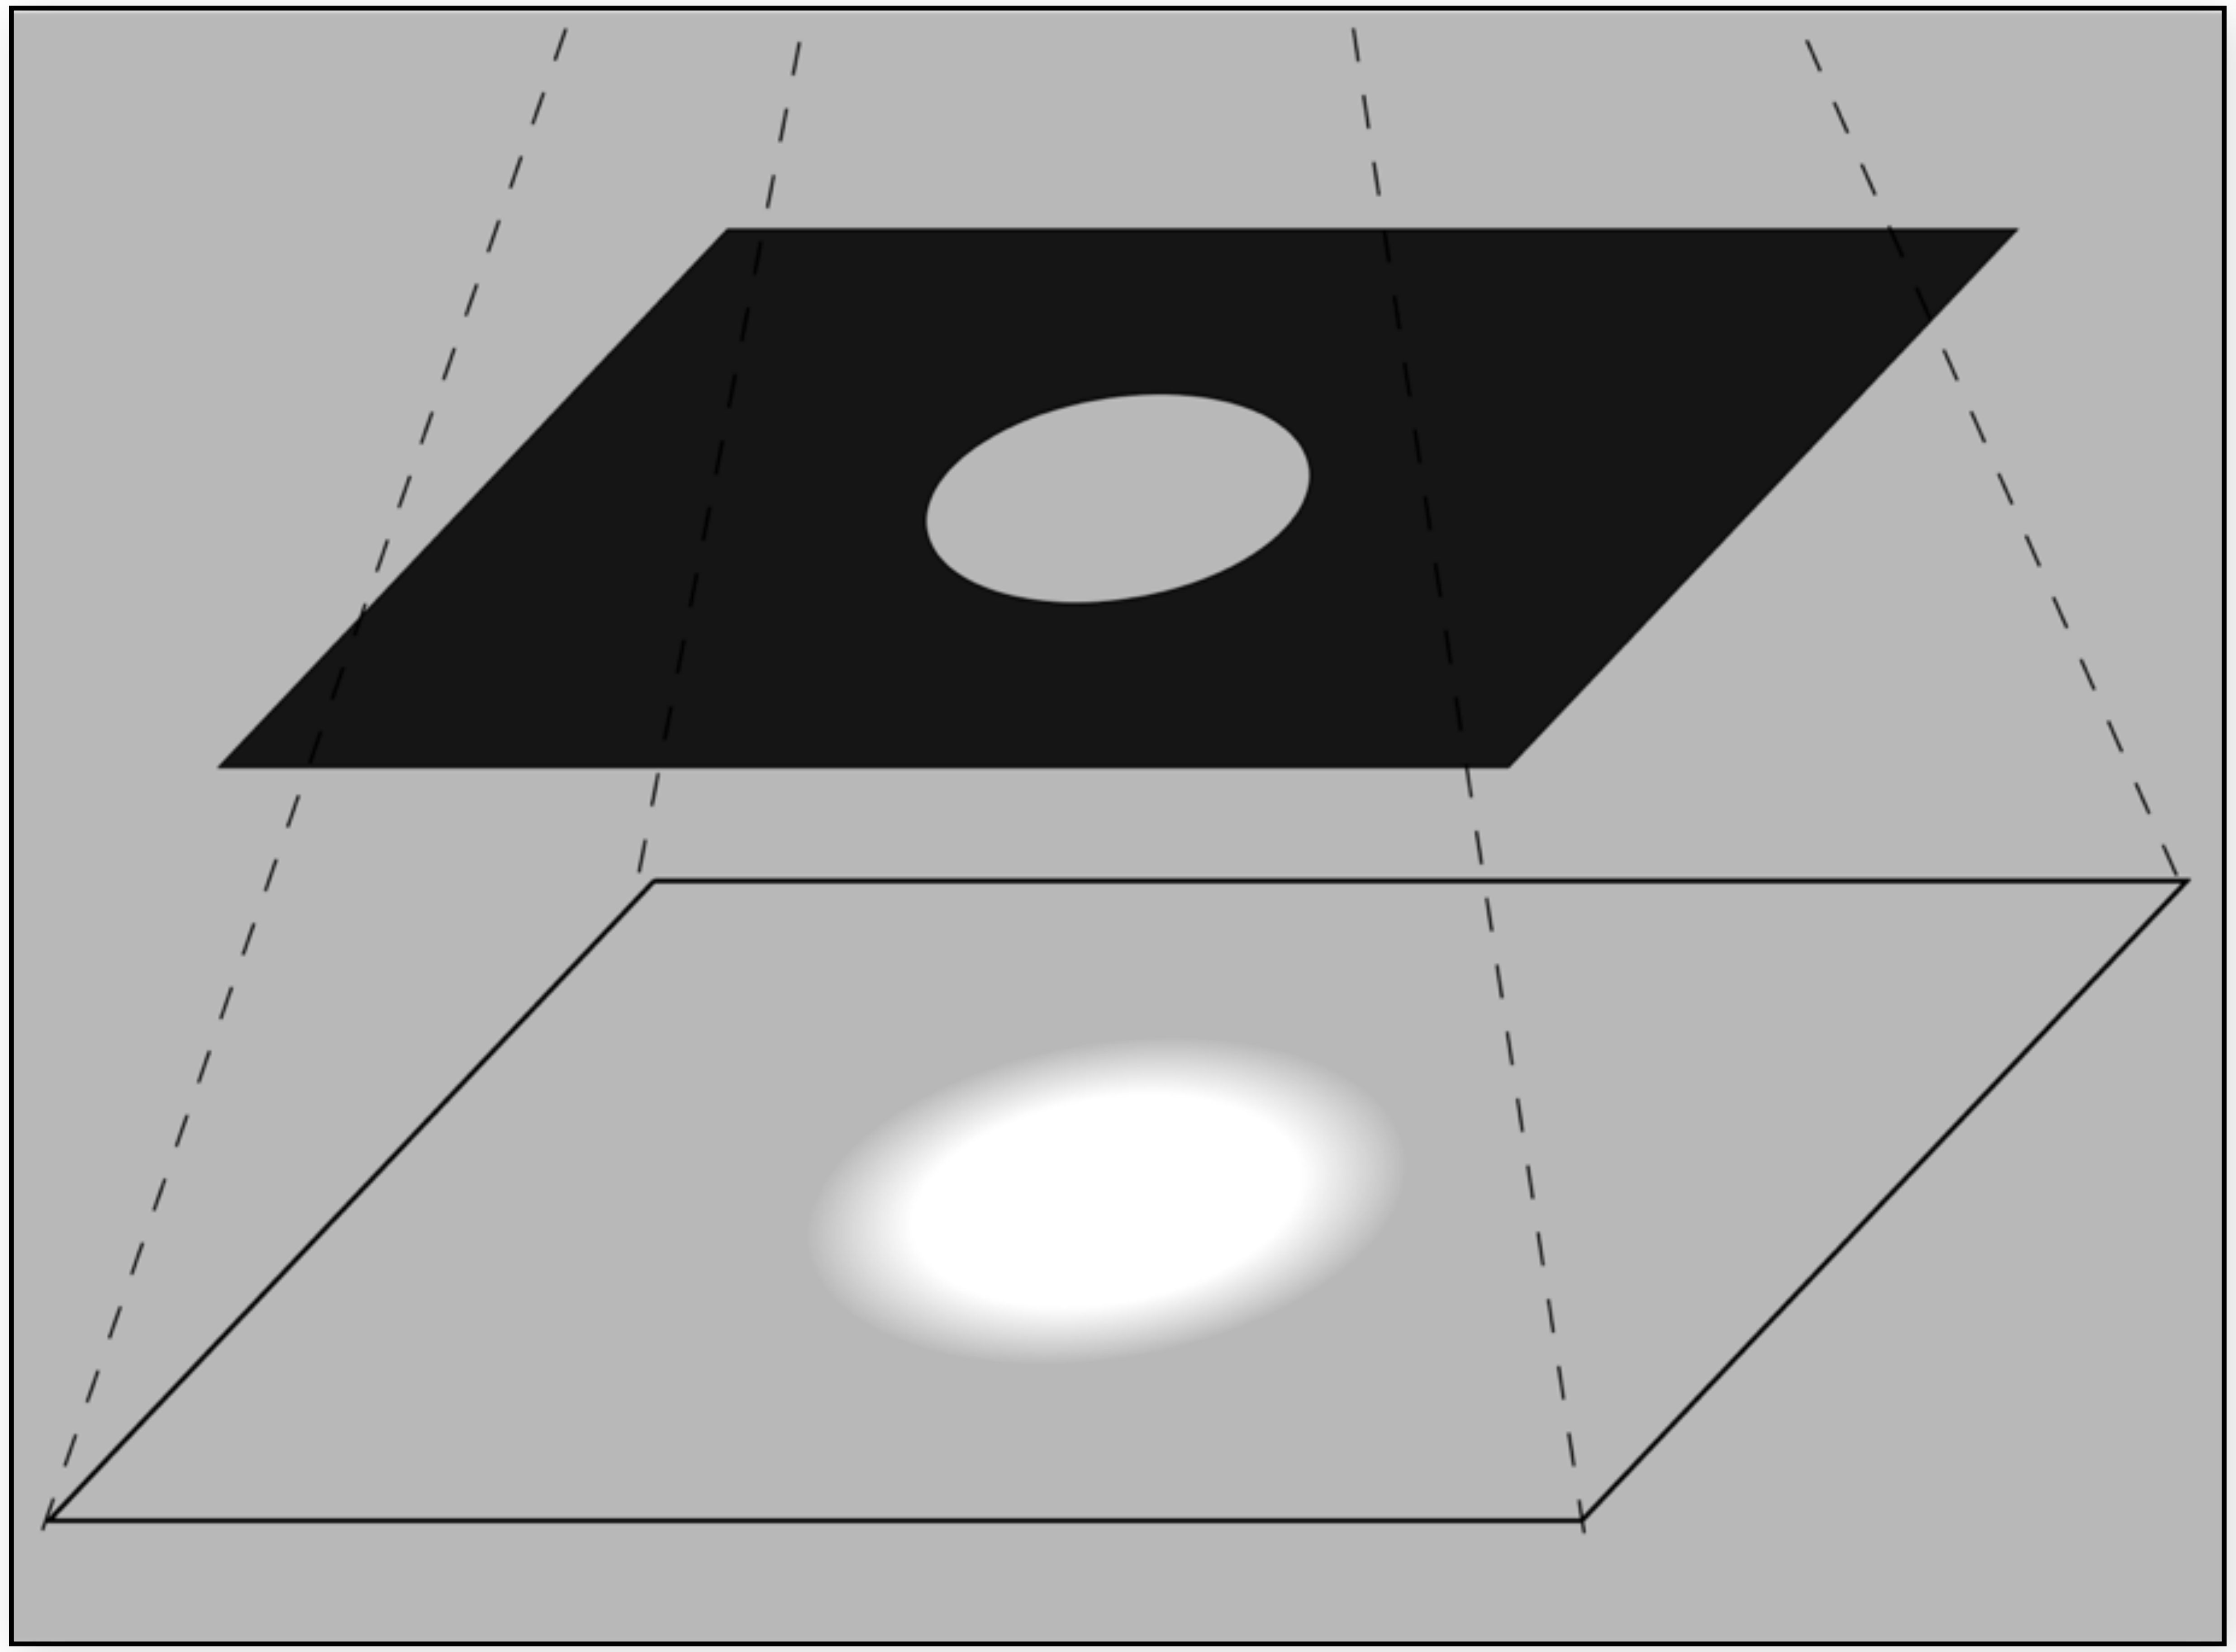

Burning
(makes print darker)

Dodging & Burning

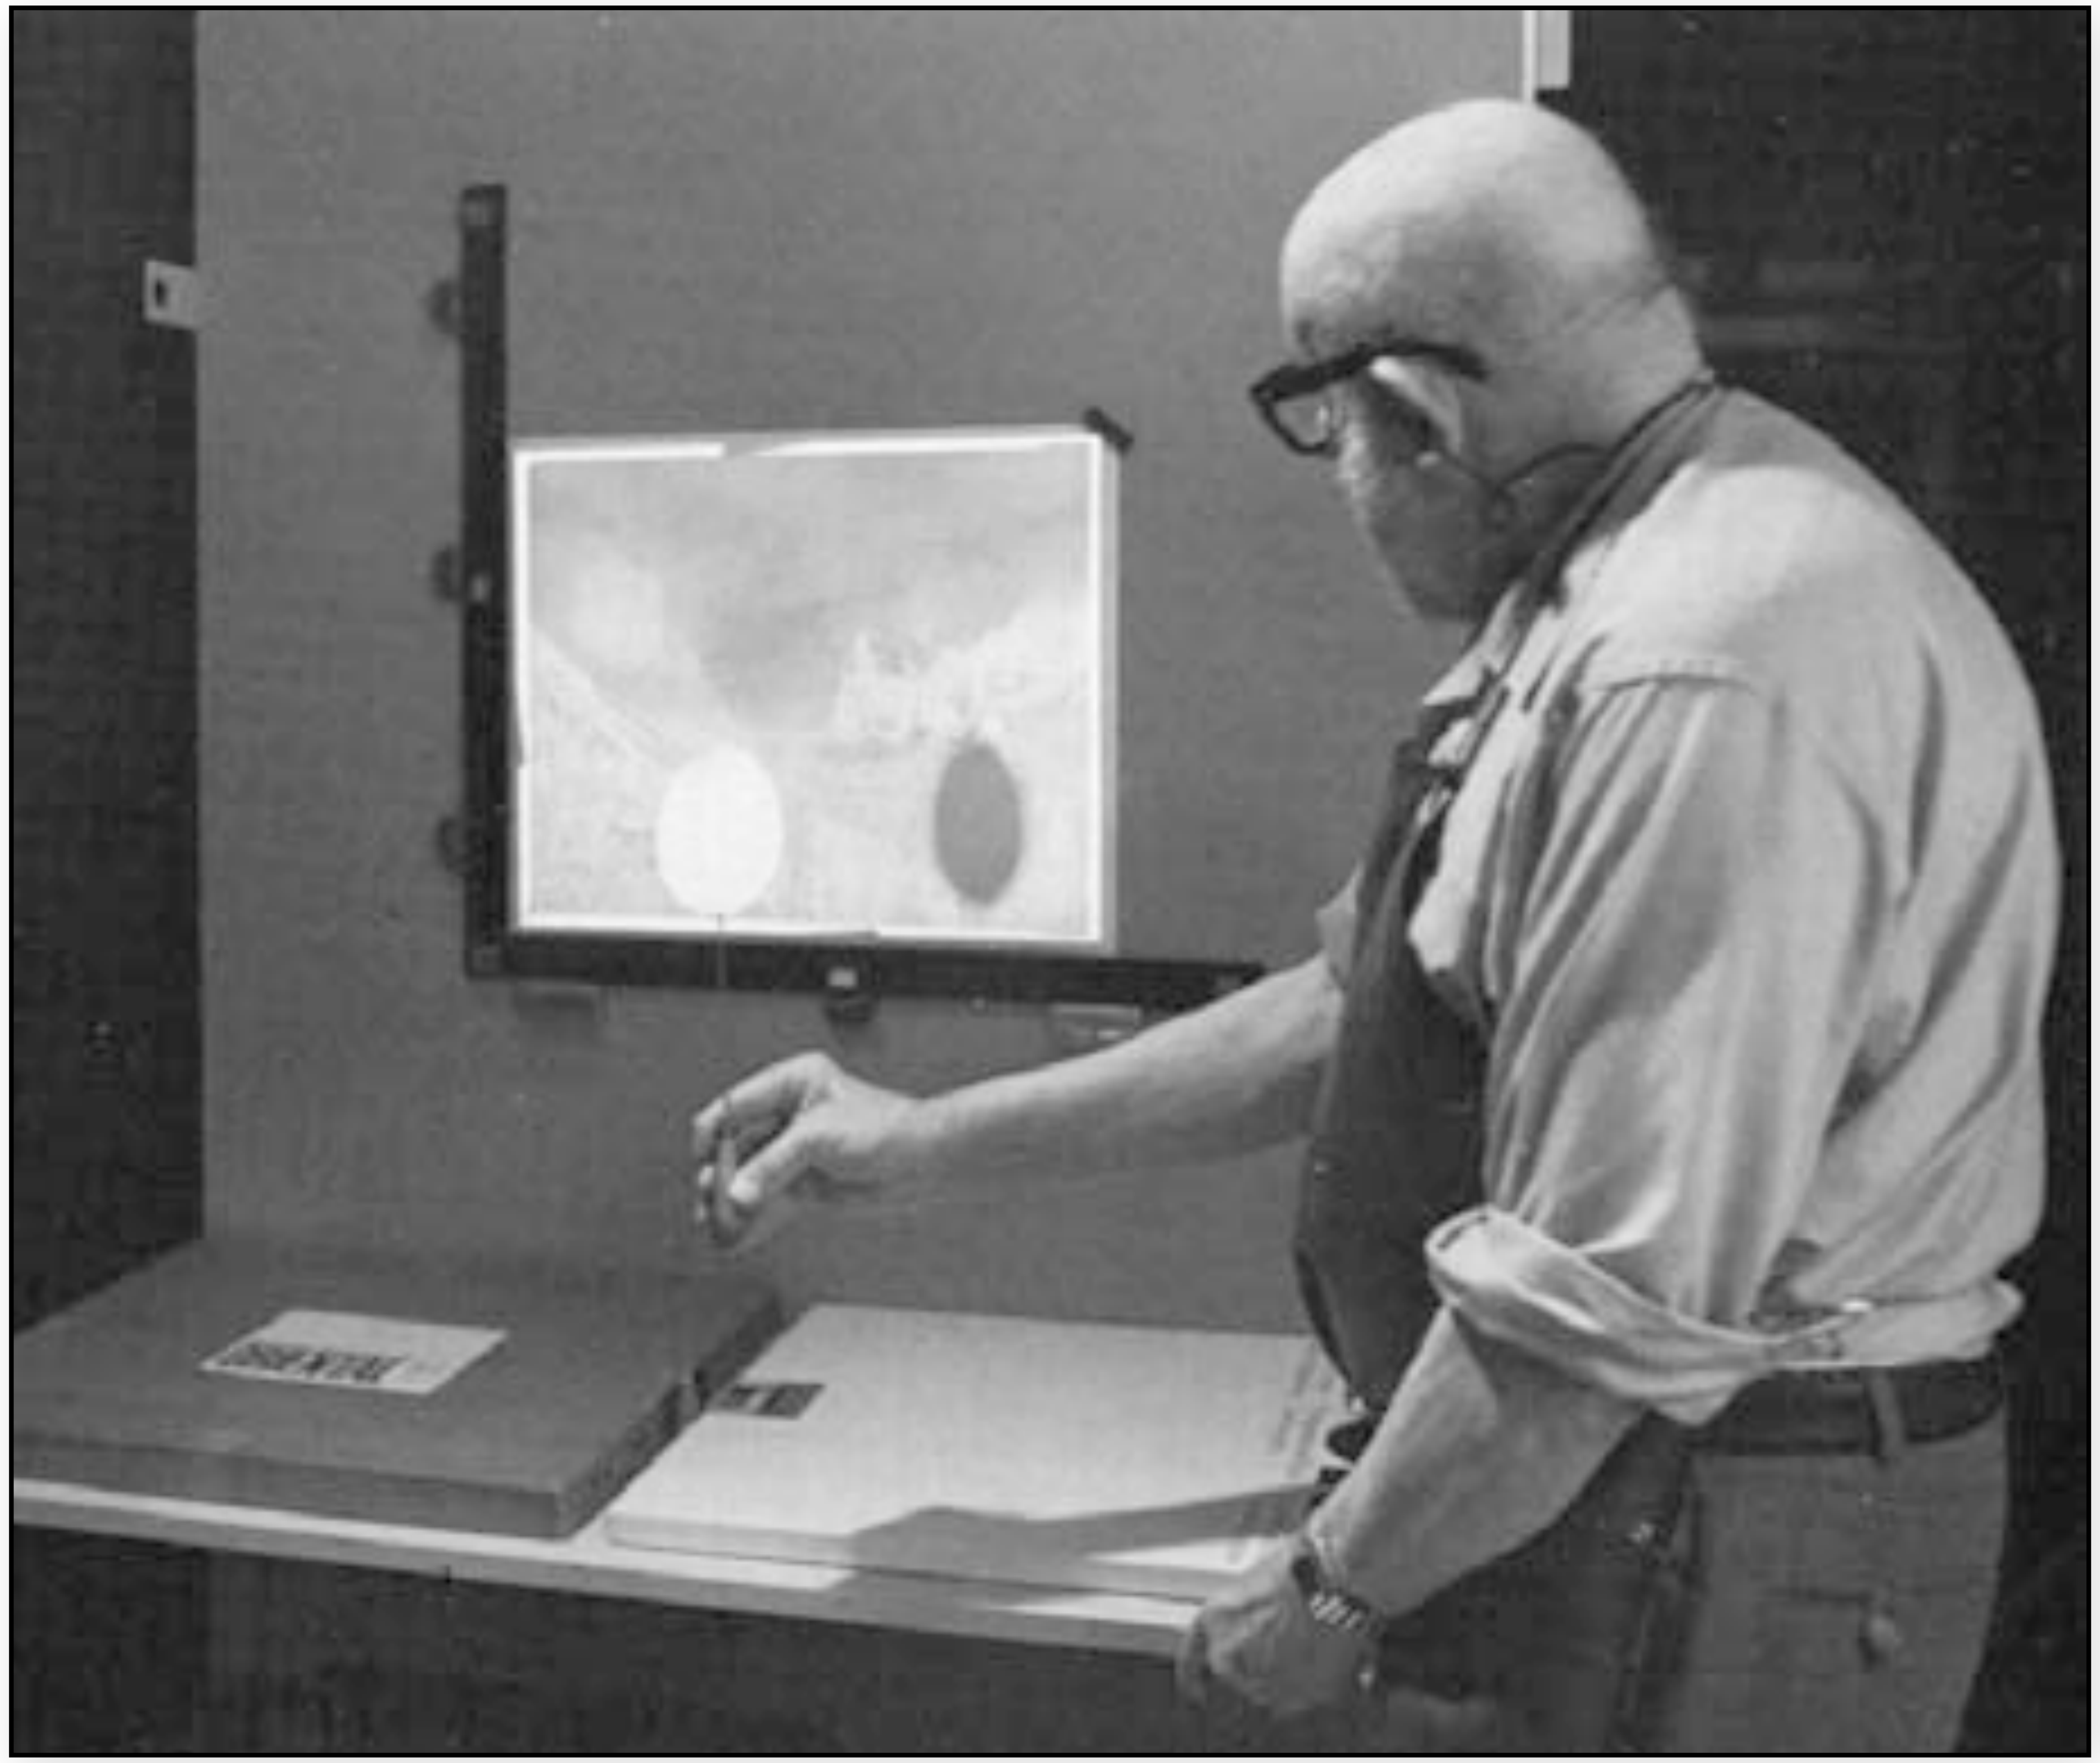
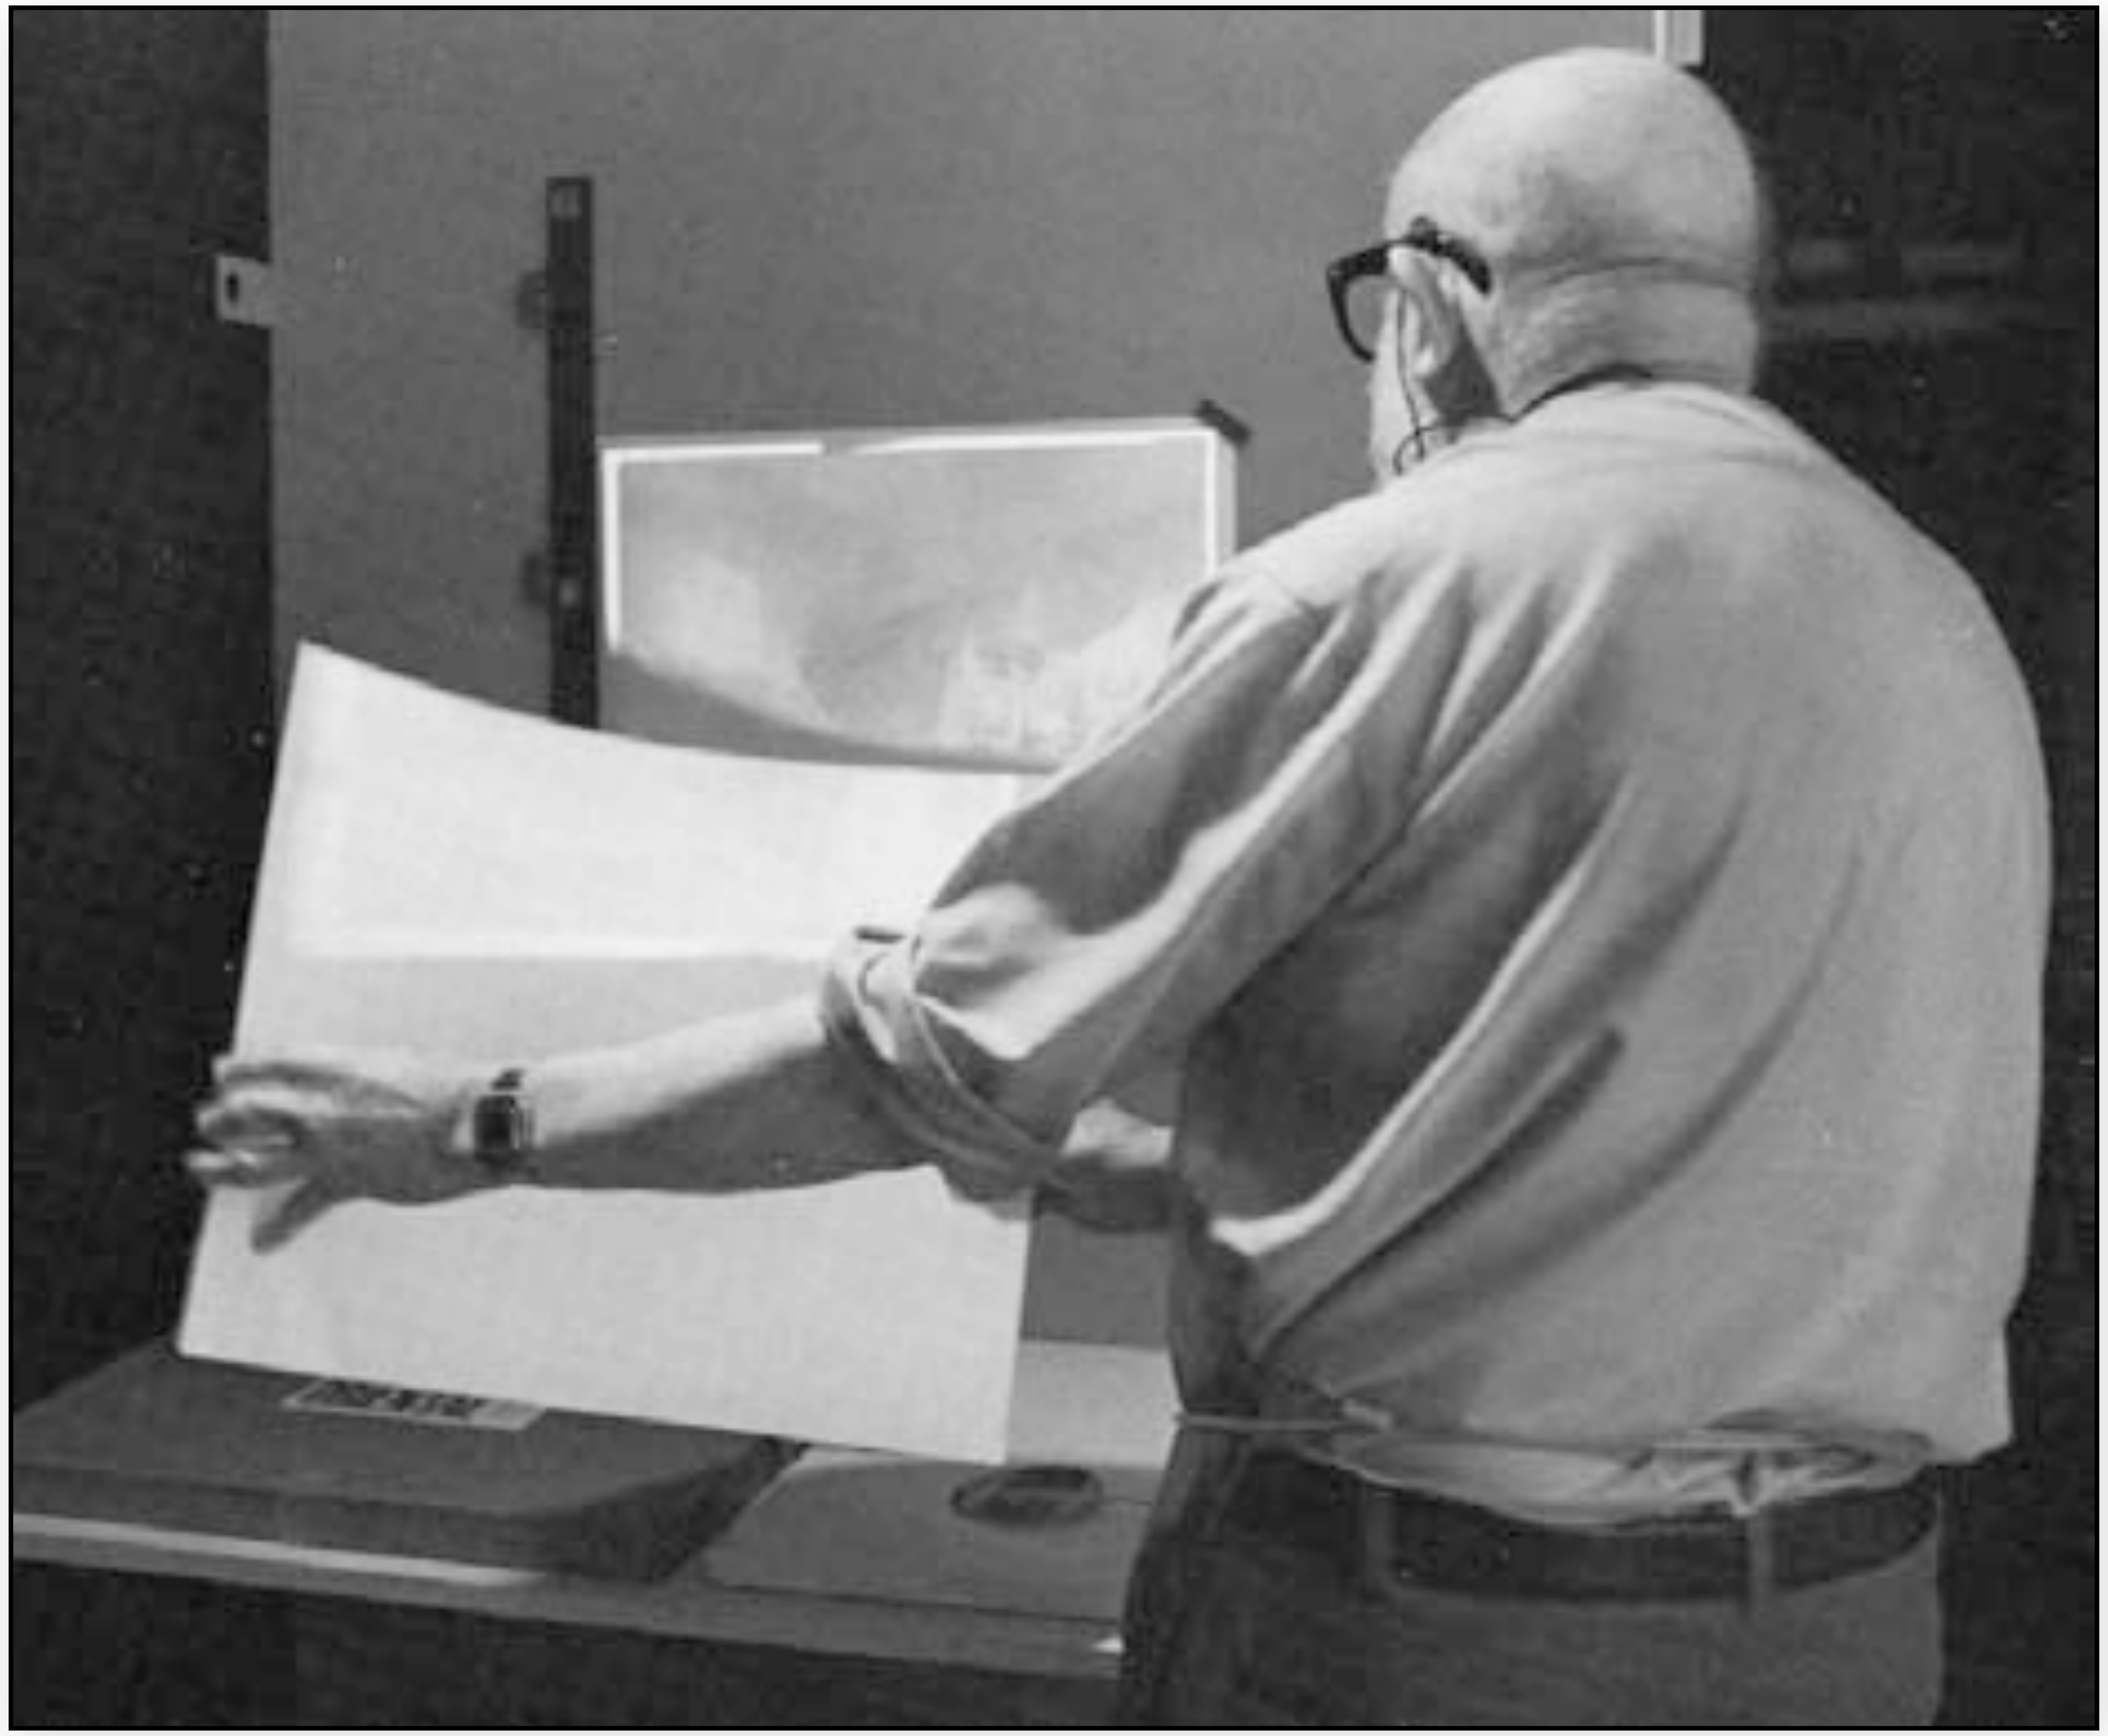

straight
print

Ansel Adams, Clearing Winter Storm, 1942

toned
print

Ansel Adams, Clearing Winter Storm, 1942

Oppenheim 1968, Chiu et al. 1993

Reduce contrast of low-frequencies, preserve high frequencies

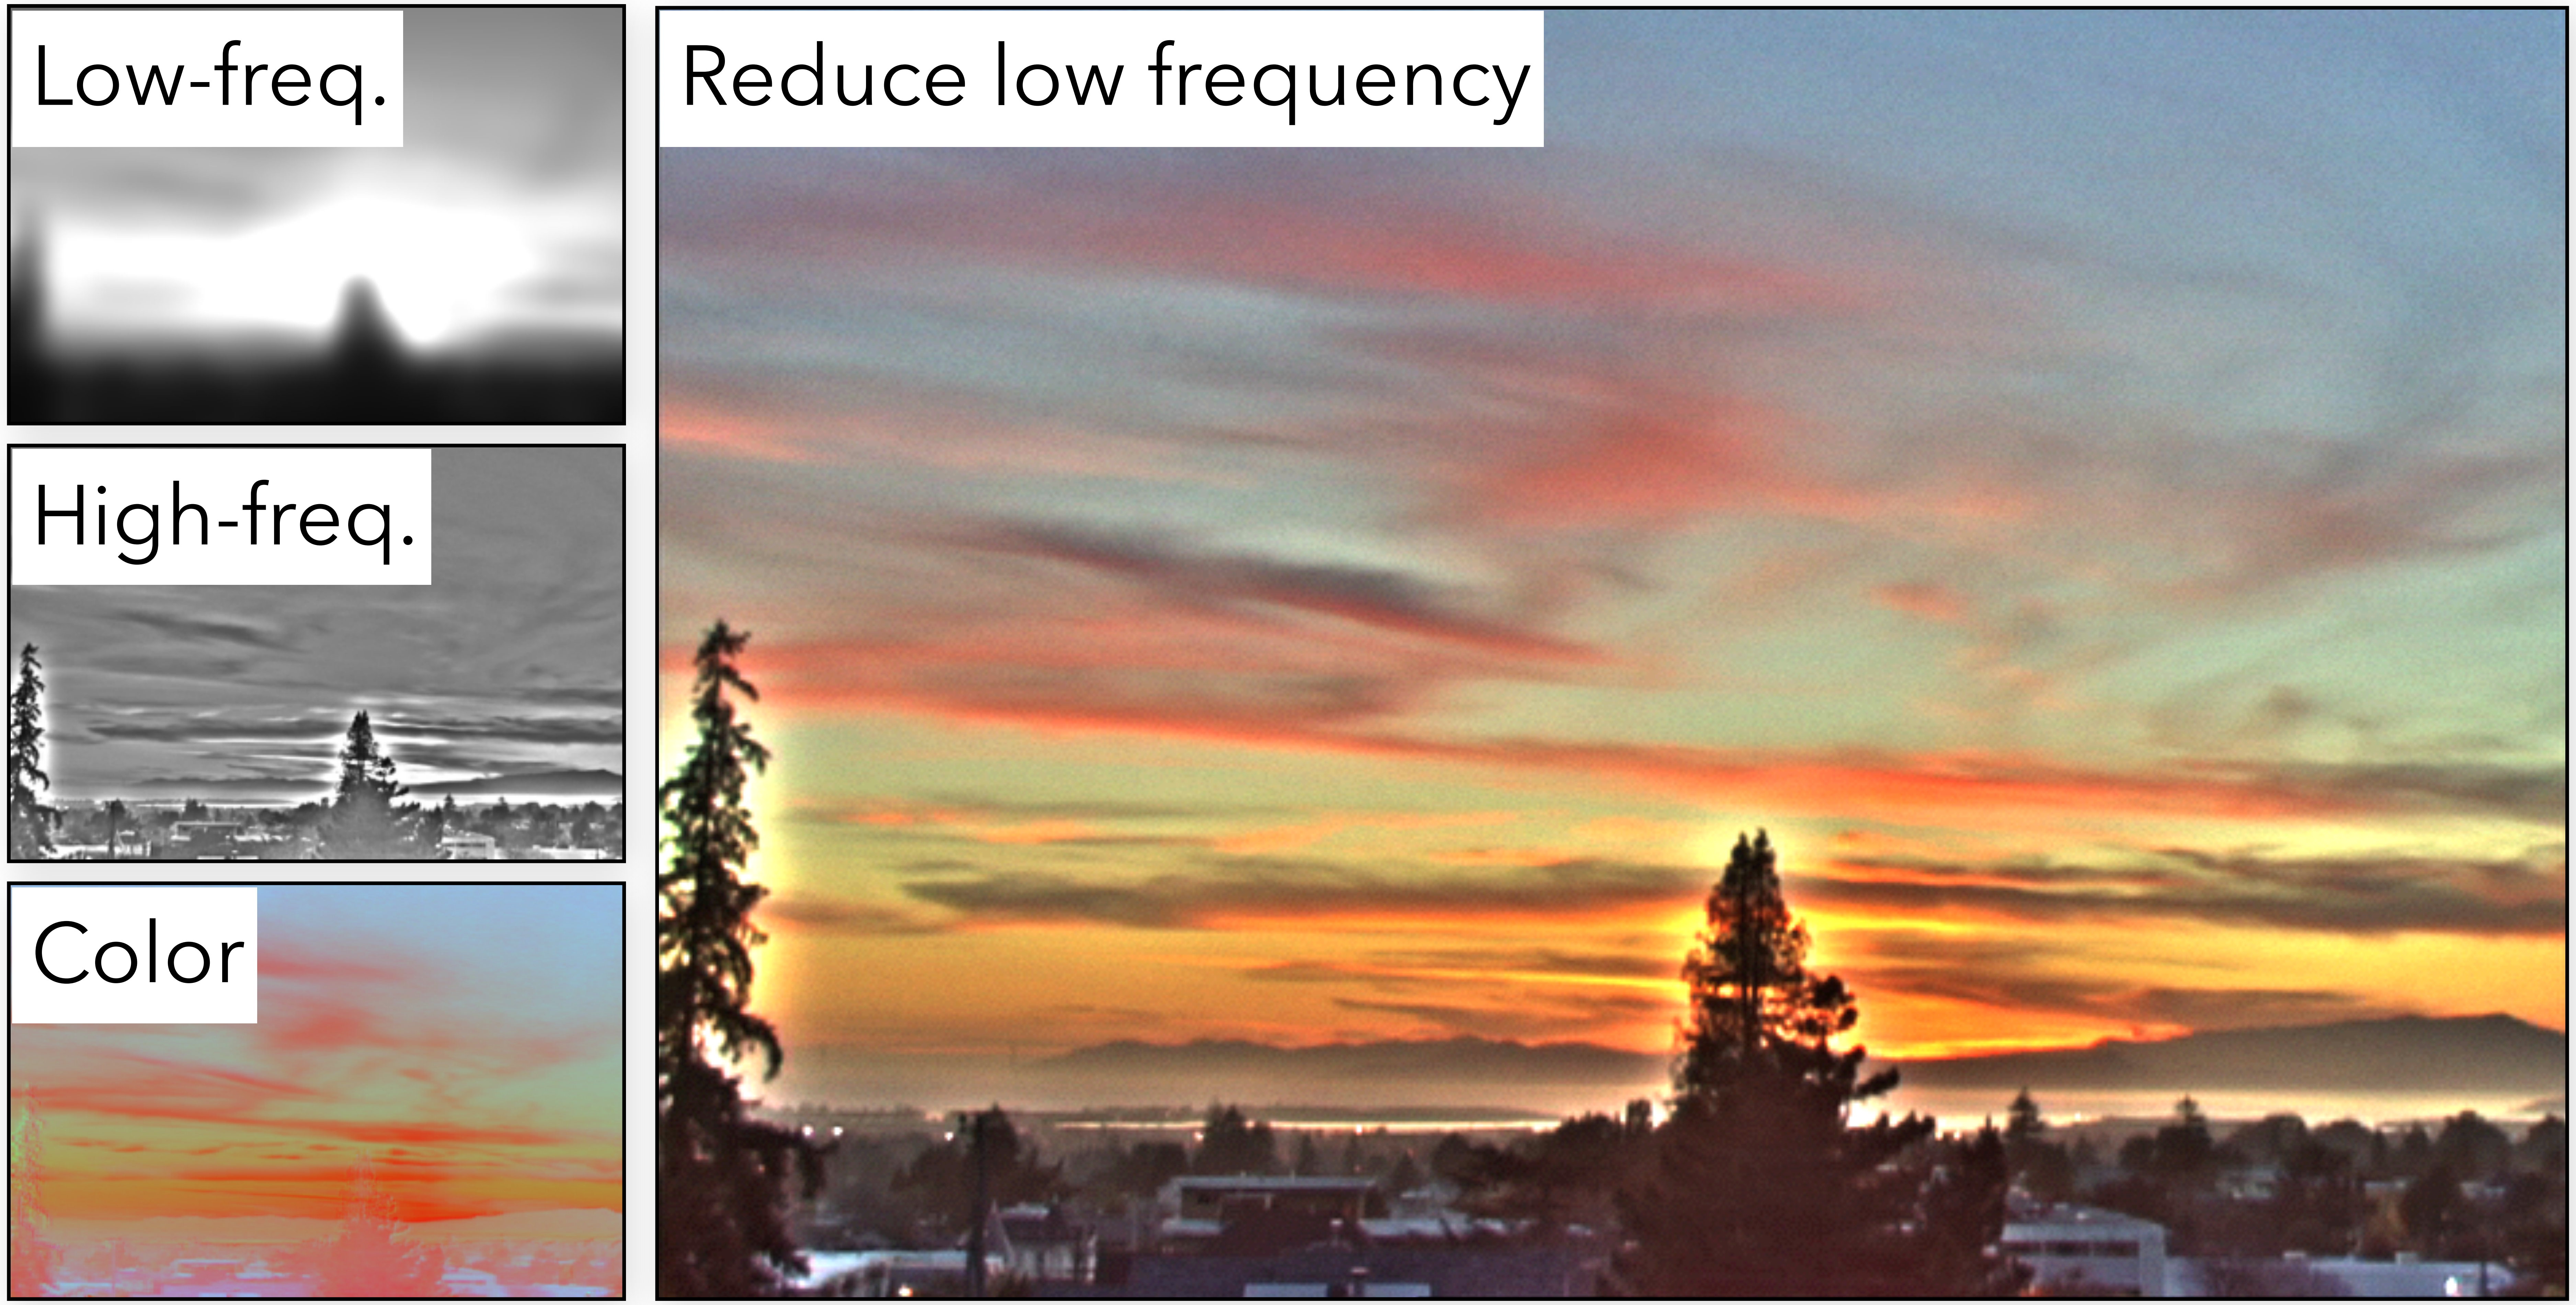

After a slide by Frédo Durand

Homomorphic filtering

Oppenhein, in the sixties

Images are the product of illumination and albedo

Illumination is usually slow-varying

Perform albedo-illumination separation using low-pass filtering of the log image

http://www.cs.sfu.ca/~stella/papers/blairthesis/main/node33.html

After a slide by Frédo Durand

The halo nightmare

For strong edges; because they contain high frequency

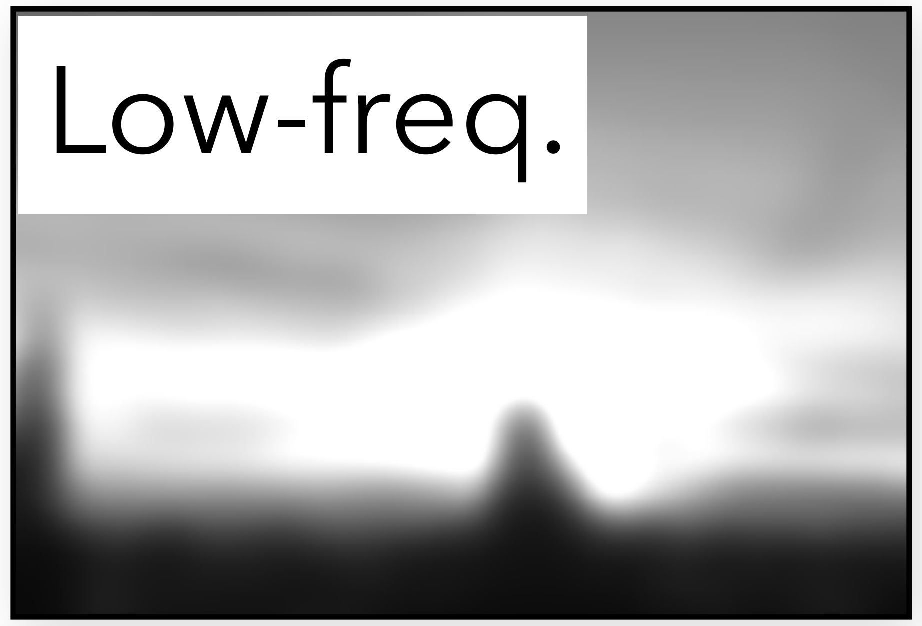

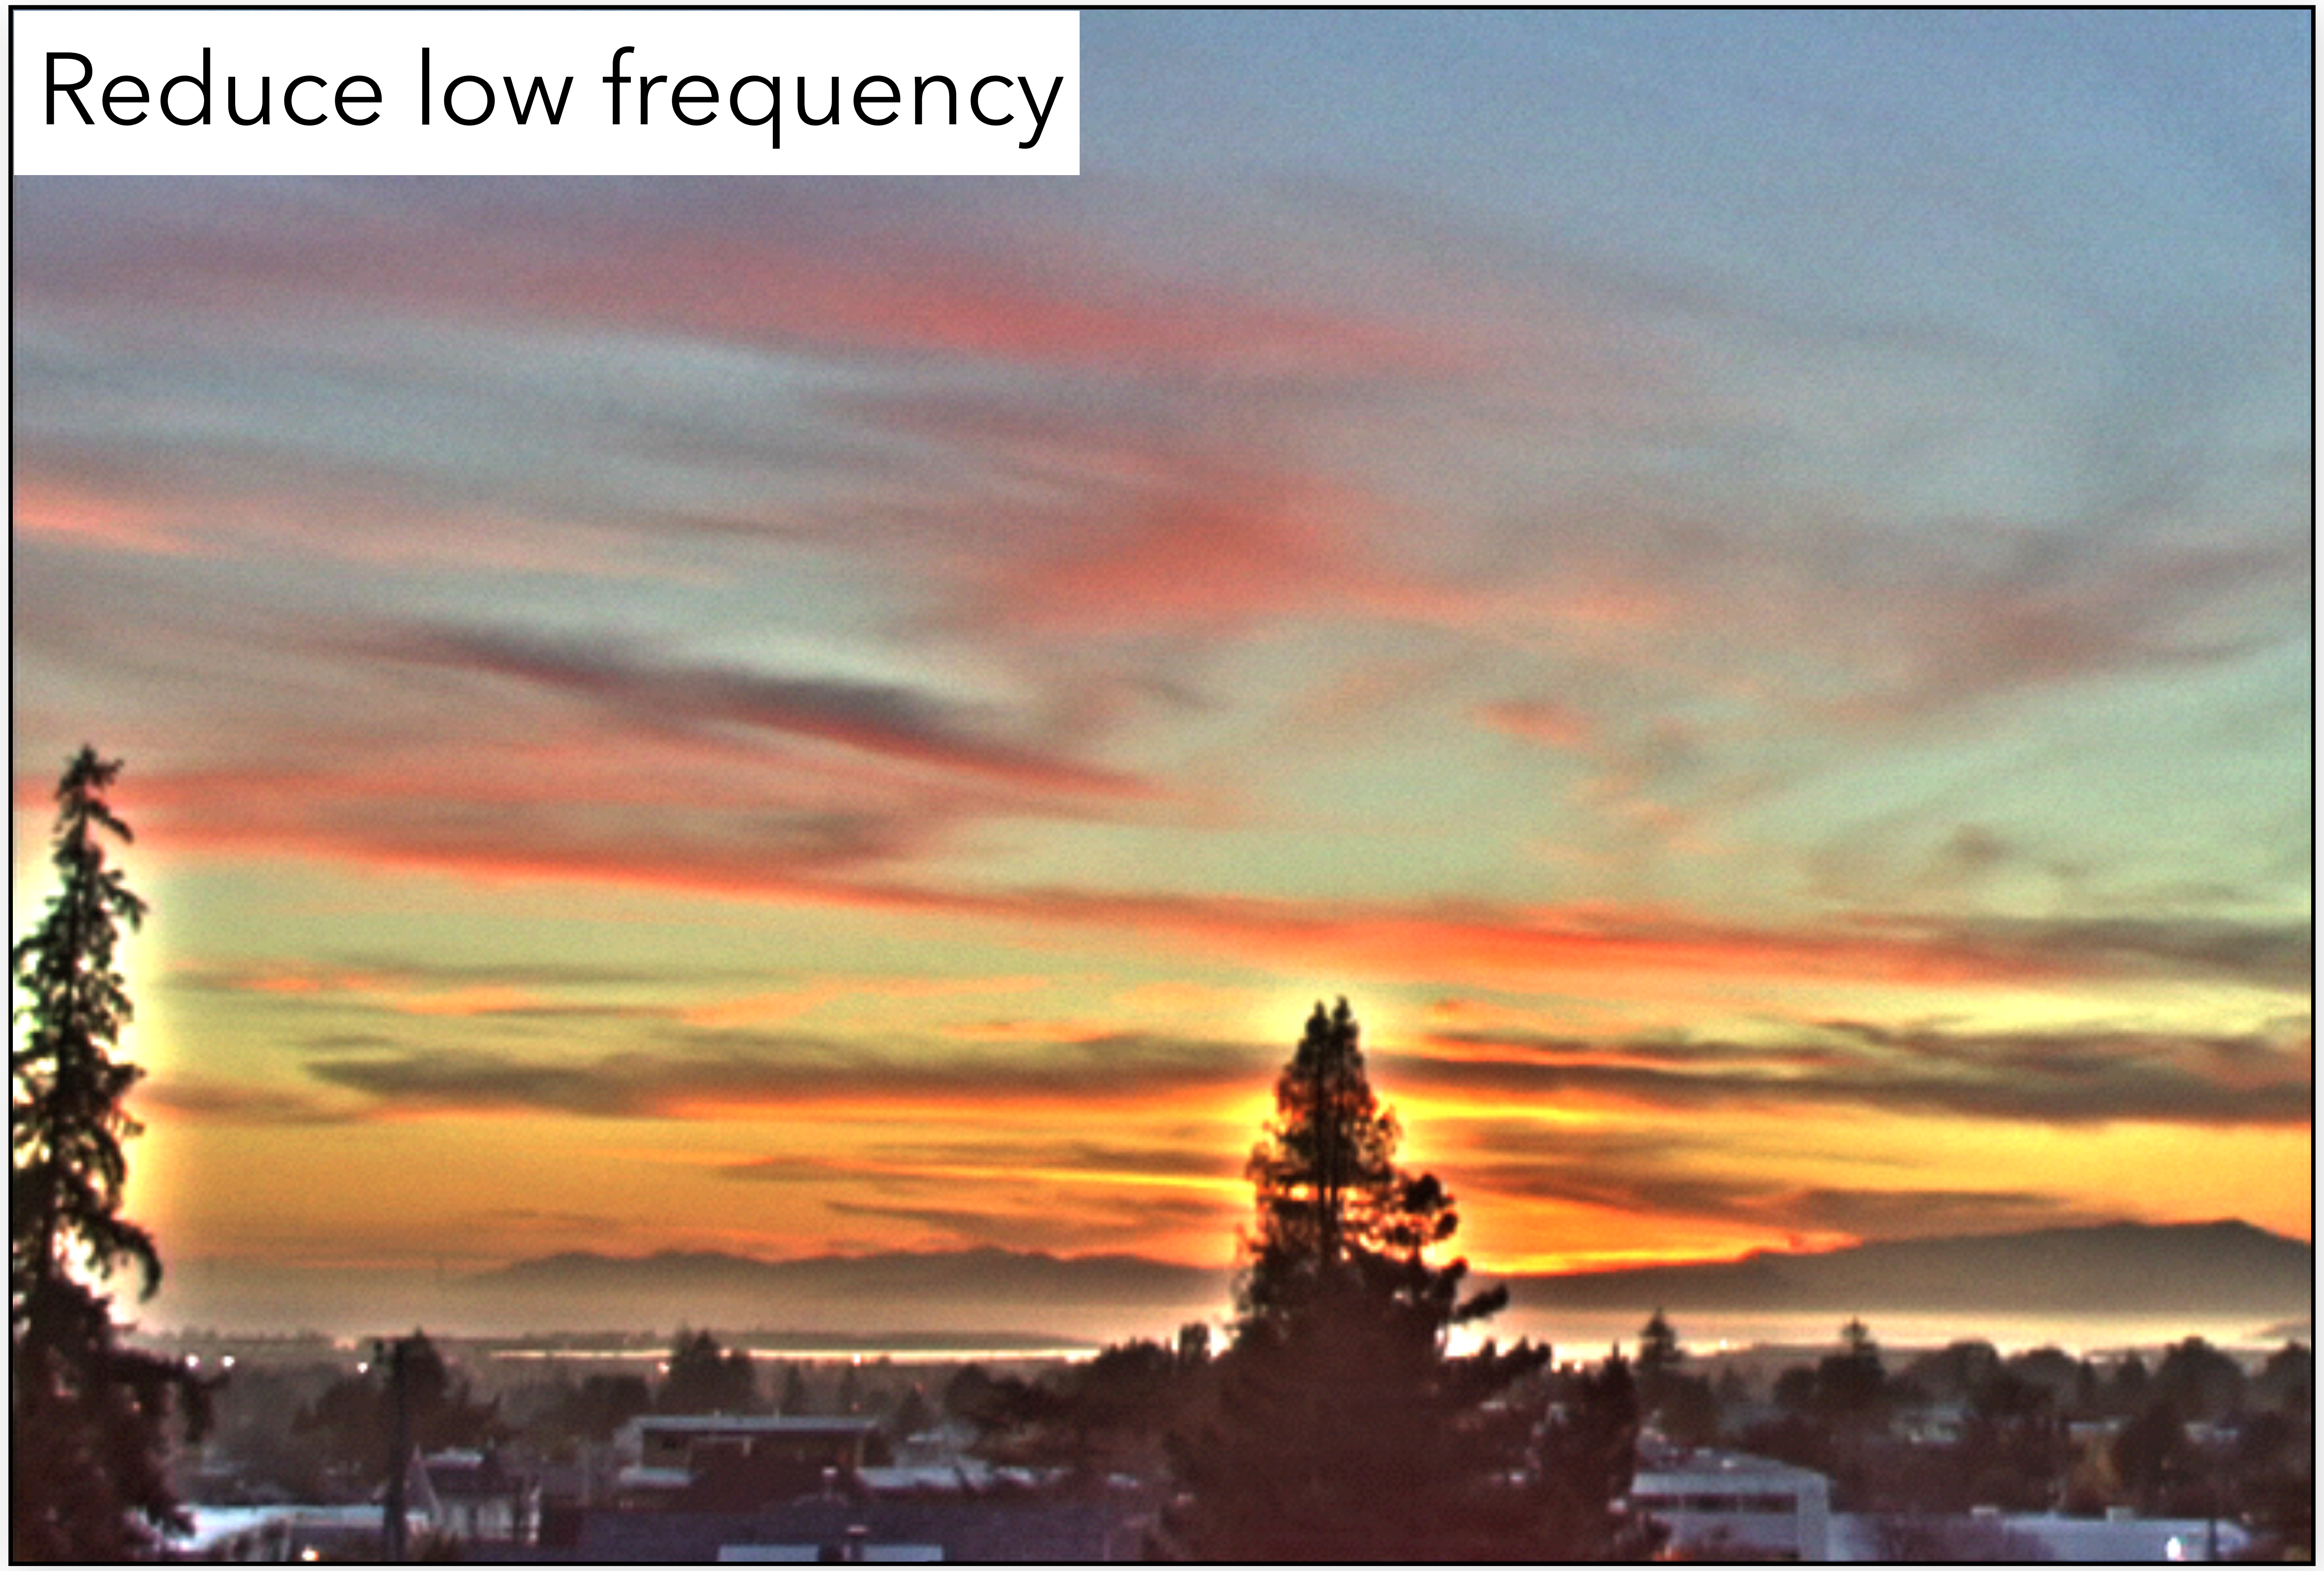

After a slide by Frédo Durand

The halo nightmare

Similar to unsharp mask of luminance in log domain

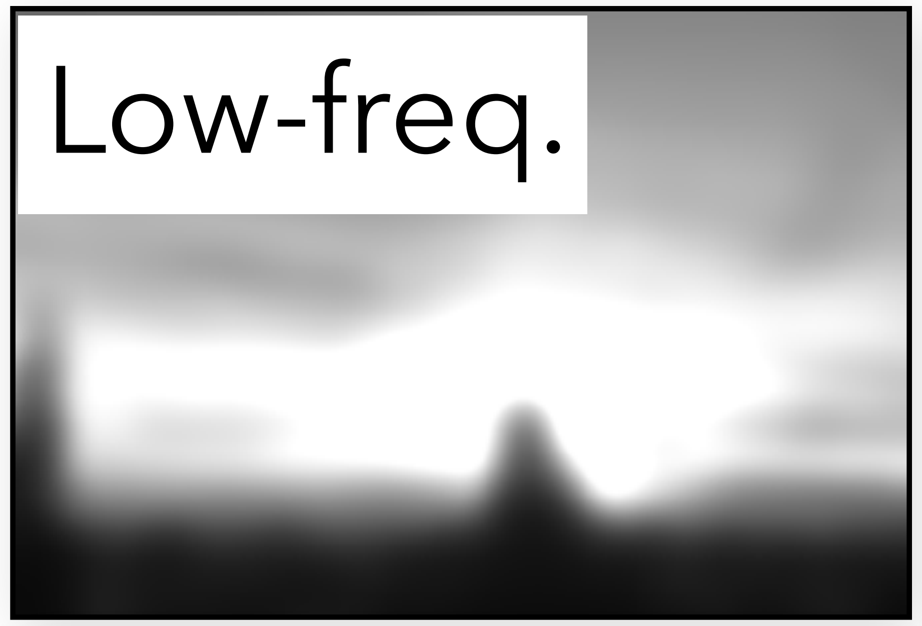

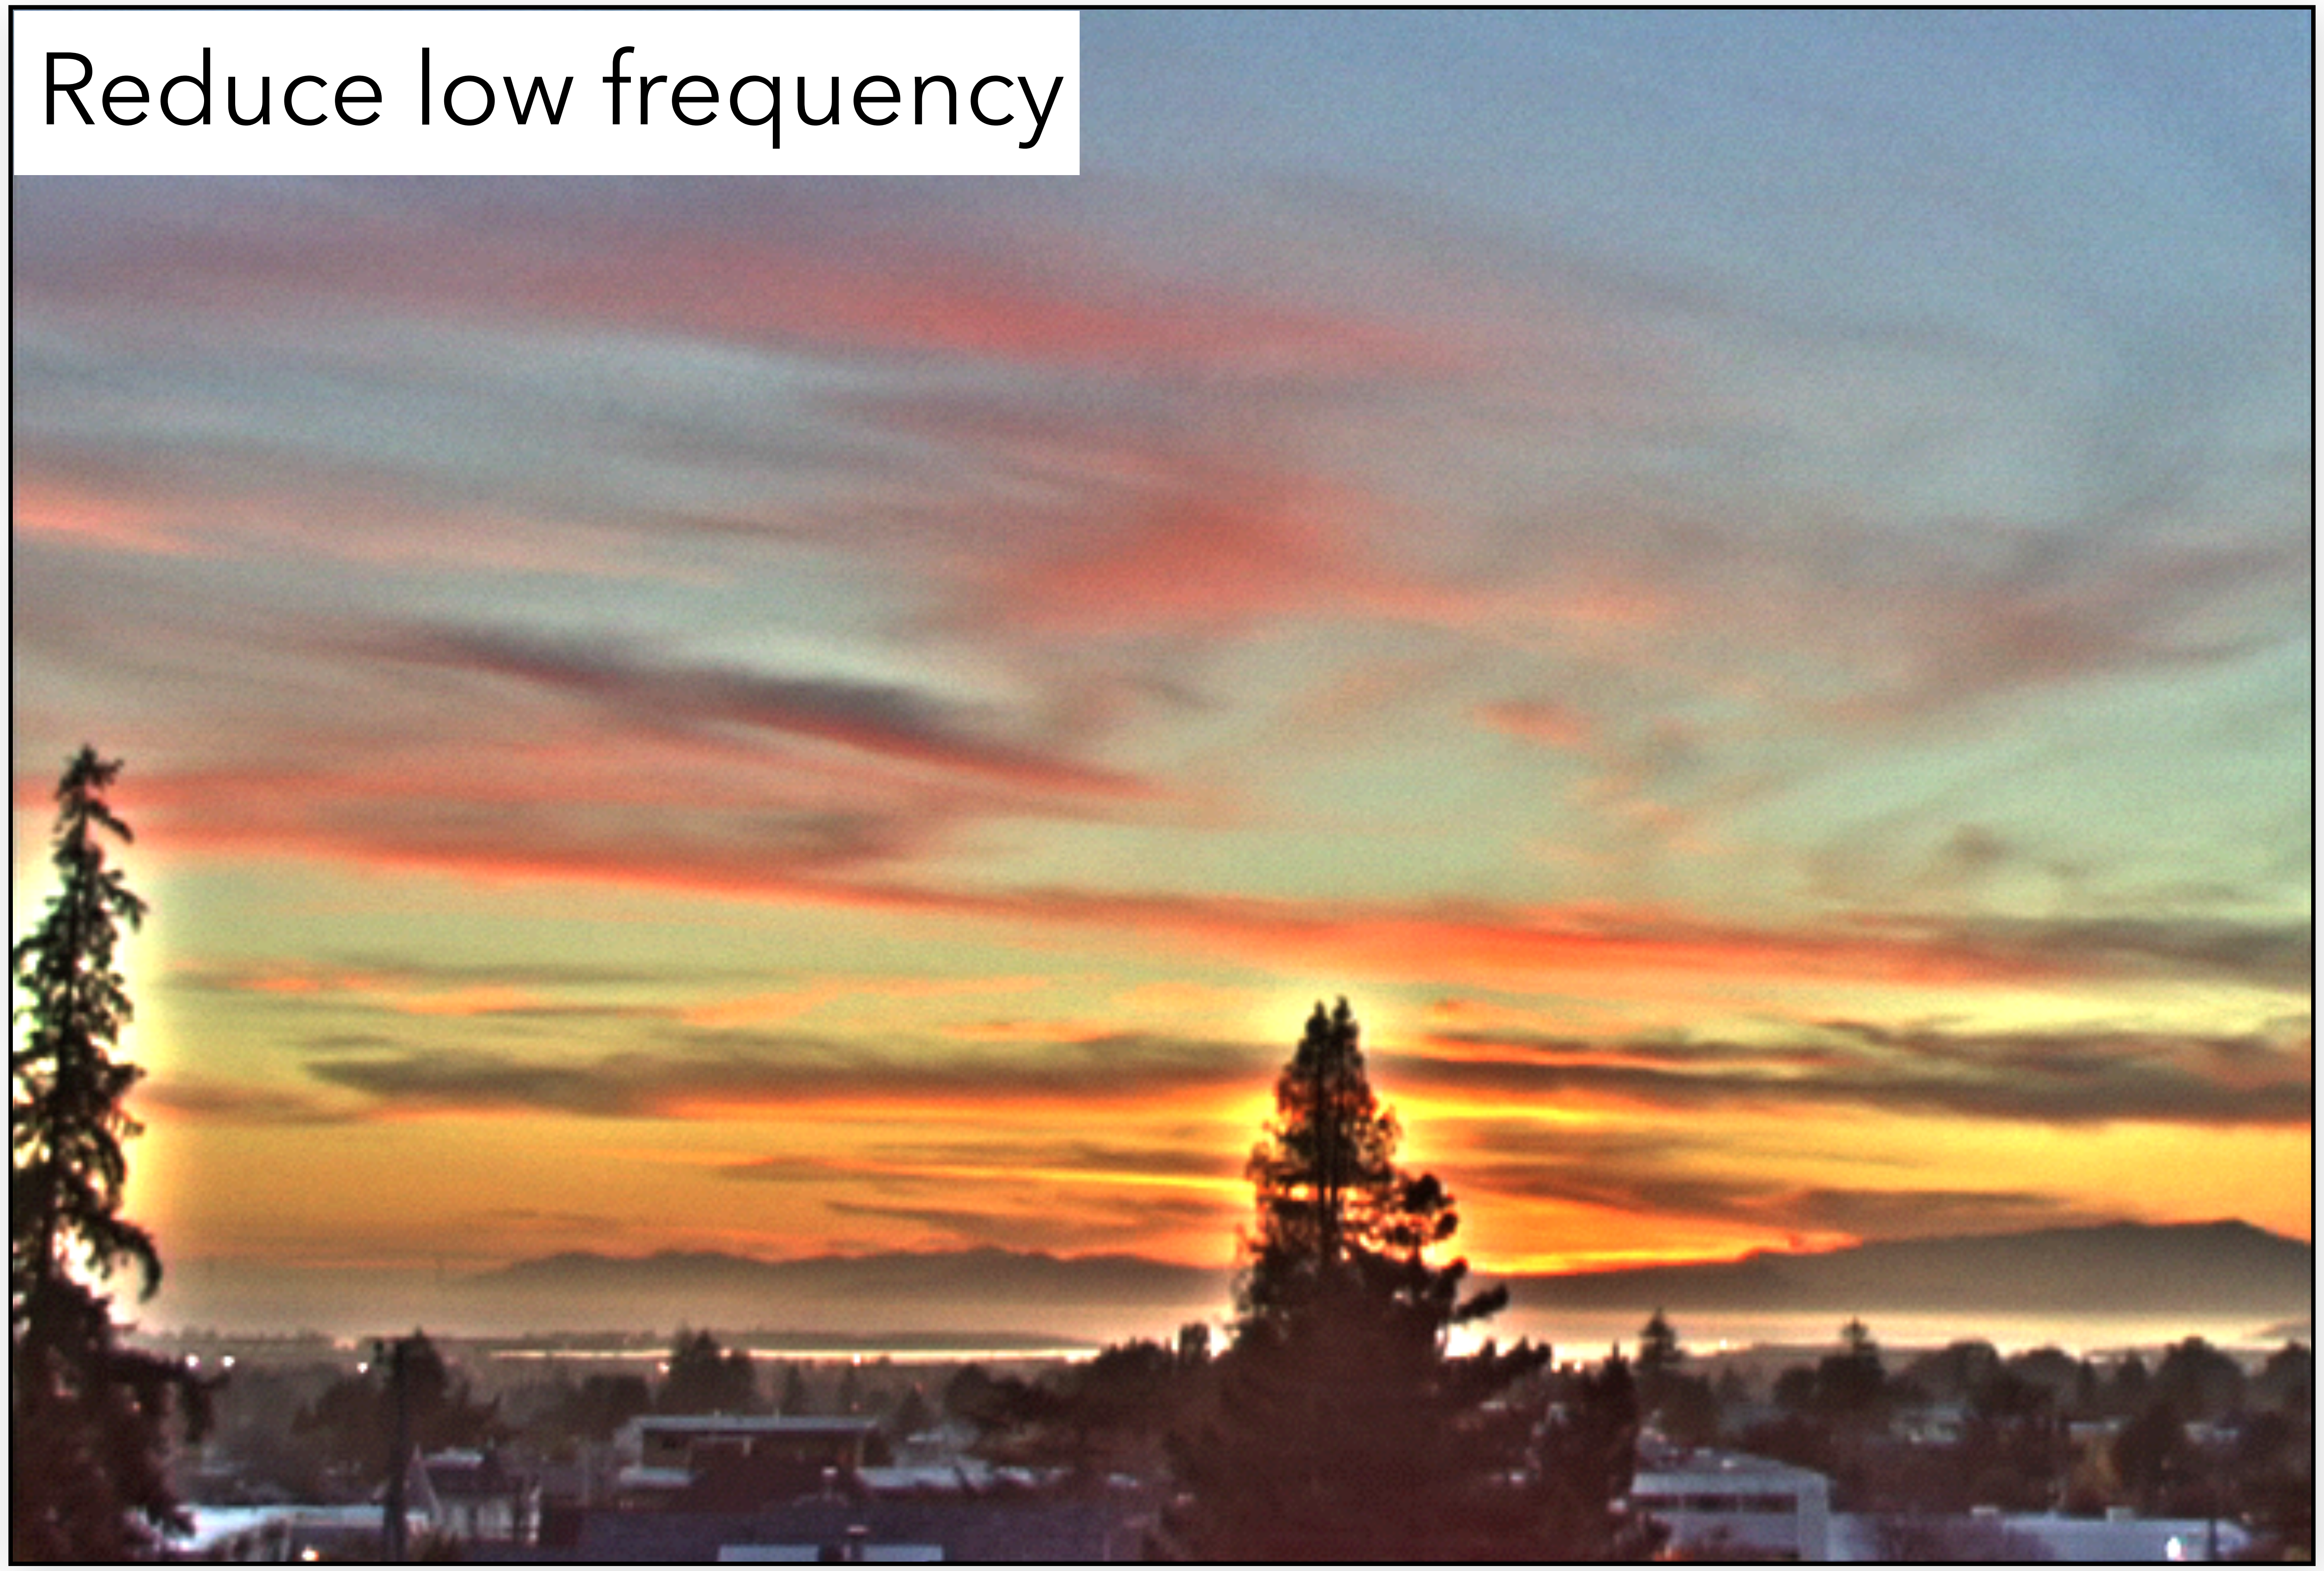

Durand and Dorsey 2002

Don’t blur across edges, decompose using bilateral filter

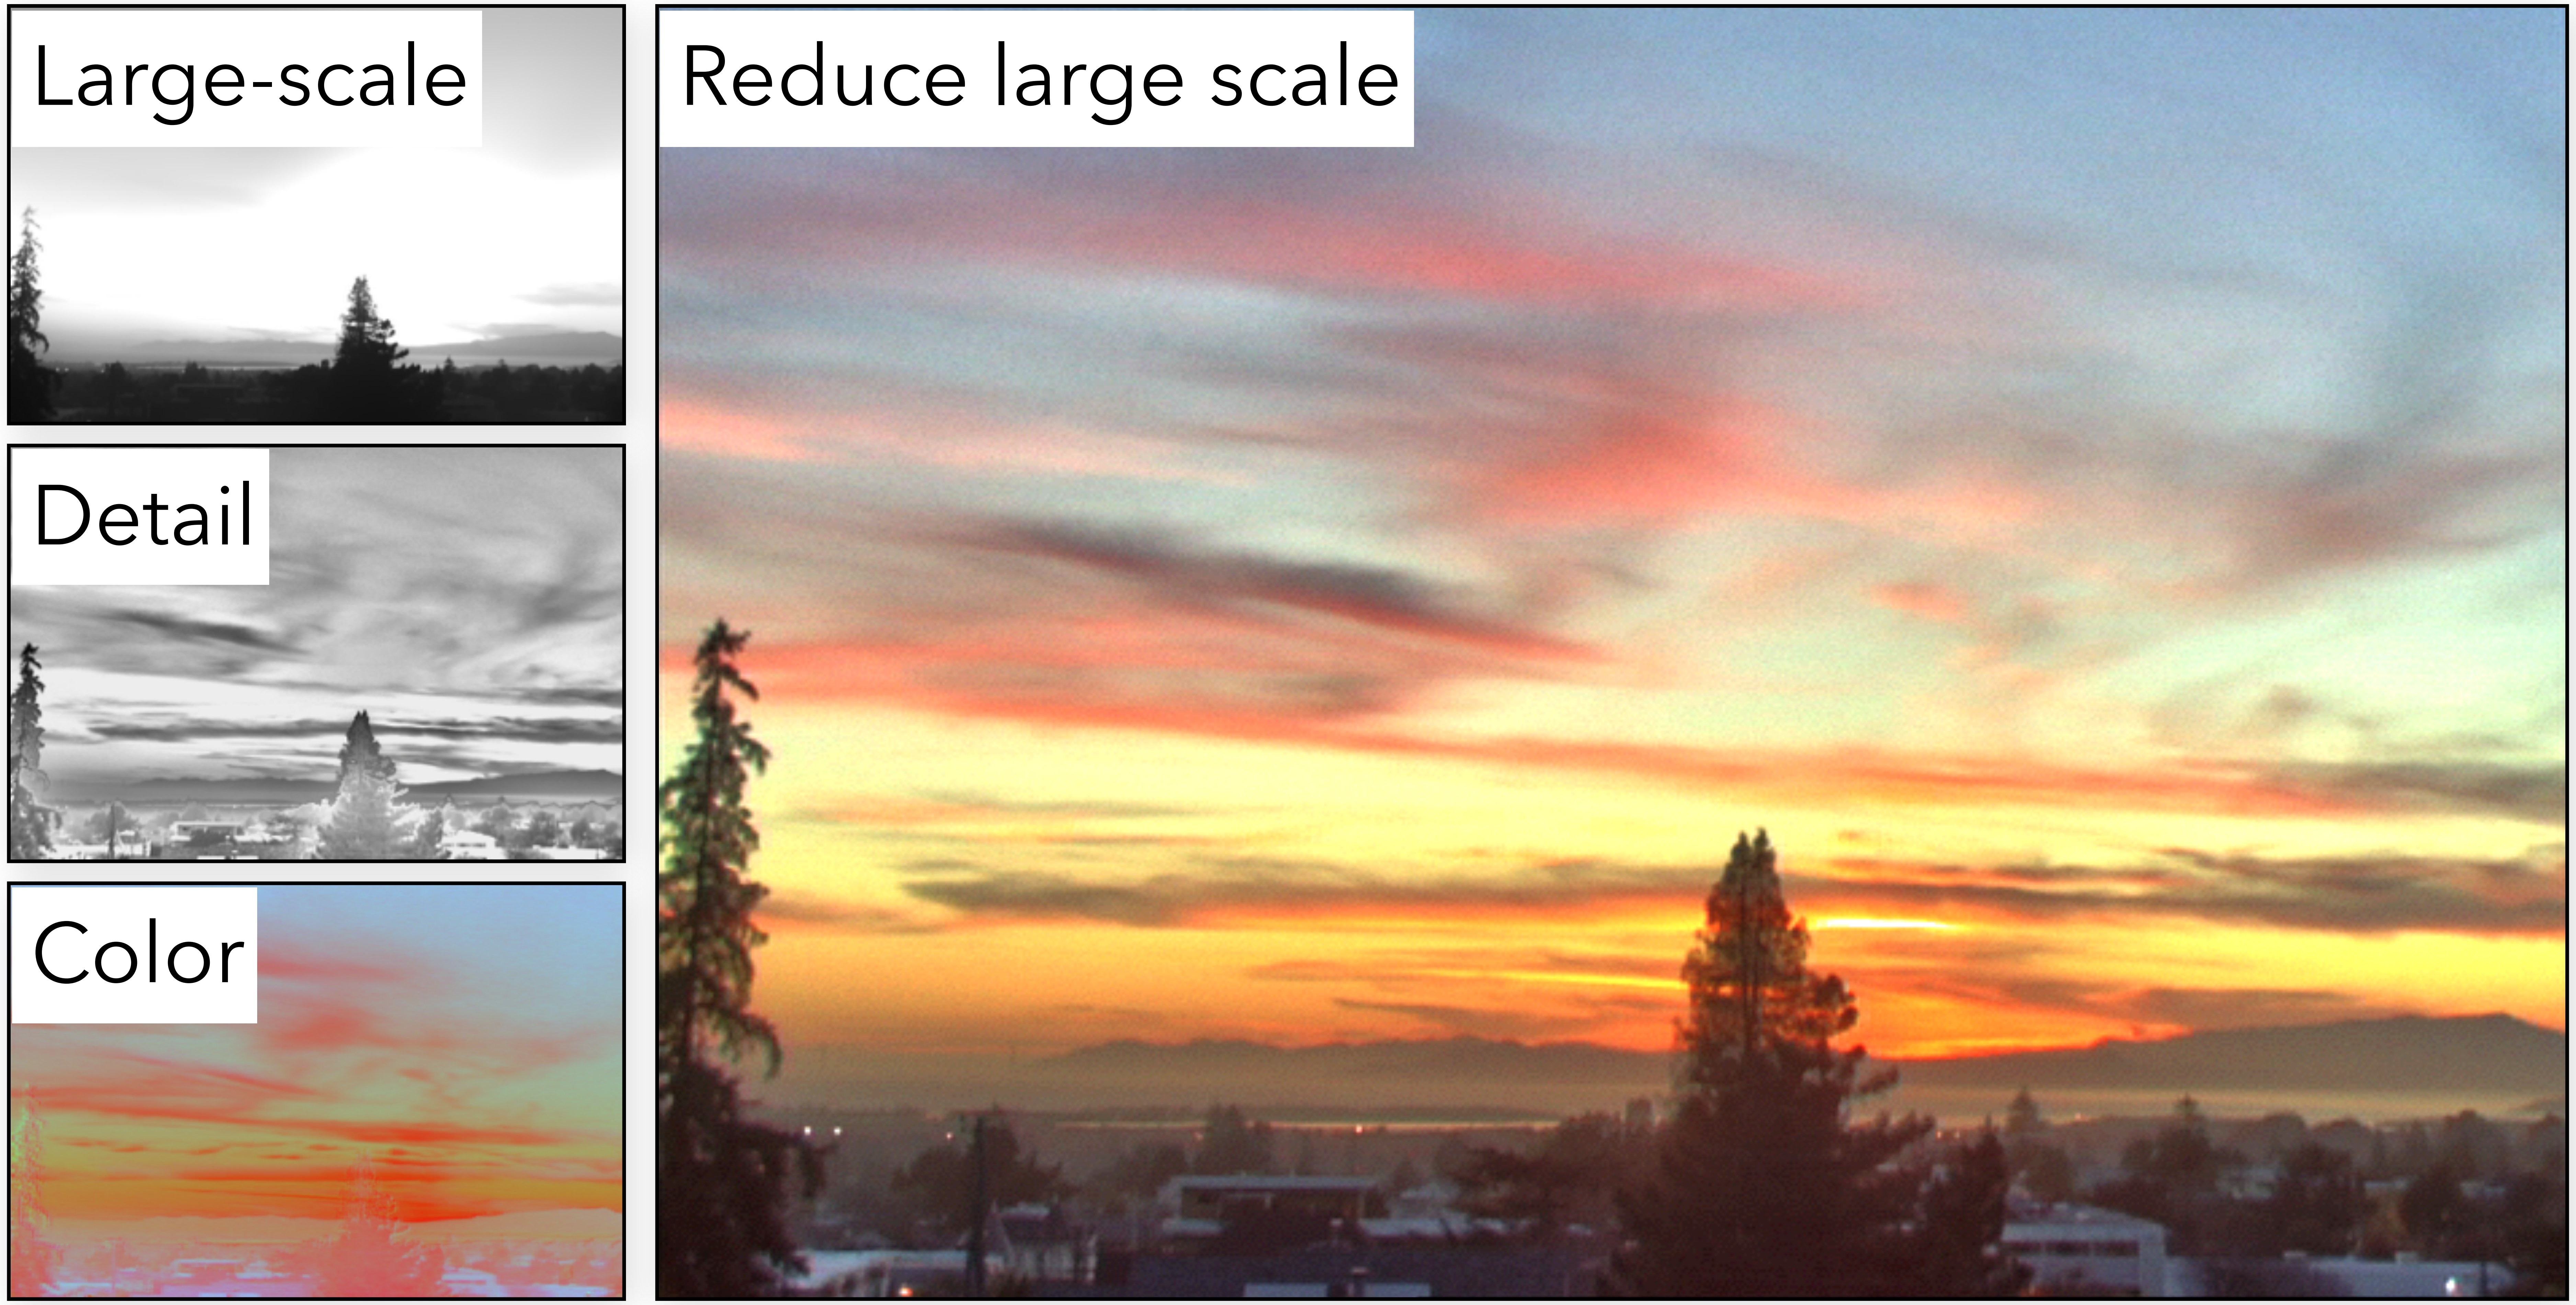

After a slide by Frédo Durand

Contrast reduction

Contrast
too high!

After a slide by Frédo Durand

Contrast reduction

intensity
= 0.4R+0.7G+0.01B

R'=R/intensity
G'=G/intensity
B'=B/intensity

important to use ratios (makes it luminance invariant)

After a slide by Frédo Durand

Contrast reduction

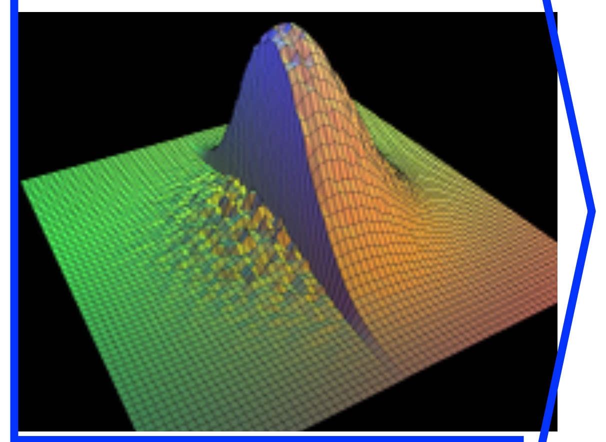

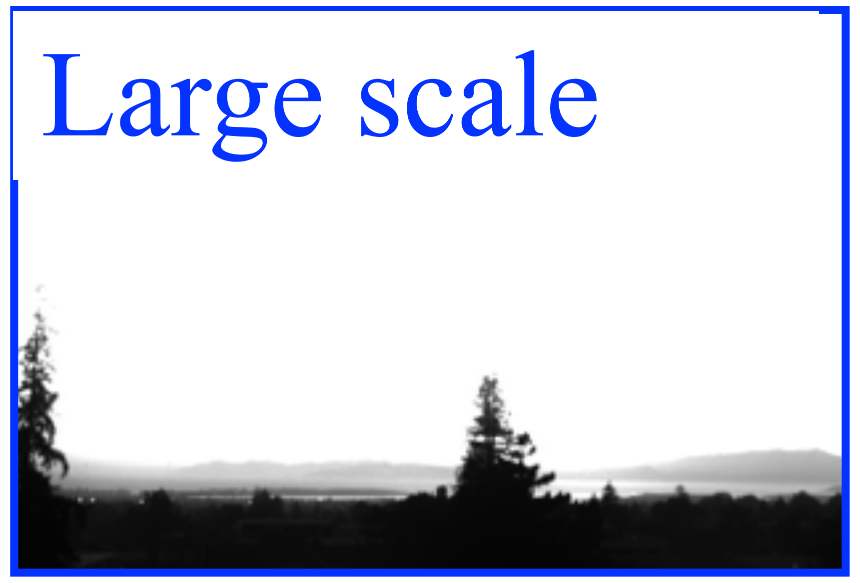

Bilateral Filter

in log

Spatial sigma: 2 to 5% image size
Range sigma: 0.4 (in log 10)

After a slide by Frédo Durand

Contrast reduction

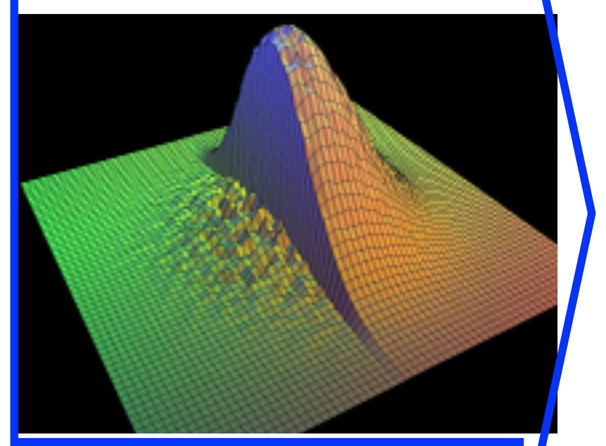

Bilateral Filter

in log

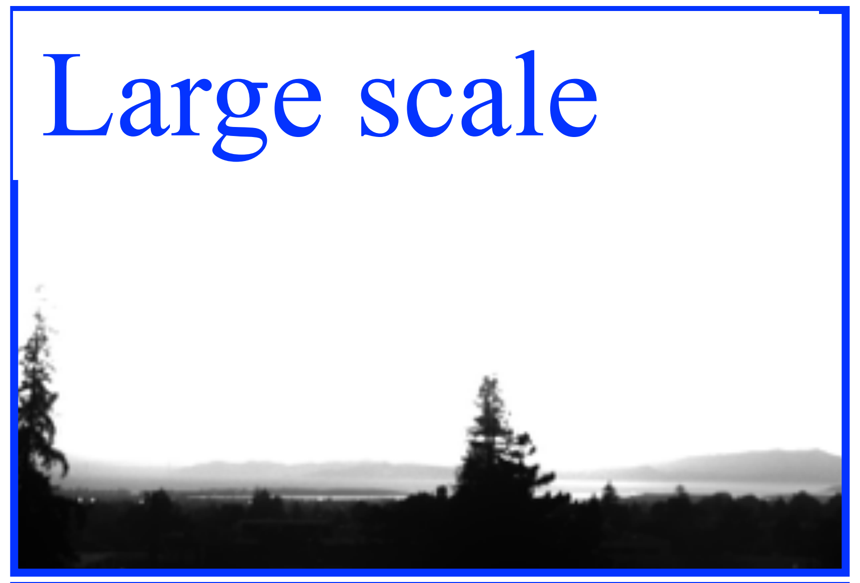

Detail = log intensity - large scale (residual)

After a slide by Frédo Durand

Contrast reduction

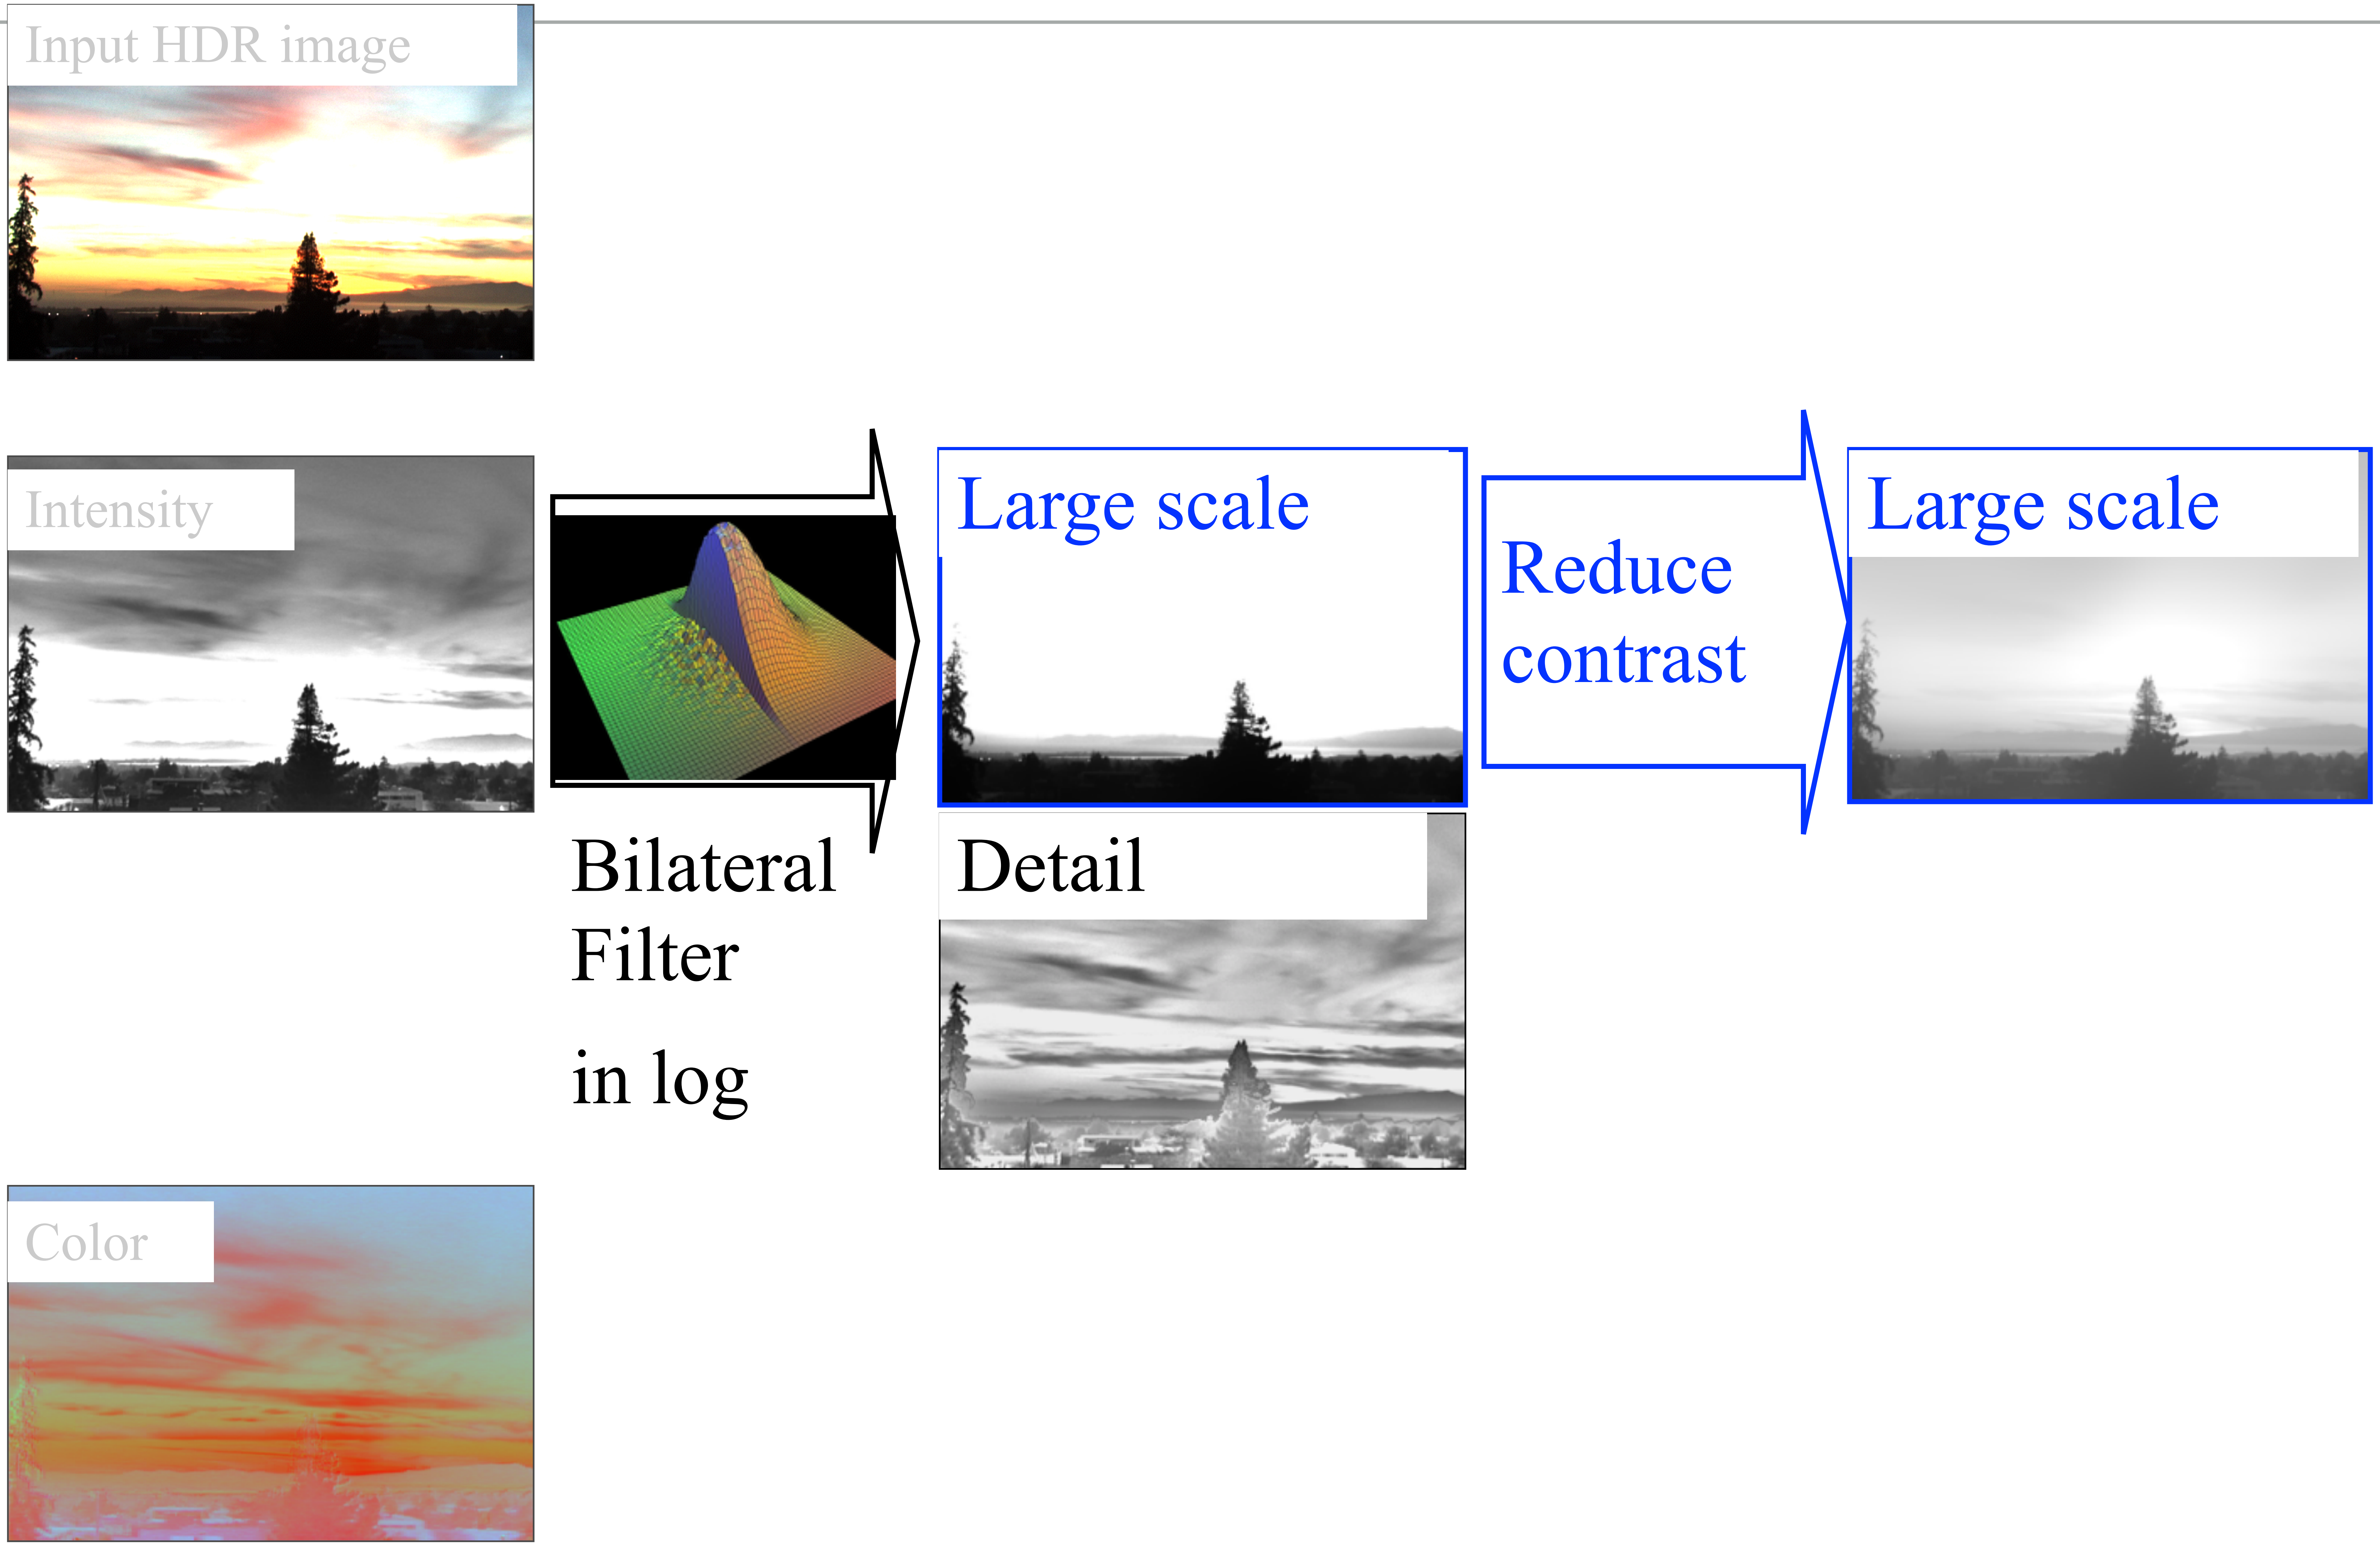

After a slide by Frédo Durand

Contrast reduction

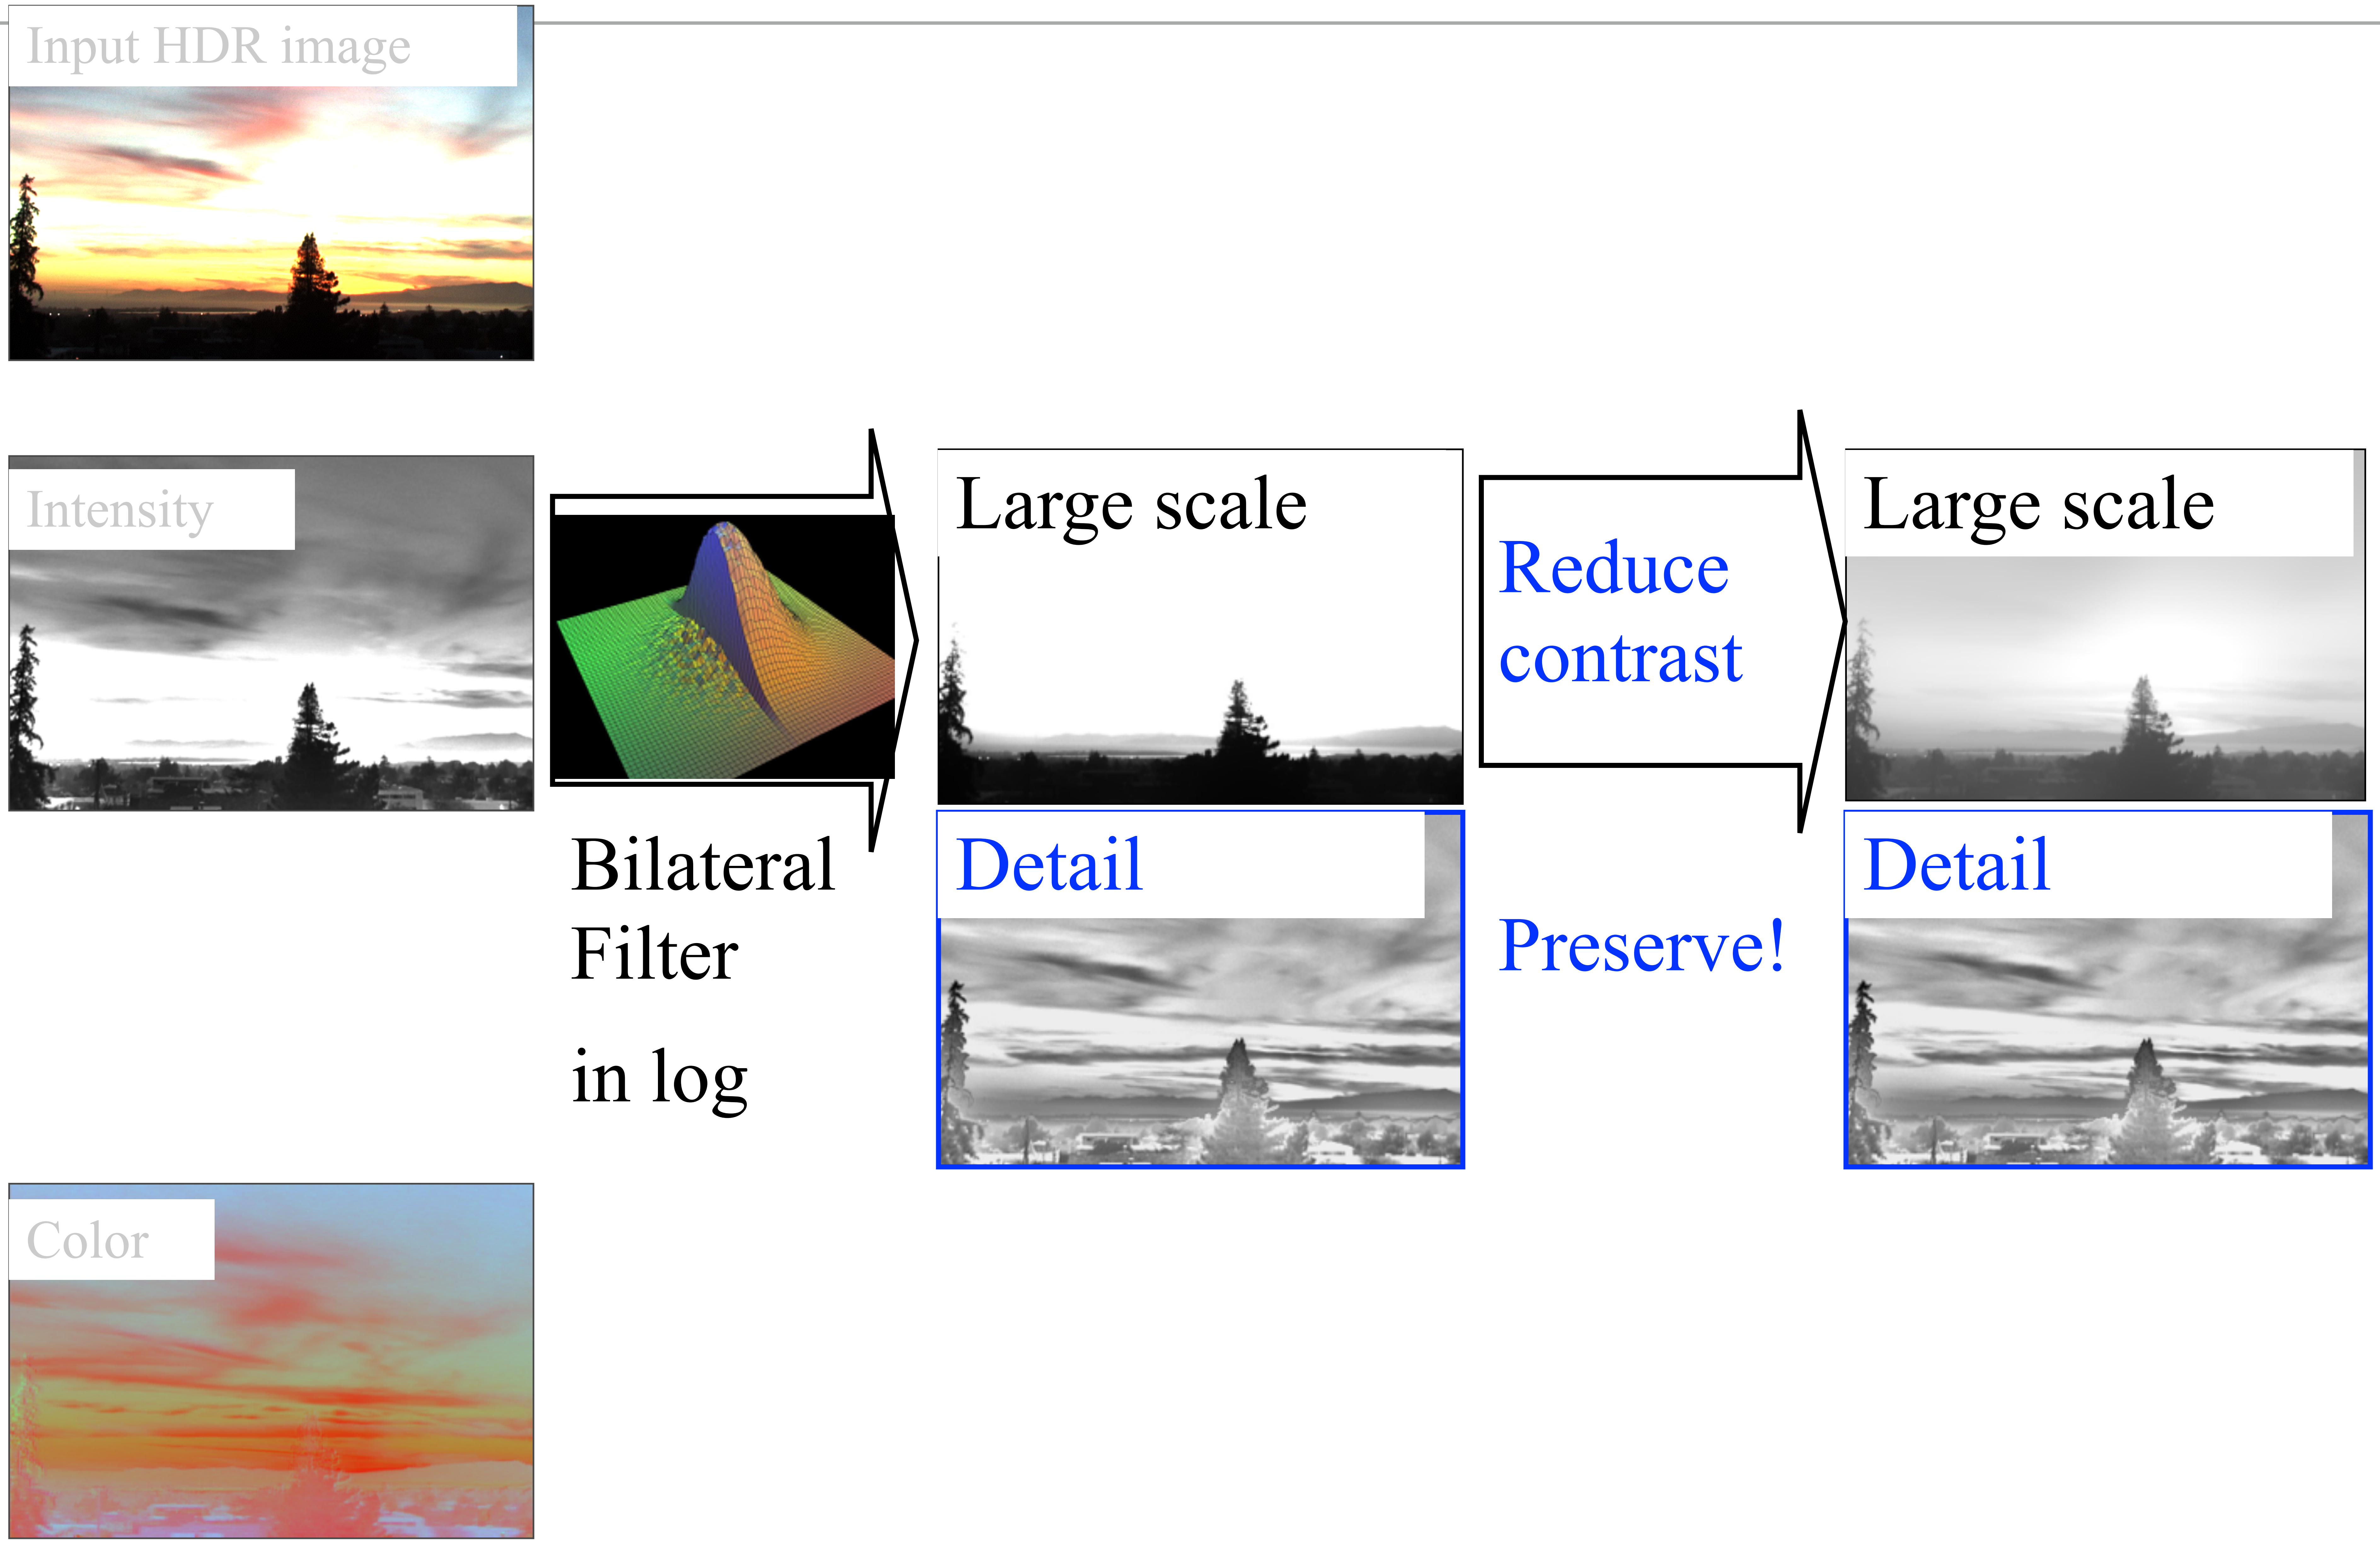

After a slide by Frédo Durand

Contrast reduction

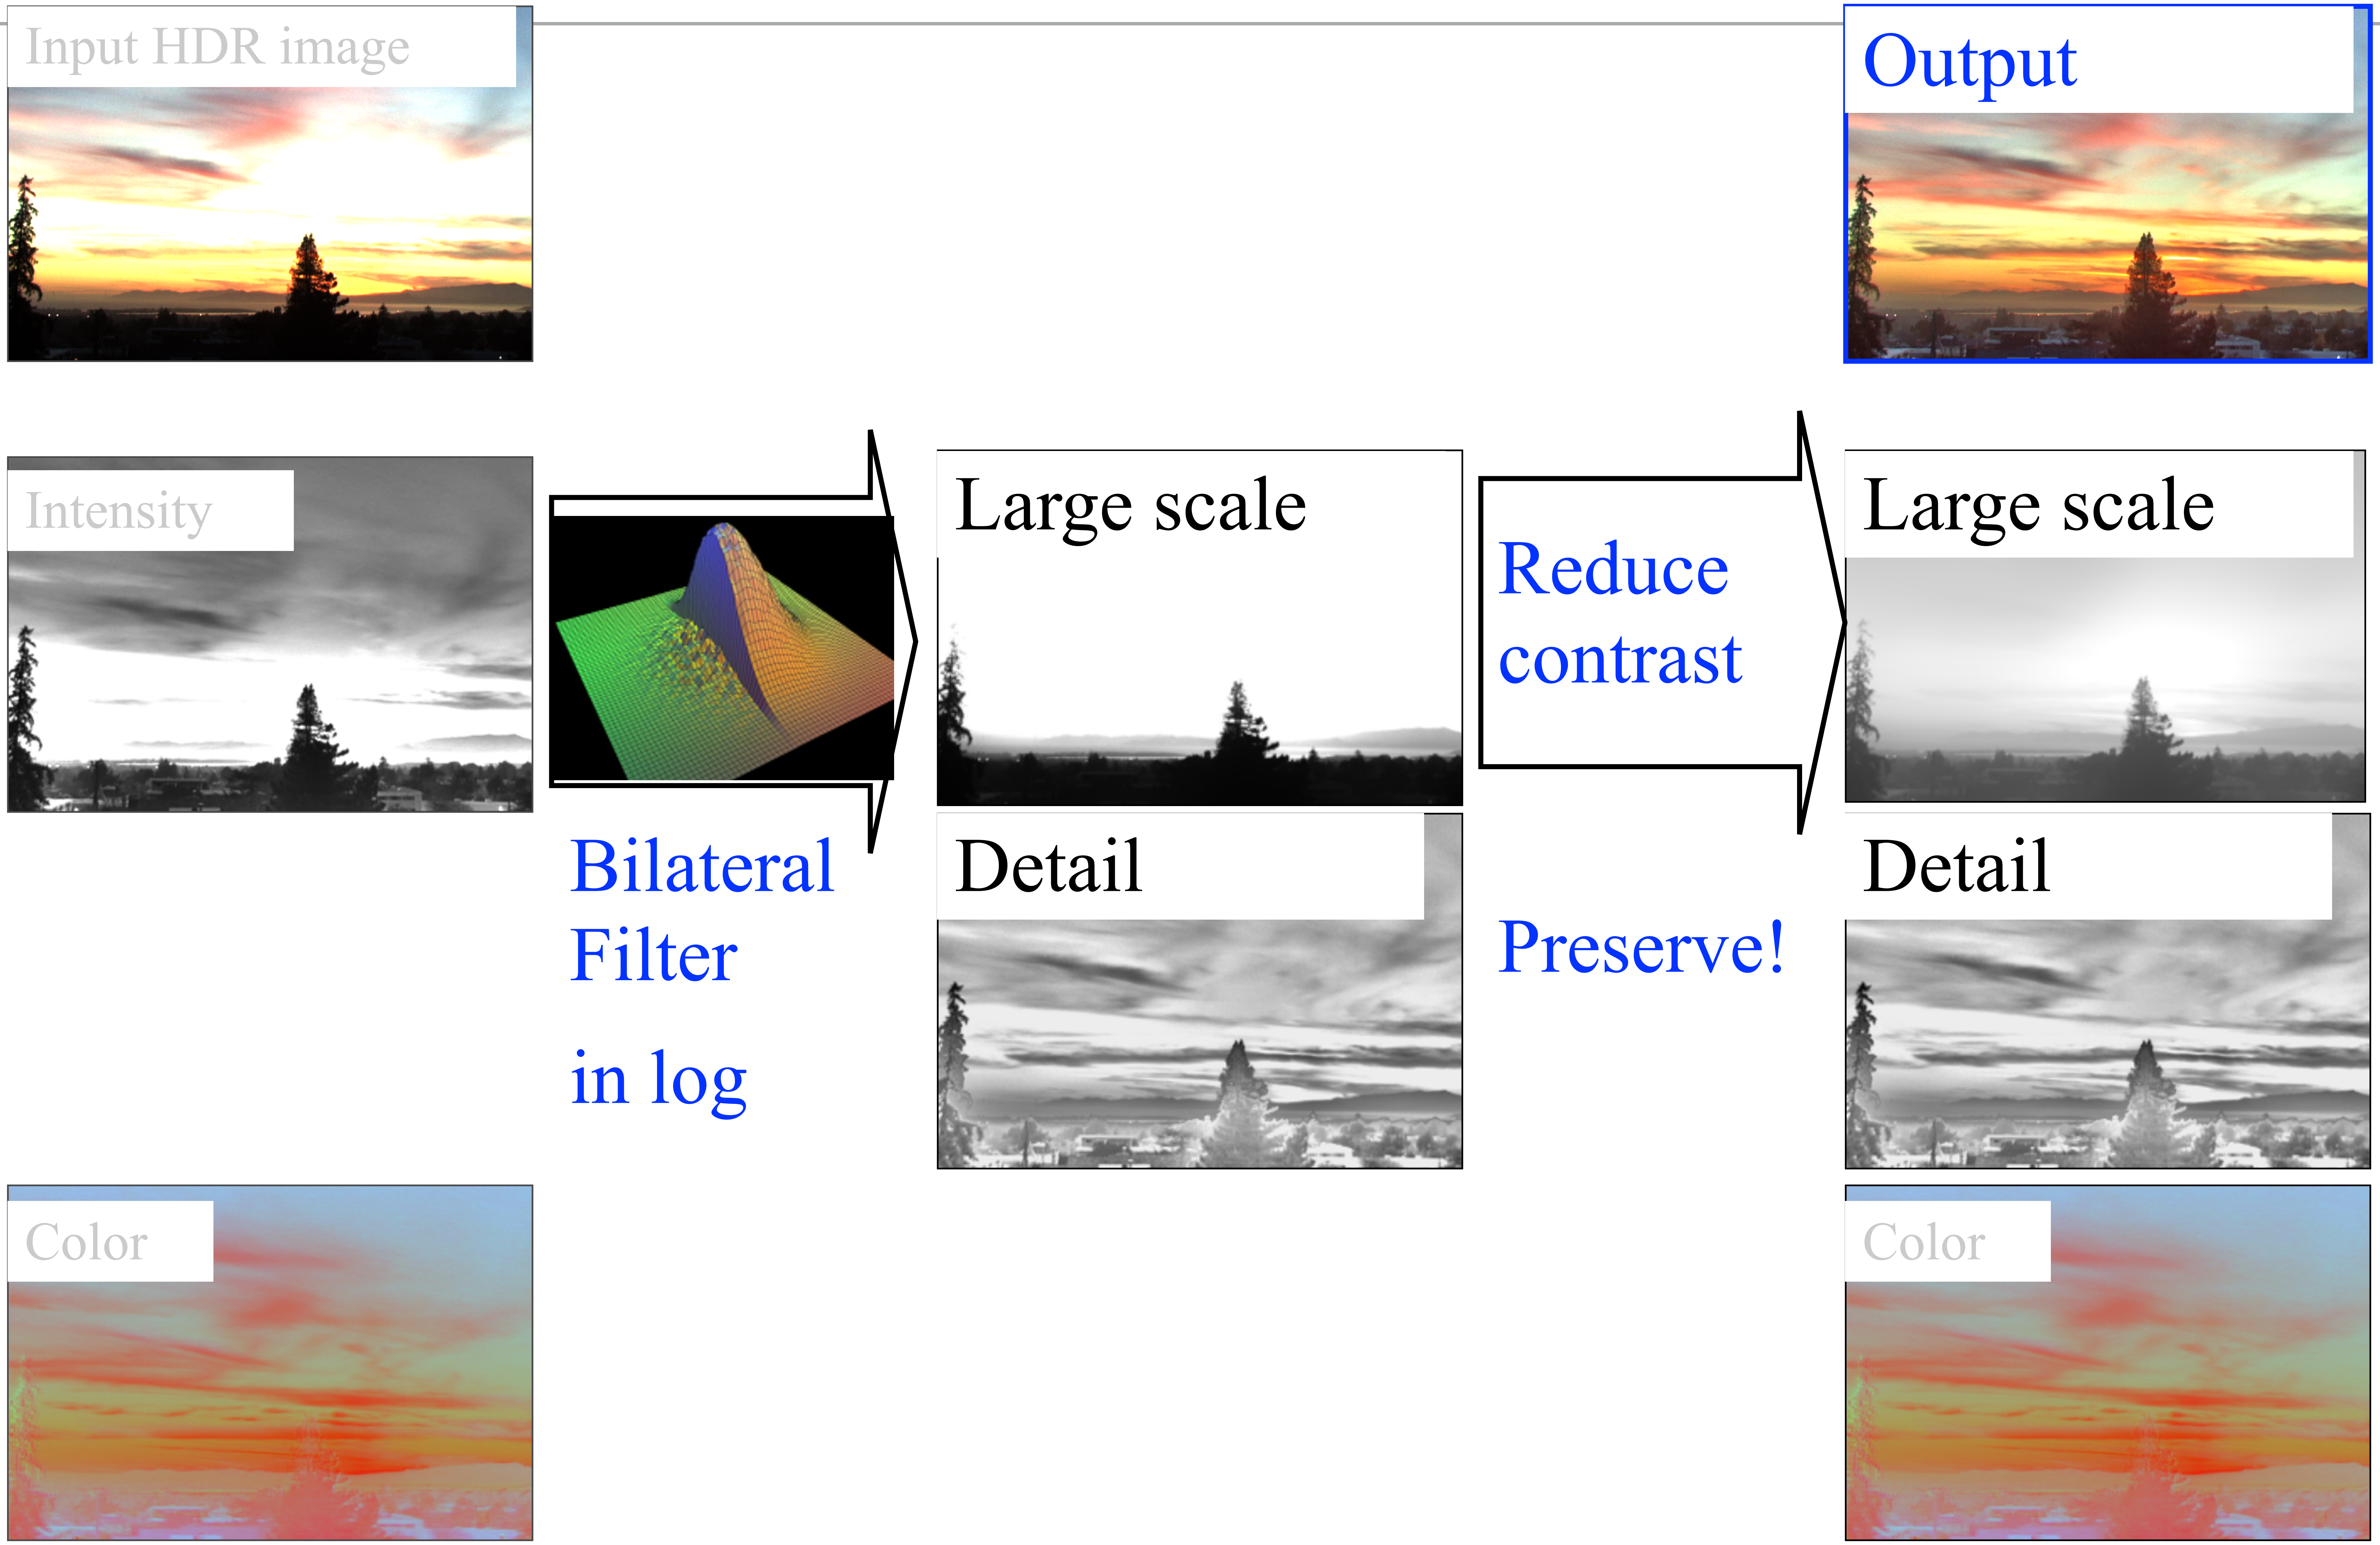

After a slide by Frédo Durand

Log domain

Very important to work in the log domain

Recall: humans are sensitive to multiplicative contrast

With log domain, our notion of "strong edge" always corresponds to the same contrast

After a slide by Frédo Durand

Contrast reduction in log domain

Set target large-scale contrast (e.g. targetRange = $\log_{10}(100)$)

- i.e. in **linear** output, we want 1:100 contrast for large scale

Compute range of input's large-scale layer:

- largeRange = max(inLogLarge) – min(inLogLarge)

Scale factor k = targetRange / largeRange

Normalize so that the biggest value is 0 in log

Optional: amplify detail by detailAmp

After a slide by Frédo Durand

outLog = detailAmp*inLogDetail + k(inLogLarge – max(inLogLarge))

Final output

outLog = detailAmp*inLogDetail + k(inLogLarge – max(inLogLarge))

outIntensity = 10^{outLog}

Recall that R',G',B' is the intensity-normalized RGB color

- outR=outIntensity * R'
- outG=outIntensity * G'
- outB=outIntensity * B'

After a slide by Frédo Durand

Recap

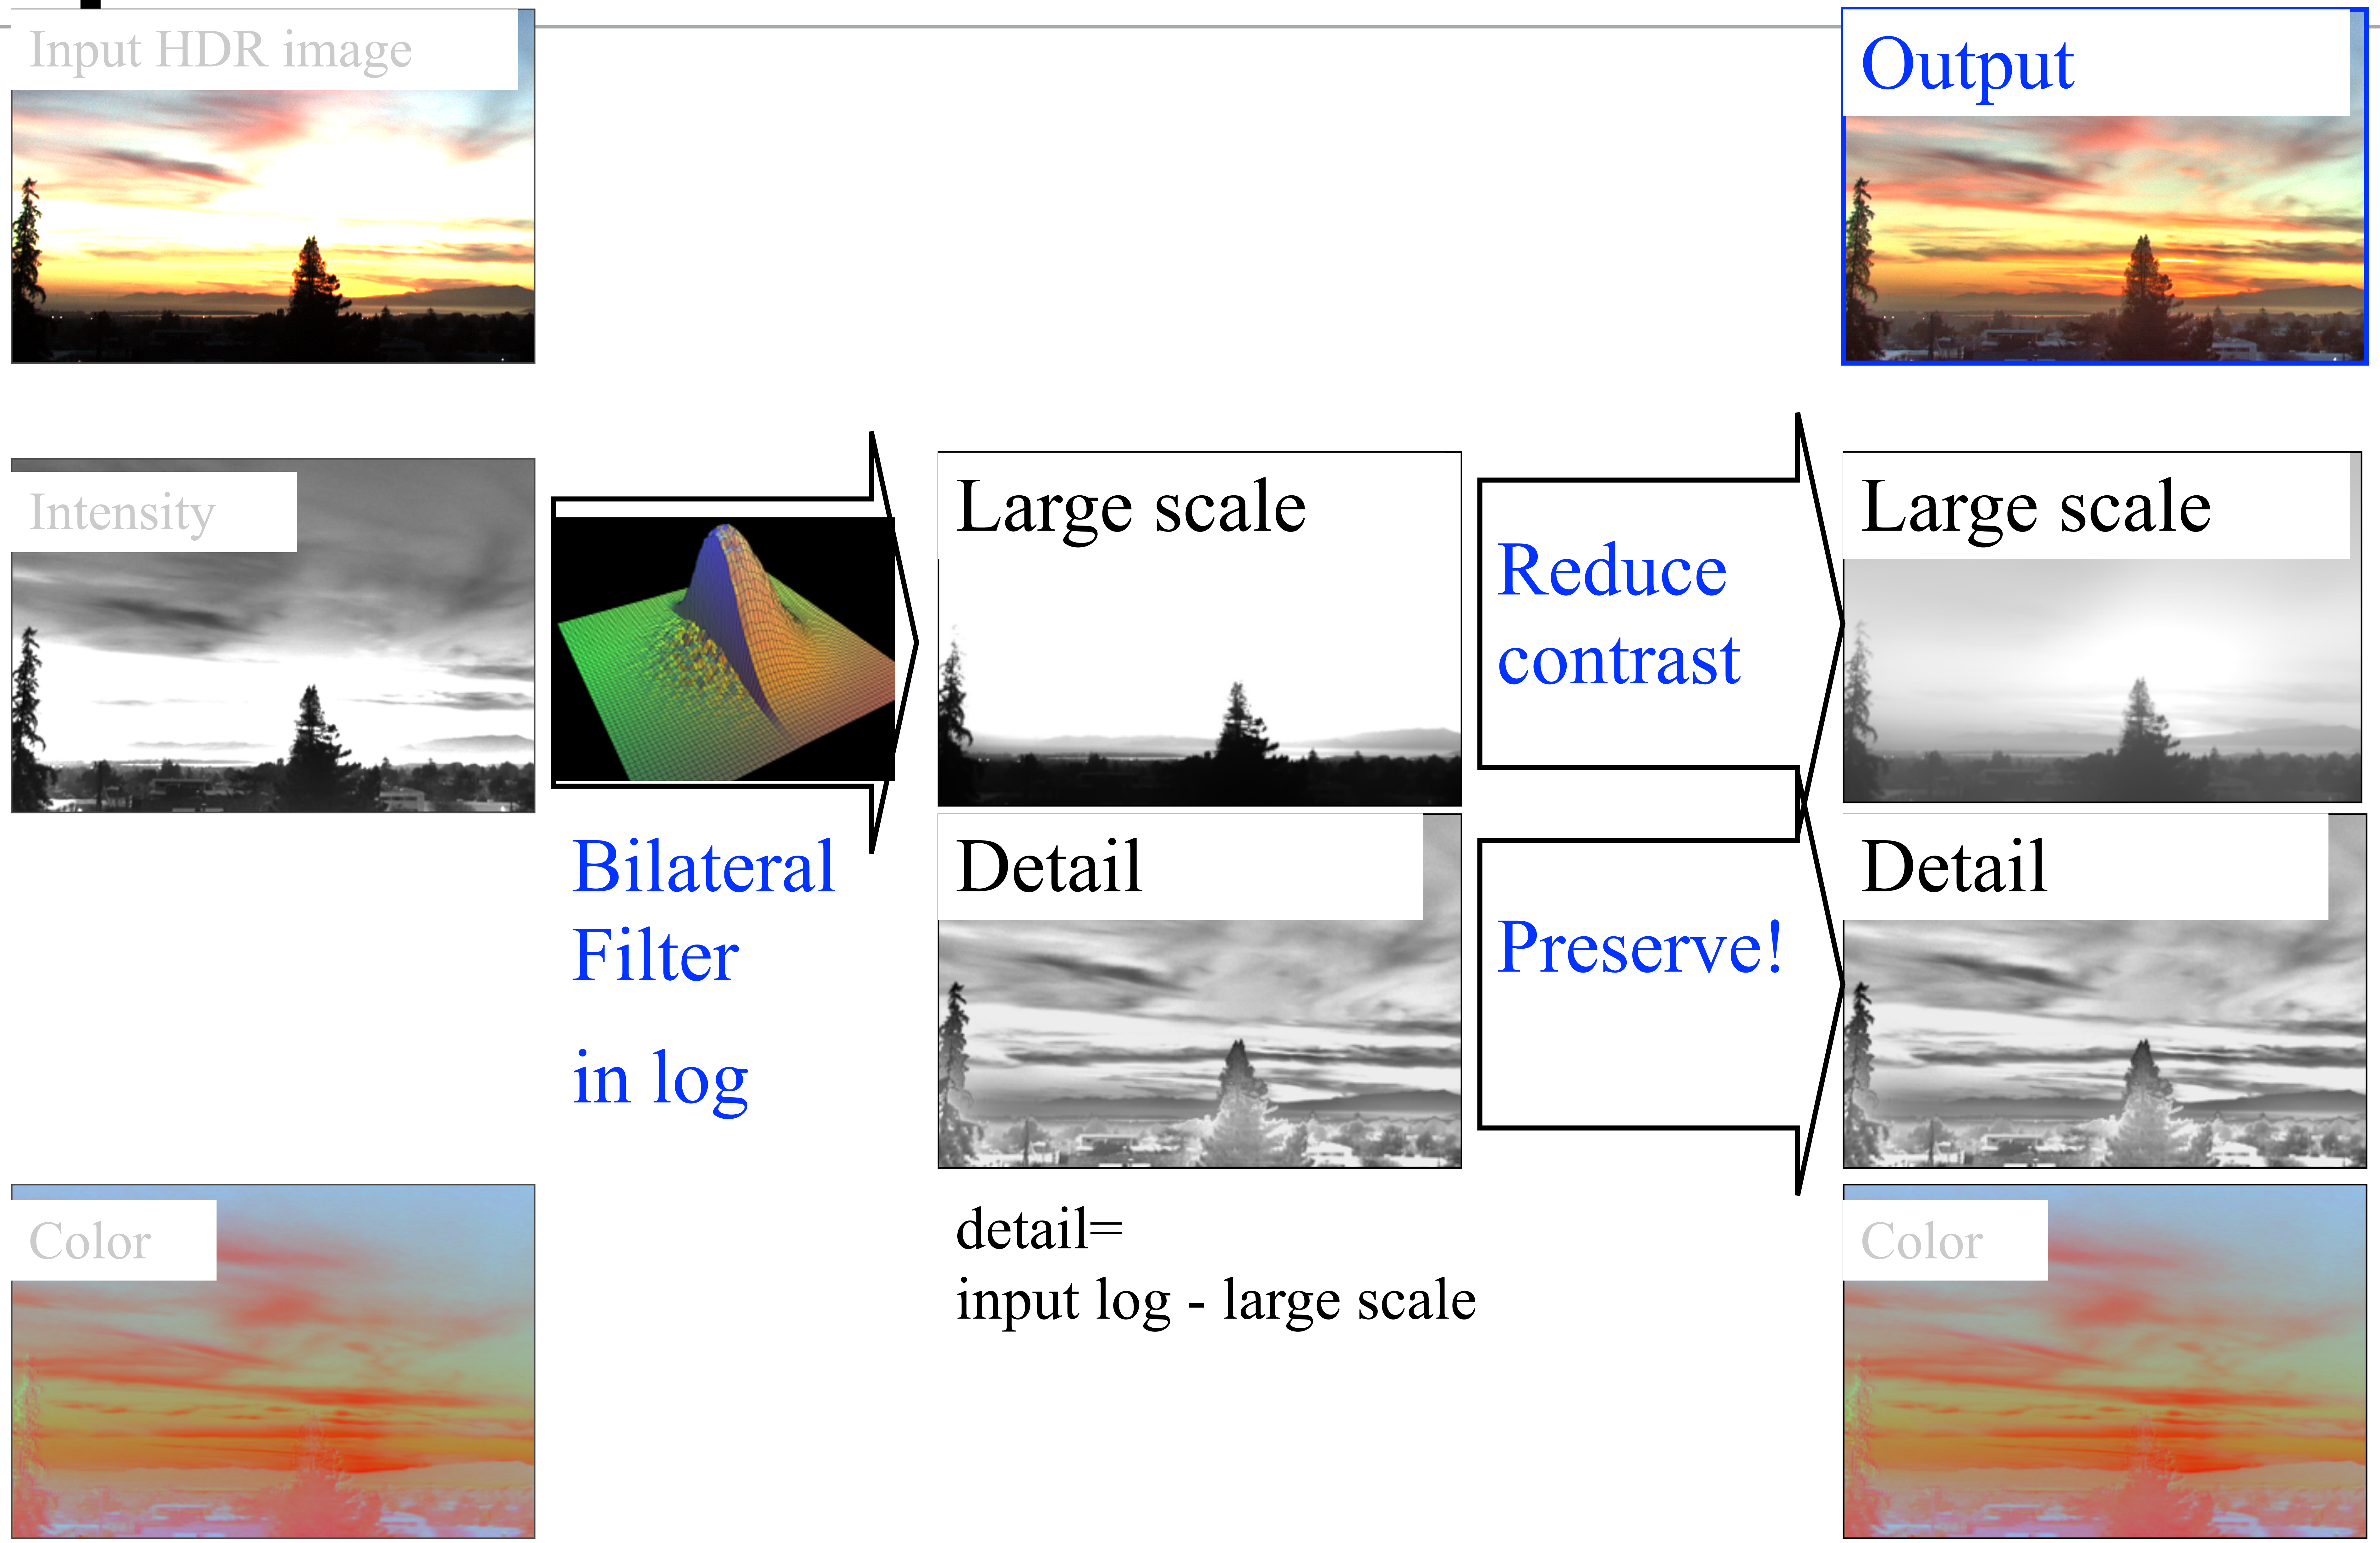

After a slide by Frédo Durand

Bells and whistles: increase detail

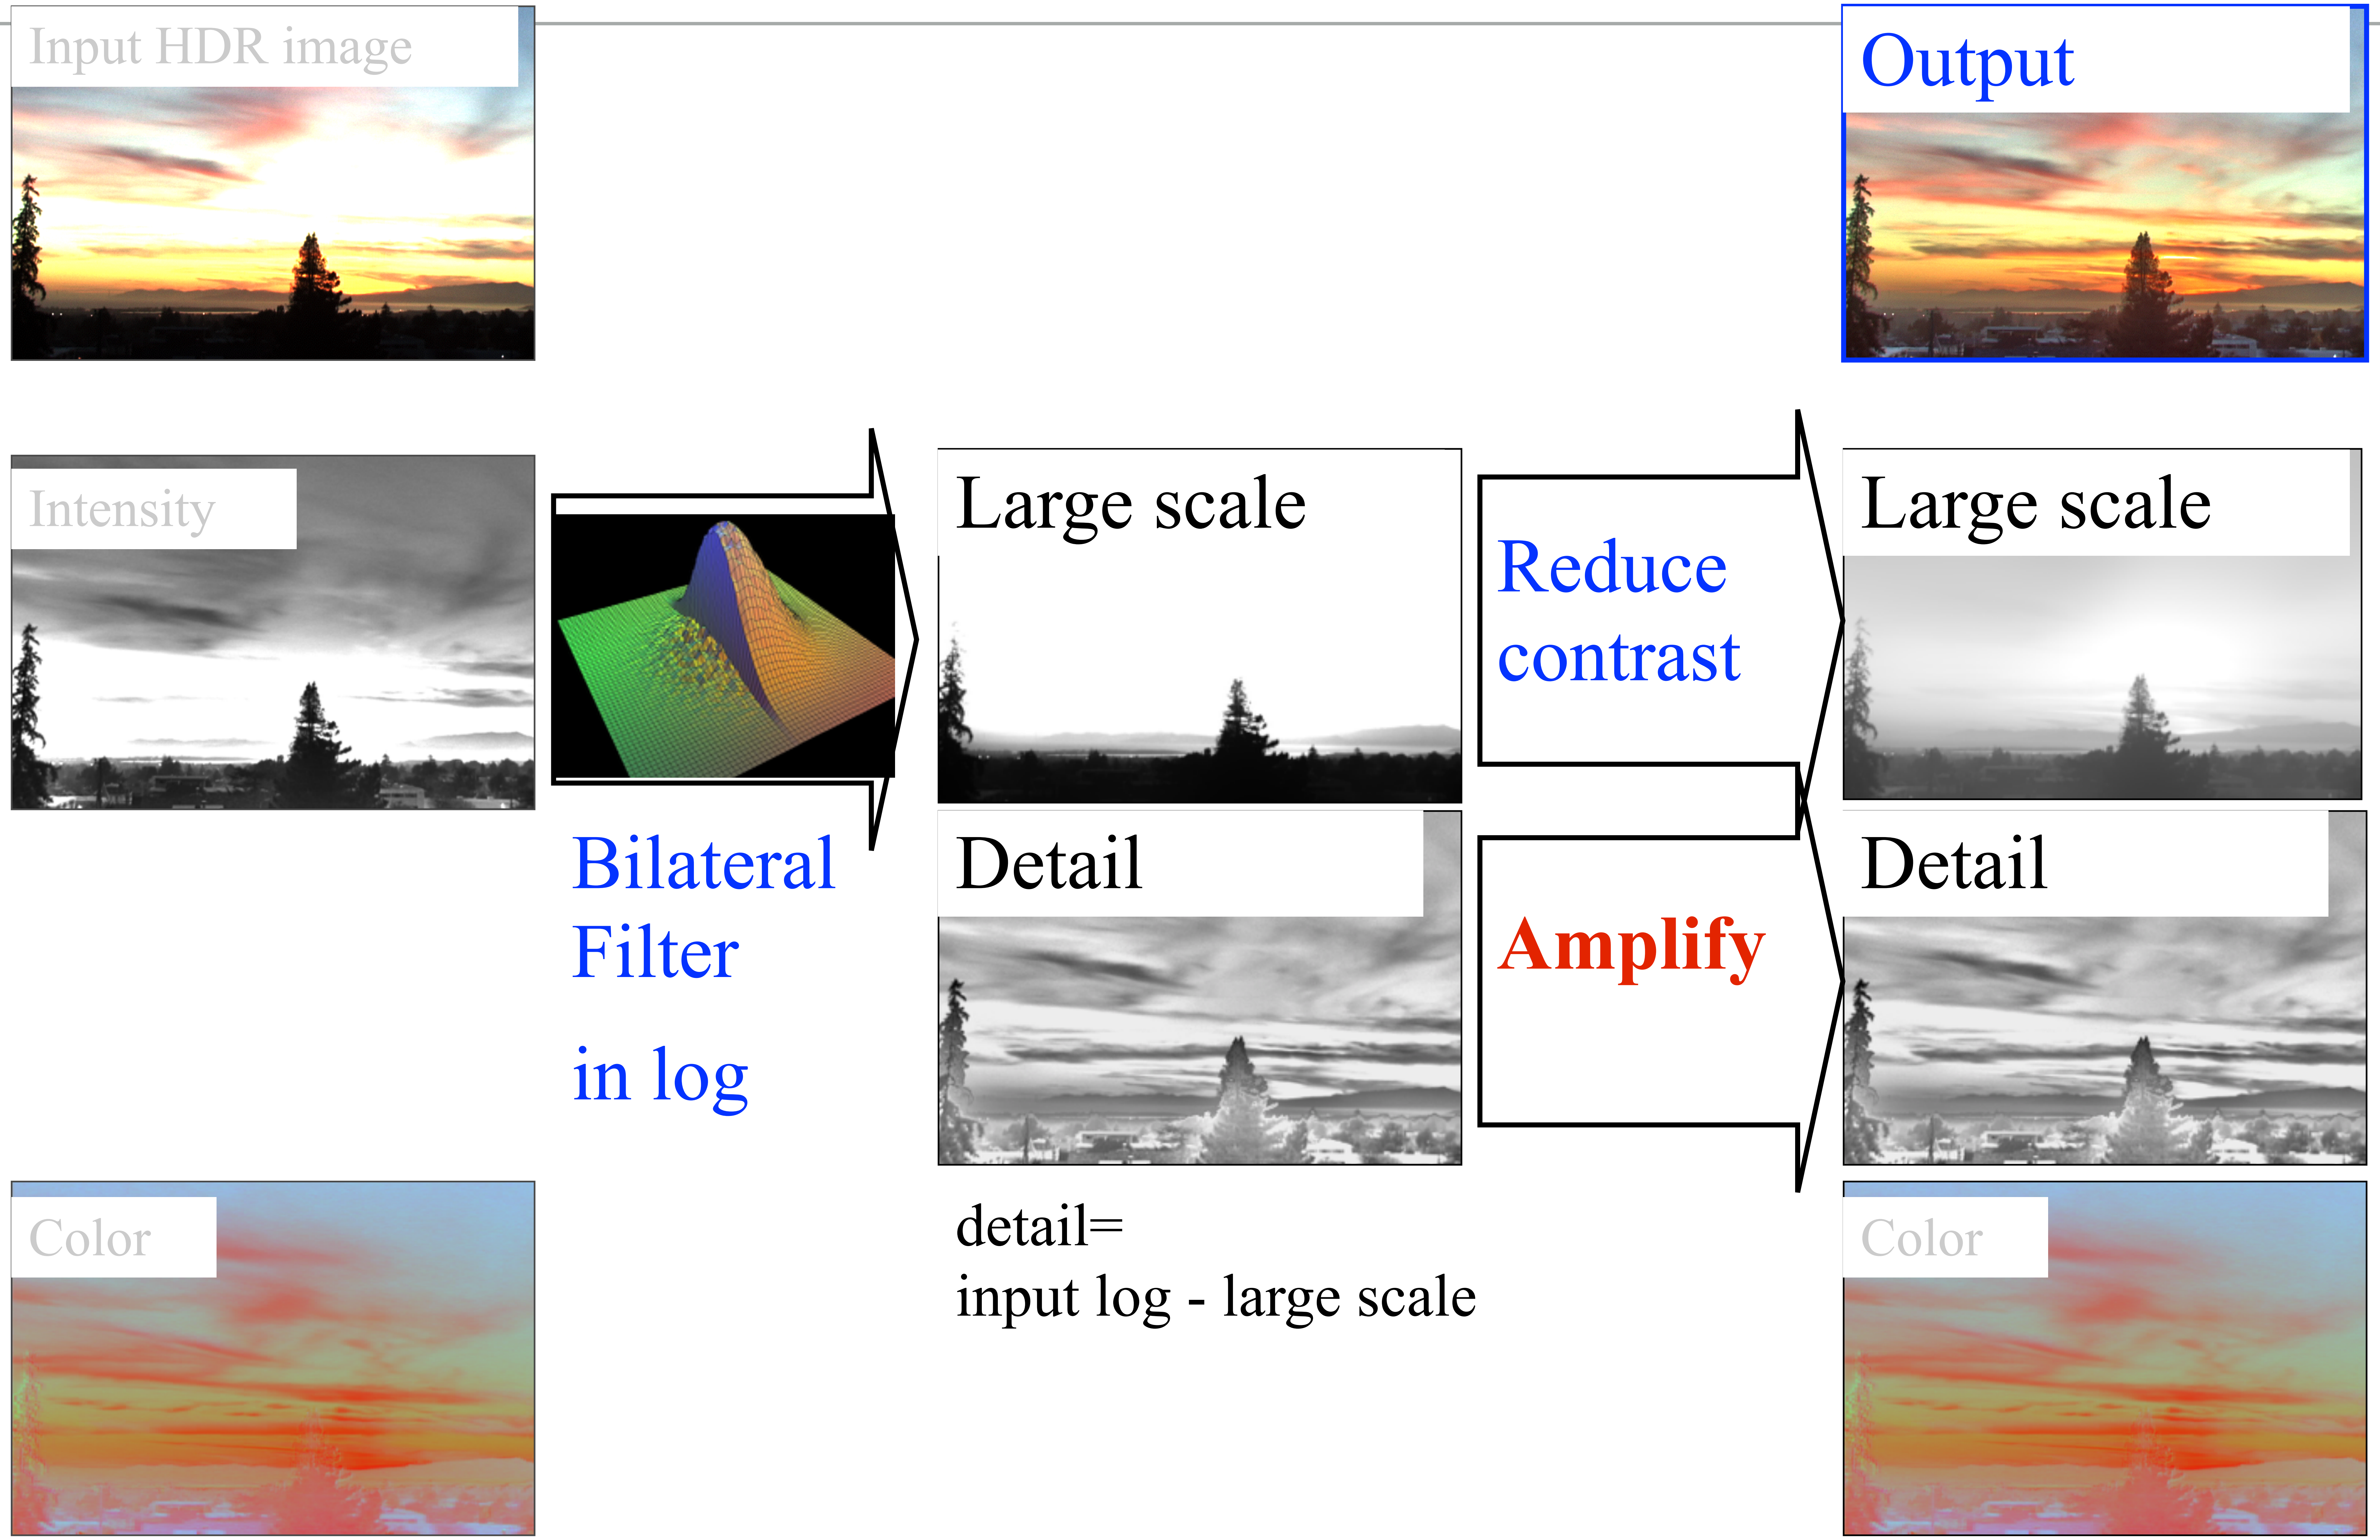

After a slide by Frédo Durand

After a slide by Frédo Durand

What matters

Spatial sigma: not very important

Range sigma: quite important

Use of the log domain for range: critical

- Because HDR and because perception sensitive to multiplicative contrast

After a slide by Frédo Durand

Speed

Direct bilateral filtering is slow (minutes)

Fast algorithm: bilateral grid

- http://groups.csail.mit.edu/graphics/bilagrid/
- http://people.csail.mit.edu/sparis/publi/2009/ijcv/Paris_09_Fast_Approximation.pdf
- http://graphics.stanford.edu/papers/gkdtrees/

After a slide by Frédo Durand

Questions?

Related tools

Photoshop "Local adaptation"

Lightroom "Fill Light"

- or "Shadows"

Photomatix "Details Enhancer"

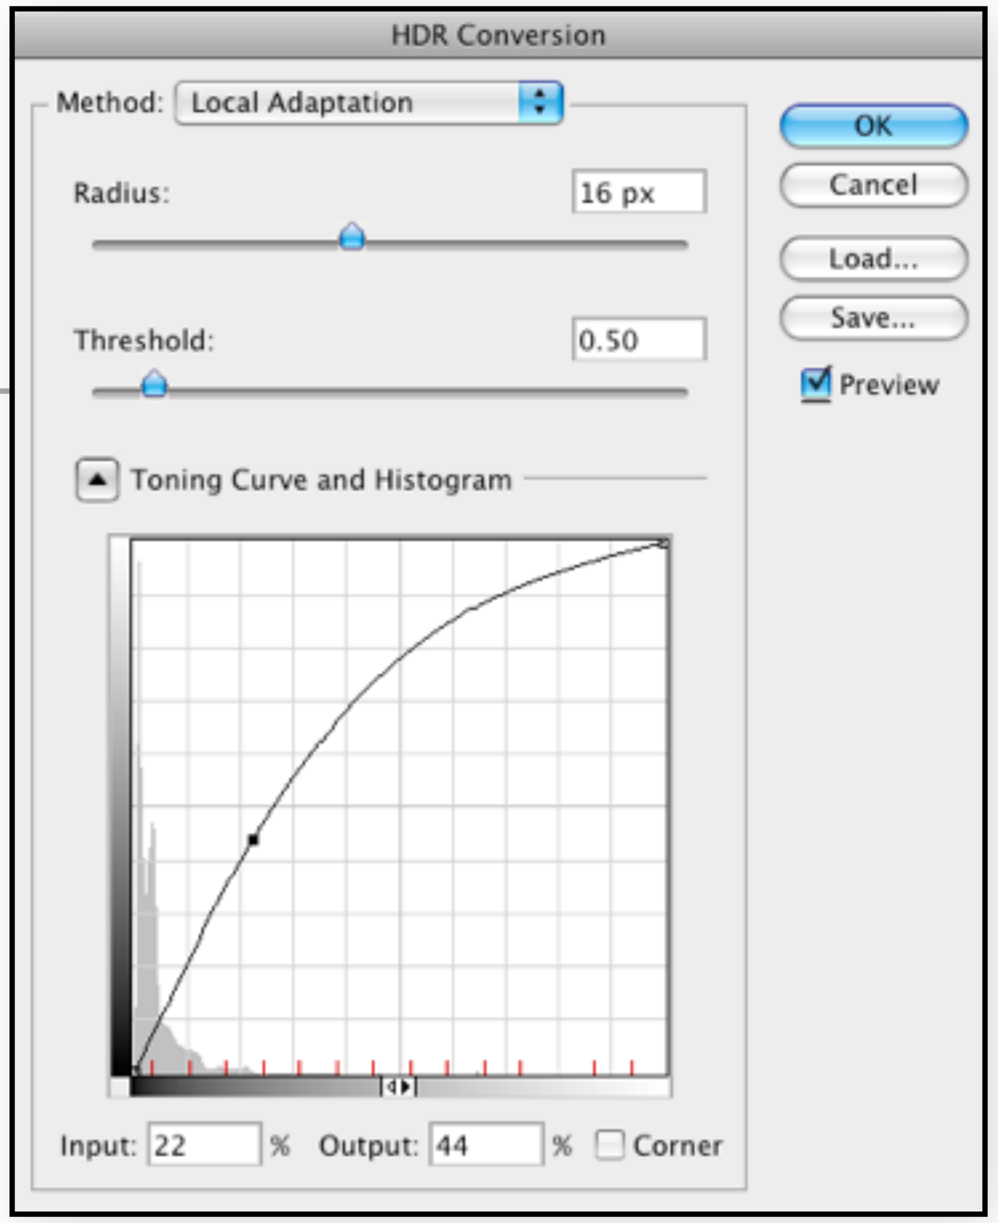

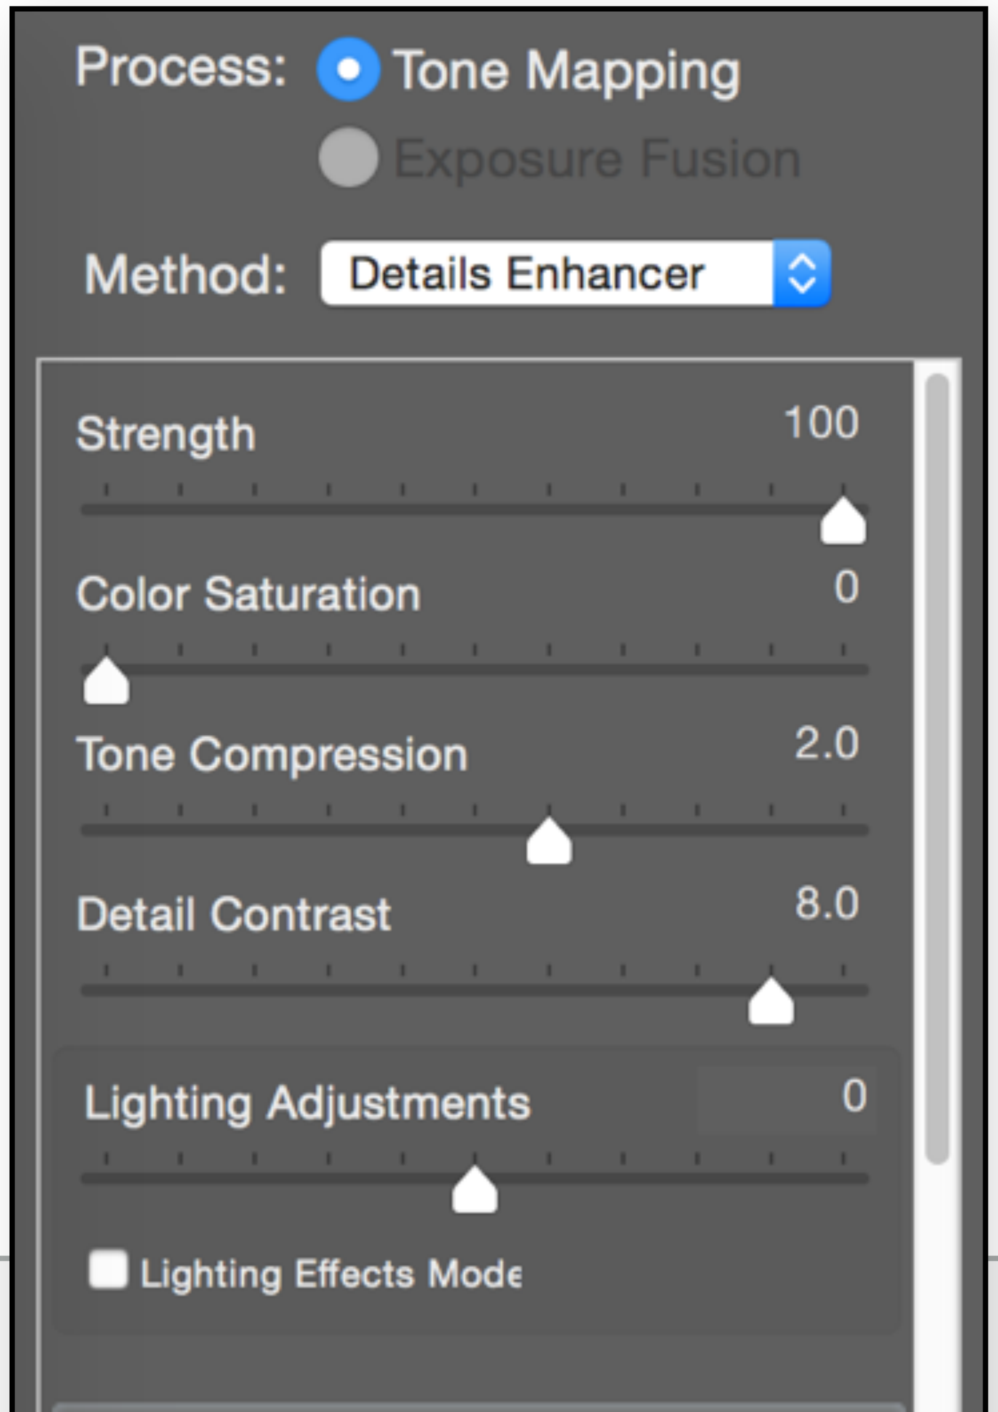

After a slide by Frédo Durand

Slide credits

Frédo Durand

Marc Levoy