HIGH-DYNAMIC-RANGE PHOTOGRAPHY + TONE MAPPING II

Dartmouth

Wojciech Jarosz
wojciech.k.jarosz@dartmouth.edu

CS 89.15/189.5, Fall 2015

Last time

The dynamic range challenge

Applications of HDR photography

Capturing HDR images

Merging bracketed exposures

Tone mapping

- global operators
- local operators

Examples

Sun overexposed

Foreground too dark

Inside is too dark

Outside is too bright

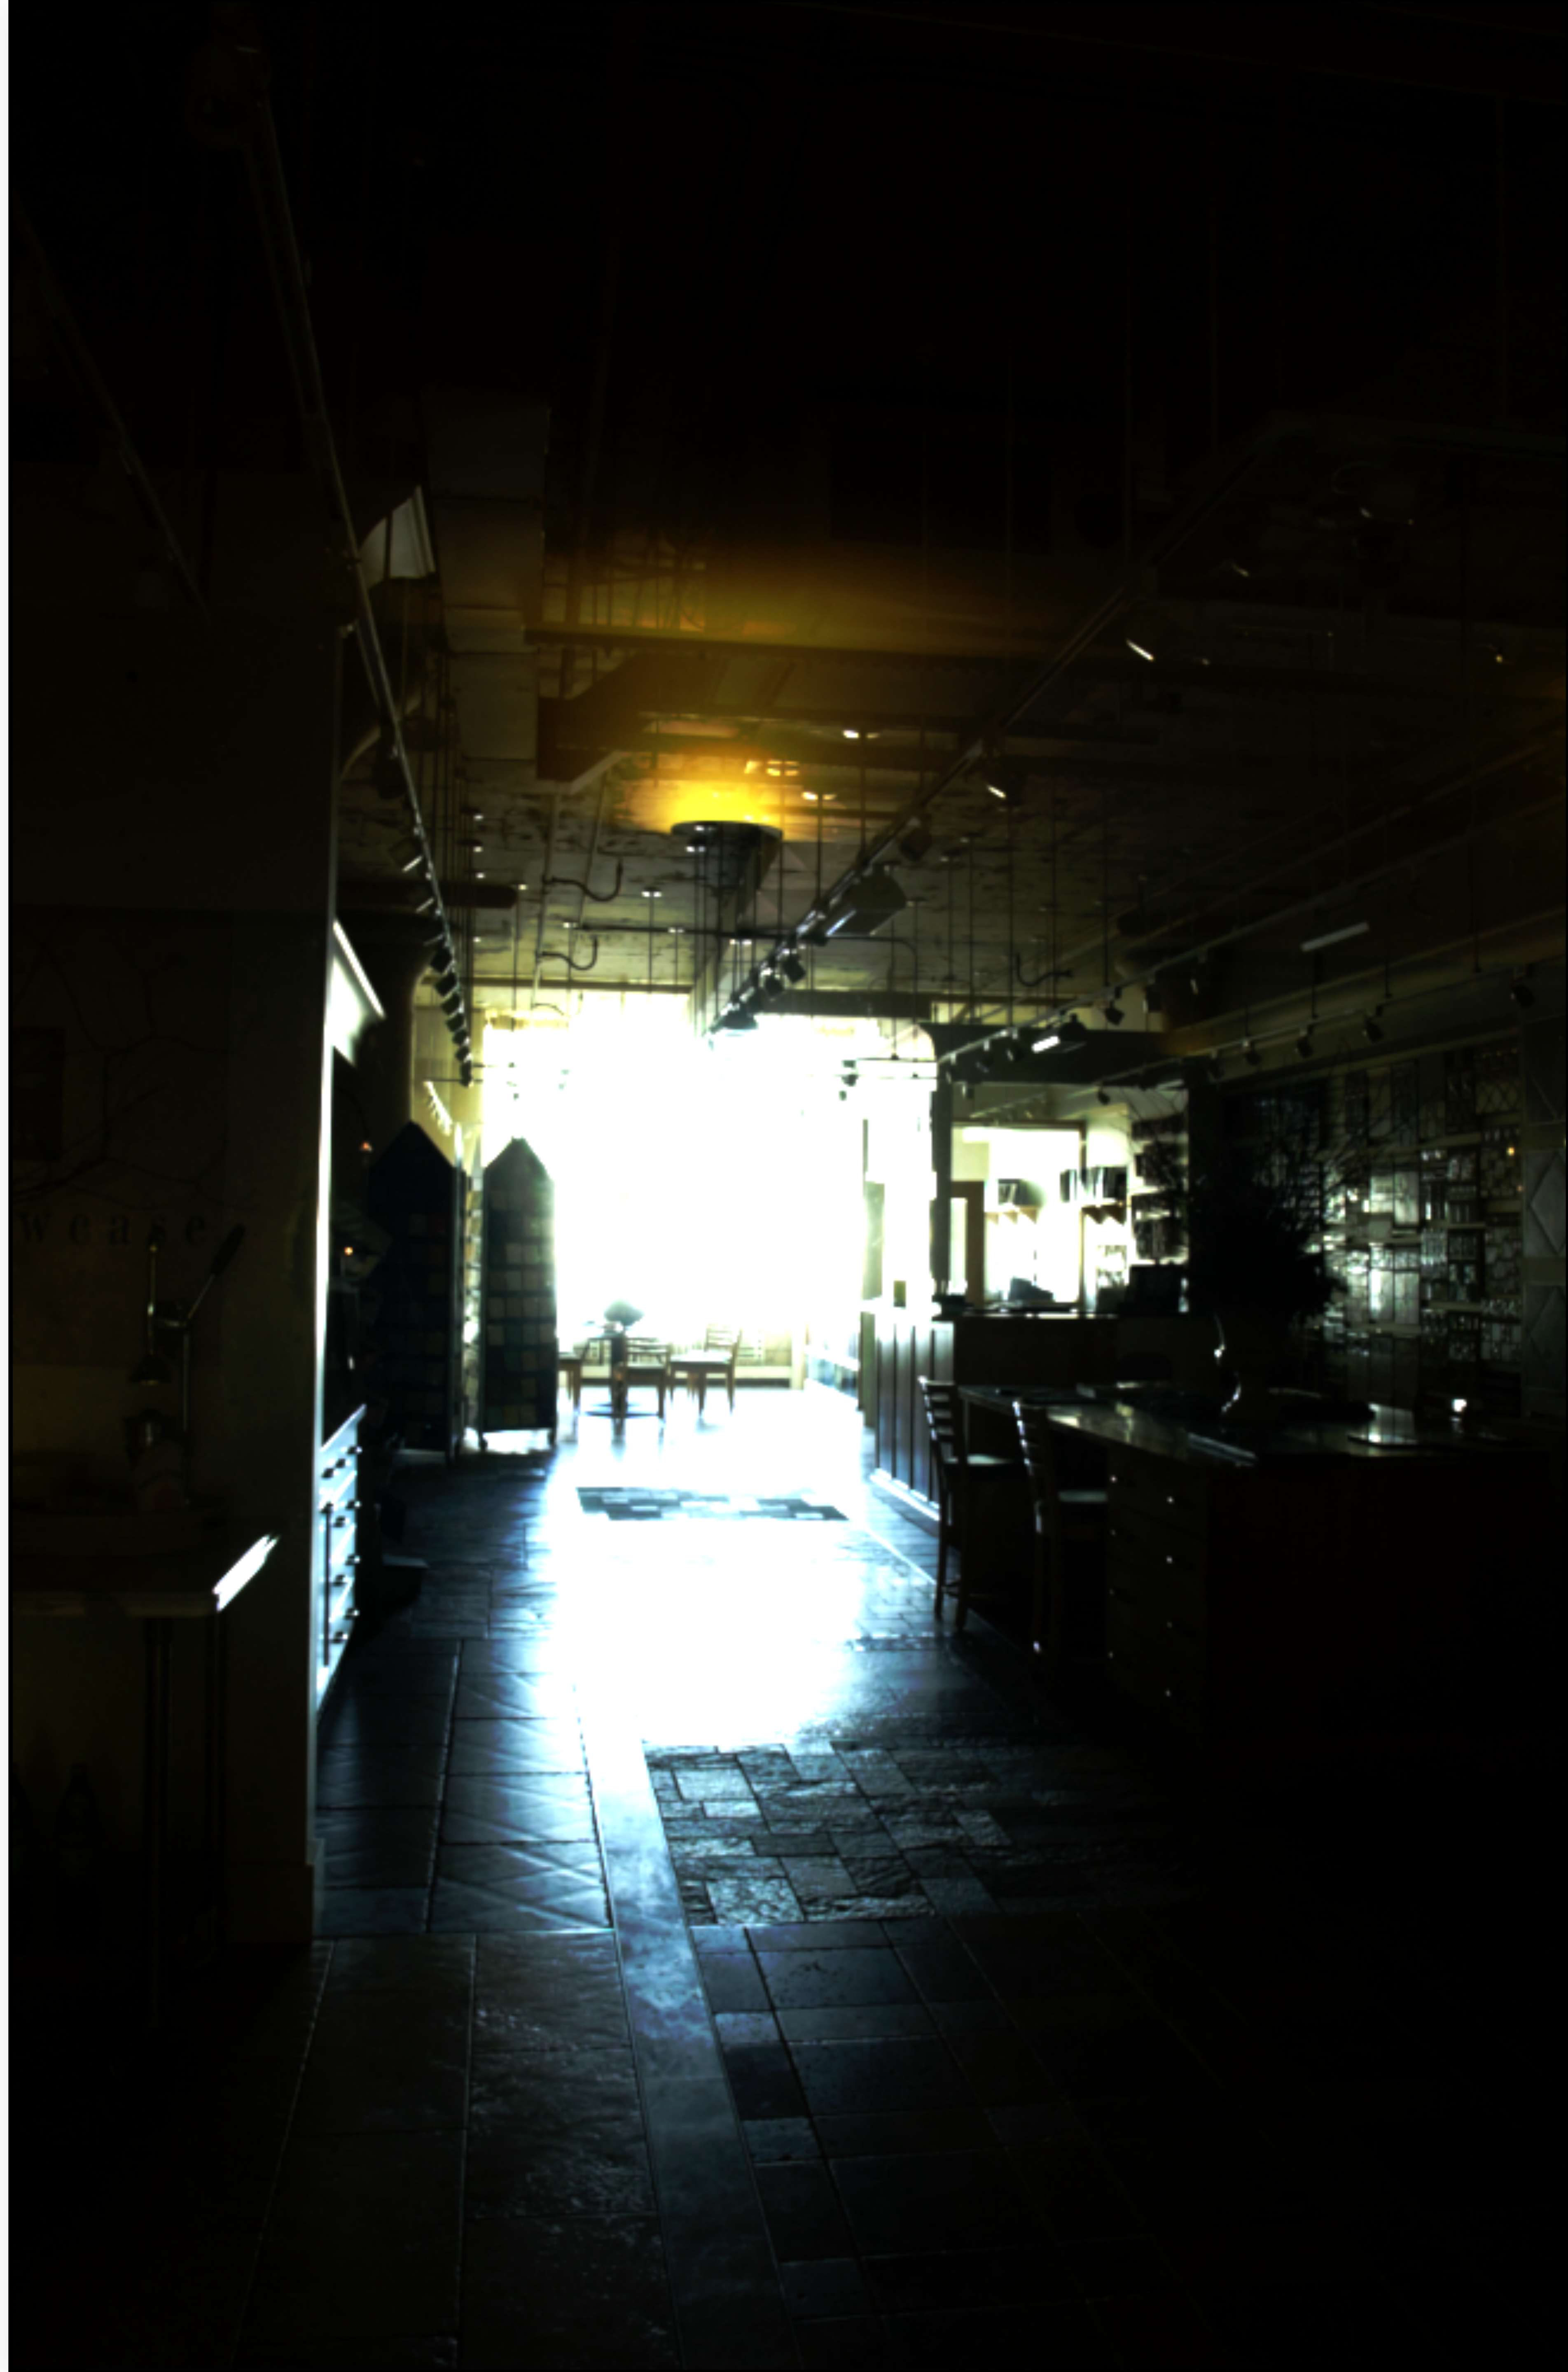

After a slide by Frédo Durand

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

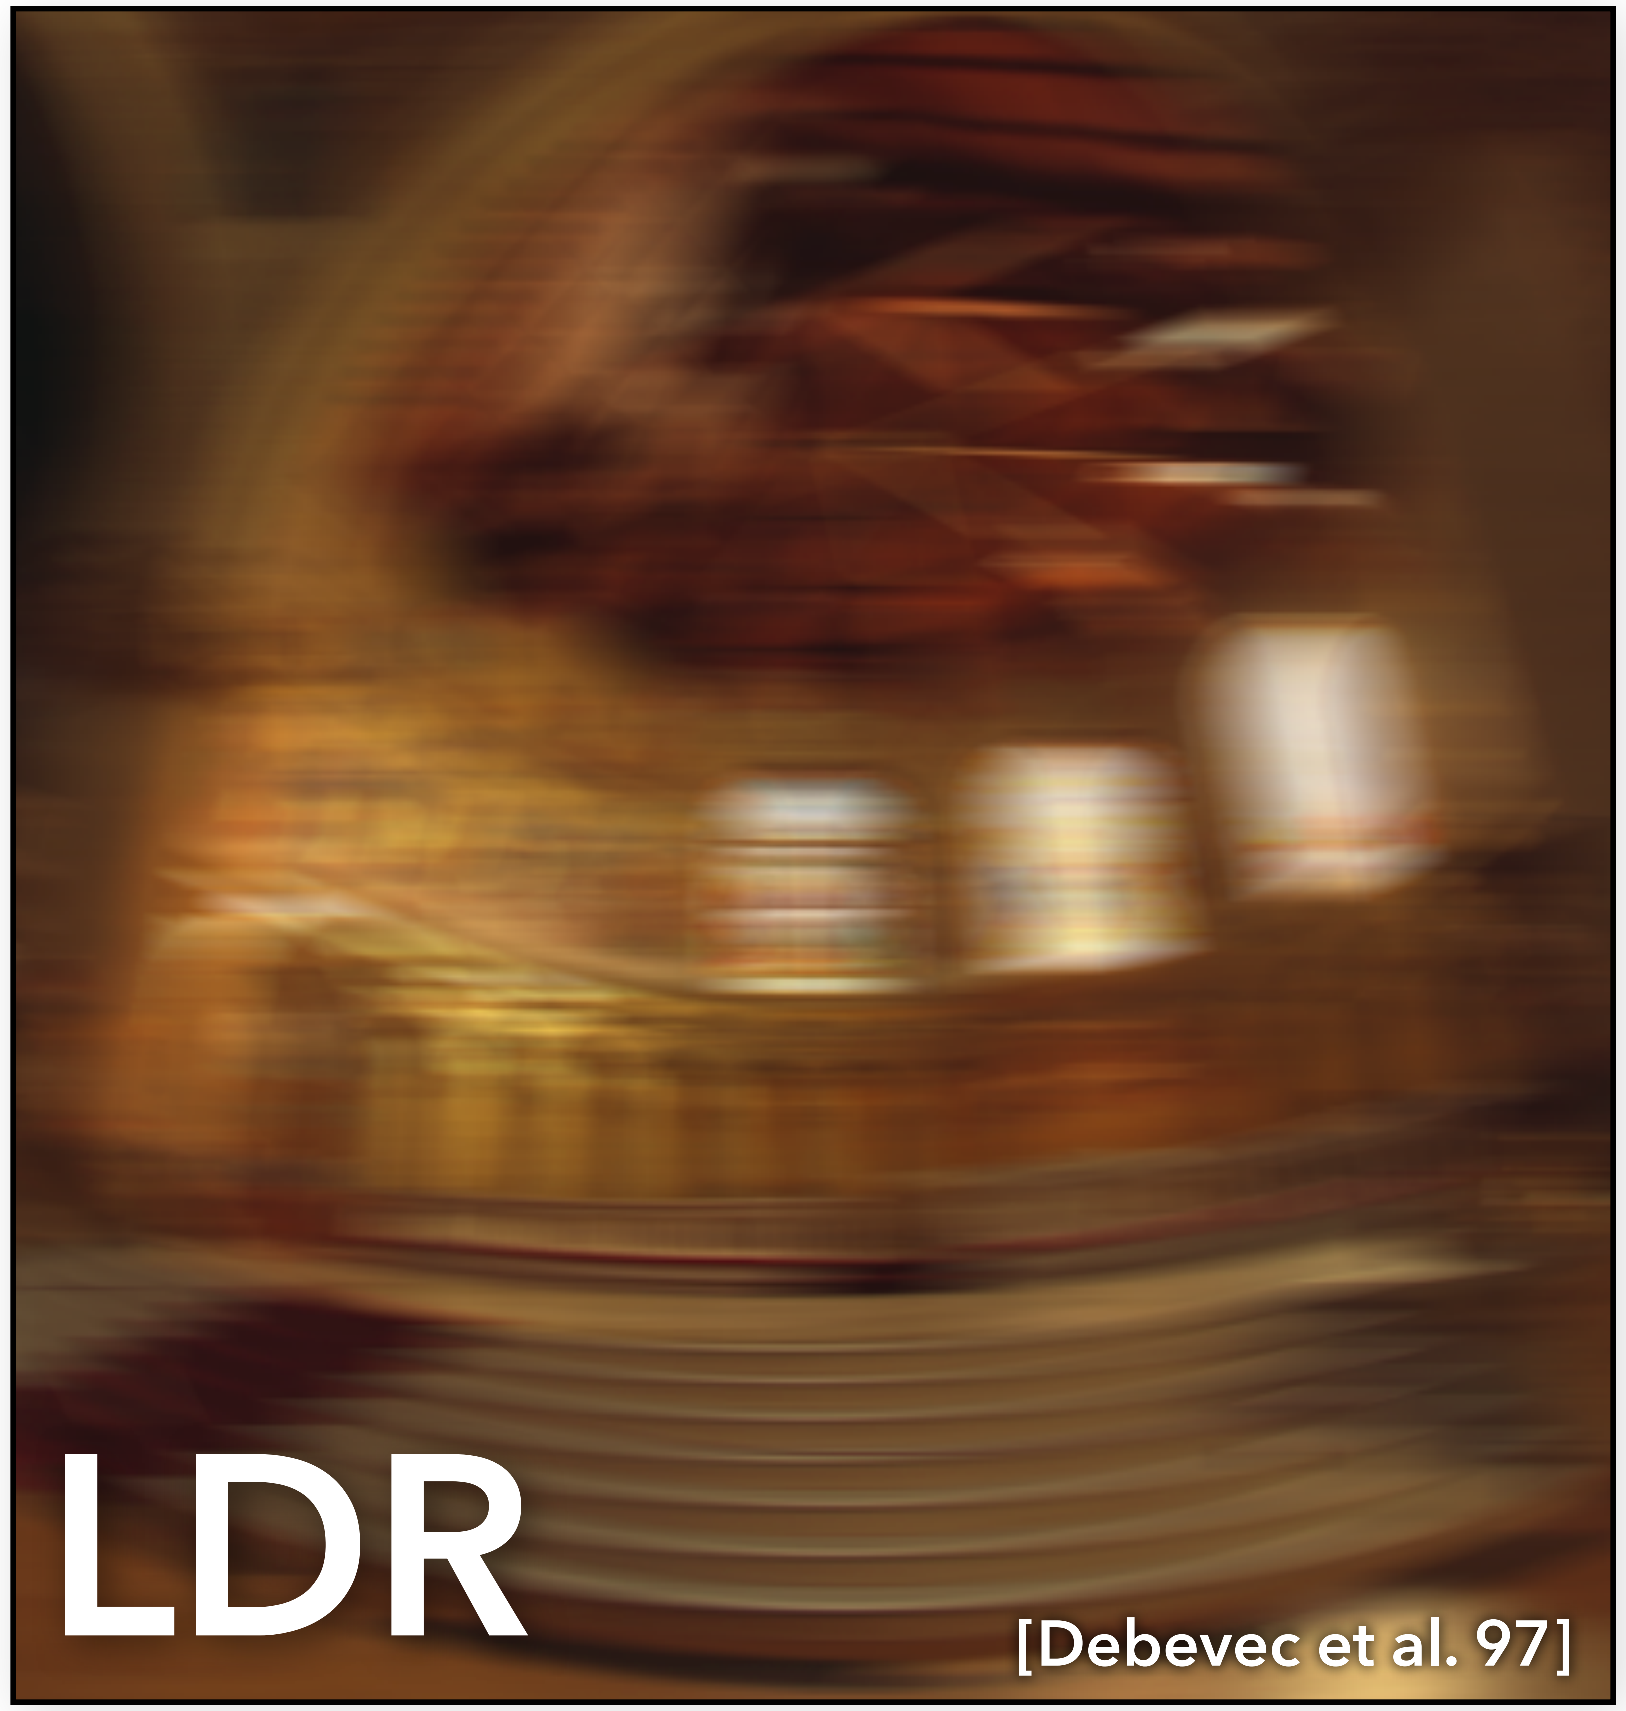

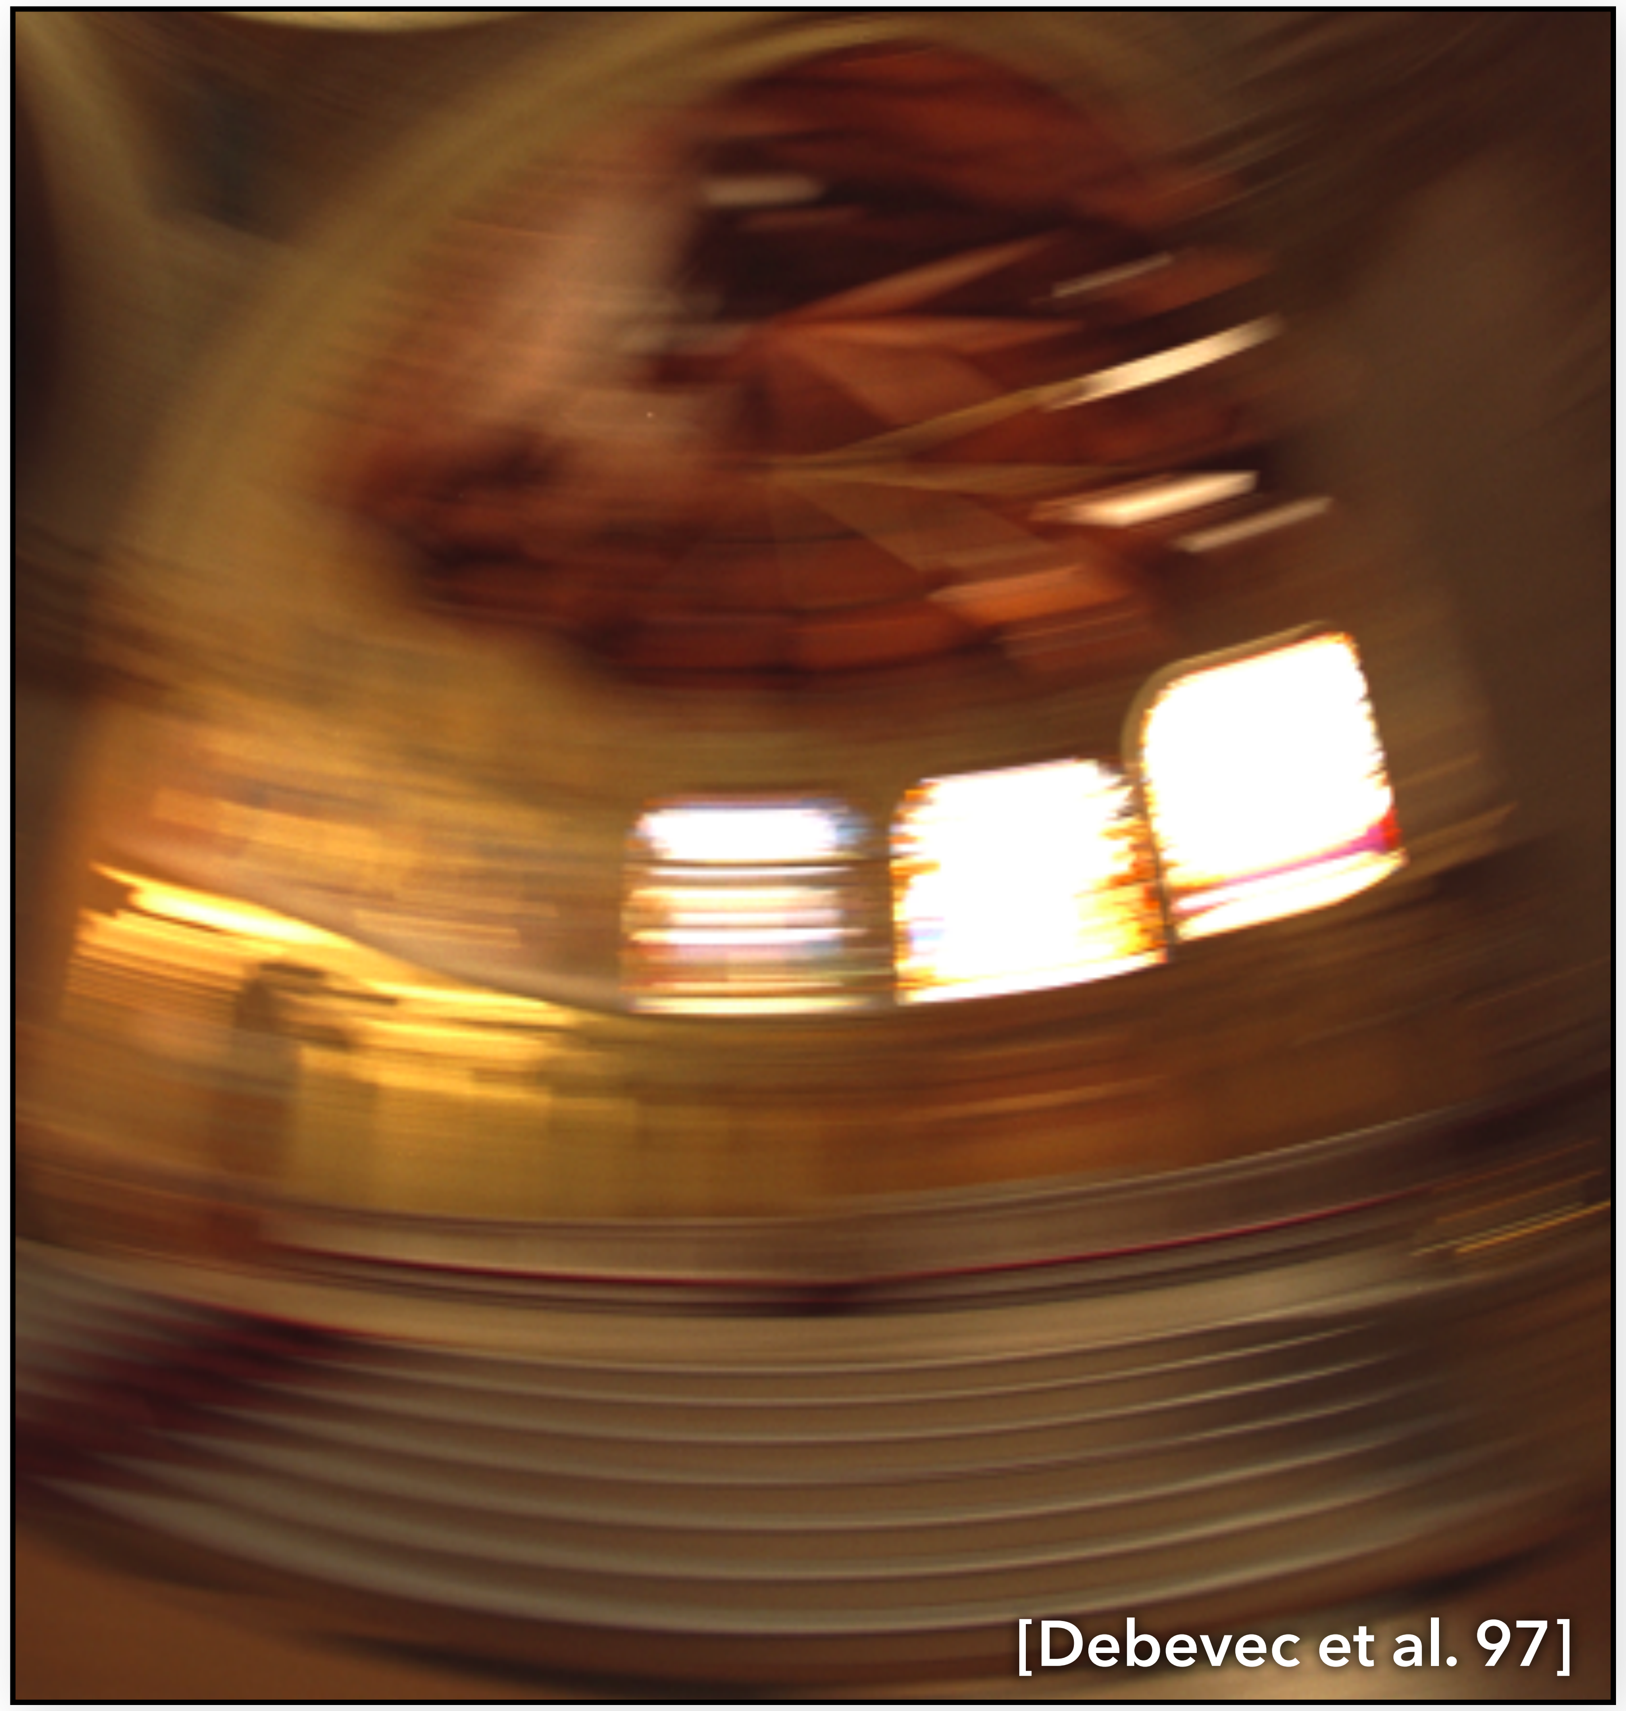

Application: Inserting Synthetic Objects


[Debevec 98]

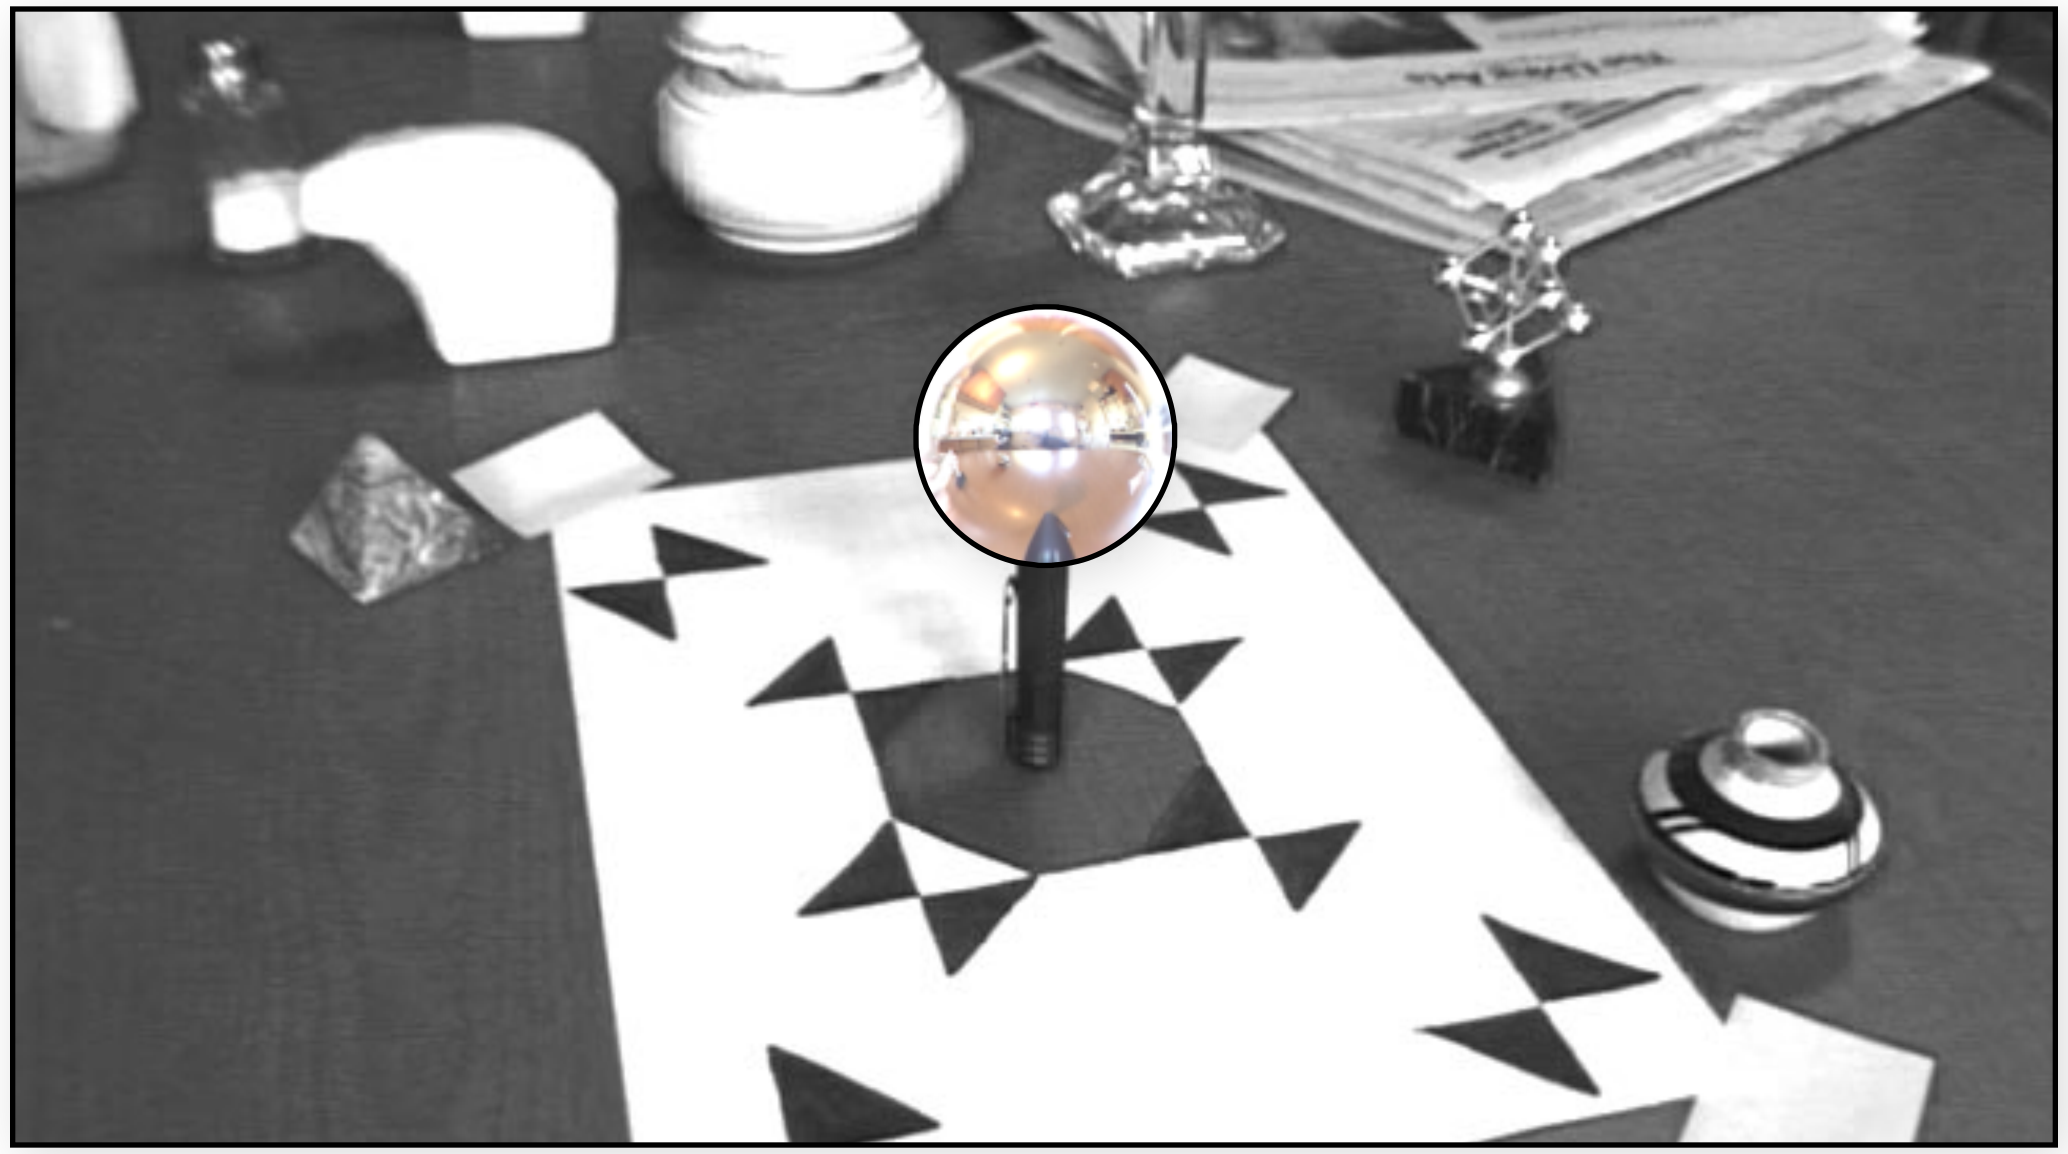

Application: Inserting Synthetic Objects


[Debevec 98]

Problem 1: Record the information

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

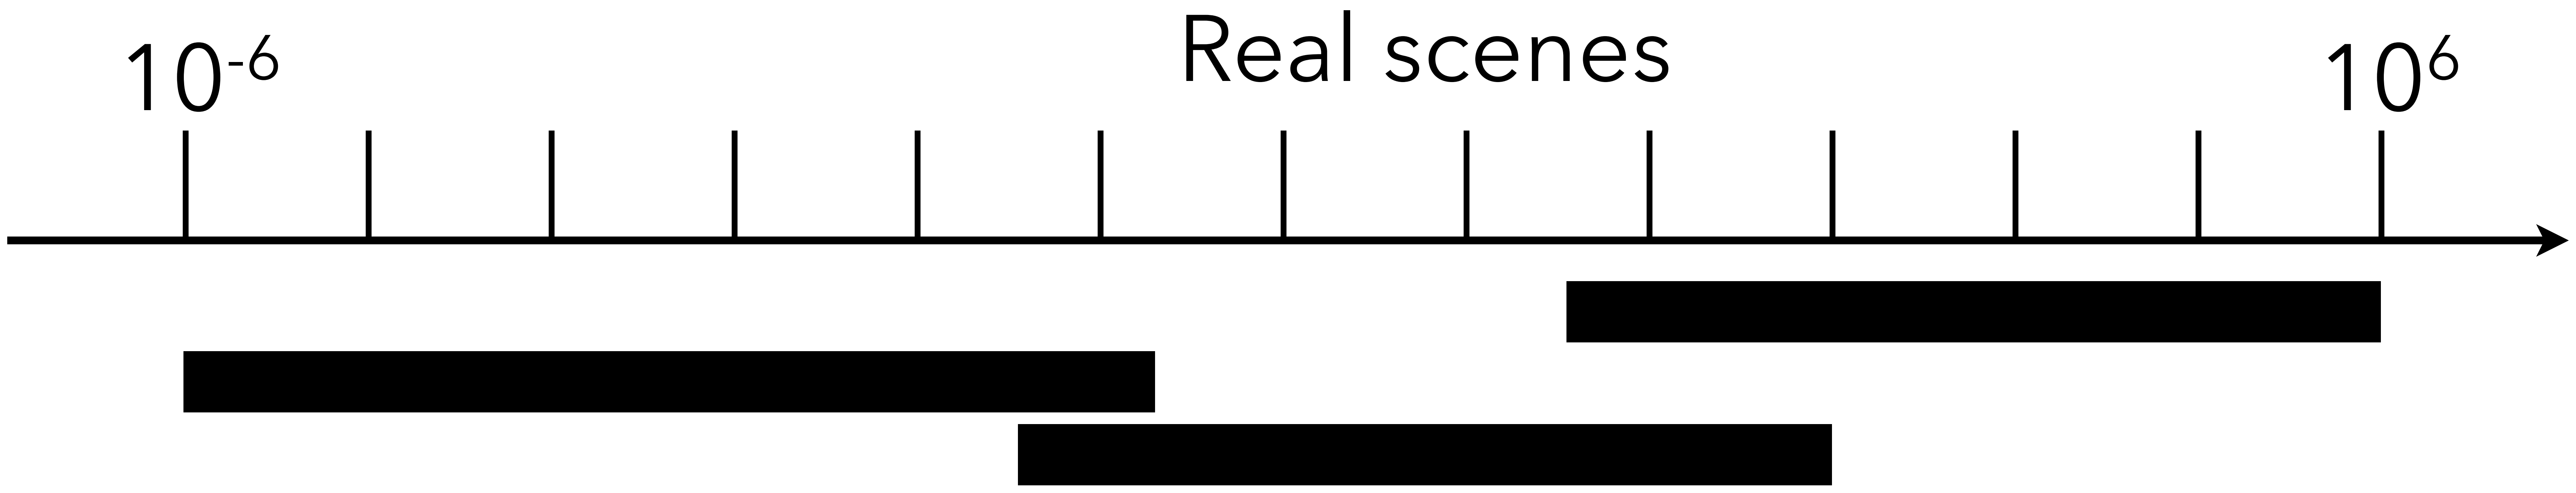

Film/sensors can record 2-3 orders of magnitude

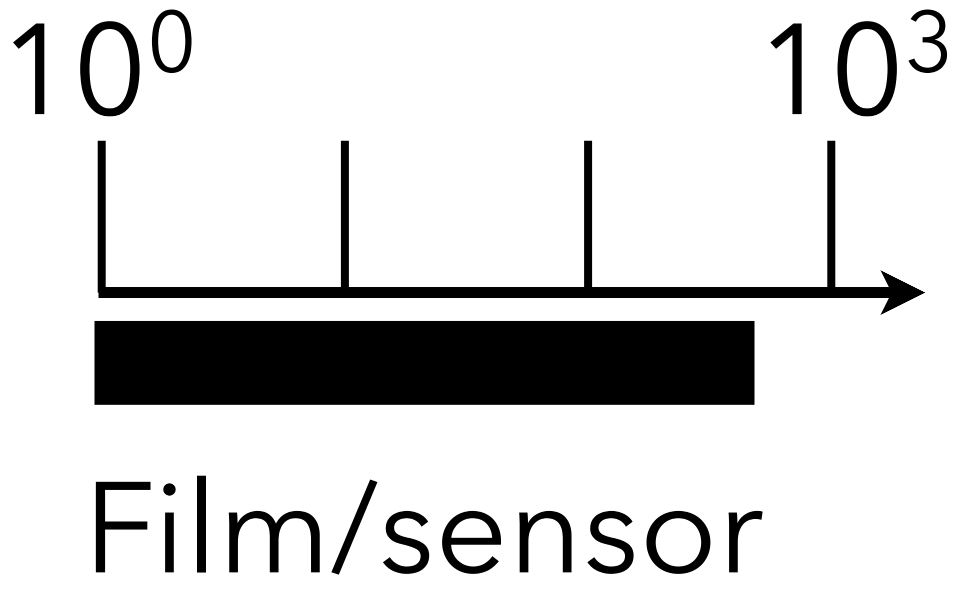

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

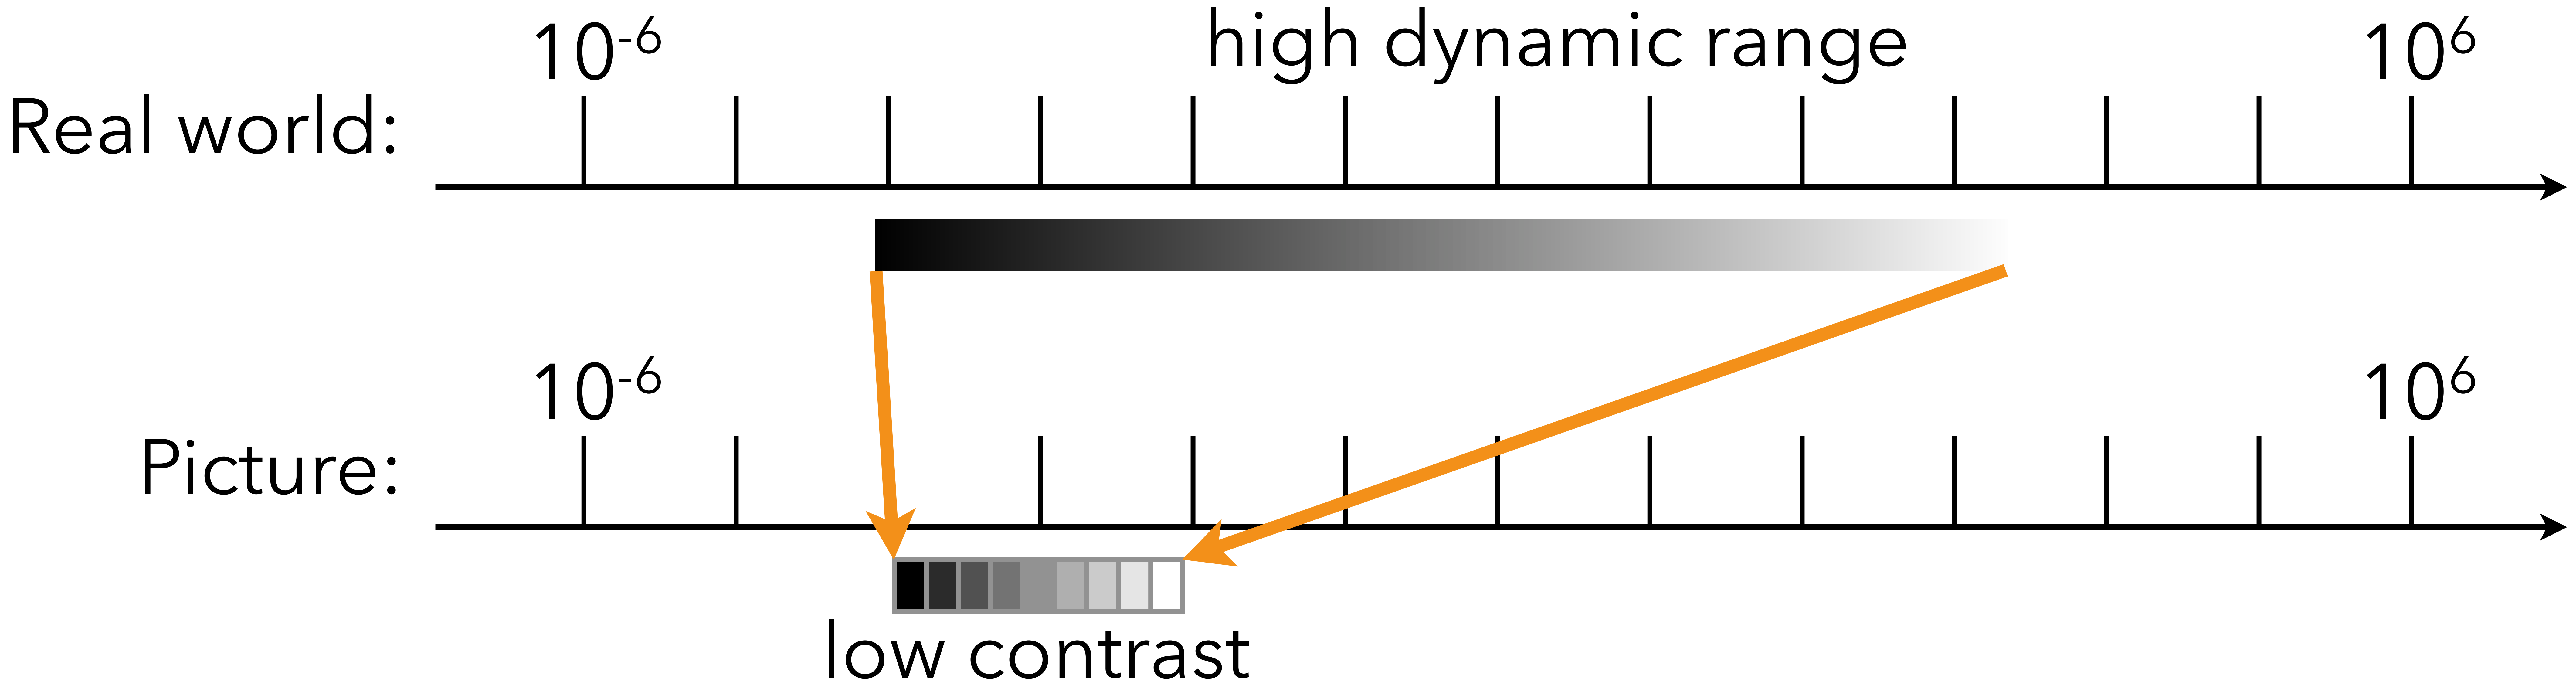

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

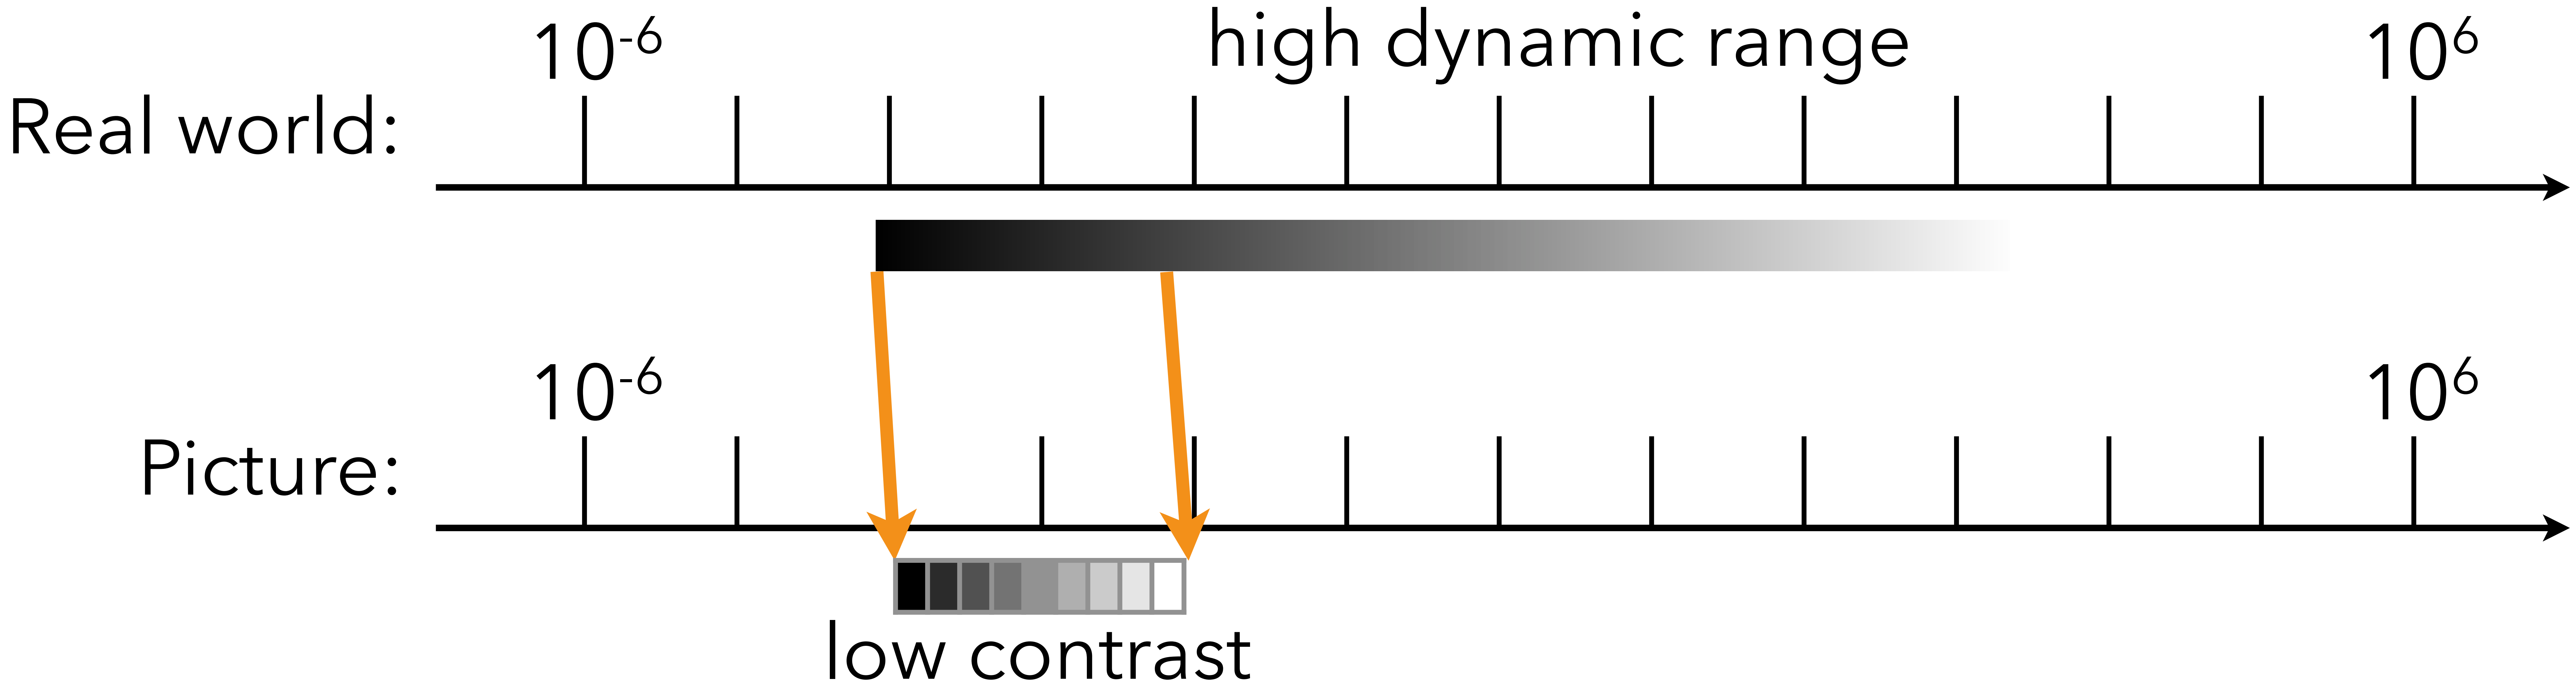

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

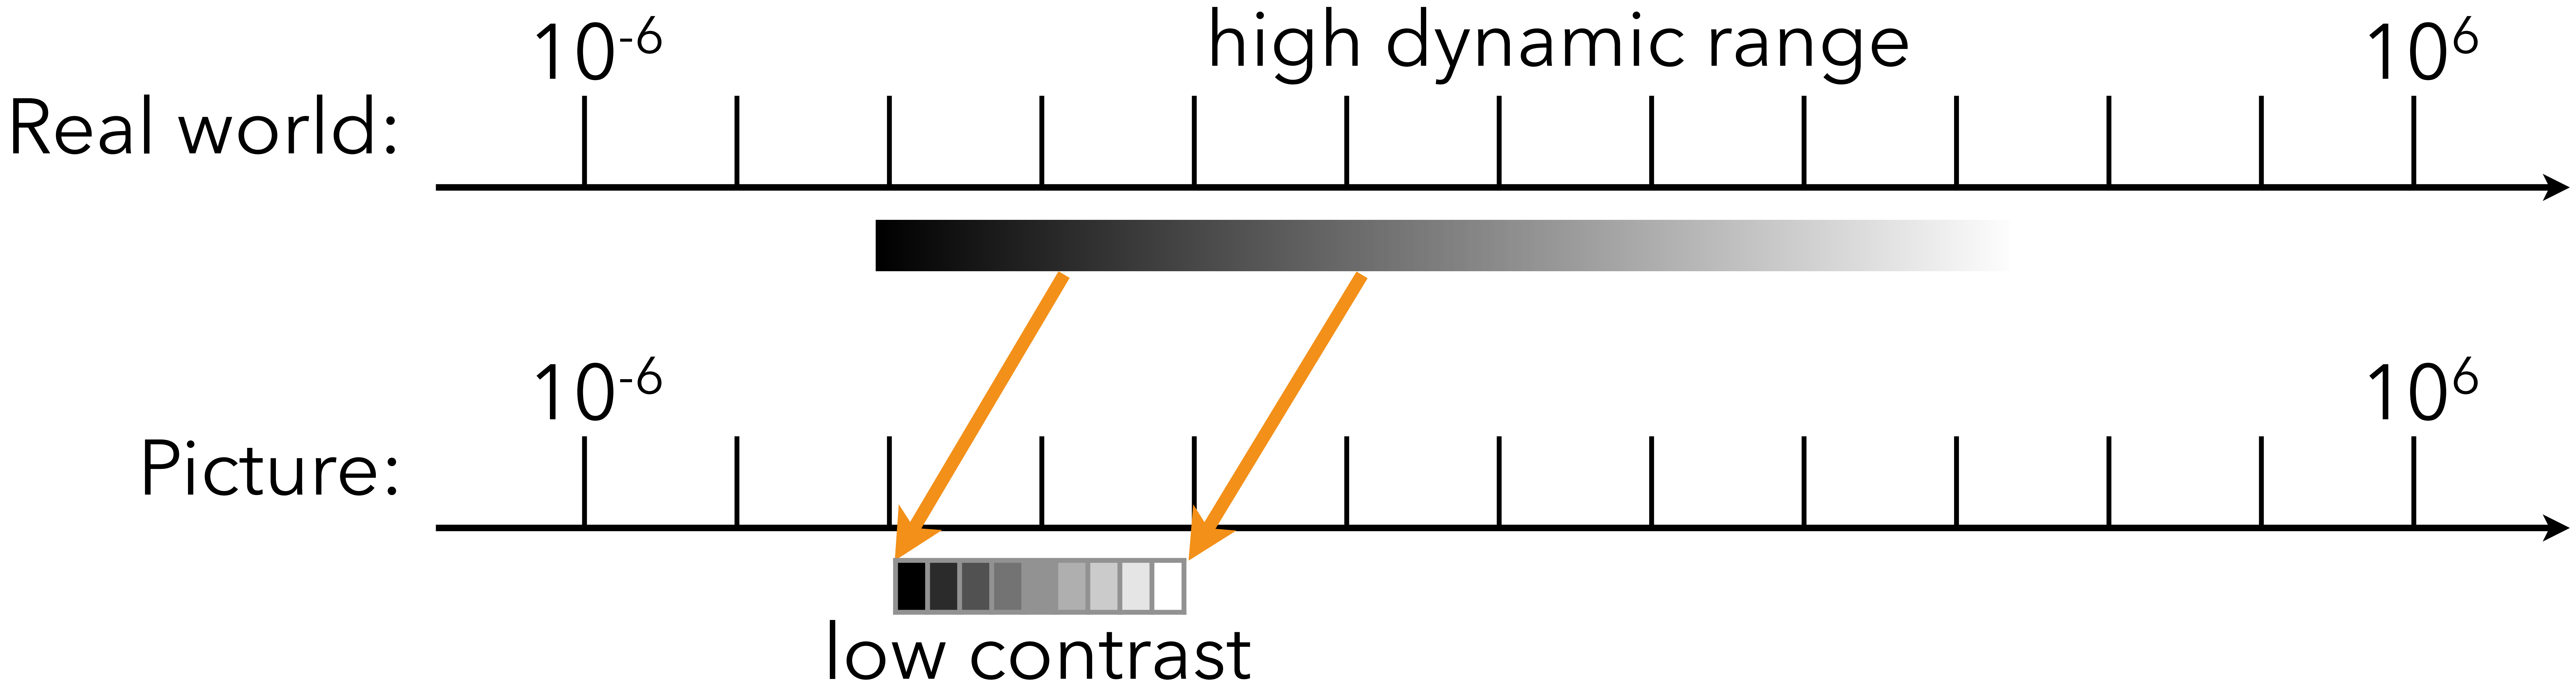

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

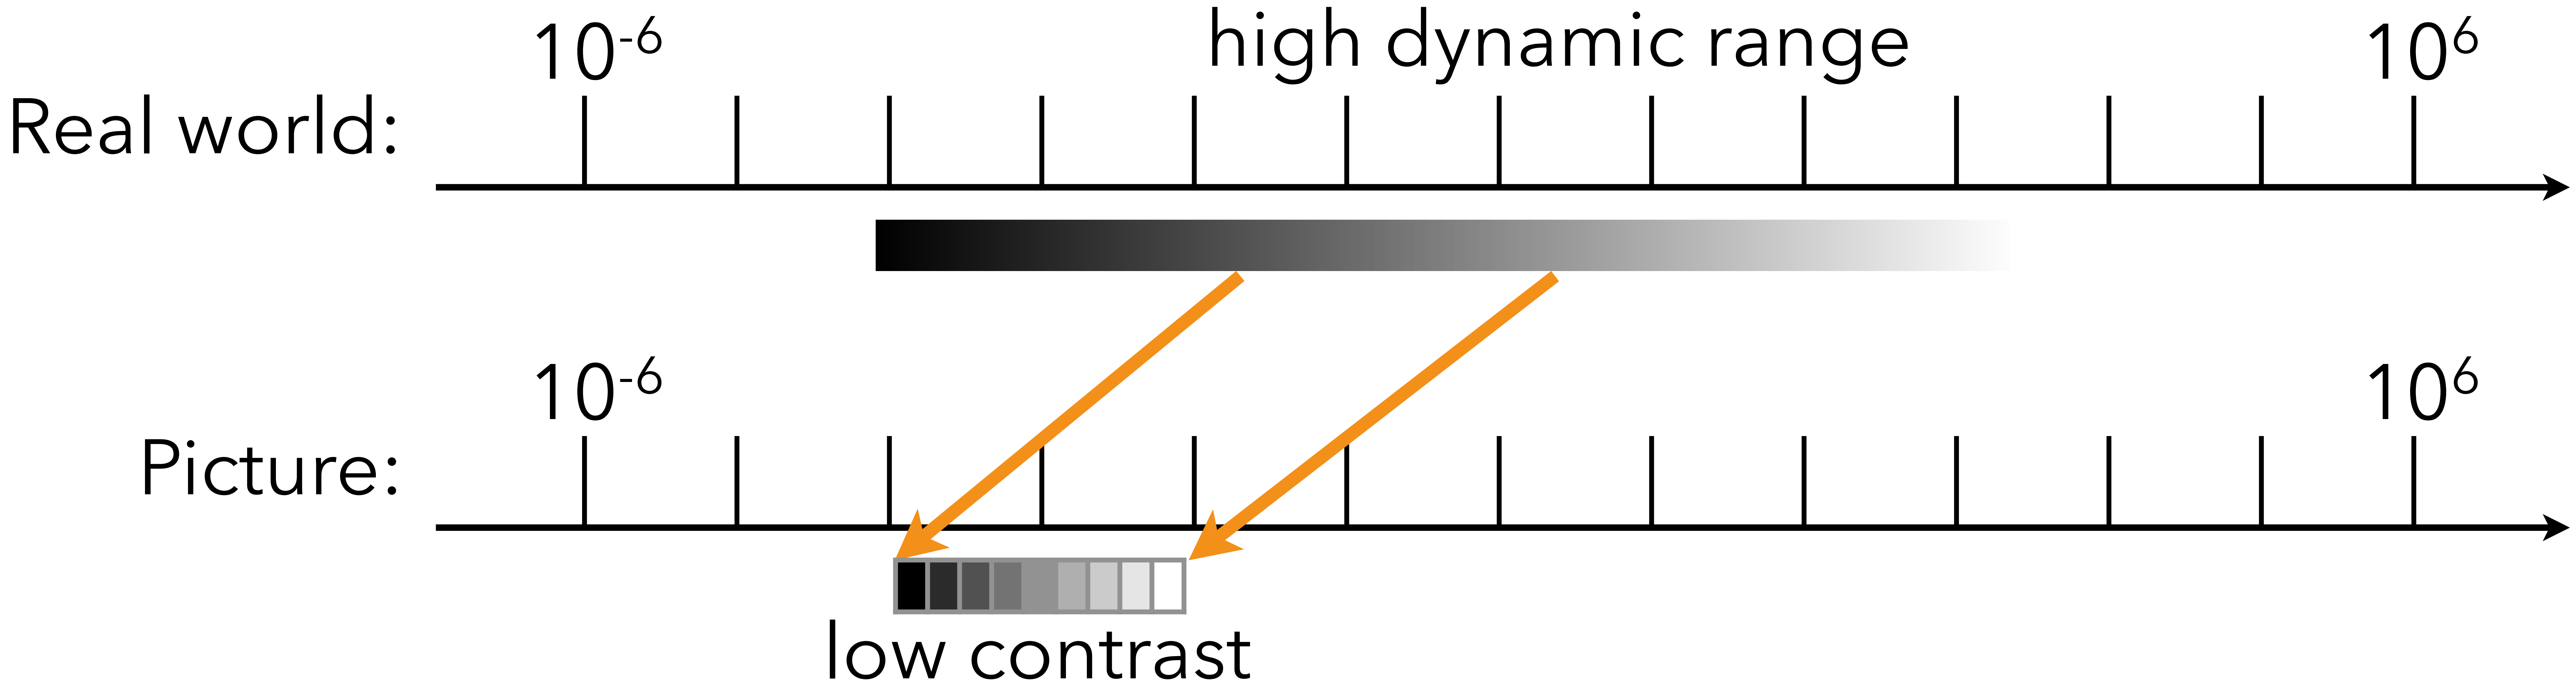

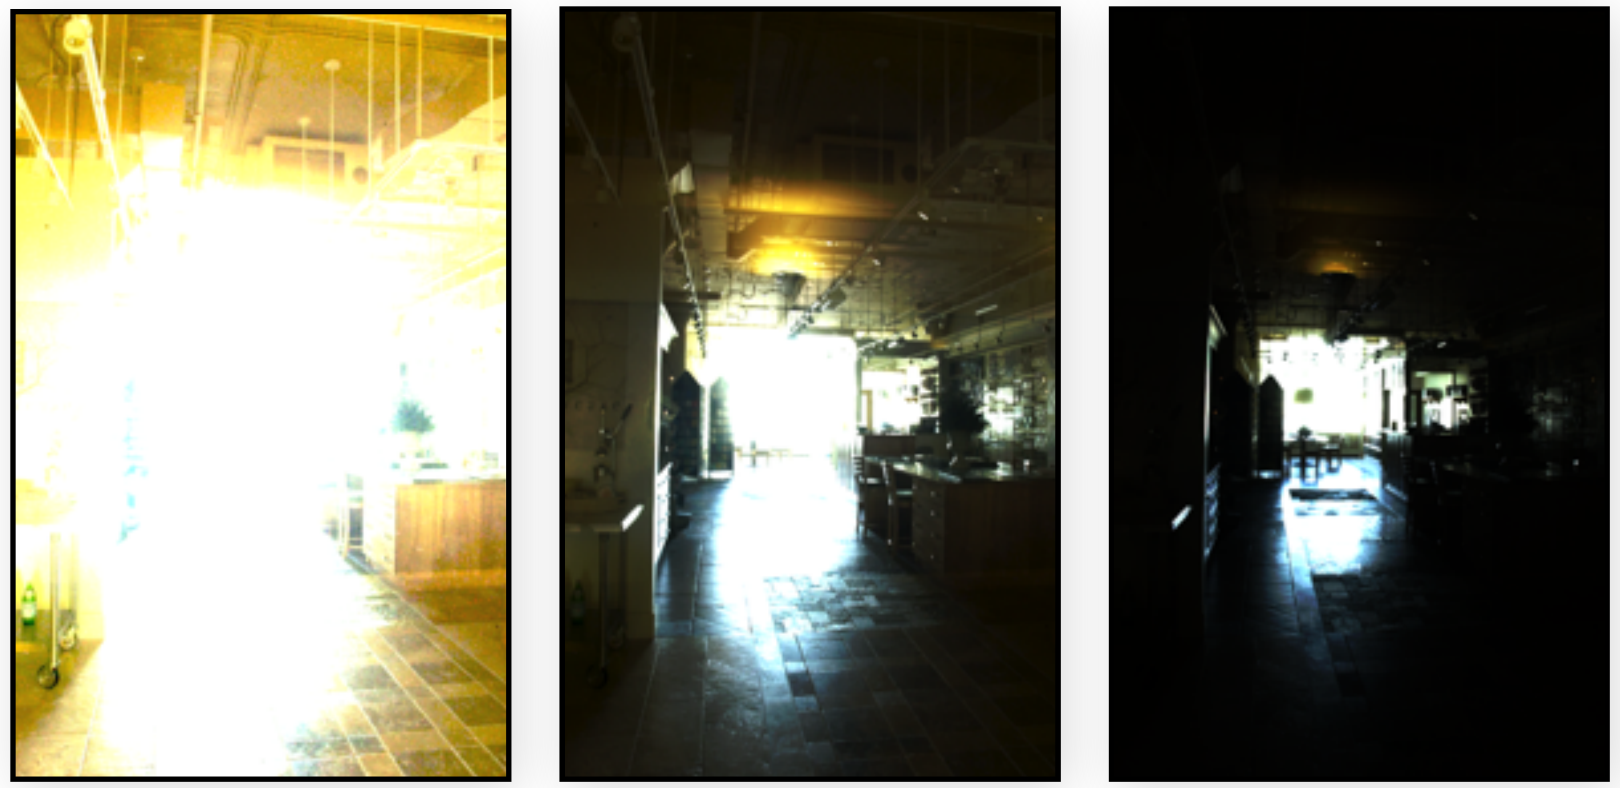

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

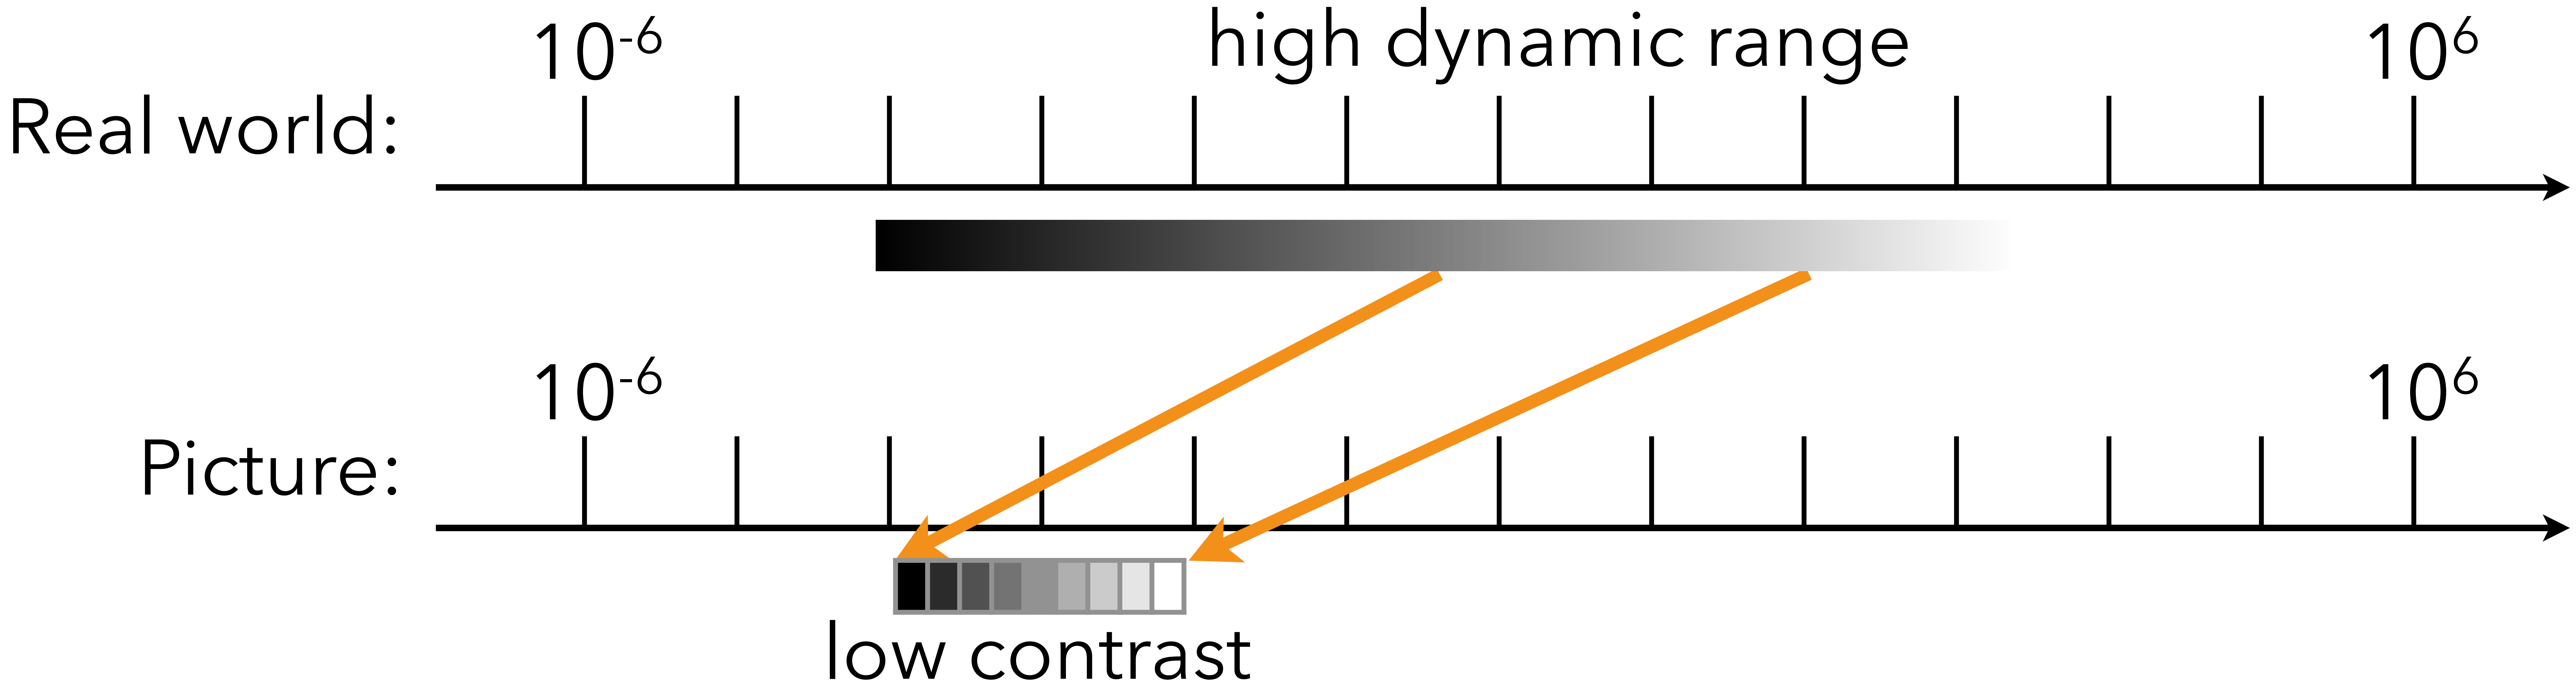

After a slide by Frédo Durand

Multiple exposure photography

Sequentially measure all segments of the range

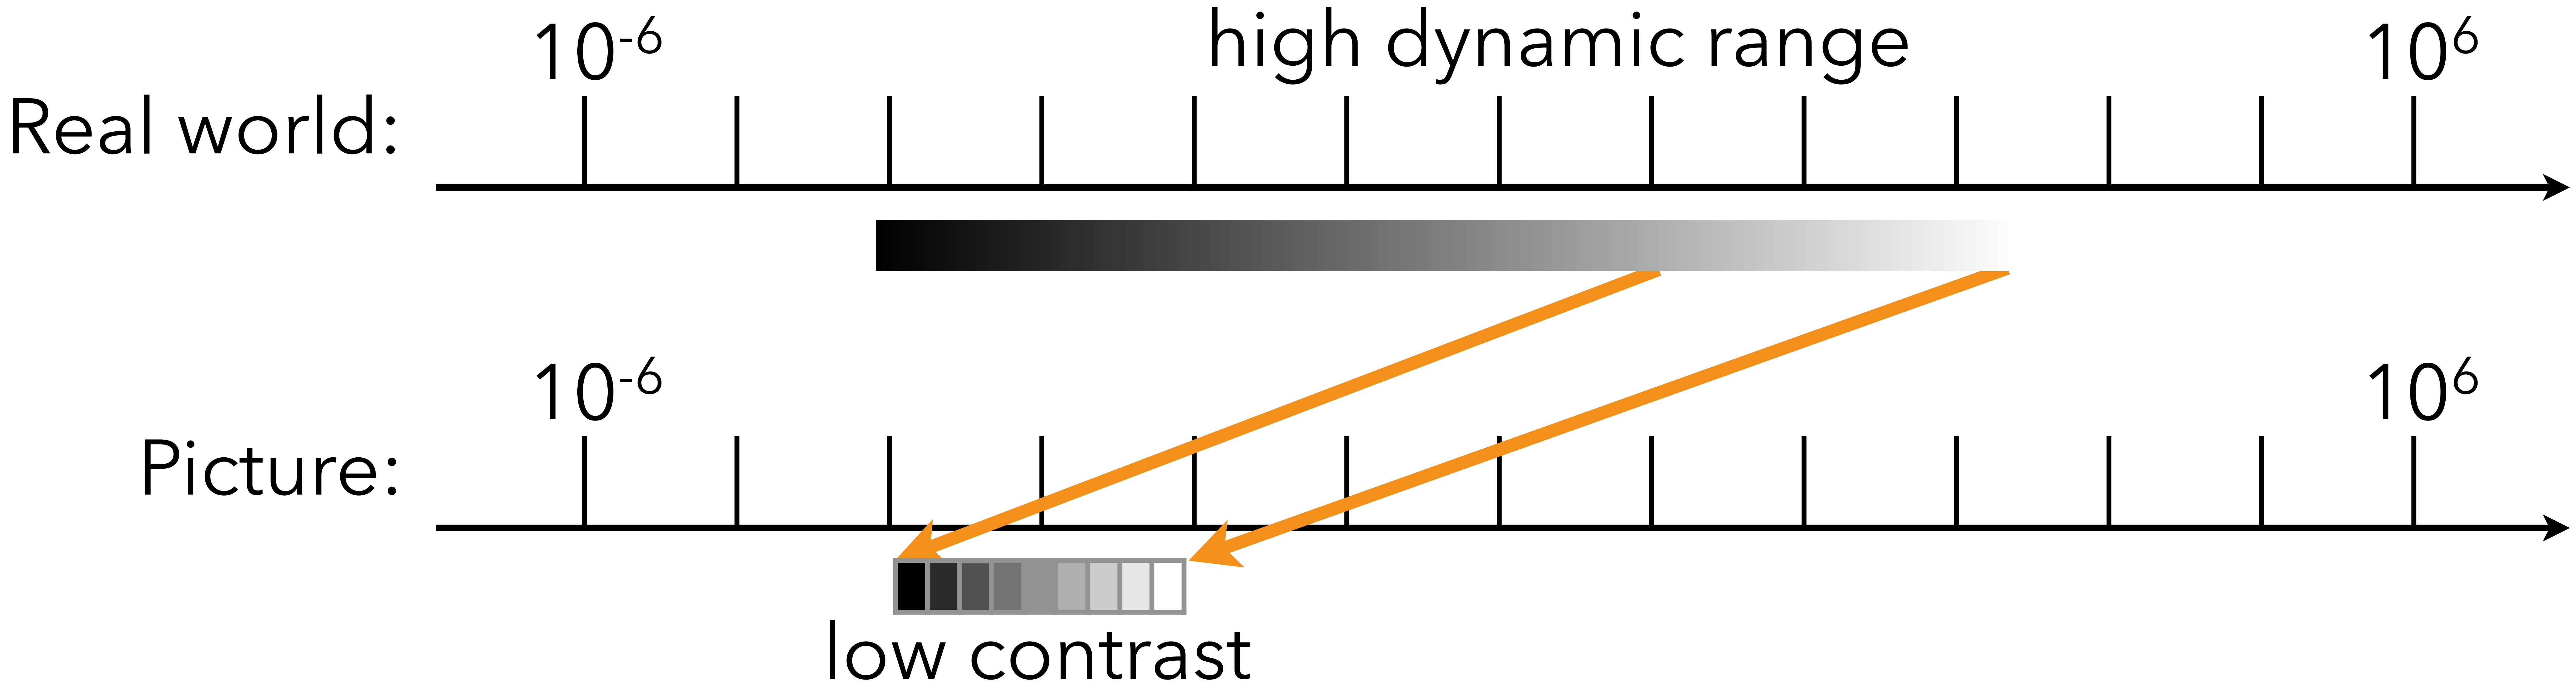

After a slide by Frédo Durand

Problem 2: Display the information

Match limited contrast of the medium while preserving details: **the tone mapping problem**

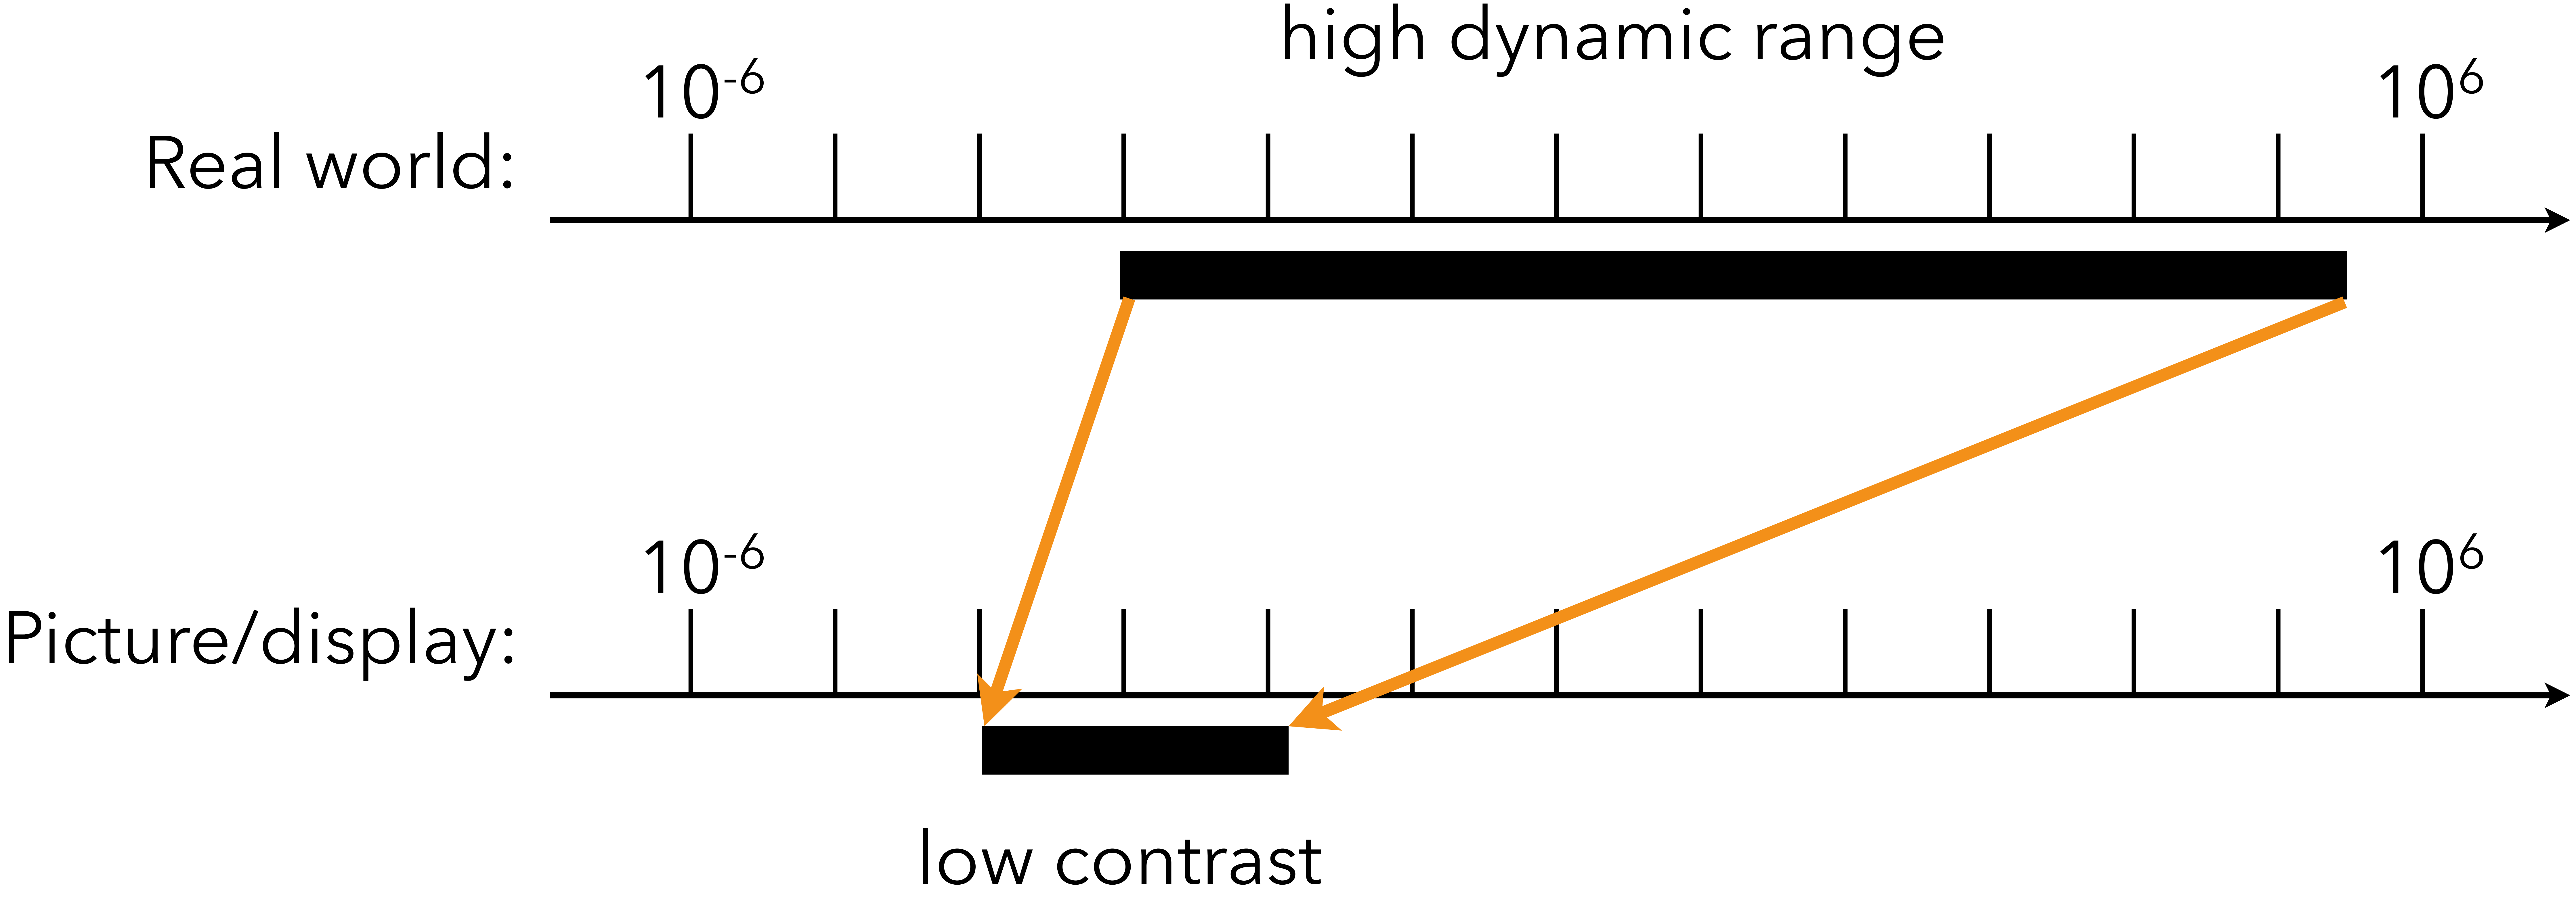

After a slide by Frédo Durand

“Wielki Staw Polski” [Wojciech Jarosz 2011]

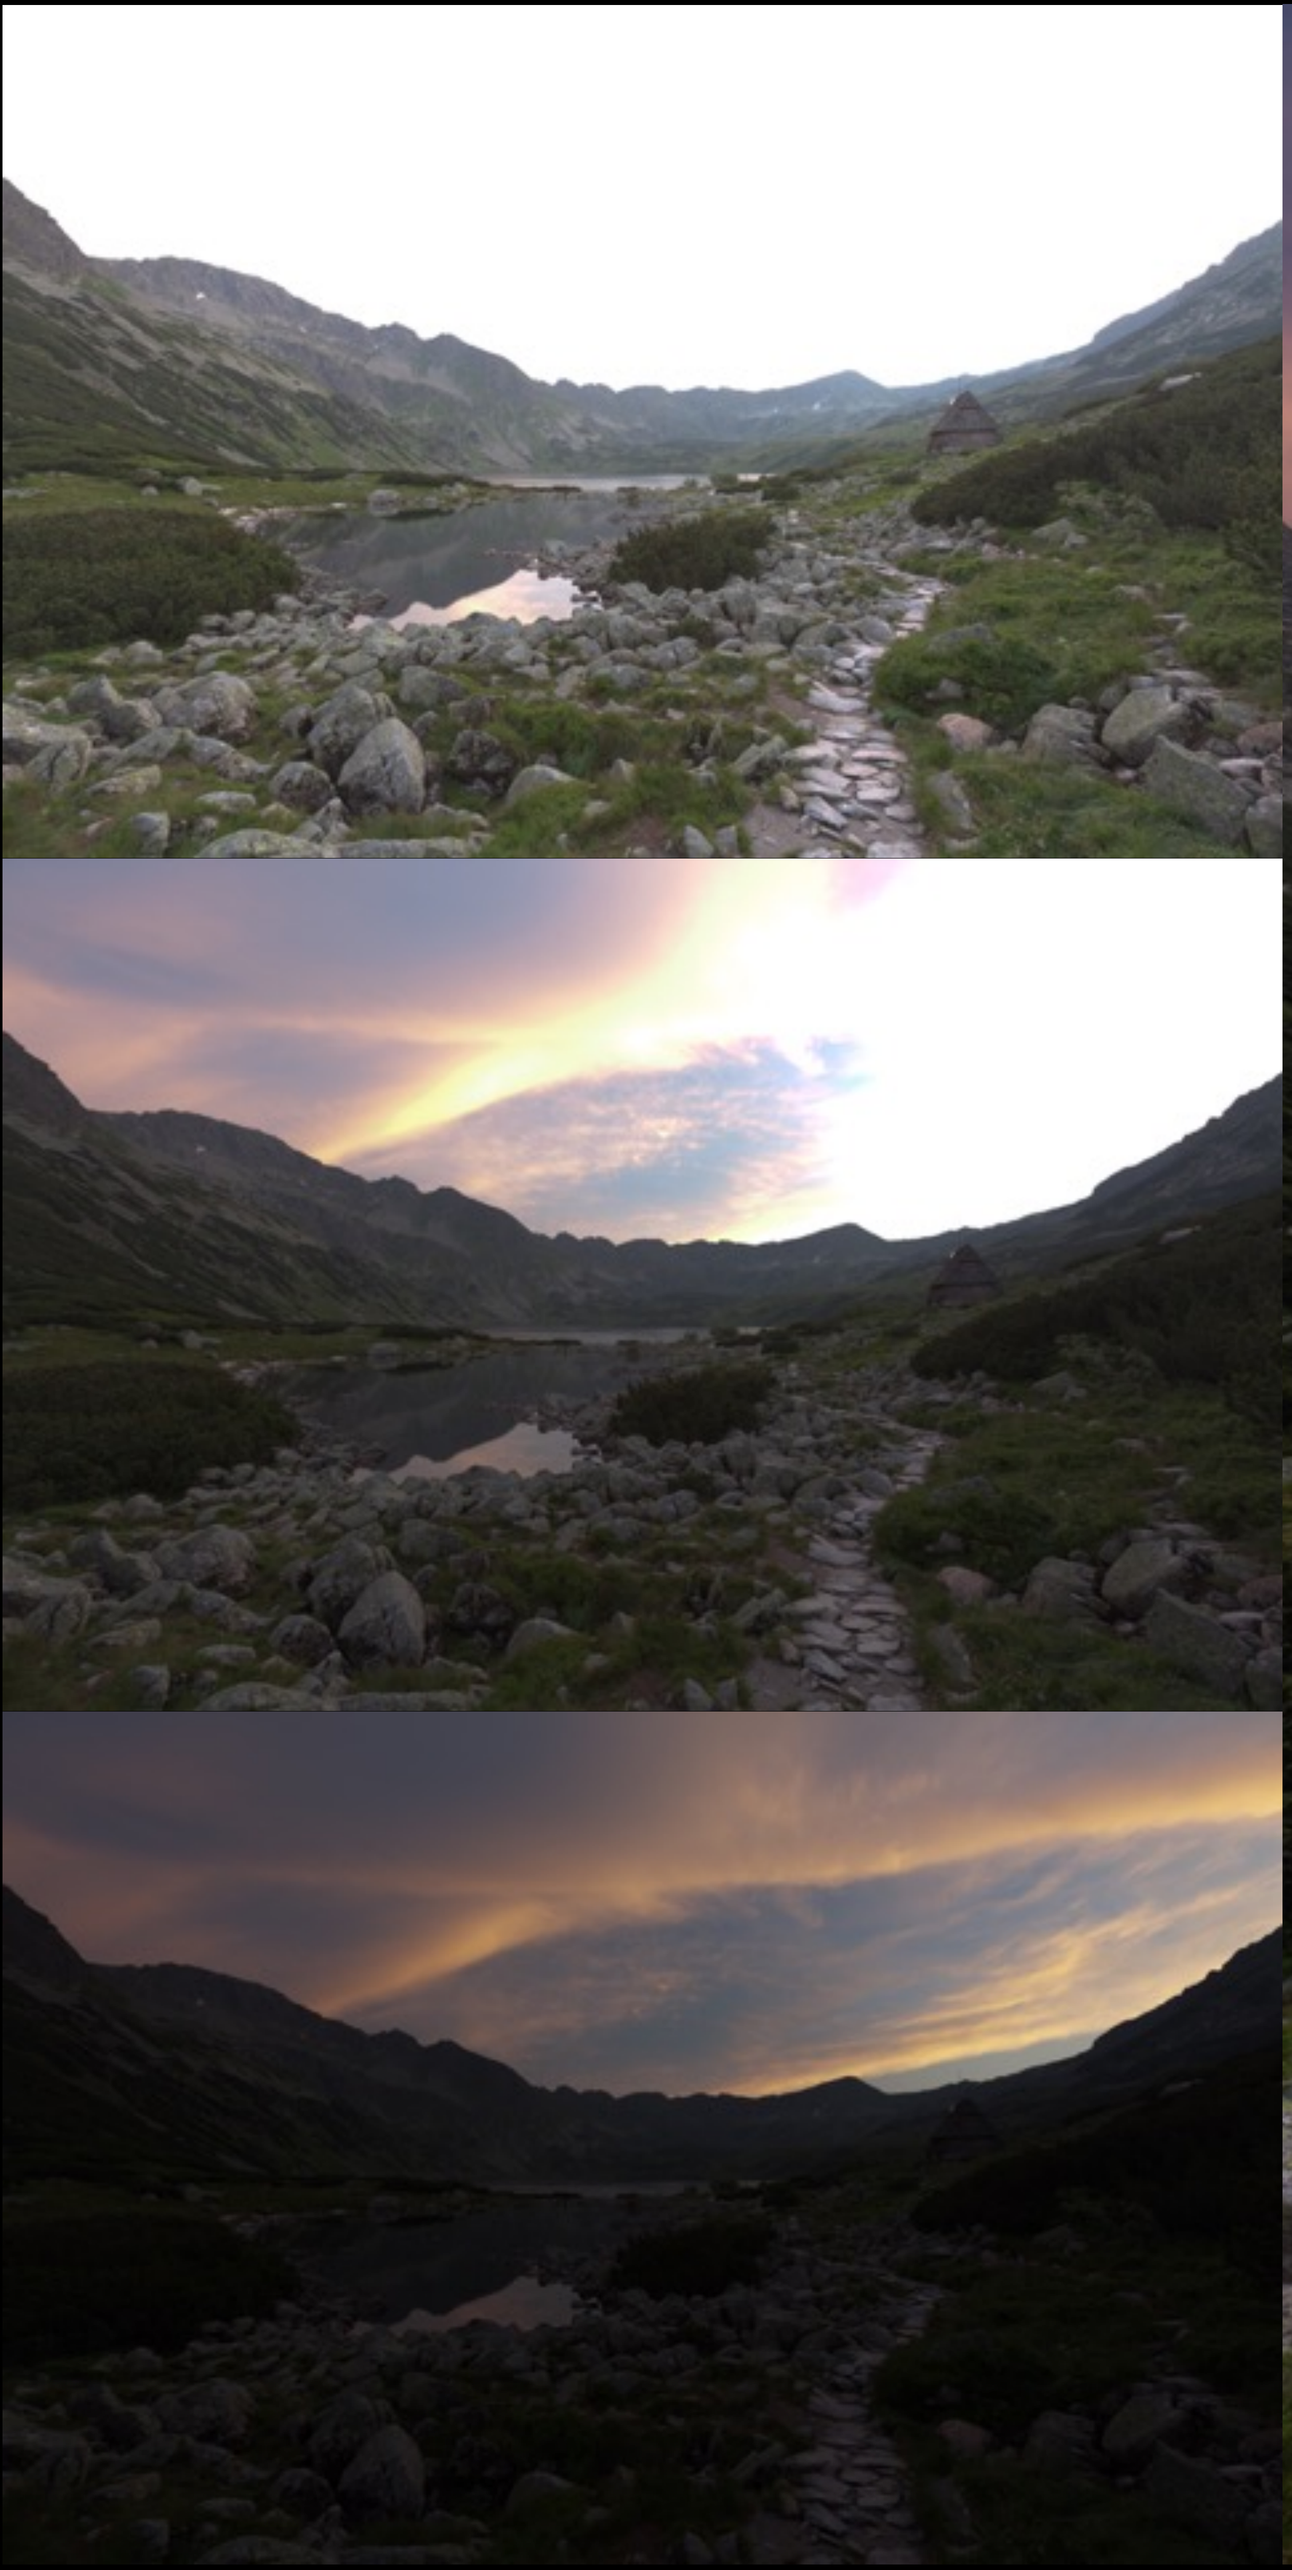

"Wielki Staw Polski"

[Wojciech Jarosz 2011]

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

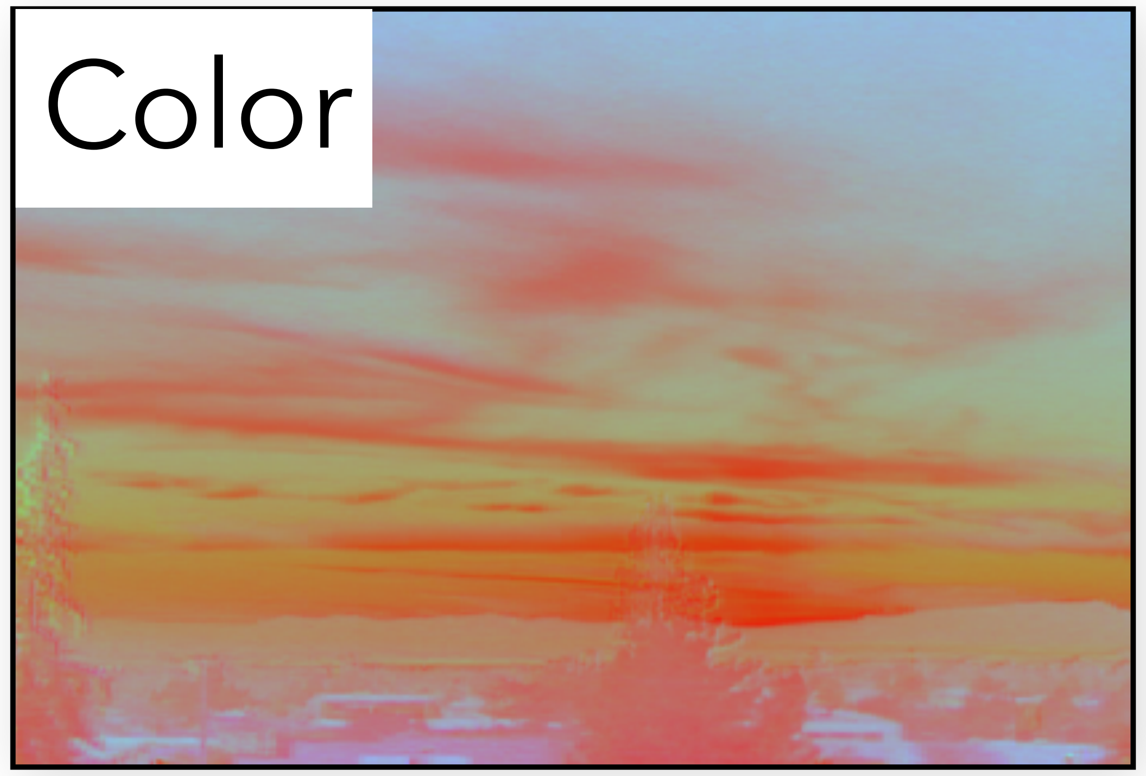

After a slide by Frédo Durand

The importance of local contrast

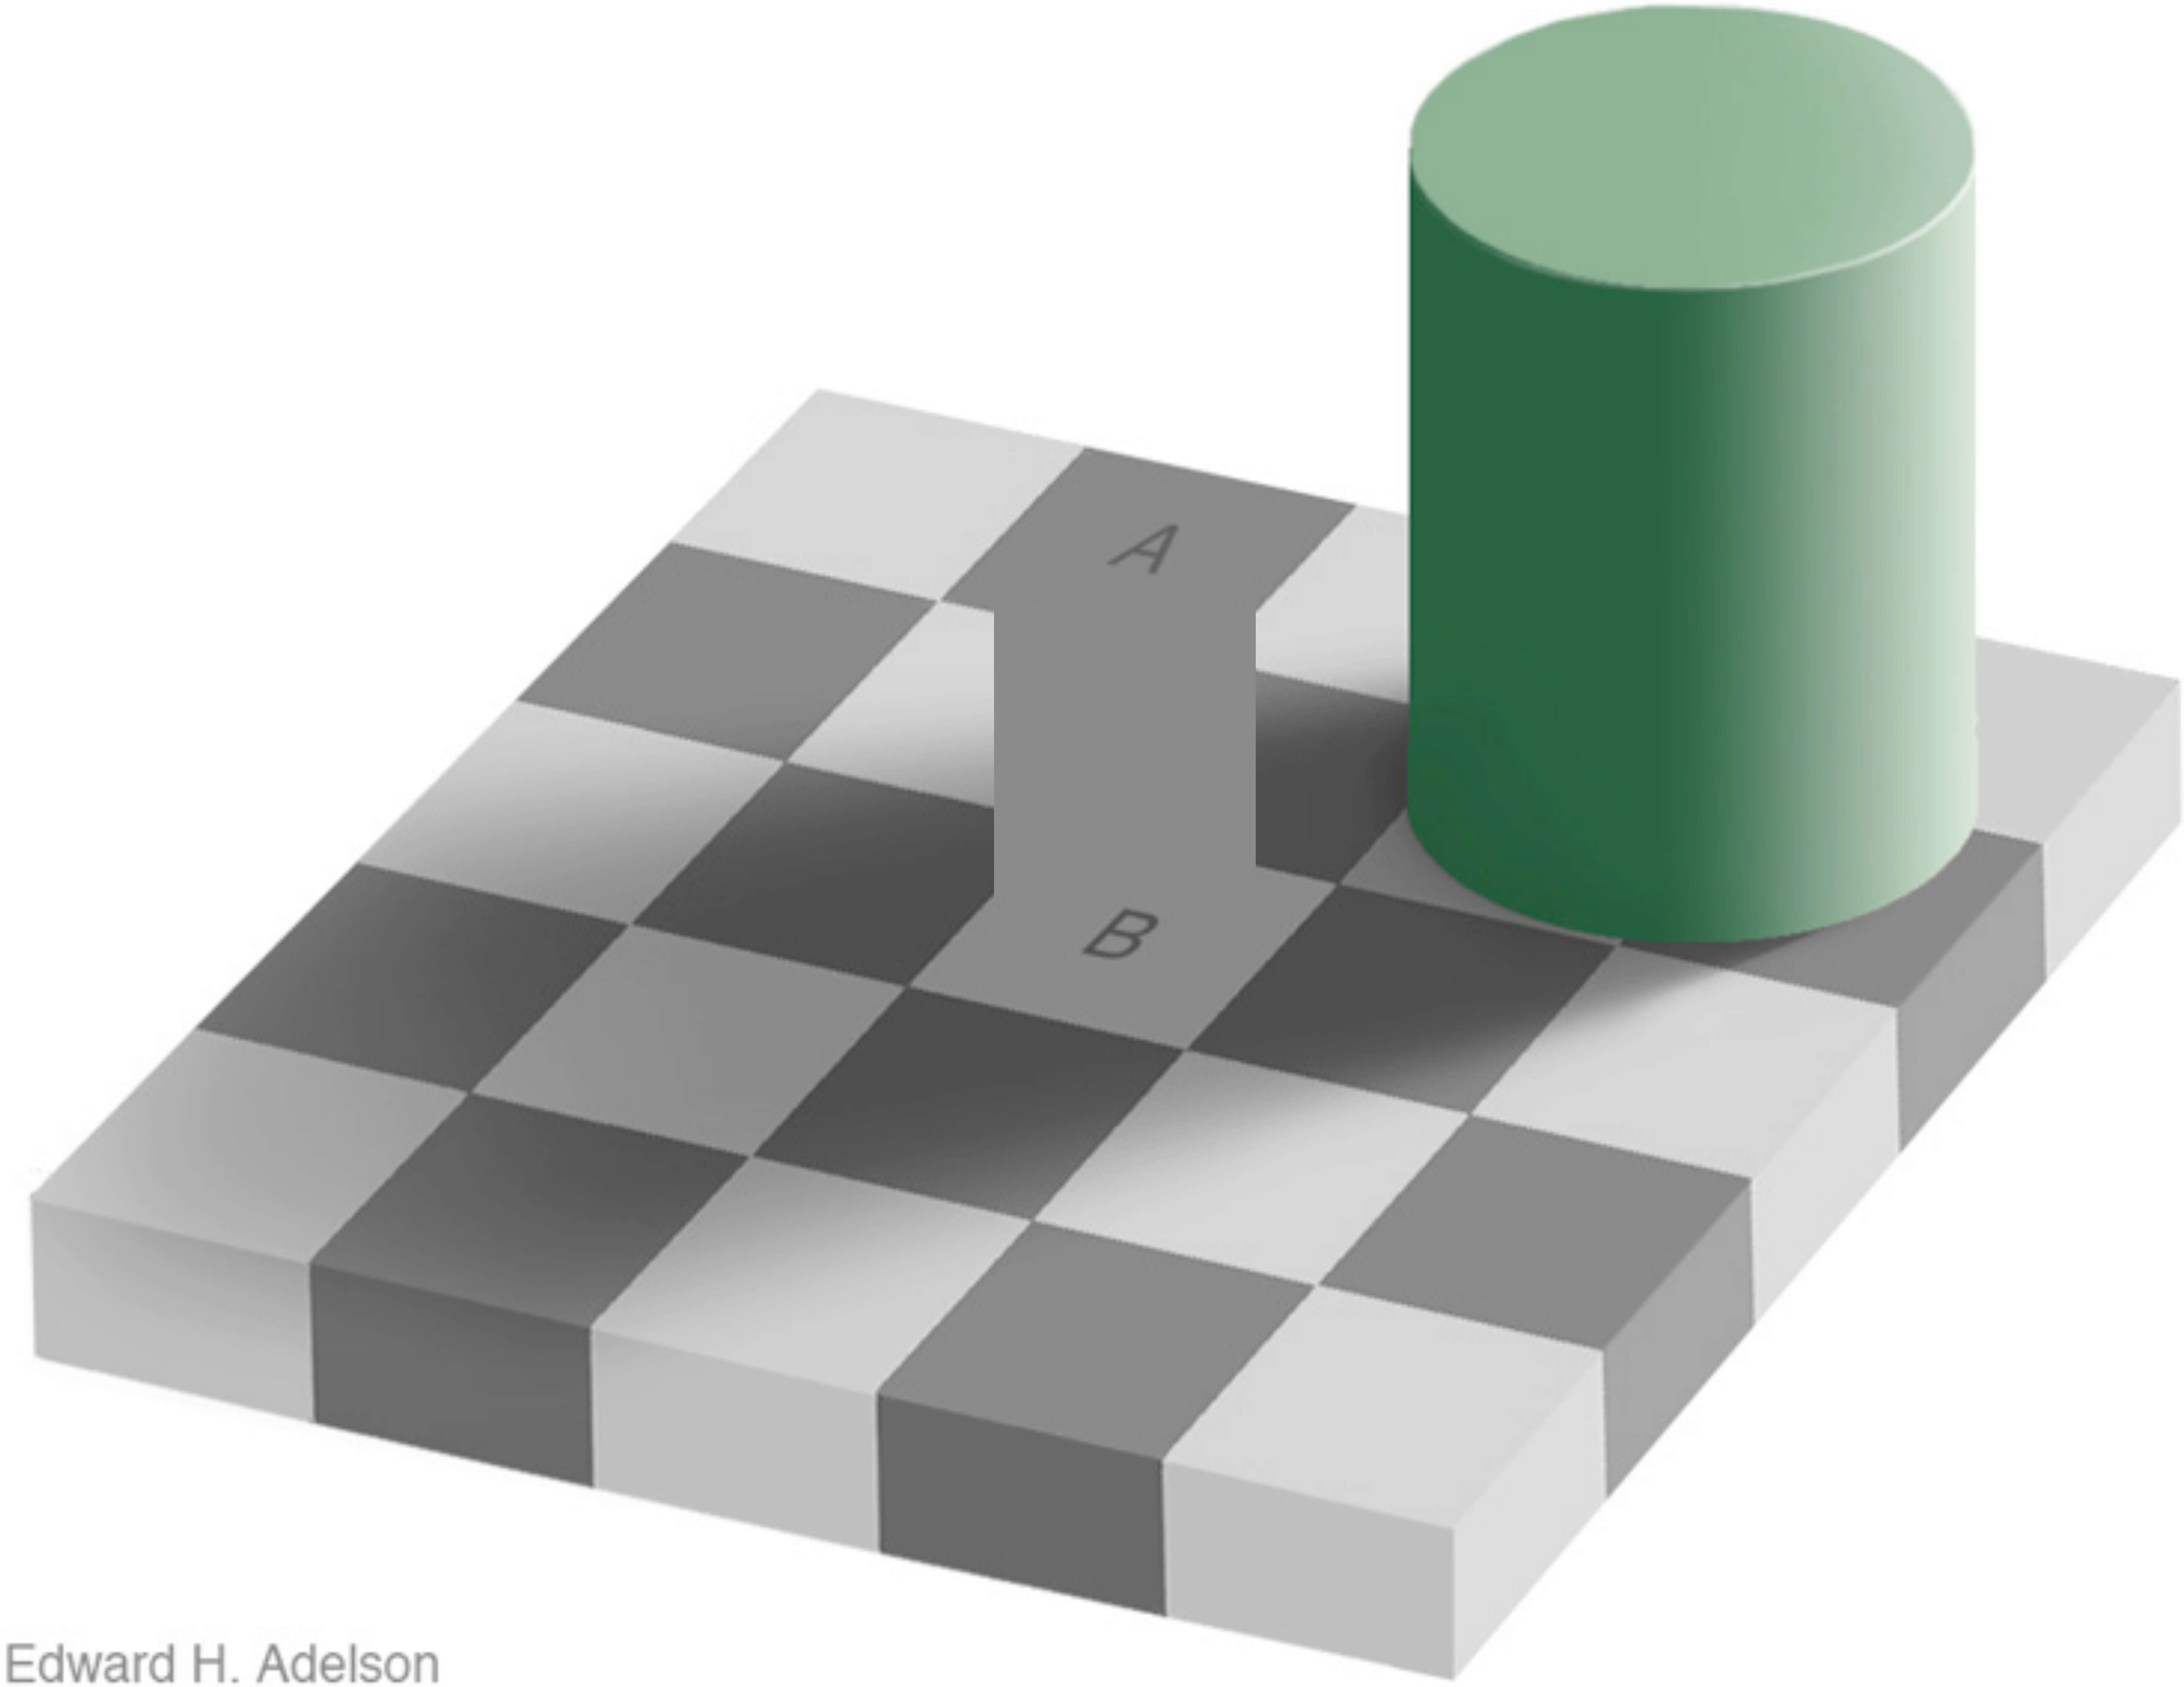

Edward H. Adelson

straight print

toned print

Ansel Adams, Clearing Winter Storm, 1942

La Grande Jatte, Georges Seurat, 1884

After a slide by Marc Levoy

Oppenheim 1968, Chiu et al. 1993

Reduce contrast of low-frequencies, preserve high frequencies

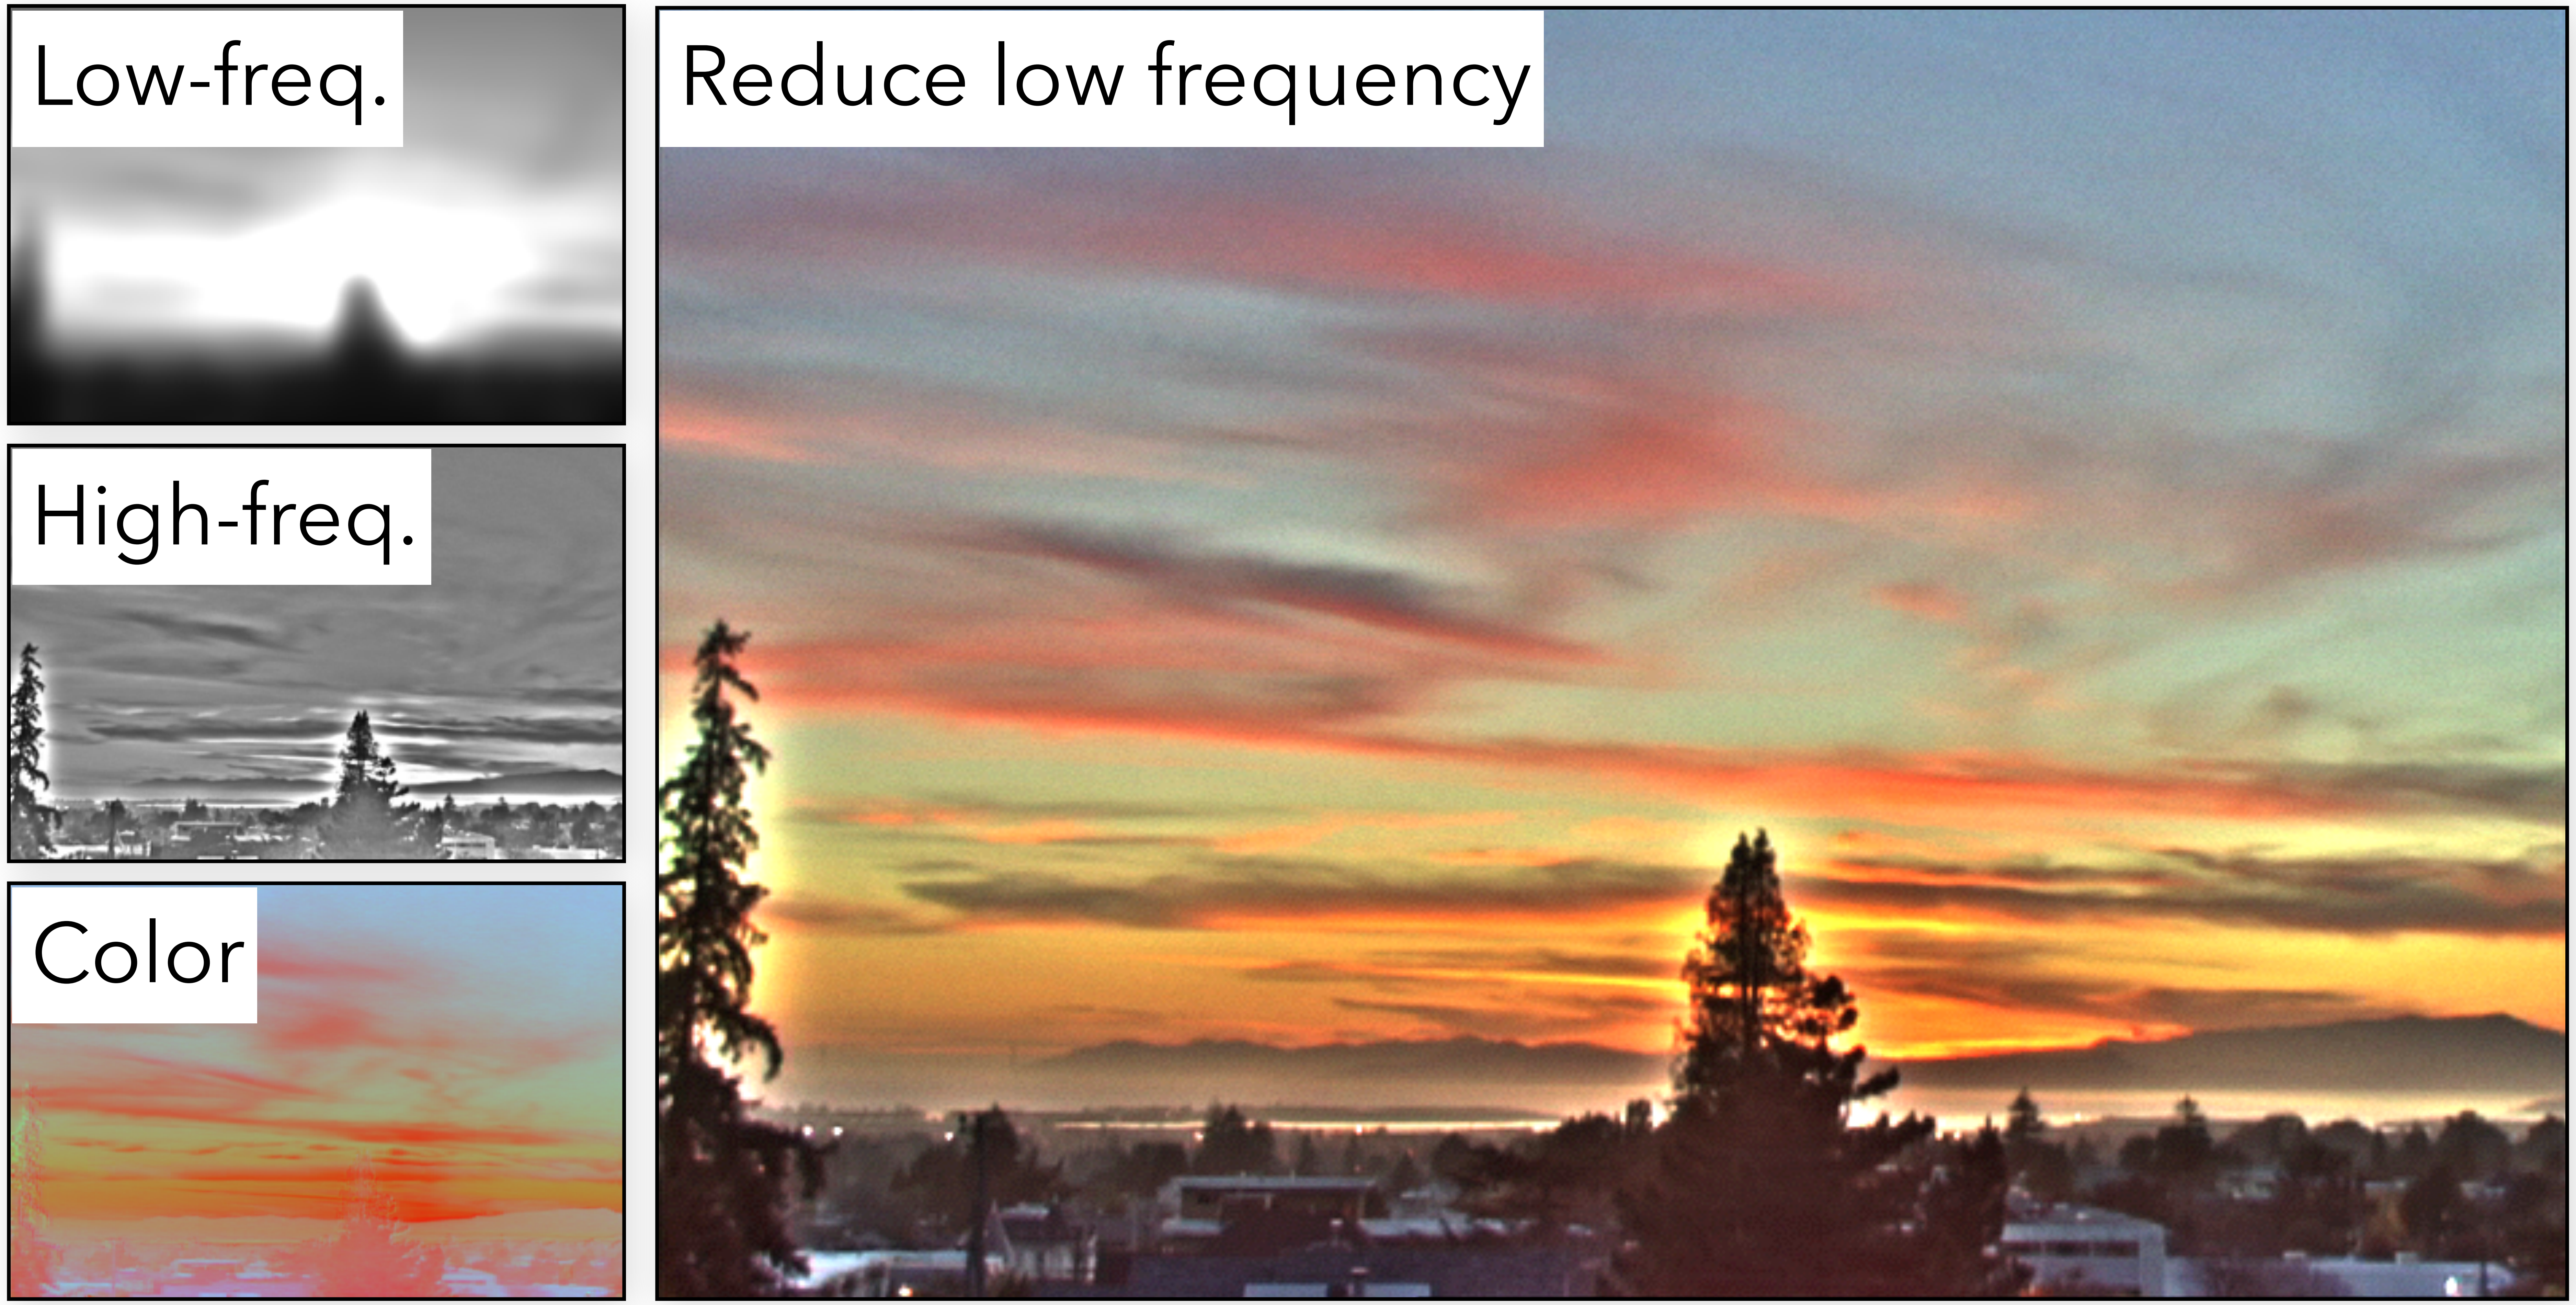

After a slide by Frédo Durand

The halo nightmare

For strong edges; because they contain high frequency

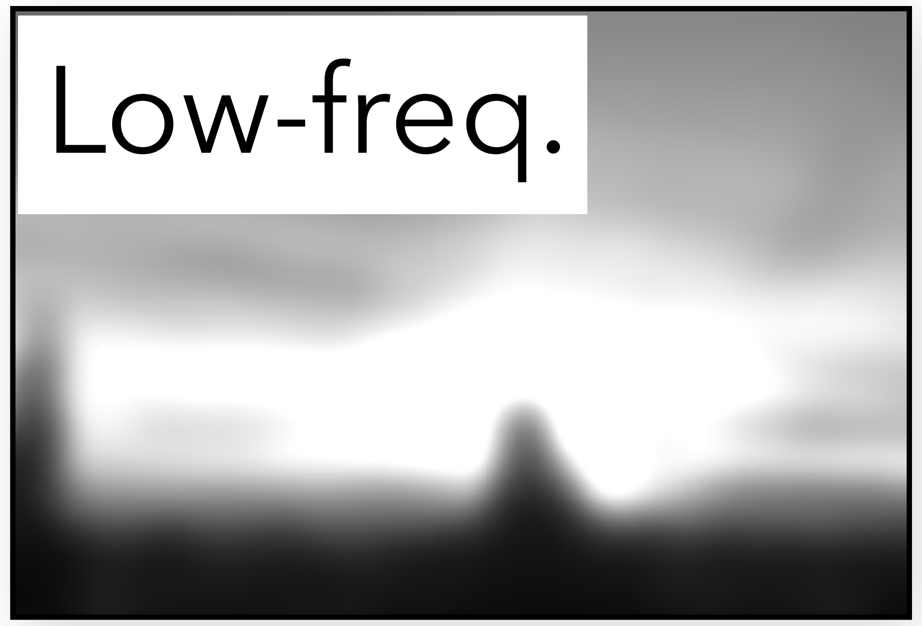

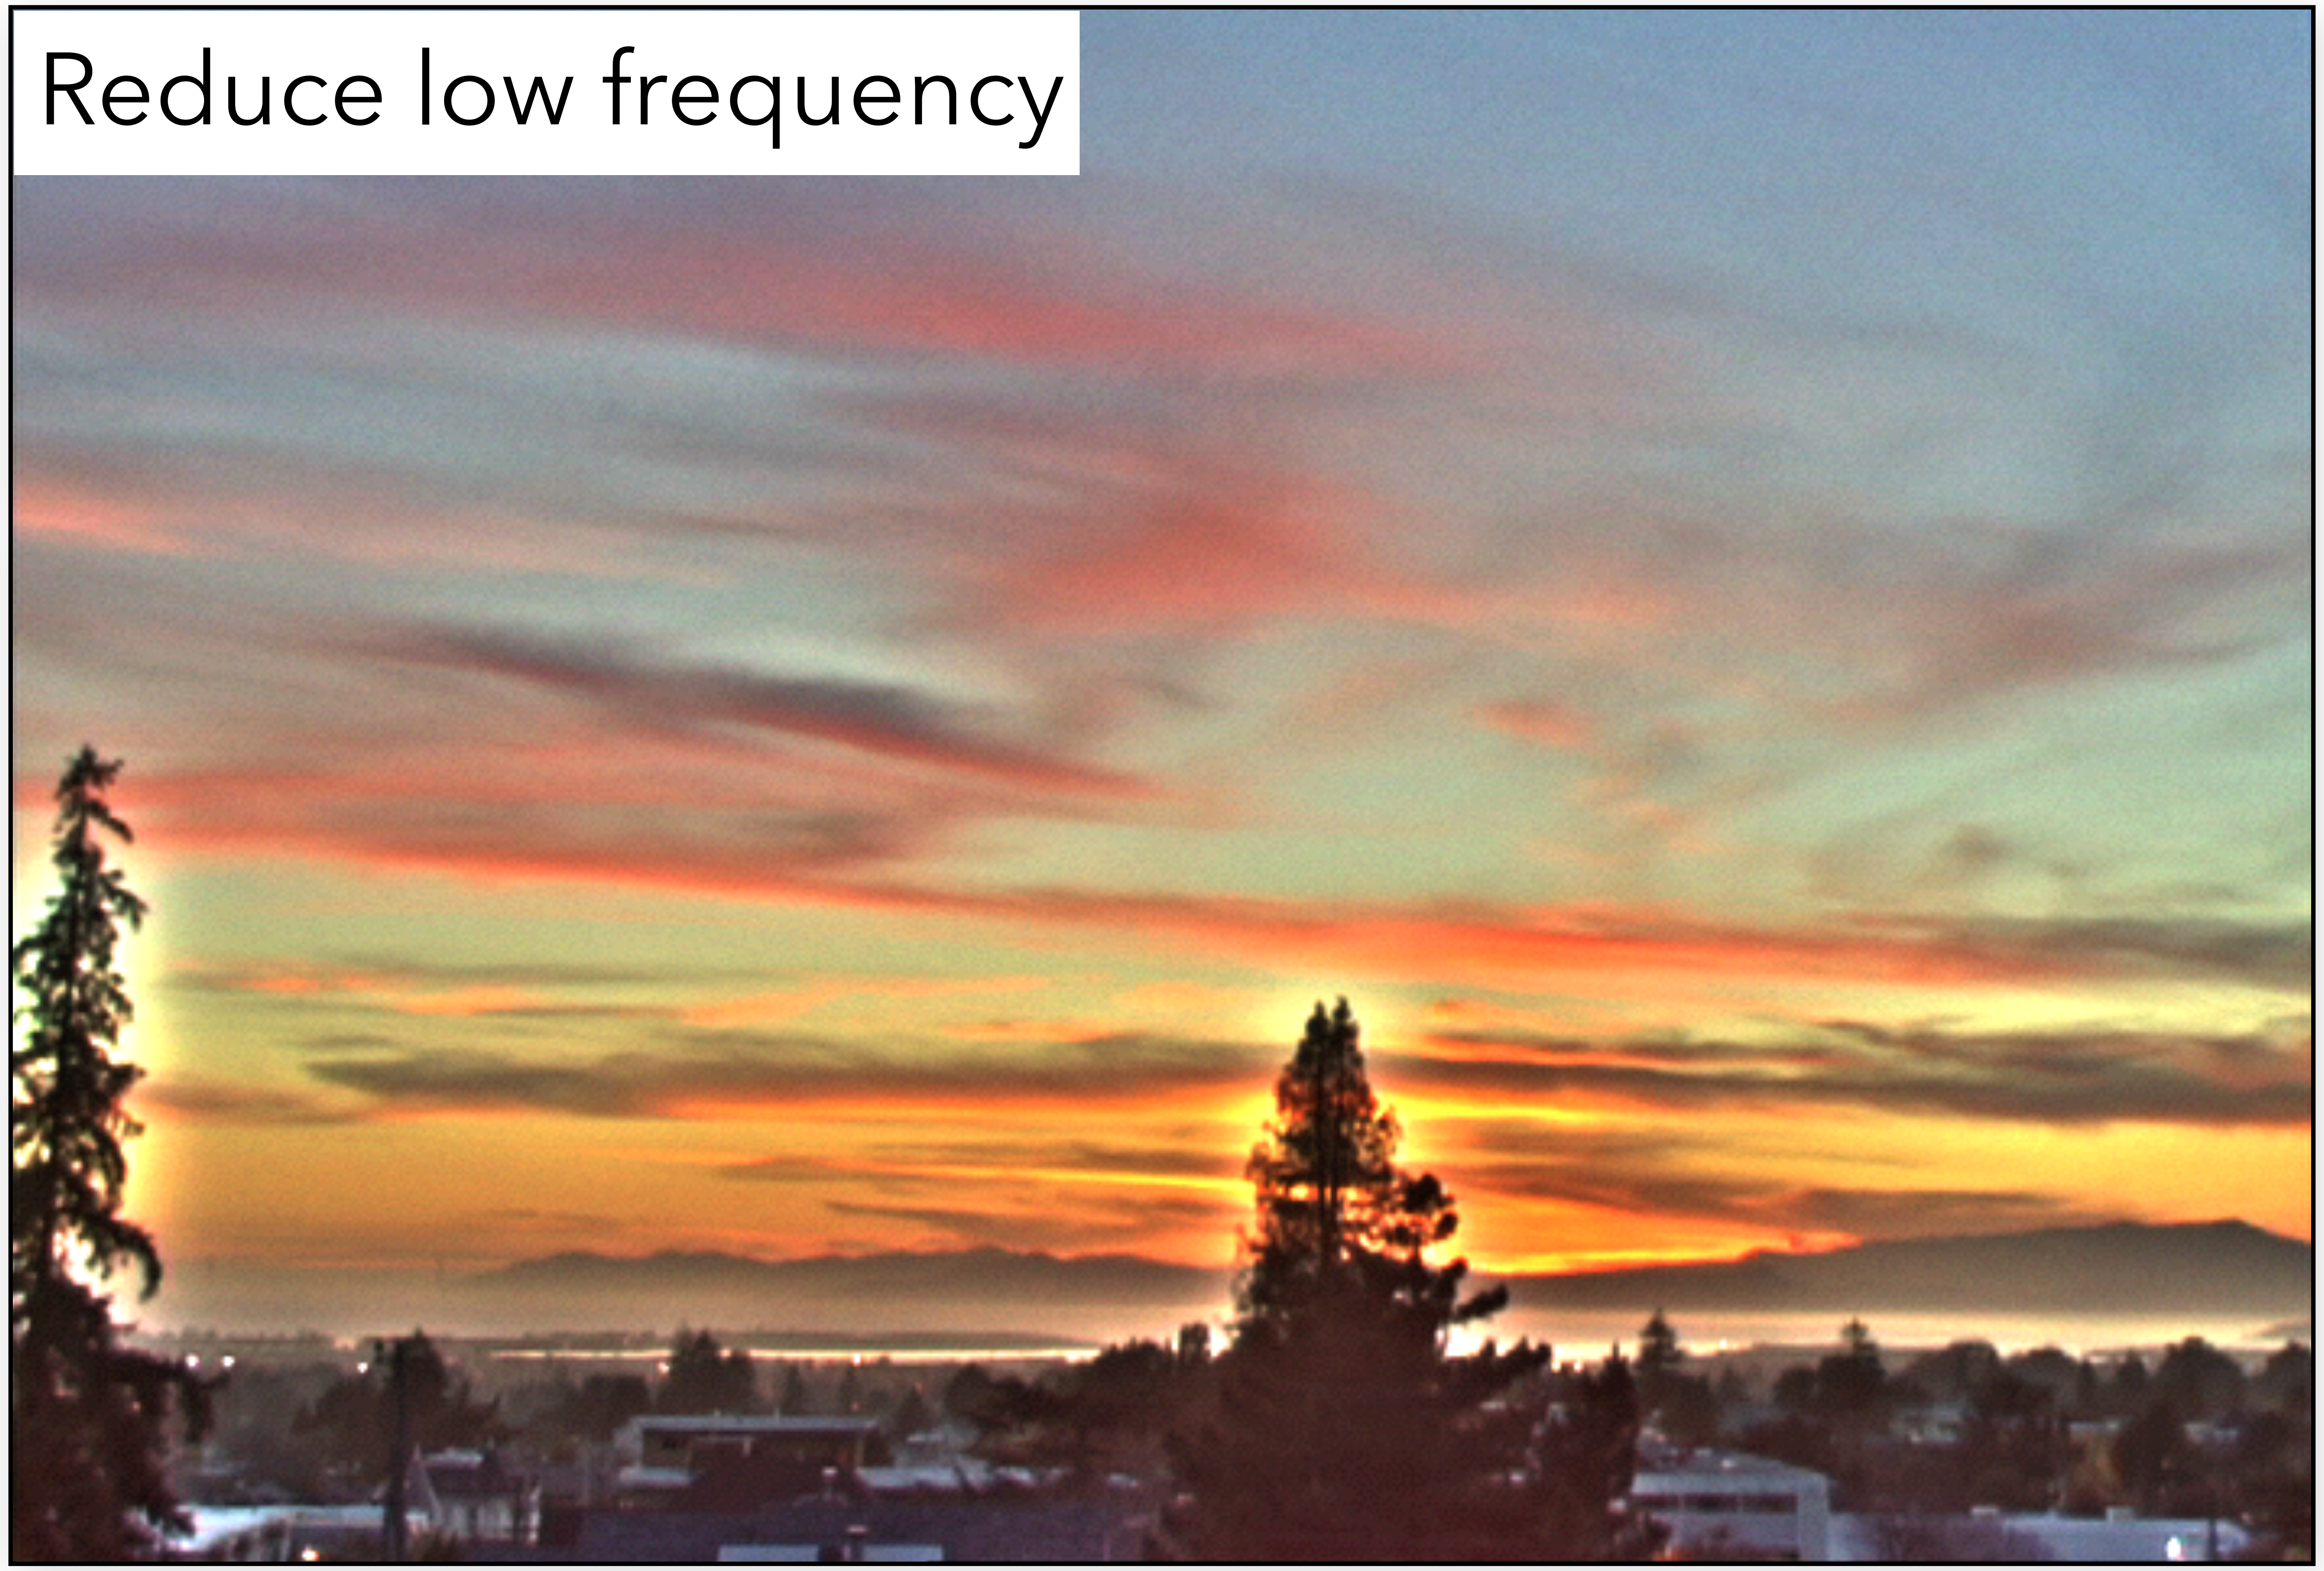

After a slide by Frédo Durand

Gaussian vs. Bilateral filter

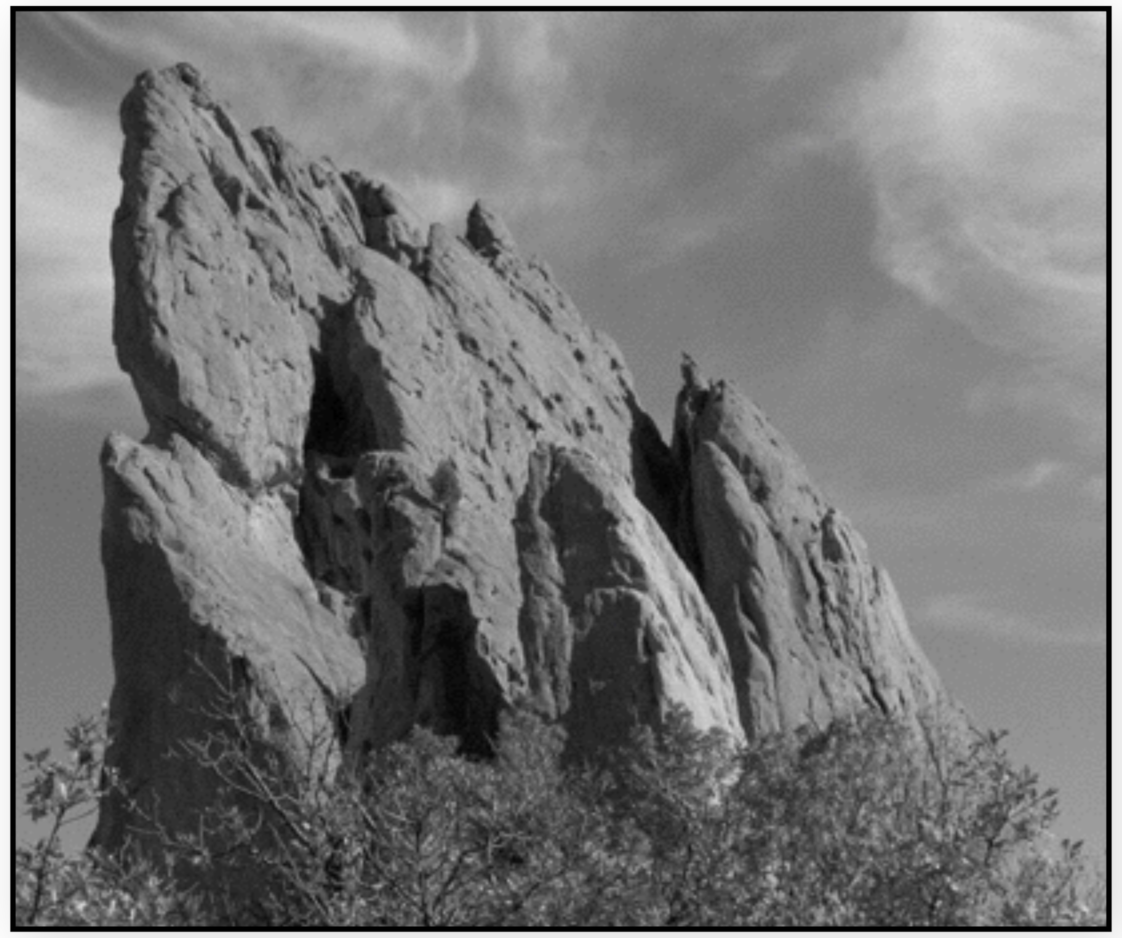

Input

$\sigma_s = 2$

$\sigma_s = 6$

$\sigma_s = 18$

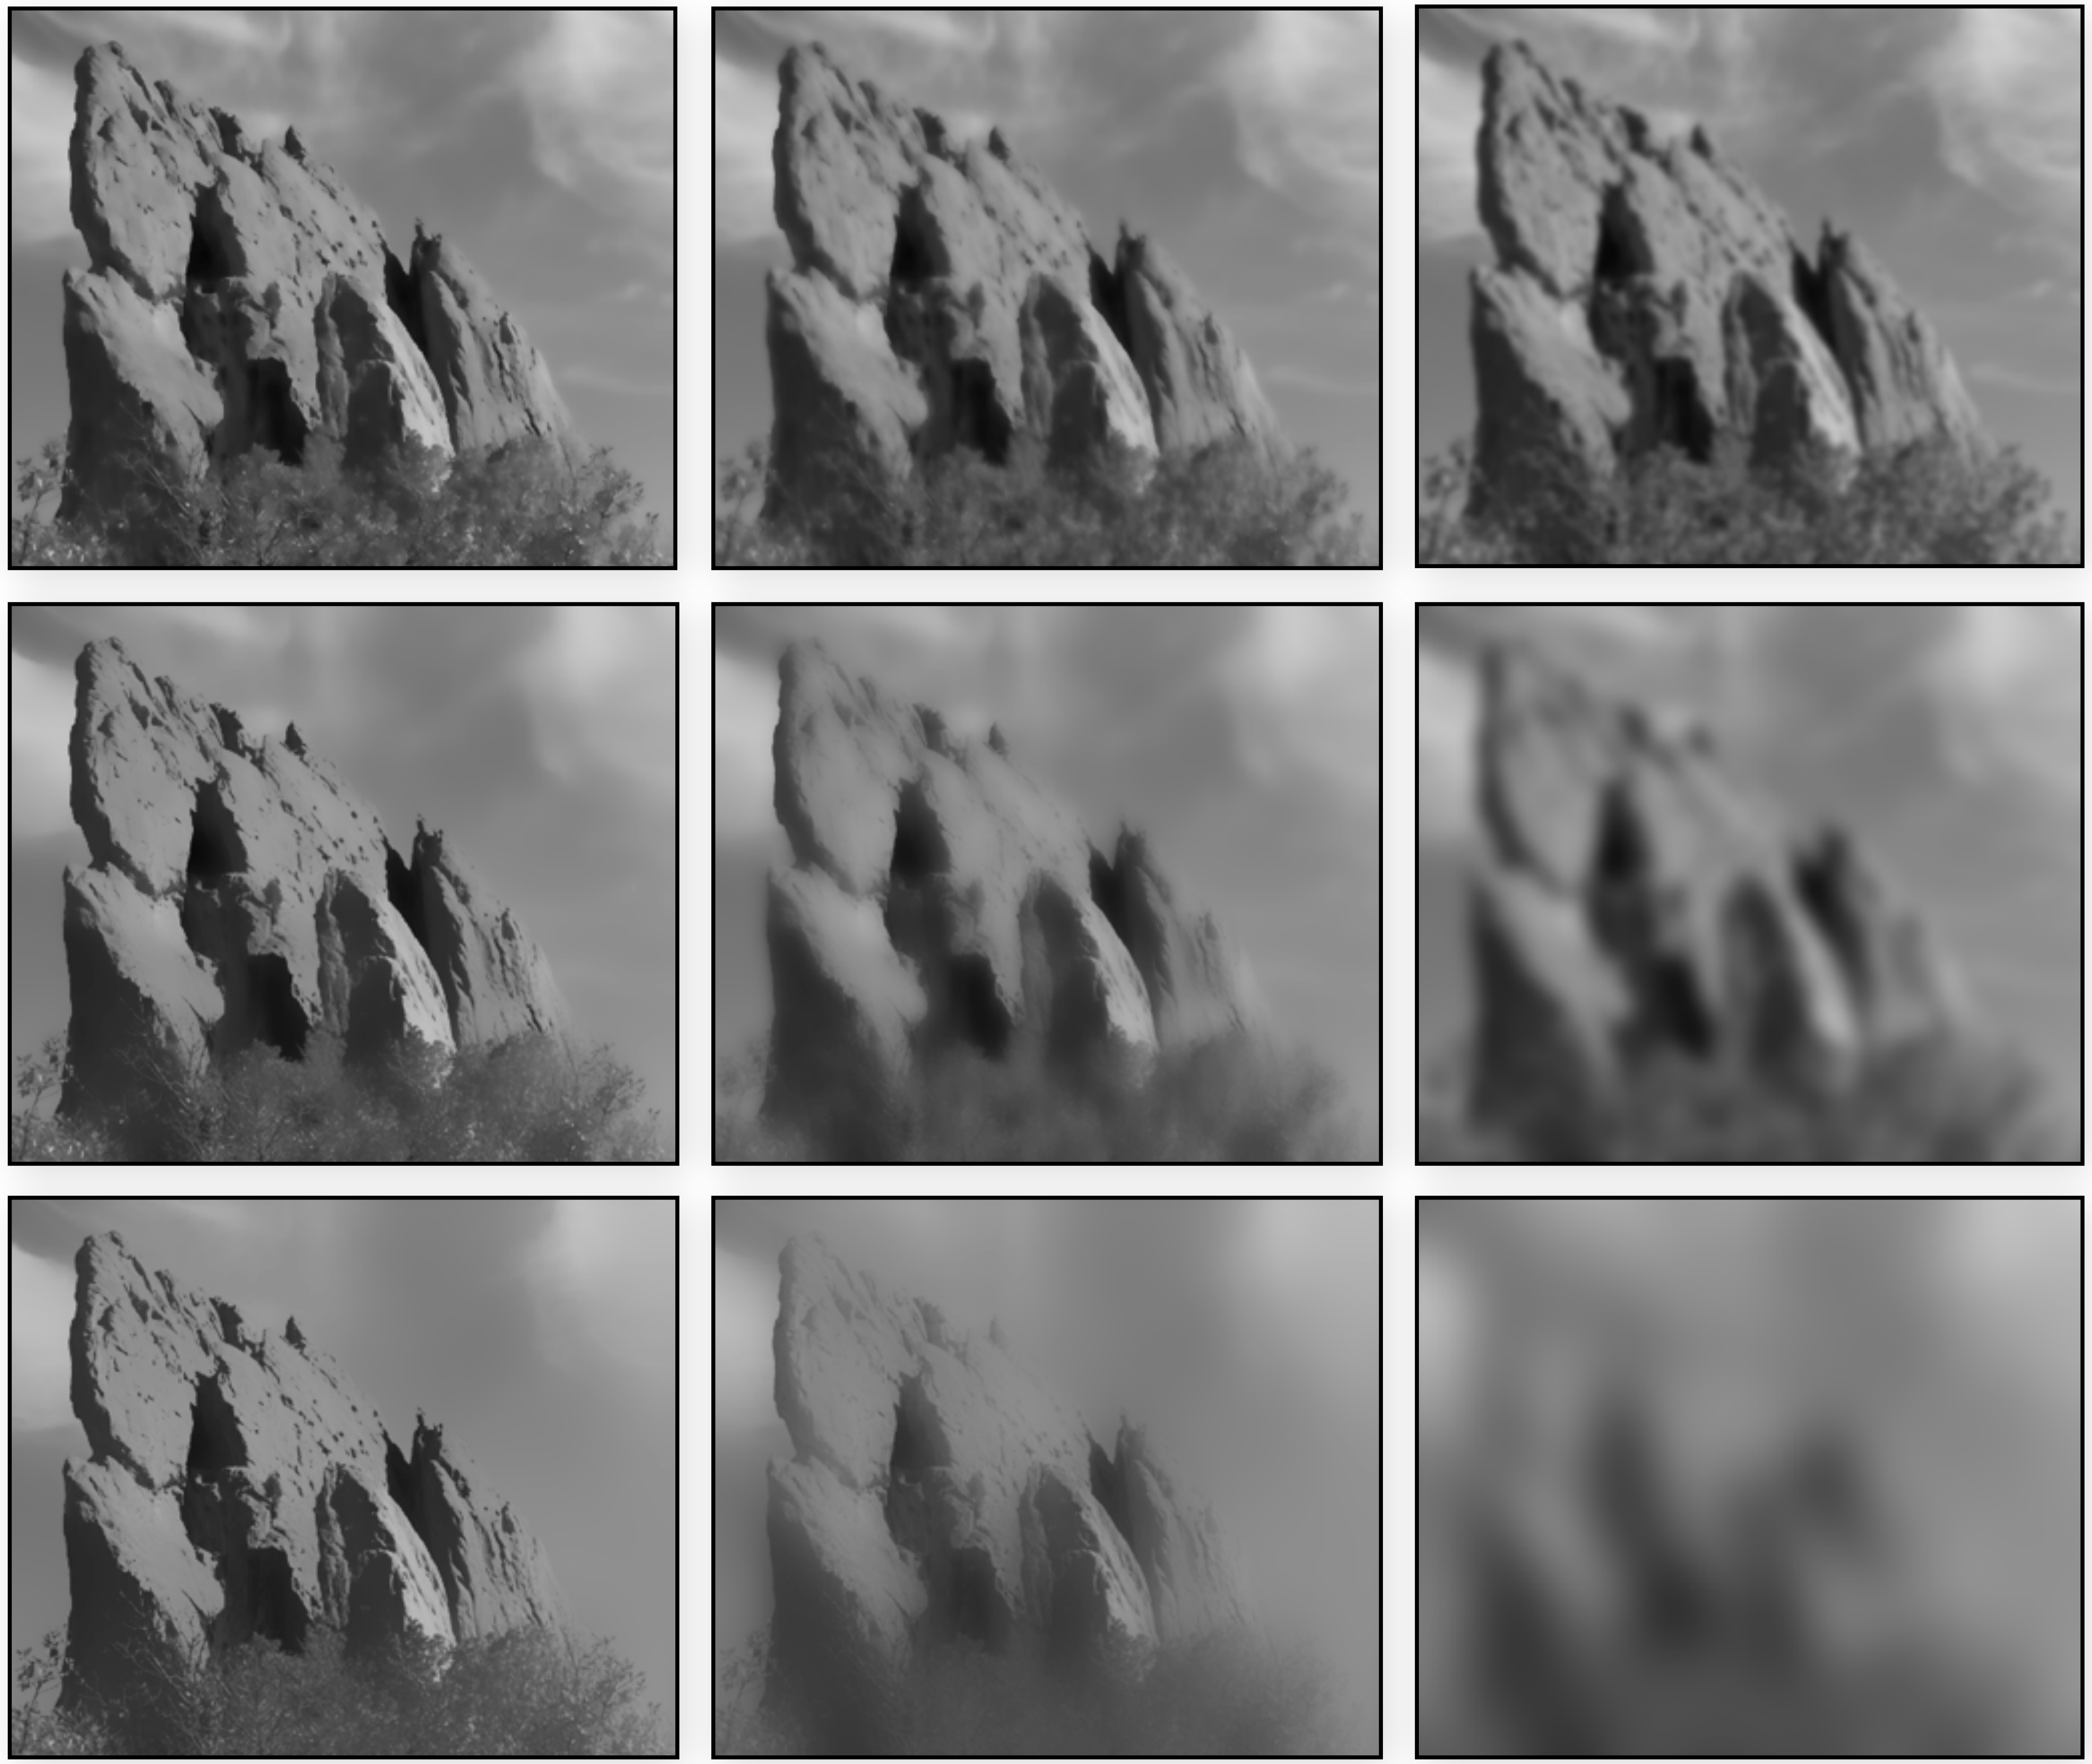

$\sigma_r = 0.1$ $\sigma_r = 0.25$ $\sigma_r = \infty$ (Gaussian blur)

http://people.csail.mit.edu/sparis/bf_course/slides/03_definition_bf.pdf

Durand and Dorsey 2002

Don't blur across edges, decompose using bilateral filter

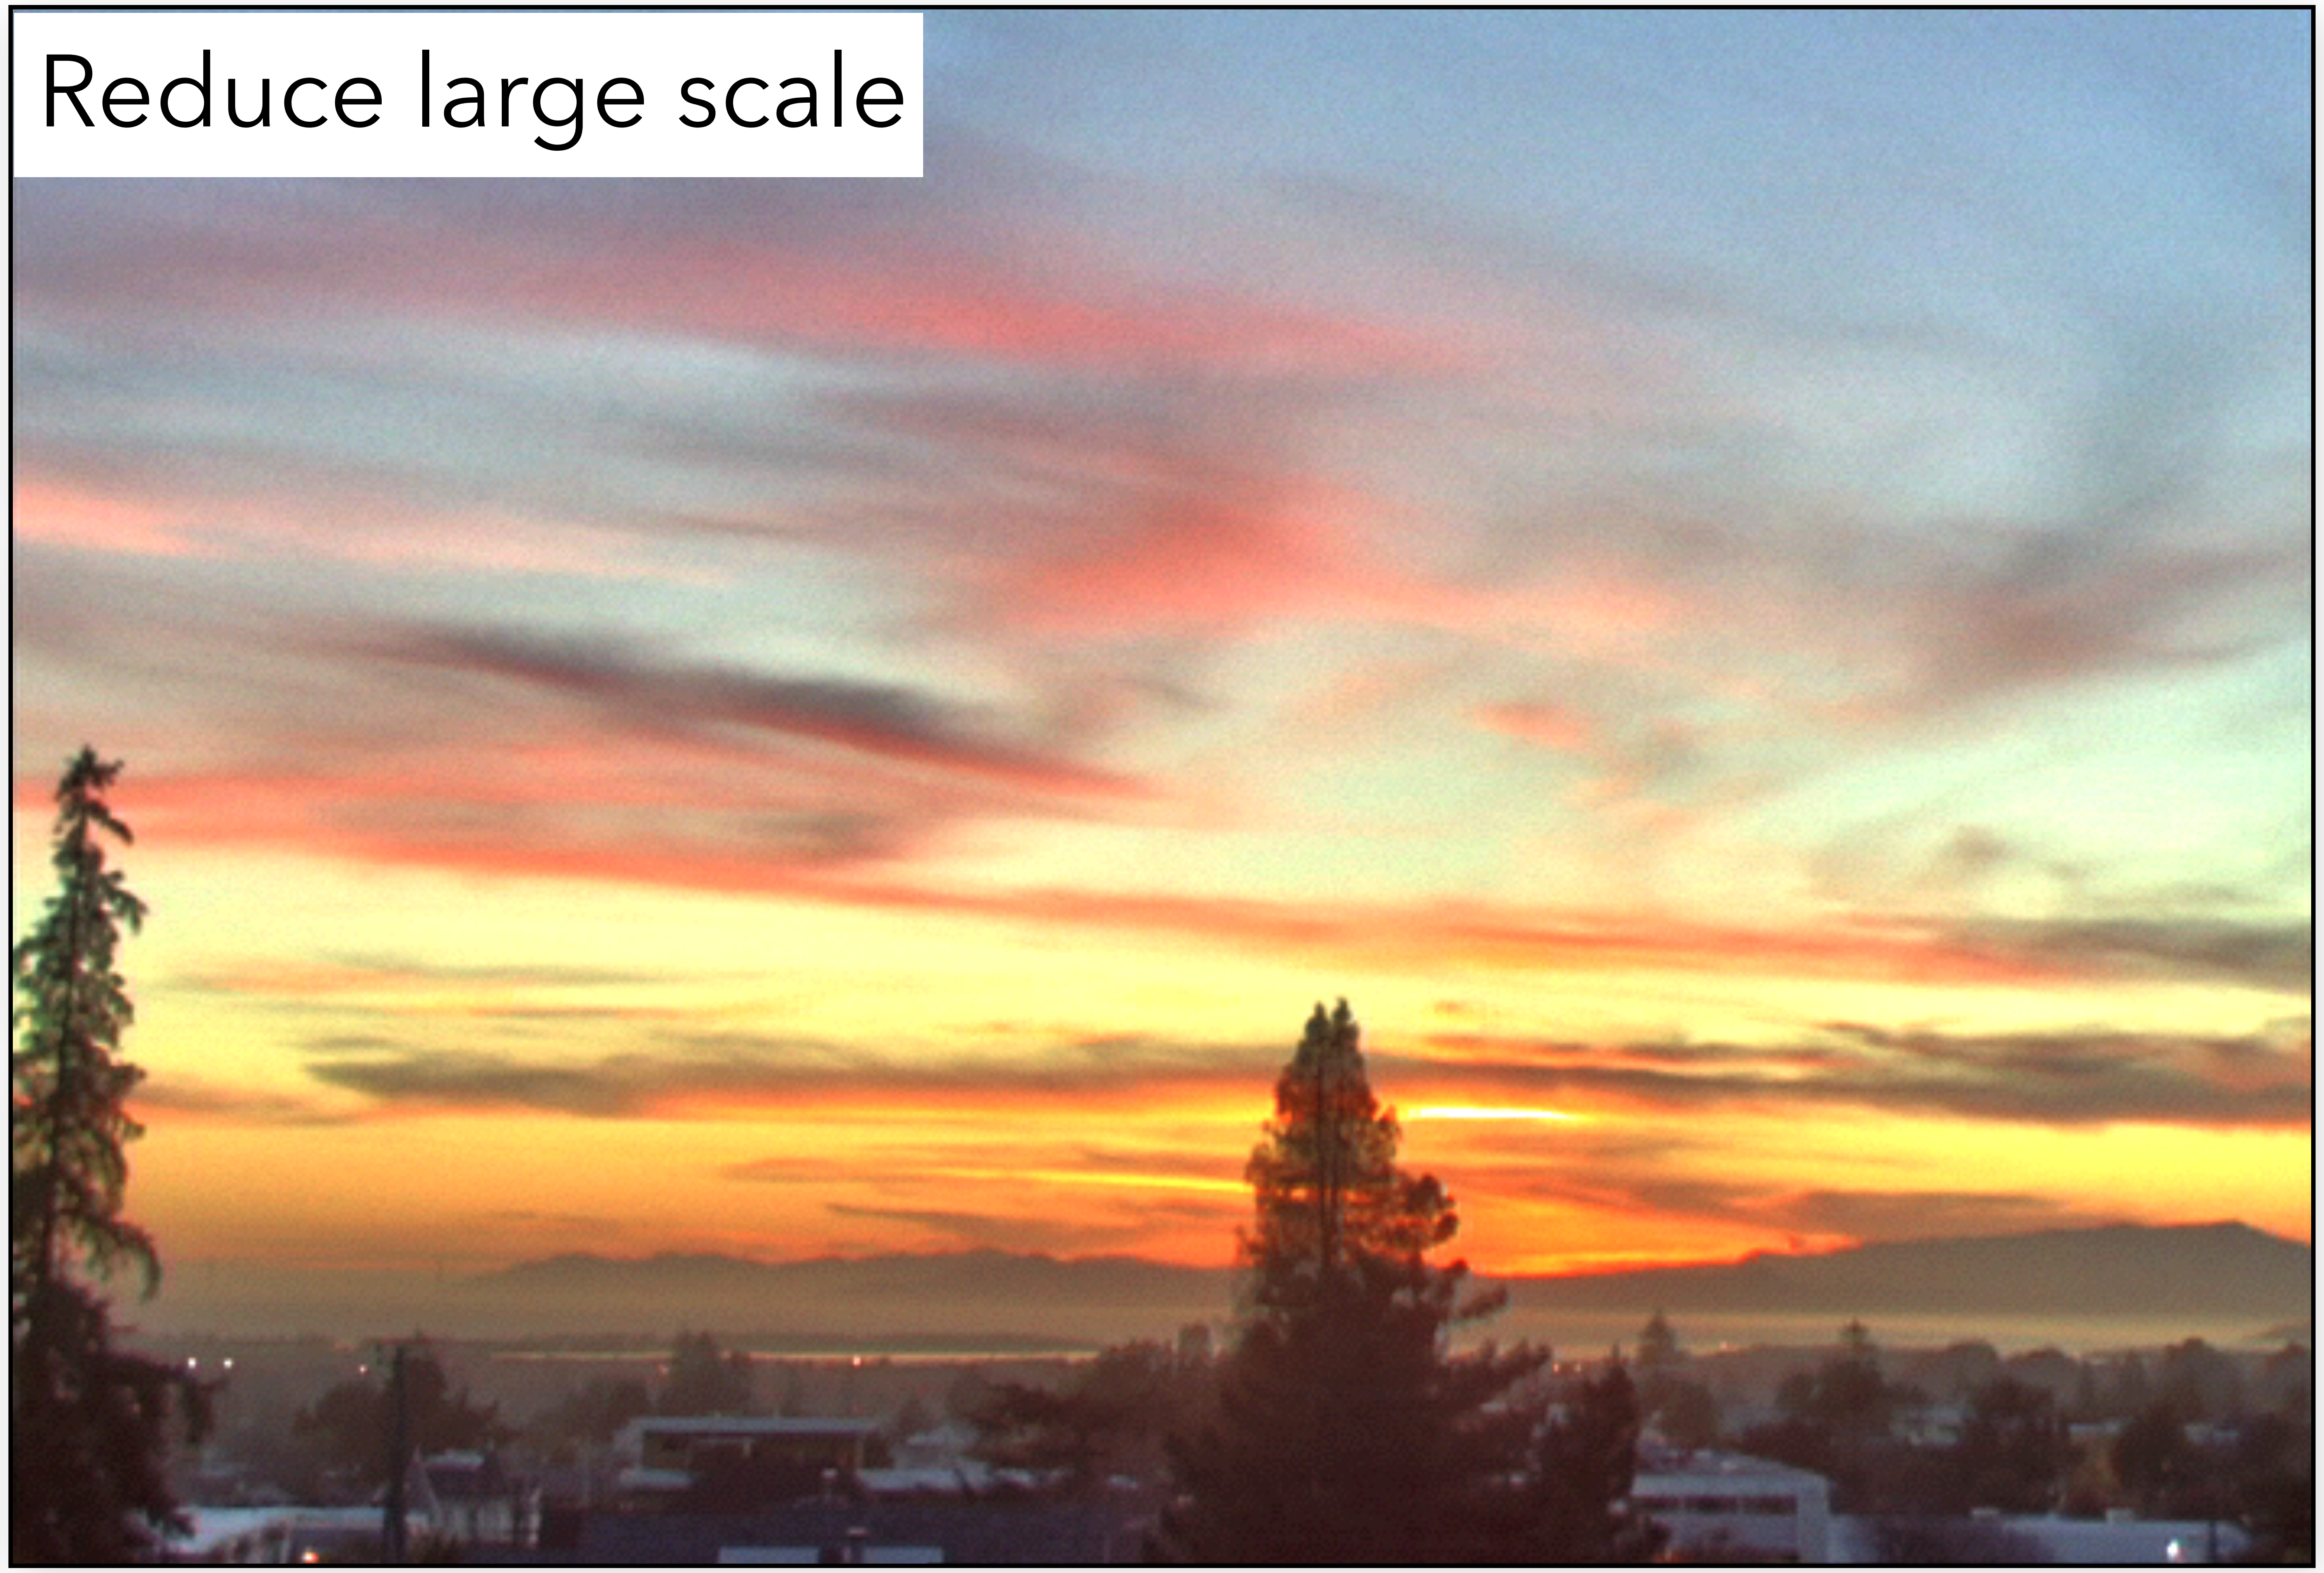

After a slide by Frédo Durand

Today

Recap of bilateral tone mapping

Variance-optimized weights for HDR merging

HDR merging/tone mapping in practice

Contrast reduction

Input HDR image

Contrast
too high!

After a slide by Frédo Durand

Contrast reduction

intensity
$= 0.4R + 0.7G + 0.01B$

$R' = R/\text{intensity}$
$G' = G/\text{intensity}$
$B' = B/\text{intensity}$

important to use ratios (makes it luminance invariant)

After a slide by Frédo Durand

Contrast reduction

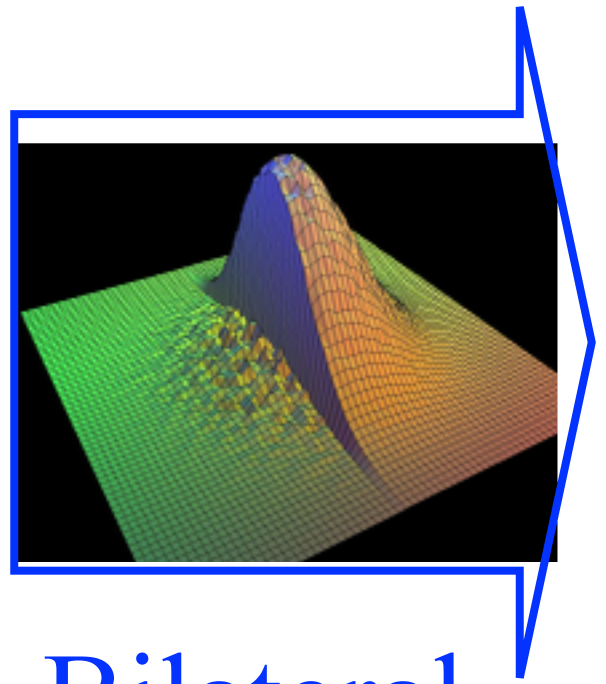

Bilateral
Filter

in log

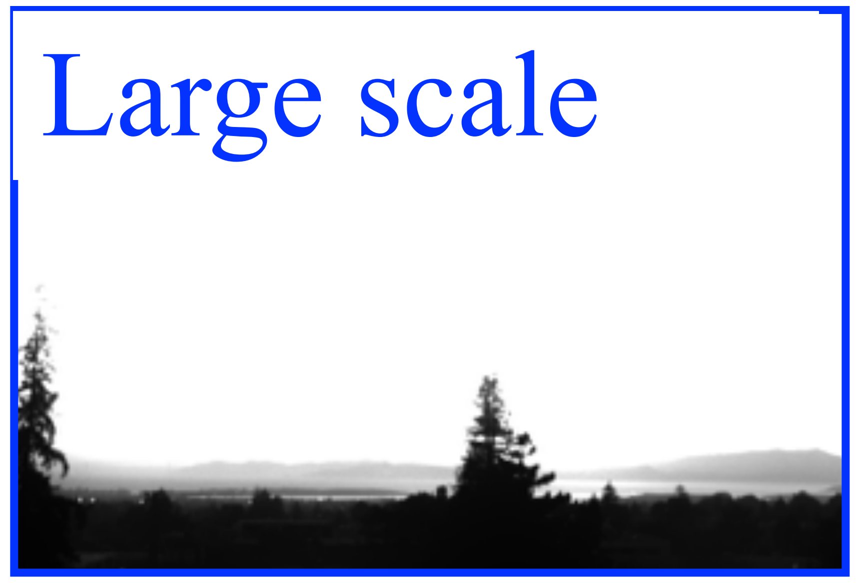

Spatial sigma: 2 to 5% image size
Range sigma: 0.4 (in log 10)

After a slide by Frédo Durand

Contrast reduction

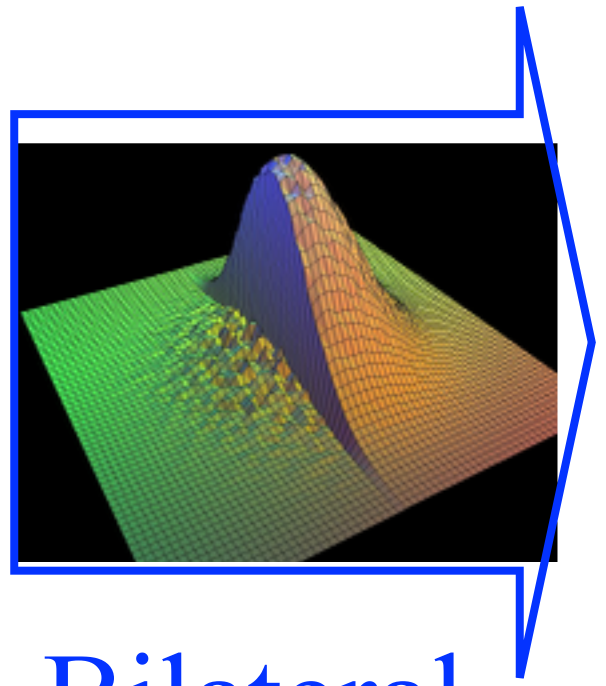

Bilateral Filter

in log

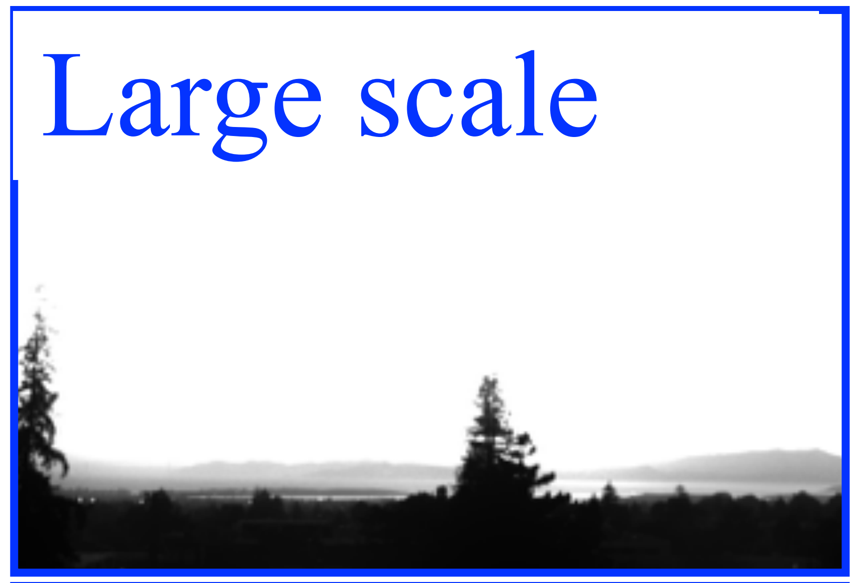

Detail = log intensity - large scale (residual)

After a slide by Frédo Durand

Contrast reduction

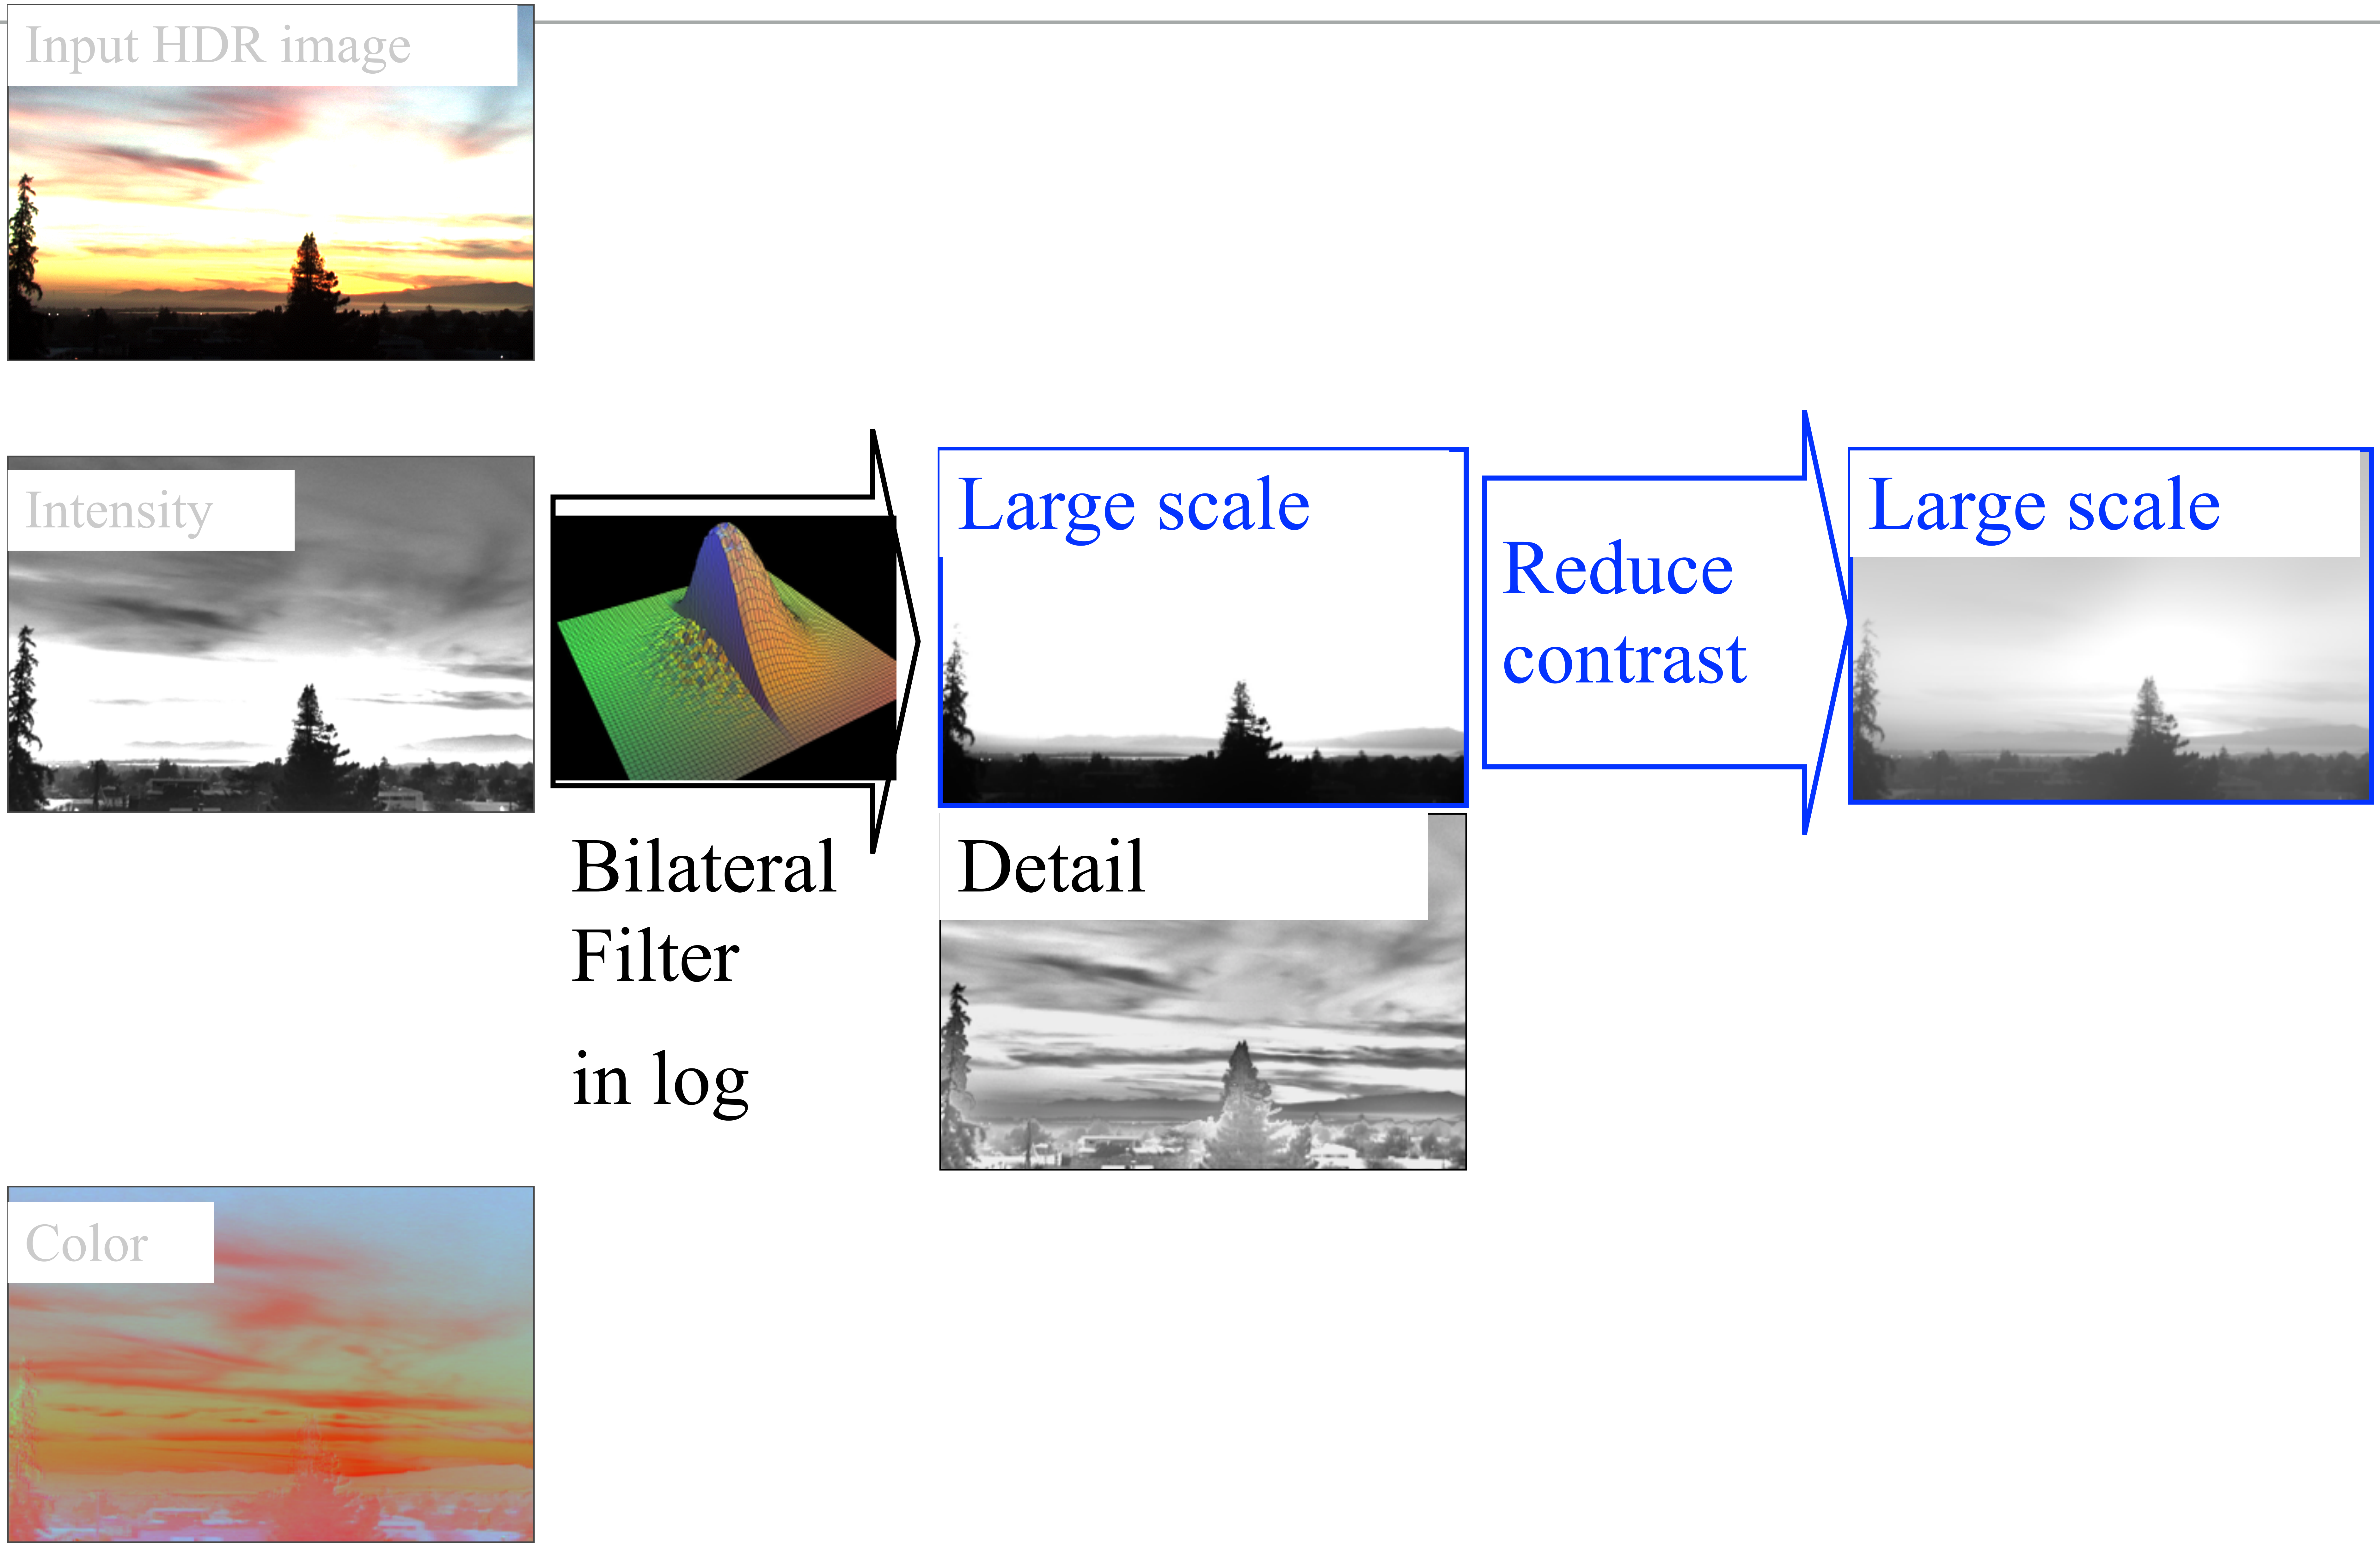

After a slide by Frédo Durand

Contrast reduction

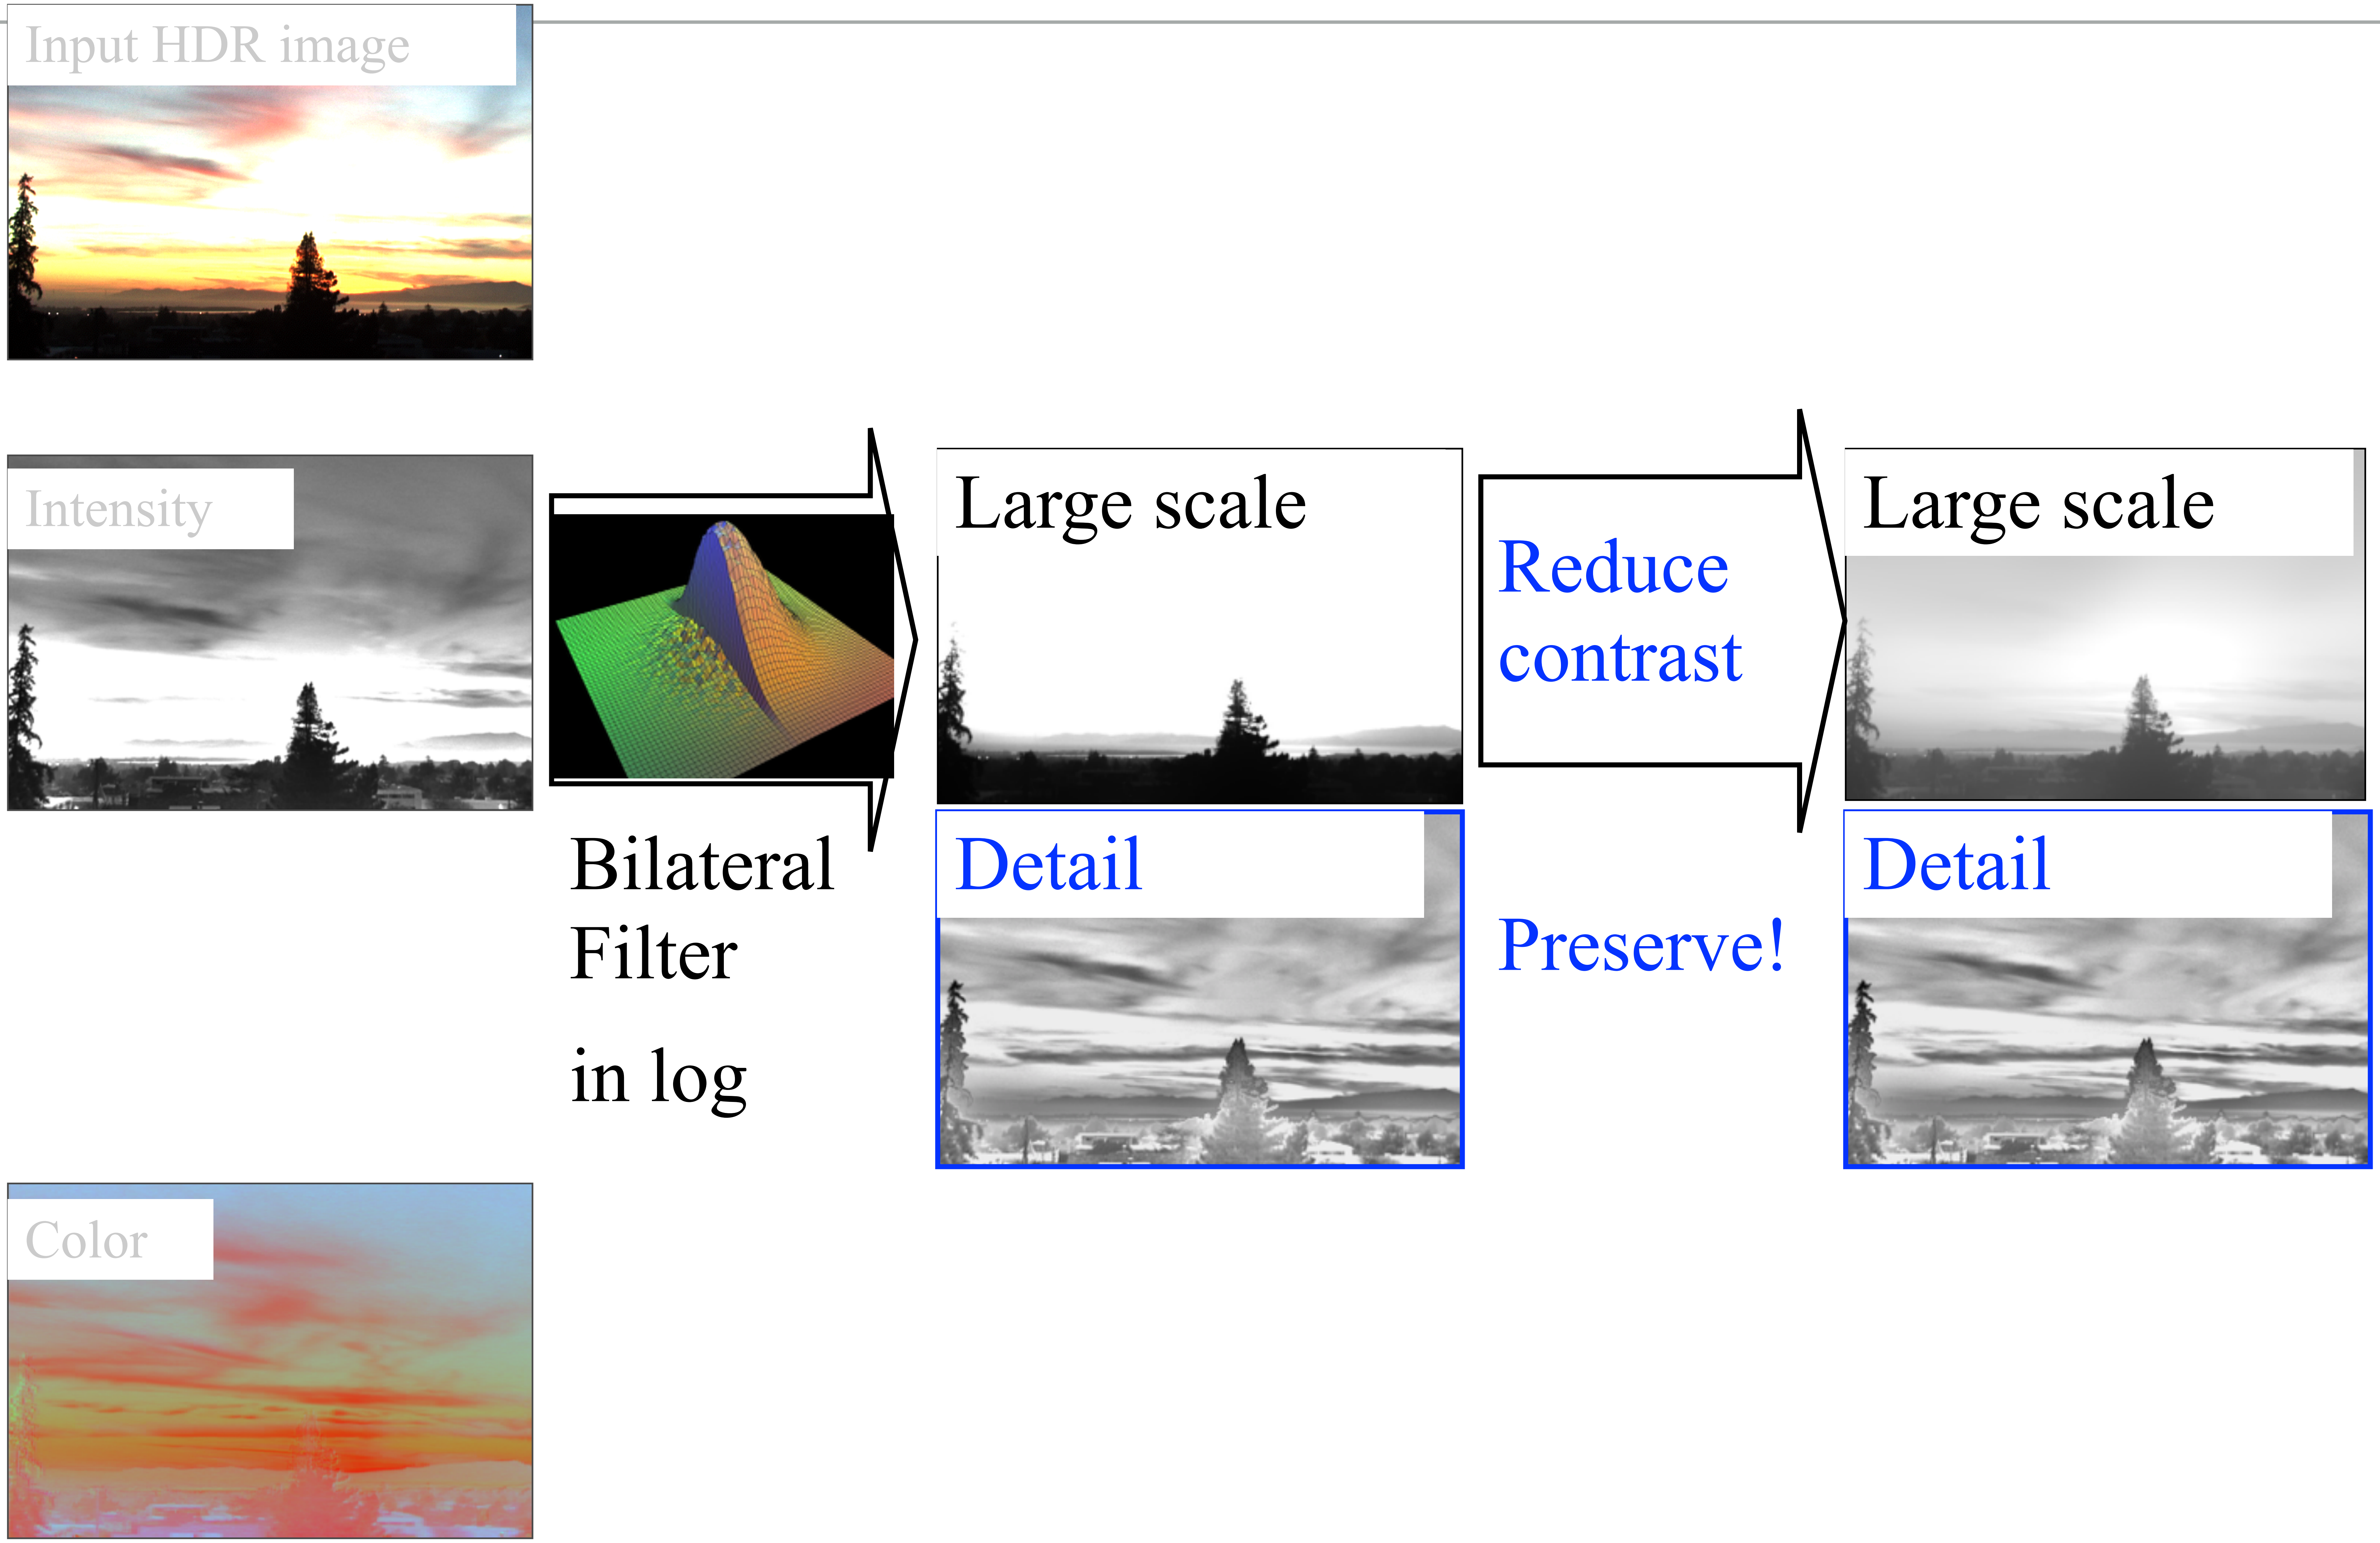

After a slide by Frédo Durand

Contrast reduction

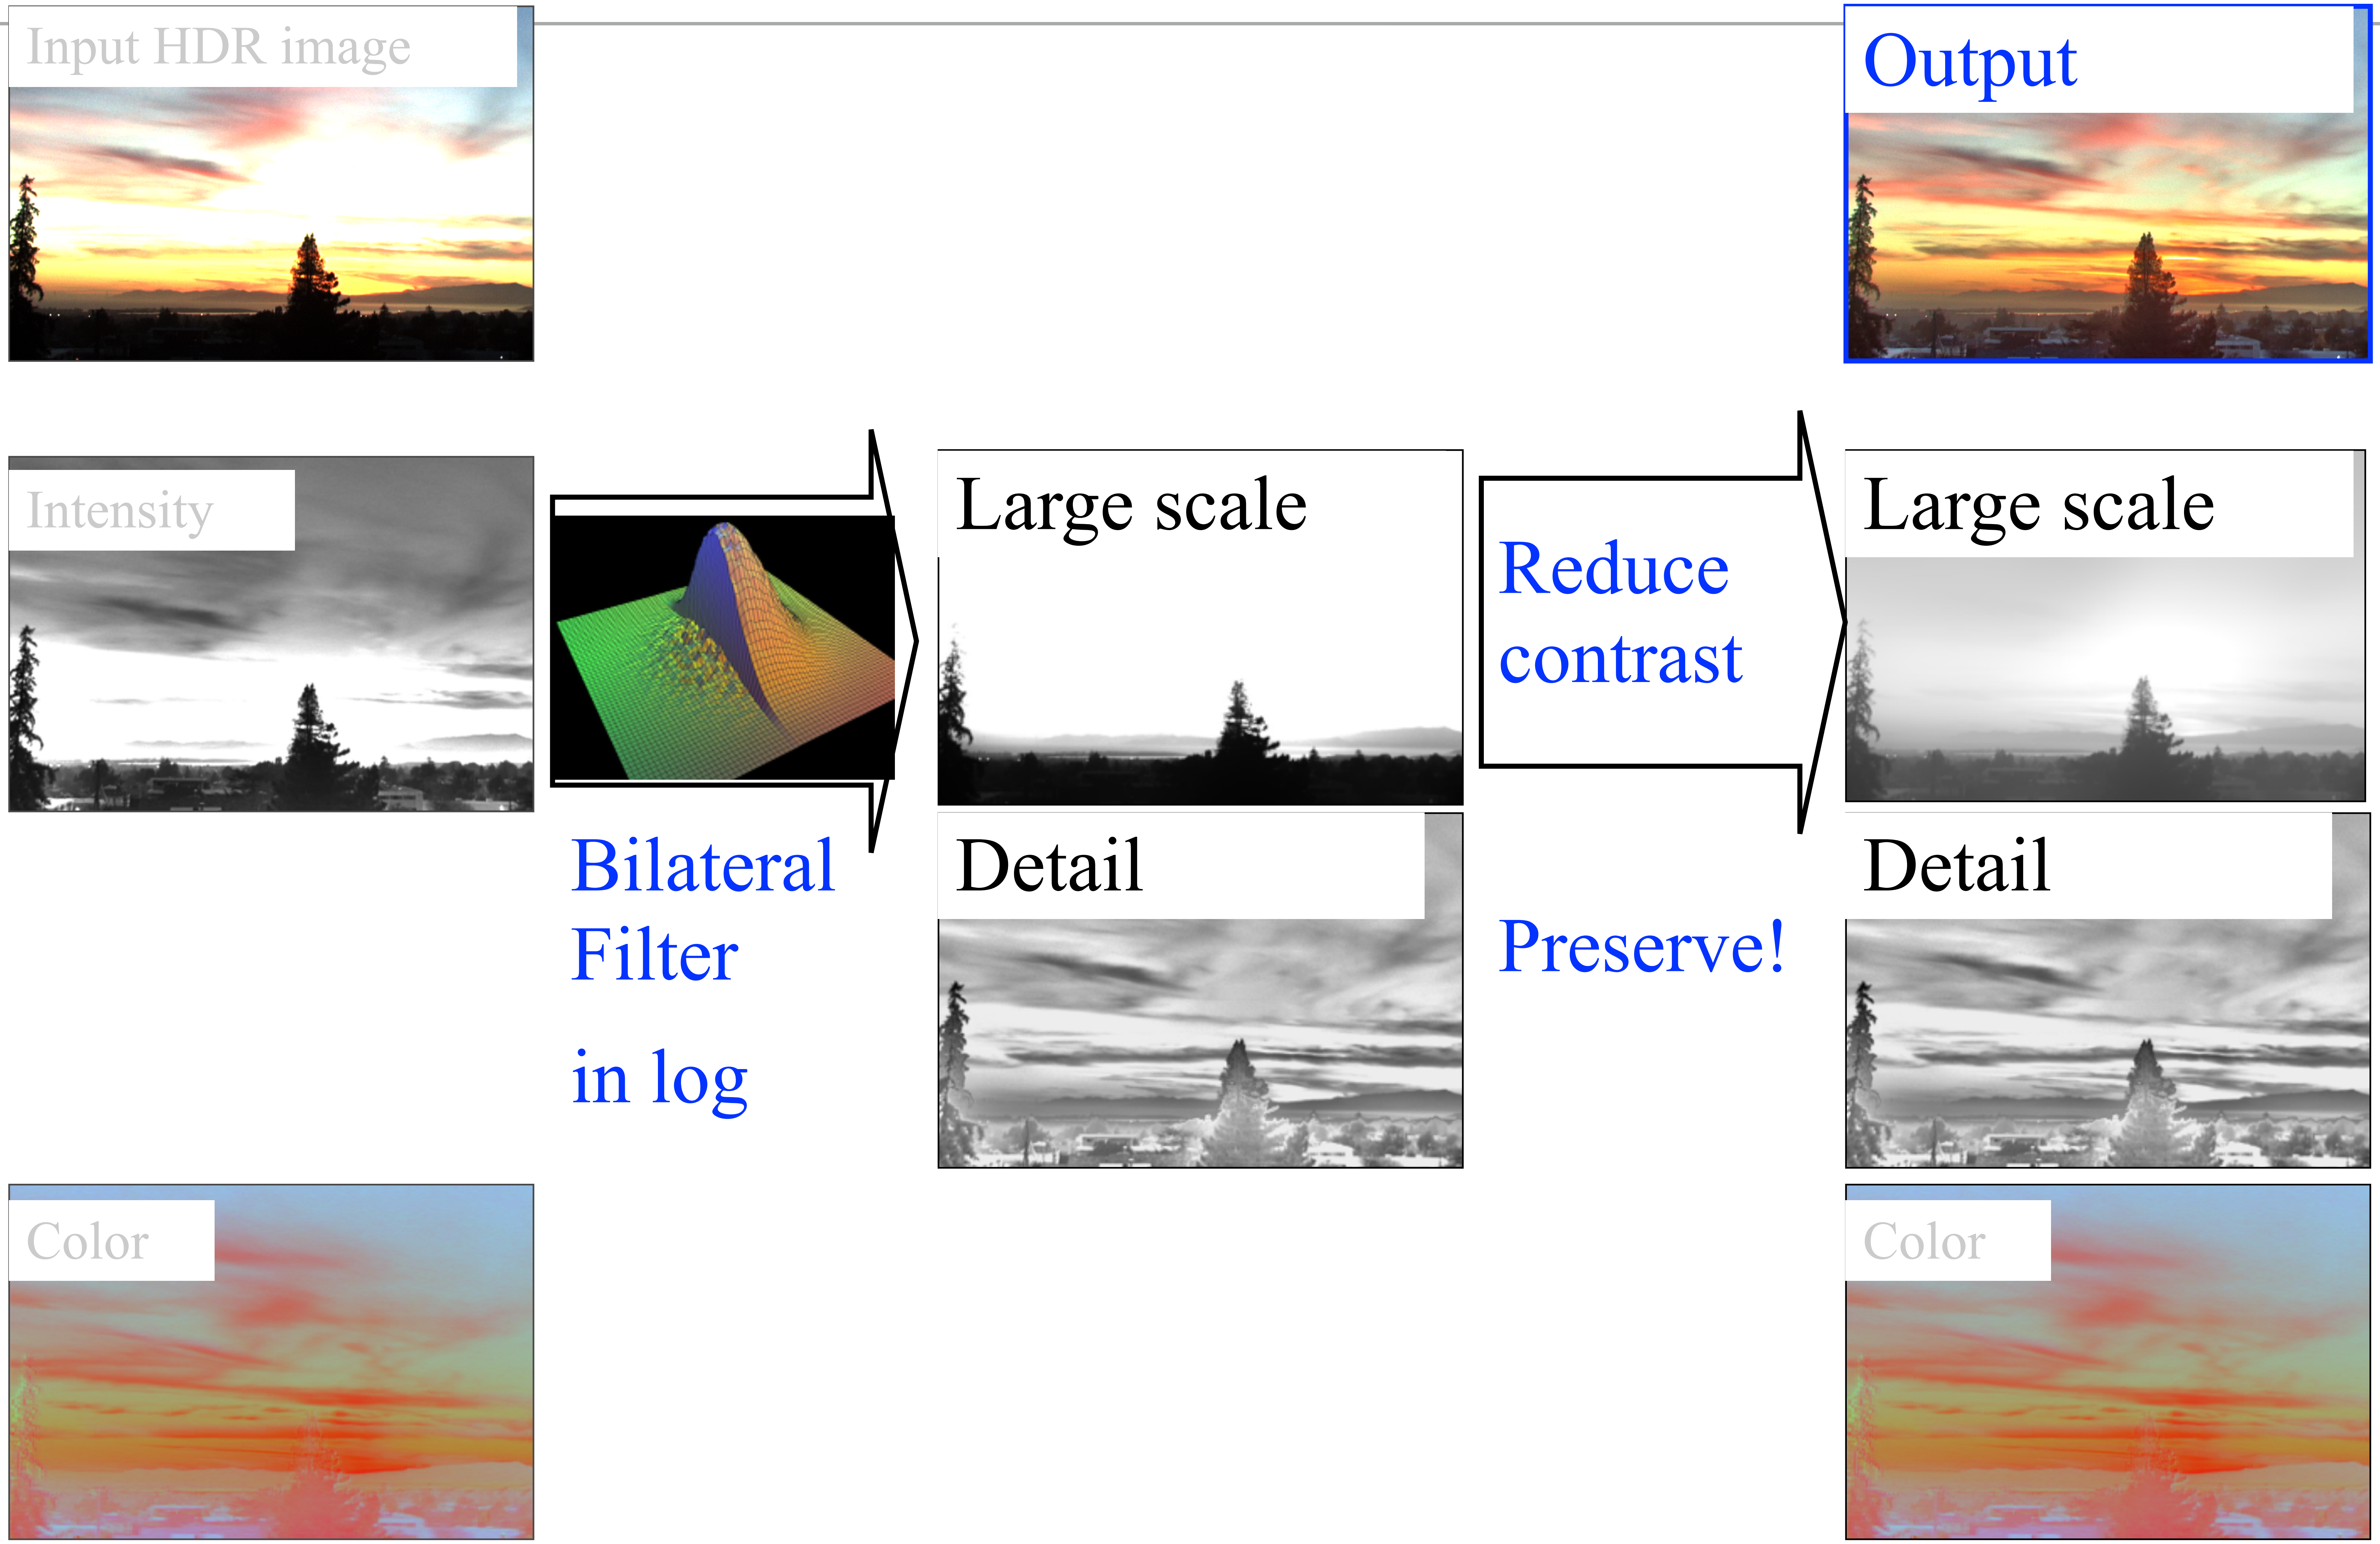

After a slide by Frédo Durand

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

Scale decomposition in log domain

inLog = $\log_{10}$(intensity)

inLogLarge = bilateralFilter(inLog)

inLogDetail = inLog - inLogLarge

hence:

- inLog = inLogDetail + inLogLarge, or
- intensity = $10^{\text{inLogDetail}} * 10^{\text{inLogLarge}}$

Now manipulate large-scale and detail separately

After a slide by Frédo Durand

Contrast reduction in log domain

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

Normalize so that the biggest value is 0 in log

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

After a slide by Frédo Durand

Contrast reduction in log domain

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

Normalize so that the biggest value is 0 in log

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

Optional: amplify detail by detailAmp

After a slide by Frédo Durand

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

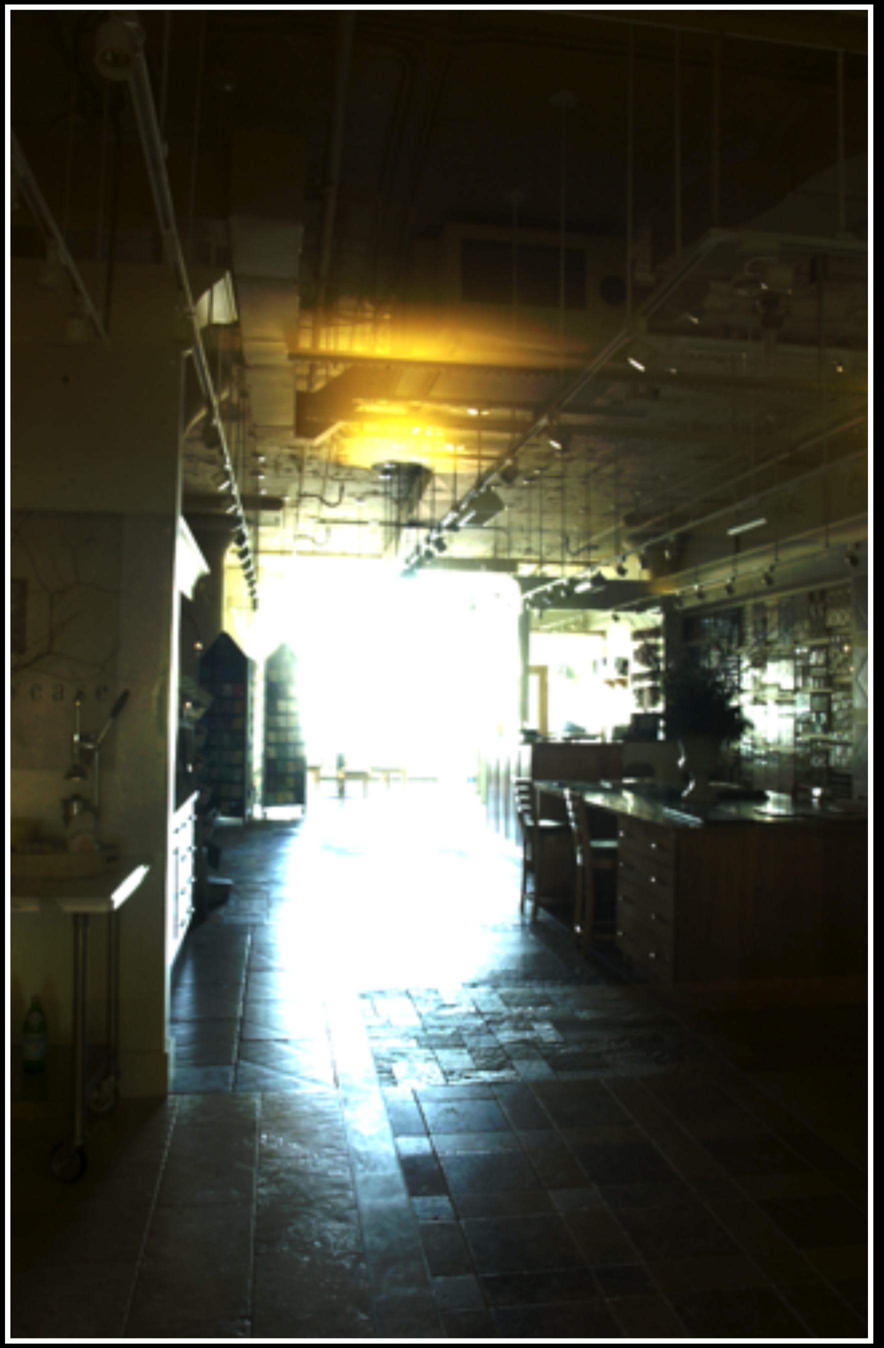

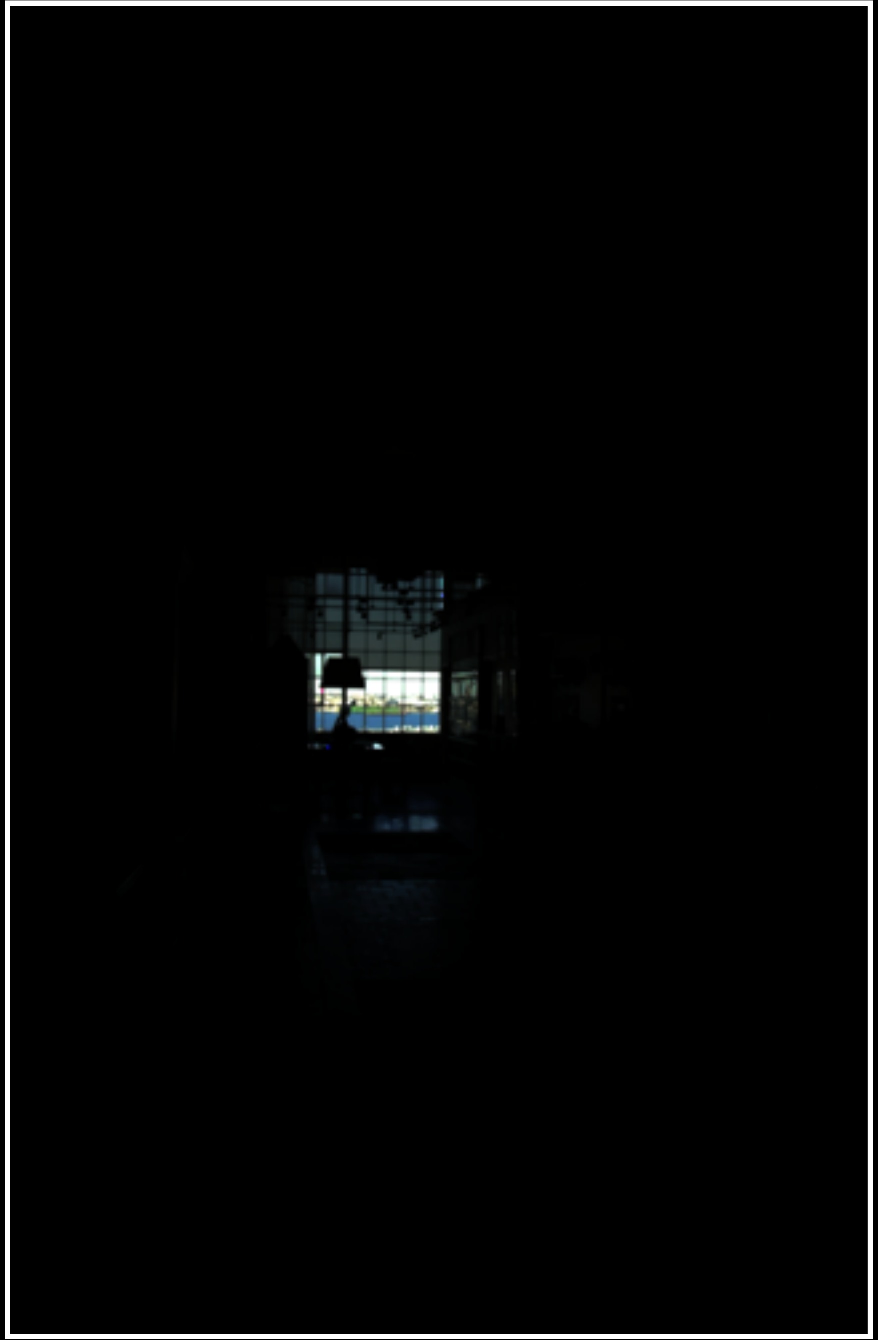

After a slide by Frédo Durand

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” (but with Local Laplacian Filter)

Photomatix “Details Enhancer”

After a slide by Frédo Durand

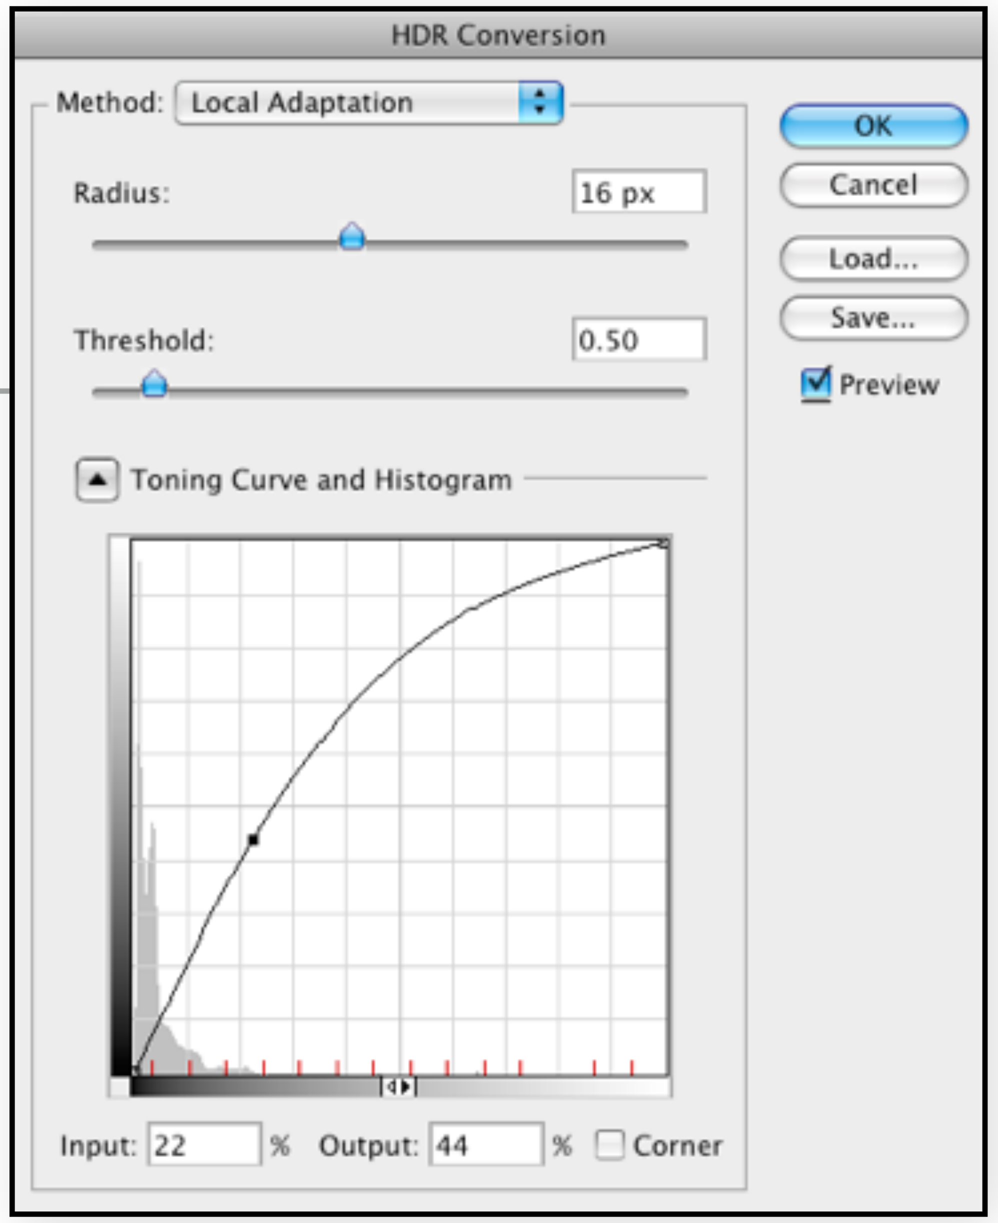

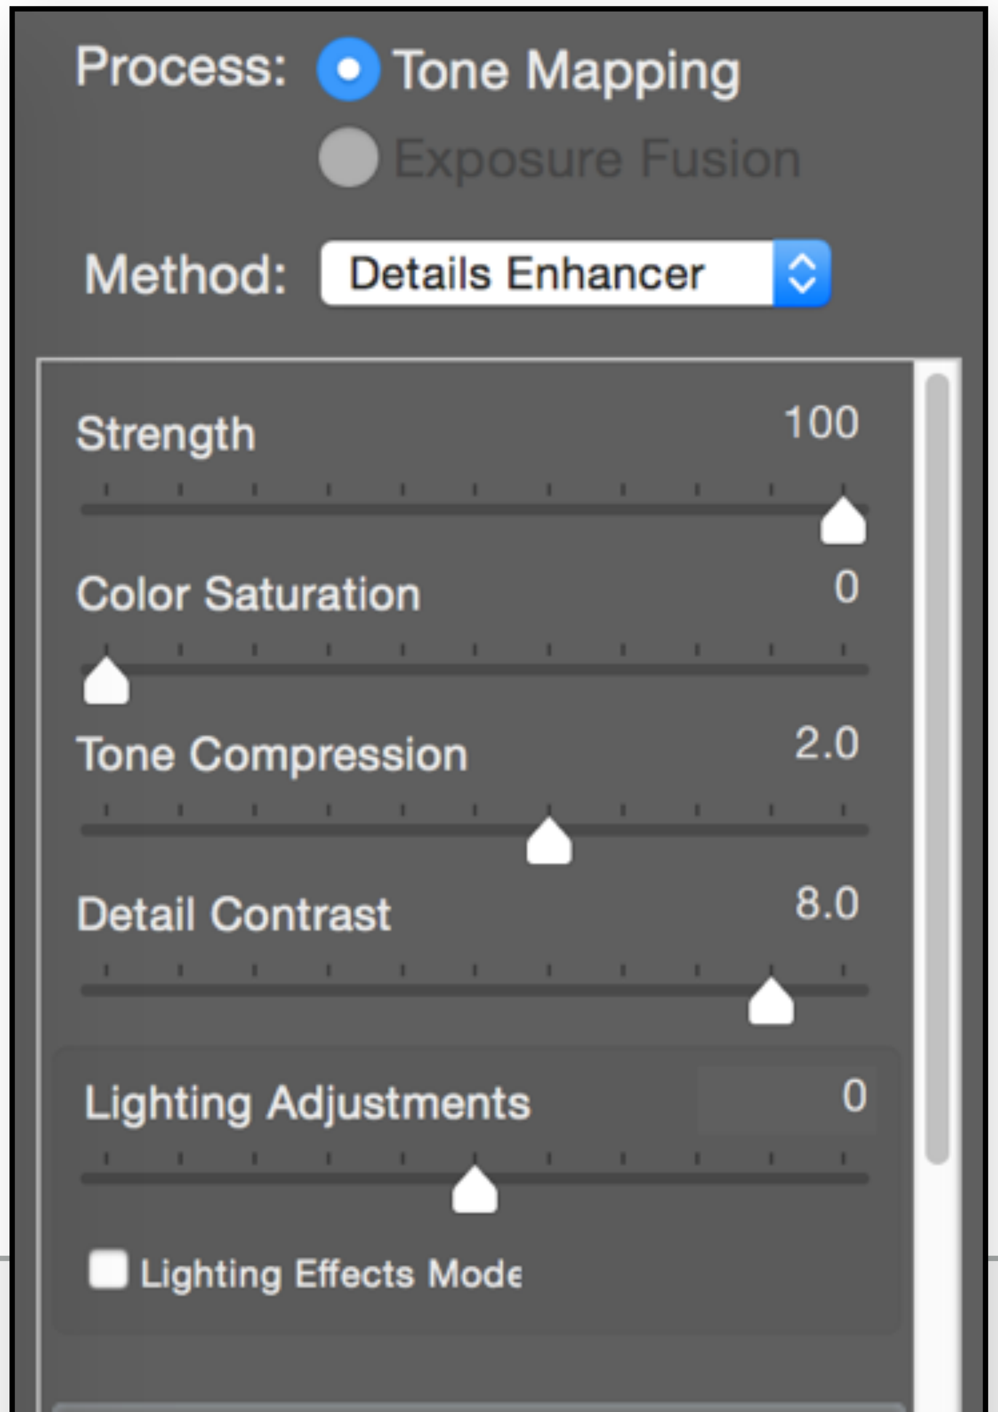

HDR vs. Tone mapping vs. Developing

Tone mapping is not something new or unique to HDR

- dodge-burn, dark room development, RAW processing

Your camera does tone mapping even if you don't do HDR

Film & digital SLRs have more dynamic range than display

- Enhance/preserve details, map larger range to smaller

HDR tone mapping and RAW processing (e.g. Lightroom Develop module) are doing fundamentally similar things

"Abandoned Ship at Point Reyes"

[Wojciech Jarosz 2014]

“Abandoned Ship at Point Reyes”

[Wojciech Jarosz 2014]

Up to a point

Noise gets amplified in the shadows

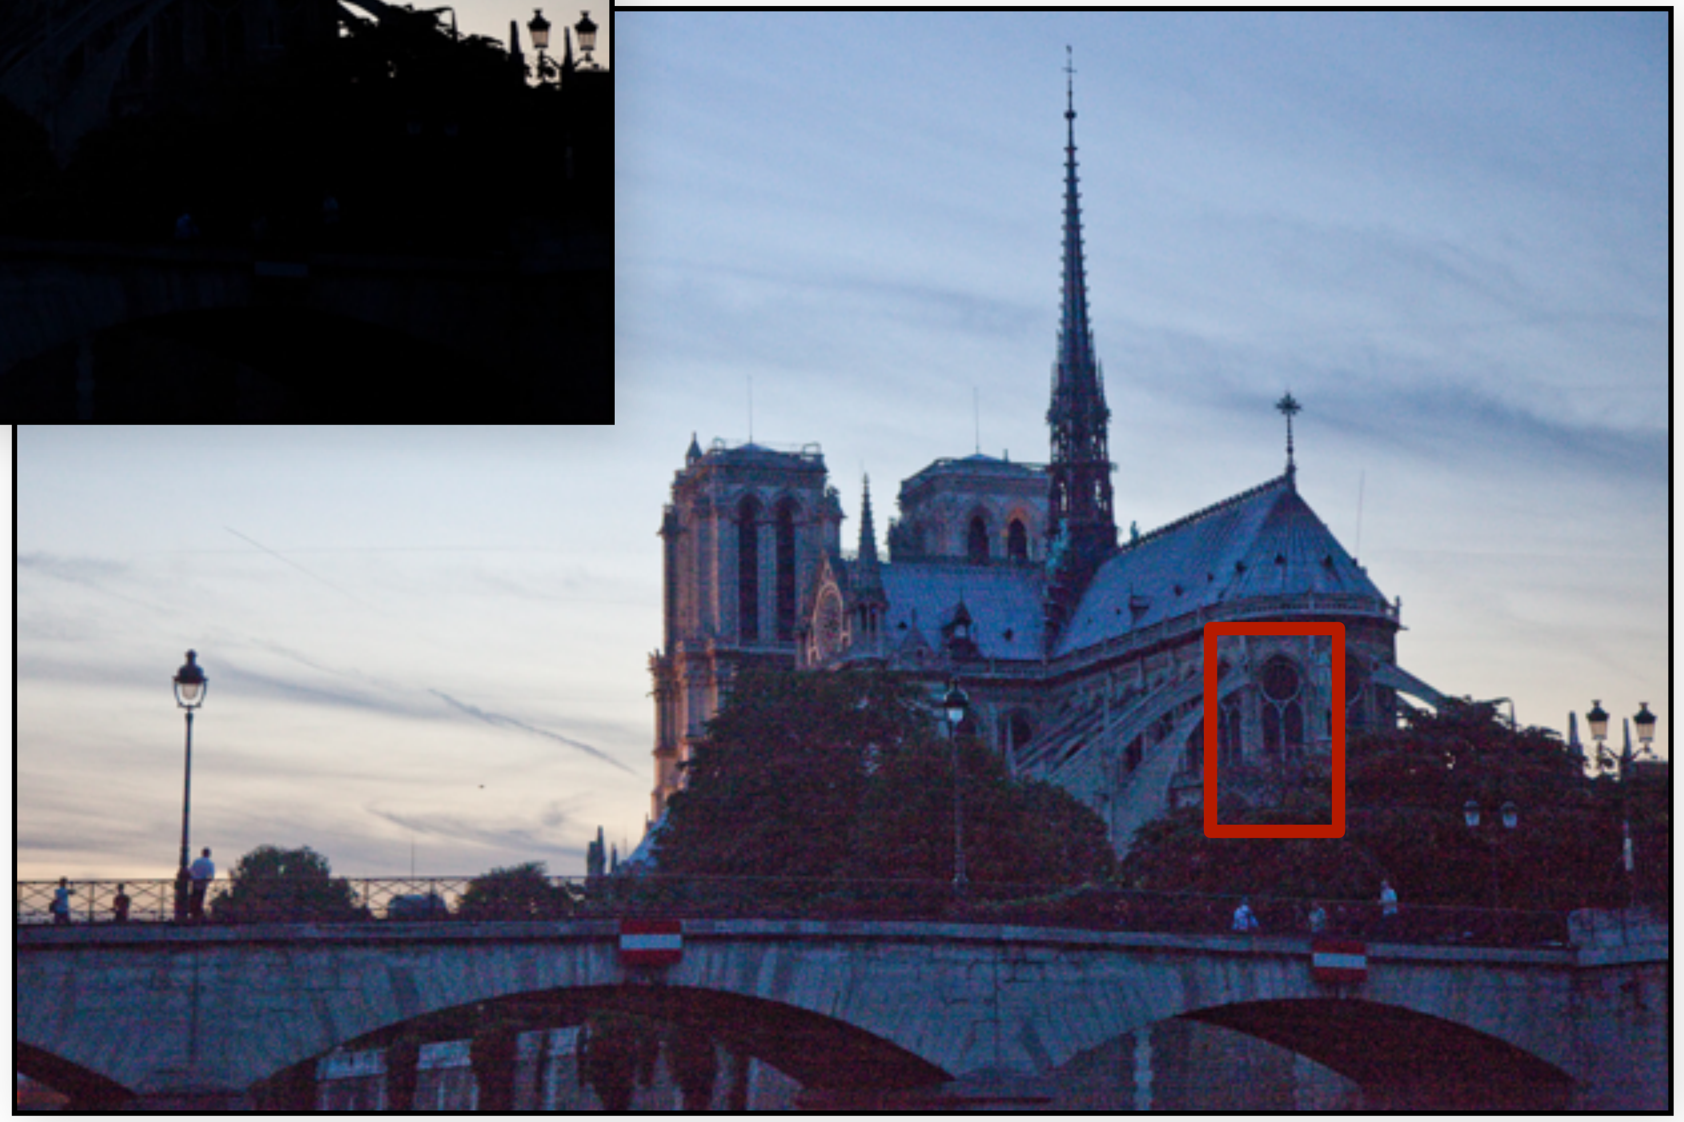

Optimal Weights

Problem setup

We may have multiple valid observations for a given pixel

They have different noise characteristics

How can we combine them optimally?

After a slide by Frédo Durand

Simple cases

They have the same noise

- Just take the average: see assignment 3
- Noise reduced by sqrt(N)

If one is a lot more noisy?

- Probably focus on the other one
- But if we only use the less noisy one, we don't get any noise reduction
- There has to be a way to use the second one at least a little bit

After a slide by Frédo Durand

Simple case

Two observations x & y of the same quantity

- with given variances $\sigma^2[x]$ $\sigma^2[y]$

Compute estimate as ax + (1-a)y

What is the optimal a?

After a slide by Frédo Durand

Minimize variance

Variance of the combination:

$$\sigma^2[ax + (1-a)y] = a^2\sigma^2[x] + (1-a)^2\sigma^2[y]$$

To minimize: set derivative to zero

$$2a\sigma^2[x] - 2(1-a)\sigma^2[y] = 0$$

$$a\left(\sigma^2[x] + \sigma^2[y]]\right) = \sigma^2[y]$$

$$a = \frac{\sigma^2[y]}{\sigma^2[x] + \sigma^2[y]]}$$

After a slide by Frédo Durand

Optimal combination

$$\frac{\sigma^2[y]}{\sigma^2[x]+\sigma^2[y]]}x+\frac{\sigma^2[x]}{\sigma^2[x]+\sigma^2[y]]}y$$

re-arrange:

$$\frac{\sigma^2[x]\sigma^2[y]}{\sigma^2[x]+\sigma^2[y]}\left(\frac{1}{\sigma^2[x]}x+\frac{1}{\sigma^2 y}y\right)$$

normalization term

The optimal combination should scale estimators according to the inverse of their variance

After a slide by Frédo Durand

Verify for same variance

$$\frac{\sigma^2[x]\sigma^2[y]}{\sigma^2[x] + \sigma^2[y]} \left(\frac{1}{\sigma^2[x]} x + \frac{1}{\sigma^2 y} y \right)$$

After a slide by Frédo Durand

General formula

Weight each estimator by the inverse variance

$$\frac{1}{\sum \frac{1}{\sigma^2[x_i]}} \sum \frac{x_i}{\sigma^2[x_i]}$$

After a slide by Frédo Durand

Optimal HDR

Recall: Assembling HDR (linear case)

Figure out scale factor between images

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

Compute weight map w_i for each image

- binary so far

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

Pixel noise and variance

Recall: noise is characterized by its variance

- i.e. each pixel value comes from a true value plus some noise added

We can calibrate this noise by taking multiple exposures, or we can derive variance equations using pen and paper

After a slide by Frédo Durand

Sources of noise

Photon noise

- variance proportional to signal
- dominates for dark pixels

Read noise

- constant variance
- dominates for dark pixels

For a pixel value x: $\sigma^2[x] \approx ax + \sigma^2_{read}$

- where a and σ^2_{read} depend on the camera and ISO

After a slide by Frédo Durand

Optimal weights

Recall irradiance formula: $I_i(x, y) = clip(k_i * L(x,y)+n)$

and HDR merging formula:

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

1/k amplifies signal and noise:

$$\sigma^2[\frac{1}{k_i} I_i(x,y)] = \frac{1}{k_i^2} \left[a I_i(x,y) + \sigma_{read}^2\right]$$

replace I by irradiance

$$\sigma^2[\frac{1}{k_i} I_i(x,y)] = \frac{1}{k_i^2} \left[a k_i L(x,y) + \sigma_{read}^2\right]$$

After a slide by Frédo Durand

Variance of one scaled image

$$\sigma^2[\frac{1}{k_i} I_i(x,y)] = \frac{1}{k_i^2}\left[aI_i(x,y) + \sigma_{read}^2\right]$$

$$\sigma^2[\frac{1}{k_i} I_i(x,y)] = \frac{1}{k_i^2}\left[ak_iL(x,y) + \sigma_{read}^2\right]$$

$$\sigma^2[\frac{1}{k_i} I_i(x,y)] = \frac{a}{k_i} L(x,y) + \frac{1}{k_i^2}\sigma_{read}^2$$

If we only look at photon noise, mostly proportional to scale factor

Ideally, should all be calibrated

Note that we ignored ISO variations

After a slide by Frédo Durand

Improved weight maps

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

Variance per pixel per image: $\sigma^2[\frac{1}{k_i} I_i(x,y)] \approx \frac{a}{k_i} L(x,y)$

replace w_i by w_i'

- still use w_i to reject dark and bright pixels
- but also weight by inverse variance

$$w_i'(x,y) = w_i(x,y) / \frac{a}{k_i} L(x,y)$$

After a slide by Frédo Durand

Improved weight maps

New formula: $out(x,y) = \frac{1}{\sum w'_i(x,y)} \sum w'_i(x,y) \frac{1}{k_i} I_i(x,y)$

with $w'_i(x,y) = w_i(x,y) / \frac{a}{k_i} L(x,y)$

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

- the two k_i in the main sum cancel each other
- a and L(x,y) are constant per pixel and present both in the main sum and the normalization. Get rid of them.

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

After a slide by Frédo Durand

New formula

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

The radiant power reaching the pixel has disappeared. All pixels of a given exposure are weighted the same.

This is because the relative photon noise changes similarly for all pixels between a pair of exposures

- Would be different with read noise

ki has disappeared from the main sum. The images are not really rescaled to scene radiant power

- But they indirectly are because of the normalization
- Recall that ki and 1/ki used to both appear

After a slide by Frédo Durand

Naive weights

Weighted by 1/ki

After a slide by Frédo Durand

References

http://www.macs.hw.ac.uk/bmvc2006/papers/372.pdf

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

- full noise model
- exploit ISO
- Also optimizes the set of exposures

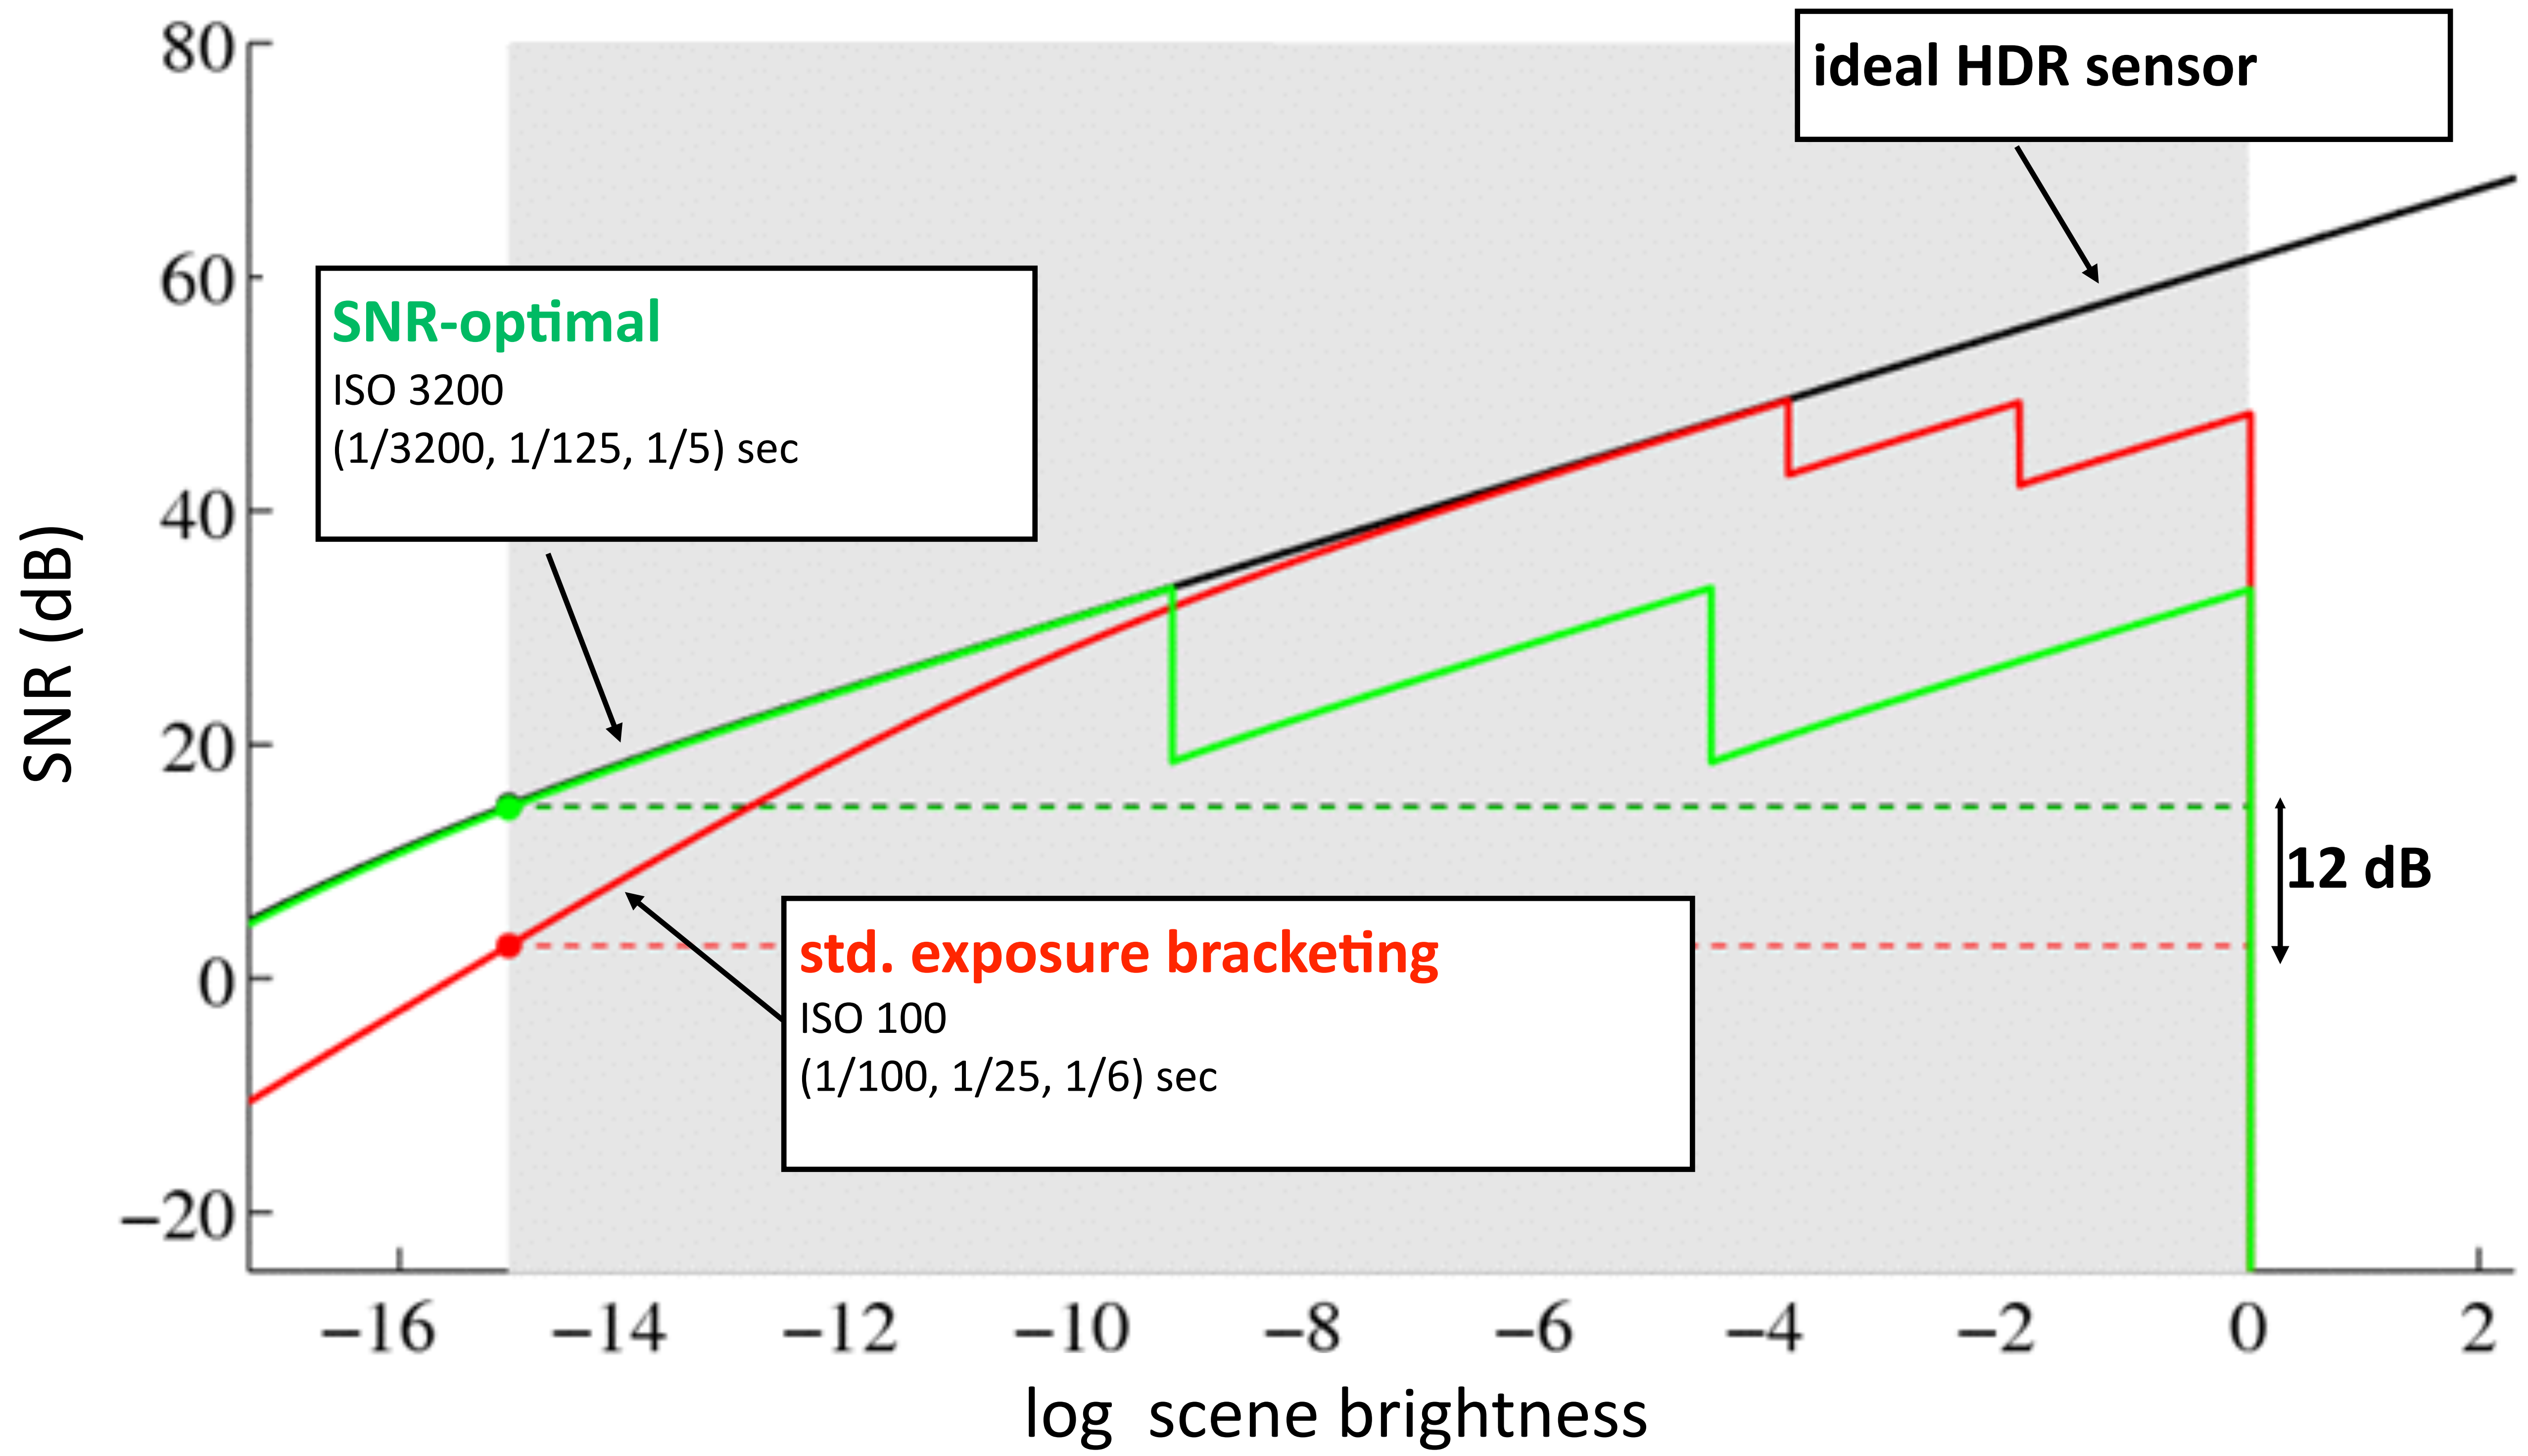

After a slide by Frédo Durand

HDR file formats

Storing/Encoding HDR images

Most formats are lossless

Portable float map (*.PFM)*:

- straight dump of 32-bit floating-point pixel values
- fwrite(buffer, 3 * sizeof(float), buffer.size(), f);
- quick and dirty, common in research

Radiance image (*.PIC*, *.HDR*):

- 8 bits per r,g,b as usual, plus 8 bits of shared exponent (rgbe)
- Introduced by Greg Ward for Radiance (light simulation)

Float-point TIFF (*.TIFF*):

- log encoded 24- or 32-bit values

OpenEXR (*.EXR*):

- by Industrial Light & Magic, also standard in graphics hardware
- 16- or 32-bit floating-point per channel
- popular in movie industry, not as much in photography

Adobe Digital NeGative (*.DNG*)

- Specific for RAW files, intended to avoid lock-in to undocumented, proprietary formats

HDR formats

Summary of all HDR encoding formats (Greg Ward):

- http://www.anyhere.com/gward/hdrenc/hdr_encodings.html

Greg's notes:

- http://www.anyhere.com/gward/pickup/CIC13course.pdf

http://www.openexr.com/

High Dynamic Range Video Encoding (MPI)

- http://www.mpi-sb.mpg.de/resources/hdrvideo/

HDR photography advice

"Good" vs "Bad" HDR/tone mapping

[Tommy Clark]

[Tommy Clark]

"Good" vs "Bad" HDR/tone mapping

Generally avoid:

- halos
- excessive noise
- black clouds
- HDR on people or out-of-focus regions

Not all images require HDR tone mapping!

My opinion: HDR tone mapping is like special effects, it's better if you can't tell it's being used

How to take your own HDR photos

Keep in mind...

HDR + tone mapping/processing won't make a boring photo interesting

Do as much as possible "in-camera"

Learn composition (rule of thirds, line, form, texture, etc)

Shoot in the right light

- golden hour

“Kiliçlar Valley”

[Wojciech Jarosz 2012]

"Kiliçlar Valley"

[Wojciech Jarosz 2012]

Use the histogram

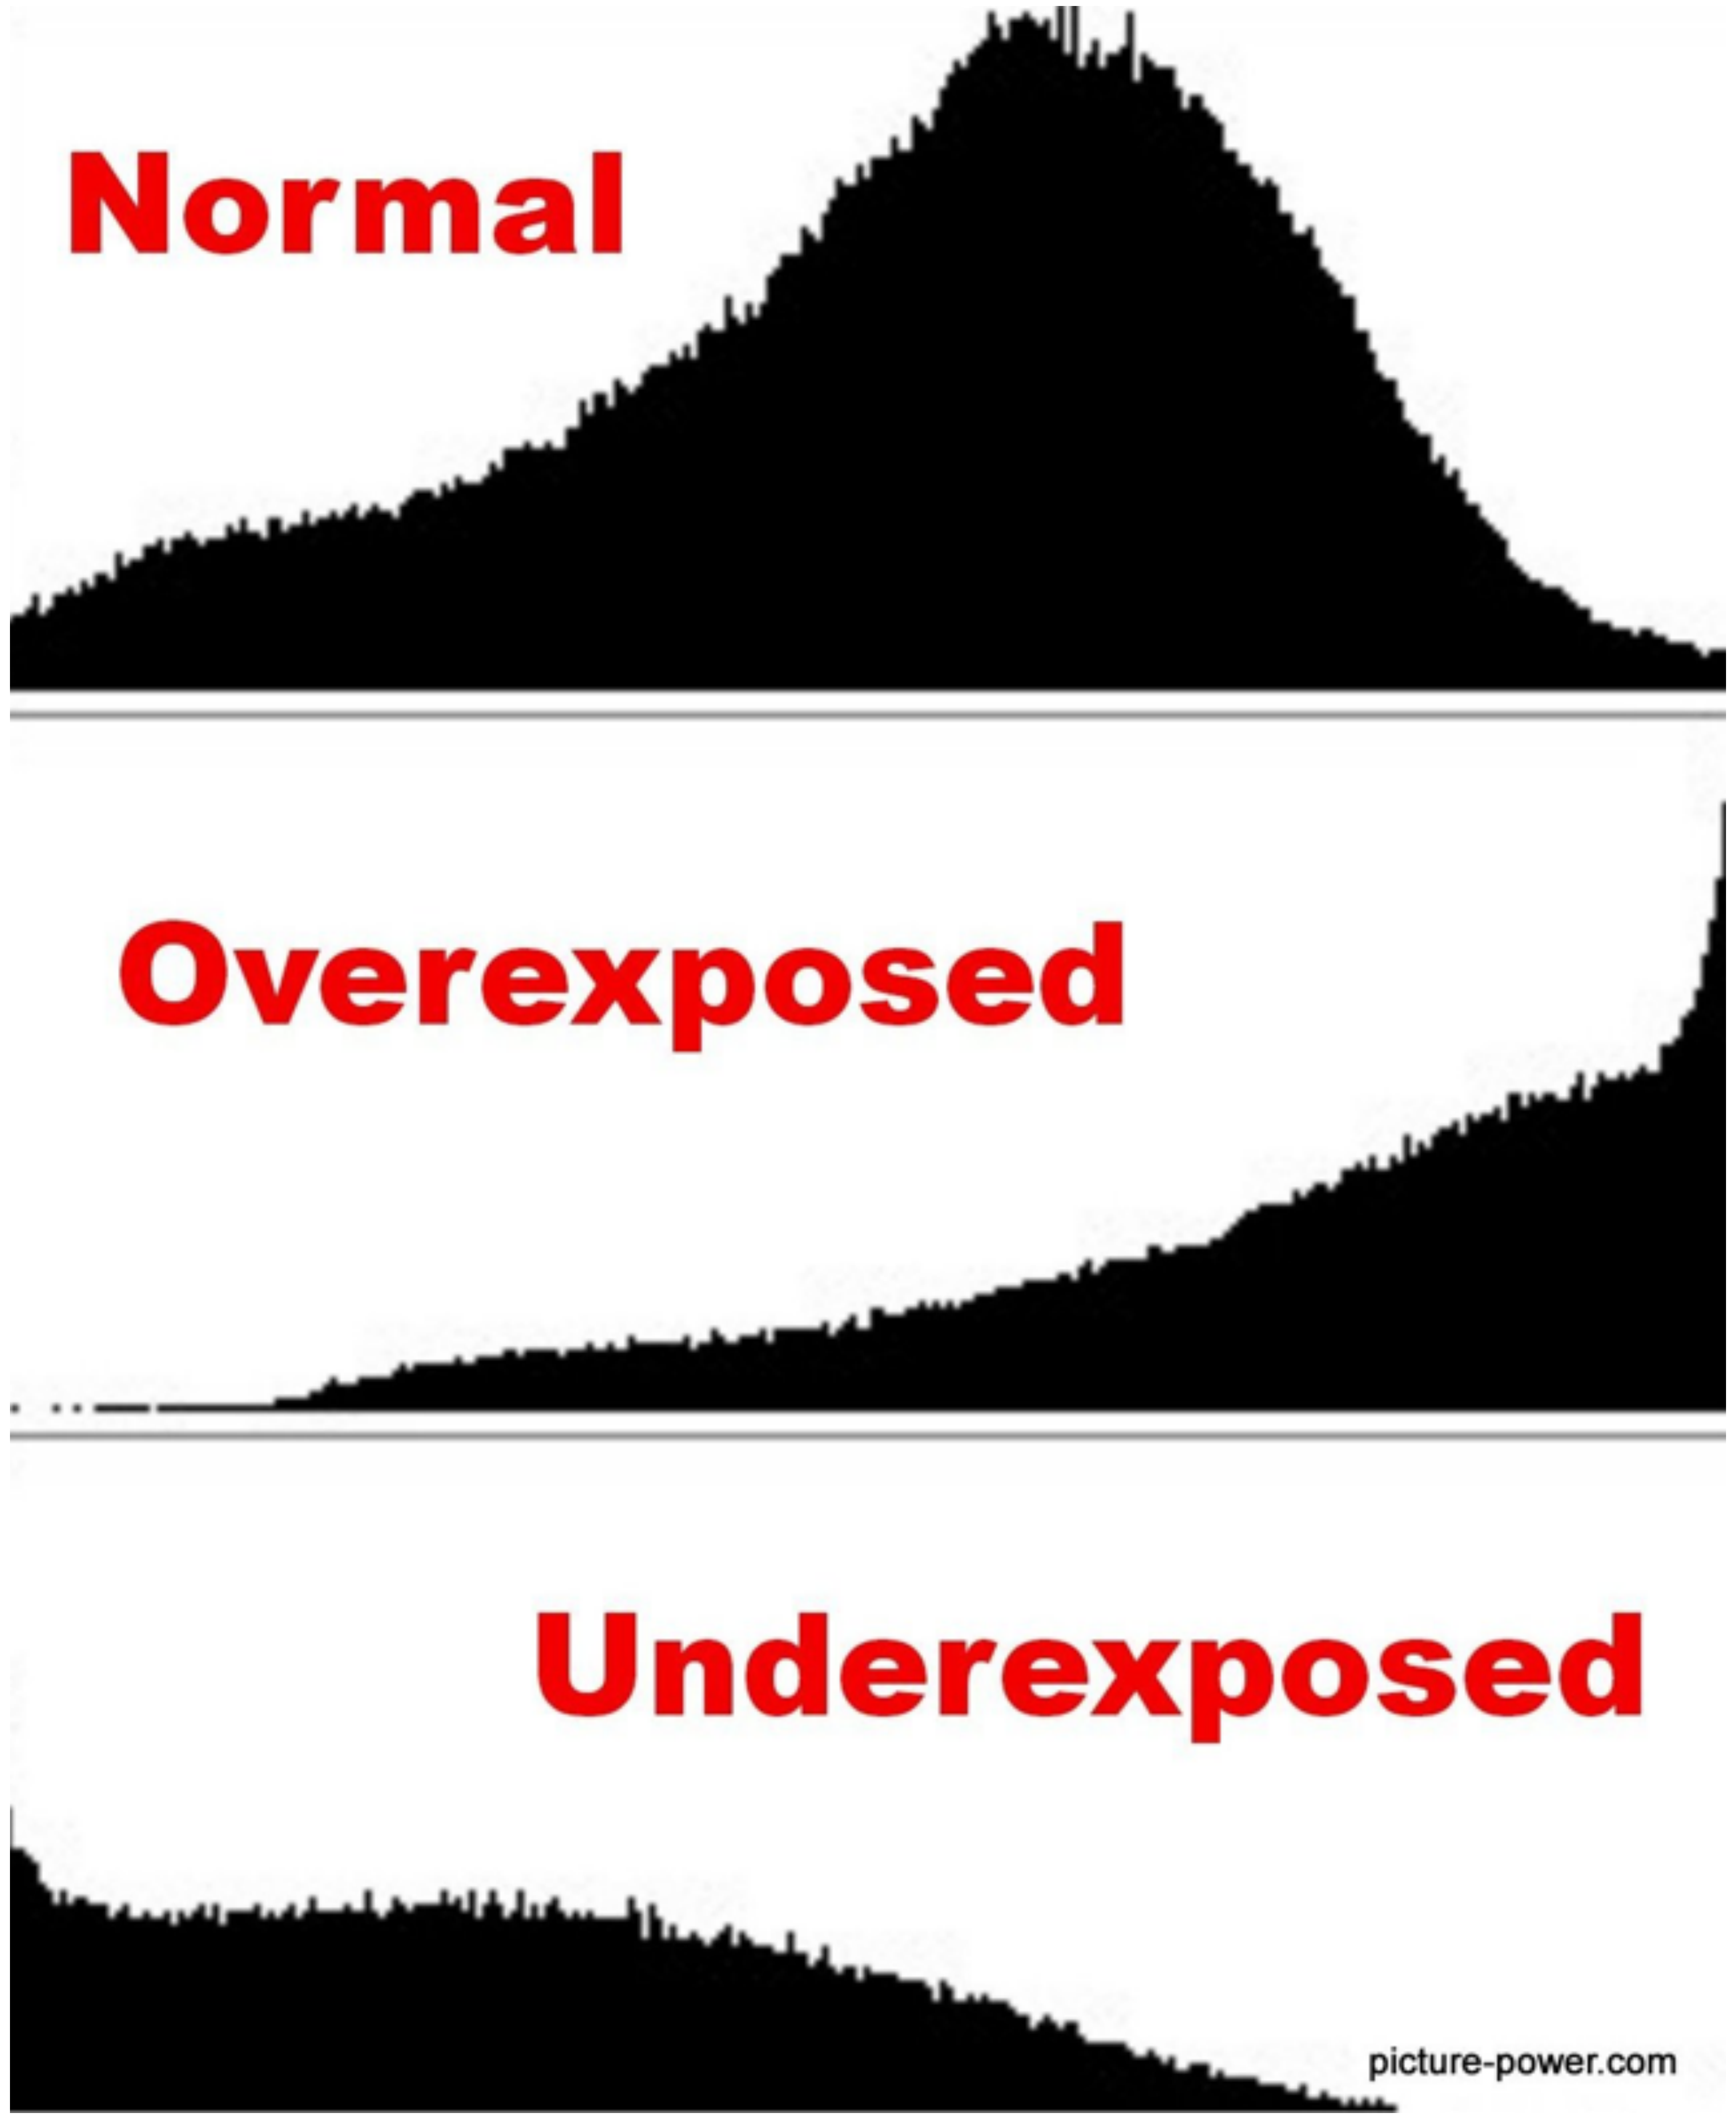

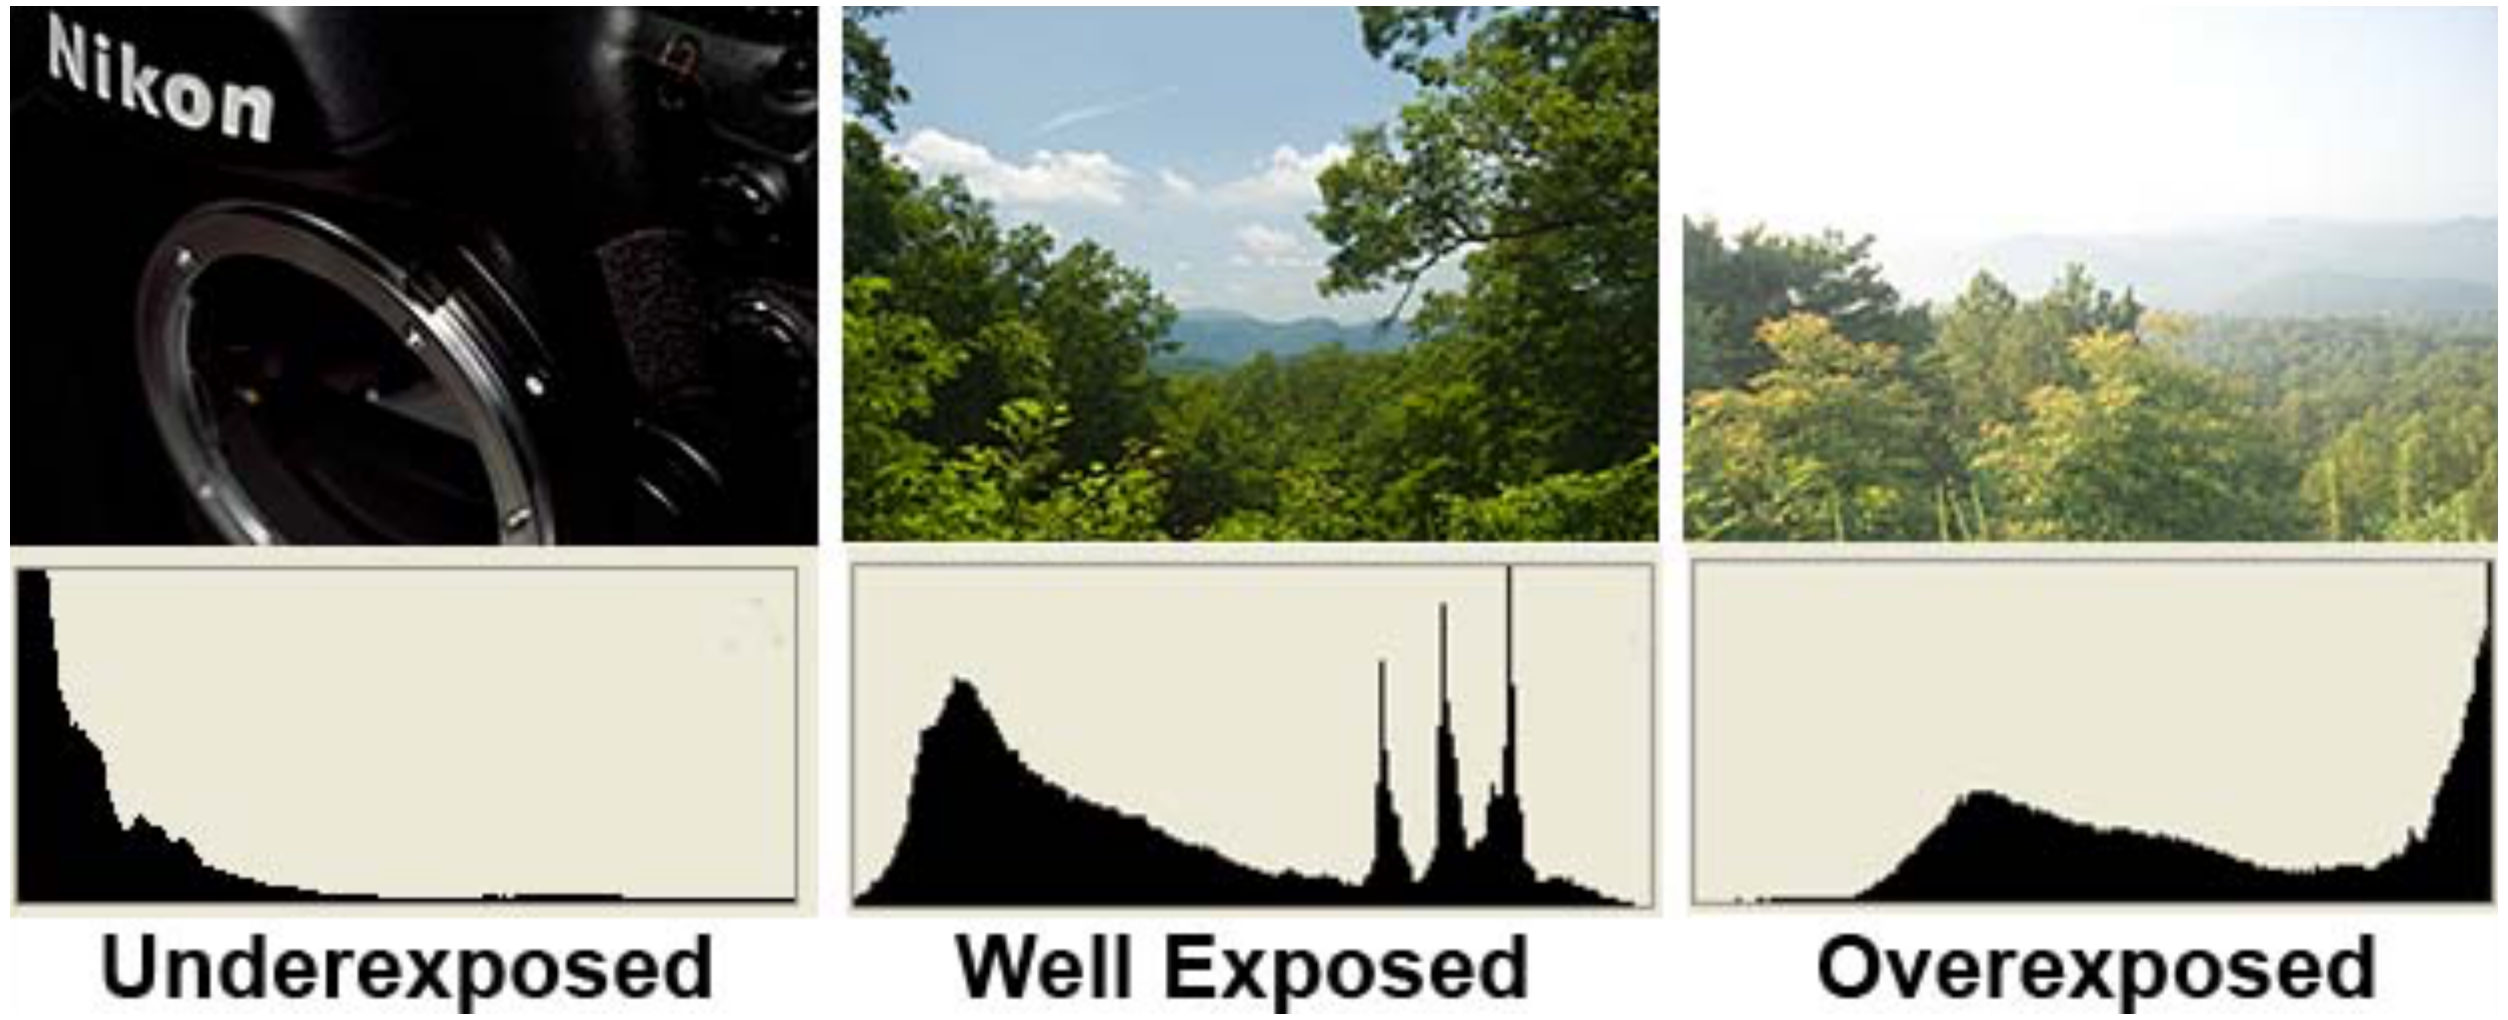

What's happening in this photo?

source: flickr user neofob

Polarizing Filter

Polarization

Without Polarizer

With Polarizing Filter

Polarization

Without Polarizer With Polarizing Filter

Effect of Polarization


[Wojciech Jarosz]

Effect of Polarization

Exposure Bracketing

I typically:

- Use aperture for desired depth of field
- Use lowest possible ISO (reduce noise)
- Control exposure with shutter speed
- Use a tripod **and** a remote trigger (mirror lock-up)

“Lake Bled”

[Wojciech Jarosz 2012]

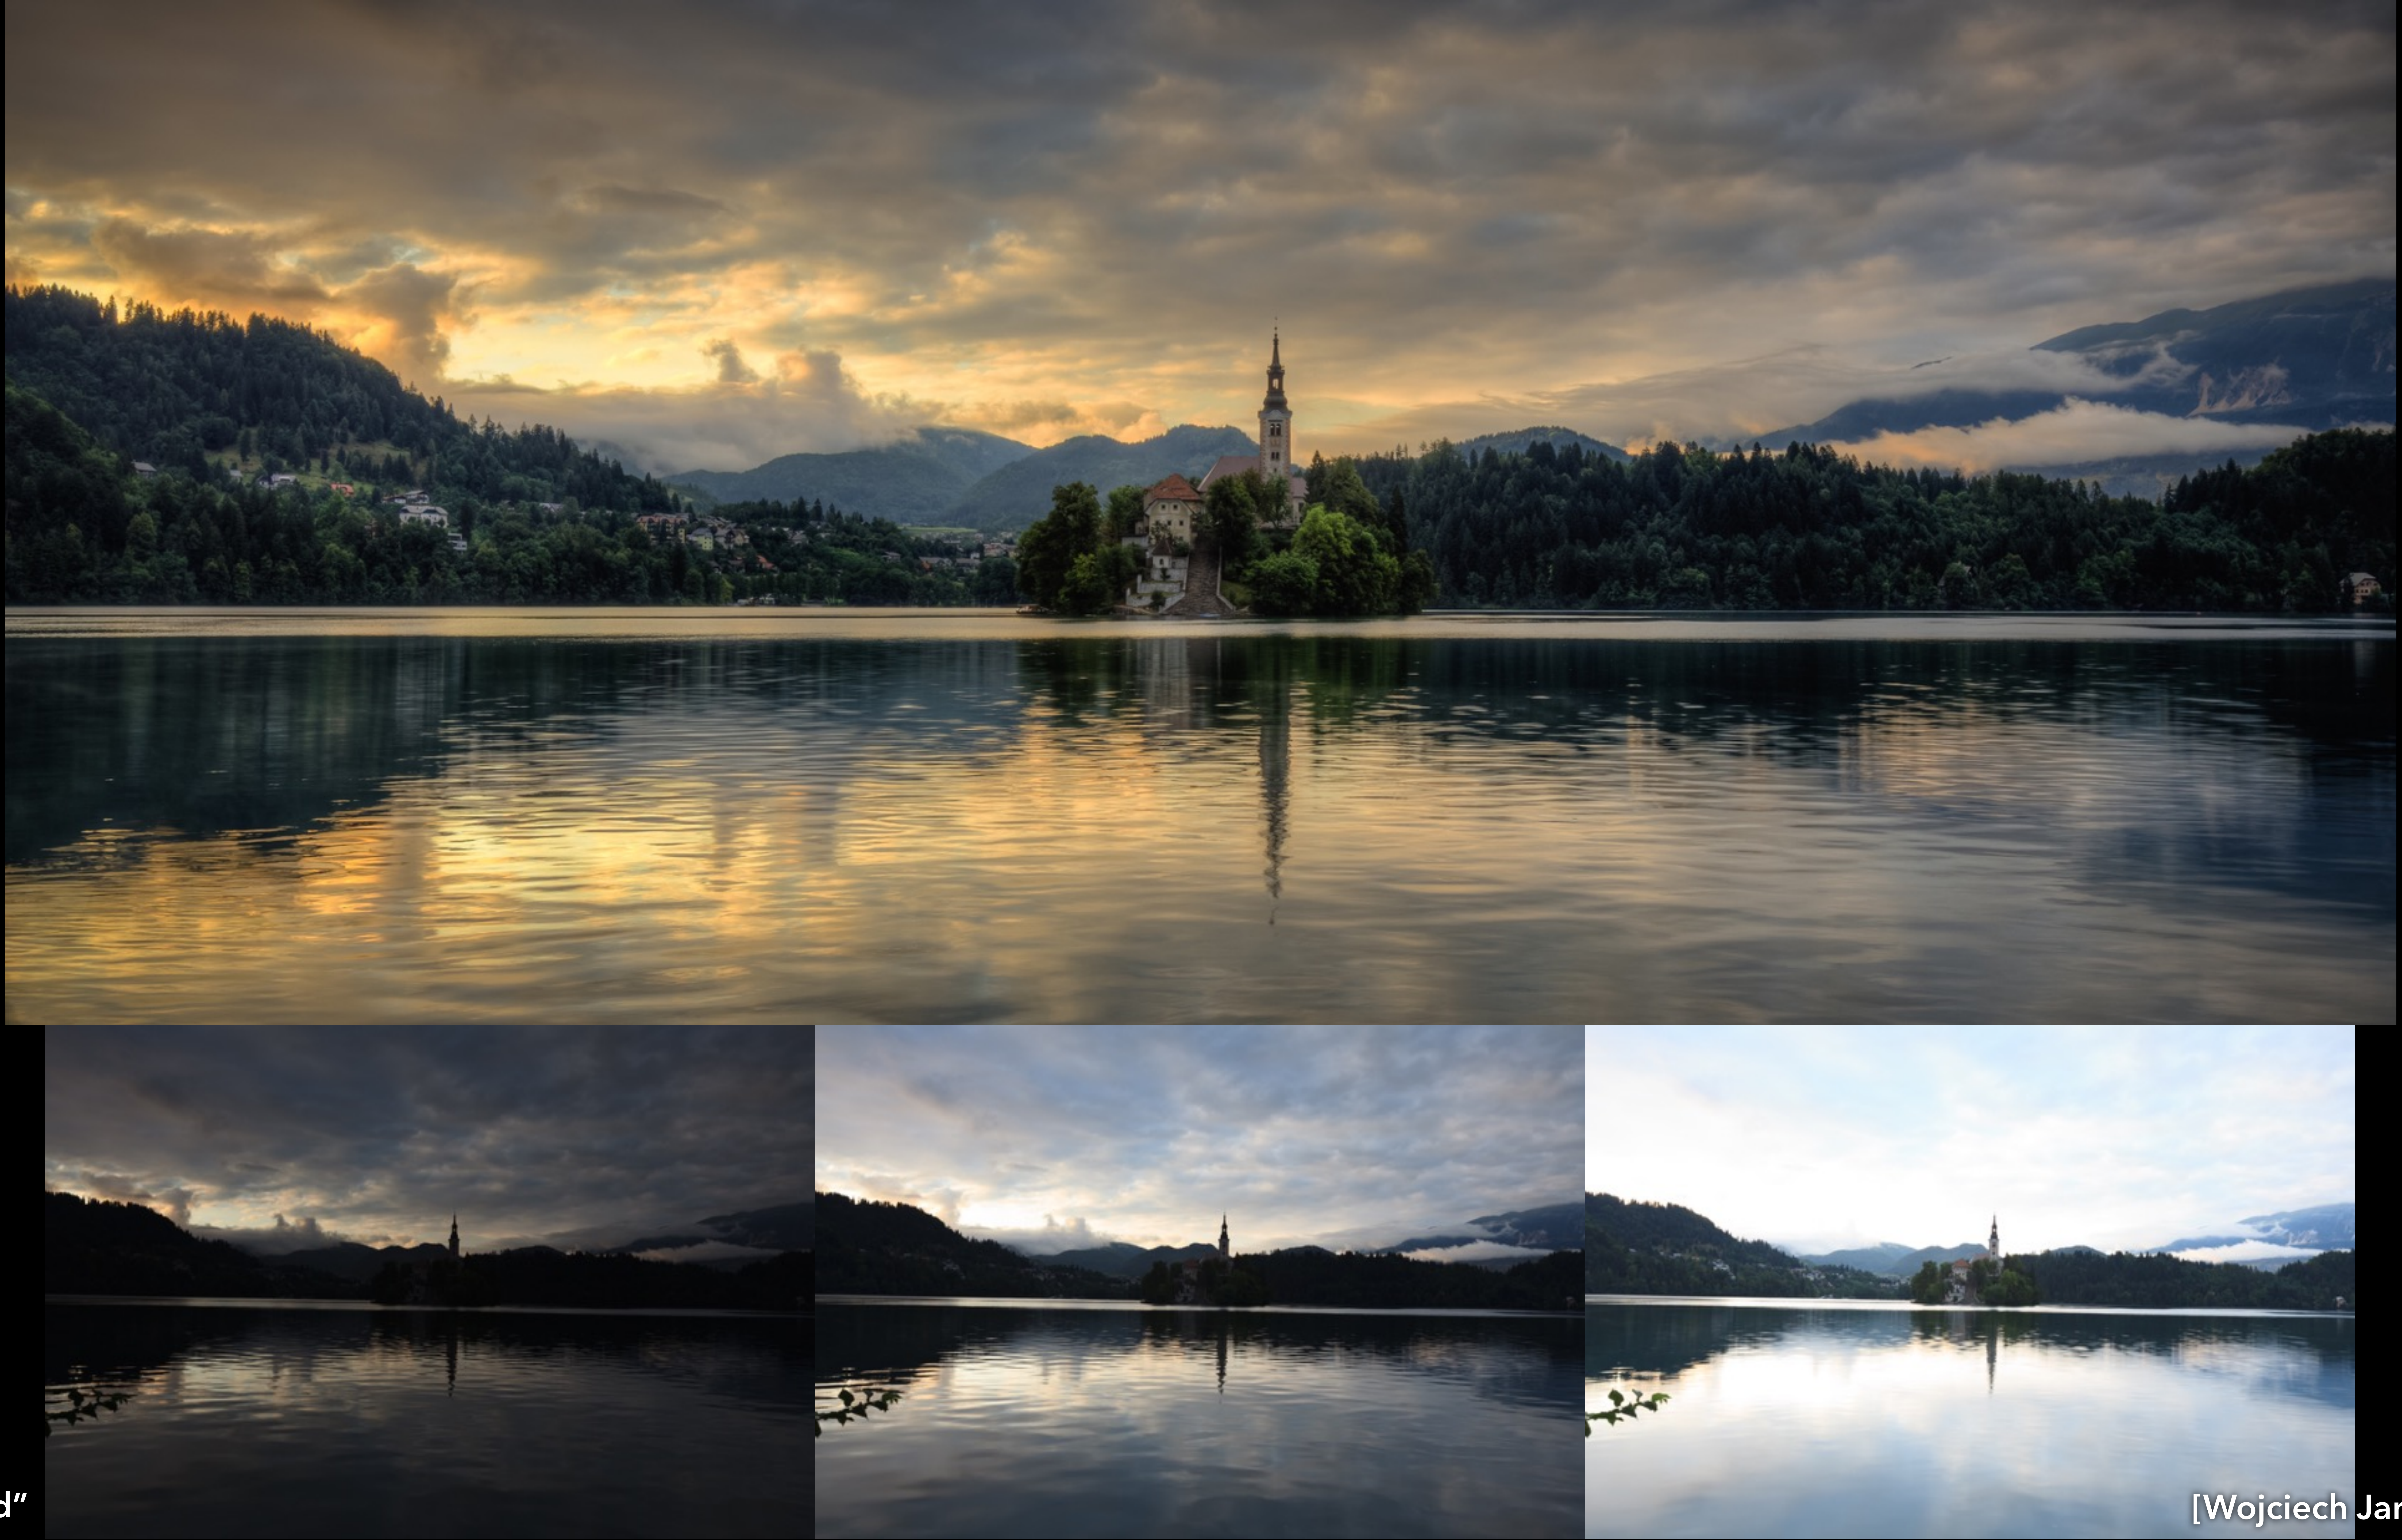

“Lake Bled”

[Wojciech Jarosz 2012]

“La Jolla Hospitals Reef”

[Wojciech Jarosz 2013]

“La Jolla Hospitals Reef”

[Wojciech Jarosz 2013]

[Wojciech Jarosz 2015]

[Wojciech Jarosz 2015]

Recap

High dynamic range (HDR) imaging is useful, and a new aesthetic

- but is not necessary in all photographic situations

Low dynamic range (LDR) tone mapping methods can also be applied to HDR scenes

- but reducing very HDR scenes to 8 bits for JPEG using only global methods is hard

Local methods reduce large-scale luminance changes (across the image) while preserving local contrast (across edges)

- use edge-preserving filters to avoid halos

After a slide by Marc Levoy

HDR practice

Slide credits

Frédo Durand

Marc Levoy