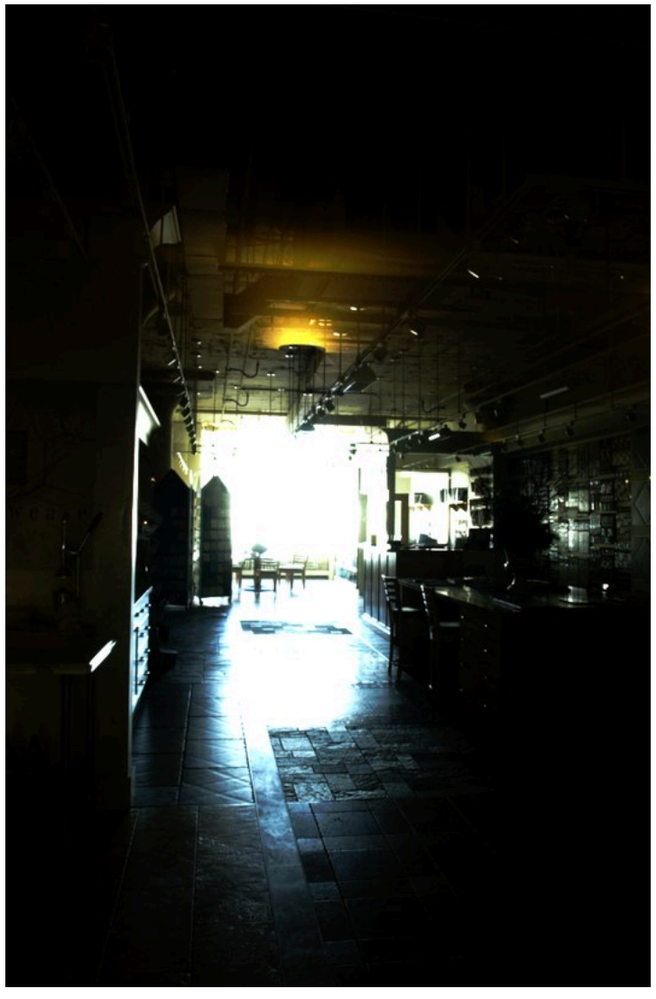
High-Dynamic-Range Imaging & Tone Mapping

Fredo Durand
MIT EECS 6.8370/6.8371

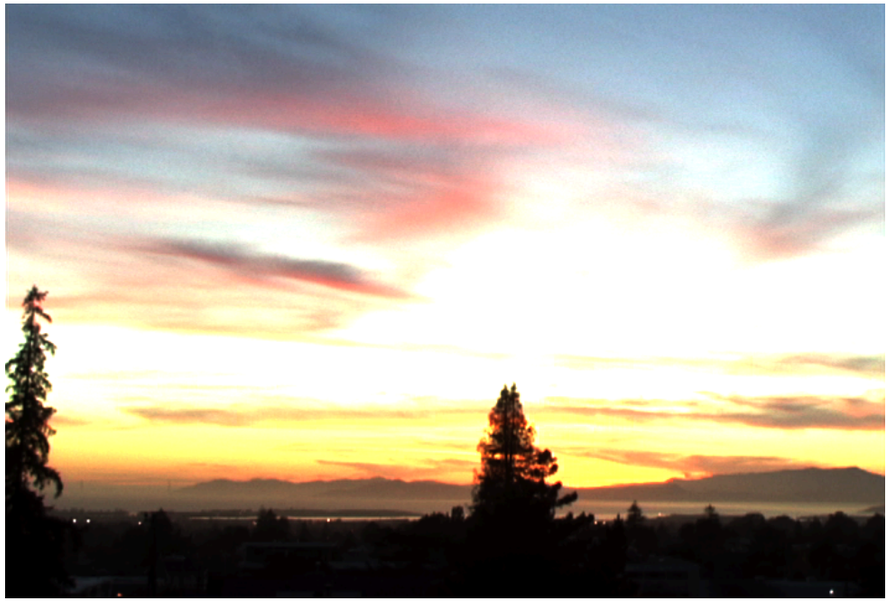
The dynamic range challenge

Examples

- **Inside is too dark**
- **Outside is too bright**

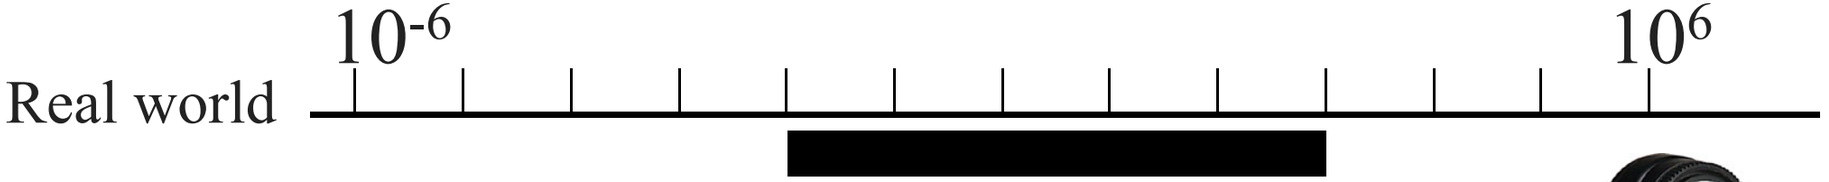


- **Sun overexposed**
- **Foreground too dark**



Real world dynamic range

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

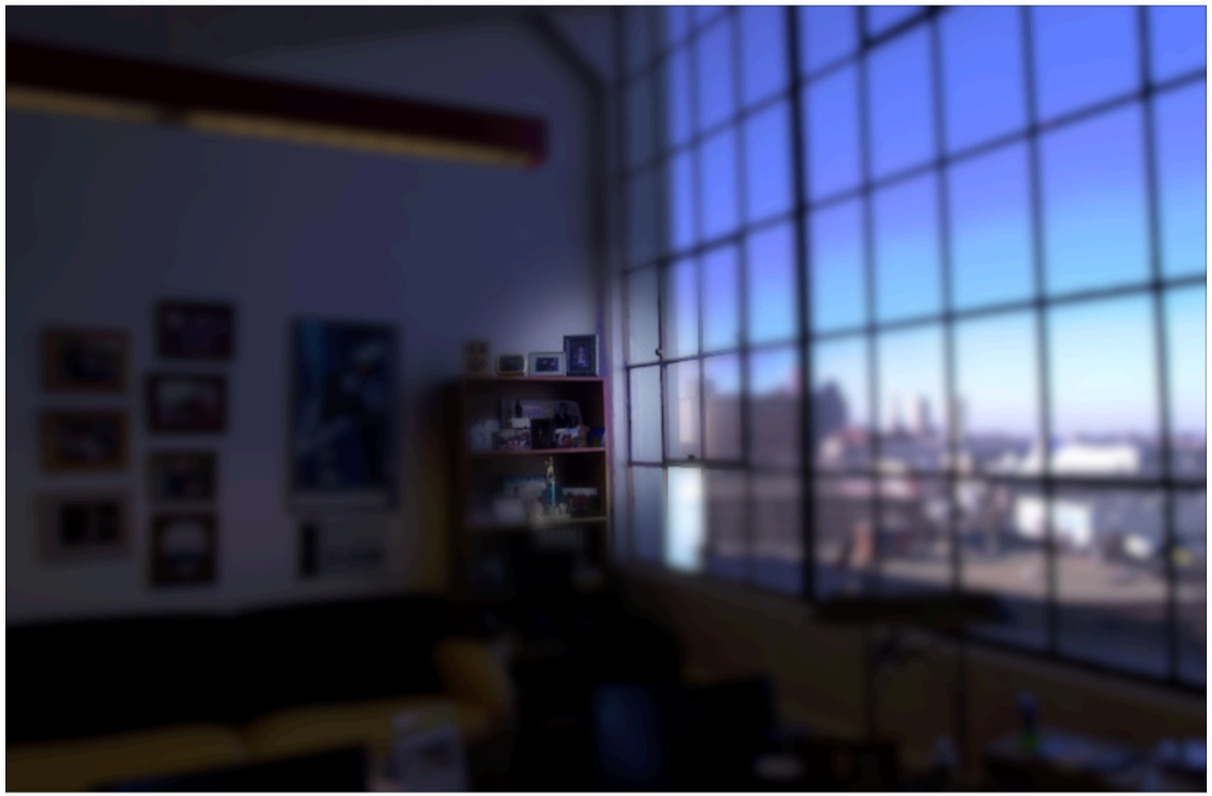


High dynamic range

spotmeter



Demo: One gain per fixation



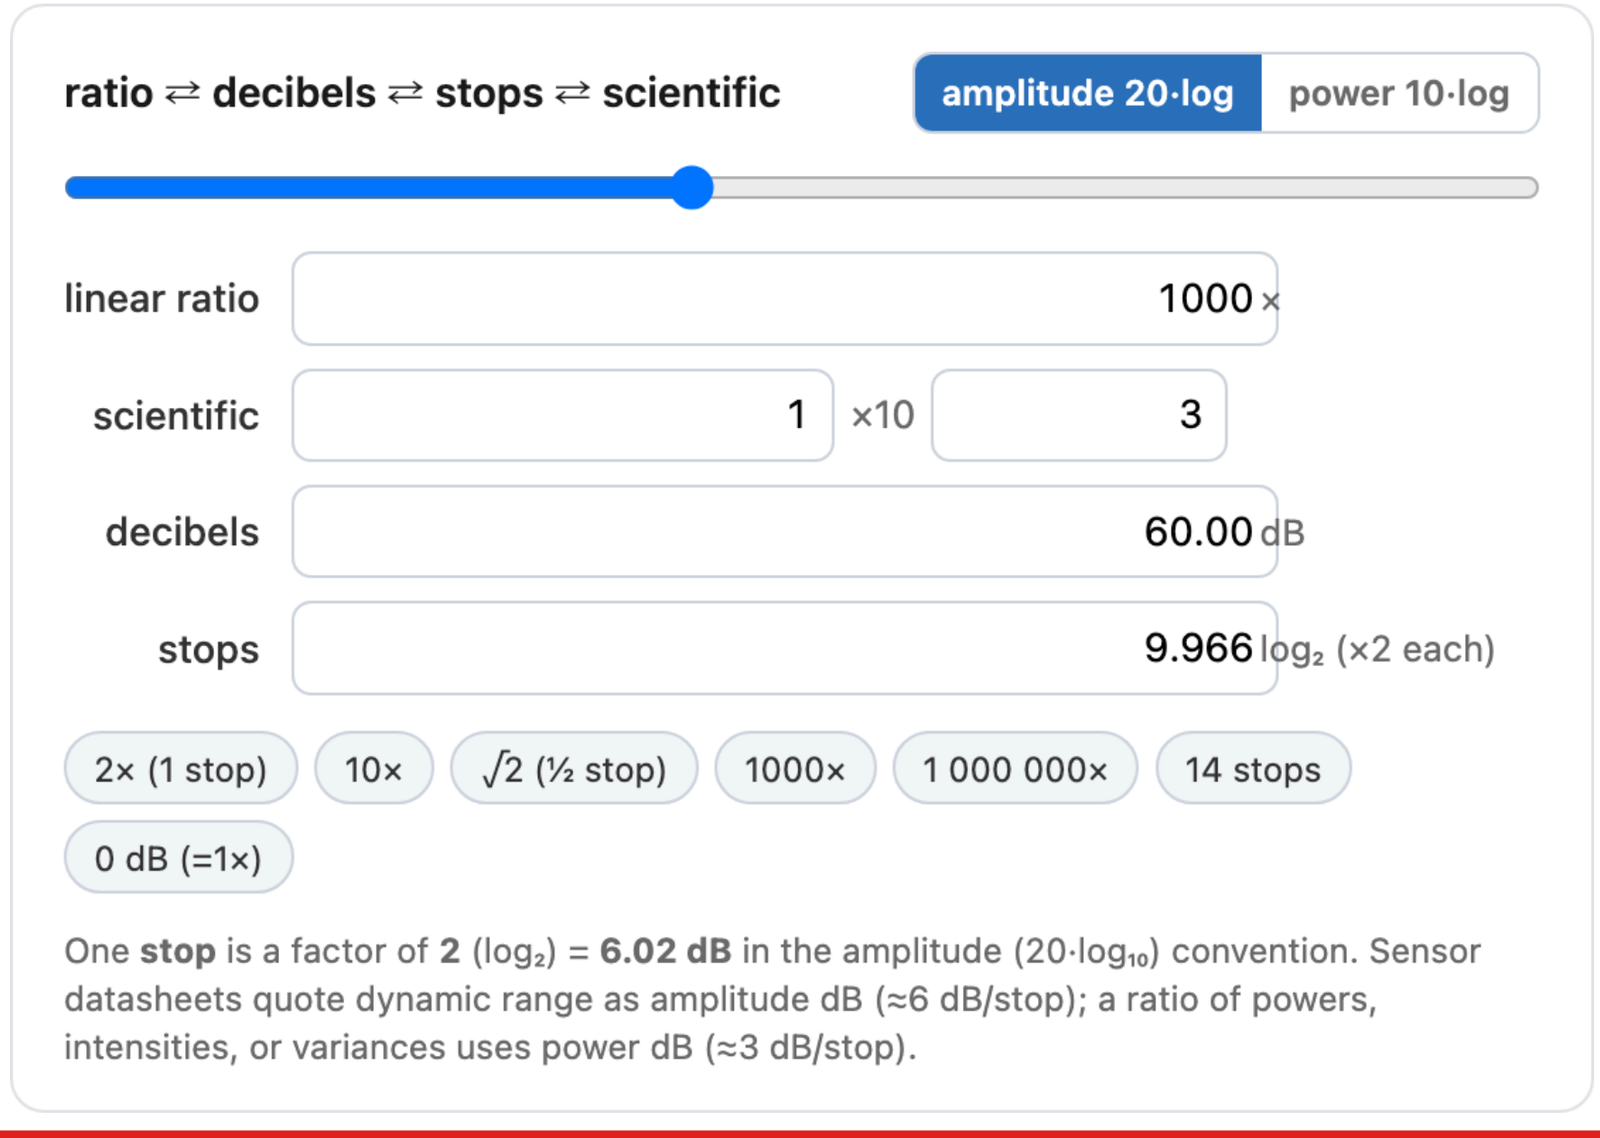
adapted to −0.7 stops relative to the scene's mean

one exposure loses 0% of this patch to white and 12% to black · the patch spans 9.9 stops, the picture 11.4

Saturating: the same multiplier plus a roll-off, which is why the ends survive.

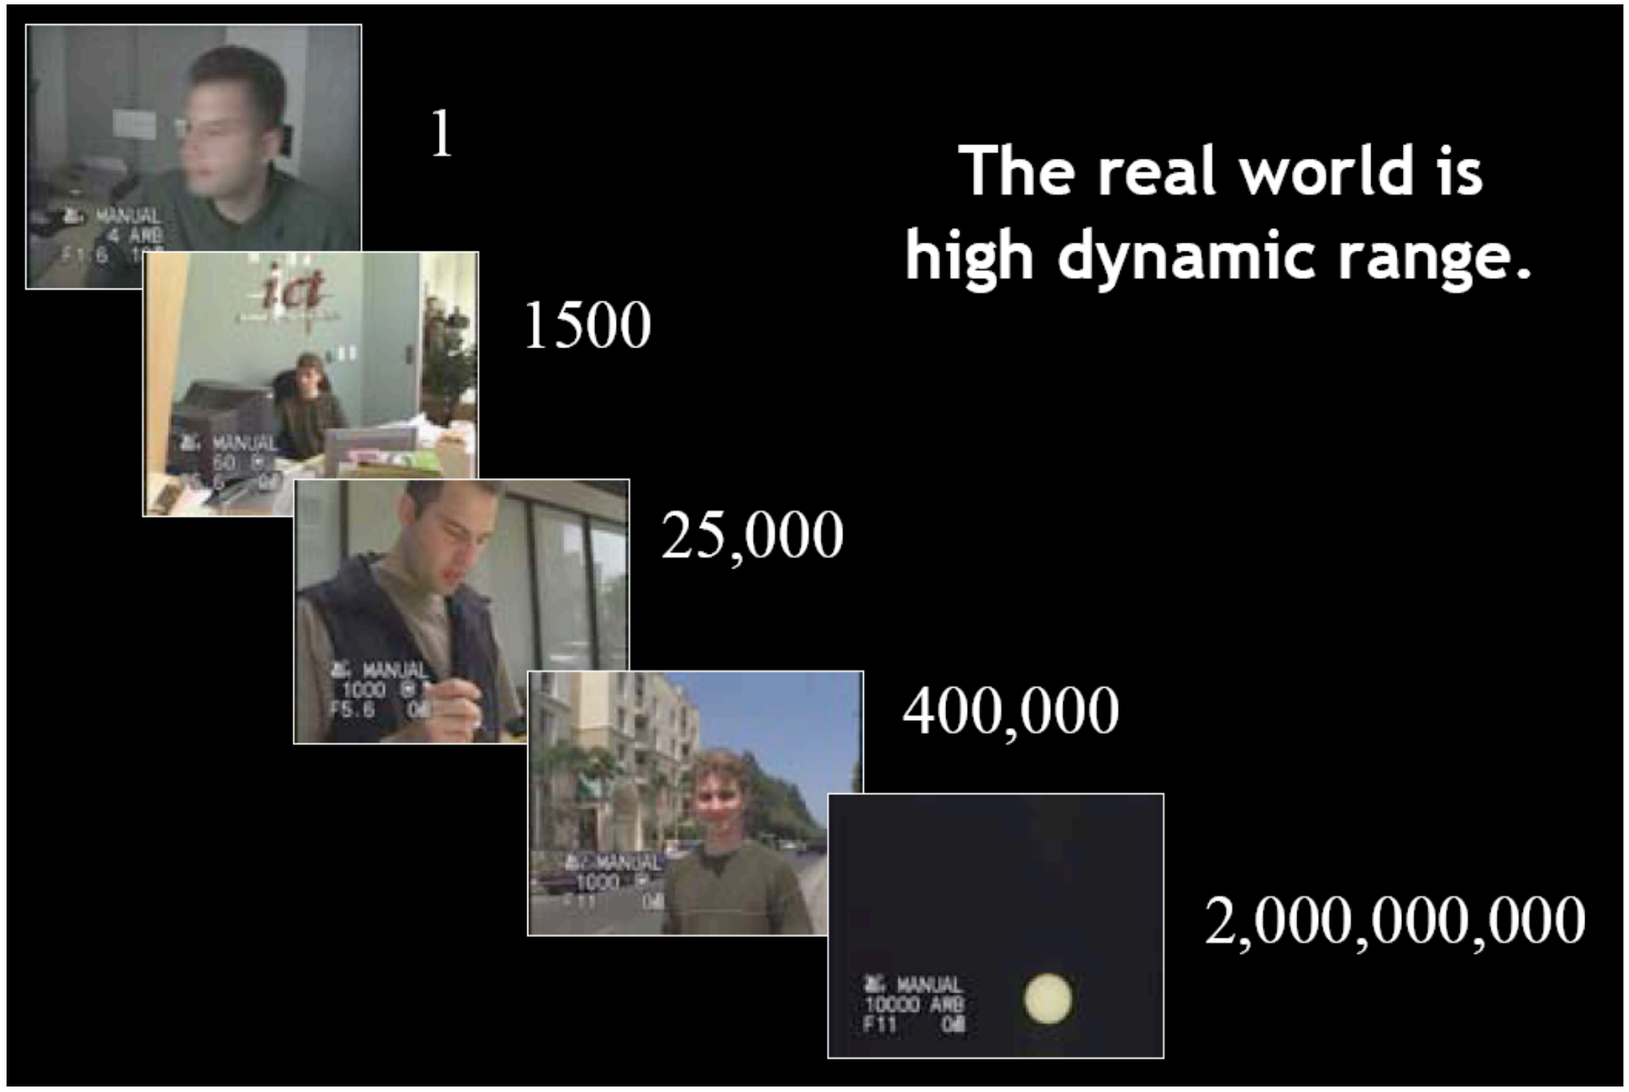
radiance map: office window — Byong Mok Oh

Demo: Ratio units converter



The world is high dynamic range

- **Slide from Paul Debevec**



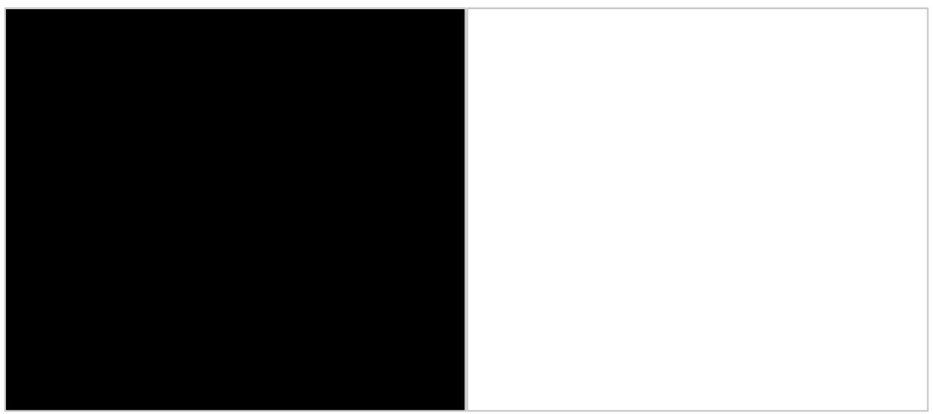
Demo: One frame per rung of light

The range of everyday light levels: each frame auto-exposed to mid-tones, so the light level shows only in the settings and the numbers

Cave / near-dark	City at night	Dim interior	Lit interior	Overcast day	Full sunlight
					
ISO 6400 · 1/20 s · f/2.2	ISO 6400 · 1/20 s · f/5.6	ISO 4000 · 1/1000 s · f/2	ISO 1600 · 1/400 s · f/5.6	ISO 640 · 1/500 s · f/9	ISO 125 · 1/1000 s · f/8
3.3 lux	**21 lux**	**218 lux**	**1,700 lux**	**13,800 lux**	**111,700 lux**
EV_{100} 0.6	EV_{100} 3.3	EV_{100} 6.6	EV_{100} 9.6	EV_{100} 12.6	EV_{100} 15.6
0.19 cd/m² metered	1.2 cd/m² metered	12 cd/m² metered	98 cd/m² metered	791 cd/m² metered	6400 cd/m² metered
$\sim 6 \times 10^3$ photons/px/s	$\sim 4 \times 10^4$ photons/px/s	$\sim 4 \times 10^5$ photons/px/s	$\sim 3 \times 10^6$ photons/px/s	$\sim 2 \times 10^7$ photons/px/s	$\sim 2 \times 10^8$ photons/px/s

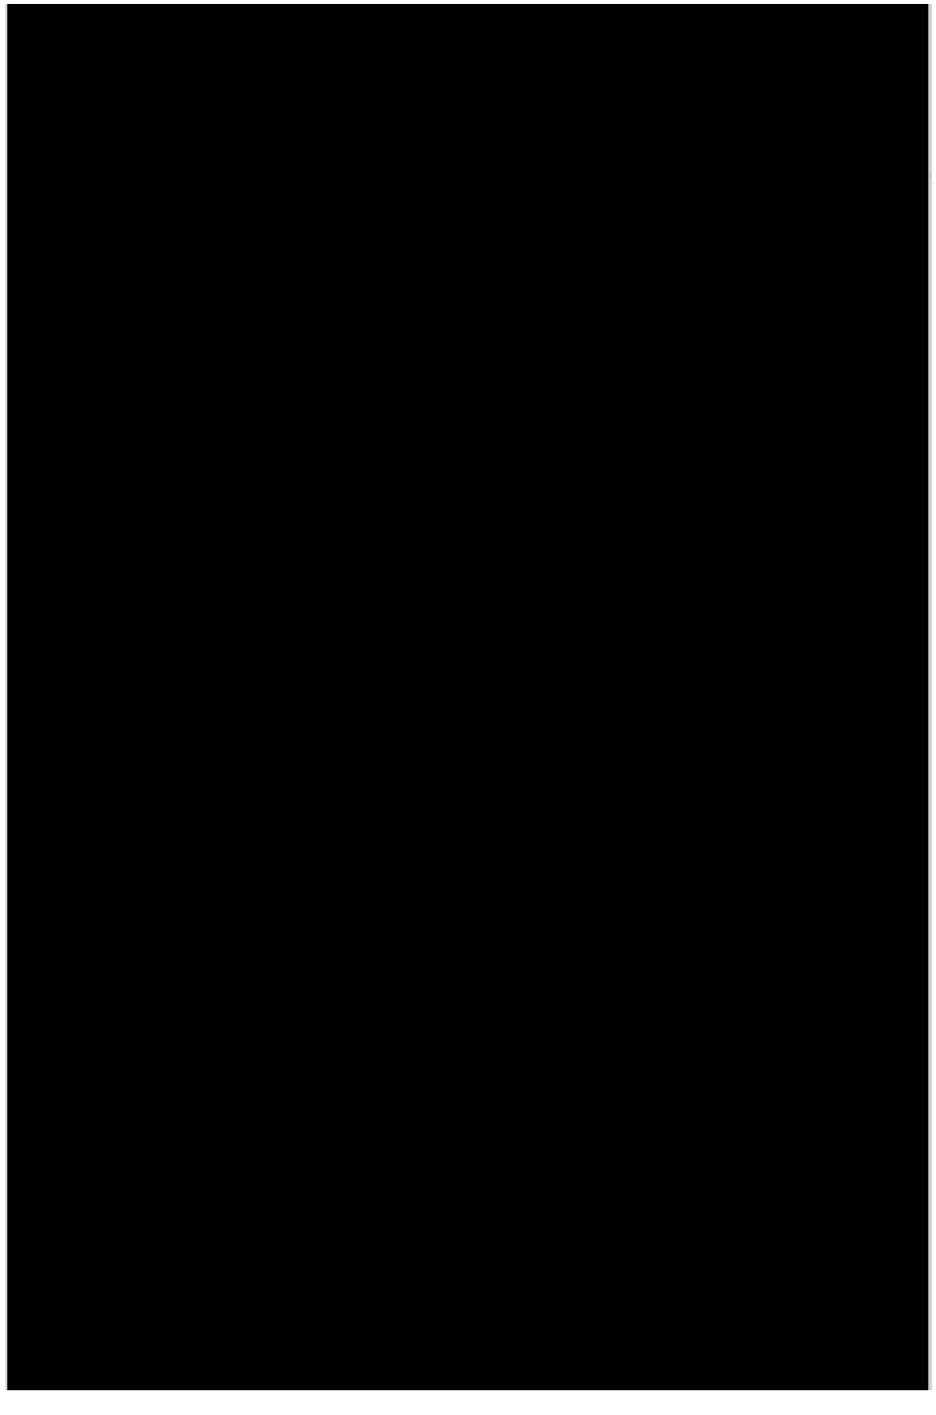
darker → scene light level spans ~15 stops → brighter (top line: the exposure the camera chose; light level as incident lux at 18% reflectance and as EV at ISO 100; photon rate at a fixed f/2, 4µm-pixel reference)

Picture dynamic range: Guess!



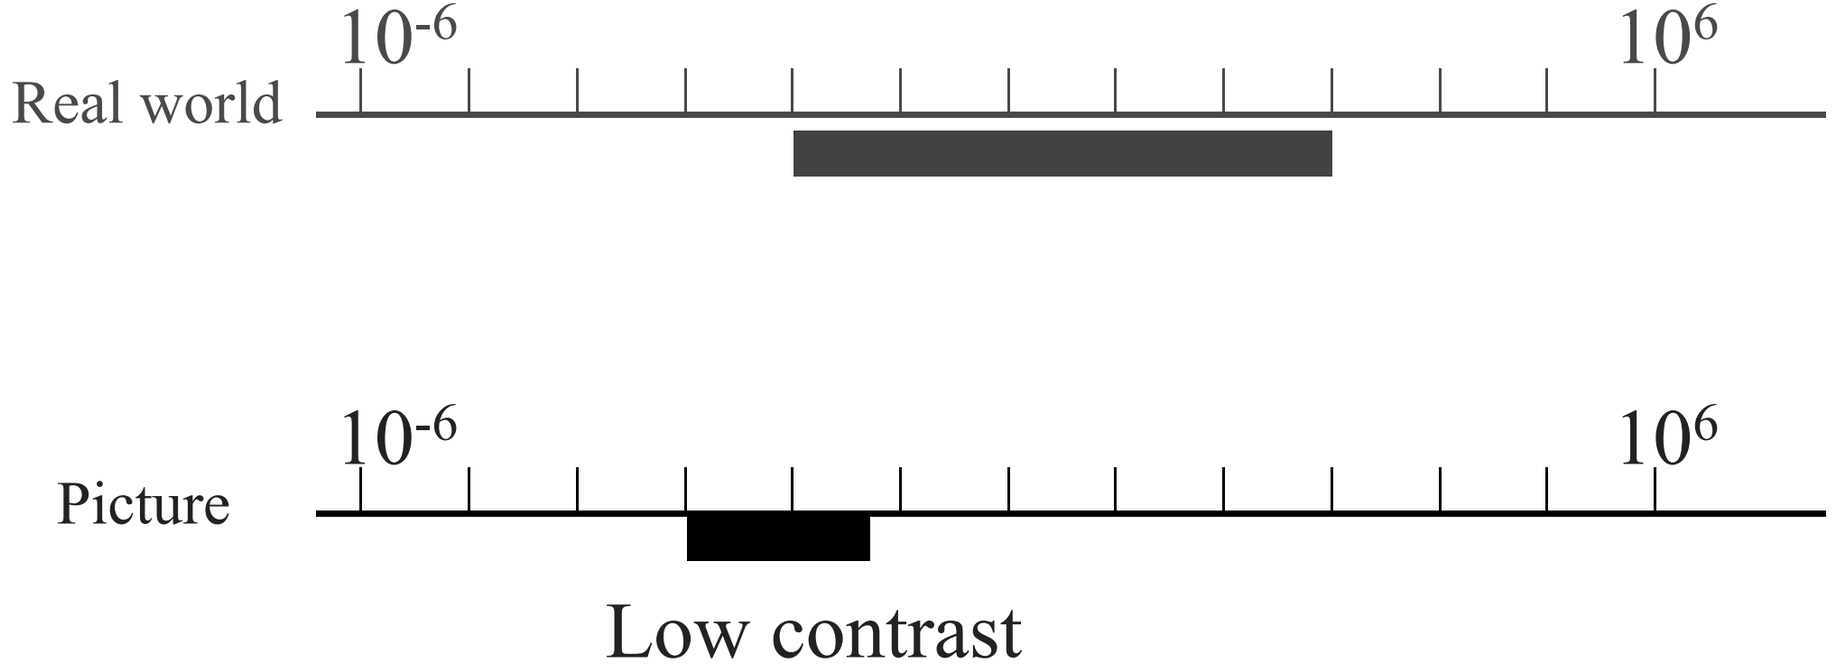






Picture dynamic range

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



The dynamic range problem

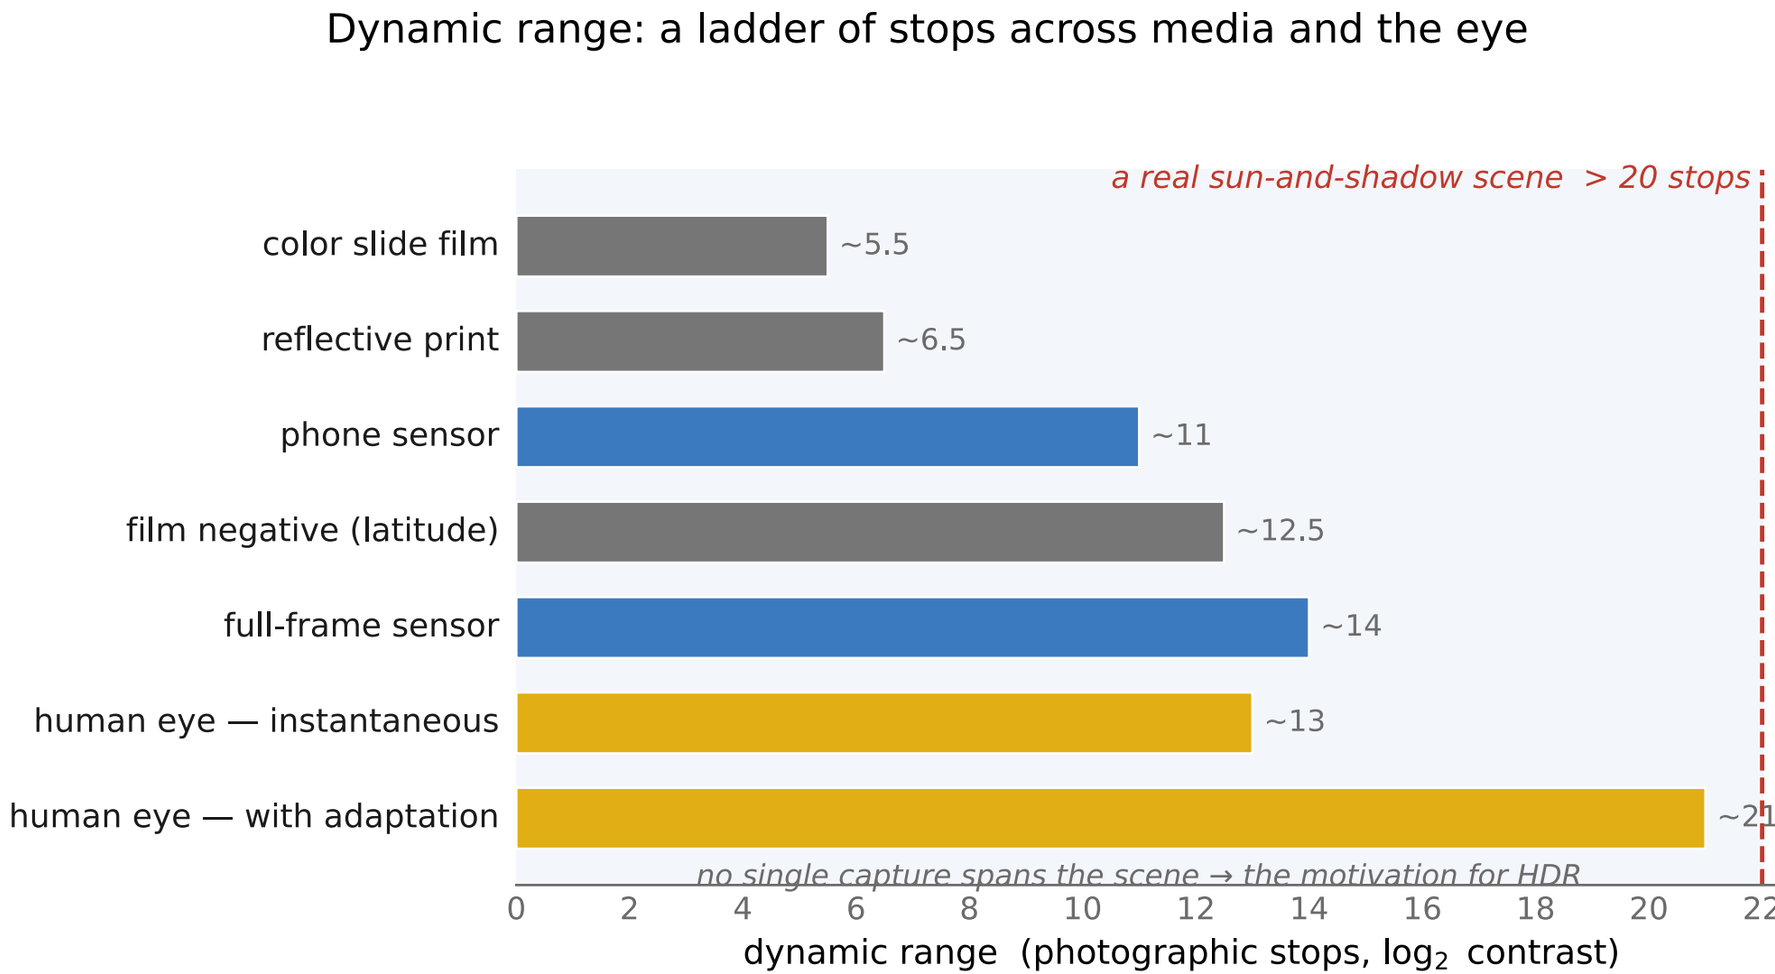
- media (approximate and debatable)

10:1	photographic print (higher for glossy paper)
20:1	artist's paints
200:1	slide film
500:1	negative film
1000:1	LCD display
2000:1	digital SLR (~11 bits)

- challenges
  - choosing which 6-12 bits of the world to include in your photograph (cell phone to professional SLR, respectively)
  - metering the world to help you make this decision, since the world has more dynamic range than any light meter
  - compressing 12 bits into 4 bits for print, or 10 for LCD
    - this is the *tone mapping* problem

© 2010 Marc Levoy

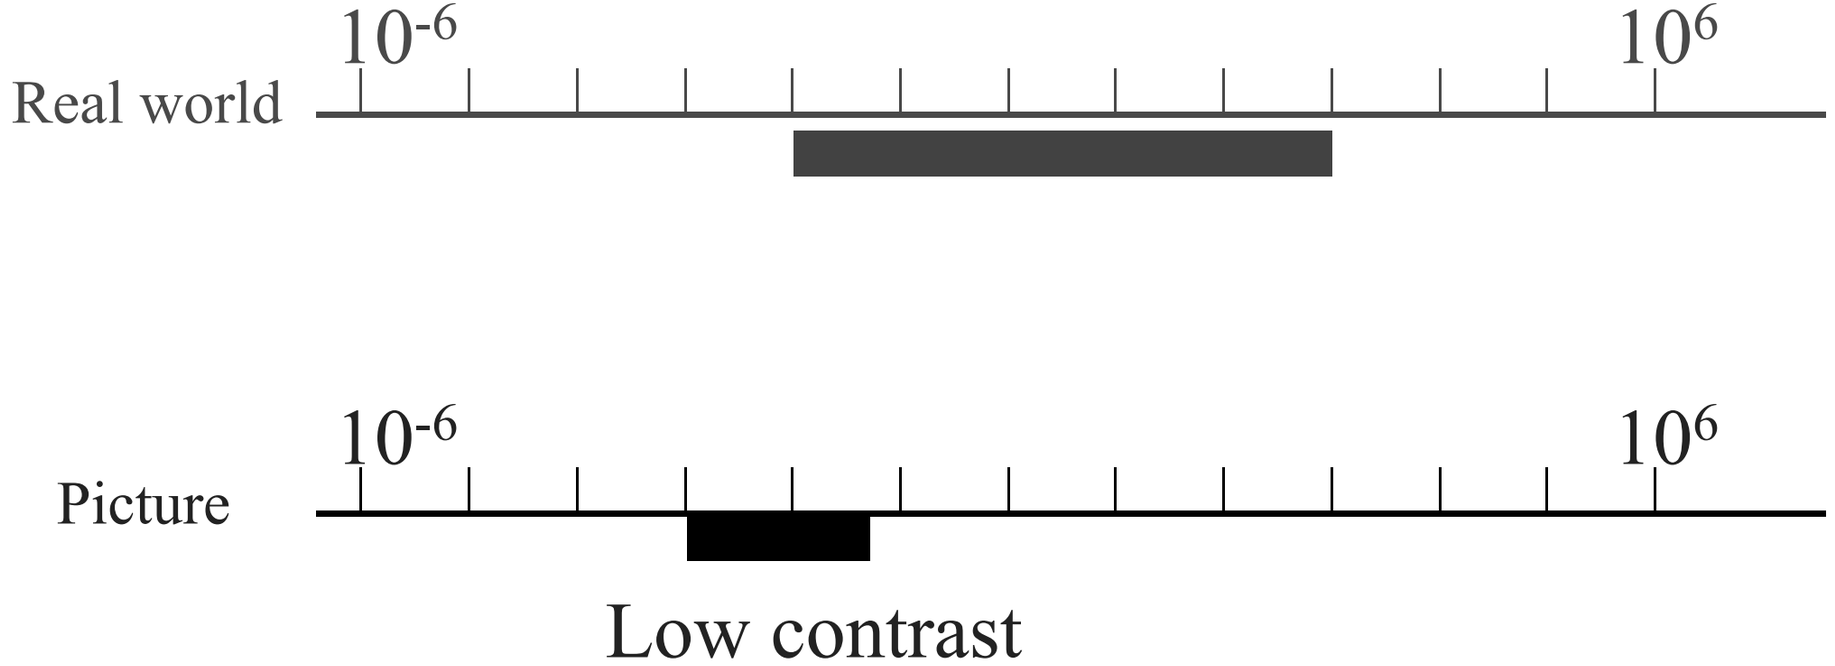
Figure 2.8.15: *Dynamic range across media and eyes, as a ladder of stops: color…



*Dynamic range across media and eyes, as a ladder of stops: color slide (~5–6 stops), reflective print (~6–7), film negative (~12–13 latitude), phone sensor (~10–12)…

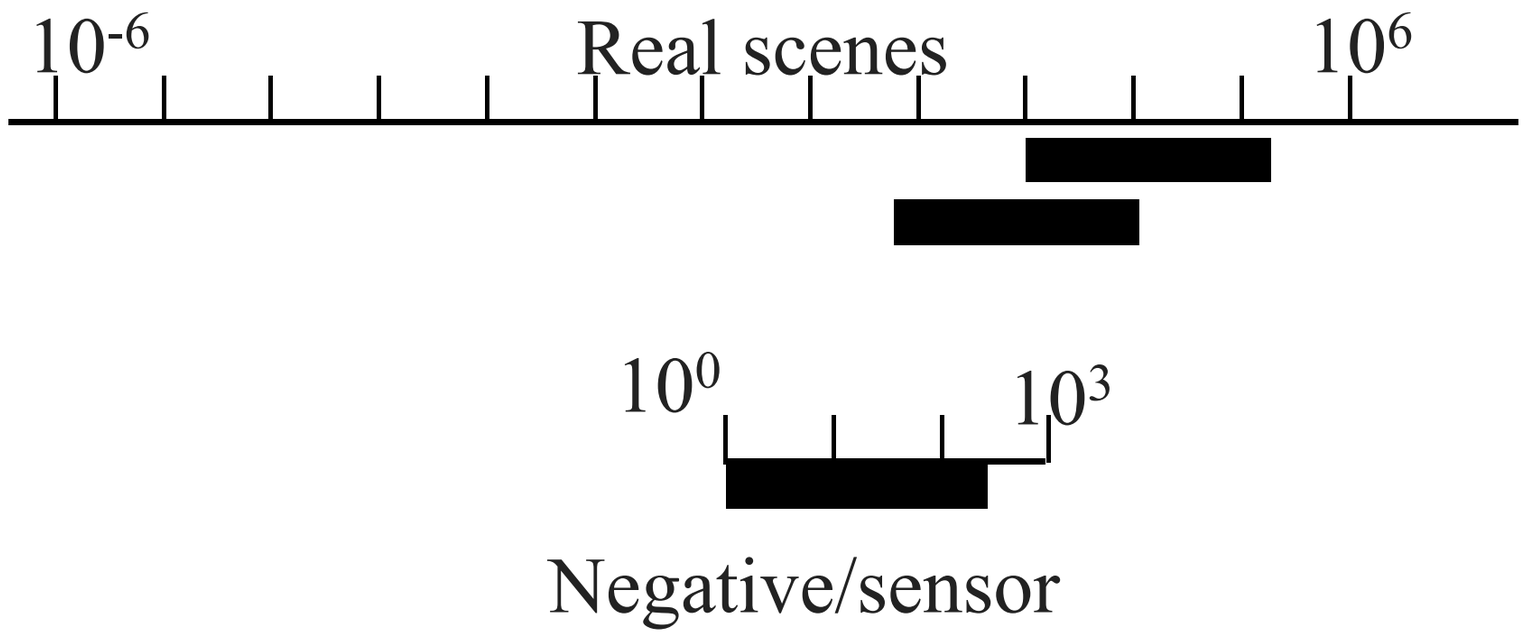
Picture dynamic range

- **Typically 1: 20 or 1:50**
  - Black  is ~ 50x darker than white

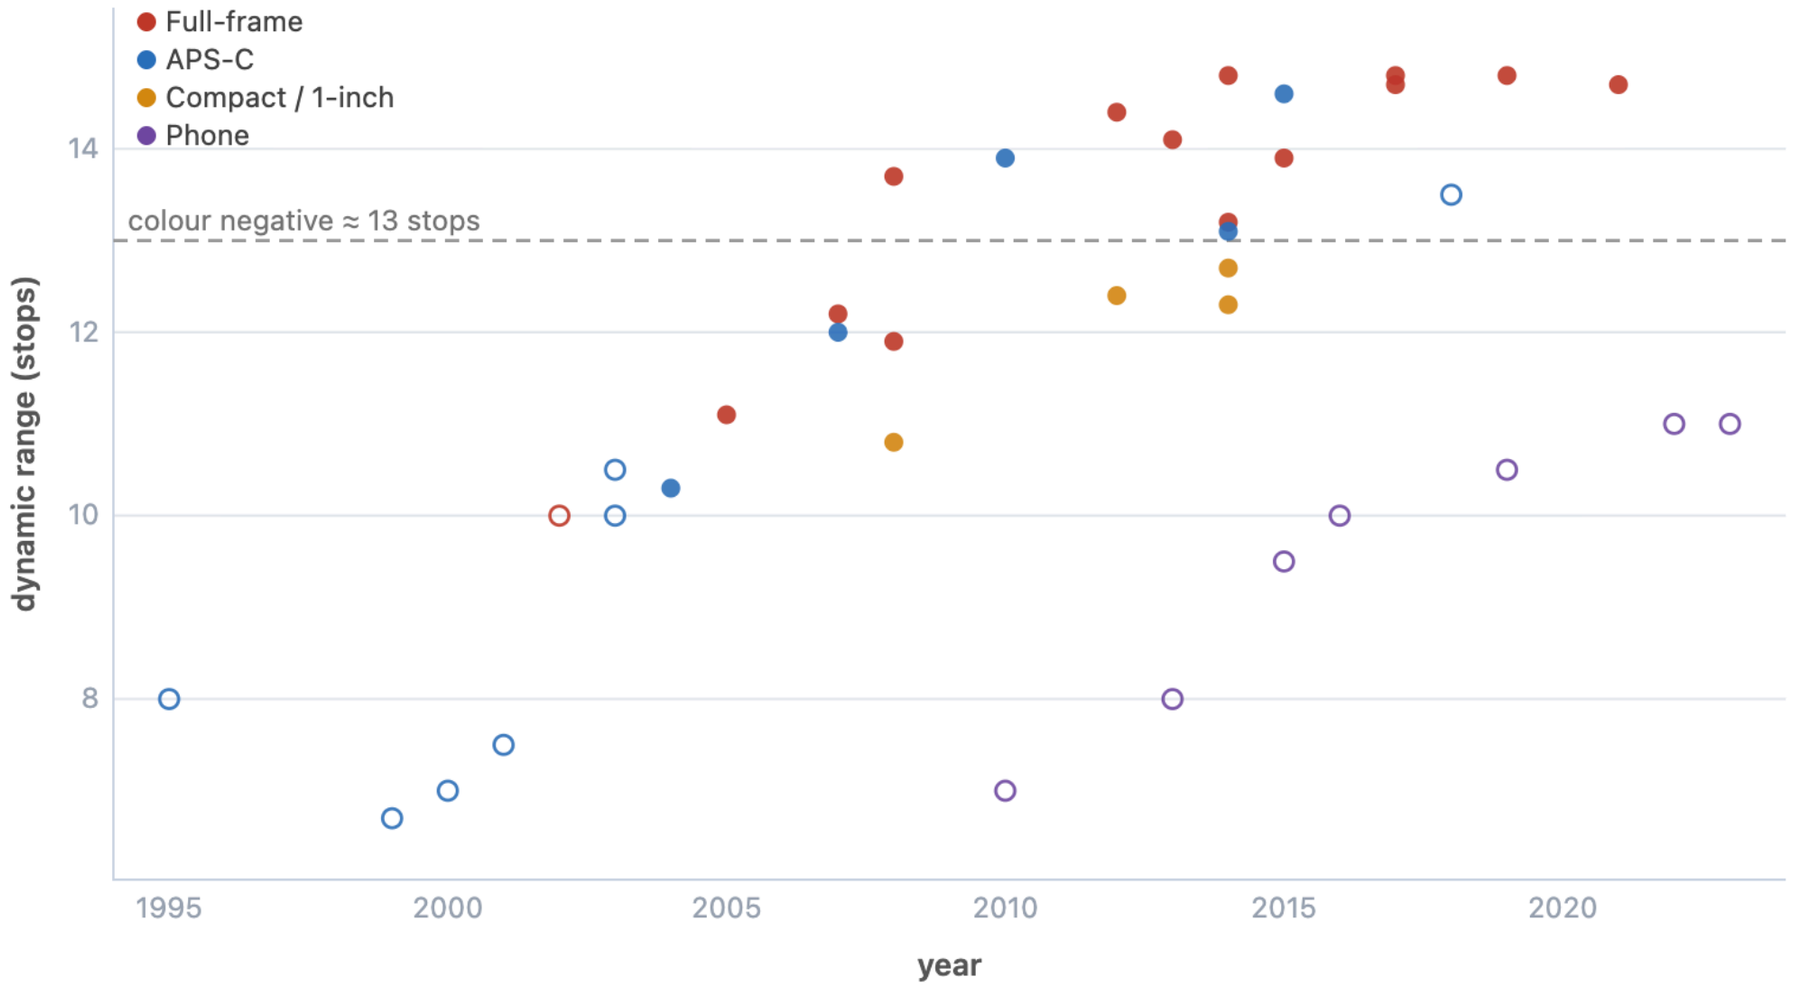


Problem 1: record the information

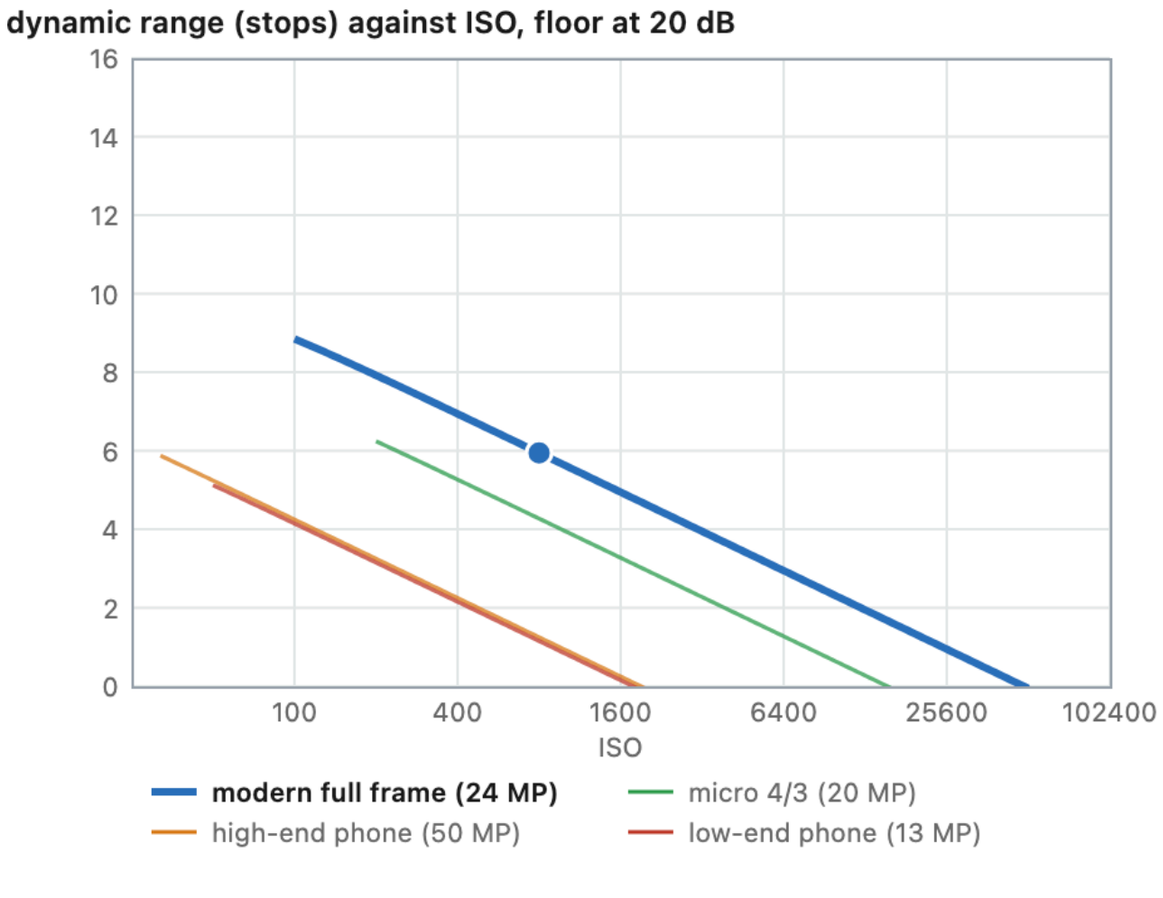
- **The range of illumination levels that we encounter is 10 to 12 orders of magnitudes**
- **Negatives/sensors can record 2 to 3 orders of magnitude**



Demo: Sensor progress



Demo: Dynamic range from the noise model



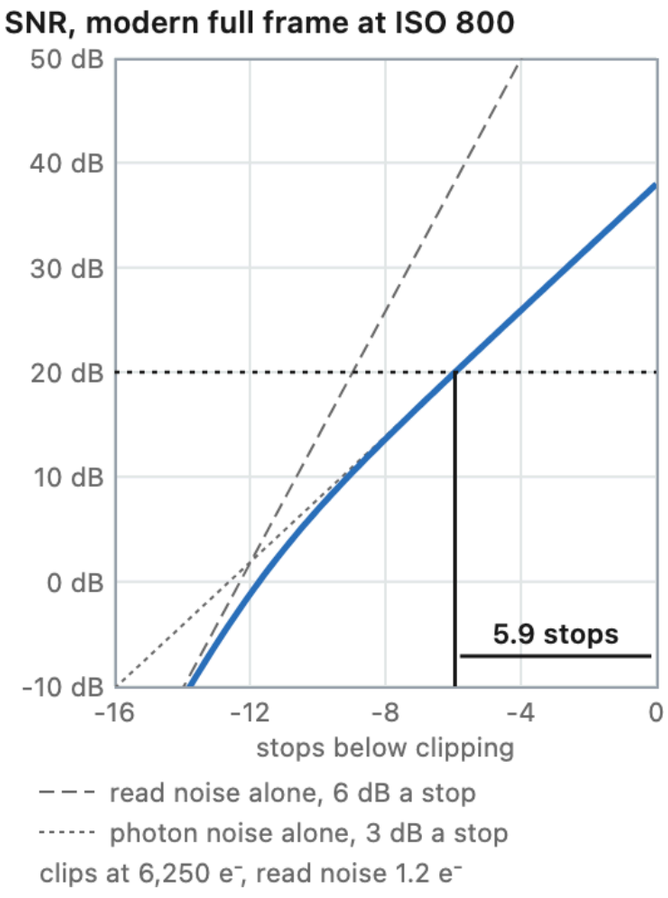
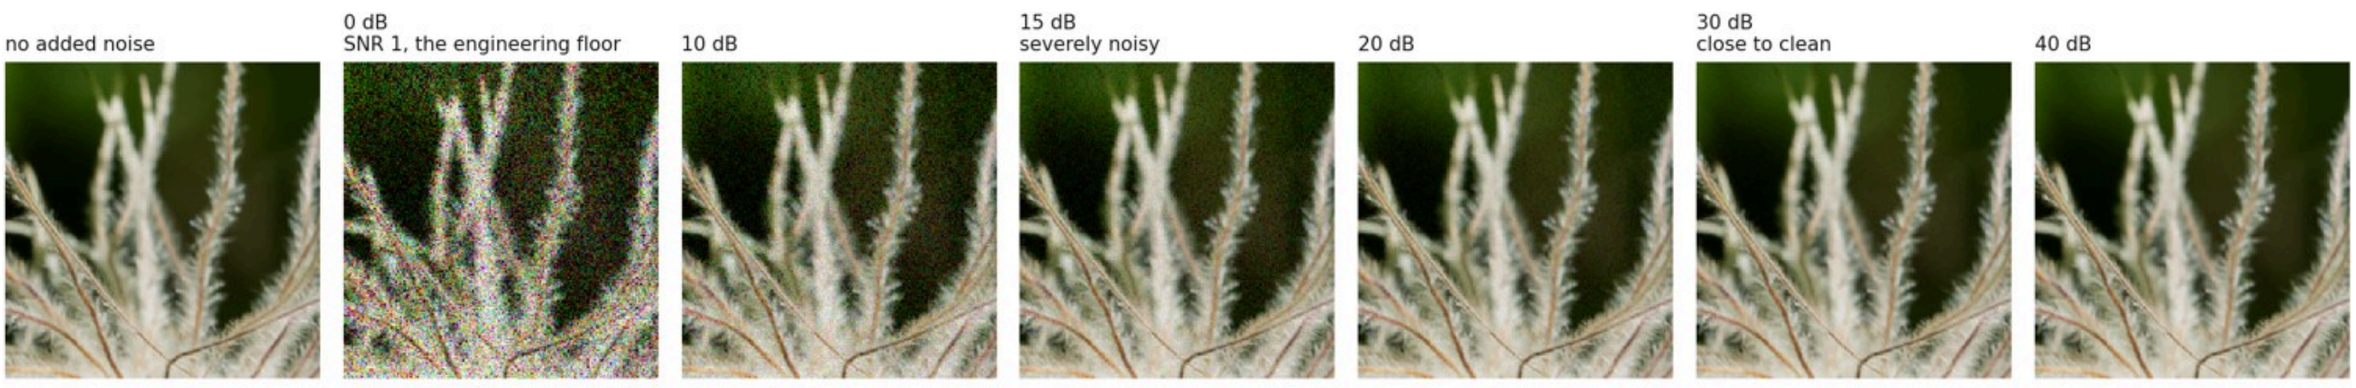


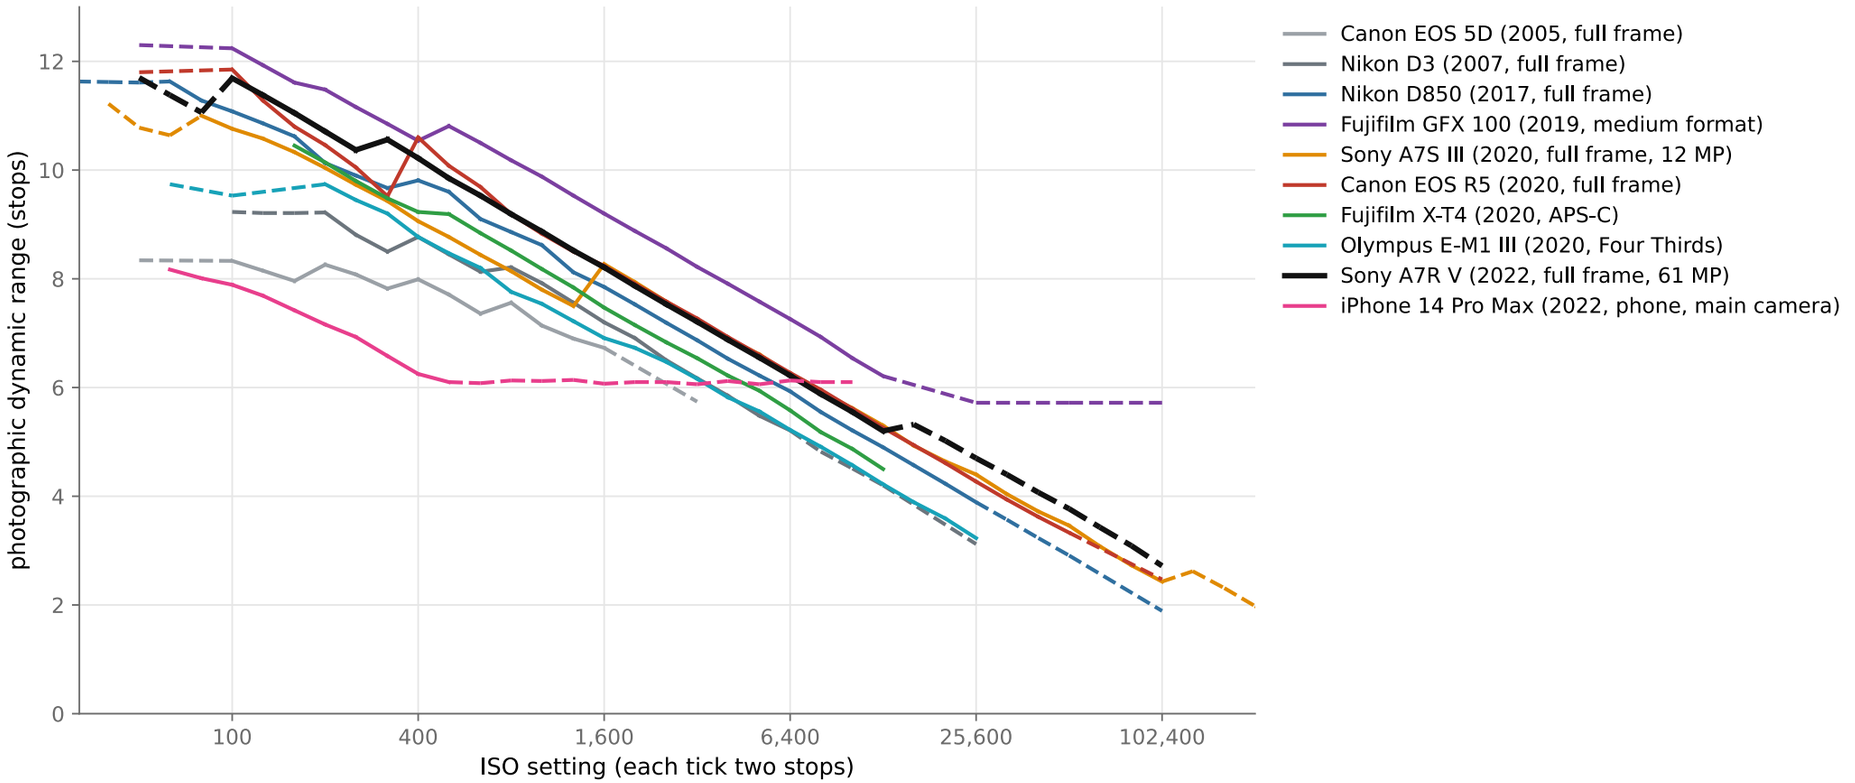
Figure 17.1.1: What a signal-to-noise ratio looks like: a 1:1 crop of a photograph…



1:1 crops; SNR stated at an 18% gray; noise grows with the square root of the signal. Photograph © Frédo Durand.

What a signal-to-noise ratio looks like: a 1:1 crop of a photograph, one image pixel per screen pixel, with noise at 0, 10, 15, 20, 30 and 40 dB, where the SNR in…

Figure 17.1.2: Dynamic range against ISO for ten cameras, from a 2005 full-frame…

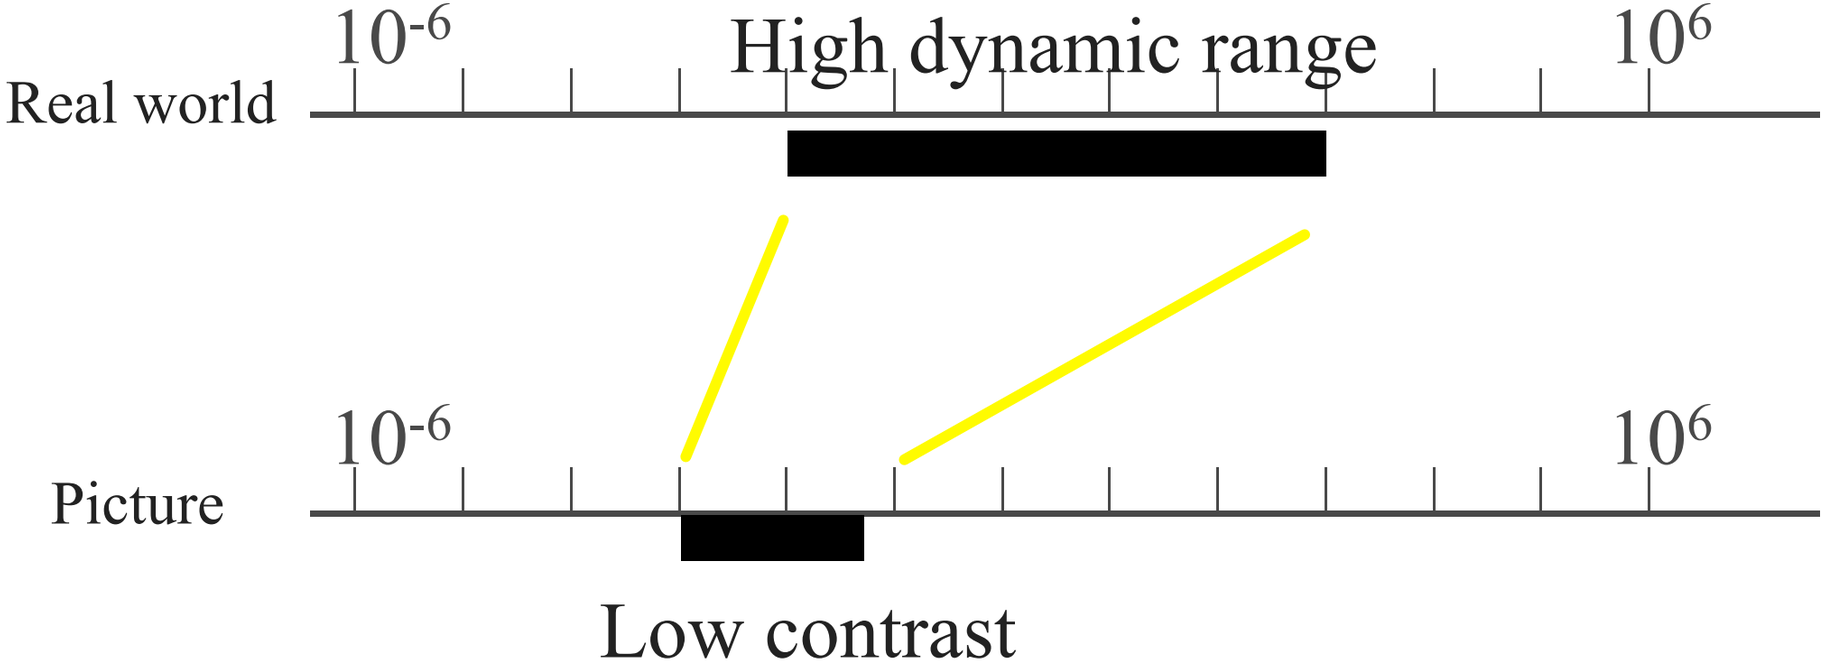


Data: William J. Claff, photonstophotos.net (photographic dynamic range; dashed: extended ISO).

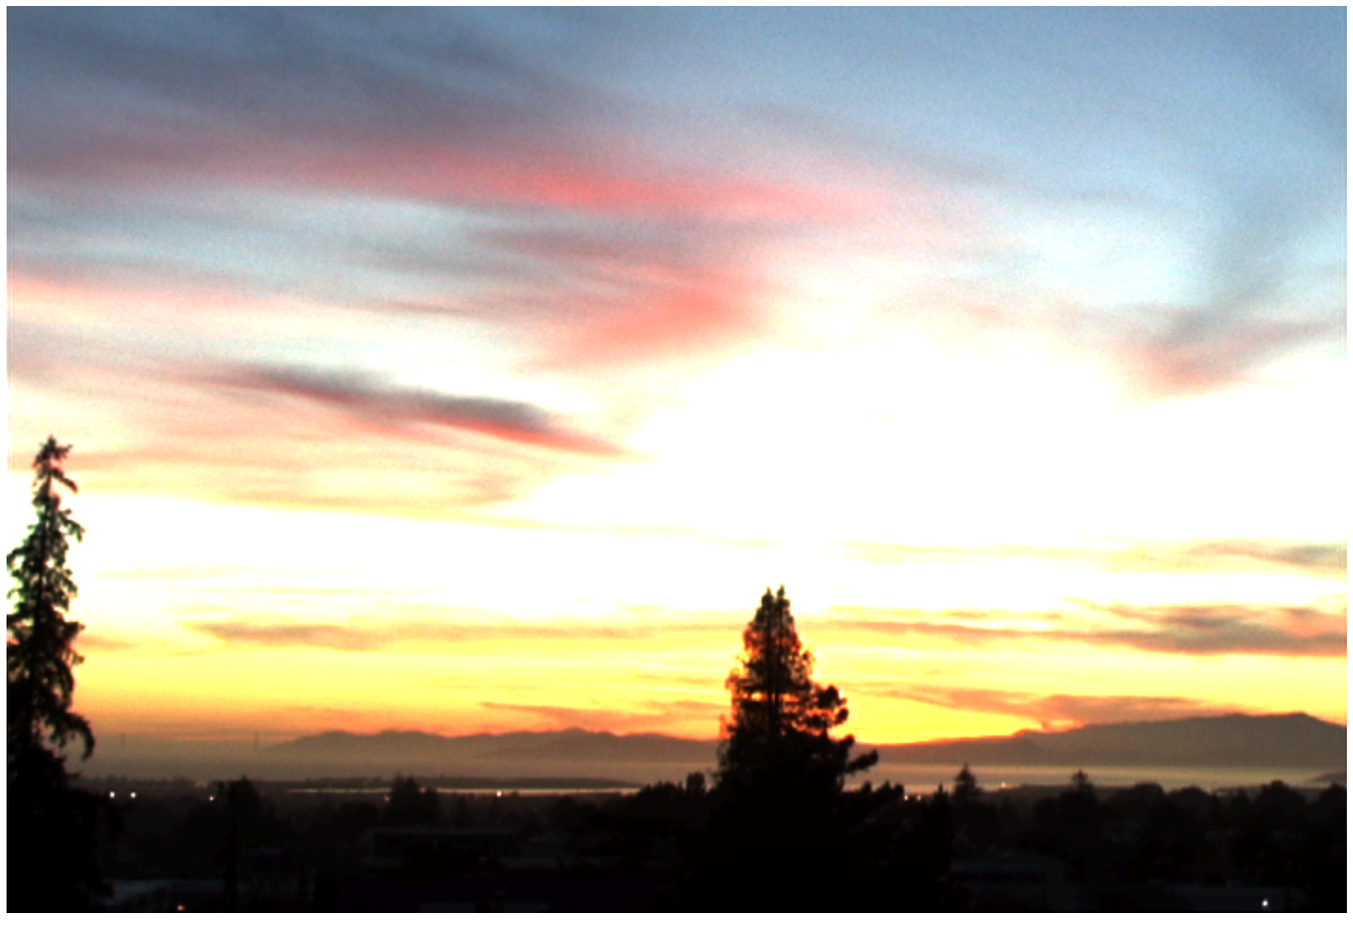
Dynamic range against ISO for ten cameras, from a 2005 full-frame DSLR to a 2022 phone, measured by William Claff as photographic dynamic range: from clipping down to…

Problem 2: Display the information

- **Match limited contrast of the medium**
- **Preserve details**



Without HDR & tone mapping



With HDR & tone mapping



- **http://www.stuckincustoms.com/2011/04/30/hdr-before-after/**





- **http://fstoppers.com/best-technique-for-shooting-interiors-hdr-or-flash**

Best Technique for Shooting Interiors: HDR or Flash?

Have you ever tried to shoot an interior photograph and have it look like the shots in magazines or high end property brochures? If so then you probably know there are two routes to go: HDR or Flash. Photographer **Dom Bower** recently made a video showing the differences in both techniques and how you can combine them both to create a sort of hybrid image. Keep in mind that Dom is only using one single speedlight directly above the camera. Many of the amazing images you see for high end hotels and expensive properties often have dozens of light sources accenting very specific elements in the image. What techniques have you guys used in your interior photos? If you have examples, feel free to post your images in the comments below and check out Dom's final photos in the **full post**.

32 Tweet 71 Share



Can be extreme

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



Not always cheesy



By Alexandre Buisse

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

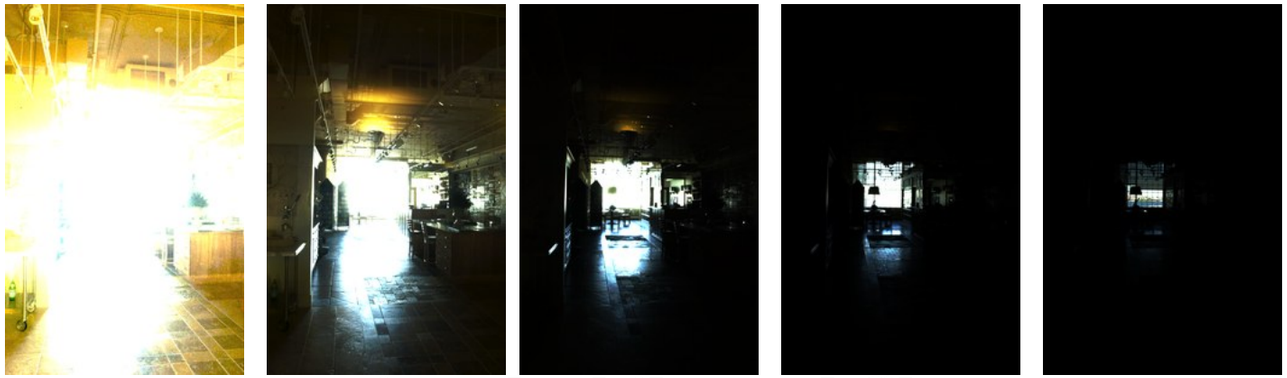
Not always cheesy



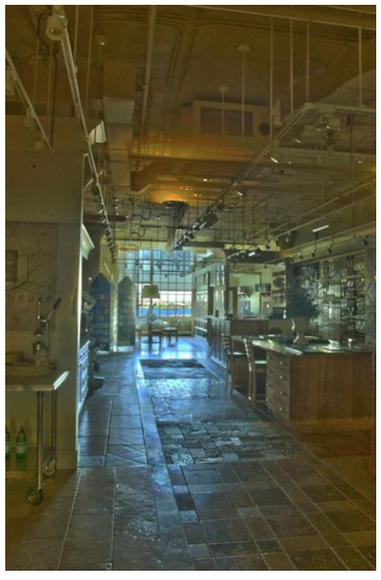
By Alexandre Buisse
http://luminous-landscape.com/essays/hdr-plea.shtml

Plan

- **Multiple-exposure High-Dynamic-Range imaging**



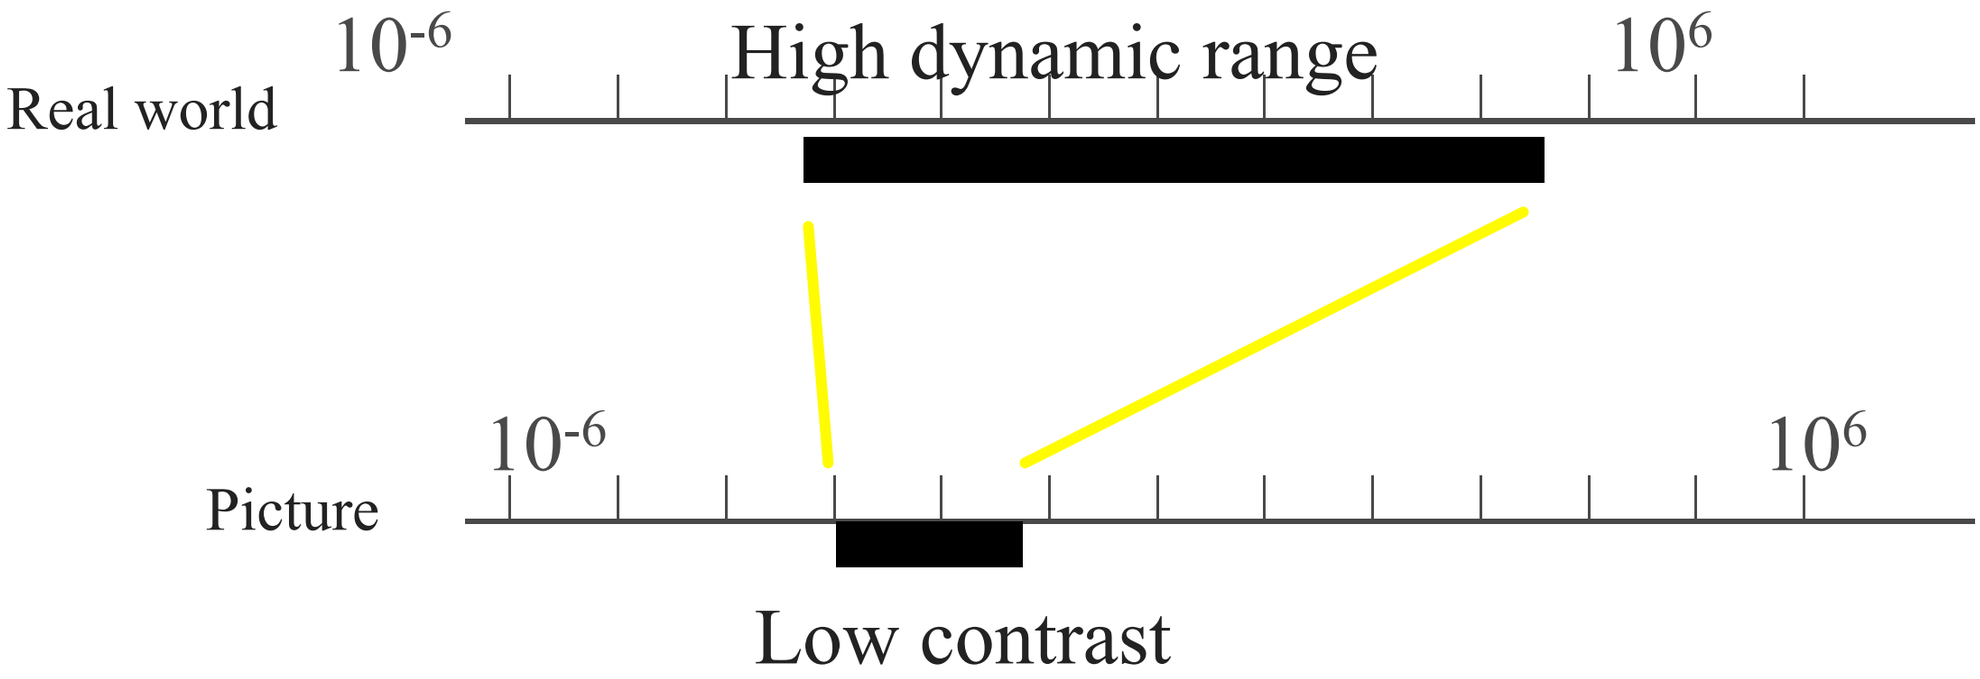
- **Tone mapping using the bilateral filter**



Multiple exposures

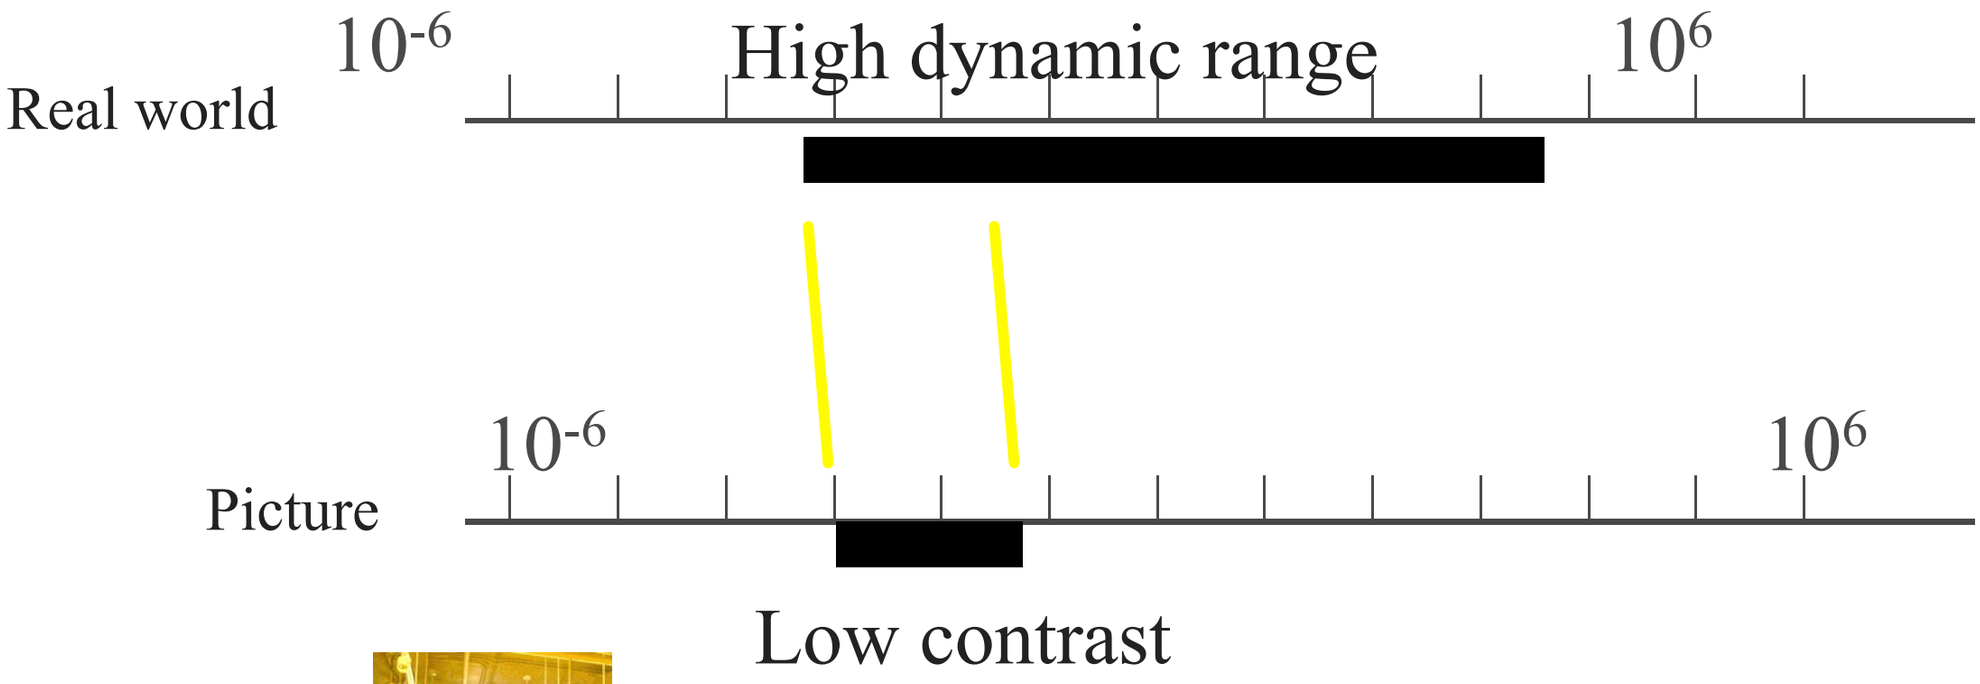
Multiple exposure photography

- **Sequentially measure all segments of the range**



Multiple exposure photography

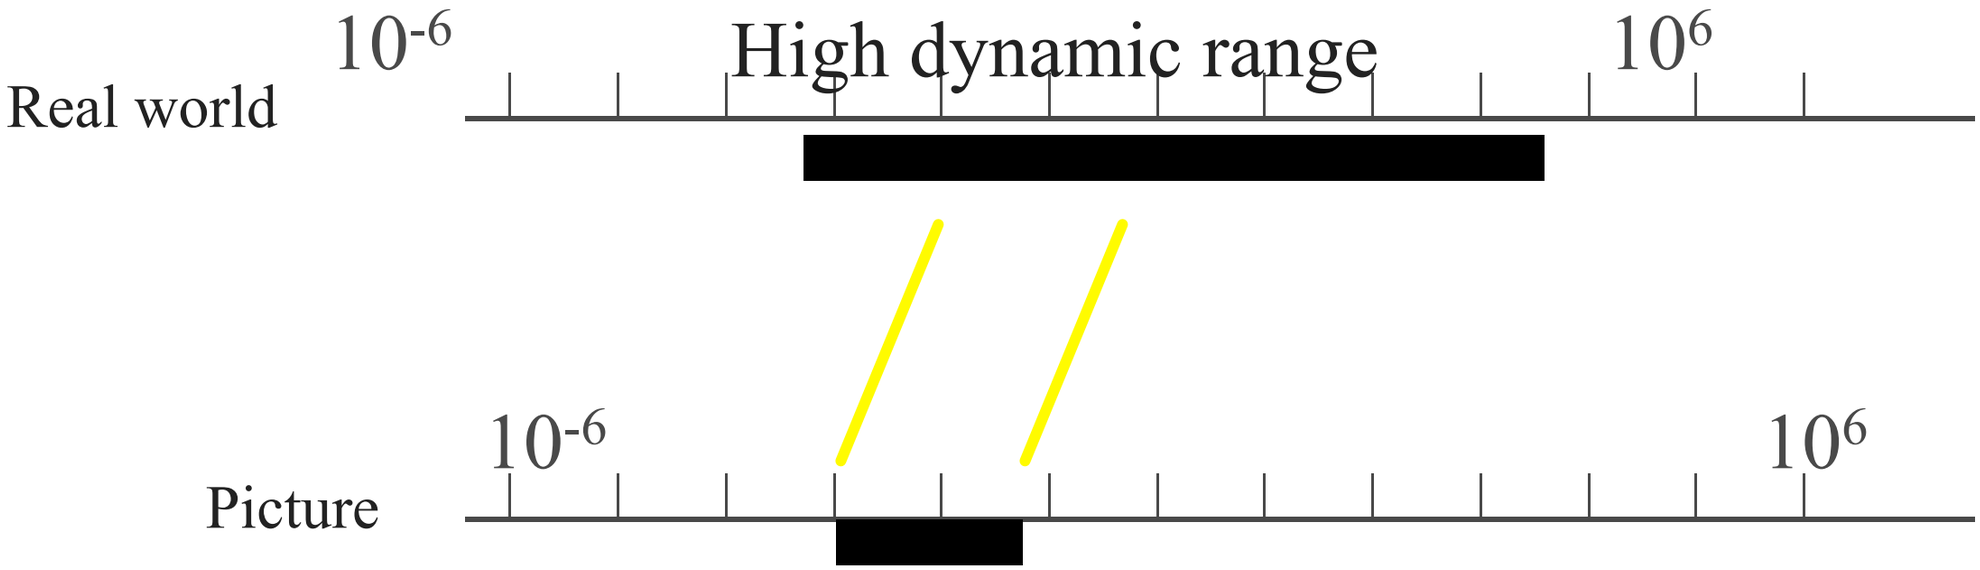
- **Sequentially measure all segments of the range**

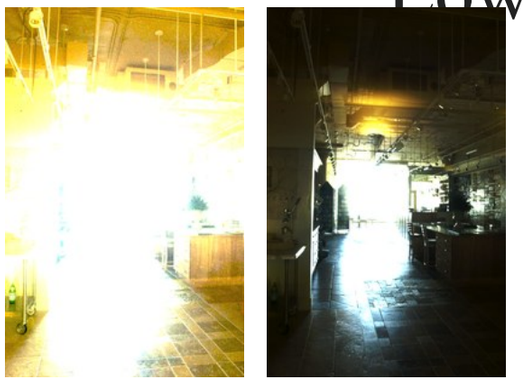




Multiple exposure photography

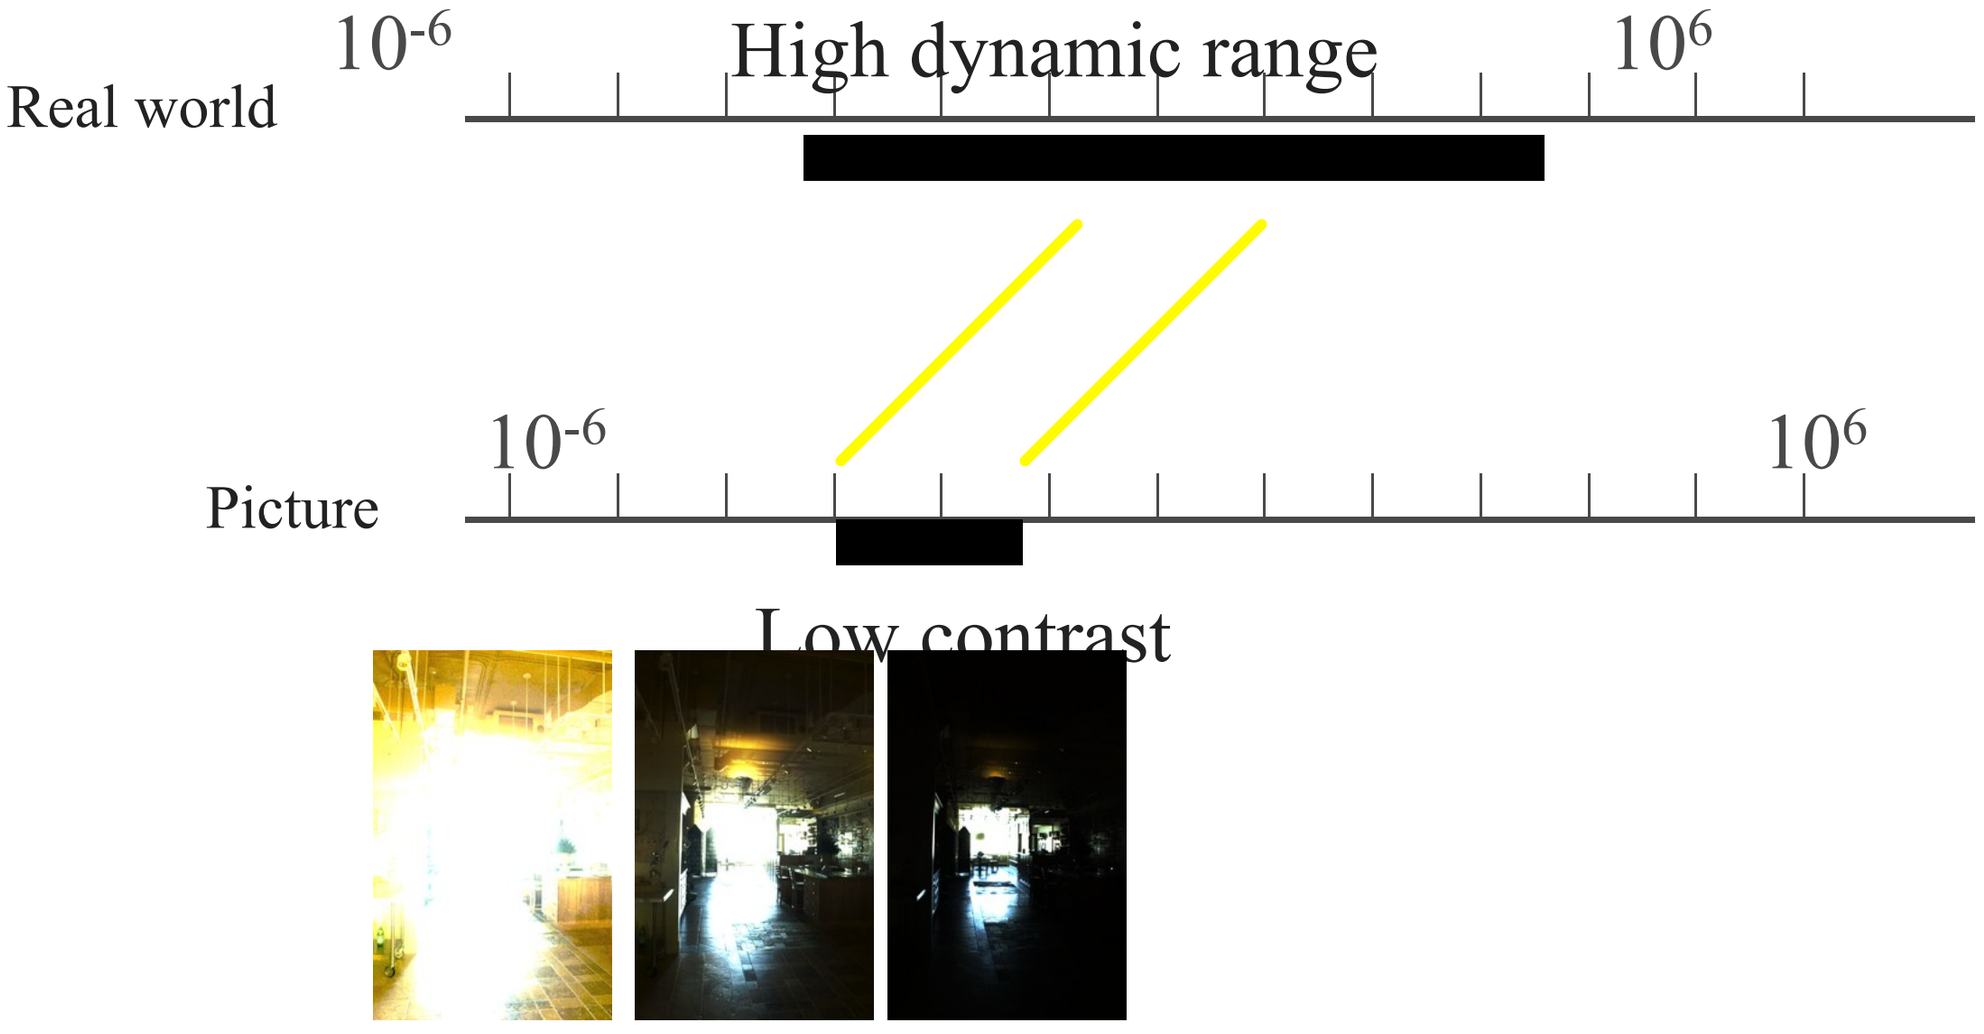
- **Sequentially measure all segments of the range**





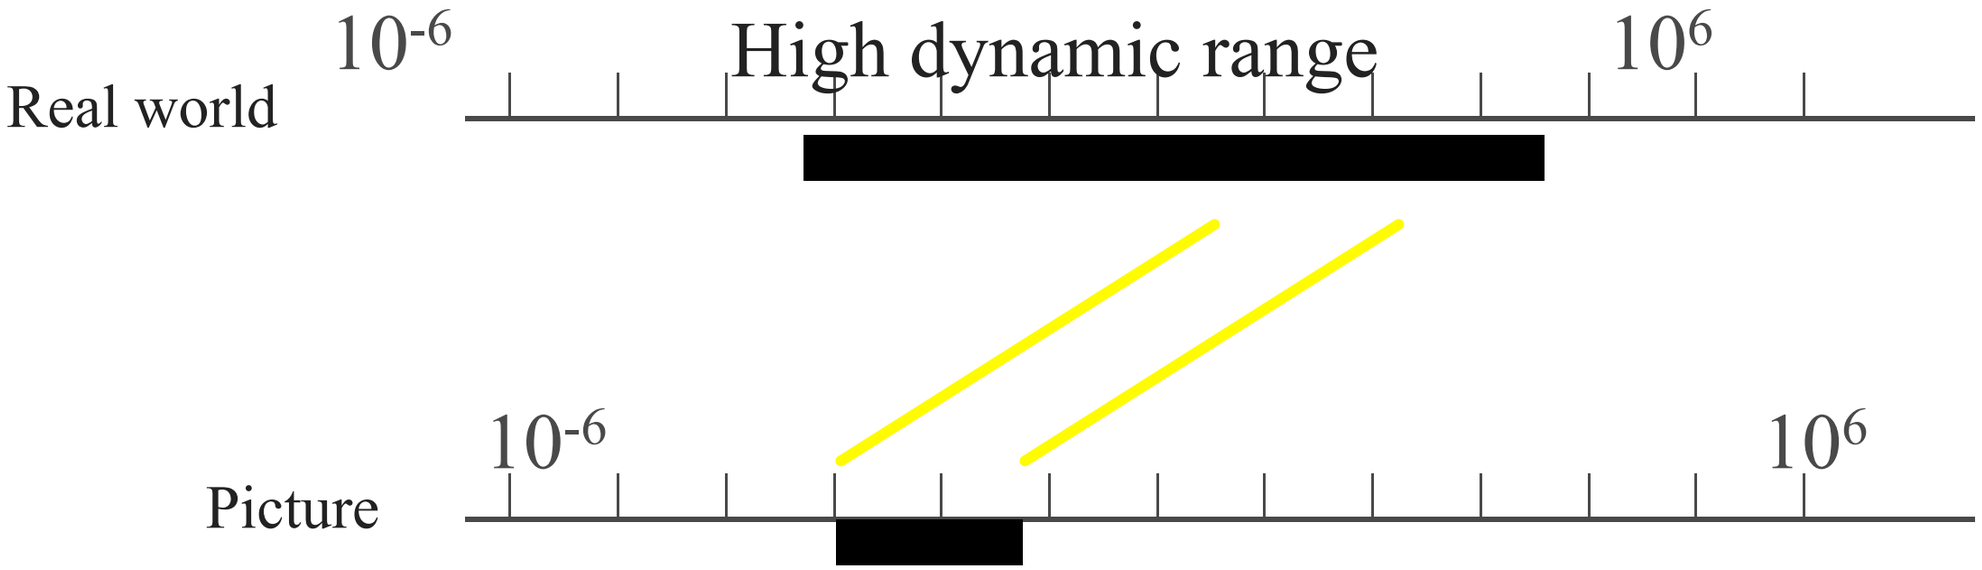
Multiple exposure photography

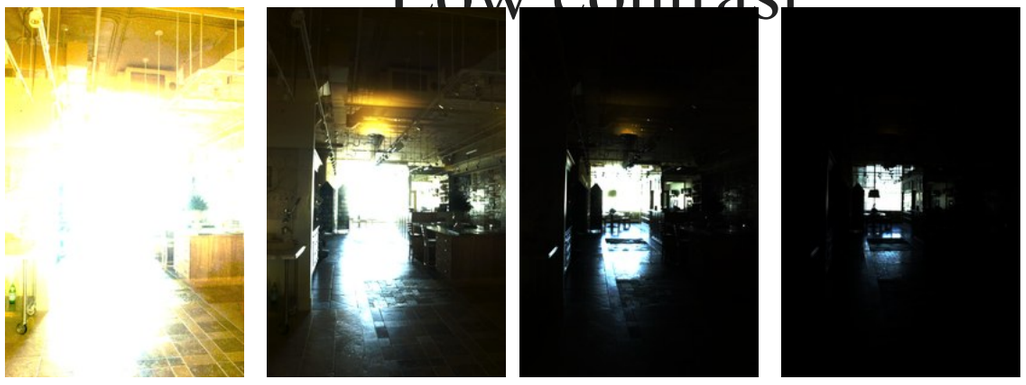
- **Sequentially measure all segments of the range**



Multiple exposure photography

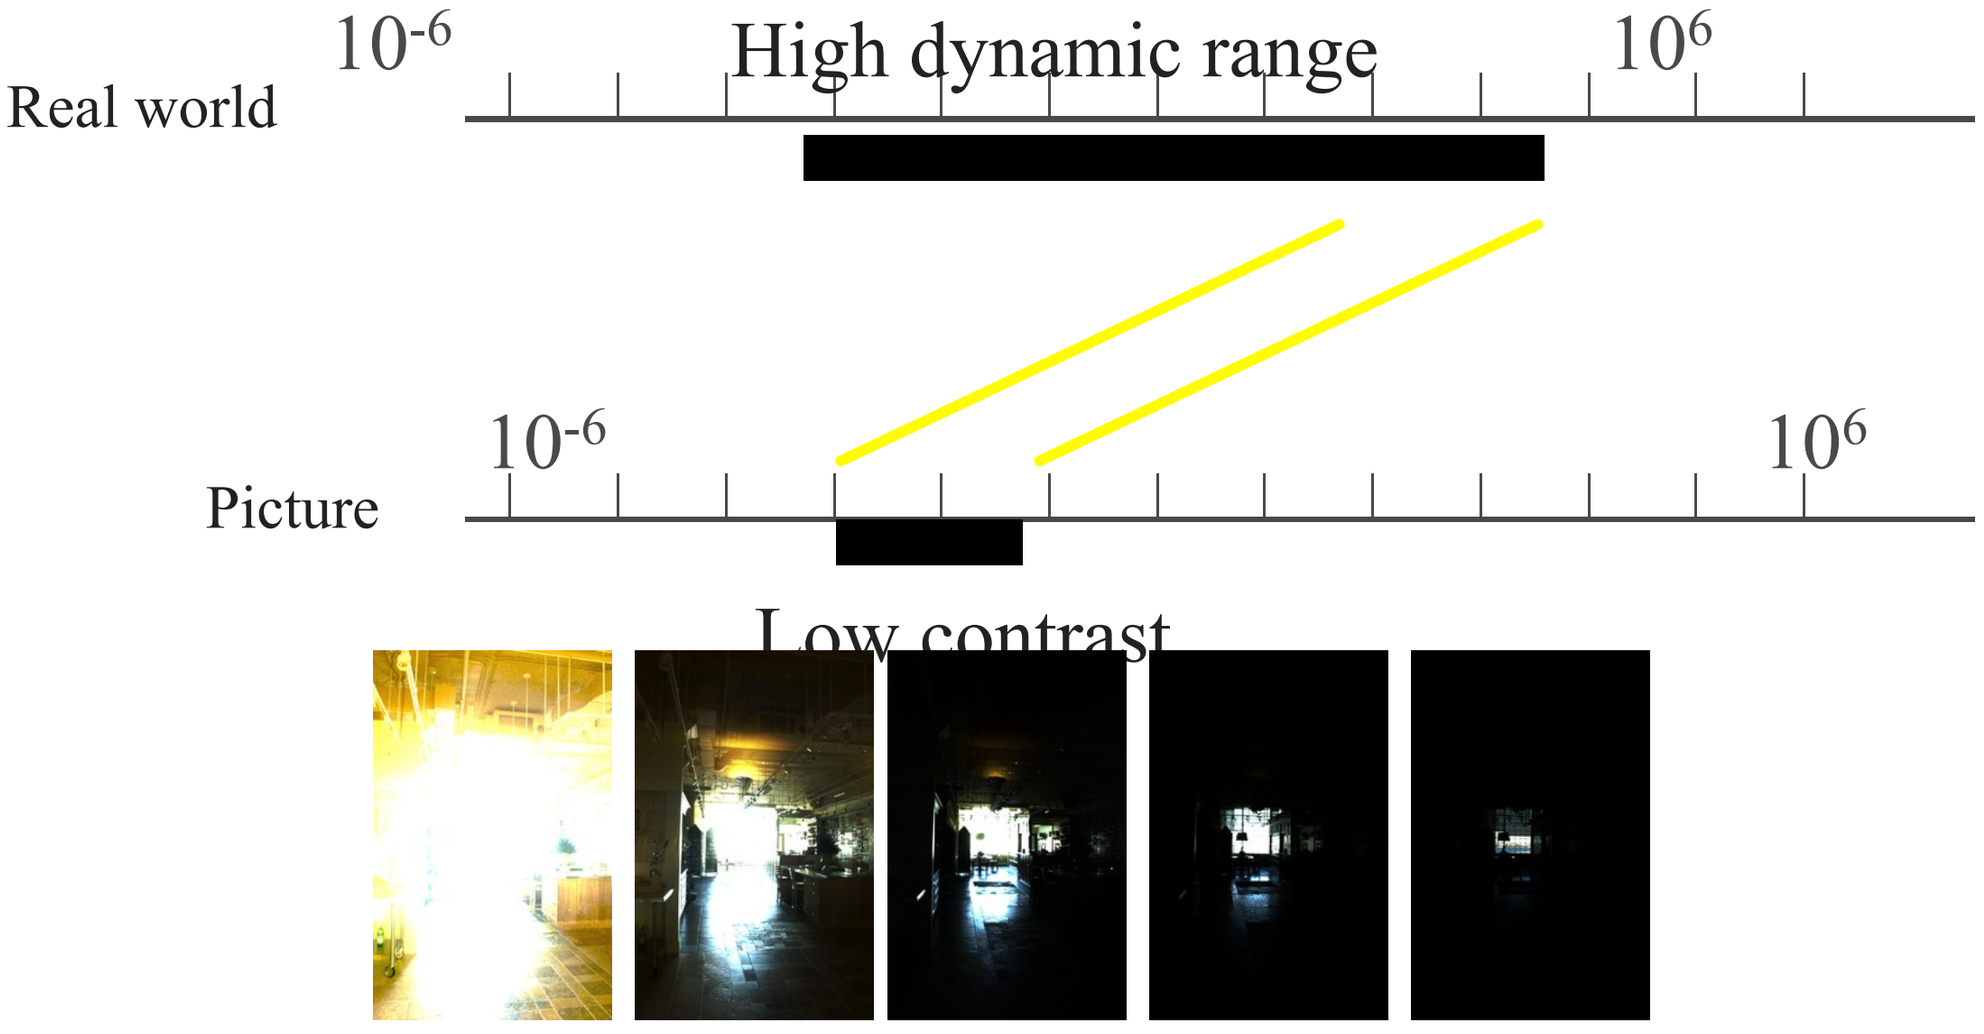
- **Sequentially measure all segments of the range**





Multiple exposure photography

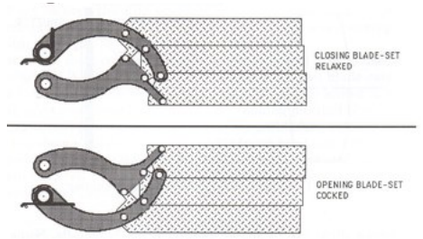
- **Sequentially measure all segments of the range**

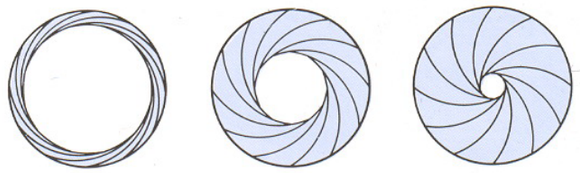


How do we vary exposure?

How do we vary exposure?

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



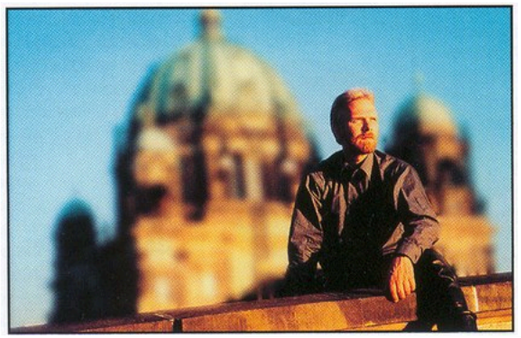


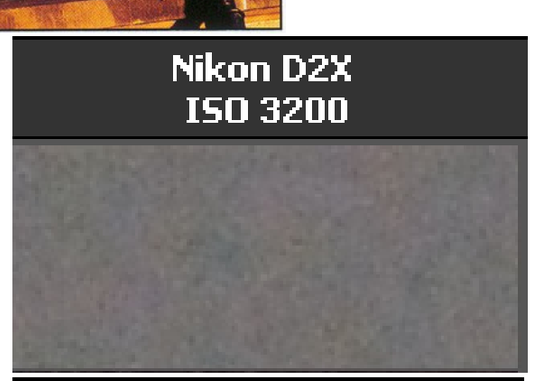


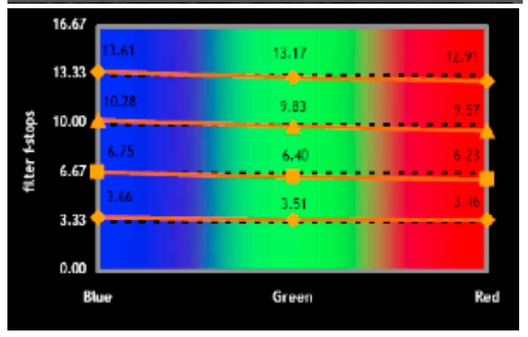
Slide inspired by Siggraph 2005 course on HDR

Tradeoffs

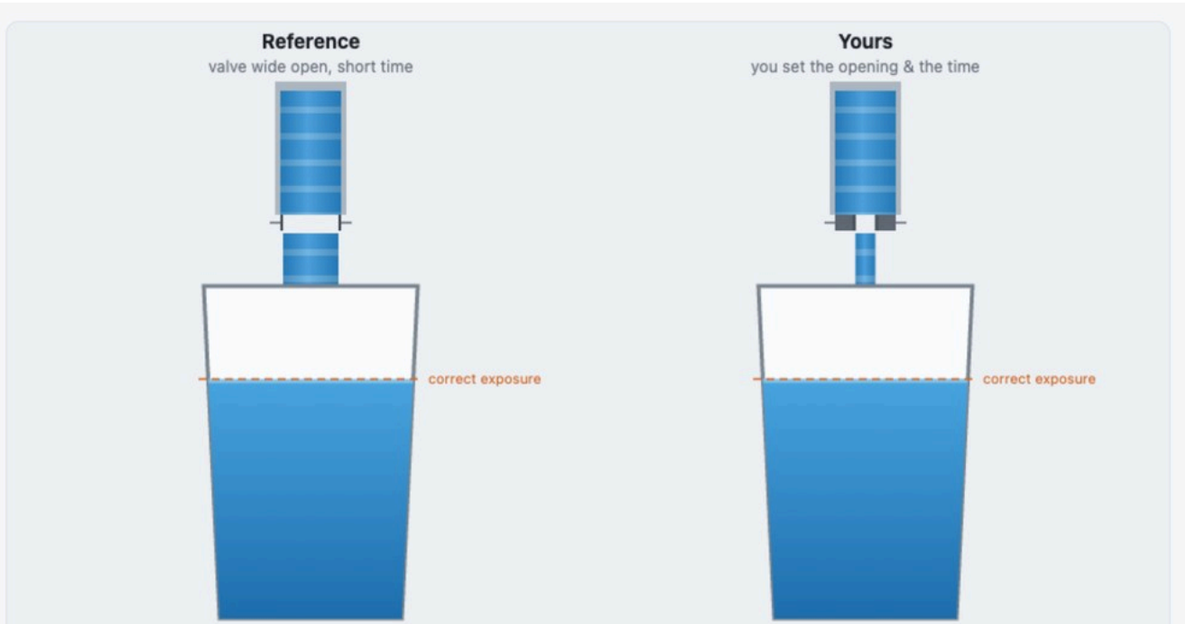
- **Shutter speed**
  - Range: ~30 sec to 1/4000sec (6 orders of magnitude)
  - Pros: reliable, linear
  - Cons: sometimes noise for long exposure
- **Aperture**
  - Range: ~f/1.4 to f/22 (2.5 orders of magnitude)
  - Cons: changes depth of field
  - Useful when desperate
- **ISO**
  - Range: ~100 to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







Demo: Exposure bucket analogy



▶ Fill the buckets | reset

your valve opening 2.83

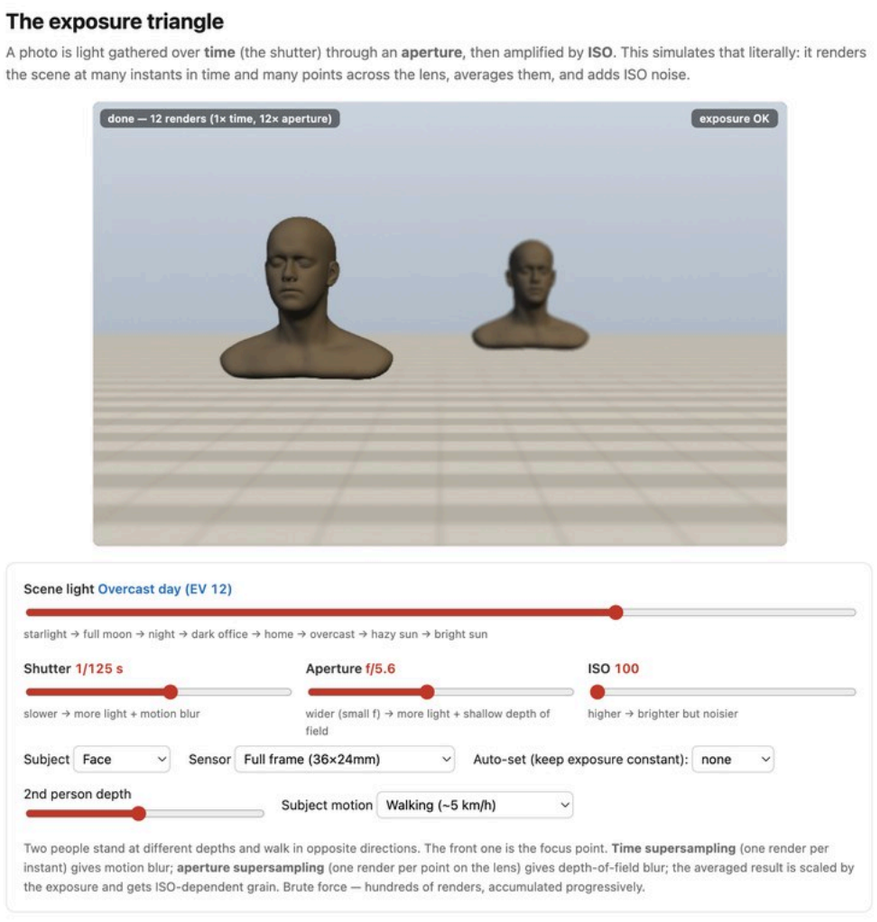
your fill duration 4.00 s

	valve opening	flow (∝ opening²)	duration	water = flow×time	result
Reference (left)	8.0	64	0.50 s	32	correct ✓
Yours (right)	2.83	8.0	4.00 s	32.0	correct ✓

Versus the reference: aperture **-3.0 stops** (smaller opening lets in less light), shutter **+3.0 stops** (longer lets in more) — net **+0.0 stops** of exposure. At the default they cancel exactly: reciprocity.

The classic exposure analogy. A bucket must catch the same amount of water for a correct **exposure**. The supply **pipe is always the same big bore**; what sets the flow is **how far the valve is open** — that is the **aperture** (widening the pipe's *nozzle* alone would not change the flow, so the valve is what throttles it). Flow goes as the opening's *area*, ∝ opening². The **fill duration** is the **shutter time** (how long it pours). The left bucket is the reference: the valve wide open for a short time. By default the right bucket is the opposite — a nearly closed valve open a long time — yet it catches the *same* water: three stops smaller opening cancels three stops longer time. That trade is **reciprocity**. Change either slider and the right bucket over- or under-fills (over- or under-exposes), past the dashed line for a correct exposure.

Demo: Exposure-triangle simulator



Gustave Le Gray, 1857

- http://www.metmuseum.org/collection/the-collection-online/search/283124



Figure 17.0.1: Gustave Le Gray, Mediterranean Sea, Sète (left) and Mediterranean…

Mediterranean Sea, Sète, 1857



Mediterranean with Mount Agde, 1857



Gustave Le Gray. The Metropolitan Museum of Art, CC0.

Gustave Le Gray, Mediterranean Sea, Sète (left) and Mediterranean with Mount Agde (right), both 1857, albumen prints.

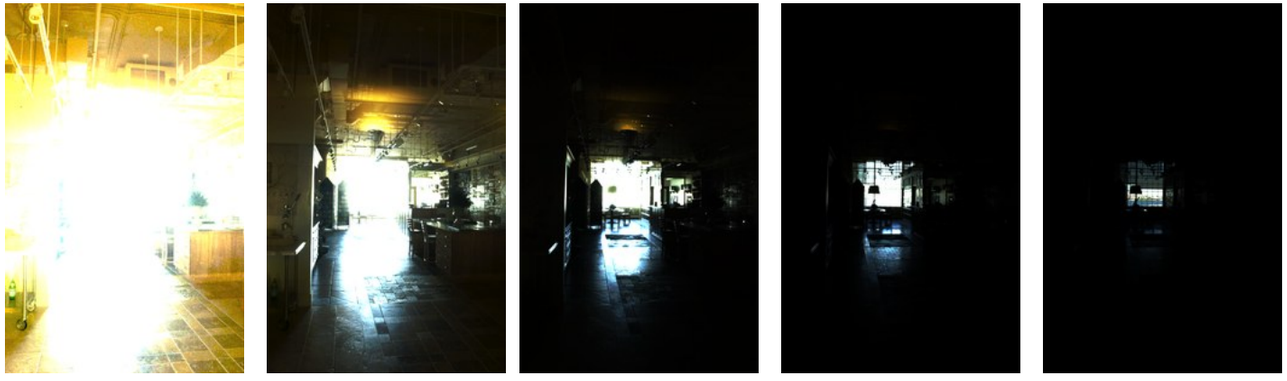
Questions?



HDR calibration: linear case

Problem statement

- **We have N images**
  - Images are encoded linearly
  - Only the exposure changes: no motion
- **We want 1 single HDR image**
  - encoded with same Image class
  - one value per x, y, c
  - values may be >1



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

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



Image formation

- 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, ISO, etc
- Noise n gets added
- values above 1 get clipped
  - depends on photosite well capacity

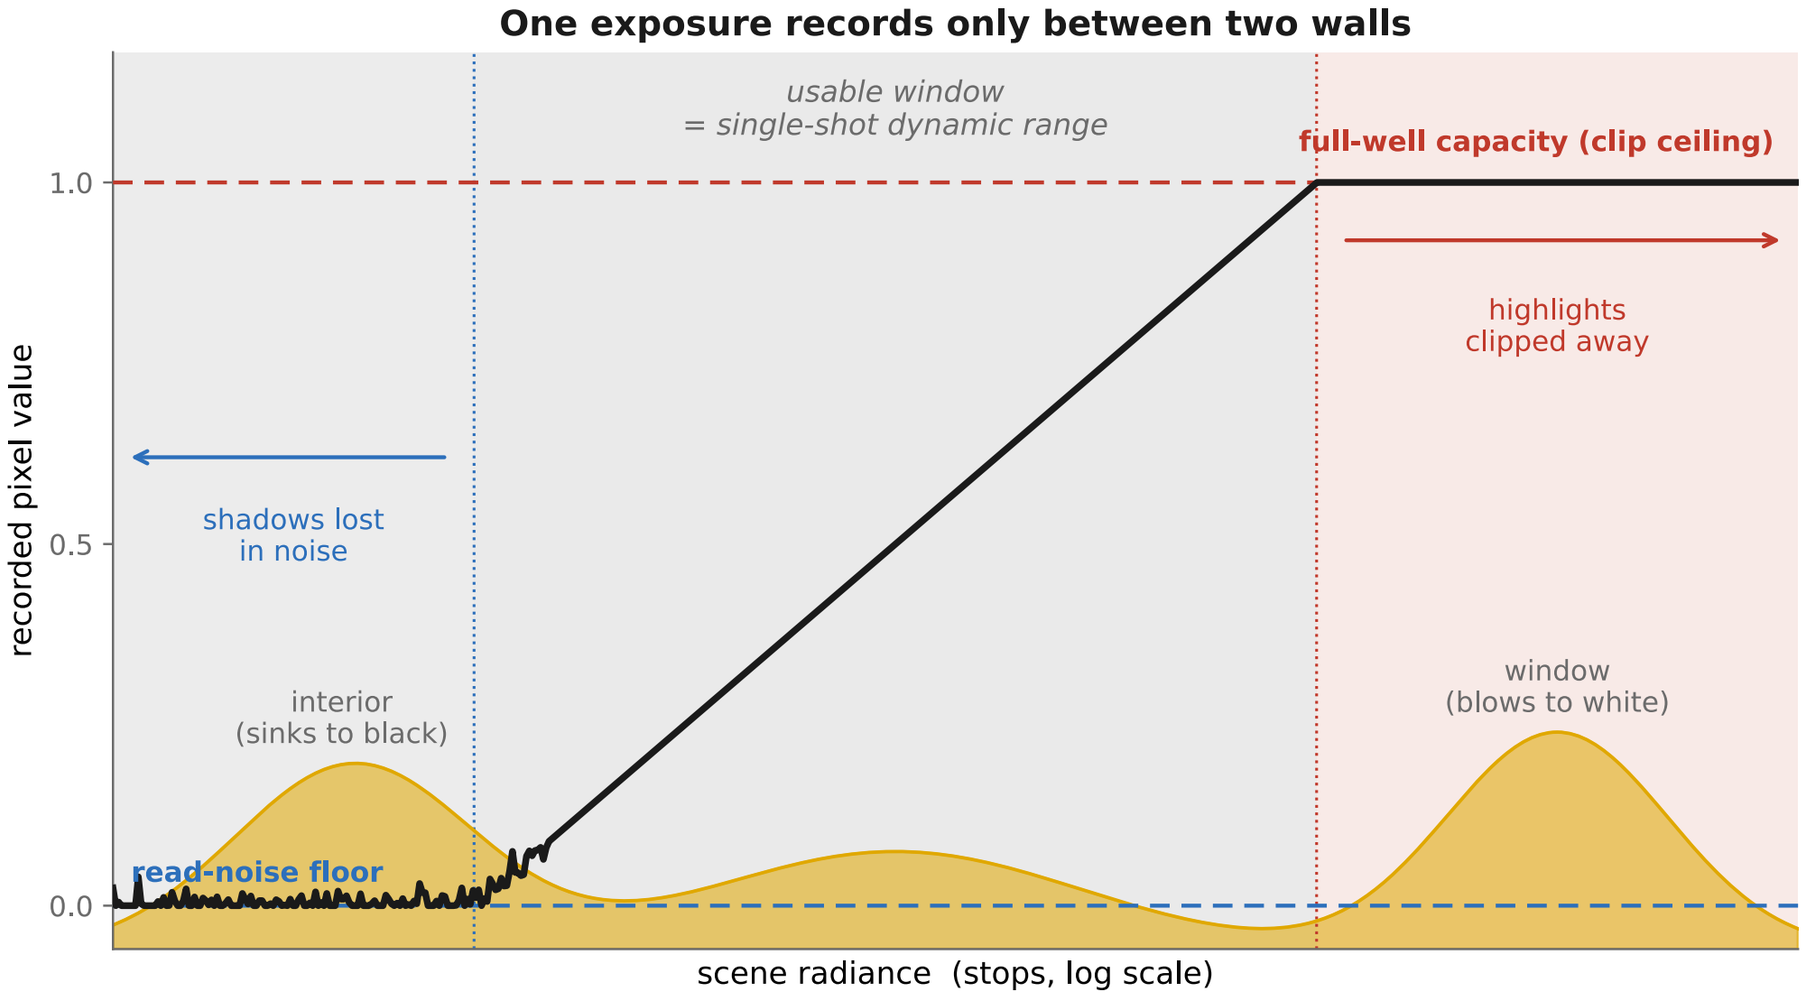
$$I_i(x, y)=clip(k_i\, L(x,y)+n)$$

Dynamic range

- In the highlights, we are limited by clipping
  - Value >1 saturate to 1
- In the shadows, we are limited by noise



Figure 17.1.4: One scene, one exposure, and the two walls of the single-shot dynamic…



One scene, one exposure, and the two walls of the single-shot dynamic range.

Merging HDR



Merging HDR

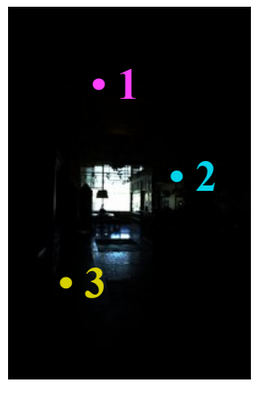


- For each pixel
  - figure out which images are useful
  - scale values appropriately according, ideally, to k_i
  - Voila!











Questions?



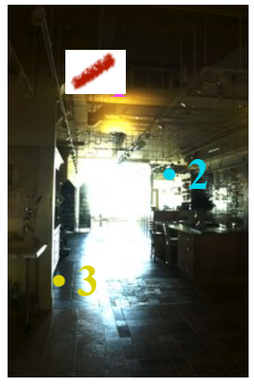
- For each pixel
  - figure out which images are useful
  - scale values appropriately according, ideally, to k_i
  - Voila!



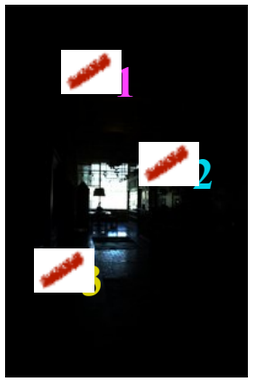
Which images are useful

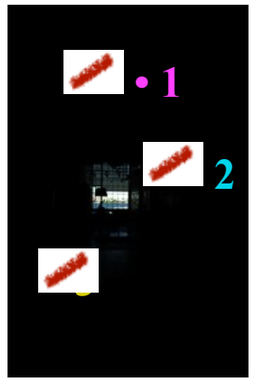
- Eliminate clipped pixels
  - e.g. >0.99
- Eliminate pixels that are too dark / too noisy
  - e.g. <0.002











Eliminate bad pixels

- In PSet 4
- We compute a weight map for each image
- We use binary weights (0 or 1)
  - But can be extended to full scalar
  - In particular for better handling noise
- Weights can be different for different channels

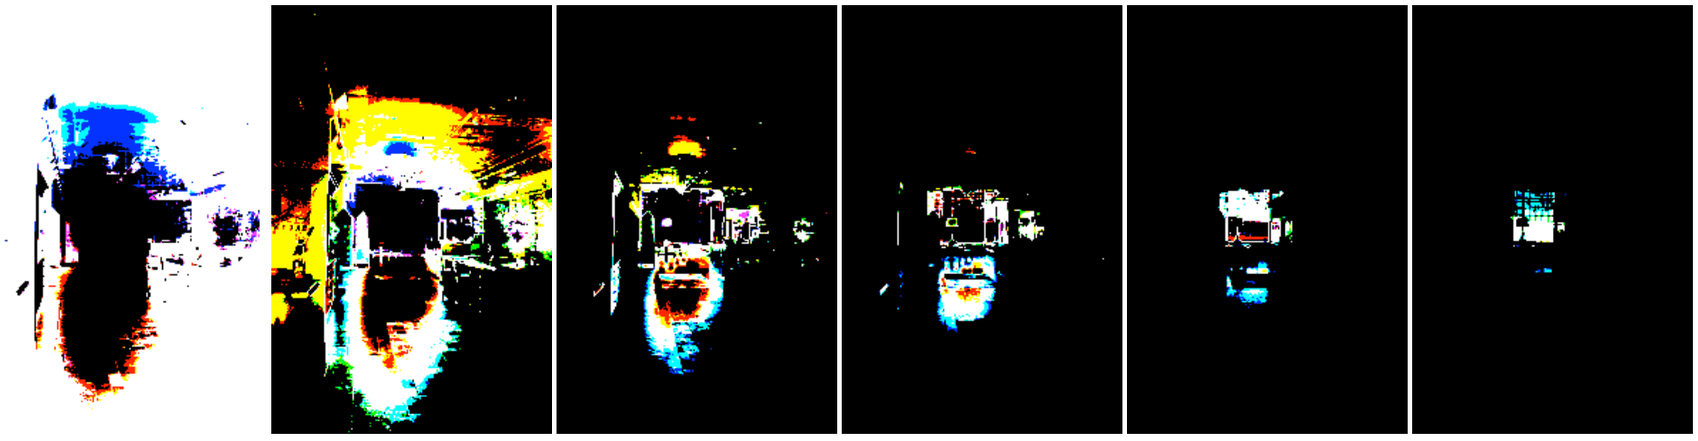
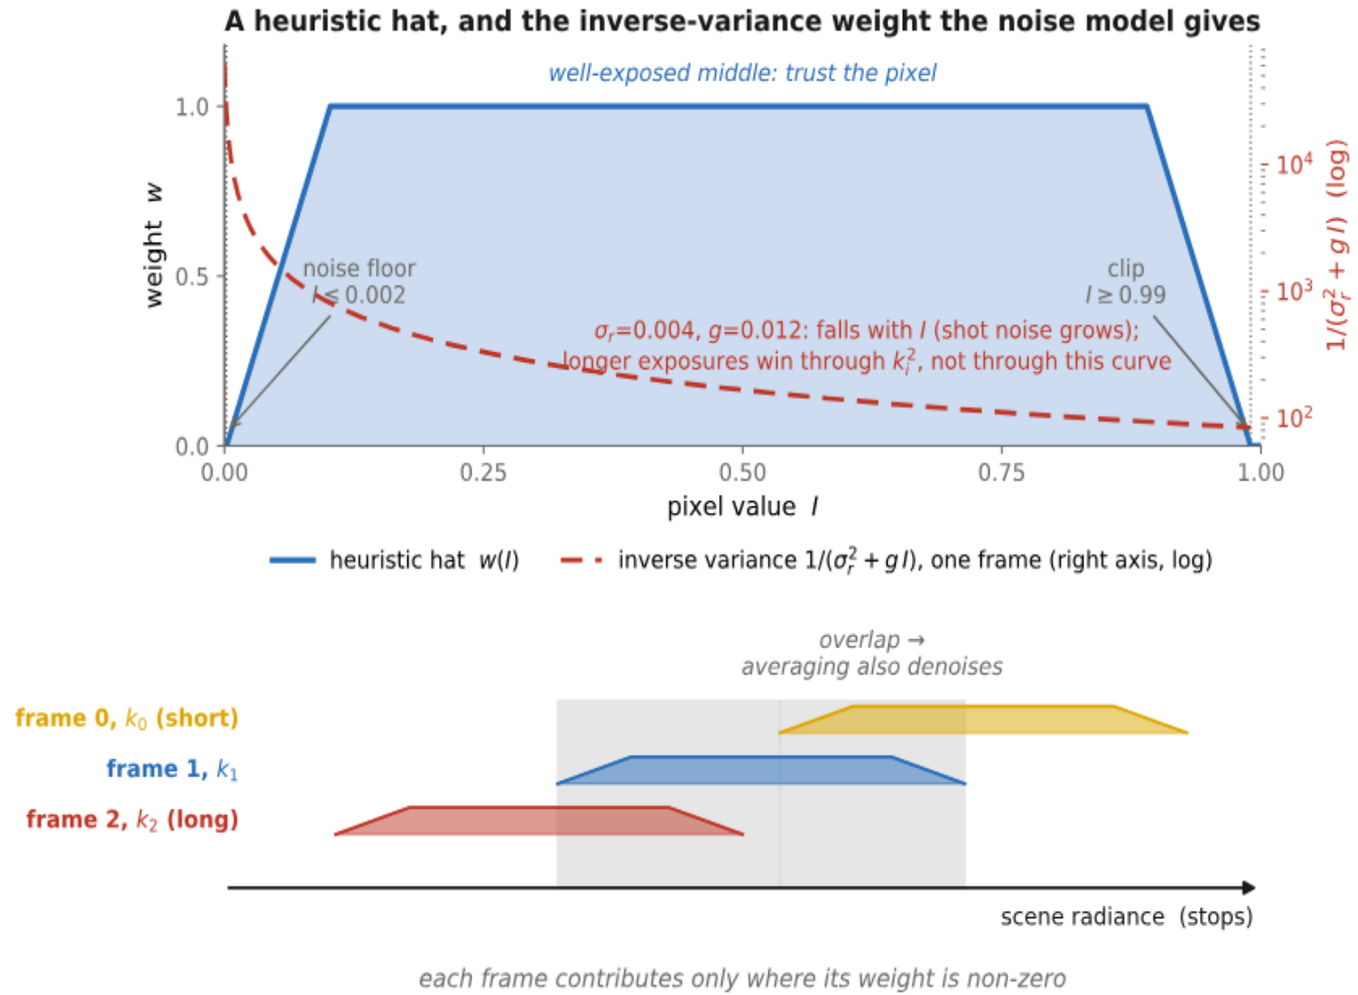


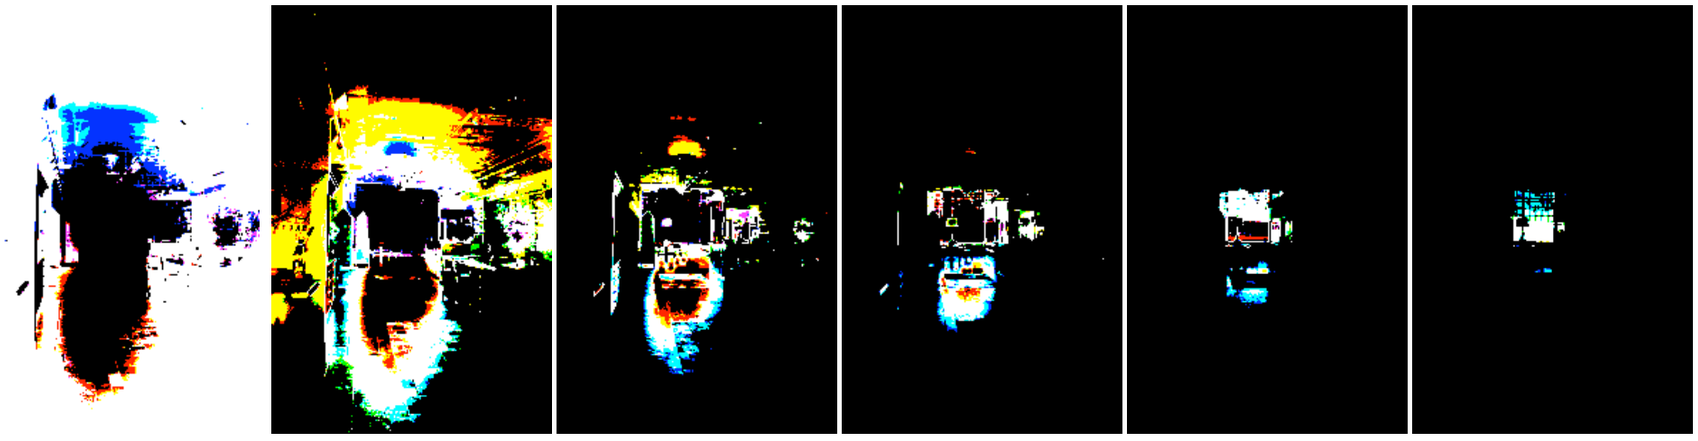
Figure 17.3.1: The reliability weight w(I) as a function of pixel value



The reliability weight w(I) as a function of pixel value.

Weight map

- We compute a weight map for each image



Questions?



Merge

Assembling HDR

Assembling HDR

- **Figure out scale factor between images**
  - From exposure data
  - Or by looking at ratios $I_i(x,y)/I_j(x,y)$
    - but only when both are good
- **Compute weight map w_i for each image**
- **Reconstruct full image using weighted combination per pixel**

$$\hat{L}(x,y) = \frac{\sum_i w_i(x,y)\, I_i(x,y)/k_i}{\sum_i w_i(x,y)}$$

Computing k_i



Computing k_i

- compute k_i/k_j for adjacent exposures
- just look at ratio of corresponding pixels
  - only for useful pixels



Computing k_i in pset 4

- 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

Computing k_i in pset 4

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

- 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

Cleaner way

- You could use a (robust) regression to avoid computing ratios
- Would also make it easy to directly model uncertainty due to noise
- Also generalizes to full non-linear response
  - See Debevec and Malik 97

Special cases

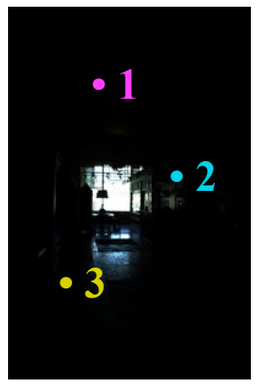
- **Some pixels might be underexposed or overexposed in all images**
- **Simple solution: don't eliminate dark pixels in the brightest image or bright pixels in the darkest one.**
  - Pay attention: dark pixels in the brightest (and vice versa), not other way around.
  - Careful is you implement it by changing the thresholds that values 0.0 (resp. 1.0) are actually kept

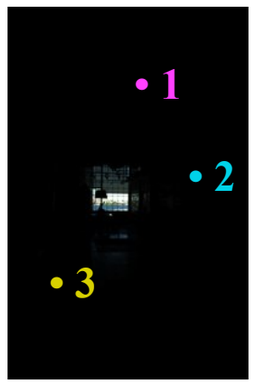


I call this the brightest image









I call this the darkest image

In the end: HDR image

- Encoded with same image class
- One single float value per x, y, c
- numbers may be >1

Demo: HDR merge

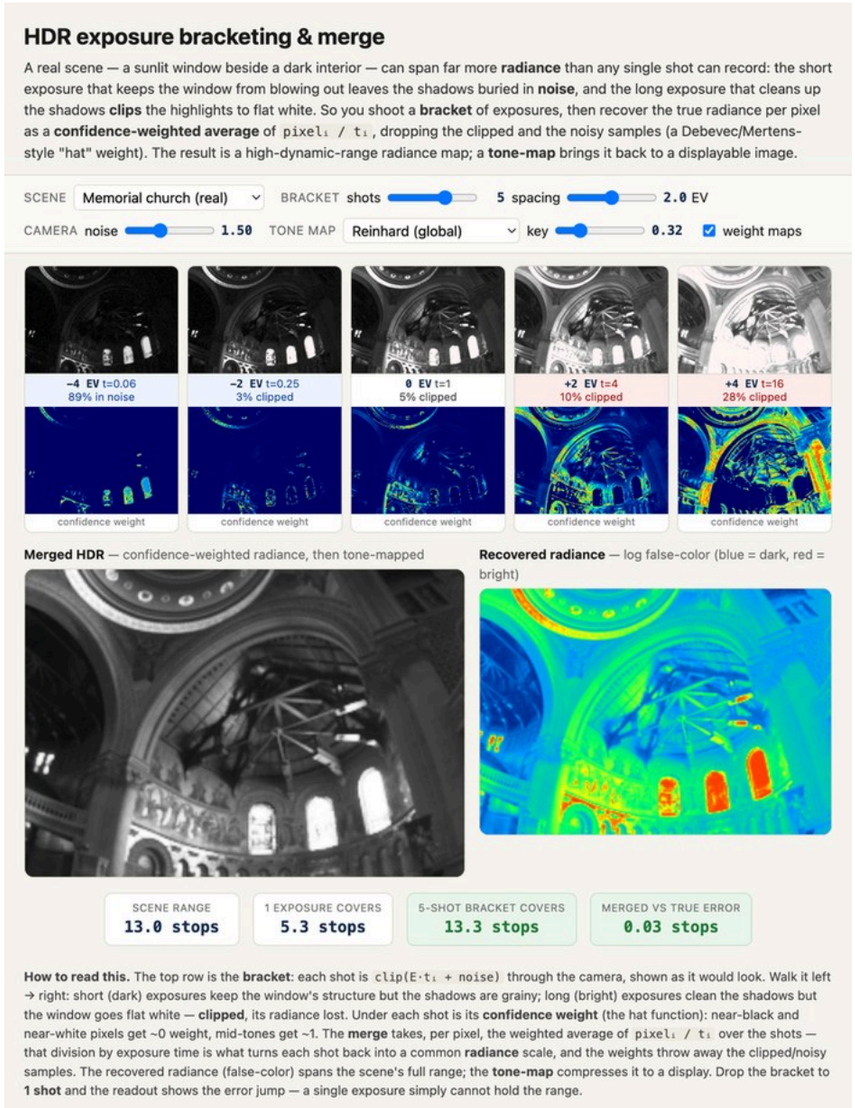
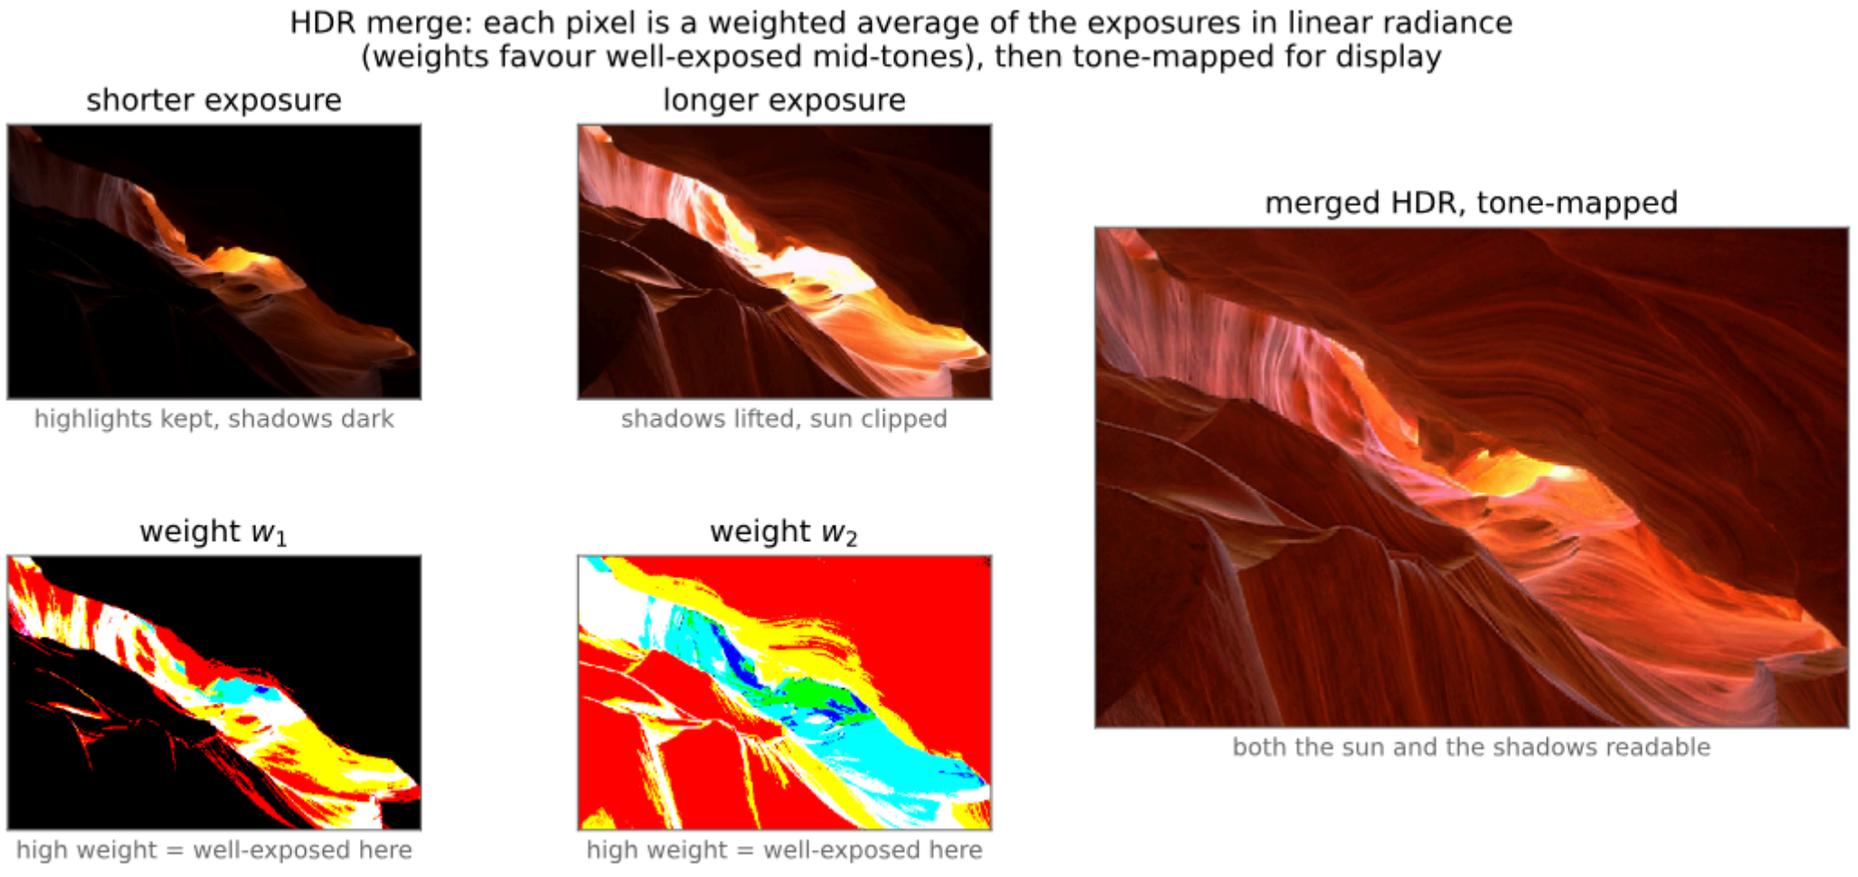
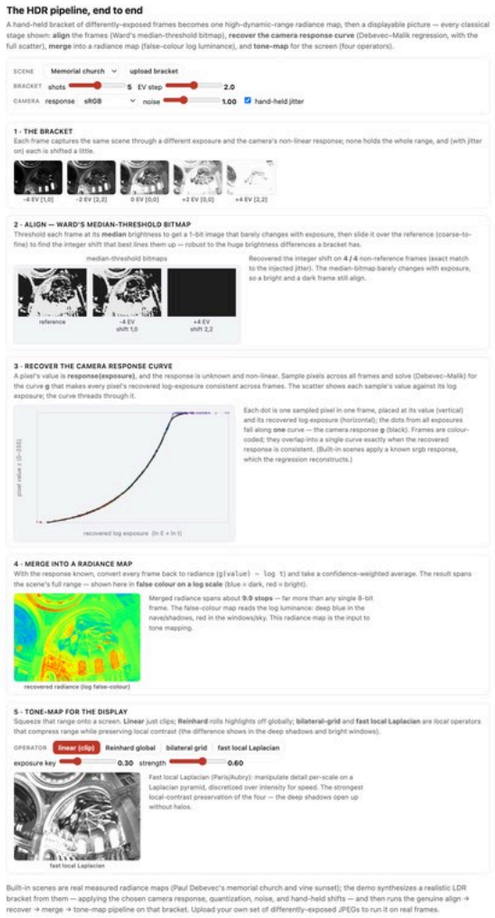


Figure 17.3.2: The merge on a real bracket (Antelope Canyon)



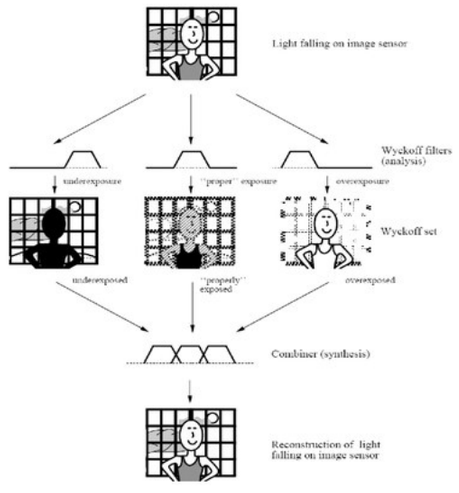
A shorter exposure keeps the bright sunlit slot but buries the rock in noise-dark shadow; a longer exposure lifts the shadows but clips the sun.

Demo: Full HDR pipeline



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/**



Advanced issues

Calibrating the camera response

- Most cameras and JPEGs are non-linear: $I_i = f(k_i\ L)$
- Debevec and Malik 97: a pixel seen at several exposures constrains one radiance L_p
  - solve for the log inverse $g = \ln f^{-1}$, one value per pixel level
  - linear least squares, with smoothness on g and weights w(I)
  - known only up to a global scale
- Mitsunaga and Nayar 99: polynomial inverse response, recovers the exposure ratios too

$$g(I_{p,i}) = \ln L_p + \ln t_i$$

Demo: Recovering the response curve

the bracket, and the 120 sample pixels



1/500 s, ISO 6400
ln t = -2.06

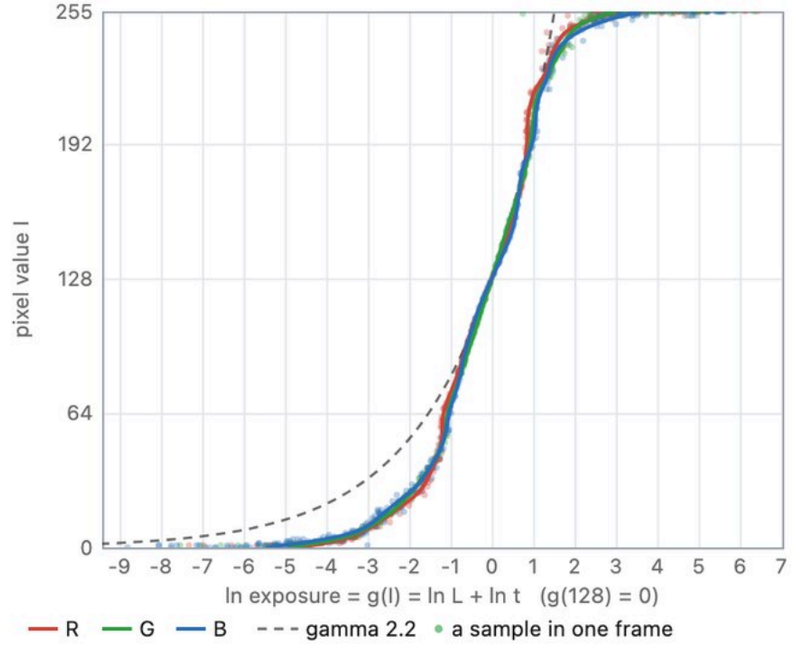
1/60 s, ISO 6400
ln t = 0.06 (+3.1 stops)

1/8 s, ISO 6400
ln t = 2.08 (+6.0 stops)

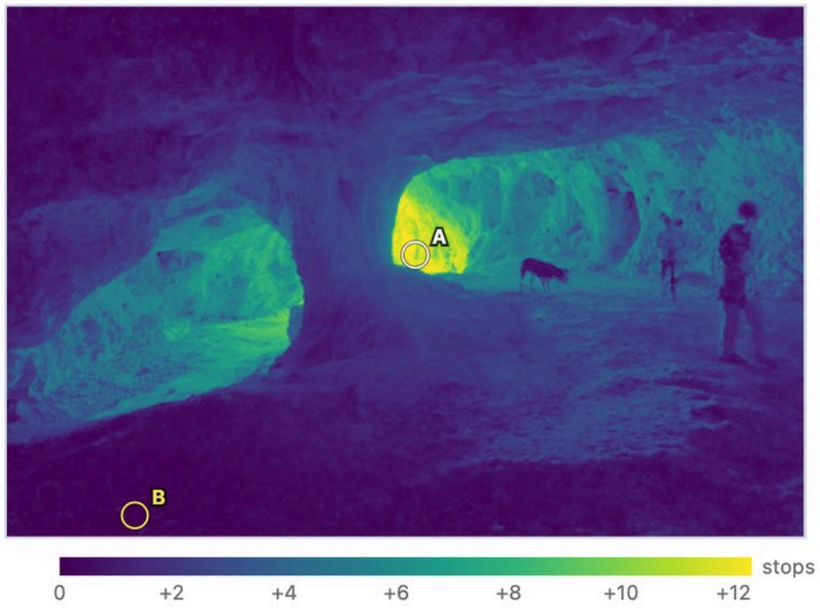
Debevec-Malik: least squares, per channel

376 unknowns: 256 values of g, one ln L per sample; 360 data rows, 254 smoothness rows, and g(128) = 0
fit residual 0.040 stops RMS
after the merge, the frames disagree by
0.231 stops with the recovered curve
1.071 stops with gamma 2.2
the map spans 12.3 stops, 1 : 5,069
(0.5th to 99.9th percentile of luminance)

the recovered response, Debevec-Malik: I against ln exposure

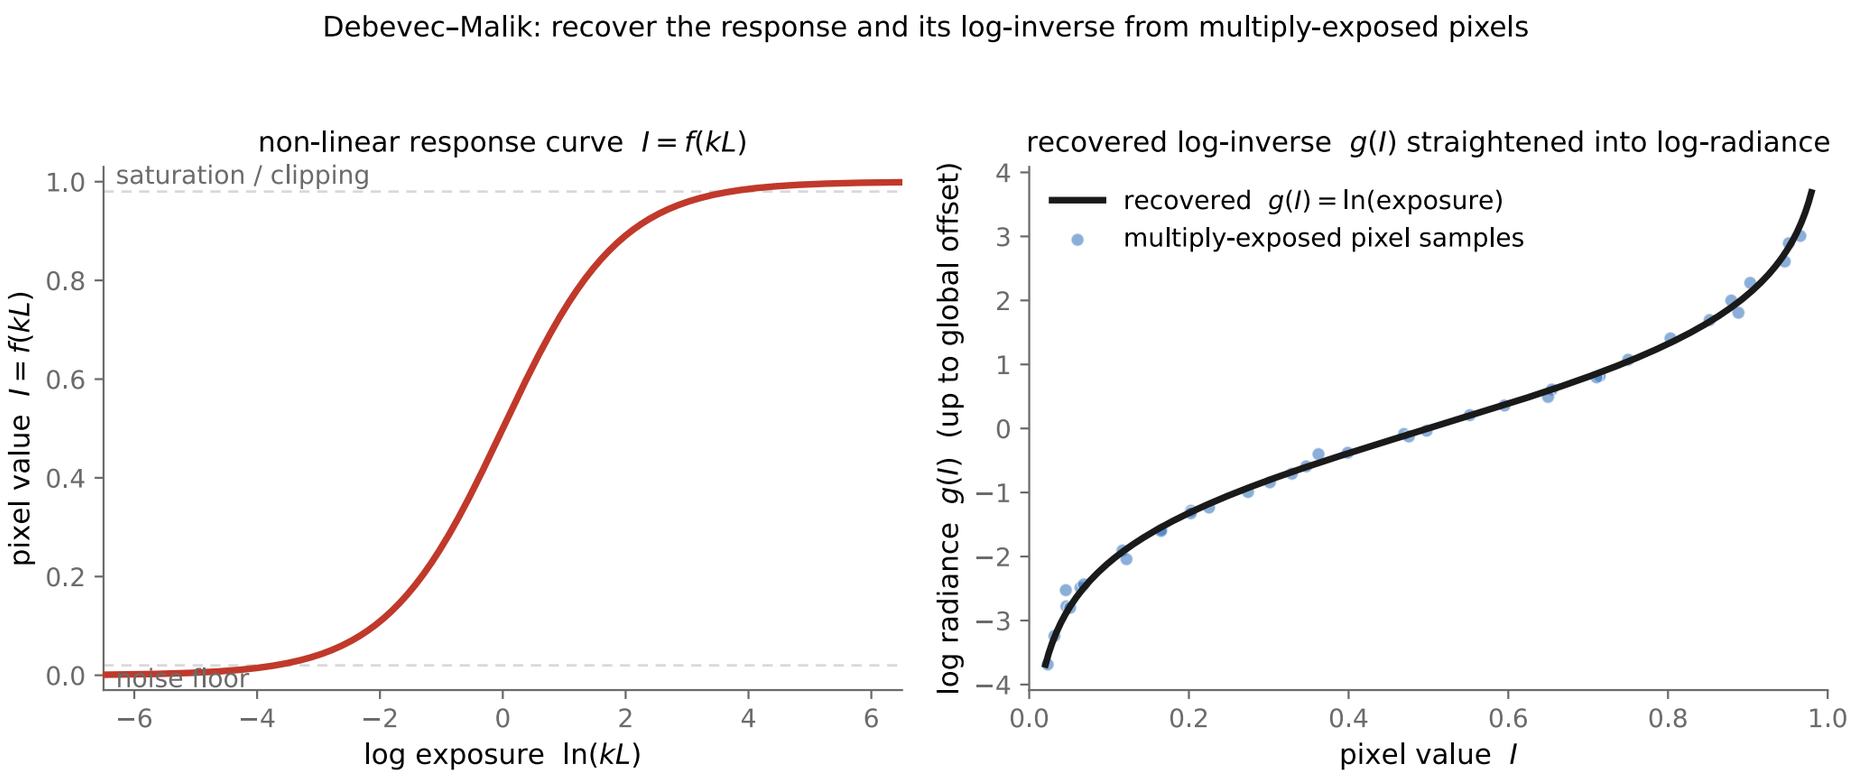


the radiance map, log luminance, Debevec-Malik



radiance at A over radiance at B: 2,096 : 1 (11.0 stops)
drag A and B; the map is known only up to one scale, so ratios are what it measures

Figure 17.2.2: The recovered camera response (Debevec–Malik)



Left: the camera's non-linear response curve I = f(kL) (pixel value against log exposure), an S-shaped film/JPEG curve.

Absolute calibration: from exposure to cd/m²

- The merge gives radiance only up to one global scale
- A reflected-light meter fixes the scale for the metered patch
  - N: f-number, t: exposure time, S: ISO, K: the meter's constant
  - K = 12.5, f/2.8, 1/100 s, ISO 400 gives $L_v \approx 24.5$ cd/m²
- On the sensor: illuminance $E_v \approx \pi\tau\, L_v / 4N^2$, exposure $H_v = E_v\, t$
- Counting photons also needs the spectrum and the quantum efficiency
- EXIF alone calibrates nothing: measure black level, gain, linearity, vignetting

$$L_v = \frac{K\,N^2}{t\,S}$$

Demo: Beauty-curve playground

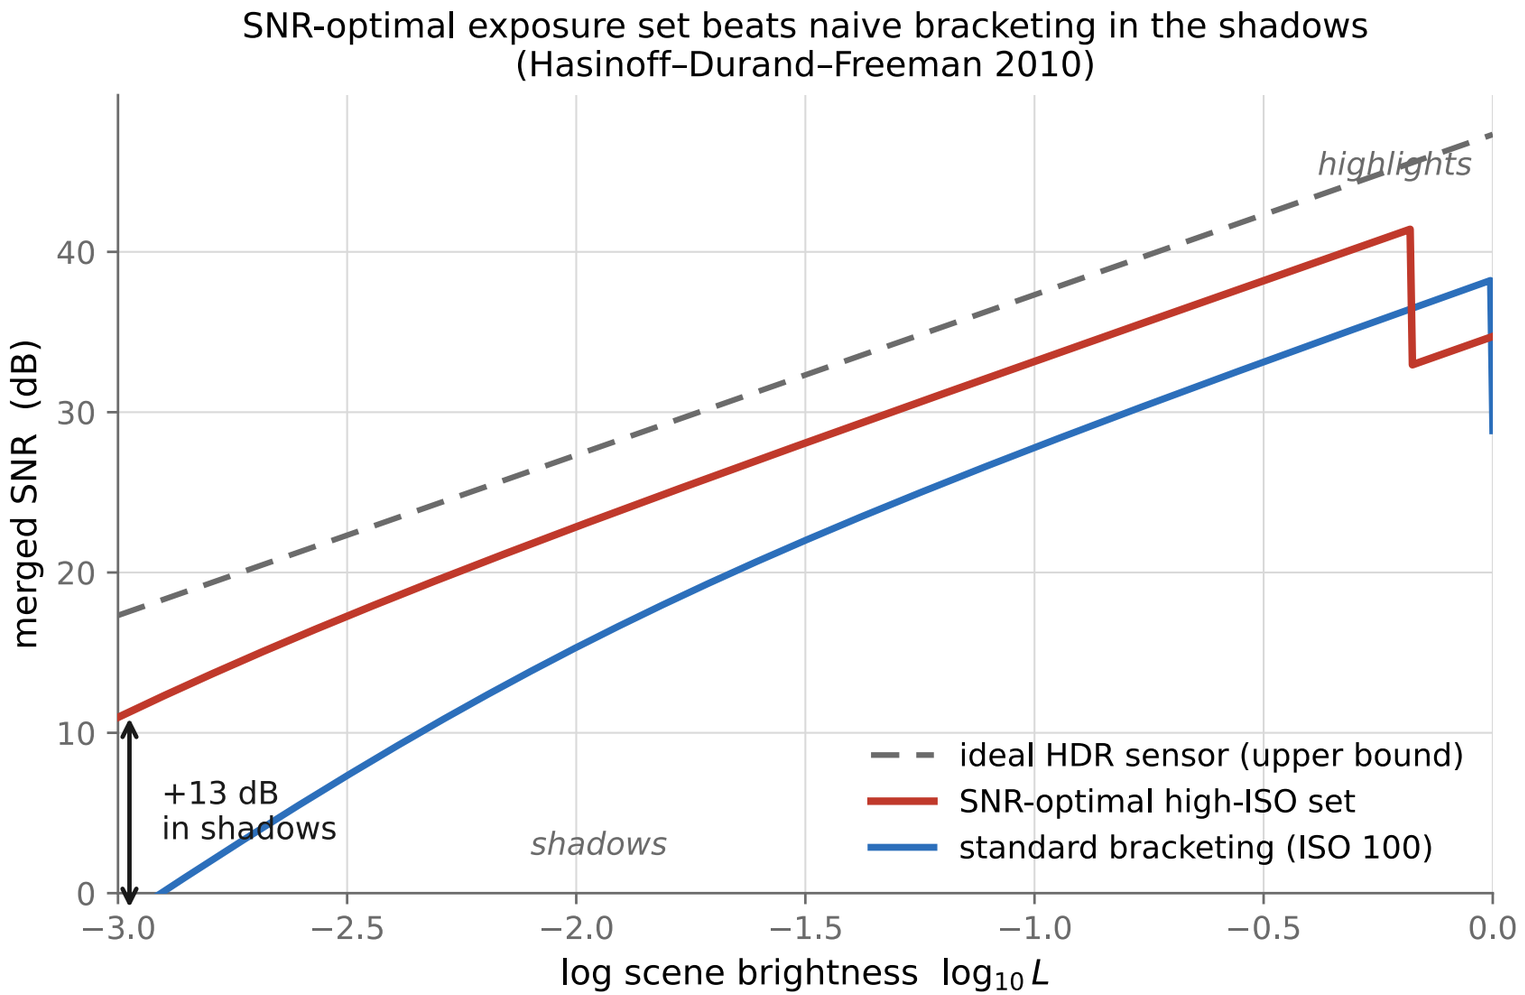


Optimal merge and capture

- Weight each frame by the inverse variance of its radiance estimate I_i/k_i
  - shot noise grows with the signal, read noise is constant
  - on the radiance scale the long exposure is more precise, until it clips
- Capture: for a fixed time budget, choose shutter times and ISO to maximize the worst-case SNR
  - Hasinoff, Durand and Freeman 2010: about 12 dB better than a plain bracket in their example

$$\hat{L} = \frac{\sum_i (I_i/k_i)/\sigma_i^2}{\sum_i 1/\sigma_i^2}, \qquad \sigma_i^2 = \frac{a\,k_i L + \sigma_{\text{read}}^2}{k_i^2}$$

Figure 17.3.6: The SNR-optimal exposure set (Hasinoff–Durand–Freeman)



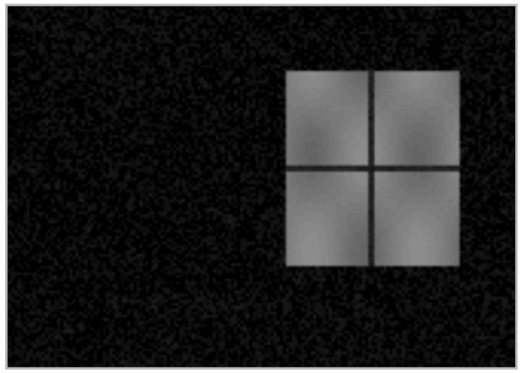
SNR (in dB) is plotted against log scene brightness for three strategies: standard bracketing at ISO 100, the SNR-optimal high-ISO short-exposure set, and the ideal HDR…

Motion: ghosts and alignment

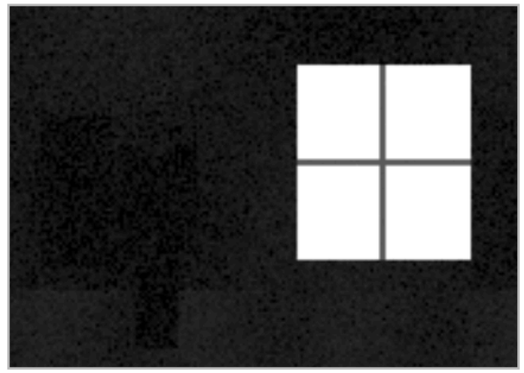
- Anything that moves between exposures appears several times: a ghost
- Align first: translation or homography, e.g. Ward's median threshold bitmap
  - parallax and moving subjects remain
- Then reject: pick a reference, drop samples that disagree beyond the noise
  - the reference decides the composition
  - rejection costs range exactly where it was needed
- Alternatives: optical flow, patch-based reconstruction, learned merges

$$d_i = \frac{|q_i - q_r|}{\sqrt{v_i + v_r + \epsilon}}$$

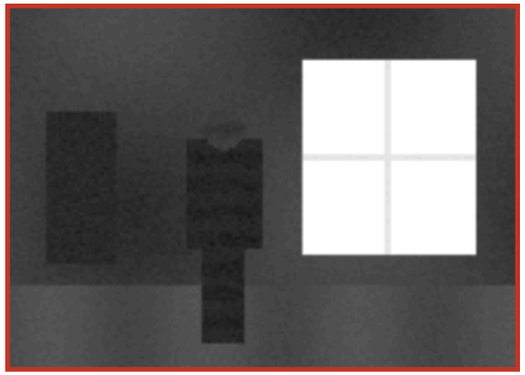
Demo: HDR ghosting



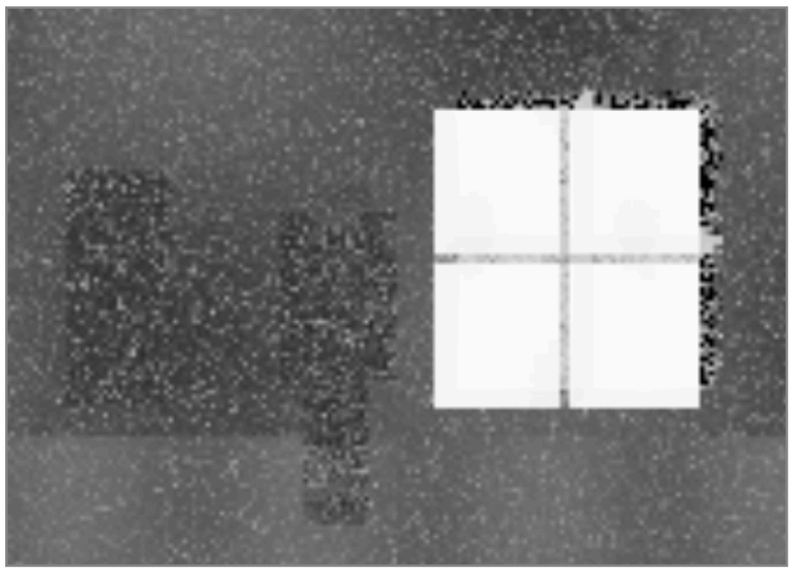
1/125 s



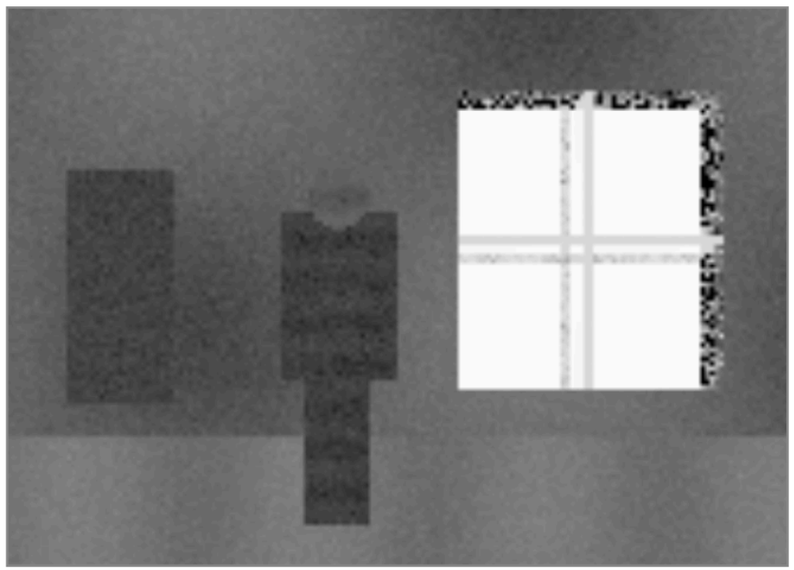
1/15 s



1/2 s reference



naive average



rejection, for comparison

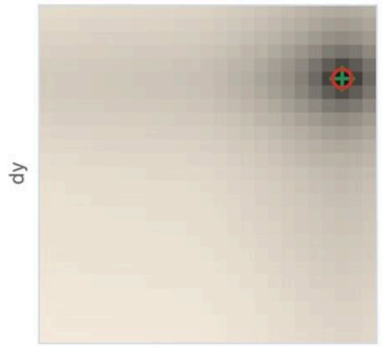
where another frame held the subject, this merge is off by 0.61 stops

Demo: Aligning a hand-held burst

frame 6, as it came



SSD against frame 1, every shift

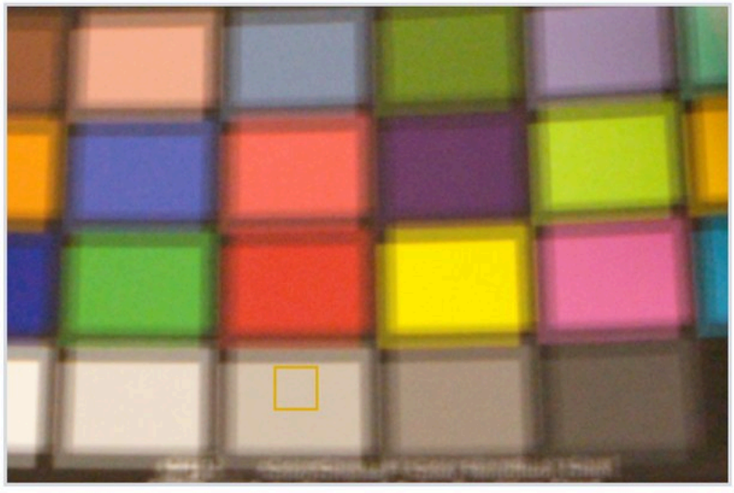


+ true shift
○ smallest SSD

frame	true shift	found
2	(0, 6)	(0, 6) ✓
3	(6, 4)	(6, 4) ✓
4	(4, 1)	(4, 1) ✓
5	(7, 1)	(7, 1) ✓
6	(10, -7)	(10, -7) ✓
7	(12, -8)	(12, -8) ✓
8	(12, -8)	(12, -8) ✓

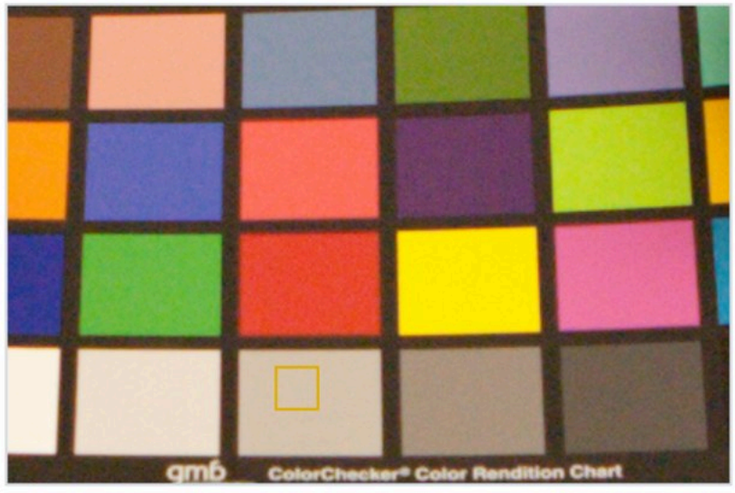
7 of 7 found

8 frames, averaged as they came



noise in the grey patch 1.48 · error against the truth 33.19

8 frames, aligned, then averaged



noise in the grey patch 0.95 · error against the truth 1.85

one frame: noise 2.81, so 8 aligned frames should reach 1.06 (1/√N, plus the truth's own noise). Codes out of 255; the truth is the mean of all 50 tripod frames.

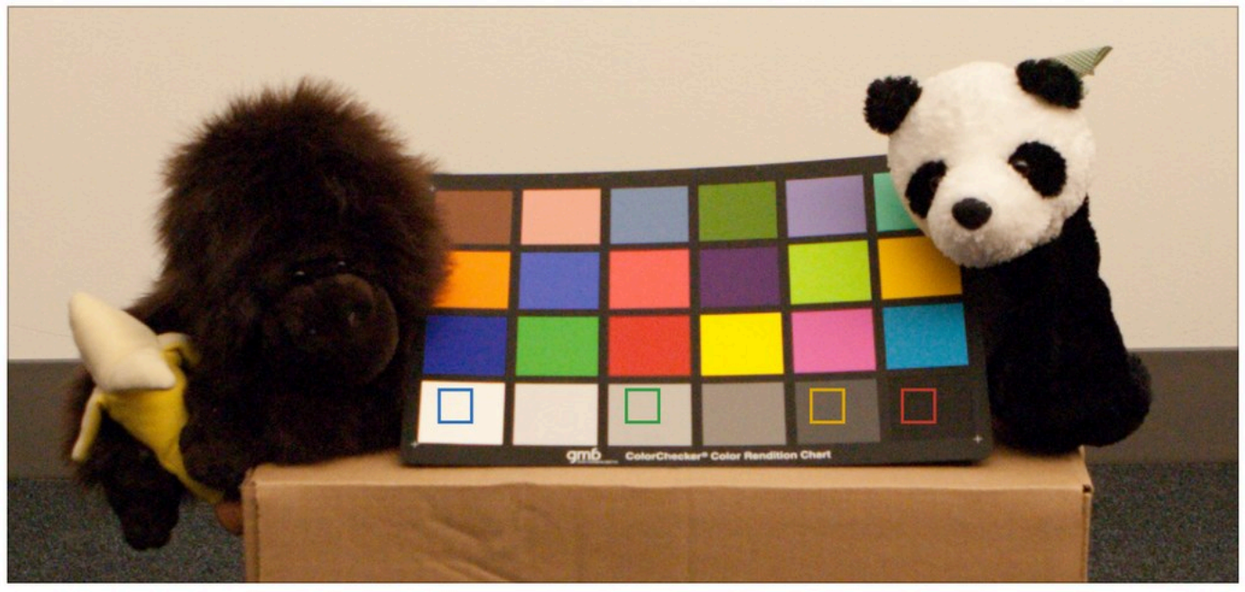
HDR+: a burst at one short exposure

- Phones capture a burst of raw frames at one short exposure that protects the highlights
- Merging N frames reduces shadow noise
  - it does not extend the highlights: a lamp clipped in every frame stays clipped
- Align tiles locally, merge robustly in raw against a reference, then demosaic
- Hasinoff et al. 2016; Night Sight and Super Res Zoom build on it

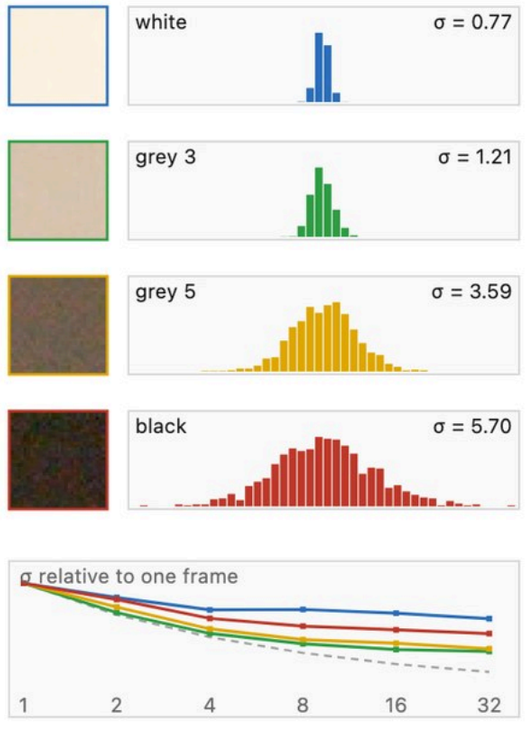
$$\sigma_{\text{merged}} \approx \frac{\sigma_{\text{single}}}{\sqrt{N}}$$

Demo: Averaging a burst: noise falls as 1/√N

the average of 8 frames



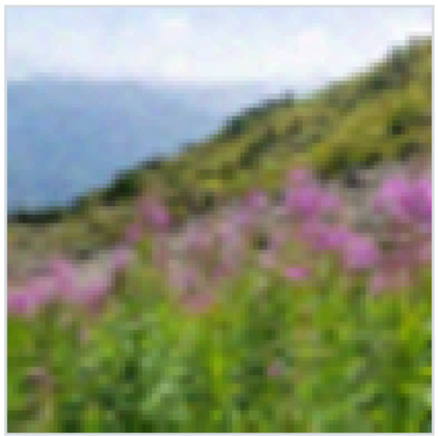
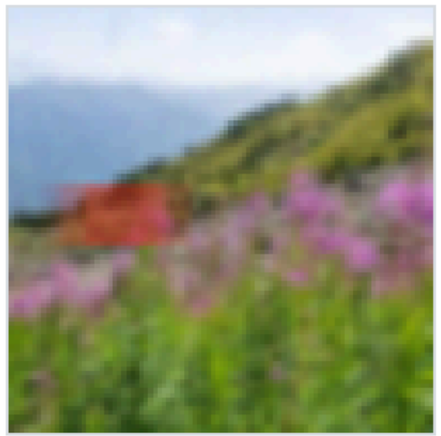
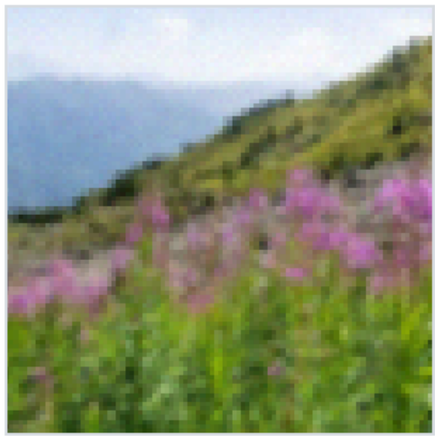
ISO 3200, aligned in advance; the same scene the problem set uses · hover to magnify



Demo: Robust burst merge

one burst, three merges, with something moving in it

9 frames · tremor ±1.4 px · subject at 3.0 px/frame

one frame	average, no fine grid	merge, no gate	merge, gated
			
24.87 dB	22.79 dB	24.11 dB	25.94 dB

2.3% of in-bounds samples rejected · gate 0.25

Gated minus ungated PSNR: 1.83 dB. Texture and noise can also trigger rejection.

HDR+ on Android

- http://static.googleusercontent.com/media/hdrplusdata.org/en//hdrplus.pdf
- https://ai.googleblog.com/2018/02/introducing-hdr-burst-photography.html

Burst photography for high dynamic range and low-light imaging on mobile cameras

Samuel W. Hasinoff Dillon Sharlet Ryan Geiss Andrew Adams
Jonathan T. Barron Florian Kainz Jiawen Chen Marc Levoy
Google Research

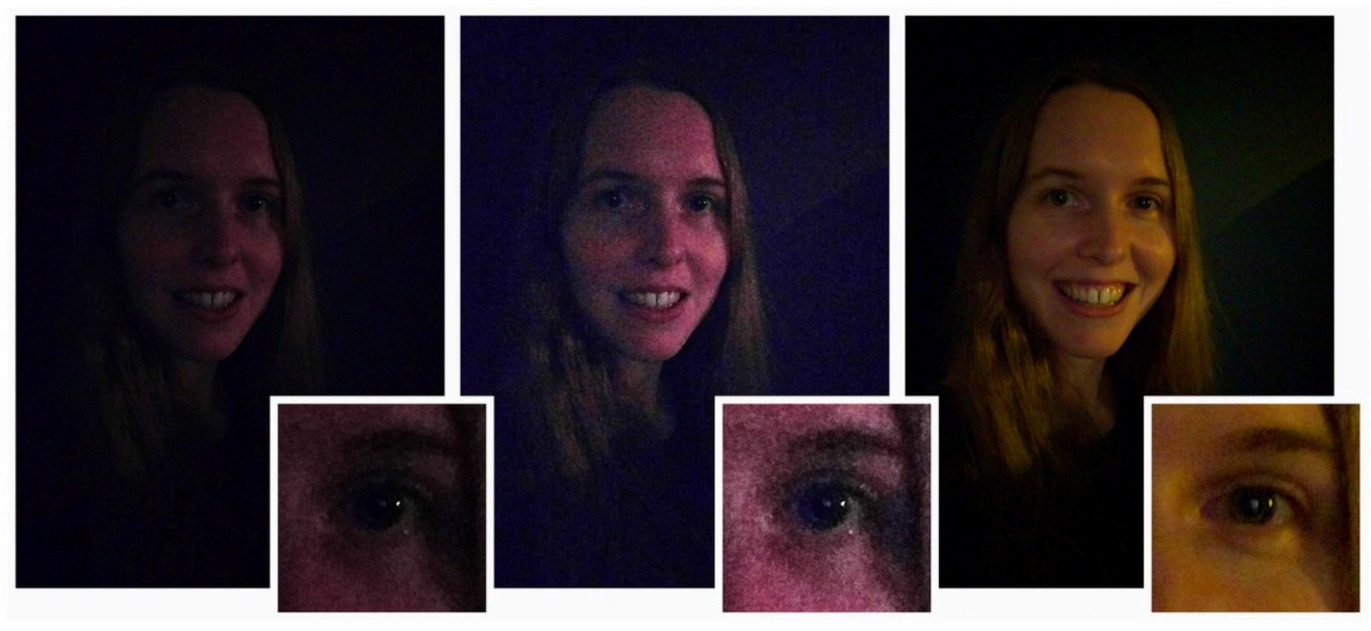
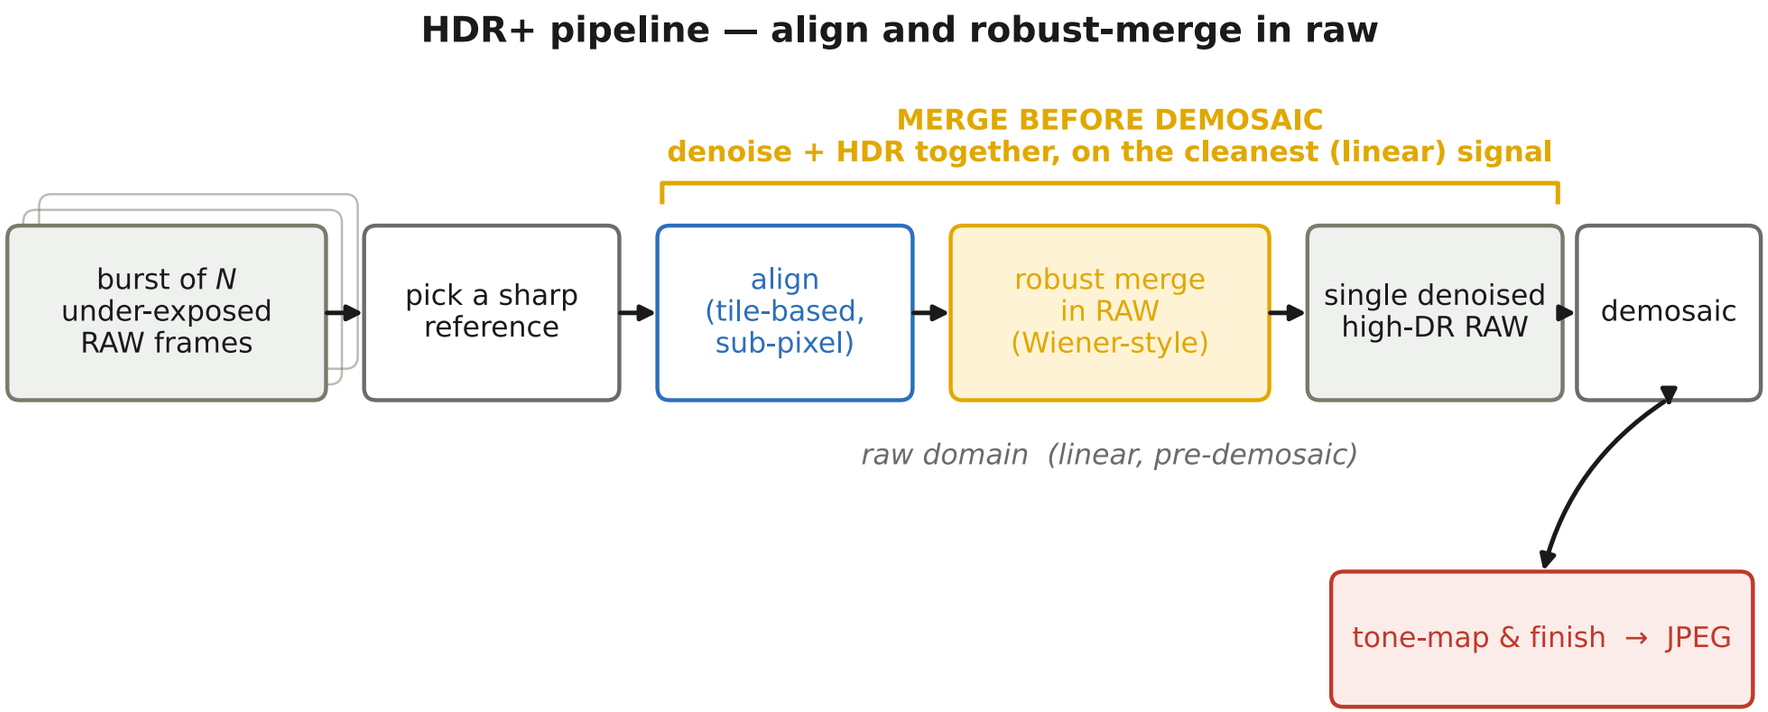


Figure 17.7.5: The HDR+ flow



Left to right: a burst of N identical under-exposed raws → pick a sharp reference frame → align every other frame to it (tile-based, sub-pixel) → robust merge in raw…

HDR in the sensor

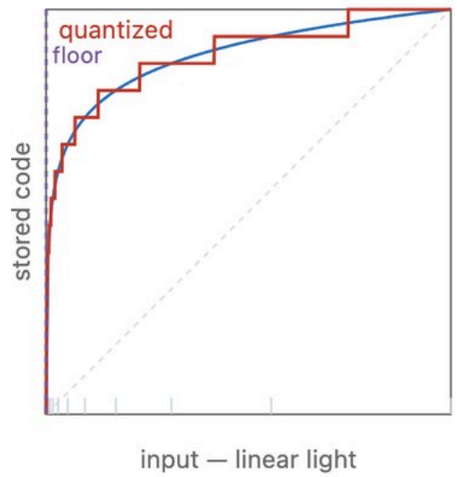
- One capture, several measurements per pixel: no bracket to align
- Time: staggered short and long exposures, read out as one frame
- Gain: two conversion gains read at once
  - low gain keeps the highlights, high gain lowers the shadow noise floor
- Space: split pixels (a large and a small photodiode), or an overflow capacitor (LOFIC) that stores extra charge
- The measurements can still differ in time: motion and LED flicker
  - an LED pulsed at a few hundred hertz can be off for a whole 1/2000 s exposure

Storing HDR images

- The merge outputs scene-linear values, often well above 1
- OpenEXR: half or full floating point per channel
- Radiance RGBE: one shared exponent per pixel
- Integer encodings need a curve: log, PQ or HLG spread levels more evenly in stops
  - linear integers waste levels in the highlights and starve the shadows
- Gain-map JPEG (Ultra HDR): an SDR picture plus a map to rebuild the HDR one

Demo: Linear, gamma and log storage

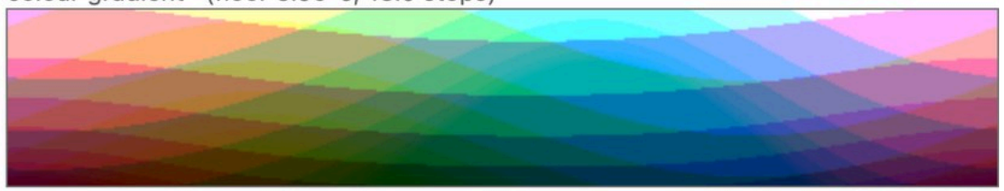
log storage, floor 1.5e-4 (12.7 stops below white) · 4-bit → 16 levels · display = sRGB



grey ramp (floor 8.1e-5, 13.6 stops)



colour gradient (floor 3.9e-6, 18.0 stops)



photograph (floor 1.5e-4, 12.7 stops)



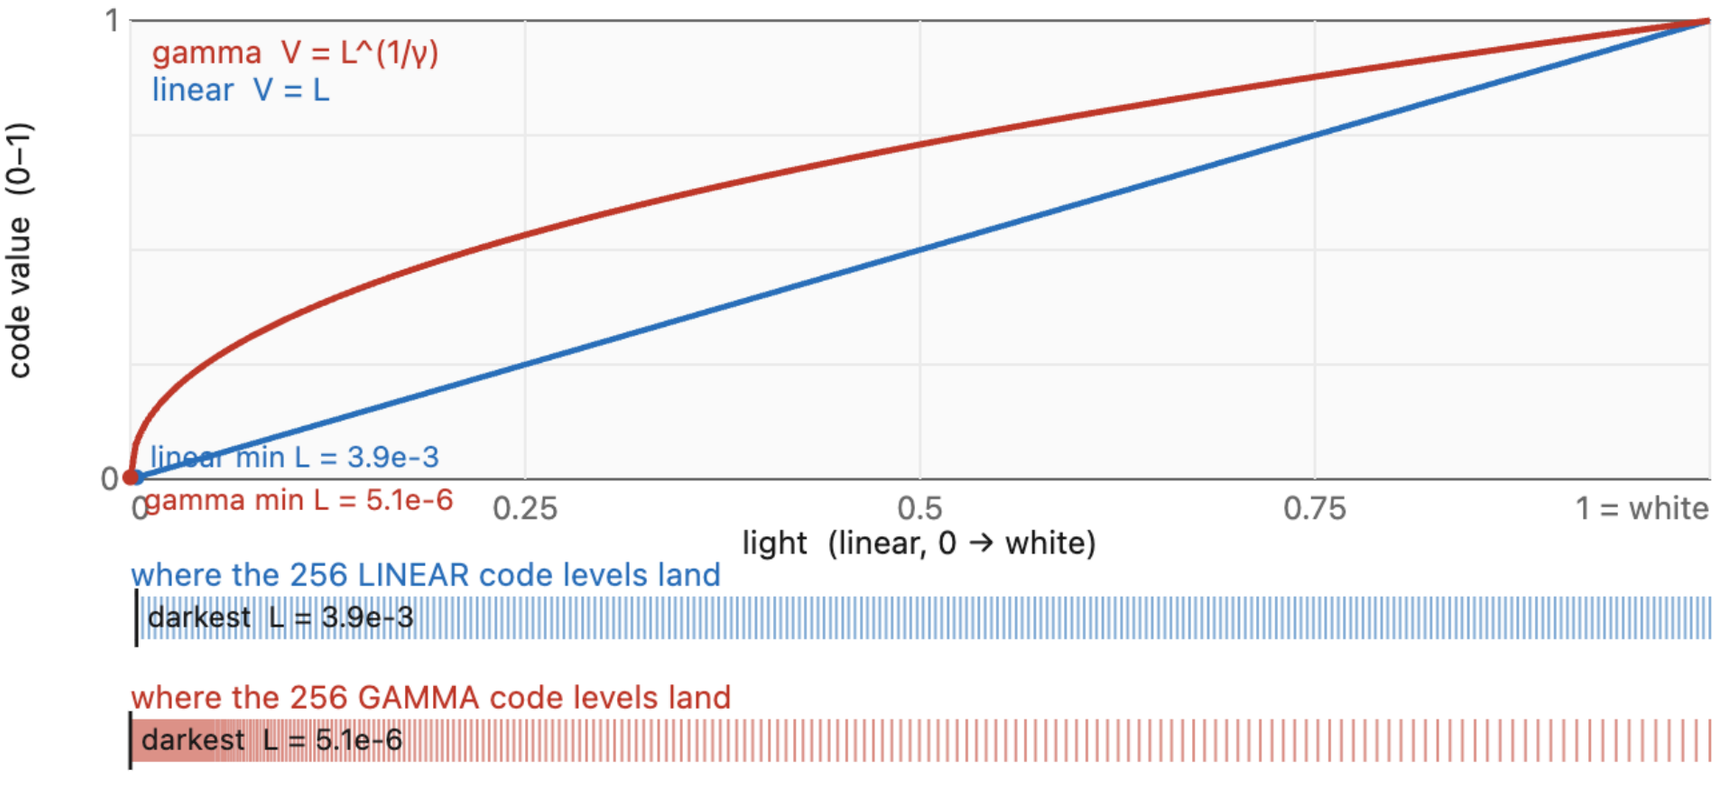
Demo: Encoding dynamic range

Smallest nonzero level and dynamic range (relative to white = 1.0):

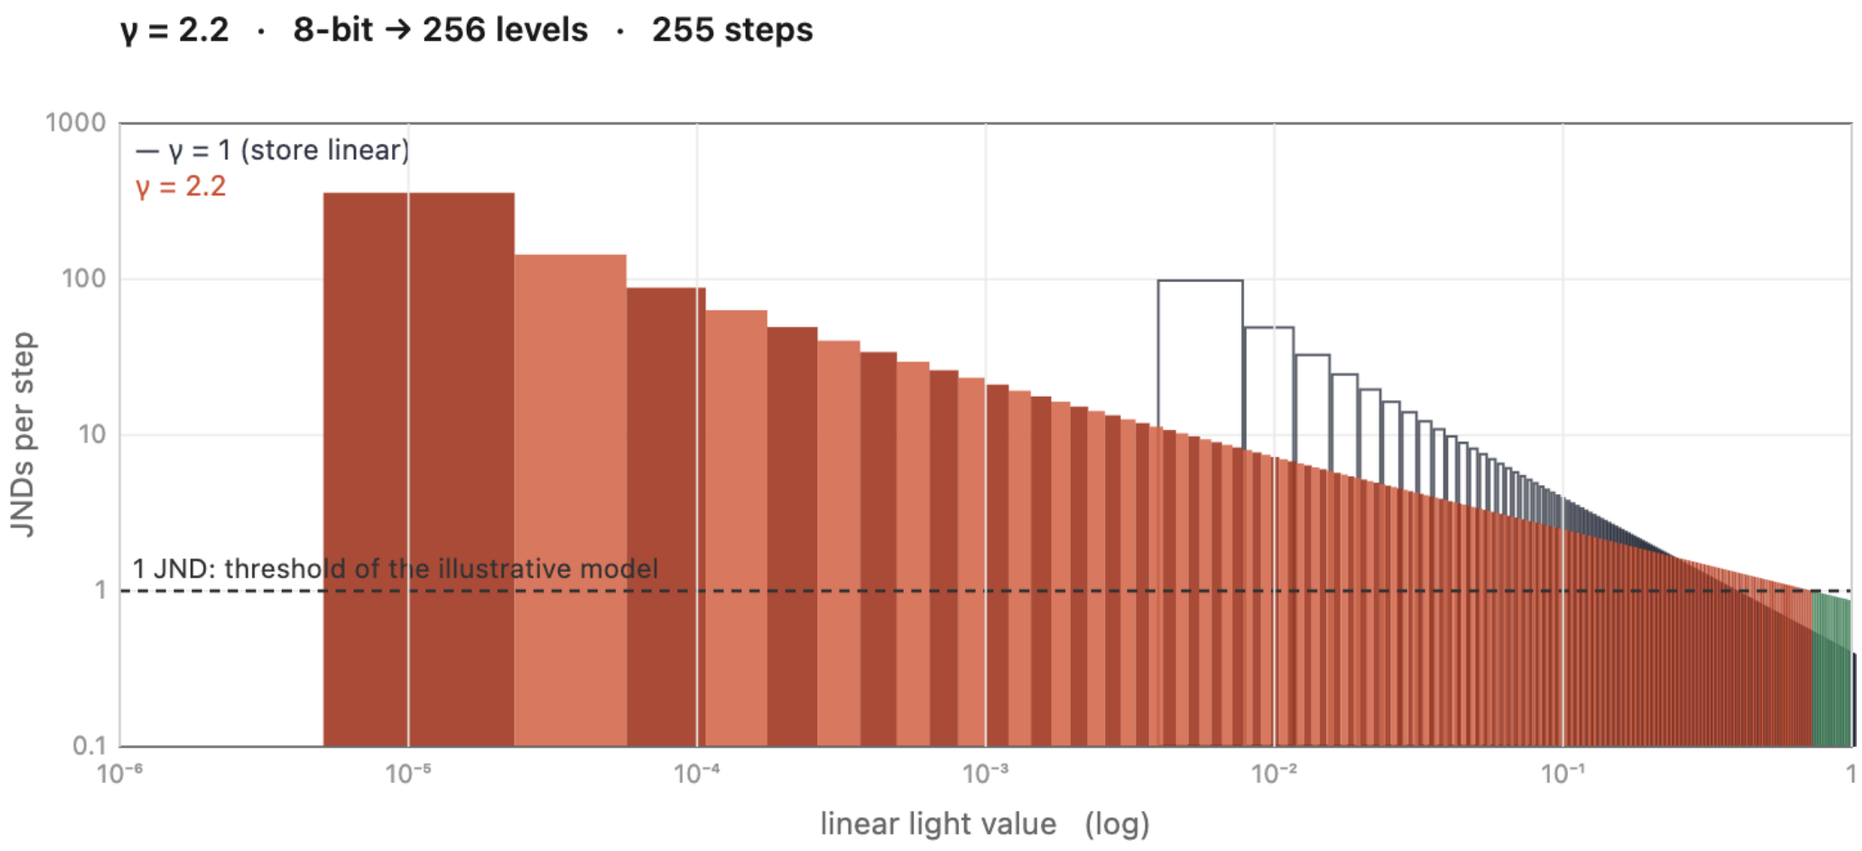
linear: min = 1/(2^N−1) = 3.9e-3 → 8.0 stops

gamma: min = (1/(2^N−1))^γ = 5.1e-6 → 17.6 stops

gamma buys ≈ γ× the stops (2.20×) for the same 8 bits, but noise, not bits, sets the real floor.



Demo: Just-noticeable differences per code step



γ = 2.2: 221/255 steps over 1 JND, worst ≈ 359 · γ = 1: 100/255, worst ≈ 100

First step (starts at zero, infinite under strict Weber, not drawn): ∞ JNDs; linear ∞

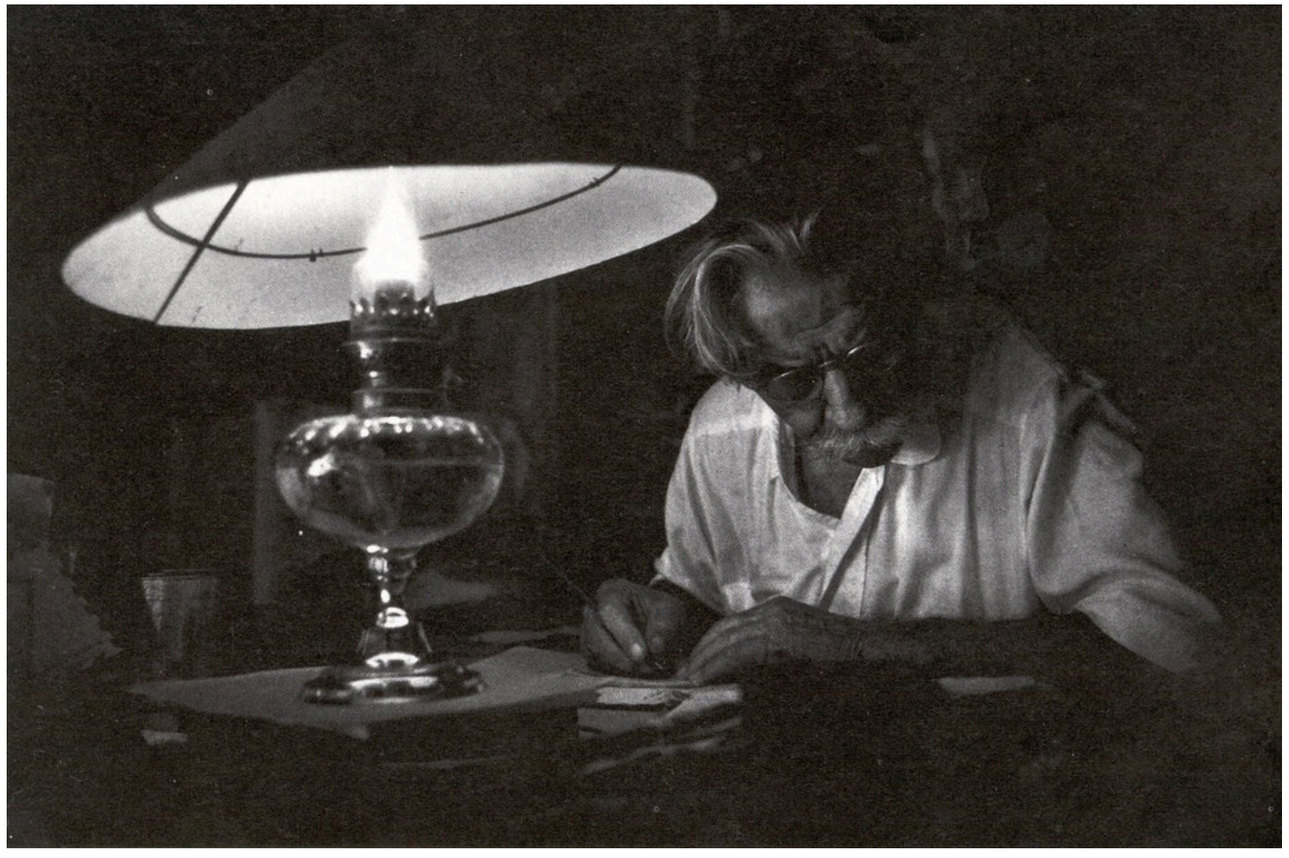
Questions?



Traditional Darkroom

Limited dynamic range can be good!

- **W. Eugene Smith photo of Albert Schweitzer**
- **5 days to print!**
- **Things can be related because the intensity is more similar**
- **Balance, composition**



Demo: When the limitation is the point

The medium's limitations are not only obstacles to beat — they can be the picture

Accepting limited contrast



shadows crushed to black → the lit water becomes a graphic

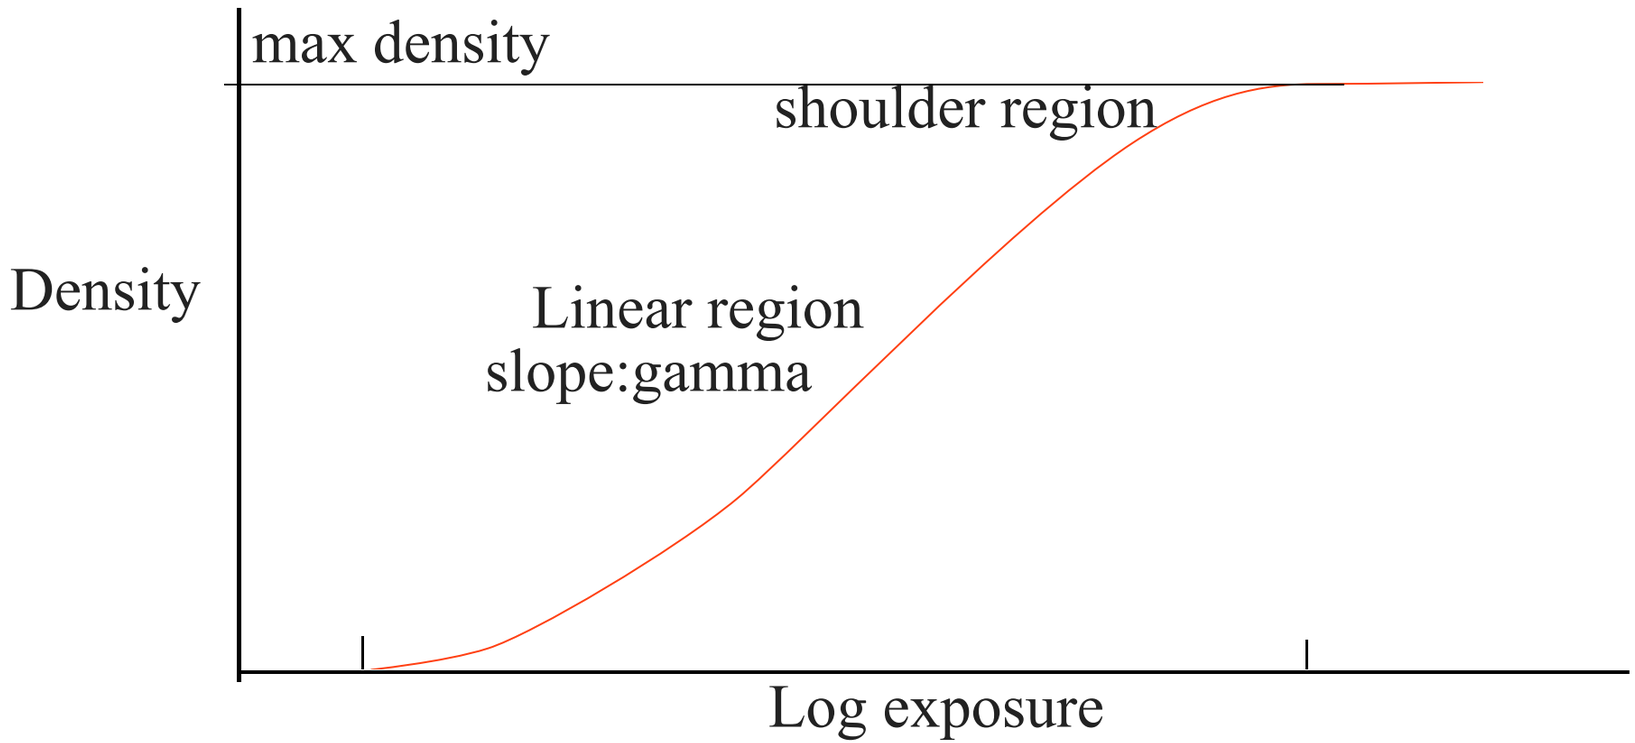
Accepting a single time slice



a fraction-of-a-second instant, held still to contemplate

Negative and response curve

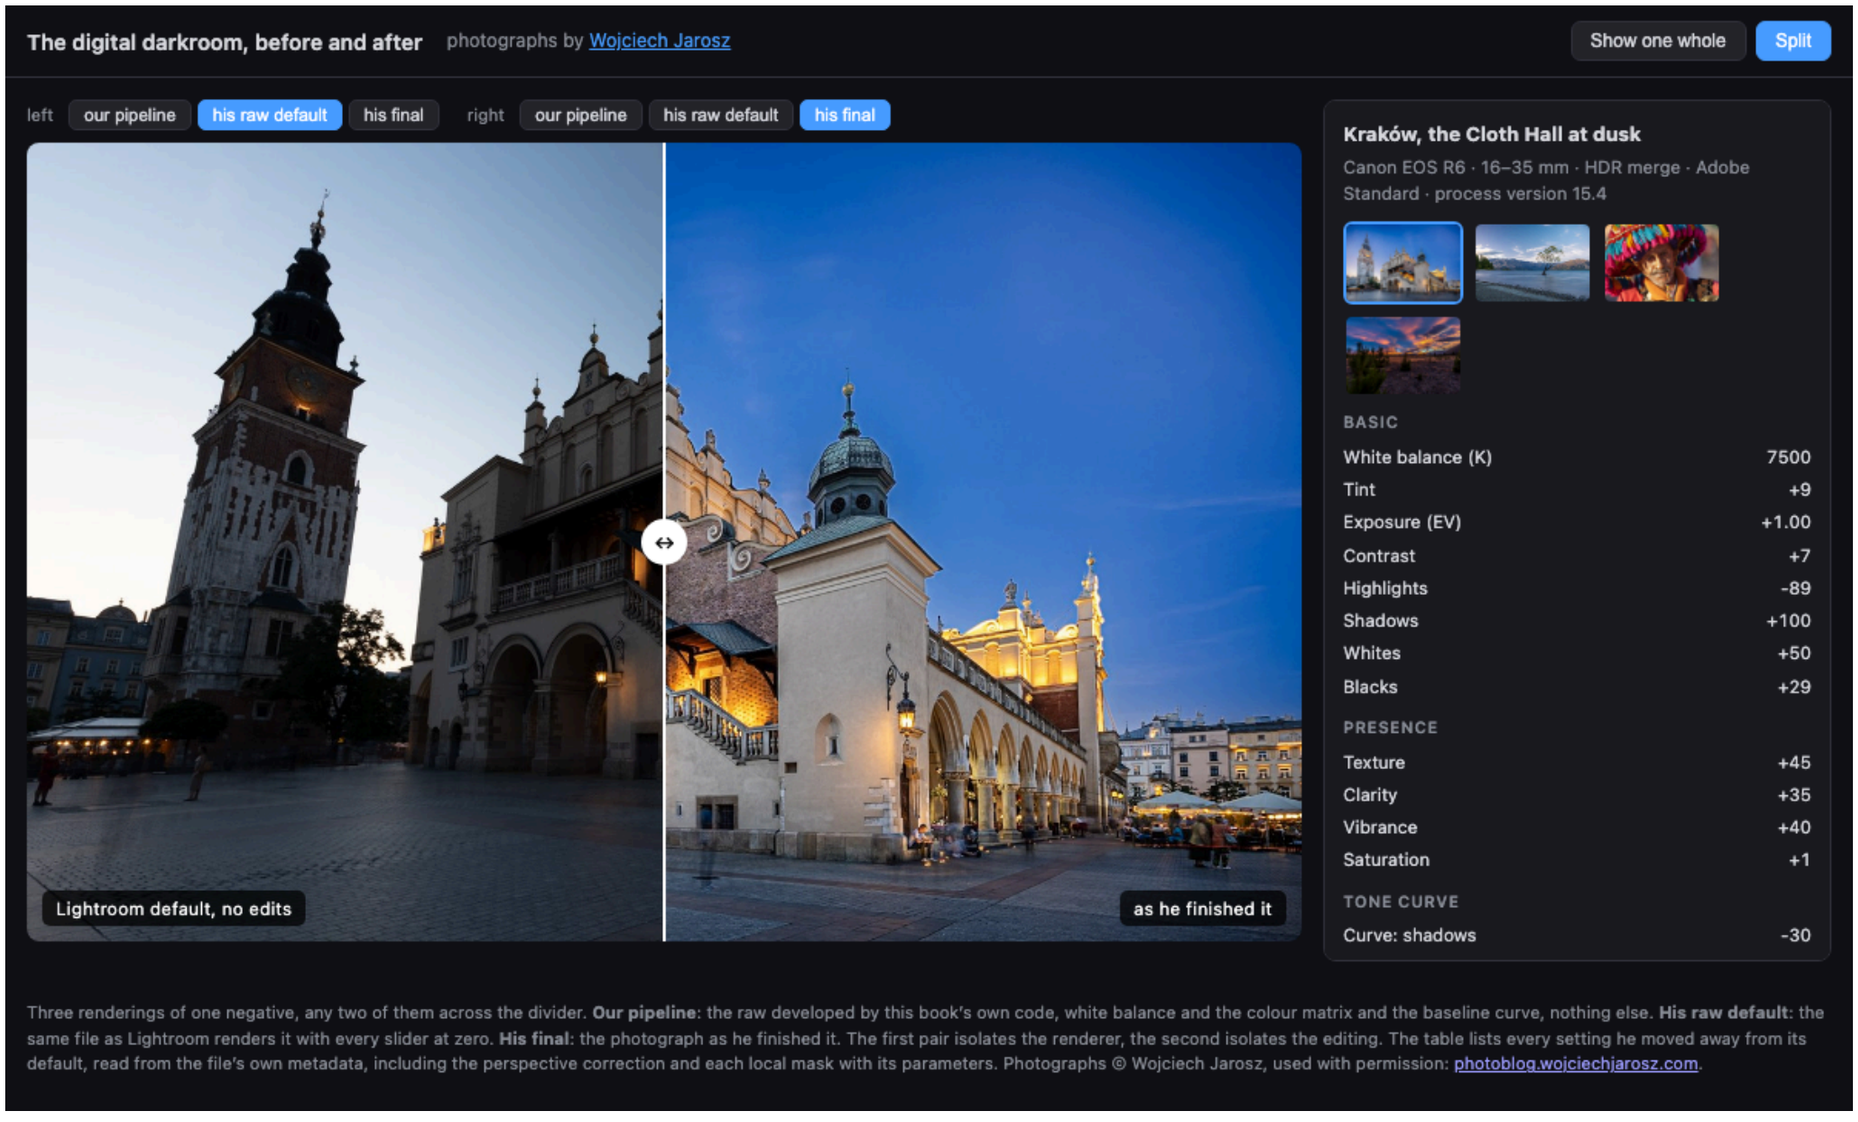
- **Negatives typically afford 3 orders of magnitude**
- **More than printing paper**
- **shoulder region provides a little more dynamic range (but less precision there)**



Response curve manipulation

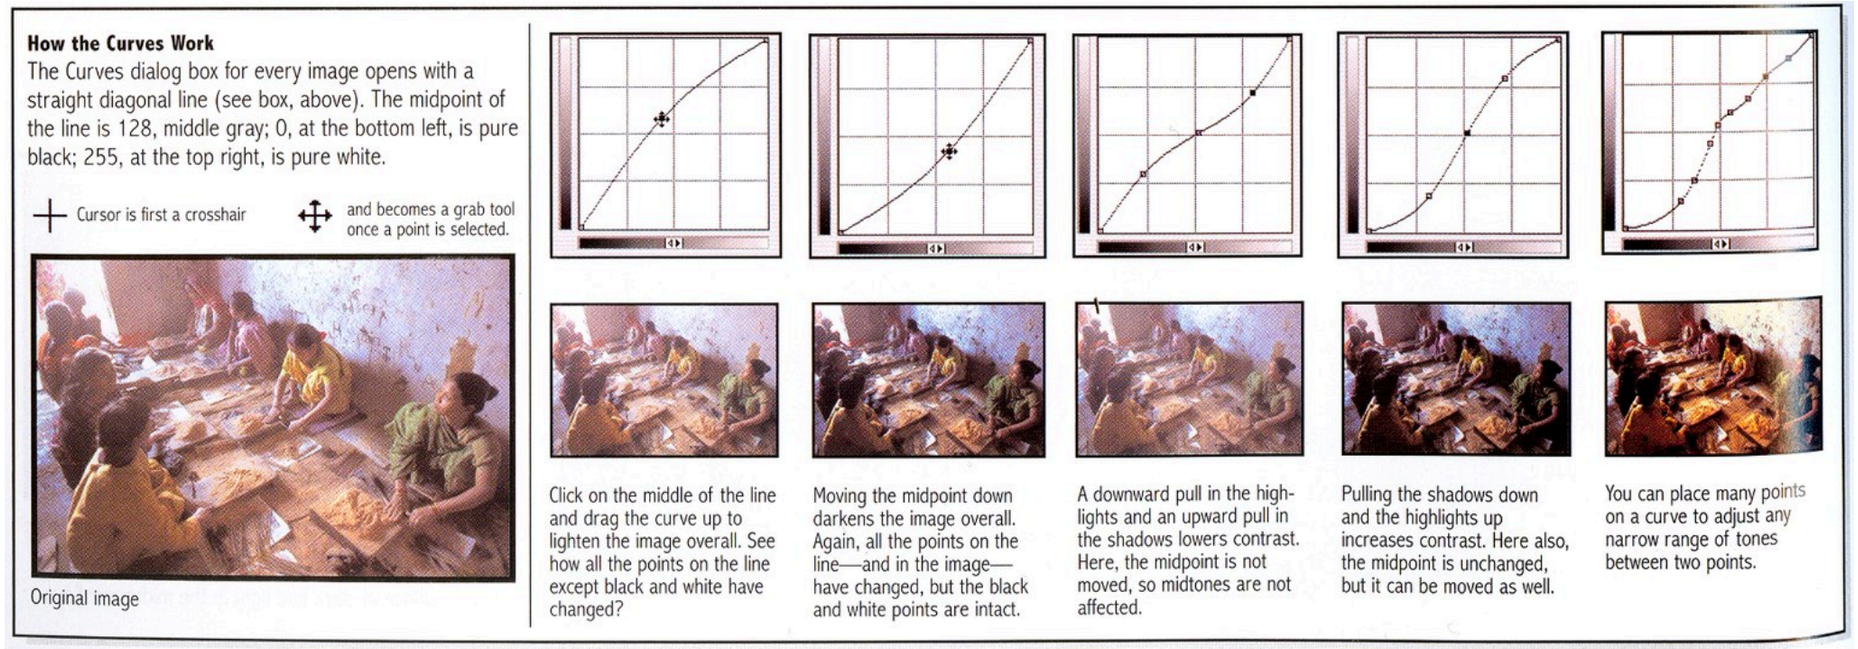
- **Traditional photography**
  - Chemicals and duration of development
  - Paper grade (~γ)
  - Flashing the paper before printing
  - Various chemicals on paper
  - Note: you have one curve for negative, one for paper
    Usually they have inverse gamma (but not strictly, enables contrast control)
- **Digital**
  - Curve tool

Demo: Four photographs, three renderings



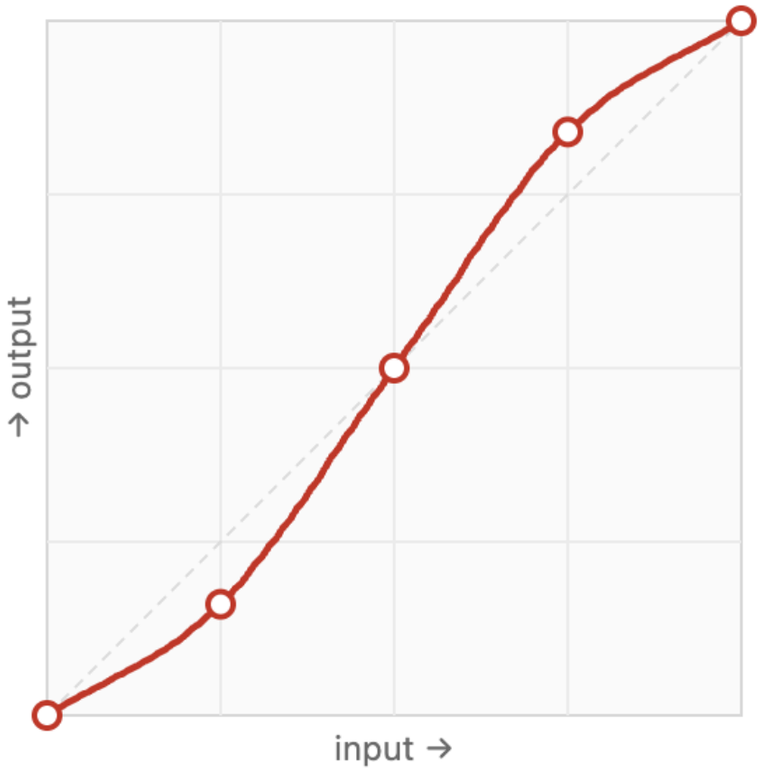
Photoshop curves

- **Specify an arbitrary remapping curve**
- **Especially useful for black and white**



From Photography by London et al.

Demo: Tone curve







Chemical curves: developt time

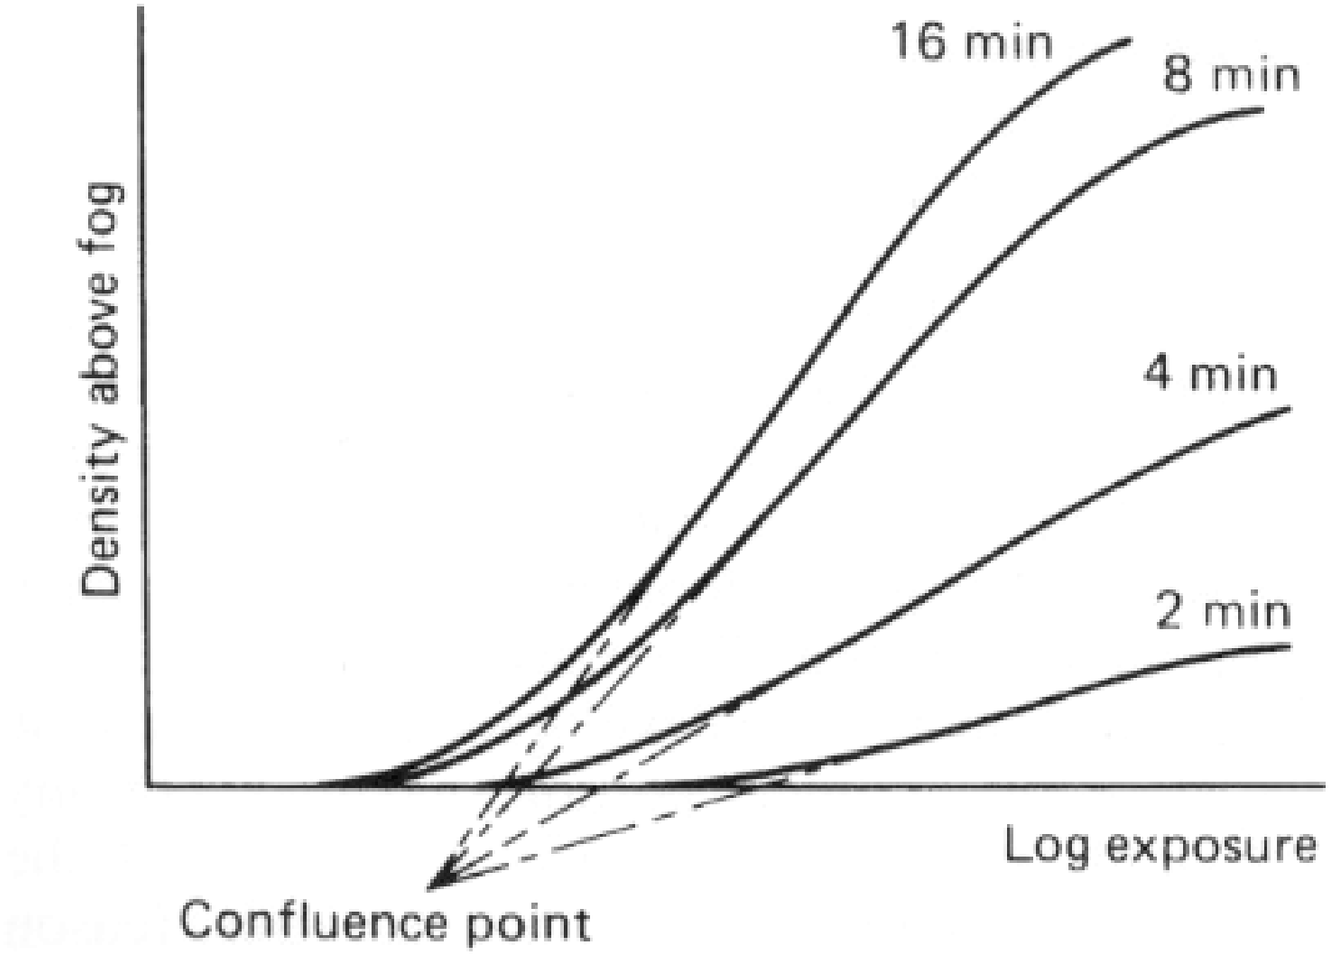
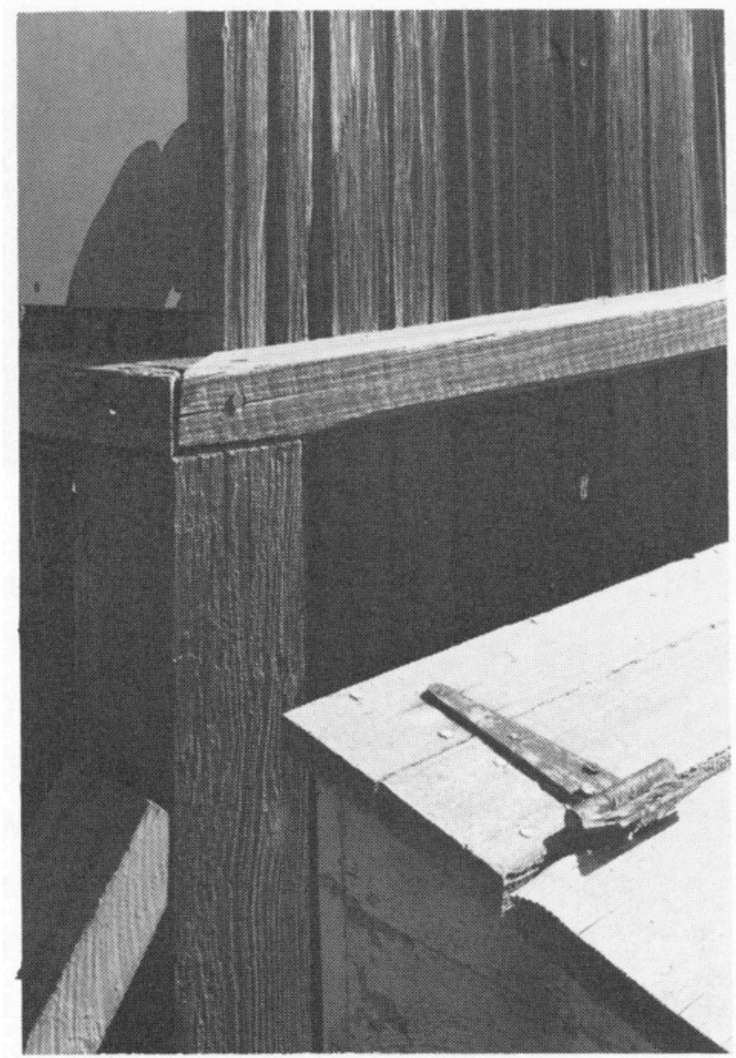


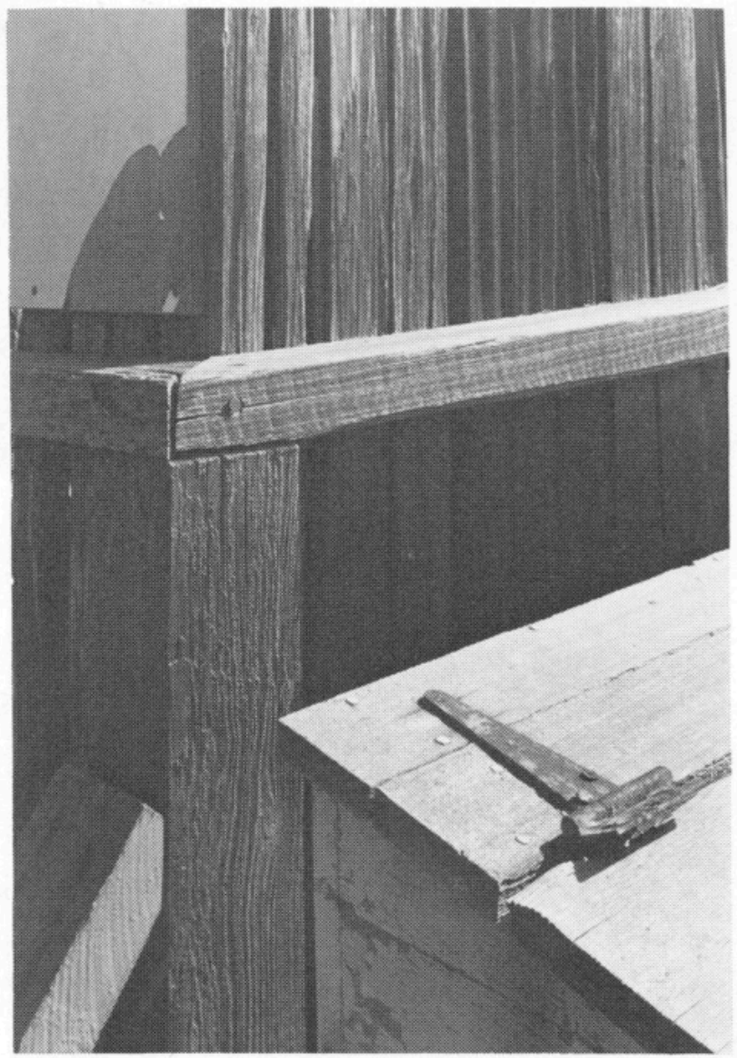
Figure 15.8 Effect of development time on characteristic curve of current sensitive materials

From The Manual of Photography, Jacobson et al.

Reduced development



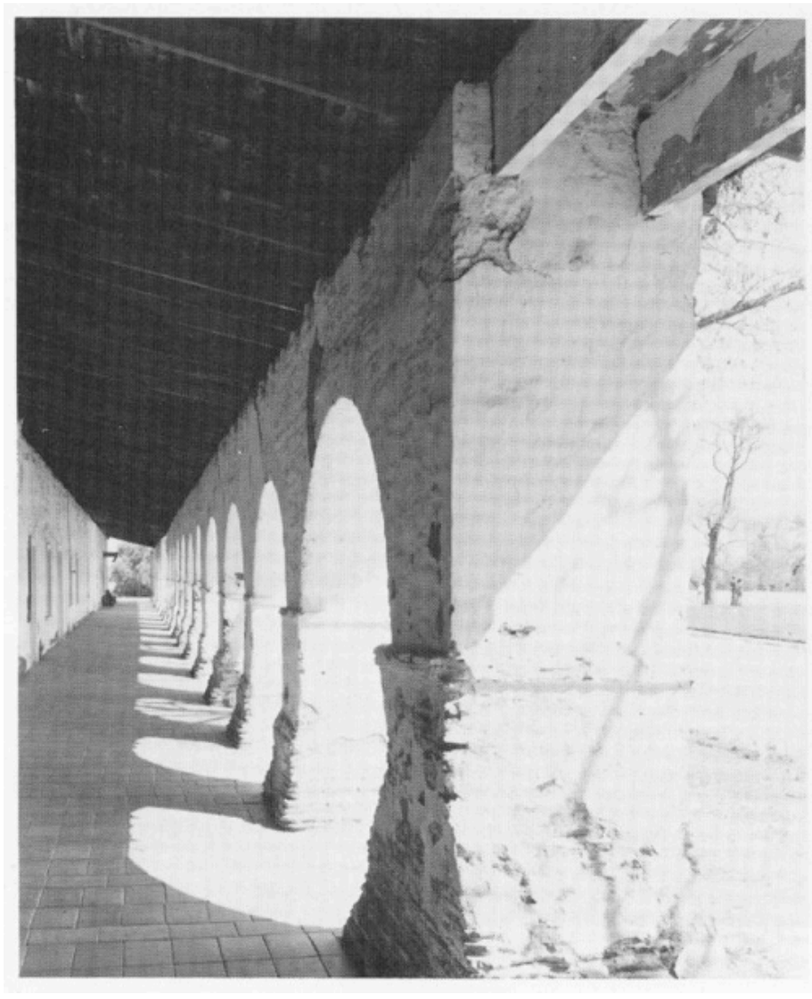
Normal development



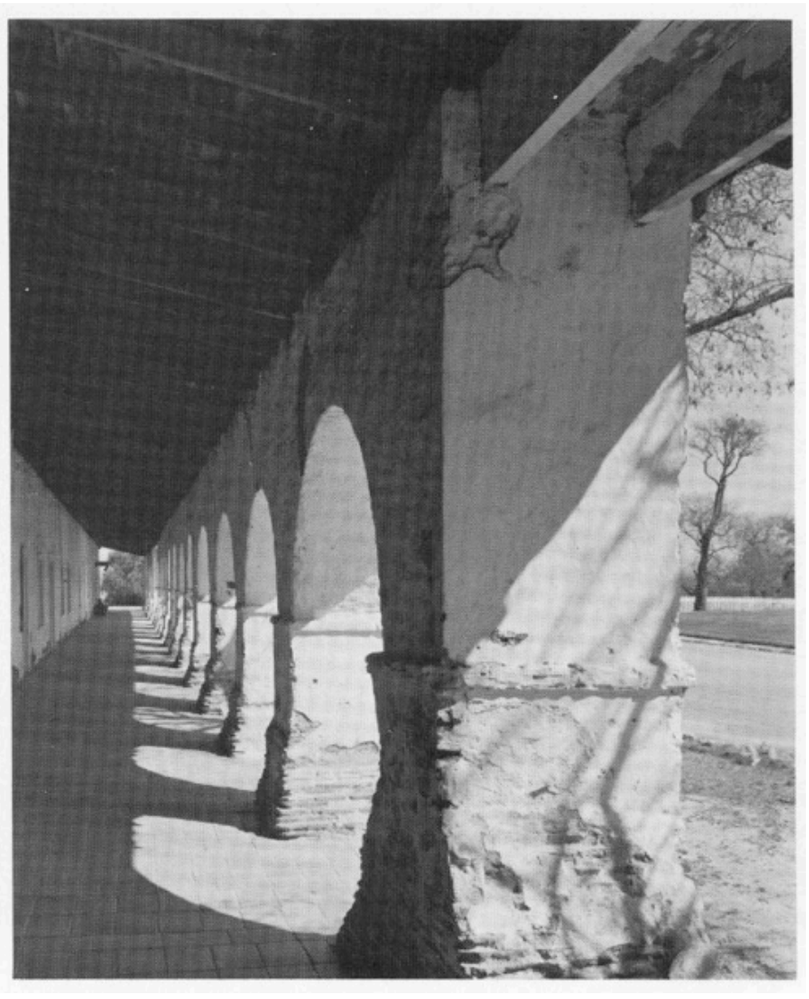
Contraction (short development)

Source: Ansel Adams

Two solutions



One development solution

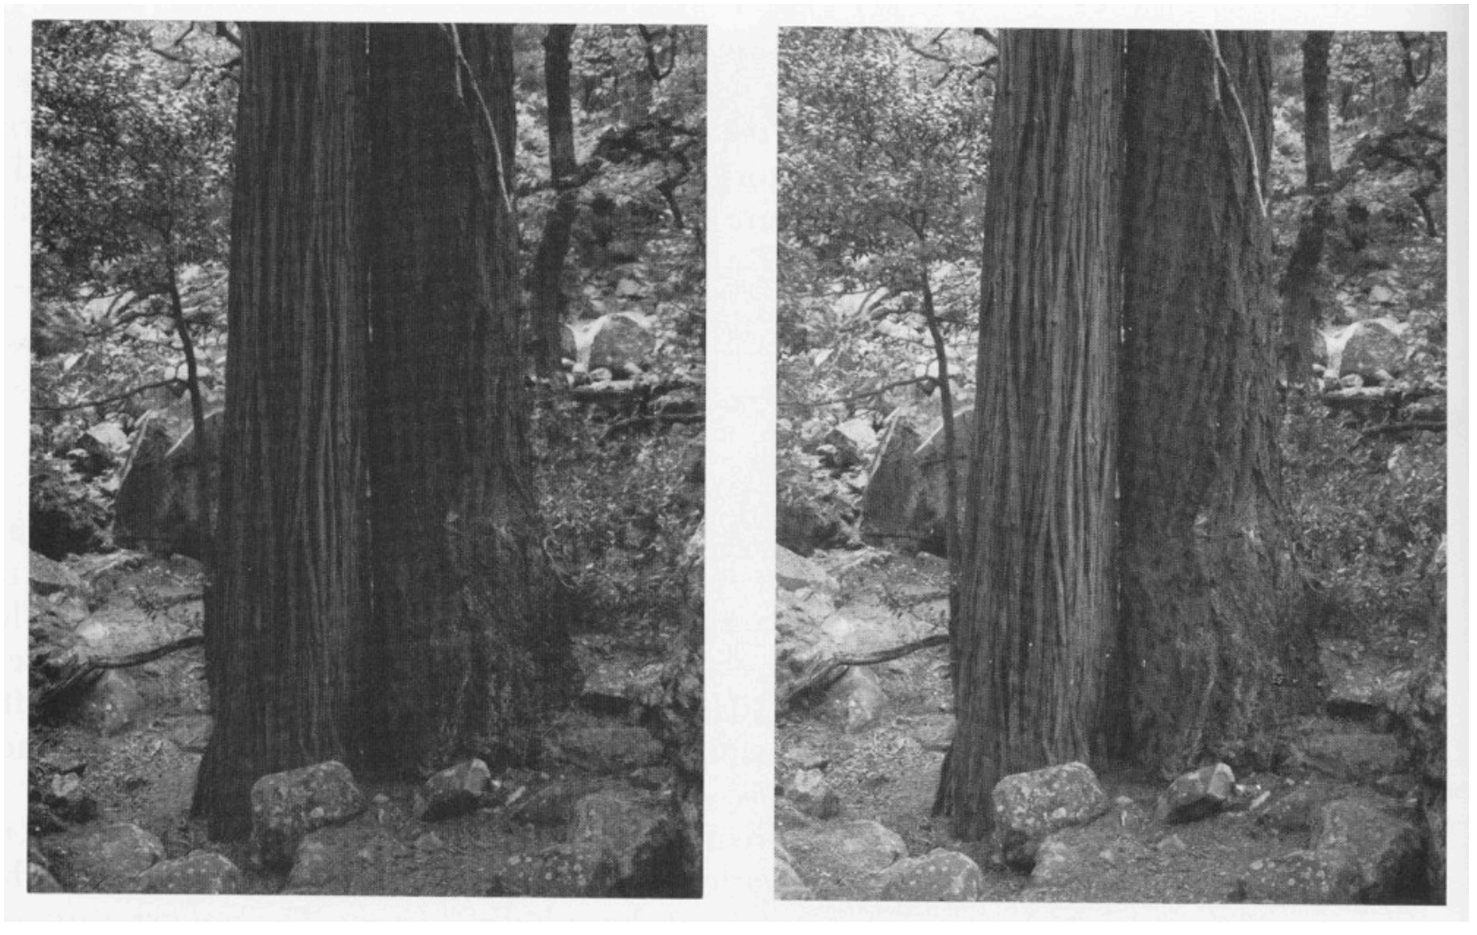


Two development solution:
the dark areas are the same,
but bright areas are different

Source: Ansel Adams

Pre-exposure

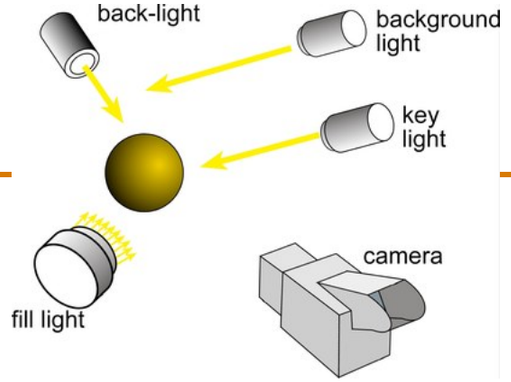
- **Briefly expose negative to a uniform light**
- **Raises the values of everything (in particular puts dark values above the low-contrast toe of response curve)**



Source: Ansel Adams

Lighting

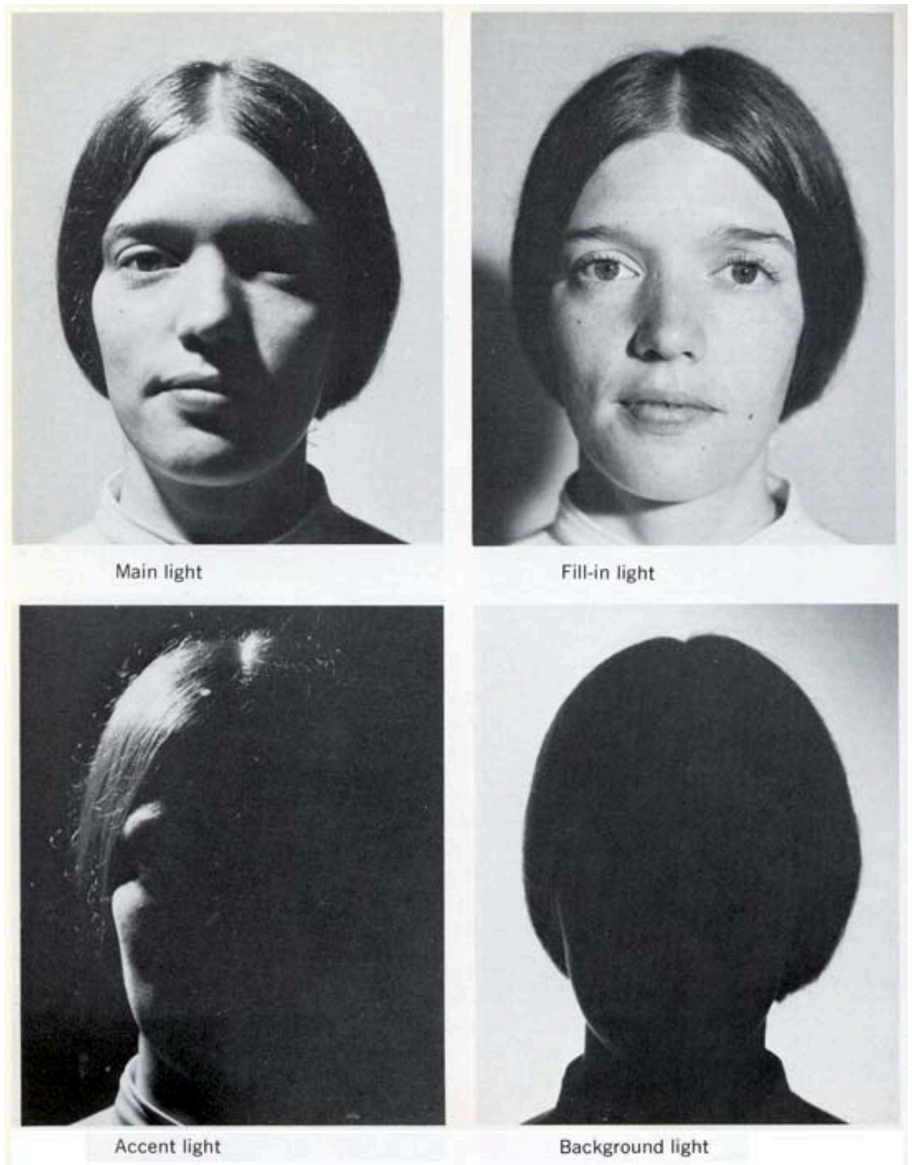
- **E.g. 3-point lighting**
  - Reduce dynamic range
  - Emphasize silhouettes ! 3D cues



- **Goals of lighting:**
  - Manage dynamic range
  - Reveal shape, layout, material
  - Tell story

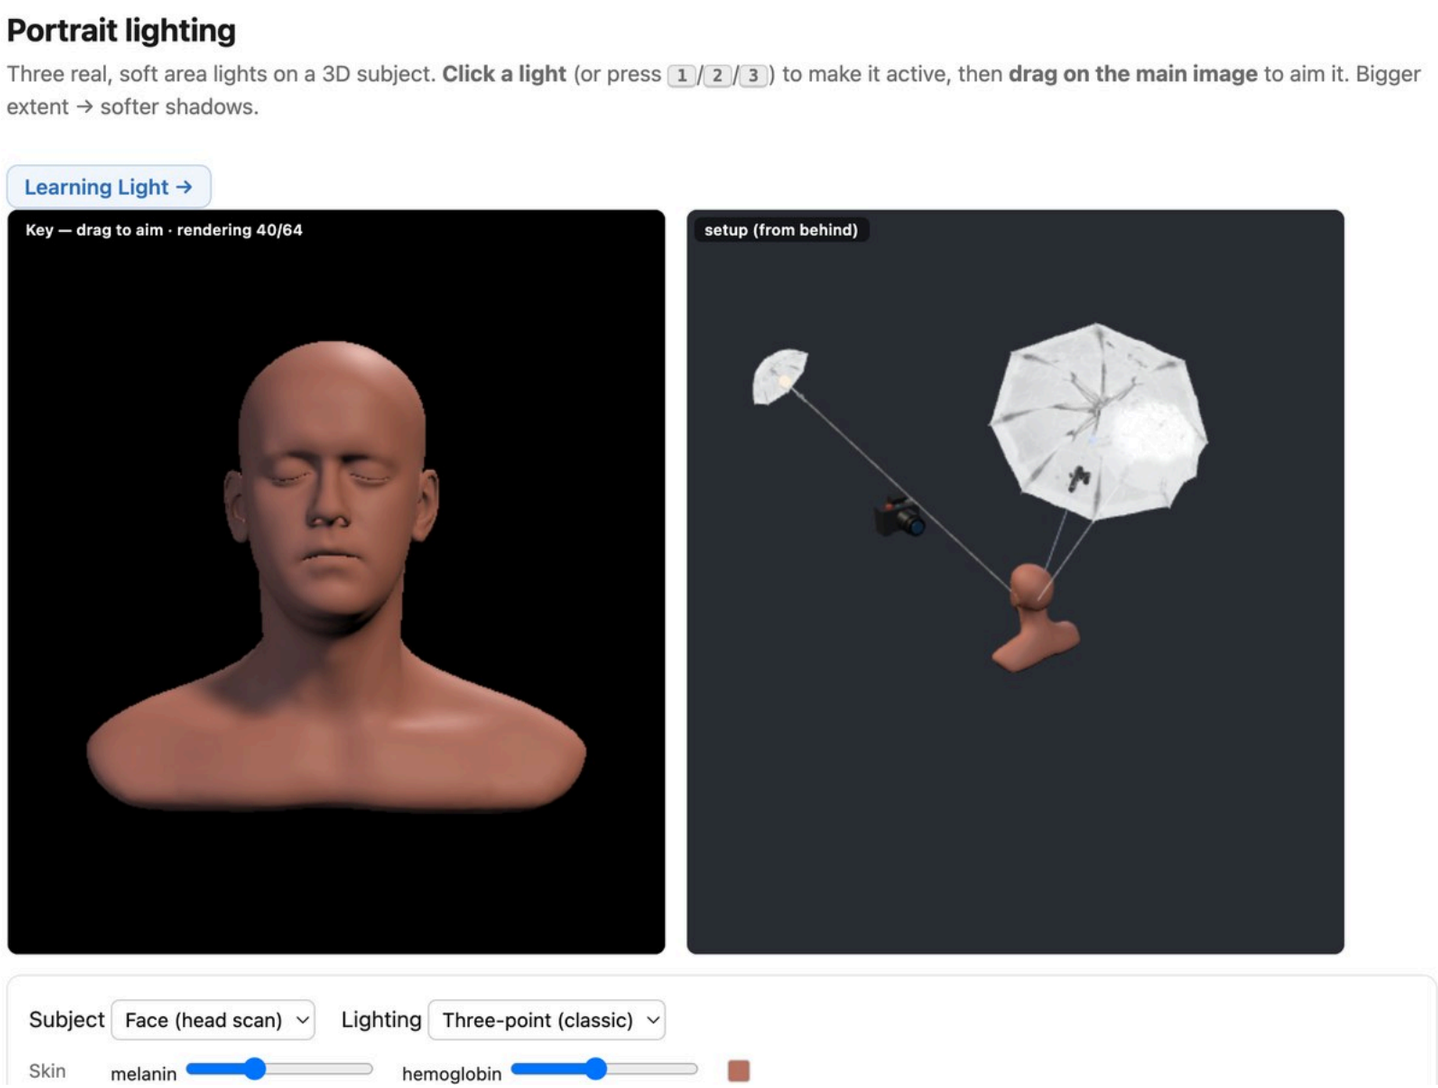


Portrait lighting





Demo: Portrait-lighting simulator



Without flash



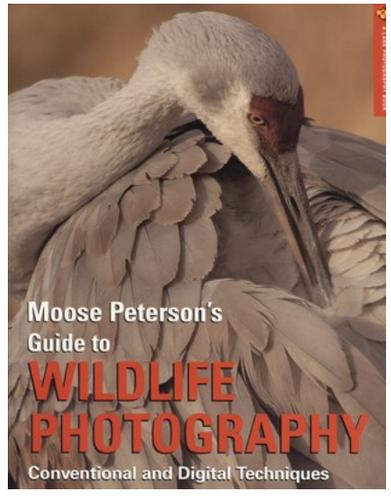
With fill flash



Fill-in flash





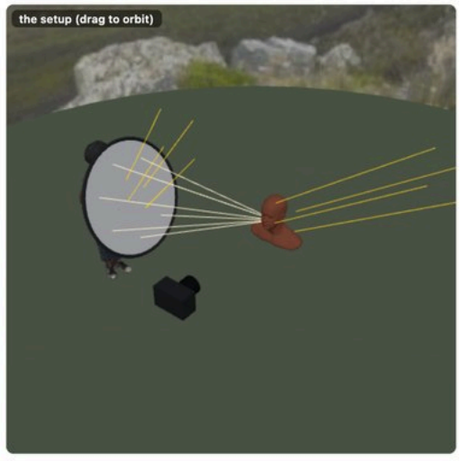


Demo: Reflector: bouncing the sun

The reflector: bouncing the sun

A face is lit by the **sun**, a single hard, parallel light. It keys one cheek and drops the other into shadow. A photographer's assistant holds a large white **reflector** to the side: it catches the sun and throws it back as a big, soft second source, filling the shadow without adding another lamp. This is global illumination put to work: the fill is the sun, arriving by way of one bounce. Move the assistant around the subject and in or out, and watch the fill soften and shift.





Preset: reflector fills the shadow

copy link

REFLECTOR (THE ASSISTANT)

- Hold up the reflector
- azimuth -62°
- distance 1.70
- size 0.63
- surface: white (soft neutral fill)
- assistant yaw -29°
- assistant back -0.20

SUN

- azimuth 55°
- elevation 38°

The sun is a hard, parallel source. High and to one side it makes a strong key and a deep shadow, which is exactly what the reflector is there to open up.

SUBJECT SKIN

- melanin 0.50
- hemoglobin 0.45

BACKGROUND

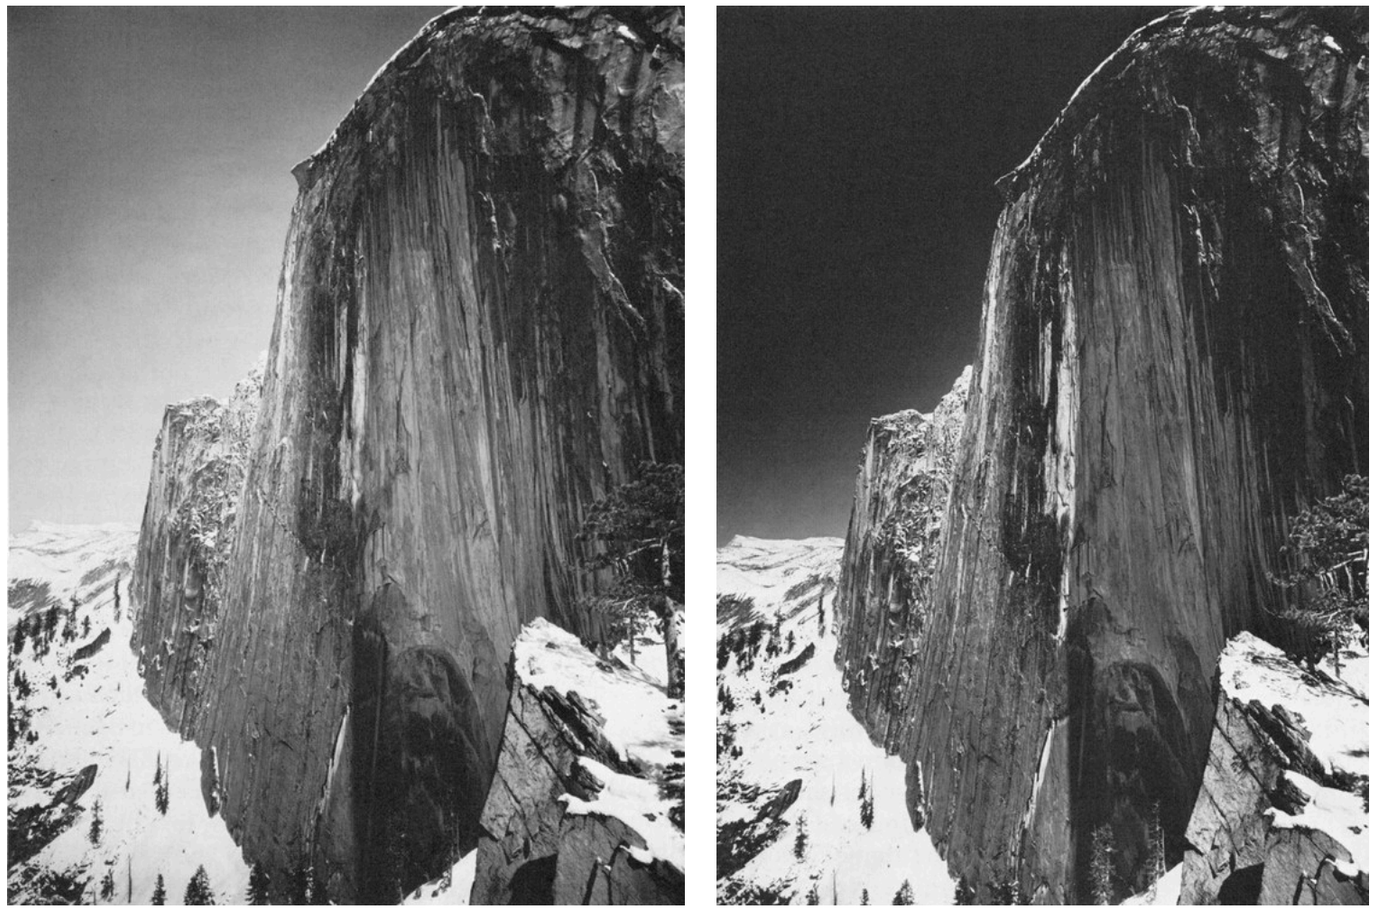
- scene: coastal

An outdoor environment map lights the scene faintly all around; the sun and the bounced reflector do the real modeling.

A reflector adds no power of its own: it simply gives the sun a second, much larger place to come from. The catch is a real one, the bigger and closer the reflector, the softer and stronger the fill, but move it too far and it does almost nothing. A white reflector fills neutrally; a gold one warms the skin (photographers use it for sunlit portraits); a silver one is brighter and more specular. The same physics is at work when you bounce a flash off a wall (the wall is the reflector) or shoot into a studio umbrella.

Filtering: black and white

- **Red/orange/yellow filters darken the sky**



No filter

With red filter

Source: Ansel Adams

Graduated neutral density



No filter: sky is too bright



Vertical neutral density gradient

Graduated ND & landscape

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



Portage Lake, Horizontal

http://www.artwolfe.com/

Graduated ND & landscape

- **Art Wolfe: Here I had to use a combination of filters and settings that greatly reduced my chance of success. I used my zoom to bring in Denali and the moose. A polarizing filter brought out the rich colors of the tundra and darkening the sky and a graduated, neutral-density filter to bring the entire scene into the same exposure.**



Denali Autumn

http://www.artwolfe.com/

Alternative: Black card

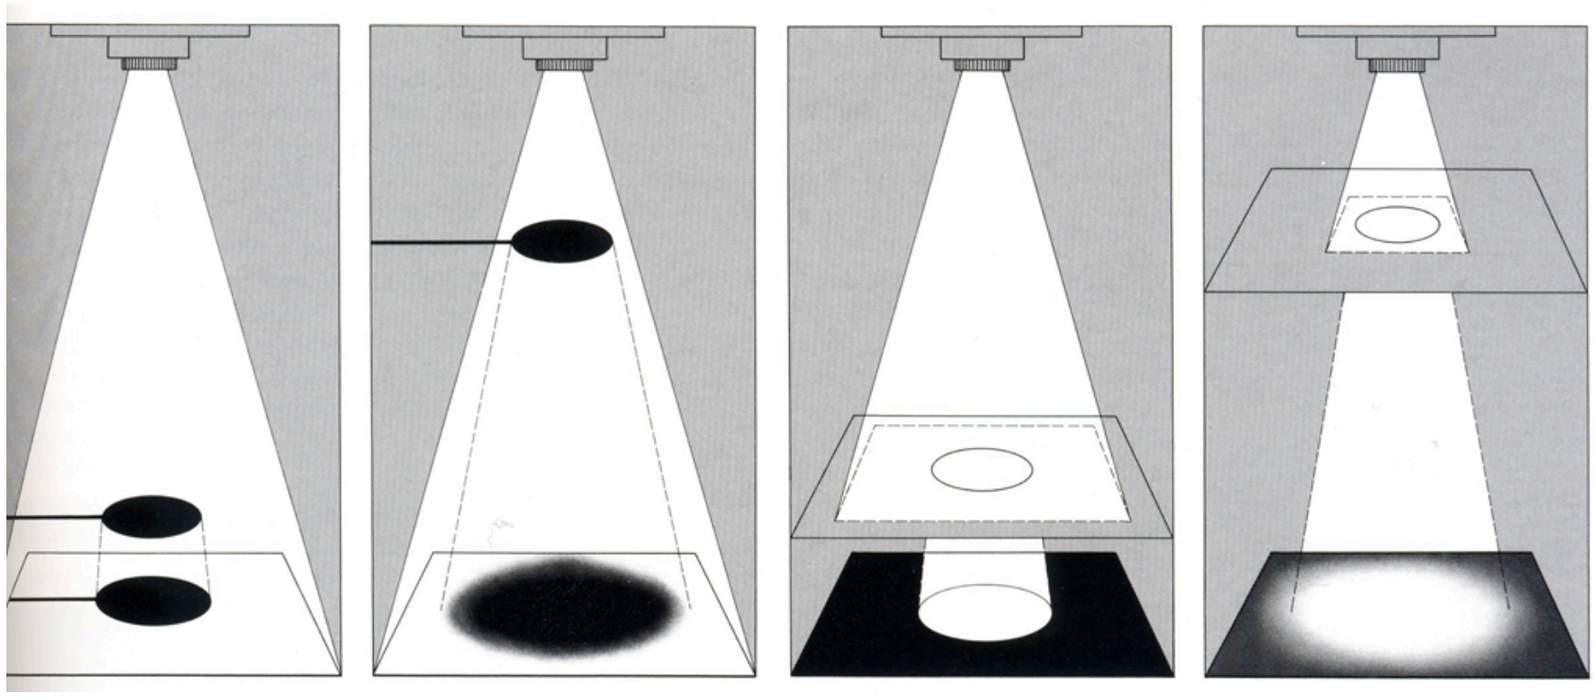
- **http://hanjies.blogspot.com/2009/10/black-card-photography-part-i.html**





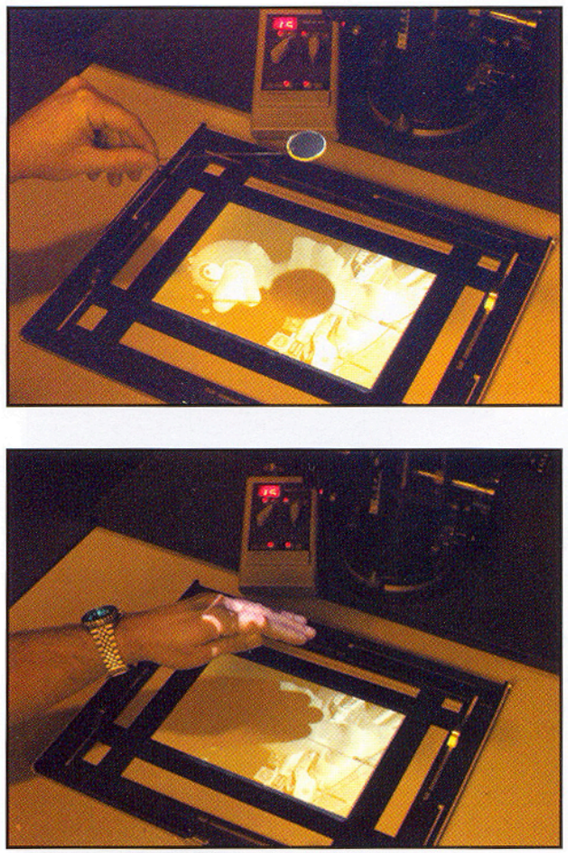
Dodging and burning

- **During the print**
- **Hide part of the print during exposure**
  - Makes it brighter

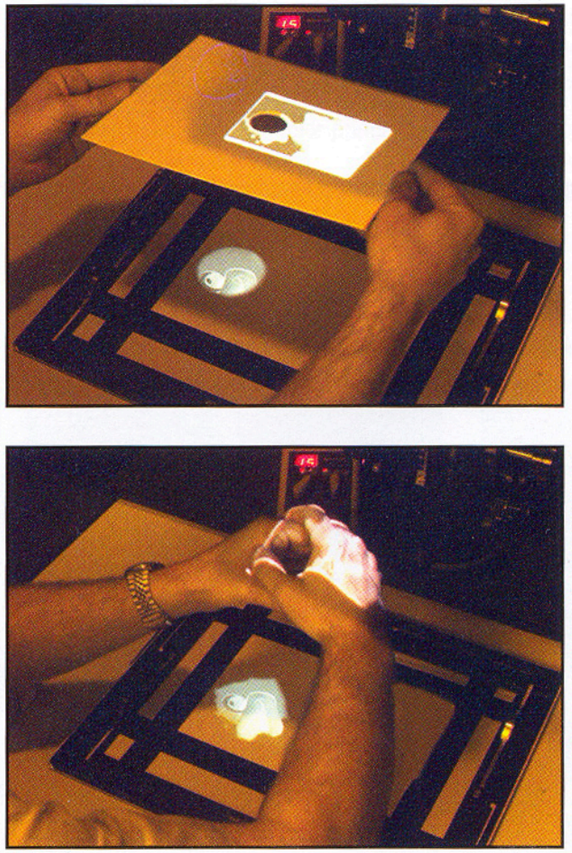


From The Master Printing Course, Rudman

Dodging and burning



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

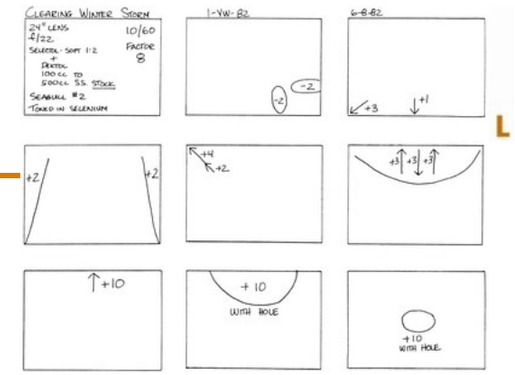


Burning *adds light after the basic exposure to darken an area.*

From Photography by London et al.

Dodging and burning

- **Must be done for every single print!**



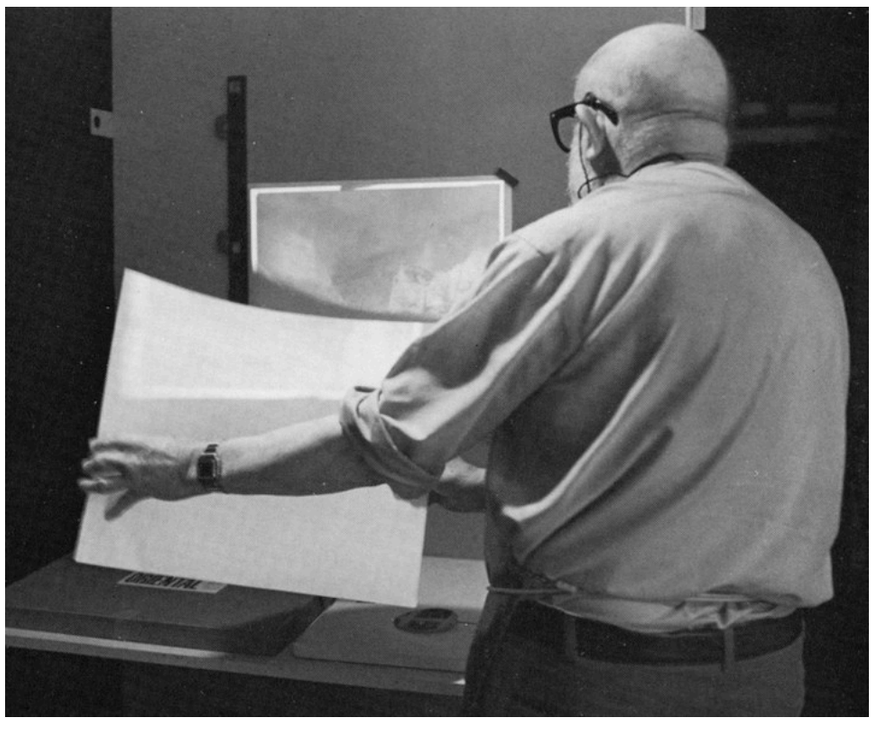


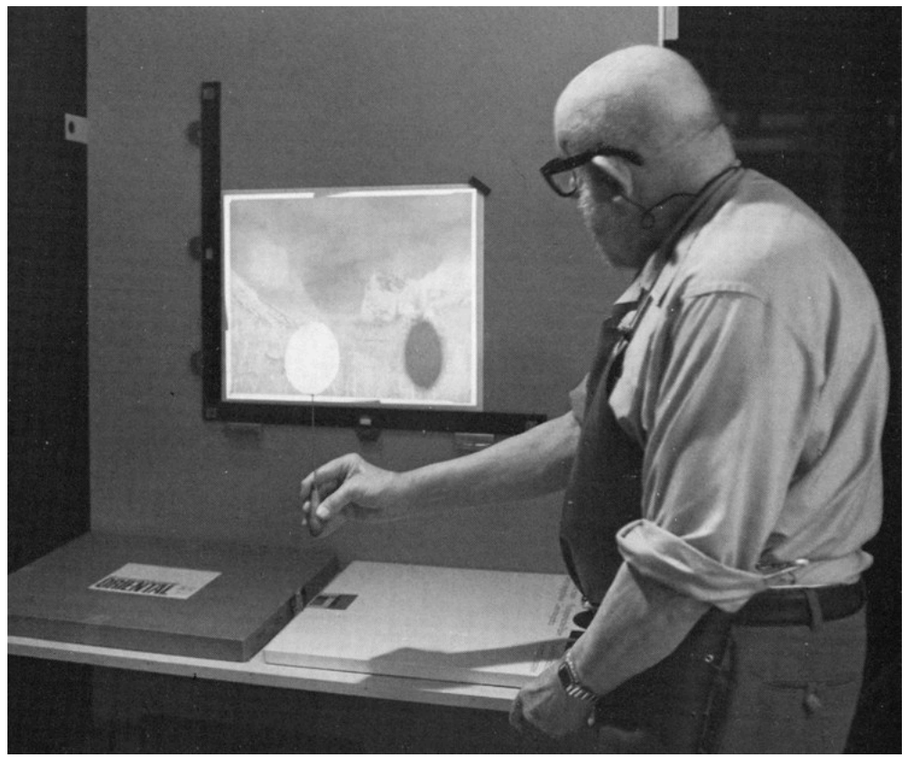
Straight print



After dodging and burning

Dodging and burning



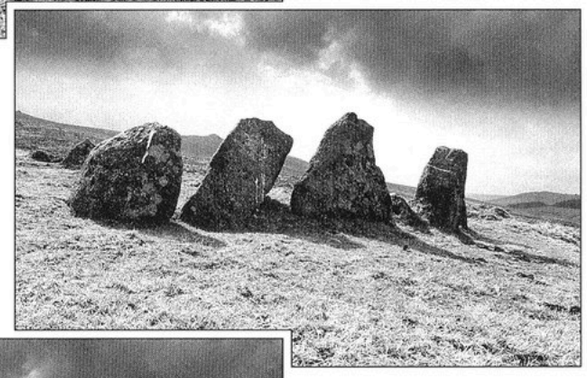


Source: Ansel Adams

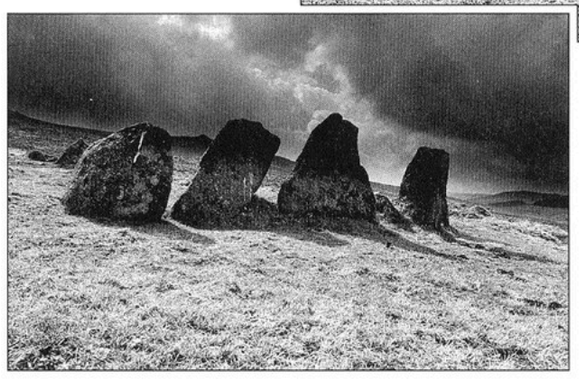
Dodging & burning is difficult!



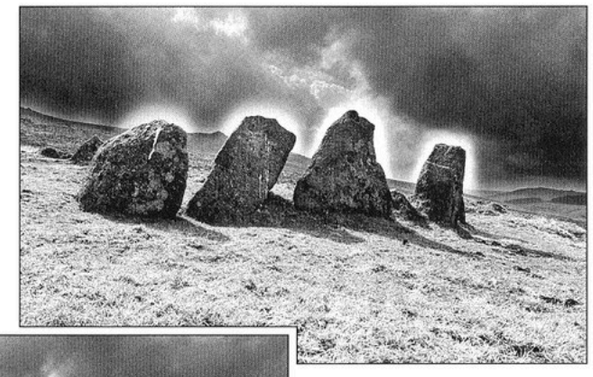
A *The straight work print without additional burning-in.*



B *This print shows the result of trying to mask off the foreground by using a moving card. An even more obvious light band will appear in the sky if the card is not kept moving.*



C *In order to remove the light band in fig B, the mask has been lowered. This, however, has caused parts of the horizon to become black.*



D *The halo effect, here deliberately exaggerated, resulted from dodging the stones during the second exposure while burning-in the sky.*



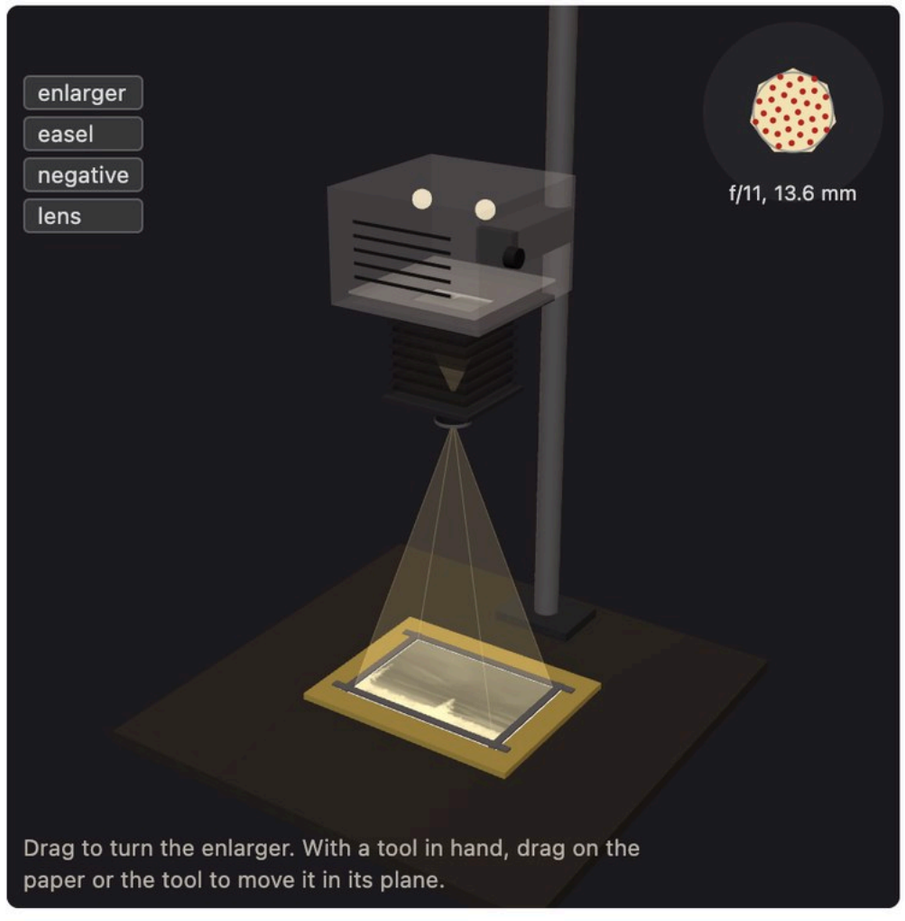
E *It is very difficult to cut a dodging card with precision, especially for a relatively small print like this. As a result, parts of the sky at the horizon are white, although careful spotting can disguise this problem when it is small. But parts of the mid-grey hill tops have gone jet black, which is less easy to rectify.*

Demo: Darkroom simulator

scene Vine sunset or open a .hdr Choose File No file chosen

1. Camera | 2. Film development | 3. Print

Vine sunset, 714 x 480, HDR by Paul Debevec



Print | Zones | Exposure | Negative | Mask



Rendered with the paper's white as the screen's white and Dmax 2.05 as 1.0% of it.

Questions?

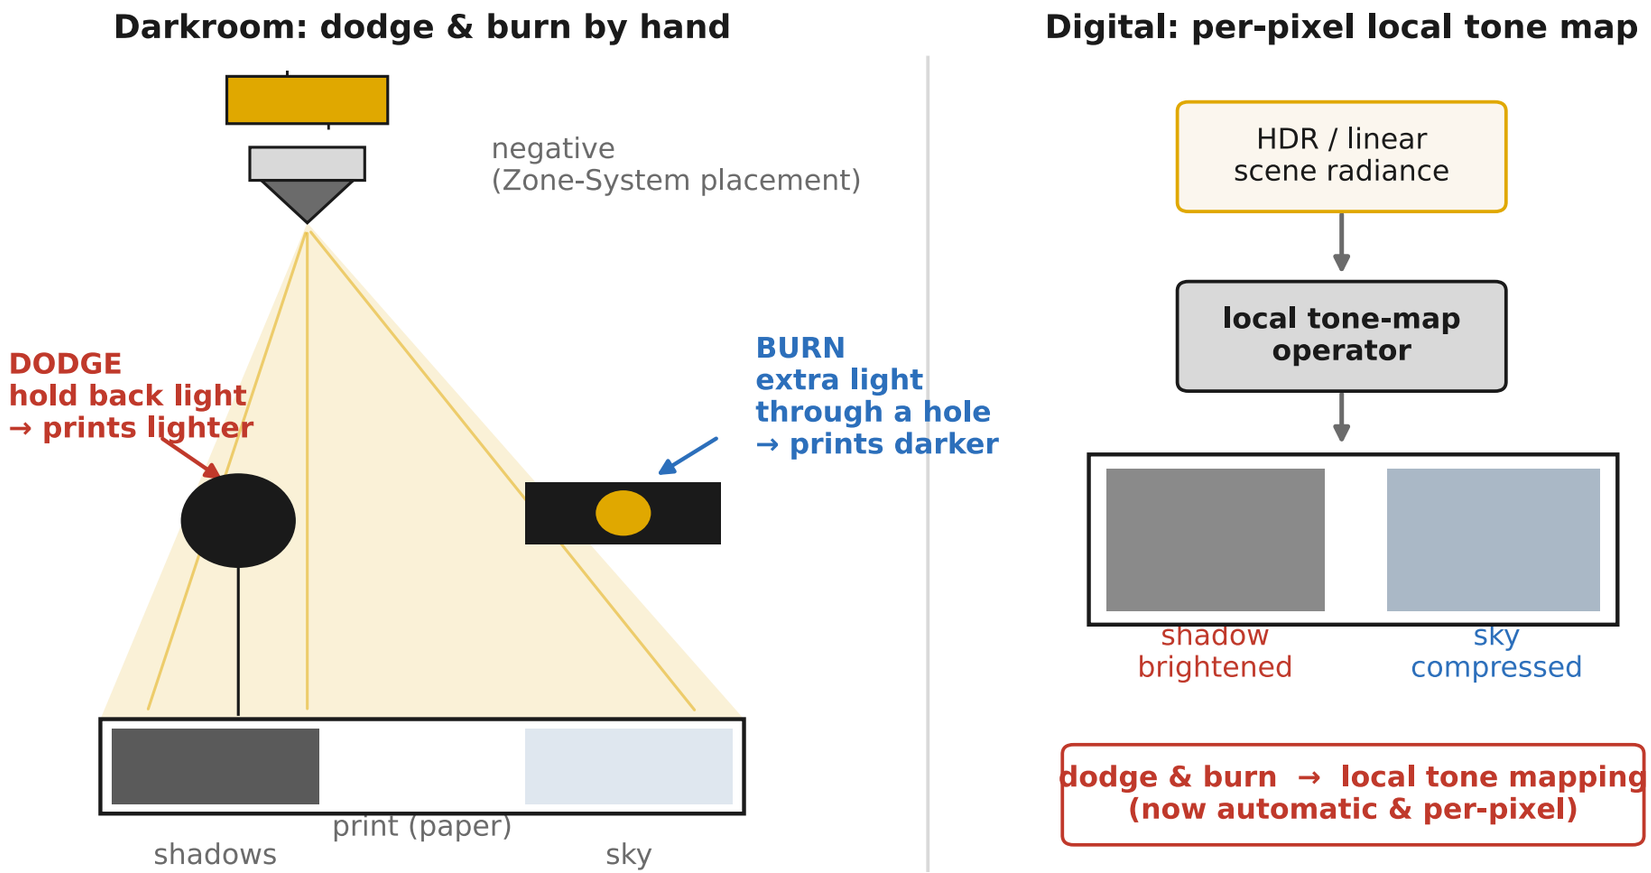
- **Gordon Parks**



Digital dodge-burn and graduated ND

- **Use adjustment layer and gradient tool**
  - Use curve adjustment layer
  - Modulate its effect using the layer mask
- **Just paint in black**
  - On a separate layer
  - With a low opacity
- **Multiple exposure photography**
  - Use a tripod
  - “Bracket” your exposure
  - Stack exposures as layers in photoshop
  - Use layer masks to select which region comes from which exposure

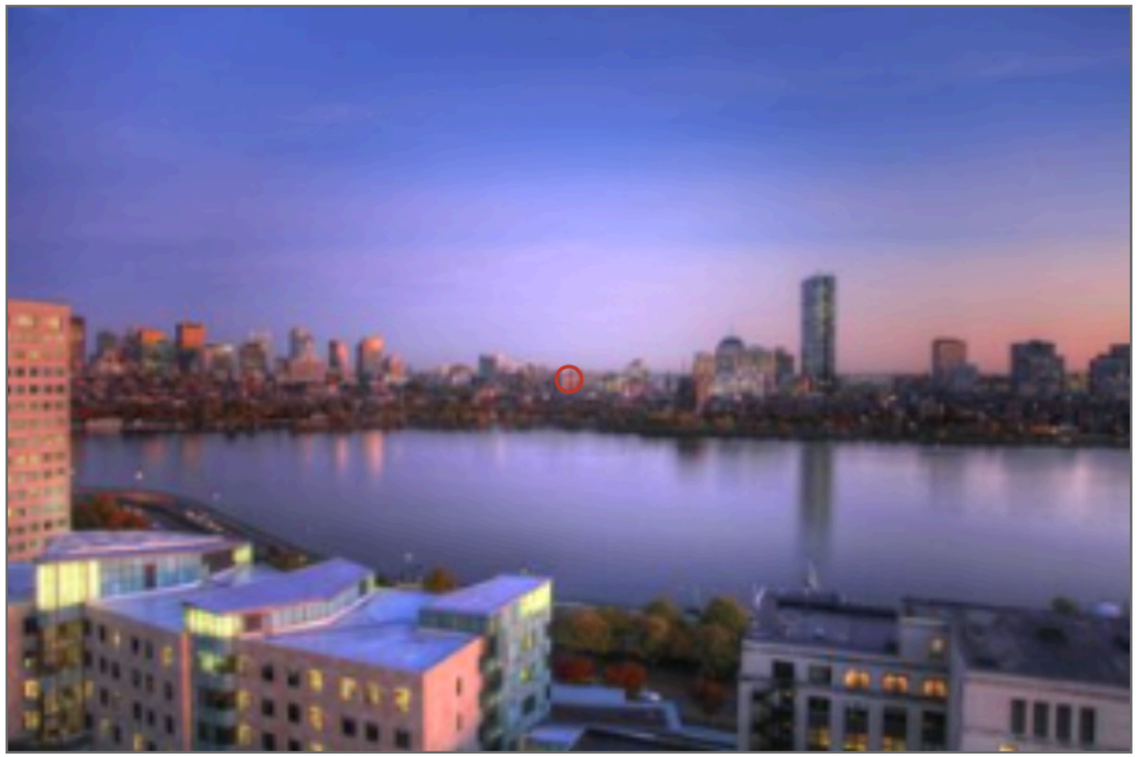
Figure 18.1.4: The darkroom beside its digital descendant



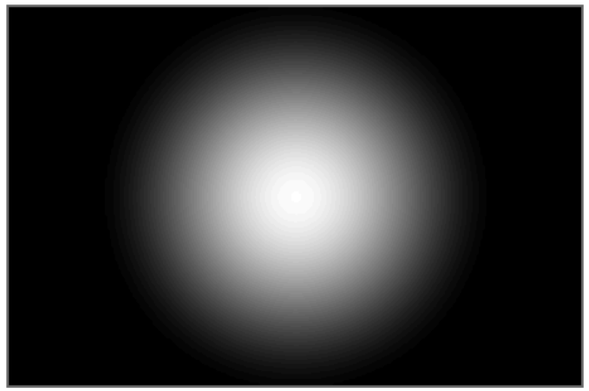
Under the enlarger, a card or hand dodges the shadows so they print lighter and burns the sky so it prints darker; a digital local tone map does the same region by…

Demo: Local point operations

+1.2 EV at full mask



mask



drag the handle on the image

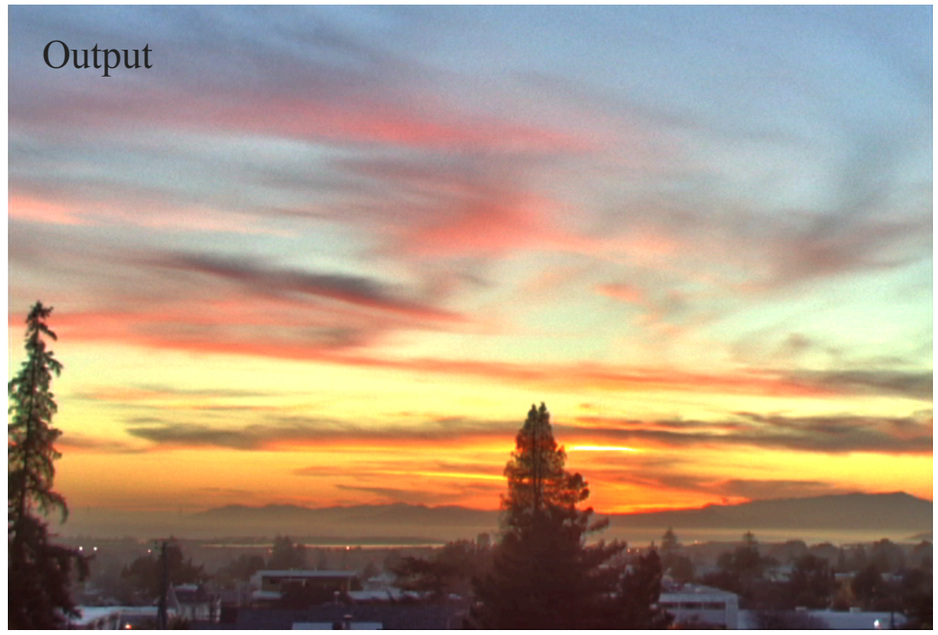
What have we learnt?

- **Log is good**
- **Luminance is different from chrominance**
- **Separate components:**
  - Low and high frequencies
- **Strong edges are important**

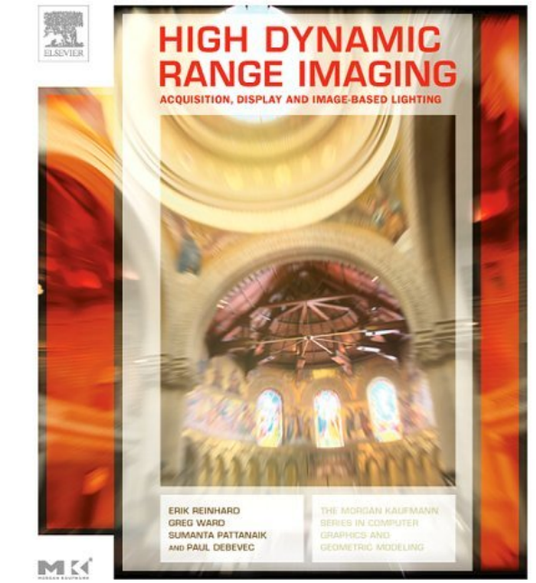


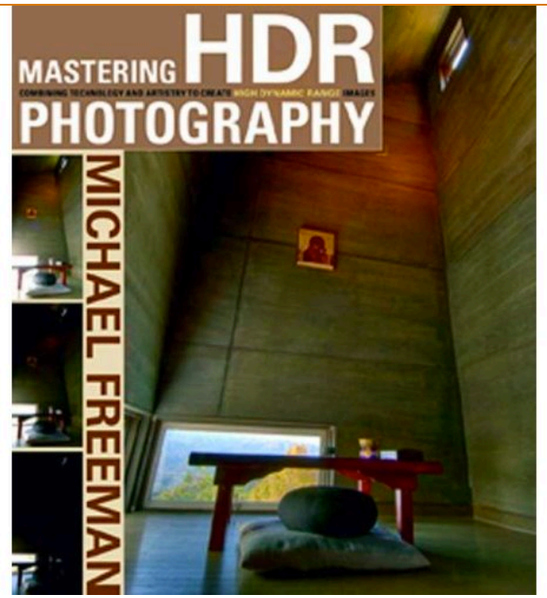


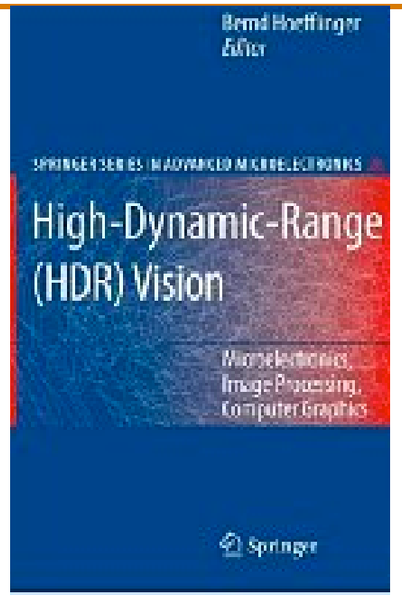


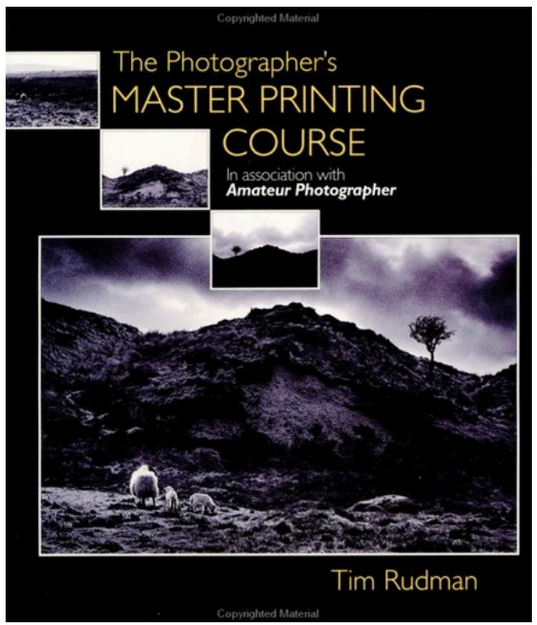


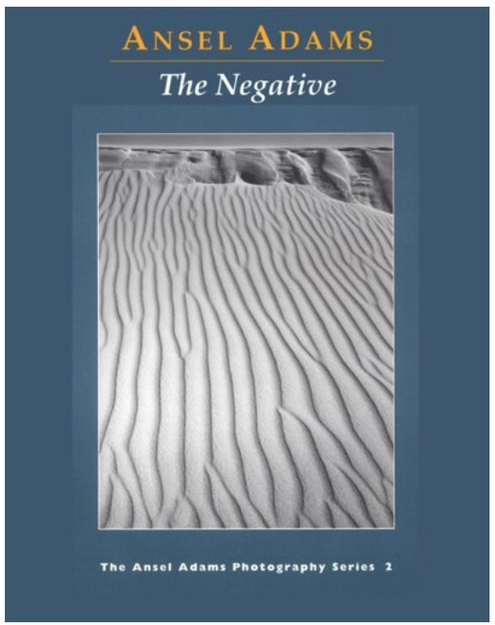
References

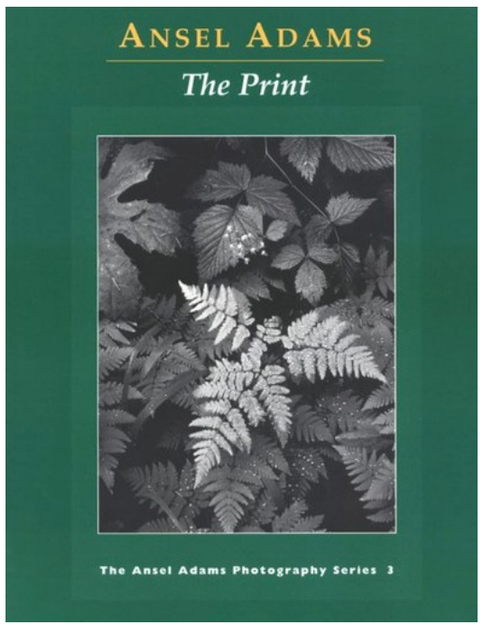

HIGH DYNAMIC RANGE IMAGING
ACQUISITION, DISPLAY AND IMAGE-BASED LIGHTING
ERIK REINHARD
GREG WARD
SUMANTA PATTANAIK
AND PAUL DEBEVEC


MASTERING HDR PHOTOGRAPHY
MICHAEL FREEMAN


Bernd Hoefflinger
Editor
High-Dynamic-Range (HDR) Vision
Springer


Copyrighted Material
The Photographer's
MASTER PRINTING COURSE
In association with
Amateur Photographer
Tim Rudman
Copyrighted Material


ANSEL ADAMS
The Negative
The Ansel Adams Photography Series 2


ANSEL ADAMS
The Print
The Ansel Adams Photography Series 3

Other tone mapping references

- J. DiCarlo and B. Wandell, Rendering High Dynamic Range Images http://www-isl.stanford.edu/%7Eabbas/group/papers_and_pub/spie00_jeff.pdf
- Choudhury, P., Tumblin, J., "The Trilateral Filter for High Contrast Images and Meshes". http://www.cs.northwestern.edu/~jet/publications.html
- Tumblin, J., Turk, G., "Low Curvature Image Simplifiers (LCIS): A Boundary Hierarchy for Detail-Preserving Contrast Reduction." http://www.cs.northwestern.edu/~jet/publications.html
- Tumblin, J., "Three Methods For Detail-Preserving Contrast Reduction For Displayed Images" http://www.cs.northwestern.edu/~jet/publications.html
- Photographic Tone Reproduction for Digital Images Erik Reinhard, Mike Stark, Peter Shirley and Jim Ferwerda http://www.cs.utah.edu/%7Ereinhard/cdrom/
- Ashikhmin, M. ``A Tone Mapping Algorithm for High Contrast Images'' http://www.cs.sunysb.edu/~ash/tm.pdf
- Retinex at Nasa http://dragon.larc.nasa.gov/retinex/background/retpubs.html
- Gradient Domain High Dynamic Range Compression Raanan Fattal, Dani Lischinski, Michael Werman http://www.cs.huji.ac.il/~danix/hdr/
- Li et al. : Wavelets and activity maps http://web.mit.edu/yzli/www/hdr_companding.htm

Tone mapping code

- **http://www.mpi-sb.mpg.de/resources/pfstools/**
- **http://scanline.ca/exrtools/**
- **http://www.cs.utah.edu/~reinhard/cdrom/source.html**
- **http://www.cis.rit.edu/mcsl/icam/hdr/**

Refs

http://people.csail.mit.edu/sparis/bf_course/

http://people.csail.mit.edu/fredo/PUBLI/Siggraph2002/

http://www.hdrsoft.com/resources/dri.html

http://www.clarkvision.com/imagedetail/dynamicrange2/

http://www.debevec.org/HDRI2004/

http://www.luminous-landscape.com/tutorials/hdr.shtml

http://www.anyhere.com/gward/hdrenc/

http://www.debevec.org/IBL2001/NOTES/42-gward-cic98.pdf

http://www.openexr.com/

http://gl.ict.usc.edu/HDRShop/

http://www.dpreview.com/learn/?/Glossary/Digital_Imaging/Dynamic_Range_01.htm

http://www.normankoren.com/digital_tonality.html

http://www.anyhere.com/

http://www.cybergrain.com/tech/hdr/

Available in HDRShop



Slide from Siggraph 2005 course on HDR

Questions?

Speaker notes