Applications

SIGGRAPH2007

Frédo Durand

MIT CSAIL

Overview

- Denoising
- Tone mapping
- Relighting & texture editing

Overview

- **Denoising**
  Not most powerful application
  Not best denoising, but good & simple

- Tone mapping

- Relighting & texture editing

Basic denoising

Noisy input

Bilateral filter 7x7 window

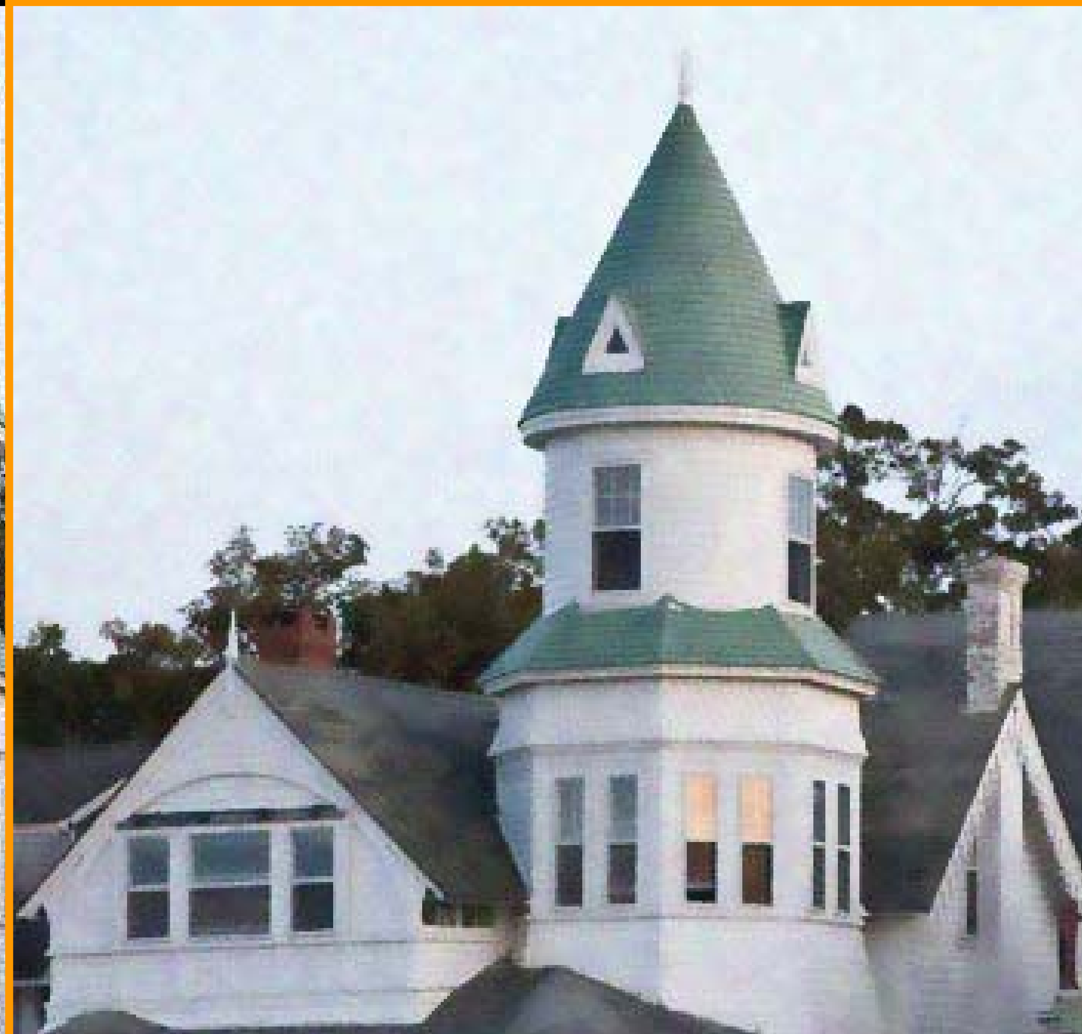

Basic denoising

Bilateral filter

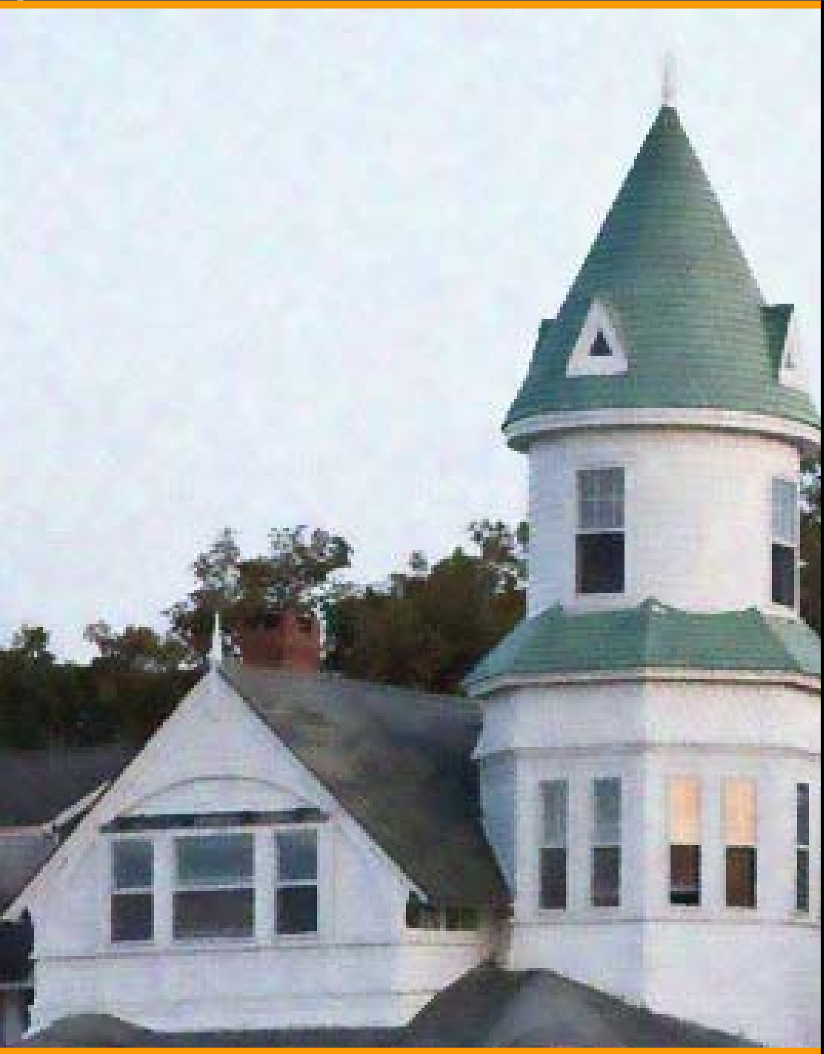

Median 3x3

Basic denoising

Bilateral filter

Median 5x5

Basic denoising

Bilateral filter

Bilateral filter – lower sigma

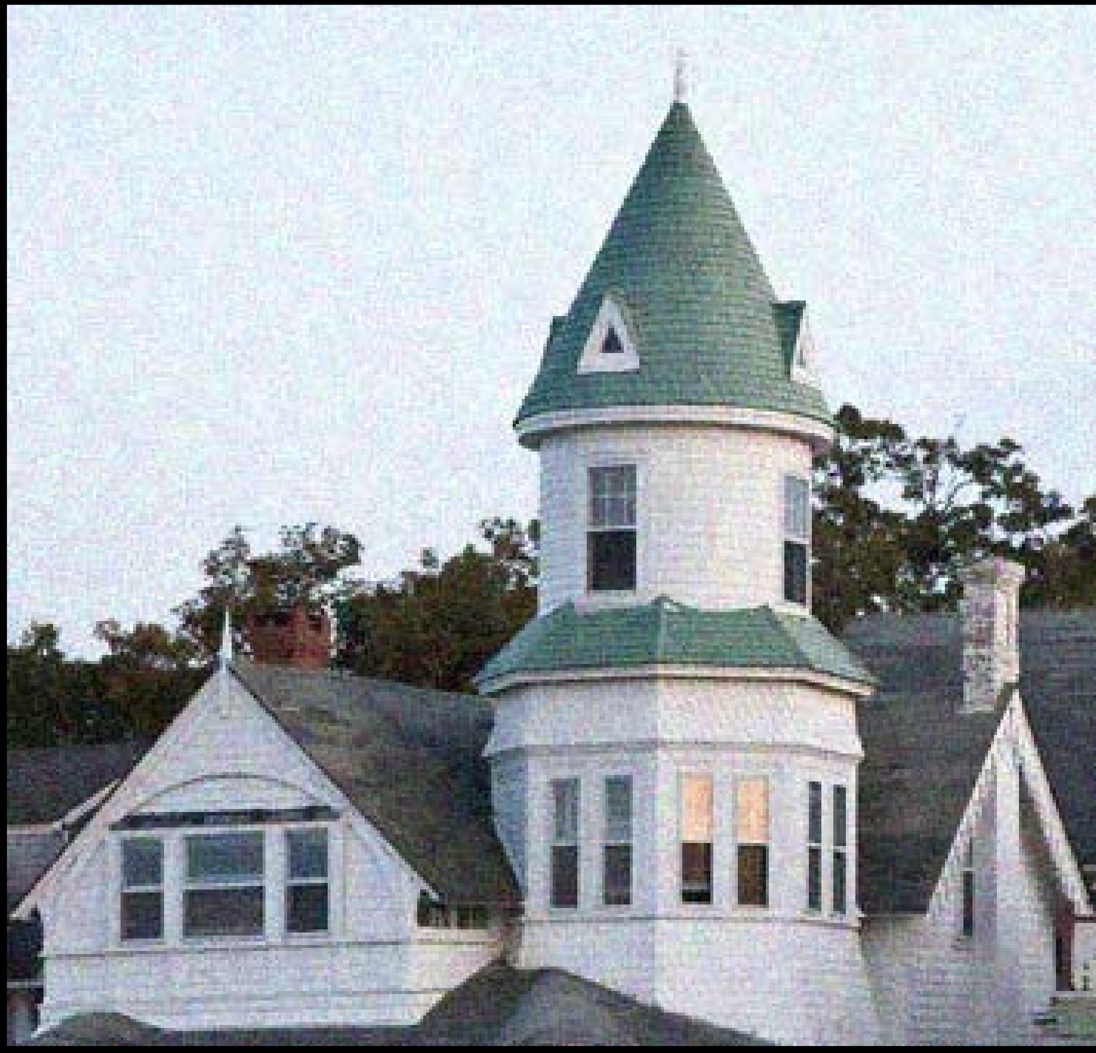

Basic denoising

Bilateral filter

Bilateral filter – higher sigma

Denoising

- Small spatial sigma (e.g. 7x7 window)
- Adapt range sigma to noise level
- Maybe not best denoising method, but best simplicity/quality tradeoff
  - No need for acceleration (small kernel)
  - But the denoising feature in e.g. Photoshop is better

Overview

- Denoising

- **Tone mapping**

- Relighting & texture editing

Real world dynamic range

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

Picture dynamic range

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

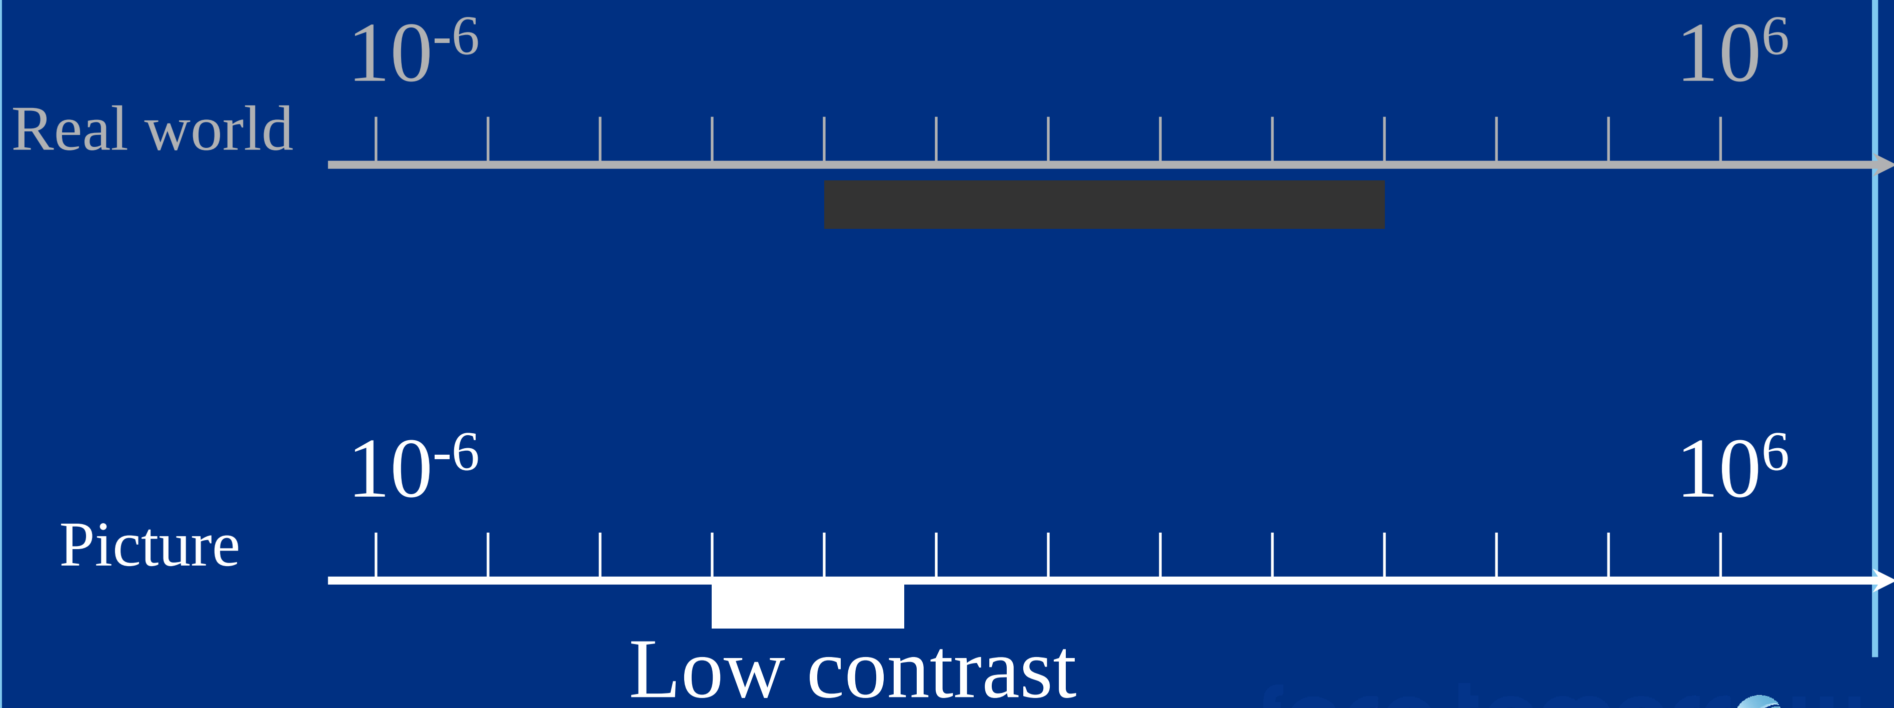

Multiple exposure photography

- Merge multiple exposure to cover full range

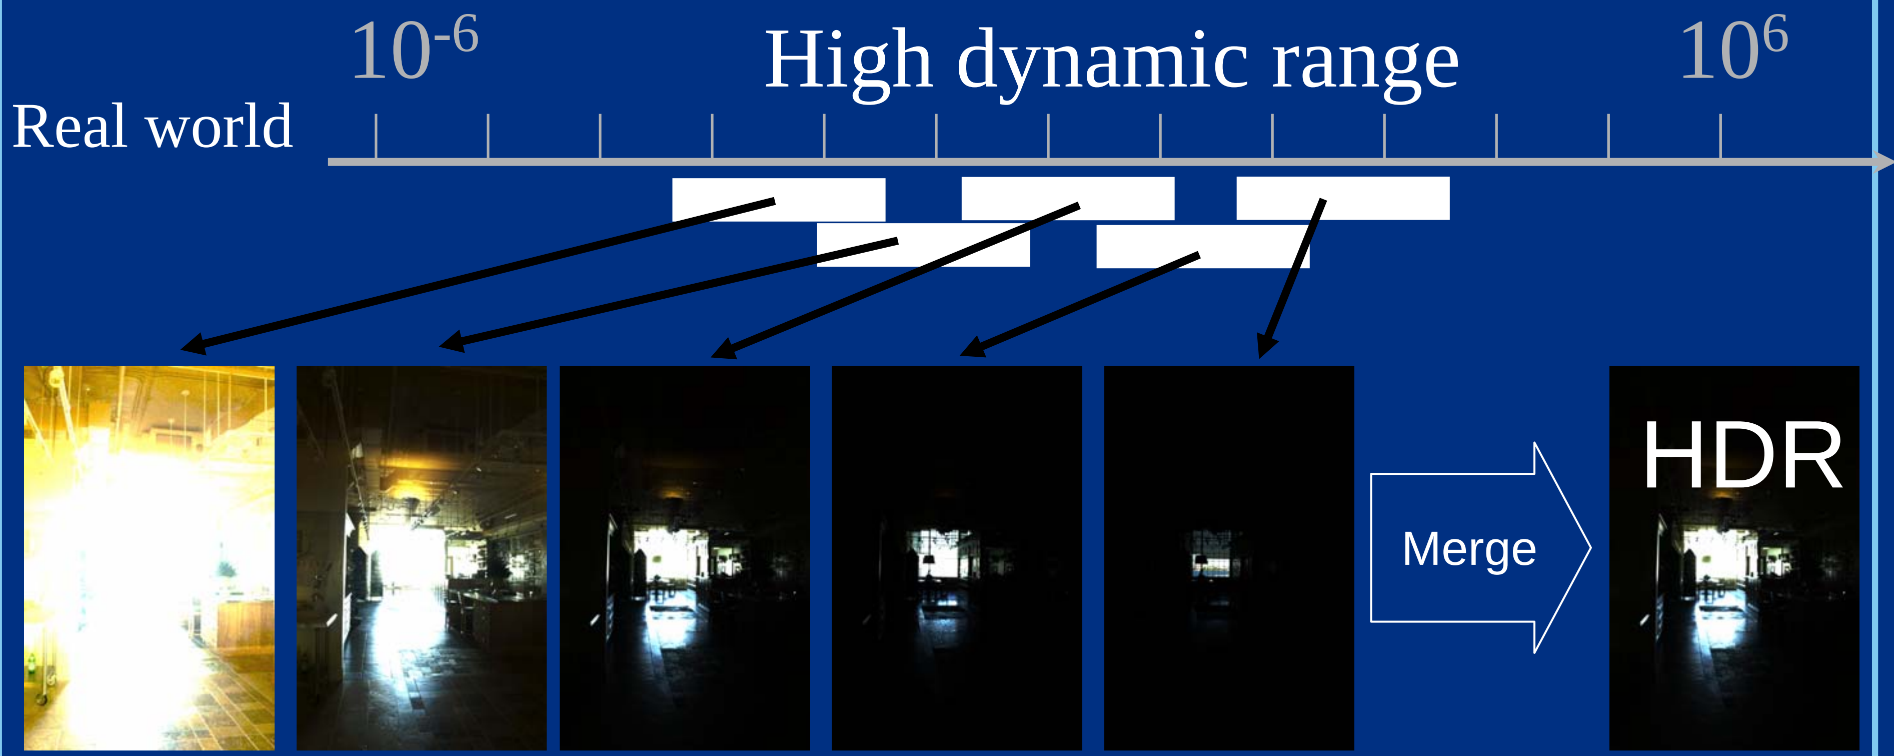

- We obtain one single image with floats per pixel
  - But we still can't display it

The future: HDR Cameras

- HDR sensors using CMOS
  - Use a log response curve
  - e.g. SMaL,

- Assorted pixels
  - Fuji
  - Nayar et al.

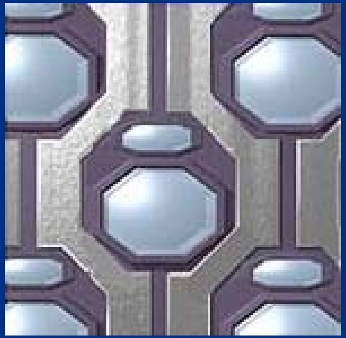

Fuji SuperCCD

- Per-pixel exposure
  - Filter
  - Integration time

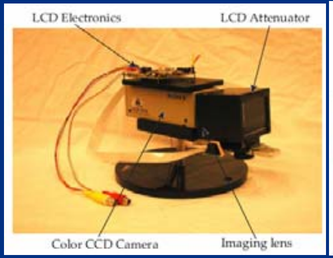

- Multiple cameras using beam splitters

Problem: Contrast reduction

- Match limited contrast of the medium
- Preserve details

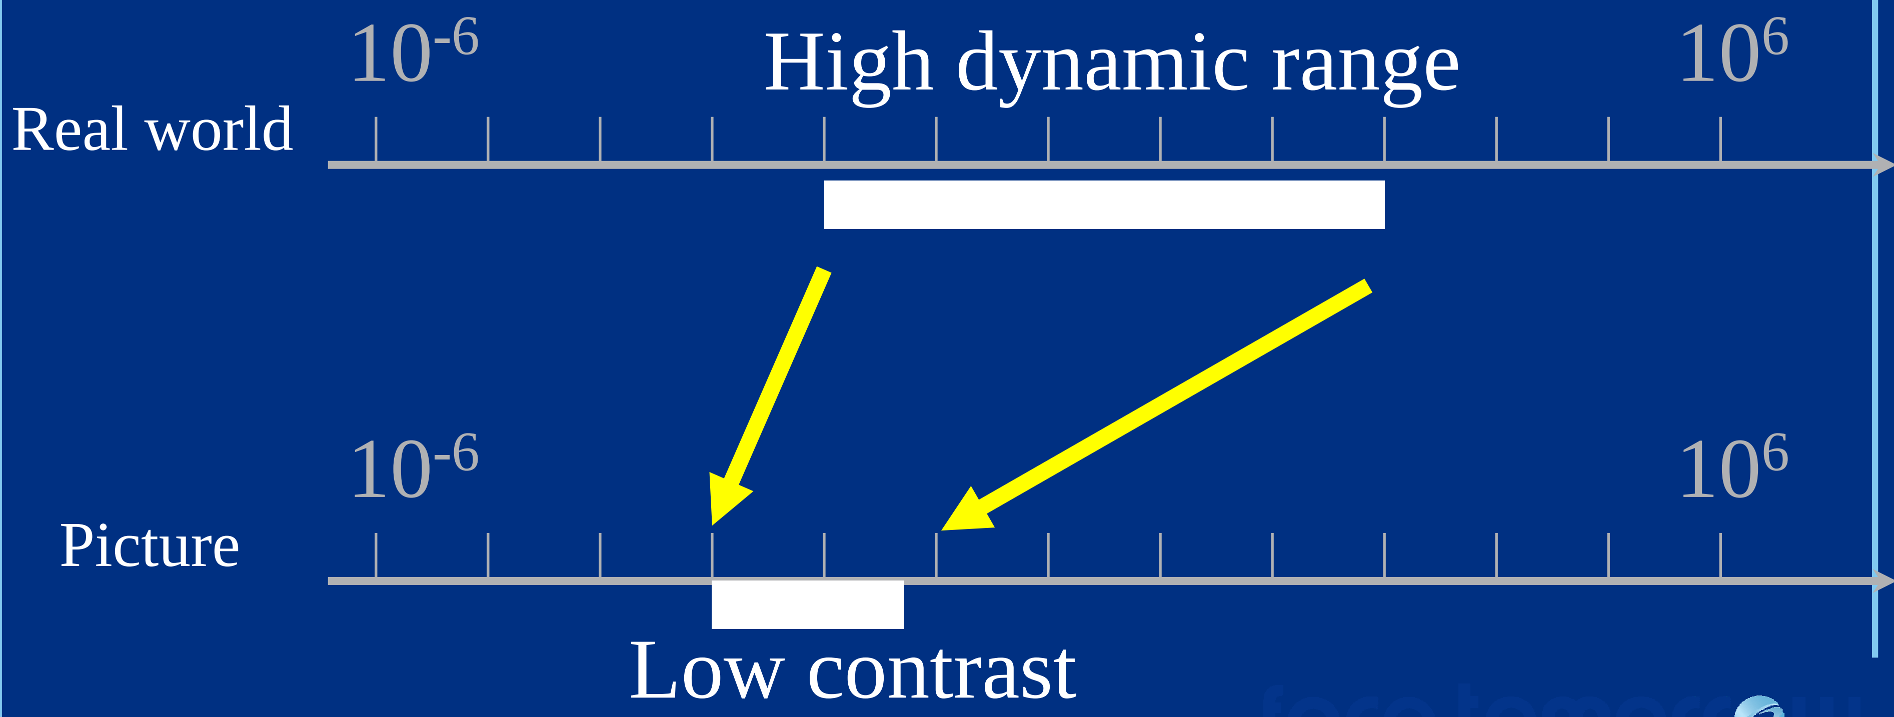

Tone mapping

- Input: high-dynamic-range image
  - (floating point per pixel)

Naïve technique

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

Naïve: Gamma compression

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

Gamma compression on intensity

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

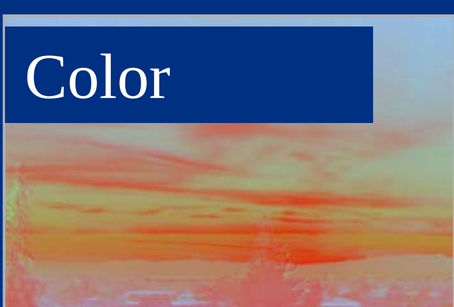

Oppenheim 1968, Chiu et al. 1993

- Reduce contrast of low-frequencies (log domain)
- Keep high frequencies

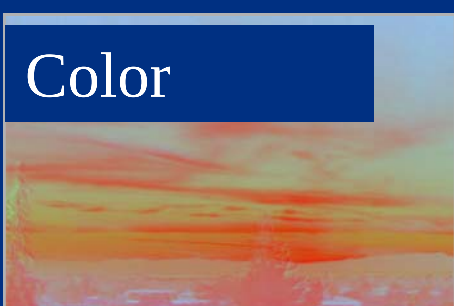

The halo nightmare

- For strong edges
- Because they contain high frequency

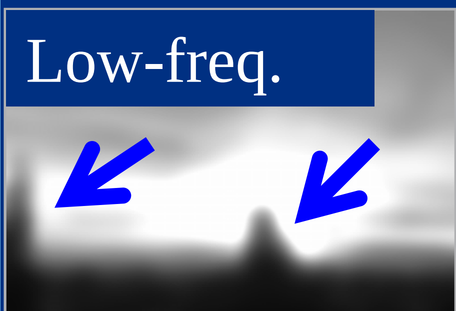

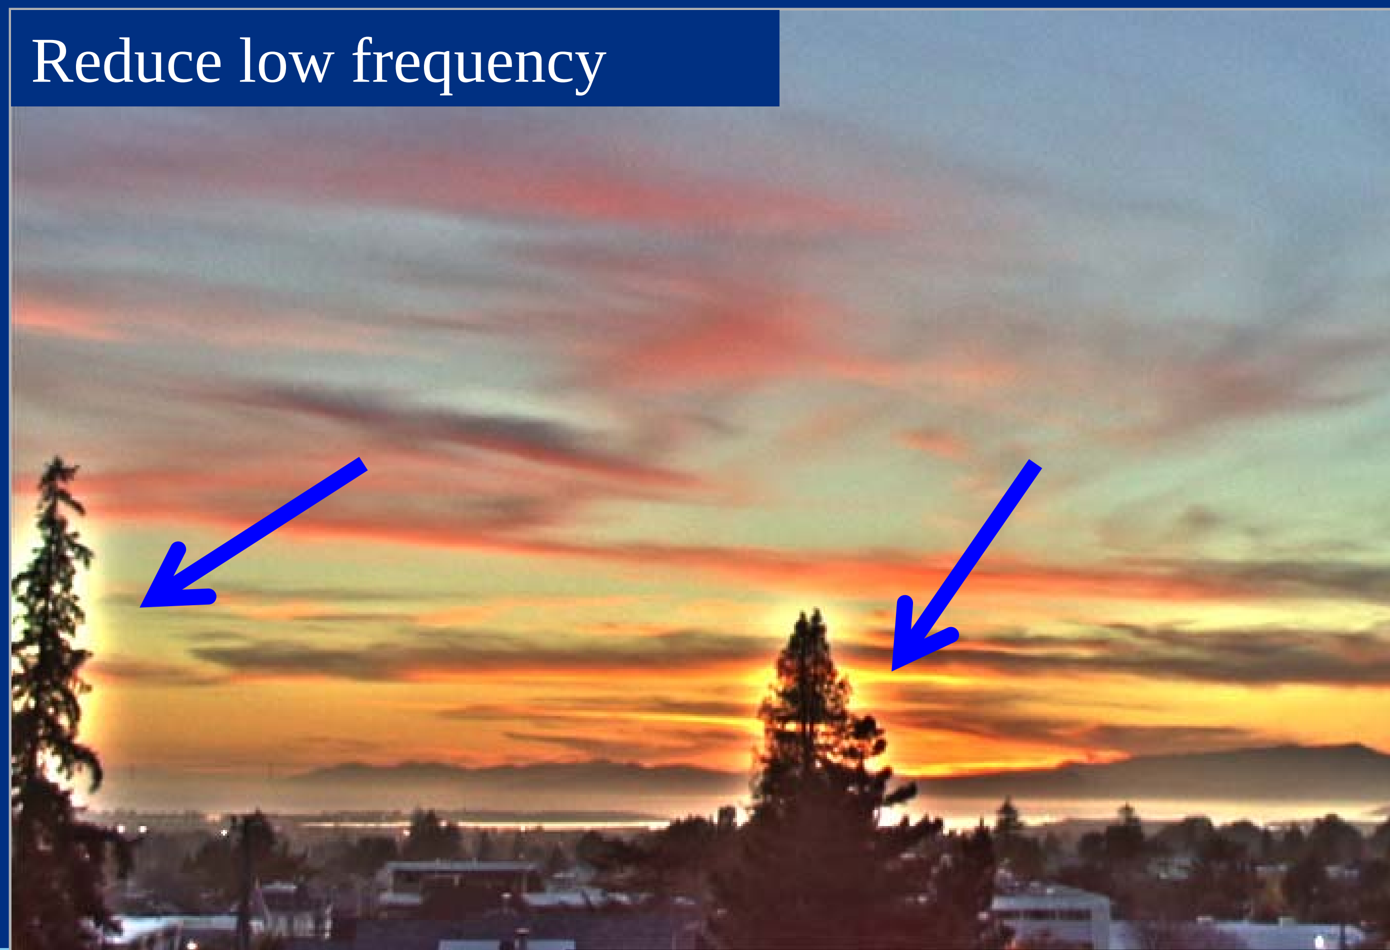

Bilateral filtering to the rescue

- Large scale = bilateral (log intensity)
- Detail = residual

[Durand & Dorsey 2002]

Contrast reduction

Input HDR image

Contrast
too high!

Contrast reduction

Contrast reduction

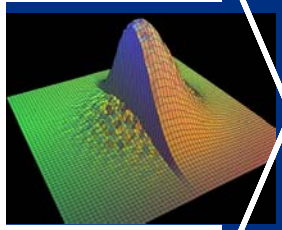

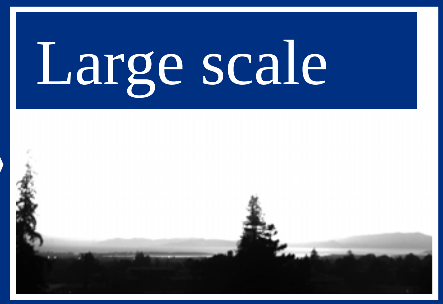

Bilateral
Filter
(in log domain!)

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

Contrast reduction

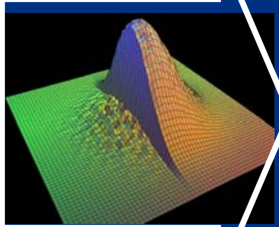

Bilateral Filter

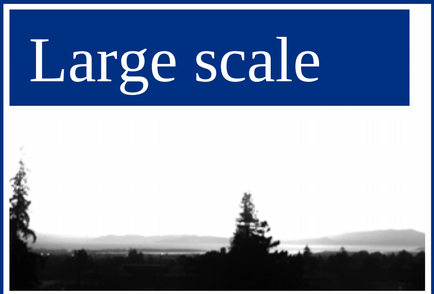

Detail = log intensity –large scale (residual)

Contrast reduction

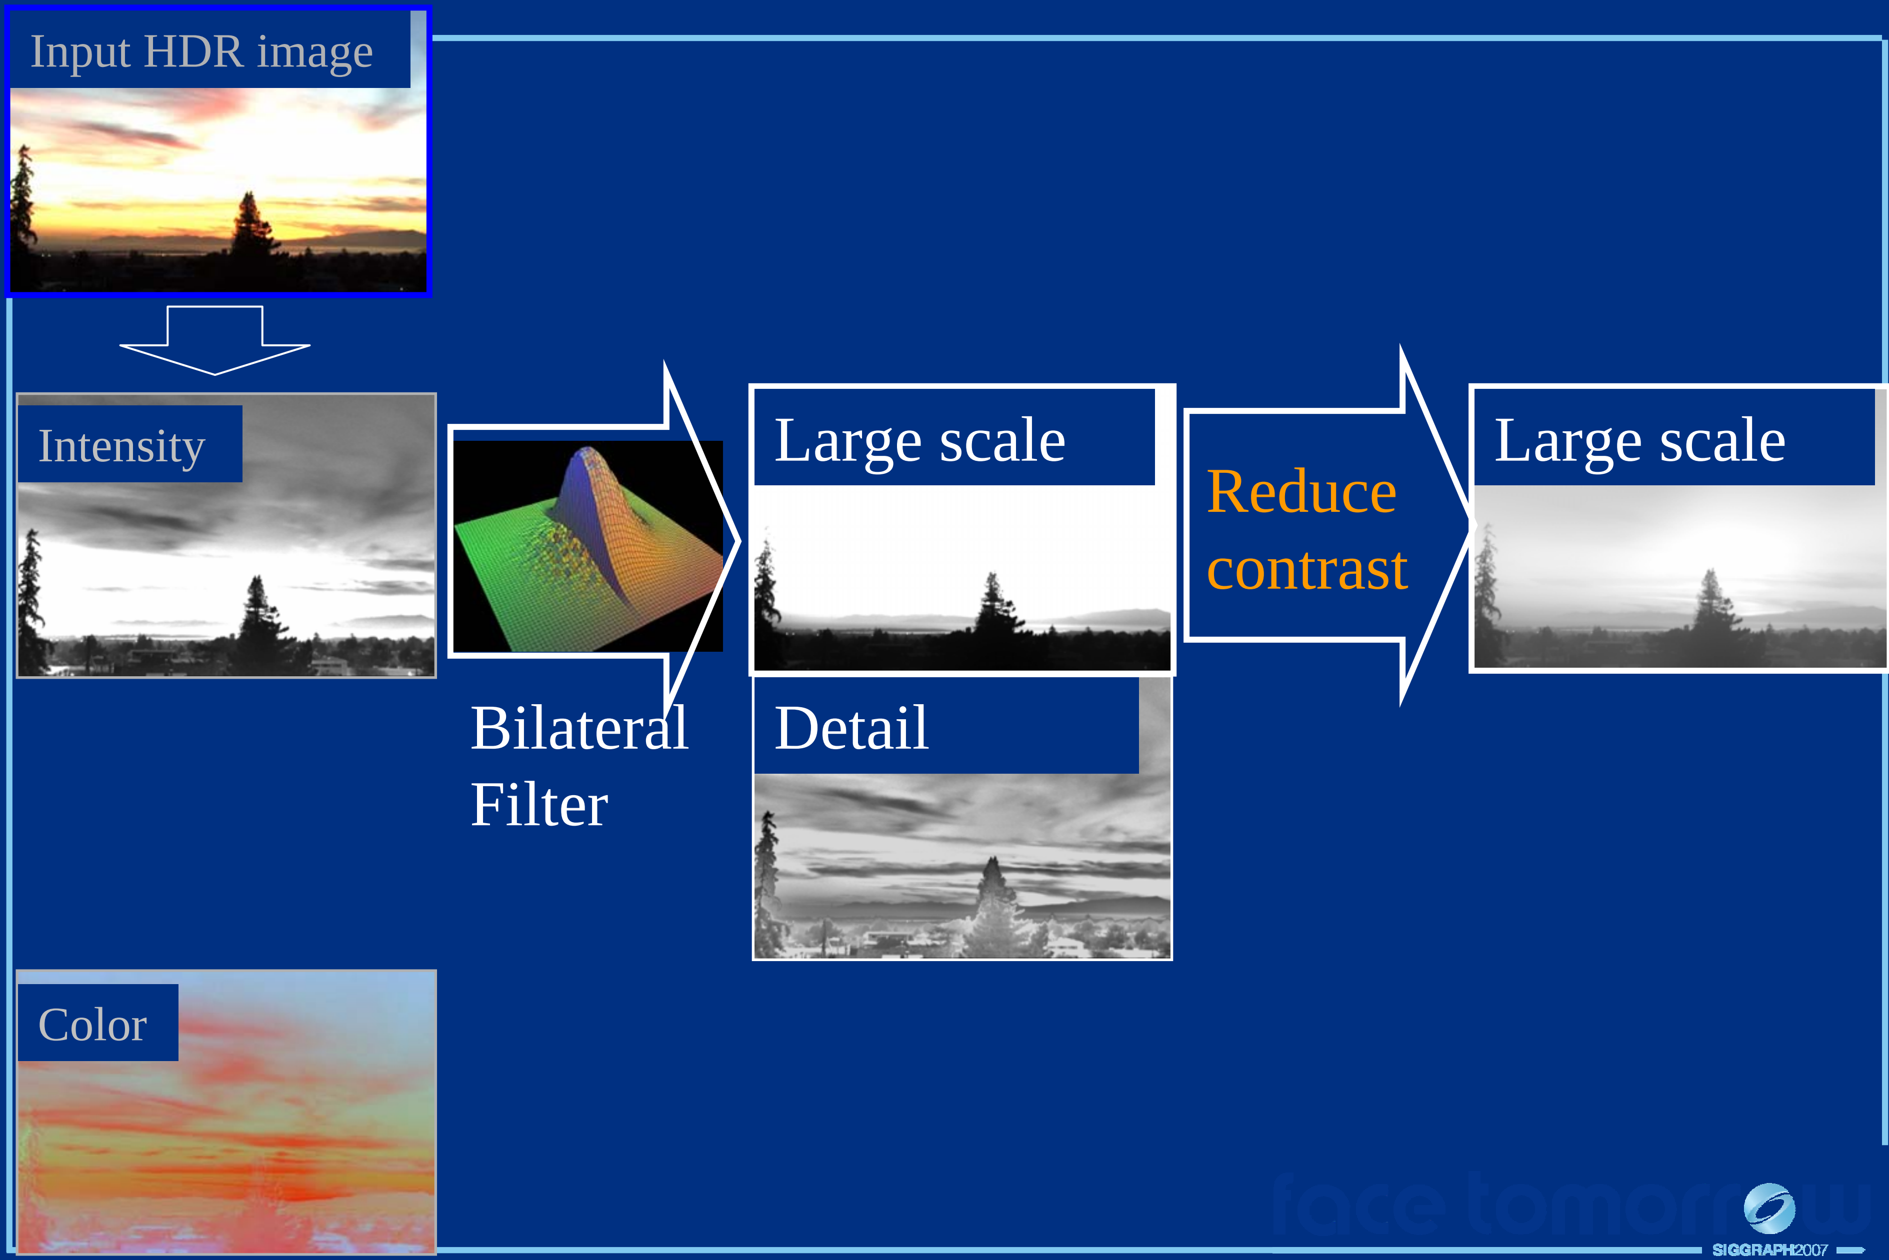

Contrast reduction

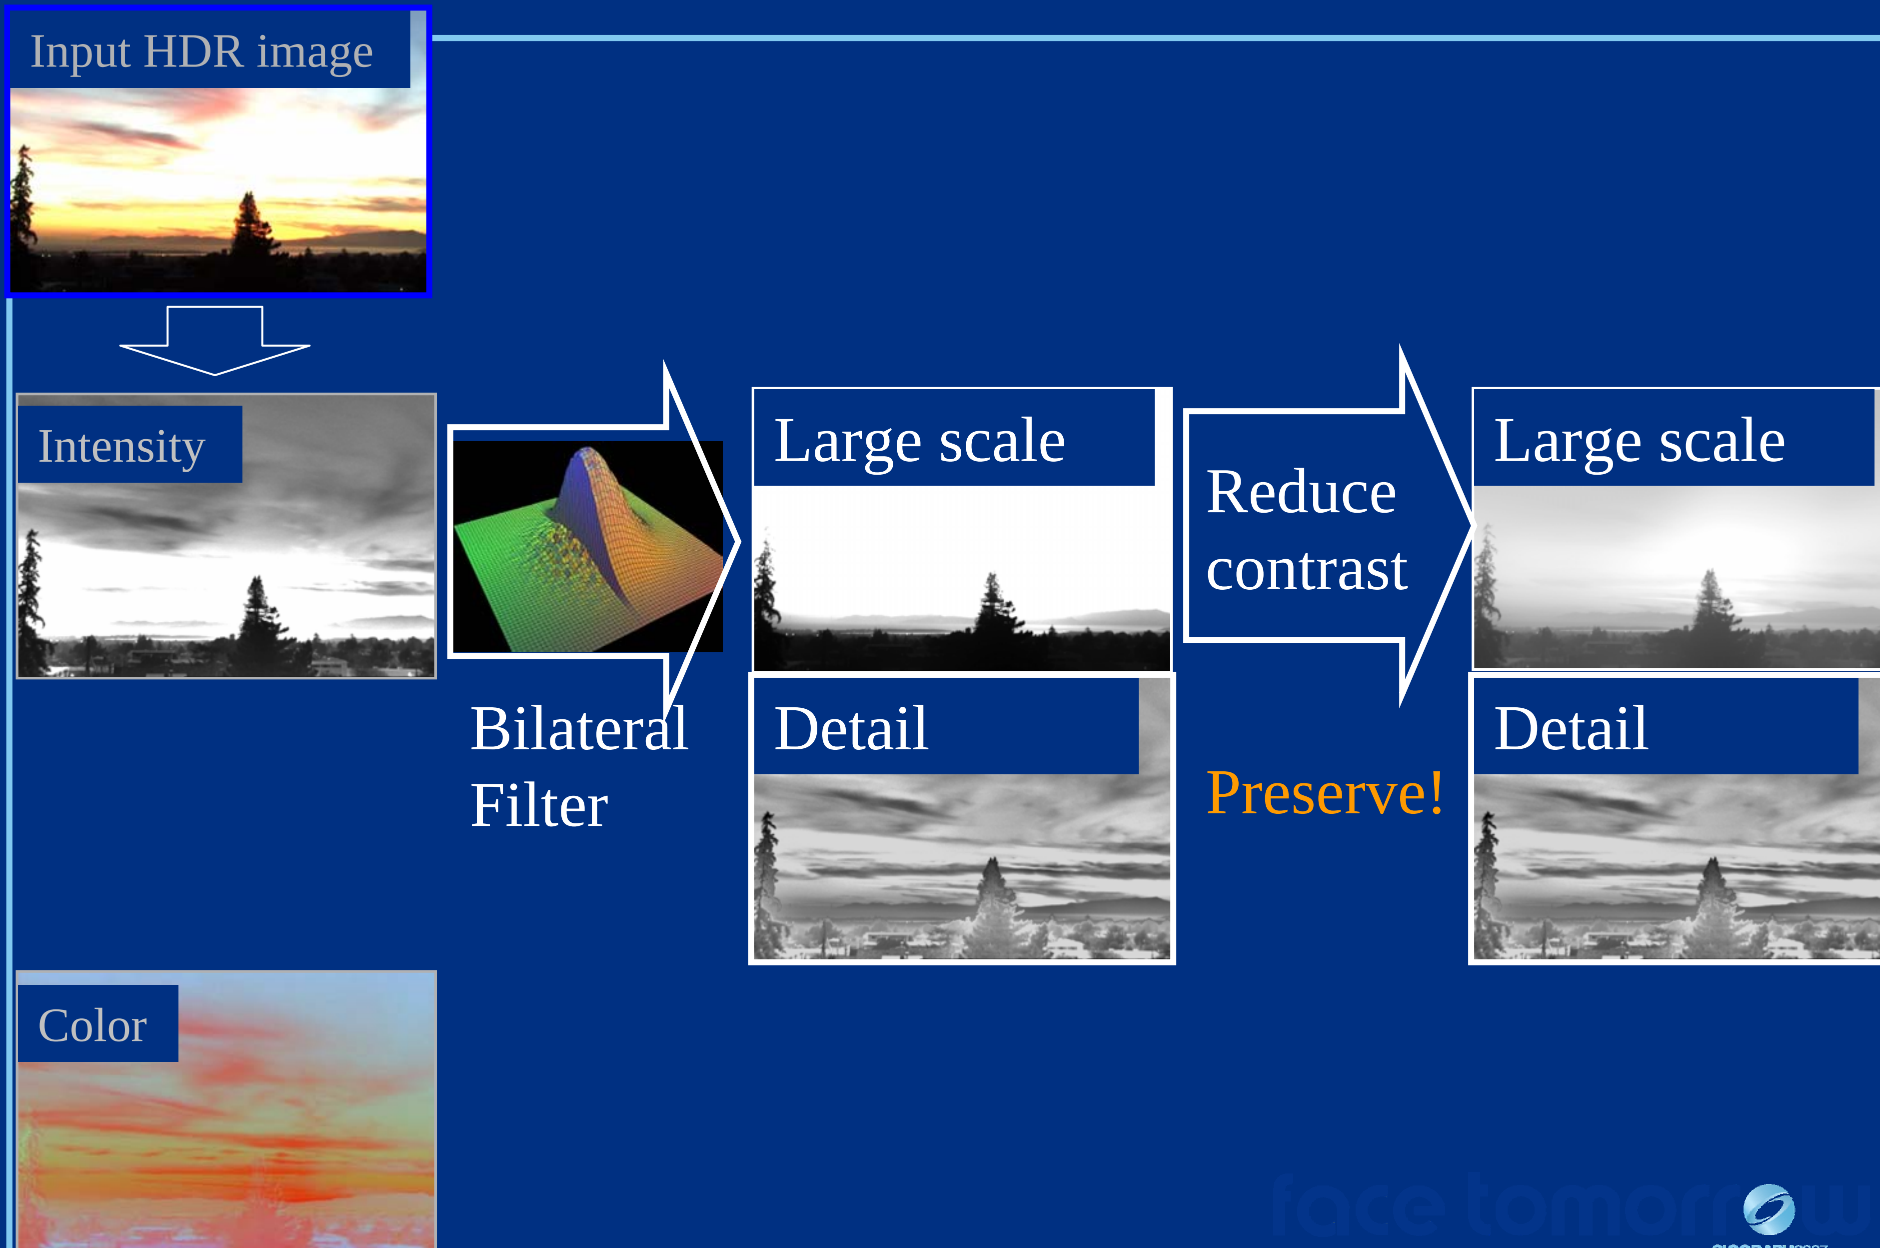

Contrast reduction

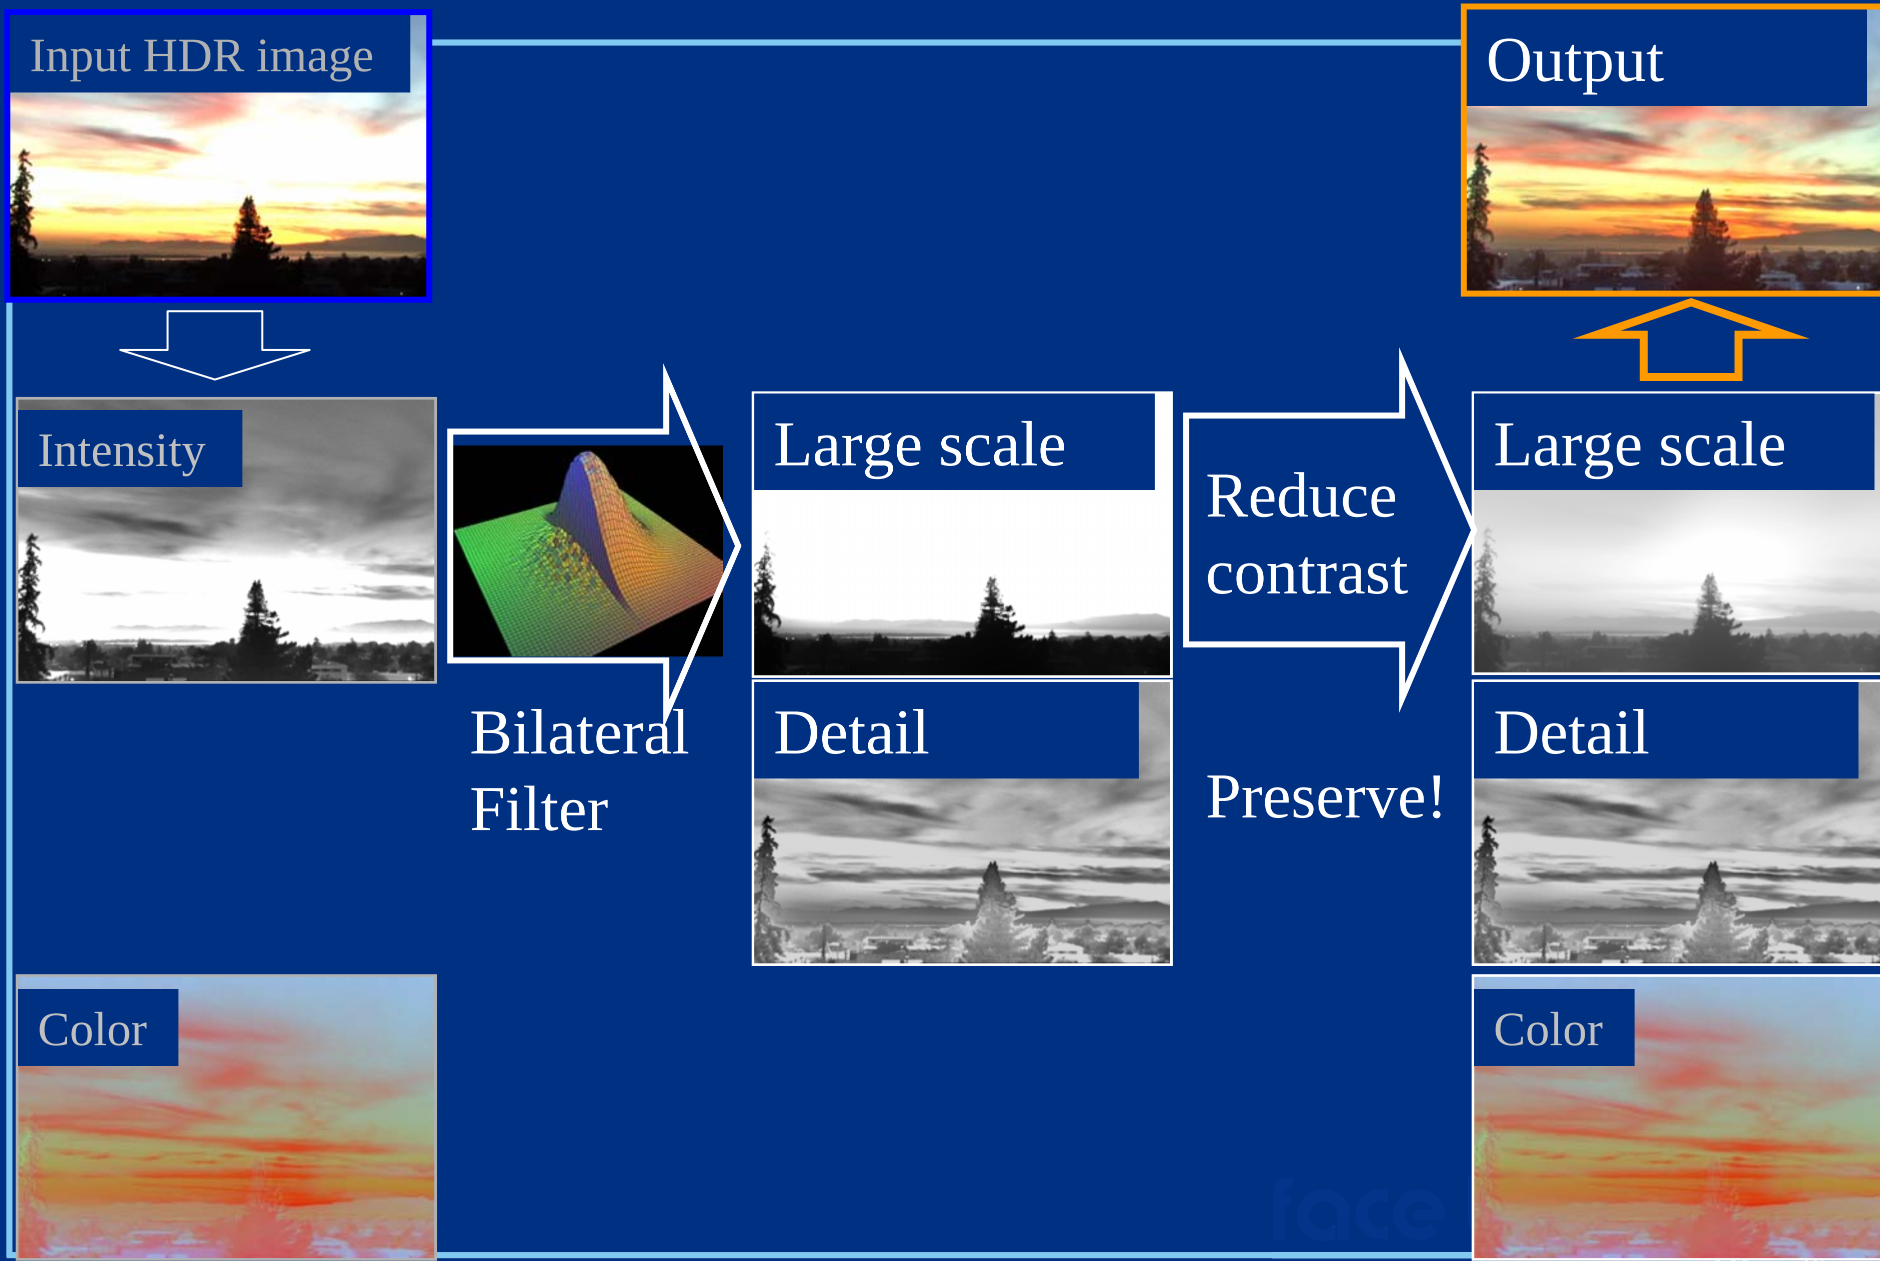

Contrast reduction in log domain

- Set target large-scale contrast (e.g. $\log_{10} 10$)
  - In **linear** output, we want 1:10 contrast for large scale
- Compute range of input large scale layer:
  - largeRange = max(inLogLarge) – min (inLogLarge)
- Scale factor $k = \log_{10}(10)$ / largeRange
- Normalize so that the biggest value is 0 in log

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

Alternative explanation

- Explanation 1 (previous slides):
  - outLog = *k* inLogLarge + inLogDetail (ignoring offset)
- Explanation 2
  - outLog = *k* inLogIntensity + (1-*k*) detail
  - Reduce contrast of full intensity layer
  - Add back some detail
- Same final effect since
  - inLogDetail+inLogLarge scale = inLogIntensity
  - But different philosophy:
    decomposition vs. add back detail

Demo

- Don’t try at home without FAST bilateral filtering

Denoising vs. tone mapping

- Denoising:
  - decompose into noise+signal
  - Throw away noise, keep signal
  - Small kernel
- Tone mapping
  - Decompose into large scale + detail
  - Preserve detail, reduce large scale
  - Large kernel
    - because detail=high+medium frequency
    - ➔computation challenge

Crossing lines

- The bilateral filter is influenced by pixels across thin line
- Good for tone mapping

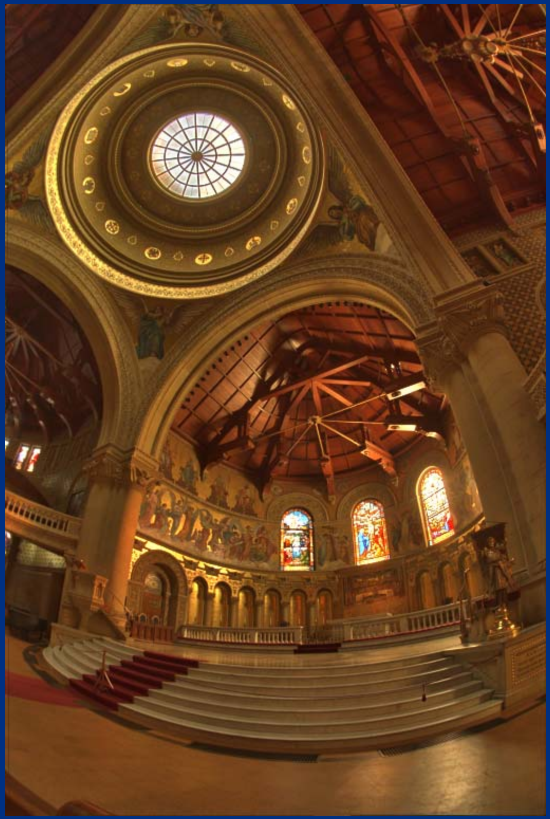

What matters

- 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
  - CIELab might be better for other applications
- Luminance computation
  - Not critical, but has influence
  - see our Flash/no-flash paper [Eisemann 2004] for smarter function

Tone mapping evaluation

- Recent user experiments to evaluate competing tone mapping
  - Ledda et al. 2005
    http://www.cs.bris.ac.uk/Publications/Papers/2000255.pdf
  - Kuang et al. 2004
    http://www.cis.rit.edu/fairchild/PDFs/PRO22.pdf
- Interestingly, the former concludes bilateral is the worst, the latter that it is the best!
  - They choose to test a different criterion: fidelity vs. preference
- More importantly, they focus on algorithm and ignore parameters

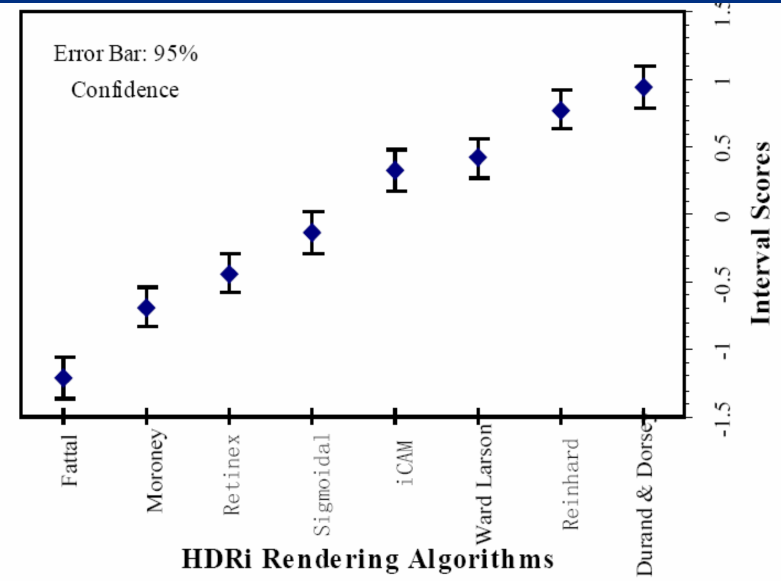

From Kuang et al.

	1st	2nd	3rd	4th	5th	6th
Scene 1	*P*	*B*	*A*	*H*	*I*	*L*
Scene 2	*I*	*P*	*H*	*A*	*B*	*L*
Scene 3	*P*	*I*	*A*	*H*	*L*	*B*
Scene 4	*P*	*L*	*I*	*A*	*H*	*B*
Scene 5	*I*	*H*	*A*	*P*	*L*	*B*
Scene 6	*I*	*H*	*A*	*P*	*L*	*B*
Scene 7	*I*	*A*	*P*	*H*	*B*	*L*
Scene 8	*I*	*P*	*A*	*H*	*L*	*B*
Scene 9	*P*	*A*	*L*	*H*	*B*	*I*

Adapted from Ledda et al.

Alternative explanation

- Contrast reduction w/ intrinsic layers [Tumblin et al. 1999]
- For 3D scenes: Reduce only illumination layer

Illumination layer
Compressed

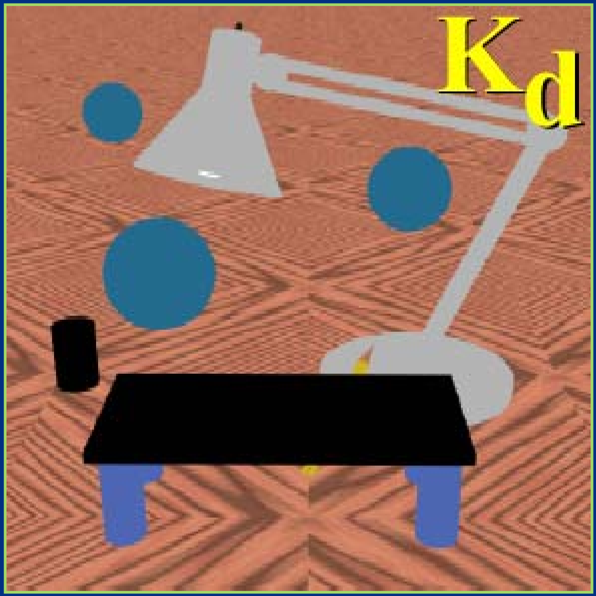

Reflectance layer

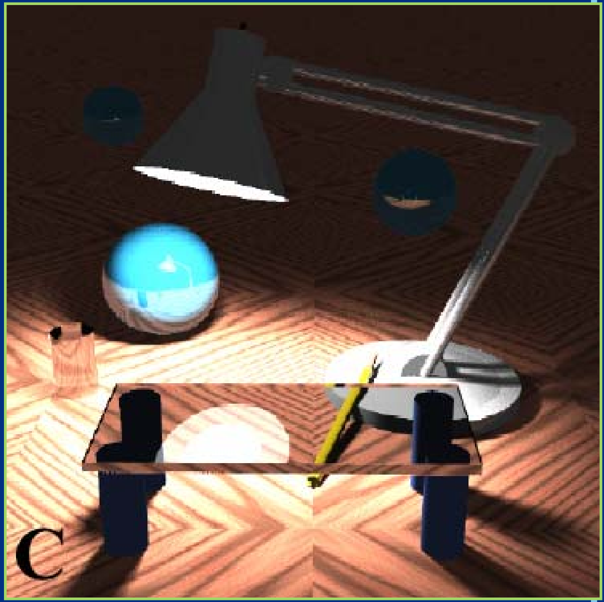

Output

Dirty vision for cool graphics

Three wrongs make one right

- Analyze image
  - Intrinsic image: albedo & illumination
  - Simple bilateral filter

- Modify
  - In our case, reduce contrast of large-scale (illumination)

- Recombine
  - Get final image

Overview

- Denoising

- Tone mapping

- **Relighting & texture editing**

Discounting Existing Lighting

- Motivation
  - Relighting
  - Image manipulation (e.g. clone brush, texture synthesis)
- Context:
  - The following slides are from a project dealing with images +depth

Inverse Lighting Simulation

- Physically-based approaches

  e.g. [Fournier et al.93, Drettakis et al.97, Debevec.98, Yu et al.99, Loscos et al. 99, Loscos et al.00]

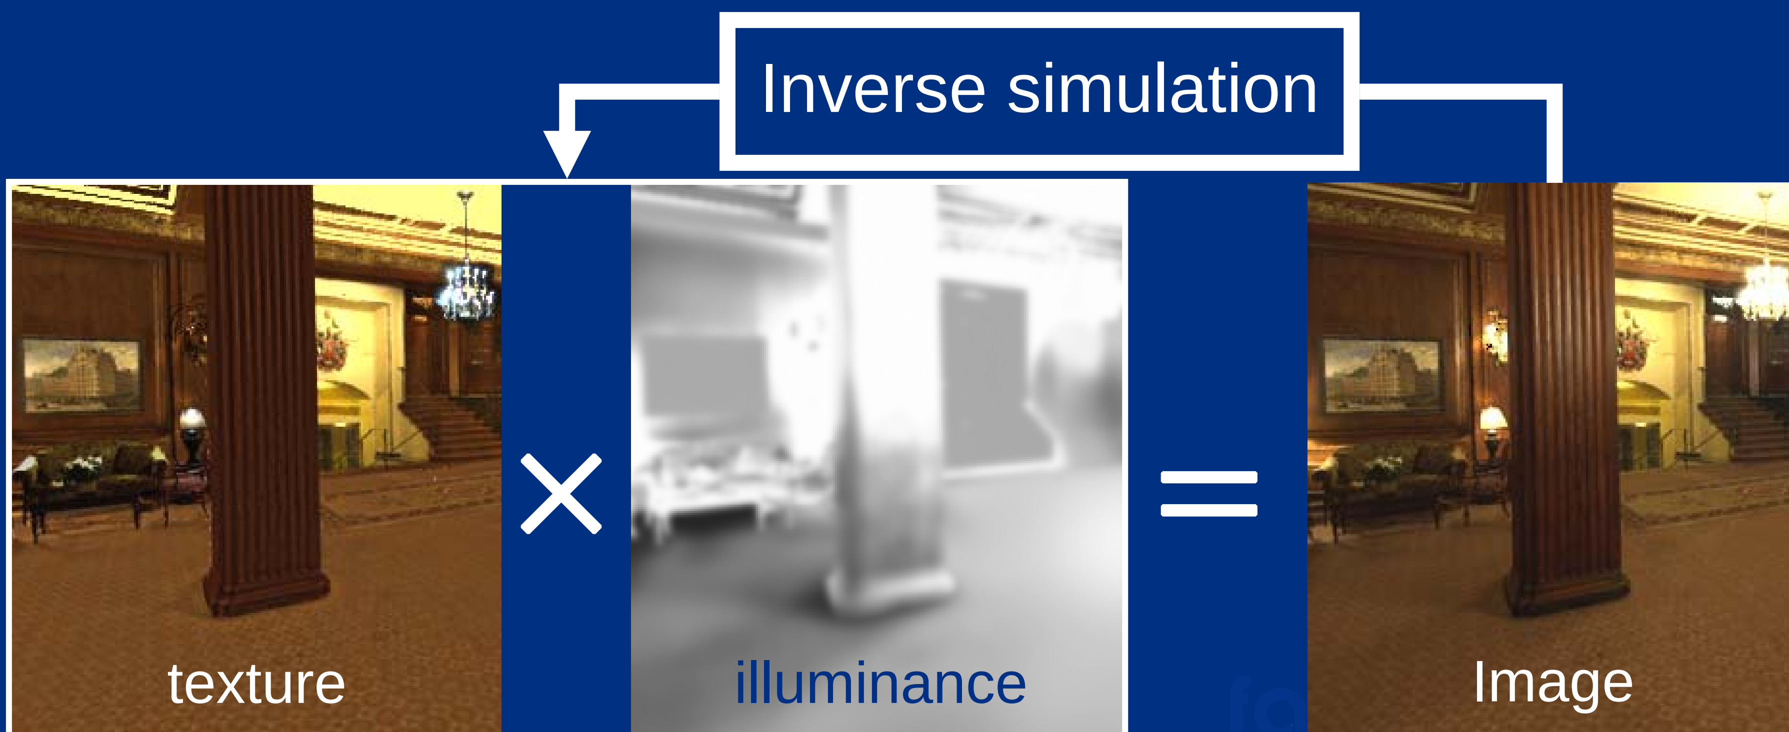

Texture-Illuminance Decoupling

- Not physically based
  - Our “texture” and “illuminance” are reasonable estimates

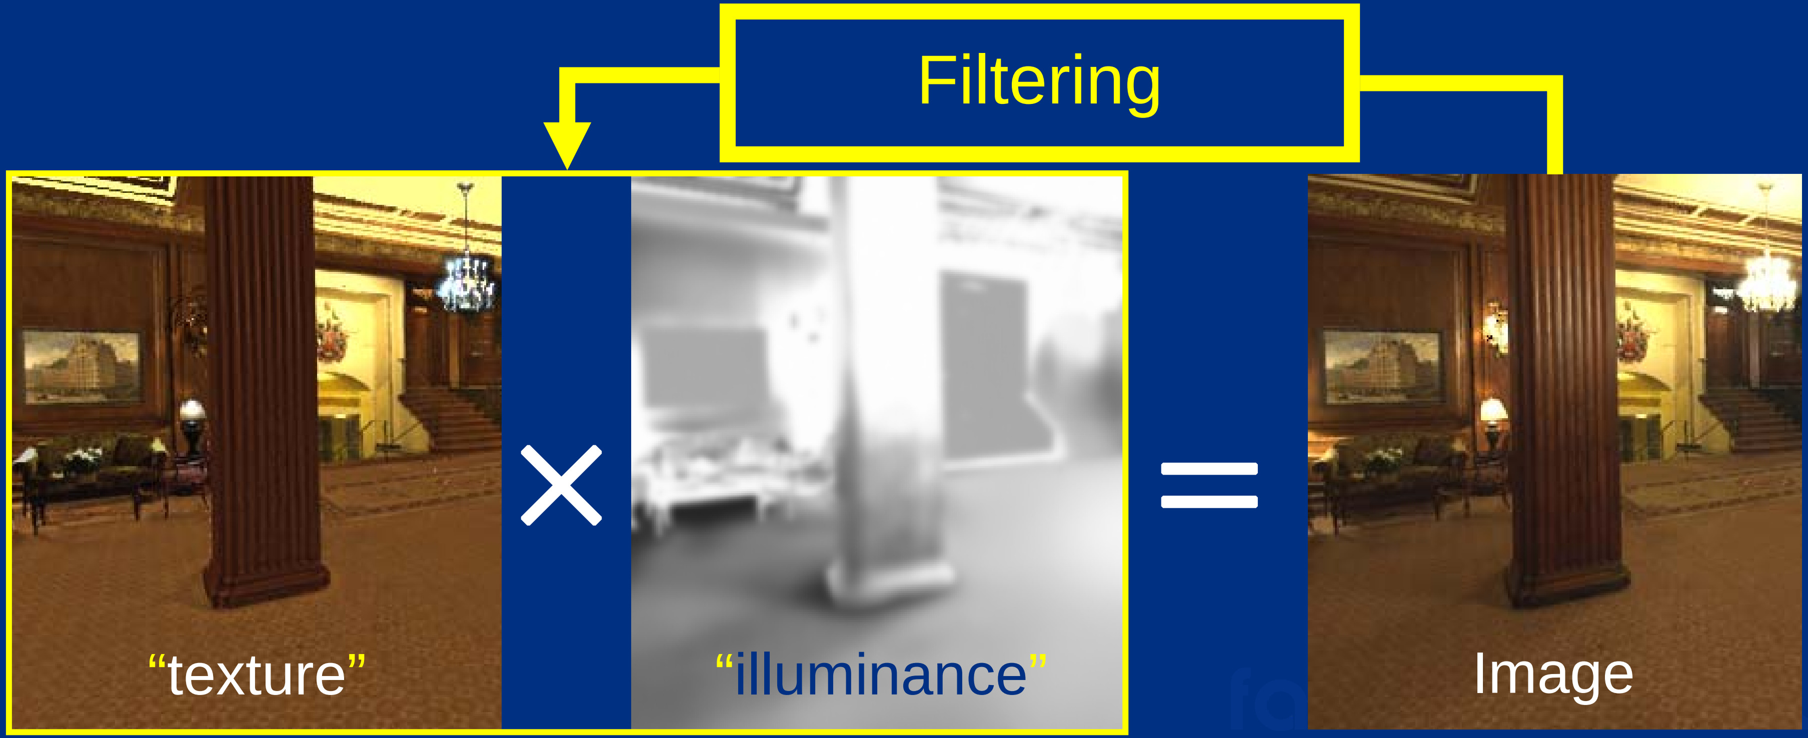

Texture-Illuminance Decoupling

- Not physically based: Filtering
- Assumptions:
  - Small-scale features → “texture”
  - Large-scale features → “illuminance”

General Idea: A Naïve Approach

- Large-scale features using low-pass filter
  - Color is assumed to be from texture

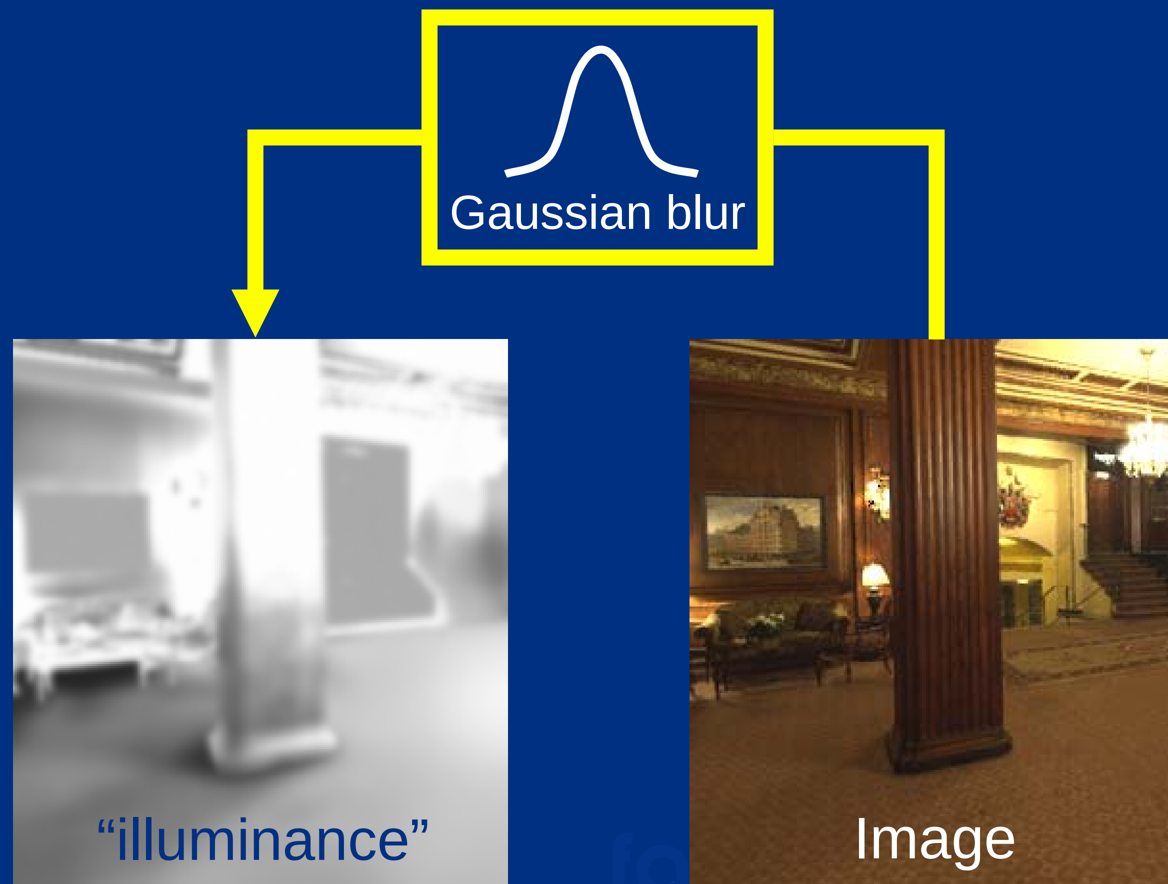

General Idea: A Naïve Approach

- Extract texture from illuminance and input image

Problems with the Naïve Approach

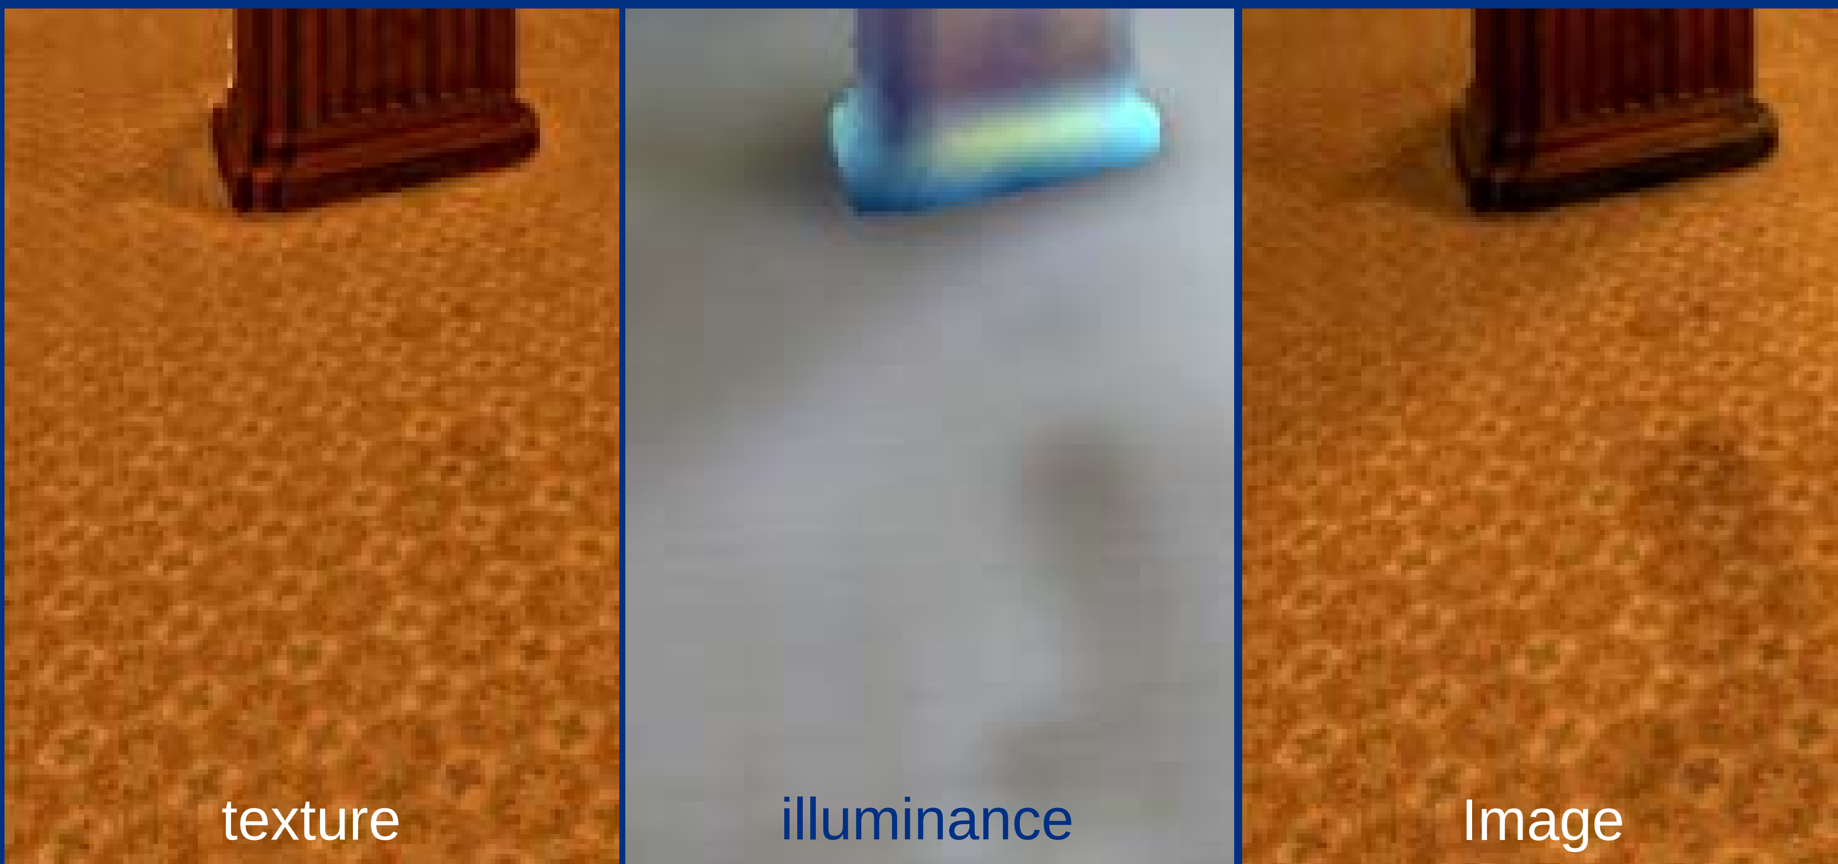

Problems with the Naïve Approach

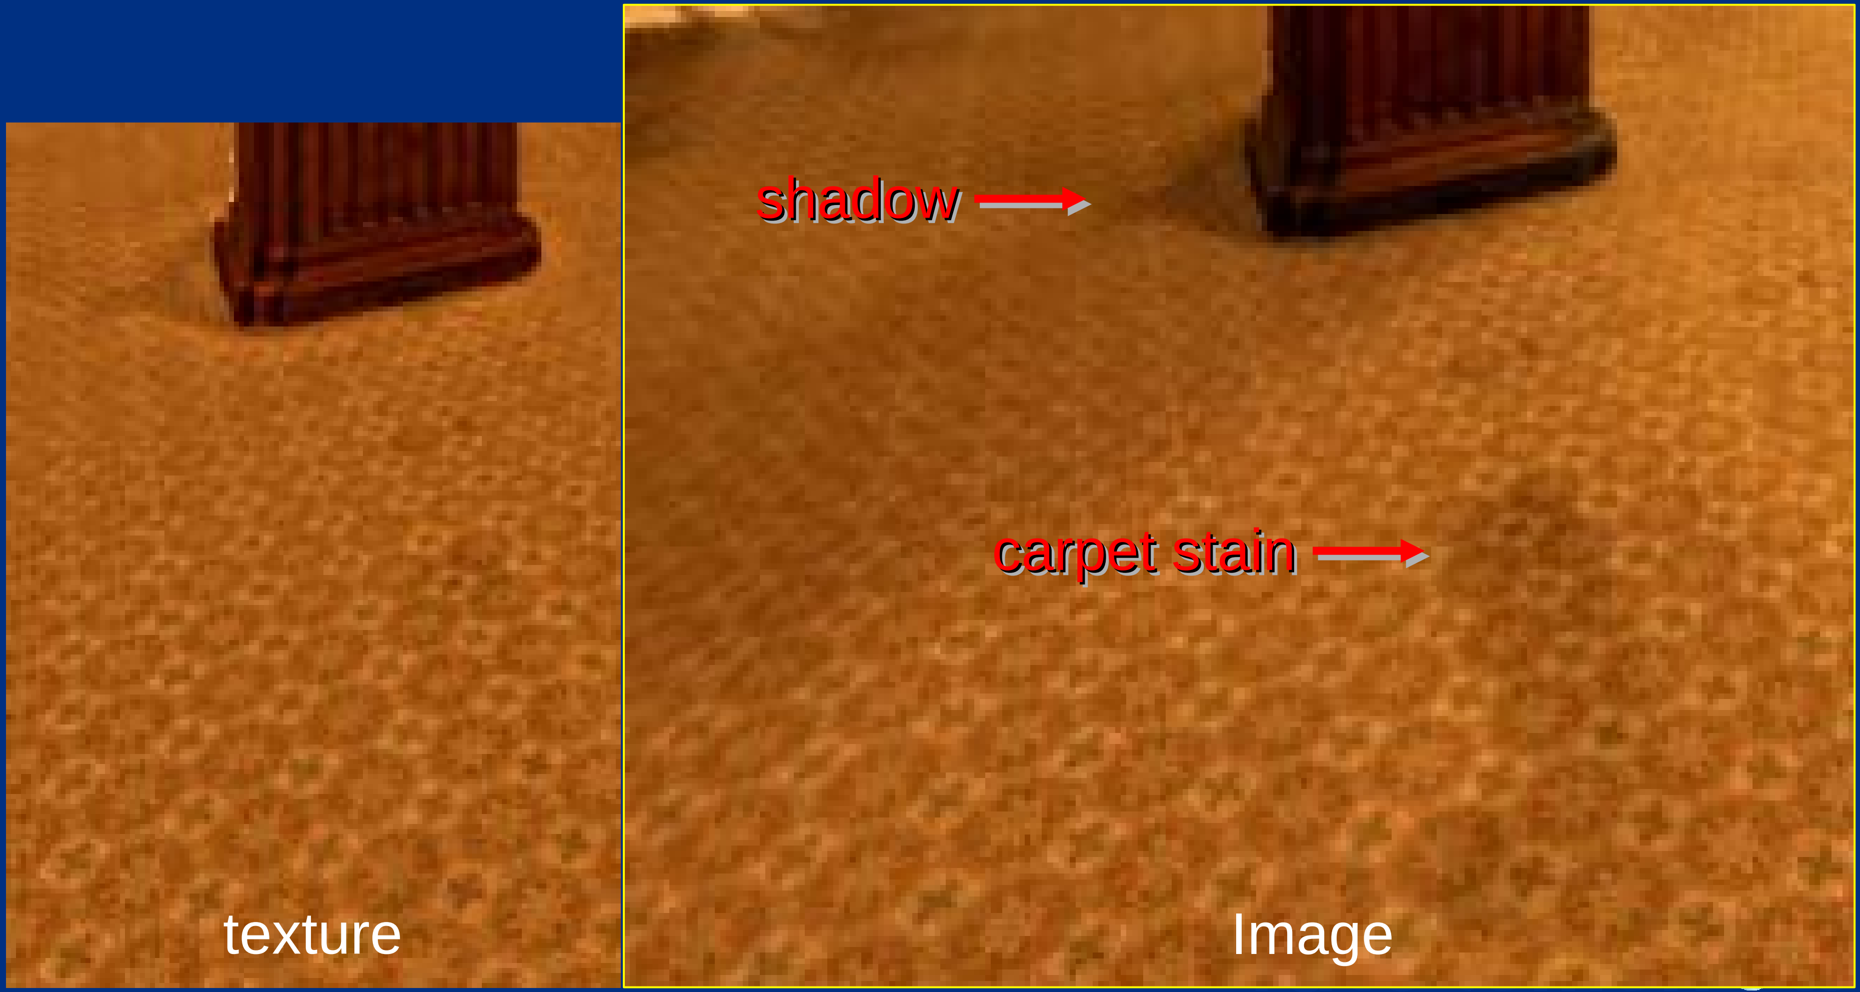

Problems with the Naïve Approach

Problems with the Naïve Approach

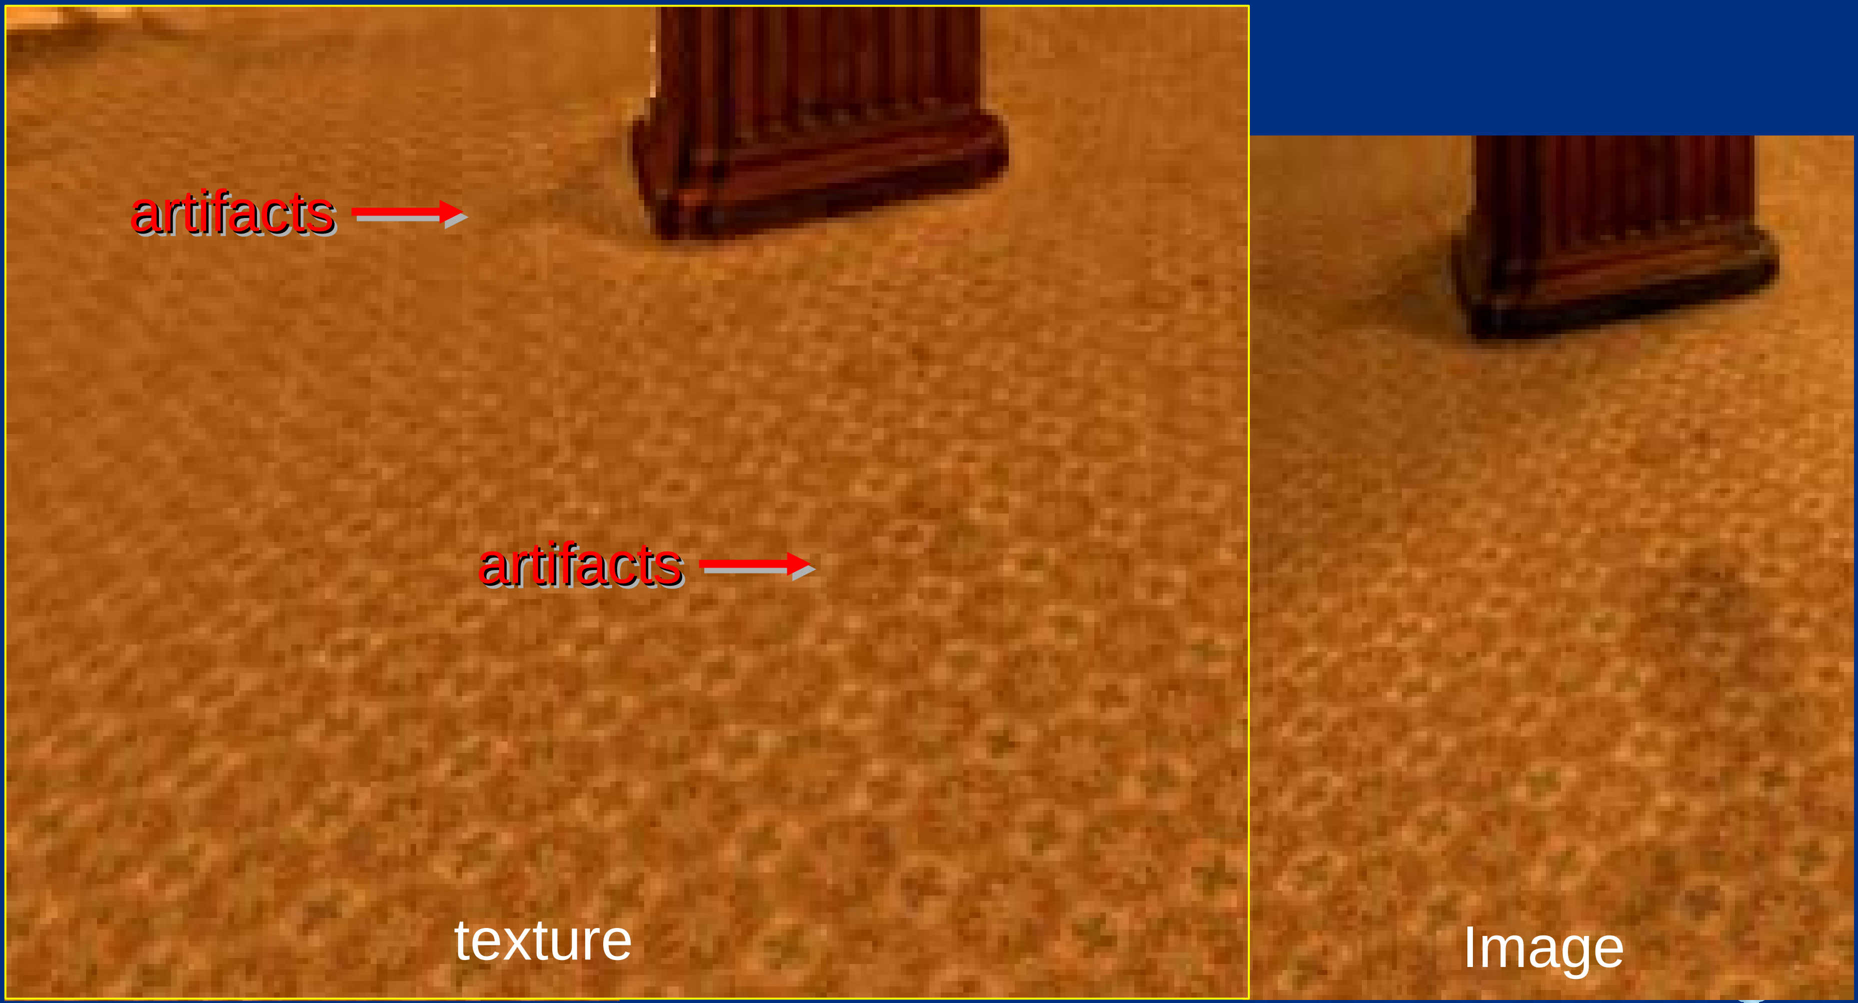

Problems with the Naïve Approach

- Failure due to texture foreshortening
- Artifacts at shadow boundaries

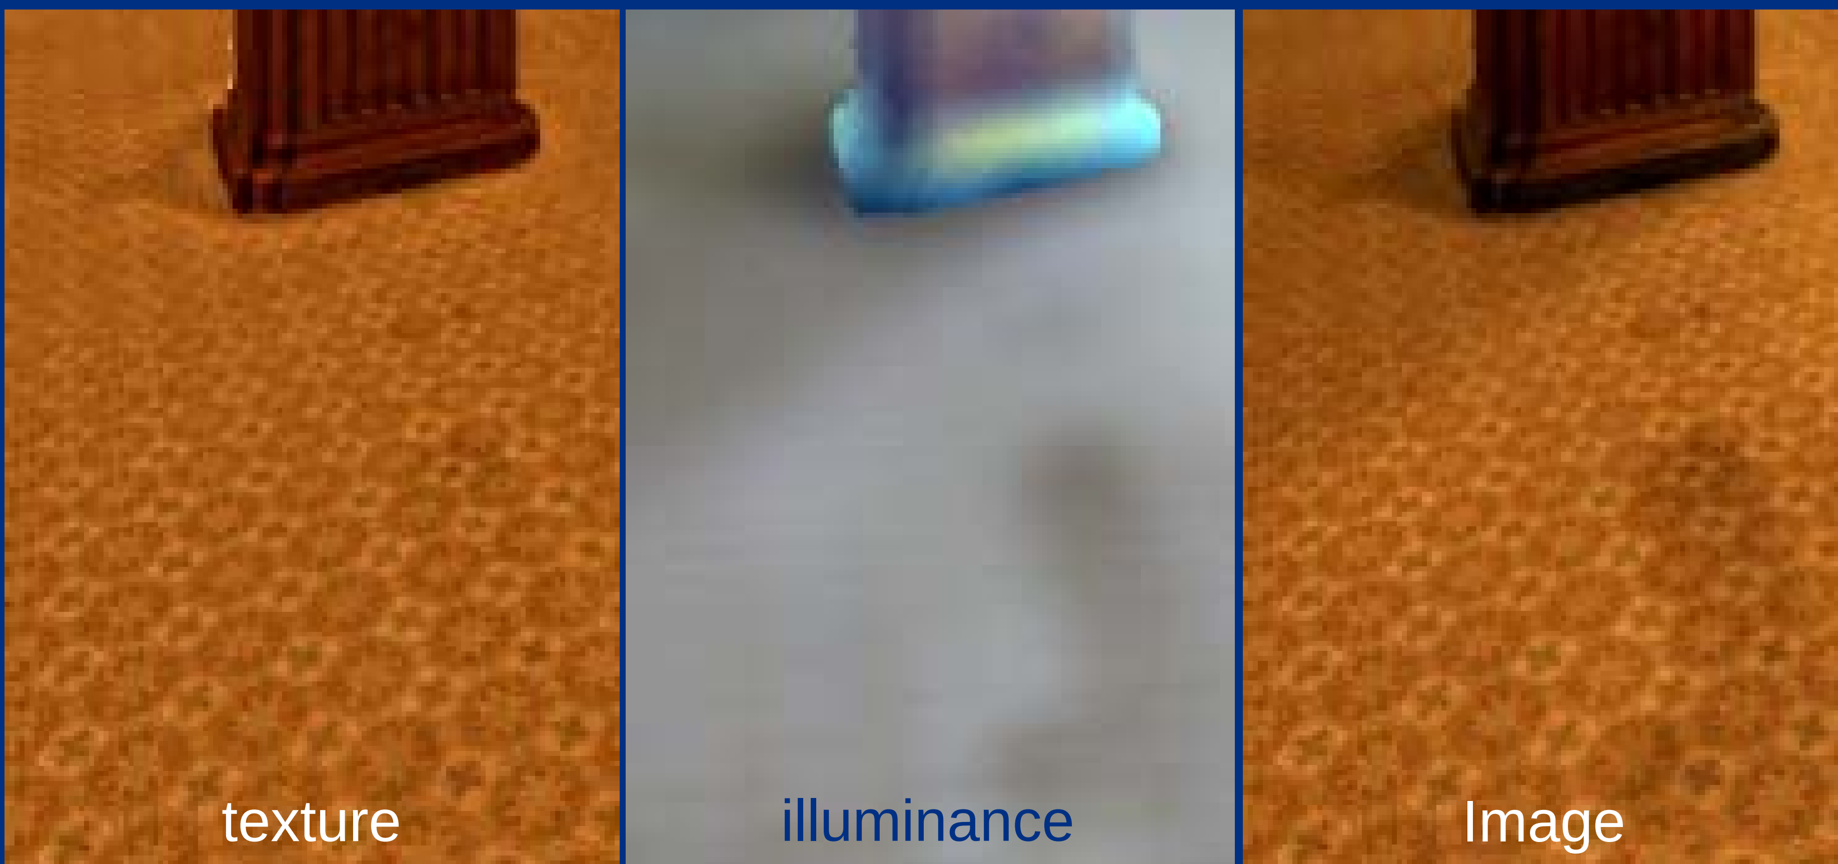

Treatment of Foreshortening

pixel

kernel size

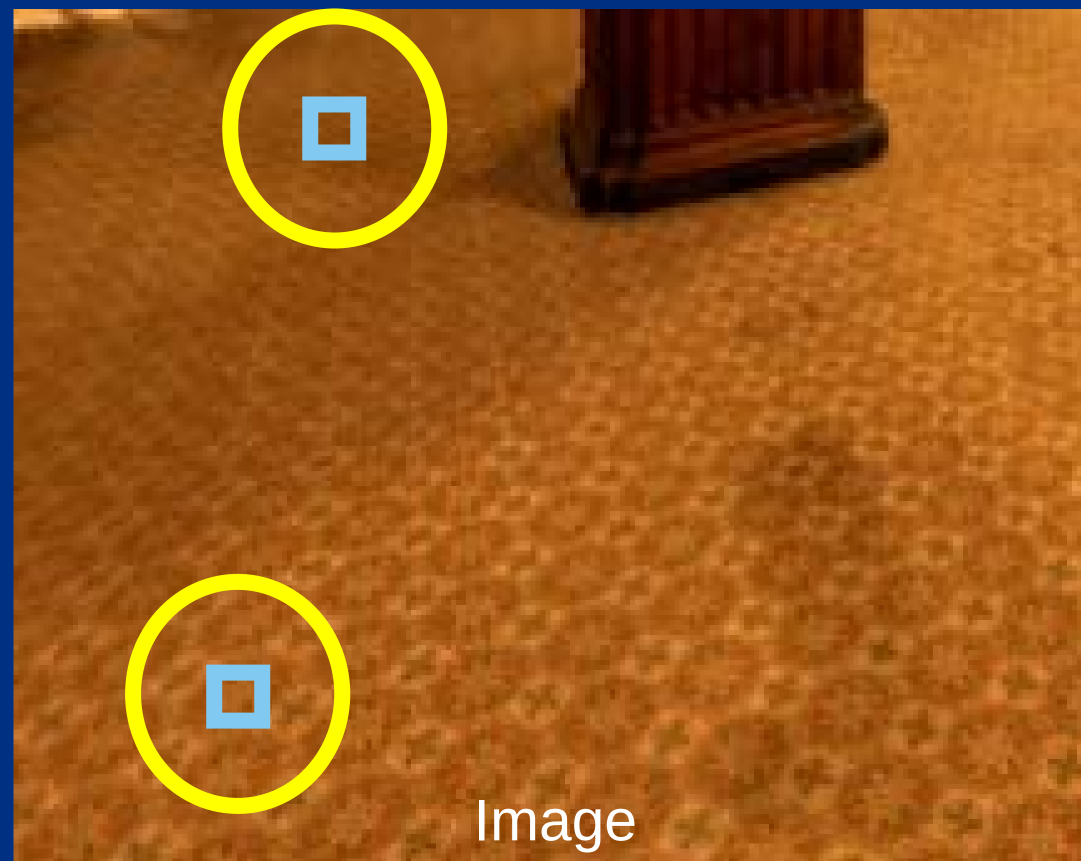

Treatment of Foreshortening

- Blurring depends on distance and orientation

pixel

kernel size

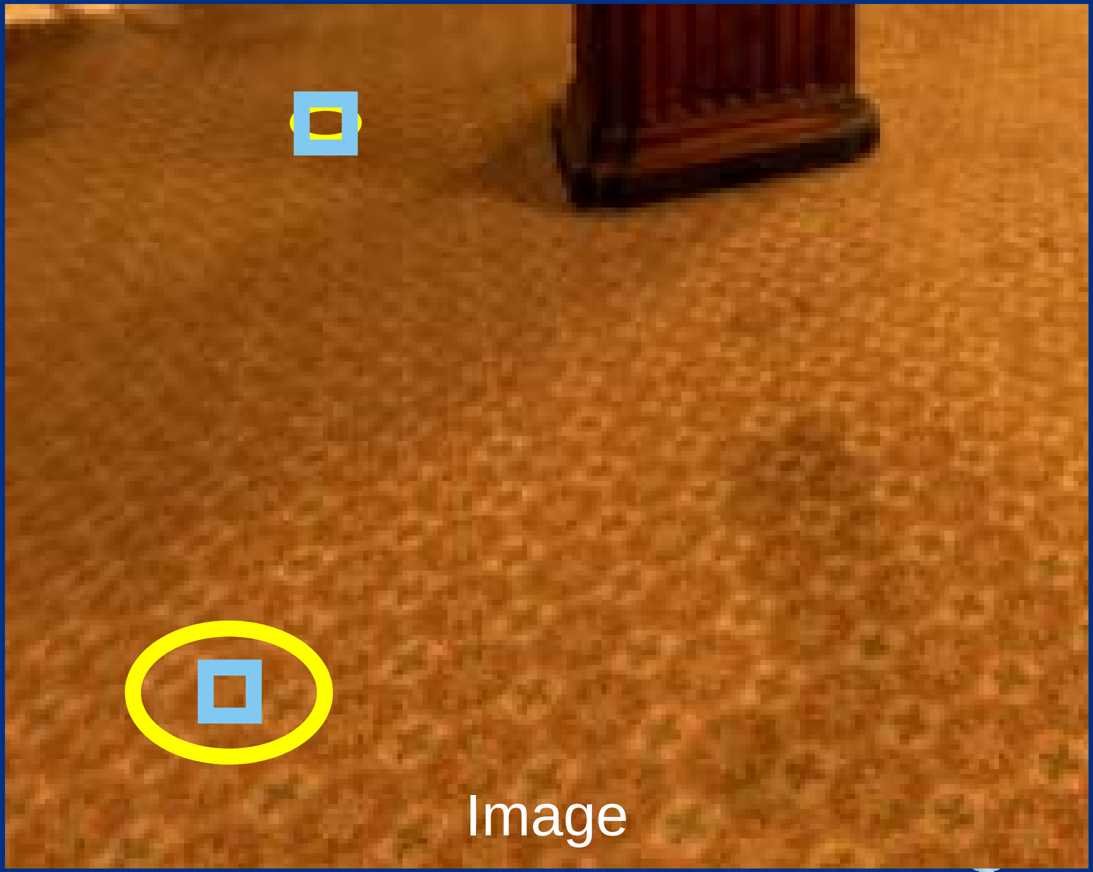

Edge-Preserving Filter

texture | illuminance

Naïve

Bilateral +
foreshortening

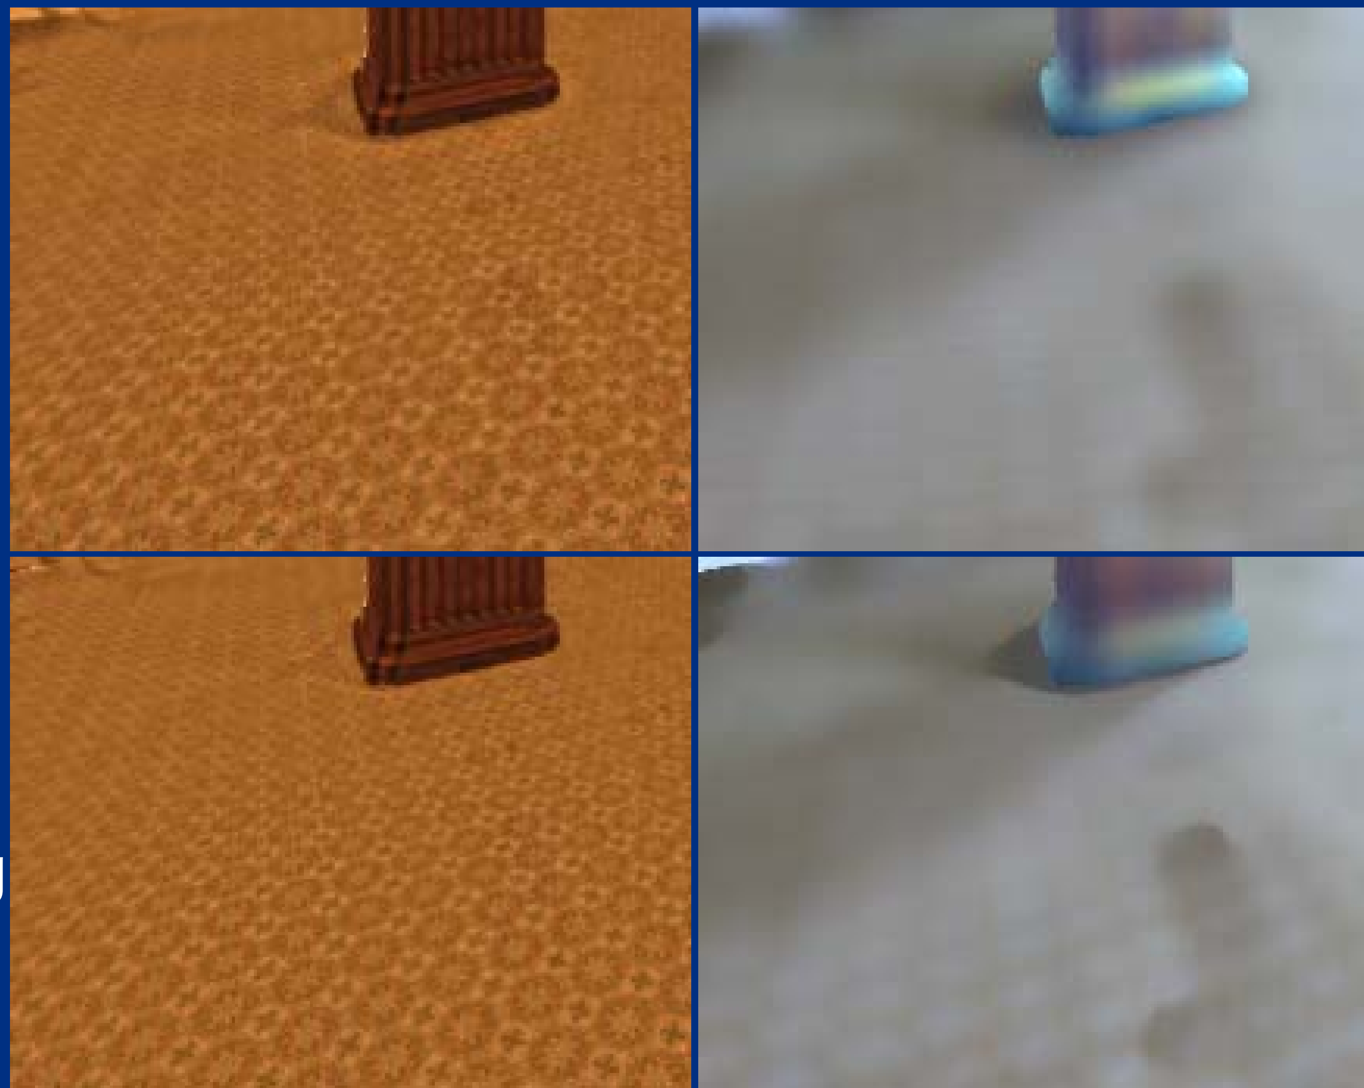

A Simple Relighting Example

Examples

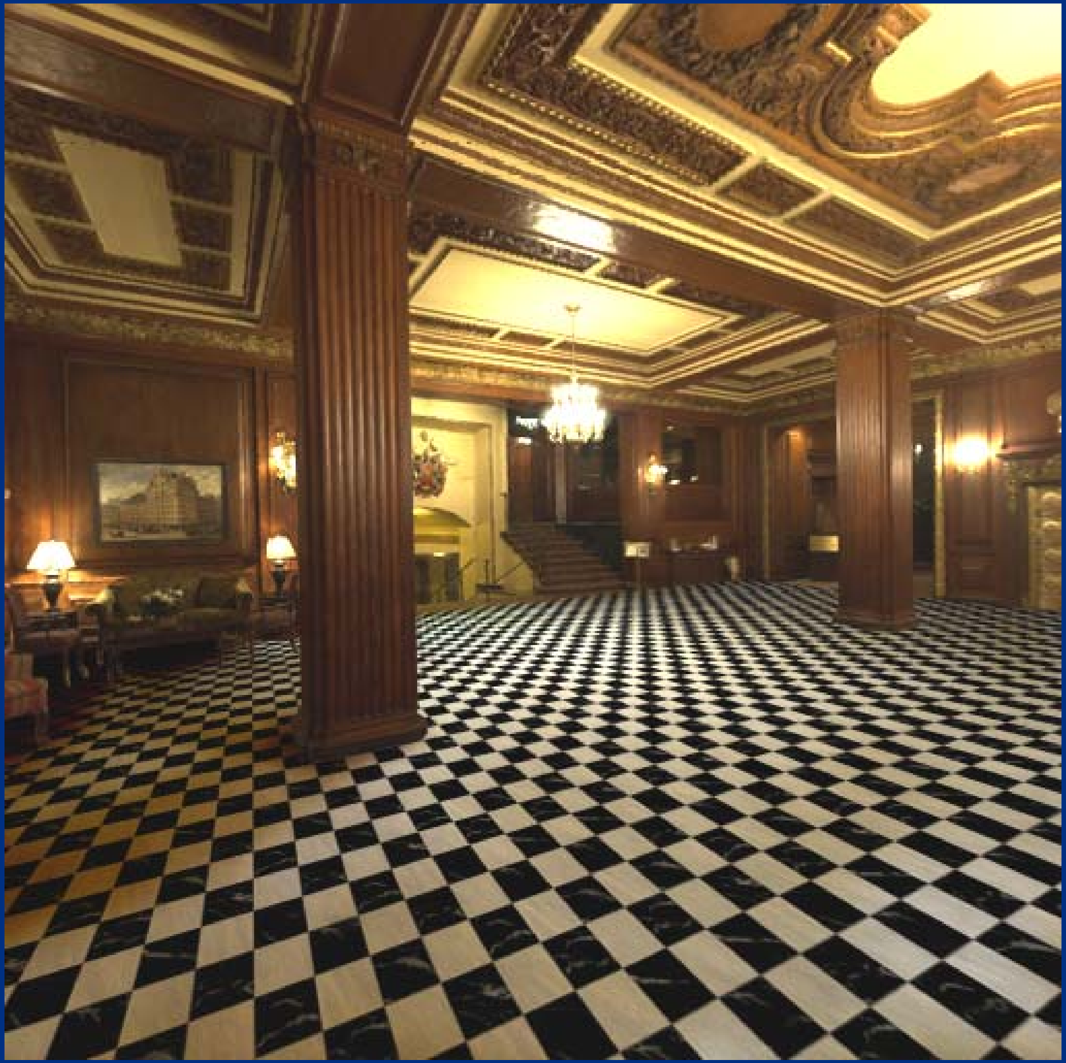

HDR hallucination

- *[Lvdi Wang, Liyi Wei, Kun Zhou, Baining Guo, Heung-Yeung Shum, EGSR 2007]*
- Low-dynamic-range images have under- and over-exposed parts
  - Information missing

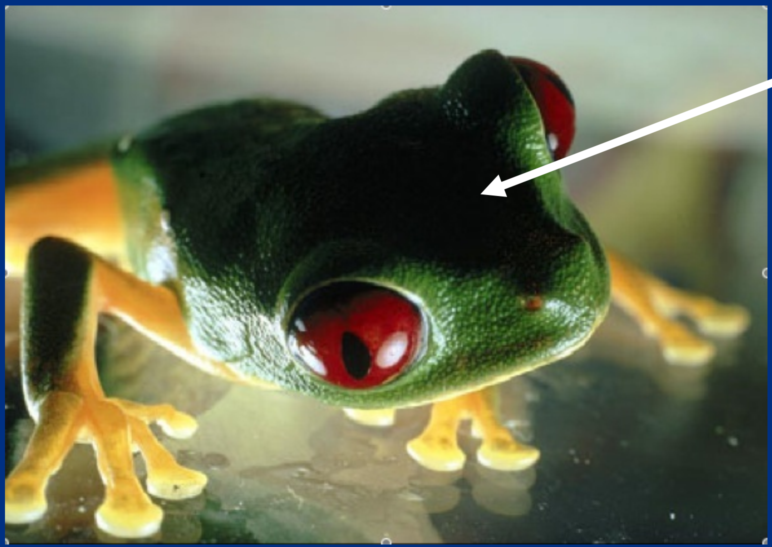

Under-
exposed

Over-
exposed

HDR hallucination

- Separate illumination and texture (Bilateral!)

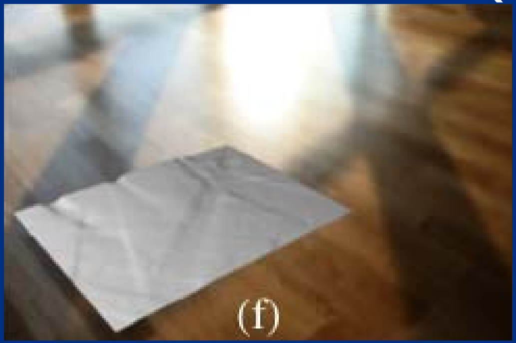

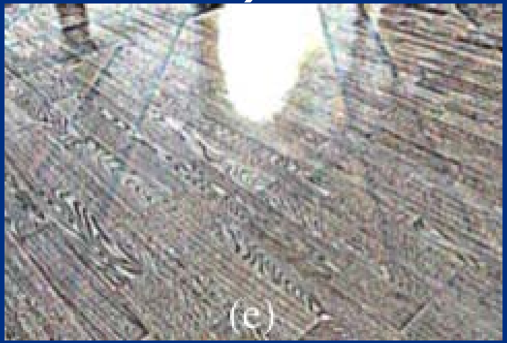

- Fit smooth function to illumination

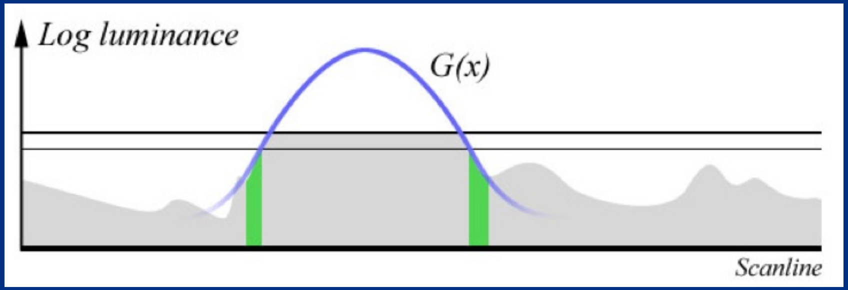

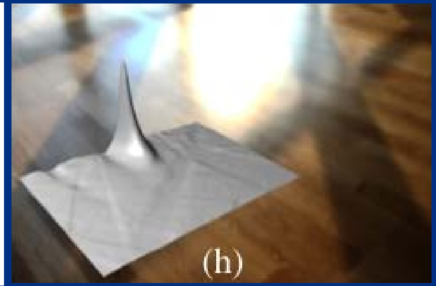

- Use texture synthesis

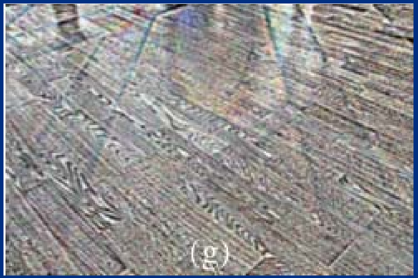

HDR hallucination

input

output

HDR hallucination

input

output

Recap

- Decompose into
  - Large scale (with bilateral filter)
  - Detail (residual: medium+high frequencies)
    - Use big kernels
- Use appropriate domain (log for HDR)
- Manipulate/process independently
- Tone mapping
- Relighting, HDR hallucination
- HDR hallucination

SIGGRAPH2007

Denoising with the BF

Input

output

Images courtesy of Ce Liu

Noise level adaptation [Liu et al. 06]

- The noise of digital cameras varies with intensity
- Adapt sigma range to noise level
- Main topic of the paper: noise level inference

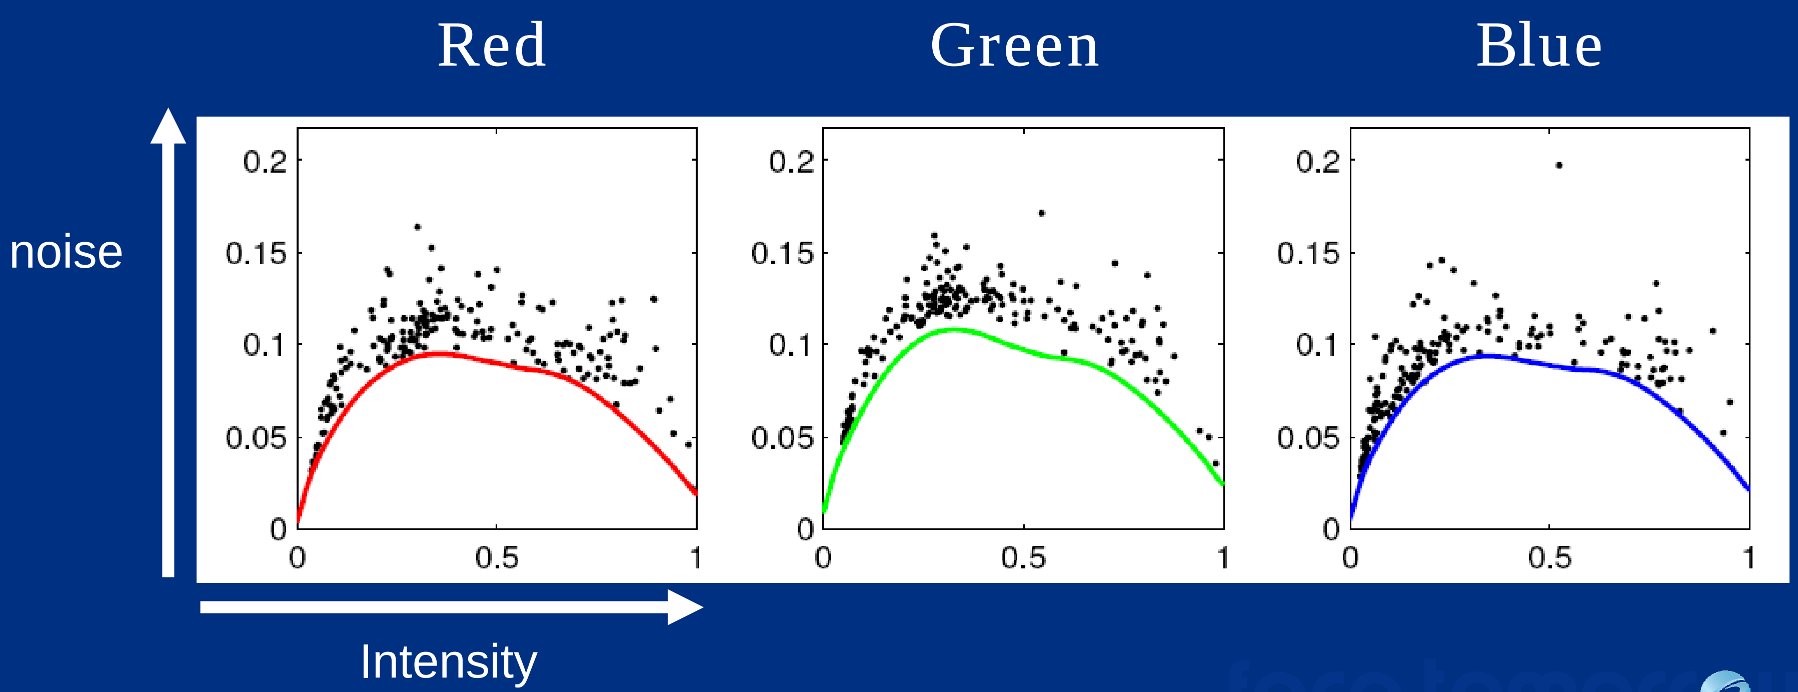

Adaptive bilateral filter

- Adapt sigma range to intensity

Input | Bilateral | Adaptive bilateral

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/
- http://people.csail.mit.edu/sparis/bf/#code