CS6640 Computational Photography

13. Bilateral filtering and HDR tone mapping

© 2012 Steve Marschner

Plan

- **Bilateral filtering**

  another way to think about large and small scale detail

- **HDR tone mapping techniques**

  a case study in applying many different notions of “large” and “small”

slide by Frédo Durand, MIT

Denoising from 1 image

- We can't take average over multiple images

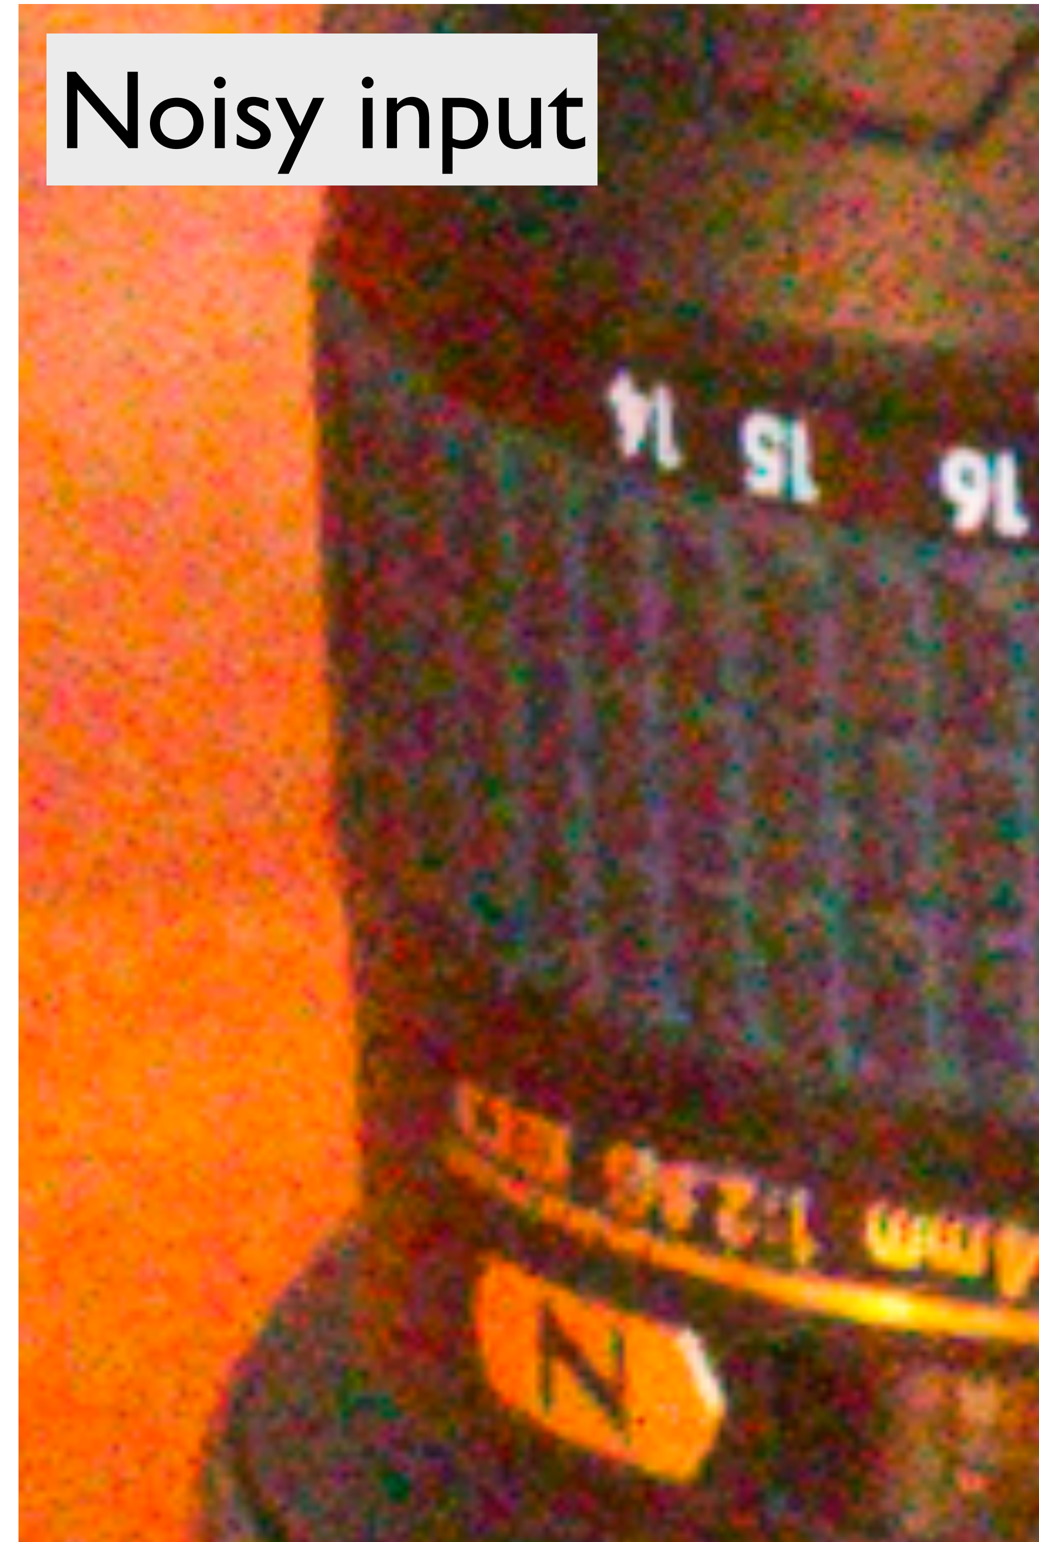

slide by Frédo Durand, MIT

Denoising from 1 image

- We can't take average over multiple images
- Idea 1: take a spatial average
  - Most pixels have roughly teh same color as their neighbor
  - Noise looks high frequency => do a low pass
- Here: Gaussian blur

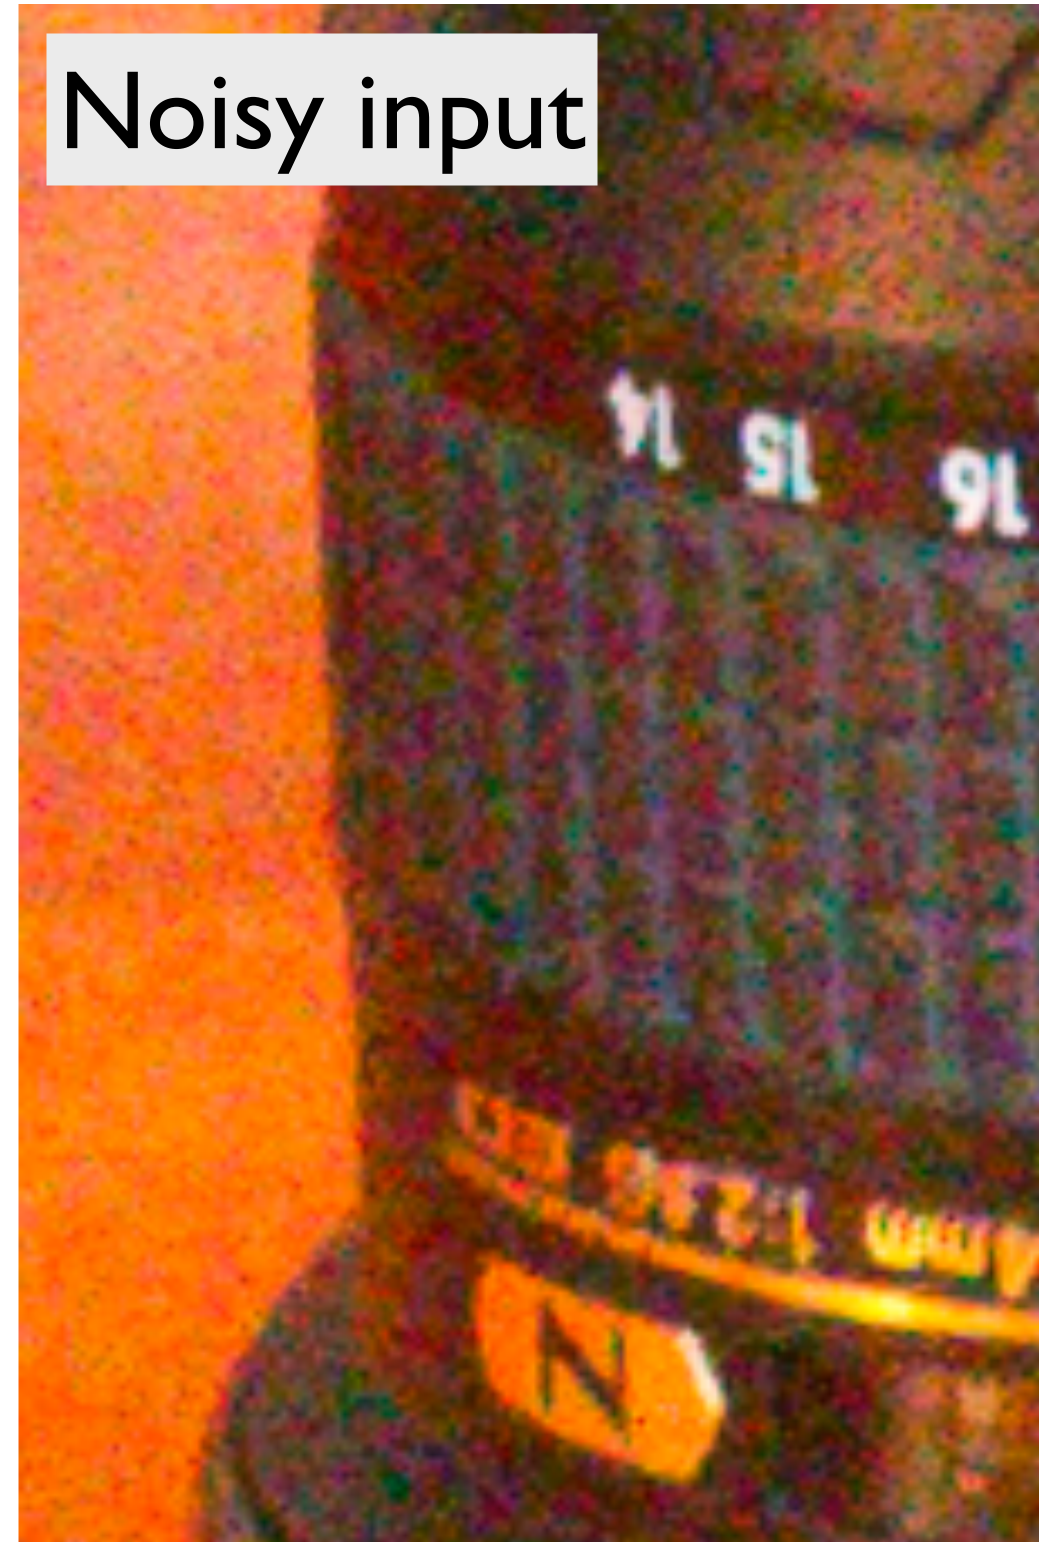

slide by Frédo Durand, MIT

Gaussian blur

- Noise is mostly gone
- But image is blurry
  - duh!

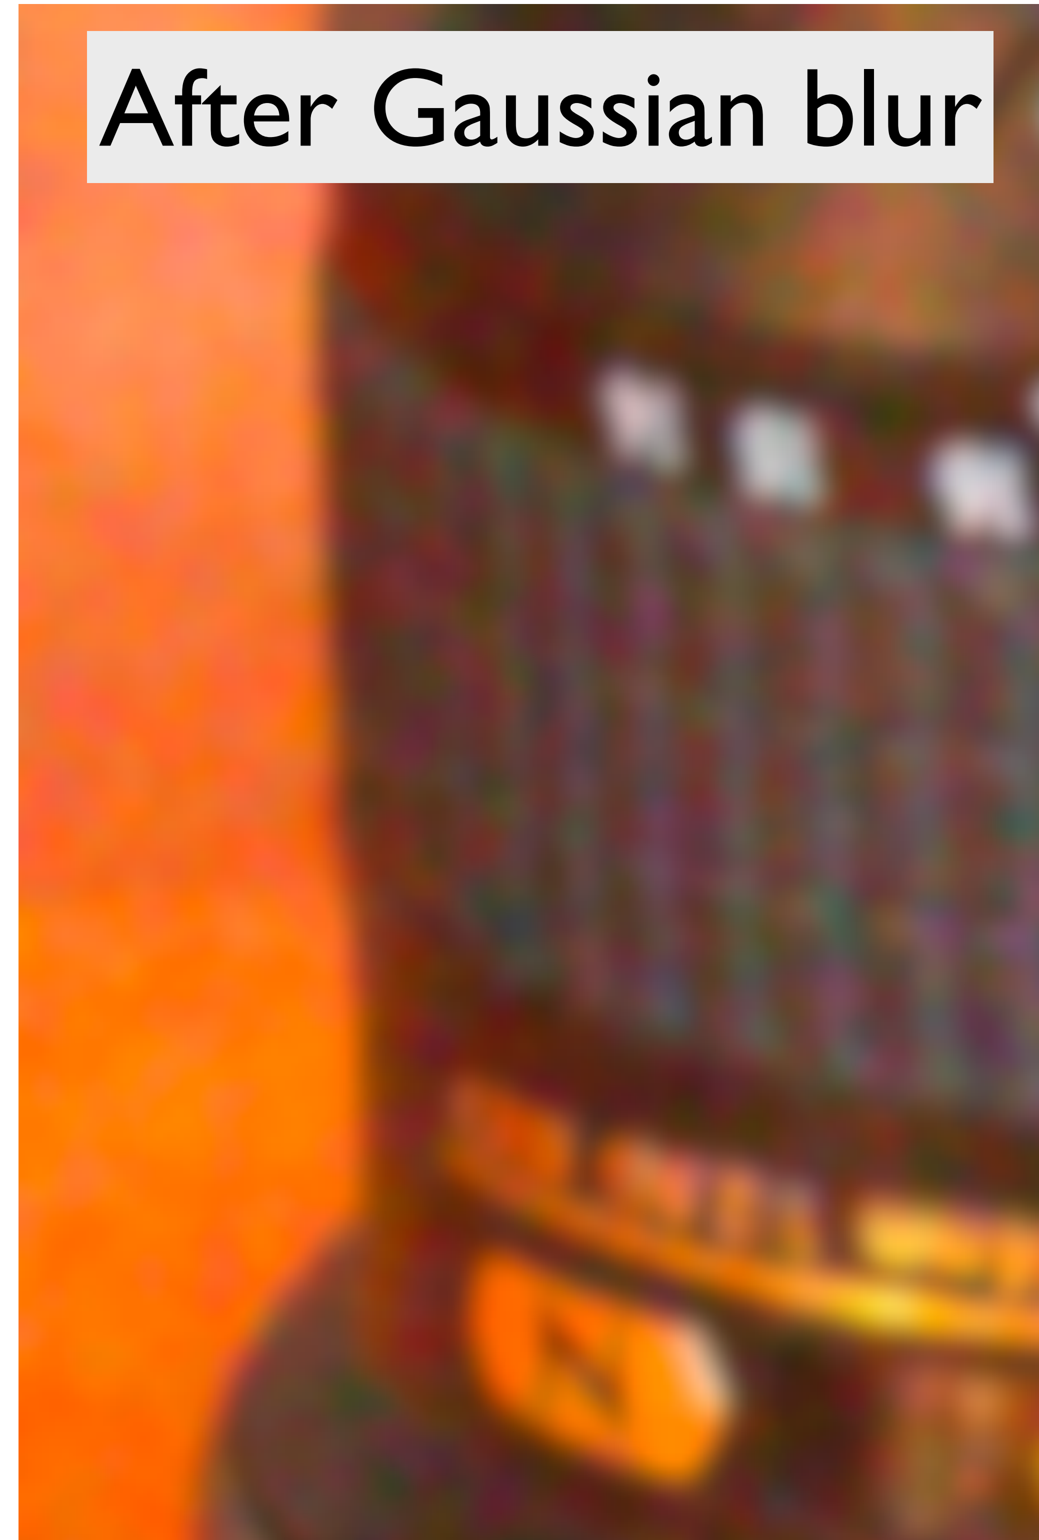

adapted from slide by Frédo Durand, MIT

Gaussian blur

- Noise is mostly gone
- But image is blurry
  - duh!

- Question: how to blur/smooth/abstract image, but without destroying important features?

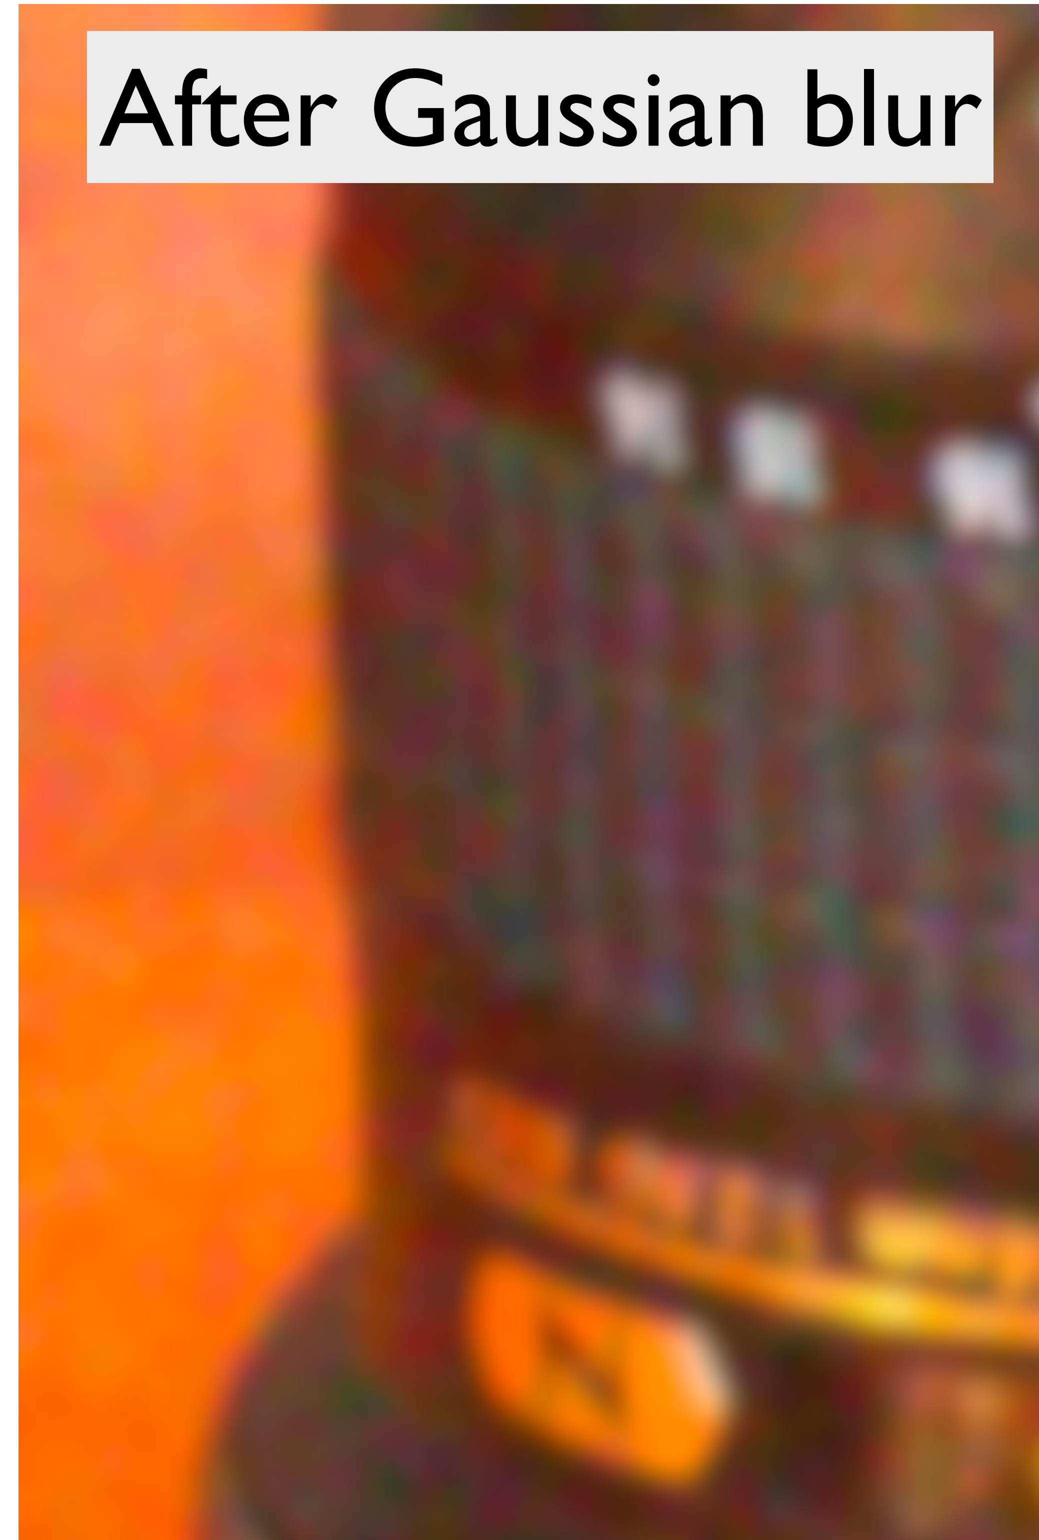

Anisotropic diffusion

- **A diffusion equation will evolve an image to become smoother**

$$I_t(x, y) = -C\,(\nabla^2 I)(x, y)$$

↑ "heat conductance"

- **Make conductance variable to smooth some features less**

$$I_t(x, y) = -C(x, y)\,(\nabla^2 I)(x, y)$$

↑ conductance inversely related to "edginess"

- **Calculate "edginess" from current image**

$$I_t(x, y) = -C(x, y, t)\,(\nabla^2 I)(x, y)$$

- **Result: positive reinforcement of edges, producing shocks**

Anisotropic Diffusion

- **Smoothing with different weights in "edginess" calculation leads to progressive abstraction—a scale space**

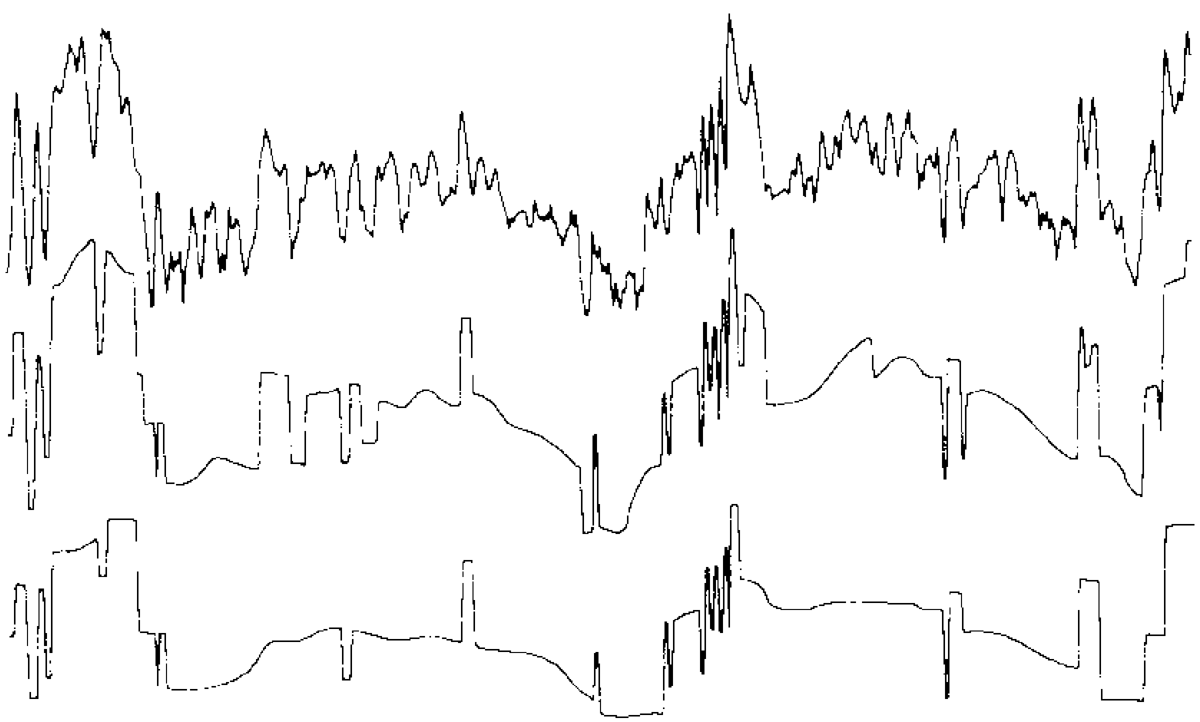

Fig. 11. Scale space obtained with anisotropic diffusion. The diffusion was performed in 2-D on the Canaletto image of which one line (the horizontal line number 400 out of 480—just above the gondola) is shown. Notice that the edges remain sharp until their disappearance.

Bilateral filter

- **[Tomasi and Manduci 1998]**
  - http://www.cse.ucsc.edu/~manduchi/Papers/ICCV98.pdf
- **Developed for denoising**
- **Related to**
  - SUSAN filter [Smith and Brady 95] http://citeseer.ist.psu.edu/smith95susan.html
  - Digital-TV [Chan, Osher and Chen 2001] http://citeseer.ist.psu.edu/chan01digital.html
  - sigma filter http://www.geogr.ku.dk/CHIPS/Manual/f187.htm
- **Full survey: http://people.csail.mit.edu/sparis/publi/2009/fntcgv/Paris_09_Bilateral_filtering.pdf**

slide by Frédo Durand, MIT

Start with Gaussian filtering

- **Here, input is a step function + noise**

$$J \quad = \quad f \quad \otimes \quad I$$

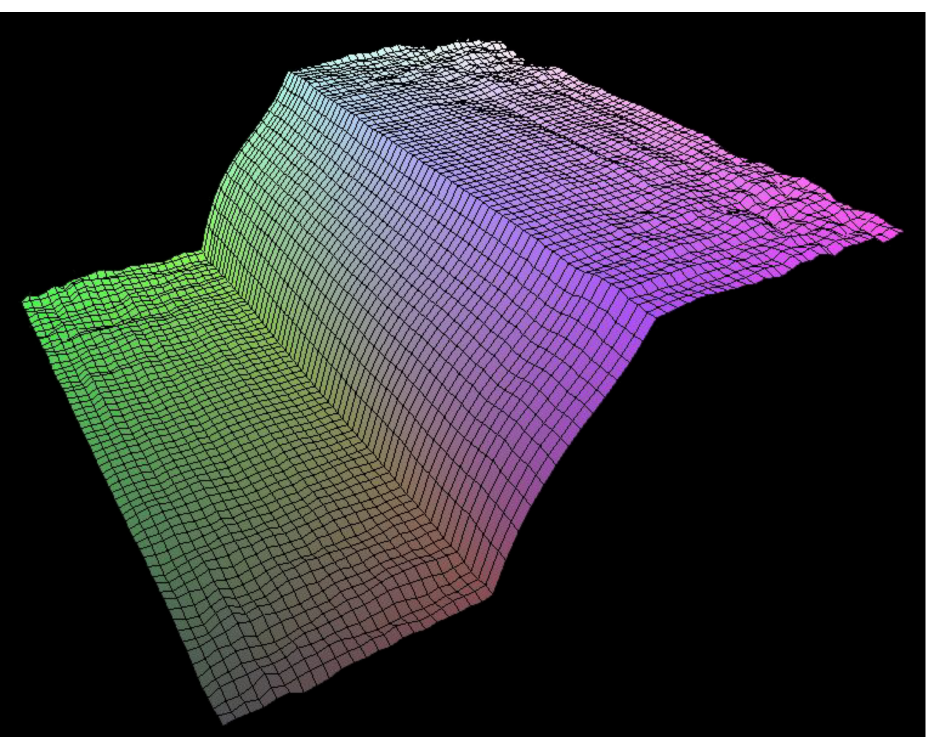

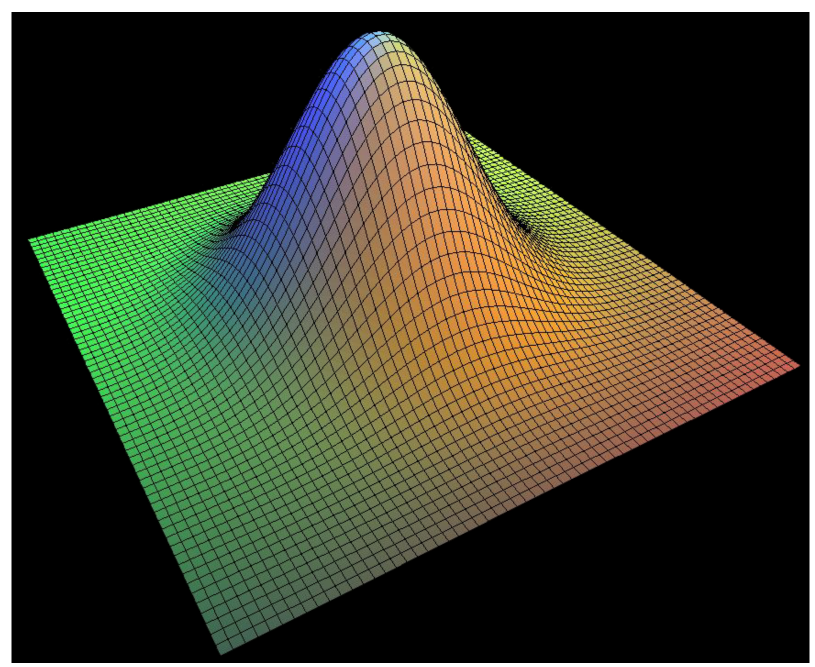

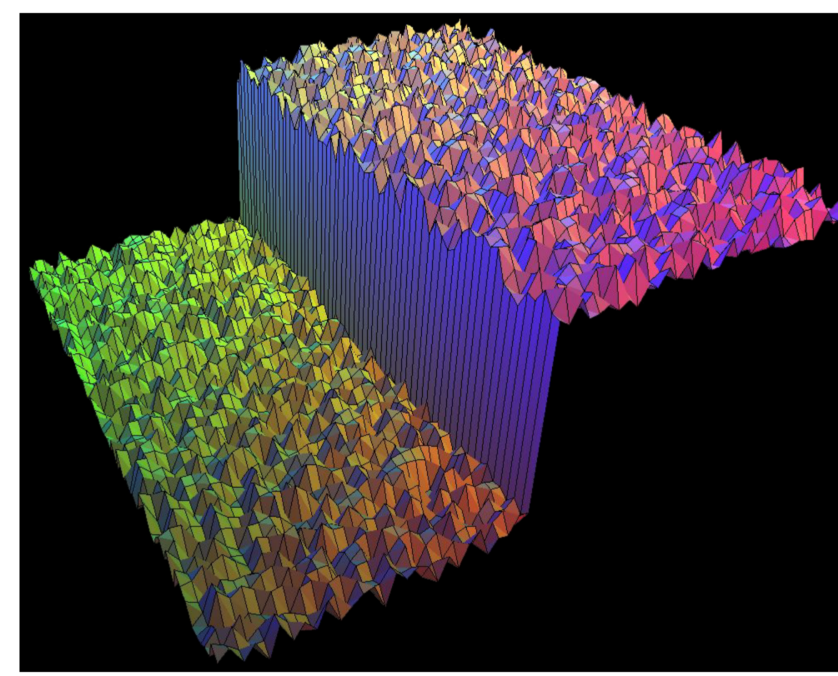

output — input

urand, MIT

Gaussian filter as weighted average

- **Weight of ξ depends on distance to x**

$$J(x) = \sum_{\xi} f(x,\xi) \quad I(\xi)$$

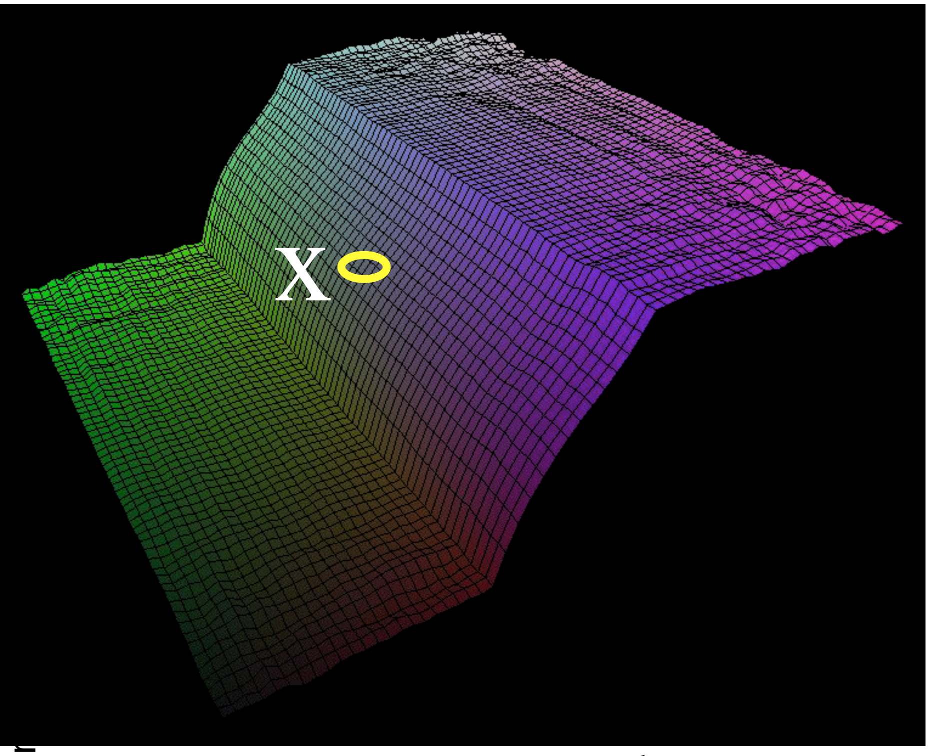

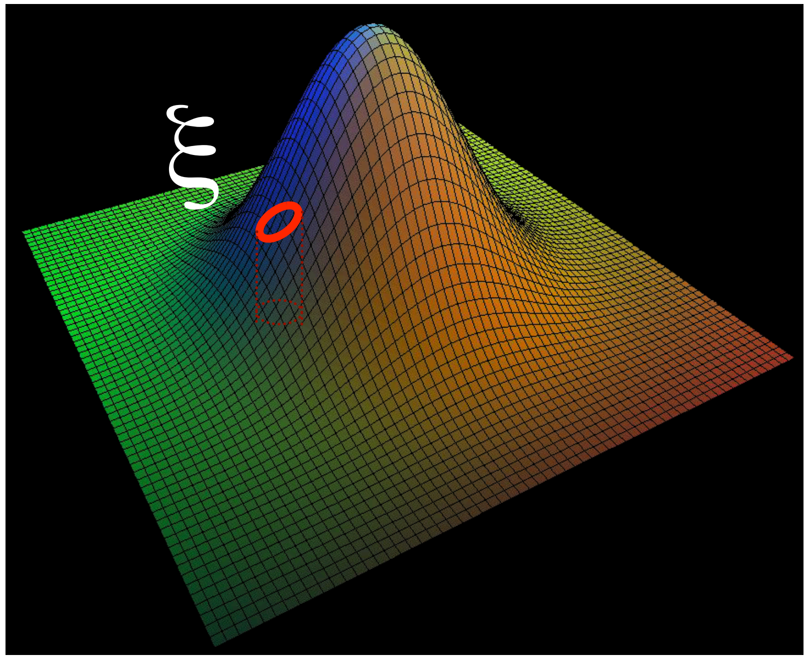

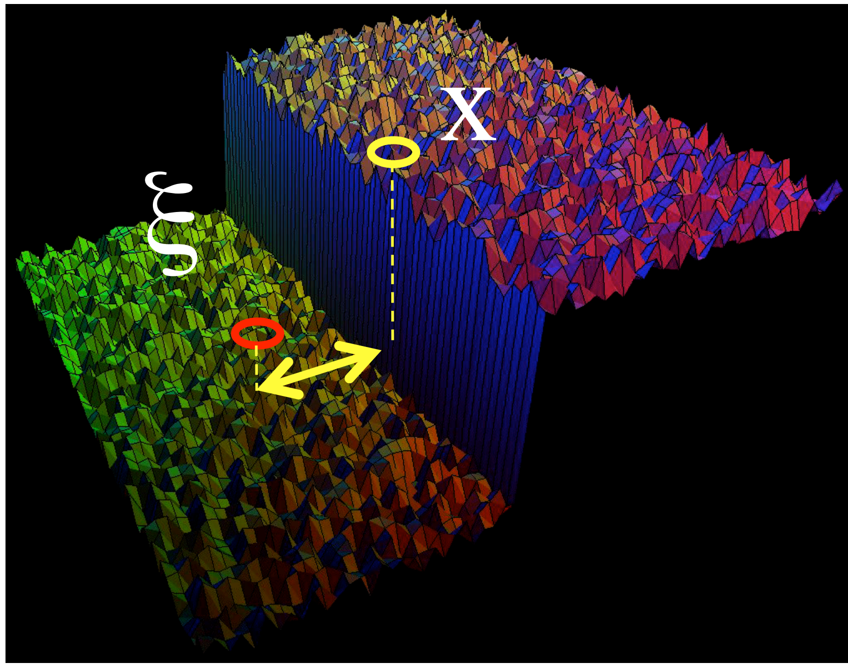

output ⟵ input

urand, MIT

The problem of edges

- Here, I(ξ) "pollutes" our estimate J(x)
- It is too different

$$J(x) = \sum_{\xi} f(x,\xi) \quad I(\xi)$$

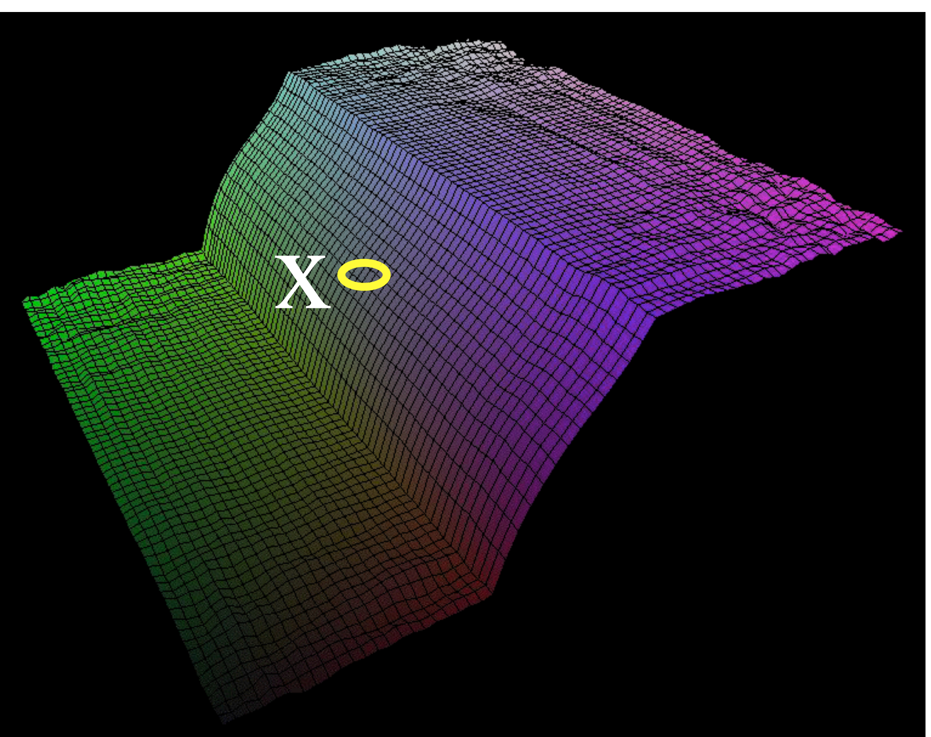

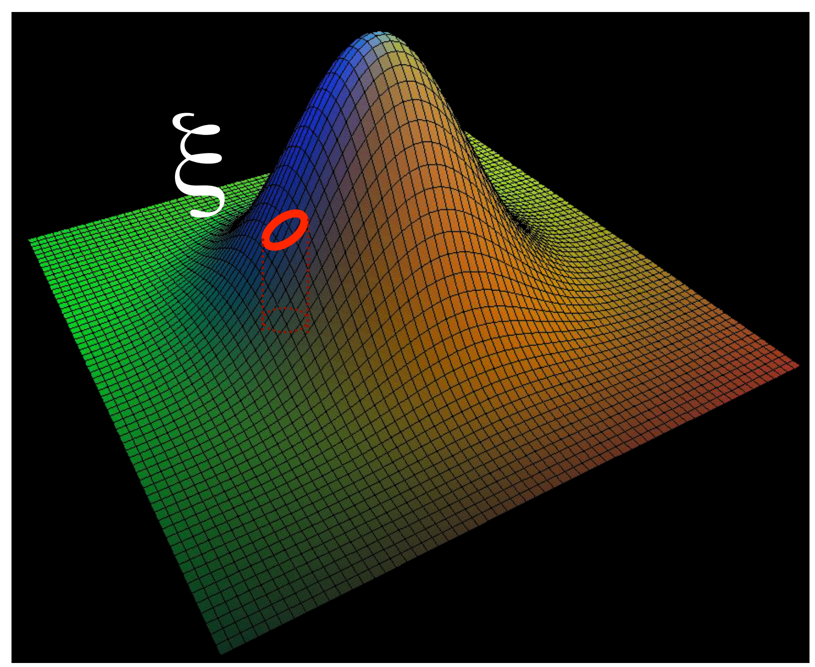

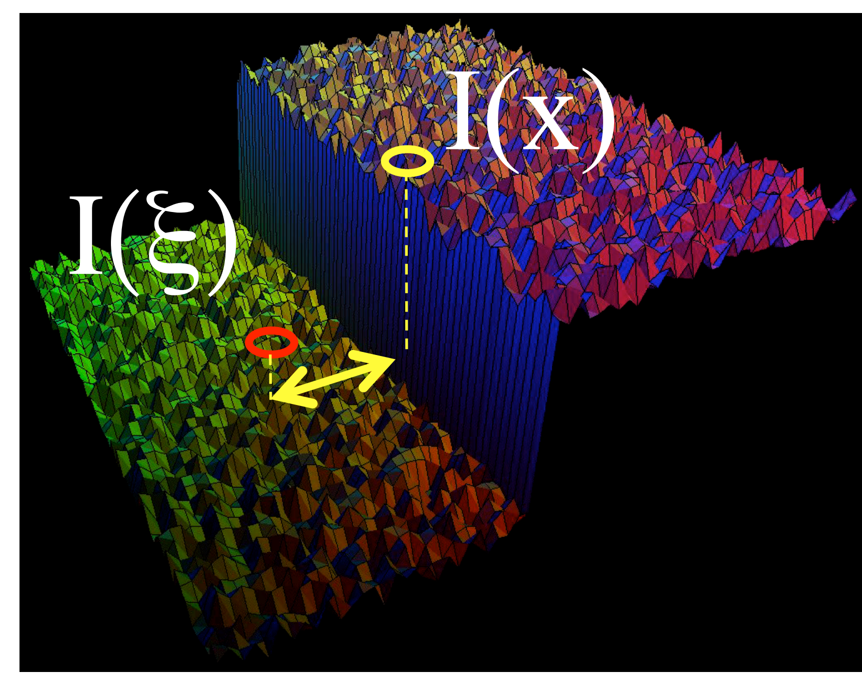

output ⟵ input

Principle of Bilateral filtering

[Tomasi and Manduchi 1998]

- **Penalty g on the intensity difference**

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x,\xi) \quad g(I(\xi) - I(x)) \quad I(\xi)$$

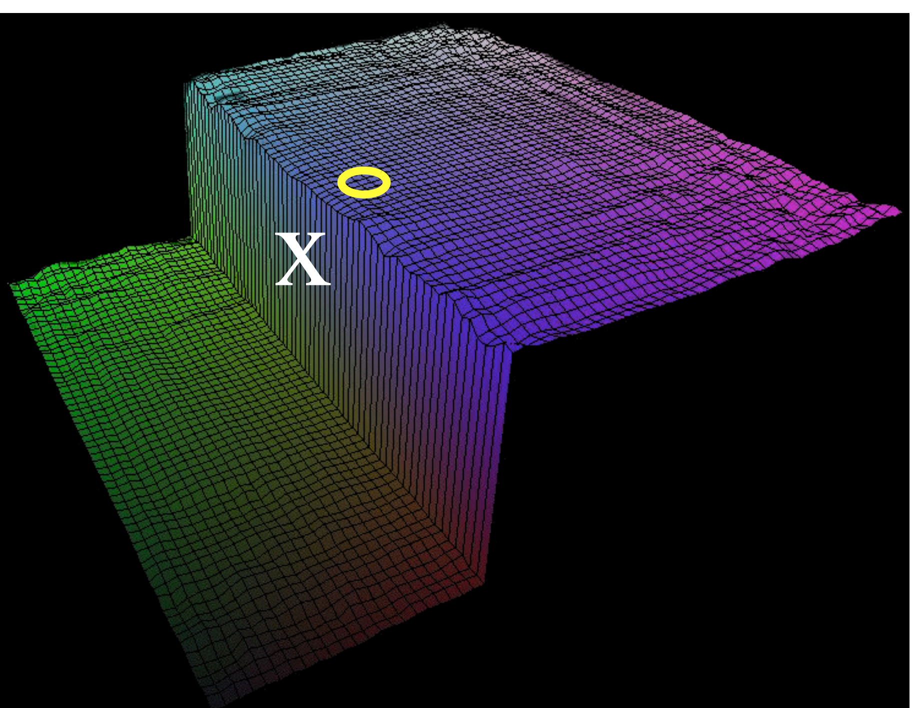

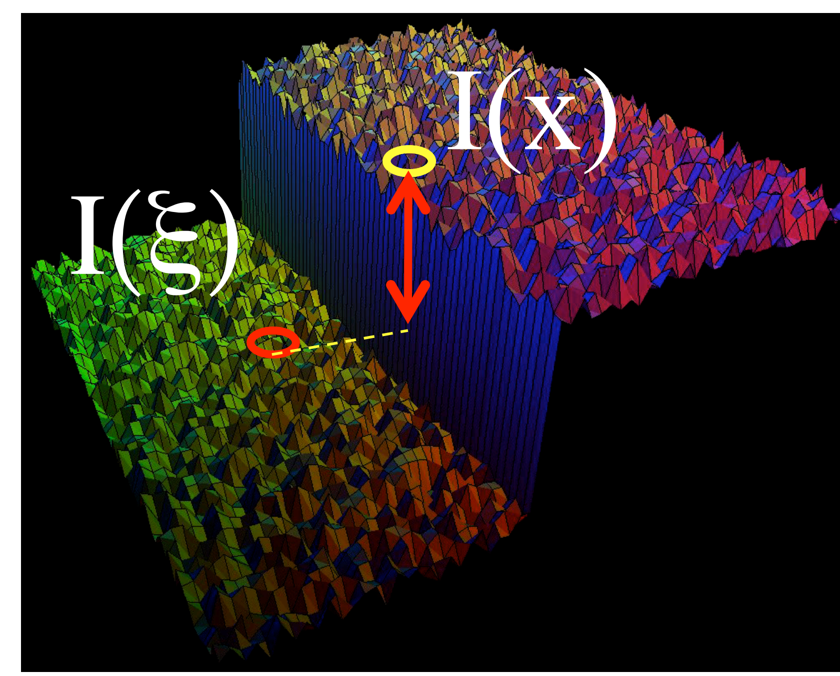

output ⟵ input

urand, MIT

Bilateral filtering

[Tomasi and Manduchi 1998]

- **Spatial Gaussian f**

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x,\xi)\, g(I(\xi) - I(x))\quad I(\xi)$$

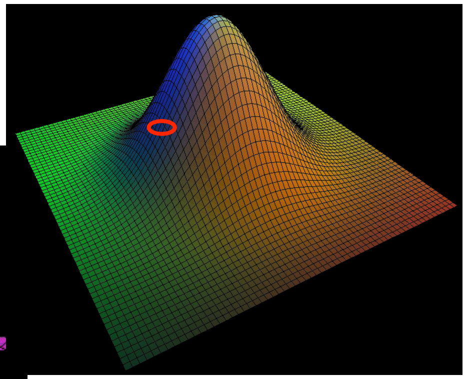

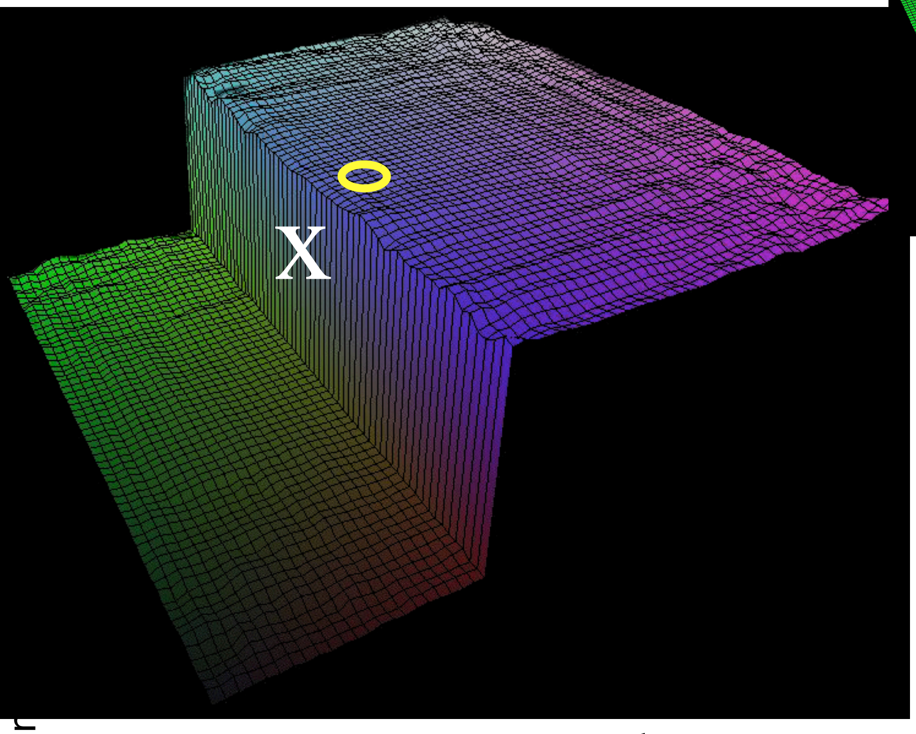

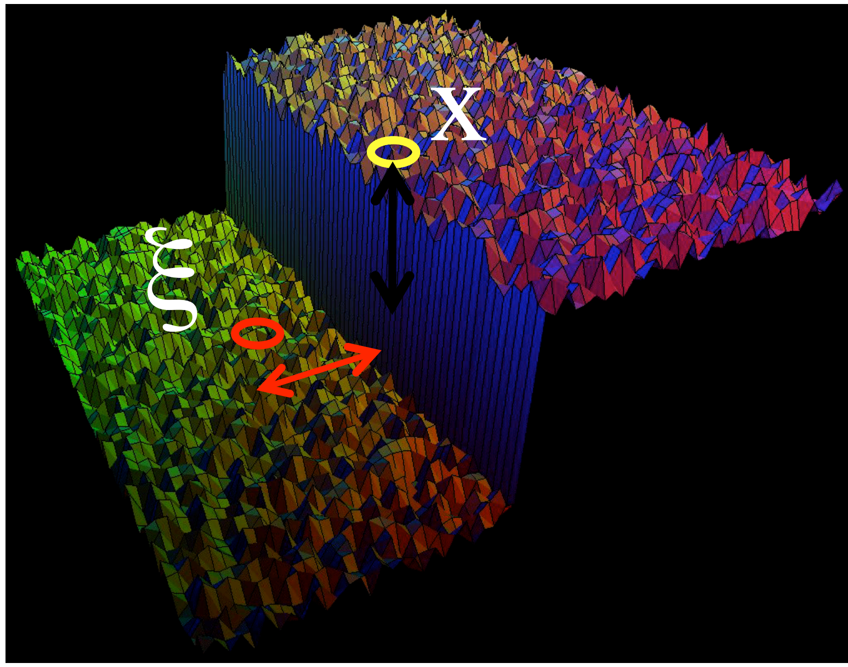

output ⟸ input

Bilateral filtering

[Tomasi and Manduchi 1998]

- **Spatial Gaussian f**
- **Gaussian g on the intensity difference**

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x,\xi)\, g(I(\xi) - I(x))\, I(\xi)$$

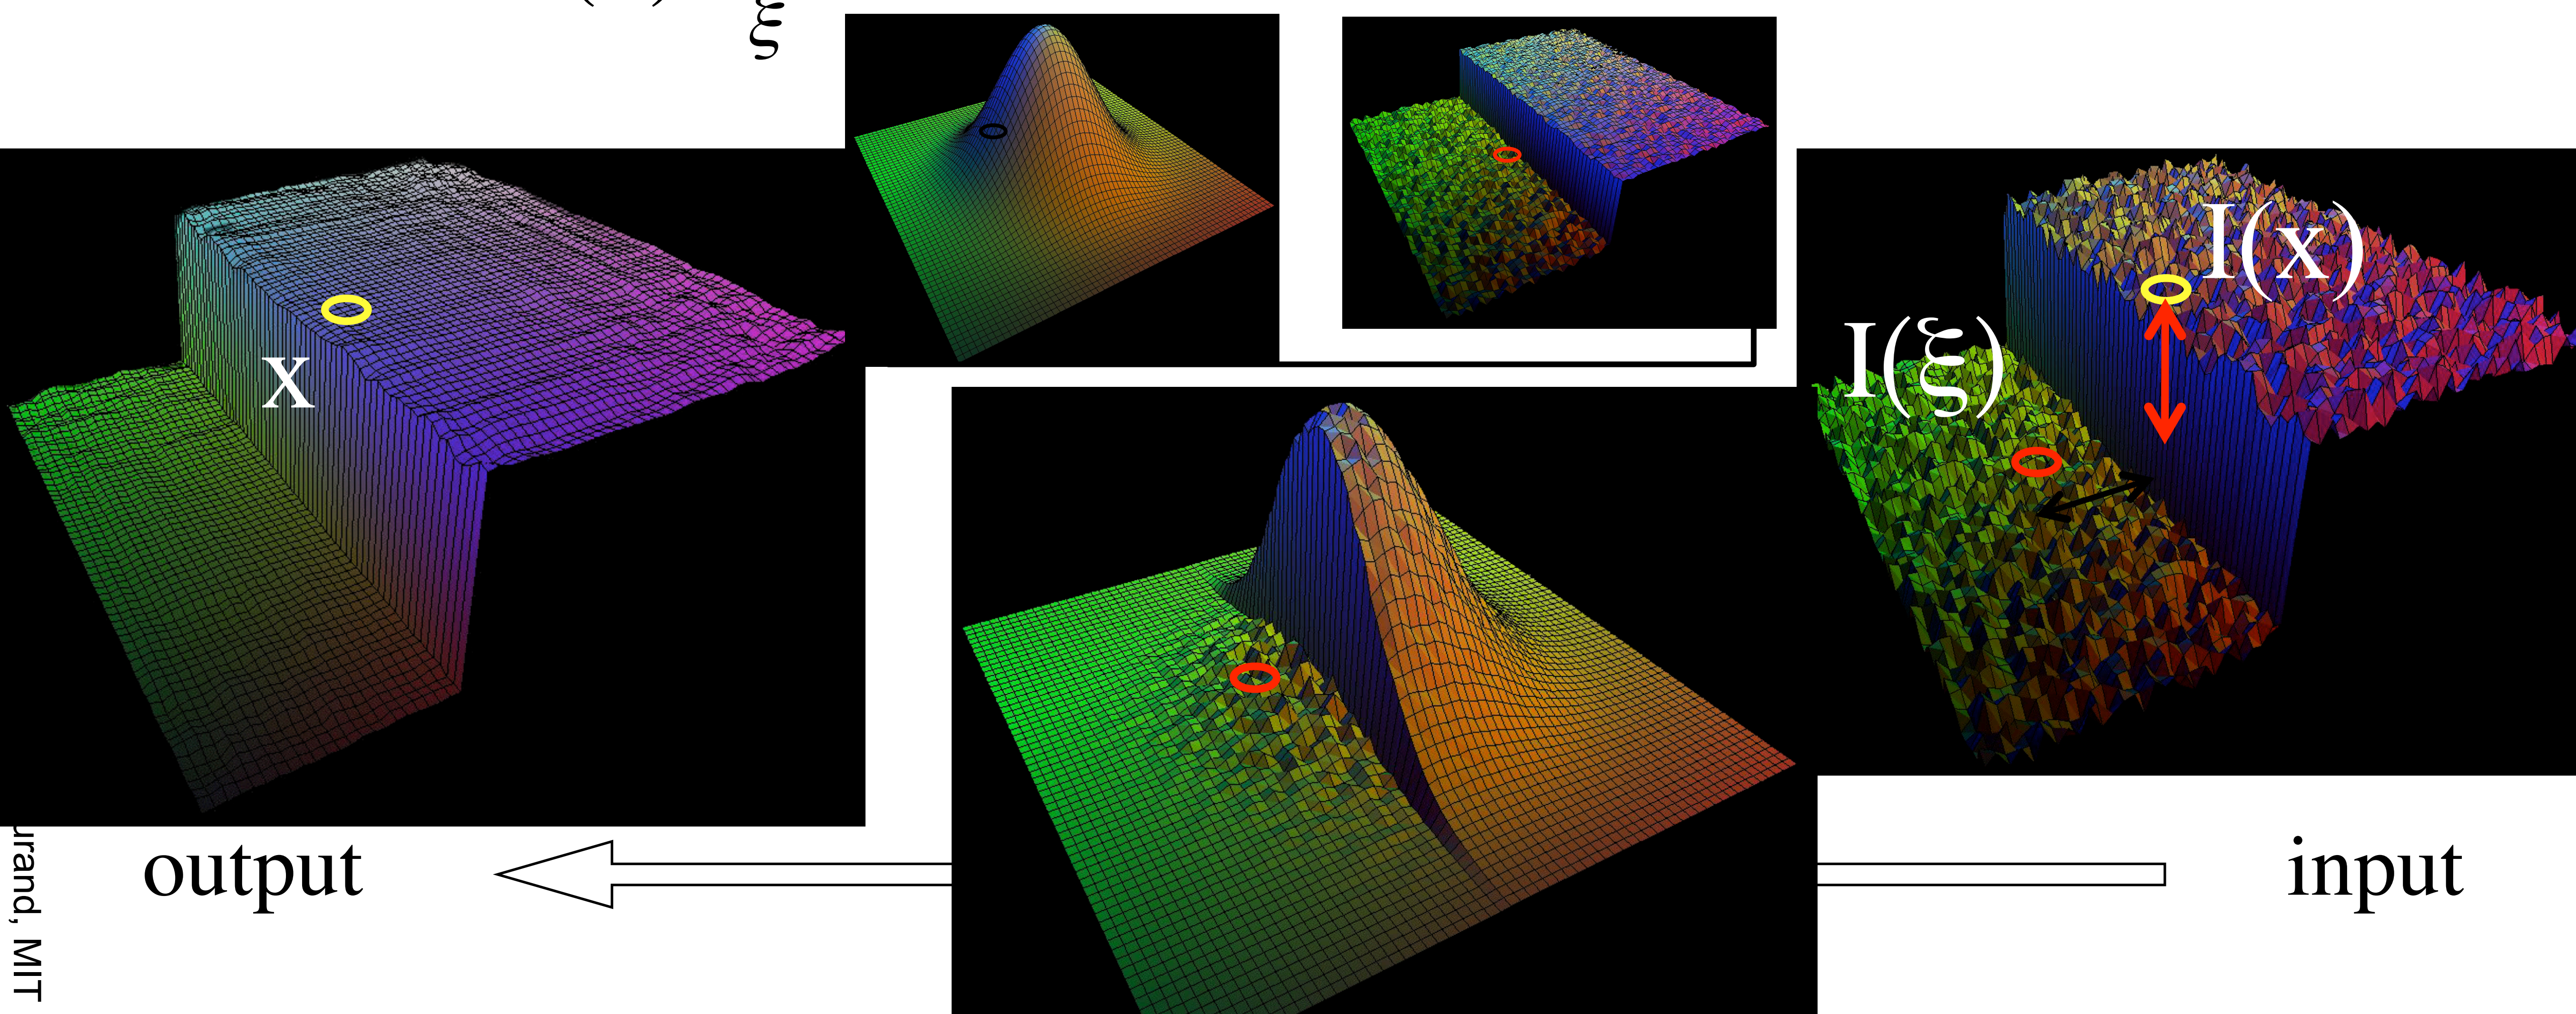

Normalization factor

[Tomasi and Manduchi 1998]

- **k(x)=** $\sum_{\xi} f(x,\xi) \quad g(I(\xi) - I(x))$

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x,\xi) \quad g(I(\xi) - I(x)) \quad I(\xi)$$

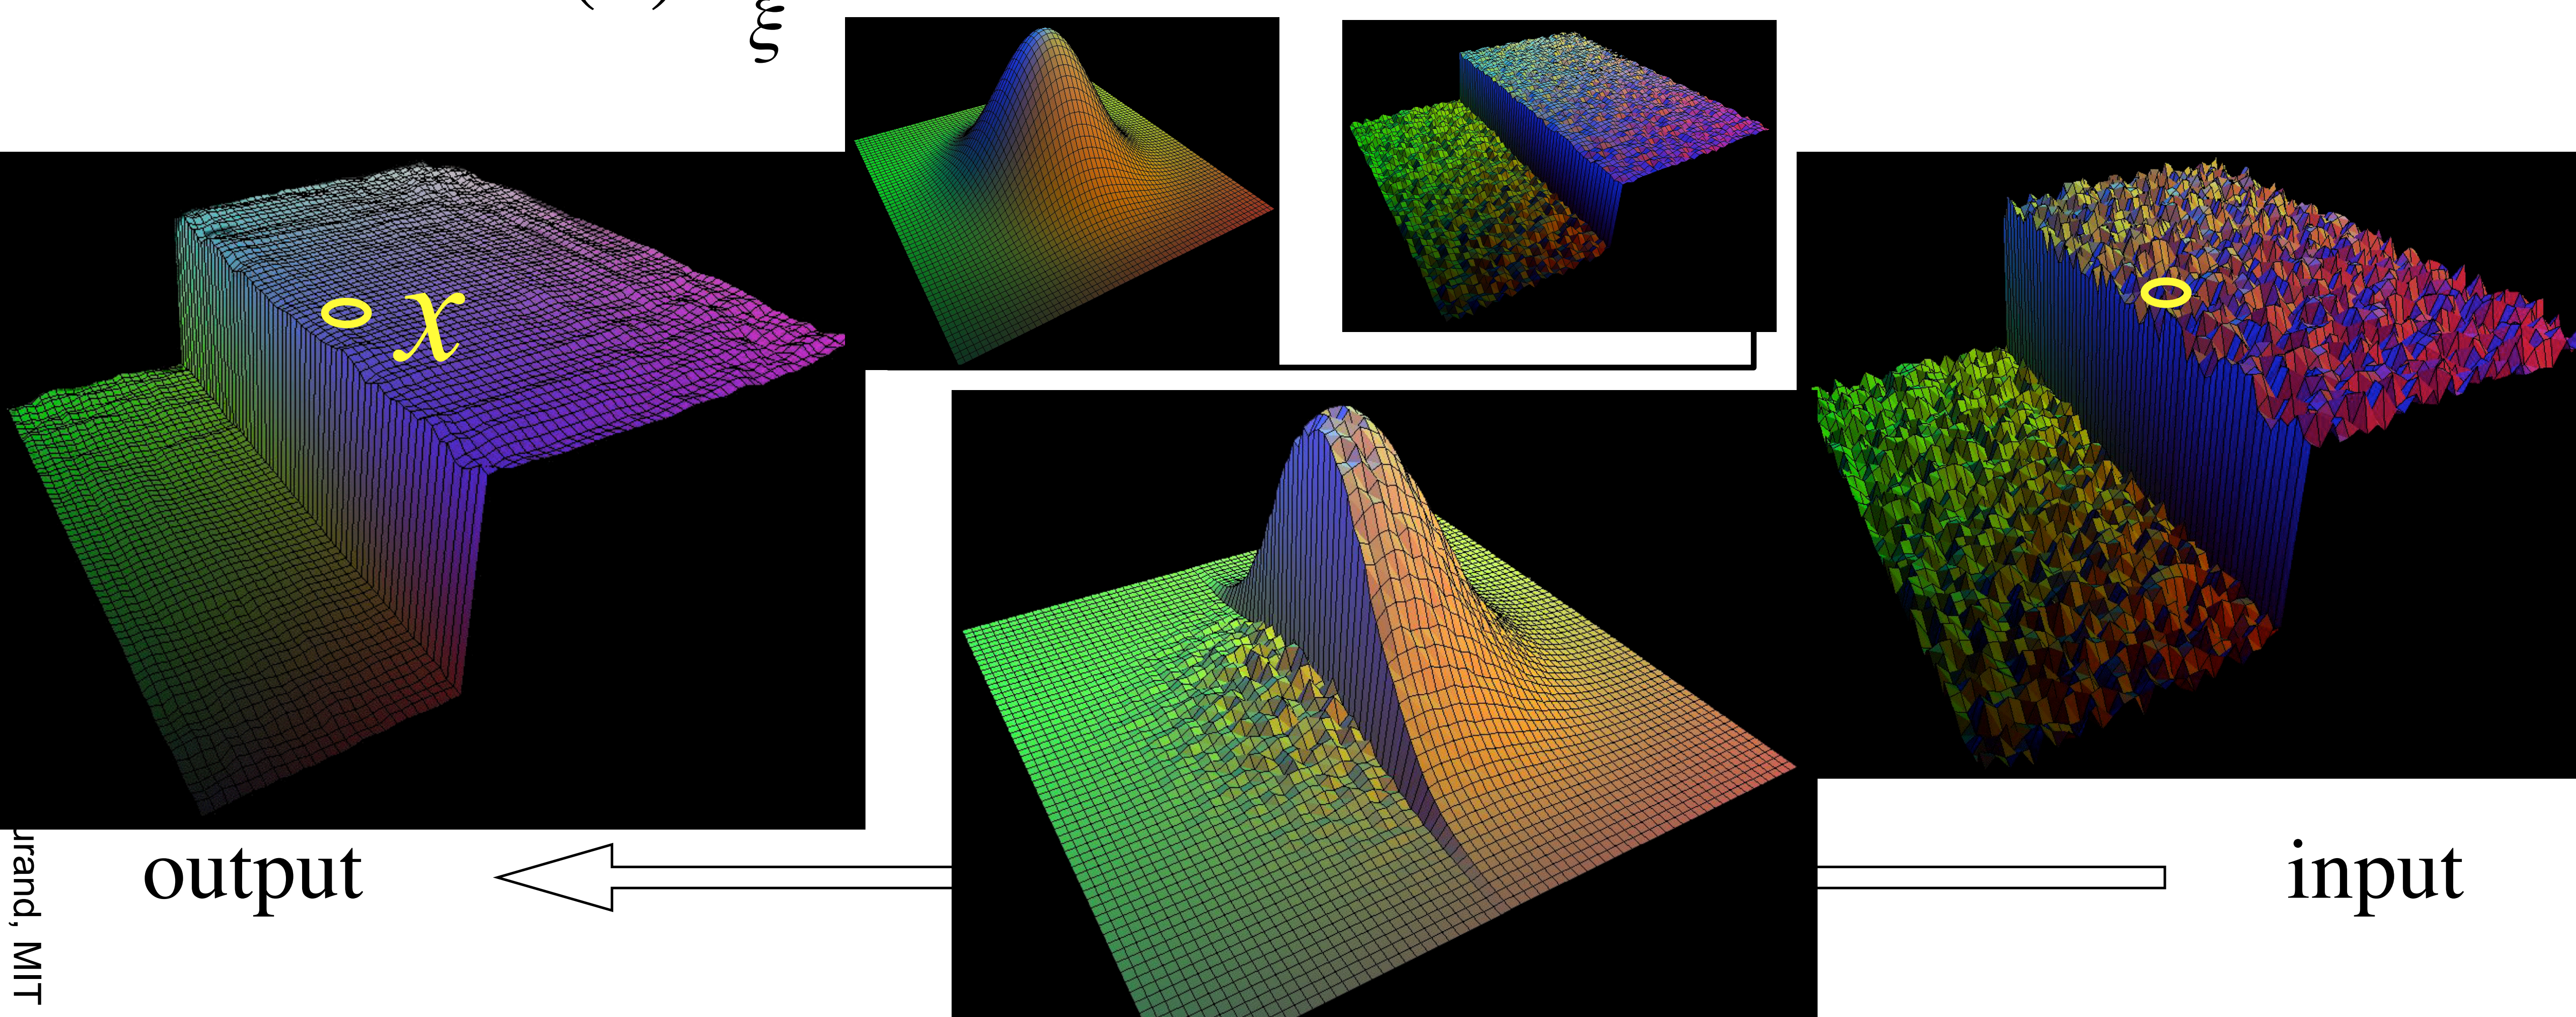

urand, MIT

Bilateral filtering is non-linear

[Tomasi and Manduchi 1998]

- The weights are different for each output pixel

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x,\xi) \quad g(I(\xi) - I(x)) \quad I(\xi)$$

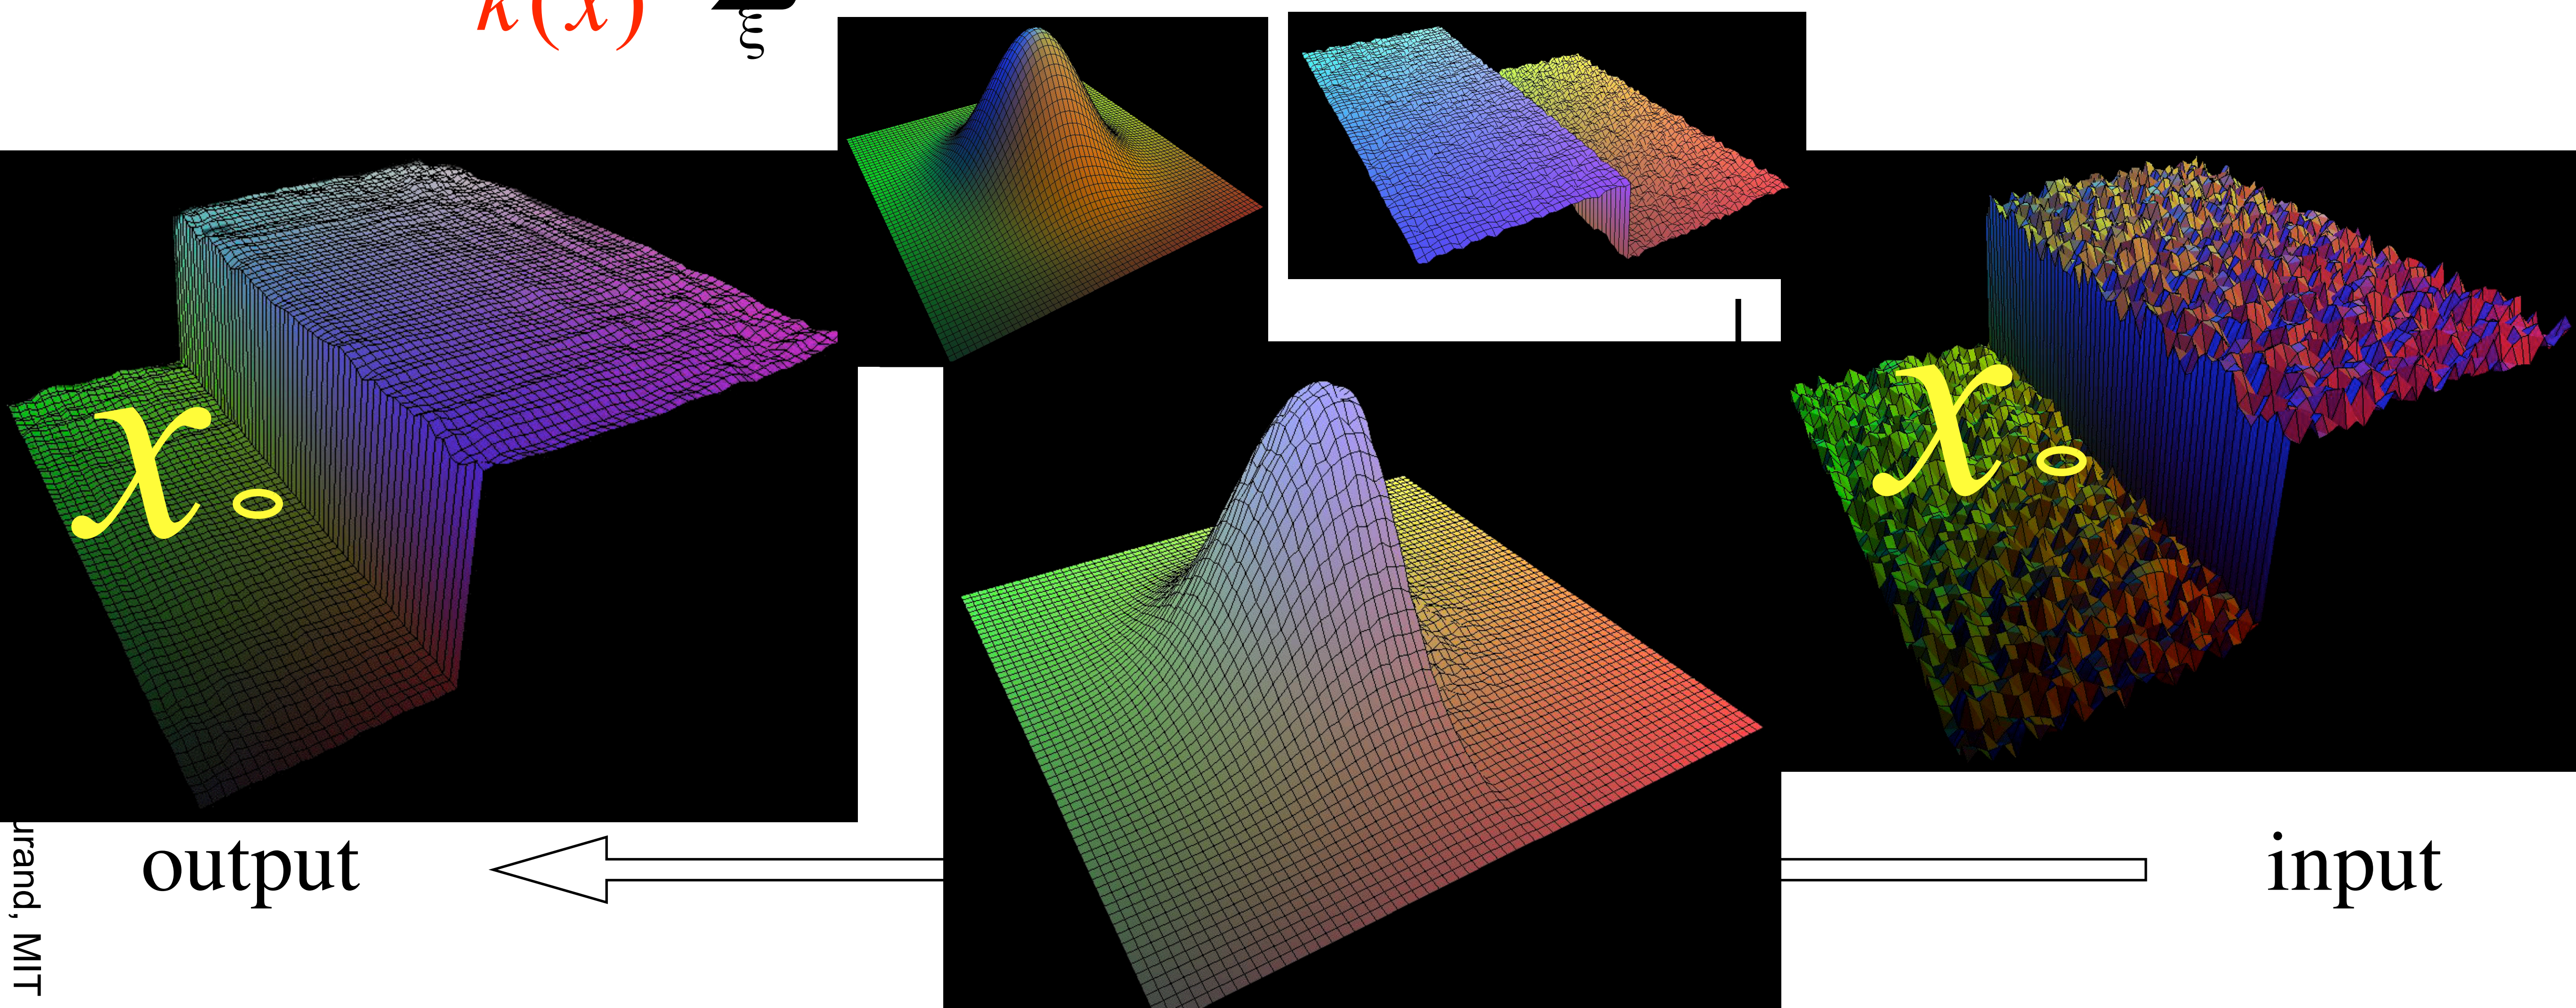

urand, MIT

Effects of bilateral filter

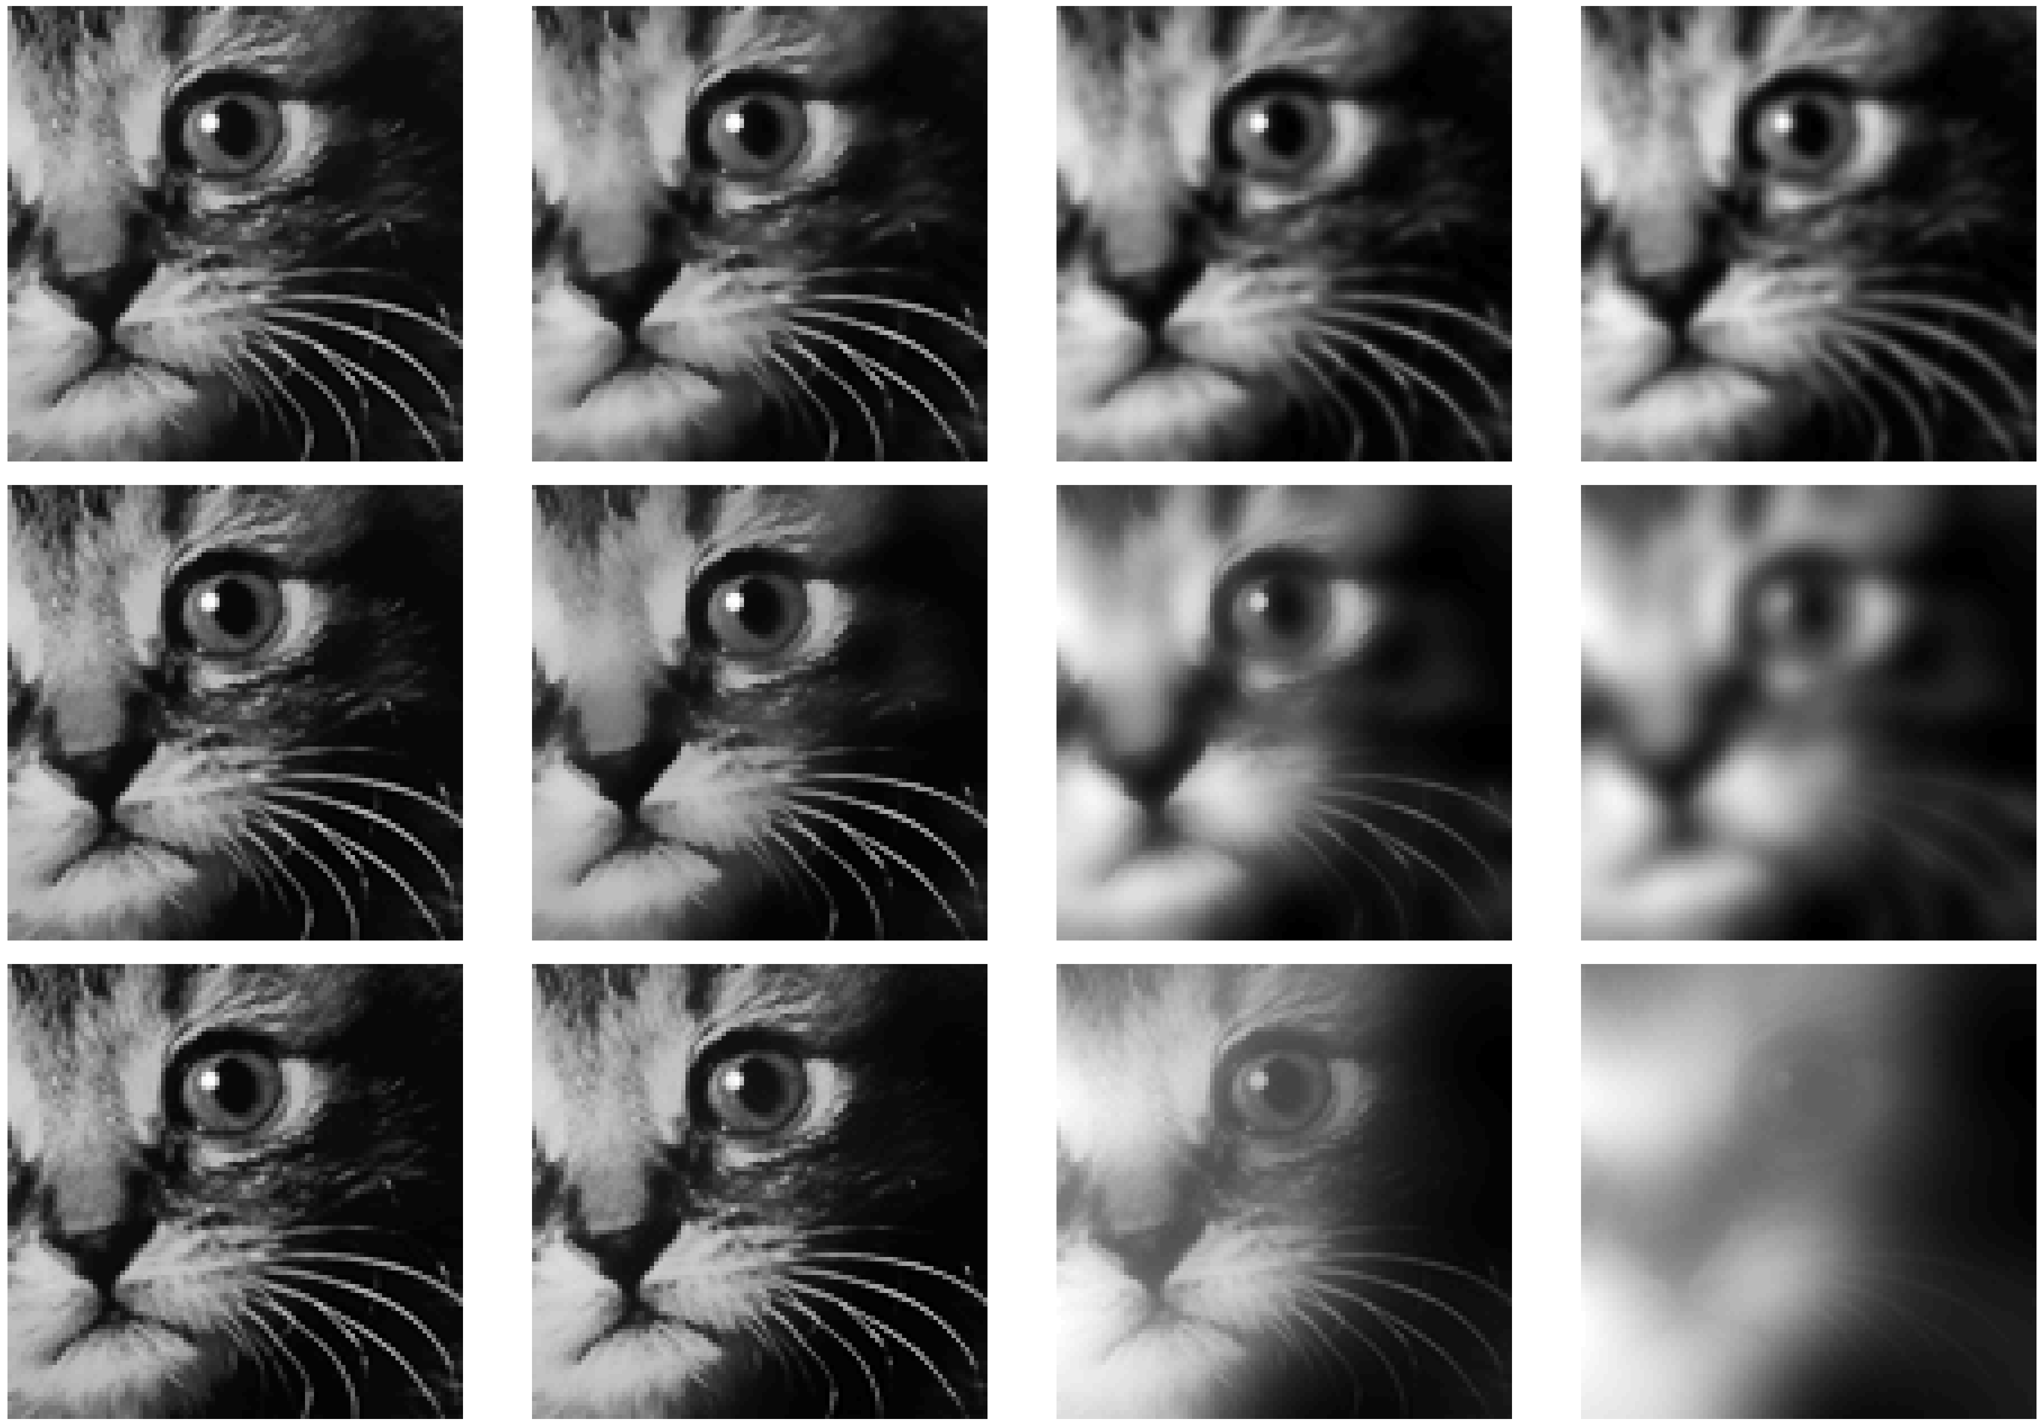

$\sigma_d = 1$

$\sigma_d = 3$

$\sigma_d = 10$

size of **domain** filter ⟶

$\sigma_r = 10$ $\sigma_r = 30$ $\sigma_r = 100$ $\sigma_r = 300$

size of **range** filter ⟶

[Tomasi & Manduchi 1998]

Bilateral filter

Noisy input

After gaussian blur

adapted from slide by Frédo Durand, MIT

Bilateral filter

Noisy input

After bilateral filter

slide by Frédo Durand, MIT

slide by Frédo Durand, MIT

Can we do better?

Noisy input

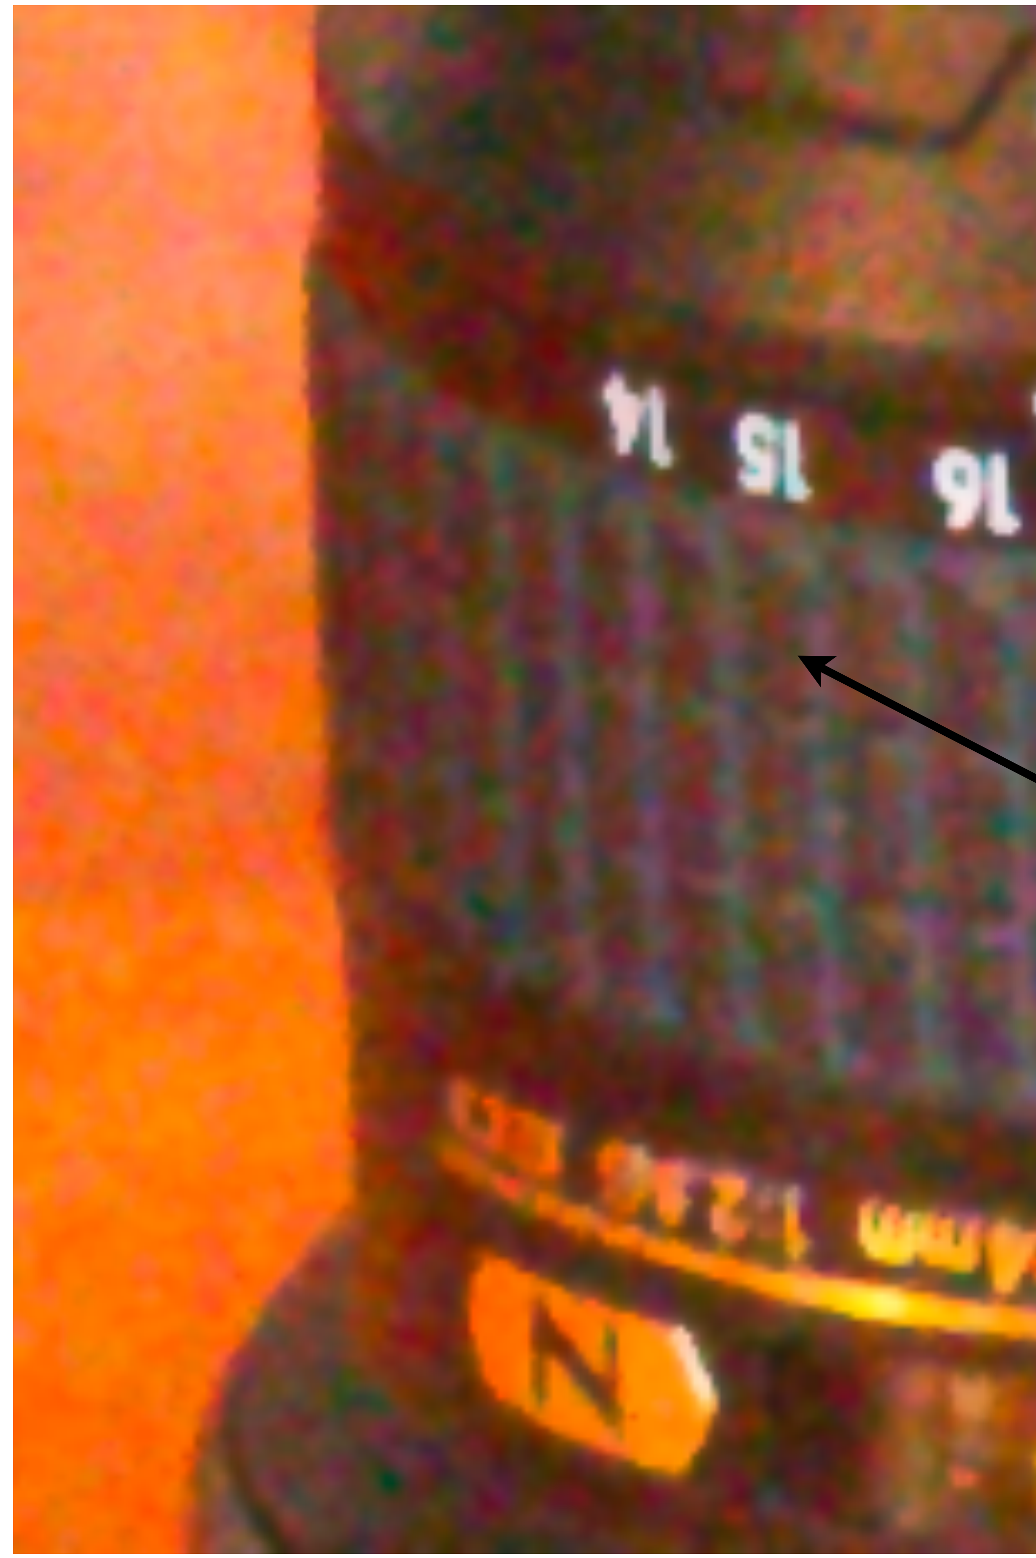

chroma noise

After bilateral filter

slide by Frédo Durand, MIT

Chroma noise

- Our visual system has different spatial frequency response to chrominance vs. luminance
- Perform Biateral filtering in YUV
- Bigger spatial filter in U & V

slide by Frédo Durand, MIT

Normal RGB Bilateral filter

Noisy input

After bilateral filter

slide by Frédo Durand, MIT

YUV bilateral filter

Noisy input

After YUV bilateral filter

slide by Frédo Durand, MIT

Comparison

Noisy input

Bilateral filter

YUV bilateral filter

Bilateral filtering

slide by Frédo Durand, MIT

- Also used to remove skin blemishes in portraits
  - Surface blur in photoshop
    - Although box spatial kernel instead of Gaussian
- Useful for lots of other things
  - In particular, tone mapping for contrast reduction and high-dynamic-range imaging

slide by Frédo Durand, MIT

Photoshop surface blur

- Note the radius and threshold controls
  - same as sigma_domain and sigma_range

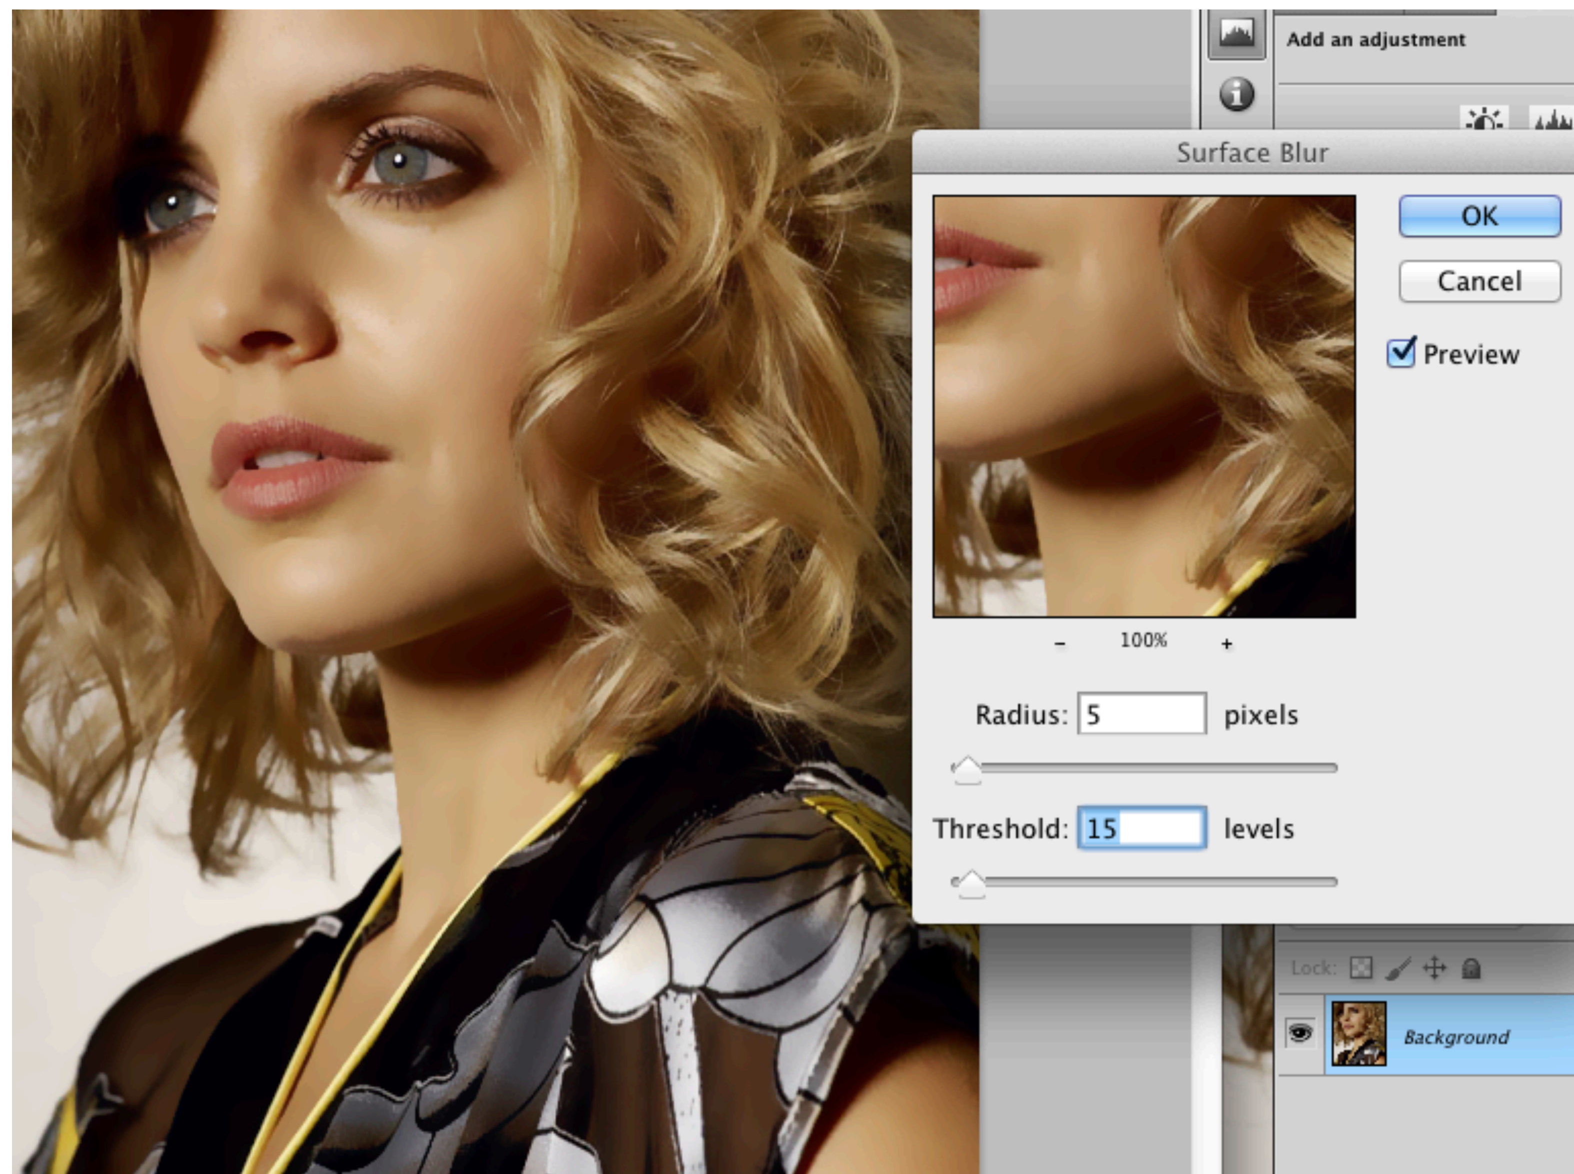

HDR tone mapping

- **Problem: measured dynamic range does not fit into displayable dynamic range**
- **Goal: compress range of image without producing low contrast**
- **Approach: divide into detail and large-scale, compress only large-scale**

  question: how to differentiate “detail” from “large scale”

  answer has gotten more nuanced over time

Dynamic range

- In the highlights, we are limited by clipping
- In the shadows, we are limited by noise

Examples

- Inside is too dark
- Outside is too bright

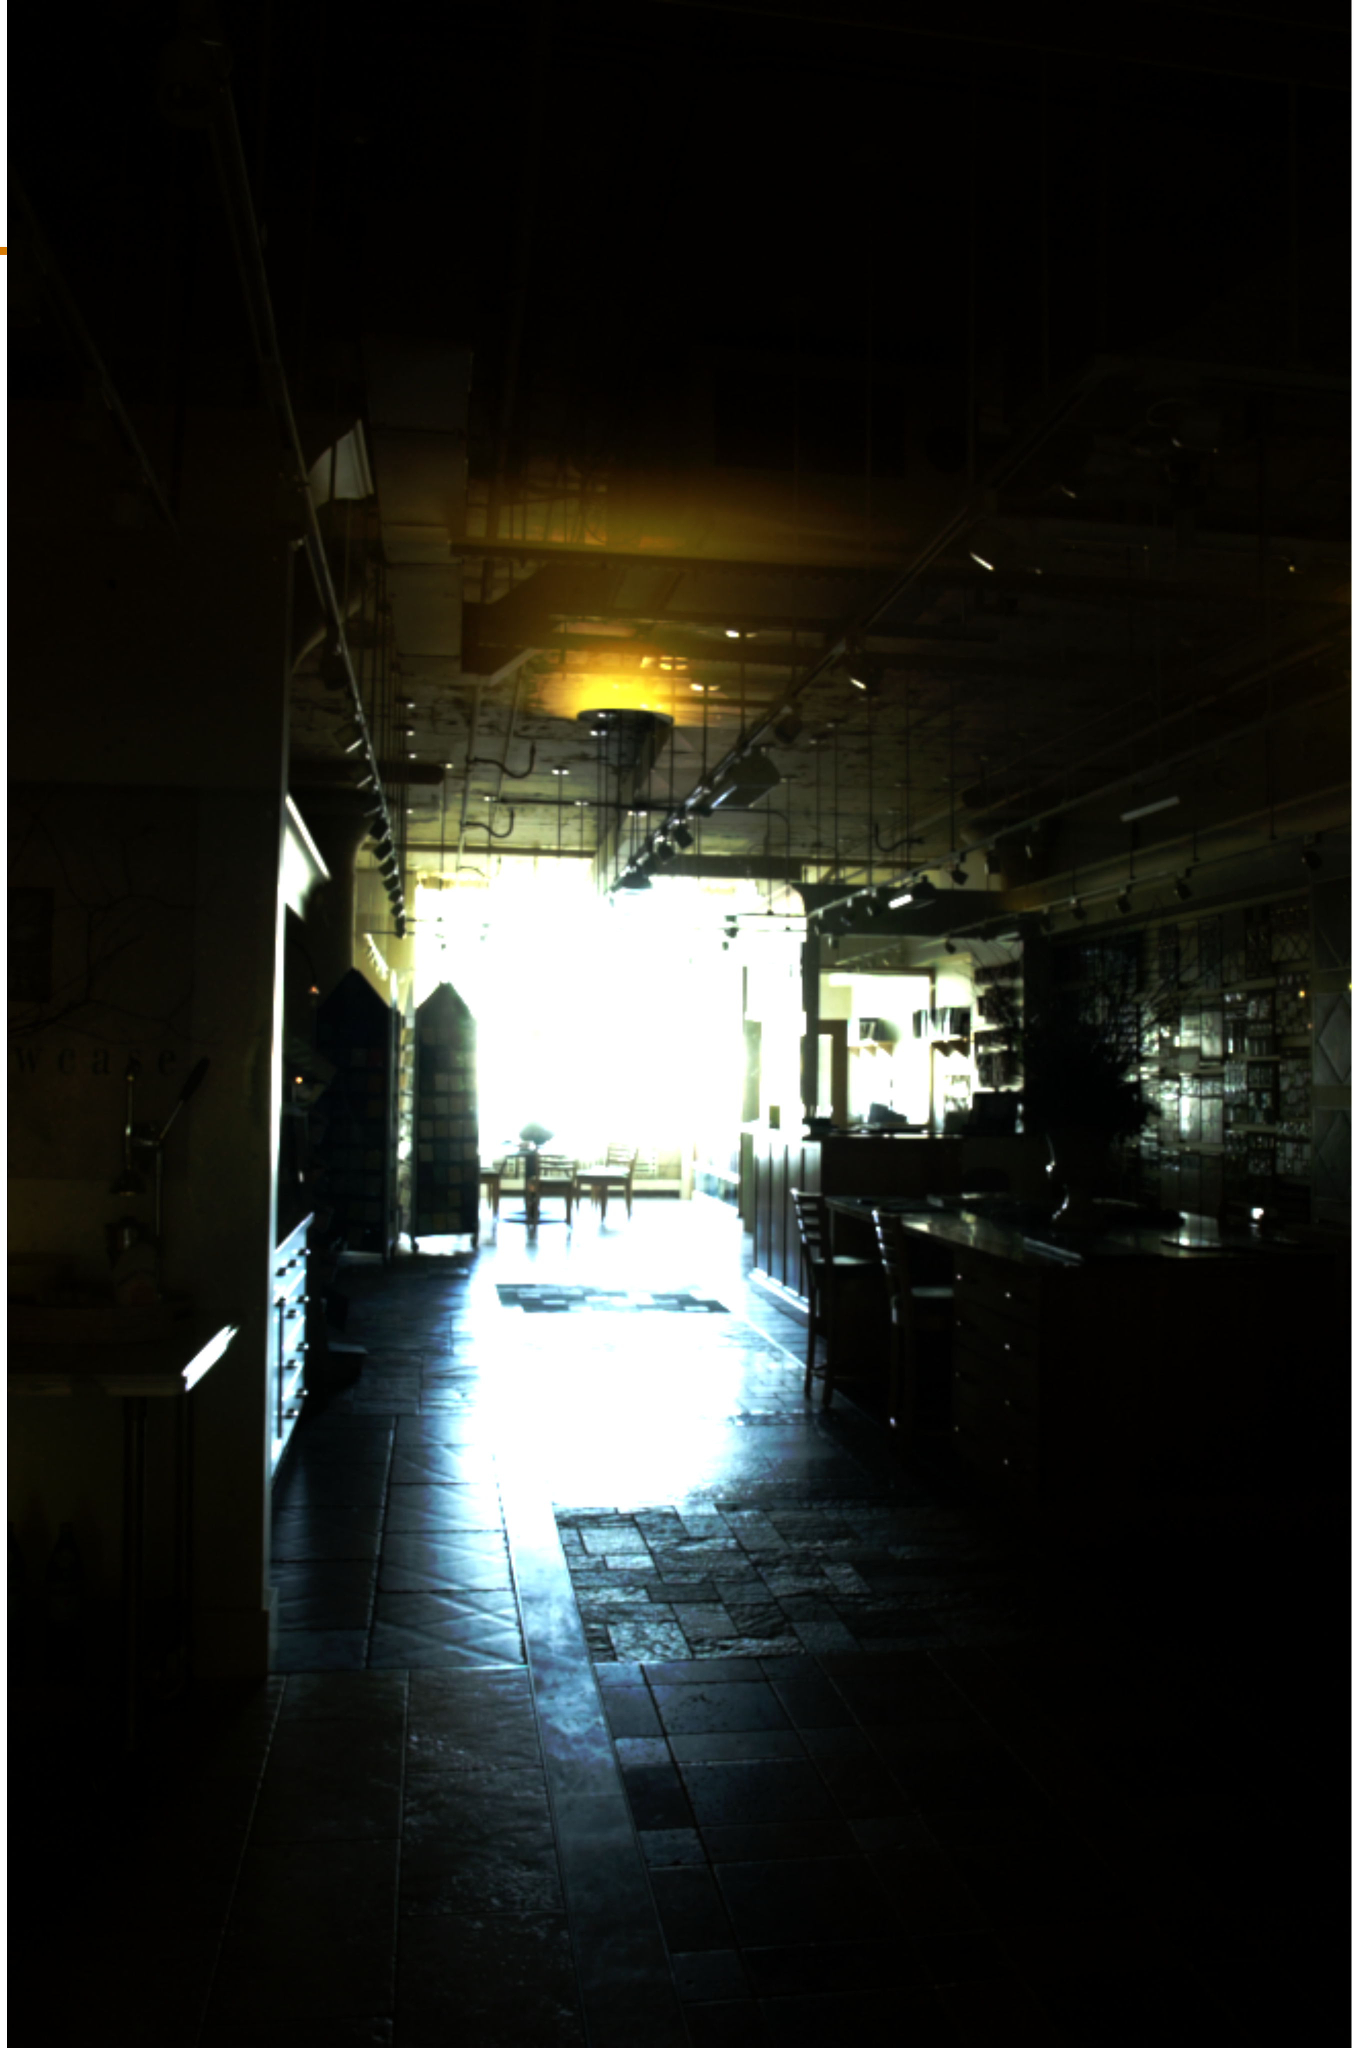

- Sun overexposed
- Foreground too dark

Dean S. Pemberton

Masato OHTA


Alexandre Buisse

Alexandre Buisse

Pointwise transforms

- **Tumblin & Rushmeier 1993**
- **Ward et al. 1997**

Ward et al. 1997

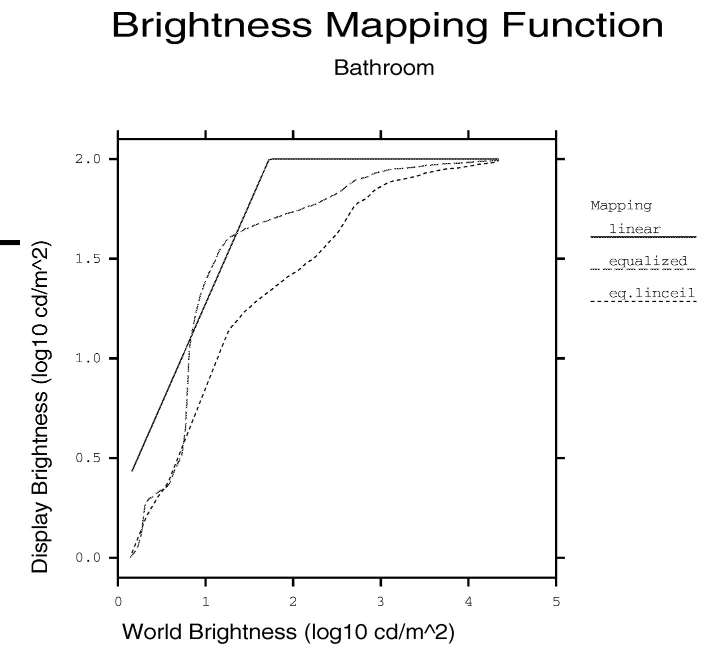

- **Histogram equalization**
- **Linear ceiling (prevent contrast expansion)**

linear

histogram equalized

with linear ceiling constraint

Pointwise transforms

- **Tumblin & Rushmeier 1993**
- **Ward et al. 1997**
- **problem: limited compression achievable**

Need for spatially varying operators

- **Chiu et al. 1993**
- **Pattanaik et al. 1998**
- **Reinhard 2002**
- **Ashikhmin 2002**
- **Scale defined by linear filtering**

Chiu et al. 1993

- **Locally varying contrast enhancement**
- **Prone to haloes**

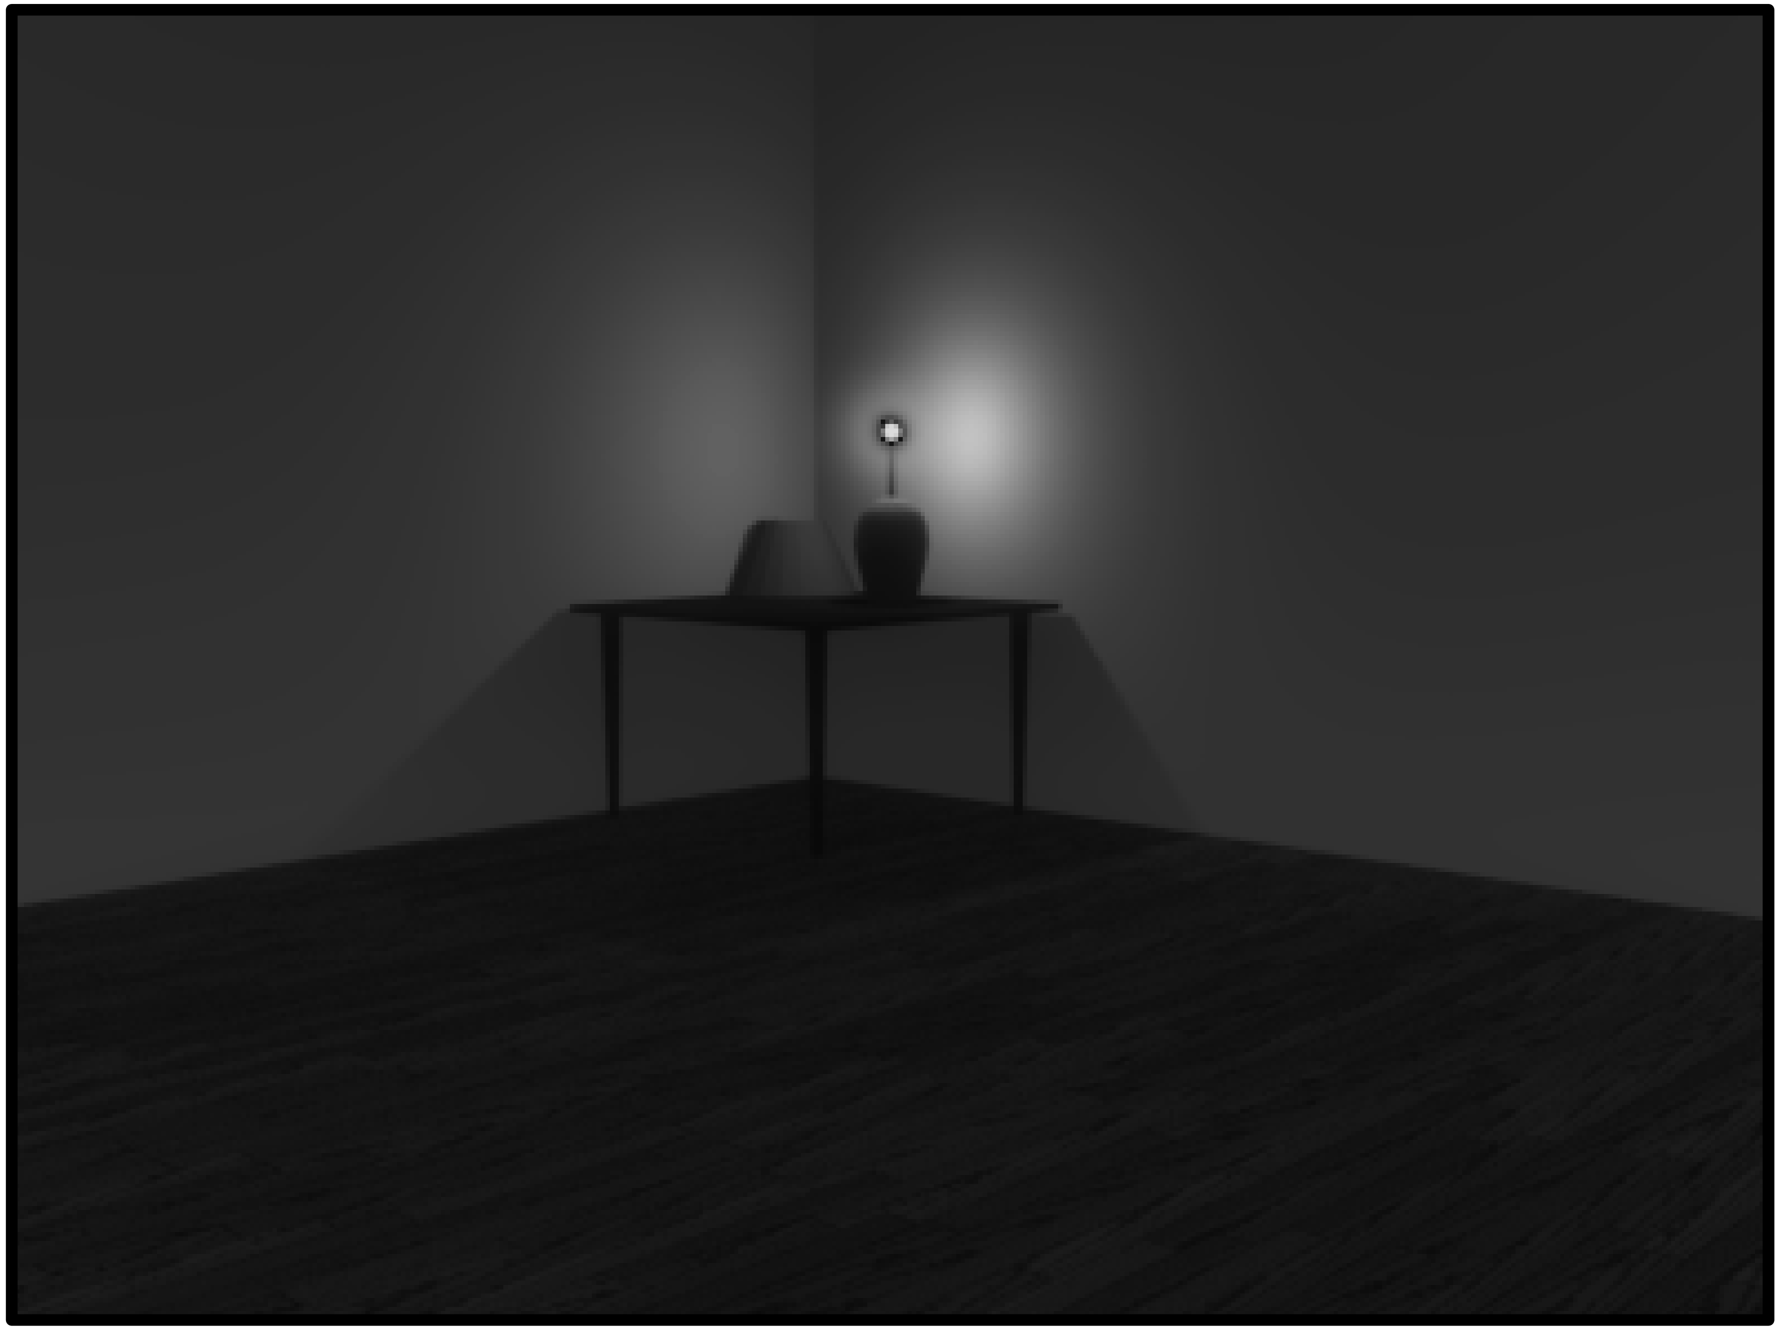

Pattanaik et al. 1998

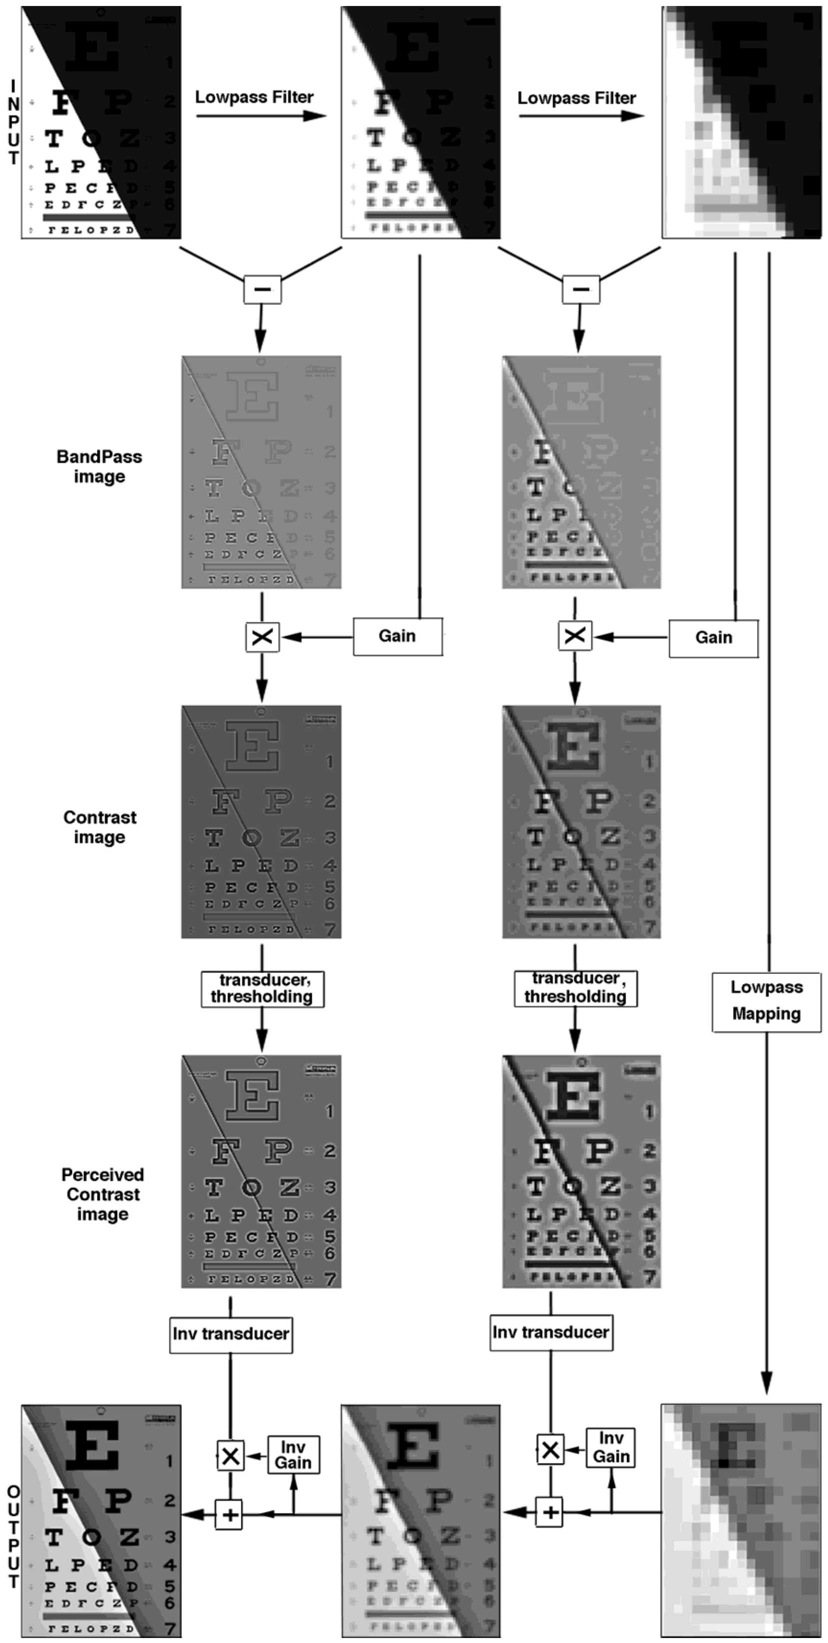

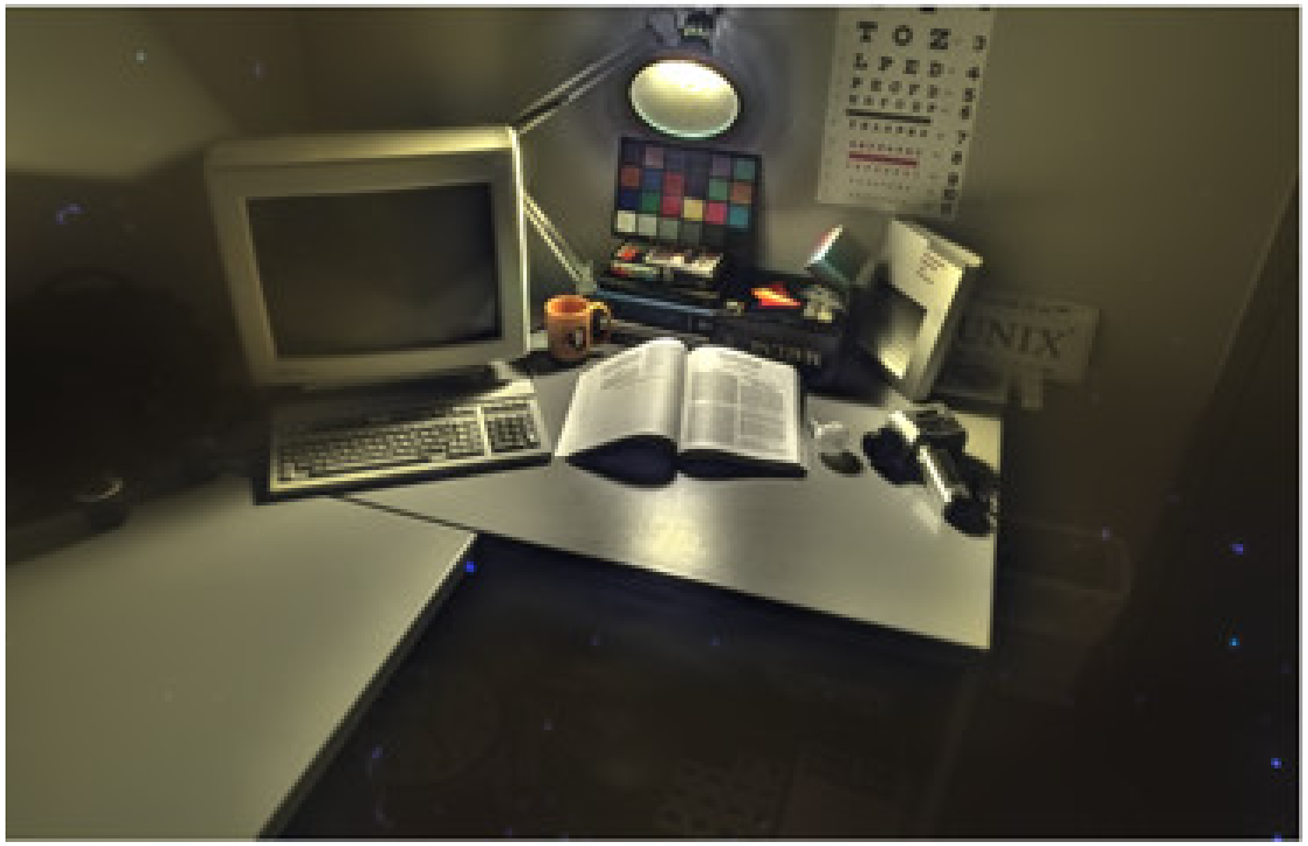

Reinhard et al. 2002

- **Careful tuning of filter scales to avoid haloes**

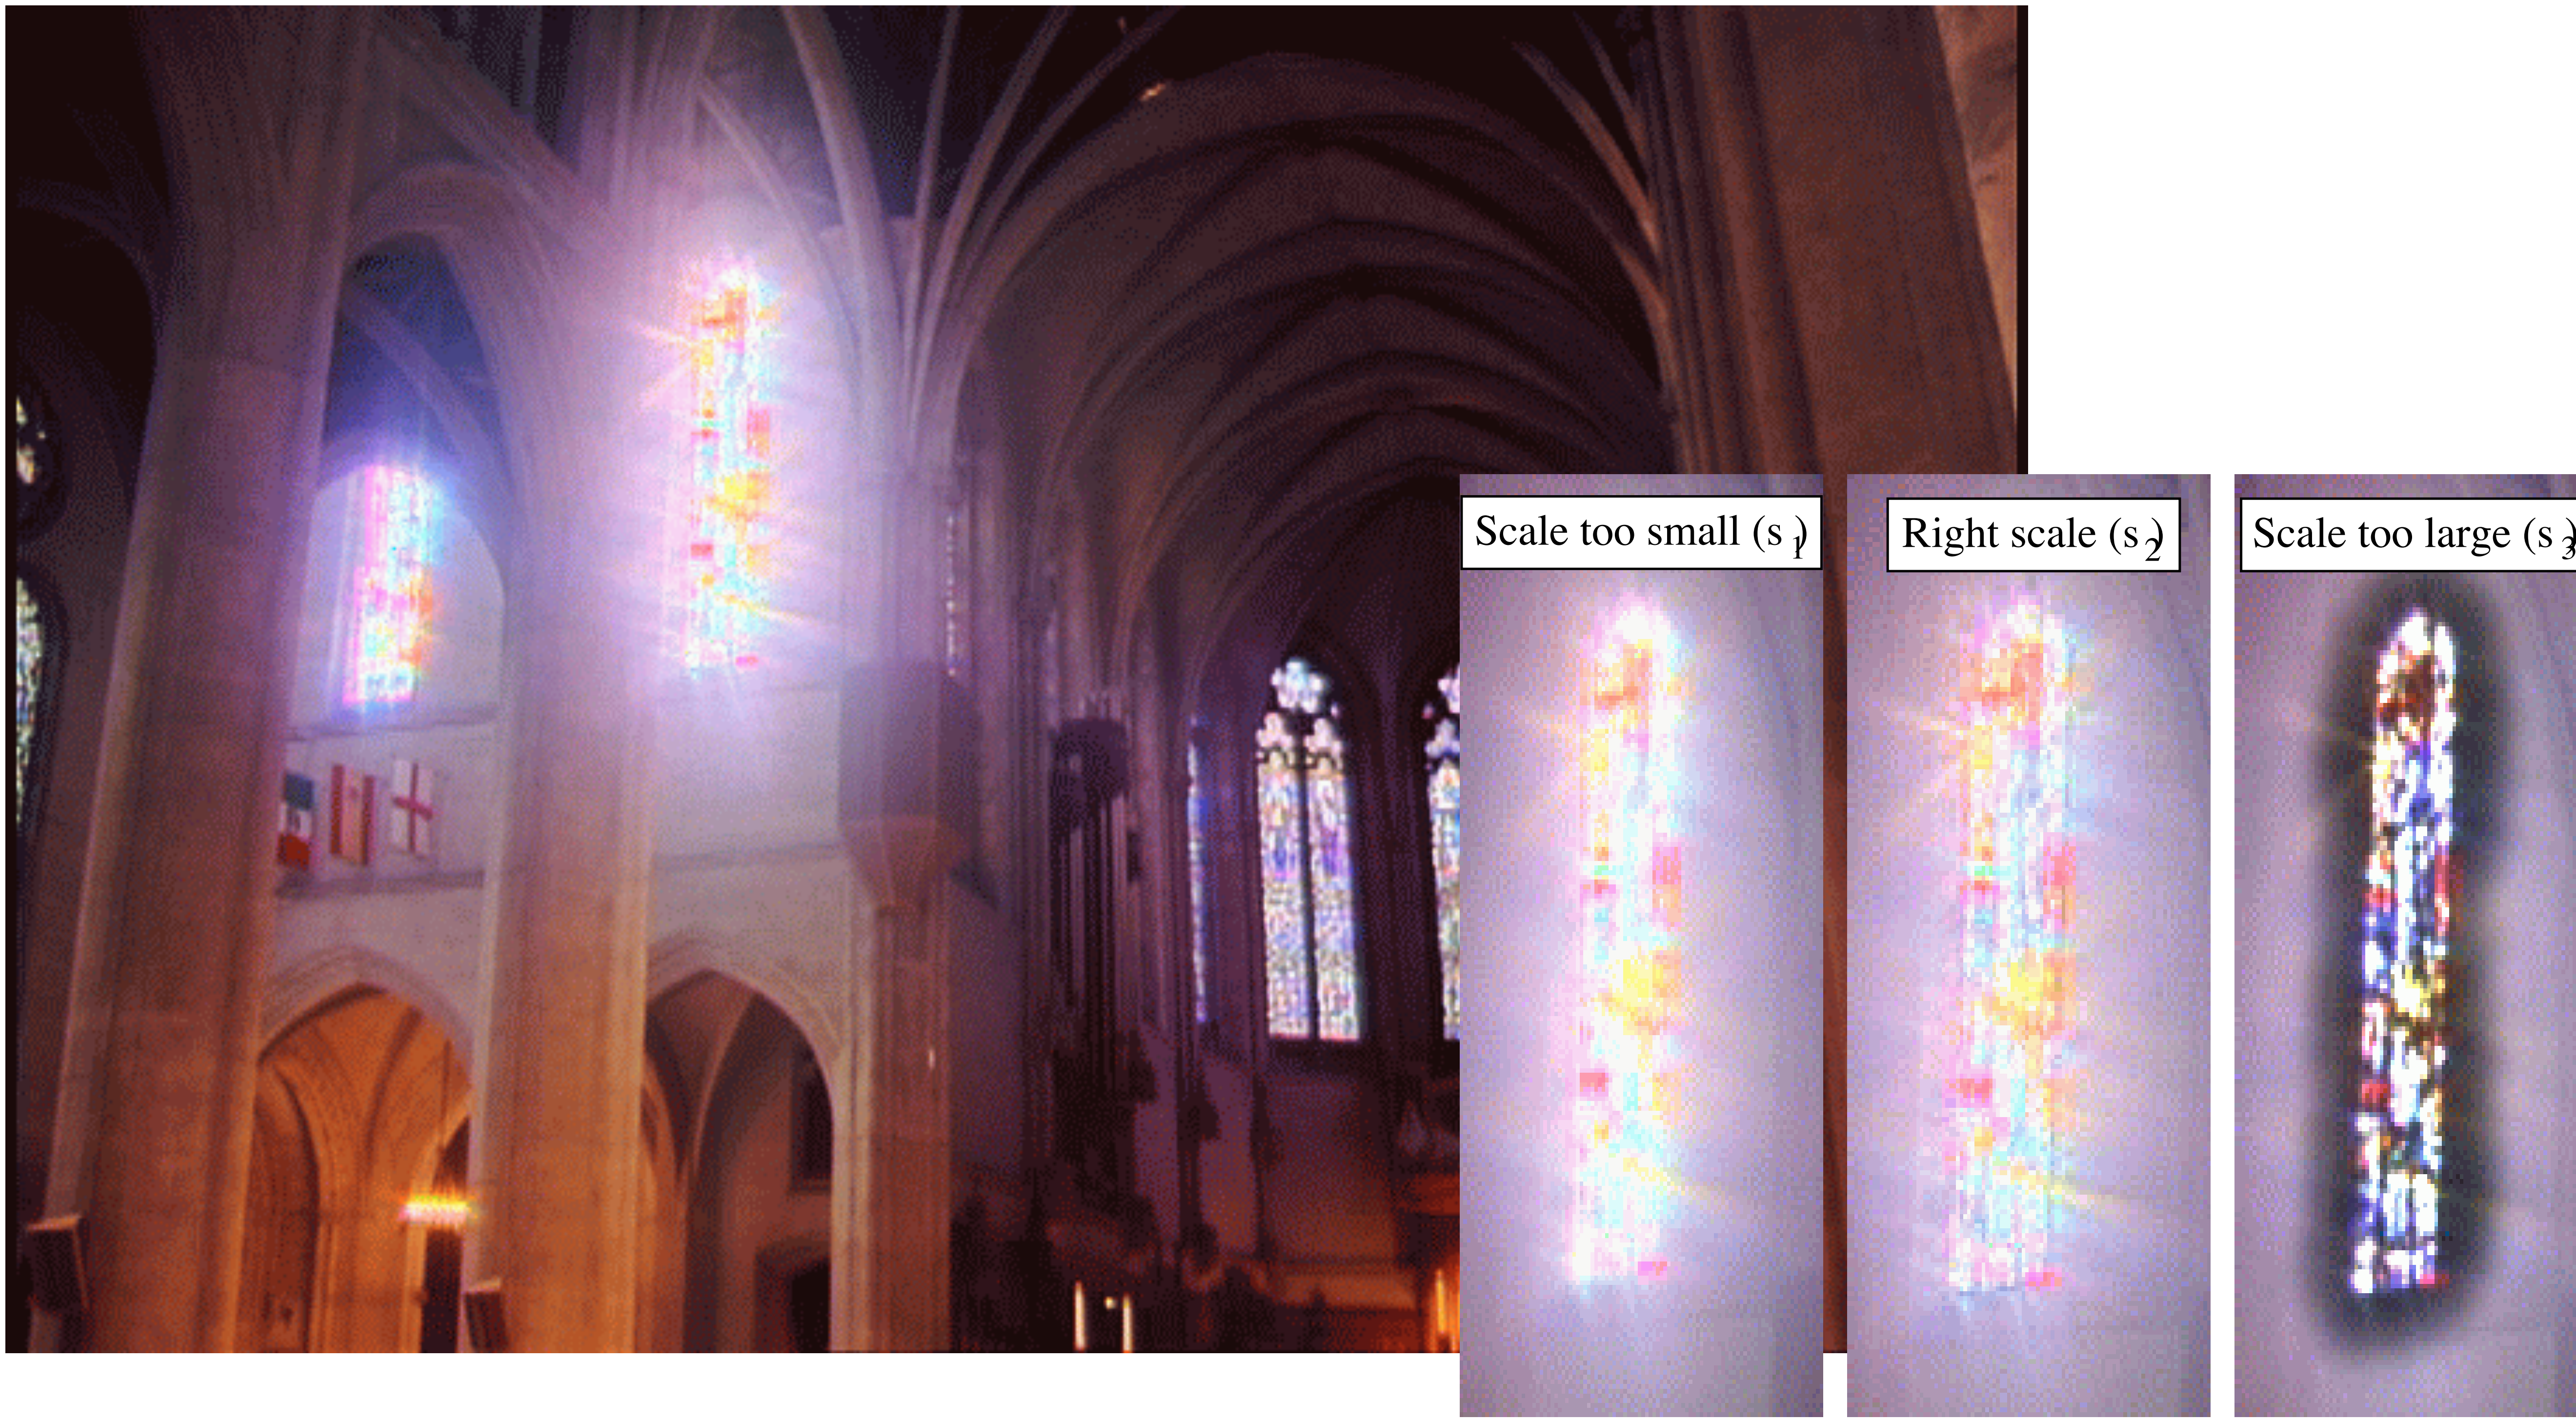

Ashikhmin 2002

- **Similarly careful adjustment of filter radius**
- **Results in edge-preserving lowpass filter**

adaptation image

detail image

Need for spatially varying operators

- **Chiu et al. 1993**
- **Reinhard 2002**
- **Ashikhmin 2002**
- **Scale defined by linear filtering**
- **Problem: haloes are tricky to avoid**

Novel notions of scale

- **Tumblin & Turk 1997 (LCIS)**
- **Durand & Dorsey 2002 (bilateral filter)**
- **Fattal et al. 2002 (gradient domain)**

Tumblin & Turk 1997

- **LCIS “Low Curvature Image Simplifier.”**
- **Uses diffusion to compute scale space decomposition**

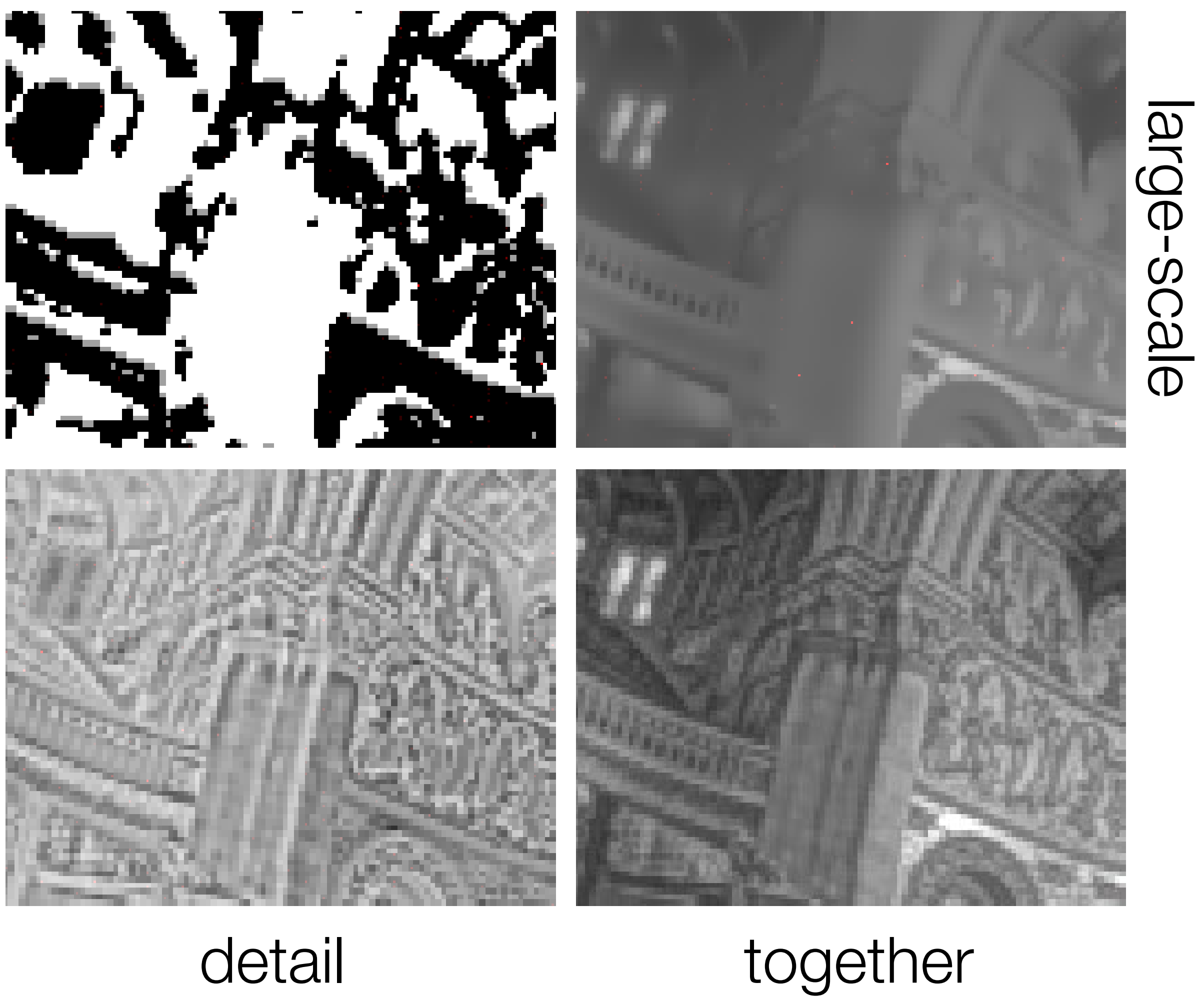

Tumblin & Turk 1997

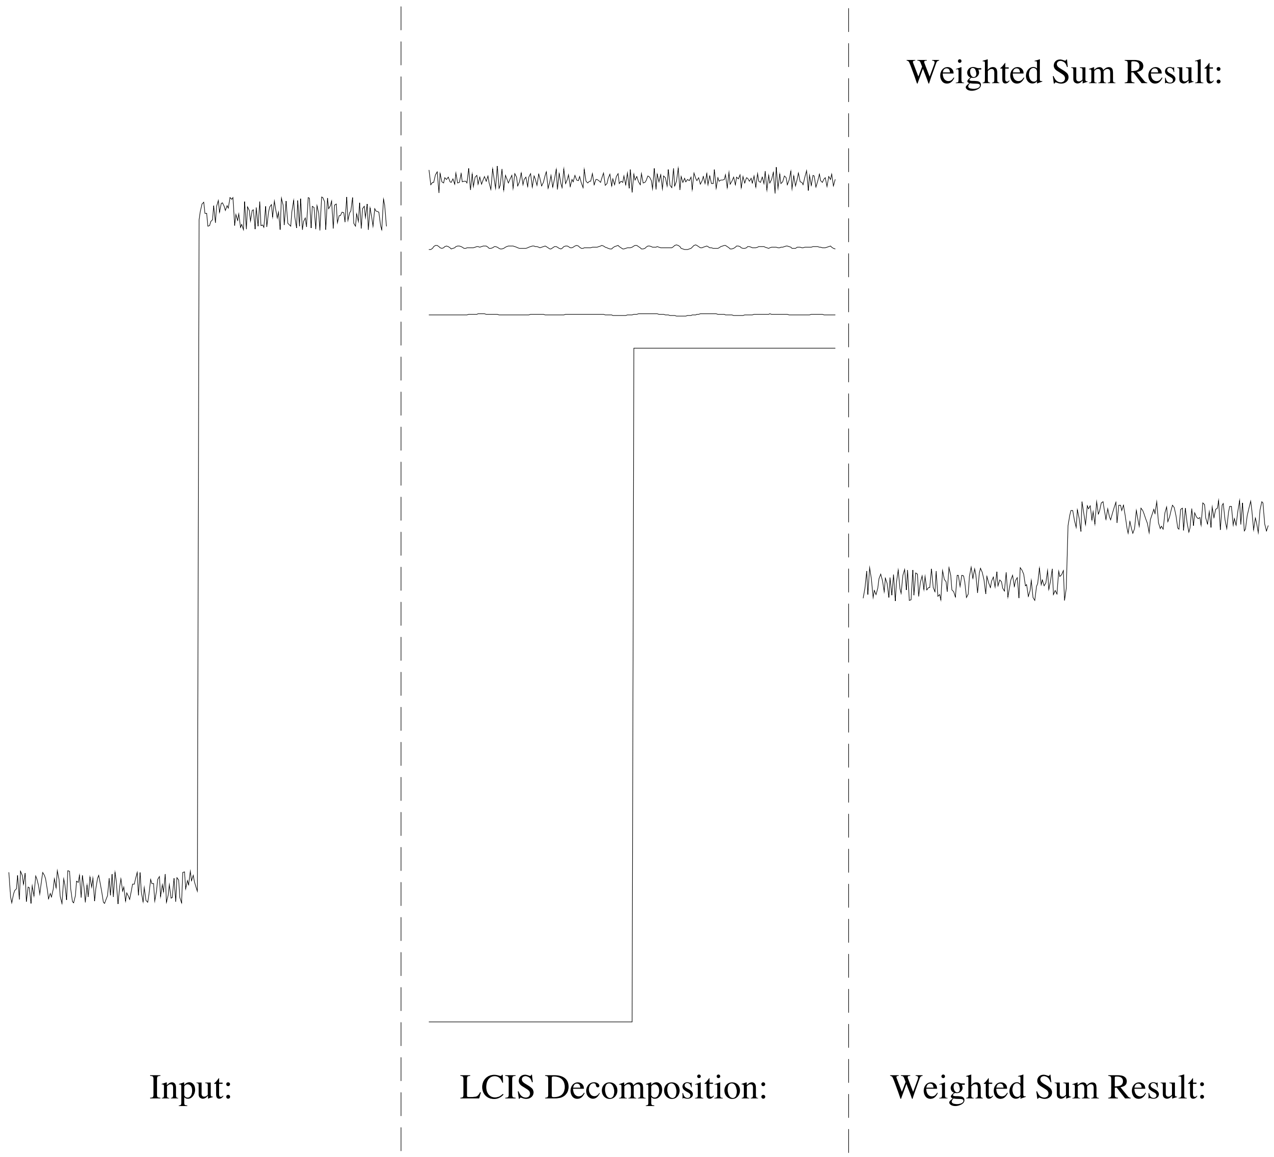

Filtering using LCIS

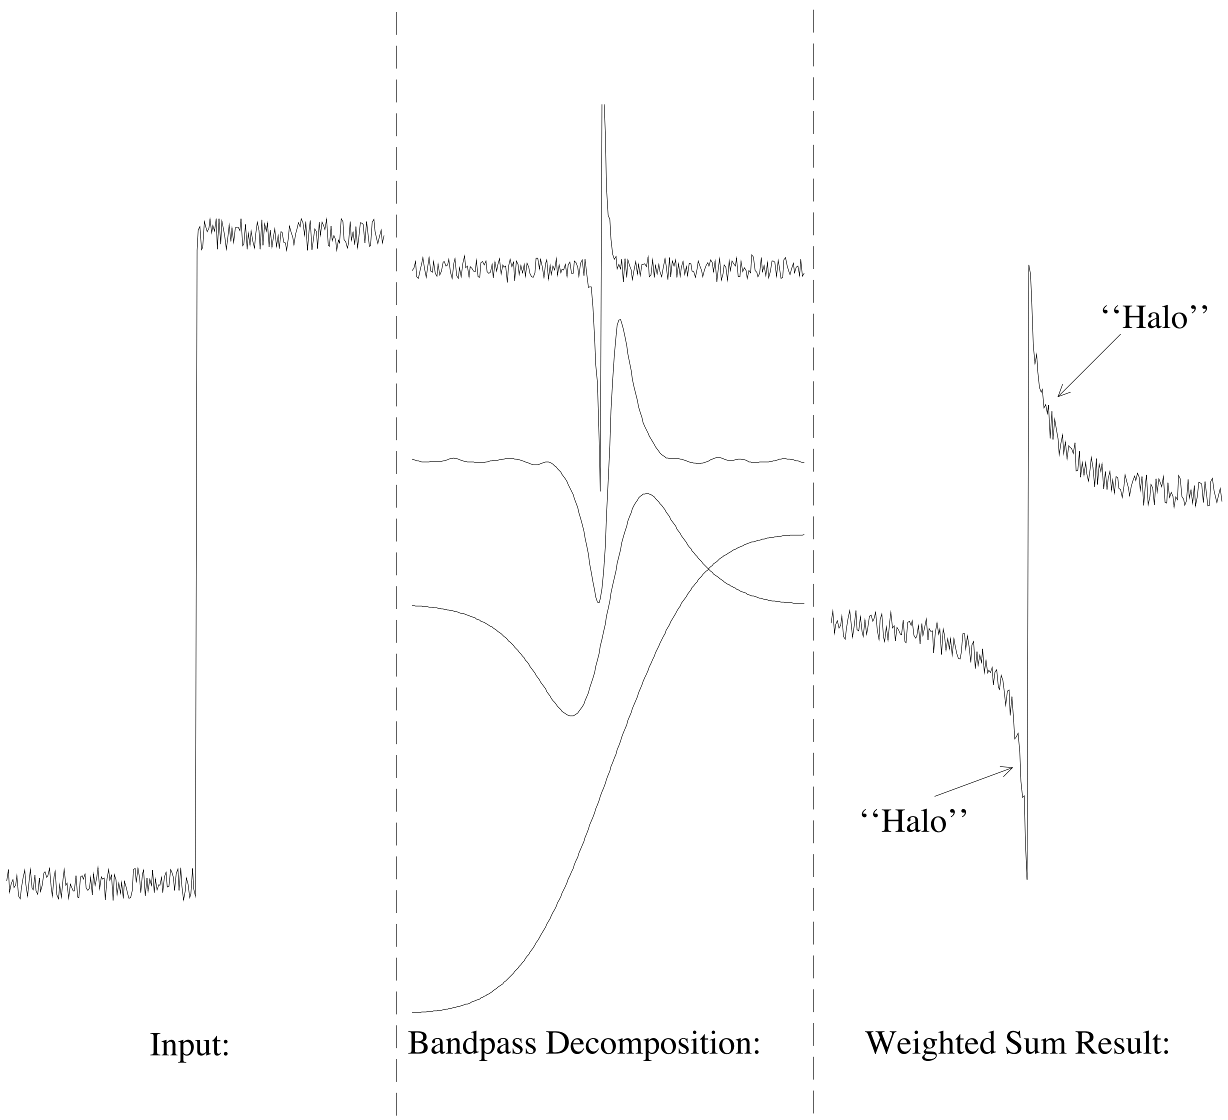

Filtering using usual bandpass filters

Durand & Dorsey 2002

- **Separate detail from large-scale using bilateral filter**

Fattal et al. 2002

- **Separate detail from large-scale using gradient in log space**

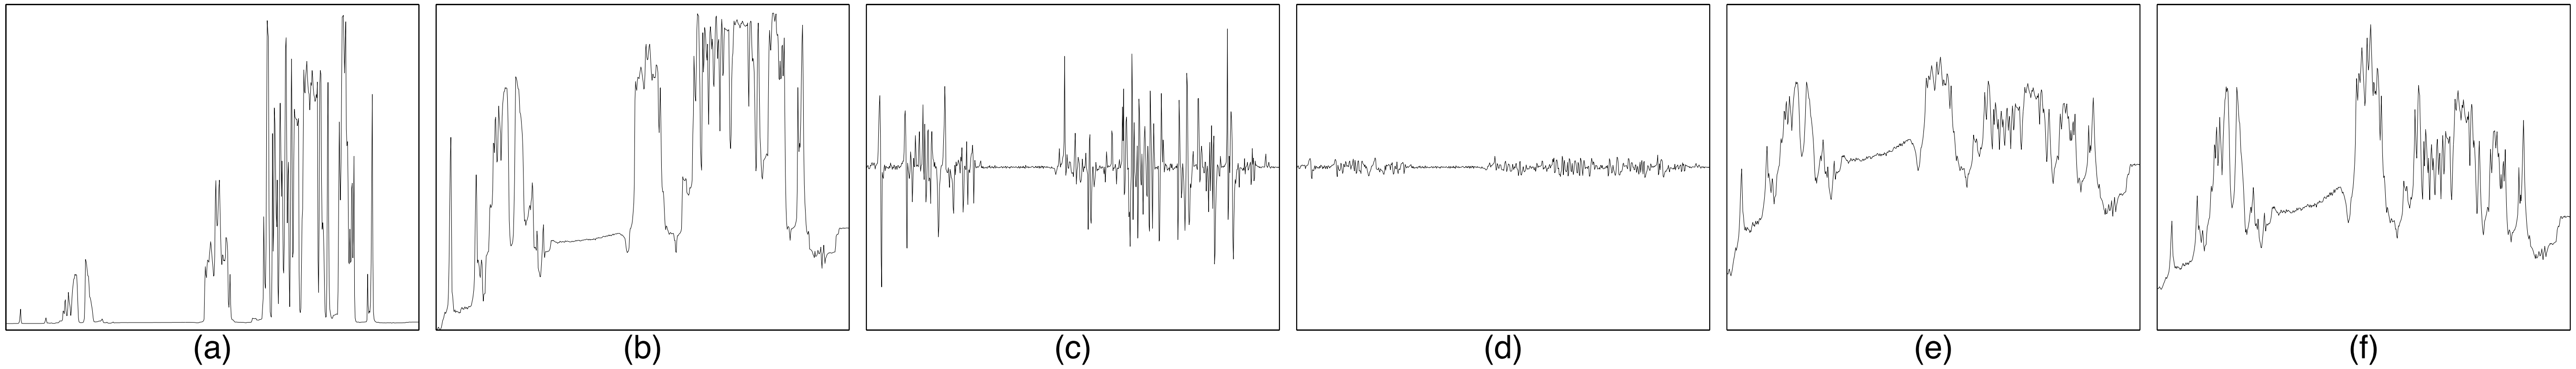

Figure 3: (a) An HDR scanline with dynamic range of 2415:1. (b) $H(x) = \log(\text{scanline})$. (c) The derivatives $H'(x)$. (d) Attenuated derivatives $G(x)$; (e) Reconstructed signal $I(x)$ (as defined in eq. 1); (f) An LDR scanline $\exp(I(x))$: the new dynamic range is 7.5:1. Note that each plot uses a different scale for its vertical axis in order to show details, except (c) and (d) that use the same vertical axis scaling in order to show the amount of attenuation applied on the derivatives.

Figure 4: Gradient attenuation factors used to compress the Belgium House HDR radiance map (Figure 2). Darker shades indicate smaller scale factors (stronger attenuation).

Ward | Fattal | LCIS

Bibliography

P. Perona & J. Malik, **Scale-Space and Edge Detection Using Anisotropic Diffusion**, *PAMI* 12:7 (1990).

C. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, & K. Zimmerman, **Spatially Nonuniform Scaling Functions for High Contrast Images**, *Graphics Interface 1993*.

J. Tumblin & H. Rushmeier, **Tone Reproduction for Realistic Images**, *IEEE CG&A* 13:6 (1993).

G.W. Larson, H. Rushmeier, & C. Piatko. **A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes**, *IEEE TVCG* 3:4 (1997).

S. Pattanaik, J. Ferwerda, M. Fairchild, & D. Greenberg, **A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display**, *SIGGRAPH 1998*.

C. Tomasi & R. Manduchi, **Bilateral Filtering for Gray and Color Images**, *ICCV 1998*.

J. Tumblin & G. Turk, **LCIS: A Boundary Hierarchy For Detail-Preserving Contrast Reduction**, *SIGGRAPH 1999*.

F. Durand & J. Dorsey, **Fast Bilateral Filtering for the Display of High-Dynamic-Range Images**, *SIGGRAPH 2002*.

R. Fattal, D. Lischinski, & M. Werman, **Gradient Domain High Dynamic Range Compression**, *SIGGRAPH 2002*.

E. Reinhard, M. Stark, P. Shirley, & J. Ferwerda, **Photographic Tone Reproduction for Digital Images**, *SIGGRAPH 2002*.

M. Ashikhmin, **A Tone Mapping Algorithm for High Contrast Images**, *Eurographics Workshop on Rendering 2002*.