

CS5630 Physically Based Realistic Rendering

17. HDR tone mapping and Bilateral filtering

Plan

- **HDR imaging**

how to deal with images with a wide range of brightness

- **Tone mapping**

a fundamental problem of both photography and rendering

- **Bilateral filtering**

another way to think about large and small scale detail

- **Spatially-varying tone mapping**

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

HDR tone mapping

- **Problem: measured dynamic range does not fit into displayable dynamic range**
- **Goal: compress range of image without producing low contrast**
- **One approach: apply a mapping to pixel values that “squeezes” a wide range of luminances into the [0,1] range**

question: how to balance contrast in darks, mids, brights

- **Another 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




Brightened
many times



Examples

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

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





Dean S. Pemberton







Baseline tone mapping method

- **The simplest plausible thing is to scale and clamp**
 1. define a reference white luminance
 2. map colors to RGB with ref_white going to [1,1,1]
 3. clamp values with $\min(1, \dots)$
 4. apply color space nonlinearity
- **Problem: lose all detail above reference white!**

Pointwise transforms

- **Tumblin & Rushmeier 1993**
- **Ward et al. 1997**
- **ACES “filmic” tone curve**

Tumblin & Rushmeier

- **1993 paper first pointed out the problem**
- **Proposed a way to reason about solving it**

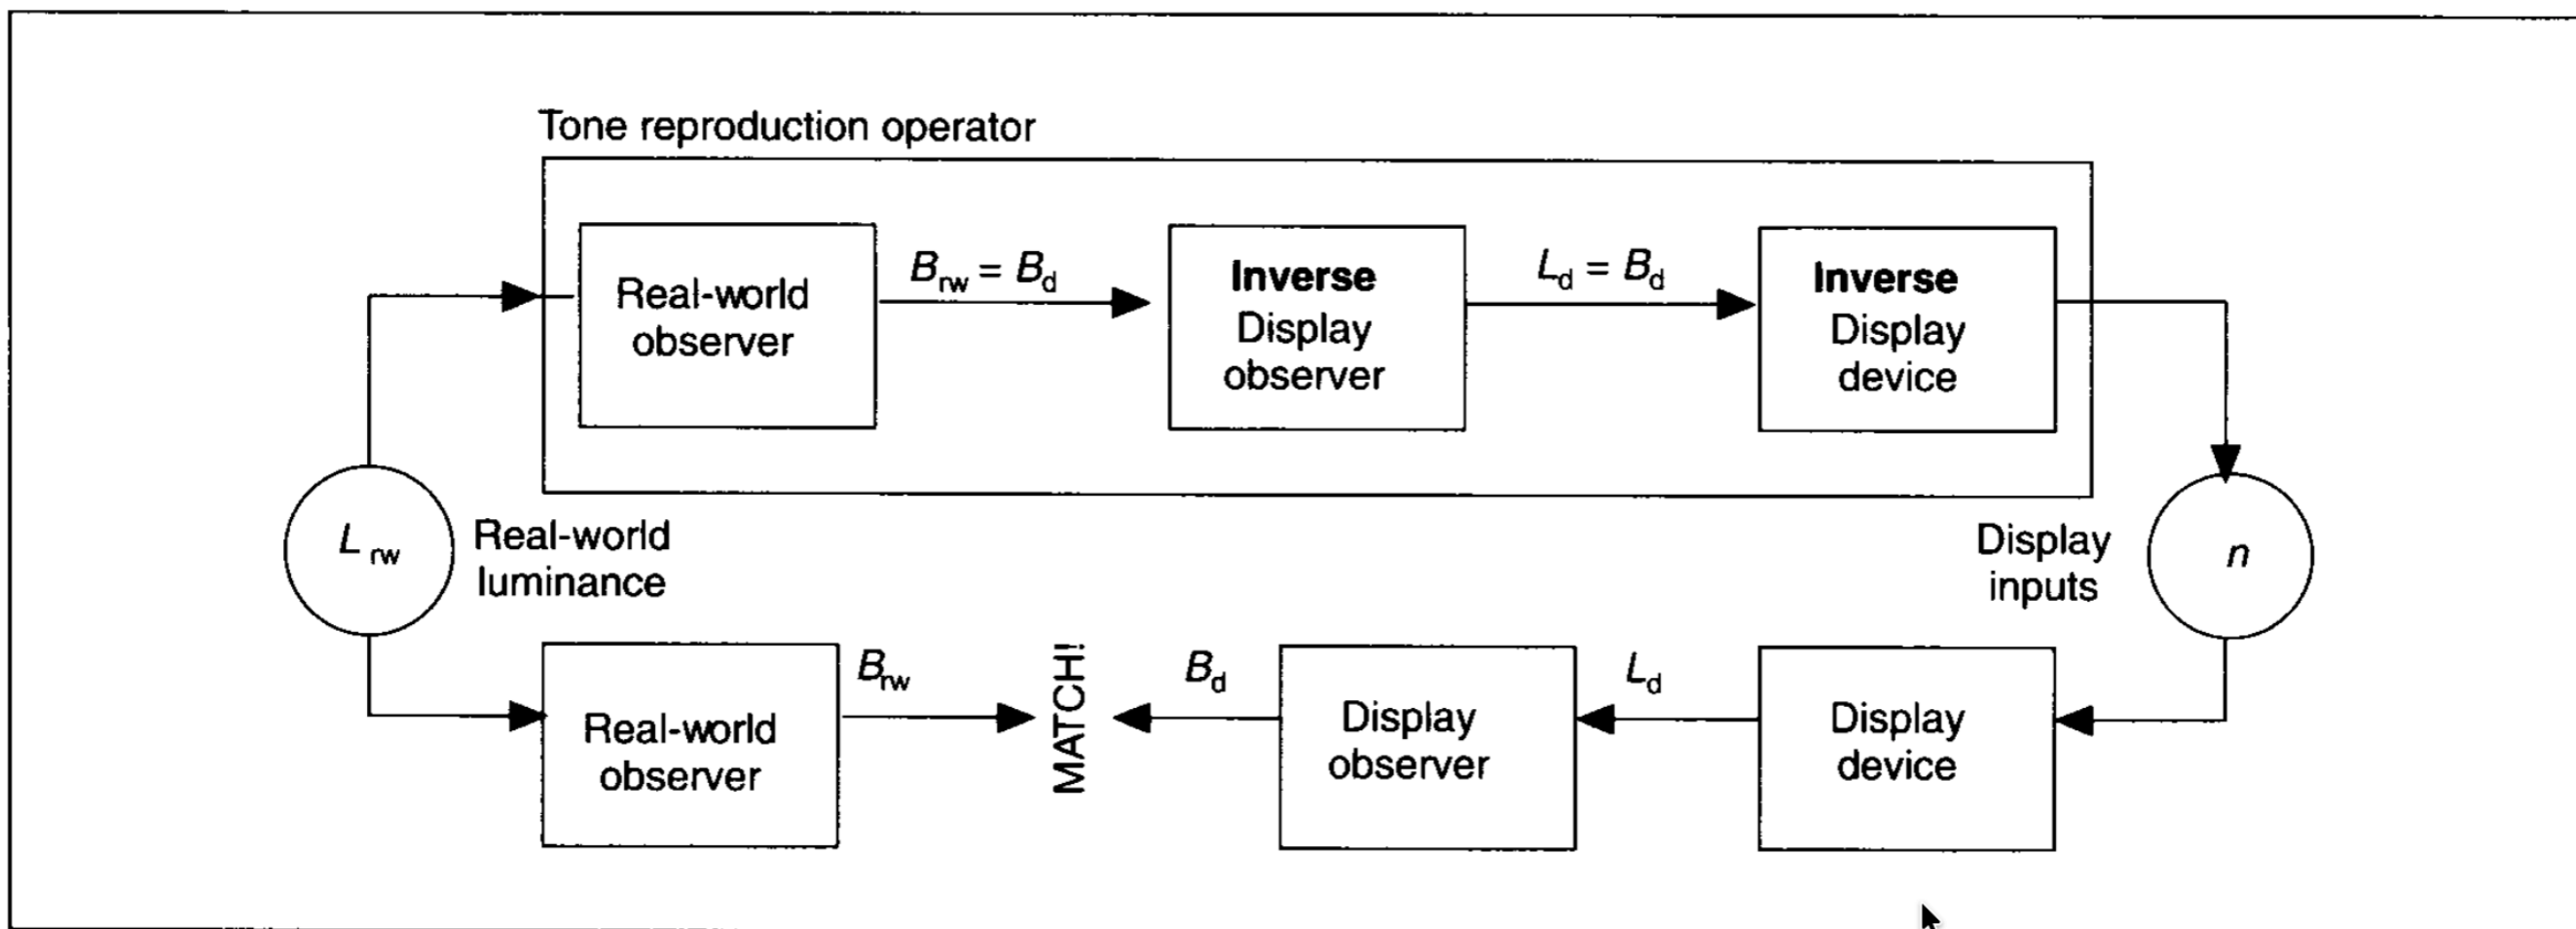
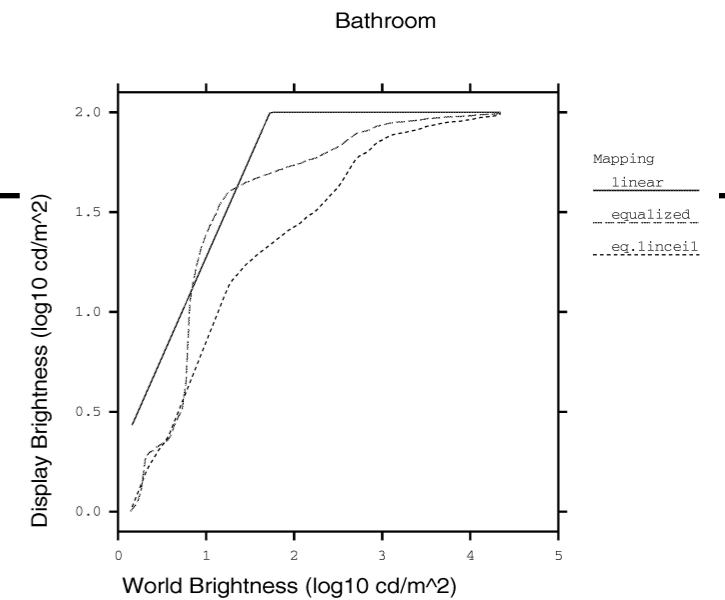


Figure 2. A simple solution: Make a tone reproduction operator by concatenating a real-world observer model, an inverse display observer model, and an inverse display device model.

Ward et al. 1997

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

Brightness Mapping Function



linear



histogram equalized



with linear ceiling constraint

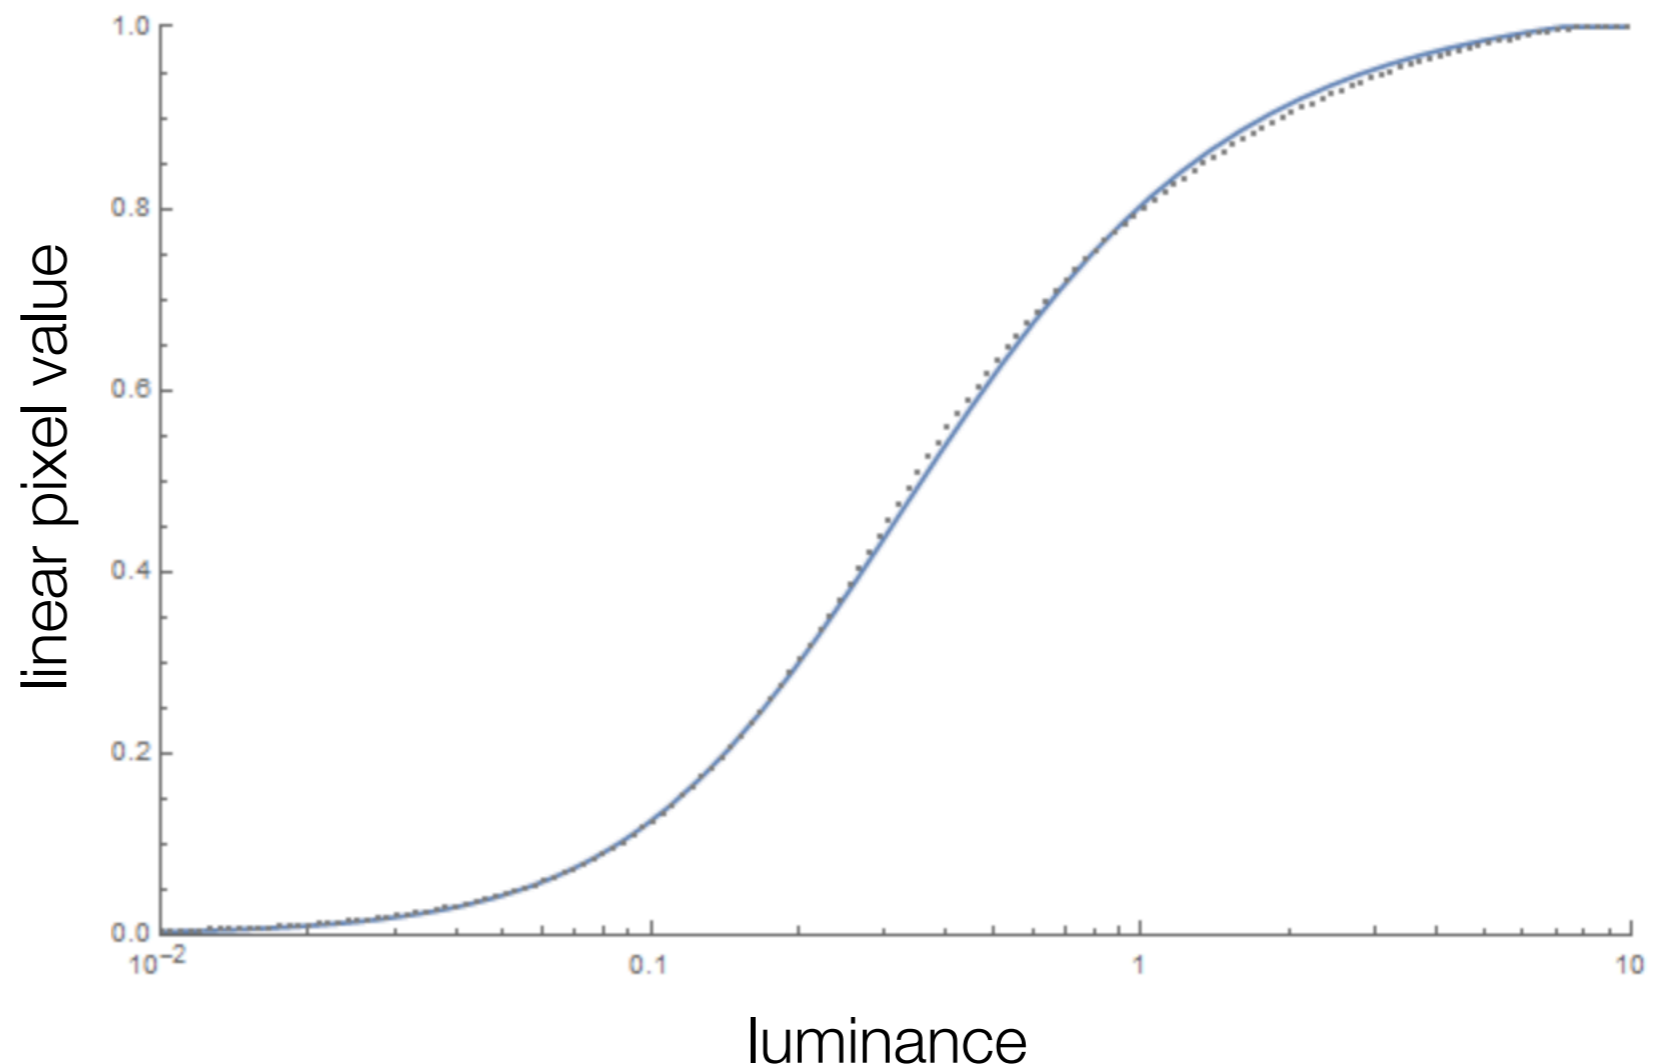
ACES “filmic” tonemapping

- **Principle**

want to smoothly decrease contrast for bright pixels
they already solved this problem in the film era; do it like that

- **ACES standard**

same standard as
ACES color space
renderers often
borrow the
(empirical) tone
curve



High exposure
Linear sRGB



High exposure
ACES sRGB



Low exposure
Linear sRGB



Low exposure
ACES sRGB



Pointwise transforms


- **Tumblin & Rushmeier 1993**
- **Ward et al. 1997**
- **ACES “filmic” tone curve**
- **problem: limited compression achievable**

Need for spatially varying operators


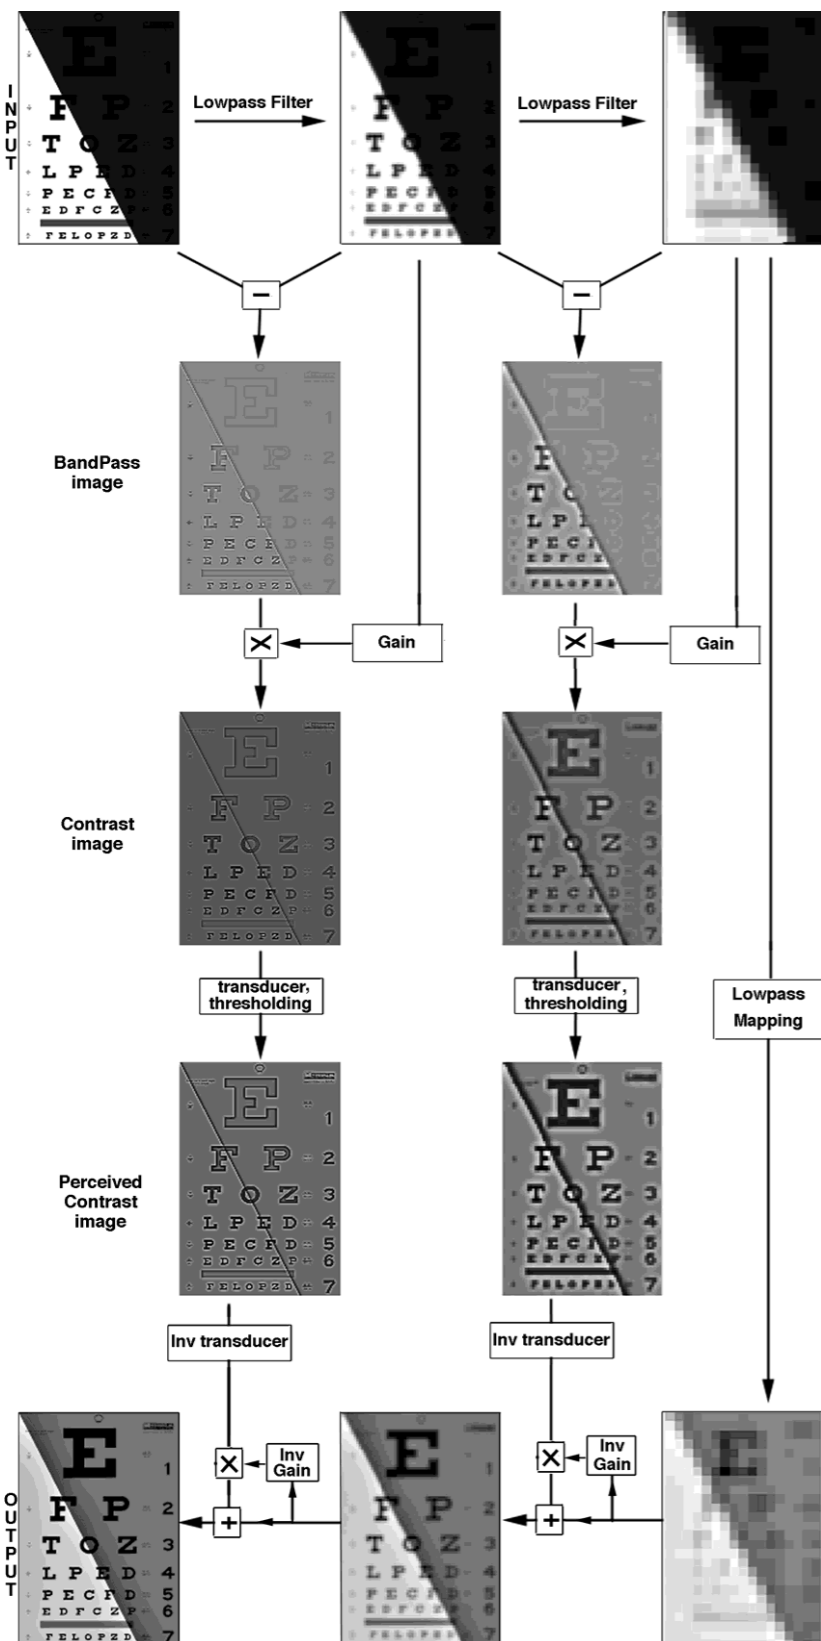
- **Chiu et al. 1993**
- **Pattanaik et al. 1998**
- **Reinhard 2002**
- **Ashikhmin 2002**
- **Scale defined by linear filtering**

Chiu et al. 1993

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

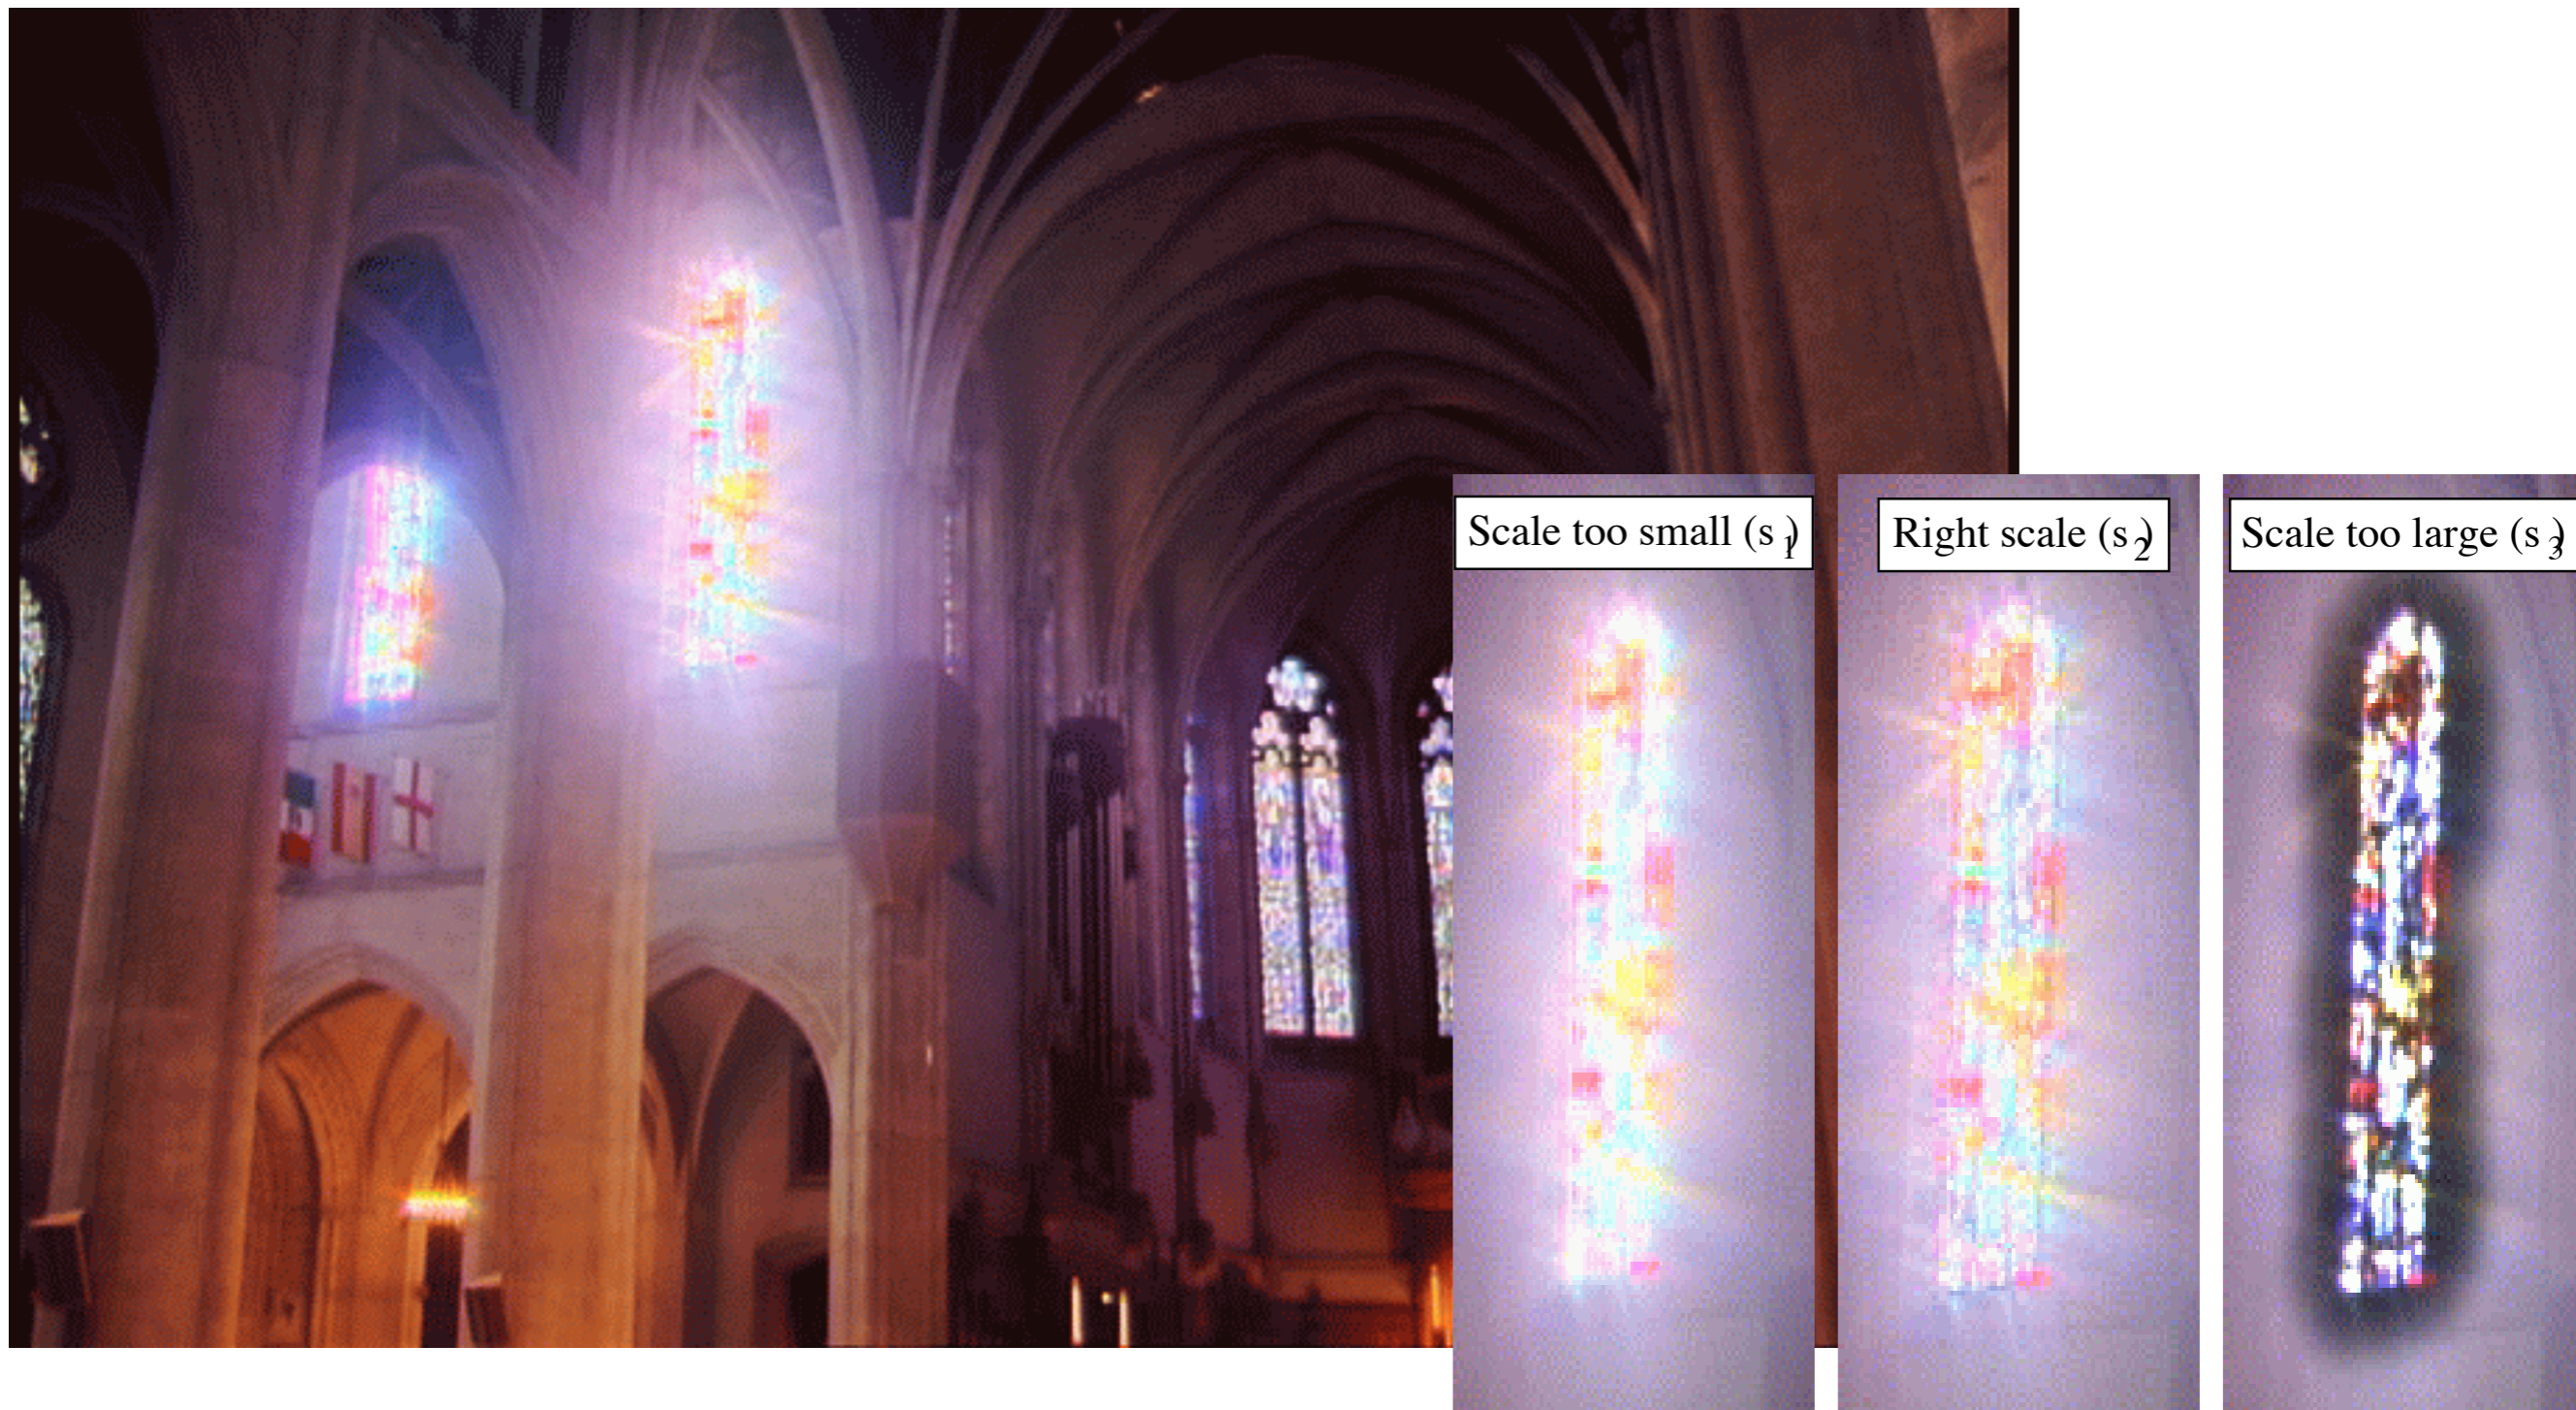


Pattanaik et al. 1998



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)**

- **...but first, bilateral filtering**

Denoising from 1 image

- We can't take average over multiple images



Denoising from 1 image


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



Gaussian blur

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

After Gaussian blur



Gaussian blur

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

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

After Gaussian blur



Anisotropic diffusion

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

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

“heat conductance”

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

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

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**

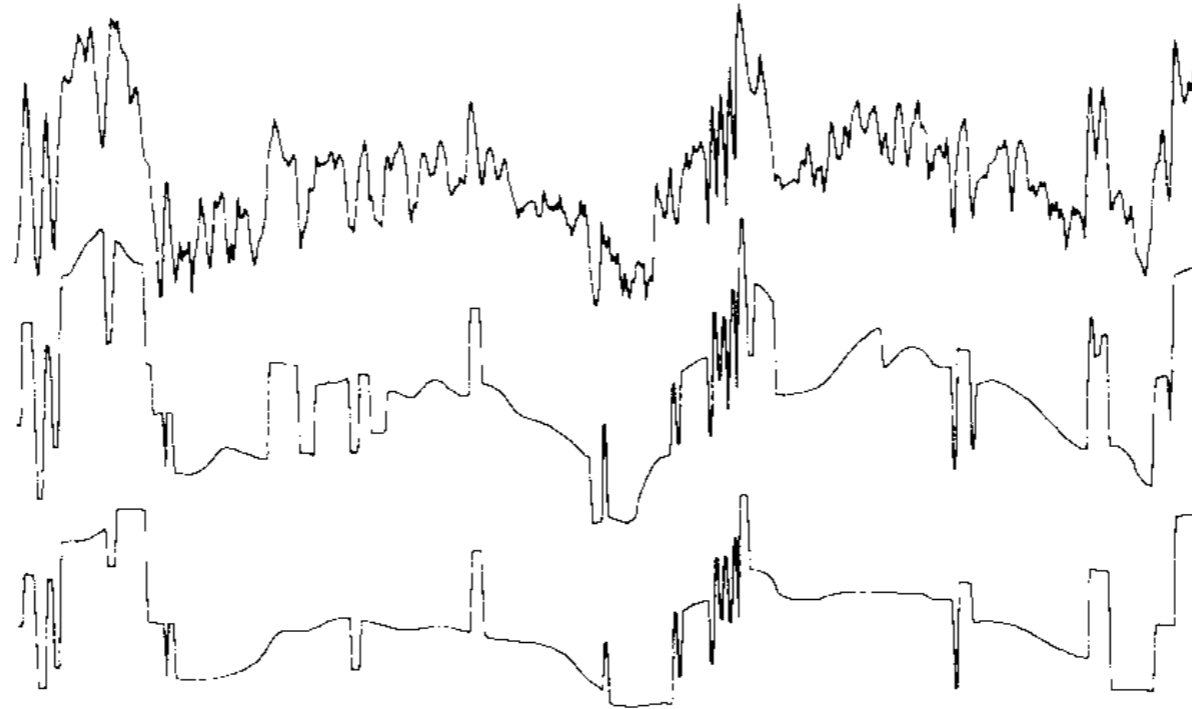


Fig. 11. Scale space obtained with anisotropic diffusion. The diffusion was performed in 2-D on the Canaletto image of which one line (the horizontal line number 400 out of 480—just above the gondola) is shown. Notice that the edges remain sharp until their disappearance.

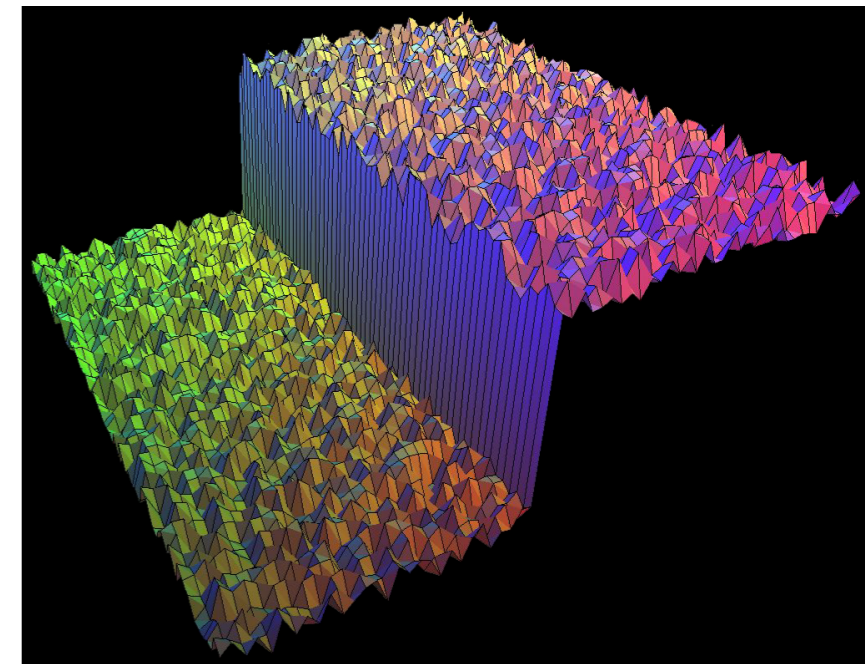


Bilateral filter

- **[Tomasi and Manduchi 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**


Start with Gaussian filtering

- Here, input is a step function + noise

$$J = f \otimes I$$



output

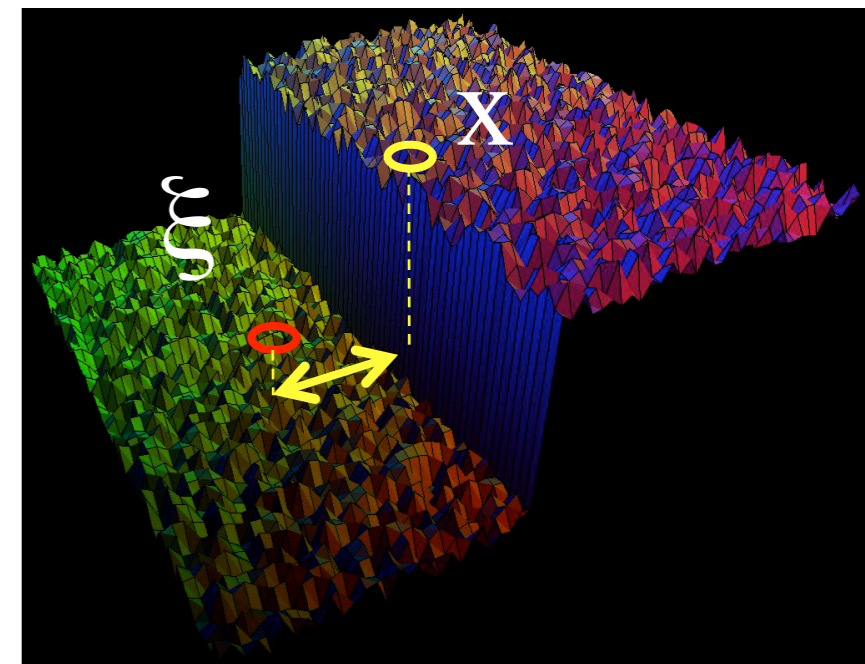
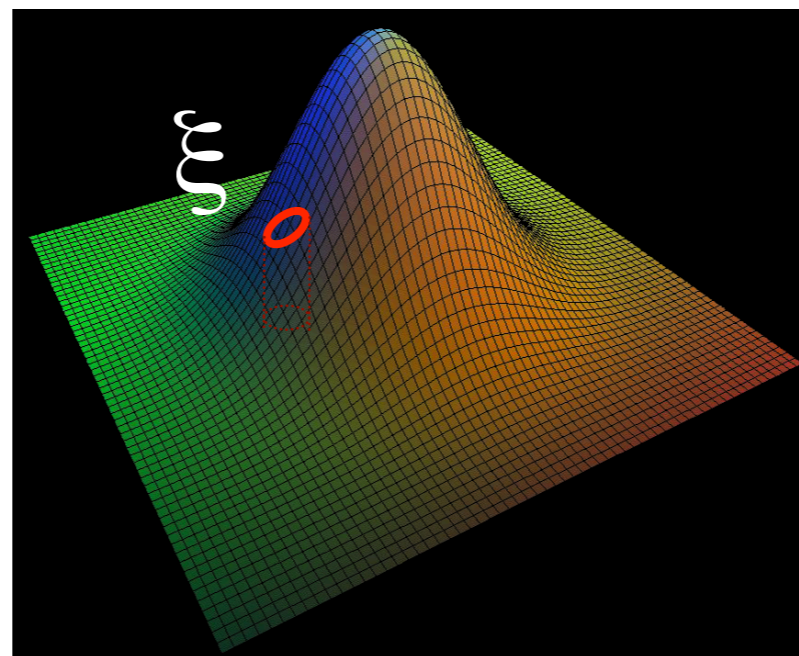



input


Gaussian filter as weighted average

- Weight of ξ depends on distance to x

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



output

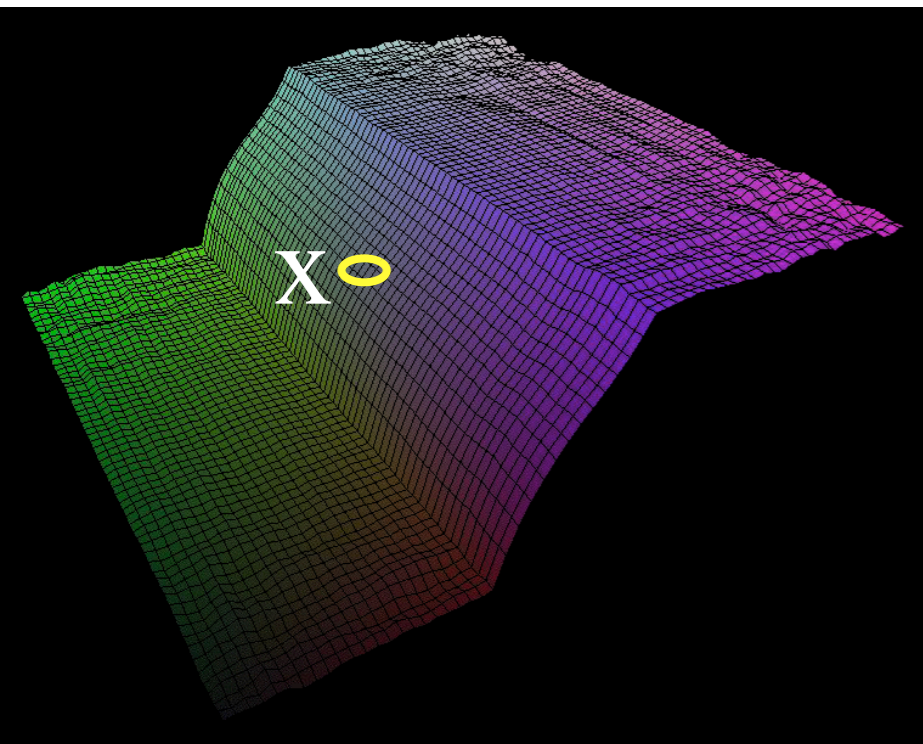


input



The problem of edges

- Here, $I(\xi)$ “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

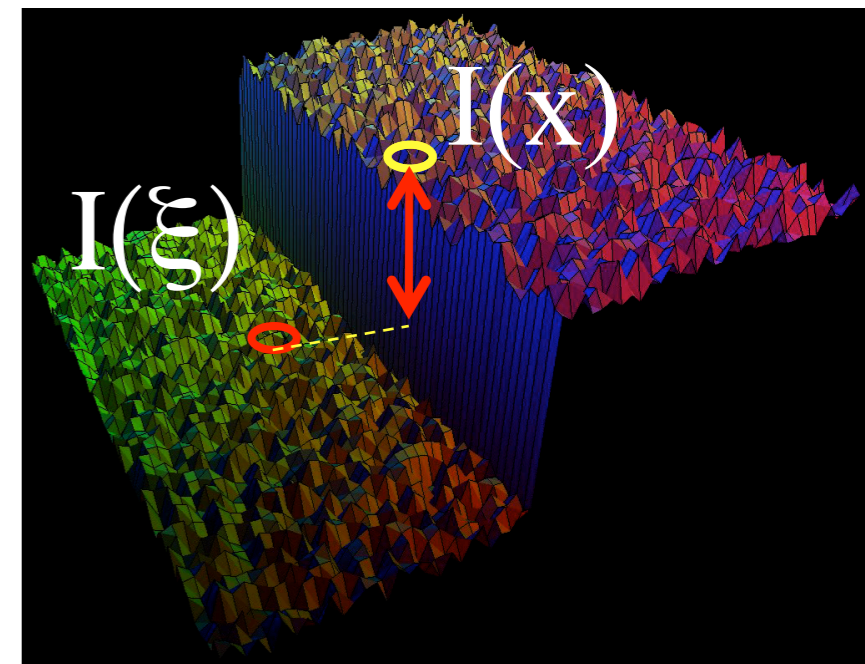
[Tomasi and Manduchi 1998]

- 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




Bilateral filtering

[Tomasi and Manduchi 1998]

- **Spatial Gaussian f**

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




Bilateral filtering

[Tomasi and Manduchi 1998]

- Spatial Gaussian f
- Gaussian g on the intensity difference

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

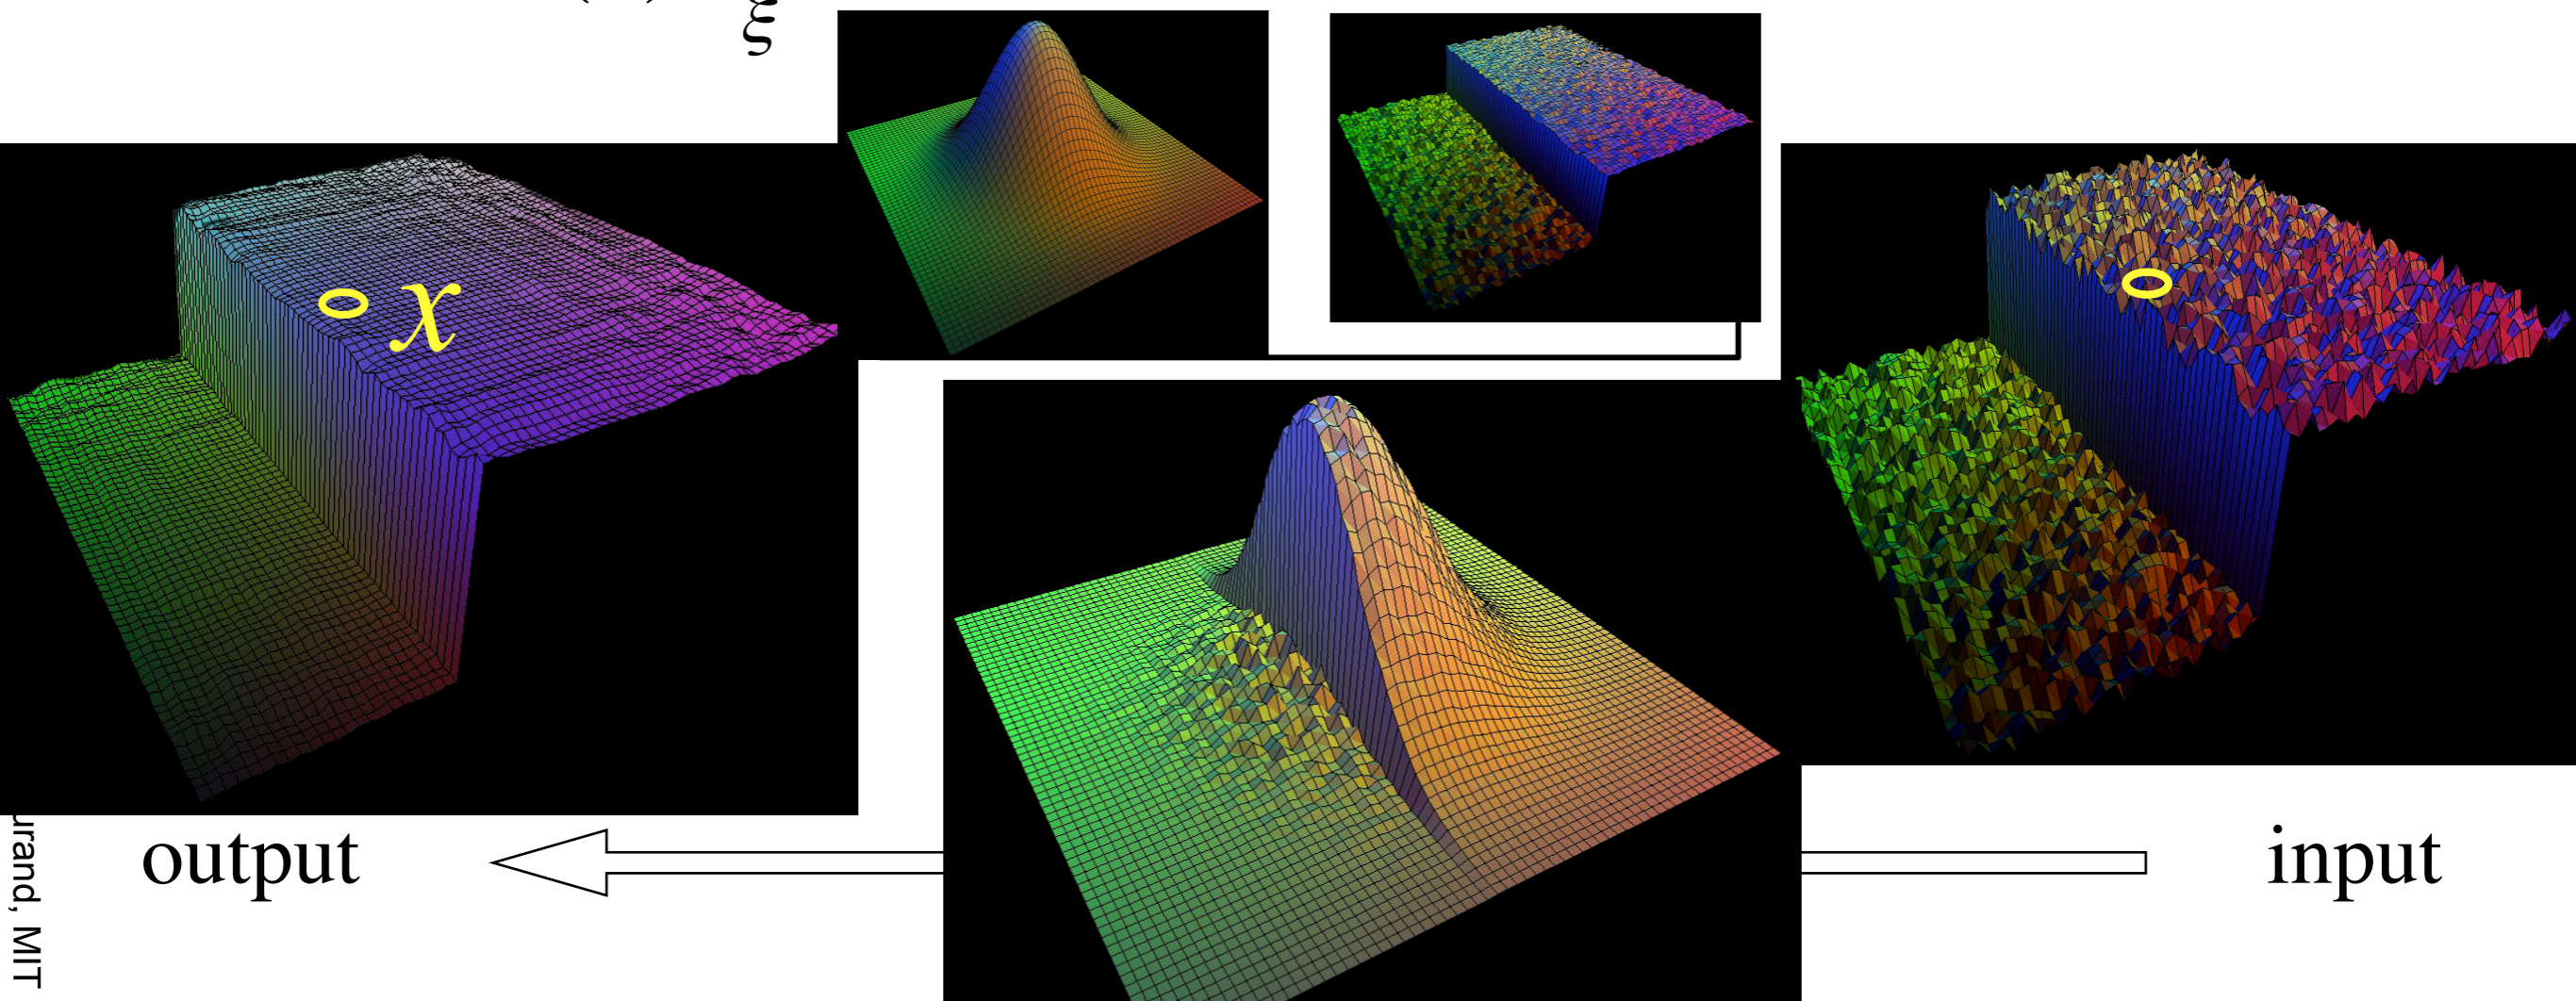


Normalization factor

[Tomasi and Manduchi 1998]

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

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




Bilateral filtering is non-linear

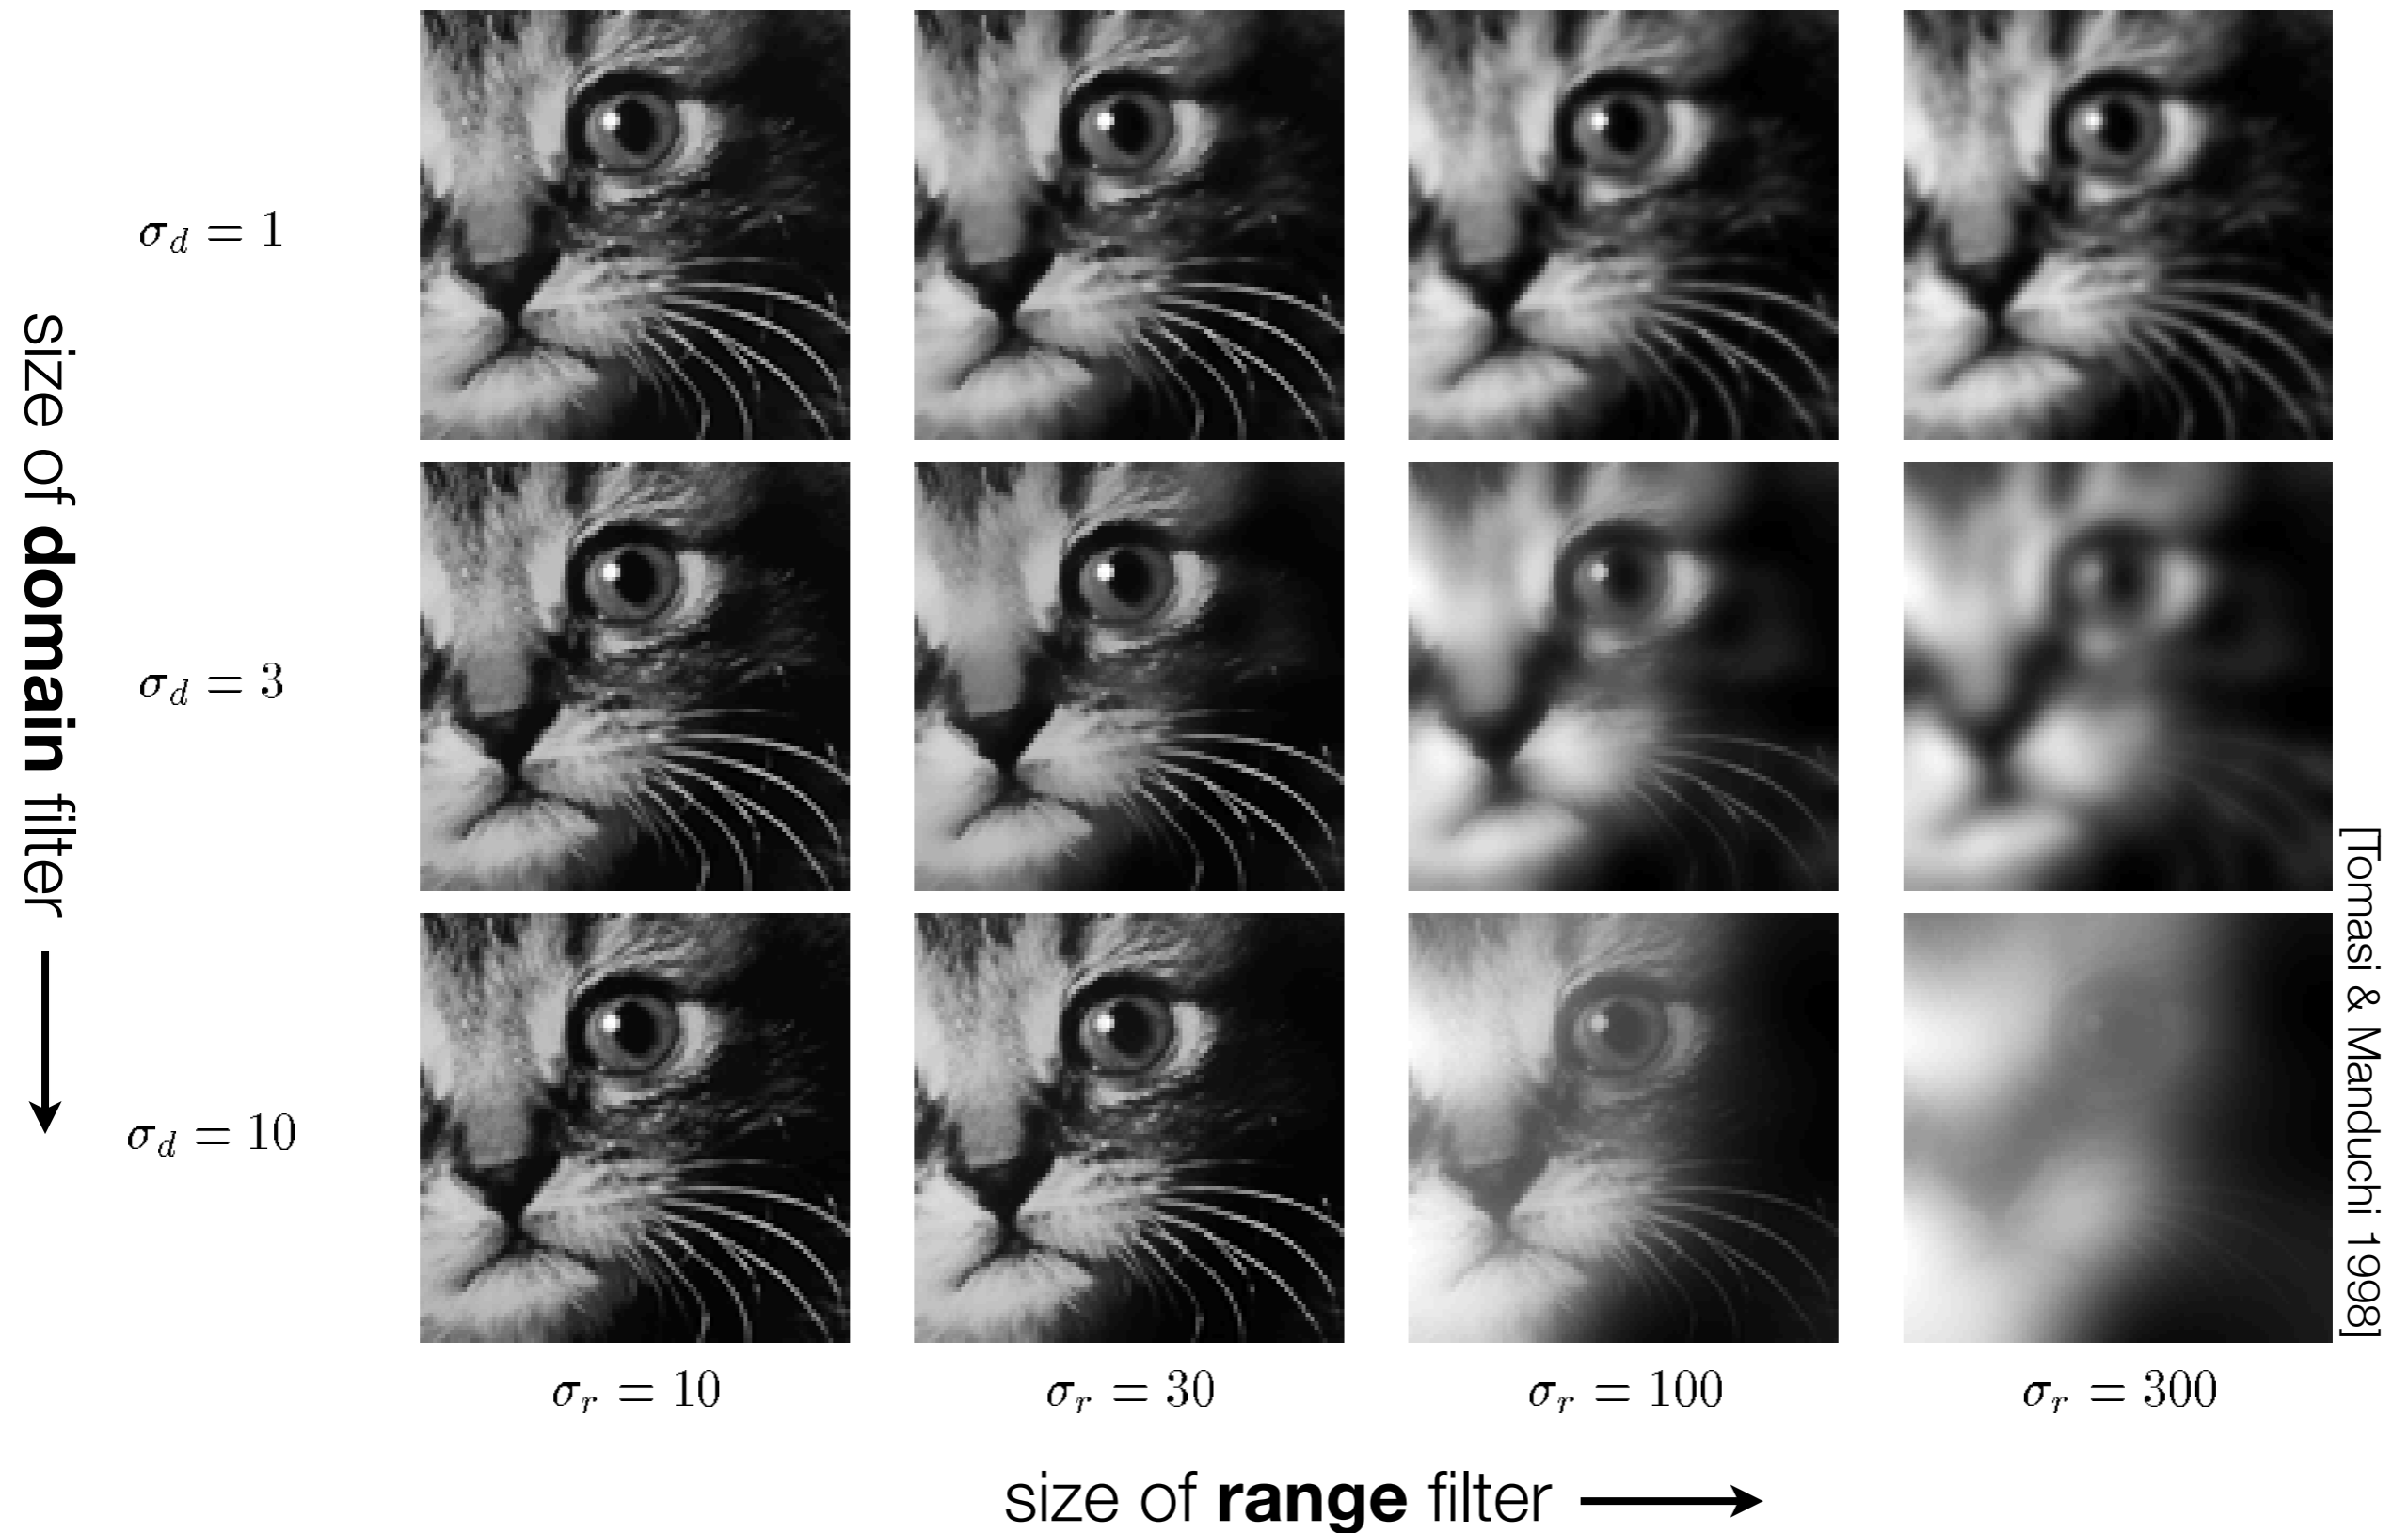
[Tomasi and Manduchi 1998]

- The weights are different for each output pixel

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




Effects of bilateral filter



Bilateral filter



Noisy input



After gaussian blur

Bilateral filter



Noisy input



After bilateral filter

Can we do better?



Noisy input



chroma
noise

After bilateral filter

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

Normal RGB Bilateral filter



Noisy input



After bilateral filter

YUV bilateral filter



Noisy input



After YUV bilateral filter

Comparison



Noisy input



Bilateral filter



YUV bilateral filter

Bilateral filtering

- **Useful baseline noise-reduction filter for photos**
- **The basis of a baseline noise-reduction filter for rendered images**
- **A handy image processing tool**
 - e.g. Surface blur in photoshop
- **Can be used to separate details**
 - In particular, tone mapping for contrast reduction and high-dynamic-range imaging

Tumblin & Turk 1997

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



large-scale



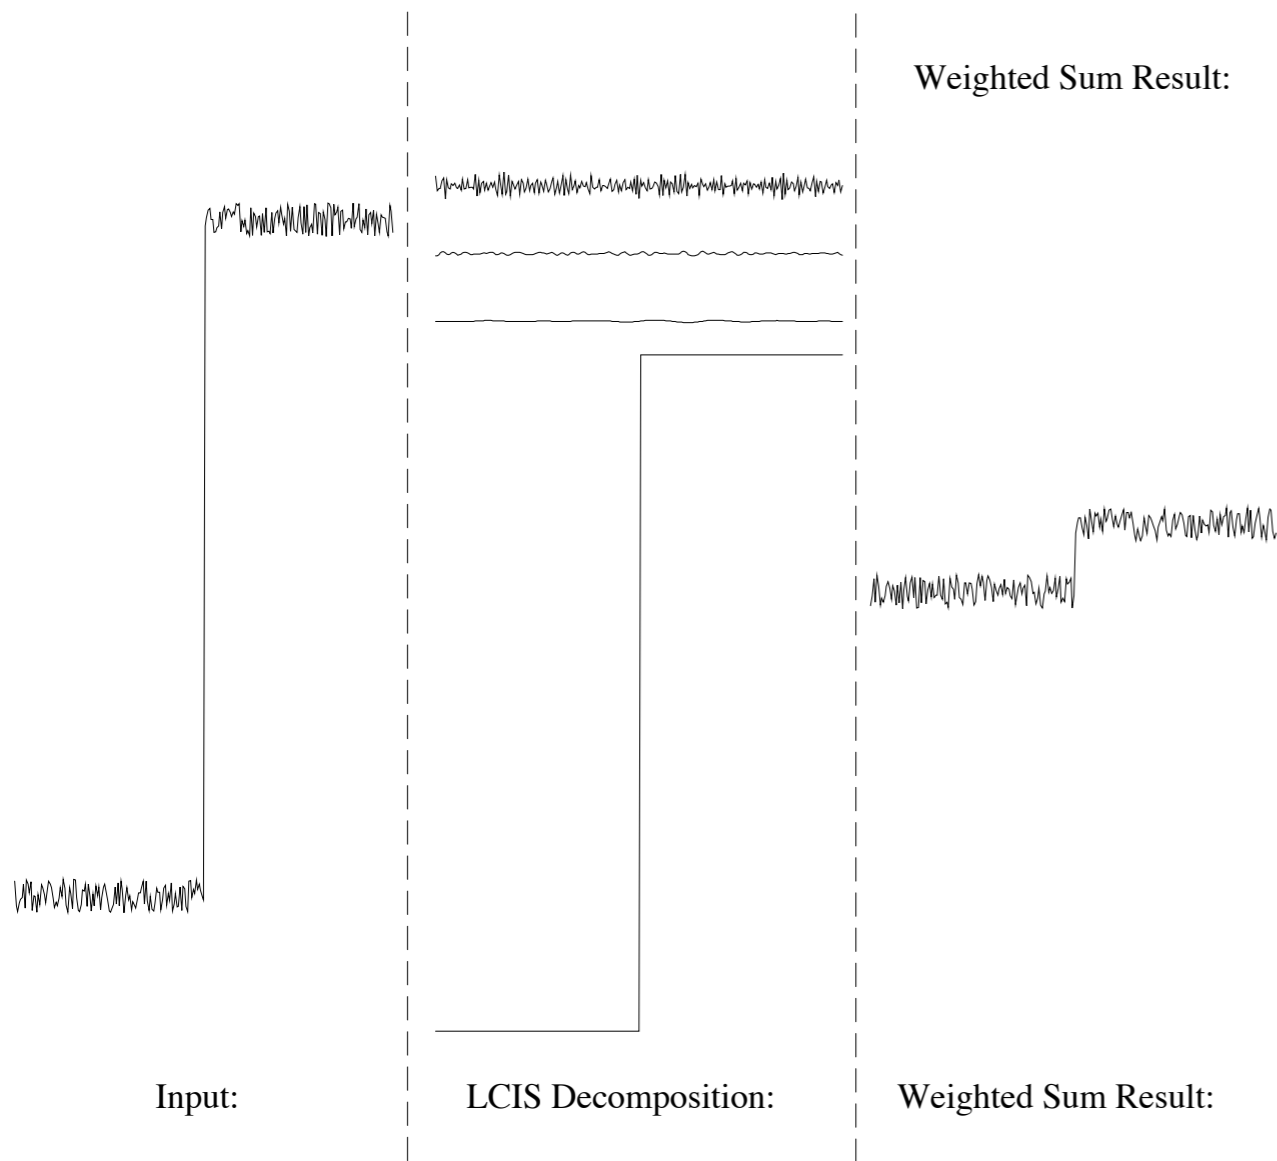
detail



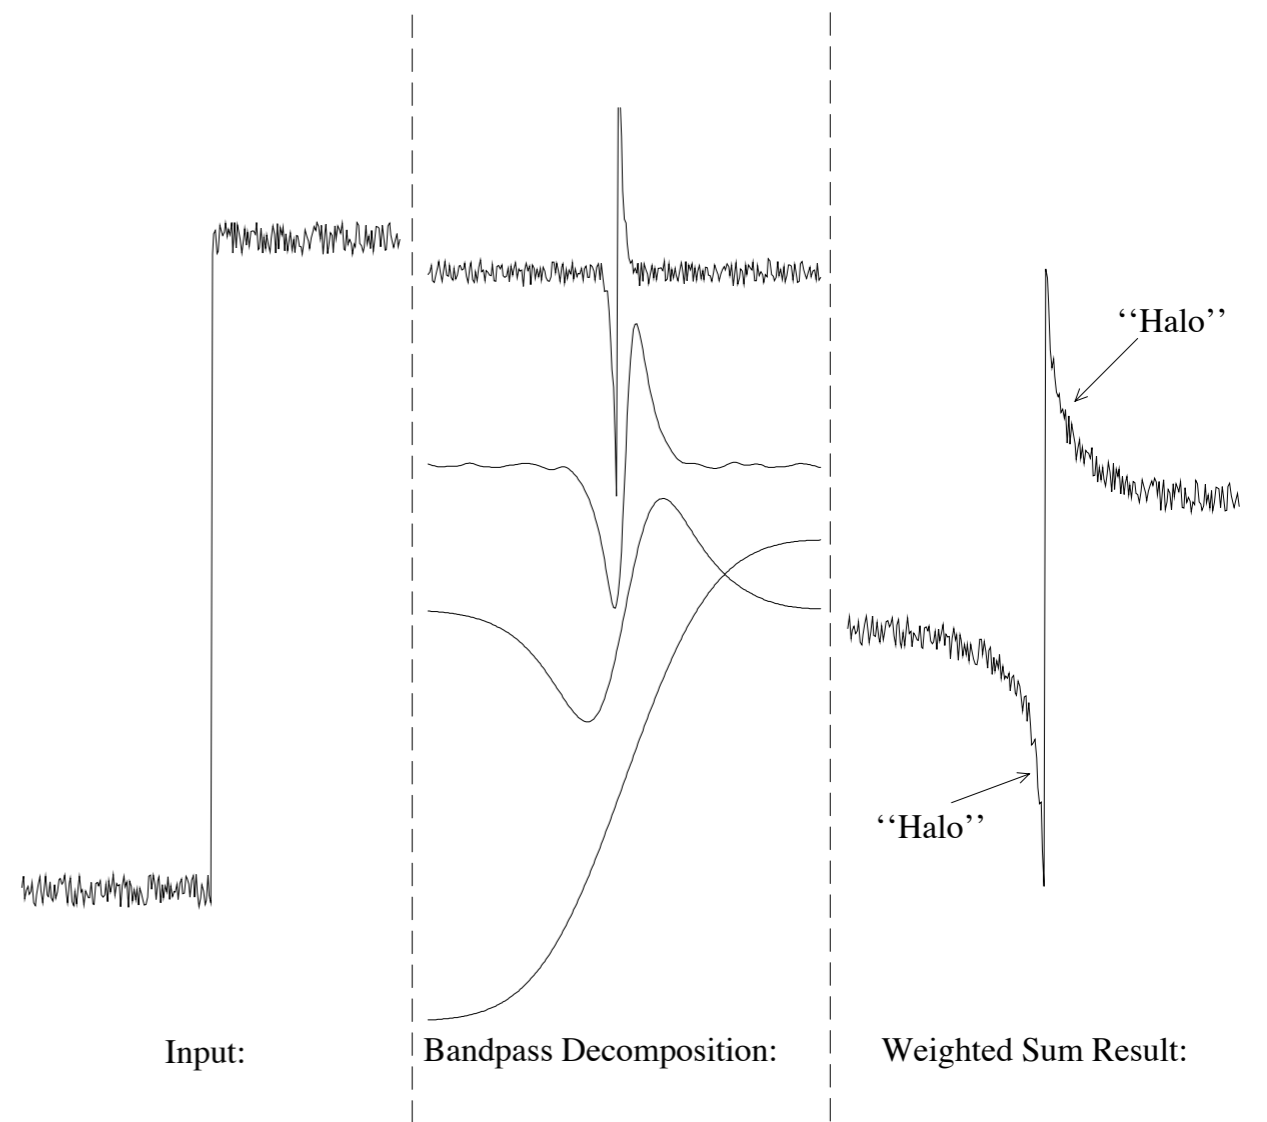
together



Tumblin & Turk 1997



Filtering using LCIS



Filtering using usual bandpass filters

Durand & Dorsey 2002

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



- **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).



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.