Guided Image Filtering

Kaiming He — The Chinese University of Hong Kong
Jian Sun — Microsoft Research Asia
Xiaoou Tang — The Chinese University of Hong Kong

Introduction

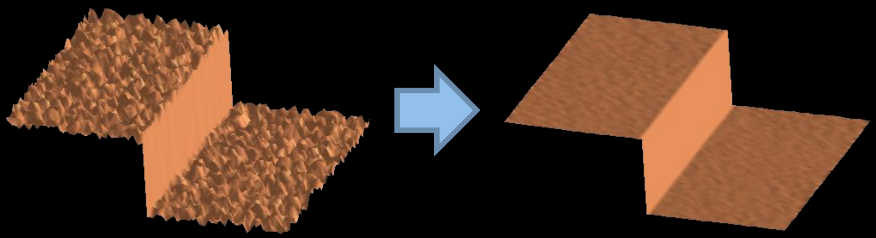

- Edge-preserving filtering
  - An important topic in computer vision
    - Denoising, image smoothing/sharpening, texture decomposition, HDR compression, image abstraction, optical flow estimation, image super-resolution, feature smoothing…
  - Existing methods
    - Weighted Least Square [Lagendijk et al. 1988]
    - Anisotropic diffusion [Perona and Malik 1990]
    - Bilateral filter [Aurich and Weule 95], [Tomasi and Manduchi 98]
    - Digital TV (Total Variation) filter [Chan et al. 2001]

Introduction

- Bilateral filter

$$q_i = \sum_{j \in N(i)} W_{ij}(p) p_j$$

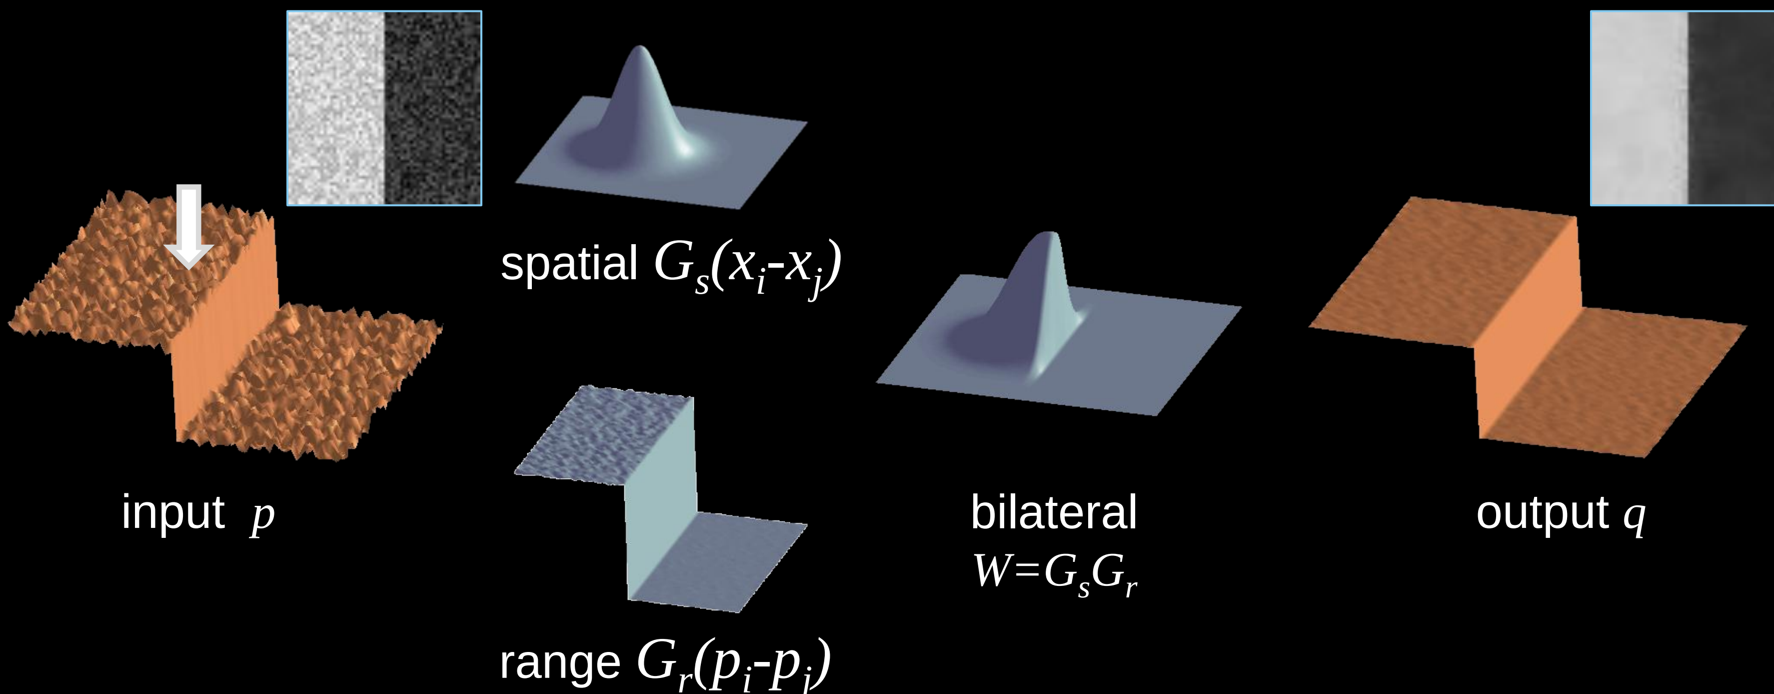

Introduction

- Joint bilateral filter [Petschnigg et al. 2004]

$$q_i = \sum_{j \in N(i)} W_{ij}(I) p_j$$

bilateral filter: $I=p$

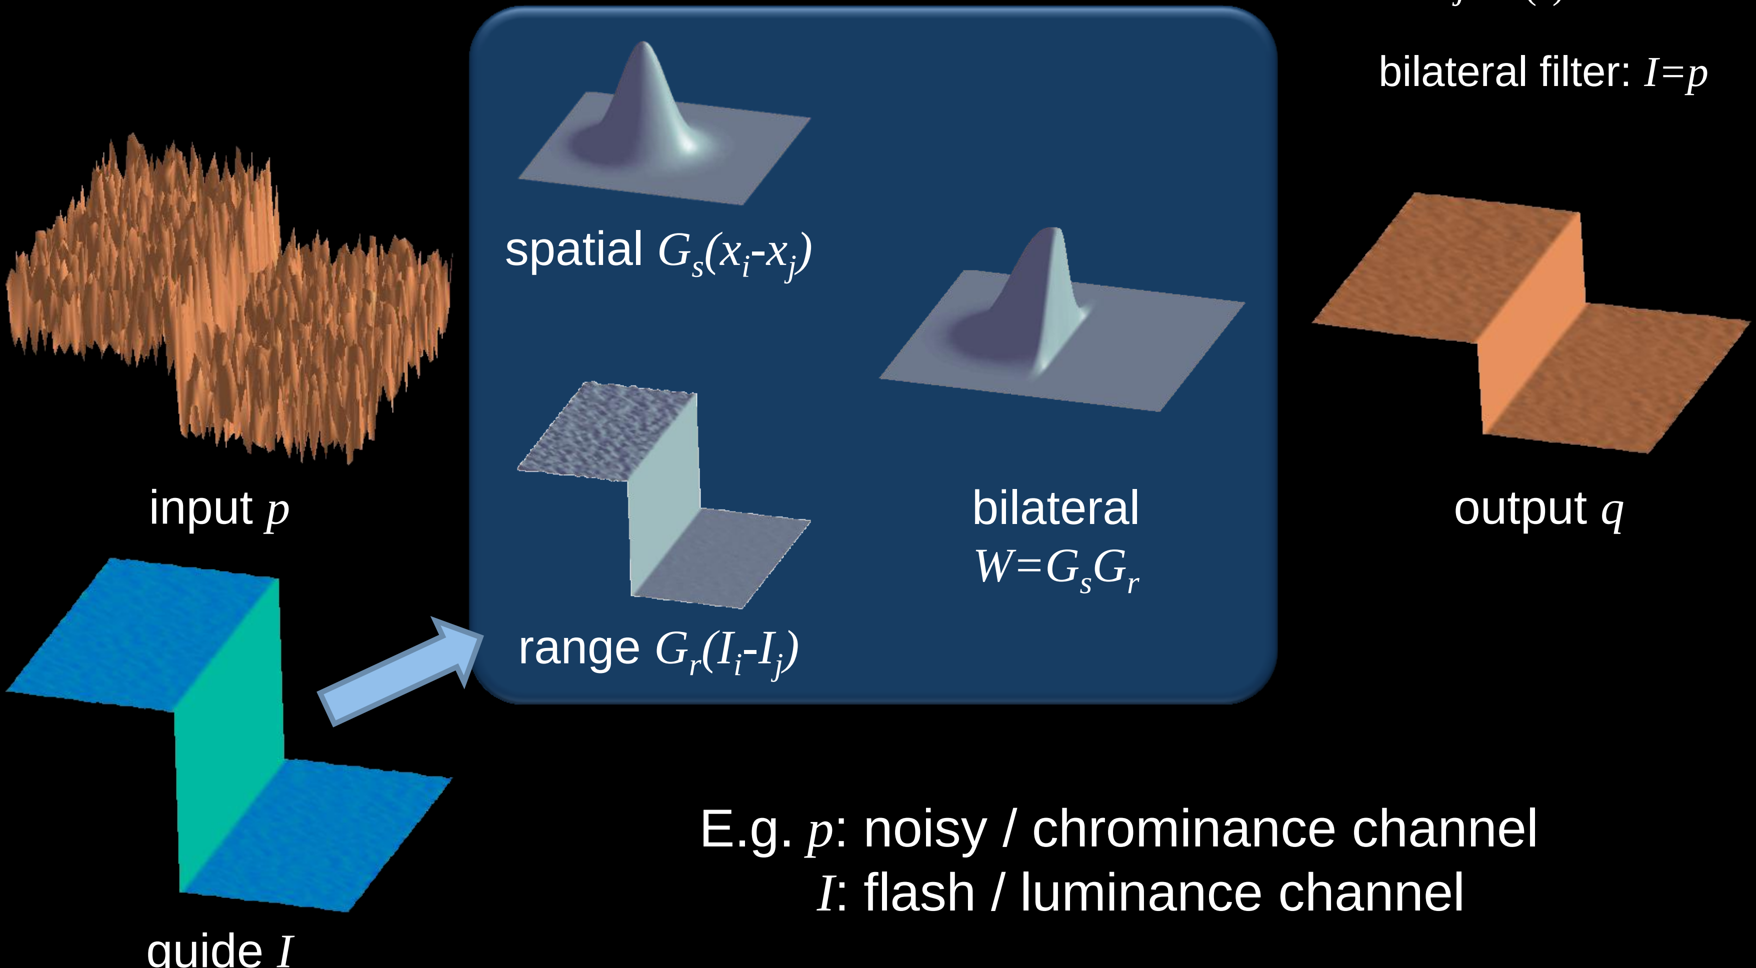

E.g. p: noisy / chrominance channel
I: flash / luminance channel

Introduction

- Advantages of bilateral filtering
  - Preserve edges in the smoothing process
  - Simple and intuitive
  - Non-iterative

Introduction

- Problems in bilateral filtering
  - Complexity
    - Brute-force: O(r^2)
    - Distributive histogram: O($\log r$) [Weiss 06]
    - Bilateral grid: band-dependent [Paris and Durand 06], [Chen et al. 07]
    - Integral histogram: O(1) [Porikli 08], [Yang et al. 09]

Approximate (quantized)

Introduction

- Problems in bilateral filtering
  - Complexity
  - Gradient distortion
    - Preserves edges, but not gradients

Example: detail enhancement

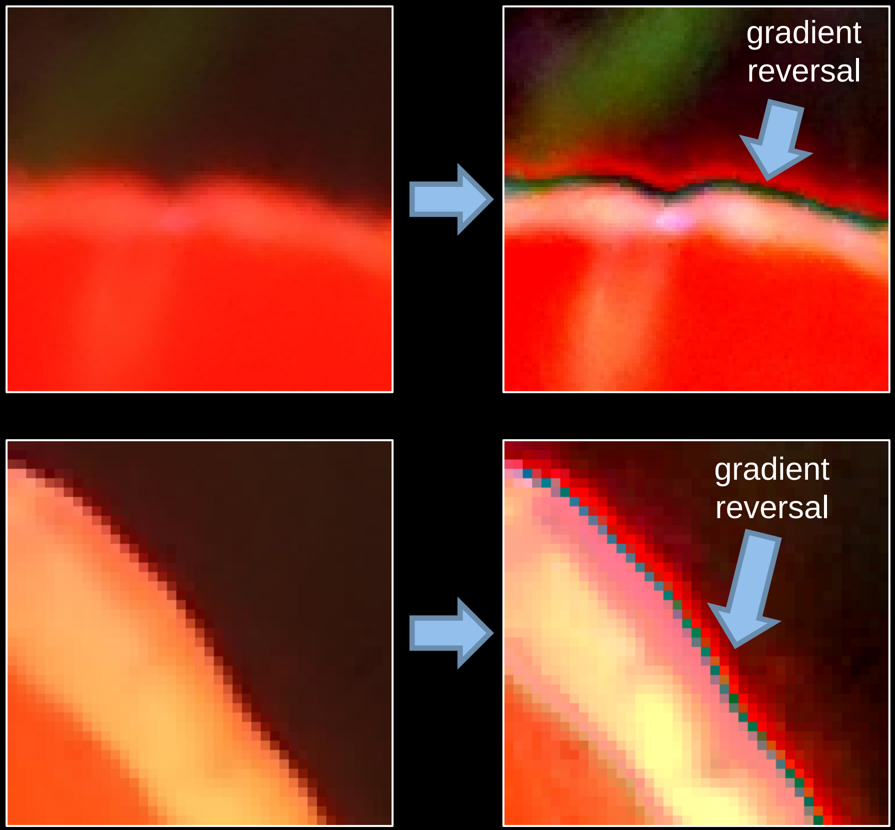

input enhanced

Introduction

- Our target - to design a new filter
  - Edge-preserving filtering
  - Non-iterative
  - O(1) time, fast and non-approximate
  - No gradient distortion

Advantages of bilateral filter (Edge-preserving filtering; Non-iterative)

Overcome bilateral filter's problems (O(1) time, fast and non-approximate; No gradient distortion)

Guided filter

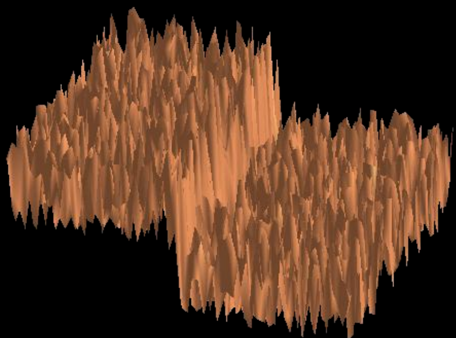

input p

$$q_i = p_i - n_i$$

n_i - noise / texture

$$\min_{(a,b)} \sum_i (aI_i + b - p_i)^2 + \varepsilon a^2$$

Linear regression

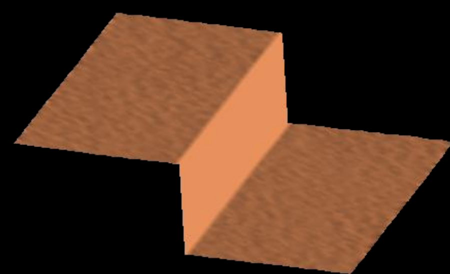

output q

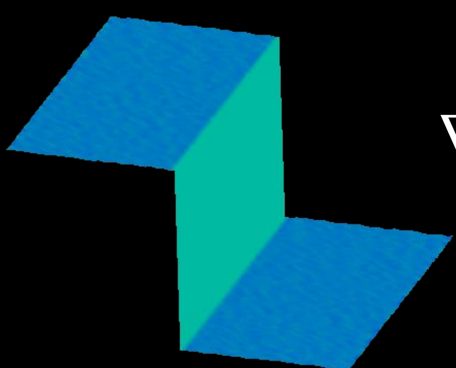

guide I

$$\nabla q_i = a \nabla I_i$$

$$q_i = aI_i + b$$

$$\begin{cases} a = \dfrac{\operatorname{cov}(I,p)}{\operatorname{var}(I) + \varepsilon} \\ b = \bar{p} - a\bar{I} \end{cases}$$

Bilateral/joint bilateral filter does not have this linear model

Guided filter

- Extend to the entire image
  - In all local windows ω_k ,compute the linear coefficients
  - Compute the average of $a_k I_i + b_k$ in all ω_k that covers pixel q_i

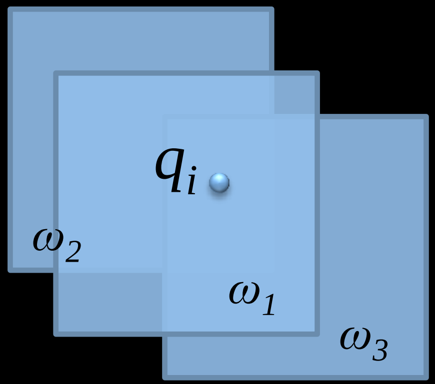

Definition

$$a_k = \frac{\text{cov}_k(I, p)}{\text{var}_k(I) + \varepsilon}$$

$$b_k = \bar{p}_k - a\bar{I}_k$$

$$q_i = \frac{1}{|\omega|} \sum_{k | i \in \omega_k} (a_k I_i + b_k)$$

$$= \bar{a}_i I_i + \bar{b}_i$$

Guided filter

- Parameters
  - Window radius r
  - Regularization ε

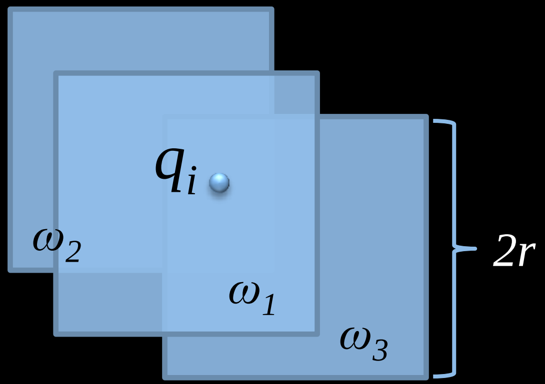

Definition

$$a_k = \frac{\mathrm{cov}_k(I, p)}{\mathrm{var}_k(I) + \varepsilon}$$

$$b_k = \bar{p}_k - a\bar{I}_k$$

$$q_i = \frac{1}{|\omega|} \sum_{k|i \in \omega_k} (a_k I_i + b_k)$$

$$= \bar{a}_i I_i + \bar{b}_i$$

Guided filter: smoothing

a cascade of mean filters

$$a = \frac{\mathrm{cov}(I, p)}{\mathrm{var}(I) + \varepsilon}$$

$$b = \bar{p} - a\bar{I}$$

$\mathrm{var}(I) << \varepsilon$

$\mathrm{cov}(I, p) << \varepsilon$

$$a \approx 0$$

$$b \approx \bar{p}$$

$$q_i = \bar{a} I_i + \bar{b} \approx \bar{\bar{p}}$$

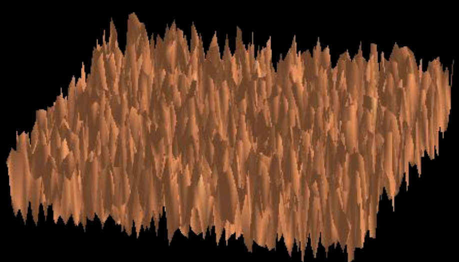

input p

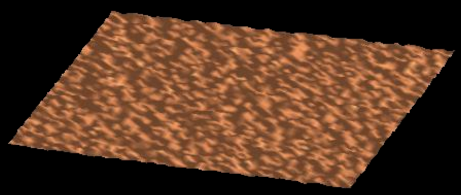

output q

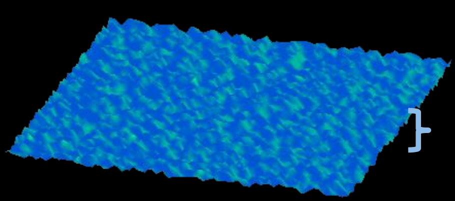

$\mathrm{var}(I) << \varepsilon$

guide I

r : determines band-width (like σ_s in BF)

Guided filter: edge-preserving

$$q_i = \bar{a} I_i + \bar{b} \quad \Rightarrow \quad \nabla q_i = \bar{a} \nabla I_i + I_i \nabla \bar{a} + \nabla \bar{b}$$

$$\varepsilon \downarrow \quad \Rightarrow \quad a = \frac{\mathrm{cov}(I, p)}{\mathrm{var}(I) + \varepsilon} \uparrow$$

ε : degree of edge-preserving (like σ_r in BF)

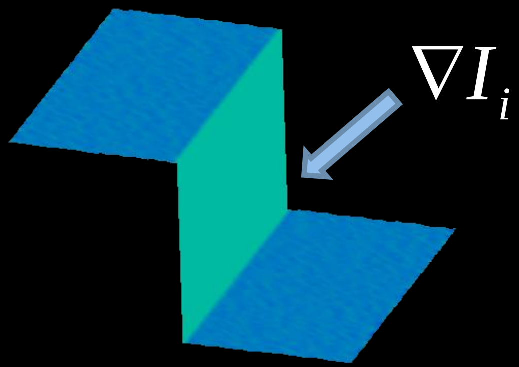

guide I

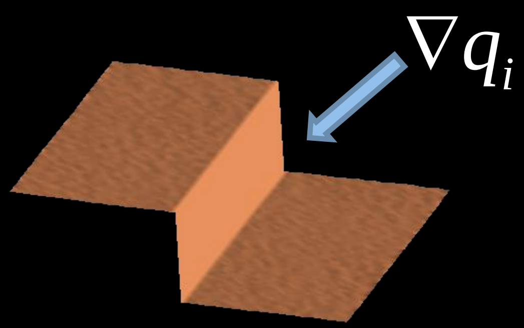

output q

Example – edge-preserving smoothing

input & guide

guided filter (let $I=p$)

$r=4,\ \varepsilon=0.1^2$

$r=4,\ \varepsilon=0.2^2$

$r=4,\ \varepsilon=0.4^2$

bilateral filter

$\sigma_s=4,\ \sigma_r=0.1$

$\sigma_s=4,\ \sigma_r=0.2$

$\sigma_s=4,\ \sigma_r=0.4$

- Our target - to design a new filter
  - Edge-preserving filtering
  - Non-iterative
  - O(1) time, fast and non-approximate
  - No gradient distortion

Advantages of bilateral filter

Overcome bilateral filter's problems

Complexity

- *mean*, *var*, *cov* in all local windows
- Integral images [Franklin 1984]
  - O(1) time – independent of *r*
  - Non-approximate

Definition

$$a_k = \frac{\mathrm{cov}_k(I, p)}{\mathrm{var}_k(I) + \varepsilon}$$

$$b_k = \bar{p}_k - a\bar{I}_k$$

$$q_i = \bar{a}_i I_i + \bar{b}_i$$

O(1) bilateral
(32-bin, 40ms/M)
[Porikli 08]

O(1) bilateral
(64-bin, 80ms/M)

O(1) guided
(exact, 80ms/M)

Gradient Preserving

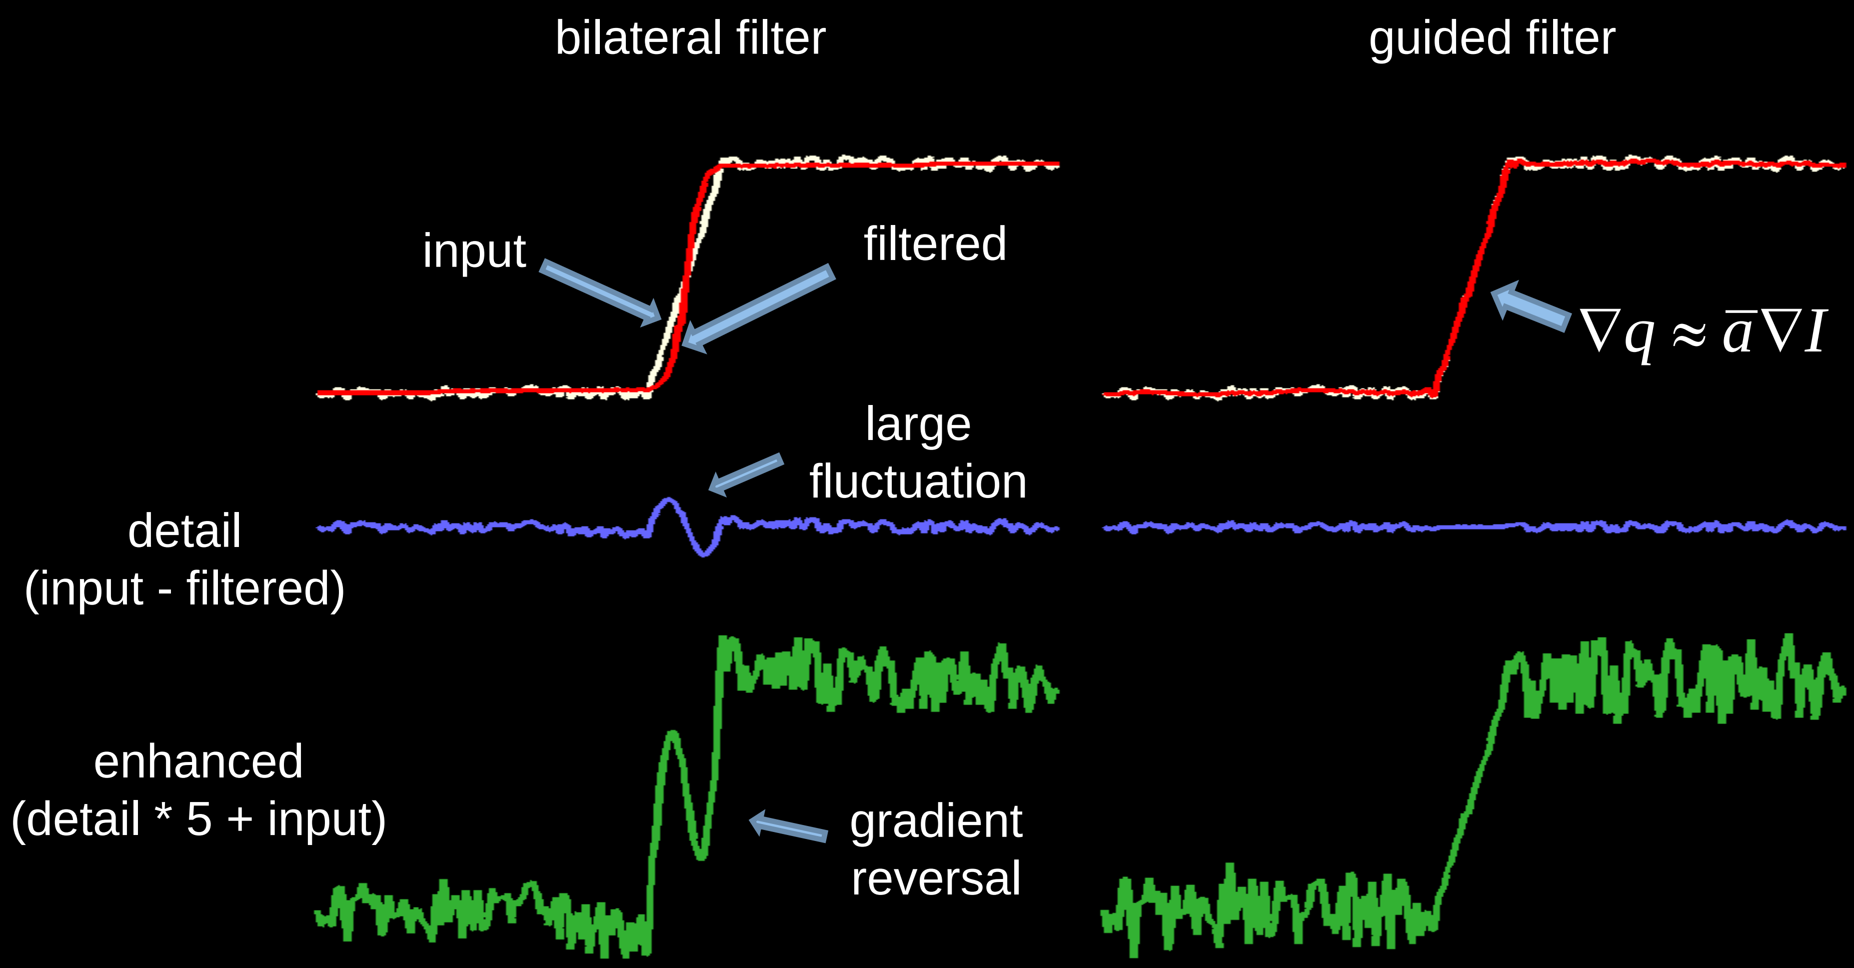

Example – detail enhancement

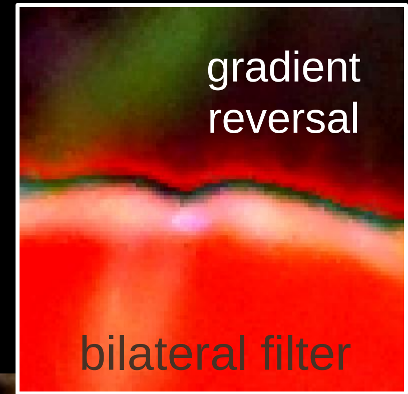

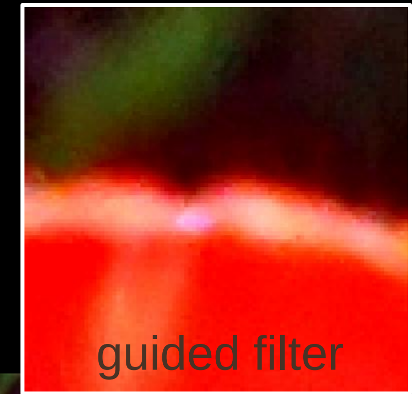

input ($I=p$)

bilateral filter
$\sigma_s=16, \sigma_r=0.1$

guided filter
$r=16, \varepsilon=0.1^2$

Example – detail enhancement

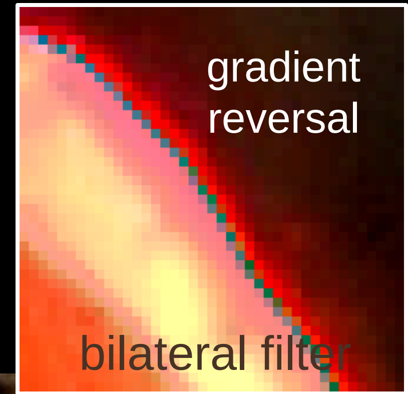

input (I=p)

bilateral filter
σ_s=16, σ_r=0.1

guided filter
r=16, ε=0.1^2

Example – HDR compression

input HDR

bilateral filter

guided filter

bilateral filter

σ_s=15, σ_r=0.12

guided filter

r=15, ε=0.12^2

Example – flash/no-flash denoising

input p
(no-flash)

joint bilateral filter
$\sigma_s=8, \sigma_r=0.02$

guide I
(flash)

guided filter
$r=8, \varepsilon=0.02^2$

joint bilateral

guided filter

Beyond smoothing

- Applications: feathering/matting, haze removal

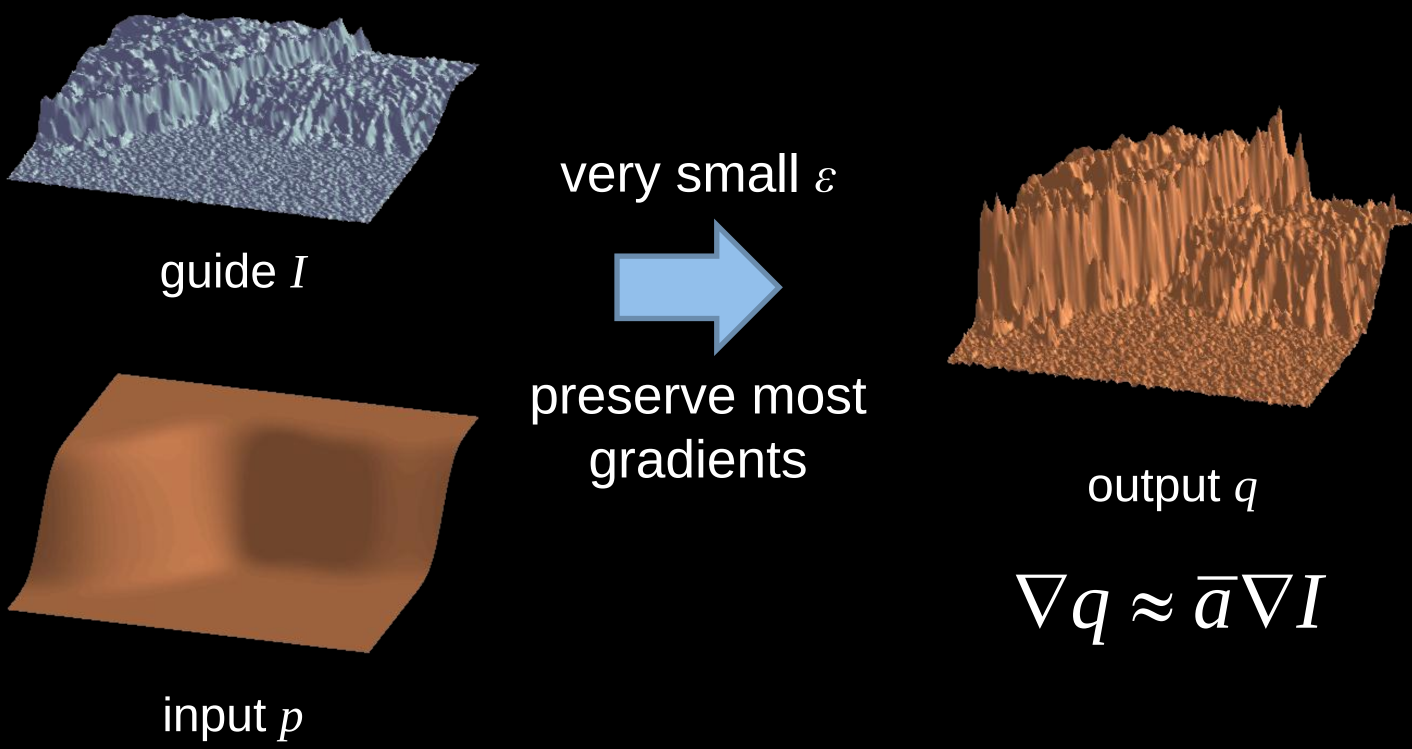

$$\nabla q \approx \bar{a} \nabla I$$

Example – feathering

guide I
(size 3000x2000)

Example – feathering

filter input p (binary segmentation)

Example – feathering

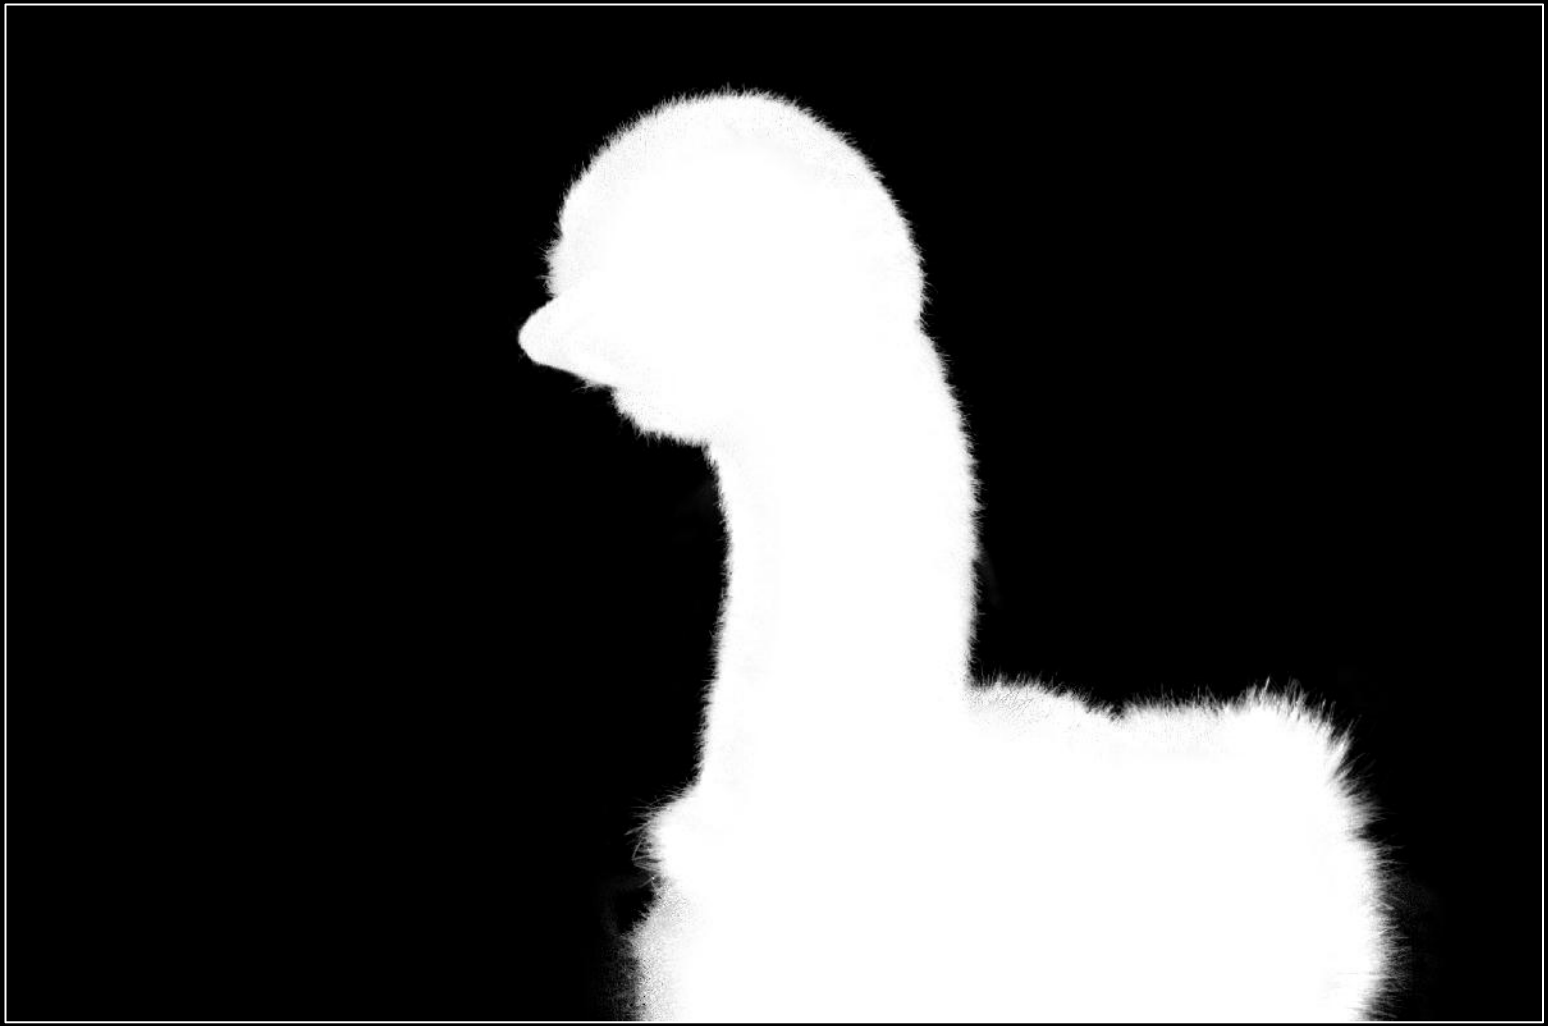

filter output q (alpha matte)

Example – feathering

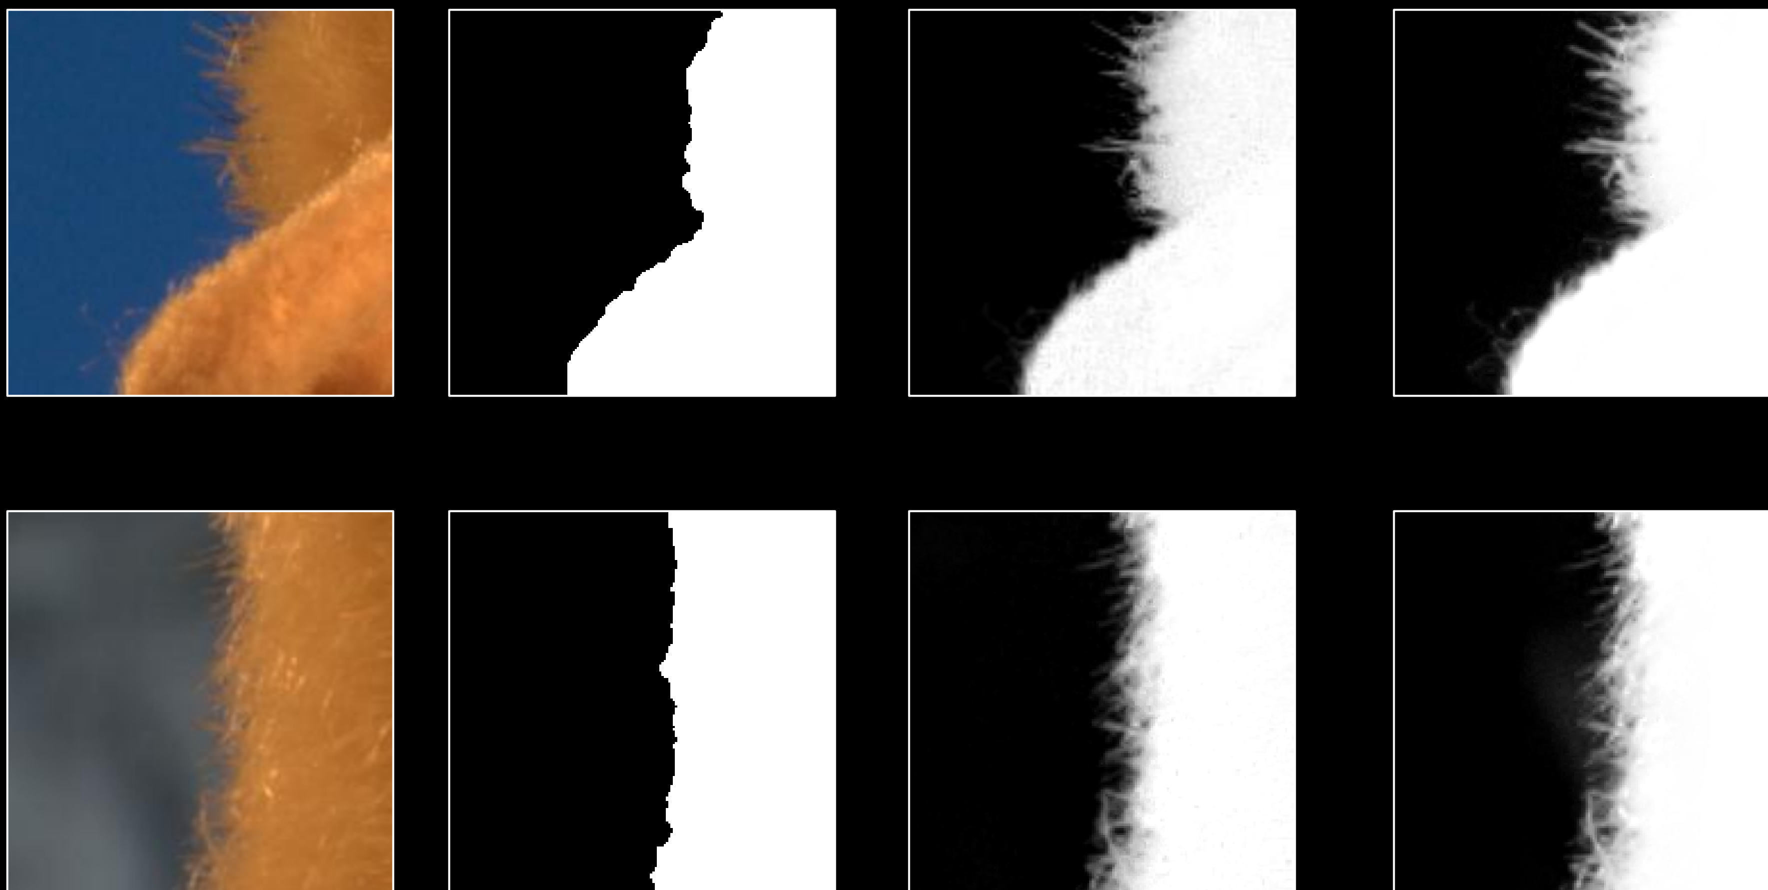

guide I

filter input p

filter output q
0.3s
image size 6M

matting Laplacian
[Levin et al. 06]
2 min

Example – haze removal

guide I

filter input p
(dark channel prior
[He et al. 09])

filter output q

Example – haze removal

guide I

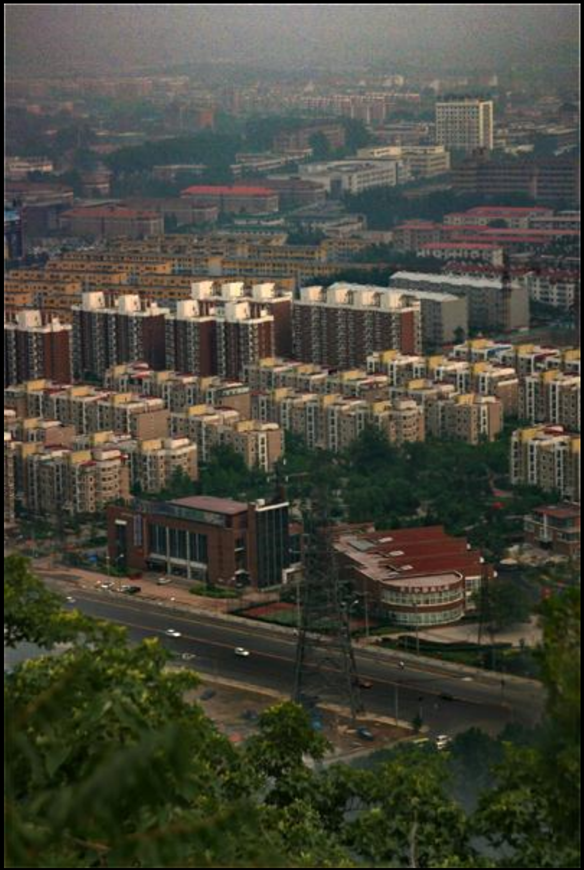

guided filter
(<0.1s, 600x400p)

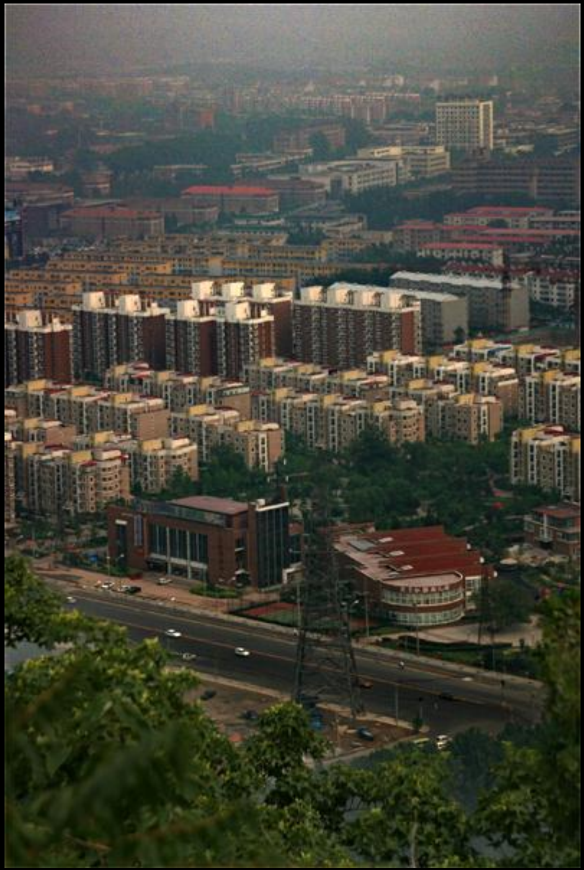

global optimization
(10s)

Limitation

- "What is an edge" – inherently ambiguous, context-dependent

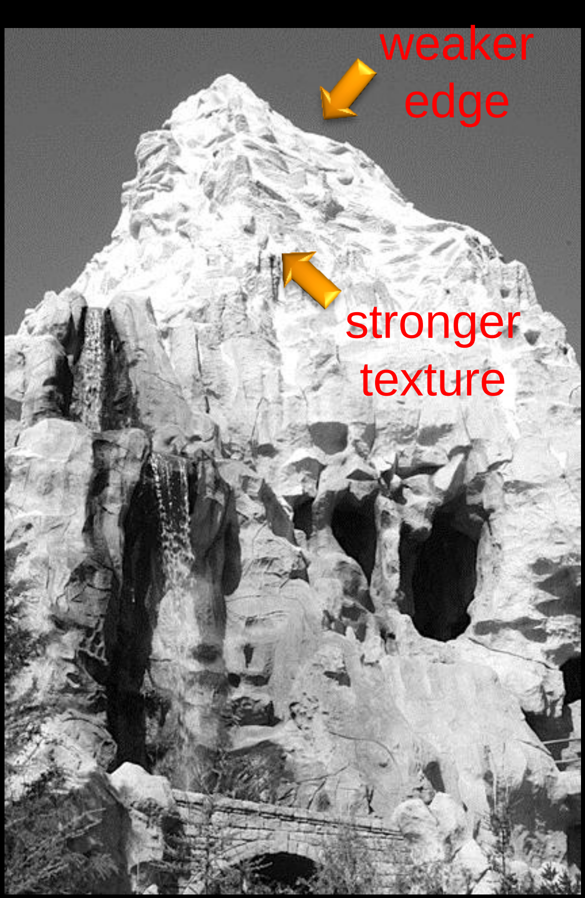

Input

Bilateral filter
$\sigma_s=16, \sigma_r=0.4$

Guided filter
$r=16, \varepsilon=0.4^2$

Conclusion

- We go from “BF” to “GF”
  - Edge-preserving filtering
  - Non-iterative
  - O(1) time, fast, accurate
  - Gradient preserving
  - More generic than “smoothing”

Thank you!