

# Bilateral Filters

Digital Visual Effects, Spring 2008

*Yung-Yu Chuang*

2008/5/27

*with slides by Fredo Durand, Ramesh Raskar, Sylvain Paris, Soonmin Bae*

# Announcements

---

- Final project proposal
- Project #3 artifacts voting

# Bilateral filtering



[Ben Weiss, Siggraph 2006]

# Image Denoising

---



noisy image



naïve denoising  
Gaussian blur



better denoising  
edge-preserving filter

Smoothing an image without blurring its edges.

# A Wide Range of Options

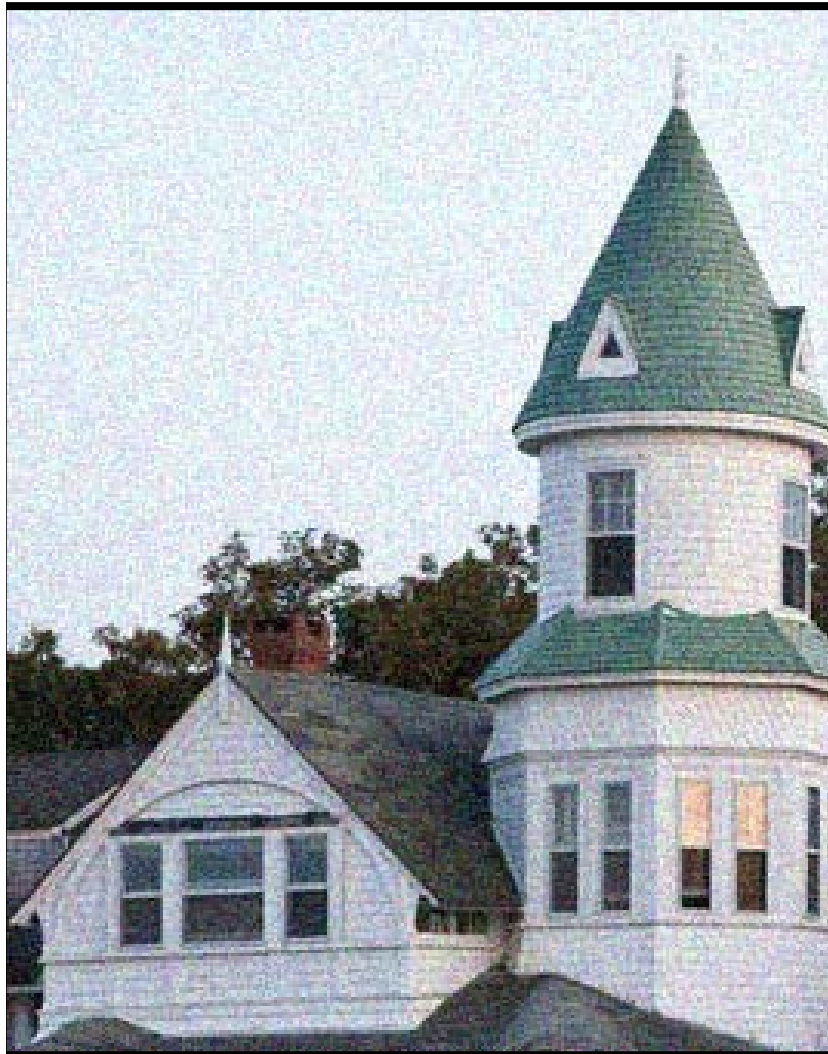
---

- Diffusion, Bayesian, Wavelets...
  - All have their pros and cons.
  
- Bilateral filter
  - not always the best result [Buades 05] but often good
  - easy to understand, adapt and set up

# Basic denoising

---

Noisy input



Median 5x5

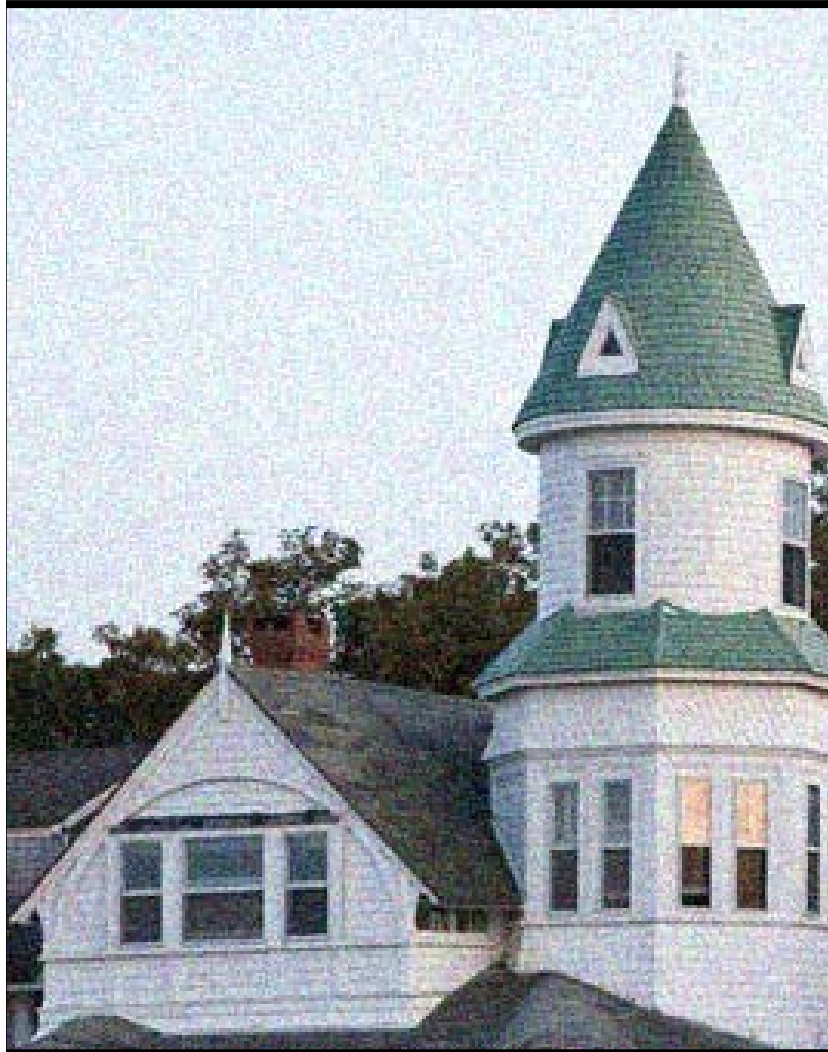


# Basic denoising

---

Noisy input

Bilateral filter 7x7 window



# Tone Mapping

[Durand 02]



HDR input



# Tone Mapping

[Durand 02]



output

# Photographic Style Transfer

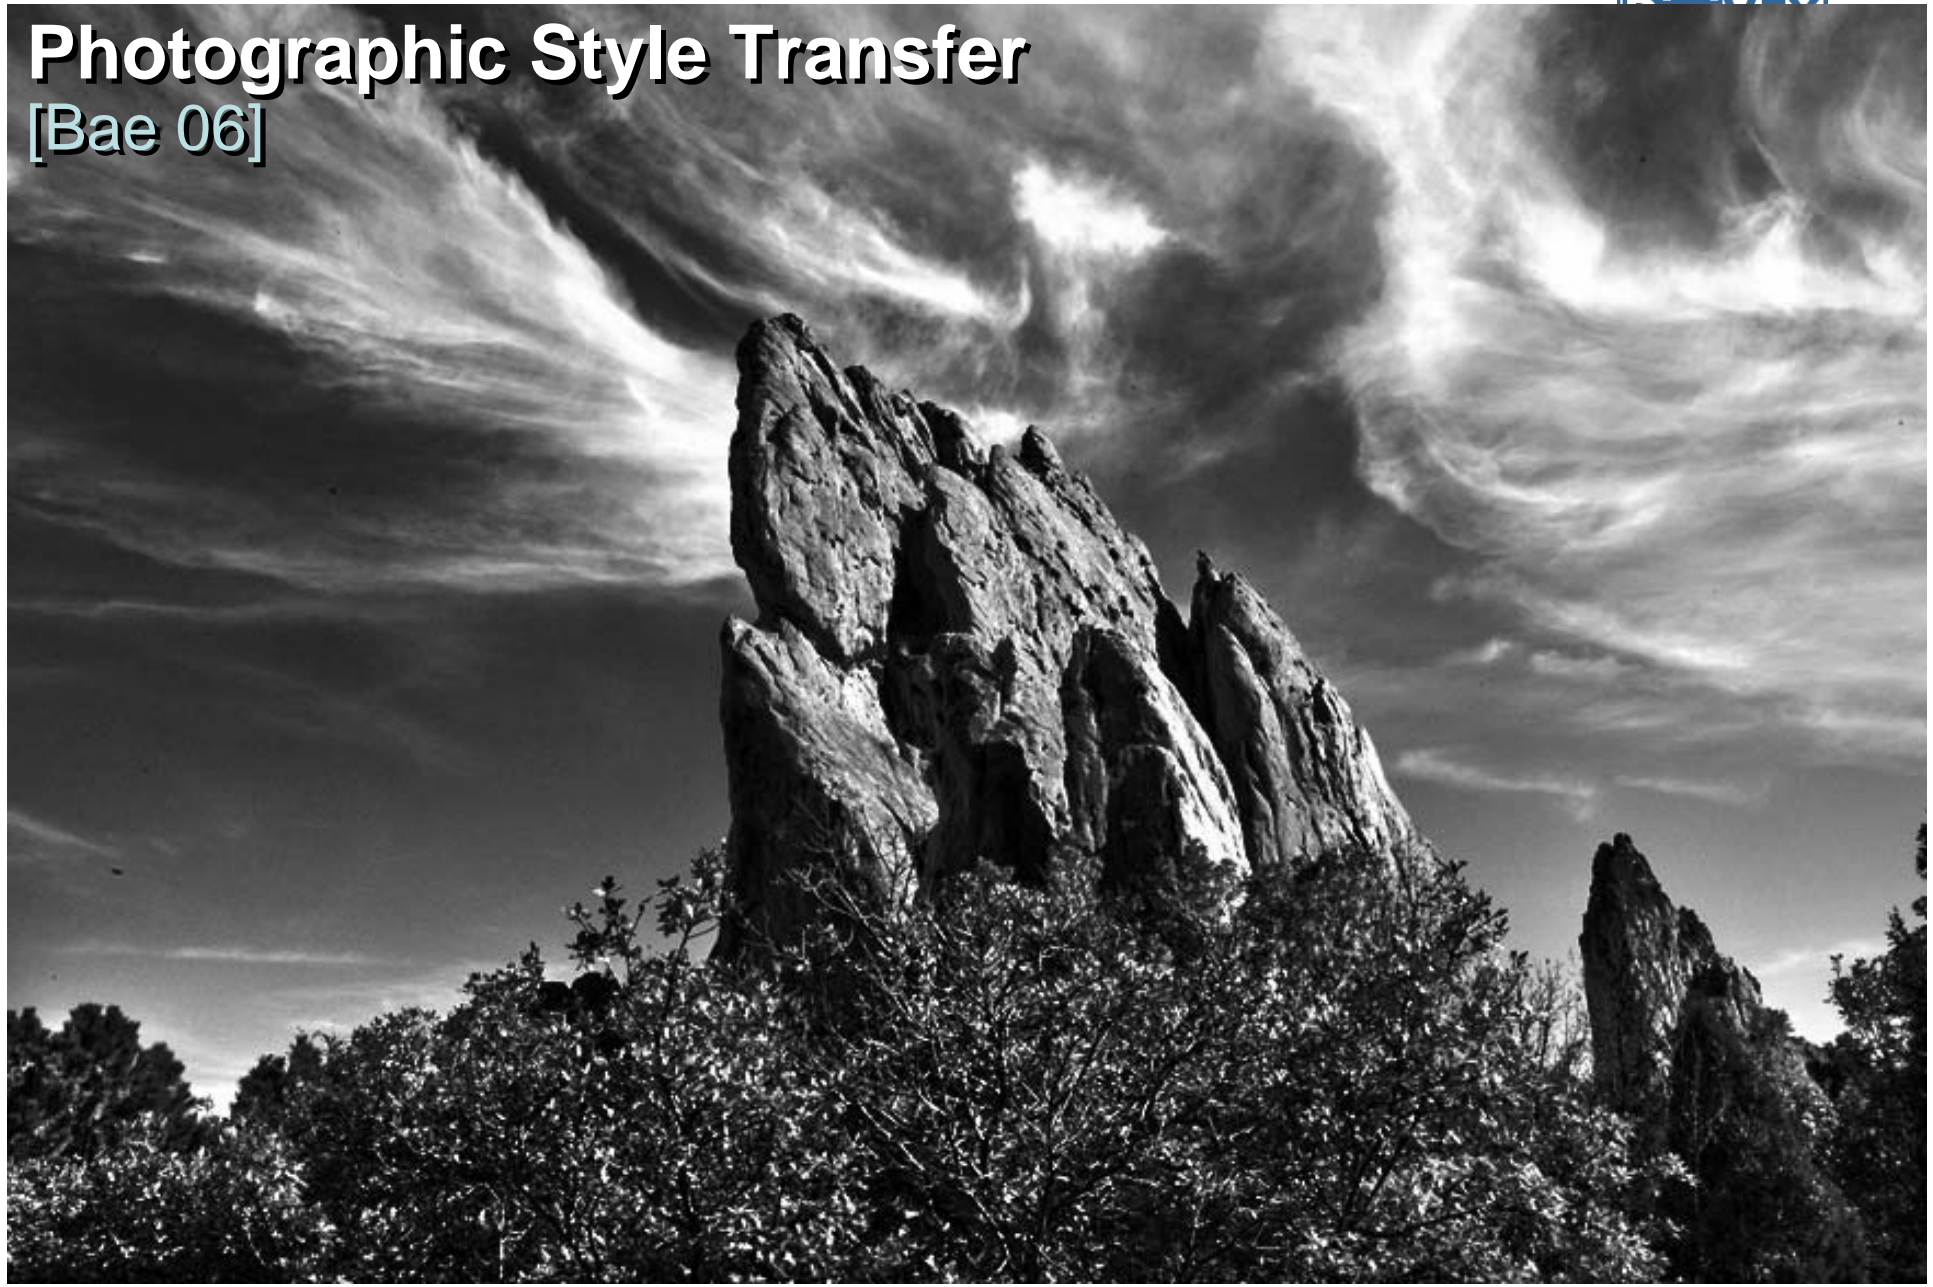
[Bae 06]



input

# Photographic Style Transfer

[Bae 06]



output

# Cartoon Rendition

[Winnemöller 06]

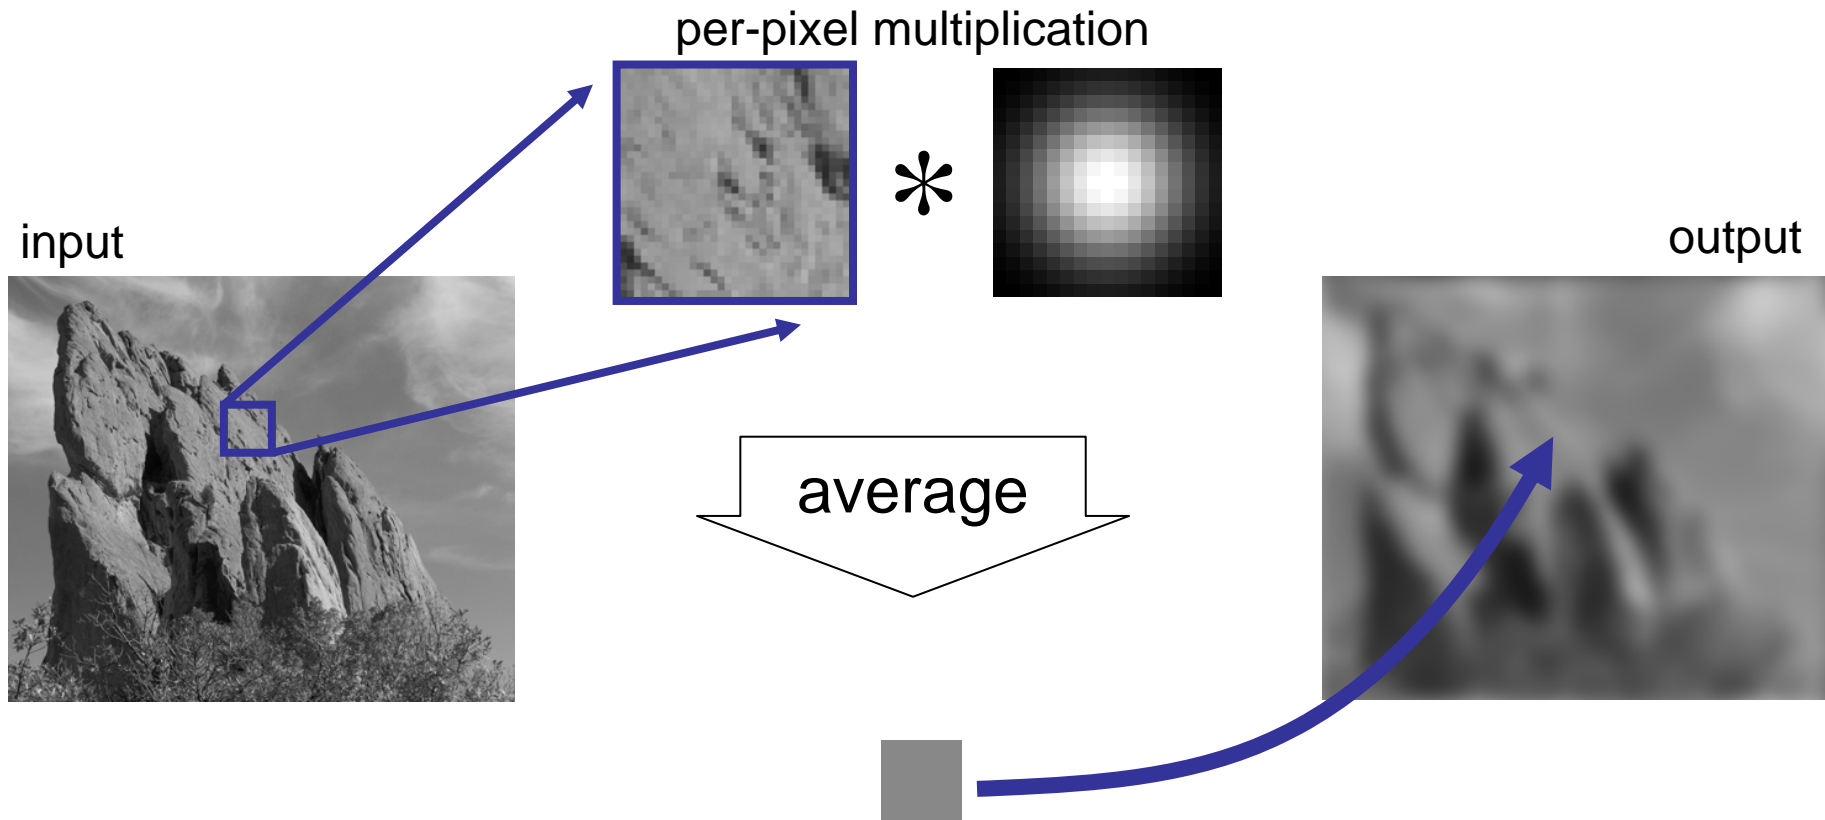


input

Cartoon Rendition  
[Koenigsmüller 06]

**6 papers at  
SIGGRAPH'07**

# Gaussian Blur



**input**



**box average**



## Gaussian blur

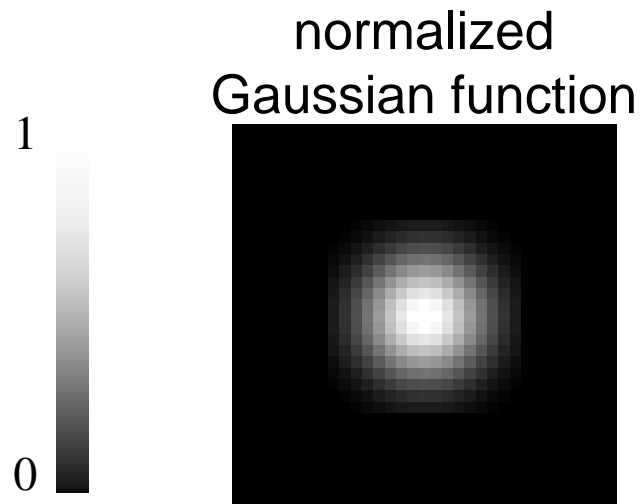


# Equation of Gaussian Blur

---

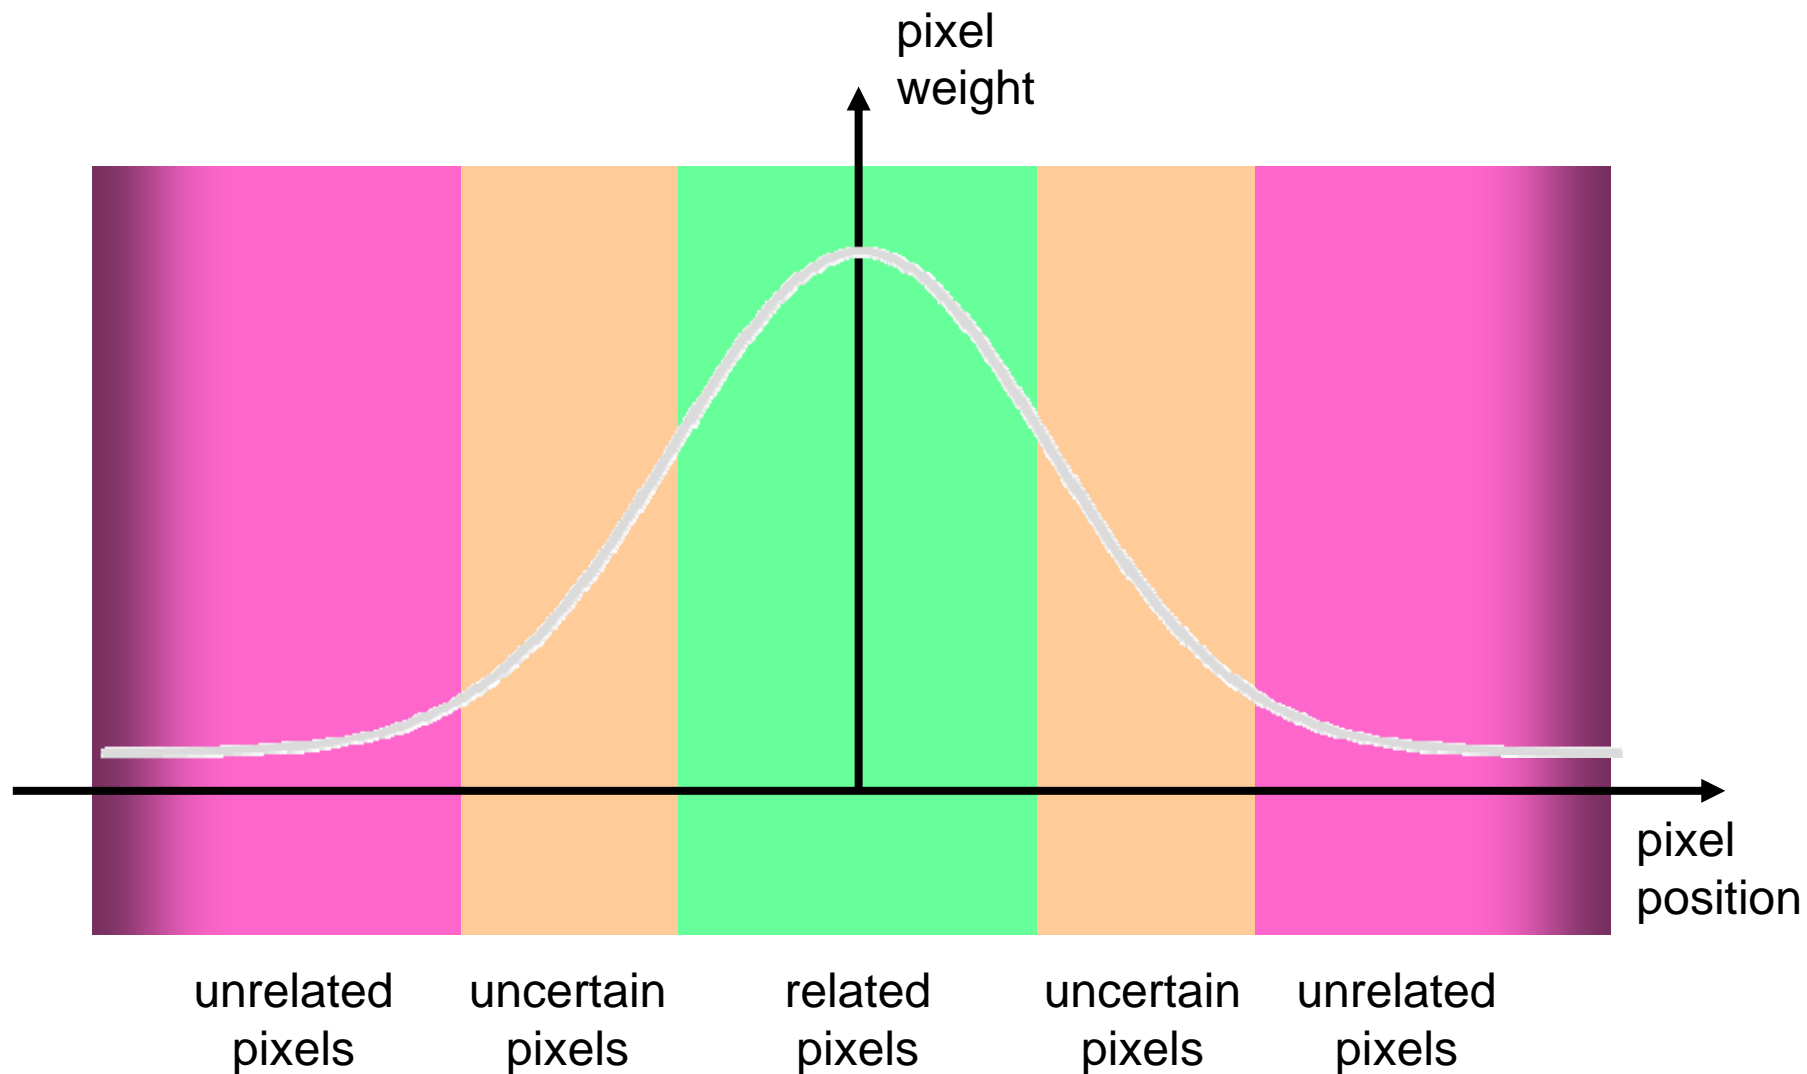
Same idea: **weighted average of pixels.**

$$GB[I]_p = \sum_{q \in S} G_{\sigma}(\|p - q\|) I_q$$



# Gaussian Profile

$$G_{\sigma}(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right)$$



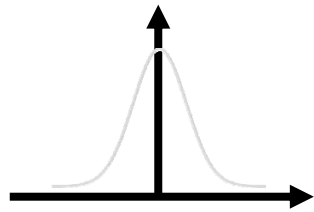
# Spatial Parameter



input

$$GB[I]_p = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma}(\|\mathbf{p} - \mathbf{q}\|) I_{\mathbf{q}}$$

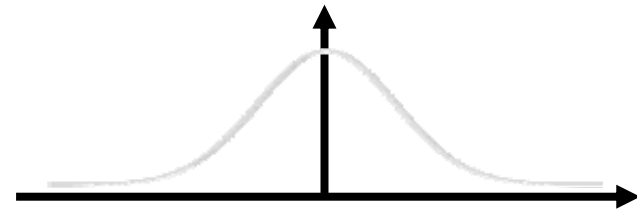
$\sigma$   
size of the window



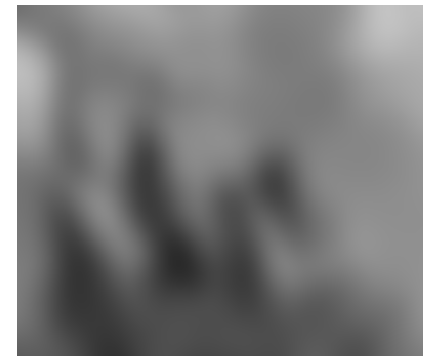
small  $\sigma$



limited smoothing



large  $\sigma$



strong smoothing

# How to set $\sigma$

---

- Depends on the application.
- Common strategy: proportional to image size
  - e.g. 2% of the image diagonal
  - property: independent of image resolution

# Properties of Gaussian Blur

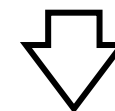
---

- Weights independent of spatial location
  - linear convolution
  - well-known operation
  - efficient computation (recursive algorithm, FFT...)

# Properties of Gaussian Blur

- Does smooth images
- But smooths too much:  
edges are **blurred**.
  - Only spatial distance matters
  - No edge term

input



output

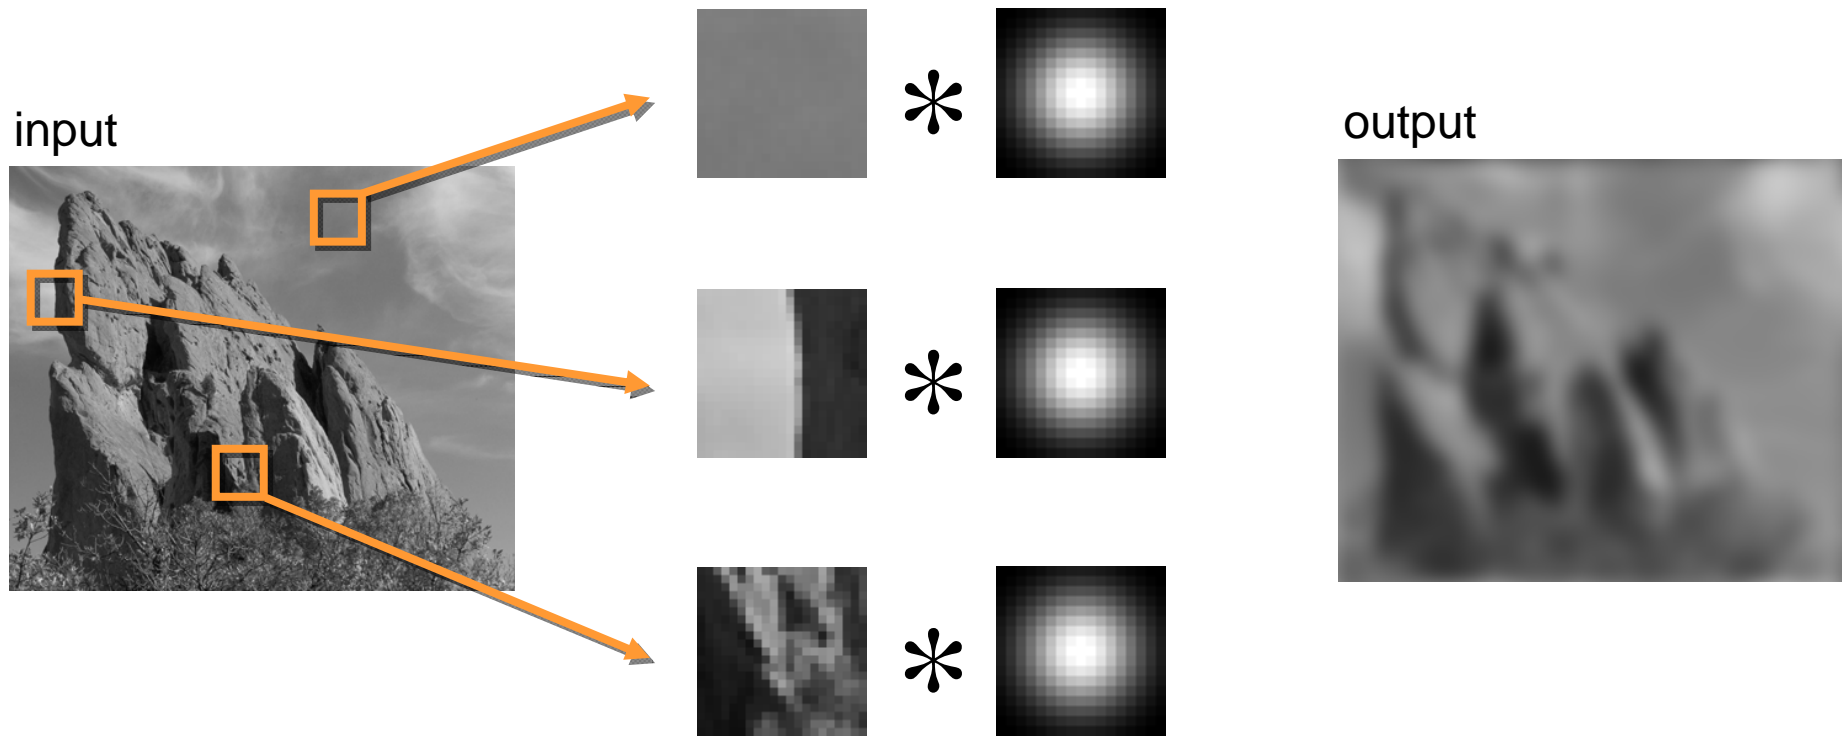


$$GB[I]_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma}(\|\mathbf{p} - \mathbf{q}\|) I_{\mathbf{q}}$$

space

# Blur Comes from Averaging across Edges DigiVFX

---



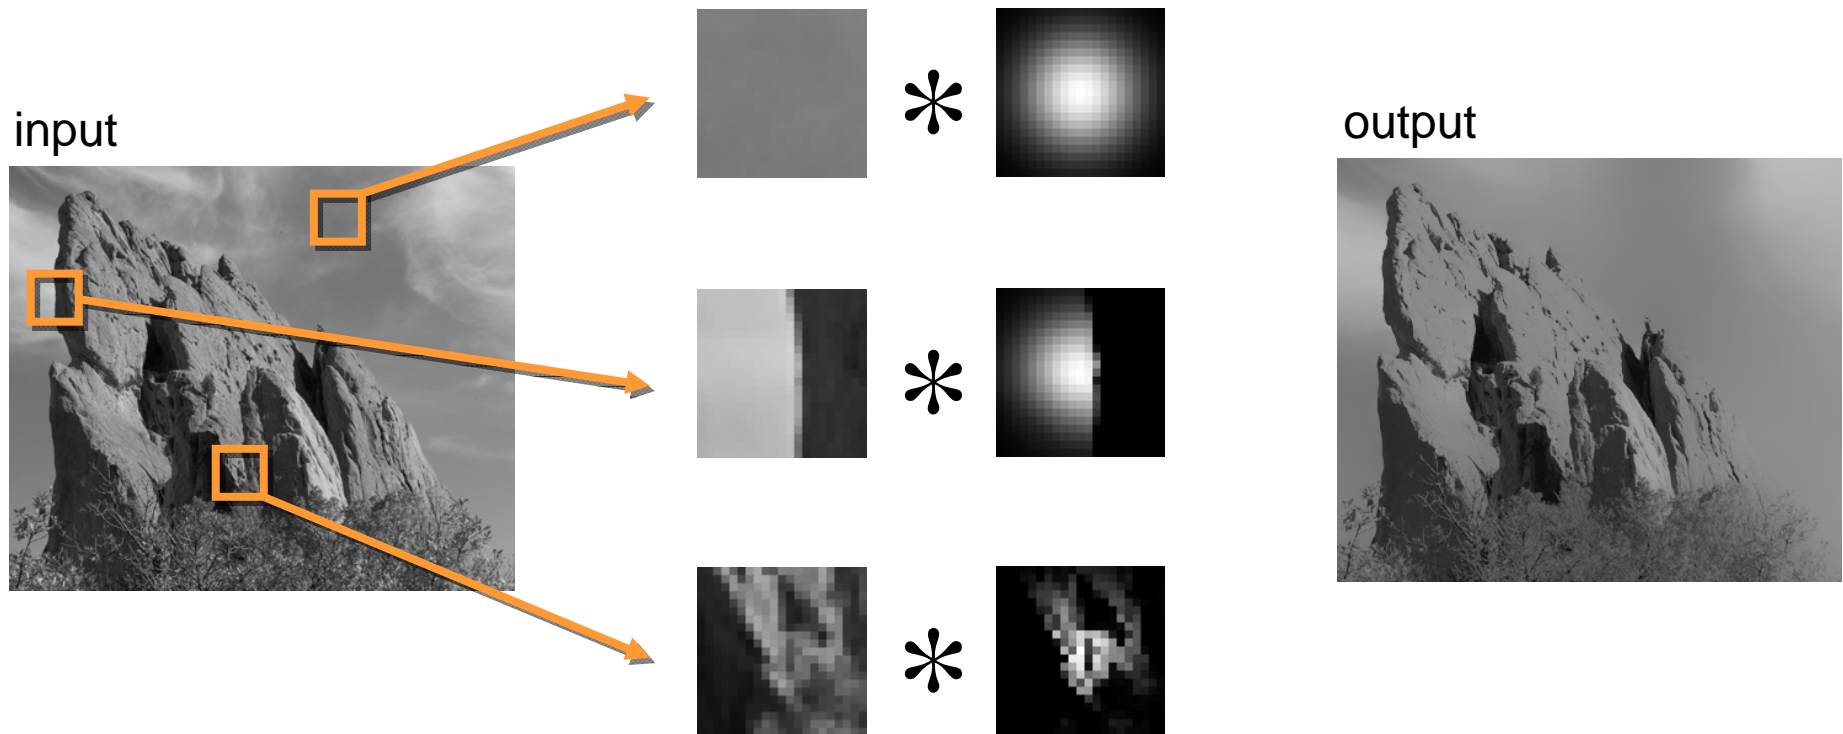
Same Gaussian kernel everywhere.



# Bilateral Filter No Averaging across Edges

[Aurich 95, Smith 97, Tomasi 98]

---



The kernel shape depends on the image content.

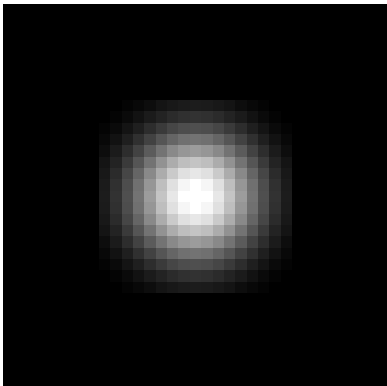
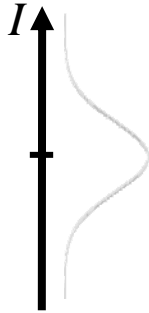
# Bilateral Filter Definition

Same idea: **weighted average of pixels.**

$$BF [I]_p = \frac{1}{W_p} \sum_{q \in S} G_{\sigma_s} (\| \mathbf{p} - \mathbf{q} \|) G_{\sigma_r} (| I_p - I_q |) I_q$$

new  
not new  
new

↓ normalization factor  
↓ **space** weight  
↓ **range** weight

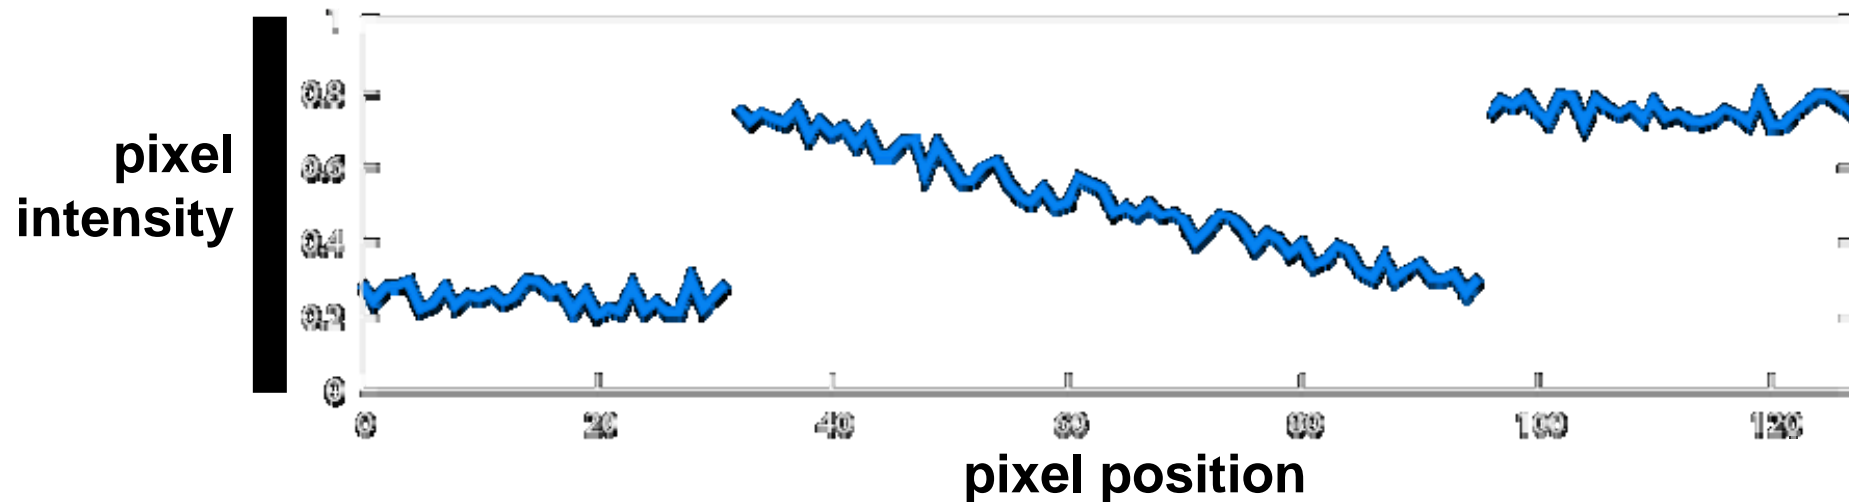
# Illustration a 1D Image

---

- 1D image = line of pixels

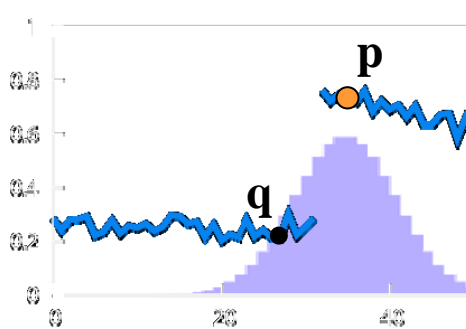


- Better visualized as a plot



# Gaussian Blur and Bilateral Filter DigiVFX

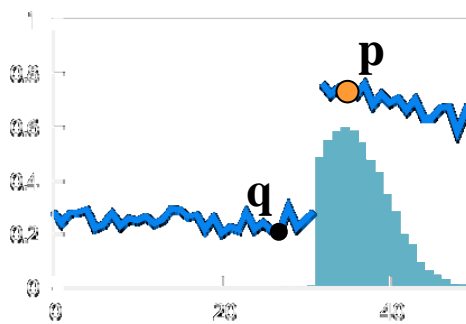
## Gaussian blur



← space →

## Bilateral filter

[Aurich 95, Smith 97, Tomasi 98]

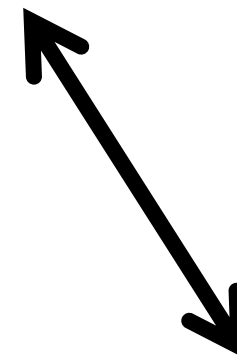
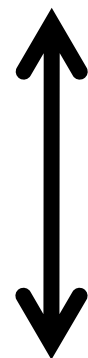


← space →

↑ range ↓

$$GB[I]_p = \sum_{q \in S} G_{\sigma}(\|p - q\|) I_q$$

space



$$BF[I]_p = \frac{1}{W_p} \sum_{q \in S} G_{\sigma_s}(\|p - q\|) G_{\sigma_r}(|I_p - I_q|) I_q$$

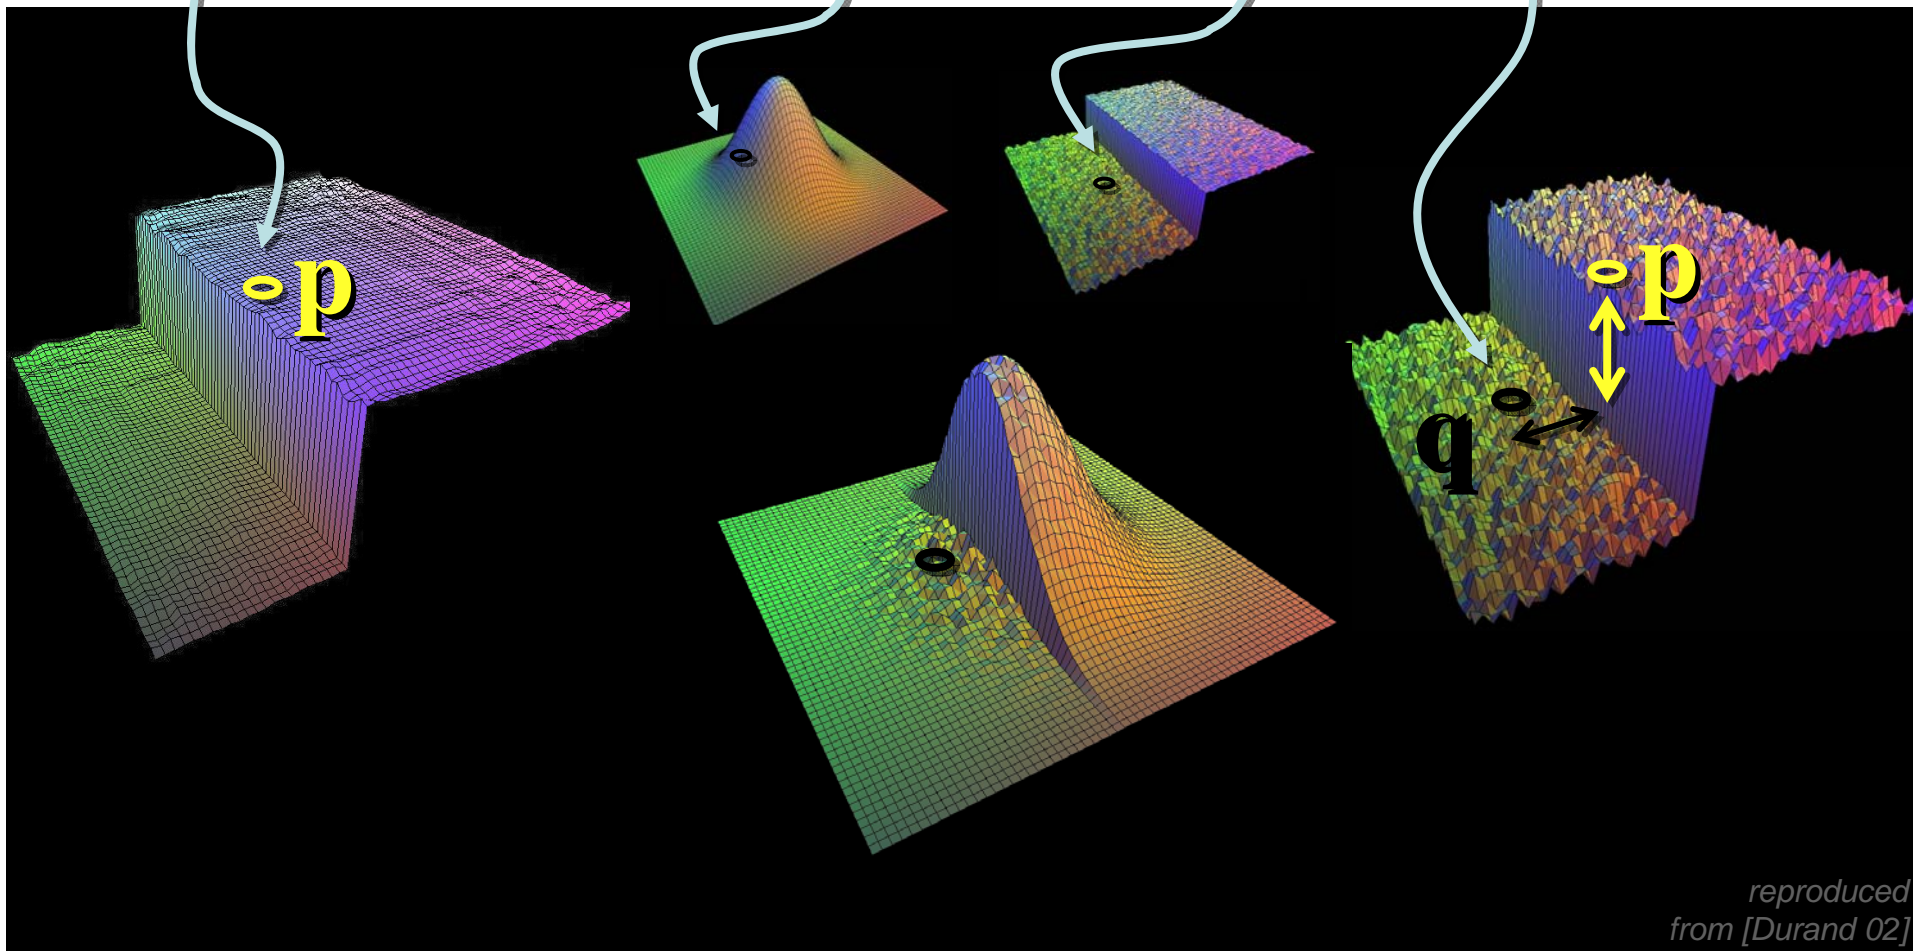
normalization

space

range


# Bilateral Filter on a Height Field

$$BF [I]_p = \frac{1}{W_p} \sum_{q \in S} \underbrace{G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|)}_{\text{spatial}} \underbrace{G_{\sigma_r}(\|I_p - I_q\|)}_{\text{range}} I_q$$



# Space and Range Parameters

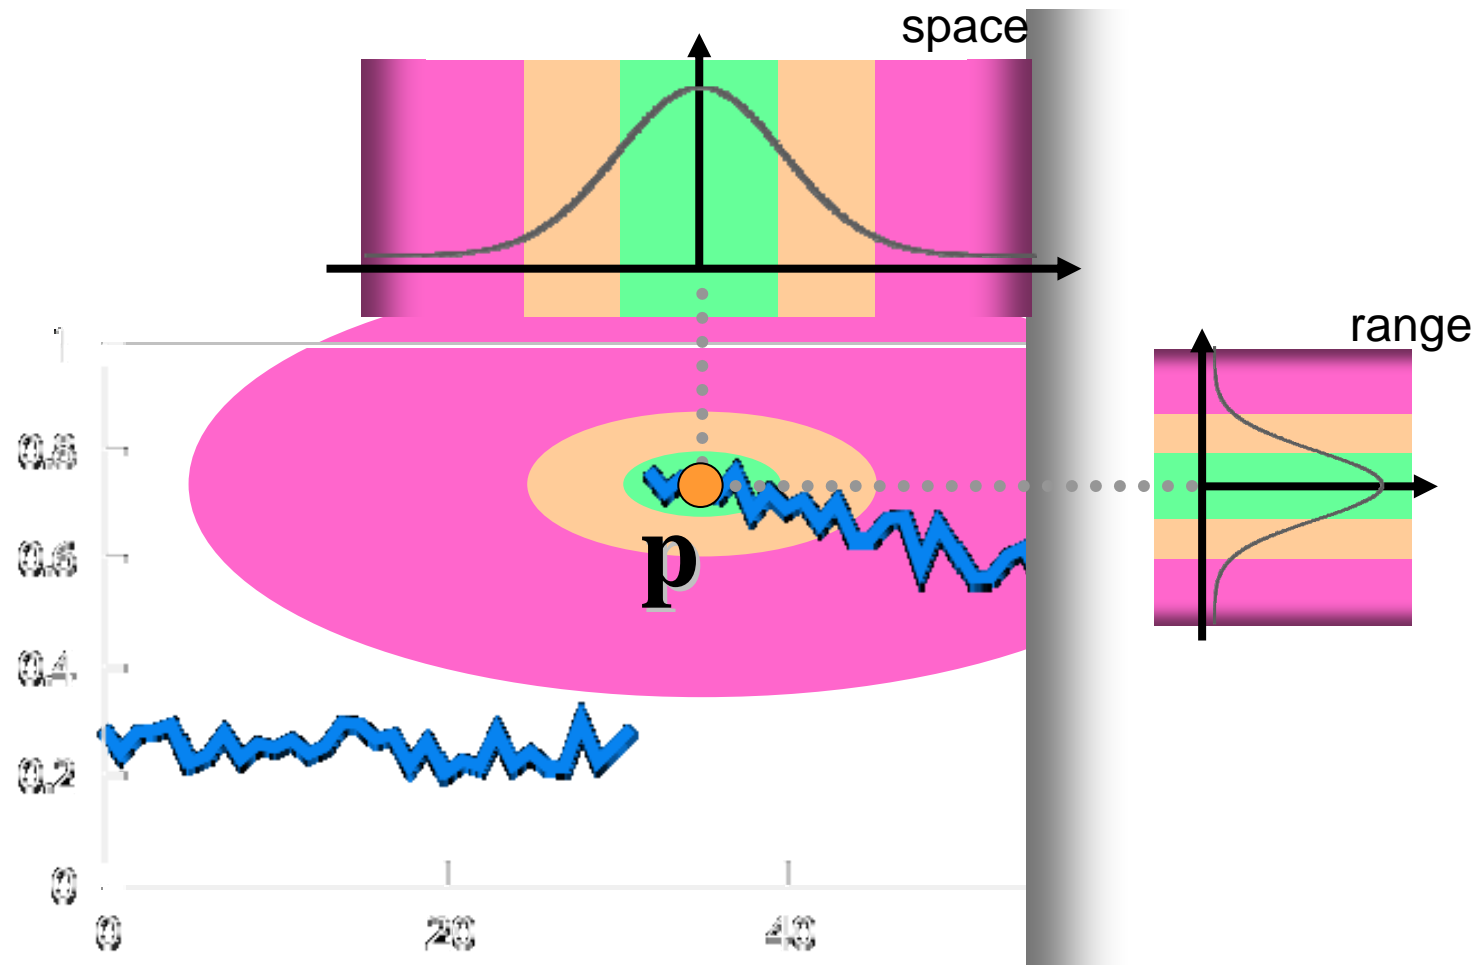
---

$$BF [I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(\|I_{\mathbf{p}} - I_{\mathbf{q}}\|) I_{\mathbf{q}}$$


- space  $\sigma_s$  : spatial extent of the kernel, size of the considered neighborhood.
- range  $\sigma_r$  : “minimum” amplitude of an edge

# Influence of Pixels

Only pixels close in space and in range are considered.



# Exploring the Parameter Space



input

$\sigma_s = 2$

$\sigma_r = 0.1$



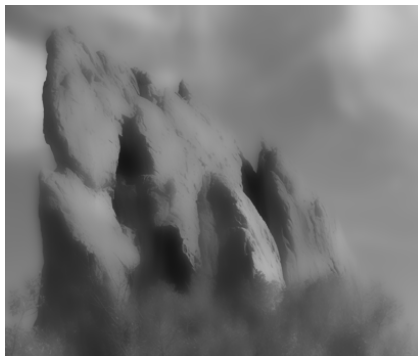
$\sigma_r = 0.25$



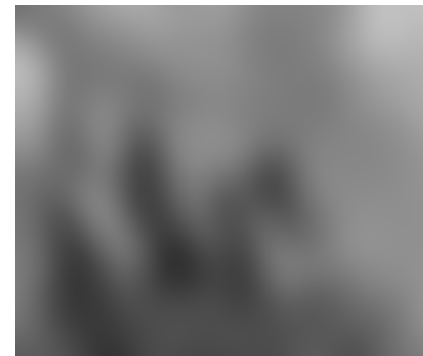
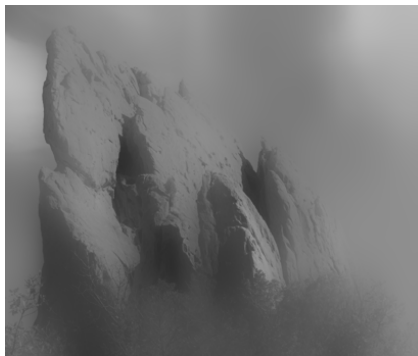
$\sigma_r = \infty$   
(Gaussian blur)



$\sigma_s = 6$



$\sigma_s = 18$





# Varying the Range Parameter



input

$\sigma_r = 0.1$

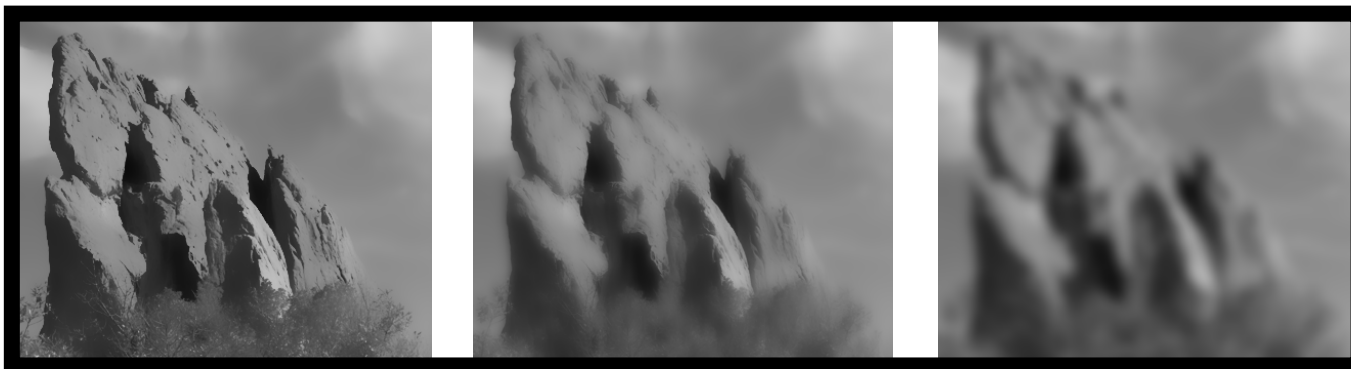
$\sigma_r = 0.25$

$\sigma_r = \infty$   
(Gaussian blur)

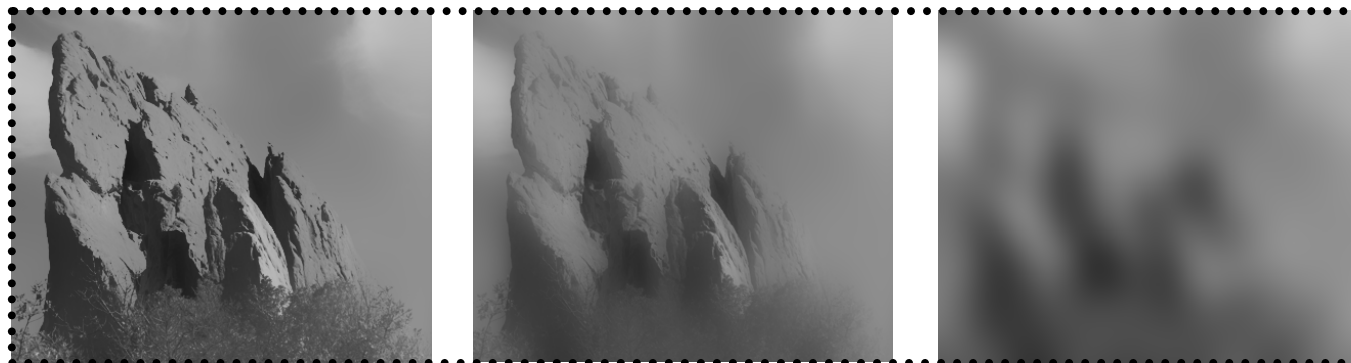
$\sigma_s = 2$



$\sigma_s = 6$



$\sigma_s = 18$



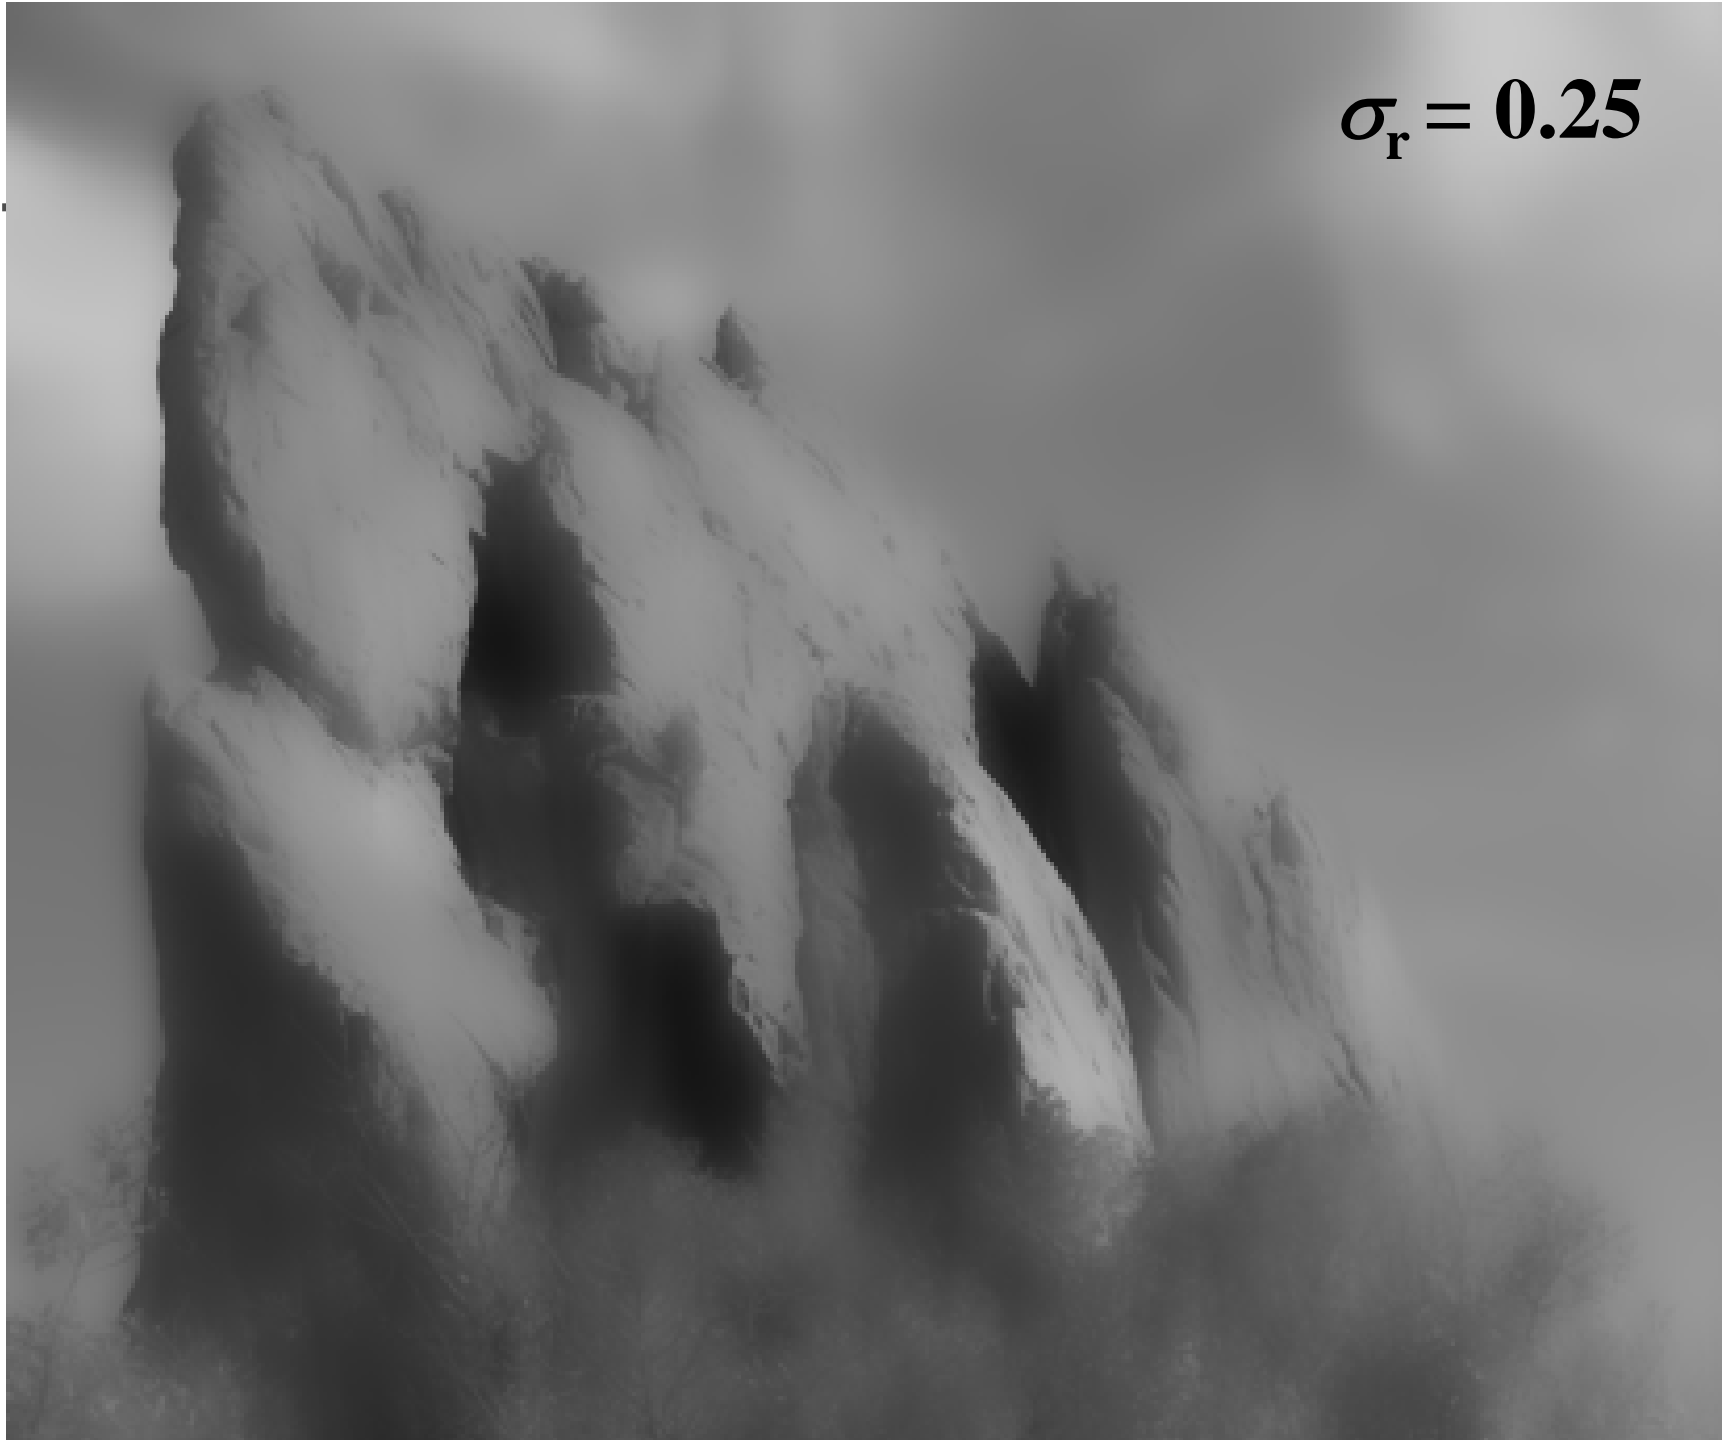
**input**



$$\sigma_r = 0.1$$



$$\sigma_r = 0.25$$



$\sigma_r = \infty$   
(Gaussian blur)



# Varying the Space Parameter



input

$\sigma_s = 2$

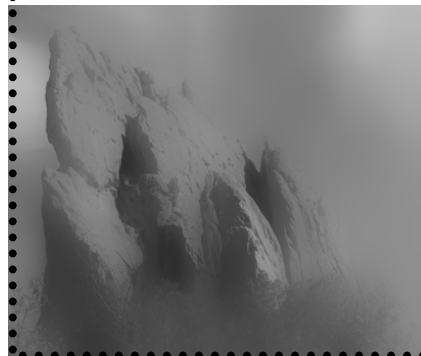
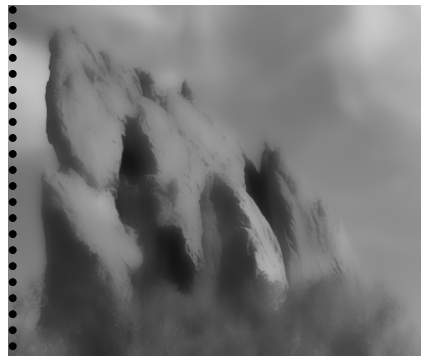
$\sigma_s = 6$

$\sigma_s = 18$

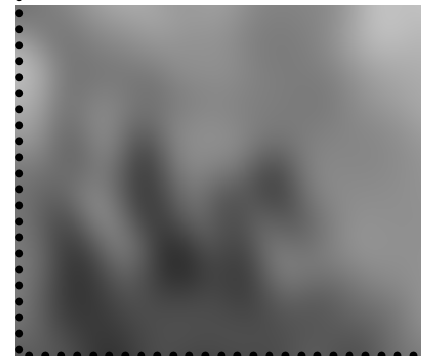
$\sigma_r = 0.1$



$\sigma_r = 0.25$



$\sigma_r = \infty$   
(Gaussian blur)



**input**



$$\sigma_s = 2$$





$$\sigma_s = 6$$



$$\sigma_s = 18$$



# How to Set the Parameters

---

Depends on the application. For instance:

- space parameter: proportional to image size
  - e.g., 2% of image diagonal
- range parameter: proportional to edge amplitude
  - e.g., mean or median of image gradients
- independent of resolution and exposure

# Iterating the Bilateral Filter

---

$$I_{(n+1)} = BF[I_{(n)}]$$

- Generate more piecewise-flat images
- Often not needed in computational photo, but could be useful for applications such as NPR.

**input**



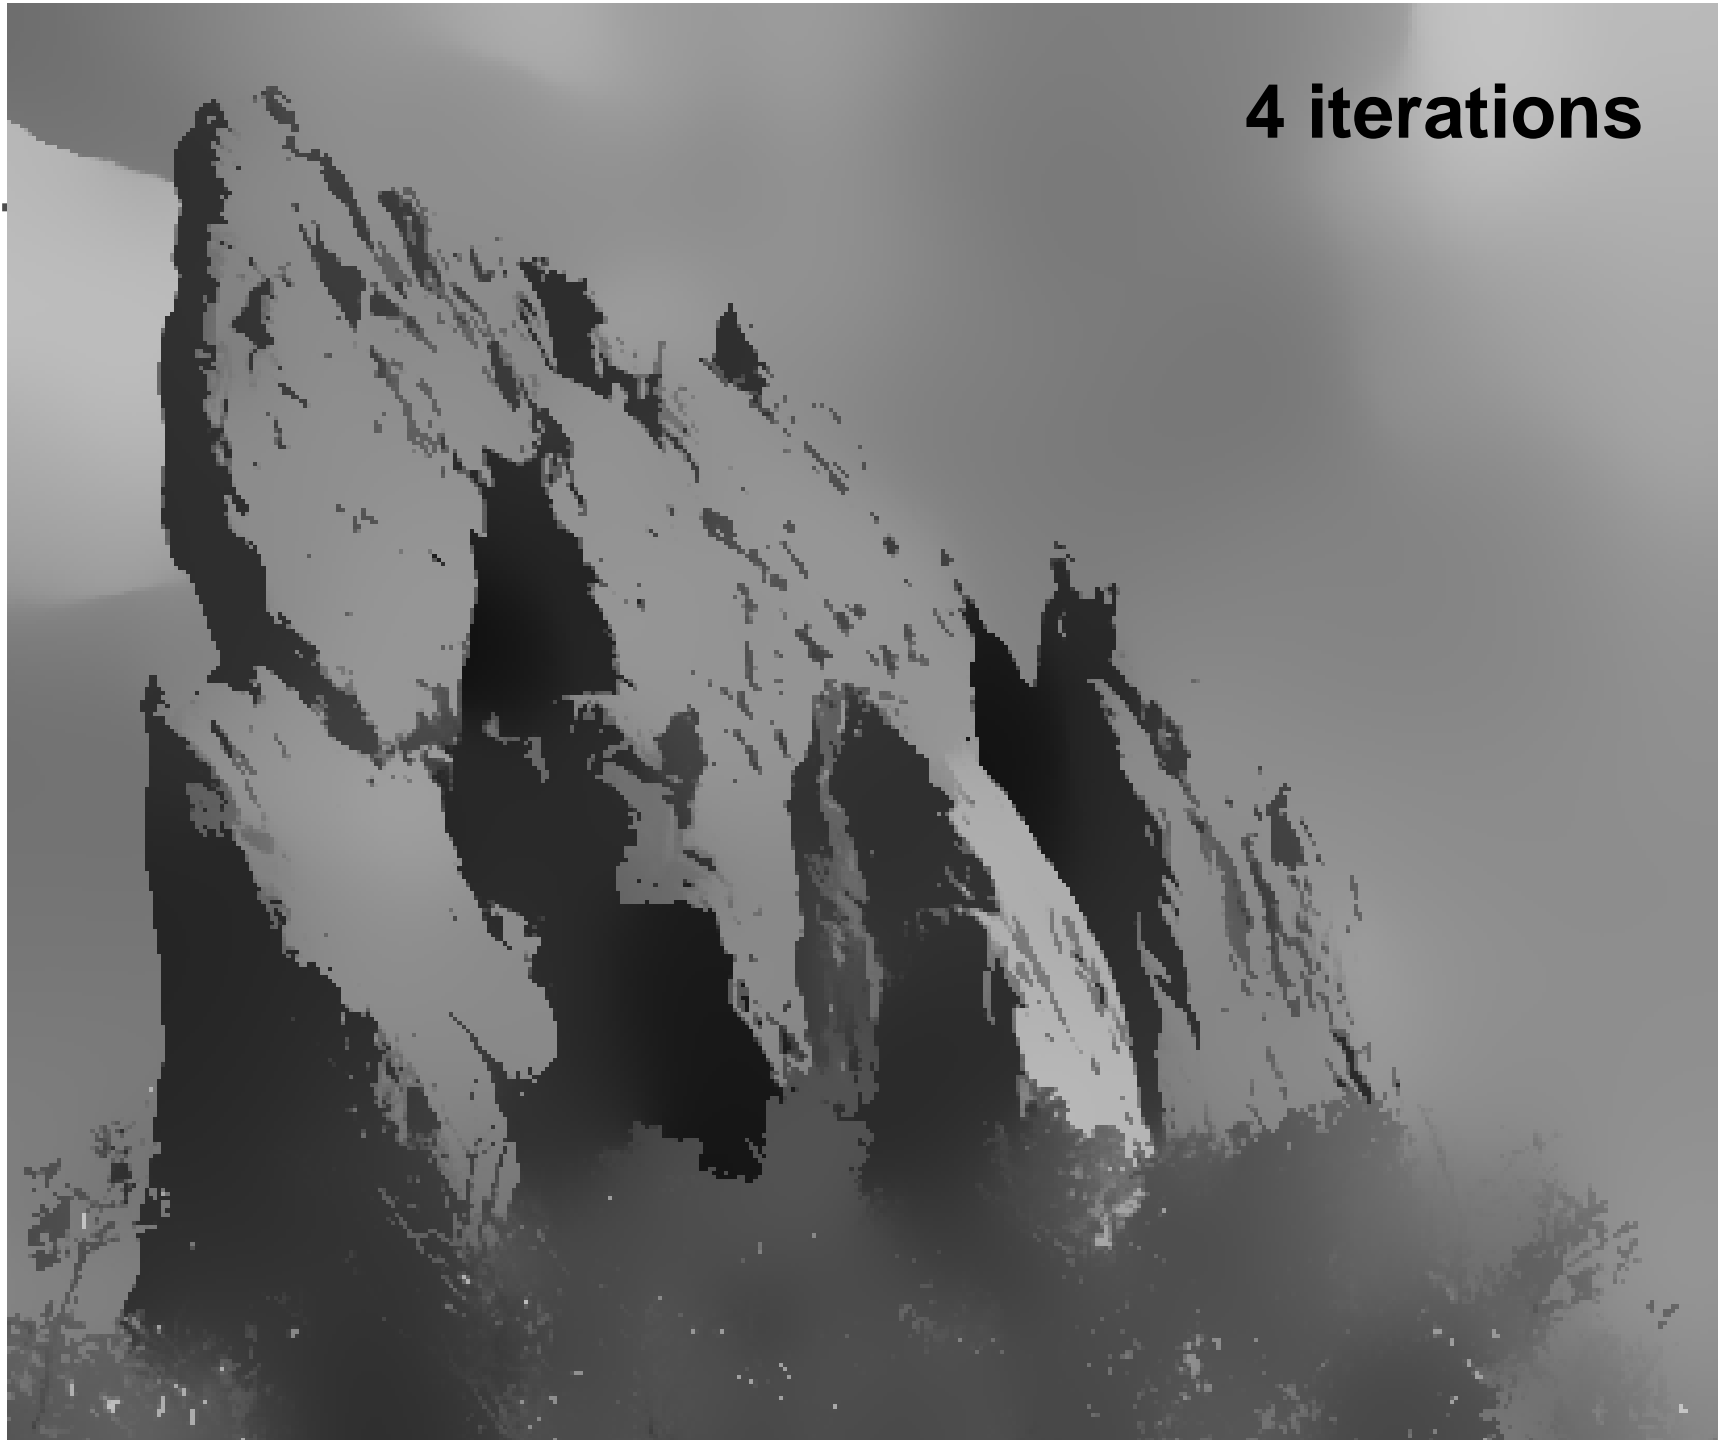
**1 iteration**



**2 iterations**



**4 iterations**





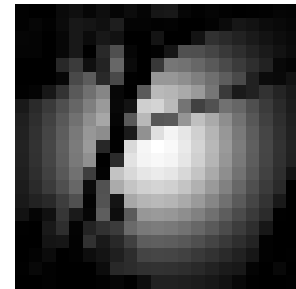
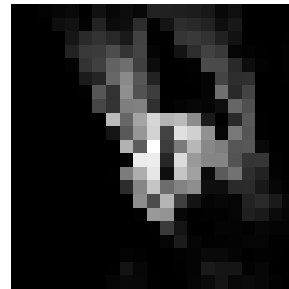
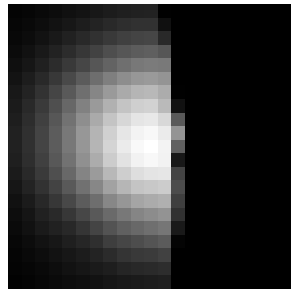
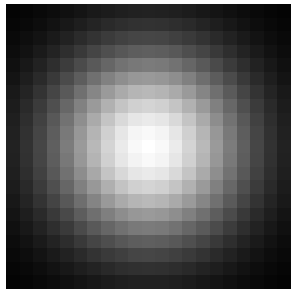
# Advantages of Bilateral Filter

---

- Easy to understand
  - Weighted mean of nearby pixels
- Easy to adapt
  - Distance between pixel values
- Easy to set up
  - Non-iterative

# Hard to Compute

- Nonlinear 
$$BF [I]_p = \frac{1}{W_p} \sum_{q \in S} G_{\sigma_s} (\| \mathbf{p} - \mathbf{q} \|) G_{\sigma_r} (| I_p - I_q |) I_q$$
- Complex, spatially varying kernels
  - Cannot be precomputed, no FFT...



- Brute-force implementation is slow > 10min

# But Bilateral Filter is Nonlinear

---

- Slow but some accelerations exist:
  - [Elad 02]: Gauss-Seidel iterations
    - Only for many iterations
  - [Durand 02, Weiss 06]: fast approximation
    - No formal understanding of accuracy versus speed
    - [Weiss 06]: Only box function as spatial kernel

# A Fast Approximation of the Bilateral Filter using a Signal Processing Approach

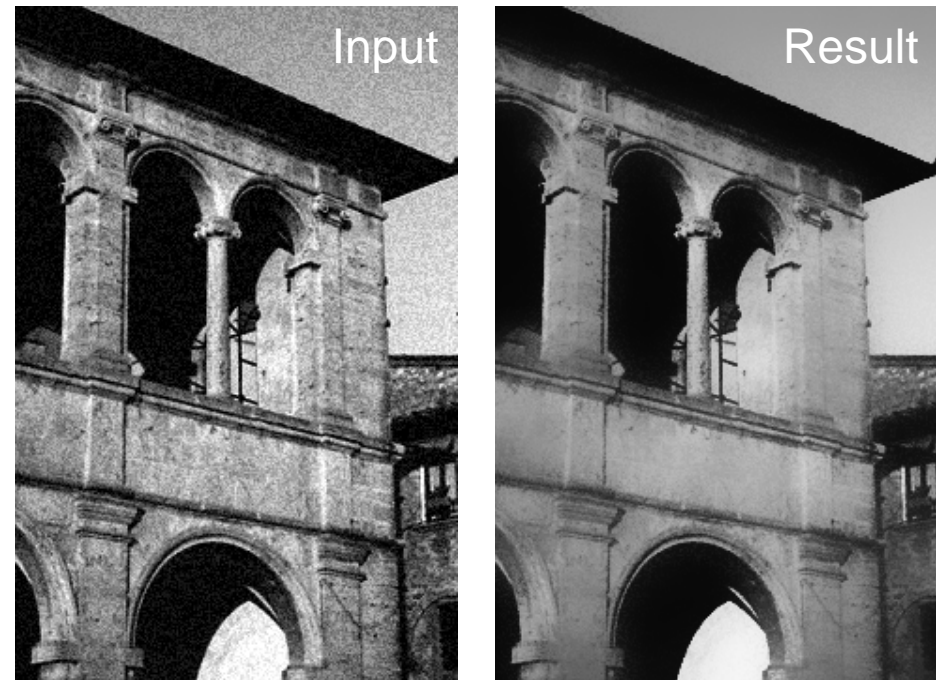
Sylvain Paris and Frédo Durand

Computer Science and Artificial Intelligence Laboratory  
Massachusetts Institute of Technology



# Definition of Bilateral Filter

- [Smith 97, Tomasi 98]
- Smooths an image and preserves edges
- Weighted average of neighbors
- Weights
  - Gaussian on *space* distance
  - Gaussian on *range* distance
  - sum to 1



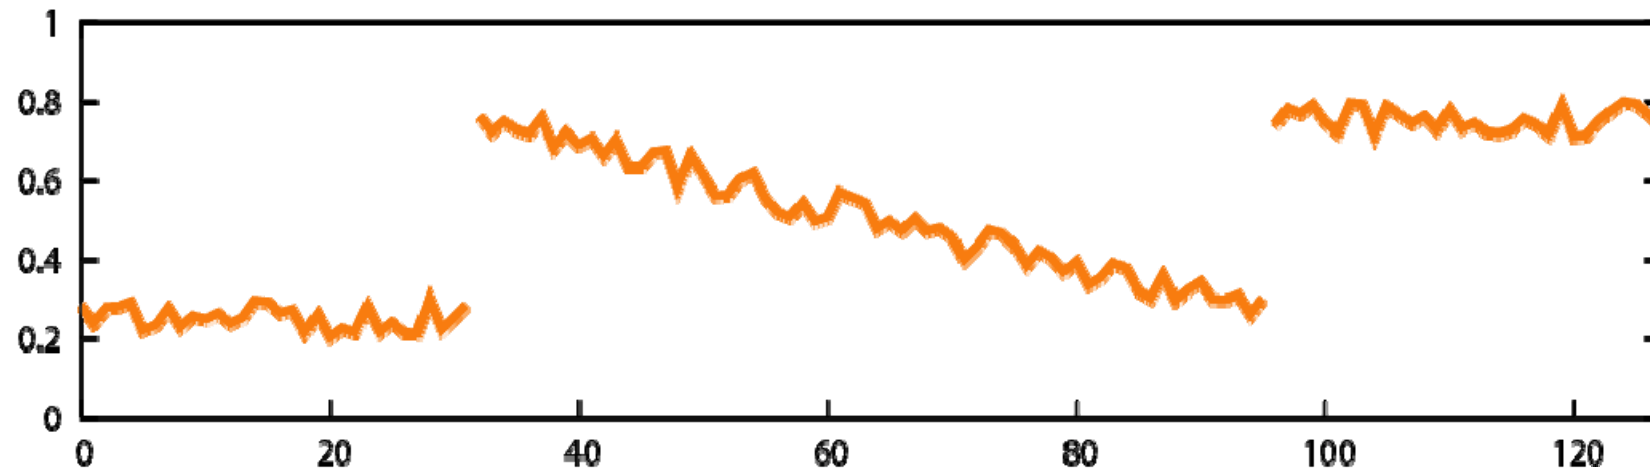
$$I_{\mathbf{p}}^{\text{bf}} = \frac{1}{W_{\mathbf{p}}^{\text{bf}}} \sum_{\mathbf{q} \in \mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|)}_{\text{space}} \underbrace{G_{\sigma_r}(\|I_{\mathbf{p}} - I_{\mathbf{q}}\|)}_{\text{range}} I_{\mathbf{q}}$$

# Contributions

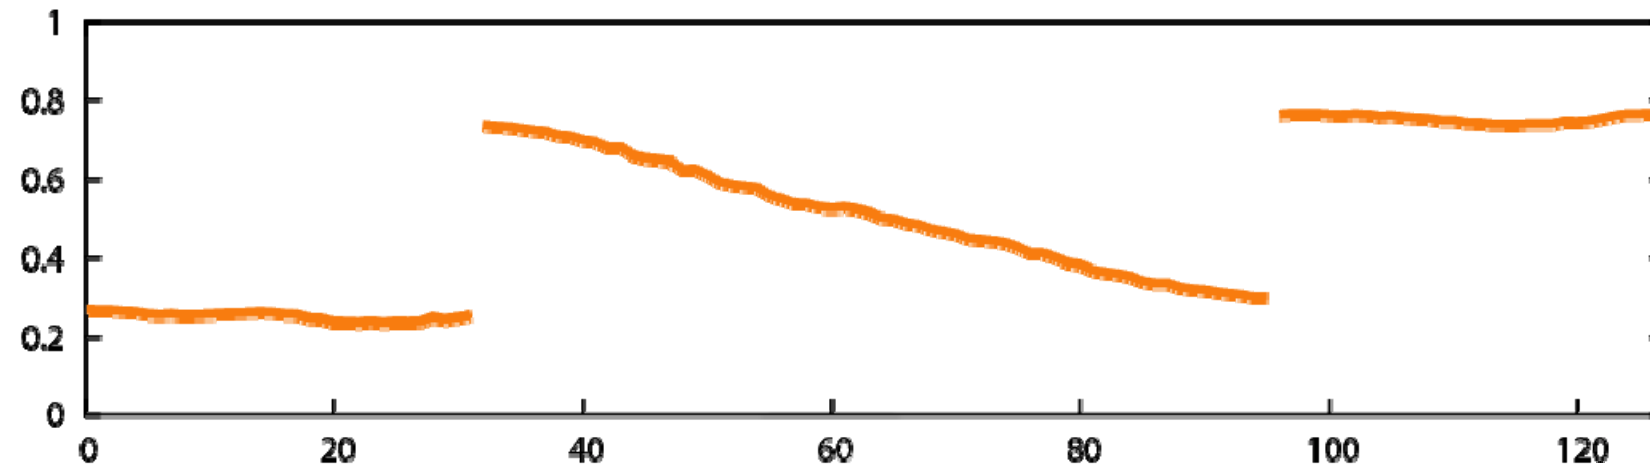
---

- Link with **linear filtering**
- **Fast** and **accurate** approximation

# Intuition on 1D Signal



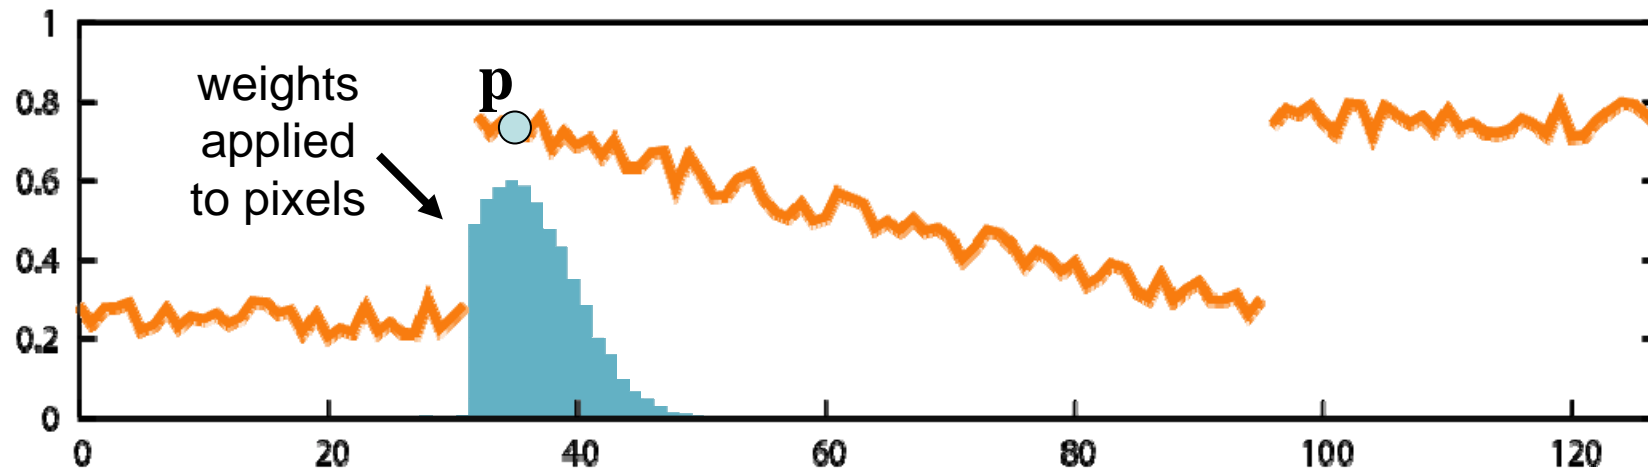
BF



# Intuition on 1D Signal

## Weighted Average of Neighbors

---

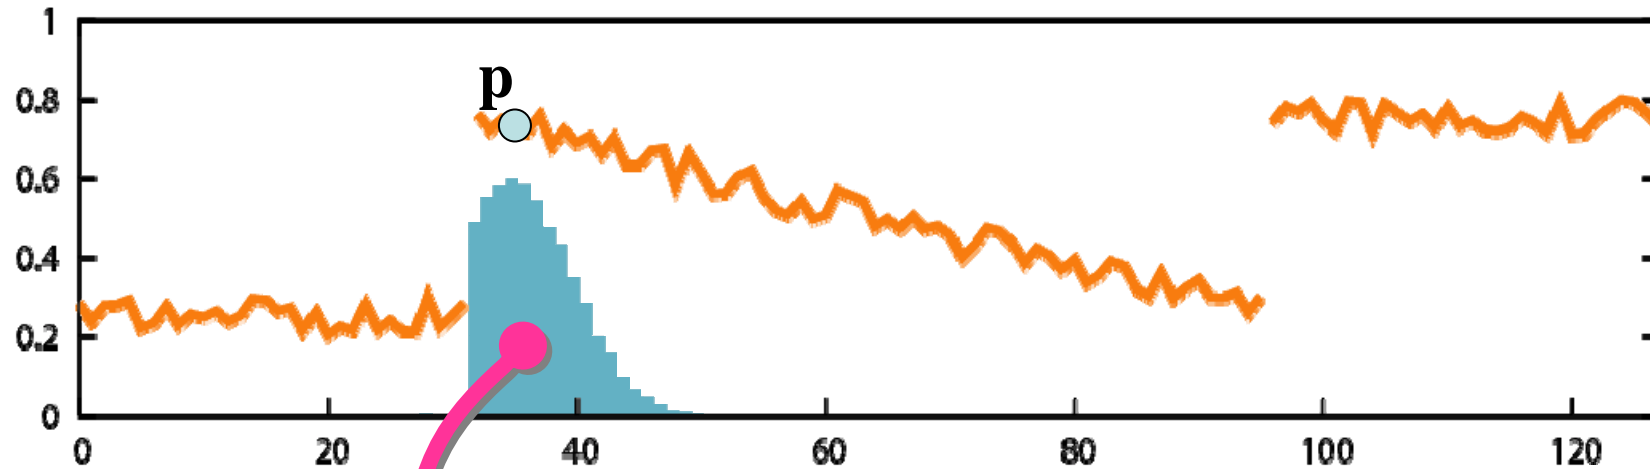


- Near and similar pixels have influence.
- Far pixels have no influence.
- Pixels with different value have no influence.



# Link with Linear Filtering

## 1. Handling the Division



sum of weights

$$I_{\mathbf{p}}^{\text{bf}} = \frac{1}{W_{\mathbf{p}}^{\text{bf}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) I_{\mathbf{q}}$$

Handling the division with a **projective space**.

# Formalization: Handling the Division

$$I_{\mathbf{p}}^{\text{bf}} = \frac{1}{W_{\mathbf{p}}^{\text{bf}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) I_{\mathbf{q}}$$
$$W_{\mathbf{p}}^{\text{bf}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$$

- Normalizing factor as homogeneous coordinate
  - Multiply both sides by  $W_{\mathbf{p}}^{\text{bf}}$

$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} & I_{\mathbf{p}}^{\text{bf}} \\ & W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) \begin{pmatrix} I_{\mathbf{q}} \\ 1 \end{pmatrix}$$

# Formalization: Handling the Division

$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} & I_{\mathbf{p}}^{\text{bf}} \\ & W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(\|I_{\mathbf{p}} - I_{\mathbf{q}}\|) \begin{pmatrix} W_{\mathbf{q}} & I_{\mathbf{q}} \\ & W_{\mathbf{q}} \end{pmatrix} \text{ with } W_{\mathbf{q}}=1$$

- Similar to homogeneous coordinates in projective space
- Division delayed until the end
- Next step: Adding a dimension to make a convolution appear

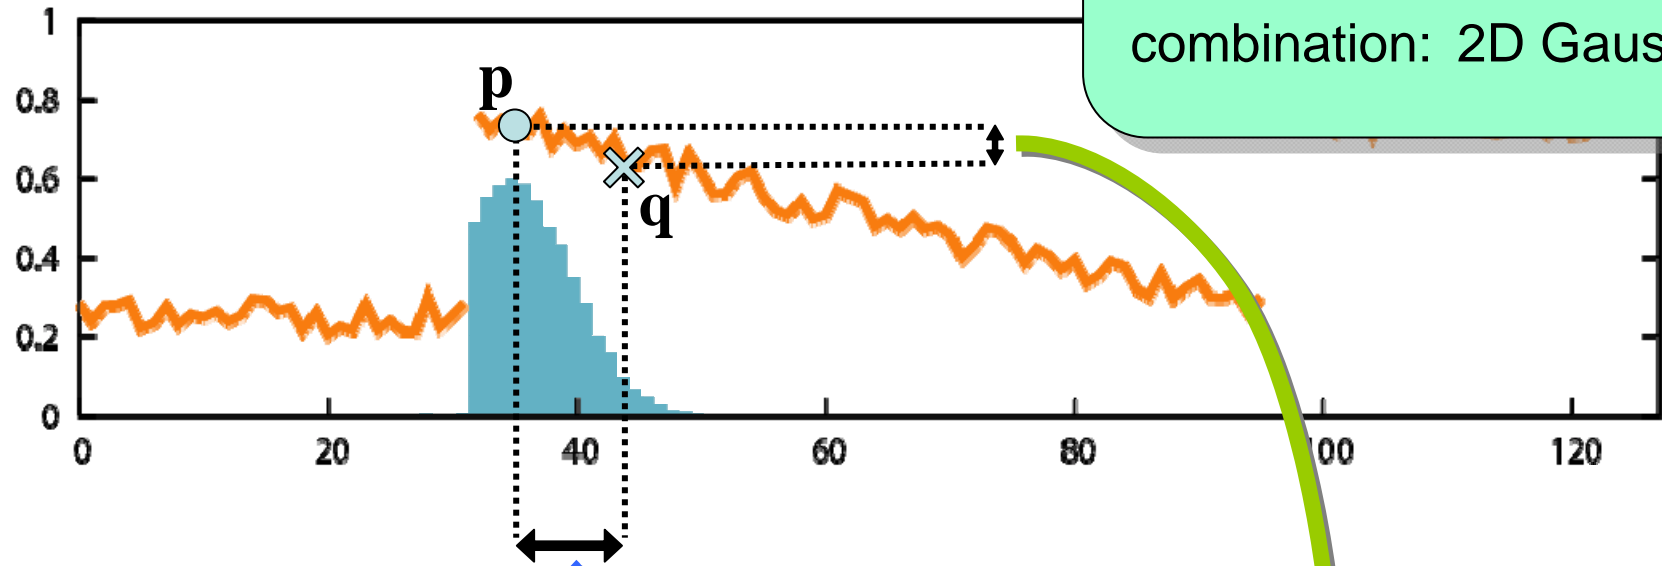
# Link with Linear Filtering

## 2. Introducing a Convolution

space: 1D Gaussian  
 × range: 1D Gaussian  


---

 combination: 2D Gaussian



$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} & I_{\mathbf{p}}^{\text{bf}} \\ W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{\mathbf{q} \in \mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|)}_{\text{space}} \underbrace{G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)}_{\text{range}} \begin{pmatrix} W_{\mathbf{q}} & I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

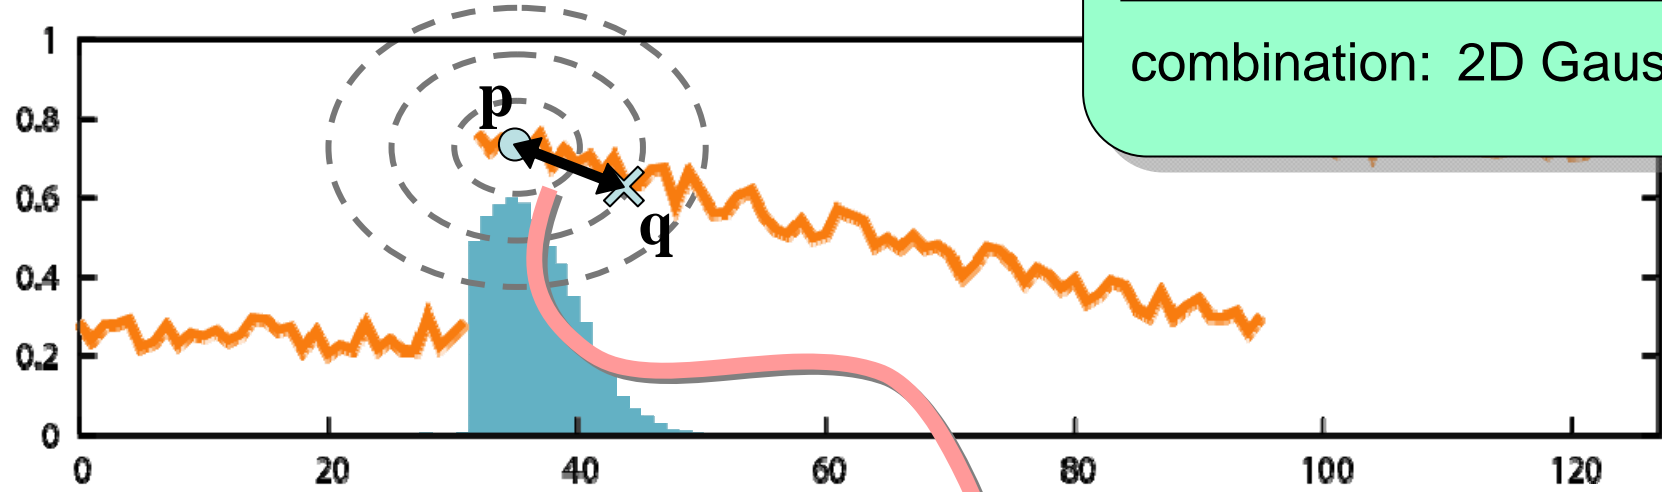
# Link with Linear Filtering

## 2. Introducing a Convolution

space: 1D Gaussian  
 × range: 1D Gaussian

---

combination: 2D Gaussian

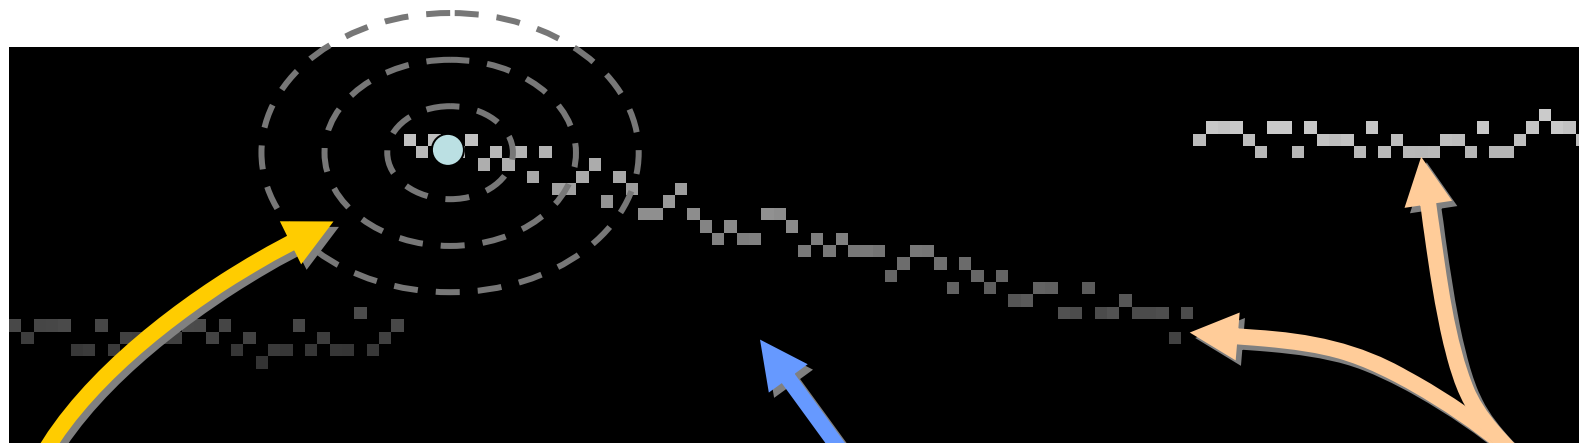


$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} & I_{\mathbf{p}}^{\text{bf}} \\ W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{\mathbf{q} \in \mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)}_{\text{space x range}} \begin{pmatrix} W_{\mathbf{q}} & I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

Corresponds to a 3D Gaussian on a 2D image.

# Link with Linear Filtering

## 2. Introducing a Convolution



sum all values

black = zero

$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} & I_{\mathbf{p}}^{\text{bf}} \\ W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{(\mathbf{q}, \zeta) \in \mathcal{S} \times \mathcal{R}} \begin{pmatrix} W_{\mathbf{q}} & I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

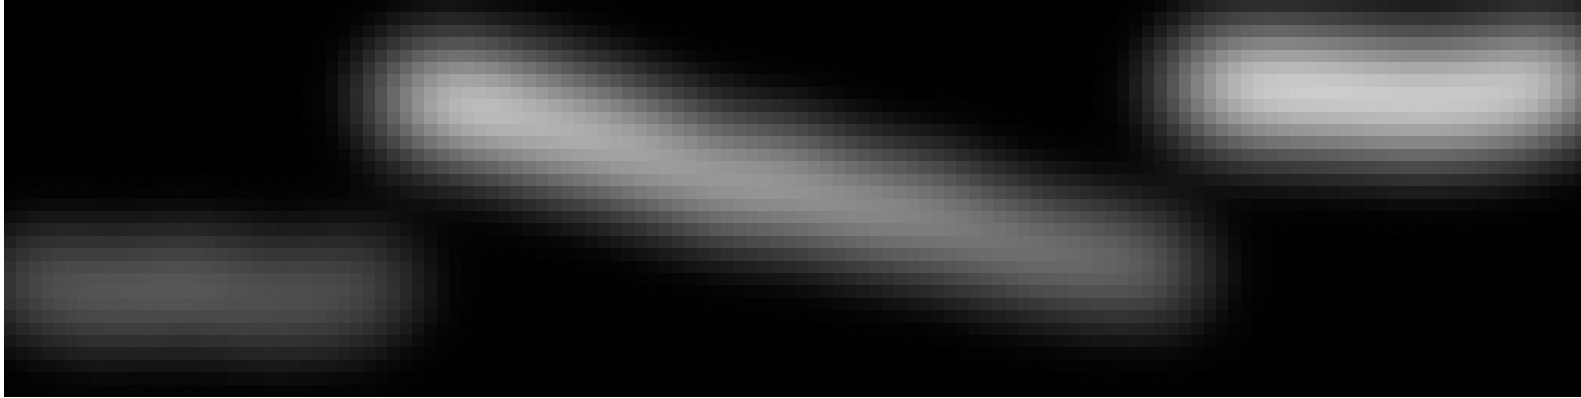
space-range Gaussian

sum all values multiplied by kernel  $\Rightarrow$  convolution

# Link with Linear Filtering

## 2. Introducing a Convolution

---



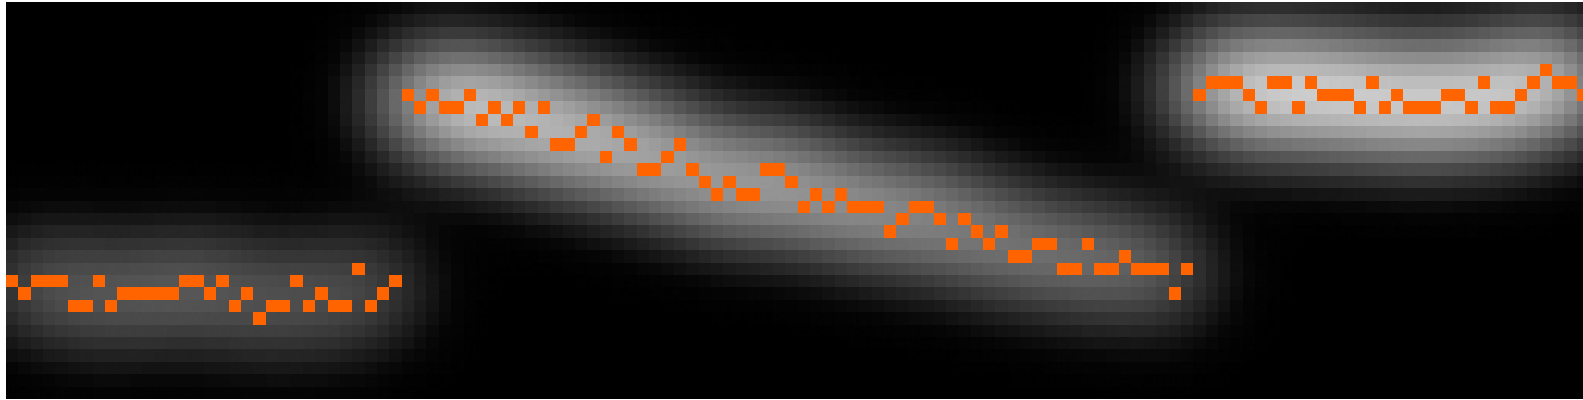
result of the convolution

$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} & I_{\mathbf{p}}^{\text{bf}} \\ W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{(\mathbf{q}, \zeta) \in \mathcal{S} \times \mathcal{R}} \text{space-range Gaussian} \begin{pmatrix} W_{\mathbf{q}} & I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

# Link with Linear Filtering

## 2. Introducing a Convolution

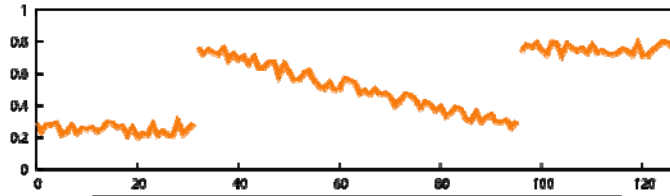
---



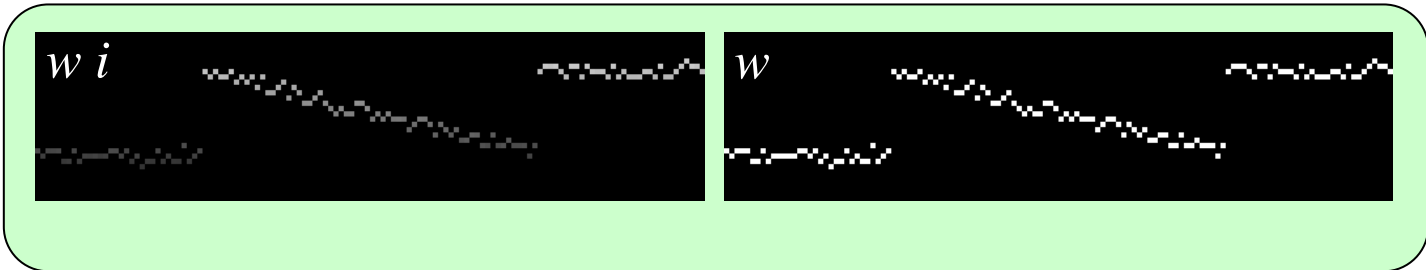
result of the convolution

$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} & I_{\mathbf{p}}^{\text{bf}} \\ W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{(\mathbf{q}, \zeta) \in \mathcal{S} \times \mathcal{R}} \text{space-range Gaussian} \begin{pmatrix} W_{\mathbf{q}} & I_{\mathbf{q}} \\ W_{\mathbf{q}} \end{pmatrix}$$

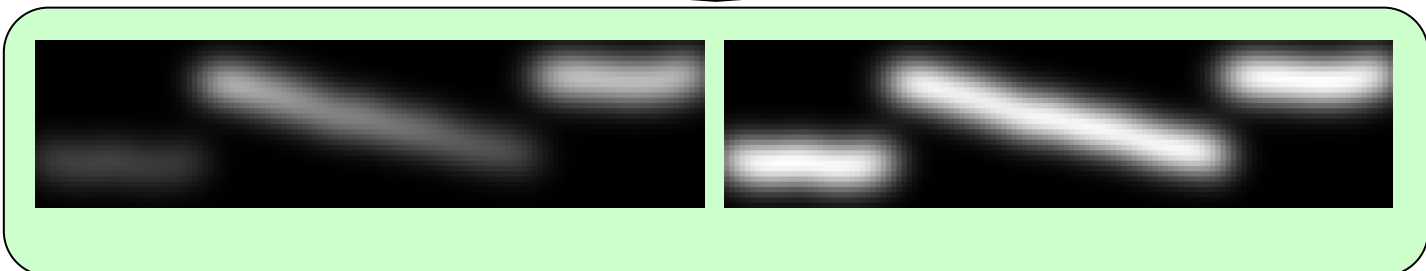




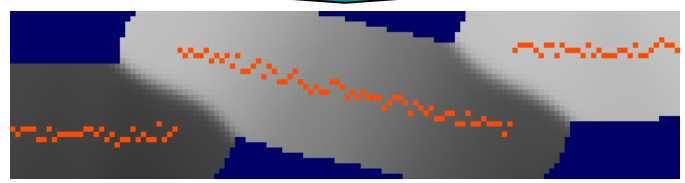
higher dimensional functions



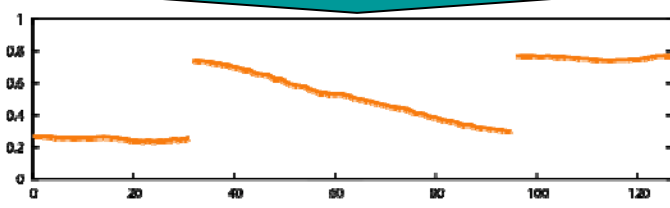
Gaussian convolution



division



slicing



# Reformulation: Summary

$$\text{linear:} \quad (w^{\text{bf}} \ i^{\text{bf}}, w^{\text{bf}}) = g_{\sigma_s, \sigma_r} \otimes (wi, w)$$

$$\text{nonlinear:} \quad I_{\mathbf{p}}^{\text{bf}} = \frac{w^{\text{bf}}(\mathbf{p}, I_{\mathbf{p}}) \ i^{\text{bf}}(\mathbf{p}, I_{\mathbf{p}})}{w^{\text{bf}}(\mathbf{p}, I_{\mathbf{p}})}$$

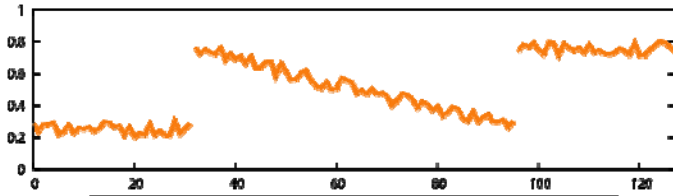
## 1. Convolution in higher dimension

- expensive but well understood (linear, FFT, etc)

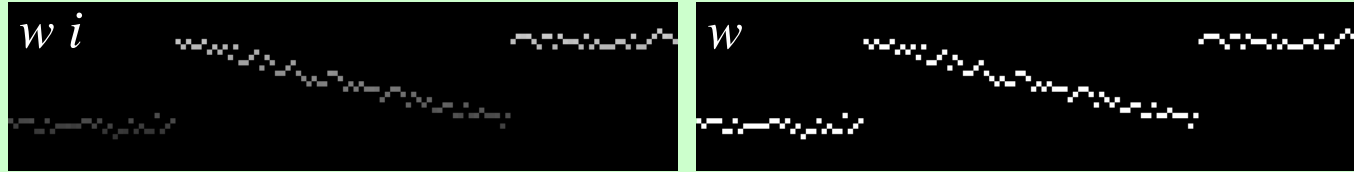
## 2. Division and slicing

- nonlinear but simple and pixel-wise

**Exact reformulation**

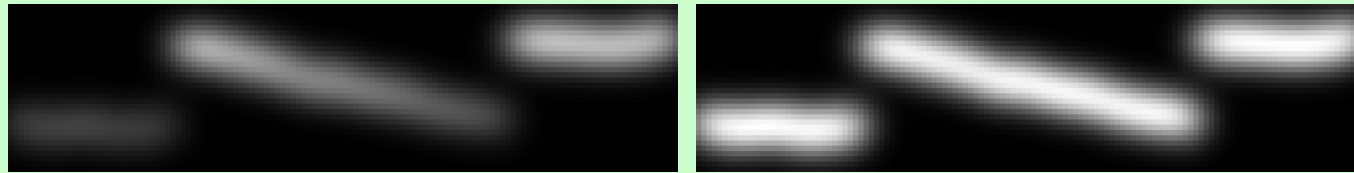


higher dimensional functions



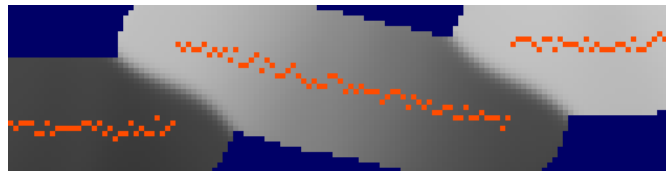
Low-pass filter

Gaussian convolution

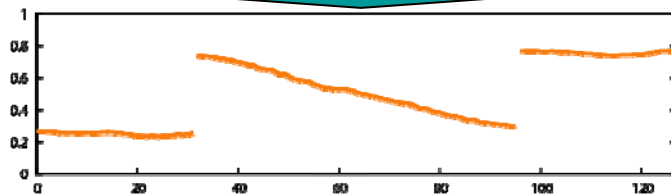


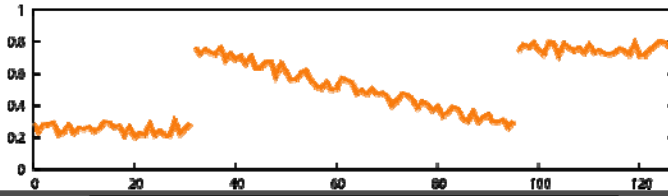
Almost only  
low freq.  
            
High freq.  
negligible

division

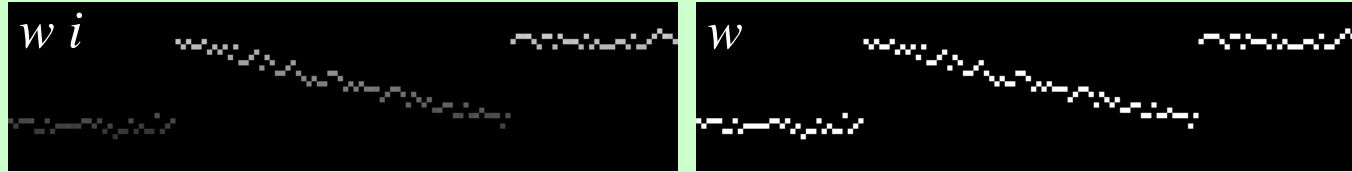


slicing



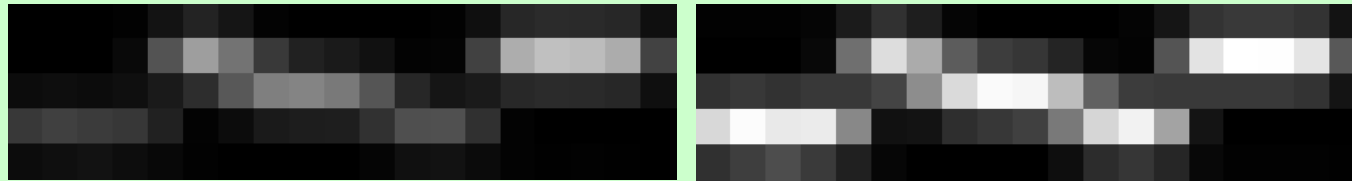


higher dimensional functions



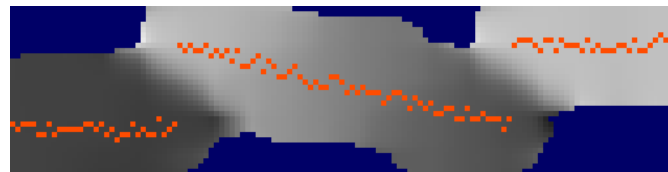
DOWNSAMPLE

Gaussian convolution

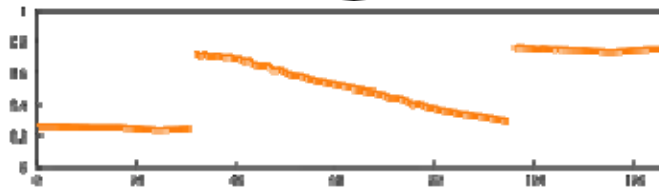


UPSAMPLE

division



slicing



Almost no  
information  
loss

# Fast Convolution by Downsampling

---

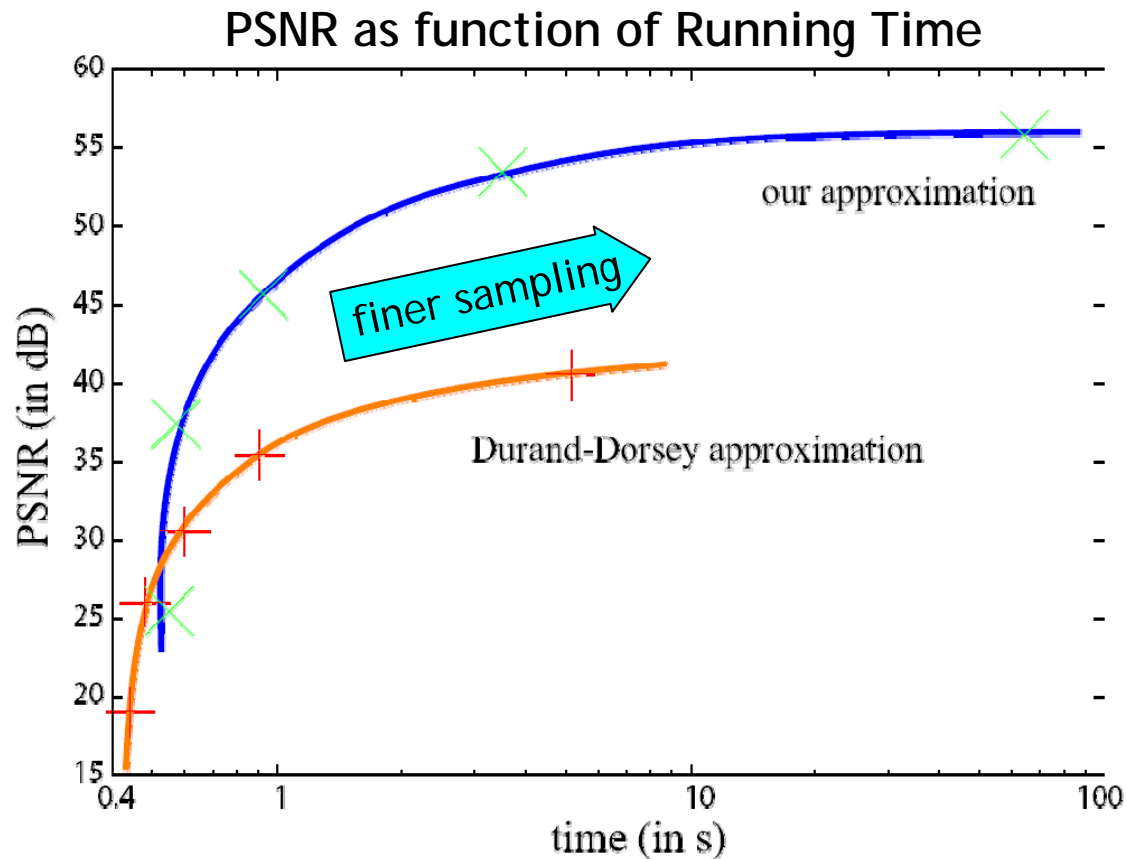
- Downsampling cuts frequencies above Nyquist limit
  - Less data to process
  - But induces error
- Evaluation of the approximation
  - Precision versus running time
  - Visual accuracy

# Accuracy versus Running Time

- Finer sampling increases accuracy.
- More precise than previous work.

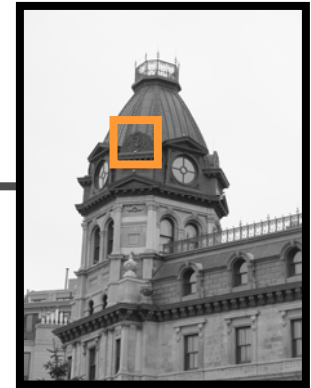


Digital photograph  
1200 × 1600



Straightforward implementation is over 10 minutes.

# Visual Results



1200 × 1600

- Comparison with previous work [Durand 02]
  - running time = 1s for both techniques

input



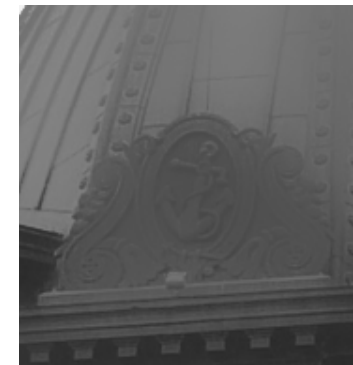
exact BF



our result



prev. work



difference  
with exact  
computation  
(intensities in [0:1])



# Conclusions

---

higher dimension  $\Rightarrow$  “better” computation

## Practical gain

- Interactive running time
- Visually similar results
- Simple to code (100 lines)

## Theoretical gain

- Link with linear filters
- Separation linear/nonlinear
- Signal processing framework



# Two-scale Tone Management for Photographic Look

Soonmin Bae, Sylvain Paris, and Frédo Durand

MIT CSAIL

SIGGRAPH2006

# Ansel Adams

---



Ansel Adams, *Clearing Winter Storm*

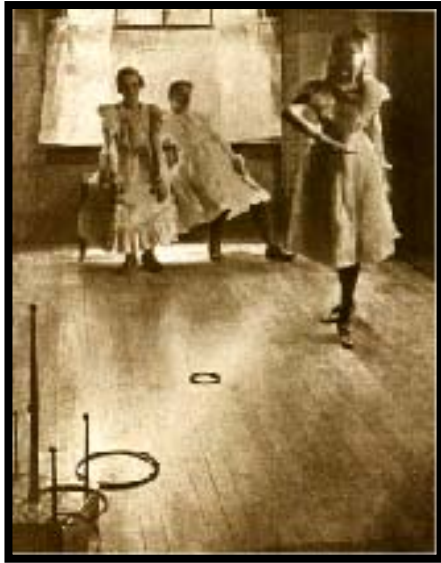
# An Amateur Photographer

---



# A Variety of Looks

---



# Goals

---

- Control over photographic look
- Transfer “look” from a model photo

For example,

we want



with the look of



# Aspects of Photographic Look

---

- Subject choice
- Framing and composition
- ➔ Specified by input photos
  
- Tone distribution and contrast
- ➔ Modified based on model photos



Input



Model

# Tonal Aspects of Look

---



Ansel Adams



Kenro Izu

# Tonal aspects of Look - Global Contrast DigiVFX

---



Ansel Adams



Kenro Izu

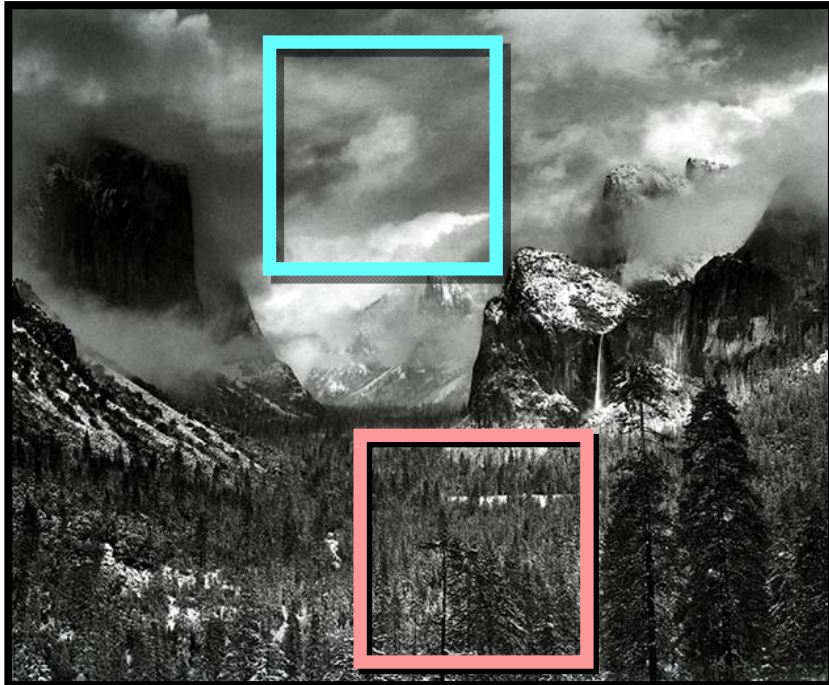
**High Global Contrast**

**Low Global Contrast**



# Tonal aspects of Look - Local Contrast DigiVFX

---



Ansel Adams

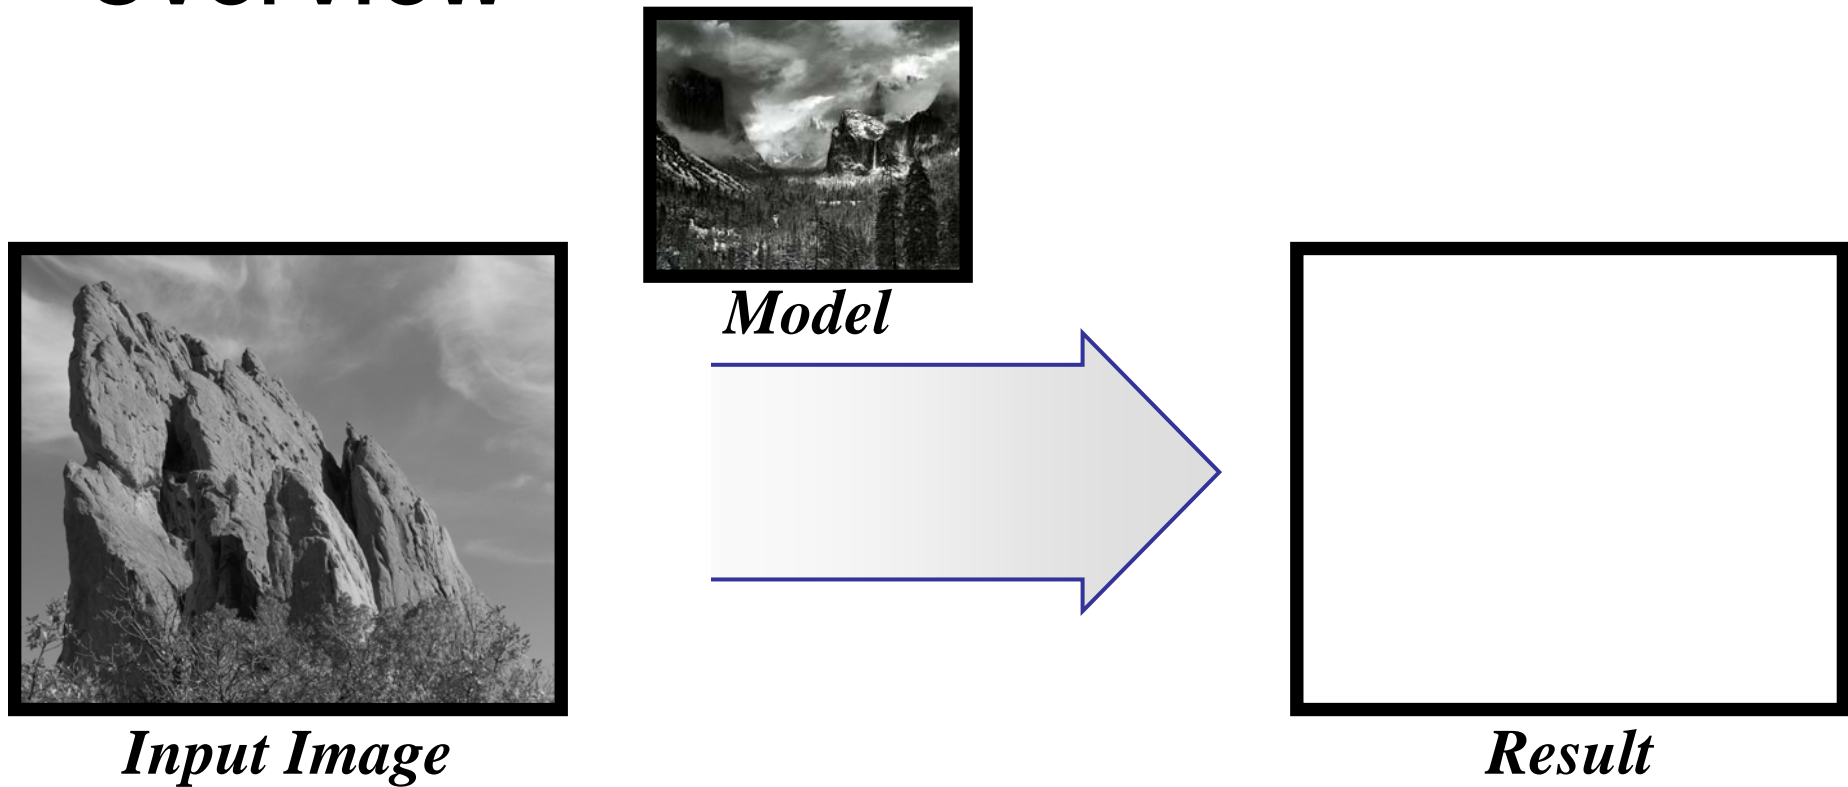


Kenro Izu

**Variable amount of texture**

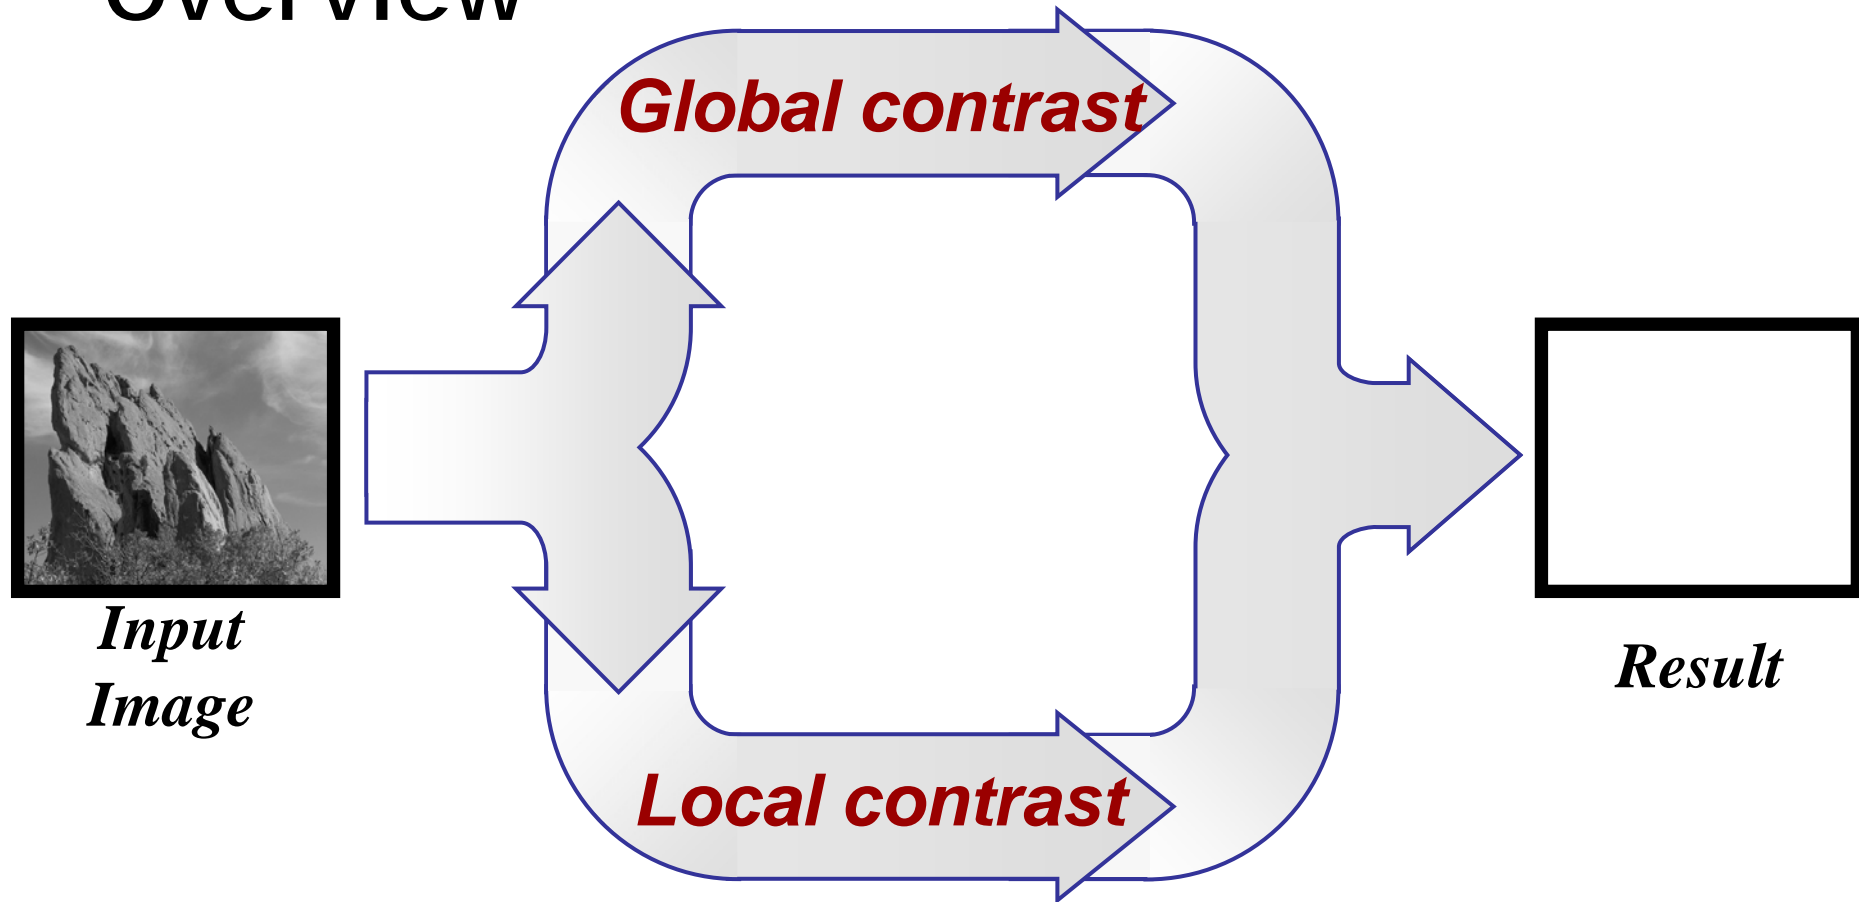
**Texture everywhere**

# Overview



- Transfer look between photographs
  - Tonal aspects

# Overview

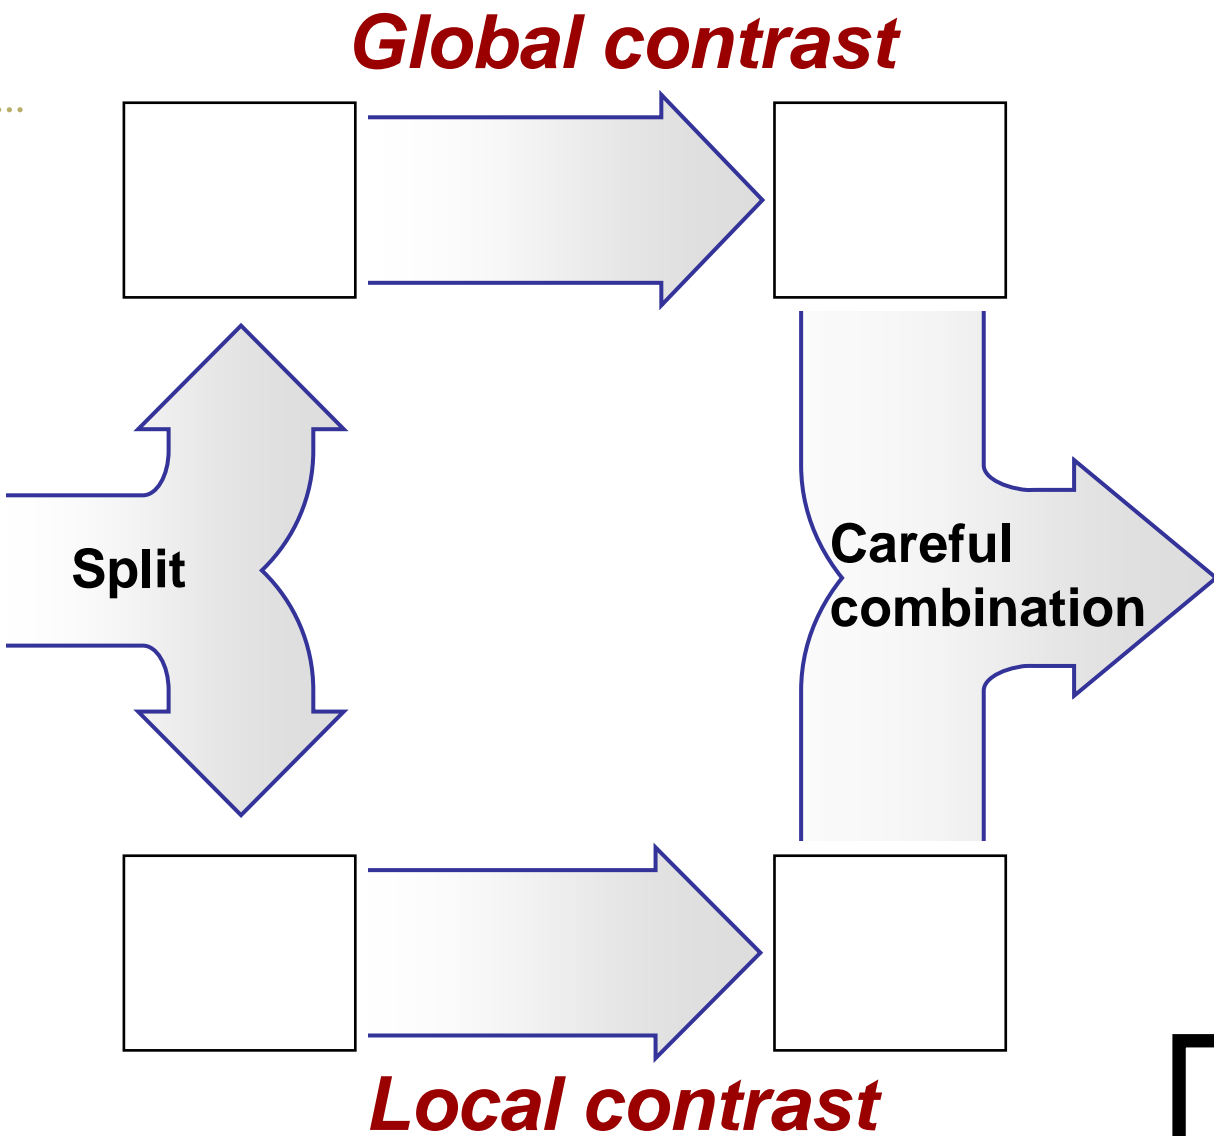


- Separate global and local contrast

# Overview



*Input Image*

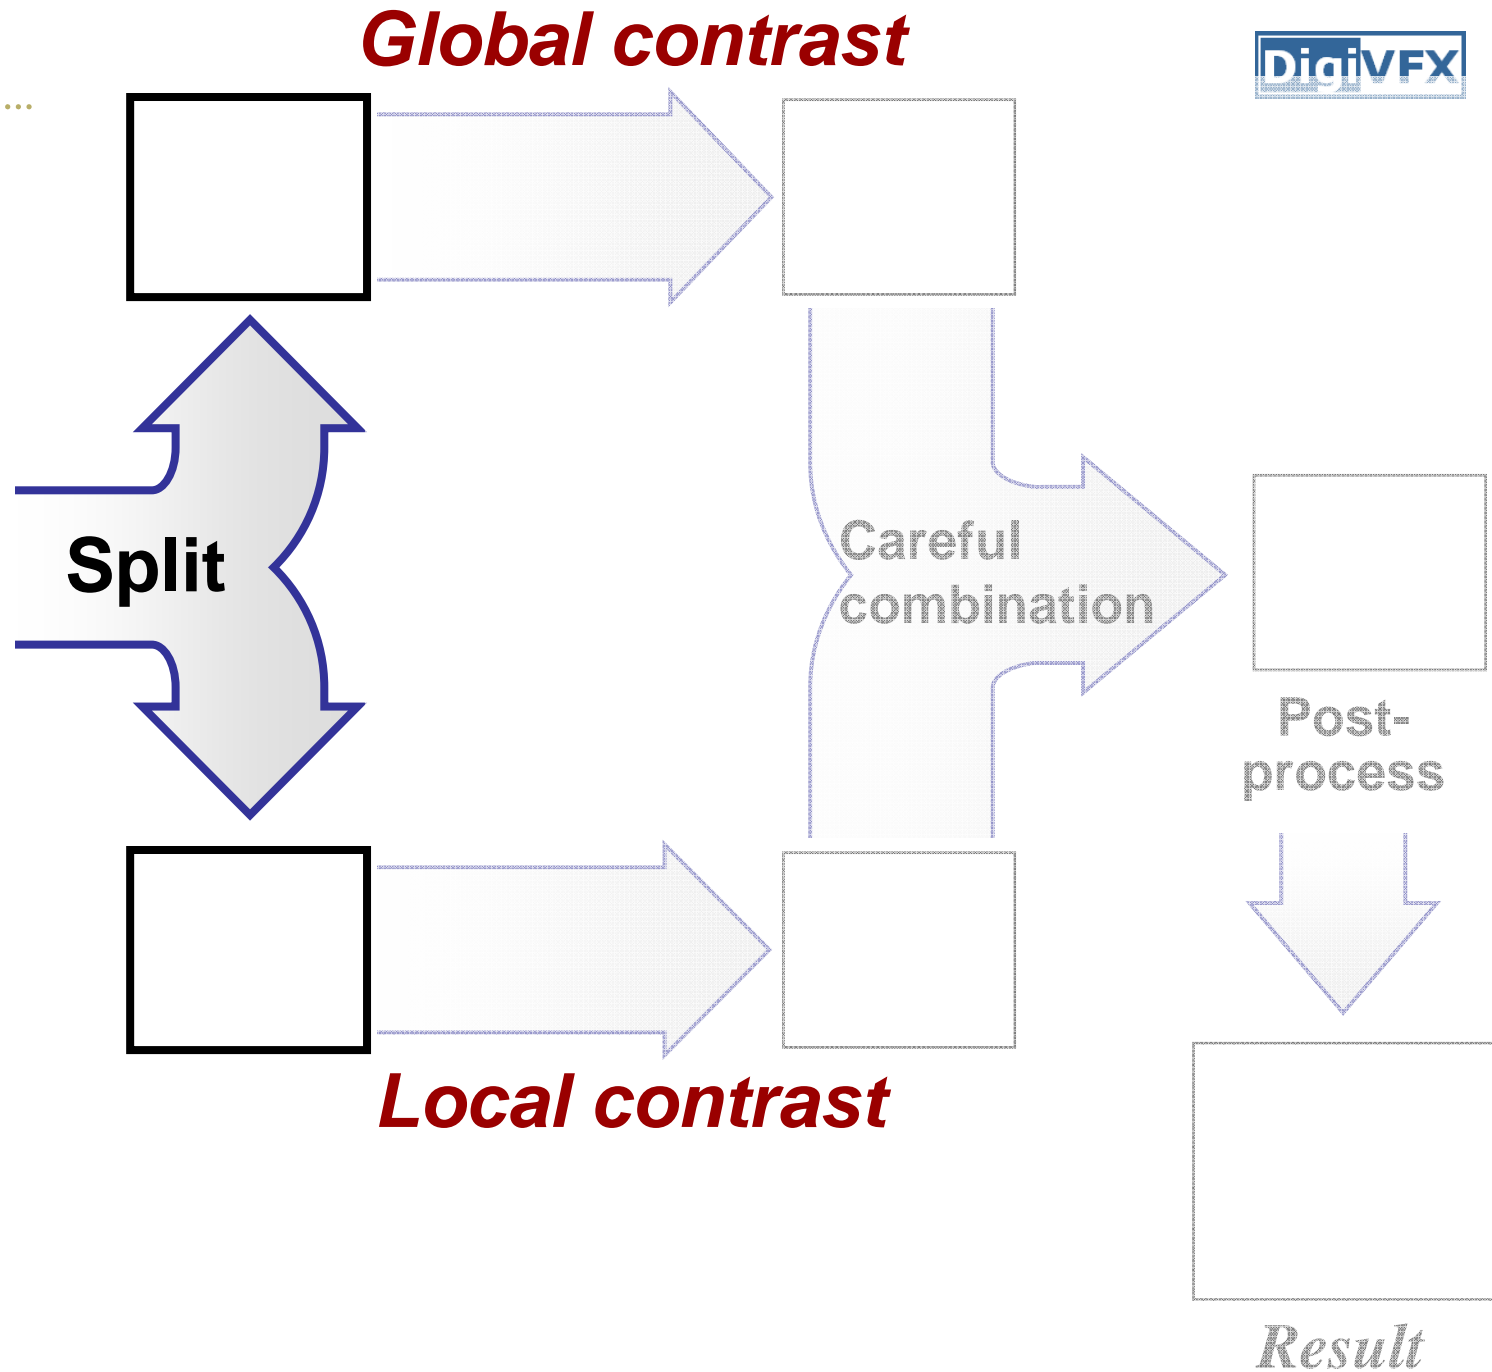


*Result*

# Overview



*Input Image*

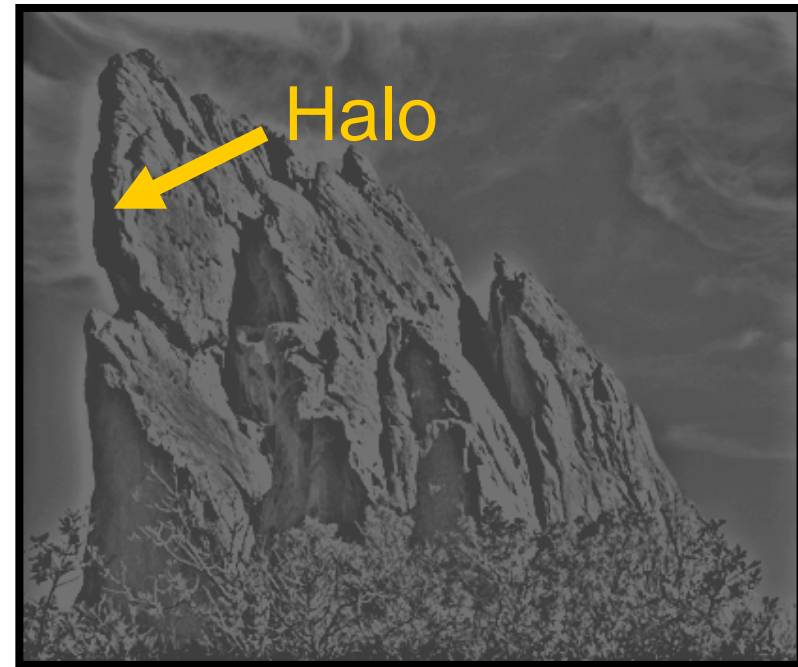


# Split Global vs. Local Contrast

- Naive decomposition: low vs. high frequency
  - Problem: introduce blur & halos



**Low frequency**  
***Global contrast***



**High frequency**  
***Local contrast***

# Bilateral Filter

---

- Edge-preserving smoothing [Tomasi 98]
- We build upon tone mapping [Durand 02]



After bilateral filtering  
*Global contrast*



Residual after filtering  
*Local contrast*

# Bilateral Filter

---

- Edge-preserving smoothing [Tomasi 98]
- We build upon tone mapping [Durand 02]



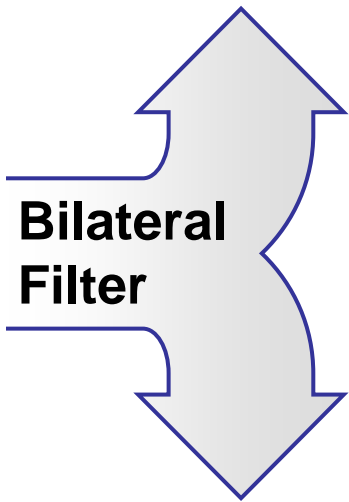
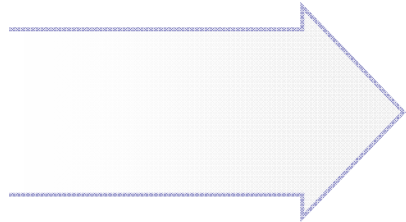
After bilateral filtering  
*Global contrast*



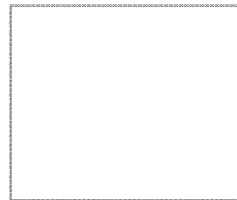
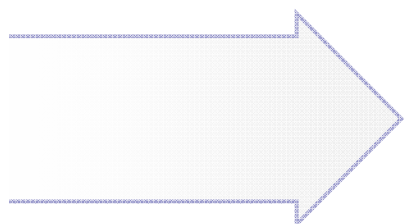
Residual after filtering  
*Local contrast*



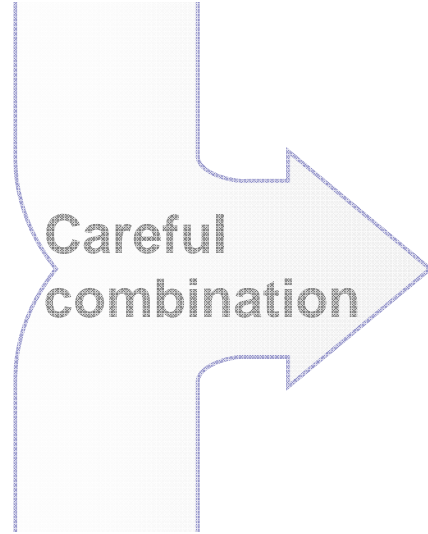
**Global contrast**



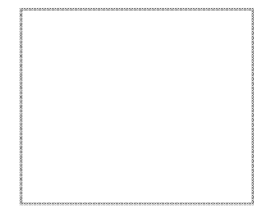
**Bilateral Filter**



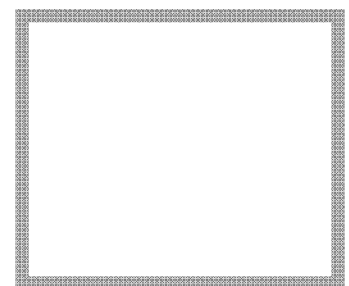
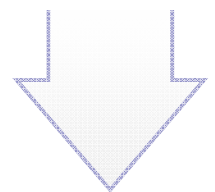
**Local contrast**



**Careful combination**



**Post-process**

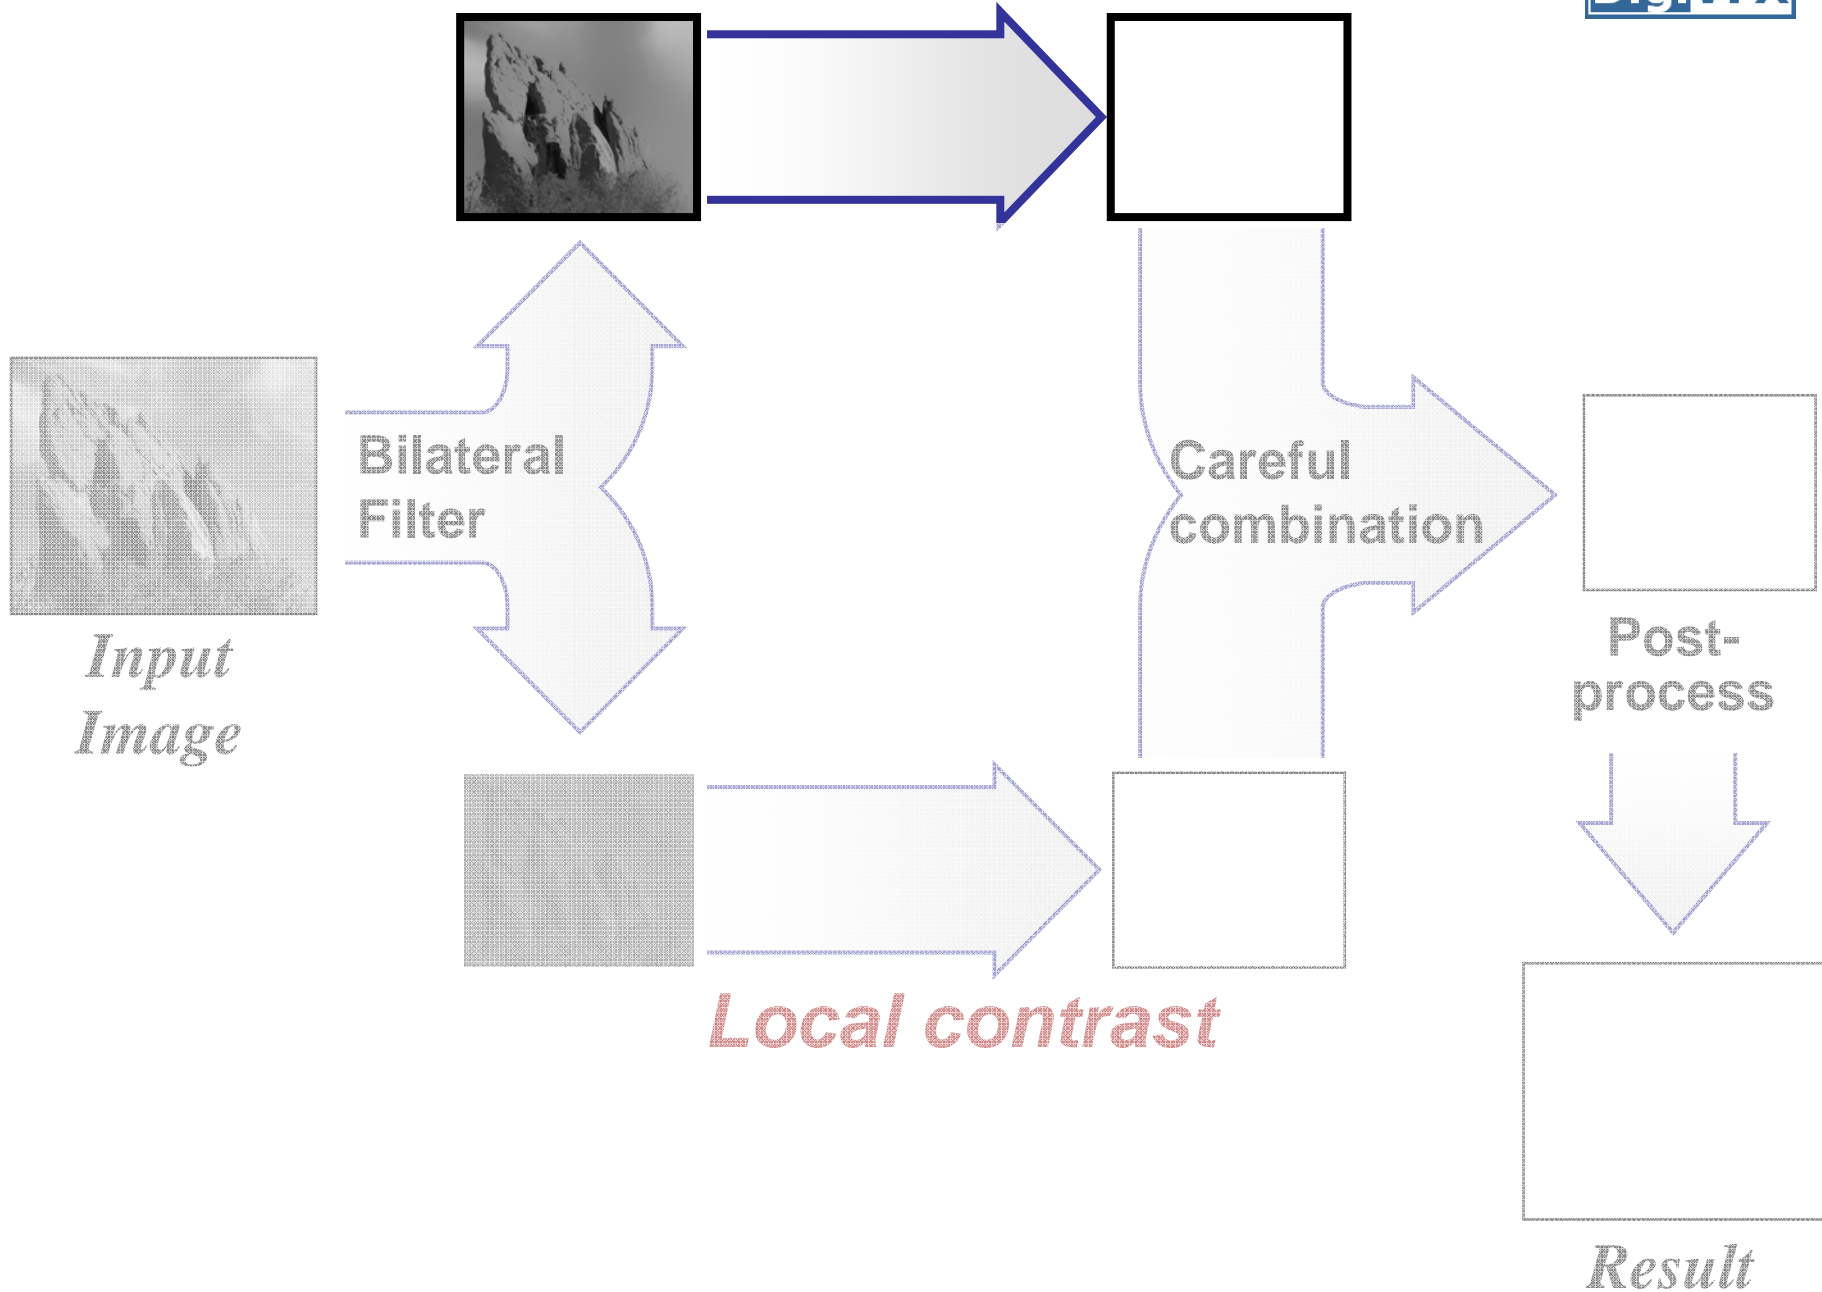


**Result**



**Input Image**

# Global contrast



# Global Contrast

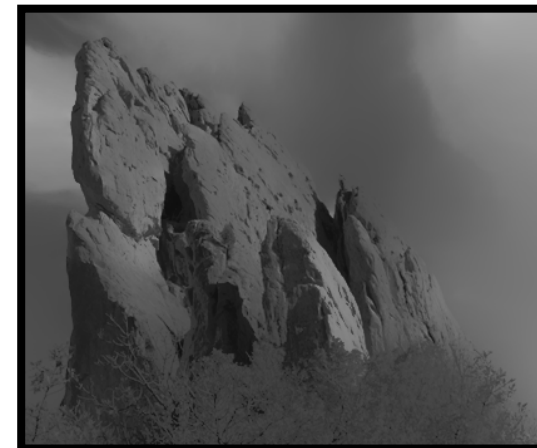
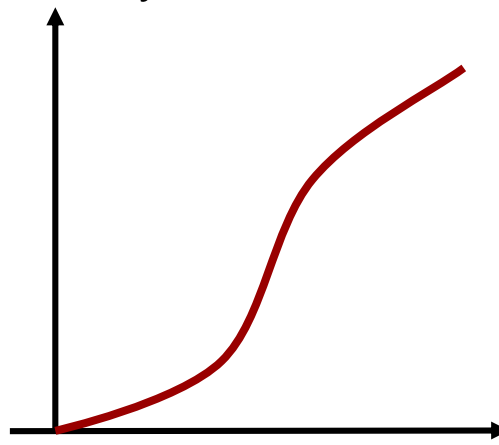
---

- Intensity remapping of base layer



Input base

Remapped  
intensity

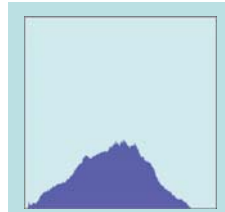


After remapping

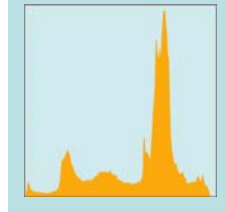
# Global Contrast (Model Transfer)



Model  
base

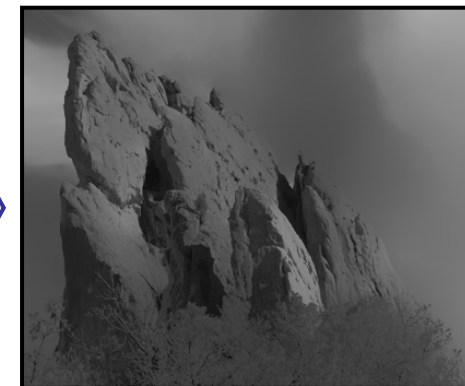
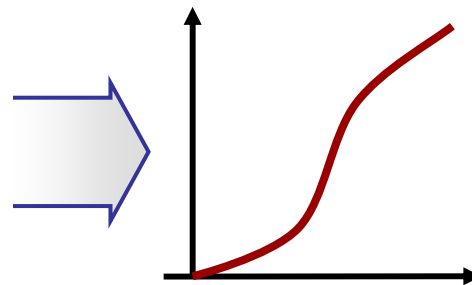
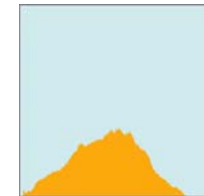


Input  
base

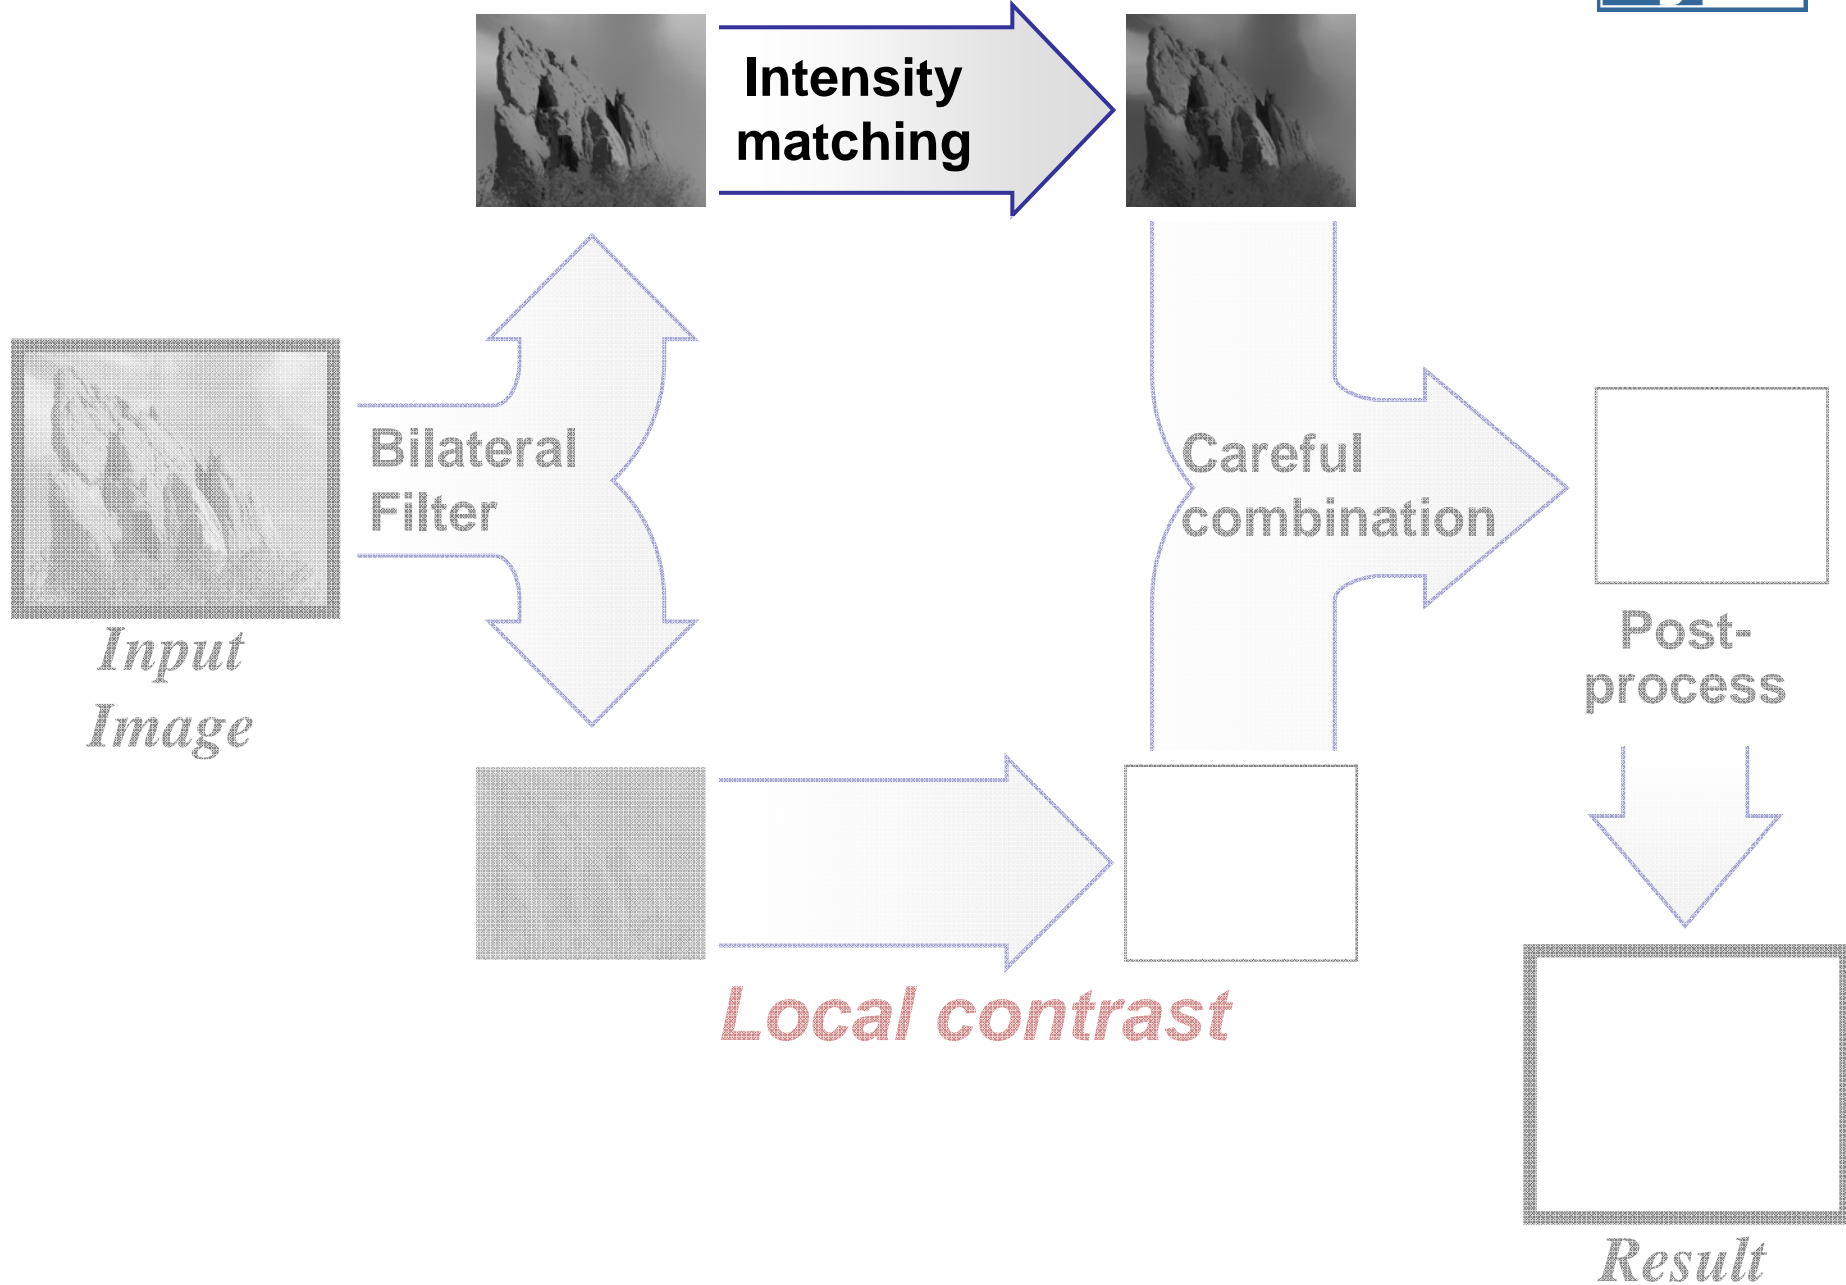


- Histogram matching
  - Remapping function given input and model histogram

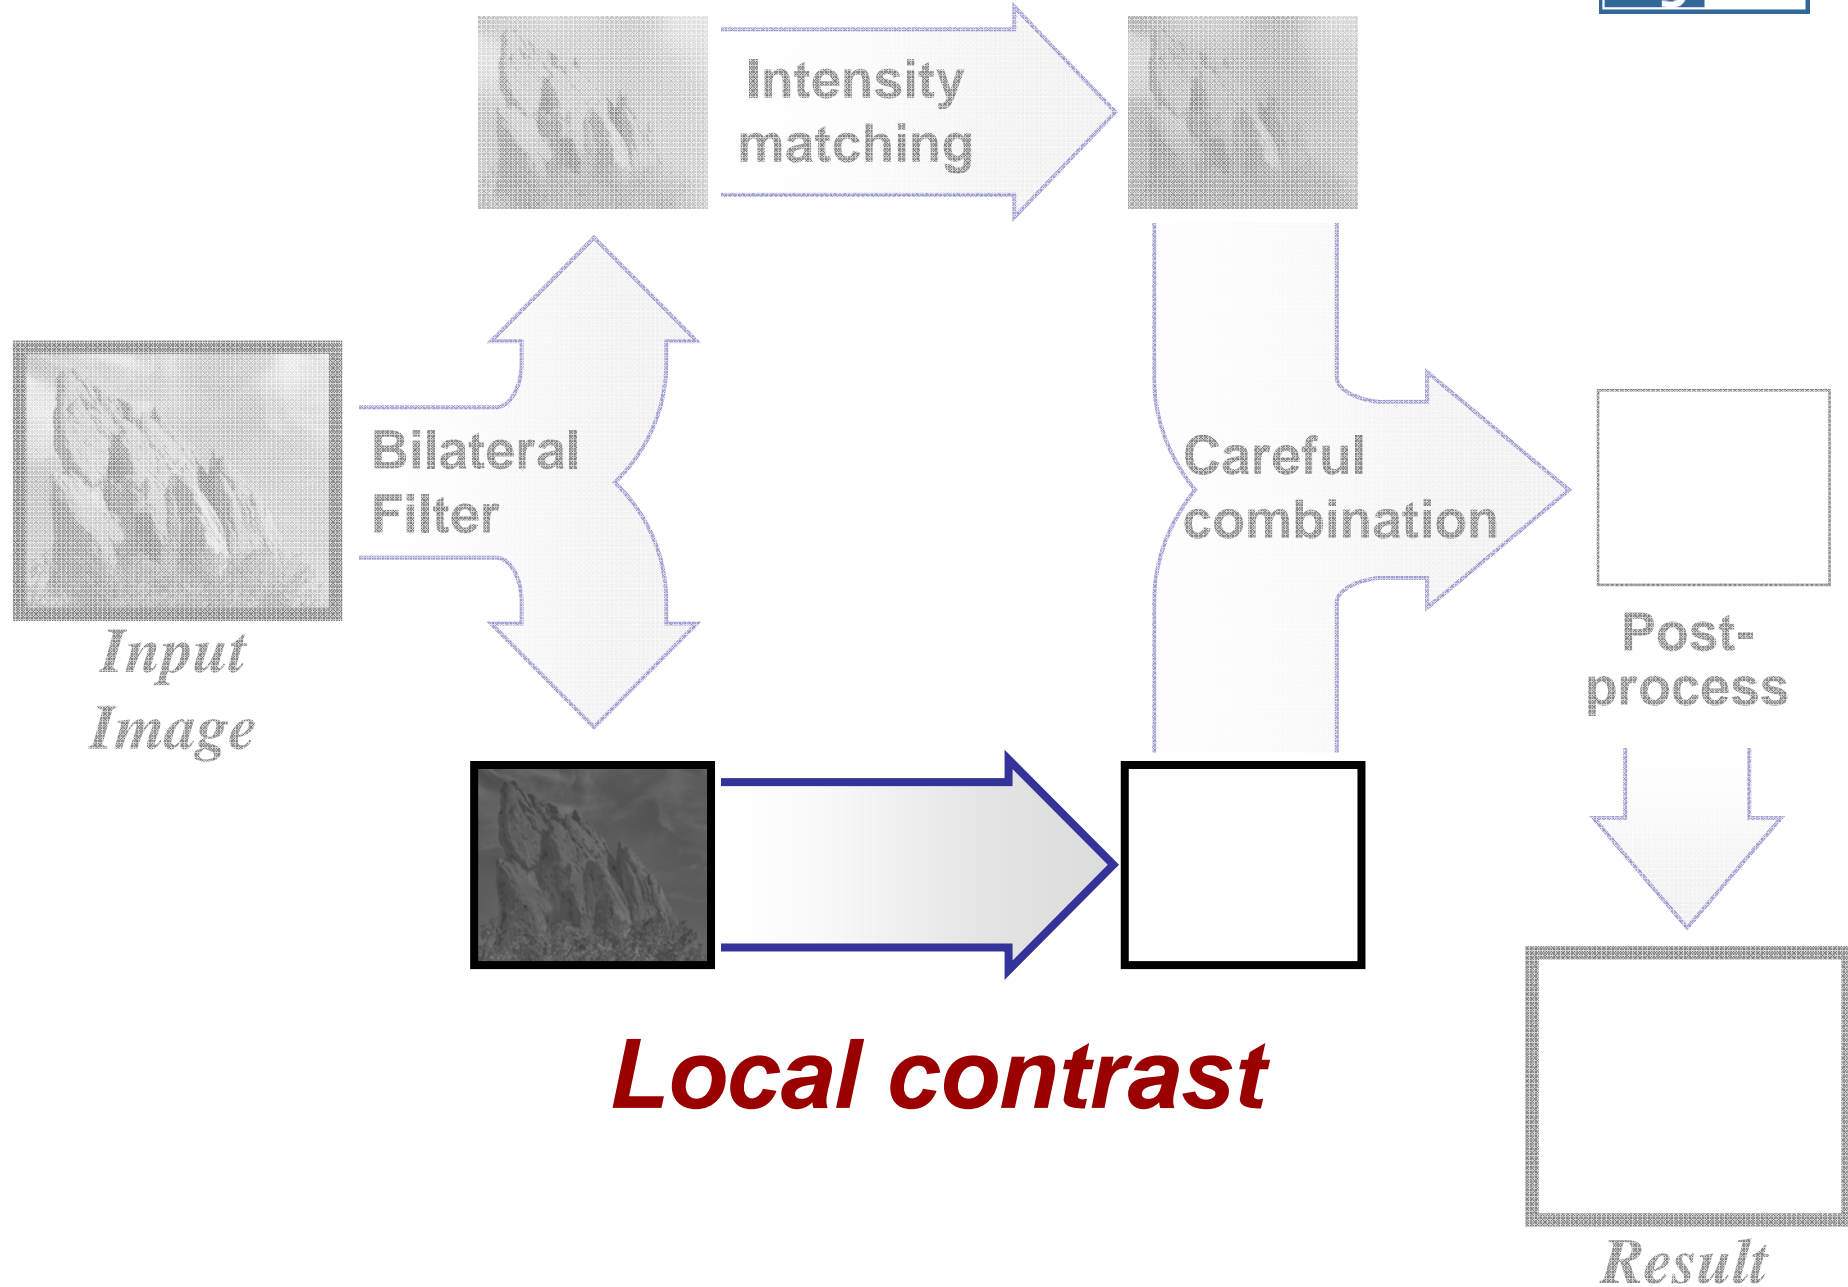
Output  
base



# Global contrast



**Global contrast**



**Local contrast**

*Result*

# Local Contrast: Detail Layer

---

- Uniform control:
  - Multiply all values in the detail layer



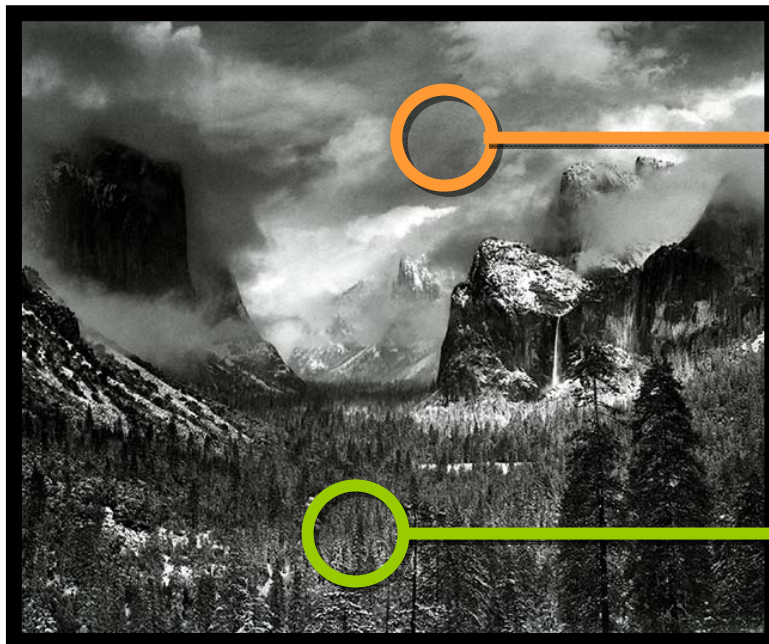
Input



Base + 3 × Detail

# The amount of local contrast is not uniform

---



Smooth region

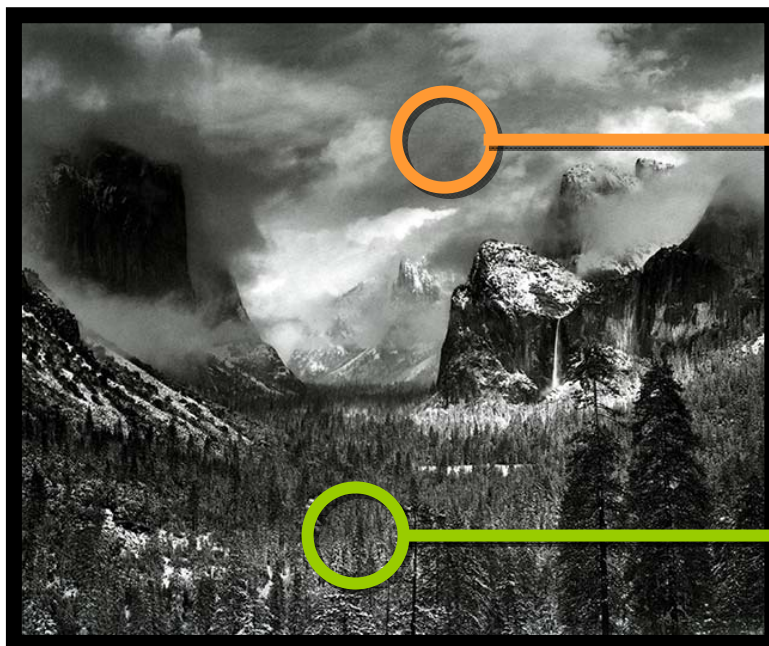
Textured region



# Local Contrast Variation

---

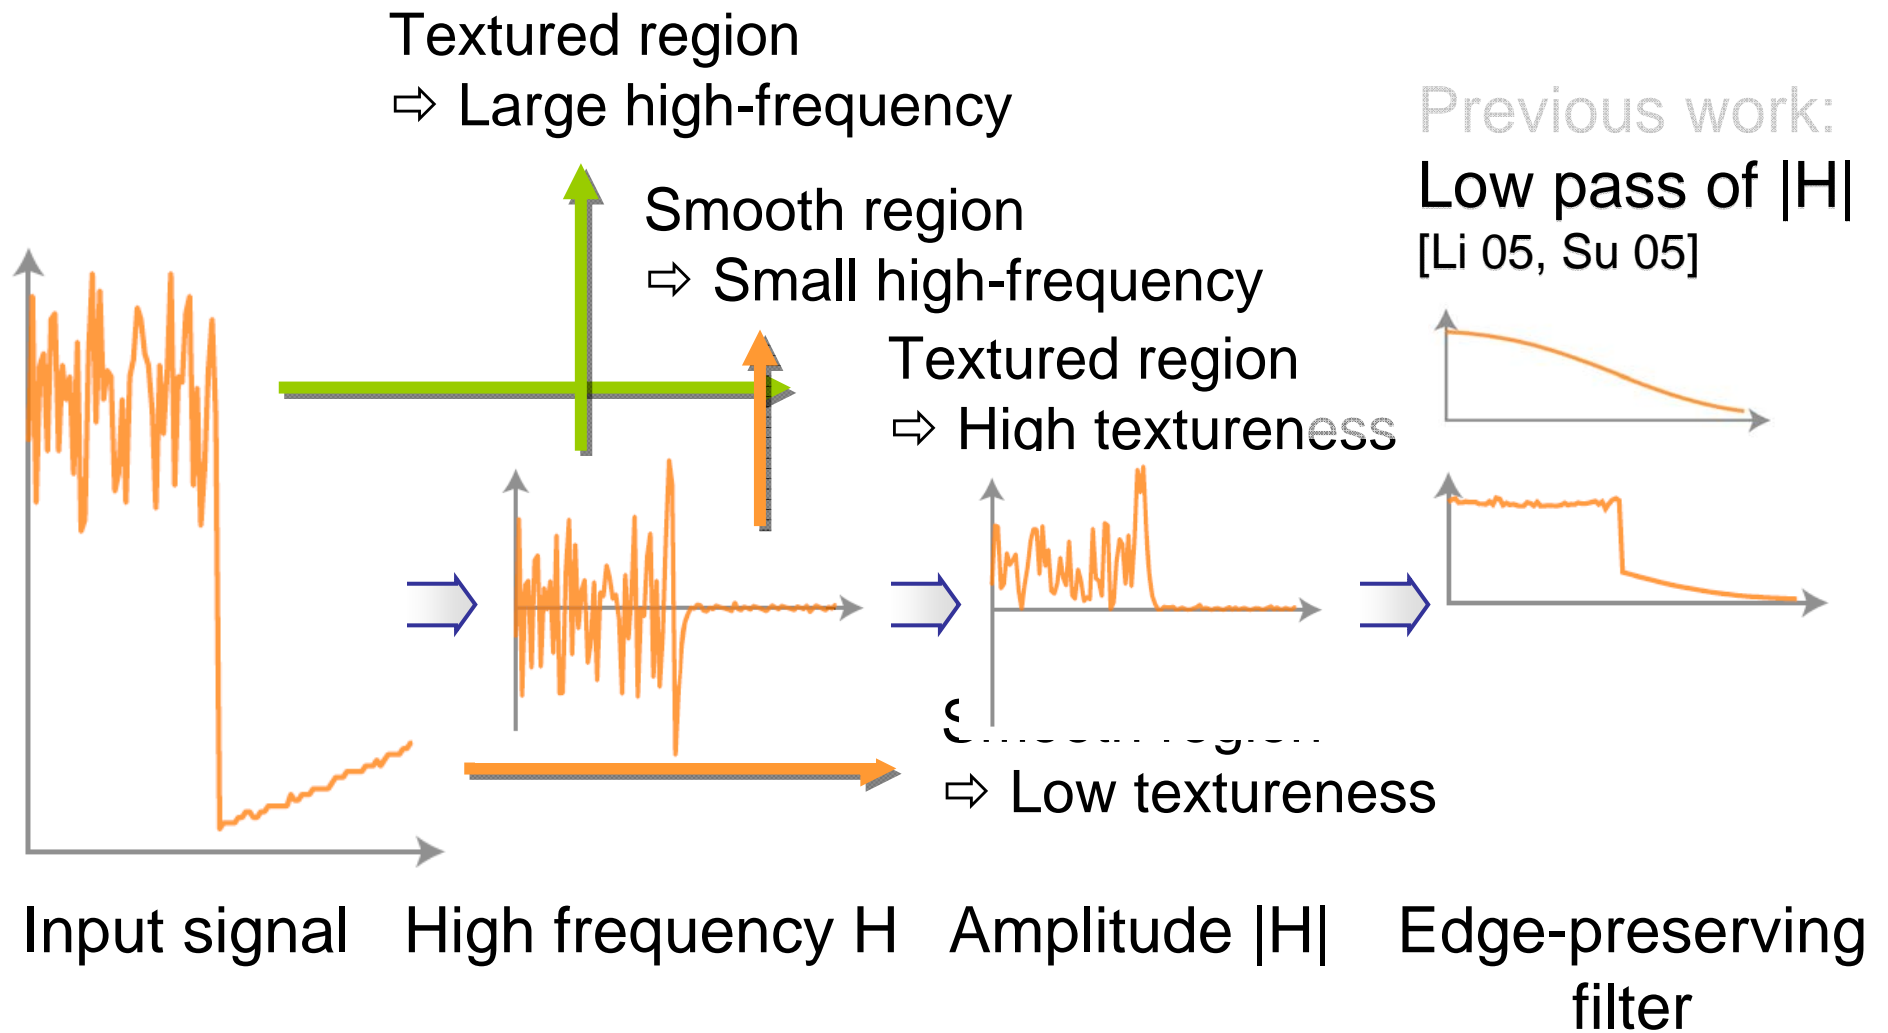
- We define “textureness”: amount of local contrast
  - at each pixel based on surrounding region



Smooth region  
⇒ Low textureness

Textured region  
⇒ High textureness

# “Textureness”: 1D Example

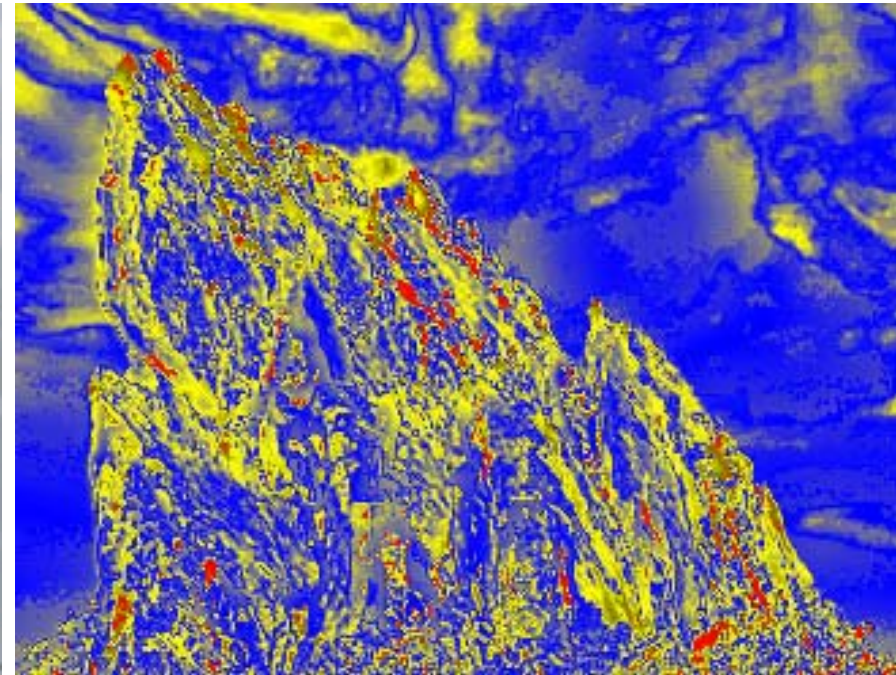


# Textureness

---



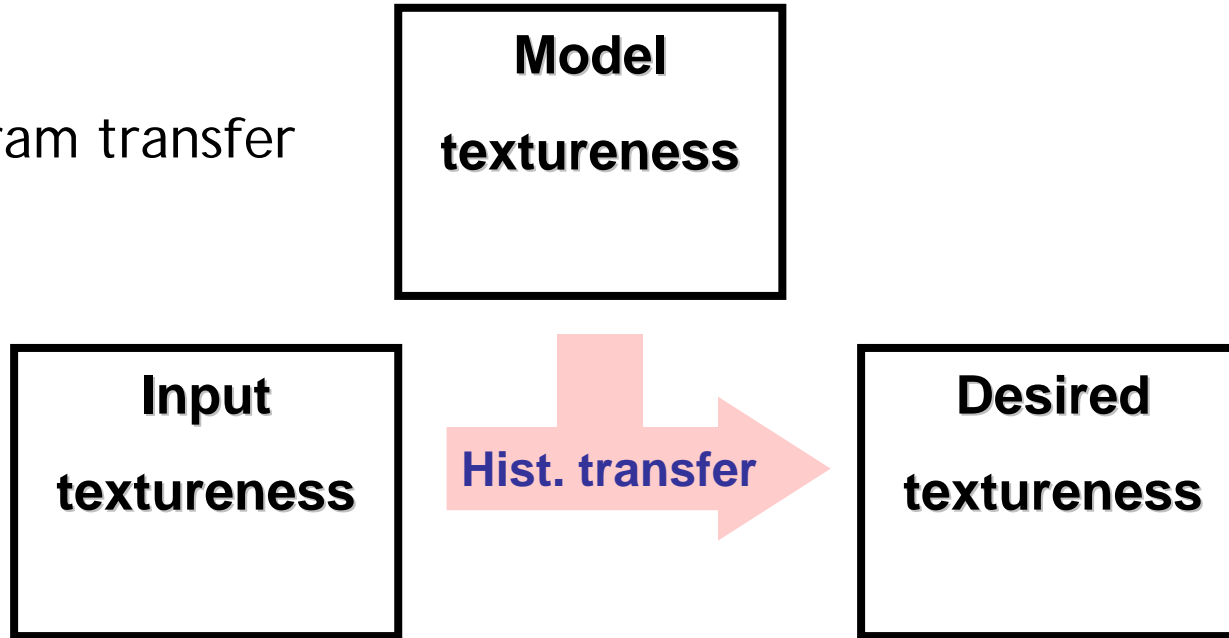
Input



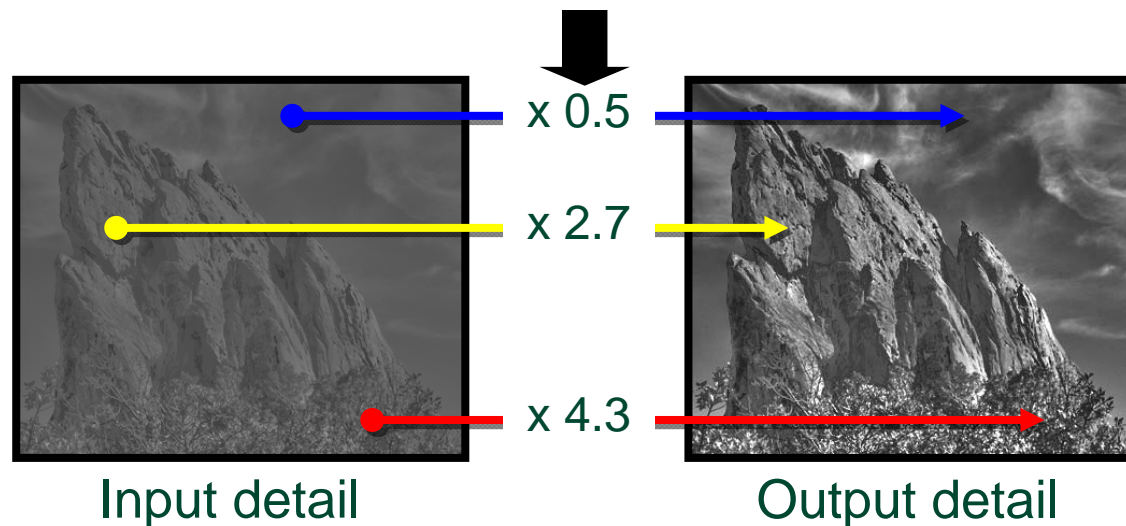
Textureness

# Textureness Transfer

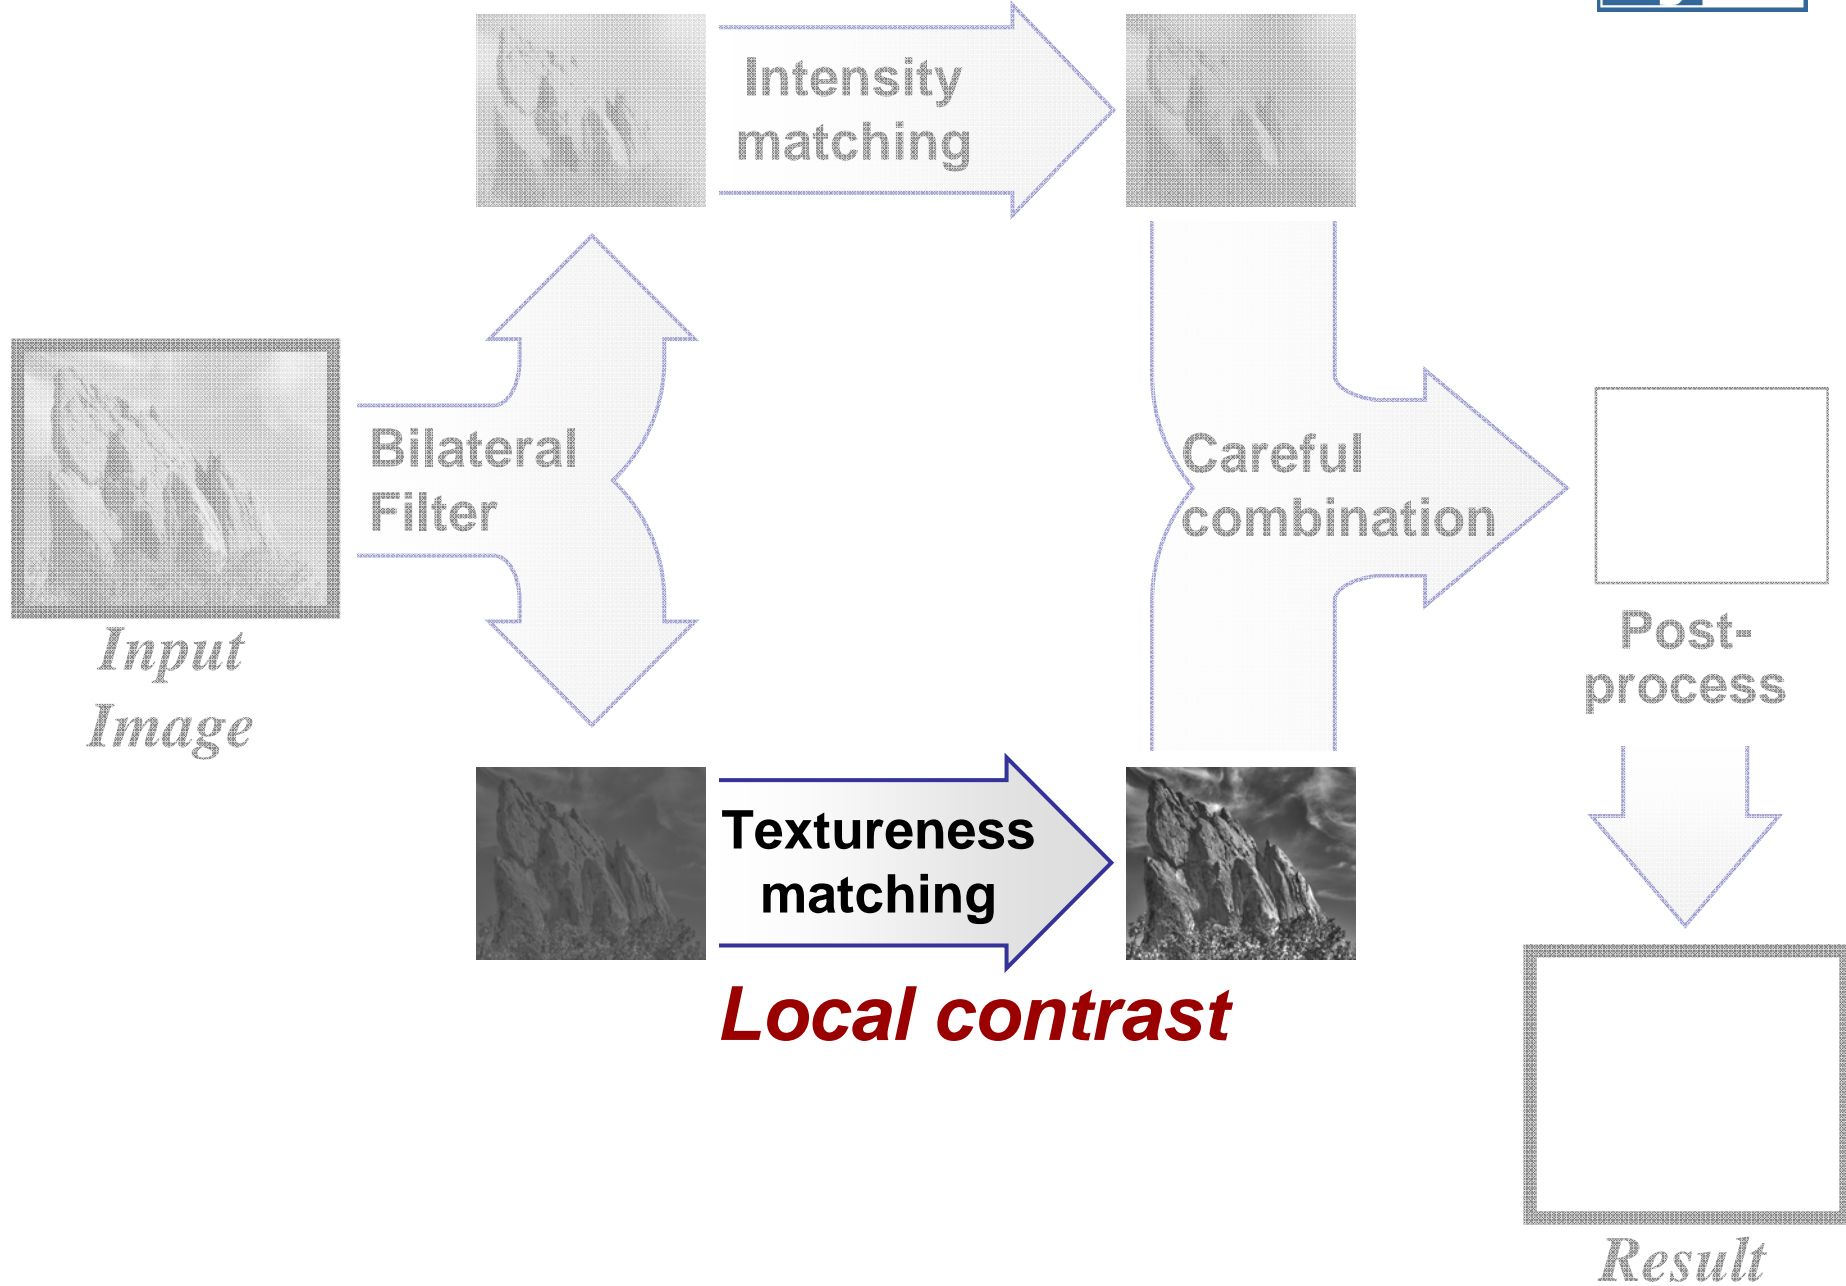
Step 1:  
Histogram transfer



Step 2:  
Scaling detail layer  
(per pixel) to match  
desired textureness



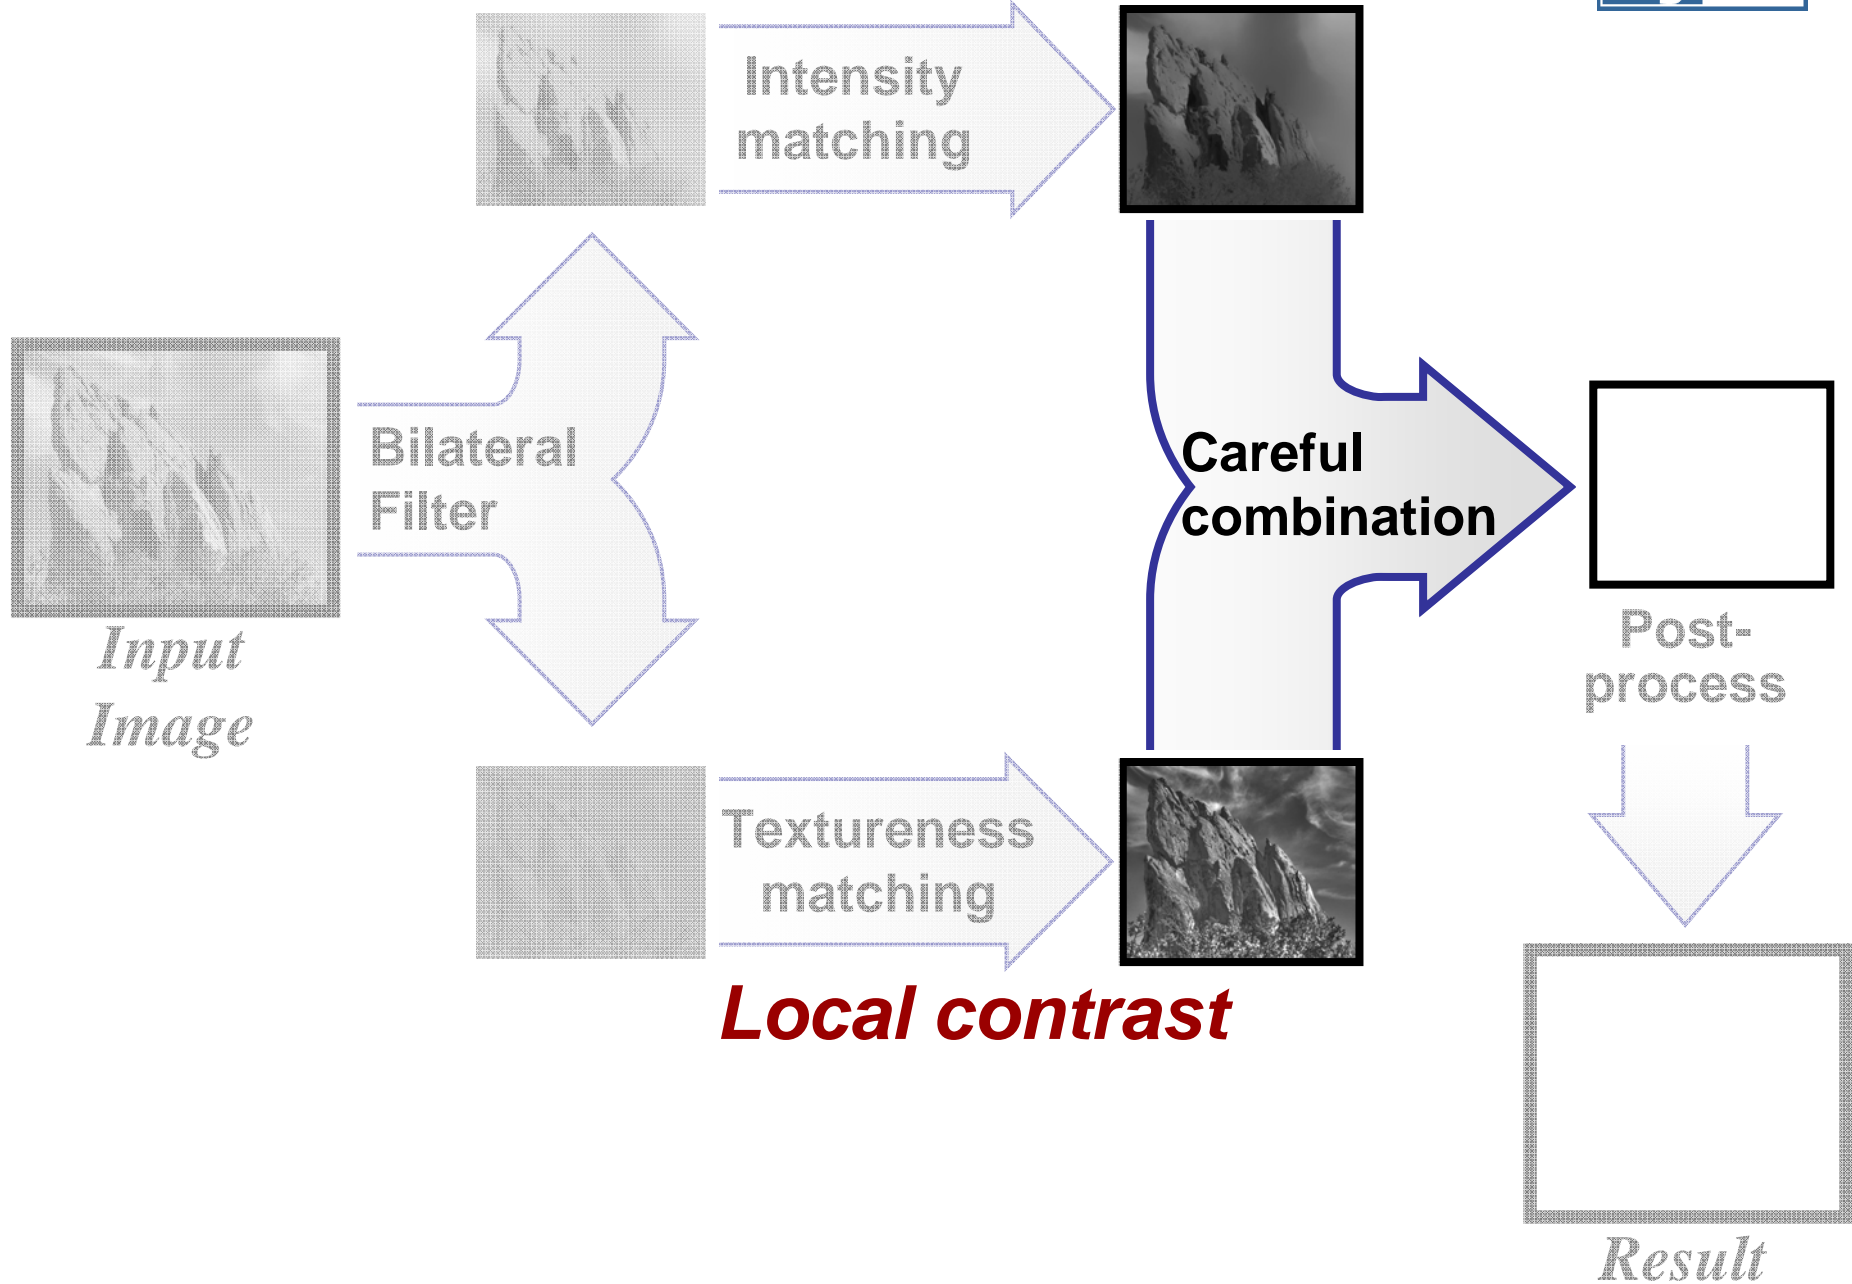
**Global contrast**



**Local contrast**

*Result*

# Global contrast



# A Non Perfect Result

- Decoupled and large modifications (up to 6x)  
→ Limited defects may appear

input (HDR)



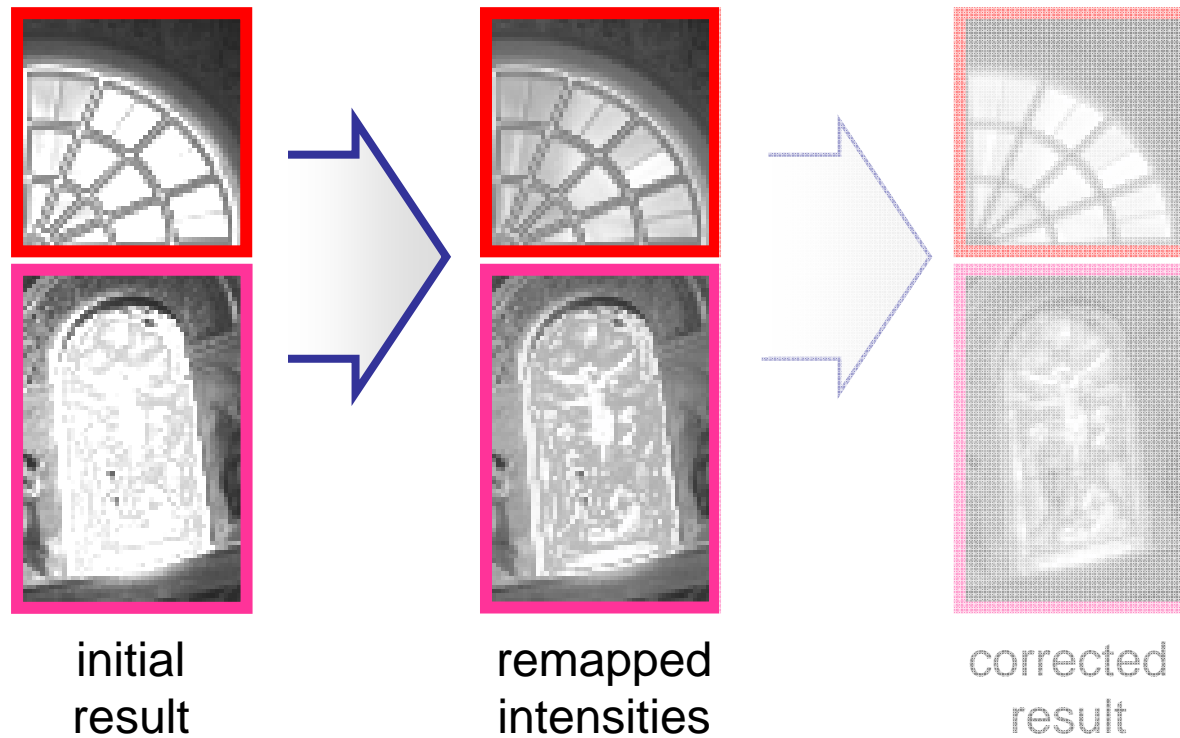
result after  
global and local adjustments



# Intensity Remapping

---

- Some intensities may be outside displayable range.
- ➔ Compress histogram to fit visible range.

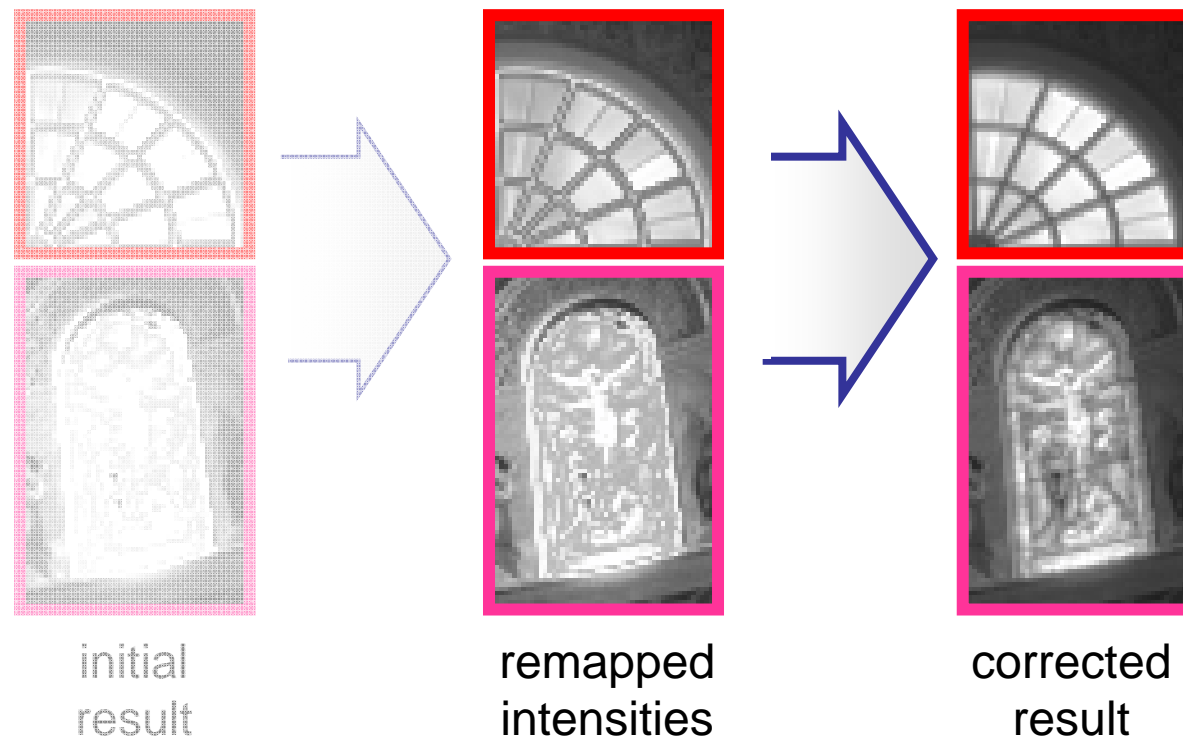




# Preserving Details

---

1. In the gradient domain:
  - Compare gradient amplitudes of input and current
  - Prevent extreme reduction & extreme increase
2. Solve the Poisson equation.



# Effect of Detail Preservation

---

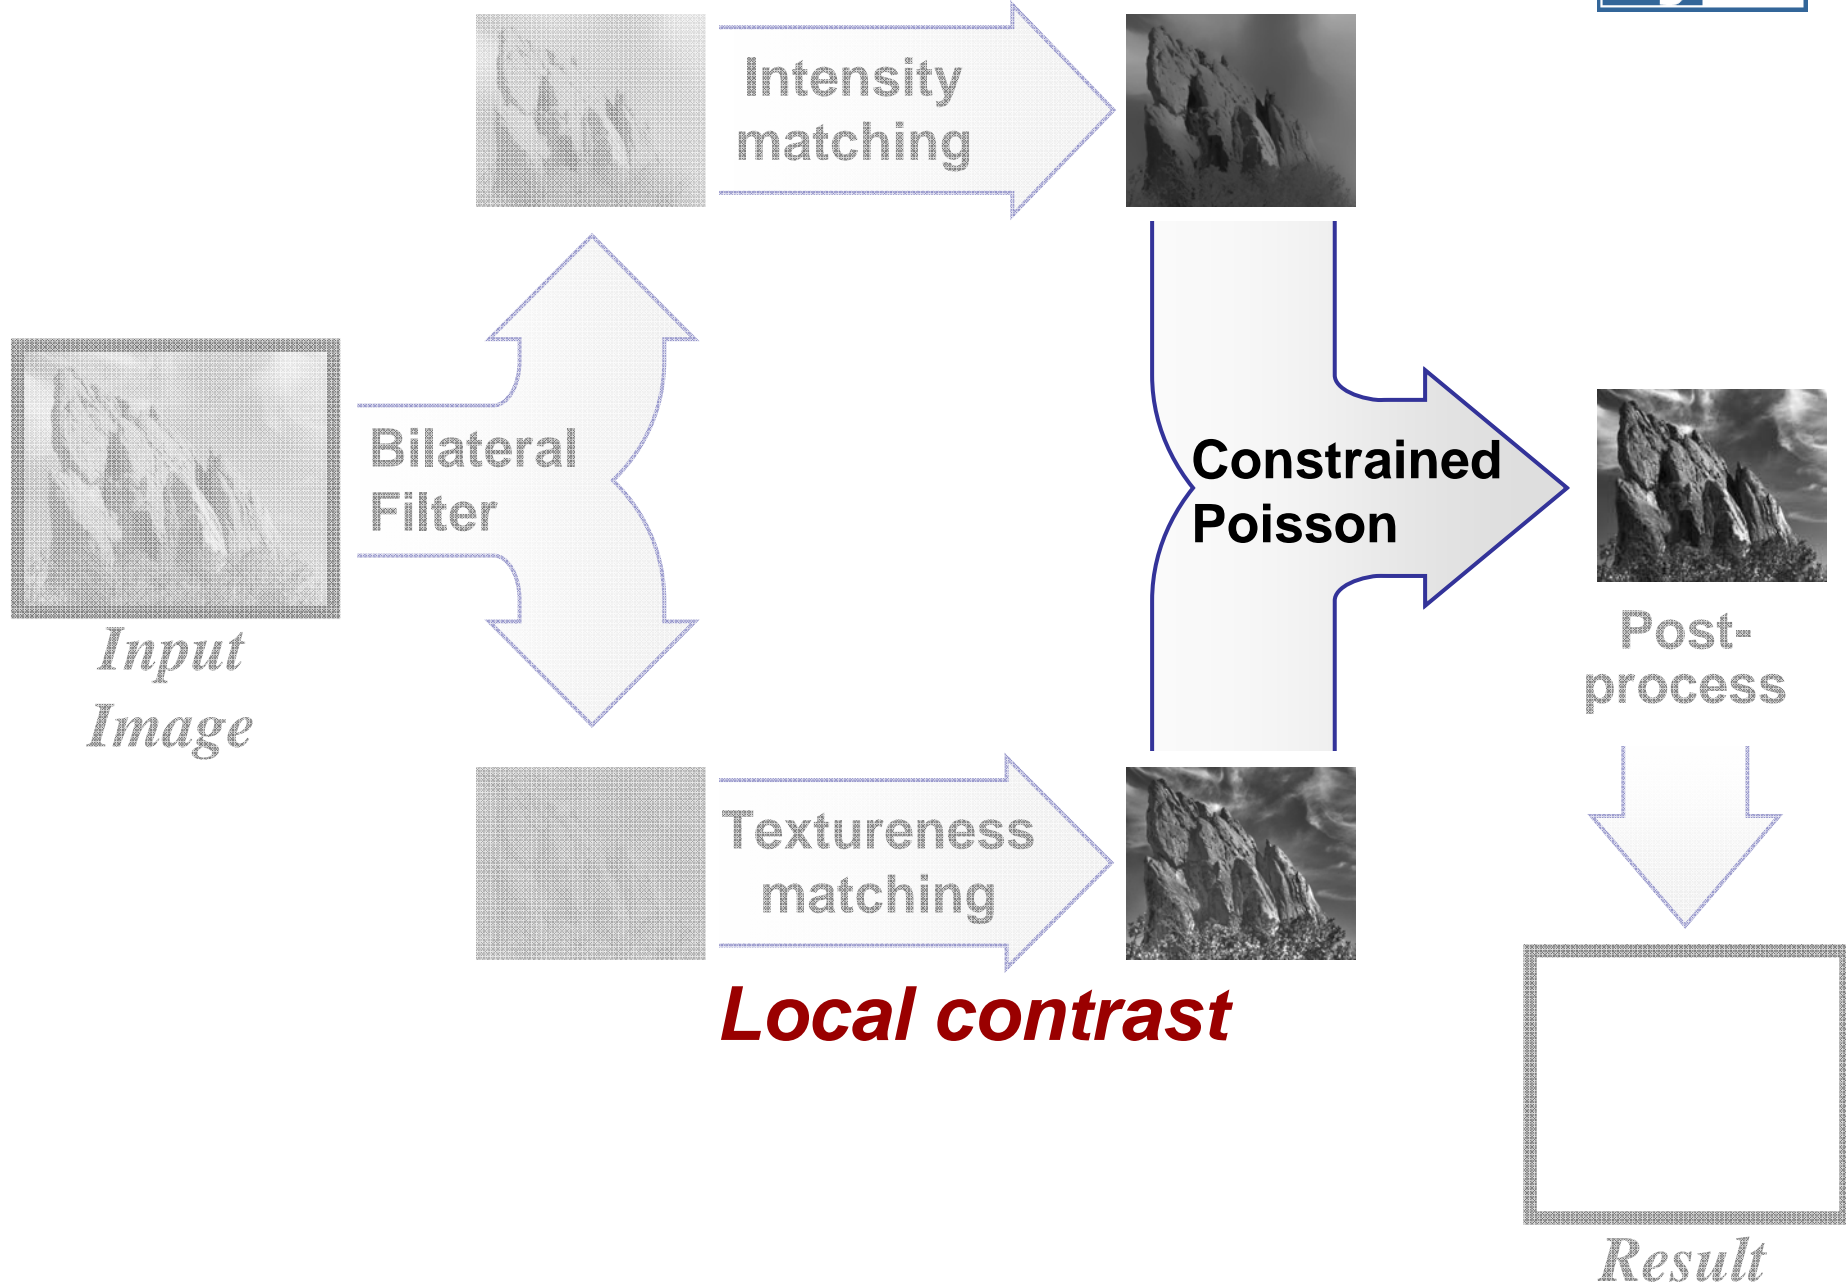
uncorrected result



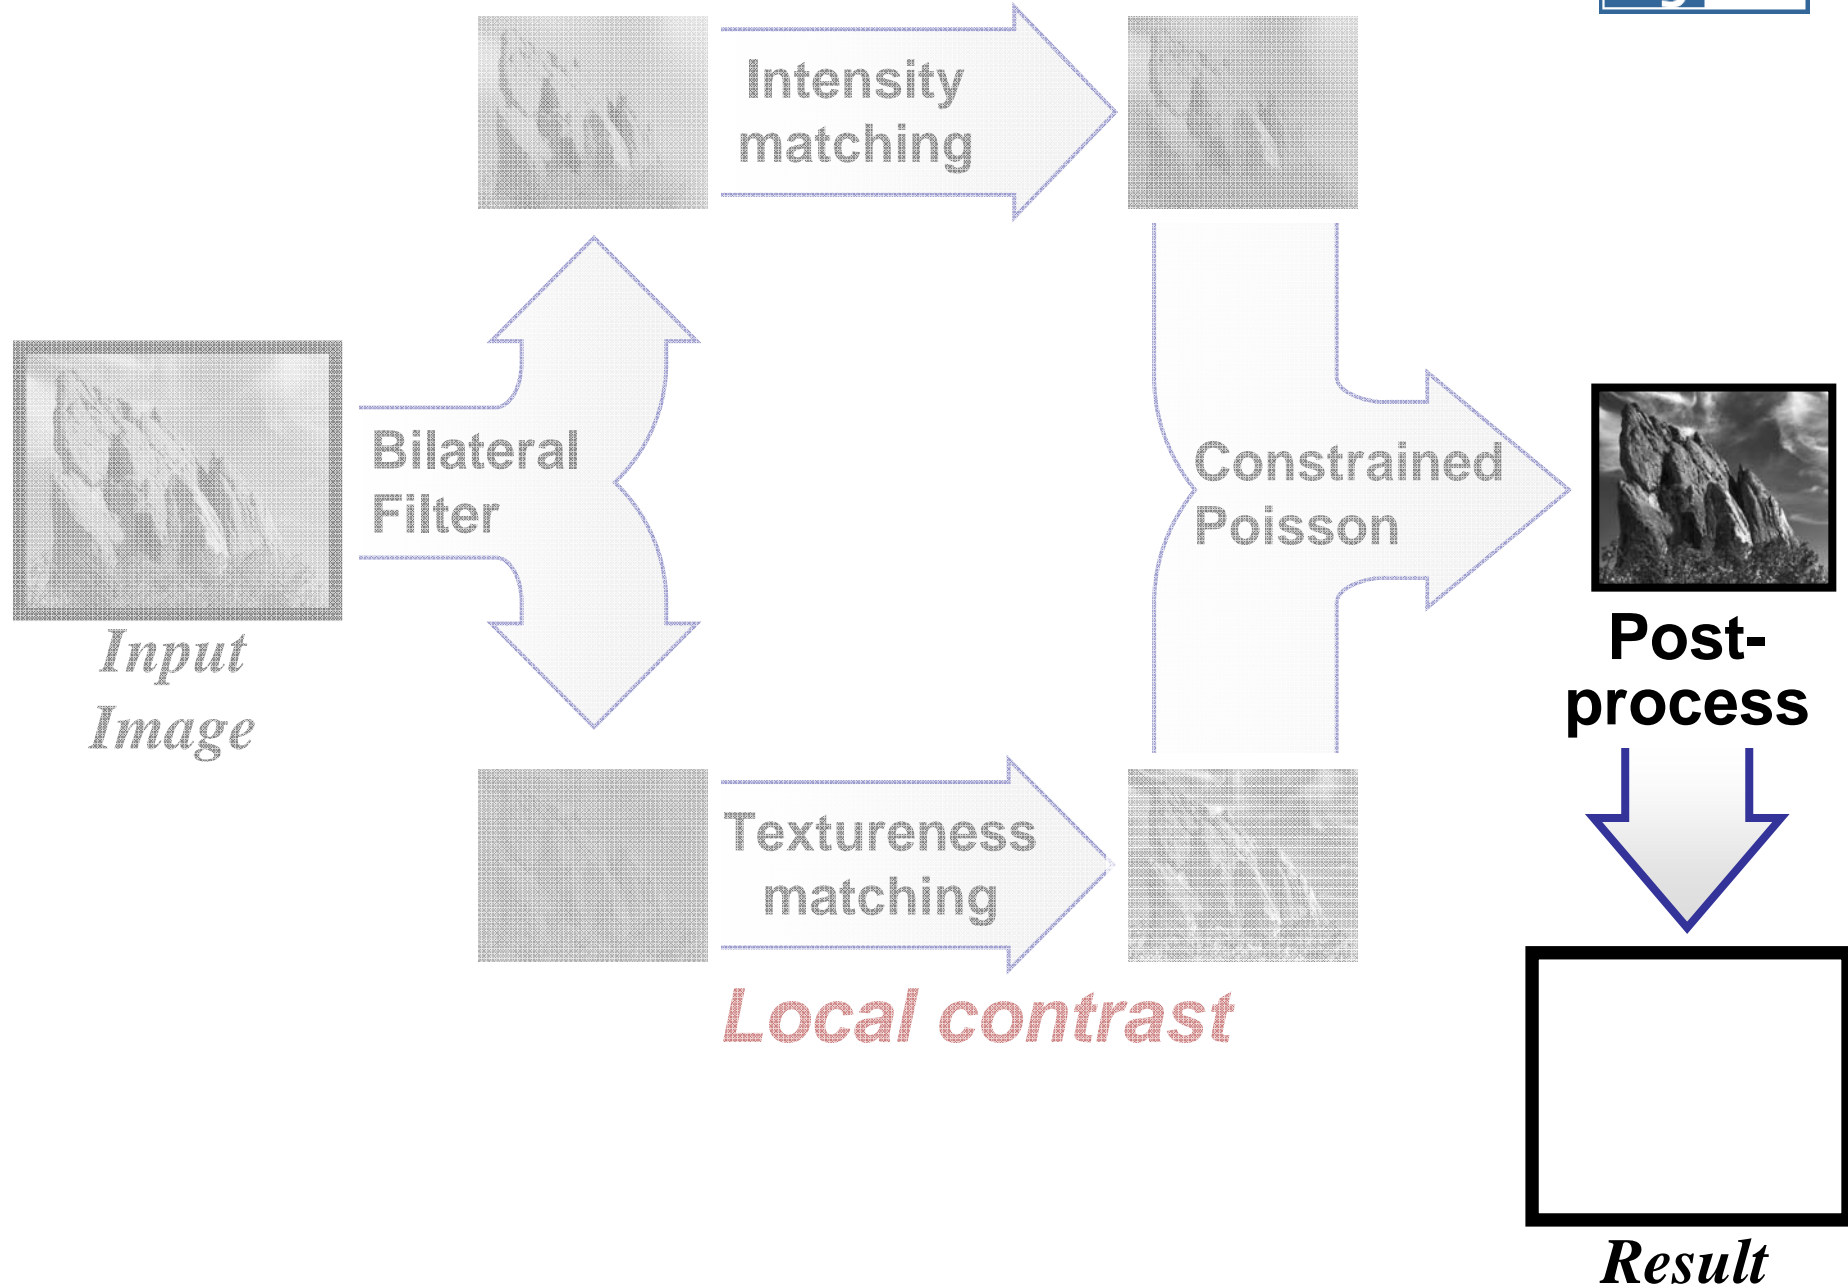
corrected result



# Global contrast



*Global contrast*



# Additional Effects

model



- Soft focus (high frequency manipulation)
- Film grain (texture synthesis [Heeger 95])
- Color toning (chrominance =  $f$ (luminance))

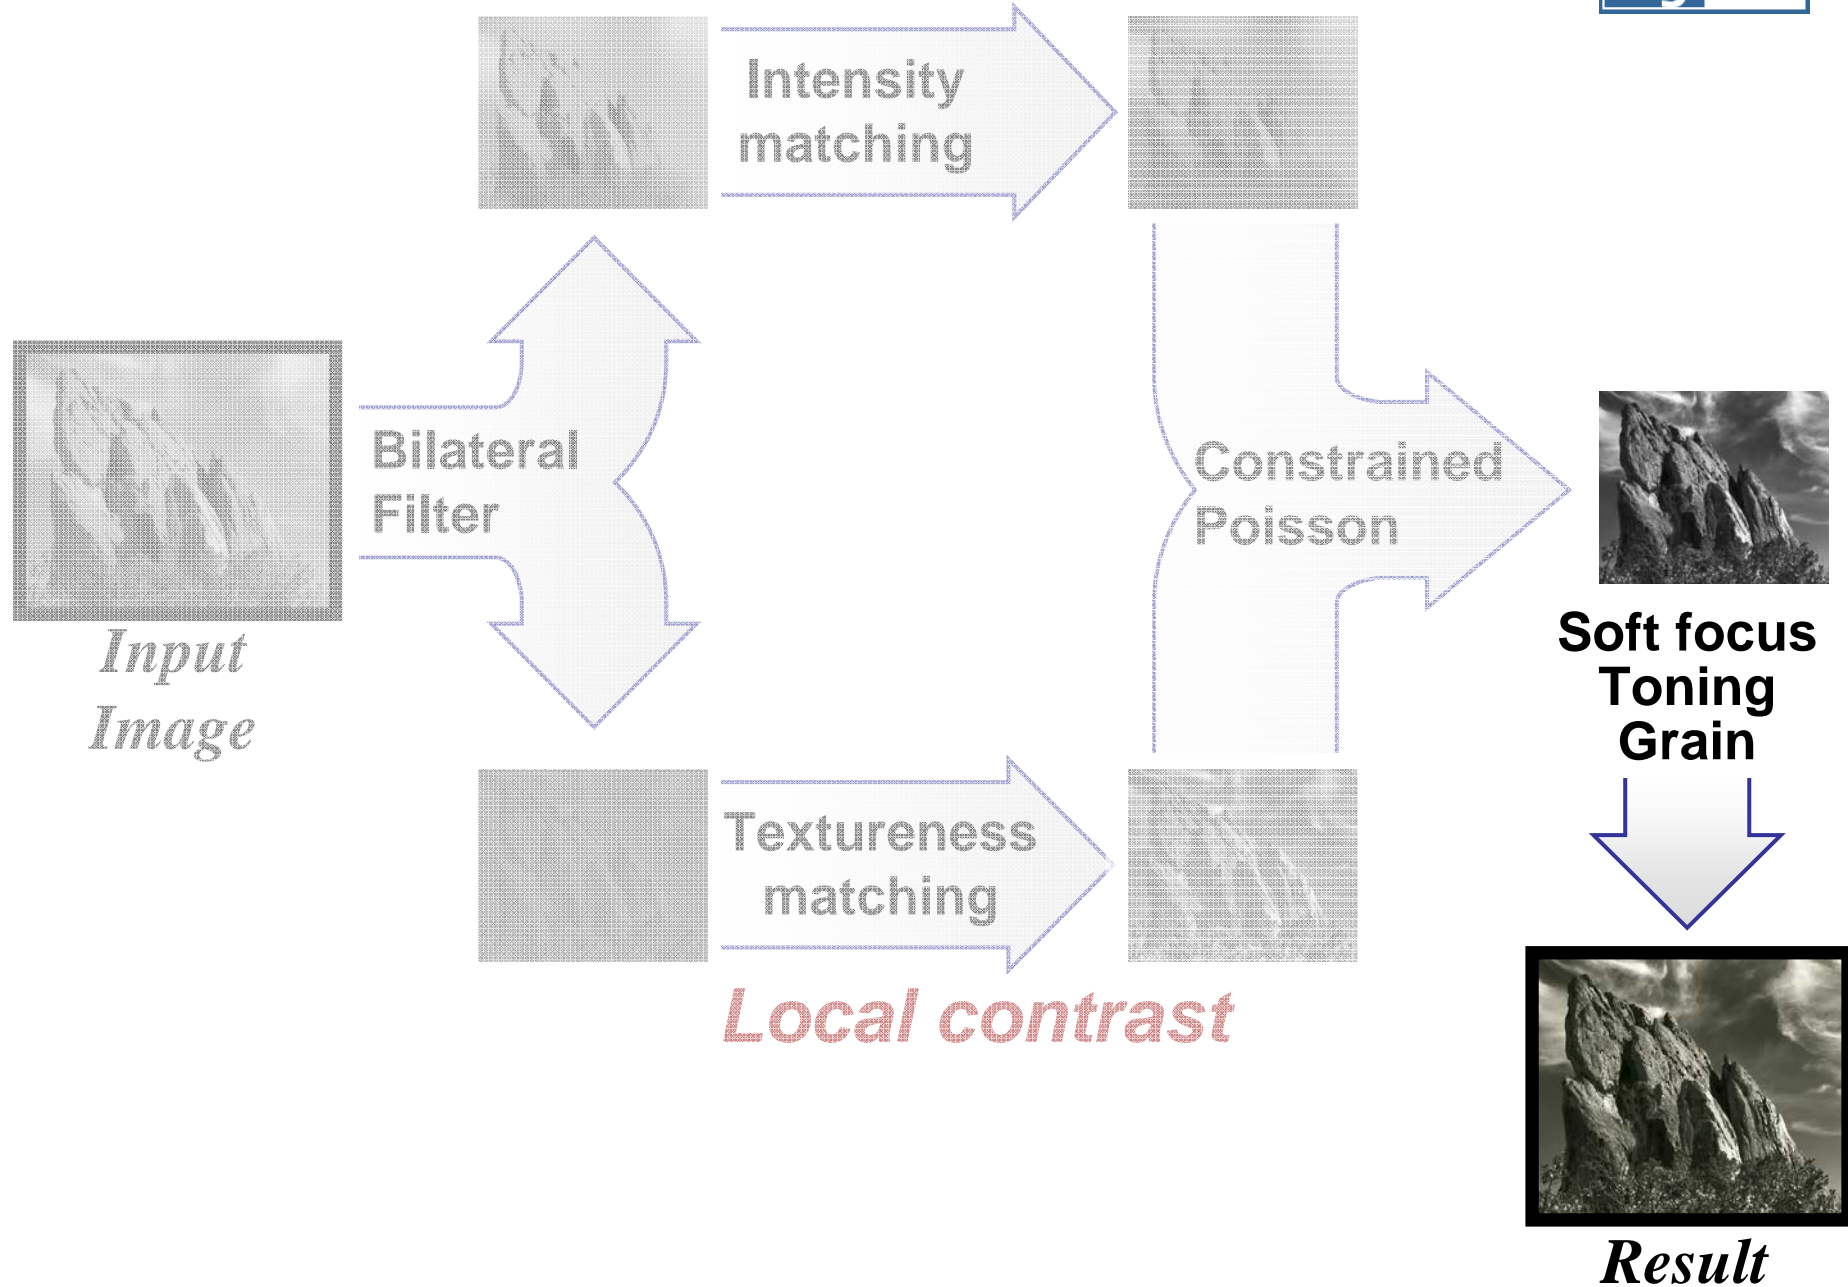
before  
effects



after  
effects



*Global contrast*

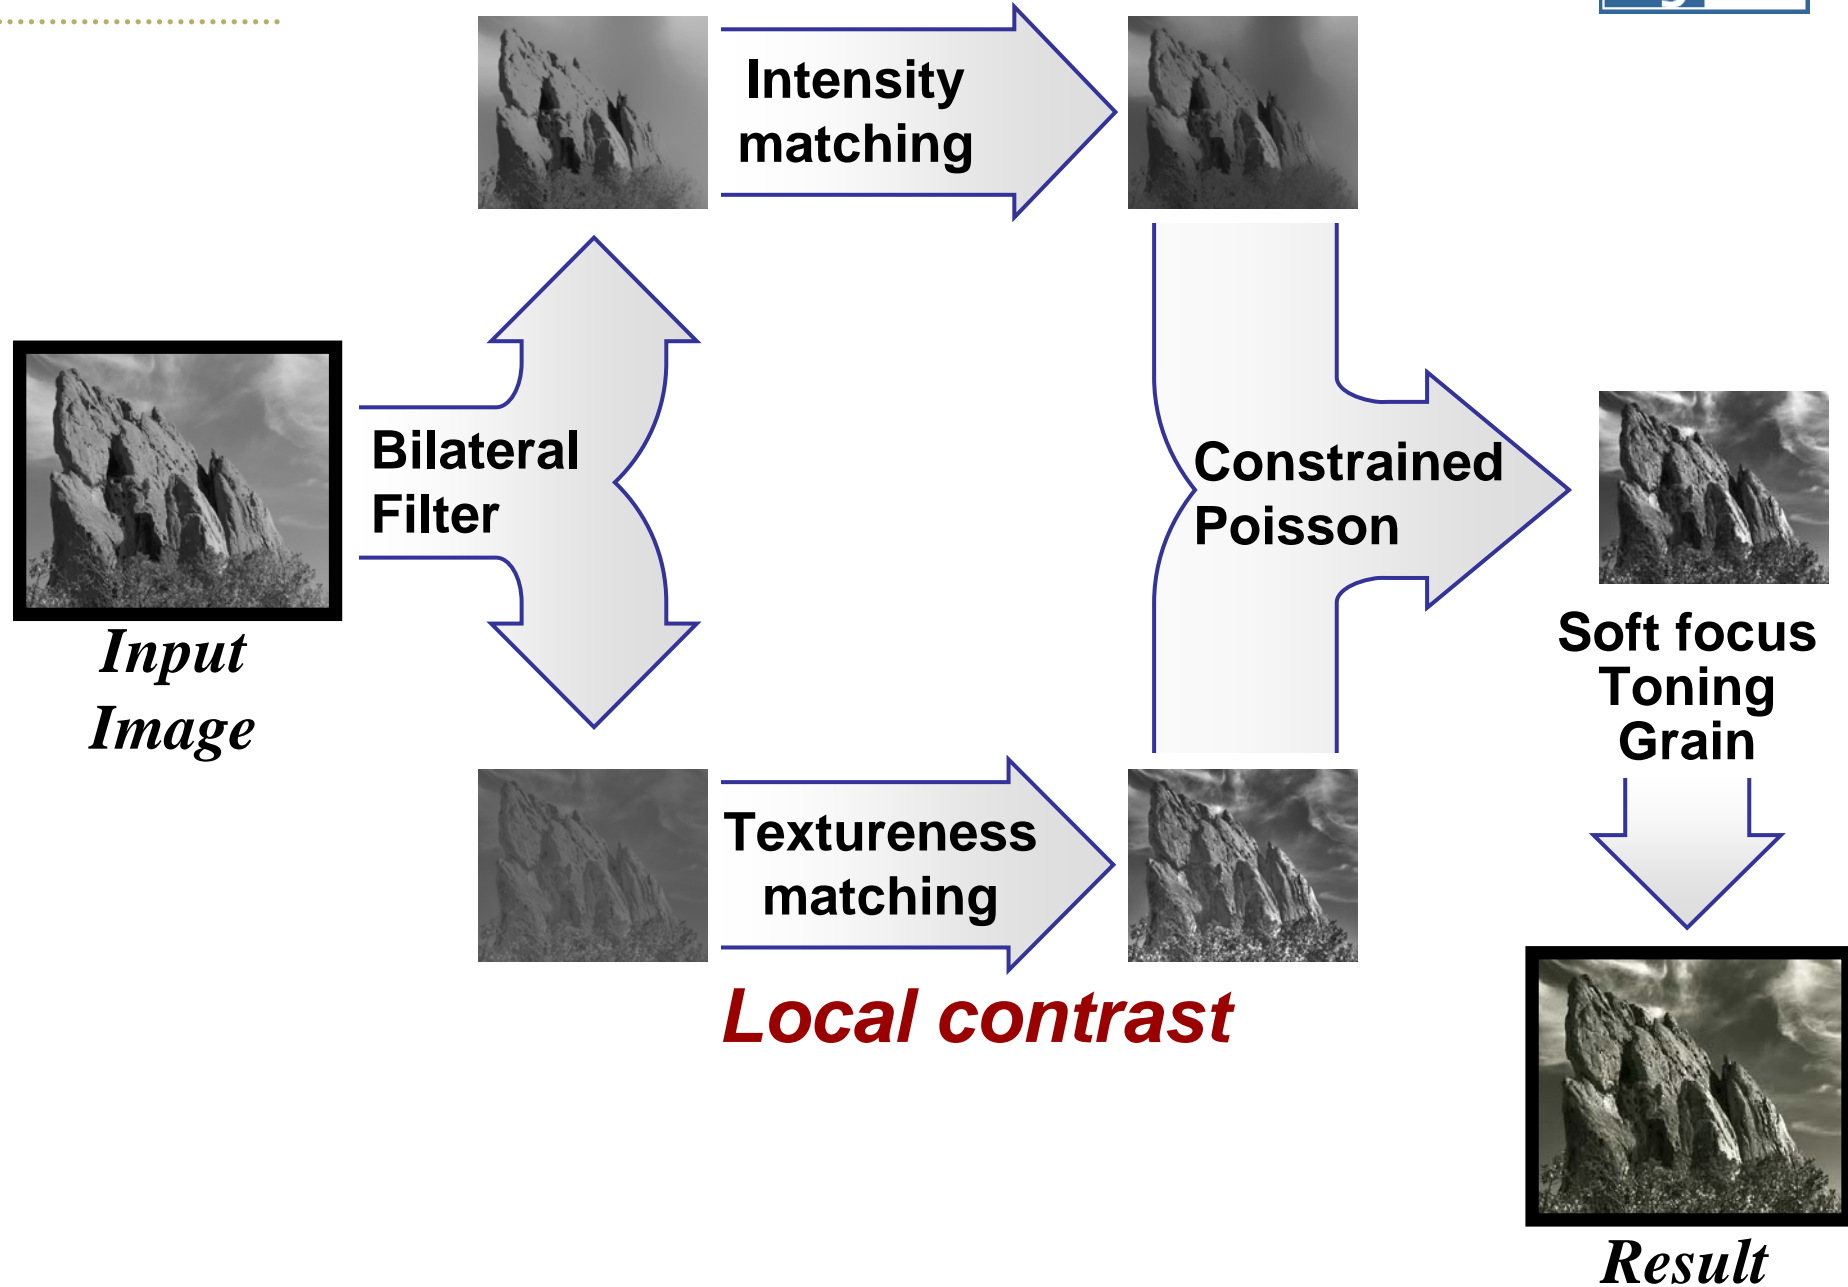


*Local contrast*

*Result*

Recap

**Global contrast**



**Local contrast**

# Results

---

User provides input and model photographs.

→ Our system **automatically** produces the result.

Running times:

- 6 seconds for 1 MPixel or less
- 23 seconds for 4 MPixels
- multi-grid Poisson solver and fast bilateral filter [Paris 06]



# Result

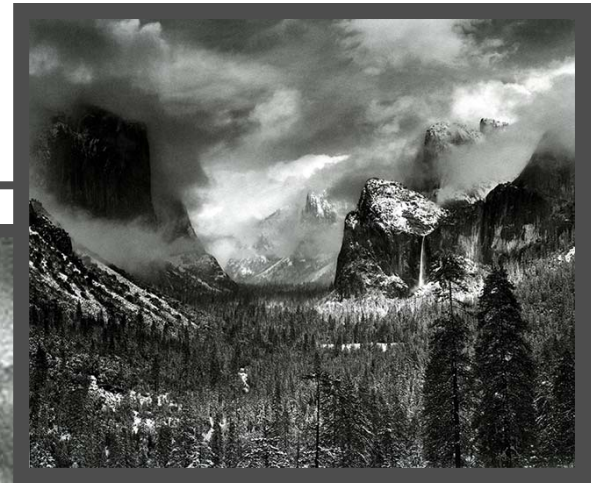
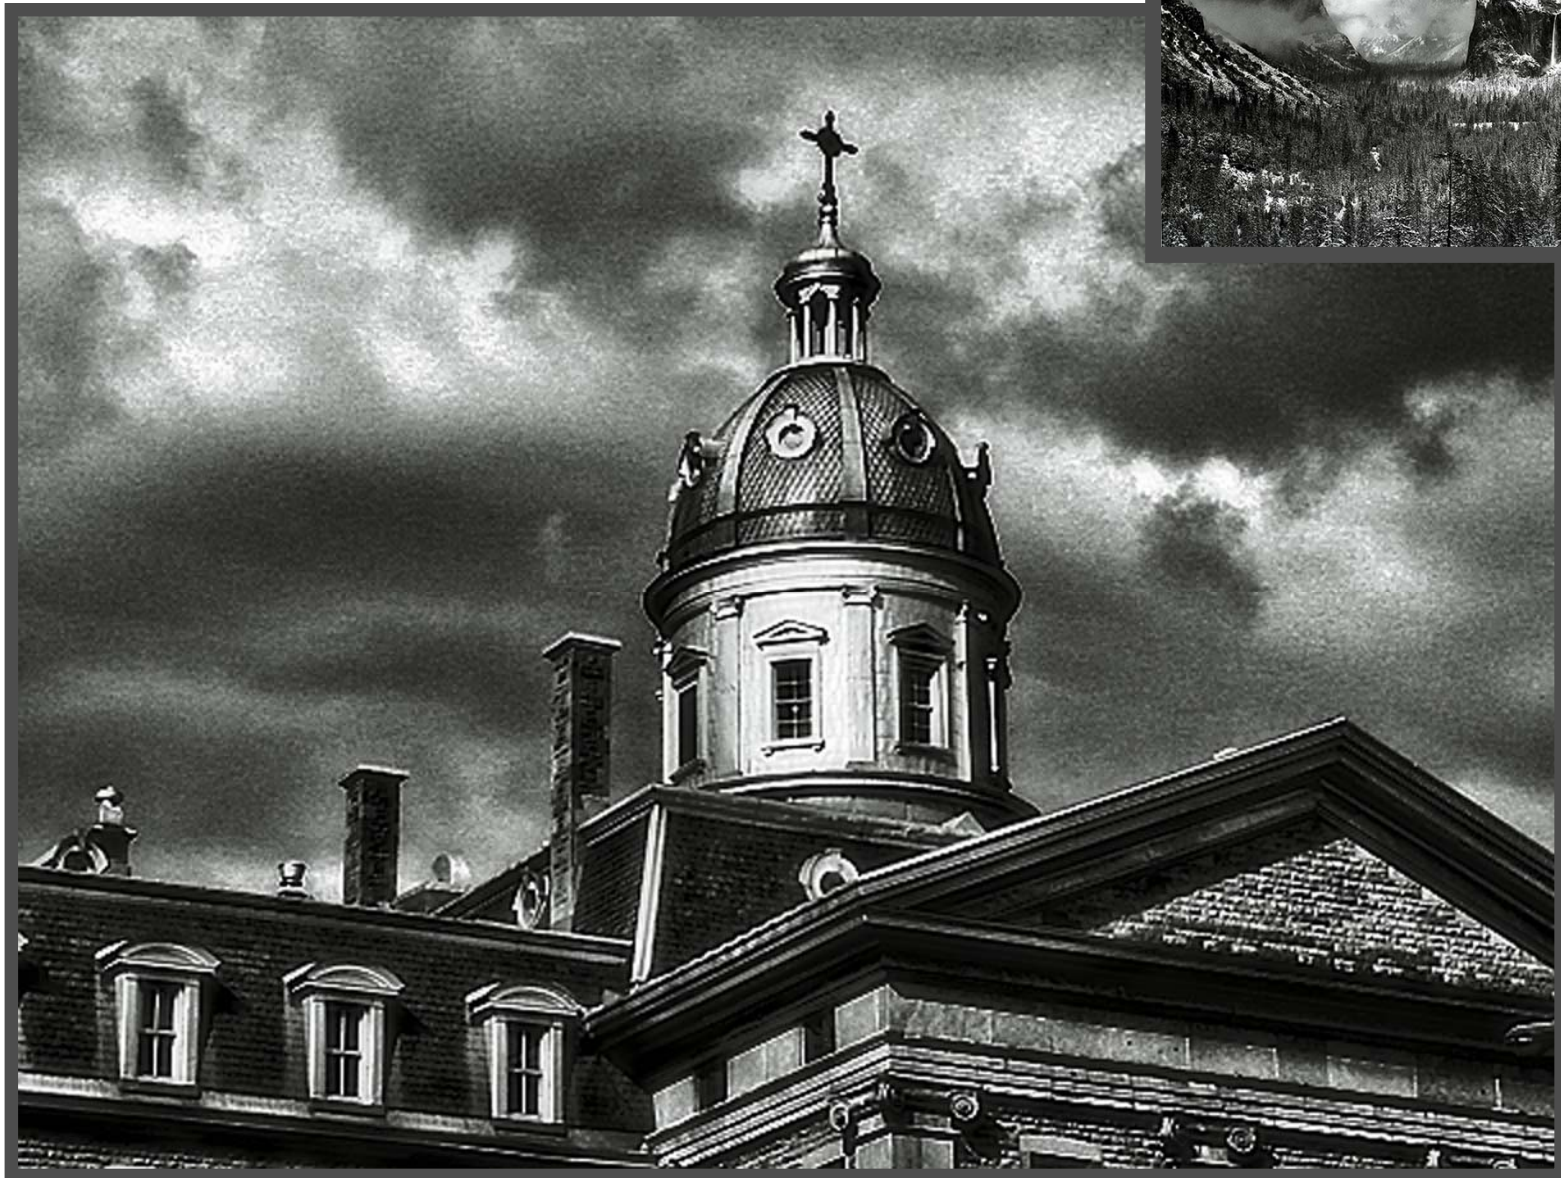
---

# Model



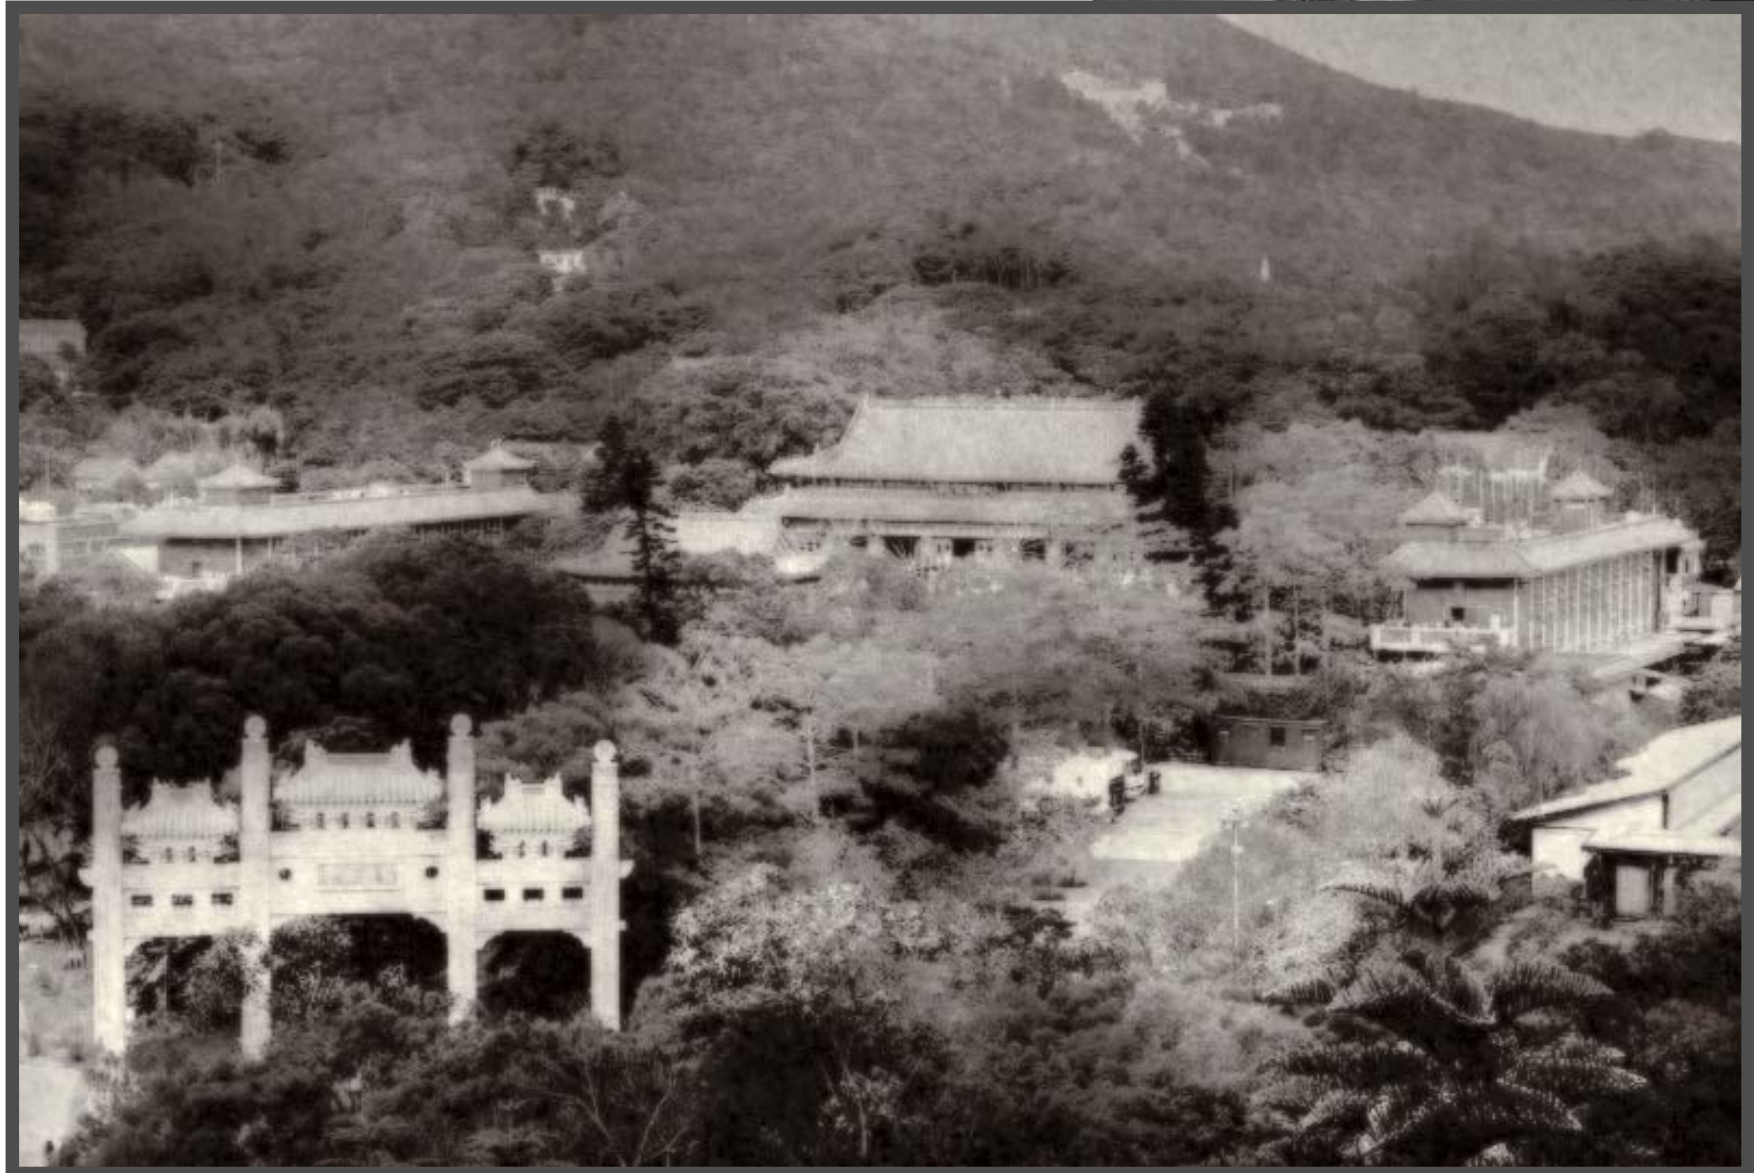
# Result

---

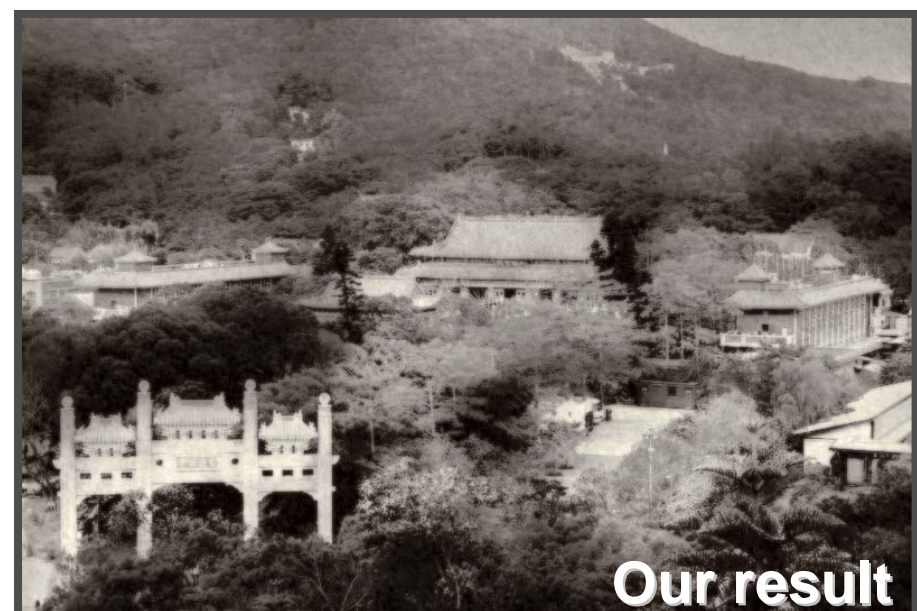
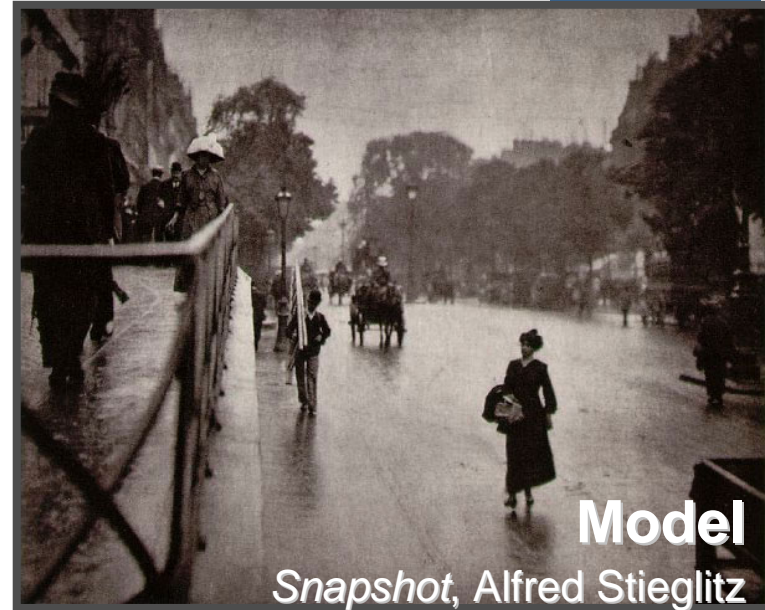


# Result

# Model

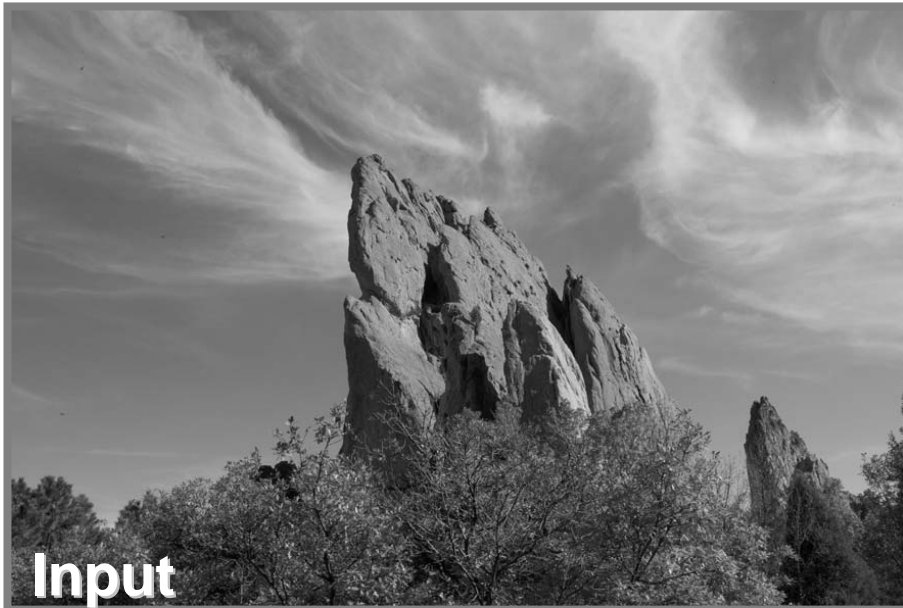


# Comparison with Naïve Histogram Matching



Local contrast, sharpness unfaithful

# Comparison with Naive Histogram Matching

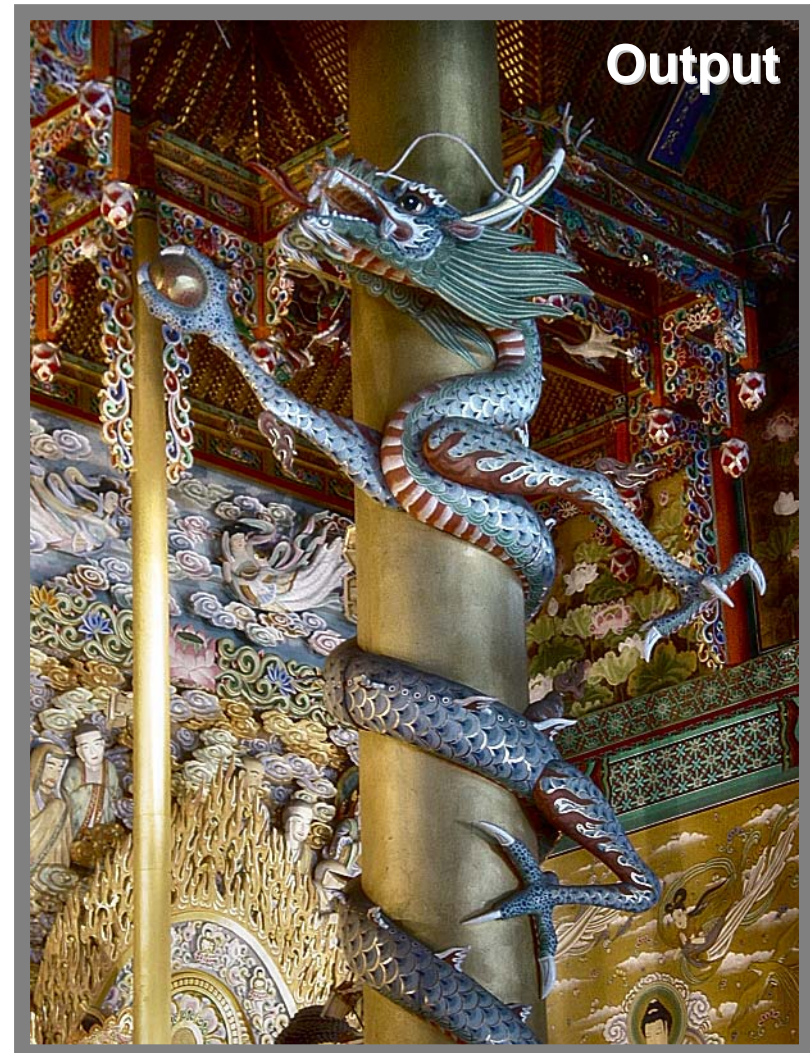


Local contrast too low



# Color Images

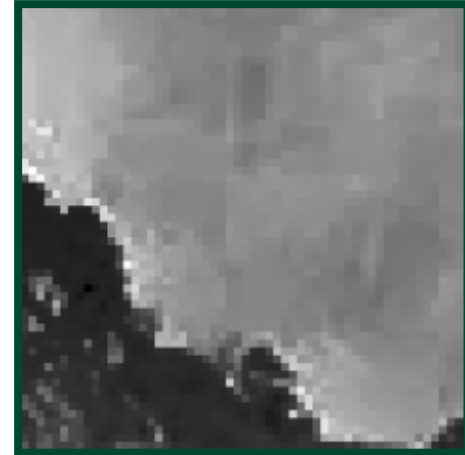
- Lab color space: modify only luminance



# Limitations

---

- Noise and JPEG artifacts
  - amplified defects
  
- Can lead to unexpected results if the image content is too different from the model
  - Portraits, in particular, can suffer



# Conclusions

---

- Transfer “look” from a model photo
- Two-scale tone management
  - Global and local contrast
  - New edge-preserving texture
  - Constrained Poisson reconstruction
  - Additional effects

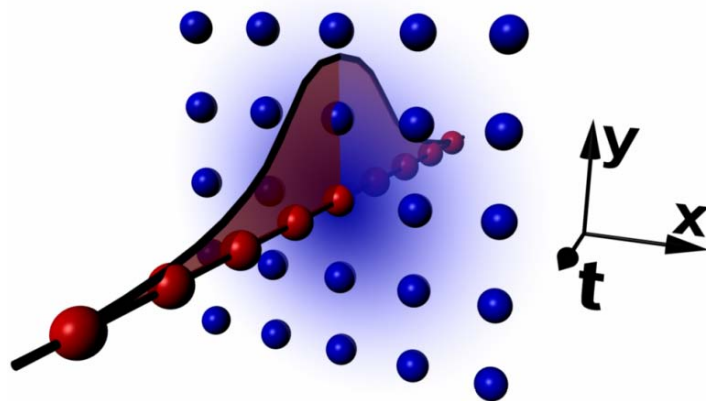


# Video Enhancement Using Per Pixel Exposures (Bennett, 06)



From this video:

ASTA: Adaptive  
Spatio-  
Temporal  
Accumulation Filter



# Joint bilateral filtering

---

$$J_p = \frac{1}{k_p} \sum_{q \in \Omega} I_q f(\|p - q\|) g(\|I_p - I_q\|)$$

$$J_p = \frac{1}{k_p} \sum_{q \in \Omega} I_q f(\|p - q\|) g(\|\tilde{I}_p - \tilde{I}_q\|)$$

# Flash / No-Flash Photo Improvement (Petschnigg04) (Eisemann04)

---

Merge best features: warm, cozy candle light (no-flash)  
low-noise, detailed flash image



# Overview

---

Basic approach of both flash/noflash papers

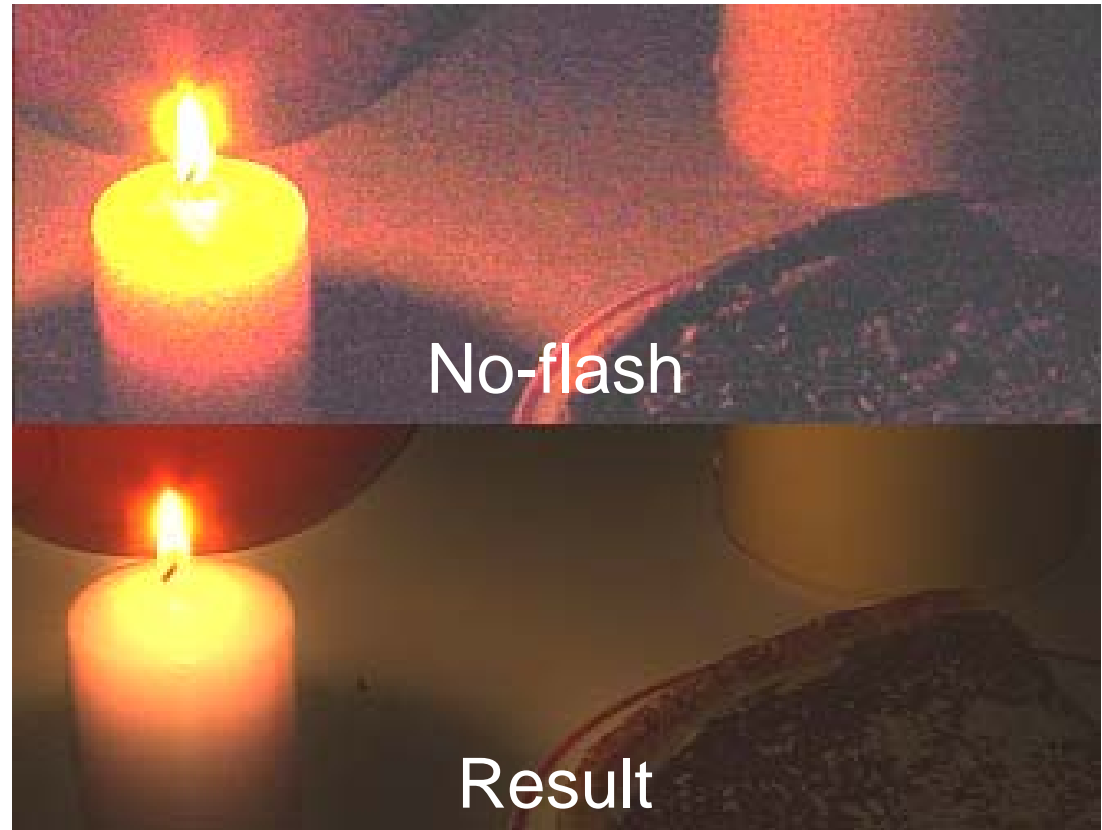
Remove noise + details from image A,

Keep as image A Lighting

-----

Obtain noise-free details from image B,

Discard Image B Lighting



# Petschnigg:

---

- Flash



# Petschnigg:

---

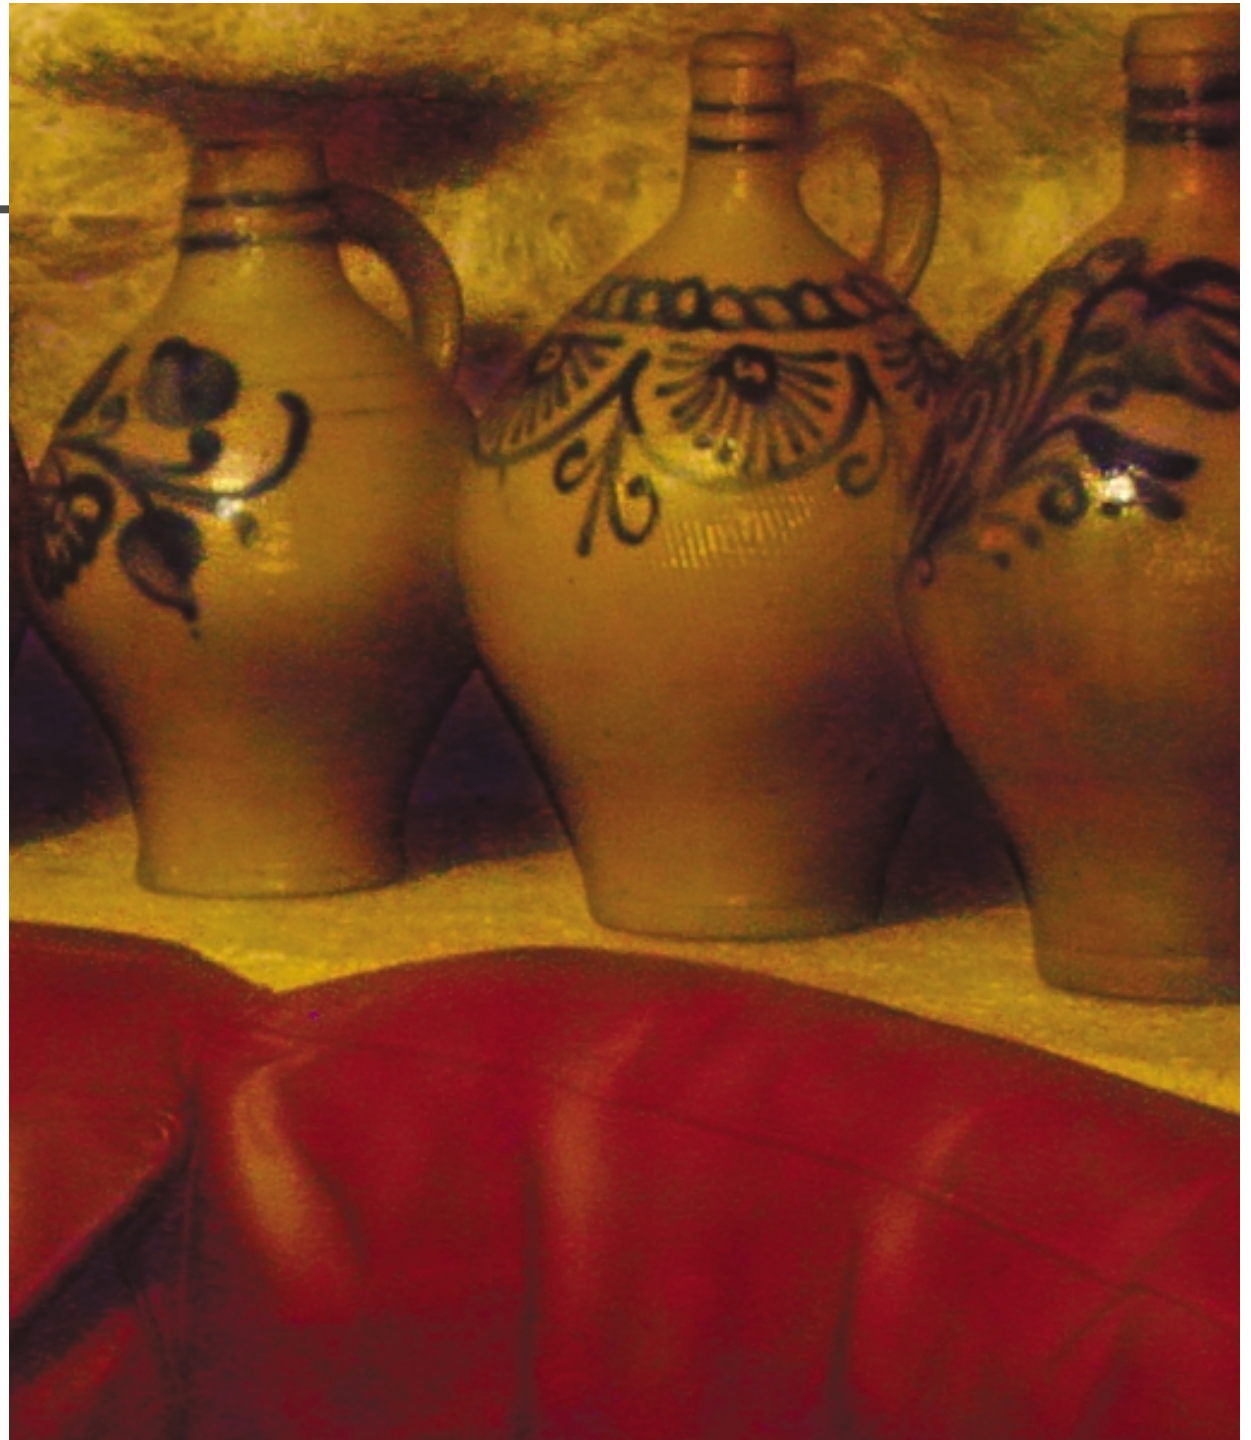
- No Flash,



# Petschnigg:

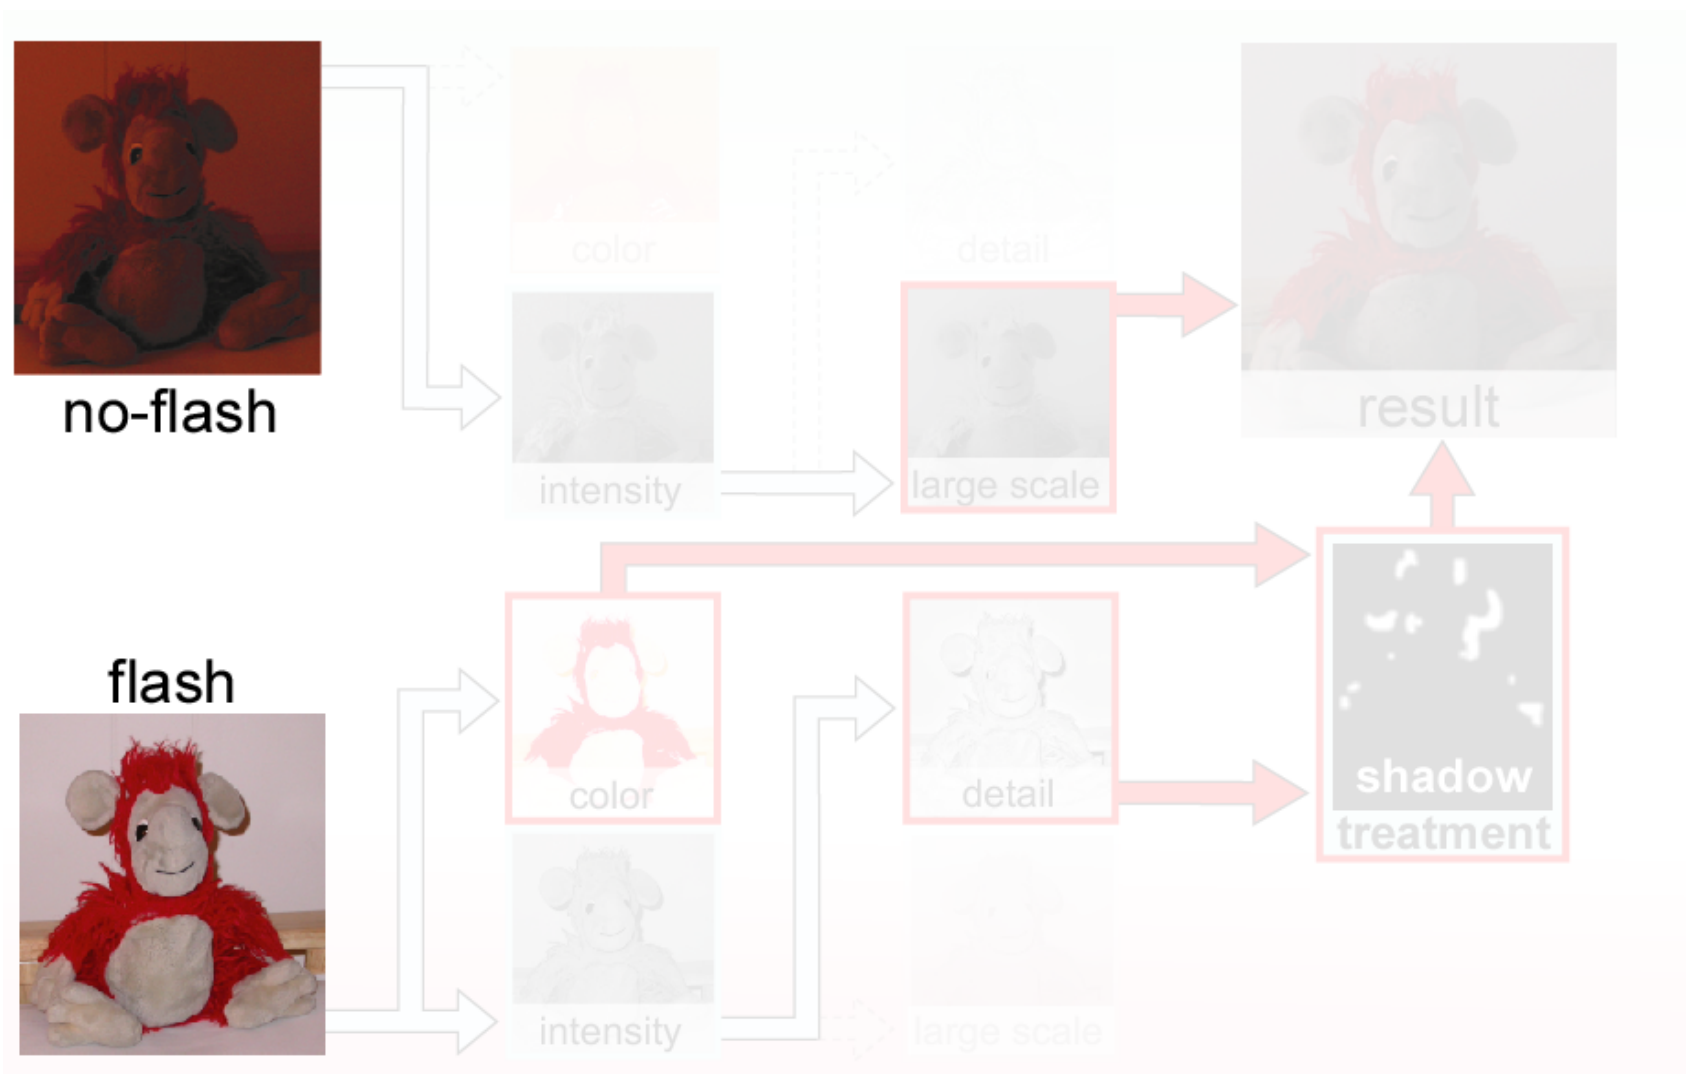
---

- Result



# Our Approach

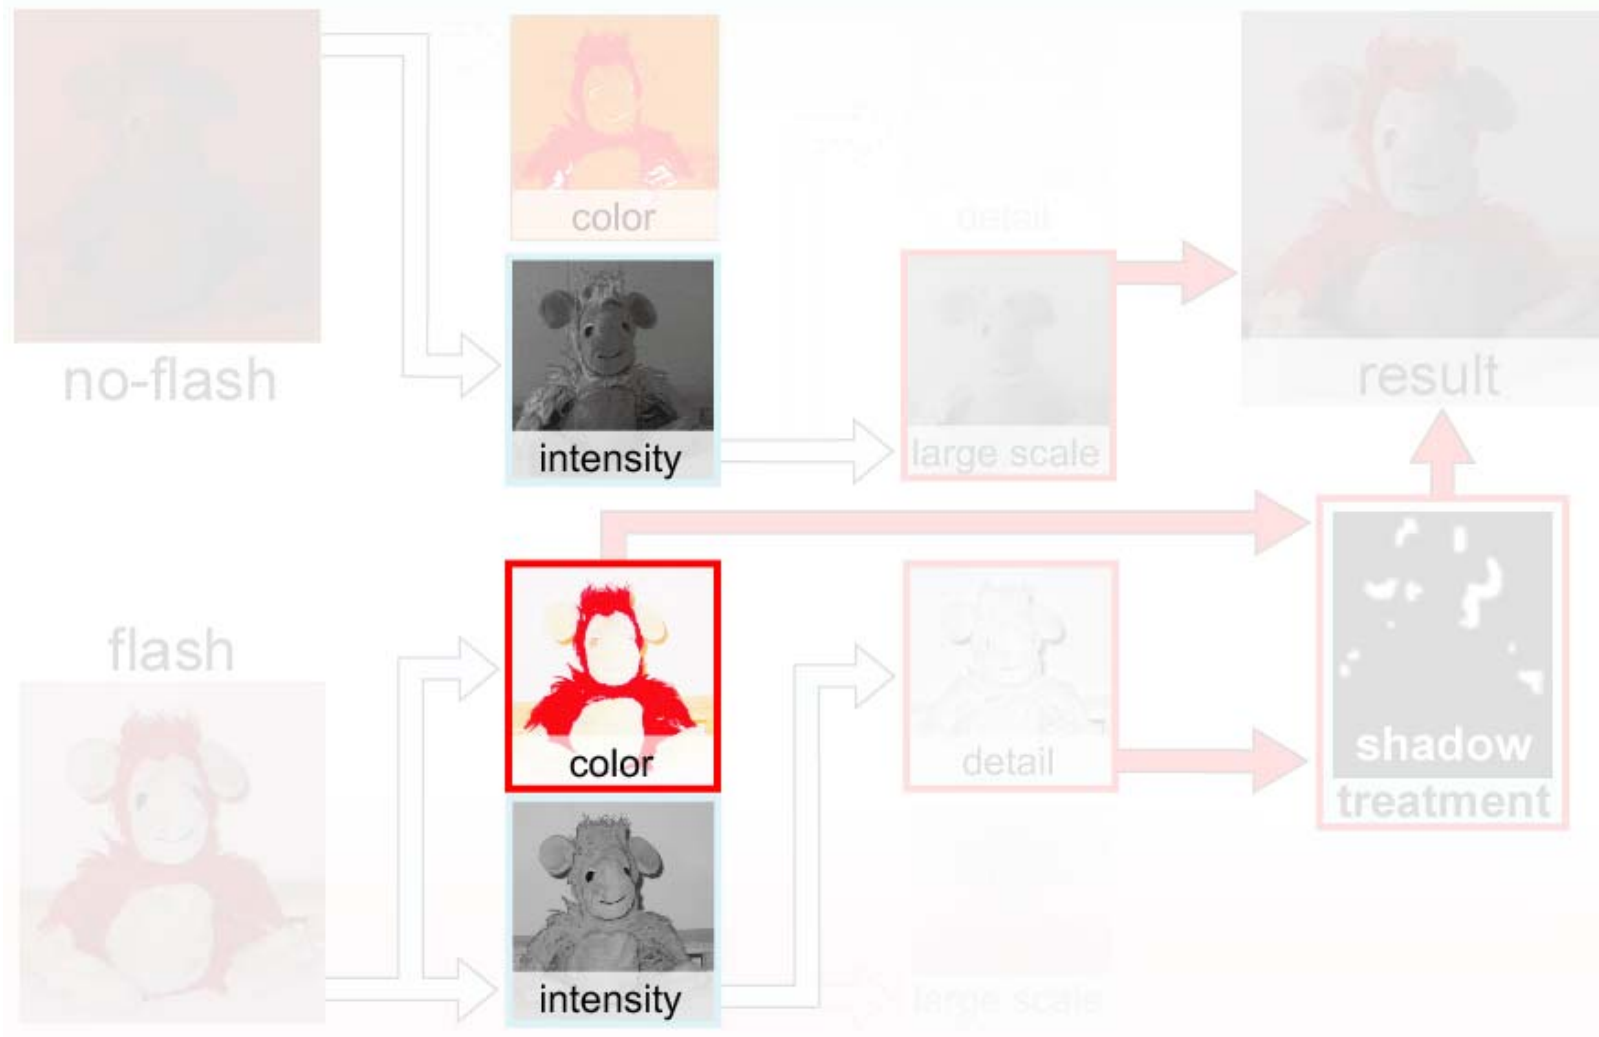
## Registration





# Our Approach

## Decomposition



# Decomposition

---

Color / Intensity:



original

=



intensity

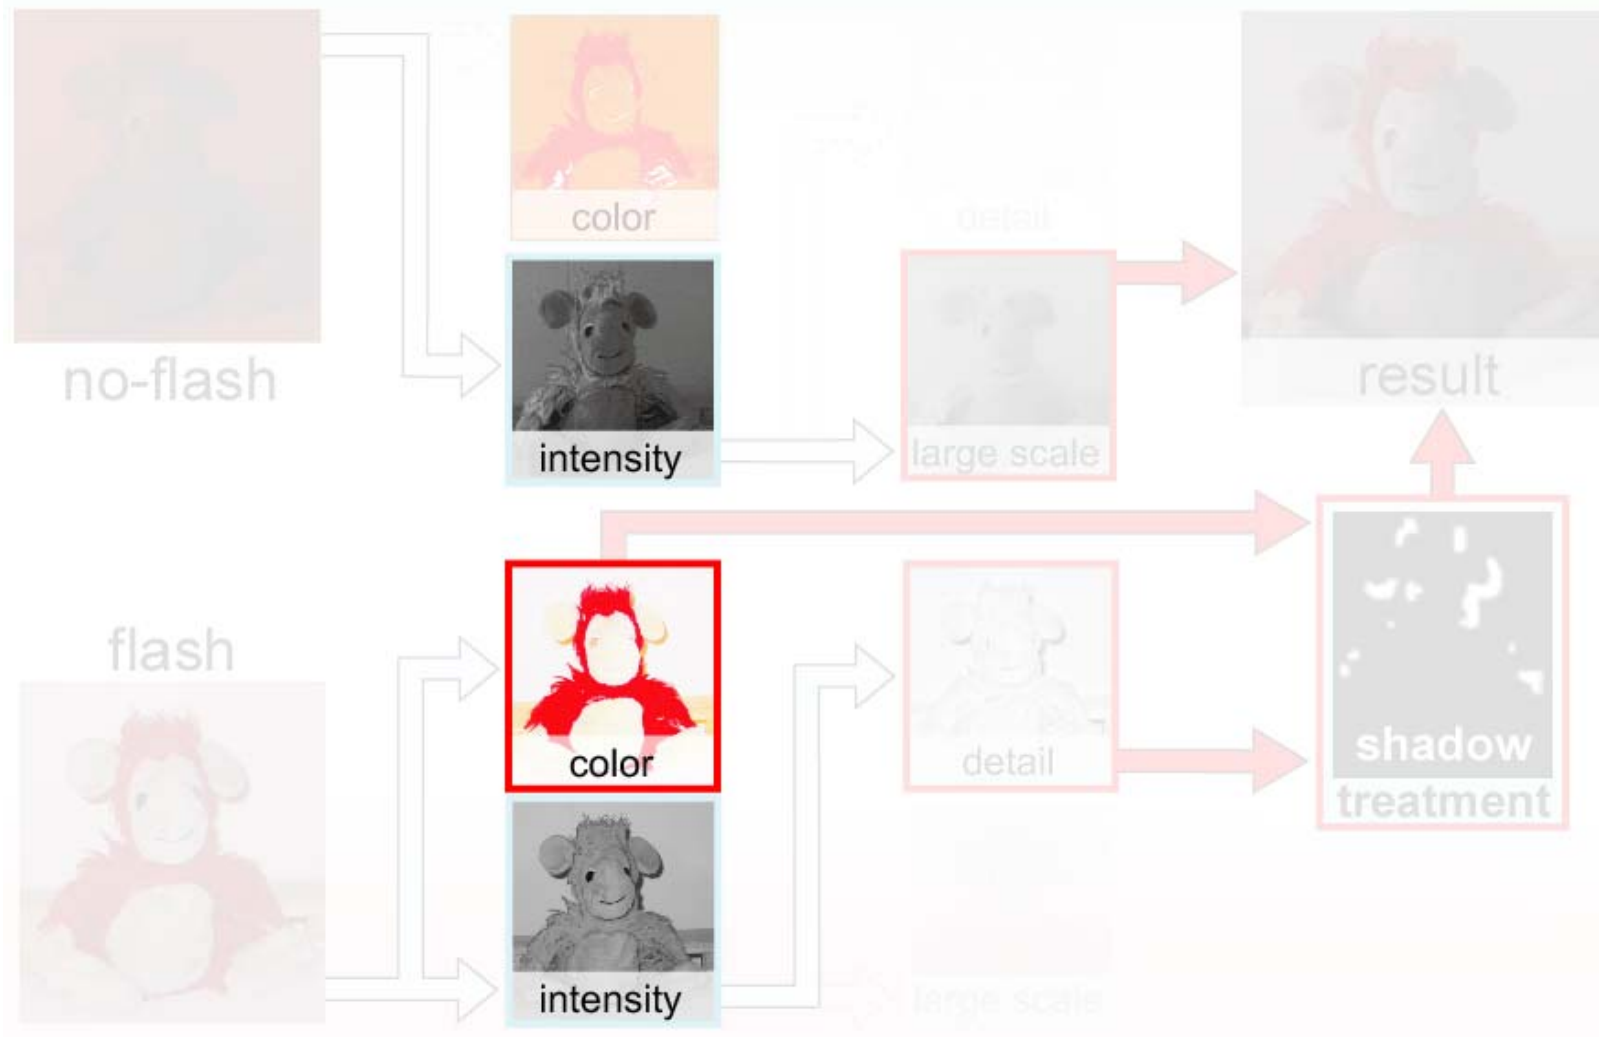
\*



color

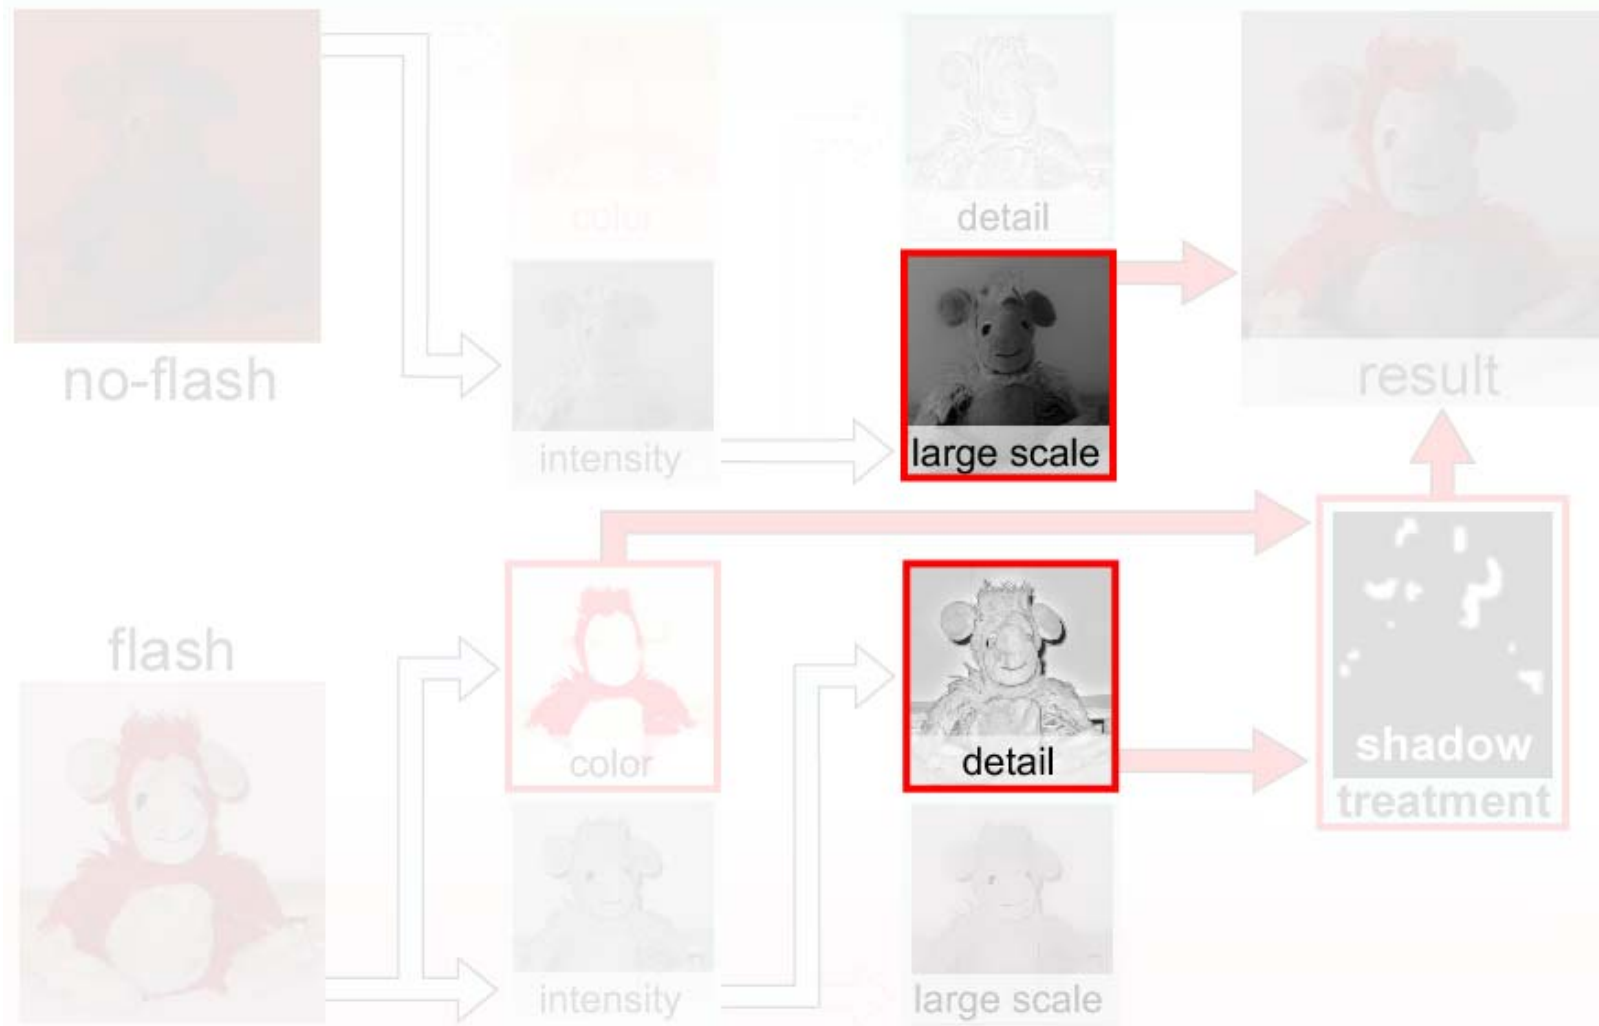
# Our Approach

## Decomposition



# Our Approach

## Decoupling



# Decoupling

---

- Lighting : Large-scale variation
- ~~Texture : Small-scale variation~~  
• Lighting : Large-scale variation
- Texture : Small-scale variation



Lighting

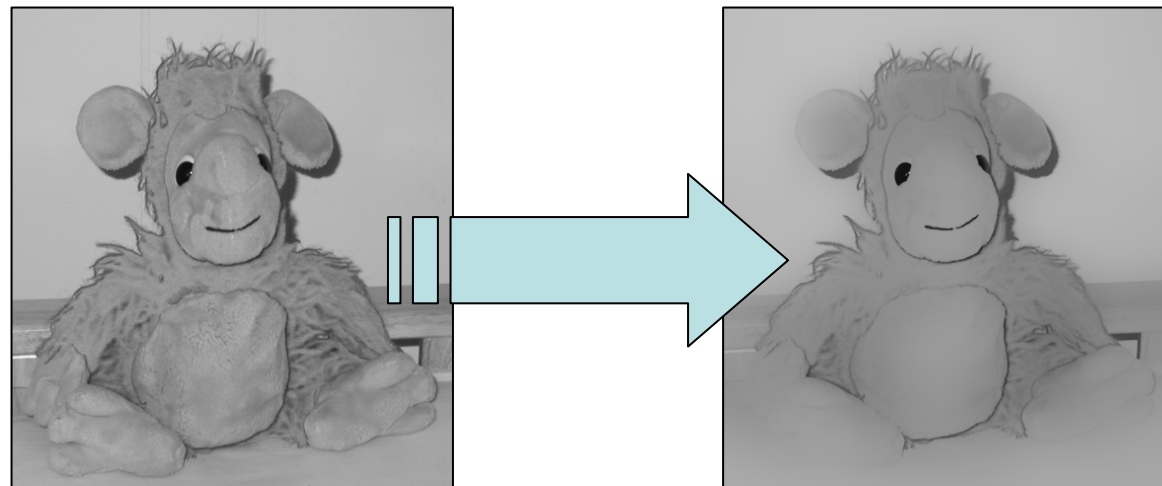
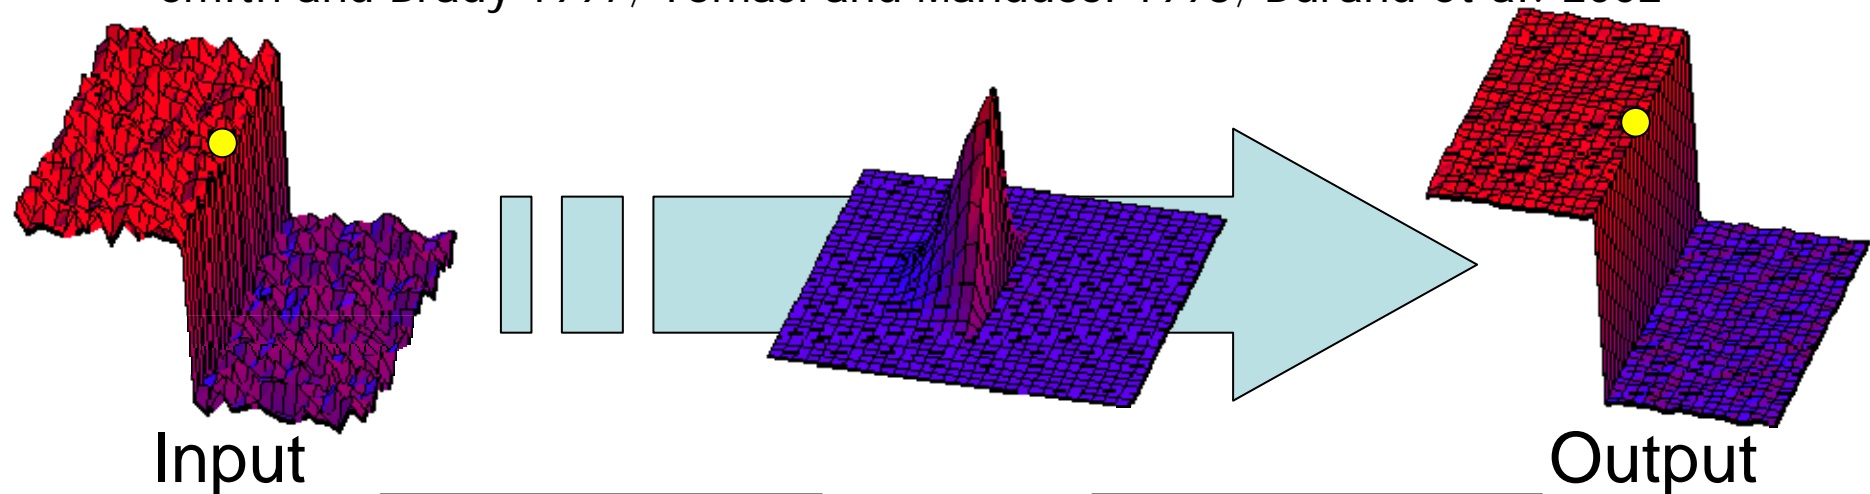


Texture

# Large-scale Layer

- **Bilateral filter** – edge preserving filter

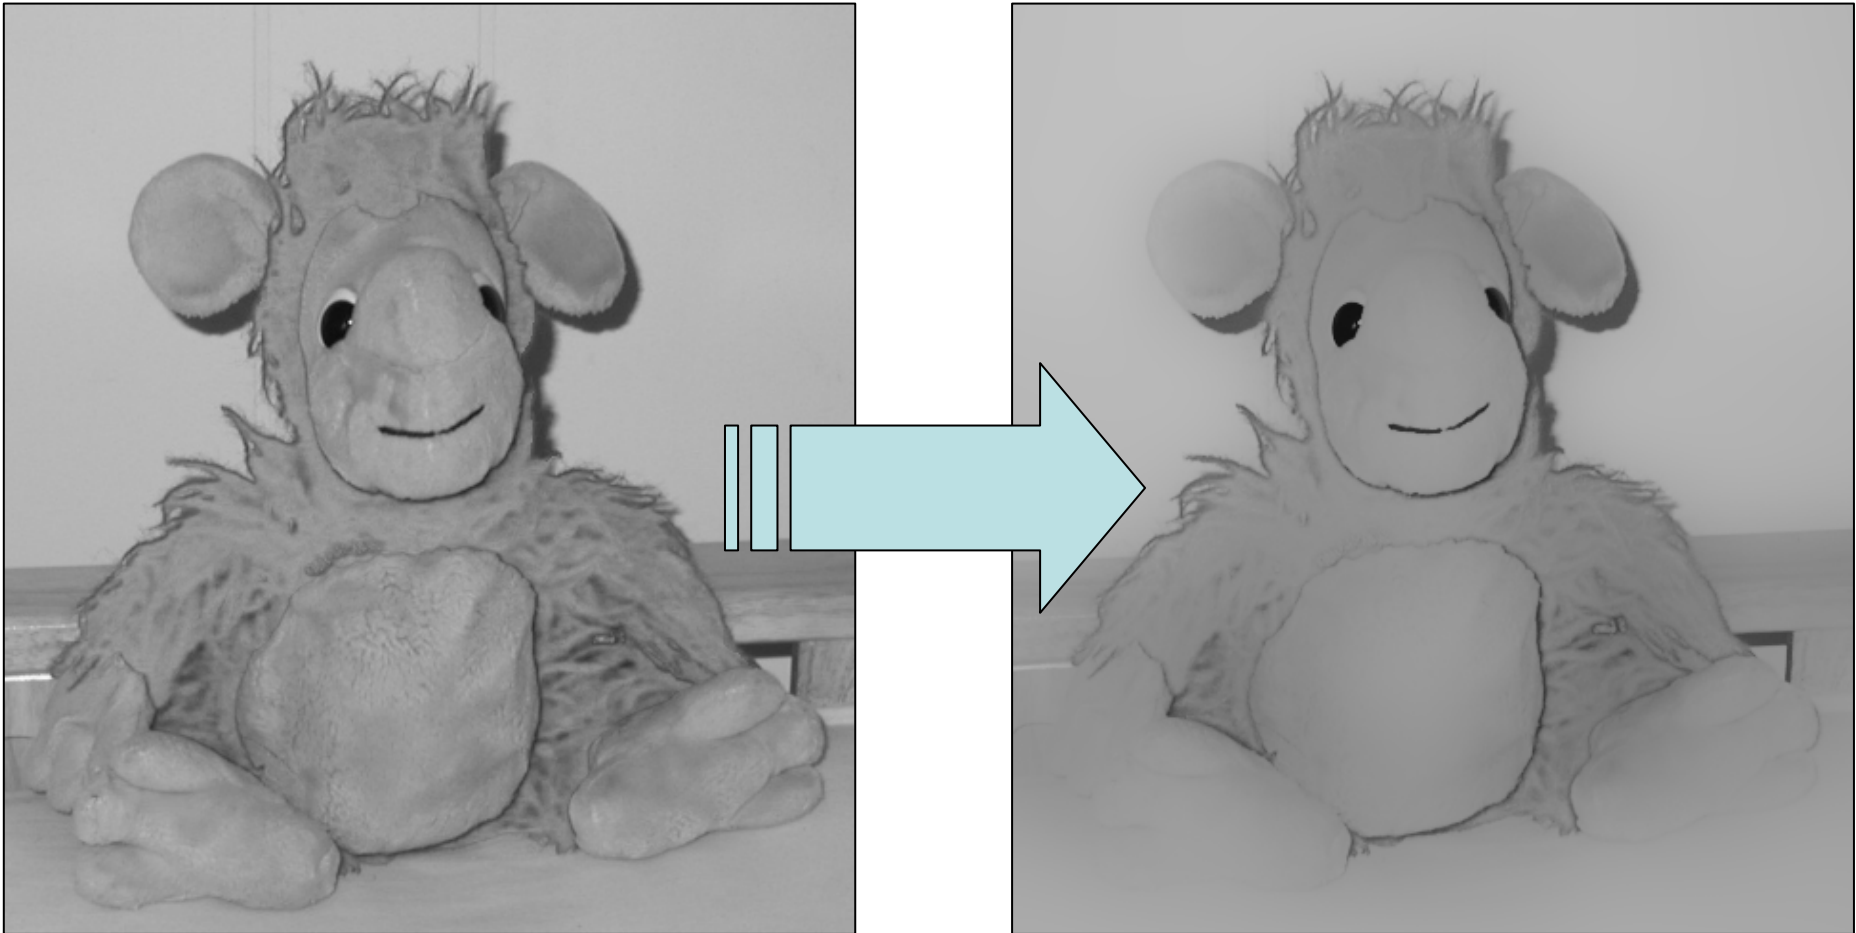
Smith and Brady 1997; Tomasi and Manducci 1998; Durand et al. 2002



# Large-scale Layer

---

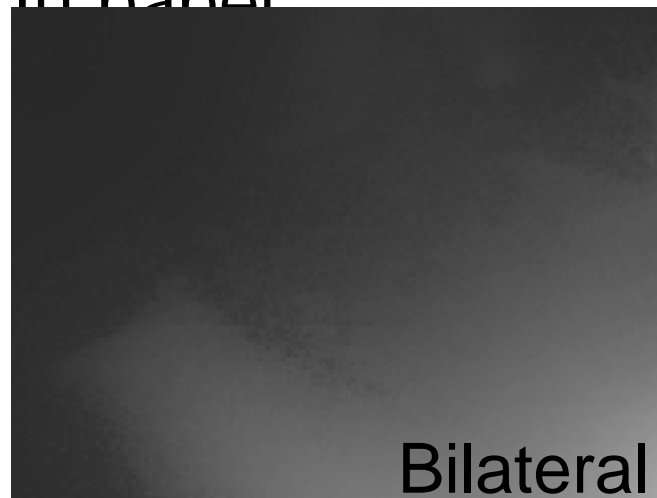
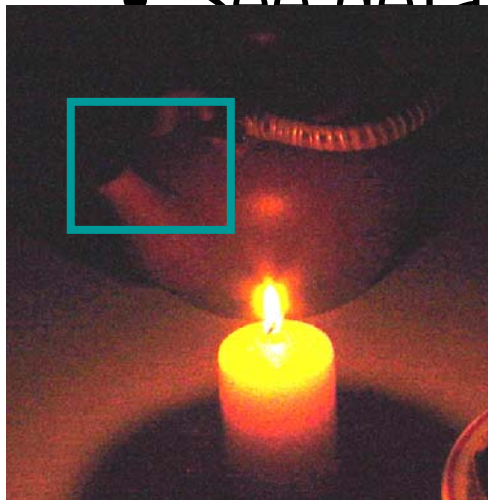
- Bilateral filter



# Cross Bilateral Filter

---

- Similar to joint bilateral filter by Petschnigg et al.
- When no-flash image is too noisy
- Borrow similarity from flash image
  - edge stopping from flash image
- See detail in paper



Bilateral



Cross Bilateral



# Detail Layer

---



Intensity



Large-scale



Detail

Recombination: Large scale \* Detail = Intensity

# Recombination

---



Large-scale  
No-flash

\*



Detail  
Flash

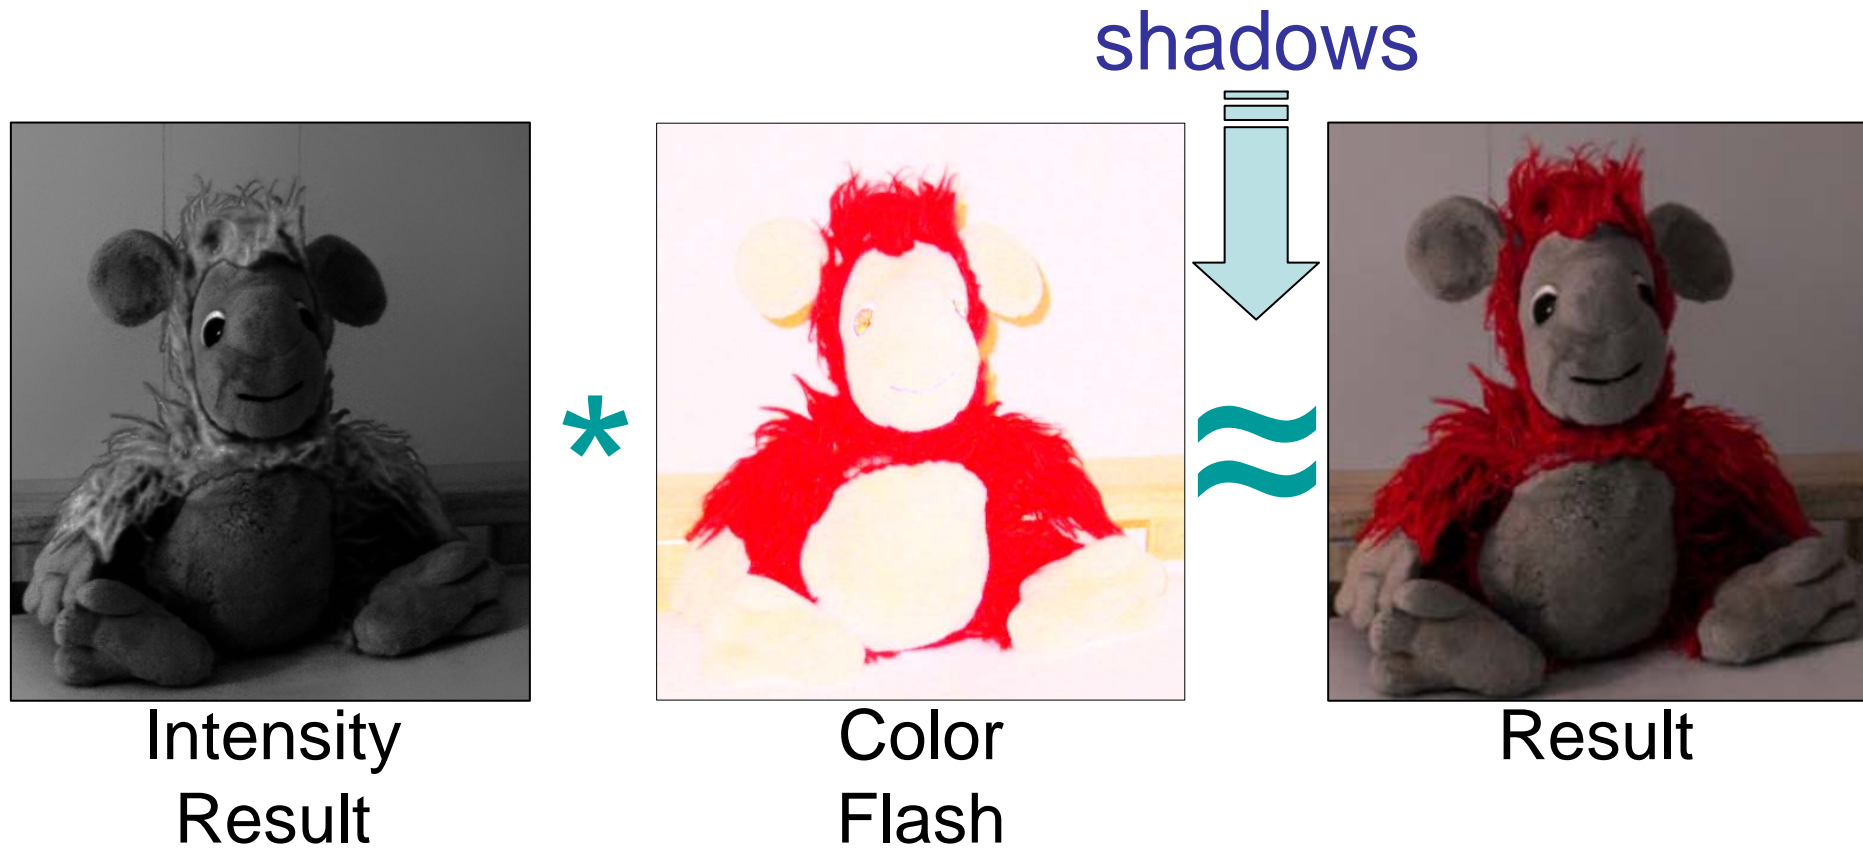
=



Intensity  
Result

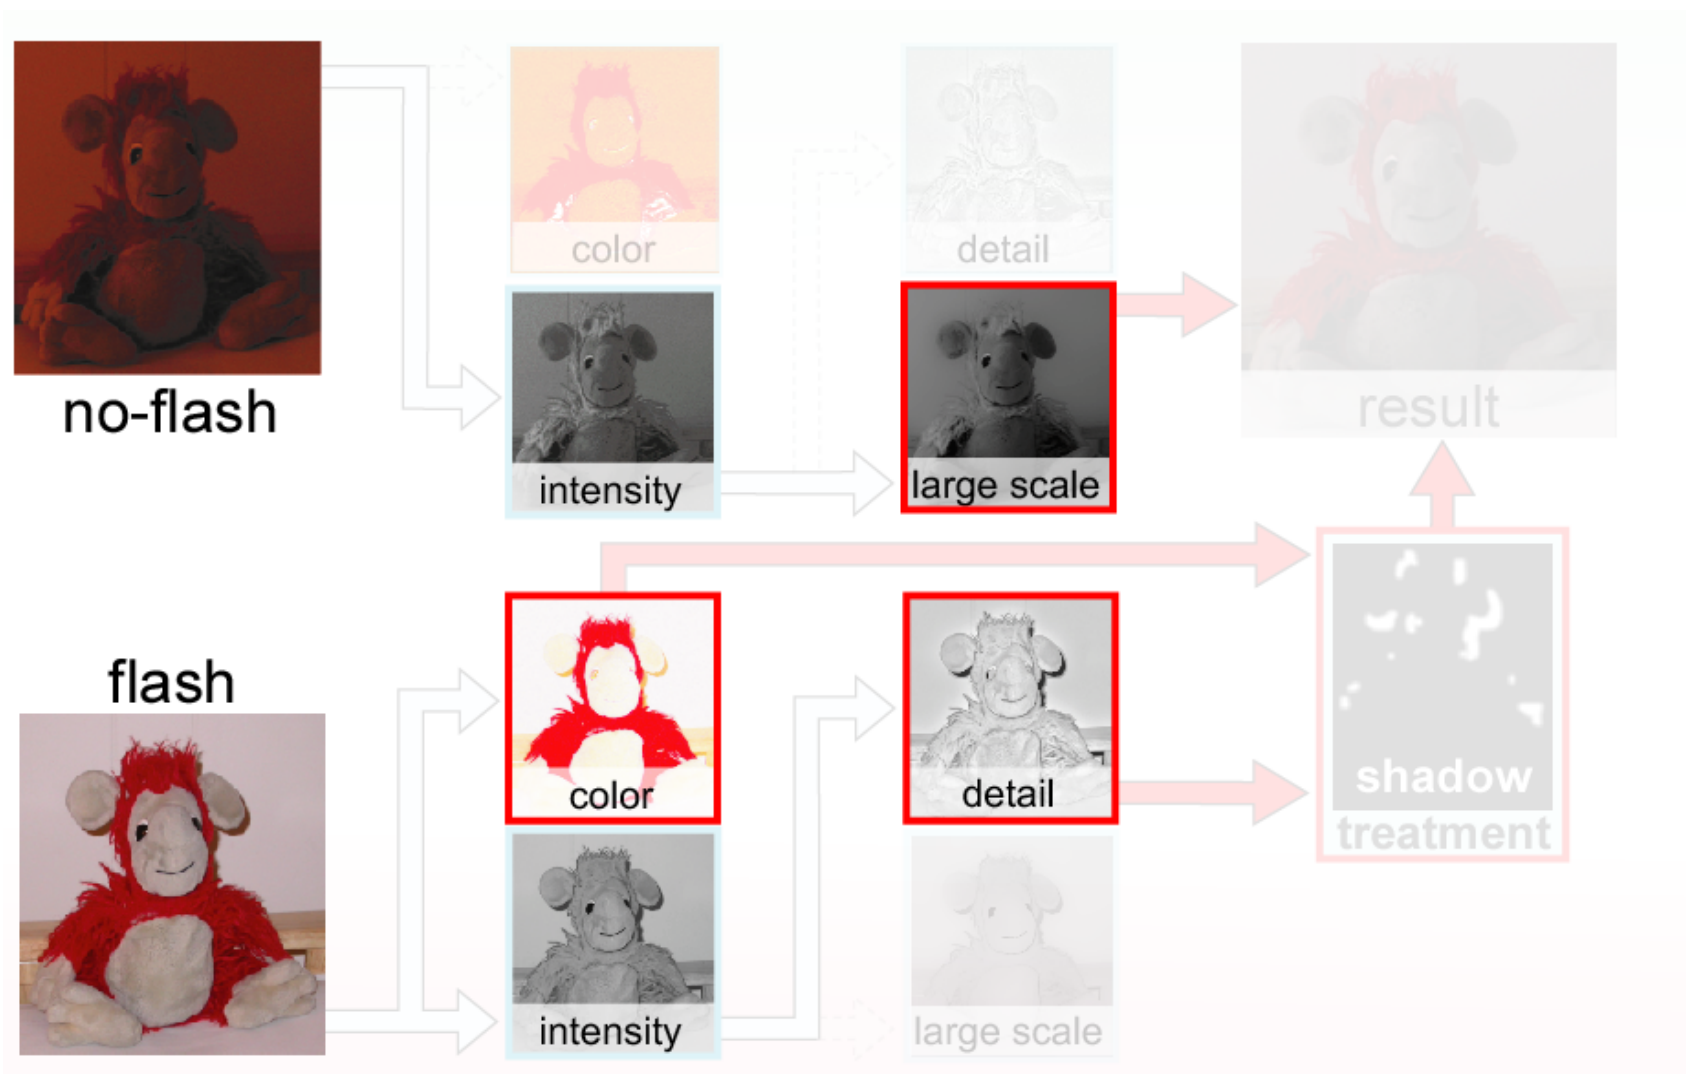
Recombination: Large scale \* Detail = Intensity

# Recombination



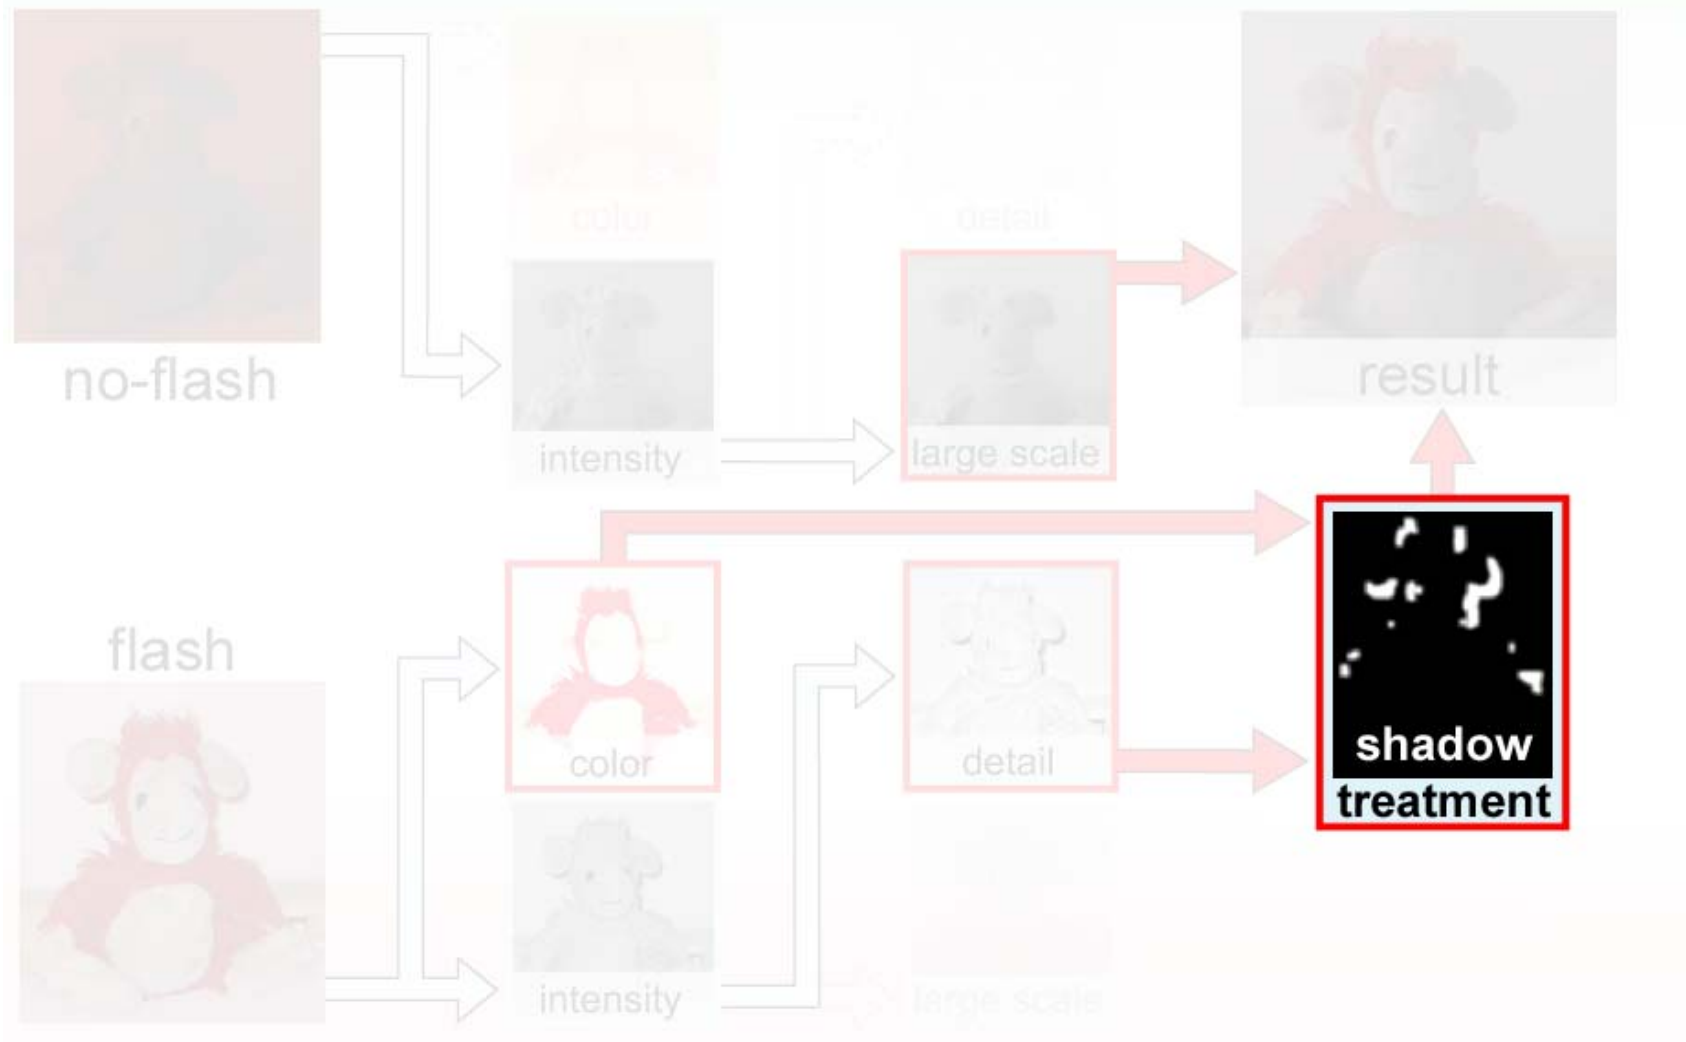
Recombination: Intensity \* Color = Original

# Our Approach



# Our Approach

## Shadow Detection/Treatment



# Results



No-flash



Flash



# Joint bilateral upsampling

---

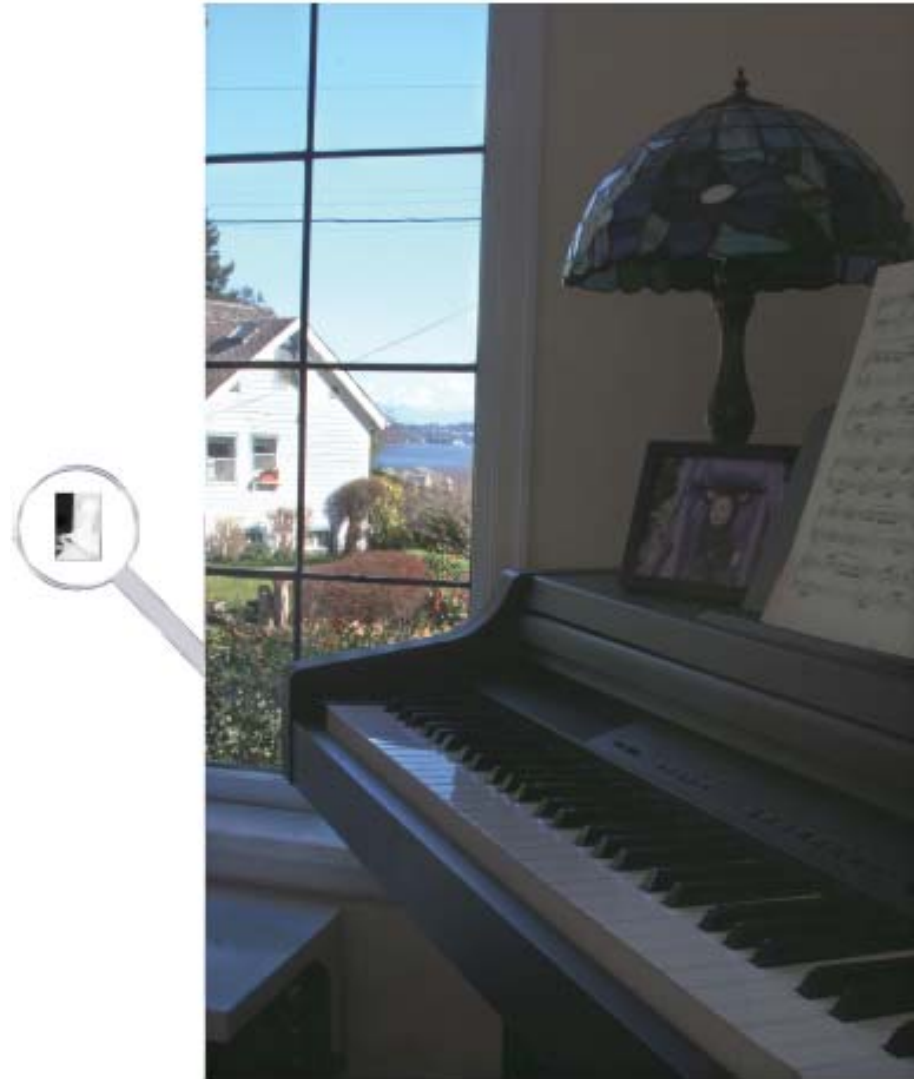
$$J_p = \frac{1}{k_p} \sum_{q \in \Omega} I_q f(\|p - q\|) g(\|I_p - I_q\|)$$

$$J_p = \frac{1}{k_p} \sum_{q \in \Omega} I_q f(\|p - q\|) g(\|\tilde{I}_p - \tilde{I}_q\|)$$

$$\tilde{S}_p = \frac{1}{k_p} \sum_{q_{\downarrow} \in \Omega} S_{q_{\downarrow}} f(\|p_{\downarrow} - q_{\downarrow}\|) g(\|\tilde{I}_p - \tilde{I}_q\|)$$

# Joint bilateral upsampling

---



Upsampled Result



# Joint bilateral upsampling

---



Nearest Neighbor

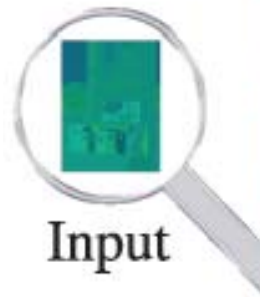
Bicubic

Gaussian

Joint Bilateral

Ground Truth

# Joint bilateral upsampling



# Joint bilateral upsampling

---



Nearest Neighbor Upsampling



Bicubic Upsampling



Gaussian Upsampling



Joint Bilateral Upsampling

# Joint bilateral upsampling



Input Images



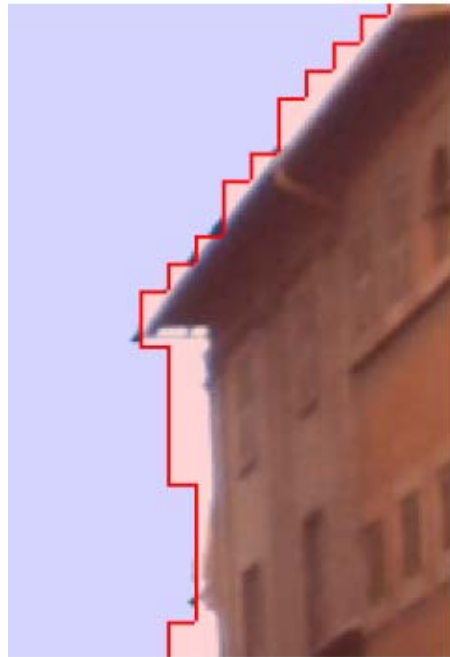
Downsampled



Input Solution

# Joint bilateral upsampling

---



Nearest Neighbor



Bicubic



Gaussian



Joint Bilateral

# Joint bilateral upsampling

---



Upsampled Result

# References

---

- Patrick Perez, Michel Gangnet, Andrew Blake, [Poisson Image Editing](#), SIGGRAPH 2003.
- Dani Lischinski, Zeev Farbman, Matt Uytendaele and Richard Szeliski. [Interactive Local Adjustment of Tonal Values](#). SIGGRAPH 2006.
- Carsten Rother, Andrew Blake, Vladimir Kolmogorov, [GrabCut - Interactive Foreground Extraction Using Iterated Graph Cuts](#), SIGGRAPH 2004.
- Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David H. Salesin, Michael F. Cohen, [Interactive Digital Photomontage](#), SIGGRAPH 2004.
- Sylvain Paris and Fredo Durand. [A Fast Approximation of the Bilateral Filter using a Signal Processing Approach](#). ECCV 2006.
- Soonmin Bae, Sylvain Paris and Fredo Durand. [Two-scale Tone Management for Photographic Look](#). SIGGRAPH 2006.