

Image Denoising



noisy image

naive denoising
Gaussian blurbetter denoising
edge-preserving filter

Smoothing an image without blurring its edges.

Bilateral Filters

Digital Visual Effects

Yung-Yu Chuang

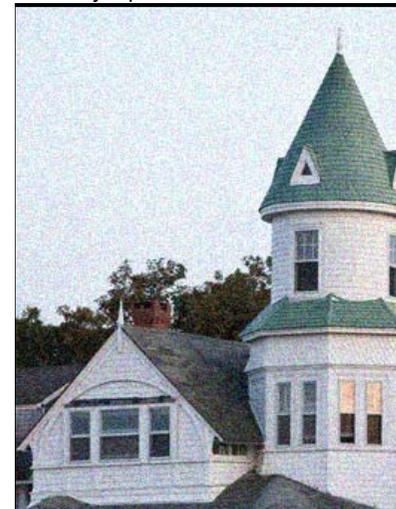
with slides by Fredo Durand, Ramesh Raskar, Sylvain Paris, Soonmin Bae

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



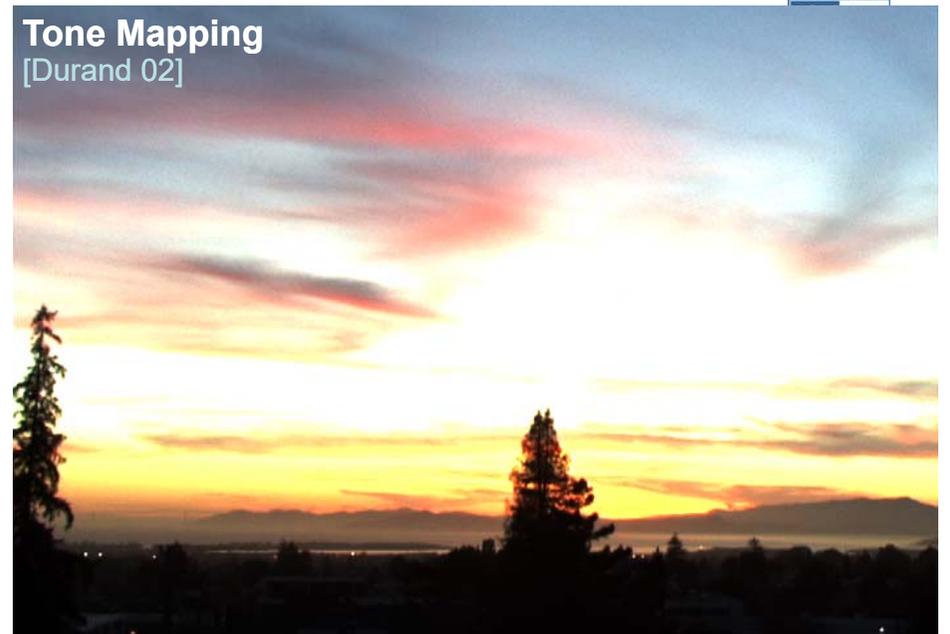
Noisy input

Bilateral filter 7x7 window



Tone Mapping

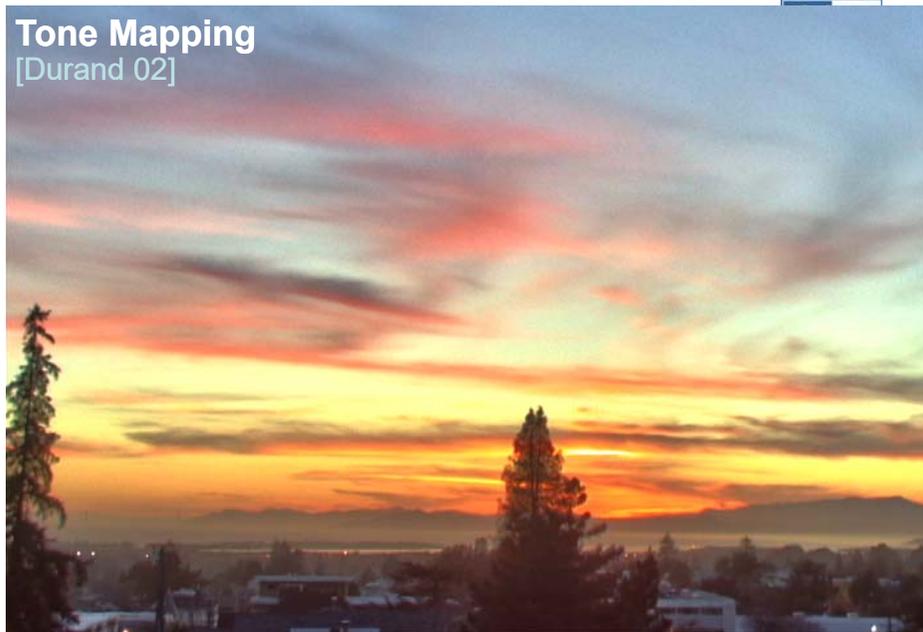
[Durand 02]



HDR input

Tone Mapping

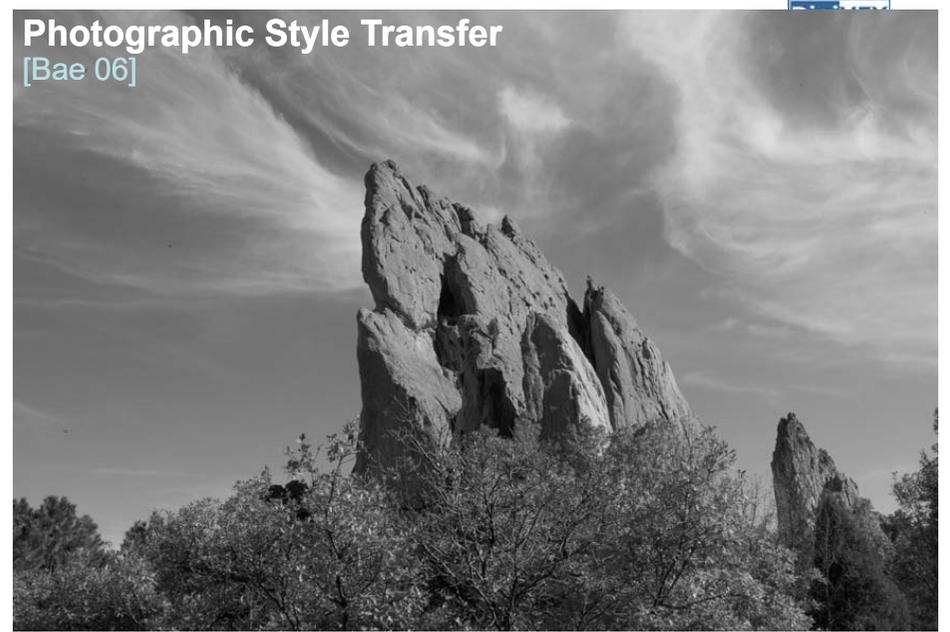
[Durand 02]



output

Photographic Style Transfer

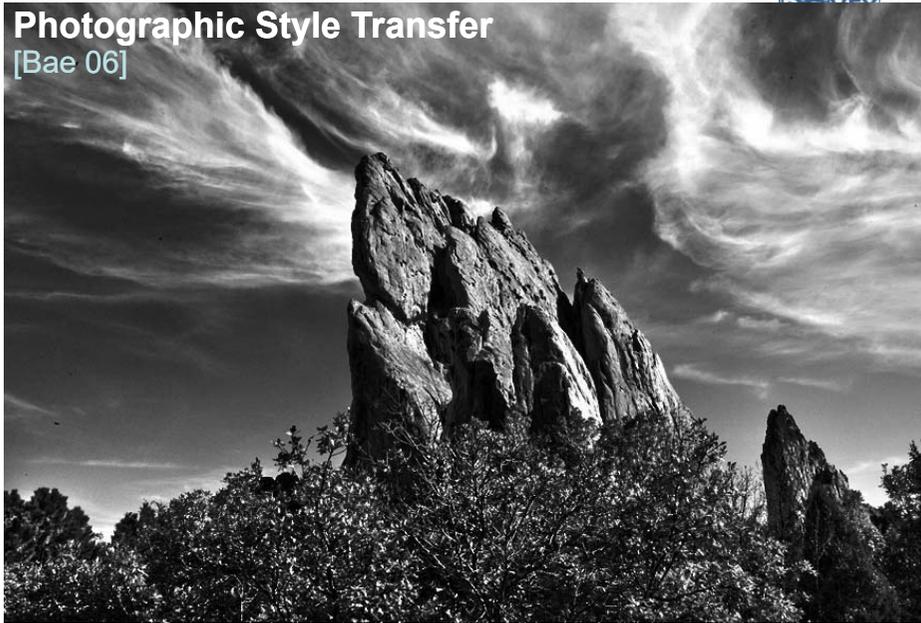
[Bae 06]



input

Photographic Style Transfer

[Bae 06]



output

Cartoon Rendition

[Winnemöller 06]



input

Cartoon Rendition

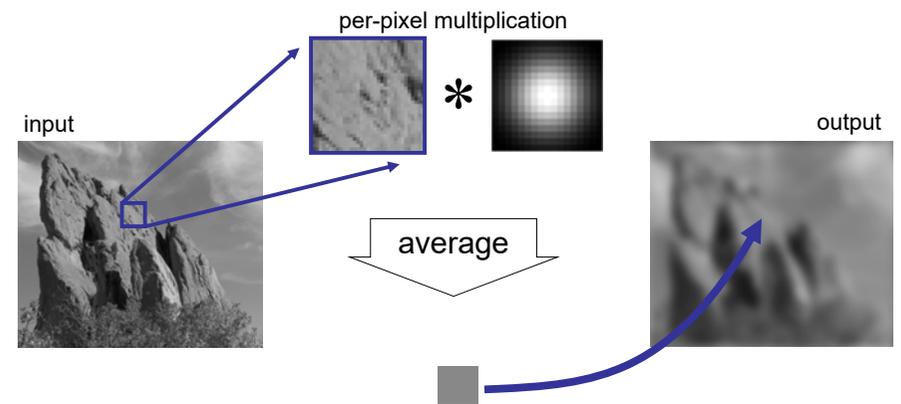
[Winnemöller 06]



output

Gaussian Blur

DigiVFX

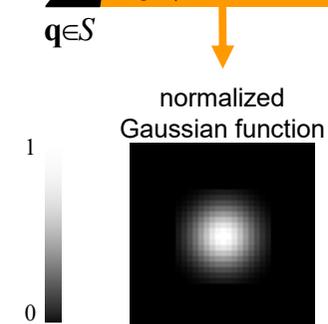




Equation of Gaussian Blur

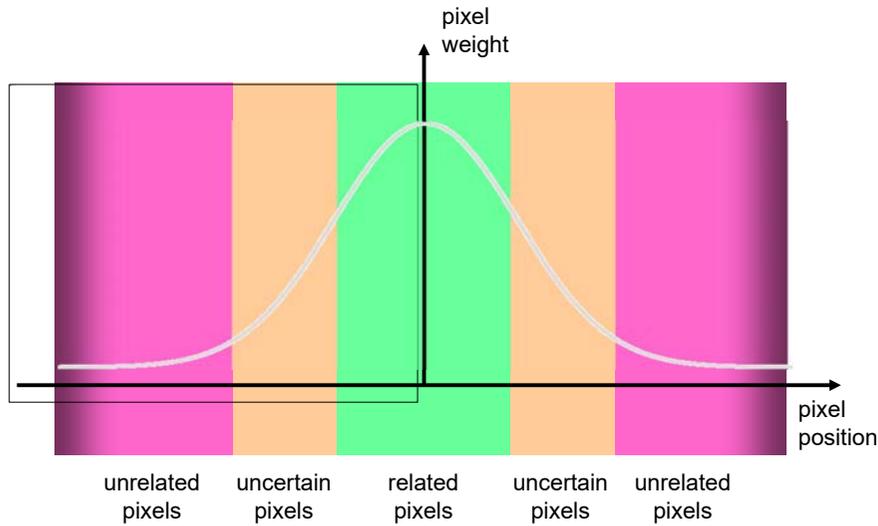
Same idea: **weighted average of pixels.**

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



Gaussian Profile

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



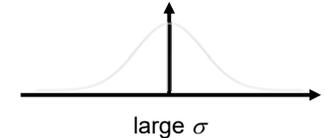
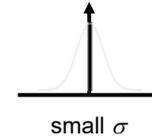
Spatial Parameter



input

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

σ
 size of the window



Properties of Gaussian Blur

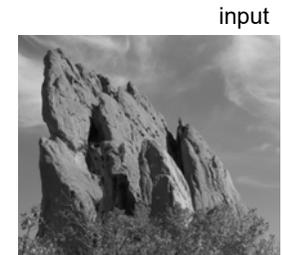


- Weights independent of spatial location
 - linear convolution
 - well-known operation
 - efficient computation (recursive algorithm, FFT...)

Properties of Gaussian Blur



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



input



output

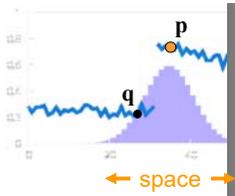


$$GB[I]_p = \sum_{q \in S} G_{\sigma}(\|p - q\|) I_q$$

σ
 space

Gaussian Blur and Bilateral Filter DigiVFX

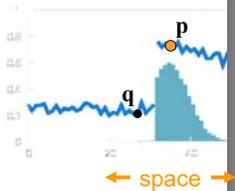
Gaussian blur



$$GB[I]_p = \sum_{q \in S} \underbrace{G_{\sigma_s}(\|p - q\|)}_{\text{space}} I_q$$

Bilateral filter

[Aurich 95, Smith 97] Tomasi 98

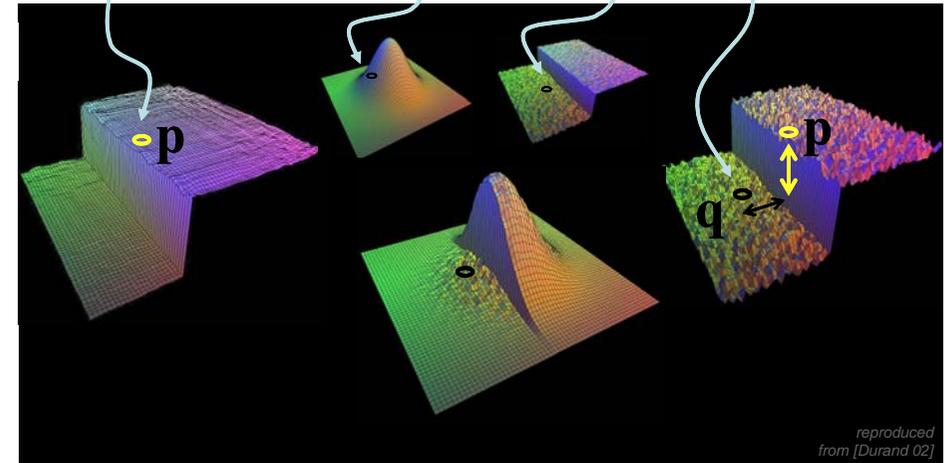


$$BF[I]_p = \frac{1}{W_p} \sum_{q \in S} \underbrace{G_{\sigma_s}(\|p - q\|)}_{\text{space}} \underbrace{G_{\sigma_r}(\|I_p - I_q\|)}_{\text{range}} I_q$$

normalization

Bilateral Filter on a Height Field DigiVFX

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



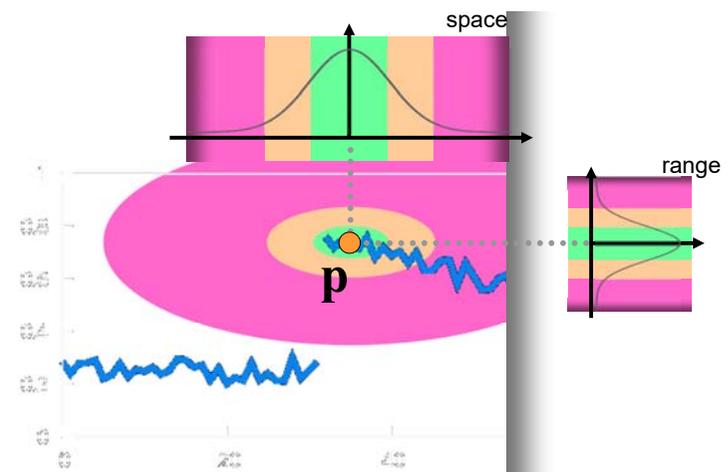
Space and Range Parameters DigiVFX

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

- space σ_s : spatial extent of the kernel, size of the considered neighborhood.
- range σ_r : "minimum" amplitude of an edge

Influence of Pixels DigiVFX

Only pixels close in space and in range are considered.



Iterating the Bilateral Filter

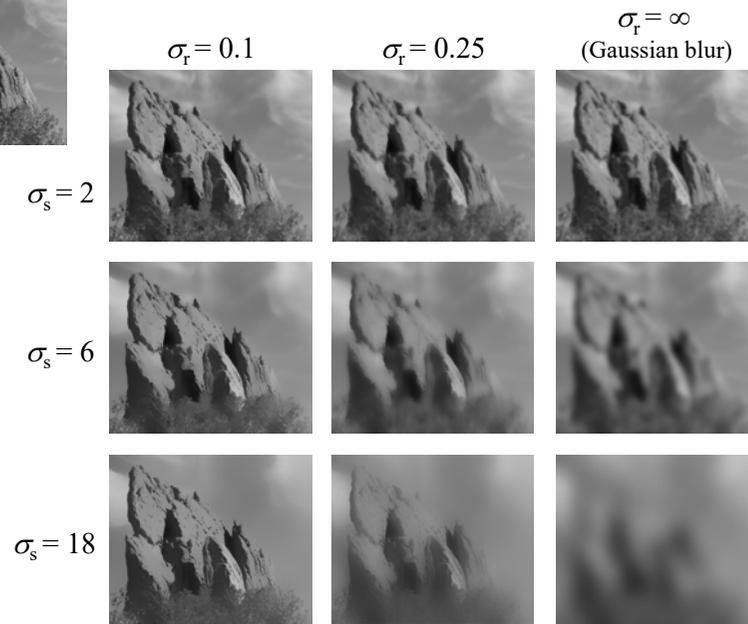
$$I_{(n+1)} = BF[I_{(n)}]$$

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



input

Exploring the Parameter Space



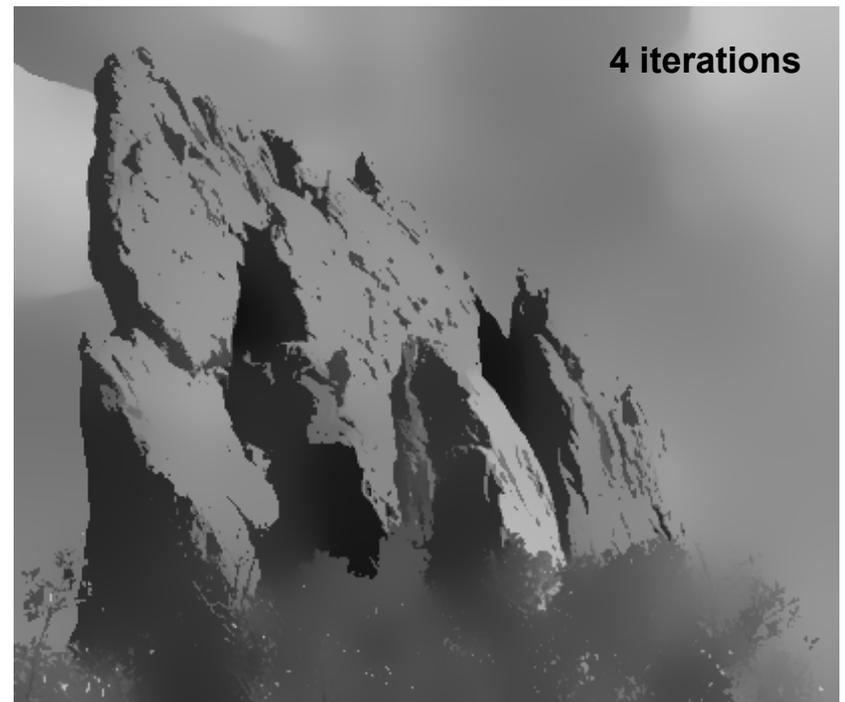
input



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

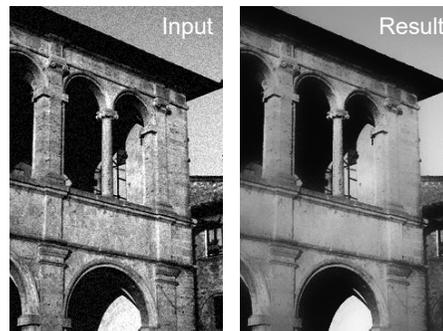
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology



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



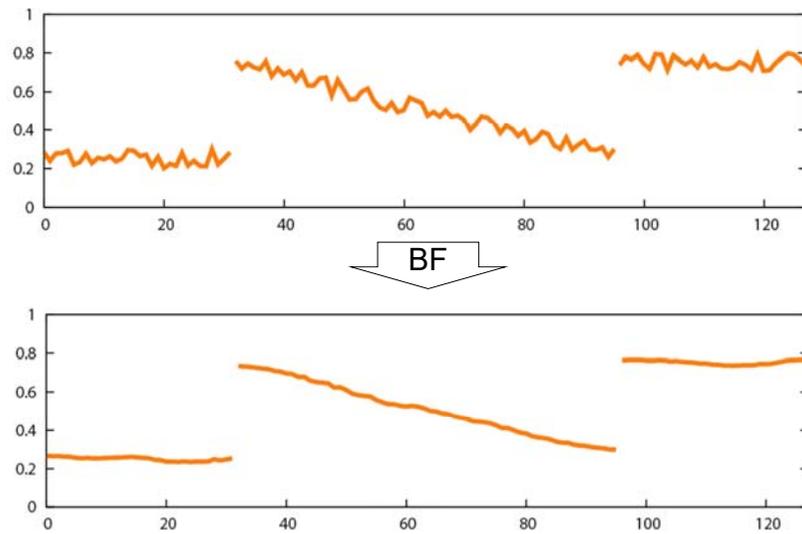
$$I_{\mathbf{p}}^{\text{bf}} = \frac{1}{W_{\mathbf{p}}^{\text{bf}}} \sum_{\mathbf{q} \in \mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|)}_{\text{space}} \underbrace{G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)}_{\text{range}} I_{\mathbf{q}}$$

Contributions

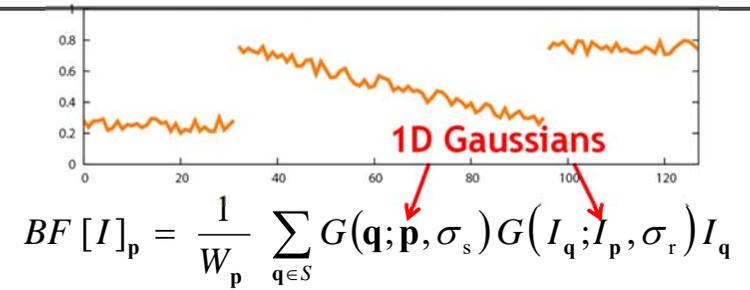


- Link with **linear filtering**
- **Fast and accurate** approximation

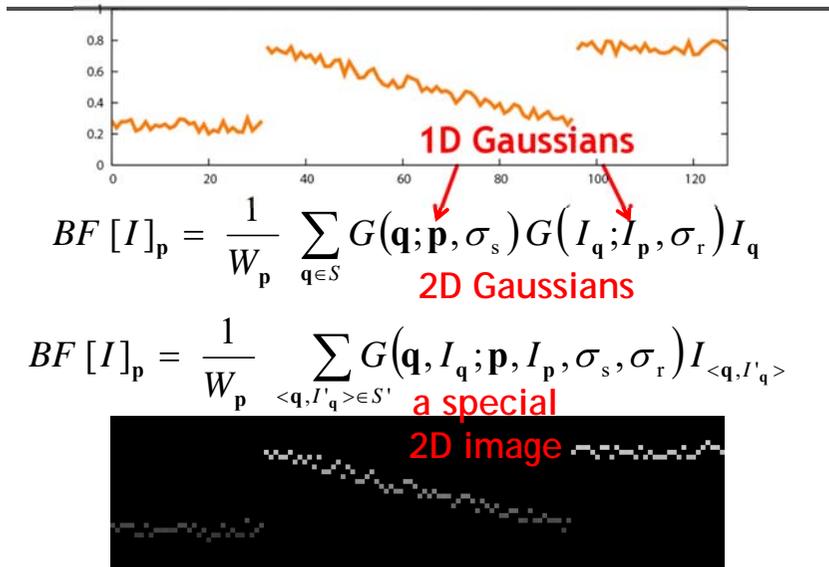
Intuition on 1D Signal



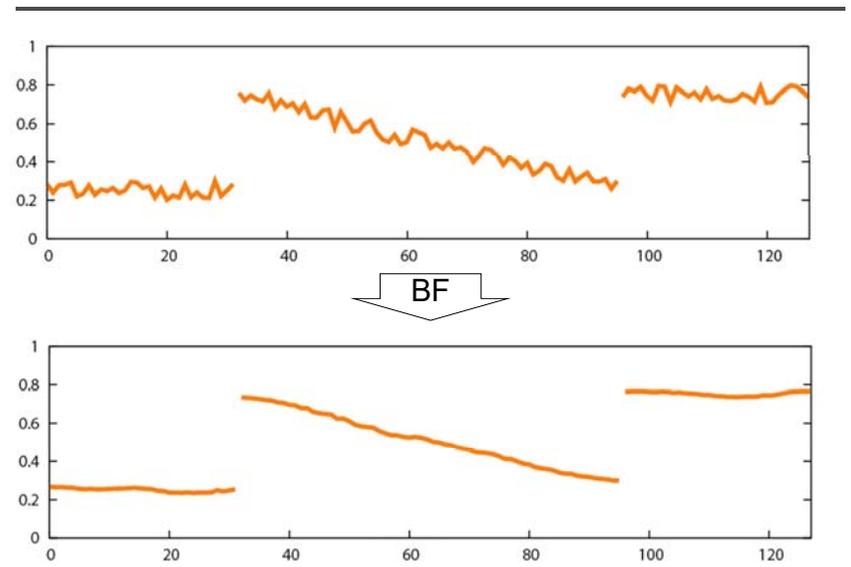
Basic idea



Basic idea



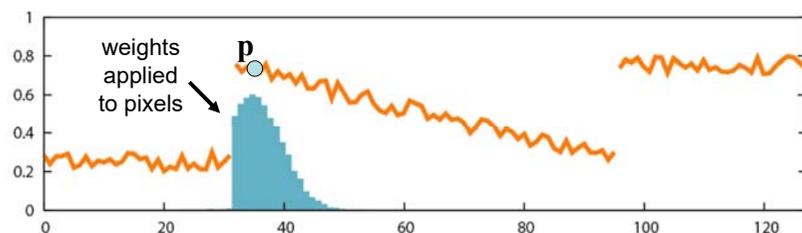
Intuition on 1D Signal



Intuition on 1D Signal

Weighted Average of Neighbors

DigiVFX

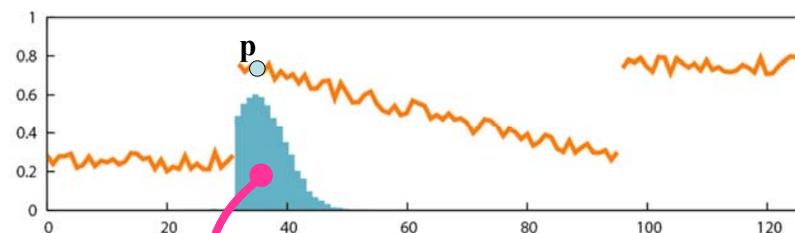


- Near and similar pixels have influence.
- Far pixels have no influence.
- Pixels with different value have no influence.

Link with Linear Filtering

1. Handling the Division

DigiVFX



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

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

Formalization: Handling the Division

DigiVFX

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

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

- 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

DigiVFX

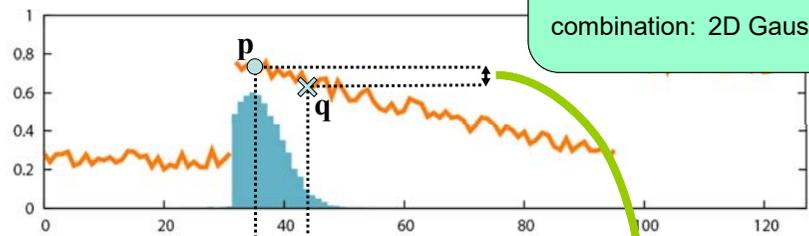
$$\begin{pmatrix} W_{\mathbf{p}}^{\text{bf}} I_{\mathbf{p}}^{\text{bf}} \\ W_{\mathbf{p}}^{\text{bf}} \end{pmatrix} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) \begin{pmatrix} W_{\mathbf{q}} I_{\mathbf{q}} \\ W_{\mathbf{q}}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

- 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

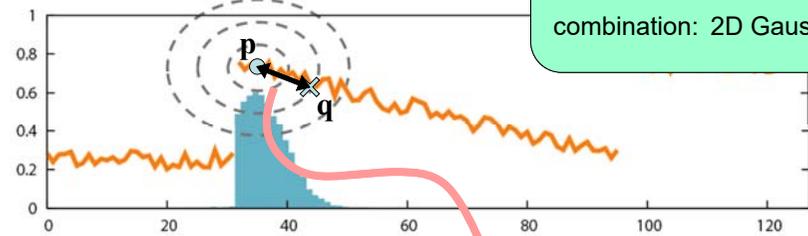


$$\begin{pmatrix} W_{\mathbf{p}}^{bf} & I_{\mathbf{p}}^{bf} \\ W_{\mathbf{p}}^{bf} & \end{pmatrix} = \sum_{\mathbf{q} \in \mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|)}_{\text{space}} \underbrace{G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)}_{\text{range}} \begin{pmatrix} W_{\mathbf{q}}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

Link with Linear Filtering 2. Introducing a Convolution

space: 1D Gaussian
× range: 1D Gaussian

combination: 2D Gaussian

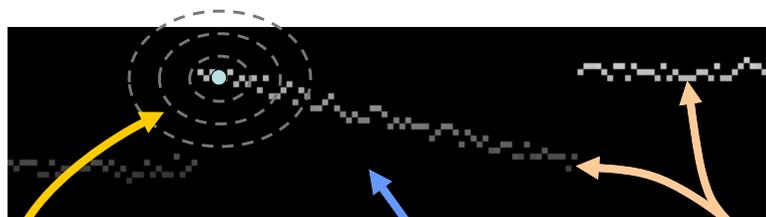


$$\begin{pmatrix} W_{\mathbf{p}}^{bf} & I_{\mathbf{p}}^{bf} \\ W_{\mathbf{p}}^{bf} & \end{pmatrix} = \sum_{\mathbf{q} \in \mathcal{S}} \underbrace{G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)}_{\text{space x range}} \begin{pmatrix} W_{\mathbf{q}} & I_{\mathbf{q}}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

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

Link with Linear Filtering 2. Introducing a Convolution

DigiVFX



sum all values

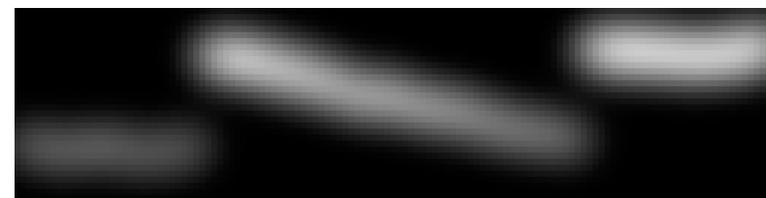
black = zero

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

sum all values multiplied by kernel ⇒ convolution

Link with Linear Filtering 2. Introducing a Convolution

DigiVFX



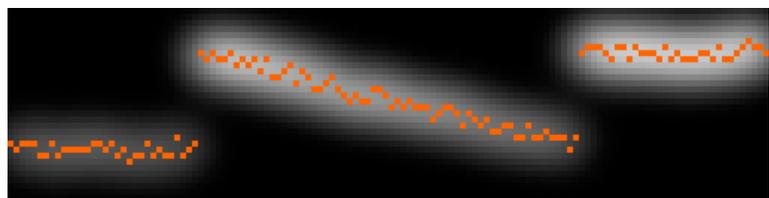
result of the convolution

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

Link with Linear Filtering

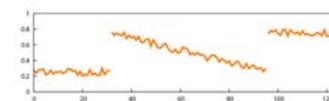
2. Introducing a Convolution

DigiVFX

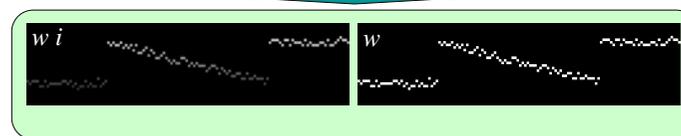


result of the convolution

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



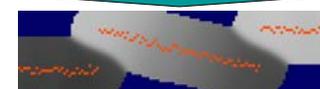
higher dimensional functions



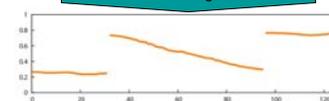
Gaussian convolution



division



slicing



Reformulation: Summary

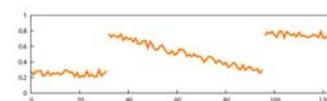
DigiVFX

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

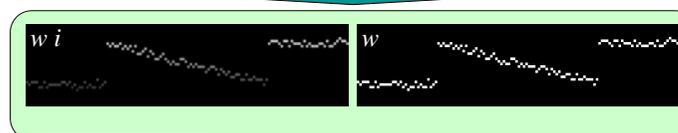
nonlinear: $I_{\mathbf{p}}^{\text{bf}} = \frac{w^{\text{bf}}(\mathbf{p}, I_{\mathbf{p}}) i^{\text{bf}}(\mathbf{p}, I_{\mathbf{p}})}{w^{\text{bf}}(\mathbf{p}, I_{\mathbf{p}})}$

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

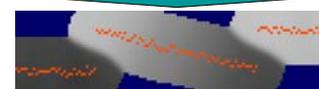


Low-pass filter

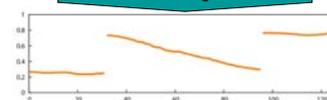
Gaussian convolution

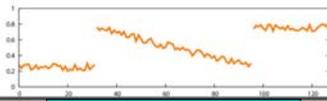


division

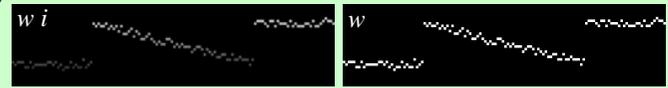


slicing





higher dimensional functions



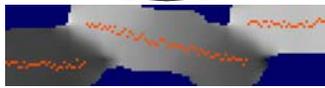
DOWNSAMPLE

Gaussian convolution

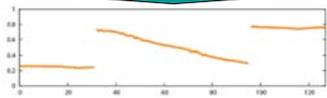


UPSAMPLE

division



slicing



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

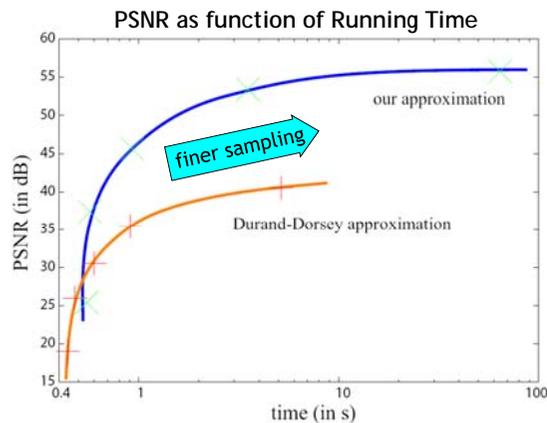
Accuracy versus Running Time

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



Digital photograph
1200 × 1600

Straightforward implementation is over 10 minutes.

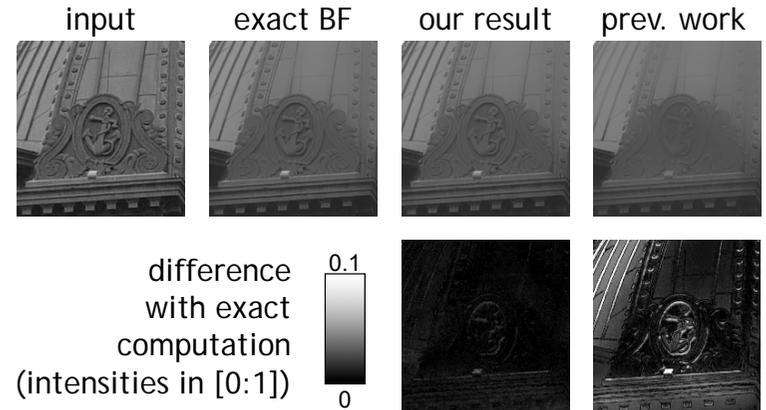


Visual Results

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



1200 × 1600

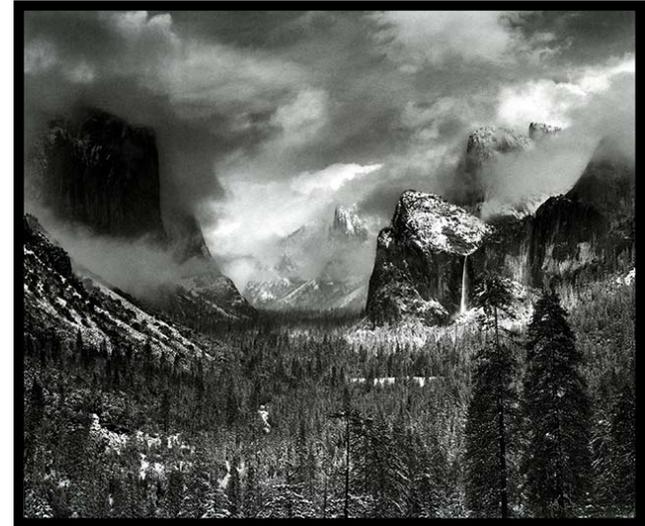


Two-scale Tone Management for Photographic Look

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

DigiVFX

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

For example,

we want



with the look of



Aspects of Photographic Look

DigiVFX

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



Input



Model

Tonal Aspects of Look

DigiVFX



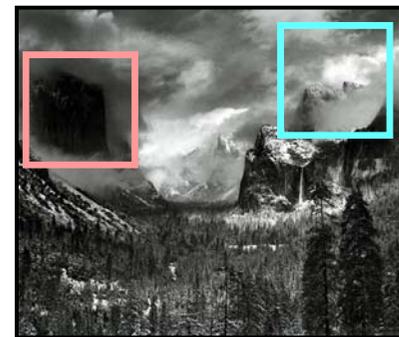
Ansel Adams



Kenro Izu

Tonal aspects of Look - Global Contrast

DigiVFX



Ansel Adams

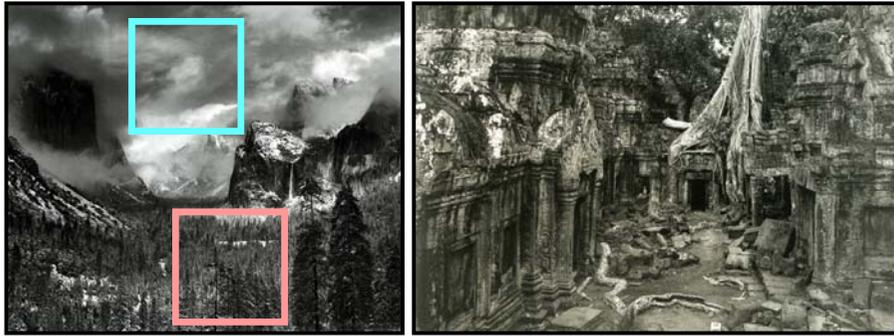


Kenro Izu

High Global Contrast

Low Global Contrast

Tonal aspects of Look - Local Contrast

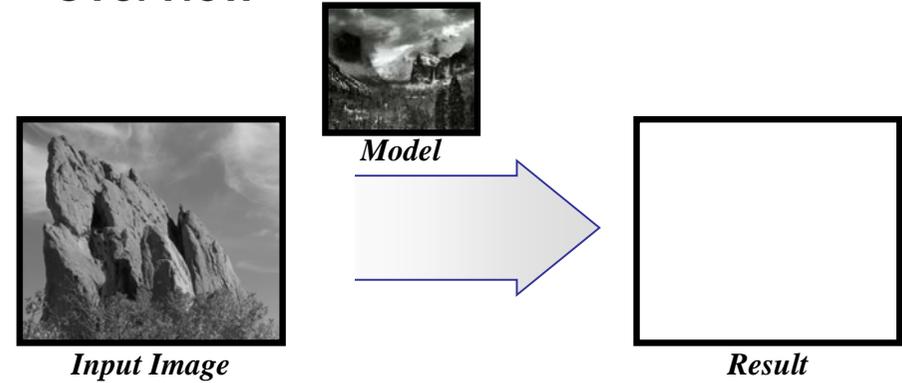


Ansel Adams

Kenro Izu

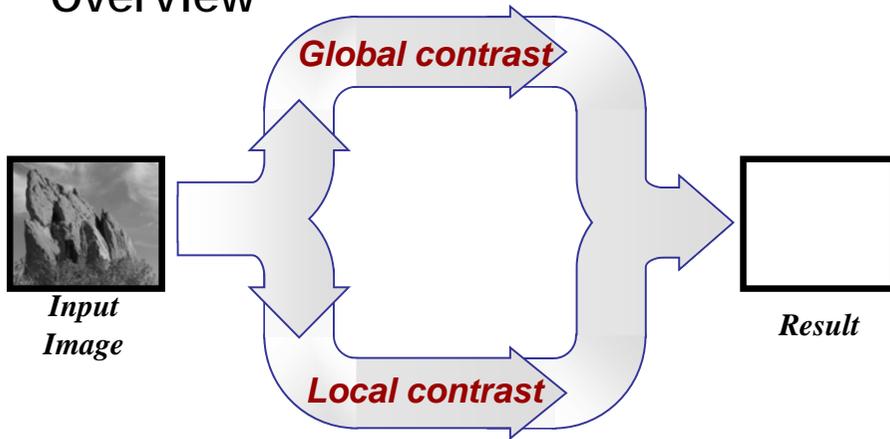
Variable amount of texture **Texture everywhere**

Overview



- Transfer look between photographs
 - Tonal aspects

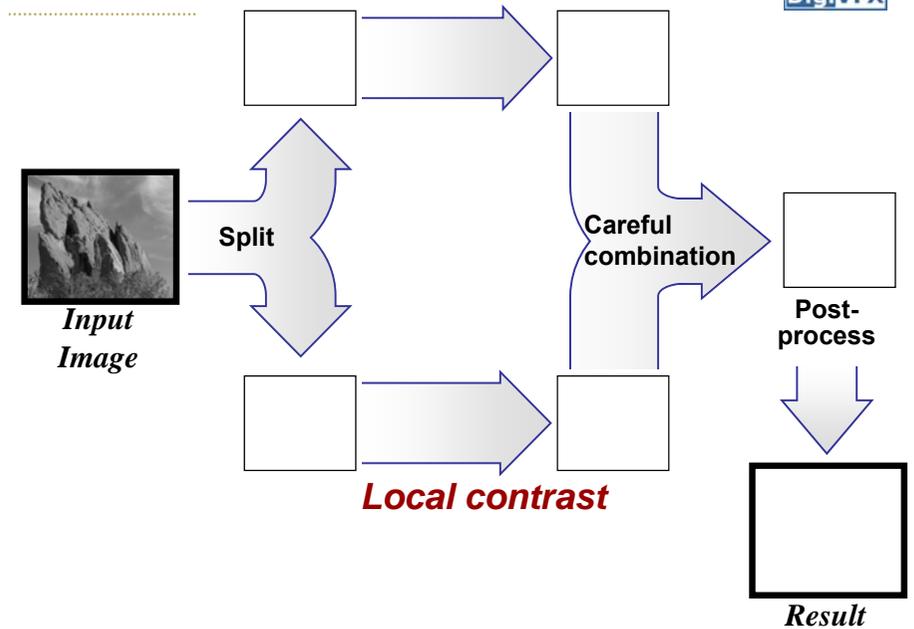
Overview



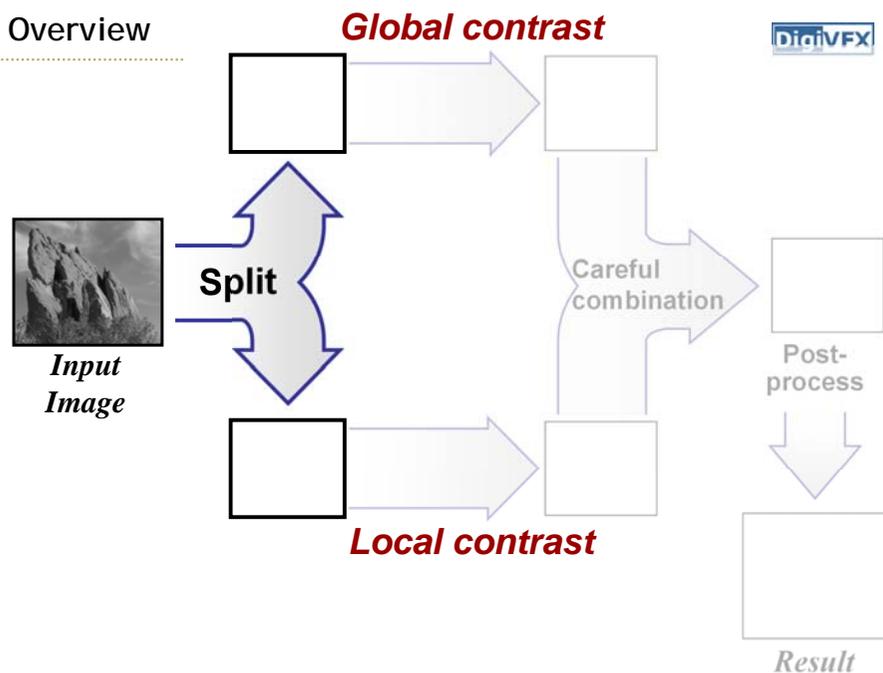
- Separate global and local contrast

Overview

Global contrast



Overview



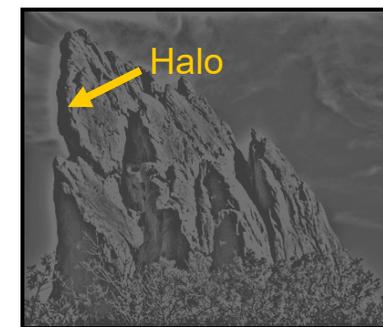
Split Global vs. Local Contrast

DigiVFX

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



Low frequency
Global contrast



High frequency
Local contrast

Bilateral Filter

DigiVFX

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



After bilateral filtering
Global contrast



Residual after filtering
Local contrast

Bilateral Filter

DigiVFX

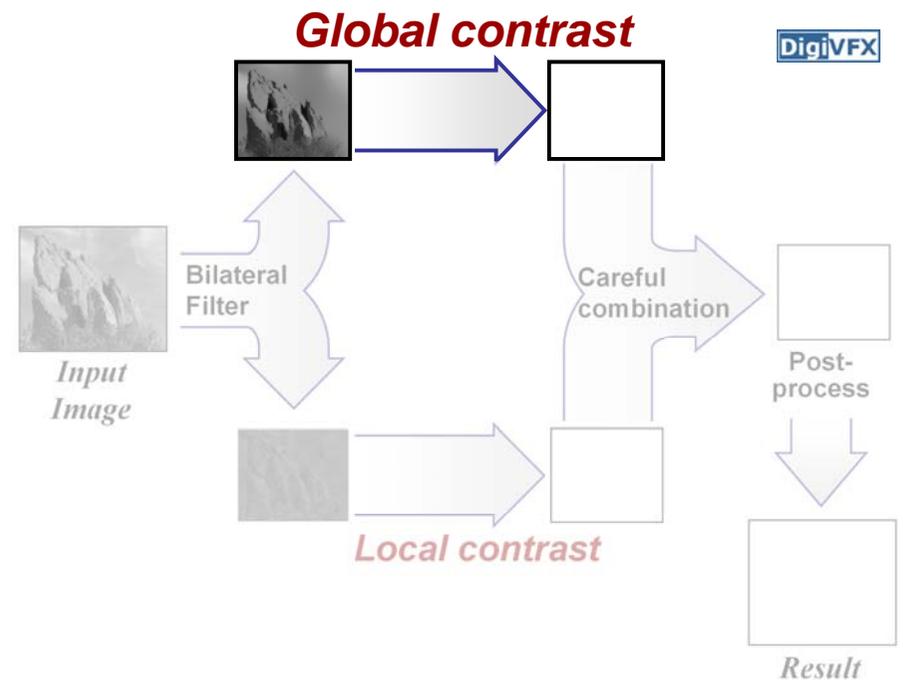
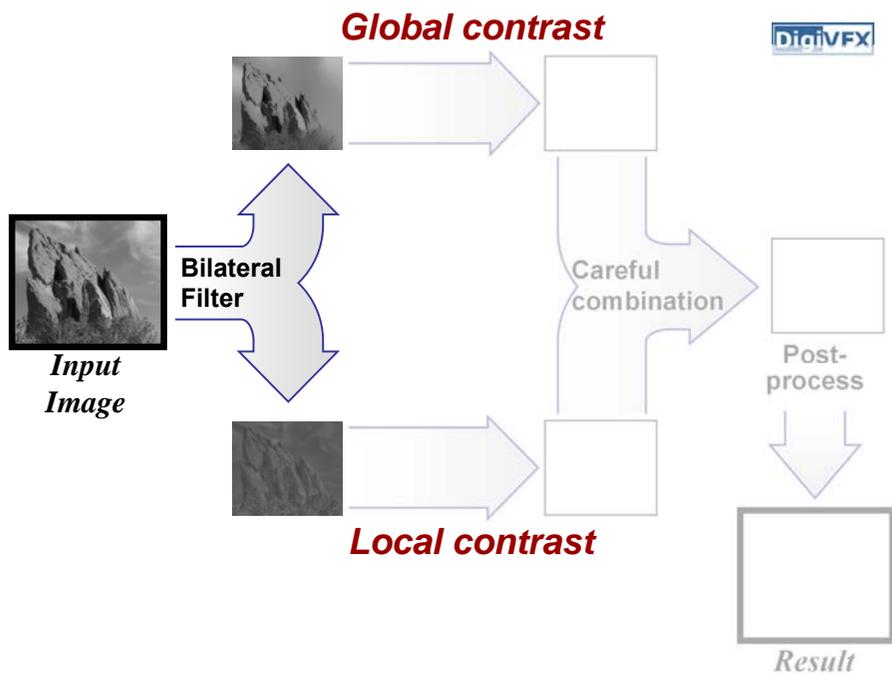
- Edge-preserving smoothing [Tomasi 98]
- We build upon tone mapping [Durand 02]



After bilateral filtering
Global contrast



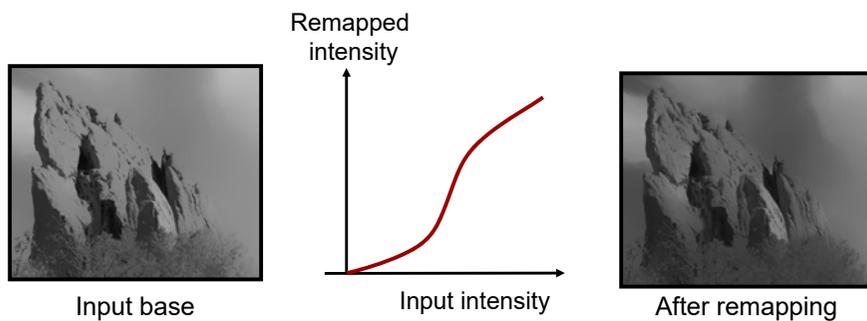
Residual after filtering
Local contrast



Global Contrast

DigiVFX

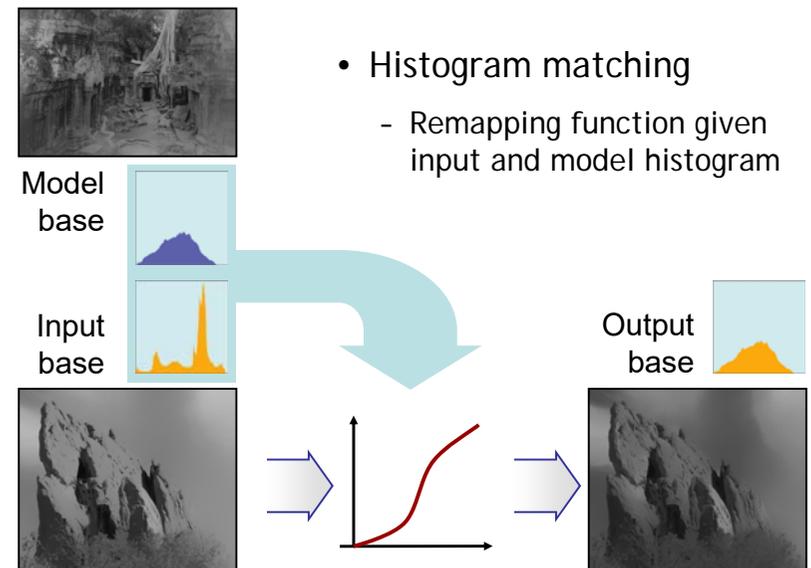
- Intensity remapping of base layer

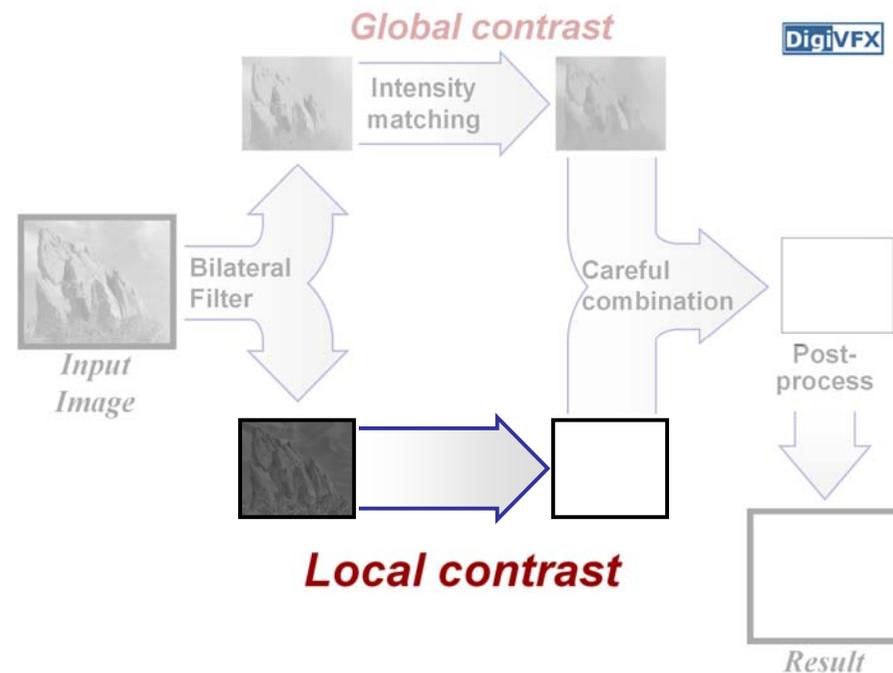
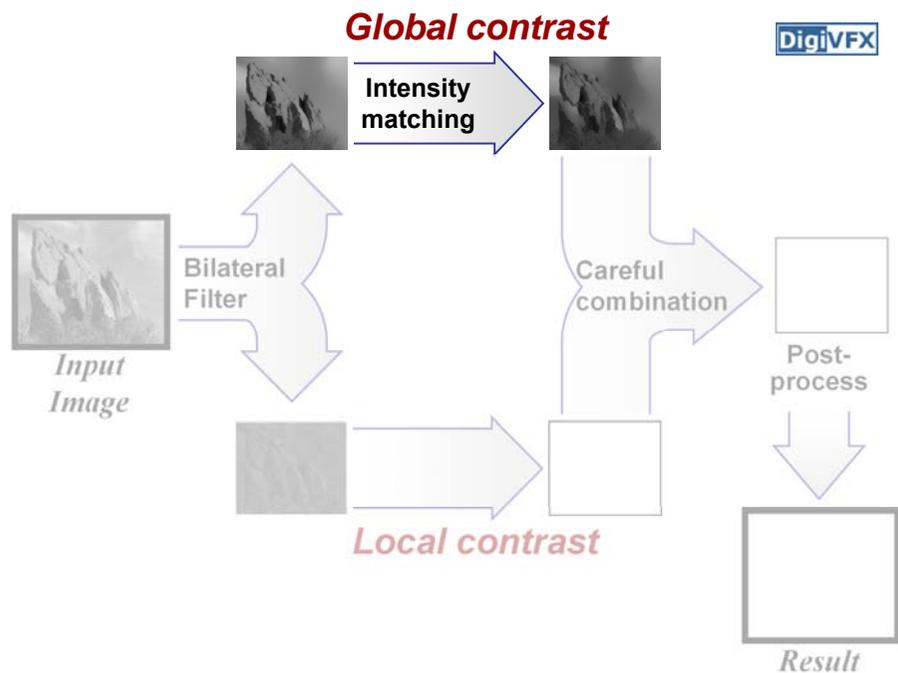


Global Contrast (Model Transfer)

DigiVFX

- Histogram matching
 - Remapping function given input and model histogram





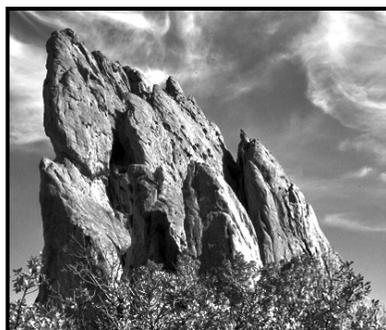
Local Contrast: Detail Layer

DigiVFX

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



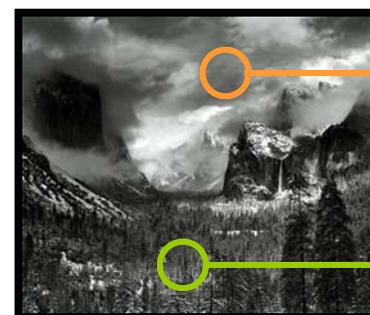
Input



Base + 3 × Detail

The amount of local contrast is not uniform

DigiVFX

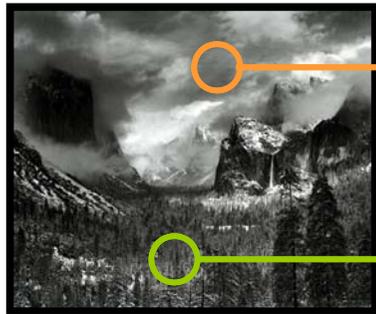


Smooth region

Textured region

Local Contrast Variation

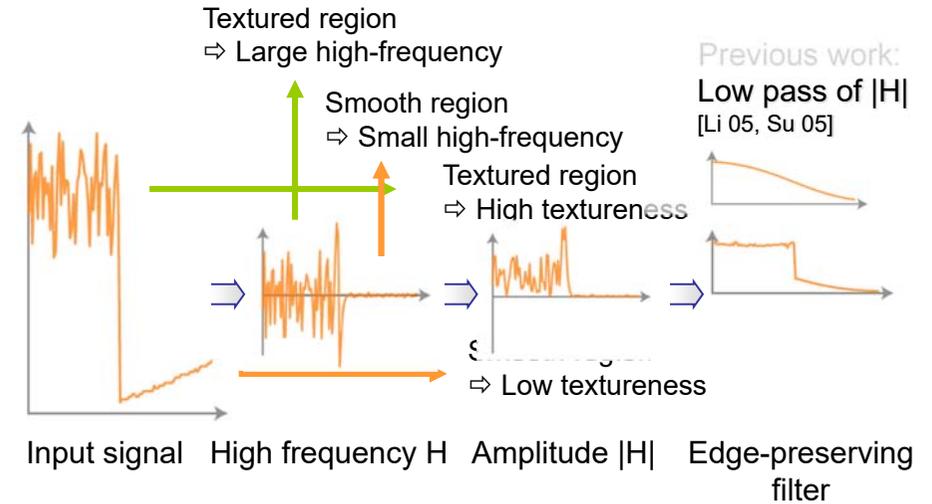
- We define "textureness": amount of local contrast
 - at each pixel based on surrounding region



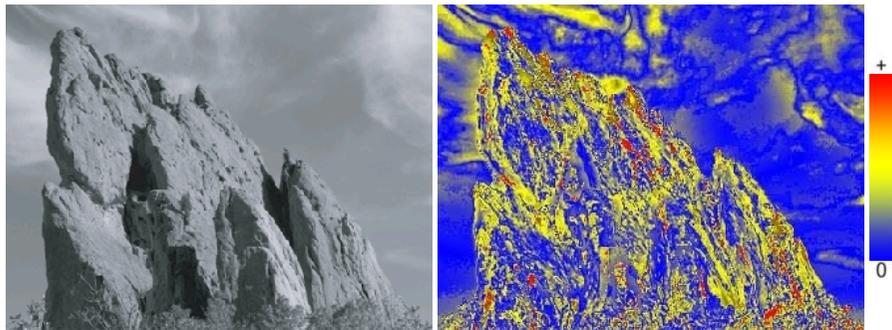
Smooth region
⇒ Low textureness

Textured region
⇒ High textureness

"Textureness": 1D Example



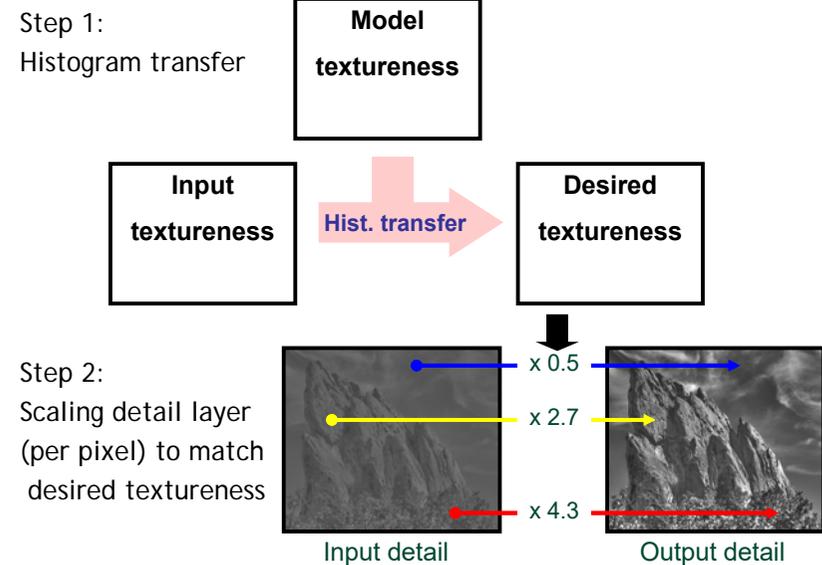
Textureness

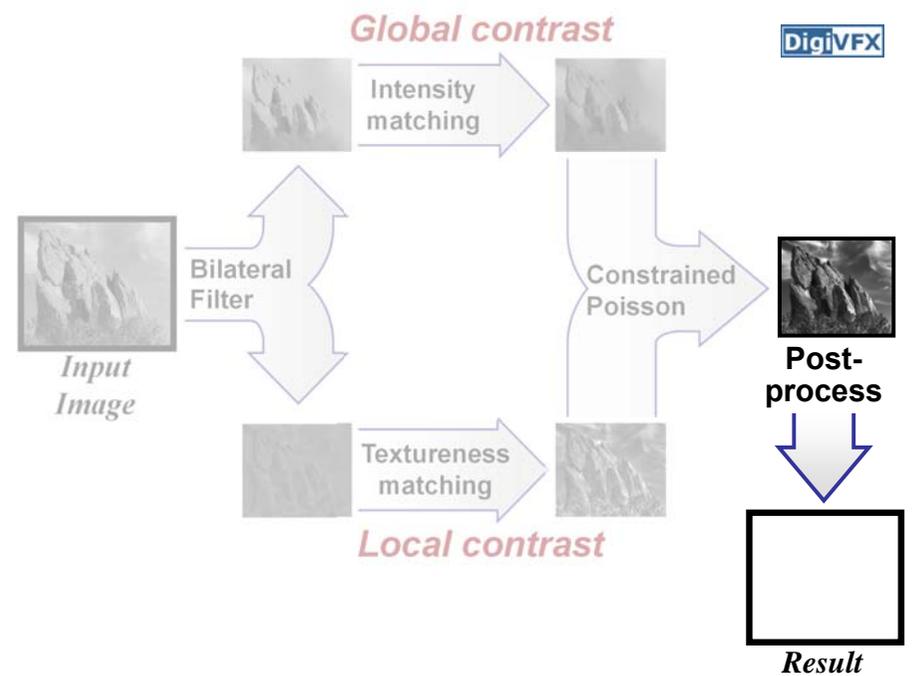
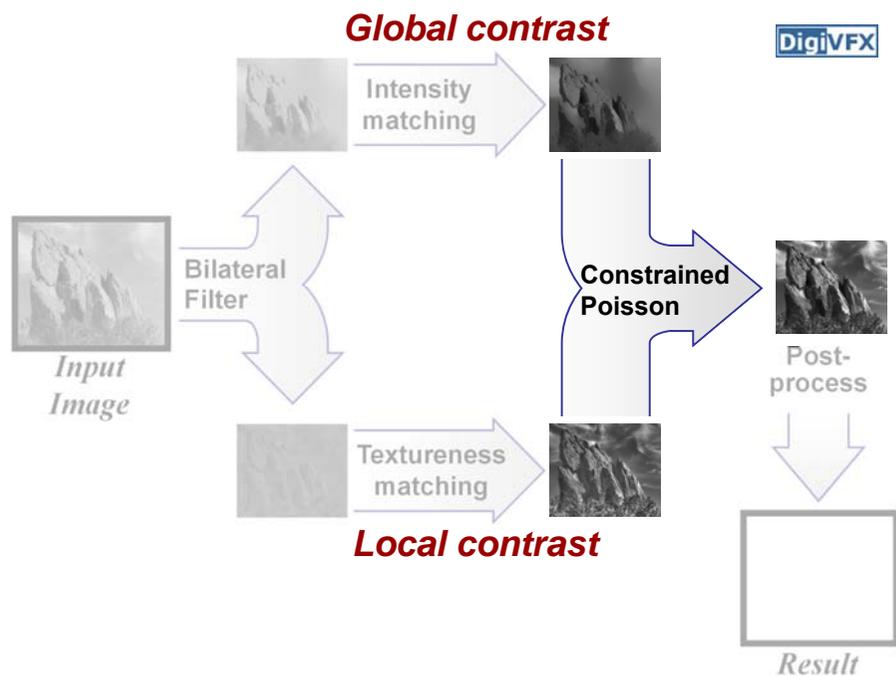
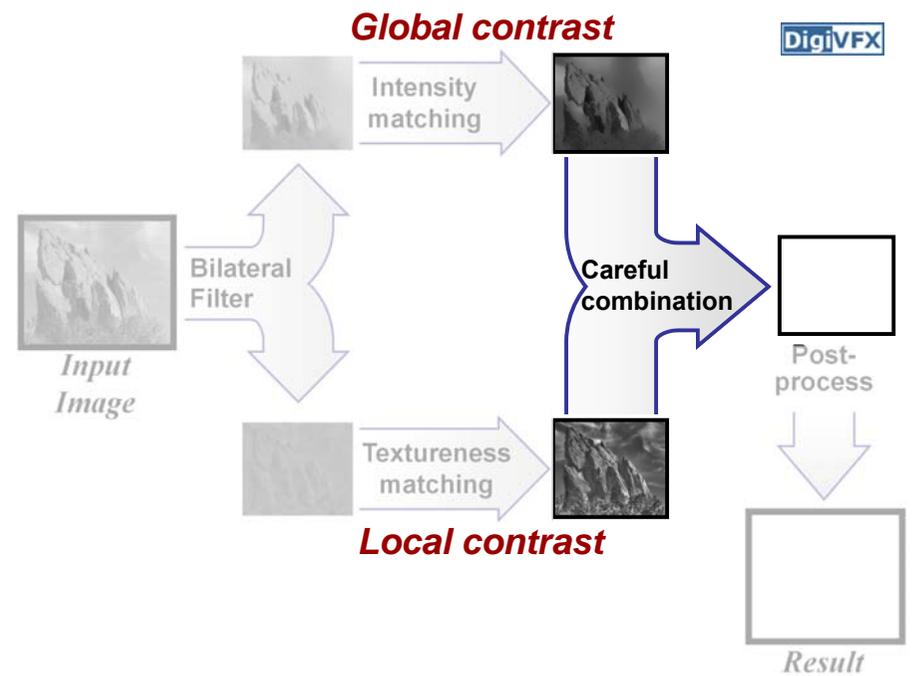
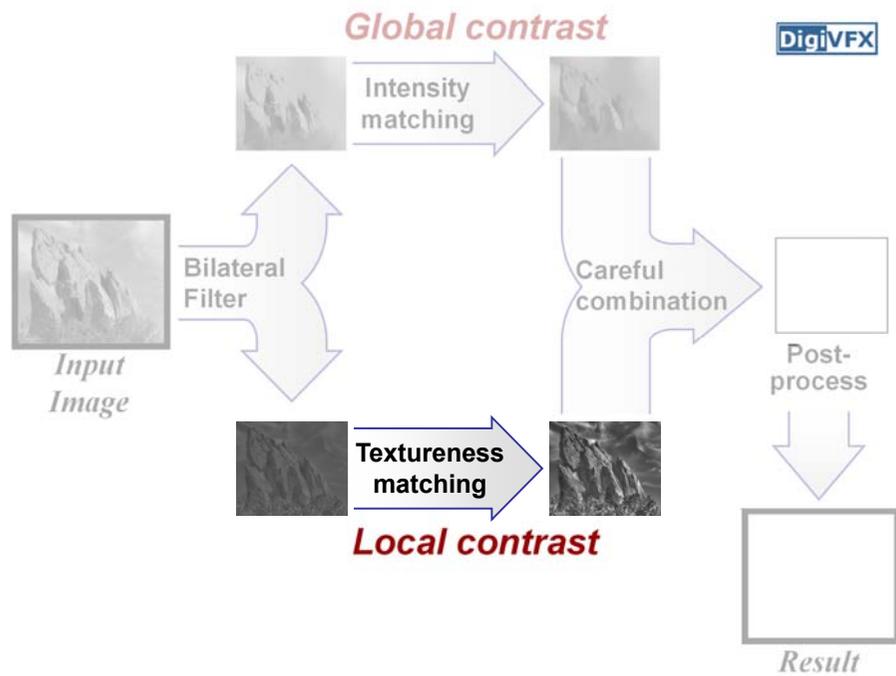


Input

Textureness

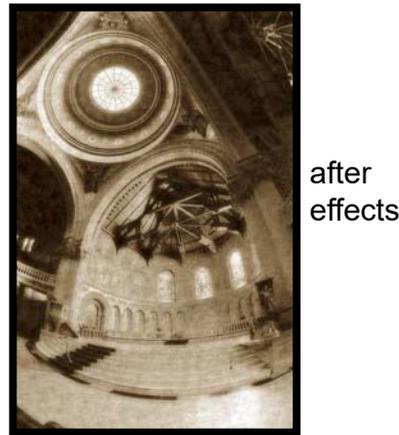
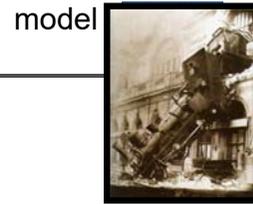
Textureness Transfer



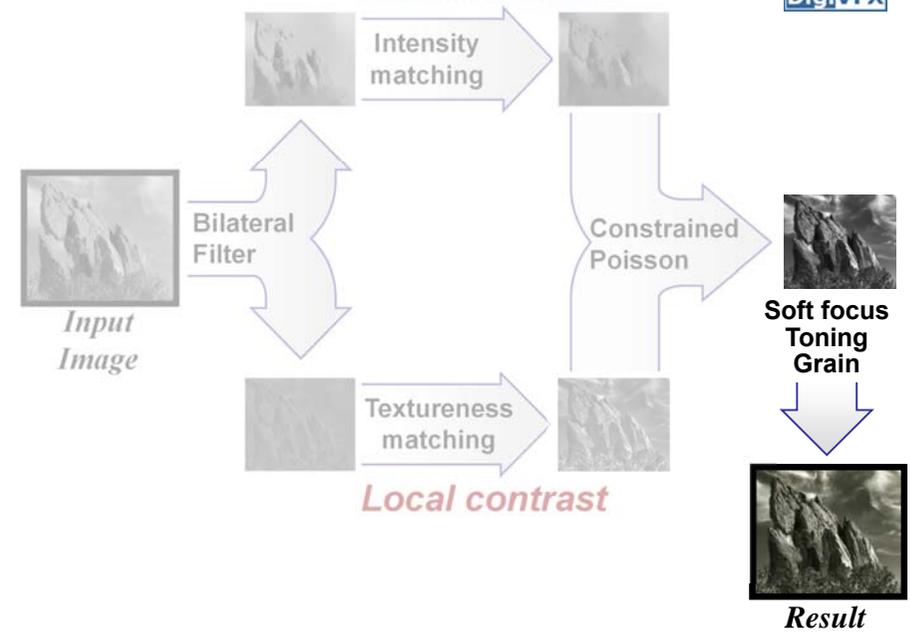


Additional Effects

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

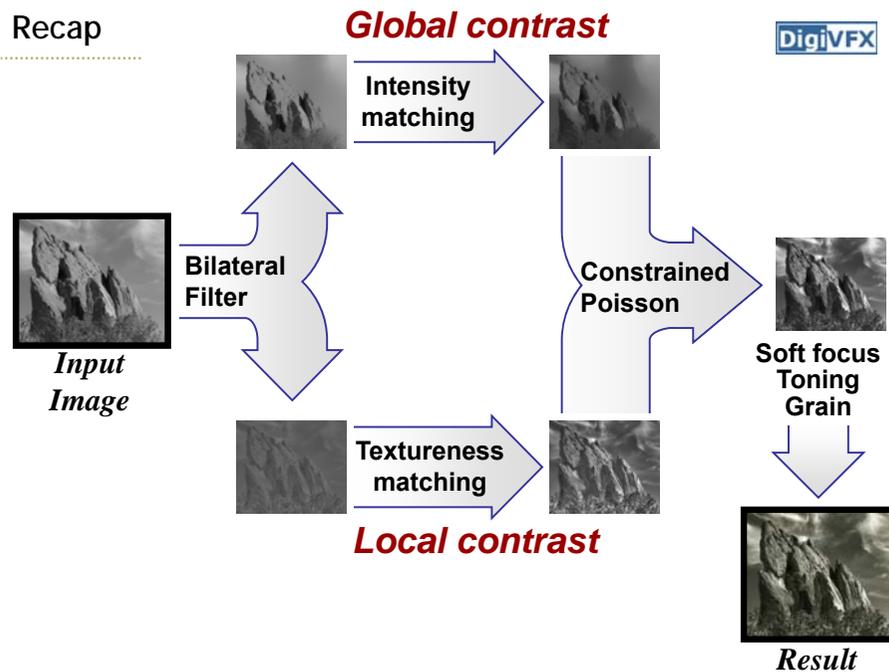


Global contrast



DigiVFX

Recap



DigiVFX

Results

DigiVFX

User provides input and model photographs.

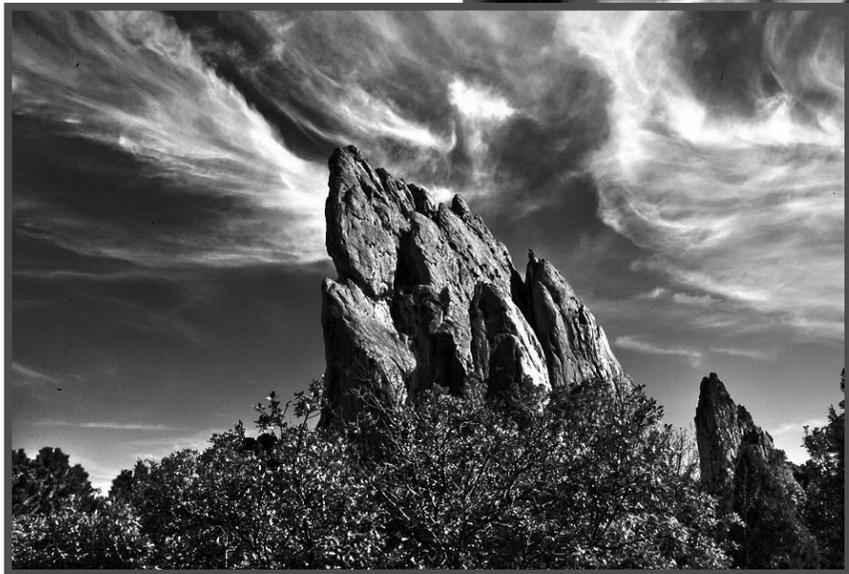
➔ Our system automatically produces the result.

Running times:

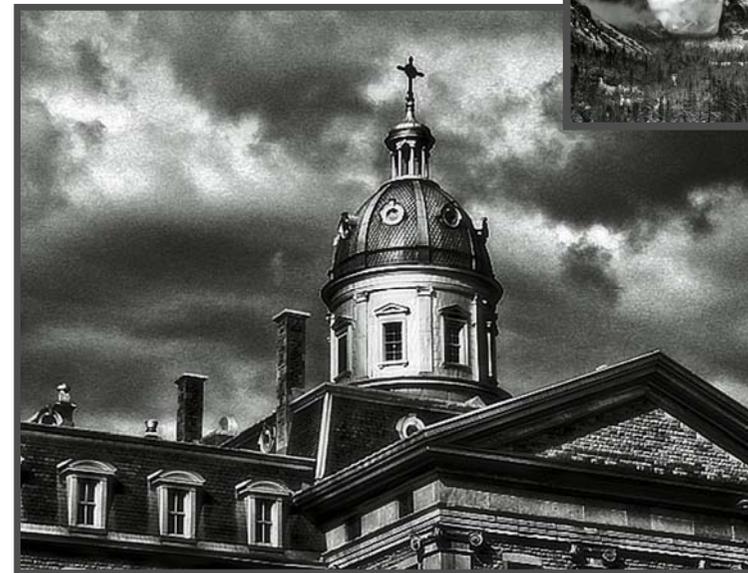
- 6 seconds for 1 MPixel or less
- 23 seconds for 4 MPixels
- multi-grid Poisson solver and fast bilateral filter [Paris 06]

Result

Model

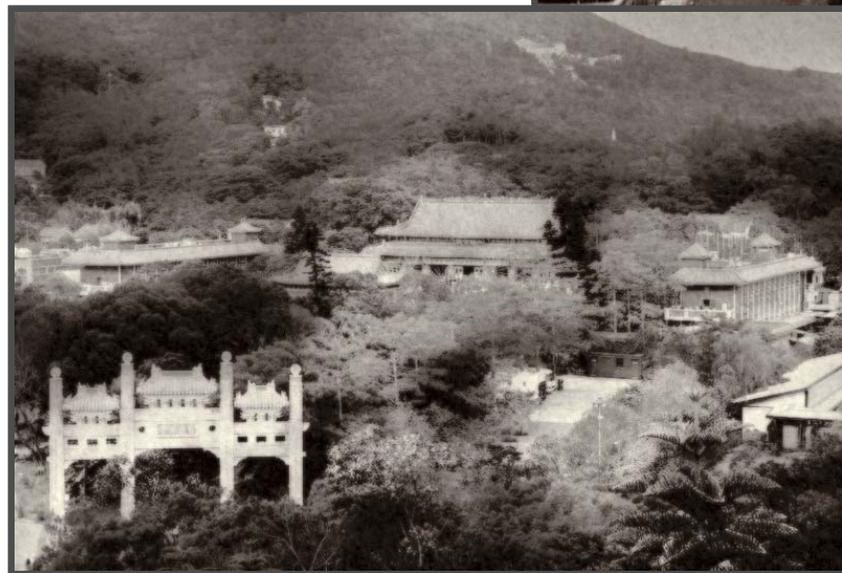


Result



Result

Model



Comparison with Naïve Histogram Matching



Local contrast, sharpness unfaithful

Comparison with Naïve Histogram Matching

DigiVFX



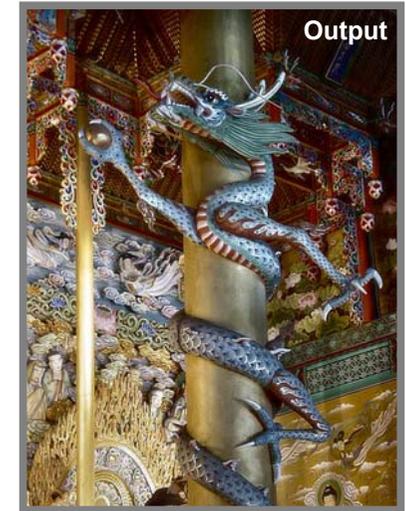
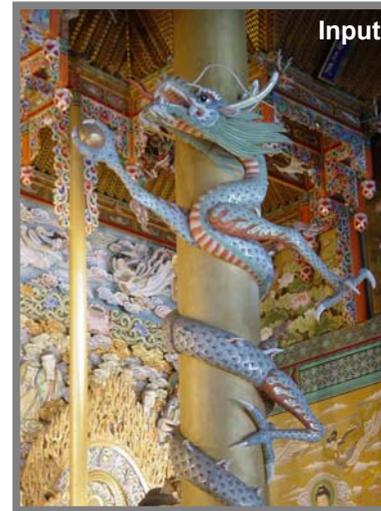
Local contrast too low



Color Images

DigiVFX

- Lab color space: modify only luminance

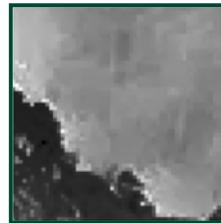


Limitations

DigiVFX

- 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

DigiVFX

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

Joint bilateral filtering

DigiVFX

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

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

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

DigiVFX

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



Overview

DigiVFX

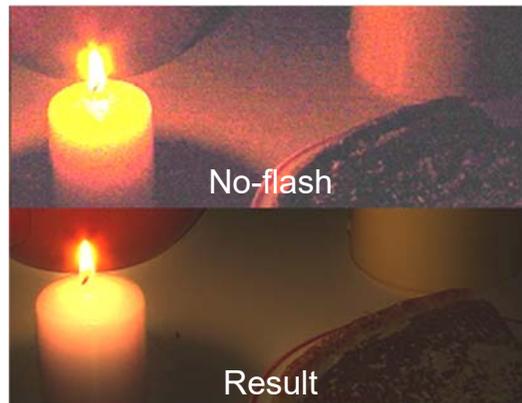
Basic approach of both flash/noflash papers

Remove noise + details
from image A,

Keep as image A Lighting

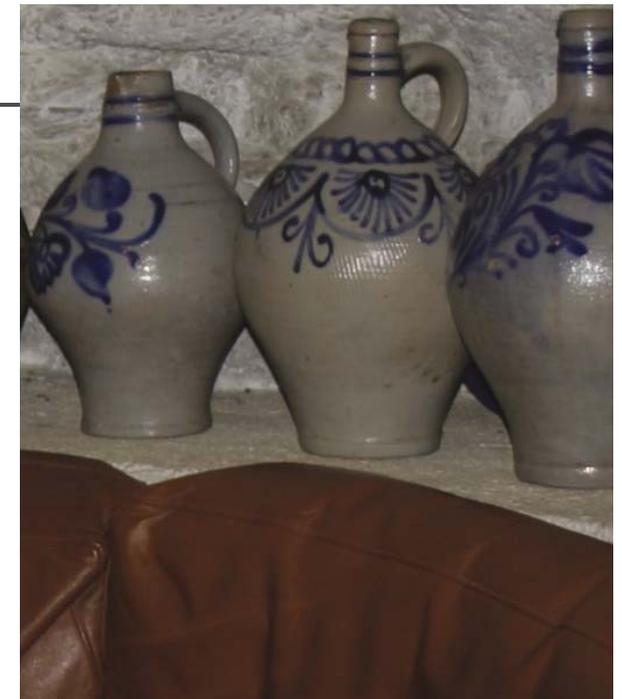
Obtain noise-free details
from image B,

Discard Image B Lighting



Petschnigg:

- Flash



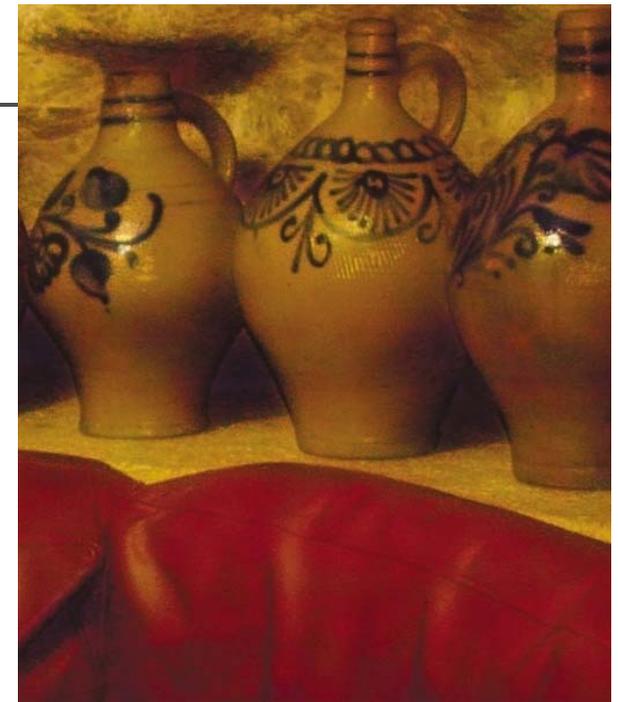
Petschnigg:

- No Flash,



Petschnigg:

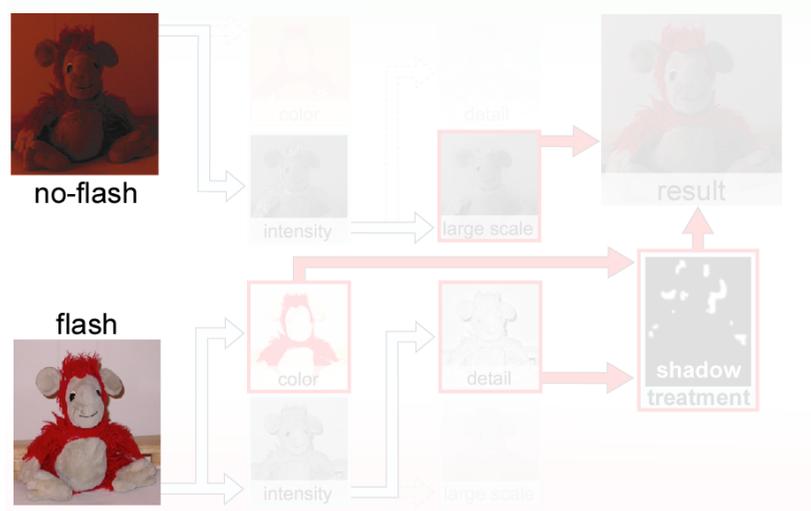
- Result



Our Approach

DigiVFX

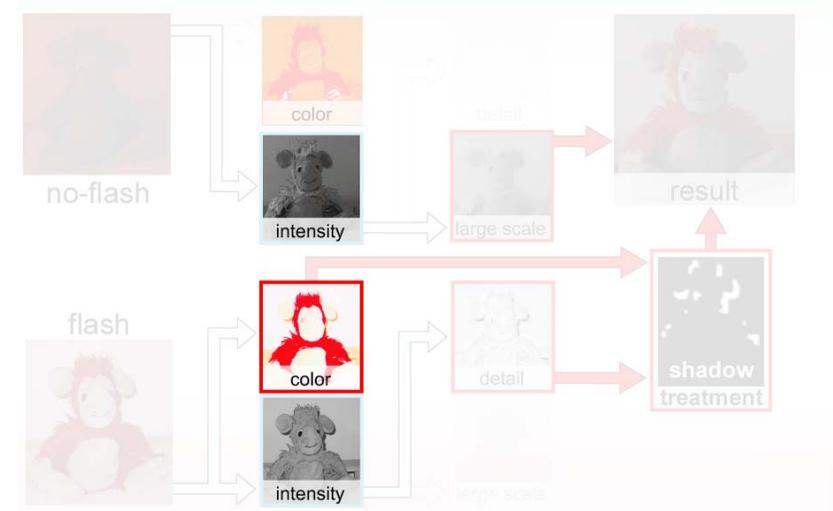
Registration



Our Approach

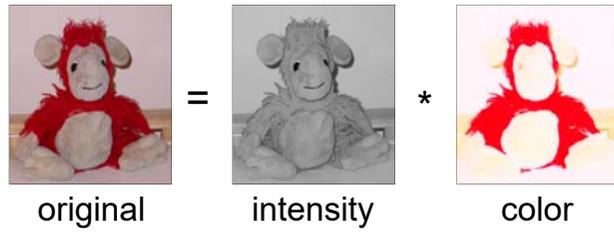
DigiVFX

Decomposition



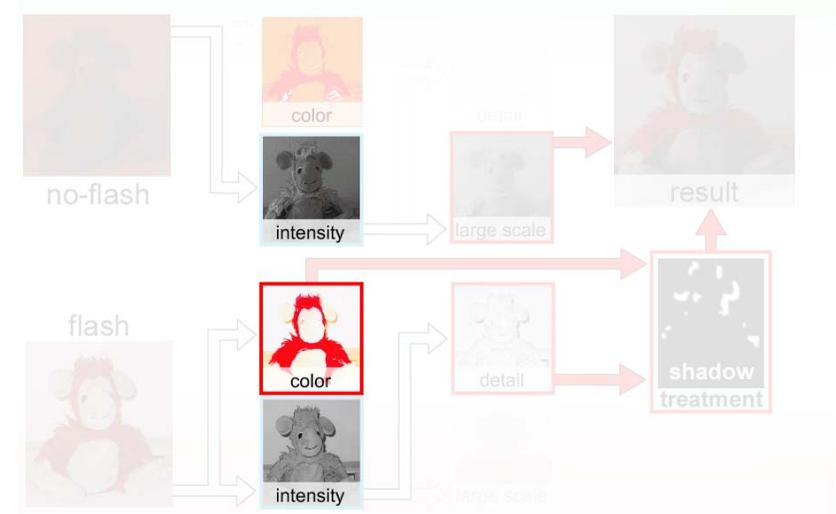
Decomposition

Color / Intensity:



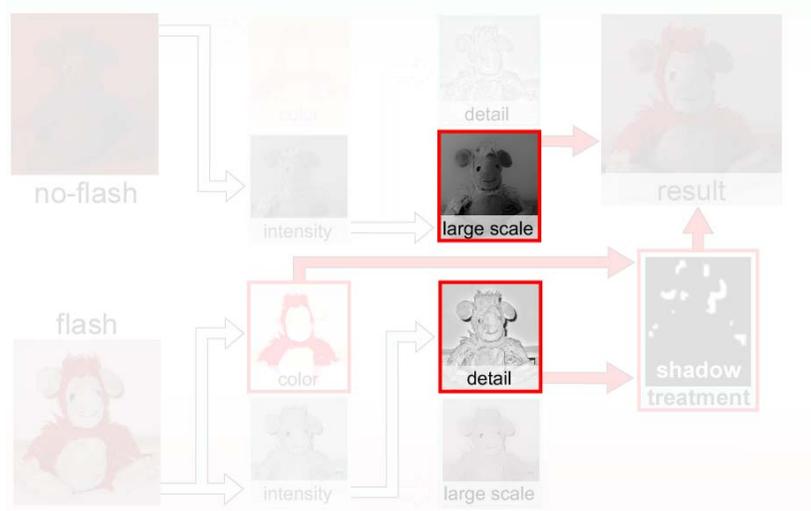
Our Approach

Decomposition



Our Approach

Decoupling



Decoupling

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

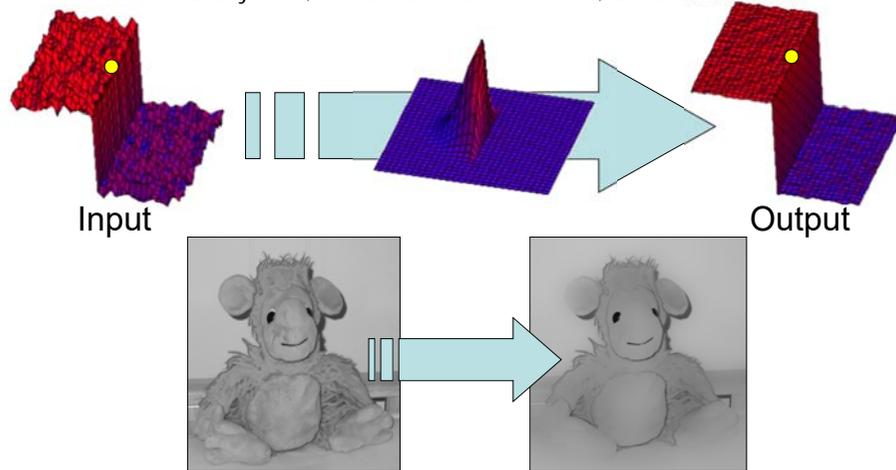


Large-scale Layer

DigiVFX

- Bilateral filter – edge preserving filter

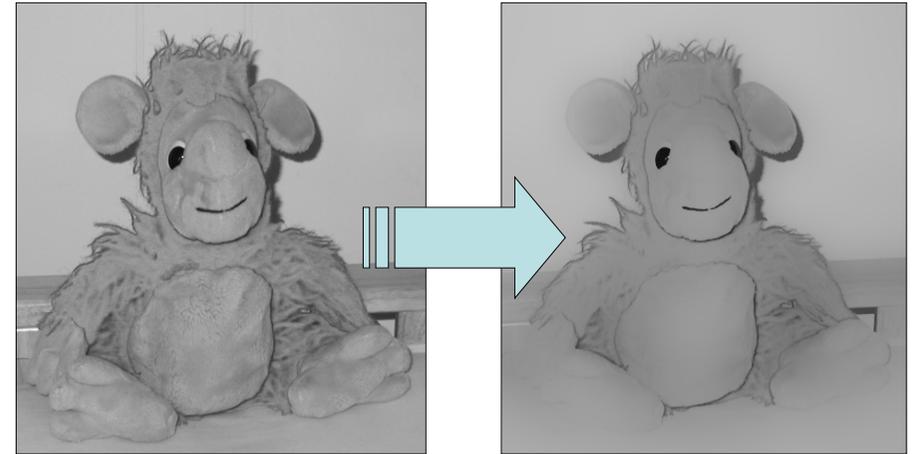
Smith and Brady 1997; Tomasi and Manducci 1998; Durand et al. 2002



Large-scale Layer

DigiVFX

- Bilateral filter



Cross Bilateral Filter

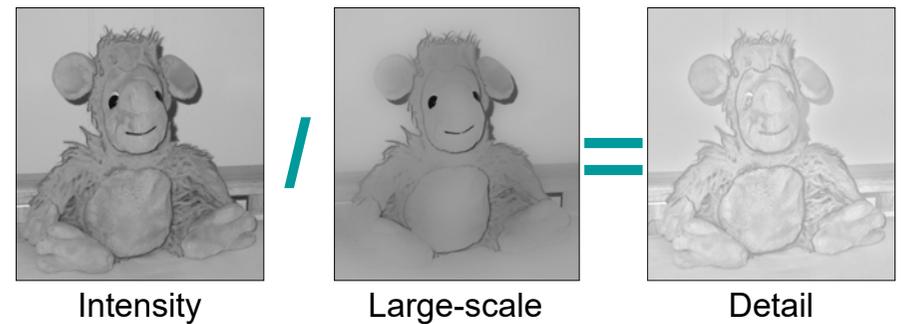
DigiVFX

- 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



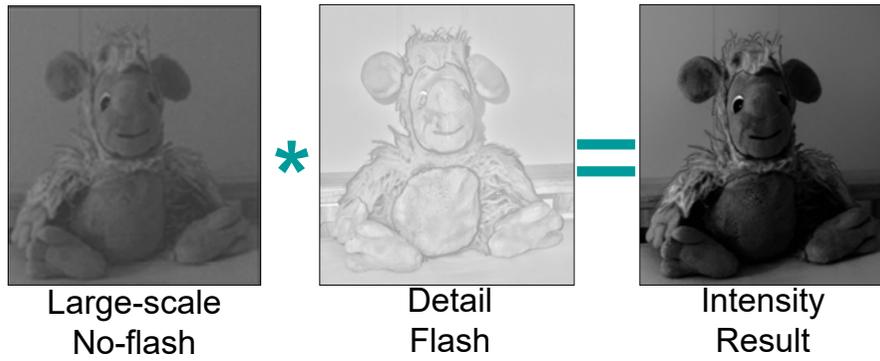
Detail Layer

DigiVFX



Recombination: Large scale * Detail = Intensity

Recombination



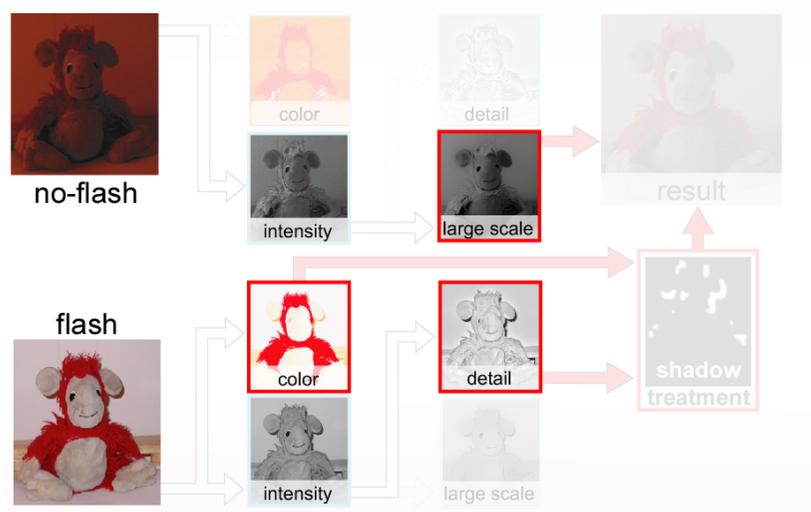
Recombination: Large scale * Detail = Intensity

Recombination



Recombination: Intensity * Color = Original

Our Approach



Results



No-flash



Flash

