




# 4. Filters



# Image Filtering

- An image is a 2D array of pixel values



[105	112	108	111	104	98	106	99	96	103	112	115	104	97	93	87]
[ 91	98	102	106	104	79	98	103	99	105	123	136	118	105	94	85]
[ 76	85	98	105	128	105	87	96	95	99	115	112	105	103	99	85]
[ 99	81	81	93	120	131	127	100	95	98	102	99	95	93	101	94]
[106	91	61	64	69	91	88	85	101	107	109	98	75	84	96	95]
[114	108	85	55	55	69	64	54	64	87	112	129	98	74	84	91]
[133	137	147	103	45	81	88	65	52	54	74	84	102	93	83	82]
[128	137	144	140	109	95	85	70	62	65	63	63	68	73	85	101]
[125	133	148	137	119	121	117	94	65	70	89	65	54	64	72	98]
[127	125	131	147	133	127	126	131	111	96	89	75	61	64	72	84]
[115	114	109	123	150	148	131	118	113	109	100	92	74	65	72	78]
[ 89	93	98	97	100	147	131	118	113	114	113	109	106	95	77	80]
[ 63	77	86	81	77	79	102	123	117	115	117	125	125	138	115	87]
[ 62	65	82	89	78	71	88	101	124	126	119	101	107	114	131	119]
[ 63	65	75	88	89	71	62	91	120	130	135	105	81	90	110	110]
[ 87	65	71	97	106	95	69	45	76	130	126	107	92	94	105	112]
[118	97	82	86	117	123	116	66	42	51	95	93	89	95	102	107]
[164	146	112	80	82	120	124	104	76	48	45	66	88	101	102	109]
[157	170	157	120	93	86	114	132	112	97	69	55	70	82	99	94]
[130	128	134	161	139	100	109	118	121	134	114	87	65	53	69	96]
[128	112	98	117	158	144	128	115	104	107	102	93	87	81	72	79]
[123	107	96	86	83	112	153	149	122	100	104	75	88	107	112	96]
[122	121	102	80	82	86	94	117	145	148	153	102	58	78	92	107]
[122	164	148	103	71	56	78	83	93	103	119	139	102	61	69	84]

- Filtering:
  - Replace each pixel by a *linear* combination of its neighbors
  - The combination is determined by the filter's *kernel*
  - Often spatially-invariant, the same kernel is applied to all pixel locations

$$g[\cdot, \cdot] = \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

Gaussian Filter

$$g[\cdot, \cdot] = \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

Box Filter

# Box filter example

$$g[\cdot, \cdot]$$

kernel

1	1	1
1	1	1
1	1	1

$\frac{1}{9}$

image  $f[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

output  $h[\cdot, \cdot]$

	0	10	20						

$$h[m, n] = \sum_{k, l} g[k, l] f[m + k, n + l]$$

output  $k, l$  filter image (signal)

# Box filter example

$$g[\cdot, \cdot]$$

kernel

1	1	1
1	1	1
1	1	1

$\frac{1}{9}$

image  $f[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0


output  $h[\cdot, \cdot]$

	0	10	20	30					

$$h[m, n] = \sum_{k, l} g[k, l] f[m + k, n + l]$$

output  $k, l$  filter image (signal)

# Gaussian vs box filtering



original

Which blur do you like better?



7x7 Gaussian



7x7 box



# How to create a soft shadow effect?


**CMU**



**CMU**

shift + overlay

**CMU**



Gaussian blur

# Fun with filters



input



filter

0	0	0
0	1	0
0	0	0

output



unchanged

input



filter

0	0	0
0	0	1
0	0	0

output



shift to left  
by one

# Detecting Edges

- How would you detect edges in an image (i.e., discontinuities in a function)?
  - Take derivatives: derivatives are large at discontinuities

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x + 0.5h) - f(x - 0.5h)}{h}$$

- How do you differentiate a discrete image (or any other discrete signal)?
  - Use finite differences

-1	0	1
----	---	---

1	0	-1
---	---	----


Remove limit and set  $h = 2$



# But Images are Noisy




intensity plot



Using a derivative filter:


derivative plot




# Blur Before Taking Derivatives

- When using derivative filters, it is critical to blur first!


input




Gaussian



blurred




derivative of  
blurred




# Derivative of Gaussian (DoG) filter

- Derivative theorem of convolution:  $\frac{\partial}{\partial x}(h \star f) = (\frac{\partial}{\partial x}h) \star f$
- Applying DoG = Applying blur first and then taking derivative


input




derivative of  
Gaussian



output (same  
as before)




# 2D Gaussian filters



$$h_{\sigma}(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$

Gaussian



$$\frac{\partial}{\partial x} h_{\sigma}(u, v)$$

Derivative of Gaussian



# The Sobel filter

Horizontal Sobel filter:

1	0	-1
2	0	-2
1	0	-1

Sobel filter

=

1
2
1

Blurring

\*

1	0	-1
---	---	----

1D derivative  
filter

Vertical Sobel filter:

1	2	1
0	0	0
-1	-2	-1

=

1
0
-1

\*


1	2	1
---	---	---



# Sobel filter example



original




horizontal Sobel filter




vertical Sobel filter




# Sobel filter example



original



horizontal Sobel filter



vertical Sobel filter

# Computing Image Gradients

1. Select a derivative filters (there are other similar filters, e.g. Scharr)

$$S_x = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}$$

$$S_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

2. Convolve with the image to compute derivatives.

$$\frac{\partial f}{\partial x} = S_x \otimes f$$

$$\frac{\partial f}{\partial y} = S_y \otimes f$$

3. Form the image gradient, and compute its direction and amplitude.

$$\nabla f = \left[ \frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right]$$

gradient

$$\theta = \tan^{-1} \left( \frac{\partial f}{\partial y} / \frac{\partial f}{\partial x} \right)$$

direction


$$\|\nabla f\| = \sqrt{\left( \frac{\partial f}{\partial x} \right)^2 + \left( \frac{\partial f}{\partial y} \right)^2}$$

amplitude




# Image gradient example


original



vertical derivative



gradient  
amplitude



horizontal  
derivative



# Questions?




# Bilateral filter

- Proposed by Tomasi and Manduci in ICCV 1998
  - “Bilateral Filtering for Gray and Color Images”
- A very good survey by Sylvain Paris et al. 2009
  - Published at Foundations and Trends in Computer Graphics and Vision
  - “Bilateral Filtering: Theory and Applications”
- Related to
  - SUSAN filter  
[Smith and Brady 95] <http://citeseer.ist.psu.edu/smith95susan.html>
  - Digital-TV [Chan, Osher and Chen 2001]  
<http://citeseer.ist.psu.edu/chan01digital.html>
  - sigma filter <http://www.geogr.ku.dk/CHIPS/Manual/f187.htm>

# Start with Gaussian Filter

- Here, input is a step function + noise
- Spatial Gaussian filter  $f$
- Output is blurred








# Gaussian Filter as Weighted Average

- Weight of  $\xi$  depends on its distance to  $x$

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



output

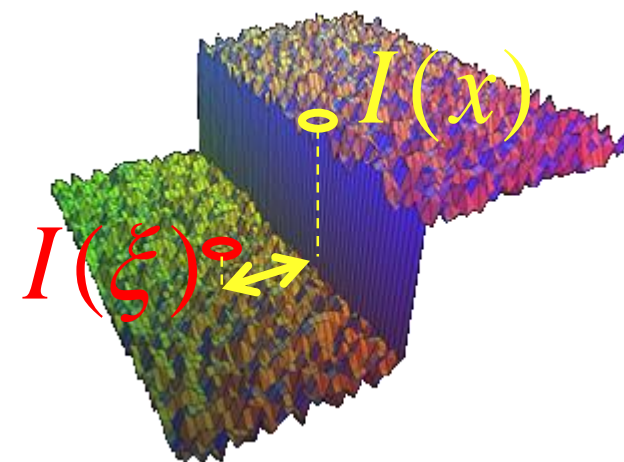




input


# The Problem of Edges

- Here,  $I(\xi)$  “pollutes” our estimated  $J(x)$
- It is too different to be averaged together

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



output




input

# Principle of Bilateral Filter


[Tomasi and Manduchi 1998]

- Penalty Gaussian  $g$  on the intensity difference

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



output




input

# Bilateral Filter


[Tomasi and Manduchi 1998]

- Spatial Gaussian  $f$

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



output



input




# Bilateral Filter

[Tomasi and Manduchi 1998]

- Combined weight

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




# Normalization Factor

[Tomasi and Manduchi 1998]

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

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






# Bilateral Filter is Spatially-Variant

[Tomasi and Manduchi 1998]


- The weights are different for each output pixel

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




# Other Explanation

- The bilateral filter uses the 3D distance to compute the weight



# Bilateral Filter vs Gaussian Filter

- Bilateral Filter – fix to the Gaussian filter
- Gaussian filtering applies the same filter everywhere



Same Gaussian kernel everywhere.


30



Sylvain  
Paris

# Bilateral Filter vs Gaussian Filter

- Adjust kernel based on image content



The kernel shape depends on the image content.

# Results: Denoise



noisy image



naïve denoising  
Gaussian blur



better denoising  
edge-preserving filter

Smoothing an image without blurring its edges.



# More Examples



(a) *Photograph*



(b) *Edge-aware smoothing*




(c) *Detail enhancement*



(d) *Stylization*



(e) *Recoloring*



(f) *Pencil drawing*



(g) *Depth-of-field*



# Questions?





# Guided Image Filtering

Kaiming He<sup>1</sup>, Jian Sun<sup>2</sup>, and Xiaoou Tang<sup>1,3</sup>

<sup>1</sup> Department of Information Engineering, The Chinese University of Hong Kong

<sup>2</sup> Microsoft Research Asia

<sup>3</sup> Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China

ECCV 2010

# Formulation

- Given an input image  $P$ , the output image  $Q$  is computed as,

$$Q_i = \sum_j W_{ij}(I) P_j$$

- The weights are computed from the guide image  $I$ ,


$$W_{ij}(I) = \frac{1}{|\omega|^2} \sum_{k:(i,j) \in \omega_k} \left( 1 + \frac{(I_i - \mu_k)(I_j - \mu_k)}{\sigma_k^2 + \epsilon} \right)$$

the total number of such windows containing both  $i$  and  $j$

the  $k$ -th local window containing both  $i$  and  $j$

$\mu_k, \sigma_k$  are the mean and variance of pixel values in the window  $\omega_k$

# Spatially-Variant Kernels



**Fig. 3.** Filter kernels. Top: a step edge (guided filter:  $r = 7, \epsilon = 0.1^2$ , bilateral filter:  $\sigma_s = 7, \sigma_r = 0.1$ ). Bottom: a textured patch (guided filter:  $r = 8, \epsilon = 0.2^2$ , bilateral filter:  $\sigma_s = 8, \sigma_r = 0.2$ ). The kernels are centered at the pixels denote by the red dots.


# Explanation

- When  $i$  and  $j$  are from different sides of an edge,

$$\frac{(I_i - \mu_k)(I_j - \mu_k)}{\sigma_k^2} \sim -1$$

- When  $i$  and  $j$  are from the same side,

$$\frac{(I_i - \mu_k)(I_j - \mu_k)}{\sigma_k^2} \sim 1$$



$$W_{ij}(I) = \frac{1}{|\omega|^2} \sum_{k:(i,j) \in \omega_k} \left( 1 + \frac{(I_i - \mu_k)(I_j - \mu_k)}{\sigma_k^2 + \epsilon} \right)$$



# Better Than Bilateral Filter

- $O(N)$  time,  $N$  is the number of pixels
  - Bilateral filter is  $O(Nr^2)$ ,  $r$  is the local window size

$$Q_i = \sum_j W_{ij}(I) P_j$$

$$W_{ij}(I) = \frac{1}{|\omega|^2} \sum_{k:(i,j) \in \omega_k} \left( 1 + \frac{(I_i - \mu_k)(I_j - \mu_k)}{\sigma_k^2 + \epsilon} \right)$$

$\mu_k, \sigma_k$  can be computed with box filter ( $O(N)$ ).