

Filters

CS 554 – Computer Vision



Pinar Duygulu

Bilkent University

Today's topics

- Image Formation
- Image filters in spatial domain
 - Filter is a mathematical operation of a grid of numbers
 - Smoothing, sharpening, measuring texture
- Image filters in the frequency domain
 - Filtering is a way to modify the frequencies of images
 - Denoising, sampling, image compression
- Templates and Image Pyramids
 - Filtering is a way to match a template to the image
 - Detection, coarse-to-fine registration

Images as functions



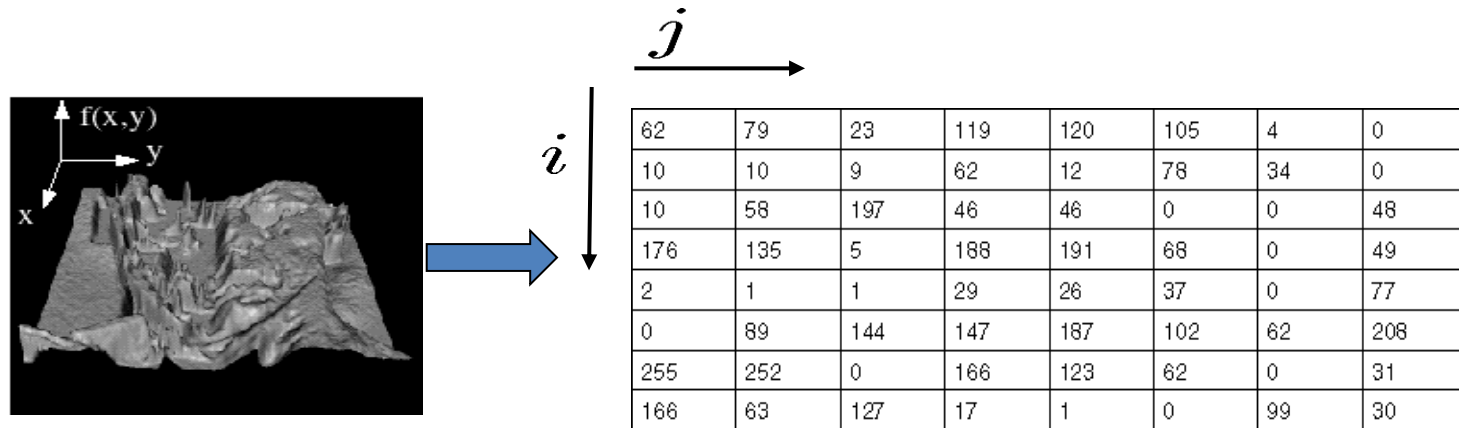
Images as functions

- We can think of an image as a function, f , from \mathbb{R}^2 to \mathbb{R} :
 - $f(x, y)$ gives the intensity at position (x, y)
 - Realistically, we expect the image only to be defined over a rectangle, with a finite range:
 - $f: [a, b] \times [c, d] \rightarrow [0, 255]$
- A color image is just three functions pasted together. We can write this as a “vector-valued” function:


$$f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$$

Digital images

- In computer vision we operate on **digital (discrete)** images:
 - **Sample** the 2D space on a regular grid
 - **Quantize** each sample (round to nearest integer)
- Image thus represented as a matrix of integer values.



2D



1D

Images as discrete functions

- Cartesian Coordinates

$$f[n, m] = \begin{bmatrix} \ddots & & \vdots & & \\ & f[-1, 1] & f[0, 1] & f[1, 1] & \\ \dots & f[-1, 0] & \underline{f[0, 0]} & f[1, 0] & \dots \\ & f[-1, -1] & f[0, -1] & f[1, -1] & \\ & & \vdots & & \ddots \end{bmatrix}$$

Today's topics

- Image Formation
- Image filters in spatial domain
 - Filter is a mathematical operation of a grid of numbers
 - Smoothing, sharpening, measuring texture
- Image filters in the frequency domain
 - Filtering is a way to modify the frequencies of images
 - Denoising, sampling, image compression
- Templates and Image Pyramids
 - Filtering is a way to match a template to the image
 - Detection, coarse-to-fine registration

Zebras vs. Dalmatians




Both zebras and dalmatians have black and white pixels in about the same number

- if we shuffle the images point-wise processing is not affected

Need to measure properties relative to small *neighborhoods* of pixels

- find different image patterns

Filtering





We want to remove unwanted sources of variation, and keep the information relevant for whatever task we need to solve



Filters

- **Filtering:**
 - Form a new image whose pixels are a combination of original pixel values
 - compute function of local neighborhood at each position
- **Goals:**
- Extract useful information from the images
Features (textures, edges, corners, distinctive points, blobs...)
- Modify or enhance image properties:
super-resolution; in-painting; de-noising, resizing
- Detect patterns
Template matching



Smooth/Sharpen Images...



Find edges...




Find waldo...

De-noising



Salt and pepper noise

Super-resolution



In-painting



Bertamio et al

Common types of noise

- **Salt and pepper noise:** random occurrences of black and white pixels
- **Impulse noise:** random occurrences of white pixels
- **Gaussian noise:** variations in intensity drawn from a Gaussian normal distribution



Original



Salt and pepper noise






Impulse noise



Gaussian noise

Gaussian noise



$$f(x, y) = \underbrace{\hat{f}(x, y)}_{\text{Ideal Image}} + \underbrace{\eta(x, y)}_{\text{Noise process}}$$

Gaussian i.i.d. ("white") noise:
 $\eta(x, y) \sim \mathcal{N}(\mu, \sigma)$

```
>> noise = randn(size(im)).*sigma;
```


```
>> output = im + noise;
```

First attempt at a solution

- Let's replace each pixel with an average of all the values in its neighborhood
- Assumptions:
 - Expect pixels to be like their neighbors
 - Expect noise processes to be independent from pixel to pixel


First attempt at a solution

- Let's replace each pixel with an average of all the values in its neighborhood
- Moving average in 1D:




Weighted Moving Average

- Can add weights to our moving average
- *Weights* $[1, 1, 1, 1, 1] / 5$



Weighted Moving Average

- Non-uniform weights $[1, 4, 6, 4, 1] / 16$



Moving Average In 2D

$F[x, y]$

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

$G[x, y]$

	0								

Moving Average In 2D

$F[x, y]$

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

$G[x, y]$

	0	10							

Moving Average In 2D

$F[x, y]$

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

$G[x, y]$

	0	10	20						

Moving Average In 2D

$$F[x, y]$$

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

$$G[x, y]$$

	0	10	20	30					

Moving Average In 2D

$$F[x, y]$$

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

$$G[x, y]$$

	0	10	20	30	30				

Moving Average In 2D

$$F[x, y]$$

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

$$G[x, y]$$

	0	10	20	30	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	60	90	90	90	60	30	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	10	20	30	30	30	30	20	10	
	10	10	10	0	0	0	0	0	

Correlation filtering

Say the averaging window size is $2k+1 \times 2k+1$:

$$G[i, j] = \underbrace{\frac{1}{(2k+1)^2}}_{\text{Attribute uniform weight to each pixel}} \underbrace{\sum_{u=-k}^k \sum_{v=-k}^k F[i+u, j+v]}_{\text{Loop over all pixels in neighborhood around image pixel } F[i,j]}$$

Now generalize to allow different weights depending on neighboring pixel's relative position:

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k \underbrace{H[u, v]}_{\text{Non-uniform weights}} F[i+u, j+v]$$

Correlation filtering

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i + u, j + v]$$

This is called cross-correlation, denoted $G = H \otimes F$


Filtering an image: replace each pixel with a linear combination of its neighbors.

The filter “kernel” or “mask” $H[u, v]$ is the prescription for the weights in the linear combination.

Averaging Filter



original




?

Averaging Filter




original

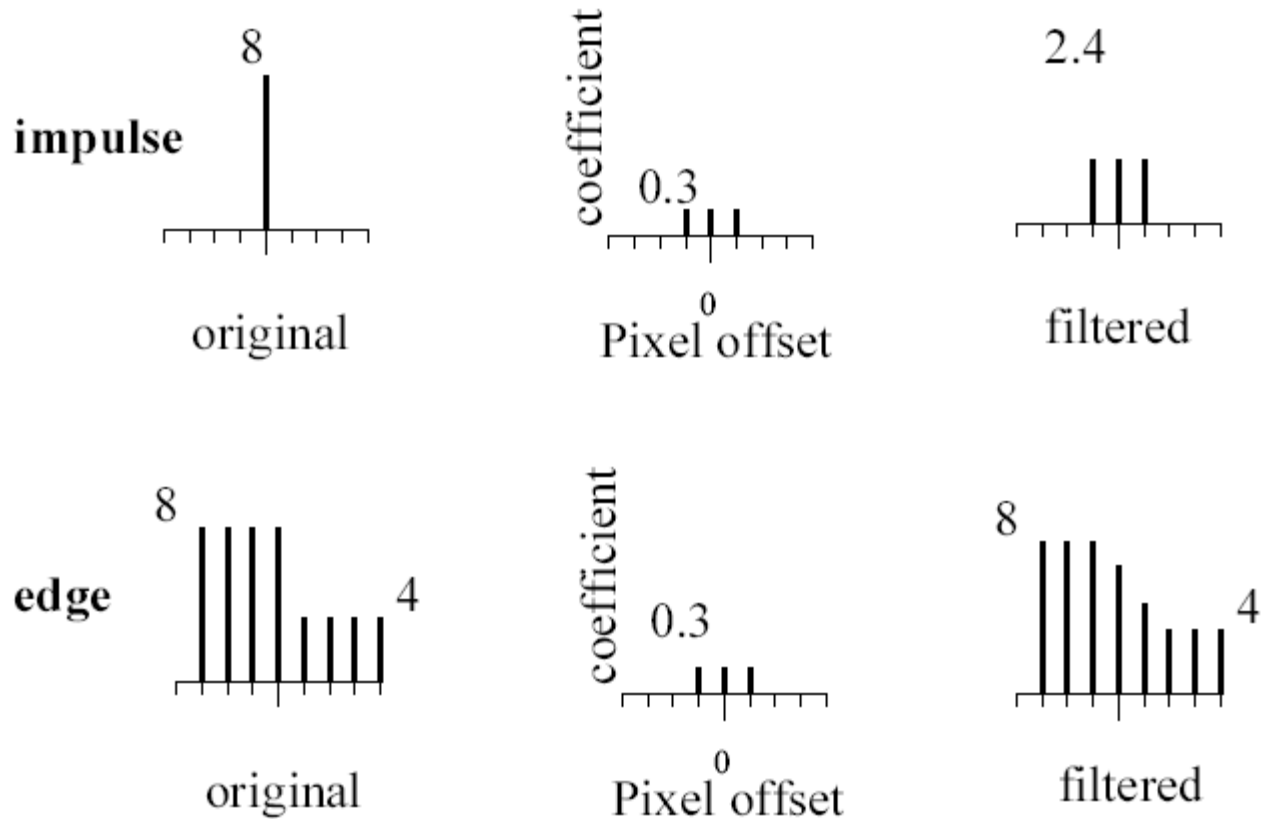


Blurred (filter applied in both dimensions).

Averaging Filter




Averaging Filter



Averaging filter

- What values belong in the kernel H for the moving average example?



$$G = H \otimes F$$


Smoothing by averaging



depicts box filter:
white = high value, black = low value



original



filtered


Example



Example




Example




Smoothing by averaging



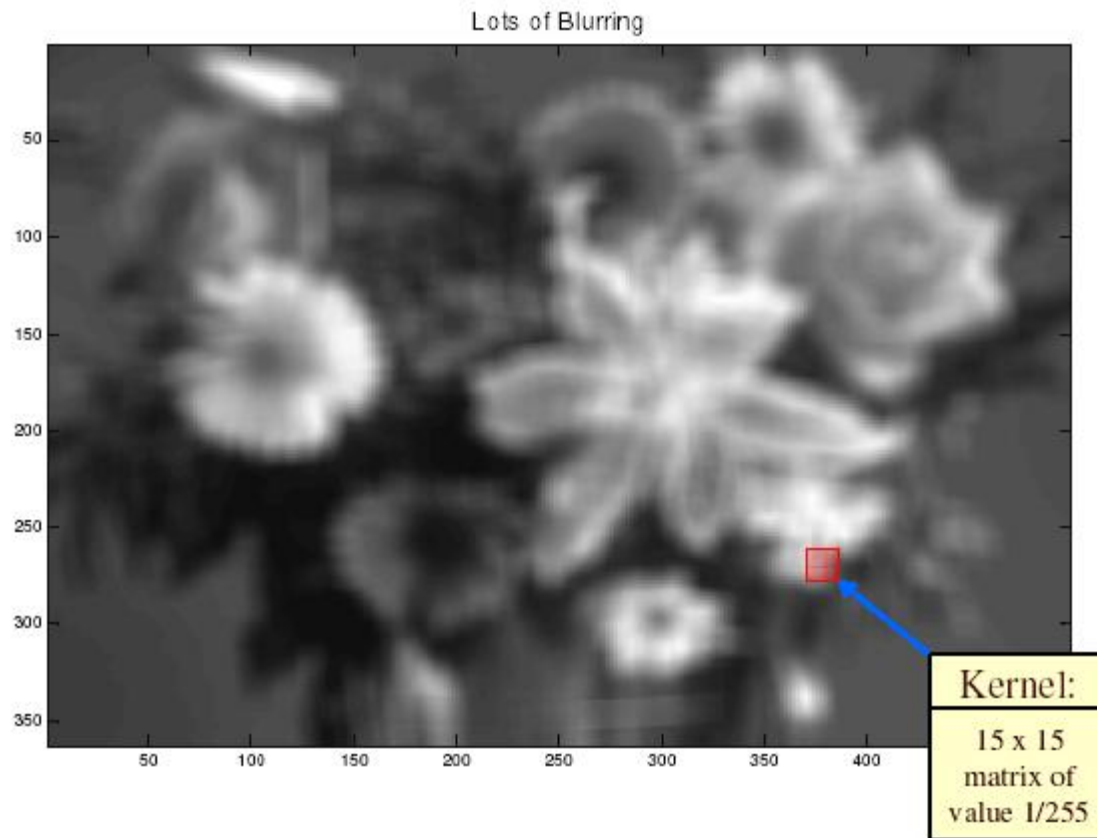
Smoothing by averaging



Smoothing by averaging



Smoothing by averaging



Gaussian filter

- What if we want nearest neighboring pixels to have the most influence on the output?

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

$F[x, y]$


$\frac{1}{16}$

1	2	1
2	4	2
1	2	1


$H[u, v]$

This kernel is an approximation of a Gaussian function:


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



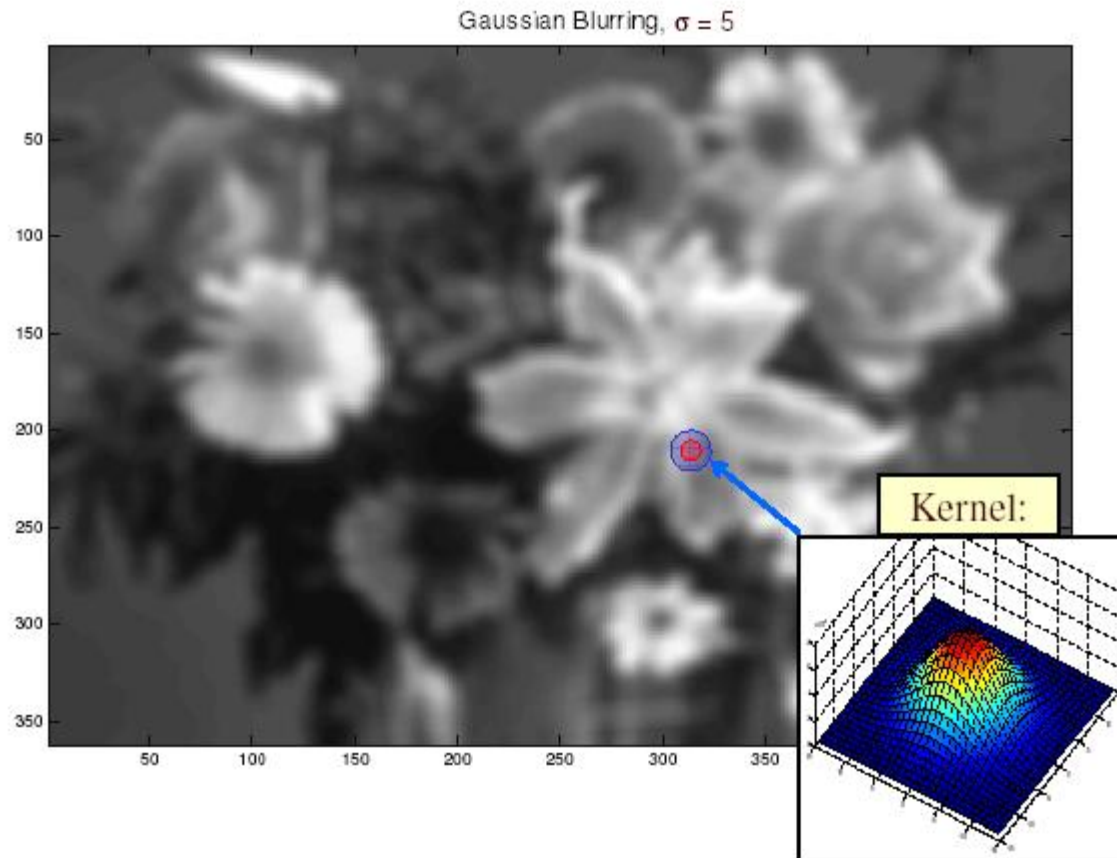
A weighted average that weights pixels at its center much more strongly than its boundaries



Smoothing with a Gaussian



Smoothing with a Gaussian




Smoothing with a Gaussian




Result of blurring using
a uniform local model

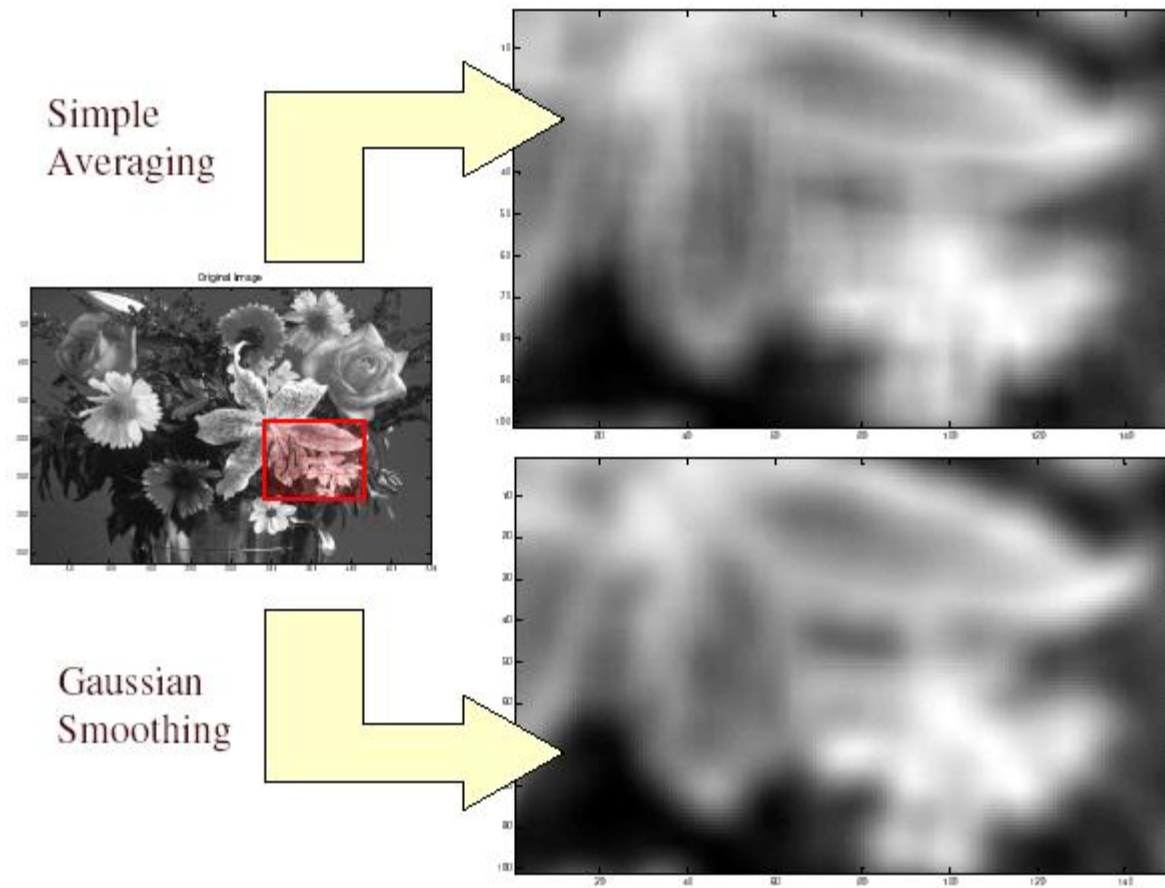
Produces a set of
narrow vertical
horizontal and vertical
bars – ringing effect



Result of blurring
using a set of
Gaussian weights




Smoothing with a Gaussian




Gaussian filters

- What parameters matter here?
- **Size** of kernel or mask
 - Note, Gaussian function has infinite support, but discrete filters use finite kernels




$\sigma = 5$ with 10
x 10 kernel




$\sigma = 5$ with 30
x 30 kernel

Gaussian filters

- What parameters matter here?
- **Variance** of Gaussian: determines extent of smoothing



$\sigma = 2$ with 30
x 30 kernel




$\sigma = 5$ with 30
x 30 kernel

Smoothing with a Gaussian

If σ is small : the smoothing will have little effect



If σ is larger : neighboring pixels will have larger weights resulting in consensus of the neighbors

If σ is very large : details will disappear along with the noise




Gaussian filter


$$G(x, y; \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$



$\sigma=1$






$\sigma=2$



$\sigma=4$

Gaussian smoothing to remove noise


Image
Noise




$$f(x, y) = \overbrace{\hat{f}(x, y)}^{\text{Ideal Image}} + \overbrace{\eta(x, y)}^{\text{Noise process}}$$

Gaussian i.i.d. ("white") noise:
 $\eta(x, y) \sim \mathcal{N}(\mu, \sigma)$


Gaussian smoothing to remove noise



No smoothing




$\sigma = 2$



$\sigma = 4$

Smoothing with a Gaussian




The effects of smoothing

Each row shows smoothing with gaussians of different width; each column shows different realizations of an image of gaussian noise.

Smoothing with a Gaussian

- Filtered noise is sometimes useful
 - looks like some natural textures, can be used to simulate fire, etc.



Gaussian kernel

$$g(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(x^2 + y^2)}{2\sigma^2}\right)$$

$$\begin{bmatrix} 0.0751 & 0.1238 & 0.0751 \\ 0.1238 & 0.242 & 0.1238 \\ 0.0751 & 0.1238 & 0.0751 \end{bmatrix}$$

Gaussian is an approximation to the binomial distribution.

Can approximate Gaussian using binomial coefficients.

$$a_{nr} \equiv \frac{n!}{r!(n-r)!} \equiv \binom{n}{r}$$

n = number of elements in the 1D filter minus 1

r = position of element in the filter kernel (0, 1, 2...)

1X3 filter: $n=(3-1)=2$, $r=0,1,2$

$$g = 1/4 \quad \begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$


$$\begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$

$$g' \cdot g = \begin{bmatrix} 0.0625 & 0.1250 & 0.0625 \\ 0.1250 & 0.2500 & 0.1250 \\ 0.0625 & 0.1250 & 0.0625 \end{bmatrix}$$

Matlab

```
>> hsize = 10;  
>> sigma = 5;  
>> h = fspecial('gaussian' hsize, sigma);
```


```
>> mesh(h);
```



```
>> imagesc(h);
```





```
>> outim = imfilter(im, h);  
>> imshow(outim);
```




outim

Smoothing with a Gaussian

Parameter σ is the “scale” / “width” / “spread” of the Gaussian kernel, and controls the amount of smoothing.



...



```
for sigma=1:3:10
    h = fspecial('gaussian', fsize, sigma);
    out = imfilter(im, h);
    imshow(out);
    pause;
end
```



Convolution

- Convolution:
 - Flip the filter in both dimensions (bottom to top, right to left)
 - Then apply cross-correlation

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i - u, j - v]$$

$$G = H \star F$$

↑
*Notation for
convolution
operator*



Convolution vs. correlation

Convolution

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i - u, j - v]$$

$$G = H \star F$$

Cross-correlation


$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i + u, j + v]$$


$$G = H \otimes F$$


For a Gaussian or box filter, how will the outputs differ?

If the input is an impulse signal, how will the outputs differ?

Predict the filtered outputs


$$\begin{matrix} * & \begin{matrix} \begin{matrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{matrix} \end{matrix} & = ? \end{matrix}$$



$$\begin{matrix} * & \begin{matrix} \begin{matrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{matrix} \end{matrix} & = ? \end{matrix}$$


$$\begin{matrix} * & \begin{matrix} \begin{matrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{matrix} \end{matrix} & - \frac{1}{9} & \begin{matrix} \begin{matrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{matrix} \end{matrix} & = ? \end{matrix}$$

Practice with linear filters



original




?

Practice with linear filters



original



Filtered
(no change)

Practice with linear filters



Original


0	0	0
0	1	0
0	0	0



Filtered
(no change)

Impulse

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




$g[m,n]$

\otimes

0	0	0	0	0
0	0	0	0	0
0	0	1	0	0
0	0	0	0	0
0	0	0	0	0

$h[m,n]$

=



$f[m,n]$

Practice with linear filters



Original

0	0	0
0	0	1
0	0	0

?

Practice with linear filters



Original


0	0	0
0	0	1
0	0	0



Shifted left
by 1 pixel
with
correlation

Shifts

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




$g[m,n]$

\otimes

0	0	0	0	0
0	0	0	0	0
0	0	0	0	1
0	0	0	0	0
0	0	0	0	0

$h[m,n]$

=




$f[m,n]$


Practice with linear filters



original




—



?


Practice with linear filters




Practice with linear filters



original



—



?


Sharpening



original



—





Sharpened
original

Sharpening



original



Sharpened
(differences are
accentuated; constant
areas are left untouched).

Practice with linear filters



Original

0	0	0
0	2	0
0	0	0


— $\frac{1}{9}$

1	1	1
1	1	1
1	1	1




Sharpening filter


- Accentuates differences with
local average



Filtering examples: sharpening



before



after

Rectangular filter



$g[m,n]$




$h[m,n]$

=



$f[m,n]$

What does blurring take away?



-



=




- Let's add it back:



+ a



=



Rectangular filter



$g[m,n]$



$h[m,n]$



$f[m,n]$

Rectangular filter



$g[m,n]$



$h[m,n]$



$f[m,n]$

Integral image



Shift invariant linear system

- **Shift invariant:**

- Operator behaves the same everywhere, i.e. the value of the output depends on the pattern in the image neighborhood, not the position of the neighborhood.

- **Linear:**

- Superposition: $h * (f_1 + f_2) = (h * f_1) + (h * f_2)$
- Scaling: $h * (k f) = k (h * f)$

Properties of convolution

- Linear & shift invariant

- Commutative:

$$f * g = g * f$$

- Associative

$$(f * g) * h = f * (g * h)$$

- Identity:

$$\text{unit impulse } e = [\dots, 0, 0, 1, 0, 0, \dots]. \quad f * e = f$$

- Differentiation:


$$\frac{\partial}{\partial x}(f * g) = \frac{\partial f}{\partial x} * g$$

Separability

- In some cases, filter is separable, and we can factor into two steps:
 - Convolve all rows
 - Convolve all columns

Separability

- In some cases, filter is separable, and we can factor into two steps: e.g.,



What is the computational complexity advantage for a separable filter of size $k \times k$, in terms of number of operations per output pixel?

$$\begin{array}{rcl}
 \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline 1 \\ \hline \end{array} \times \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array} & = & \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline 2 & 3 & 3 \\ \hline 3 & 5 & 5 \\ \hline 4 & 4 & 6 \\ \hline \end{array} \\
 & & \begin{array}{l} = 2 + 6 + 3 = 11 \\ = 6 + 20 + 10 = 36 \\ = 4 + 8 + 6 = 18 \\ \hline 65 \end{array}
 \end{array}$$

Advantages of separability

First convolve the image with a one dimensional horizontal filter

Then convolve the result of the first convolution with a one dimensional vertical filter

For a $k \times k$ Gaussian filter, 2D convolution requires k^2 operations per pixel

But using the separable filters, we reduce this to $2k$ operations per pixel.

Seperable Gaussian

$$g(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-x^2 / (2\sigma^2))$$

$$g(y) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-y^2 / (2\sigma^2))$$

Product?

$$g(x, y) = \frac{1}{2\pi\sigma^2} \exp(-(x^2 + y^2) / (2\sigma^2))$$

Advantages of Gaussians

- Convolution of a Gaussian with itself is another Gaussian
 - so we can first smooth an image with a small Gaussian
 - then, we convolve that smoothed image with another small Gaussian and the result is equivalent to smoothing the original image with a larger Gaussian.
 - If we smooth an image with a Gaussian having sd σ twice, then we get the same result as smoothing the image with a Gaussian having standard deviation $(2\sigma)^{1/2}$

Effect of smoothing filters

5x5




Additive Gaussian noise



Salt and pepper noise


Median filter





- No new pixel values introduced
- Removes spikes: good for impulse, salt & pepper noise

Median filter

Salt and
pepper noise →




← Median
filtered



Plots of a row of the image




Median filter

- Median filter is edge preserving




Boundary issues

- What is the size of the output?
- MATLAB: `filter2(g, f, shape)`
 - *shape* = 'full': output size is sum of sizes of *f* and *g*
 - *shape* = 'same': output size is same as *f*
 - *shape* = 'valid': output size is difference of sizes of *f* and *g*



Boundary issues


- What about near the edge?
 - the filter window falls off the edge of the image
 - need to extrapolate
 - methods:
 - clip filter (black)
 - wrap around
 - copy edge
 - reflect across edge




Boundary issues

- What about near the edge?
 - the filter window falls off the edge of the image
 - need to extrapolate
 - methods (MATLAB):
 - clip filter (black): `imfilter(f, g, 0)`
 - wrap around: `imfilter(f, g, 'circular')`
 - copy edge: `imfilter(f, g, 'replicate')`
 - reflect across edge: `imfilter(f, g, 'symmetric')`

Borders




zero




wrap




clamp




mirror




blurred: zero



normalized zero



clamp



mirror