

CAP 4453

Robot Vision

Dr. Gonzalo Vaca-Castaño

Gonzalo.vacacastano@ucf.edu




Administrative details

- COVID case in class. Mask is mandatory.
- Thursday class will be via zoom due to football game
- Homework 1 issues ?
- Linear algebra



Questions?



Robot Vision

3. Image Filtering



Credits

- Some slides comes directly from:
 - Yogesh S Rawat (UCF)
 - Noah Snavely (Cornell)
 - Ioannis (Yannis) Gkioulekas (CMU)
 - Mubarak Shah (UCF)
 - S. Seitz
 - James Tompkin
 - Ulas Bagci



Outline


- ~~Image as a function~~
- Extracting useful information from Images
 - Histogram
 - Filtering (linear)
 - Smoothing/Removing noise
 - Convolution/Correlation
 - Image Derivatives/Gradient
 - Edges
- Colab Notes/ homeworks
- Read Szeliski, Chapter 3.
- Read/Program CV with Python, Chapter 1.

What is an image?

- We can think of a (grayscale) image as a **function, f** , from \mathbb{R}^2 to \mathbb{R} :
 - $f(x, y)$ gives the **intensity** at position (x, y)



snoop



3D view

- A **digital** image is a discrete (**sampled, quantized**) version of this function

Image transformations

- As with any function, we can apply operators to an image



$$g(x,y) = f(x,y) + 20$$

$$g(x,y) = f(-x,y)$$

- Today we'll talk about a special kind of operator, *convolution* (linear filtering)




Filters

- Filtering
 - Form a new image whose pixels are a combination of the original pixels
- Why?
 - To get useful information from images
 - E.g., extract edges or contours (to understand shape)
 - To enhance the image
 - E.g., to remove noise
 - E.g., to sharpen and “enhance image” a la CSI
 - A key operator in Convolutional Neural Networks

Question: Noise reduction


- Given a camera and a still scene, how can you reduce noise?



Take lots of images and average them!

Question: Noise reduction

- Given a camera and a still scene, how can you reduce noise?



Take lots of images and average them!

Can we something else?

CAP4453


Thresholding !



$$g(m, n) = \begin{cases} 255, & f(m, n) > A \\ 0 & \textit{otherwise} \end{cases}$$

Question: Noise reduction

- This is not a gray scale image



```
import cv2
import os
import numpy as np
import matplotlib.pyplot as plt

folder='C:/Users/gonza/OneDrive/Teaching/CAP4453/class3/'
list_dir = [fil for fil in os.listdir(folder) if fil[-3:]=='jpg']

for iFile, fname in enumerate(list_dir):
    if iFile == 0:
        sumFile = cv2.imread(folder + fname)
        sumFile = sumFile.astype(np.float)
    else:
        sumFile = sumFile + cv2.imread(folder + fname).astype(np.float)

sumFile = sumFile/len(list_dir)
sumFile[sumFile>90]=255
sumFile[sumFile<=90]=0

plt.imshow(sumFile.astype(np.uint8))
```



Image noise

- Light Variations
- Camera Electronics
- Surface Reflectance
- Lens

- Noise is random,
 - it occurs with some probability
 - It has a distribution

Additive Noise

$$I_{observed}(x, y) = I_{original}(x, y) + n(x, y)$$

True pixel value at x,y

Noise at x,y



Multiplicative Noise

$$I_{observed}(x, y) = I_{original}(x, y) \times n(x, y)$$


True pixel value at x,y

Noise at x,y



Gaussian Noise


$$n(x, y) \approx g(n) = e^{\frac{-n^2}{2\sigma^2}}$$



Probability Distribution
 n is a random variable



Gaussian function



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


Salt and pepper noise

- Each pixel is randomly made black or white with a uniform probability distribution



Salt-pepper

Uniform distribution





Noise implementation

```
#Parameters
#-----
#image : ndarray
#   Input image data. Will be converted to float.
#mode : str
#   One of the following strings, selecting the type of noise to add:

#   'gauss'      Gaussian-distributed additive noise.
#   'poisson'    Poisson-distributed noise generated from the data.
#   's&p'        Replaces random pixels with 0 or 1.
#   'speckle'    Multiplicative noise using out = image + n*image,where
#               n,is uniform noise with specified mean & variance.
import numpy as np
import os
import cv2

def noisy(noise_typ,image):
    if noise_typ == "gauss":
        row,col,ch= image.shape
        mean = 0
        var = 1
        sigma = var**0.5
        gauss = np.random.normal(mean,sigma,(row,col,ch))
        gauss = gauss.reshape(row,col,ch)
        noisy = image + gauss
        return noisy
    elif noise_typ == "s&p":
        row,col,ch = image.shape
        s_vs_p = 0.5
        amount = 0.004
        out = image
        # Salt mode
        num_salt = np.ceil(amount * image.size * s_vs_p)
        coords = [np.random.randint(0, i - 1, int(num_salt))
                  for i in image.shape]
        out[coords] = 1

        # Pepper mode
        num_pepper = np.ceil(amount* image.size * (1. - s_vs_p))
        coords = [np.random.randint(0, i - 1, int(num_pepper))
                  for i in image.shape]
        out[coords] = 0
        return out
    elif noise_typ == "poisson":
        vals = len(np.unique(image))
        vals = 2 ** np.ceil(np.log2(vals))
        noisy = np.random.poisson(image * vals) / float(vals)
        return noisy
    elif noise_typ == "speckle":
        row,col,ch = image.shape
        gauss = np.random.randn(row,col,ch)
        gauss = gauss.reshape(row,col,ch)
        noisy = image + image * gauss
        return noisy
```



Outline

- ~~Image as a function~~
- Extracting useful information from Images
 - Histogram
 - **Filtering (linear)**
 - Smoothing/Removing noise
 - Convolution/Correlation
 - Image Derivatives/Gradient
 - Edges
- Colab Notes/ homeworks
- Read Szeliski, Chapter 3.
- Read/Program CV with Python, Chapter 1.



Linear shift-invariant image filtering

- Replace each pixel by a *linear* combination of its neighbors (and possibly itself).
- The combination is determined by the filter's *kernel*.
- The same kernel is *shifted* to all pixel locations so that all pixels use the same linear combination of their neighbors.

Filtering

- Modify pixels based on some function of neighborhood




Image filtering

- Image filtering: compute function of local neighborhood at each position

h=output f=filter I=image

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

2d coords=k,l 2d coords=m,n

[] [] []



Image filtering

- Image filtering: compute function of local neighborhood at each position
- Enhance images
 - Denoise, resize, increase contrast, etc.
- Extract information from images
 - Texture, edges, distinctive points, etc.
- Detect patterns
 - Template matching

Let's run the box filter

Box filter




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 $h[\cdot, \cdot]$

note that we assume that the kernel coordinates are centered

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

output
 k, l
filter
image (signal)


Let's run the box filter



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

output
 k, l filter
image (signal)

Let's run the box filter




shift-invariant:
as the pixel
shifts, so does
the kernel

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

output
 k, l
filter
image (signal)


Let's run the box filter



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

output
 k, l filter
image (signal)


Let's run the box filter



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

output
 k, l filter
image (signal)

Let's run the box filter



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

output
 k, l filter
image (signal)

Let's run the box filter

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

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)

Let's run the box filter

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

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)


Let's run the box filter



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

output
 k, l filter
image (signal)

Let's run the box filter



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

output
 k, l filter
image (signal)

Let's run the box filter

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

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	30	30	20		

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

output
 k, l
filter
image (signal)

Let's run the box filter

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

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	30	30	20	10	

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

output
 k, l
filter
image (signal)

Let's run the box filter

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

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	30	30	20	10	
	0								

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

output
 k, l
filter
image (signal)

Let's run the box filter



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

output
 k, l filter
image (signal)

Let's run the box filter

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

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	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	30	30	20	10	
	0	20	40						

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

output
 k, l filter
image (signal)

Let's run the box filter

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

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	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0								

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

output
 k, l
filter
image (signal)

Let's run the box filter

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

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	30	30	20	10	
	0	20	40	60	60	60	40	20	
0	30								

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

output
 k, l
filter
image (signal)

Let's run the box filter

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

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

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

output
 k, l
filter
image (signal)

Let's run the box filter

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

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

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

output
 k, l
filter
image (signal)

... and the result is

$g[\cdot, \cdot]$
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	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	0	10	20	30	30	30	20	10	
	10	10	10	10	0	0	0	0	
	10	10	10	10	0	0	0	0	

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


output
 k, l filter
image (signal)

Correlation (linear relationship)

$$f \otimes h = \sum_k \sum_l f(k,l)h(k,l)$$

f = Image

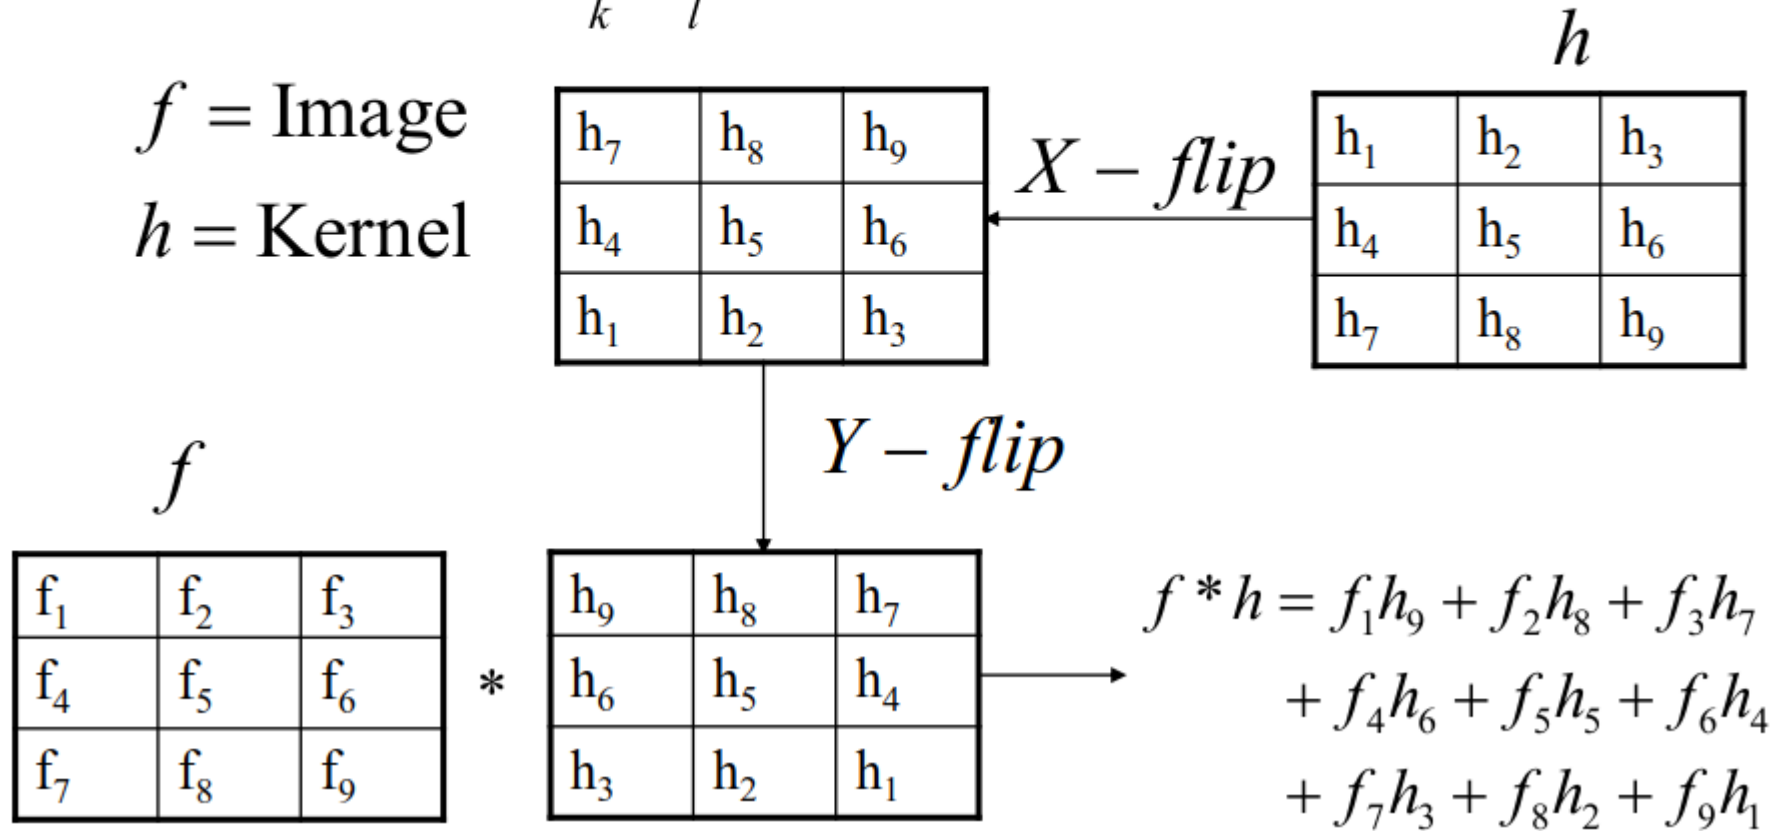
h = Kernel



Convolution

$$f * h = \sum_k \sum_l f(k, l) h(-k, -l)$$

$f = \text{Image}$
 $h = \text{Kernel}$





Correlation and Convolution

- **Convolution** is a filtering operation
 - expresses **the amount of overlap** of one function as it is shifted over another function
- **Correlation** compares the similarity of two sets of data
 - relatedness of the signals!



Key properties of linear filters

Linearity:

$$\text{filter}(f_1 + f_2) = \text{filter}(f_1) + \text{filter}(f_2)$$

Shift invariance: same behavior regardless of pixel location

$$\text{filter}(\text{shift}(f)) = \text{shift}(\text{filter}(f))$$

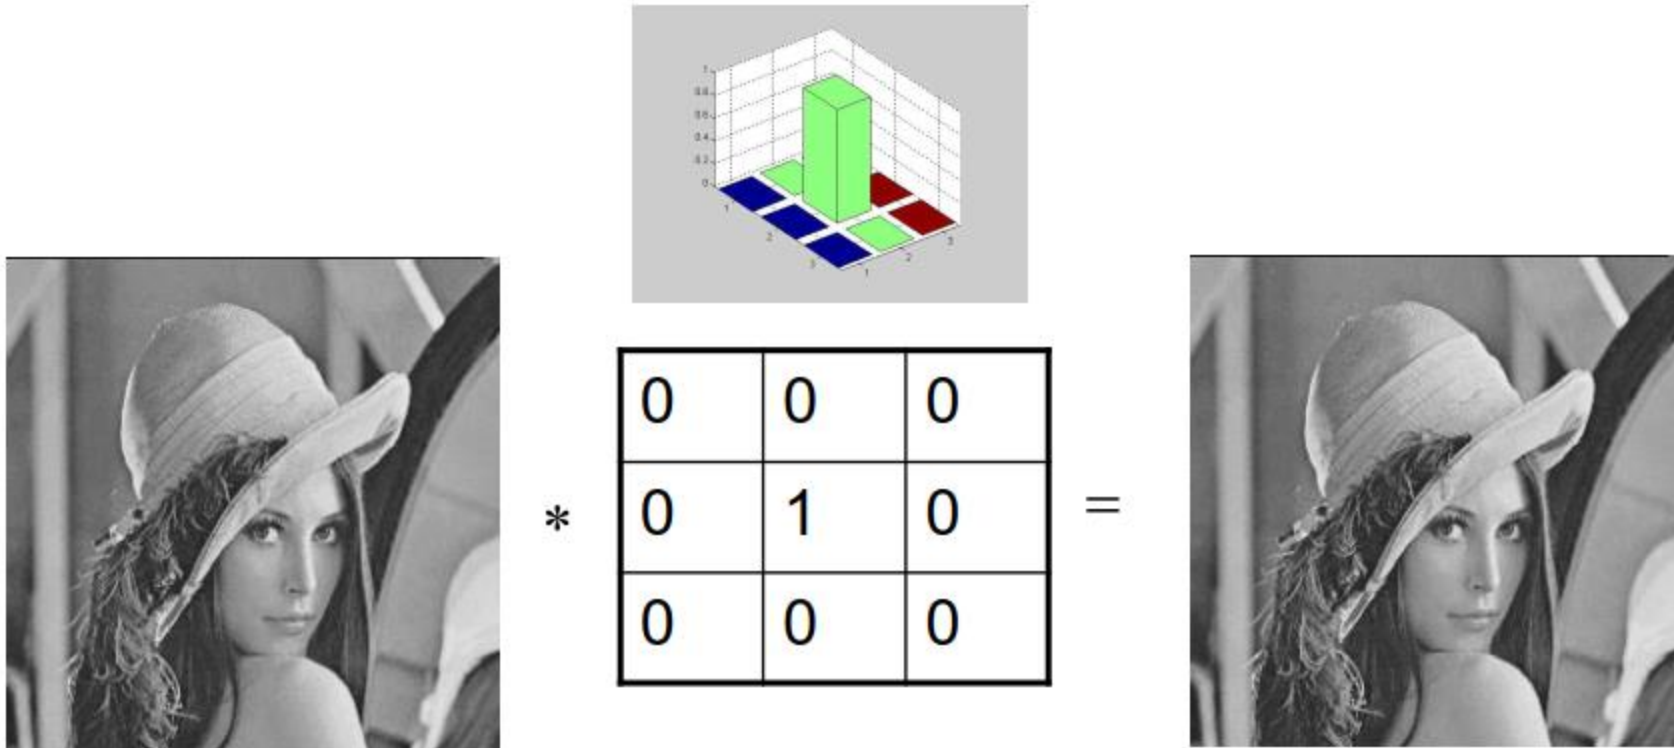
Any linear, shift-invariant operator can be represented as a convolution




More properties

- Commutative: $a * b = b * a$
 - Conceptually no difference between filter and signal
 - But particular filtering implementations might break this equality
- Associative: $a * (b * c) = (a * b) * c$
 - Often apply several filters one after another: $((a * b_1) * b_2) * b_3$
 - This is equivalent to applying one filter: $a * (b_1 * b_2 * b_3)$
- Distributes over addition: $a * (b + c) = (a * b) + (a * c)$
- Scalars factor out: $ka * b = a * kb = k(a * b)$
- Identity: unit impulse $e = [0, 0, 1, 0, 0]$,
 $a * e = a$

Filtering Examples - 1



Filtering Examples - 2




*

0	0	0
1	0	0
0	0	0

=

Filtering Examples - 2



*

0	0	0
1	0	0
0	0	0

=



Example: box filter

What does it do?


- Replaces each pixel with an average of its neighborhood

Average: mean

- Dividing the sum of N values by N

$$\frac{1}{9} \begin{matrix} & & g[\cdot, \cdot] \\ \begin{matrix} 1 \\ | \\ 9 \end{matrix} & \begin{matrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{matrix} \end{matrix}$$

Filtering Examples - 3




$\ast \frac{1}{9}$

1	1	1
1	1	1
1	1	1

=

Filtering Examples - 3



$\ast \frac{1}{9}$

1	1	1
1	1	1
1	1	1

=



Example: box filter

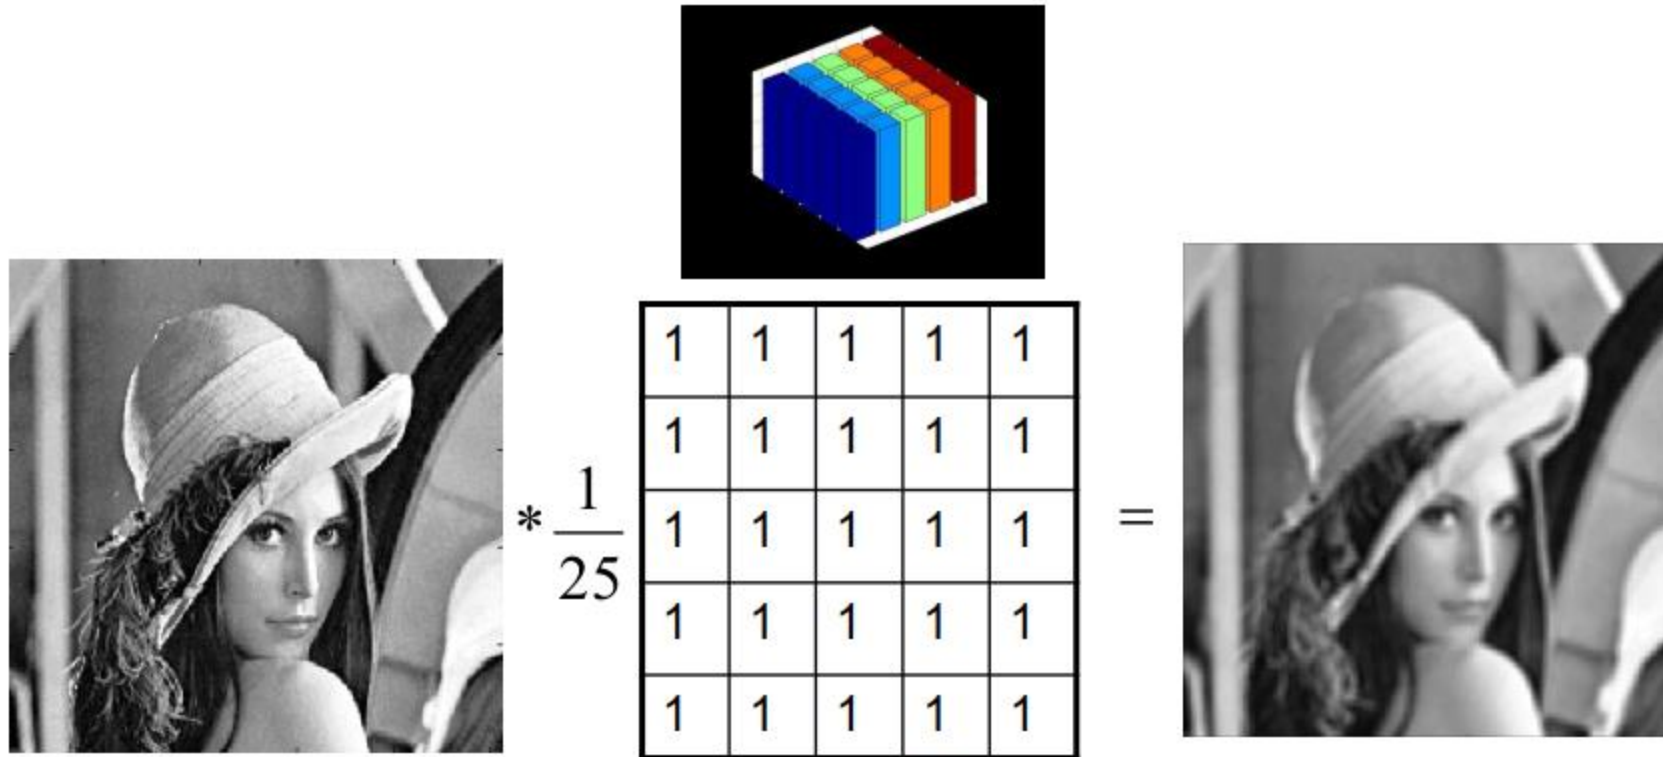
What does it do?

- Replaces each pixel with an average of its neighborhood
- Achieve smoothing effect (remove sharp features)

$$\frac{1}{9} g[\cdot, \cdot]$$

1	1	1
1	1	1
1	1	1

Filtering Examples - 4





Separable filters

A 2D filter is separable if it can be written as the product of a “column” and a “row”.

example:
box filter

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$$

column row

What is the rank of this filter matrix?



Separable filters

A 2D filter is separable if it can be written as the product of a “column” and a “row”.

example:
box filter

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$$

column row

Why is this important?

Separable filters

A 2D filter is separable if it can be written as the product of a “column” and a “row”.

example:
box filter

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$$

column row

2D convolution with a separable filter is equivalent to two 1D convolutions (with the “column” and “row” filters).

Separable filters

A 2D filter is separable if it can be written as the product of a “column” and a “row”.

example:
box filter

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$$

column row

2D convolution with a separable filter is equivalent to two 1D convolutions (with the “column” and “row” filters).

If the image has $M \times M$ pixels and the filter kernel has size $N \times N$:

- What is the cost of convolution with a non-separable filter?

Separable filters

A 2D filter is separable if it can be written as the product of a “column” and a “row”.

example:
box filter

1	1	1
1	1	1
1	1	1

=

1
1
1

column

*

1	1	1
---	---	---

row

2D convolution with a separable filter is equivalent to two 1D convolutions (with the “column” and “row” filters).

If the image has $M \times M$ pixels and the filter kernel has size $N \times N$:

- What is the cost of convolution with a non-separable filter? $\longrightarrow M^2 \times N^2$
- What is the cost of convolution with a separable filter?

Separable filters

A 2D filter is separable if it can be written as the product of a “column” and a “row”.

example:
box filter

$$\begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} = \begin{array}{|c|} \hline 1 \\ \hline 1 \\ \hline 1 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline \end{array} \text{row}$$

column

2D convolution with a separable filter is equivalent to two 1D convolutions (with the “column” and “row” filters).

If the image has $M \times M$ pixels and the filter kernel has size $N \times N$:

- What is the cost of convolution with a non-separable filter? $\longrightarrow M^2 \times N^2$
- What is the cost of convolution with a separable filter? $\longrightarrow 2 \times N \times M^2$


The Gaussian filter

- named (like many other things) after Carl Friedrich Gauss
- kernel values sampled from the 2D Gaussian function:

$$f(i, j) = \frac{1}{2\pi\sigma^2} e^{-\frac{i^2+j^2}{2\sigma^2}}$$

- weight falls off with distance from center pixel
- theoretically infinite, in practice truncated to some maximum distance

Any heuristics for selecting where to truncate?



The Gaussian filter


- named (like many other things) after Carl Friedrich Gauss
- kernel values sampled from the 2D Gaussian function:

$$f(i, j) = \frac{1}{2\pi\sigma^2} e^{-\frac{i^2+j^2}{2\sigma^2}}$$

- weight falls off with distance from center pixel
- theoretically infinite, in practice truncated to some maximum distance

Any heuristics for selecting where to truncate?

- usually at $2-3\sigma$



Is this a separable filter?

kernel $\frac{1}{16}$

1	2	1
2	4	2
1	2	1

The Gaussian filter


- named (like many other things) after Carl Friedrich Gauss
- kernel values sampled from the 2D Gaussian function:

$$f(i, j) = \frac{1}{2\pi\sigma^2} e^{-\frac{i^2+j^2}{2\sigma^2}}$$

- weight falls off with distance from center pixel
- theoretically infinite, in practice truncated to some maximum distance

Any heuristics for selecting where to truncate?

- usually at $2-3\sigma$




Is this a separable filter? Yes!


kernel $\frac{1}{16}$

1	2	1
2	4	2
1	2	1

The Gaussian Filter



$$g(x) = e^{\frac{-x^2}{2\sigma^2}}$$




$$g(x, y) = e^{\frac{-(x^2 + y^2)}{2\sigma^2}}$$


$$g(x) = [.011 \quad .13 \quad .6 \quad 1 \quad .6 \quad .13 \quad .011]$$

$$\sigma = 1$$


Gaussian filters



$\sigma = 1$ pixel



$\sigma = 5$ pixels



$\sigma = 10$ pixels




$\sigma = 30$ pixels

Filtering Examples - 5




*



=

Filtering Examples - 5



Gaussian Smoothing

Filtering Examples - 6



Gaussian Smoothing



Smoothing by Averaging

Filtering Examples - 7



After additive
Gaussian Noise



After Averaging



After Gaussian Smoothing

Filtering Examples – 8

Sharpening

input



filter

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} - \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

output


(Note that filter sums to 1)

- do nothing for flat areas
- stress intensity peaks

Filtering Examples – 8

Sharpening

Accentuates differences with local average



(Note that filter sums to 1)


- do nothing for flat areas
- stress intensity peaks

Sharpening


- What does blurring take away?



−



=




(This “detail extraction” operation is also called a *high-pass filter*)

Let's add it back:



+




=




Sharpening

- What does blurring take away?




-



(This “detail extraction” operation is also called a *high-pass filter*)



+



=



Sharpening



- What does blurring take away?

2 times original


0	0	0
0	2	0
0	0	0

-


Smoothed

$$\frac{1}{9} \begin{array}{|c|c|c|} \tr & 1 & 1 & 1 \\ \tr & 1 & 1 & 1 \\ \tr & 1 & 1 & 1 \end{array}$$



(This “detail extraction” operation is also called a *high-pass filter*)



+



-



=





Sharpening

- What does blurring take away?



(This “detail extraction” operation is also called a *high-pass filter*)

Sharpening examples





Median Filter


- A **Median Filter** operates over a window by selecting the median intensity in the window.



Median Filter

- A **Median Filter** operates over a window by selecting the median intensity in the window.
- Great to deal with salt and pepper noise !

Median Filter



Practical matters

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



Source: S. Marschner



Questions?