


Lecture 1 Recap

Phong Illumination Model


- Includes ambient, diffuse and specular reflection

$$I = k_a i_a + k_d i_d \cos \theta + k_s i_s \cos^\alpha \phi$$



Diffuse and Specular Reflection

- A sphere lit with ambient, +diffuse, +specular reflectance




Ambient

+Diffuse

+Specular

Pinhole Camera


f' is the **focal length** of the camera



Note: In a pinhole camera we can adjust the focal length, all this will do is change the **size** of the resulting image

Perspective Projection: Matrix Form

Camera Matrix



$$\mathbf{C} = \begin{bmatrix} f' & 0 & 0 & 0 \\ 0 & f' & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

3D object point

Forsyth & Ponce (1st ed.) Figure 1.4

$$P = \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

projects to 2D image point

$$P' = \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix}$$

where

$$sP' = \mathbf{C}P$$


(s is a scale factor)



2.4


Why **Not** a Pinhole Camera?

- If pinhole is **too big** then many directions are averaged, blurring the image
- If pinhole is **too small** then diffraction becomes a factor, also blurring the image
- Generally, pinhole cameras are **dark**, because only a very small set of rays from a particular scene point hits the image plane
- Pinhole cameras are **slow**, because only a very small amount of light from a particular scene point hits the image plane per unit time



Reason for **Lenses**


A real camera must have a finite aperture to get enough light, but this causes blur in the image




Solution: use a **lens** to focus light onto the image plane

Reason for **Lenses**

A real camera must have a finite aperture to get enough light, but this causes blur in the image




The role of a lens is to **capture more light** while preserving, as much as possible, the abstraction of an ideal pinhole camera.




Solution: use a **lens** to focus light onto the image plane

Snell's Law



$$n_1 \sin \alpha_1 = n_2 \sin \alpha_2$$


Snell's Law



$$n_1 \sin \alpha_1 = n_2 \sin \alpha_2$$

Lens Basics

- A lens focuses rays from infinity at the focal length of the lens
- Points passing through the centre of the lens are not bent




2.6

- We can use these 2 properties to find the **thin** lens equation

Lens Basics

- A lens focuses rays from infinity at the focal length of the lens
- Points passing through the centre of the lens are not bent




2.6

- We can use these 2 properties to find the **thin** lens equation

Lens Basics

- A lens focuses rays from infinity at the focal length of the lens
- Points passing through the centre of the lens are not bent



2.6


- We can use these 2 properties to find the **thin** lens equation

Lens Basics

- A 50mm lens is focussed at infinity. It now moves to focus on something 5m away. How far does the lens move?




Pinhole Model **with Lens**




Lens Basics

- Lenses focus all rays from a plane in the world




- Objects off the plane are blurred depending on distance

Effect of Aperture Size



Smaller aperture \Rightarrow smaller blur, larger **depth of field**



Depth of Field

- Photographers use large apertures to give small depth of field




Aperture size = f/N , \Rightarrow large N = small aperture

Real Lenses



- Real Lenses have multiple stages of positive and negative elements with differing refractive indices
- This can help deal with issues such as chromatic aberration (different colours bent by different amounts), vignetting (light fall off at image edge) and sharp imaging across the zoom range

Spherical Aberration




Forsyth & Ponce (1st ed.) Figure 1.12a

Spherical **Aberration**

Un-aberrated image




Image from lens with Spherical Aberration



Vignetting

Vignetting in a two-lens system



Forsyth & Ponce (2nd ed.) Figure 1.12

The shaded part of the beam **never reaches** the second lens

Vignetting



Chromatic **Aberration**

- Index of **refraction depends on wavelength**, λ , of light
- Light of different colours follows different paths
- Therefore, not all colours can be in equal focus





Image Credit: Trevor Darrell

Lens Distortion

Fish-eye Lens



Szeliski (1st ed.) Figure 2.13

Lines in the world are no longer lines on the image, they are curves!


Other (Possibly Significant) **Lens Effects**

Scattering at the lens surface

- Some light is reflected at each lens surface

There are other **geometric phenomena/distortions**

- pincushion distortion
- barrel distortion



Parametric calibration errors

[Schöps et al., 2020]

Image from [Schöps et al., 2019]. Reproduced for educational purposes.

<https://www.flickr.com/photos/nragsdale/3192314056/>

Lecture **Summary**

- We discussed a “physics-based” approach to image formation. Basic abstraction is the **pinhole camera**.
- **Lenses overcome limitations** of the pinhole model while trying to preserve it as a useful abstraction
- Projection equations: **perspective**, weak perspective, orthographic
- Thin lens equation
- Some “aberrations and **distortions**” persist (e.g. spherical aberration, vignetting)

Course **logistics**

Times: Mon, Wed 3:30-5:00pm

Locations: Friedman (FRDM), Room 153

Instructor: Kwang



Fri. (ICCS 115)
1 — 2 pm

Fred



Mon. (Zoom)
5 — 6 pm


Teaching Assistants

Ramin



Tues. (Room TBA)
5 — 6 pm

Bicheng



Wed. (Zoom)
5 — 6 pm


Rayat



Thurs. (ICCS X239)
2:30 — 3:30 pm



CPSC 425: Computer Vision



Lecture 3: Image Filtering

(unless otherwise stated slides are taken or adopted from **Bob Woodham, Jim Little** and **Fred Tung**)

This Lecture

Topics: Image Filtering

- **Image** as a **function**
- **Linear** filters
- **Correlation / Convolution**

Readings:

- **Today's** Lecture: Szeliski 3.1-3.3, Forsyth & Ponce (2nd ed.) 4.1, 4.5

Reminders:

- Complete **Assignment 1** is out! Due 29th

Goal

1. Learn how to mathematically describe image processing
2. Basic building blocks


Image as a **2D Function**

A (grayscale) image is a 2D function



grayscale image

$$I(X, Y)$$



What is the **range** of the image function?

$$I(X, Y) \in [0, 255] \in \mathbb{Z}$$

domain: $(X, Y) \in ([1, width], [1, height])$

Adding two Images

Since images are functions, we can perform operations on them, e.g., **average**



$I(X, Y)$



$G(X, Y)$




$$\frac{I(X, Y)}{2} + \frac{G(X, Y)}{2}$$

Adding two Images



$$a = \frac{I(X, Y)}{2} + \frac{G(X, Y)}{2}$$



$$b = \frac{I(X, Y) + G(X, Y)}{2}$$

Adding two Images




$$a = \frac{I(X, Y)}{2} + \frac{G(X, Y)}{2}$$

Question:

$$a = b$$

$$a > b$$

$$a < b$$



$$b = \frac{I(X, Y) + G(X, Y)}{2}$$


Adding two Images



Red pixel in camera man image = 98

Red pixel in moon image = 200

$$\frac{98}{2} + \frac{200}{2} = 49 + 100 = 149$$



$$\frac{98 + 200}{2} = \frac{\lfloor 298 \rfloor}{2} = \frac{255}{2} = 127$$


Question:

$$a = b$$

$$a > b$$

$$a < b$$

Adding two Images



It is often convenient to convert images to **doubles** when doing processing

In Python

```
from PIL import Image
img = Image.open('cameraman.png') ←
import numpy as np
imgArr = np.asfarray(img)

# Or do this or "imgArr=np.array(img).astype(np.float32)/255.0"
import matplotlib.pyplot as plt
camera = plt.imread('cameraman.png');
```

What types of **transformations** can we do?

$I(X, Y)$



Filtering



$I'(X, Y)$



changes range of image function


$I(X, Y)$



Warping




$I'(X, Y)$



changes domain of image function

What types of **filtering** can we do?

Point Operation



point processing

Examples of Point Processing

original



$$I(X, Y)$$

darken



$$I(X, Y) - 128$$

lower contrast



$$\frac{I(X, Y)}{2}$$

non-linear lower contrast



invert



lighten



raise contrast




non-linear raise contrast



Brightness v.s. Contrast

Brightness: all pixels get lighter/darker, relative difference between pixel values stays the same

Contrast: relative difference between pixel values becomes higher / lower



Examples of Point Processing

original



$$I(X, Y)$$

darken



$$I(X, Y) - 128$$

lower contrast



$$\frac{I(X, Y)}{2}$$

non-linear lower contrast



$$\left(\frac{I(X, Y)}{255}\right)^{1/3} \times 255$$

invert



$$255 - I(X, Y)$$

lighten



$$I(X, Y) + 128$$

raise contrast



$$I(X, Y) \times 2$$

non-linear raise contrast



$$\left(\frac{I(X, Y)}{255}\right)^2 \times 255$$

Examples of Point Processing

original



$$I(X, Y)$$

darken



$$I(X, Y) - 128$$

lower contrast



$$\frac{I(X, Y)}{2}$$

non-linear lower contrast




$$\left(\frac{I(X, Y)}{255}\right)^{1/3} \times 255$$

invert




$$255 - I(X, Y)$$

lighten



$$I(X, Y) + 128$$

raise contrast



$$I(X, Y) \times 2$$


non-linear raise contrast



$$\left(\frac{I(X, Y)}{255}\right)^2 \times 255$$


What types of **filtering** can we do?

Point Operation



point processing

Neighborhood Operation



“filtering”

Linear Neighborhood Operators (Filtering)



Original Image



blur



sharpen



edge filter

Non-Linear Neighborhood Operators (Filtering)



Original Image



edge preserving
smoothing



median




canny edges

Linear **Filters**

Let $I(X, Y)$ be an $n \times n$ digital image (for convenience we let width = height)

Let $F(X, Y)$ be another $m \times m$ digital image (our “**filter**” or “**kernel**”)




For convenience we will assume m is odd. (Here, $m = 5$)

Linear **Filters**


For a give X and Y , superimpose the filter on the image centered at (X, Y)

Compute the new pixel value, $I'(X, Y)$, as the sum of $m \times m$ values, where each value is the product of the original pixel value in $I(X, Y)$ and the corresponding values in the filter




Linear **Filters**

The computation is repeated for each
 (X, Y)




Linear Filter Example




$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

output
filter
image (signal)




Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


output
filter
image (signal)

Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


output
filter
image (signal)

Linear Filter Example




$$\underbrace{I'(X, Y)}_{\text{output}} = \sum_{j=-k}^k \sum_{i=-k}^k \underbrace{F(i, j)}_{\text{filter}} \underbrace{I(X + i, Y + j)}_{\text{image (signal)}}$$

Linear Filter Example



$$\underbrace{I'(X, Y)}_{\text{output}} = \sum_{j=-k}^k \sum_{i=-k}^k \underbrace{F(i, j)}_{\text{filter}} \underbrace{I(X + i, Y + j)}_{\text{image (signal)}}$$


Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


output
filter
image (signal)

Linear Filter Example




$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

output
filter
image (signal)


Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


output
filter
image (signal)

Linear Filter Example




$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$


Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

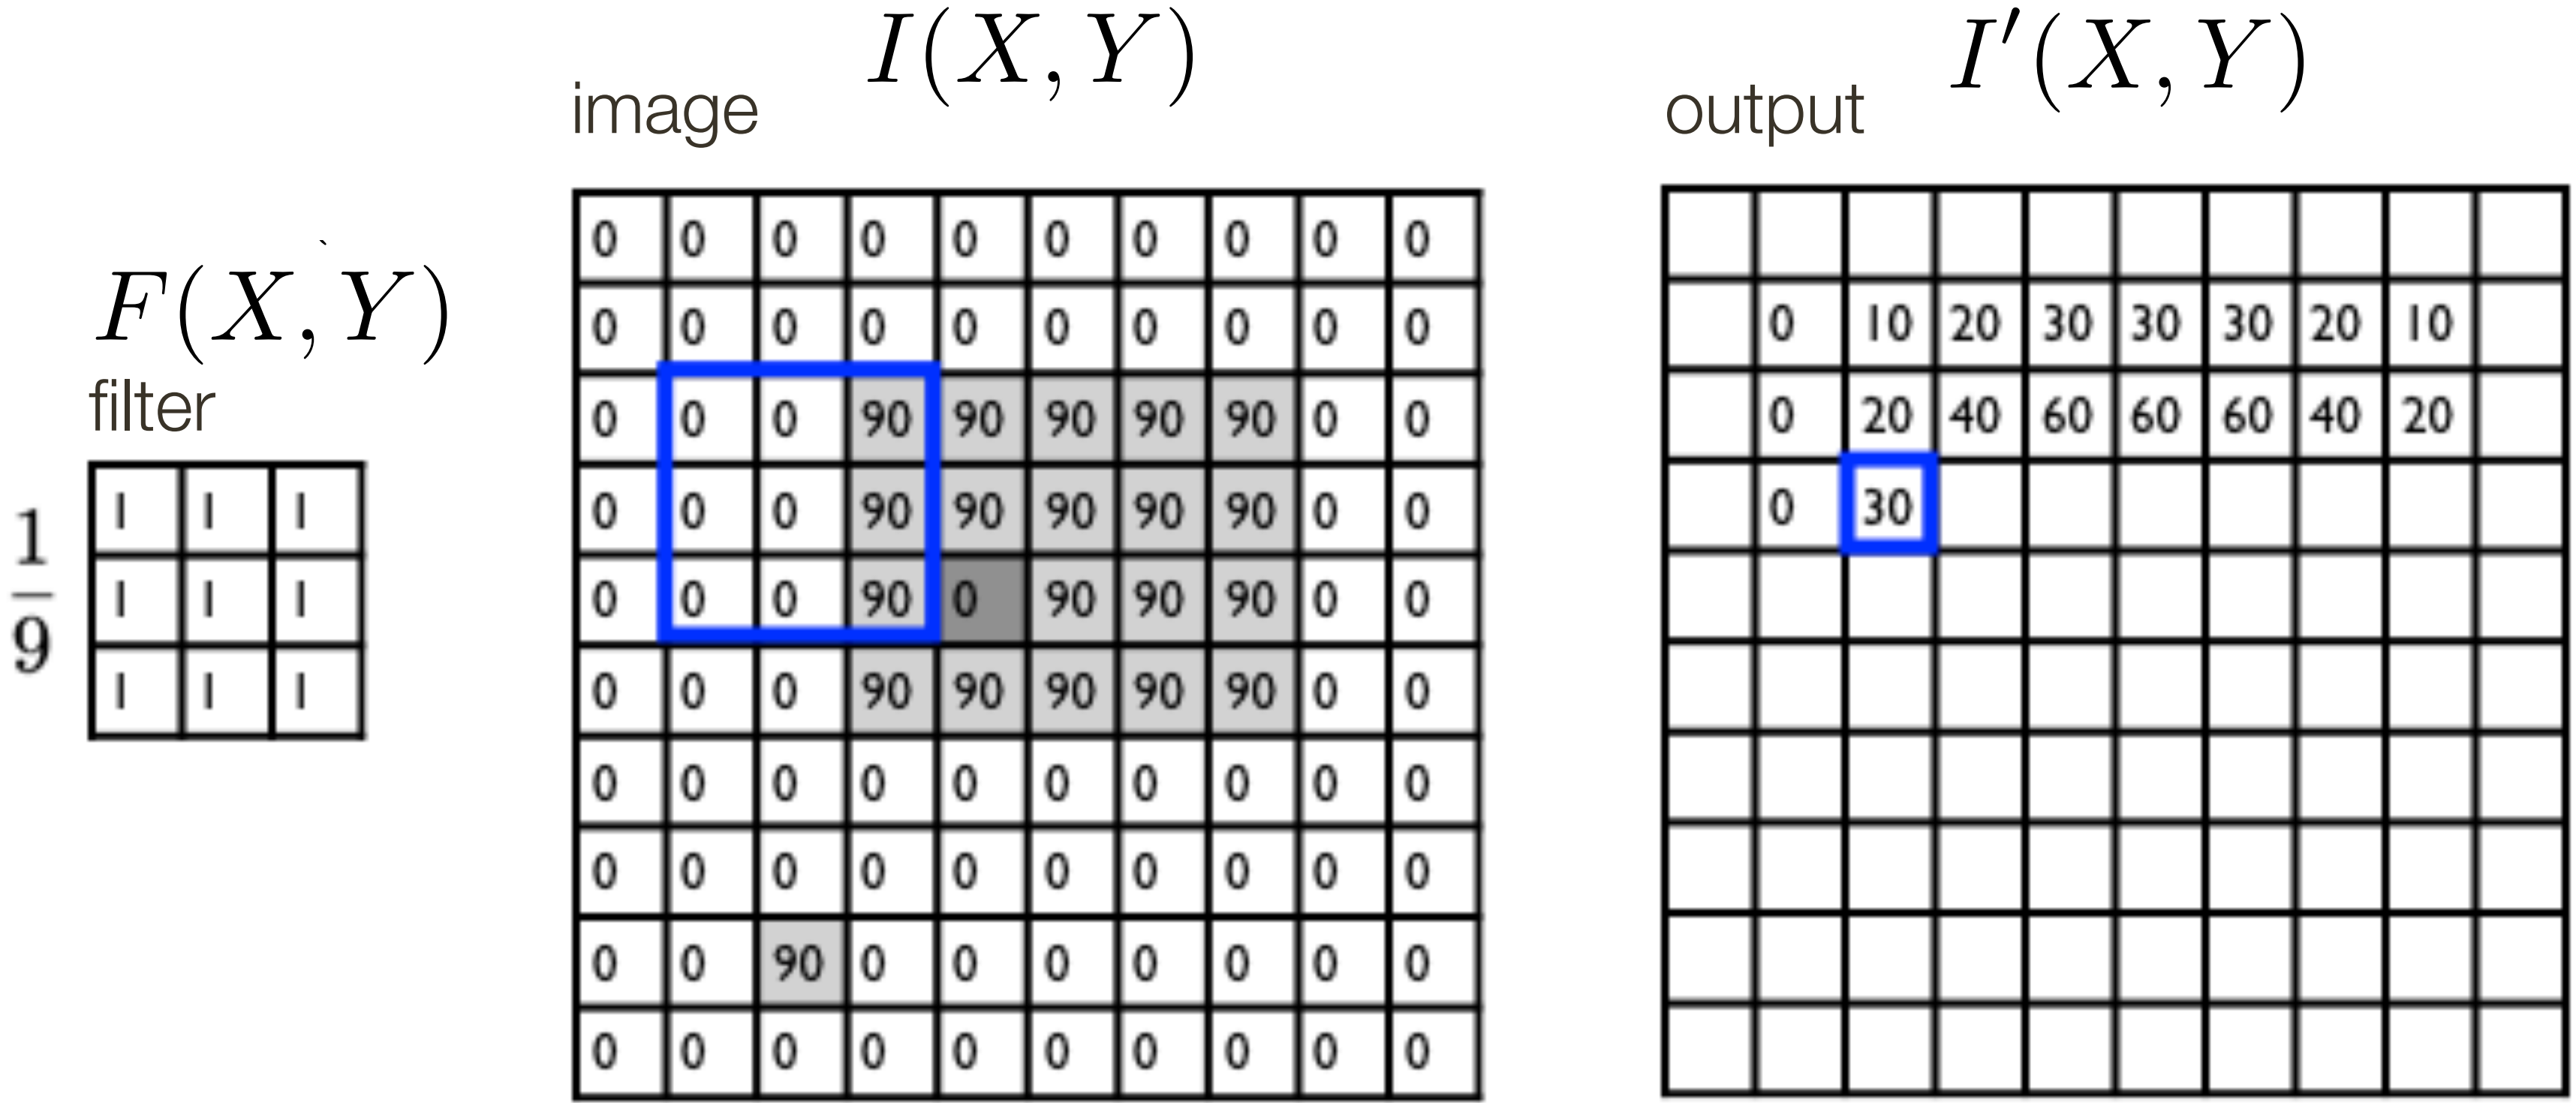
output
filter
image (signal)

Linear Filter Example



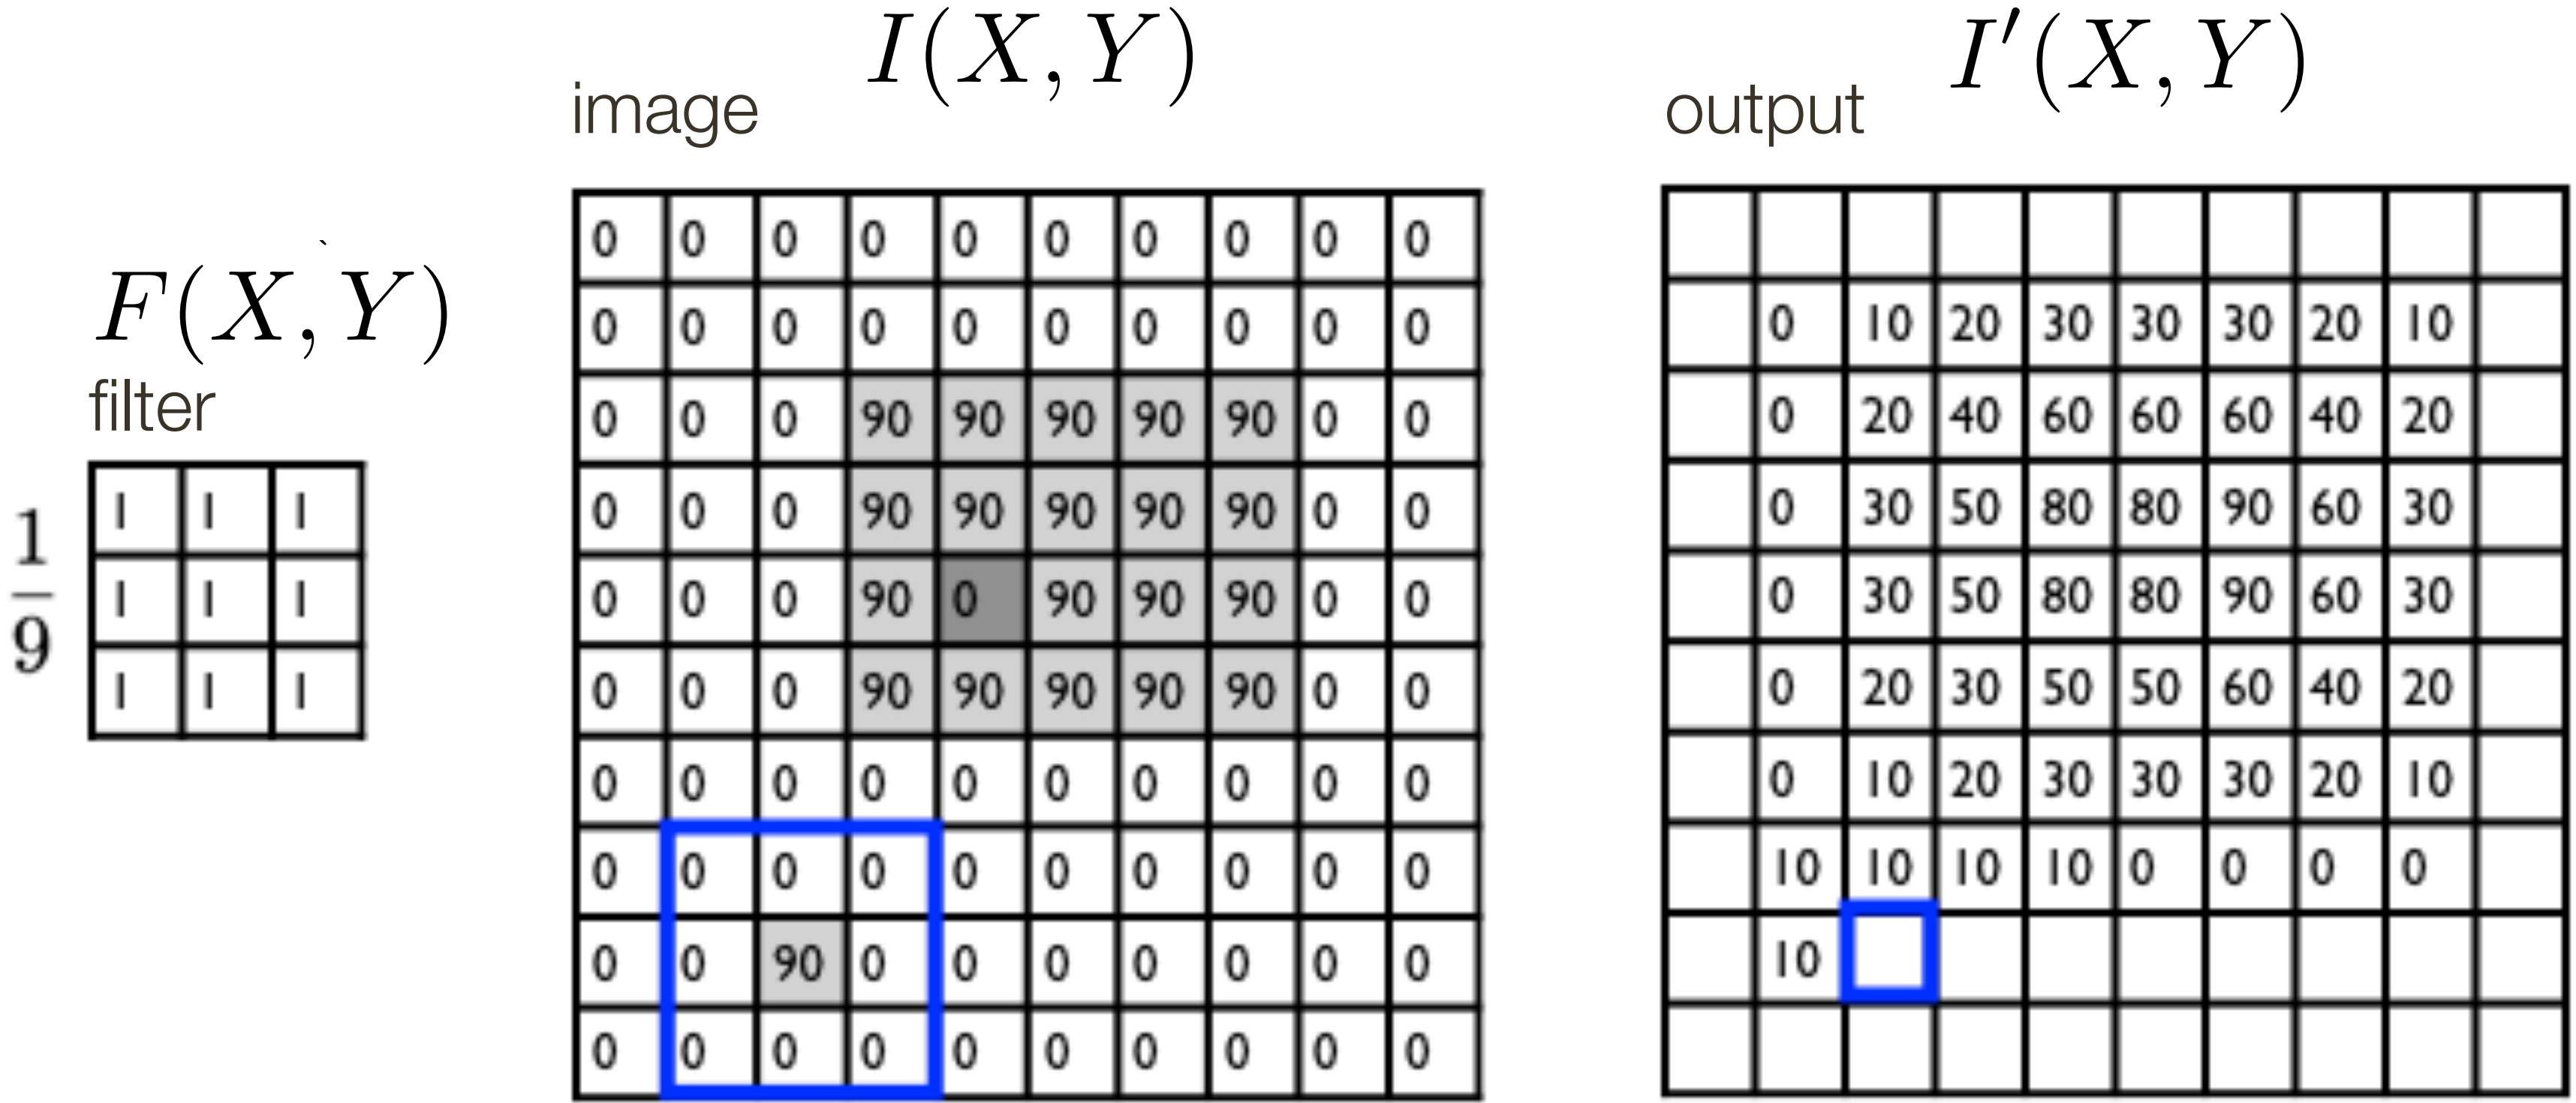
$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

Linear Filter Example




$$\underbrace{I'(X, Y)}_{\text{output}} = \sum_{j=-k}^k \sum_{i=-k}^k \underbrace{F(i, j)}_{\text{filter}} \underbrace{I(X + i, Y + j)}_{\text{image (signal)}}$$

Linear Filter Example



$$\underbrace{I'(X, Y)}_{\text{output}} = \sum_{j=-k}^k \sum_{i=-k}^k \underbrace{F(i, j)}_{\text{filter}} \underbrace{I(X + i, Y + j)}_{\text{image (signal)}}$$


Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

output
filter
image (signal)

Linear Filter Example



$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

output
 filter
 image (signal)

Linear **Filters**

$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X + i, Y + j)$$

output filter image (signal)

For a given X and Y , superimpose the filter on the image centered at (X, Y)

Compute the new pixel value, $I'(X, Y)$, as the sum of $m \times m$ values, where each value is the product of the original pixel value in $I(X, Y)$ and the corresponding values in the filter

Linear Filters

Let's do some accounting ...

$$\begin{array}{c} I'(X, Y) \\ \text{output} \end{array} = \sum_{j=-k}^k \sum_{i=-k}^k \begin{array}{c} F(i, j) \\ \text{filter} \end{array} \begin{array}{c} I(X + i, Y + j) \\ \text{image (signal)} \end{array}$$


At each pixel, (X, Y) , there are $m \times m$ multiplications

There are $n \times n$ pixels in (X, Y)


Total: $m^2 \times n^2$ multiplications

When m is fixed, small constant, this is $\mathcal{O}(n^2)$. But when $m \approx n$ this is $\mathcal{O}(m^4)$.

Linear Filters: **Boundary** Effects




Linear Filters: **Boundary** Effects



Three standard ways to deal with boundaries:

1. **Ignore these locations:** Make the computation undefined for the top and bottom k rows and the leftmost and rightmost k columns
2. **Pad the image with zeros:** Return zero whenever a value of I is required at some position outside the defined limits of X and Y


Linear Filters: **Boundary** Effects



Linear Filters: **Boundary** Effects




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



Notice **decrease** in brightness at edges


Linear Filters: **Boundary** Effects




Three standard ways to deal with boundaries:

1. **Ignore these locations:** Make the computation undefined for the top and bottom k rows and the leftmost and rightmost k columns
2. **Pad the image with zeros:** Return zero whenever a value of I is required at some position outside the defined limits of X and Y
3. **Assume periodicity:** The top row wraps around to the bottom row; the leftmost column wraps around to the rightmost column


Linear Filters: **Boundary** Effects



Linear Filters: **Boundary** Effects




Linear Filters: **Boundary** Effects




Four standard ways to deal with boundaries:

1. **Ignore these locations:** Make the computation undefined for the top and bottom k rows and the leftmost and rightmost k columns
2. **Pad the image with zeros:** Return zero whenever a value of I is required at some position outside the defined limits of X and Y
3. **Assume periodicity:** The top row wraps around to the bottom row; the leftmost column wraps around to the rightmost column
4. **Reflect boarder:** Copy rows/columns locally by reflecting over the edge

Linear Filters: **Boundary** Effects



Linear Filters: **Boundary** Effects



Four standard ways to deal with boundaries:

1. **Ignore these locations:** Make the computation undefined for the top and bottom k rows and the leftmost and rightmost k columns
2. **Pad the image with zeros:** Return zero whenever a value of I is required at some position outside the defined limits of X and Y
3. **Assume periodicity:** The top row wraps around to the bottom row; the leftmost column wraps around to the rightmost column
4. **Reflect boarder:** Copy rows/columns locally by reflecting over the edge