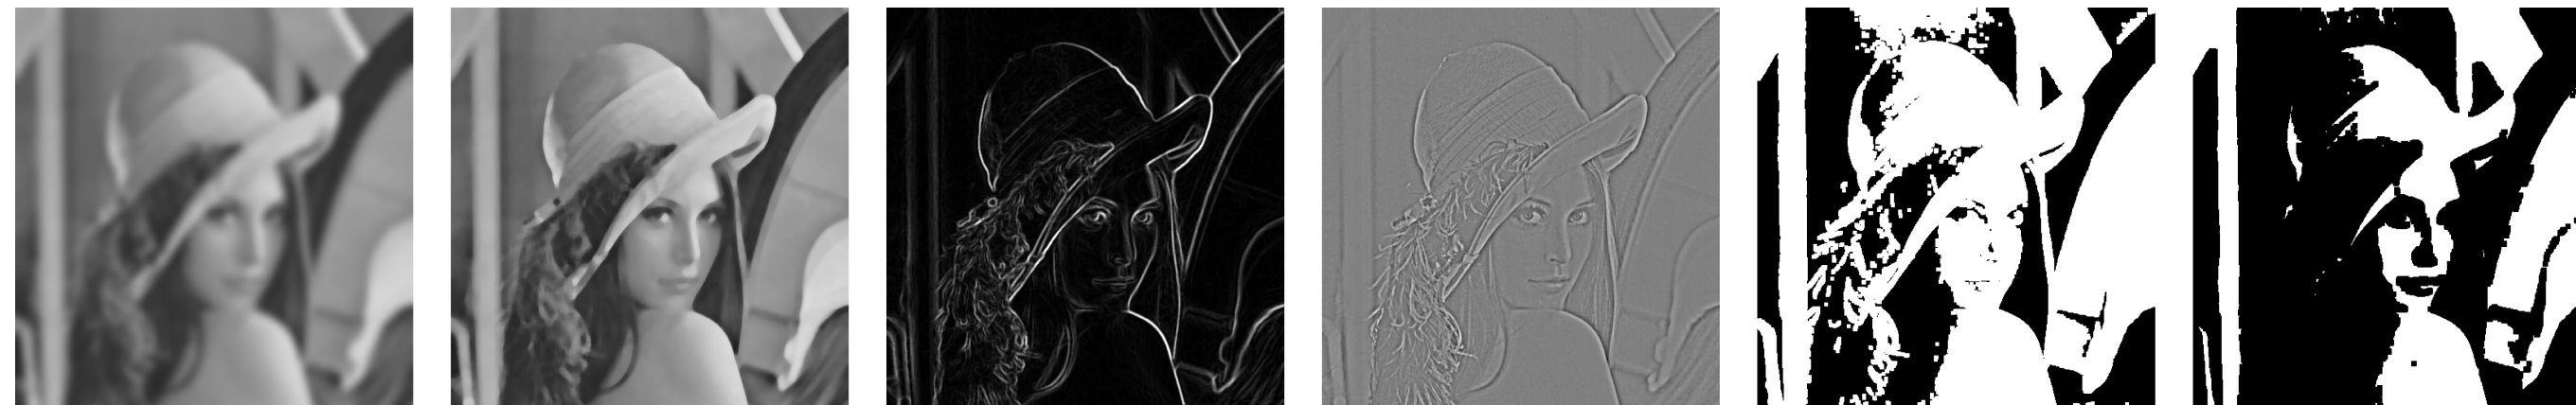




CPSC 425: Computer Vision



Lecture 4: Image Filtering (continued)

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

Menu for Today (September 13, 2018)

Topics:

- Linear filters
- Linear filter properties
- Correlation / Convolution
- Filter examples: Box, Gaussian, ...

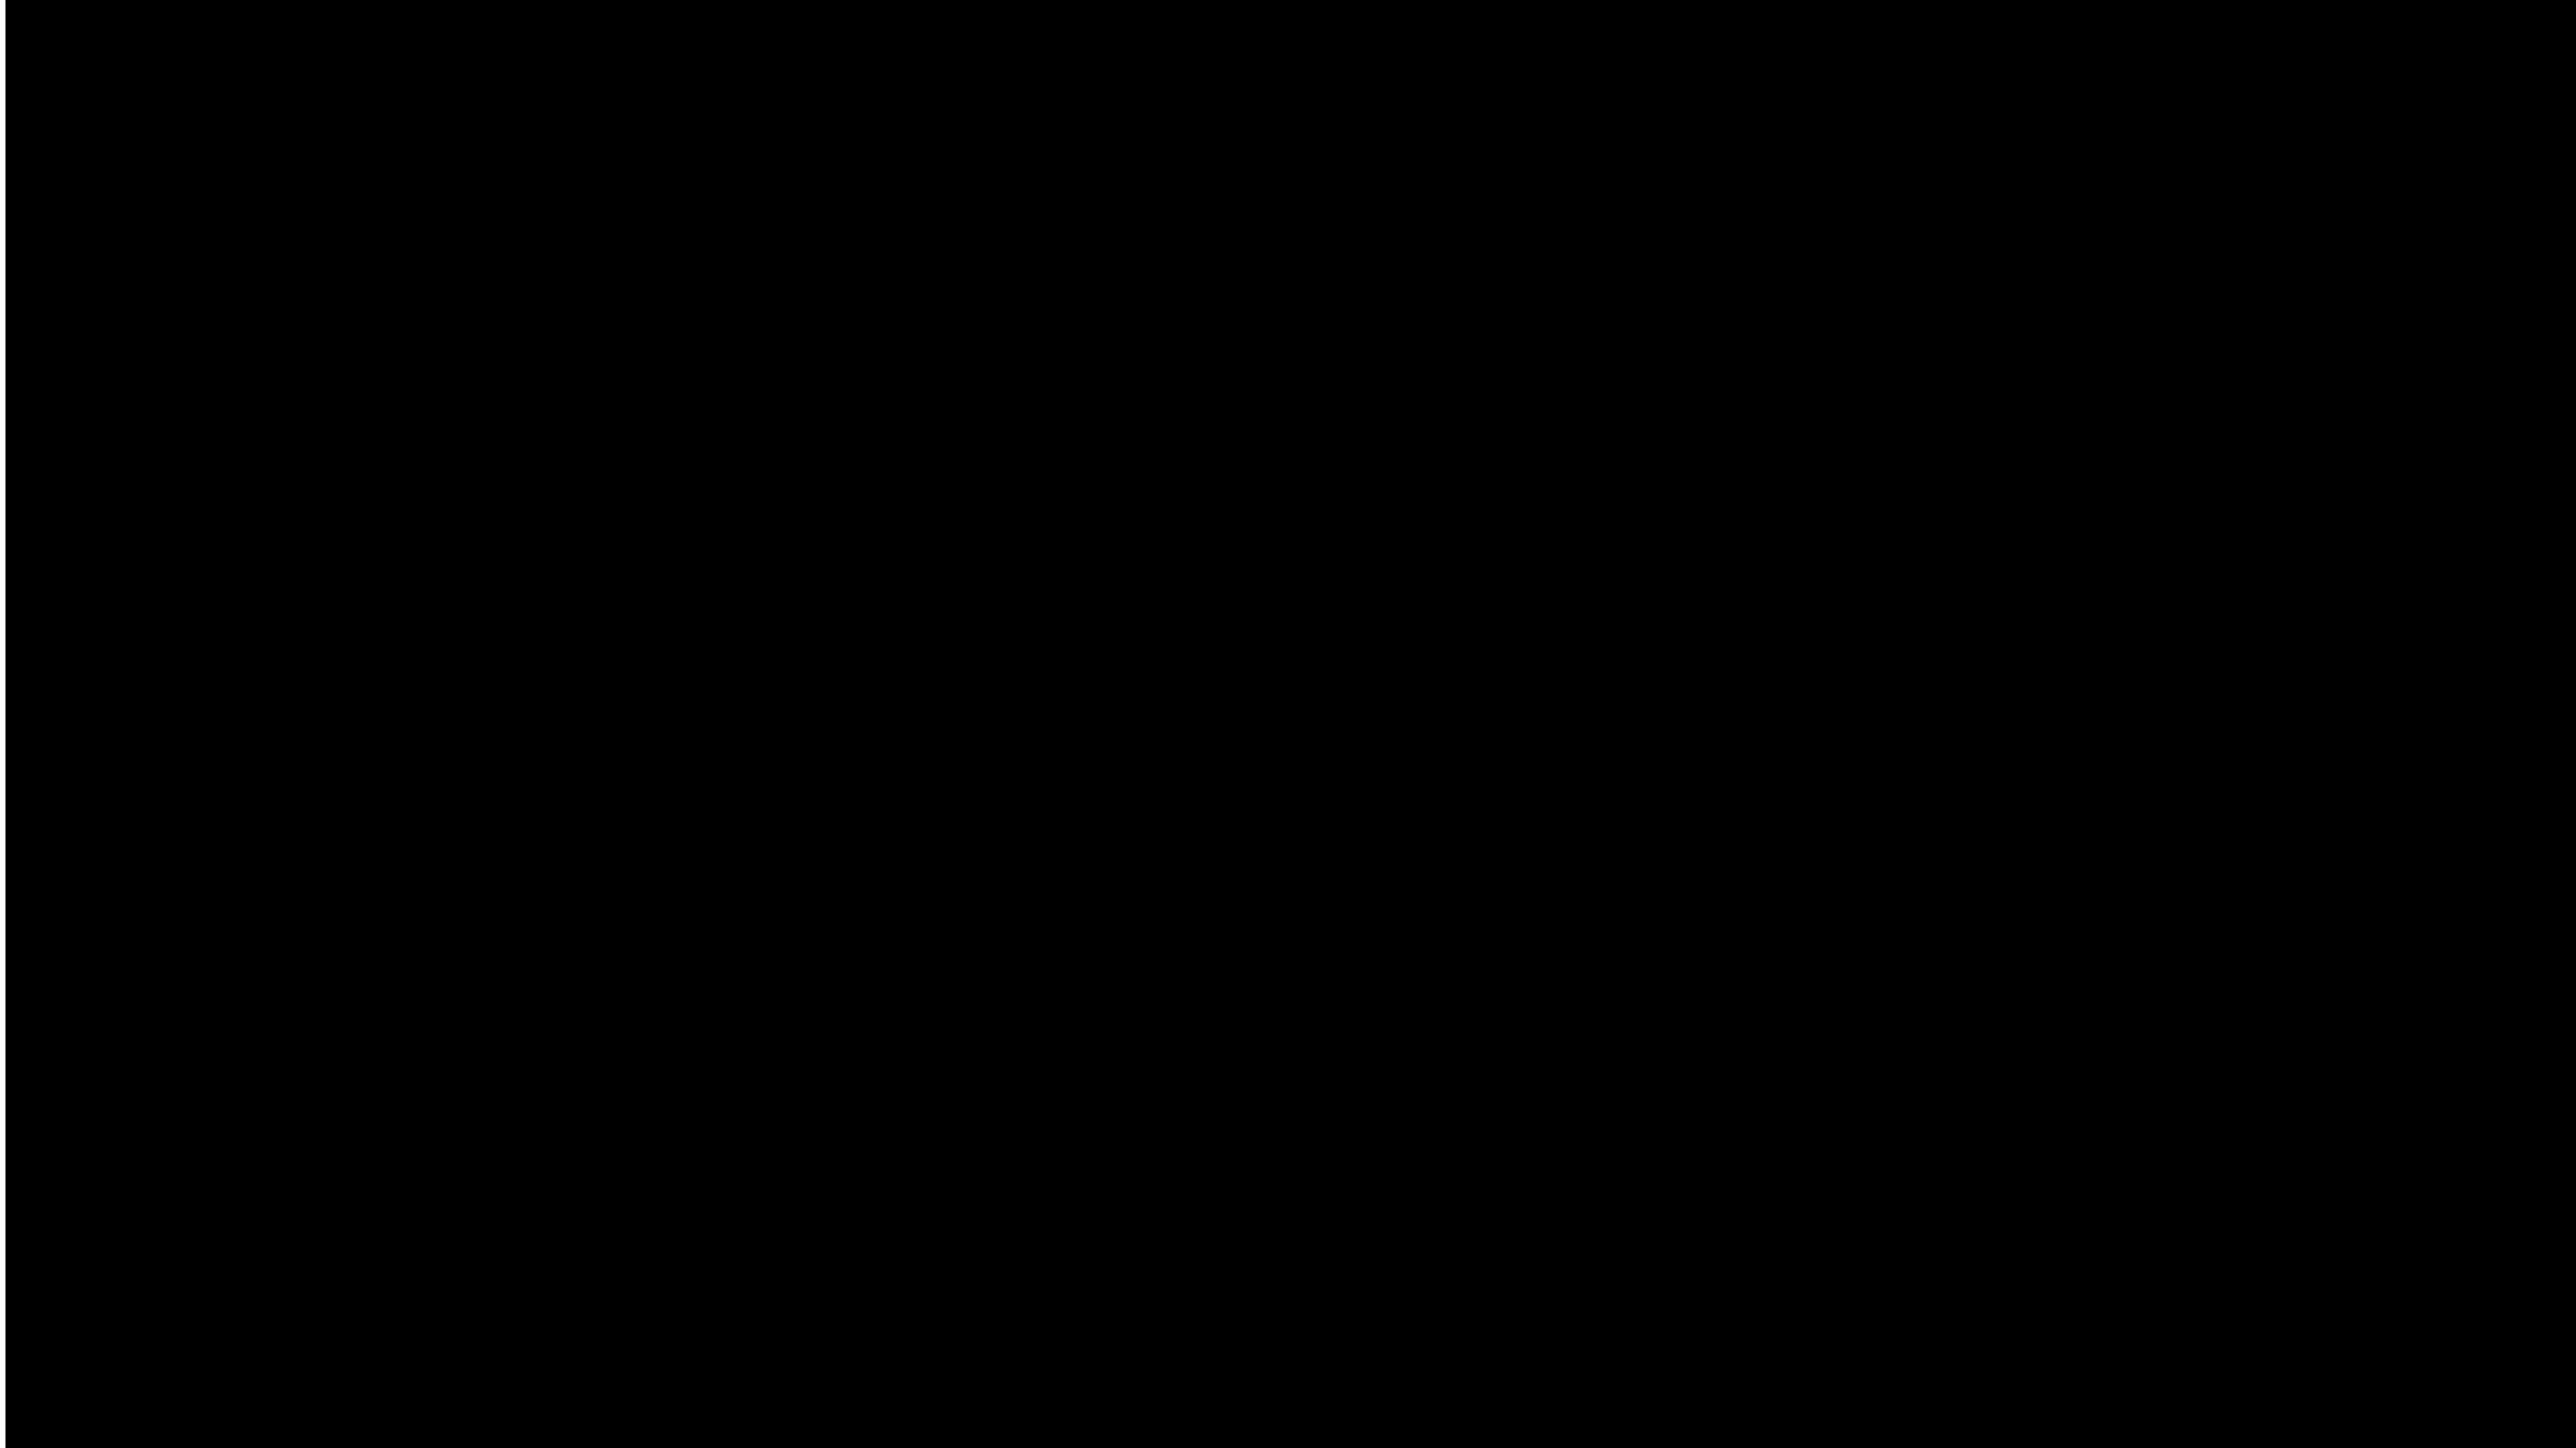
Readings:

- **Today's** Lecture: Forsyth & Ponce (2nd ed.) 4.1, 4.5
- **Next** Lecture: none

Reminders:

- **Assignment 1:** Image Filtering and Hybrid Images is out
- Conveniently, office hours of TAs who are responsible for this assignment are on **Thursday** (Siddhesh 11-noon) and **Friday** (Borna 9-10am)

Today's **“fun”** Example:



Lecture 3a: Re-cap Lenses

We take a “physics-based” approach to image formation

- Treat camera as an instrument that takes measurements of the 3D world

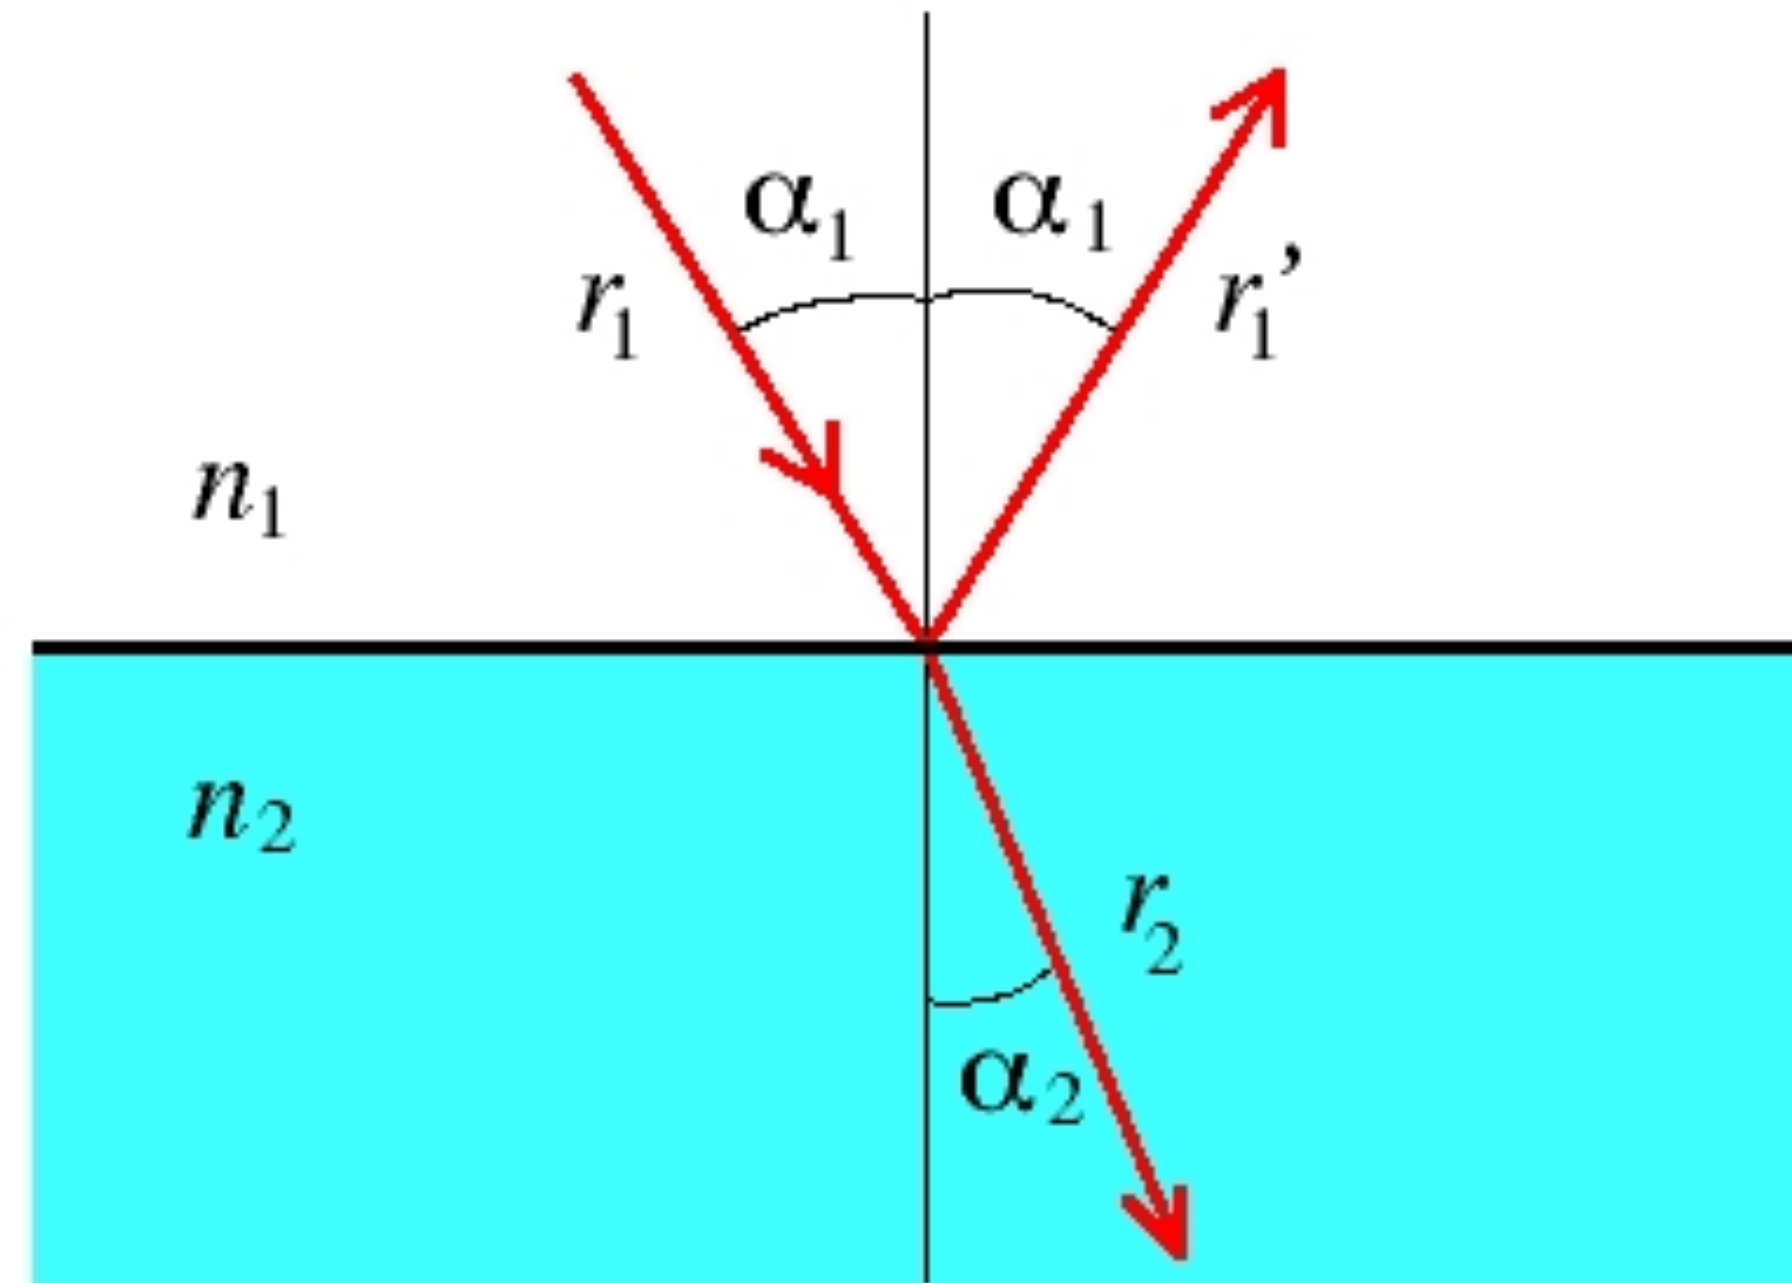
Basic abstraction is the **pinhole camera**

Lenses overcome limitations of the pinhole model while trying to preserve it as a useful abstraction

When **maximum accuracy** required, it is necessary to model additional details of each particular camera (and camera setting)

- Aside: This is called camera calibration

Lecture 3a: Re-cap Snell's Law



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

Lecture 3a: Re-cap Lenses

Thin **lens equation**

$$\frac{1}{z'} - \frac{1}{z} = \frac{1}{f}$$

characterizes the relationship between f , z and z'

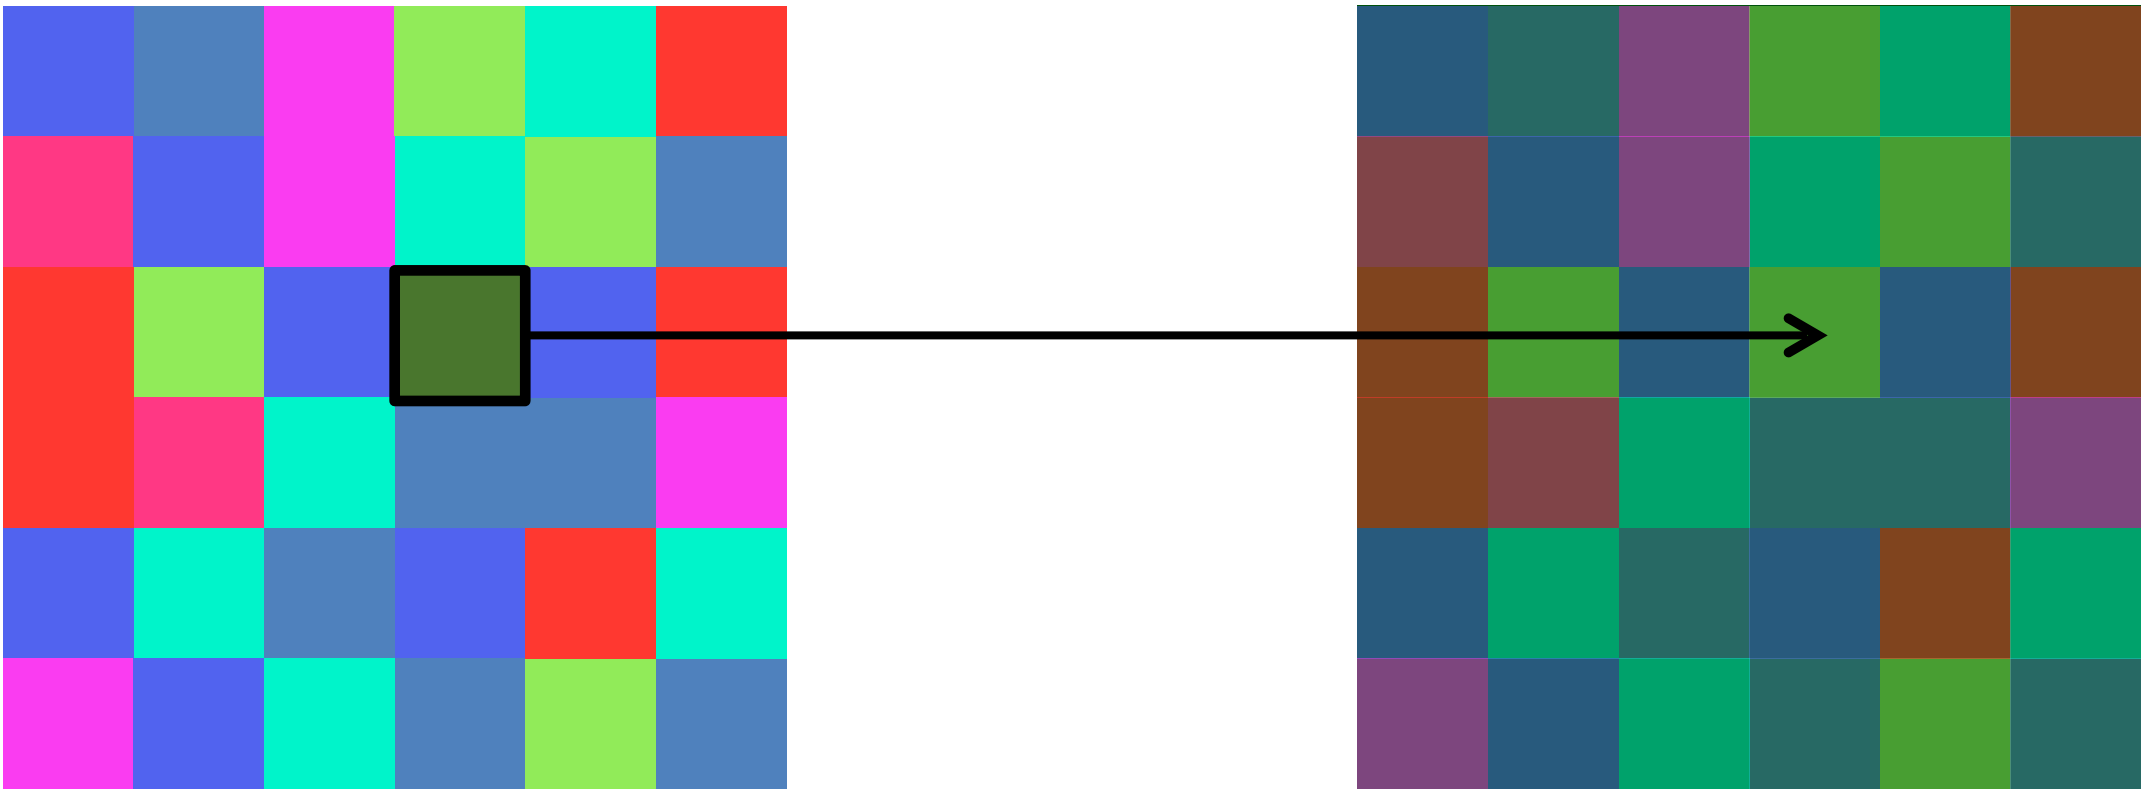
Some “**aberrations** and **distortions**” persist. For example:

- index of refraction depends on wavelength, λ , of light
- vignetting reduces image brightness (gradually) away from the image center

The human eye functions much like a camera

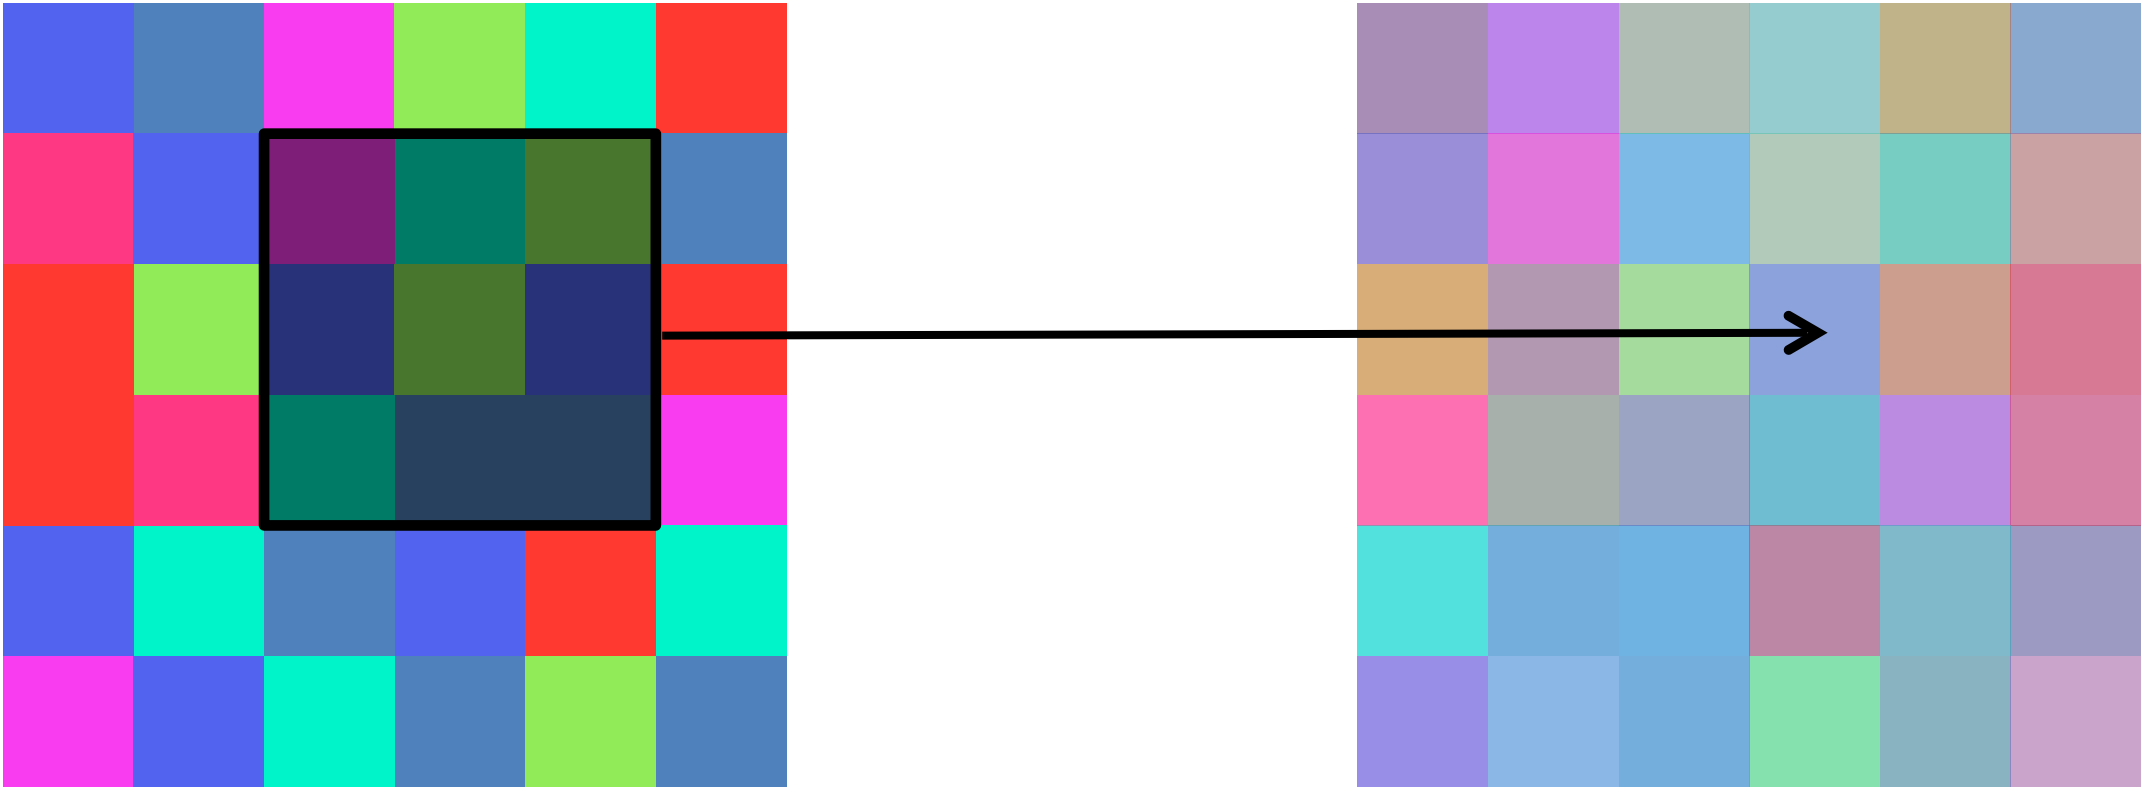
Lecture 3b: Introduction to Filterings

Point Operation



point processing

Neighborhood Operation

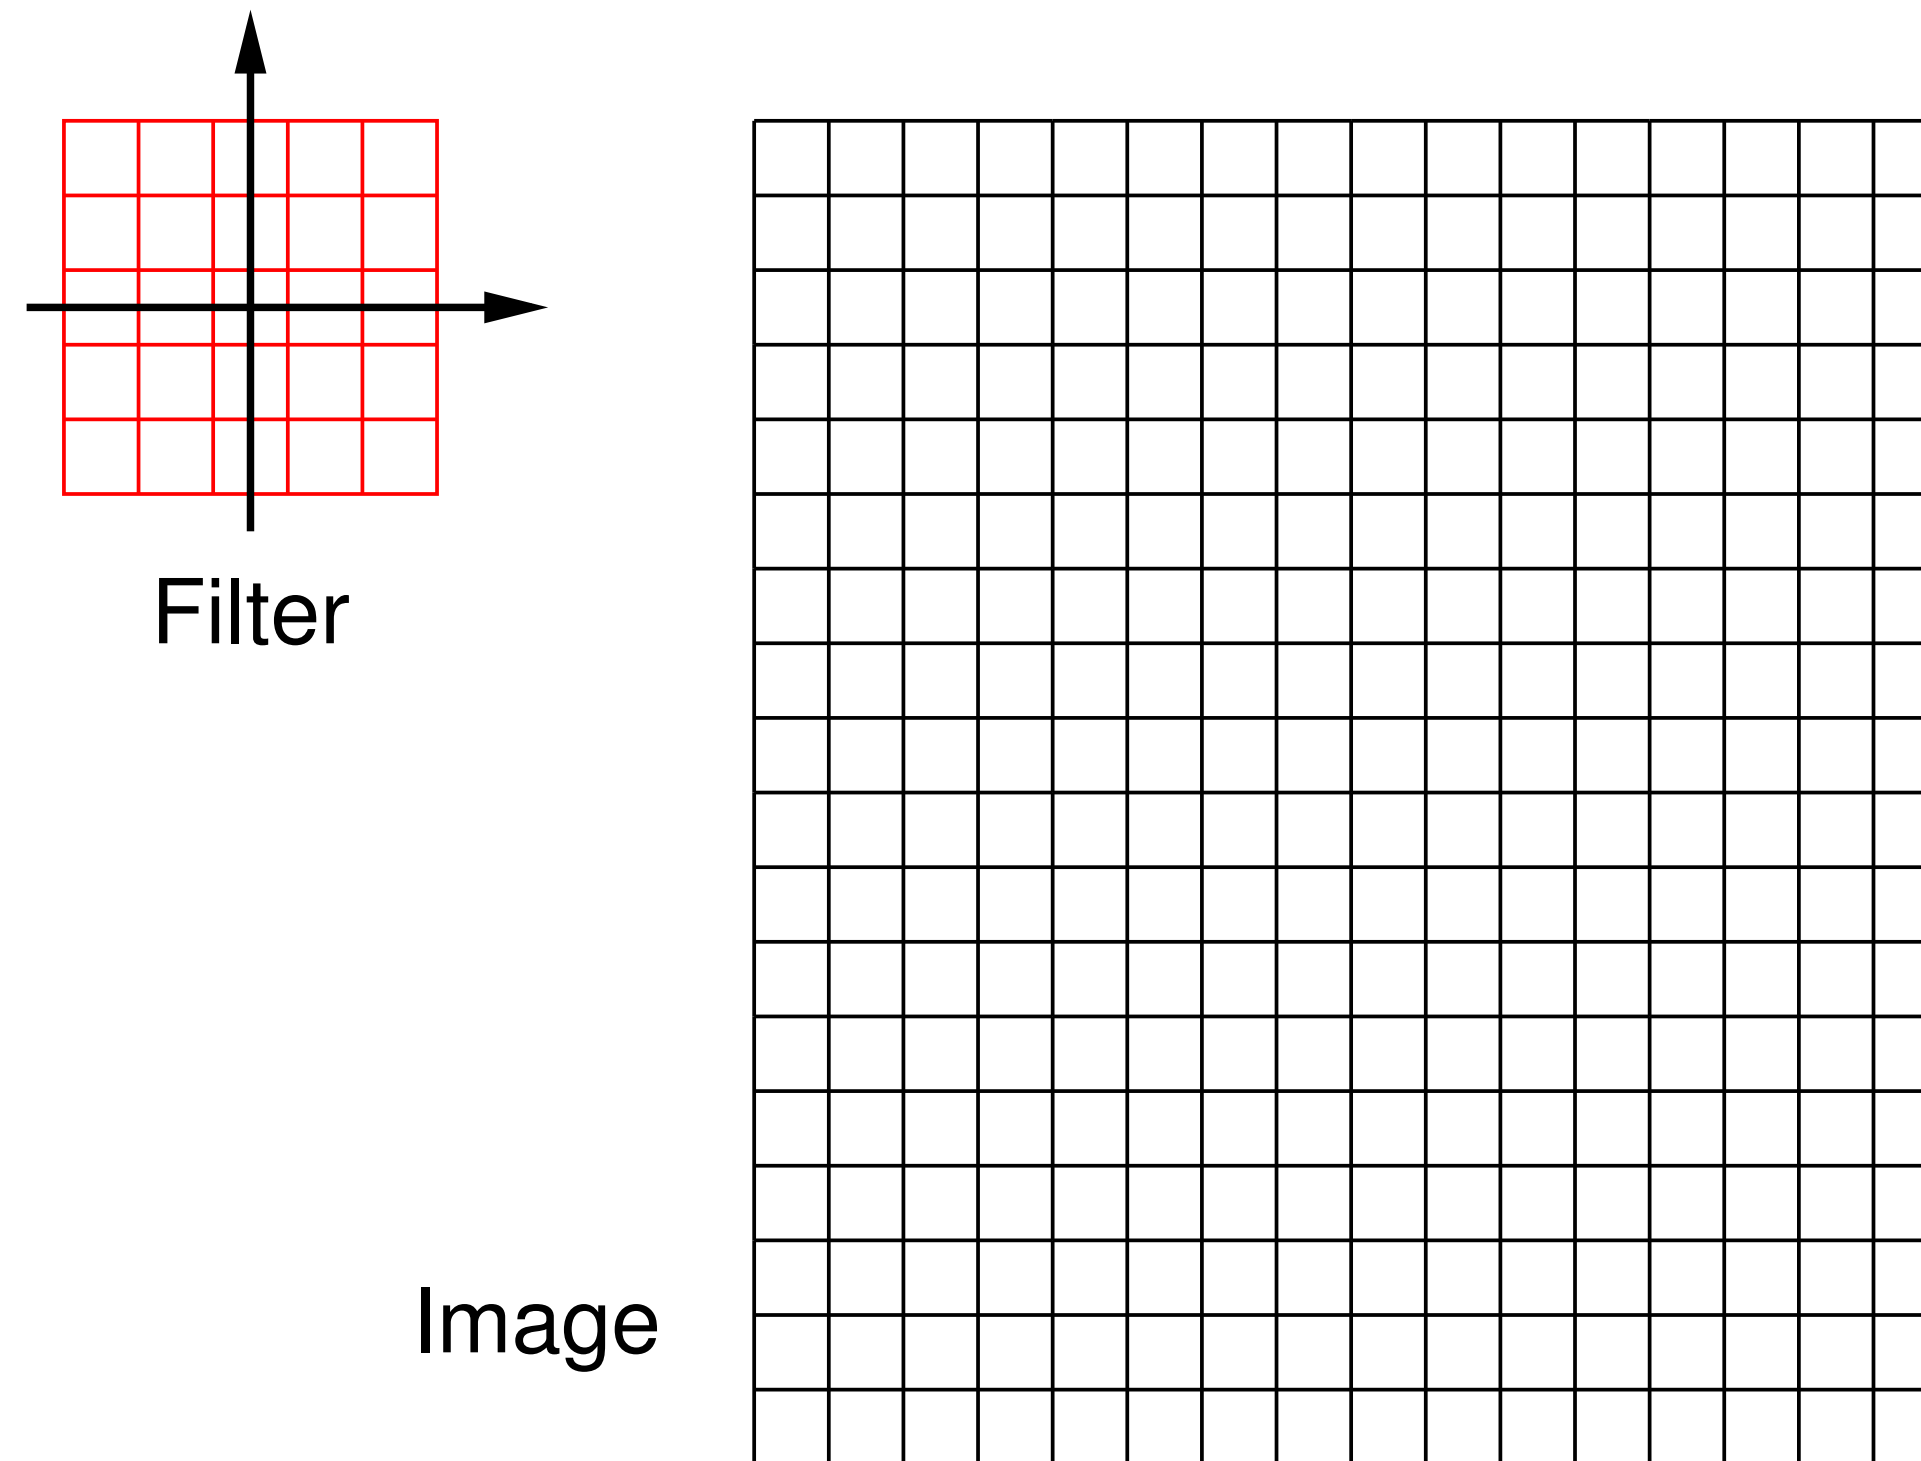


“filtering”

Linear **Filters**

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

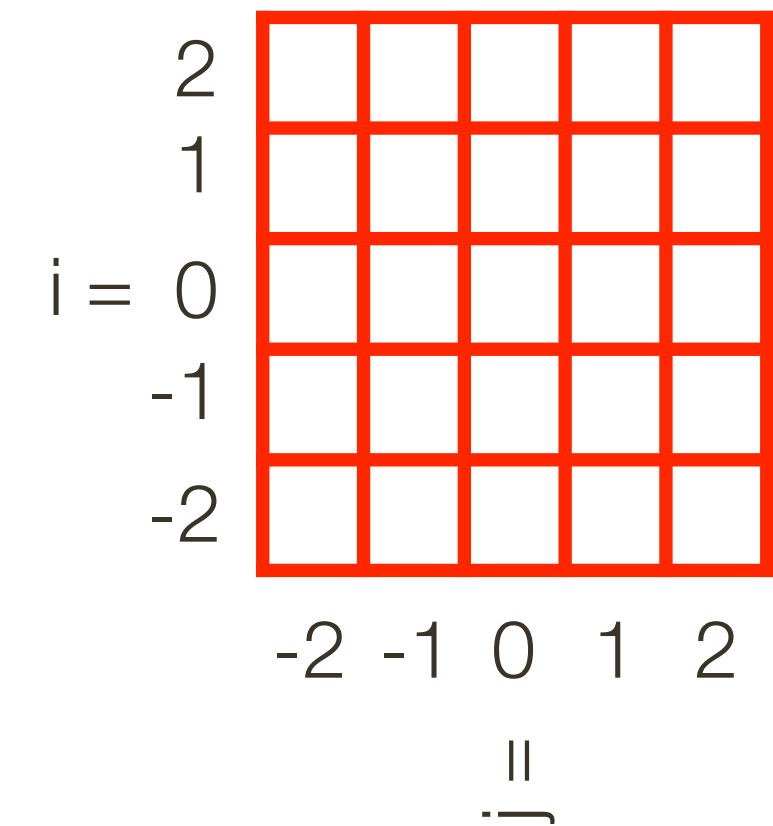
Let $F(X, Y)$ be another $m \times m$ digital image (our “**filter**” or “**kernel**”)



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

Linear Filters

$$\text{Let } k = \left\lfloor \frac{m}{2} \right\rfloor$$



Compute a new image, $I'(X, Y)$, as follows

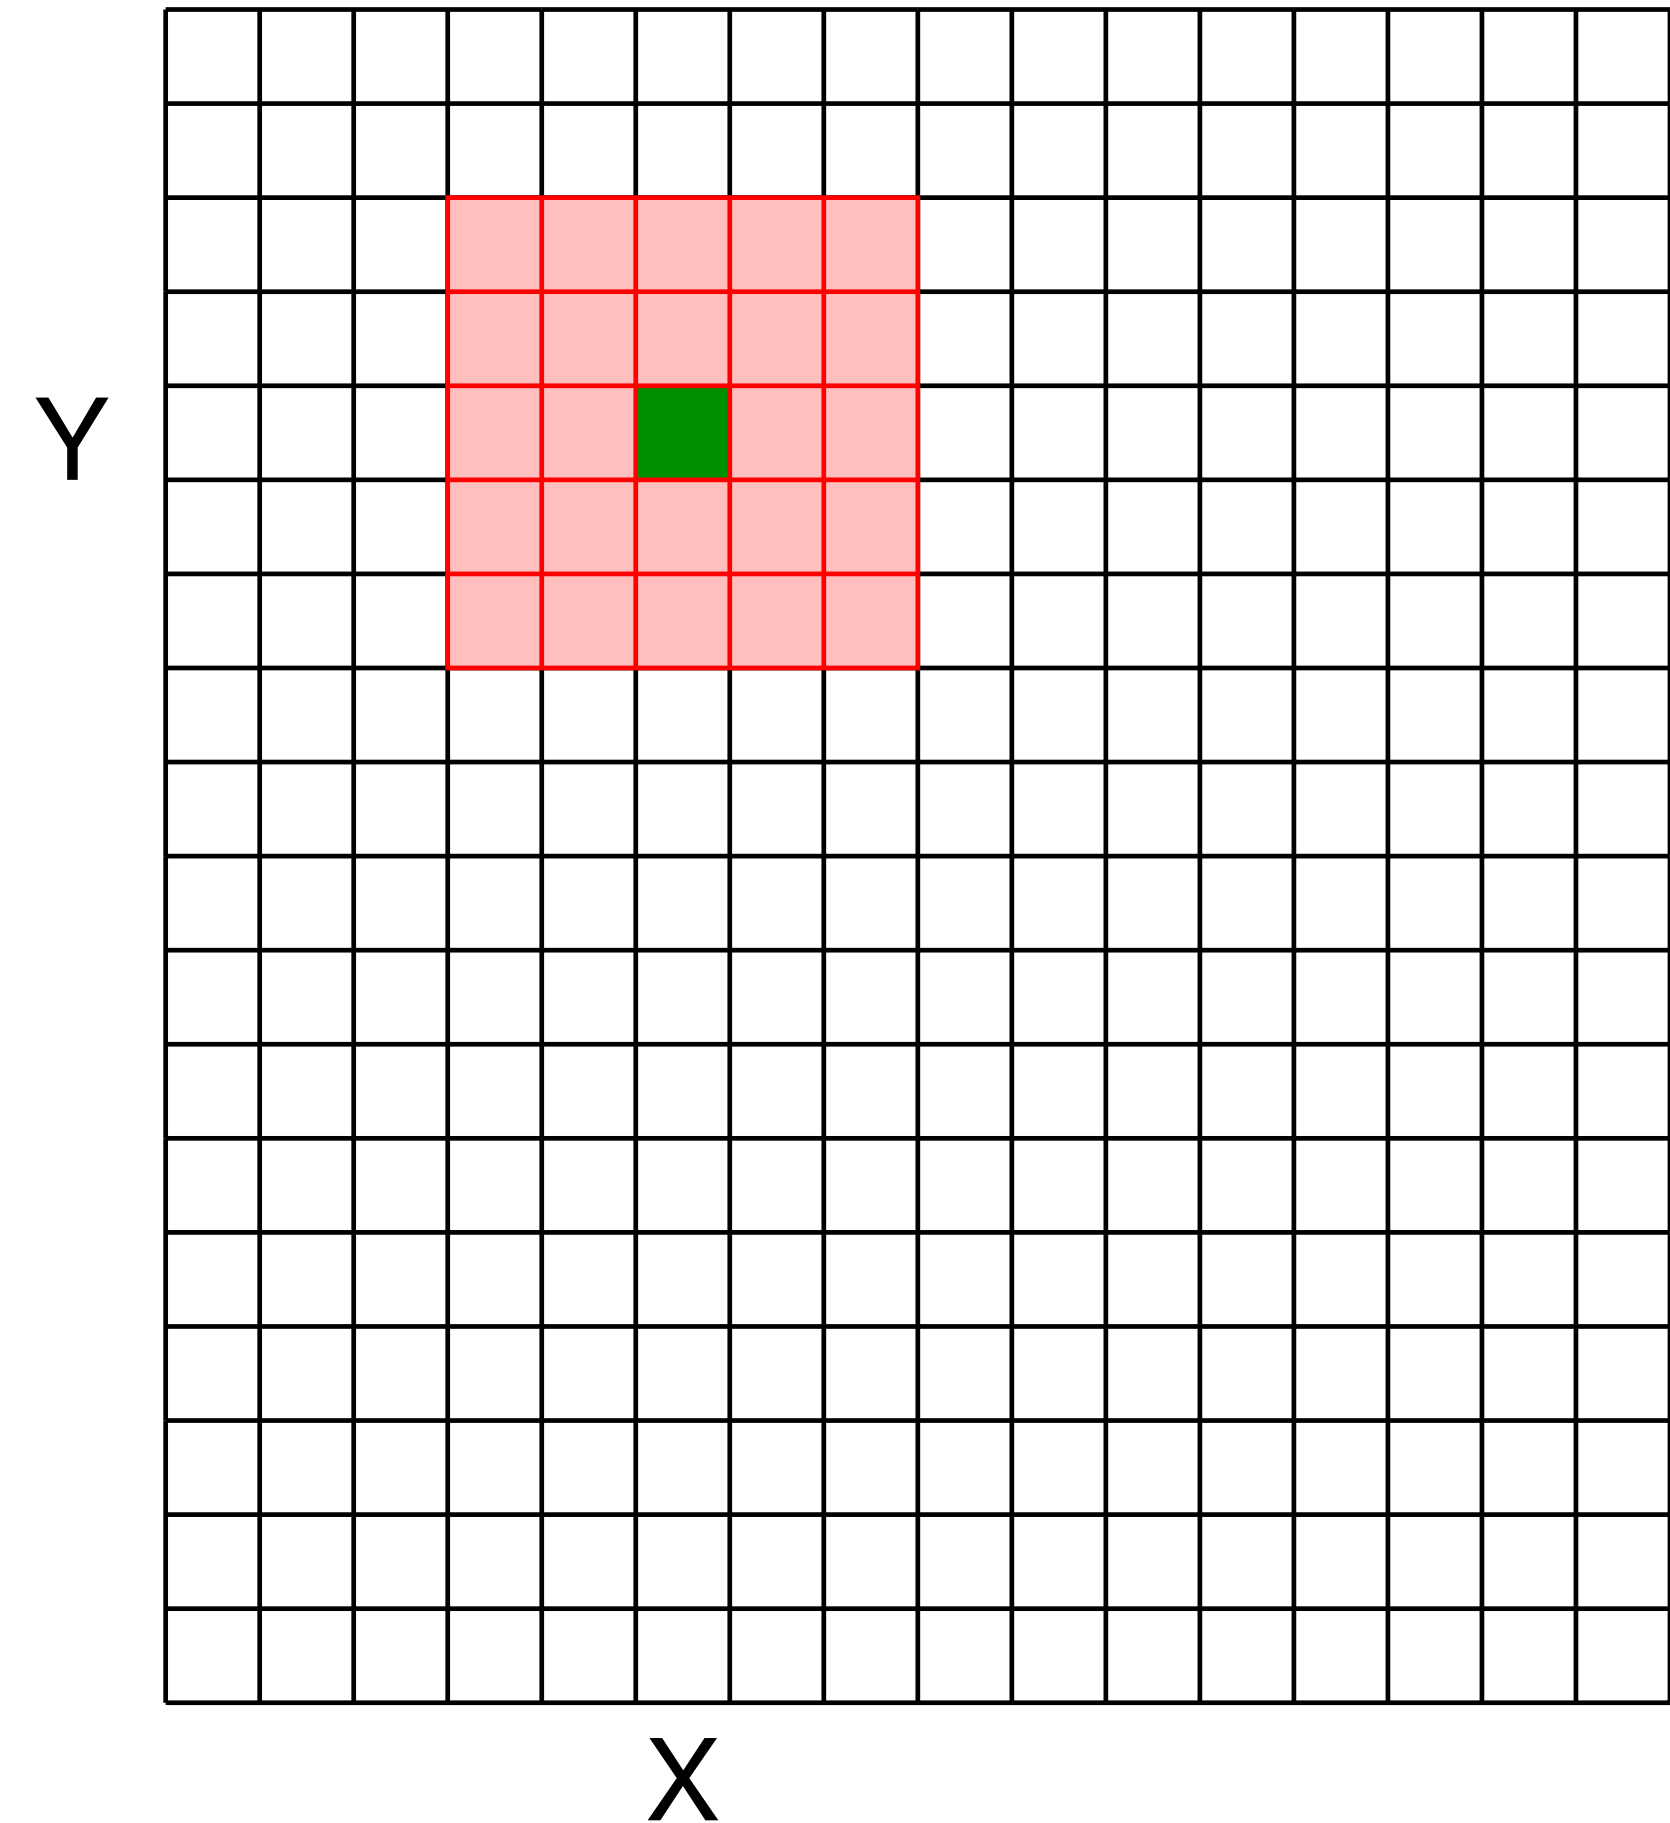
$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(I, J) I(X + i, Y + j)$$

The equation is annotated with colored boxes: a light green box around $I'(X, Y)$ labeled "output", a light blue box around $F(I, J)$ labeled "filter", and a light purple box around $I(X + i, Y + j)$ labeled "image (signal)".

Intuition: each pixel in the output image is a linear combination of the same pixel and its neighboring pixels in the original image

Linear **Filters**

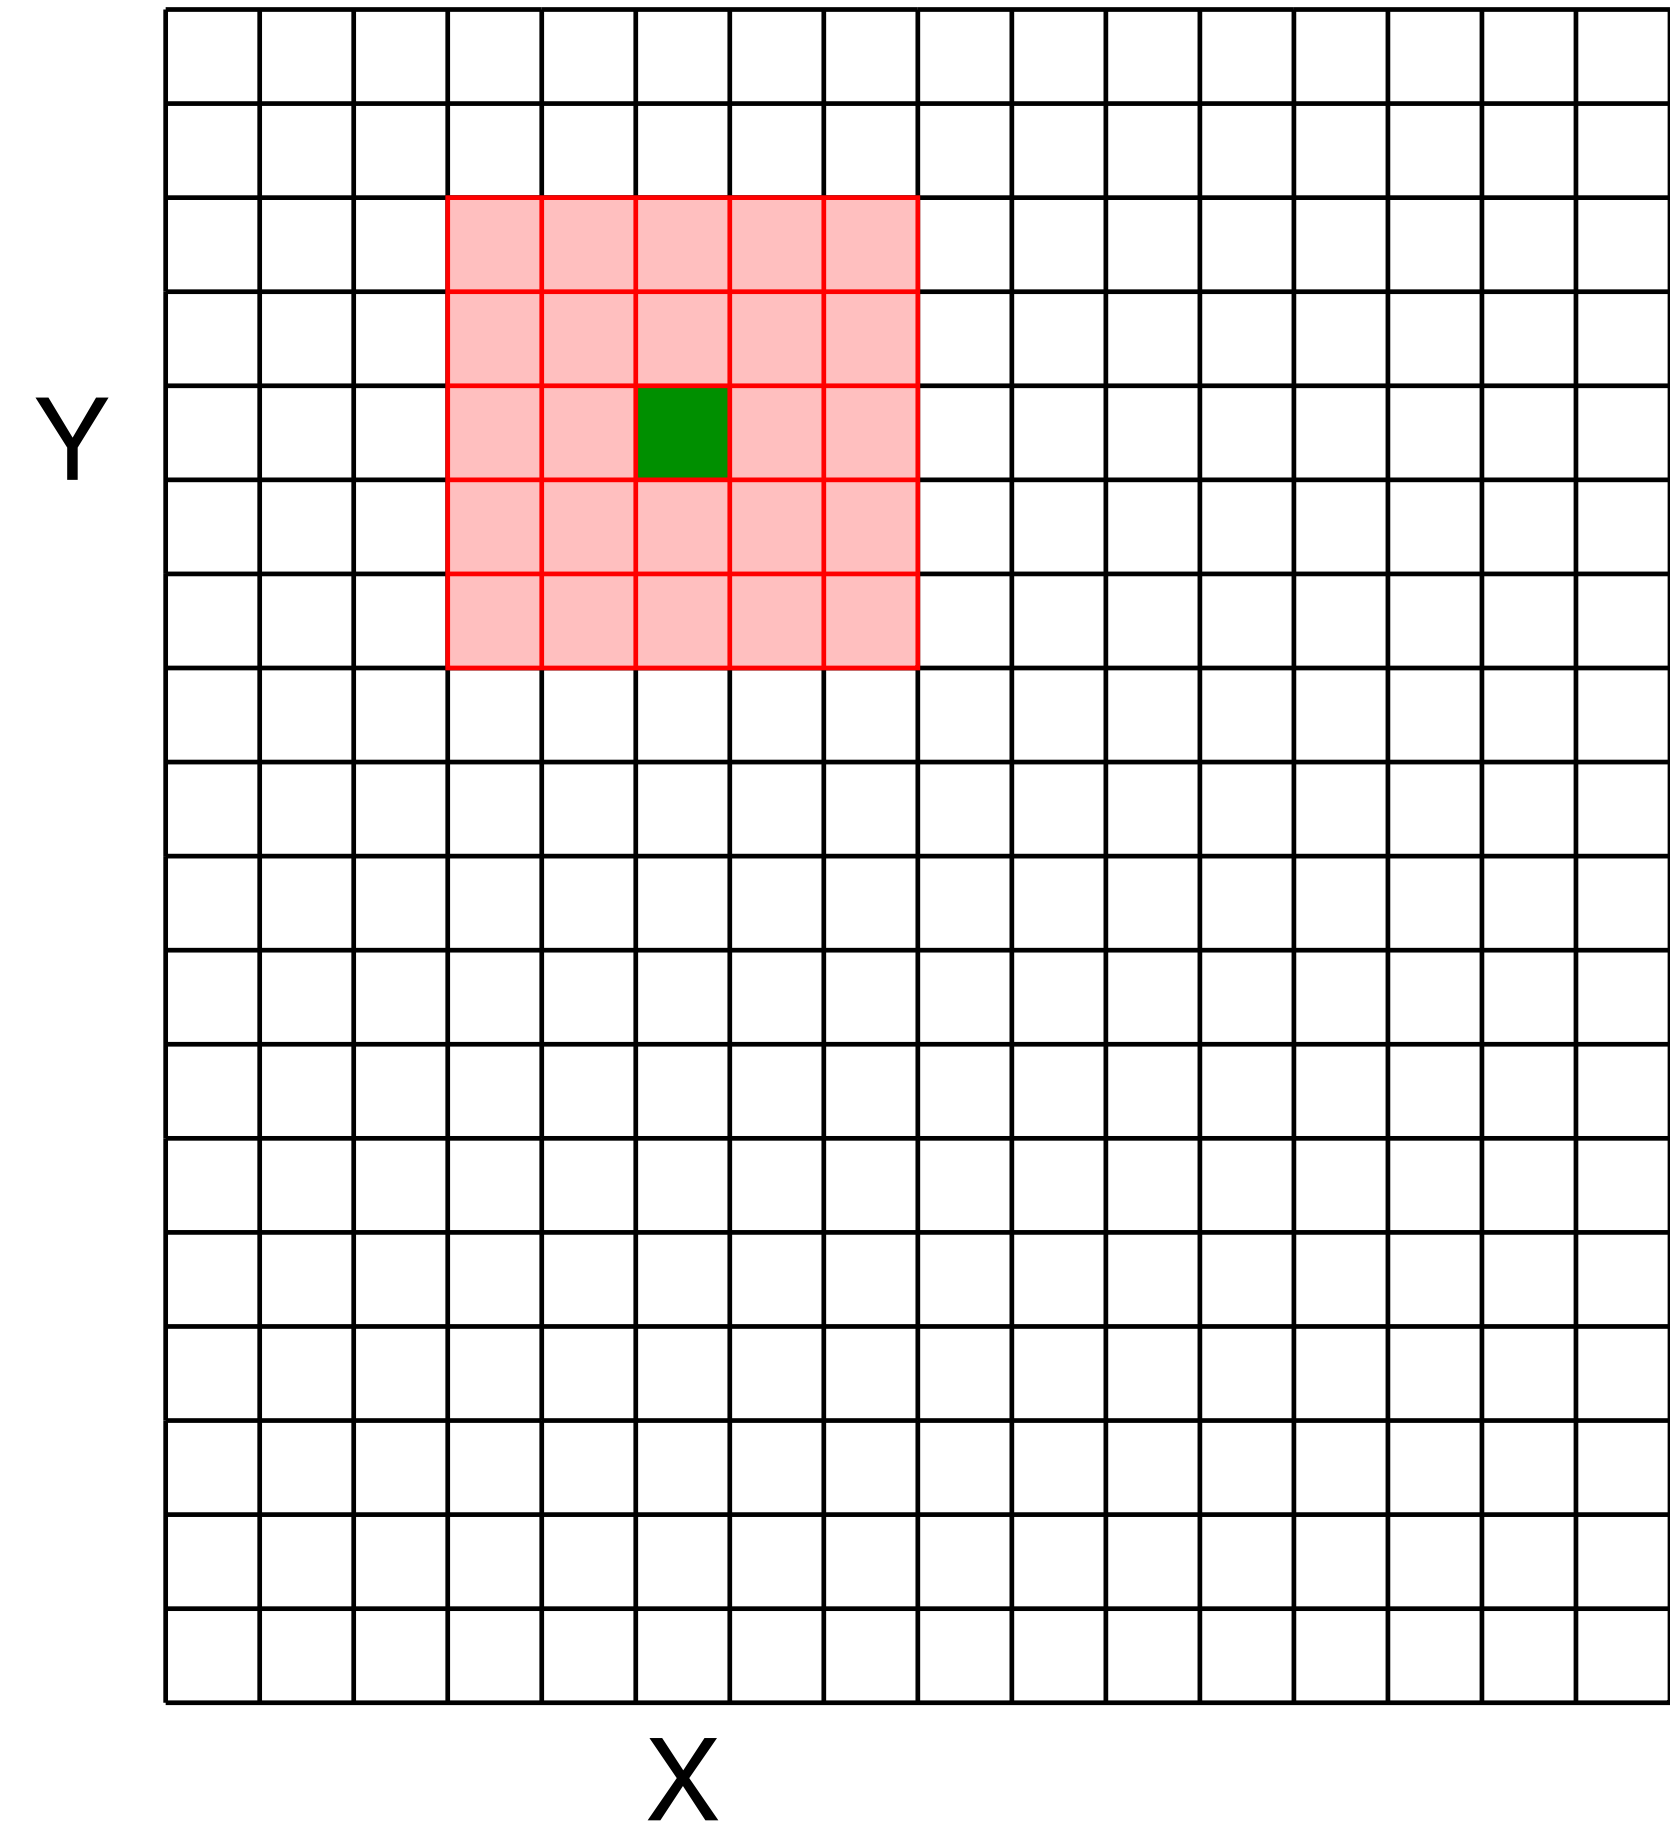
For a give X and Y , superimpose the filter on the image centered at (X, Y)



Linear **Filters**

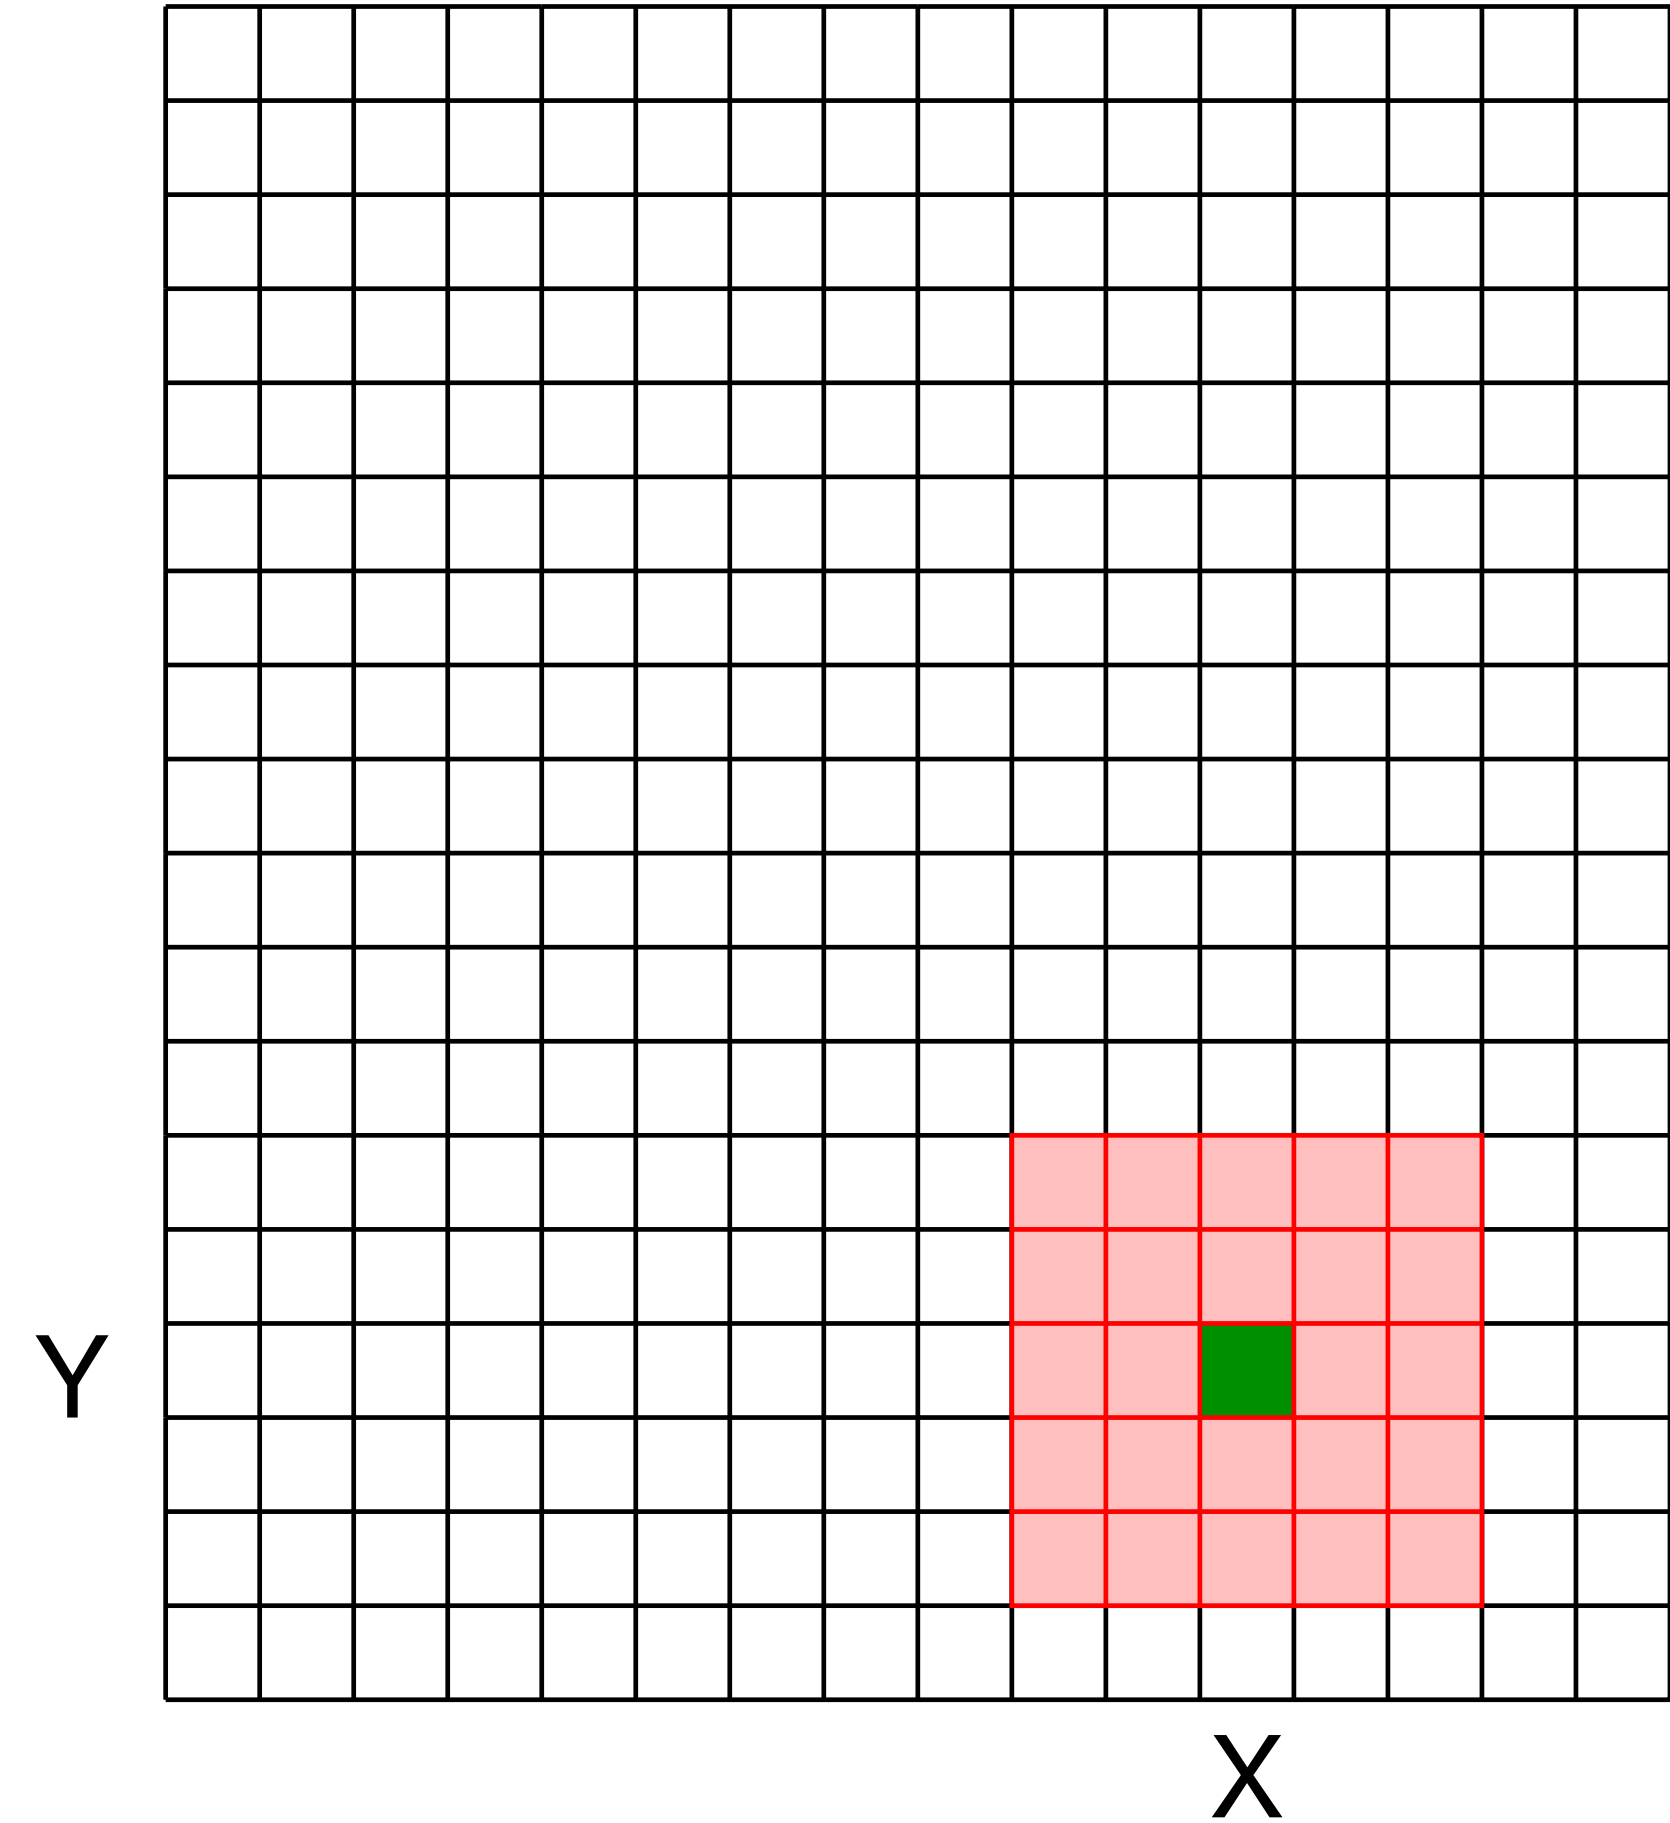
For a give X and Y , superimpose the filter on the image centered at (X, Y)

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

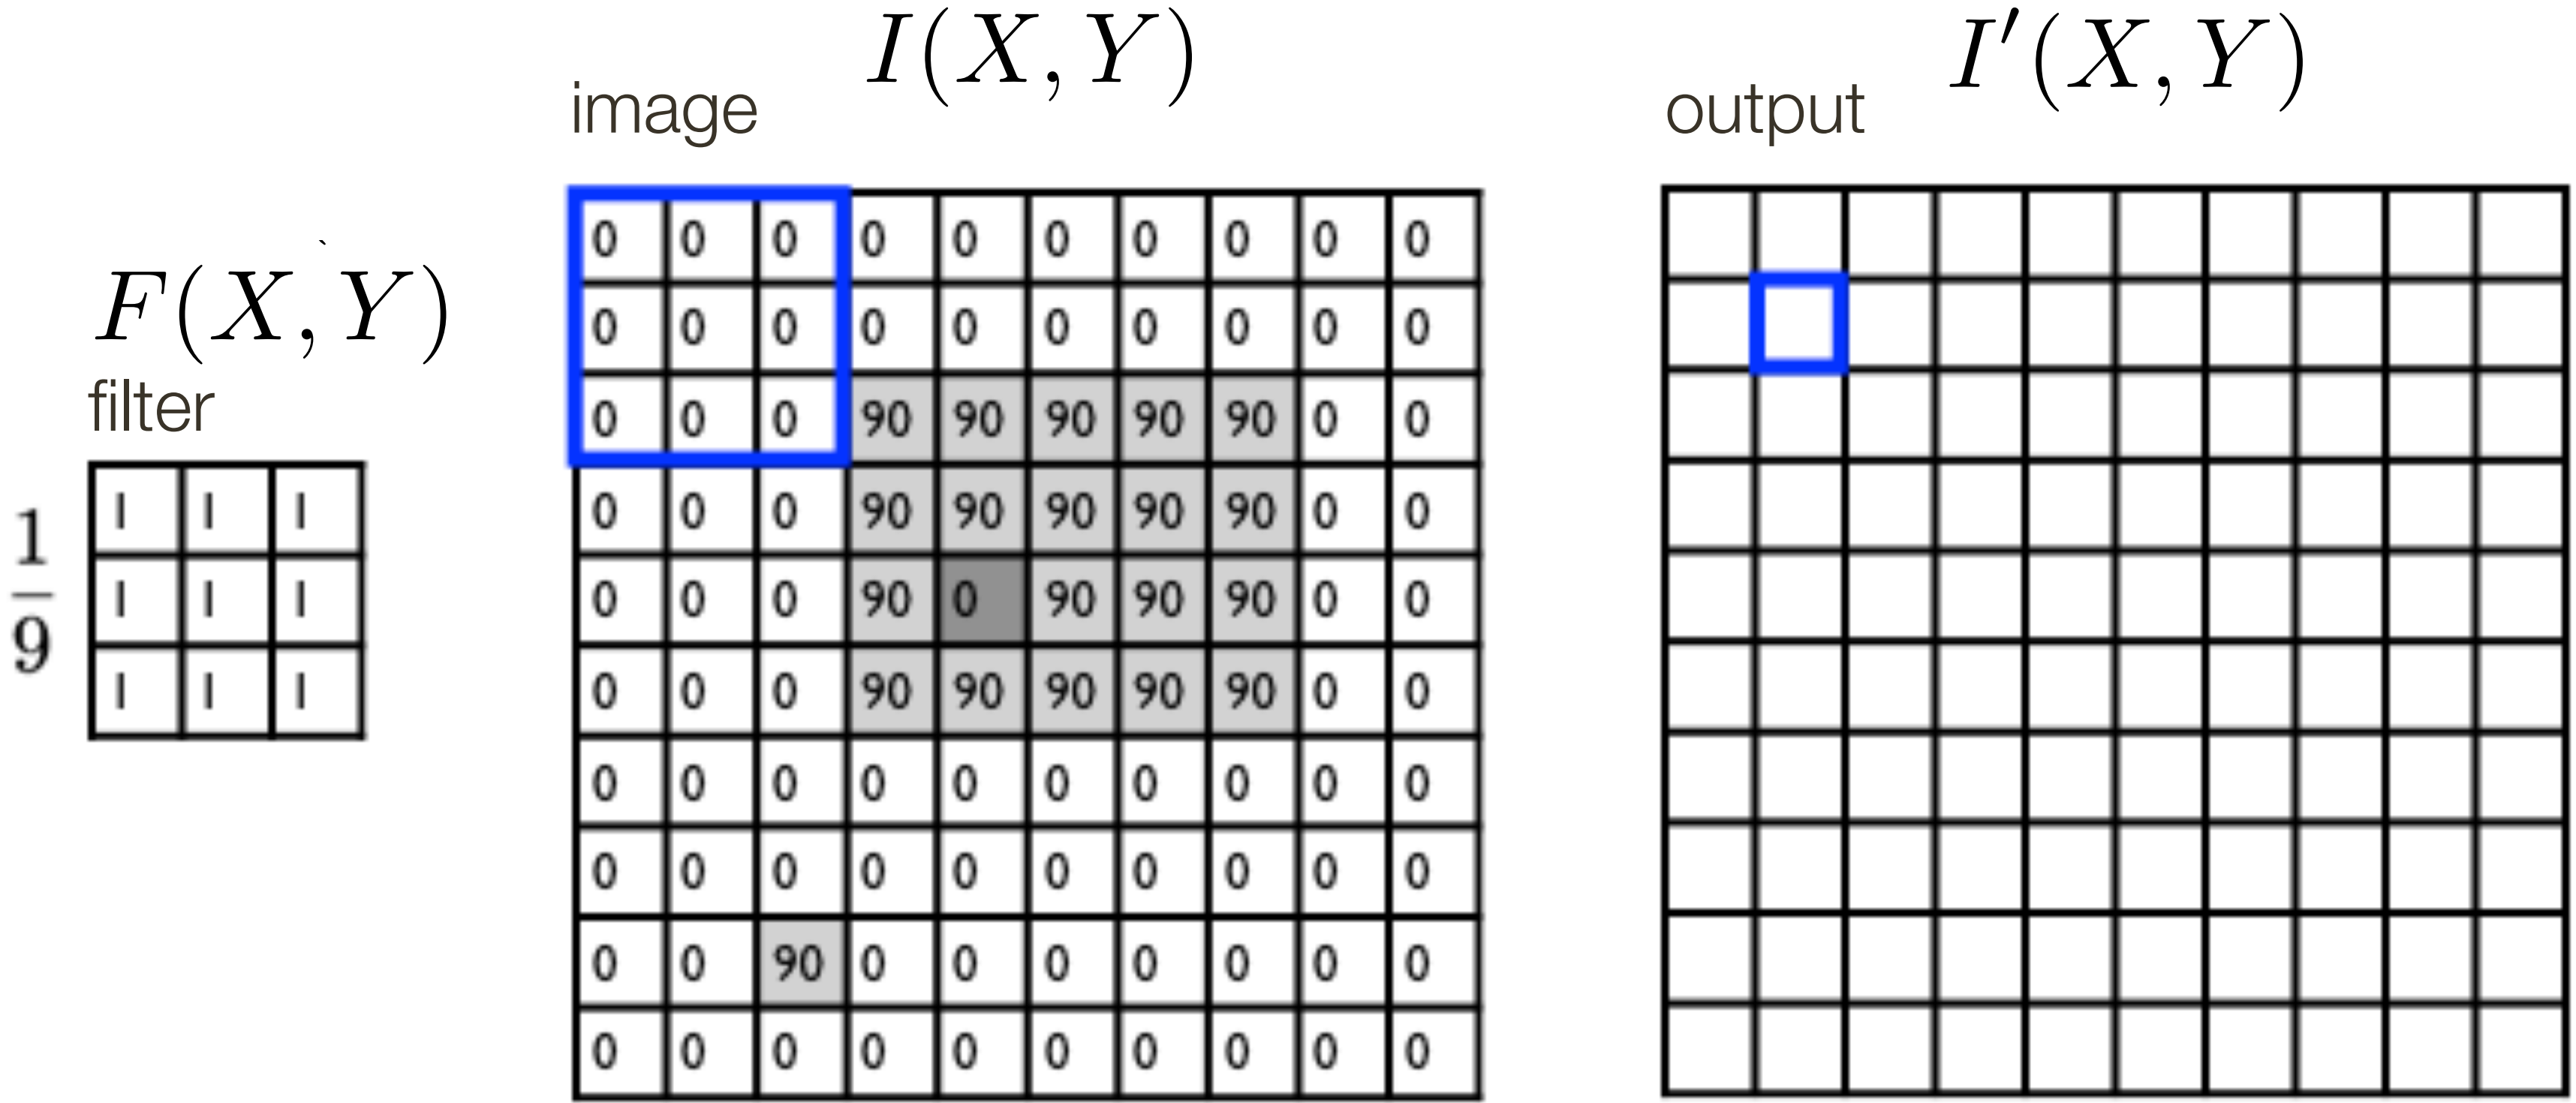


Linear **Filters**

The computation is repeated for each
 (X, Y)

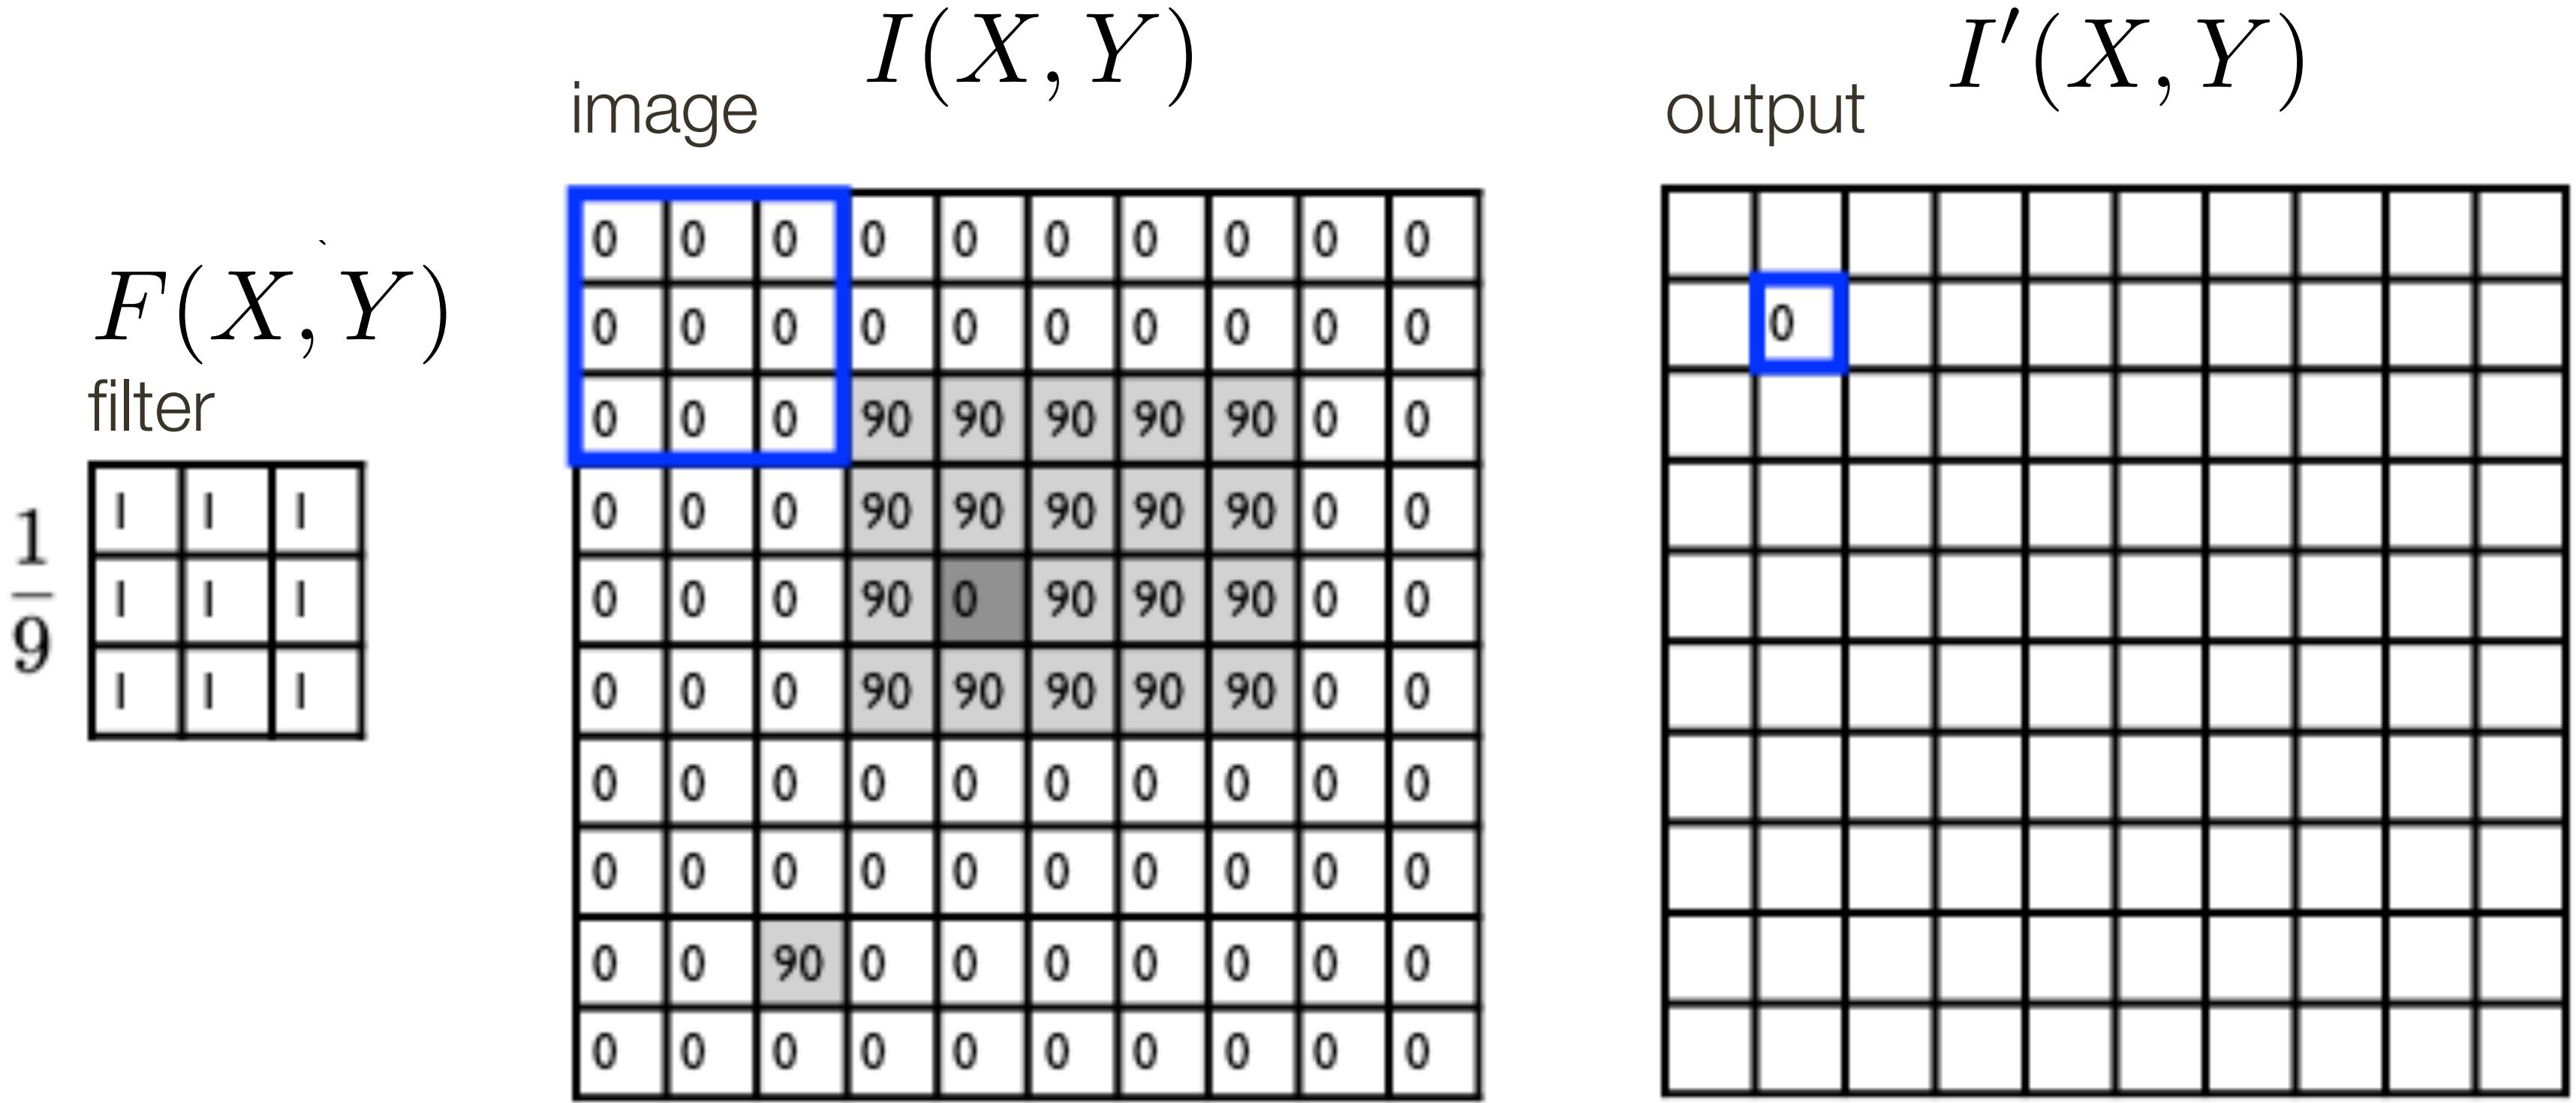


Linear Filter Example



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

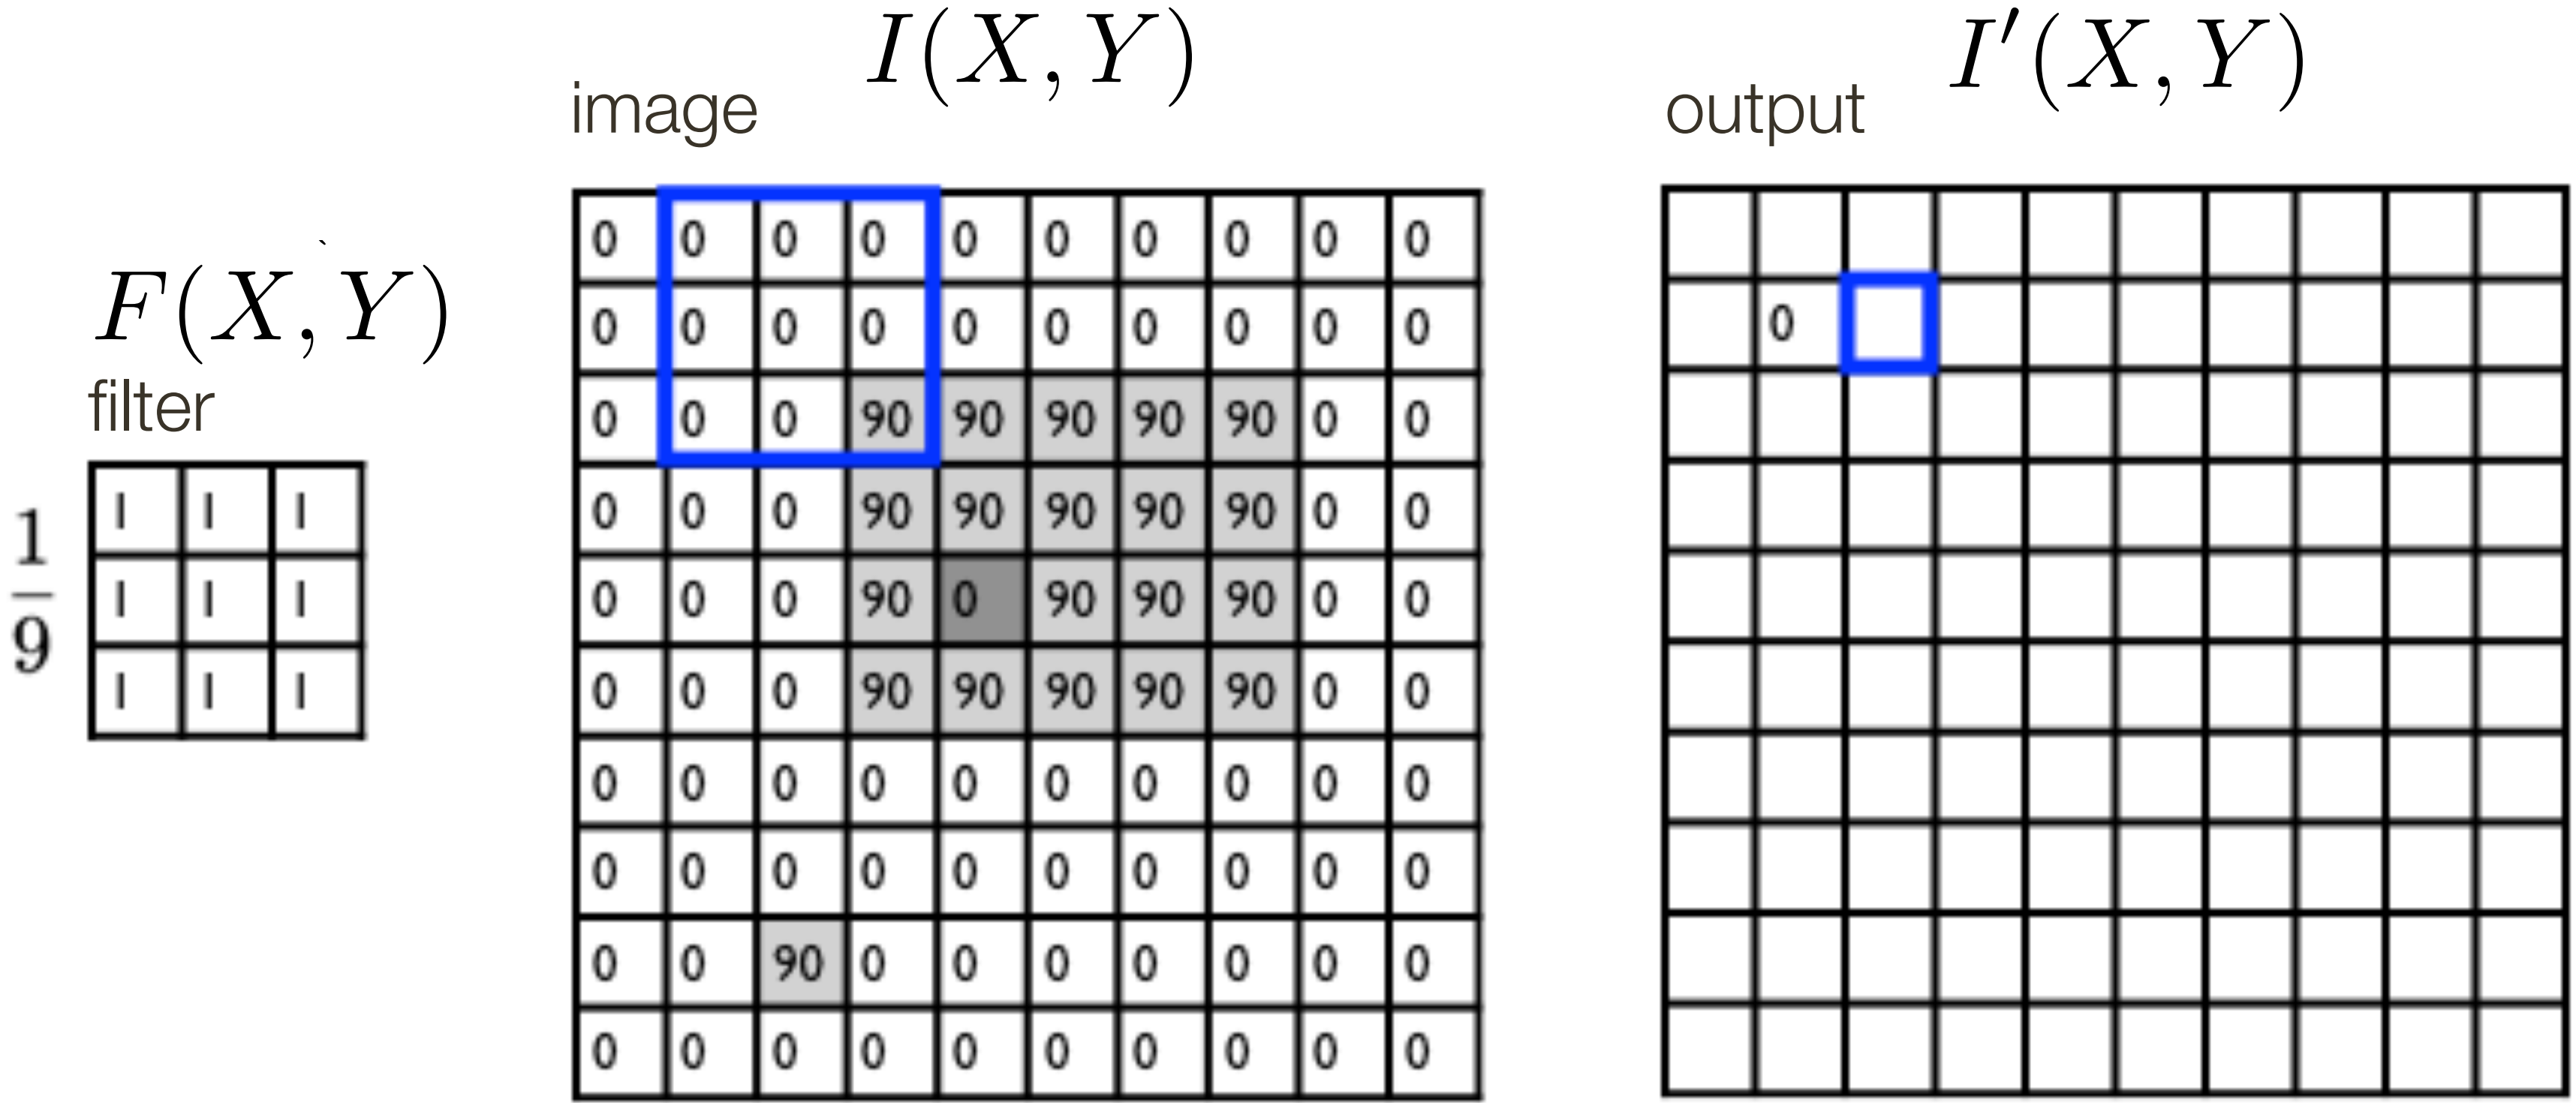
Linear Filter Example



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

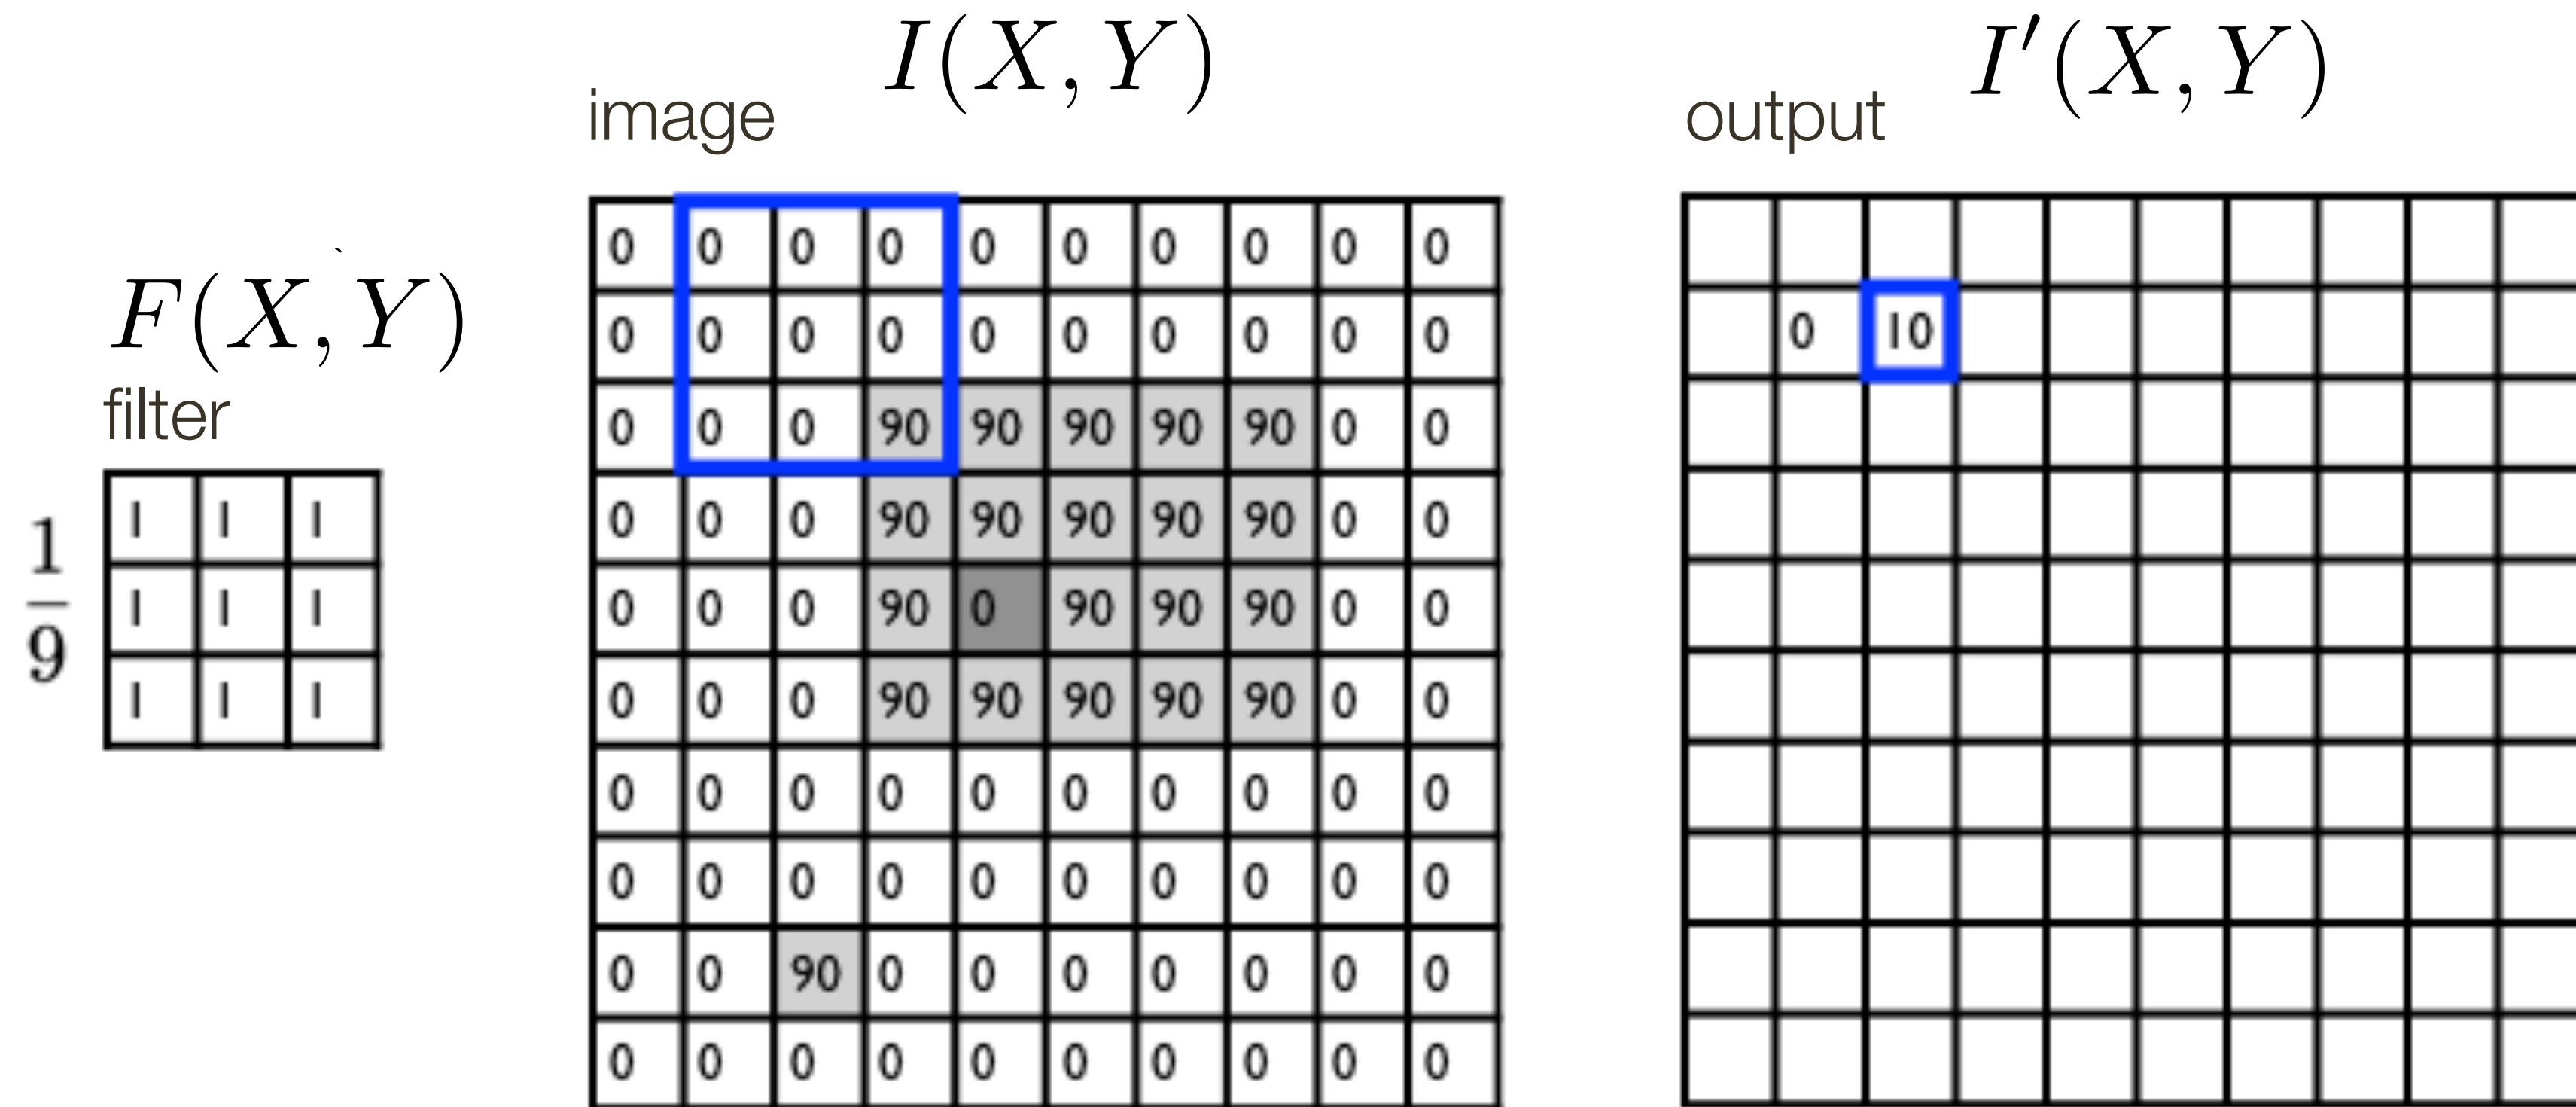
output
filter
image (signal)

Linear Filter Example



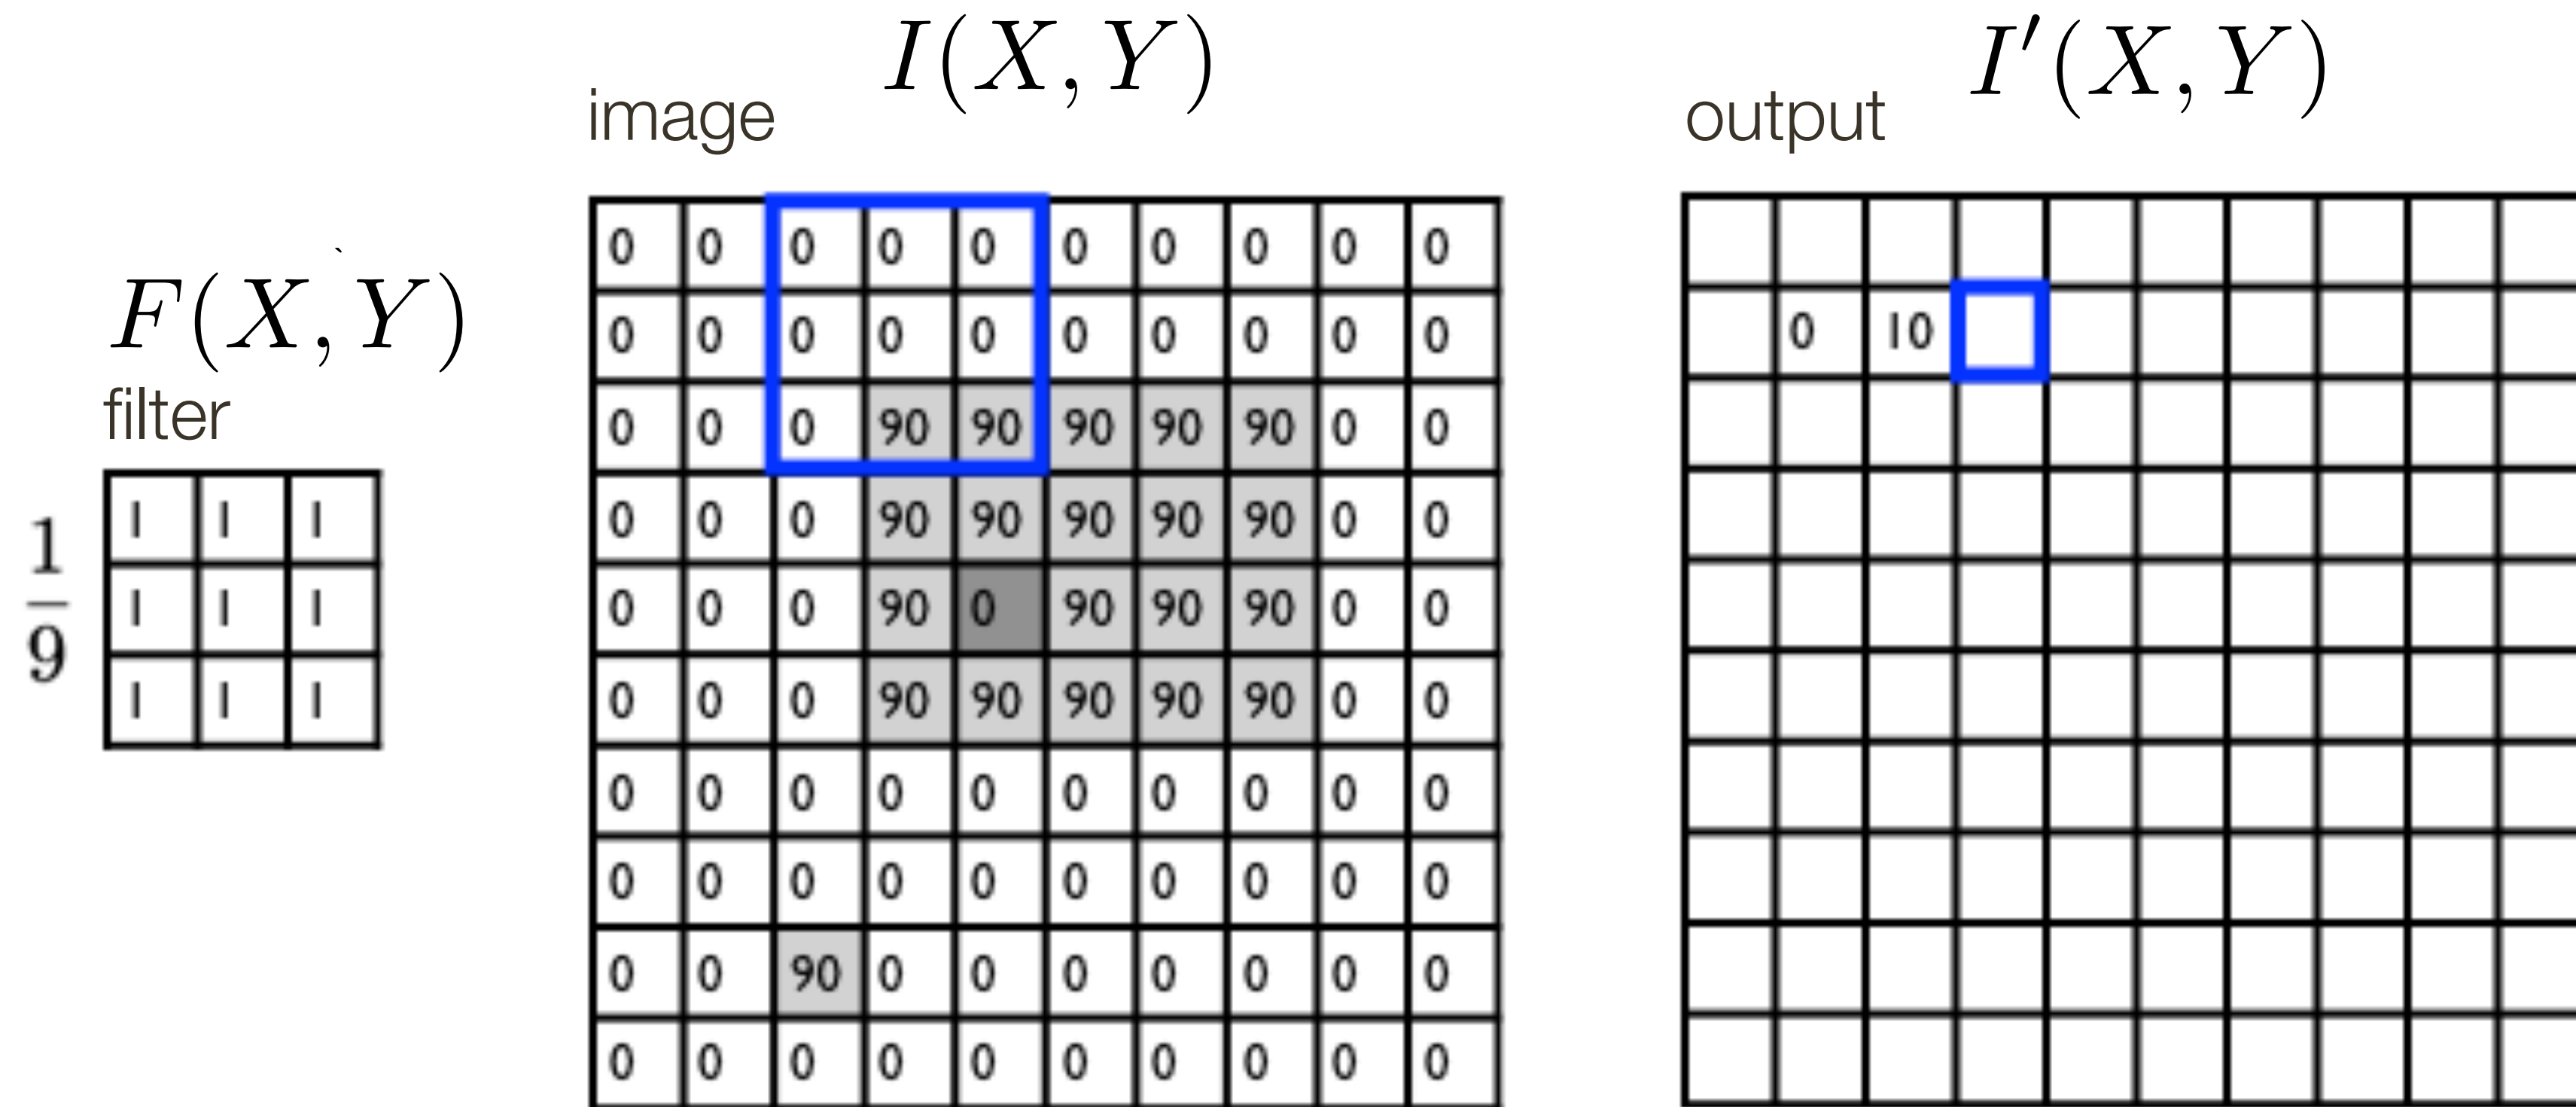
$$\underbrace{I'(X, Y)}_{\text{output}} = \sum_{j=-k}^k \sum_{i=-k}^k \underbrace{F(I, J)}_{\text{filter}} \underbrace{I(X + i, Y + j)}_{\text{image (signal)}}$$

Linear Filter Example



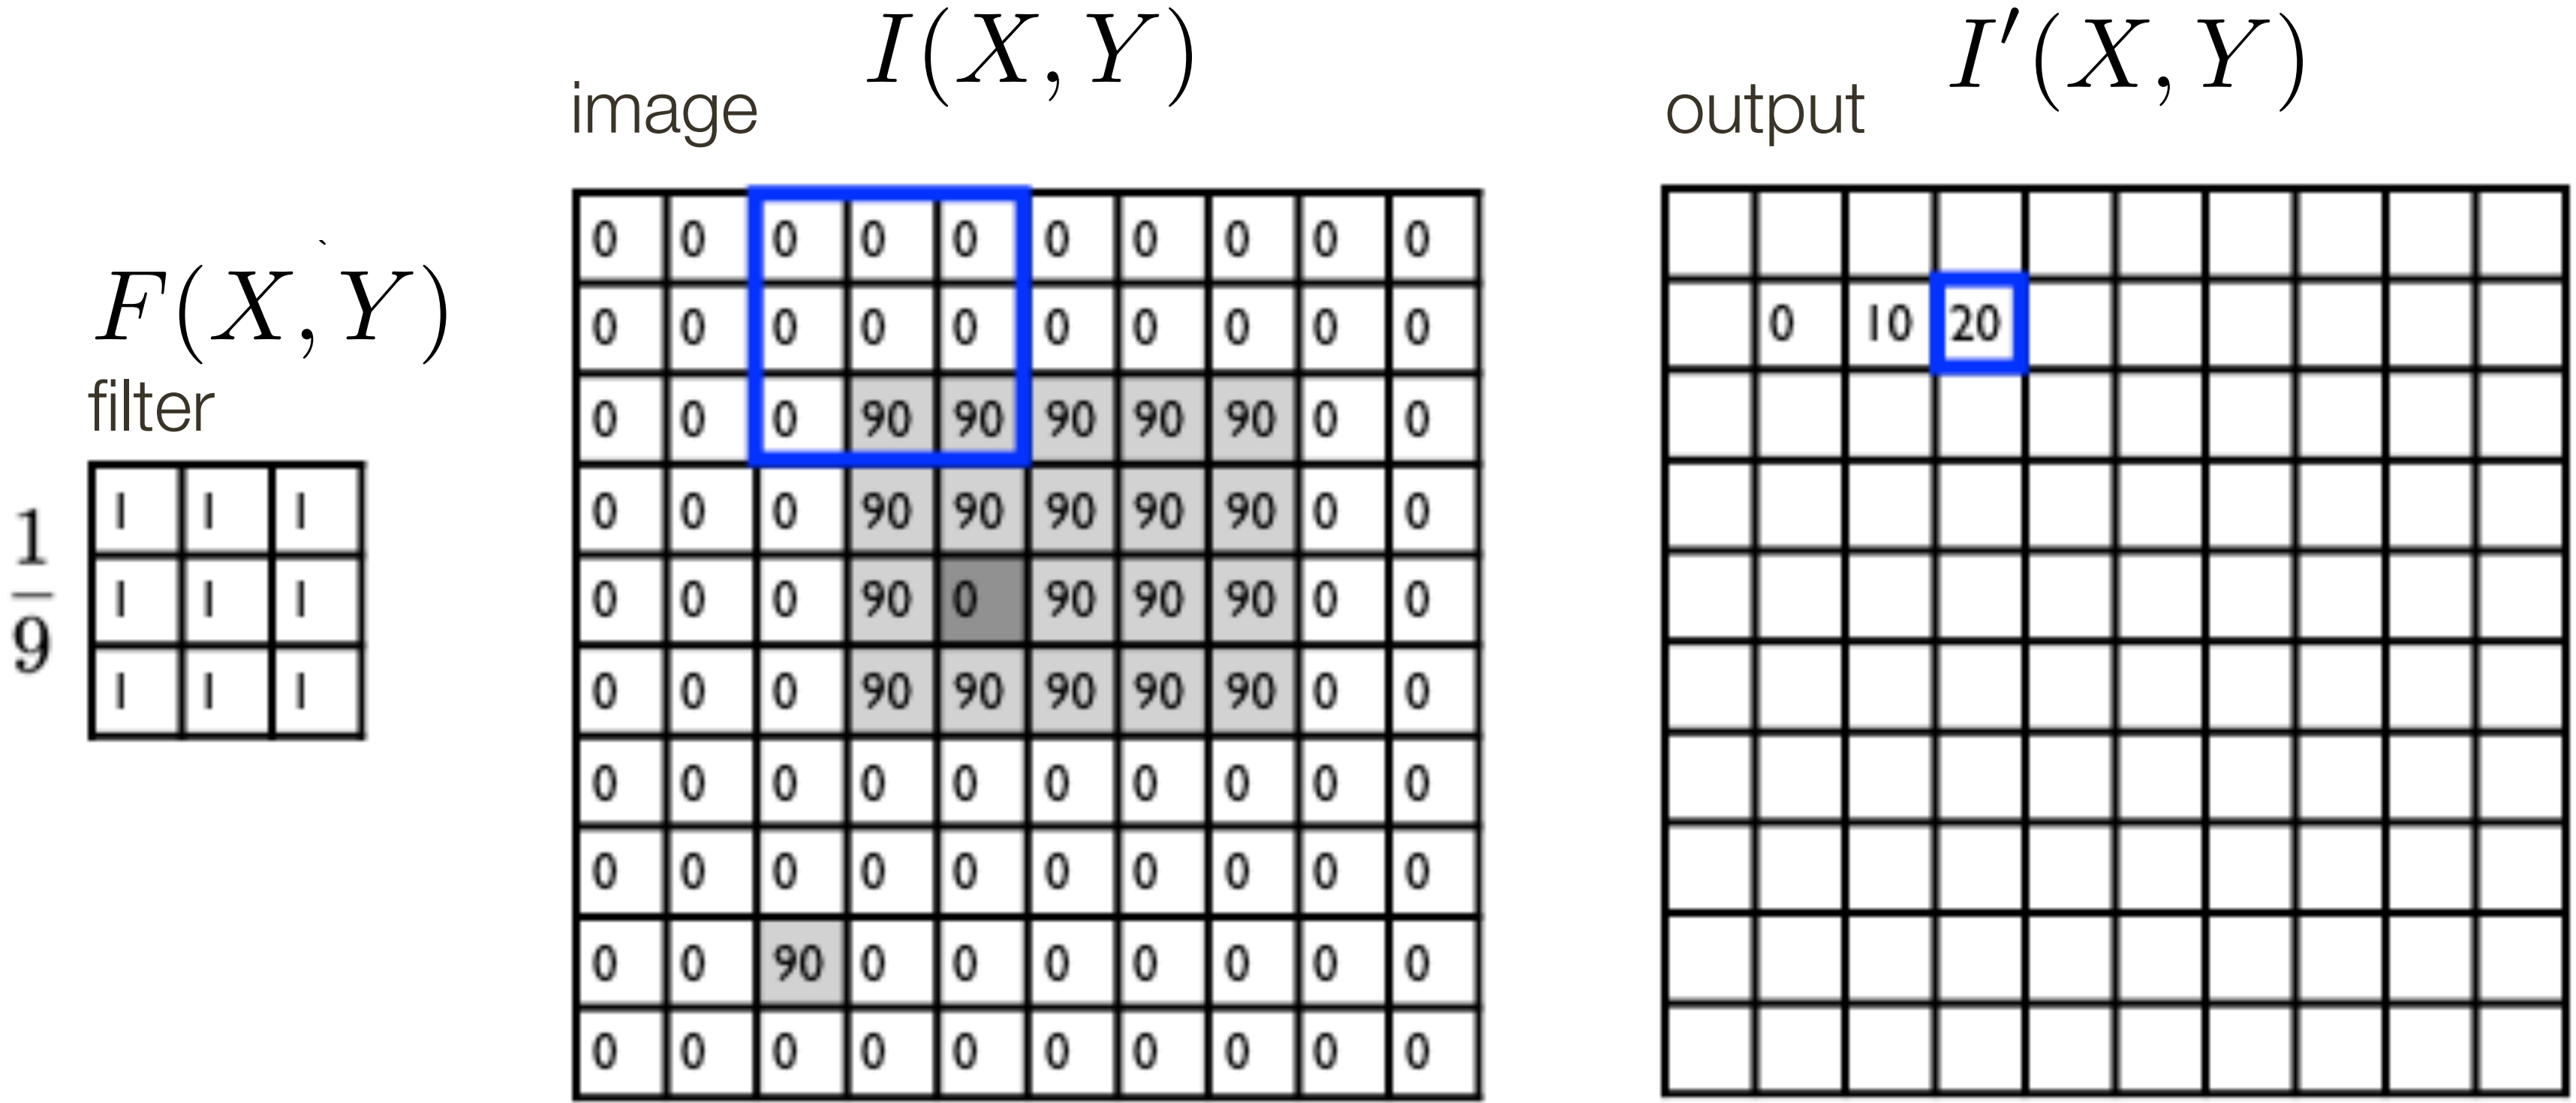
$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(I, J) I(X + i, Y + j)$$

Linear Filter Example



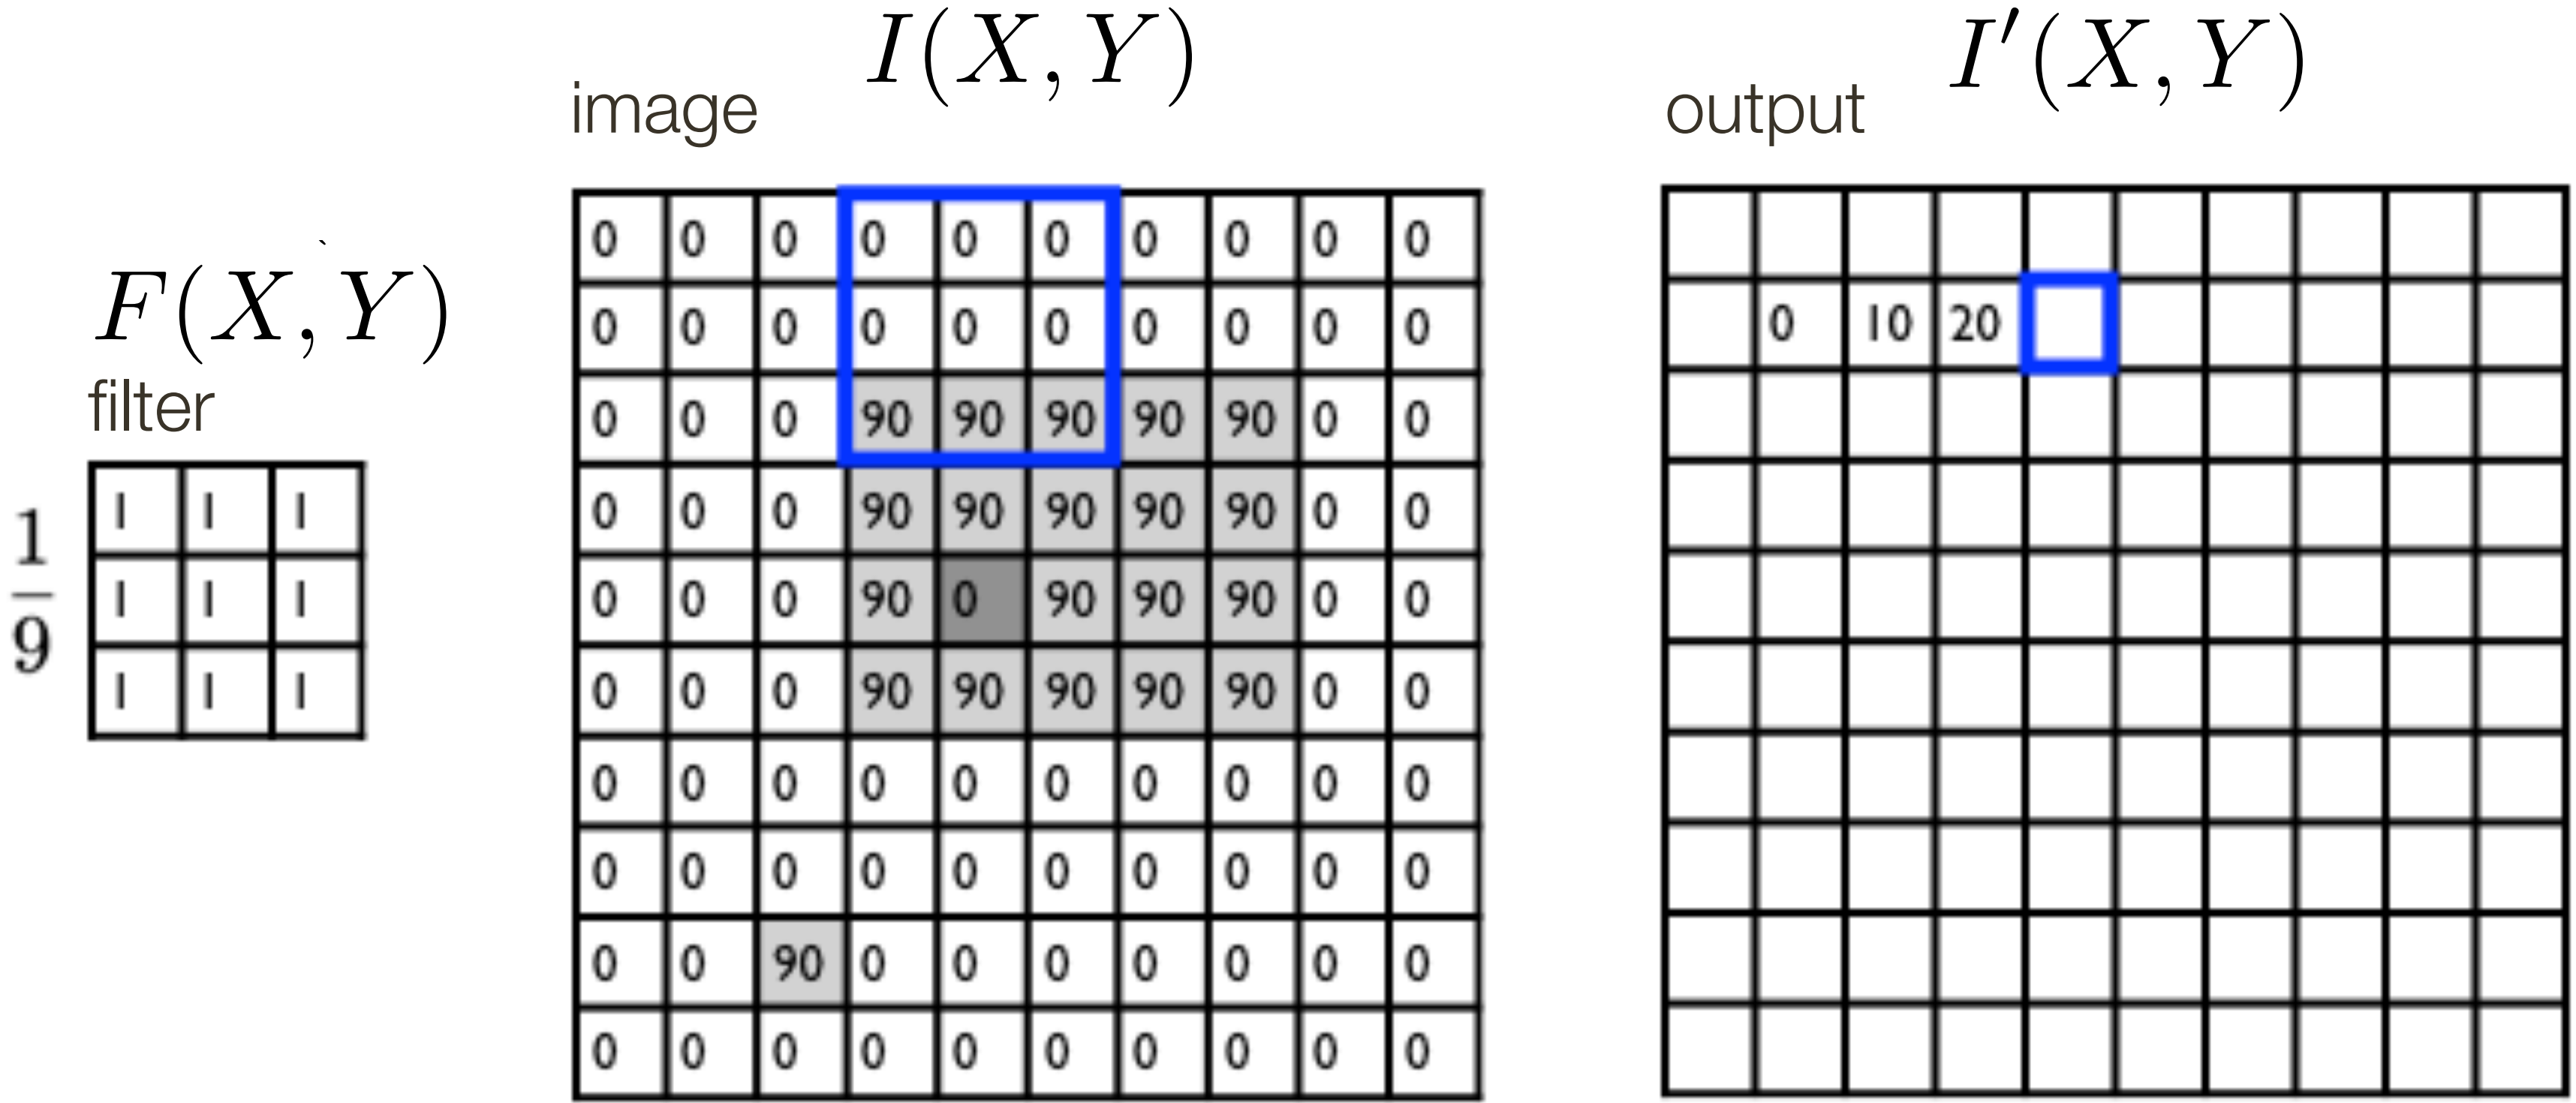
$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(I, J) I(X + i, Y + j)$$

Linear Filter Example



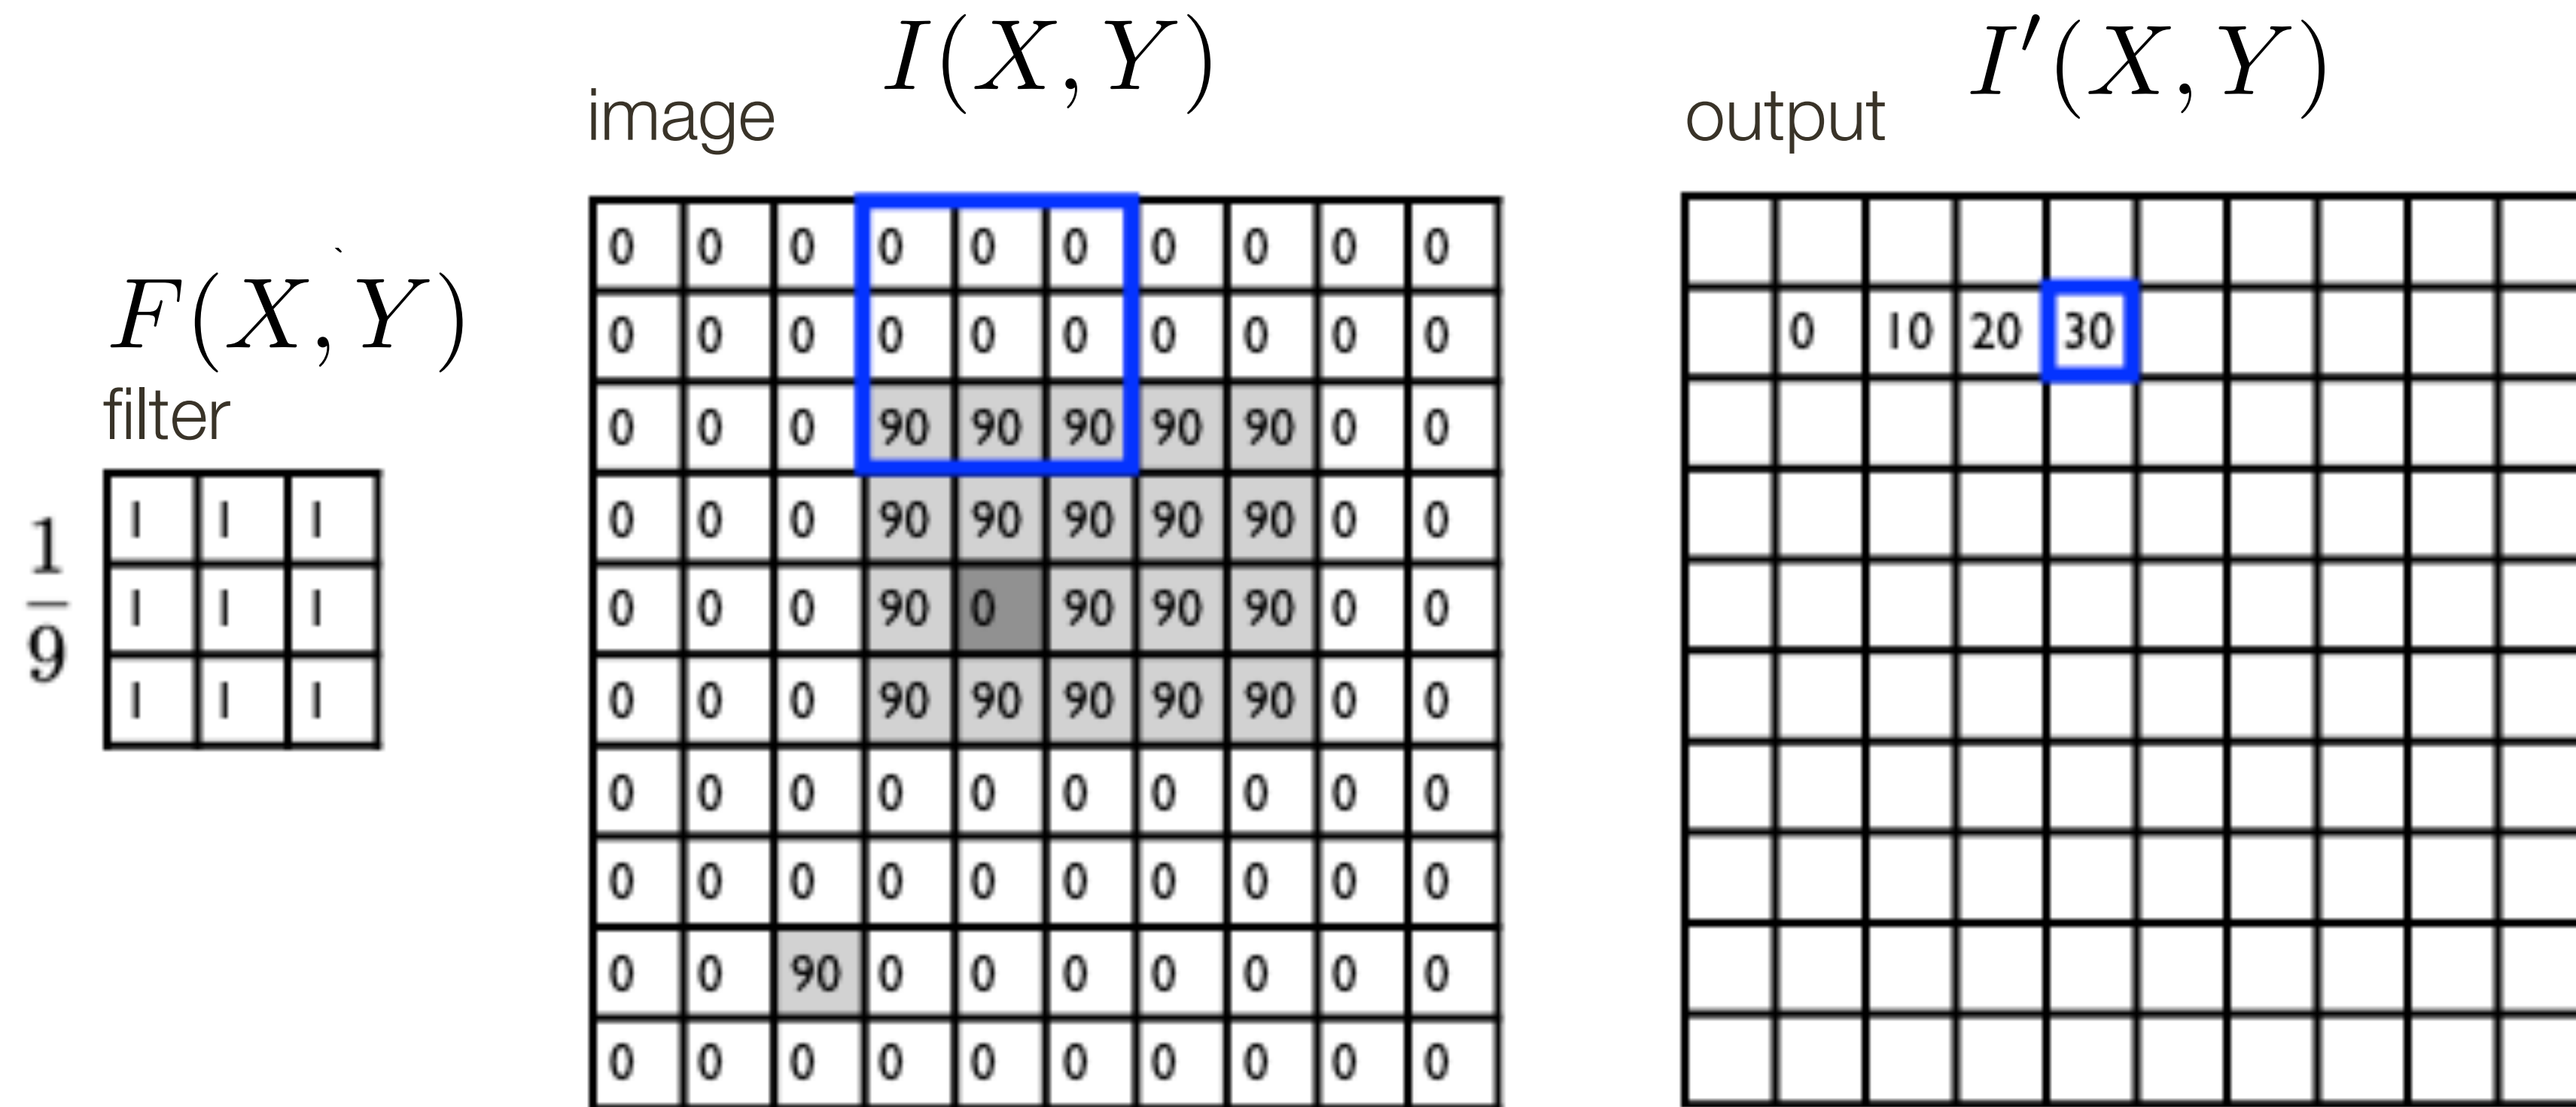
$$\underbrace{I'(X, Y)}_{\text{output}} = \sum_{j=-k}^k \sum_{i=-k}^k \underbrace{F(I, J)}_{\text{filter}} \underbrace{I(X + i, Y + j)}_{\text{image (signal)}}$$

Linear Filter Example



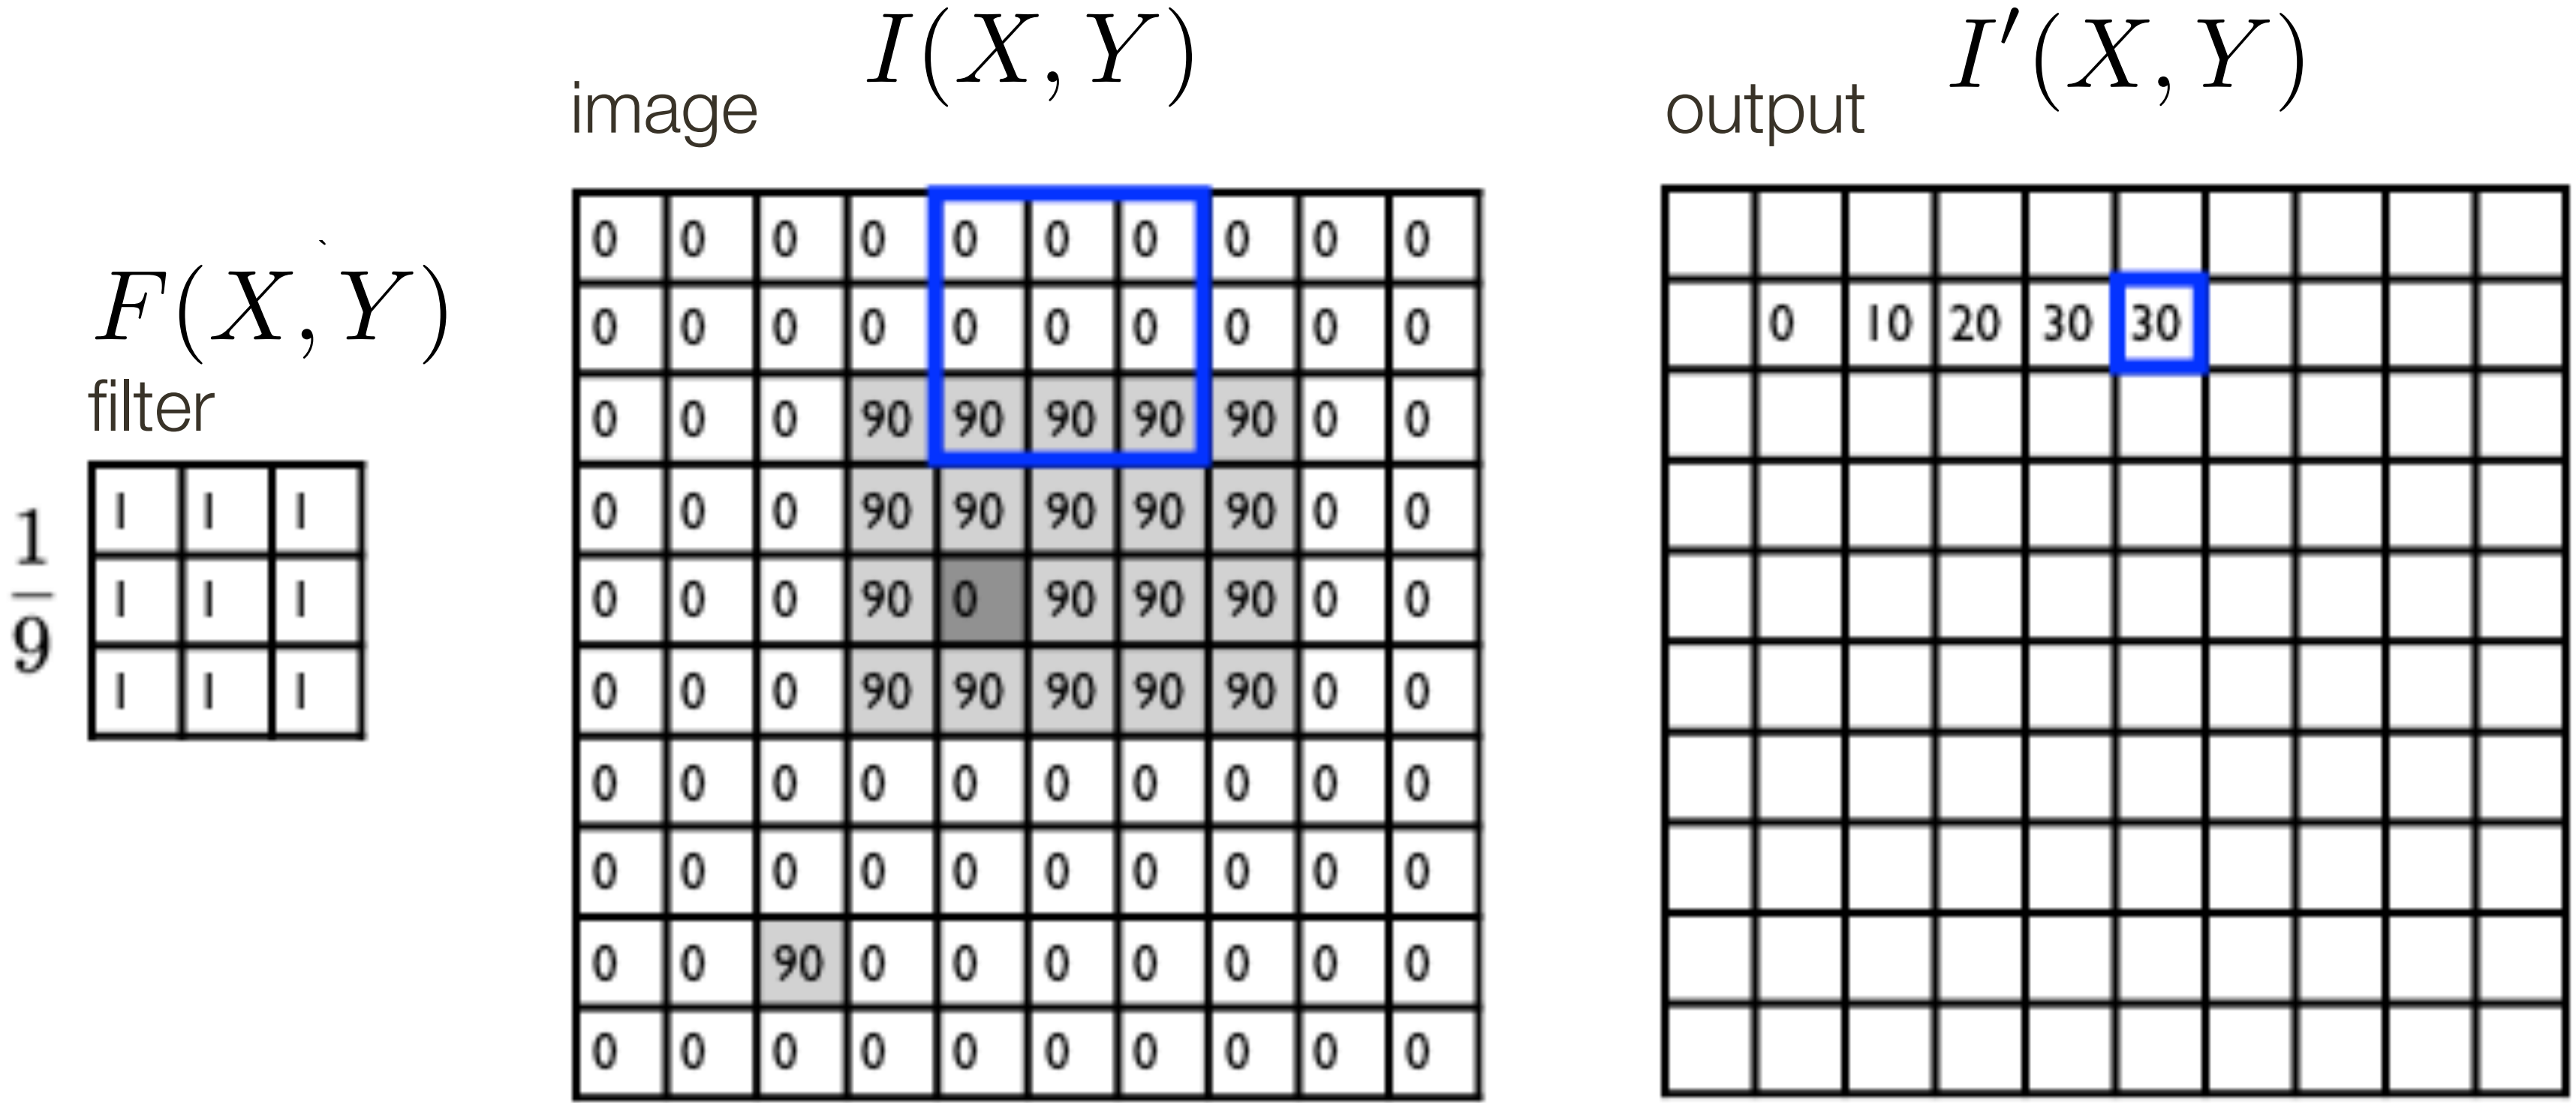
$$\underbrace{I'(X, Y)}_{\text{output}} = \sum_{j=-k}^k \sum_{i=-k}^k \underbrace{F(I, J)}_{\text{filter}} \underbrace{I(X + i, Y + j)}_{\text{image (signal)}}$$

Linear Filter Example



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

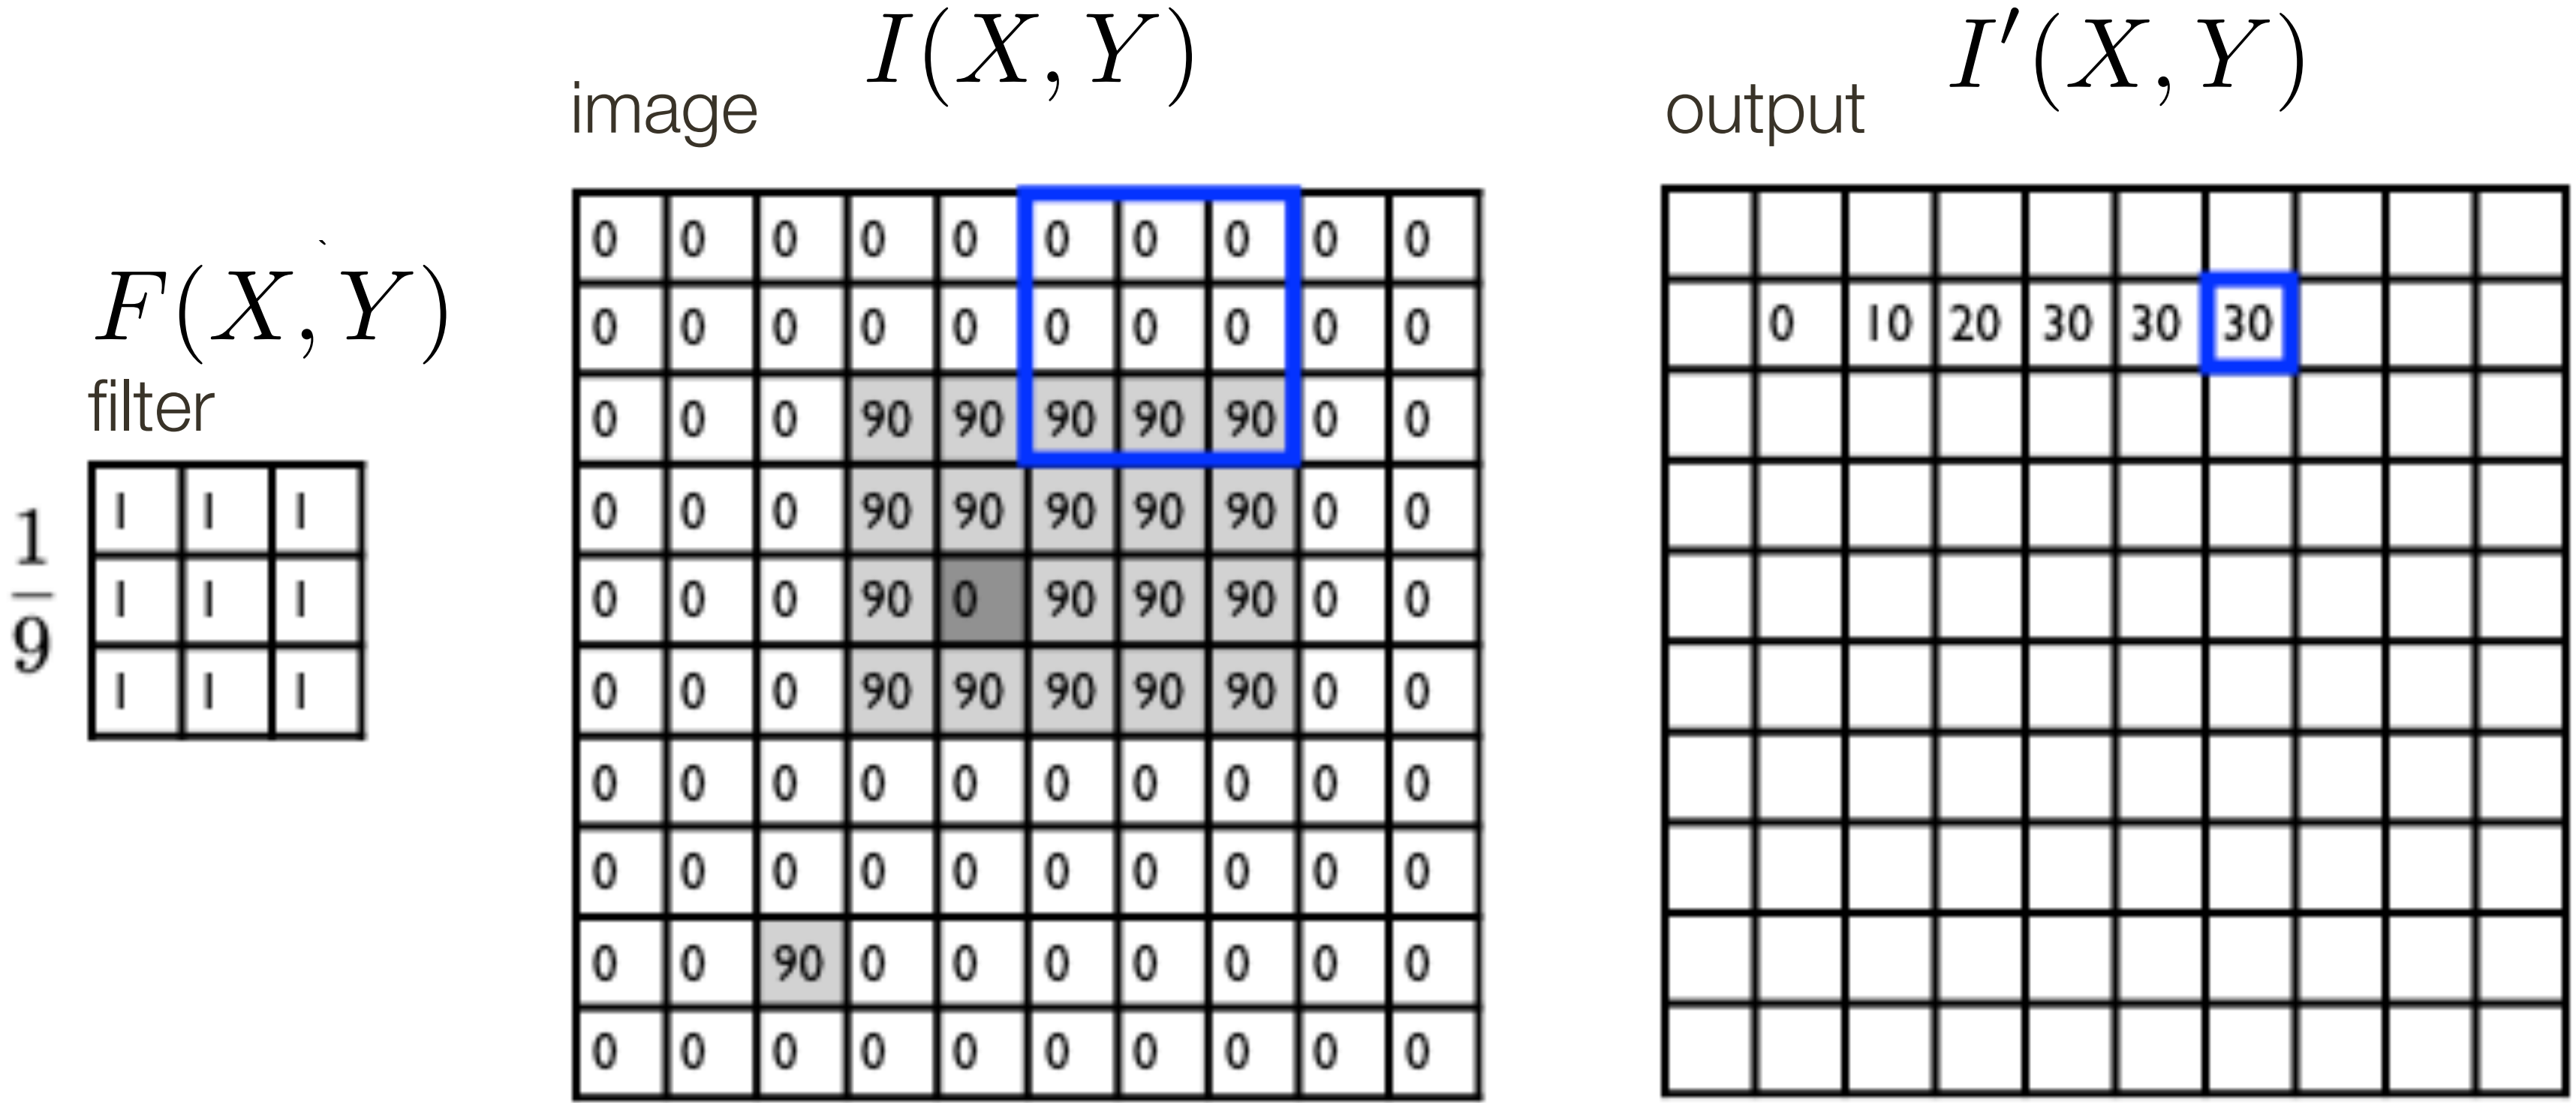
Linear Filter Example



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

output
filter
image (signal)

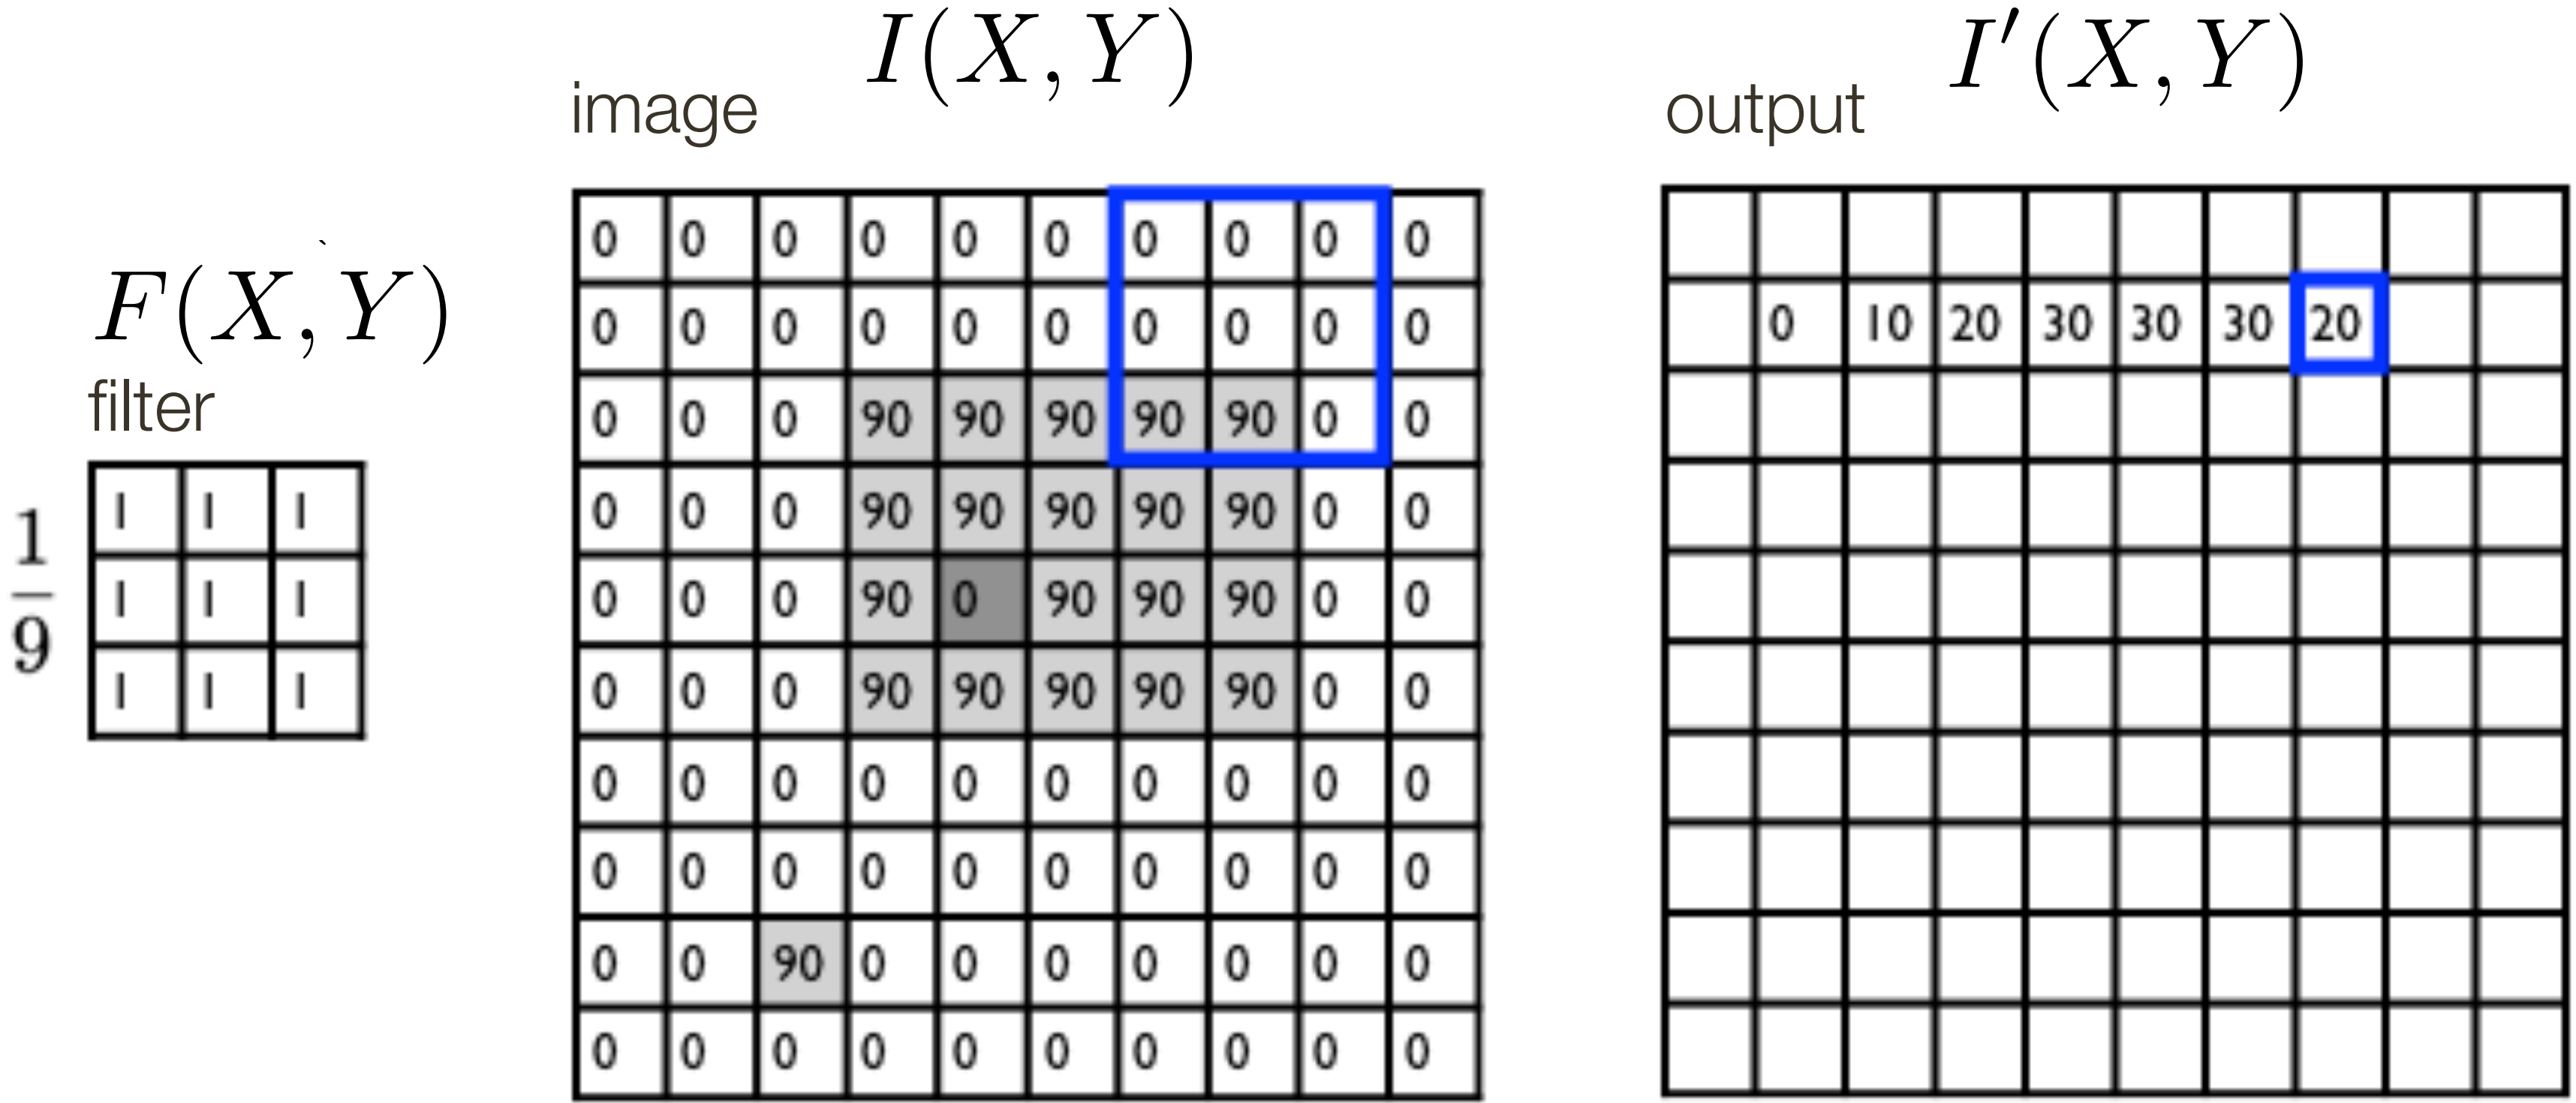
Linear Filter Example



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

output
filter
image (signal)

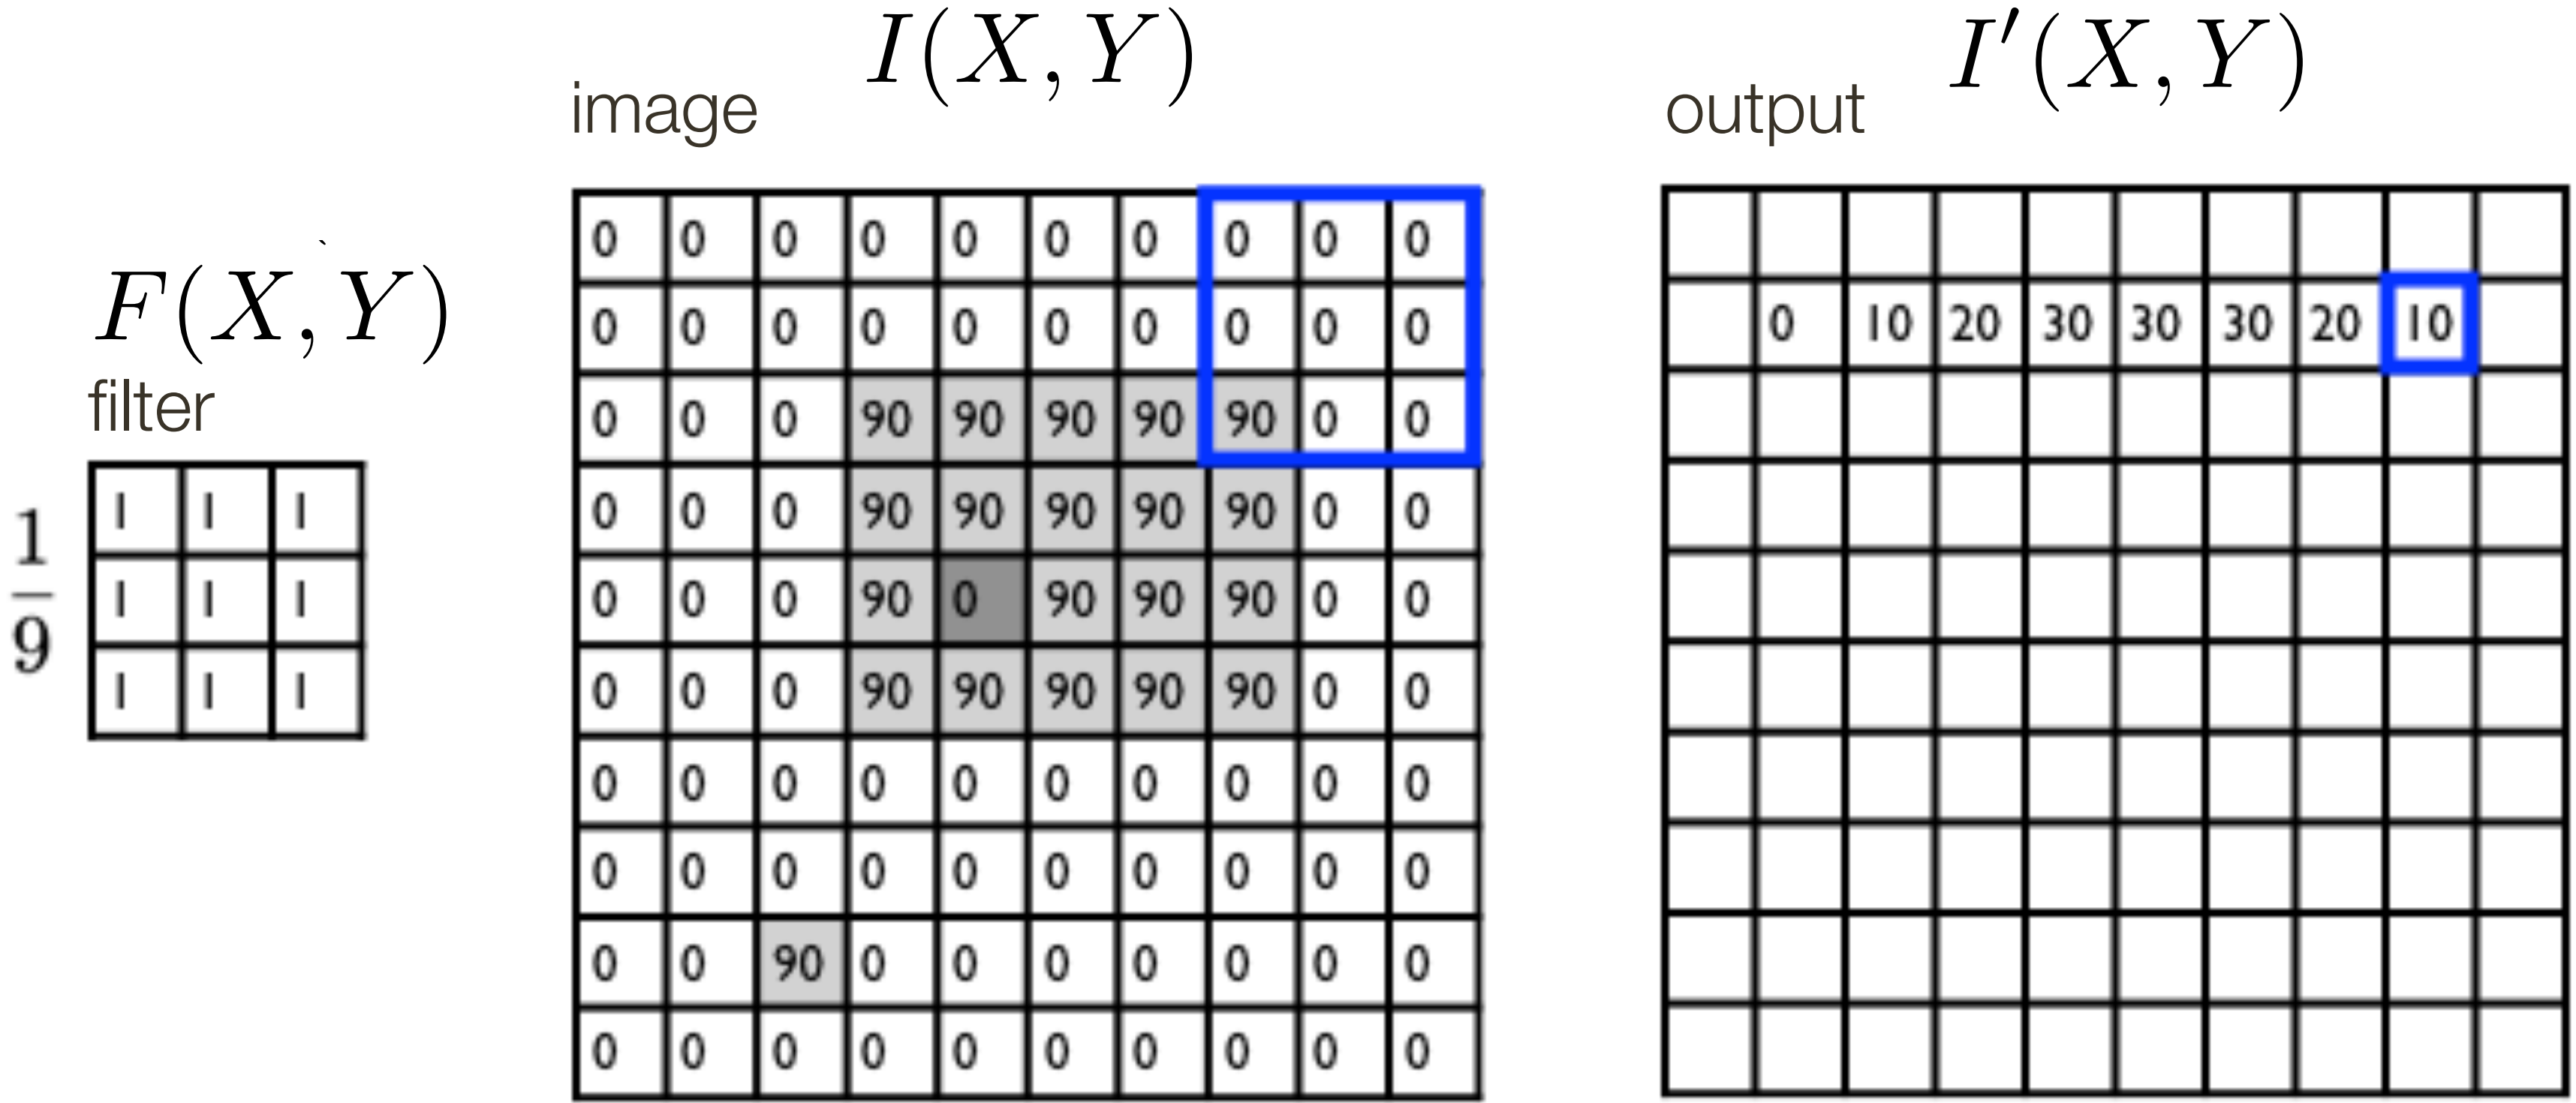
Linear Filter Example



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

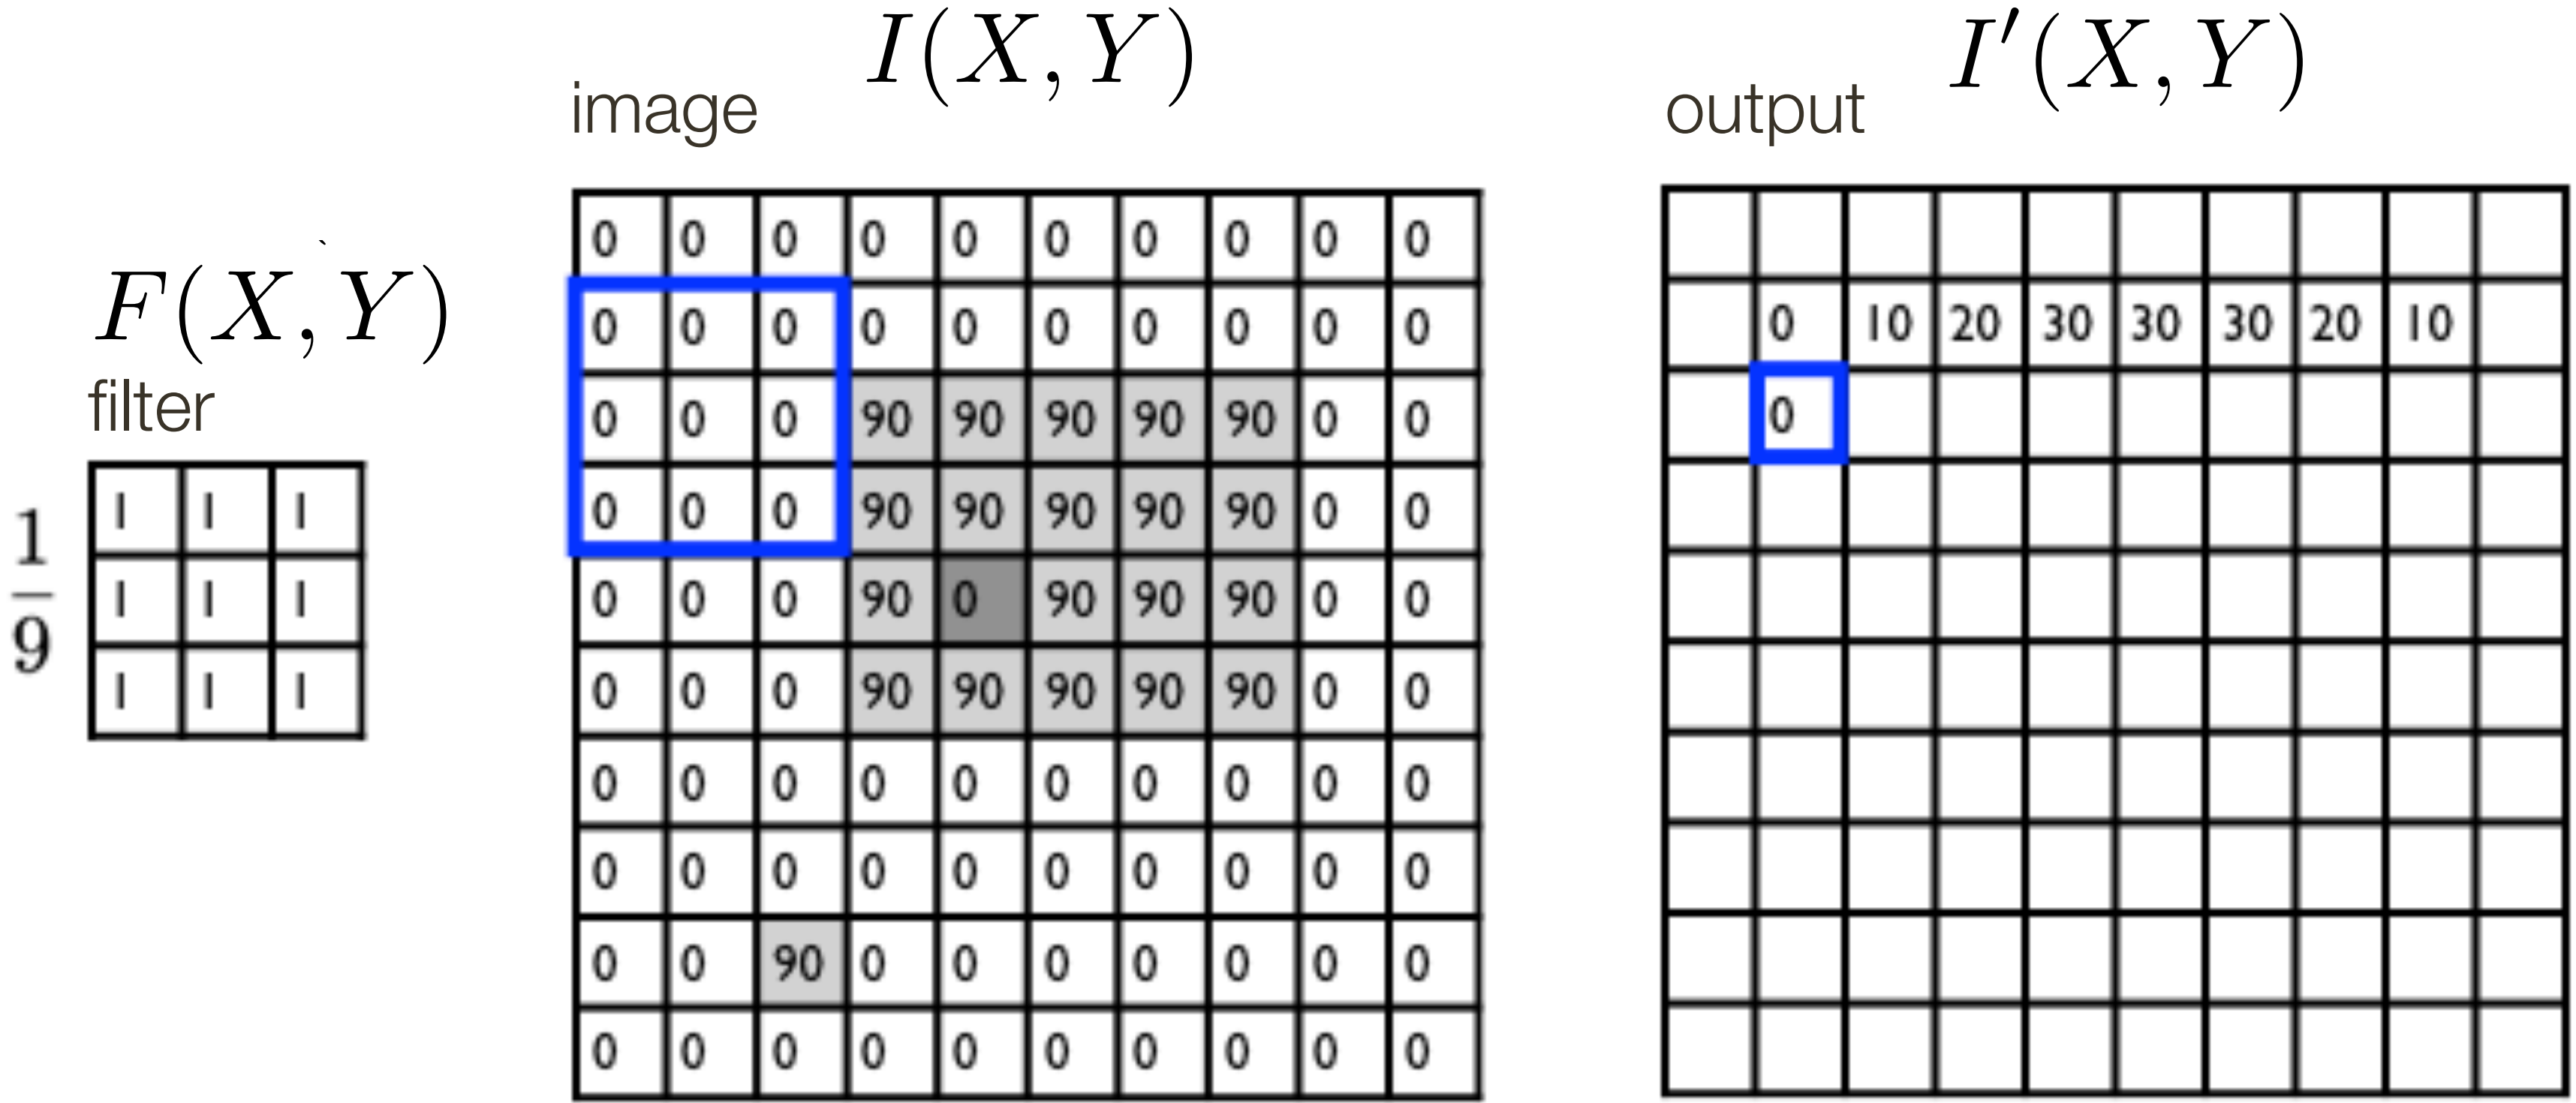
output
filter
image (signal)

Linear Filter Example



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

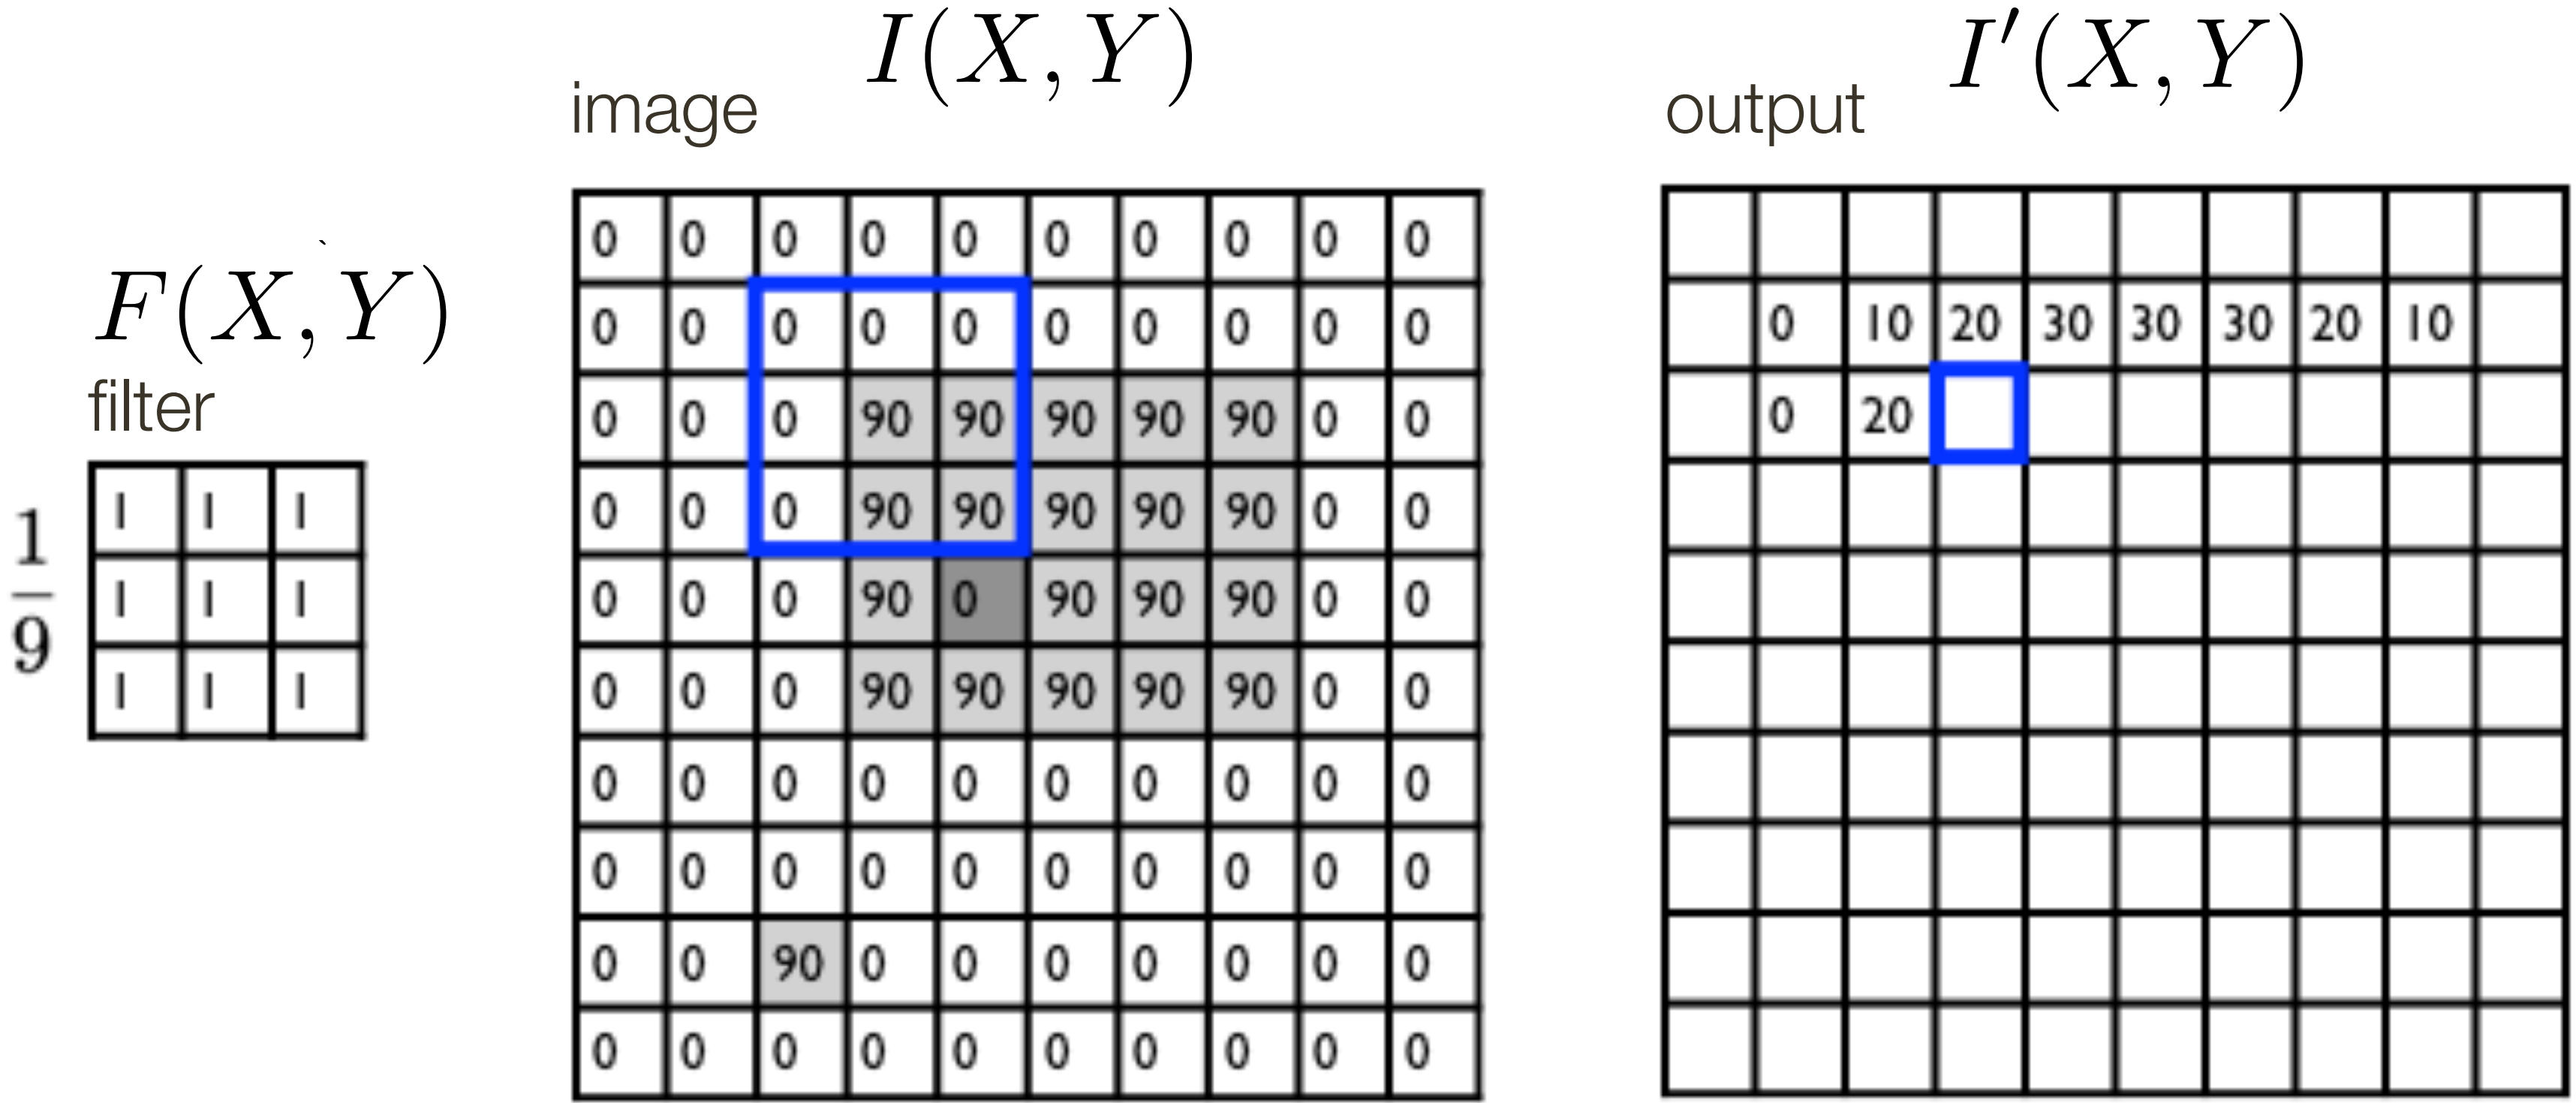
Linear Filter Example



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

output
filter
image (signal)

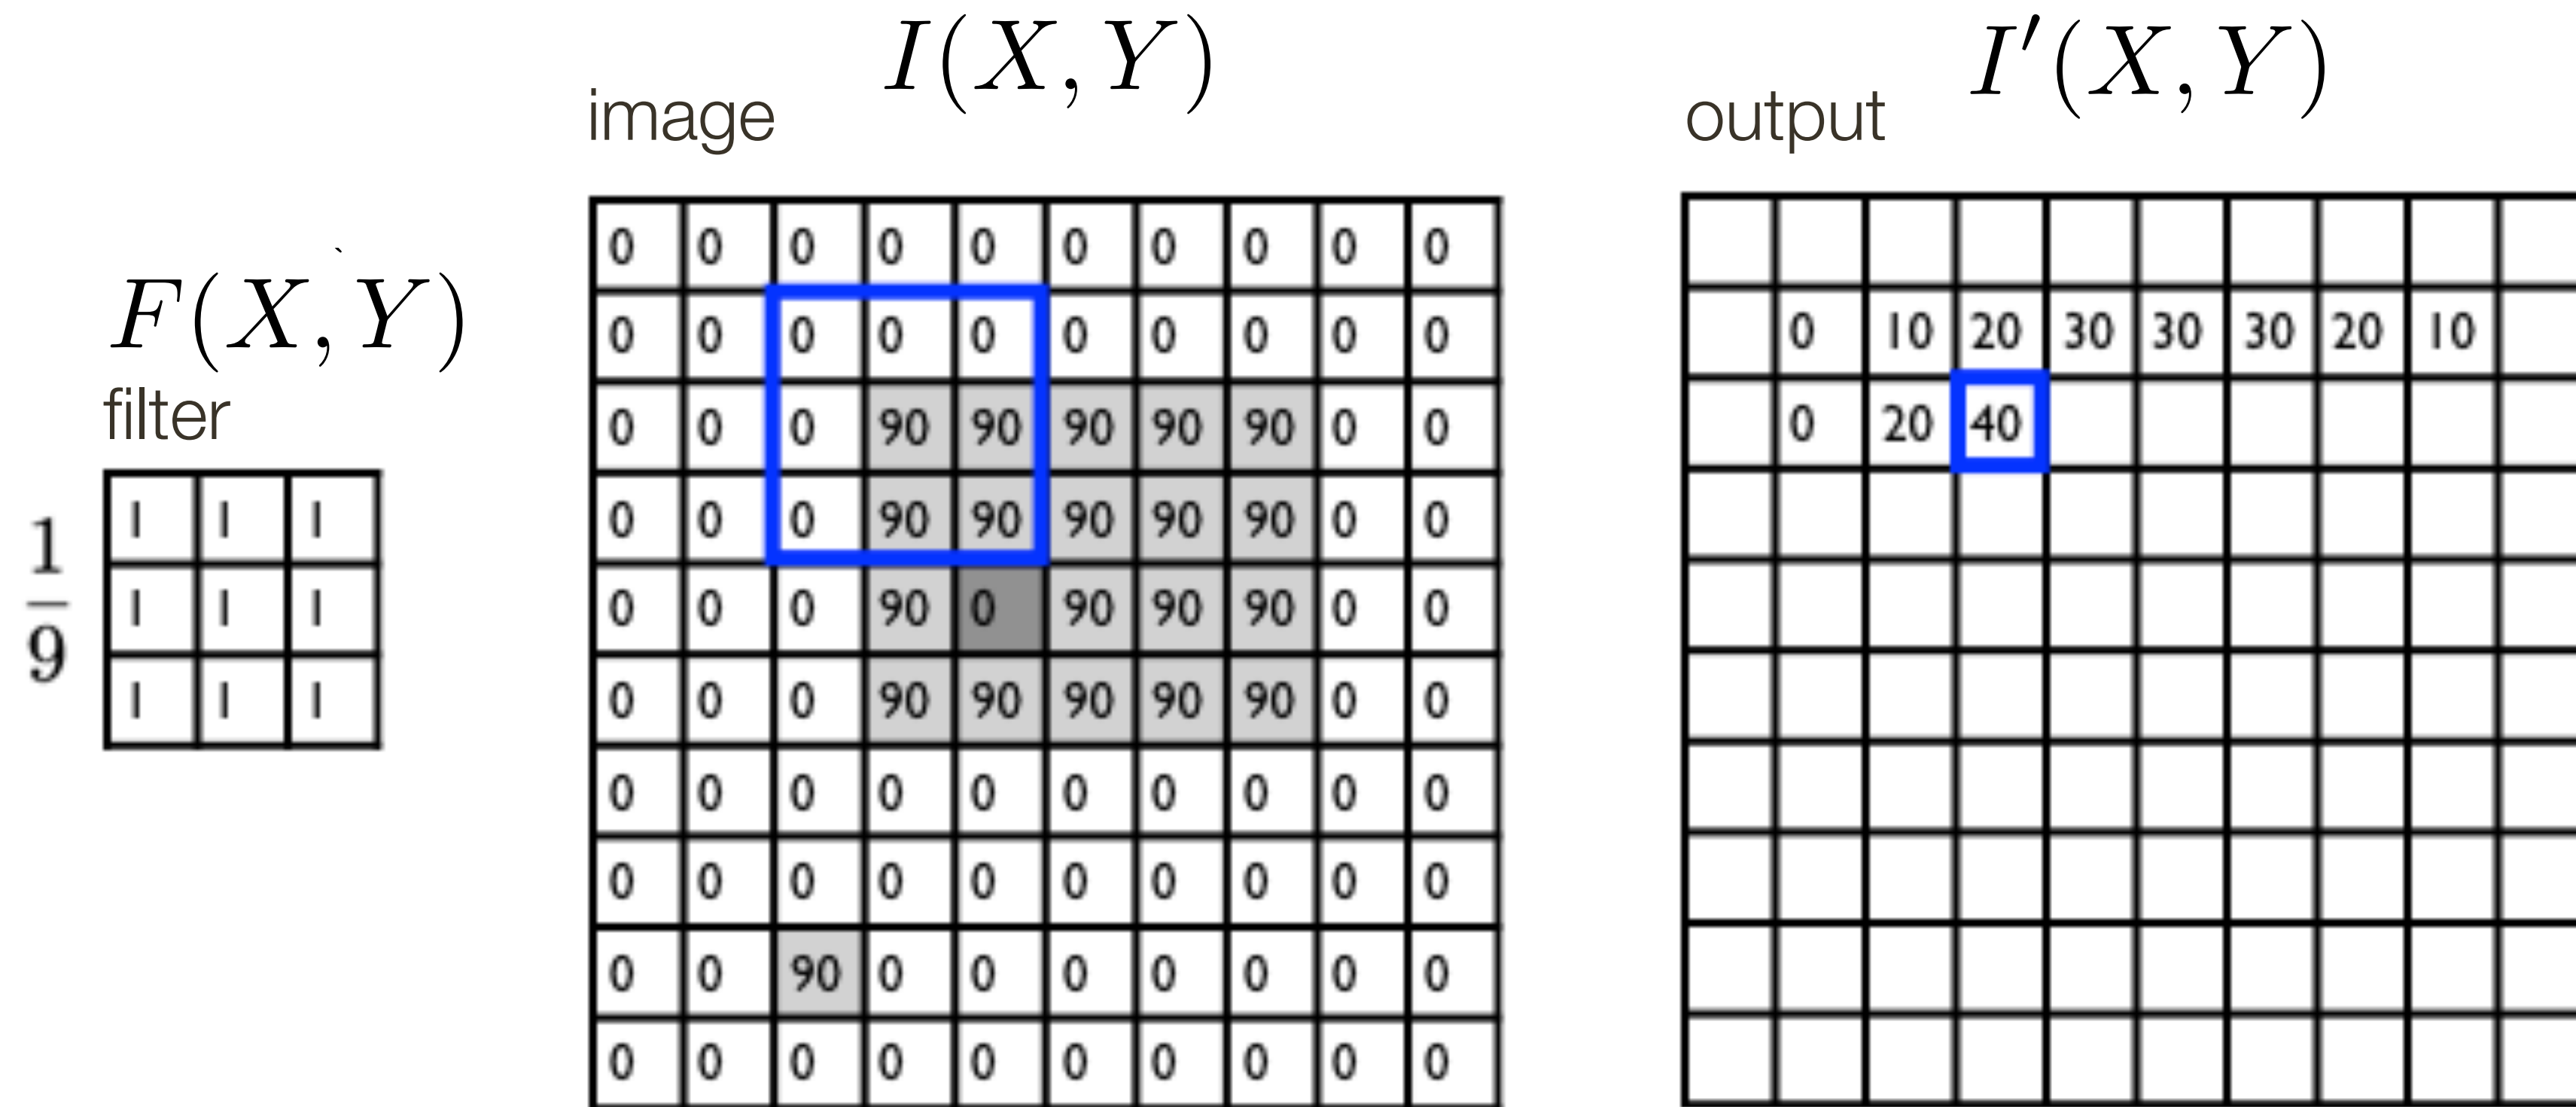
Linear Filter Example



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

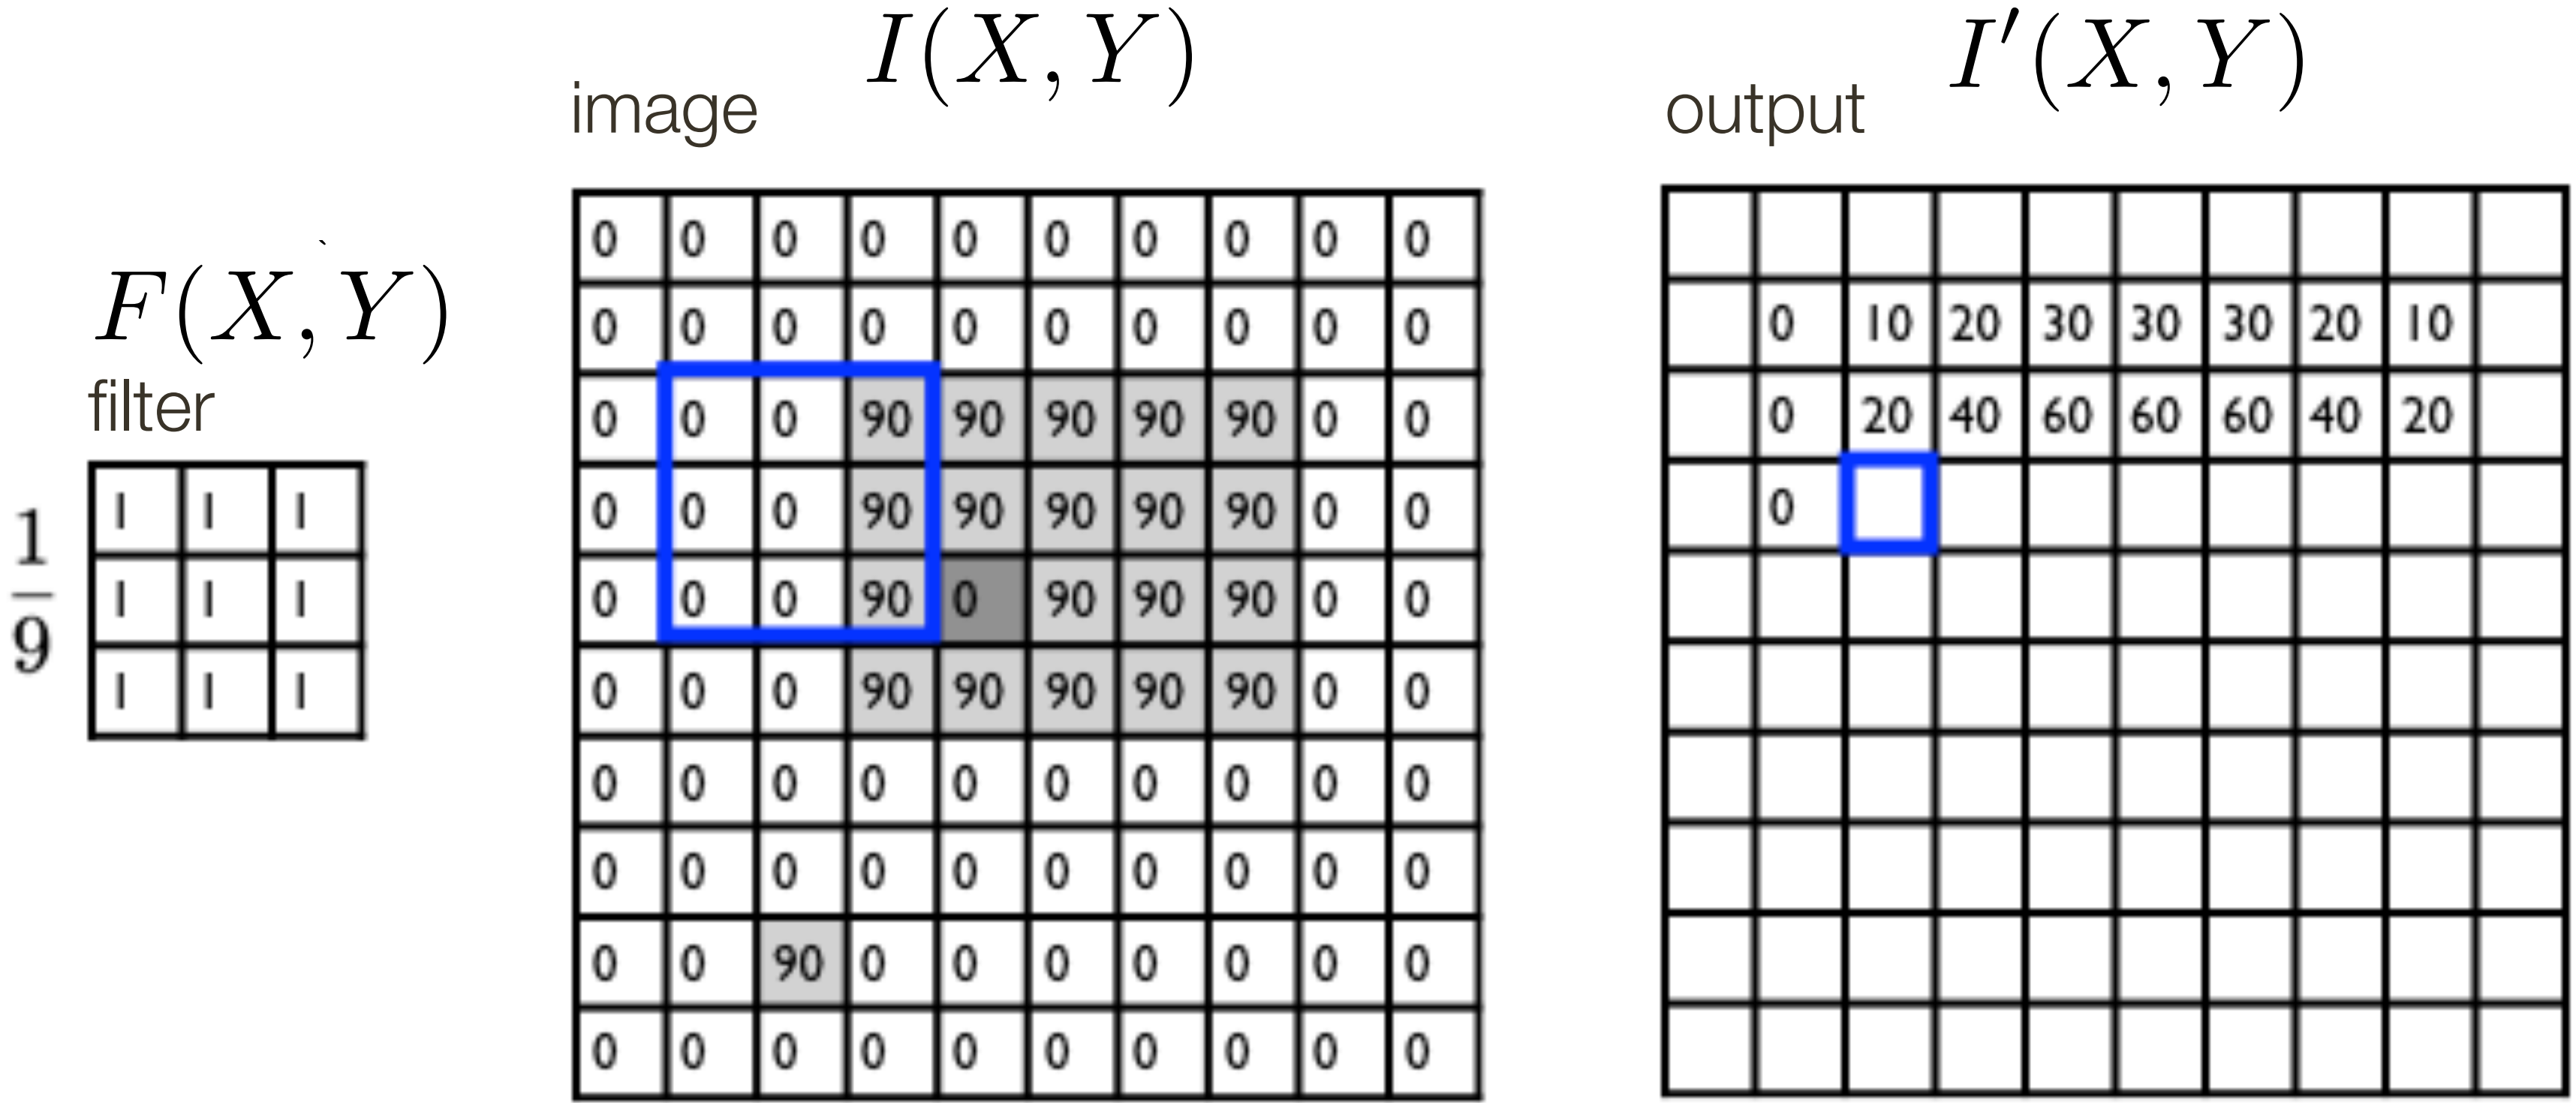
output
filter
image (signal)

Linear Filter Example



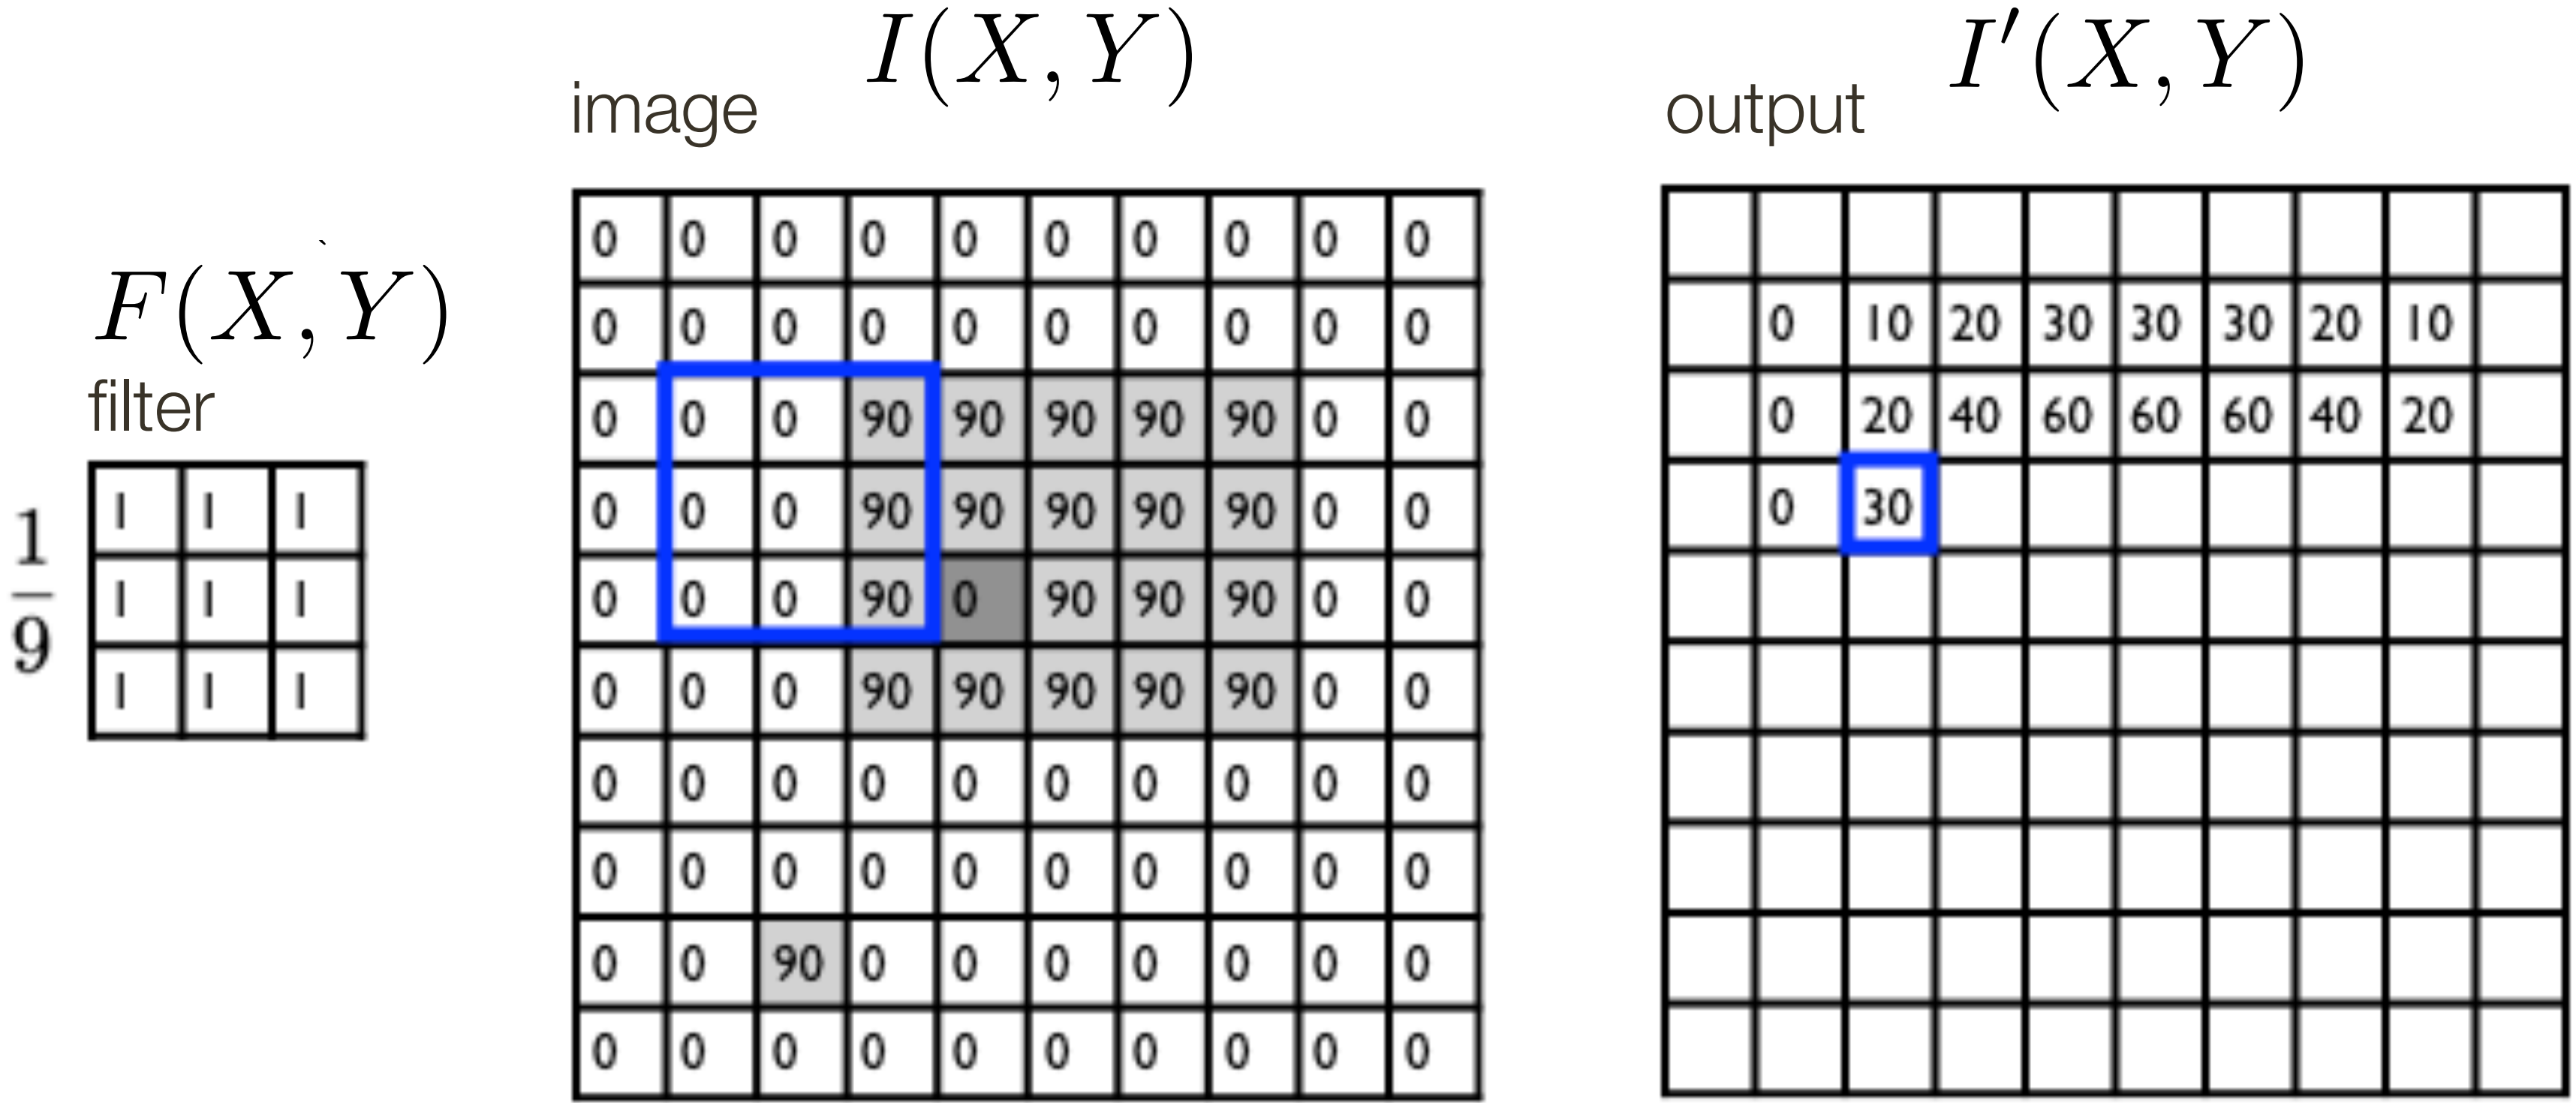
$$I'(X, Y) = \sum_{j=-k}^k \sum_{i=-k}^k F(I, J) I(X + i, Y + j)$$

Linear Filter Example



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

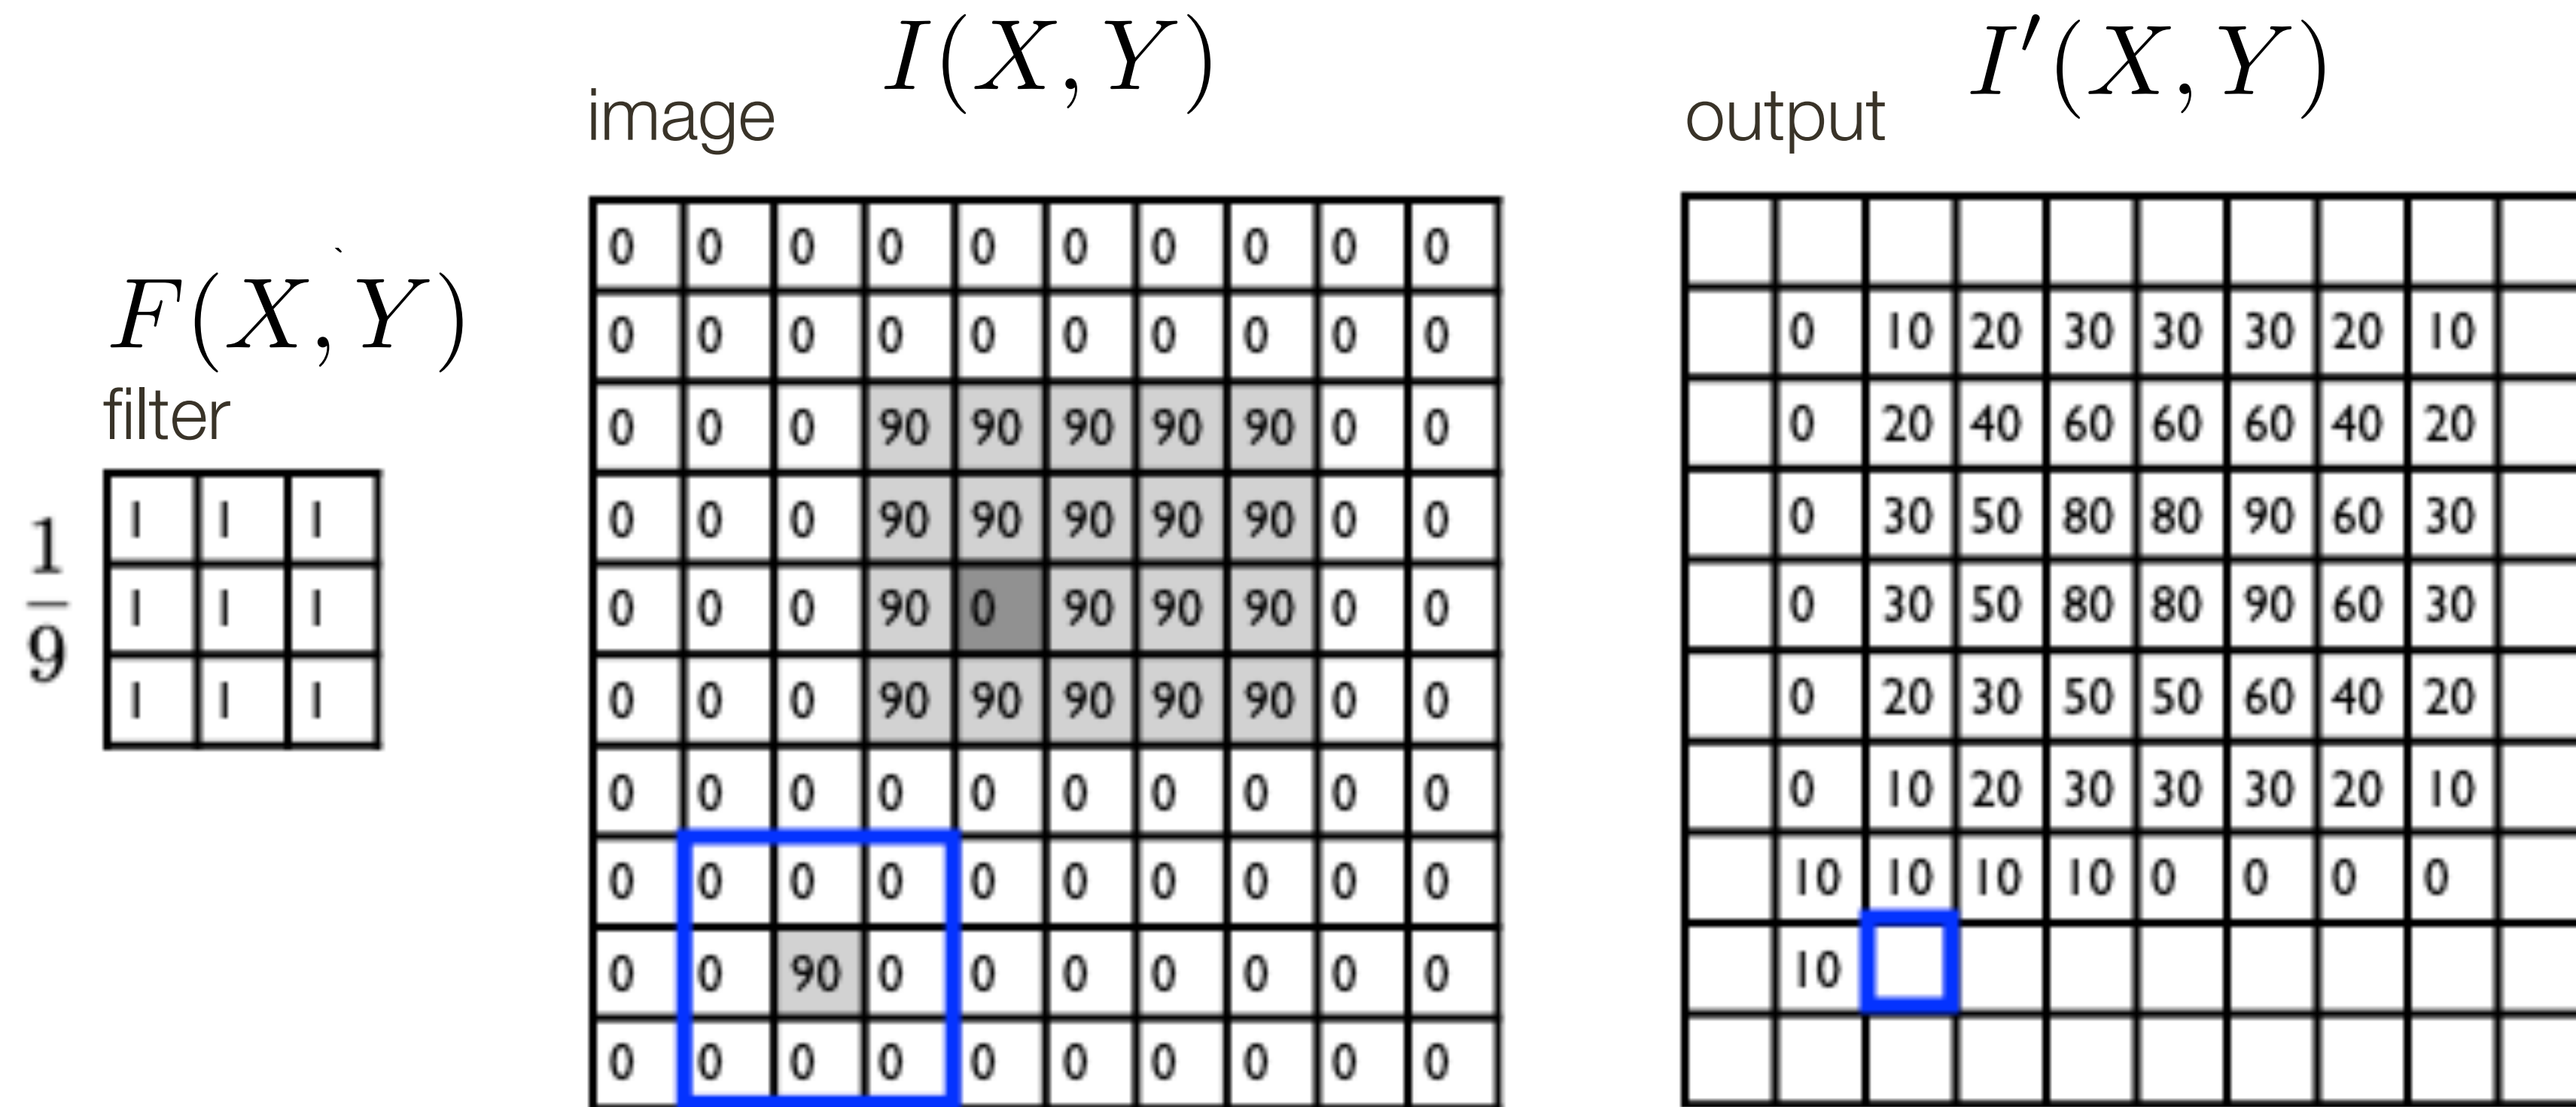
Linear Filter Example



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

output
filter
image (signal)

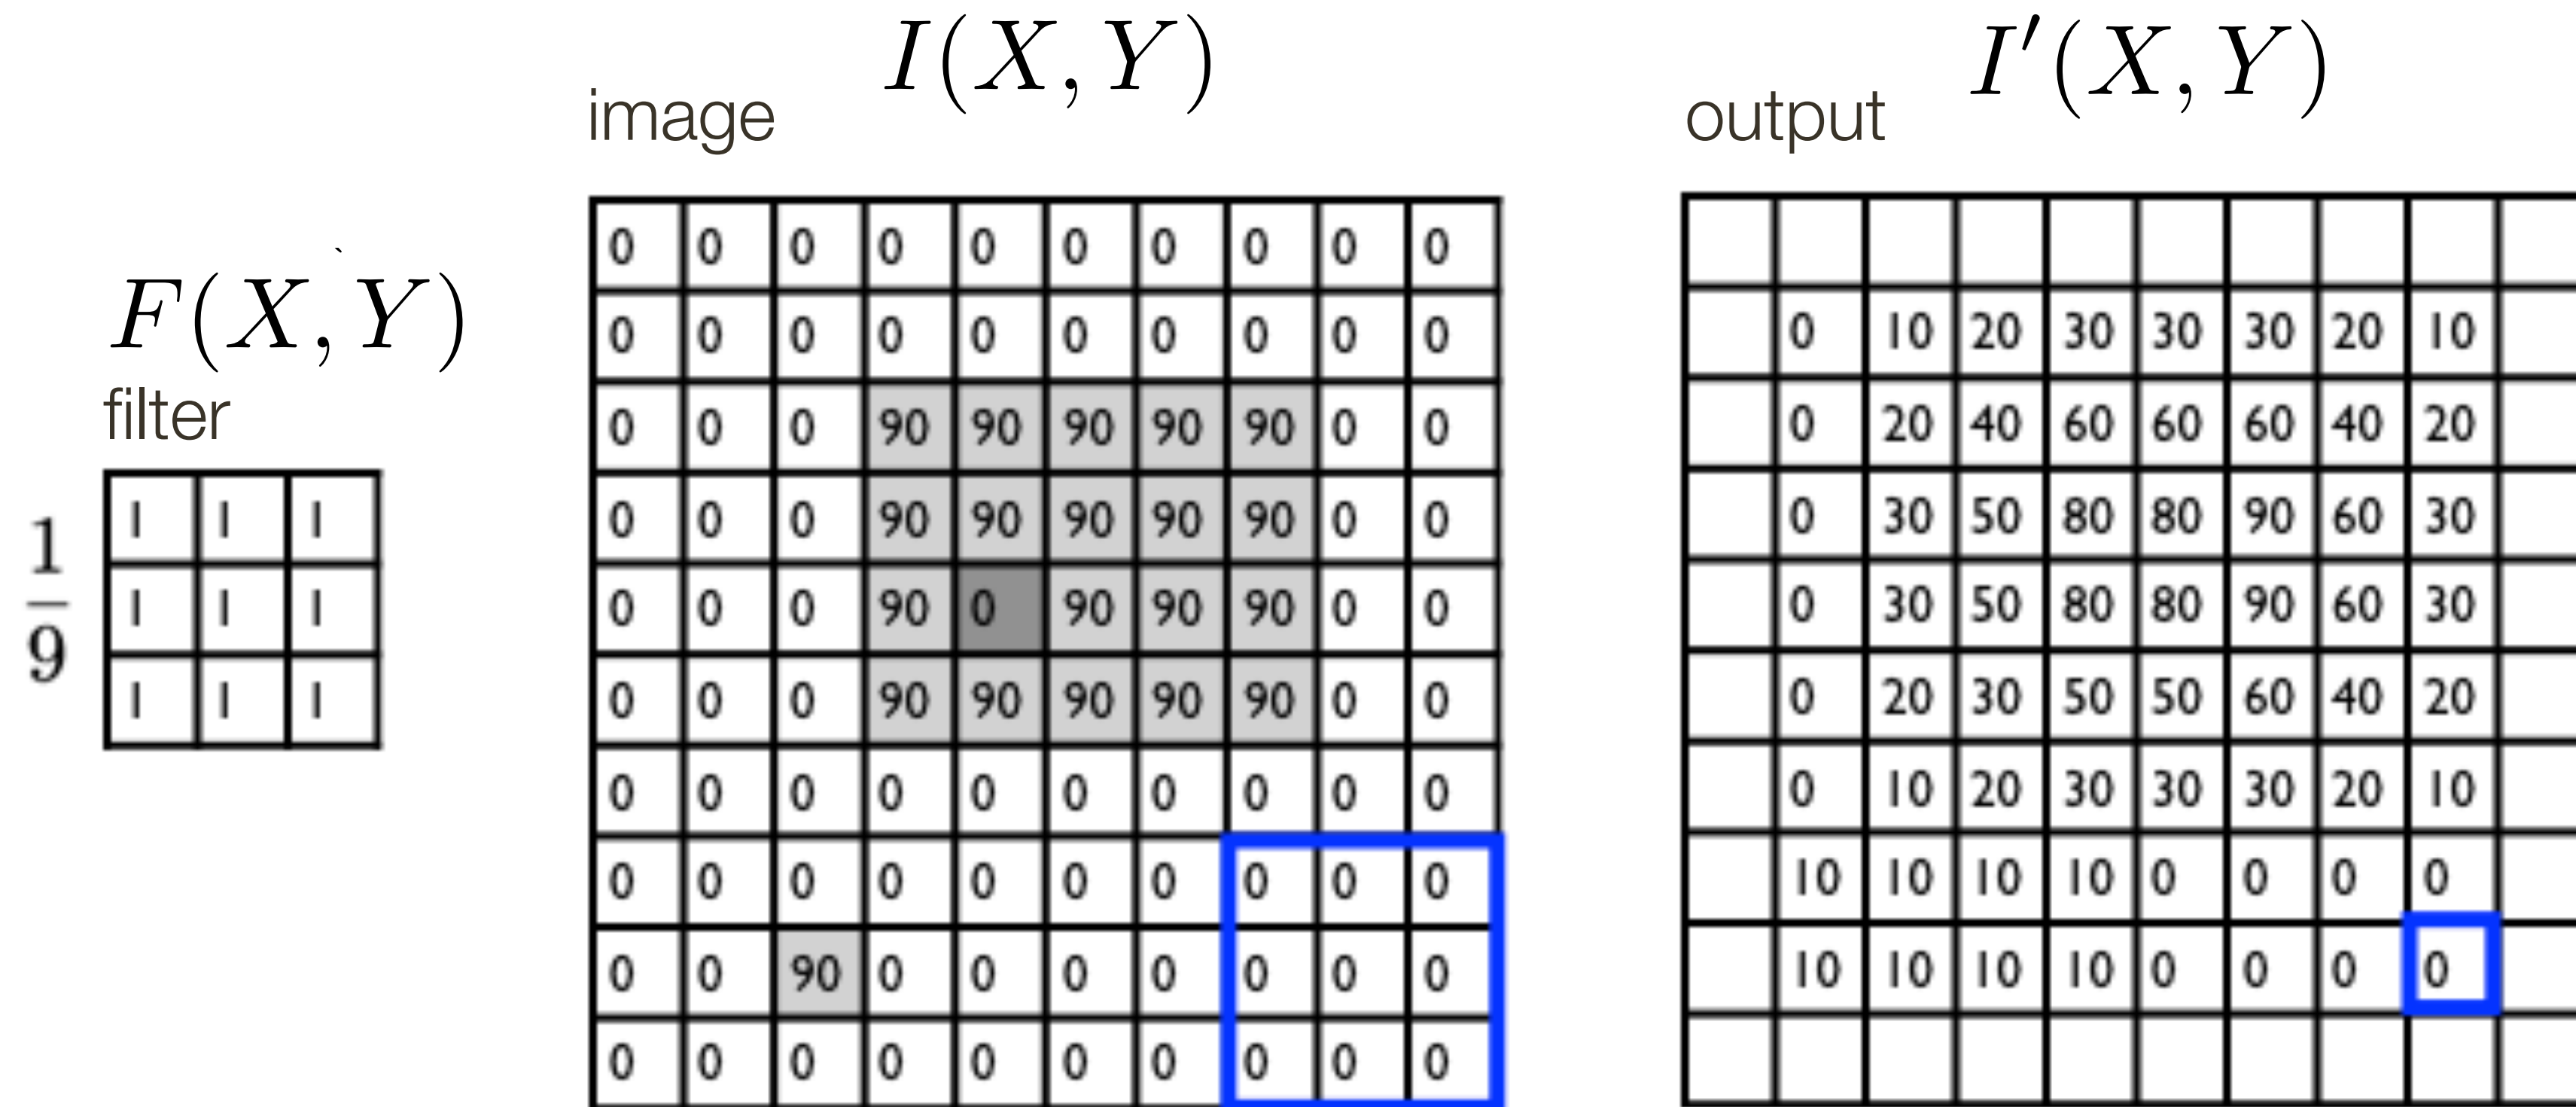
Linear Filter Example



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

output
 filter
 image (signal)

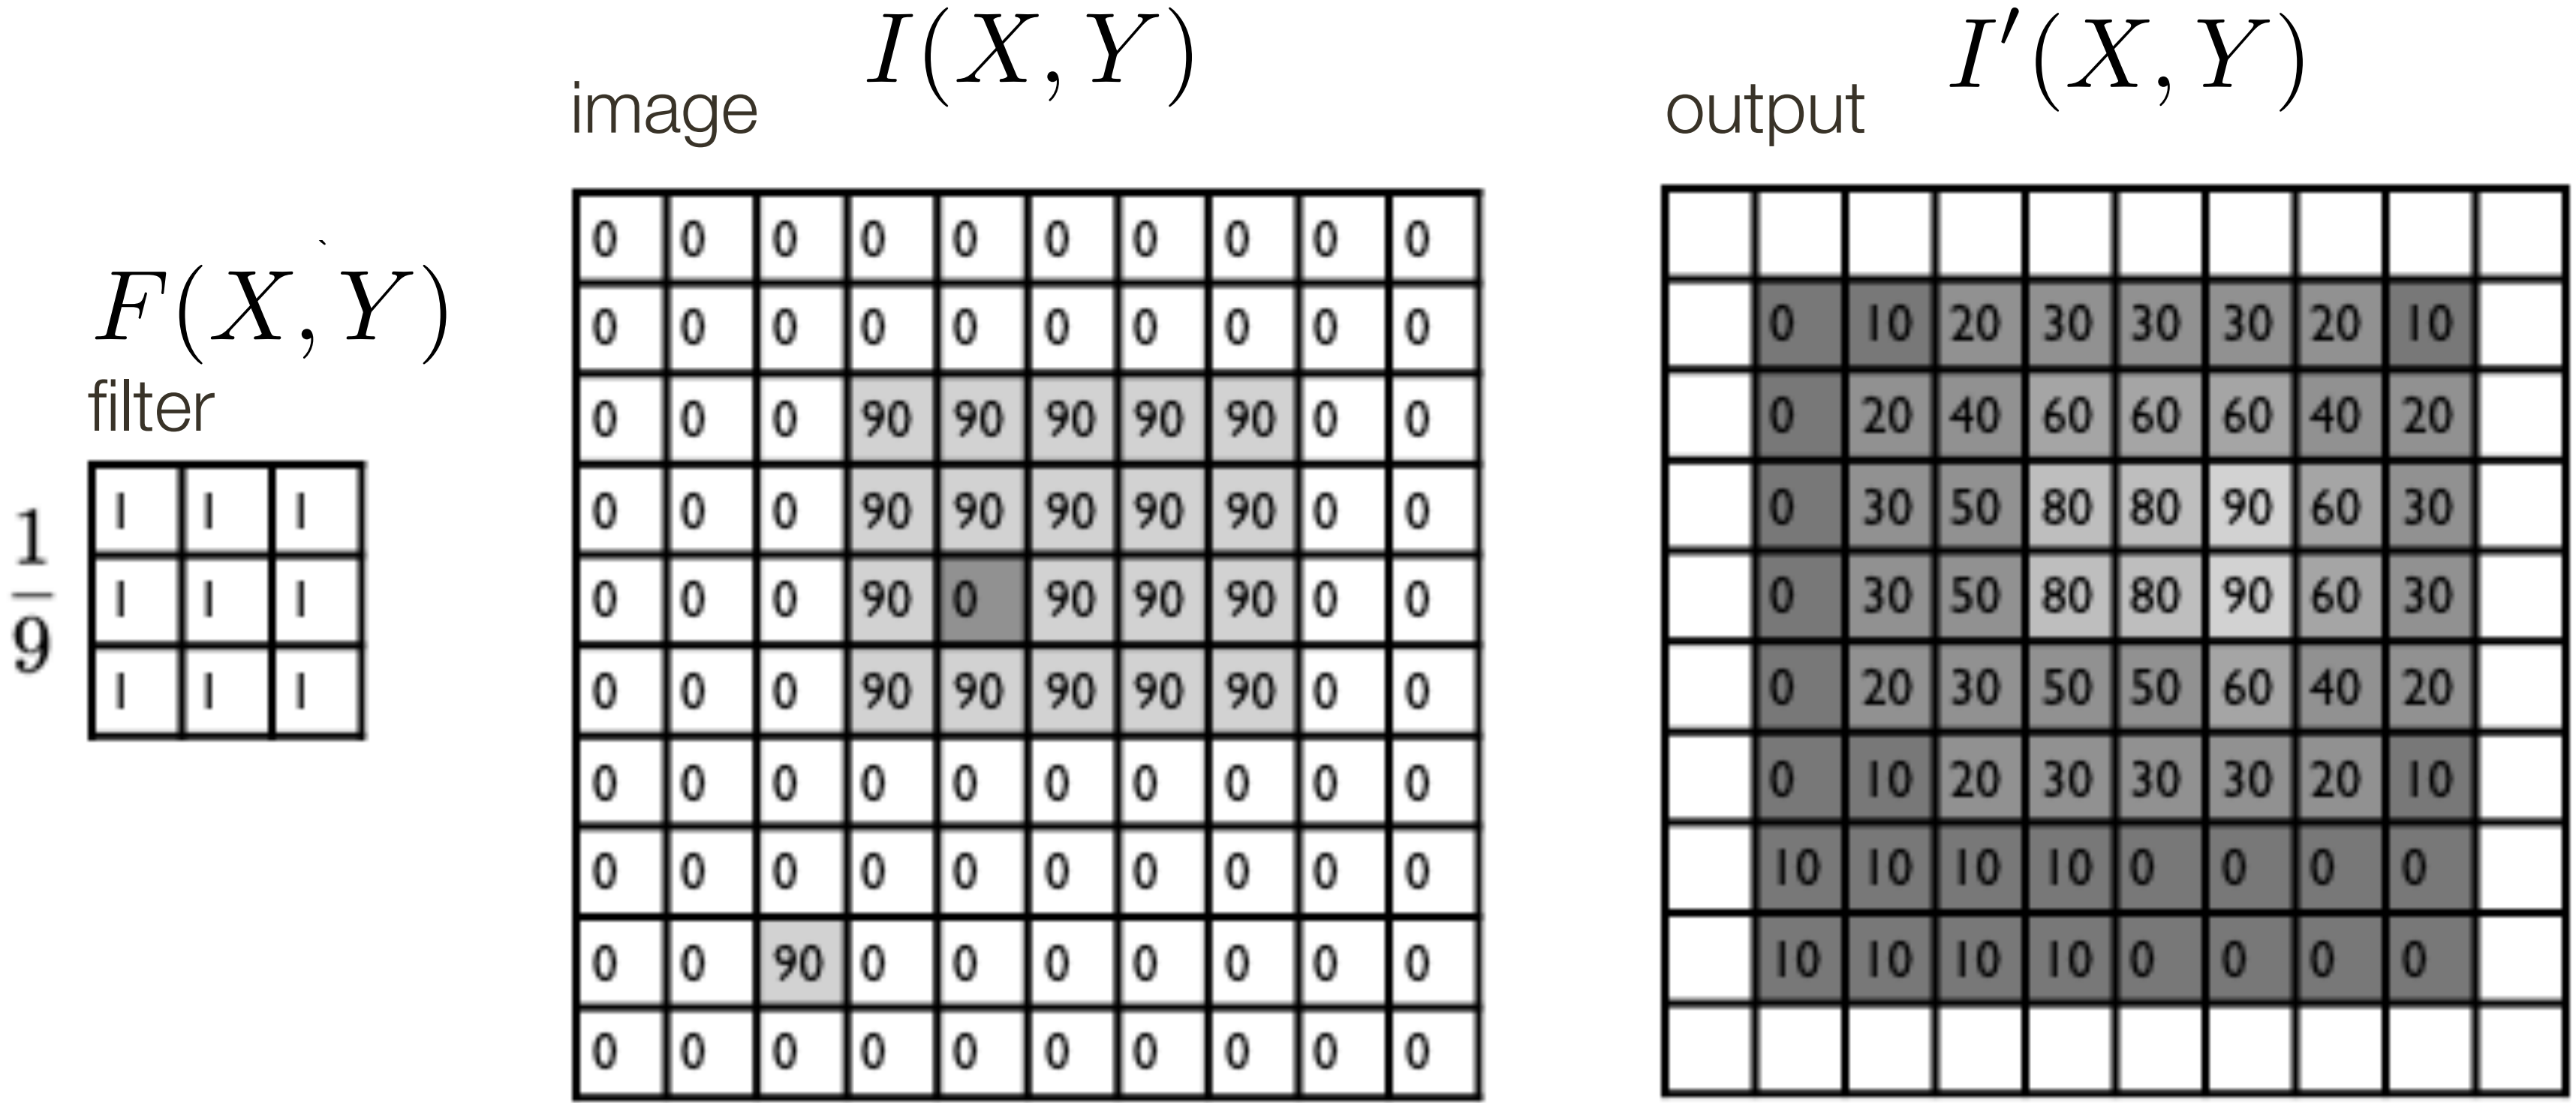
Linear Filter Example



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

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 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

Let's do some accounting ...

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

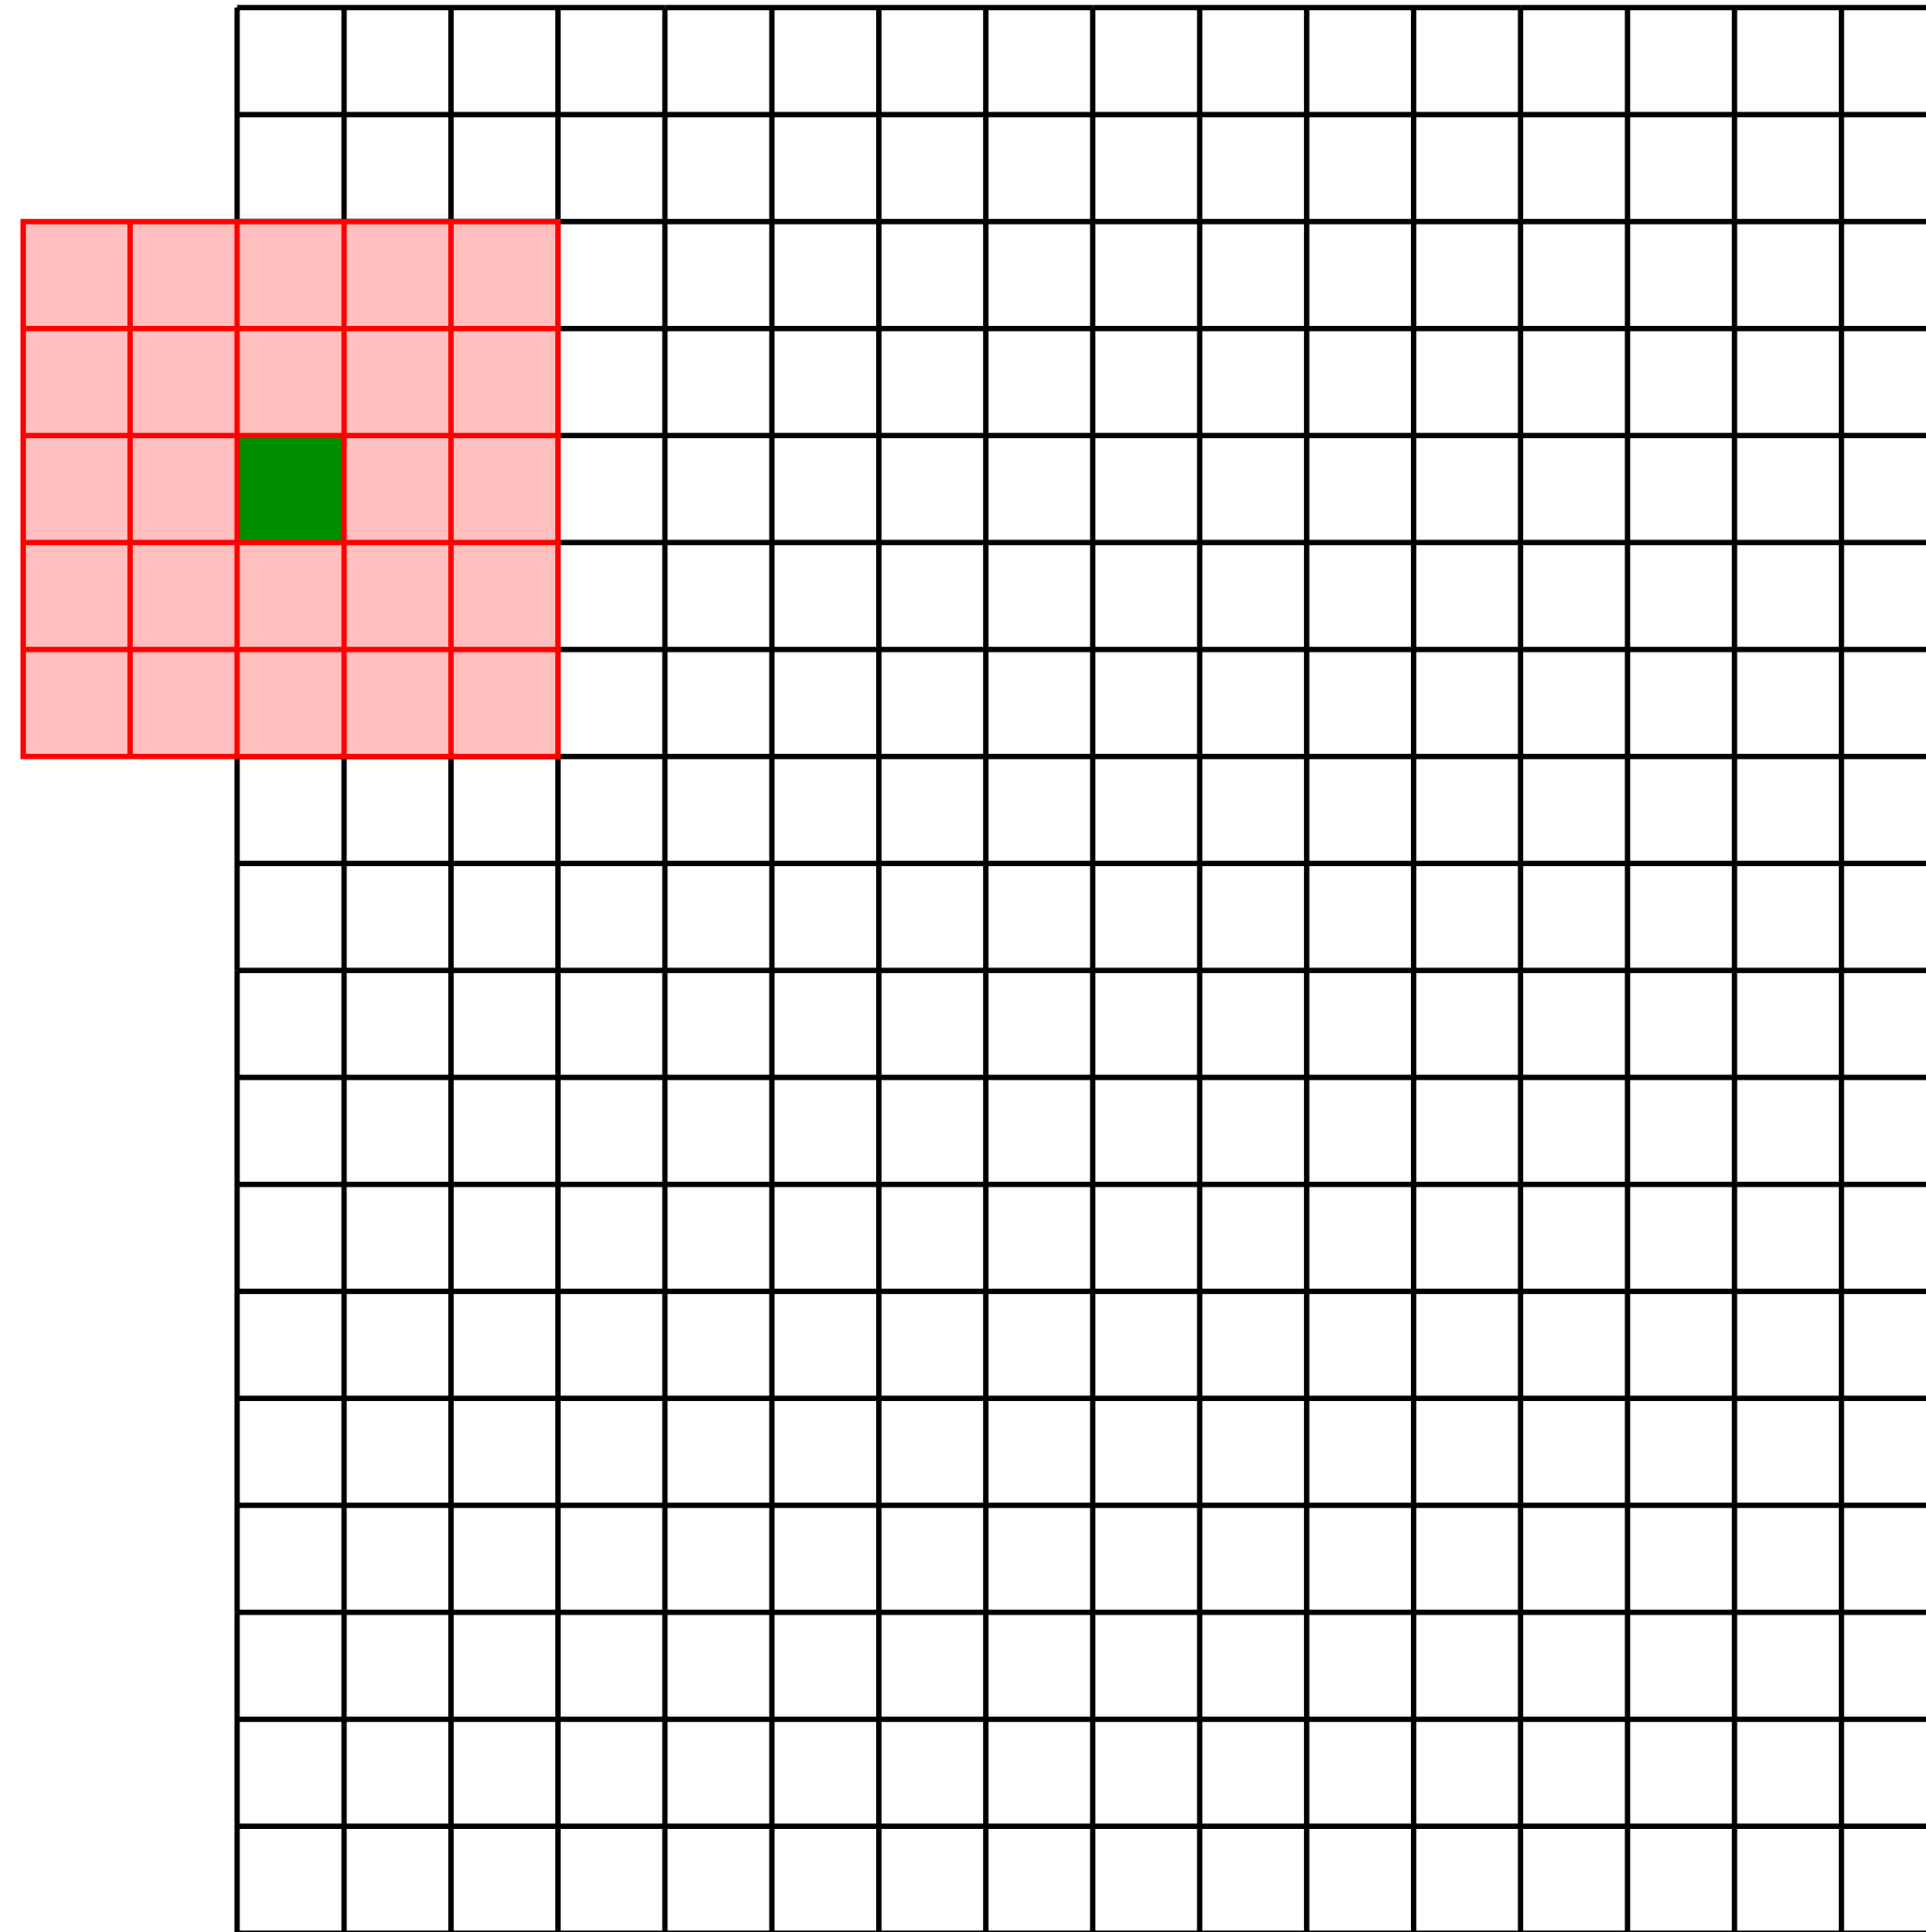
At each pixel, (X, Y) , there are $m \times m$ multiplications

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

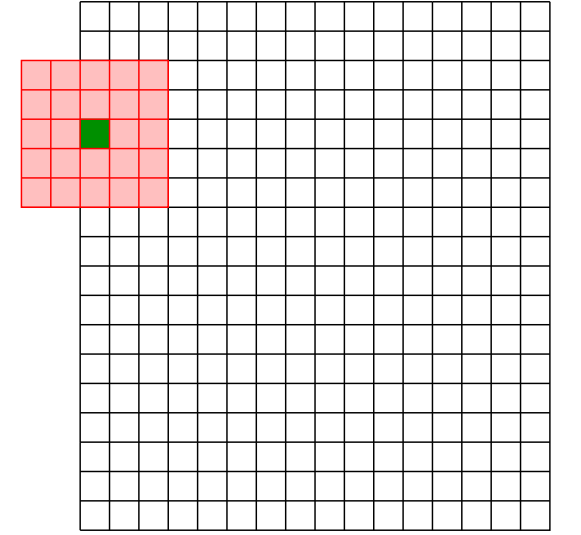
Total: $m^2 \times n^2$ multiplications

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



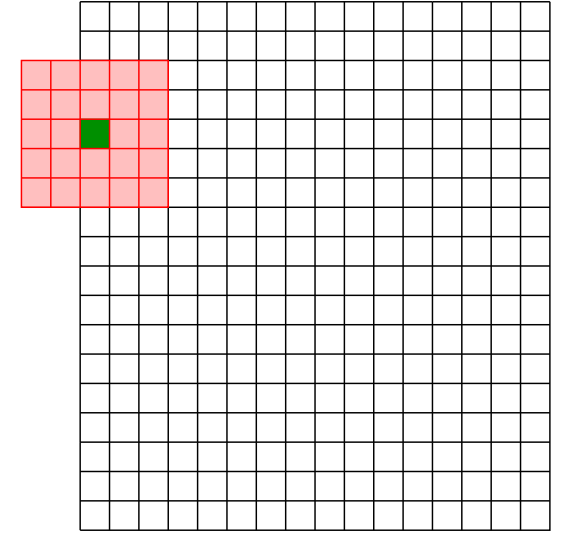
Linear Filters: **Boundary** Effects



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

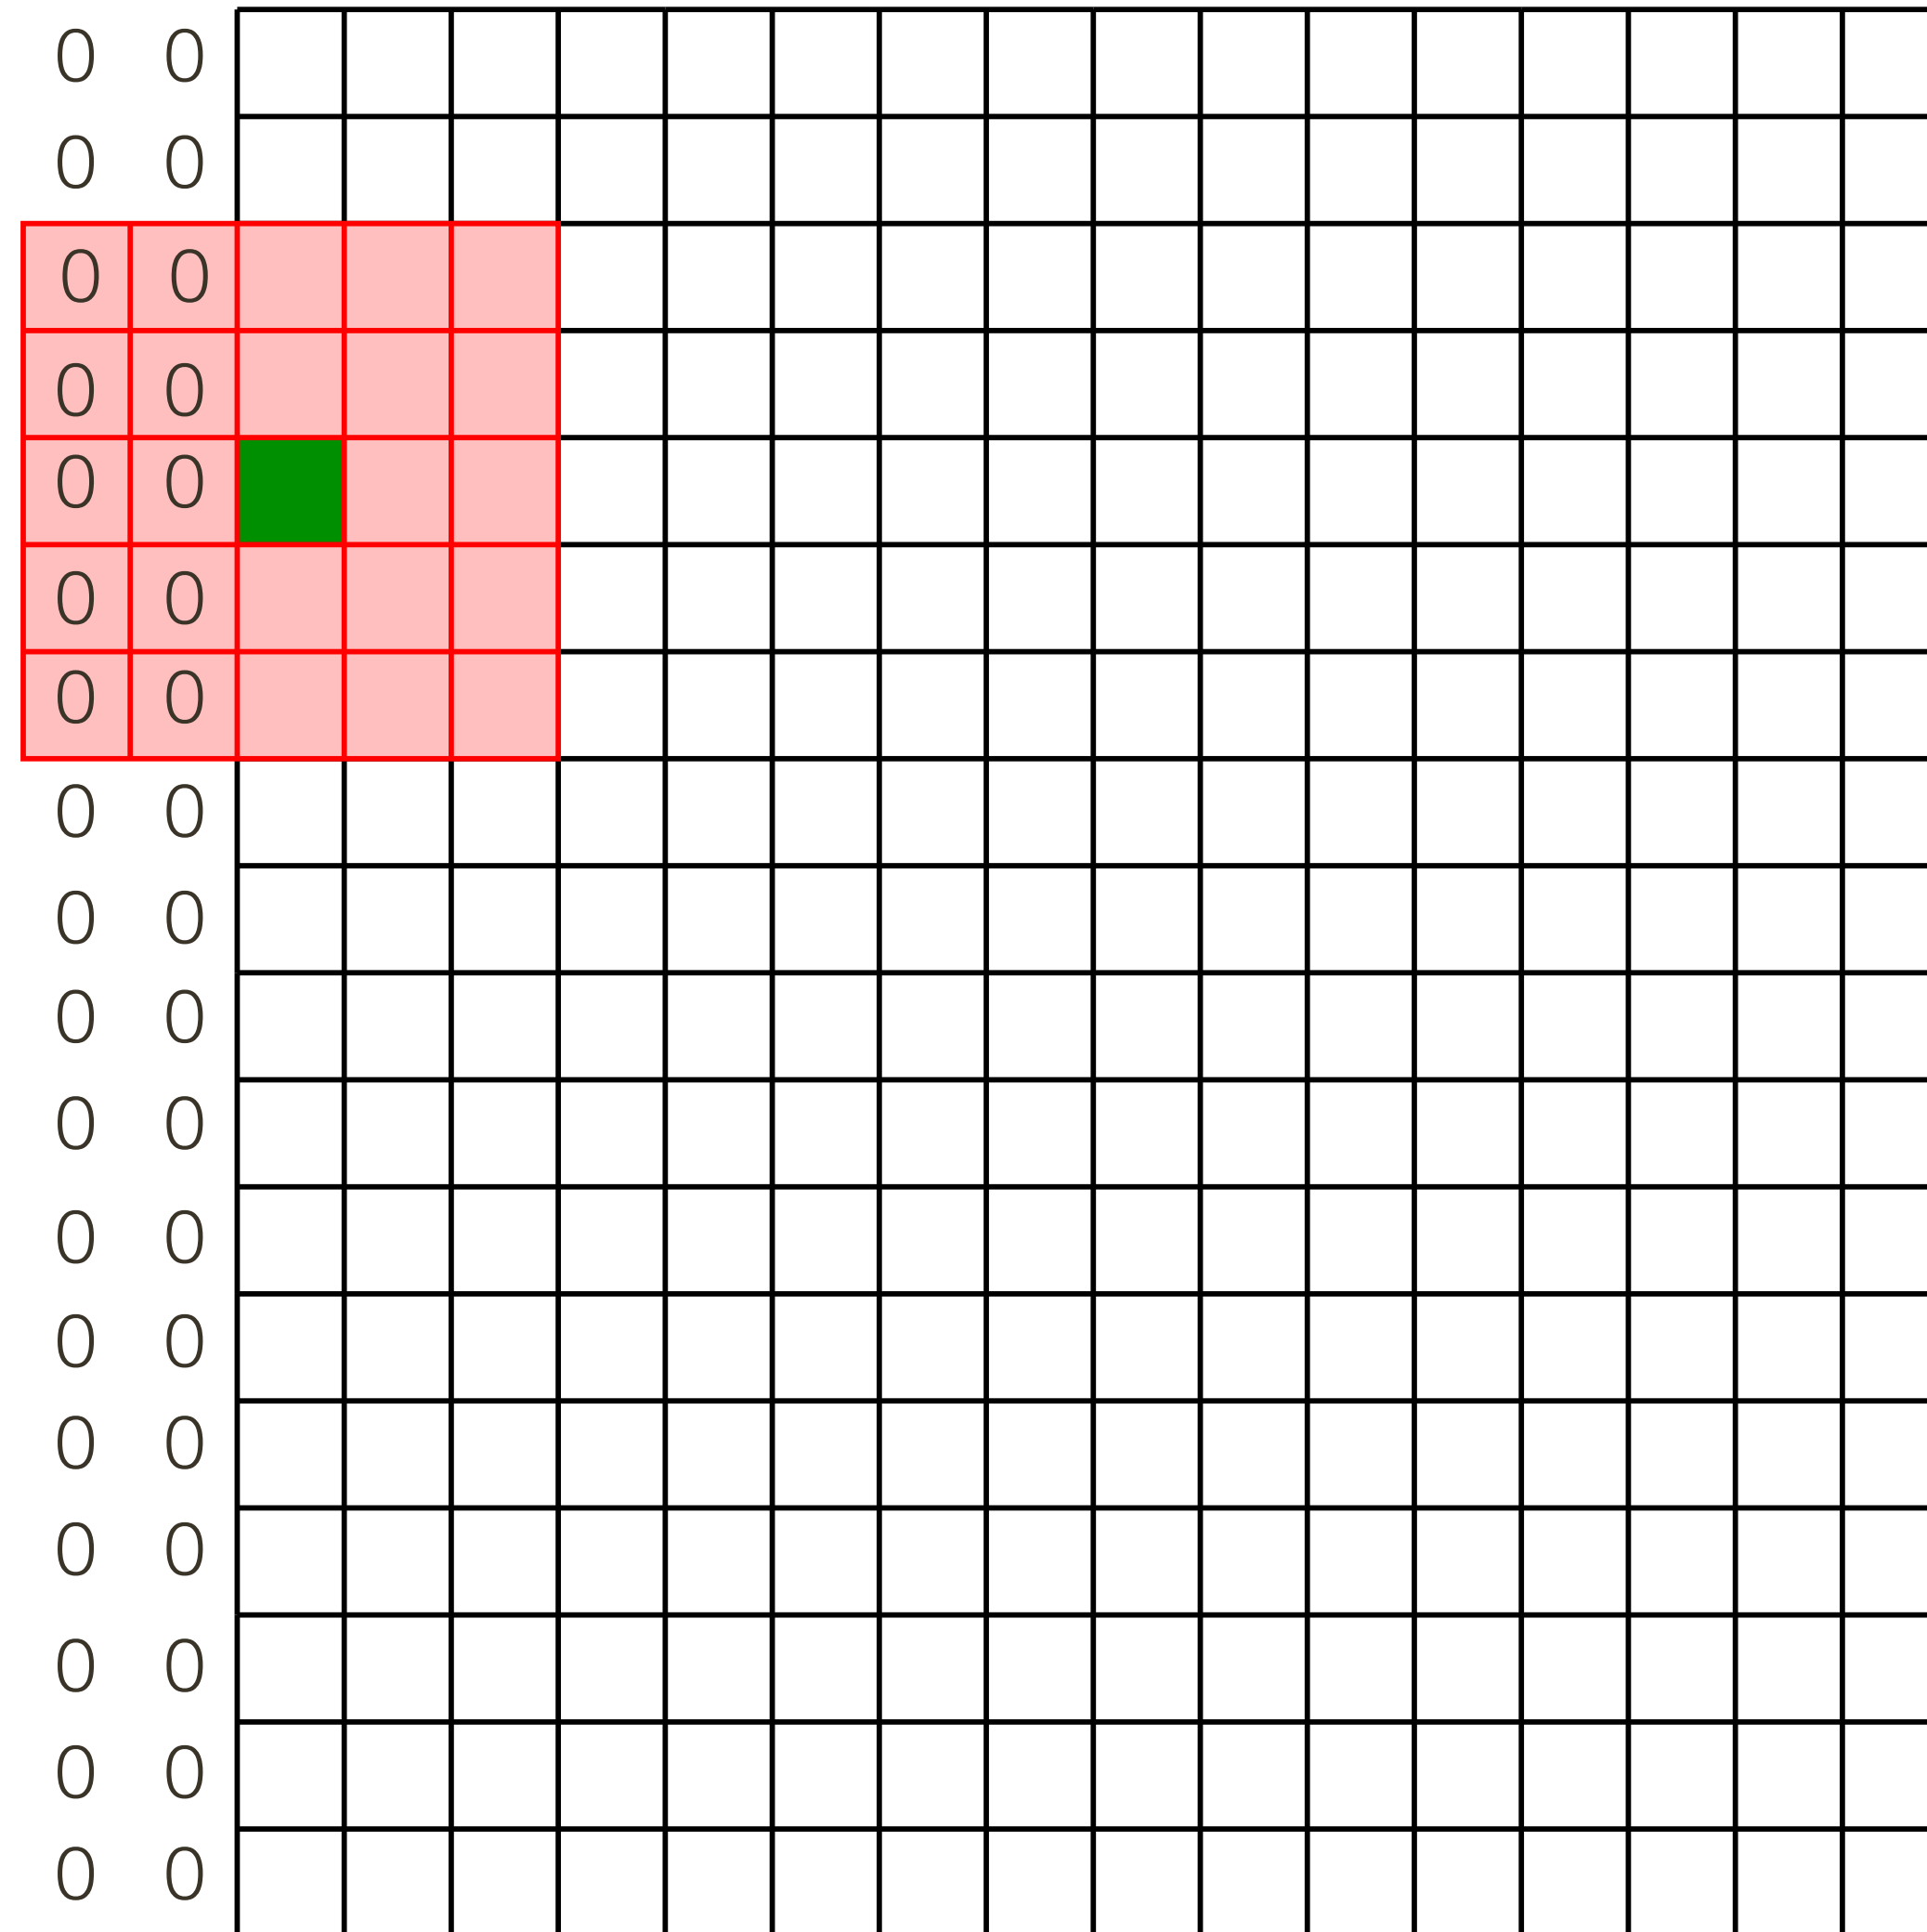
Linear Filters: **Boundary** Effects



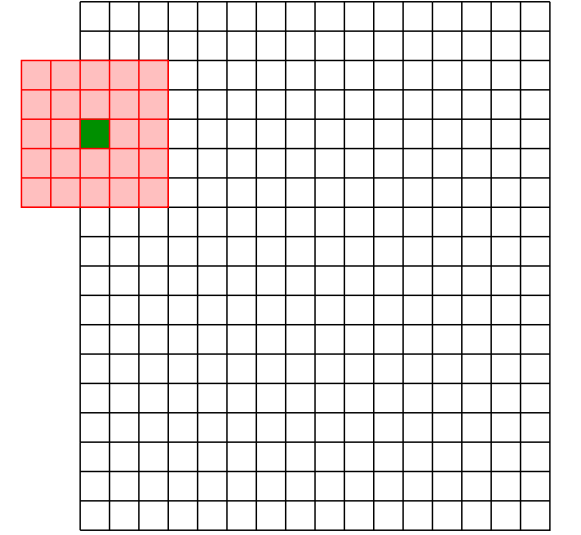
Three standard ways to deal with boundaries:

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



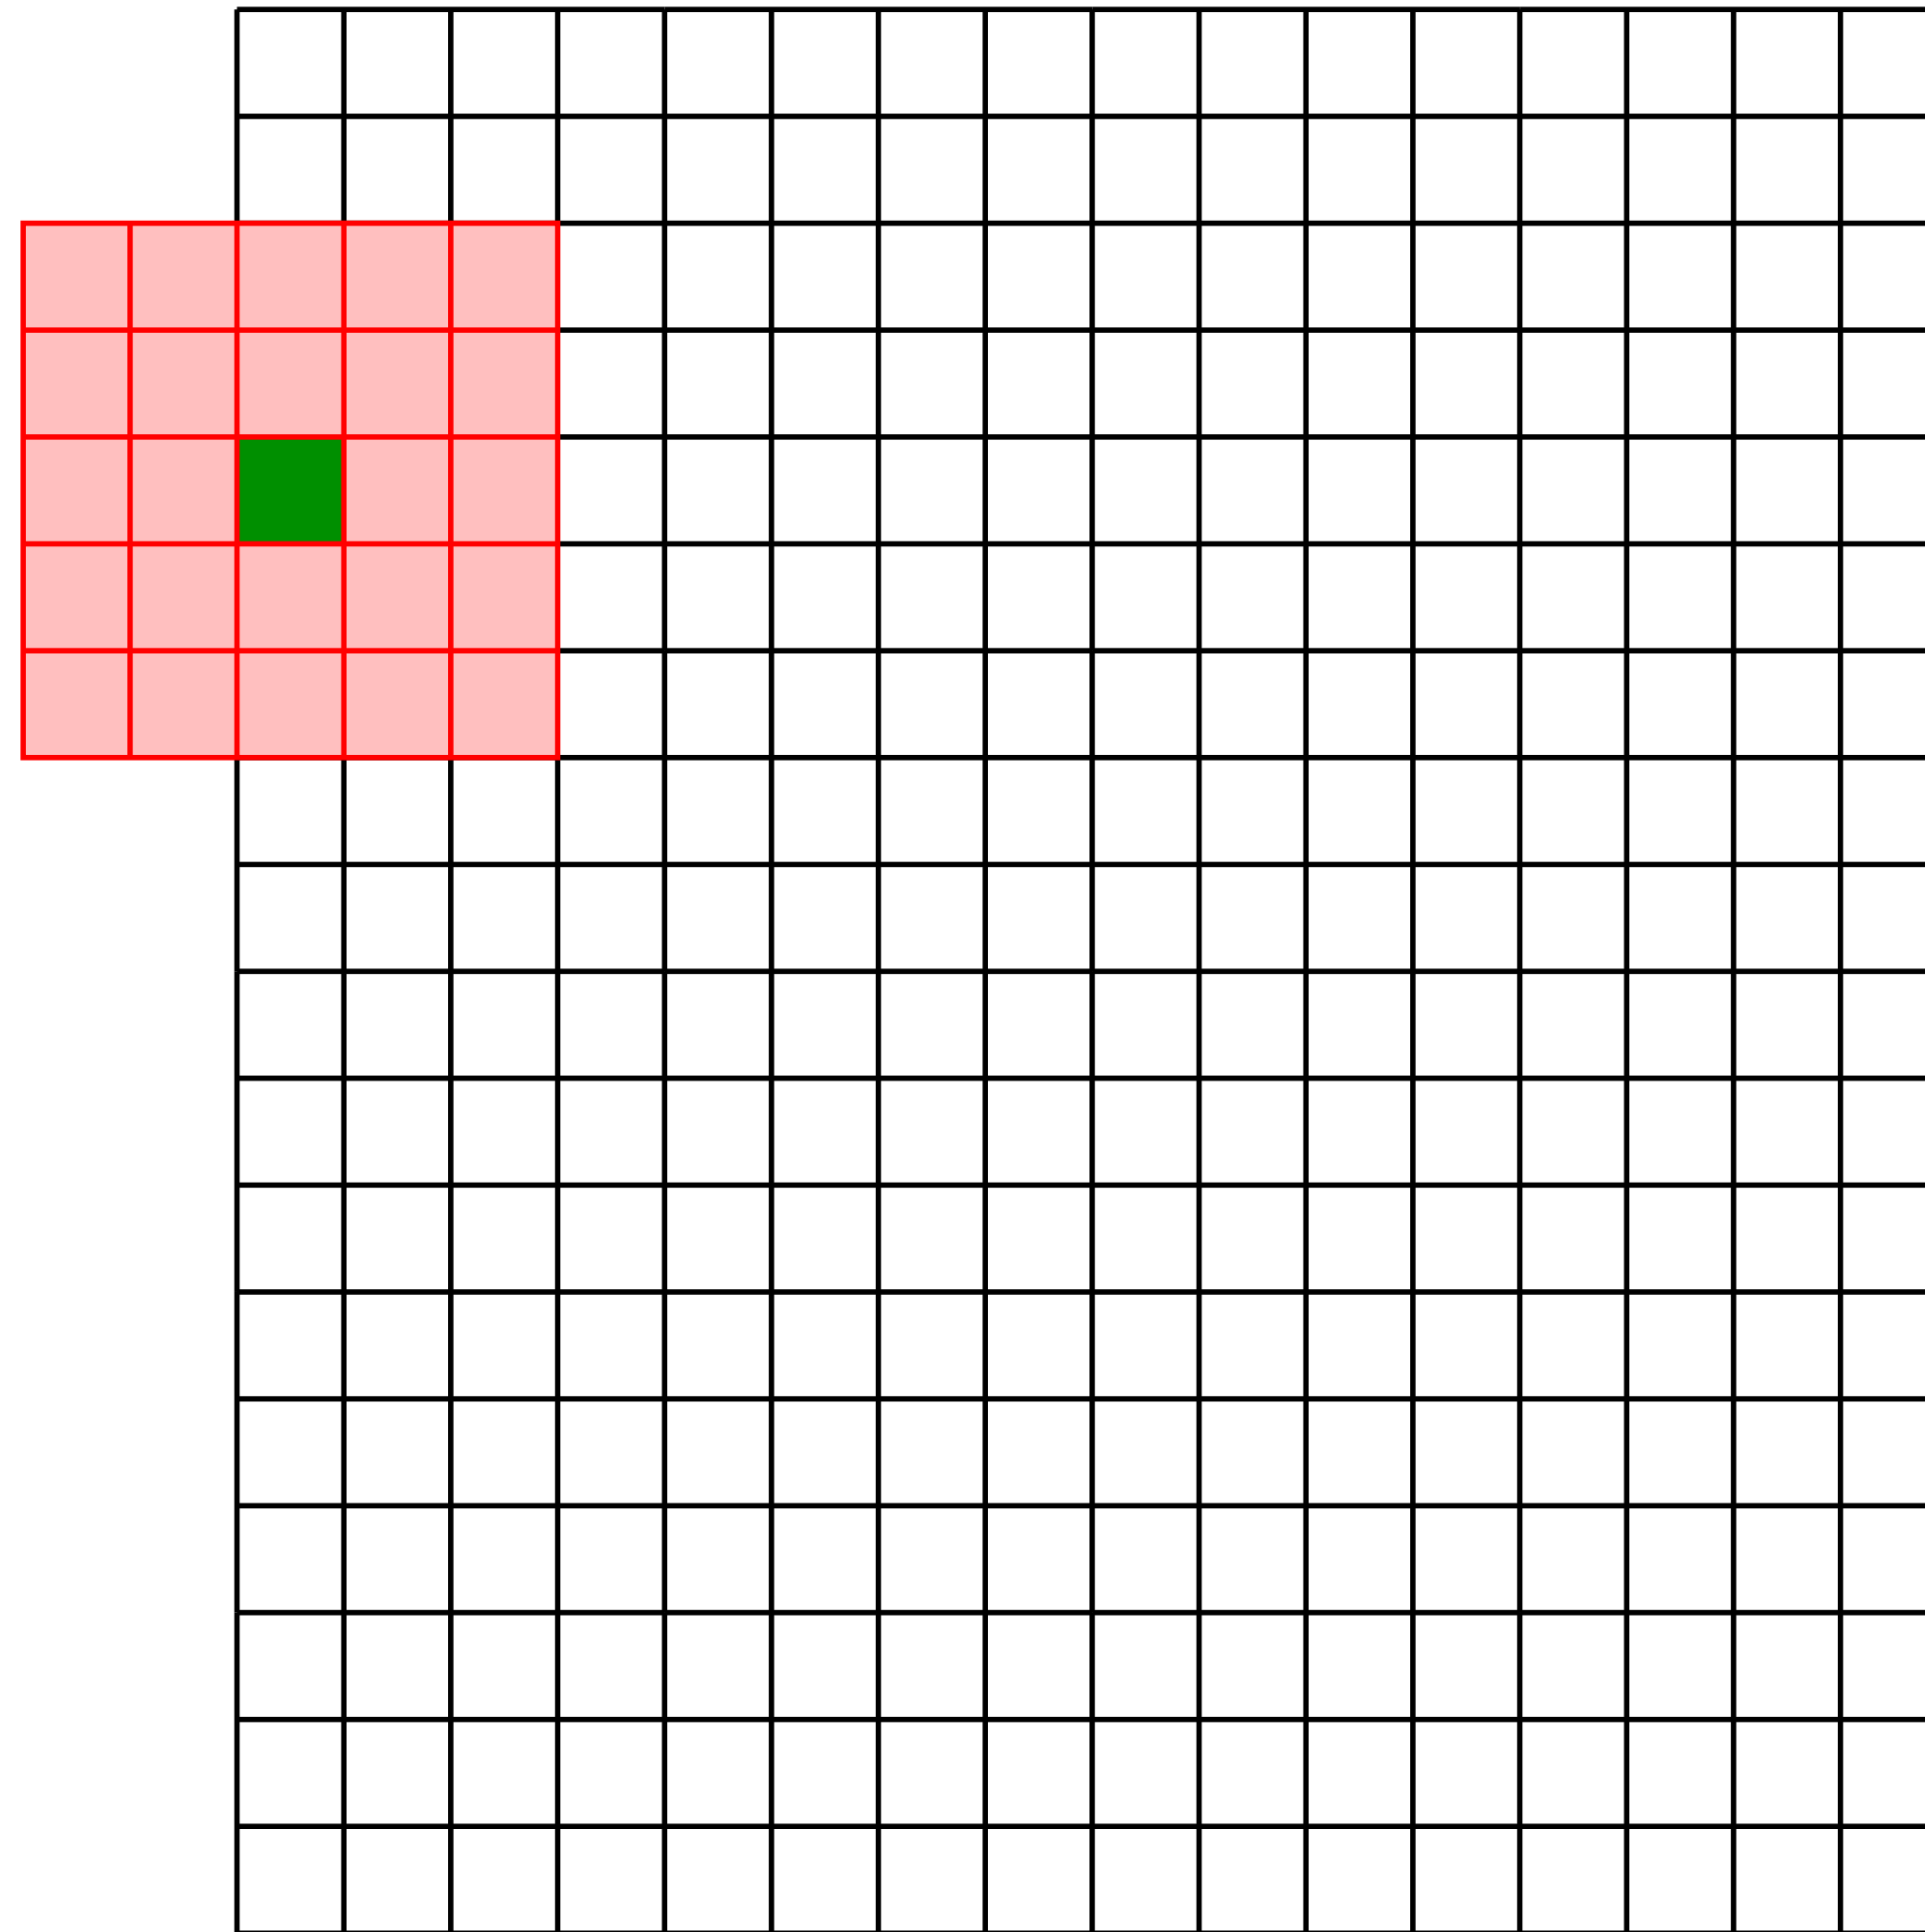
Linear Filters: **Boundary** Effects



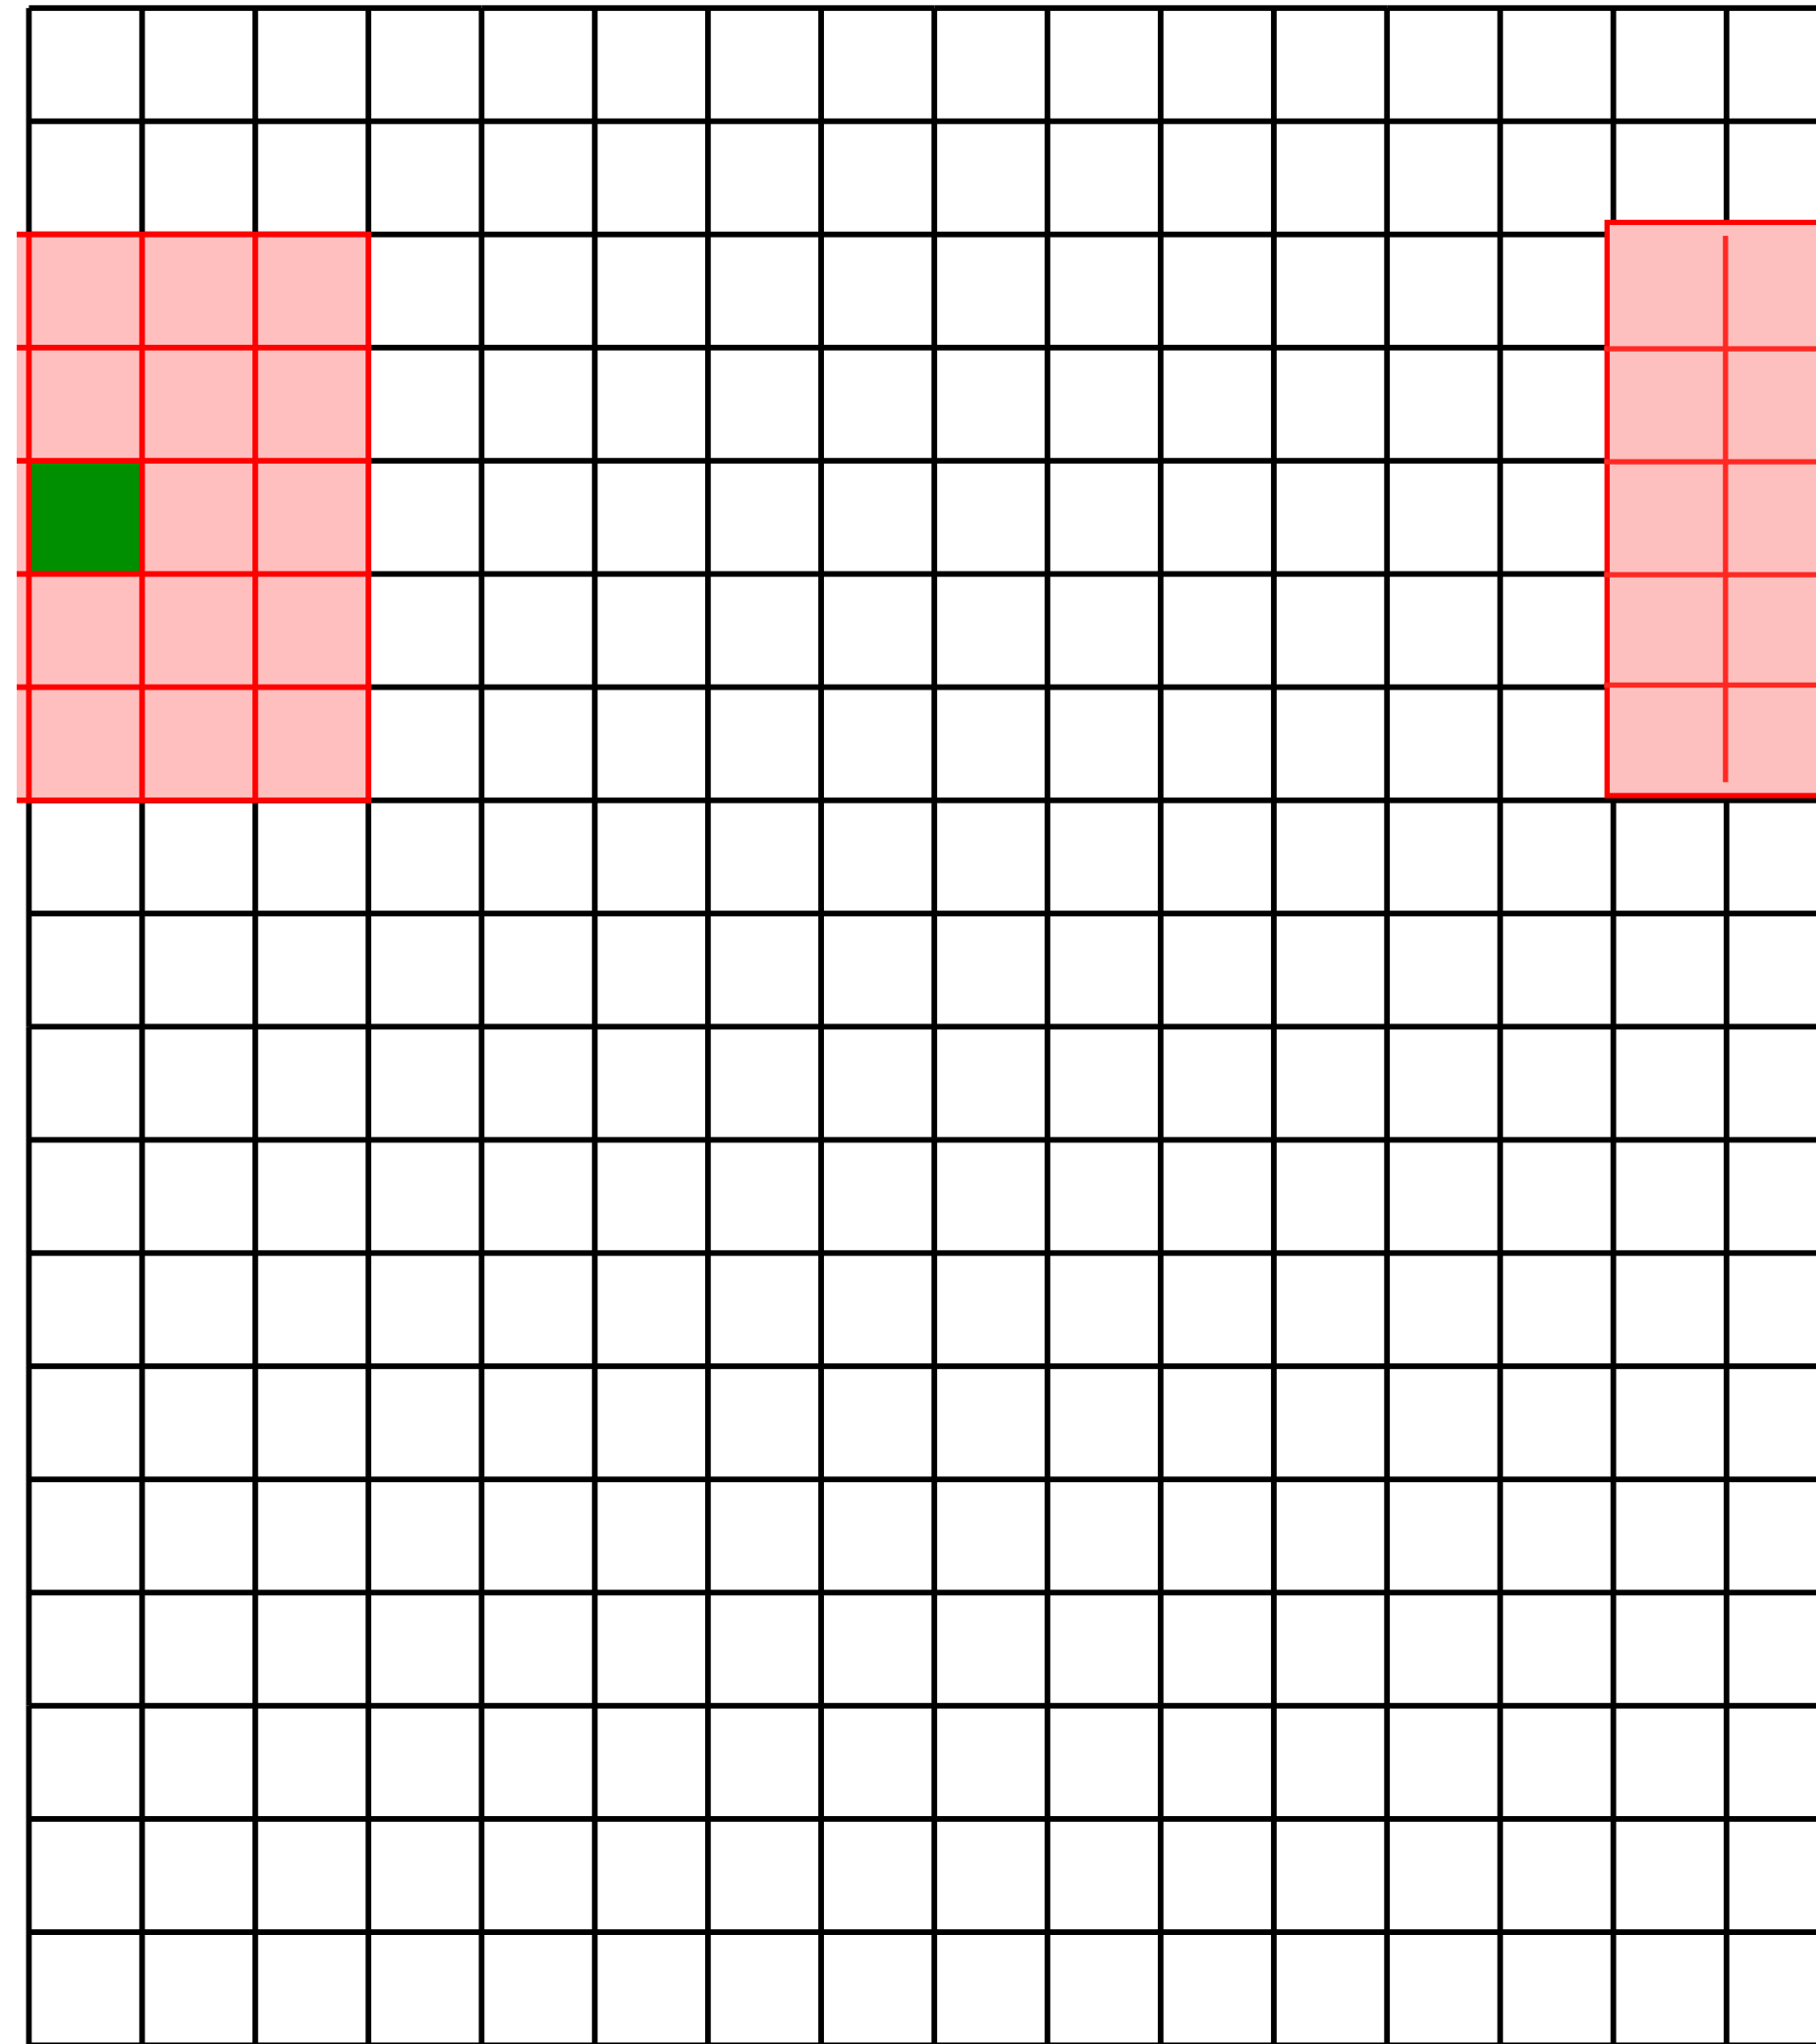
Three standard ways to deal with boundaries:

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



A short exercise ...

Example 1: Warm up



Original

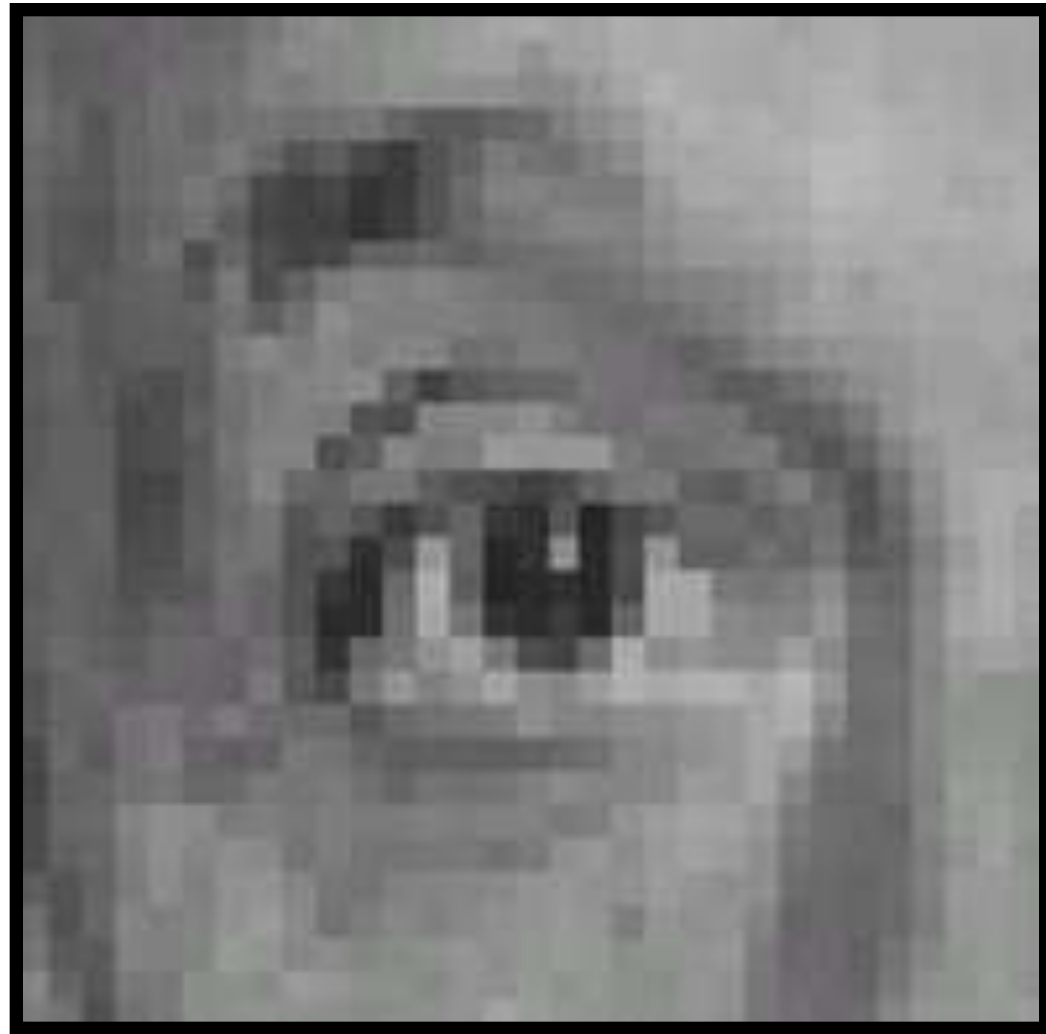
0	0	0
0	1	0
0	0	0

Filter



Result

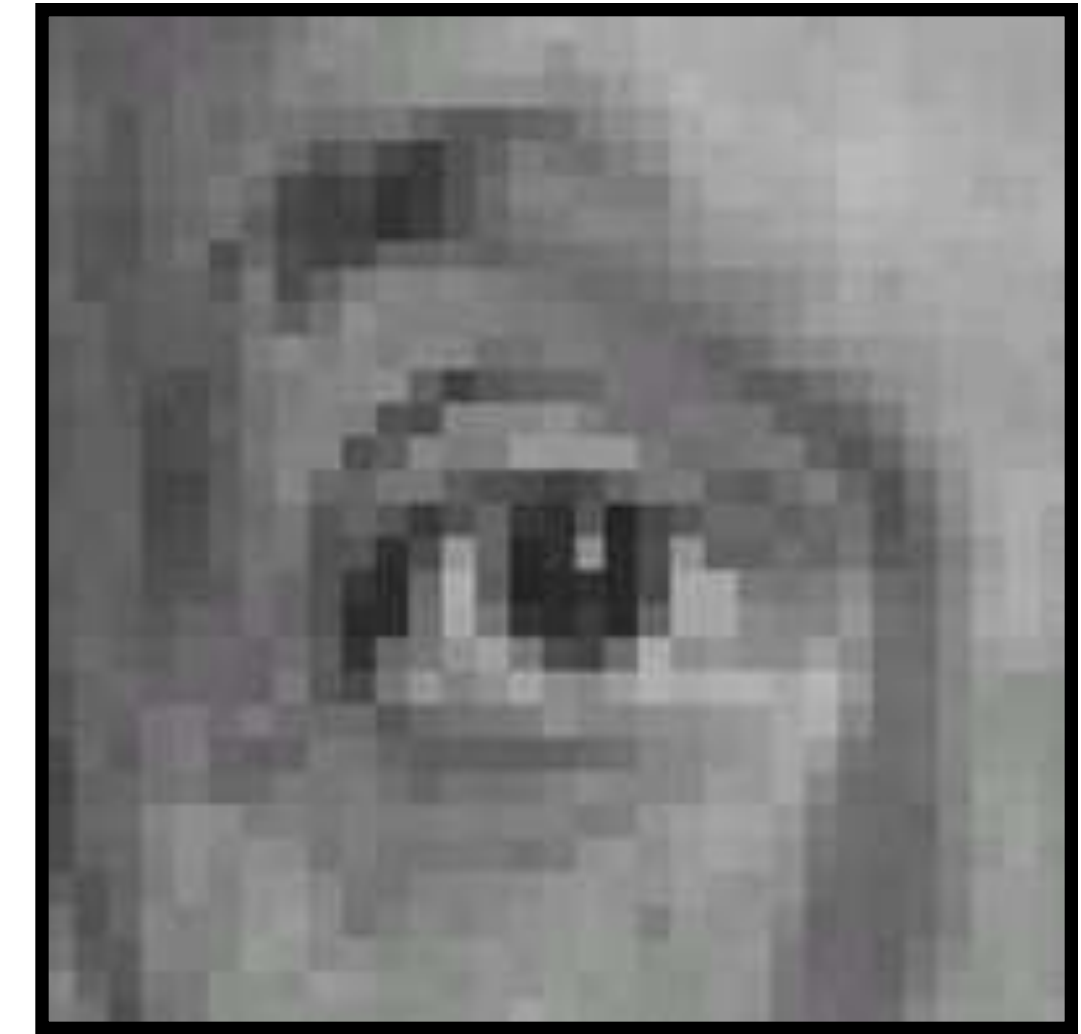
Example 1: Warm up



Original

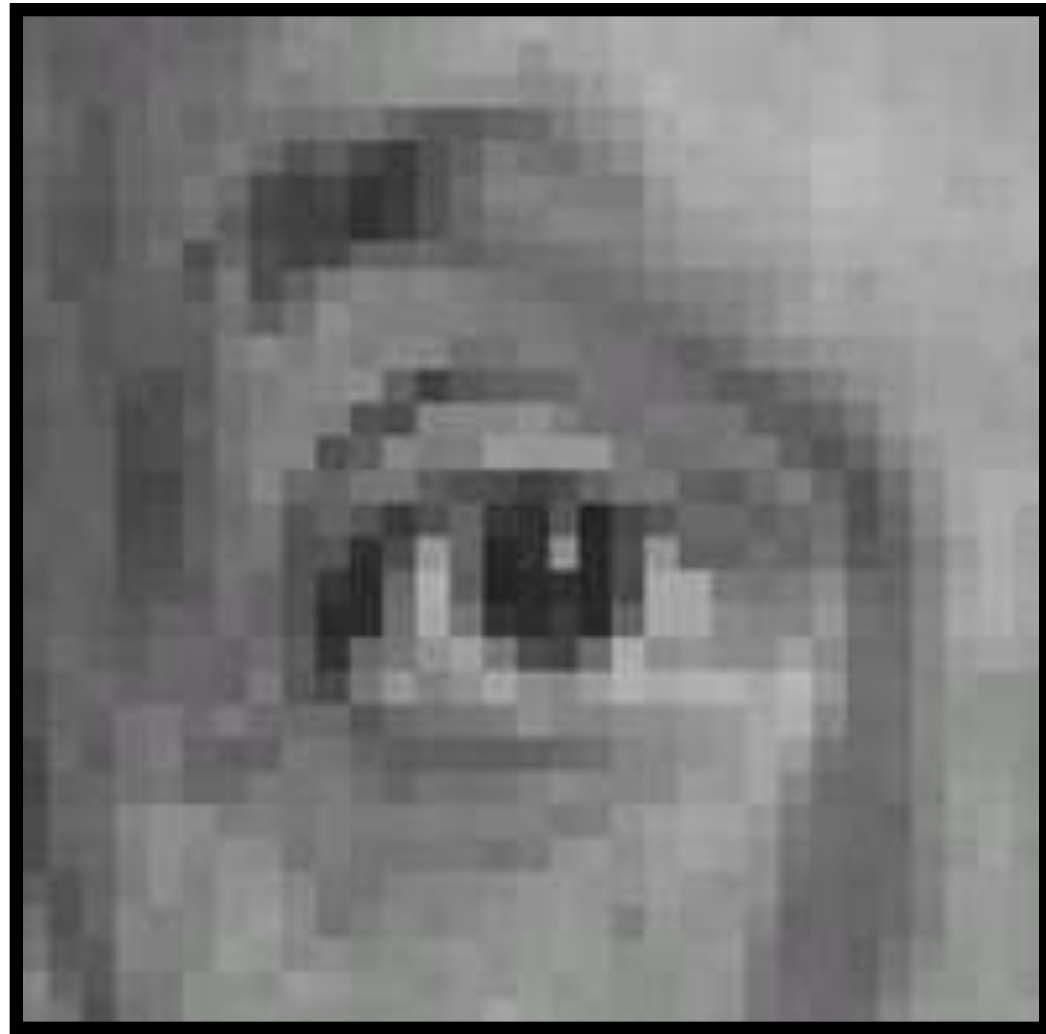
0	0	0
0	1	0
0	0	0

Filter



Result
(no change)

Example 2:



Original

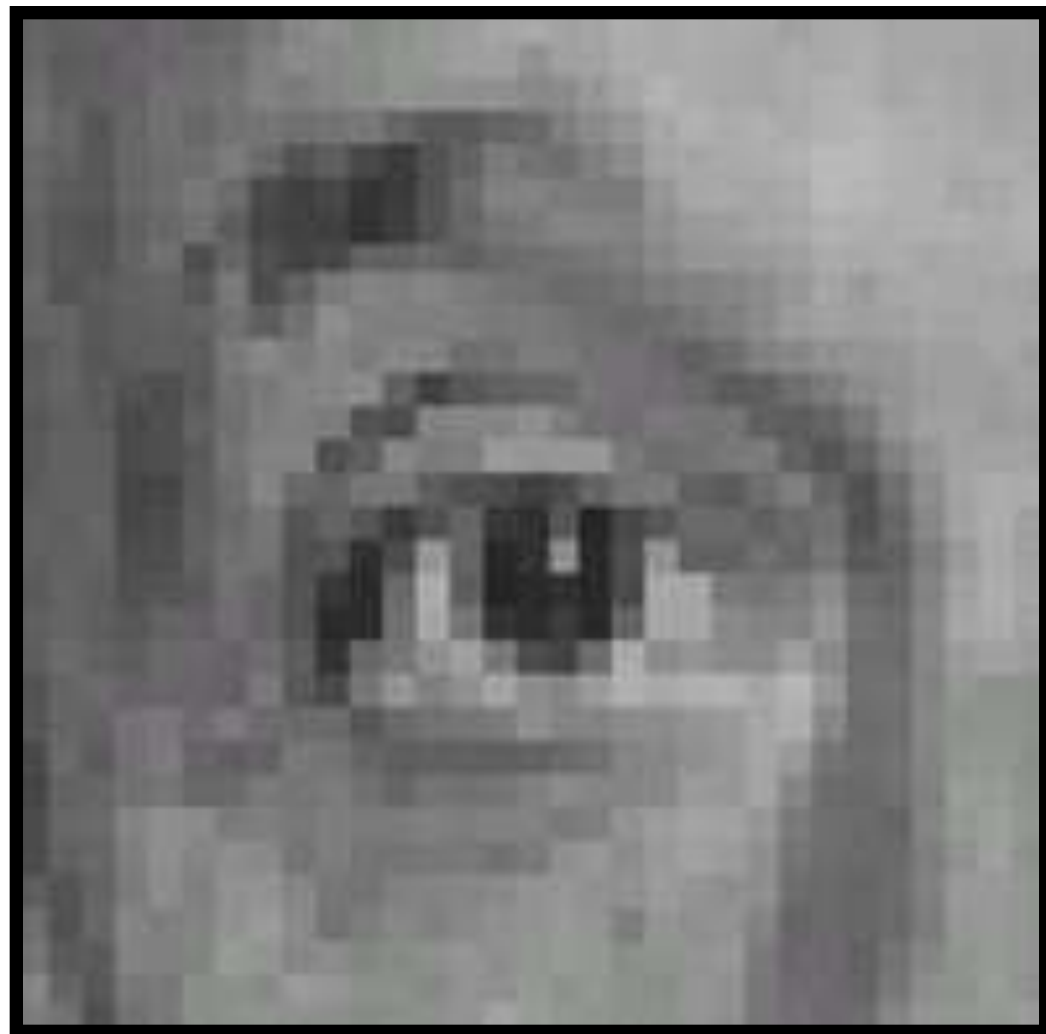
0	0	0
0	0	1
0	0	0

Filter



Result

Example 2:



Original

0	0	0
0	0	1
0	0	0

Filter



Result
(sift left by 1 pixel)

Example 3:



Original

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

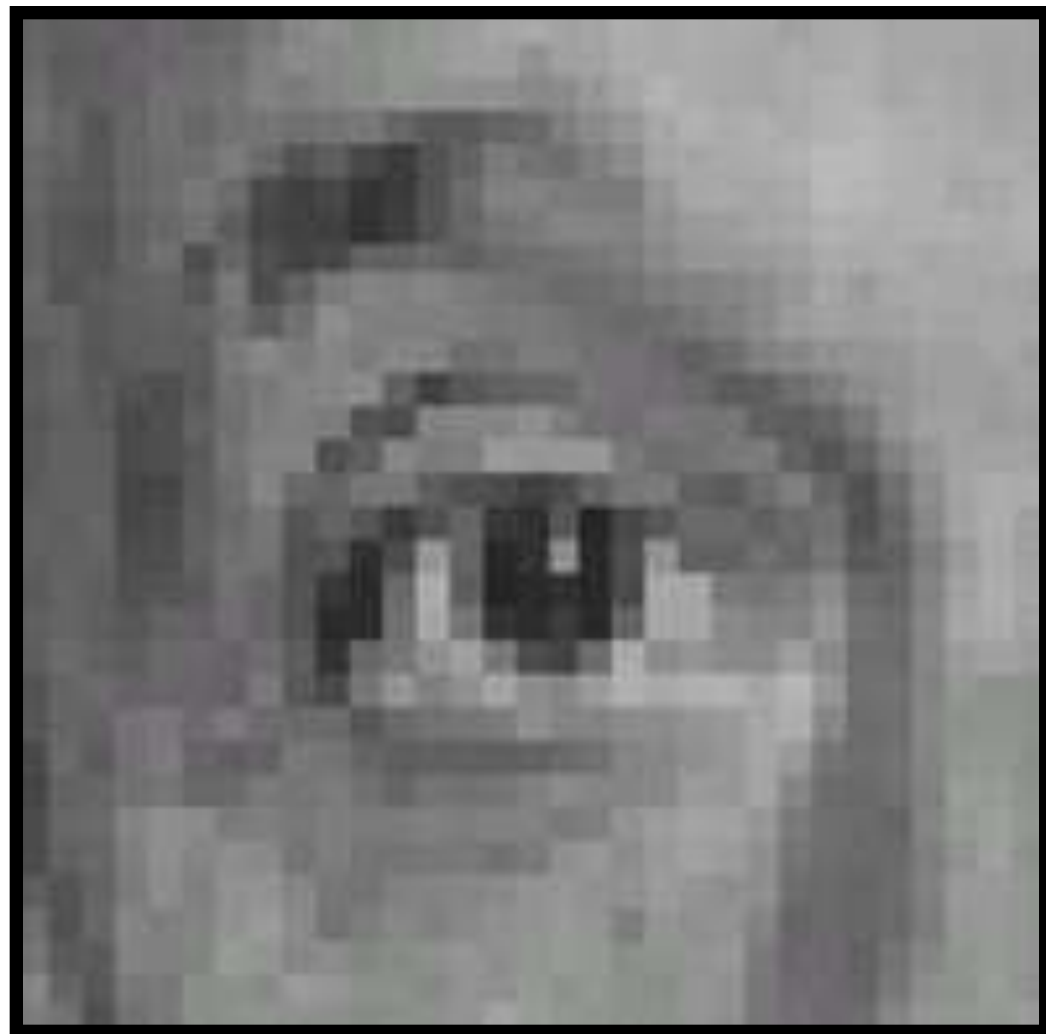
Filter

(filter sums to 1)



Result

Example 3:



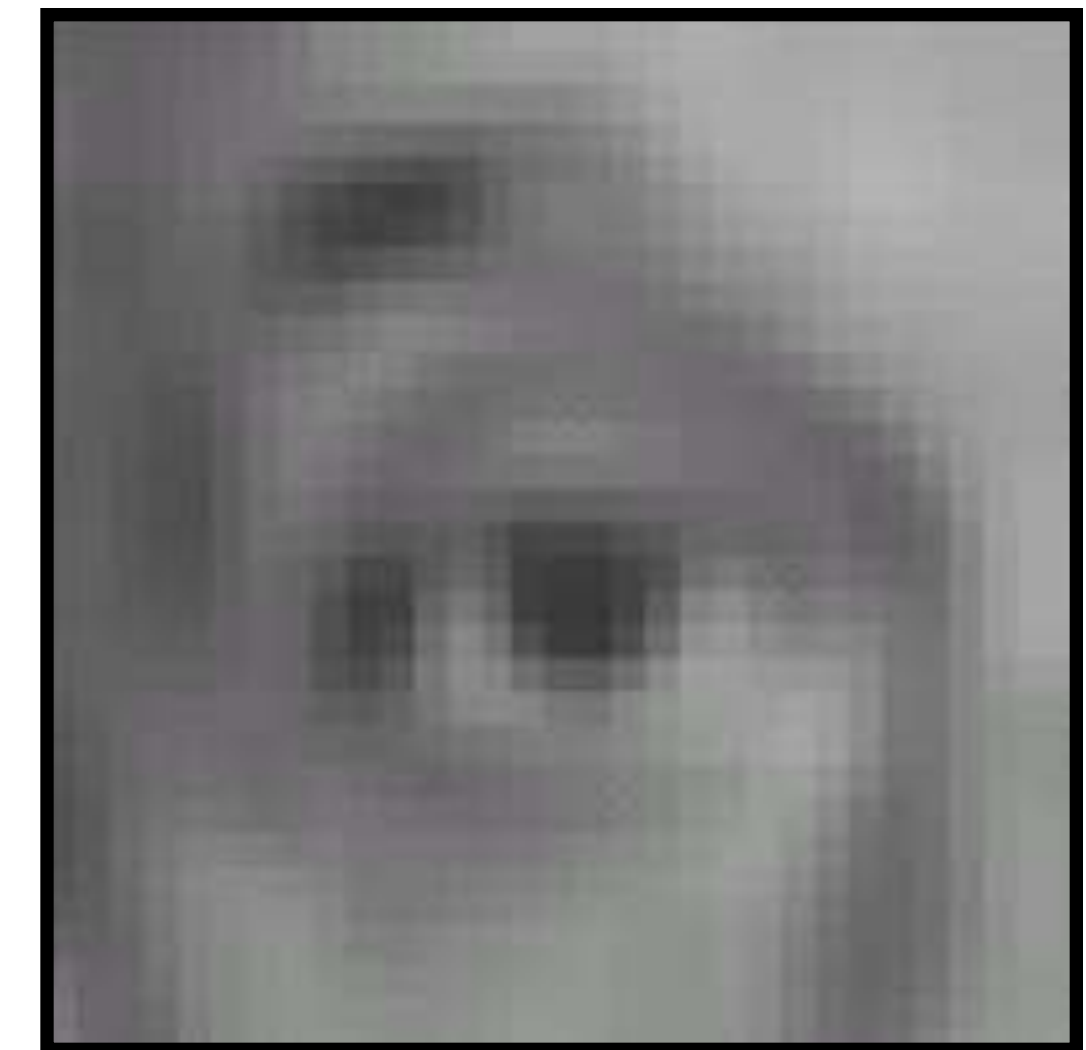
Original

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

1	1	1
1	1	1
1	1	1

Filter

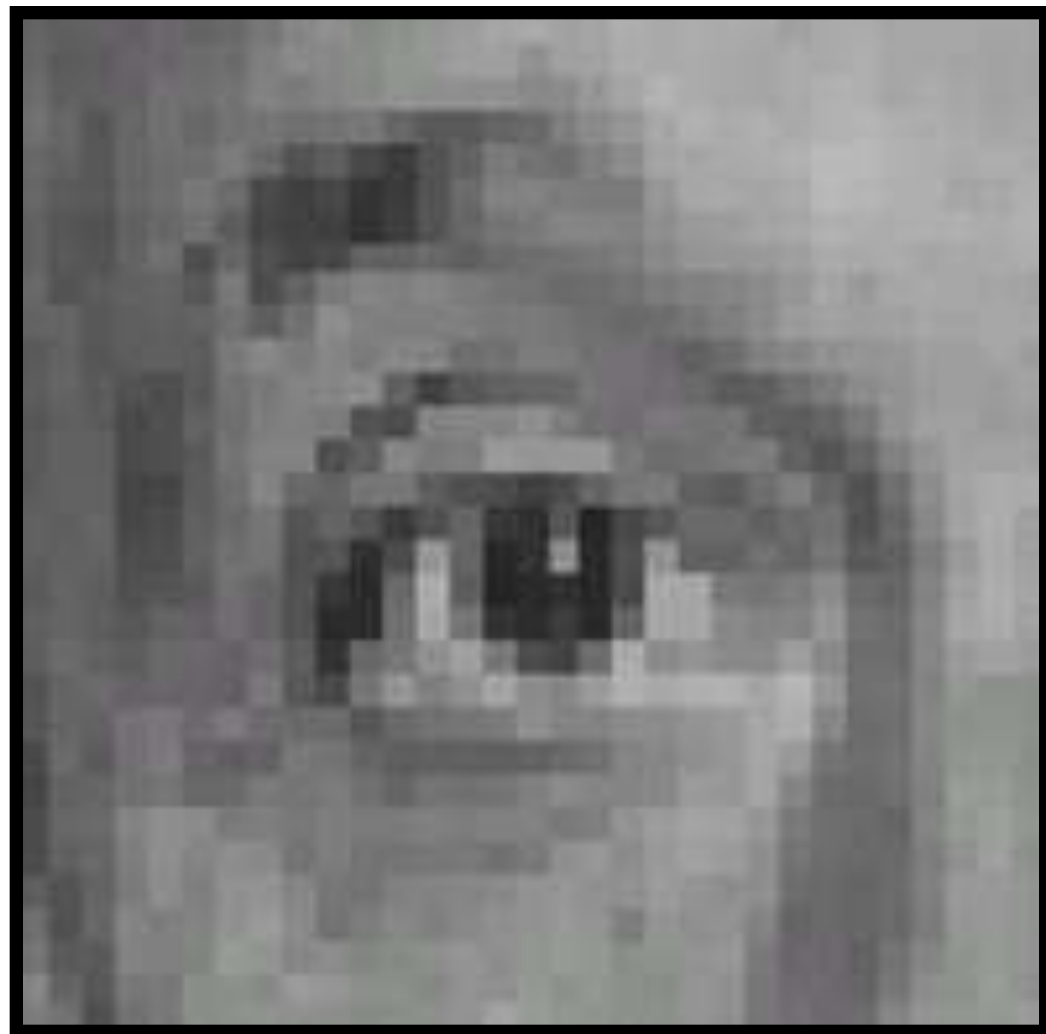
(filter sums to 1)



Result

(blur with a box filter)

Example 4:



Original

$$\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 2 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$

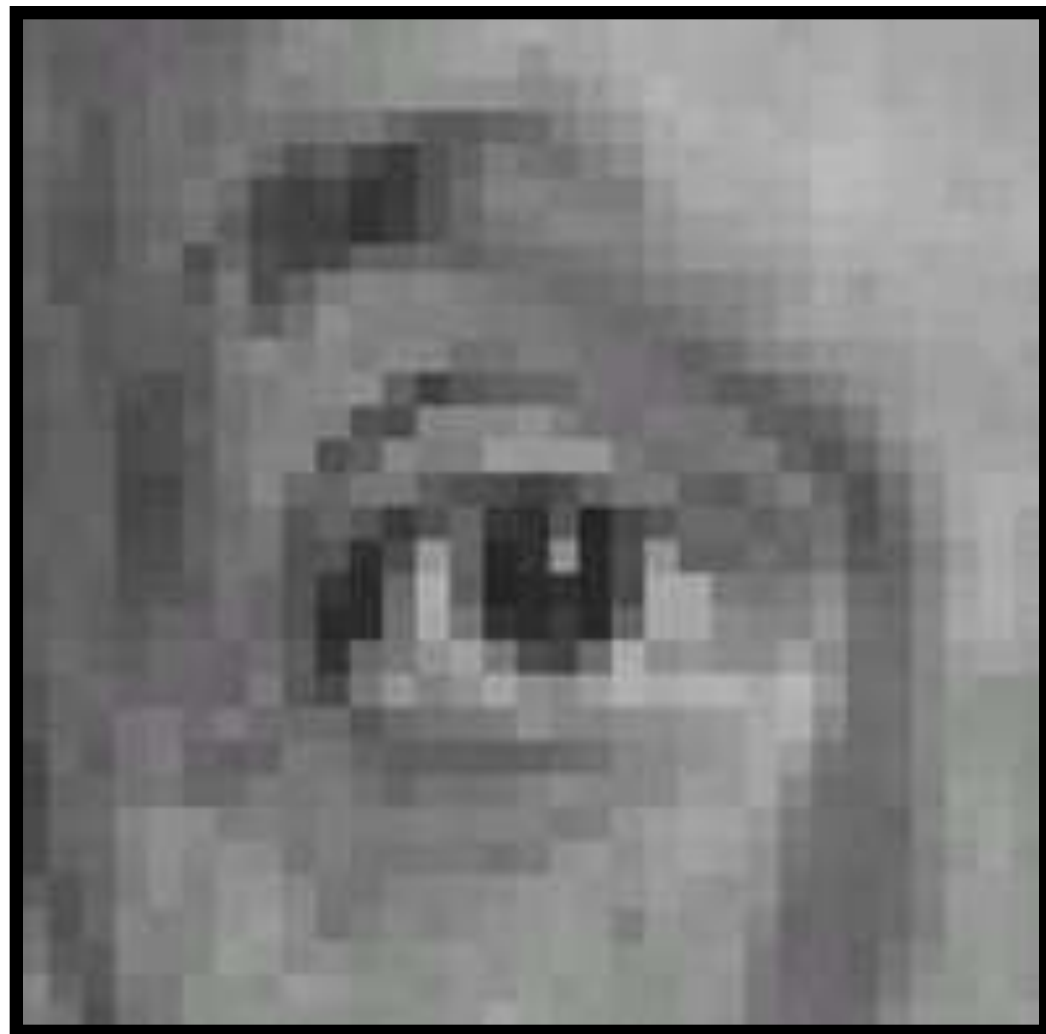
Filter

(filter sums to 1)



Result

Example 4:



Original

$$\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 2 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$

Filter

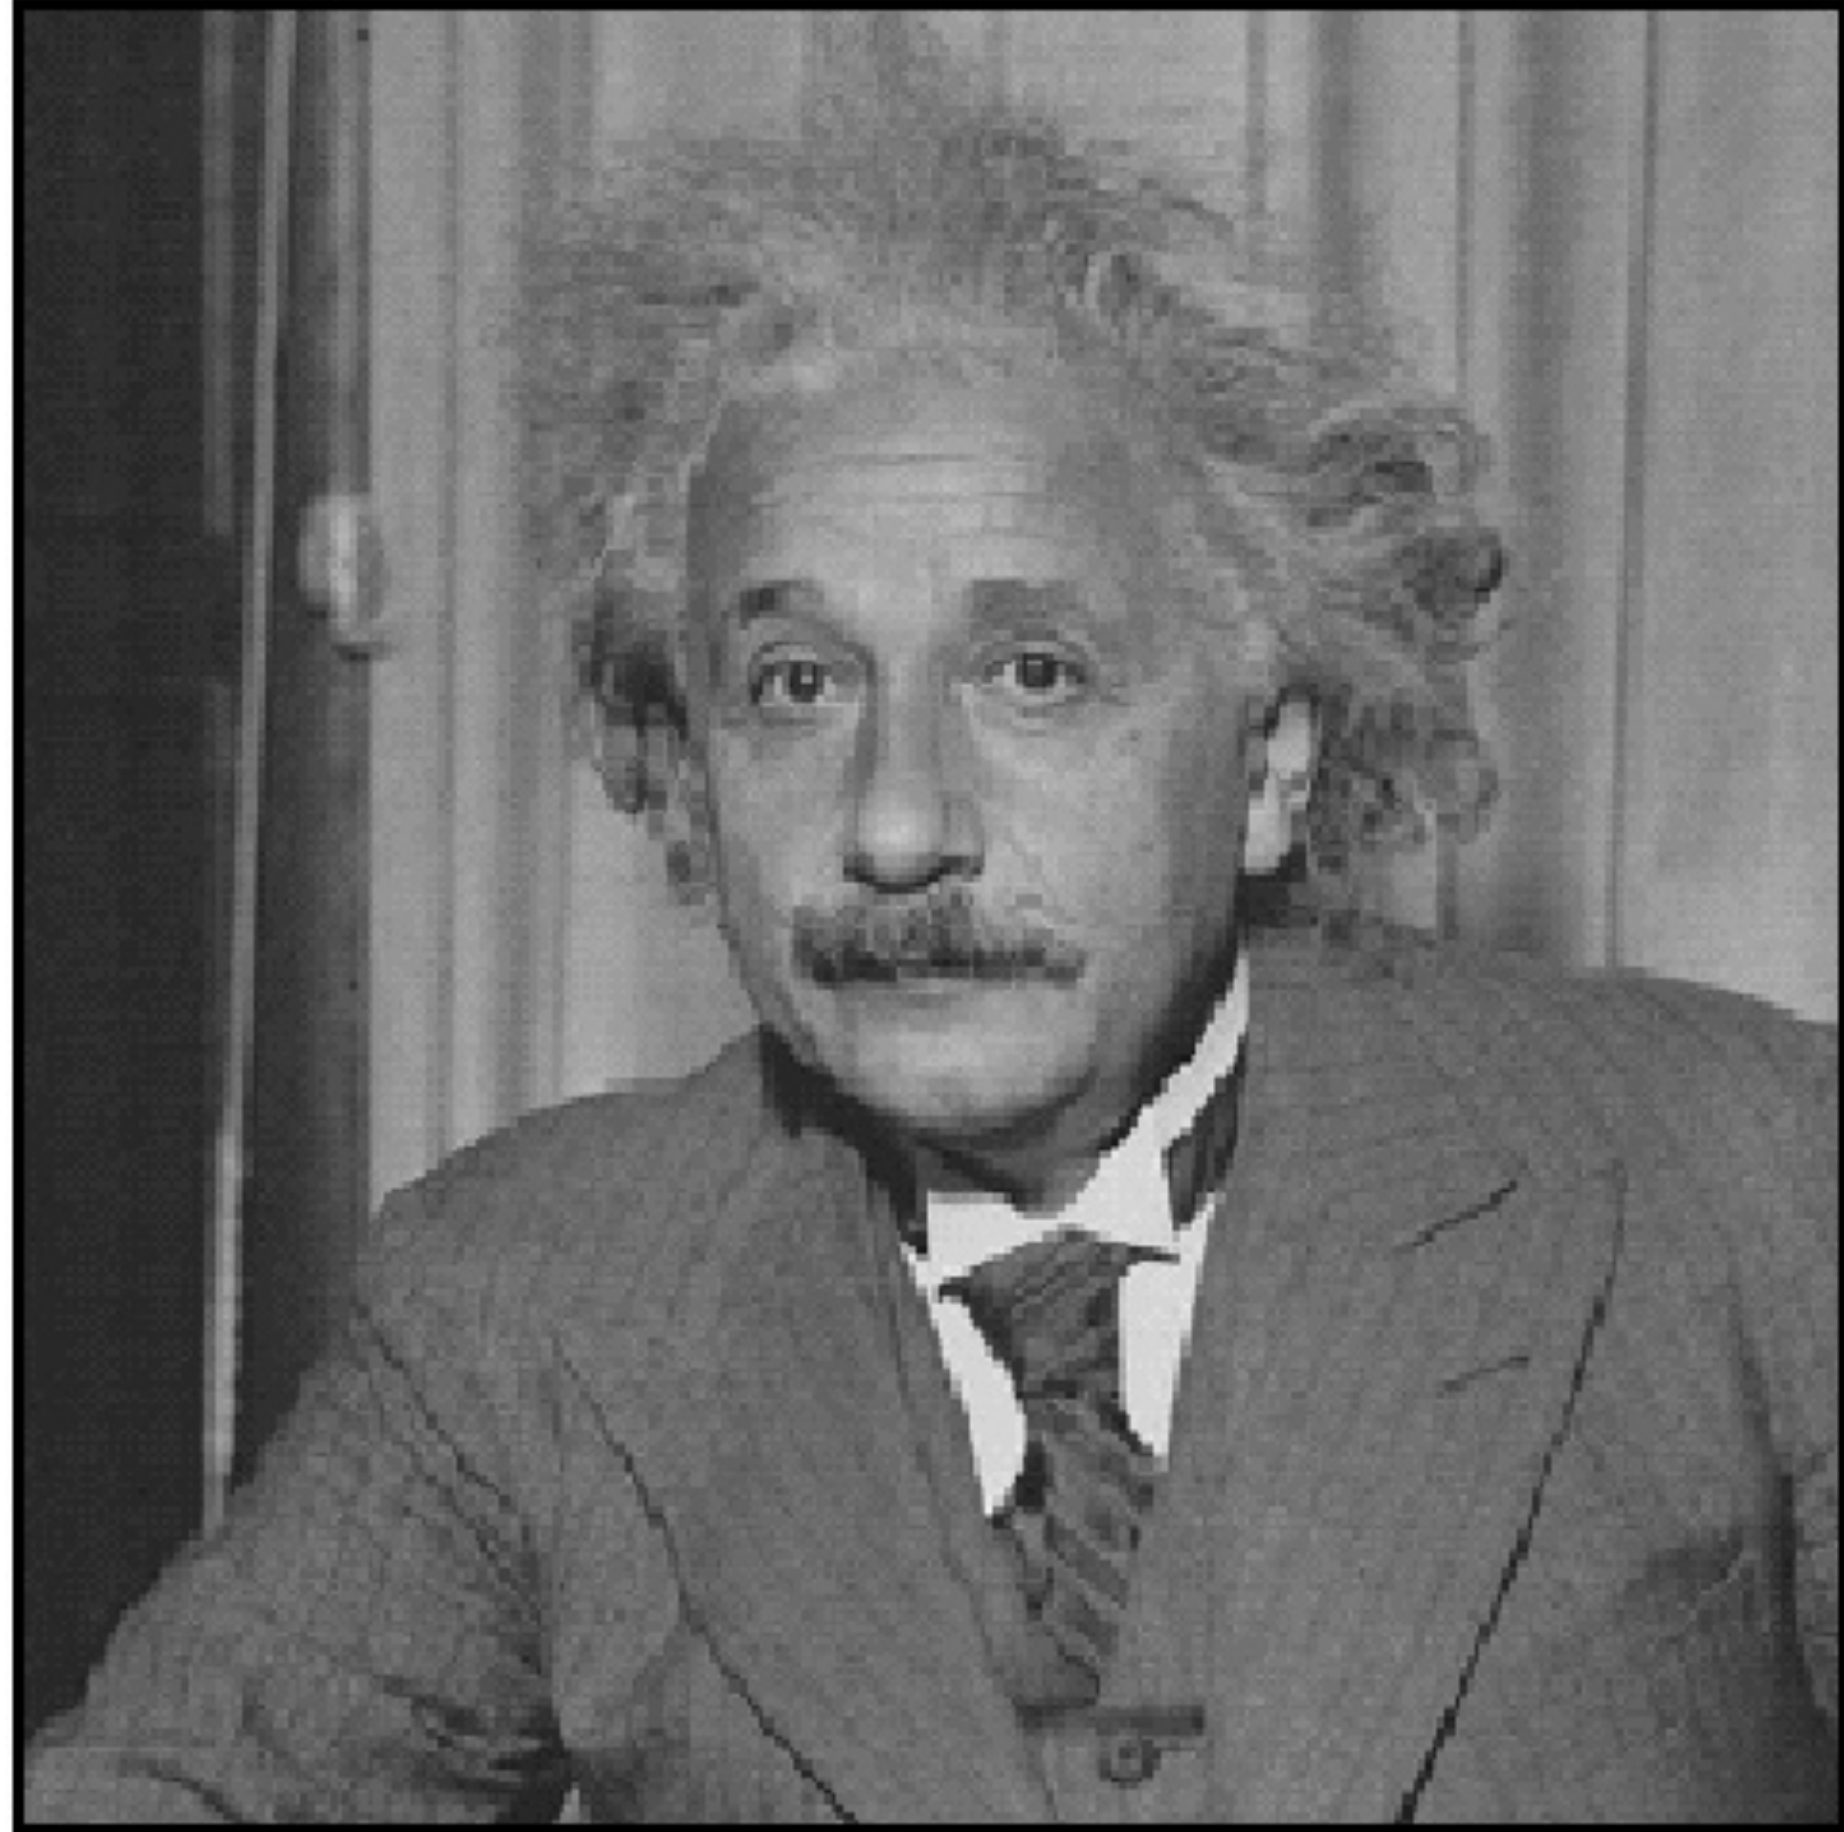
(filter sums to 1)



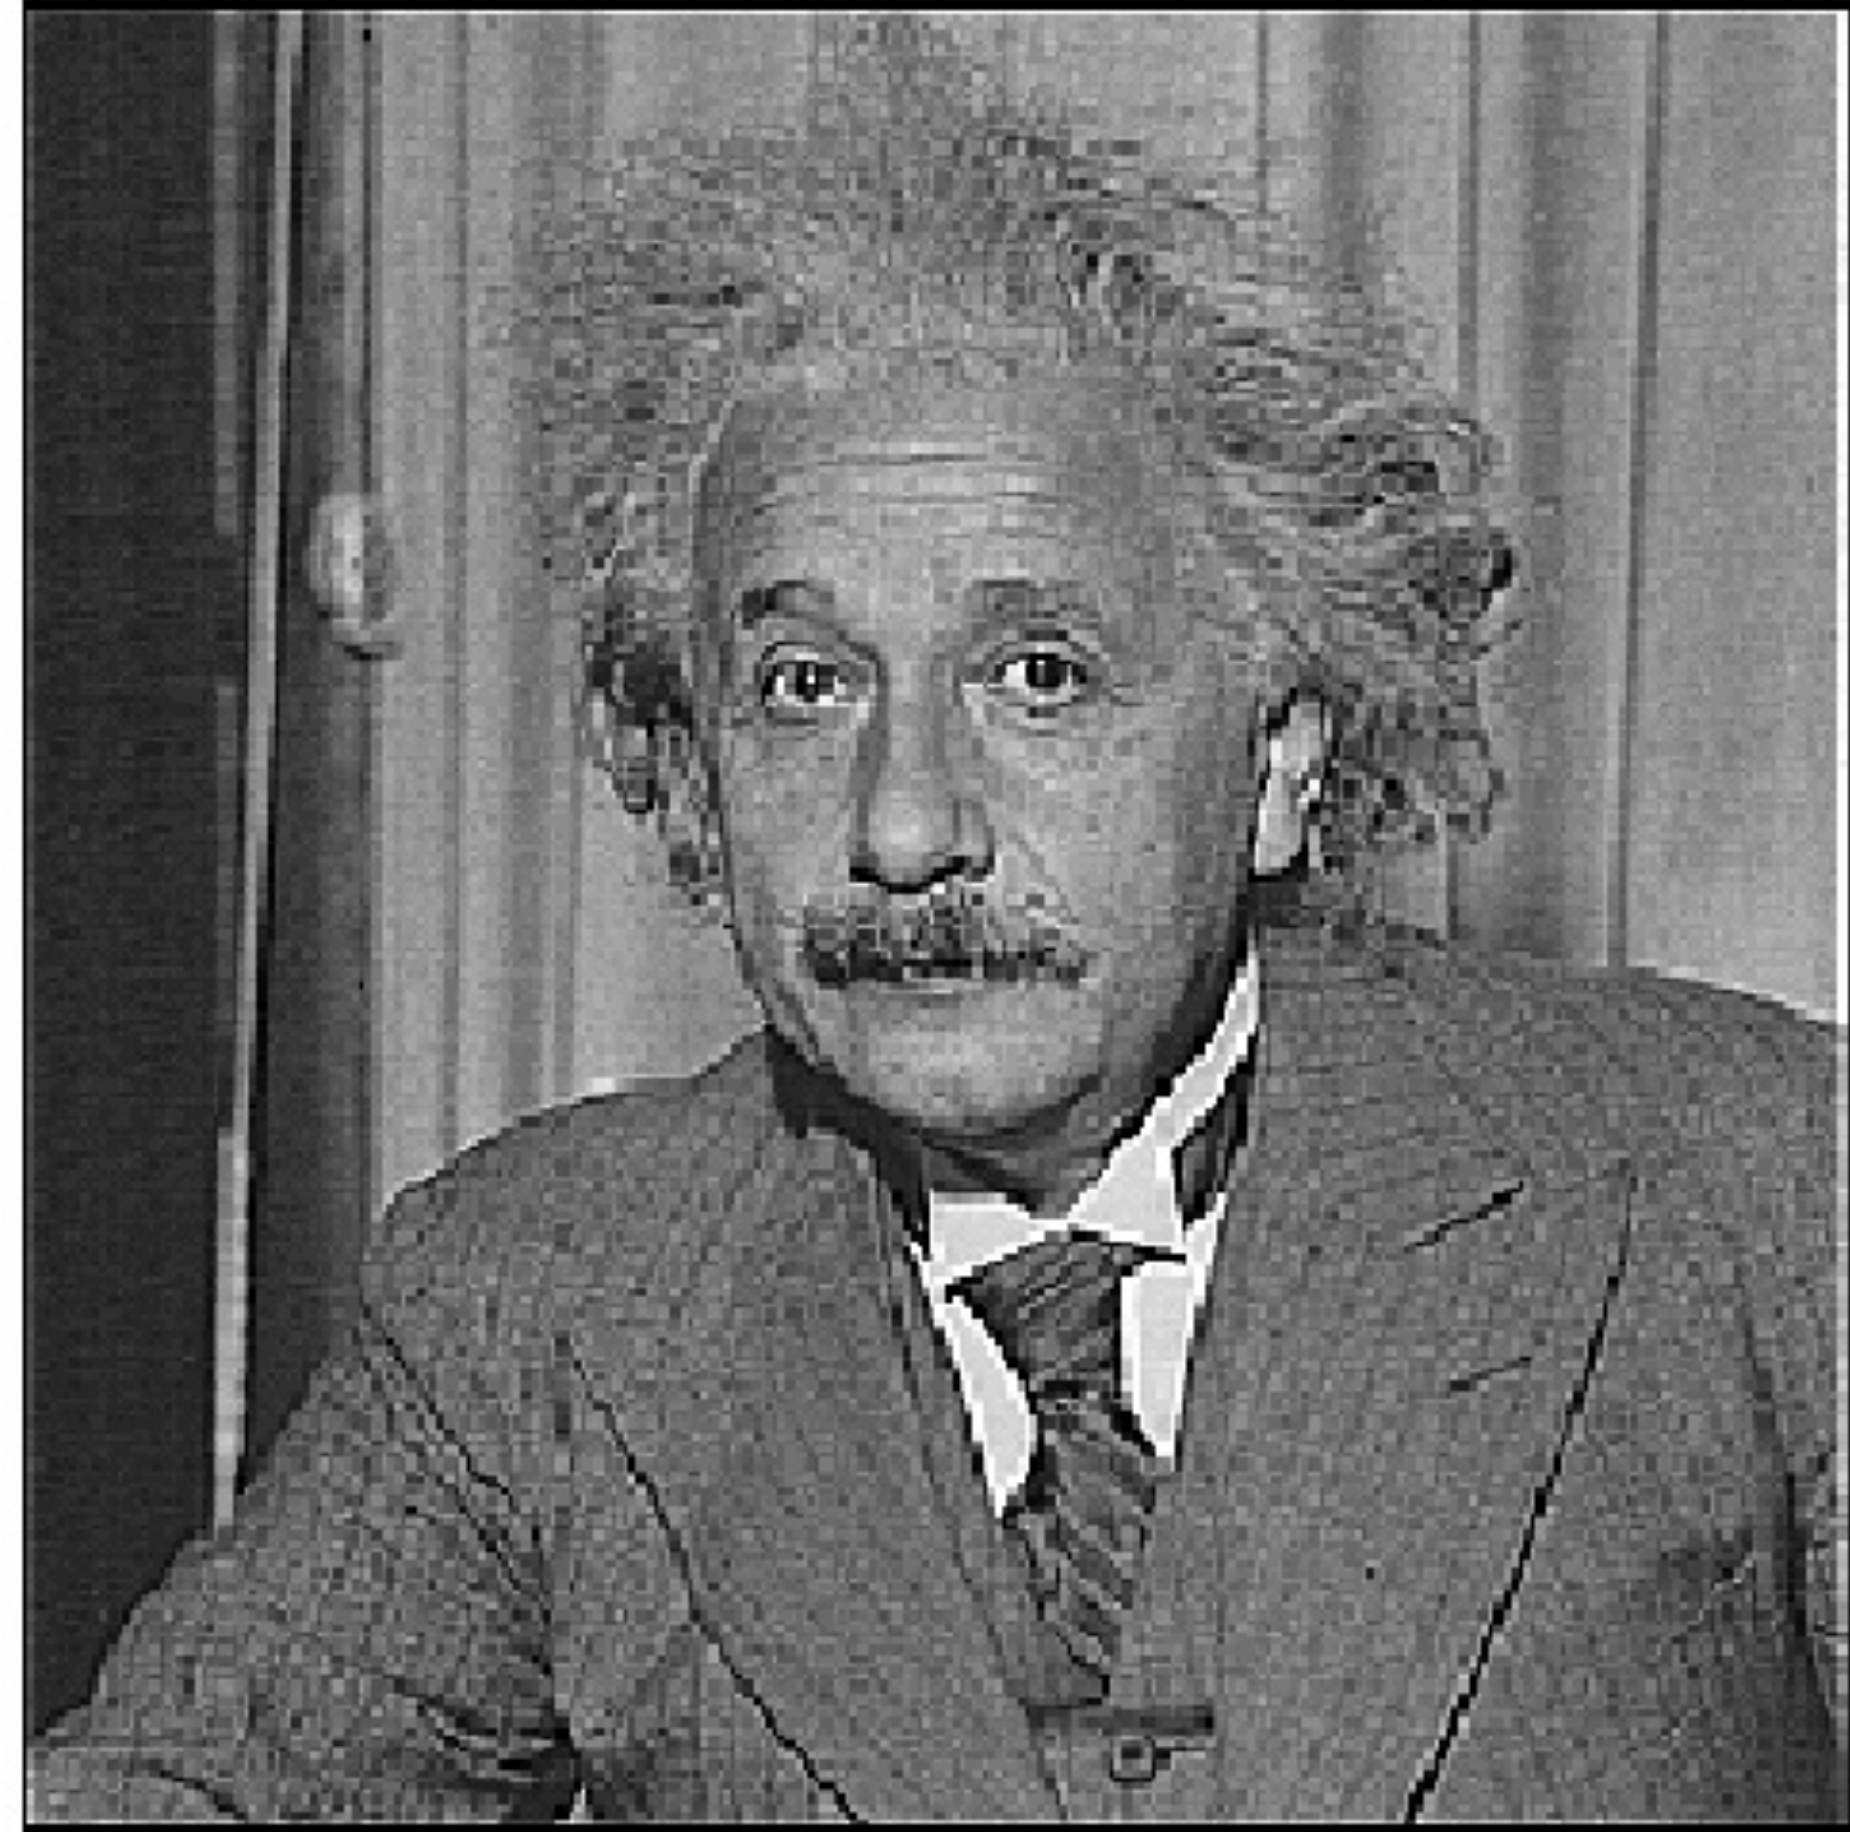
Result

(sharpening)

Example 4: Sharpening



Before



After

Example 4: Sharpening



Before



After

Linear Filters: Correlation vs. Convolution

Definition: **Correlation**

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

Linear Filters: Correlation vs. Convolution

Definition: **Correlation**

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

Definition: **Convolution**

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

Linear Filters: Correlation vs. Convolution

Definition: **Correlation**

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

a	b	c
d	e	f
g	h	i

Filter

1	2	3
4	5	6
7	8	9

Image

Output

$$\begin{aligned} &= 1a + 2b + 3c \\ &\quad + 4d + 5e + 6f \\ &\quad + 7g + 8h + 9i \end{aligned}$$

Linear Filters: Correlation vs. Convolution

Definition: **Correlation**

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

Definition: **Convolution**

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

a	b	c
d	e	f
g	h	i

Filter

1	2	3
4	5	6
7	8	9

Image

Output

$$\begin{aligned} &= 9a + 8b + 7c \\ &\quad + 6d + 5e + 4f \\ &\quad + 3g + 2h + 1i \end{aligned}$$

Linear Filters: Correlation vs. Convolution

Definition: **Correlation**

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

Definition: **Convolution**

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

Filter

(rotated by 180)

!	4	6
7	9	p
c	q	a

a	b	c
d	e	f
g	h	i

Filter

1	2	3
4	5	6
7	8	9

Image

Output

$$\begin{aligned} &= 9a + 8b + 7c \\ &\quad + 6d + 5e + 4f \\ &\quad + 3g + 2h + 1i \end{aligned}$$

Linear Filters: Correlation vs. Convolution

Definition: **Correlation**

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

Definition: **Convolution**

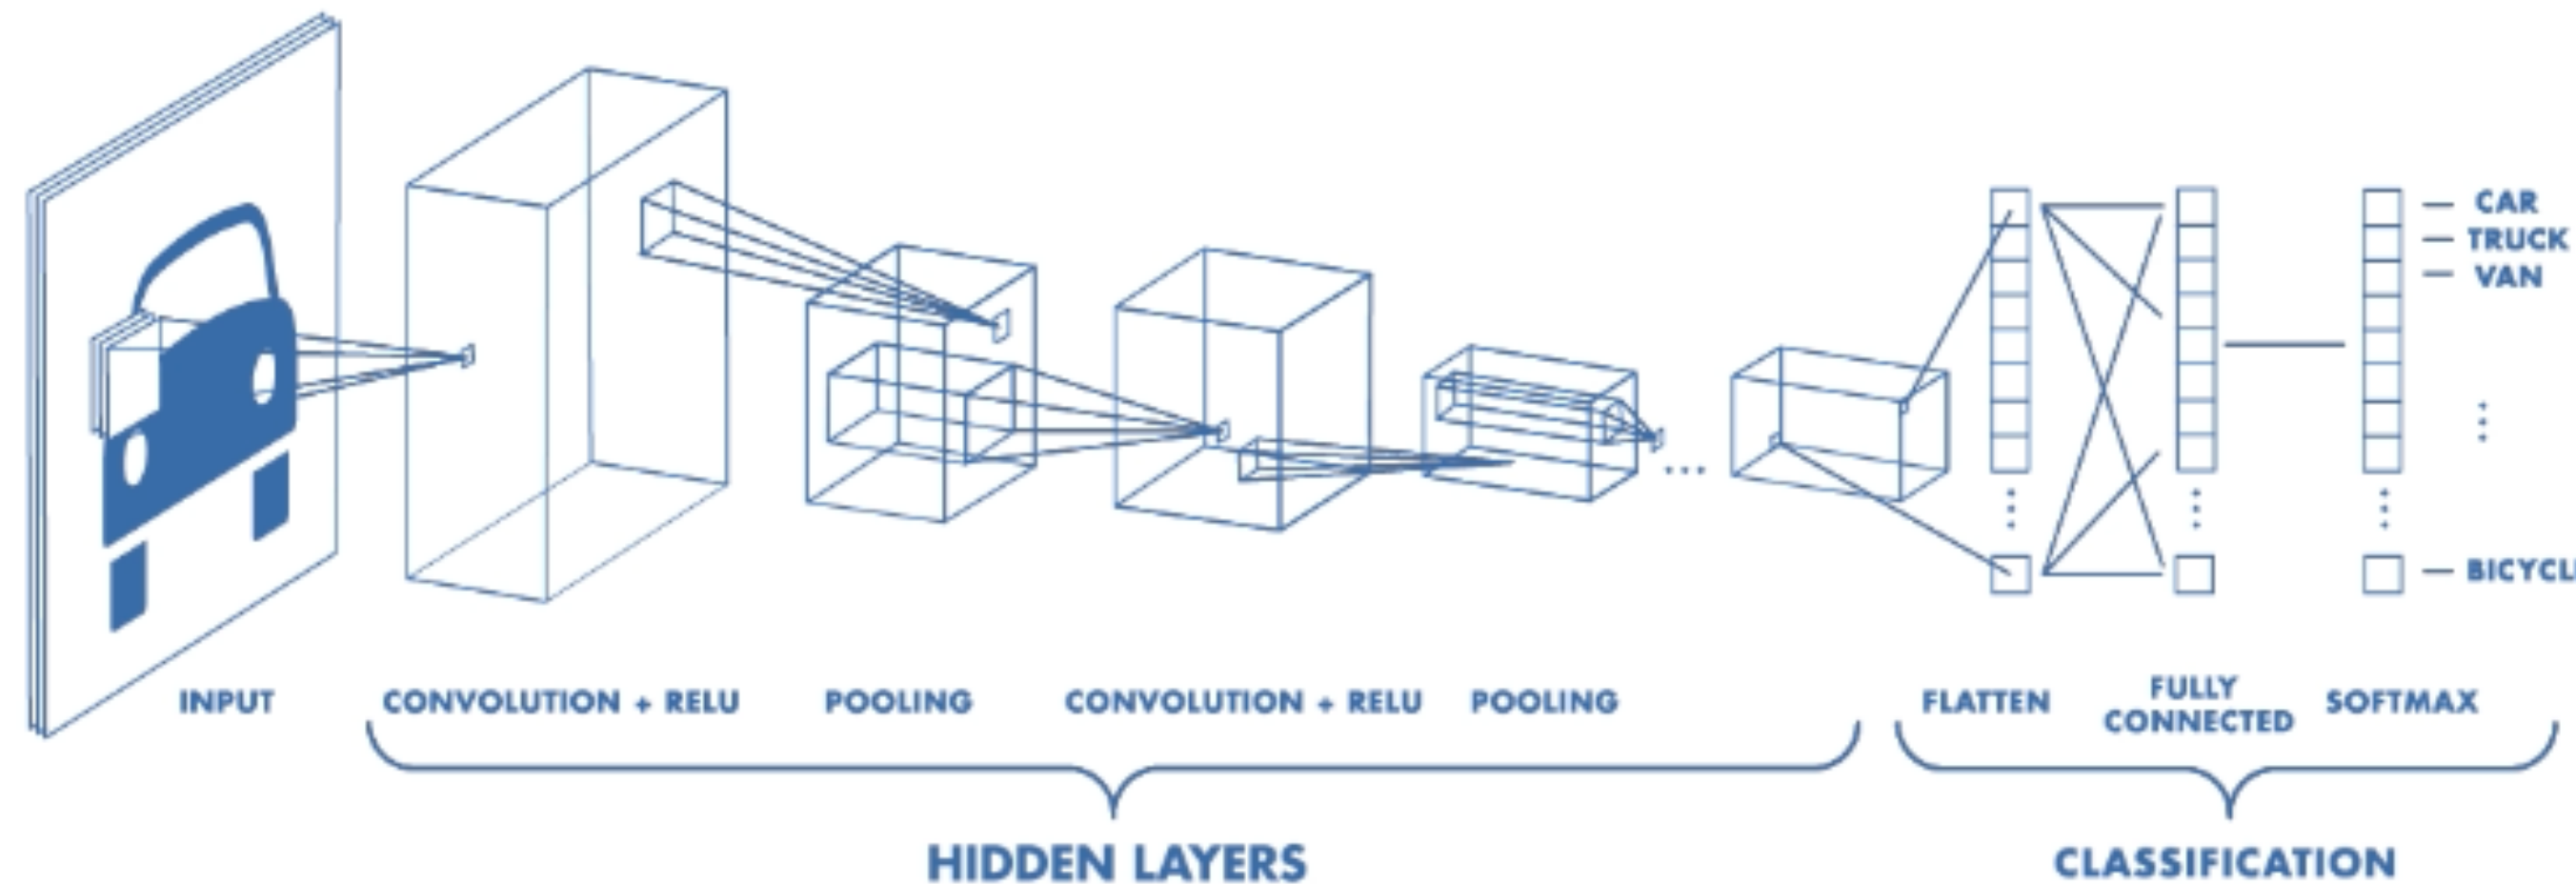
$$\begin{aligned} I'(X, Y) &= \sum_{j=-k}^k \sum_{i=-k}^k F(i, j) I(X - i, Y - j) \\ &= \sum_{j=-k}^k \sum_{i=-k}^k F(-i, -j) I(X + i, Y + j) \end{aligned}$$

Note: if $F(X, Y) = F(-X, -Y)$ then correlation = convolution.

Preview: Why convolutions are important?

Who has heard of **Convolutional Neural Networks** (CNNs)?

What about **Deep Learning**?



Basic operations in CNNs are convolutions (with learned linear filters) followed by non-linear functions.

Note: This results in non-linear filters.

Linear Filters: **Properties**

Let \otimes denote convolution. Let $I(X, Y)$ be a digital image

Superposition: Let F_1 and F_2 be digital filters

$$(F_1 + F_2) \otimes I(X, Y) = F_1 \otimes I(X, Y) + F_2 \otimes I(X, Y)$$

Linear Filters: **Properties**

Let \otimes denote convolution. Let $I(X, Y)$ be a digital image

Superposition: Let F_1 and F_2 be digital filters

$$(F_1 + F_2) \otimes I(X, Y) = F_1 \otimes I(X, Y) + F_2 \otimes I(X, Y)$$

Scaling: Let F be digital filter and let k be a scalar

$$(kF) \otimes I(X, Y) = F \otimes (kI(X, Y)) = k(F \otimes I(X, Y))$$

Linear Filters: **Properties**

Let \otimes denote convolution. Let $I(X, Y)$ be a digital image

Superposition: Let F_1 and F_2 be digital filters

$$(F_1 + F_2) \otimes I(X, Y) = F_1 \otimes I(X, Y) + F_2 \otimes I(X, Y)$$

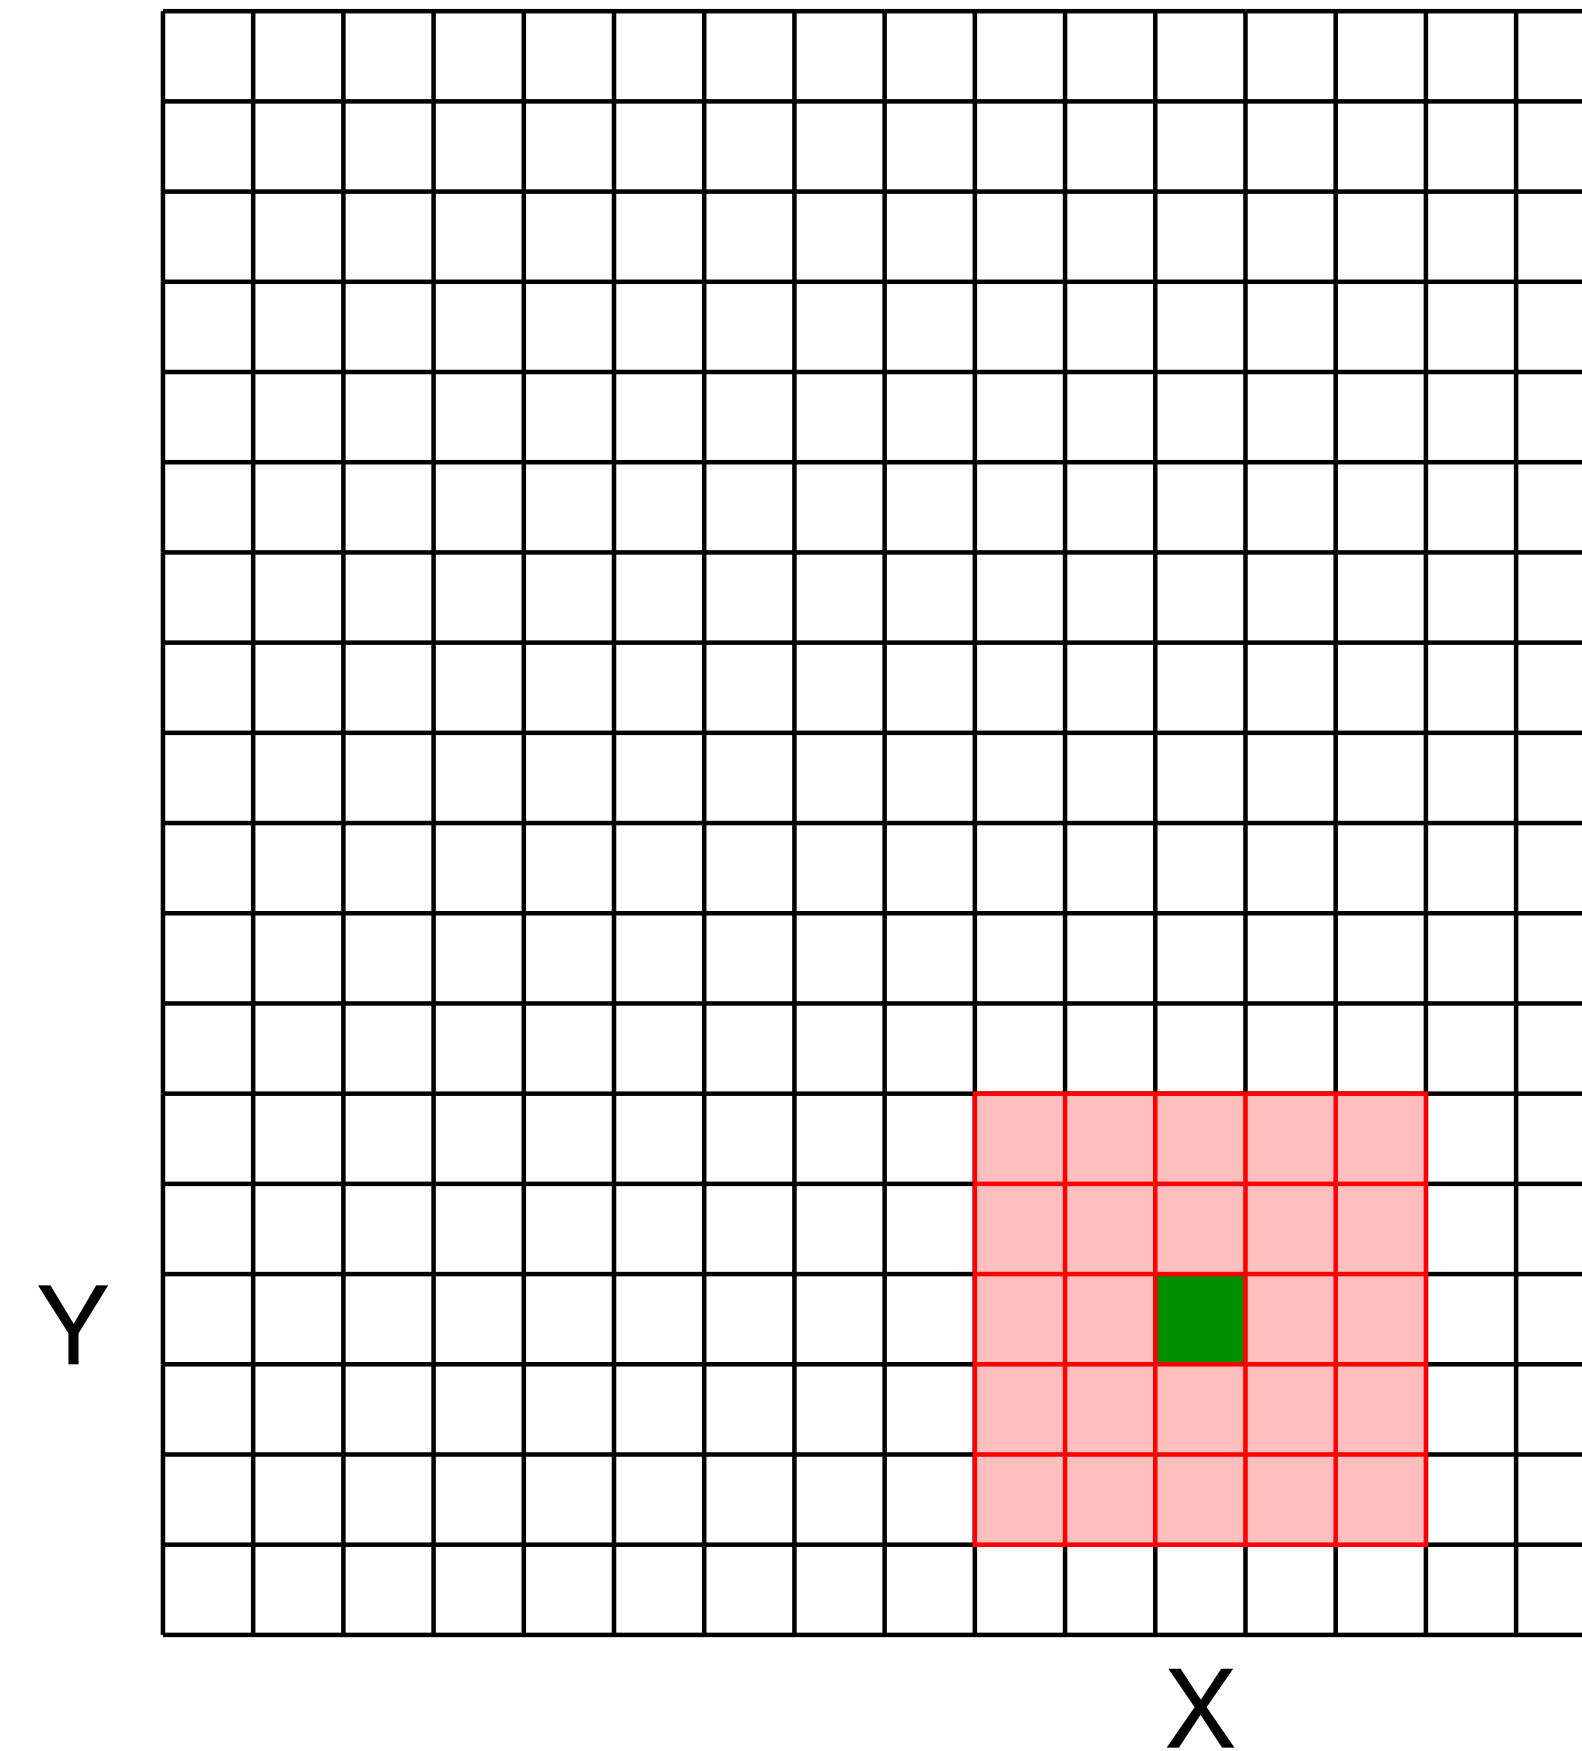
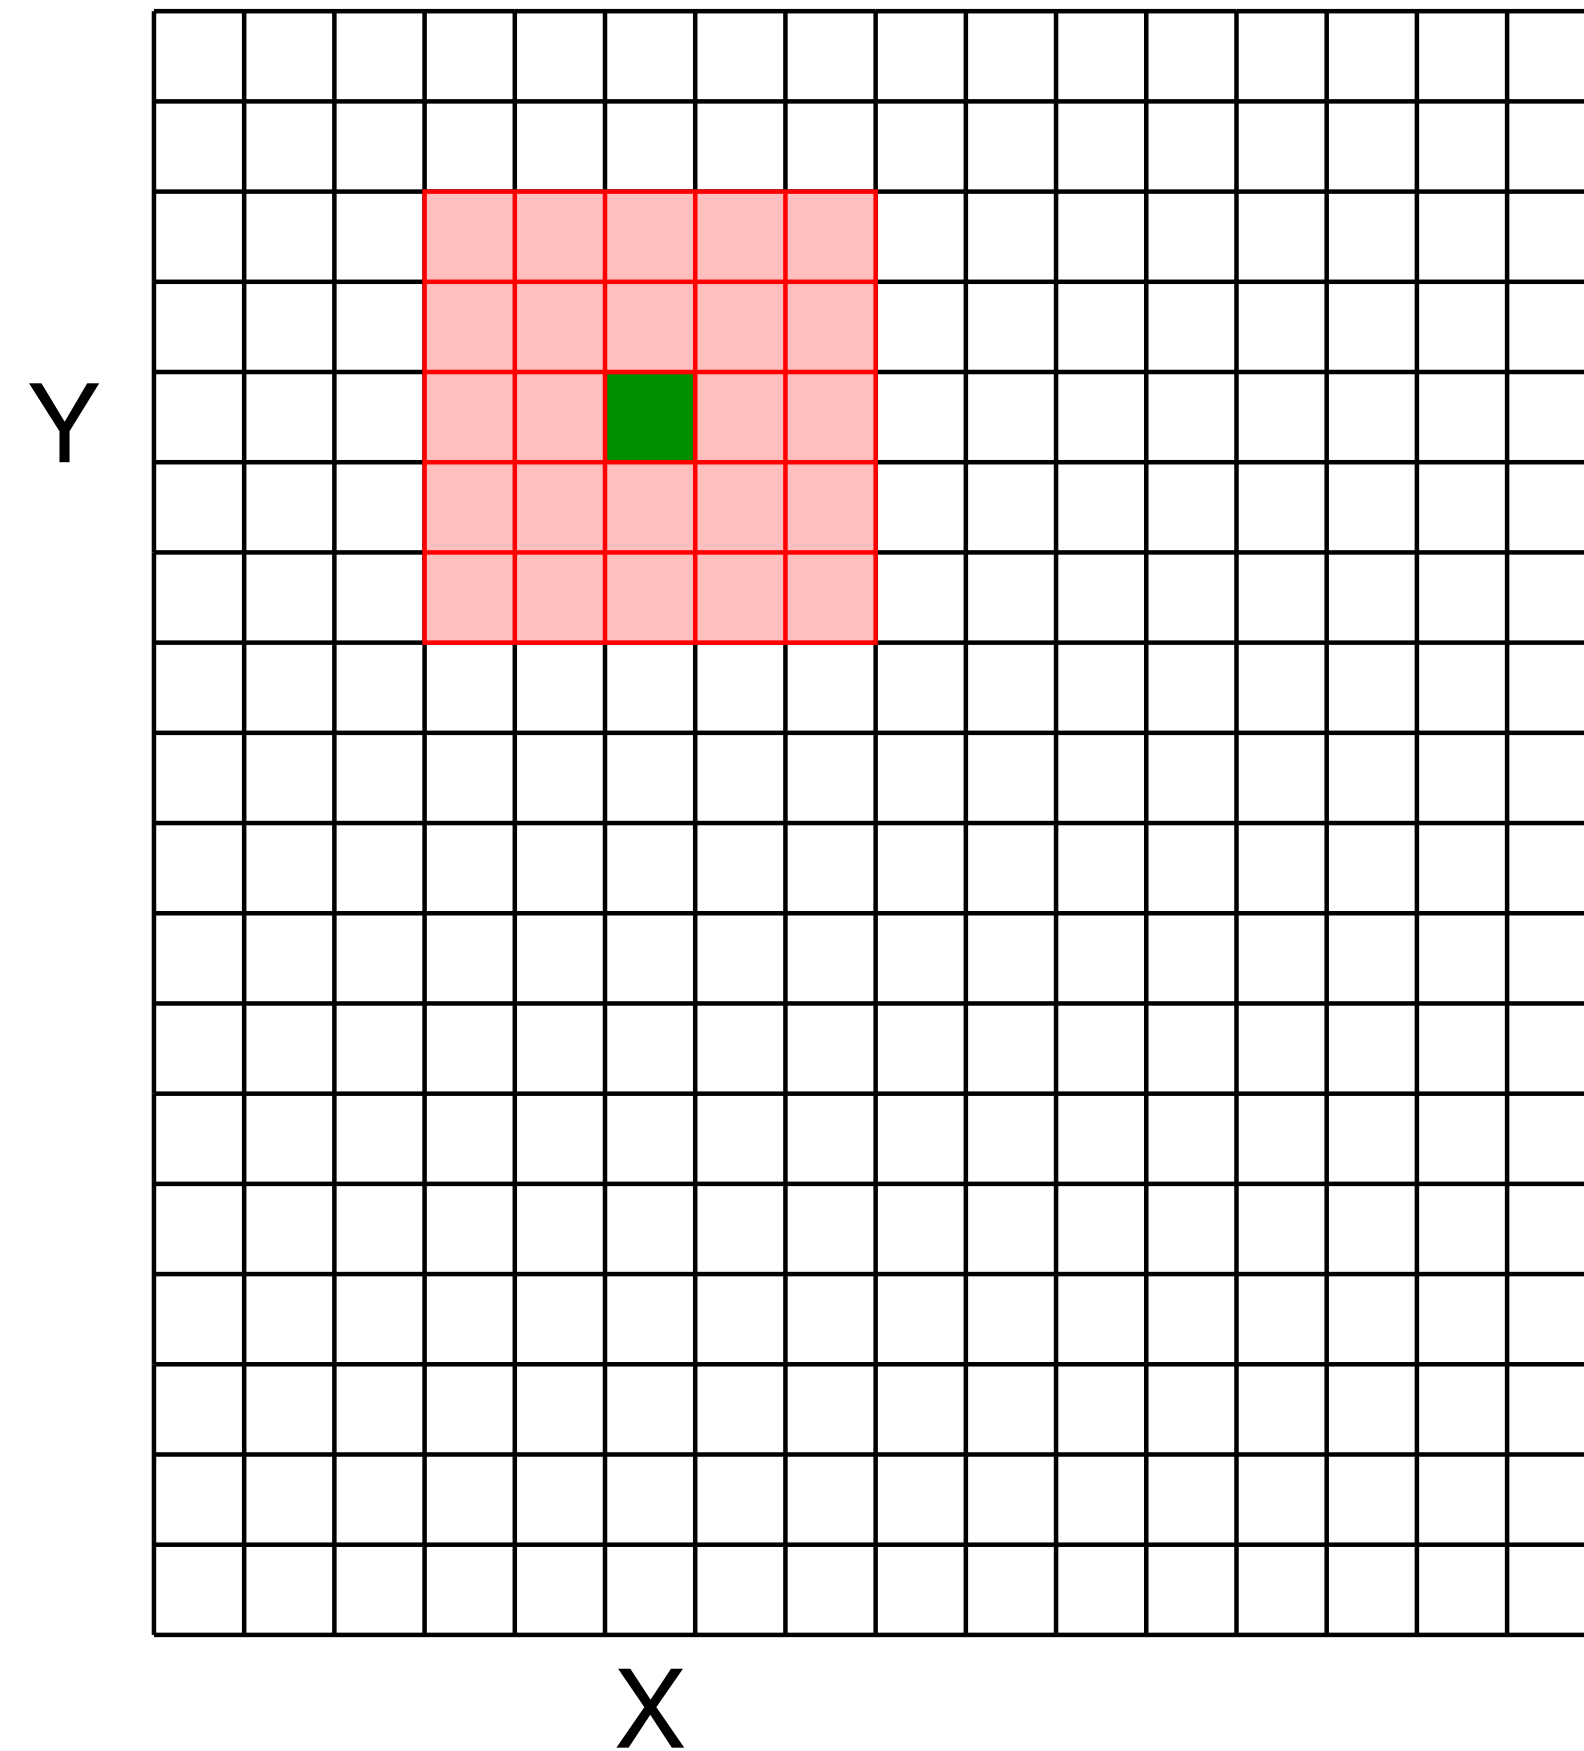
Scaling: Let F be digital filter and let k be a scalar

$$(kF) \otimes I(X, Y) = F \otimes (kI(X, Y)) = k(F \otimes I(X, Y))$$

Shift Invariance: Output is local (i.e., no dependence on absolute position)

Linear Filters: Shift Invariance

Output does **not** depend on absolute position



Linear Filters: **Properties**

Let \otimes denote convolution. Let $I(X, Y)$ be a digital image

Superposition: Let F_1 and F_2 be digital filters

$$(F_1 + F_2) \otimes I(X, Y) = F_1 \otimes I(X, Y) + F_2 \otimes I(X, Y)$$

Scaling: Let F be digital filter and let k be a scalar

$$(kF) \otimes I(X, Y) = F \otimes (kI(X, Y)) = k(F \otimes I(X, Y))$$

Shift Invariance: Output is local (i.e., no dependence on absolute position)

An operation is **linear** if it satisfies both **superposition** and **scaling**

Linear Systems: Characterization Theorem

Any linear, shift invariant operation can be expressed as convolution

Example 5: Smoothing with a Box Filter

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



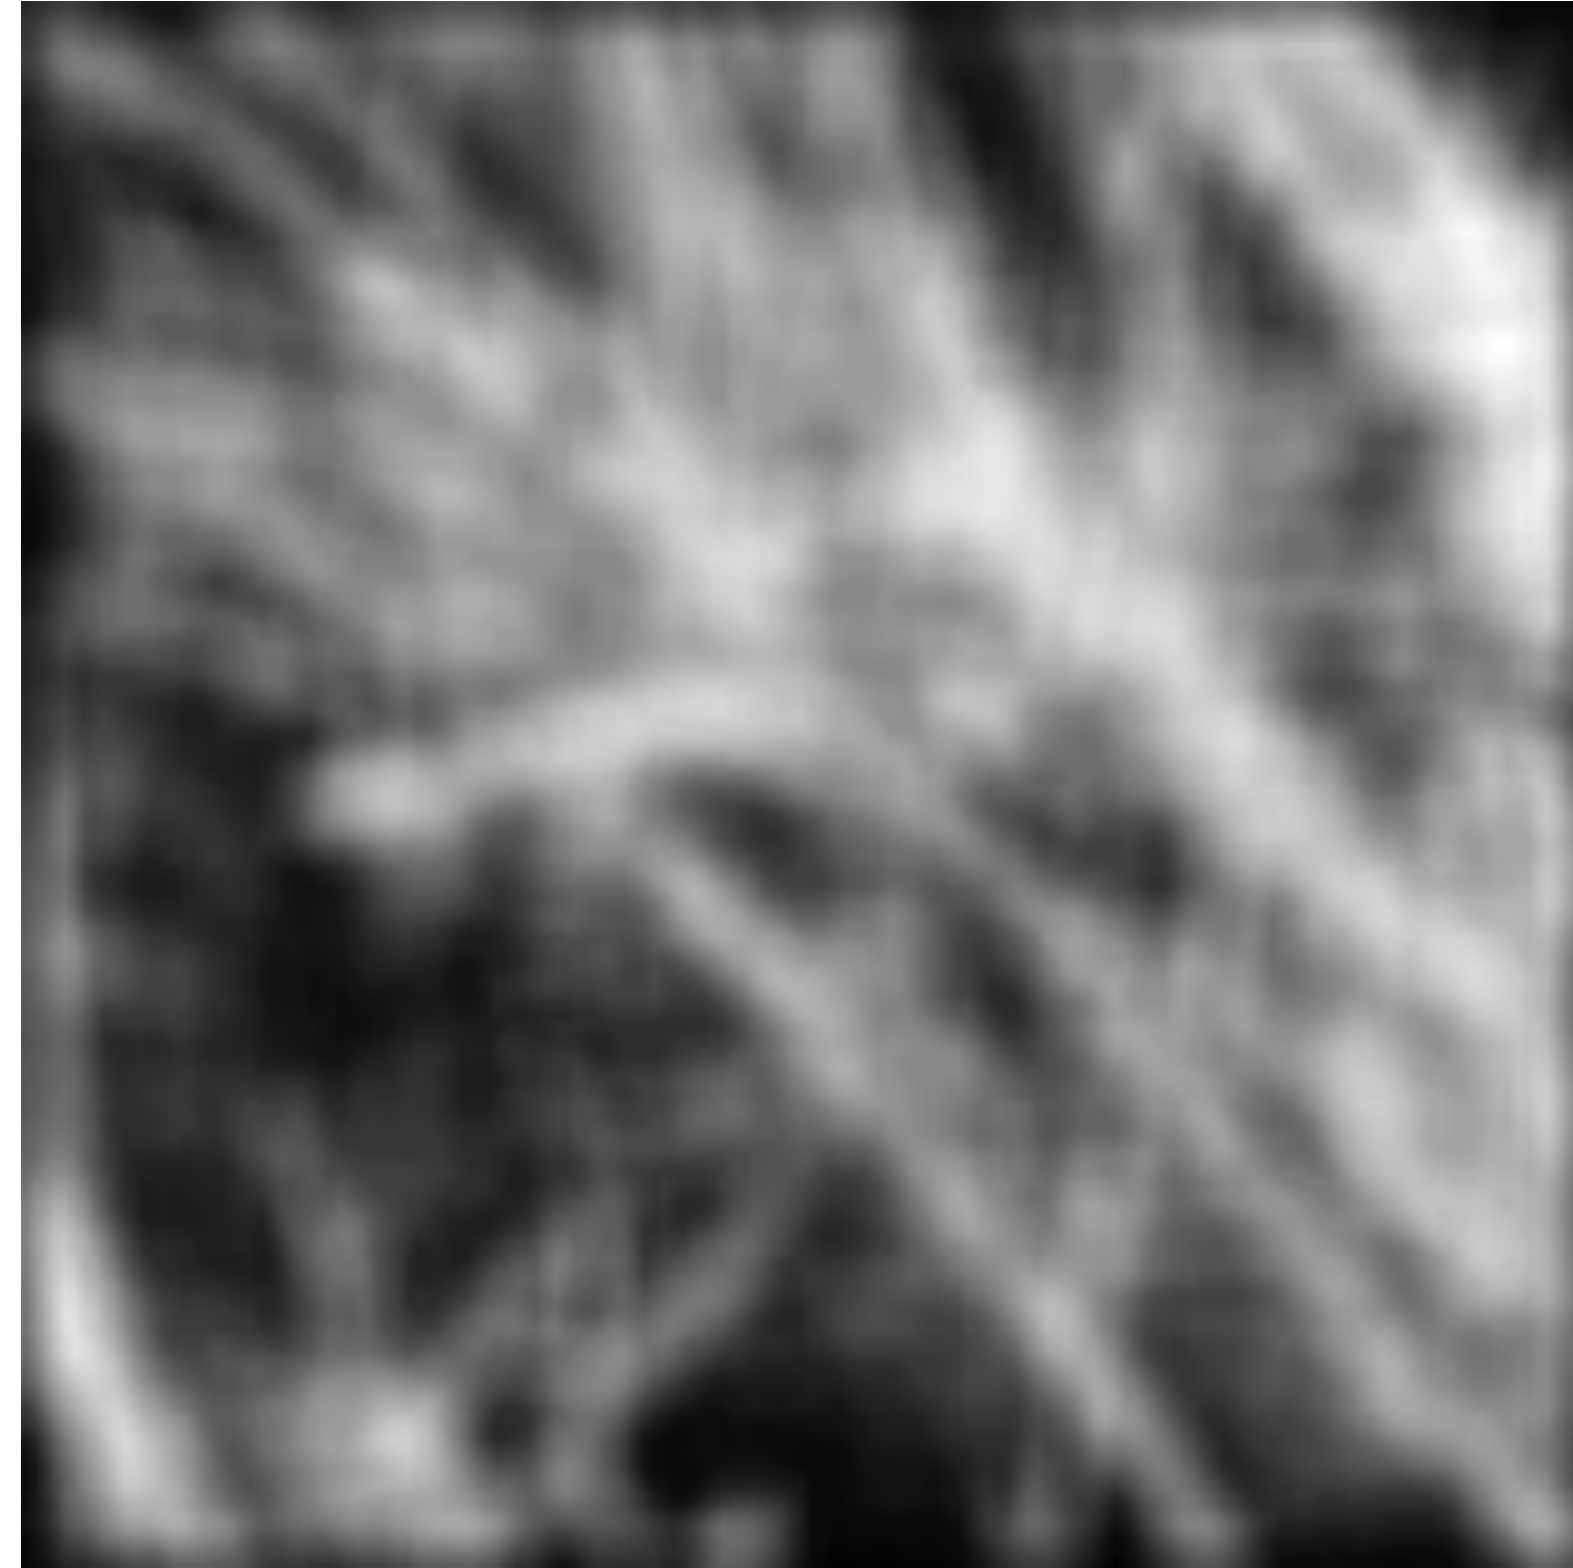
Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

Filter has equal positive values that sum up to 1

Replaces each pixel with the average of itself and its local neighborhood

— Box filter is also referred to as **average filter** or **mean filter**

Example 5: Smoothing with a Box Filter

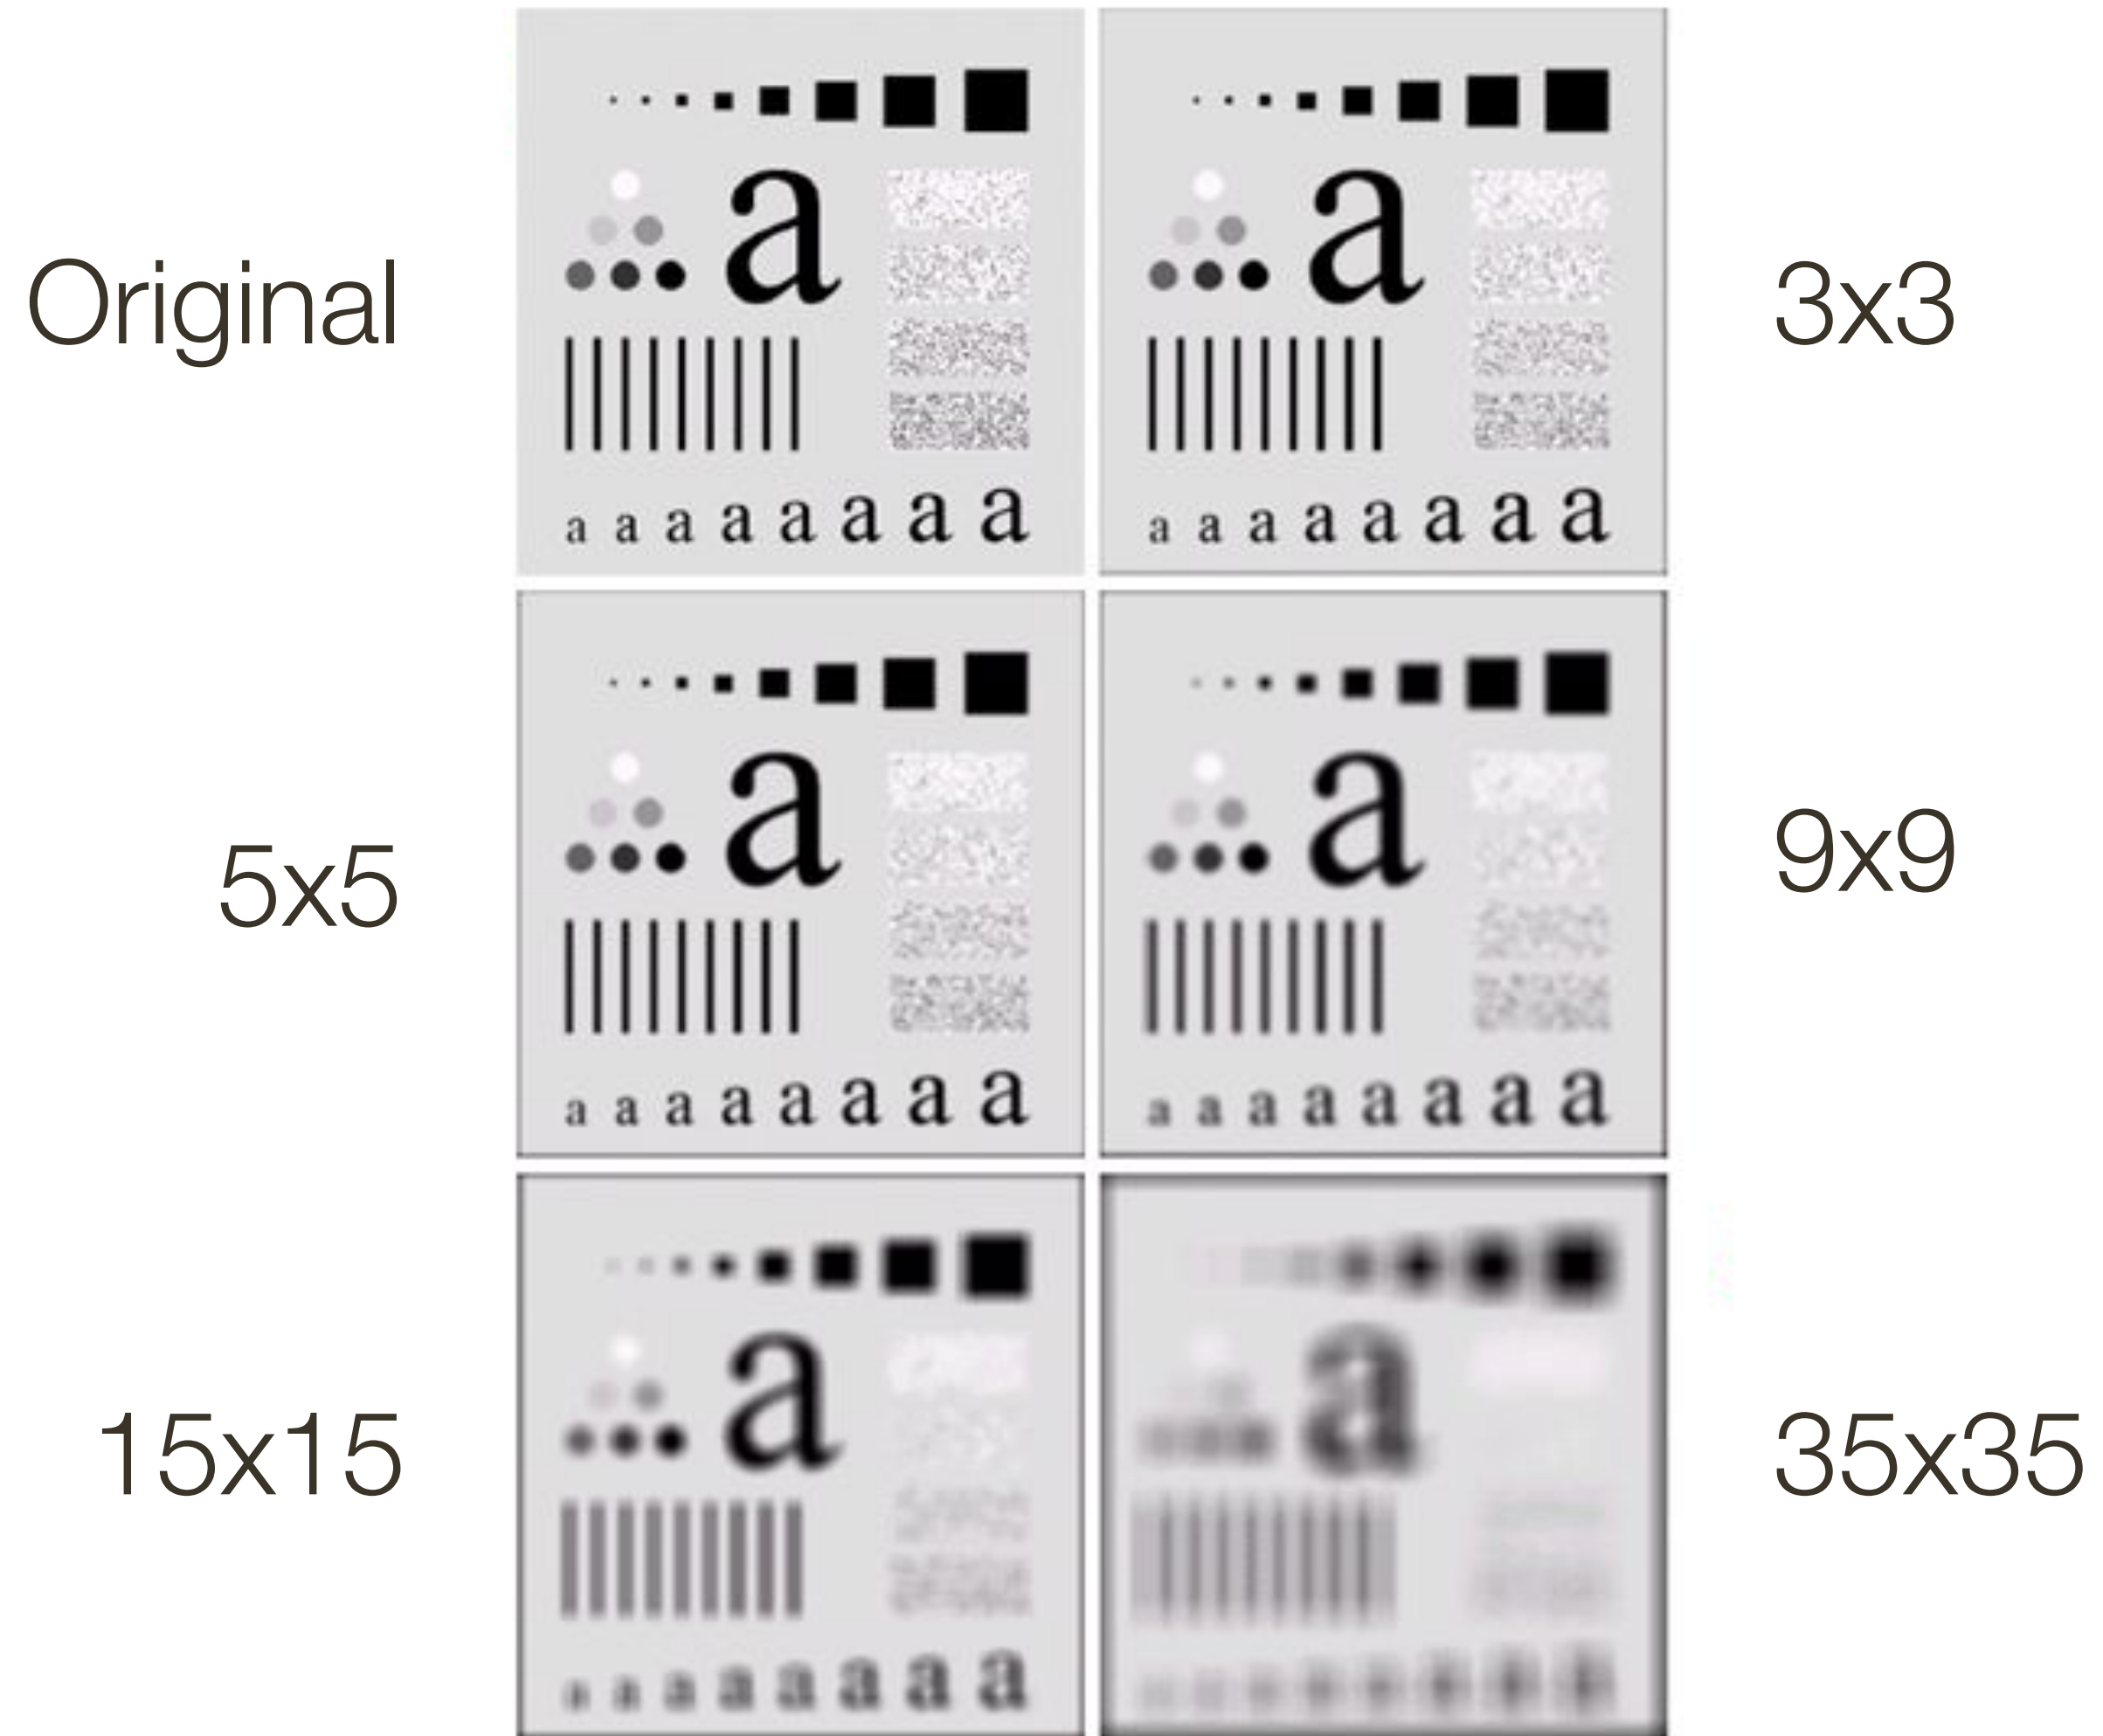


Forsyth & Ponce (2nd ed.) Figure 4.1 (left and middle)

Example 5: Smoothing with a Box Filter

What happens if we increase the width (size) of the box filter?

Example 5: Smoothing with a Box Filter



Gonzales & Woods (3rd ed.) Figure 3.3

Example 6: Smoothing with a Gaussian

Smoothing with a box **doesn't model lens defocus** well

- Smoothing with a box filter depends on direction
- Image in which the center point is 1 and every other point is 0

Example 6: Smoothing with a Gaussian

Smoothing with a box **doesn't model lens defocus** well

- Smoothing with a box filter depends on direction
- Image in which the center point is 1 and every other point is 0

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

1	1	1
1	1	1
1	1	1

Filter

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

Image

Example 6: Smoothing with a Gaussian

Smoothing with a box **doesn't model lens defocus** well

- Smoothing with a box filter depends on direction
- Image in which the center point is 1 and every other point is 0

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

1	1	1
1	1	1
1	1	1

Filter

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

Image

0	0	0	0	0
0	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	0
0	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	0
0	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	0
0	0	0	0	0

Result

Example 6: Smoothing with a Gaussian

Smoothing with a box **doesn't model lens defocus** well

- Smoothing with a box filter depends on direction
- Image in which the center point is 1 and every other point is 0

Smoothing with a (circular) **pillbox** is a better model for defocus (in geometric optics)

The **Gaussian** is a good general smoothing model

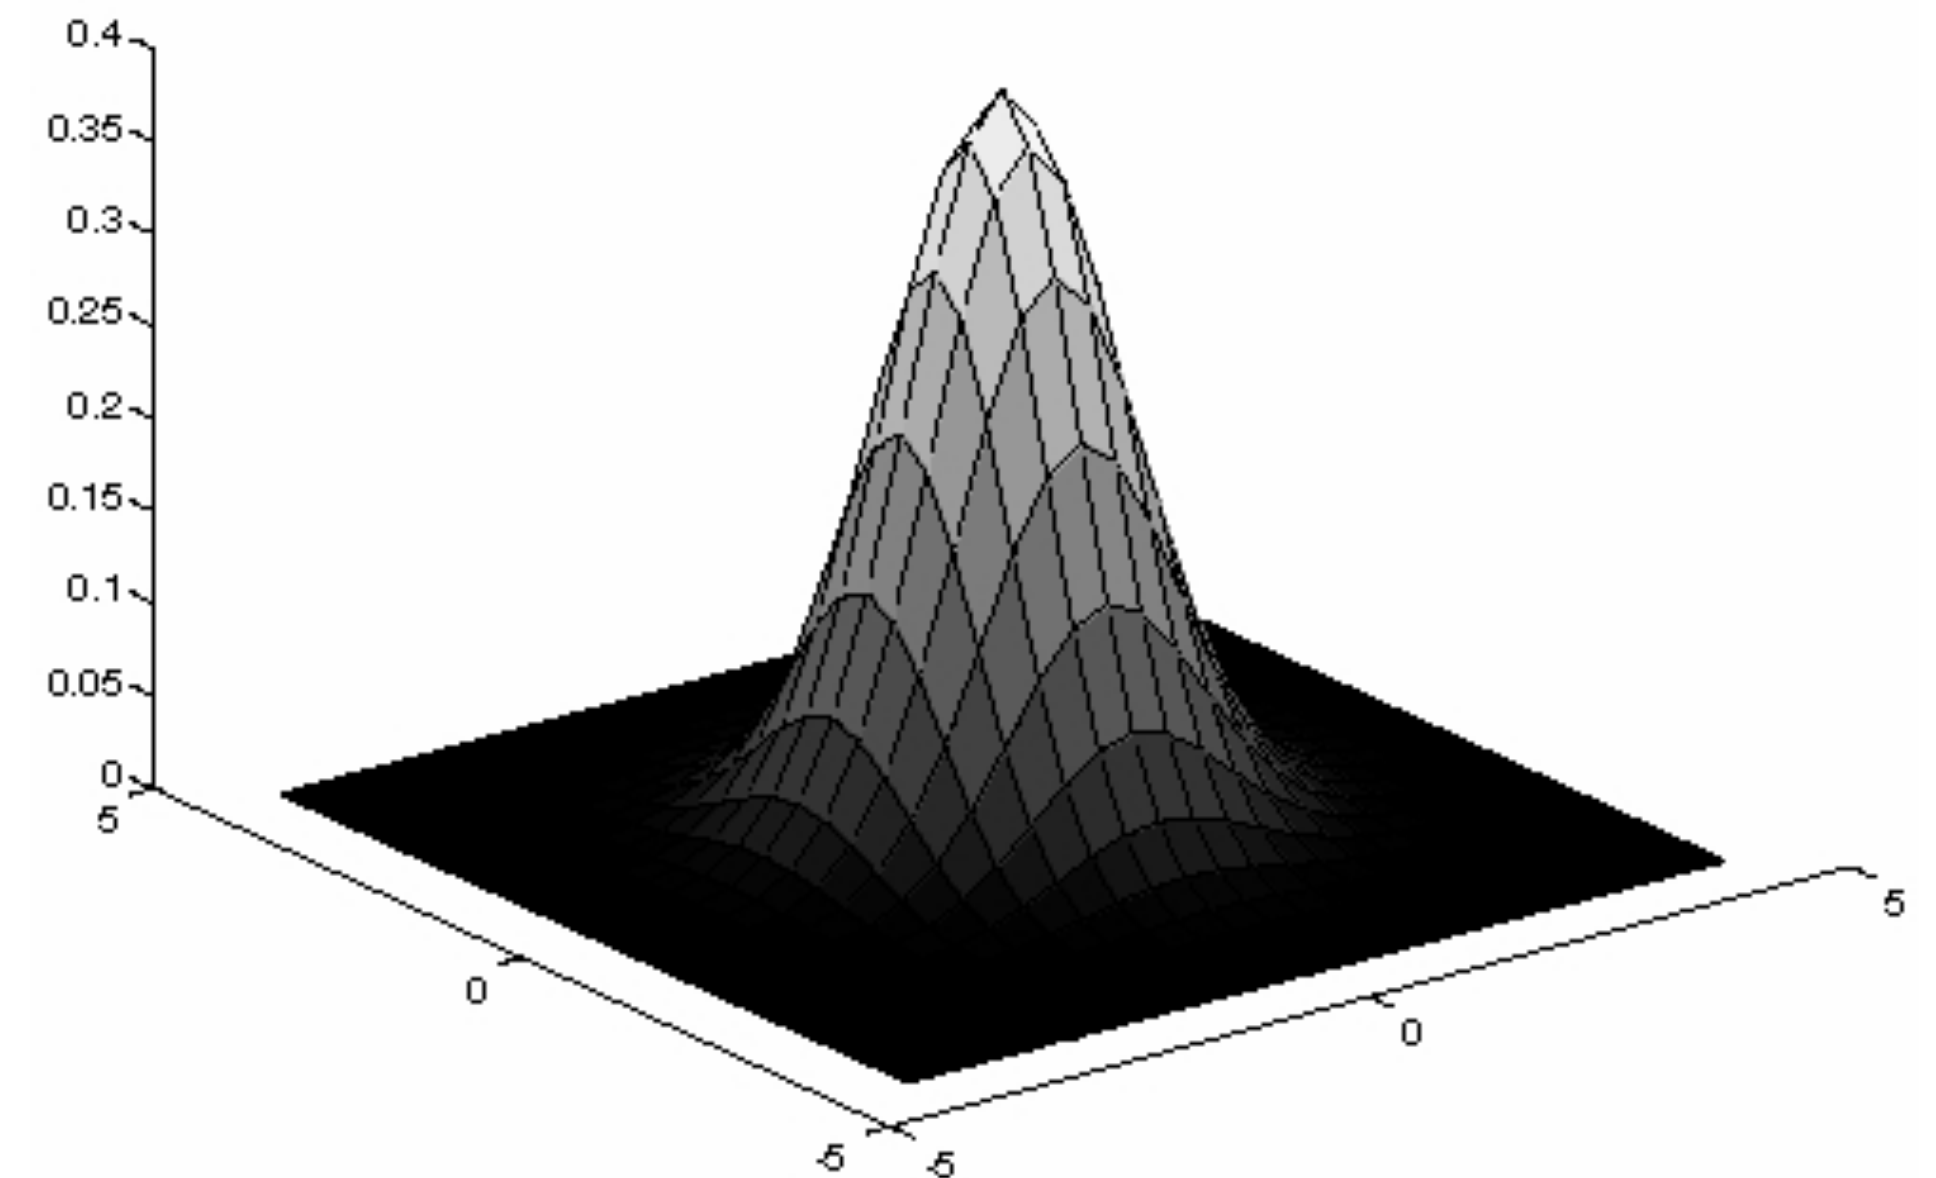
- for phenomena (that are the sum of other small effects)
- whenever the Central Limit Theorem applies

Example 6: Smoothing with a Gaussian

Idea: Weight contributions of pixels by spatial proximity (nearness)

2D **Gaussian** (continuous case):

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



Forsyth & Ponce (2nd ed.)

Figure 4.2

Summary

- The **correlation** of $F(X, Y)$ and $I(X, Y)$ is:

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

- **Visual interpretation:** Superimpose the filter F on the image I at (X, Y) , perform an element-wise multiply, and sum up the values
- **Convolution** is like **correlation** except filter “flipped”
 - if $F(X, Y) = F(-X, -Y)$ then correlation = convolution.
- **Characterization Theorem:** Any linear, spatially invariant operation can be expressed as a convolution