# Deep Learning (1470)

Randall Balestriero

Class 7: Convolutions

# Multilayer Perceptrons

## And why we need better

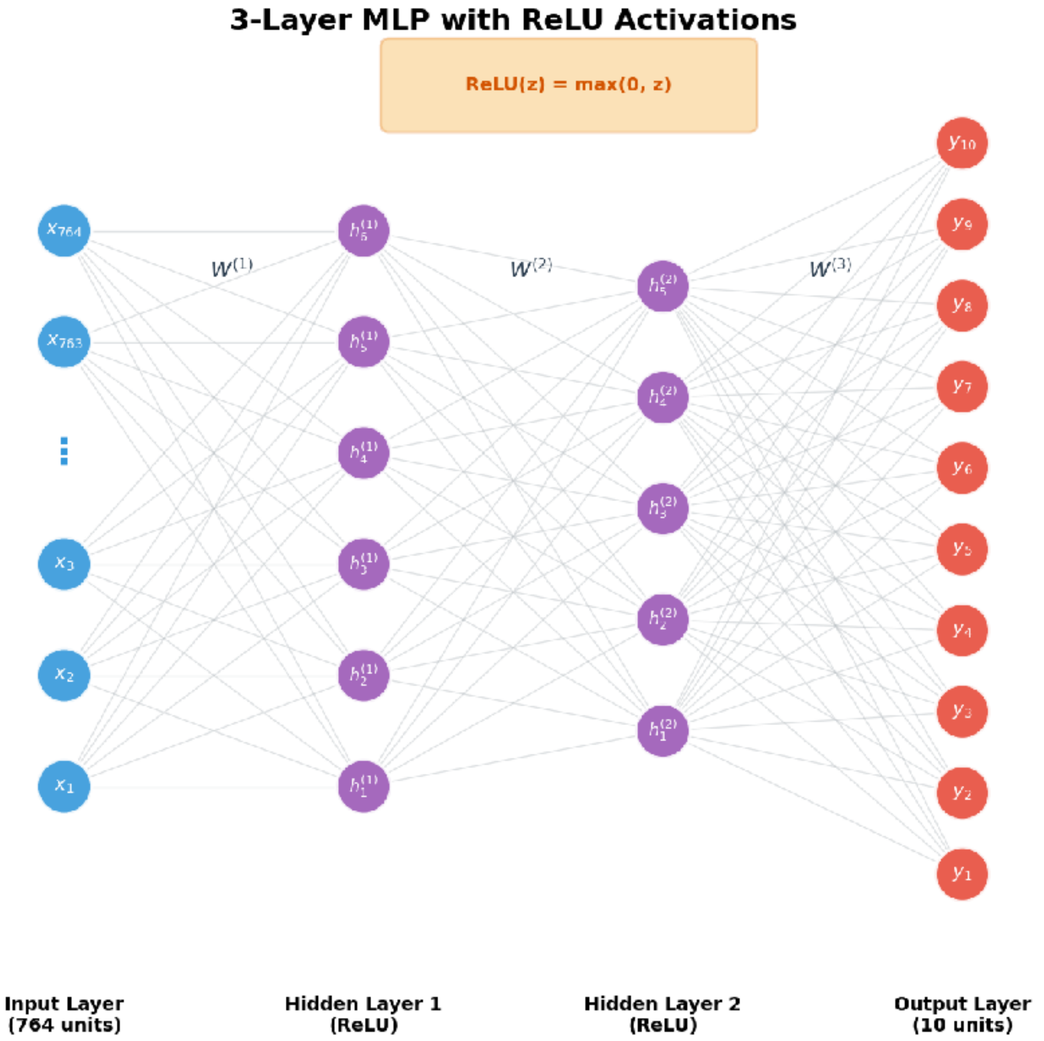

# Multilayer Perceptrons

## And why we need better

Extracting local features
at each layer is enough!

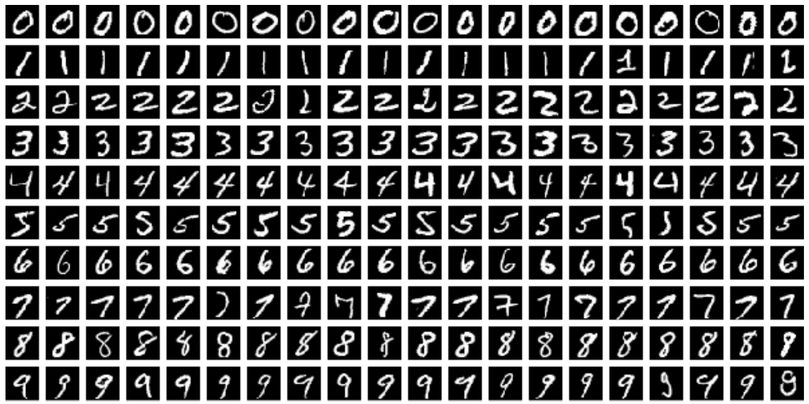

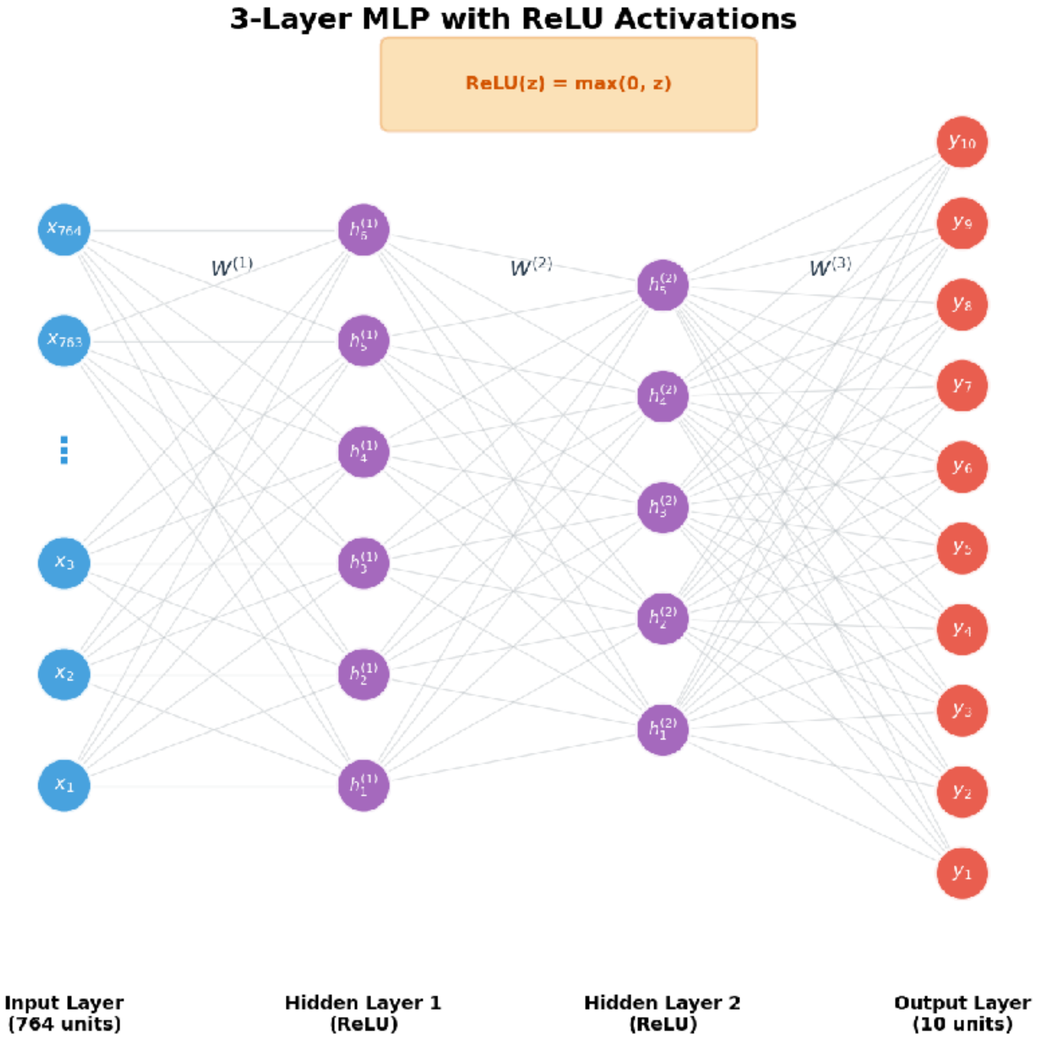

# Multilayer Perceptrons

## And why we need better

Extracting local features
at each layer is enough!

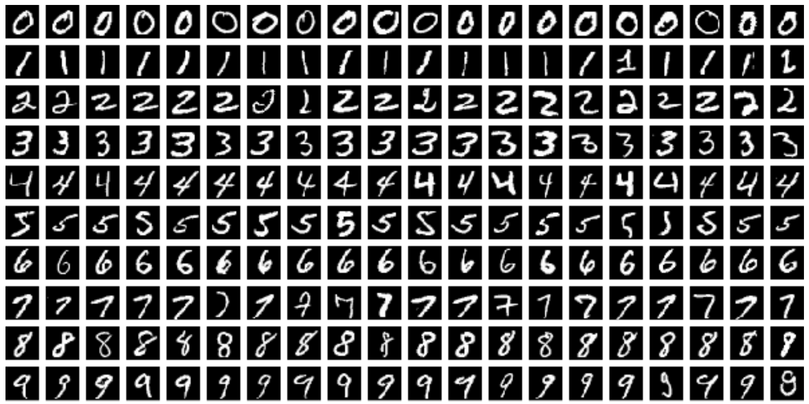

Not only local features… but
translation invariant features!

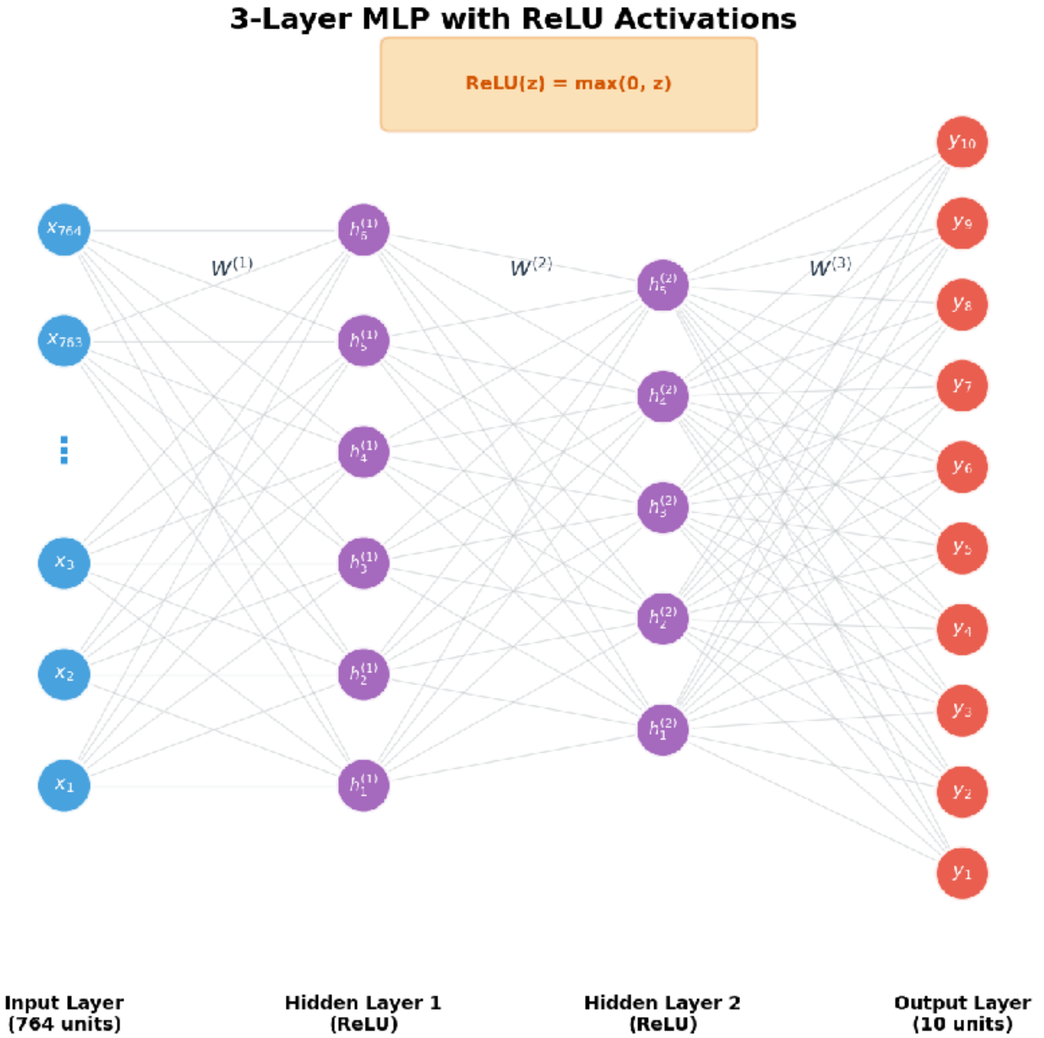

# The Main Building Block: Convolution

- A convolution is a linear operator
- Convolution is an operation that takes two inputs

(1) An image (2D – B/W)

(2) A filter (also called a kernel)

| | | |
|---|---|---|
| 1 | 1 | 1 |
| 0 | 0 | 0 |
| −1 | −1 | −1 |

2D array of numbers; could be any values

# What Convolution Does (Visually)

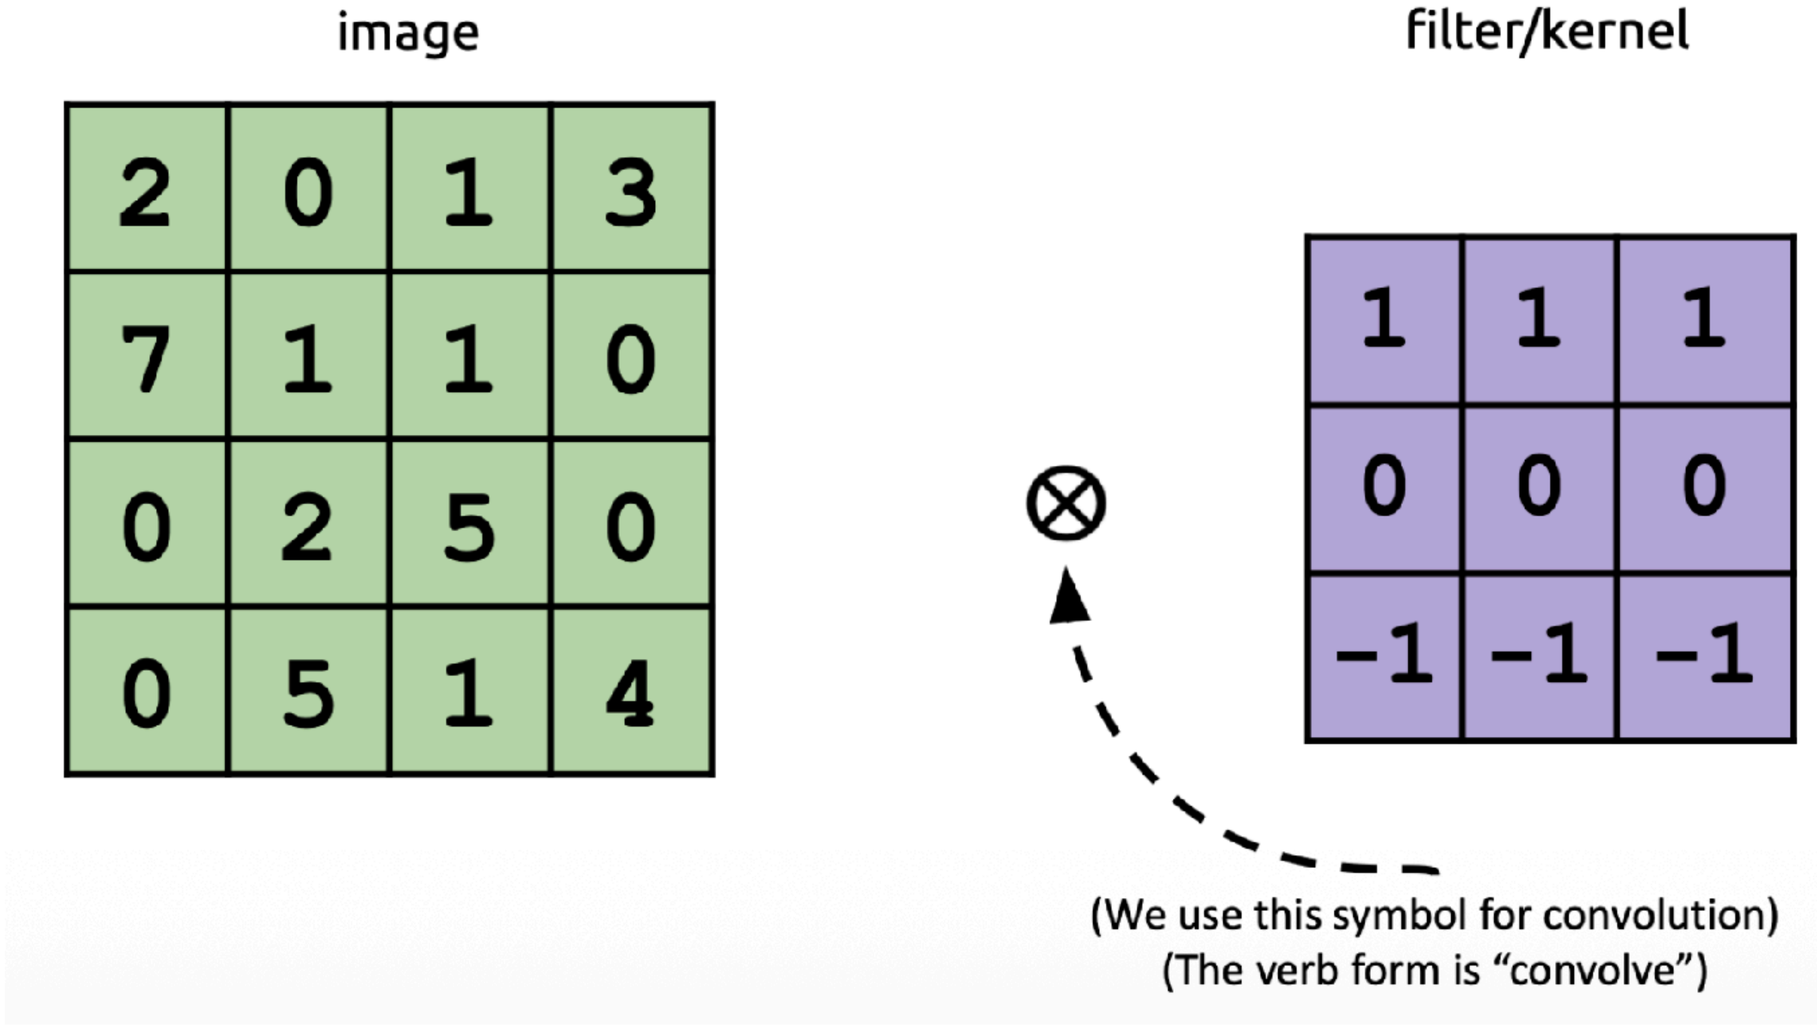

# What Convolution Does (Visually)

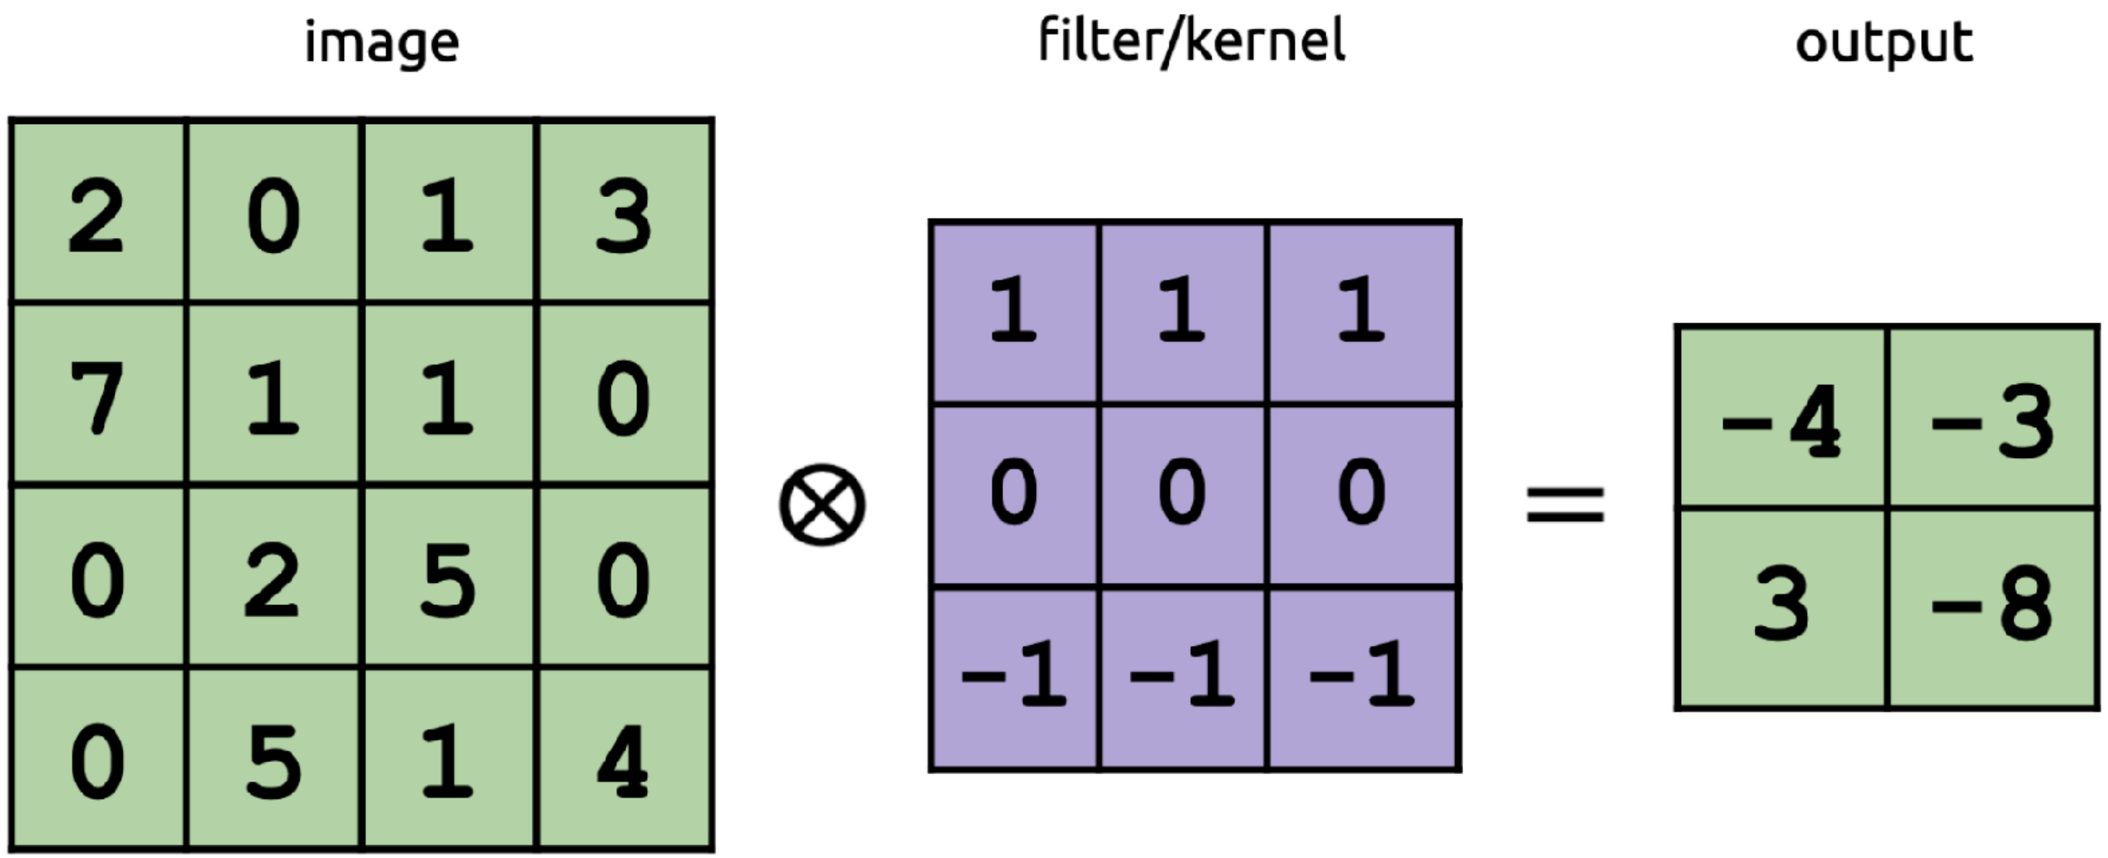

# What Convolution Does (Visually)

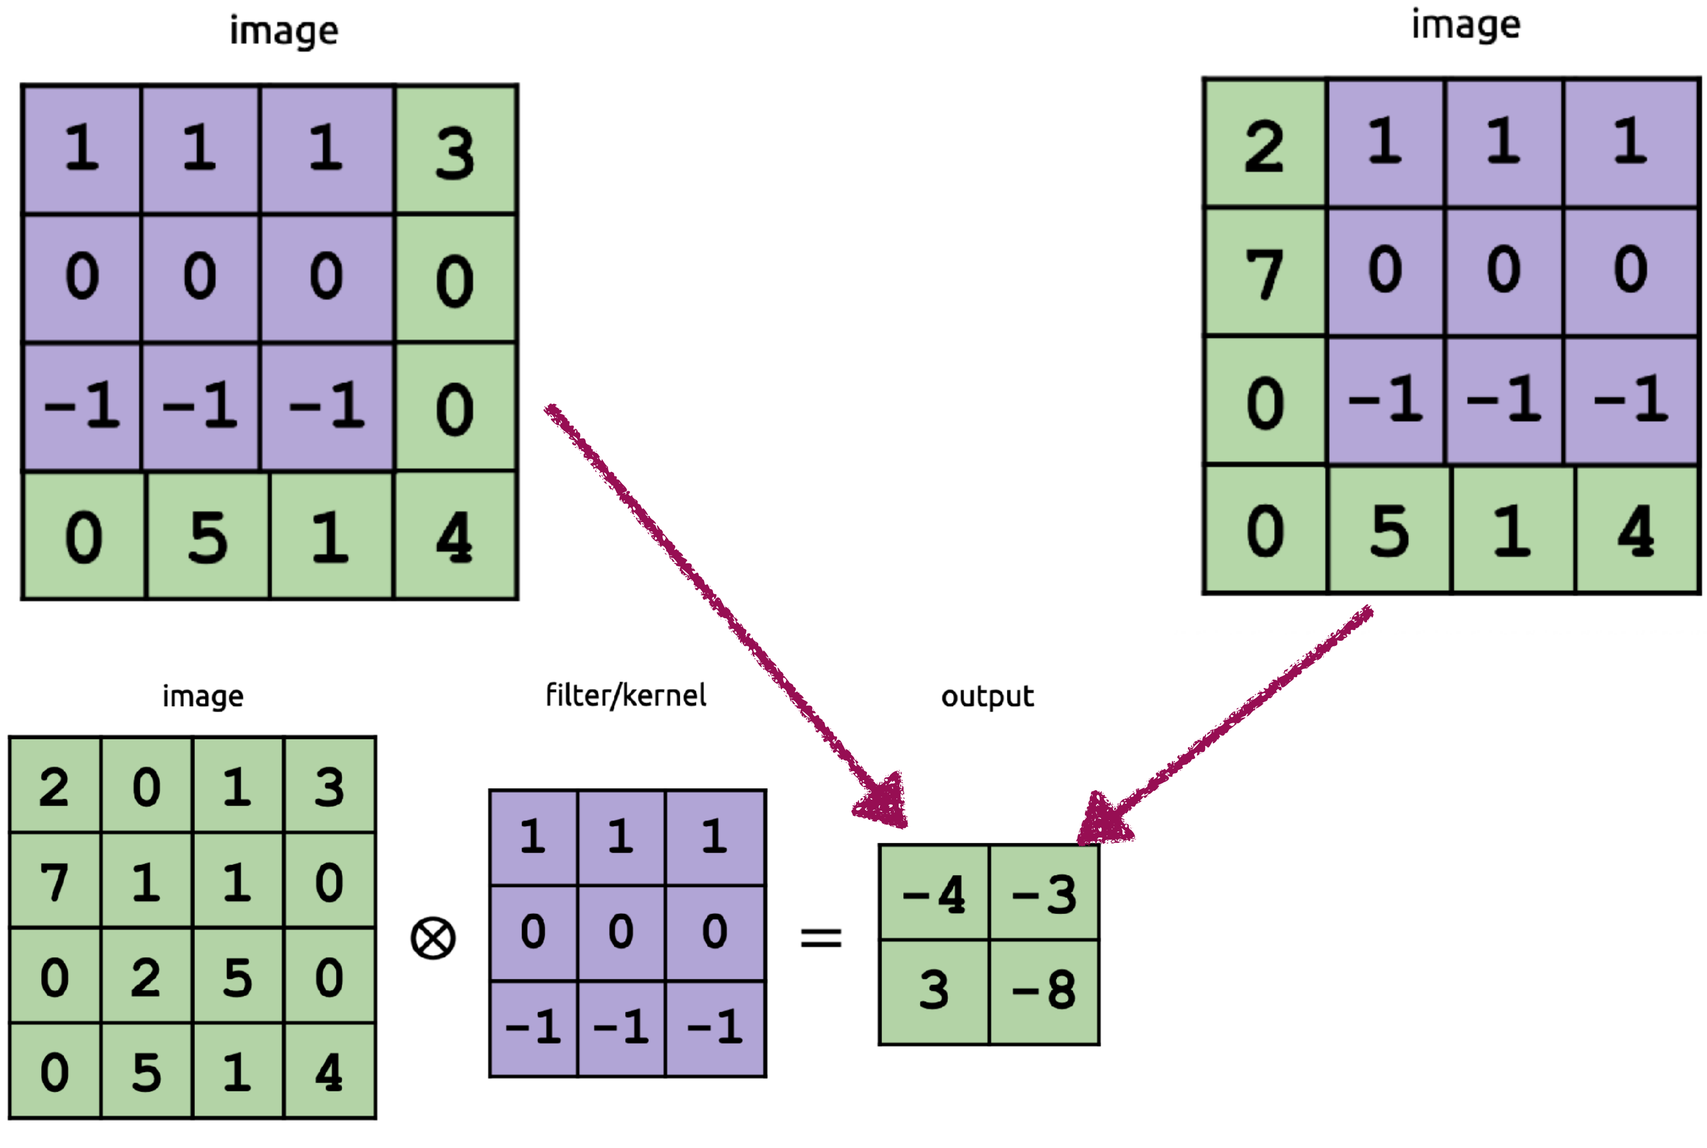

# What Convolution Does (Visually)

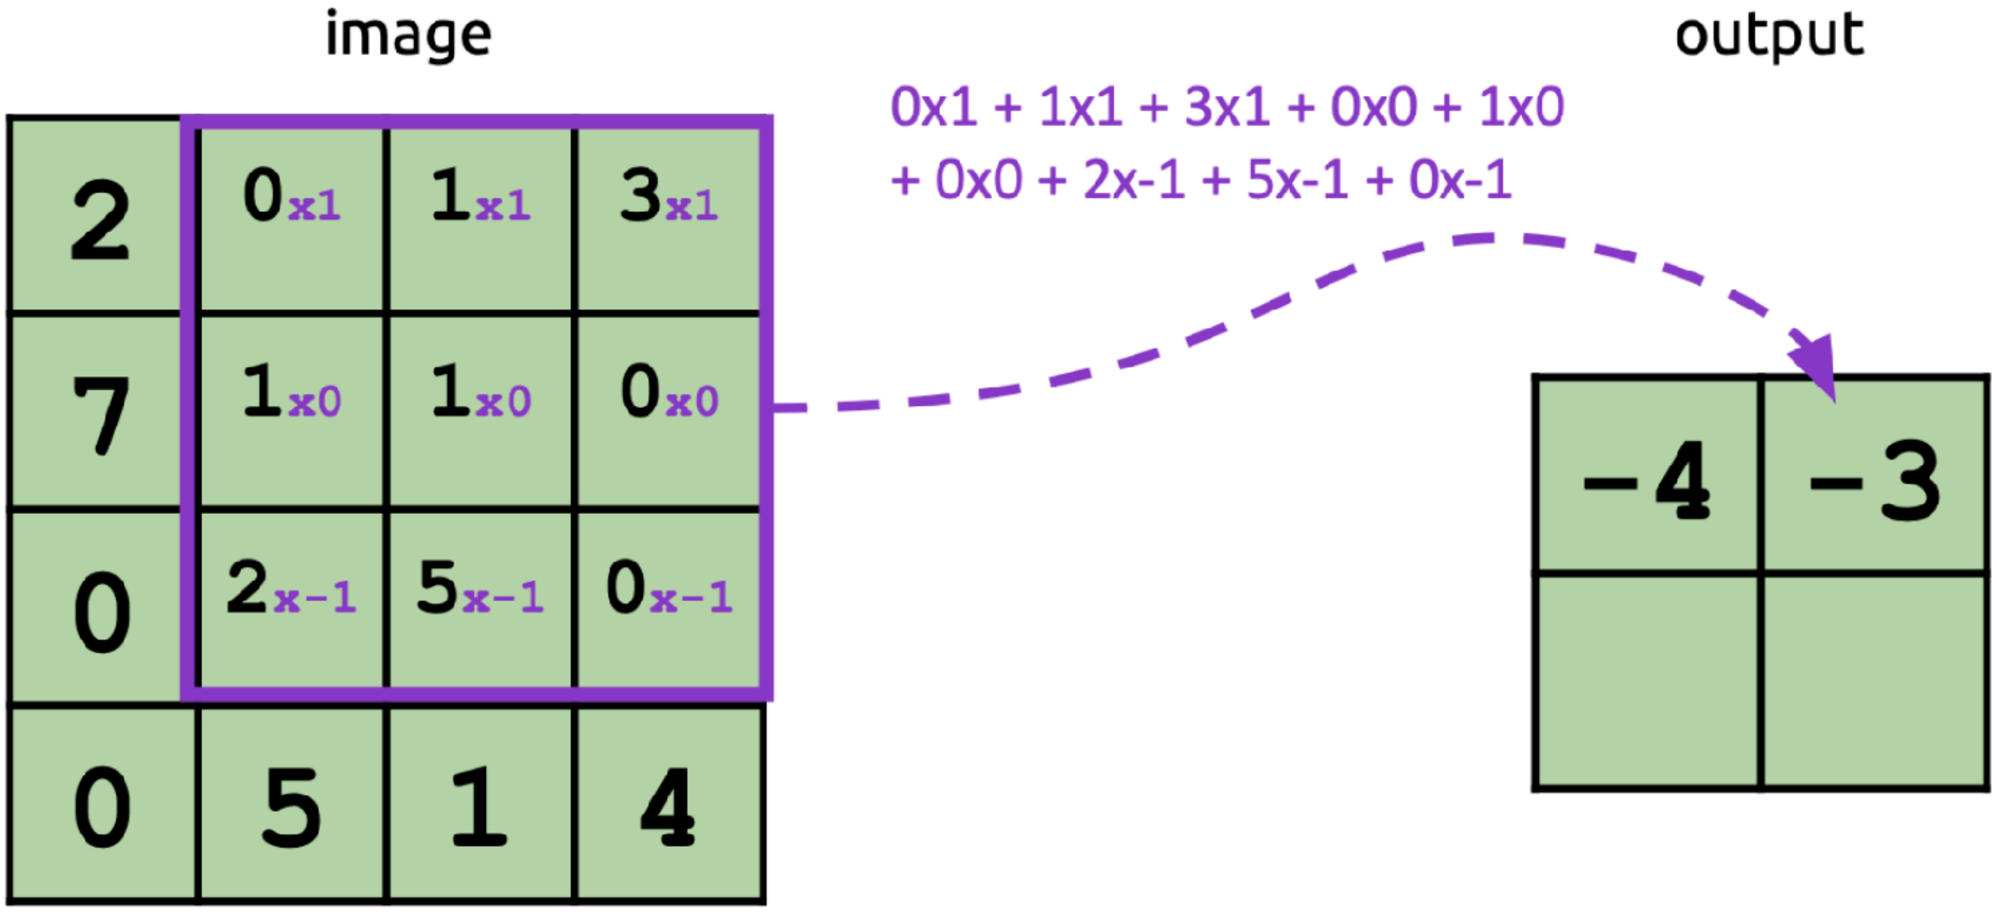

# Example

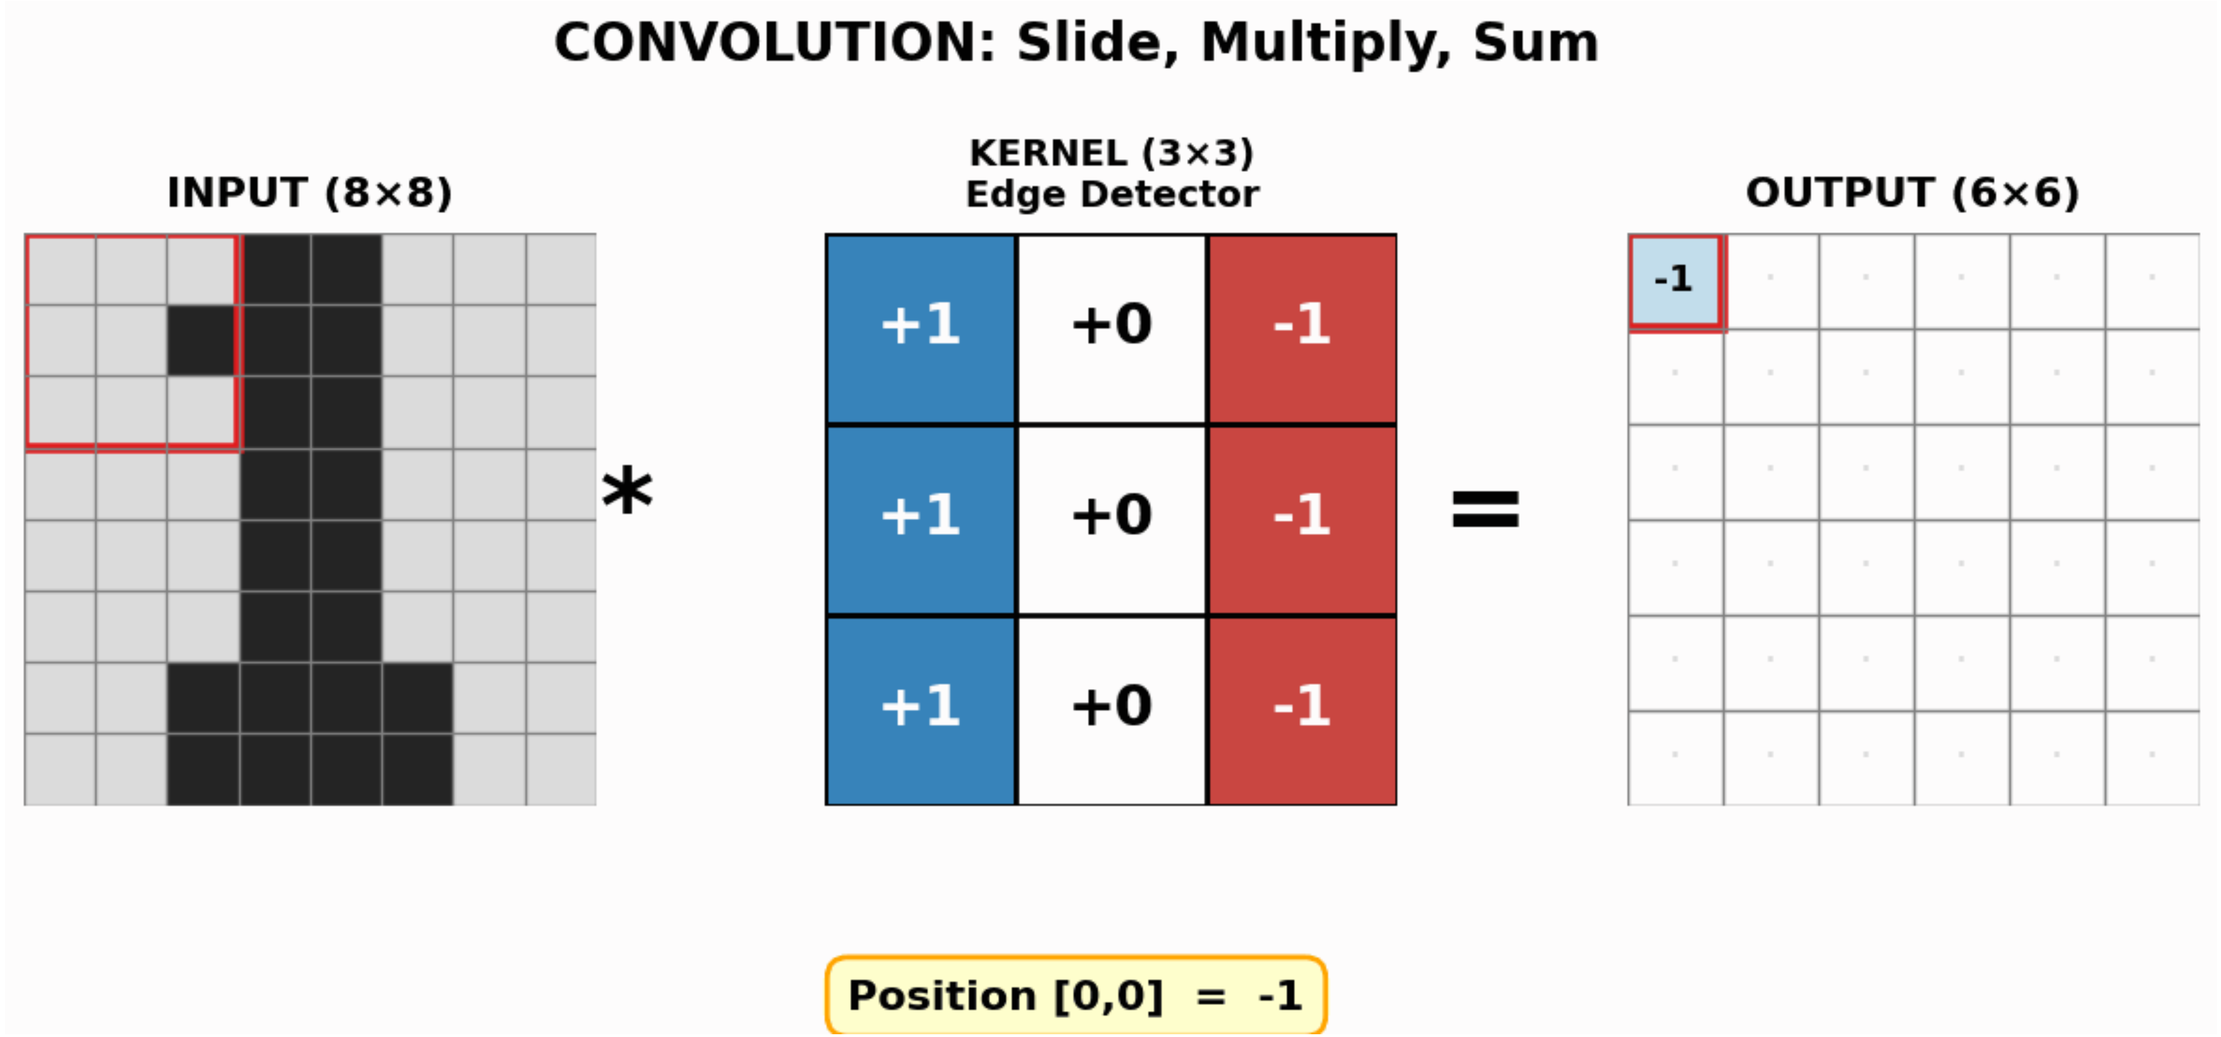
CONVOLUTION: Slide, Multiply, Sum
INPUT (8×8)
KERNEL (3×3)
Edge Detector
OUTPUT (6×6)
+1
+0
-1
+1
+0
-1
+1
+0
-1
-1
Position [0,0] = -1

# Example

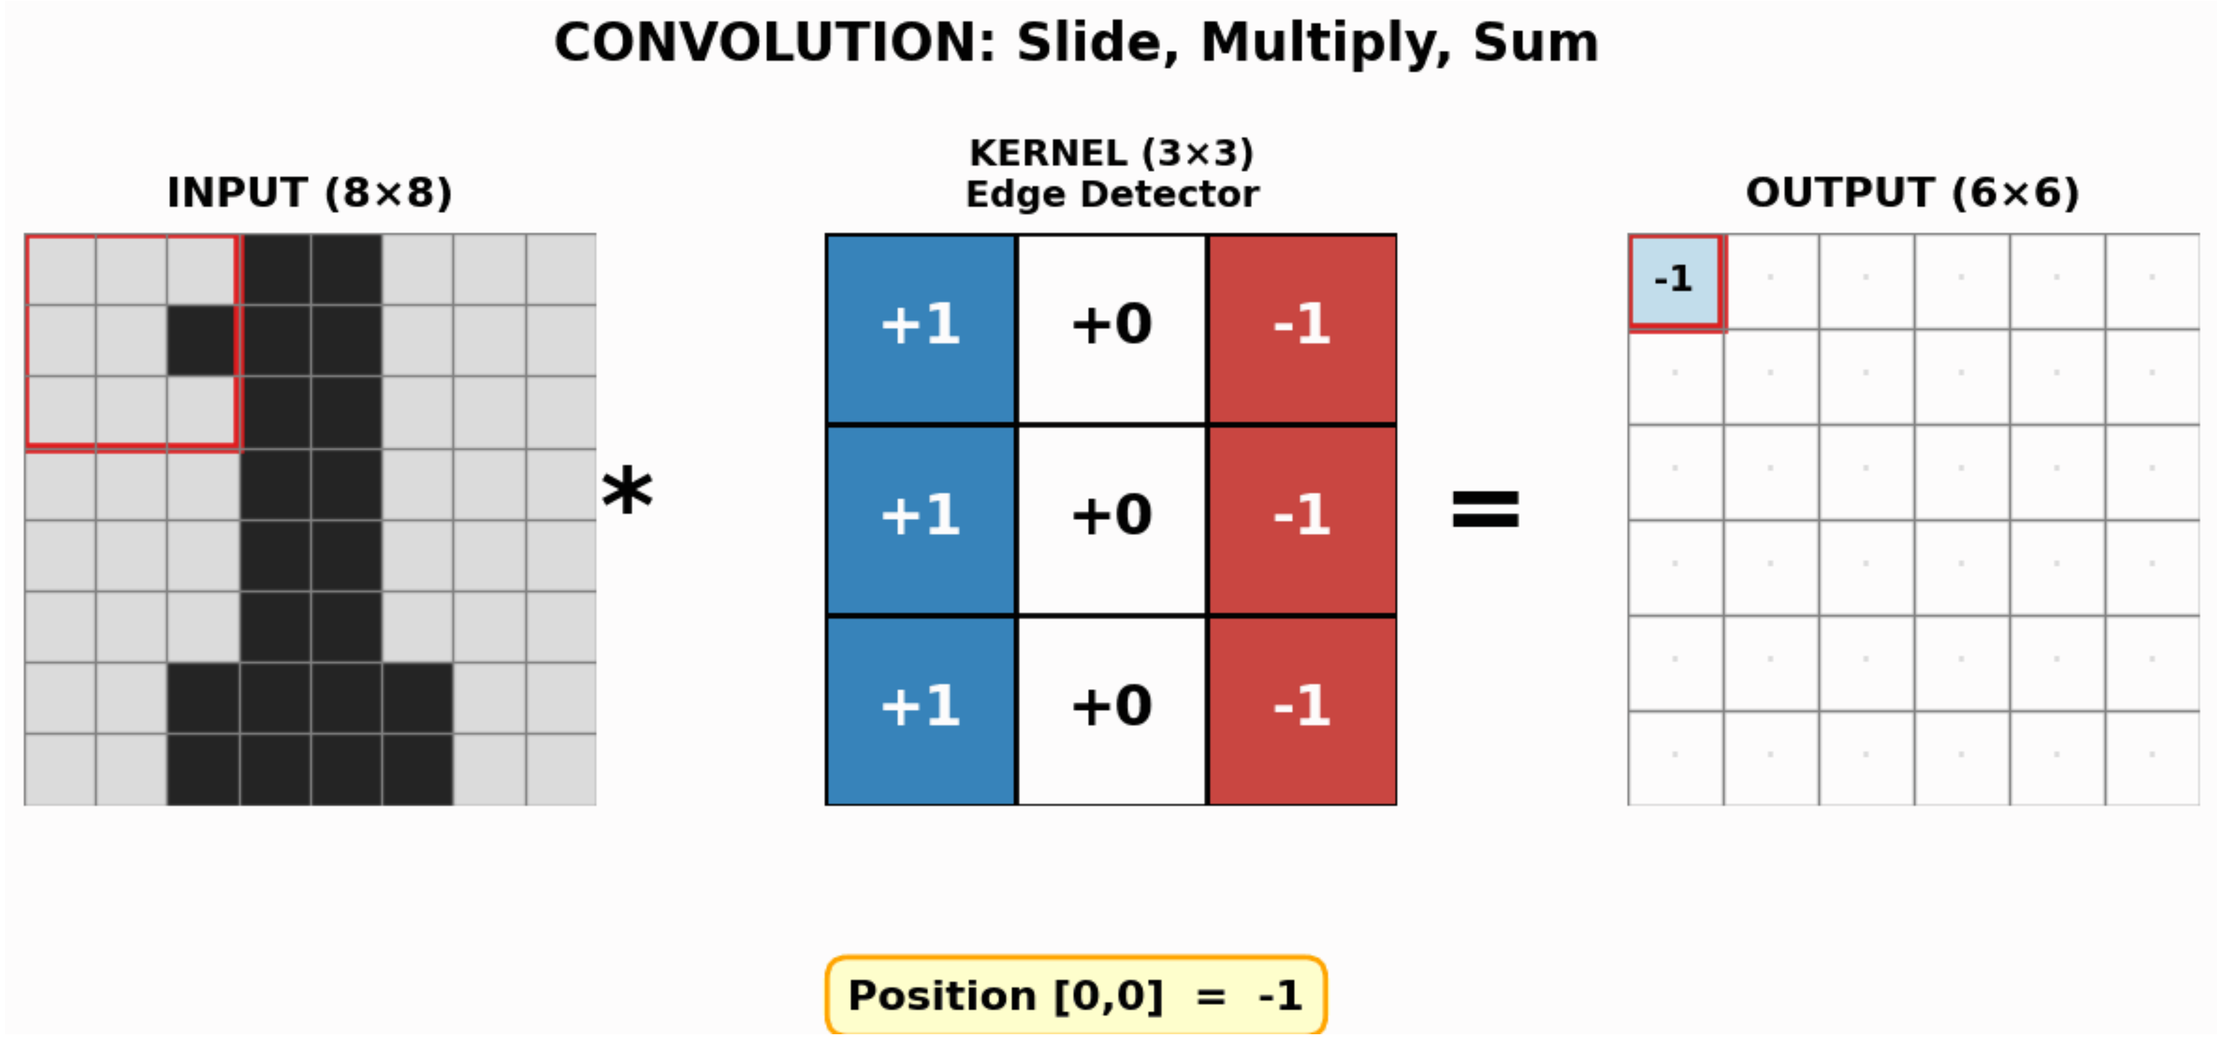
CONVOLUTION: Slide, Multiply, Sum
INPUT (8×8)
KERNEL (3×3)
Edge Detector
OUTPUT (6×6)
+1
+0
-1
+1
+0
-1
+1
+0
-1
-1
Position [0,0] = -1

# This was all a lie!

- All those example are actually ``cross-correlation'', not ``convolution''
- It is what is implemented in all those frameworks and called ``convolution''
- True convolution is almost that, but you need to flip the filter before applying

# This was all a lie!

- All those example are actually ``cross-correlation'', not ``convolution''
- It is what is implemented in all those frameworks and called ``convolution''
- True convolution is almost that, but you need to flip the filter before applying

Exercise: think why when starting from random init, it doesn't matter!

# Example

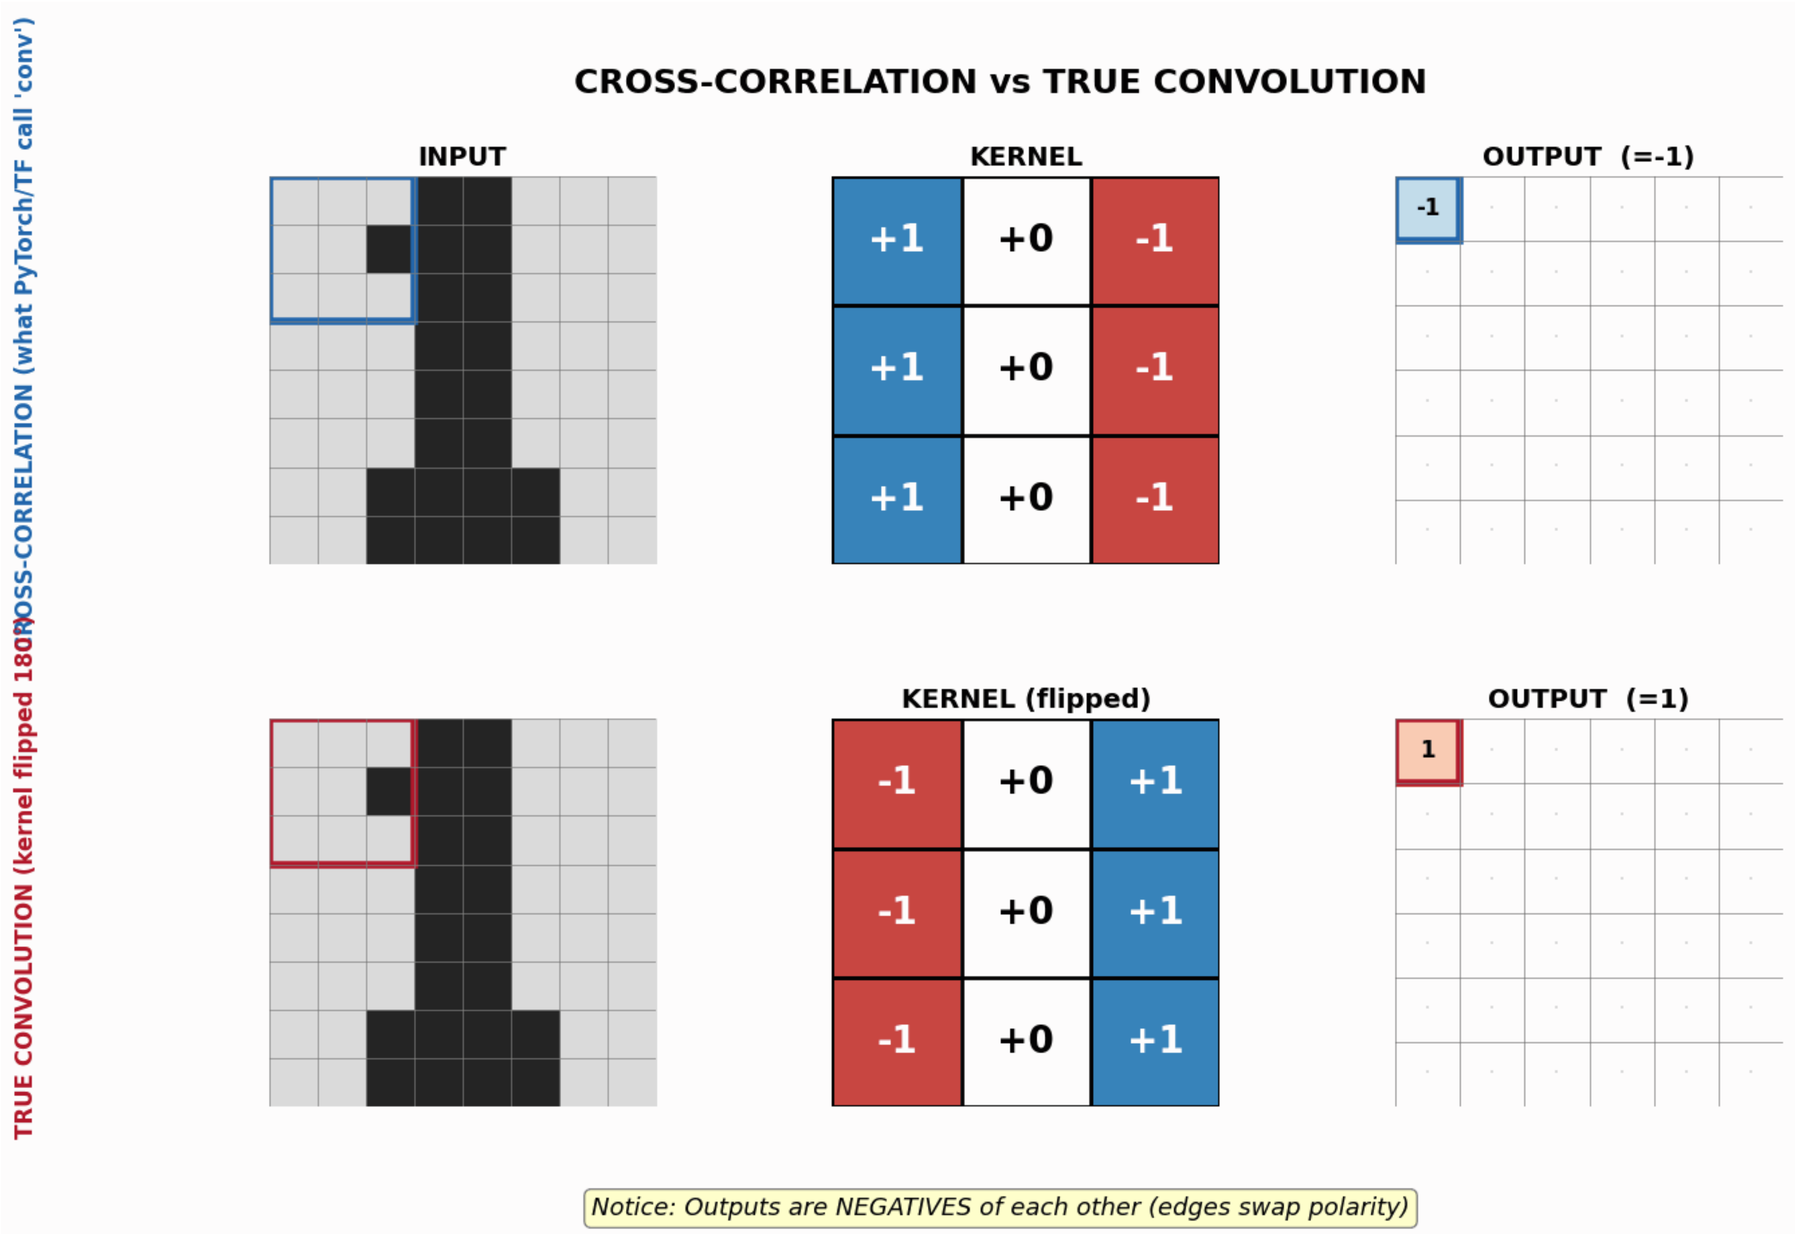
CROSS-CORRELATION vs TRUE CONVOLUTION
CROSS-CORRELATION (what PyTorch/TF call 'conv')
TRUE CONVOLUTION (kernel flipped 180°)
INPUT
KERNEL
+1
+0
-1
+1
+0
-1
+1
+0
-1
OUTPUT (=-1)
-1
KERNEL (flipped)
-1
+0
+1
-1
+0
+1
-1
+0
+1
OUTPUT (=1)
1
Notice: Outputs are NEGATIVES of each other (edges swap polarity)

# Example

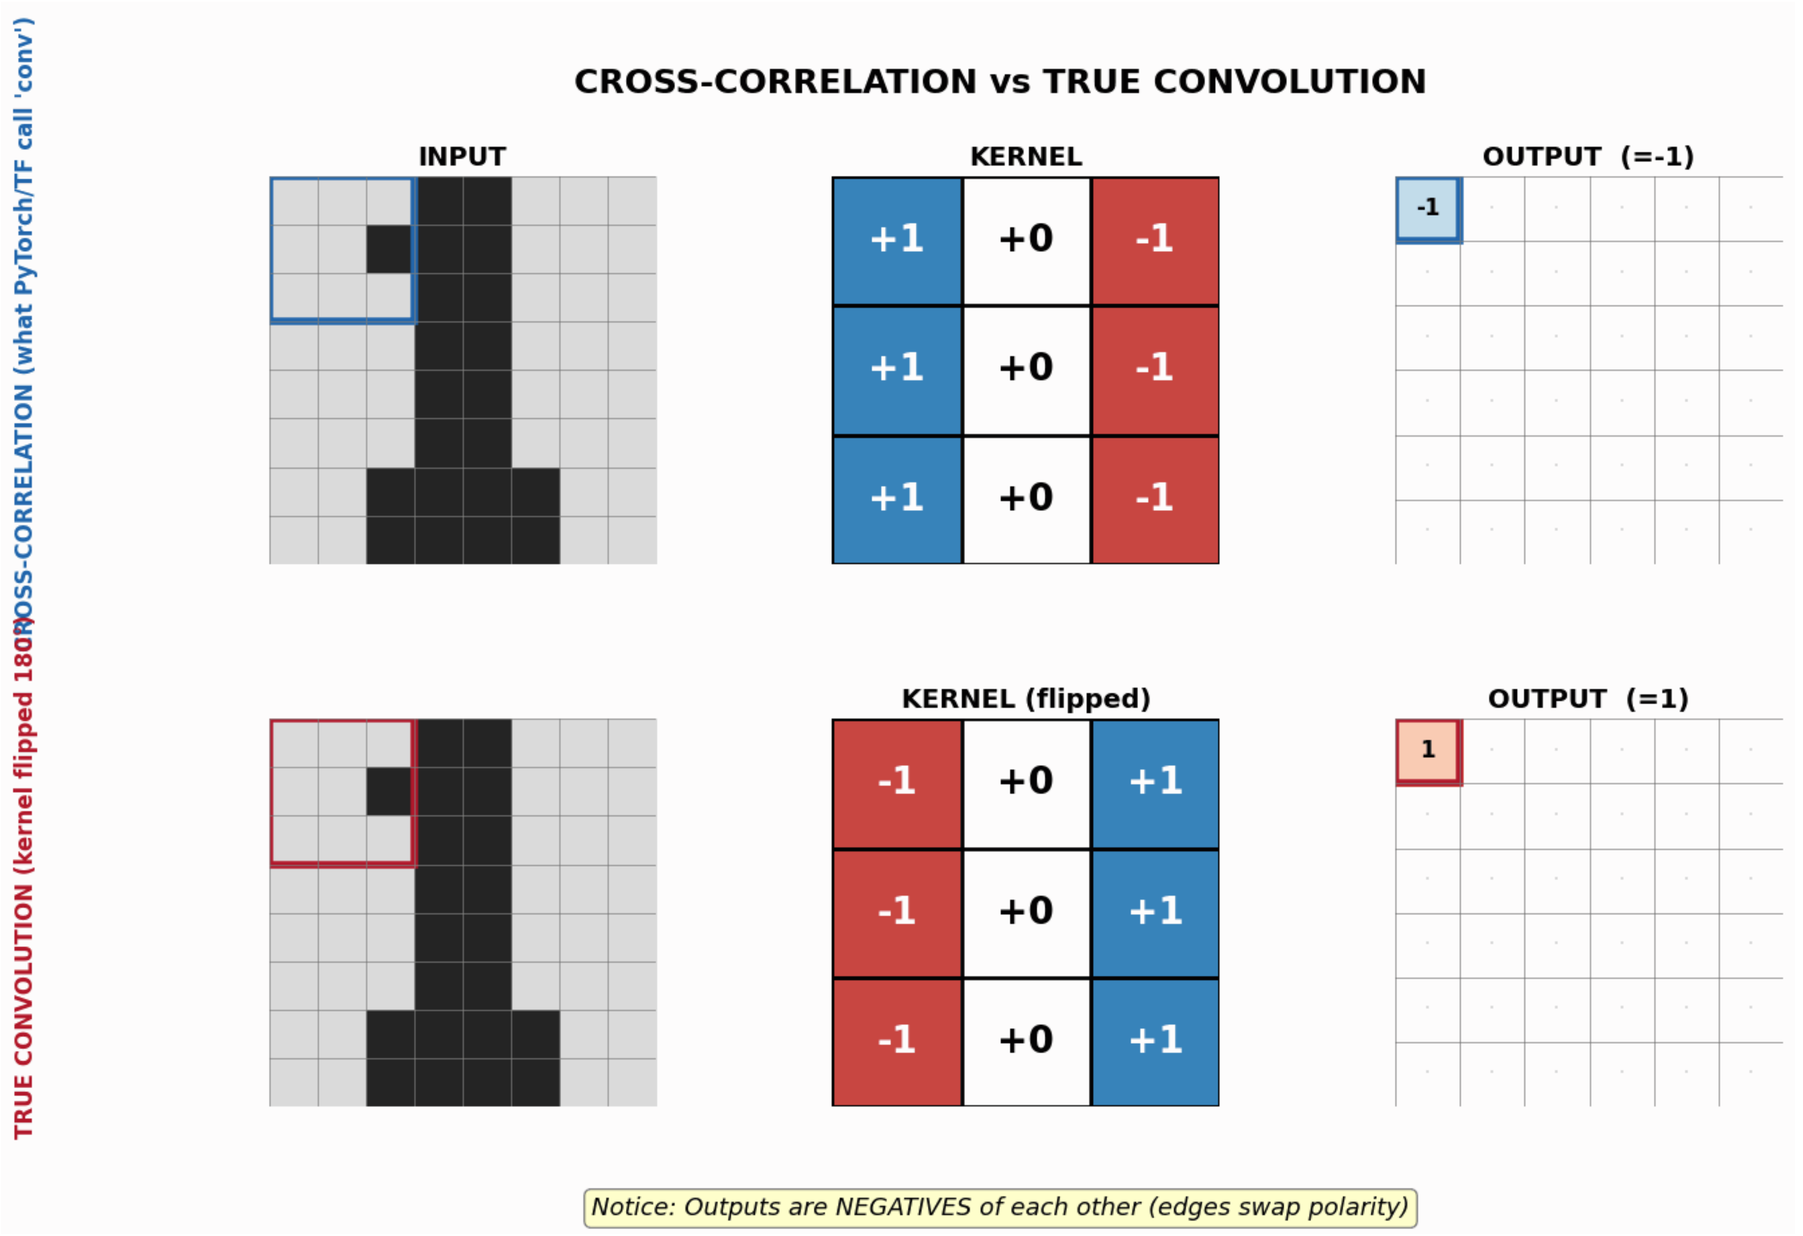
CROSS-CORRELATION vs TRUE CONVOLUTION
CROSS-CORRELATION (what PyTorch/TF call 'conv')
TRUE CONVOLUTION (kernel flipped 180°)
INPUT
KERNEL
+1
+0
-1
+1
+0
-1
+1
+0
-1
OUTPUT (=-1)
-1
KERNEL (flipped)
-1
+0
+1
-1
+0
+1
-1
+0
+1
OUTPUT (=1)
1
Notice: Outputs are NEGATIVES of each other (edges swap polarity)

# Your First Convolutional Network

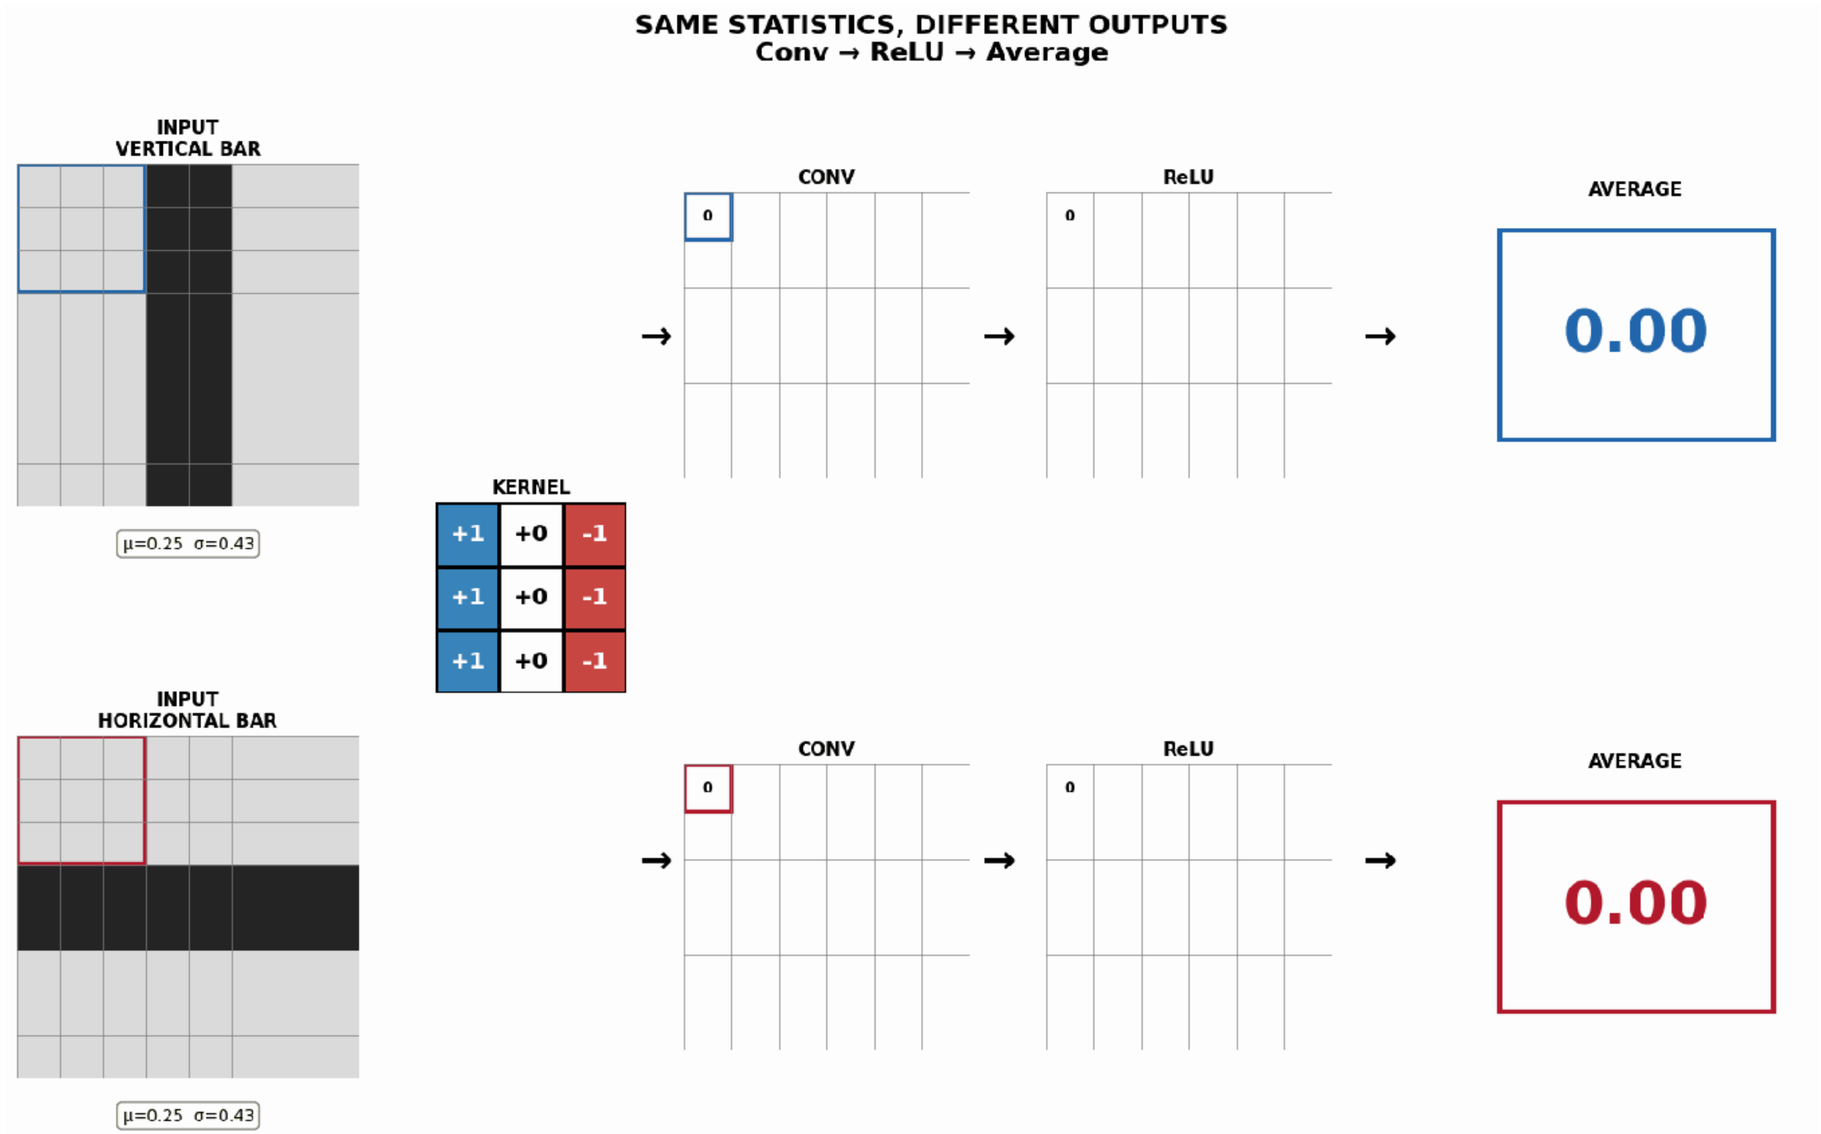

# Your First Convolutional Network

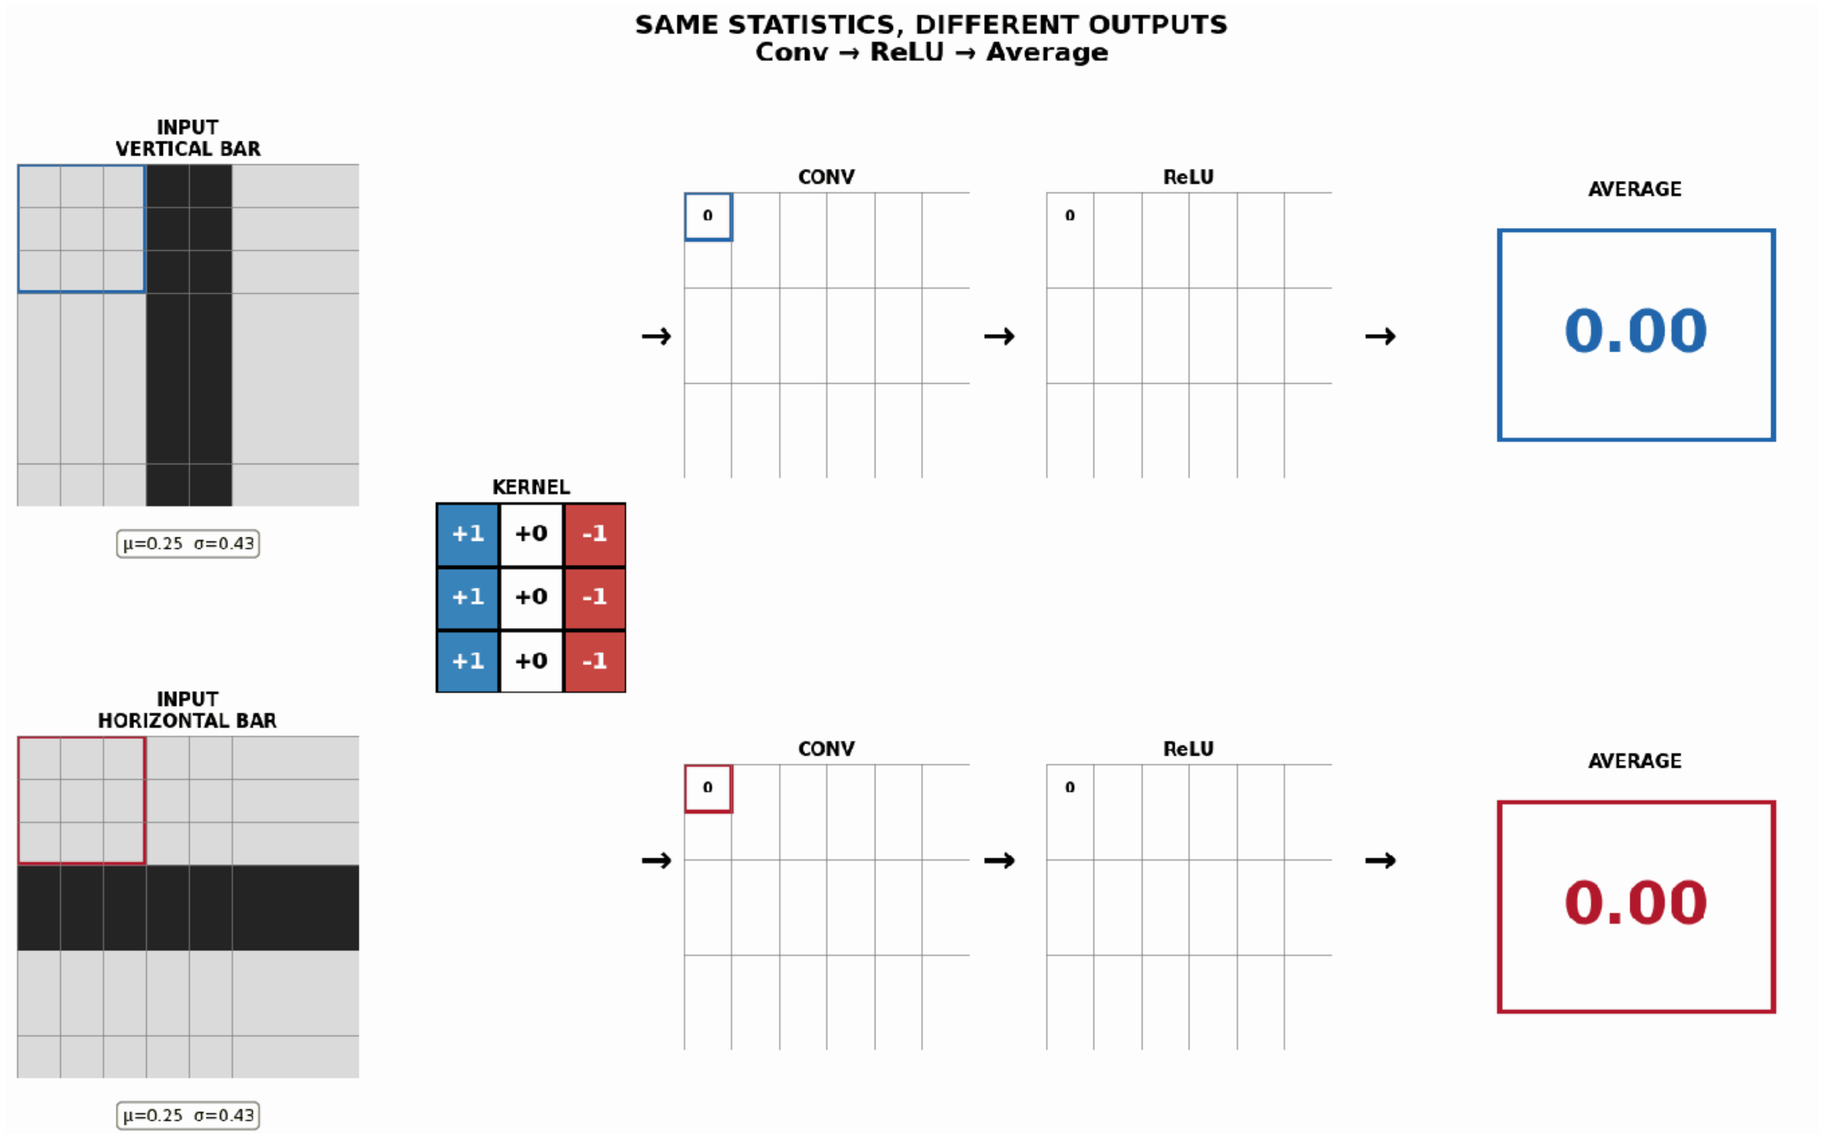

# Your First Convolutional Network

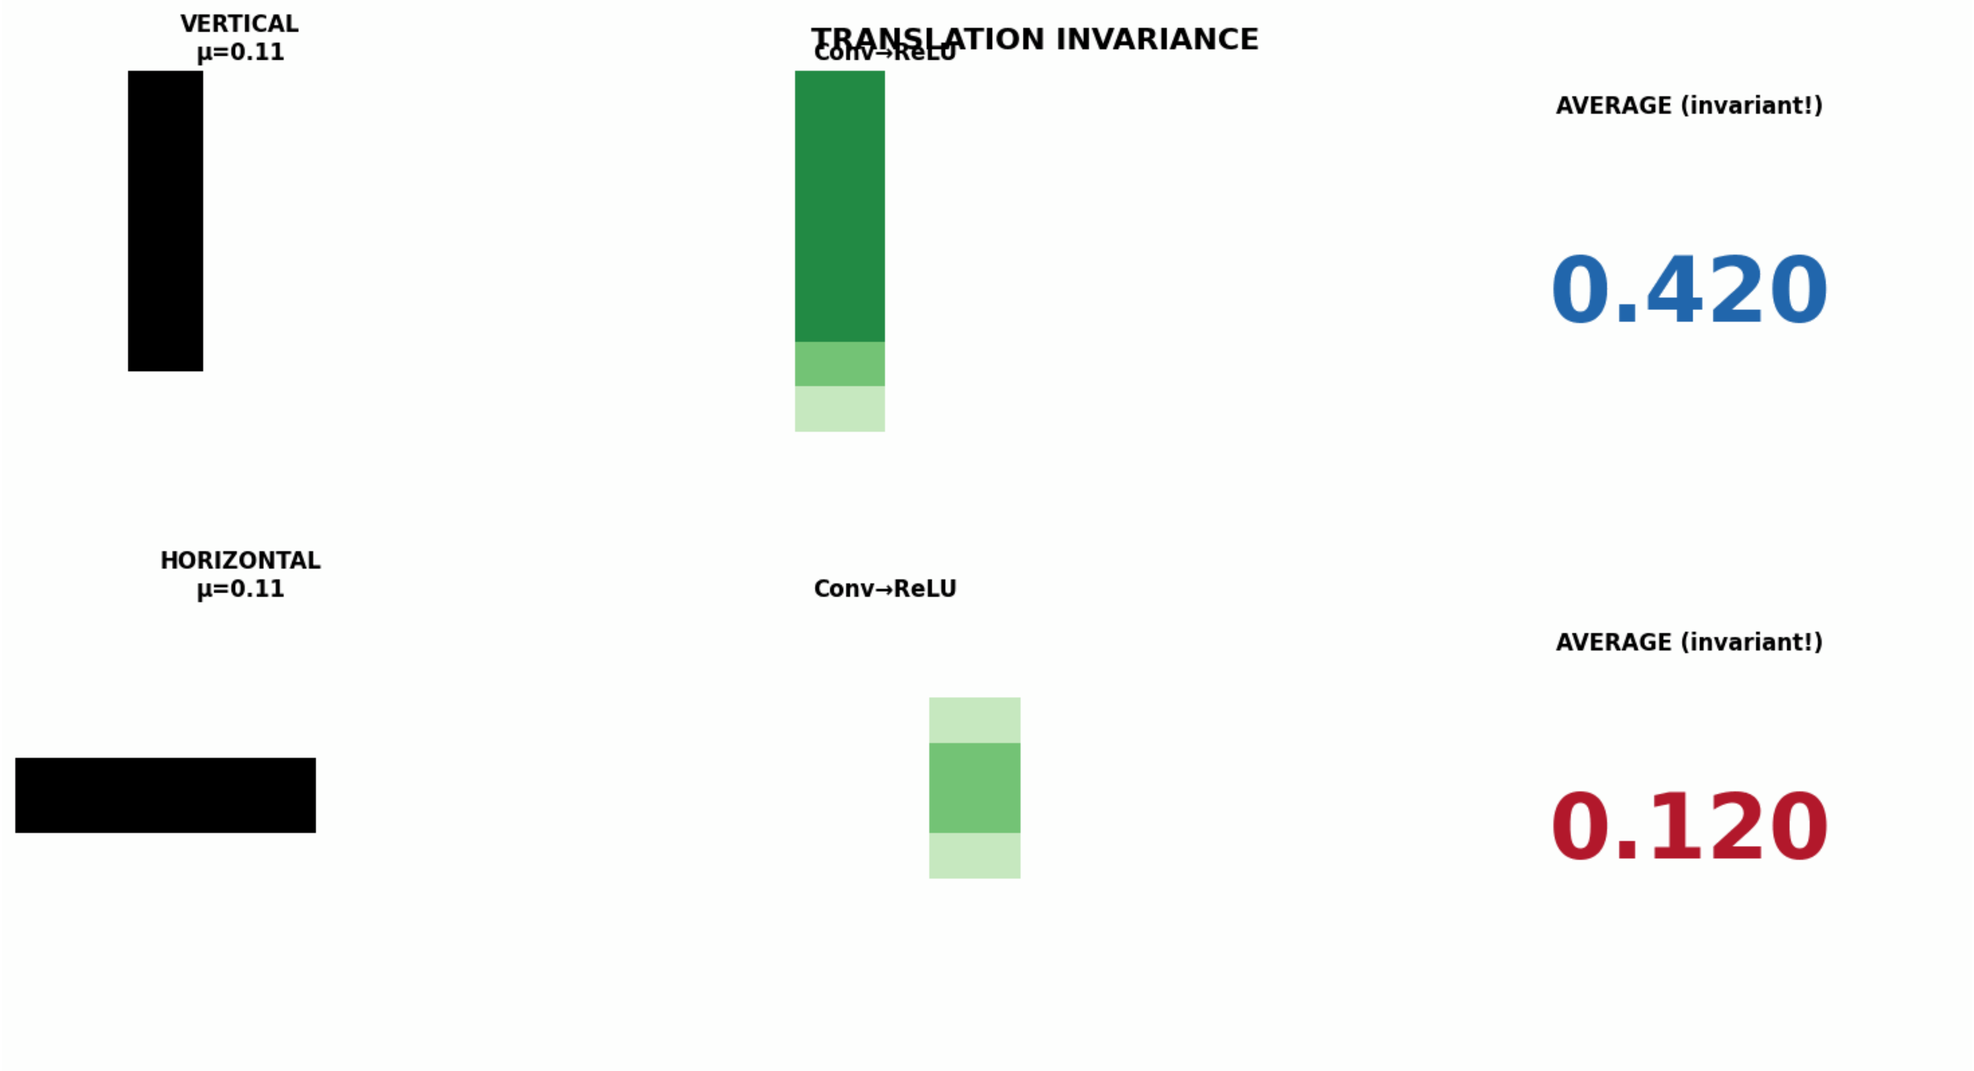

# Your First Convolutional Network

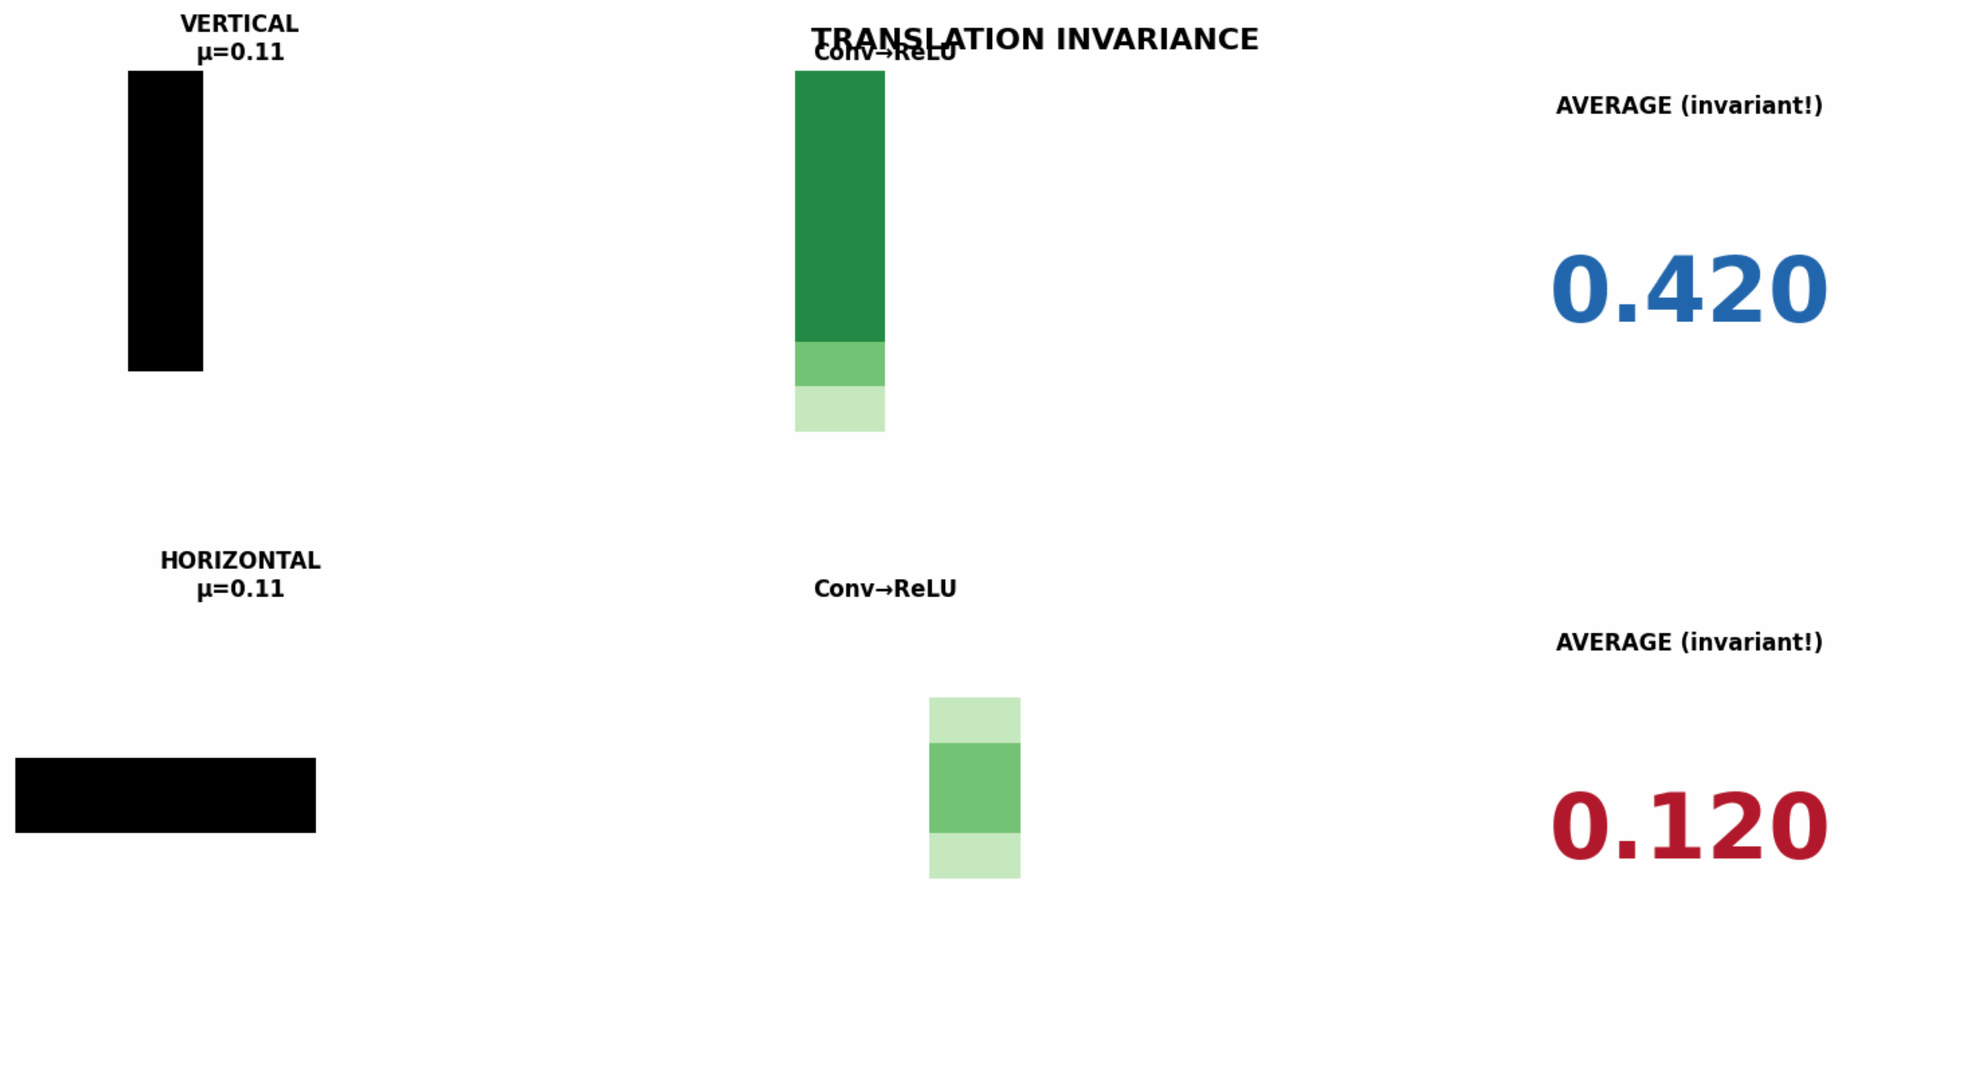

# Multiple Filters

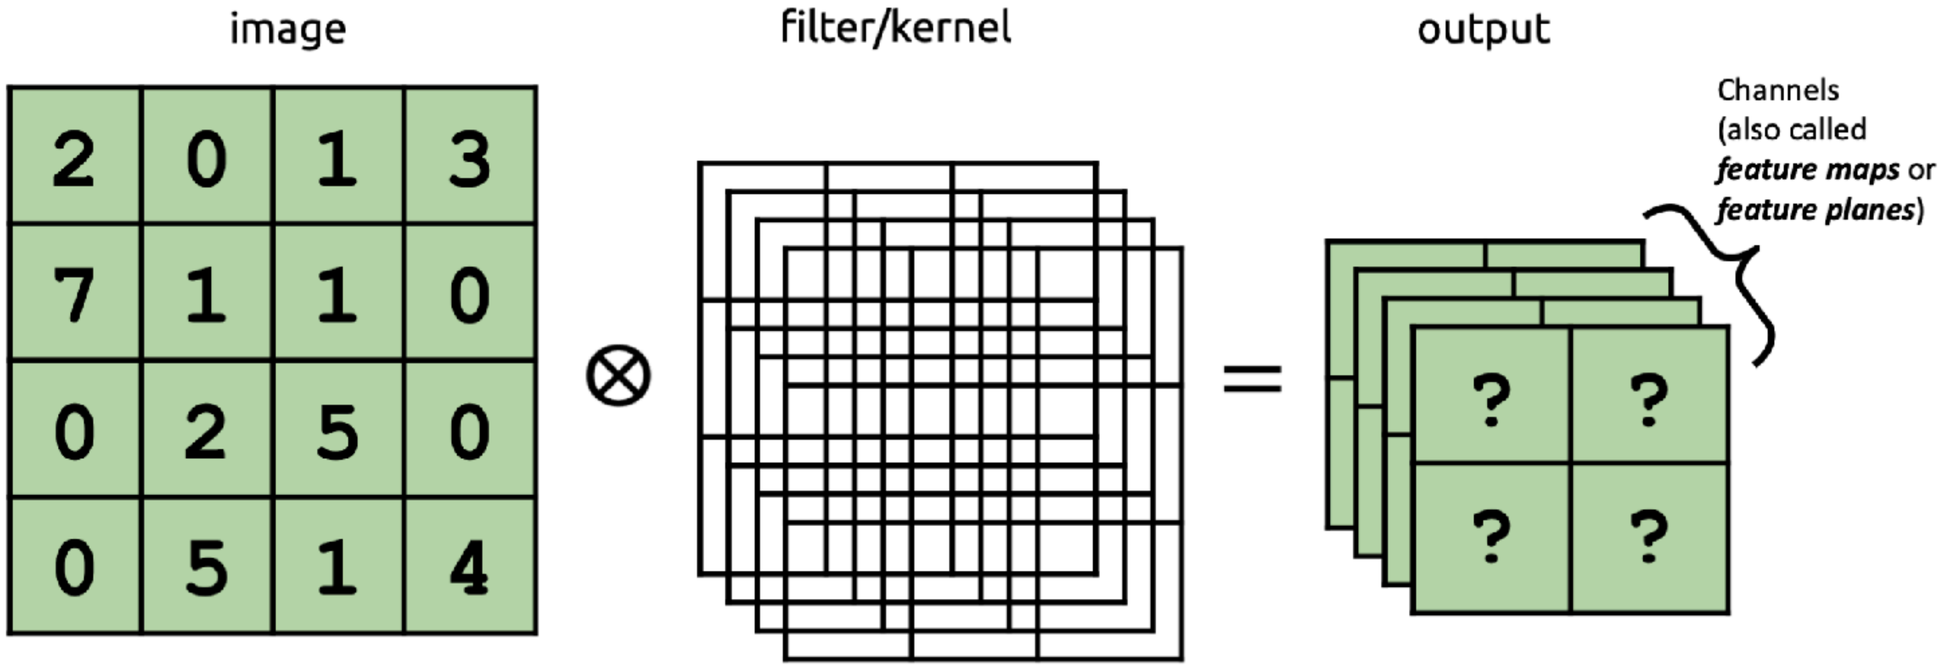

# Multiple Multi-Channel Filters

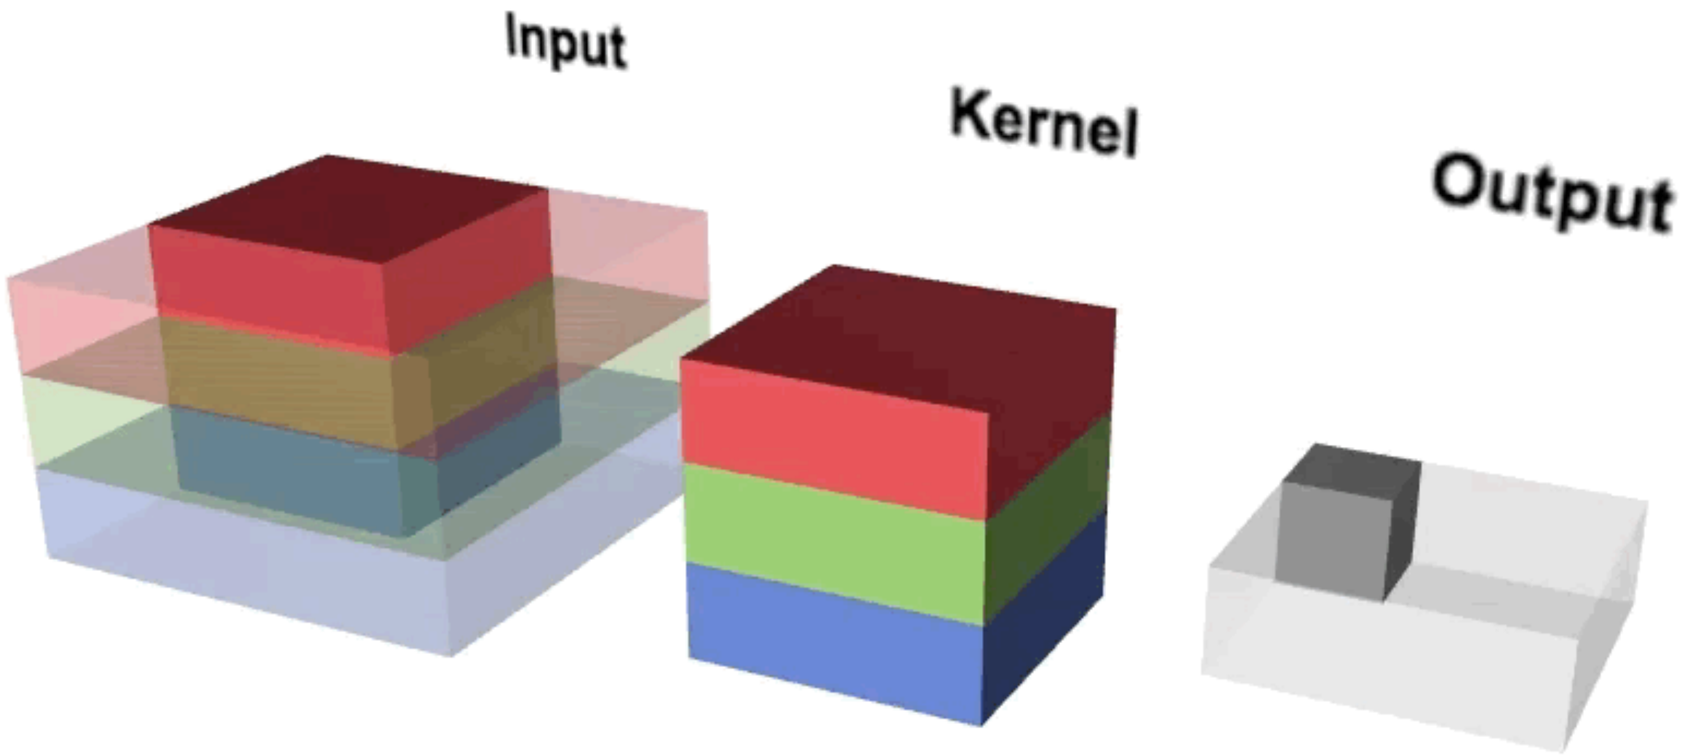

# Multiple Multi-Channel Filters

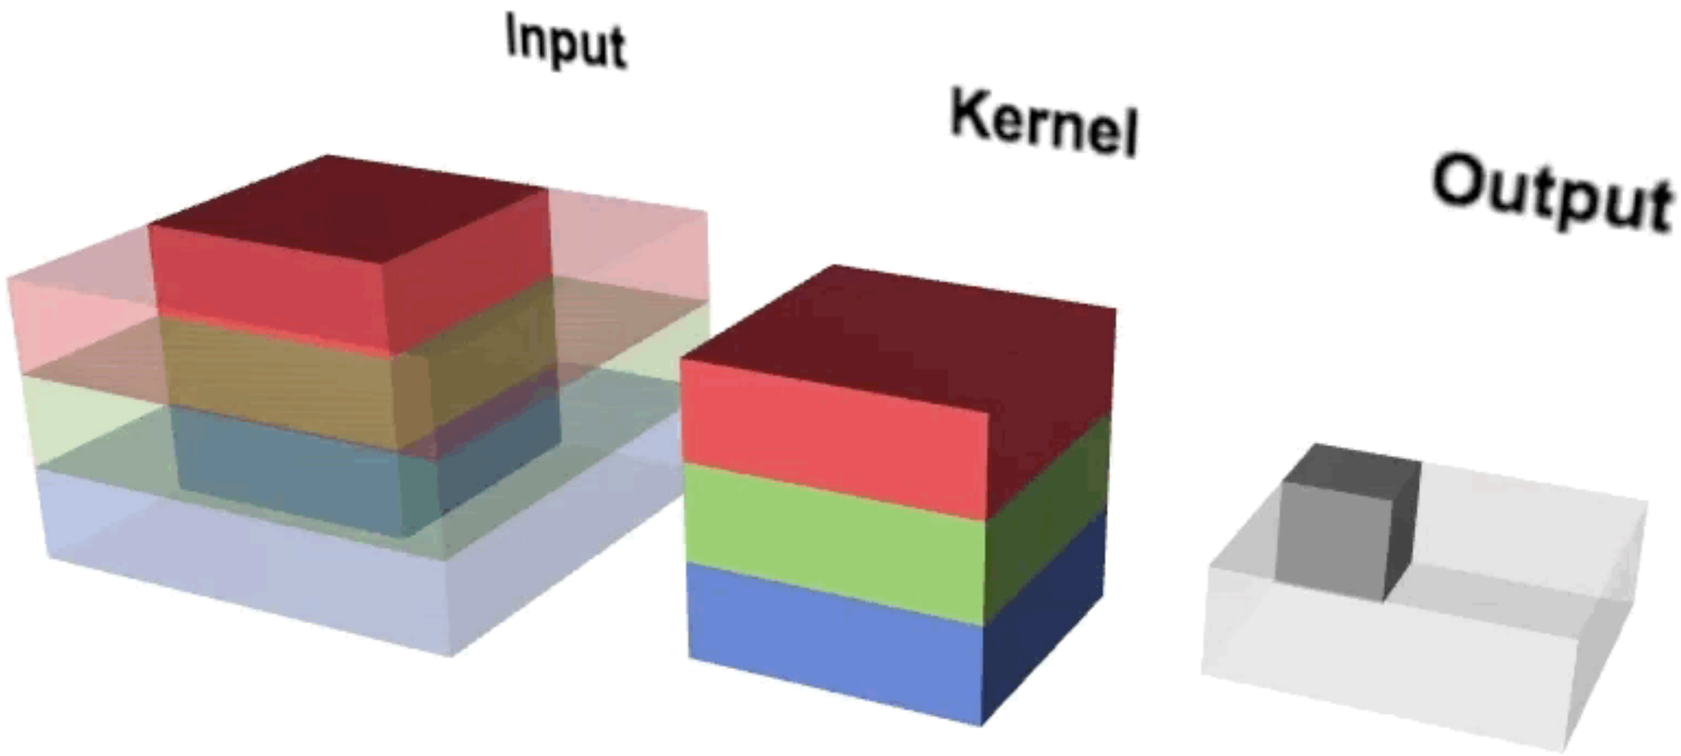

# Multiple Multi-Channel Filters

Input image

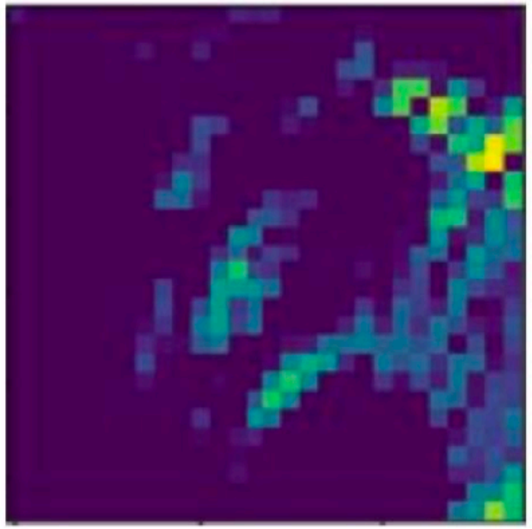

Output of filter 1

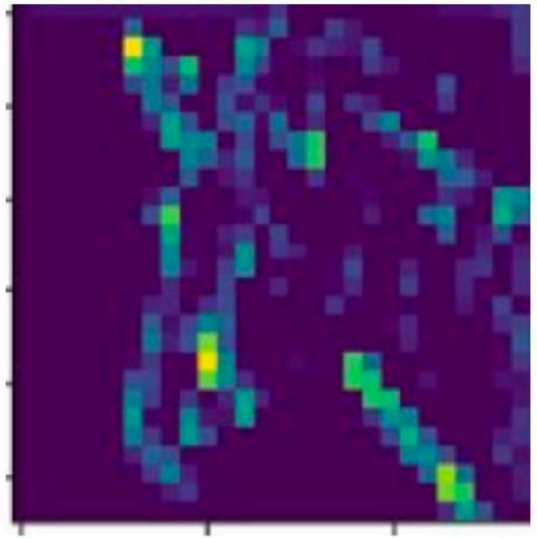

Output of filter 2

# More Filter Examples

| Operation | Filter | Convolved Image |
|---|---|---|
| Identity | $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ |  |
| Edge detection | $\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$ |  |
| | $\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$ |  |
| | $\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$ |  |
| Sharpen | $\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$ |  |
| Box blur (normalized) | $\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$ |  |
| Gaussian blur (approximation) | $\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$ |  |

# Deep Convolutional Networks

# Deep Convolutional Networks

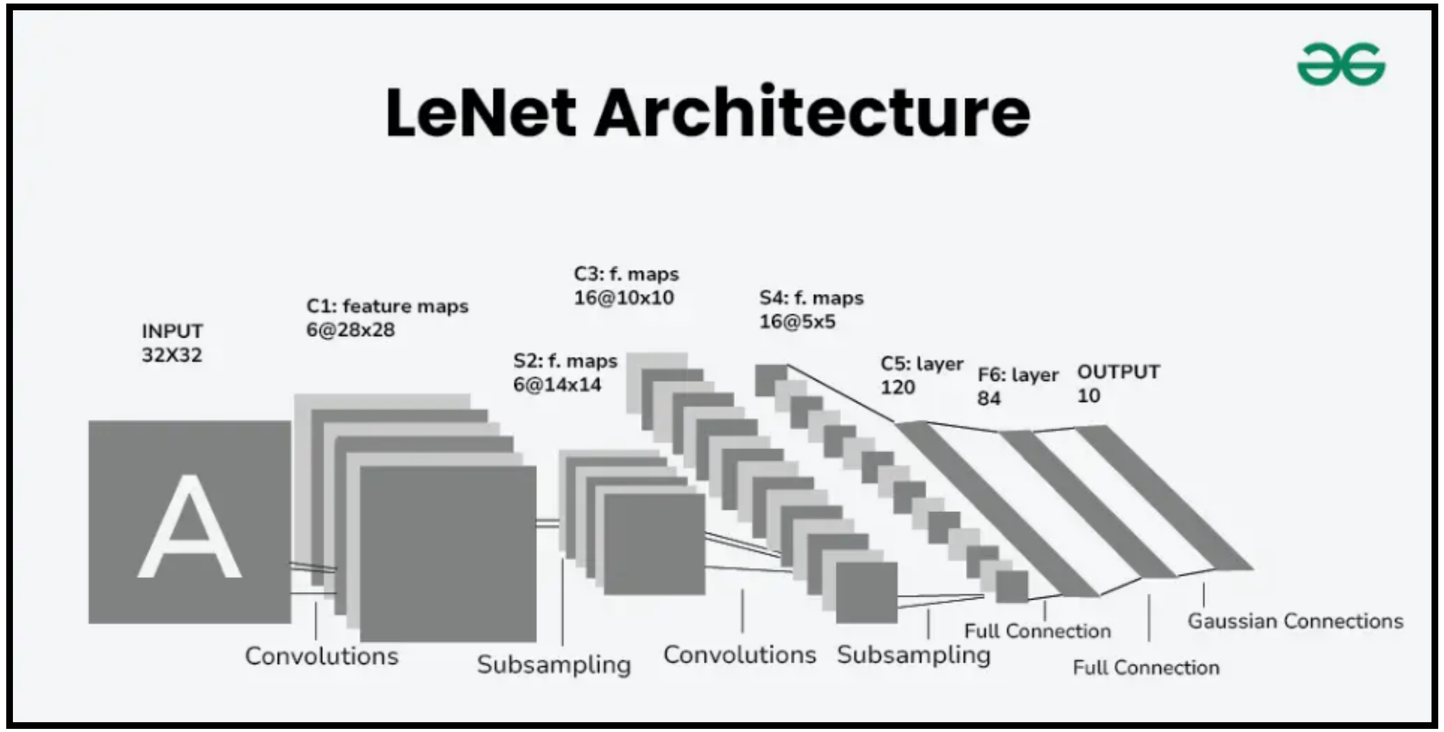

# Deep Convolutional Networks

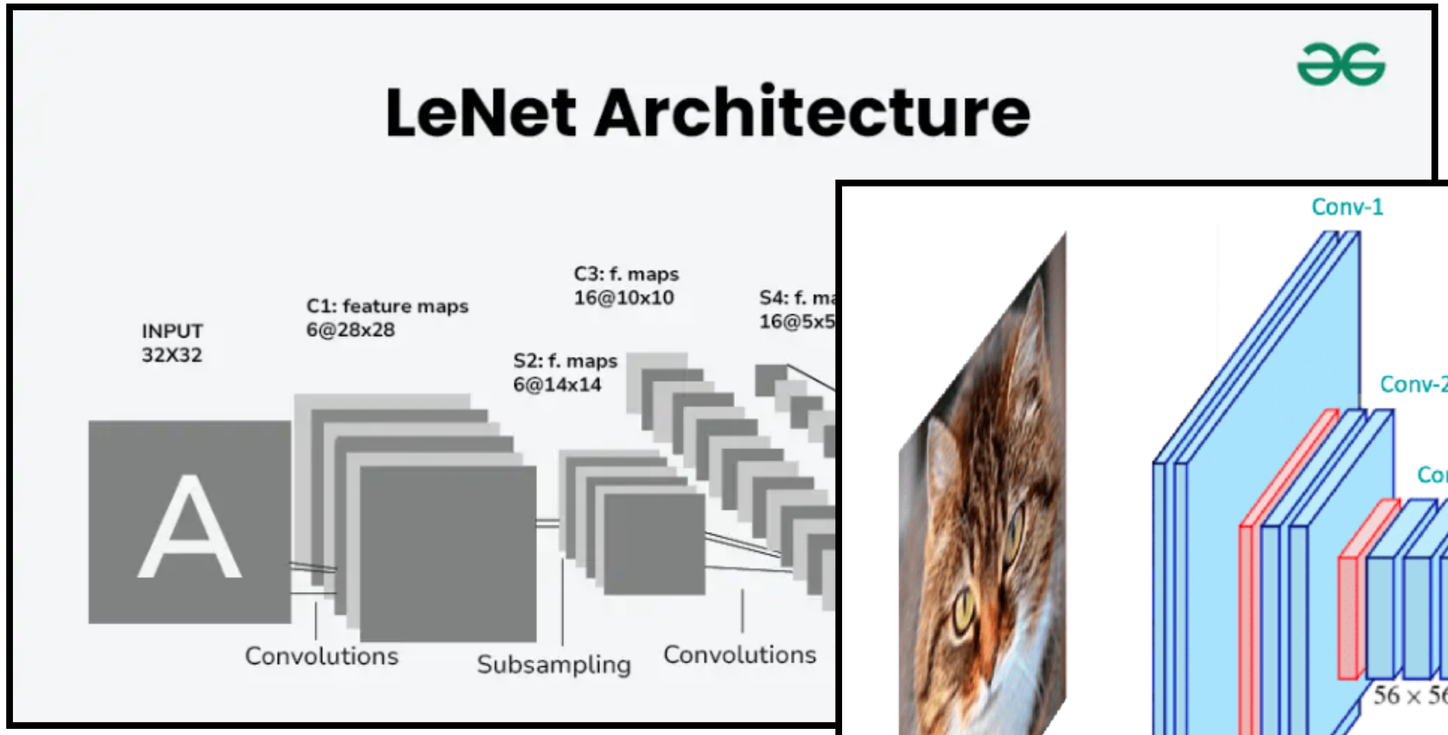


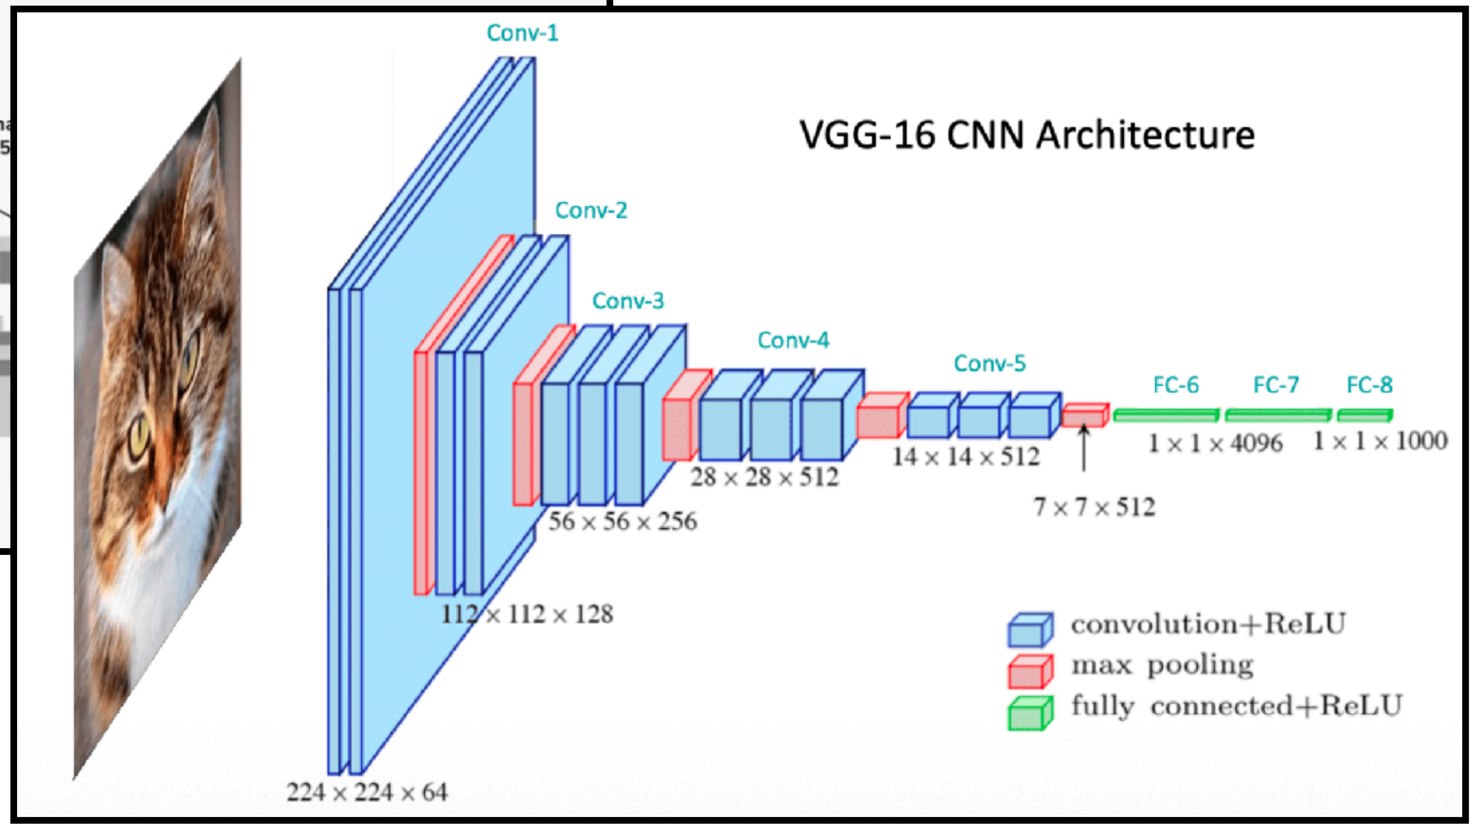

# Other Convolution Parameters

## Stride and padding!

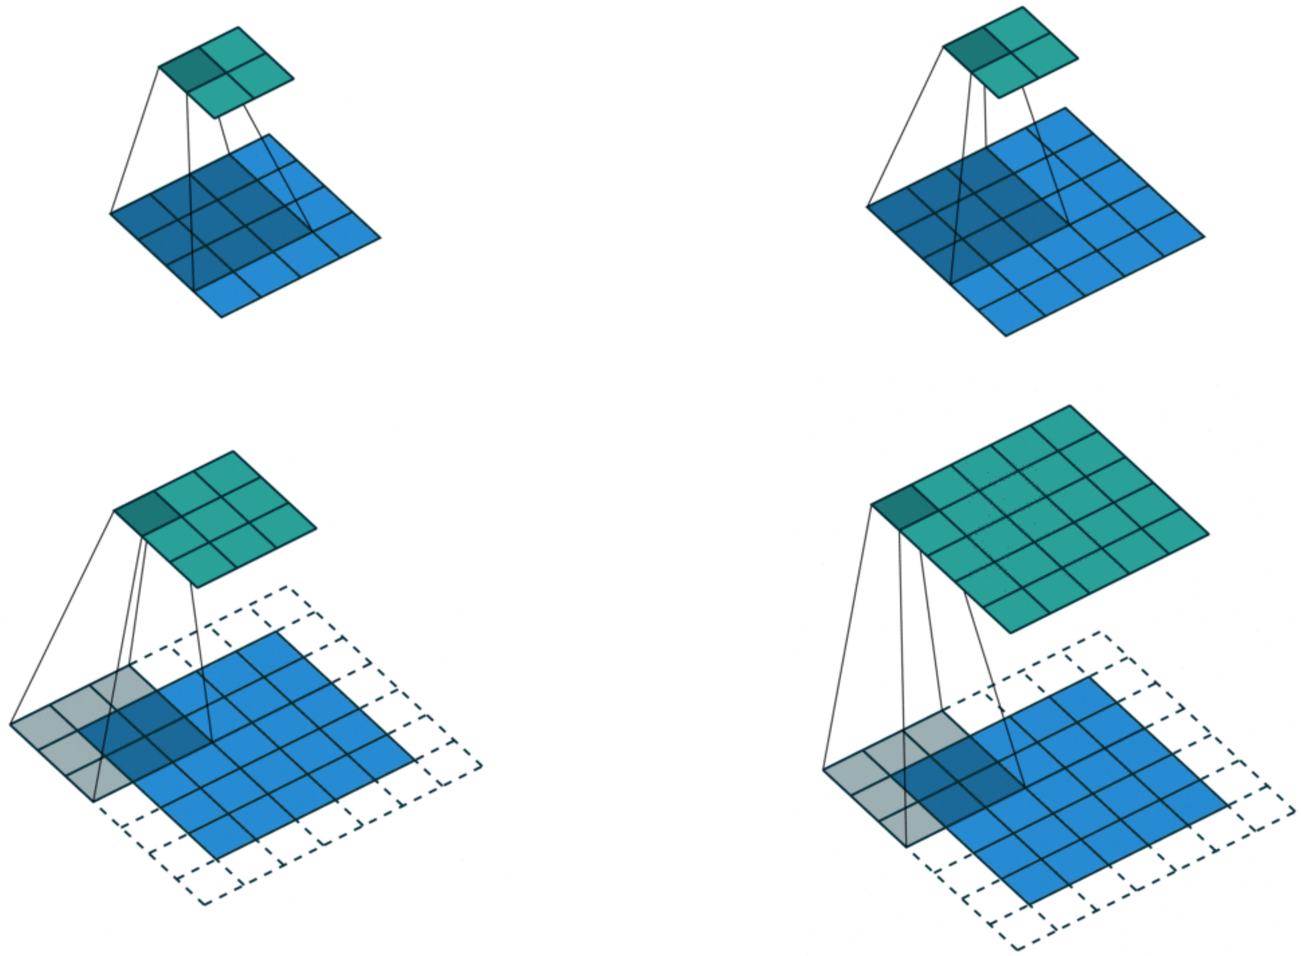

https://github.com/vdumoulin/conv_arithmetic

# Other Convolution Parameters

## Stride and padding!

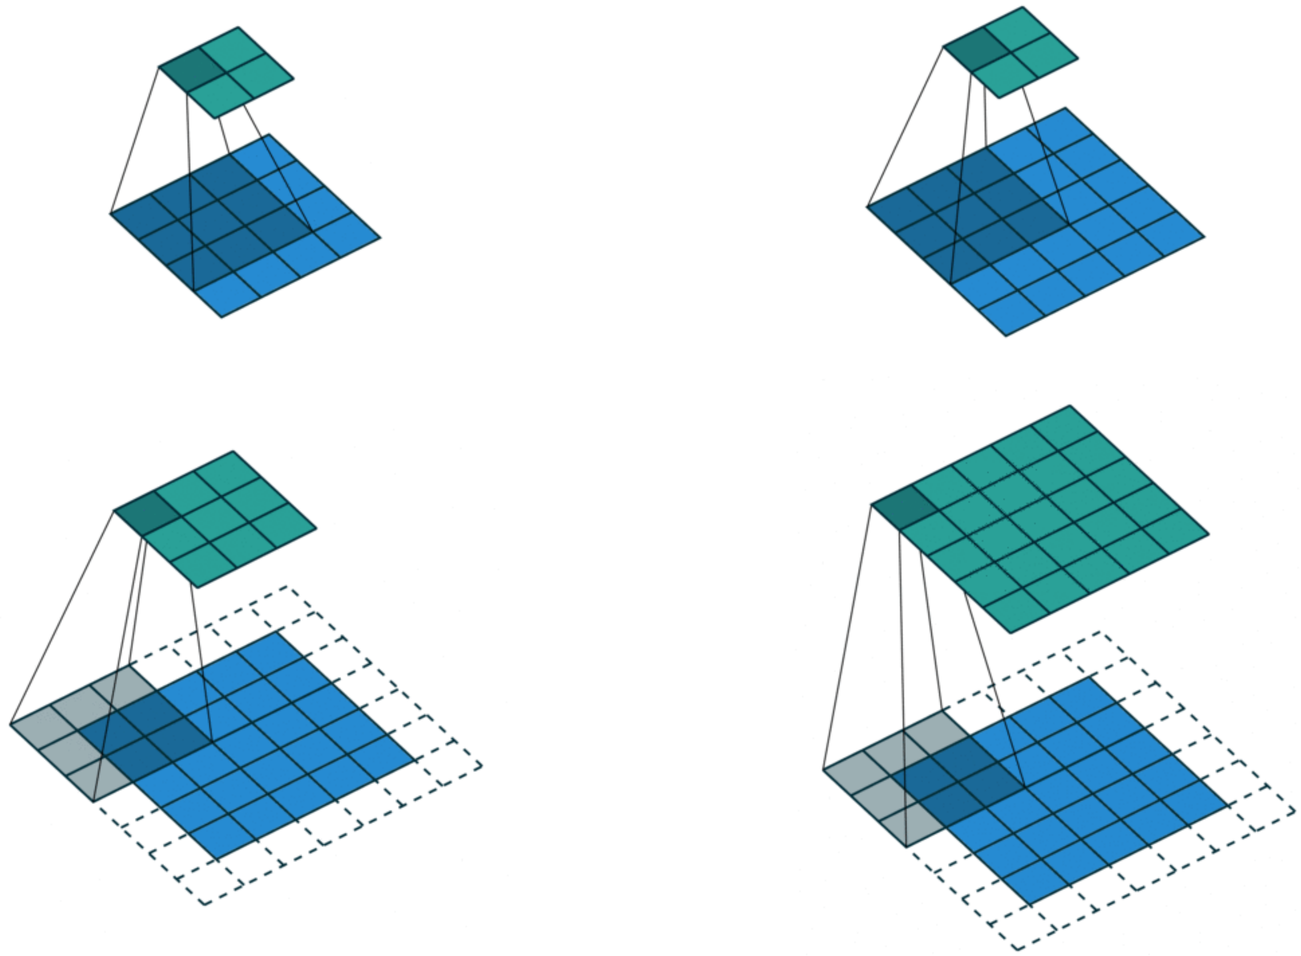

https://github.com/vdumoulin/conv_arithmetic

# Provable Benefit of Convolution

- You will encounter the same “patterns” at different translations (position)
- Holds for 1d (audio, music), 2d (images), 3d (videos)
- This is what the brain does!

# Provable Benefit of Convolution

- You will encounter the same “patterns” at different translations (position)
- Holds for 1d (audio, music), 2d (images), 3d (videos)
- This is what the brain does!

Exercise: can you think of inputs for which we DO NOT want convolutions?