

CSCI 1470

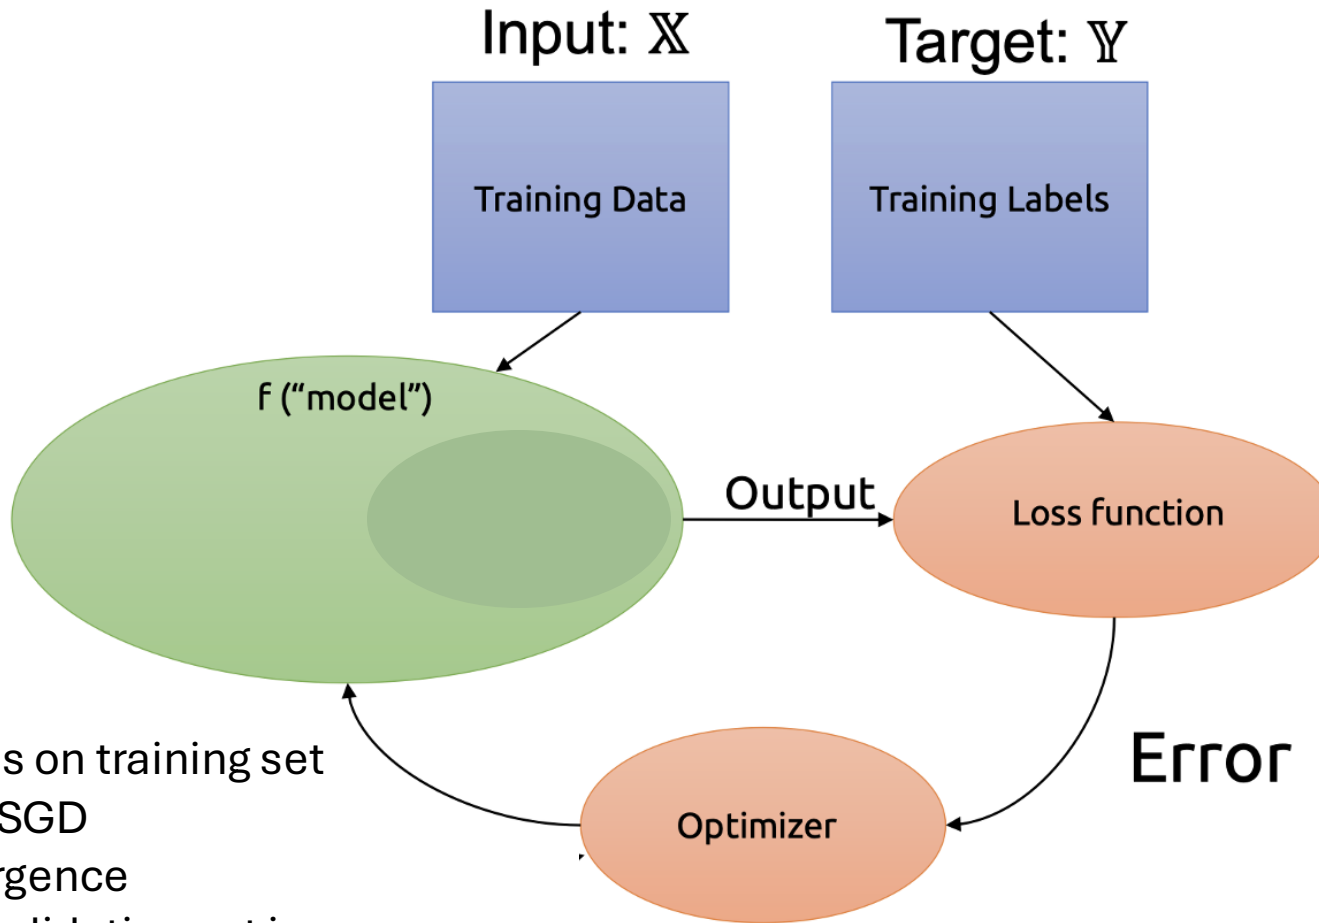
Eric Ewing

Wednesday,
2/12/25

Deep Learning


Day 10: Introduction to Convolutions

Recap: MLPs




1. Compute Error/Loss on training set
2. Run Backprop and SGD
3. Repeat until convergence
4. If performance on validation set is acceptable, terminate, else try new hyperparameters


A Brief History of AI with Deep Learning




McCulloch-Pitts




Rosenblatt Widrow-Hoff




Minsky-Papert




Rumelhart, Hinton et al.



LeCun



Hinton-Ruslan Krizhevsky et al. Vaswani



What has happened in the last 15 years?

What has changed?

1. Power and efficiency of compute (GPUs)
2. Availability of data (the internet)
3. New Architectures (e.g., CNNs, Transformers)



Issues with MLPs

1. Resource Intensive
2. Difficult to incorporate certain types of information
3. (and more)

Issues with MLPs

1. Resource Intensive
2. Difficult to incorporate certain types of information
3. (and more)

GPUs to the rescue!

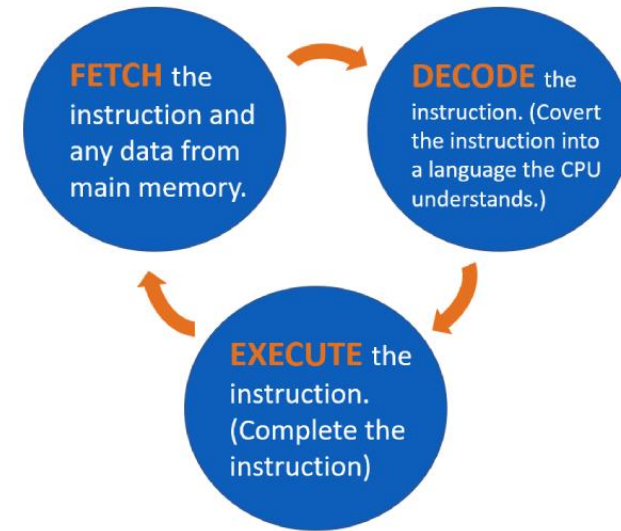
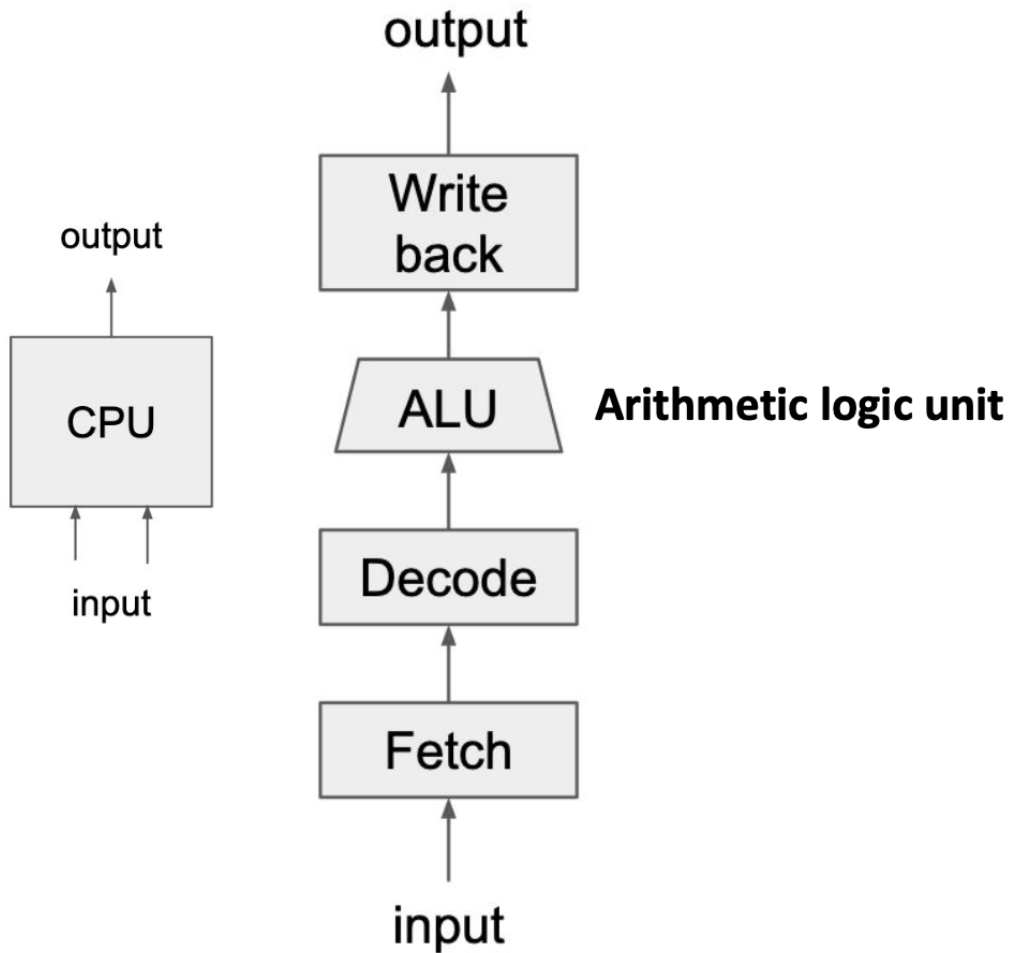


- *Graphics* Processing Units
- GPUs are really good at computing mathematical operations in parallel!
- Matrix multiplication == many **independent** multiply and add operations

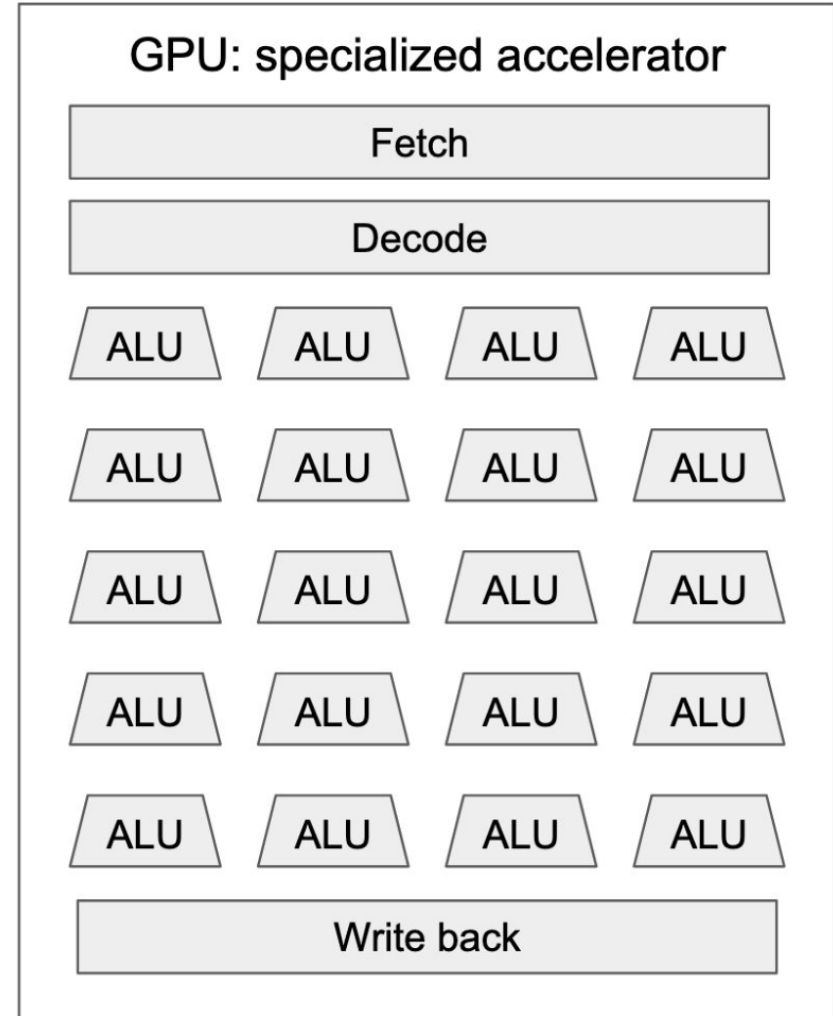
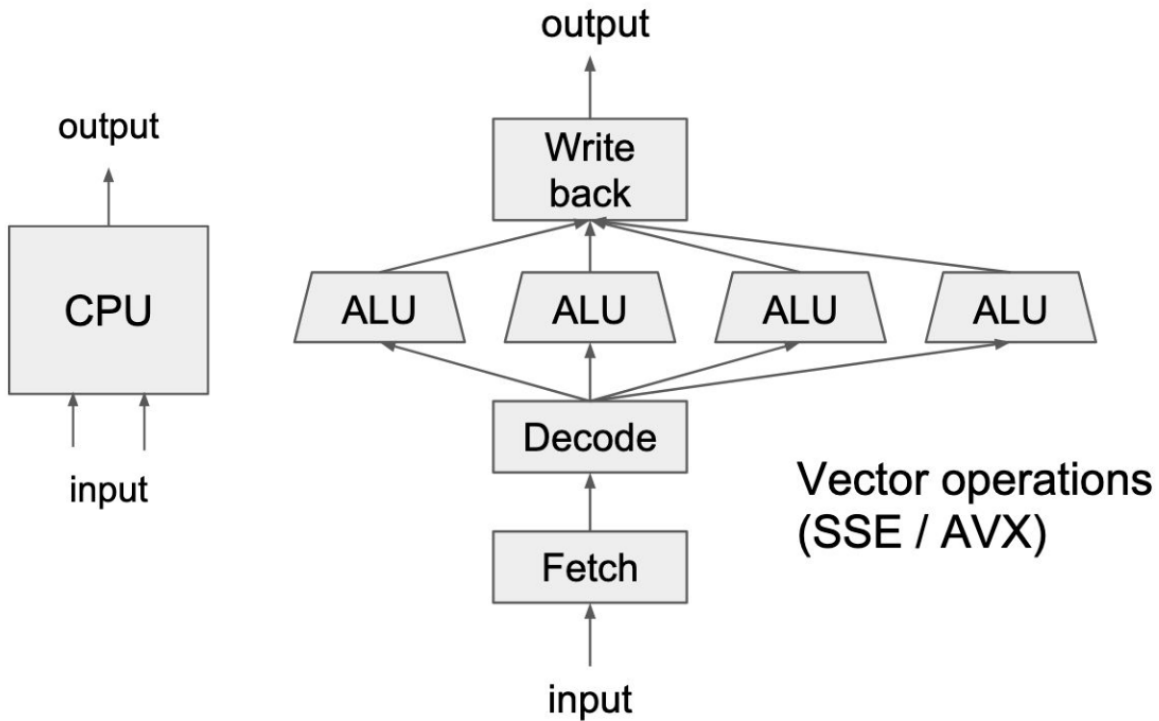
Easily parallelizable

GPUs are great for this!

CPU v/s GPU



CPU v/s GPU



GPU-Parallel Acceleration

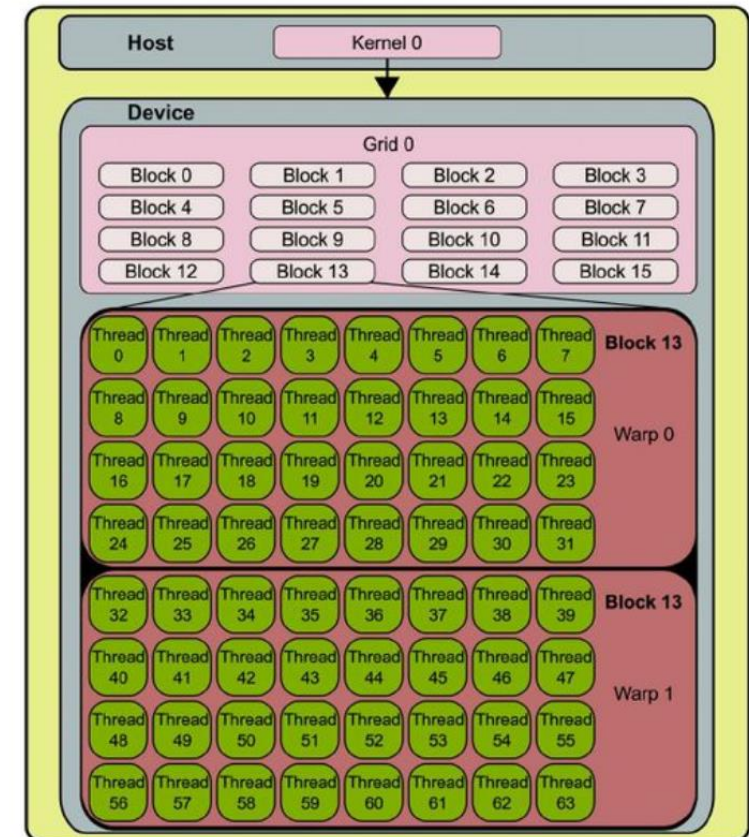
- User code (*kernels*) is compiled on the *host* (the CPU) and then transferred to the *device* (the GPU)
- Kernel is executed as a *grid*
- Each grid has multiple *thread blocks*
- Each thread block has multiple *warps*

A warp is the basic schedule unit in kernel execution

A warp consists of 32 threads

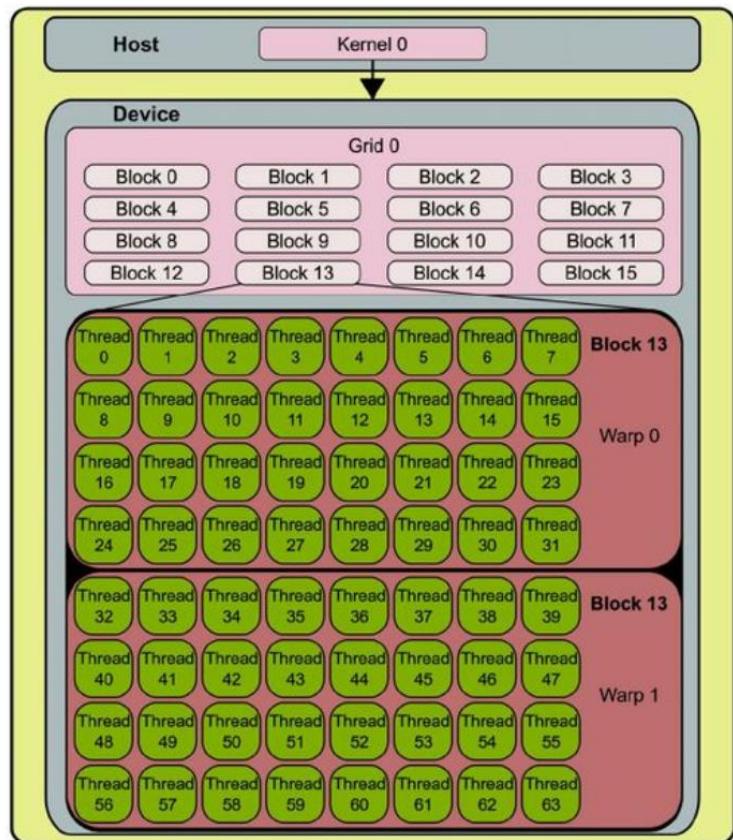
Compute Unified Device Architecture is a parallel computing platform and application programming interface (API)

CUDA compute model



GPU-Parallel Acceleration

CUDA compute model



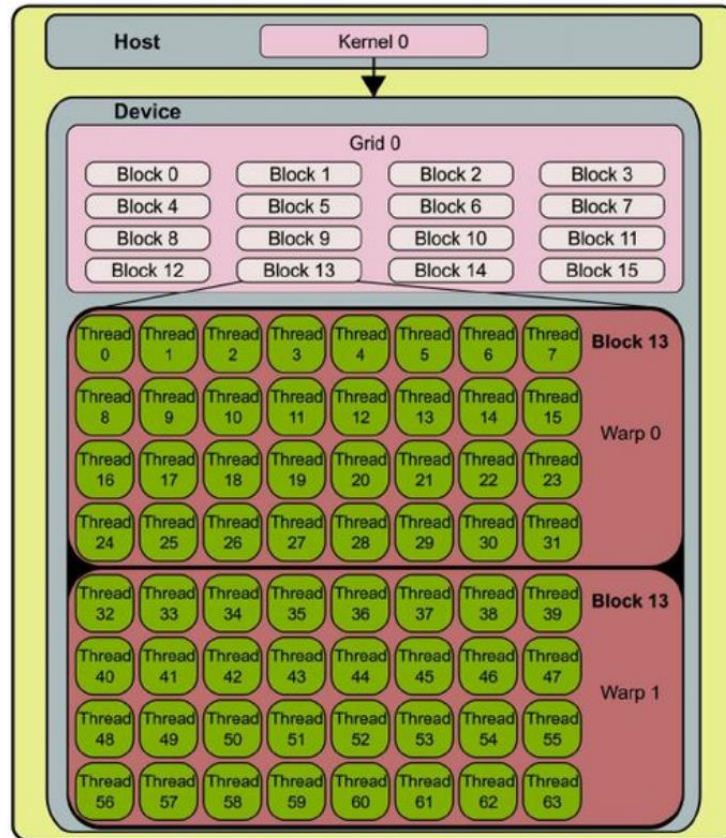
- Programmer decides how they want to parallelize the computation across grids and blocks
 - Modern deep learning frameworks take care of this for you
- CUDA compiler figures out how to schedule these units of computation on to the physical hardware

Any questions?



GPU-Parallel Acceleration

CUDA compute model




- Upshot: order of magnitude speedups!
- Example: training CNN on CIFAR-10 dataset

Device	Speed of training, examples/sec
2 x AMD <u>Opteron 6168</u>	440
i7-7500U	415
<u>GeForce 940MX</u>	1190
<u>GeForce 1070</u>	6500

From:
<https://medium.com/@andriylazorenko/tensorflow-performance-test-cpu-vs-gpu-79fcd39170c>

AMD GPUs are competitive for gaming and graphics, why not for AI?



- CUDA is far better than competitors (AMD) (With a benchmarking tool made by AMD)
 - Easier to use
 - Better optimization
- AMD makes GPUs for graphics, NVIDIA makes GPUs for AI

CUDA is Still a Giant Moat for NVIDIA

Despite everyone's focus on hardware, the software of AI is what protects NVIDIA



JAMES WANG

MAR 23, 2024

Issues with MLPs

1. Resource Intensive
2. Difficult to incorporate certain types of information

MLPs and Spatial Reasoning

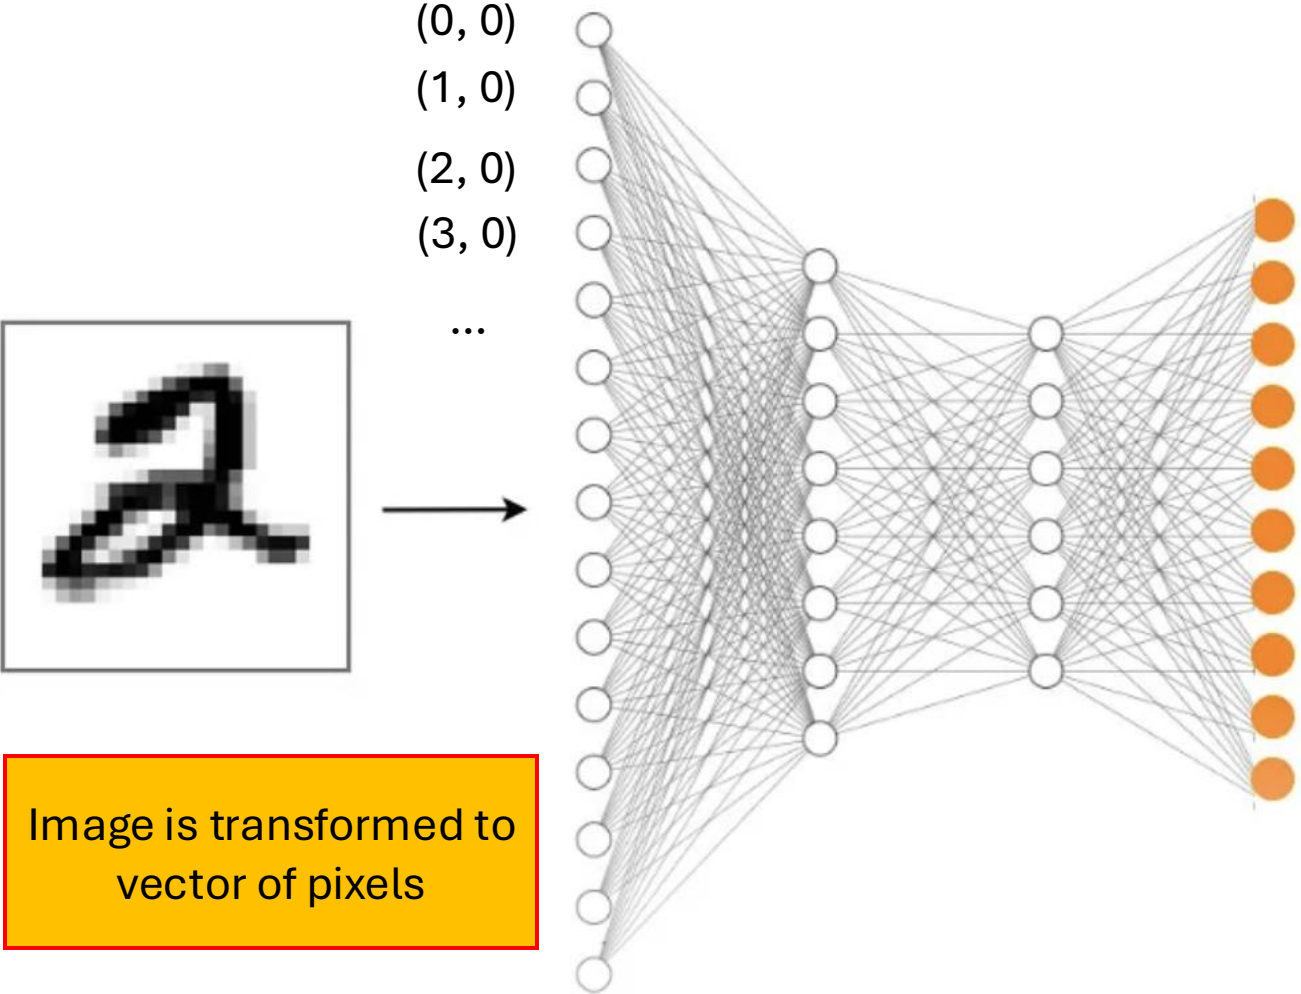


Image is transformed to vector of pixels

What would happen if we permuted the ordering of the pixels?

MLPs and Spatial Reasoning

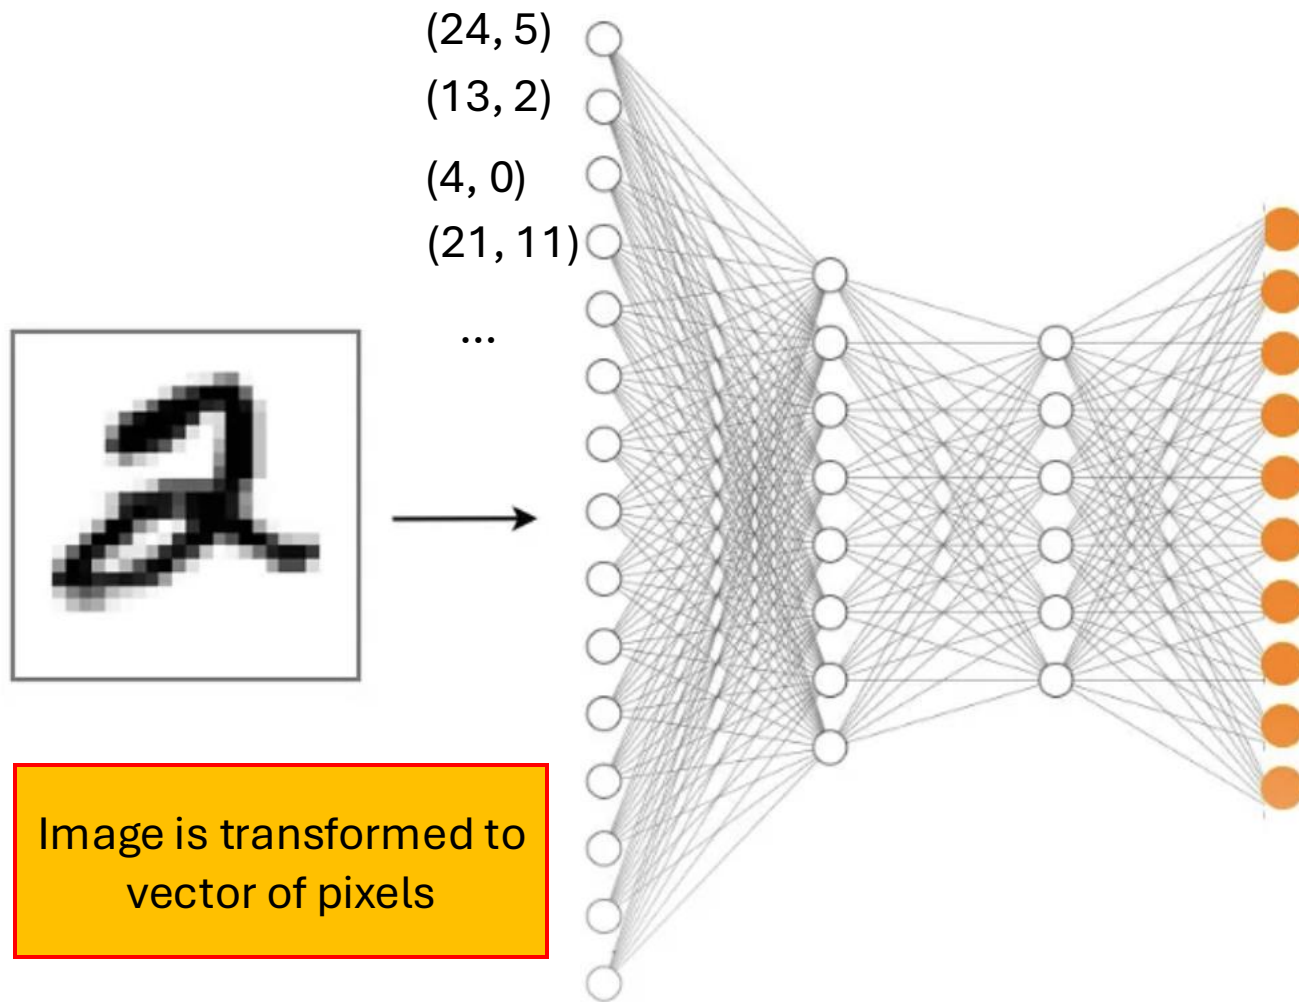


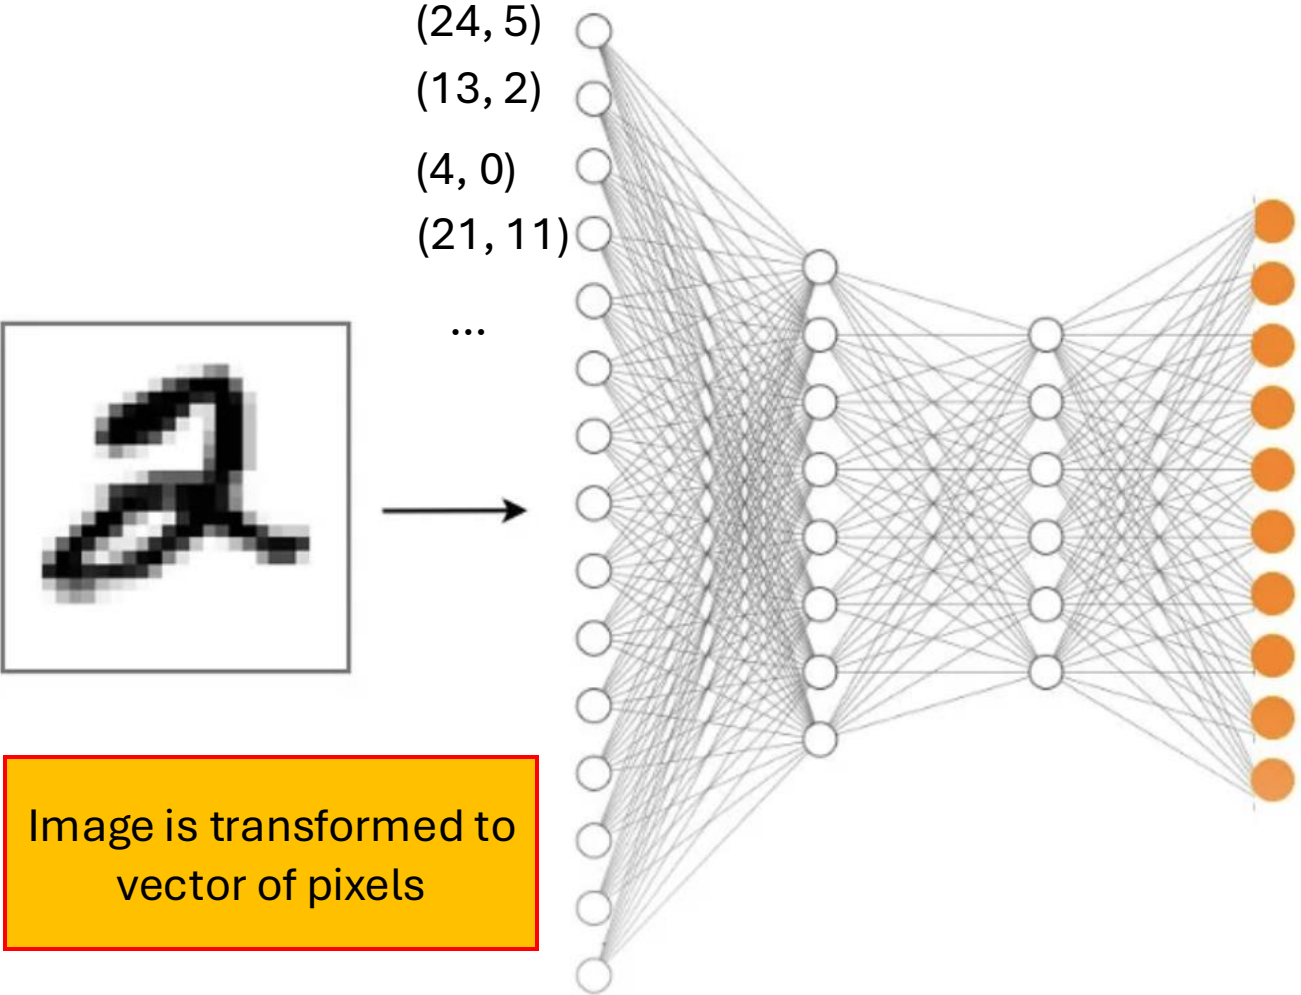
Image is transformed to vector of pixels

What would happen if we permuted the ordering of the pixels?

Will the training of the neural network differ?

No! MLPs do not use spatial information, it does not matter which order the pixels are fed in so long as it is the same ordering for every input


MLPs and Spatial Reasoning




Isn't this actually a hard problem that we are trying to learn?

Limitations of Full Connections for MNIST

If we shift the digit to the right, then a different set of weights becomes relevant $\square \rightarrow$ network might have trouble classifying this as a 1...



#1 encoded as \square



this pixel gets weight 0.6

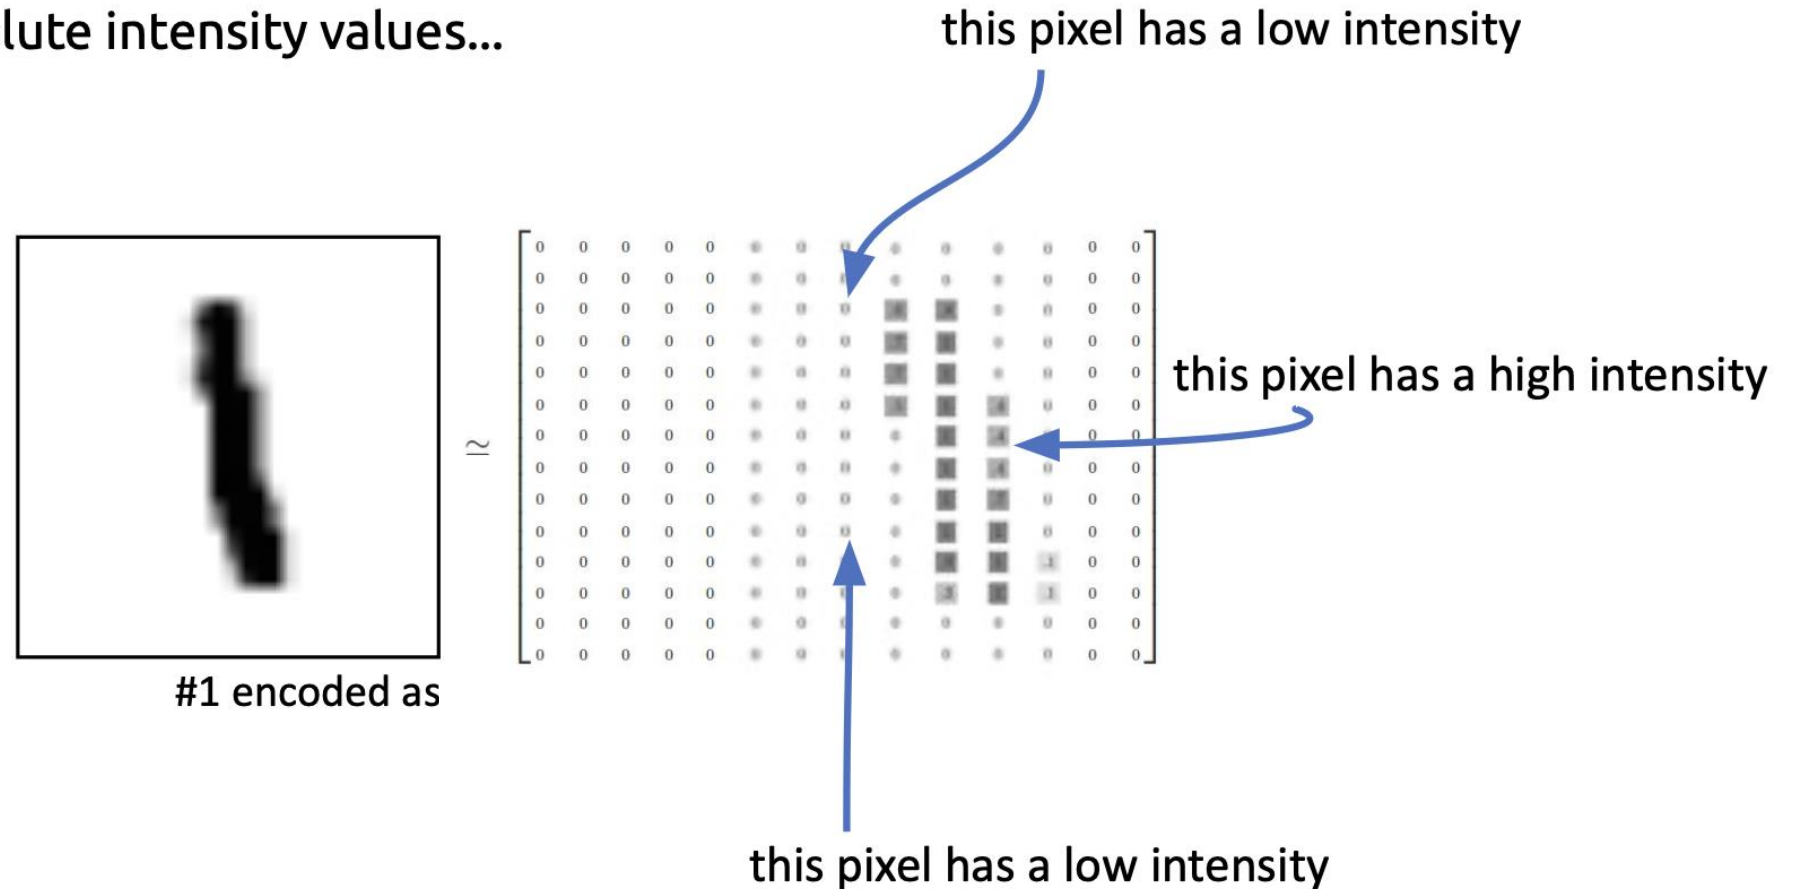
this pixel gets weight 0.1

this pixel gets weight 0.9

Can you tell this is a 1?

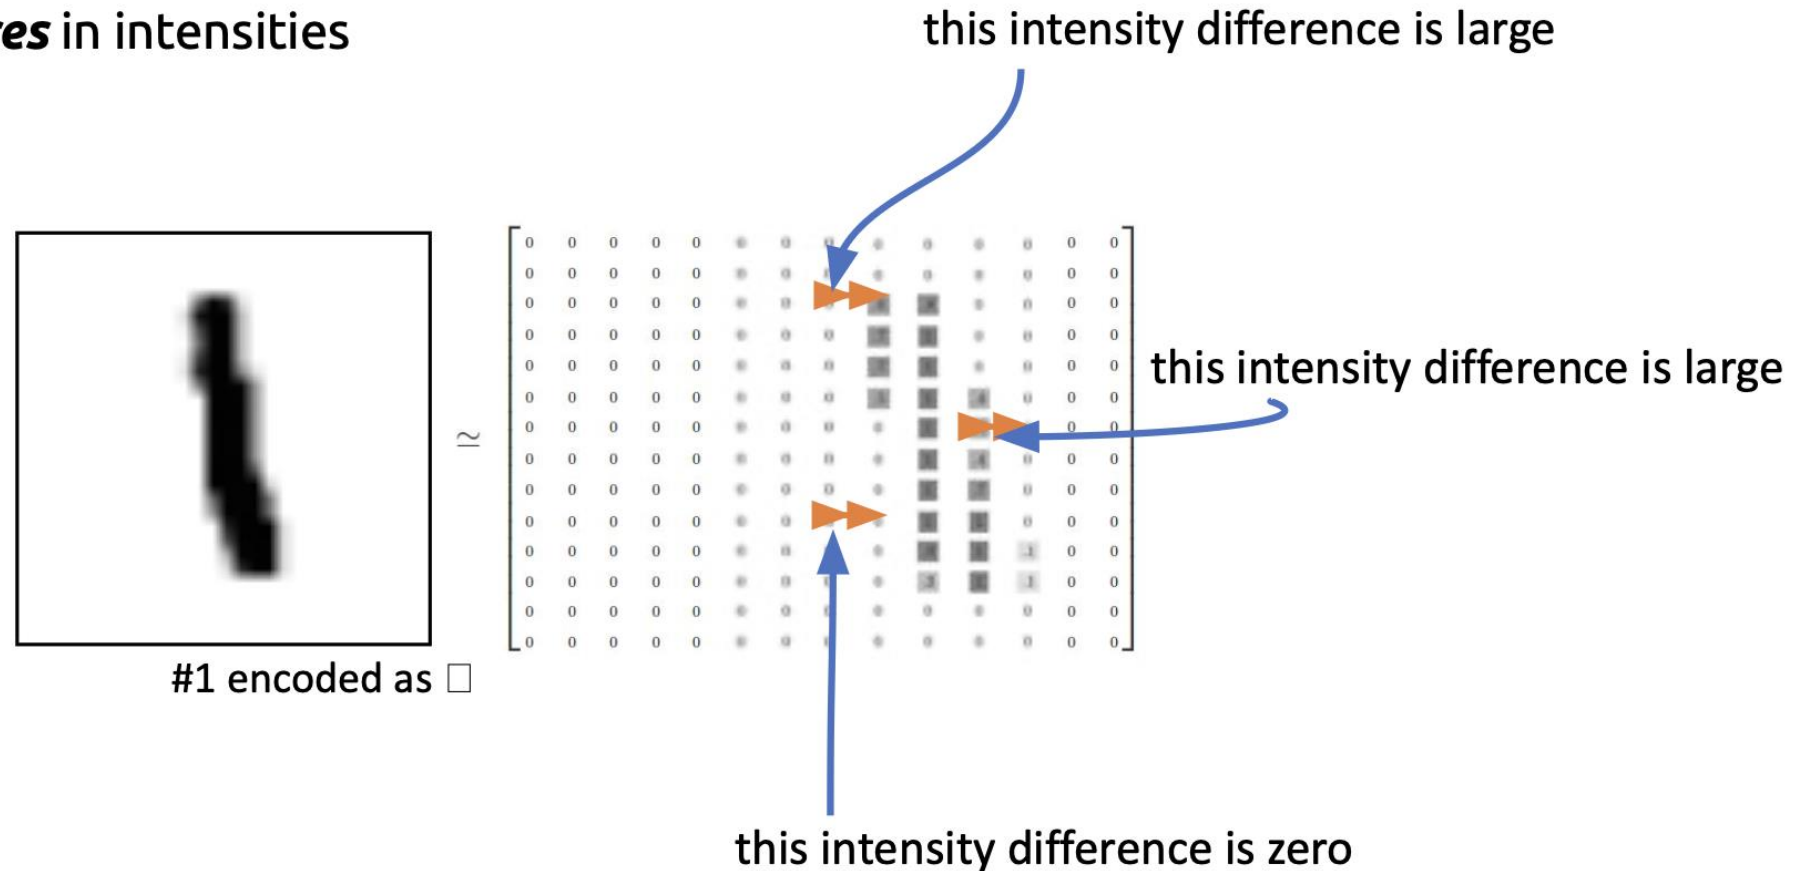
This would *not* be a problem for the human visual system

Our eyes don't look at absolute intensity values...




This would *not* be a problem for the human visual system

...but rather *local differences* in intensities



Fully Connected Nets are *not* Translationally Invariant





Sum of these three: $0.6 \cdot 0.8 + 0.1 \cdot 0 + 0.9 \cdot 1 = 1.38$

Sum of these three: $0.6 \cdot 0 + 0.1 \cdot 0.4 + 0.9 \cdot 0 = 0.4$

MLPs and Spatial Reasoning


How can we change a fully-connected network to account for spatial information?

MLPs (also called fully-connected networks) have weights from every pixel to every neuron




MLPs and Spatial Reasoning

Patches: Pixels close to each other




Advantages of Not Fully Connected Layers

- Fewer weights → Faster?
- The outputs of neurons are “features” for local “patches”
- Incorporates spatial information (pixels that are close together matter)



Disadvantages of Not Fully Connected Layers

- What happens if the image is Translated?
- The patches on the right side were never trained with 5's in that side.



Even though we include spatial **information**, we still don't have spatial **reasoning**. (Can't recognize a shifted 5 is still a 5)

What if we used the same weights for each patch? (Weight Sharing)

The Main Building Block: Convolution

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)

image

2	0	1	3
7	1	1	0
0	2	5	0
0	5	1	4

filter/kernel

1	1	1
0	0	0
-1	-1	-1



(We use this symbol for convolution)
(The verb form is "convolve")

What Convolution Does (Visually)


Overlay the filter on the image

image

1	1	1	3
0	0	0	0
-1	-1	-1	0
0	5	1	4

What Convolution Does (Visually)

Sum up multiplied values to produce output value



What Convolution Does (Visually)

Move the filter over by one pixel

image

1	1	1	3
0	0	0	0
-1	-1	-1	0
0	5	1	4

output

-4	

What Convolution Does (Visually)

Move the filter over by one pixel

image

2	1	1	1
7	0	0	0
0	-1	-1	-1
0	5	1	4

output

-4	

What Convolution Does (Visually)

Repeat (multiply, sum up)

image


2	0_{x1}	1_{x1}	3_{x1}
7	1_{x0}	1_{x0}	0_{x0}
0	2_{x-1}	5_{x-1}	0_{x-1}
0	5	1	4

output

-4	

What Convolution Does (Visually)

Repeat (multiply, sum up)



What Convolution Does (Visually)

Repeat...

image

2	0	1	3
7 _{x1}	1 _{x1}	1 _{x1}	0
0 _{x0}	2 _{x0}	5 _{x0}	0
0 _{x-1}	5 _{x-1}	1 _{x-1}	4

$$7 \times 1 + 1 \times 1 + 1 \times 1 + 0 \times 0 + 2 \times 0 + 5 \times 0 + 0 \times -1 + 5 \times -1 + 1 \times -1$$

output

-4	-3
3	

What Convolution Does (Visually)

Repeat...

image

2	0	1	3
7	1 _{x1}	1 _{x1}	0 _{x1}
0	2 _{x0}	5 _{x0}	0 _{x0}
0	5 _{x-1}	1 _{x-1}	4 _{x-1}

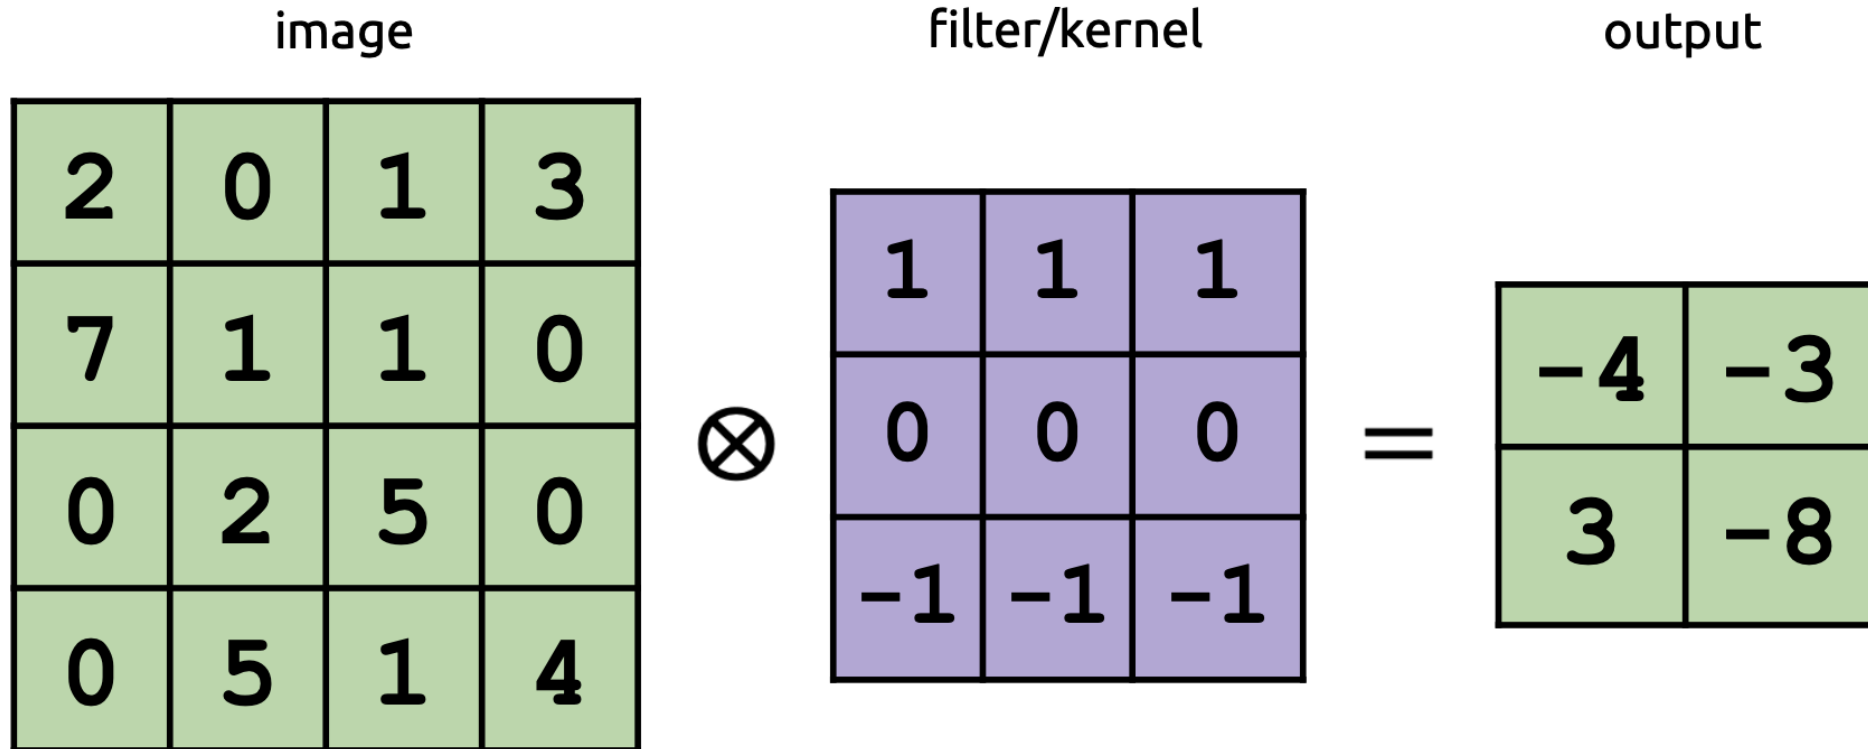
$$1 \times 1 + 1 \times 1 + 0 \times 1 + 2 \times 0 + 5 \times 0 \\ + 0 \times 0 + 5 \times -1 + 1 \times -1 + 4 \times -1$$

output

-4	-3
3	-8

What Convolution Does (Visually)

In summary:



Handmade Kernels and Filters

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Identity kernel

$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Edge detection

$$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$$







Sharpen kernel

$$\frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

Box blur

$$\frac{1}{256} \begin{bmatrix} 1 & 4 & 6 & 4 & 1 \\ 4 & 16 & 24 & 16 & 4 \\ 6 & 24 & 36 & 24 & 6 \\ 4 & 16 & 24 & 16 & 4 \\ 1 & 4 & 6 & 4 & 1 \end{bmatrix}$$

Gaussian blurr kernel

Operation	Kernel ω	Image result $g(x,y)$
Identity	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	
Ridge or edge detection	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	
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 3 x 3 (approximation)	$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$	

What Comes Next?

Can we learn a filter for our images rather than “hand crafting” one?

Recap

Fully Connected Neural Networks for images lack spatial information

Participation Quiz #3 is up!

We can add spatial reasoning by connecting pixels that are “spatially” close

Convolutions/Filters/Kernels are a technique from image processing that combine close pixels with a linear transformation