



Secure Autonomous Systems

CSCI 6907/3907 86

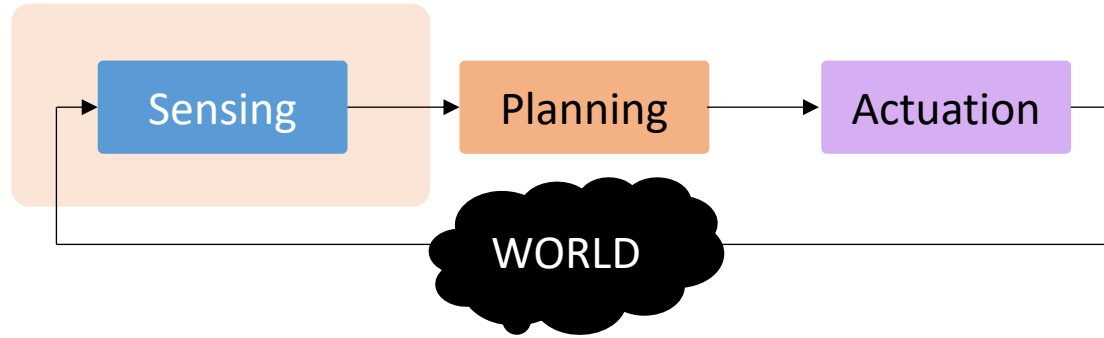
Spring 2024

Prof. Sibin Mohan

<https://bit.ly/secureauto-spring24>



Sensing, Planning, Actuation Loop



Sensors

“perception”

Interface to the external world

Sensor “plan” will balance

Function

Range

Cost

Common Sensor Types



LiDAR/Millimeter Radar



GPS



Cameras:

lane, road, traffic light,
stop line, surroundings



Stereo Vision




Sonic


LiDAR




- **Light Detection And Ranging**
- Laser scanning/3D scanning
- Uses eye-safe laser beams → create 3D representation of environment



A typical lidar sensor emits pulsed light waves into the surrounding environment



These pulses bounce off surrounding objects and return to the sensor





The sensor uses the time it took for each pulse to return to the sensor to calculate the distance it traveled

LiDAR Output

Point clouds In 3D

Range: **70-100 m**

View: **360 degrees**



Millimeter Wave Radar [mmWave]

- Radar technology
- short-wavelength electromagnetic waves
- Measures reflected radar signals

- **High accuracy**
- 76-81 GHz → detect movements in a **fraction of a millimeter!**
- Limited distance [$< 80\text{m}$]


- Also, used for in-cabin monitoring of drivers





Cameras

- Accurate way to create visual representations
- Front, left, right, rear cameras
 - to create a **360-degree view**
- Main focus → **object detection**






Cameras | Computer Vision


- Computer Vision algorithms for object detection
 1. Image classification → determine **what** the objects in an image are
 2. Image localizations → providing specific **locations** of image [bounding box]

Cameras | Image Classification

- Convolutional Neural Networks (CNNs)
 - **trained** to recognize objects like cars, pedestrians, etc.
 - performs **convolution operations at runtime**
 - to classify images from camera




Deep Learning



CNNs

- A class of deep neural networks
- Mostly applied for visual imagery
- Uses **mathematical convolutions**
 - operation on two functions
 - produces a third function
 - expresses shape of how one modifies the other

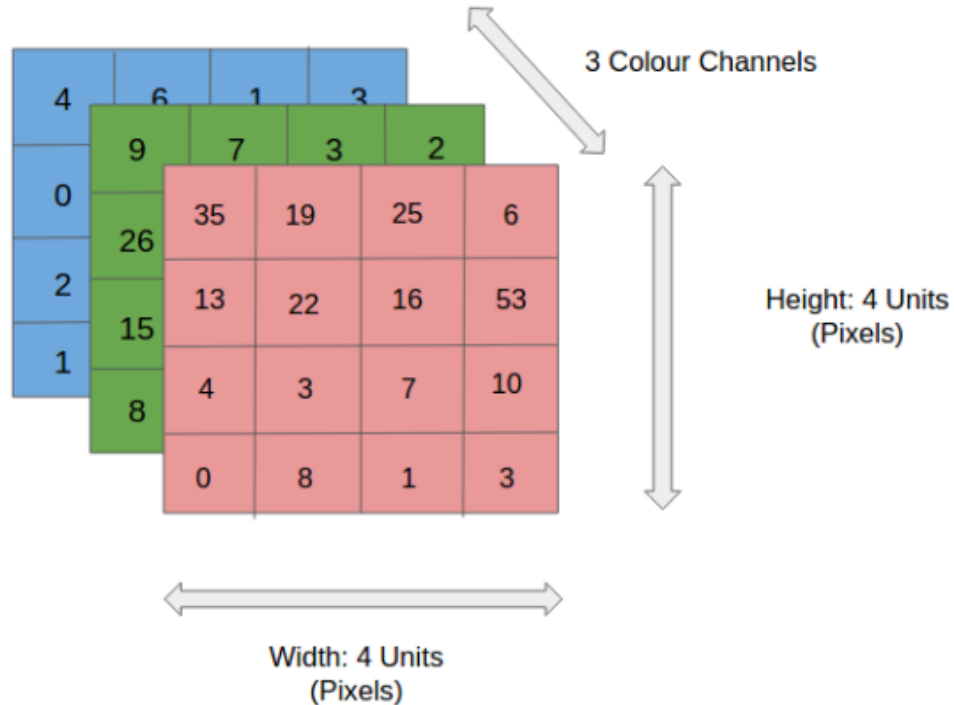



How Do Convolutions Work?

Consider an image



Input




for simplicity,
consider
a **grayscale**
version




Input

[Greyscale] Convolution





Kernel/Filters

- Filters exploit **spatial locality** of a particular image
- A small matrix used for → *blurring, sharpening, edge detection, etc.*
- **Enforcing connectivity** between neurons




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





“sharpen” kernel

Kernel/Filter Example [contd.]









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



“ridge/edge detection” kernel

Kernels
Come in
Various
Shapes



0	0	0	0	0
0	1	1	1	0
0	1	1	1	0
0	1	1	1	0
0	0	0	0	0

(a)

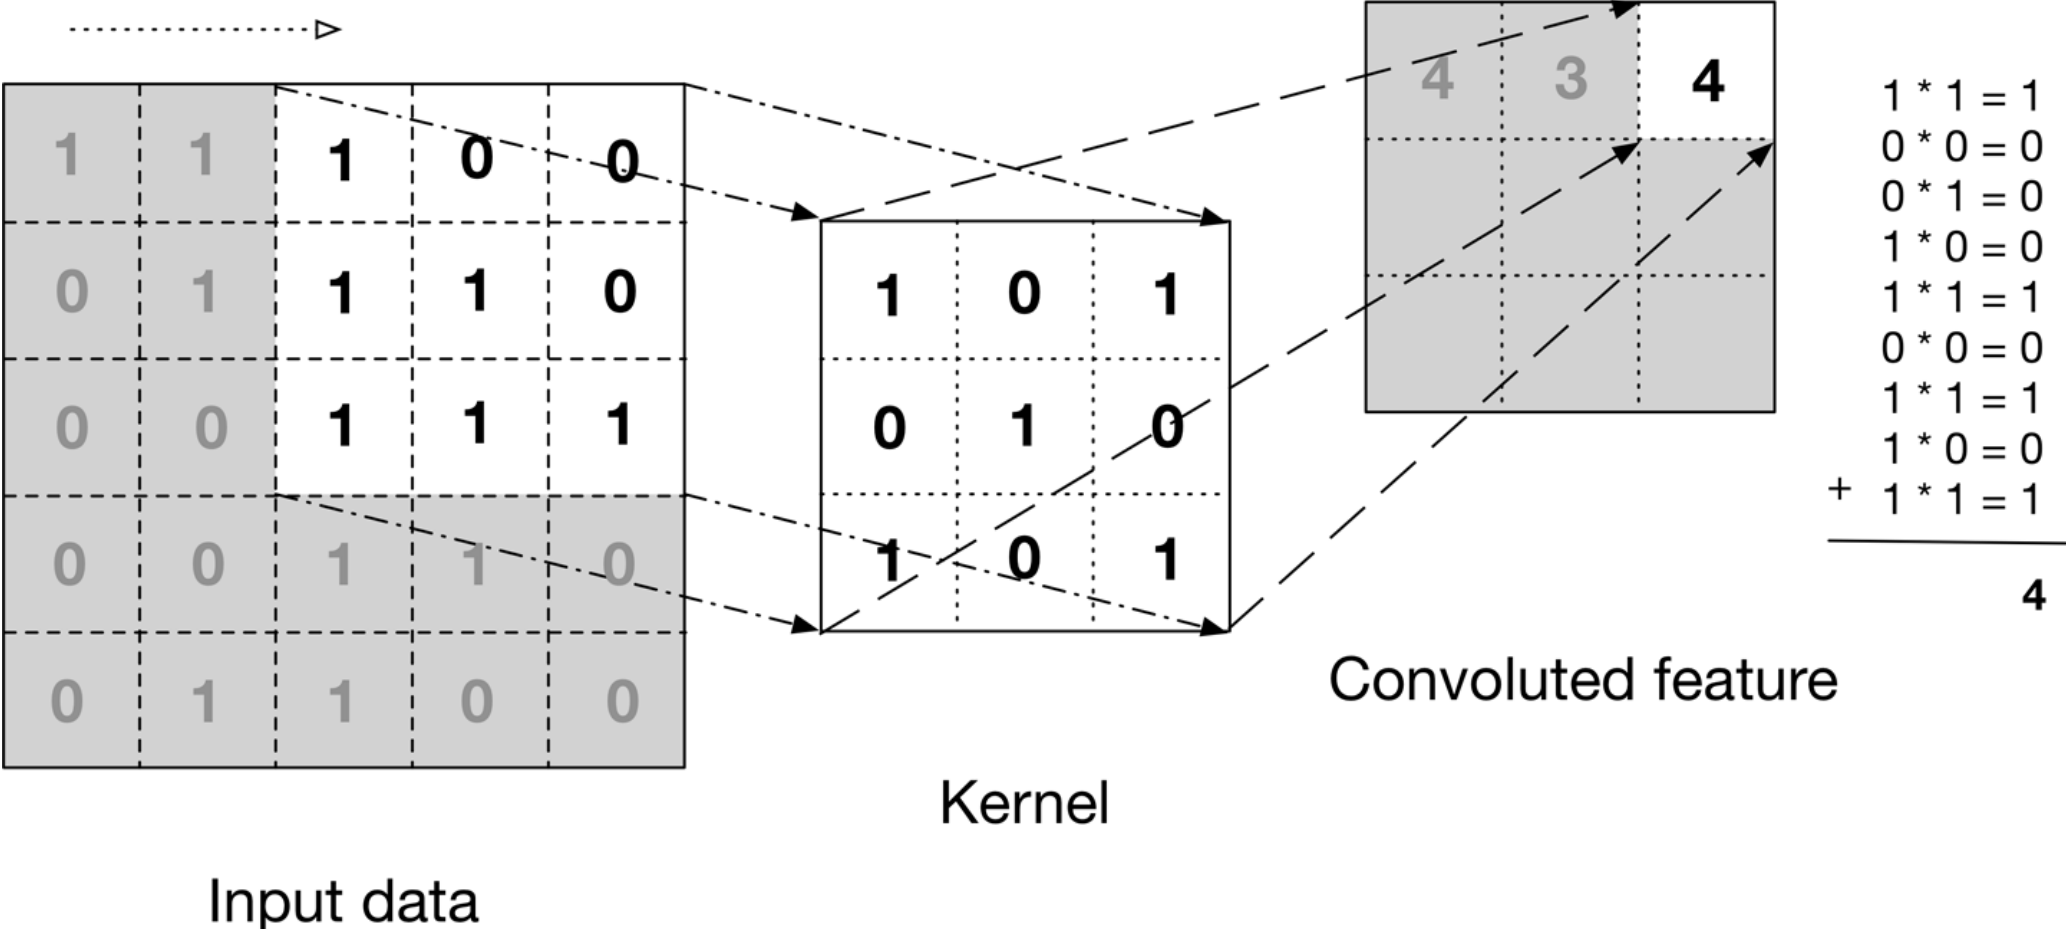
0	1	2	1	0
1	3	5	3	1
2	5	9	5	2
1	3	5	3	1
0	1	2	1	0

(b)

0	0	-1	0	0
0	-1	-2	-1	0
-1	-2	16	-2	-1
0	-1	-2	-1	0
0	0	-1	0	0

(c)

[Greyscale] Convolution



[Greyscale]
Convolution
[contd.]

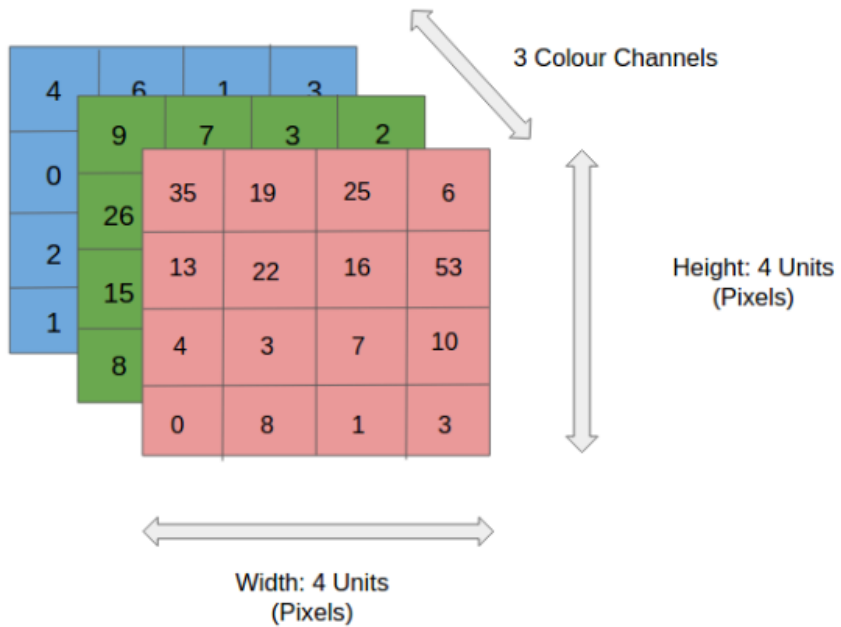
1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image


4		

Convolved
Feature

but we deal with color images!



[Color] Convolutions



So, then...what are CNNs?

- multiple layers of **artificial neurons**
- mathematical functions → calculate weighted sum of inputs/outputs
- on **activation value**




Why Multiple layers?


Each layer has a unique function



Pooling Layer?




- **reduce spatial size** of convolution layer
- **reduce computational loads**



“max” pool

Issues with CNNs

- limited to single objects taking up entire image
- multiple objects in image?
- **sliding windows!**




Cameras | Sliding Window Algorithm





Issues to Consider



what is the **window size?**

- start with large/small window and keep increasing?



accuracy?

- bounding box may not be accurate
- rectangular boxes?



computational cost!

References

- Velodyne LiDAR Video:

https://sibin.github.io/teaching/csci6907_88-gwu-secure_autonomous/fall_2022/other_docs/What-is-Lidar-video.mp4

- mmWave Documentation

https://www.ti.com/lit/wp/spyy005a/spyy005a.pdf?ts=1641417836995&ref_url=https%253A%252F%252Fwww.google.com%252F