Autonomous Car Security

Hoda A.Alkhzaimi

Director of the Center for Cyber Security

Research Assistant Professor, Department of Engineering, NYUAD

Hoda.Alkhzaimi@nyu.edu Nyuad.nyu.edu/ccs-ad

SDCONGRESS/SDDUBAI

CENTER FOR CYBER SECURITY

جامعة نيويورك ابوظبي

NYU ABU DHABI

Agenda

- Autonomous Car Security: A Global Landscape
  - Statistical information on Countries
  - Investments
- Opportunities and Risks
  - Attack landscape

Autonomous Car Security: A Global Landscape

Definitions, Statistics and Current State of Art

What is AI **Vs** Machine Learning **Vs** Deep Learning?

- English Oxford Living Dictionary : " "The t**heory and development** of computer systems able to **perform** tasks normally **requiring human intelligence**, such as **visual perception**, **speech recognition**, **decision-making**, and **translation between languages**."."
- Merriam-Webster : "A branch of computer science dealing with the **simulation of intelligent behavior** in computers. The capability of a machine to **imitate intelligent** human **behavior**."

Source: The Independent

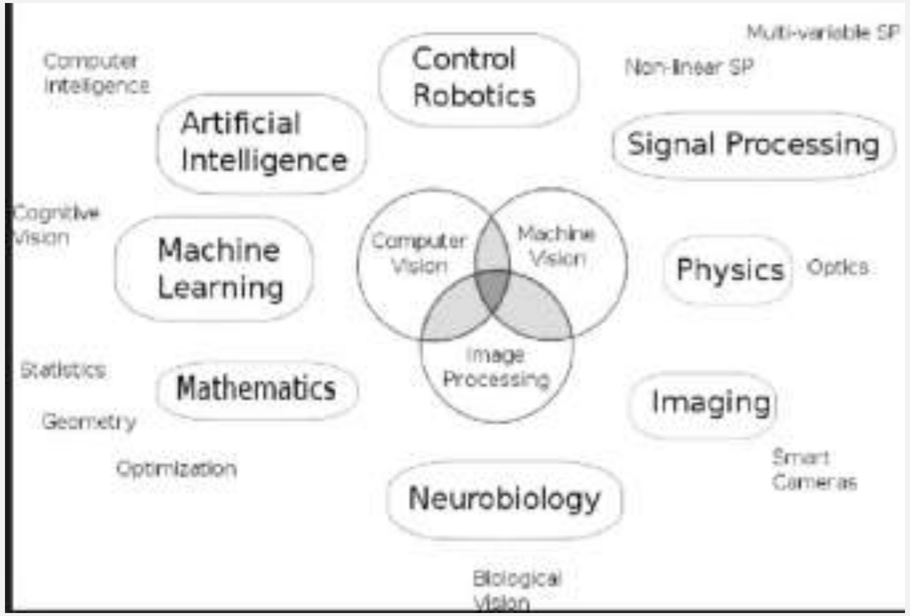

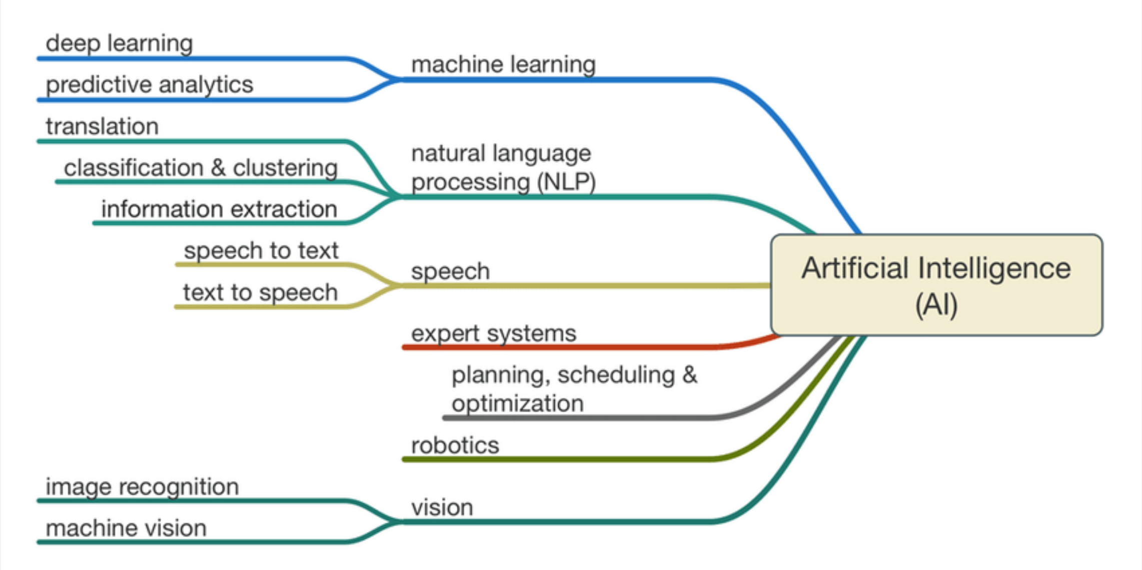

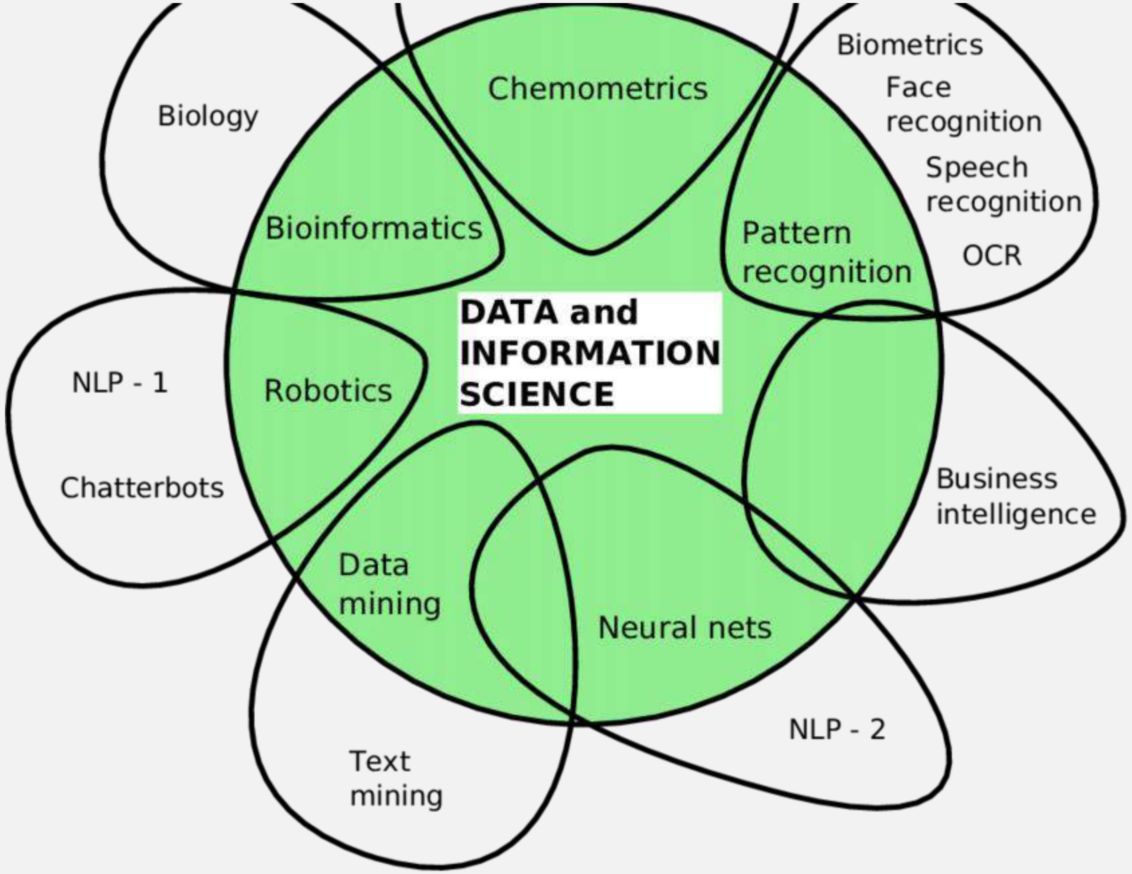

Source: researchgate.net

What is AI **Vs** Machine Learning **Vs** Deep Learning?

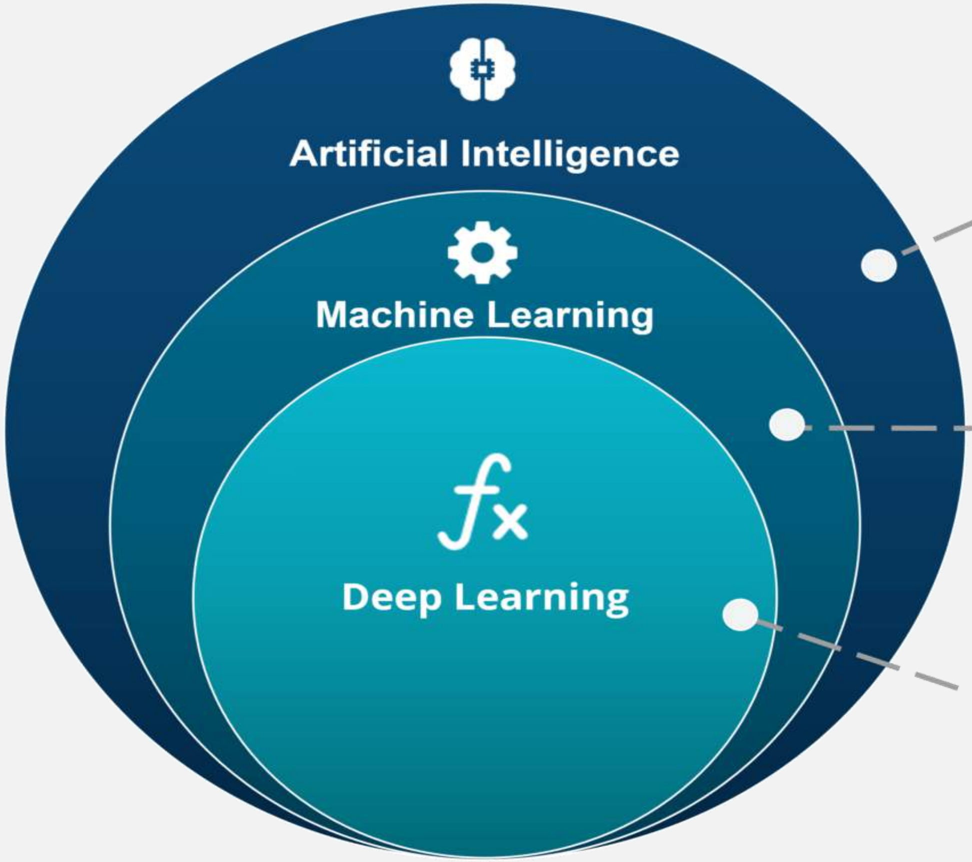

ARTIFICIAL INTELLIGENCE
A technique which enables machines to mimic human behaviour

MACHINE LEARNING
Subset of AI technique which use statistical methods to enable machines to improve with experience

DEEP LEARNING
Subset of ML which make the computation of multi-layer neural network feasible

Source: Edureka

INVESTMENTS (IN BILLIONS)

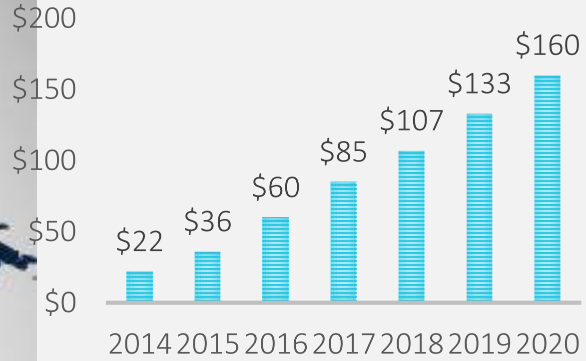

Expected Growth on Number of Cities

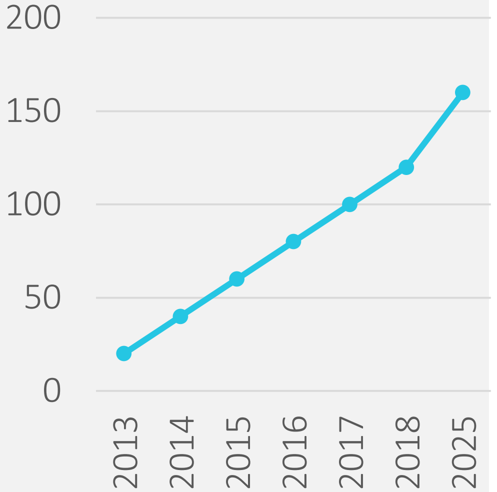

Expected Number of Connected Devices by 2020: **50 billion Devices**

Number of Cities in the world:4416→**11M/city**
Number of Countries: 195 → **257M/Country**

Smart Societies:

Numerical Outlook

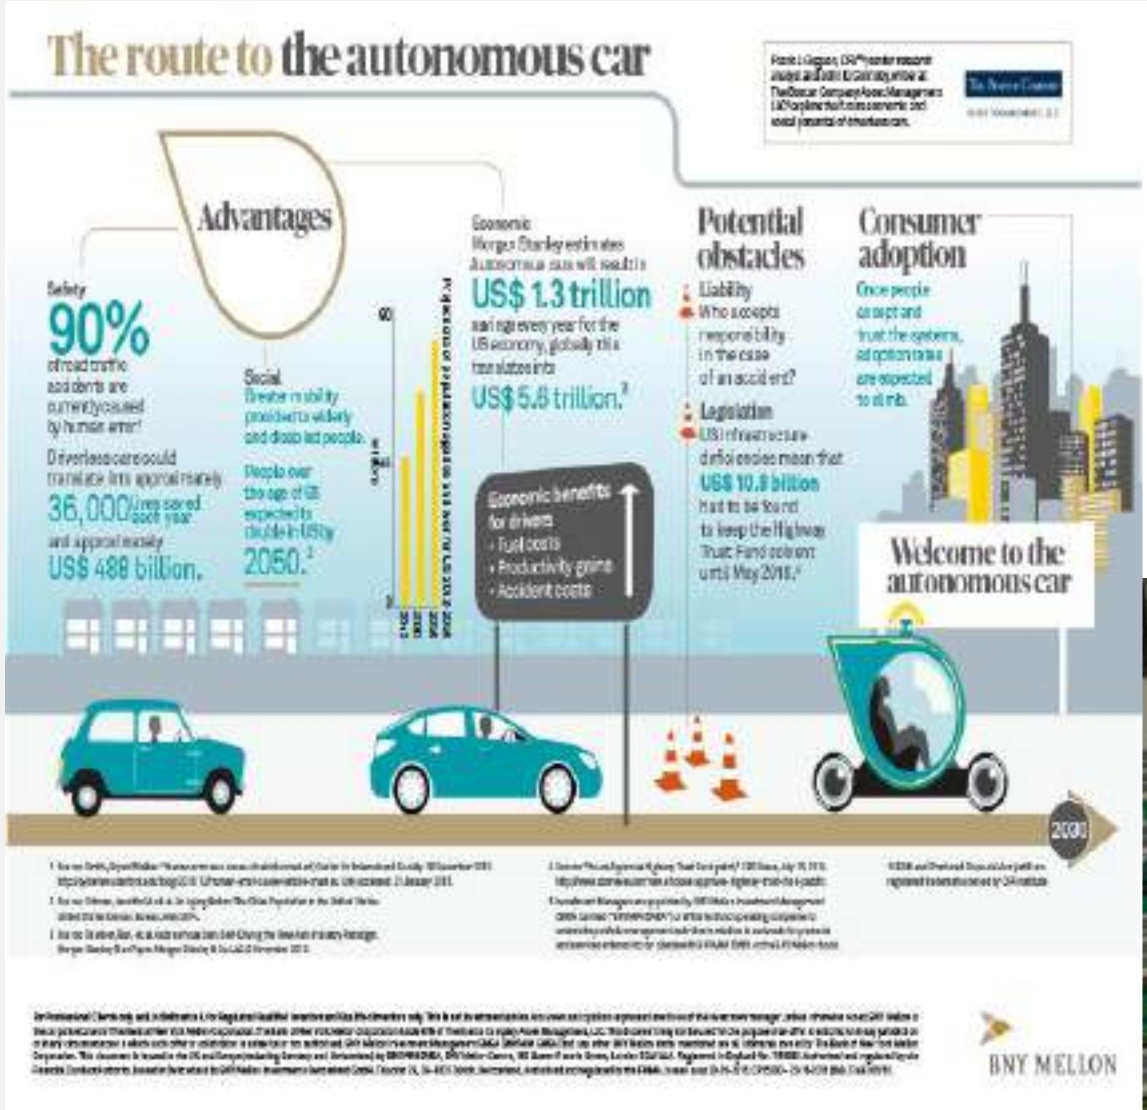

% of new-car buyers that (strongly) agree with the statement

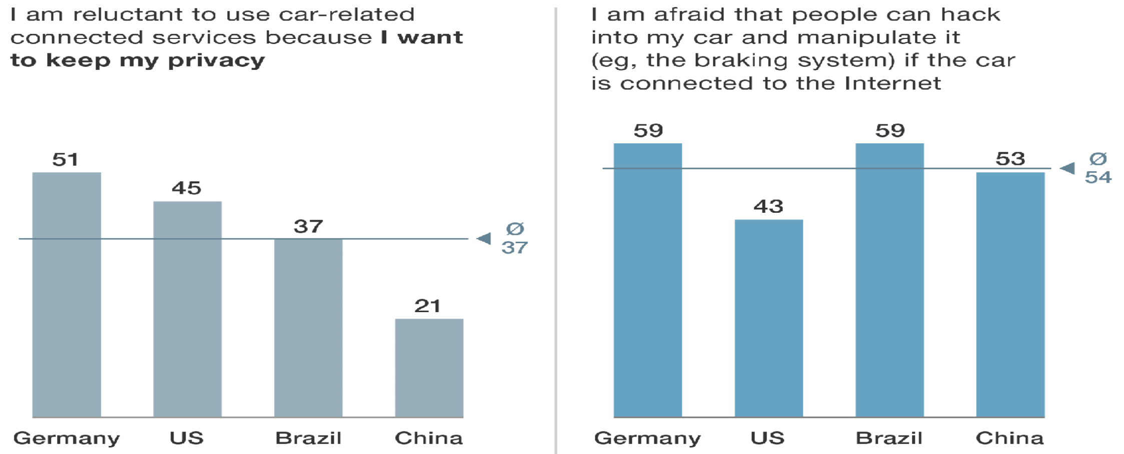

Source: McKinsey's Connected Car Consumer Survey, 2014

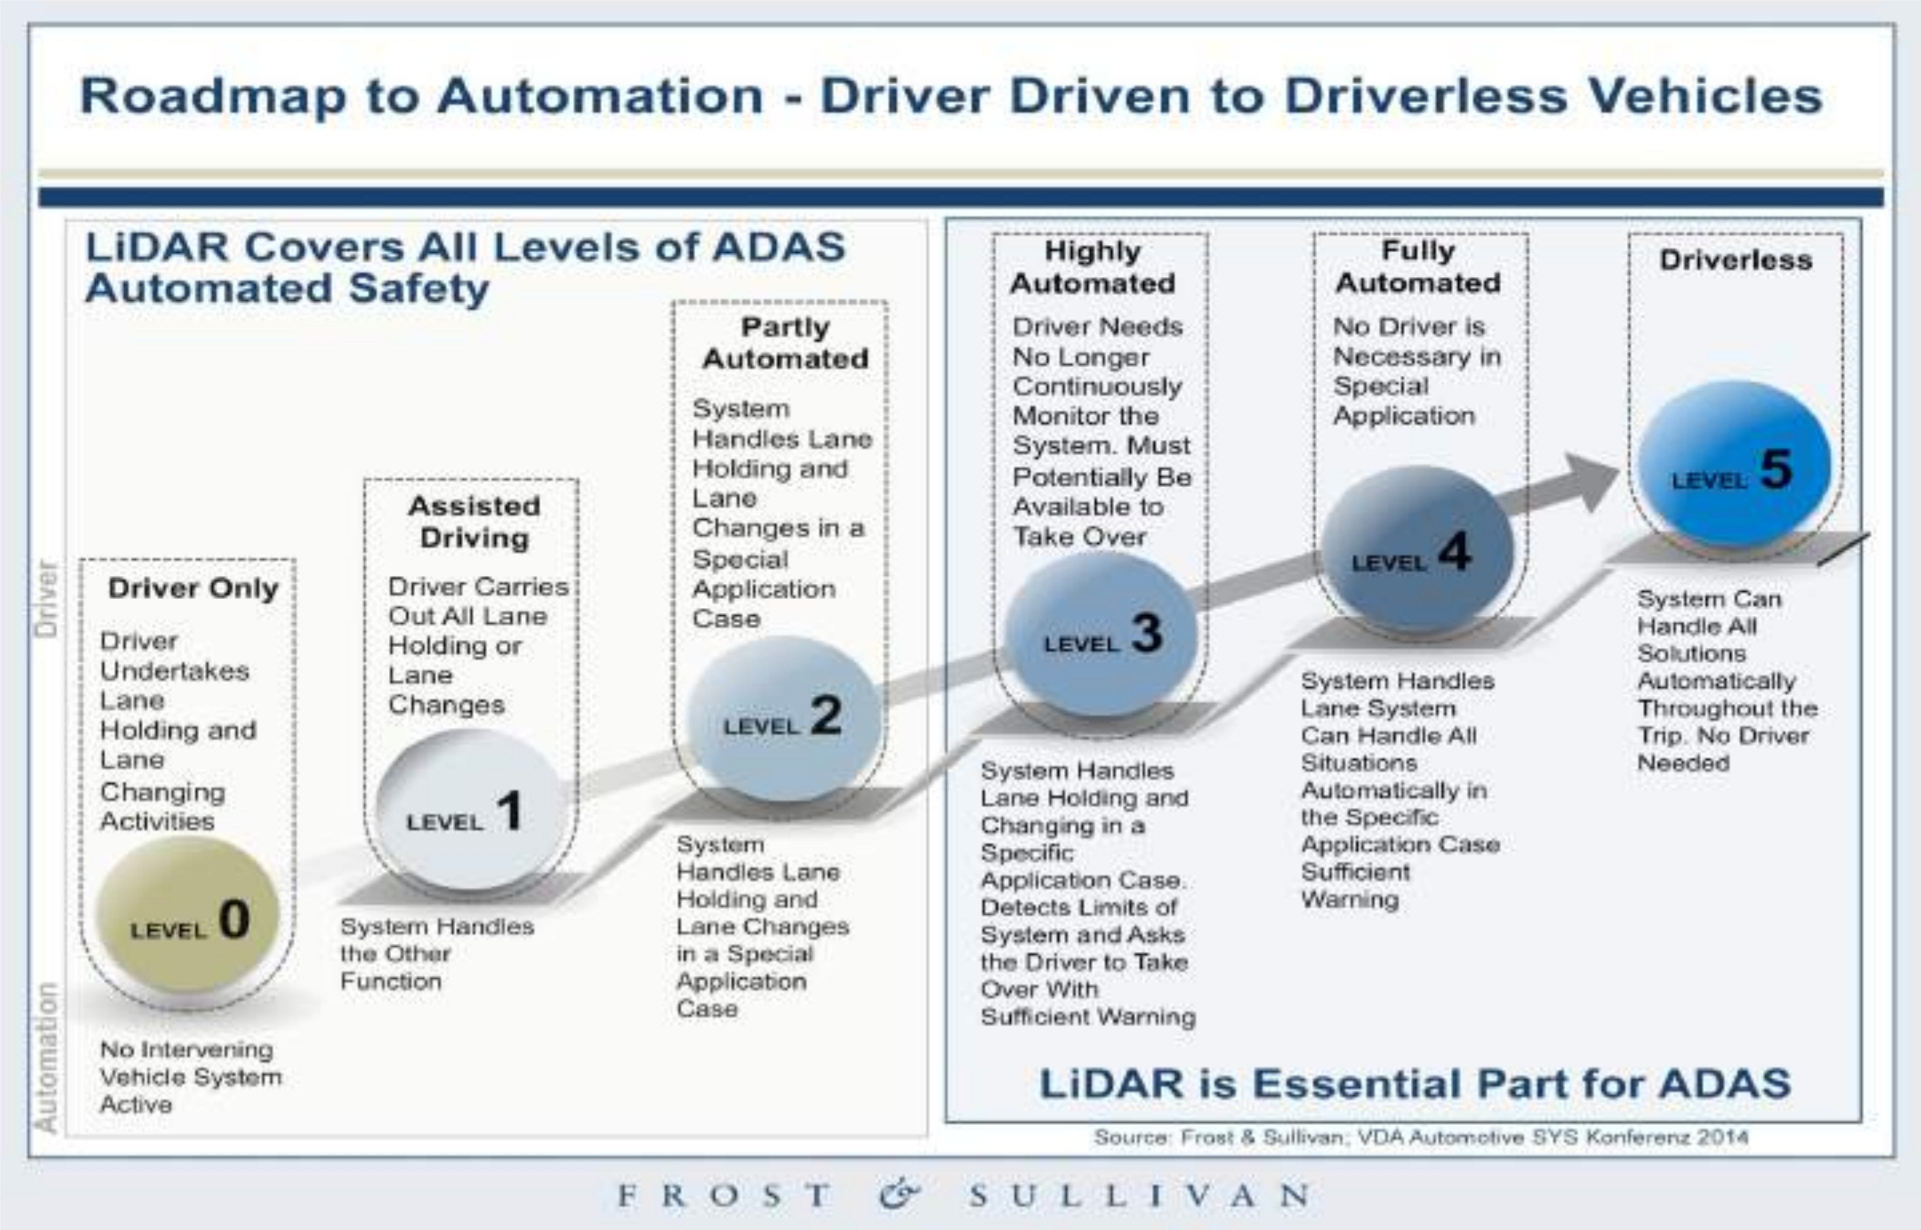
Roadmap to Automation - Driver Driven to Driverless Vehicles
LiDAR Covers All Levels of ADAS Automated Safety
Driver
Automation
Driver Only
Driver Undertakes Lane Holding and Lane Changing Activities
LEVEL 0
No Intervening Vehicle System Active
Assisted Driving
Driver Carries Out All Lane Holding or Lane Changes
LEVEL 1
System Handles the Other Function
Partly Automated
System Handles Lane Holding and Lane Changes in a Special Application Case
LEVEL 2
System Handles Lane Holding and Lane Changes in a Special Application Case
Highly Automated
Driver Needs No Longer Continuously Monitor the System. Must Potentially Be Available to Take Over
LEVEL 3
System Handles Lane Holding and Changing in a Specific Application Case. Detects Limits of System and Asks the Driver to Take Over With Sufficient Warning
Fully Automated
No Driver is Necessary in Special Application
LEVEL 4
System Handles Lane System Can Handle All Situations Automatically in the Specific Application Case Sufficient Warning
Driverless
LEVEL 5
System Can Handle All Solutions Automatically Throughout the Trip. No Driver Needed
LiDAR is Essential Part for ADAS
Source: Frost & Sullivan; VDA Automotive SYS Konferenz 2014
FROST & SULLIVAN

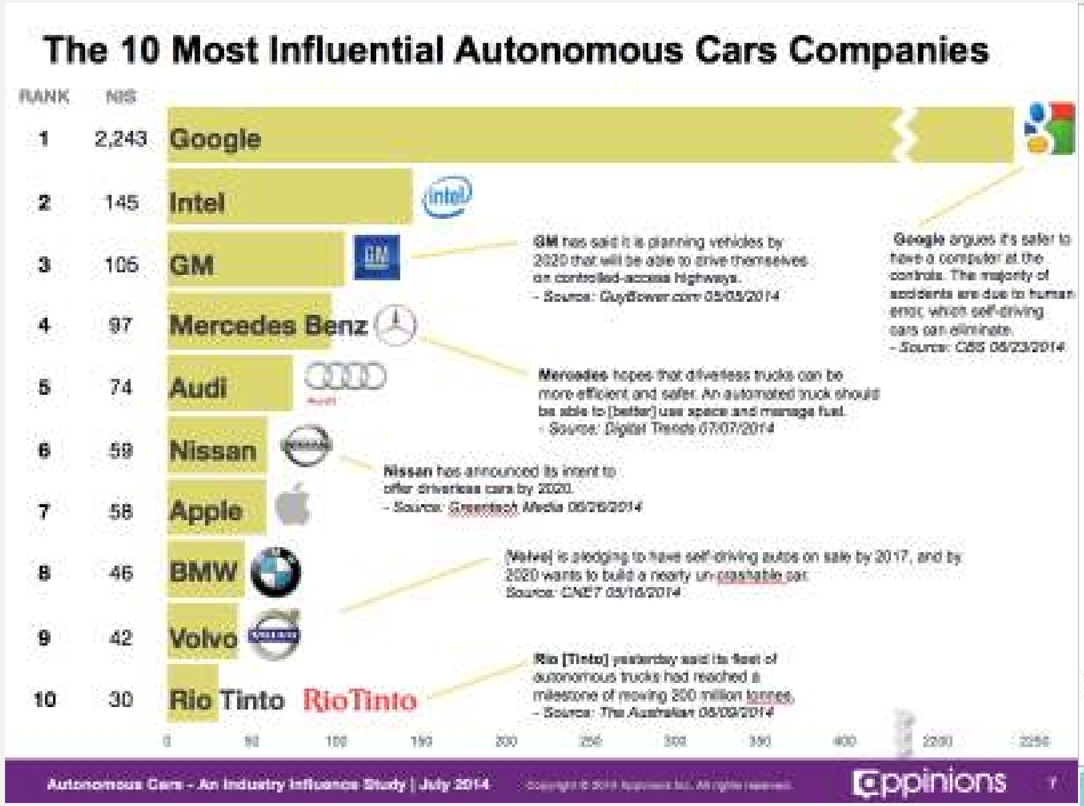
The 10 Most Influential Autonomous Cars Companies
RANK
NIS
1 2,243 Google
2 145 Intel
3 105 GM
4 97 Mercedes Benz
5 74 Audi
6 59 Nissan
7 58 Apple
8 46 BMW
9 42 Volvo
10 30 Rio Tinto
Autonomous Cars – An Industry Influence Study | July 2014
appinions

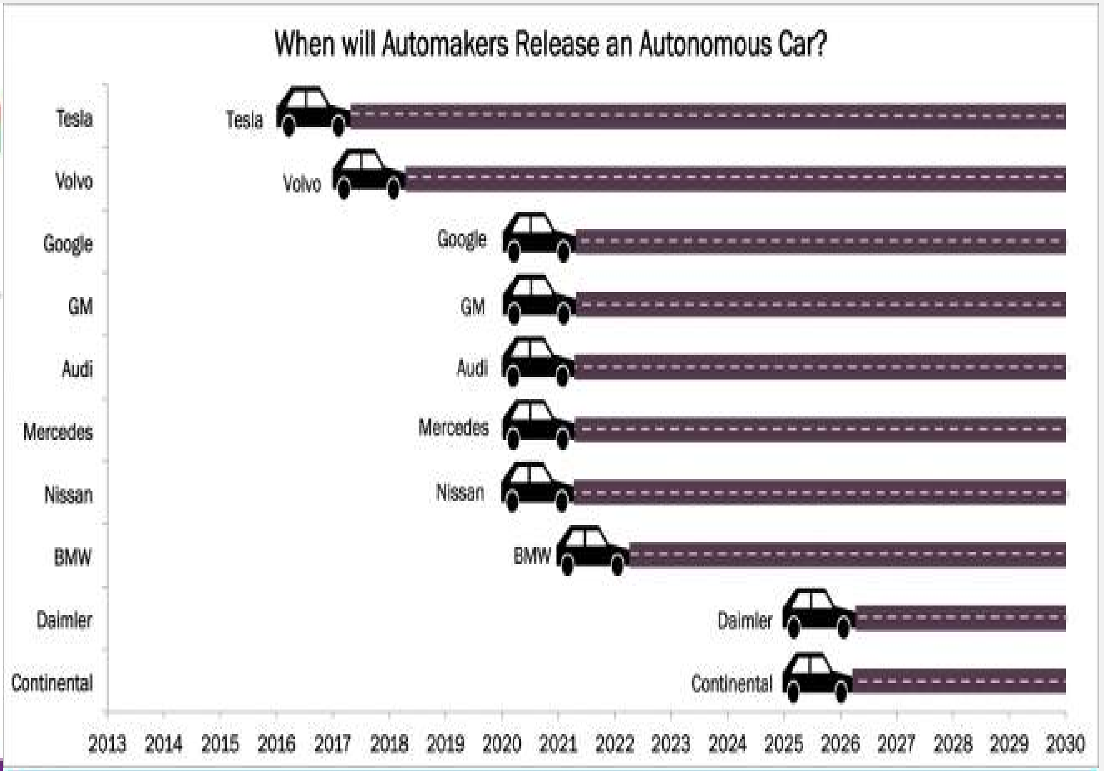
When will Automakers Release an Autonomous Car?
Tesla
Volvo
Google
GM
Audi
Mercedes
Nissan
BMW
Daimler
Continental
2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023 2024 2025 2026 2027 2028 2029 2030

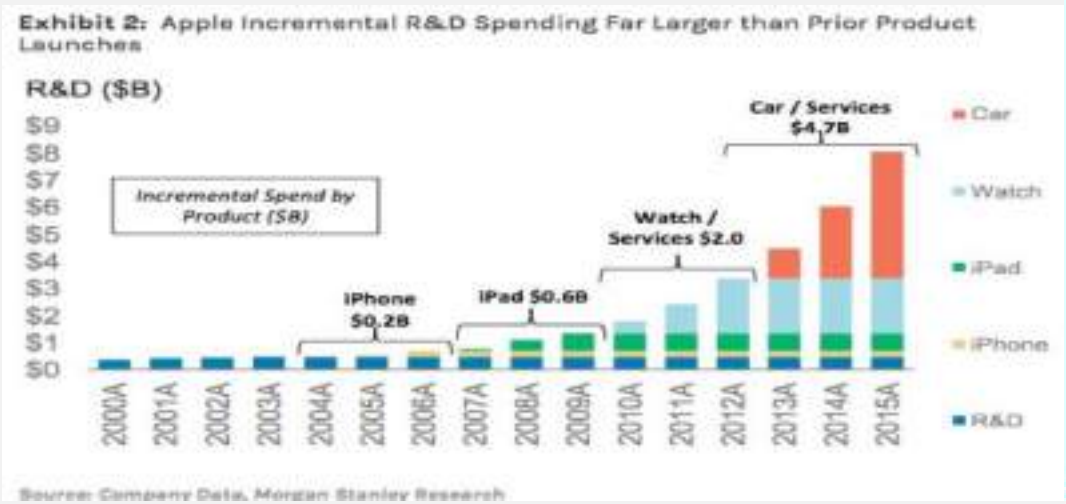
Exhibit 2: Apple Incremental R&D Spending Far Larger than Prior Product Launches
R&D ($B)
$9
$8
$7
$6
$5
$4
$3
$2
$1
$0
Incremental Spend by Product ($B)
iPhone $0.2B
iPad $0.6B
Watch / Services $2.0
Car / Services $4.7B
2000A
2001A
2002A
2003A
2004A
2005A
2006A
2007A
2008A
2009A
2010A
2011A
2012A
2013A
2014A
2015A
Car
Watch
iPad
iPhone
R&D
Source: Company Data, Morgan Stanley Research

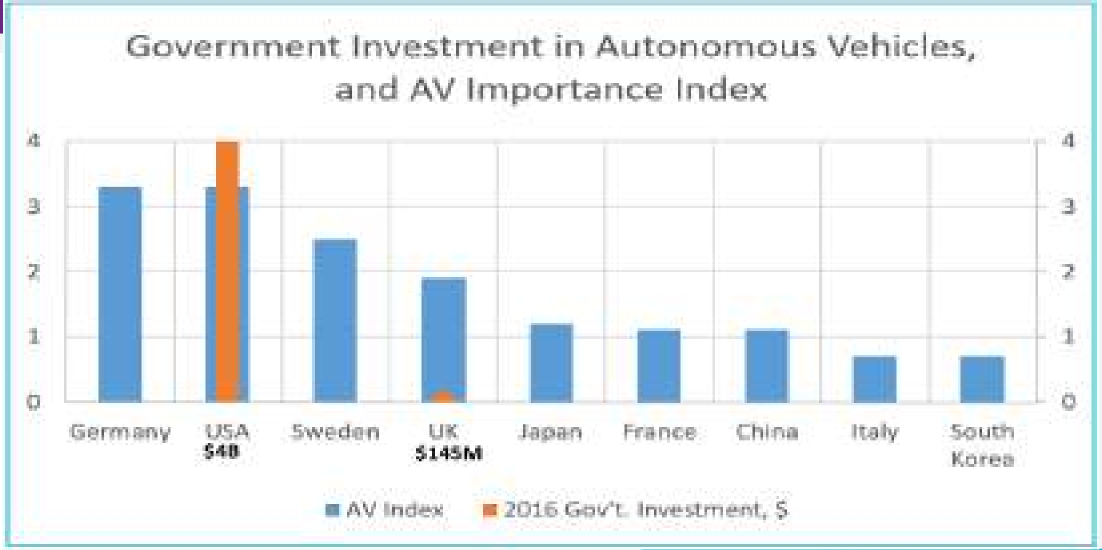
Government Investment in Autonomous Vehicles, and AV Importance Index
0
1
2
3
4
Germany
USA $4B
Sweden
UK $145M
Japan
France
China
Italy
South Korea
AV Index
2016 Gov't. Investment, $

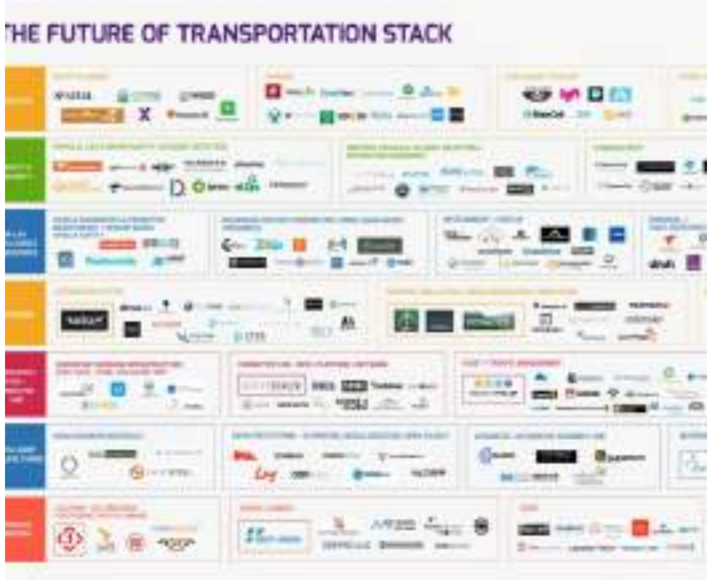

By 2035, 12 million full AV units could be sold a year globally

Market for partial and full AV features expected to grow from ~$42B in 2025 to ~$77B in 2035

In 2035, 25% of market to be AV sales with 15% partial and 10% full AV systems

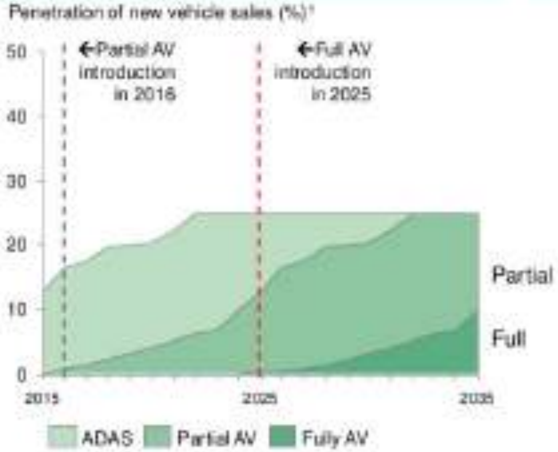

Represents 12M full AVs and ~18M partial; ~$77B market for AV features in 2035

	2025 global sales			2035 global sales		
	Share (%)	Volume (M)	Sales[3] ($B)	Share (%)	Volume (M)	Sales[3] ($B)
	Estimated global new light vehicle sales: ~111M[2]			*Estimated global new light vehicle sales: ~122M[2]*		
Partial	12.4%	13.9	36	15.0%	18.4	38
Full	0.5%	0.6	6	9.5%	12.0	39
Total			$42bn			$77bn

Opportunities and Challenges

Autonomous Car Security

The Holistic Framework

- **Stakeholders** and **Players** : (Private, Public, Education, Industry, Government)
- **Cyber Resilience**: (Security and Privacy)
- **Policies**, **Standards** and **Regulations**
- Measurable Models
- Data Management and Utilization

- Governance (Services, Happiness, Ethical Development, Social Development)

Mobility	Agriculture/ Food Security	Critical Infrastructure	Health Care	Water and Climate Management	Education

Failure and Risks


AUTOMOTIVE

Tesla self-driving car fails to detect truck in fatal crash

Tesla, based out of Palo Alto, is now under a federal investigation after the death of 40-year-old Joshua Brown, a former Navy Seal and entrepreneur.

By Jessica Castro

Friday, July 01, 2016

PALO ALTO, Calif. (KGO) -- There are new details out on Tesla's much-talked-about self-driving technology, which is now under investigation, after a man died in a car crash in Florida.

Tesla, based out of Palo Alto, is now under a federal investigation after the death of 40-

Failure and Risks

Connected Cars Threats

Car Hacking in 30 Minutes or Less

October 20, 2018 | Computer Science, Social Media, Technology

This tutorial will guide you step-by-step into one of the hottest cyber skills in the world: car hacking!


My Facebook Live video on car hacking brought in over 3,600 viewers!

Using VirtualBox and Kali Linux, you can start car hacking using completely free open-source software and tools, including can-utils, ICSim, ScanTool, Wireshark, and tcpdump. You'll create a functioning CAN (controller area network) simulator with a dashboard just like the one in your car. I'll even show you how to perform a "replay attack", recording packets from the CAN bus and replaying them to change the car's settings, sensors, and controls! When you're ready to try you skill on an actual automobile, you can buy (or build) the hardware for $20 to $75 (USD) and connect to your vehicle

I'm a Certified Ethical Hacker, and I've always worked on my own cars. But when I started car hacking a little over 10 years ago, there weren't any clear, step-by-step, easy-to-follow instructions for beginners. Ten years later, there still aren't any complete car hacking beginner-to-pro tutorials out there that show you how to get started and how to actually connect to a real car or truck, so I wrote this post based on my popular presentations at hacking conferences and workshops (GenCyber, CCERP, NCWA, STEM, and more).

In this tutorial, you'll be able to:

- install free, open-source car-hacking tools on your computer (desktop, laptop, even a Raspberry Pi, or on a Linux virtual machine),
- perform a replay attack on a simulated controller area network (CAN),
- buy or build the low-cost tools necessary to test for similar vulnerabilities in modern automobiles.

If you've got any ethical hacking experience, you can use a Kali Linux VM (the most popular ethical hacking toolkit) for your car-hacking workstation. And the only physical equipment needed to connect the software to a modern car, a USB to OBD-II cable or wireless connector, can be built or bought easily online for around $20-75 (USD).

Introduction to the CAN Bus

Car hacking itself is surprisingly similar to hacking other networked devices. We can use a network sniffer to view packets as they move across the controller area network, or CAN bus, in an automobile. The CAN (controller area network) bus enables communication between the vehicle's sensors and its various electronic control units (ECUs). Modern production cars can have as many as 70 or more ECUs controlling the engine, [redacted]ainment system, and more. One advantage of the CAN bus architecture is that it allows all of [redacted]air of shared wires running to each sensor, controller,and ECU.

[redacted]s of data inside another 8-byte package), with no addresses in the messages, just a priority [redacted]er priority than the air conditioning or audio player). The CAN bus protocol was not built with modern security in mind.

Brief History of Car Hacking

Adversarial Learning: A Deep Learning Approach for Self-Driving

1. Build an **object detector** that is able to:
   - **Efficiency**
   - Detect different objects found in **autonomous driving-car environment** (cars, pedestrians, traffic lights etc.)
   - **Scalability**: build both small and large **scale models** in multiple **orientations**.
2. Build **a generator** that can **attack** the object detector that will:
   - **Attack:** Creating **adversarial examples to fool** the deep learning-based system
   - **Defend:** Improve the object detector to be able to withstand different types of adversarial attacks.

Evaluation on VOC2007

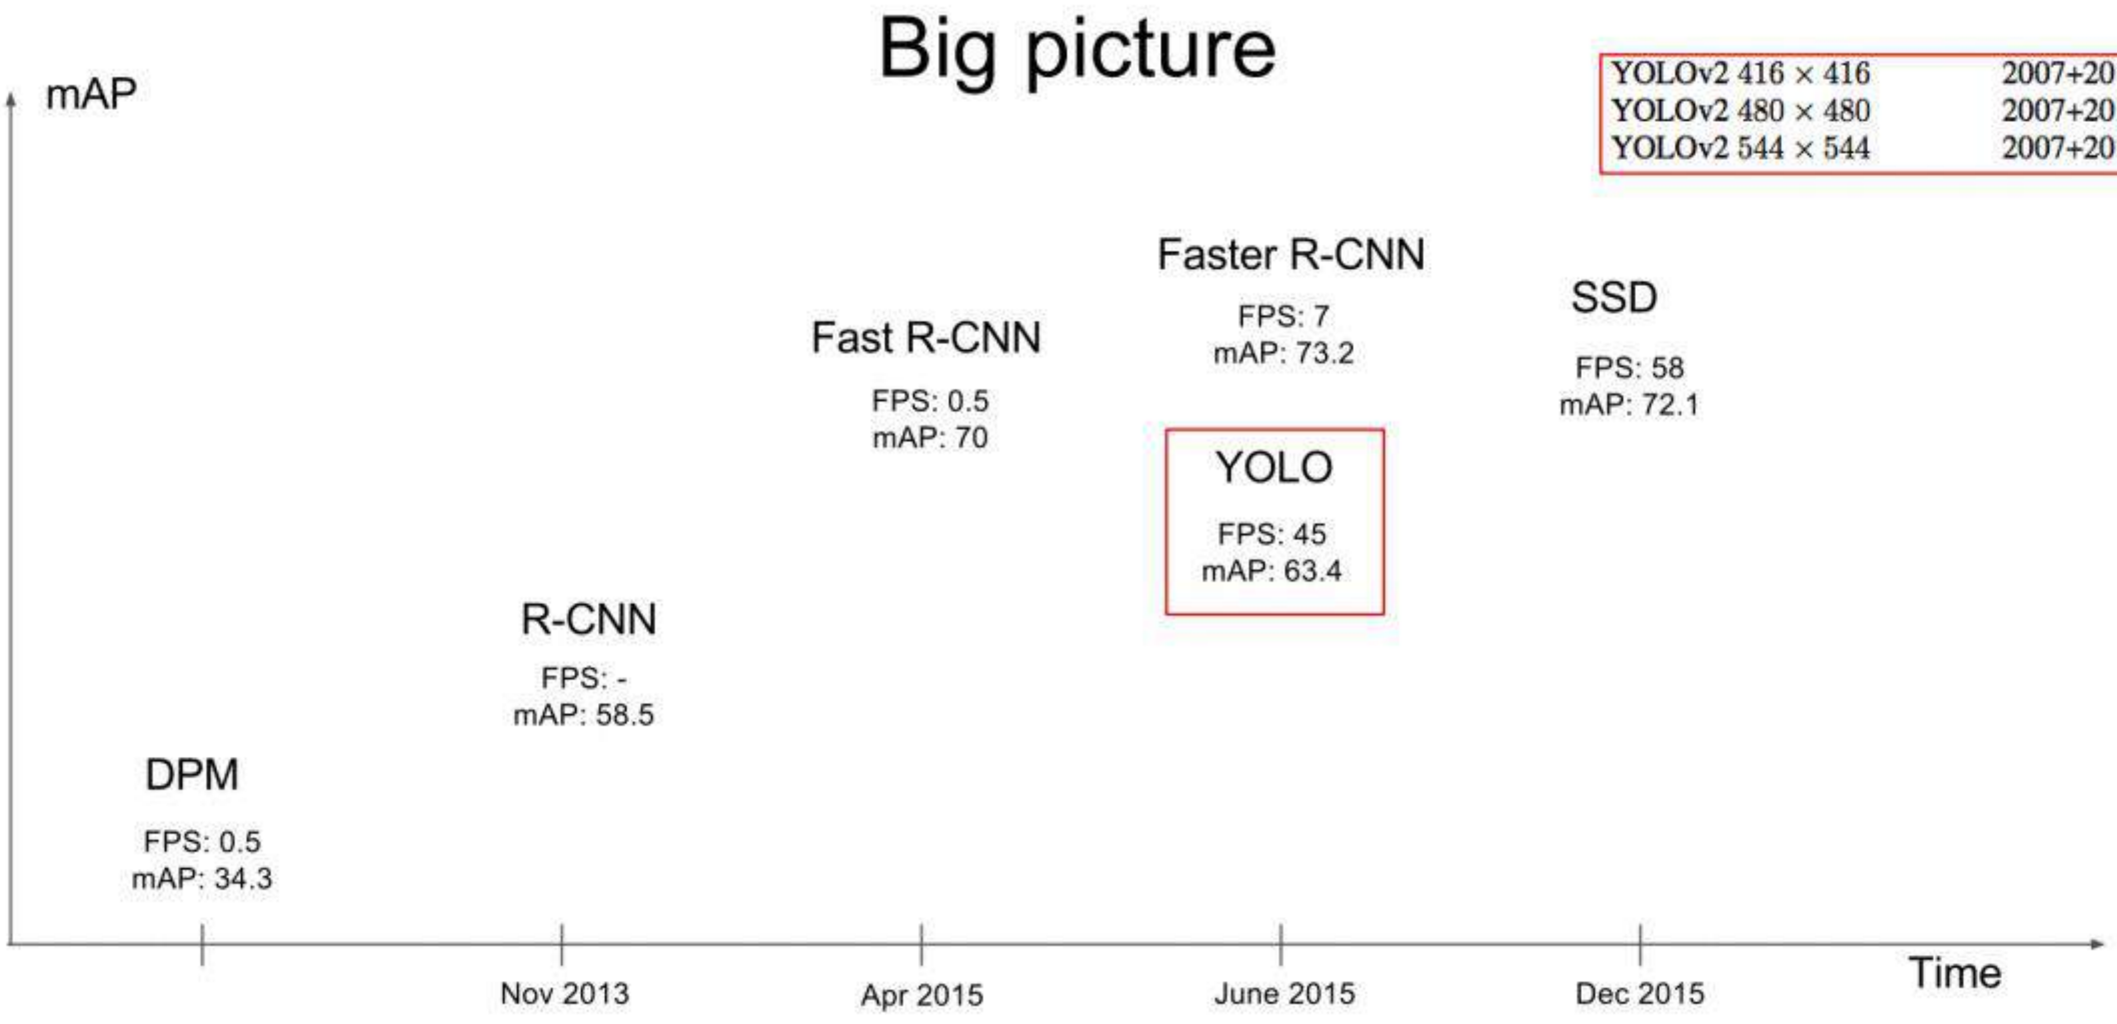

YOLOv2 416 × 416	2007+2012	76.8	67
YOLOv2 480 × 480	2007+2012	77.8	59
YOLOv2 544 × 544	2007+2012	**78.6**	40

Slide courtesy of DeepSystem.io "YOLO: You only look once Review"

Object Detection as Regression Problem

- **YOLO:** Single Regression Problem
  - Image → bounding box coordinate and class probability.
  - Extremely Fast
  - Global reasoning
  - Generalizable representation

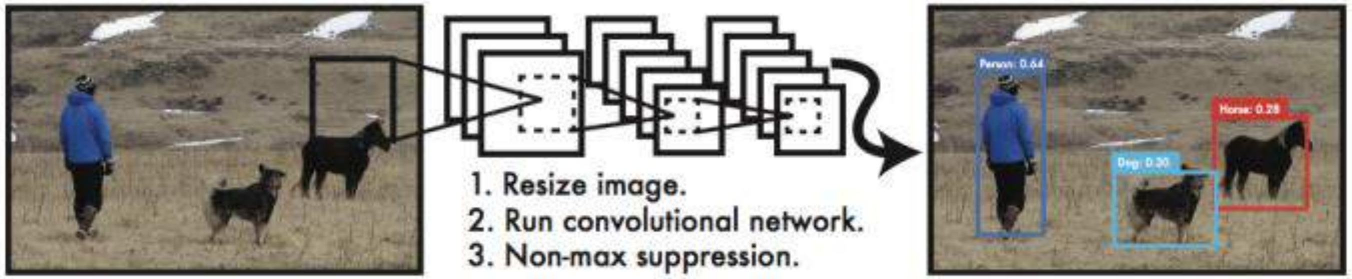

Figure 1: The YOLO Detection System. Processing images with YOLO is simple and straightforward. Our system (1) resizes the input image to 448×448, (2) runs a single convolutional network on the image, and (3) thresholds the resulting detections by the model's confidence.

Detection Procedure

We split the image into an S*S grid

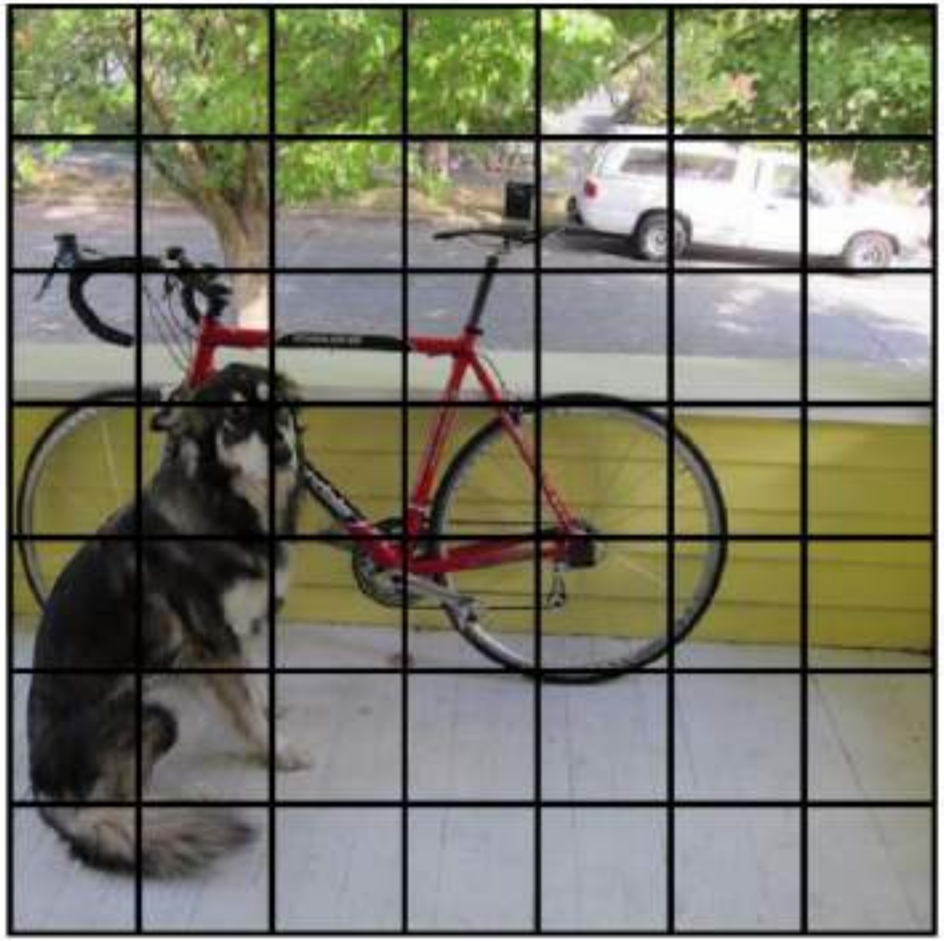

We split the image into an S*S grid

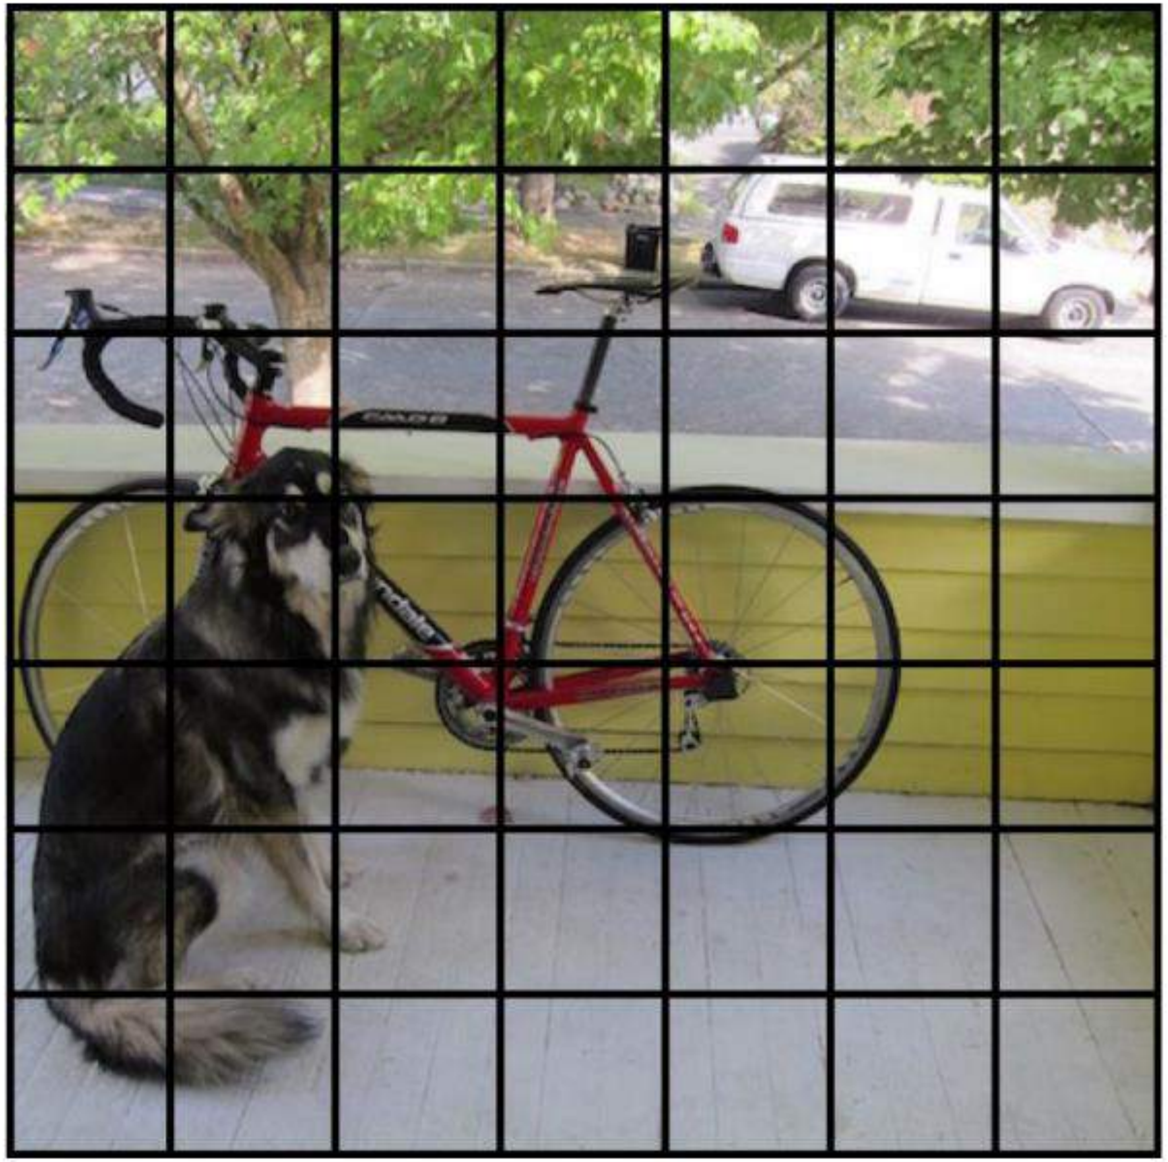

7*7 grid

Each cell predicts B boxes(x,y,w,h) and confidences of each box: P(Object)

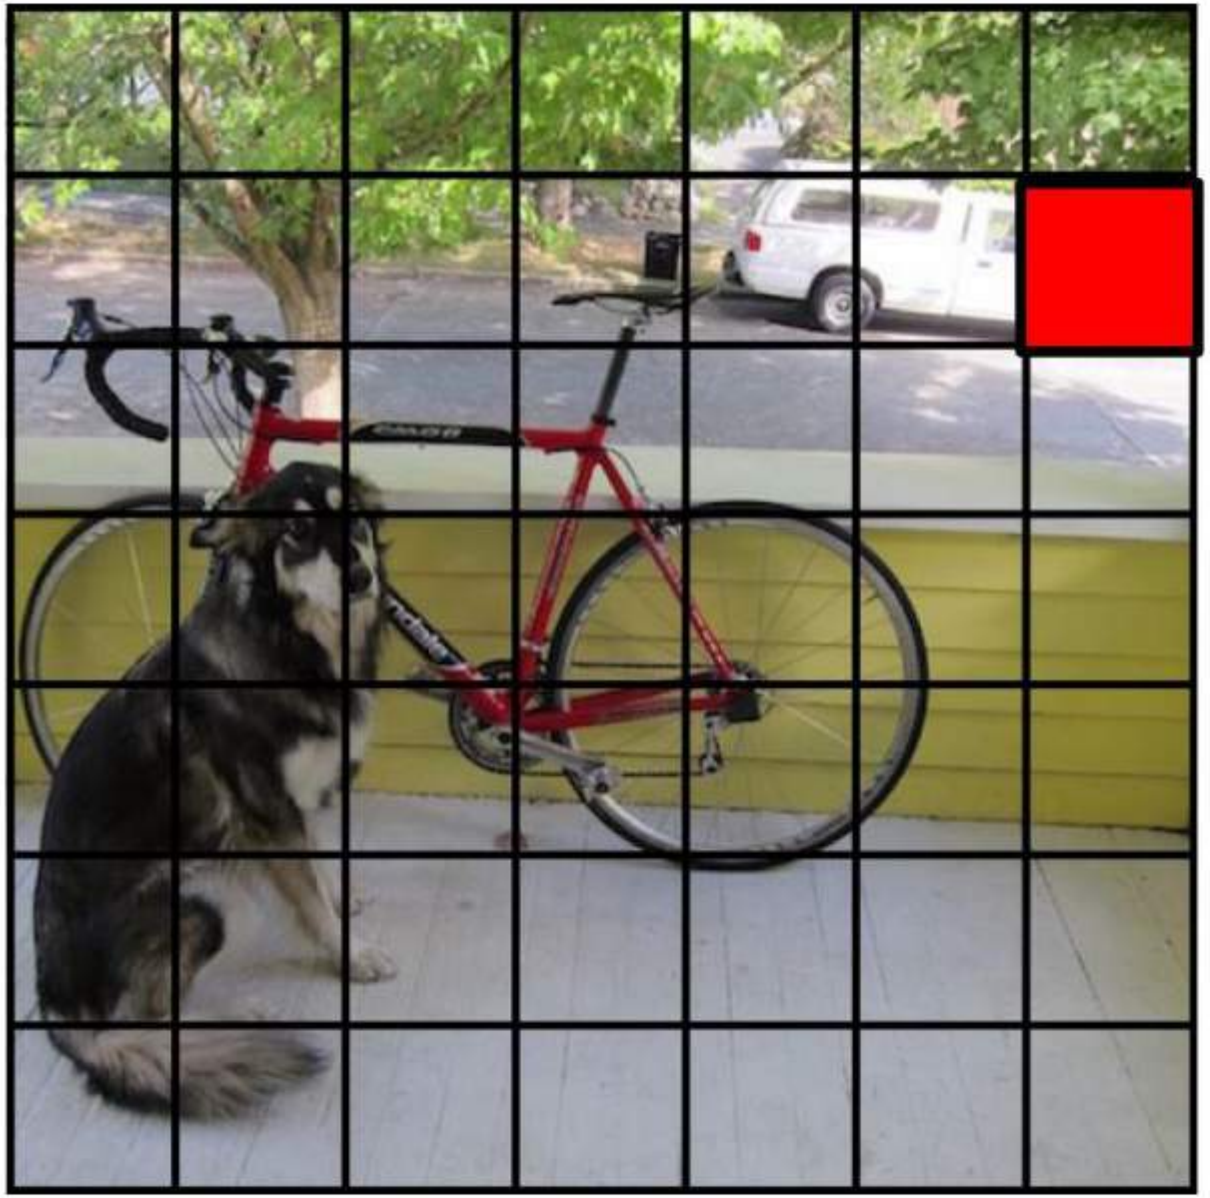

Each cell predicts B boxes(x,y,w,h) and confidences of each box: P(Object)

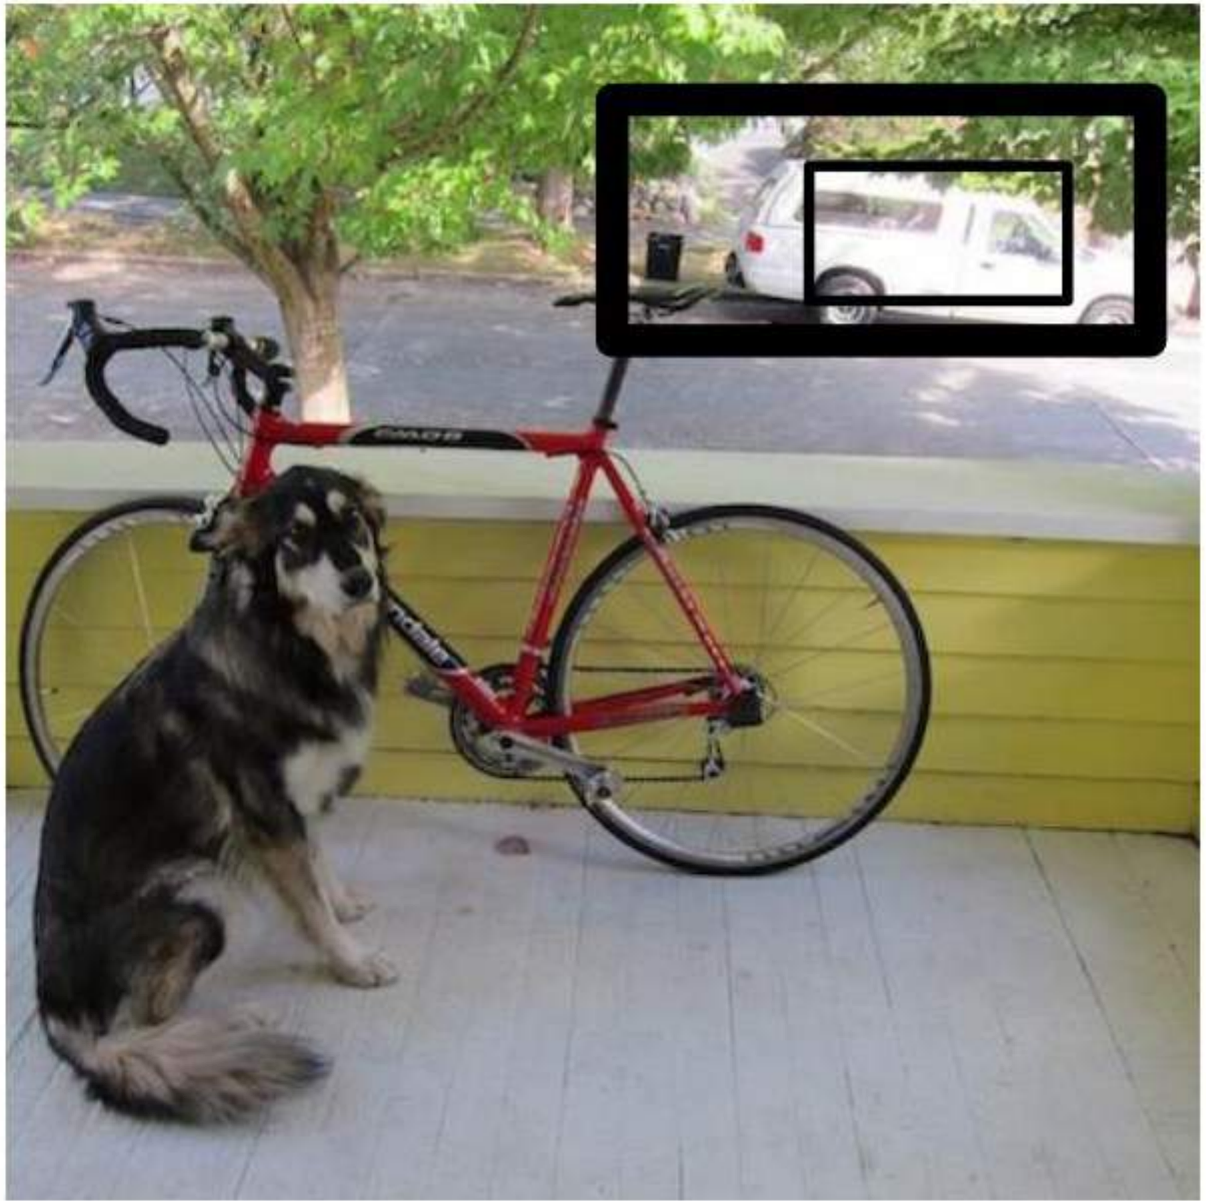

Each cell predicts B boxes(x,y,w,h) and confidences of each box: P(Object)

B = 2

each box predict:

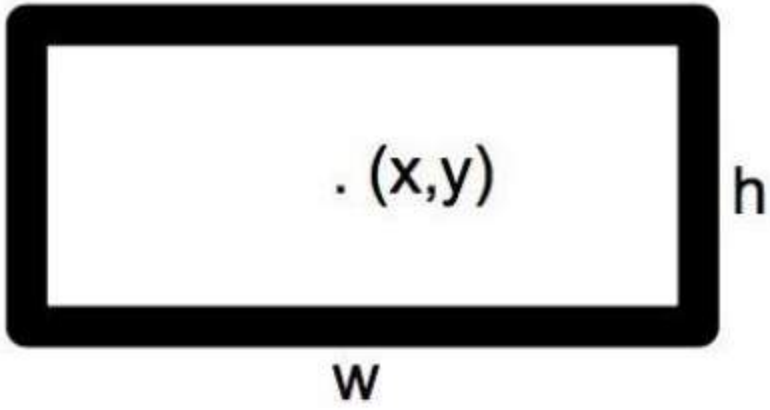

P(Object): probability that the box contains an object

Each cell predicts boxes and confidences: P(Object)

Each cell also predicts a class probability.

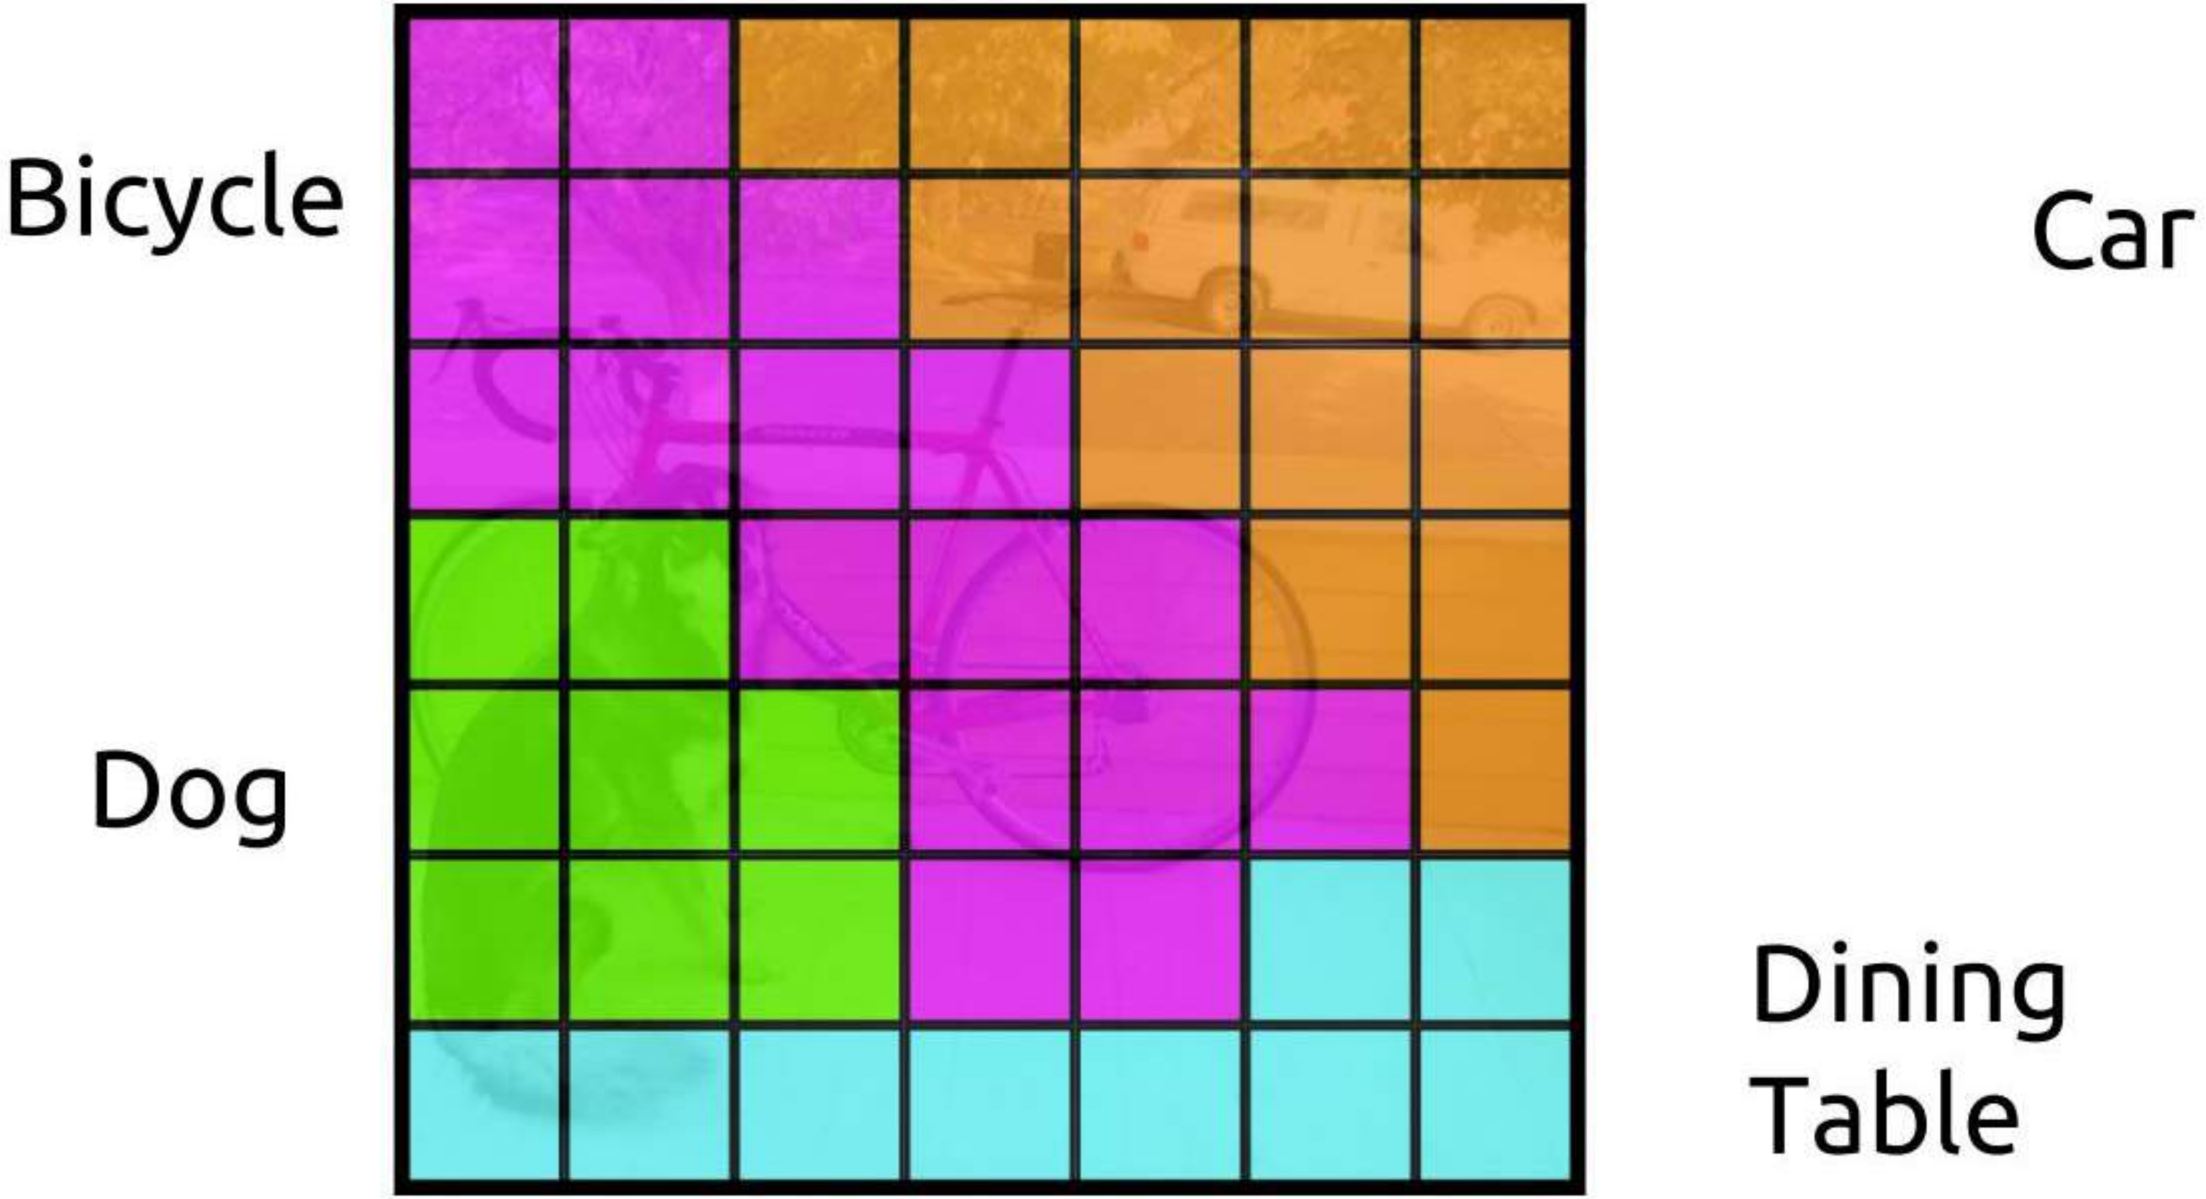

Conditioned on object: P(Car | Object)

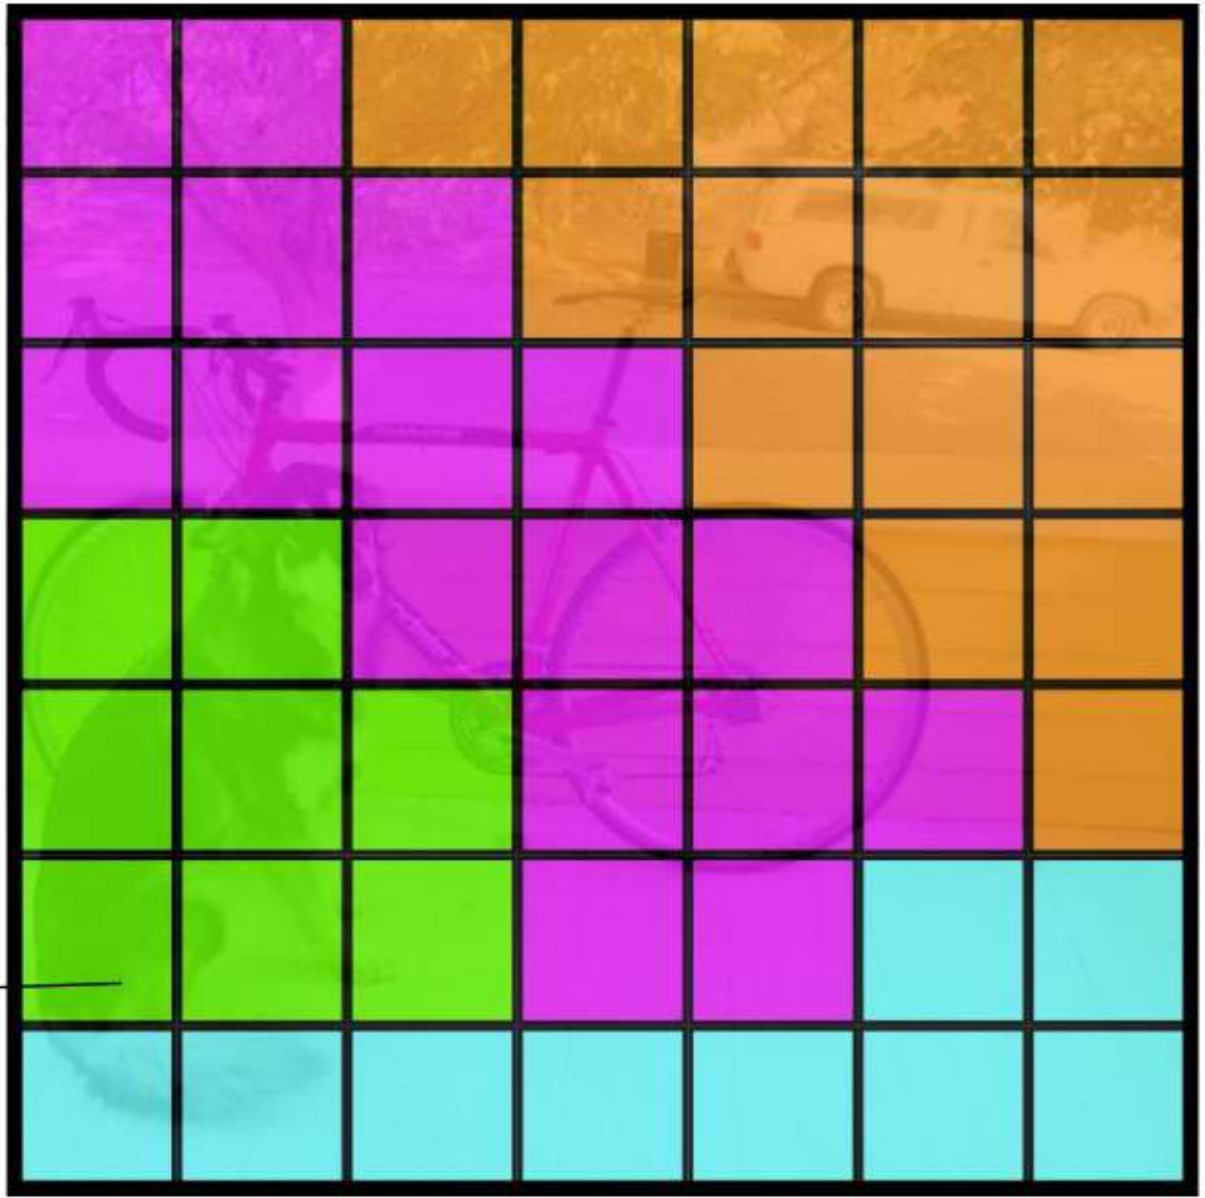

Then we combine the box and class predictions.

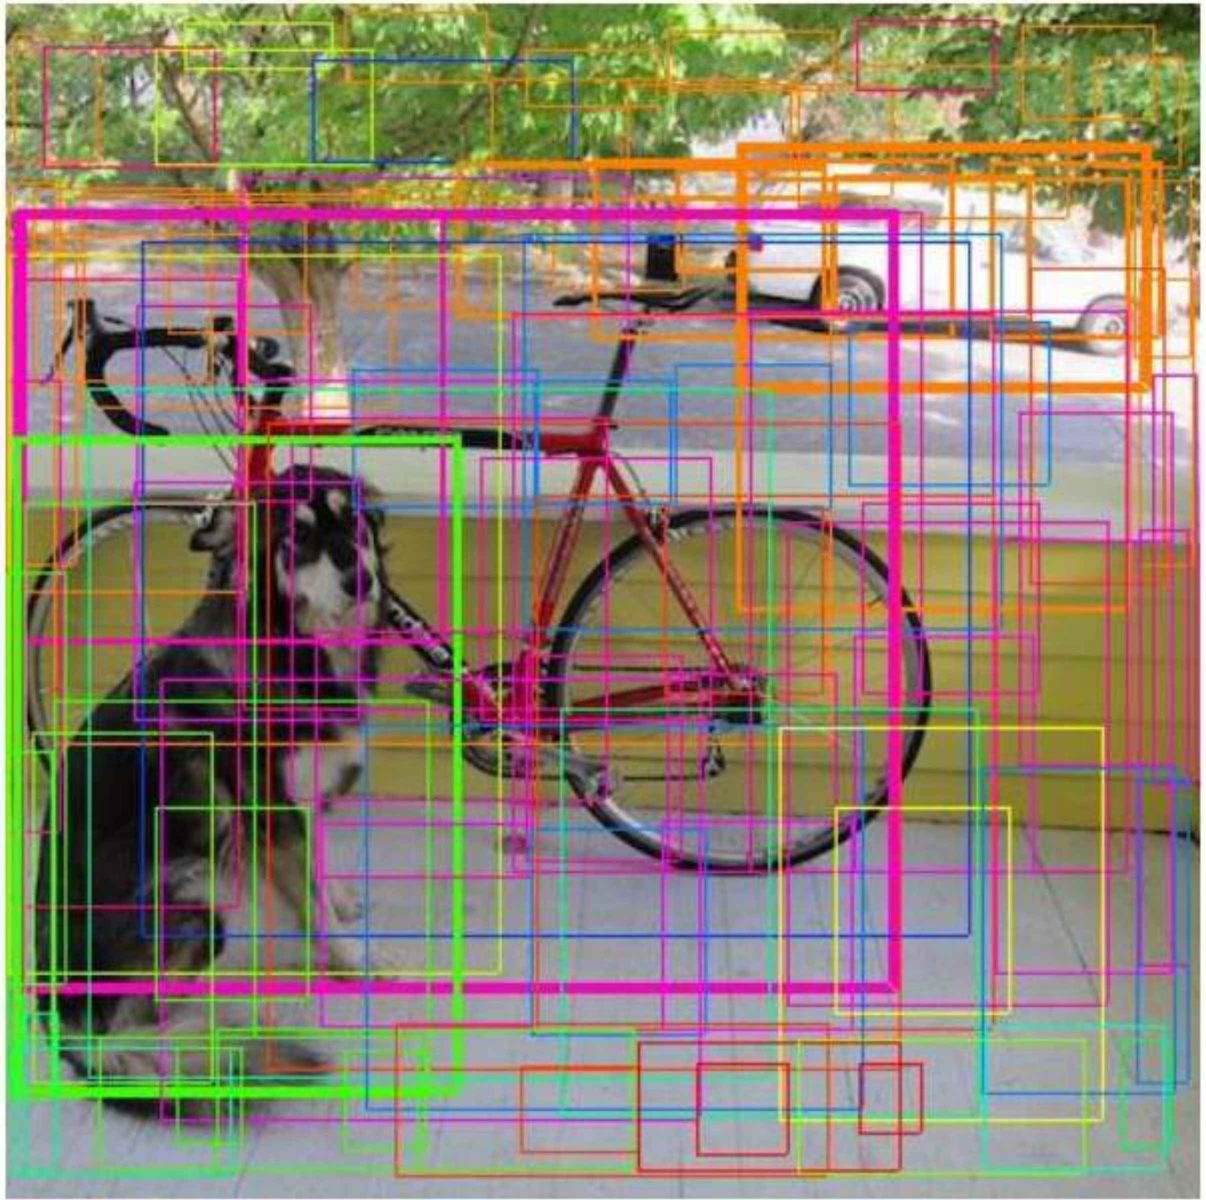

P(class|Object) * P(Object) =P(class)

Finally we do threshold detections and NMS

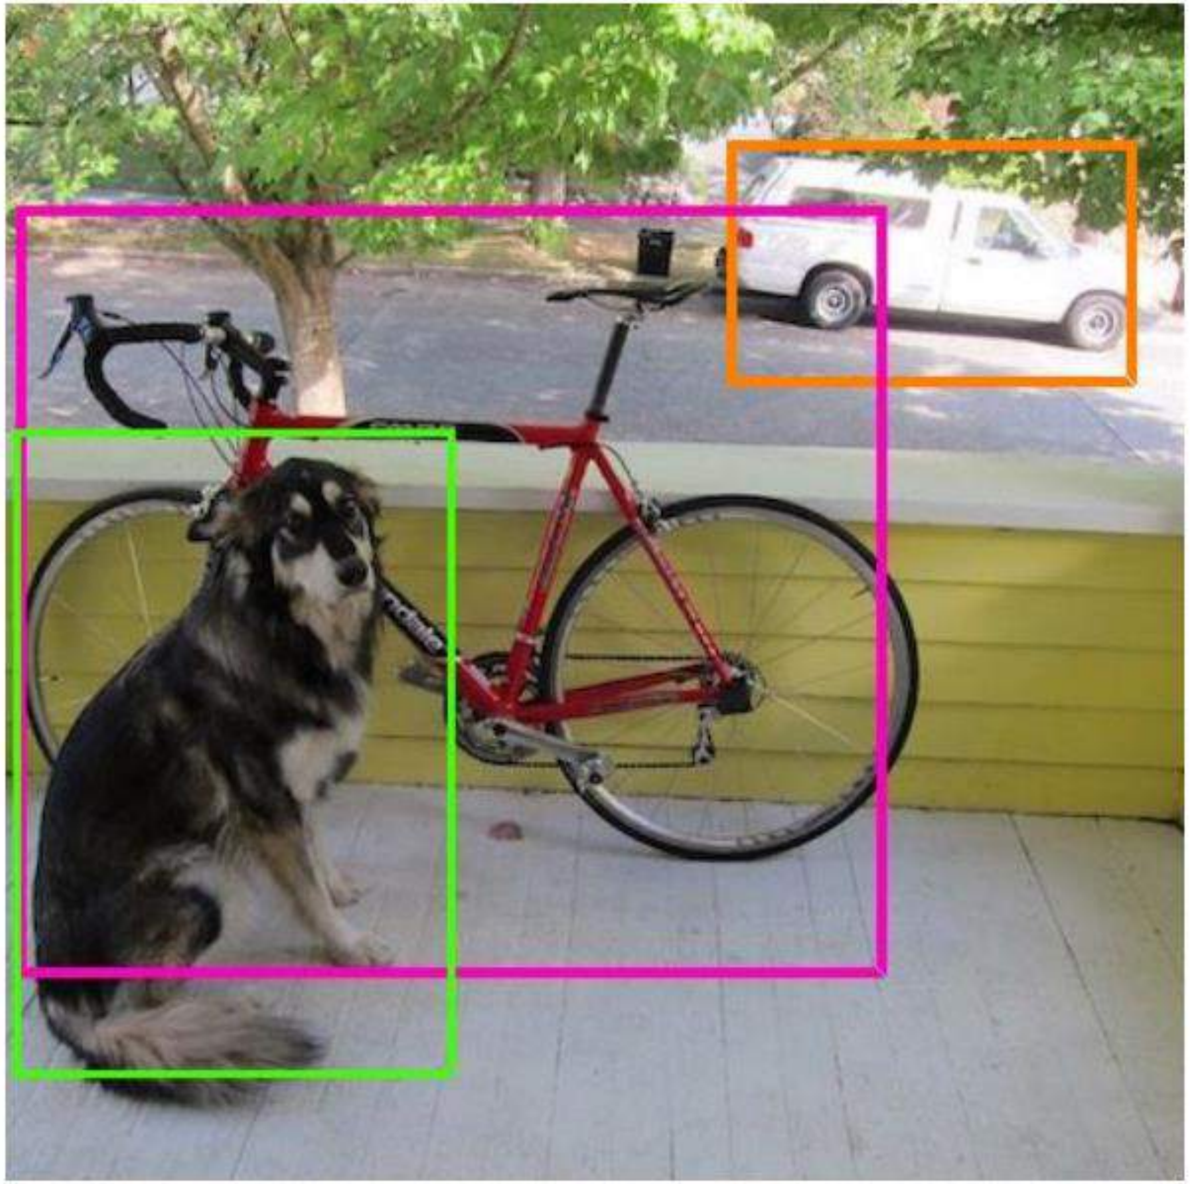

Information C

Taxonomy of Attacks

- White box testing
  - DD Attacks
  - Robust Physical Attack
  - etc
- Black box testing
  - Evolutionary Algorithms
  - Pixels Variations
  - etc

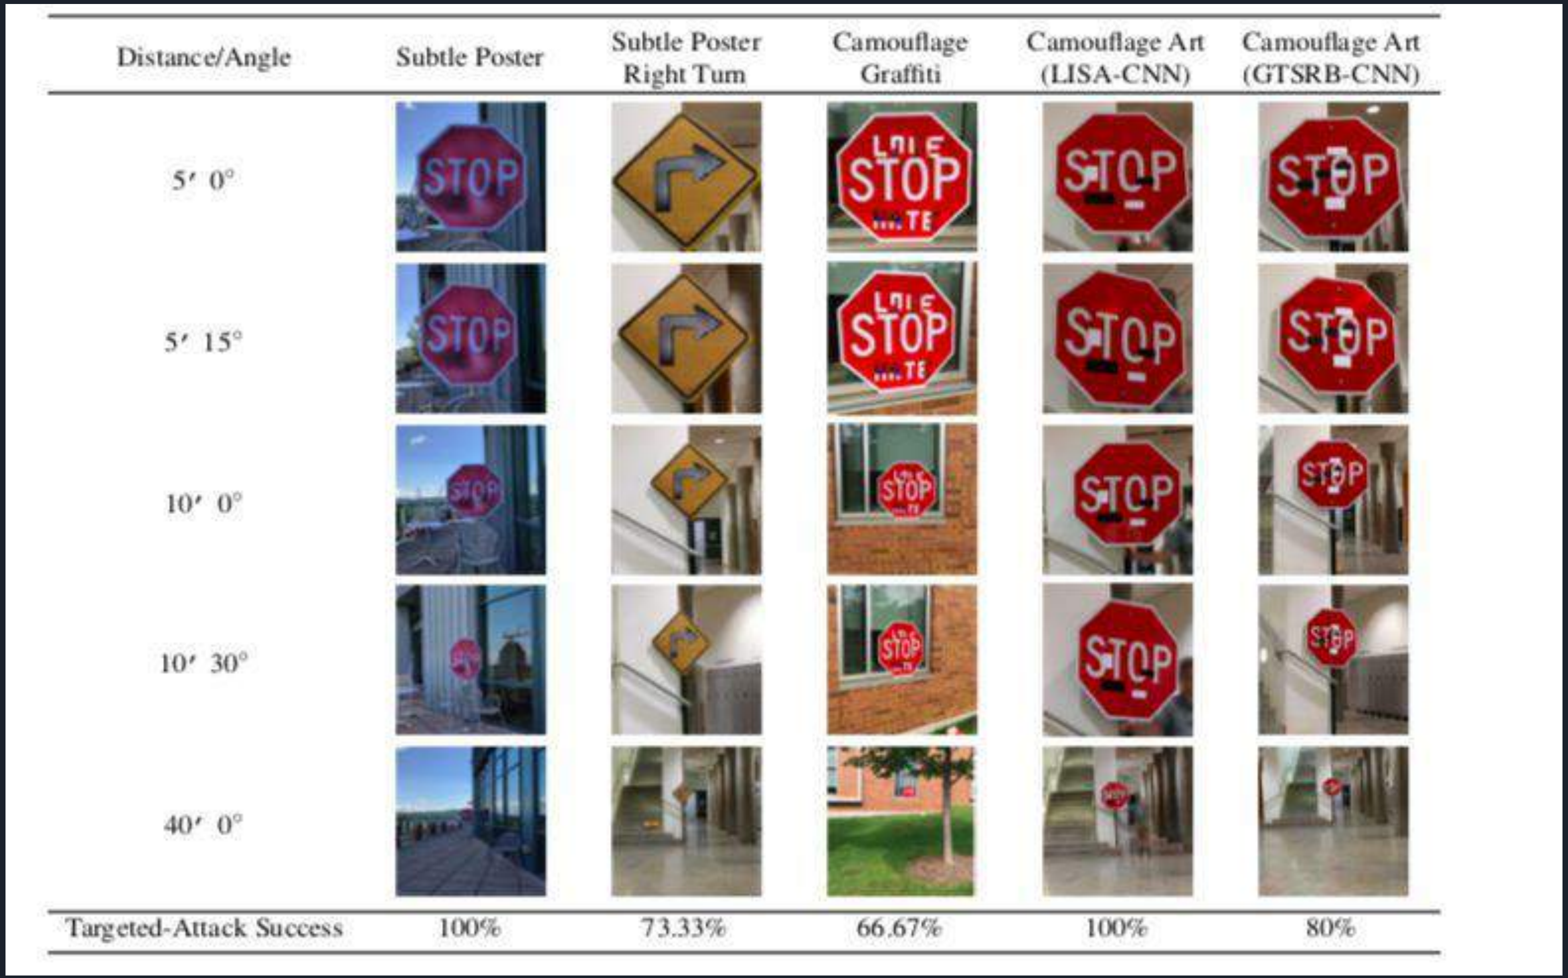

Source: Papernot et al.

Udacity

UDACITY
SELF-DRIVING CAR ENGINEER
CONTROLS
TRAINING MODE
AUTONOMOUS MODE
Select Track

Training Mode

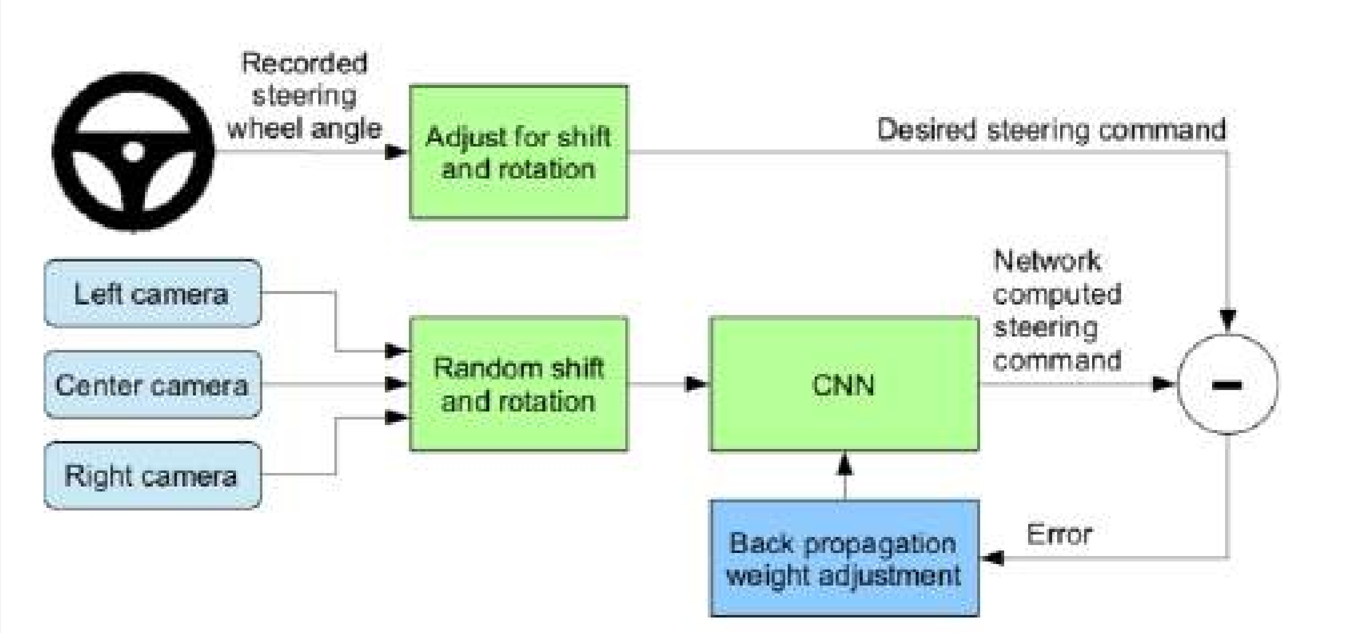

Autonomous Mode

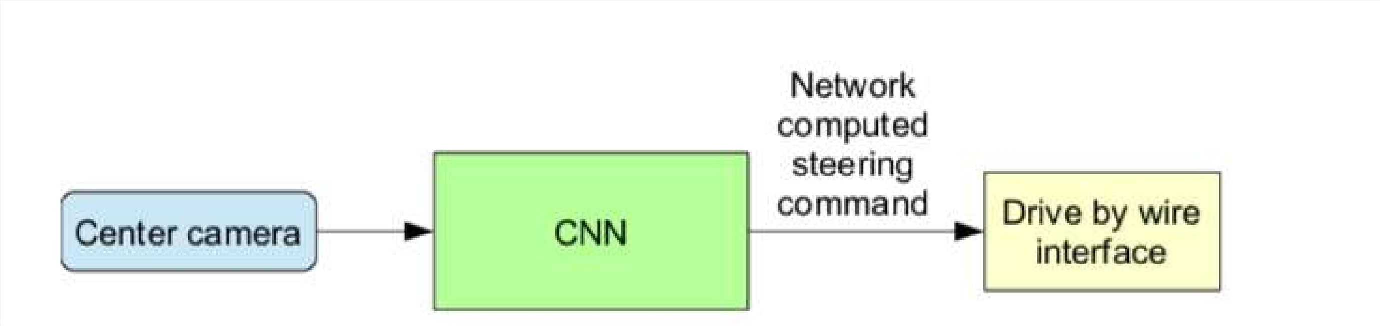

Adversarial Samples

In most of the cases, the adversary's goal is to produce a minimally altered version of the input x (image, video, text etc) such that it changes the output of the DL model, without being perceptible to the human eye.

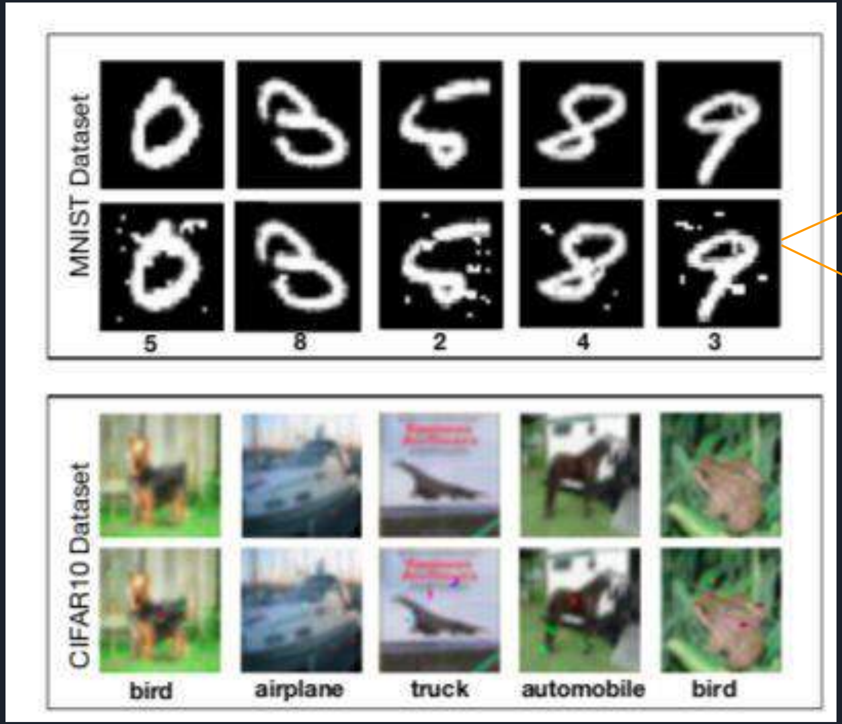

Adversarial Samples Perturbation

Source: Papernot et al.

$$\vec{x}^* = \vec{x} + \arg\min\{\vec{z} : \tilde{O}(\vec{x}+\vec{z}) \neq \tilde{O}(\vec{x})\} = \vec{x} + \delta_{\vec{x}}$$

$$\tilde{O}(\vec{x}^*) \neq \tilde{O}(\vec{x})$$

Adversarial Goals and Capabilities

Goals:

1. Confidence reduction
2. Misclassification
3. Targeted misclassification
4. Source/target misclassification

Black-Box Attacks

- In black-box attacks, the adversary doesn't have any knowledge about the model, except for the the labels

- The goal is to produce a minimal perturbation to input X, sufficient to determine the DNN to misclassify it, but imperceptible enough for the human eye

Regular Pixel Attack

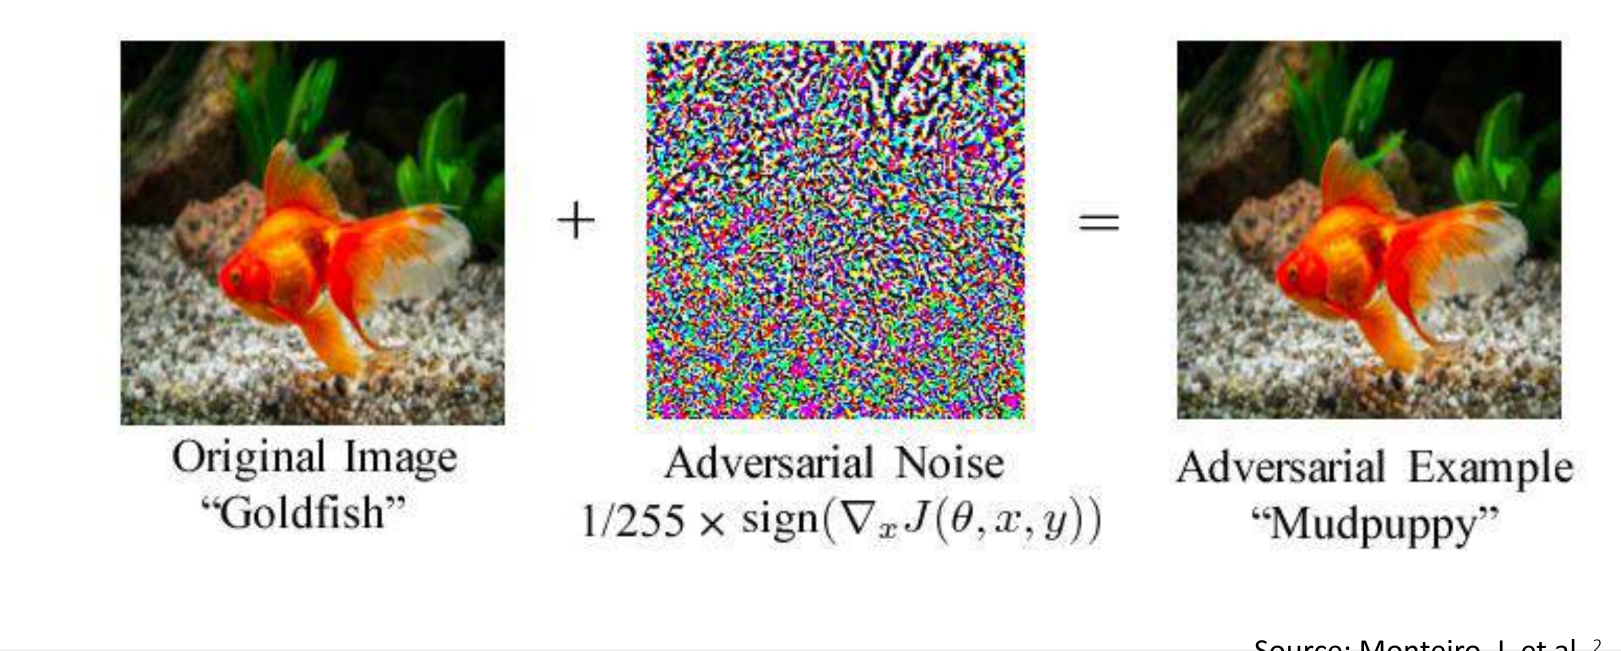

Source: Monteiro, J. et al. [2]

Two variations of the attack

- Attack on the steering angle
- Attack on the distance between two images

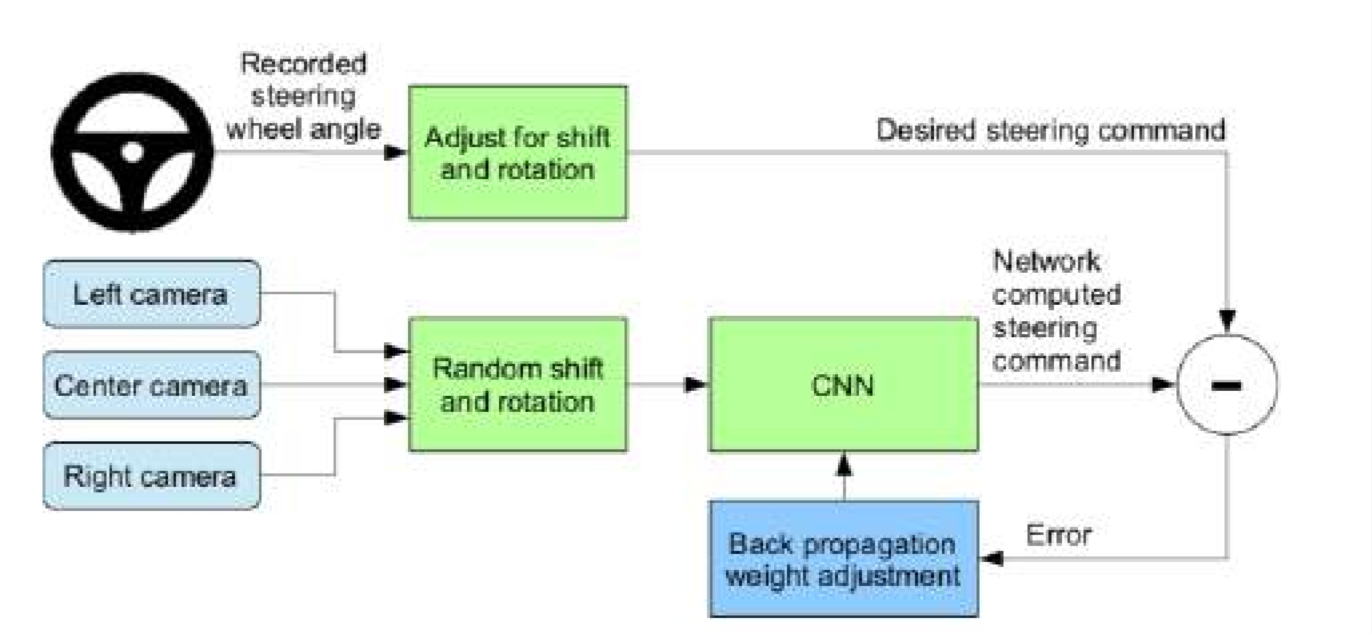

Attack Scheme

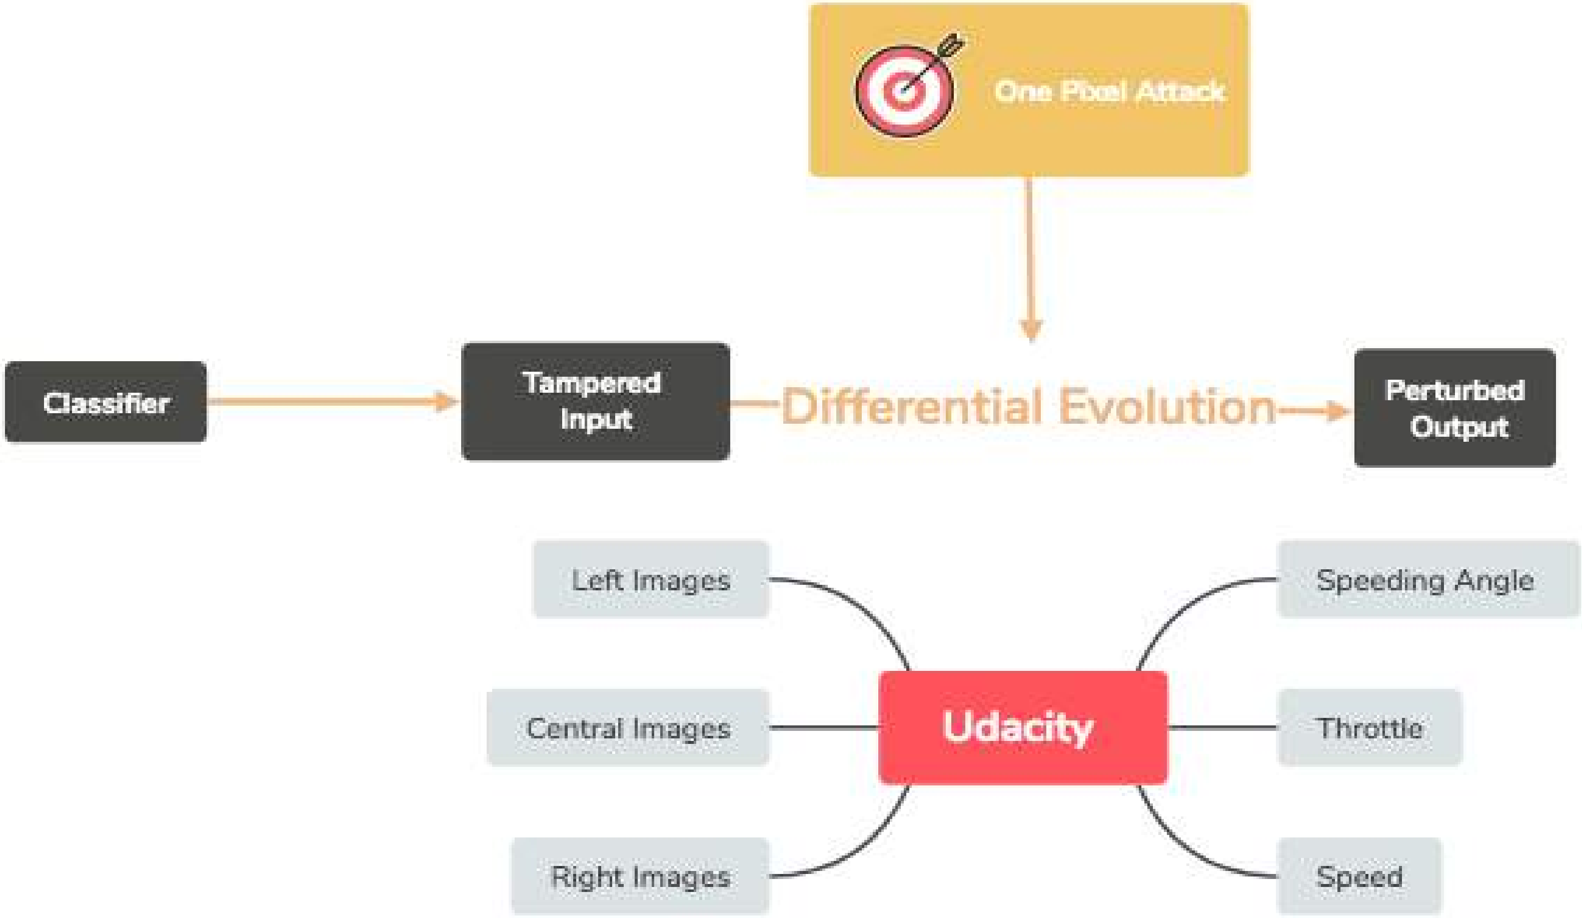

Attack on the steering angle prediction

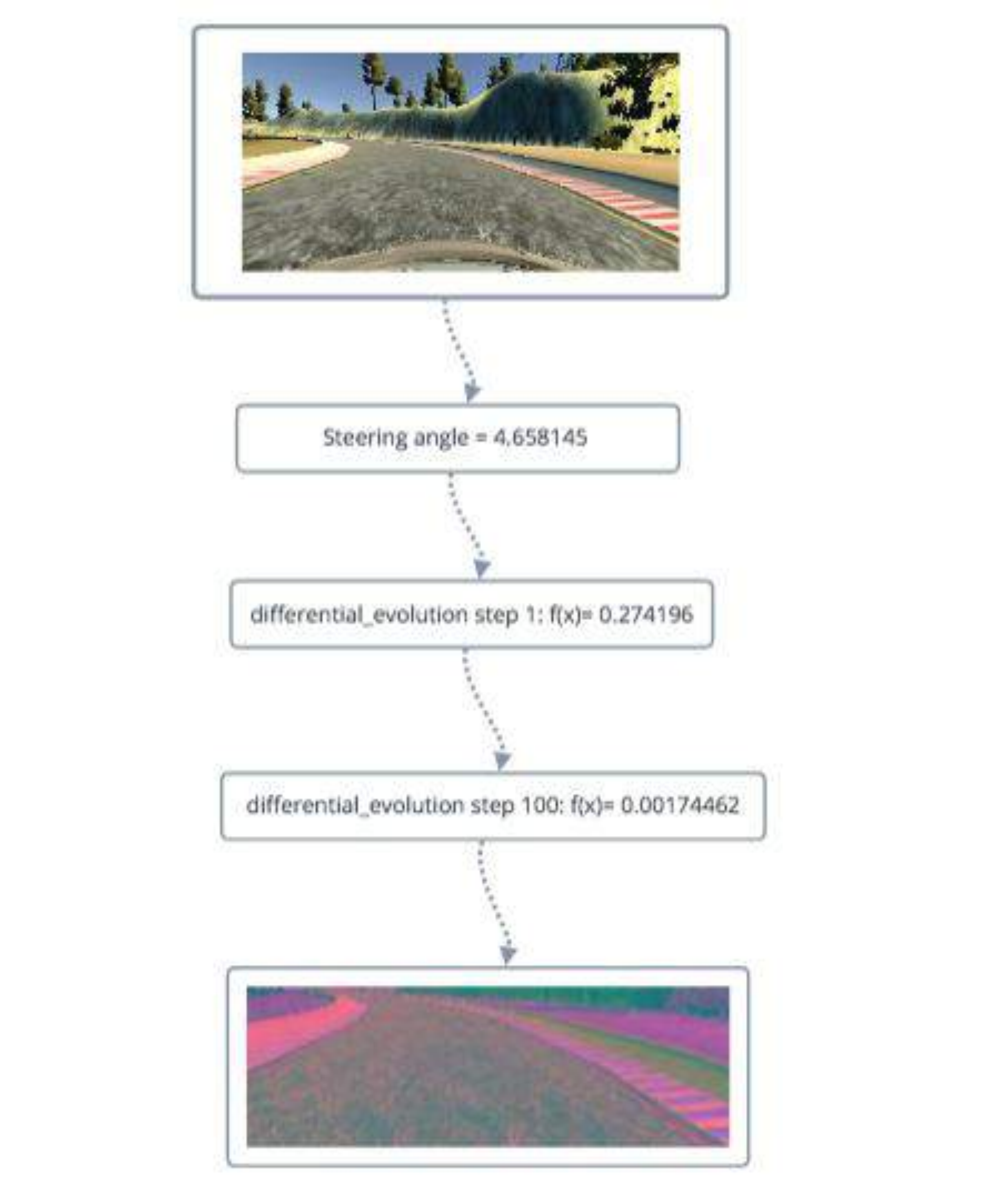

Attack on the distance between two images

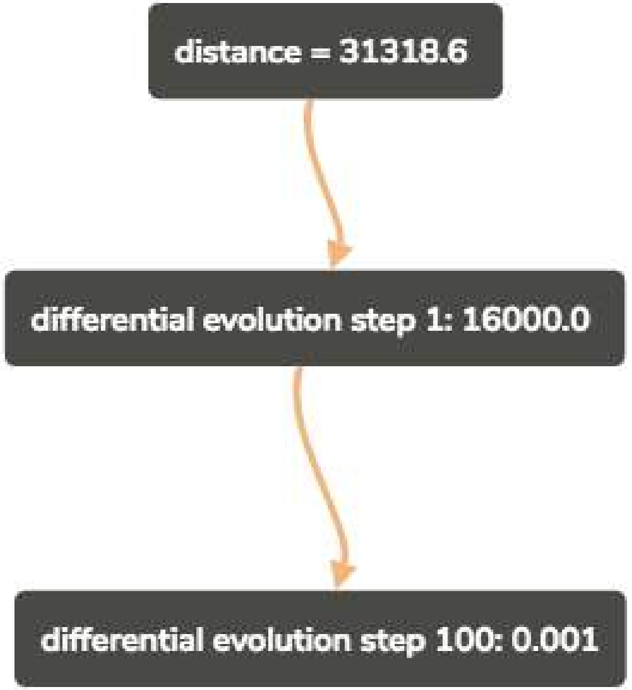

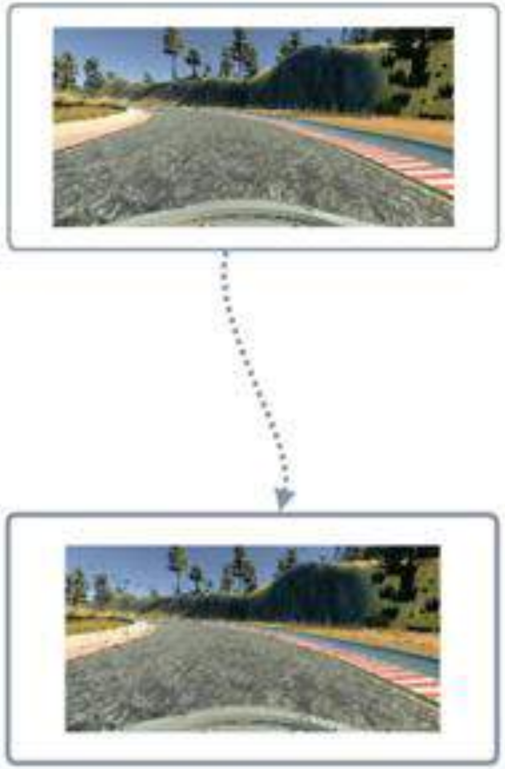

Normalization

- Account for the change in the number of pixels while minimizing the change in pixels

$$normalization = 0.5\ x\ \frac{\#\ different\ pixels}{total\ \#\ pixels} + 0.5\ x\ \frac{steering\ angle}{\max \#\ steering\ angles}$$

Future Work

- **Normalize** variables **Steering angle, Speeding angel, Speed, and Orientation of Angel**
- **Improve** the attack by tampering with the classifier and the decision making model
- **Tampering** with the mapping and routing algorithms

Acknowledgement

- The hard work of our autonomous cars research space
  - Daria Zahaleanu
  - Fathuur Fahmi Said
  - Mohamed Munaib Afzal
  - Fahad Ahmad
  - Mohamed Osama Khan

CENTER FOR
CYBER SECURITY
جامعة نيويورك ابوظبي
NYU ABU DHABI

Hoda A.Alkhzaimi
+971504330432
Hoda.Alkhzaimi@nyu.edu
www.nyuad.nyu.edu/ccs-ad