

Module 3: Automotive CPS Data driven modeling

Principles of Modeling for Cyber-Physical Systems

Instructor: Madhur Behl

Slides credits:
- Urs Muller

Everything that moves will go autonomous



Cars



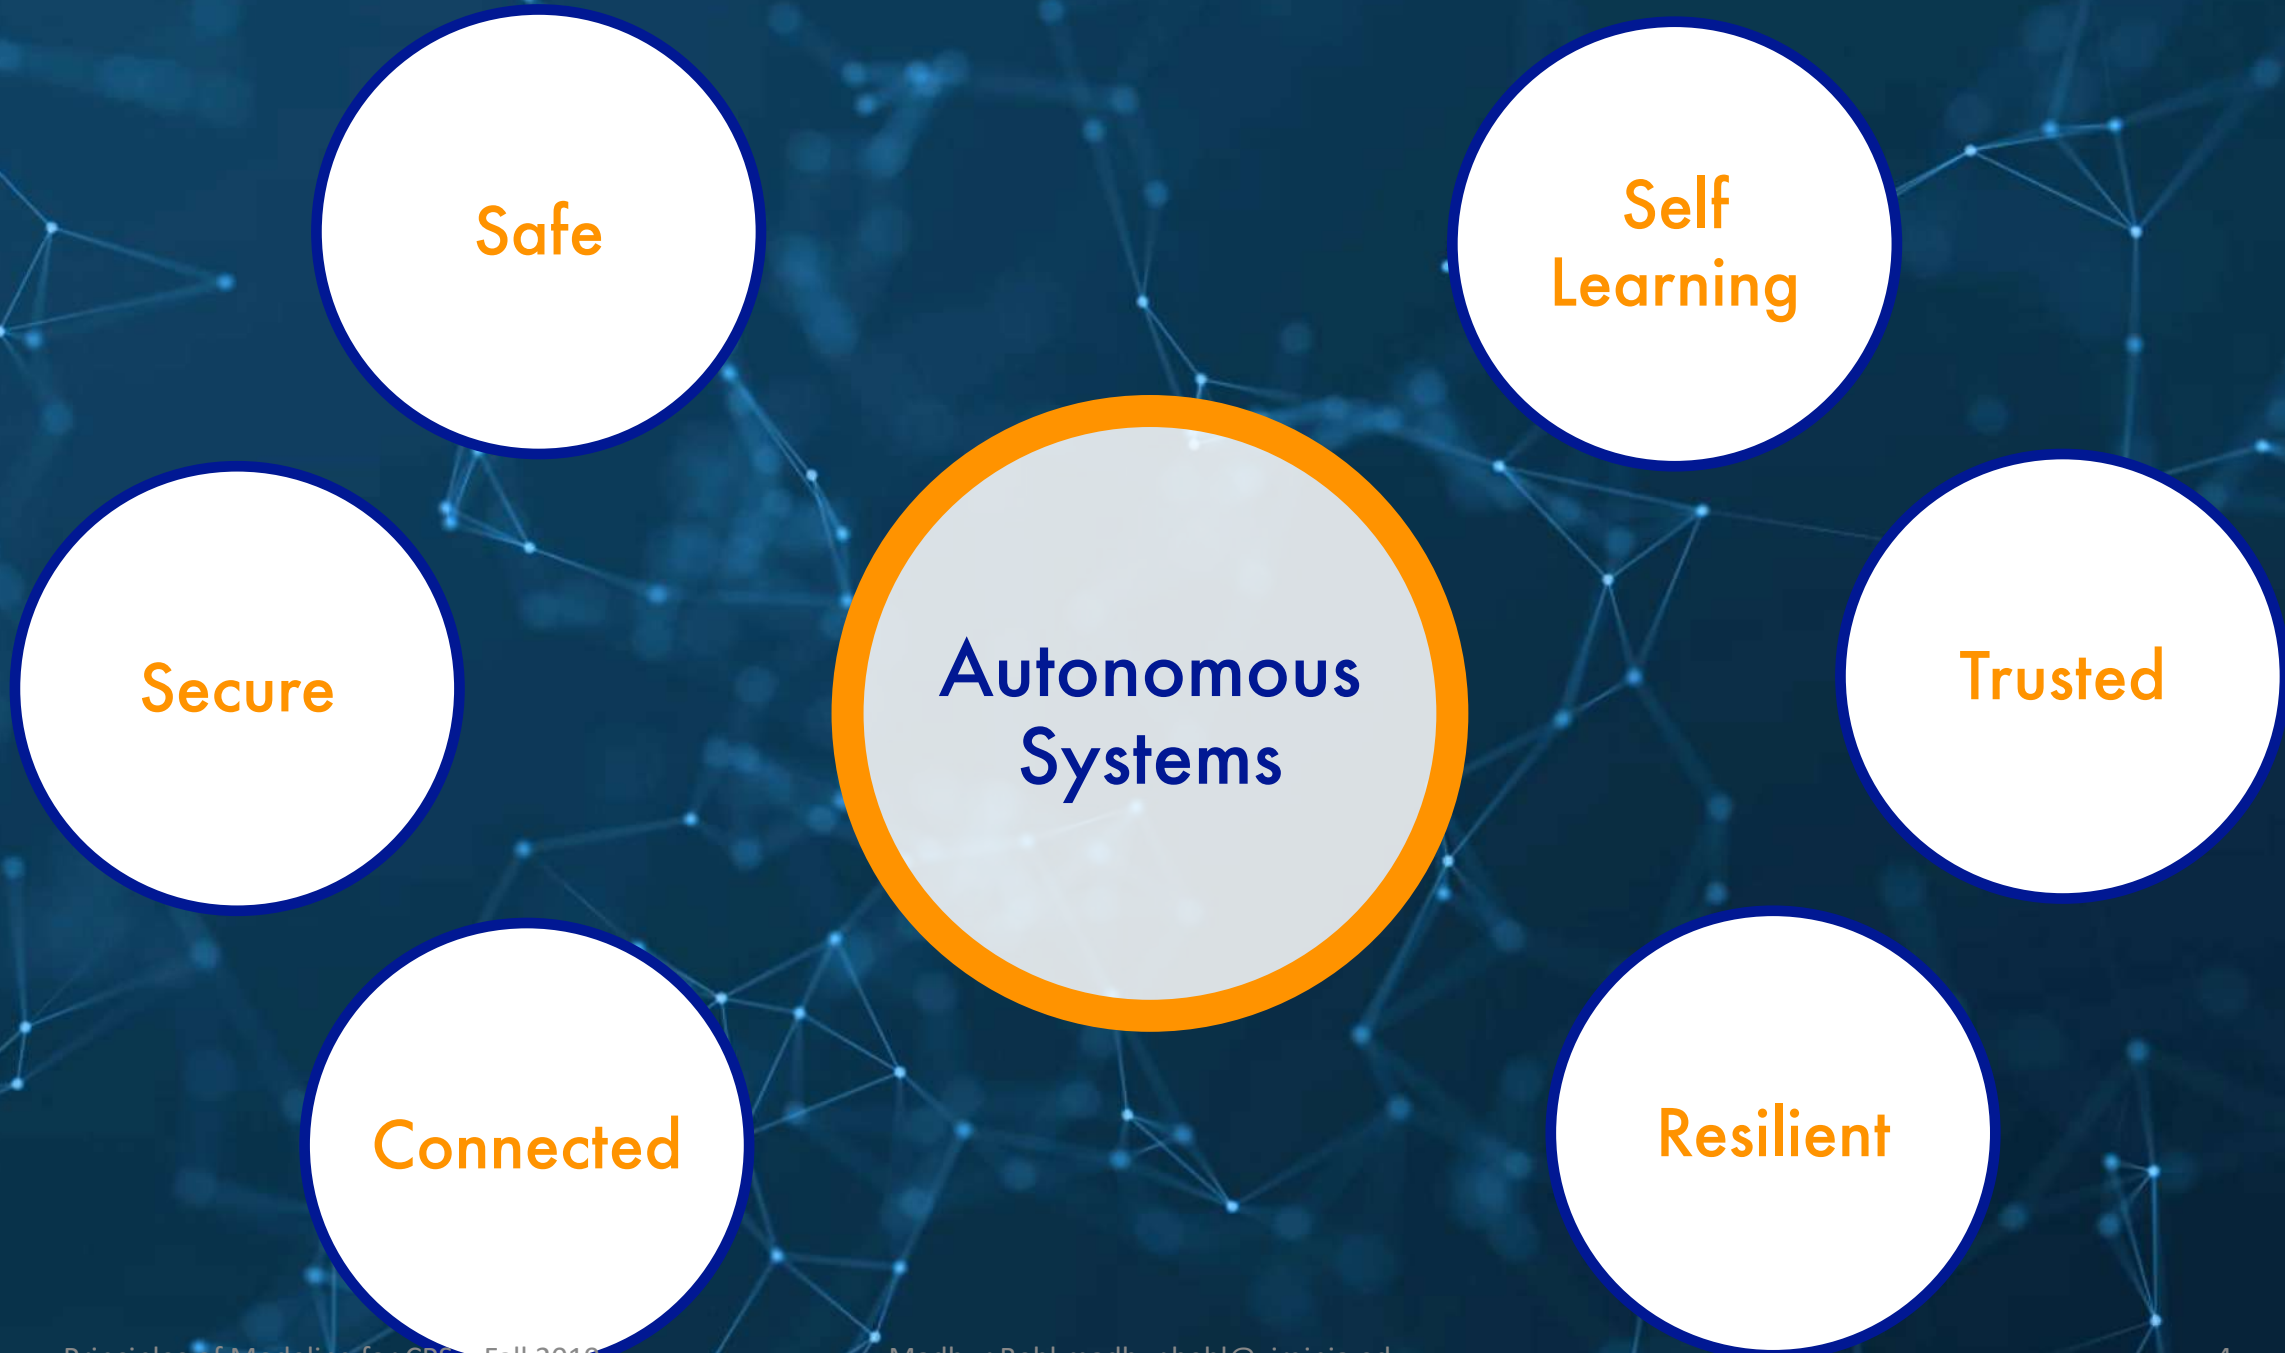
Trucks



Carts



Drones





THE FUTURE OF TRANSPORTATION STACK

<p>SERVICES</p> <p>ROUTE PLANNING</p> <p>SPATIAL</p>	<p>PARKING</p>	<p>CAR HAILING + POOLING</p>	<p>OTHER: AFTERMARKET, REPAIR, RENTAL</p>	<p>SPECIALTY VEHICLES</p> <p>2-WHEELERS</p>	
<p>SAFETY & SECURITY</p> <p>PHYSICAL CAR & DRIVER SAFETY + ACCIDENT DETECTION</p>	<p>EMOTION, FATIGUE & ALCOHOL DETECTION + DISTRACTION AVOIDANCE</p>	<p>CYBERSECURITY</p>	<p>INTRUSION, TRACKING & RECOVERY</p>	<p>PUBLIC TRANSPORT</p>	
<p>IN-CAR INTELLIGENCE + ASSISTANCE</p> <p>VEHICLE DIAGNOSTICS & PREDICTIVE MAINTENANCE + SENSOR-BASED VEHICLE SAFETY</p>	<p>PASSENGER-FOCUSED SENSORS (INCLUDING USAGE-BASED INSURANCE)</p>	<p>INFOTAINMENT + DISPLAY</p>	<p>PERSONAL / VOICE ASSISTANCE</p>	<p>NAVIGATION ASSISTANCE + PEDESTRIAN ANALYSIS & COMMUNICATIONS</p>	<p>TRUCKS / FREIGHT</p>
<p>AUTONOMY</p> <p>AUTOMATION SYSTEM</p>	<p>MAPPING, SIMULATION, & IMAGE RECOGNITION / ANNOTATION</p>	<p>AUTONOMOUS VEHICLE MAKER + TOOLS</p>	<p>FLIGHT</p>		
<p>INFRASTRUCTURE + CONNECTED CAR</p> <p>SENSOR NETWORKING INFRASTRUCTURE (V2V, V2X) - LPWA, CELLULAR, WIFI</p>	<p>CONNECTED CAR - DATA, PLATFORM, SOFTWARE</p>	<p>FLEET + TRAFFIC MANAGEMENT</p>	<p>OTA CAR SOFTWARE UPDATE + SMART PHONE ENABLED TELEMATICS</p>	<p>OTHER: HYPERLOOP, PERSONAL MOBILITY</p>	
<p>INTELLIGENT MANUFACTURING</p> <p>NEW / ADVANCED MATERIALS</p>	<p>RAPID PROTOTYPING - 3D PRINTING, MODULARIZATION, OPEN SOURCE</p>	<p>ADVANCED / AUTOMATED ASSEMBLY LINE</p>	<p>MATERIAL CHARACTERIZATION & TESTING</p>		
<p>ONBOARD SENSORS</p> <p>LOCATION - GIS, PRECISION POSITIONING, PATH PLANNING</p>	<p>VISION / CAMERA</p>	<p>LIDAR</p>	<p>RADAR</p>		



THE 6 LEVELS OF AUTONOMOUS DRIVING





0

No Automation

Zero autonomy; the driver performs all driving tasks.

1

Driver Assistance

Vehicle is controlled by the driver, but some driving assist features may be included in the vehicle design.

2

Partial Automation

Vehicle has combined automated functions, like acceleration and steering, but the driver must remain engaged with the driving task and monitor the environment at all times.

3

Conditional Automation

Driver is a necessity, but is not required to monitor the environment. The driver must be ready to take control of the vehicle at all times with notice.

4

High Automation

The vehicle is capable of performing all driving functions under certain conditions. The driver may have the option to control the vehicle.

5

Full Automation

The vehicle is capable of performing all driving functions under all conditions. The driver may have the option to control the vehicle.



1.35 million deaths worldwide due to vehicle crashes

94% of crashes involve human choice or error in the US.

3 million

Americans age 40 and older are blind or have low vision

79%

of seniors age 65 and older living in car-dependent communities

42 hours

wasted in traffic each year per person

Localization and Mapping

Where am I ?

Scene Understanding


Where/who/what/why of everyone/everything else ?

Trajectory Planning and Control

Where should I go next ?
How do I steer and accelerate ?

Human Interaction

How do I convey my intent to the passenger and everyone else ?



Localization and Mapping

Where am I ?

Scene Understanding


Where/who/what/why of
everyone/everything else ?

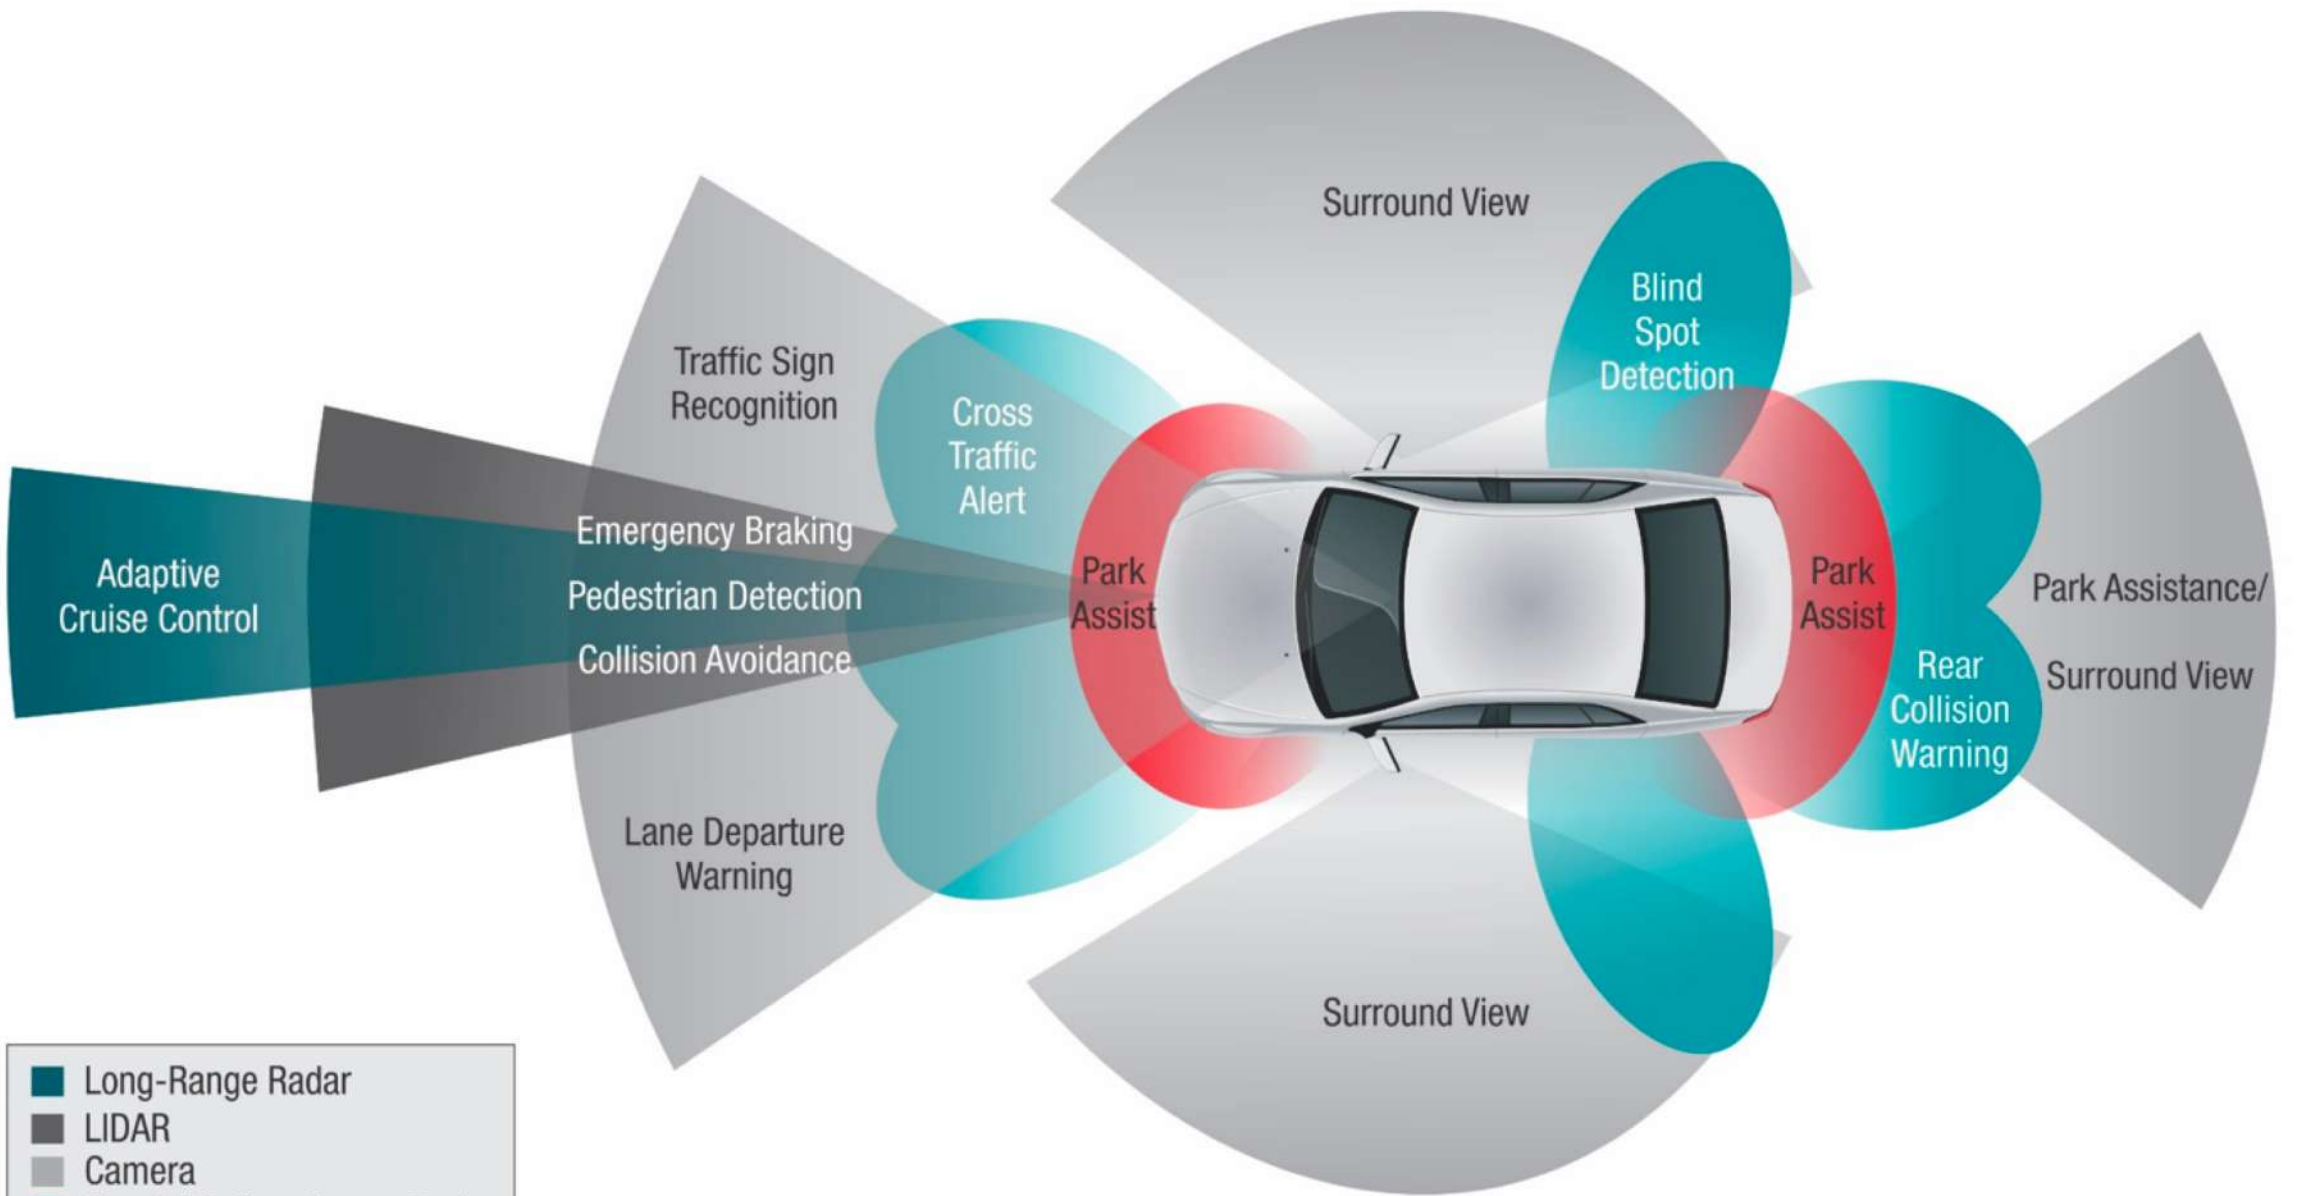
Trajectory Planning and Control

Where should I go next ?
How do I steer and accelerate ?


Human Interaction

How do I convey my intent to the
passenger and everyone else ?





- Long-Range Radar
- LIDAR
- Camera
- Short-/Medium Range Radar
- Ultrasound




Camera



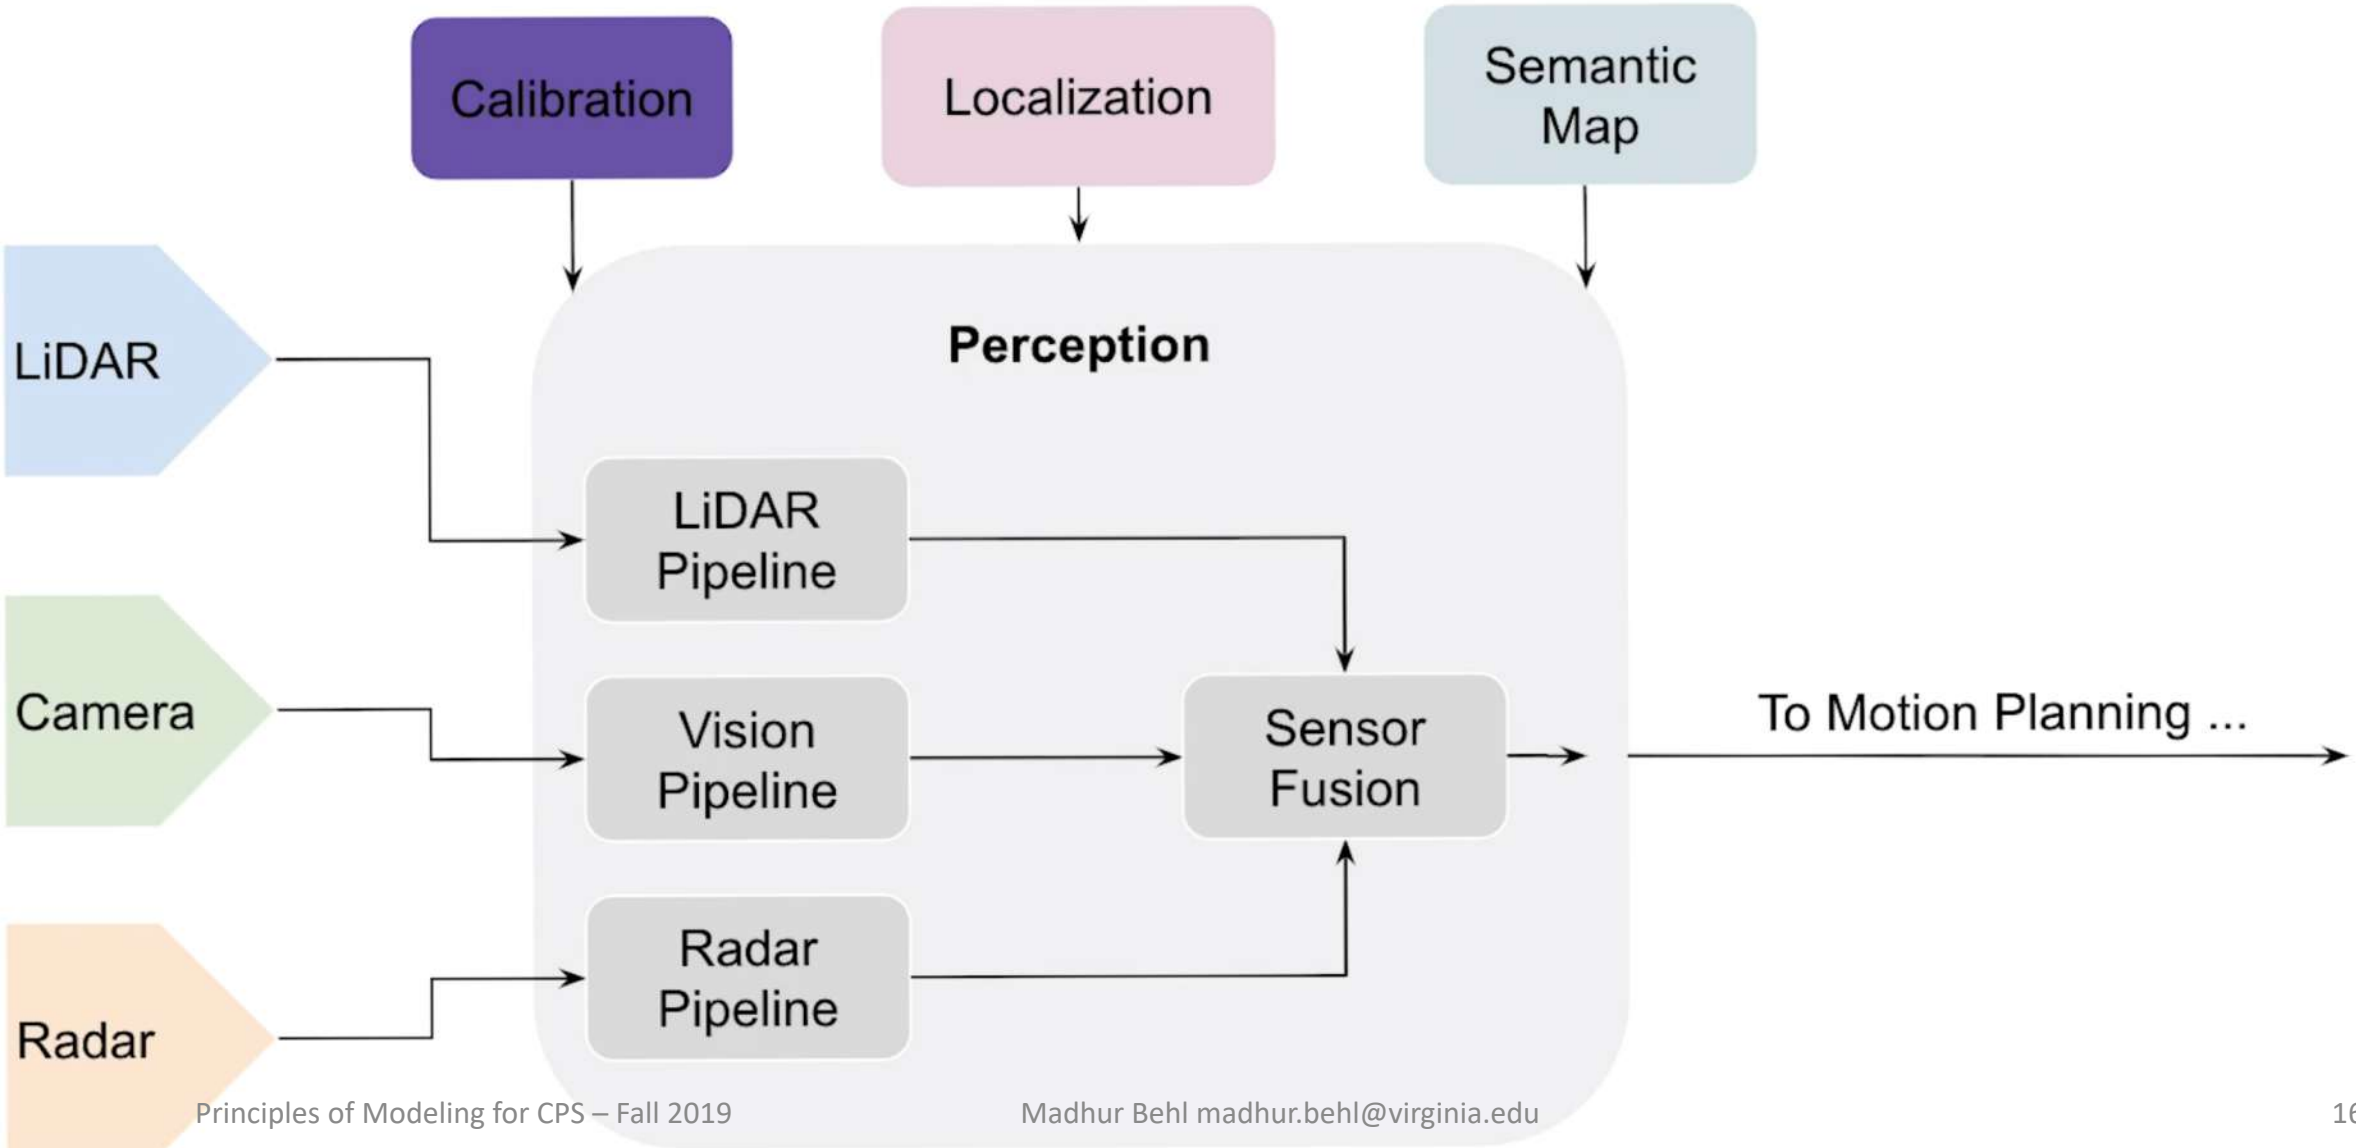
Radar



LIDAR



Perception in AV Stack




SLAM: Simultaneous Localization and Mapping

What works: SIFT and optical flow

x2



TRACKING - KFs: 17 , MPs: 3497 , Tracked: 315 + 51



Object Detection



- Past approaches: cascades classifiers (Haar-like features)
- Where deep learning can help:
recognition, classification, detection



ection - NVDriveNet detection

TECH EVENTS



suv-truck car suv-truck suv-truck suv-truck suv-truck

Front:



car car suv-truck suv-truck

Rear :






Full Driving Scene Segmentation



- Building
- Pole
- Road Marking
- Road
- Pavement
- Tree
- Sign Symbol
- Fence
- Vehicle
- Pedestrian
- Bike



Localization and Mapping

Where am I ?

Scene Understanding


Where/who/what/why of
everyone/everything else ?

Trajectory Planning and Control

Where should I go next ?
How do I steer and accelerate ?

Human Interaction

How do I convey my intent to the
passenger and everyone else ?







Detailed three-dimensional maps that highlight information such as road profiles, curbs and sidewalks, lane markers, crosswalks, traffic lights, stop signs, and other road features.

Where am I?



Scan constantly for objects around the vehicle—
pedestrians, cyclists, vehicles, road work, obstructions—
and continuously read traffic controls, from traffic light
color and railroad crossing gates to temporary stop signs.

What's around me?




Predict the movements of everything around you based on their speed and trajectory

What will happen next?





Determine the exact trajectory, speed, lane, and steering maneuvers needed to progress along the route safely


What should I do?



HD Maps: Localization








San Francisco

Localization: Scan Matching





Localization and Mapping

Where am I ?

Scene Understanding


Where/who/what/why of everyone/everything else ?

Trajectory Planning and Control

Where should I go next ?
How do I steer and accelerate ?

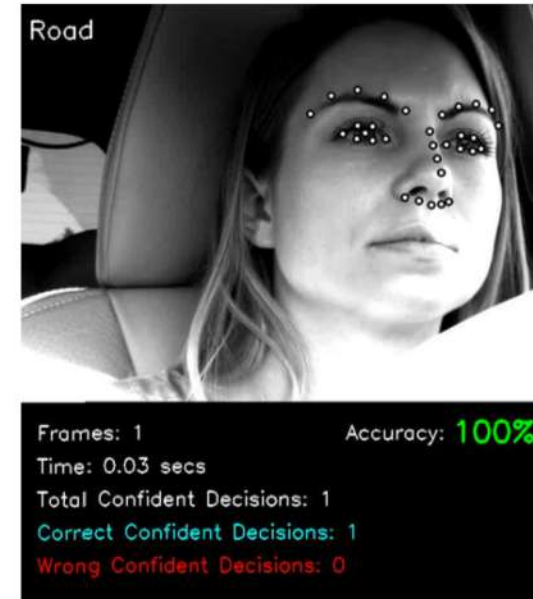
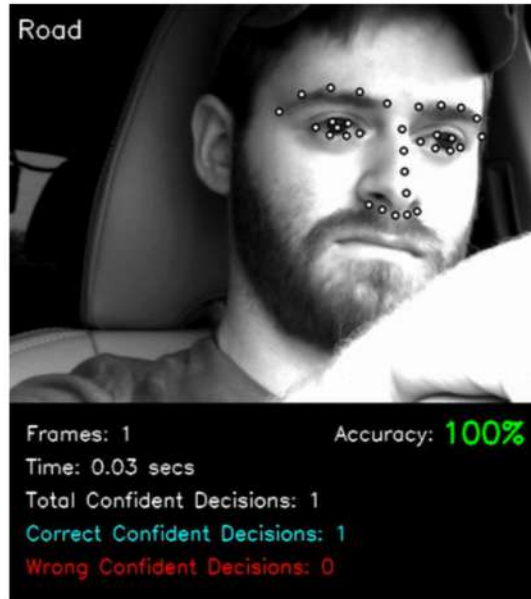
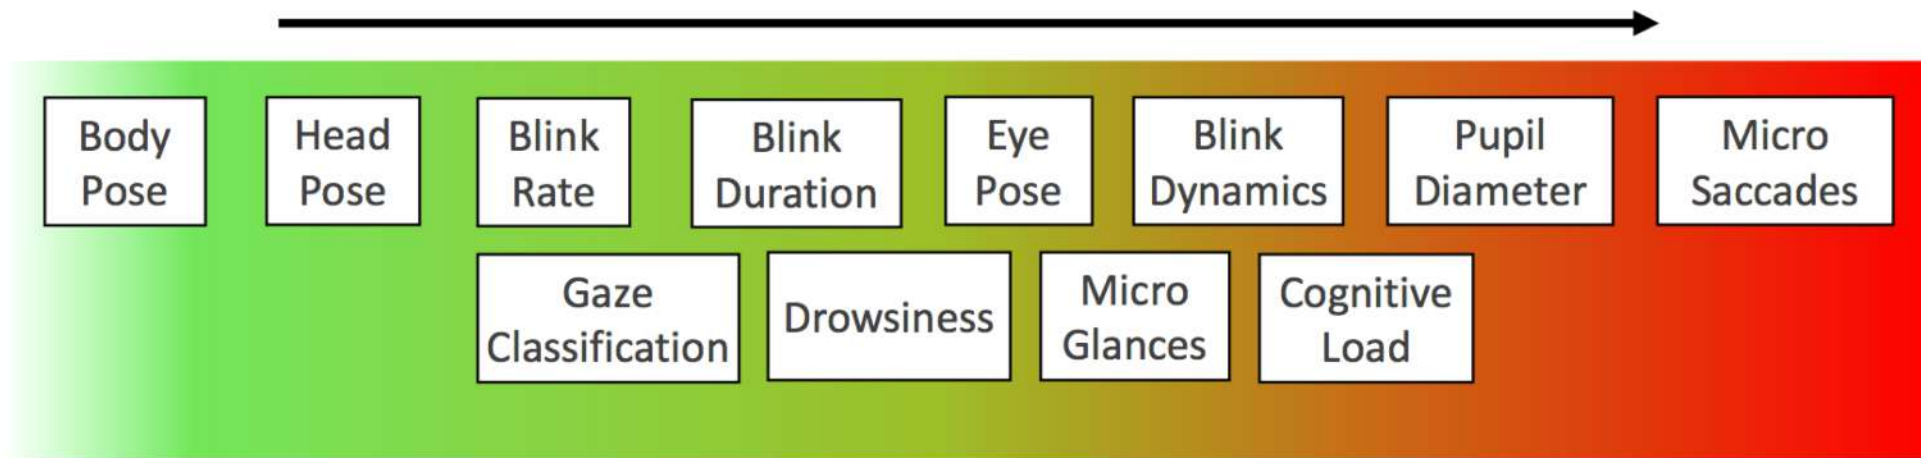
Human Interaction

How do I convey my intent to the passenger and everyone else ?



Drive State Detection: A Multi-Resolutional View

Increasing level of detection resolution and difficulty



Localization and Mapping

Where am I ?

Scene Understanding


Where/who/what/why of everyone/everything else ?

Trajectory Planning and Control

Where should I go next ?
How do I steer and accelerate ?

Human Interaction

How do I convey my intent to the passenger and everyone else ?



Localization and Mapping

Where am I ?

Scene Understanding

Where/who/what/why of everyone/everything else ?

Trajectory Planning and Control


Where should I go next ?
How do I steer and accelerate ?

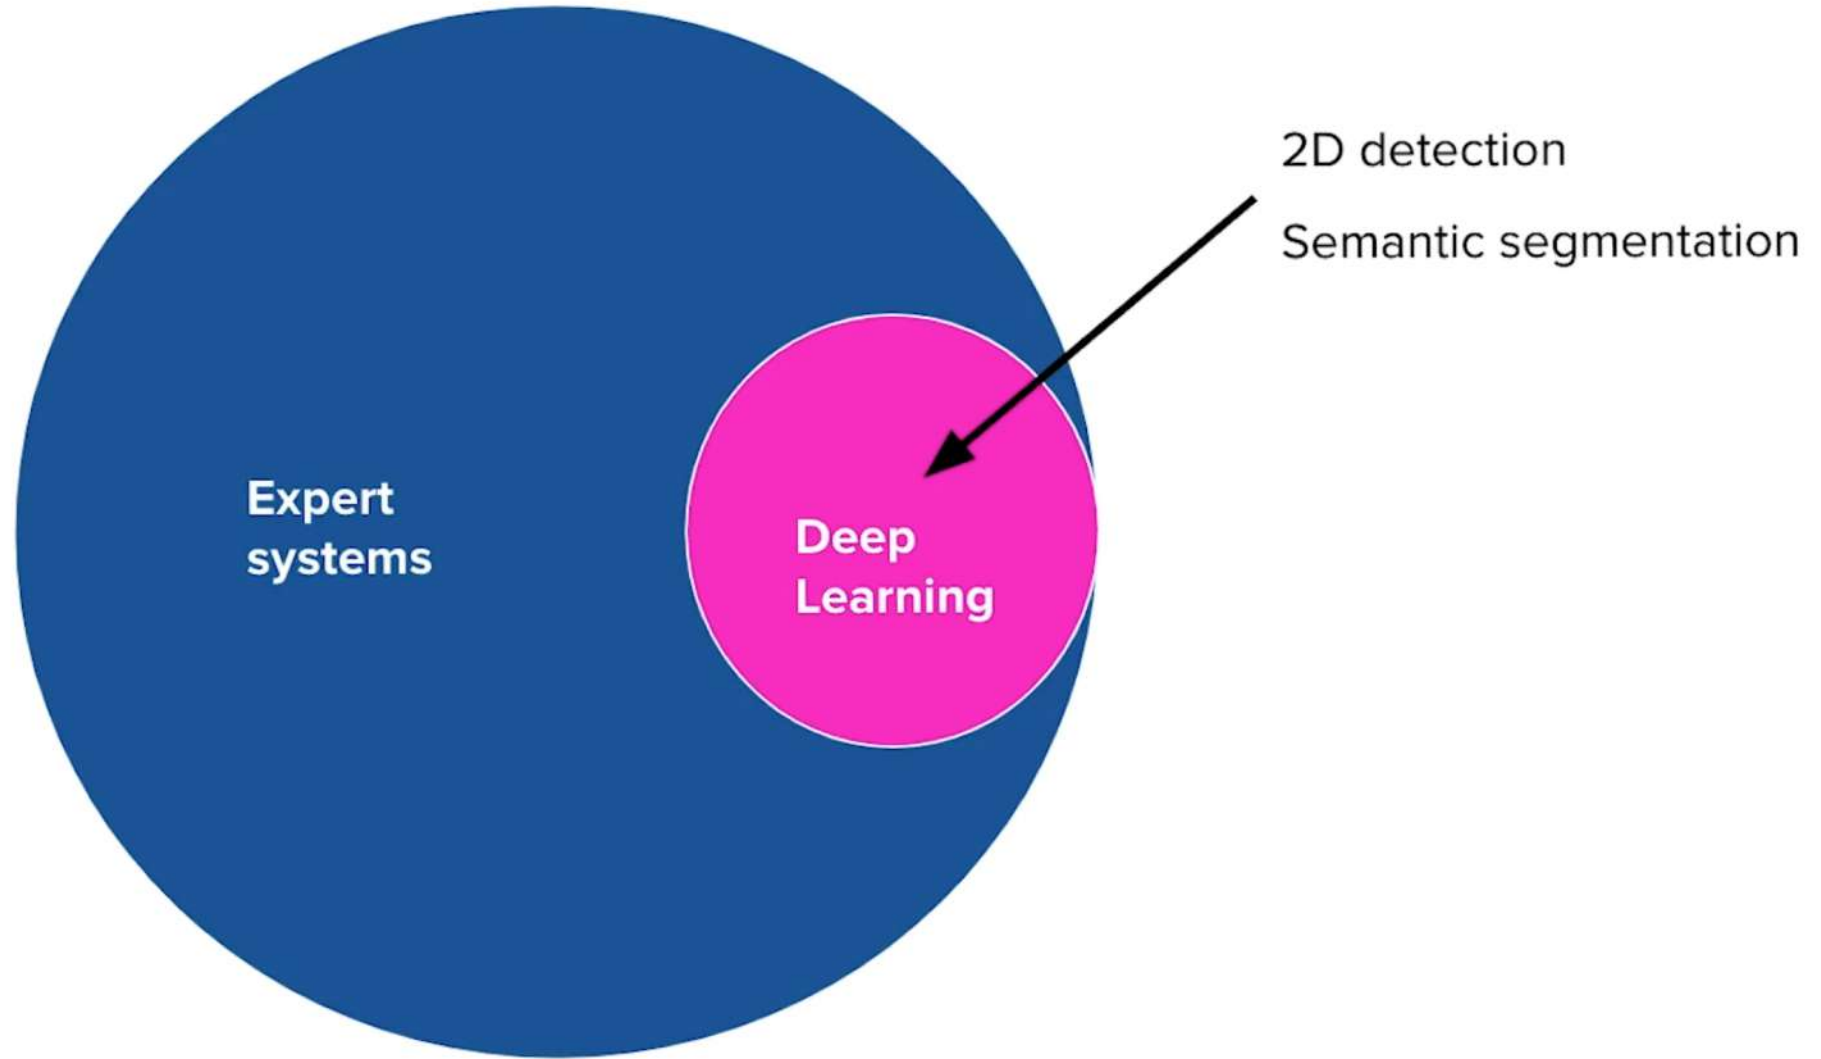
Human Interaction

How do I convey my intent to the passenger and everyone else ?

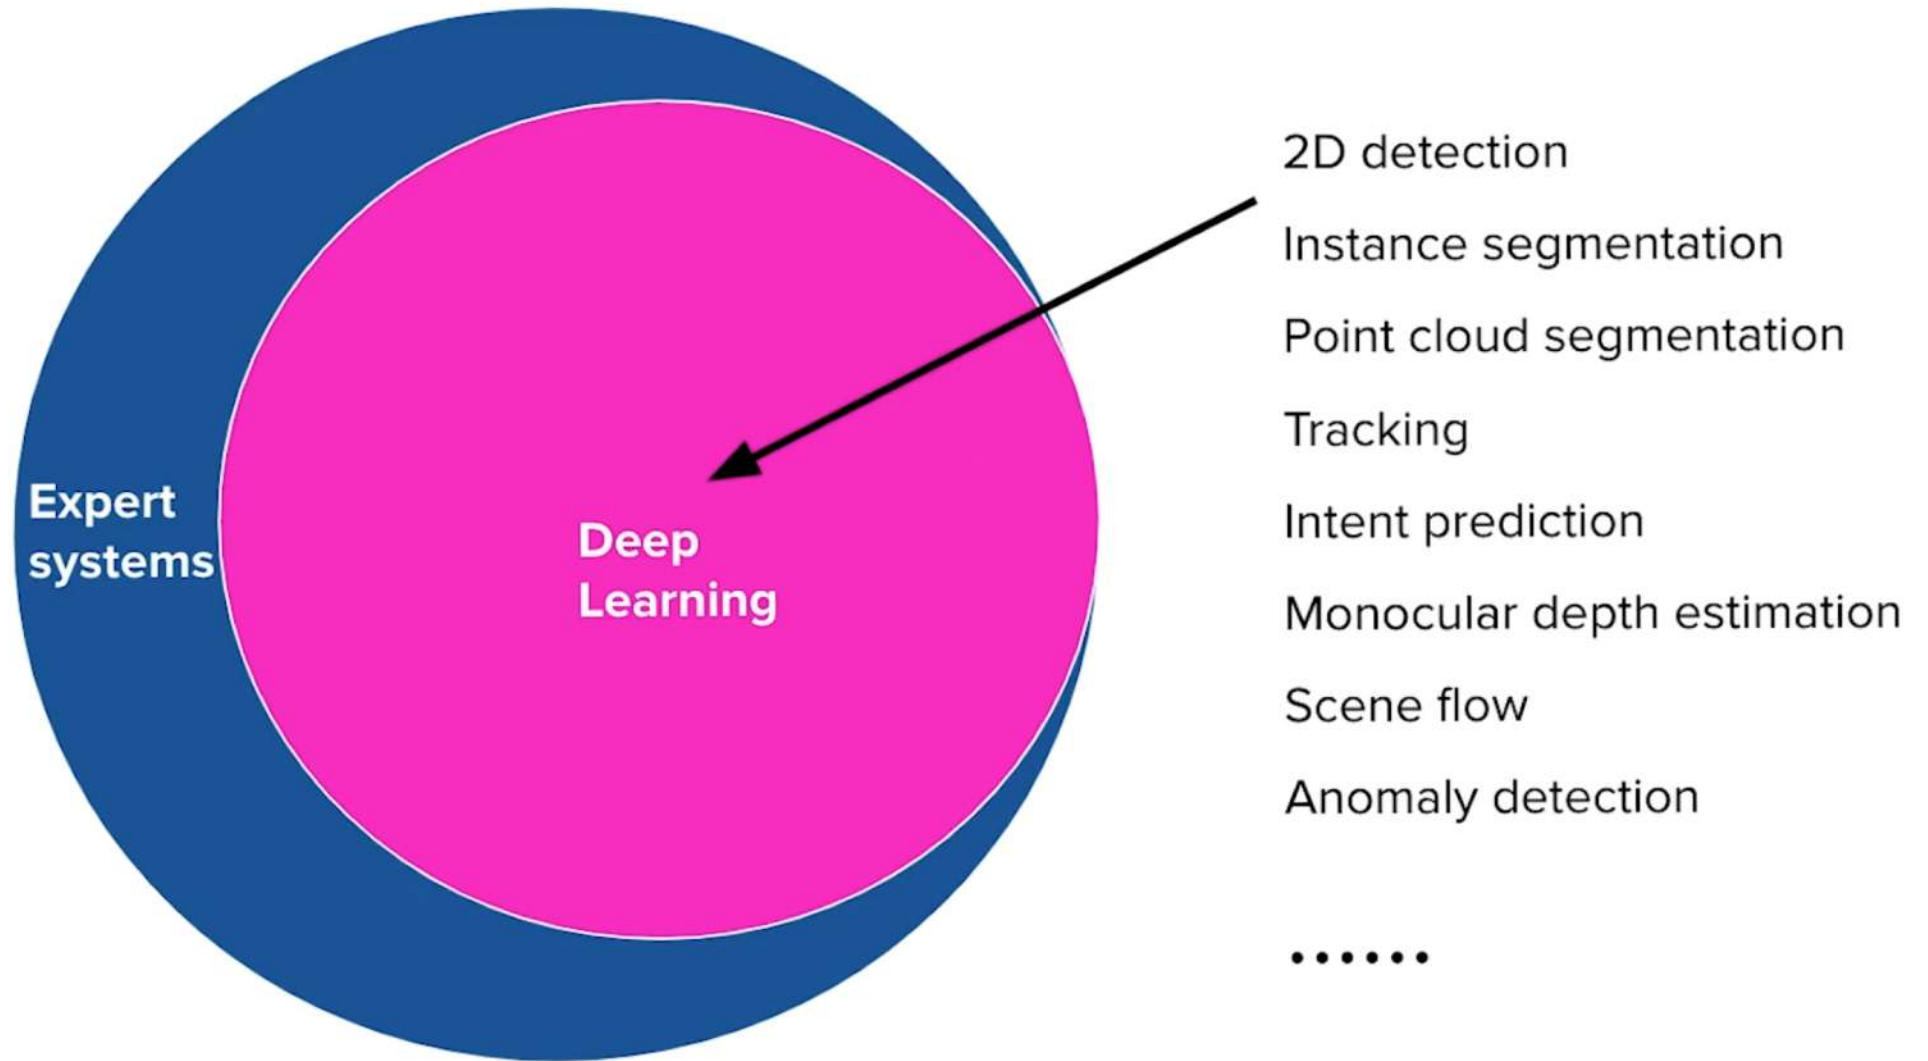
Deep Neural Networks

End-to-end learnable



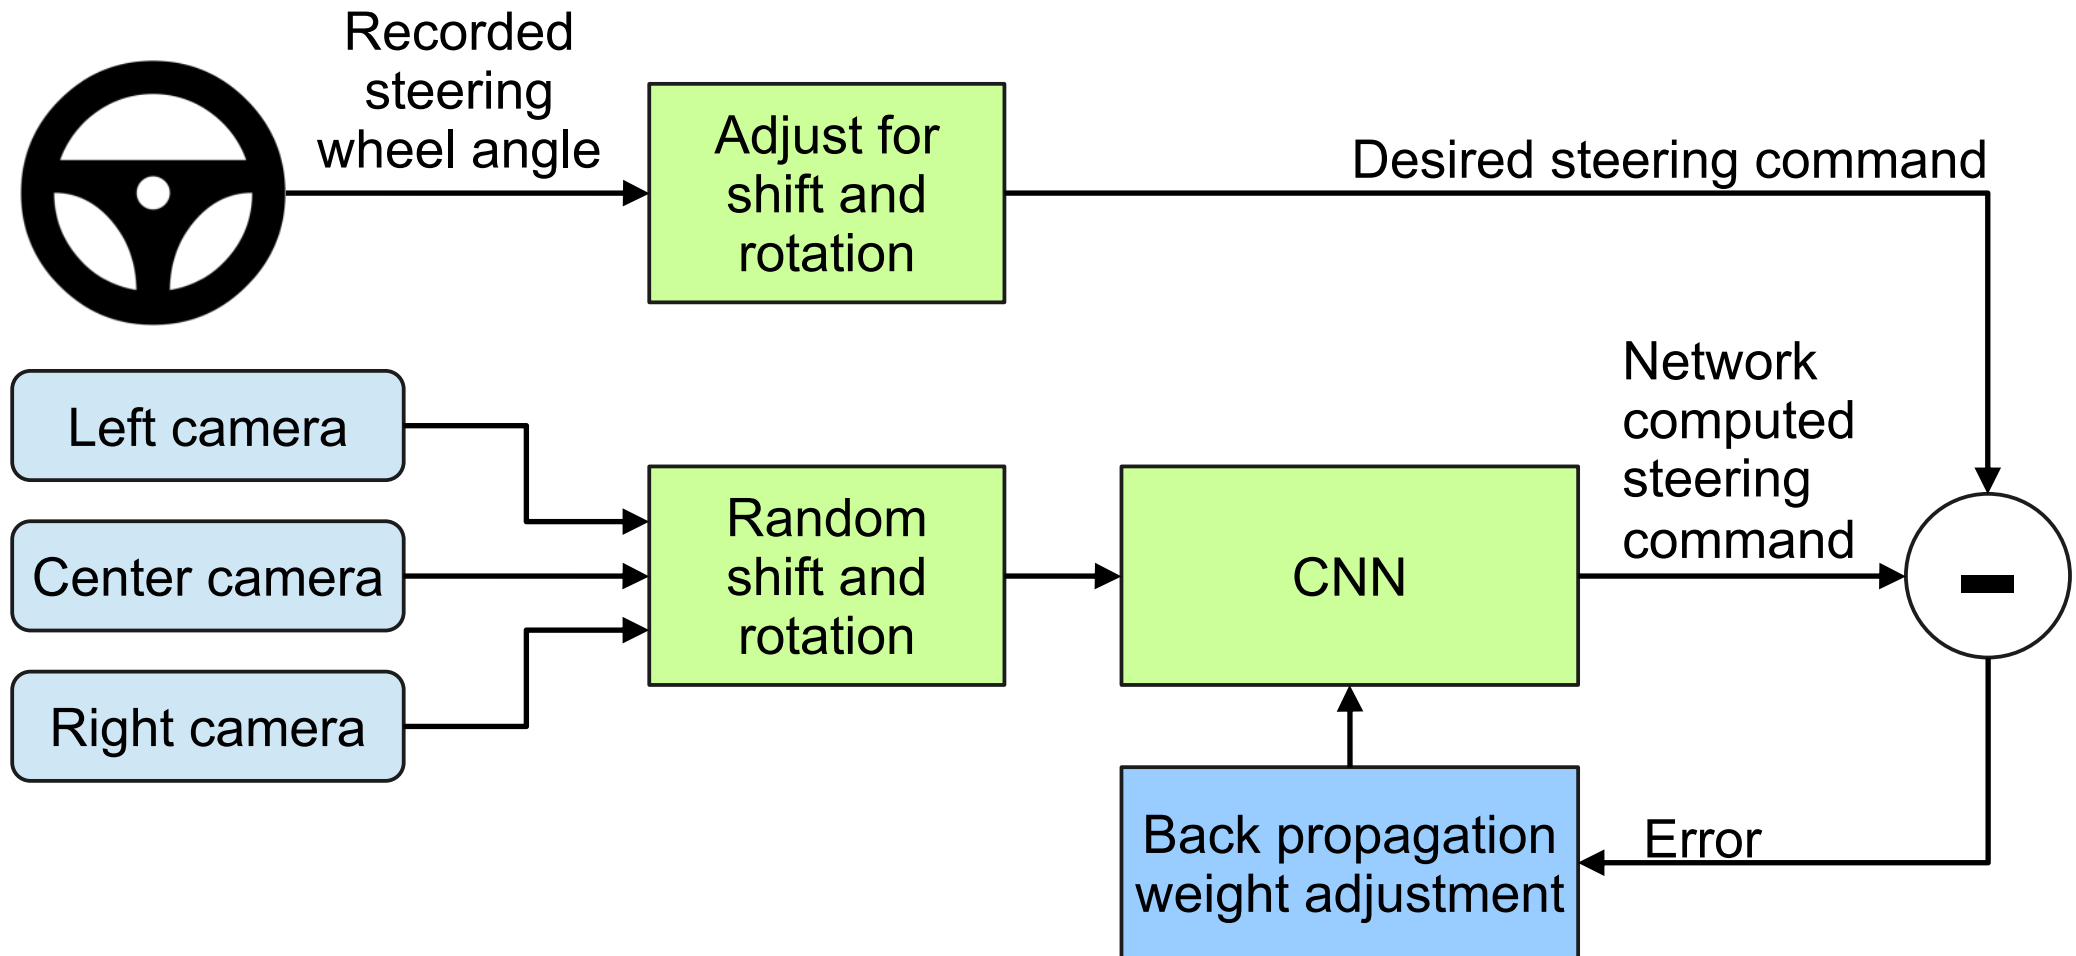


AV Perception in 2015

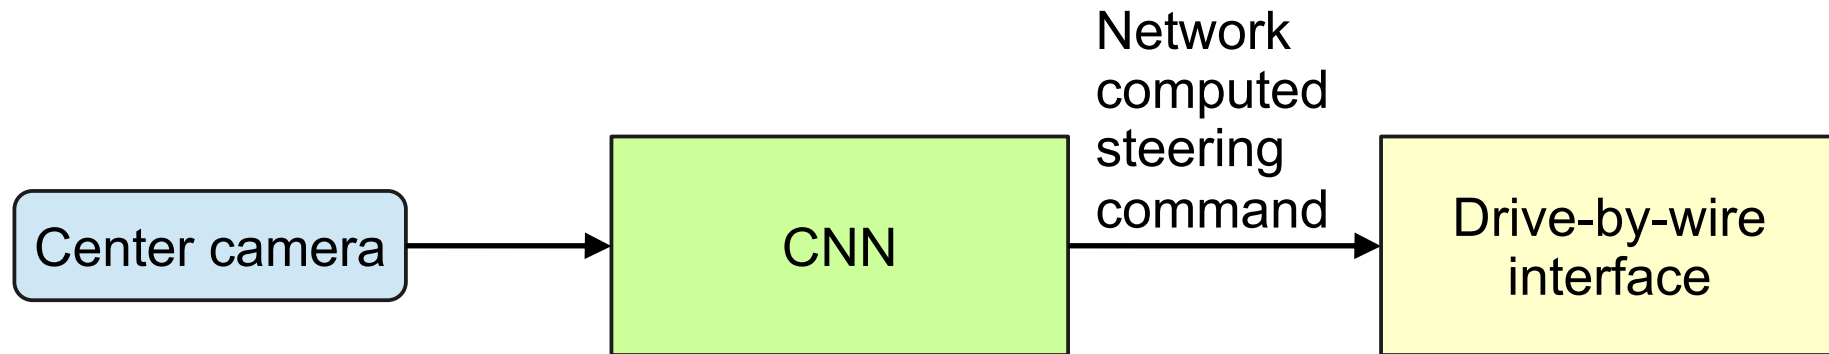


AV Perception today

End-to-End Driving: PilotNET




With a single front-facing camera



Machine intelligence is largely about training data.


When's a pedestrian not a pedestrian? When it's a decal.



One car ? or Multiple cars ?







Ramen Noodle place or Do Not Enter Sign ?



下鳥羽城/越町 付近
| 4:41:11 | 20m
現在地
AUDIO

今日 明日 16:10
現在地 12/5 9/3
気温(°C) 60 40
降水確率(%)
12°C A 77.5km 000160km







There is a bus right next to you!!









How can we ensure that an autonomous vehicle drives safely upon encountering an unusual traffic situation ?



How can CNNs help us drive ?



Tesla Control
(by Autopilot)

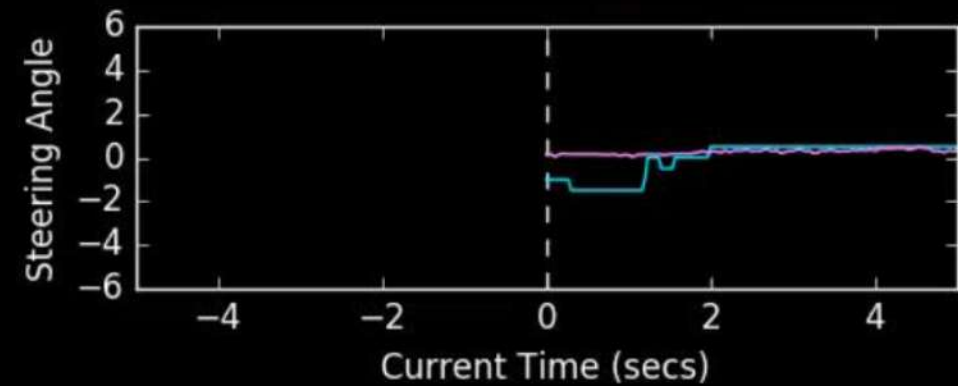



(Ground Truth)

Learned Control
(by Deep Neural Network)



Red = Disagree Green = Agreee





Actual wheel: 9.5

Predicted wheel: -2.5

Error: 12.0

Principles of Modeling for CPS – Fall 2019



Frame #: 961

Foward pass (ms): 51

MADHUR BEHL MADHUR.BEHL@VIRGINIA.EDU



Autonomous Driving: End-to-End

End to End Learning for Self-Driving Cars

Mariusz Bojarski
NVIDIA Corporation
Holmdel, NJ 07735

Davide Del Testa
NVIDIA Corporation
Holmdel, NJ 07735

Daniel Dworakowski
NVIDIA Corporation
Holmdel, NJ 07735

Bernhard Firner
NVIDIA Corporation
Holmdel, NJ 07735

Beat Flepp
NVIDIA Corporation
Holmdel, NJ 07735

Prasoon Goyal
NVIDIA Corporation
Holmdel, NJ 07735

Lawrence D. Jackel
NVIDIA Corporation
Holmdel, NJ 07735

Mathew Monfort
NVIDIA Corporation
Holmdel, NJ 07735

Urs Muller
NVIDIA Corporation
Holmdel, NJ 07735

Jiakai Zhang
NVIDIA Corporation
Holmdel, NJ 07735

Xin Zhang
NVIDIA Corporation
Holmdel, NJ 07735


Jake Zhao
NVIDIA Corporation
Holmdel, NJ 07735

Karol Zieba
NVIDIA Corporation
Holmdel, NJ 07735

Autonomous Driving: End-to-End

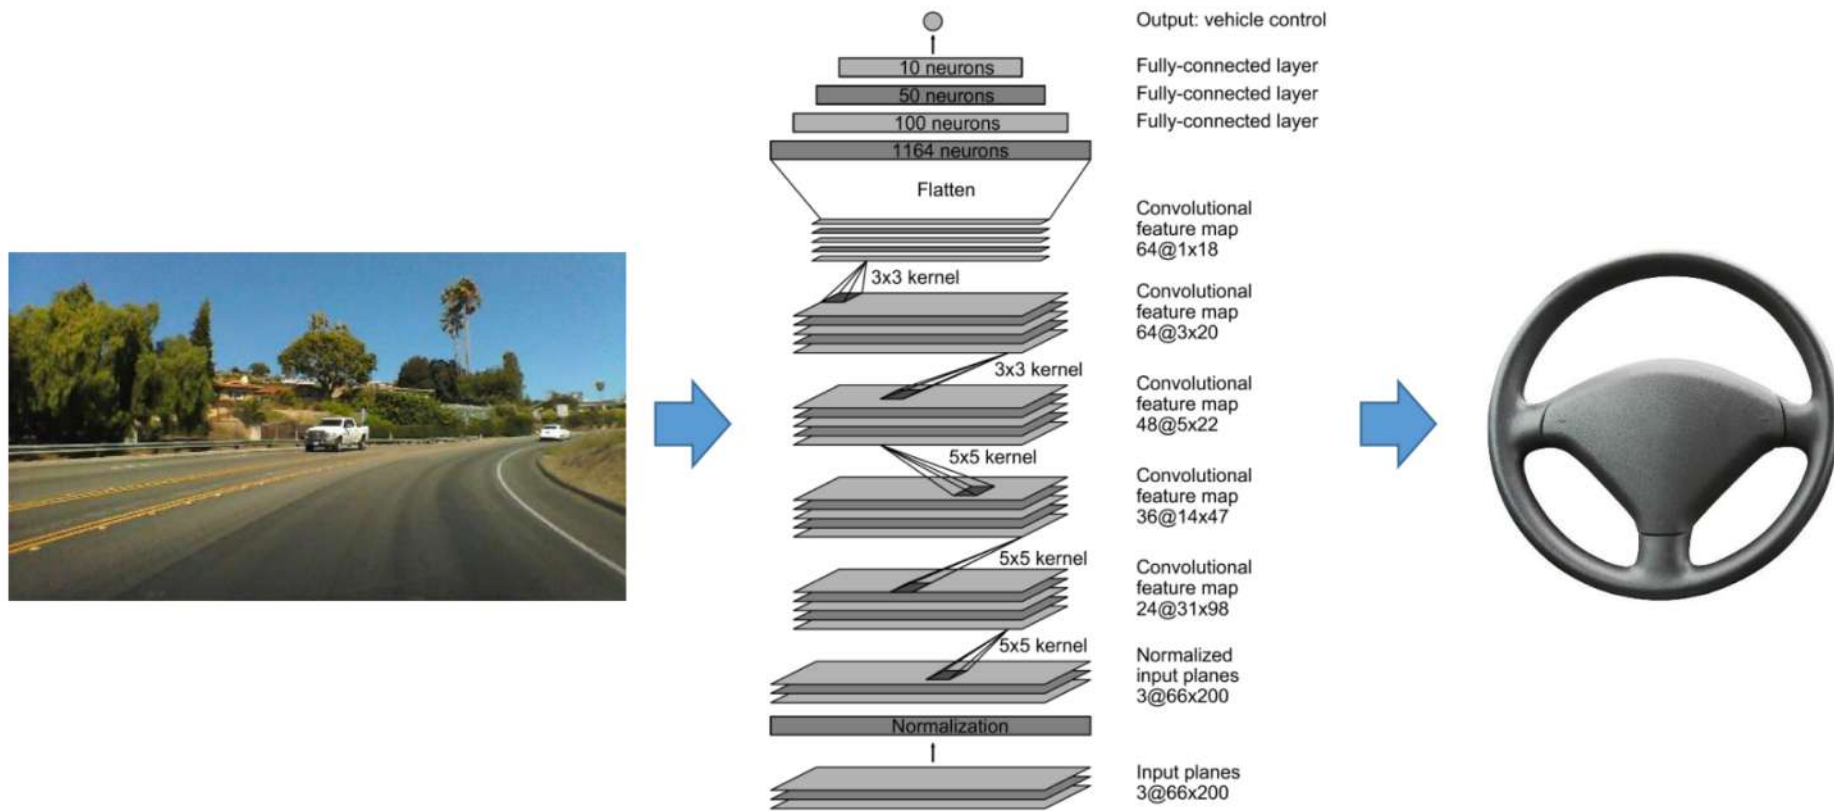


Autonomous Driving: End-to-End



- 9 layers
 - 1 normalization layer
 - 5 convolutional layers
 - 3 fully connected layers
- 27 million connections
- 250 thousand parameters

Autonomous Driving: End-to-End





F1/10 FPV Driving







Predicted: -0.04 Ground Truth: -0.00

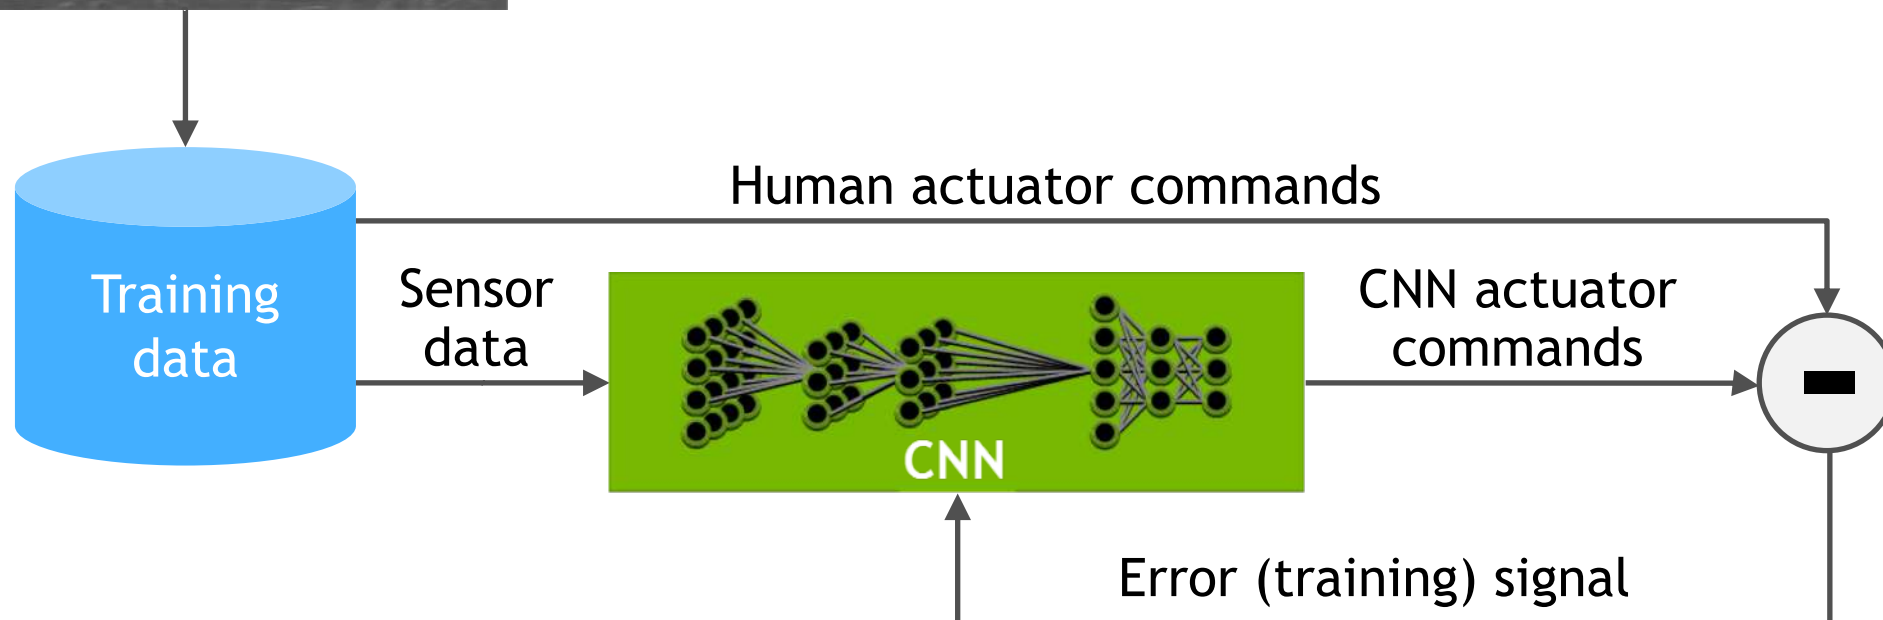
THE BASIC IDEA

Learn from human drivers




Record data from lots of humans driving their cars:

- Sensor data
- Actuator data



EARLY EXAMPLES

Of end-to-end learning



ALVINN, CMU, late 80es
(Pomerleau et Al.)

Lane following with a small 2-layer fully connected network and low-resolution video input

30x32 pixel

DAVE, Net-Scale/NYU, 2004
(LeCun et Al.)

Off-road obstacle avoidance using a convolutional network (ConvNet)

149x48 pixel

TRAINING EXAMPLES



Label: turn right



Label: go straight



Label: turn right



Label: turn left

225K images

Machine Learning

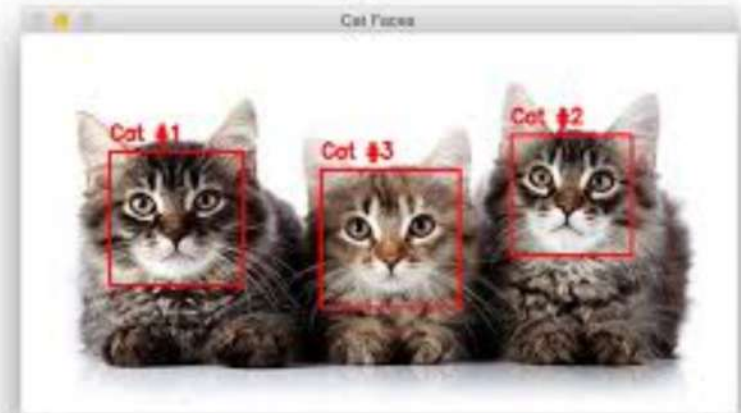
- Machine Learning is the ability to teach a computer without explicitly programming it
- Examples are used to train computers to perform tasks that would be difficult to program

First Name

L	O	R	I								
---	---	---	---	--	--	--	--	--	--	--	--




Last Name

W	A	L	T	E	R	S					
---	---	---	---	---	---	---	--	--	--	--	--



Types of machine Learning

- Supervised Learning
 - Training data is labeled
 - Goal is correctly label new data
- Reinforcement Learning
 - Training data is unlabeled
 - System receives feedback for its actions
 - Goal is to perform better actions
- Unsupervised Learning
 - Training data is unlabeled
 - Goal is to categorize the observations



Capability of Machine to imitate intelligent behavior

ARTIFICIAL INTELLIGENCE

Early artificial intelligence stirs excitement.



MACHINE LEARNING

Machine learning begins to flourish.



DEEP LEARNING

Deep learning breakthroughs drive AI boom.



1950's

1960's

1970's

1980's

1990's

2000's

2010's

Supervised learning setup

Inputs (AKA features) - real-valued vectors of data

e.g. Image pixels, audio spectrograms, character sequences

Outputs (AKA labels) - real-valued or categorical “truth” vectors

e.g. class labels for images, audio transcription, sentiment

Training data - many samples of input-output pairs

Score function (AKA model)


A function that predicts the output given an input

Example: linear regression

$$y_i = mx_i + c$$

Predicted output \rightarrow y_i \downarrow Data \rightarrow x_i \uparrow Slope m \leftarrow Intercept c

Together, m and c are called the **model parameters**




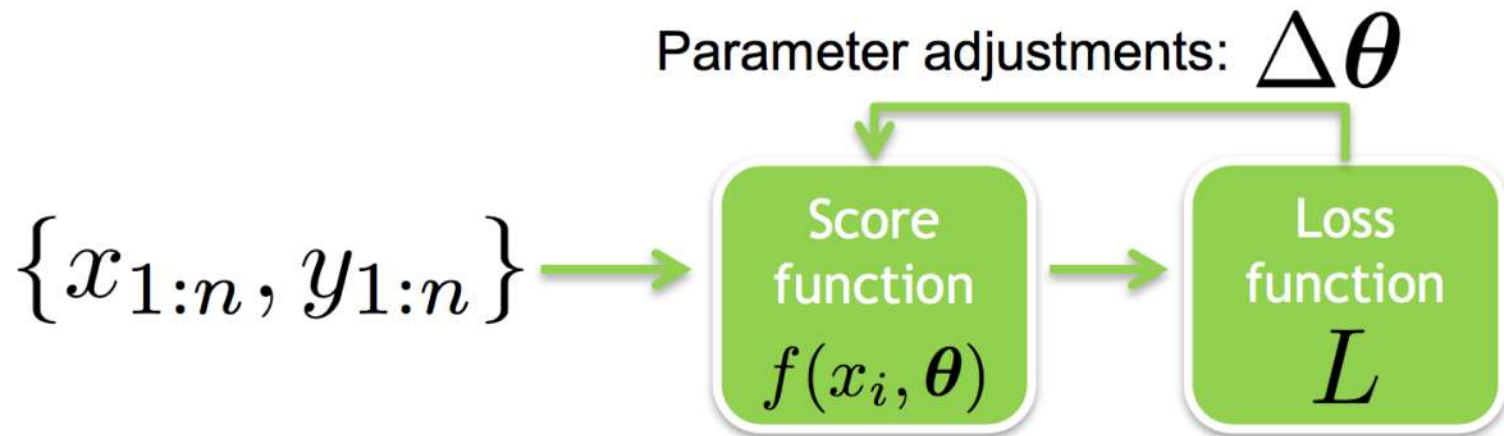
Supervised learning

How do we do this?

Repeatedly feed training data into a learning algorithm



Iteratively modify the model parameters to optimize (e.g. minimize) the loss function

Repeat until the model is “good enough”



Gradient descent

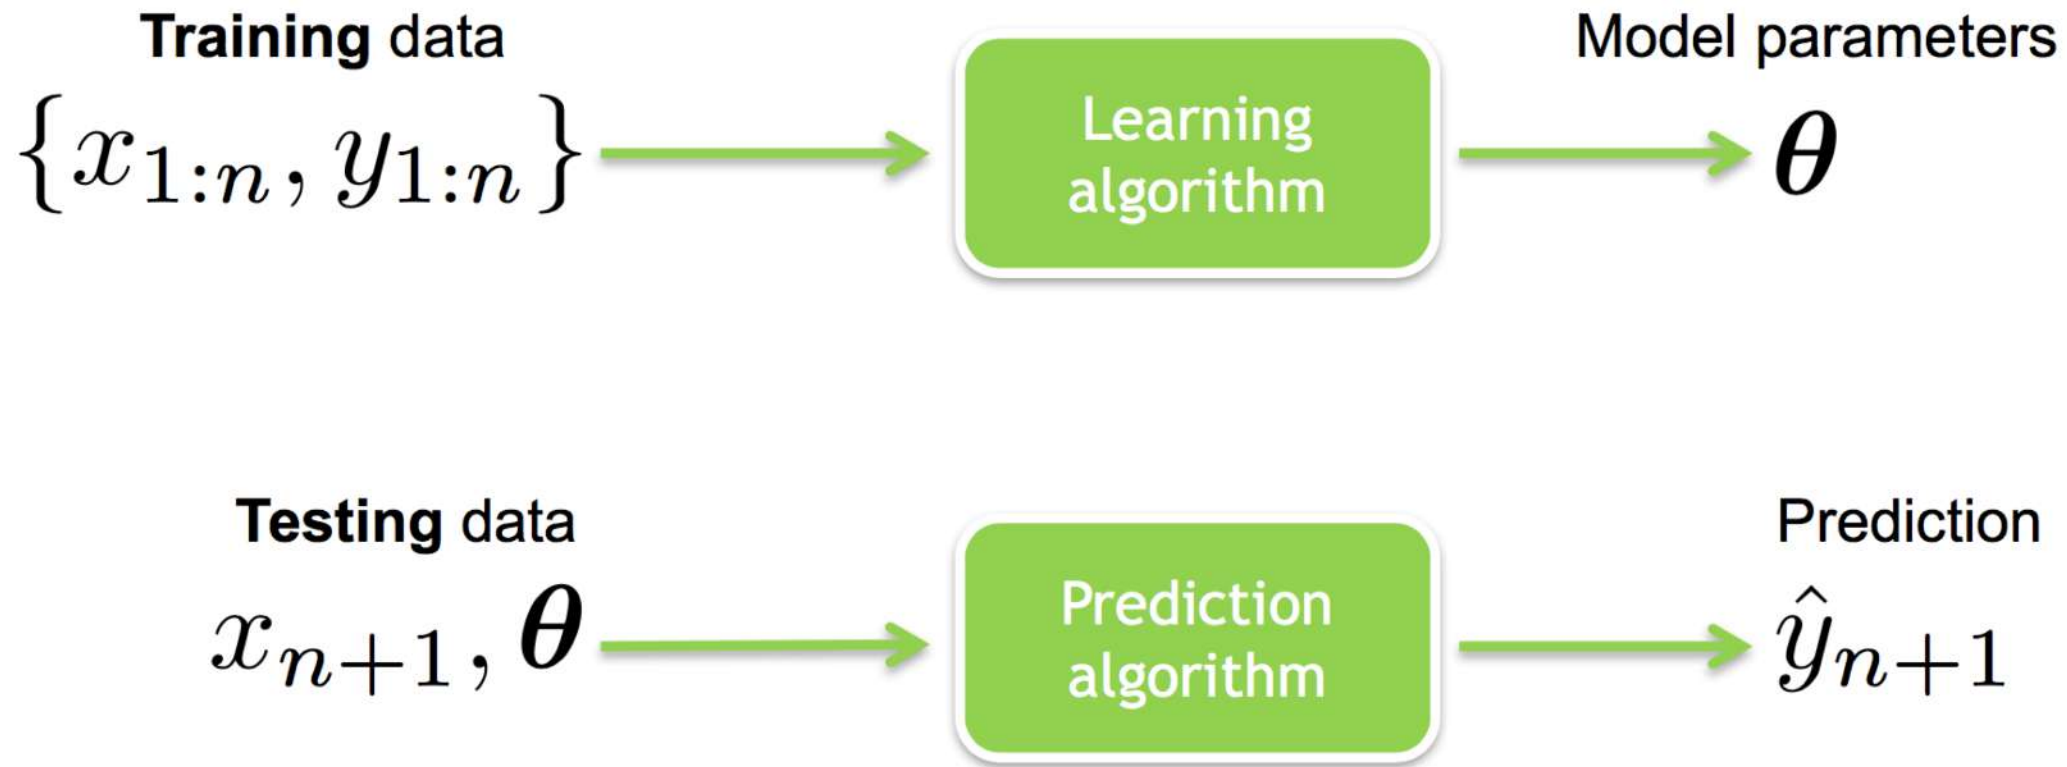
Finding the Optimal Parameters for our Hypothesis



Supervised learning


Why do we do this?

Given the **model** we can take previously unseen inputs and predict the corresponding output. We call this **testing** or **deployment**.

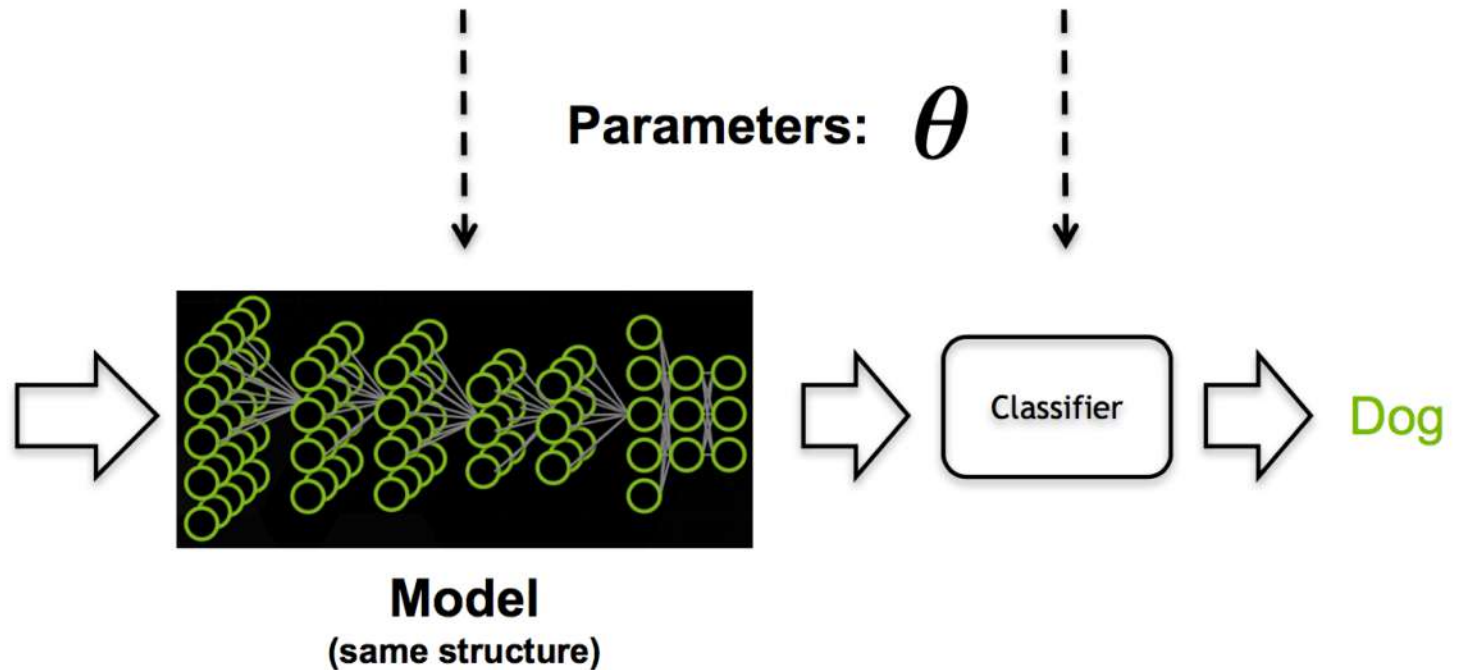


Example: ImageNet

Training:



Testing:



Deep Learning success



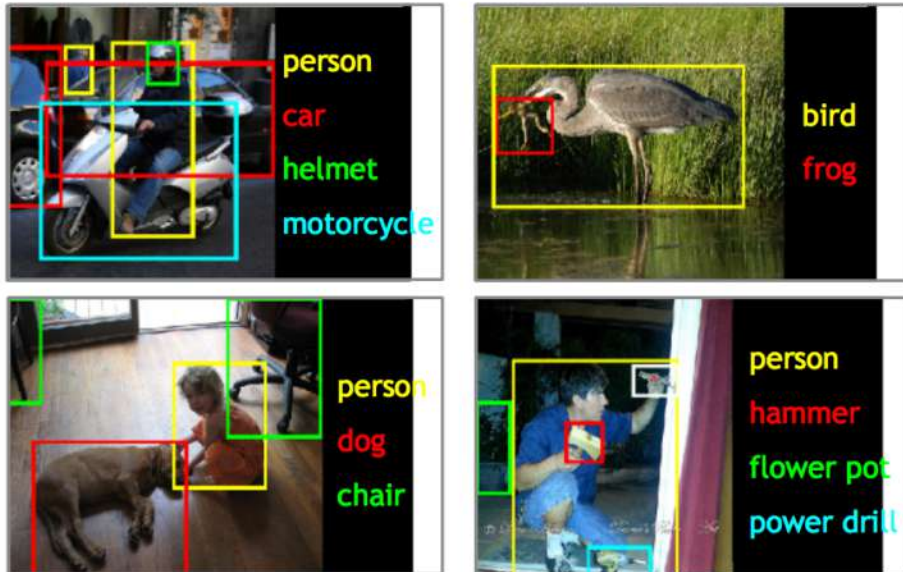
Object classification and localization in images

Image Recognition Challenge

1.2M training images • 1000 object categories

Hosted by

IMAGENET



Training problems

Two major problems

Underfitting: model is bad at it's objective for all data

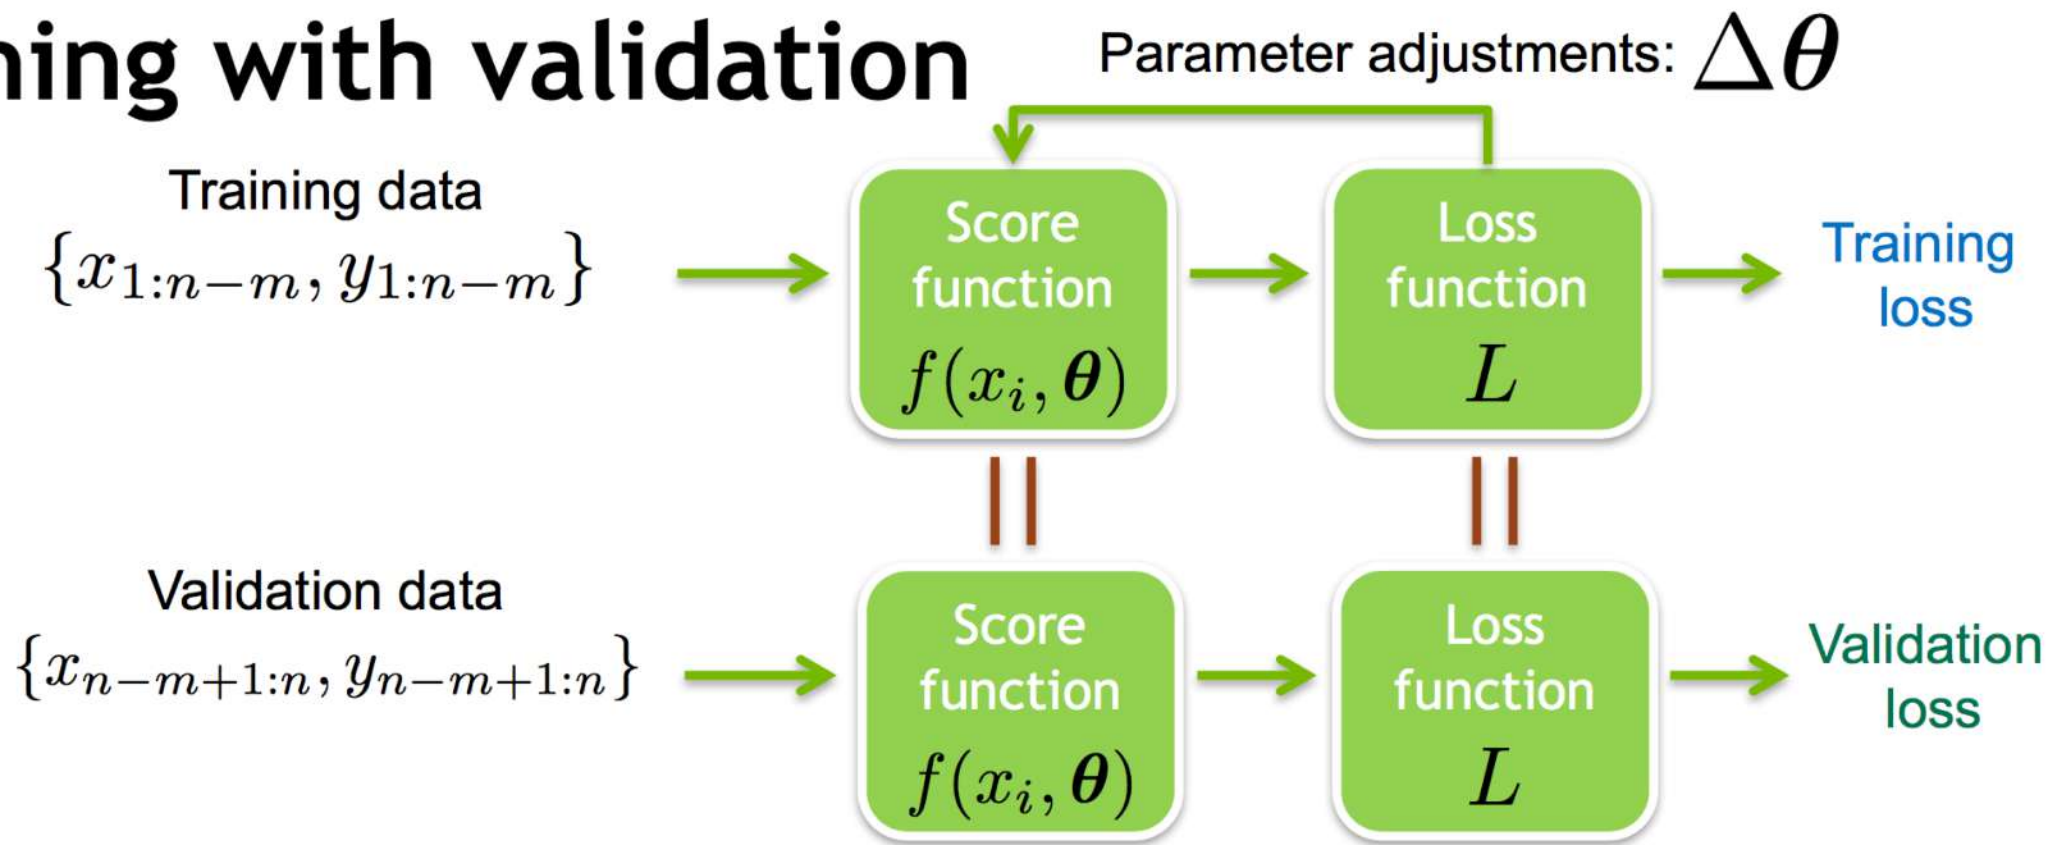
Overfitting: model is really good at the objective for the training data but bad on the testing data

First line of defense:



Break off a **validation** dataset from the training data, e.g. 25%

Use it during training to check model performance on unseen data

Training with validation



Original training data




Deep Learning

Learning Representation/Features

The traditional model of pattern recognition (since the late 50's)
Fixed/engineered features (or fixed kernel) + trainable classifier

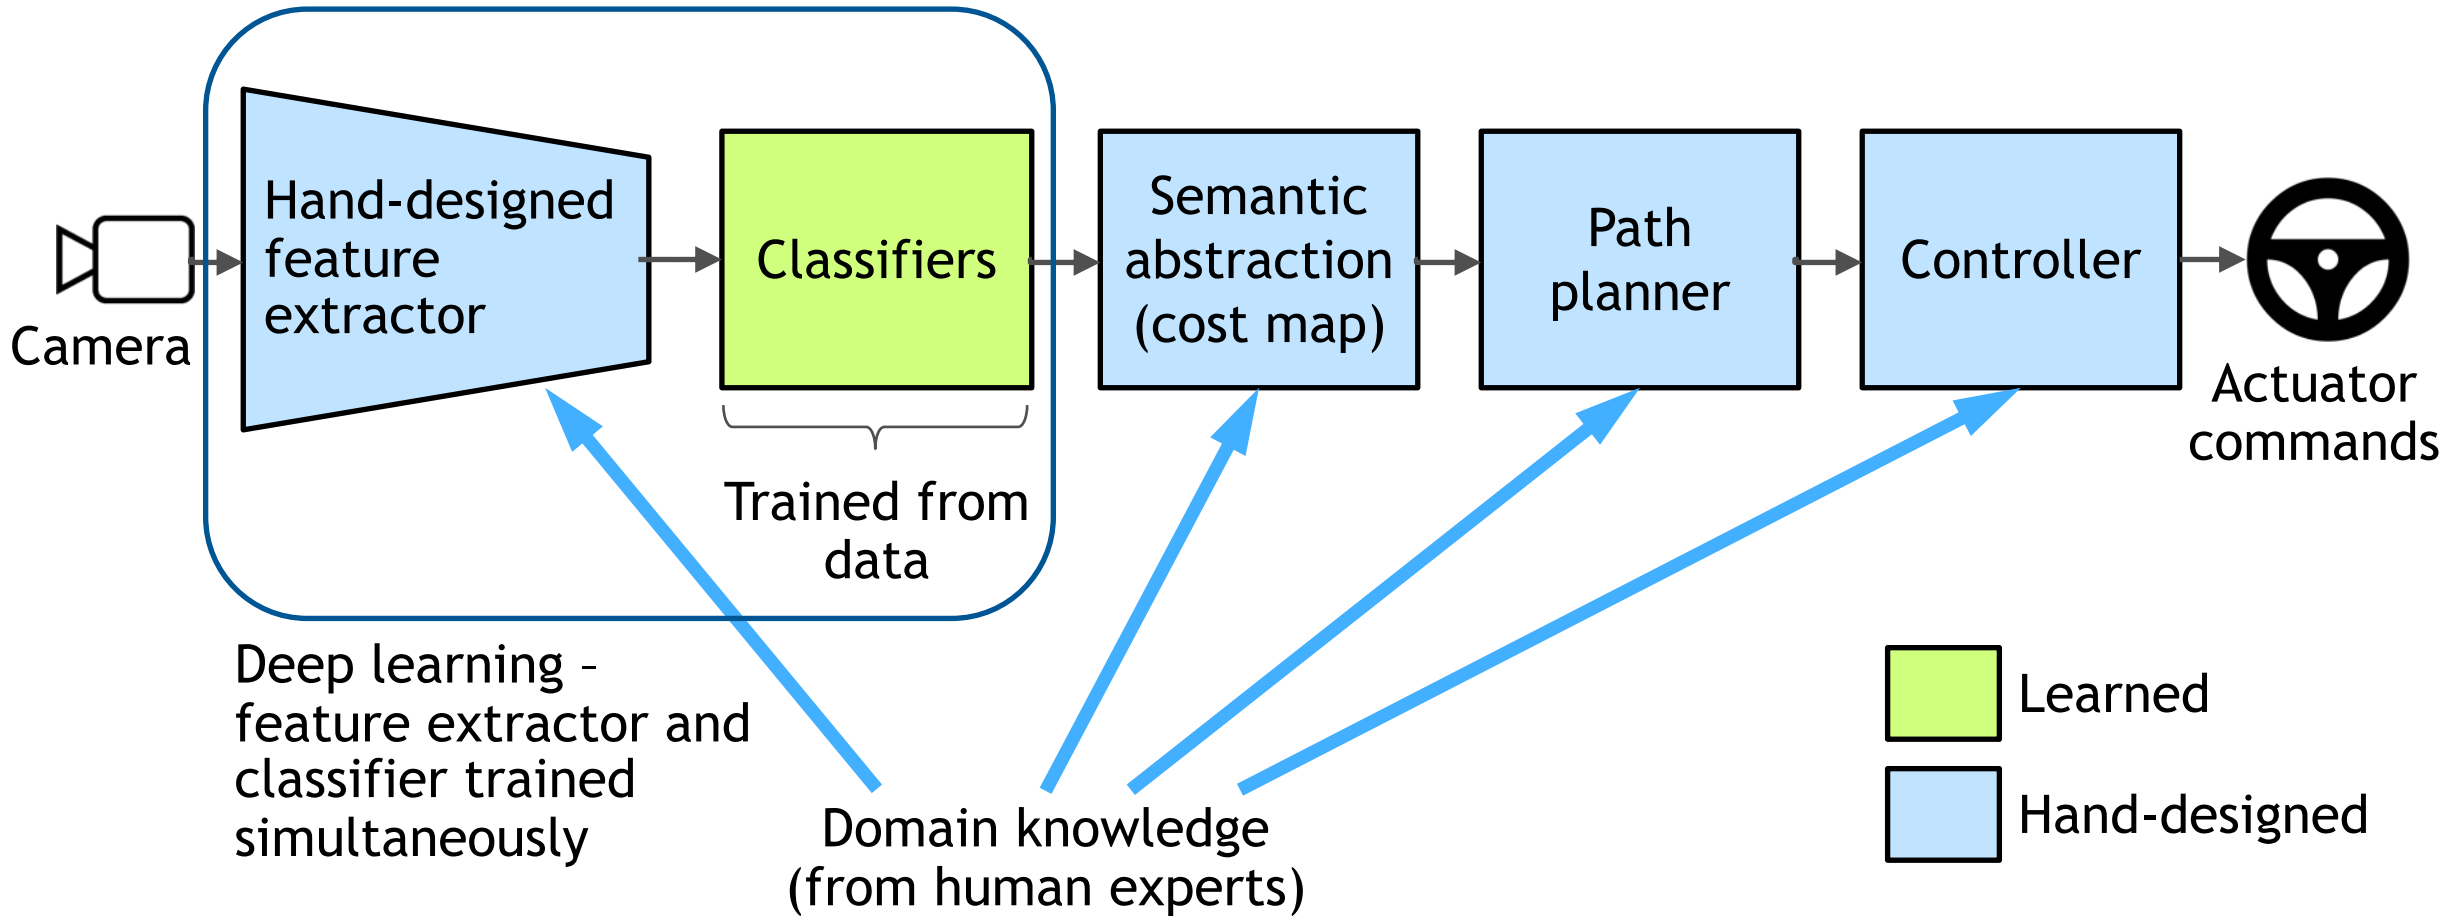


End-to-end learning / Feature learning / Deep learning
Trainable features (or kernel) + trainable classifier



TRADITIONAL DECOMPOSITION

Necessary approach when data and compute power are limited



EXAMPLE: ROAD FOLLOWING



Good quality lane markers, good driving conditions

Traditional lane detection-based systems expected to work well



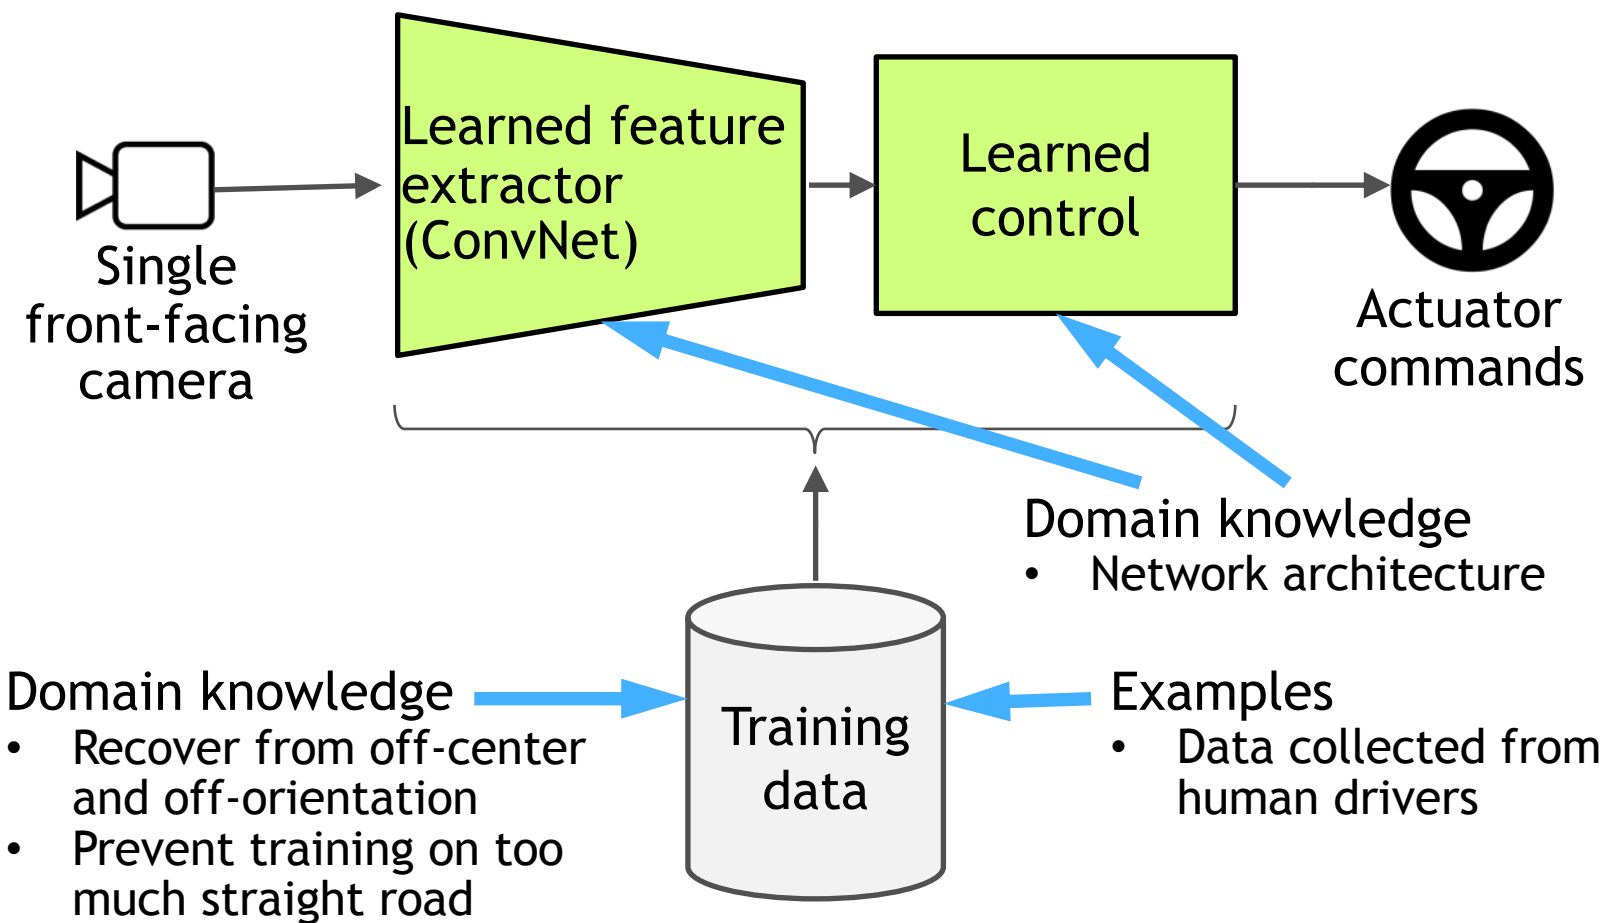
Poor quality lane markers

Lane detection-based systems struggle

End-to-end learning empowers the network to use additional cues

LEARNED ROAD FOLLOWING (PILOTNET)


Highway, local, residential - with or without lane markings



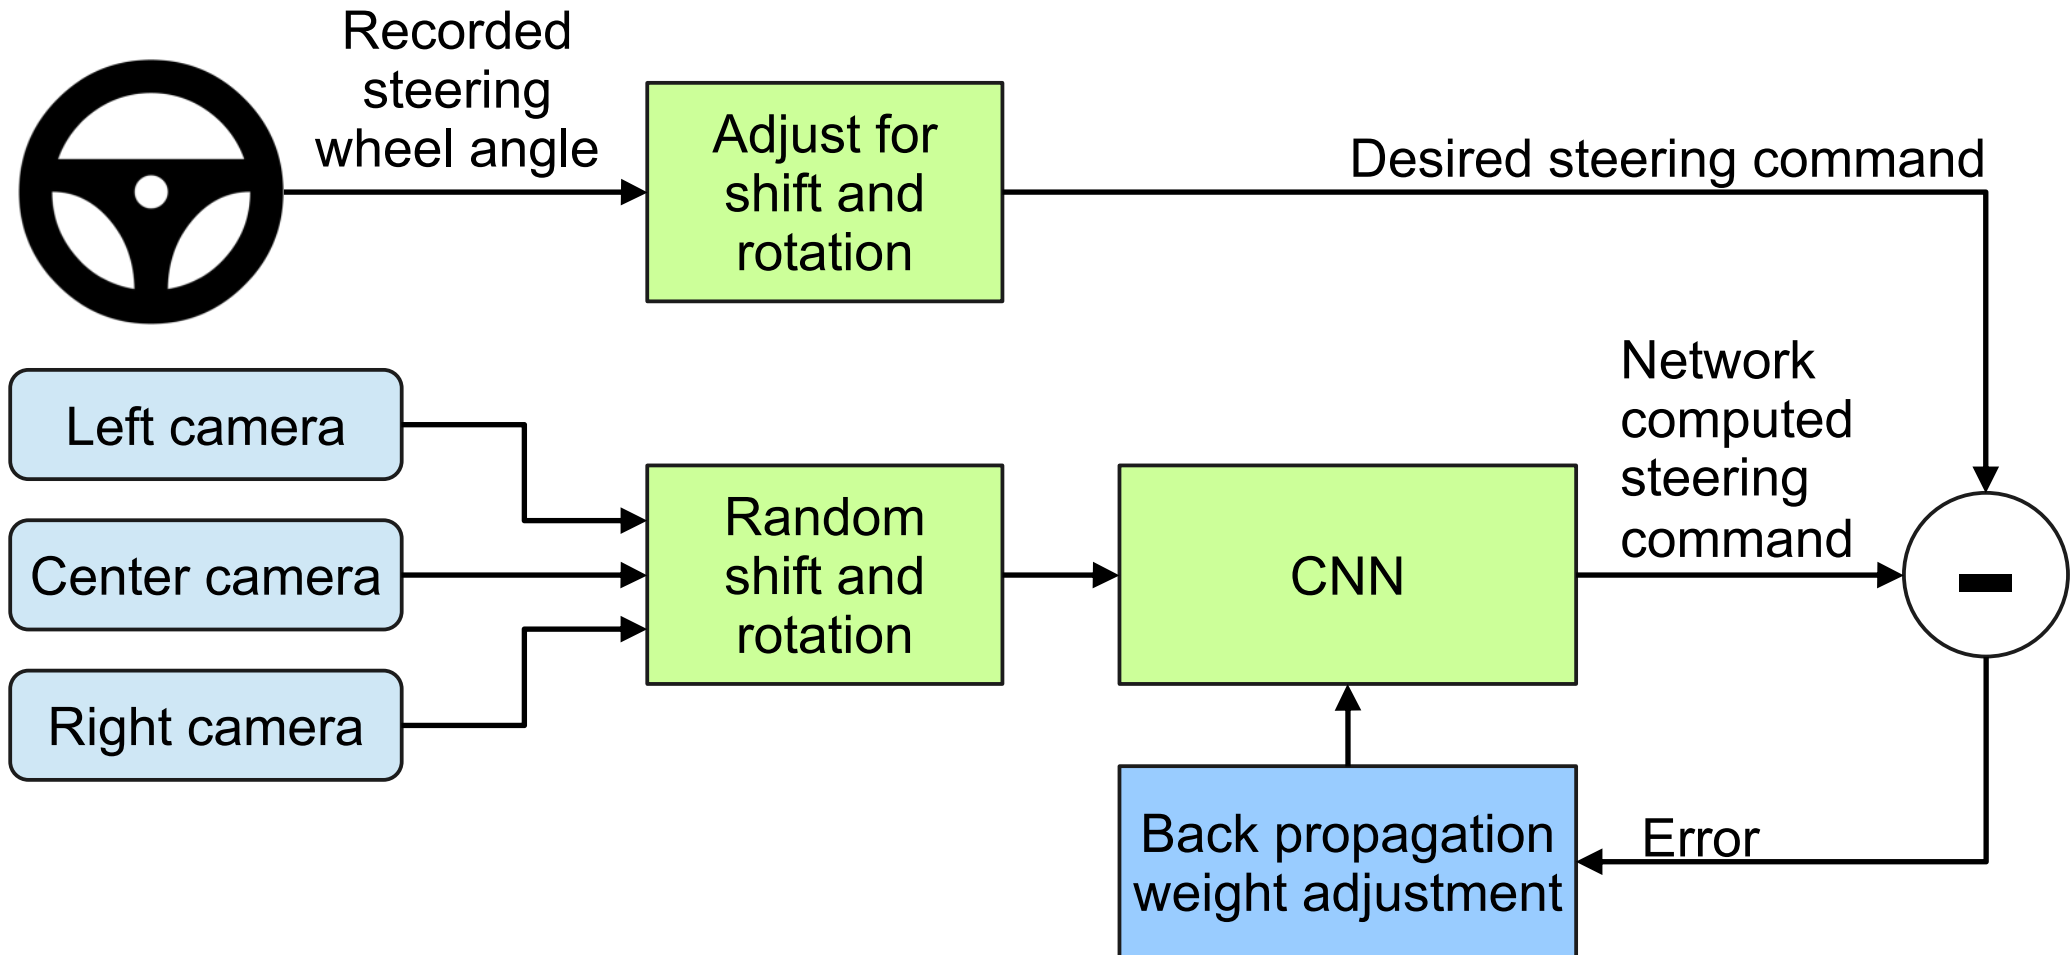
Both blocks trained simultaneously

No explicit object detection nor path planning

→ Maps pixels directly to steering




TRAINING THE NEURAL NETWORK



DRIVING

With a single front-facing camera





VISUALIZATION

What the network pays attention to