Artificial intelligence for driving

Stéphane Canu
asi.insa-rouen.fr/enseignants/~scanu
scanu@insa-rouen.fr

« ITI »

Lundi 2 septembre 2024

Road map

1. A very brief history of autonomous vehicles
2. How has this happened? (Deep Learninng)
3. Data to train the deep network
4. How Artificial Intelligence will change the Automotive Industry
5. Conclusion

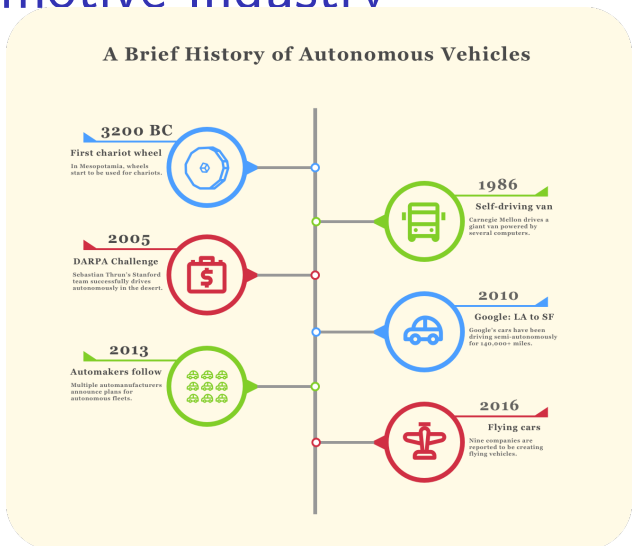

Artificial intelligence breakthroughs

2005 How to drive
DARPA Gran Challenge

2012 How to recognize objects
ImageNet competition

2016 How to play games
AlphaGO

2022 Question answering – generating natural language
GPT -> chatGPT

How?

Specific AI using a particular machine learning tool: deep learning

Artificial intelligence and autonomous vehicles

Artificial intelligence is about doing things better than human

$\rightarrow$ It can do a lot of things better than a human driver

https://www.forbes.com/sites/lauriewinkless/2016/05/02/is-tomorrows-car-just-a-case-of-history-repeating-itself

Artificial intelligence and autonomous vehicles

From Chandler to Chandler

$\rightarrow$ Can AI drive better than a human driver?

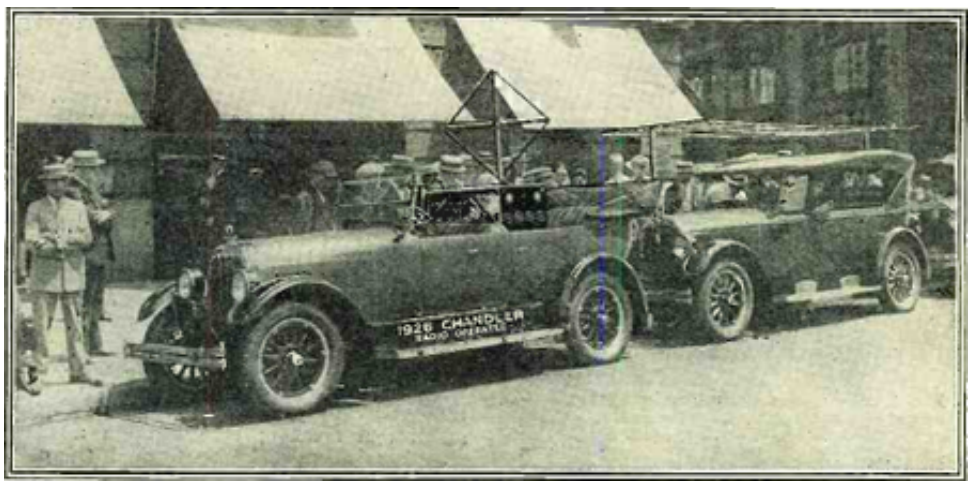

1925

2023 .

https://en.wikipedia.org/wiki/History_of_self-driving_cars

NavLab: the autonomous vehicle of the 80s

1 M $, 10 km/h

http://www.rediscoverthe80s.com/2016/11/navlab-the-selfdriving-car-of-the-80s.html/

DARPA Grand Challenge 2004

- for American autonomous vehicles only
- 1 million $
- 140 miles (225km) from Barstow, California to Primm, Neveda

None of the 15 finalists completes more than 12 km of the race

DARPA Grand Challenge 2005

- 2 million $ – 132 miles (213 km) in the desert, Primm, Neveda.

Stanford vs Carnegie Mellon

212.7 km at an average speed of 30.7 km/h

Key issue

"The specific transfer function emulates human driving characteristics, and is learned from data gathered through human driving."

CMU autonomous vehicles

Carnegie Mellon University
30 Years of Self-Driving Car Research

Happy Birthday!

1984
- The Terregator's top speed was a few centimeters per second; it could avoid obstacles.
- NavLab launched. Its goal: apply computer vision, sensors and high-speed processors to create vehicles that drive themselves.

1986
Humans or computers controlled NavLab1, a Chevy van. Top speed: 20 mph.

1990
NavLab 2, a US Army HMMWV, wrangled rough terrain at 6 mph. Highway speed: 70 mph.

1995
NavLab 5, a Pontiac Trans Sport, traveled from Pittsburgh to San Diego in the "No Hands Across America Tour."

2000
NavLab 11, a Jeep, was equipped with Virtual Valet.

2005
Sandstorm and Highlander placed 2nd and 3rd in the DARPA Grand Challenge.

2007
Carnegie Mellon's "Boss" won the DARPA Grand Urban Challenge by outmaneuvering other vehicles along the 55-mile course.

2014
Carnegie Mellon's **14th self-driving vehicle** is a Cadillac SRX that:
- avoids pedestrians and cyclists
- takes ramps and merges
- recognizes and obeys traffic lights
- looks like other Cadillac SRXs

www.engineering.cmu.edu

Autonomous vehicles: when?

Tesla : prévues en 2014, 2015, 2016, 2018, 2019, 2020, 2021 et 2022, les voitures autonomes sont maintenant promises pour 2025

Un jour, Elon Musk aura raison.

Elon Musk removing his hands from the wheel with Autopilot engaged during an interview (Bloomberg, 2014).
numerama.com/vroom/972975-tesla-prevues-en-2014-2015-2016-2018-2019-2020-2021-et-2022-les-voitures-autonomes-sont-maintenant-promises-pour-2023.html

Level 3 Partial Autonomy adoption is when the market "tips" into autonomy

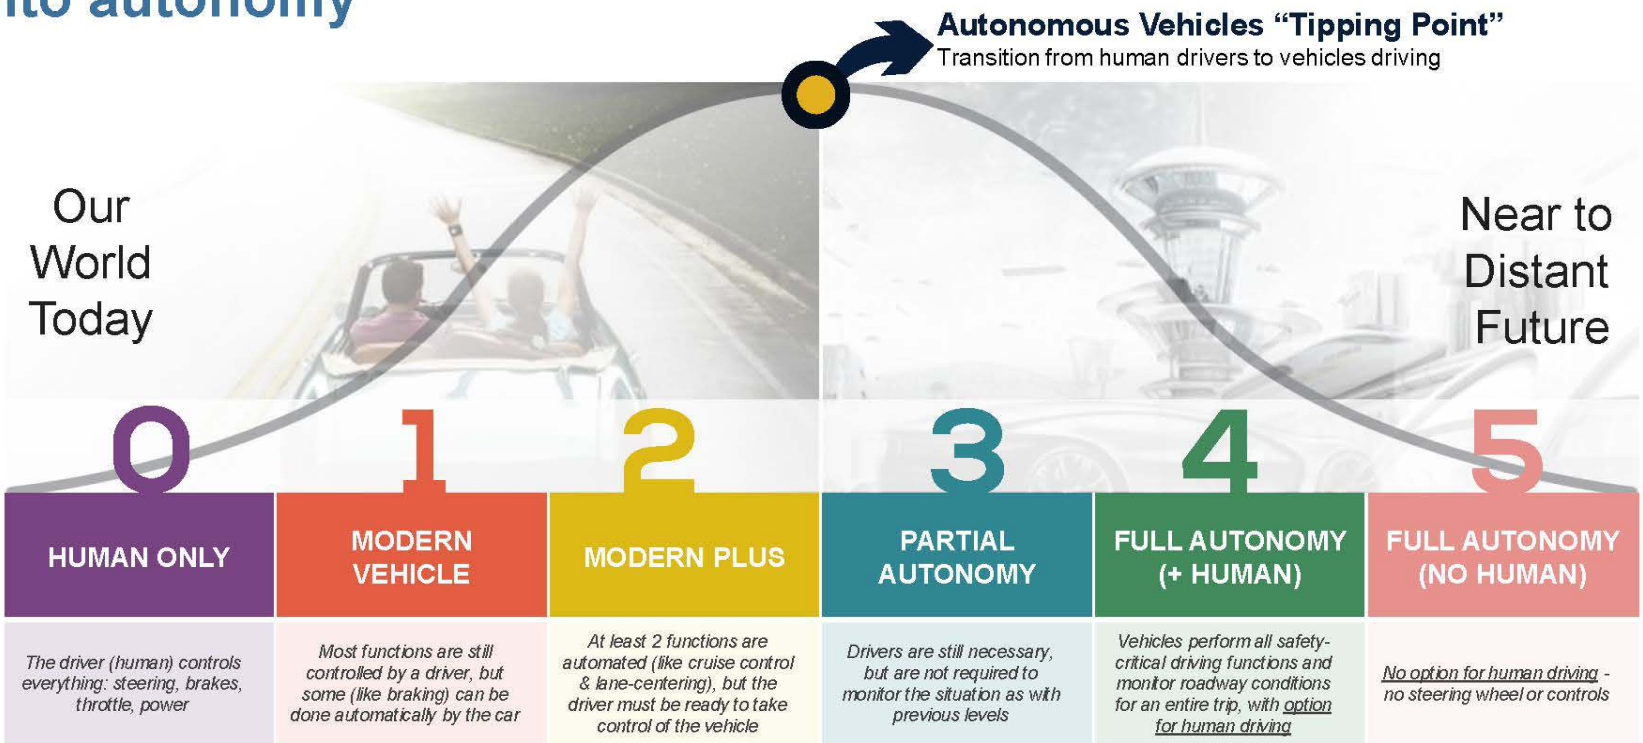

Kelley Blue Book KBB.COM The Trusted Resource

http://www.techrepublic.com/article/autonomous-driving-levels-0-to-5-understanding-the-differences/

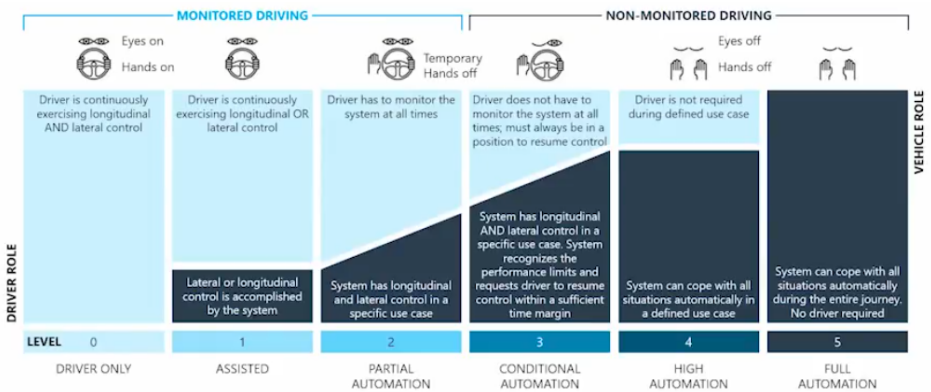

L'autonomie 2, 3, 4 et 5 chez Mobileye

https://www.mobileye.com/ces-2024/

Level 2/3 Autonomous vehicles for sale

100 000 $

Tesla Model X

vs.

Audi A8

Motivations

- Today:
  - driver comfort (12,000 €-> 7,500 €)
- Tomorrow
  - save lives (safety)
  - environmental issues

https://https://www.vehie.com/c/audi-a3-2019-sedan--tesla-model+s-2018

Level 2/3 = ADAS Ratings

Consumer Reports' for major Advanced Driver Assistance Systems (2020)

SYSTEM NAME	SCORE	CAPAB. & PERF.	KEEPING DRIVER ENGAGED	EASE OF USE	CLEAR WHEN SAFE TO USE	UNRE-SPONSIVE DRIVER
Comma Two Open Pilot	78	8	9	8	6	8
Cadillac Super Cruise	69	8	7	3	8	9
Tesla Autopilot	57	9	3	7	2	6
Ford/Lincoln Co-Pilot 360	52	8	4	3	4	5
Audi Driver Assistance Plus	48	8	3	3	2	6
Mercedes-Benz Driver Assistance	46	6	4	4	2	5
Subaru Eyesight	46	7	4	3	4	5
Hyundai Smart Sense, **Kia** Drive Wise	46	5	4	5	4	4
BMW Active Driving Assistance Pro	44	7	3	3	2	6
Porsche Active Safe	41	4	3	6	2	5
Volvo Pilot Assist	41	6	3	3	2	5
Toyota/Lexus Safety Sense 2.0	40	5	4	2	4	5
Honda/Acura Sensing	40	6	4	2	4	4
Nissan/Infiniti ProPILOT Assist	40	5	3	3	4	7
Volkswagen Driver Assistance	39	4	3	6	2	5
Land Rover Driver Assist	38	4	3	6	2	4
Buick/Chevy Driver Confidence	36	3	3	5	2	6
Mazda i-ACTIVSENSE	27	3	2	5	2	1

Some players:

- OpenPilot (open source 50 k)
- Super Cruise (Cadillac 110 k)
- AutoPilot (Tesla, 2M)
- Mobil Eye (54 M)

https://www.thedrive.com/news/37833/consumer-reports-ranks-this-aftermarket-driver-assistance-kit-above-tesla-autopilot-cadillac-super-cruise

Level 4 experences in Rouen, Phoenix, 13 cities in China...

Waymo's cars (Google) hit the 10 million-mile milestone on public roads

New uses

- public transportation (last kilometer)
- isolated people
- autonomous ride services (taxi)
- ...

pole-moveo.org/actualites/rouen-normandy-autonomous-lab-la-metropole-rouen-normandie-teste-un-service-de-mobilite-autonome/
https://waymo.com/

Autonomous vehicle performance ranking

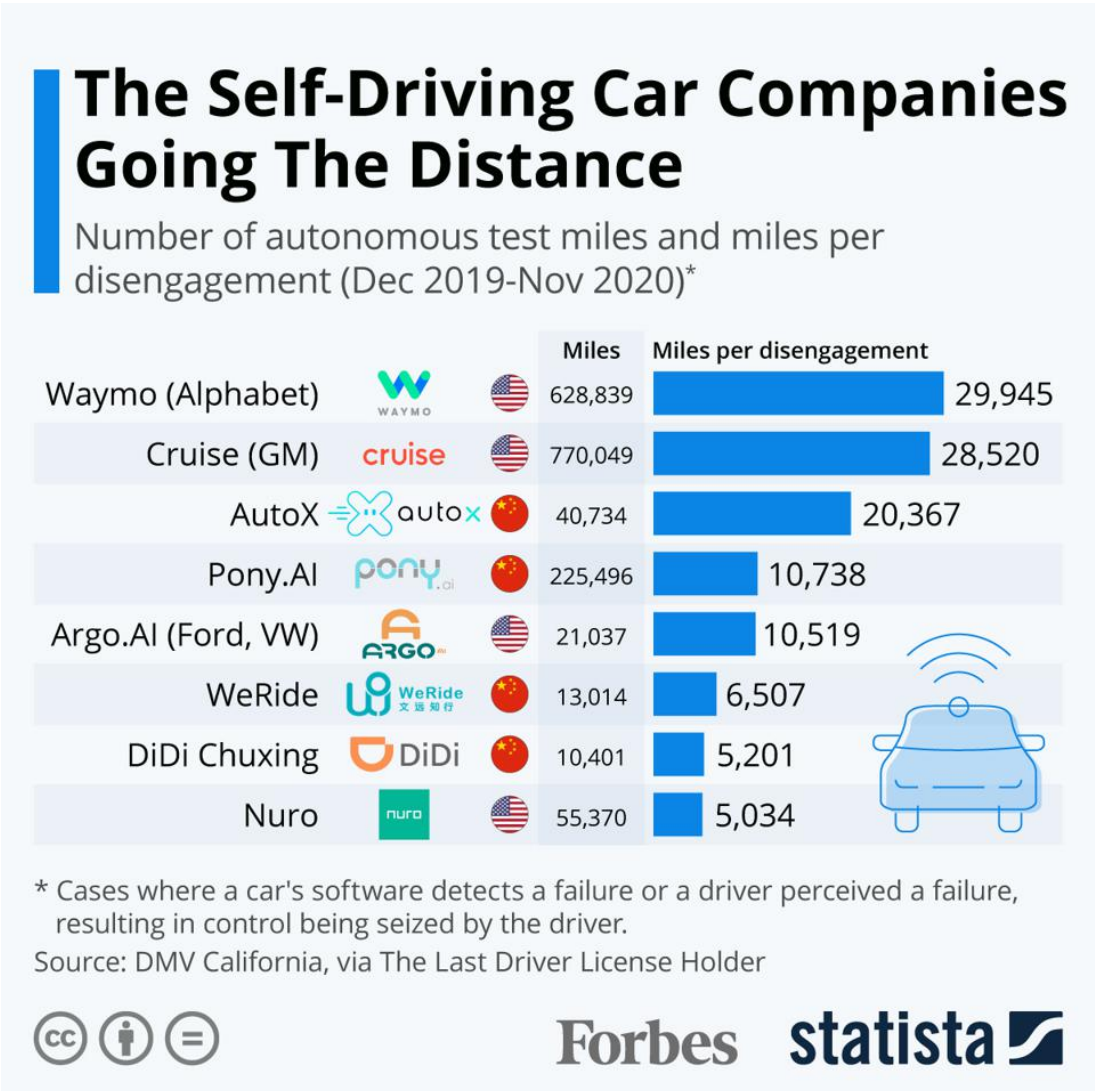

Some player:

- Waymo (Google)
- Cruise (GM)
- Apollo (Baidu)
- ...

Related initiatives:

- La stratégie nationale de développement de la mobilité routière automatisée
- L3 Pilot (European project)
- ...

forbes.com/sites/niallmccarthy/2021/02/15/the-self-driving-car-companies-going-the-distance-infographic

Two kind of AI systems for cars

1. Driver assistance — Driver is responsible
   - Level 2/3 autonomy
   - Specific intelligence
   - it works: how many seconds for take-over?

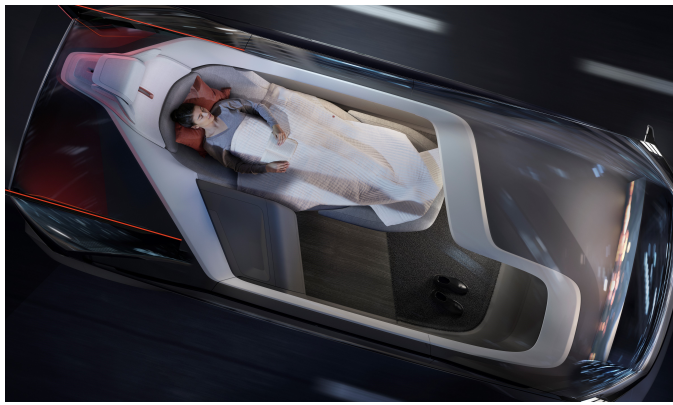

2. Full Autonomous driving — Car is responsible
   - Level 4/5 autonomy
   - Generic Intelligence
   - Experience level: it doesn't scale yet!

Lex Fridman long term vision

When will we have more than 10,000 Full Autonomous cars?

Road map

1. A very brief history of autonomous vehicles
2. How has this happened? (Deep Learninng)
3. Data to train the deep network
4. How Artificial Intelligence will change the Auto
5. Conclusion

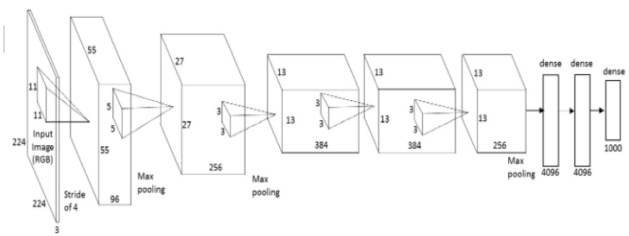

ImageNet results: from 50% to 91%

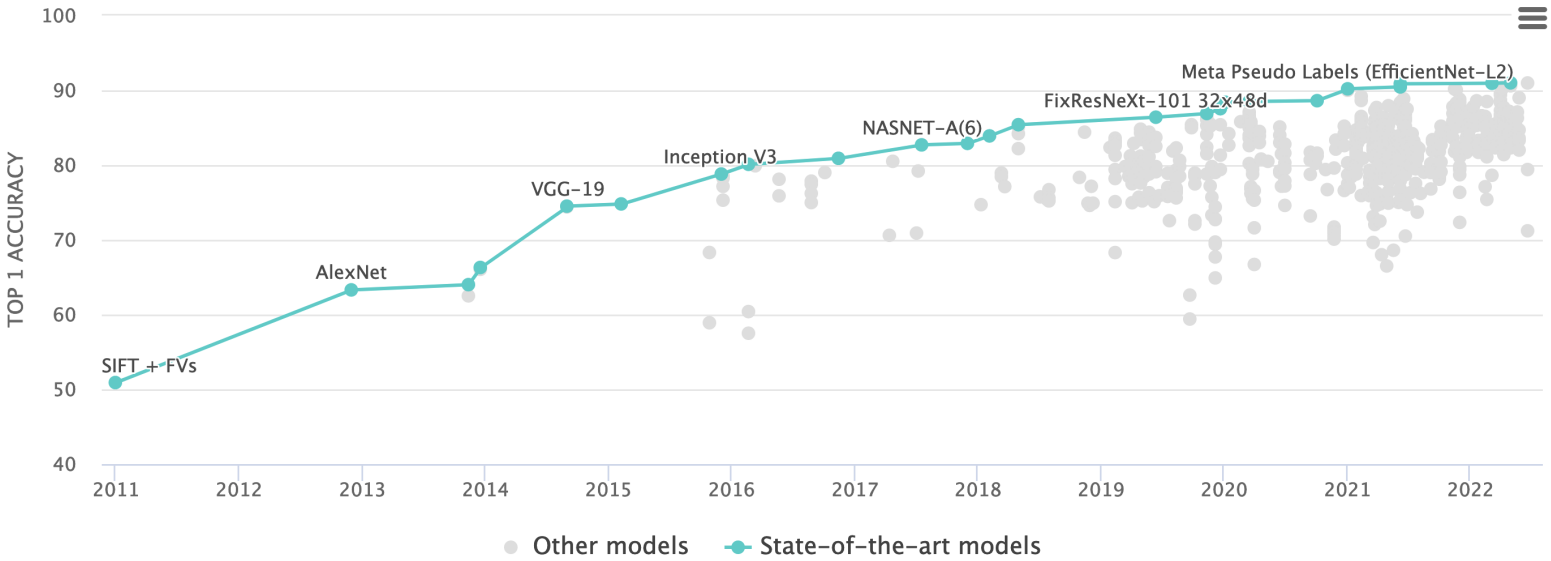

2012 Alex Net

2014 VGG

2015 GoogLeNet / Inception

2016 Residual Network

2018 NAS Network

2020 EfficientNet (Transformers)

2022 CoCa (Contrastive Captioners = Image-Text Foundation Models)

https://paperswithcode.com/sota/image-classification-on-imagenet

Detection, tracking and recognition of traffic signs (2011-13)

Recognition German Traffic Sign Recognition Benchmark (GTSRB) data set, containing 51839 labelled images of real-world traffic signs.

Detection The German Traffic Sign Detection Benchmark is a single-image detection assessment 900 images (600 for training and 300 for test)

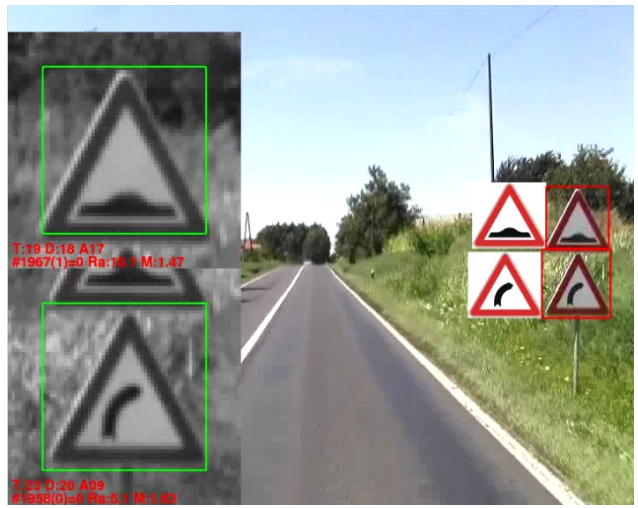

and the winner is

$\rightarrow$ Deep learning gives very good results on both tasks

Open Pilot: 2200 $

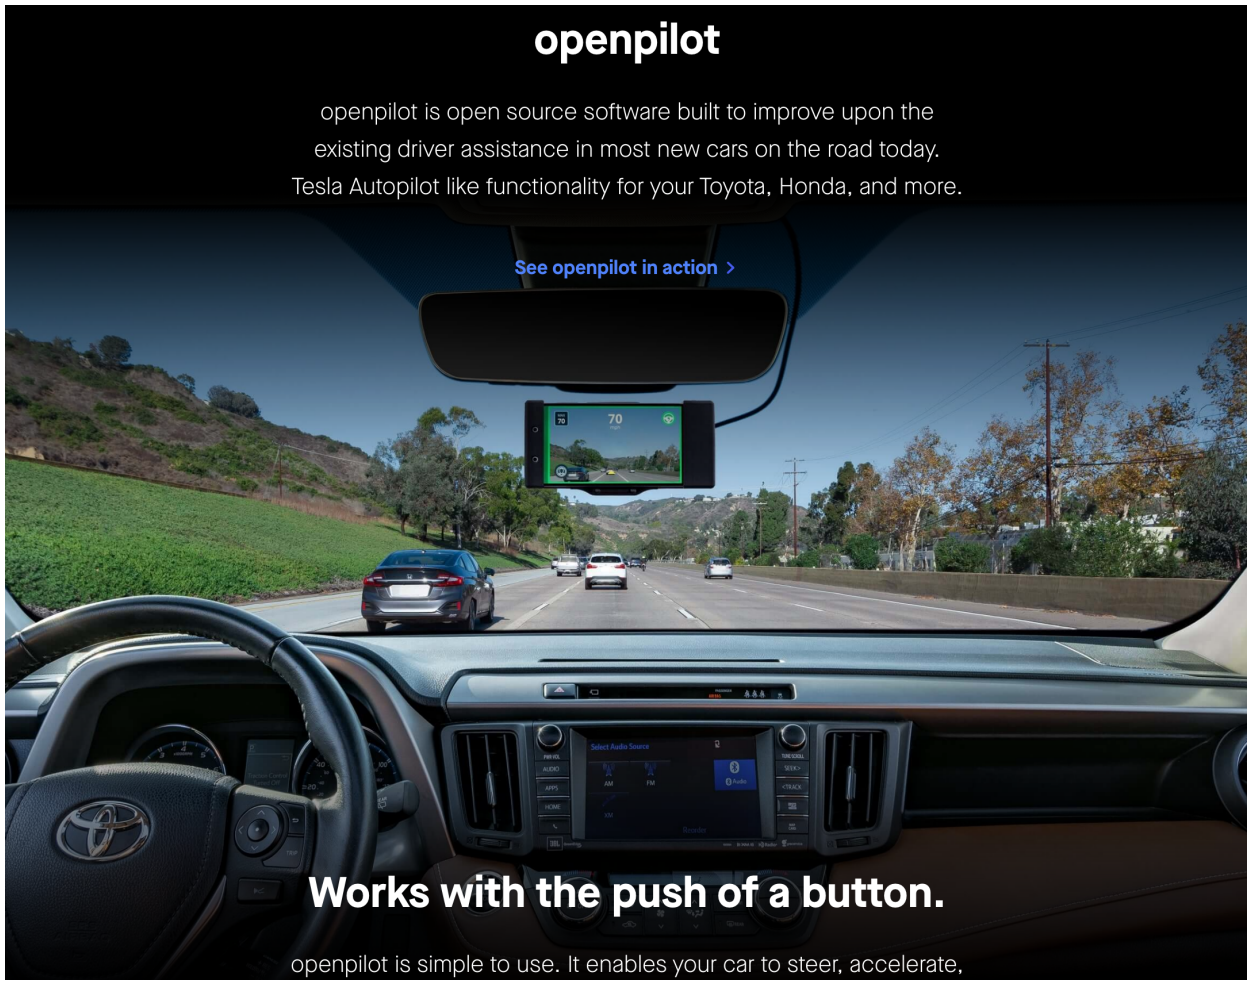

openpilot is the Android

https://github.com/commaai/openpilot

THIS IS ALPHA QUALITY SOFTWARE FOR RESEARCH PURPOSES ONLY. THIS IS NOT A PRODUCT. YOU ARE RESPONSIBLE FOR COMPLYING WITH LOCAL LAWS AND REGULATIONS.

Openpilot AI features

Two AI

- Diving agent
  - Automated lane-centering
  - Adaptive cruise control
    OpenStreetMap inside
  - Assisted lane change

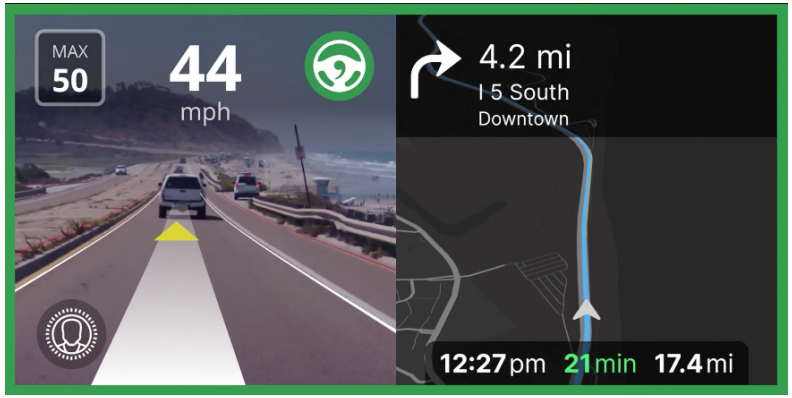

- Driver monitoring system (DMS)
  - Safety concerns

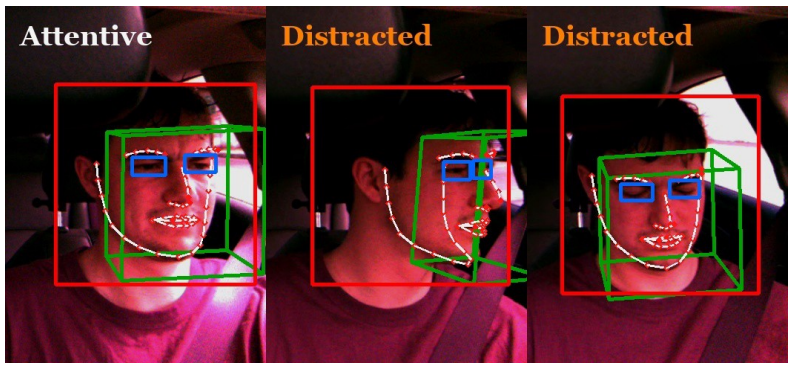

software update

https://comma-ai.medium.com/towards-a-superhuman-driving-agent-1f7391e2e8ec

Openpilot 's driver monitoring system (DMS)

Three components

- Face localization
  - openCV -> cropping
- Feature generation
  - EfficentNet b0 architecture
  - Fine tuning
- Decision module
  - Treshold based decision

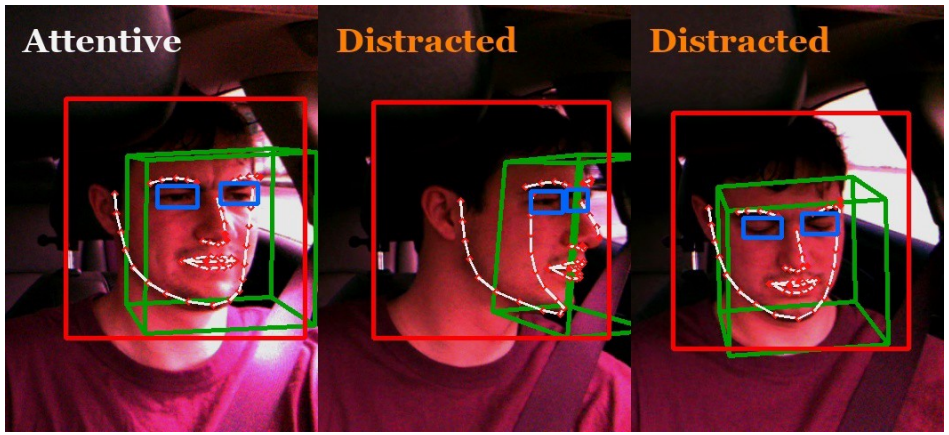

https://github.com/commaai/openpilot

Architecture of the feature generator of openpilot's DMS

- Input: YUV 420 (6 channels)
  - EfficentNet b0 architecture
  - Tan et. al. (Google), ICML 2019
- Output: 45-features (03/22)
  - Face position (12 values)
  - Eyes positions (8 values)
  - sunglasses
  - visible face probability
  - blinking
  - ...
- Training data: fine tuning
  - pytorch inside
  - Qualcomm Snapdragon 845

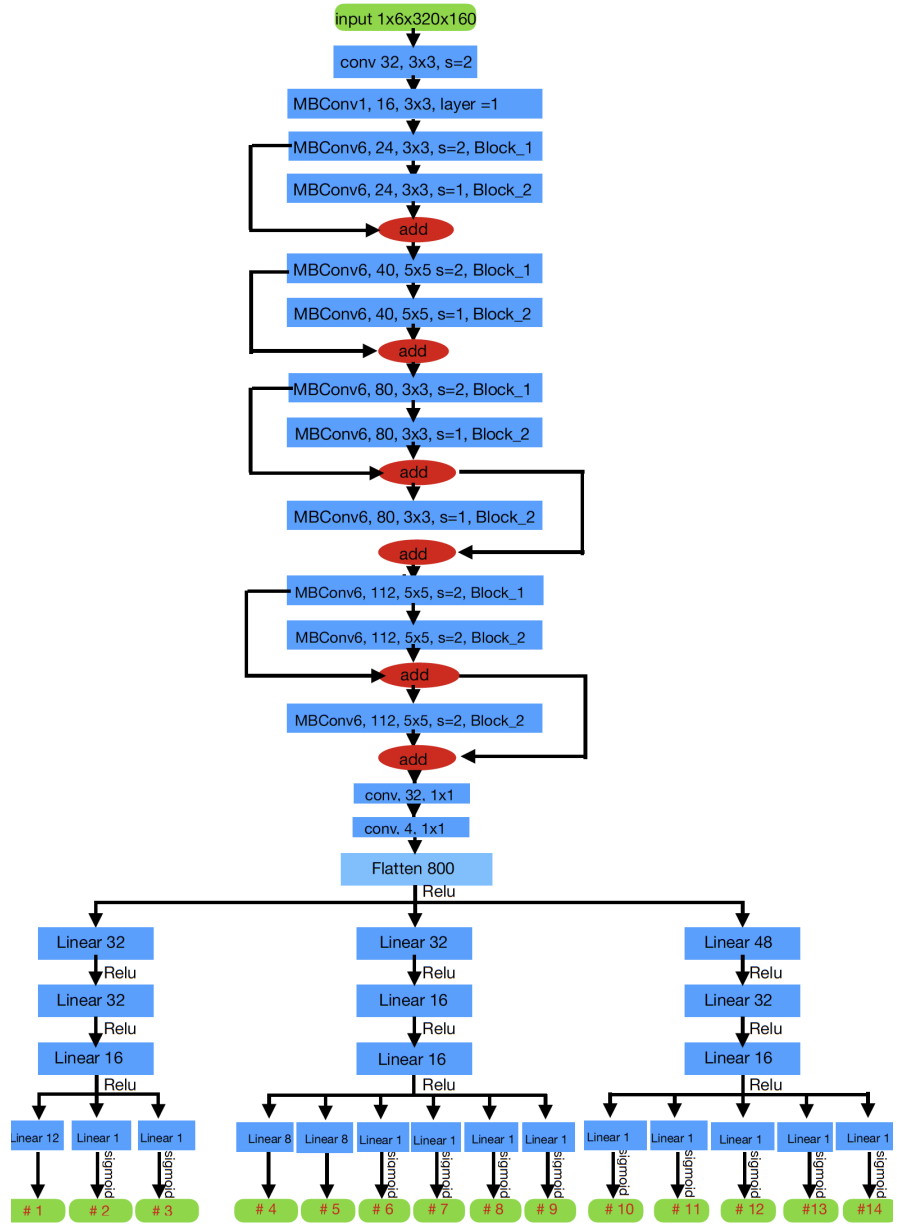

Openpilot's components

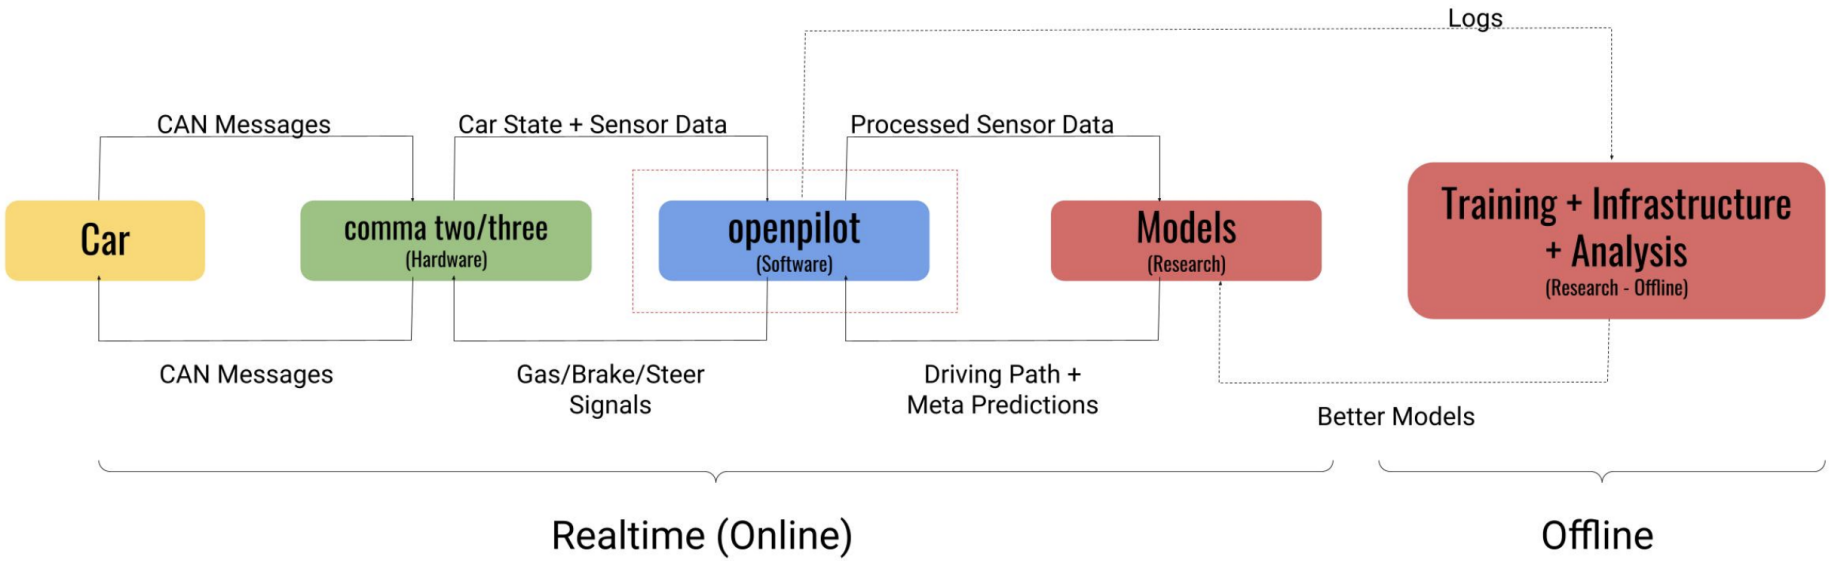

AI inside: uses data

(Deep) learning based programming

Tesla's autopilot components

- Driving agent
  - Automatic lane change
  - Adaptive cruise control
  - Autosteer
  - Navigate on Autopilot (Freeway)
  - Traffic Light and Stop Sign Control
  - ...
  - FSD (limited-access Beta)
- Parking Summon

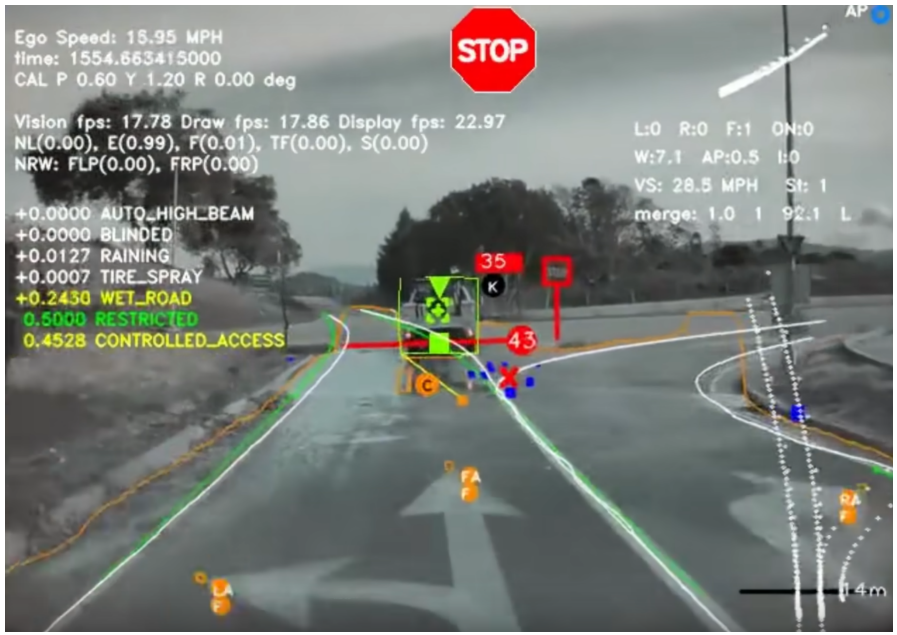

- Driver monitoring system (DMS)

Summarizing the driving agent architecture

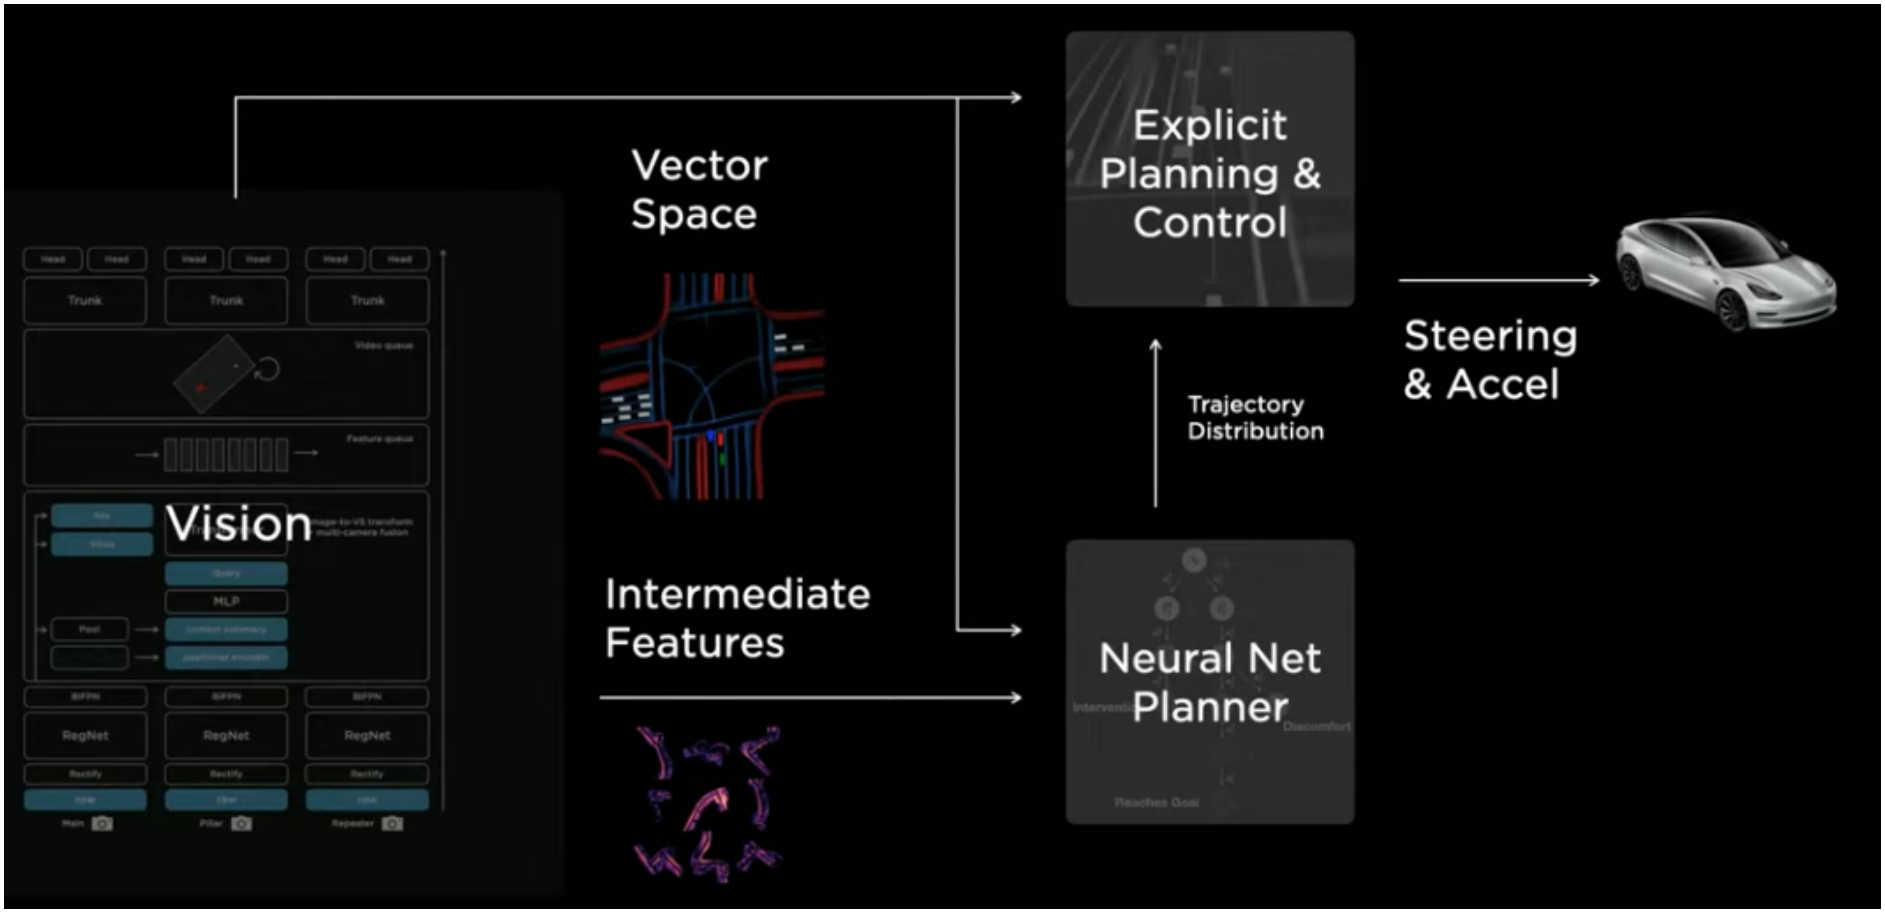

Two AI components = two deep networks

- perception module
- decision module (planner) using deep reinforcement learning

Tesla's autopilot perception module

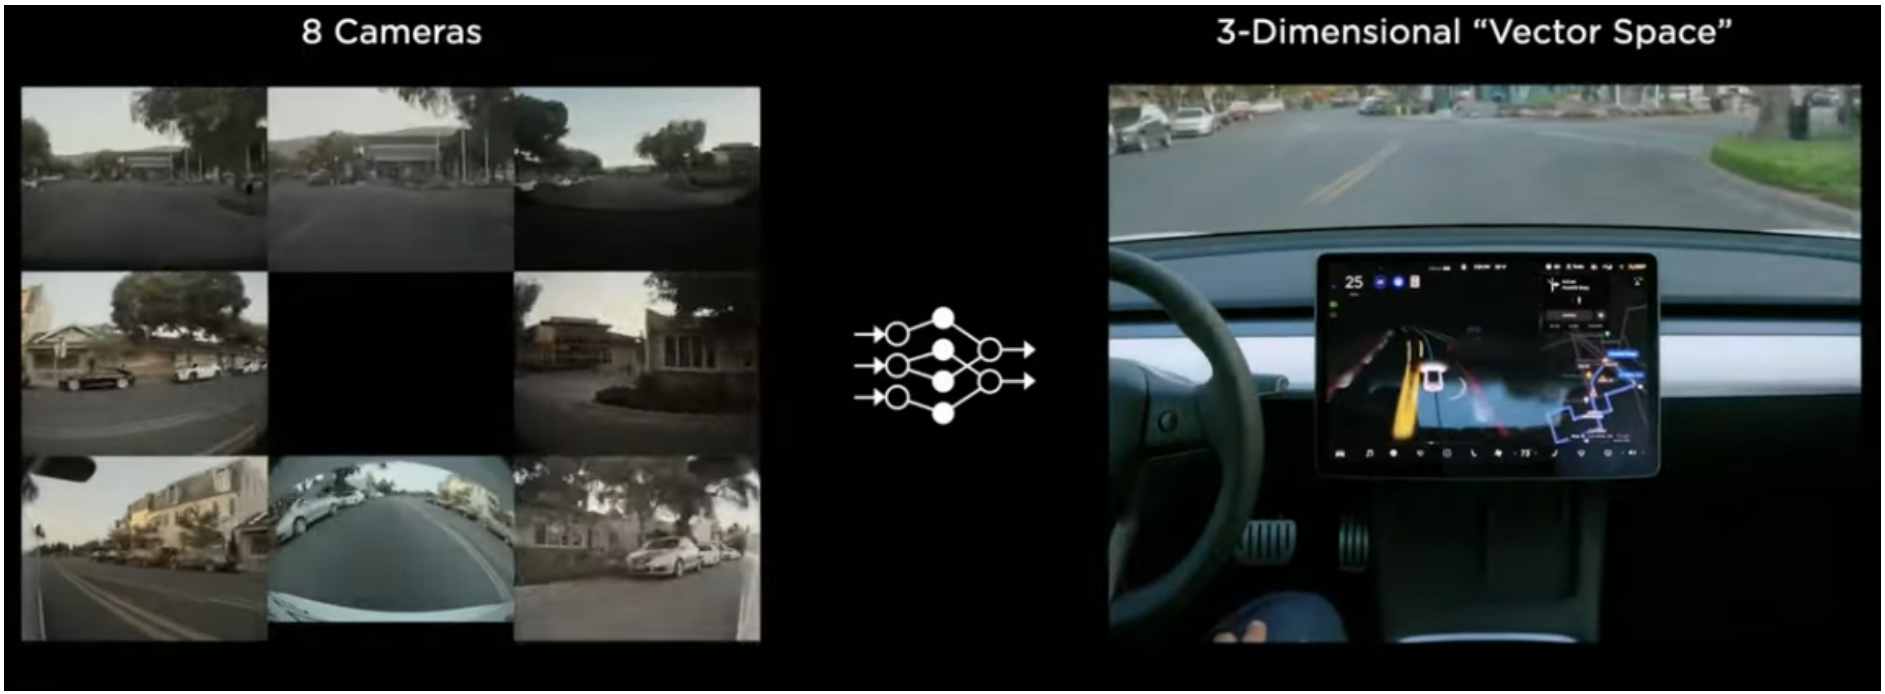

- input: 8 cameras
- output: 640×460 3D map of the surroundings

Perception is scene understanding

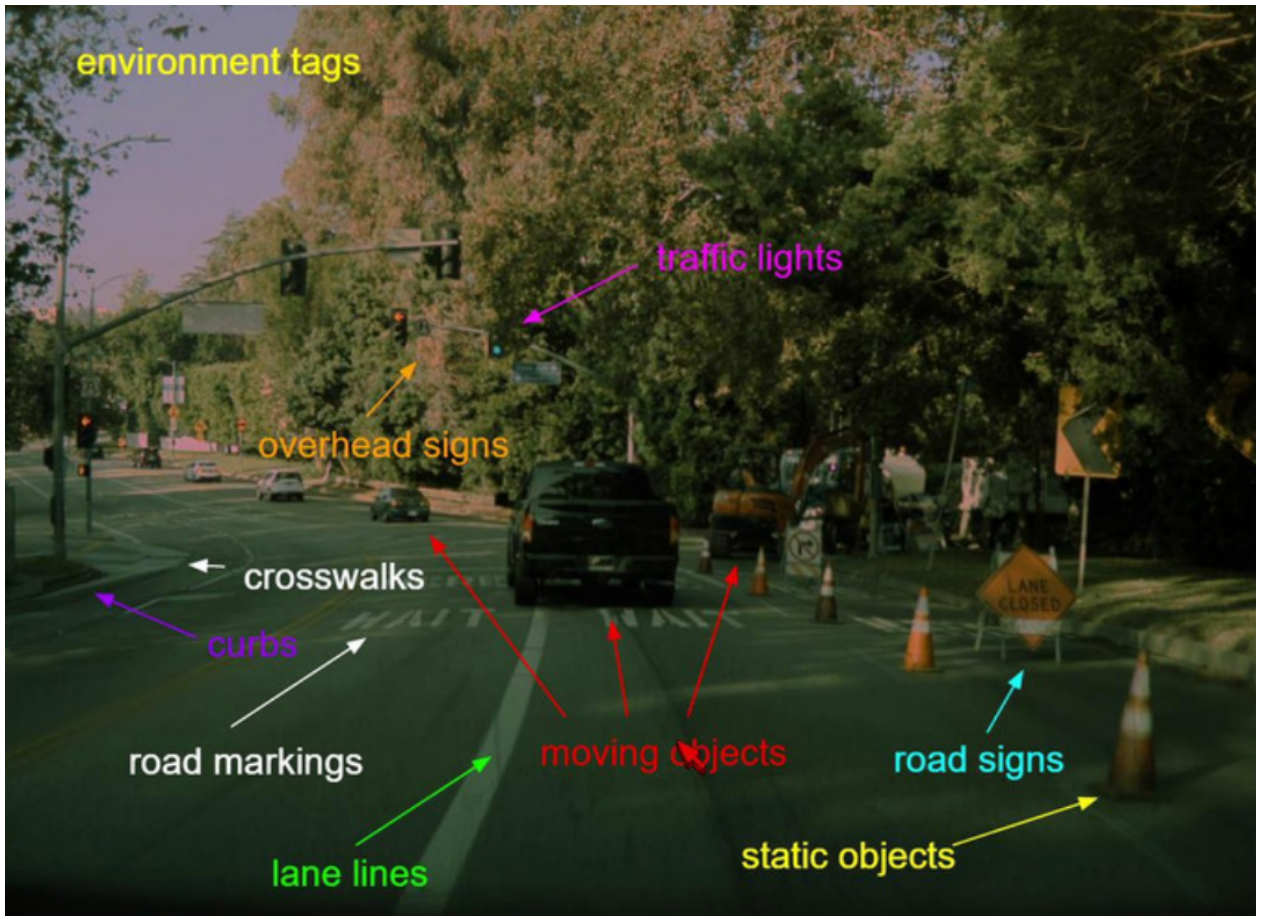

Scene understanding is

Multi-task learning

Andrej Karpathy, Multi-Task Learning in the Wilderness, ICML 2019
https://slideslive.com/38917690/multitask-learning-in-the-wilderness

The 5 components of Tesla's perception module

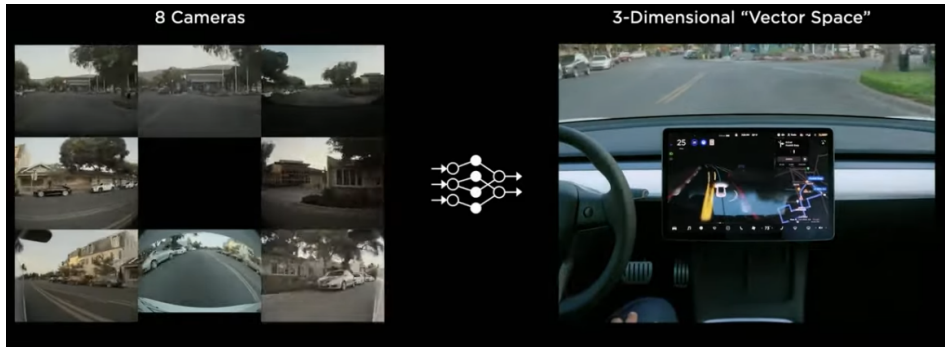

input: 8 cameras

1. feature generator: backbone
2. multi scale feature fusion
3. multi camera fusion
4. time filtering
5. multi task decision module per pixel on the output map (one per task)
   - item location (cars, pedestrian...)
   - traffic signs (Stop sign, traffic light...
   - lane prediction
   - ...

output: 640×460 3D map of the surroundings

1. and 2. Tesla's feature extractor

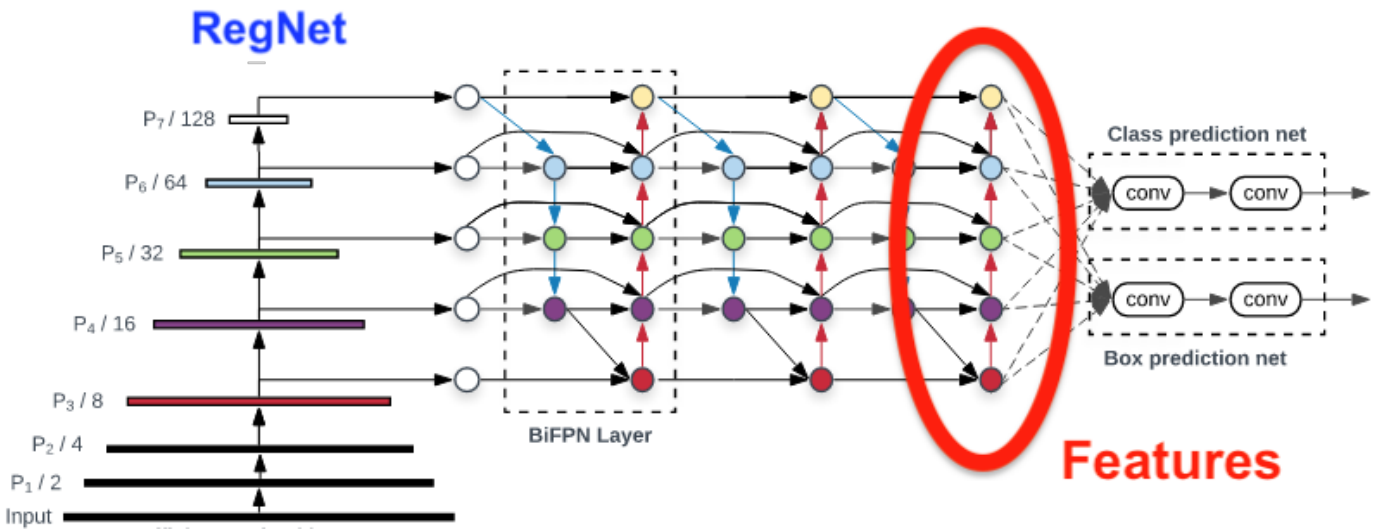

image 1280× 960 12 bits

Multiscale latency/accuracy trade off

- ResNet50 (2020), RegNet (2021) - different resolution/different scale
- EfficientDet - bi-directional feature pyramid net (BiFPN)

Andrej Karpathy , Tesla FULL self driving explained by an engineer, Tesla's AI day, Aug 20, 2021

1. and 2. Tesla's feature extractor

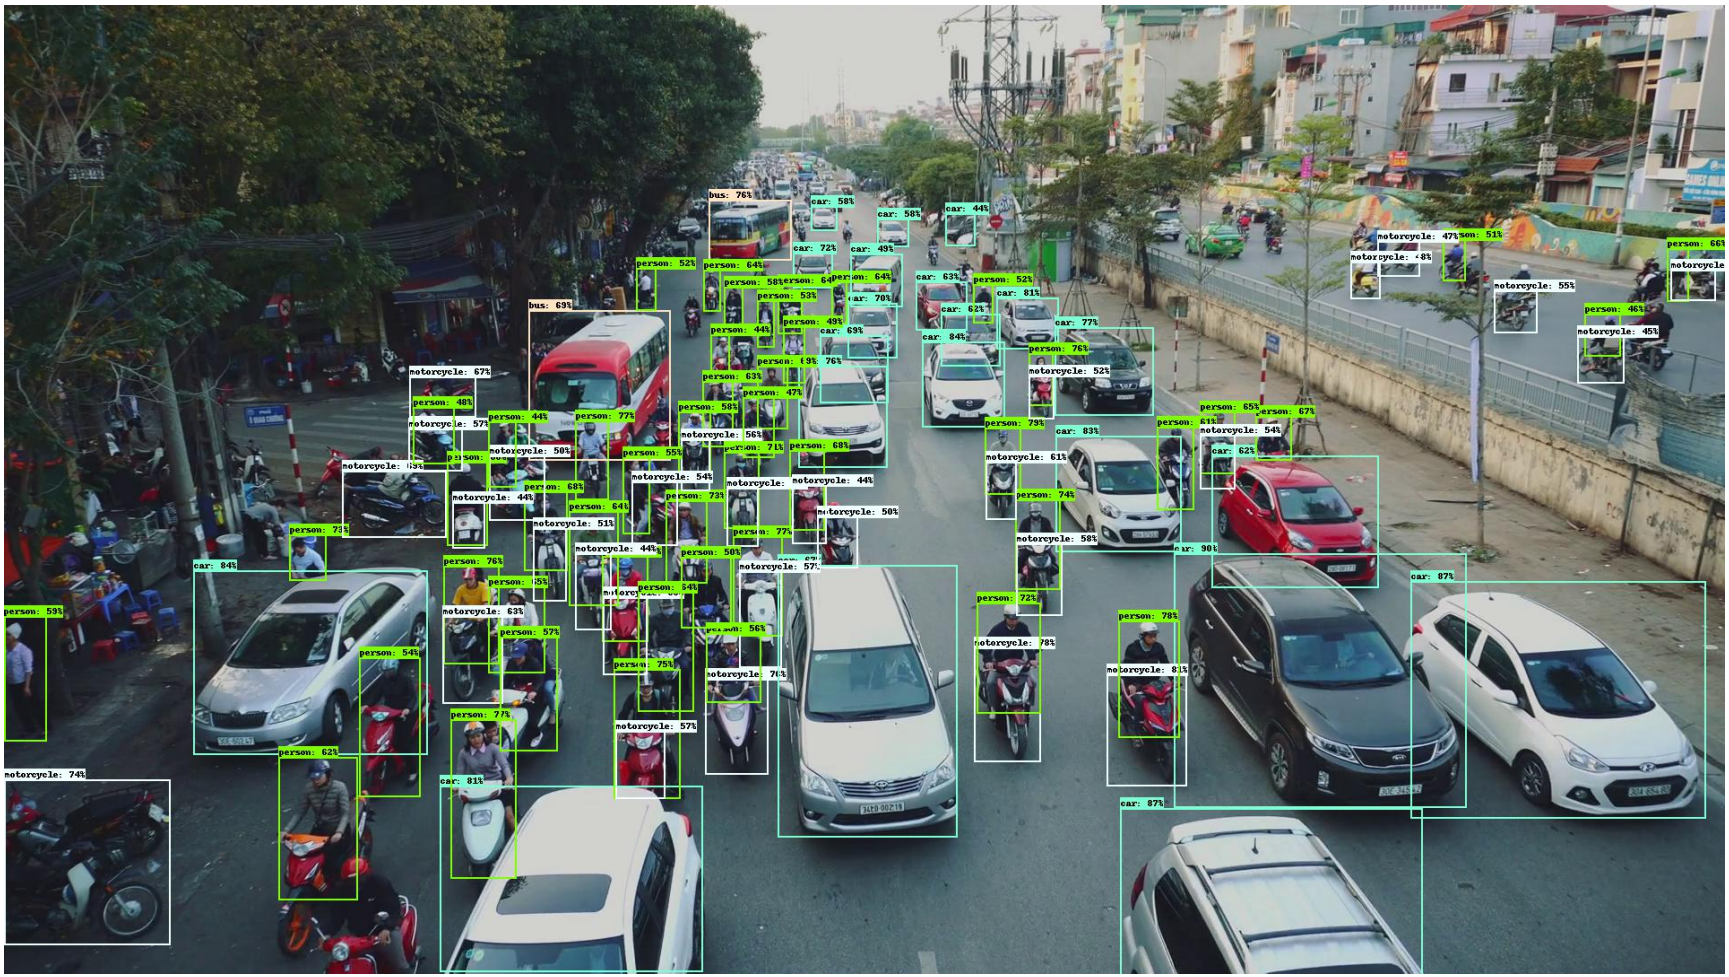

Mingxing Tan et al., EfficientDet: Scalable and Efficient Object Detection, CVPR 2020

Ok but Tesla has got 8 cameras !

3. Sensor fusion

To deal with uncertainties

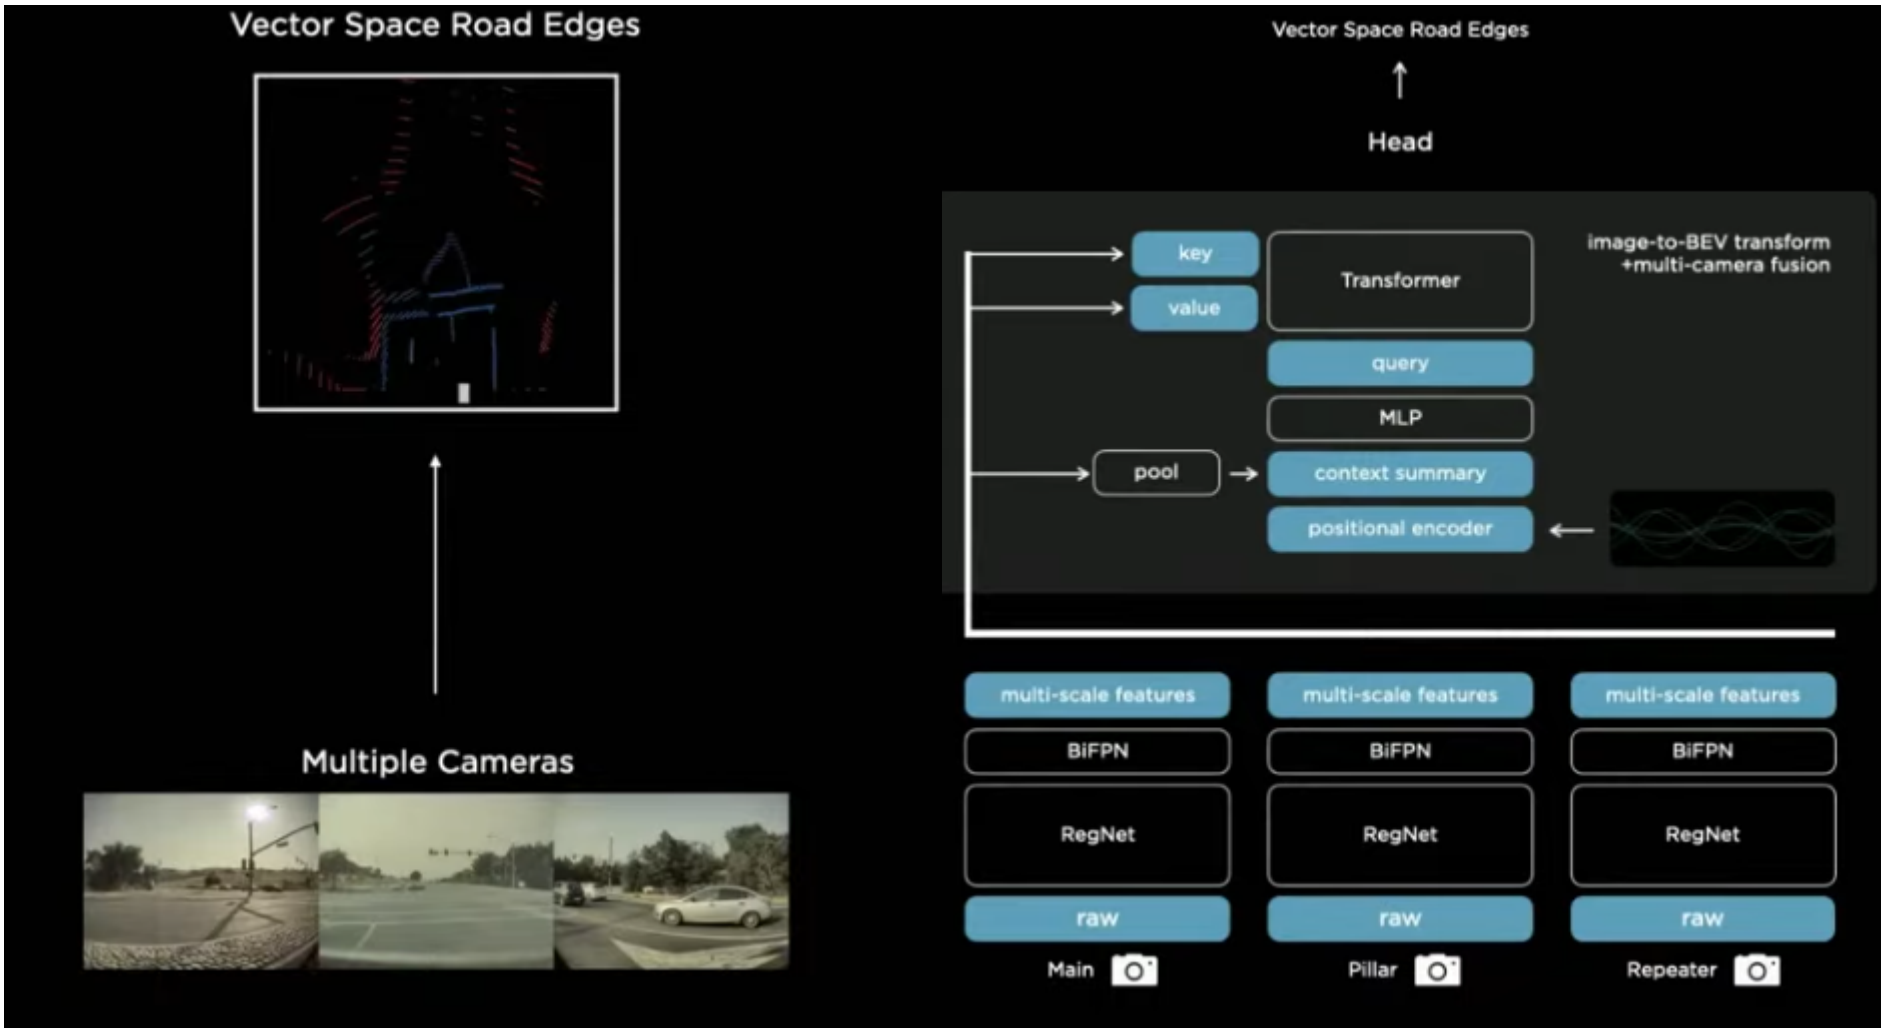

Nicolas Carion et al. "End-to-end object detection with transformers." ECCV 2020.

3. Sensor fusion results

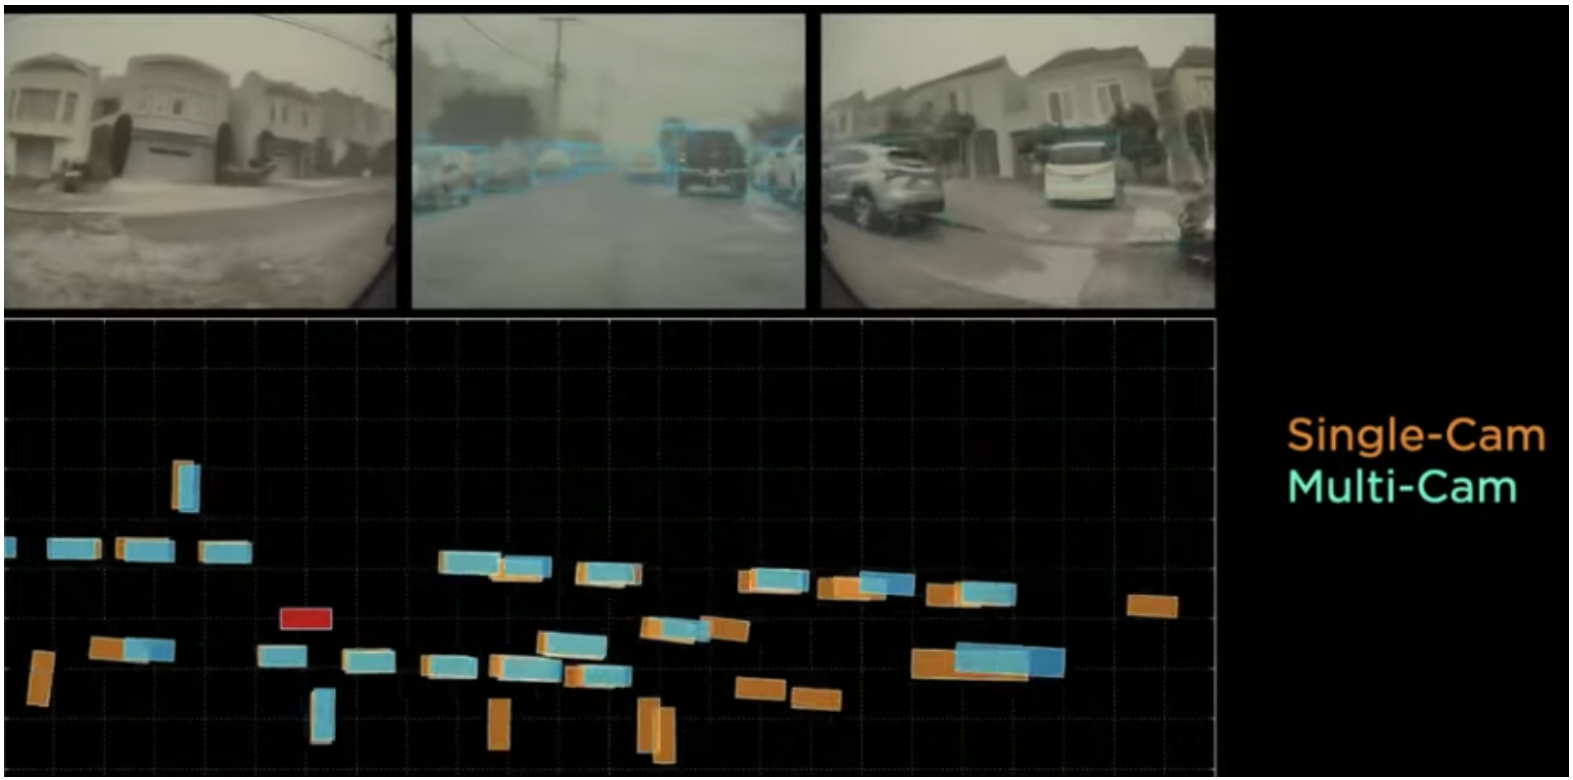

4. Time filtering

to deal with time (occlusion, past traffic signs...)

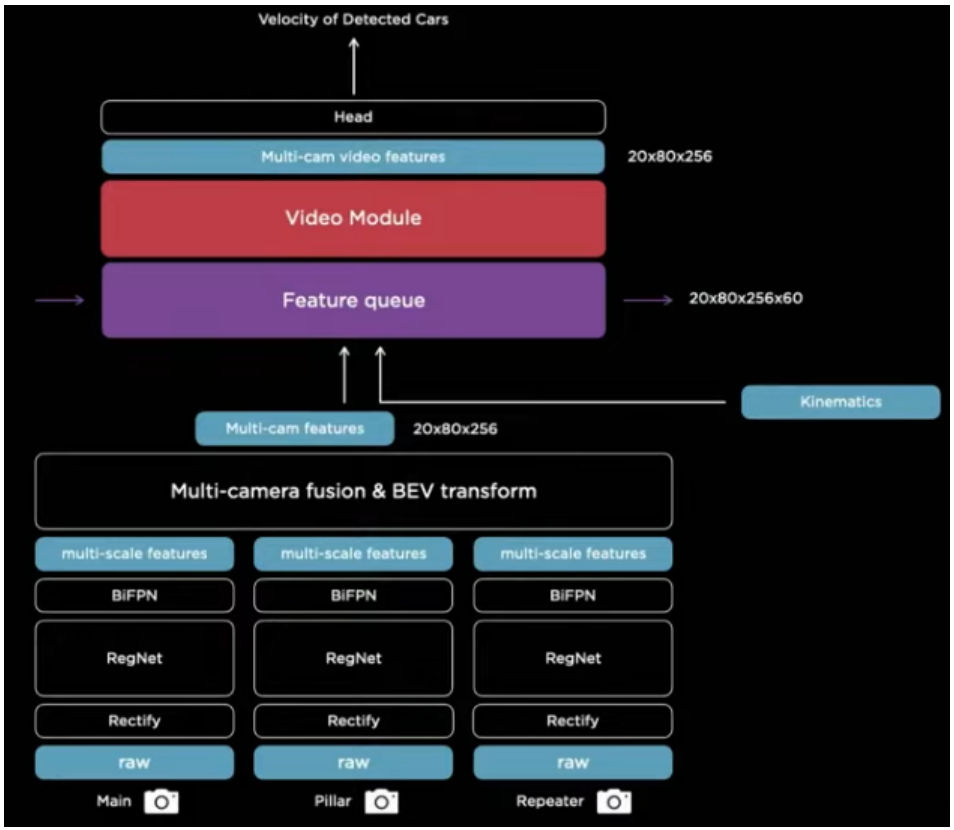

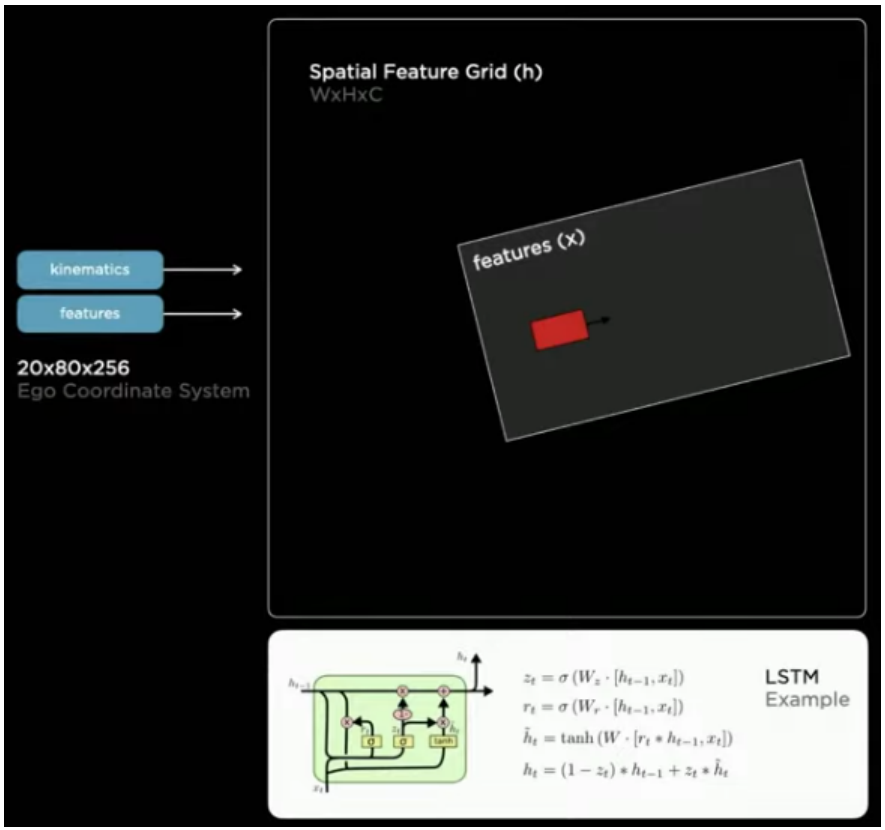

Video module preforming spatio temporal filteing

- 36 frames per second
- spatio temporal LSTM (Liu et al, ECCV 2016)

5. Decision modules

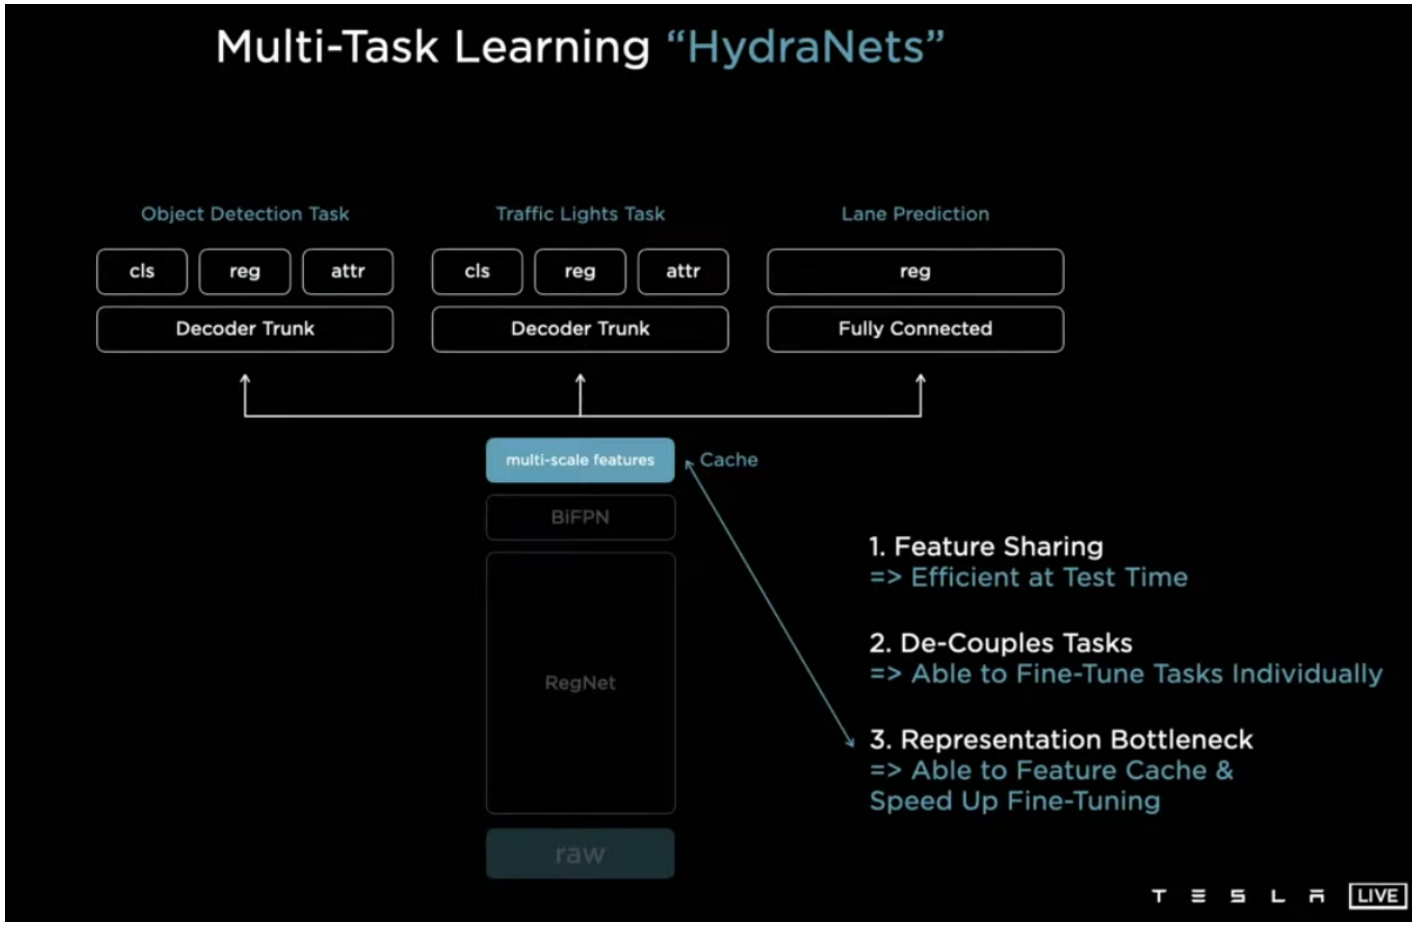

HydraNets, Mullapudi et al, 2018

- Multi task learning
- Specialized shared feature (to reduce inference computing time)

Tesla perception module

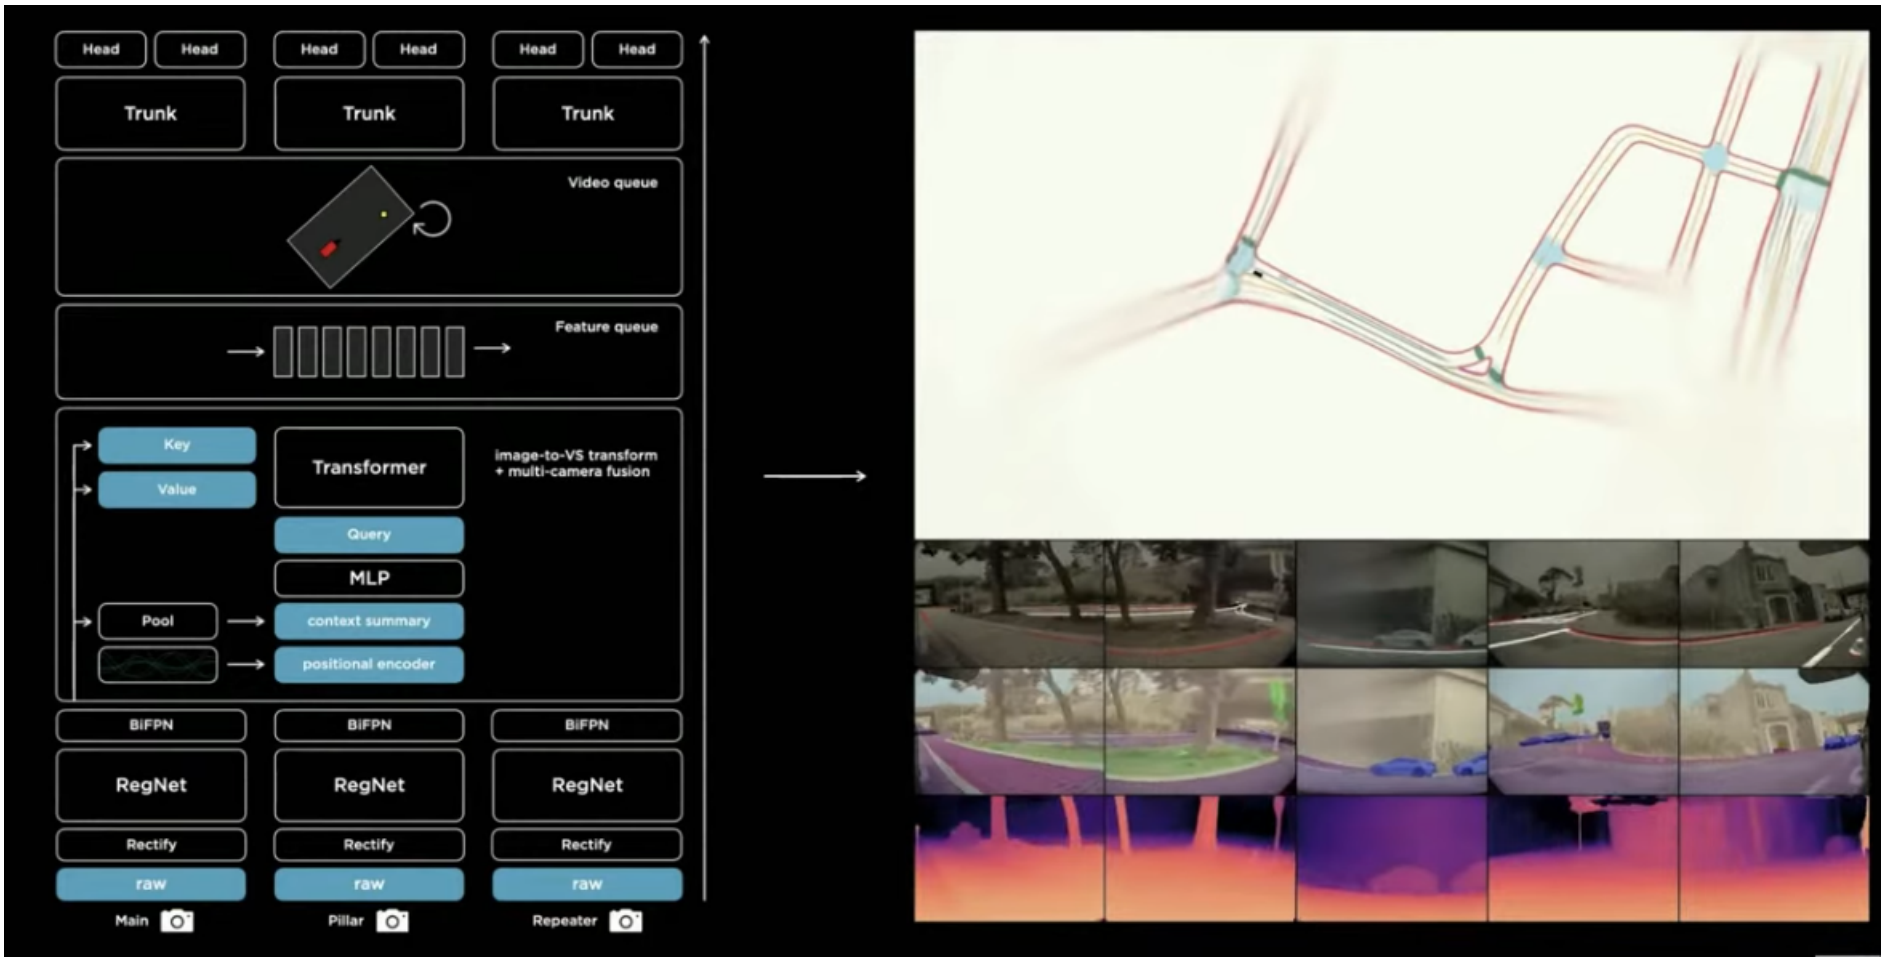

Tesla perception module

1. feature generator: backbone — ResNet50 (2020), RegNet (2021) (from a CVPR 2020 Facebook paper)
2. multi scale feature fusion — EfficientDet (from a 2019 Google paper)
3. multi camera fusion — Transformers (from a 2020 Facebook paper)
4. time filtering — LSTM (recurrent neural network)
5. multi task decision module — Hydranet
   - item location (cars, pedestrian...)
   - traffic signs (Stop sign, traffic light...
   - lane prediction
   - ...

This perception module contains

1. 48 networks, 1,000 outputs tensors, 70,000 GPU h to train
2. performs 40 prediction per second

Perception module at Waymo

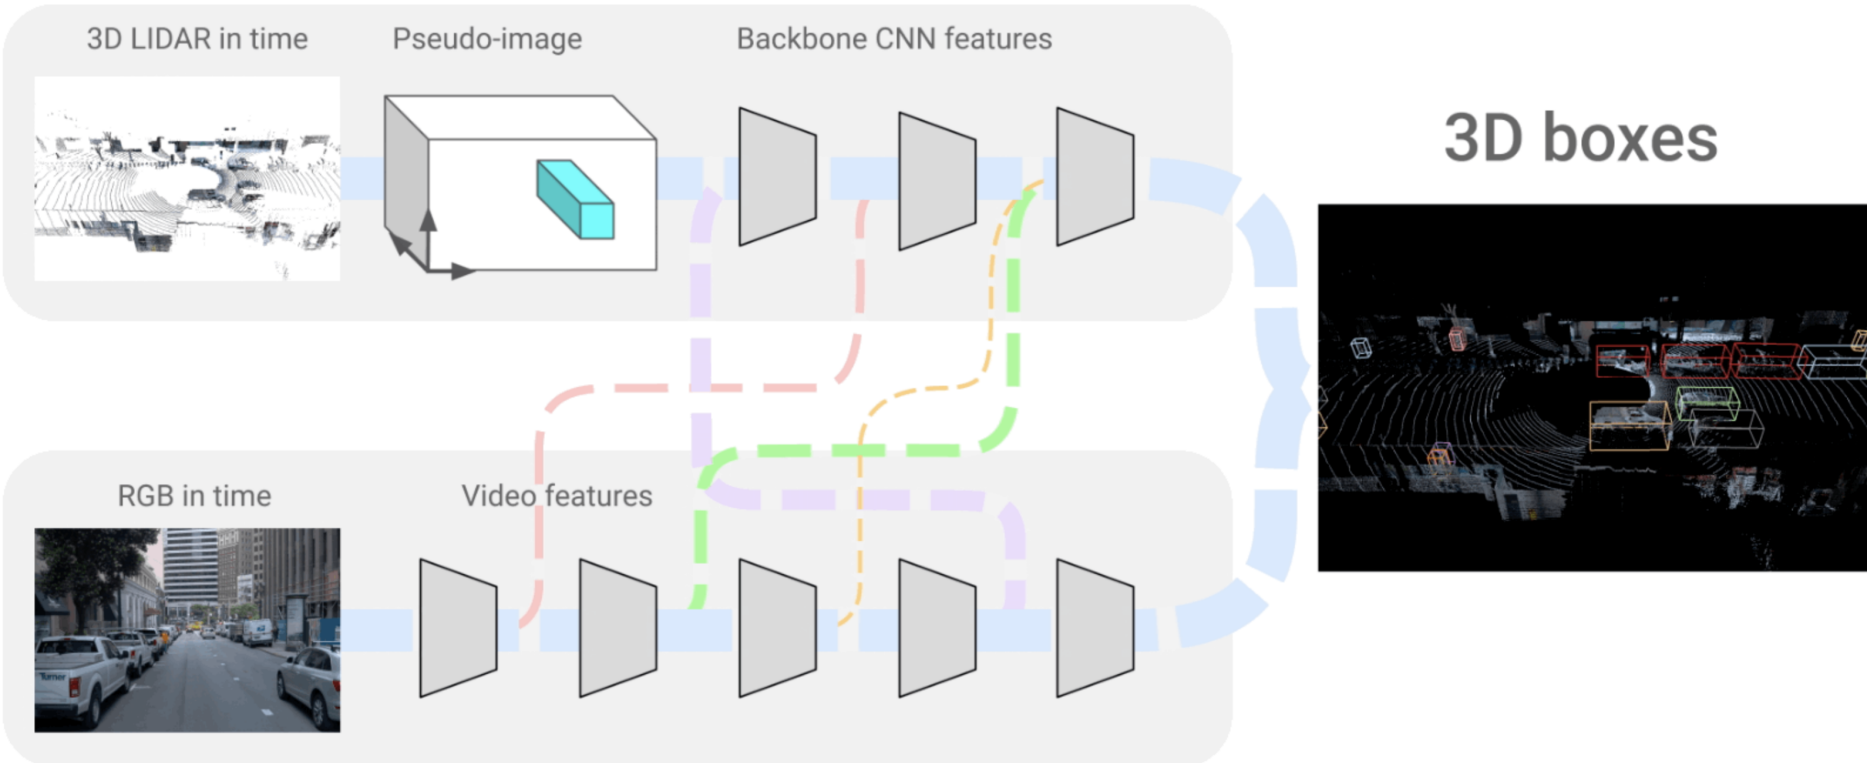

"4D-Net for Learned Multi-Modal Alignment", ICCV 2021
https://ai.googleblog.com/2022/02/4d-net-learning-multi-modal-alignment.html

Waymo's point of view

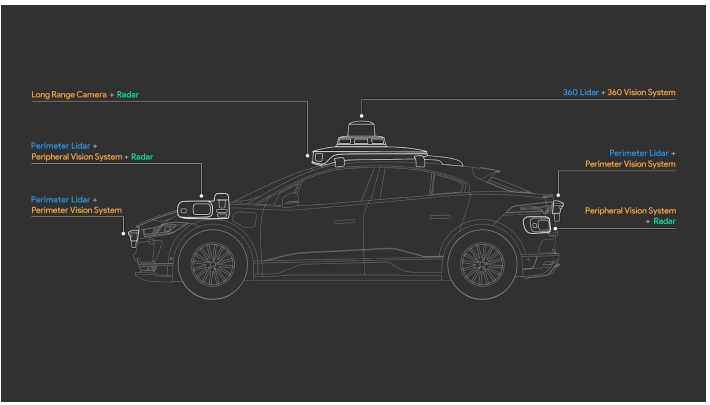

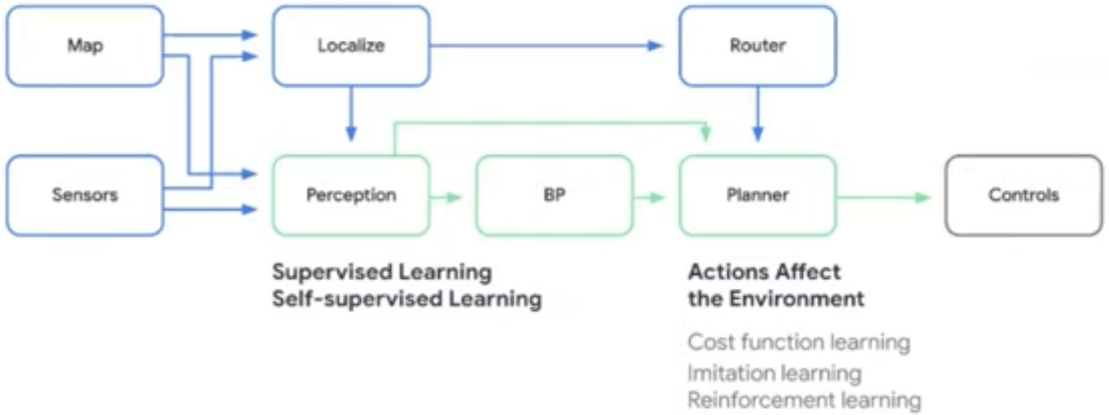

Active research

- Stinet: Spatio-temporal-interactive network for pedestrian detection and trajectory prediction, CVPR 2020
- Vectornet: Encoding hd maps and agent dynamics from vectorized representation, CVPR 2020
- Taskology: Utilizing Task Relations at Scale, CVPR 2021
- ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst, ICML 2019

Decision making using deep reinforcement learning

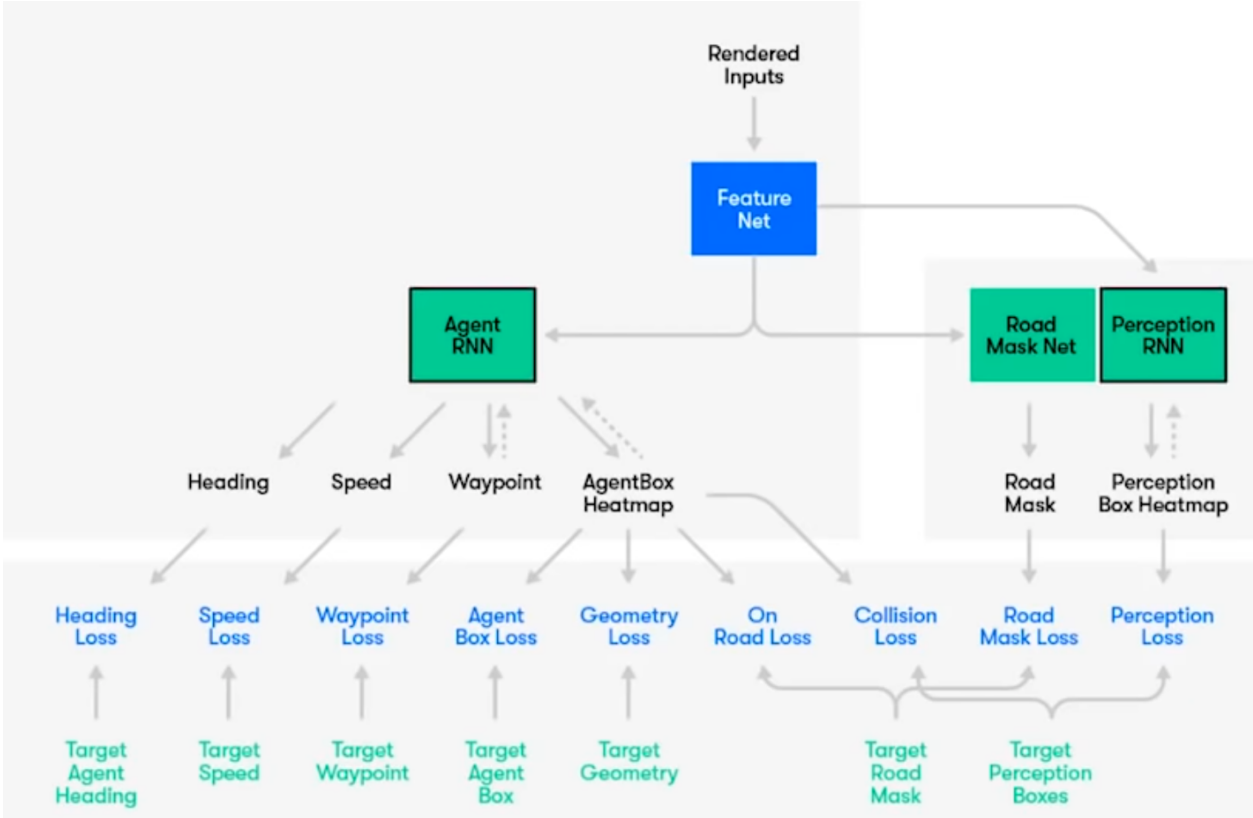

Imitation model providing safety, confort and efficiency

Multi-Task and multi objective learning

Mayank Bansal , ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst, ICML 2019
https://slideslive.com/38917927/chauffeurnet-learning-to-drive-by-imitating-the-best-and-synthesizing-the-worst

Waymo's AutoML

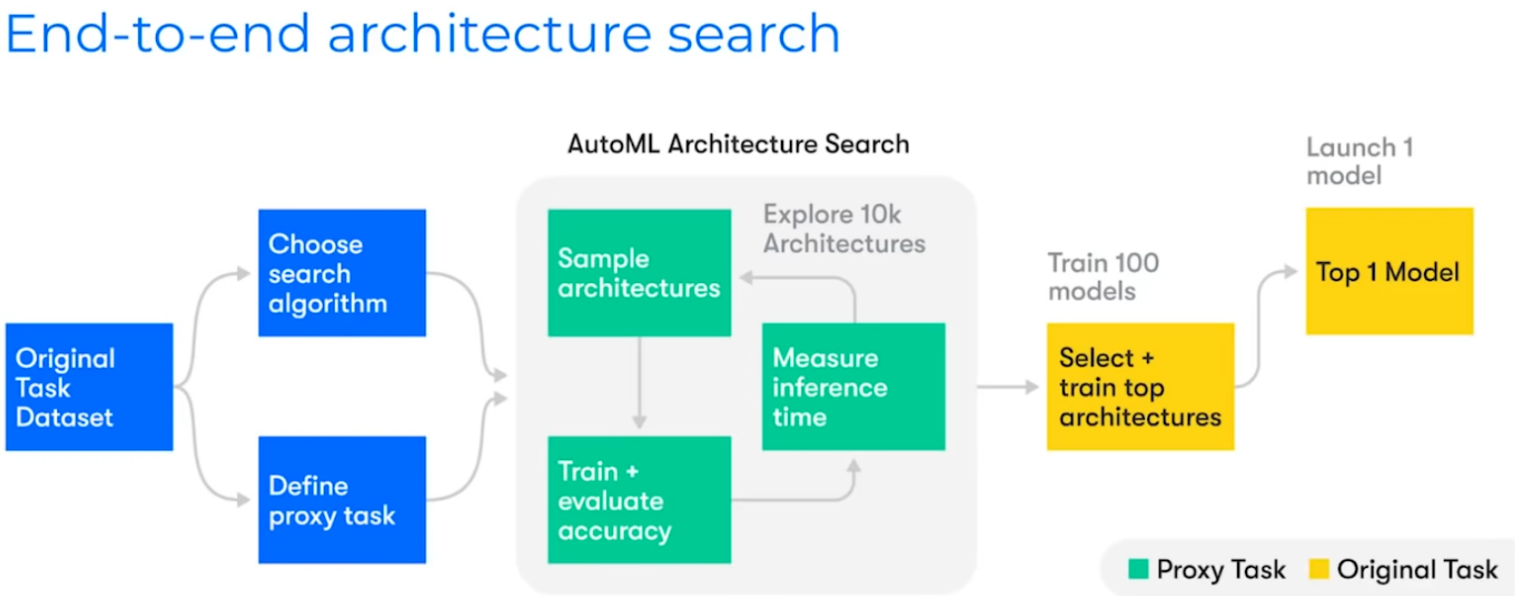

Proxy end-to-end search: Explore thousands of architecture on a scaled-down proxy task, apply the 100 best ones to the original task, validate and deploy the best of the best architectures on the car.

Drago Anguelov (Waymo) - MIT Self-Driving Cars lectures
https://medium.com/waymo/automl-automating-the-design-of-machine-learning-models-for-autonomous-driving-141a5583ec2a

Waymo's AutoML

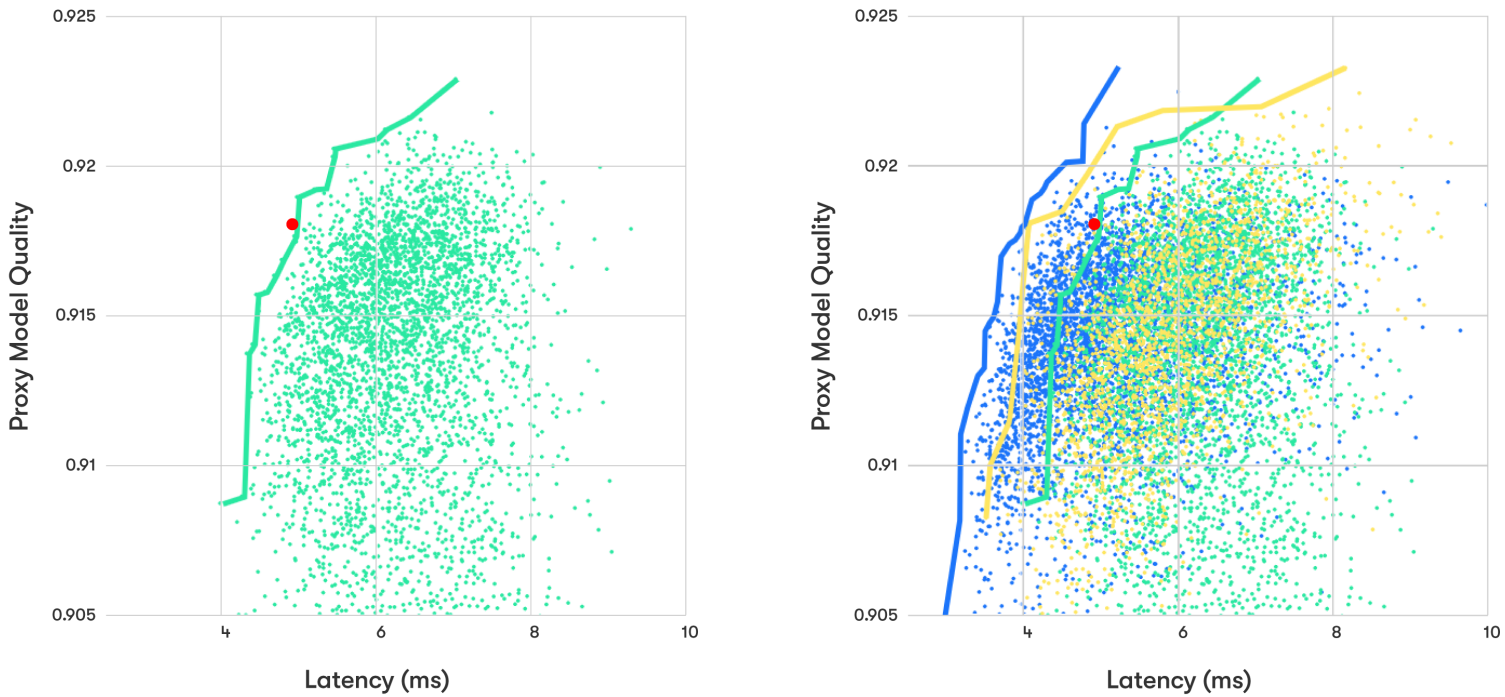

1) The first graph shows about 4,000 architectures discovered with a random search on a simple set of architectures. Each point is an architecture that was trained and evaluated. The solid line marks the best architectures at different inference time constraints. The red dot shows the latency and performance of the net built with transfer learning. In this random search, the nets were not as good as the one from transfer learning. 2) In the second graph, the yellow and blue points show the results of two other search algorithms. The yellow one was a random search on a refined set of architectures. The blue one used reinforcement learning as in [1] and explored more than 6,000 architectures. It yielded the best results. These two additional searches found nets that were significantly better than the net from transfer learning.

AutoML: Automating the design of machine learning models for autonomous driving
https://medium.com/waymo/automl-automating-the-design-of-machine-learning-models-for-autonomous-driving-141a5583e2a

Road map

1. A very brief history of autonomous vehicles
2. How has this happened? (Deep Learninng)
3. Data to train the deep network
4. How Artificial Intelligence will change the Auto
5. Conclusion

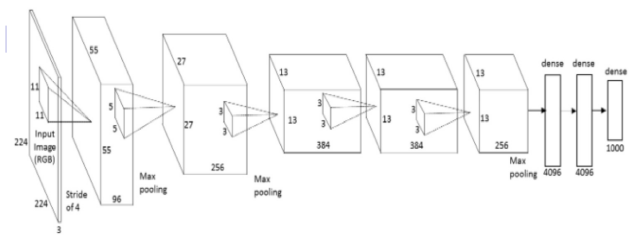

Data: the long tail of situations

Taïwan, june 2020,

Andrej Karpathy - AI for Full-Self Driving at Tesla, Scaled ML, feb 2020,

Improving the autopilot: iterative process

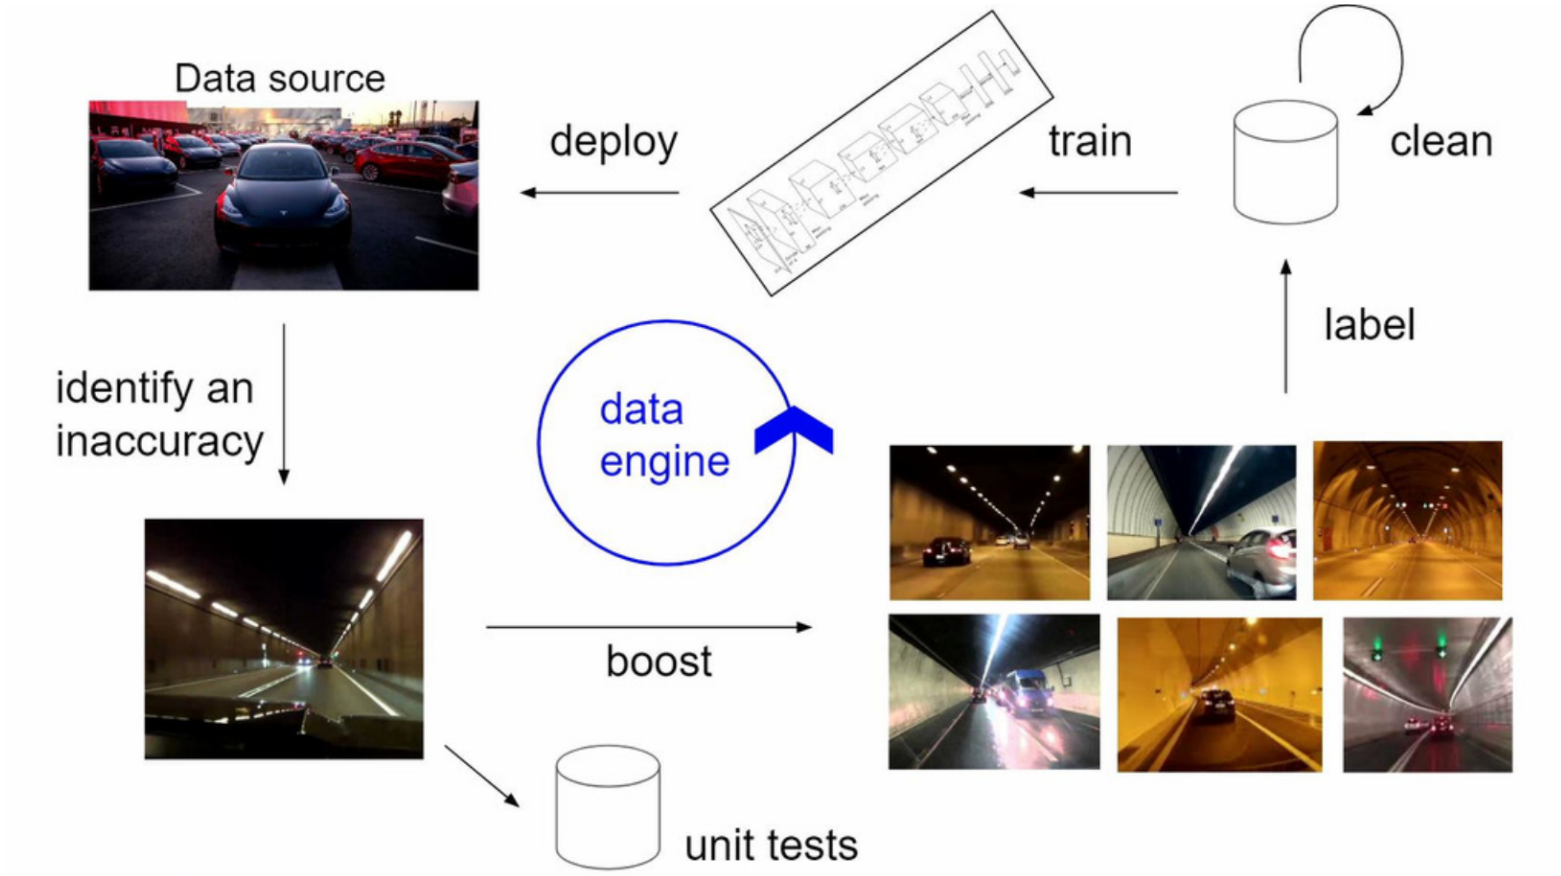

- fleet learning
- testing = shadow mode for more training data

Karpathy (Tesla) ICML 2019

Tesla's point of view on data

- Gathering process
  - 221 triggering situations
- manual labelling (1000 person)
  - 2d -> 3d
  - reconstruction labelling
- auto labelling
  - use specificly trained networks
  - human to clean
- simulation
  - rare event
  - sensor robustness
  - adversarial exemples

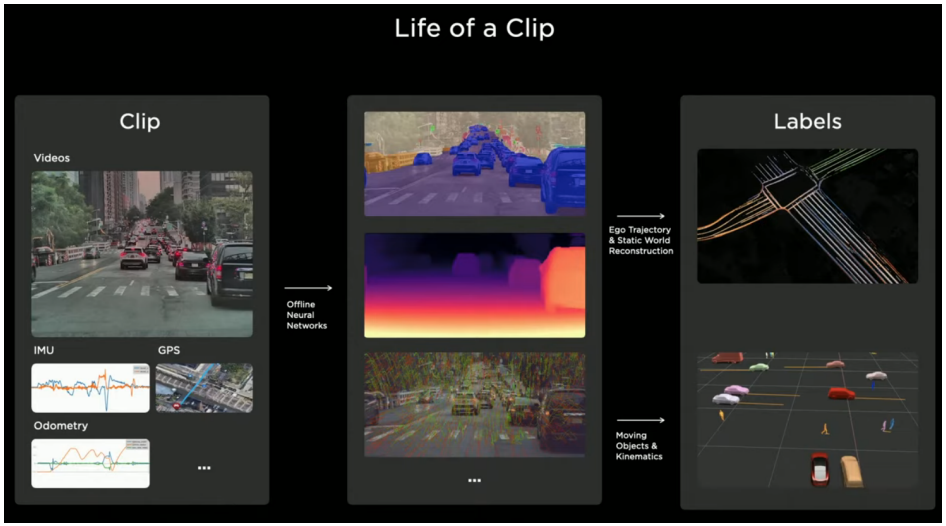

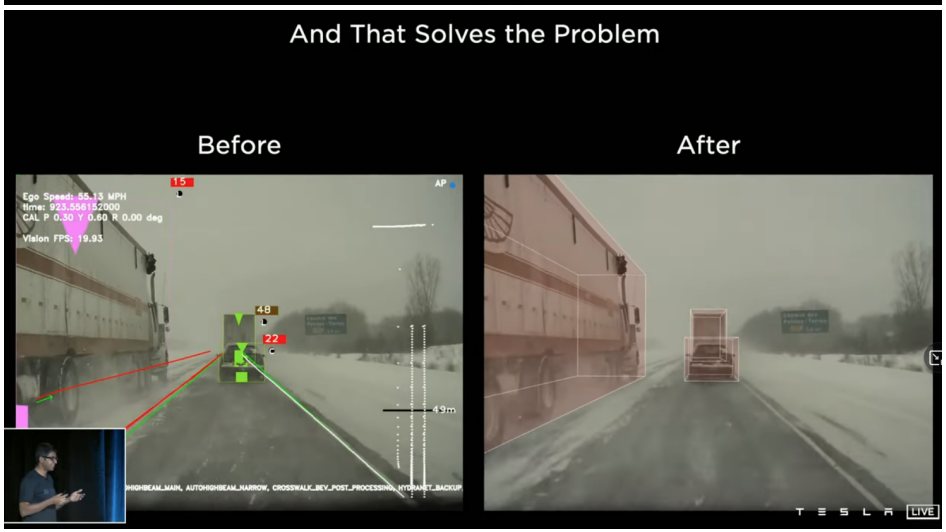

Tesla's AI day youtube.com/watch?v=j0z4FweCy4M

Openpilot : l'étiquetage des données par crowd sourcing

comma10k

Count and Percentage of Available Images Labeled 6344 out of 9874, 64.25%

This is the first 2,000 images of our internal comma10k dataset. After we clean up these new labels, we'll release more. Learn more from the Medium post, or on the comma.ai discord in the #comma-pencil channel.

It's 10,000 pngs of real driving captured from the comma fleet. It's MIT license, no academic only restrictions or

https://github.com/commaai/comma10k

Waymo's open data set

574 hours of data

https://github.com/waymo-research/waymo-open-dataset

AI issues in self driving

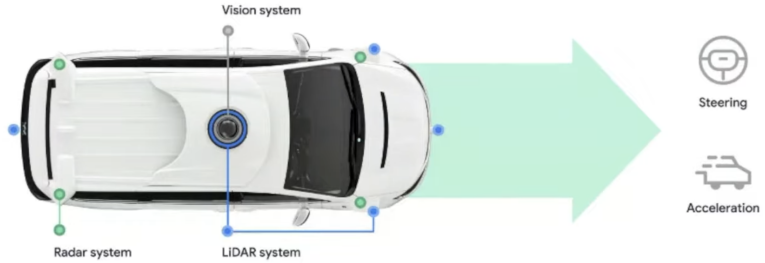

- modular end-to-end differential programming
- multi task, multi objective
- architecture design issues
- scene understanding: the never ending learning (long tails events)
- under budget

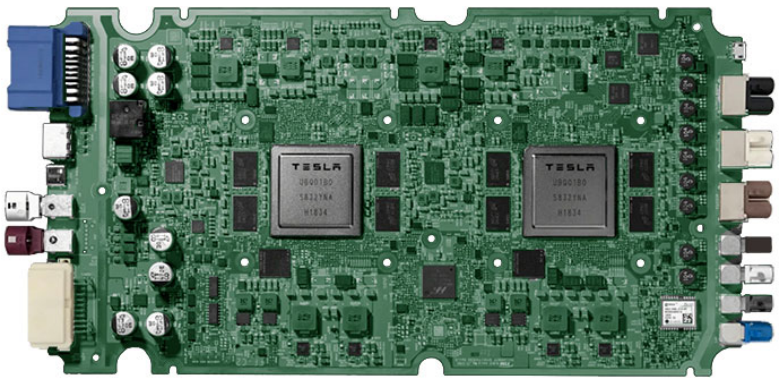

Tesla Full self-driving computer

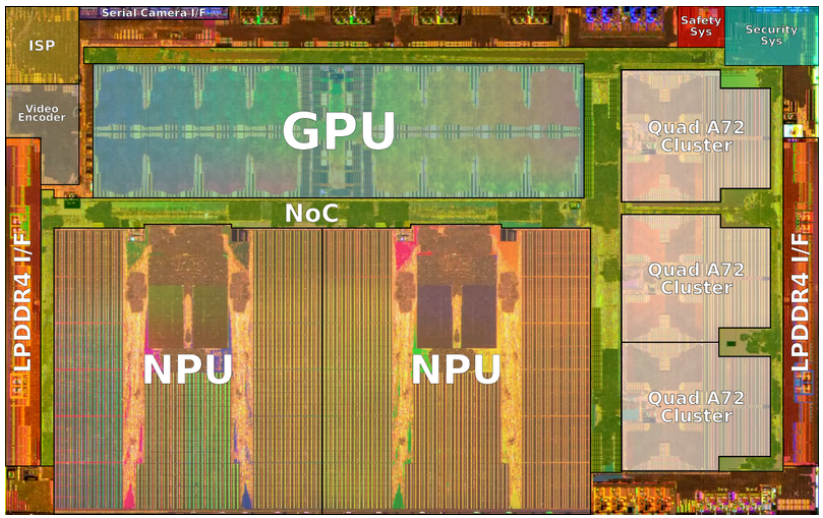

Tesla Full Self-Driving Chip

144 TOPS / 2300 Frames per second

AI hardware

Apple

2022

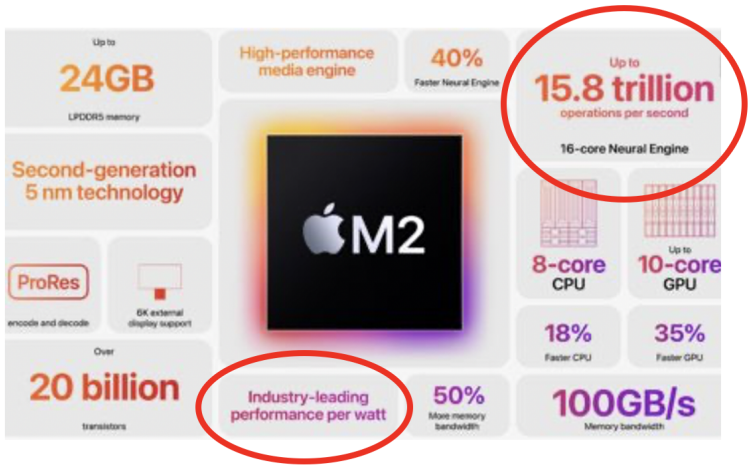

MobilEye

2018

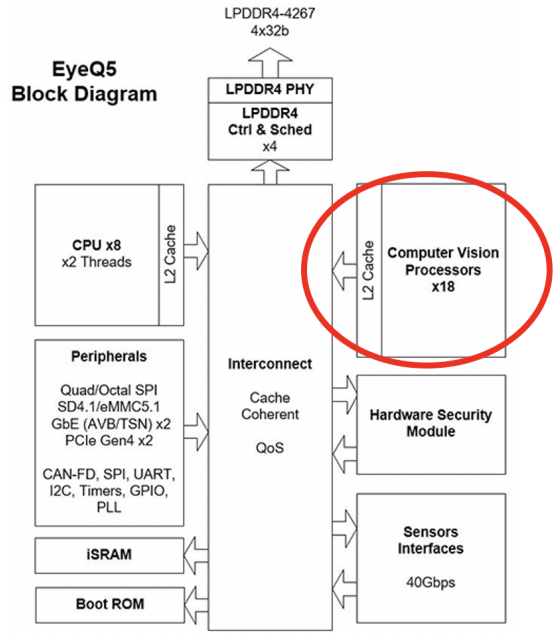

Qualcom

2019

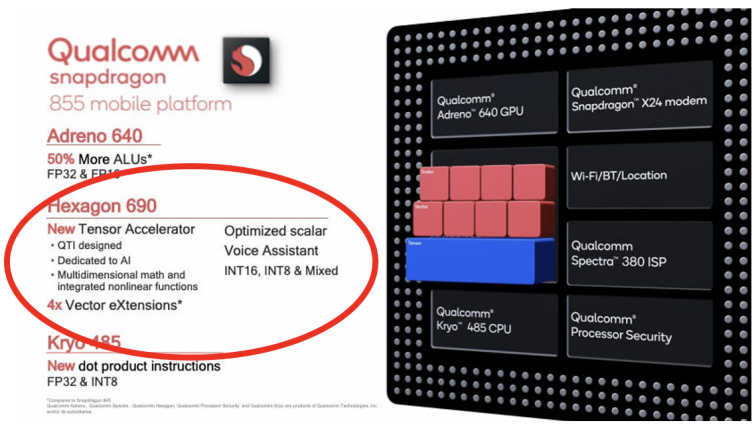

Tesla

2019

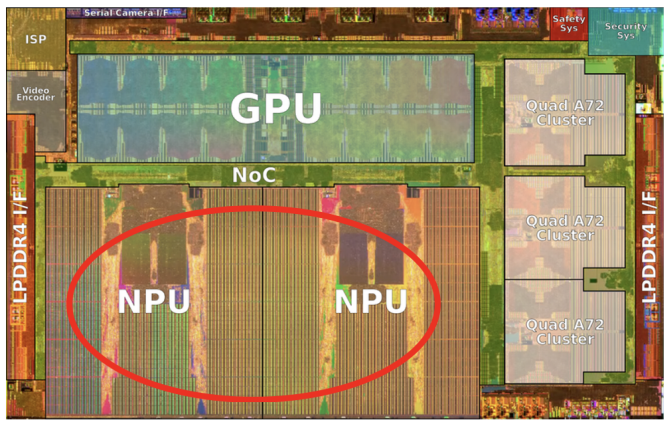

The Audi A8 hardware

Automotive tracks – Audi A8
Level 3: Aptiv zFAS controller

(Source: www.reverse-costing.com, System Plus Consulting)

➢ NVIDIA Tegra K1
Traffic sign recognition
Pedestrian detection
collision avoidance warning
Light detection
Lane recognition

➢ MobilEye EyeQ3
Traffic sign recognition
360° camera images & processing

Functions: Courtesy of Aptiv

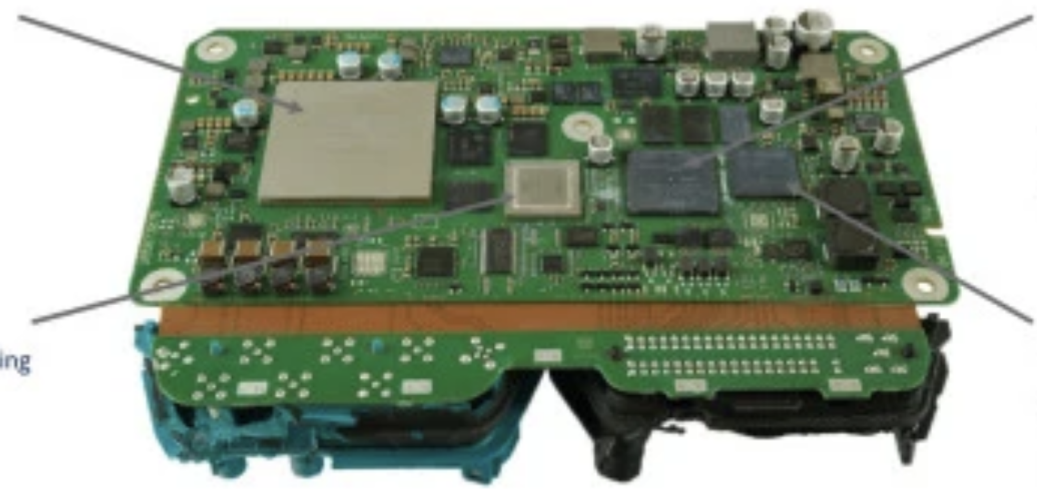

➢ Altera Cyclone
Object fusion
Map fusion
Parking pilot
Pre-crash
Sensor data pre-processing

➢ Infineon Aurix Tricore
Traffic jam pilot
Assistance systems
Matrix beam
Road graph

Comparizon

Company	DL framework	sensors	hardware (chip)
Openpilot	Meta Pytorch?	cameras + radar	Qualcomm (M1?)
Tesla	Meta Pytorch	8 cameras	Tesla's FSD chip
Mobil eye	Tensorflow on AWS	11 cameras (vidar)	ST microelectronic
Waymo	Google Tensorflow	cameras + Lidars + radars	Intel -> Samsung ?
Cruise	Microsoft Azure	4 cameras + Lidar + radar + audio	origin cruise chip

Road map

1. A very brief history of autonomous vehicles
2. How has this happened? (Deep Learninng)
3. Data to train the deep network
4. How Artificial Intelligence will change the Auto
5. Conclusion

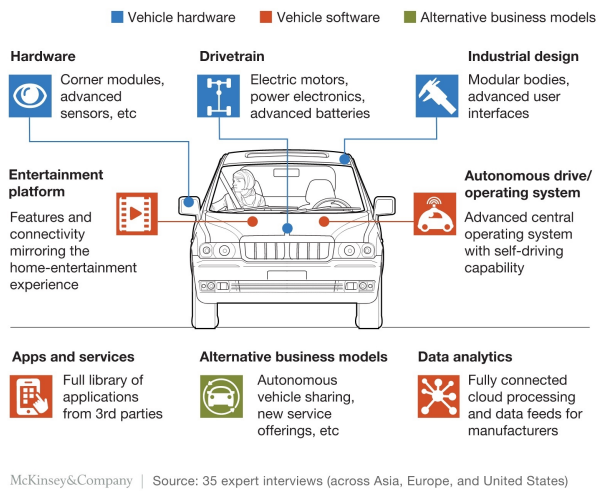

Programmation par l'exemple : le pari de Tesla & Waymo

Tesla is collecting "just over 3 million miles [of data] per day."

Waymo train the car with 6 million miles driven on public roads and 5 billion driven in simulation

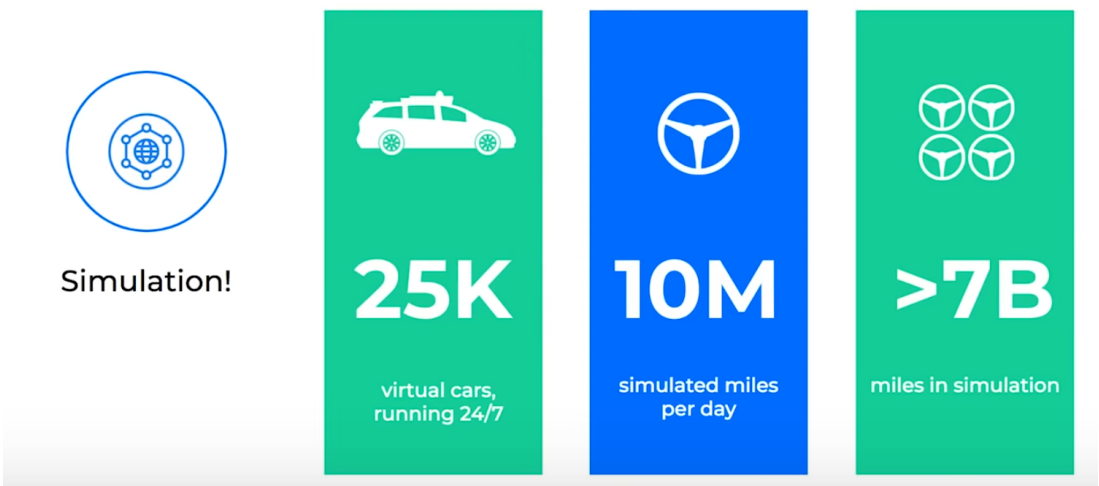

Learn agent for driving situation simulations

ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst Mayank Bansal, Alex Krizhevsky, Abhijit Ogale

Massive open data sets

BDD100K: A Large-scale Diverse Driving Video Database

Fisher Yu *May 30, 2018*

Update 06/18/2018: please also check our follow-up blog post after reading this.

TL;DR, we released the largest and most diverse driving video dataset with rich annotations called BDD100K. You can access the data for research now at http://bdd-data.berkeley.edu. We have recently released an arXiv report on it. And there is still time to participate in our CVPR 2018 challenges!

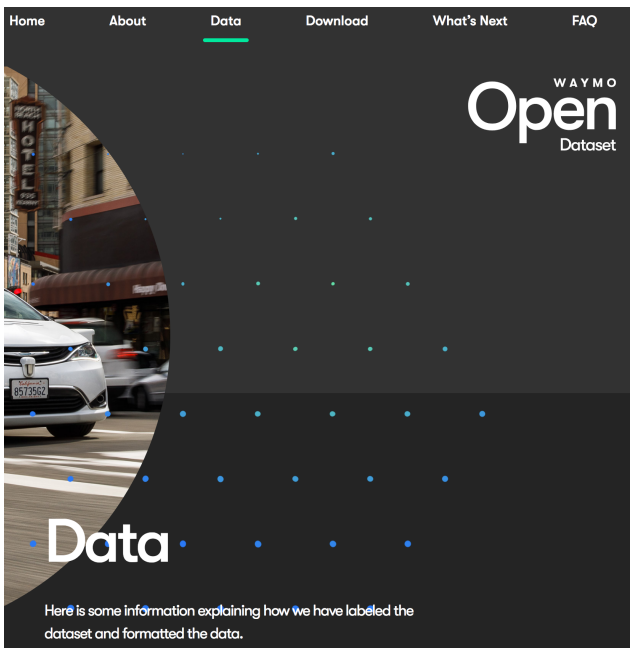

Baidu Apollo Releases Massive Self-driving Dataset; Teams Up With Berkeley DeepDrive

Baidu this Thursday announced the release of ApolloScape, billed as the world's largest open-source dataset for autonomous driving technology.

and simulators (Carla, google & microsoft)

Pour quoi faire ?

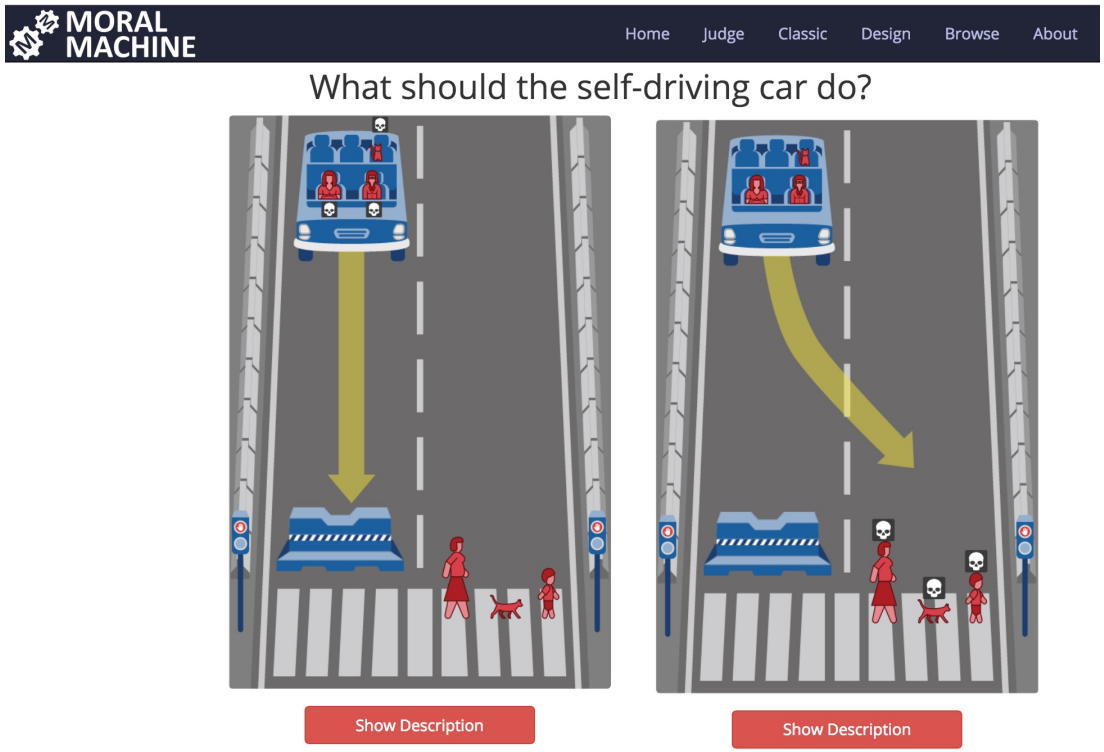

Accepter les voitures autonomes

- Une histoire de confiance
- Nous voulons comprendre les enjeux
- conduite du changement

http://moralmachine.mit.edu/

Towards scaling self driving

When will we have more than 10,000 Full Autonomous cars?

- Tesla's strategy of the little steps (improving the ADAS)
- Wyamo strategy including more areas (less specific)

- not yet: status quo
  - driving assistance (automation)
    - ★ increase safety
    - ★ reduces environmental impact
  - specific applications
  - communication and equipment

- No full autonomy unless... safety is proven
  - new solution (cf Google)

Accidents: 14 lethal since 2015 (+1 processing)

https://en.wikipedia.org/wiki/List_of_self-driving_car_fatalities

Safety Ratings

Safety Assist evaluating driver-assistance and crash-avoidance technologies.

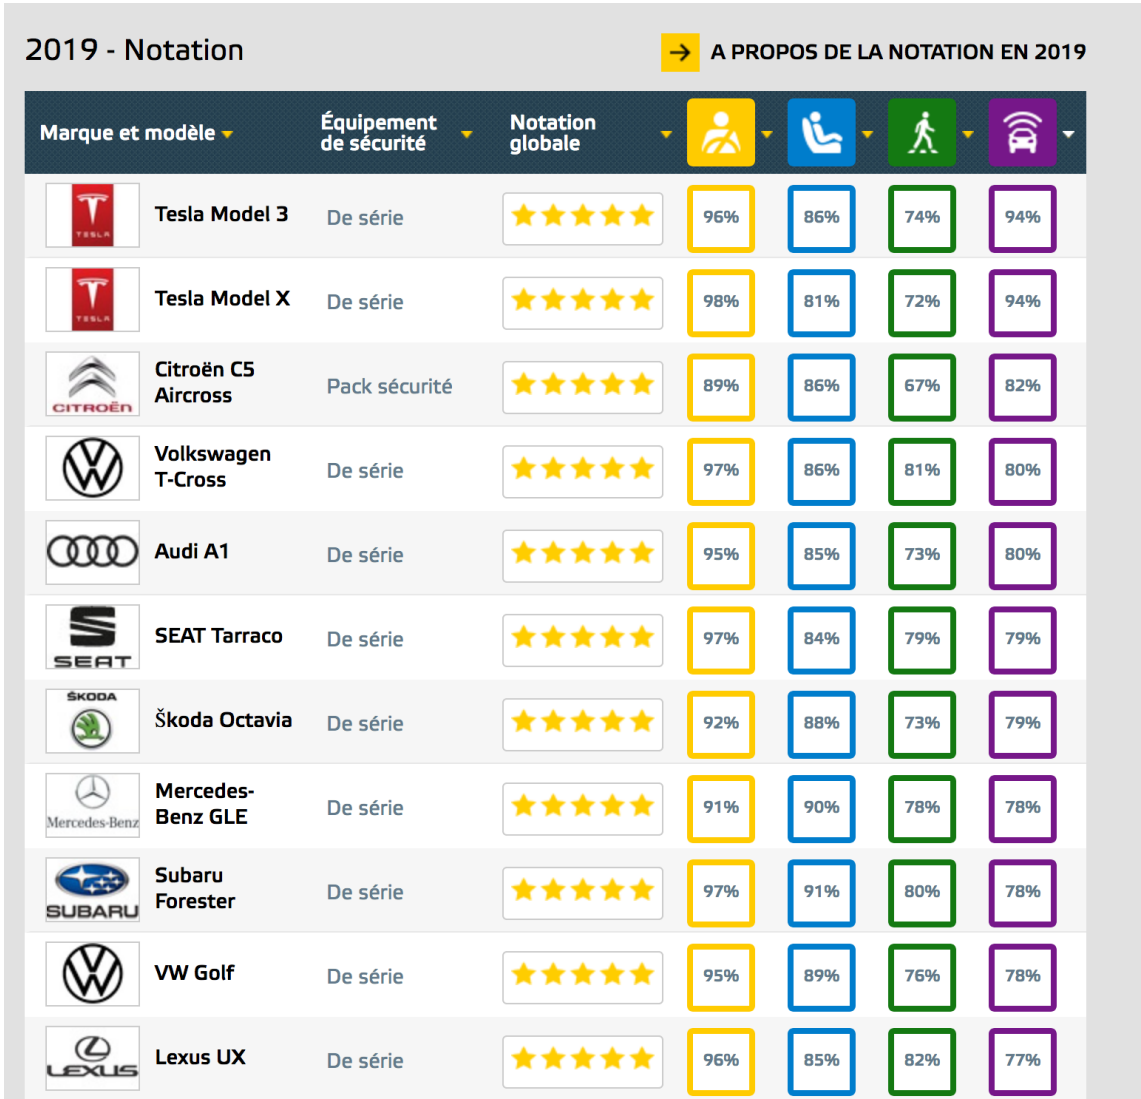

2019 - Notation

A PROPOS DE LA NOTATION EN 2019

Marque et modèle	Équipement de sécurité	Notation globale				
Tesla Model 3	De série	★★★★★	96%	86%	74%	94%
Tesla Model X	De série	★★★★★	98%	81%	72%	94%
Citroën C5 Aircross	Pack sécurité	★★★★★	89%	86%	67%	82%
Volkswagen T-Cross	De série	★★★★★	97%	86%	81%	80%
Audi A1	De série	★★★★★	95%	85%	73%	80%
SEAT Tarraco	De série	★★★★★	97%	84%	79%	79%
Škoda Octavia	De série	★★★★★	92%	88%	73%	79%
Mercedes-Benz GLE	De série	★★★★★	91%	90%	78%	78%
Subaru Forester	De série	★★★★★	97%	91%	80%	78%
VW Golf	De série	★★★★★	95%	89%	76%	78%
Lexus UX	De série	★★★★★	96%	85%	82%	77%

https://www.euroncap.com/en/ratings-rewards/latest-safety-ratings/

Attacks against autonomous vehicles

Eykholt et al, Robust Physical-World Attacks on Deep Learning Visual Classification, CVPR 2018

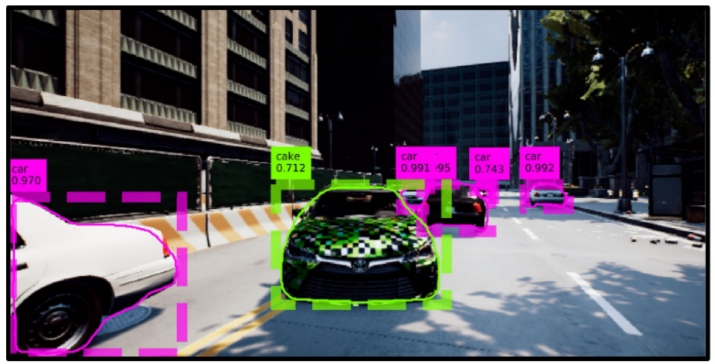

Zhang et al., CAMOU: Learning Physical Vehicle Camouflages to Adversarially Attack Detectors in the Wild, ICLR 2019

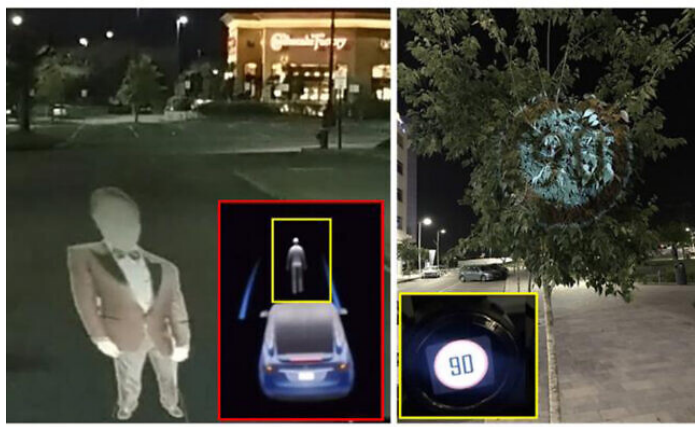

https://www.mcafee.com/blogs/other-blogs/mcafee-labs/model-hacking-adas-to-pave-safer-roads-for-autonomous-vehicles/

Nassi et al., Phantom of the ADAS: Securing Advanced Driver-AssistanceSystems from Split-Second Phantom Attacks, 2020

Qayyum, et al., Securing Connected & Autonomous Vehicles: Challenges Posed by Adversarial ML, IEEE Communications, 2019

Attacking Openpilot 's DMS

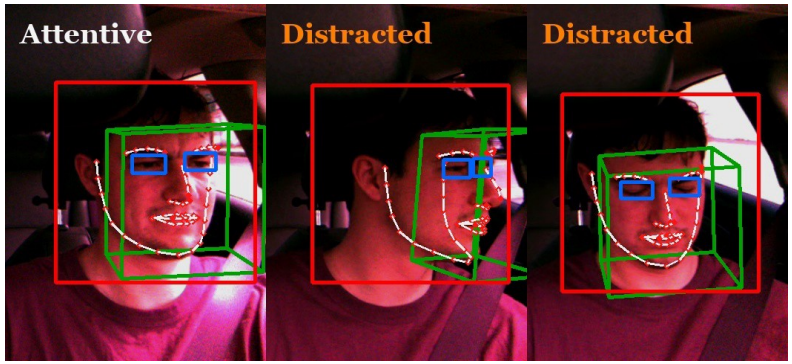

Three components

- Face localization
  - openCV -> cropping
- Feature generation
  - EfficentNet b0 architecture
  - Fine tuning
- Decision module
  - Treshold based decision

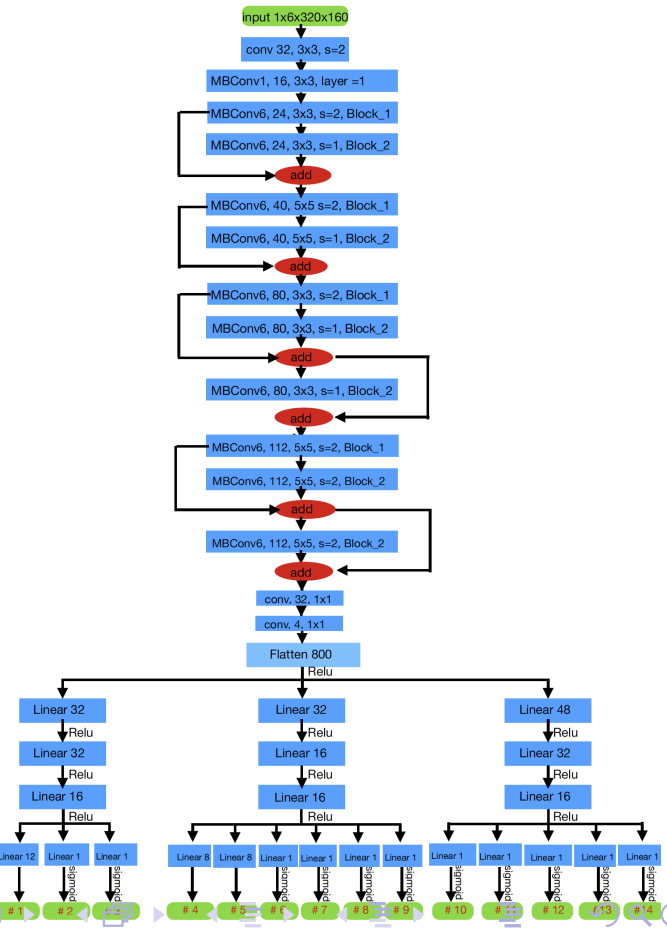

Datasets: Pandora (head pose)

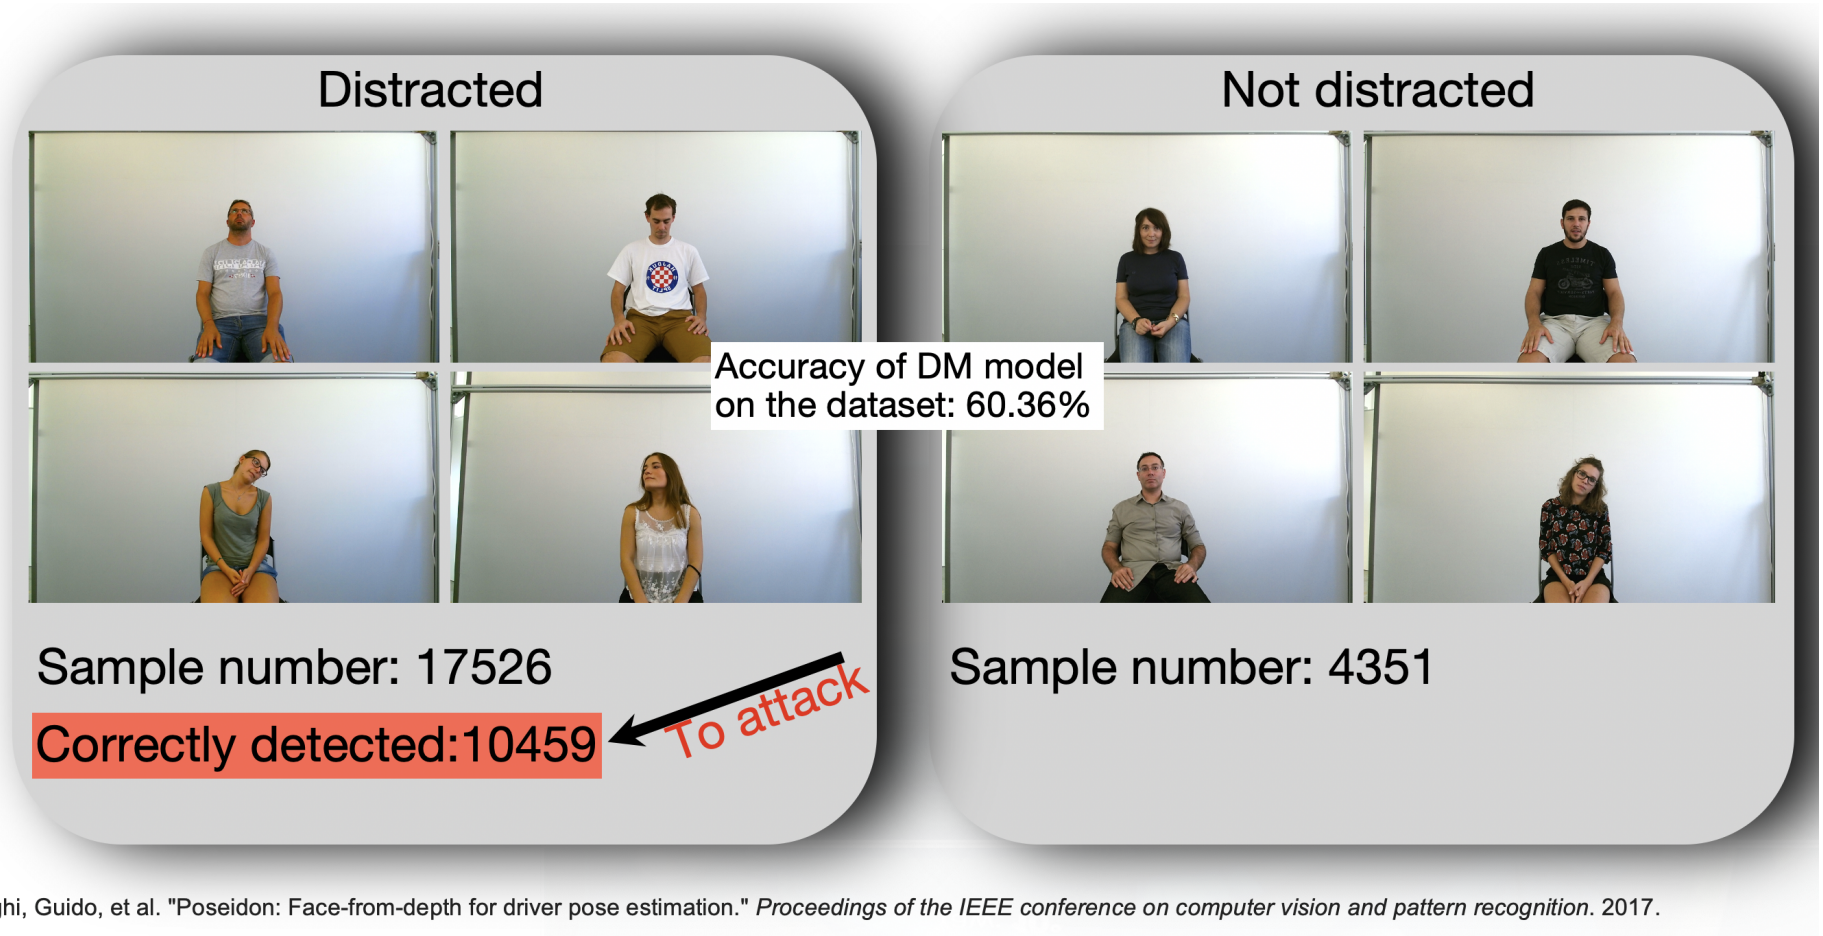

1. Borghi, Guido, et al. "Poseidon: Face-from-depth for driver pose estimation." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017.

Attack performance

- Accuracy on original data: 100%
- Attack settings:
  - torchattacks
  - c=1000 for CW
  - *steps =50 for CW and Deepfool*
  - L_∞ *10/255 = for all the others*
- Accuracy on adversarial data:

Attack models	FGSM	CW	PGD	APGD	AutoAttack	Deepfool
Accuracy(%)	81.85	21.90	13.17	0.057	**0.0**	6.39

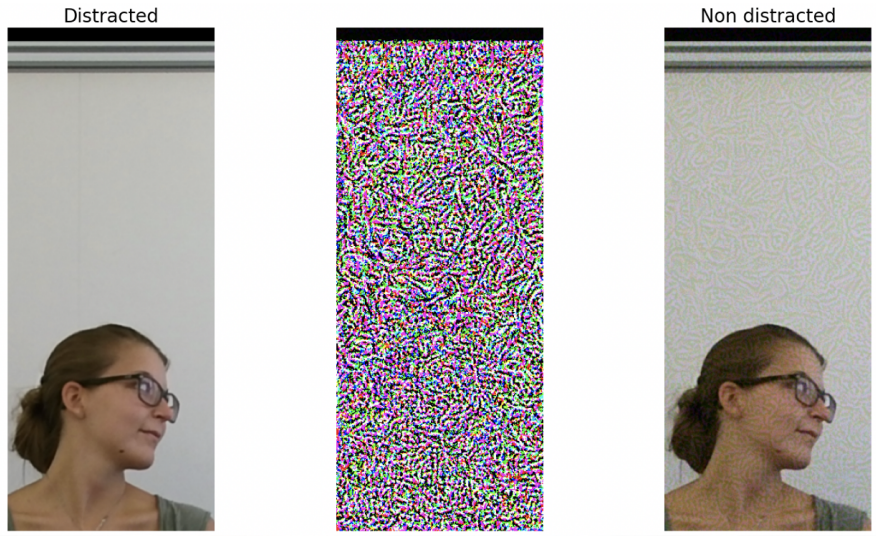

100 % Distracted

100 % Attentive

Road map (done)

1. A very brief history of autonomous vehicles
2. How has this happened? (Deep Learninng)
3. Data to train the deep network
4. How Artificial Intelligence will change the Auto
5. Conclusion

Future of AI in Automotive Industry

- The value of your data (IA fuel) — big data
- Robustness (degraded conditions) — deep learning theory
- Level 4 Autonomous driving — common sense (cf Y. LeCun) unsupervised learning
- Predictive Maintenance — data + prior knowledge
- Acceptability (safety) — Ethic Interpretable AI

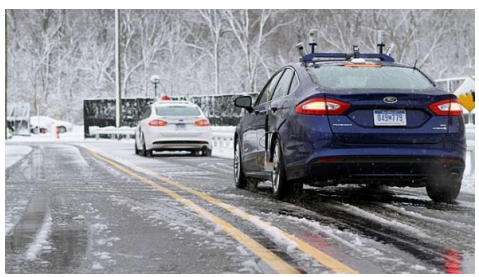

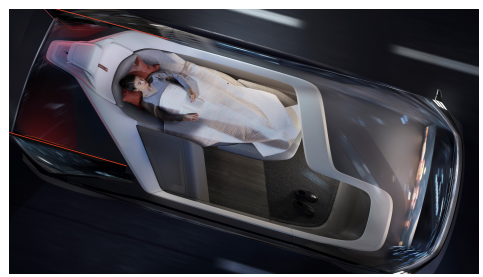

.

Les murs de l'IA

Les murs de l'IA

Les murs de l'IA

Après que l'IA s'est fracassée sur le mur,
le mur est encore là mais il n'y a plus d'IA

Les murs de l'IA

- Recherche : aujourd'hui c'est l'IA spécifique
- Confiance : Données, Validation et Sureté
- Biais - usage - éthique
- Durabilité : les questions énergétiques

Acceptabilité sociale de l'IA

Questions?

http://asi.insa-rouen.fr/enseignants/~scanu/