My Cross Sectional investigation into autonomous vehicle research

Jason Chalom

contact@jasonchalom.com

2018/06/24

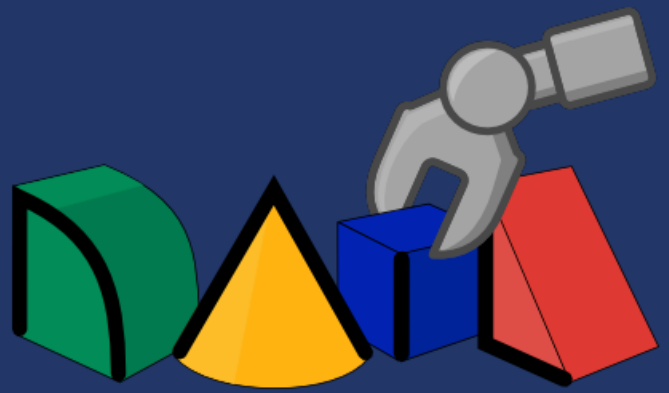

Slight Disclaimers

- My personal views etc ...
- There is just so much research in the field that this cannot be comprehensive
- Im no expert

How I got here

Who Am I?

- I am currently a part-time masters student
- I work as a software engineer at a start-up
- I have a fair bit of experience programming
- I like video games
- My honours project was in tree searching and game AIs (I did some interesting things with pseudo-RNGs)


During honours I did
Richard's Computer
Vision Course

GTA V at the time was being used as an autonomous car simulator


And together this seemed like a good masters project
MAKE GIFS AT GIFSOUP.COM

But it wasn't the best idea

Because of Lawyers

The field moves fast

And new more powerful simulators powered by video game engines have emerged

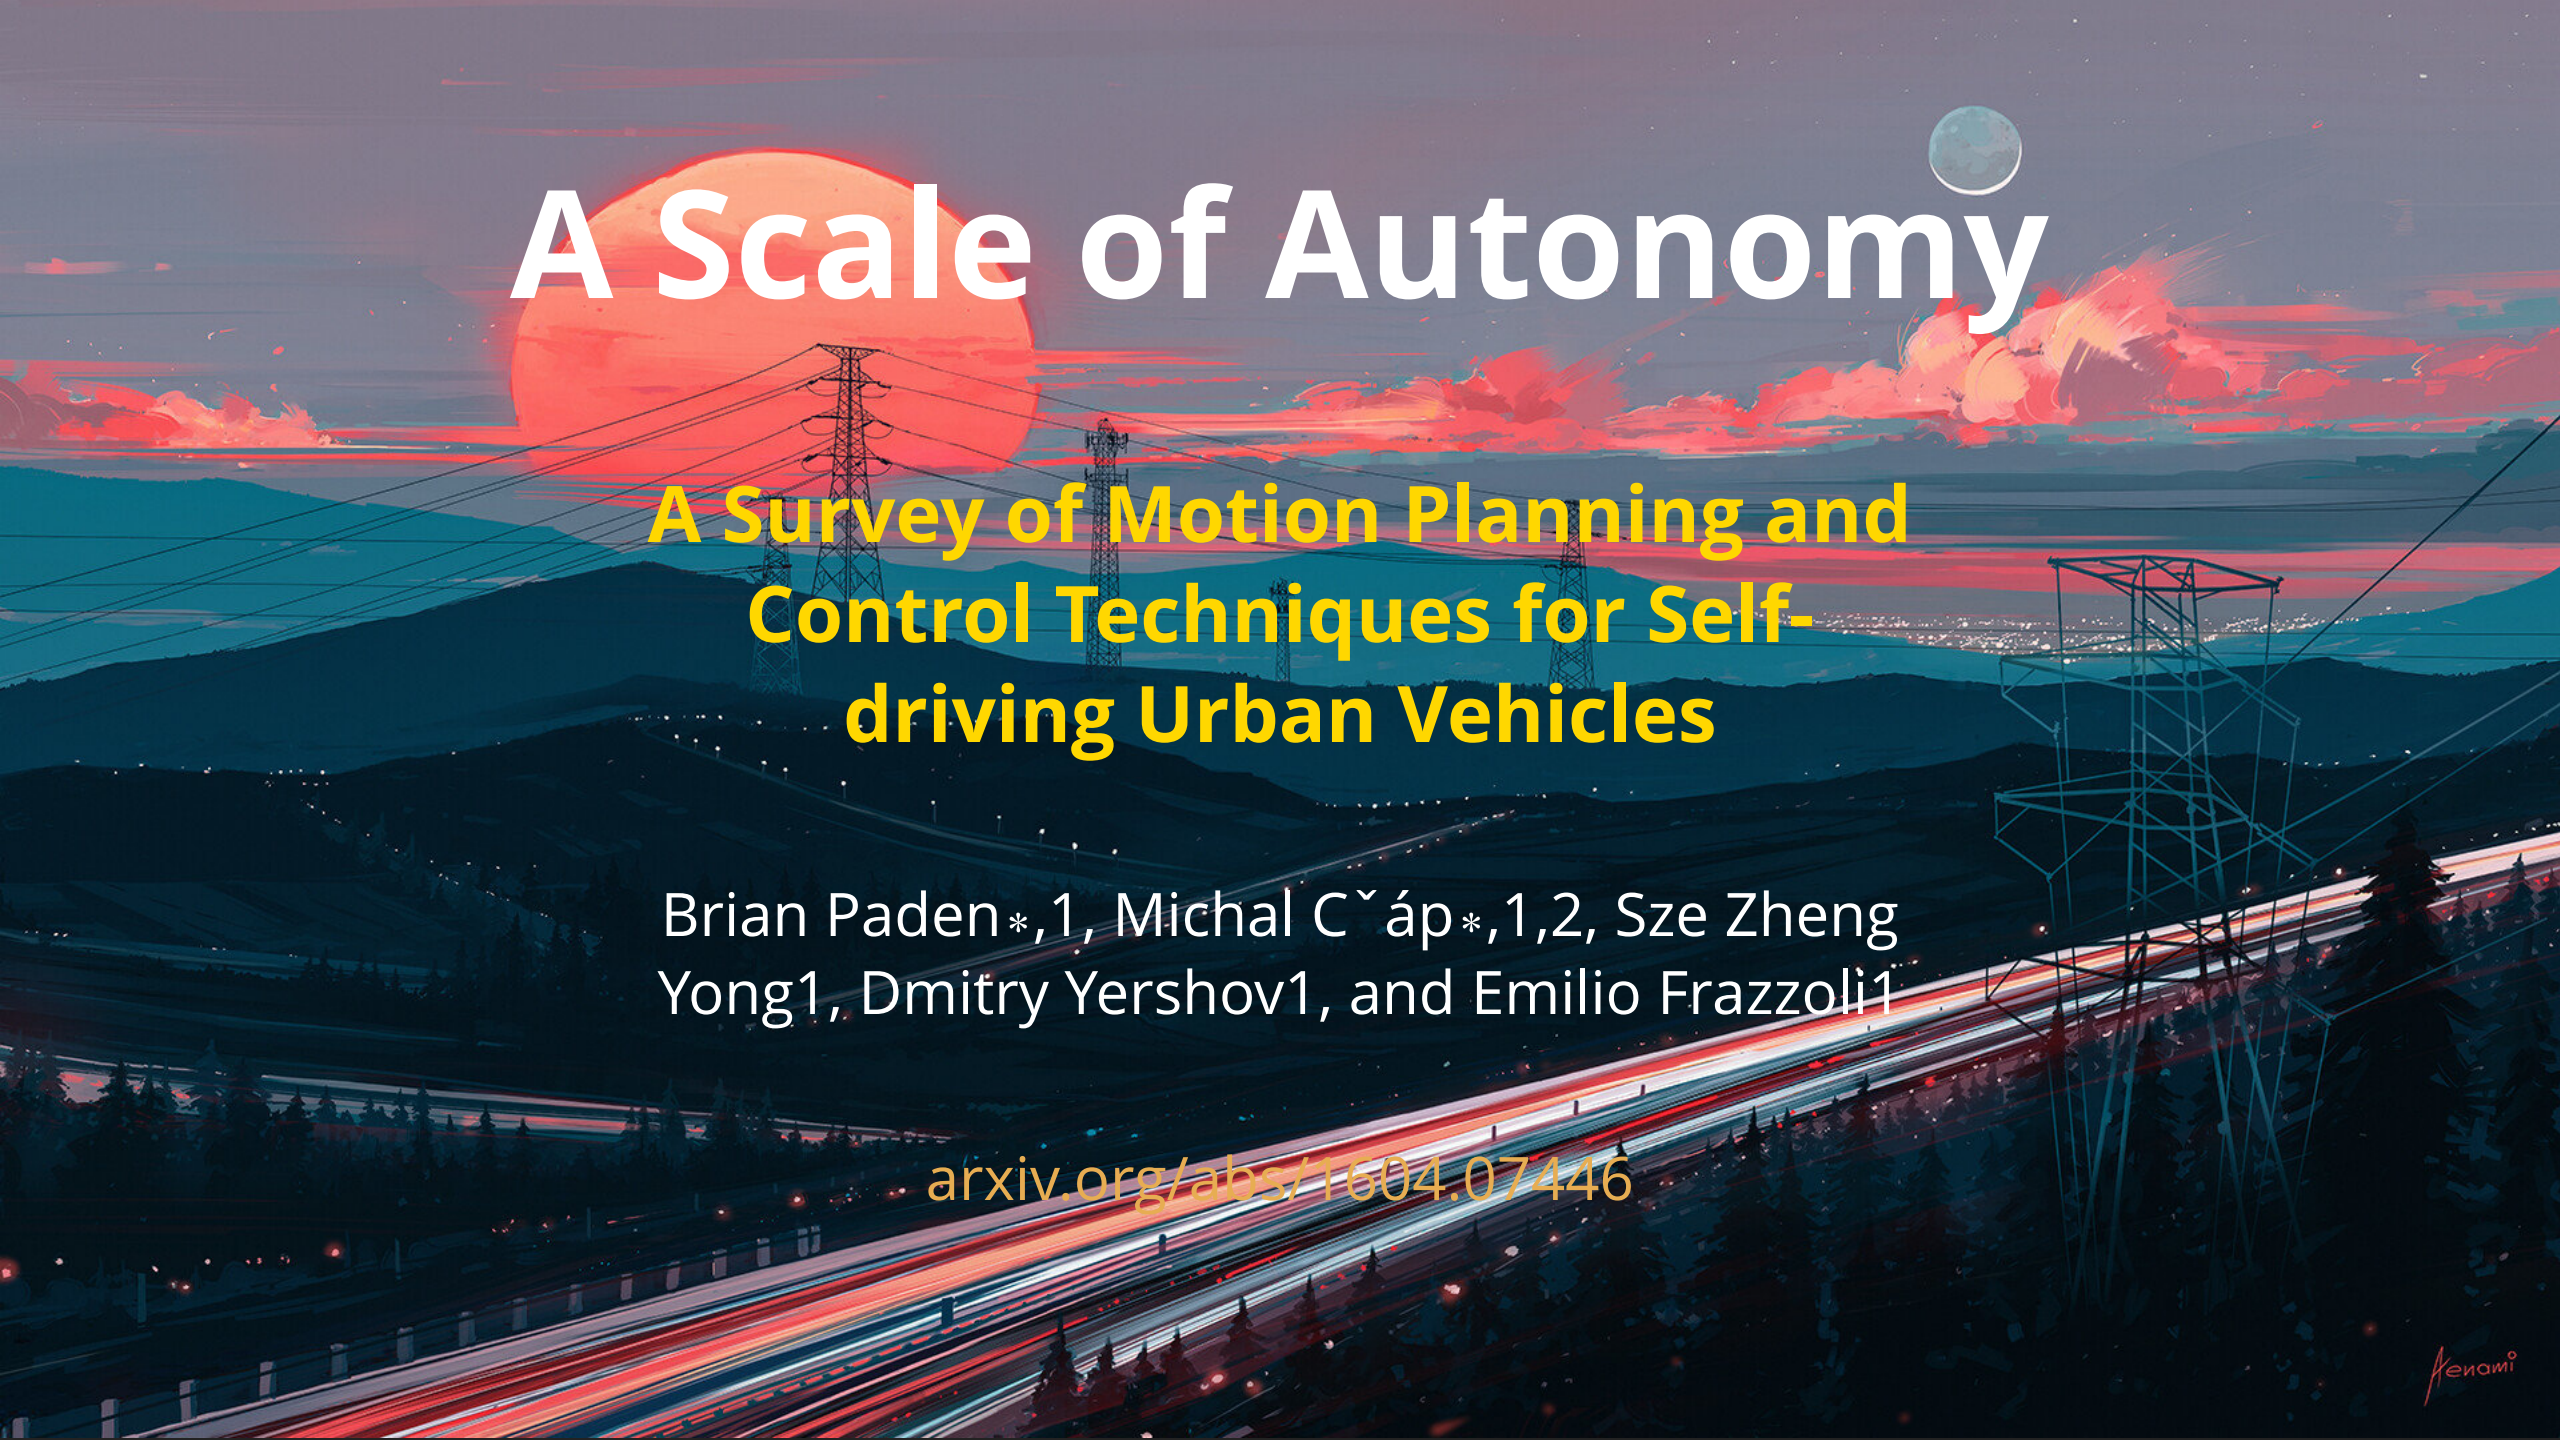

A Scale of Autonomy

A Survey of Motion Planning and Control Techniques for Self-driving Urban Vehicles

Brian Paden*,1, Michal Čáp*,1,2, Sze Zheng Yong1, Dmitry Yershov1, and Emilio Frazzoli1

arxiv.org/abs/1604.07446


Level 0 is fully human controlled
(all driving tasks)


Level 1 has basic automatic assistance such as adaptive cruise-control, anti-locking brakes, etc ...


Level 2 Has Advanced automated
control such as automatic braking

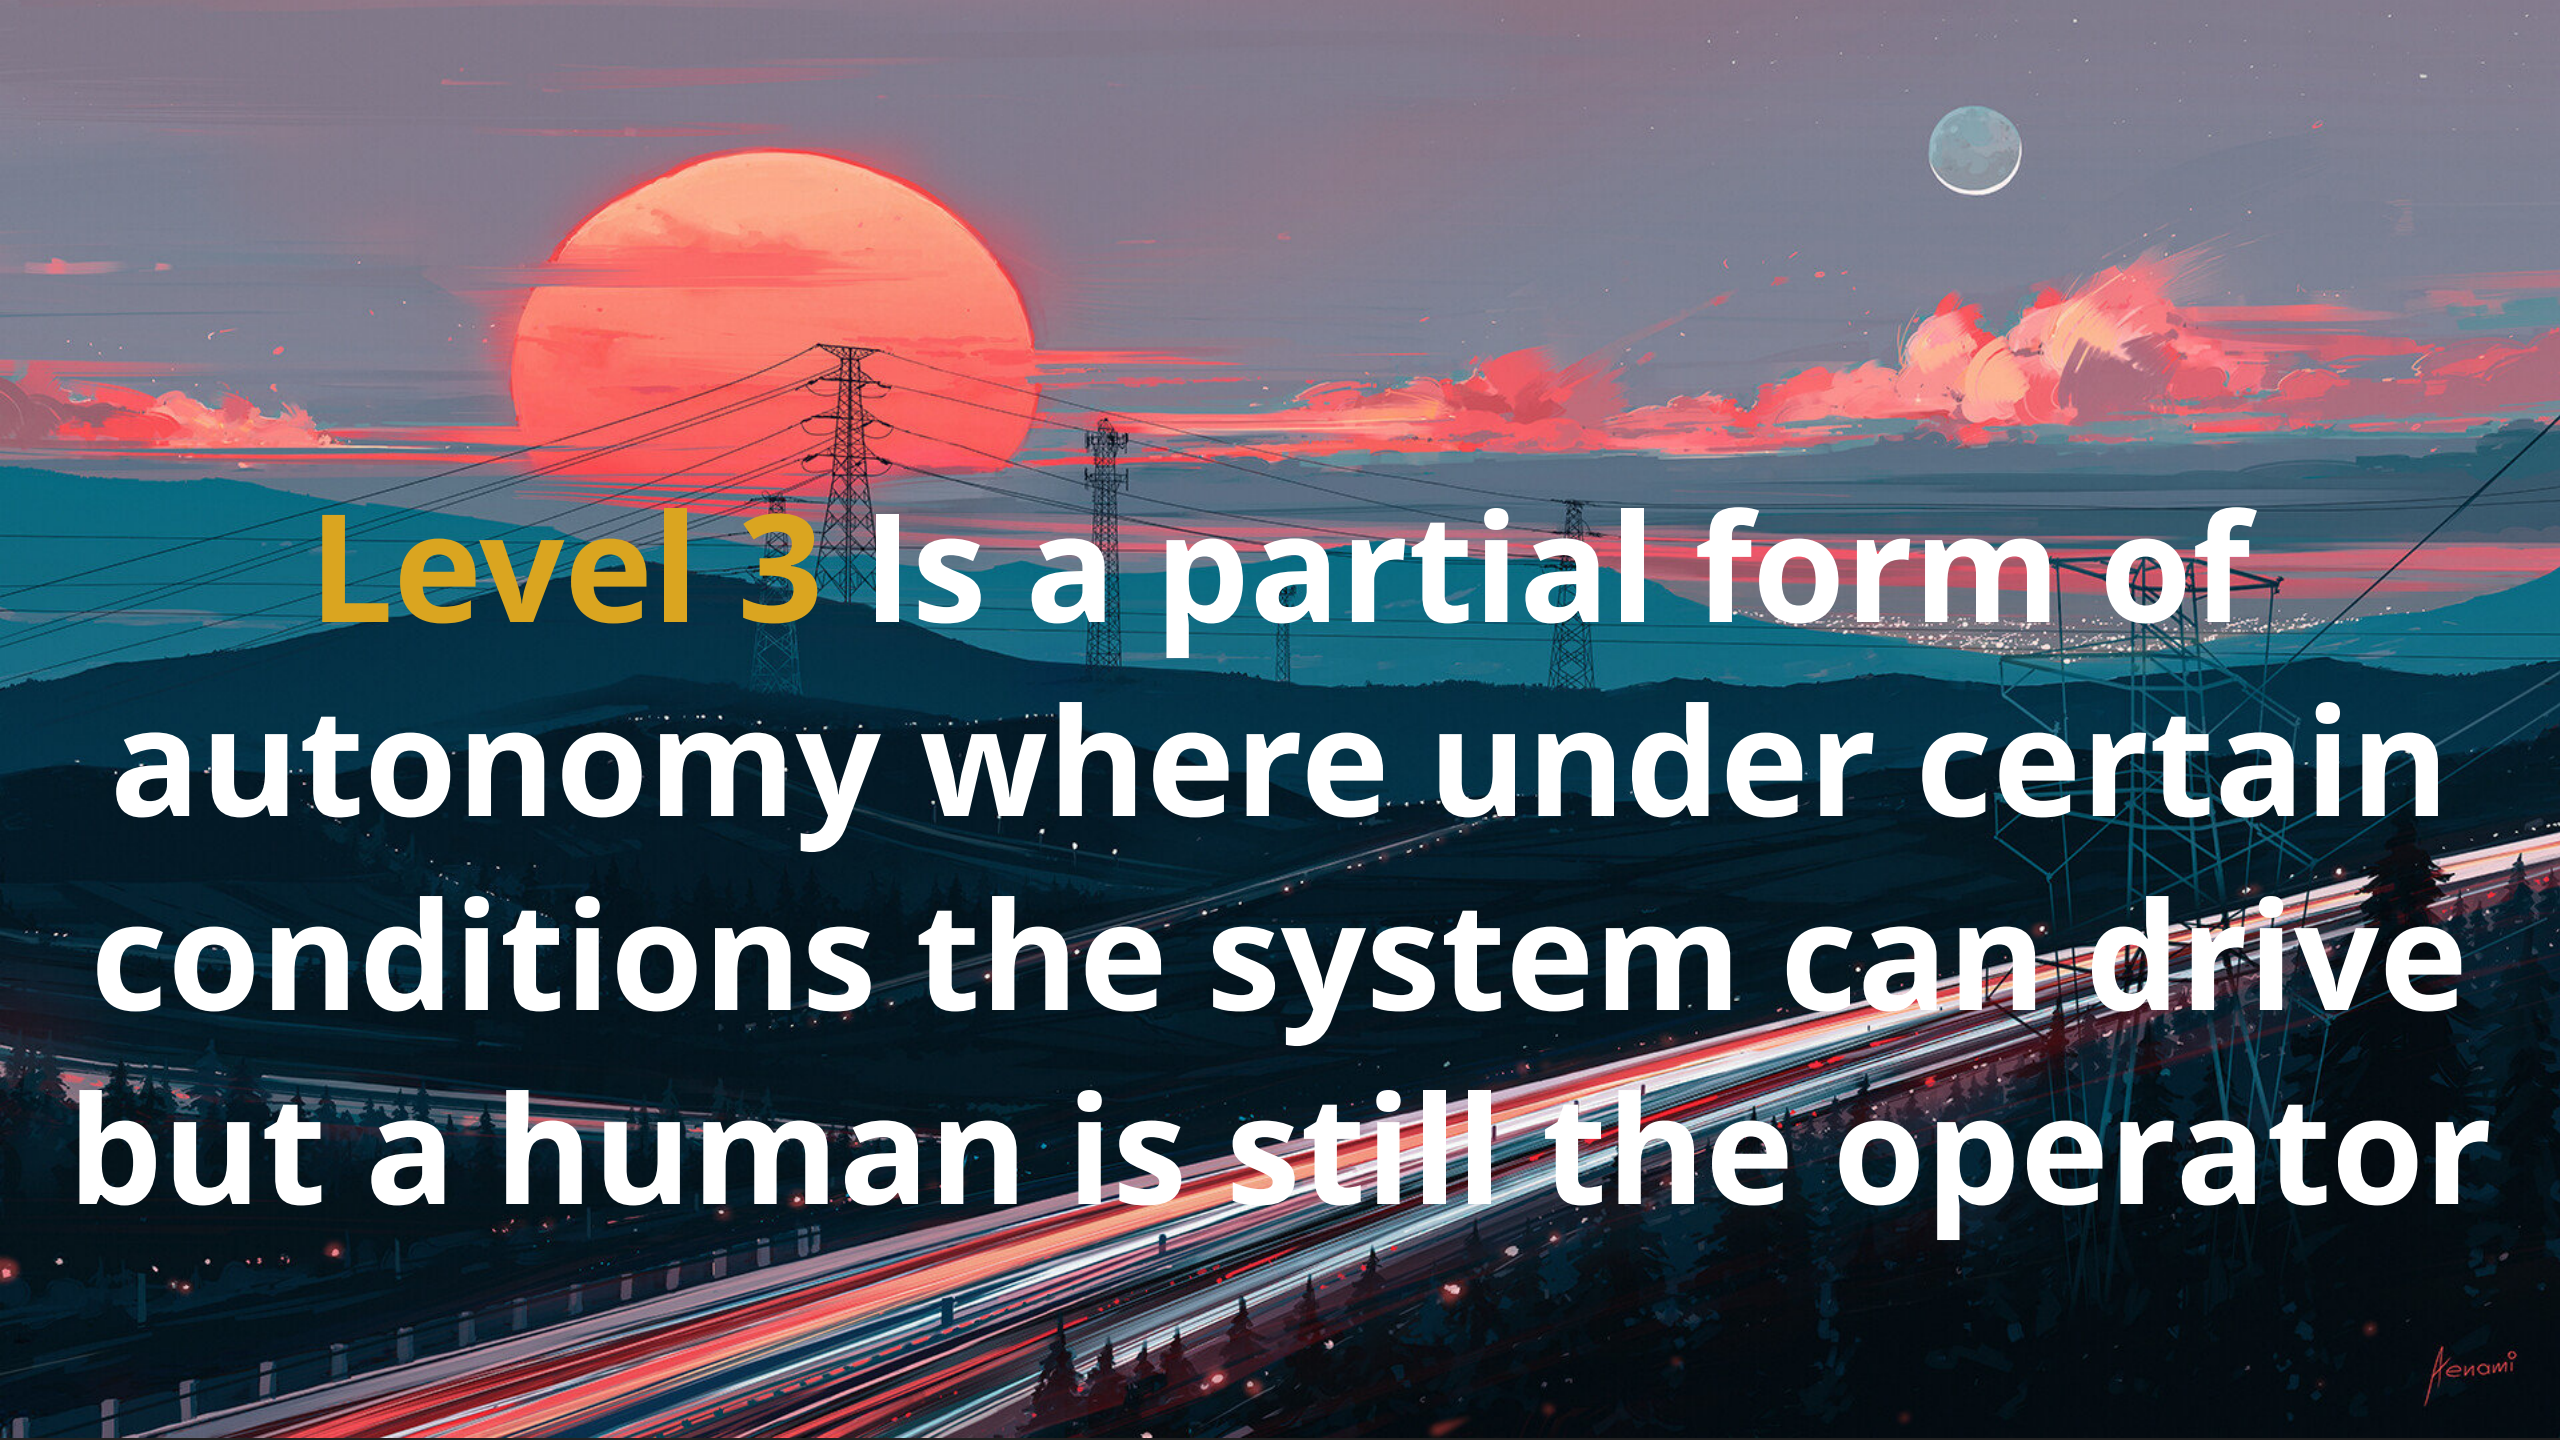
Level 3 Is a partial form of autonomy where under certain conditions the system can drive but a human is still the operator

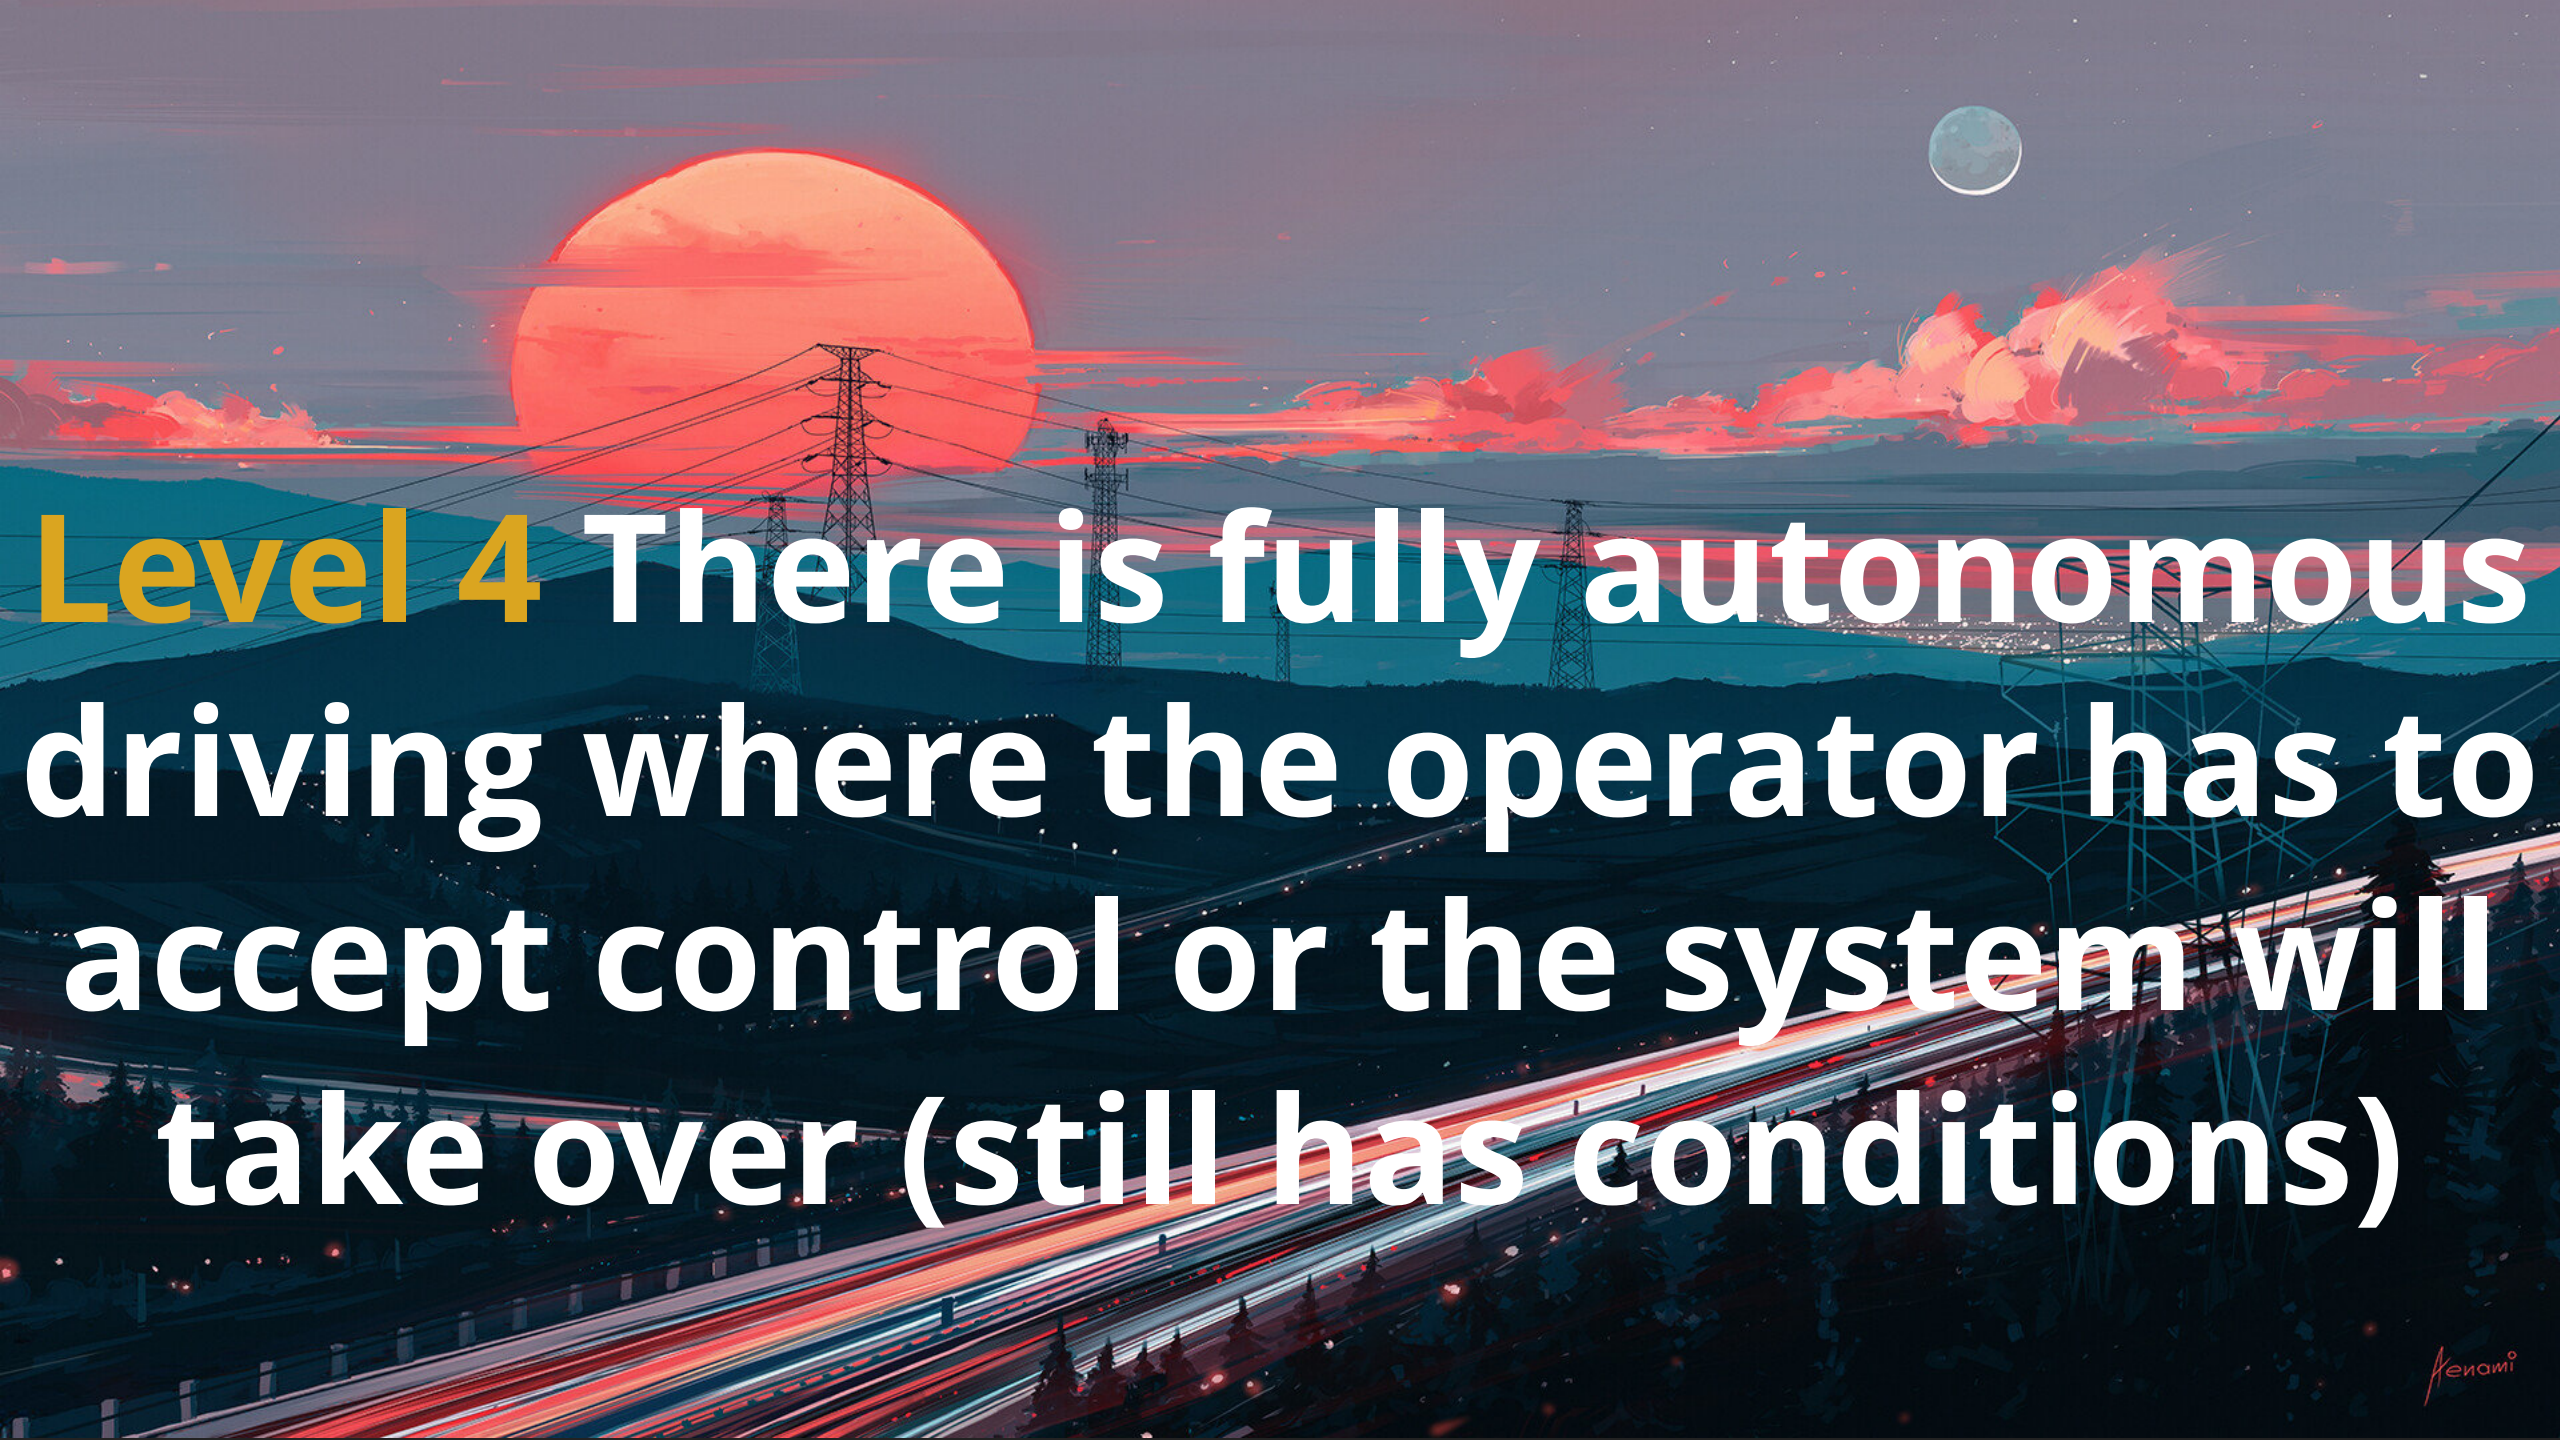
Level 4 There is fully autonomous driving where the operator has to accept control or the system will take over (still has conditions)


Level 5 System is fully
autonomous under all conditions

A measure of autonomy

Usually a percentage

A metric for the amount of autonomy given

$$autonomy = (1 - \frac{(number of interventions * 6[seconds]))}{elapsed time[seconds]}) \cdot 100$$

A bit of history

The 1920s

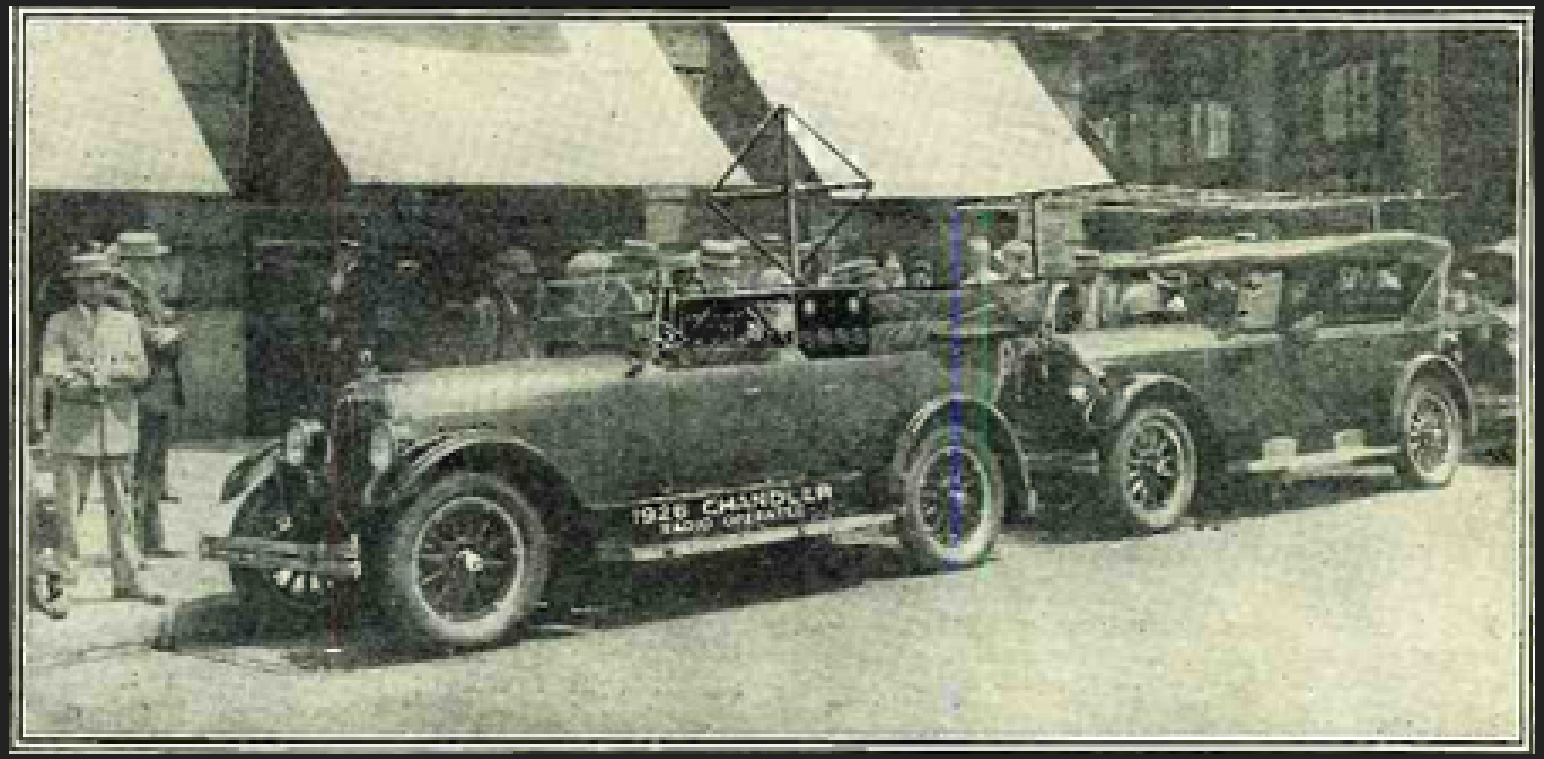

Radio Controlled Car

The Next Few Decades

- Research into automated guidance
- Better remote control (mostly military aircraft)
- Specialised roads
  - Radio controllers
  - Magnetic strips
- By the late 1960s there was the beginning of very basic computer control


1980s

Early Vision

DARPA ALV

Neural Networks

1990s

First Real Demonstrations

VaMP Driverless car (+1000 km)

US Dept Transport Funding


2000s

Military Interest Grand Challenge Better ECUs

2010s

The Decade of Hacking

Private Sector
Google + Tesla
The first trials

My Takeaways

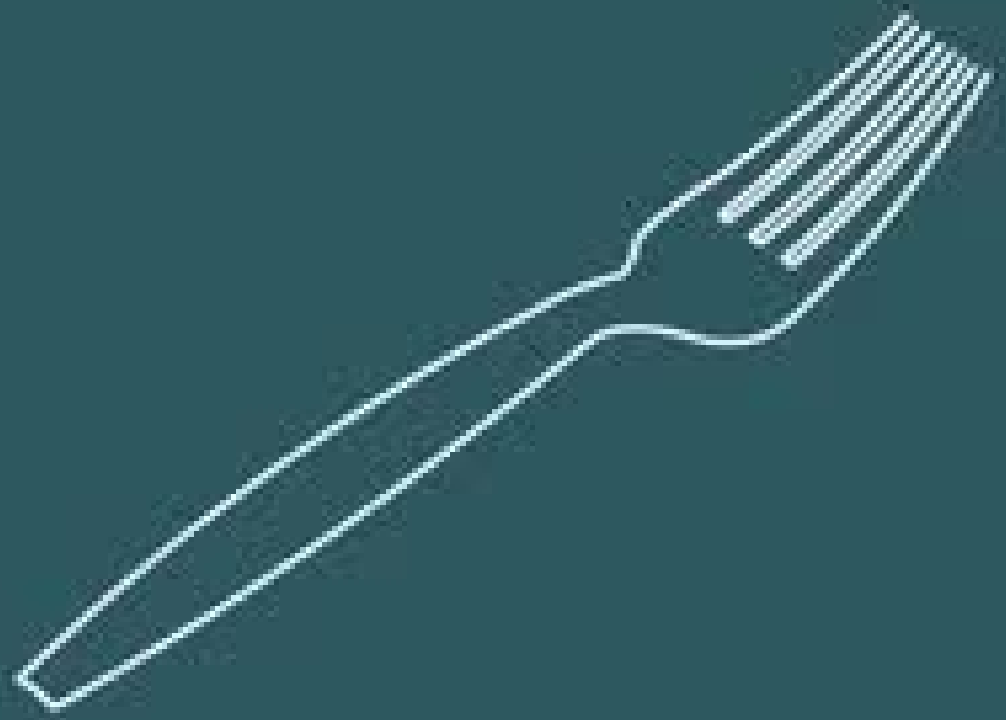

- There has always been concurrent research
- Autonomy is Easy. Safe and Effective Autonomy is Hard
- Most vehicles have autonomy

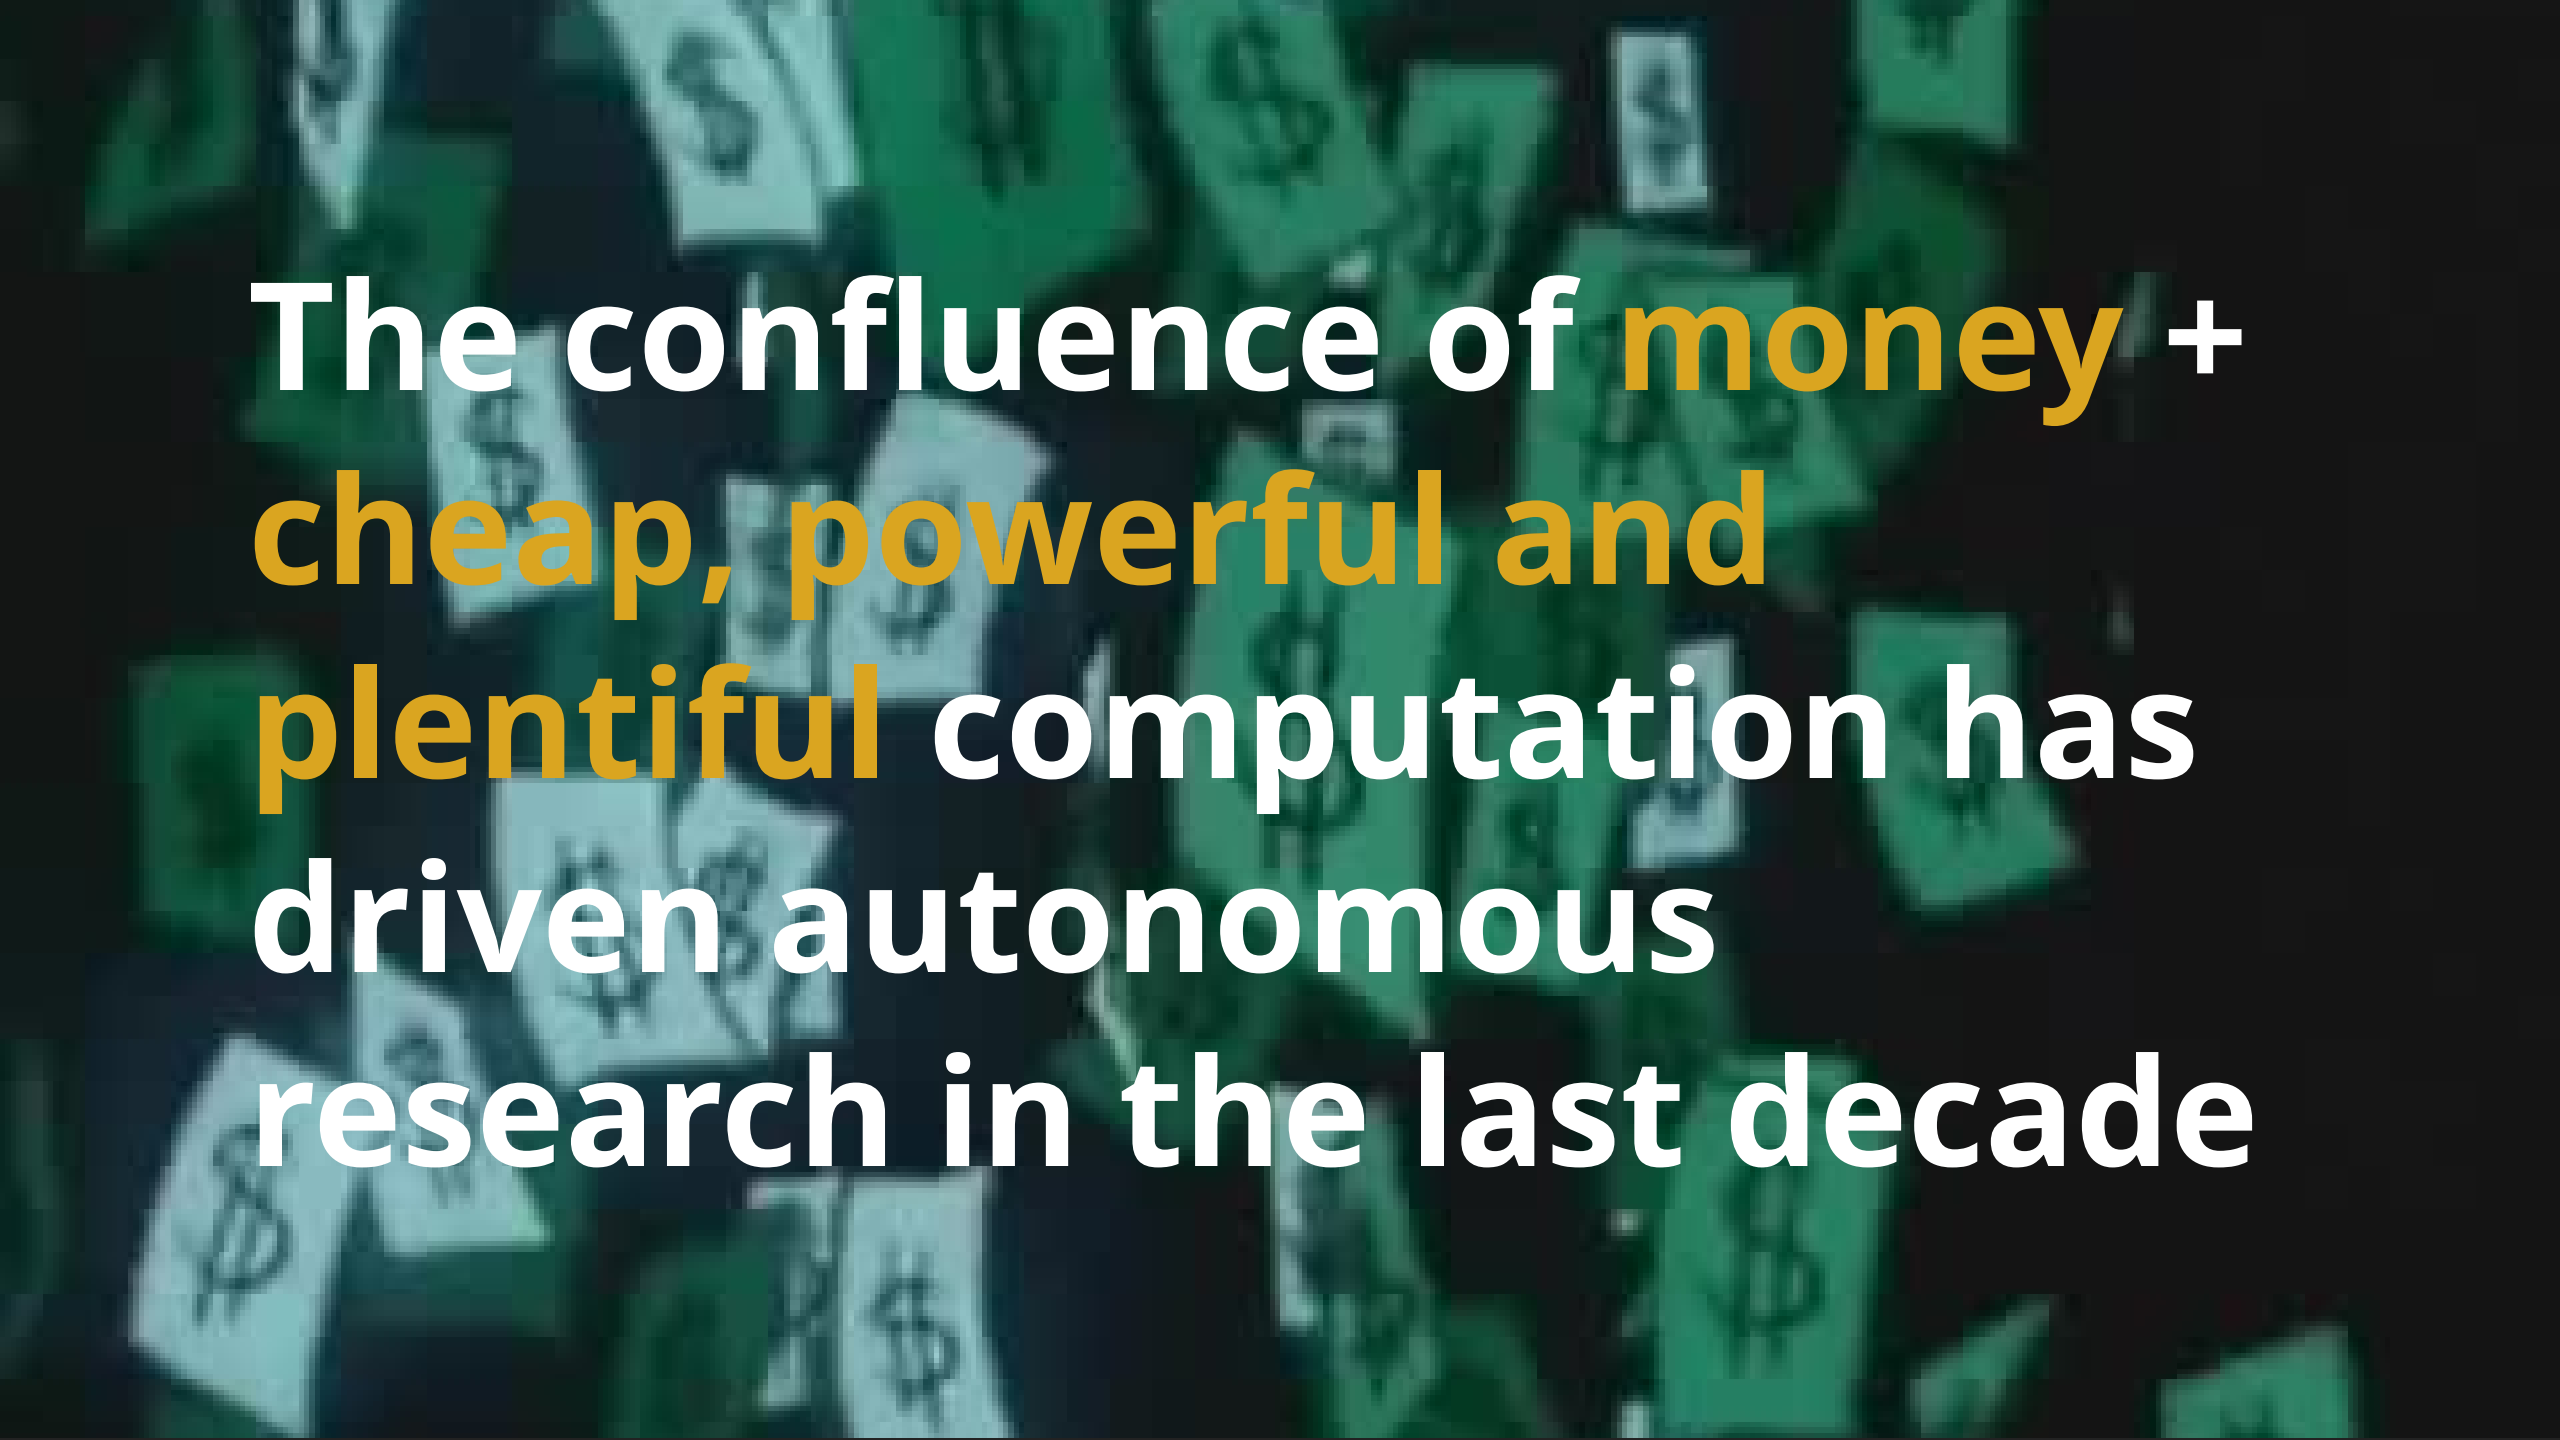
The confluence of money +
cheap, powerful and
plentiful computation has
driven autonomous
research in the last decade

The two prevailing methodologies

* From my point-of-view

Hierarchical Structures

VS

An End-to-end solution

Hierarchical Structures

- A pipeline of sub-systems which make up an autonomous system
- Decisions may be encoded
- Attribution of decisions
- Ability to apply multiple competing methods for better results

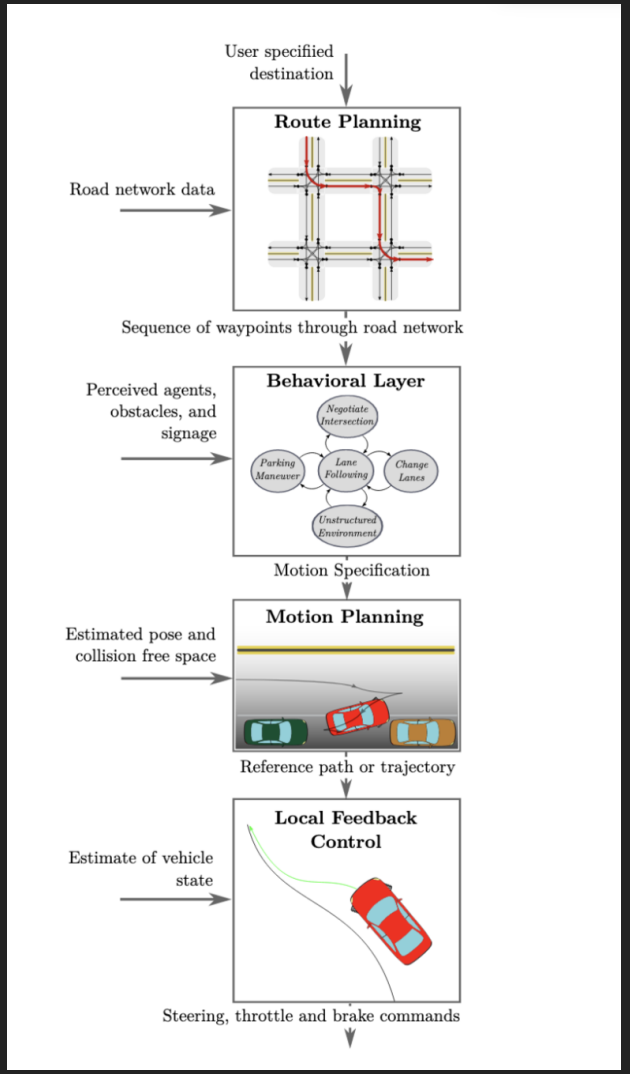

End-to-end

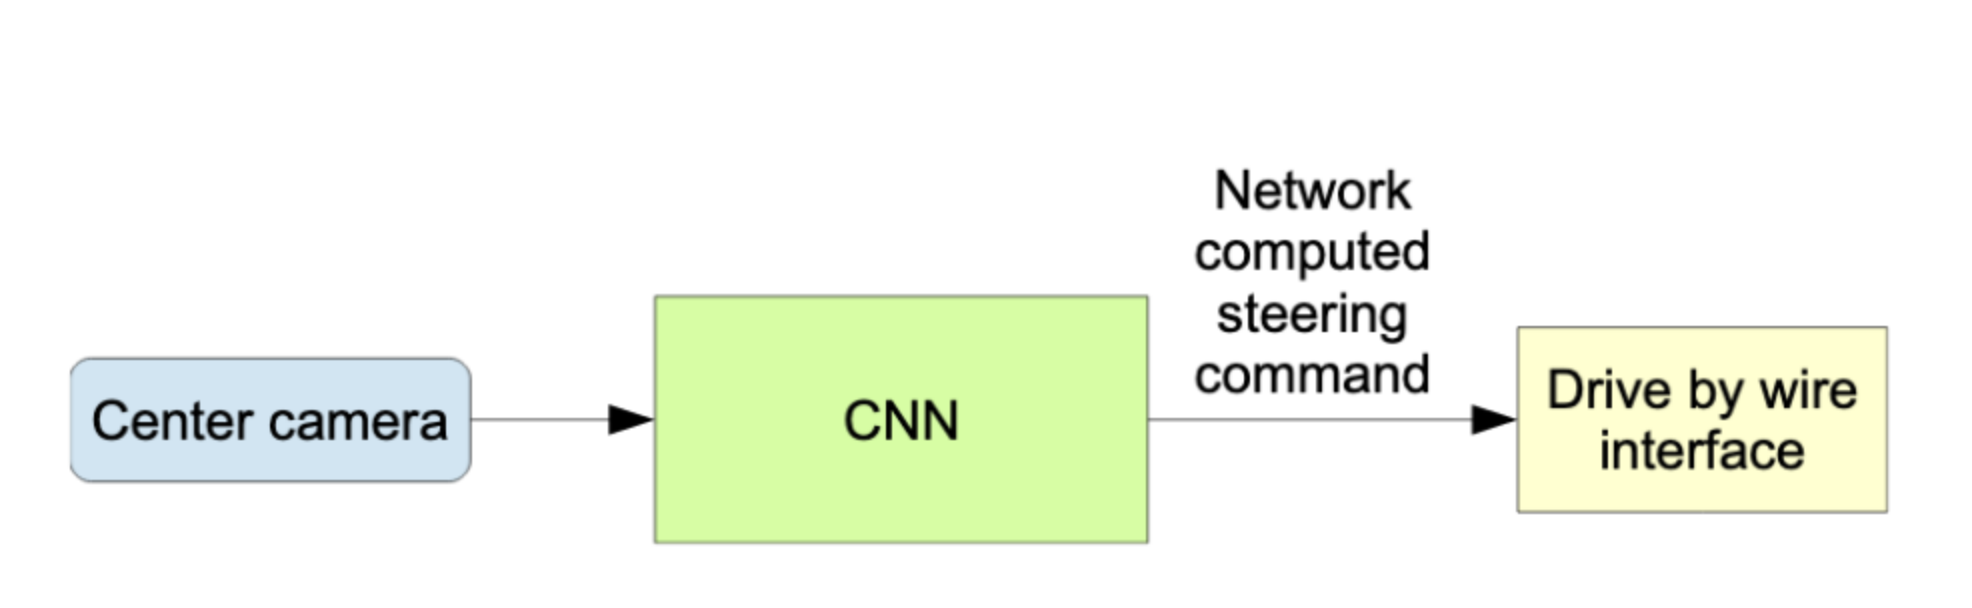

- One monolithic system which solves the problem as a whole
- Sometimes is a monolith within a greater system
- Lots of Deep learning
- Very "quick" results
- Hard to attribute decisions
- Behaviour has to be encoded into training data
- May learn invalid or undesirable behaviours


What Ive Been Looking At

Data Sets

More than 50 Open Data Sets

https://medium.com/datadriveninvestor/the-50-best-public-datasets-for-machine-learning-d80e9f030279

Berkeley DeepDrive BDD100k

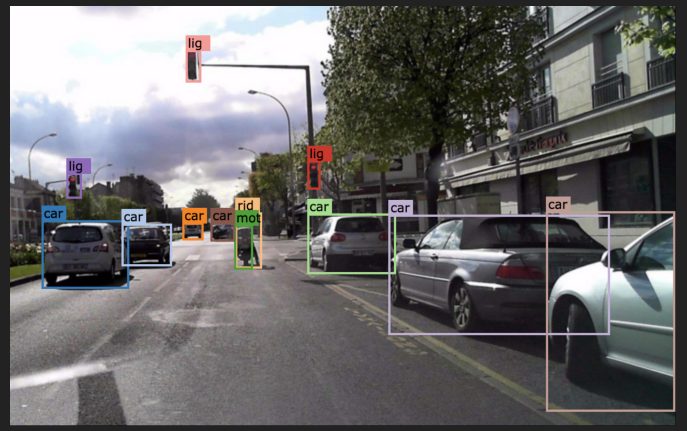

- 100000 videos (+1100 hours)
- Segmentation
- Labelling
- Path maps
- Sample data
- Other data?
- ~1.6 - 2TB

Baidu Apolloscapes

- Contains 26 distinct semantic items
- Divided into 4 categories of data
    1. Scene sparsing
    2. Car instance
    3. Lane segmentation
    4. Self localisation
    5. Stereo Data
    6. Detection / Tracking
    7. Trajectory

Comma.ai

- ~7 hours of highway driving
- Contains data on speed of car, acceleration, steering angle, gps coordinates

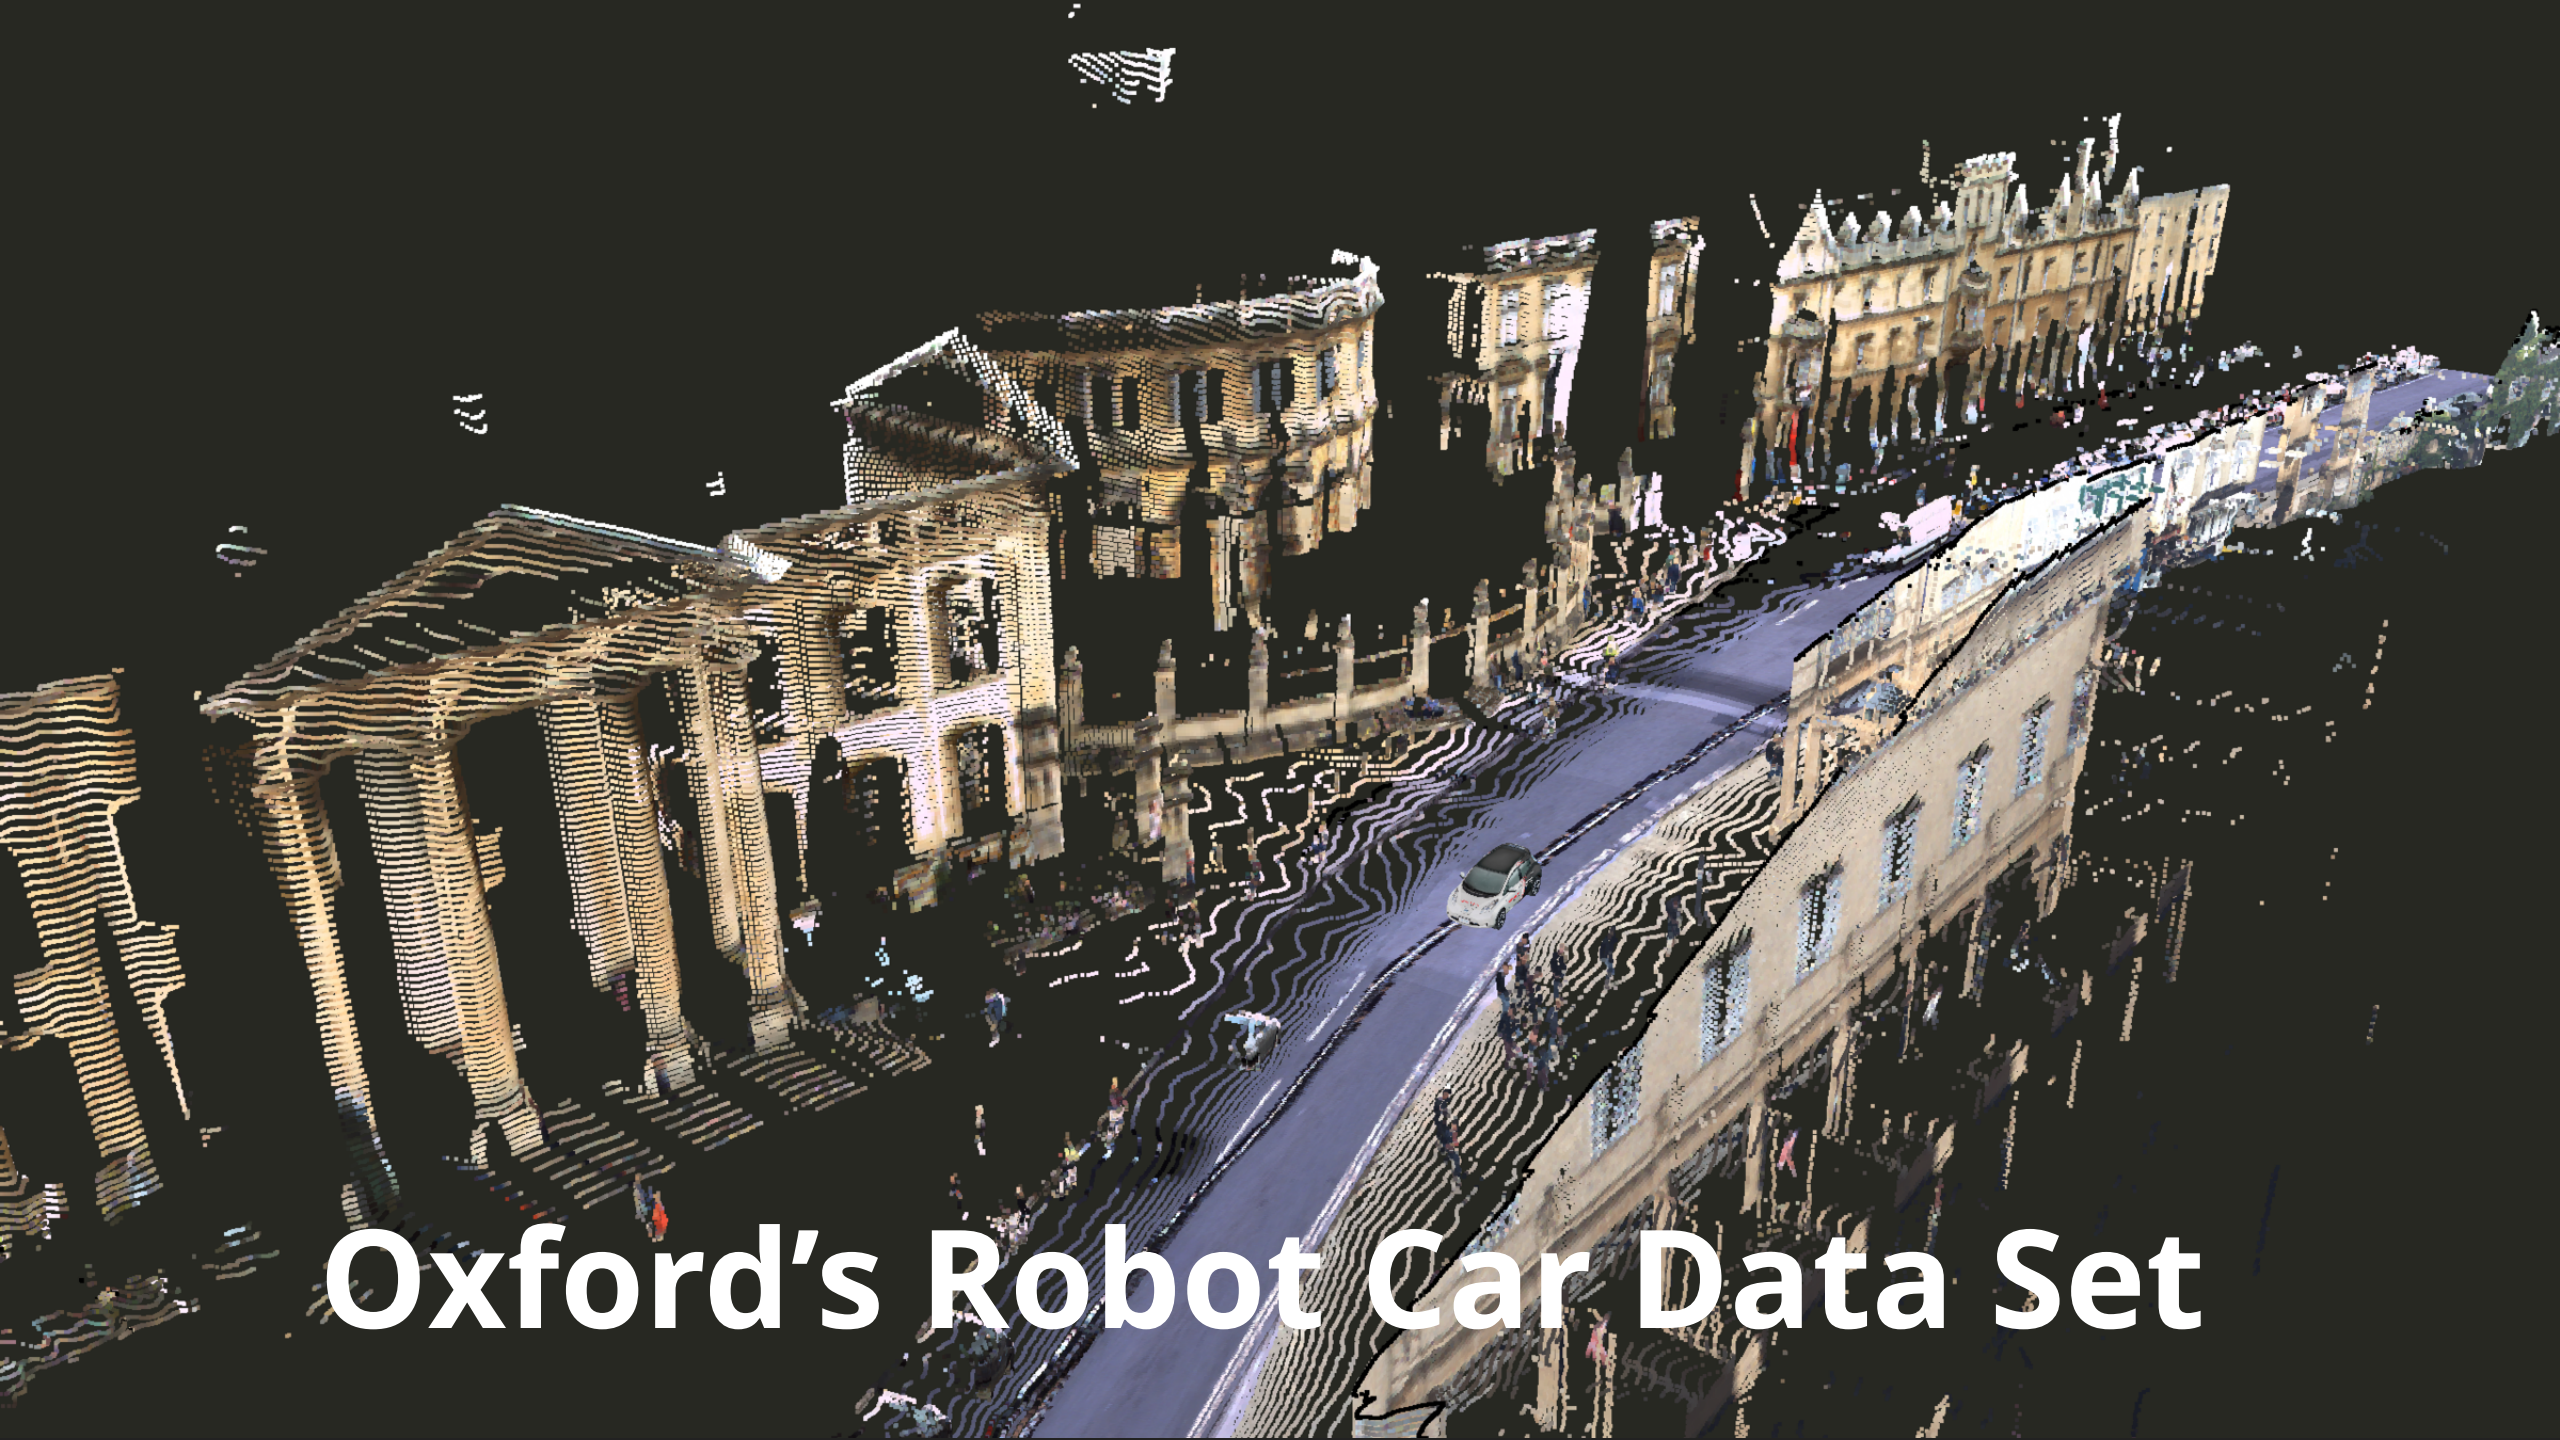

Oxford's Robot Car Data Set

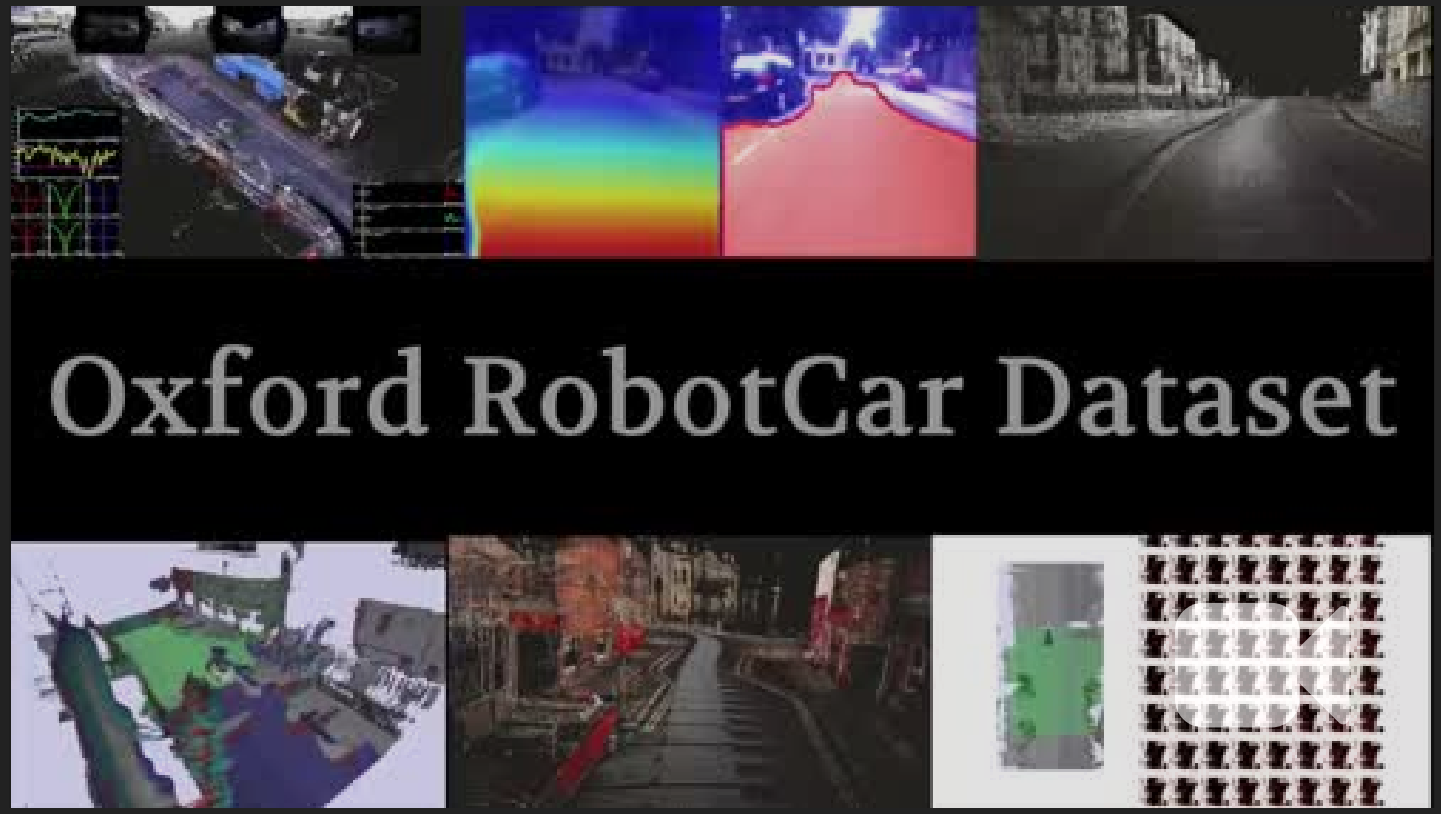

- Over 100 repetitions of the same route
- Different combinations of weather, traffic and pedestrians
- Also includes construction (changing landscape) and road-works
- Stereo Images, LIDAR and other specialised sensors

Honourable Mentions

- **Cityscapes**
- **Udacity Self-Driving Car Data Sets**

The older papers from around a decade ago +, all stated that their limiting factor was access to enough expressive data. Not really a problem anymore.

Simulators

AirSim

Realistic environments

Many environments

Can run headless

Many sensors

Choose time of day

*Traffic simulation

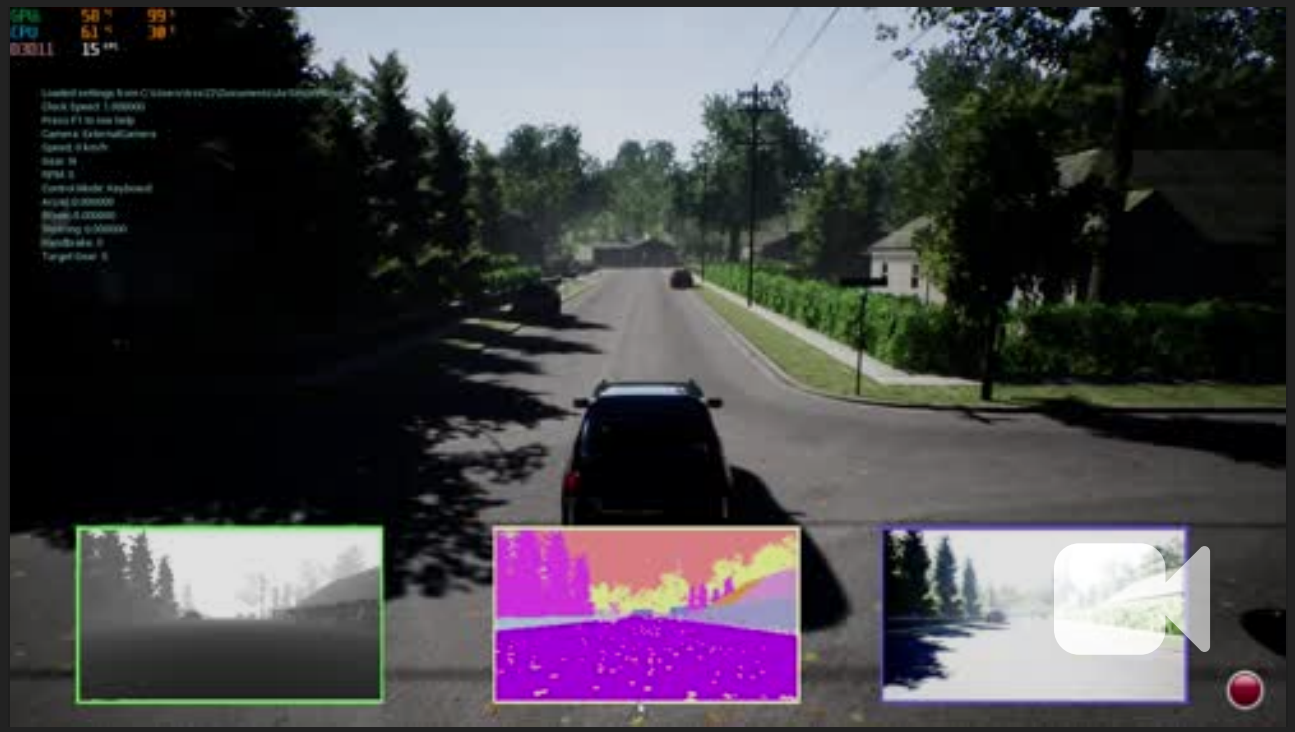

CARLA

Warrants another look

Headless

Has traffic scenarios now

Map editor

Has development road-map

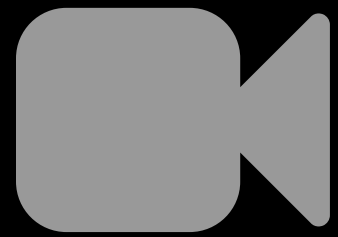

The ethics of this research (abridged)

Unpacking

- **Autonomous vehicles have created lots of ethical quandaries**
- **There are issues with replacing people's jobs**
- **The cost of research**
- **The changing of industry**
- **Dangers of Driving**
- **Attribution of decisions**
- And many others ...

A Defensible Choice

Choices don't have absolute ethical and moral certainty

https://www.cs.uct.ac.za/mit_notes/ethics/htmls/ch02s07.html

But a rational choice
Is made with understanding
of the situation

The End-to-end approach

Makes it difficult to reason

Especially when things go wrong

Attribution of decisions in deep learning

What features Do Networks Observe?

Many different methods

Saliency Maps Were Big (Now Broken)

VisualBackProp - has promise

DeepLearning against hidden layers

Feature maps

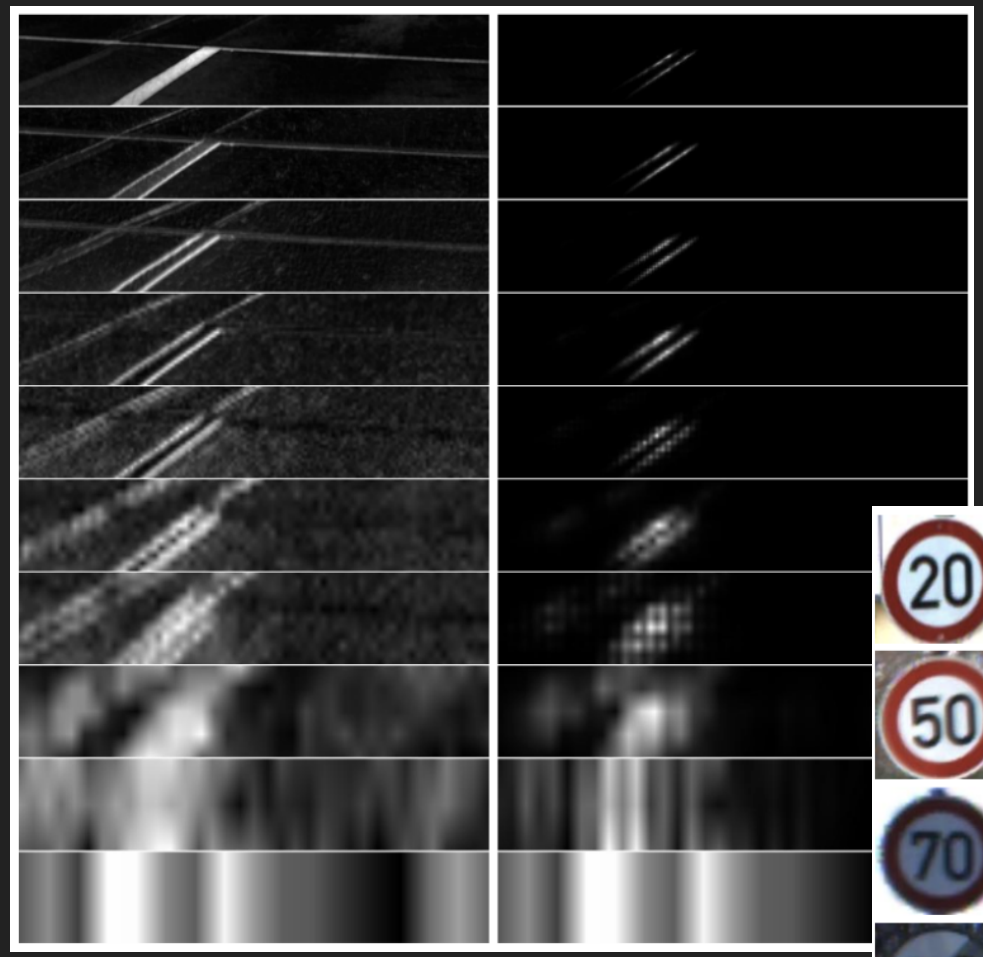

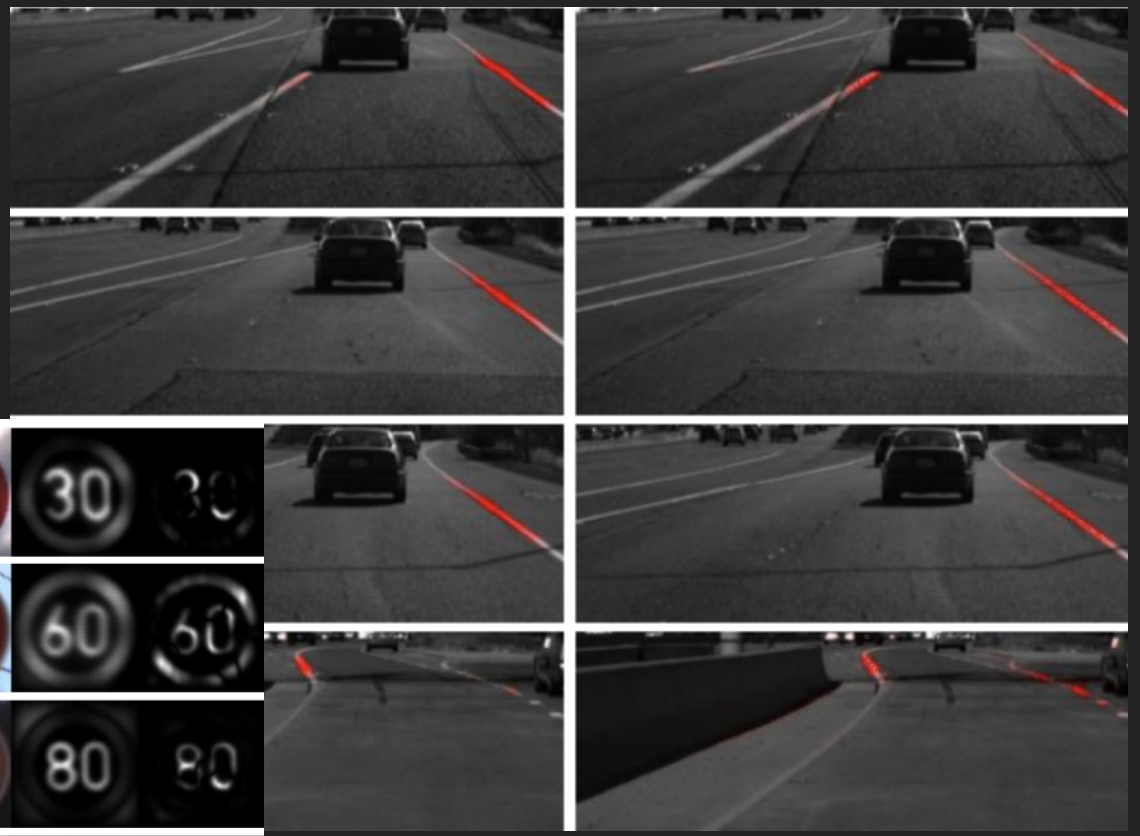

Progress Report

Progress Report

- Continued reading
- Found a lot of useful resources for the specific problem domain
- Moving prototype to PyTorch
- Starting to map out the skeleton of my proposal

Thank You.

References

- A Survey of Motion Planning and Control Techniques for Self-driving Urban Vehicles, Brian Paden∗,1, Michal Cˇáp∗,1,2, Sze Zheng Yong1, Dmitry Yershov1, and Emilio Frazzoli1, arxiv.org/abs/1604.07446
- https://www.cs.uct.ac.za/mit_notes/ethics/htmls/ch02s07.html
- Visual SLAM for Driverless Cars: A Brief Survey, German Ros∗, Angel D. Sappa†, Daniel Ponsa∗and Antonio M. Lopez∗, Proceedings of the 2012 IEEE Intelligent Vehicles Symposium Workshops
- Survey of Pedestrian Detection for Advanced Driver Assistance Systems, David Geró nimo, Antonio M. Ló pez, Angel D. Sappa, Member, IEEE, and Thorsten Graf, July 2010, IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE
- **End to End Learning for Self-Driving Cars, Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, Karol Zieba, https://arxiv.org/abs/1604.07316**

References

- **AirSim: High-Fidelity Visual and Physical Simulation for Autonomous Vehicles, Shital Shah, Debadeepta Dey, Chris Lovett, Ashish Kapoor, https://arxiv.org/abs/1705.05065**
- VisualBackProp- efficient visualization of CNNs, Mariusz Bojarski, Anna Choromanska, Krzysztof Choromanski, Bernhard Firner, Larry Jackel, Urs Muller, Karol Zieba, https://arxiv.org/abs/1611.05418
- VisualBackProp: visualizing CNNs for autonomous driving, Mariusz Bojarski, Anna Choromanska, Krzysztof Choromanski, Bernhard Firner, Larry Jackel, Urs Muller, Karol Zieba
- CARLA: An Open Urban Driving Simulator, Alexey Dosovitskiy and German Ros and Felipe Codevilla and Antonio Lopez and Vladlen Koltun, Proceedings of the 1st Annual Conference on Robot Learning
- https://en.wikipedia.org/wiki/History_of_self-driving_cars
- https://www.titlemax.com/resources/history-of-the-autonomous-car/
- Oxford Robot Car Data Set, https://robotcar-dataset.robots.ox.ac.uk/

I also made use of giphy.com and wallhaven.cc