Self Driving Cars

A study in deep learning applications

Chris Foster

What and why?

Traffic kills 1.25 million people worldwide per year

That's about one person every 25 seconds

Self driving cars are an application of deep learning that can save lives!

Moravec's paradox

"It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility"

Challenges...

- Real time requirements
- Sparse training data types
- Reflections, light changes
- Unpredictable pedestrians
- Construction zones
- Parking garages
- Weather

Human error: 0.000001%

...but it works!

Total miles driven in U.S. in 2014:
3,000,000,000,000

U.S crash fatalities:
32,675 *(1 in 90 million)*

Tesla Autopilot miles driven:
1,557,000,000

Total fatalities: 5

Technologies available

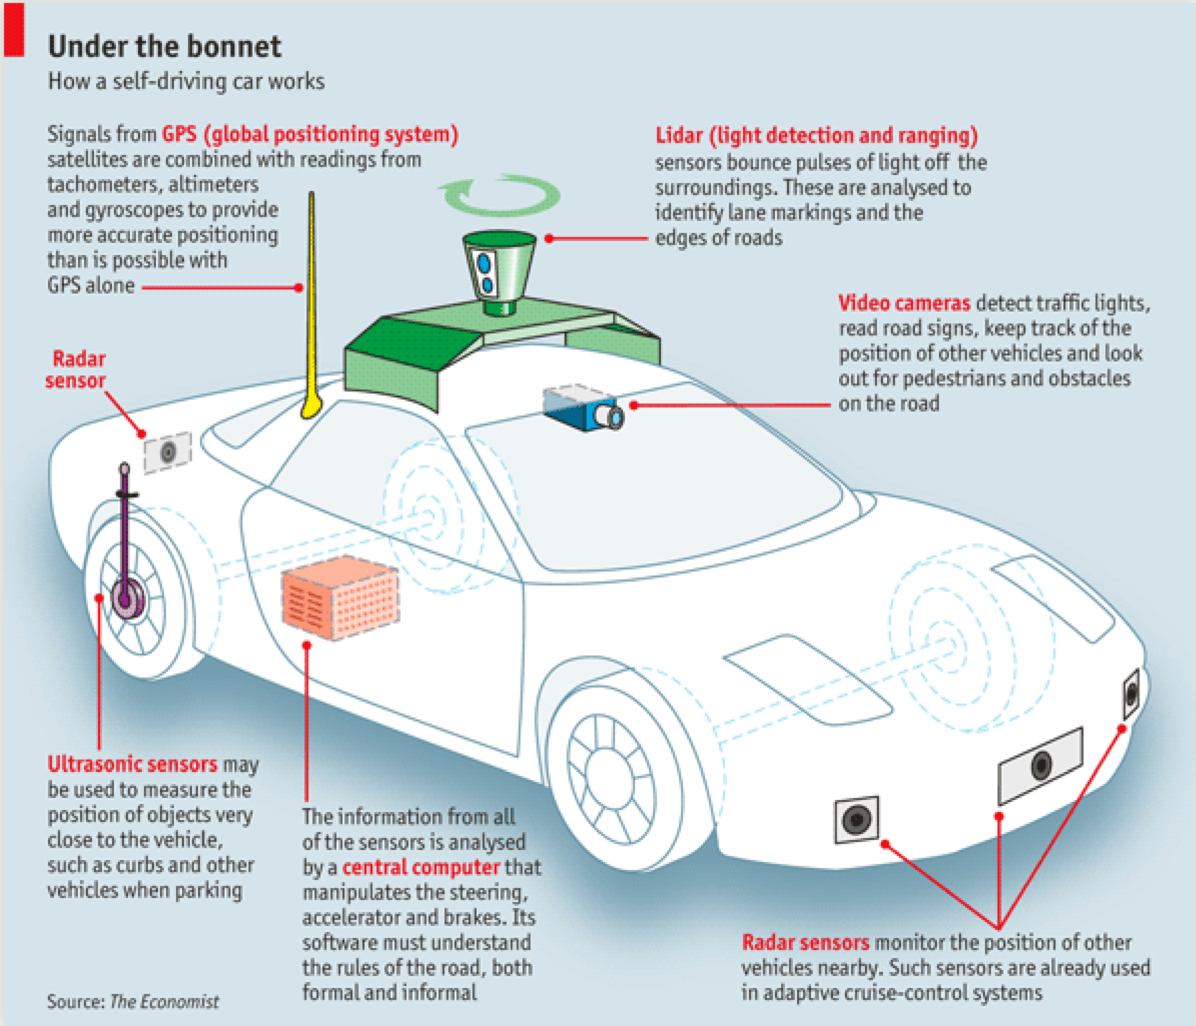

World in LIDAR

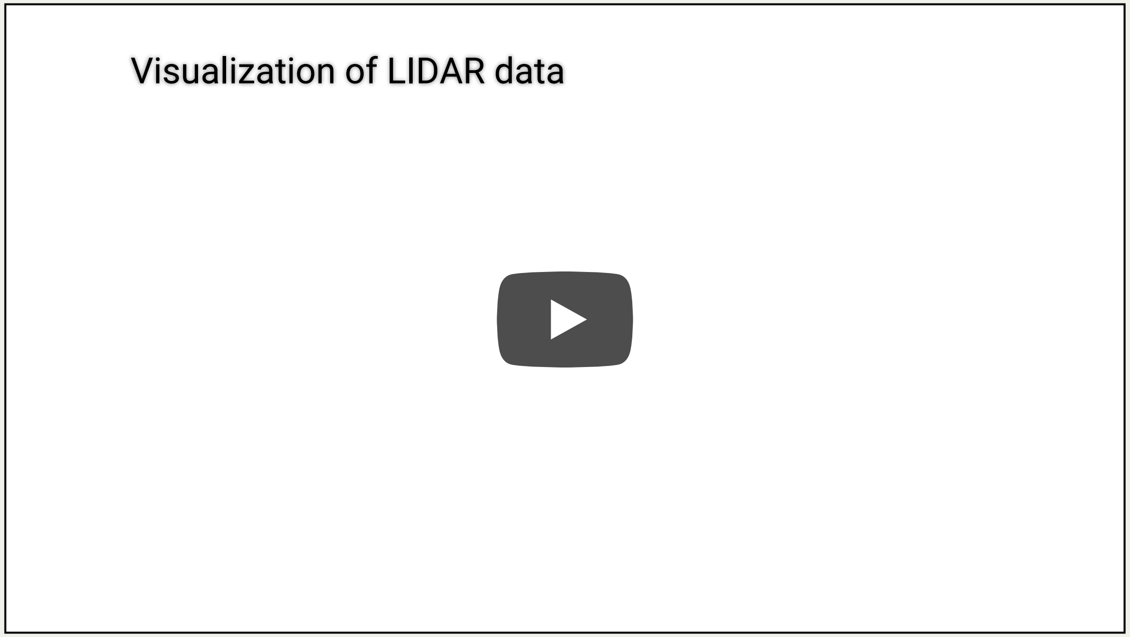

(~1GB of data per second)

Deep Learning!

These are the tools we need:

- Convolutional Neural Networks
- Recurrent Neural Networks
- Reinforcement Learning

Traditional Neural Networks

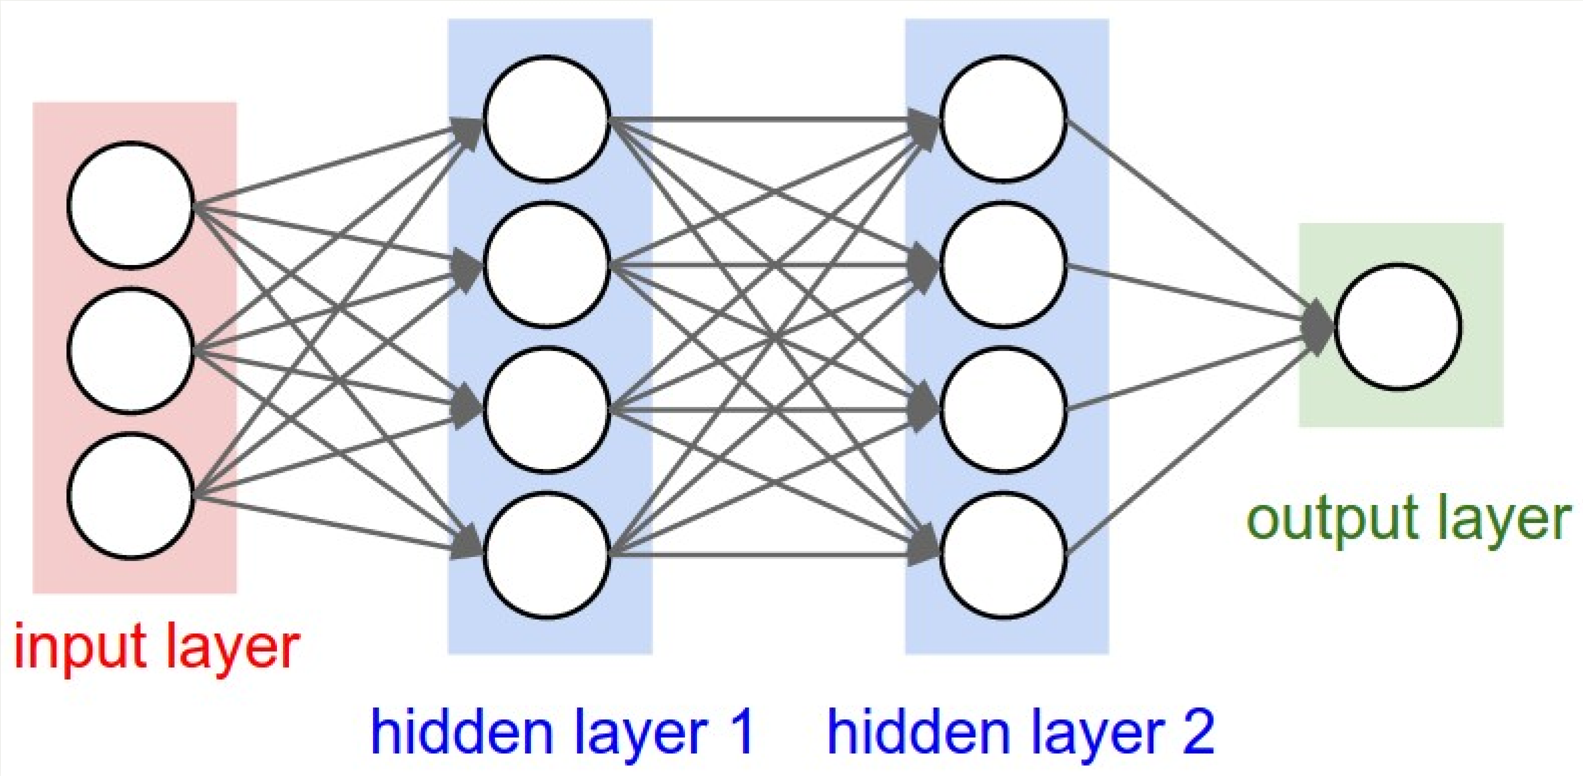

Convolutional Neural Network

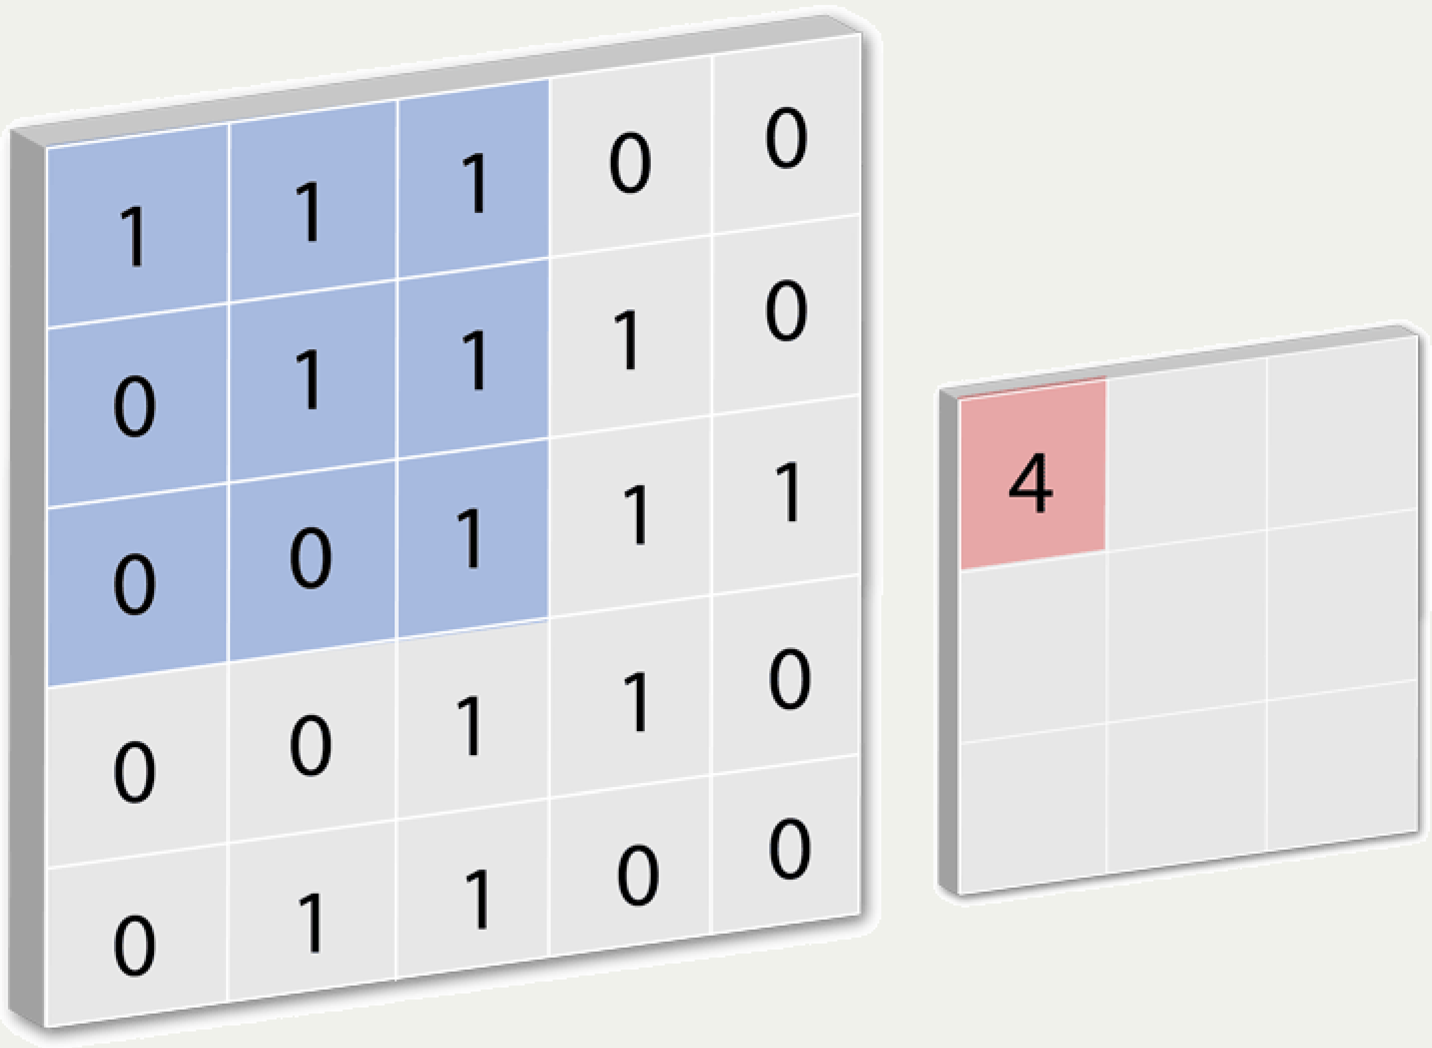

Recurrent Neural Network

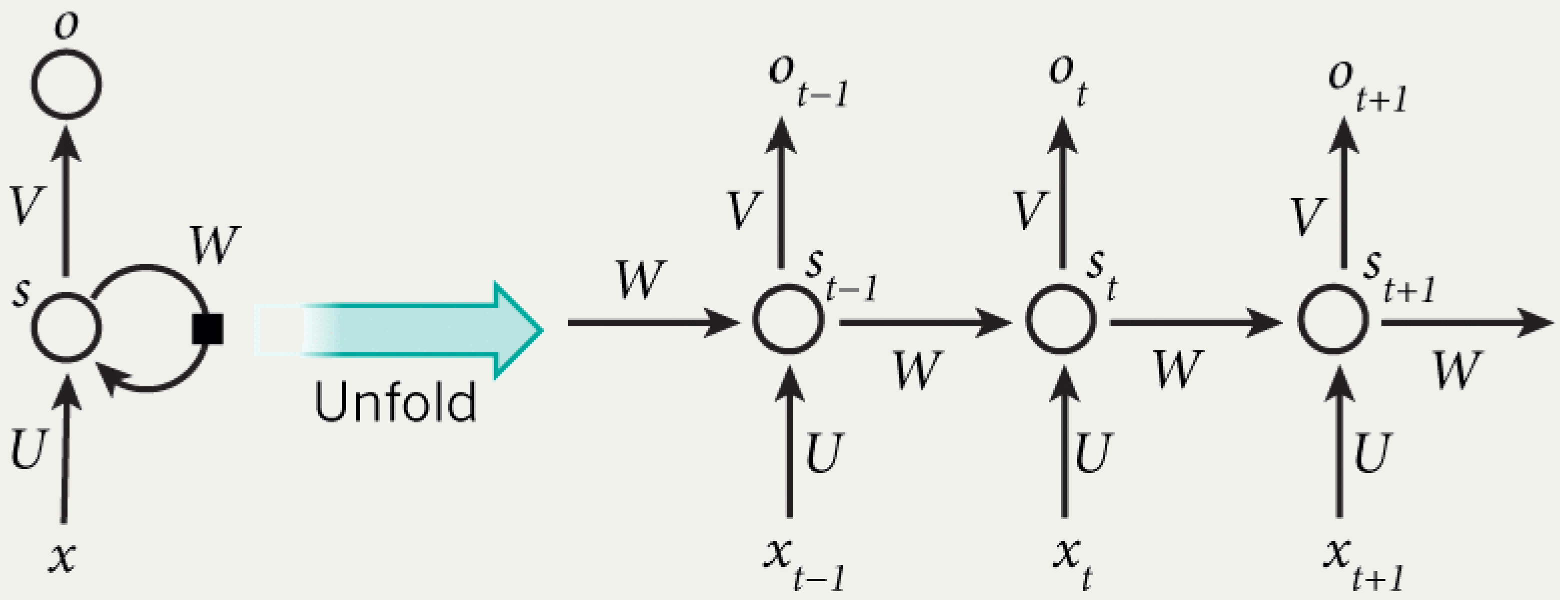

Reinforcement Learning

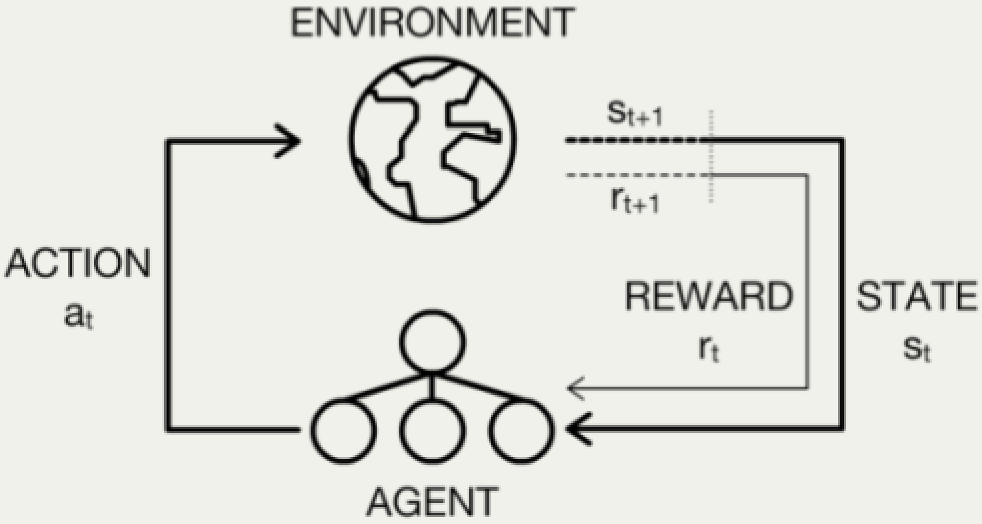

How much deep learning?

Two approaches to utilizing DL:

1. Task-orientated deep learning
2. End-to-end deep learning

Google Self Driving Car

Task-orientated

Use deep learning for specific tasks:

- Localization and mapping
- Scene understanding
- Movement planning
- Driver state

Task-orientated

Use deep learning for specific tasks:

- Localization and mapping
- Scene understanding
- Movement planning
- Driver state

Visual Odometry

100 fps & Very Low Drift Visual Odometry - New College Da...

(Cars also have GPS/accelerometer/etc..)

Traditional Approaches

1. Undistortion/rectification
2. Disparity map computation
3. Feature detection (SIFT)
4. Feature tracking (KLT)
5. Trajectory estimation

Deep Learning Approaches

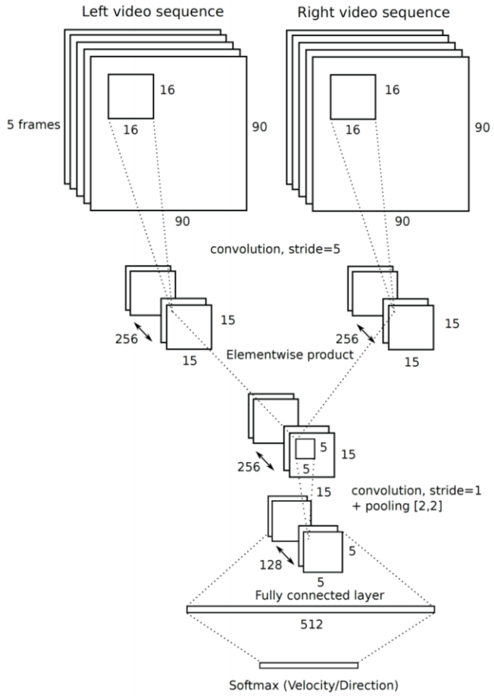

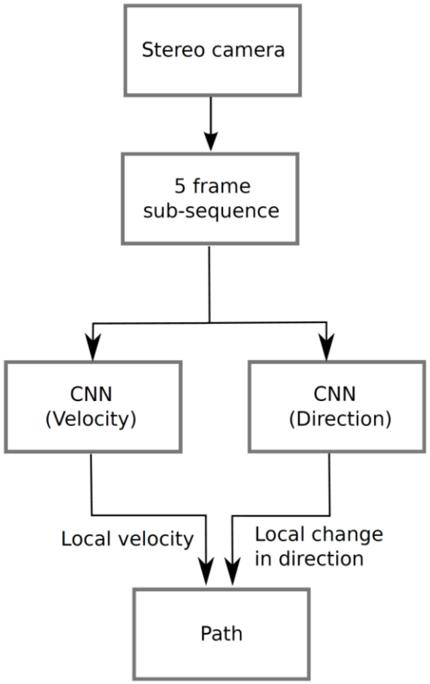

Konda, Kishore, and Roland Memisevic. "Learning visual odometry with a convolutional network." *International Conference on Computer Vision Theory and Applications*. 2015.

Task-orientated

Use deep learning for specific tasks:

- Localization and mapping
- Scene understanding
- Movement planning
- Driver state

SegNet: Road Scene Segmentation

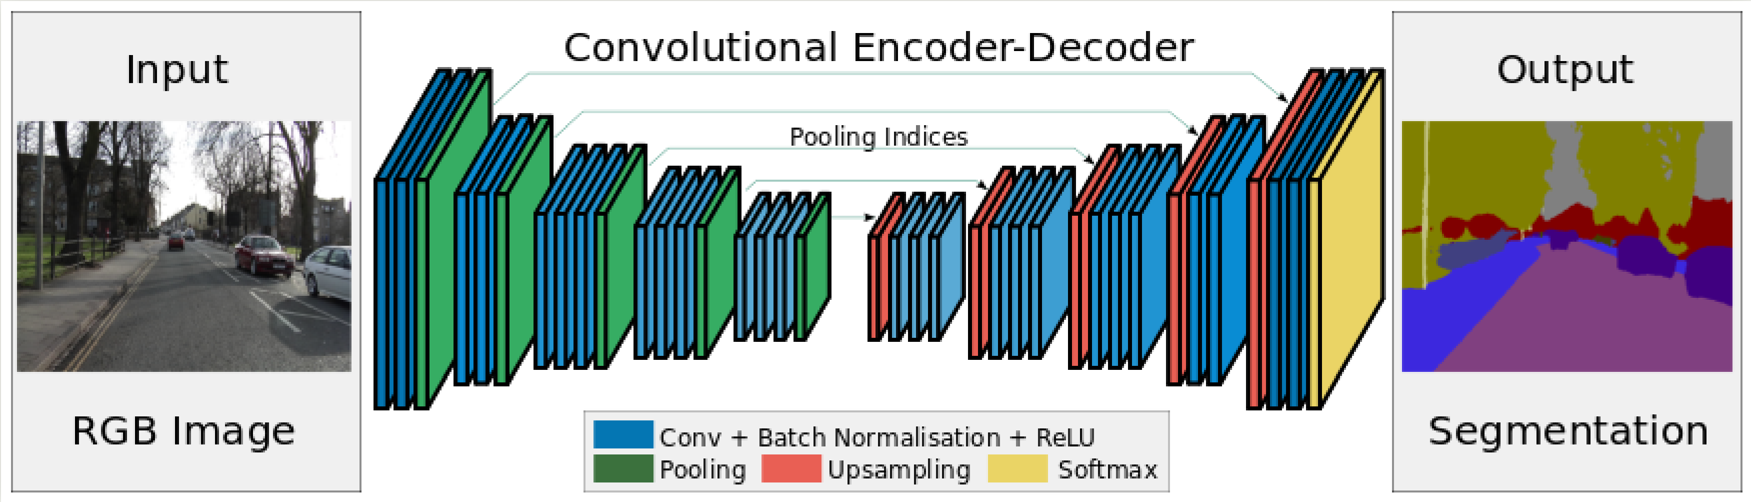

SegNet: Road Scene Segmentation

LIDAR object classification

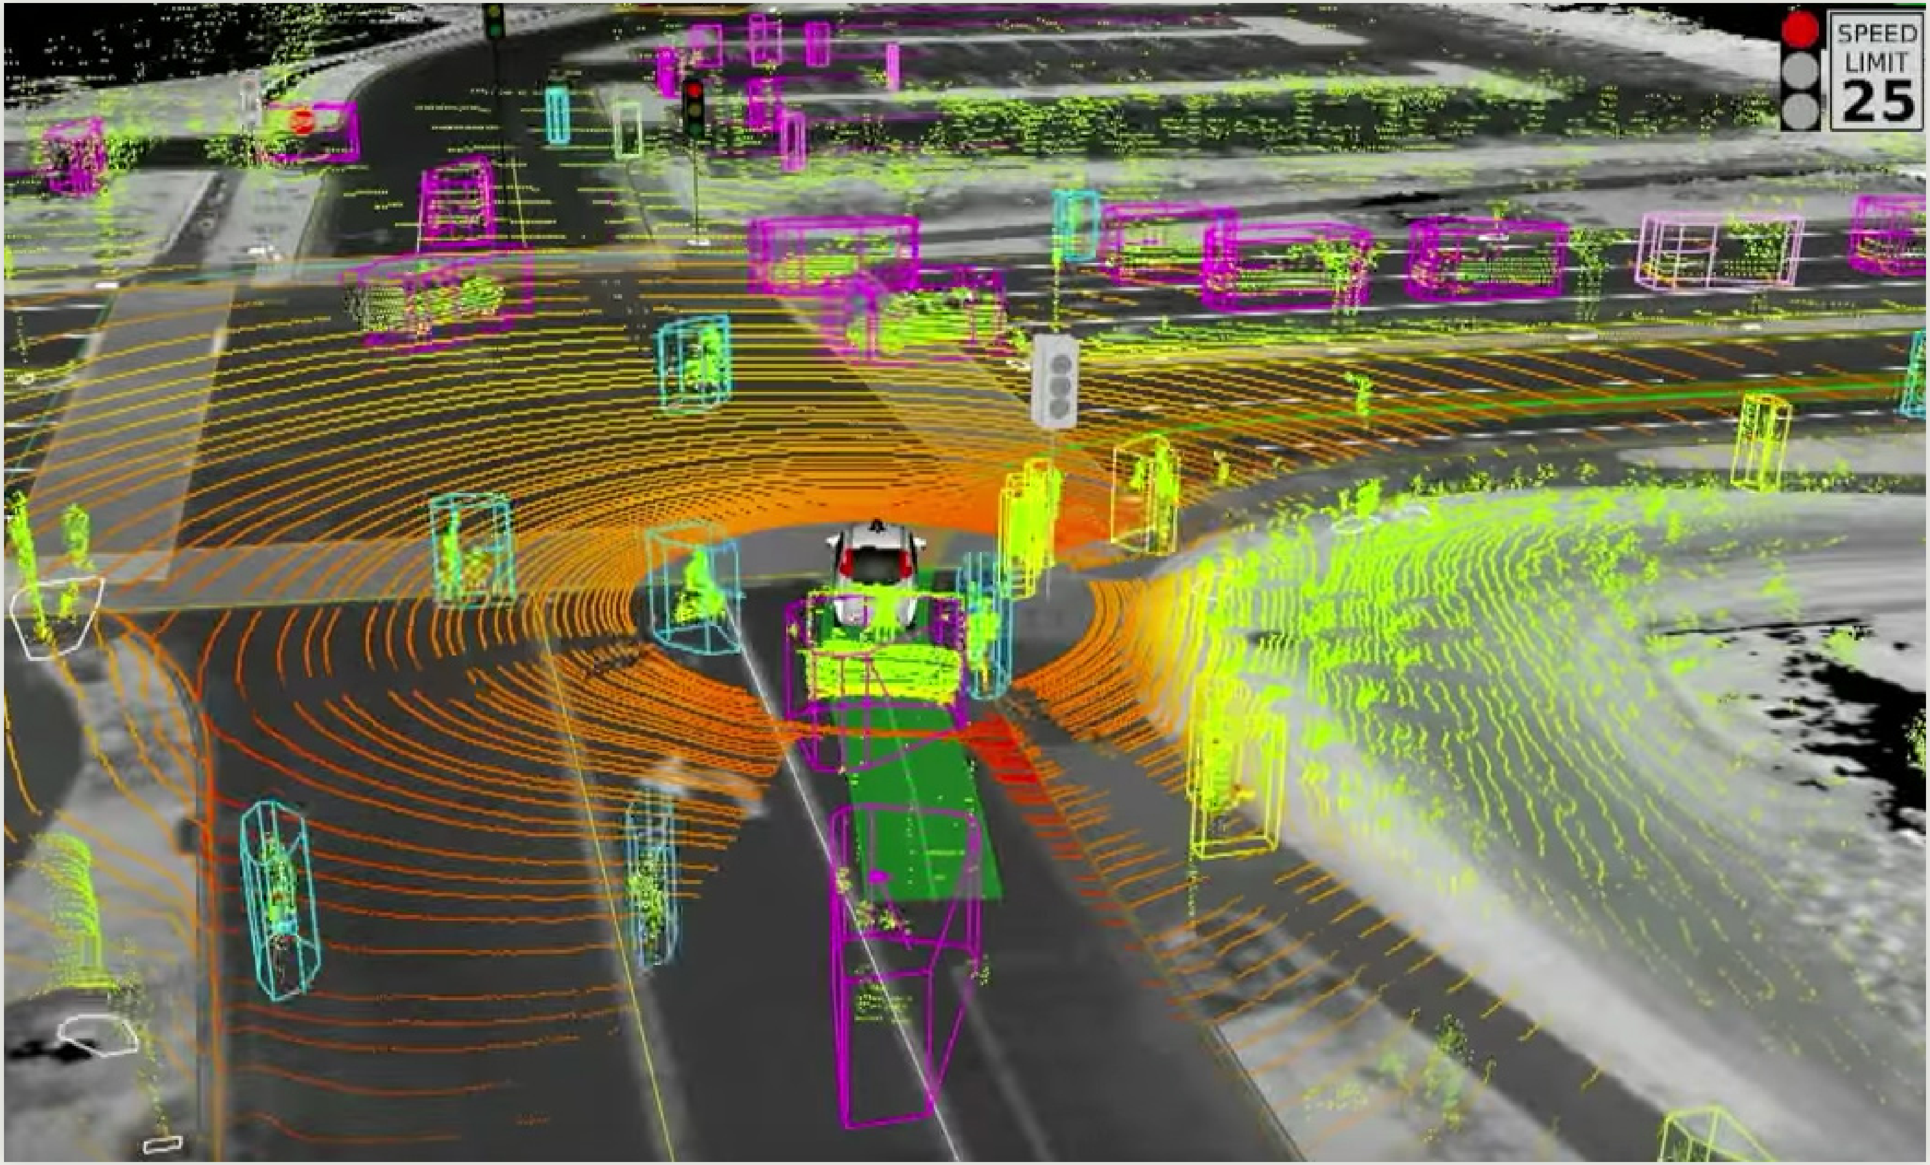

Object identification

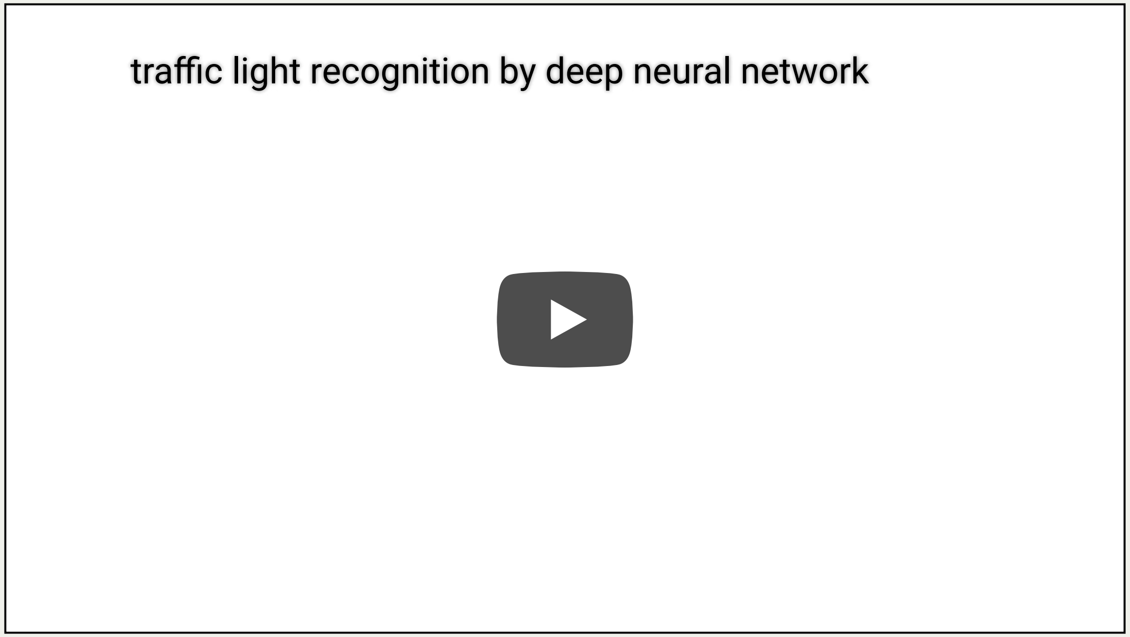

Wet road classification

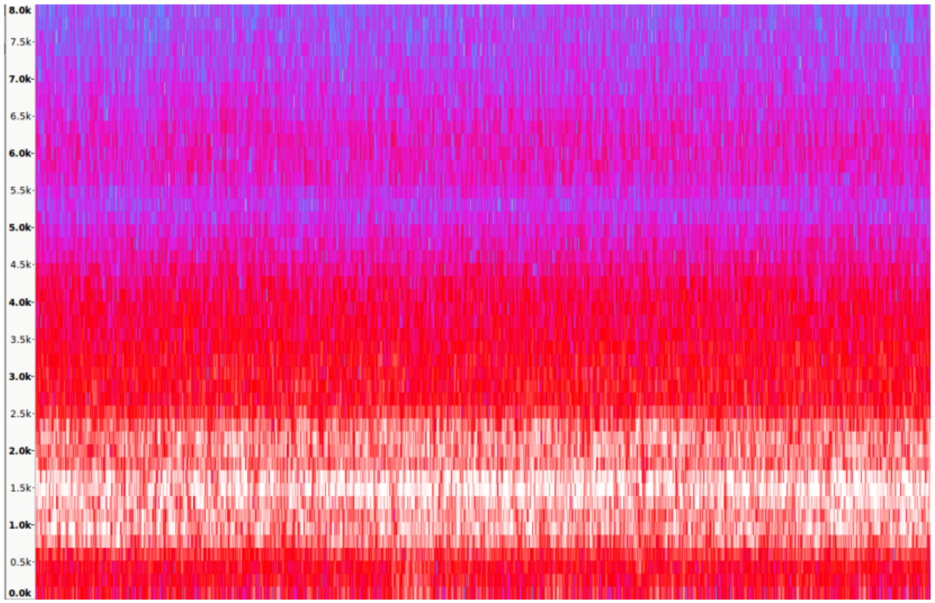

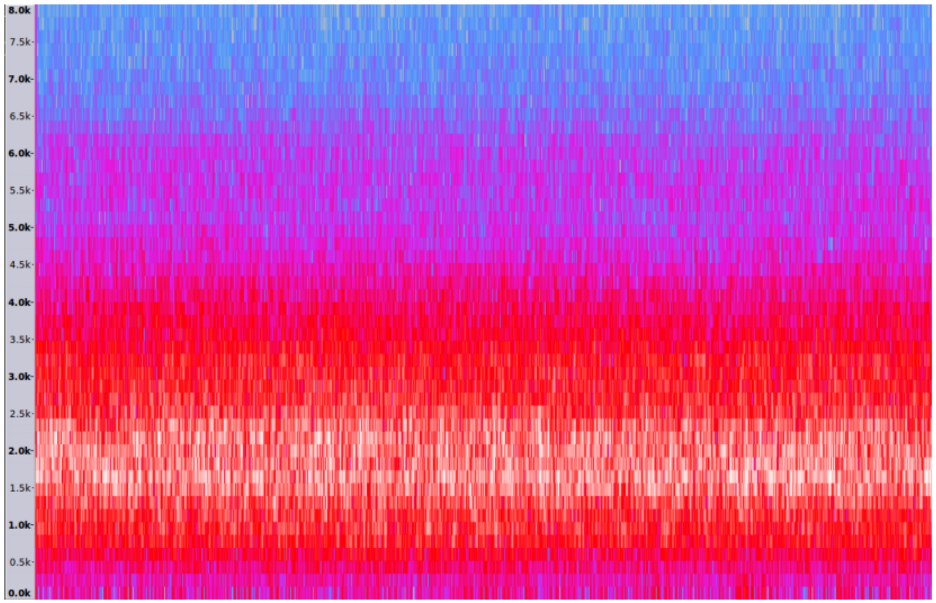

Task-orientated

Use deep learning for specific tasks:

- Localization and mapping
- Scene understanding
- Movement planning
- Driver state

Movement planning problems:

- Optimal lane to use
- How far/close to drive to others
- Poorly defined conditions

Deep Traffic

Distance

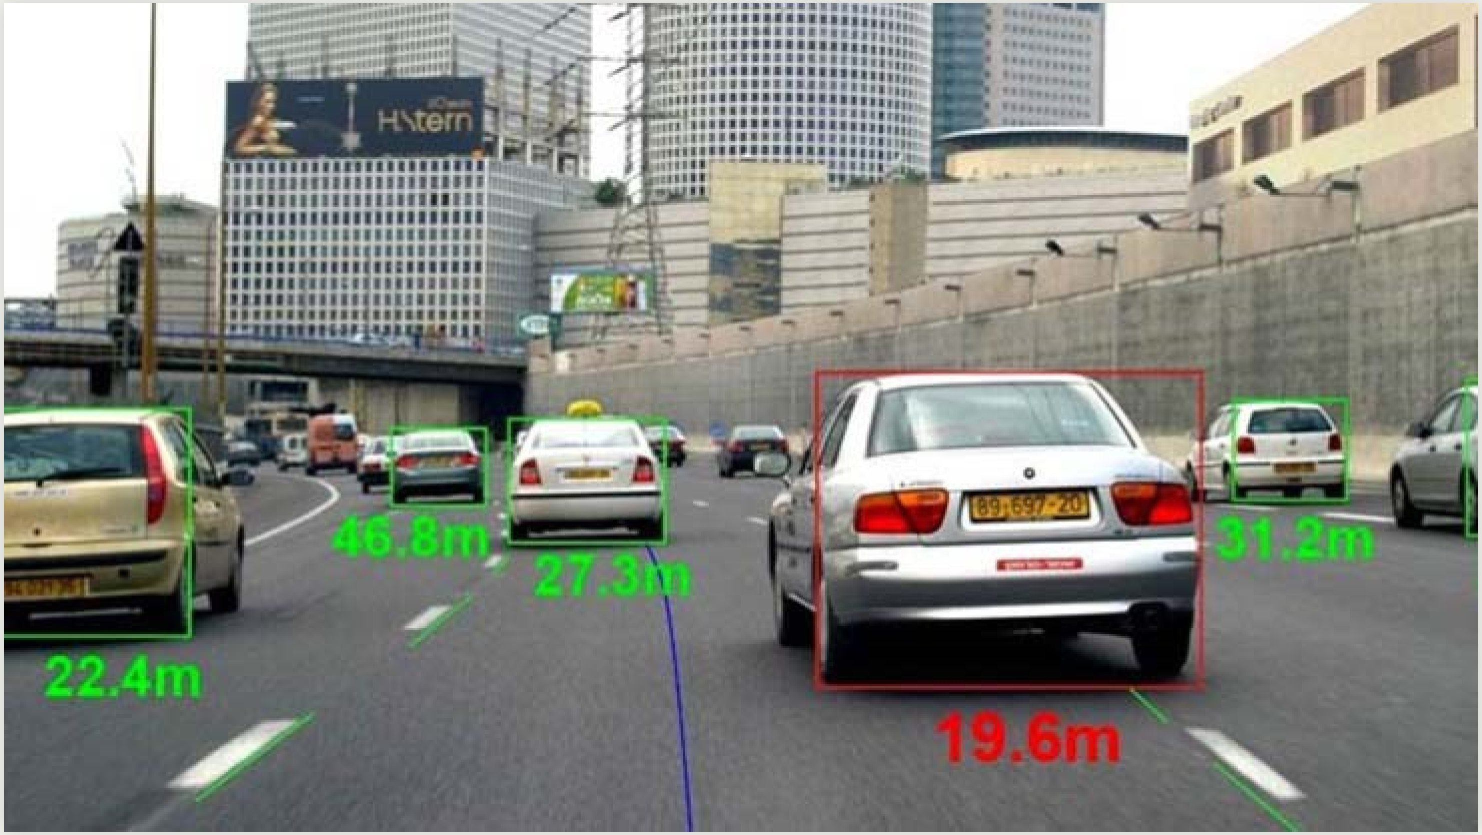

Object avoidance

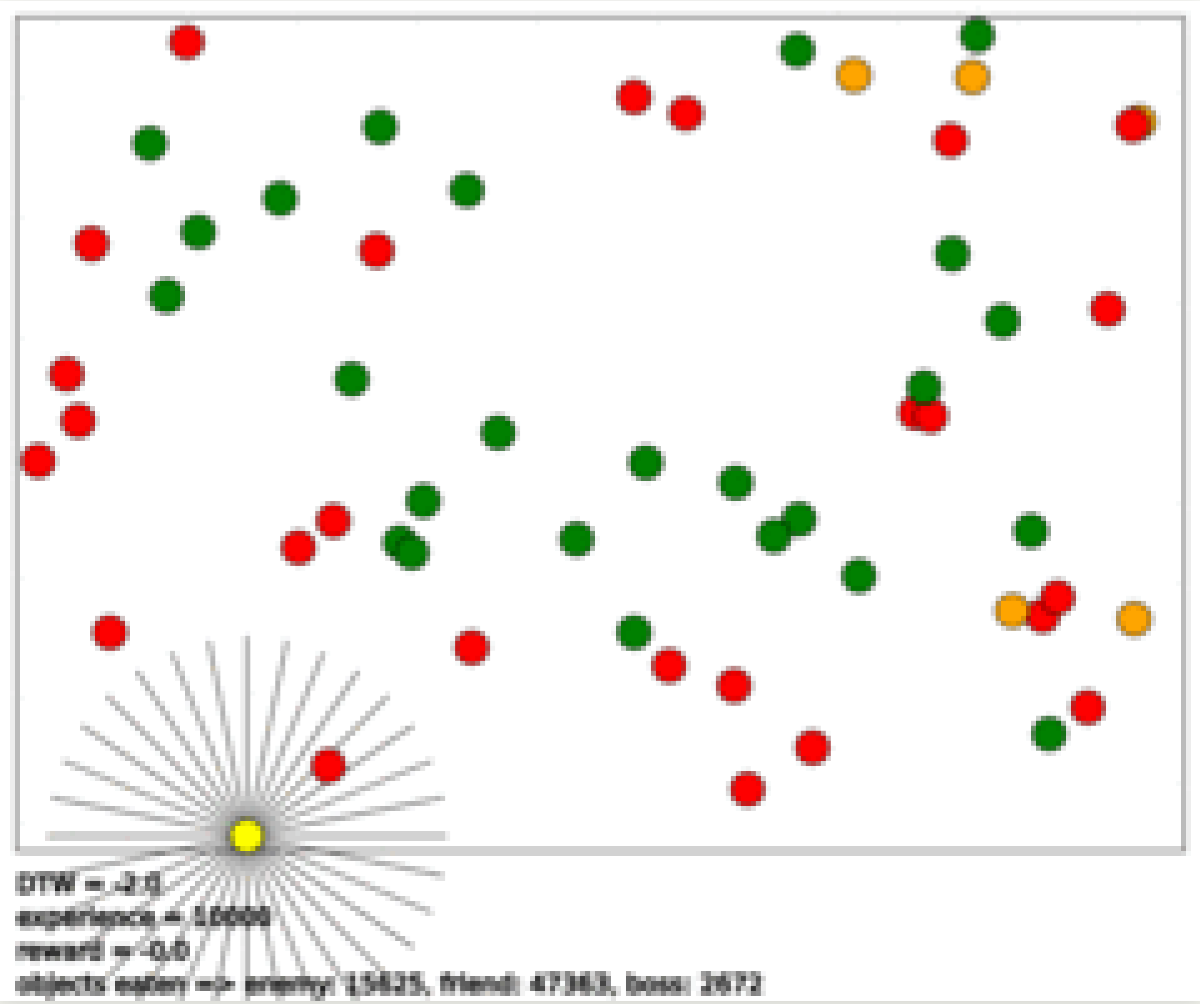

Task-orientated

Use deep learning for specific tasks:

- Localization and mapping
- Scene understanding
- Movement planning
- Driver state

NHTSA Car Classifications

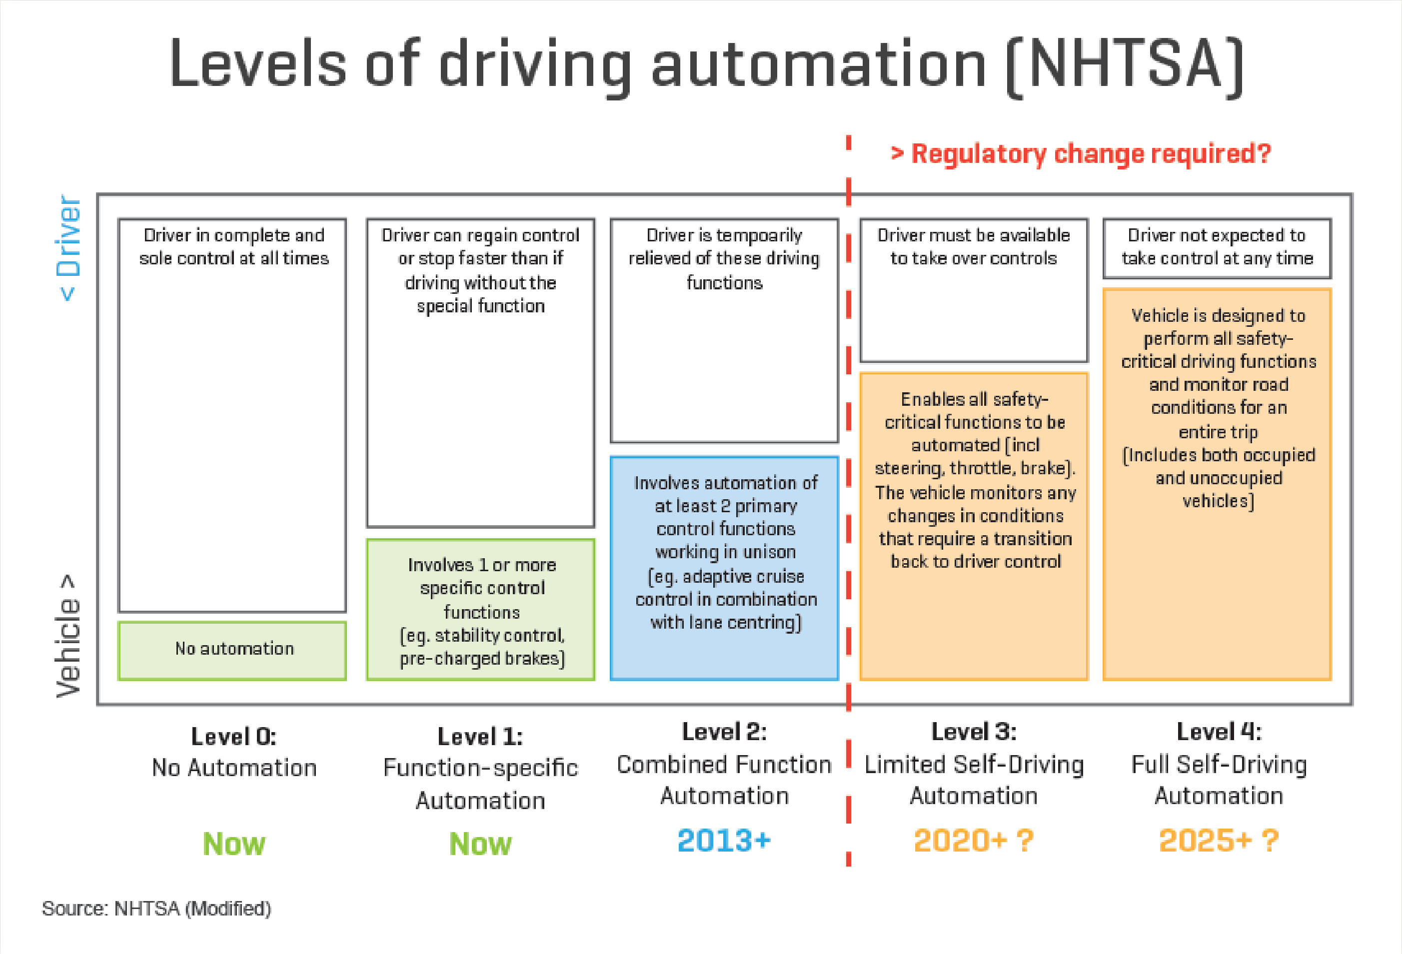

Building Blocks

Self driving cars are going to be an incremental process

Some people love to drive their own cars!

Can deep learning still help?

Where is the driver looking?

How is the driver sitting?

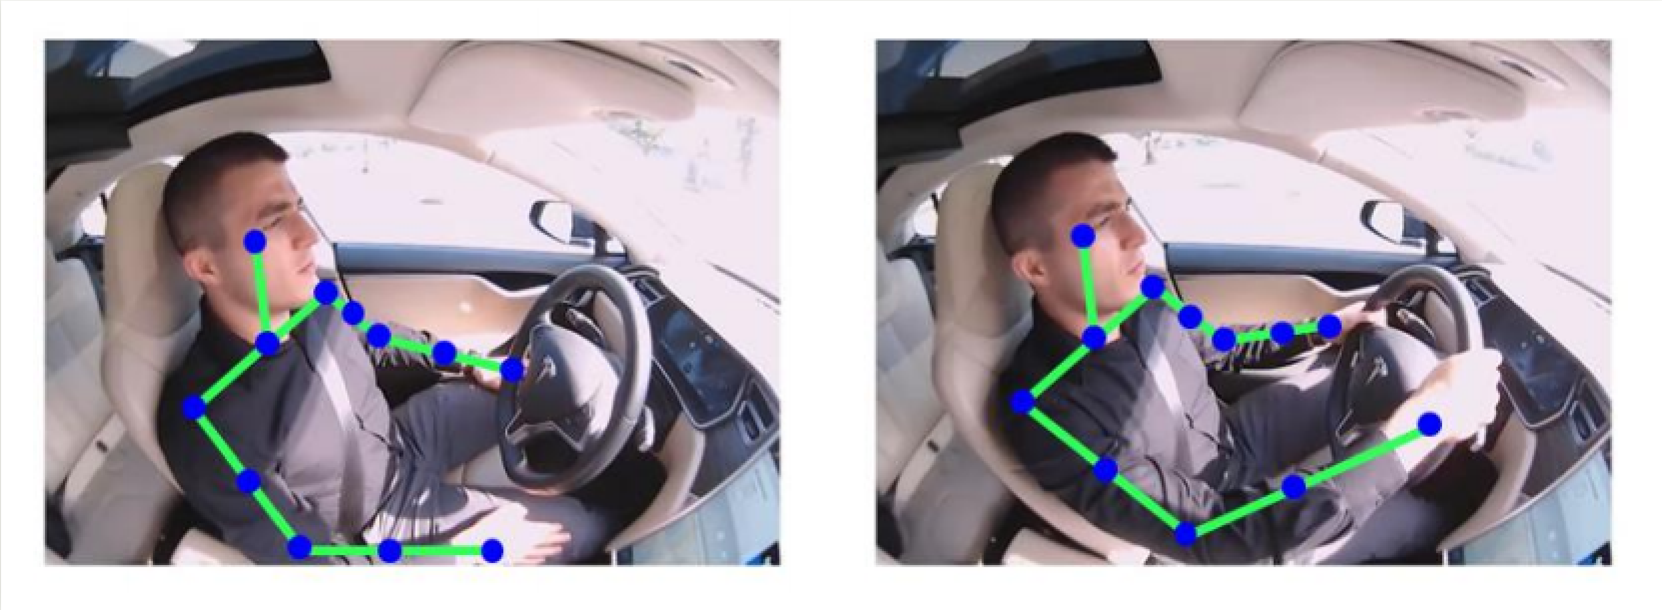

Is the driver tired?

NVIDIA Self Driving Car

End to End Deep Learning

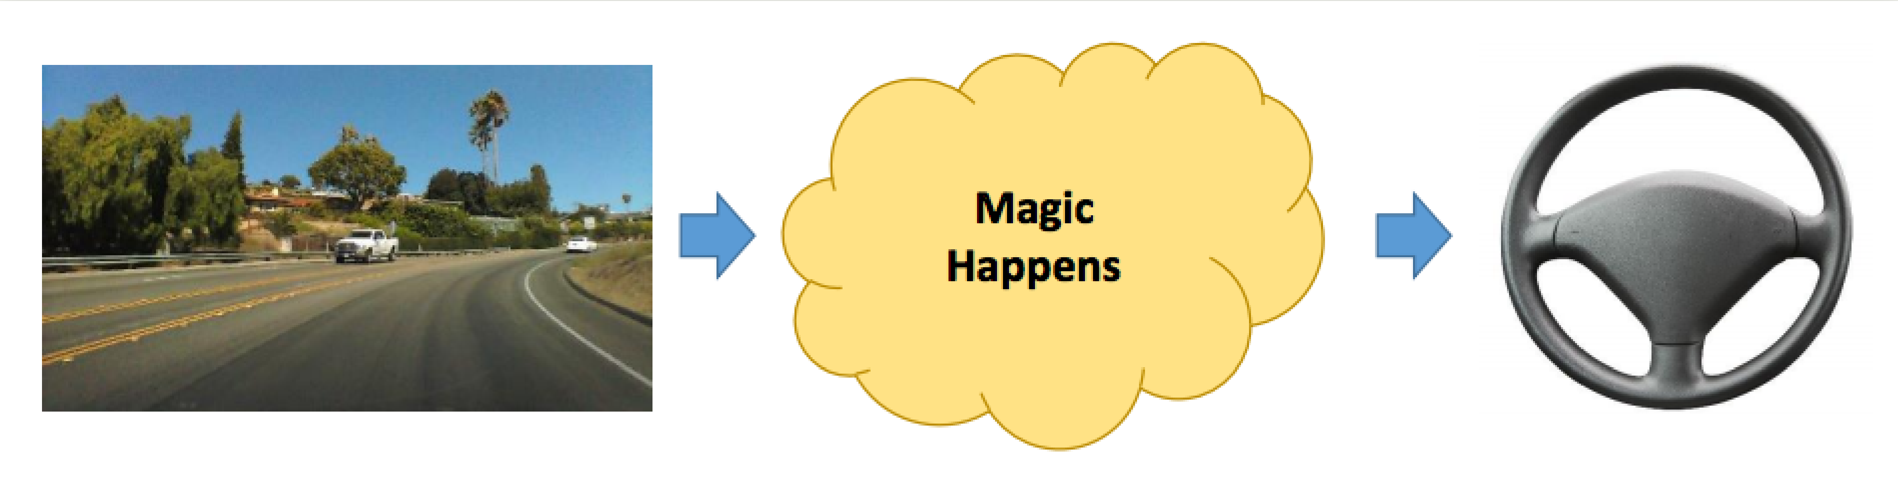

End-to-end deep learning maps sensor data directly to vehicle controls

End to End Deep Learning

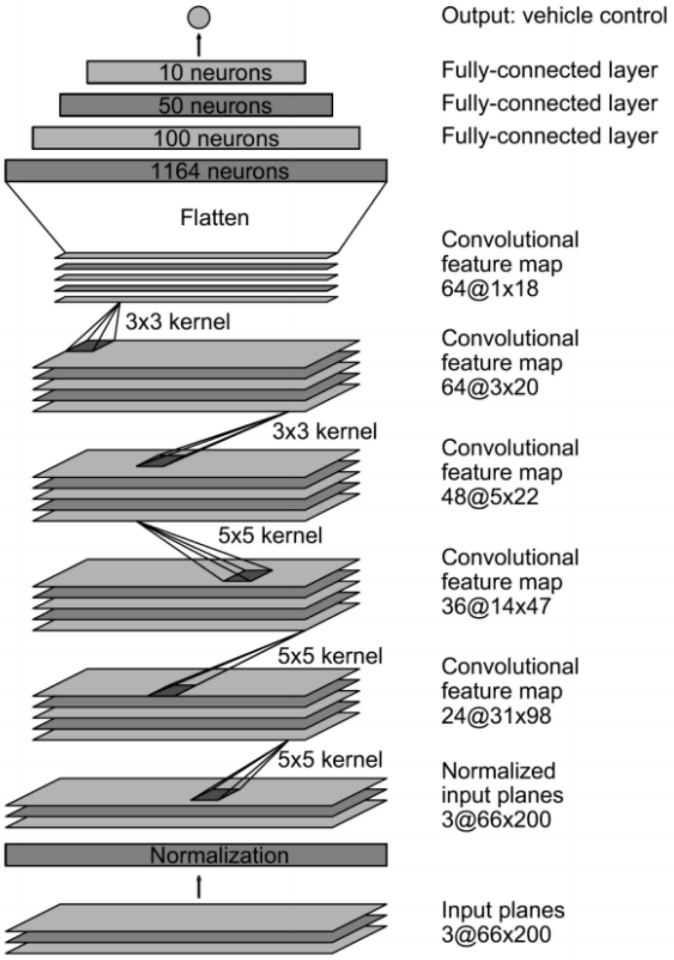

- 9 layers
  - 1 normalization layer
  - 5 convolutional layers
  - 3 fully connected layers
- 27 million connections
- 250 thousand parameters

Prediction targets:

- Steering wheel position
- Accelerator or break strength
- On/off of various signals

CNN Applications

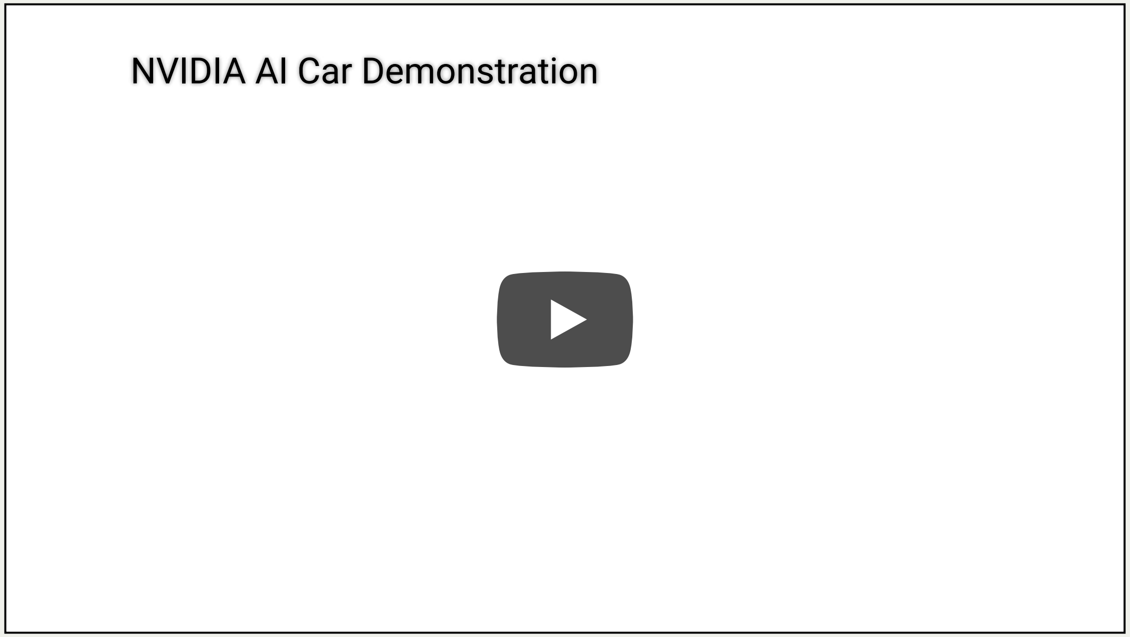

RNN Applications

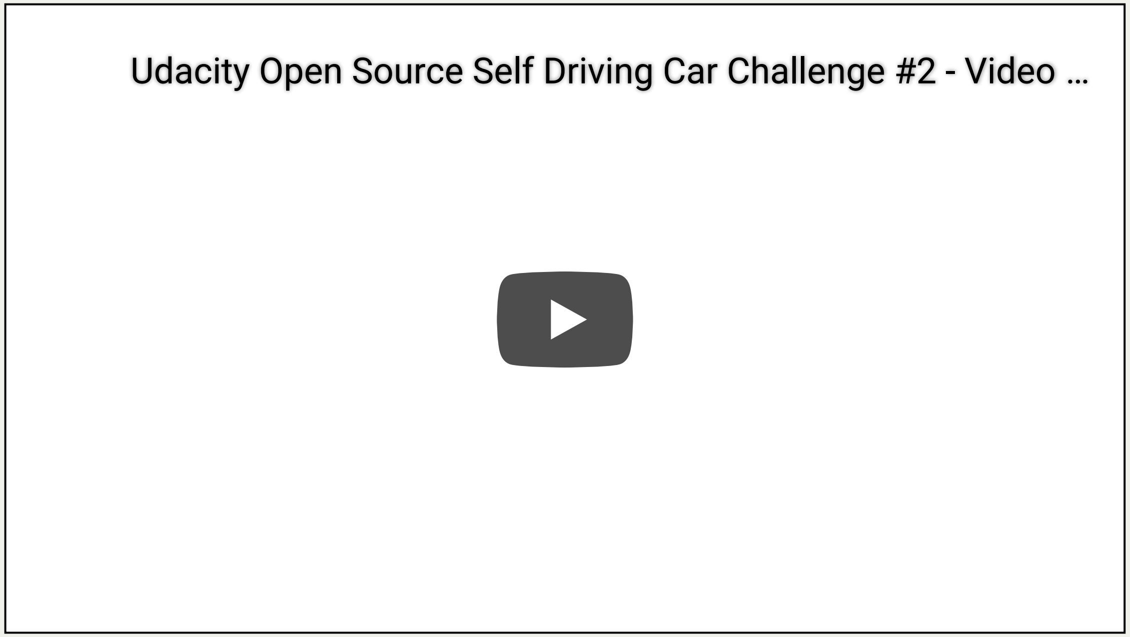

Deep Tesla

Summary

- Deep Learning is an important tool in self driving
- Self driving cars can use it end-to-end, or as a tool
- We have the building blocks: CNN's, RNN's, and RL
- This is an important problem for deep learning!

Thanks!