Applying Deep Learning to Car Data Logging (CDL) and Driver Assessor (DA)

October

Kyungwon Byun(Steve Byun)
FAE
NVIDIA

AGENDA

1. What is Deep Learning?
2. Deep Learning software
3. Deep Learning deployment
4. CDL using Deep Learning
5. Driver Workload Assessor using Deep Learning

Deep Learning & Artificial Intelligence

Deep Learning has become the most popular approach to developing Artificial Intelligence (AI) – machines that perceive and understand the world

The focus is currently on specific perceptual tasks, and there are many successes.

Today, some of the world's largest internet companies, as well as the foremost research institutions, are using GPUs for deep learning in research and production

Practical DEEP LEARNING Examples

Image Classification, Object Detection, Localization,
Action Recognition, Scene Understanding

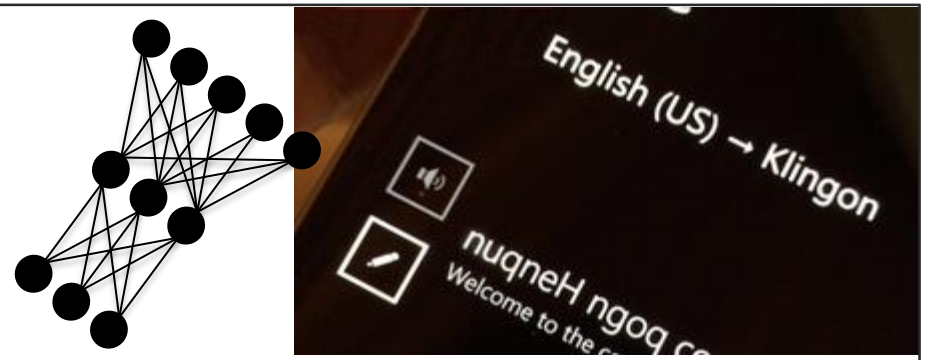

Speech Recognition, Speech Translation,
Natural Language Processing

Pedestrian Detection, Traffic Sign Recognition

Breast Cancer Cell Mitosis Detection,
Volumetric Brain Image Segmentation

Traditional MACHINE PERCEPTION – hand tuned features

Raw data	Feature extraction	Classifier/ detector	Result
		SVM, shallow neural net, …	
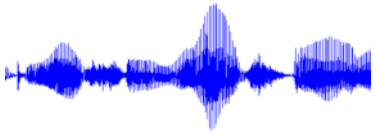	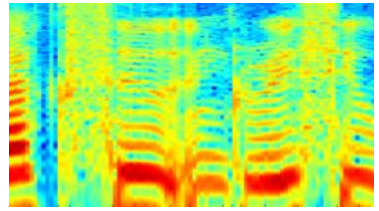	HMM, shallow neural net, …	Speaker ID, speech transcription, …
	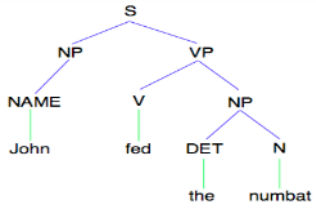	Clustering, HMM, LDA, LSA …	Topic classification, machine translation, sentiment analysis…

Deep Learning Approach

Train:

Errors

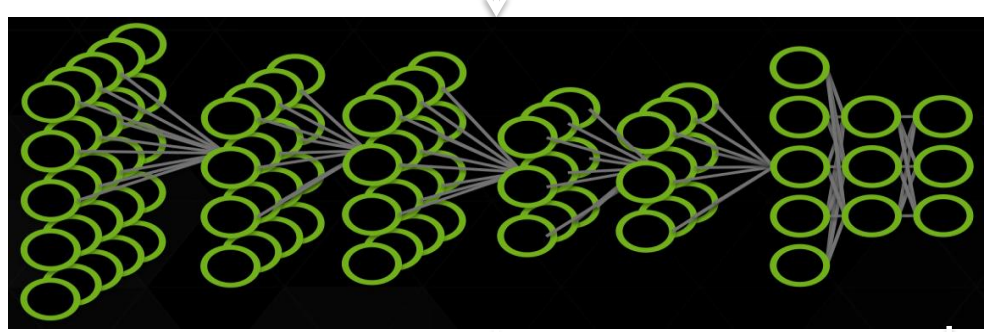

Dog
Cat
Raccoon

Deploy:

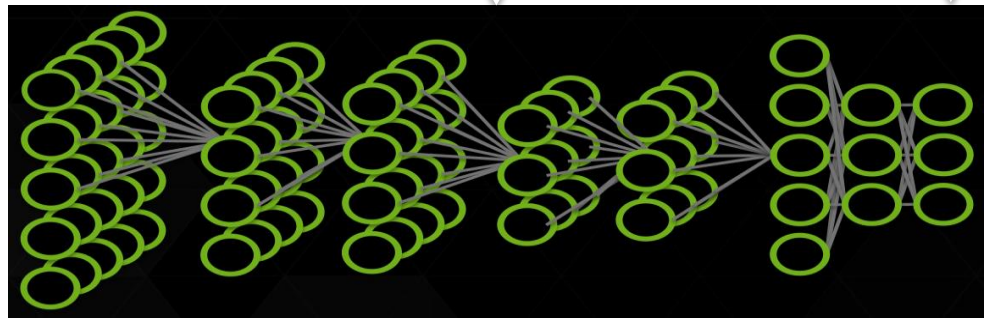

Dog

SOME DEEP LEARNING Use Cases

Input	Output
Pixels: 	“lion”
Audio: 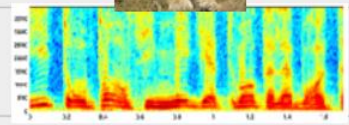	“see at tuhl res taur aun ts”
<query, doc>	P(click on doc)
“Hello, how are you?”	“Bonjour, comment allez-vous?”
Pixels: 	“A close up of a small child holding a stuffed animal”

Artificial Neural Network (ANN)

A collection of simple, trainable mathematical units that collectively learn complex functions

Biological neuron

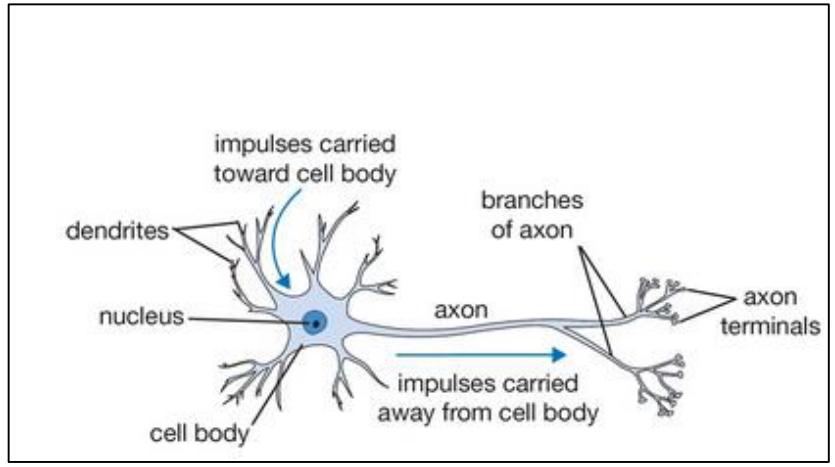

From Stanford cs231n lecture notes

Artificial neuron

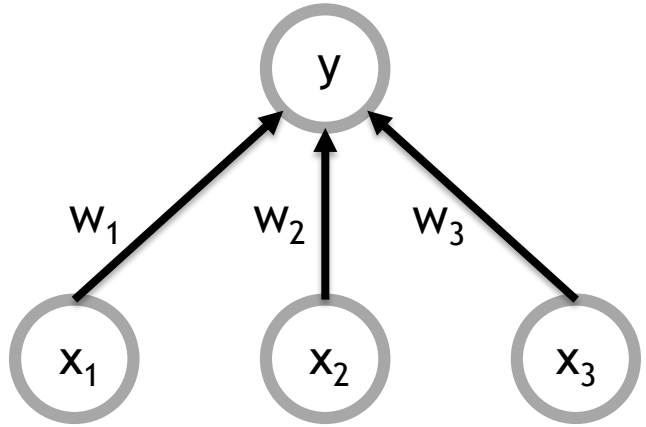

$$y=F(w_1x_1+w_2x_2+w_3x_3)$$

$$F(x)=\max(0,x)$$

Artificial Neural Network (Ann)

A collection of simple, trainable mathematical units that collectively learn complex functions

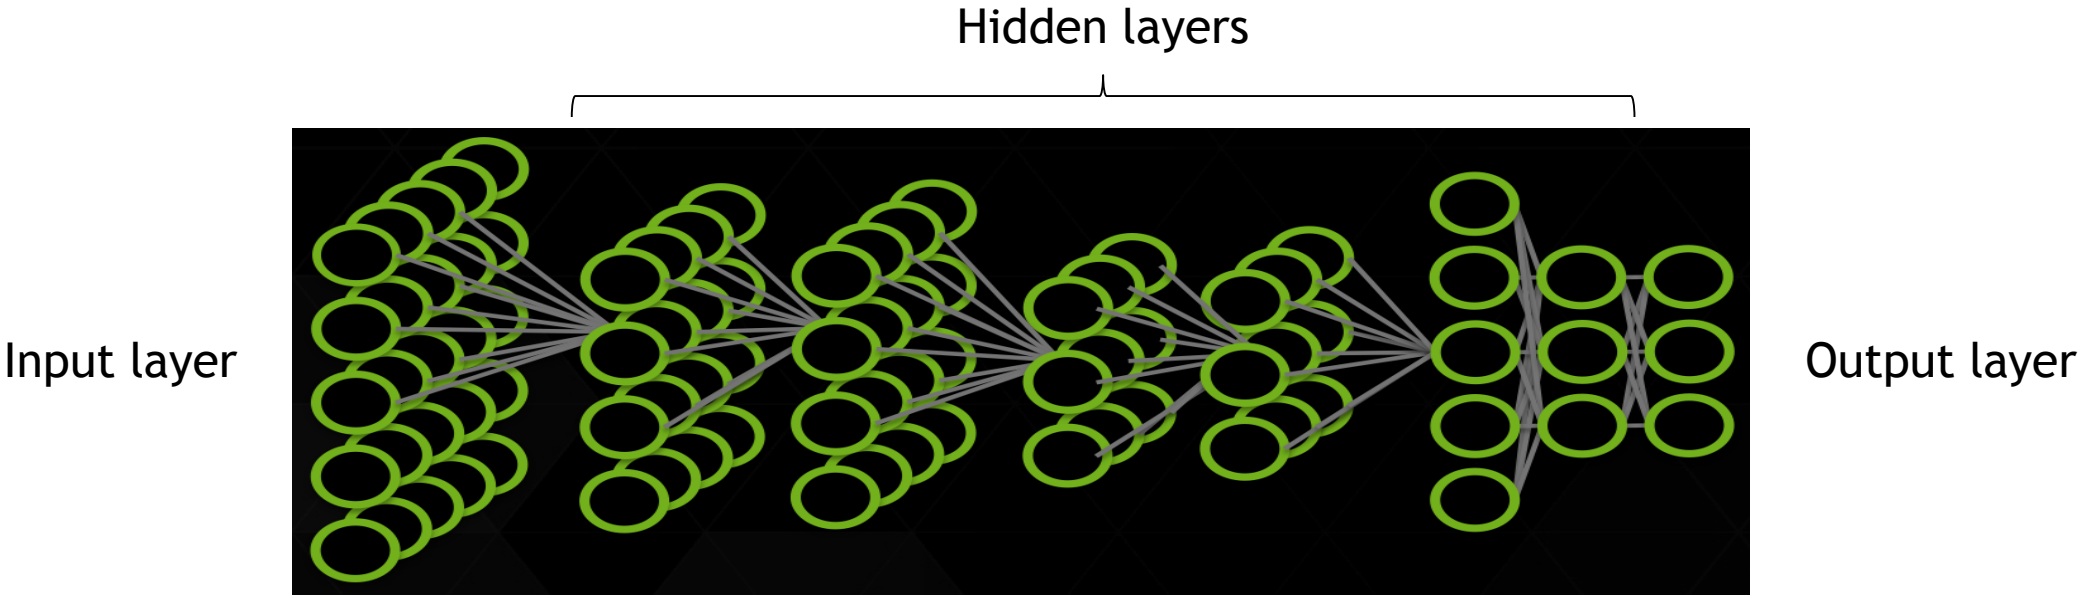

Given sufficient training data an artificial neural network can approximate very complex functions mapping raw data to output decisions

Deep Neural Network (DNN)

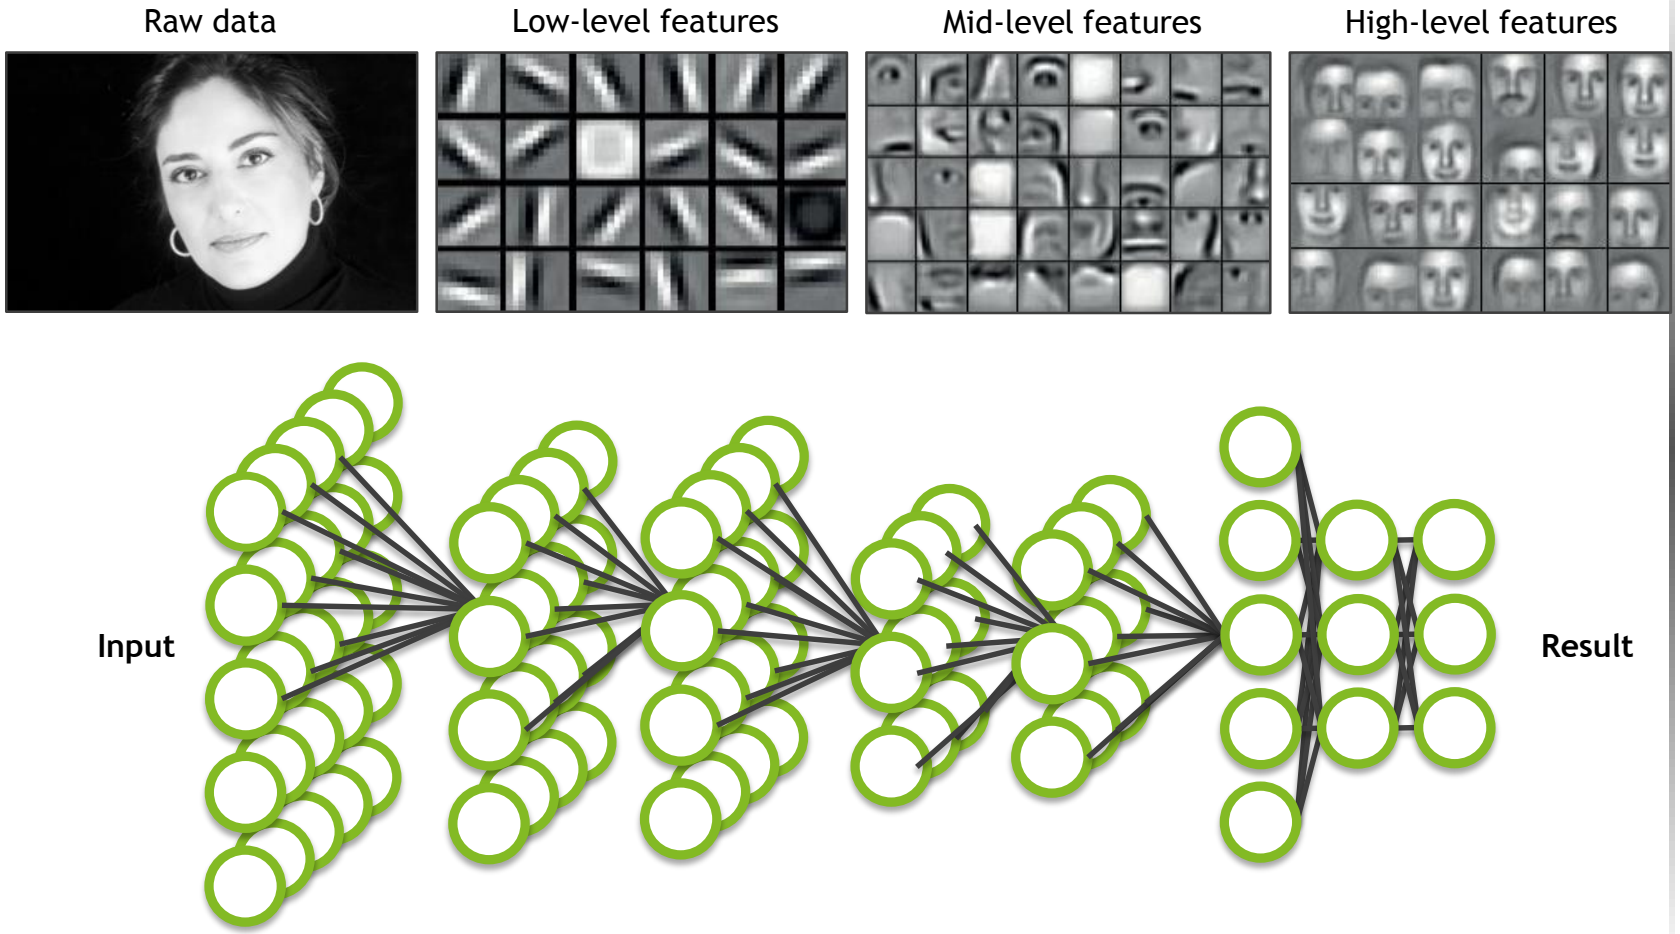

Application components:

Task objective
- e.g. Identify face

Training data
- 10-100M images

Network architecture
- ~10 layers
- 1B parameters

Learning algorithm
- ~30 Exaflops
- ~30 GPU days

Deep Learning Advantages

- **Robust**
  - No need to design the features ahead of time – features are automatically learned to be optimal for the task at hand
  - Robustness to natural variations in the data is automatically learned
- **Generalizable**
  - The same neural net approach can be used for many different applications and data types
- **Scalable**
  - Performance improves with more data, method is massively parallelizable

Convolutional Neural Network (cnn)

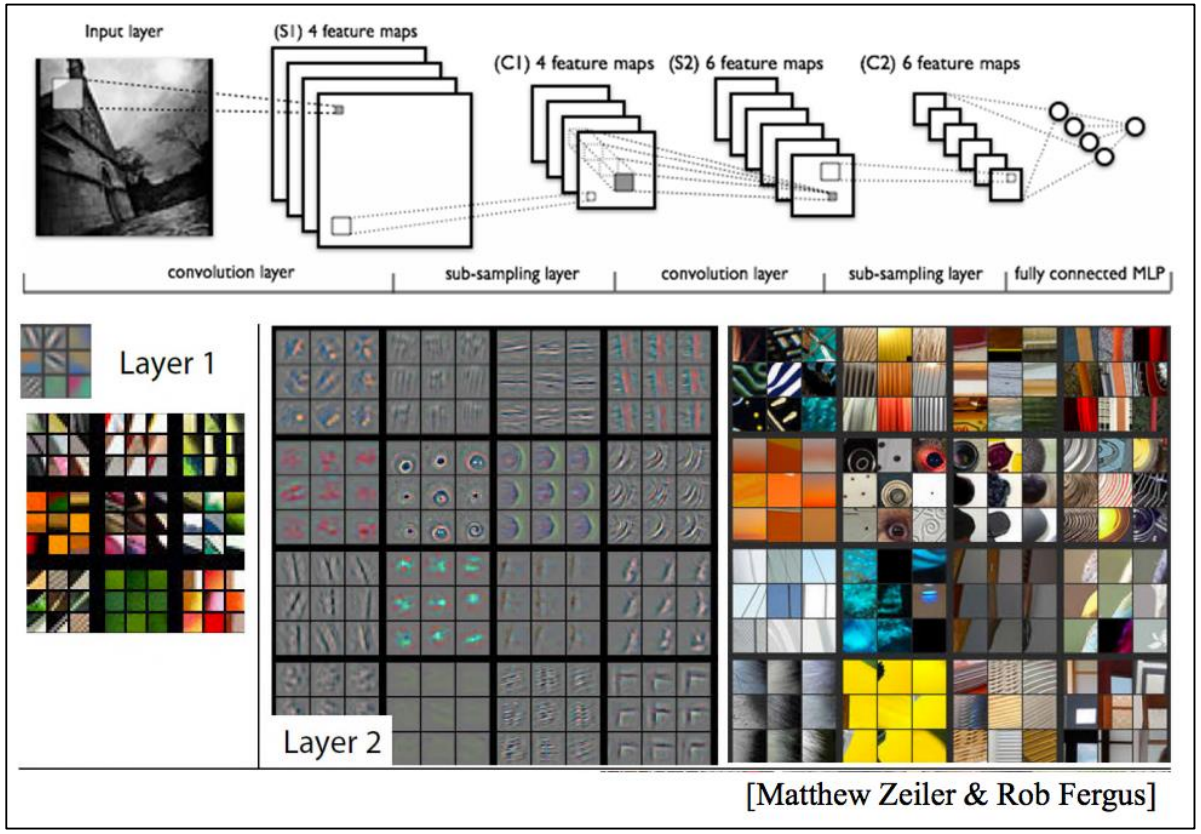

- Inspired by the human visual cortex
- Learns a hierarchy of visual features
- Local pixel level features are scale and translation invariant
- Learns the “essence” of visual objects and generalizes well

Convolutional Neural Network (cnn)

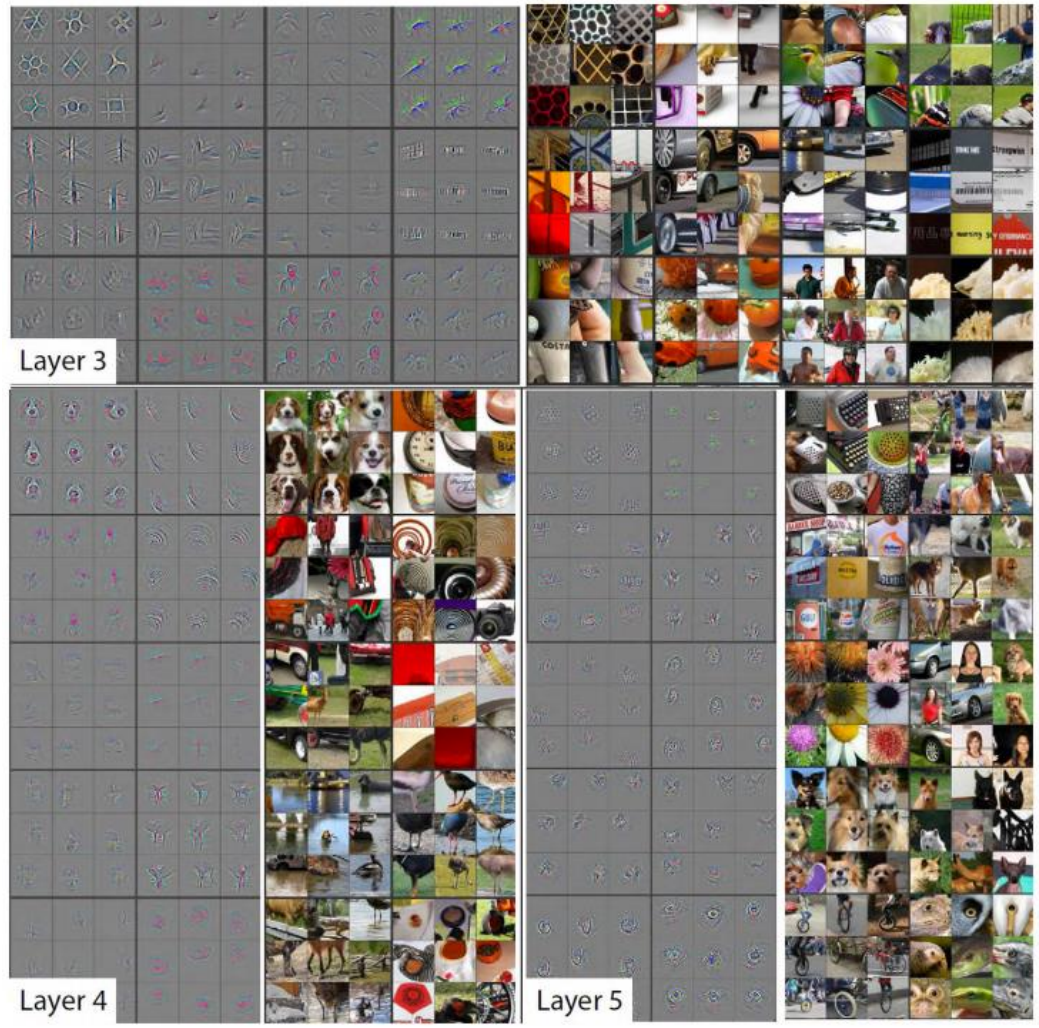

Recurrent Neural Network (RNN)

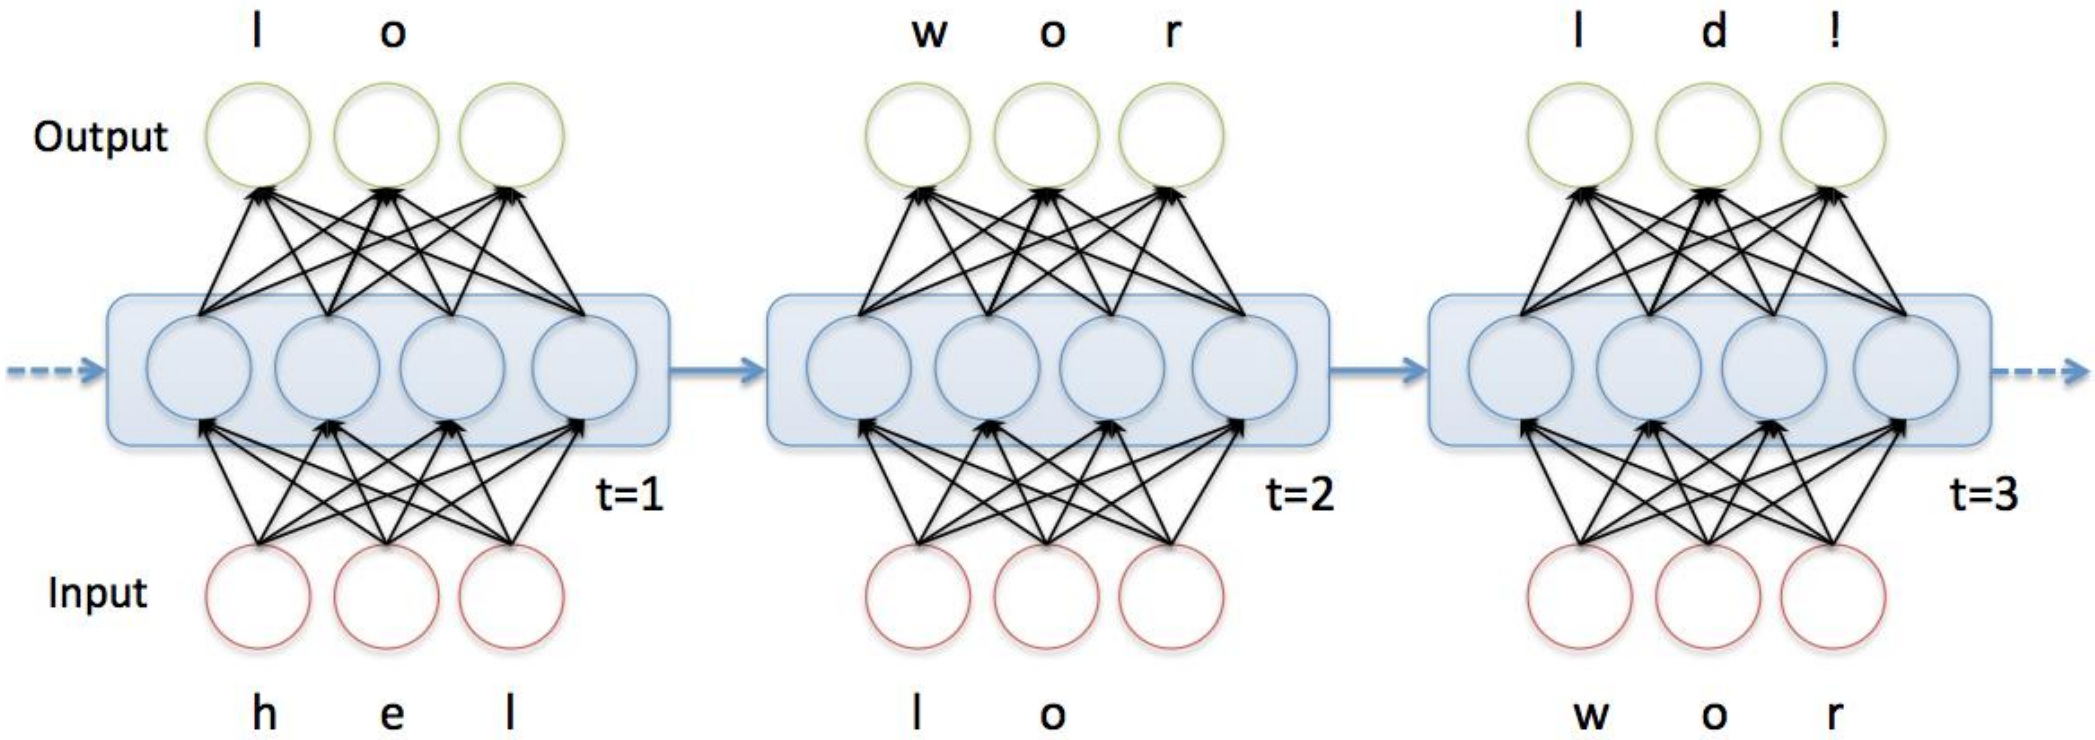

DNNs dominate in perceptual tasks

- Handwriting recognition MNIST (many), Arabic HWX (IDSIA)
- OCR in the Wild [2011]: StreetView House Numbers (NYU and others)
- Traffic sign recognition [2011] GTSRB competition (IDSIA, NYU)
- Asian handwriting recognition [2013] ICDAR competition (IDSIA)
- Pedestrian Detection [2013]: INRIA datasets and others (NYU)
- Volumetric brain image segmentation [2009] connectomics (IDSIA, MIT)
- Human Action Recognition [2011] Hollywood II dataset (Stanford)
- Object Recognition [2012] ImageNet competition (Toronto)
- Scene Parsing [2012] Stanford bgd, SiftFlow, Barcelona datasets (NYU)
- Scene parsing from depth images [2013] NYU RGB-D dataset (NYU)
- Speech Recognition [2012] Acoustic modeling (IBM and Google)
- Breast cancer cell mitosis detection [2011] MITOS (IDSIA)

Slide credit: Yann Lecun, Facebook & NYU

Why is Deep learning hot *now*?

Three Driving Factors…

Big Data Availability

facebook — 350 millions images uploaded per day

Walmart — 2.5 Petabytes of customer data hourly

YouTube — 100 hours of video uploaded every minute

New DL Techniques

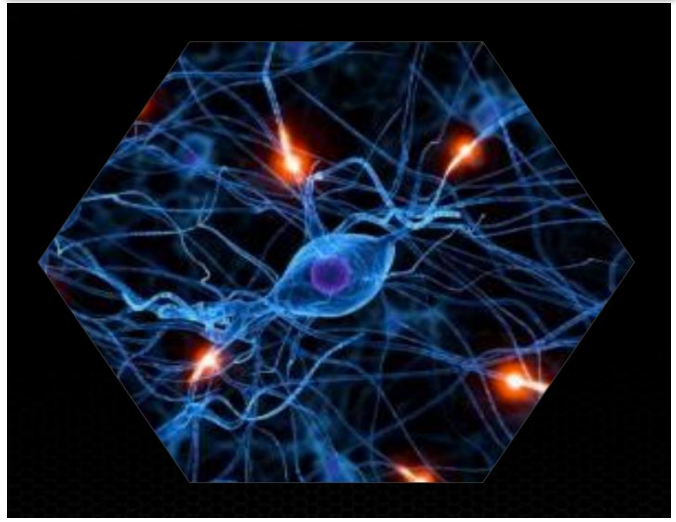

GPU acceleration

The platform for Deep Learning

Image Recognition Challenge

1.2M training images • 1000 object categories

Hosted by

IMAGENET

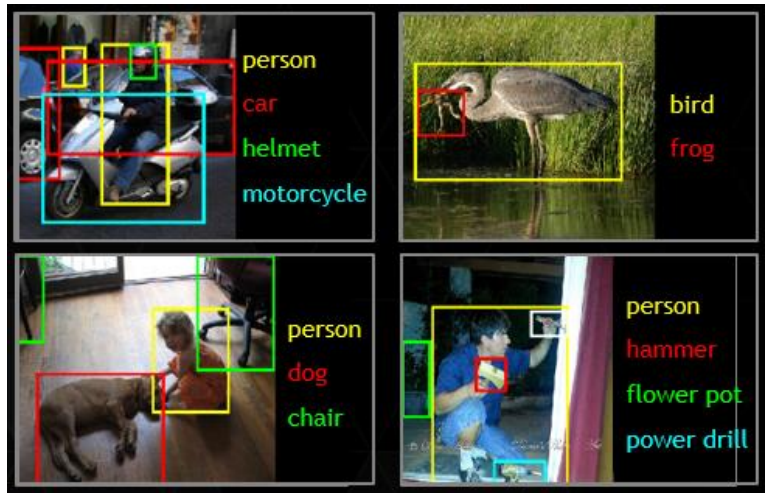

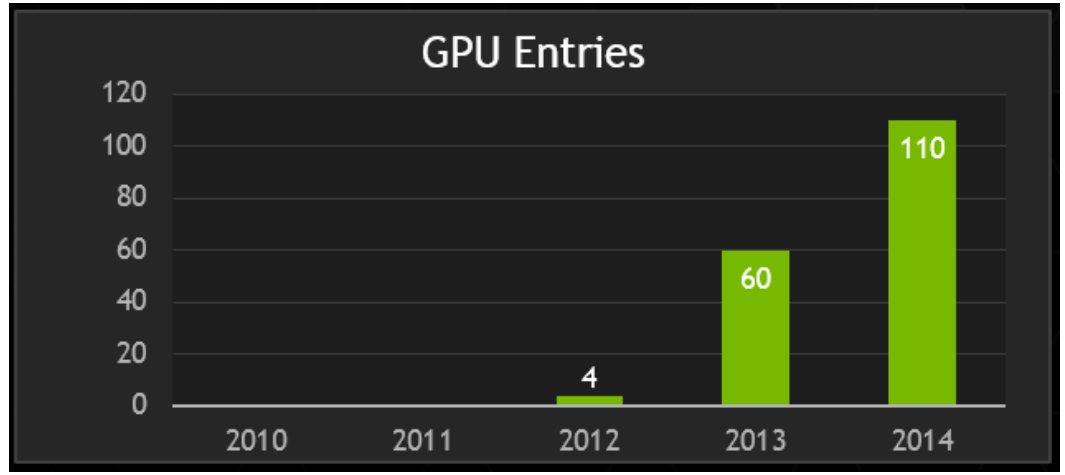

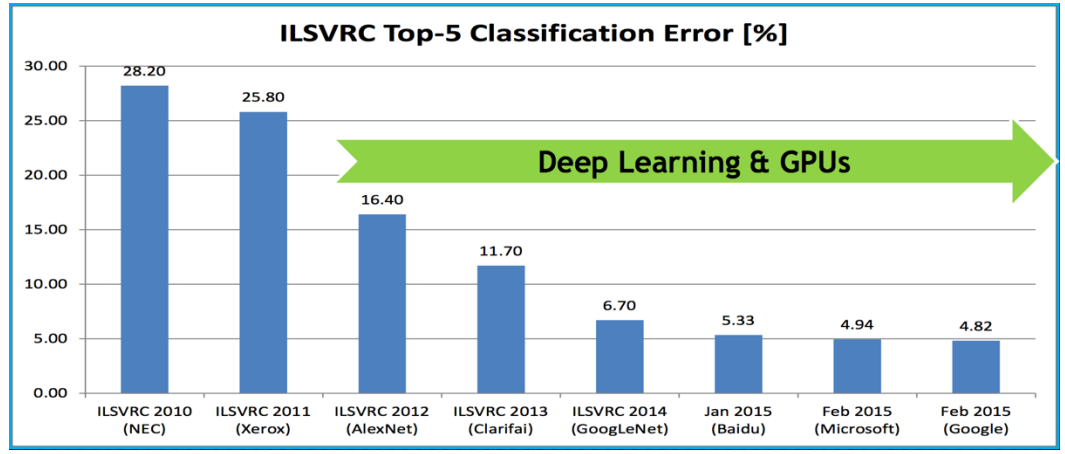

How to write applications using DL

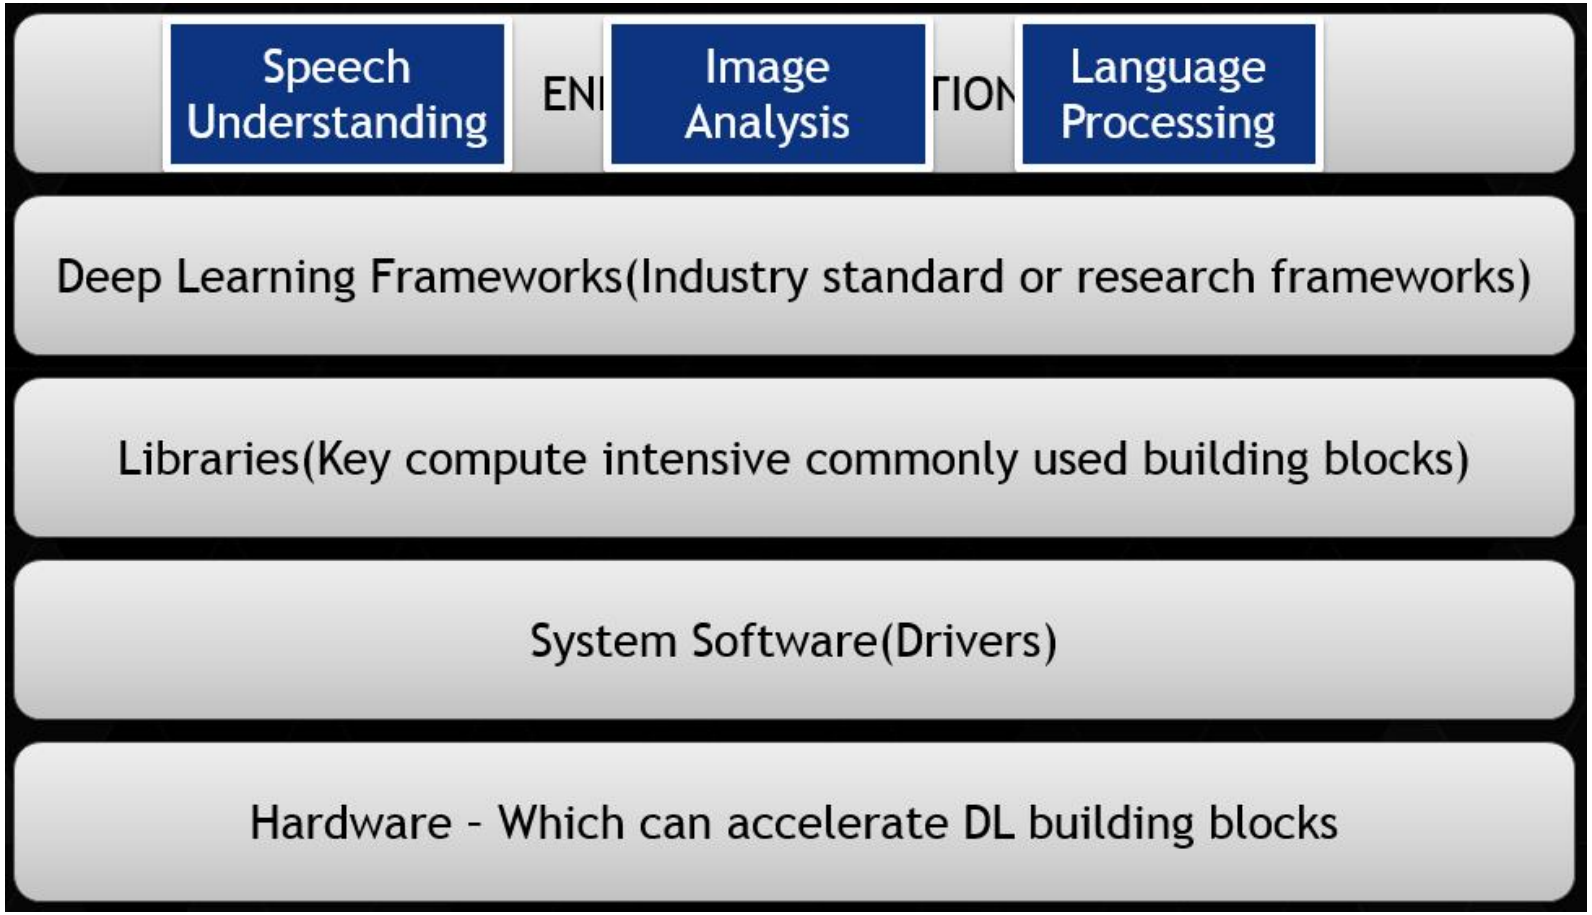

GPU-Accelerated Deep Learning Frameworks

	CAFFE	TORCH	THEANO	KALDI
Domain	Deep Learning Framework	Scientific Computing Framework	Math Expression Compiler	Speech Recognition Toolkit
cuDNN	2.0	2.0	2.0	--
Multi-GPU	via DIGITS 2	In Progress	In Progress	✓(**nnet2**)
Multi-CPU	×	×	×	✓(**nnet2**)
License	BSD-2	GPL	BSD	Apache 2.0
Interface(s)	Command line, Python, MATLAB	Lua, Python, MATLAB	Python	C++, Shell scripts
Embedded (TK1)	✓	✓	×	×

Deep Learning deployment workflow

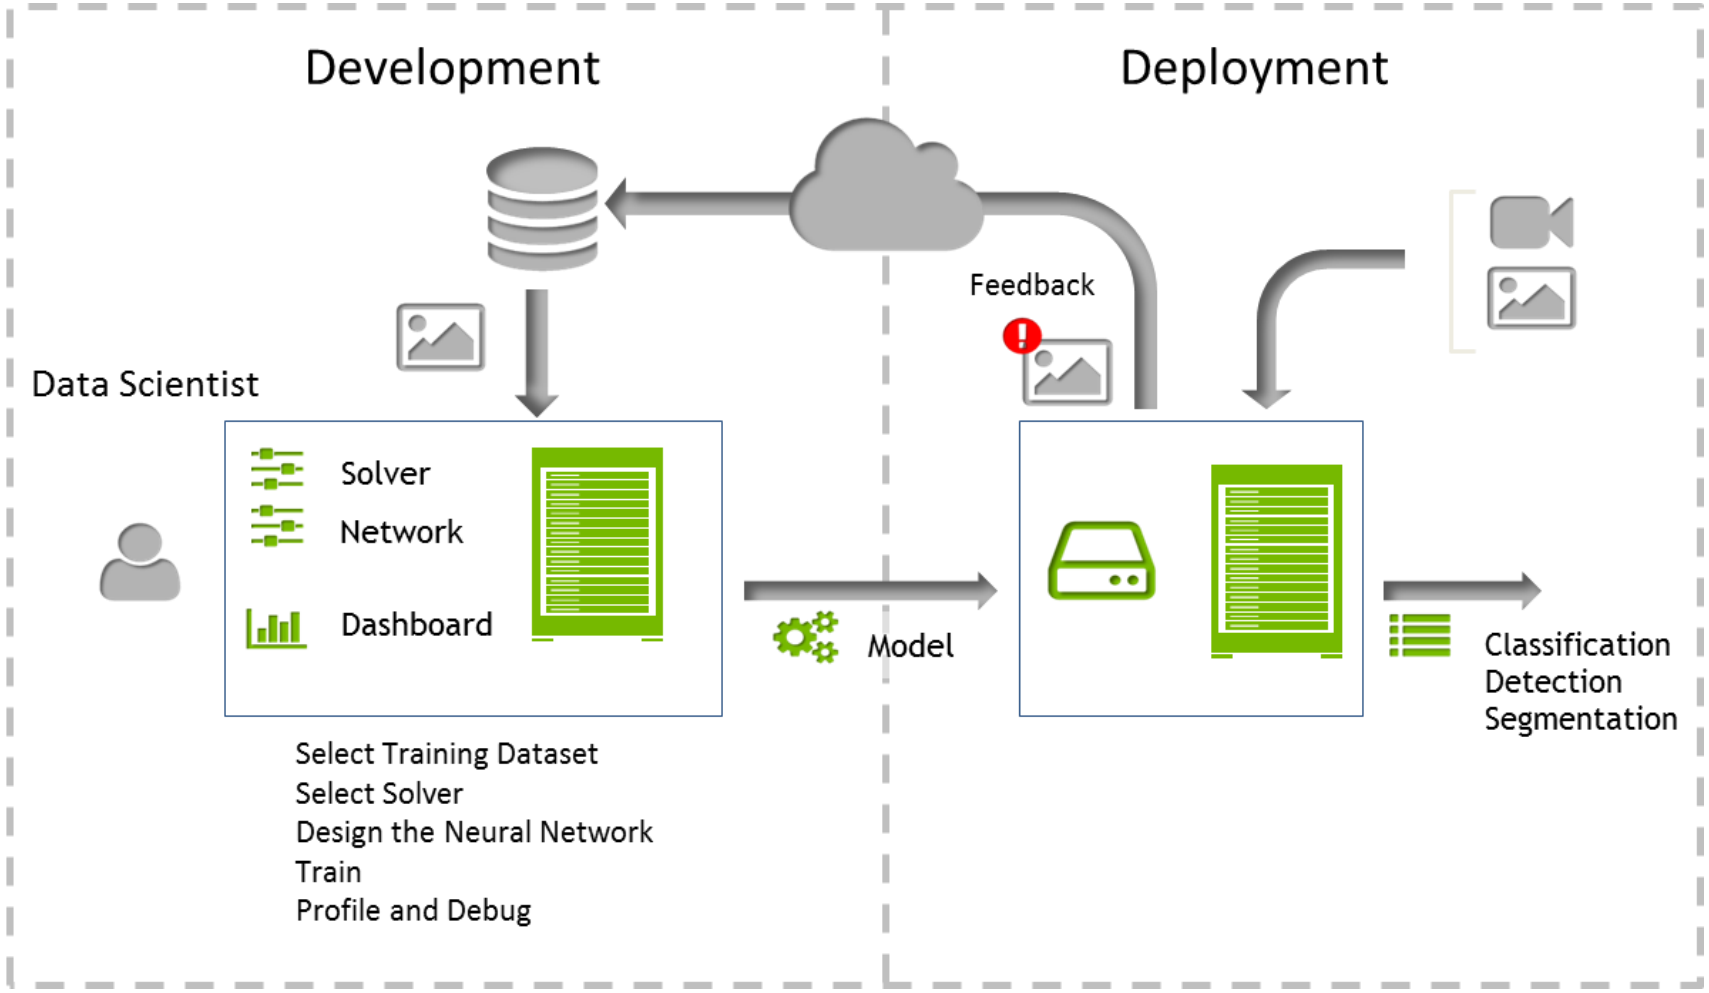

Project DAVE - DARPA autonomous vehicle

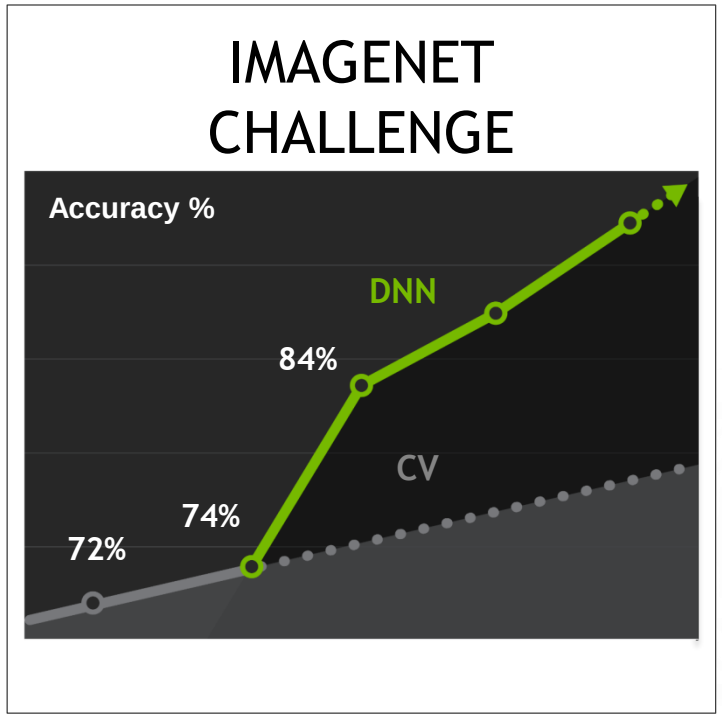

DNN-based self-driving robot

Training data by human driver

No hand-coded CV algorithms

Project DAVE - DARPA autonomous vehicle




TRAINING DATA

225K Images

Project DAVE - DARPA autonomous vehicle

TEST DRIVE

No training

Project DAVE - DARPA autonomous vehicle

TEST DRIVE

Partially Trained
(52K images)

Project DAVE - DARPA autonomous vehicle

TEST DRIVE

Fully Trained
(225K images)

Project DAVE - DARPA autonomous vehicle

DAVE IN ACTION

Project DAVE - DARPA autonomous vehicle

CAR classification using DL

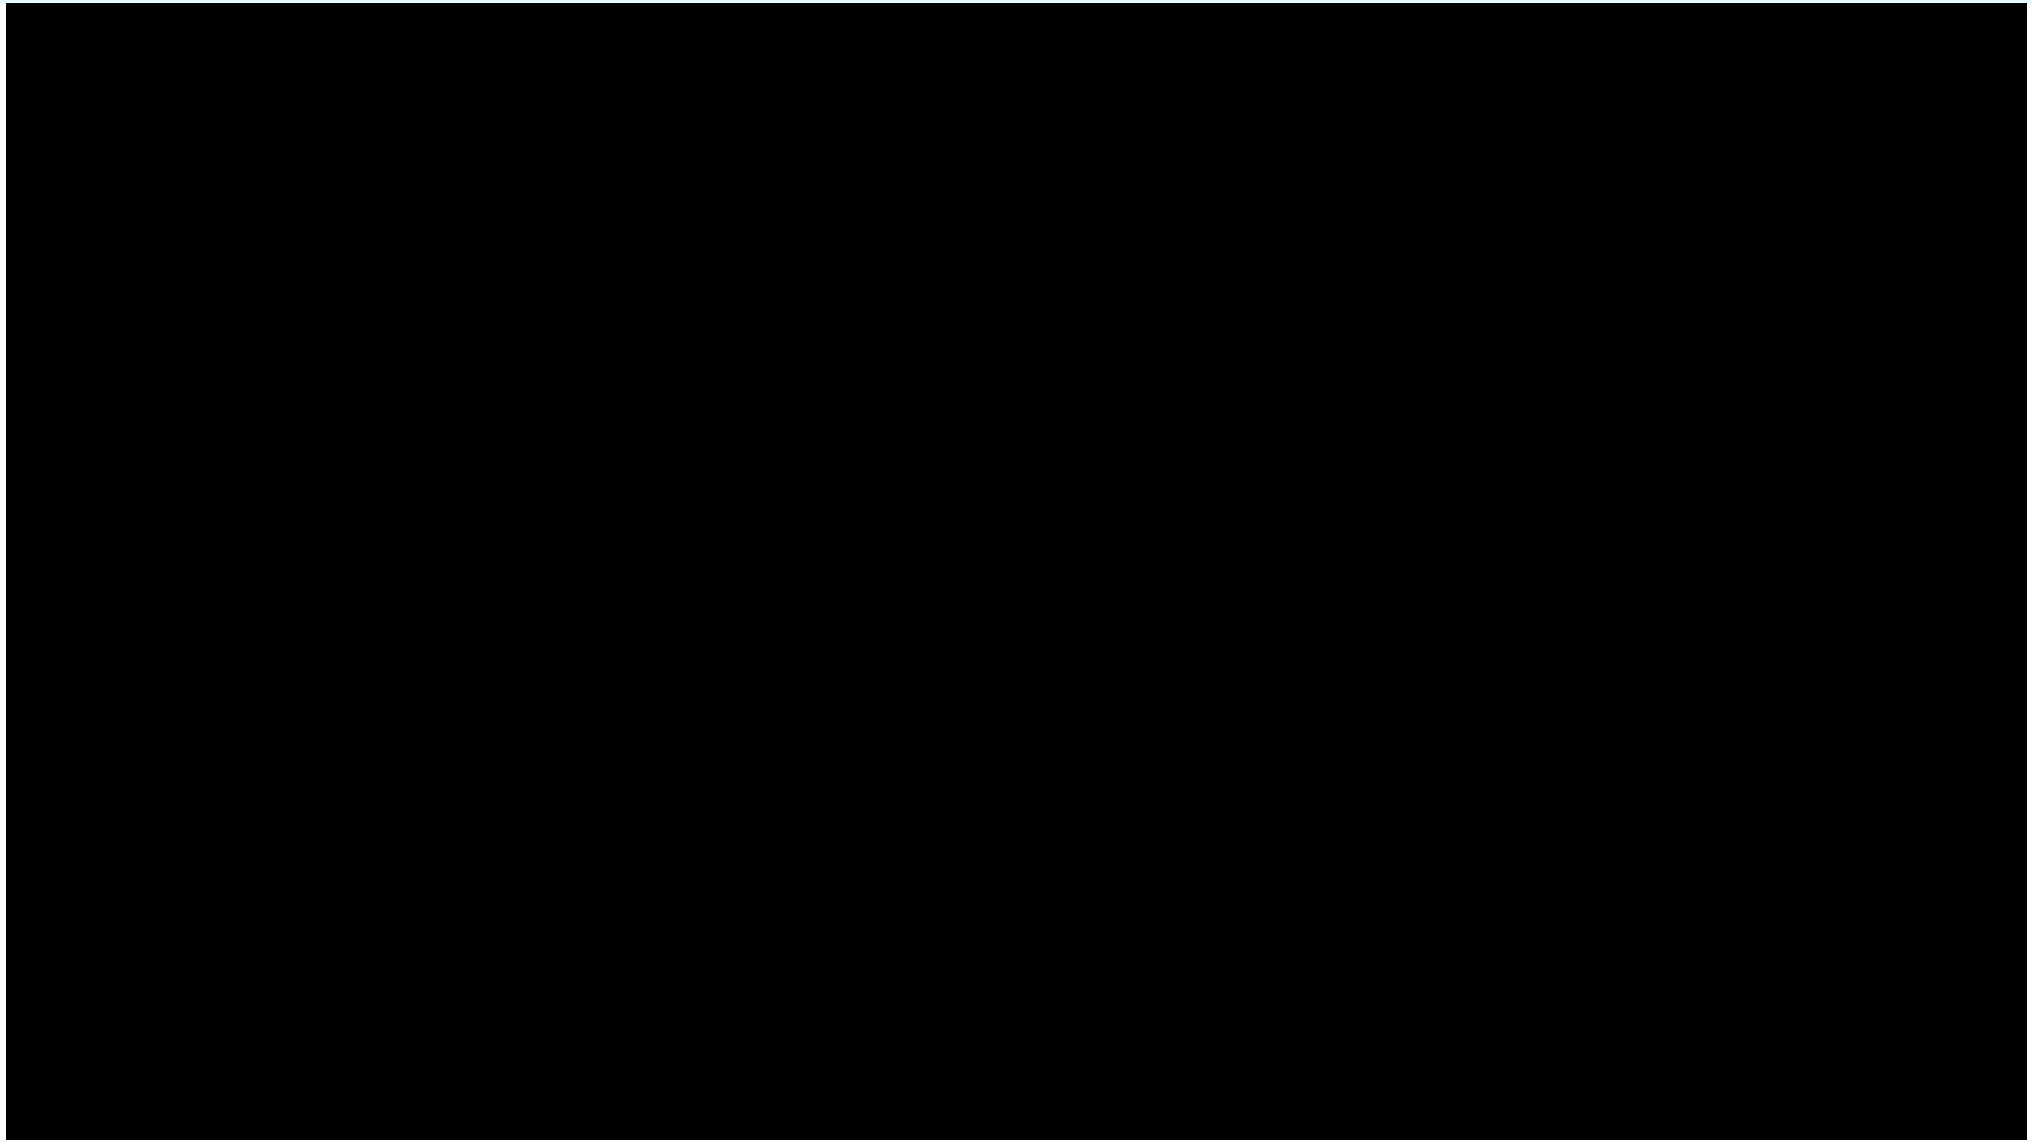

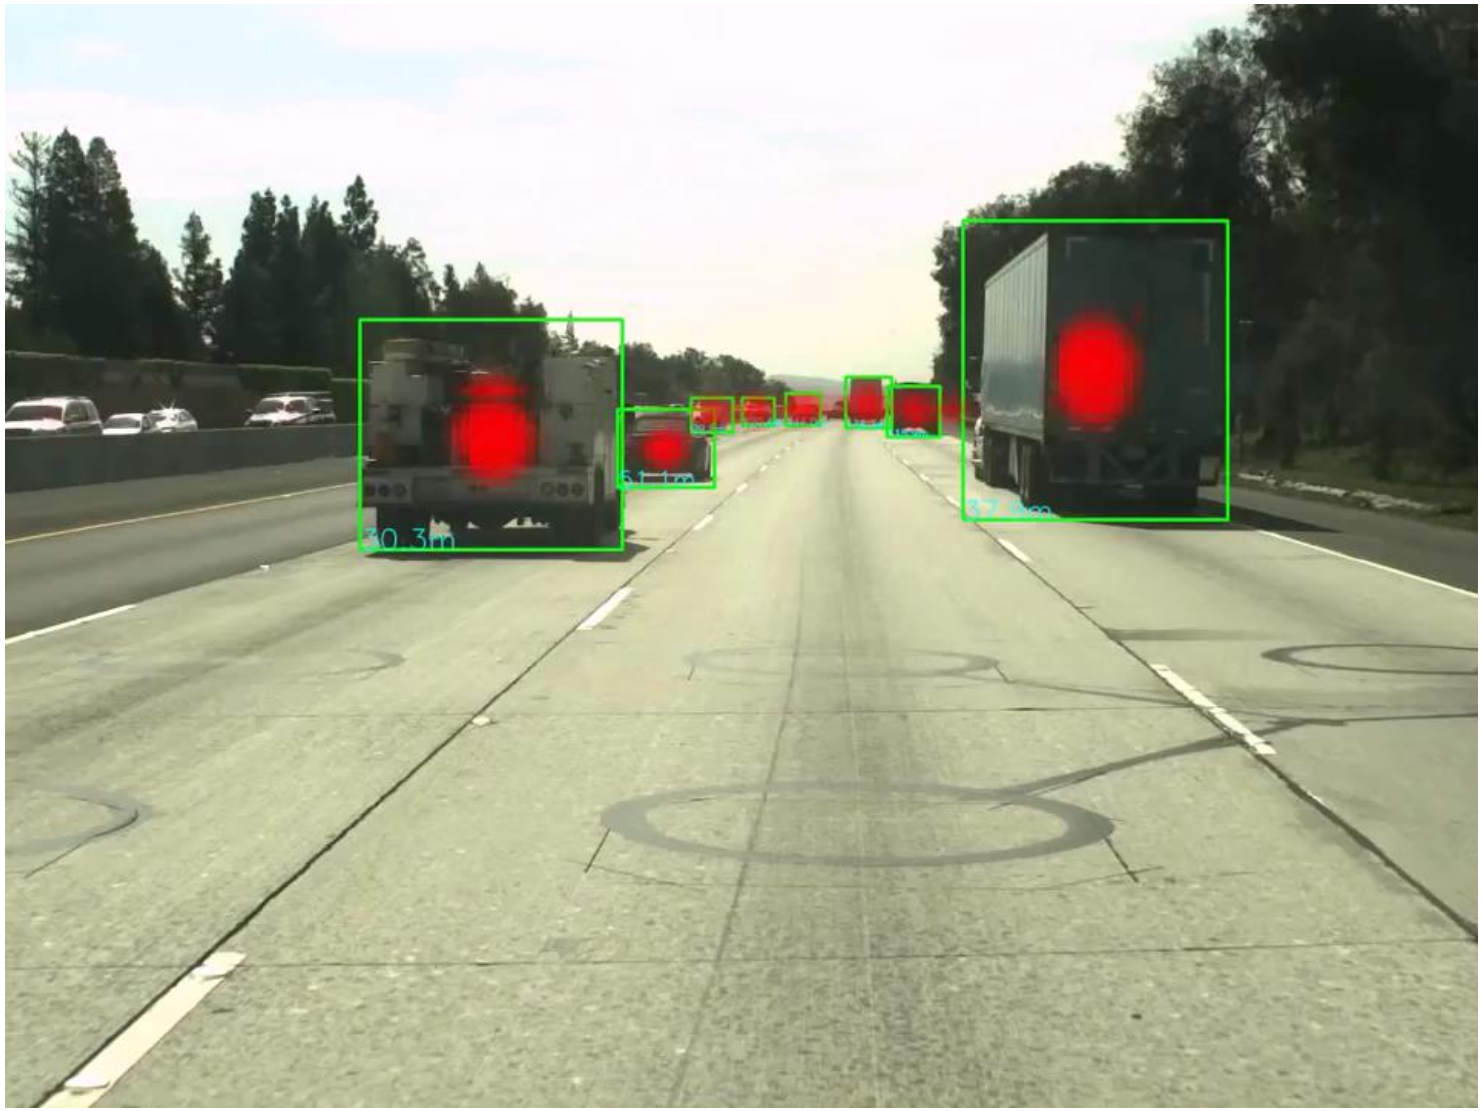

From: An Empirical Evaluation of Deep Learning on Highway Driving

CDL(Car Data Logging)

- Logging Car Information
  - Diagnostic data: engine parameter, speedometer etc.
  - Navigation data: GPS position, travel route etc.
  - Infotainment data: video, audio, music etc.

- Automatic Logging & Filtering for Real-Time Car Information
  - Interworking with DLT (Diagnostic Log & Trace) system on GENIVI platform
  - Using CAN/IPC/Serial/Ethernet protocols

- Logging data saving and sending data to off-board server
  - Saving Log data after Filtering and File format conversions.
  - Uploading Log data to off-board server

CDL(Car Data Logging)

- Why does CDL need?
  - The need to collect amount of Data generated by vehicle.(e.g. CAN data collection, Analysis)
  - Mass production stage for gathering real field data.
  - Reference Usage: Real time CAN Data Monitoring, Data Analysis

- CDL Scope
  - Specify what kind of data we have to collect and how to store it.
  - Specify how often do we have to collect data.
  - Specify what method we use for data transferring from car to server.

CDL(Car Data Logging) Arch

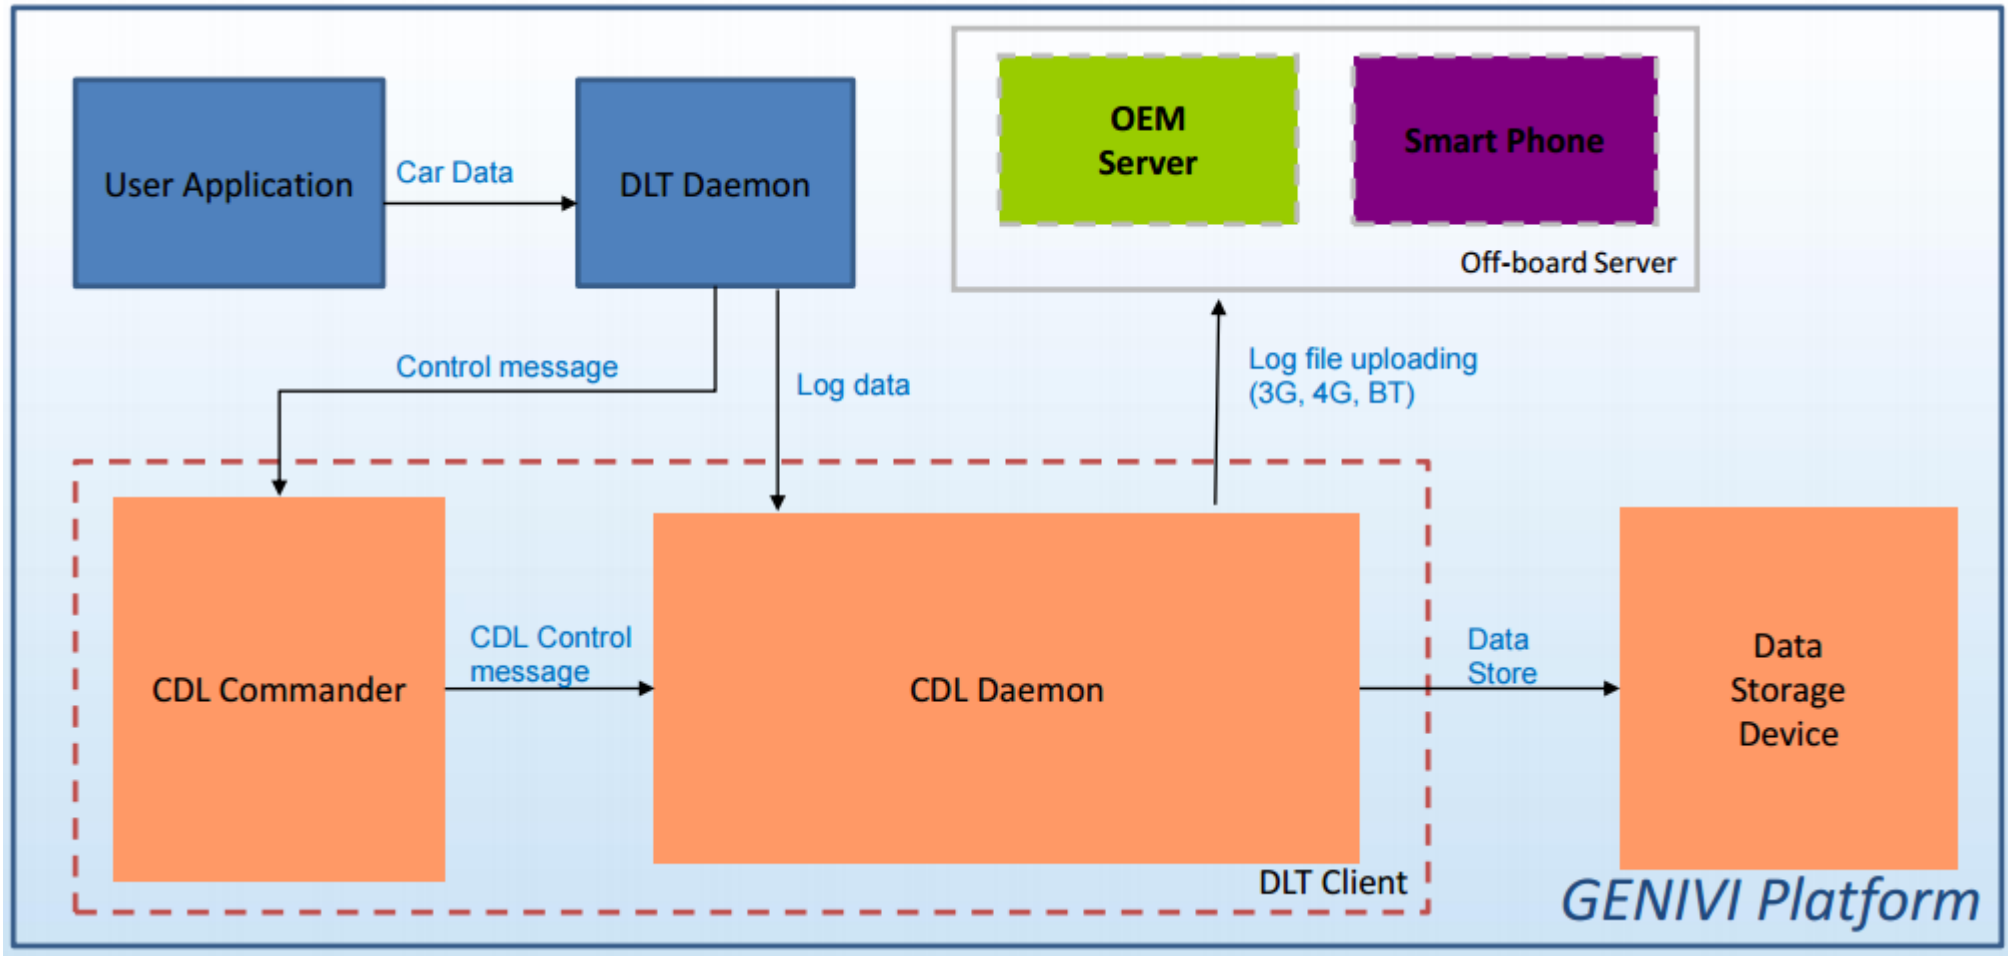

Deep Learning to Car Data Logging (CDL)

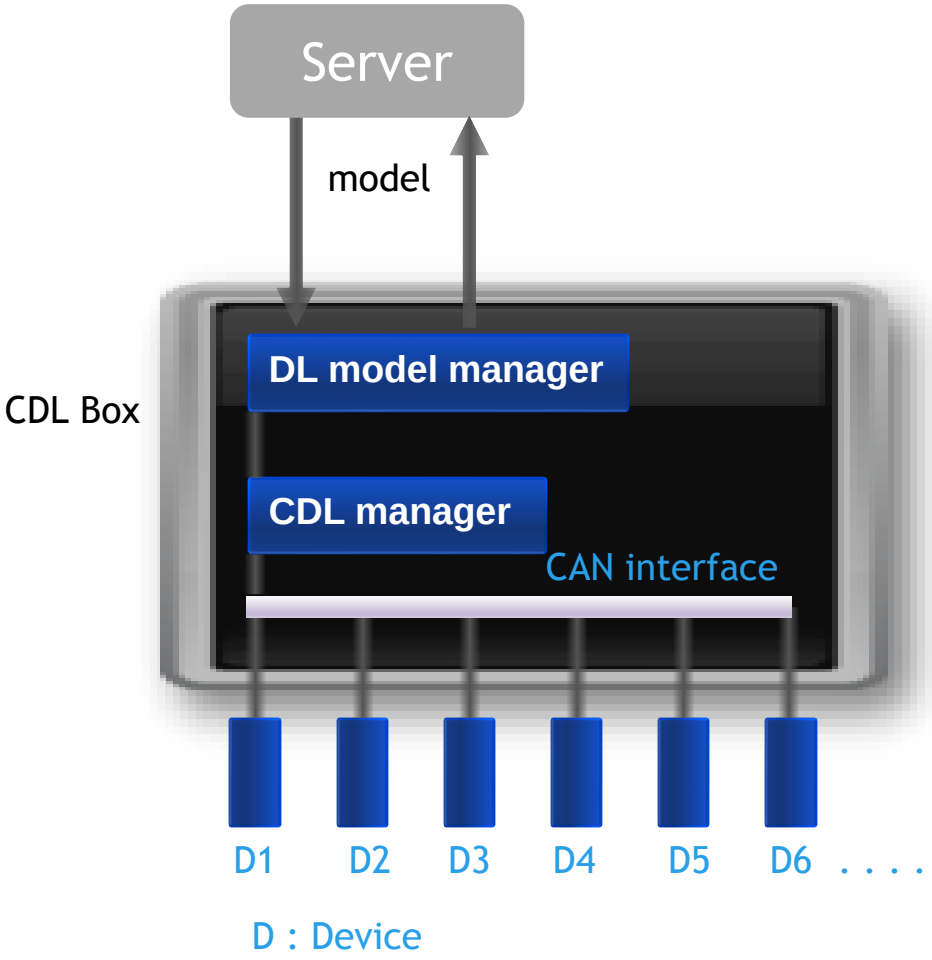

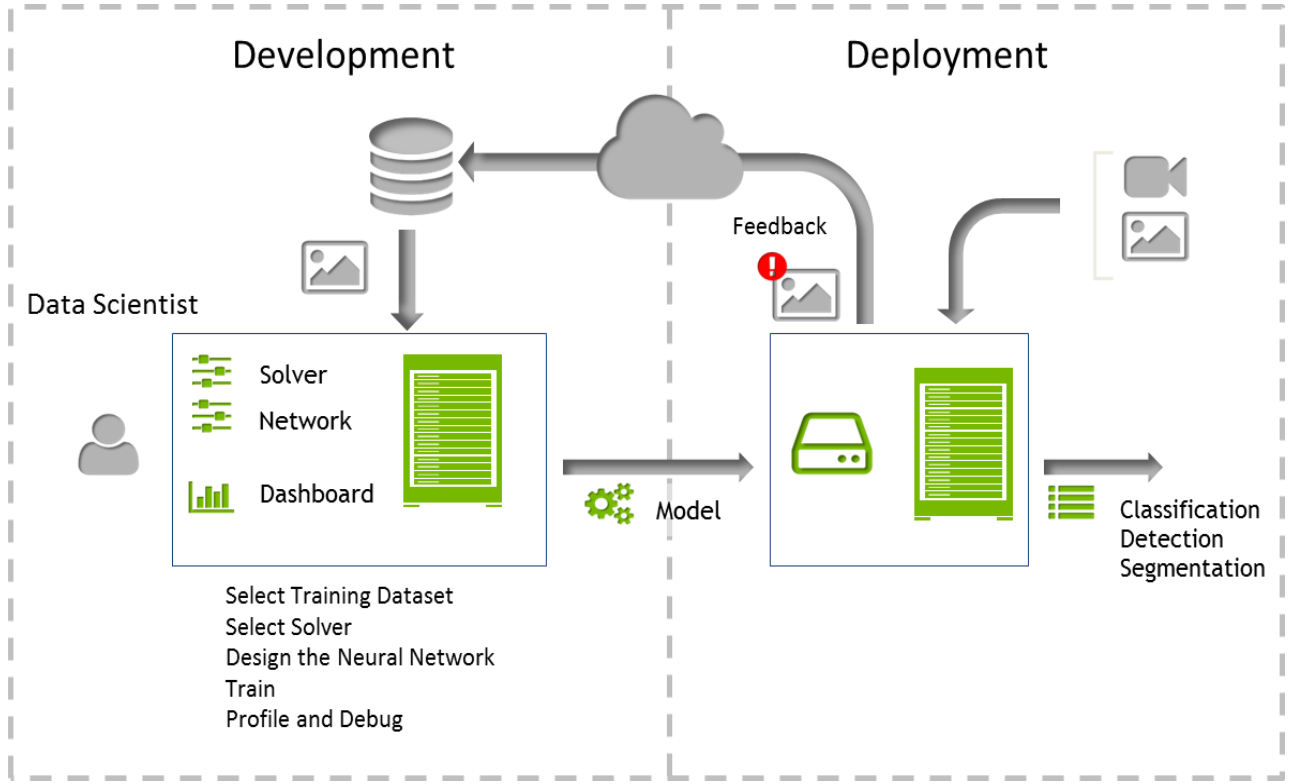

Driver Workload Assessor Diagram

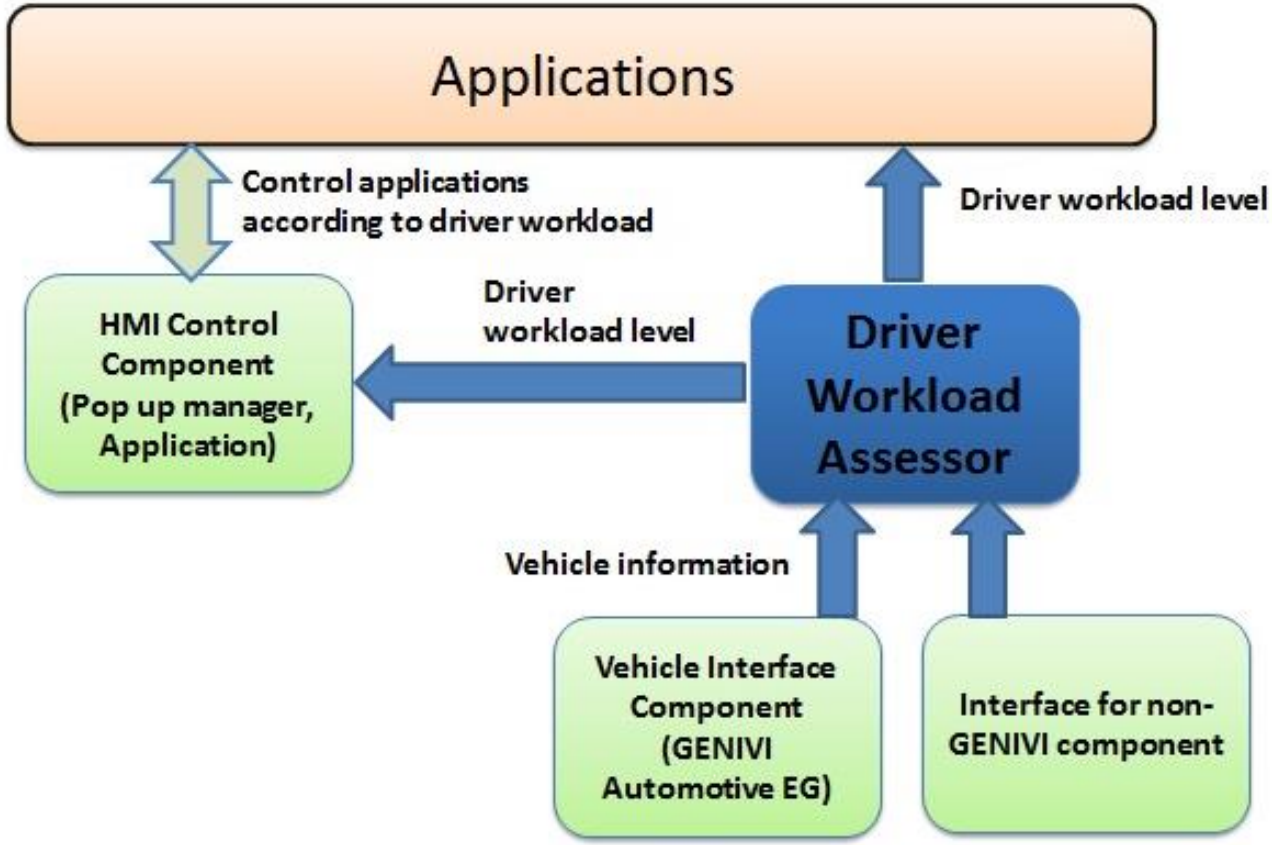

Driver Workload Assessor using DL

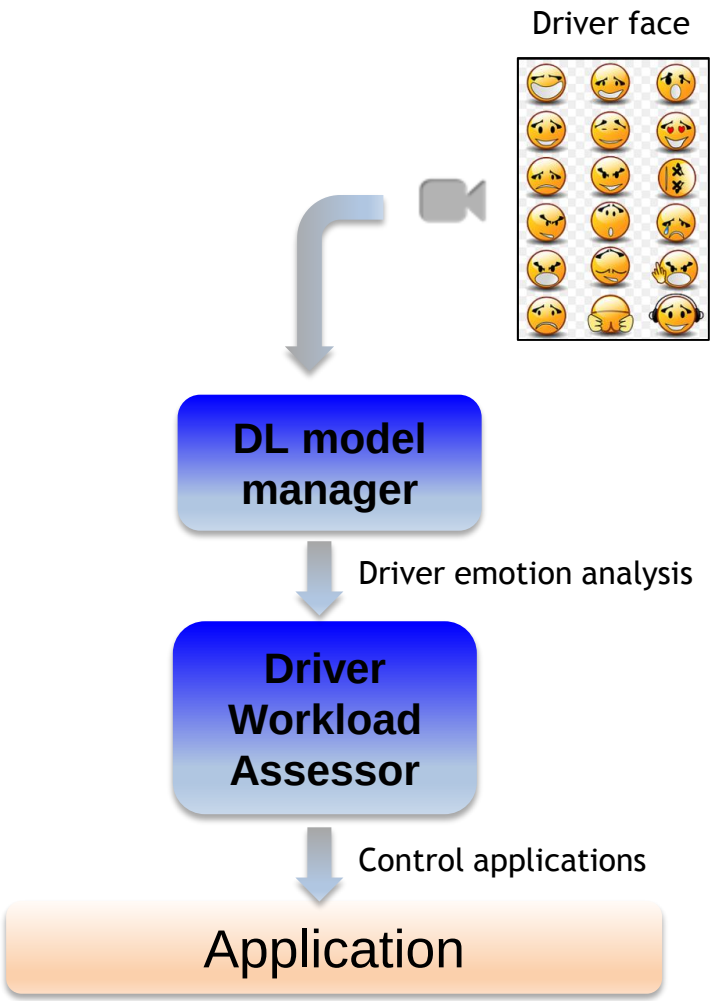

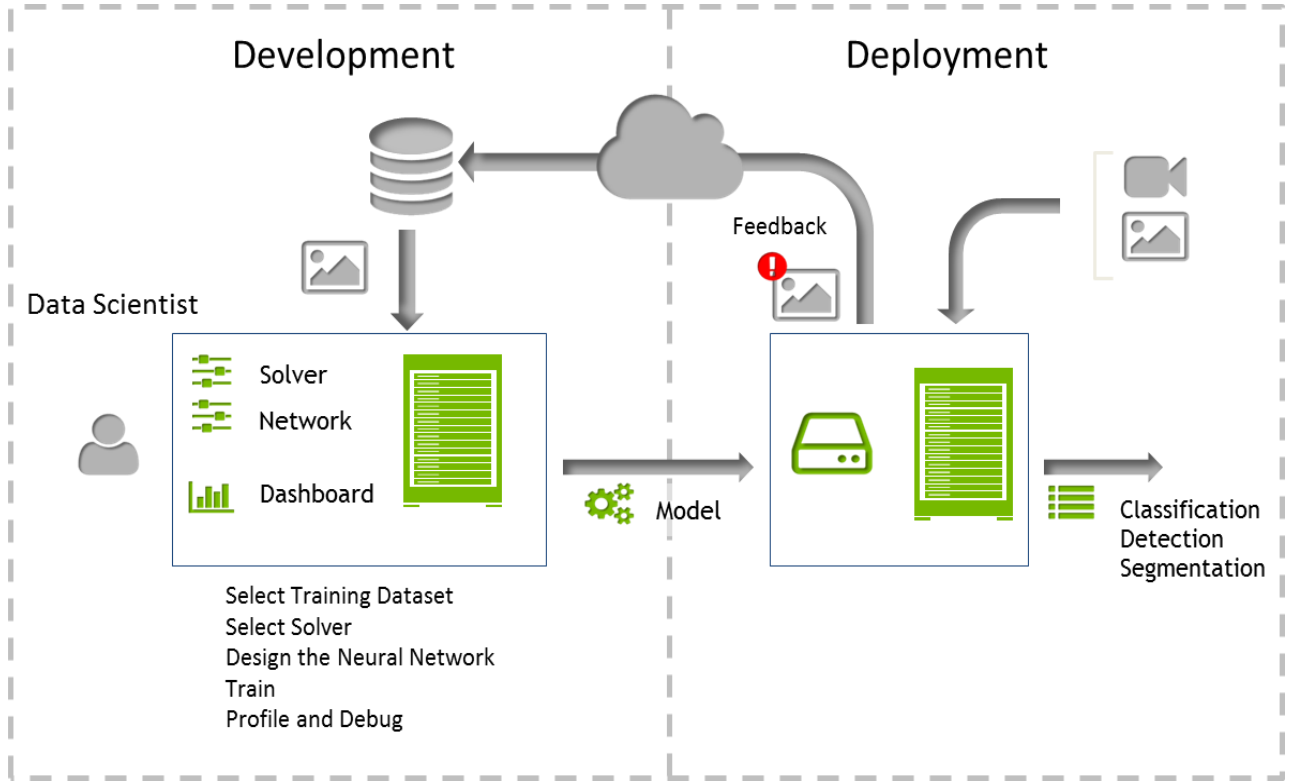

THANK YOU

sbyun@nvidia.com