MATLAB EXPO 2018

Demystifying Deep Learning

Paola.Jaramillo@mathworks.nl Application Engineer, MathWorks Benelux

Matthias.Sommer@mathworks.ch Application Engineer, MathWorks Switzerland

Agenda

- What’s Deep Learning and why should I care?
- A practical approach to Deep Learning (for images)
  – Transfer the learning from an expert model to your own application
- Building a Deep Learning network from scratch
  – Deep Learning for time series and text data
- Key learnings of the session and cool features

Deep learning with MATLAB is easy and accessible!

Use Less Data and Less Time with Transfer Learning

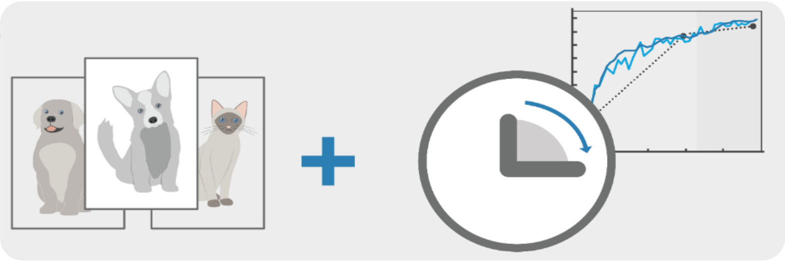

Automatically generate code!

Image Pre-processing | CNN | Image Post-processing

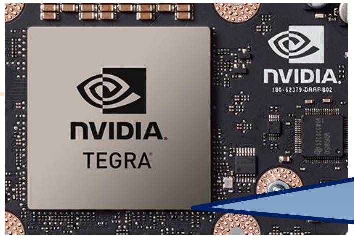

Prototype and productize your entire workflow on embedded GPUs

Apply Deep Learning not only on images but also on text and time series data

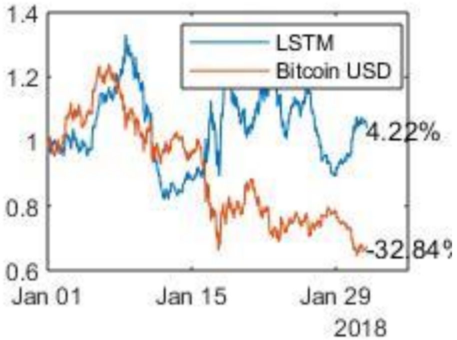

Why Machine Learning or Deep Learning?

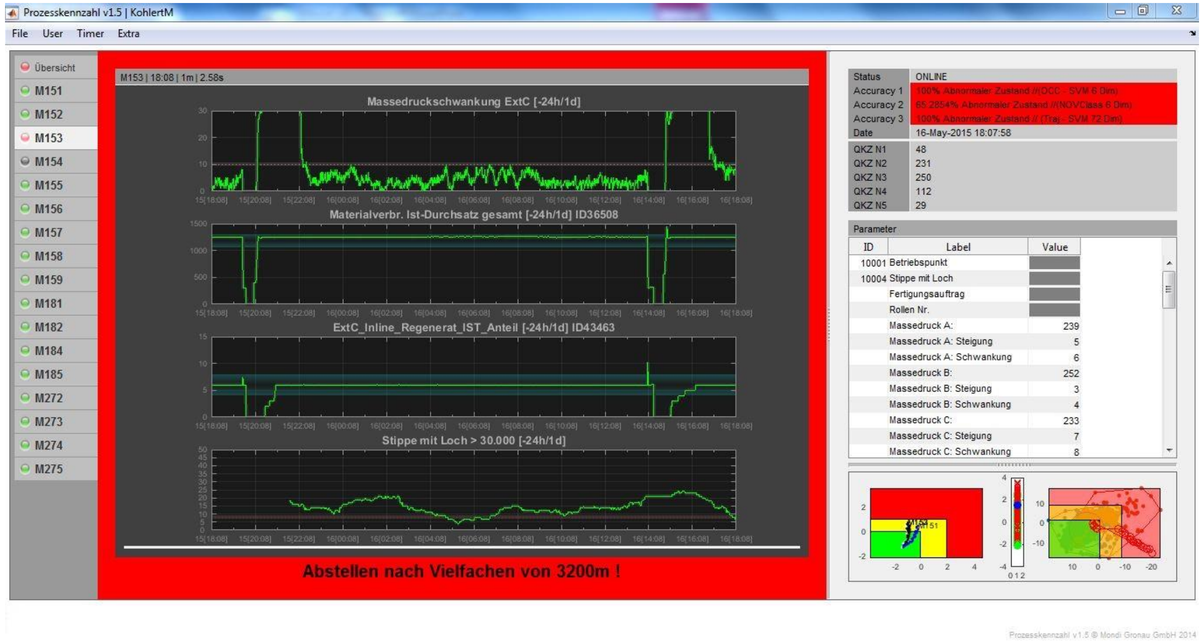

Transformational technology

Close (better) than human accuracy for specific tasks

Performance scales with data

It is hard to use, it is challenging

Enables
engineers, researchers and other domain experts
to create
products and applications with more
built-in intelligence

Artificial Intelligence, Machine Learning and Deep Learning

Application Breadth

Artificial Intelligence

Machine Learning

Deep Learning

AlphaGO	Lee Se-dol
1202 CPUs, 176 GPUs,	1 Human Brain
100+ Scientists.	1 Coffee.

1950s 1997 2016 Today

Timeline

Machine Learning vs Deep Learning

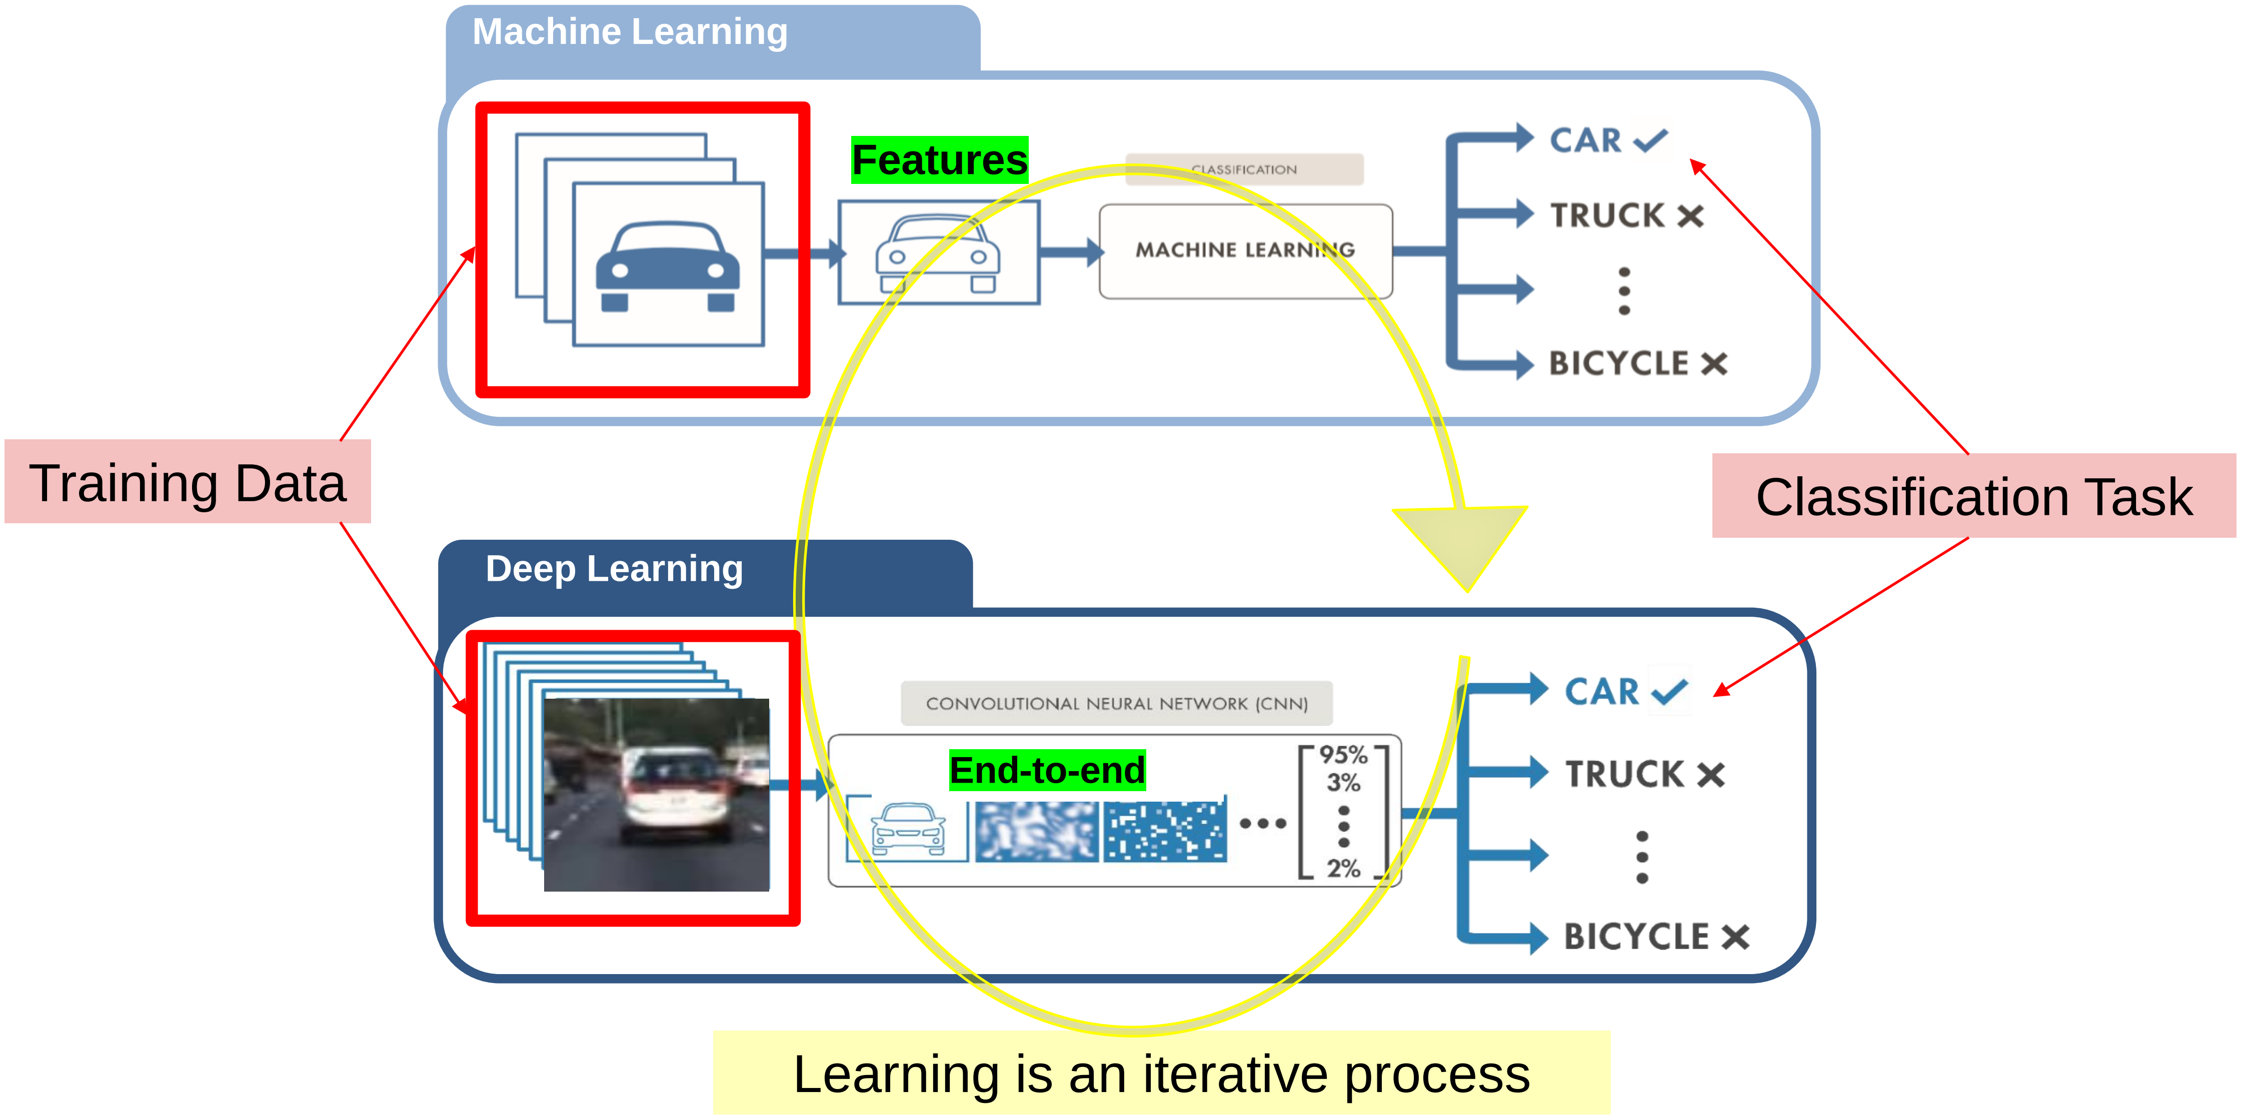

Deep Learning Common Workflow

ACCESS AND EXPLORE DATA → **LABEL AND PREPROCESS DATA** → **DEVELOP PREDICTIVE MODELS** → **INTEGRATE MODELS WITH SYSTEMS**

Access and Explore Data

Files

Databases

Sensors

Label and Preprocess Data

Data Augmentation/ Transformation

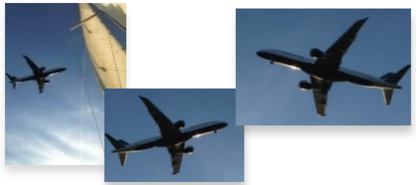

Labeling Automation

Import Reference Models

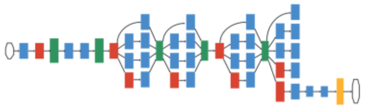

Develop Predictive Models

Hardware-Accelerated Training

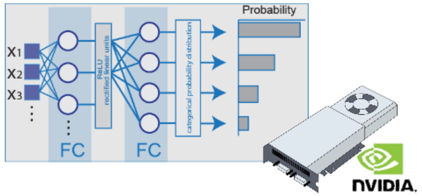

Hyperparameter Tuning

Network Visualization

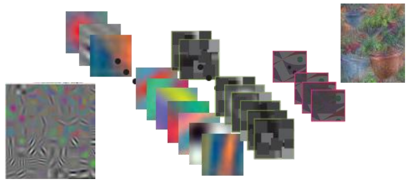

Integrate Models with Systems

Desktop Apps

Enterprise Scale Systems

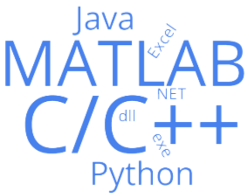

Embedded Devices and Hardware

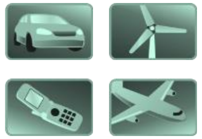

Deep learning is usually implemented using a neural network architecture

- The term “deep” refers to the number of layers in the network—the more layers, the deeper the network.
- Data flows through network in layers, which provide transformation of data

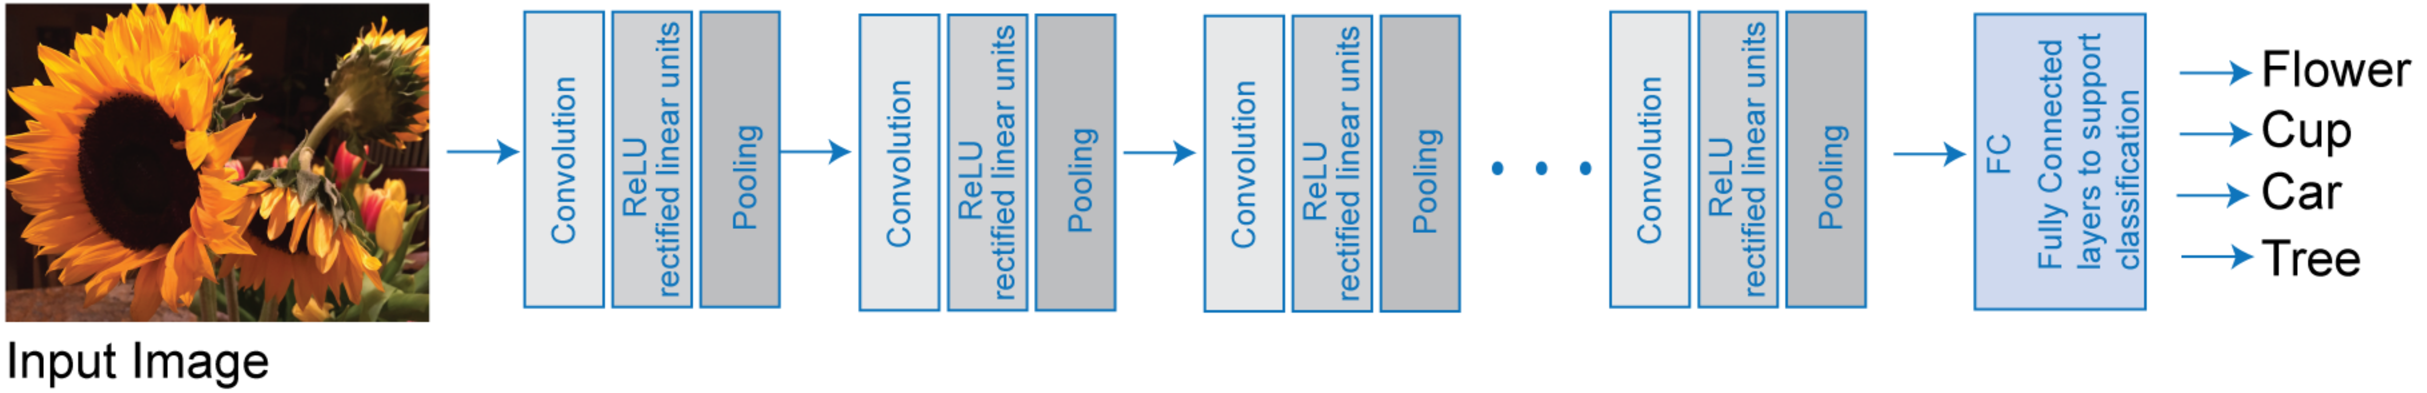

Convolutional Neural Network

Convolutional Neural Network is a popular Deep Learning architecture

Shallow Neural Network

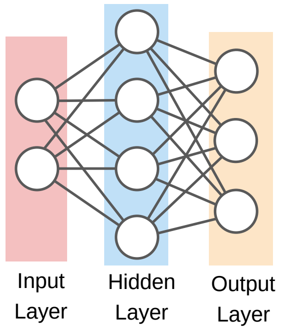

Every input neuron
Connects to every neuron
in the hidden layer

Convolutional Neural Network

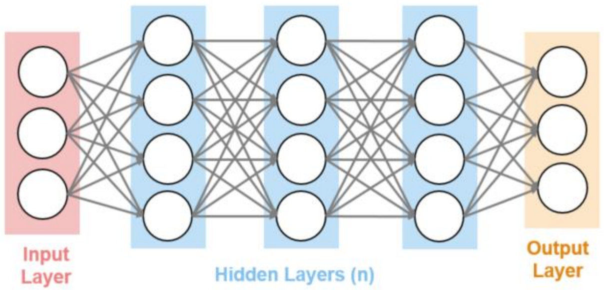

Local receptive fields
connect to neurons in the hidden layer

Translate across an image
create a feature map
efficiently with convolution

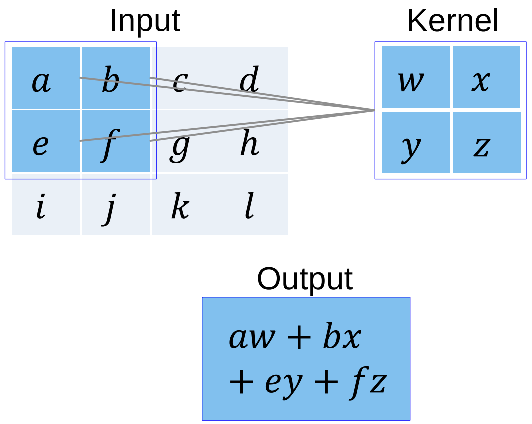

Deep Learning: different types of network architectures

Type	Series Network	Directly Acyclic Graph Network (DAG)	Recurrent Neural Network (RNN)
Basic Architecture	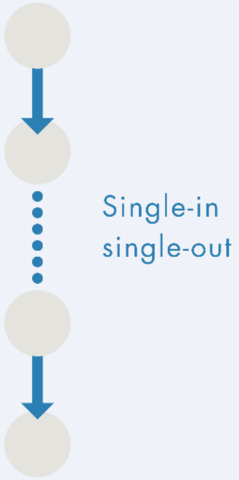	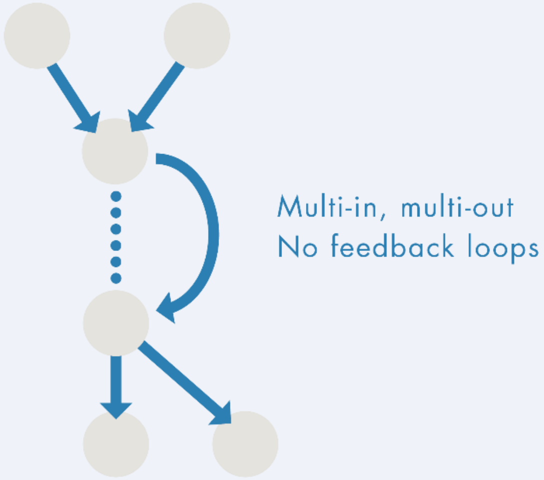	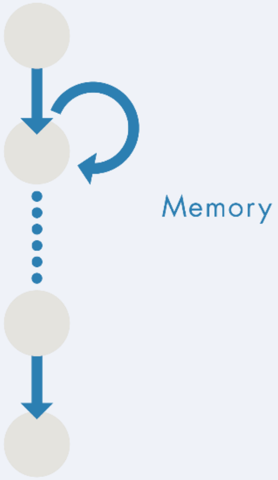
Network Example	AlexNet, VGG	R-CNN (fast, faster), GoogLeNet, SegNet	LSTM
Application	Object Recognition: Cars, Lane, Pedestrian	Object Detection: Cars, Traffic Signs, Scene (Semantic Segmentation)	Sequential data: time series, signals

Agenda

- What’s Deep Learning and why should I care?
- A practical approach to Deep Learning (for images)
  – Transfer the learning from an expert model to your own application
- Building a Deep Learning network from scratch
  – Deep Learning for time series and text data
- Key learnings of the session and cool features

Deep Learning can be complex and challenging to apply

Design Deep Learning & Vision Algorithms

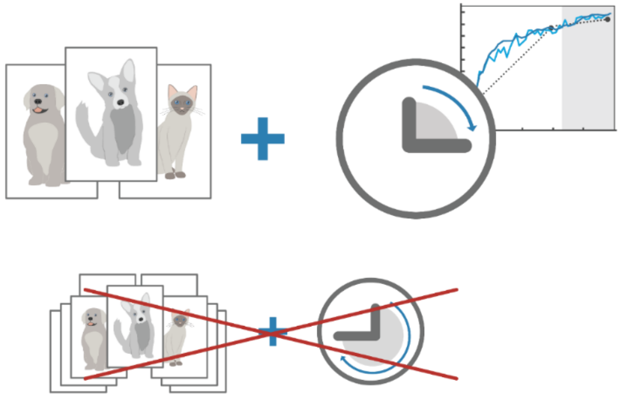

Accelerate and Scale Training

High Performance Deployment

Main Challenges

- Handle large image sets
- Image labeling is tedious
- Have access to models

Main Challenges

- Capability of training with multiple GPUs
- Capability of training in the cloud

Main Challenges

- Convert models to CUDA code
- Compress models to fit into embedded GPUs.

Object Recognition using Deep Learning in MATLAB

Supervised Learning

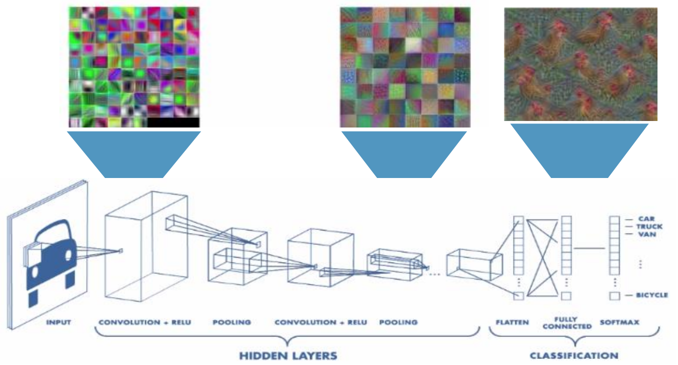

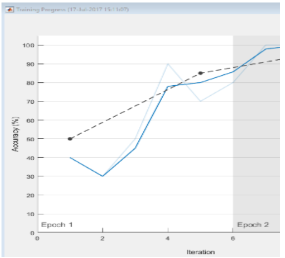

Stapler

CoffeCoaster

Highlighter

RubikCube

AlexNet

PRETRAINED MODEL

Transfer Knowledge

OfficeNet

NEW MODEL

Where to start?
Two Approaches for Deep Learning

Train from Scratch

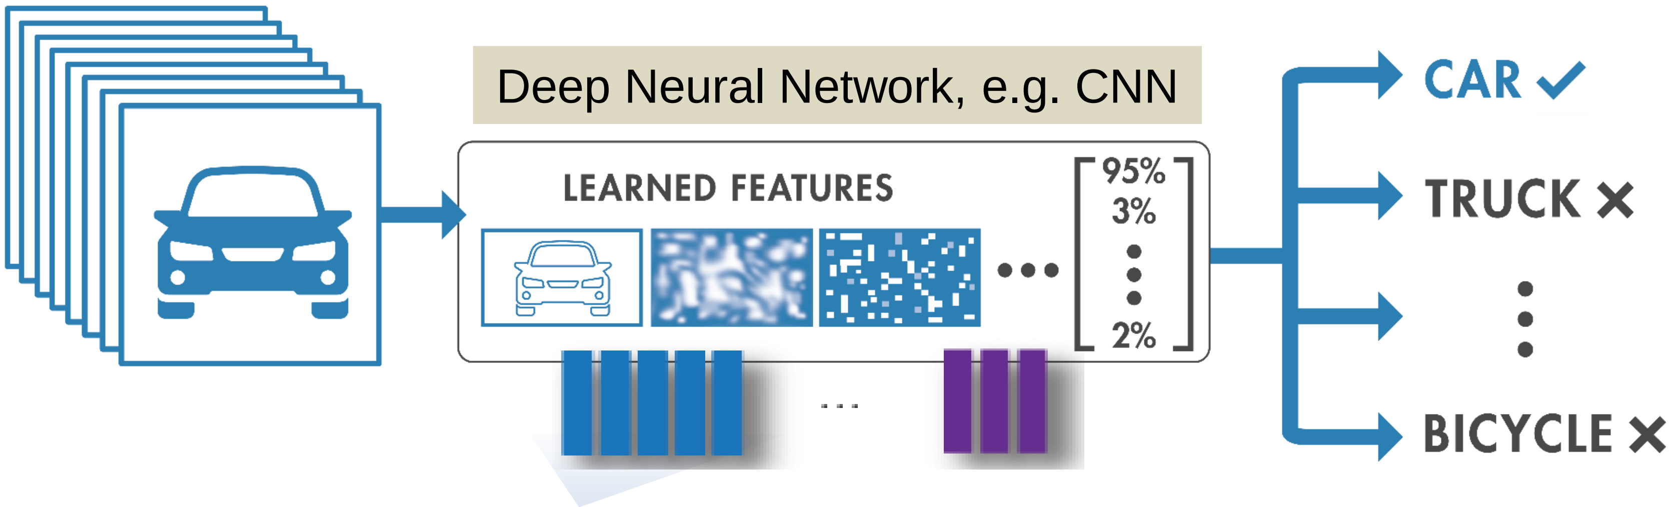

Transfer Learning

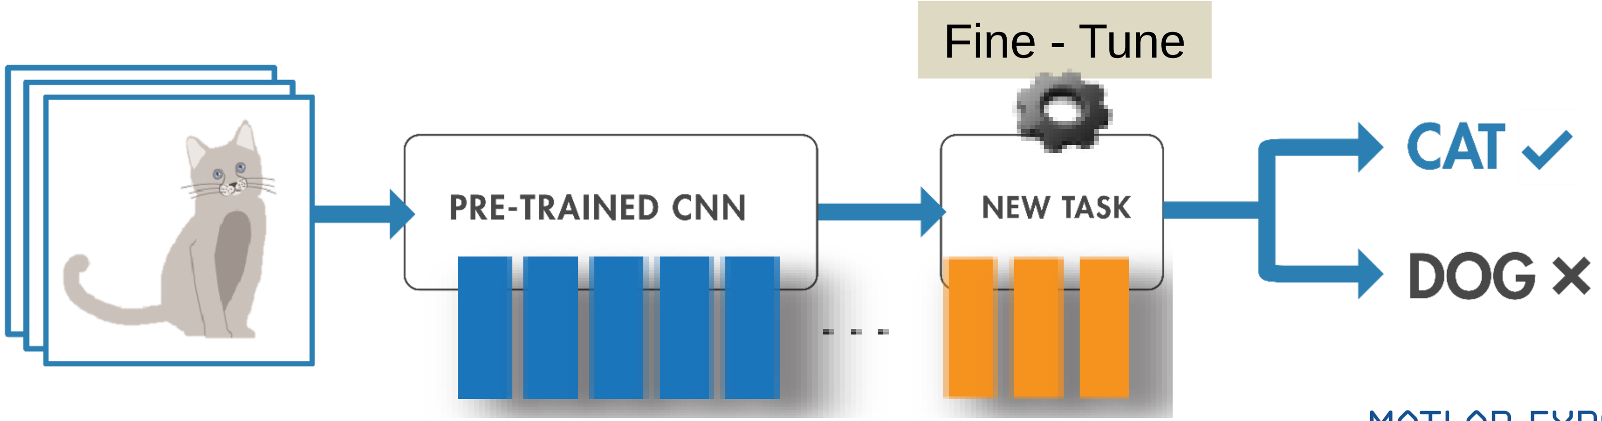

Why Perform Transfer Learning?

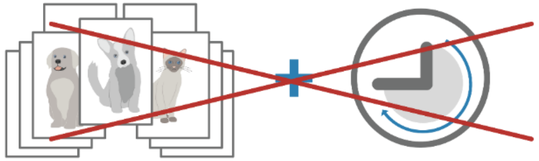

- Less data
- Less training time

- Leverage best network types from top researchers
- Reference models are great feature extractors

AlexNet PRETRAINED MODEL

VGG-16 PRETRAINED MODEL

ResNet PRETRAINED MODEL

Caffe MODELS

GoogLeNet PRETRAINED MODEL

TensorFlow/Keras MODELS

Transfer Learning Workflow

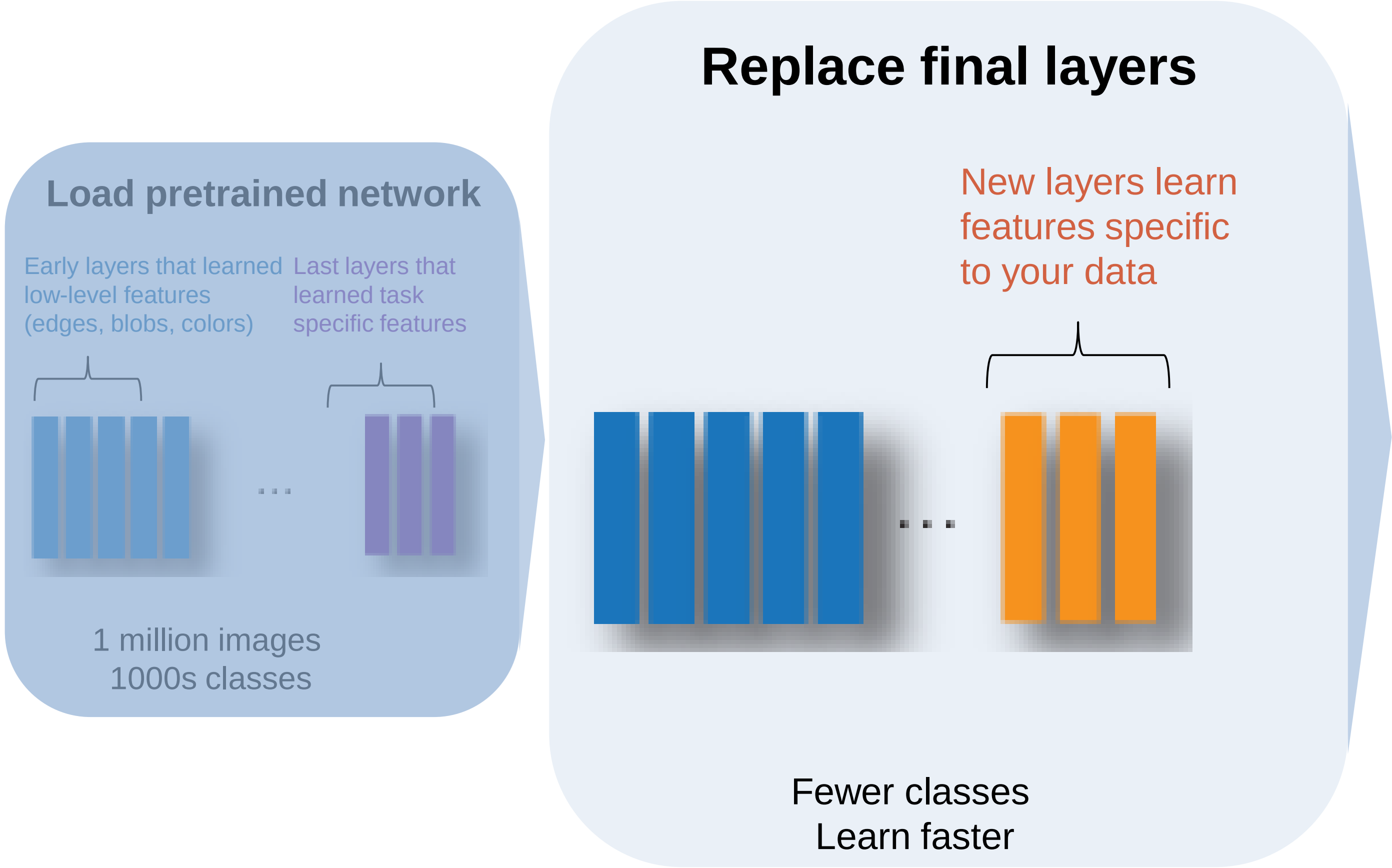

Transfer Learning Workflow

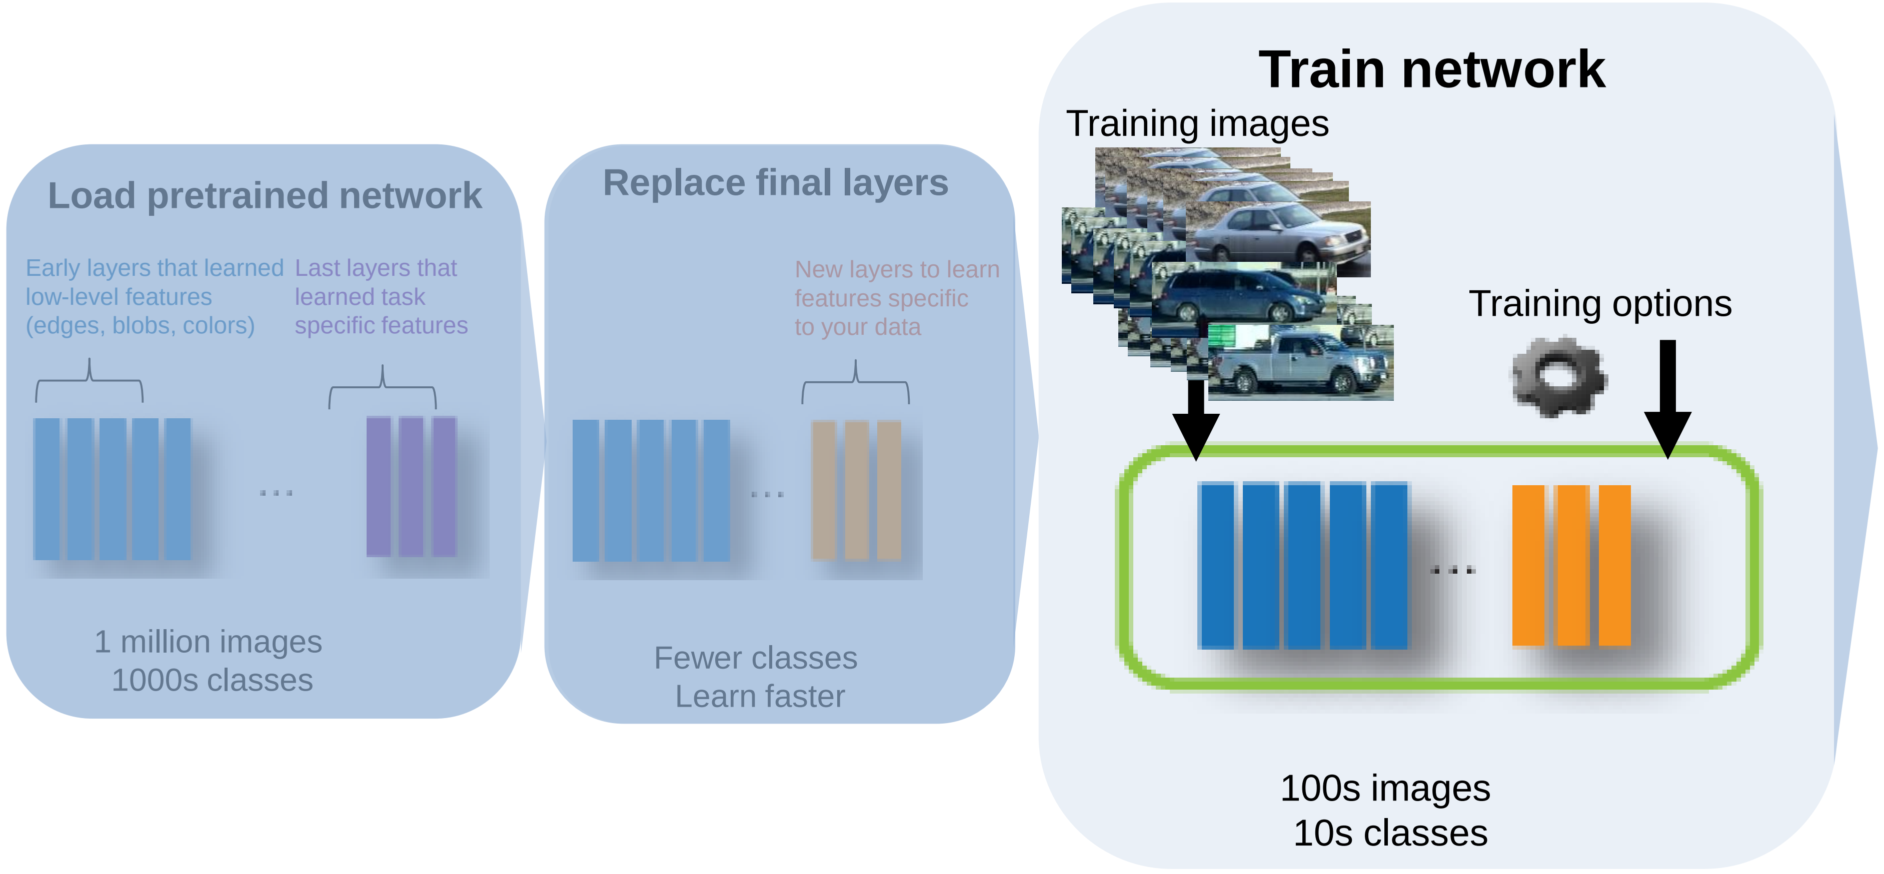

Transfer Learning Workflow

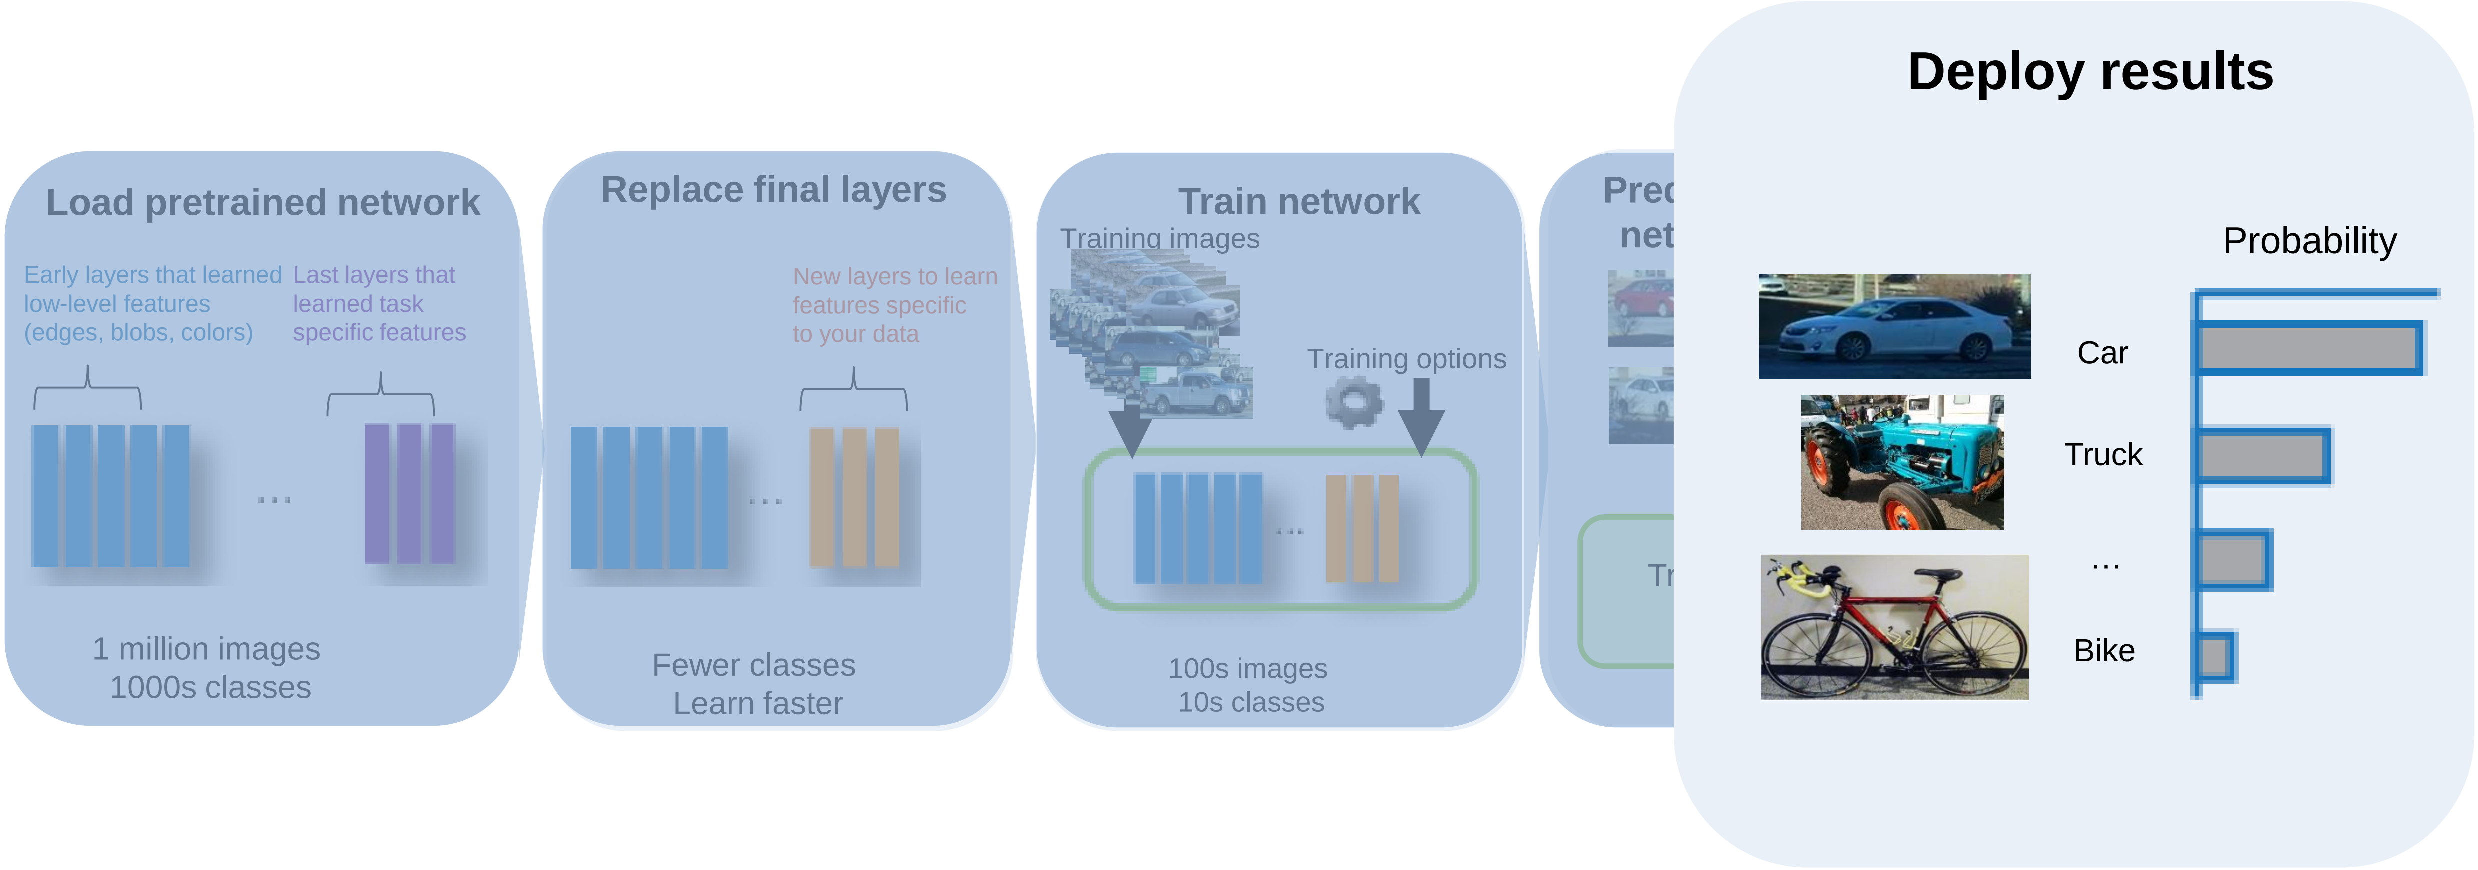

Transfer Learning Workflow

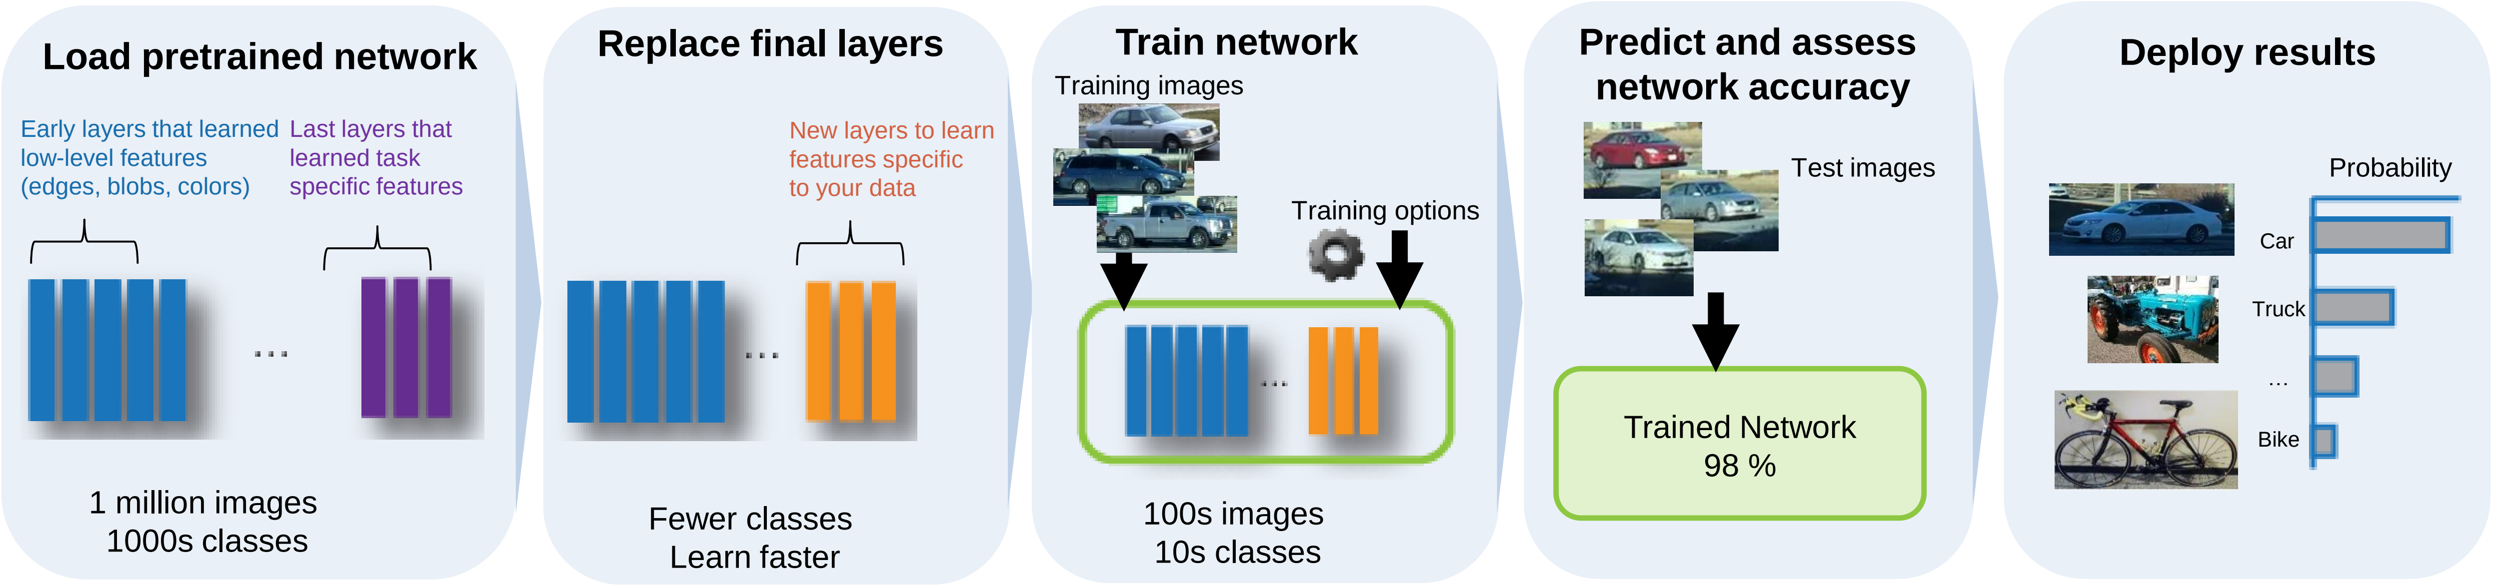

Accelerate training and prediction!

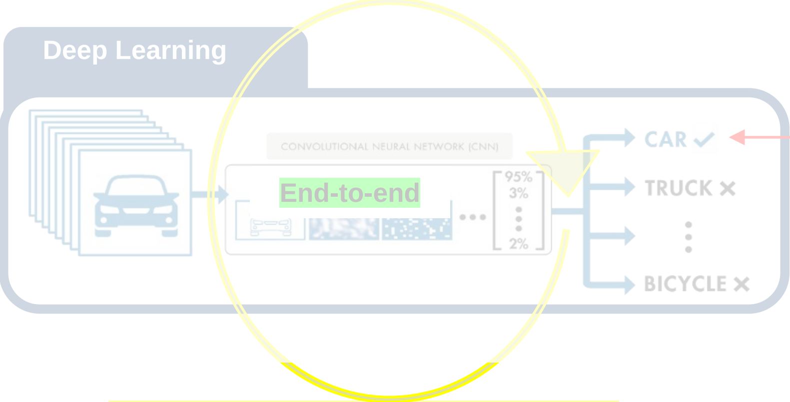

Prediction

Learning is an iterative process

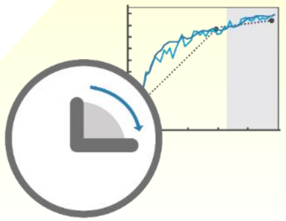

```
opts = trainingOptions('sgdm', ...
    'MaxEpochs', 100, ...
    'MiniBatchSize', 250, ...
    'InitialLearnRate', 0.00005, ...
    'ExecutionEnvironment', 'auto' );
```

```
'ExecutionEnvironment', 'multi-gpu' );
```

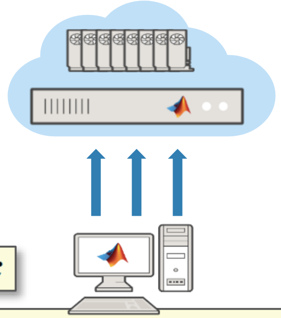

```
'ExecutionEnvironment', 'parallel' );
```

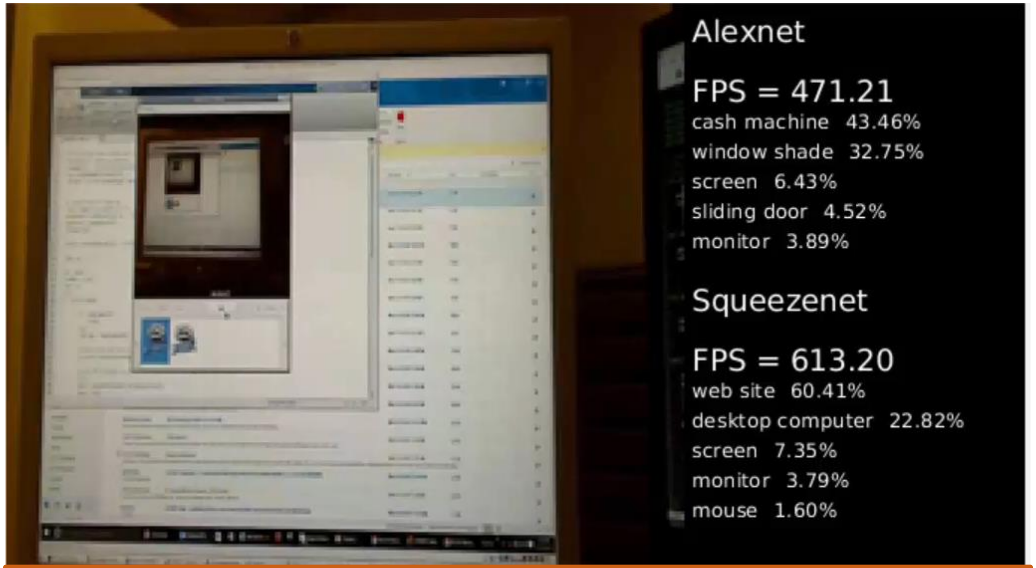

Alexnet vs Squeezenet

More GPUs

[Whitepaper: Deep Learning in the Cloud]

Alexnet Inference on NVIDIA Titan Xp

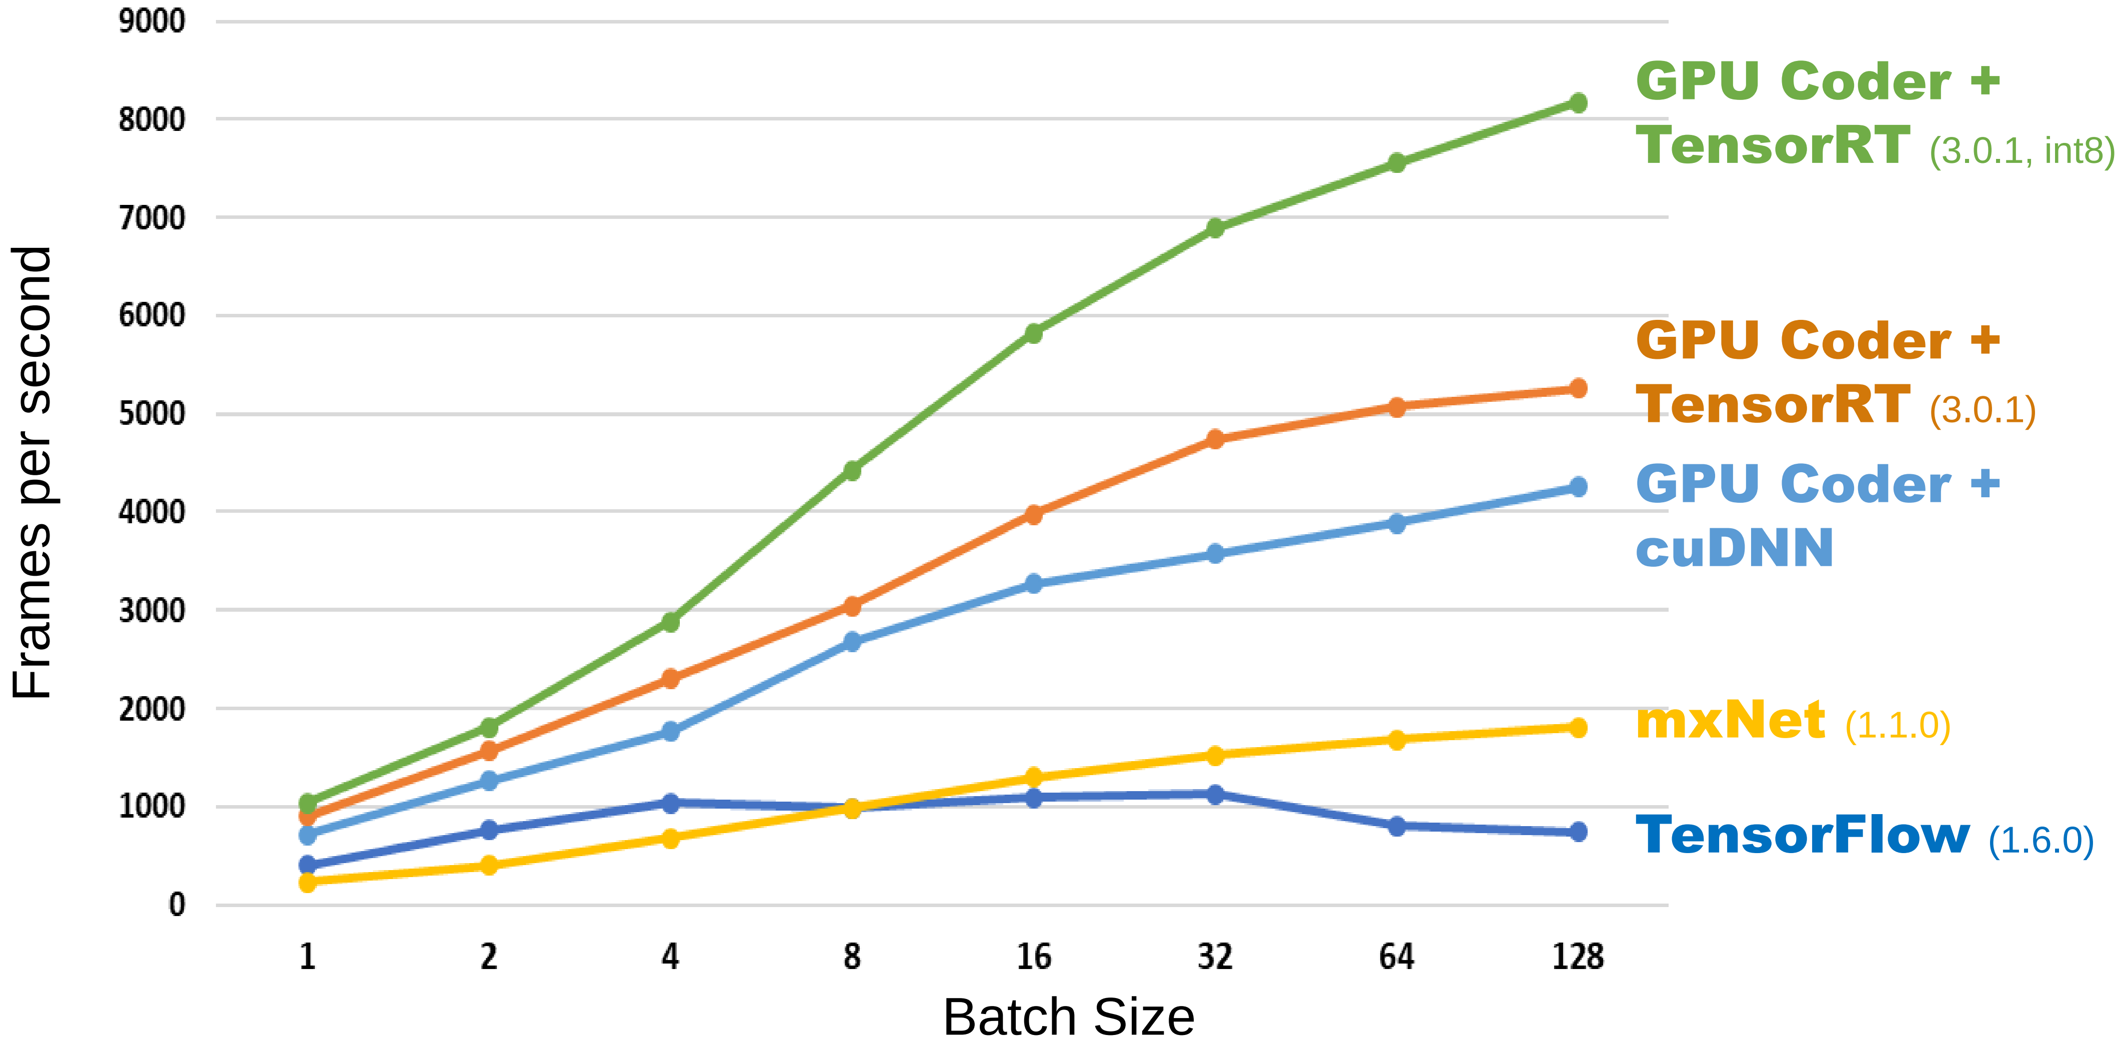

Testing platform

CPU	Intel(R) Xeon(R) CPU E5-1650 v4 @ 3.60GHz
GPU	Pascal Titan Xp
cuDNN	v7

Deep Learning is easy and accessible with MATLAB!

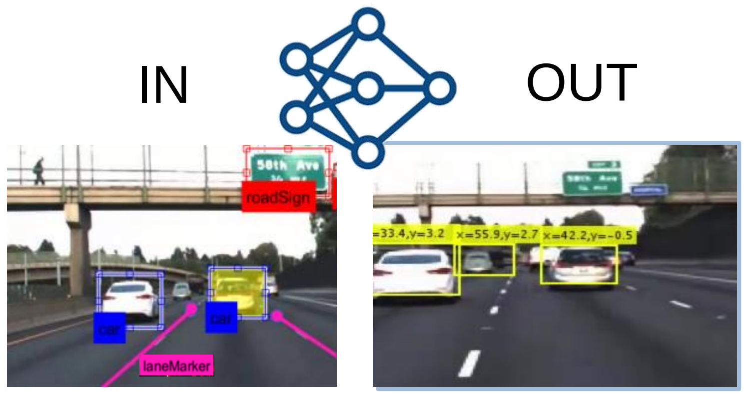

Design Deep Learning & Vision Algorithms

Accelerate and Scale Training

MKL-DNN

TensorRT & cuDNN

Neon

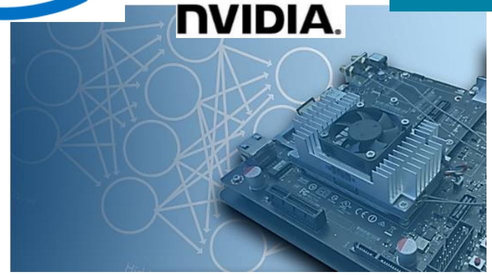

High Performance Deployment

Highlights

- **Datastores** for large image sets
- **Automate** image labeling
- **Direct access** to models within MATLAB with support packages
- **Import** Tensor Flow Keras and Caffe networks

Highlights

- **Single line of code** to:
- **Accelerate** training with multiple GPUs or
- **Scale** to clusters

Highlights

- **Automate compilation** with GPU Coder
- **1.4x speedup** over C++ Caffe on Jetson TX2

Agenda

- What's Deep Learning and why should I care?
- A practical approach to Deep Learning (for images)
  - Transfer the learning from an expert model to your own application
- Building a Deep Learning network from scratch
  - Deep Learning for time series and text data
- Key learnings of the session and cool features

Deep Learning for Time Series, Sequences and Text

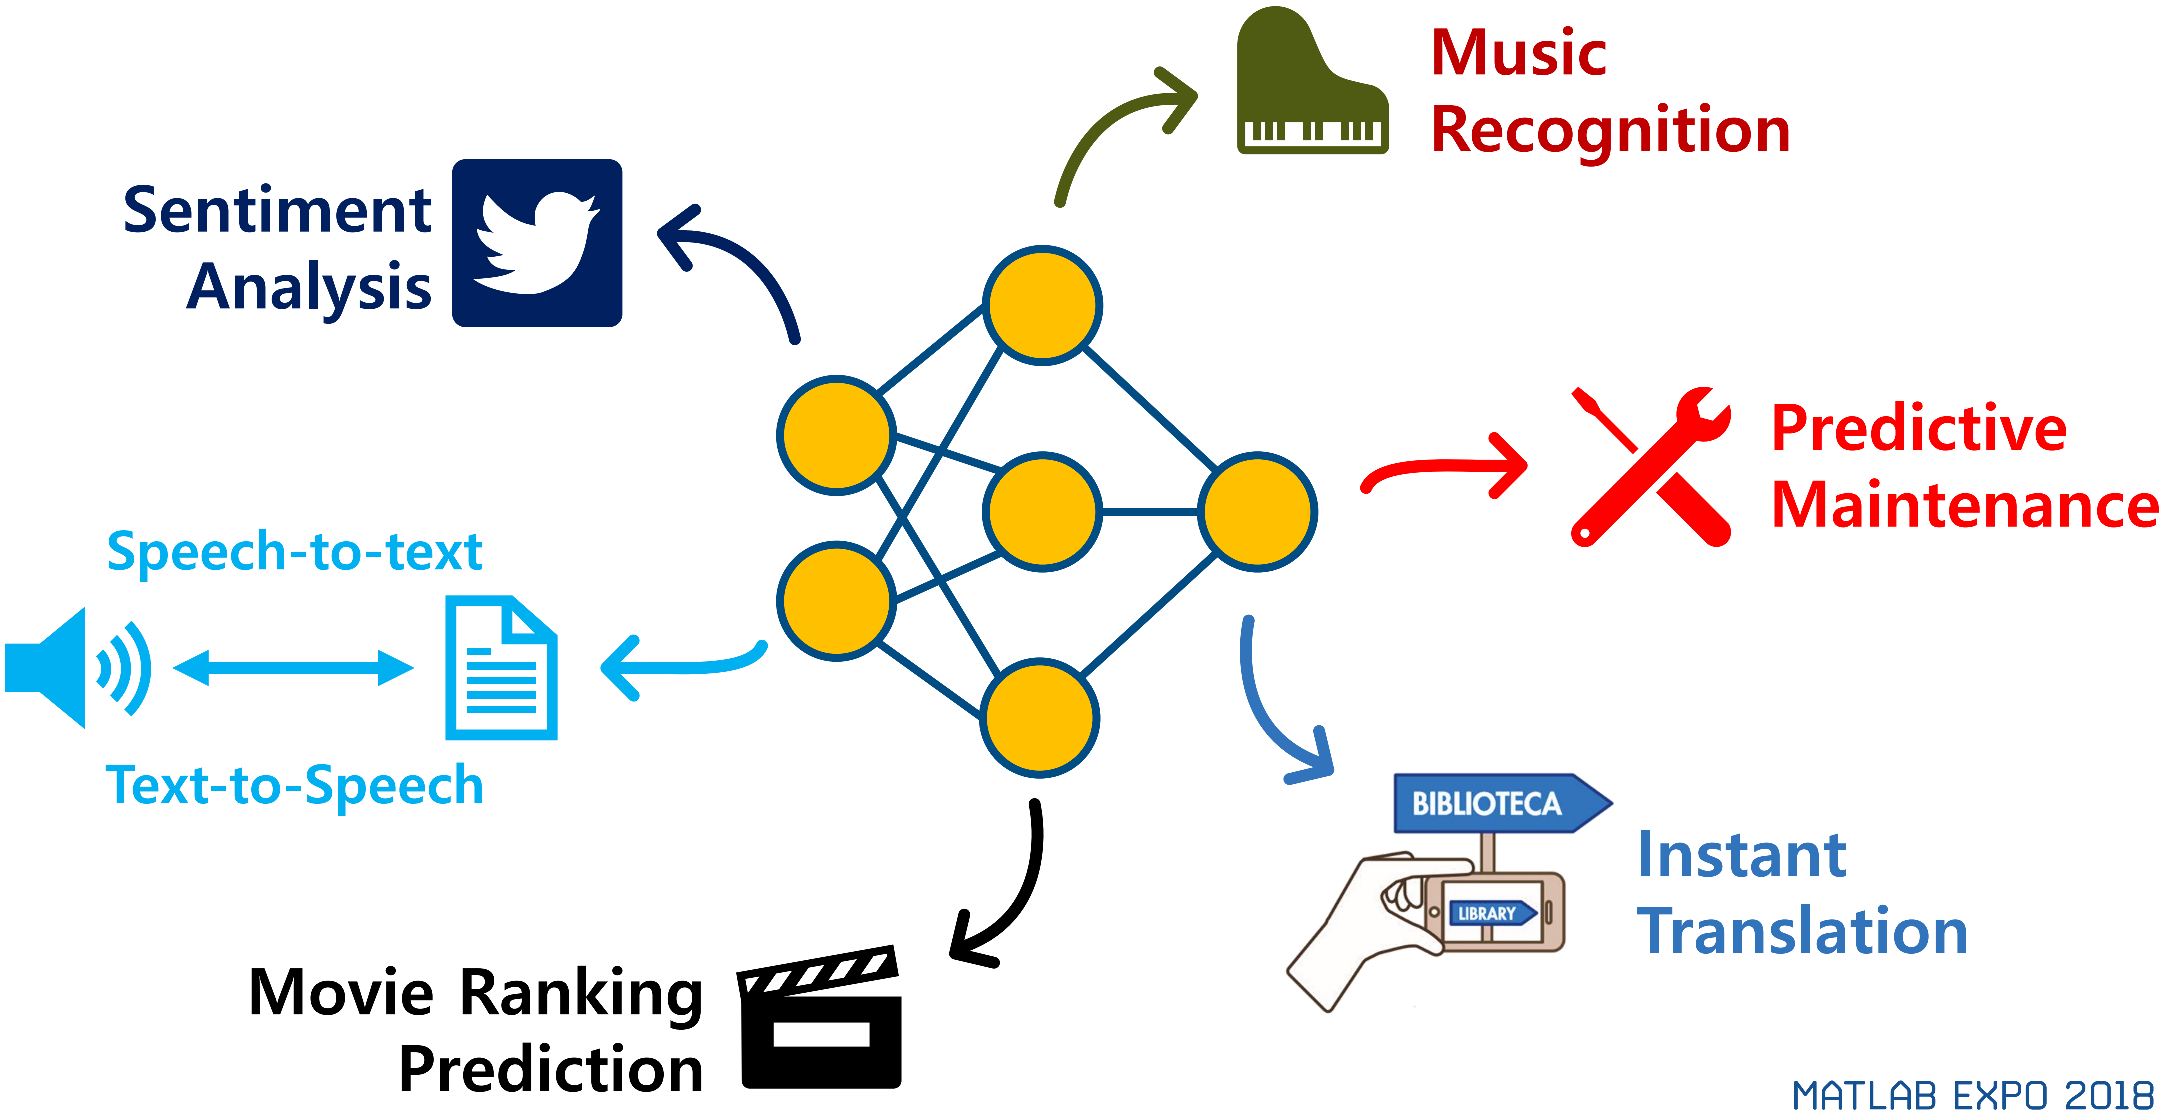

Deep Learning for Time Series

Example: Seizure prediction (time series classification)

THE UNIVERSITY OF MELBOURNE

Goal: Predict seizures in long-term

Dataset: iEGG time series

Data size: 20GB

Output: Classification
before or between seizure

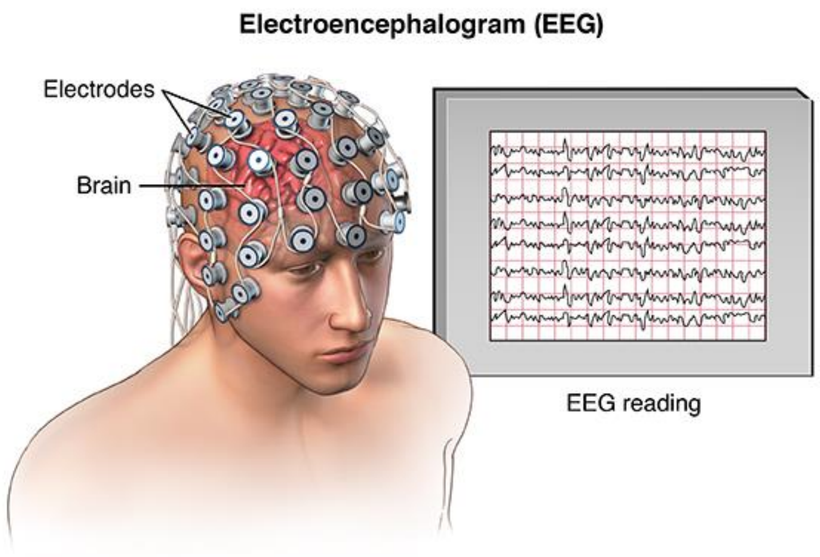

[Link for melbourne-university-seizure-prediction]()

Types of Datasets

Numeric Data

ID	WC_TA	RE_TA	EBIT_TA	MVE_BVTD	S_TA	Industry	Rating
62394	0.013	0.104	0.036	0.447	0.142	3	BB
48608	0.232	0.335	0.062	1.969	0.281	8	A
42444	0.311	0.367	0.074	1.935	0.366	1	A
48631	0.194	0.263	0.062	1.017	0.228	4	BBB
43768	0.121	0.413	0.057	3.647	0.466	12	AAA
39255	-0.117	-0.799	0.01	0.179	0.082	4	CCC
62236	0.087	0.158	0.049	0.816	0.324	2	BBB
39354	0.005	0.181	0.034	2.597	0.388	7	AA
40326	0.47	0.752	0.07	11.596	1.12	8	AAA
51681	0.11	0.337	0.045	3.835	0.812	4	AAA

Long Short Term Memory (LSTM)

Time Series/ Text Data

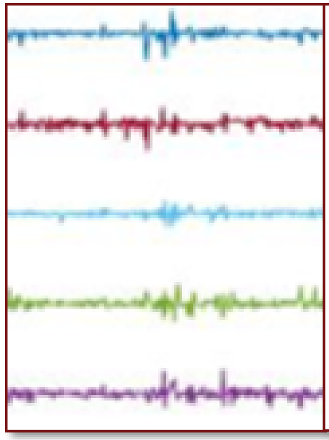

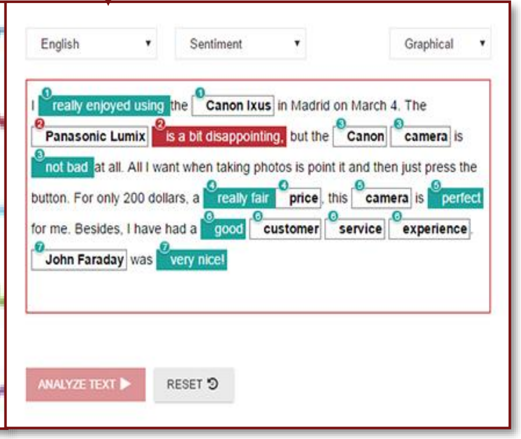

LSTM or CNN

Image Data

Convolutional Neural Networks (CNN)
Directed acyclic graph networks (DAG)

Deep Learning for Time Series

CNN: Data for Time series = Pixel for Images

0.2	-0.5	-1	-2.1
-1.3	0.8	1.1	-2
1.2	0	1.2	1.6
0.8	0.7	-0.2	-0.4

1	0	5	4
3	4	8	3
1	4	6	5
2	5	4	1

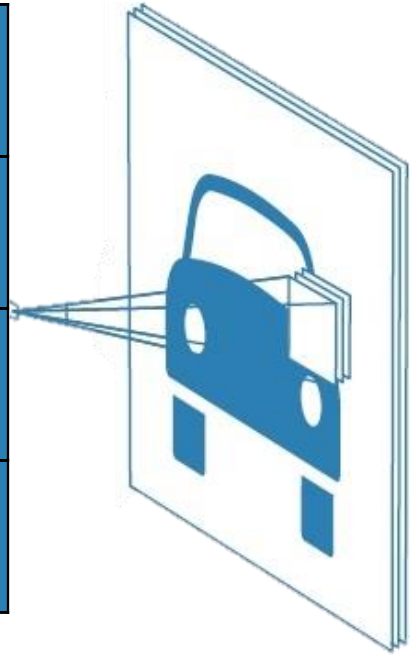

Deep Learning for Time Series

Long Short Term Memory (LSTM) Network

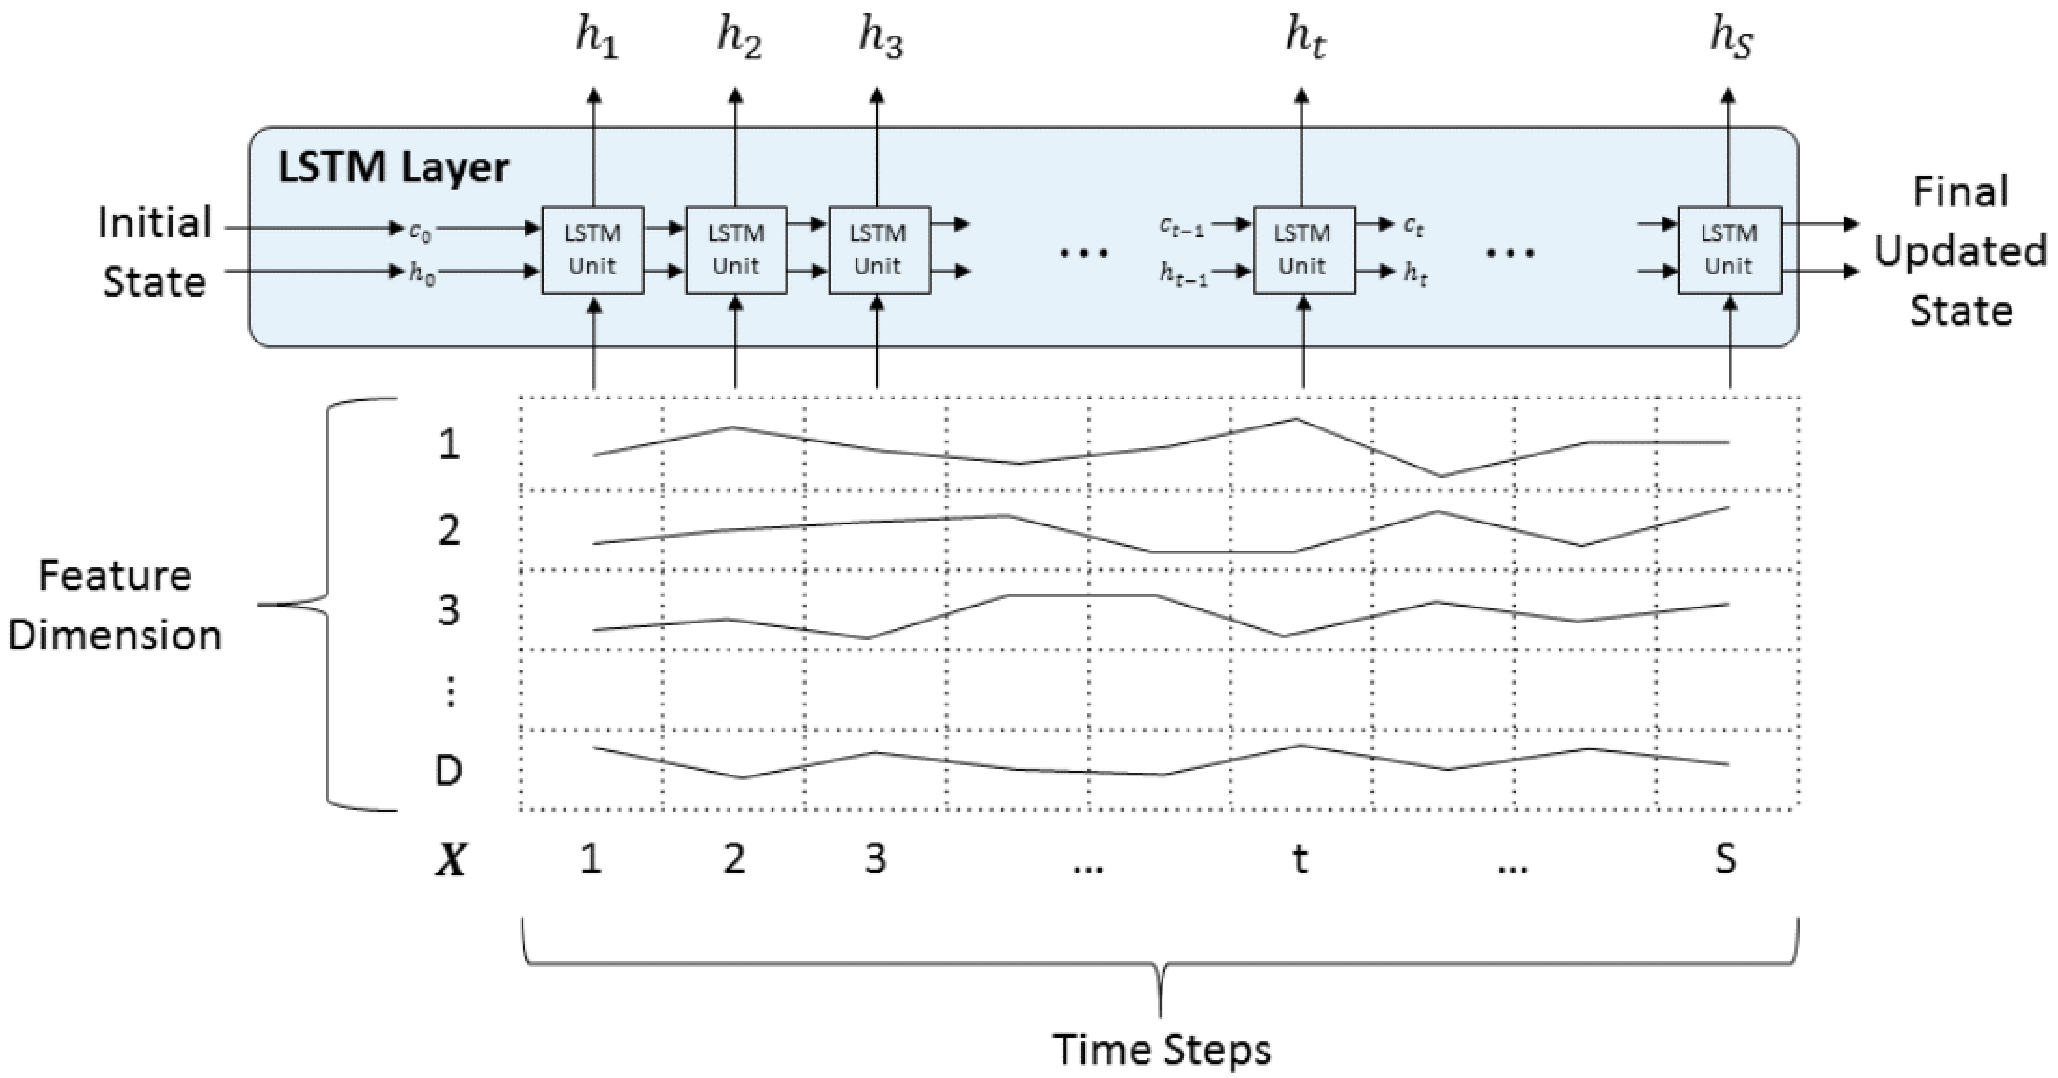

Text Analytics

Access and Explore Data → **Preprocess Data** → **Develop Predictive Models**

Clean-up Text → **Convert to Numeric**

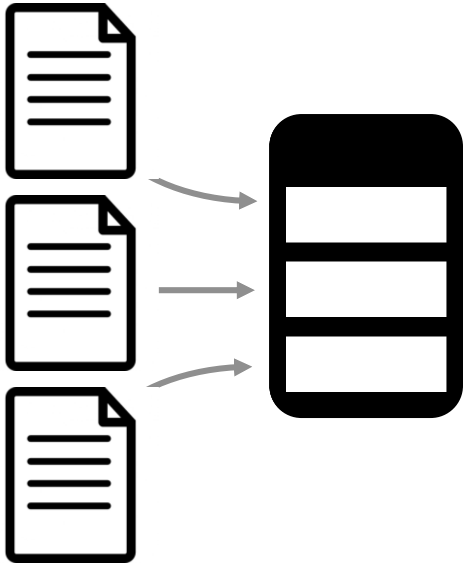

Media reported two trees blown down along I-40 in the Old Fort area.

media report two tree blown down i40 old fort area

	cat	dog	run	two
doc1	1	0	1	0
doc2	1	1	0	1

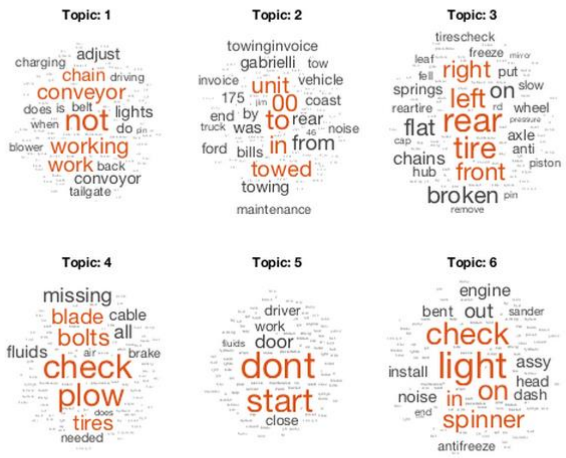

- Word Docs
- PDF’s
- Text Files

- Stop Words
- Stemming
- Tokenization

- Bag of Words
- TF-IDF
- Word Embeddings

- LSTM
- Latent Dirichlet Allocation
- Latent semantic analysis

Key Takeaways:
Deep Learning for Time Series and Text

- Applications
  - Time series
    - forecasting
    - classification (Predictive Maintenance)
  - Text
    - classification (Sentiment Analysis, Tagging)
    - clustering (Topic Modelling)
- Text Analytics: Prepend
  - Text preprocessing
  - Conversion to numeric

Agenda

- What’s Deep Learning and why should I care?
- A practical approach to Deep Learning (for images)
  – Transfer the learning from an expert model to your own application
- Building a Deep Learning network from scratch
  – Deep Learning for time series and text data
- Key learnings of the session and cool features

Deep learning with MATLAB is easy and accessible!

Use Less Data and Less Time with Transfer Learning

Automatically generate code!

Image Pre-processing

CNN

Image Post-processing

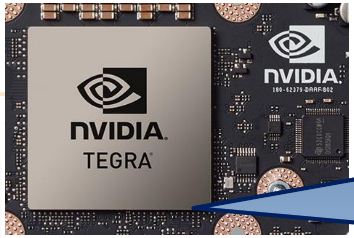

Prototype and productize your entire workflow on embedded GPUs

Apply Deep Learning not only on images but also on text and time series data

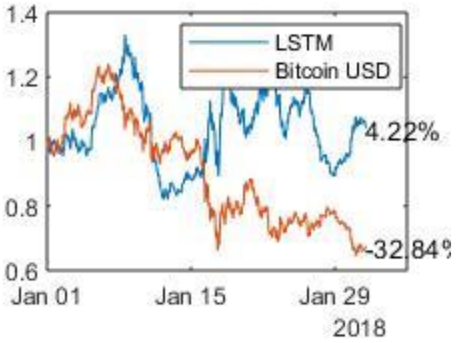

Algorithm Development, Testing & Verification

Neural Network Toolbox™
GPU Coder™

Images

Deep Learning

Plot and analyze your network using network analyzer, generate CUDA code that integrates with TensorRT, and deploy deep learning networks to Intel and ARM processors.

Semantic Scene Segmentation

Acceleration (Classification)

Deployment (Embedded)

Time Series

Mini-batchable datastore

Have a small number of large high-res images

To train, need a large number of pairs of images

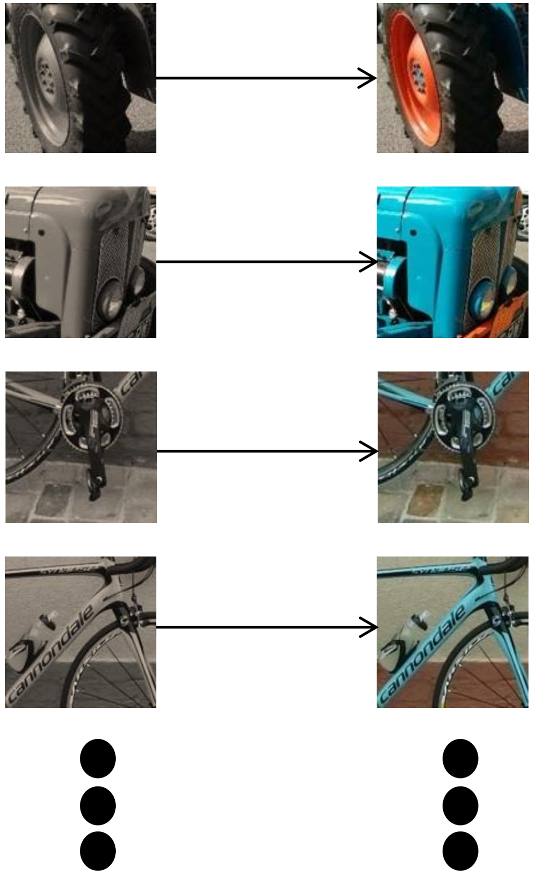

Ground truth labelling: attributes and sublabels

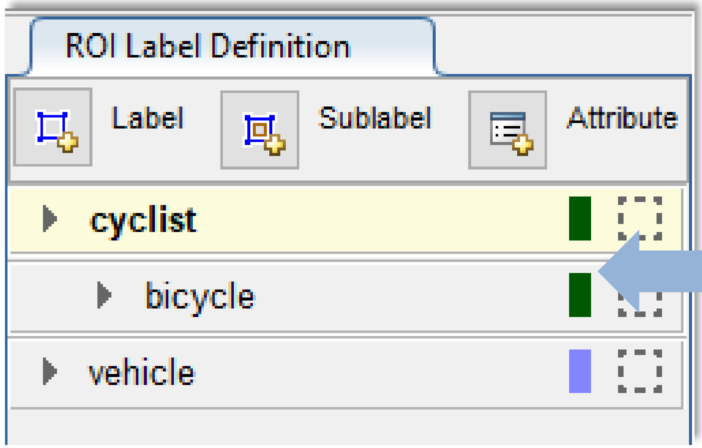

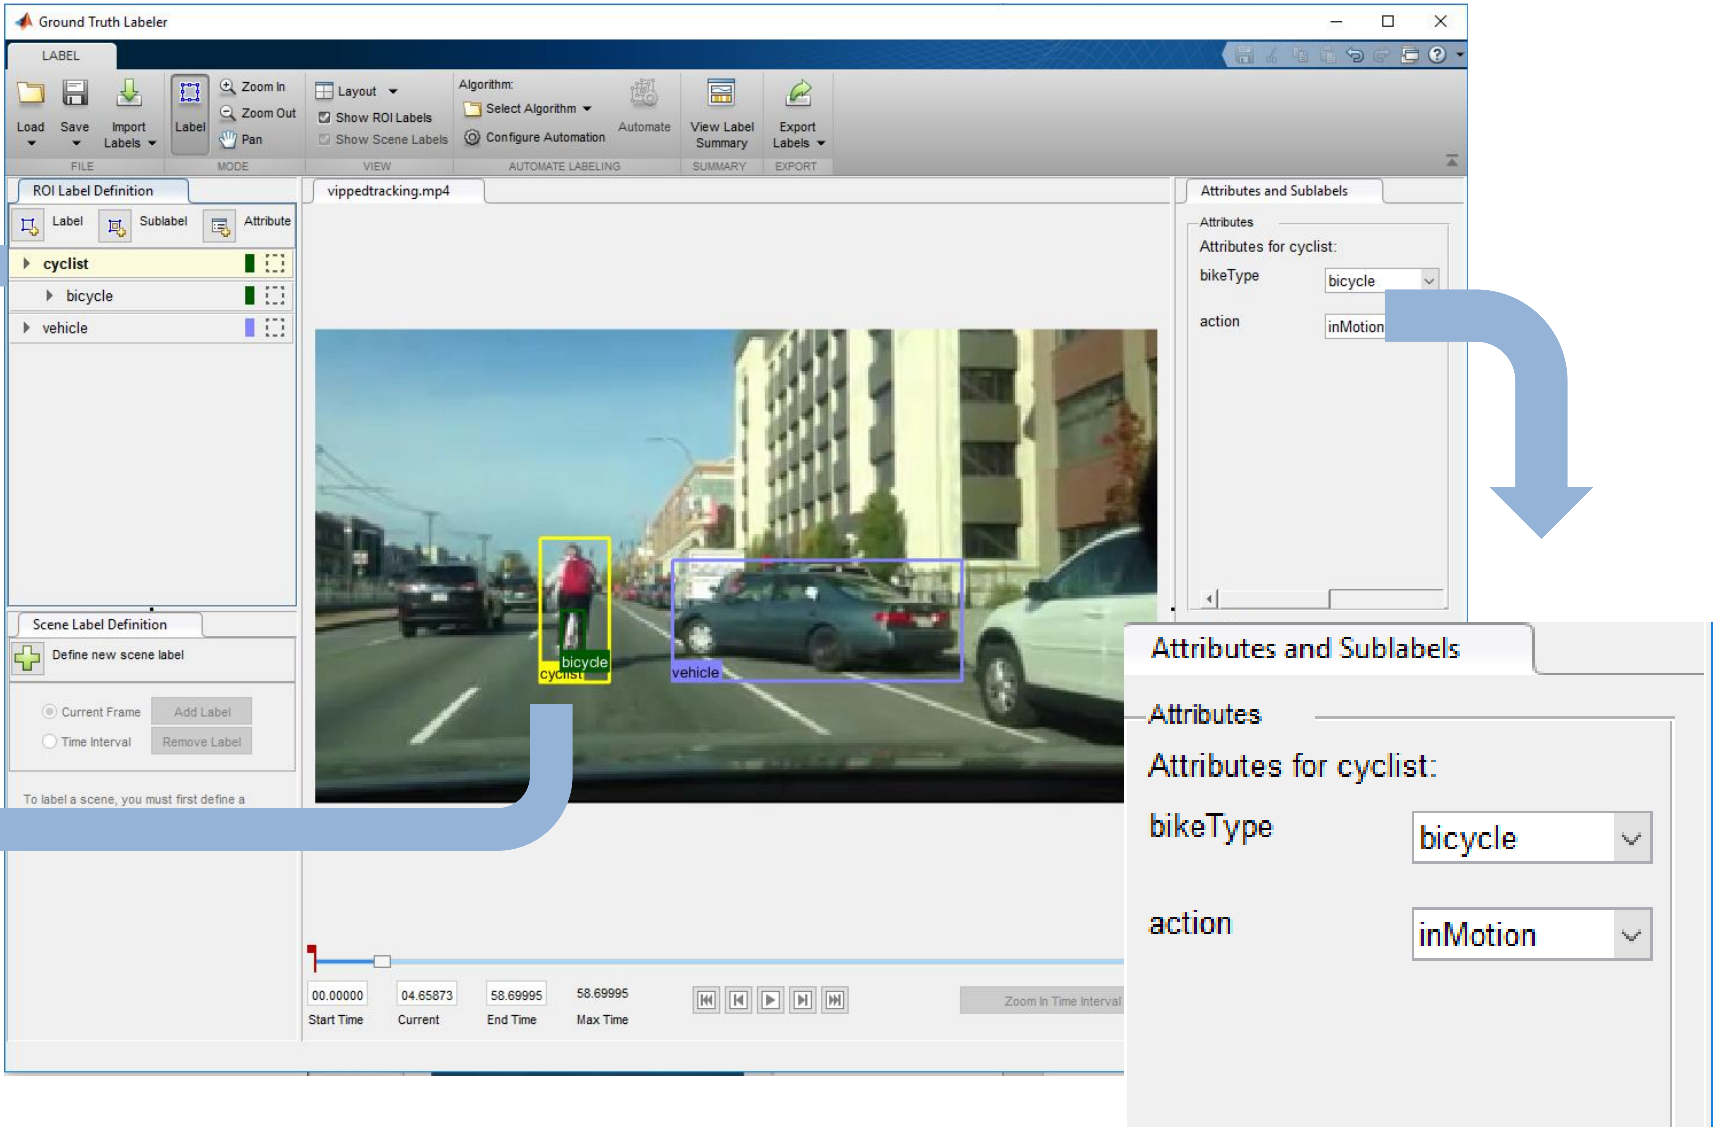

Visualize and understand the network architecture

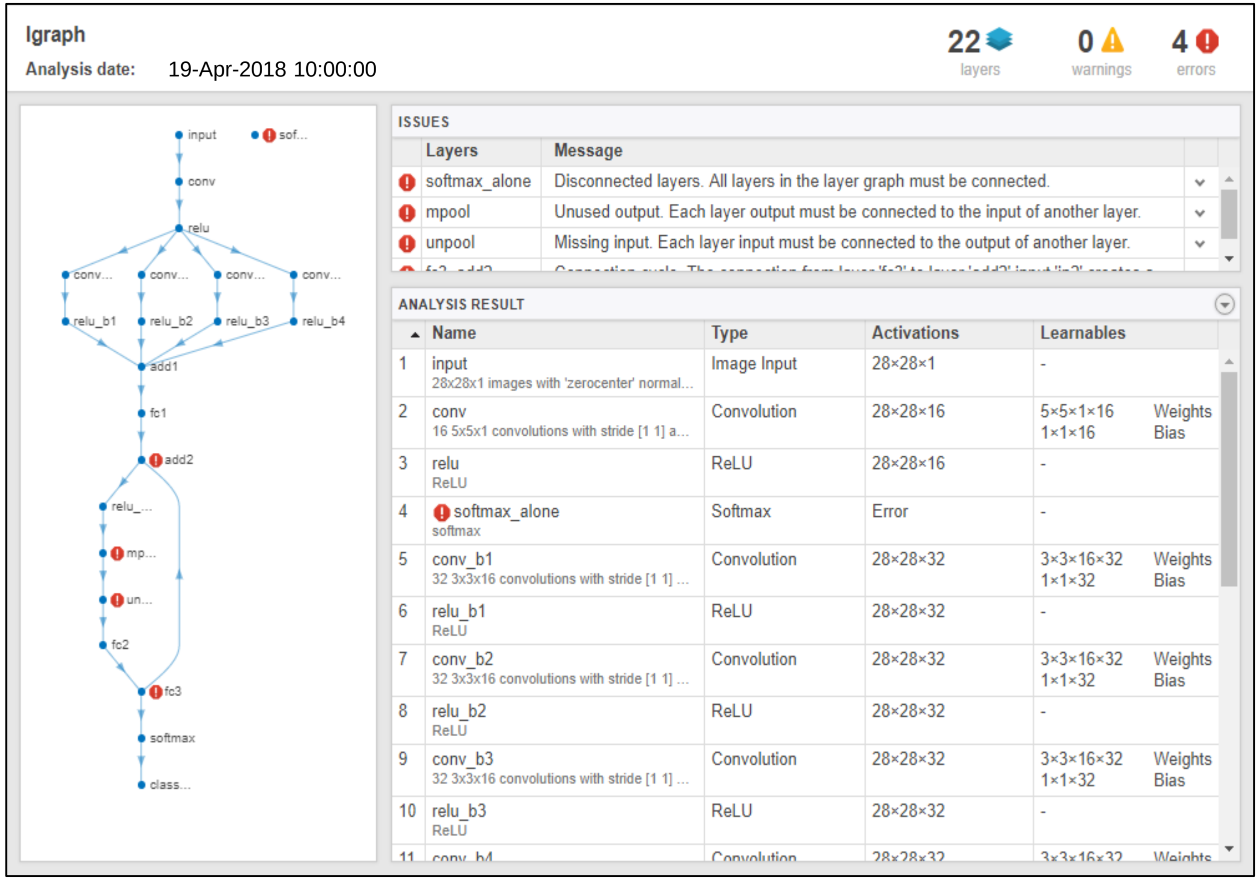

Detect problems before wasting time training!

- Missing or disconnected layers,
- Mismatching or incorrect sizes of layer inputs,
- Incorrect number of layer inputs,
- Invalid graph structures.

Network Analyzer
Visualize and Analyze network

Deep Learning Network Analyzer for Neural Network Toolbox

version 1.0 (15.1 KB) by MathWorks Neural Network Toolbox Team

Visualize and Analyze Deep Learning Networks

[Download Support Package]

Semantic Segmentation

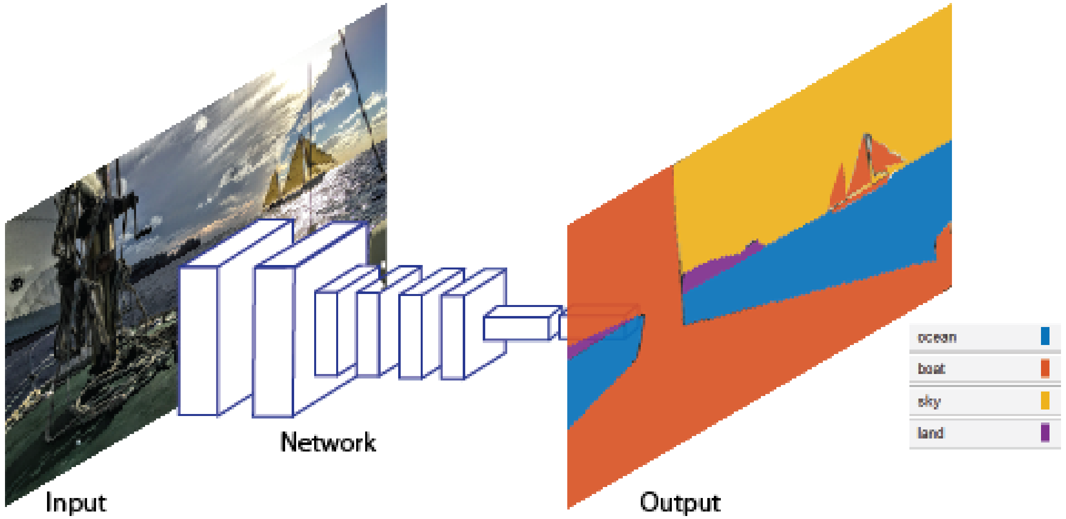

- Fully convolutional networks (FCN)
- Segmentation Networks (SegNet)
- Other directed acyclic graph (DAG)

- Manage connections, add and remove layers
- Manage label data and evaluate performance

[>>help semanticseg]()

Algorithm Development, Testing & Verification

Applicable to other sensors, e.g. LiDar

Automated Driving System Toolbox™
Computer Vision System Toolbox™
Neural Network Toolbox™

Access Data → Preprocess and Label → Object Classification

Velodyne
FileReader

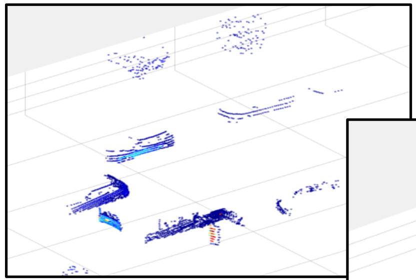

Preprocessing

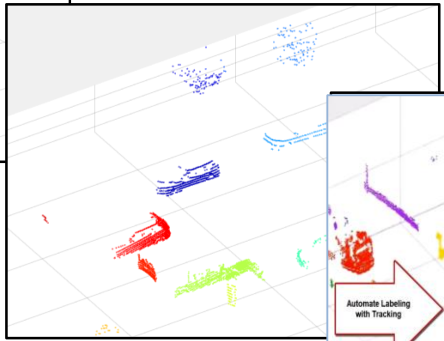

Clustering

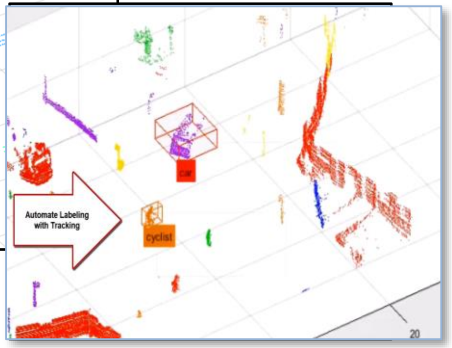

LiDAR Annotation
Ground Truth
[Under Request]

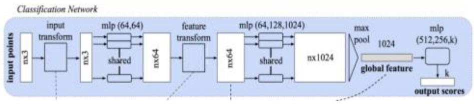

PointNet
Deep Neural Network

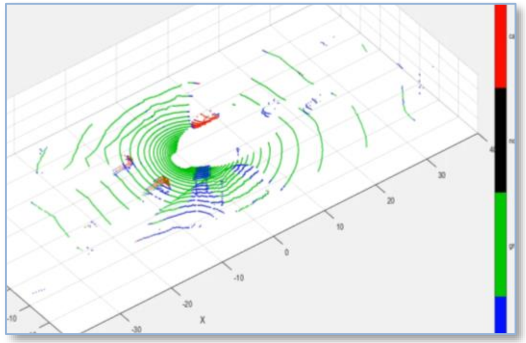

LSTM for both Classification and Regression

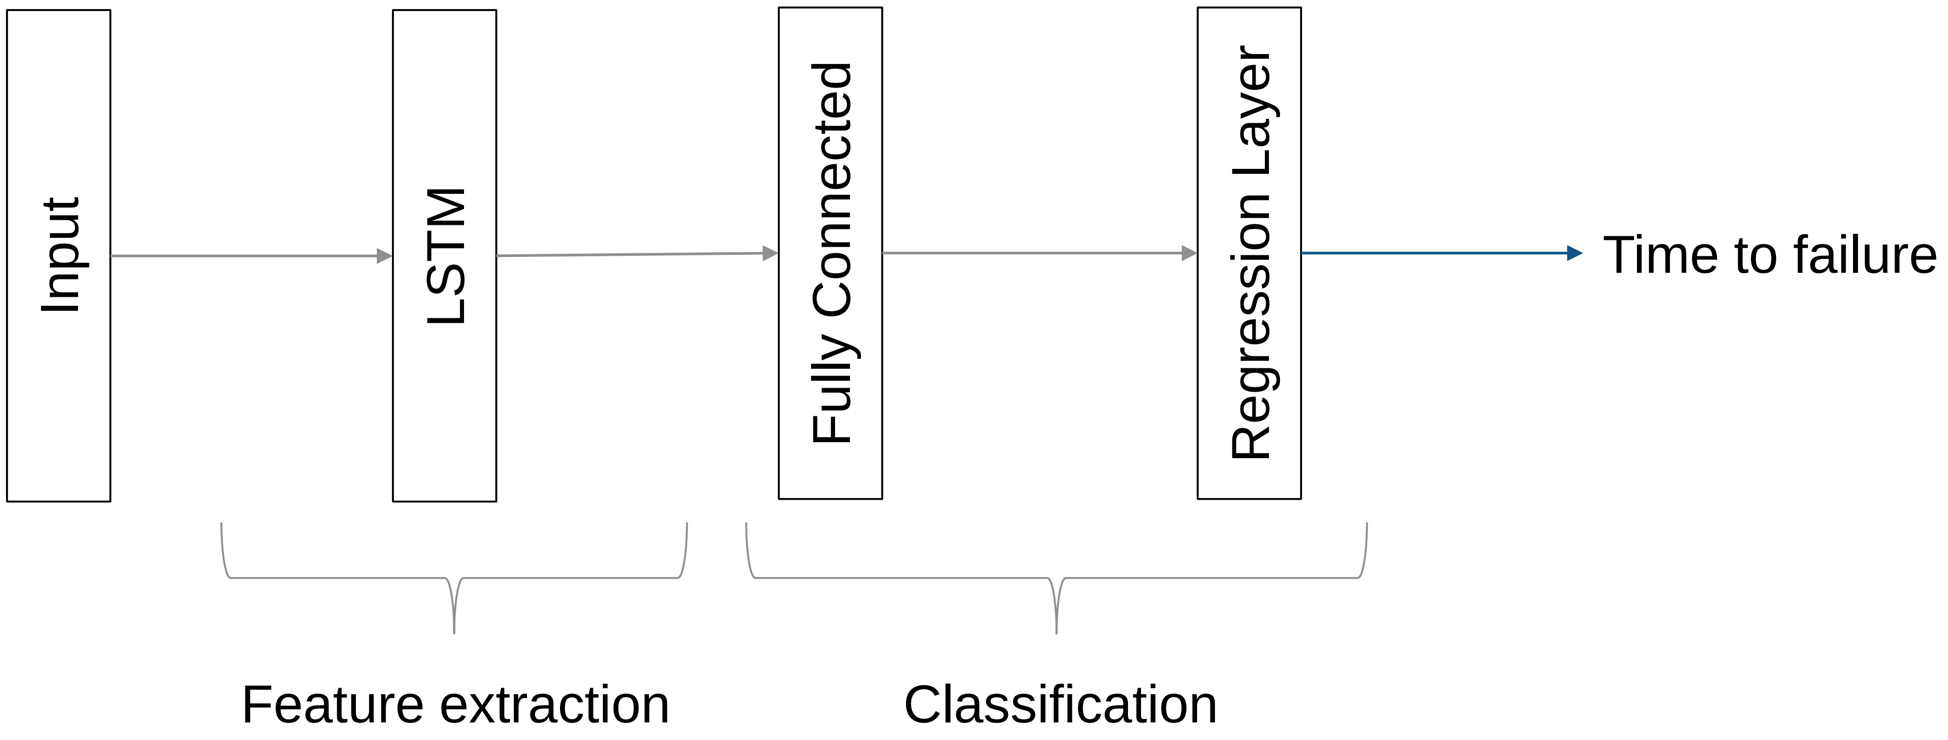

Define new operations for deep networks with 'Custom Layers'

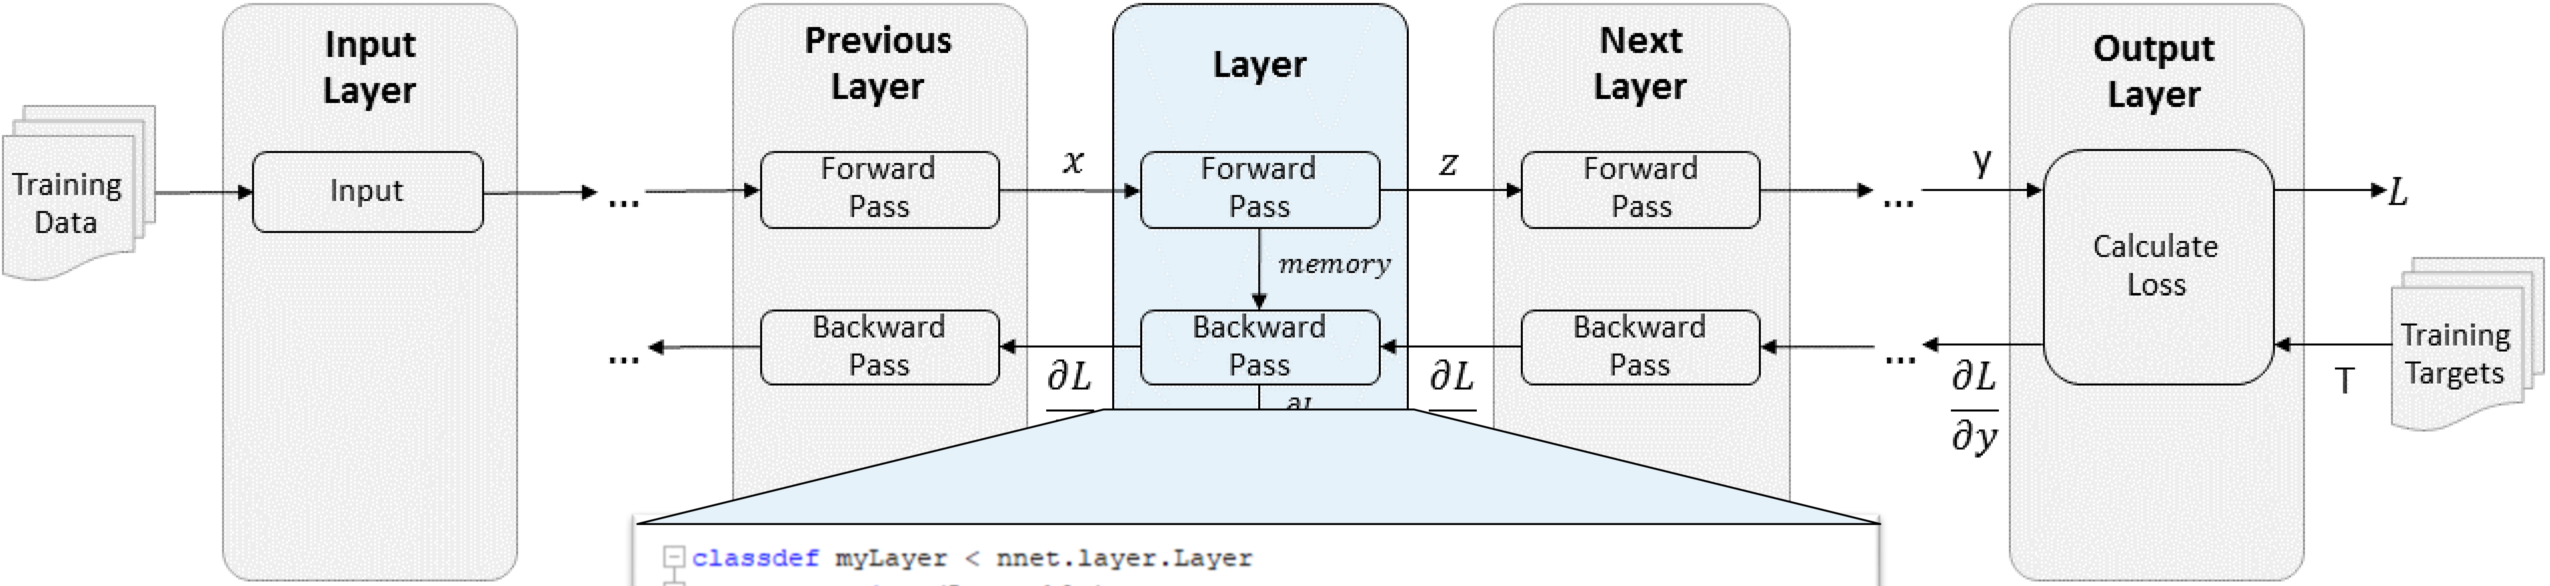


```
classdef myLayer < nnet.layer.Layer
    properties (Learnable)
        % Layer learnable parameters go here
    end
    methods
        Z = predict(layer, X)
        % Forward input data through the layer at prediction time and
        % output the result

        [dLdX, dLdW] = backward(layer, X, Z, dLdZ, memory)
        % Backward propagate the derivative of the loss function through
        % the layer
    end
end
```

THANK YOU!

Use Less Data and Less Time with Transfer Learning

Automatically generate code!

Image Pre-processing → CNN → Image Post-processing

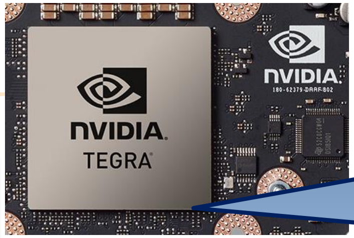

Prototype and productize your entire workflow on embedded GPUs

Apply Deep Learning not only on images but also on time series data