



# Deep Learning in MATLAB:


## A Brief Overview



**Brett Shoelson, PhD**

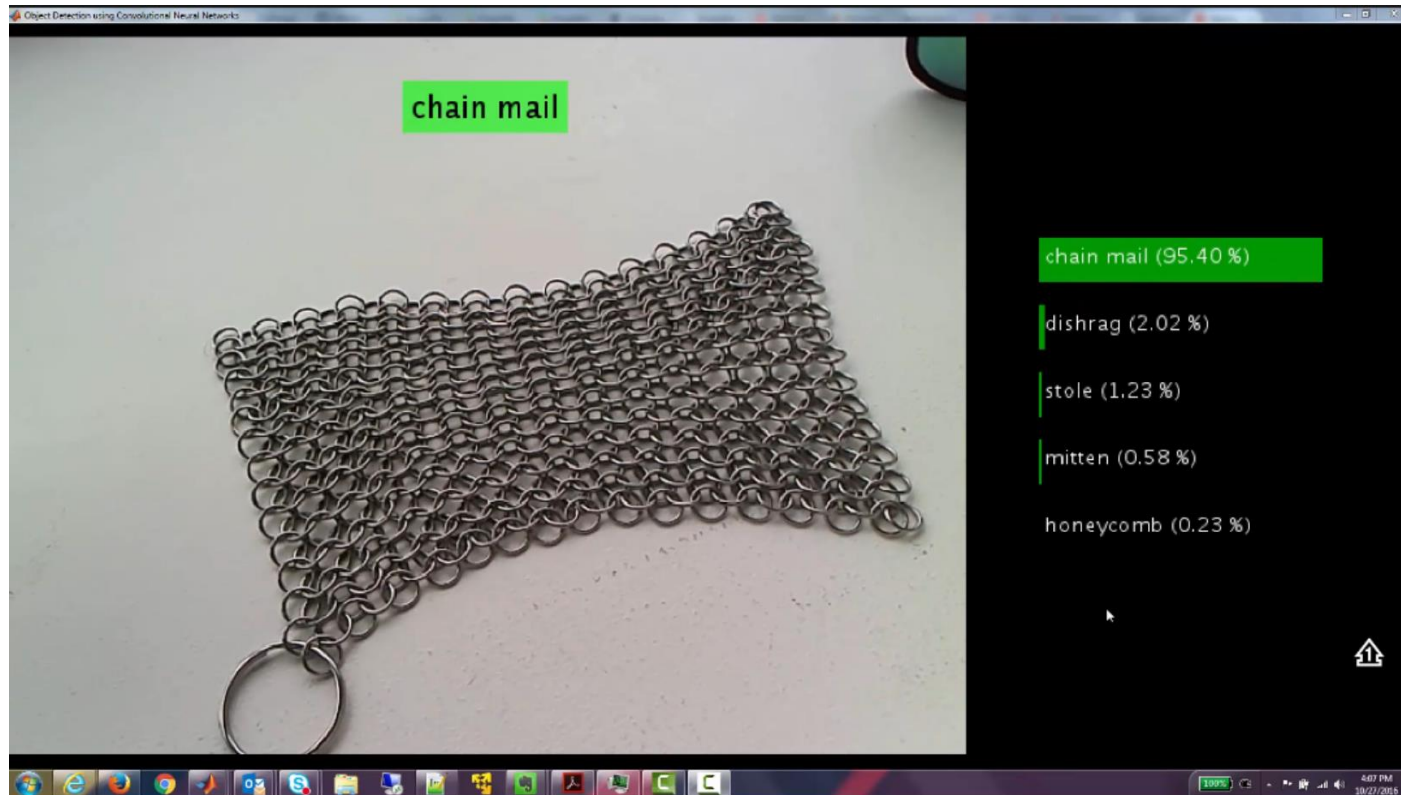
Principal Application Engineer

# What is can Deep Learning do for us? (An example)






# Object recognition using deep learning

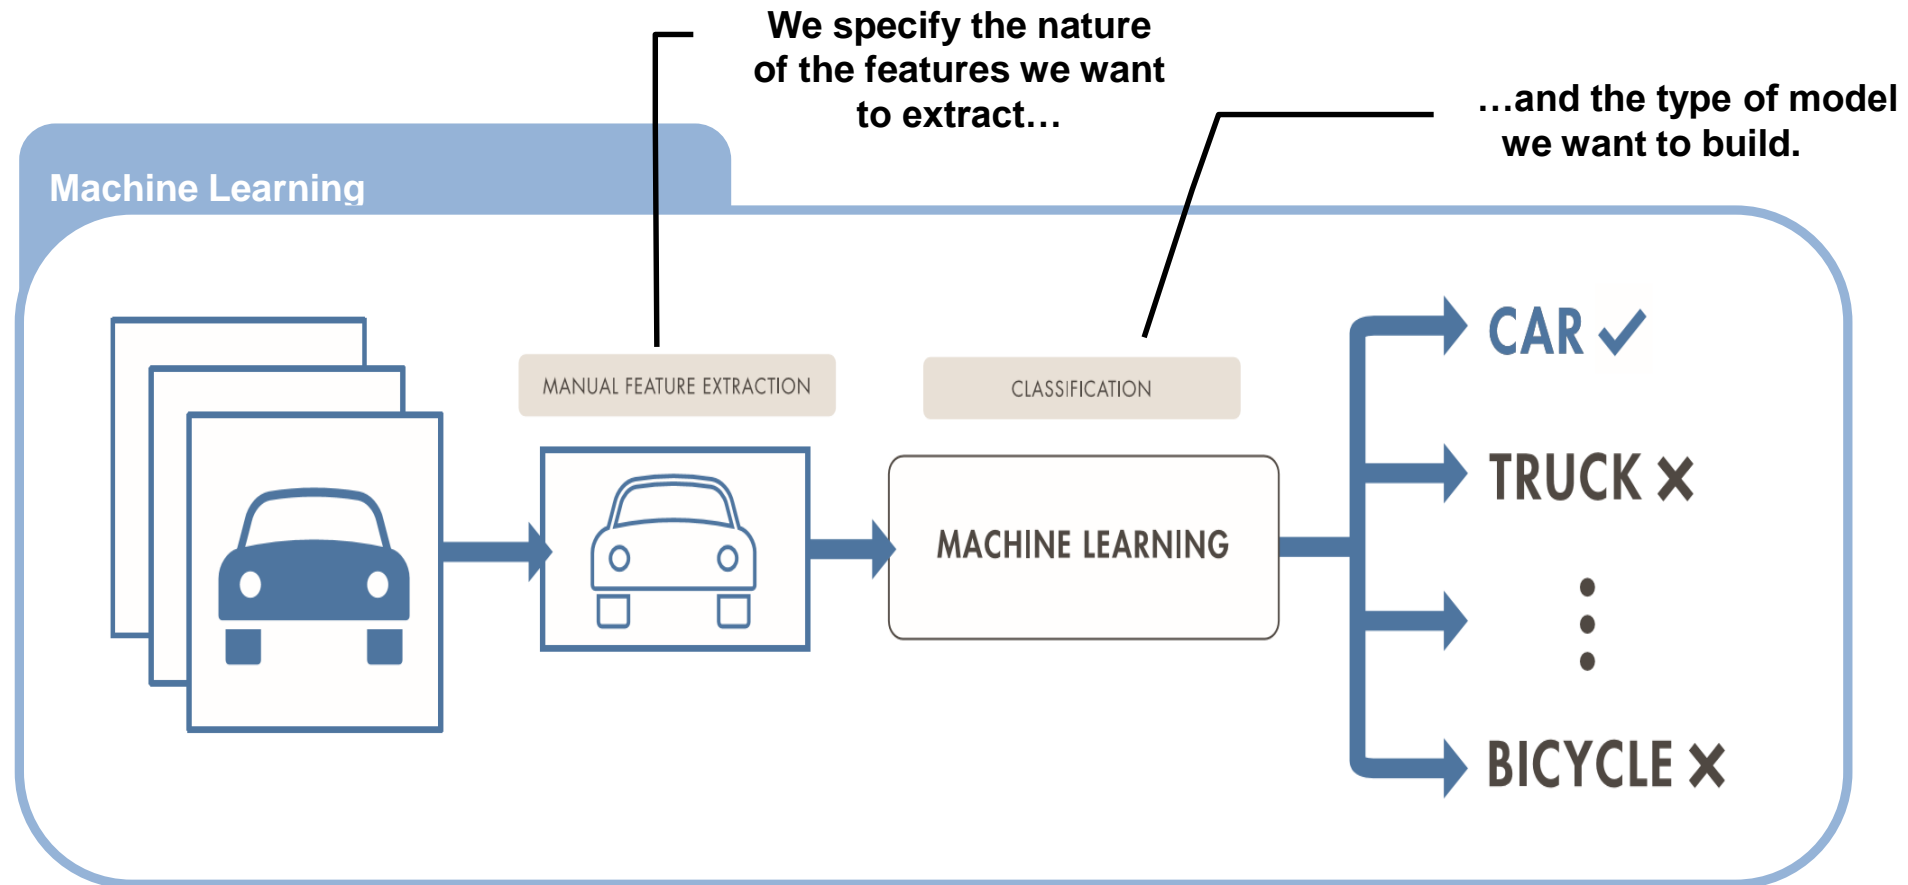


<b>Training (GPU)</b>	Millions of images from 1000 different categories
<b>Prediction</b>	Real-time object recognition using a webcam connected to a laptop

# What is Deep Learning?

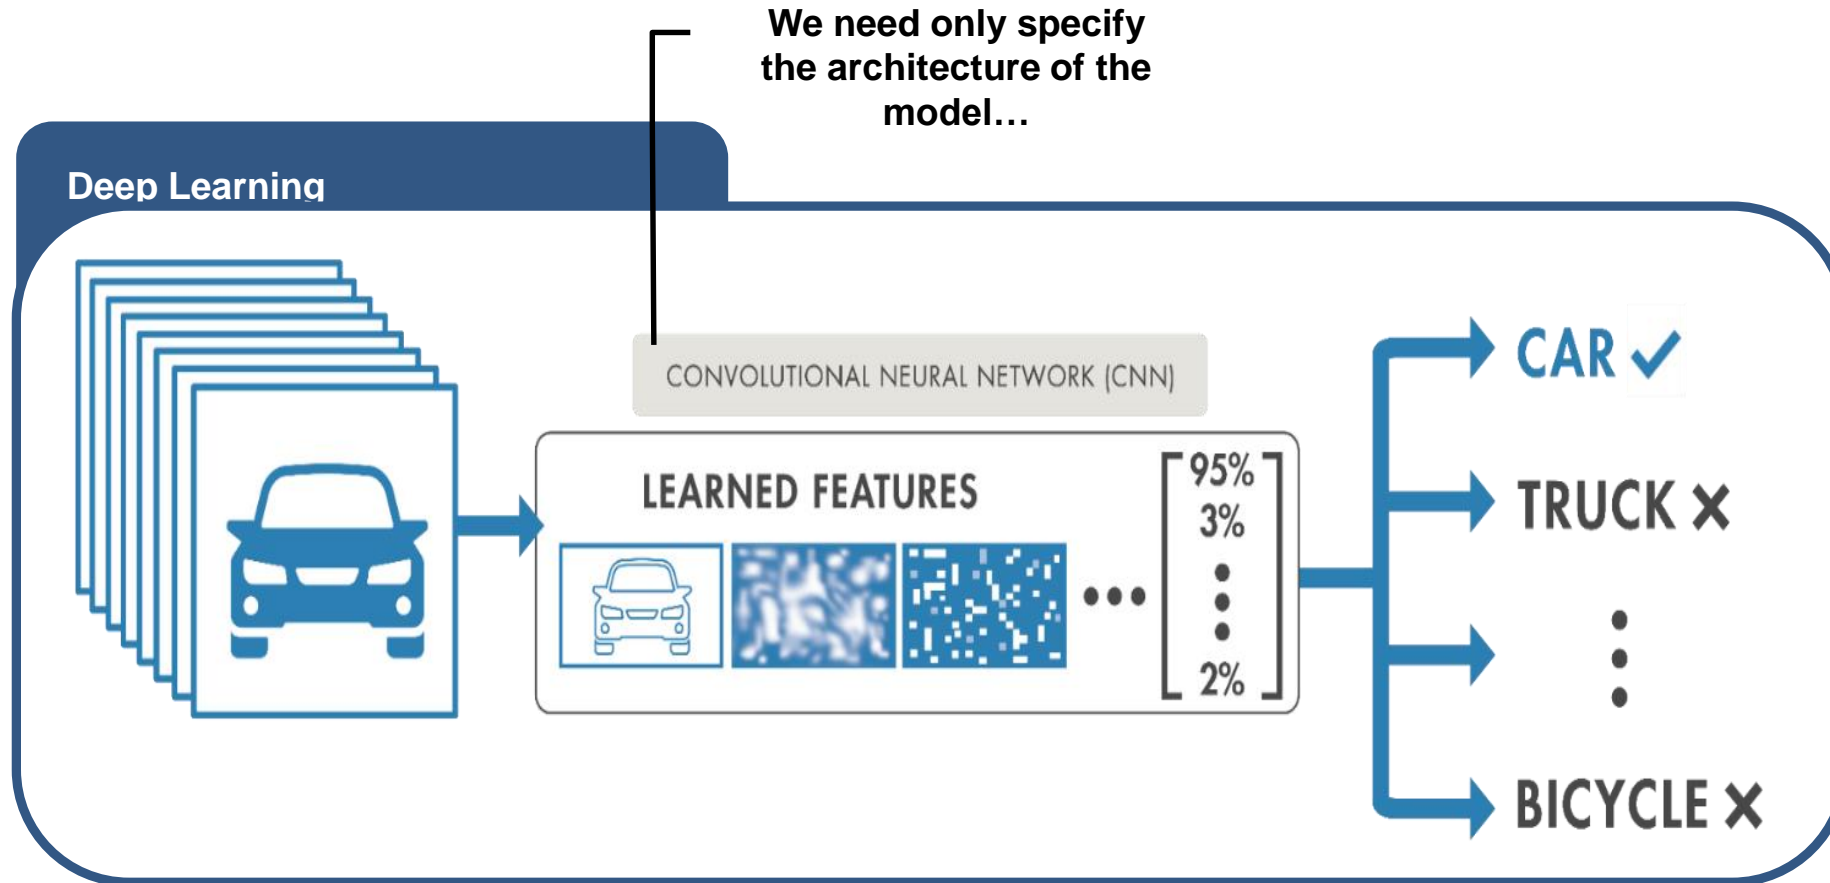


# Machine Learning vs Deep Learning






# Machine Learning vs Deep Learning




- **Deep learning** is a type of machine learning in which a model learns to perform tasks like classification – directly from images, texts, or signals.
- Deep learning performs **end-to-end learning**, and is usually implemented using a **neural network architecture**.
- Deep learning algorithms also **scale with data** – traditional machine learning **saturates**.





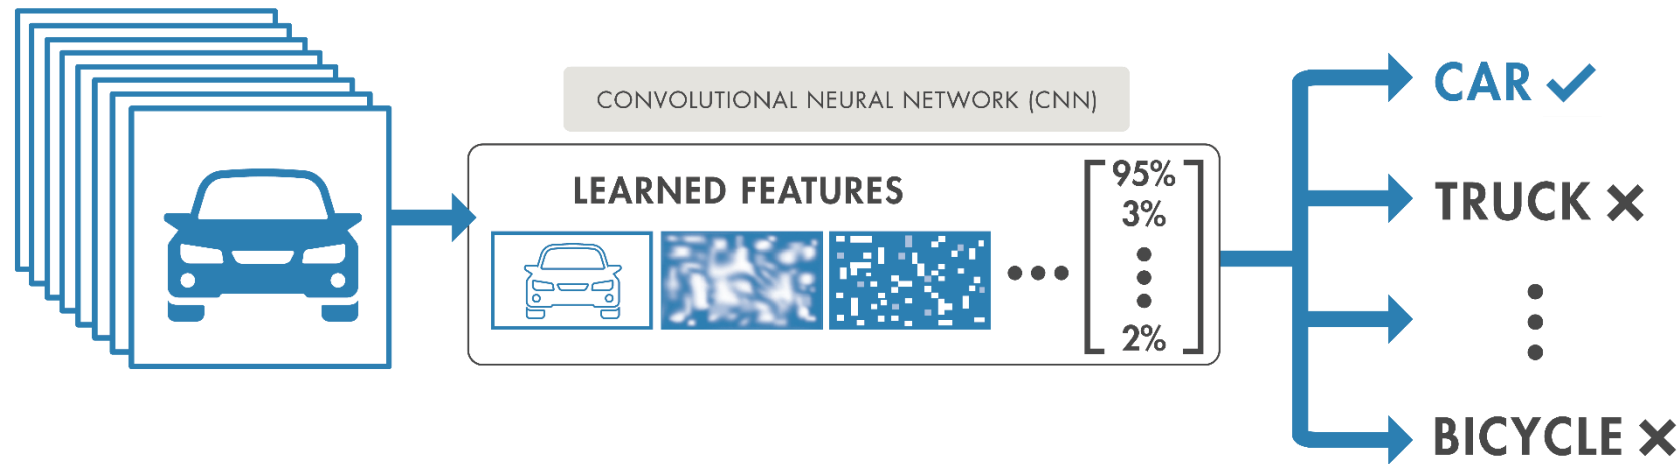
# Why is Deep Learning So Popular Now?



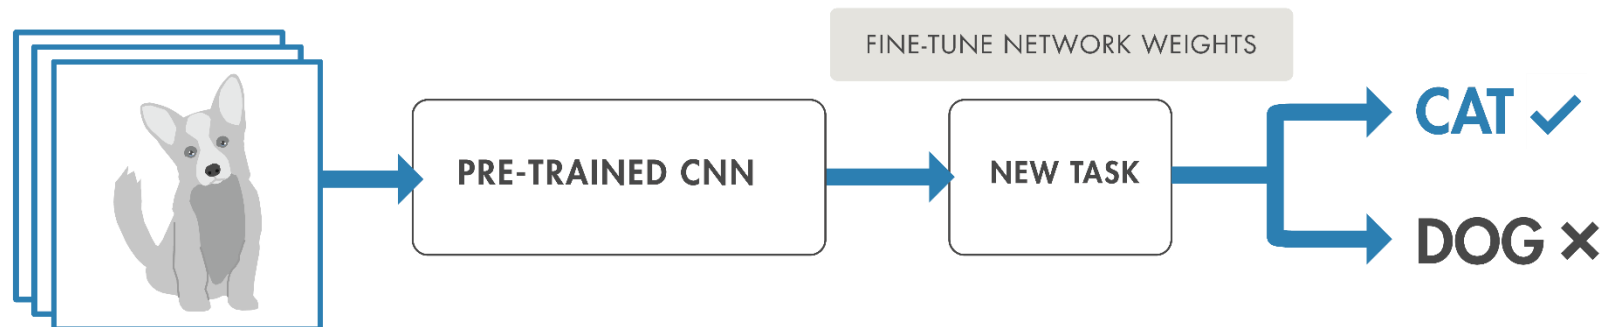
Source: ILSVRC Top-5 Error on ImageNet

# Two Approaches for Deep Learning

## 1. Train a Deep Neural Network from Scratch



## 2. Fine-tune a pre-trained model (transfer learning)



# Pains In Deep Learning

## Expertise

```
layers = [  
    imageInputLayer([28 28 1])  
  
    convolution2dLayer(3,16,'Padding',1)  
  
    batchNormalizationLayer  
    reluLayer  
  
    maxPooling2dLayer(2,'Stride',2)  
  
    convolution2dLayer(3,32,'Padding',1)  
  
    batchNormalizationLayer  
    reluLayer  
  
    maxPooling2dLayer(2,'Stride',2)  
  
    convolution2dLayer(3,64,'Padding',1)  
    batchNormalizationLayer  
    reluLayer  
  
    fullyConnectedLayer(10)  
    softmaxLayer  
    classificationLayer];
```

## Time to Train

Training iteration 278 of 14040...

### Training Time

Start time: 03-Apr-2018 16:15:30  
Elapsed time: 17 min 51 sec

## Data



# Example: Vehicle recognition using deep transfer learning



Cars



Trucks



SUVs



Big Trucks



Vans



5 Category  
Classifier

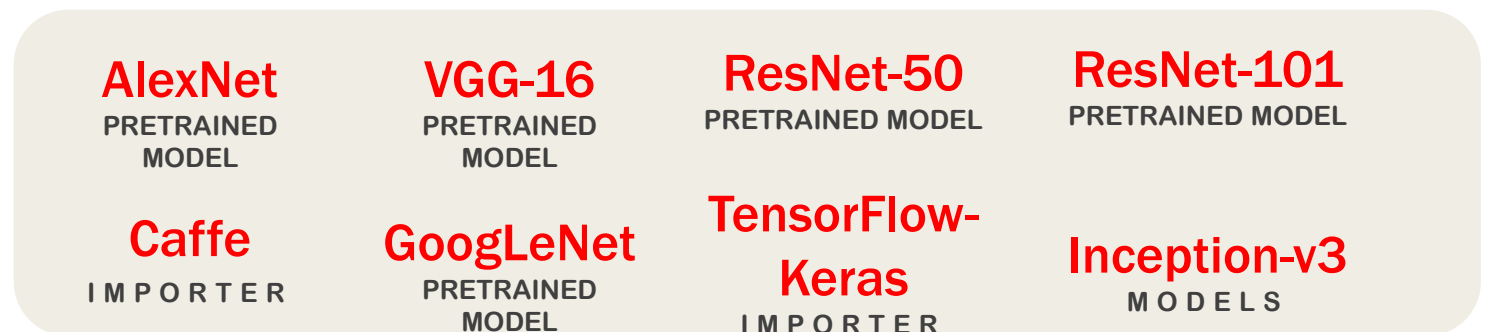
# Import the Latest Models for Transfer Learning

## Pretrained Models\*

- AlexNet
- VGG-16
- VGG-19
- GoogLeNet
- Inception-v3
- ResNet50
- ResNet-101
- Inception-resnet-v2
- SqueezeNet
- MobileNet(coming soon)


## Import Models from Frameworks

- Caffe Model Importer
- TensorFlow-Keras Model Importer
- Onnx - Importer/ Exporter (Coming Soon)



\* single line of code to access model


# Detection and localization using deep learning



**Regions with Convolutional Neural Network Features (R-CNN)**



# What is semantic segmentation?




# Localization using deep learning

Original Image




ROI detection




Pixel classification



# Semantic Segmentation Network



# Semantic Segmentation Network





# Semantic Segmentation Demo



## CamVid Dataset

1. Segmentation and Recognition Using Structure from Motion Point Clouds, ECCV 2008
2. Semantic Object Classes in Video: A High-Definition Ground Truth Database ,Pattern Recognition Letters

# Semantic Segmentation



## CamVid Dataset

1. Segmentation and Recognition Using Structure from Motion Point Clouds, ECCV 2008
2. Semantic Object Classes in Video: A High-Definition Ground Truth Database ,Pattern Recognition Letters



“I love to label and  
preprocess my data”

*~ Said no engineer, ever.*


# Ground truth Labeling

“How do I *label* my data?”

**New App for Ground Truth Labeling**

Label pixels and regions for semantic segmentation

**Data**



# Attributes and Sublabels

**NEW in**  
**R2018a**

The screenshot illustrates the Ground Truth Labeler interface, which is used for defining and labeling objects in video frames. The interface is divided into several panels:

- ROI Label Definition:** This panel allows users to define new ROI labels. It includes tabs for Label, Sublabel, and Attribute. The 'cyclist' label is selected, and its sublabels 'bicycle' and 'vehicle' are visible. The 'bicycle' sublabel is highlighted with a green arrow.
- Scene Label Definition:** This panel allows users to define new scene labels. It includes a 'Define new scene label' button and options for 'Current Frame' and 'Time Interval'.
- Video Labeling:** The central panel shows a video frame with a cyclist and a car. The cyclist is labeled 'bicycle' and the car is labeled 'vehicle'. A blue arrow points from the 'bicycle' sublabel in the ROI Label Definition panel to the 'bicycle' label in the video frame.
- Attributes and Sublabels:** This panel shows the attributes for the selected label. For the 'bicycle' label, the attributes are 'bicycle' (bikeType) and 'inMotion' (action). A blue arrow points from the 'Attributes and Sublabels' panel to the 'bicycle' label in the video frame.
- Attributes and Sublabels (Inset):** This inset panel shows the attributes for the selected label. It includes a dropdown for 'bikeType' (set to 'bicycle') and a dropdown for 'action' (set to 'inMotion').

At the bottom of the interface, there is a timeline with markers for Start Time, Current, End Time, and Max Time. The video frame shows a cyclist on a road with cars in the background.

# 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	

ML or LSTM

## Time Series/ Text Data



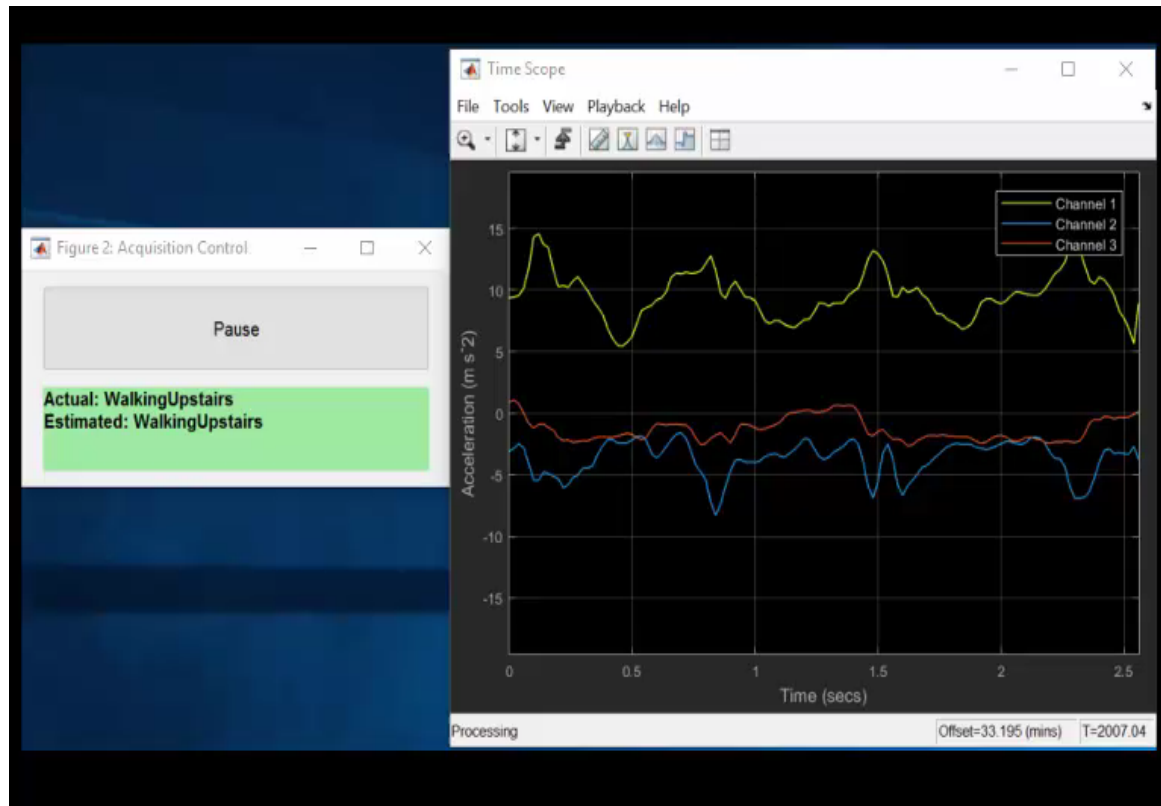
LSTM or CNN

## Image Data

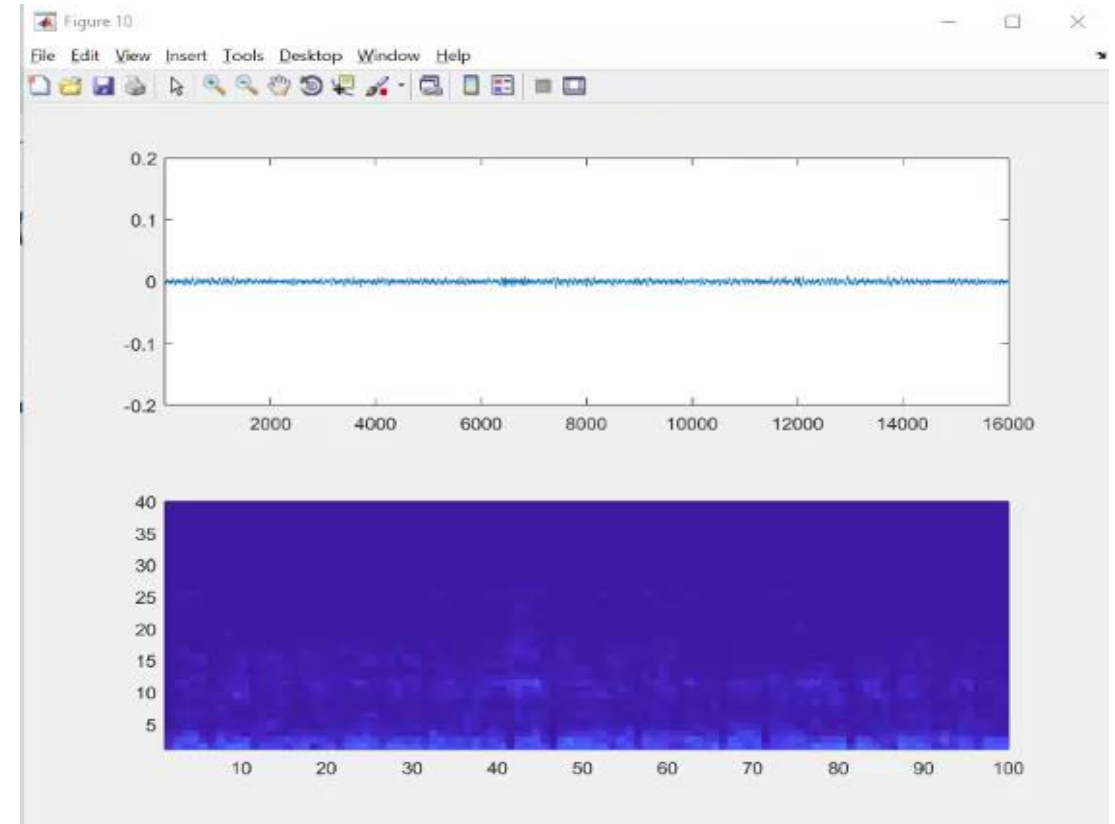


CNN

# Analyzing signal data using deep learning



**Signal Classification using LSTMs**




**Speech Recognition using CNNs**

# Deep learning features overview

- Classification
- Regression
- Semantic segmentation
- Object detection
- Scalability
  - Multiple GPUs
  - Cluster or cloud
- Custom network layers
- Import models
  - Caffe
  - Keras/TensorFlow
- Data augmentation
- Hyperparameter tuning
  - Bayesian optimization
- Python ↔ MATLAB interface
- LSTM networks
  - Time series, signals, audio
- Custom labeling
  - API for ground-truth labeling automation
  - Superpixels
- Data validation
  - Training and testing



# Prediction Performance: Fast with GPU Coder



Why is GPU Coder so fast?

- Analyzes and optimizes network architecture
- Invested 15 years in code generation

**TensorFlow**

**MATLAB**

**MXNet**

**GPU Coder**

Using CUDA v9  
and cuDNN v7


# Overview of deep learning deployment options

“How do I *deploy*  
my model?”


Introducing:

**GPU Coder-  
Convert to  
NVIDIA CUDA  
code**


**Deploy / Share**




- Create Desktop Apps



- Run Enterprise Solution



- Generate C and C++ Code

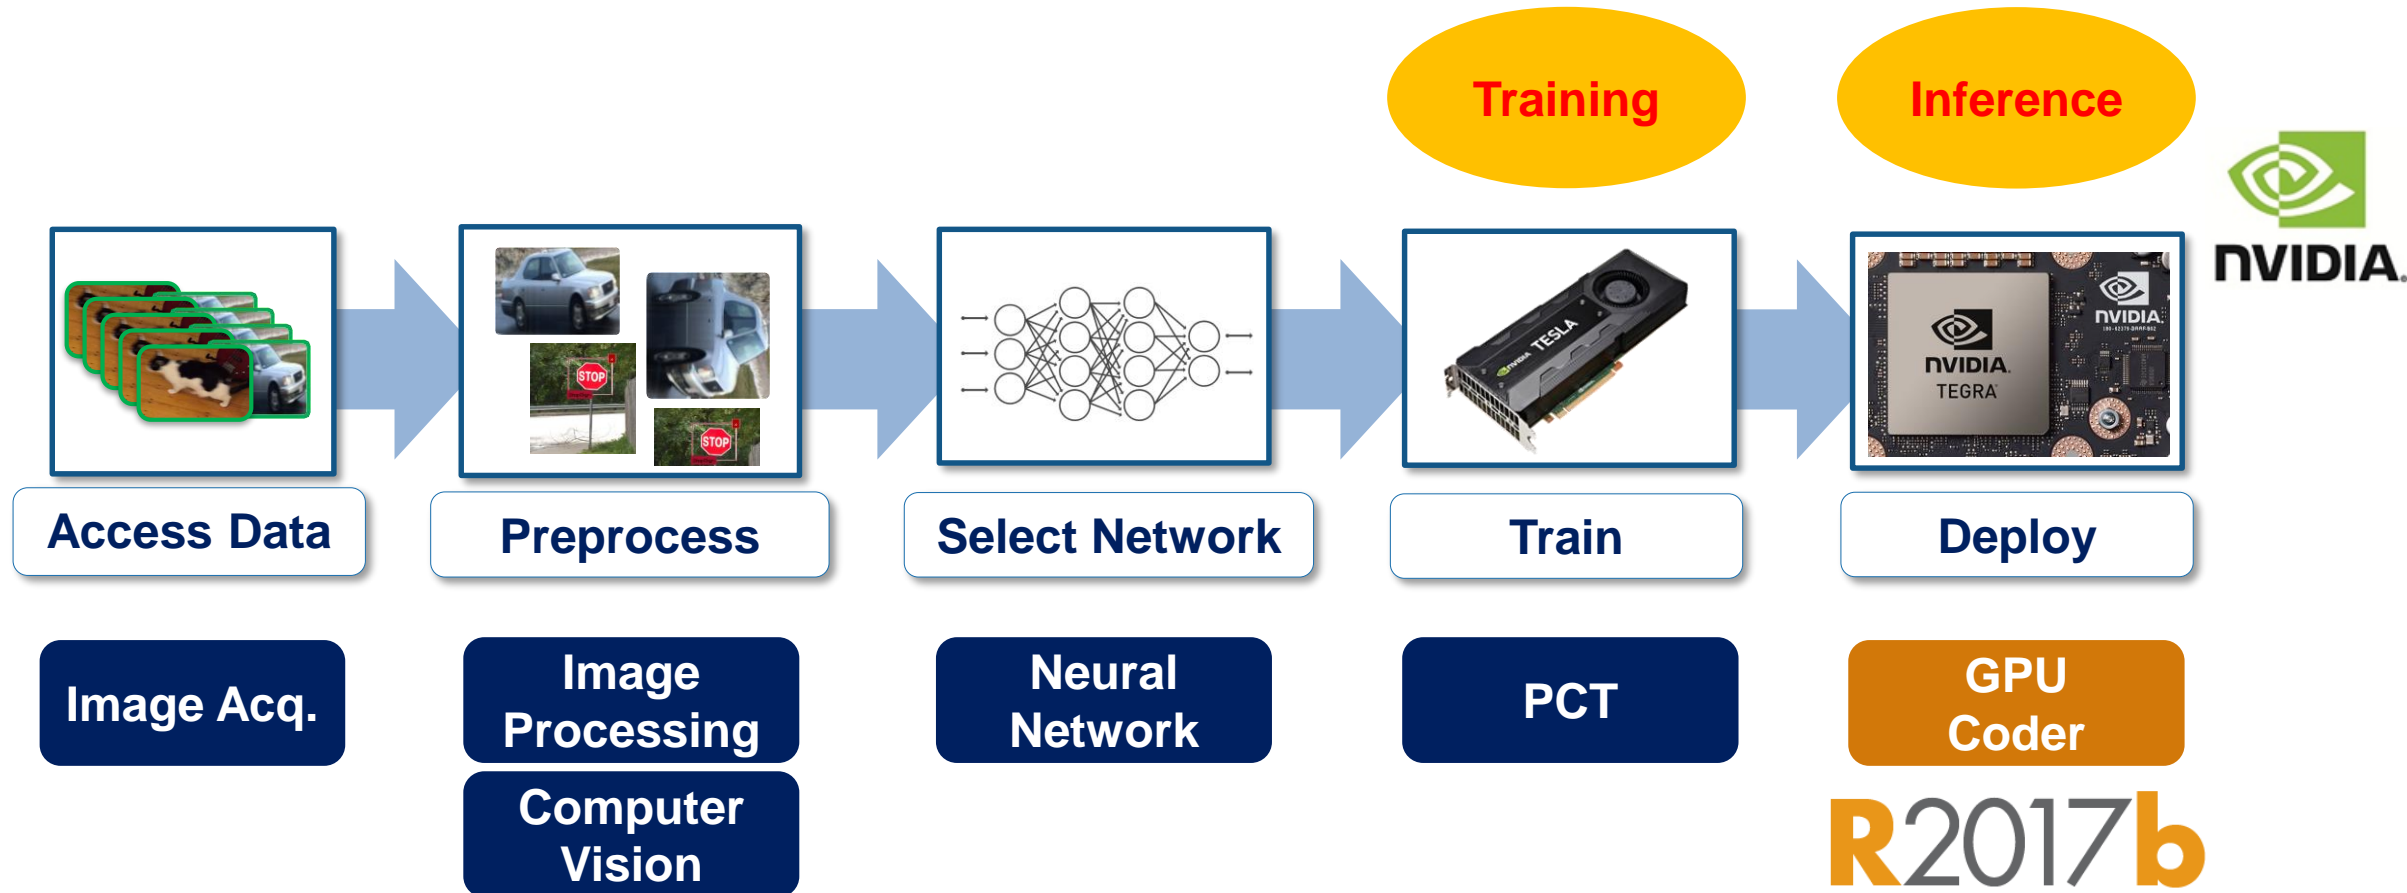


- Target GPUs
- Generate C and C++ Code

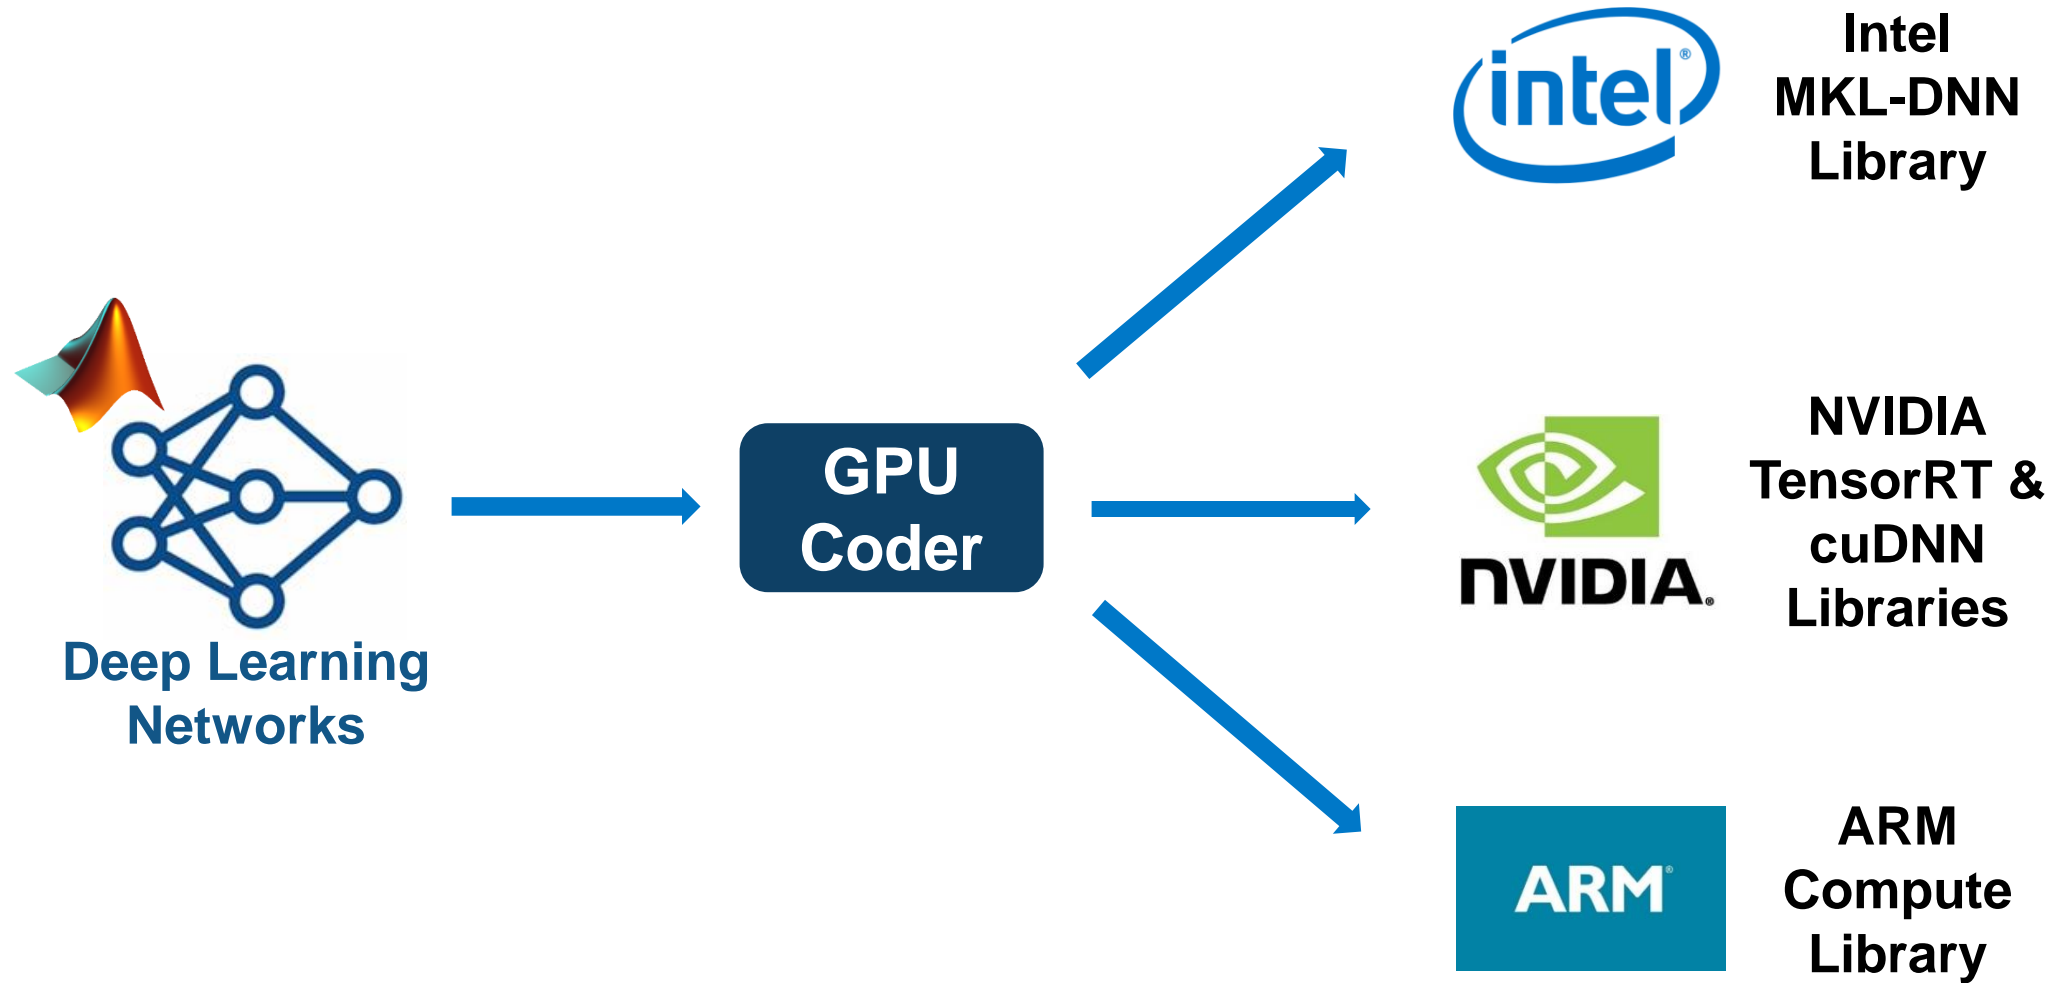
GPU Coder



# GPU Coder Fills a Gap in Our Deep Learning Solution



# Deploying to CPUs



# MATLAB products for deep learning

## Required products

- Neural Network Toolbox
- Parallel Computing Toolbox
- Image Processing Toolbox
- Computer Vision System Toolbox

## Recommended products

- Statistics and Machine Learning Toolbox
- MATLAB Coder
- GPU Coder **R2017b**
- Automated Driving System Toolbox **R2017a**

# Deep learning features overview


# R2017b

- Classification
- Regression \*
- Semantic segmentation
- Object detection \*
- Scalability \*
  - Multiple GPUs
  - Cluster or cloud
- Custom network layers \*
- Import models \*
  - Caffe
  - Keras/TensorFlow
- Data augmentation \*
- Hyperparameter tuning \*
  - Bayesian optimization
- Python ↔ MATLAB interface \*
- LSTM networks \*
  - Time series, signals, audio
- Custom labeling \*
  - API for ground-truth labeling automation
  - Superpixels
- Data validation \*
  - Training and testing





Thank you!






## Deep learning in automated driving...





# Deep Learning Onramp

- Get started using deep learning methods to perform image recognition.
- Free access for everyone
- Interactive exercises and short video demonstrations
- Work on real-life image recognition problems
- Topics include:
  - Convolutional neural networks
  - Working with pre-trained networks
  - Transfer learning
  - Evaluating network performance



# Convolutional Neural Networks (CNN)

