

# MATLAB EXPO 2018

## Demystifying Deep Learning: Image/Video Focus

*“Let the computers do the hard work”*

**Abhijit Bhattacharjee**

*Senior Application Engineer*

*MathWorks*



# Deep Learning Demo

# Image Classification

# Why MATLAB for Deep Learning?


- MATLAB is Productive
- MATLAB is Fast
- MATLAB Integrates with Open Source





# Deep Learning Applications

Voice assistants (speech to text)

Teaching character to beat video game

Automatically coloring black-and-white images





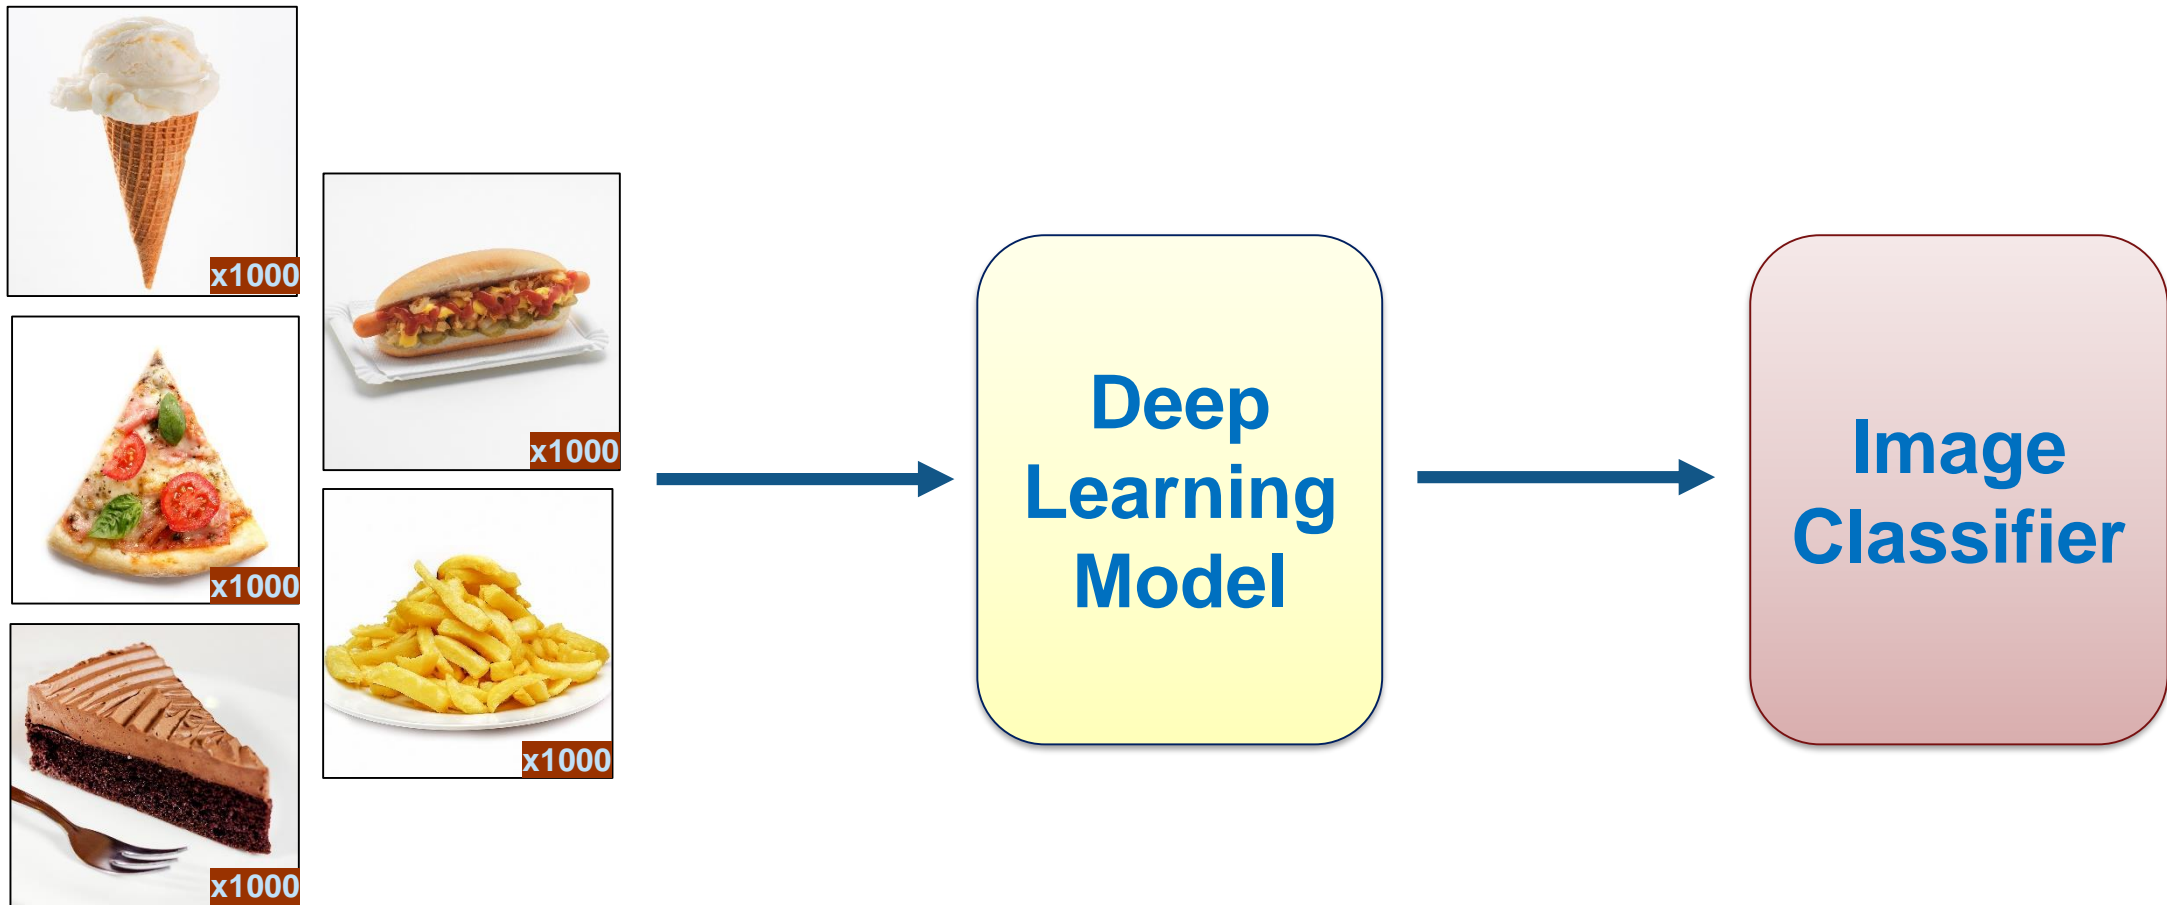
# What is Deep Learning?



12 40.0%	0 0.0%	100% 0.0%
0 0.0%	18 60.0%	100% 0.0%
100% 0.0%	100% 0.0%	100% 0.0%

# Deep Learning

Model learns to perform classification tasks directly from data.






# Deep Learning is **Versatile**



*Detection of cars and road in autonomous driving systems*



*Rain Detection and Removal<sup>1</sup>*




*Iris Recognition – 99.4% accuracy<sup>2</sup>*

1. "Deep Joint Rain Detection and Removal from a Single Image" Wenhan Yang, Robby T. Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan
2. Source: An experimental study of deep convolutional features for iris recognition Signal Processing in Medicine and Biology Symposium (SPMB), 2016 IEEE Shervin Minaee ; Amirali Abdolrashidiy ; Yao Wang; An experimental study of deep convolutional features for iris recognition


# How is deep learning performing so well?

# Deep Learning Uses a Neural Network Architecture




# Thinking about Layers

- Layers are like blocks
  - Stack on top of each other
  - Replace one block with a different one
- Each hidden layer processes the information from the previous layer

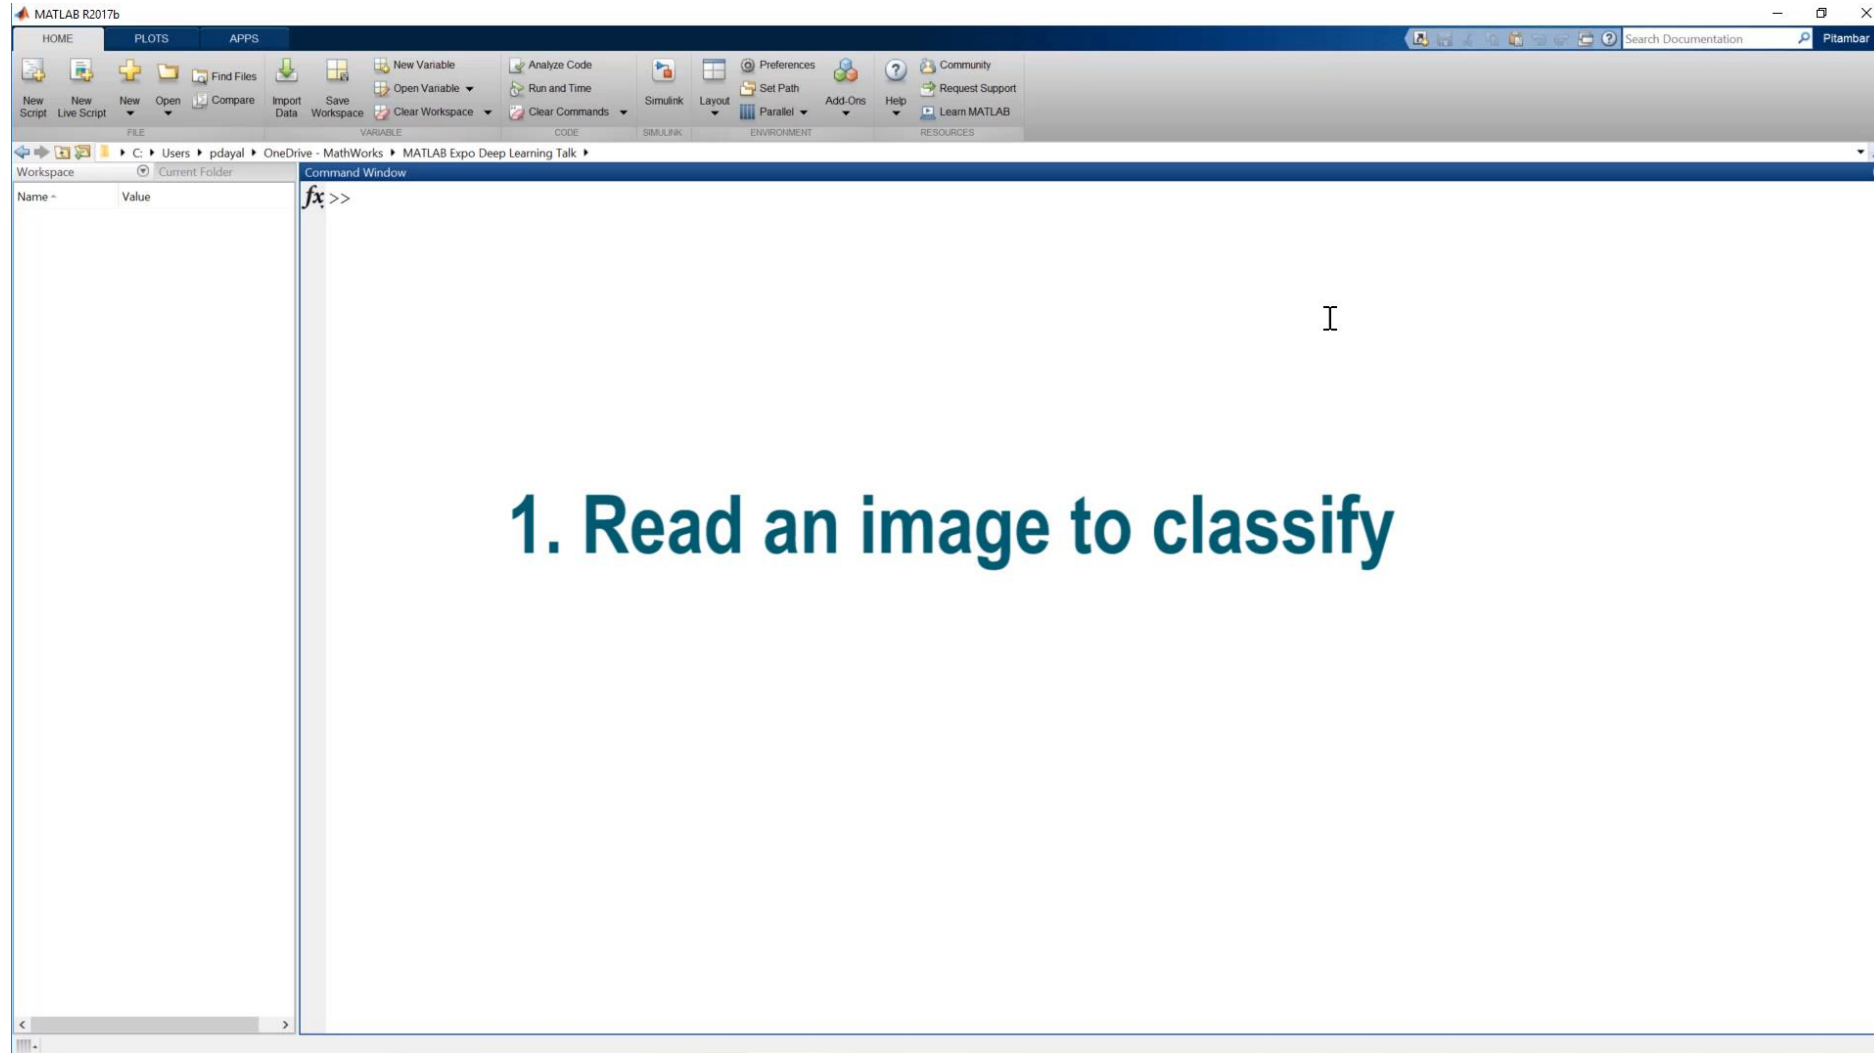


# Thinking about Layers

- Layers are like blocks
  - Stack them on top of each other
  - Replace one block with a different one
- Each hidden layer processes the information from the previous layer
- Layers can be ordered in different ways



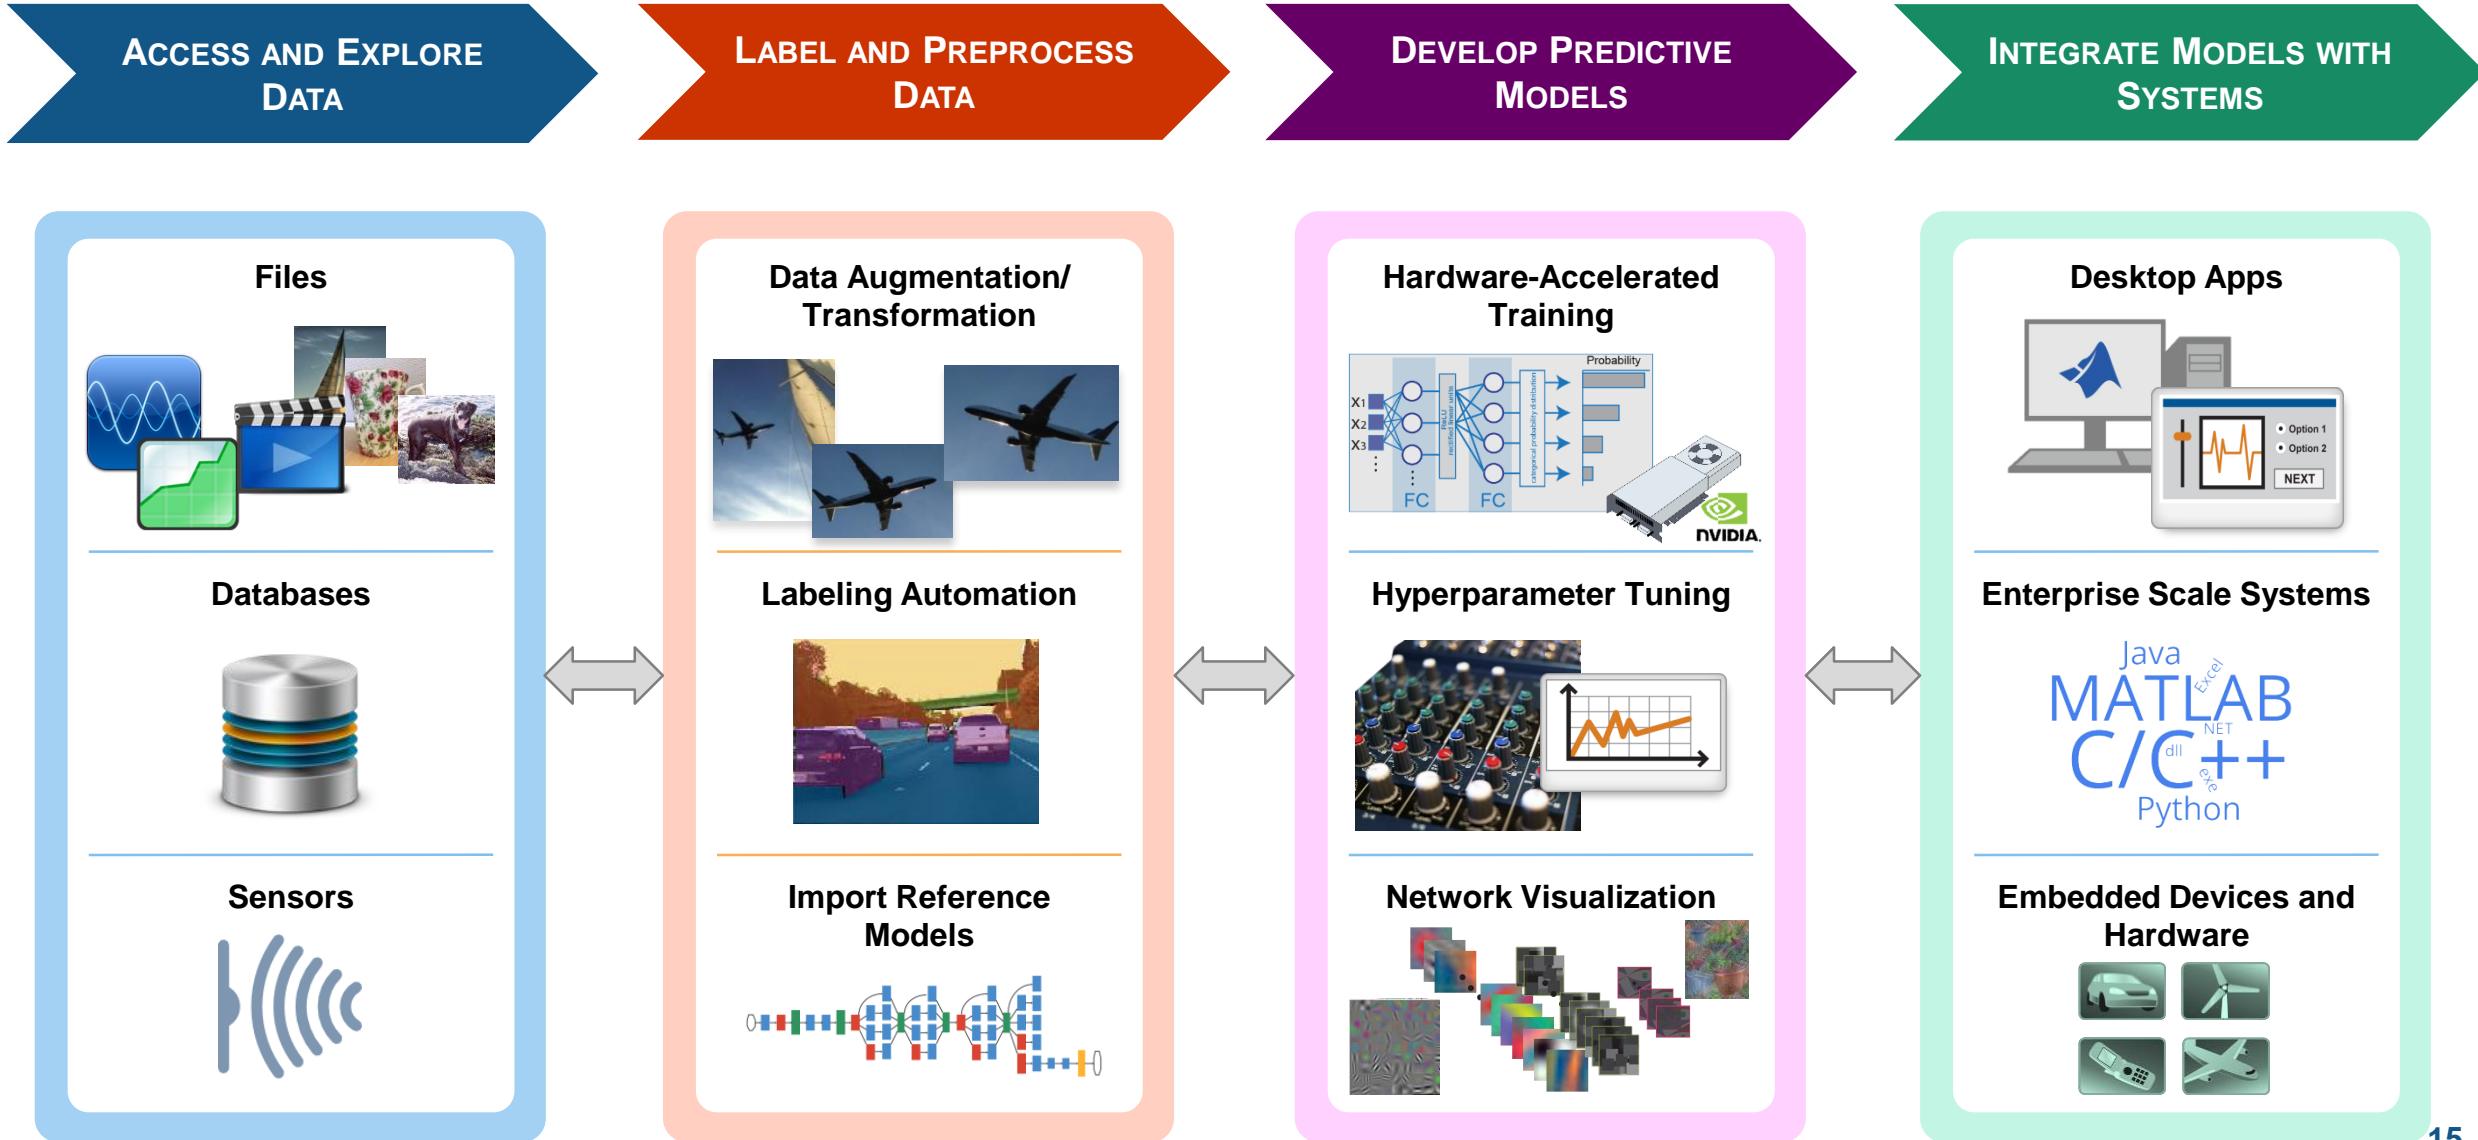
# Deep Learning in 6 Lines of MATLAB Code



# Why MATLAB for Deep Learning?

- **MATLAB is Productive**
- **MATLAB is Fast**
- **MATLAB integrates with Open Source**

# Deep Learning Workflow



**“I love to label and  
preprocess my data”**

*~ Said no engineer, ever.*


# Caterpillar Case Study



- World's leading manufacturer of construction and mining equipment.
- Similarity between these projects?
  - Autonomous haul trucks
  - Pedestrian detection
  - Equipment classification
  - Terrain mapping

# Computer Must Learn from Lots of Data

- ALL data must first be labeled to create these autonomous systems.

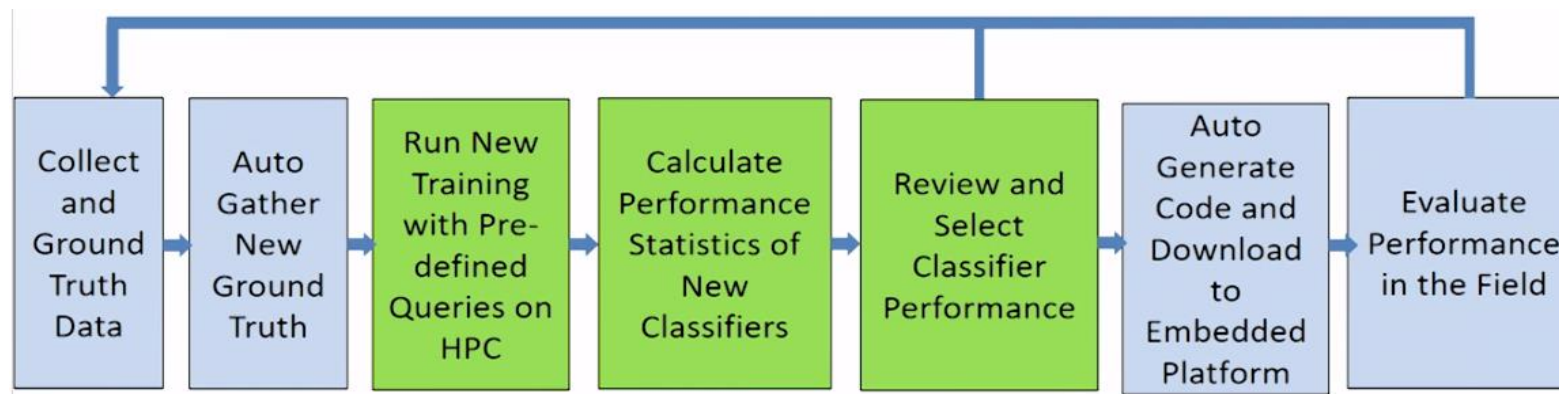


“We were spending way too much time ground-truthing [the data]”


--Larry Mianzo, Caterpillar


# How Did Caterpillar Do with Our Tools?

- Semi-automated labeling process
  - *“We go from having to label 100 percent of our data to only having to label about 80 to 90 percent”*
- Used MATLAB for entire development workflow.
  - *“Because everything is in MATLAB, development time is short”*

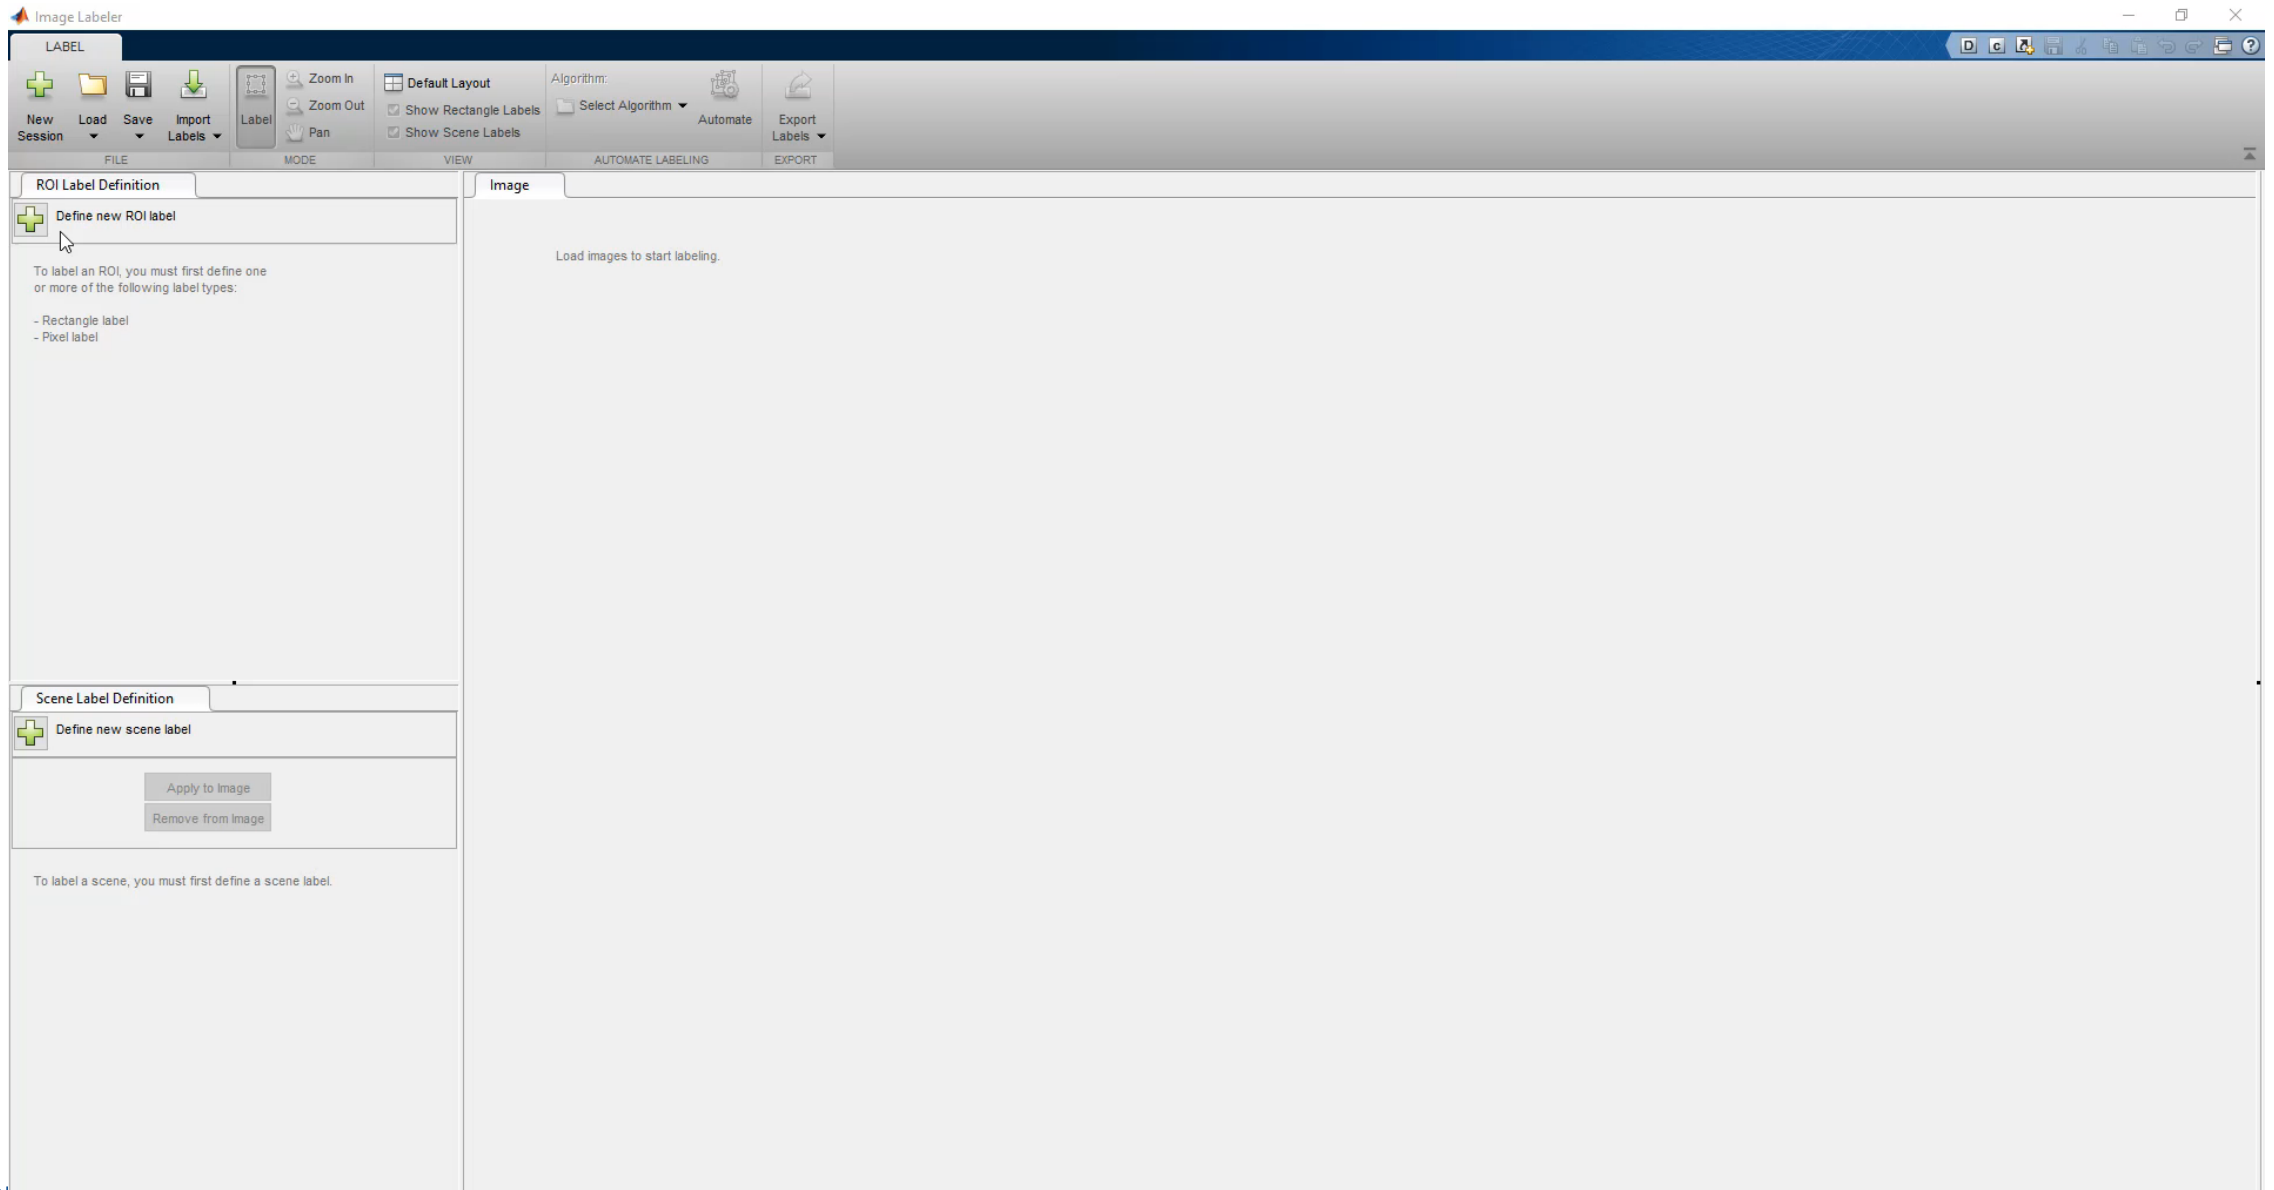


# How Does MATLAB Come into Play?





# Using image labeling tools

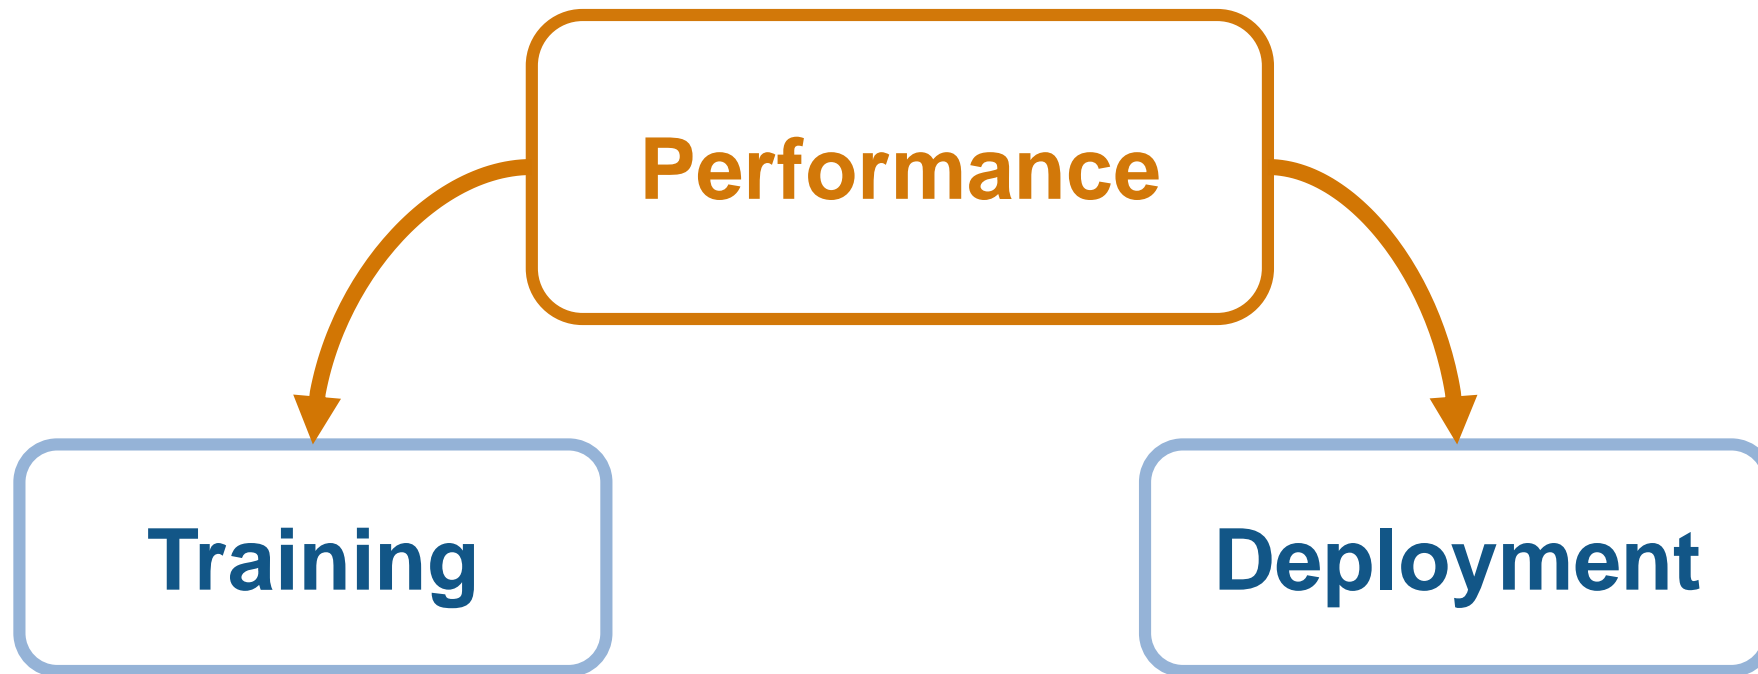




# MATLAB is Productive

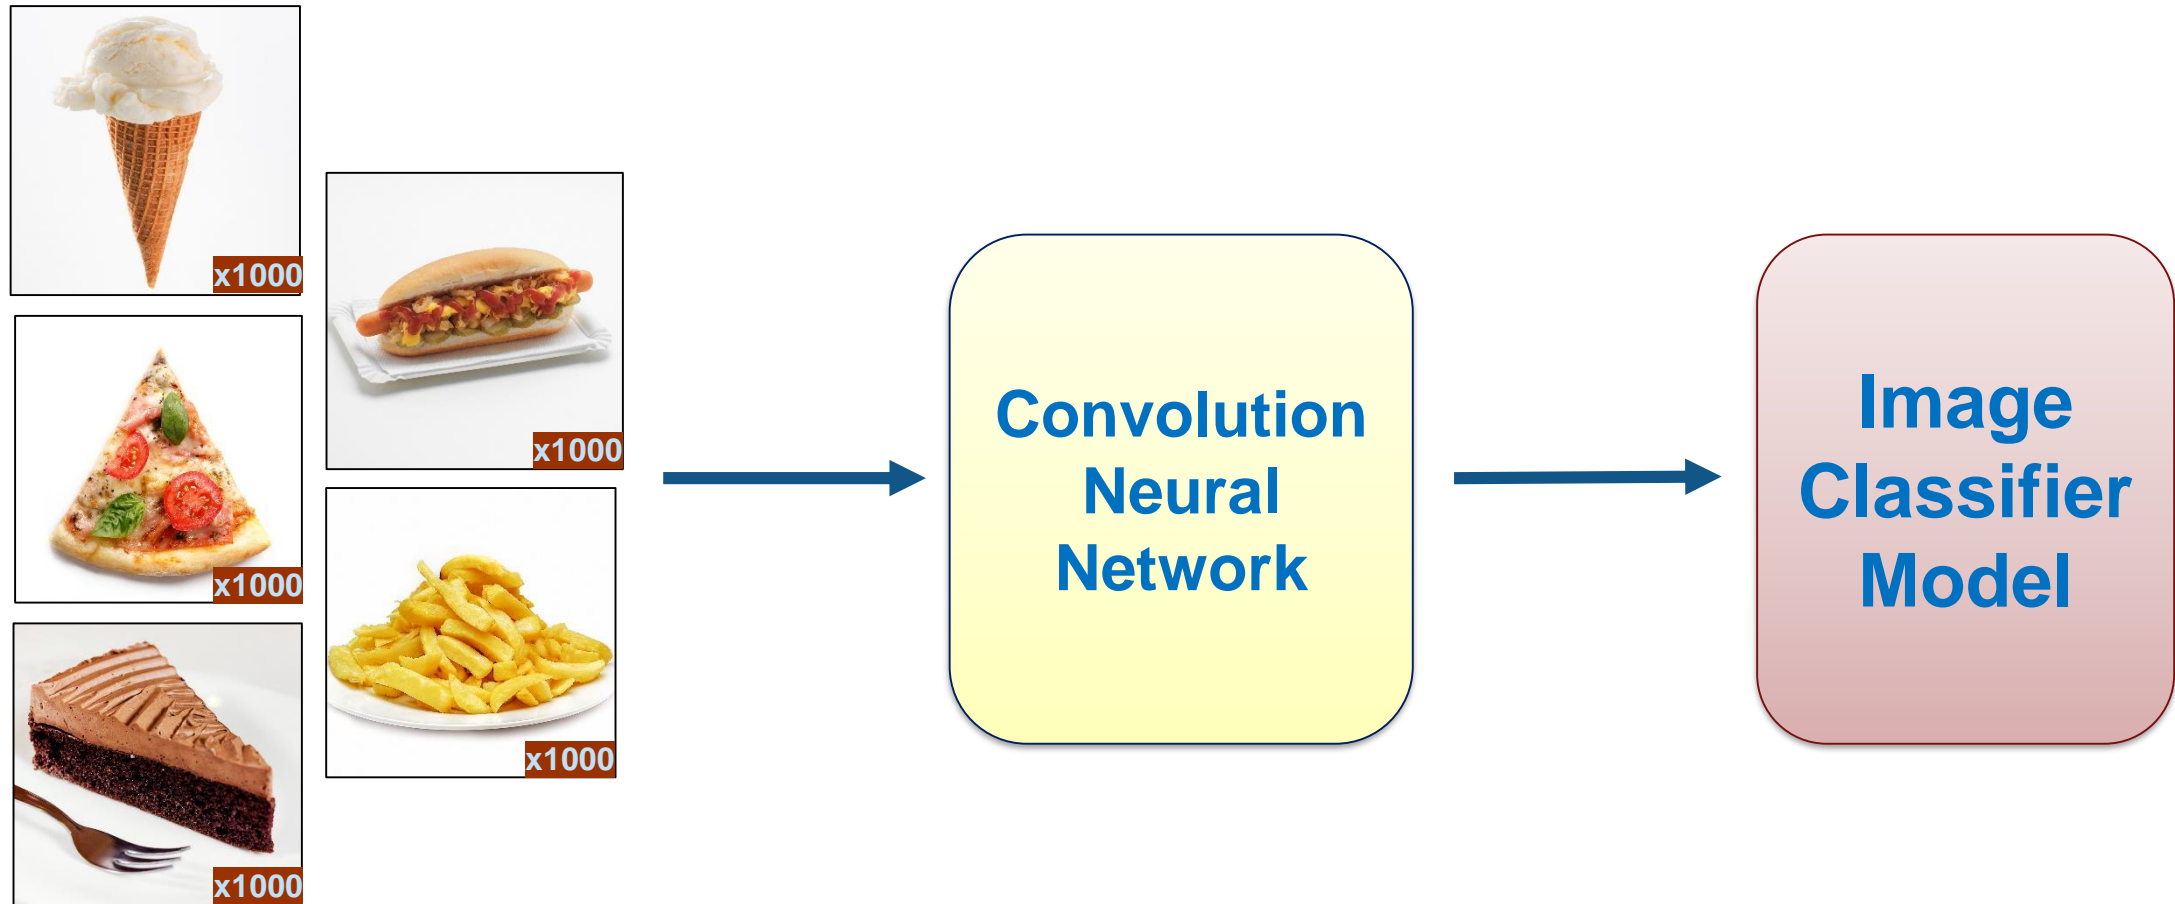
- Image Labeler App semi-automates labeling workflow
- Bootstrapping
  - Improve automatic labeling by updating algorithm as you label more images correctly.
- Easy to load metadata even when labeling manually

# MATLAB is Fast




# What is Training?


*Feed labeled data into neural network to create working model*




# Image Classification Network



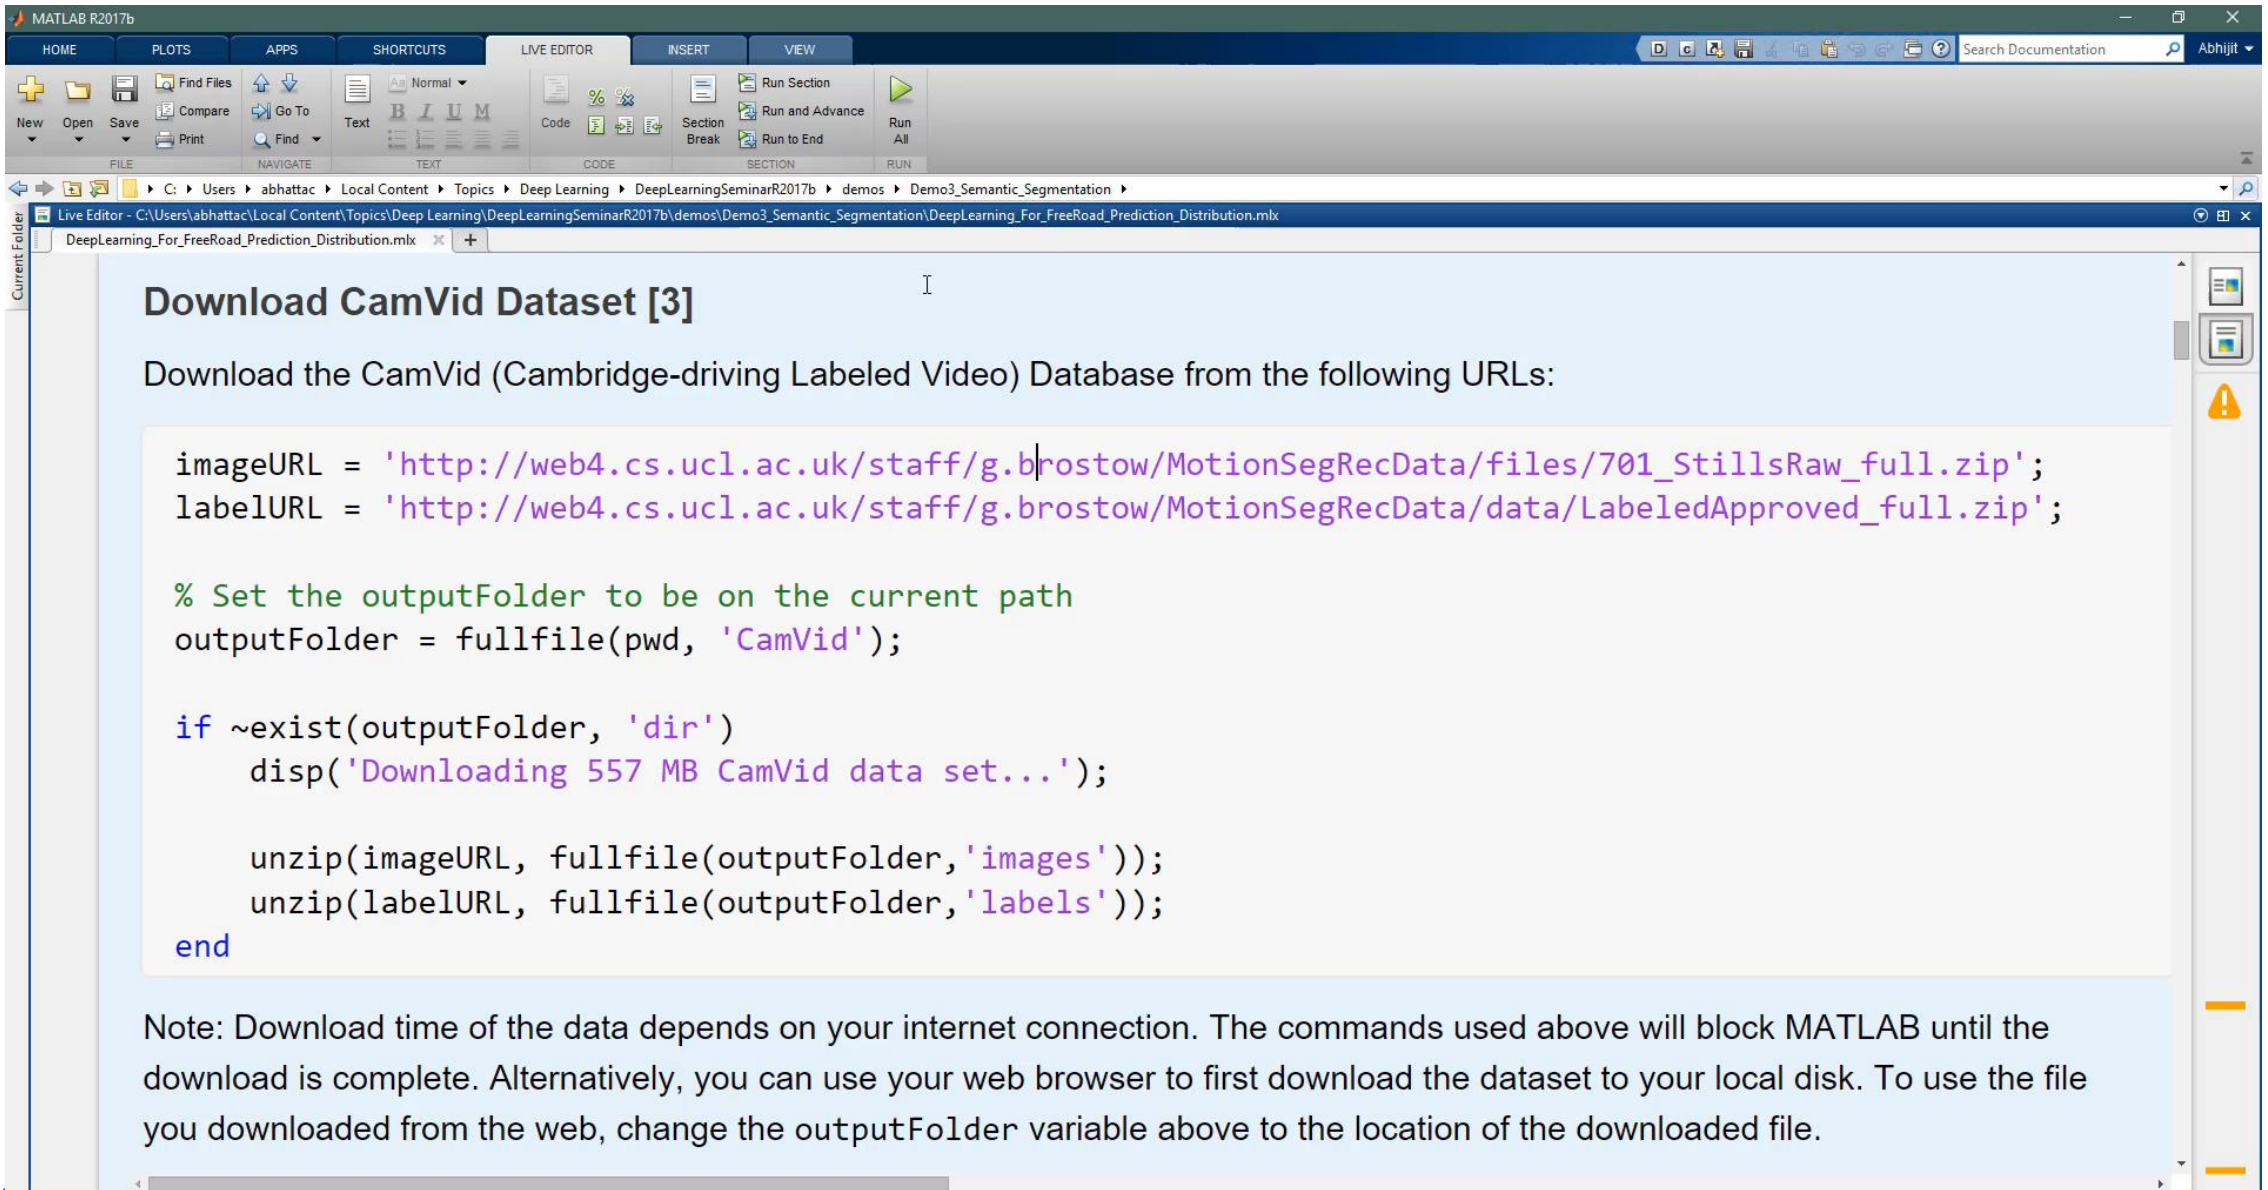
# Semantic Segmentation Network



# Access and Explore Data



# Access image/video data




The image shows the MATLAB R2017b Live Editor interface. The title bar indicates the current file is 'DeepLearning\_For\_FreeRoad\_Prediction\_Distribution.mlx'. The editor contains a script titled 'Download CamVid Dataset [3]'. The script defines two URLs for downloading image and label data, sets an output folder, and uses the 'unzip' function to extract the files. A note at the bottom explains that the download process will block MATLAB until completion and offers an alternative method using a web browser.


```
imageURL = 'http://web4.cs.ucl.ac.uk/staff/g.brostow/MotionSegRecData/files/701_StillsRaw_full.zip';  
labelURL = 'http://web4.cs.ucl.ac.uk/staff/g.brostow/MotionSegRecData/data/LabeledApproved_full.zip';  
  
% Set the outputFolder to be on the current path  
outputFolder = fullfile(pwd, 'CamVid');  
  
if ~exist(outputFolder, 'dir')  
    disp('Downloading 557 MB CamVid data set...');  
  
    unzip(imageURL, fullfile(outputFolder, 'images'));  
    unzip(labelURL, fullfile(outputFolder, 'labels'));  
end
```

Note: Download time of the data depends on your internet connection. The commands used above will block MATLAB until the download is complete. Alternatively, you can use your web browser to first download the dataset to your local disk. To use the file you downloaded from the web, change the outputFolder variable above to the location of the downloaded file.

# Access pixel label data



# Develop and Test Algorithms



# Prepare the network for training

MATLAB R2017b

HOME PLOTS APPS SHORTCUTS LIVE EDITOR INSERT VIEW

Get More Apps Install App Package App Image Viewer Video Viewer Camera Calibrator Filter Designer MATLAB Coder Application Compiler Color Thresholder Image Acquisition Image Region Analyzer Signal Analyzer Classification Learner Ground Truth Labeler

FILE APPS

C:\Users\abhatac\Local Content\Topics\Deep Learning\DeepLearningSeminarR2017b\demos\Demo3\_Semantic\_Segmentation

Live Editor - C:\Users\abhatac\Local Content\Topics\Deep Learning\DeepLearningSeminarR2017b\demos\Demo3\_Semantic\_Segmentation\DeepLearning\_For\_FreeRoad\_Prediction\_Distribution.mlx

DeepLearning\_For\_FreeRoad\_Prediction\_Distribution.mlx


Bicyclist  
 Pedestrian  
 Car  
 SignSymbol  
 Pavement  
 Lane  
 Road  
 Pole  
 Building  
 Environment

# Train and test the network

The image displays the MATLAB R2017b Live Editor interface. The top menu bar includes HOME, PLOTS, APPS, SHORTCUTS, LIVE EDITOR, INSERT, and VIEW. The toolbar contains icons for file operations (New, Open, Save, Print), navigation (Find Files, Go To, Find), text formatting (Text, Bold, Italic, Underline, Monospace), code execution (Run Section, Run and Advance, Run to End, Run All), and search (Search Documentation). The current folder is `C:\Users\abhatac\Local Content\Topics\Deep Learning\DeepLearningSeminarR2017b\demos\Demo3_Semantic_Segmentation`. The live editor shows a neural network diagram for semantic segmentation, consisting of two identical vertical chains of blocks connected by downward arrows. The left chain includes: `decoder1_unpool`, `decoder1_conv2`, `decoder1_bn_2`, `decoder1_relu_2`, `decoder1_conv1`, `decoder1_bn_1`, `decoder1_relu_1`, `softmax`, and `pixelLabels`. The right chain includes: `decoder1_unpool`, `decoder1_conv2`, `decoder1_bn_2`, `decoder1_relu_2`, `decoder1_conv1`, `decoder1_bn_1`, `decoder1_relu_1`, `softmax`, and `labels`. A warning icon is visible in the bottom right corner of the live editor window.

# Evaluate algorithm performance

Ground Truth labels vs Predicted labels



# More evaluation

MATLAB R2017b

HOME PLOTS APPS SHORTCUTS LIVE EDITOR INSERT VIEW

FILE NAVIGATE TEXT CODE SECTION RUN

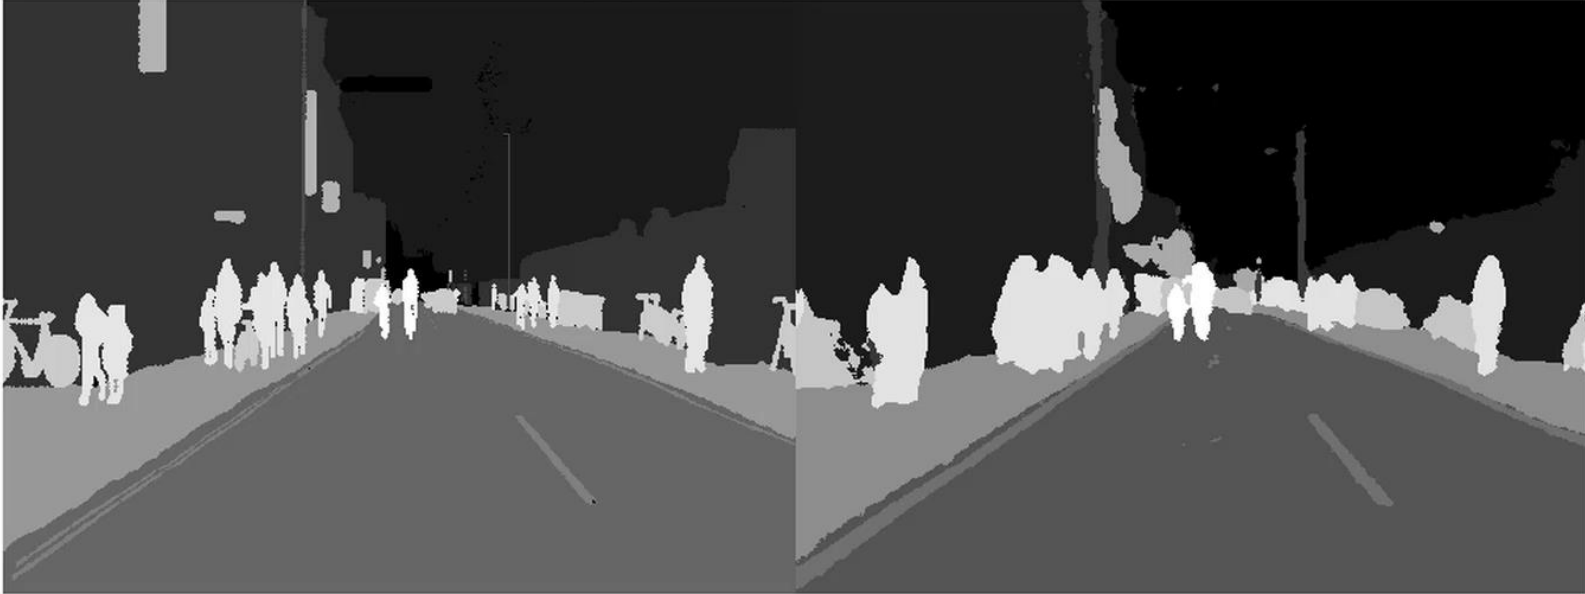
Search Documentation Abhijit

C:\Users\abhatac\Local Content\Topics\Deep Learning\DeepLearningSeminarR2017b\demos\Demo3\_Semantic\_Segmentation

Live Editor - C:\Users\abhatac\Local Content\Topics\Deep Learning\DeepLearningSeminarR2017b\demos\Demo3\_Semantic\_Segmentation\DeepLearning\_For\_FreeRoad\_Prediction\_Distribution.mlx

```
% Compare differences between images - Image Processing toolbox  
imshowpair(expected, predicted, 'montage')  
title('Ground Truth labels vs Predicted labels')
```

Ground Truth labels vs Predicted labels



MATLAB

# Challenges we addressed

- Accessing data
- Labeling and preprocessing
- Using previous research
- Developing an algorithm
- Evaluating the algorithm
- **What's next?**

# Deep Learning on CPU, GPU, Multi-GPU and Clusters


## HOW TO TARGET?




Single CPU




Single CPU  
Single GPU



Single CPU, Multiple GPUs



On-prem server with GPUs



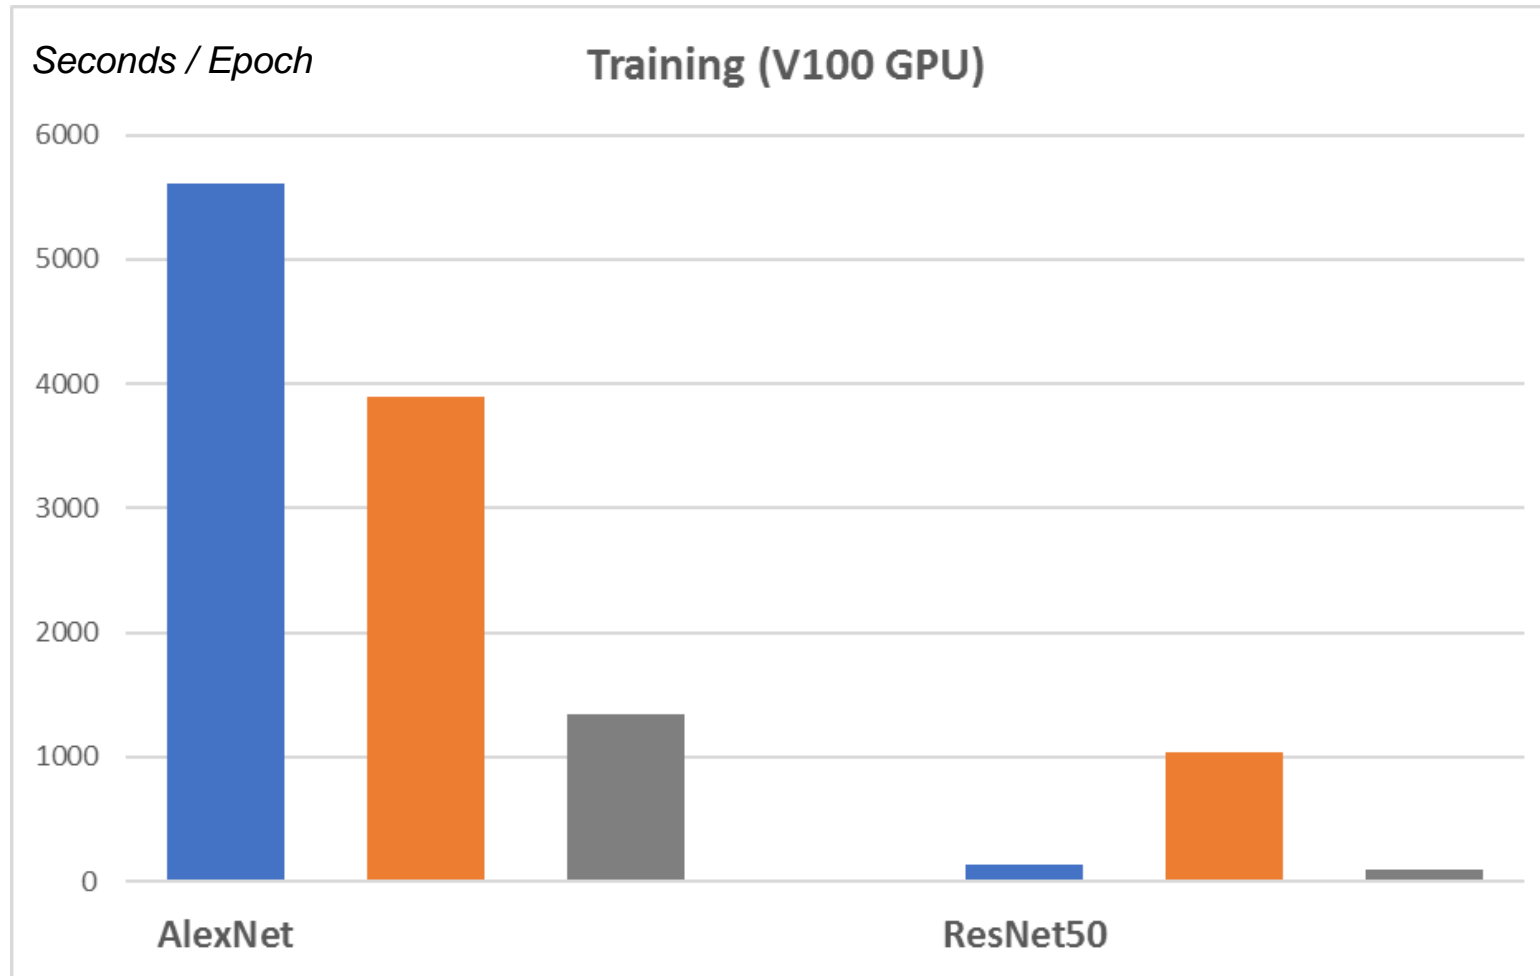
Cloud GPUs (AWS)

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

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

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

# Training Performance



**TensorFlow**

**MATLAB**

**MXNet**

*Batch size 32*


# Training is an Iterative Process

```
miniBatchSize = 128;
validationFrequency = floor(numel(YTrain)/miniBatchSize);
options = trainingOptions('adam', ...
    'InitialLearnRate',5e-4, ...
    'MaxEpochs',25, ...
    'MiniBatchSize',miniBatchSize, ...
    'Shuffle','every-epoch', ...
    'Plots','training-progress', ...
    'Verbose',false, ...
    'ValidationData',{XValidation,YValidation}, ...
    'ValidationFrequency',validationFrequency, ...
    'ValidationPatience',Inf, ...
    'LearnRateSchedule','piecewise', ...
    'LearnRateDropFactor',0.1, ...
    'LearnRateDropPeriod',20);
```

*Parameters adjusted according to performance*

# MATLAB is Fast for Deployment

- Target a GPU for optimal performance
- NVIDIA GPUs use CUDA code
- We only have MATLAB code.  
Can we translate this?




# GPU Coder

- Automatically generates **CUDA** Code from MATLAB Code
  - can be used on NVIDIA GPUs




- CUDA extends C/C++ code with constructs for parallel computing

# GPU Coder Performance



# 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**

# Why MATLAB?

- MATLAB is Productive
- MATLAB is Fast
- **MATLAB Integrates with Open Source**

# Used MATLAB and Open Source Together



- Used Caffe and MATLAB together
- Achieved significantly better results than an engineered rain model.
- Use our tools where it makes your workflow easier!

1. *Deep Joint Rain Detection and Removal from a Single Image*" Wenhan Yang, Robby T. Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan

# MATLAB Integrates with Open Source Frameworks

- Access to many pretrained models through add-ons
- Users wanted to import latest models
- Import models directly from Tensorflow, Caffe, or other frameworks using ONNX
- Export models to ONNX, which can be read by other frameworks
  - Allows for improved collaboration

## KERAS IMPORTER

Importer for TensorFlow-Keras Models


## Caffe

M O D E L S

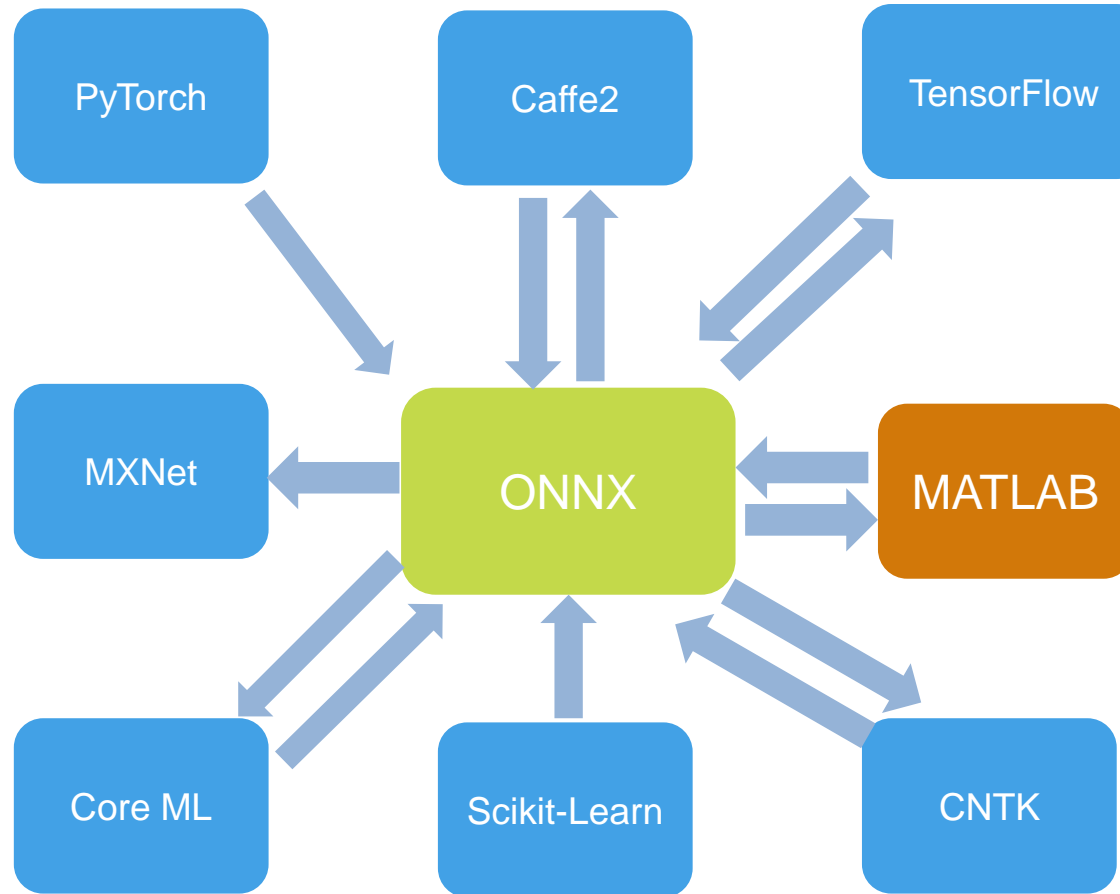
## ONNX Converter

Export to ONNX model format

# Keras-Tensorflow Importer



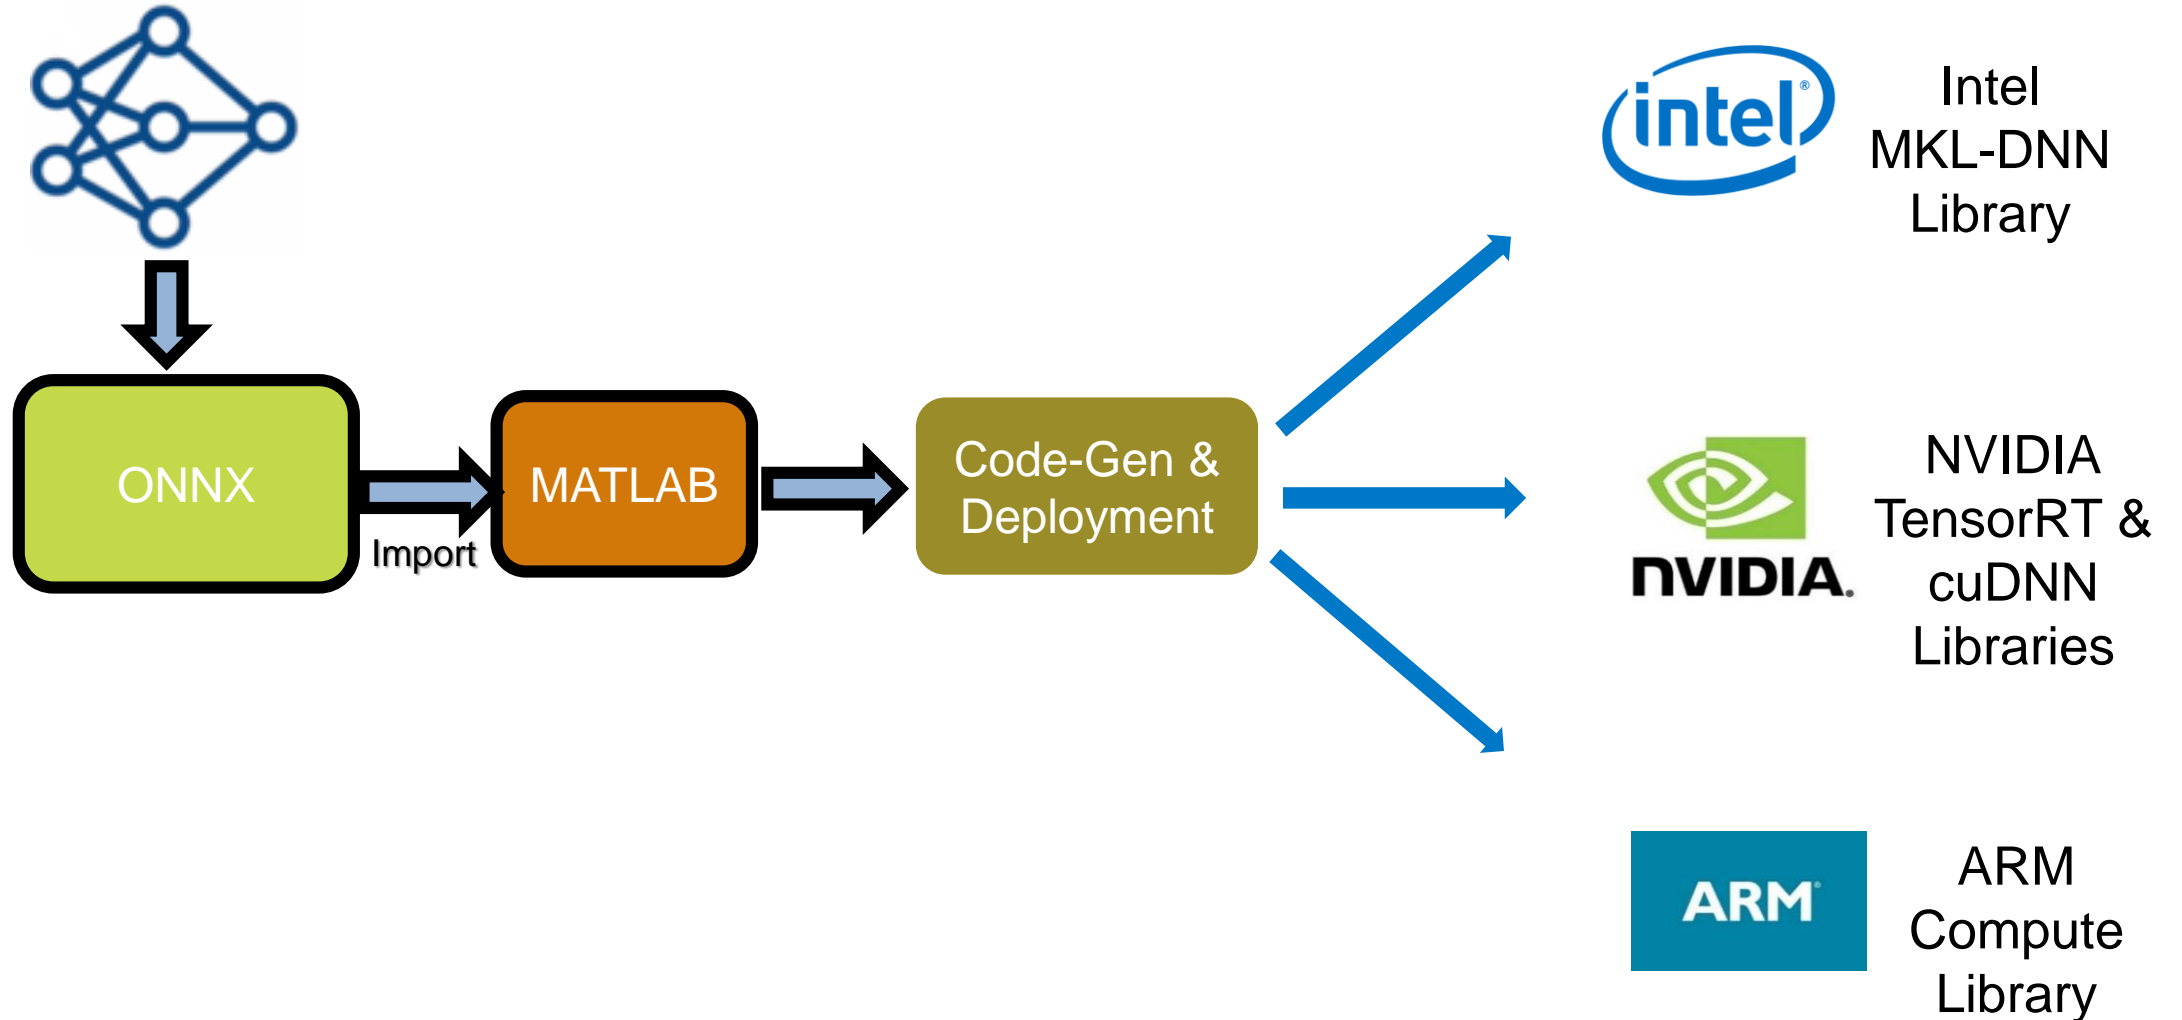
# ONNX is an open format to represent deep learning models



ONNX = Open Neural Network Exchange Format

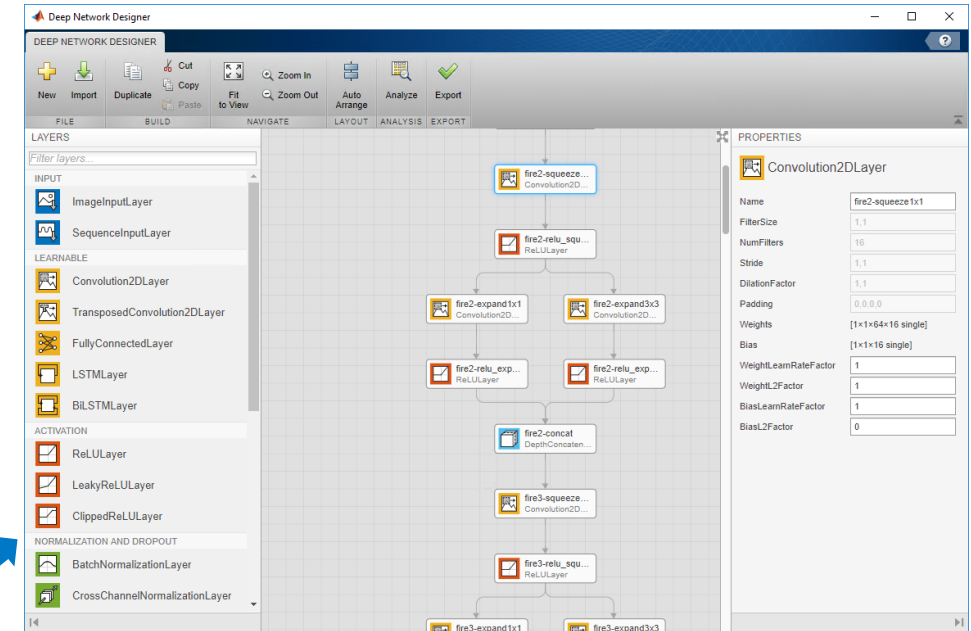
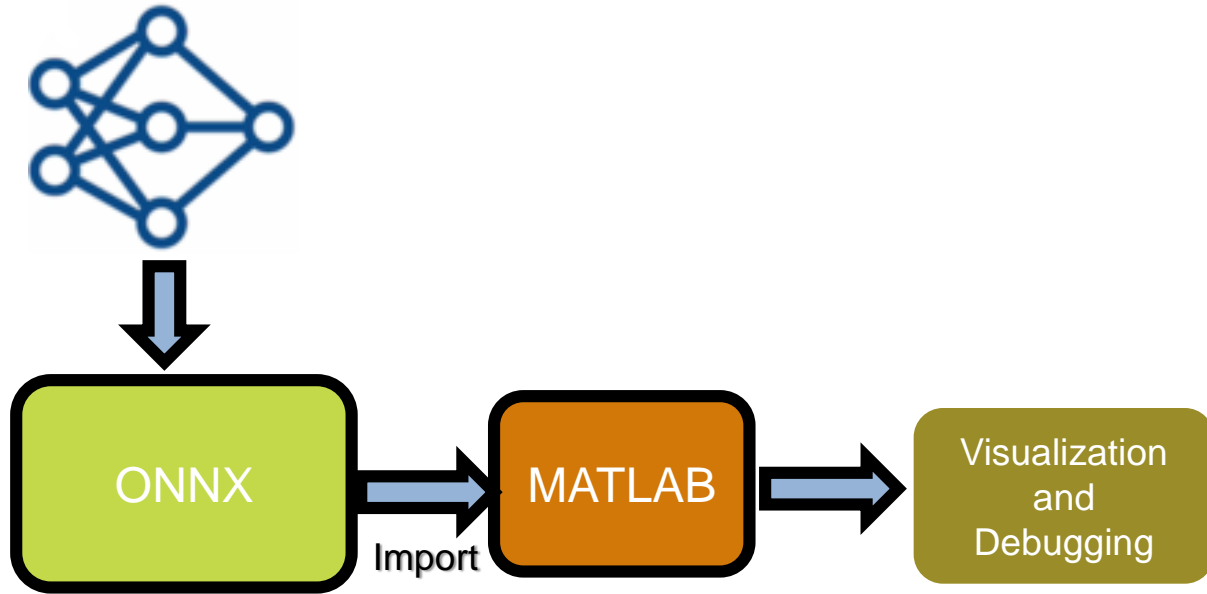
# 3 Reasons to Import ONNX models into MATLAB

## A. Code Generation and Multi-target deployment




# 3 Reasons to Import ONNX models into MATLAB

## B. Visualization and Debugging



Network Designer



Visualize Layer Activations

# 3 Reasons to Import ONNX models into MATLAB

## C. Improving Model Accuracy and Optimization

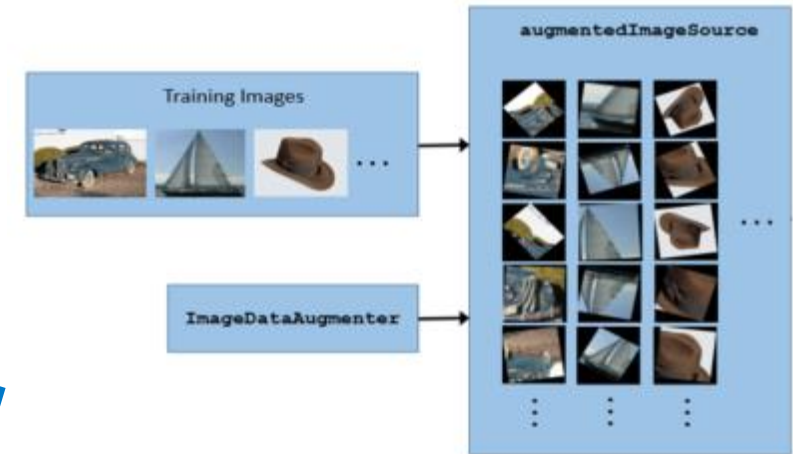
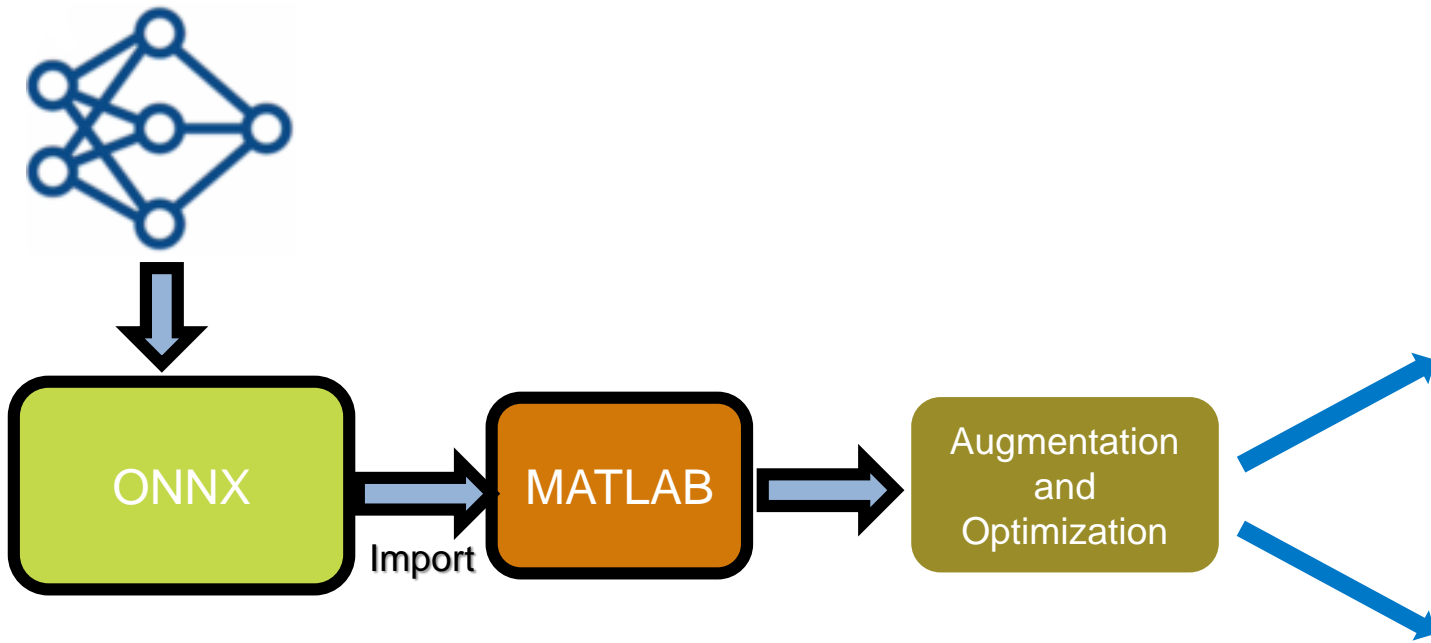
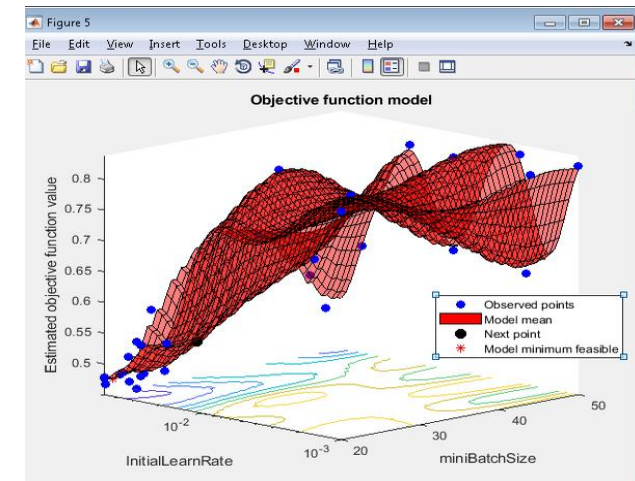


Image Augmentation




Hyperparameter tuning using Bayesian Optimization

# ONNX Converter

## Import from ONNX format

```
modelfile = 'cifarResNet.onnx';
classes = ["airplane" "automobile" "bird"];
net = importONNXNetwork(modelfile, 'Output');
analyzeNetwork(net)
```



## Export to ONNX format

```
filename = 'fishdetector.onnx';
exportONNXNetwork(net, filename)
```

net  
Analysis date: 04-Jul-2018 08:16:54

77 layers 0 warnings 0 errors

	NAME	TYPE	ACTIVATIONS	LEARNABLES
1	Input_input 32x32x3 images	Image Input	32x32x3	-
2	input_Sub Layer adds bias to the input	Bias	32x32x3	-
3	convInp 16 3x3x3 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	32x32x16	Weights 3x3x3x16 Bias 1x1x16
4	BNInp Batch normalization with 16 channels	Batch Normalization	32x32x16	Offset 1x1x16 Scale 1x1x16
5	reluInp ReLU	ReLU	32x32x16	-
6	S1U1_conv1 16 3x3x16 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	32x32x16	Weights 3x3x16x16 Bias 1x1x16
7	S1U1_BN1 Batch normalization with 16 channels	Batch Normalization	32x32x16	Offset 1x1x16 Scale 1x1x16
8	S1U1_relu1 ReLU	ReLU	32x32x16	-
9	S1U1_conv2 16 3x3x16 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	32x32x16	Weights 3x3x16x16 Bias 1x1x16
10	S1U1_BN2 Batch normalization with 16 channels	Batch Normalization	32x32x16	Offset 1x1x16 Scale 1x1x16
11	add11 Element-wise addition of 2 inputs	Addition	32x32x16	-
12	relu11 ReLU	ReLU	32x32x16	-
13	S1U2_conv1 16 3x3x16 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	32x32x16	Weights 3x3x16x16 Bias 1x1x16
14	S1U2_BN1 Batch normalization with 16 channels	Batch Normalization	32x32x16	Offset 1x1x16 Scale 1x1x16
15	S1U2_relu1 ReLU	ReLU	32x32x16	-
16	S1U2_conv2 16 3x3x16 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	32x32x16	Weights 3x3x16x16 Bias 1x1x16
17	S1U2_BN2 Batch normalization with 16 channels	Batch Normalization	32x32x16	Offset 1x1x16 Scale 1x1x16

# MATLAB Integrates with Open Source Frameworks


[Contact Us](#)
[How to Buy](#)
[Jon Cherrie](#)

[MathWorks®](#)
[Products](#)
[Solutions](#)
[Academia](#)
[Support](#)
[Community](#)
[Events](#)

**Community Profile**

[Support](#)

[MATLAB Answers](#)
[File Exchange](#)
[Cody](#)
[Blogs](#)
[ThingSpeak](#)
[Maker Community](#)
[Trial software](#)



**MathWorks Deep Learning Toolbox Team**  
STAFF  
 MathWorks  
 18 total contributions since 2017

Contributions in 
View by

↔	Import models from other frameworks via ONNX or direct from Keras-TensorFlow and Caffe	<b>ONNX Converter</b> <small>Export to ONNX model format</small>
↔	Pre-trained models	<b>KERAS IMPORTER</b> <small>Importer for TensorFlow-Keras Models</small>
↔	Pre-trained models	<b>ResNet-50</b> <small>PRETRAINED MODEL</small>
↔	Pre-trained models	<b>SqueezeNet</b> <small>PRETRAINED MODEL</small>

# MATLAB Integrates with Open Source Frameworks

- MATLAB supports entire deep learning workflow
  - Use when it is convenient for your workflow
- Access to latest models
- Improved collaboration with other users

# Why MATLAB for Deep Learning?

- MATLAB is Productive
- MATLAB is Fast (*Performance*)
- MATLAB Integrates with Open Source (*Frameworks*)