MATLAB EXPO 2018

Demystifying Deep Learning

“Let the computers do the hard work”

Pitambar Dayal
Product Marketing
Deep Learning, Image Processing


Deep Learning Demo
Image Classification

Why MATLAB for Deep Learning?

- MATLAB is Productive
- MATLAB Supports the Entire Deep Learning Workflow
- MATLAB Integrates with Open Source

What is Deep Learning?

A. Deep learning is learning done really far underground.

B. I don’t know, I’m just here for the free snacks.

C. Deep learning is machine learning with automatic feature extraction. It uses a neural network architecture to perform tasks like classification, detection, and regression.

What is Deep Learning?

- Subset of machine learning with automatic feature extraction
  - Learns features and tasks directly from data
  - More Data = better model

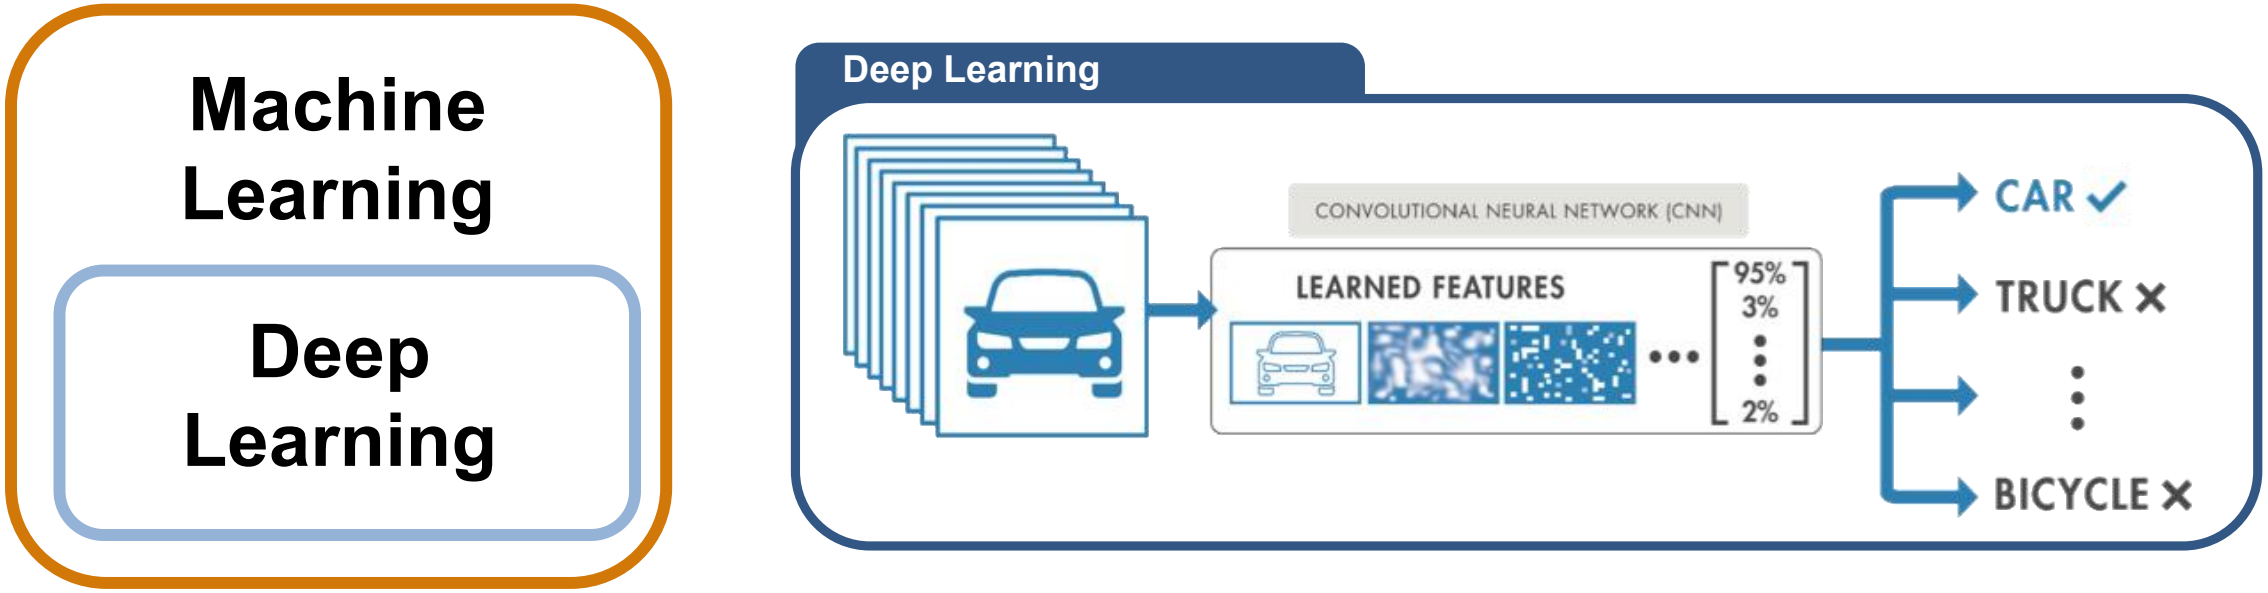

Deep Learning Uses a Neural Network Architecture

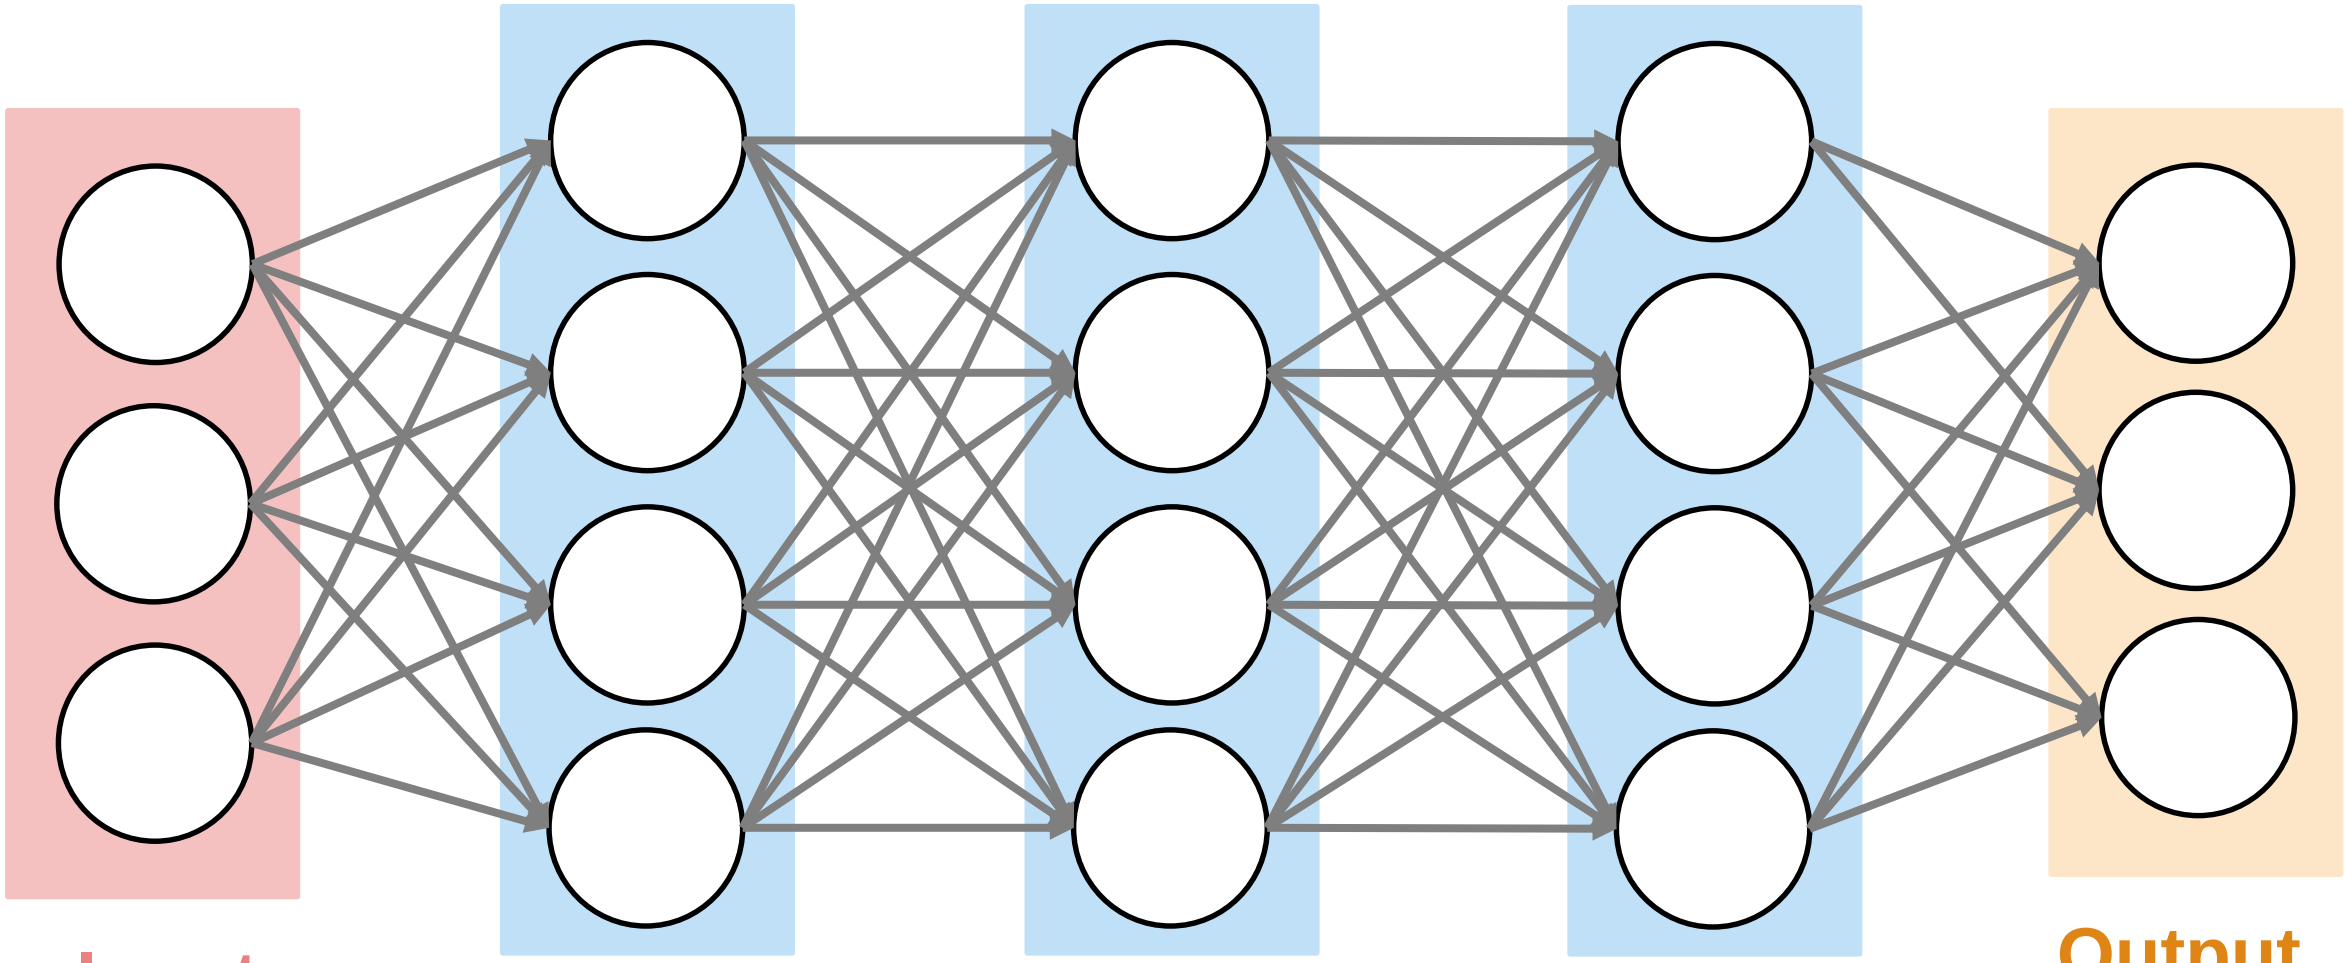

Deep Learning Datatypes

Image

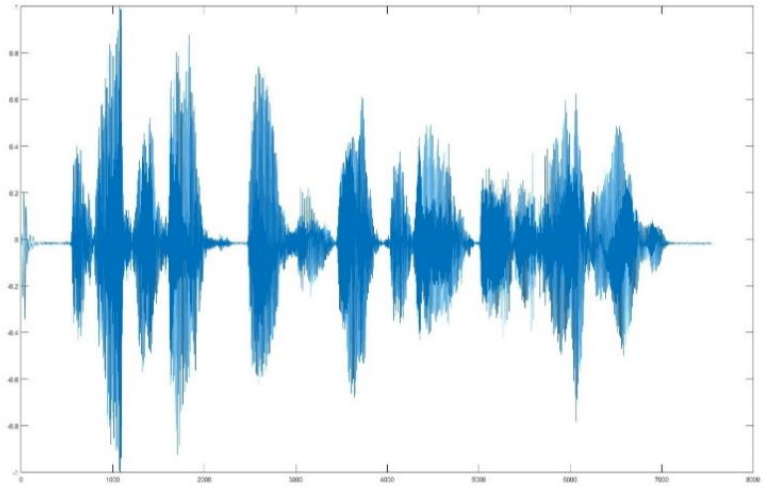

Signal

Numeric

AgeCat	WeightQ	GroupCount	mean_BloodPressure	
Under 30	Q1	6	123.17	79.667
Under 30	Q2	3	120.33	79.667
Under 30	Q3	2	127.5	86.5
Under 30	Q4	4	122	78
30-39	Q1	12	121.75	81.75
30-39	Q2	9	119.56	82.556
30-39	Q3	9	121	83.222
30-39	Q4	11	125.55	87.273
Over 40	Q1	7	122.14	84.714
Over 40	Q2	13	123.38	79.385
Over 40	Q3	14	123.07	84.643
Over 40	Q4	10	124.6	85.1

Text

Deep Learning is Versatile

Detection of cars and road in autonomous driving systems

Rain Detection and Removal[1]

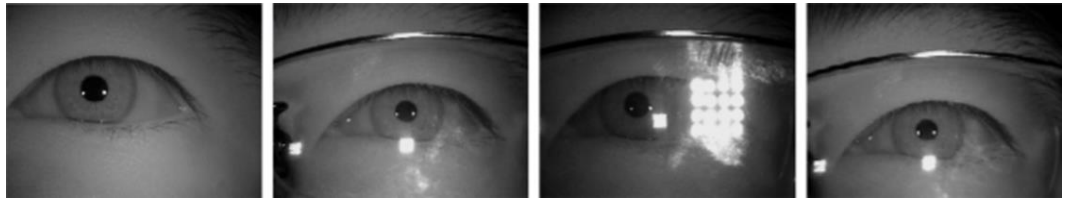

Iris Recognition – 99.4% accuracy[2]

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*

Why MATLAB for Deep Learning?

- MATLAB is Productive
- MATLAB Supports the Entire Deep Learning Workflow
- MATLAB integrates with Open Source

Deep Learning in 6 Lines of MATLAB Code

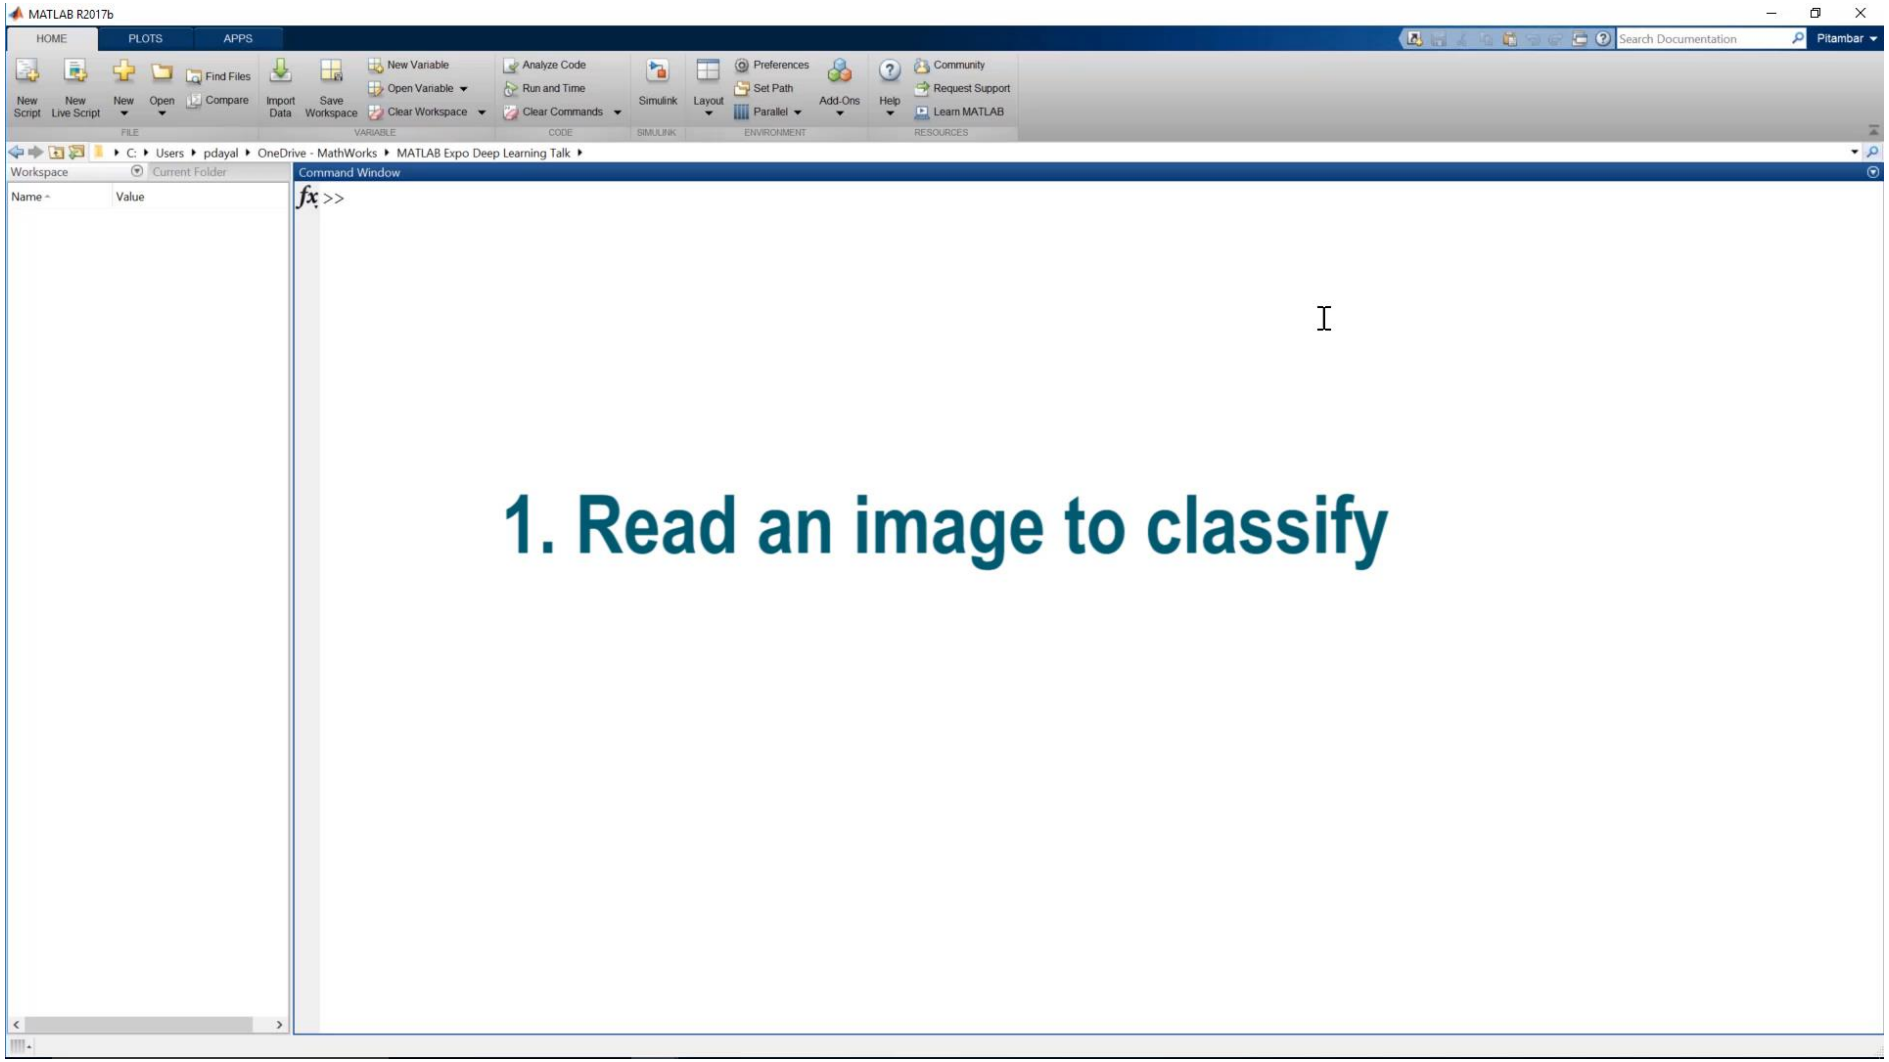

“I love to label and preprocess my data”

~ 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"*

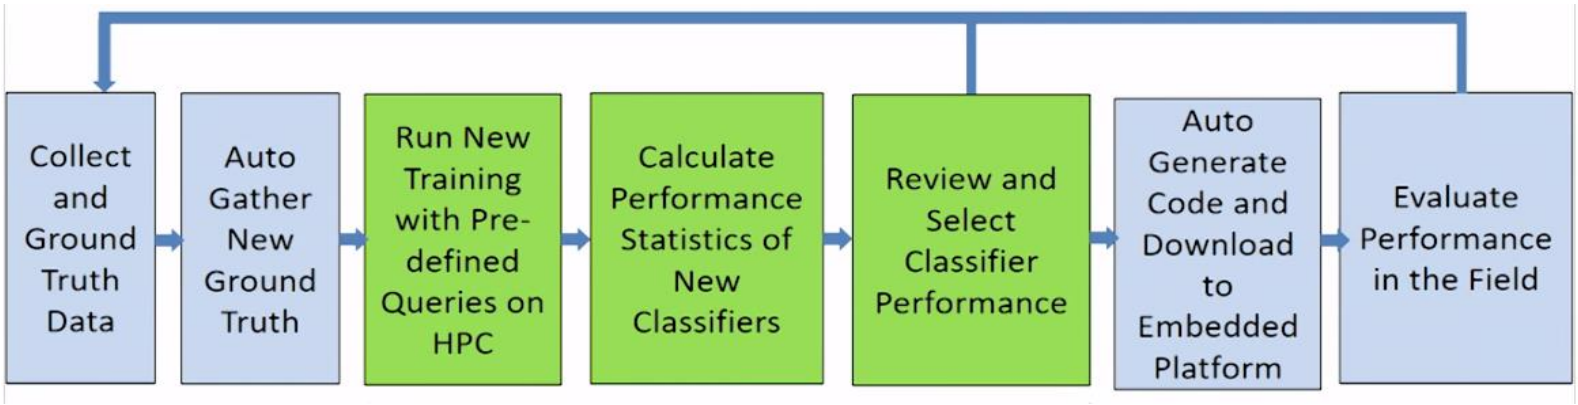

How Does MATLAB Come into Play?

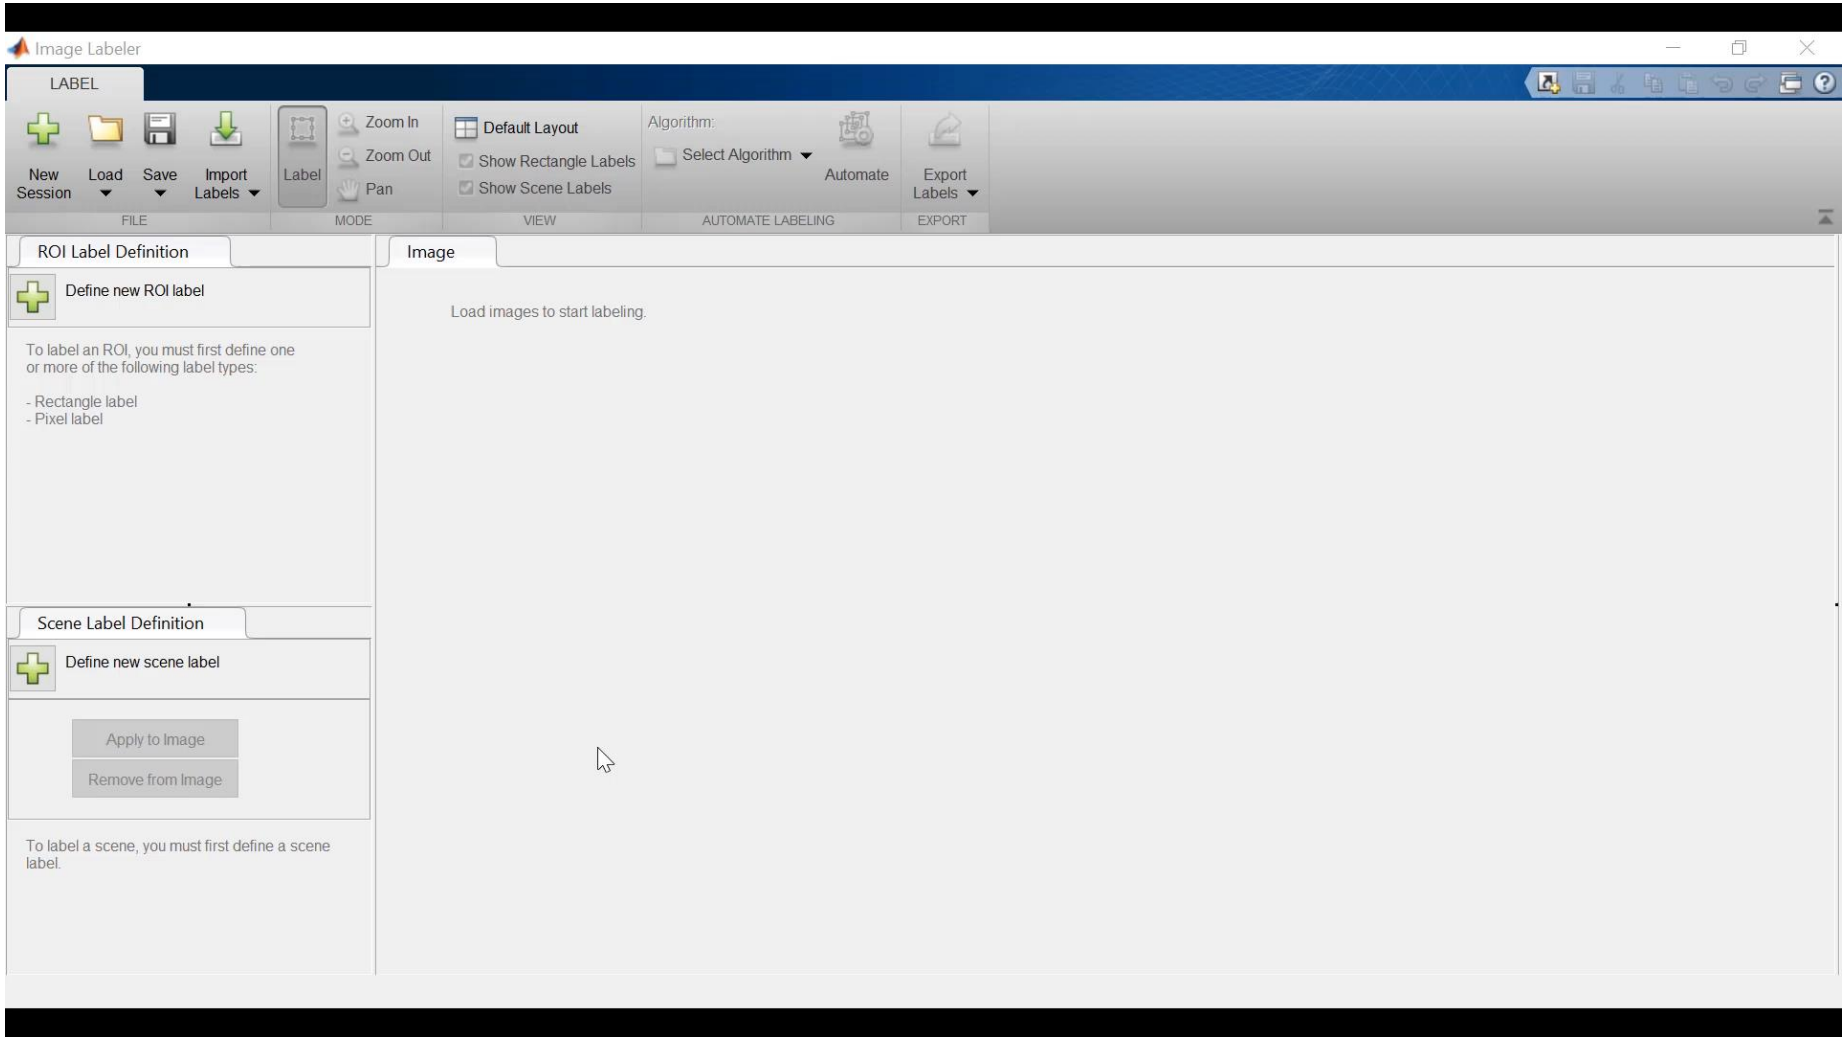

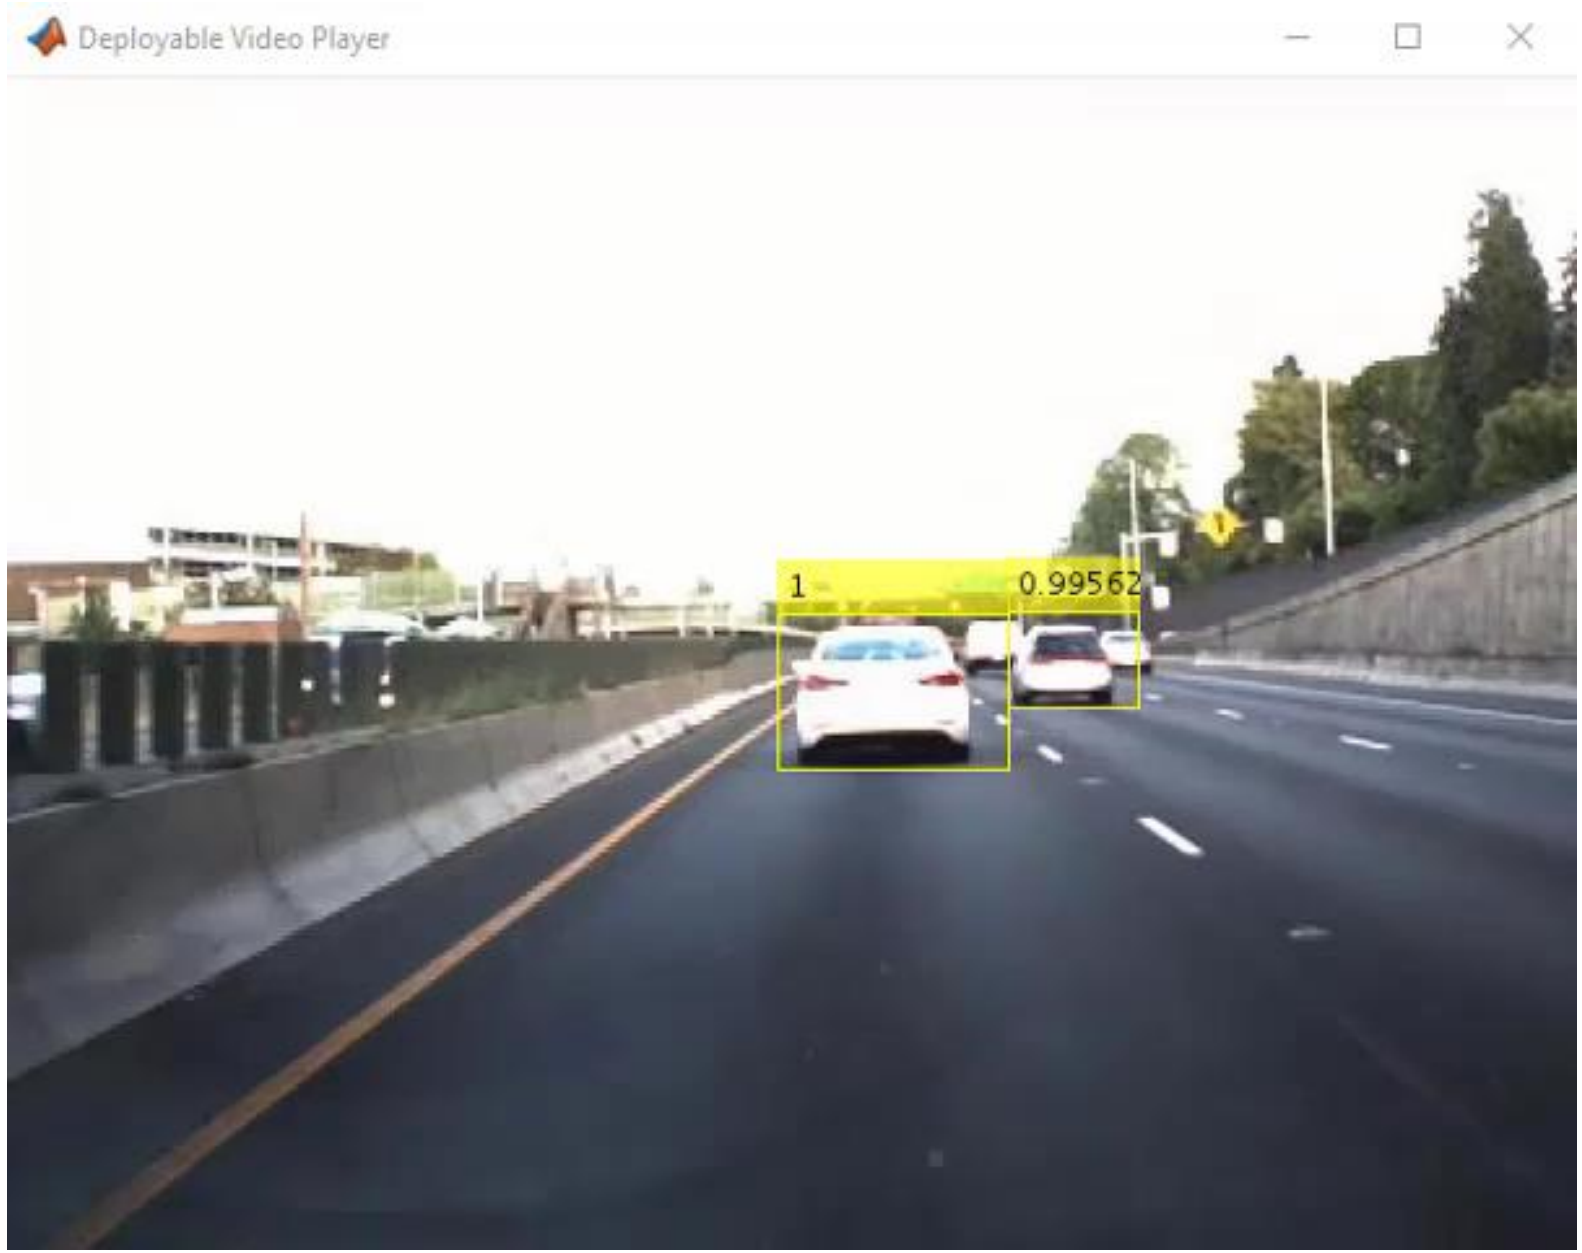
Deployable Video Player
1
0.99562

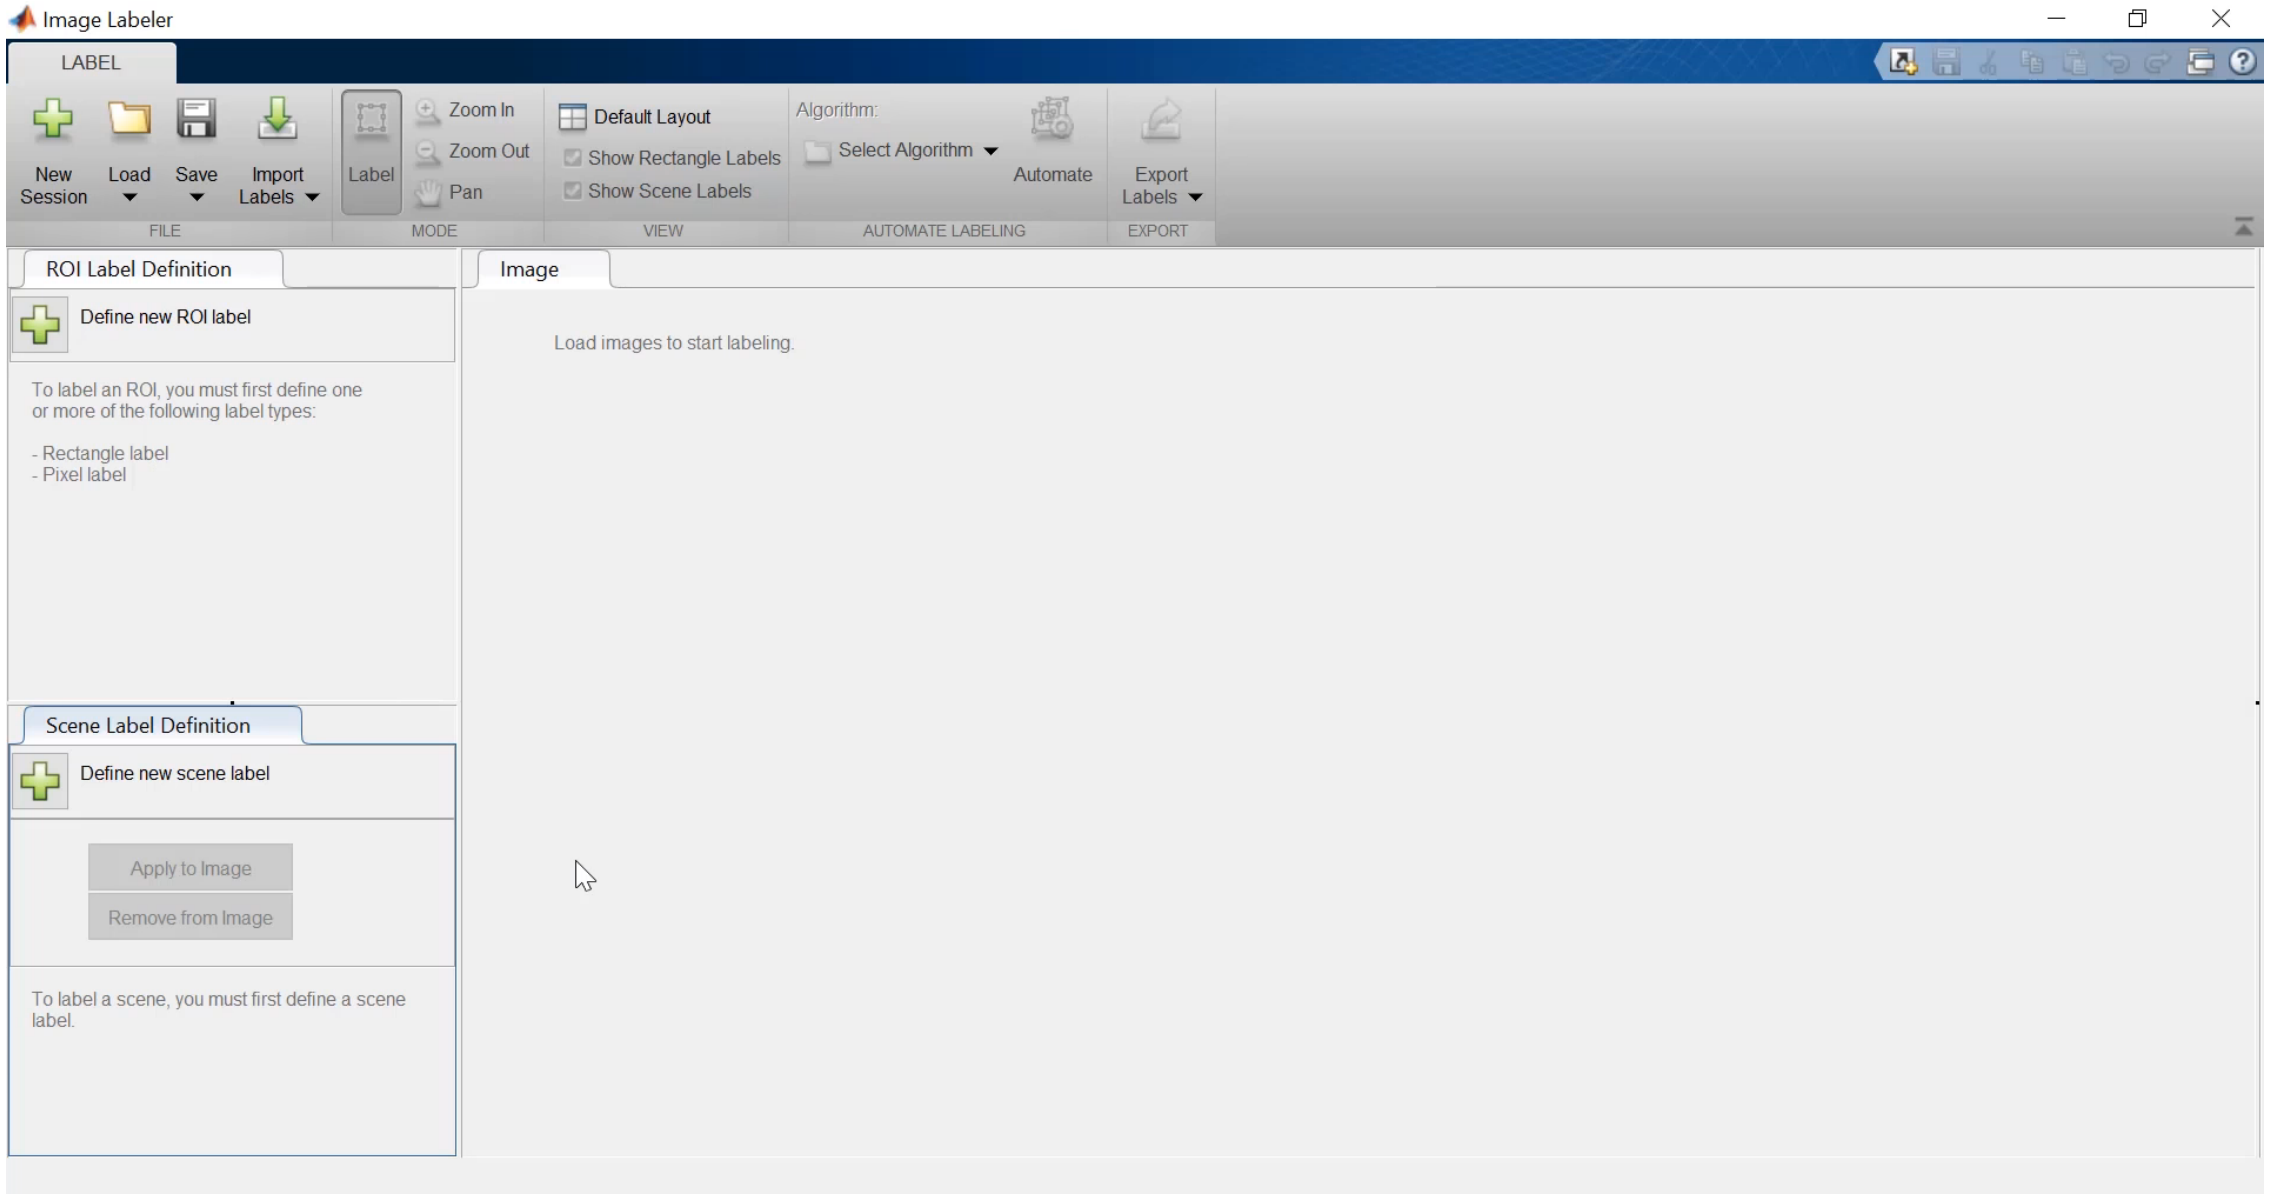
Image Labeler
LABEL
New Session
Load
Save
Import Labels
FILE
Label
Zoom In
Zoom Out
Pan
MODE
Default Layout
Show Rectangle Labels
Show Scene Labels
VIEW
Algorithm:
Select Algorithm
Automate
AUTOMATE LABELING
Export Labels
EXPORT
ROI Label Definition
Define new ROI label
To label an ROI, you must first define one or more of the following label types:
- Rectangle label
- Pixel label
Scene Label Definition
Define new scene label
Apply to Image
Remove from Image
To label a scene, you must first define a scene label.
Image
Load images to start labeling.


RICE HOUSE

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

Why MATLAB for Deep Learning?

- MATLAB is Productive
- MATLAB Supports the Entire Deep Learning Workflow
- MATLAB integrates with Open Source

Deep Learning Workflow

Access and Explore Data

Files

Databases

Sensors

Label and Preprocess Data

Data Augmentation/ Transformation

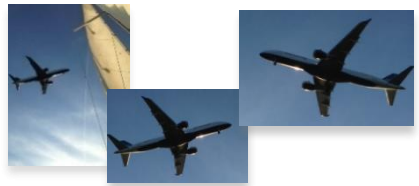

Labeling Automation

Import Reference Models

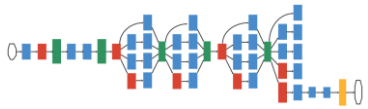

Develop Predictive Models

Hardware-Accelerated Training

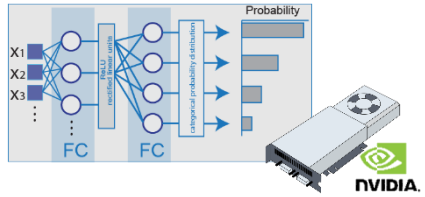

Hyperparameter Tuning

Network Visualization

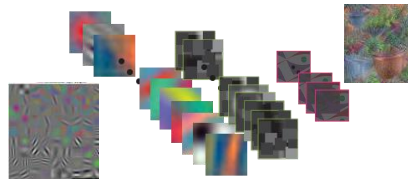

Integrate Models with Systems

Desktop Apps

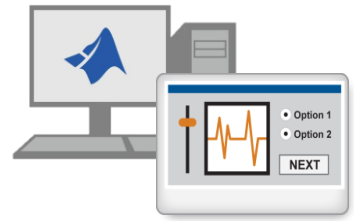

Enterprise Scale Systems

Embedded Devices and Hardware

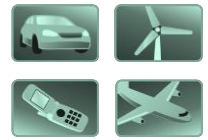

Access and Explore Data

Files

Databases

Sensors

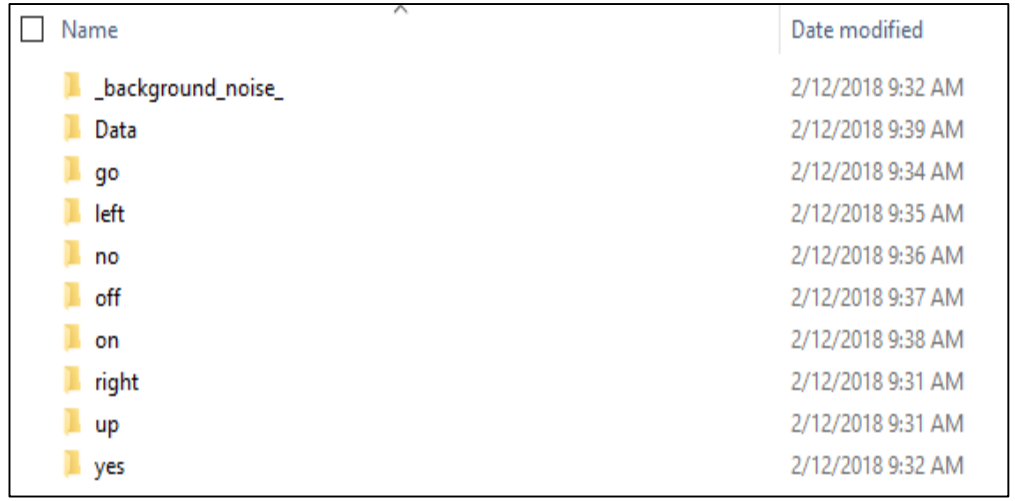


```
datafolder = fullfile(tempdir,'speech_commands_v0.01');
ads = audioDatastore(datafolder, ...
    'IncludeSubfolders',true, ...
    'FileExtensions','.wav', ...
    'LabelSource','foldernames')
```

- Datastore useful for large datasets
- Built in read support for many file types

Label and Preprocess Data

Data Augmentation/ Transformation

Labeling Automation

Import Reference Models

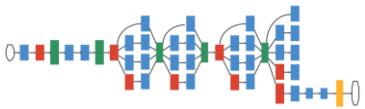

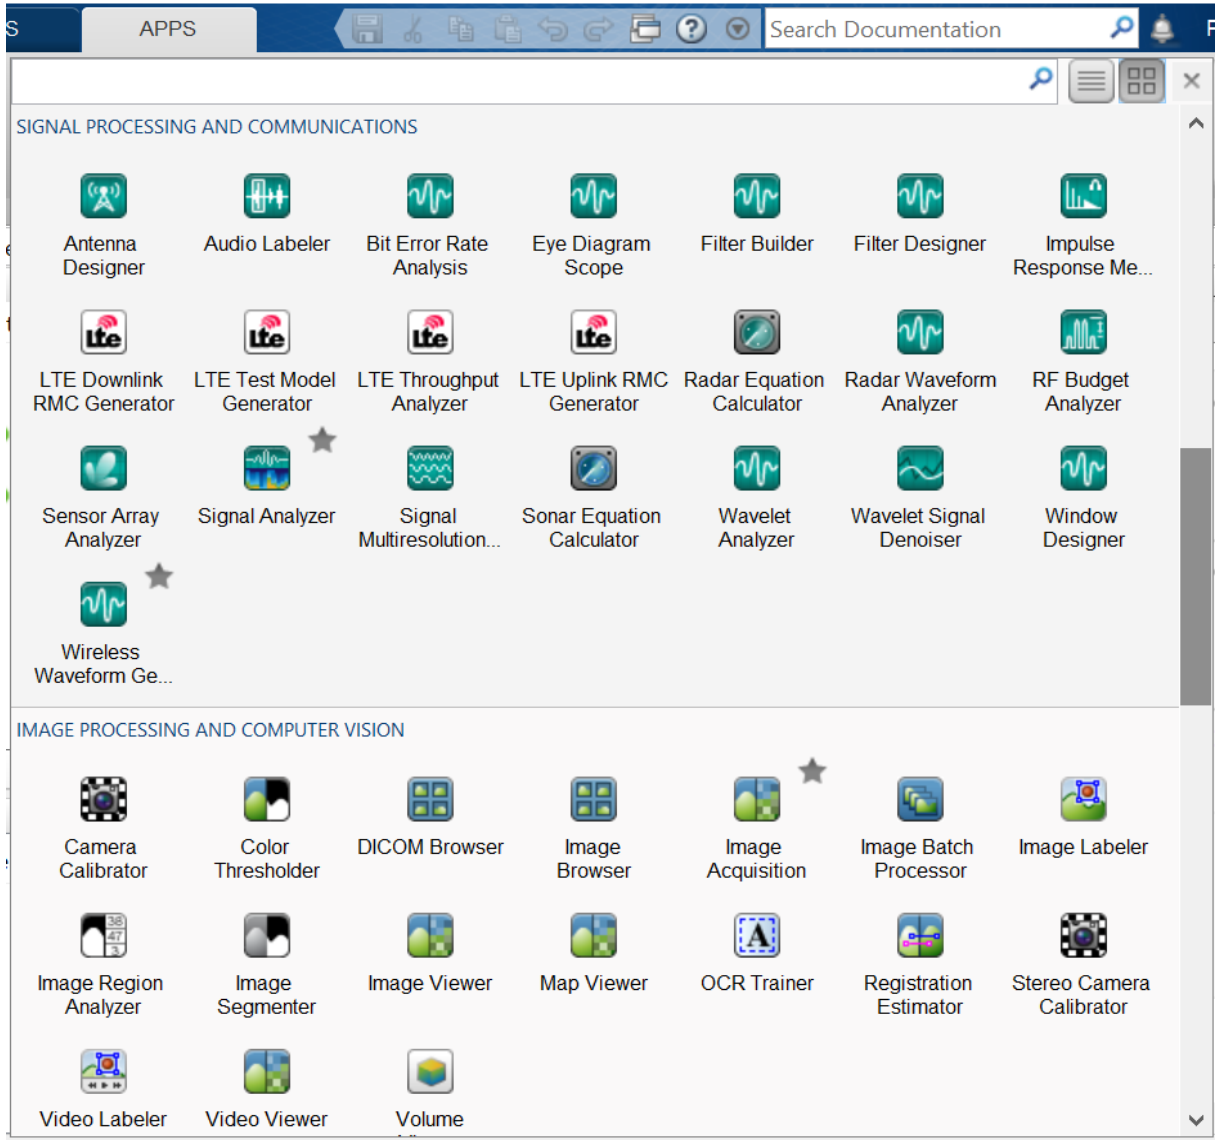

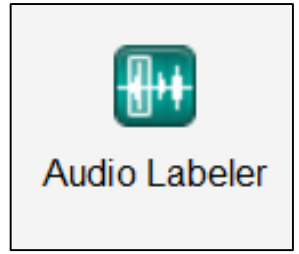

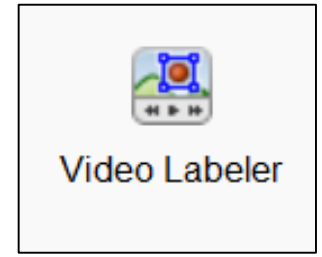

Develop Predictive Models

Hardware-Accelerated Training

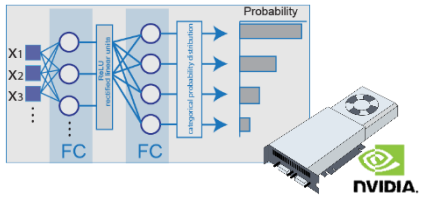

Hyperparameter Tuning

Network Visualization

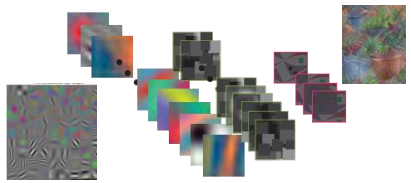

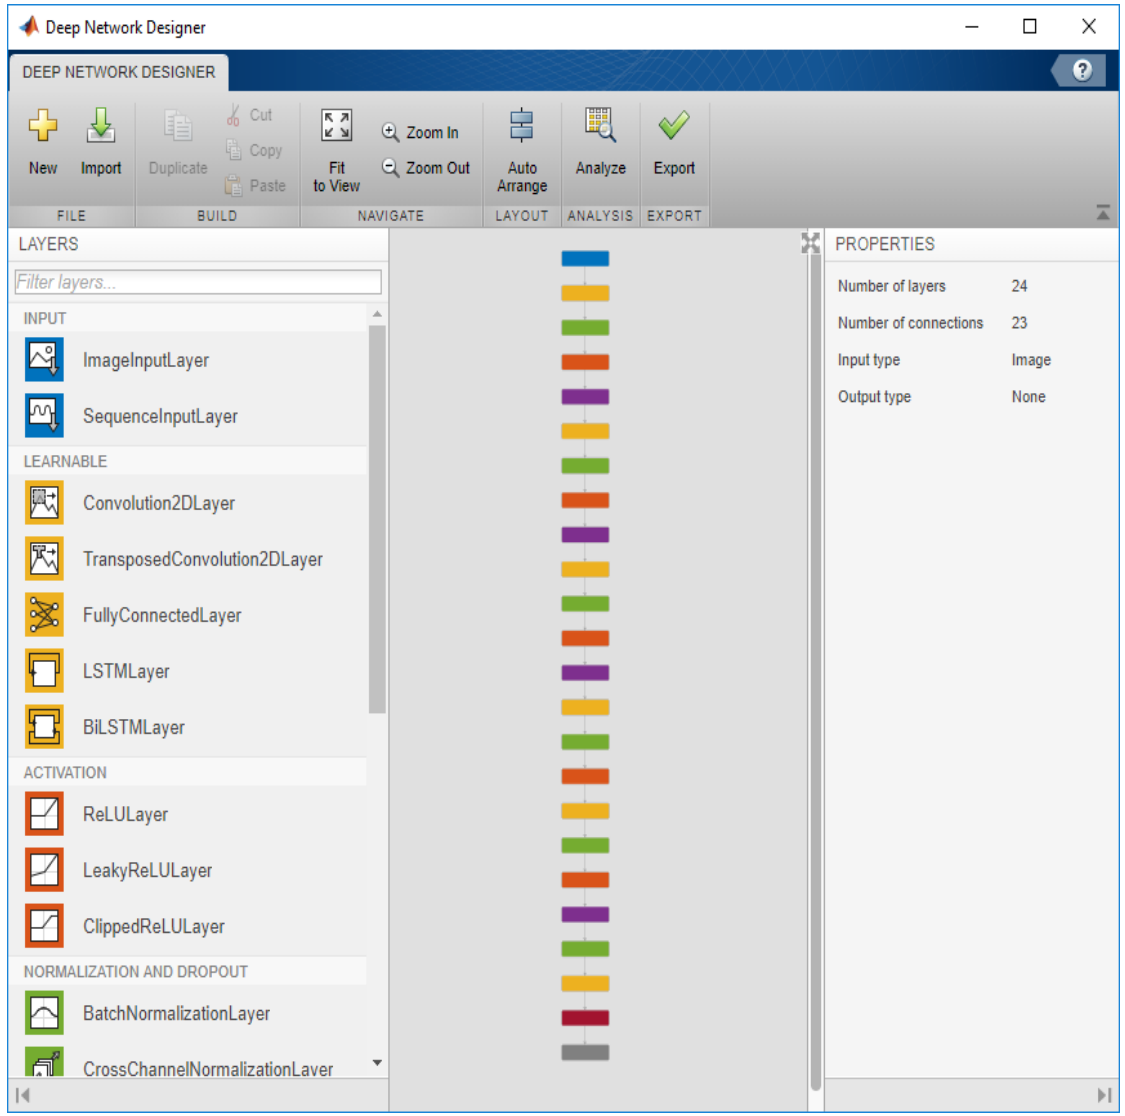

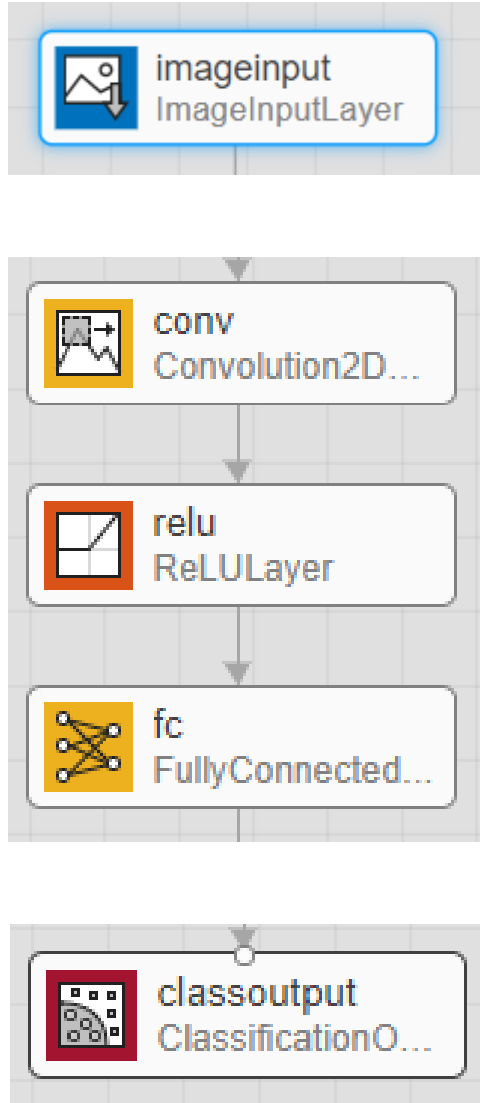

Develop Predictive Models

Hardware-Accelerated Training

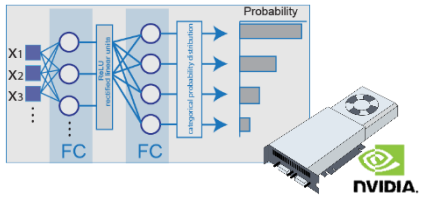

Hyperparameter Tuning

Network Visualization

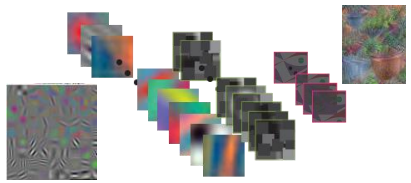

- **Inception-v3** PRETRAINED MODEL
- **ResNet-101** PRETRAINED MODEL
- **GoogLeNet** PRETRAINED MODEL
- **AlexNet** PRETRAINED MODEL
- **VGG-16** PRETRAINED MODEL
- **DenseNet-201** PRETRAINED MODEL
- **ResNet-50** PRETRAINED MODEL
- **VGG-19** PRETRAINED MODEL
- **SqueezeNet** PRETRAINED MODEL
- **ResNet-18** PRETRAINED MODEL
- **ResNet-50** PRETRAINED MODEL
- **Inception-ResNet-v2** PRETRAINED MODEL

Easy access to many pre-trained models

Develop Predictive Models

Hardware-Accelerated Training

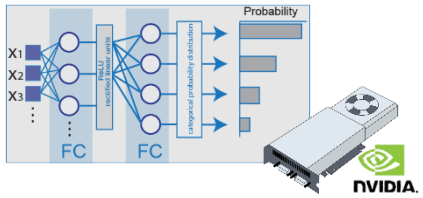

Hyperparameter Tuning

Network Visualization

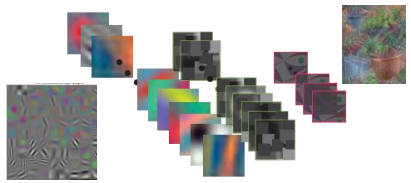

Single CPU

Single CPU
Single GPU

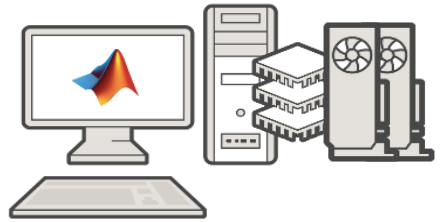

Single CPU, Multiple GPUs

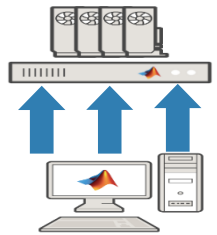

On-prem server with GPUs

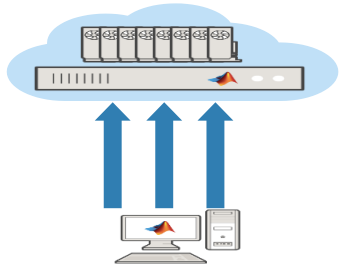

Cloud GPUs (AWS)

```
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' );
```

Integrate Models with Systems

Desktop Apps

Enterprise Scale Systems

Java
MATLAB
Excel
C/C++
.NET
dll
.exe
Python

Embedded Devices and Hardware

Coder Products

MATLAB Code

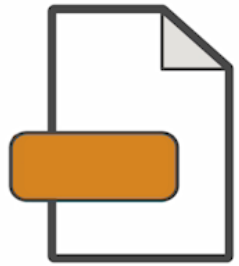

C/C++/HDL/etc.

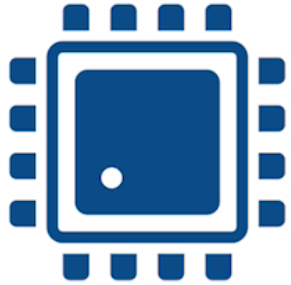

Deployment Target

Deploying Deep Learning Models for Inference

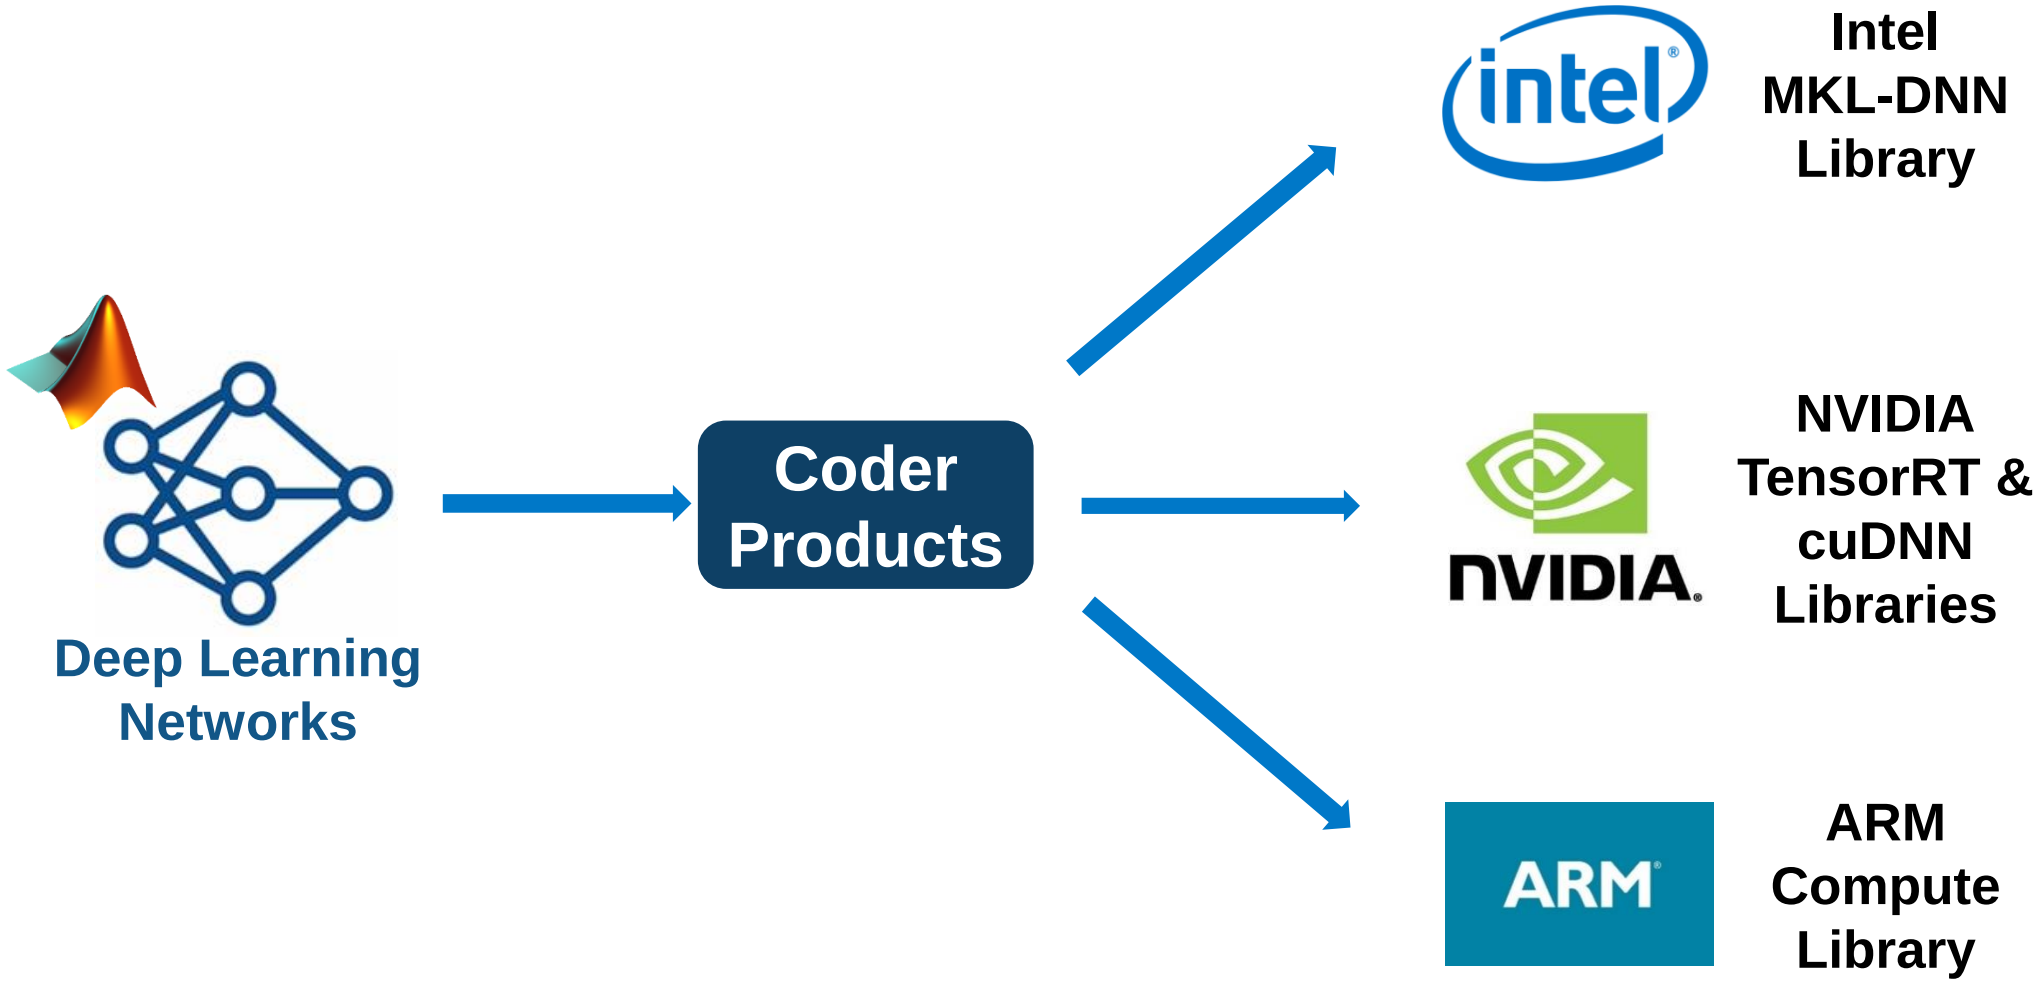

Deploying to Various Targets

R2018a

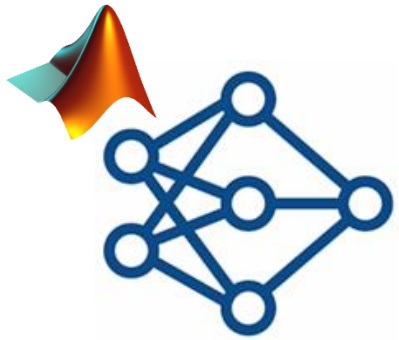

Deep Learning Networks

Coder Products

Desktop CPU

NVIDIA TensorRT & cuDNN Libraries

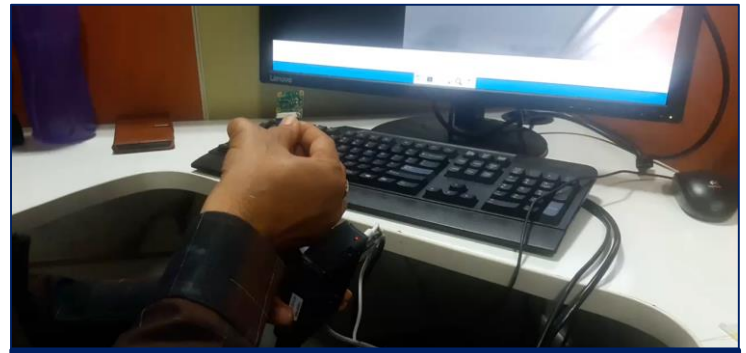

Raspberry Pi board

With GPU Coder, MATLAB is fast

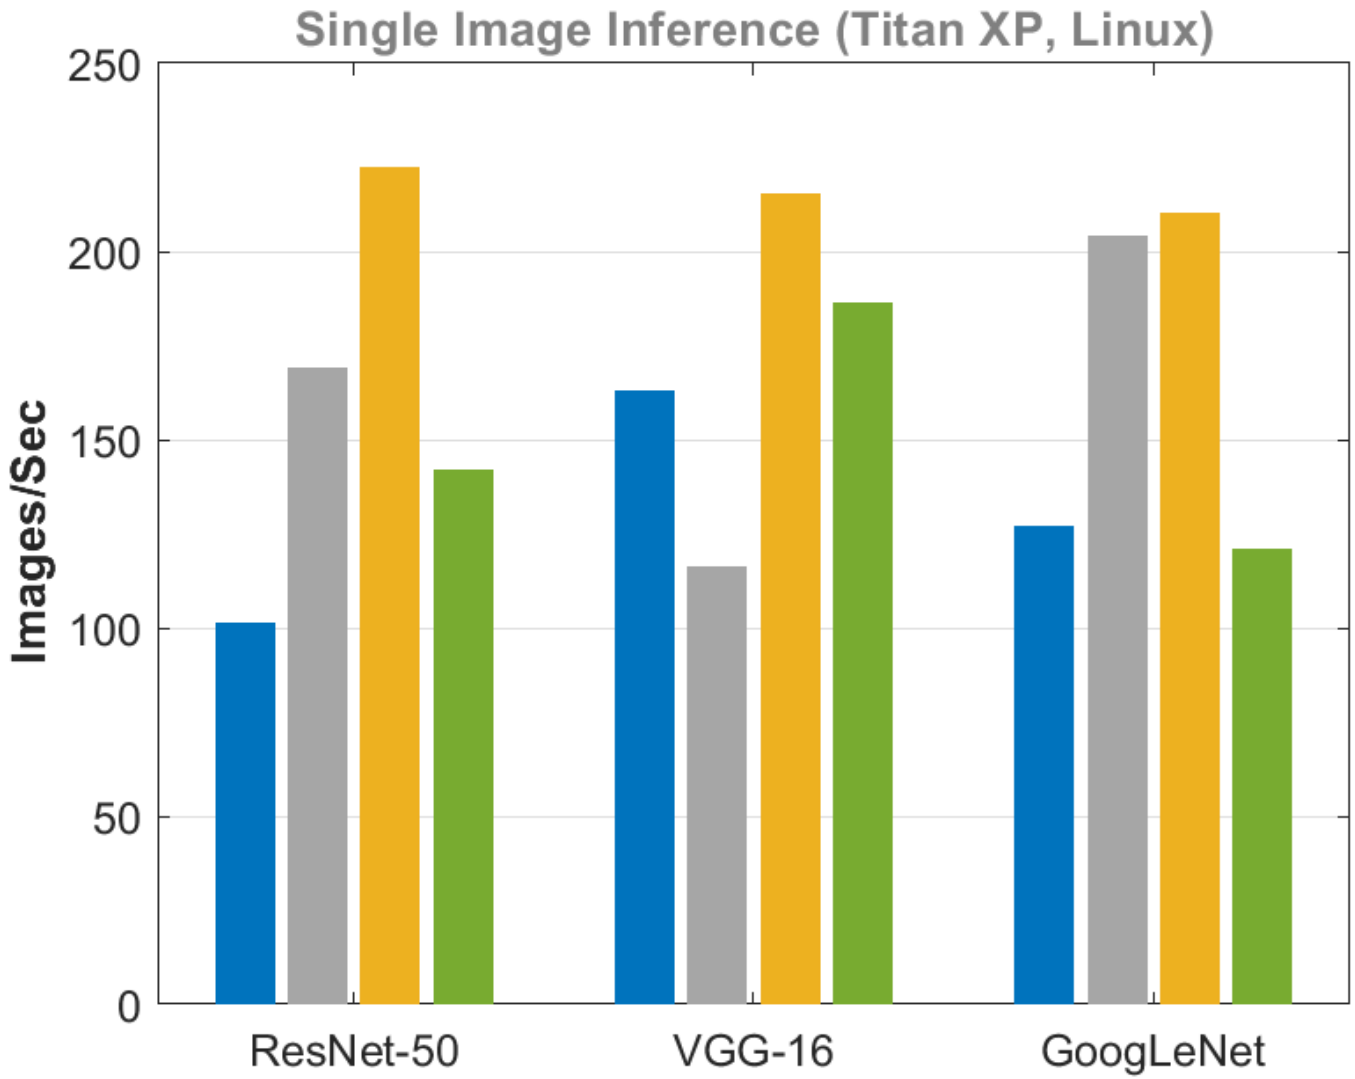

GPU Coder is faster than TensorFlow, MXNet and Pytorch

- TensorFlow
- MXNet
- GPU Coder
- PyTorch

Intel® Xeon® CPU 3.6 GHz - NVIDIA libraries: CUDA9 - cuDNN 7 - Frameworks: TensorFlow 1.8.0, MXNet 1.2.1, PyTorch 0.3.1

Deep Learning Workflow

Access and Explore Data

Files

Databases

Sensors

Label and Preprocess Data

Data Augmentation/ Transformation

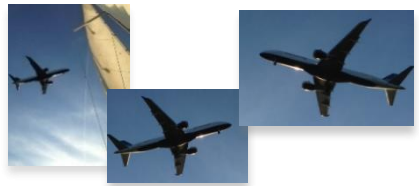

Labeling Automation

Import Reference Models

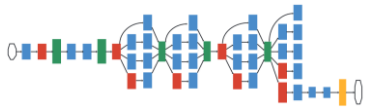

Develop Predictive Models

Hardware-Accelerated Training

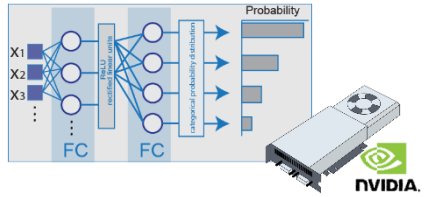

Hyperparameter Tuning

Network Visualization

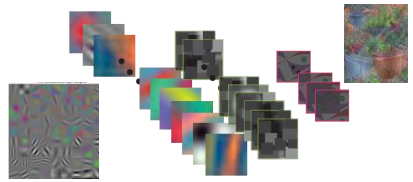

Integrate Models with Systems

Desktop Apps

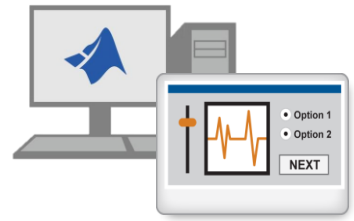

Enterprise Scale Systems

Embedded Devices and Hardware

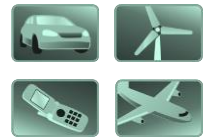

Why MATLAB?

- MATLAB is Productive
- MATLAB Supports the Entire Deep Learning Workflow
- 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

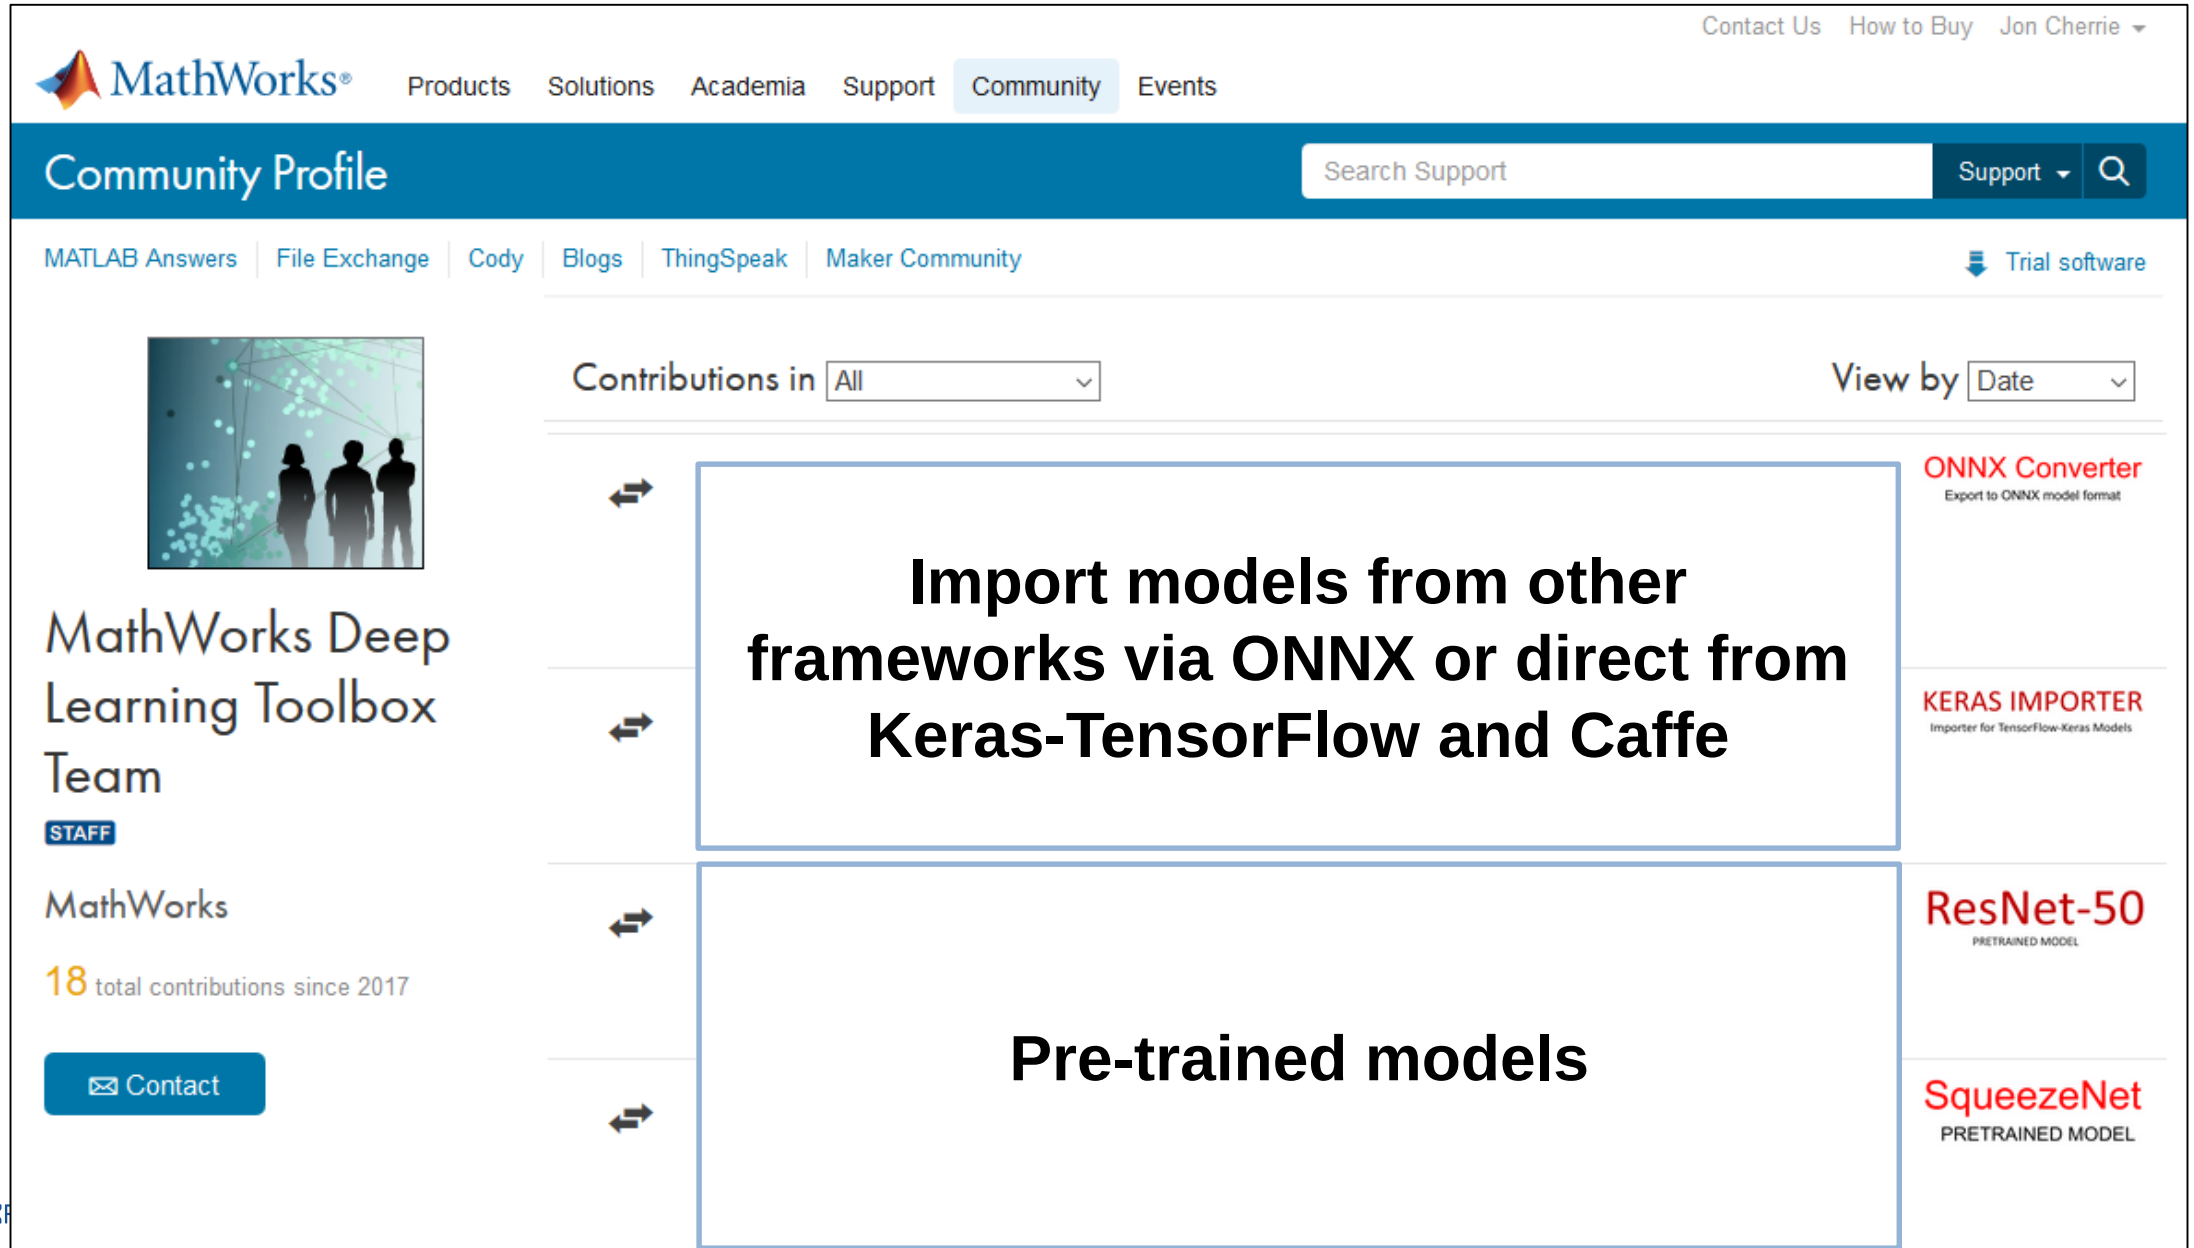

ONNX is an open format to represent deep learning models

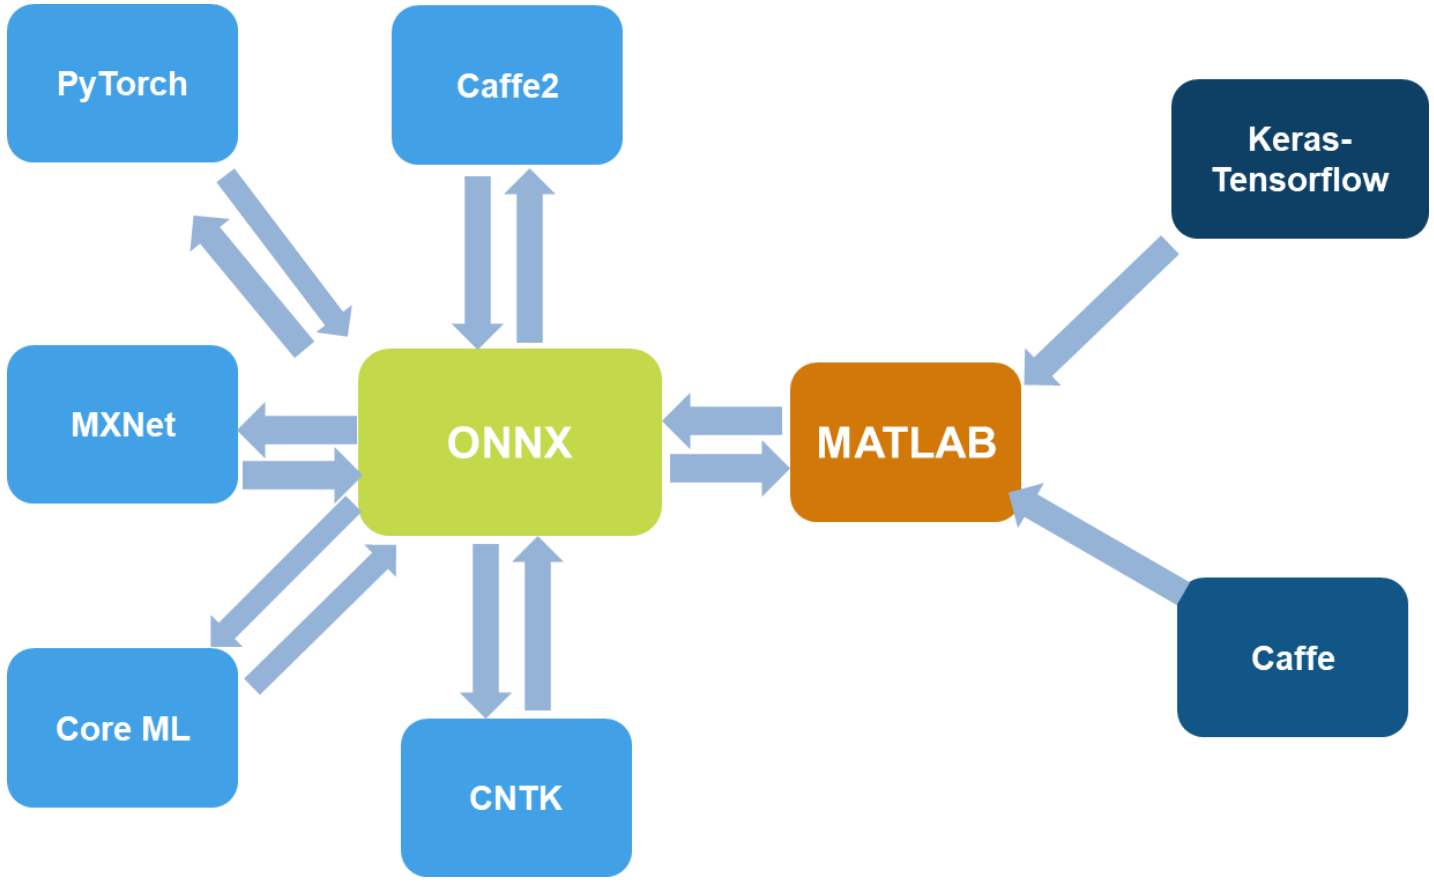

ONNX = Open Neural Network Exchange Format

ONNX Converter

Import from ONNX format

```
modelfile = 'cifarResNet.onnx';
classes = ["airplane" "automobile" "bir
net = importONNXNetwork(modelfile,'Outp
analyzeNetwork(net)
```

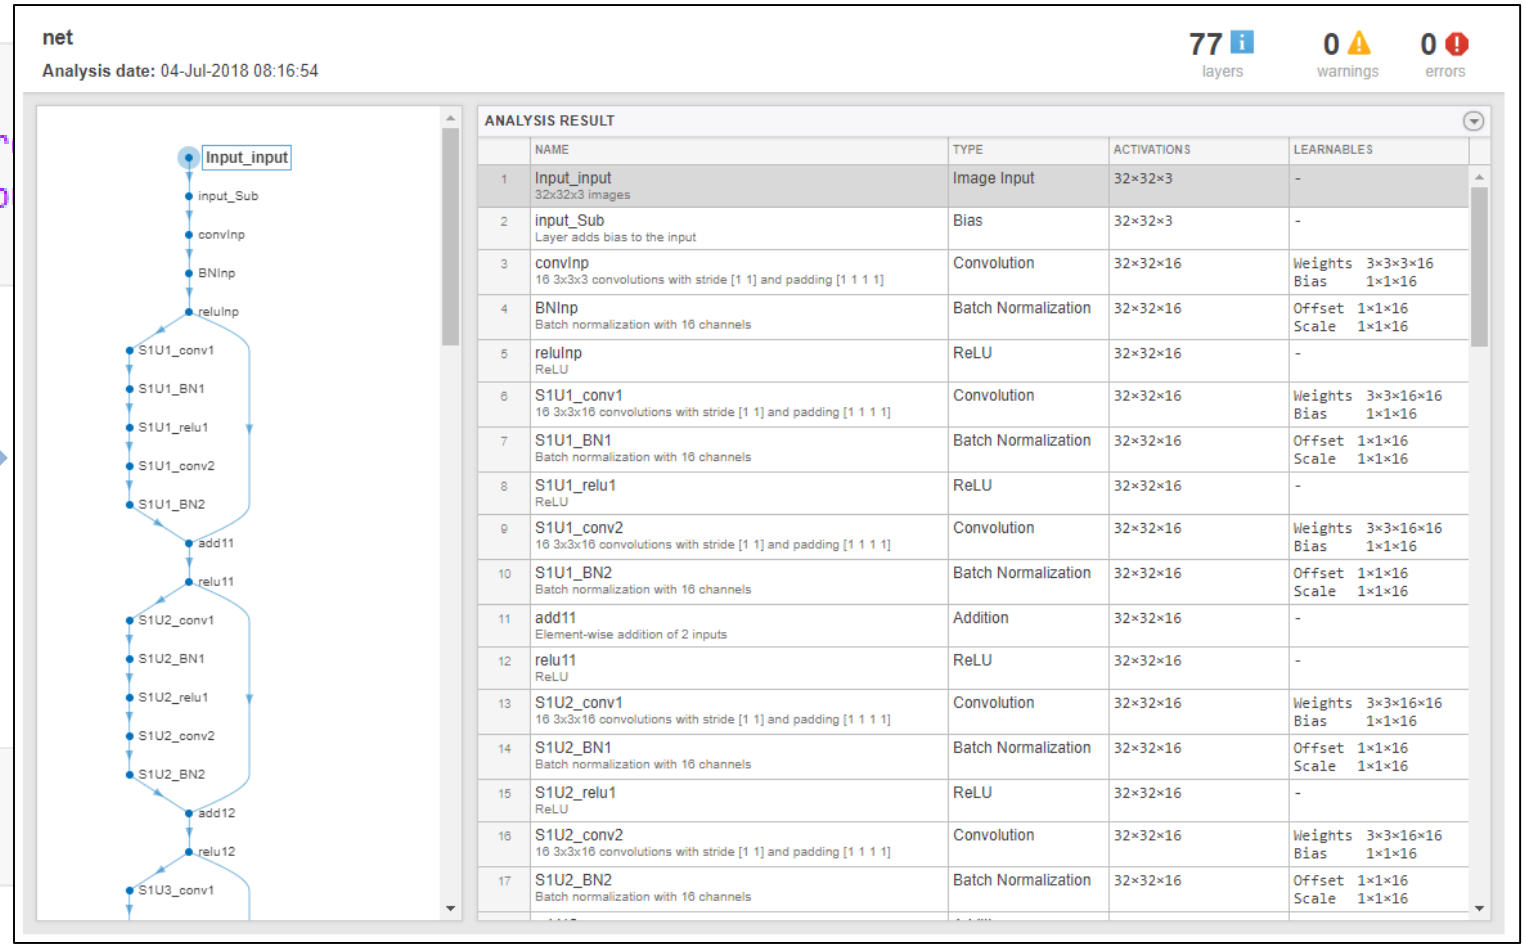

Export to ONNX format

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

MATLAB Integrates with Open Source Frameworks

- Use MATLAB for parts of workflow
- Access to latest models
- Improved collaboration with other users

Why MATLAB for Deep Learning?

- MATLAB is Productive
- MATLAB Supports Entire Deep Learning Workflow
- MATLAB Integrates with Open Source

Other Deep Learning Expo Events

- Deep Learning Workshops – 1pm, 2pm, 3:30pm
- Deep Learning Booth
- Deep Learning for Signals – 3:30pm