

Demystifying Deep Learning

“Let the computers do the hard work”


Phoebe Li, Terasoft Inc



Deep Learning Demo

Image Classification

Why MATLAB for Deep Learning?

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






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 tasks directly from data.



Data Types for Deep Learning



Signal



Text



Image


Deep Learning is **Versatile**



Detection of cars and road in autonomous driving systems



Rain Detection and Removal¹




Iris Recognition – 99.4% accuracy²



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




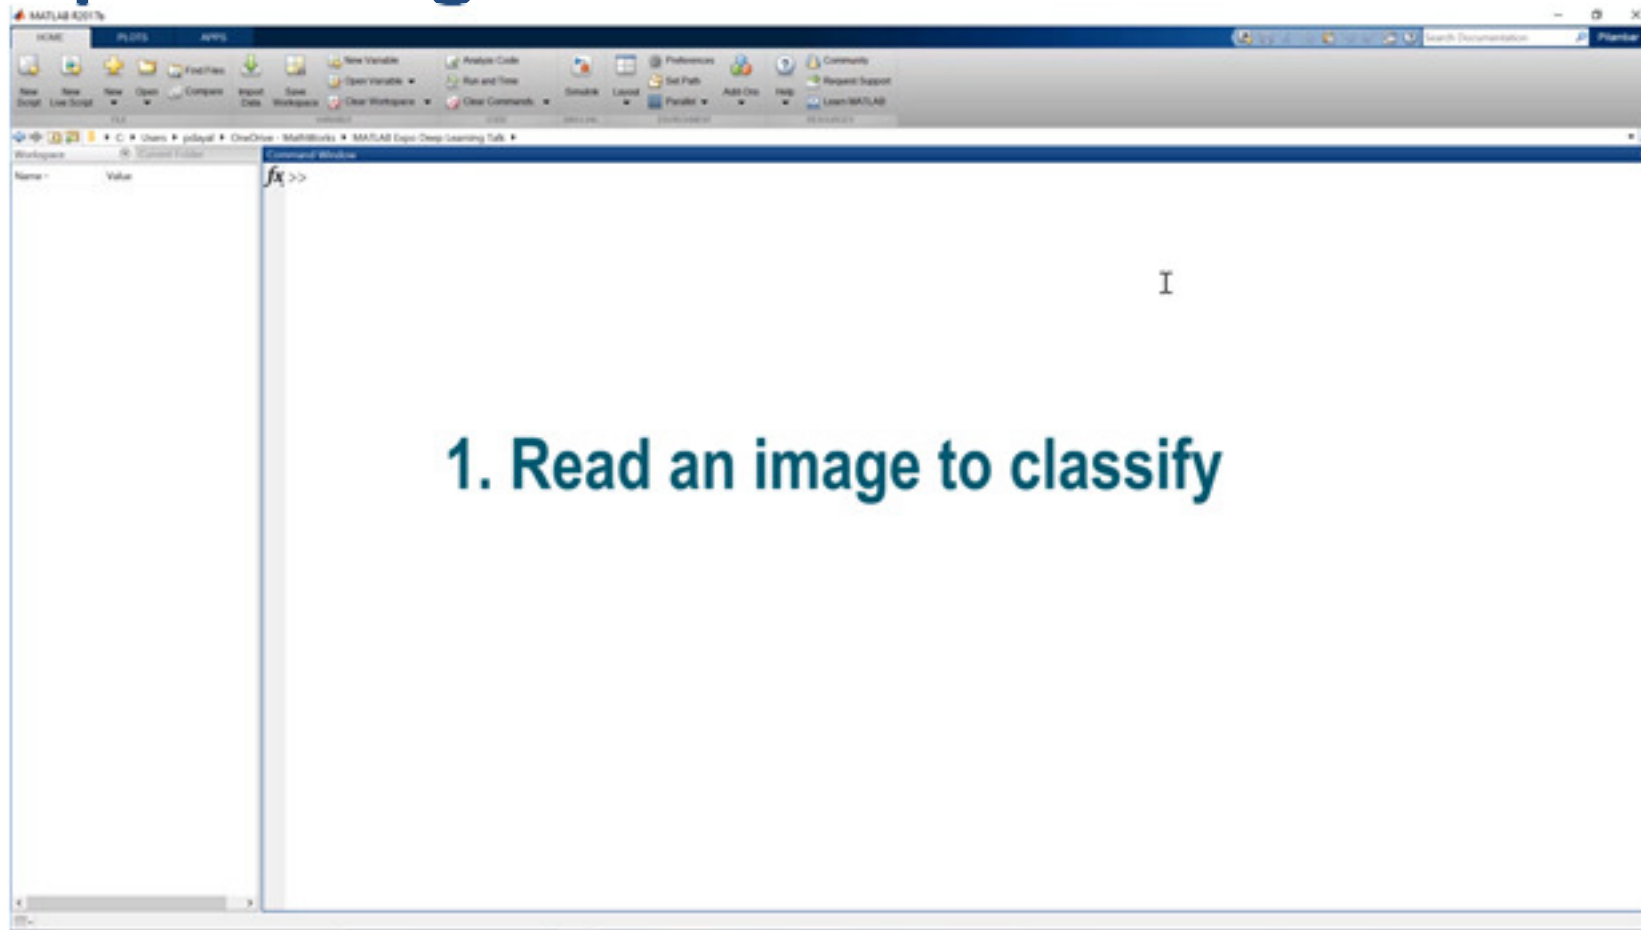
How is deep learning performing so well?



Deep Learning Uses a Neural Network Architecture




Deep Learning in 6 Lines of MATLAB Code




Why MATLAB for Deep Learning?

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



“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.



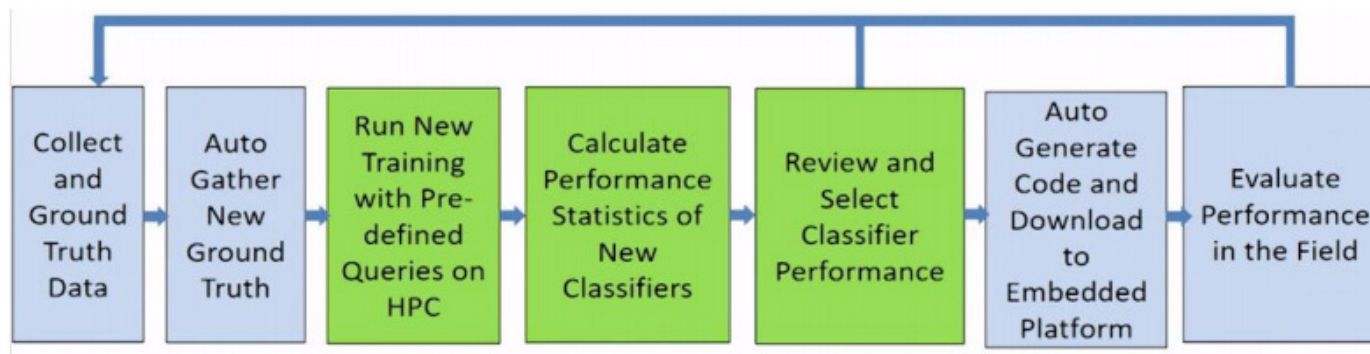
“We were spending way too much time ground-truthing [the data]”

--Larry Mianzo, Caterpillar


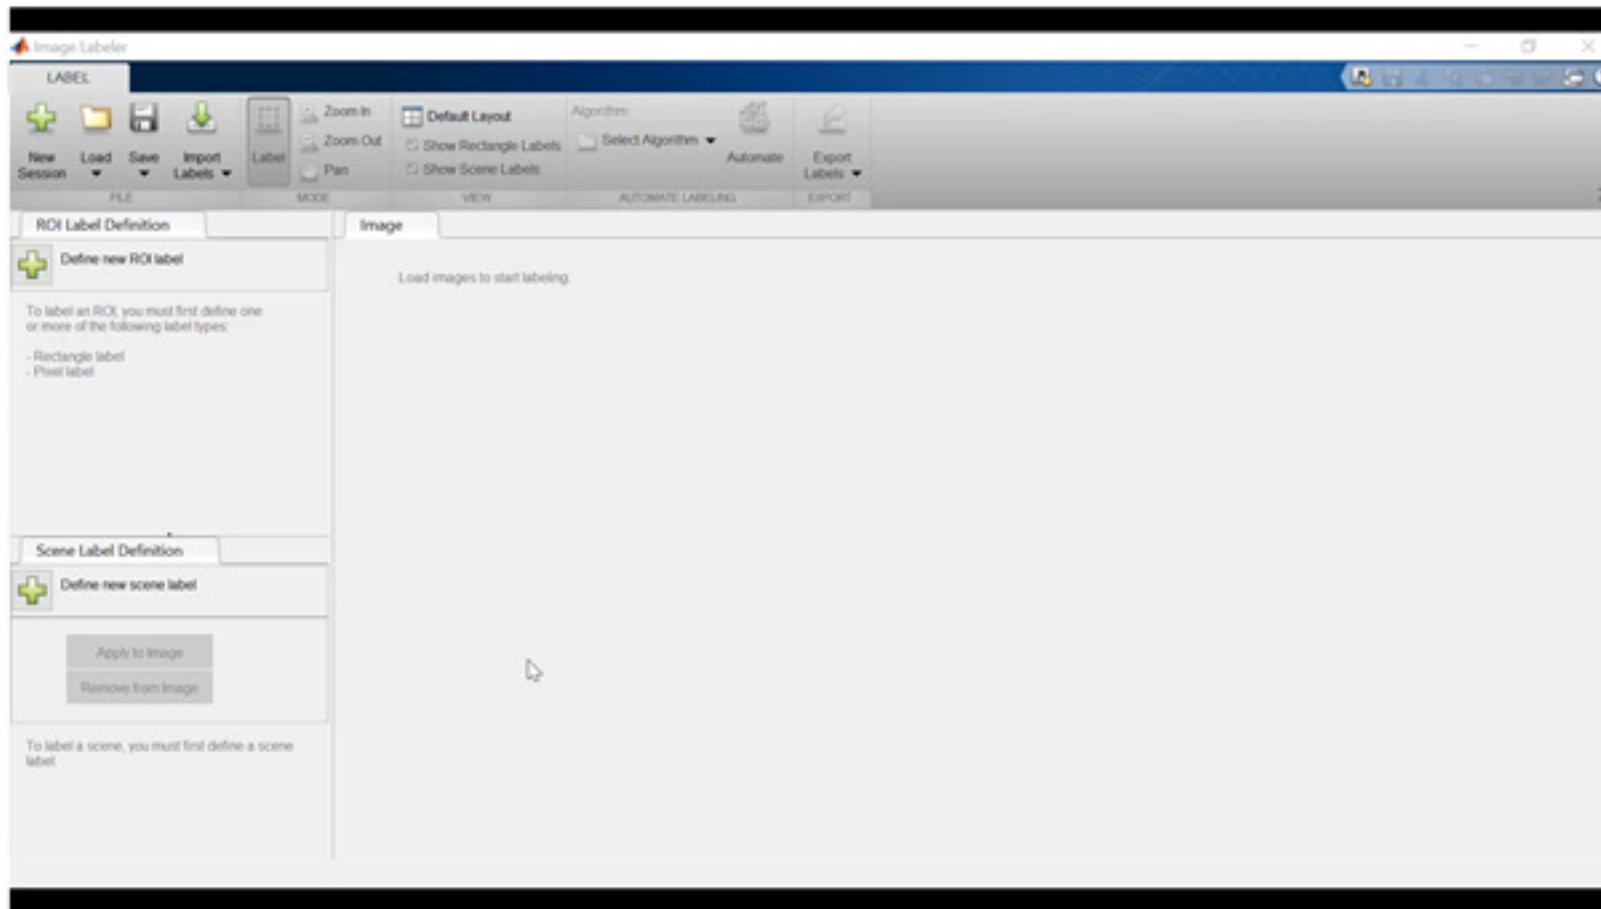



How Did Caterpillar Do with Our Tools?

- 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?






Regions with Convolutional Neural Network Features (R-CNN)


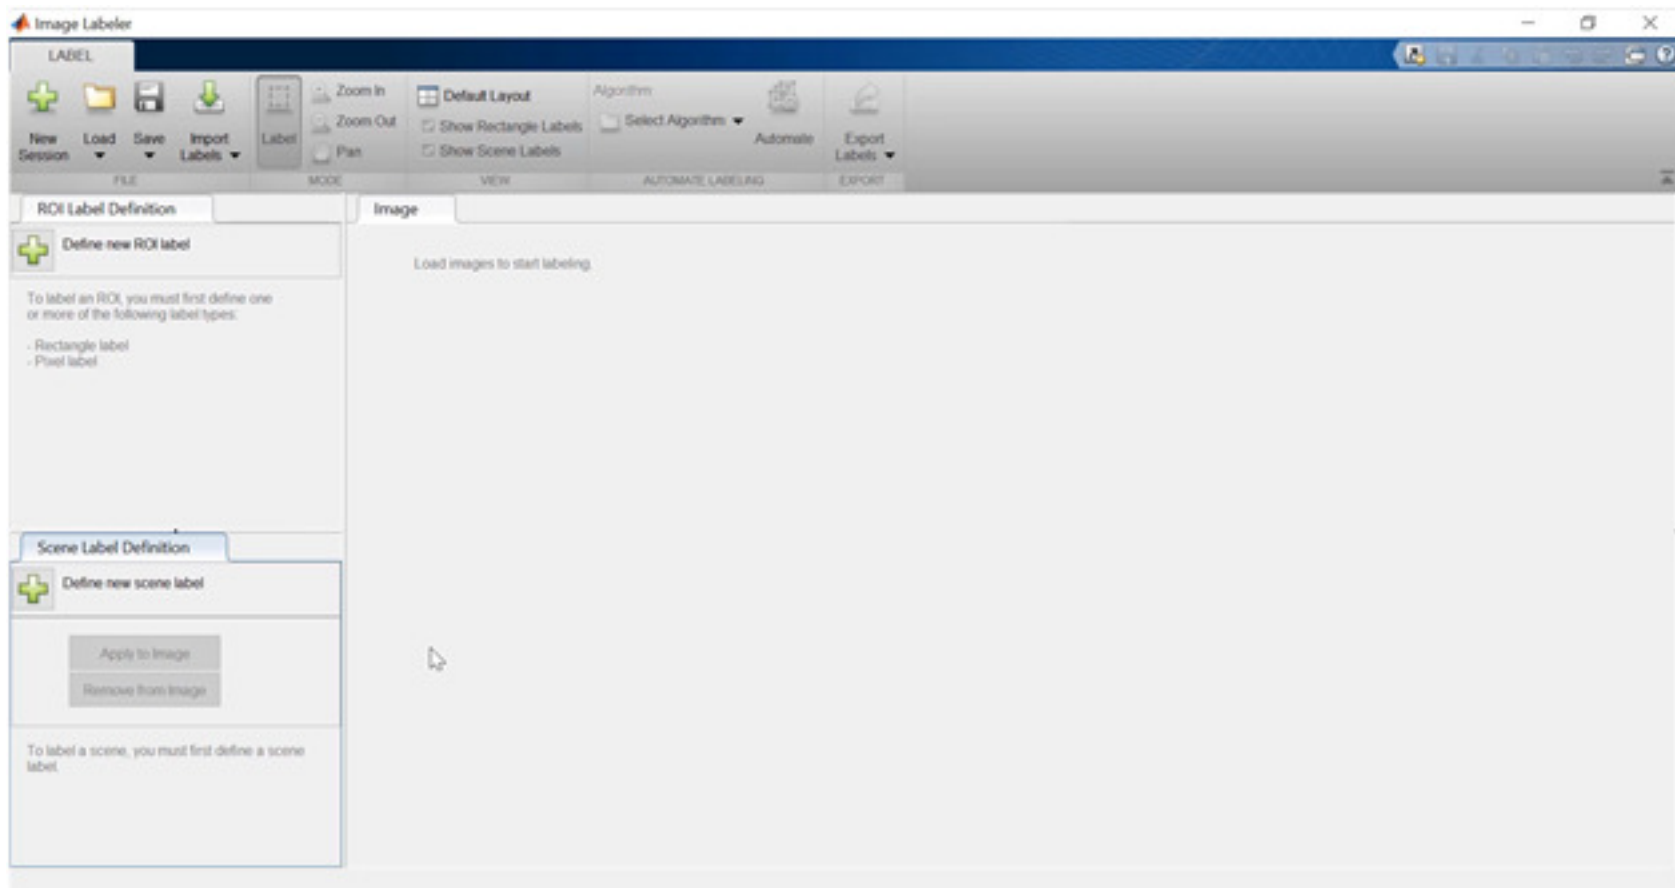


Labeling Videos: Ground Truth Labeler App

**Label
ground truth**




Labeling pixels






Semantic Segmentation using SegNet


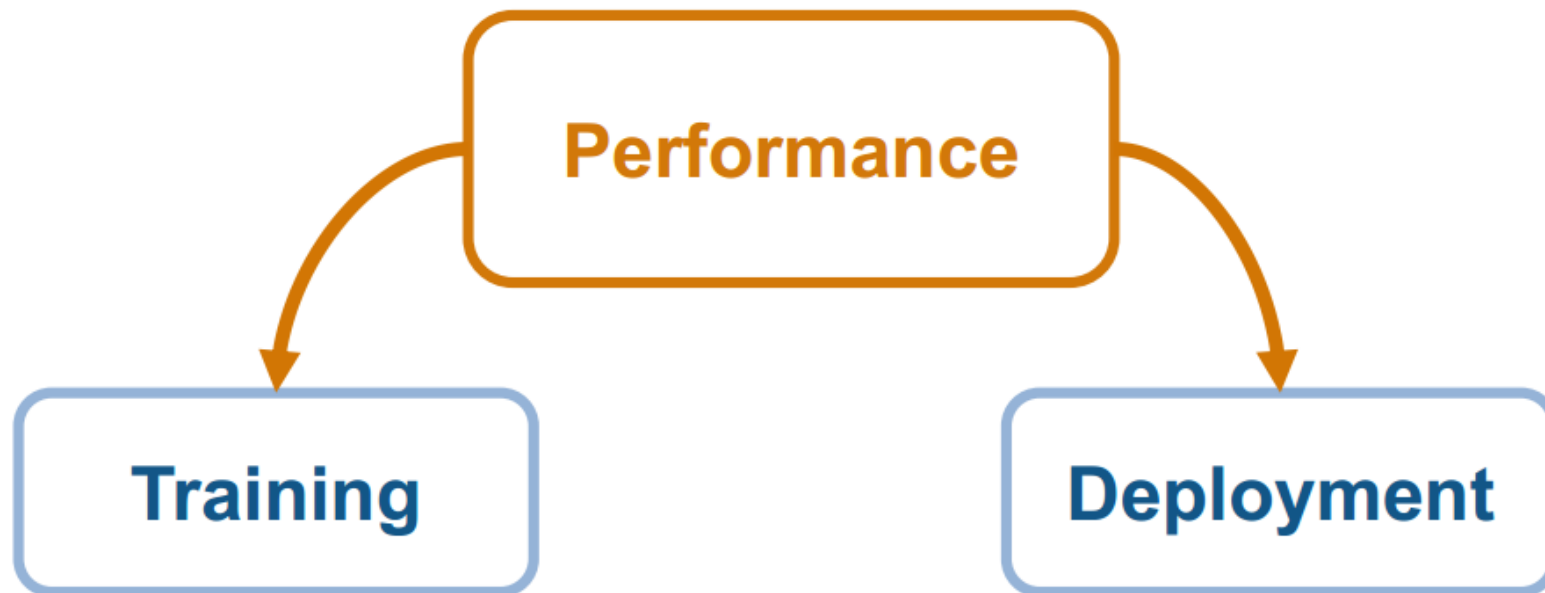


MATLAB is Productive

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







MATLAB is Fast





Speech Recognition Example

Audio signal \rightarrow Spectrogram \rightarrow Image Classification algorithm



Another Network for Signals - LSTM

- LSTM = Long Short Term Memory (Networks)
 - Signal, text, time-series data
 - Use previous data to predict new information
- I live in France. I speak _____.




1. Create Datastore


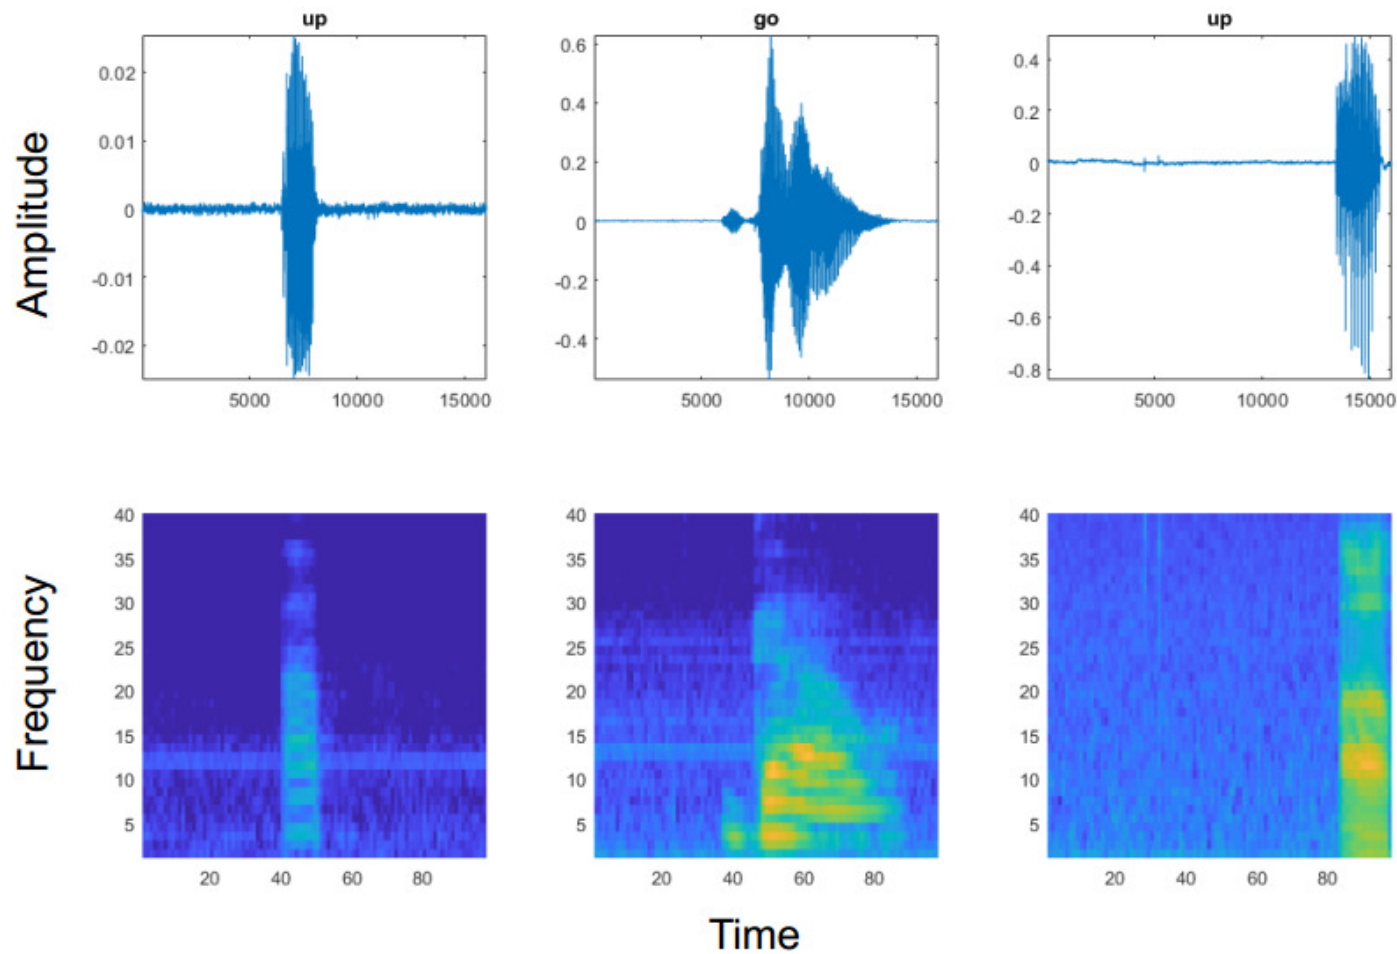
- Datastore creates reference for data
- Do not have to load in all objects into memory

<input type="checkbox"/> Name	Date modified
📁 _background_noise_	2/12/2018 9:32 AM
📁 Data	2/12/2018 9:39 AM
📁 go	2/12/2018 9:34 AM
📁 left	2/12/2018 9:35 AM
📁 no	2/12/2018 9:36 AM
📁 off	2/12/2018 9:37 AM
📁 on	2/12/2018 9:38 AM
📁 right	2/12/2018 9:31 AM
📁 up	2/12/2018 9:31 AM
📁 yes	2/12/2018 9:32 AM

```
datafolder = fullfile(tempdir, 'speech_commands_v0.01');  
  
addpath(fullfile(matlabroot, 'toolbox', 'audio', 'audiodemons'))  
ads = audioexample.Datastore(datafolder, ...  
    'IncludeSubfolders', true, ...  
    'FileExtensions', '.wav', ...  
    'LabelSource', 'foldernames', ...  
    'ReadMethod', 'File')
```



2. Compute Speech Spectrograms



3. Split datastores

Training

70%



- Trains the model
- Computer “learns” from this data

Validation

15%




- Checks accuracy of model during training

Test

15%



- Tests model accuracy
- Not used until validation accuracy is good



4. Define Architecture and Parameters


```
layers = [  
    imageInputLayer(imageSize)  
  
    convolution2dLayer(3,16,'Padding','same')  
    batchNormalizationLayer  
    reluLayer  
  
    maxPooling2dLayer(2,'Stride',2)  
  
    convolution2dLayer(3,32,'Padding','same')  
    batchNormalizationLayer  
    reluLayer  
  
    maxPooling2dLayer(2,'Stride',2,'Padding',[0,1])  
  
    dropoutLayer(dropoutProb)  
    convolution2dLayer(3,64,'Padding','same')  
    batchNormalizationLayer  
    reluLayer  
  
    dropoutLayer(dropoutProb)
```

```
convolution2dLayer(3,64,'Padding','same')  
batchNormalizationLayer  
reluLayer  
  
maxPooling2dLayer(2,'Stride',2,'Padding',[0,1])  
  
dropoutLayer(dropoutProb)  
convolution2dLayer(3,64,'Padding','same')  
batchNormalizationLayer  
reluLayer  
  
dropoutLayer(dropoutProb)  
convolution2dLayer(3,64,'Padding','same')  
batchNormalizationLayer  
reluLayer  
  
maxPooling2dLayer([1 13])  
  
fullyConnectedLayer(numClasses)  
softmaxLayer  
weightedCrossEntropyLayer(classNames,classWeights)];
```

```
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);
```

Training Parameters

Neural Network Architecture



4. How to choose the right structure to start from?

Pretrained Models for Transfer learning

- Access pretrained models from top researchers with a single line of code

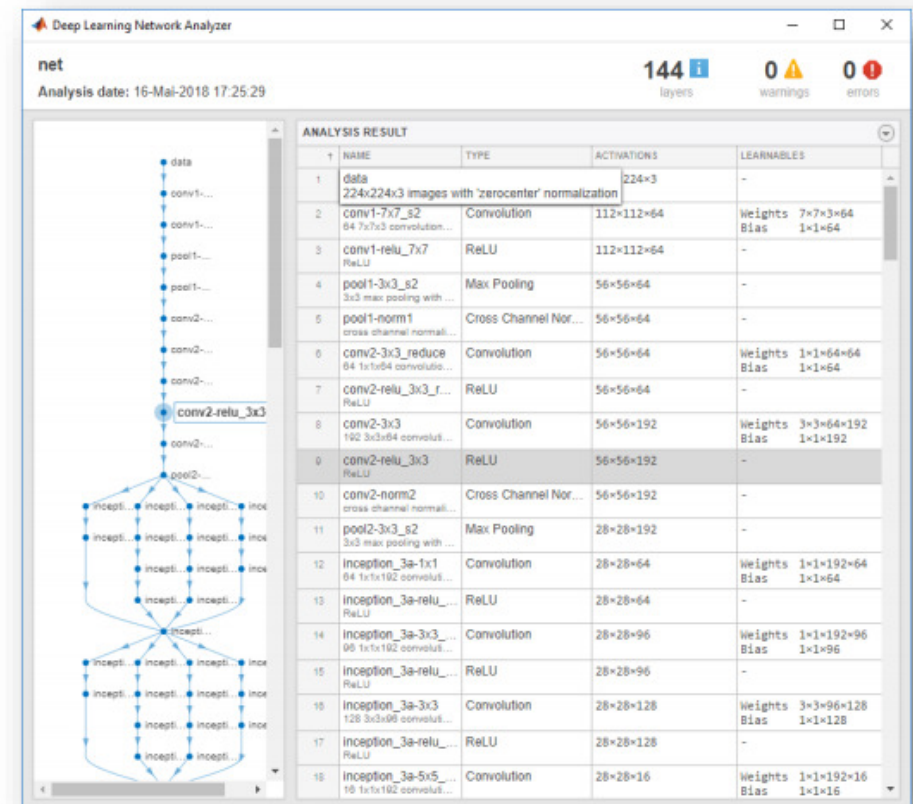
```
net = alexnet;  
net = vgg16;  
net = inceptionv3;  
net = googlenet;  
...
```

Import Models from OSS Frameworks

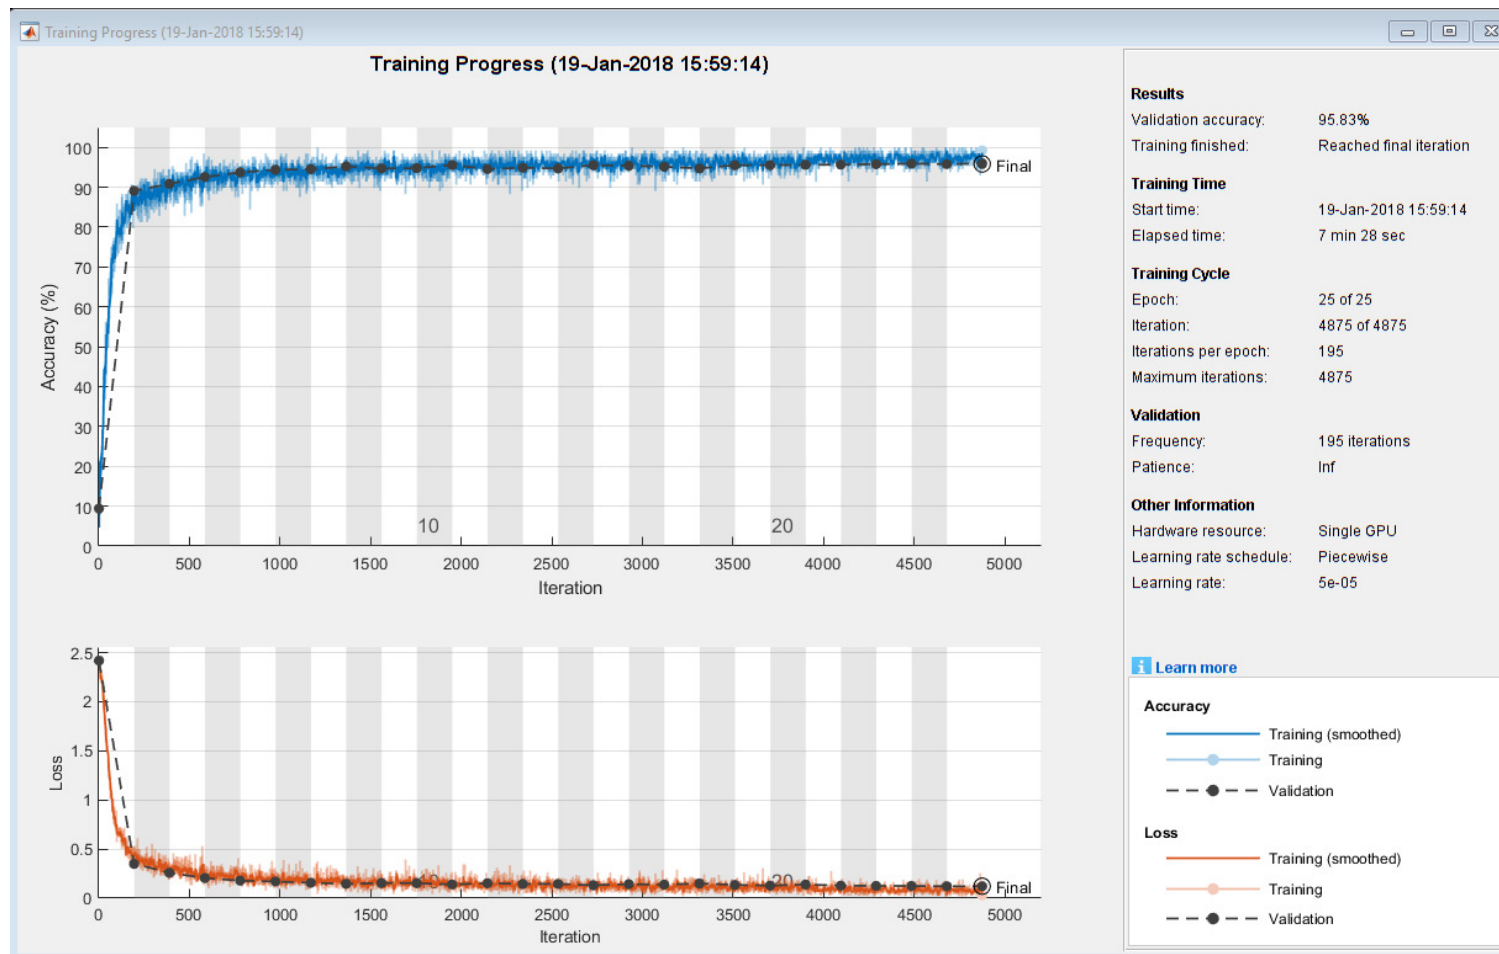
- Caffe Model Importer
- TensorFlow/Keras Model Importer

Start from scratch

- Verify structure with [Network Analyzer App](#)



5. Train Network



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)

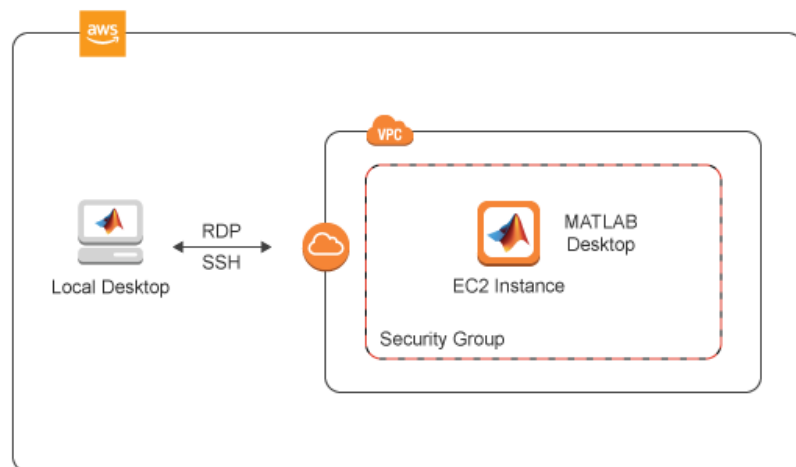
```
opts = trainingOptions('sgdm', ...  
    'MaxEpochs', 100, ...  
    'MiniBatchSize', 250, ...  
    'InitialLearnRate', 0.00005, ...  
    'ExecutionEnvironment', 'auto' );
```

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

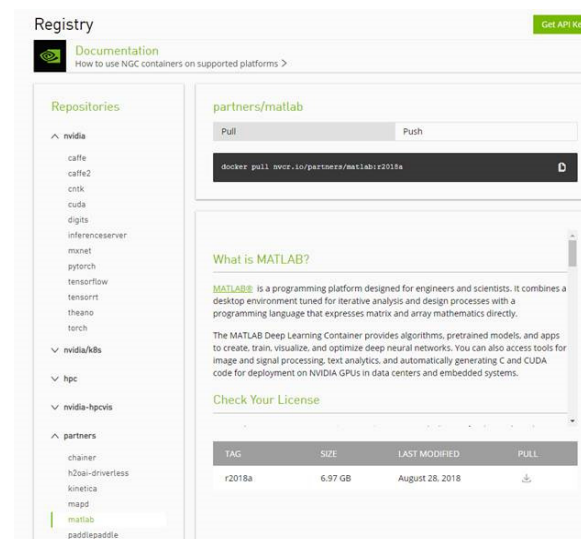
```
opts = trainingOptions('sgdm', ...  
    'MaxEpochs', 100, ...  
    'MiniBatchSize', 250, ...  
    'InitialLearnRate', 0.00005, ...  
    'ExecutionEnvironment', 'parallel' );
```



R2018b: Easier to Get Started on Cloud and HPC



Reference Architectures




Containers for Nvidia GPU Cloud (NGC).
Also supports NVIDIA DGX

Why is this important?

1. Customers want to get started in cloud instead of buying hardware (Reference Architectures)
2. Customers want to use NVIDIA DGX super computer (NGC, DGX) to accelerate training

Training Performance




TensorFlow

MATLAB


MXNet

Batch size 32



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




With GPU Coder, MATLAB is fast




Faster than TensorFlow,
MXNet, and 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




GPU Coder Performance

Inference with MATLAB



Why MATLAB?

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




Used MATLAB and Open Source Together




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


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




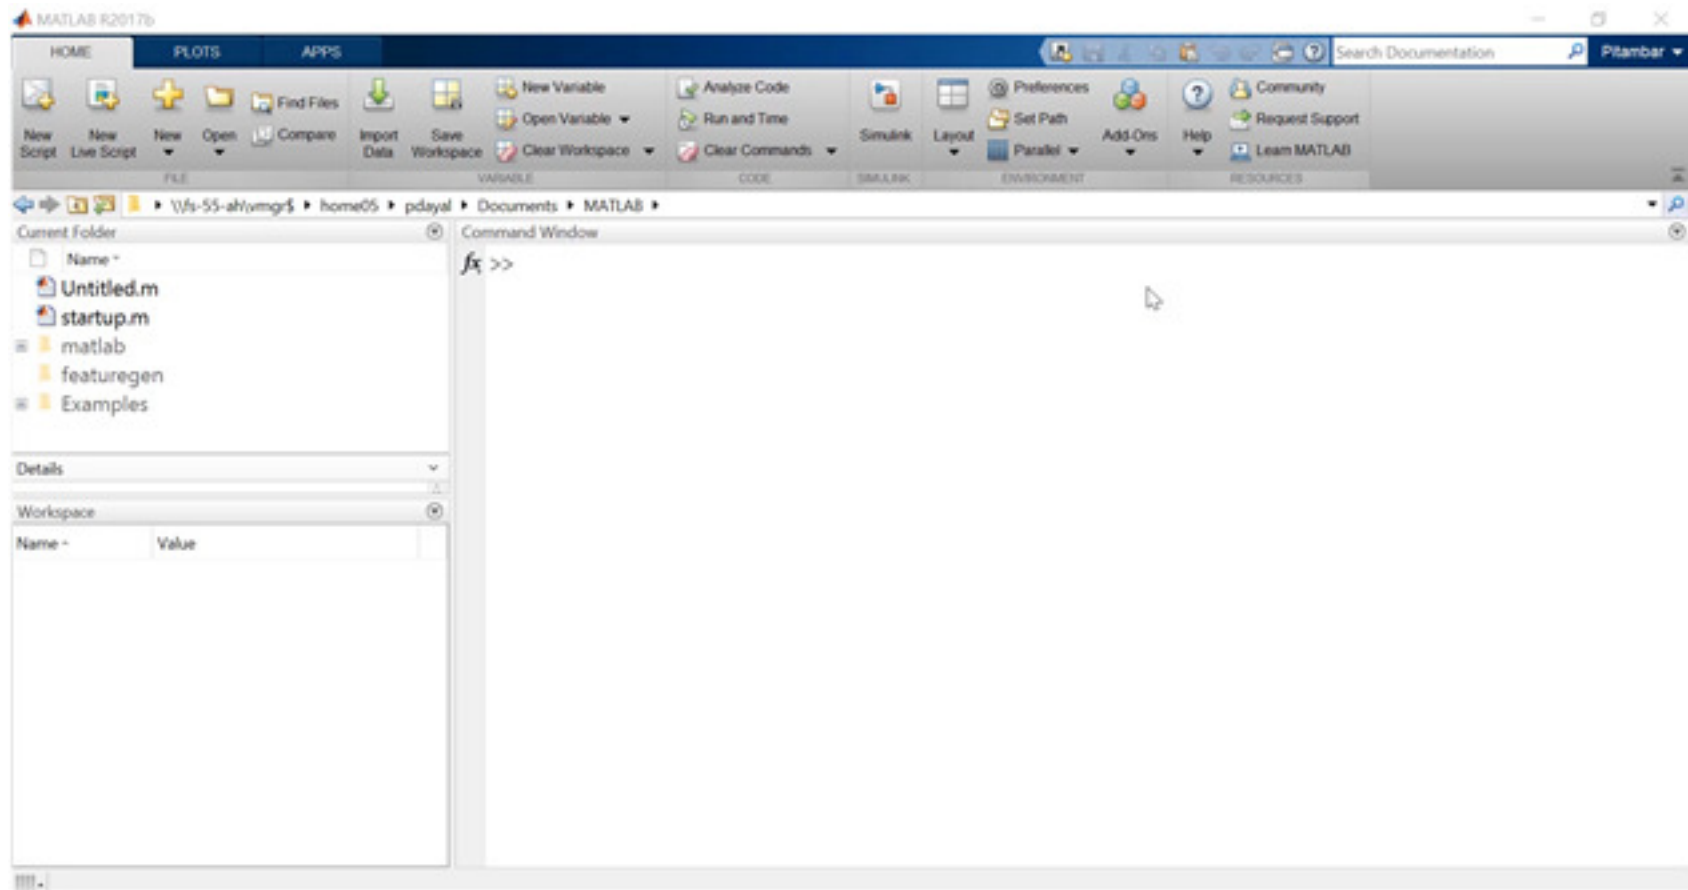
MATLAB Integrates with Open Source Frameworks



- ONNX: Input and export to several frameworks
- Model importers
 - Caffe
 - TensorFlow-Keras




Keras-Tensorflow Importer




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



MATLAB Integrates with Open Source Frameworks

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



MATLAB 课程

https://matlabacademy.mathworks.com/cn

MATLAB 课程

搜索 MathWorks.com

概述 | 课程提供 | 时间表与报名 | 自定义进度课程 | 在您的机构中提供培训 | 认证项目 | 更多

免费学习 MATLAB

实践课和演示

MATLAB 入门之旅

查看我的课程



快速入门



免费

MATLAB 入门之旅 (中文)

MATLAB 基础知识快速入门

启动 | 详细信息



免费

深度学习入门之旅 (英语)

使用深度学习方法快速入门，以执行图像识别

启动 | 详细信息

