Deep Learning for Computer Vision Applications Using MATLAB

Elza John
Training Engineer

Arun Mathamkode
Application Support Engineer

Agenda

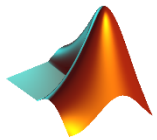

- What is deep learning?- Overview
- Why deep learning?
- Pretrained networks and Network layers
- Building your own network
- Transfer learning with CNNs

Traditional Machine Learning vs Deep Learning

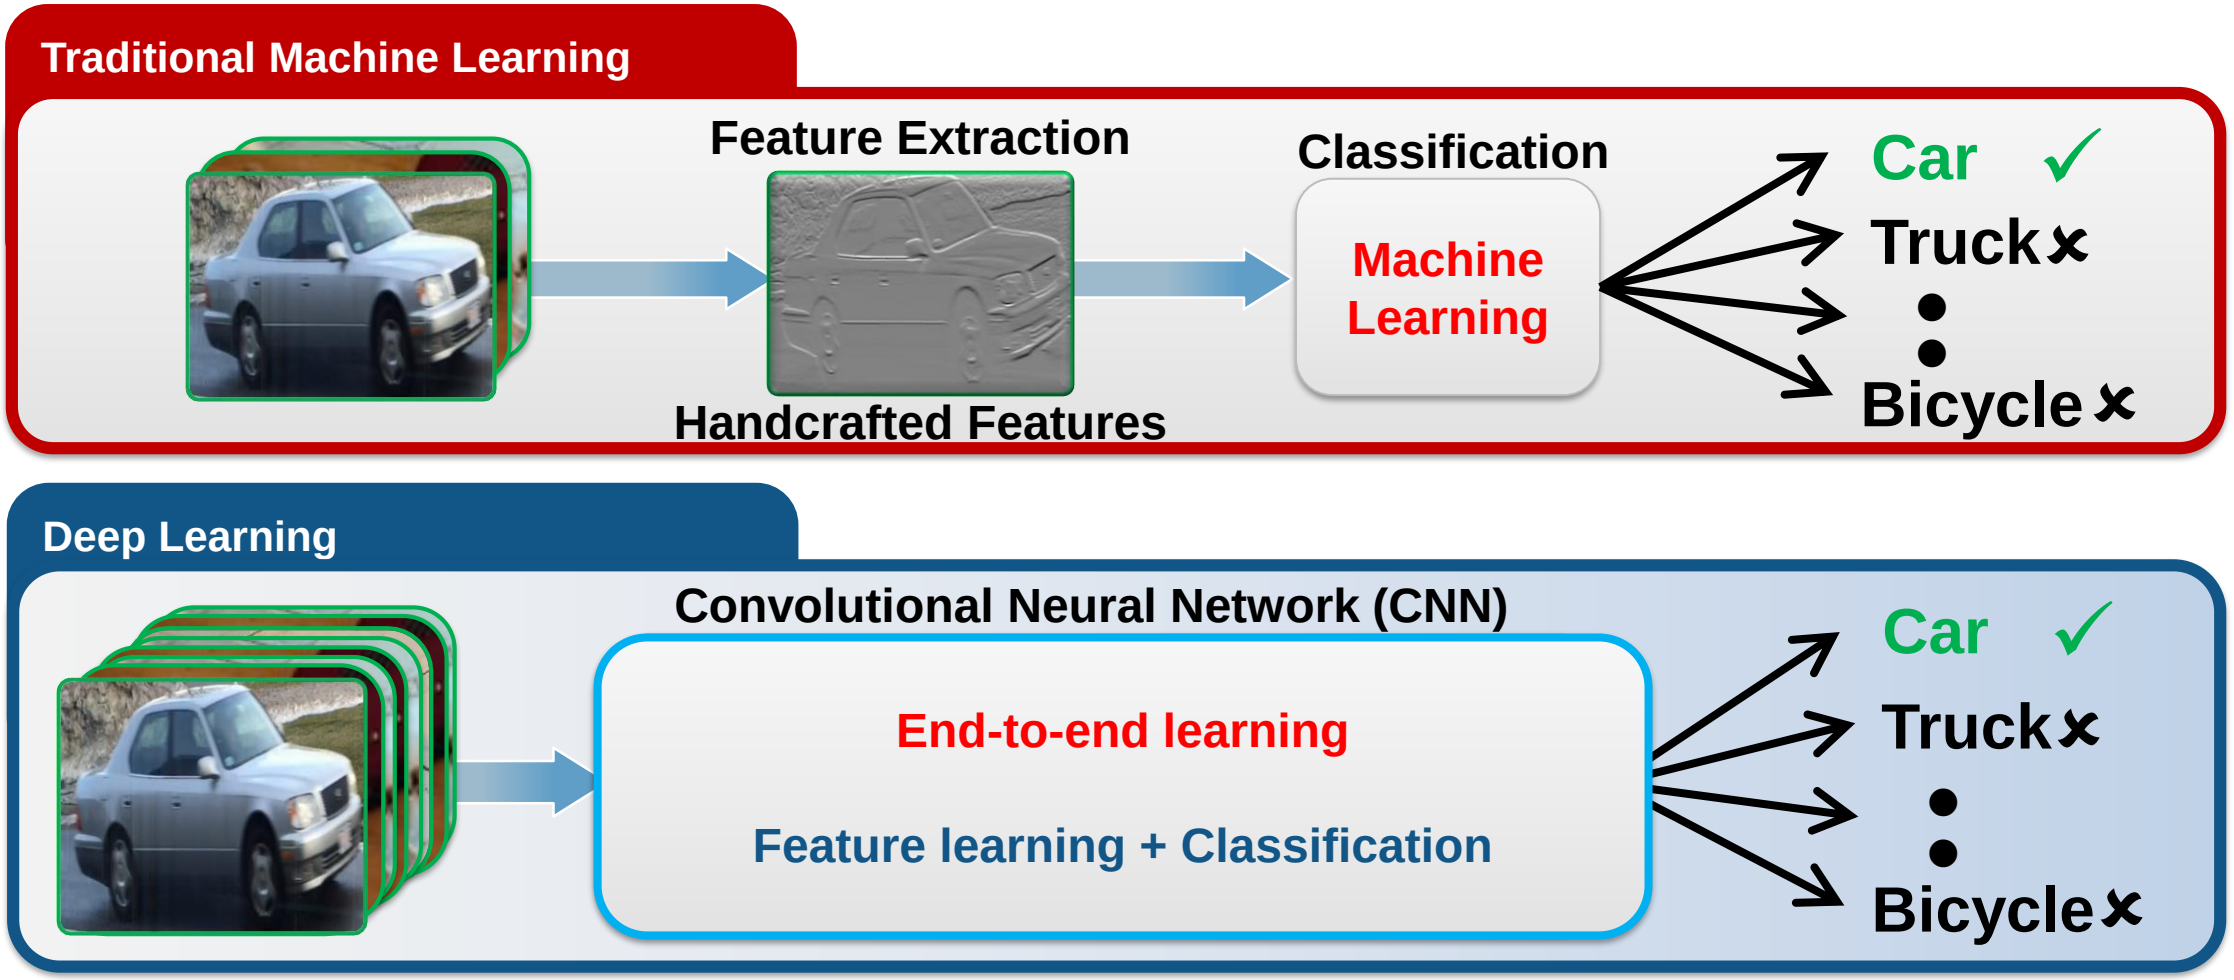

Deep learning performs **end-end learning** by learning **features, representations and tasks** directly from **images, text and sound**

Deep learning Applications

Semantic Segmentation

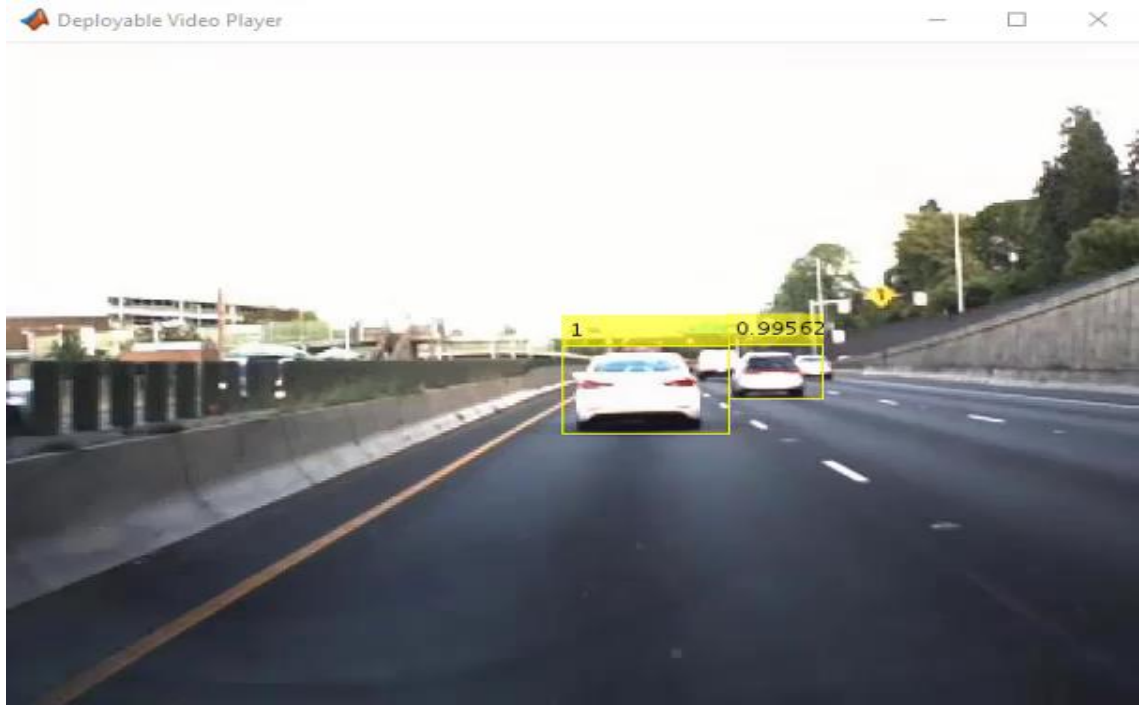

Vehicle Detection

Agenda

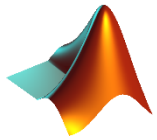

What is deep learning?- Overview

Why deep learning?

Pretrained networks and Network layers

Building your own network

Transfer learning with CNNs

Deep learning models can surpass human accuracy.

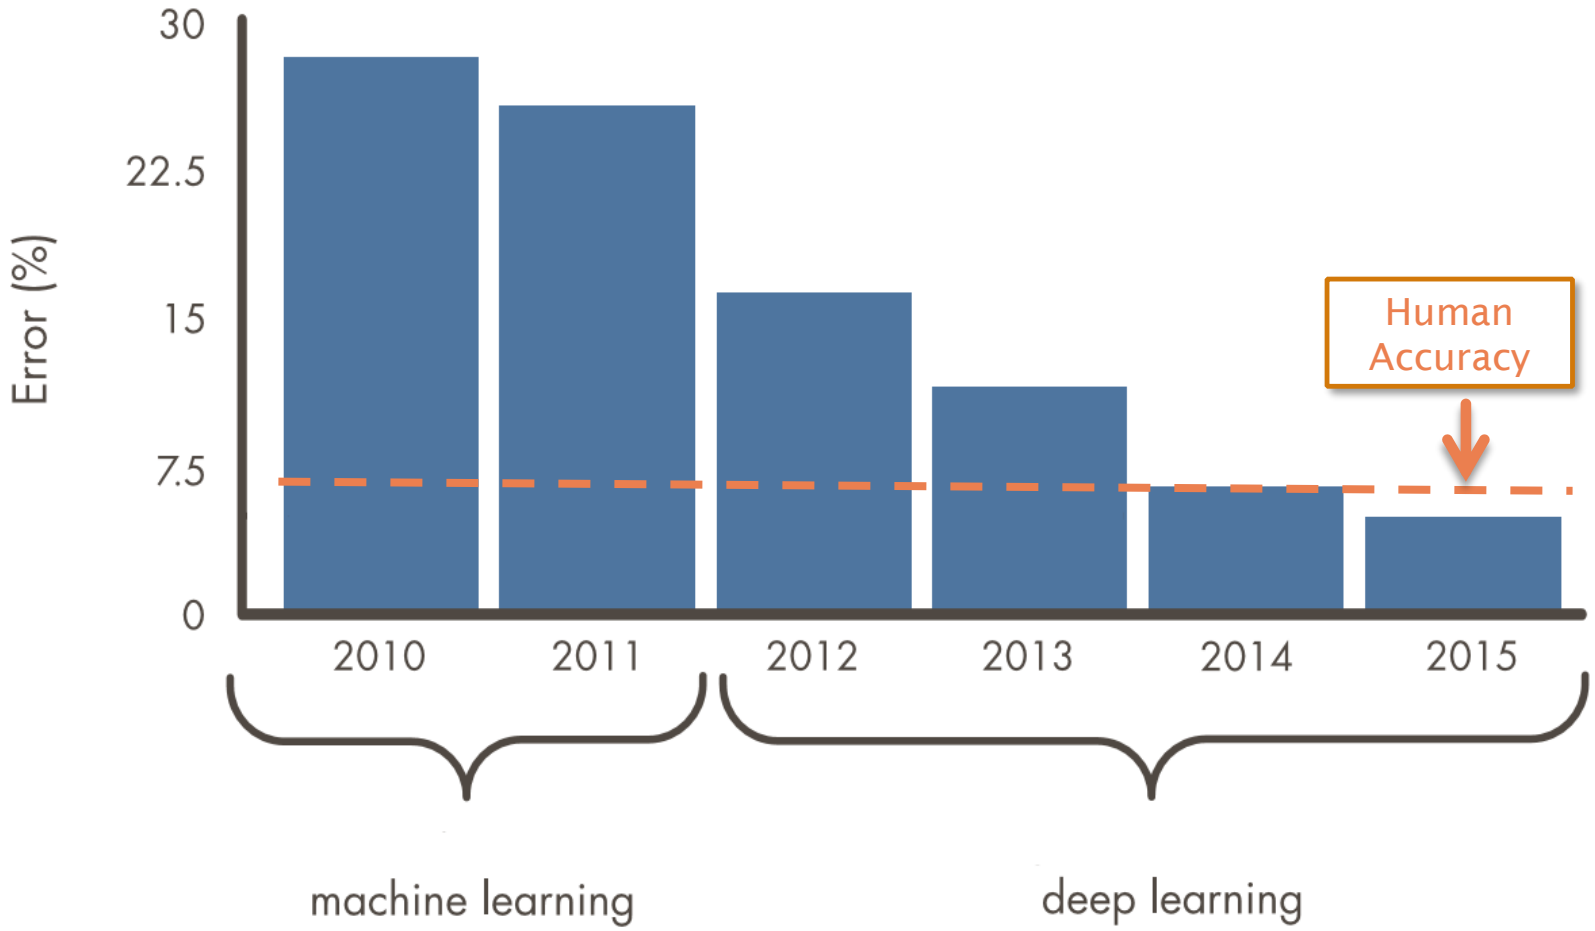

Source: ILSVRC Top-5 Error on ImageNet

Deep Learning Enablers

Labeled public datasets

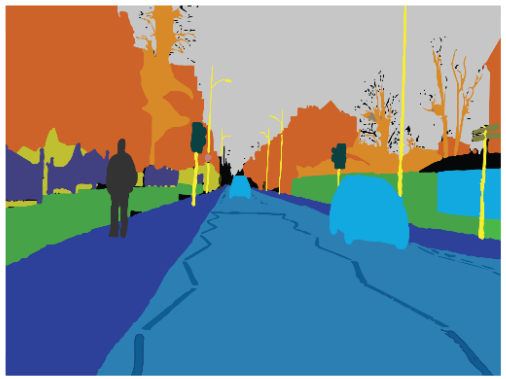

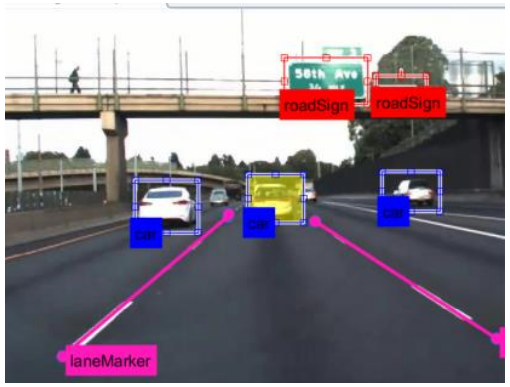

Increased GPU acceleration

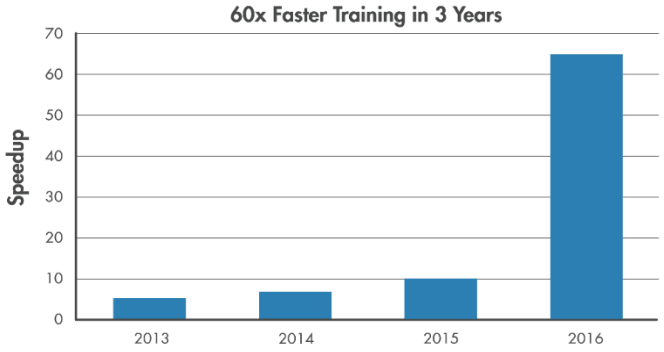

World-class models

AlexNet
PRETRAINED MODEL

VGG-16
PRETRAINED MODEL

ResNet
PRETRAINED MODEL

Caffe
MODELS

GoogLeNet
PRETRAINED MODEL

TensorFlow/Keras
MODELS

Let’s try it out!

Exercise: SS_DeepLearningIn5Lines.mlx

Agenda

What is deep learning?- Overview

Why deep learning?

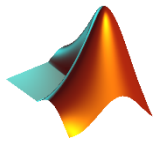

Pretrained networks and Network layers

Building your own network

Transfer learning with CNNs

Pre-Trained Models and Network Architectures

Pretrained Models

- AlexNet `net = alexnet;`
- VGG-16 `net = vgg16;`
- VGG-19 `net = vgg19;`
- GoogLeNet `net = googlenet;`
- Resnet50 `net = resnet50;`
- Resnet101 `net = resnet101`

Download from within MATLAB

Import Models from Frameworks

- Caffe Model Importer (including Caffe Model Zoo)
  - `importCaffeLayers`
  - `importCaffeNetwork`
- TensorFlow-Keras Model Importer
  - `importKerasLayers`
  - `importKerasNetwork`

Deep Learning Uses a Neural Network Architecture

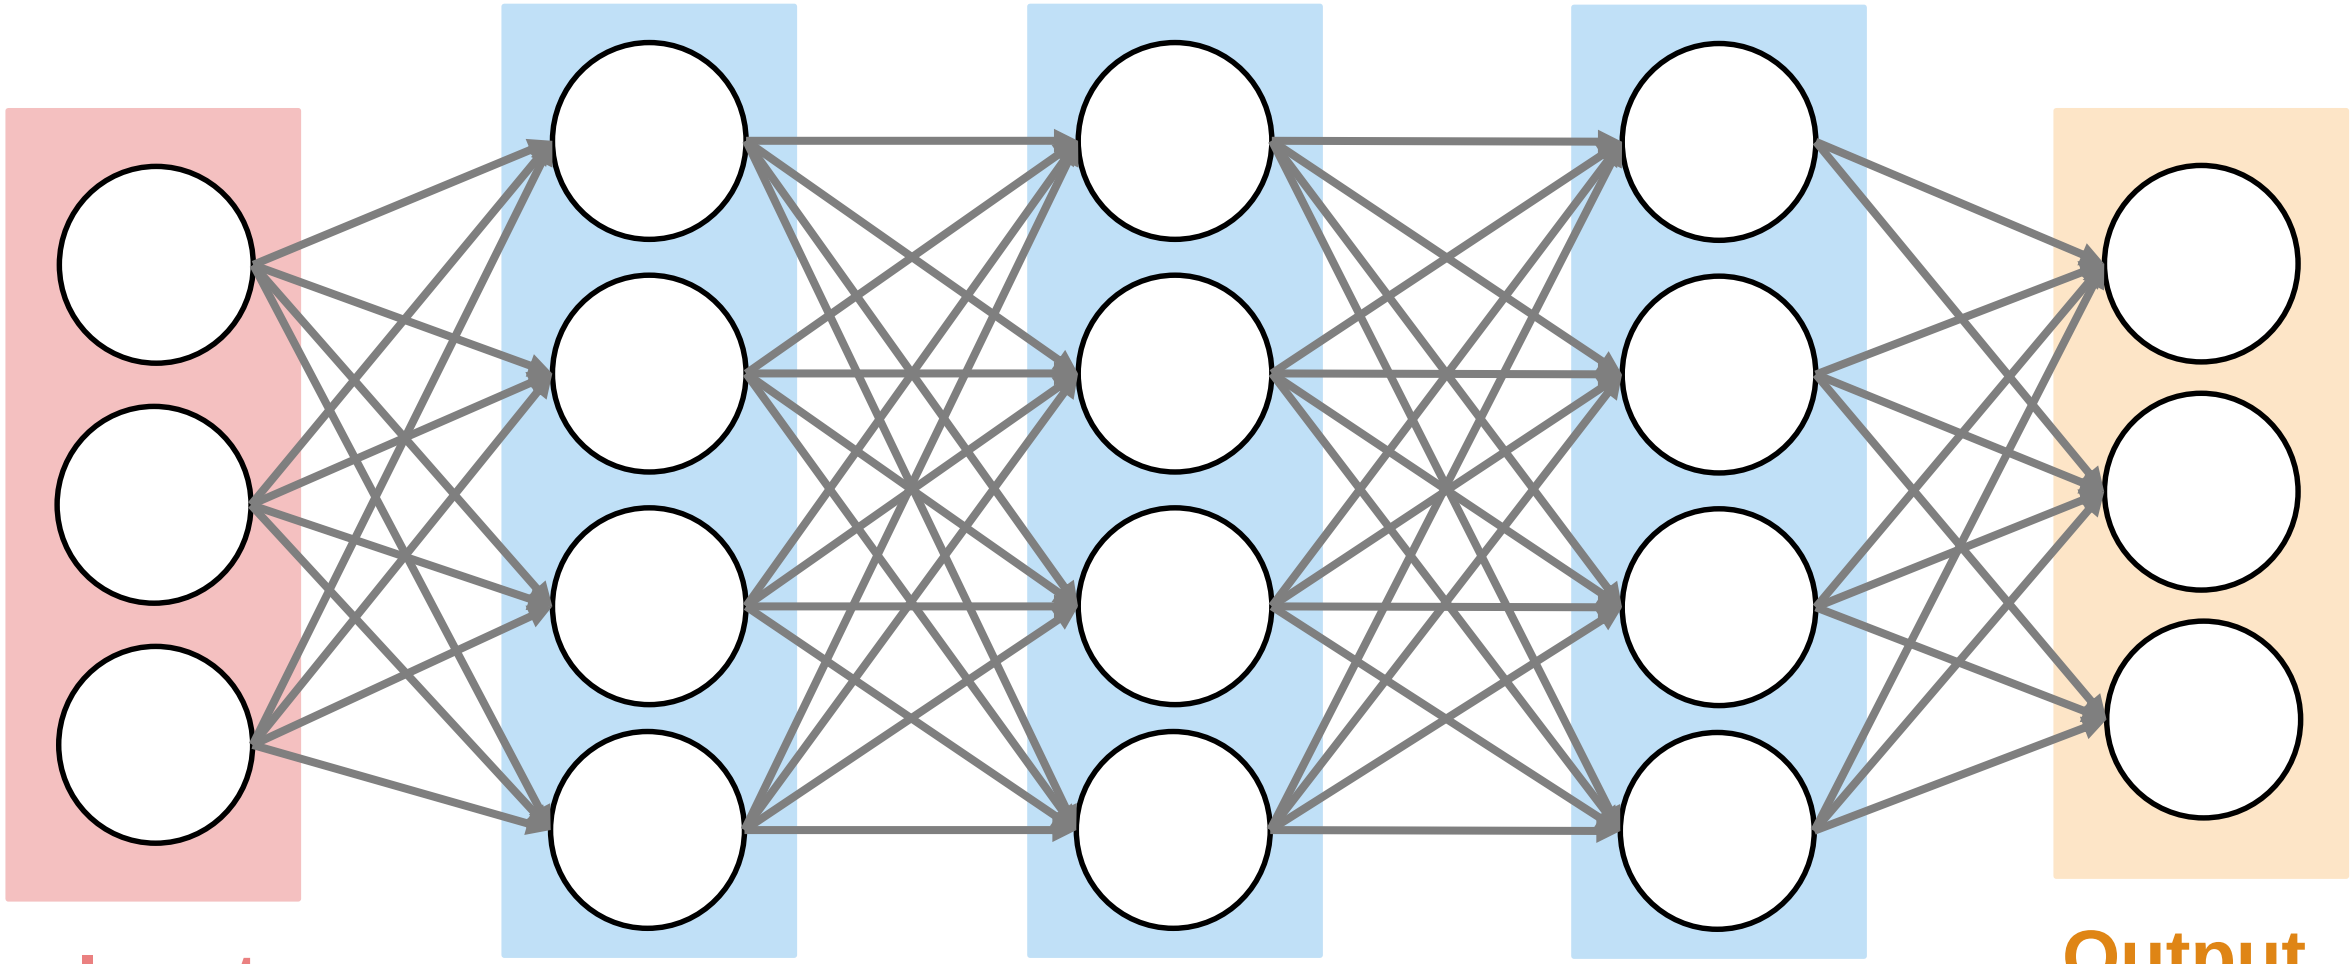

Convolutional Neural Network

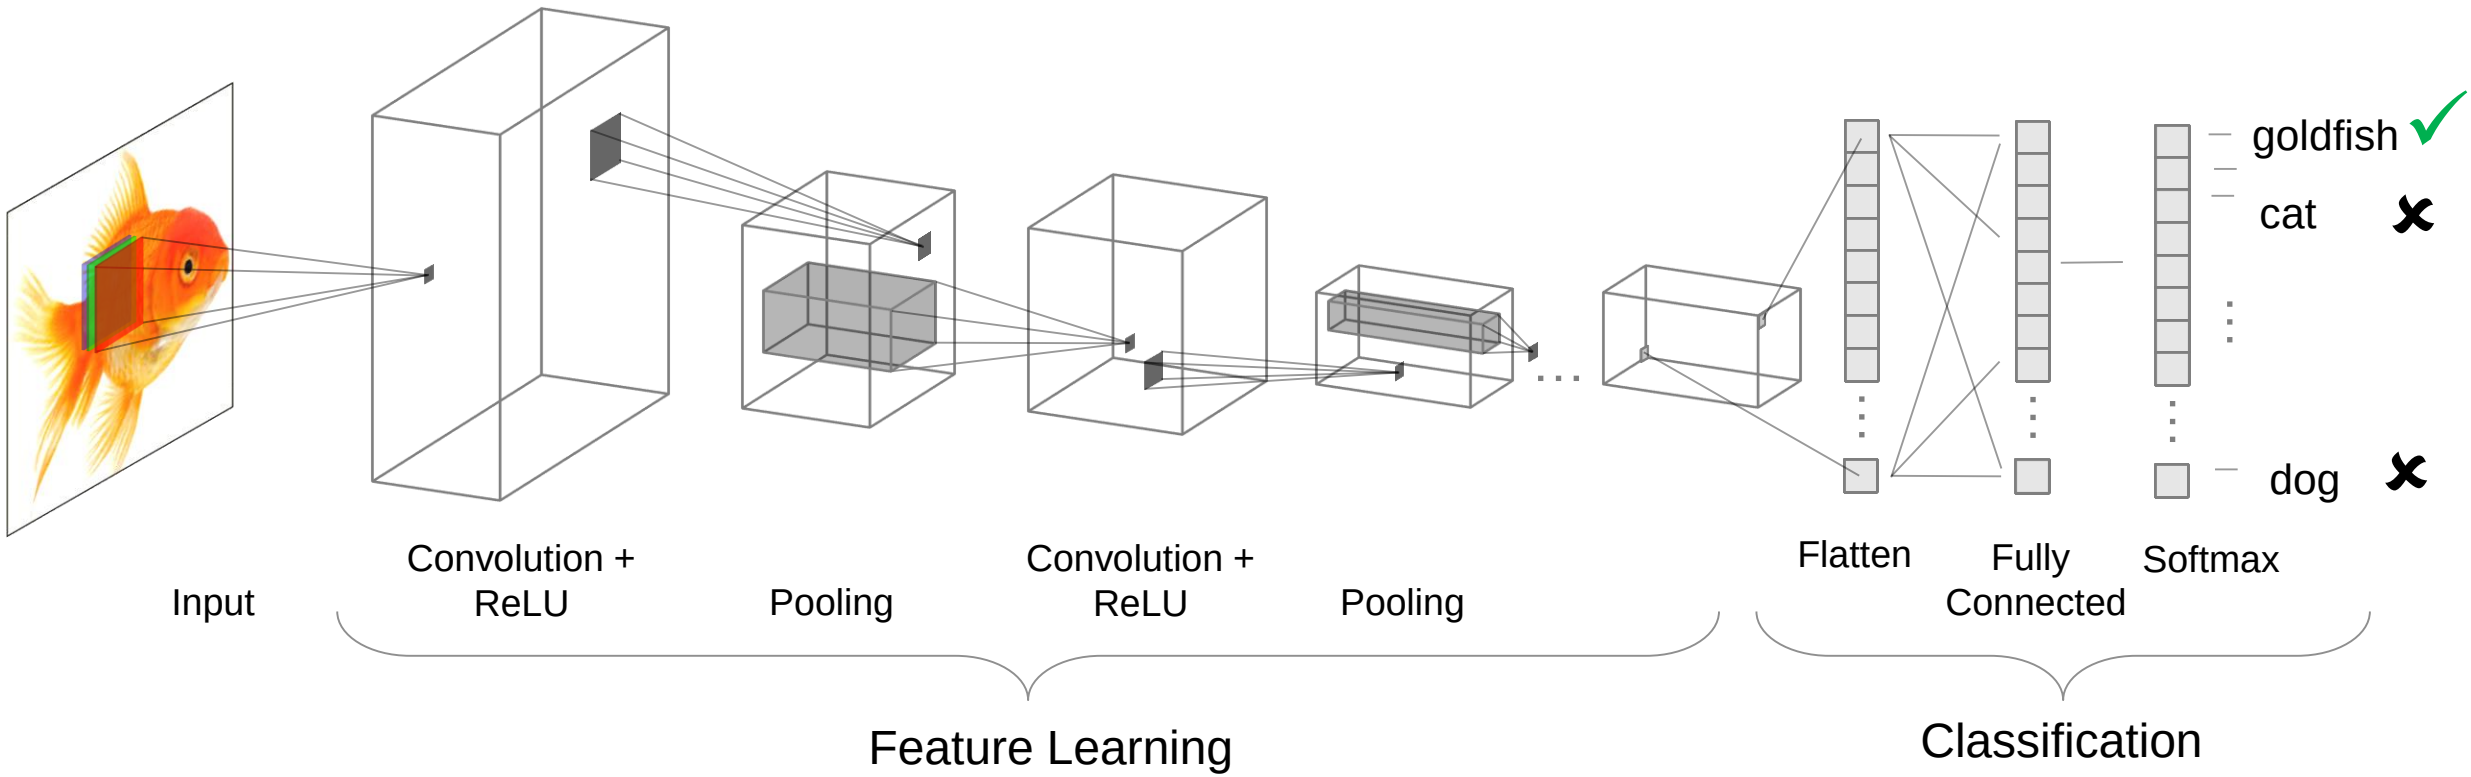

Takeaways

- Pre-trained networks have a pre-determined layer order that makes them effective for classifying images
  - Typically trained to classify lots of images
- Great starting point, but not consistently accurate
  - We'll fix this later with transfer learning!

Agenda

What is deep learning?- Overview

Why deep learning?

Pretrained networks and Network layers

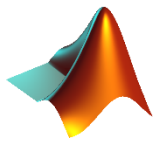

Building your own network

Transfer learning with CNNs

Deep Learning Workflow

Repeat these steps until network reaches desired level of accuracy

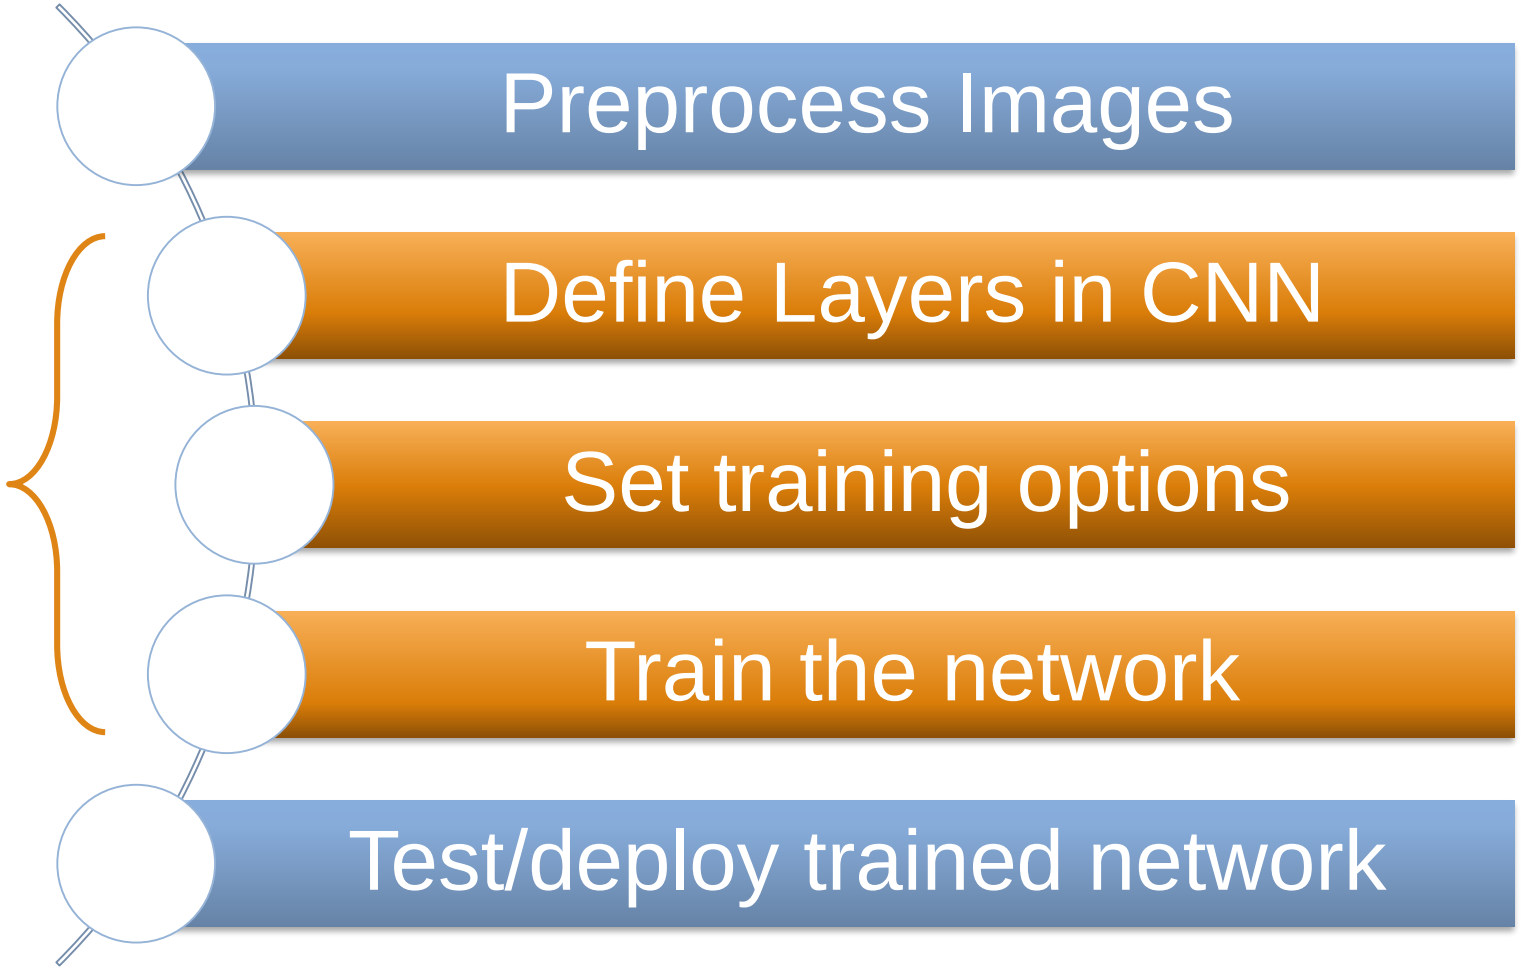

Let’s try it out!

Exercise: SS_MNIST_HandwritingRecognition.mlx

Takeaways

- Deep learning for image classification uses CNNs
- CNNs can have different combinations of initial layers but usually end with:
  - Fully Connected Layer
  - Softmax Layer
  - Classification Layer
- Important factors that affect accuracy and training time
  - Network architecture
  - Mini Batch Size
  - Initial learning rate

Agenda

What is deep learning?- Overview

Why deep learning?

Pretrained networks and Network layers

Building your own network

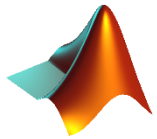

Transfer learning with CNNs

Two Approaches for Deep Learning

1. Train a Deep Neural Network from Scratch

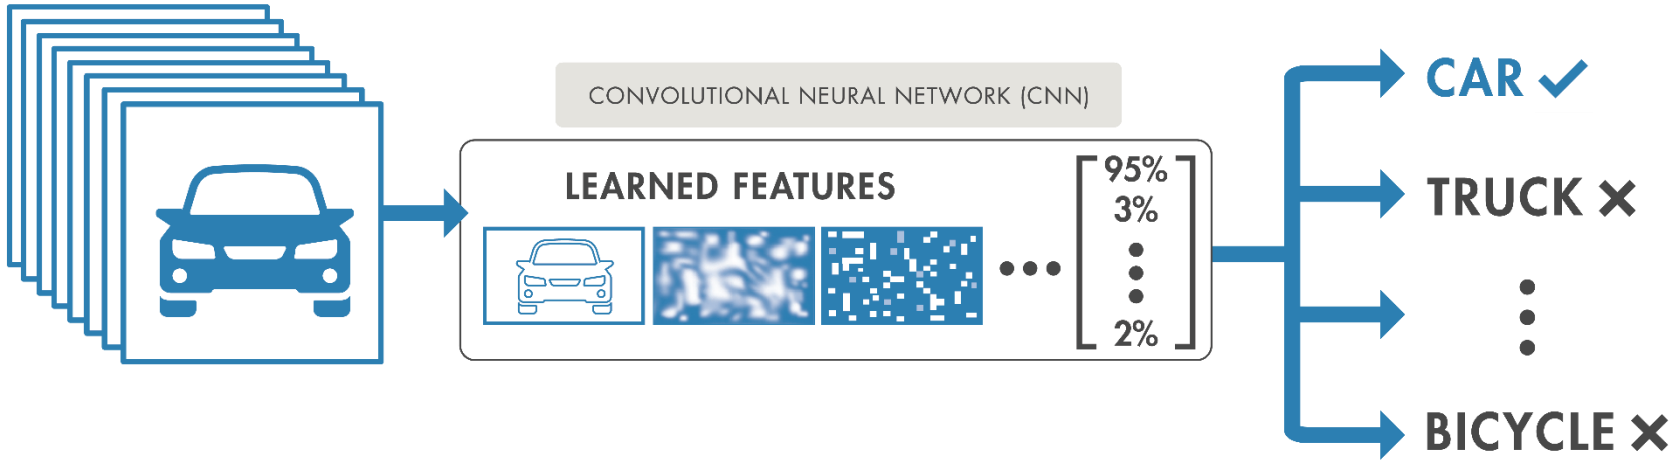

2. Fine-tune a pre-trained model (transfer learning)

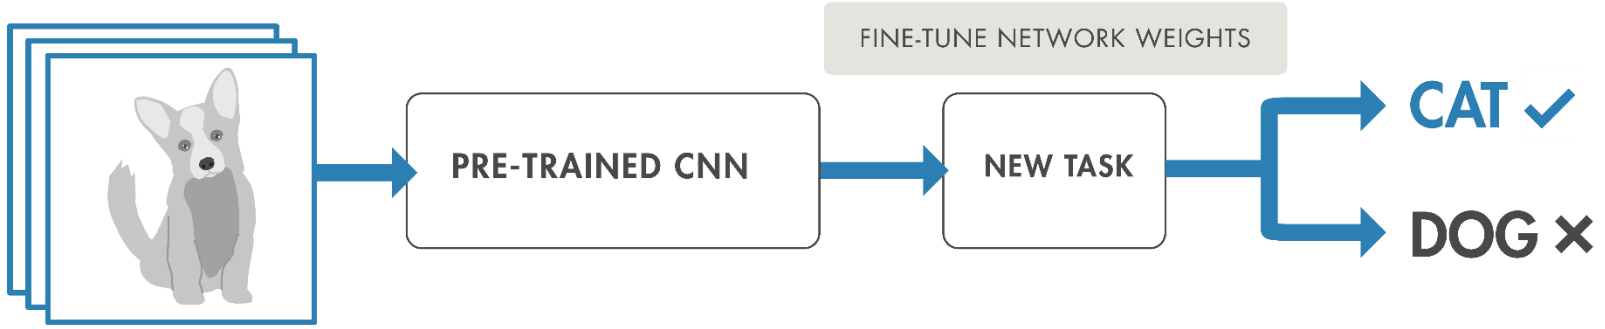

Two Approaches for Deep Learning

Fine-tune a pre-trained model (transfer learning)

Recommended when:

Training data	100s to 1000s of labeled images (small)
Computation	Moderate computation (GPU optional)
Training Time	Seconds to minutes
Model accuracy	Good, depends on the pre-trained CNN model

Train a deep neural network from scratch

Recommended when:

Training data	1000s to millions of labeled images
Computation	Compute intensive (requires GPU)
Training Time	Days to Weeks for real problems
Model accuracy	High (can over fit to small datasets)

Transfer Learning Workflow

Load pretrained network

Early layers that learned low-level features (edges, blobs, colors)

Last layers that learned task specific features

1 million images
1000s classes

Replace final layers

New layers to learn features specific to your data

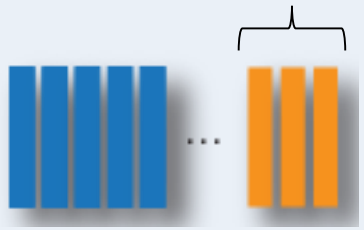

Fewer classes
Learn faster

Train network

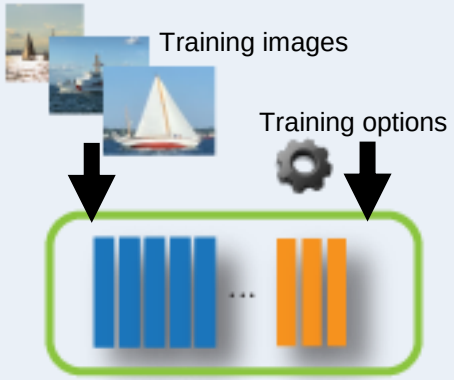

100s images
10s classes

Predict and assess network accuracy

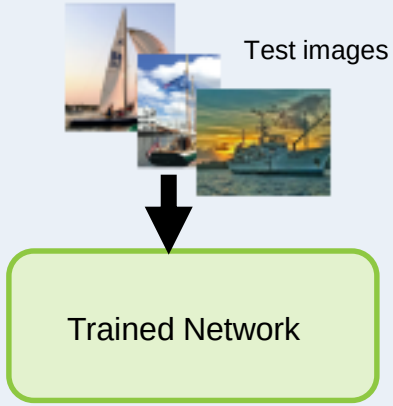

Deploy results

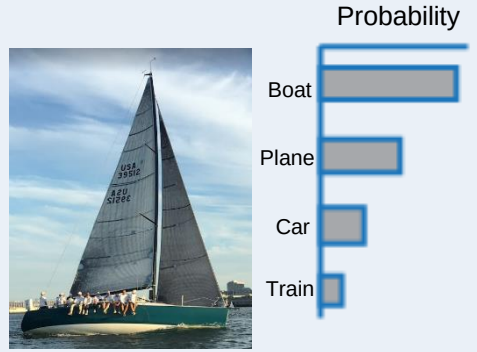

Transfer Learning Workflow – Step 1

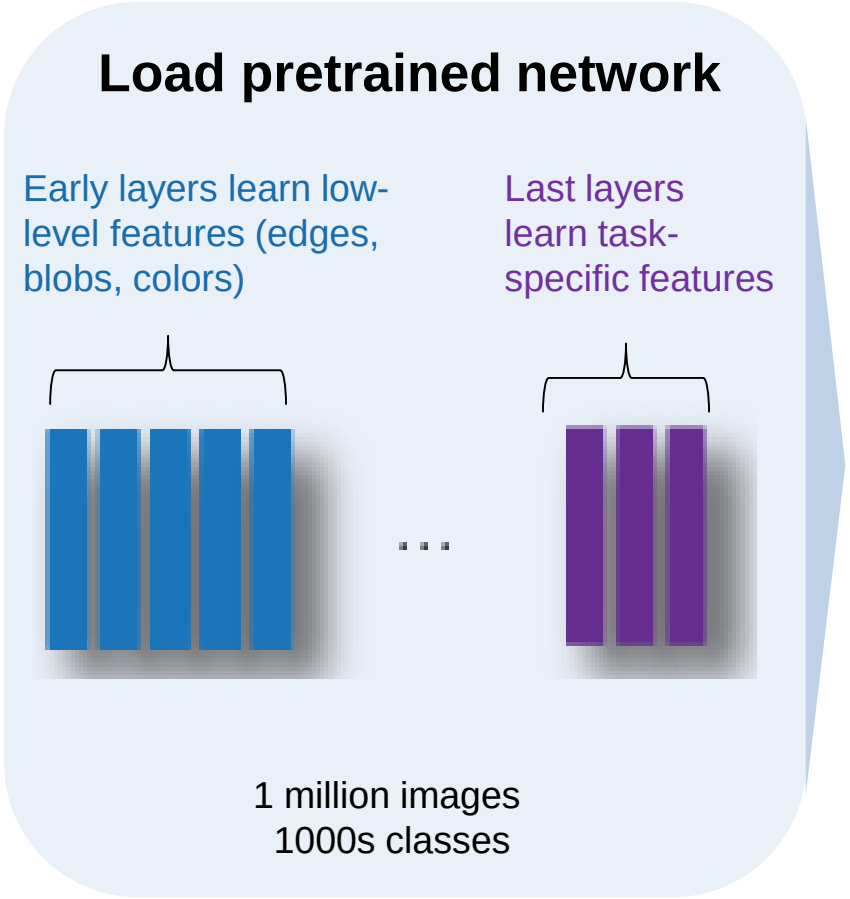

Transfer Learning Workflow – Step 2

Load pretrained network

Early layers that learned low-level features (edges, blobs, colors)

Last layers that learned task specific features

1 million images
1000s classes

Replace final layers

New layers learn features specific to your data

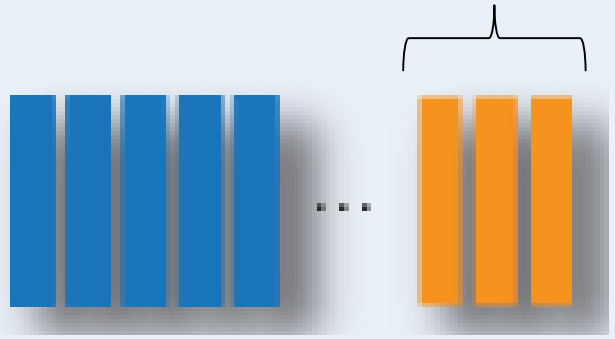

Fewer classes
Learn faster

Transfer Learning Workflow – Step 3

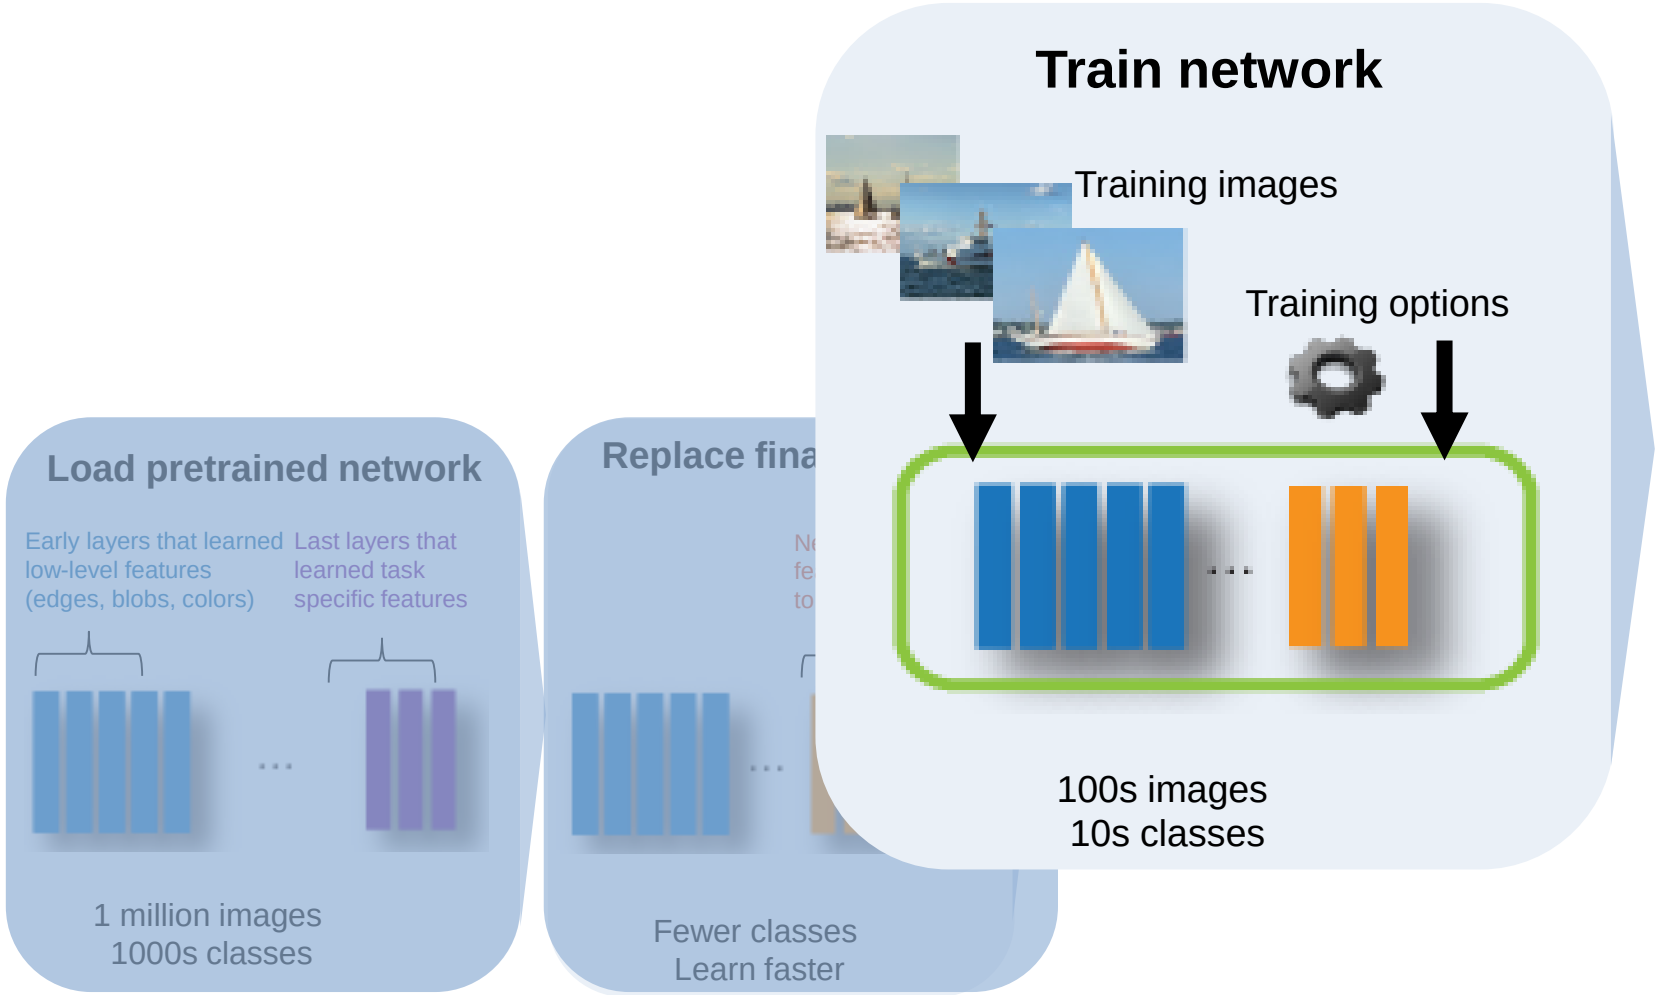

Transfer Learning Workflow – Step 4

Load pretrained network

Early layers that learned low-level features (edges, blobs, colors)

Last layers that learned task specific features

1 million images
1000s classes

Replace fina

Ne
fea
to

Fewer classes
Learn faster

100s images
10s classes

Predict and assess network accuracy

Test images

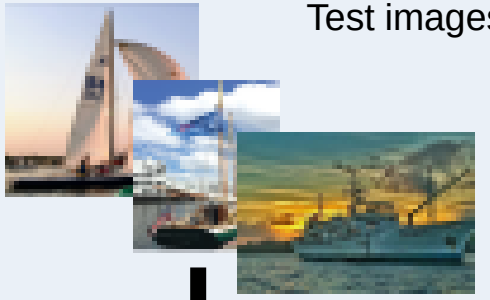

Trained Network

Transfer Learning Workflow – Step 5

Load pretrained network

Early layers that learned low-level features (edges, blobs, colors)

Last layers that learned task specific features

1 million images
1000s classes

Replace fina

Ne
fe
to

Fewer classes
Learn faster

Deploy results

Probability

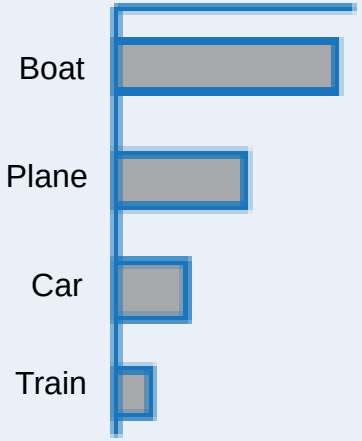

100s images
10s classes

lict and assess twork accuracy

Test images

Trained Network

Transfer Learning Workflow

Load pretrained network

Early layers that learned low-level features (edges, blobs, colors)

Last layers that learned task specific features

1 million images
1000s classes

Replace final layers

New layers to learn features specific to your data

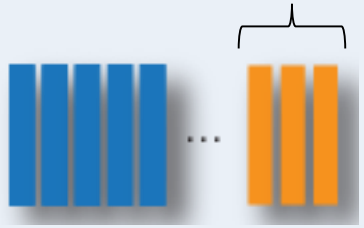

Fewer classes
Learn faster

Train network

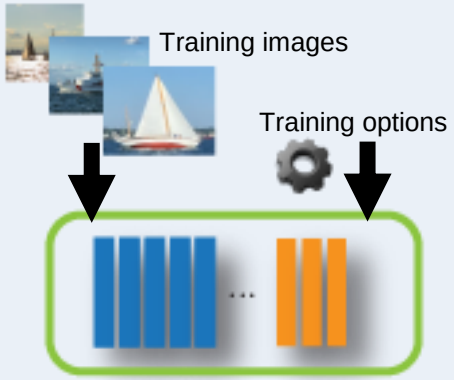

100s images
10s classes

Predict and assess network accuracy

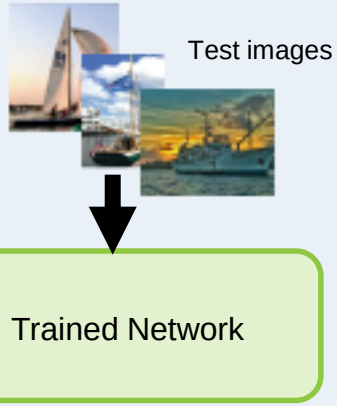

Deploy results

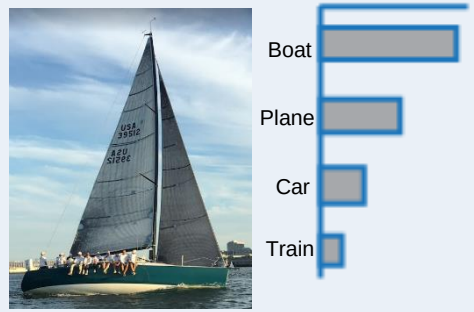

Let’s try it out!

Exercise: SS_SeeFoodTransferLearning.mlx

Takeaways – Transfer Learning

- Replace last layers with our own layers
- Efficient way to modify pre-trained models to our needs
- Use an Image datastore when working with lots of images

One Step Left – Deployment!

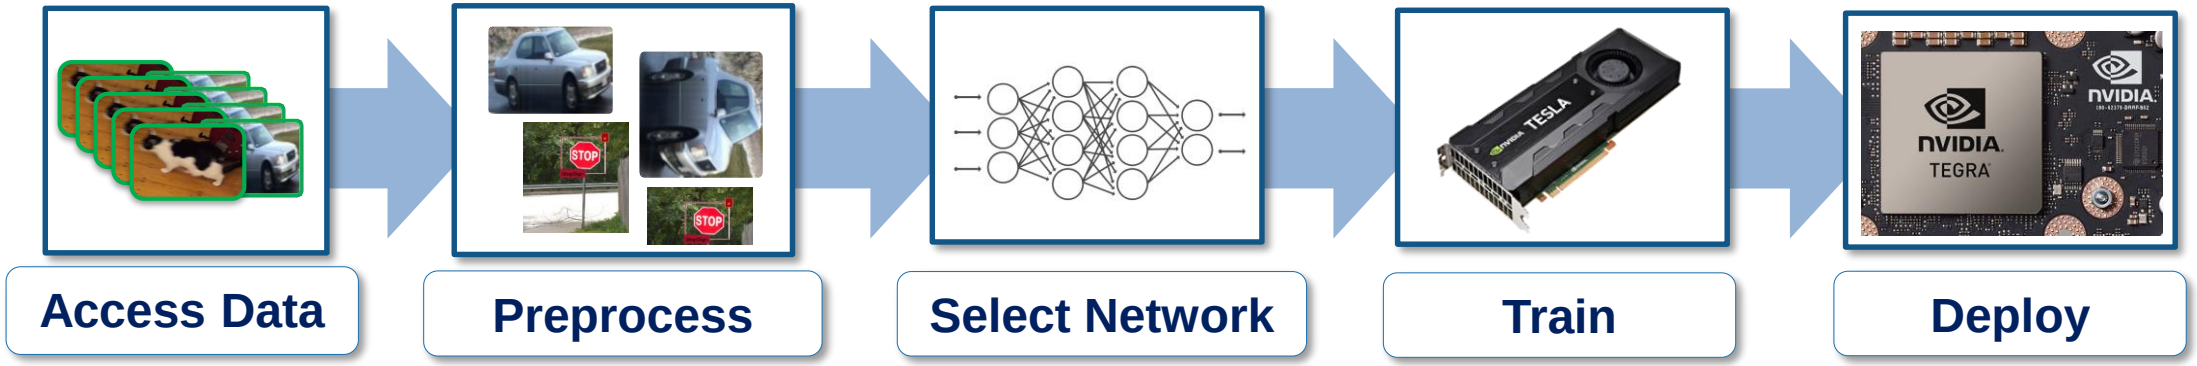

Image Acquisition

Image Processing

Computer Vision System

Neural Network

Parallel Computing

GPU Coder

R2017b

What Next?

Deep Learning Onramp

This free self-paced tutorial provides an interactive introduction to practical deep learning. It focuses on using MATLAB® to apply deep learning methods to perform image recognition. The tutorial consists of hands-on exercises and short videos. In the exercises, you will enter commands in an online version of MATLAB and receive contextual feedback that will help you correct common mistakes. Topics include:

- Convolutional neural networks
- Preprocessing images
- Using pretrained networks
- Transfer learning
- Evaluating network performance

Available in English and Japanese languages

Launch the course

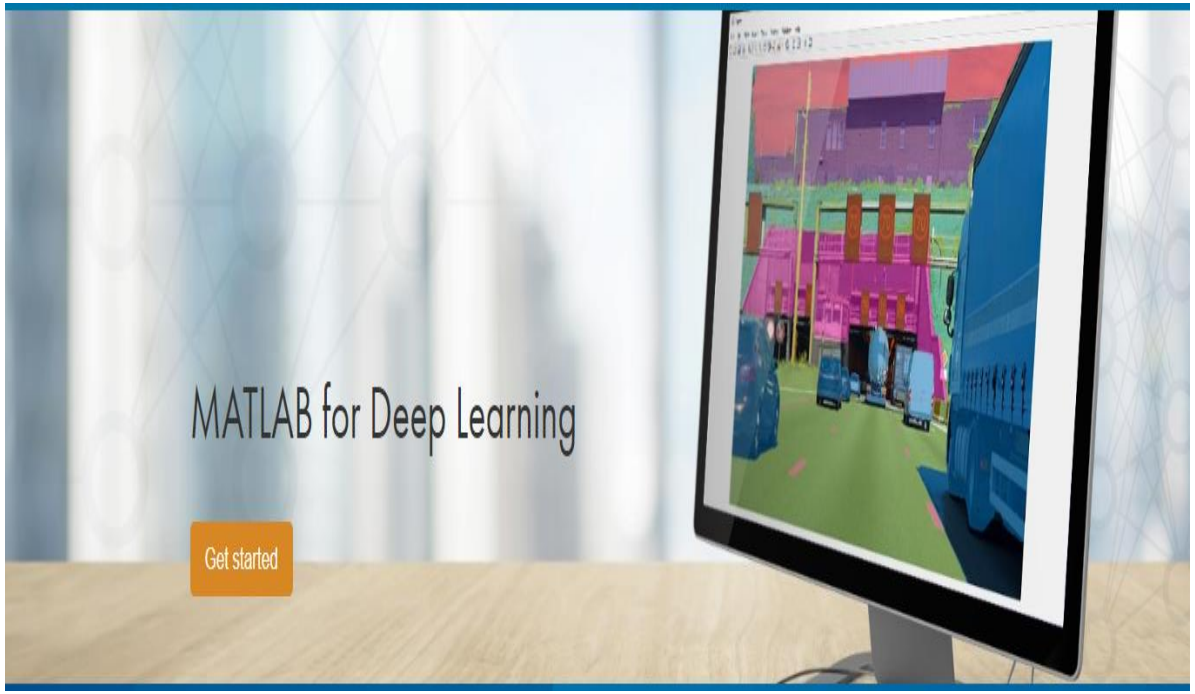

Training Services

Exploit the full potential of MathWorks products

Flexible delivery options:

- Public training available in several cities
- Onsite training with standard or customized courses
- Web-based training with live, interactive instructor-led courses

More than 48 course offerings:

- Introductory and intermediate training on MATLAB, Simulink, Stateflow, code generation, and Polyspace products
- Specialized courses in control design, signal processing, parallel computing, code generation, communications, financial analysis, and other areas

www.mathworks.in/training

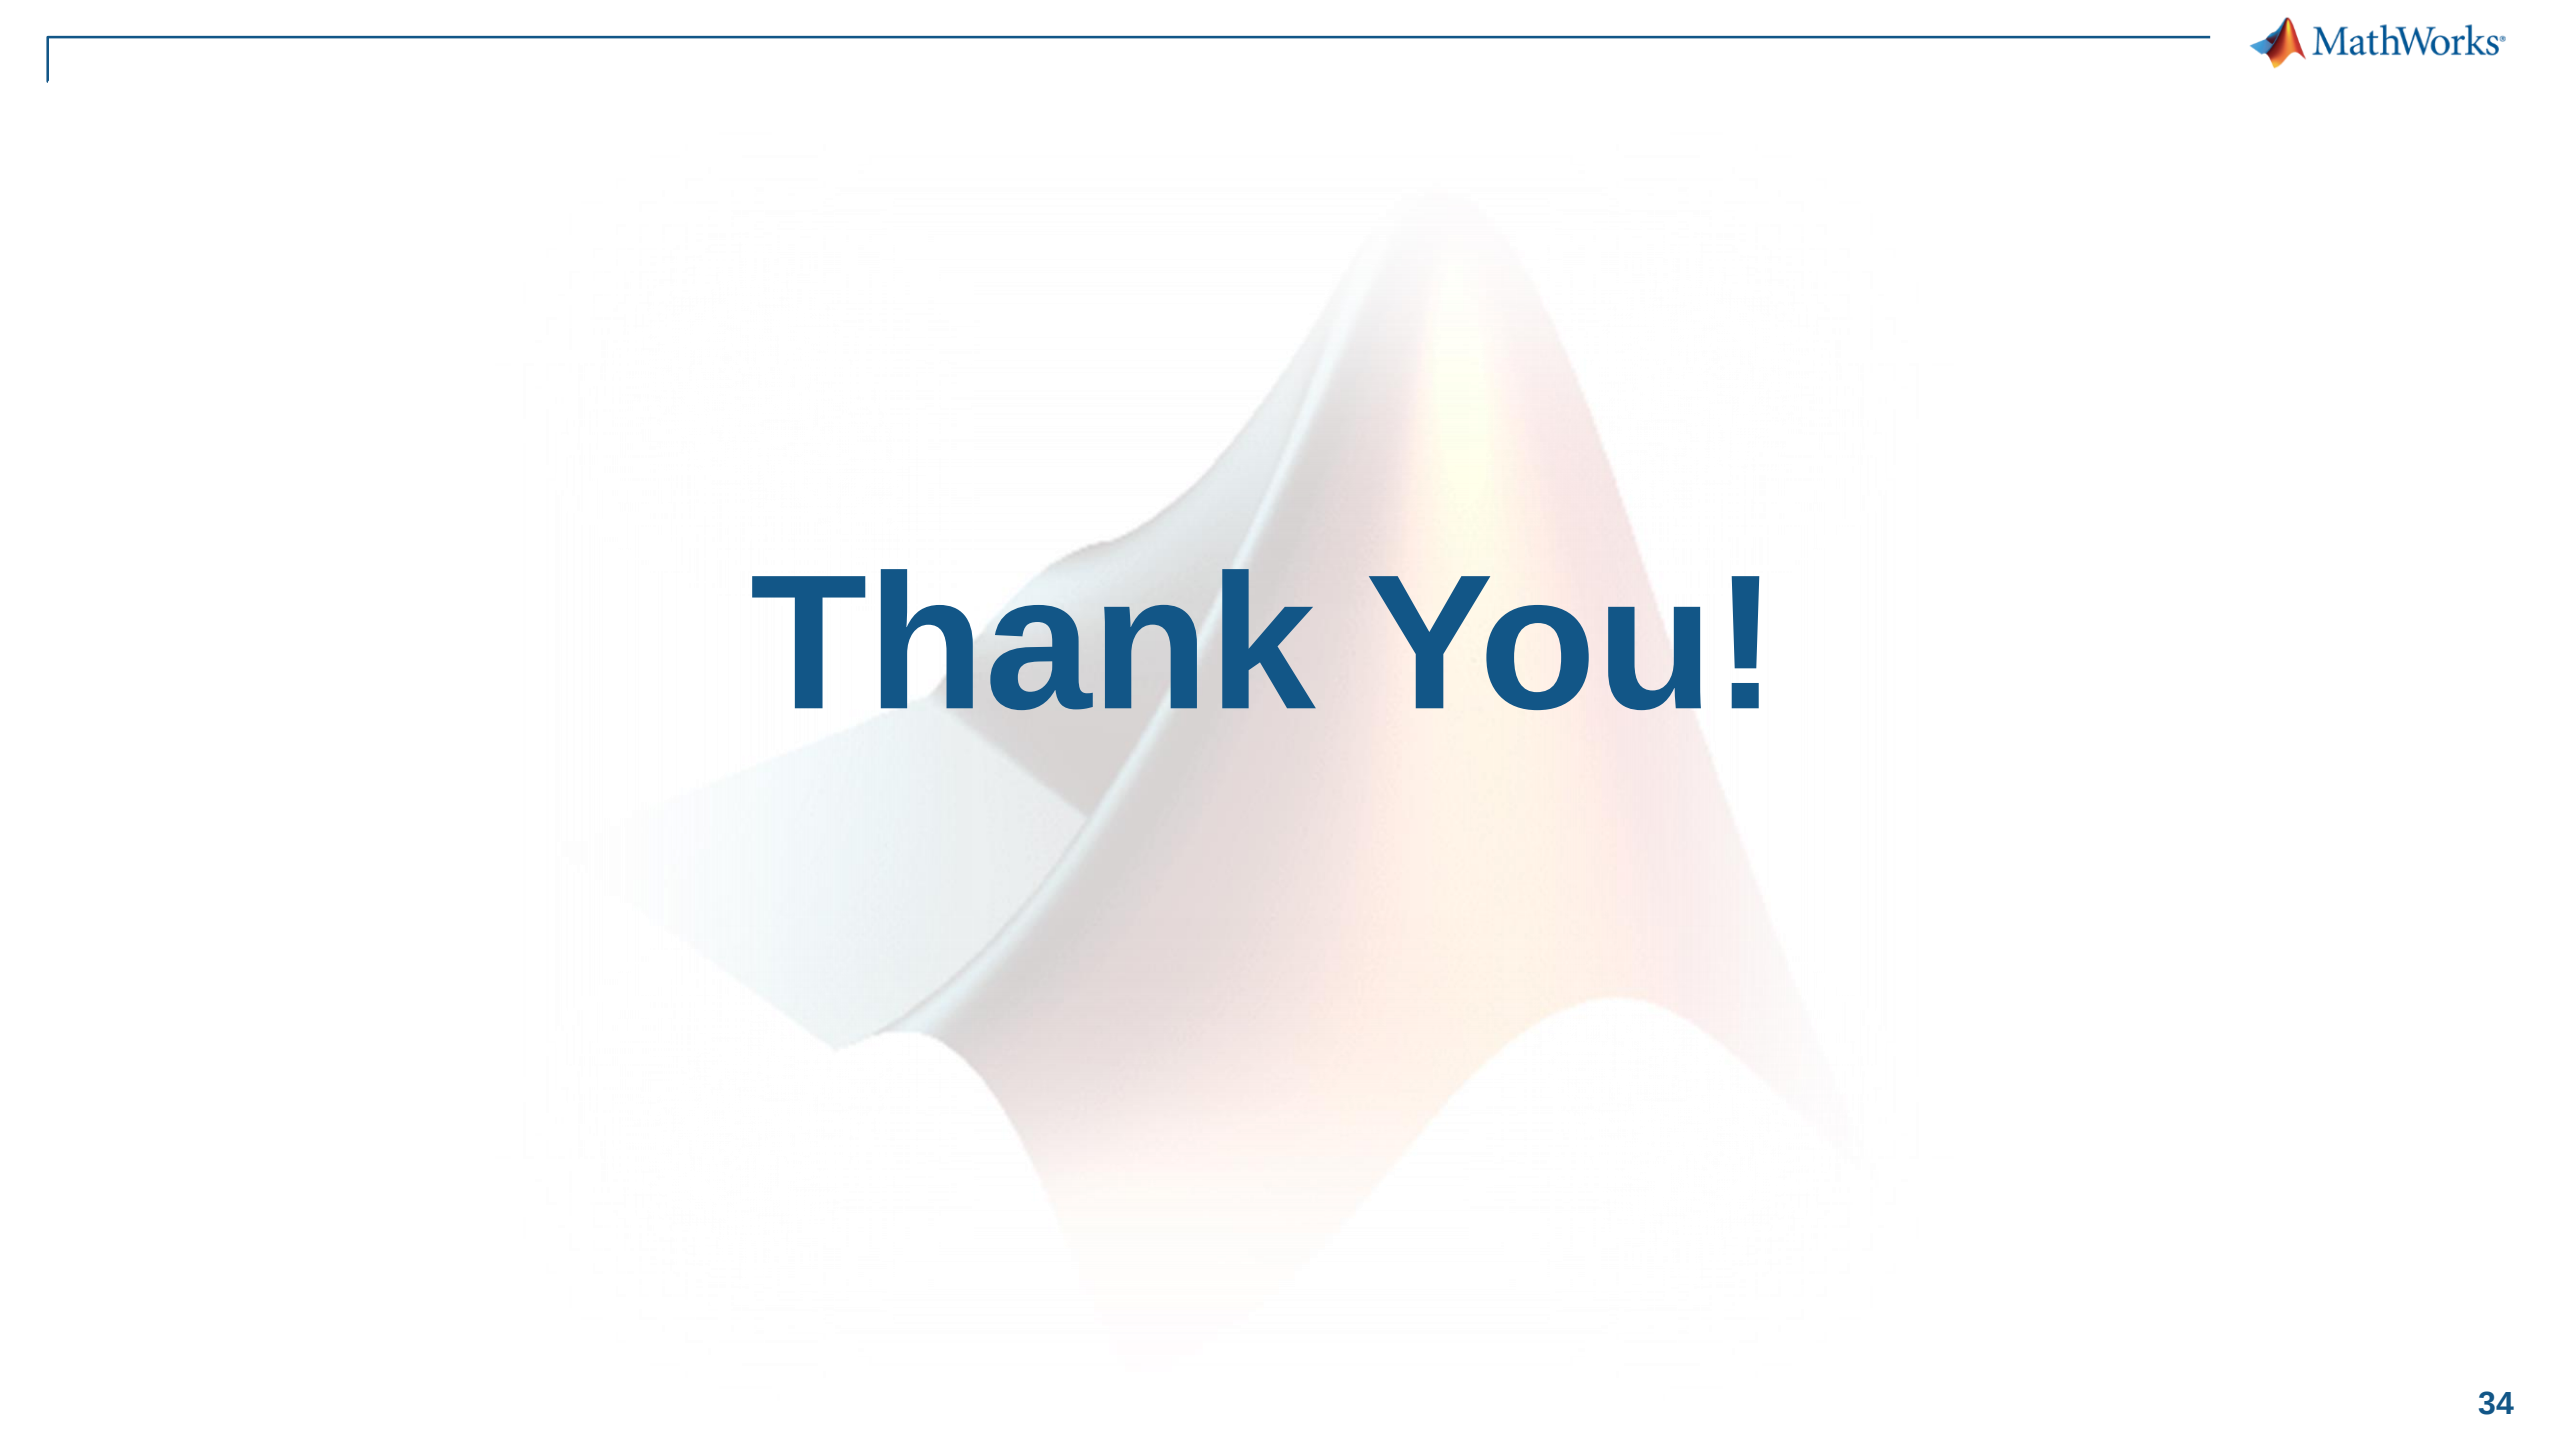
Thank You!

Questions?