Hands-on Deep Learning Workshop

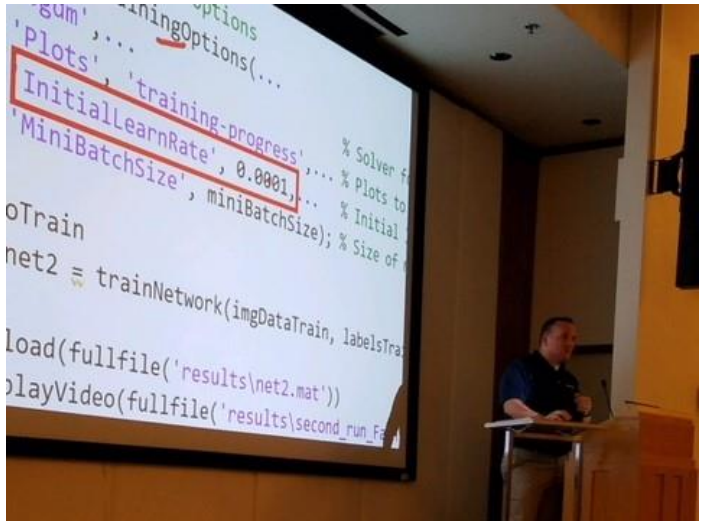

Matthias Sommer

Campus-Wide License at University of Bern

University & lab computers

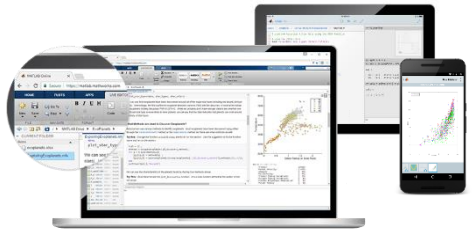

Personal Computers & Mobile Devices

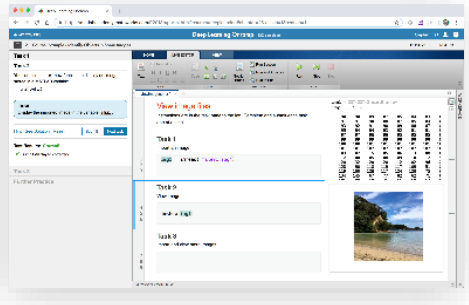

Self-paced online learning

Online access

- License covers all faculty, staff, students and their devices
- Access on campus, in lab and field, and at home
- Get MATLAB via https://ch.mathworks.com/ or IT store of the University of Bern

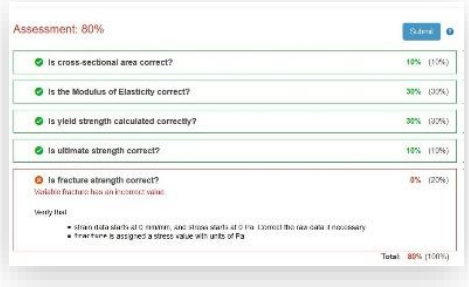

Auto-graded homework

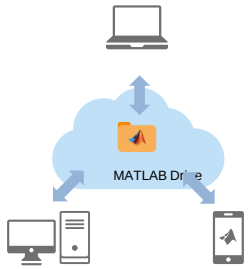

Cloud Storage & Sharing

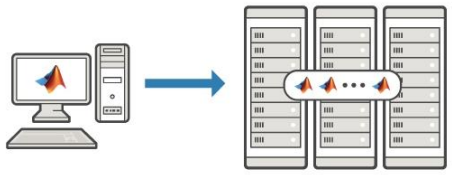

Clusters & HPC

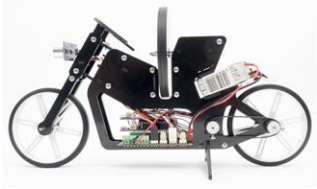

Low-cost hardware support

Self-Paced, Online Training for MATLAB & Simulink

Campus-Wide Online Training

- Hands-on MATLAB and Simulink experience
- Measurable progress report and completion certificate
- Interactive lessons with immediate feedback
- 24/7 availability

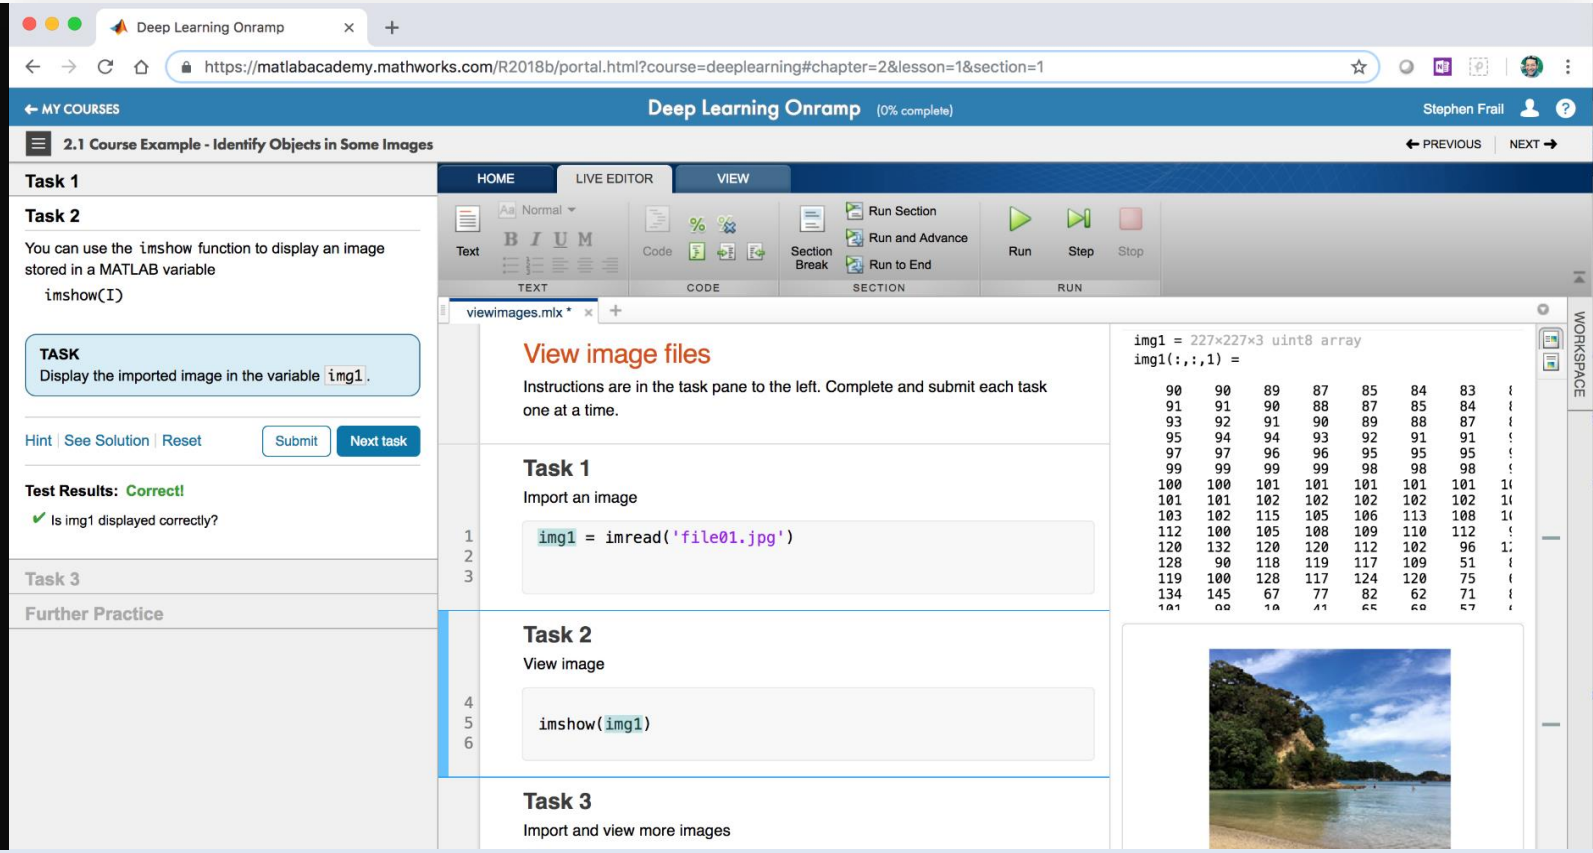

FREE — MATLAB Onramp

FREE — Simulink Onramp

FREE — Deep Learning Onramp

FREE — Stateflow Onramp

NEW FREE — Machine Learning Onramp

11 hours of FREE content available for everyone

https://matlabacademy.mathworks.com/

Set-Up Instructions

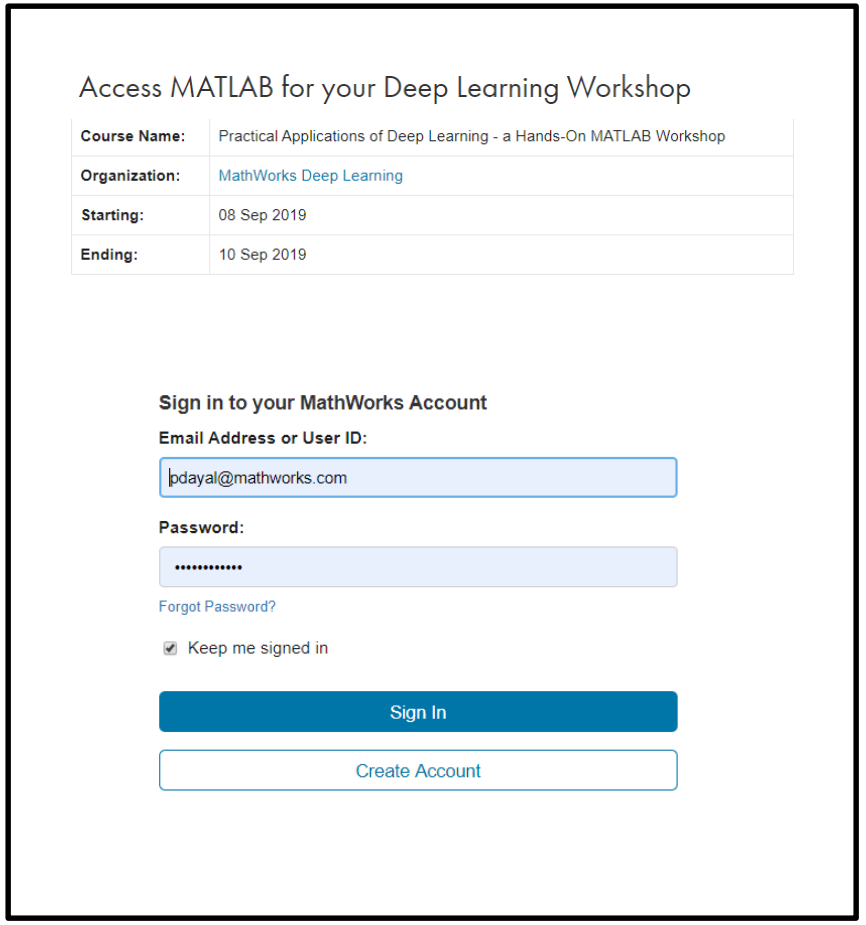

1 - Log into MATLAB Online

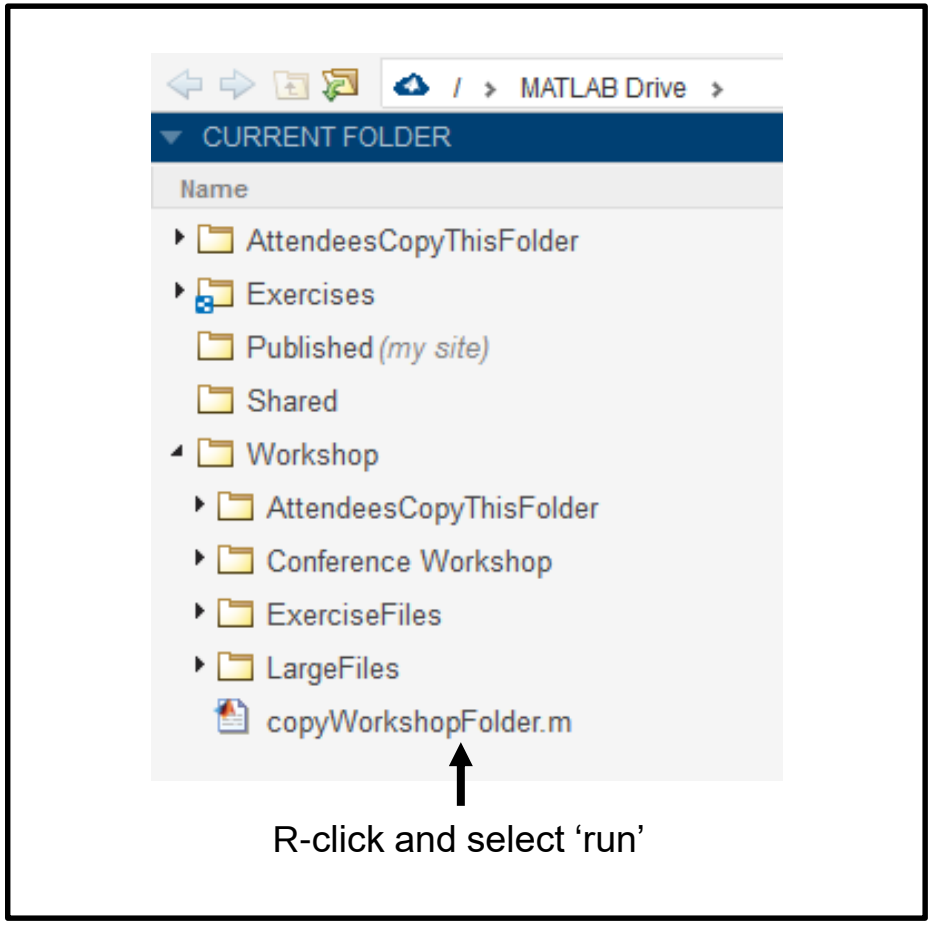

2 - Run ‘copyWorkshopFolder.m’


Deep Learning Demo
Image Classification

Agenda

Introduction

Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Exercises 2 and 3: Exploring pretrained networks/Classifying handwritten digits

Exercise 4: Transfer Learning – OR – Signal Classification Exercise

Demo: Deploying Deep Networks– OR – Improving Network Accuracy

Conclusion

What is Deep Learning?

- Subset of machine learning with automatic feature extraction
  - Learns features and tasks directly from data
- Accuracy can surpass traditional ML Algorithms

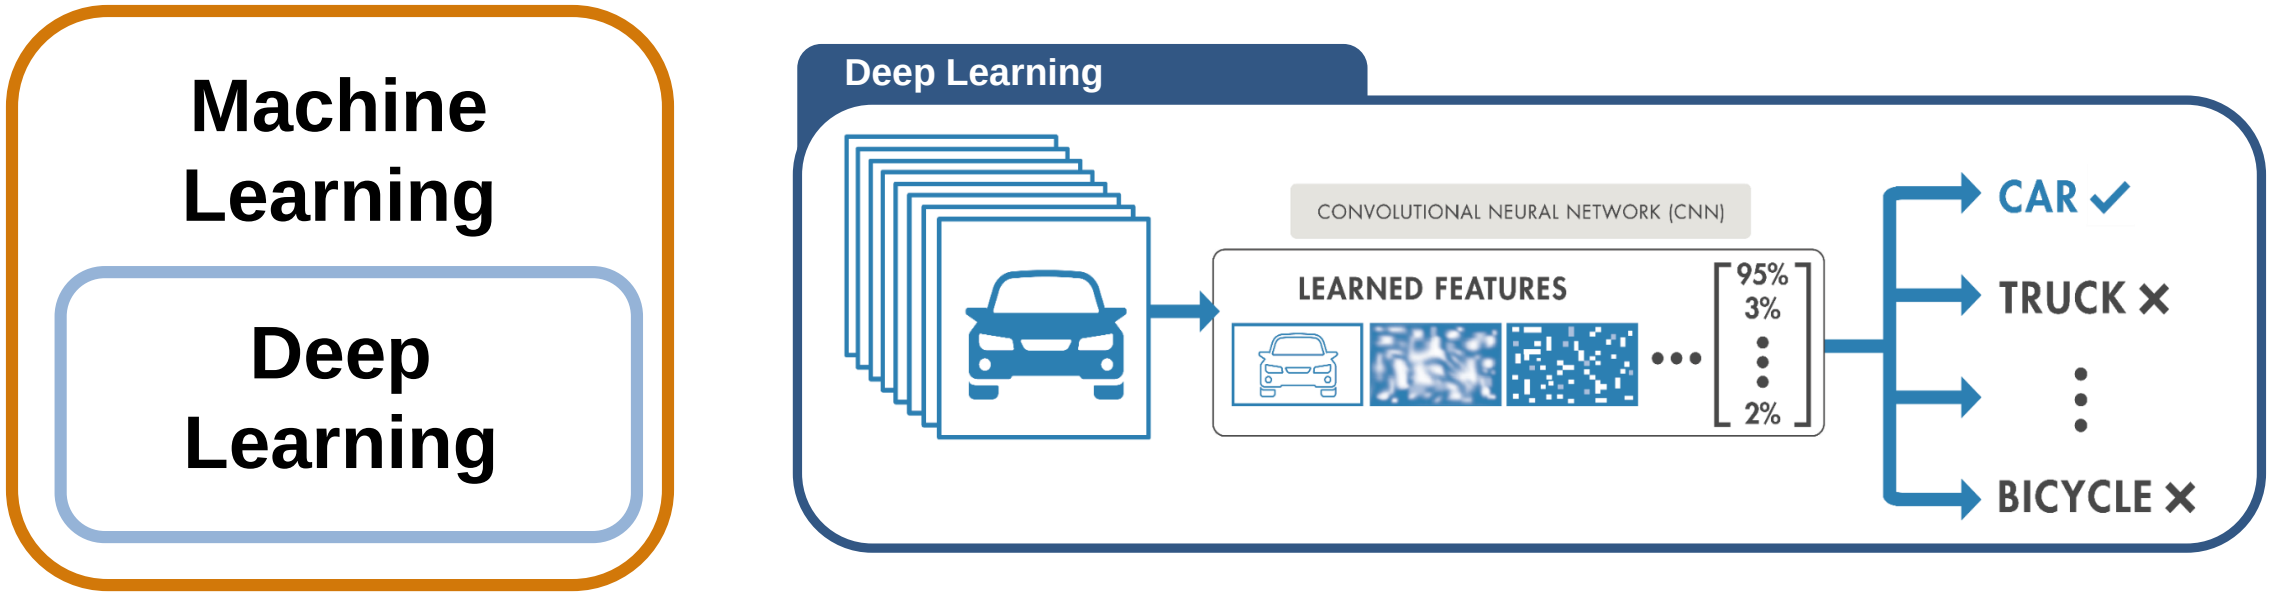

Deep Learning Workflow

PREPARE DATA

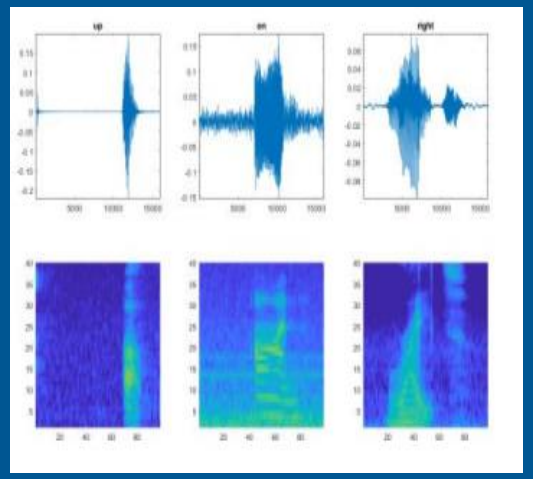

The data must be labeled and preprocessed to give accurate results

TRAIN MODEL

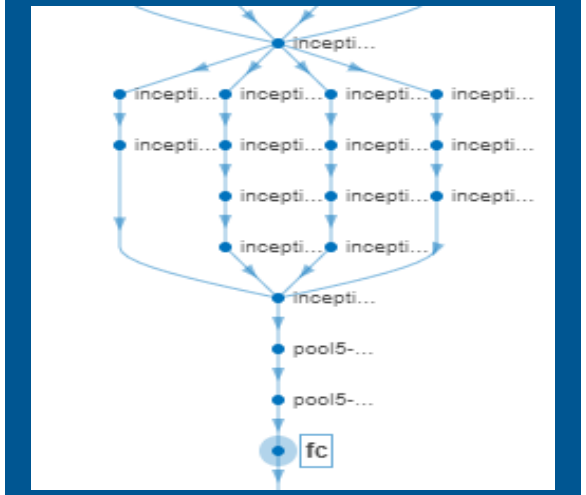

Build a neural network that learns from your dataset

DEPLOY SYSTEM

Integrate your trained model onto embedded hardware or cloud

Deep Learning Examples

Shell

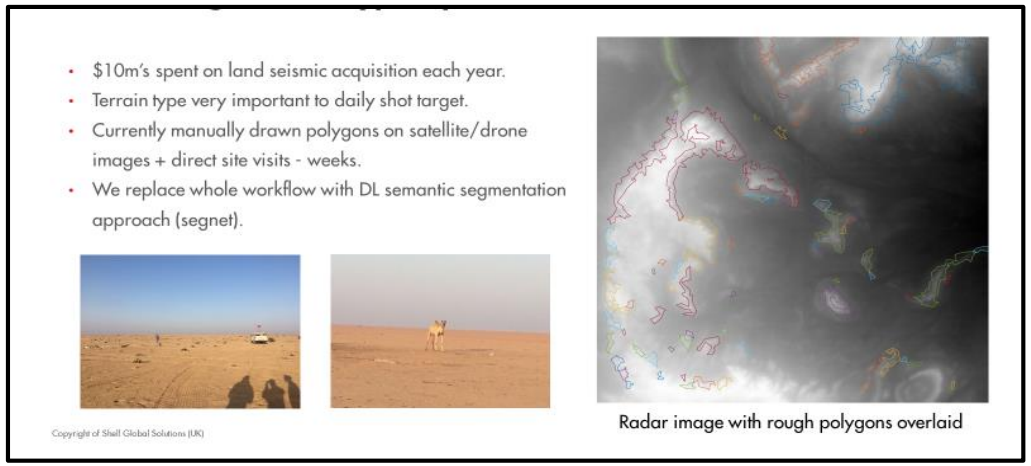

Terrain Recognition with Hyperspectral Data

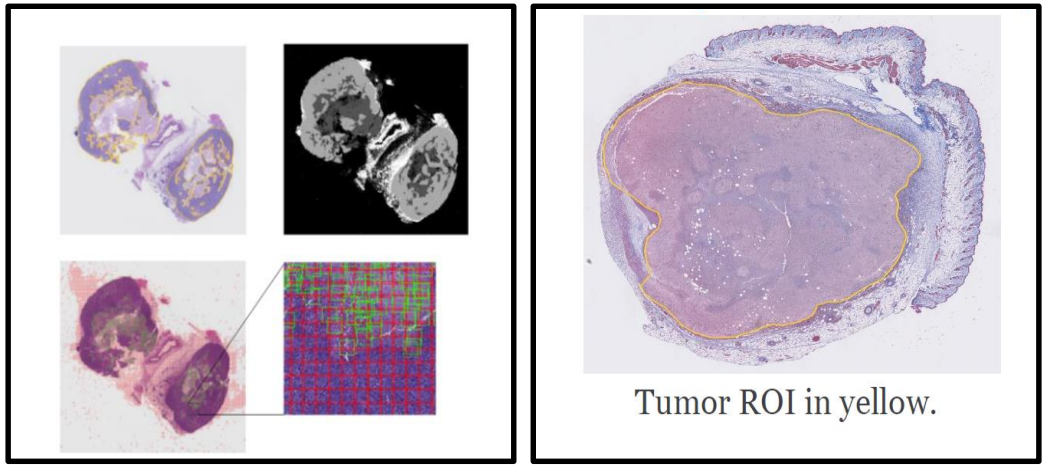

CNNs for Digital Pathology Analysis

Genentech

Veoneer

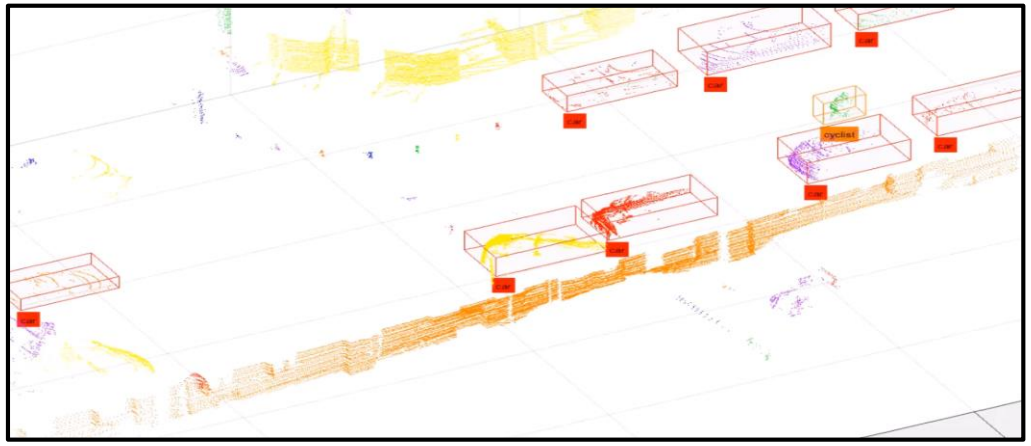

LiDAR-Based Sensor Verification

Equipment Classification

Caterpillar

MATLAB's deep learning workflows were designed for engineers and scientists in many domains

Computer Vision

Image Processing

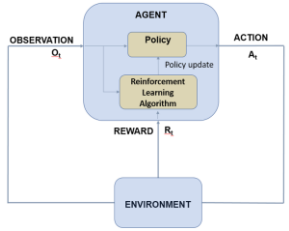

Control Design

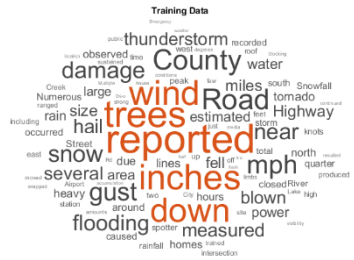

Text Analytics

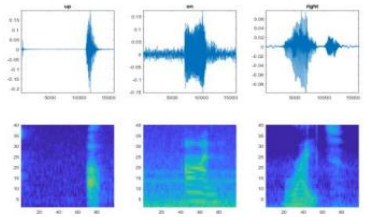

Audio Processing

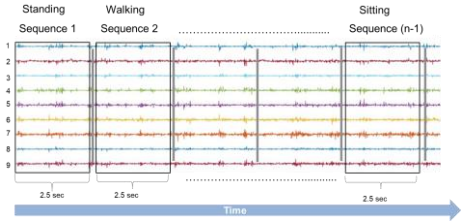

Sensor Data Analysis

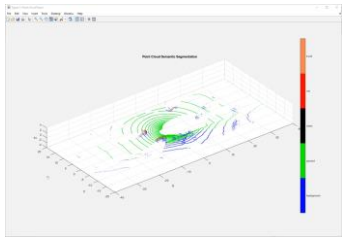

Lidar Processing

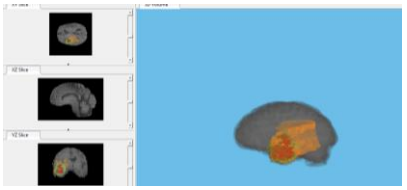

N-D Volumes

Sensor Data Analysis

Deep Learning models are Neural networks (for all data types)

- Deep neural networks have many layers
- Data is passed through the network, and the layer parameters are updated (training)

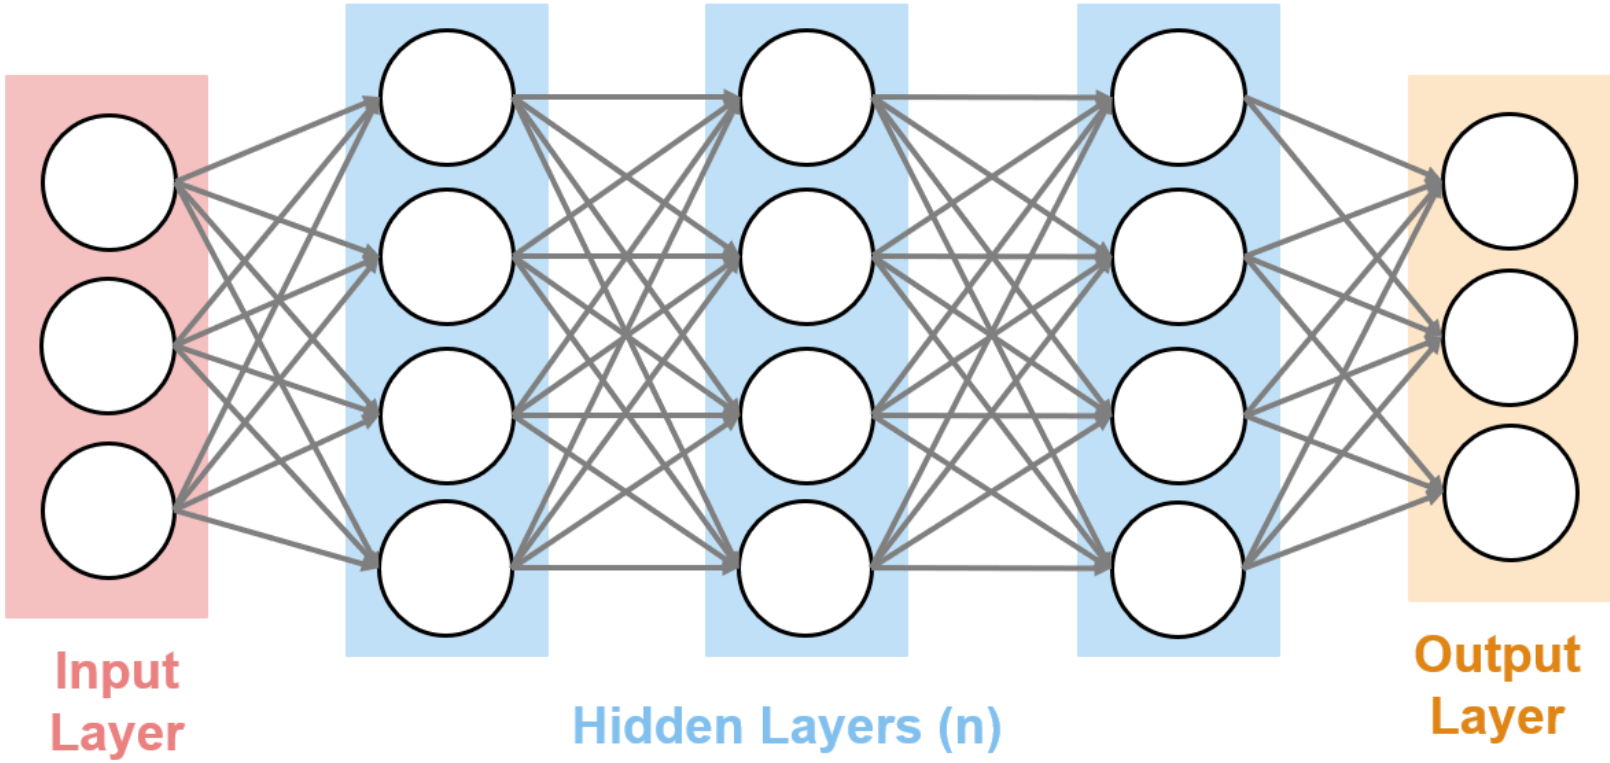

Exercise 1 – Deep Learning in 6 Lines of Code

Purpose:

- Ensure MATLAB Online is running properly
- Use a neural network to classify an image

To Do:

1. Open work_deeplearningin6lines.mlx
2. Follow along with instructor

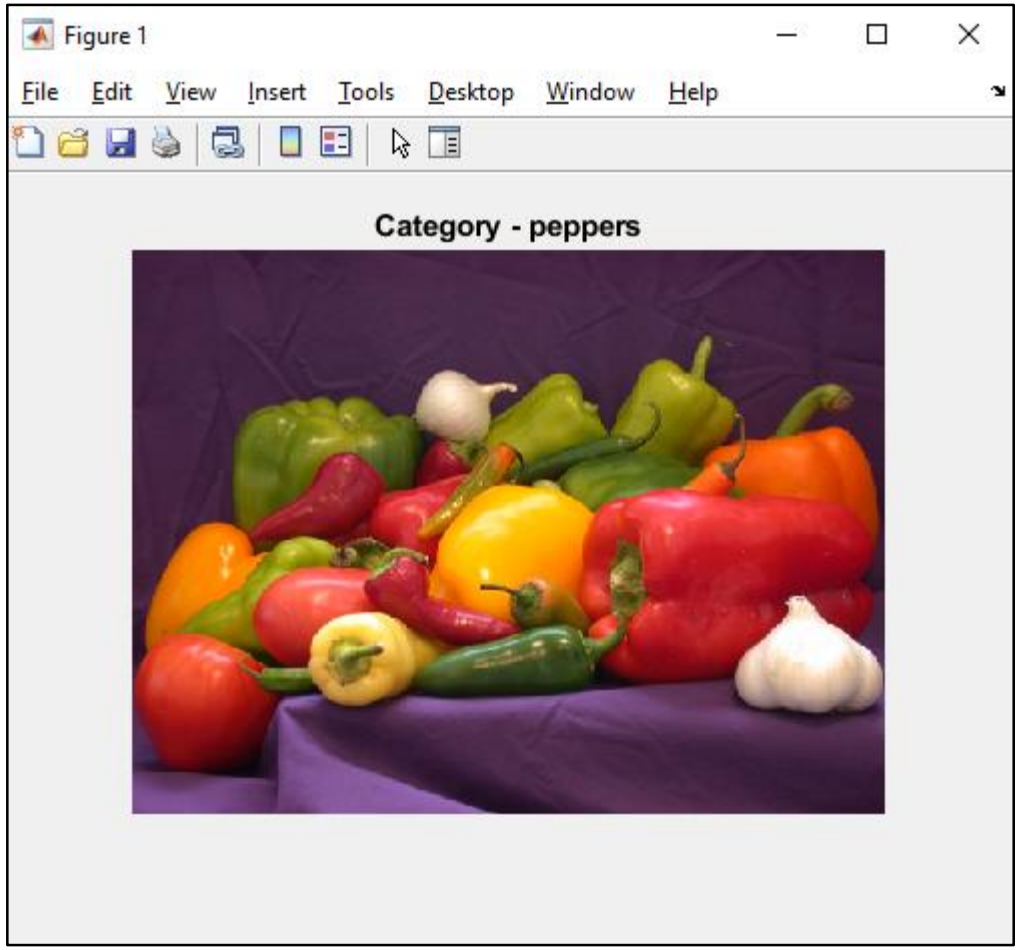

We Can Build Networks from Scratch or Use Pretrained Models

- Pretrained models have predefined layer orders and parameter values
- Can be used for inference without training

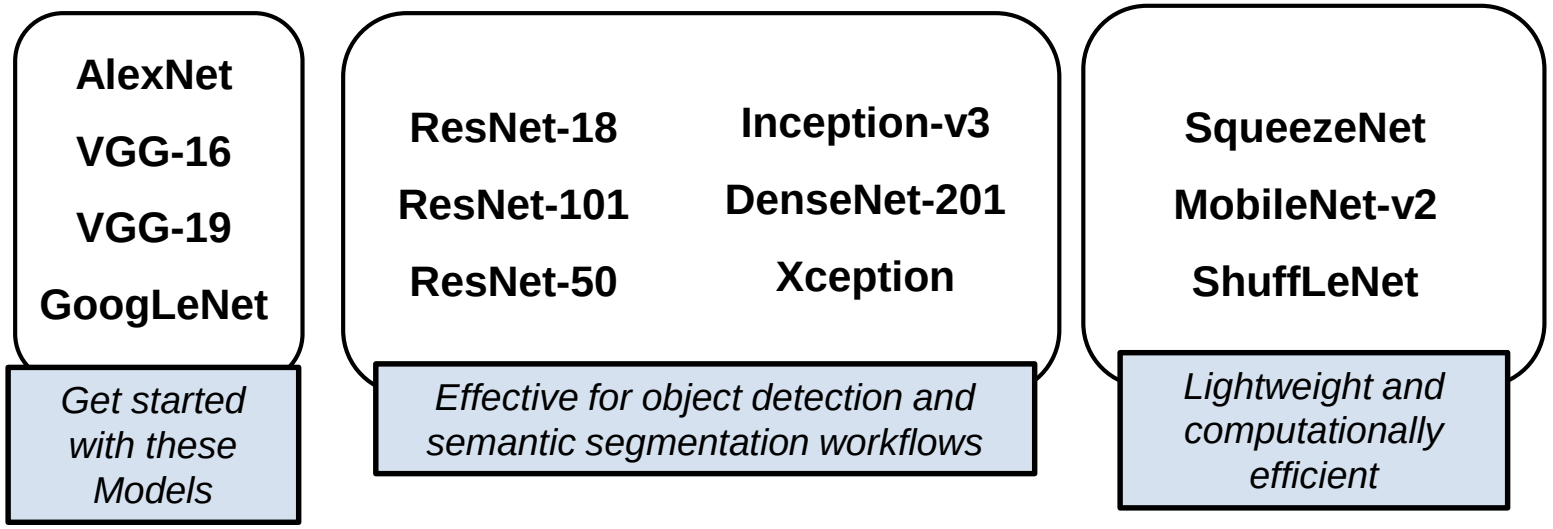

Full list of models available at: https://www.mathworks.com/help/deeplearning/ug/pretrained-convolutional-neural-networks.html

Where Can I Access Pretrained models

- Many are built into MATLAB
- Others can found on the web and imported into MATLAB

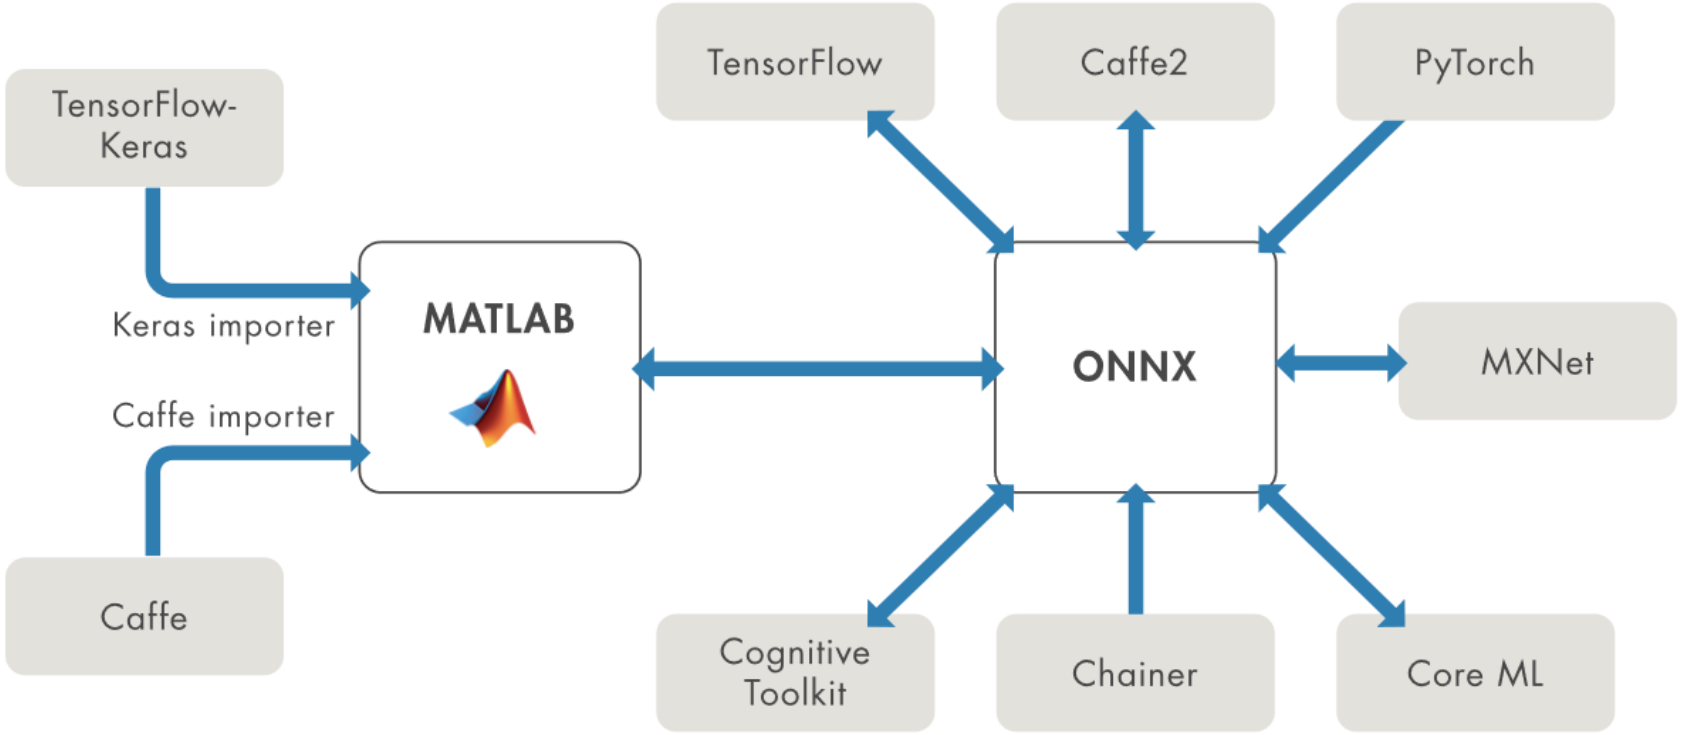

Open Neural Network Exchange

Exercise 2 – Pretrained Models

Purpose:

- Learn how to use pre-trained models in MATLAB.
- See how different network architectures affect results.
- Use datastores to access data efficiently

To Do:

1. Open work_pretrainednetworks.mlx.

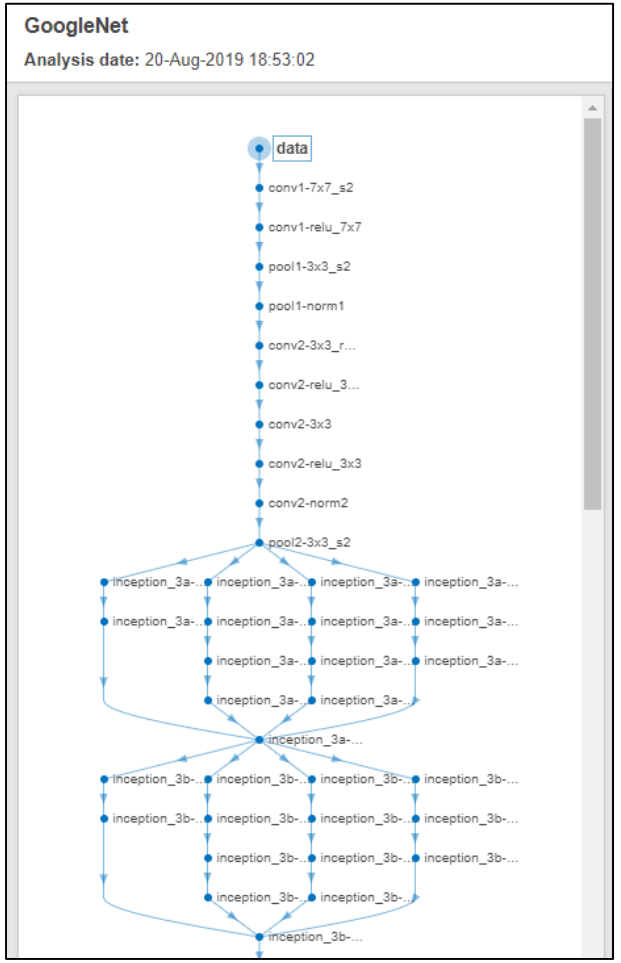

Pretrained models aren't always enough, we may have to build and train networks from scratch

PREPARE DATA

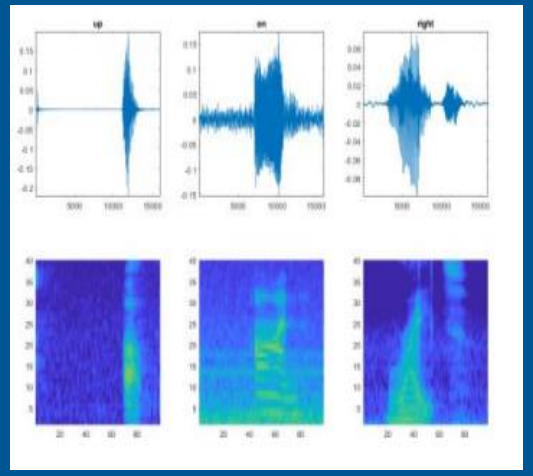

TRAIN MODEL

- Model design and tuning
- Hardware-accelerated training
- Model exchange across frameworks

DEPLOY SYSTEM

Creating Layer Architectures

- Convolution Neural Networks – CNN
- Special layer combinations that make them adept at classifying images

- Convolution Layer
- Max Pooling Layer
- ReLU Layer

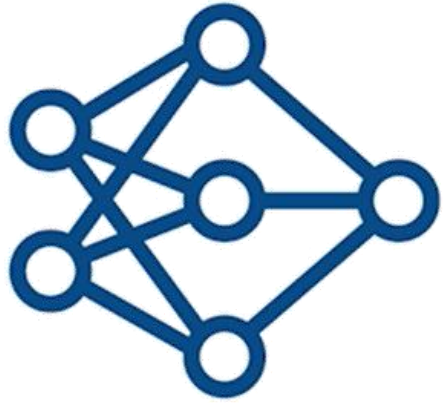

Convolution Layers Search for Patterns

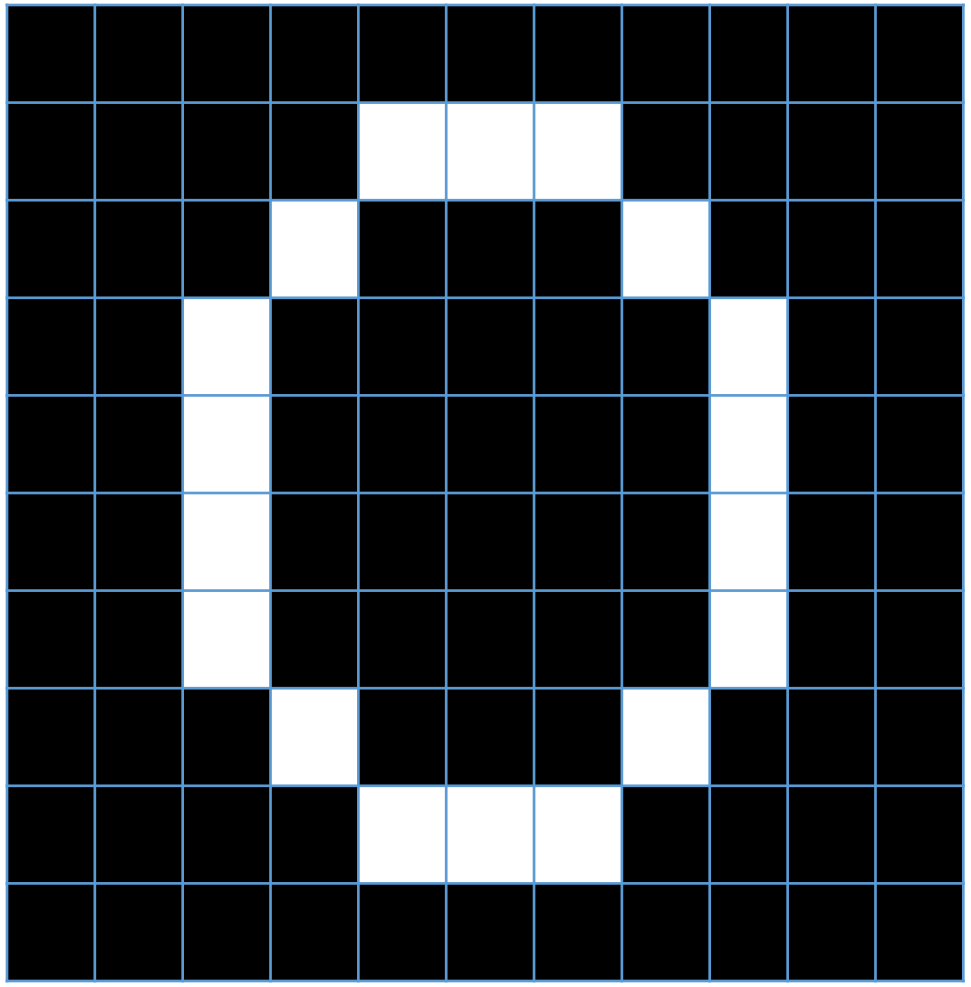

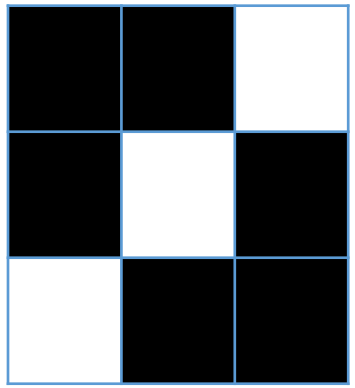

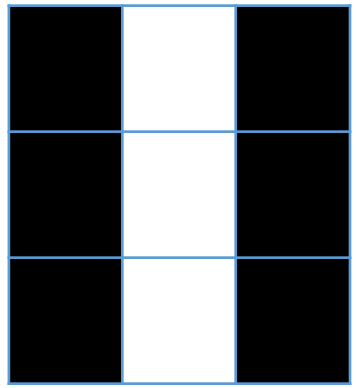

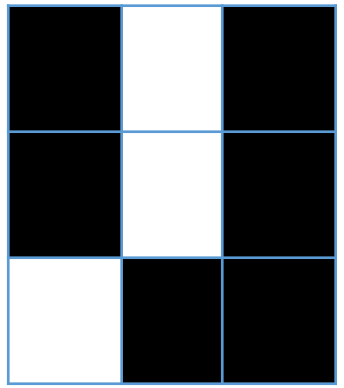

These patterns would be common in the number 0

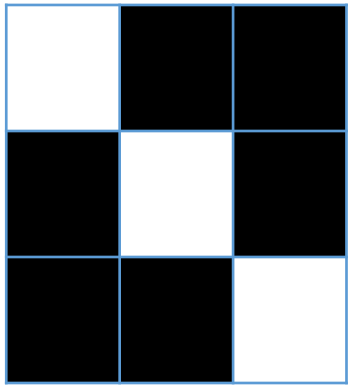

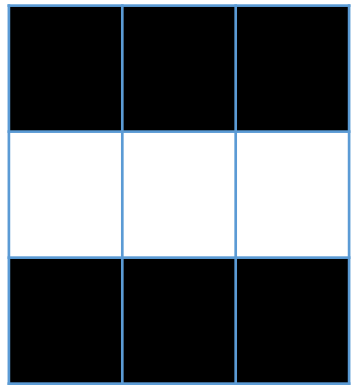

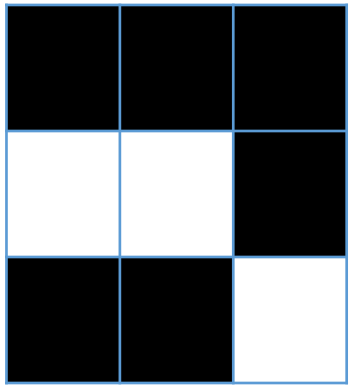

All patterns are compared to the patterns on a new image.

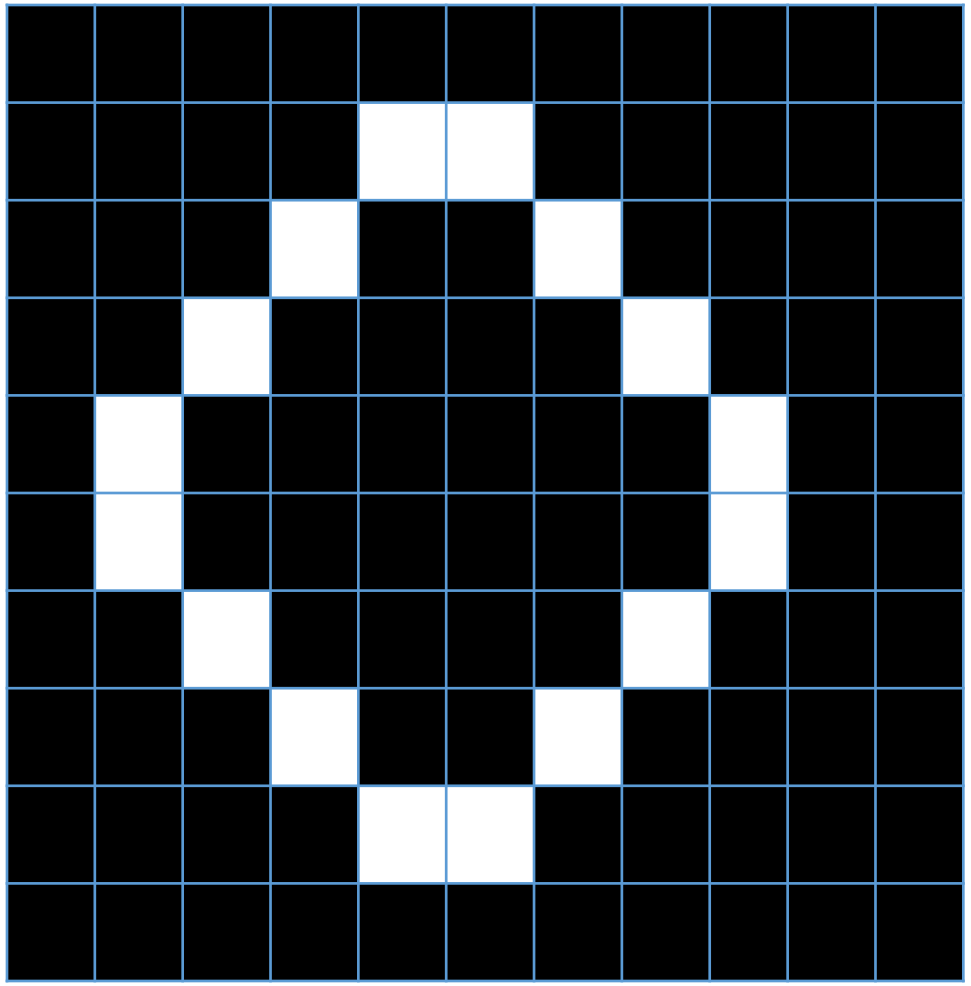

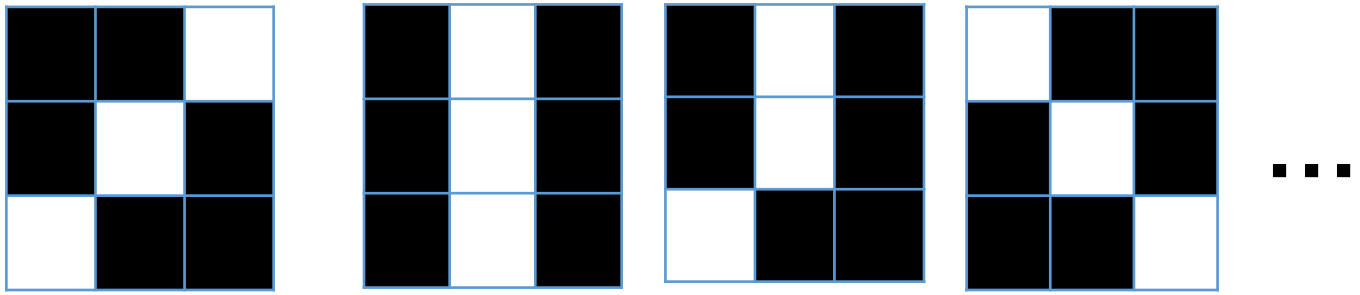

- Pattern starts at left corner
  - Perform comparison
  - Slide over one pixel
- Reach end of image
- Repeat for next pattern

Convolution Layers Search for Patterns

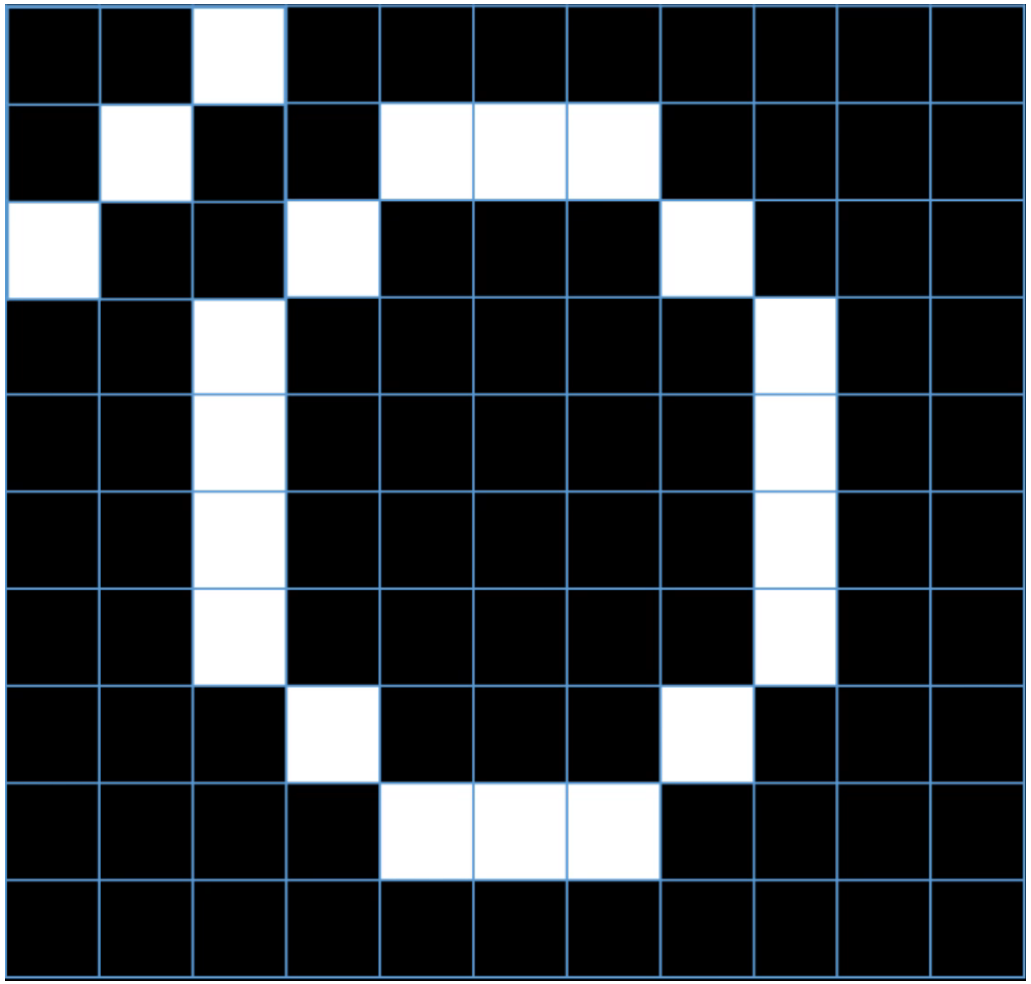

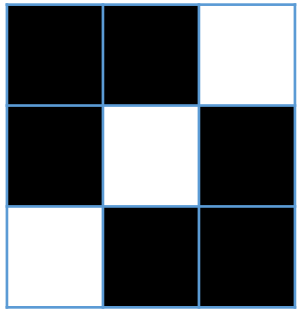

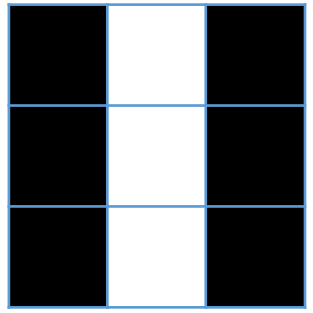

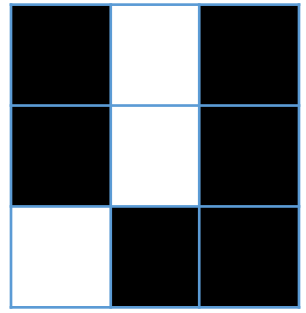

These patterns would be common in the number 0

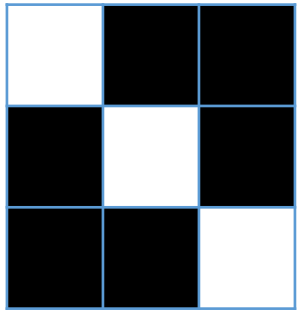

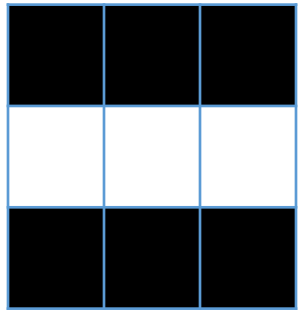

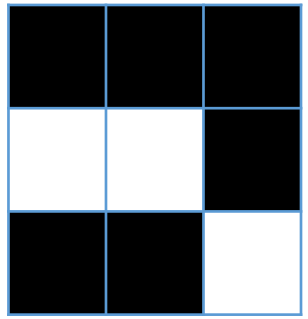

Good pattern matching in convolution improves chances that object will classify properly

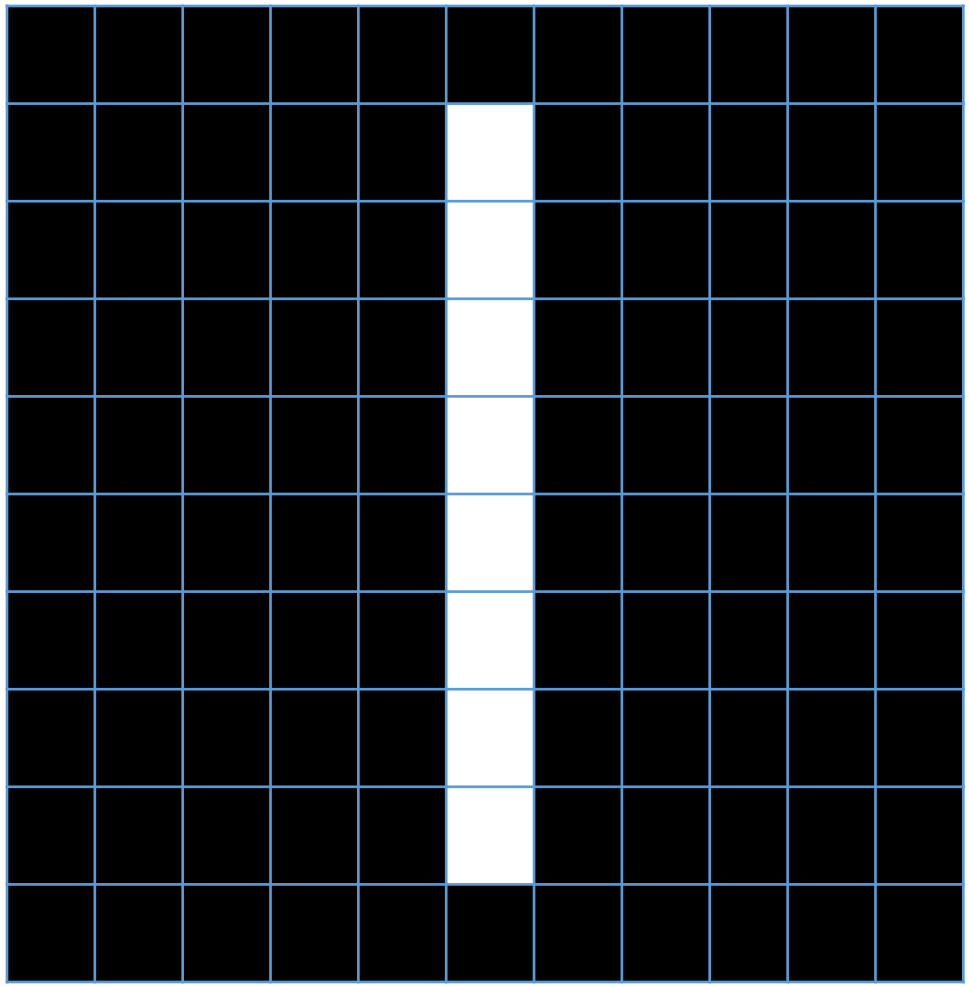

- This image would not match well against the patterns for the number zero
- It would only do very well against this pattern

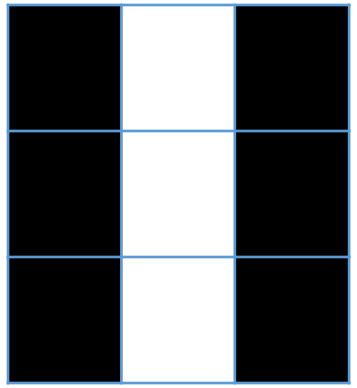

Max Pooling is a down-sampling operation

Shrink large images while preserving important information

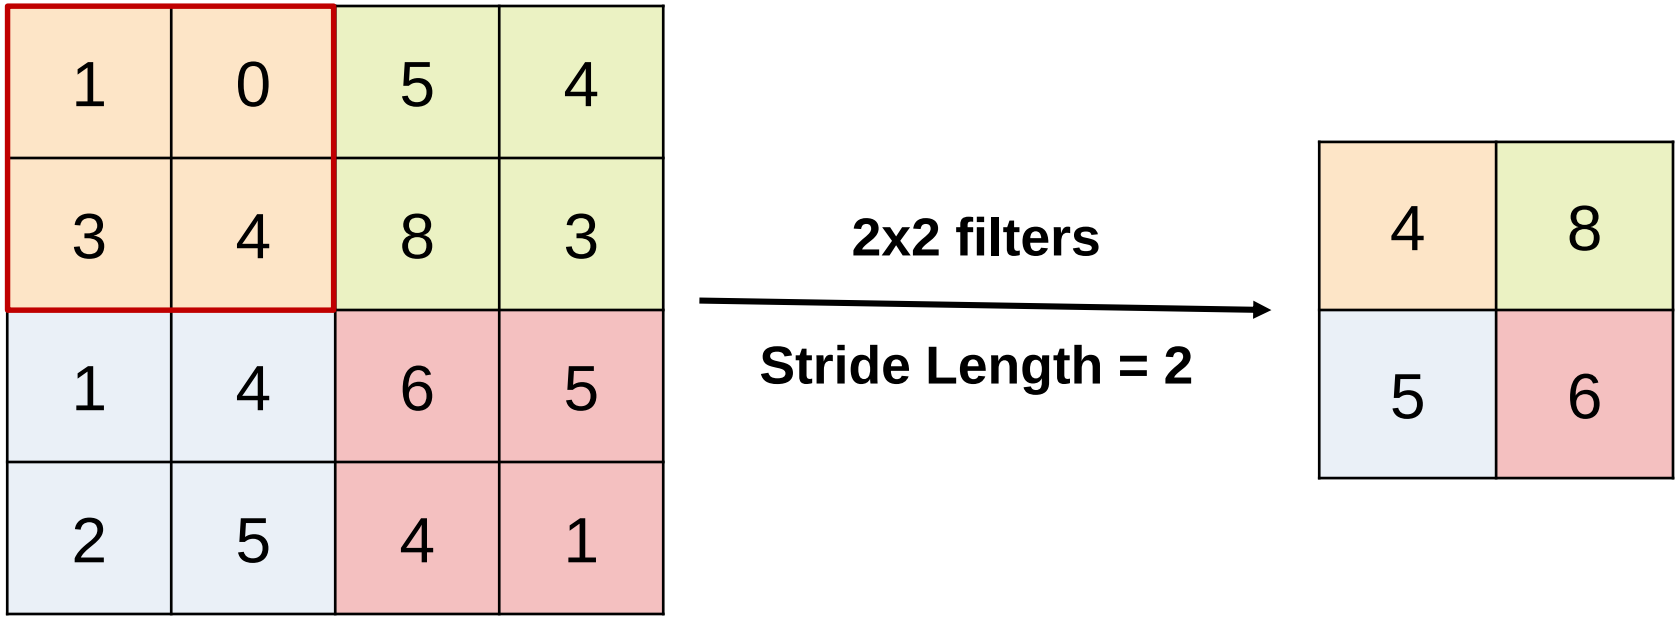

Rectified Linear Units Layer (ReLU)

Converts negative numbers to zero

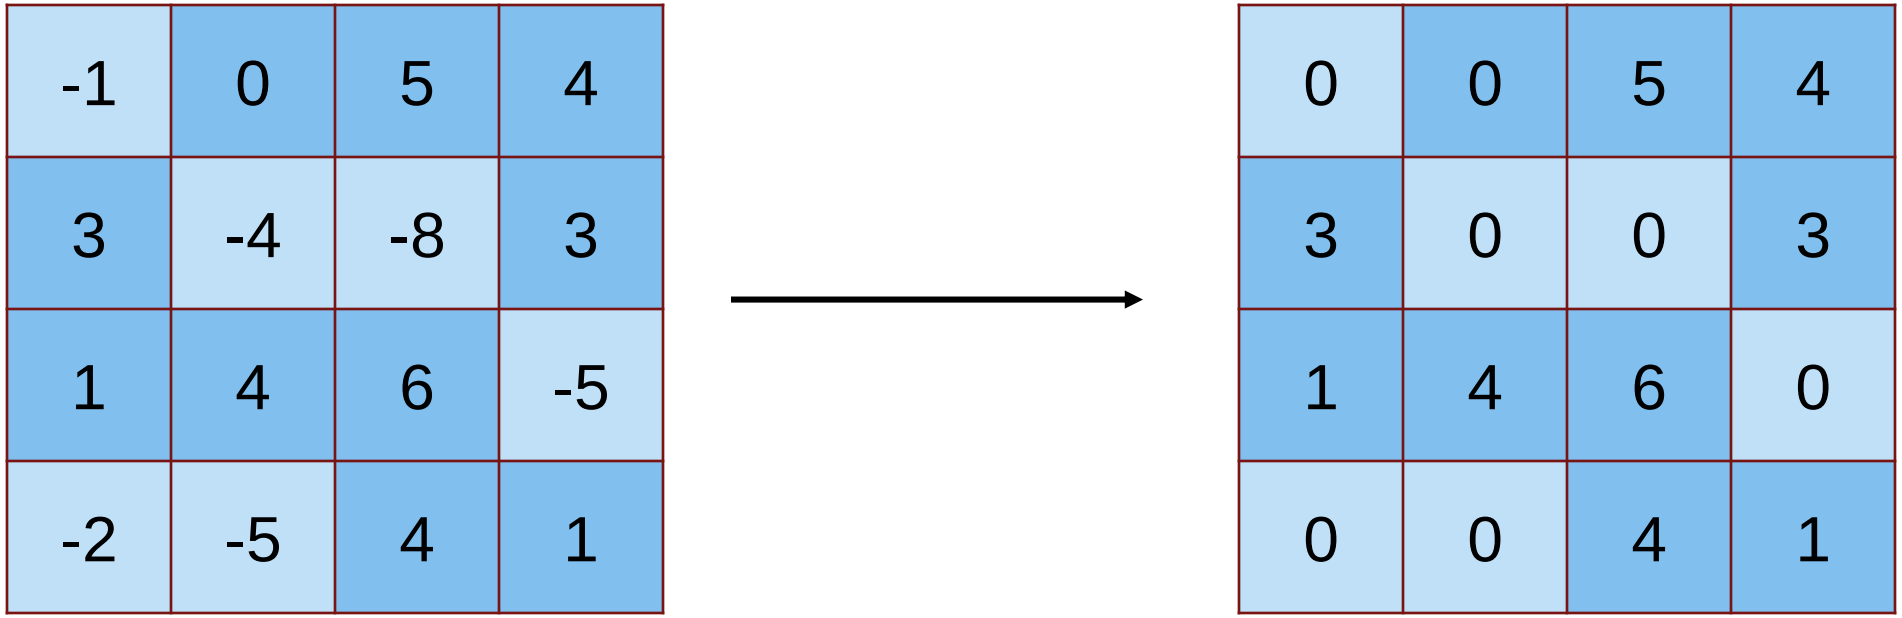

Classification Problems End with 3 Layers

- Fully Connected Layer
  - Looks at which high-level features correspond to a specific category
  - Calculates scores for each category (highest score wins)

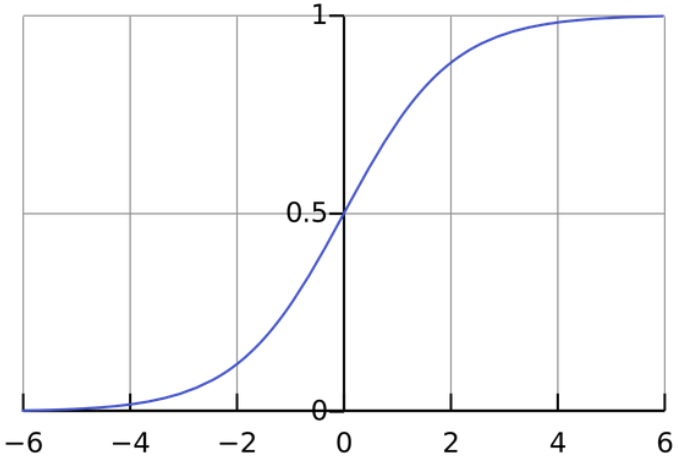

- Softmax Layer
  - Turns scores into probabilities.

- Classification Layer
  - Categorizes image into one of the classes that the network is trained on

Note: Regression problems end with a fully connected layer and regression layer

How Do I know Which Layers to Use?

Feature Extraction - Images

- 2D and 3D convolution
- Transposed convolution (…)

Activation Functions

- ReLU
- Tanh (…)

Sequence Data

Signal, Text, Numeric

- LSTM
- BiLSTM
- Word Embedding (…)

Normalization

- Dropout
- Batch normalization
- (…)

Tip: Research papers can provide guidelines for creating architecture

3 Components to Train any Network

"How much data do I need?"

It depends…but **A LOT**

Define inputs and layers for deep learning model

Influence training time and accuracy

- Solver Type
- Initial Learn Rate
- Minibatch Size
- Max Epochs
- …

Exercise 3 - MNIST

Purpose:

- Learn how to create and train deep neural network
- Use MATLAB's Deep Network Designer
- Explore hyperparameters

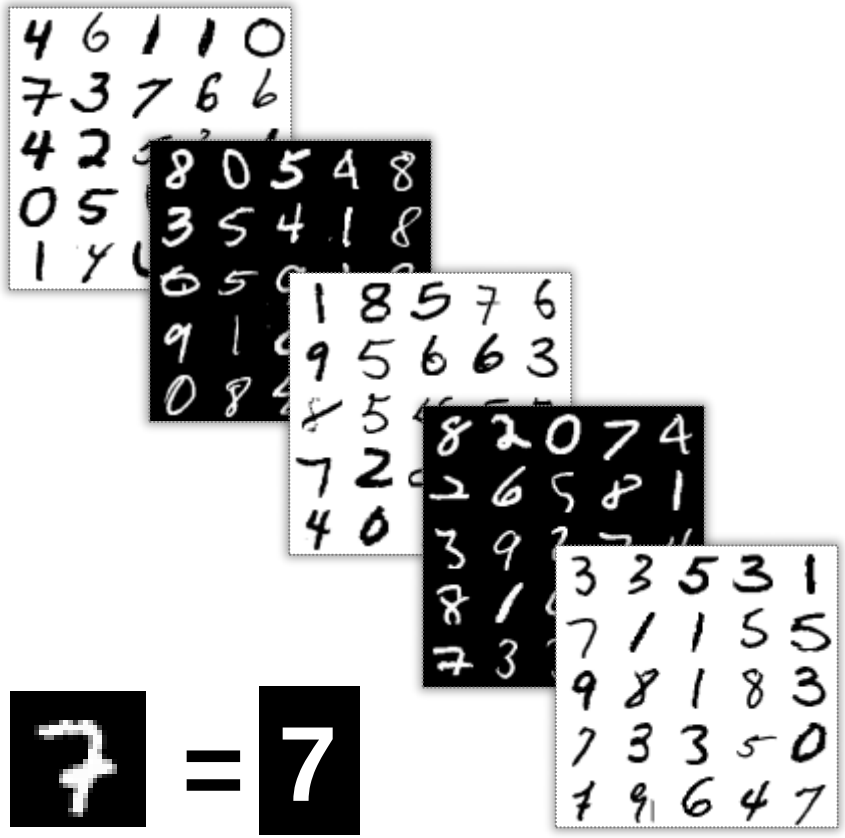

Details

- Dataset consists of handwritten digits 0-9
- 60,000 training images
- 10,000 test images

Sources: *http://yann.lecun.com/exdb/mnist/*
https://rodrigob.github.io/are_we_there_yet/build/classification_datasets_results

Network to Create for Part 1 of MNIST

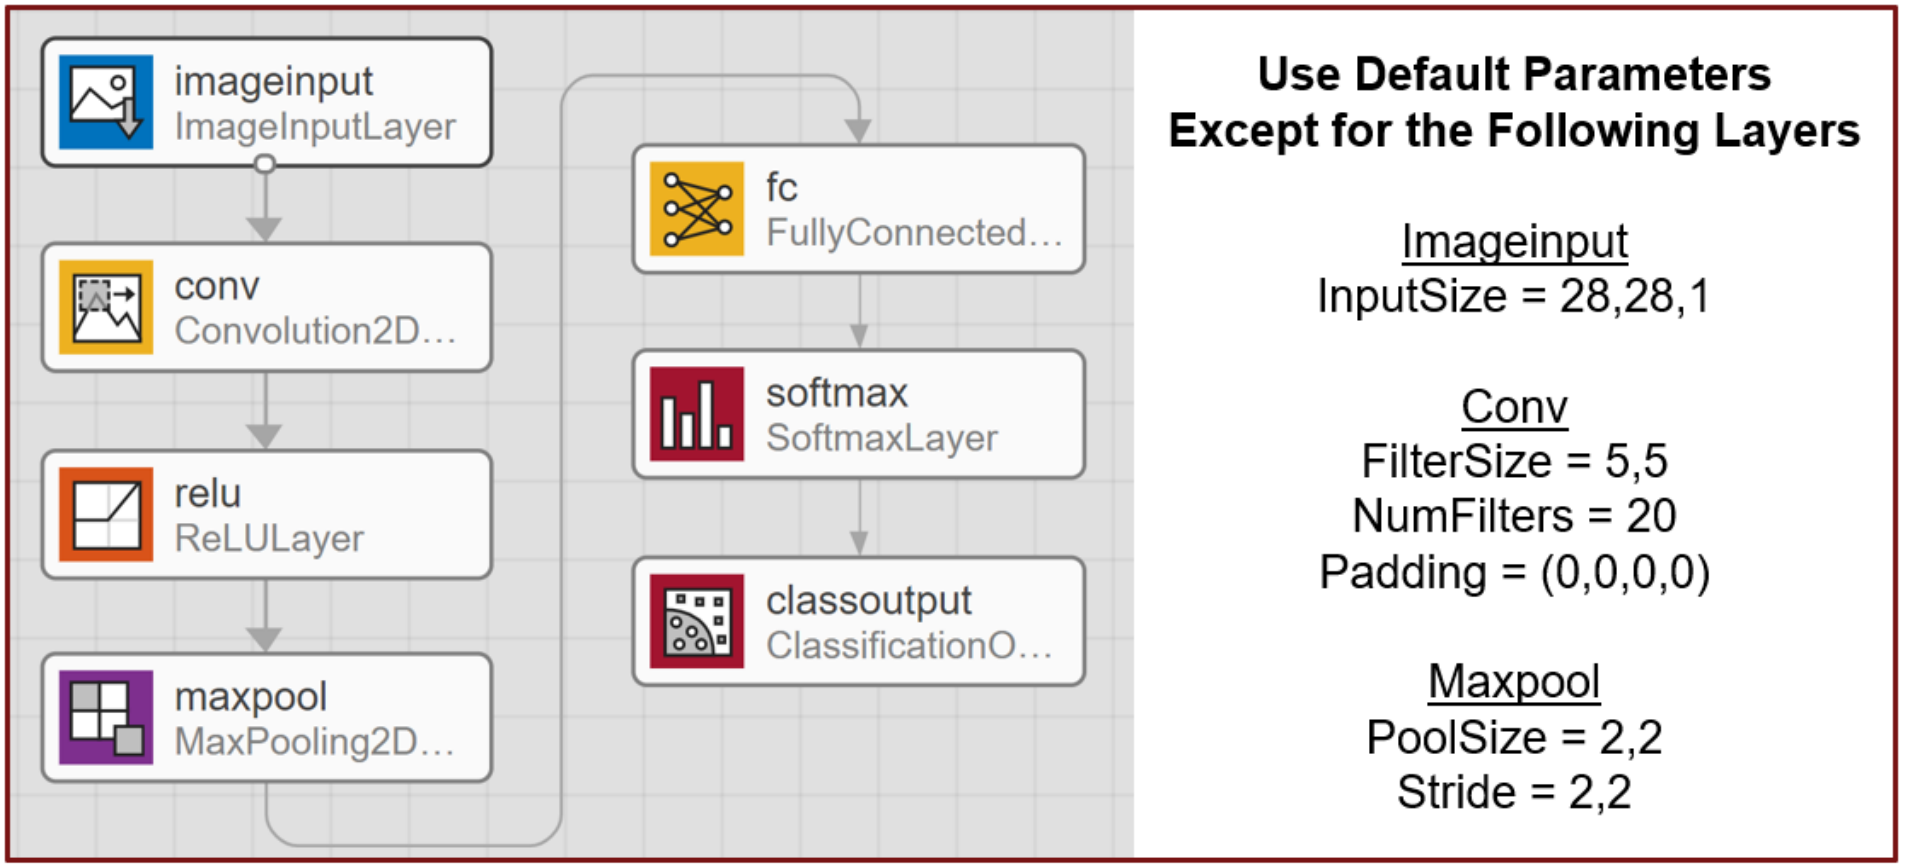

Use Default Parameters Except for the Following Layers

Imageinput
InputSize = 28,28,1

Conv
FilterSize = 5,5
NumFilters = 20
Padding = (0,0,0,0)

Maxpool
PoolSize = 2,2
Stride = 2,2

Transfer Learning Workflow

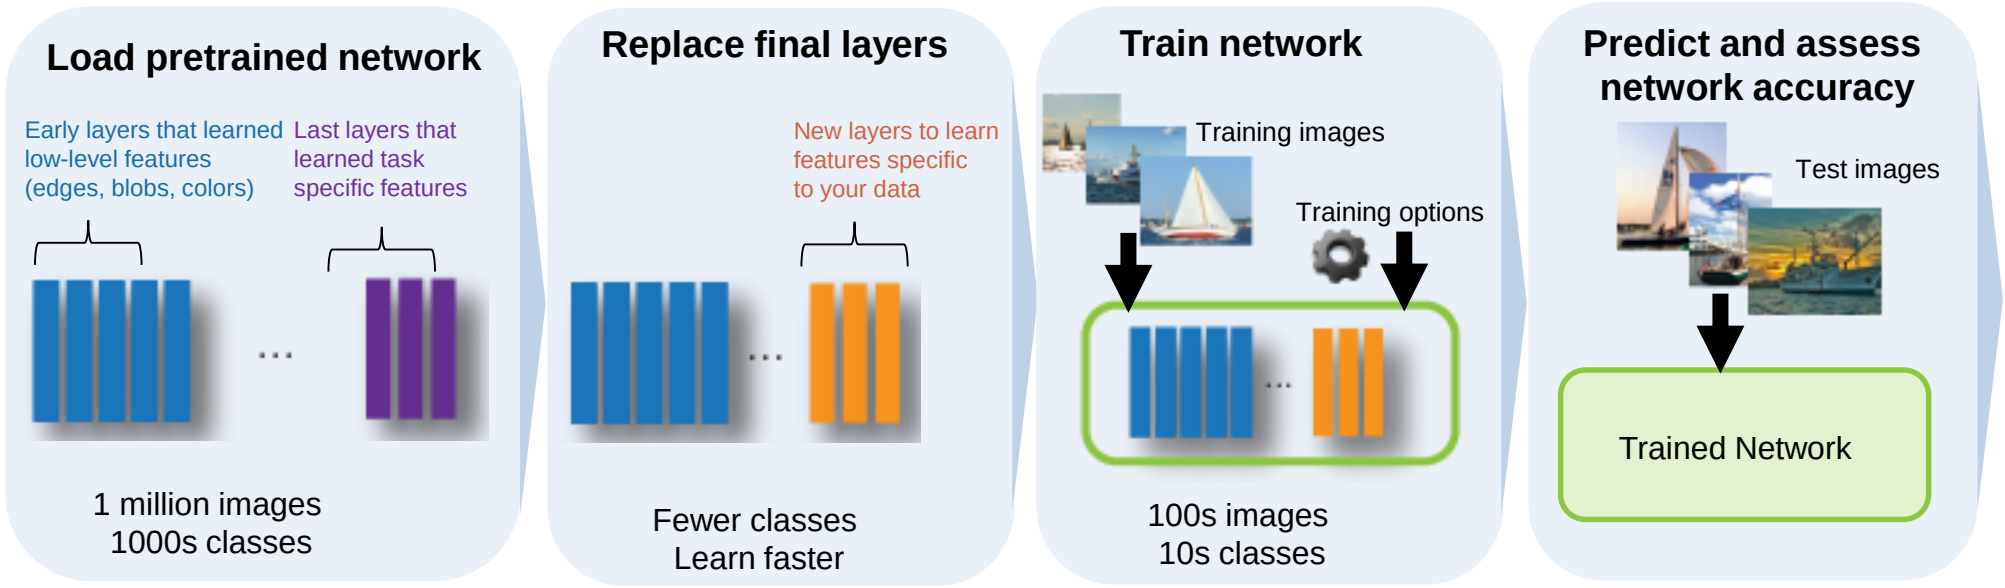

Transfer Learning Workflow – Step 1

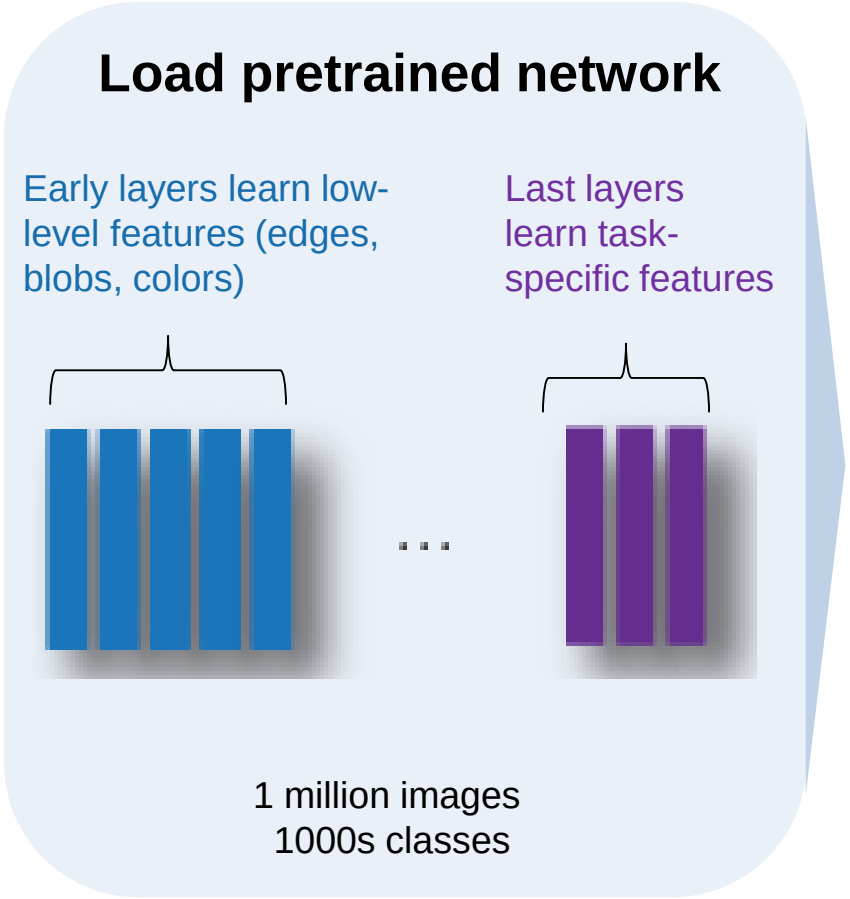

Transfer Learning Workflow – Step 2

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

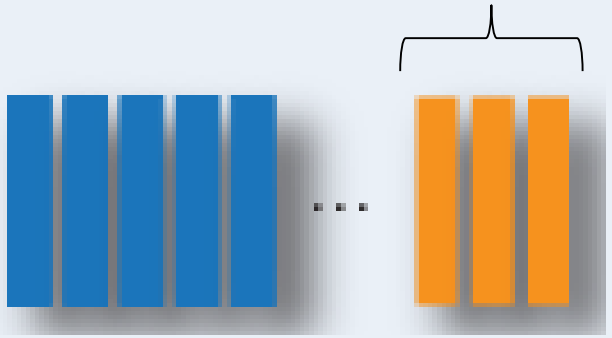

Fewer classes
Learn faster

Transfer Learning Workflow – Step 3

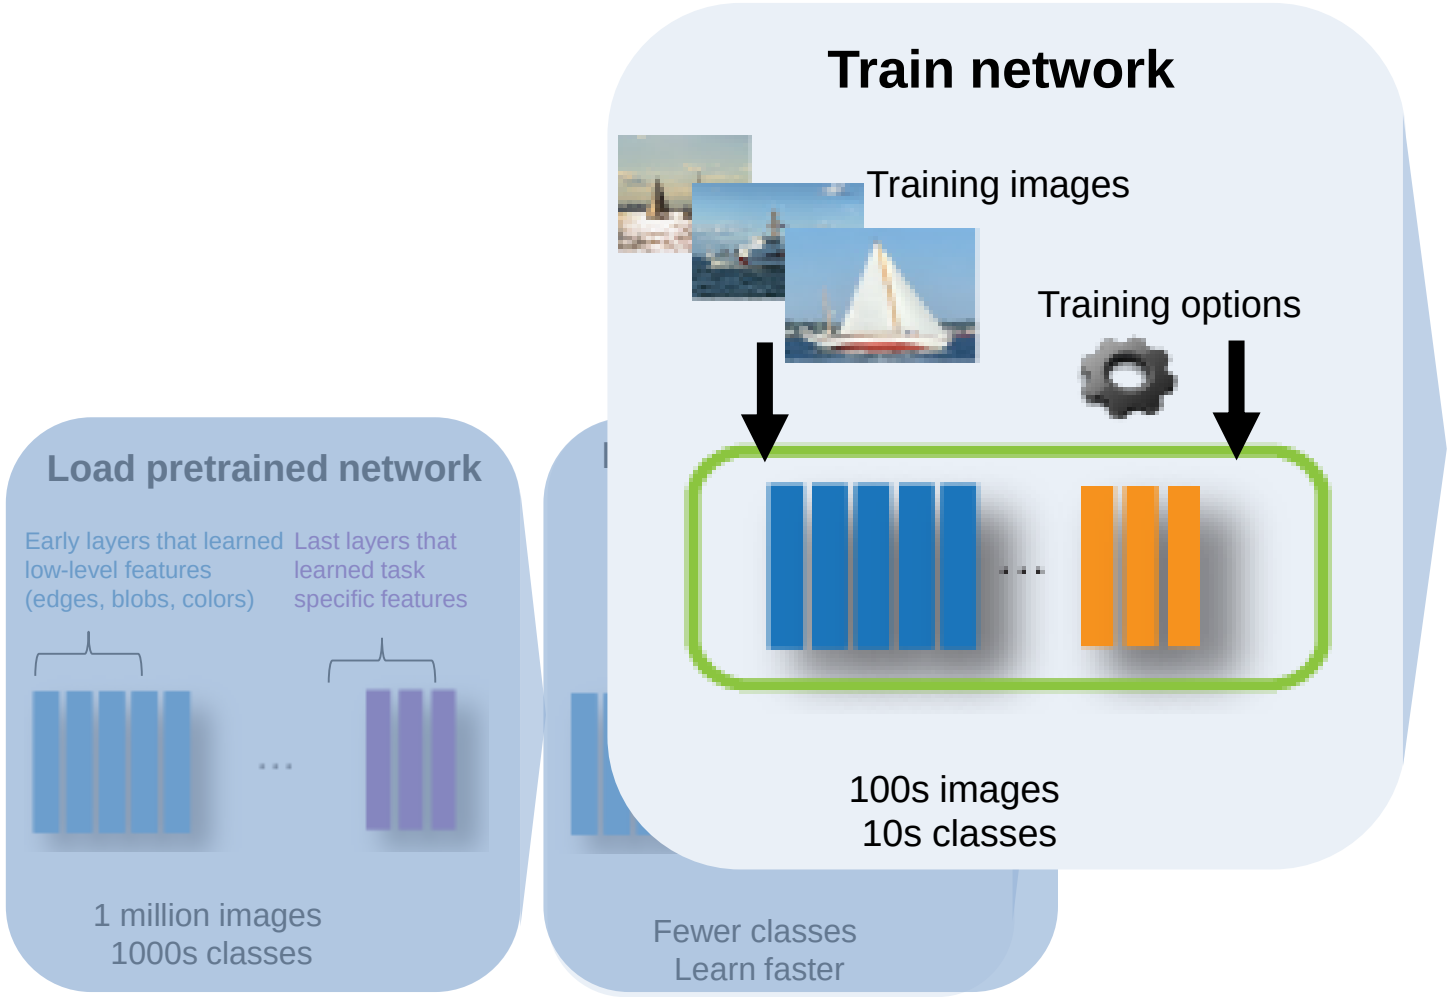

Transfer Learning Workflow – Step 4

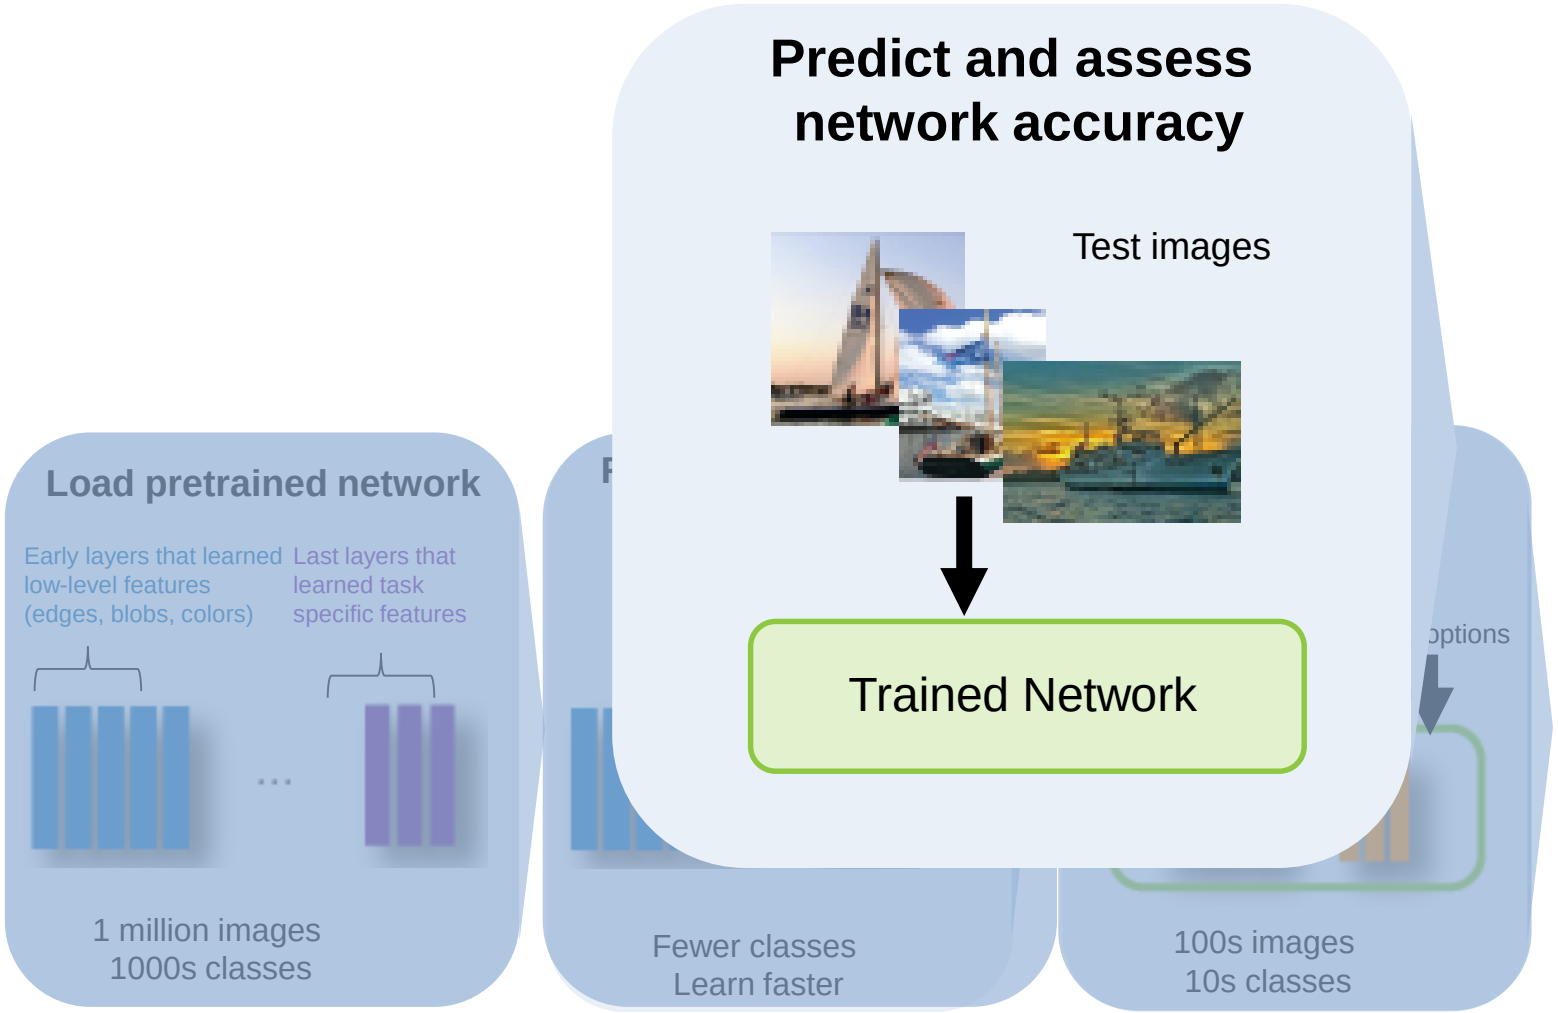

Transfer Learning Workflow

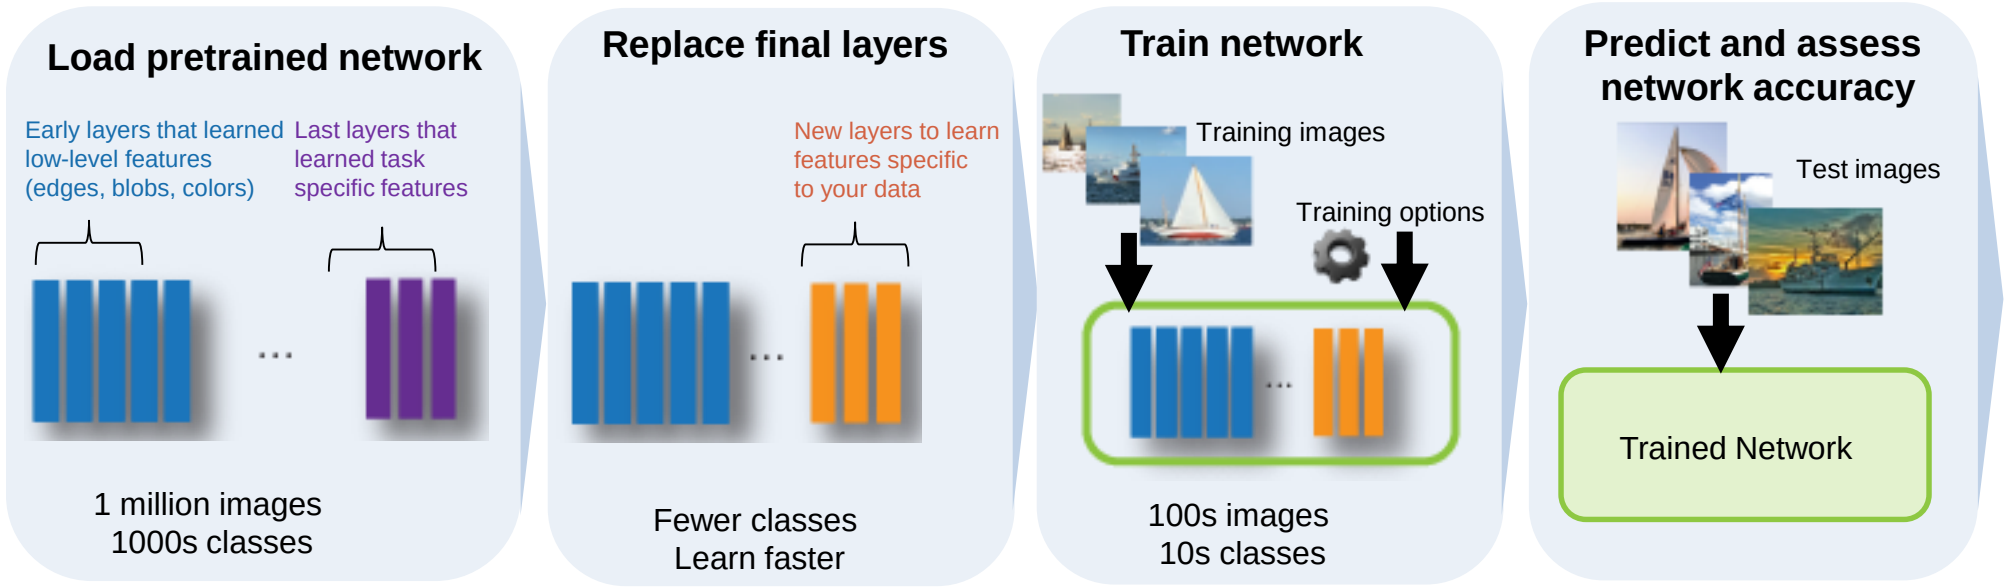

Exercise 4 – Transfer Learning

Purpose:

- Use transfer learning to leverage a pretrained model to classify 5 types of food
- Visualize activations within a network





Questions to consider once you understand training

- How do I prepare a dataset for training and testing?
- How do I create ground truth data?
- How do I deploy to hardware or the cloud?

Deep Learning Workflow – Prepare Data

PREPARE DATA

- Data access and preprocessing
- Simulation-based data generation
- Ground truth labeling

TRAIN MODEL

- Model design and tuning
- Hardware-accelerated training
- Model exchange across frameworks

DEPLOY SYSTEM

How do I label my data?

Image Labeler
+ Video labeler

Signal Labeler
+ Audio Labeler

Big-Image
Labeler

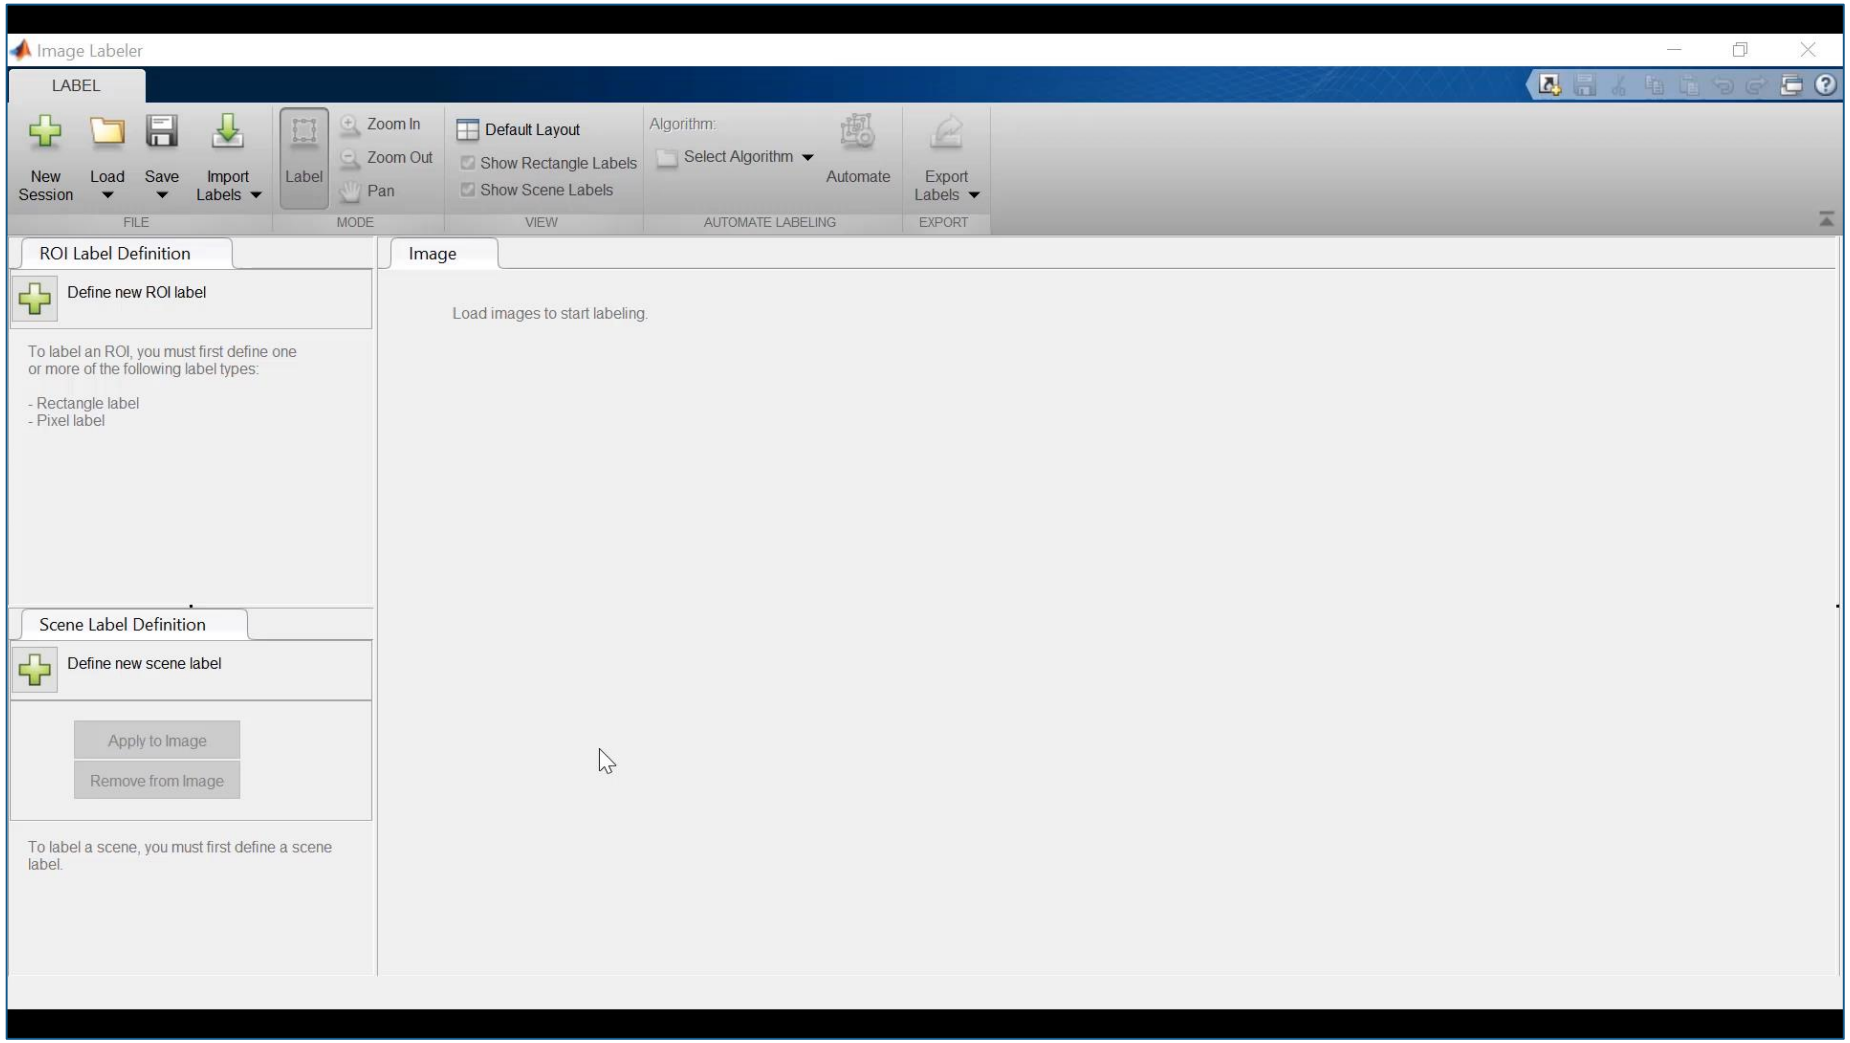

How do I label my data?

Image Labeler
+ Video labeler

Signal Labeler
+ Audio Labeler

Big-Image
Labeler

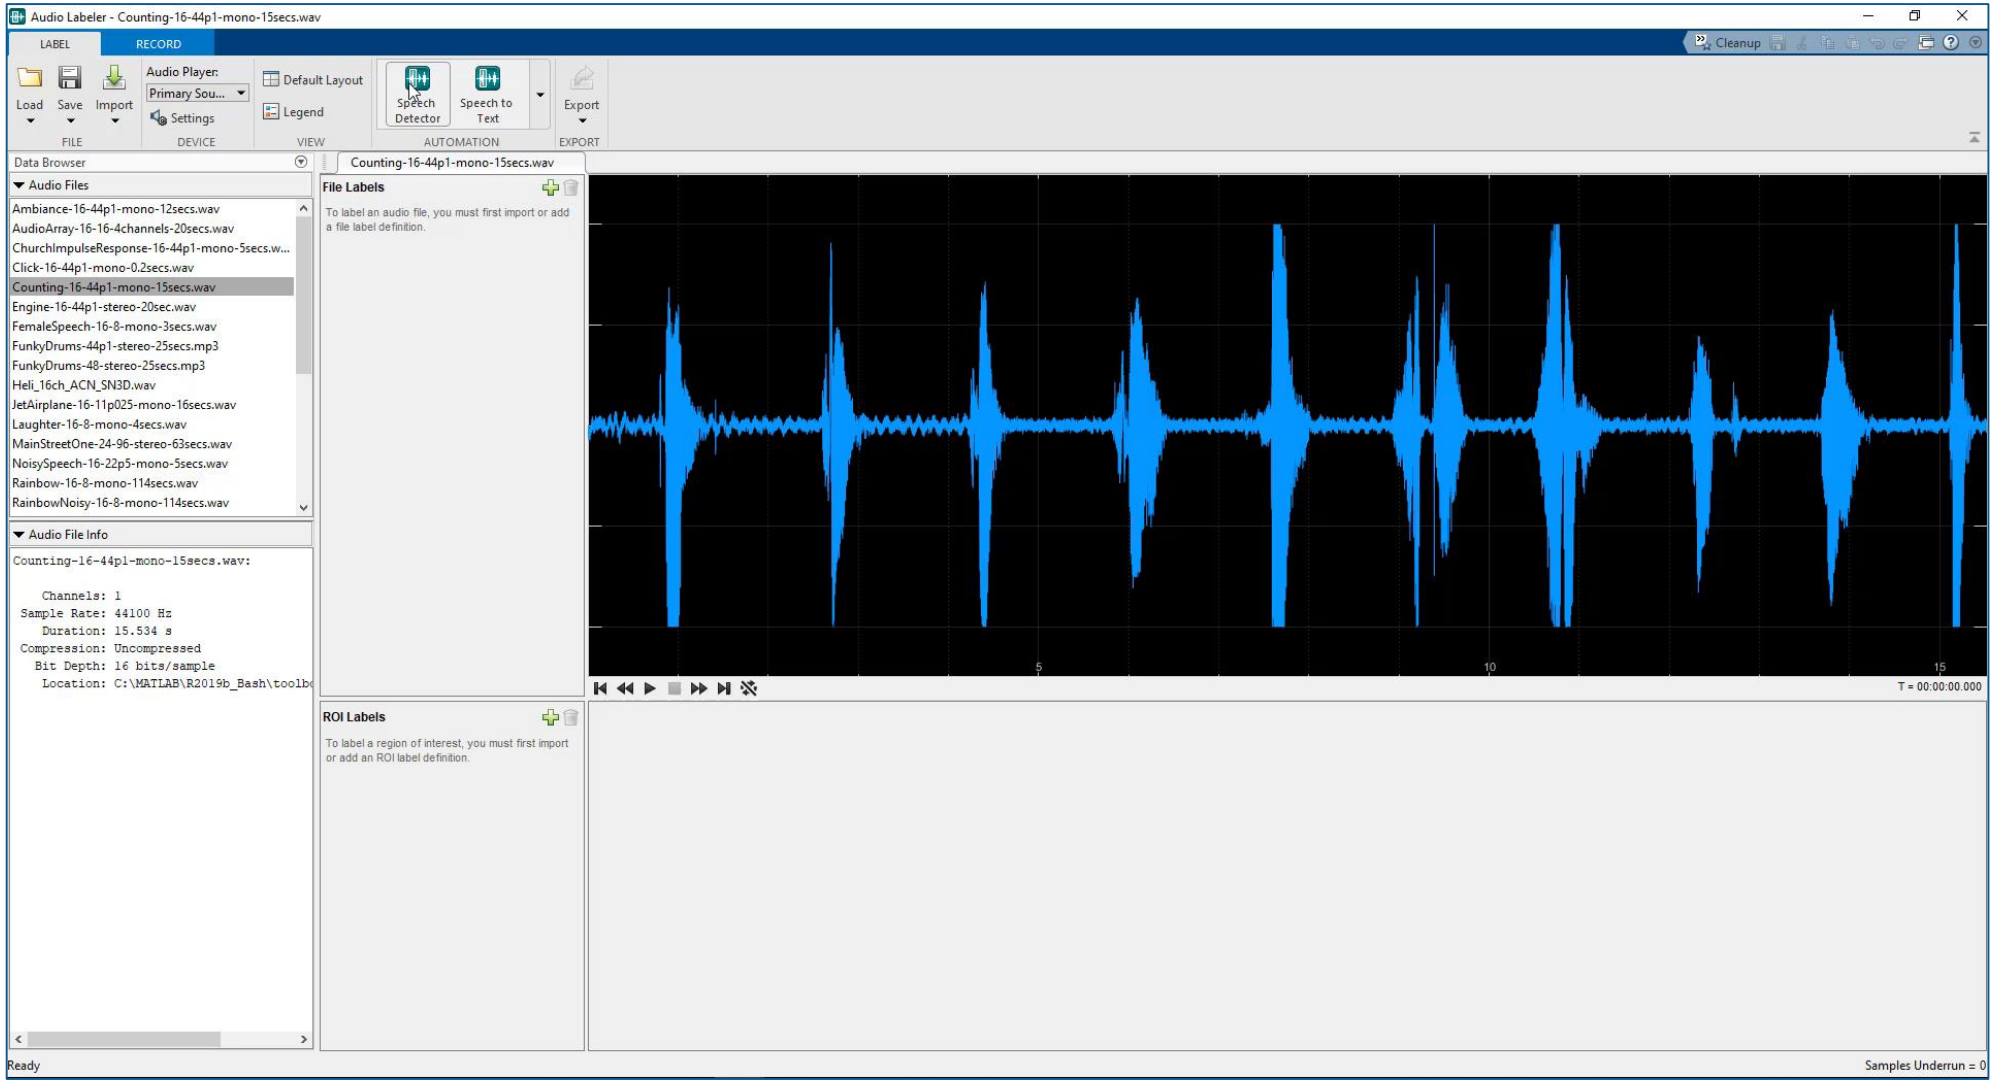

How do I label my data?

Image Labeler
+ Video labeler

Signal Labeler
+ Audio Labeler

Big-Image
Labeler

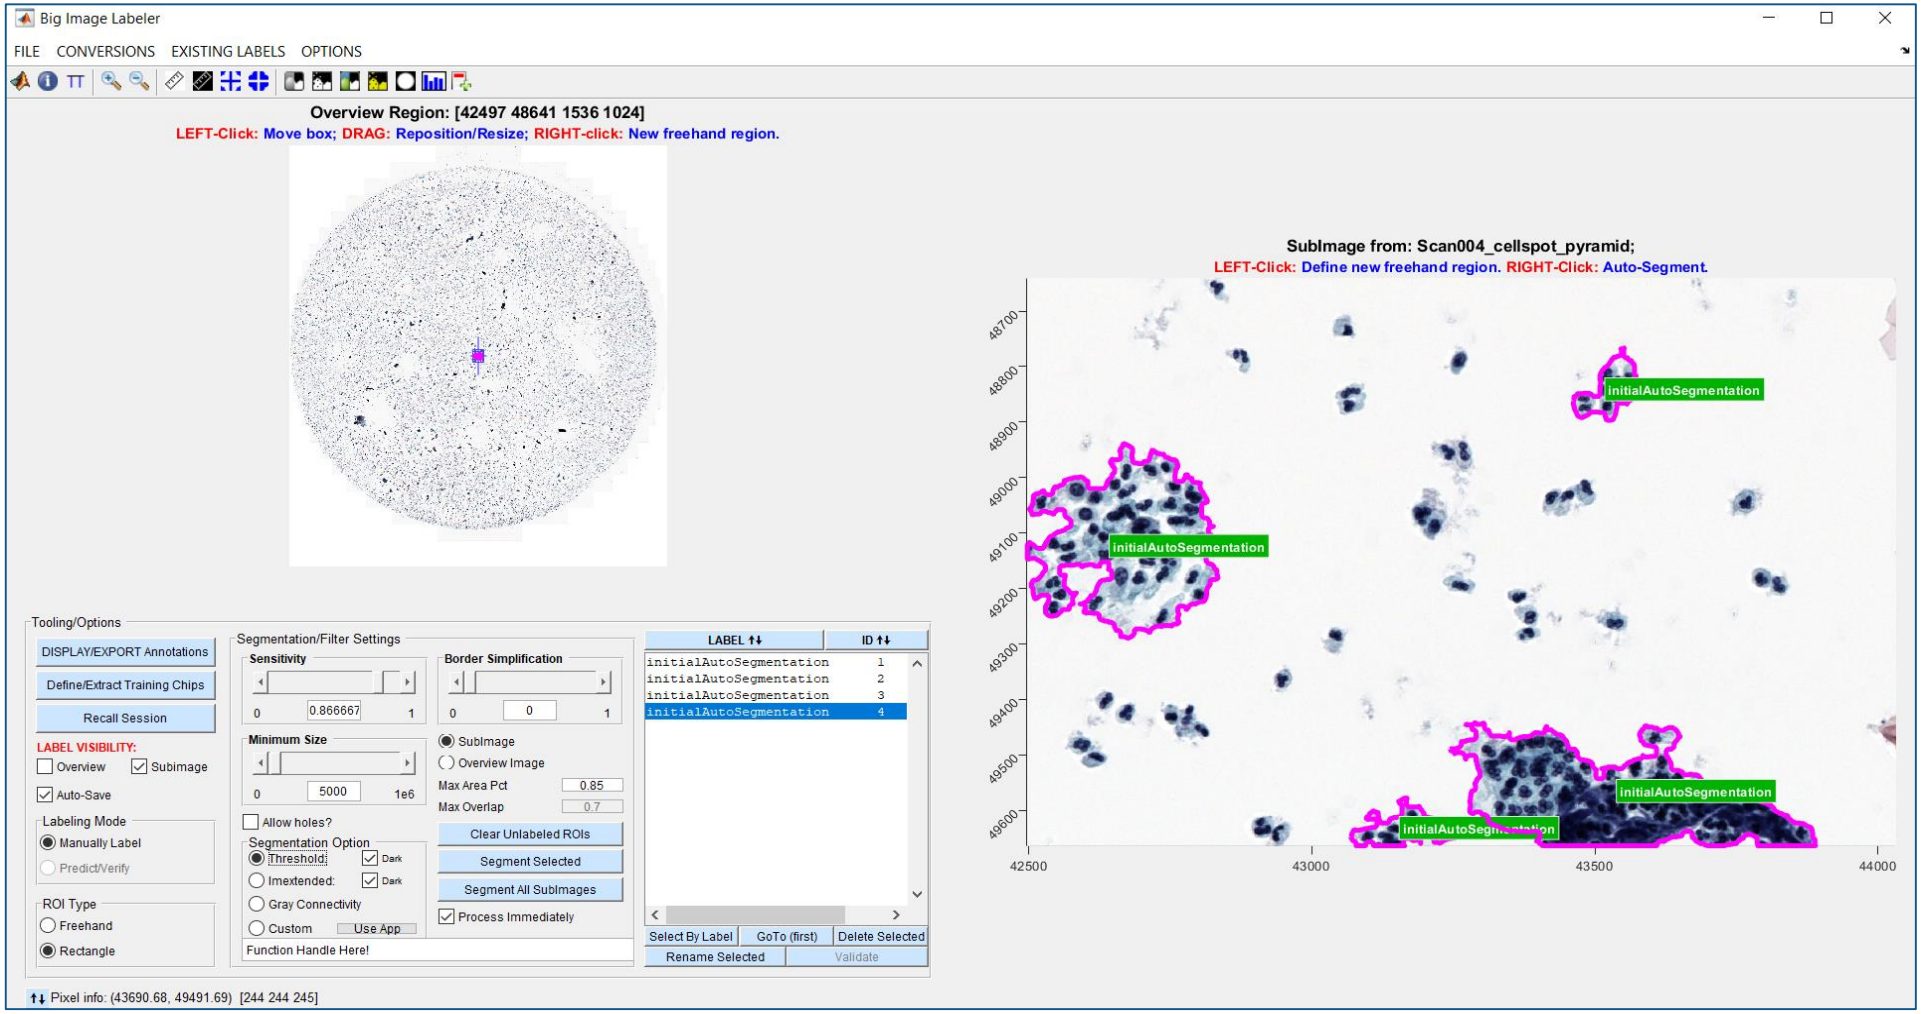

Deep Learning Workflow – Deploy System

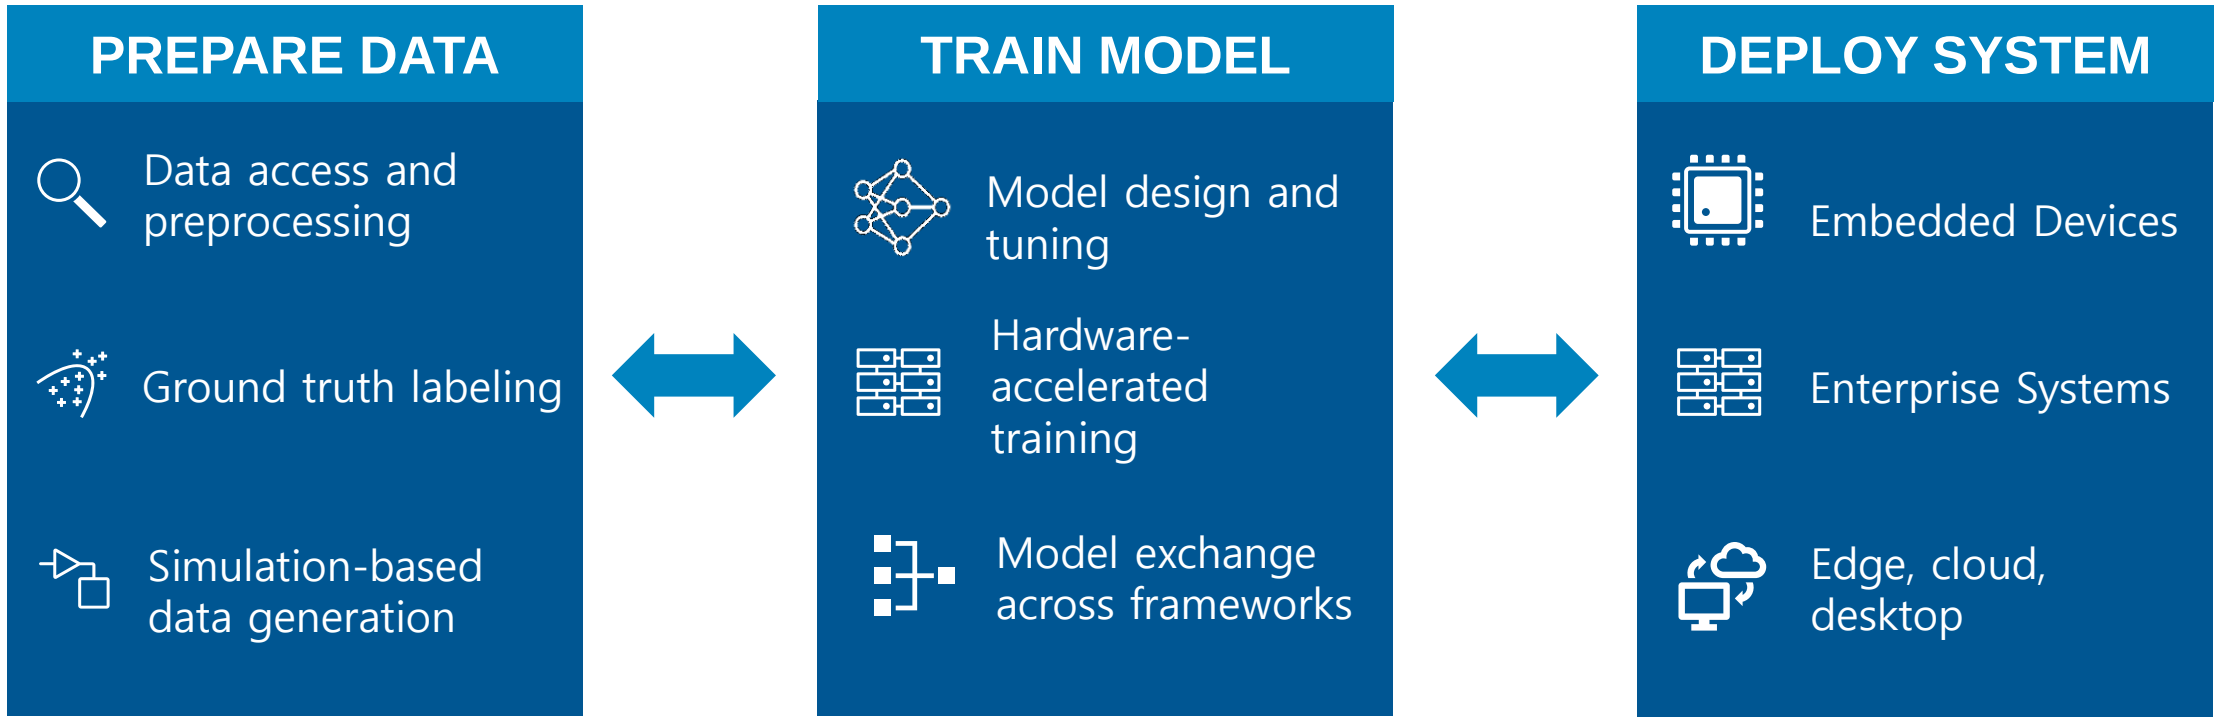

Deployment and Scaling for A.I.

MATLAB

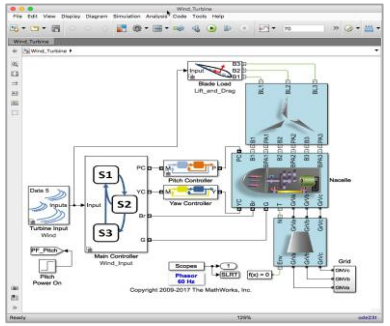

Embedded Systems

CPU GPU FPGA

Enterprise Systems

Embedded Deployment – Automatic Code Generation

MATLAB Code

Auto-generated Code (C/C++/CUDA)

Deployment Target

Deploying Models for Inference

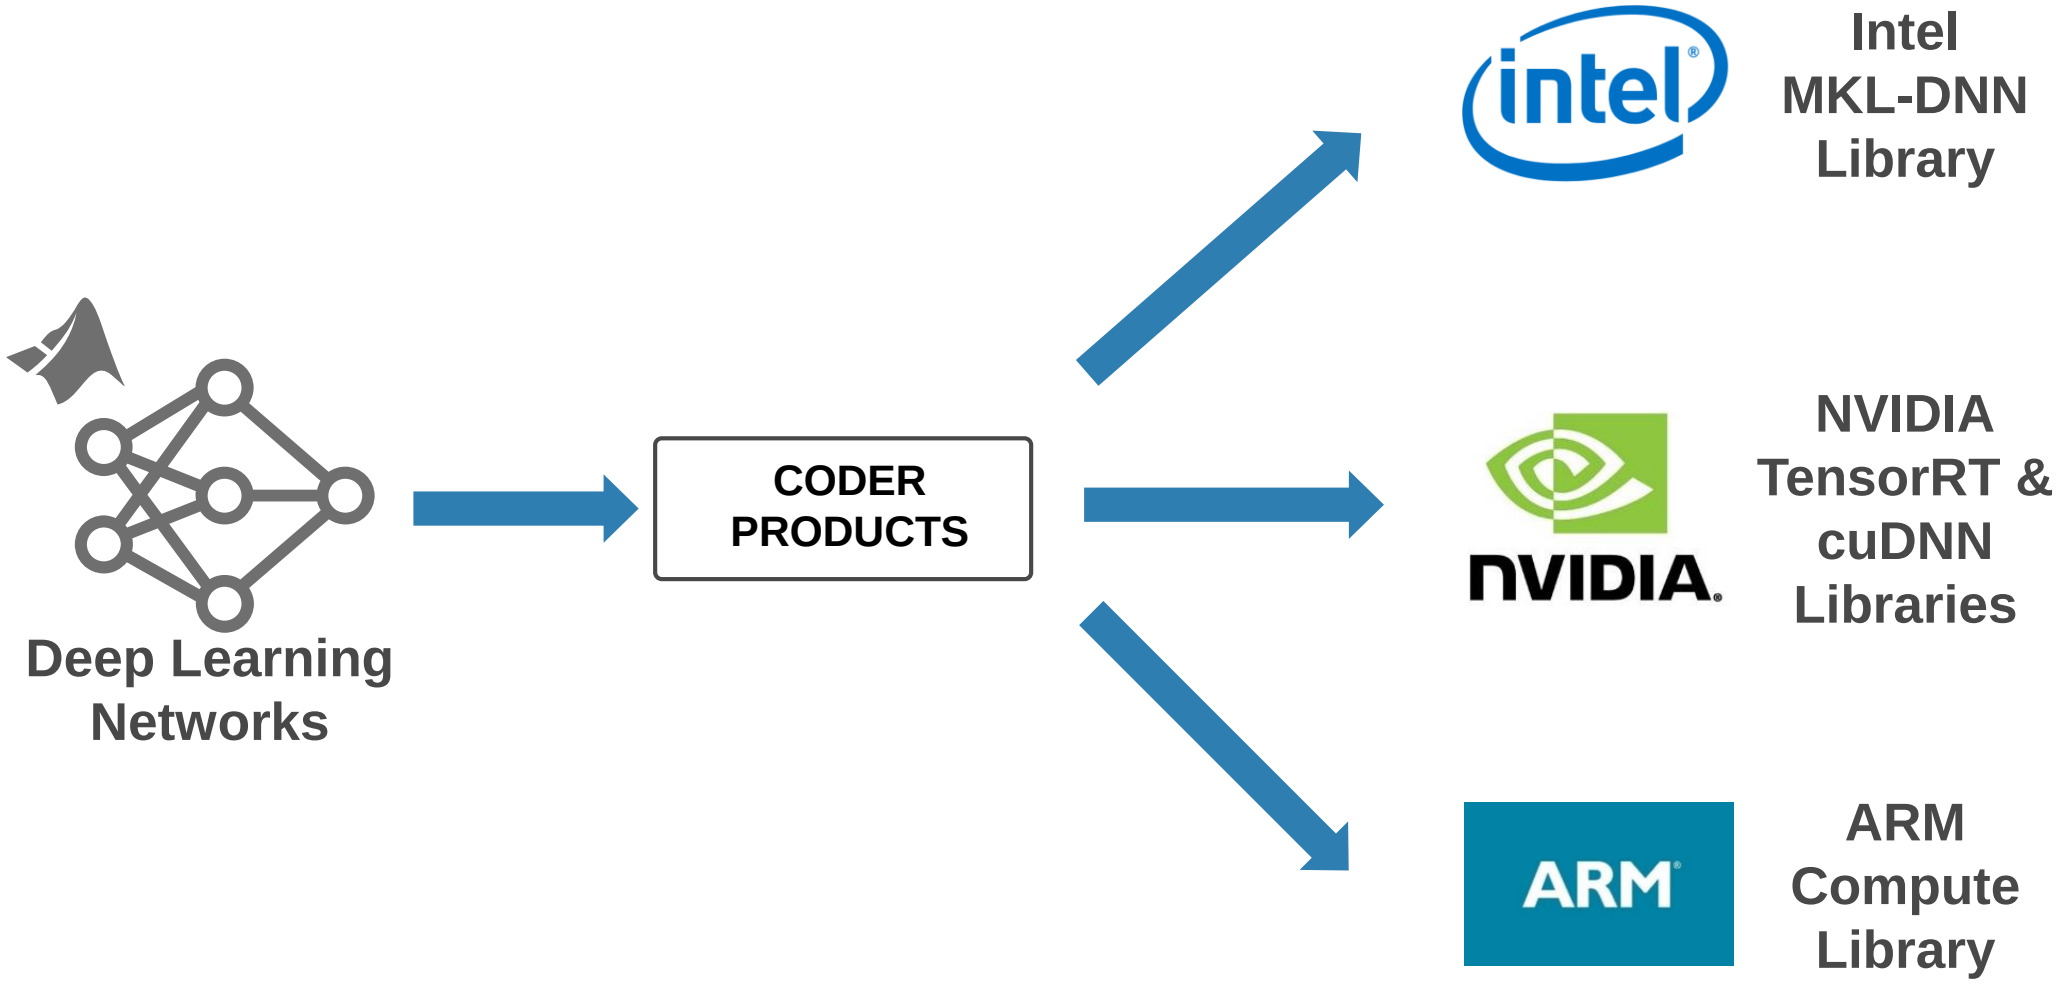

NVIDIA, the NVIDIA logo, and TensorRT are registered trademarks of NVIDIA Corporation
Intel logo is a registrered trademark of Intel Coroporation

GPU Coder is over 2x Faster Than Other Compiled Frameworks

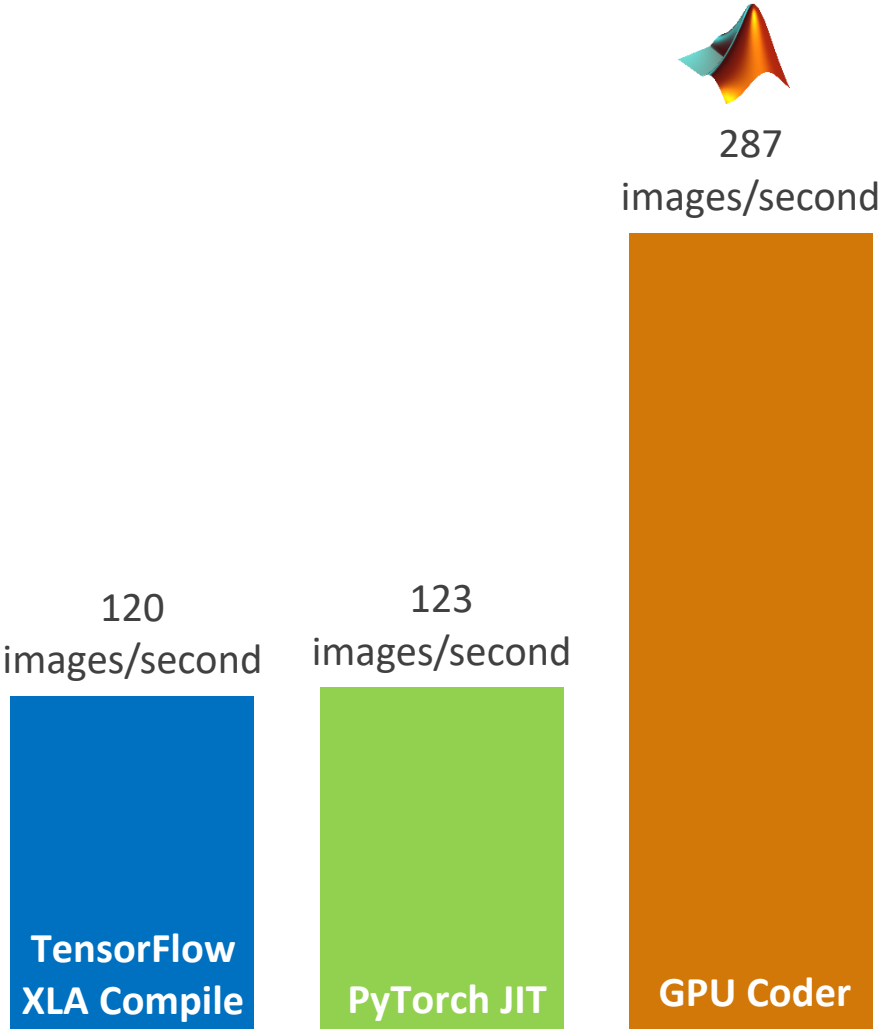

Intel® Xeon® CPU 3.6 GHz with NVIDIA® Titan V GPU - NVIDIA libraries: CUDA® – cuDNN - Frameworks: TensorFlow™ 1.13.0, PyTorch 1.1.0 – ResNet-50 pre-trained model

Deployment to the Cloud and Data Centers with MATLAB Compiler and MATLAB Production Server

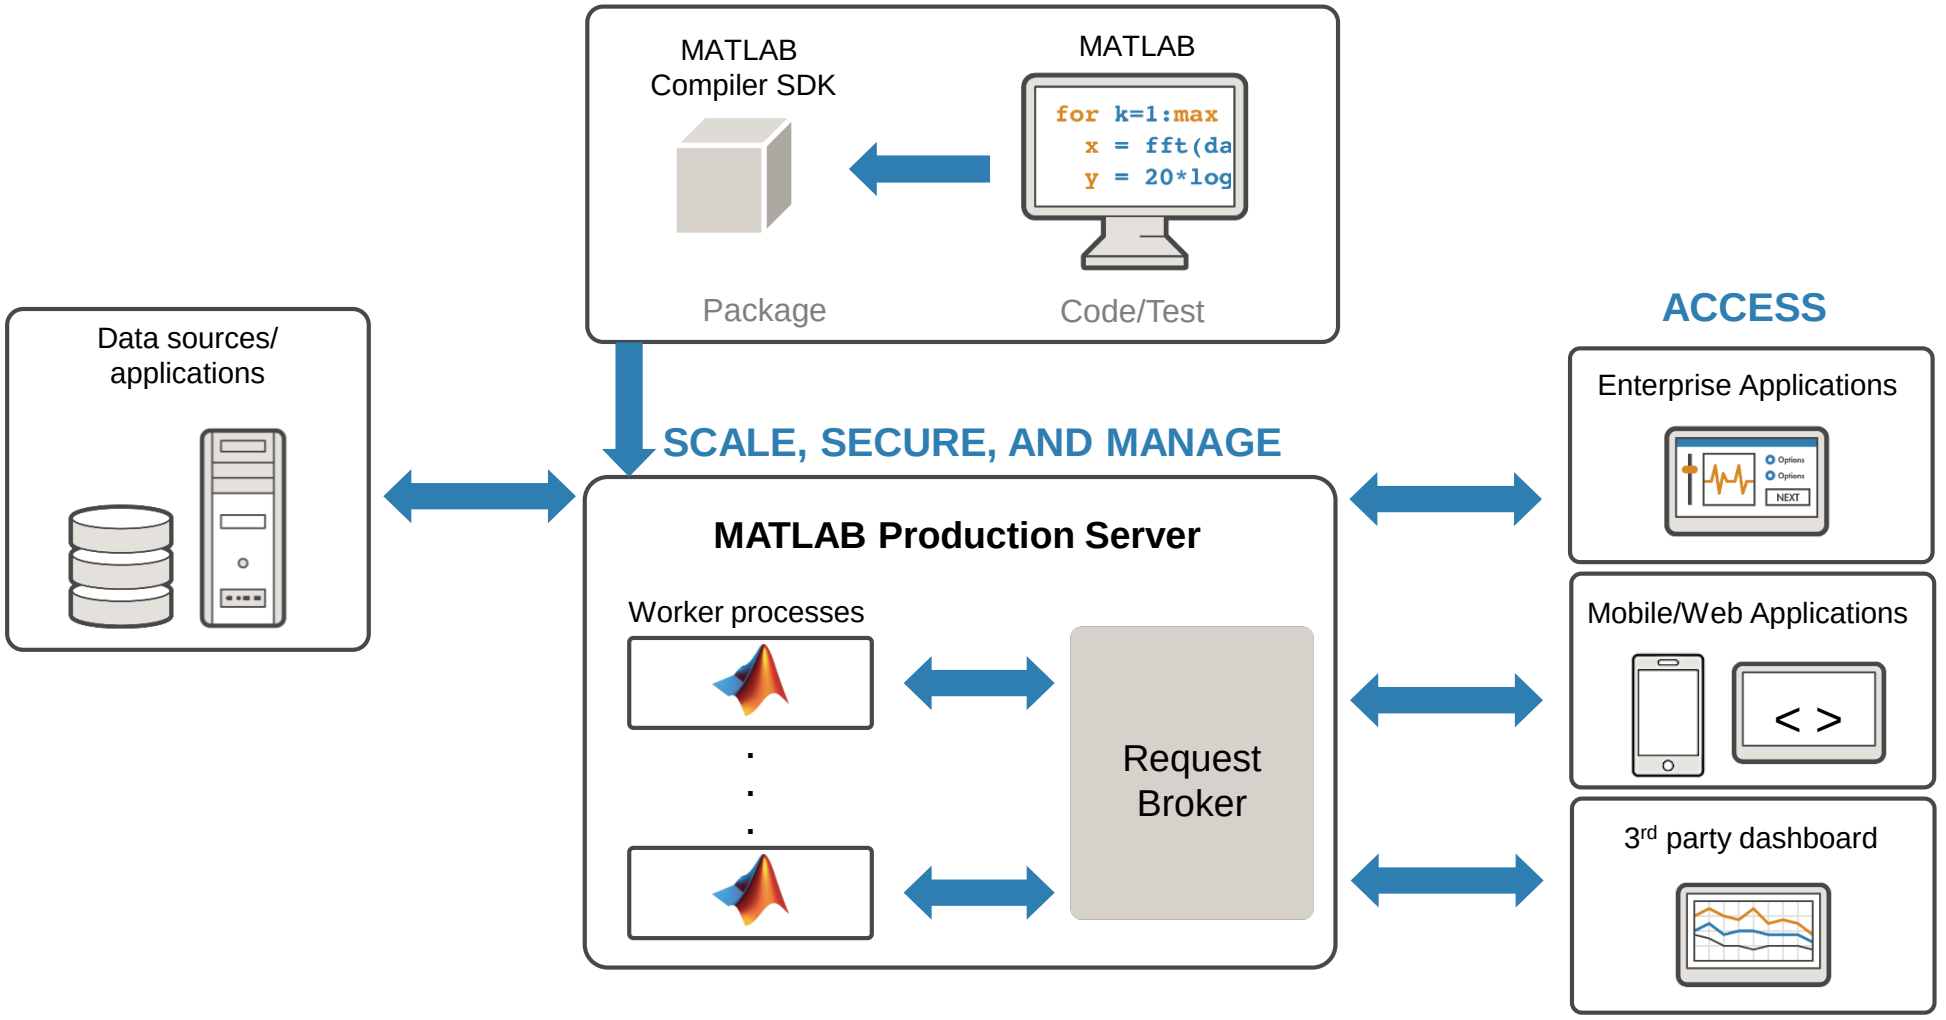

GPU Code Generation with ONNX and YOLOv2

YOLOv2

- **Y**ou **O**nly **L**ook **O**nce
- Real-time object detector
- 1000x faster than R-CNN
- Autonomous driving, traffic monitoring

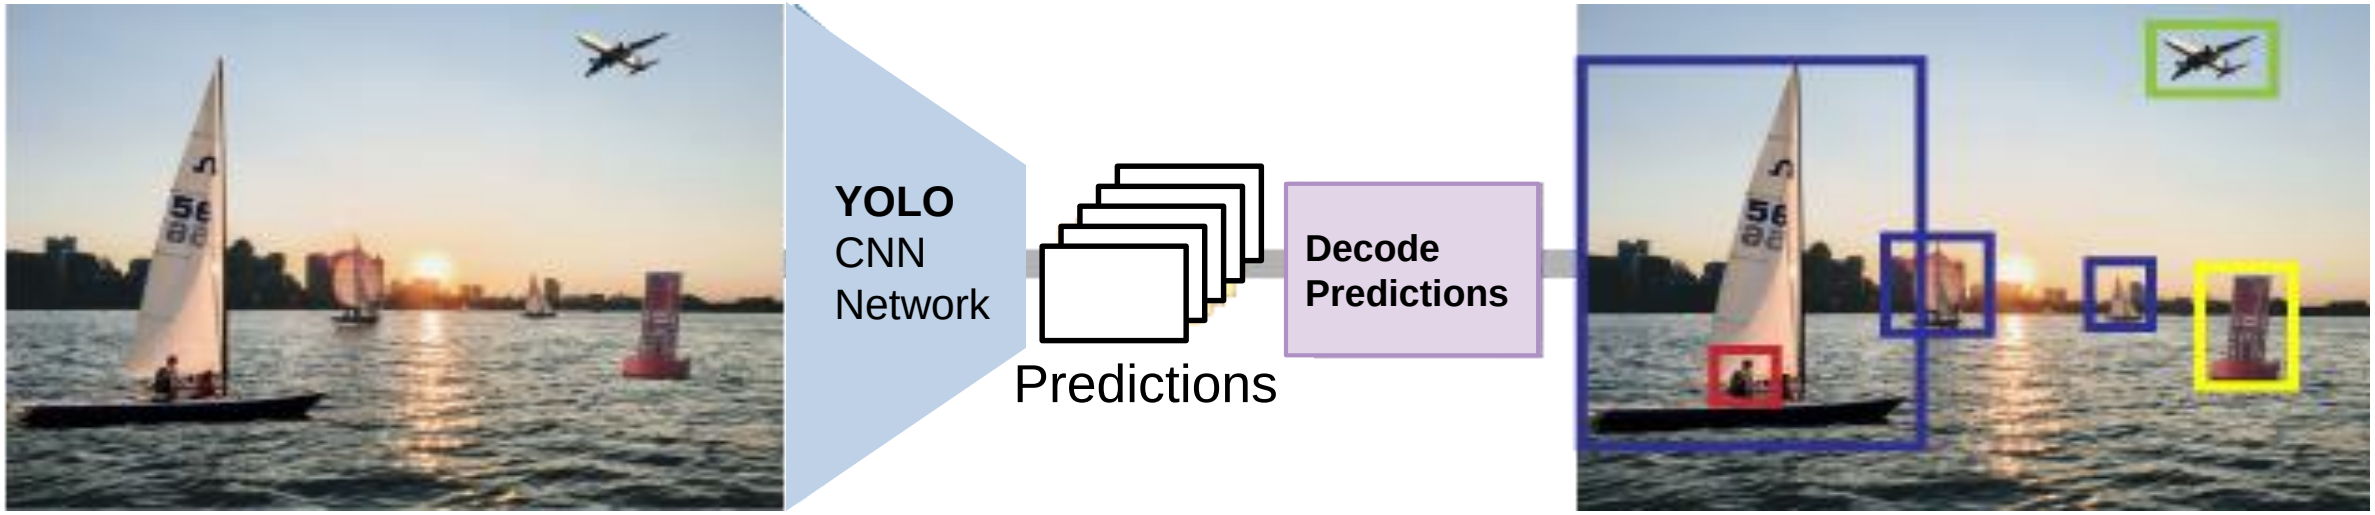

Why Use MATLAB

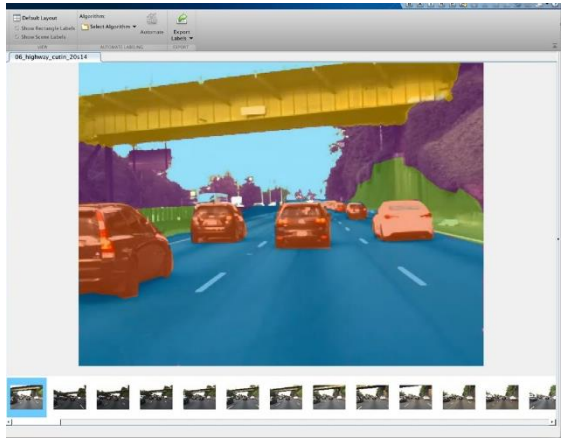

MATLAB supports the **data preparation, training, and deployment** workflow

MATLAB has specialized DL tools designed for **scientists and engineers**

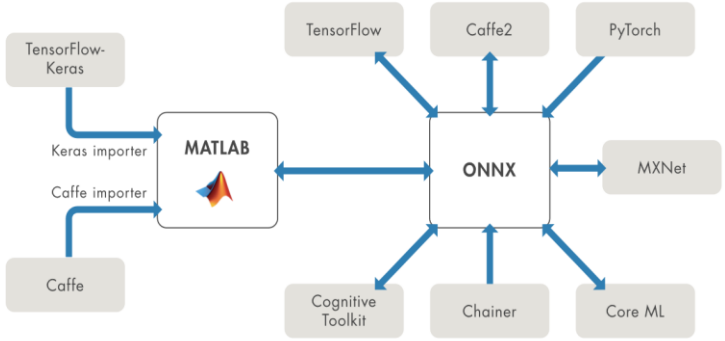

MATLAB **interoperates and enhances** Open Source frameworks

Common Network Architectures - Signal Processing

Convolutional Neural Networks (**CNN**)

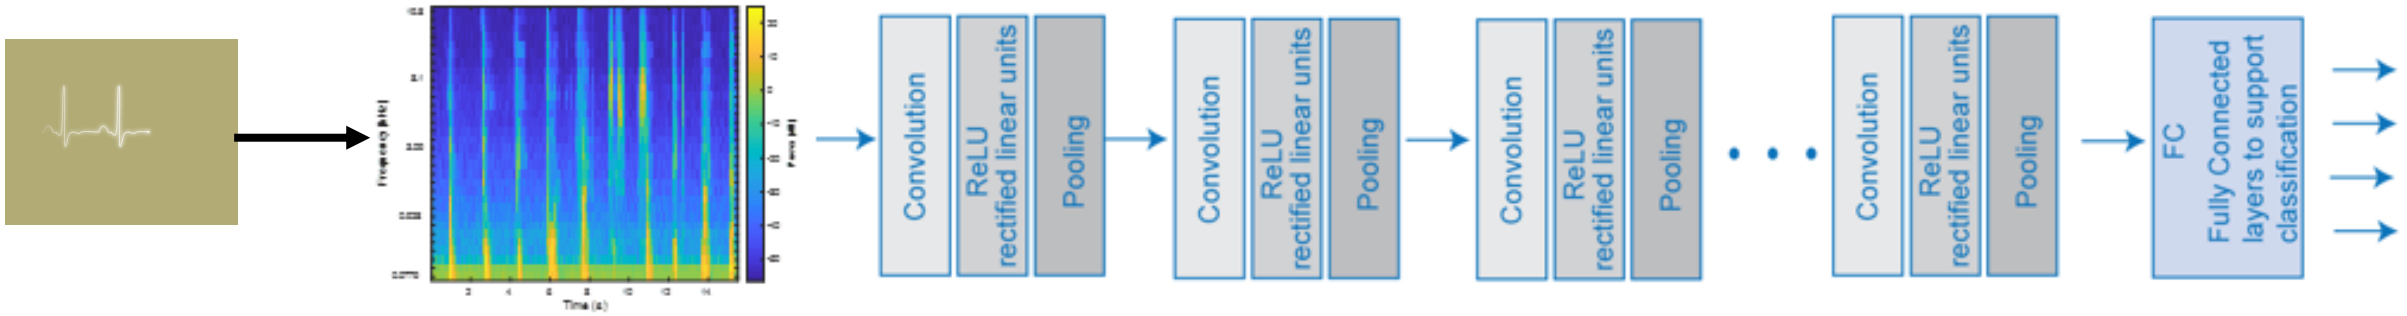

Time-Frequency Transformation

Long Short Term Memory (**LSTM**) Networks

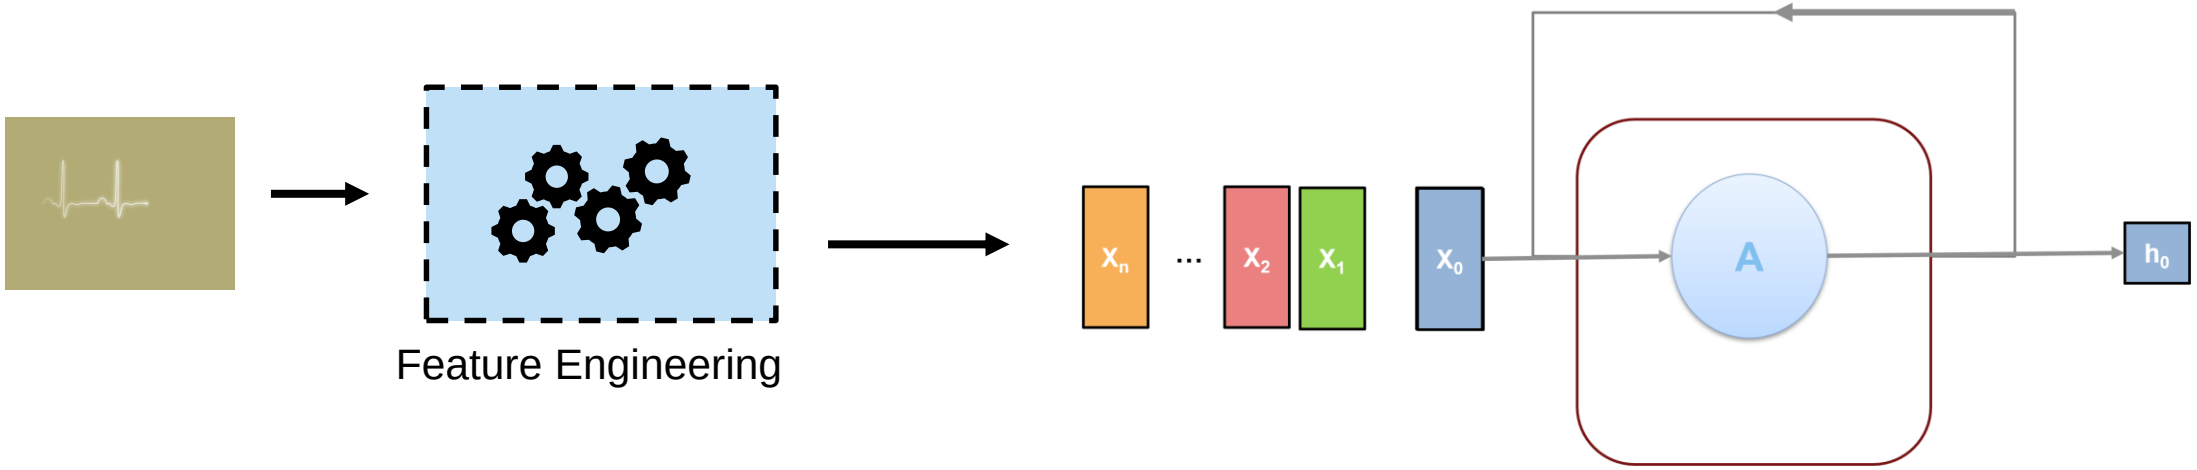

Feature Engineering

Selecting a Network Architecture

Image Data

CNN

Signal or Text Data

LSTM or CNN

LSTM = Long Short Term Series Network (more detail in later slides)

I was born in France…

… I speak _____ ?

Recurrent Neural Networks

Take into account previous data when making new predictions

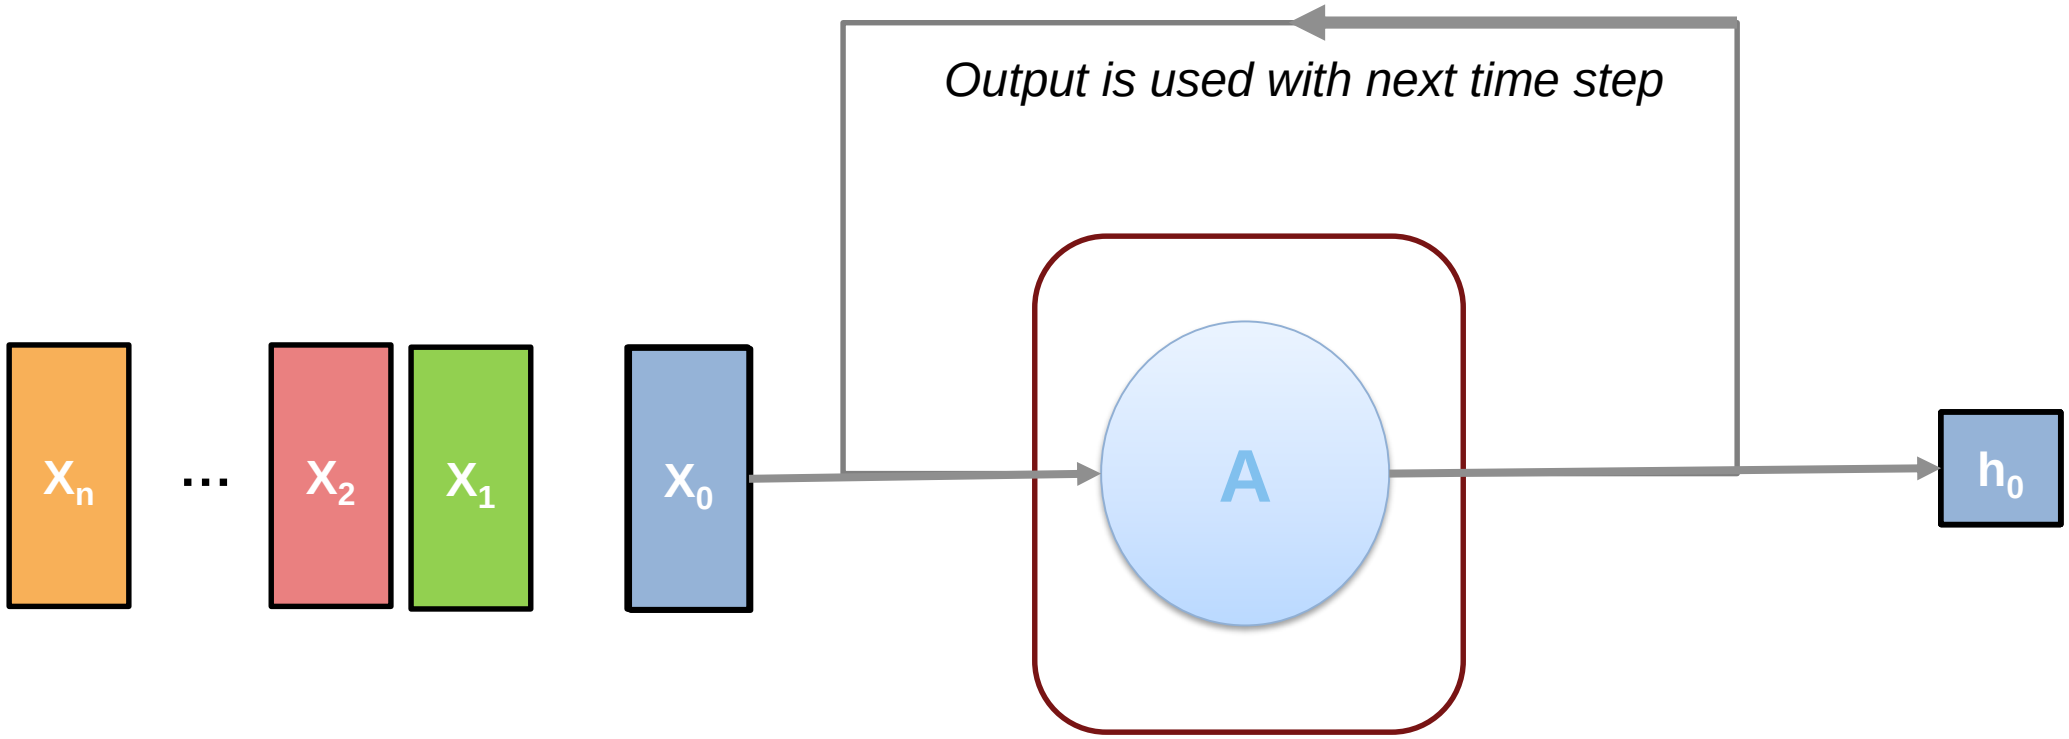

I was born in France…

[2000 words]

… I speak _____ ?

Long Short-Term Memory Network

Recurrent Neural Network that carries a memory cell (state) throughout the process

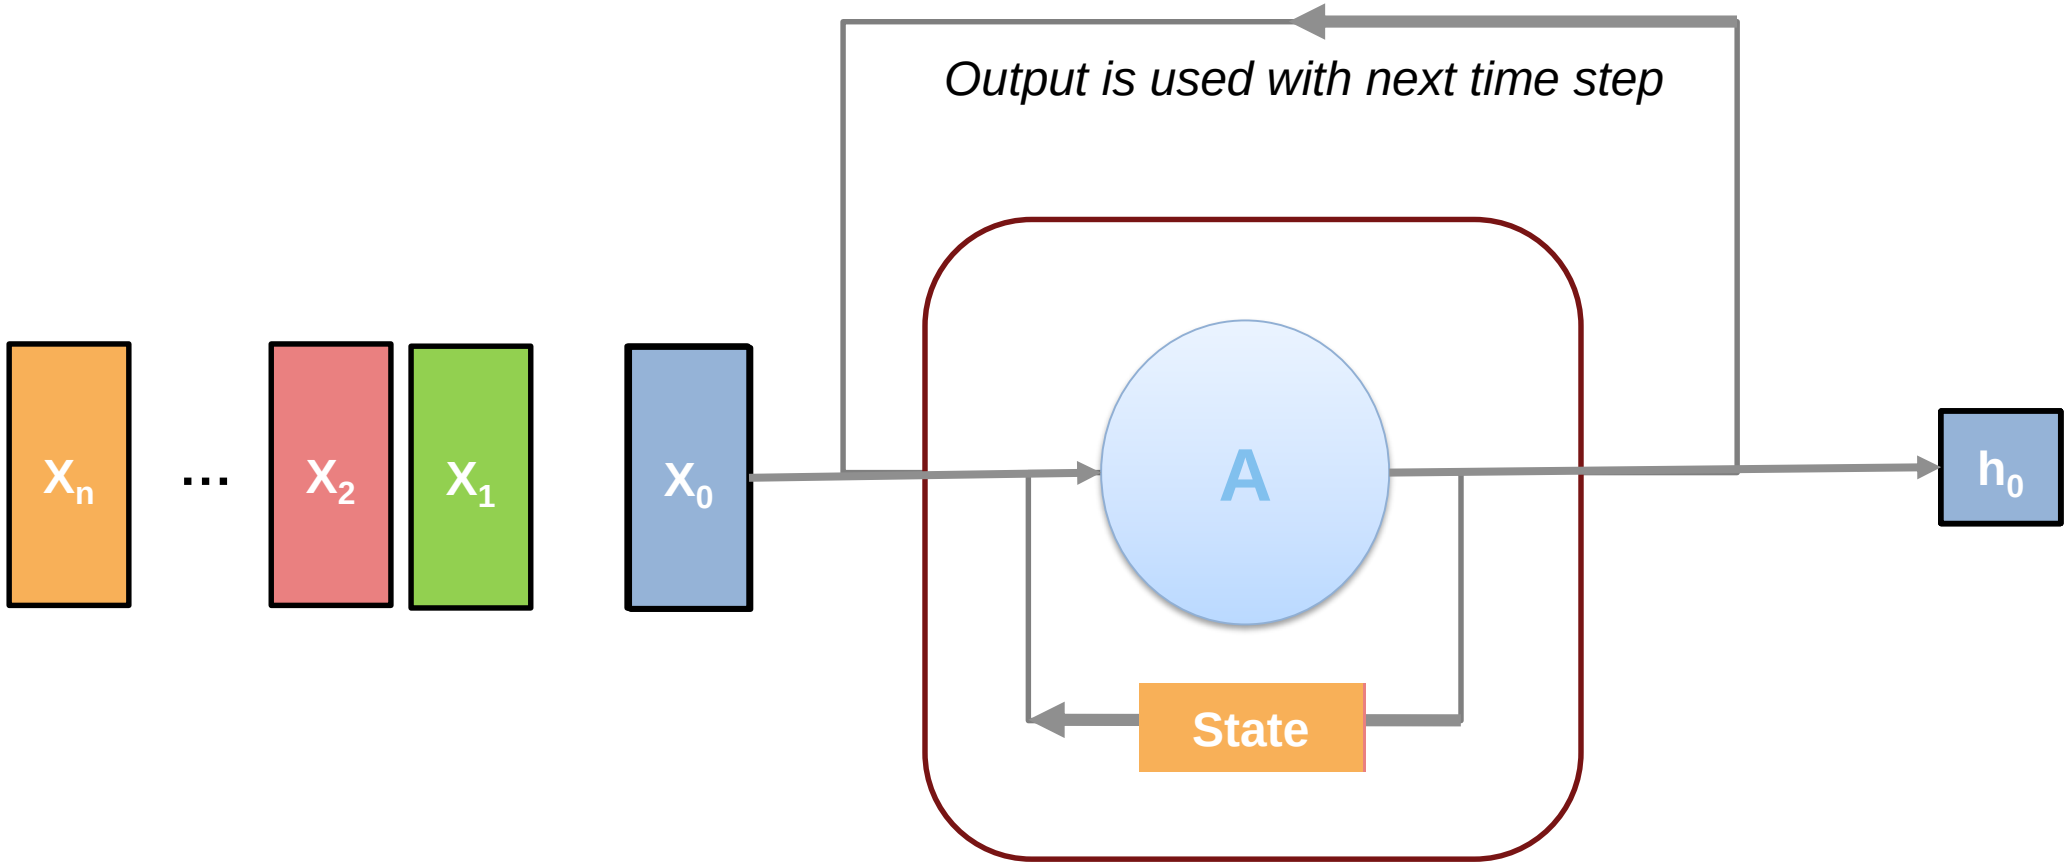

Examples in MATLAB Documentation

Sequence Classification Using Deep Learning

Classify sequence data using a long short-term memory (LSTM) network.

Open Live Script

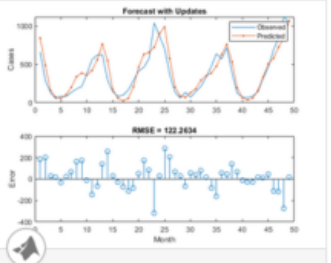

Time Series Forecasting Using Deep Learning

Forecast time series data using a long short-term memory (LSTM) network.

Open Live Script

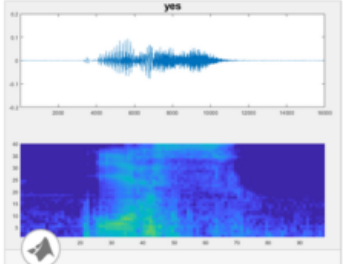

Speech Command Recognition Using Deep Learning

Train a simple deep learning model that detects the presence of speech commands in audio. The example uses the Speech Commands

Open Script

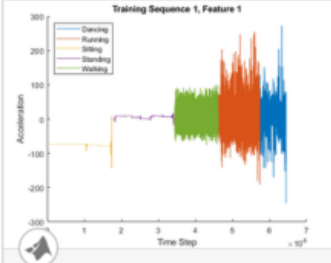

Sequence-to-Sequence Classification Using Deep Learning

Classify each time step of sequence data using a long short-term memory (LSTM) network.

Open Live Script

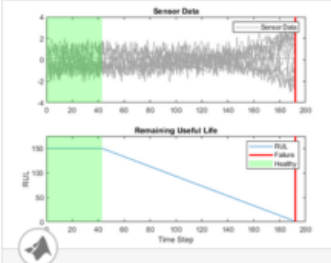

Sequence-to-Sequence Regression Using Deep Learning

Predict the remaining useful life (RUL) of engines by using deep learning.

Open Live Script

Classify Text Data Using Deep Learning

Classify text descriptions of weather reports using a deep learning long short-term memory (LSTM) network.

Open Live Script

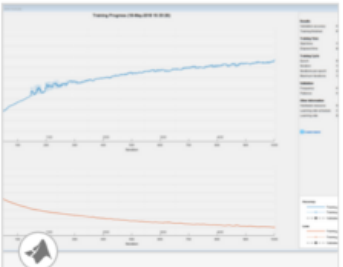

Generate Text Using Deep Learning

Train a deep learning long short-term memory (LSTM) network to generate text.

Open Live Script

Pride and Prejudice and MATLAB

Train a deep learning LSTM network to generate text using character embeddings.

Open Live Script

Word-By-Word Text Generation Using Deep Learning

Train a deep learning LSTM network to generate text word-by-word.

Open Live Script

Classify Out-of-Memory Text Data Using Custom Mini-Batch Datastore

Classify out-of-memory text data with a deep learning network using a custom mini-batch datastore.

Open Live Script

Exercise – ECG Signal Classification

Purpose:

- Use LSTM to classify ECG signal as normal heartbeat or Atrial Fibulation
- Perform preliminary feature engineering and view difference in results.

To Do:

1. Open work_ClassifyECGSignals.mlx.

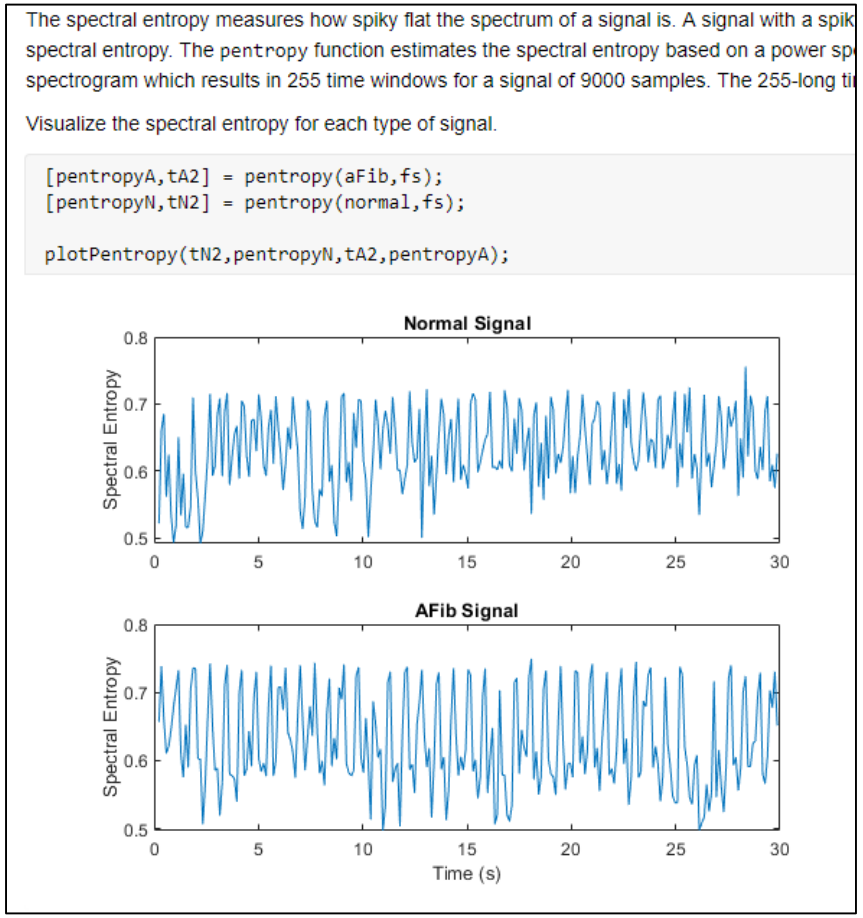

Speech Recognition Example

Audio signal → Spectrogram → Image Classification algorithm

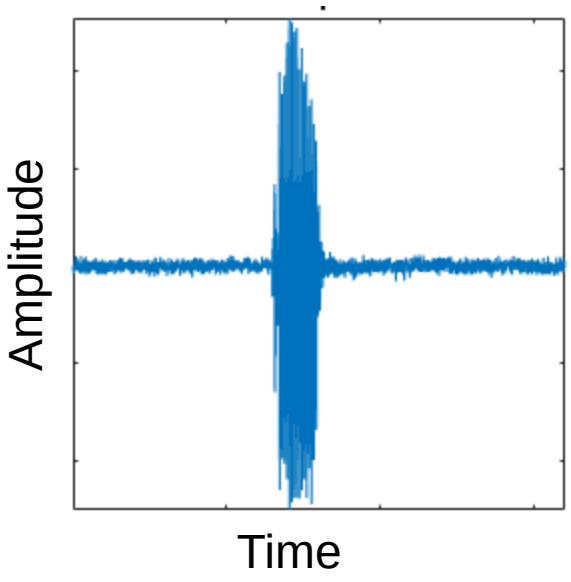

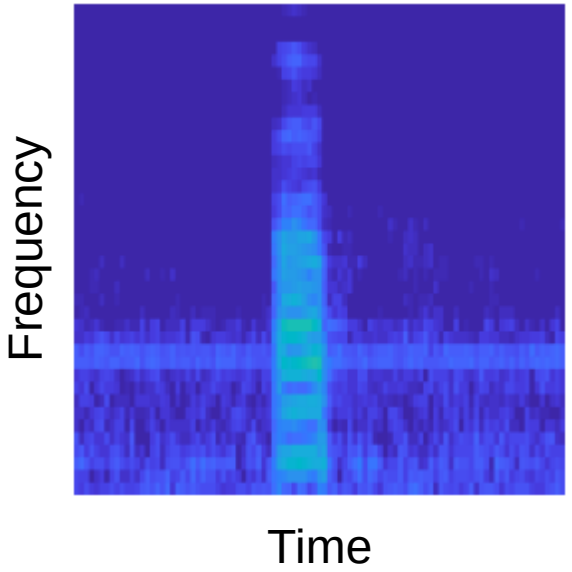

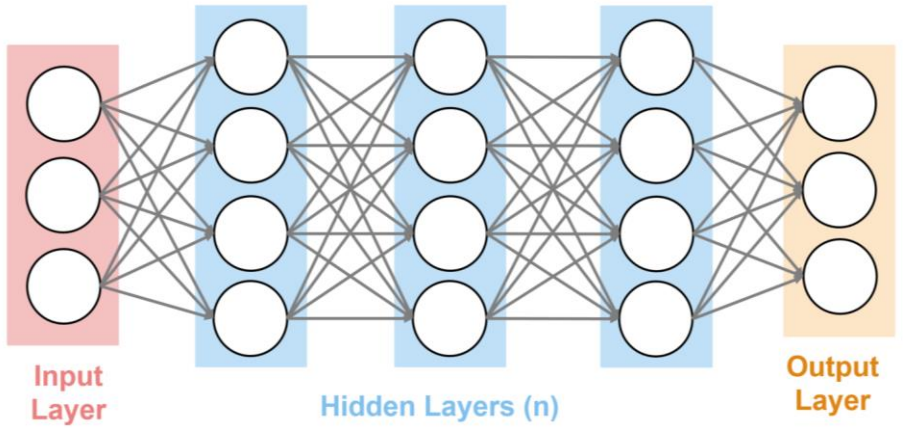

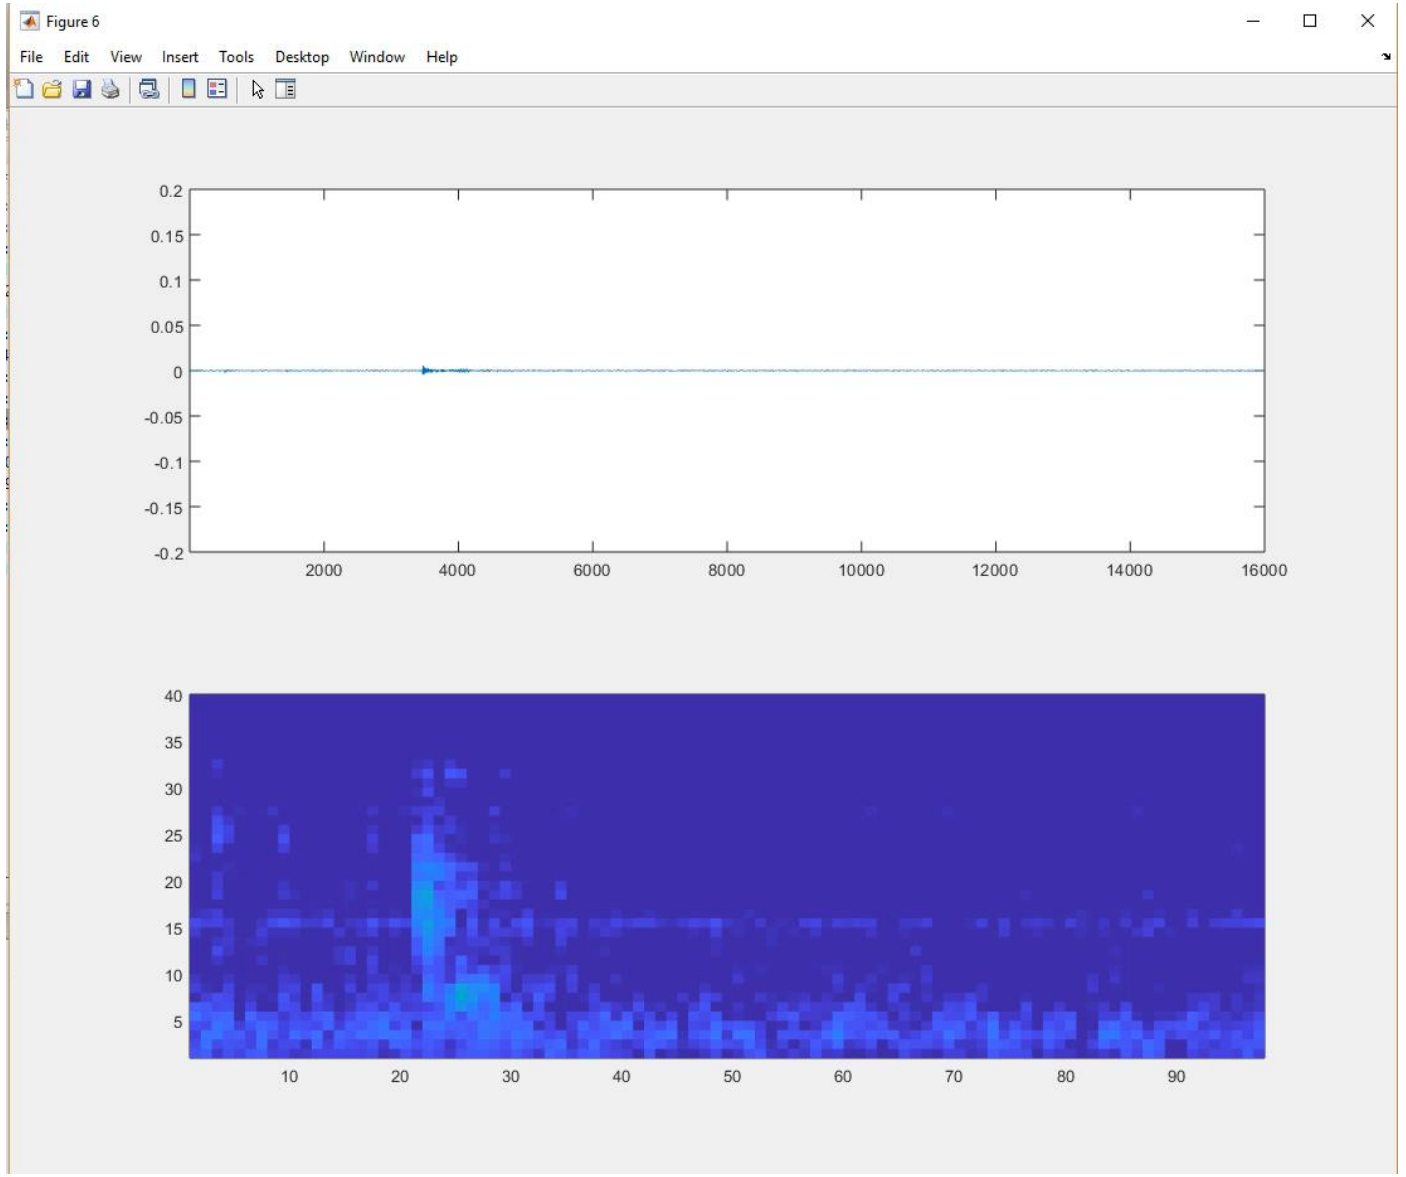

Speech Recognition using CNNs

Exercise: Speech Command Recognition with Deep Learning

- Train a Convolutional Neural Network (CNN) to recognize speech commands
- Work with Google's speech command dataset
- Leverage:
  - `audioDatastore` (Read and manage large datasets)
  - `melSpectrogram` (Transform 1D signals into 2D images using perceptually-spaced frequency scaling)
- Prototype trained network in real-time on live audio