TECHSOURCE
Dynamic Solutions Precise Results

Deeplearning with MATLAB

Ian Alferez
Application Engineering Manager
Techsource - Asia

Southeast Asia's sole distributor of
MATLAB® & SIMULINK®

Internet Retweeted

Computer Facts
@computerfact

concerned parent: if all your friends jumped off a bridge would you follow them?
machine learning algorithm: yes.

2:20 PM · Mar 15, 2018


Deep Learning Demo
Image Classification

Agenda

- Introduction
-  **Demo 1:** Deep learning in 6 lines of code
- Deep Learning Fundamentals
-  **Demo 2 and 3:** Exploring pretrained networks/Classifying handwritten digits
-  **Demo 4:** Transfer Learning – OR – Non Image Exercise
-  **Demo:** Deploying Deep Networks– OR – Improving Network Accuracy
- Conclusion

Deep Learning Applications

Voice assistants (speech to text)

Teaching character to beat video game

Automatically coloring black and white images

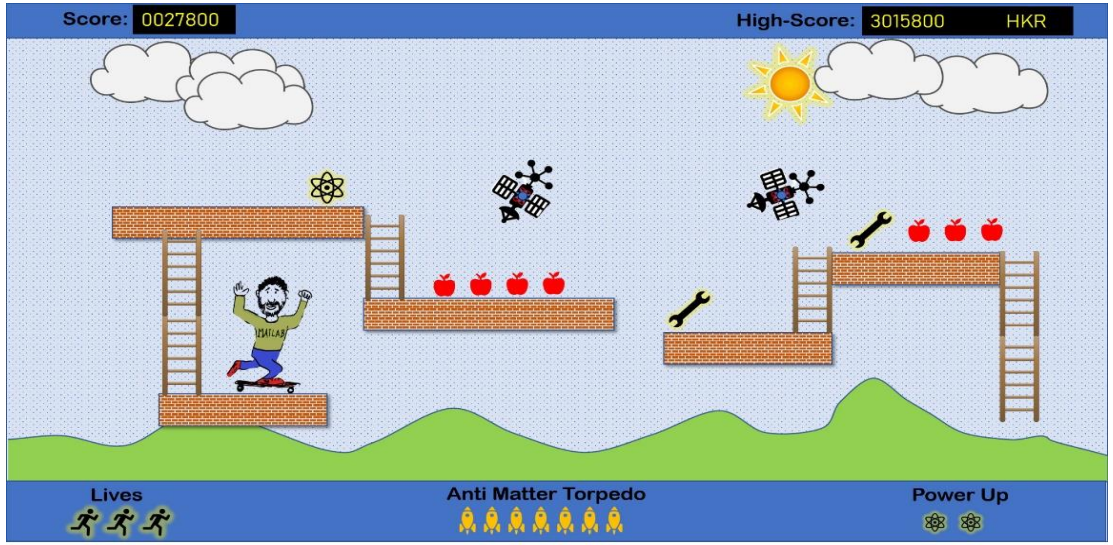

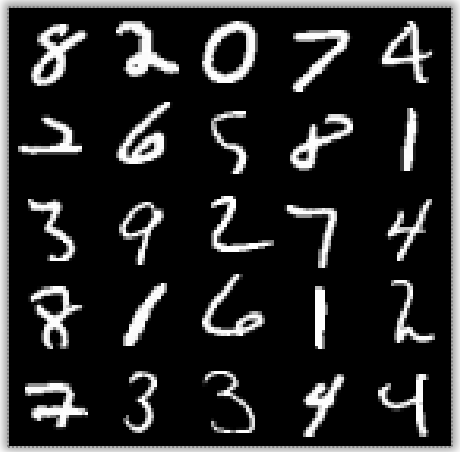

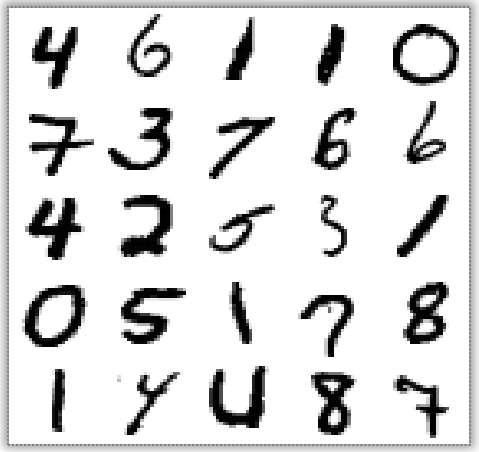

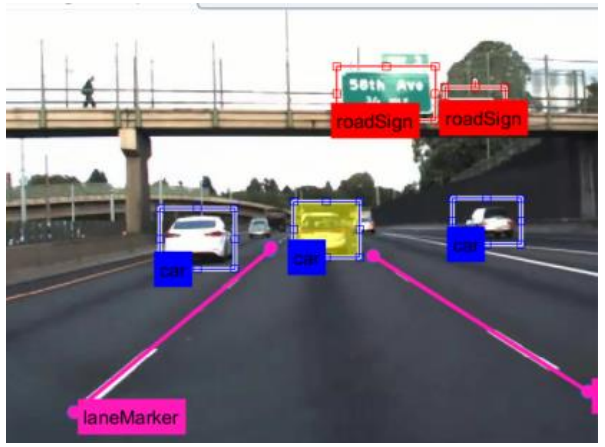

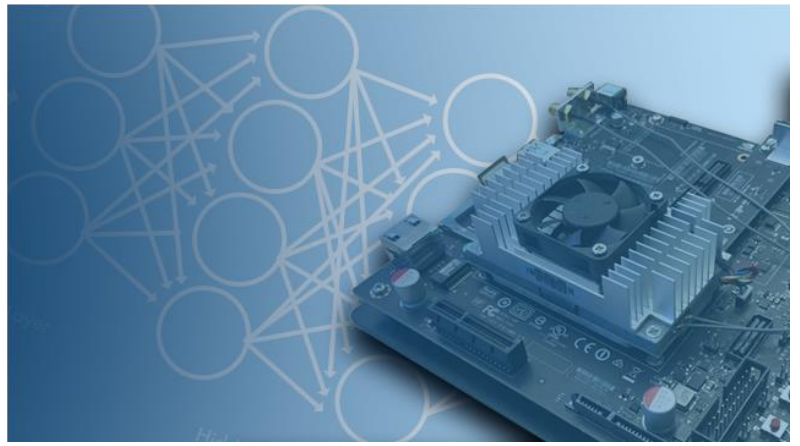

What is Deep Learning?

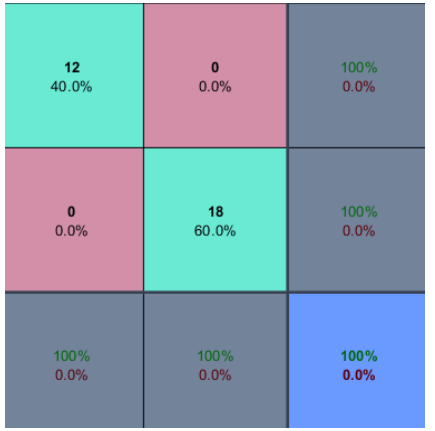

What is Deep Learning?

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

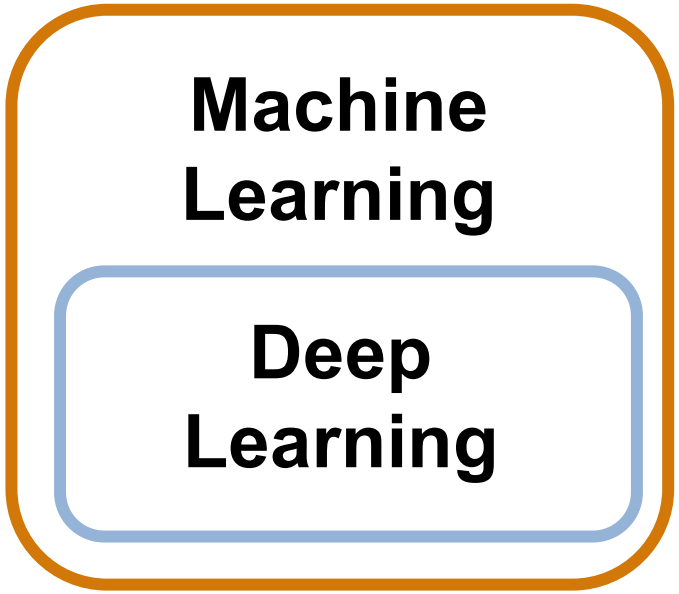

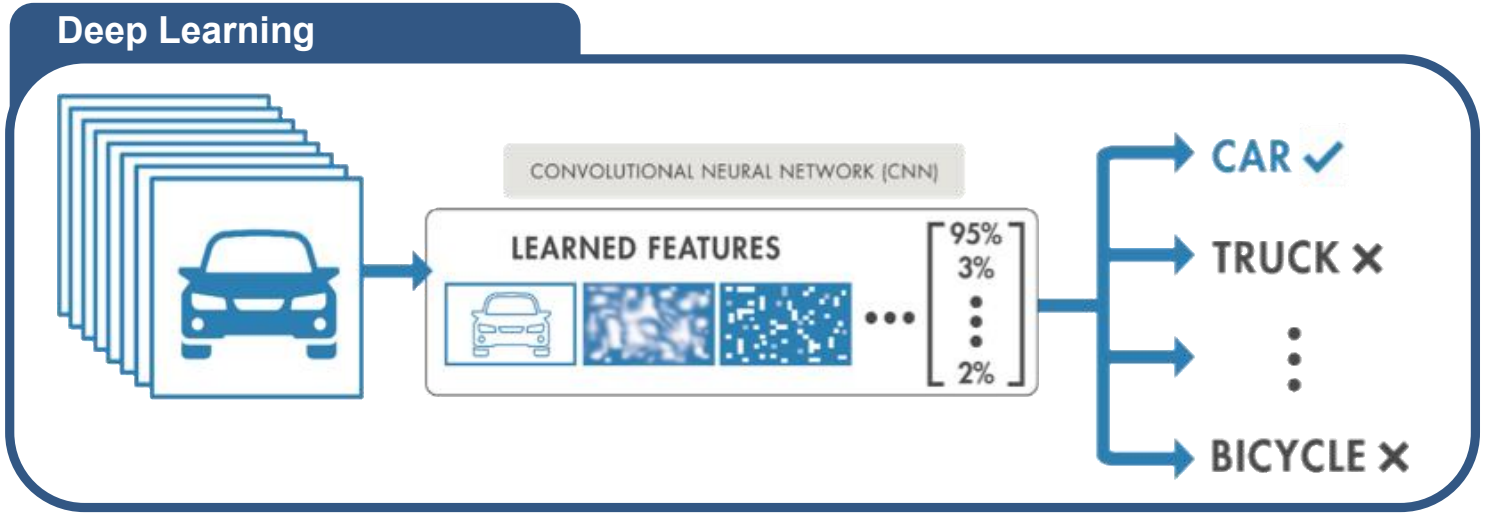

Applications of Deep Learning

Shell

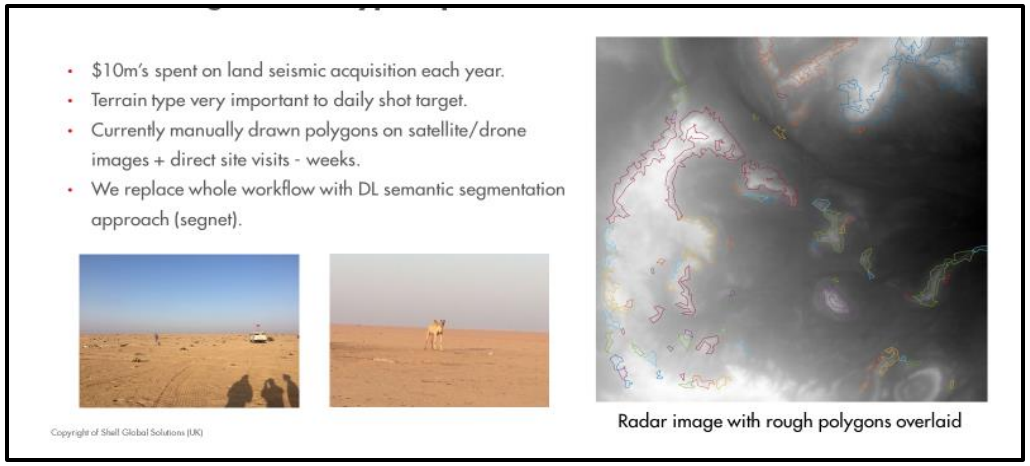

Terrain Recognition with Hyperspectral Data

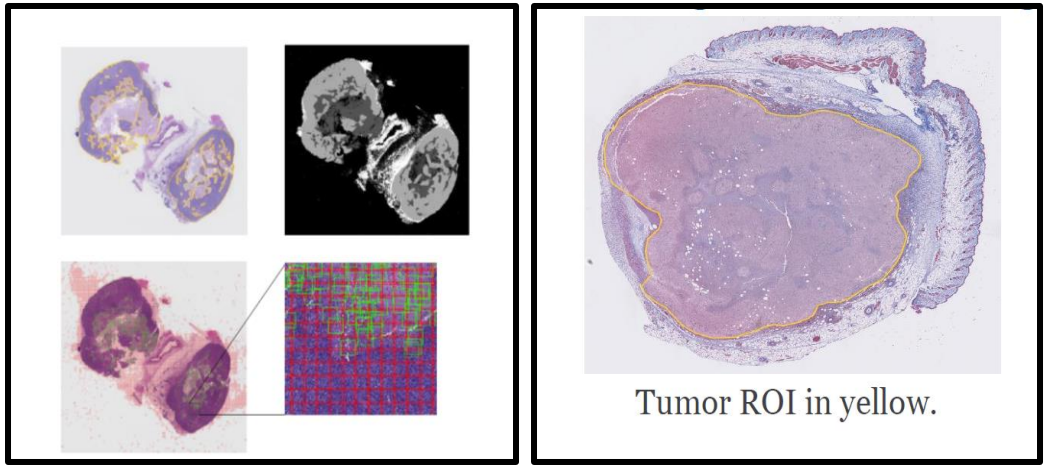

CNNs for Digital Pathology Analysis

Genentech

Veoneer

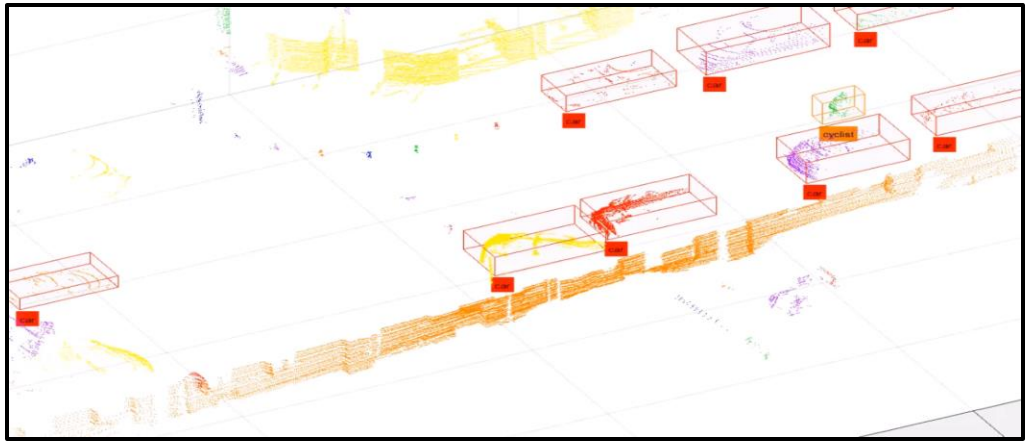

LiDAR-Based Sensor Verification

Equipment Classification

Caterpillar

Deep Learning Models can Surpass Human Accuracy.

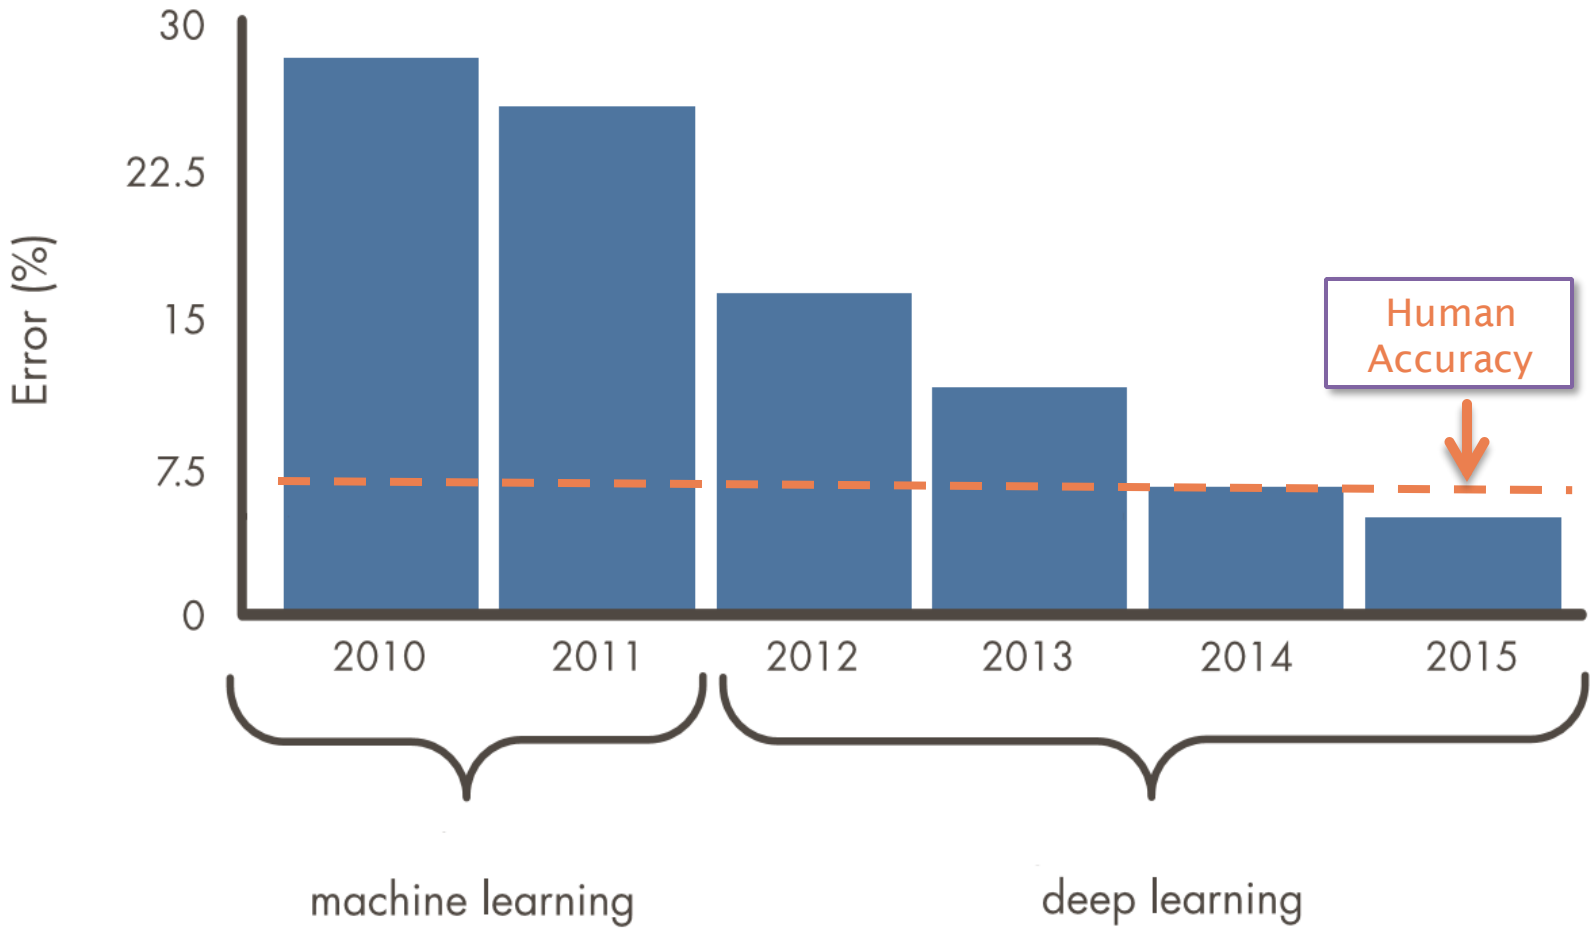

Source: ILSVRC Top-5 Error on ImageNet

Deep Learning Enablers

Increased GPU acceleration

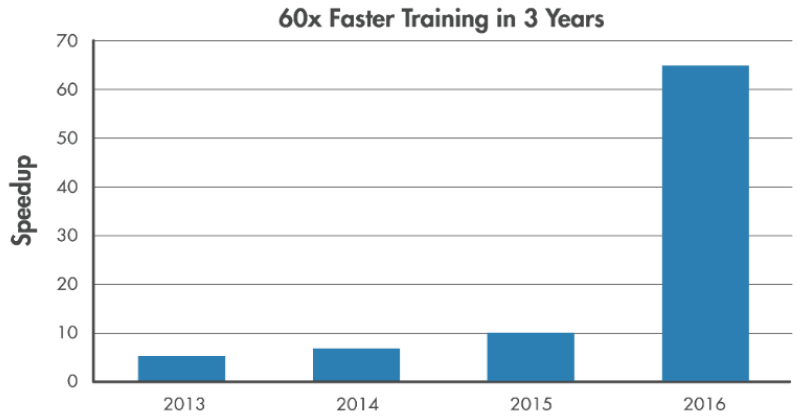

Labeled public datasets

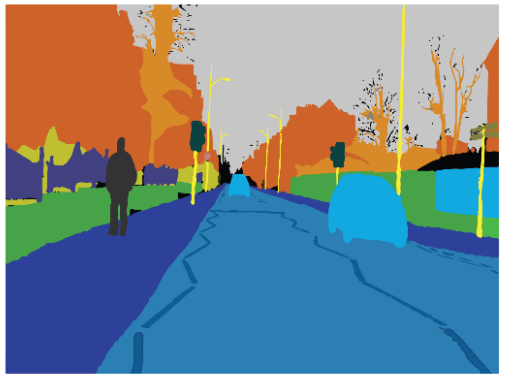

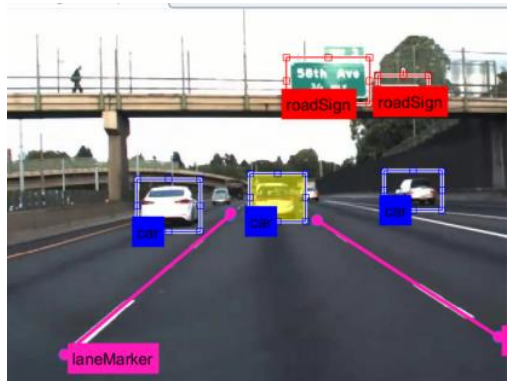

World-class models

AlexNet
PRETRAINED MODEL

VGG-16/19
PRETRAINED MODEL

Caffe
MODEL IMPORTER

ResNet
PRETRAINED MODEL

GoogLeNet
PRETRAINED MODEL

TensorFlow-Keras
MODEL IMPORTER

Deep Learning Datatypes

Image

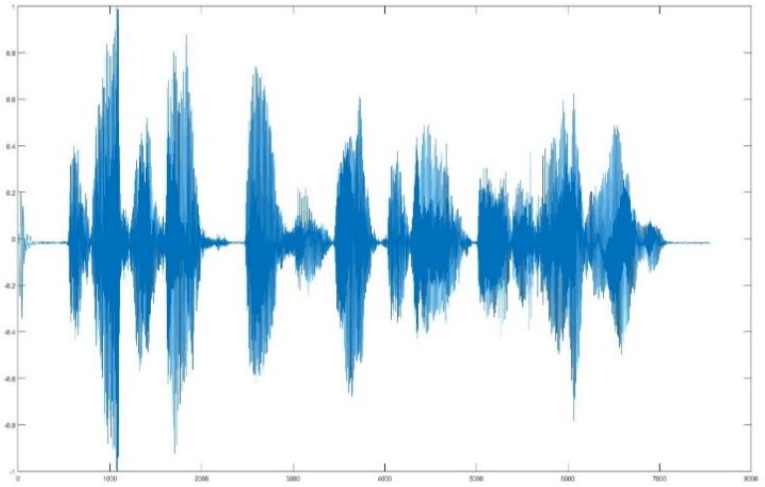

Signal

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

Wide Domain Support for Deep Learning in MATLAB

Computer Vision

Image Processing

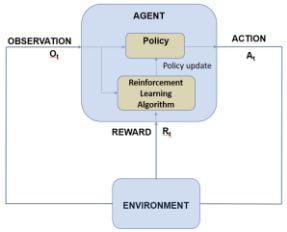

Control Design

Text Analytics

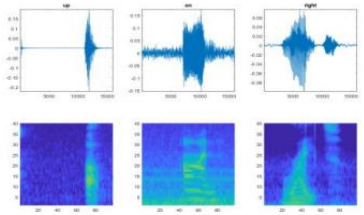

Audio Processing

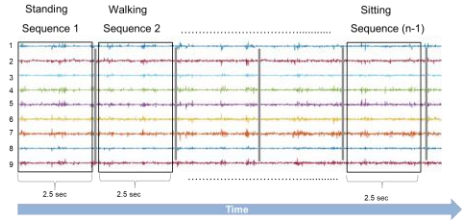

Sensor Data Analysis

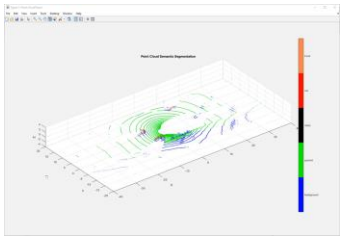

Lidar Processing

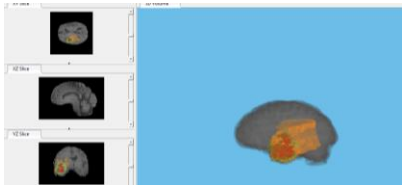

N-D Volumes

Let's try it out!

Open: ***DeepLearningIn6Lines.mlx***

in folder 01-DeepLearningIn6Lines

Deep learning is not complicated. It can be easy!

Expectation

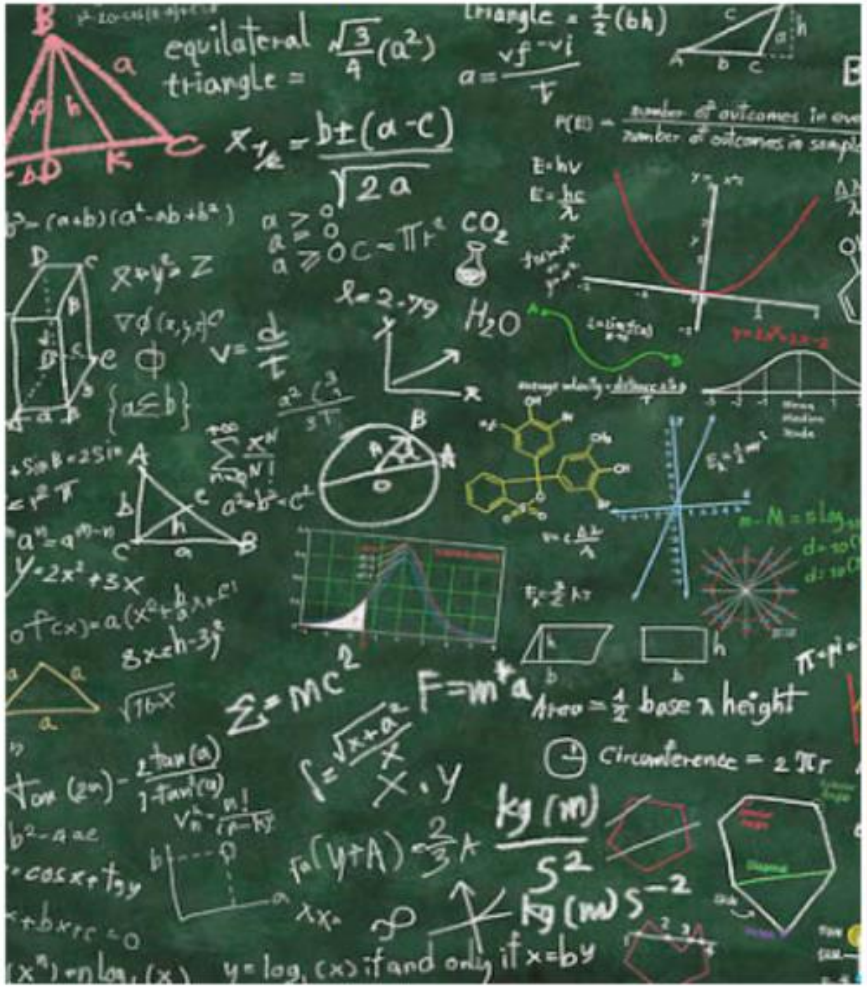

Reality

Deep Learning Uses a Neural Network Architecture

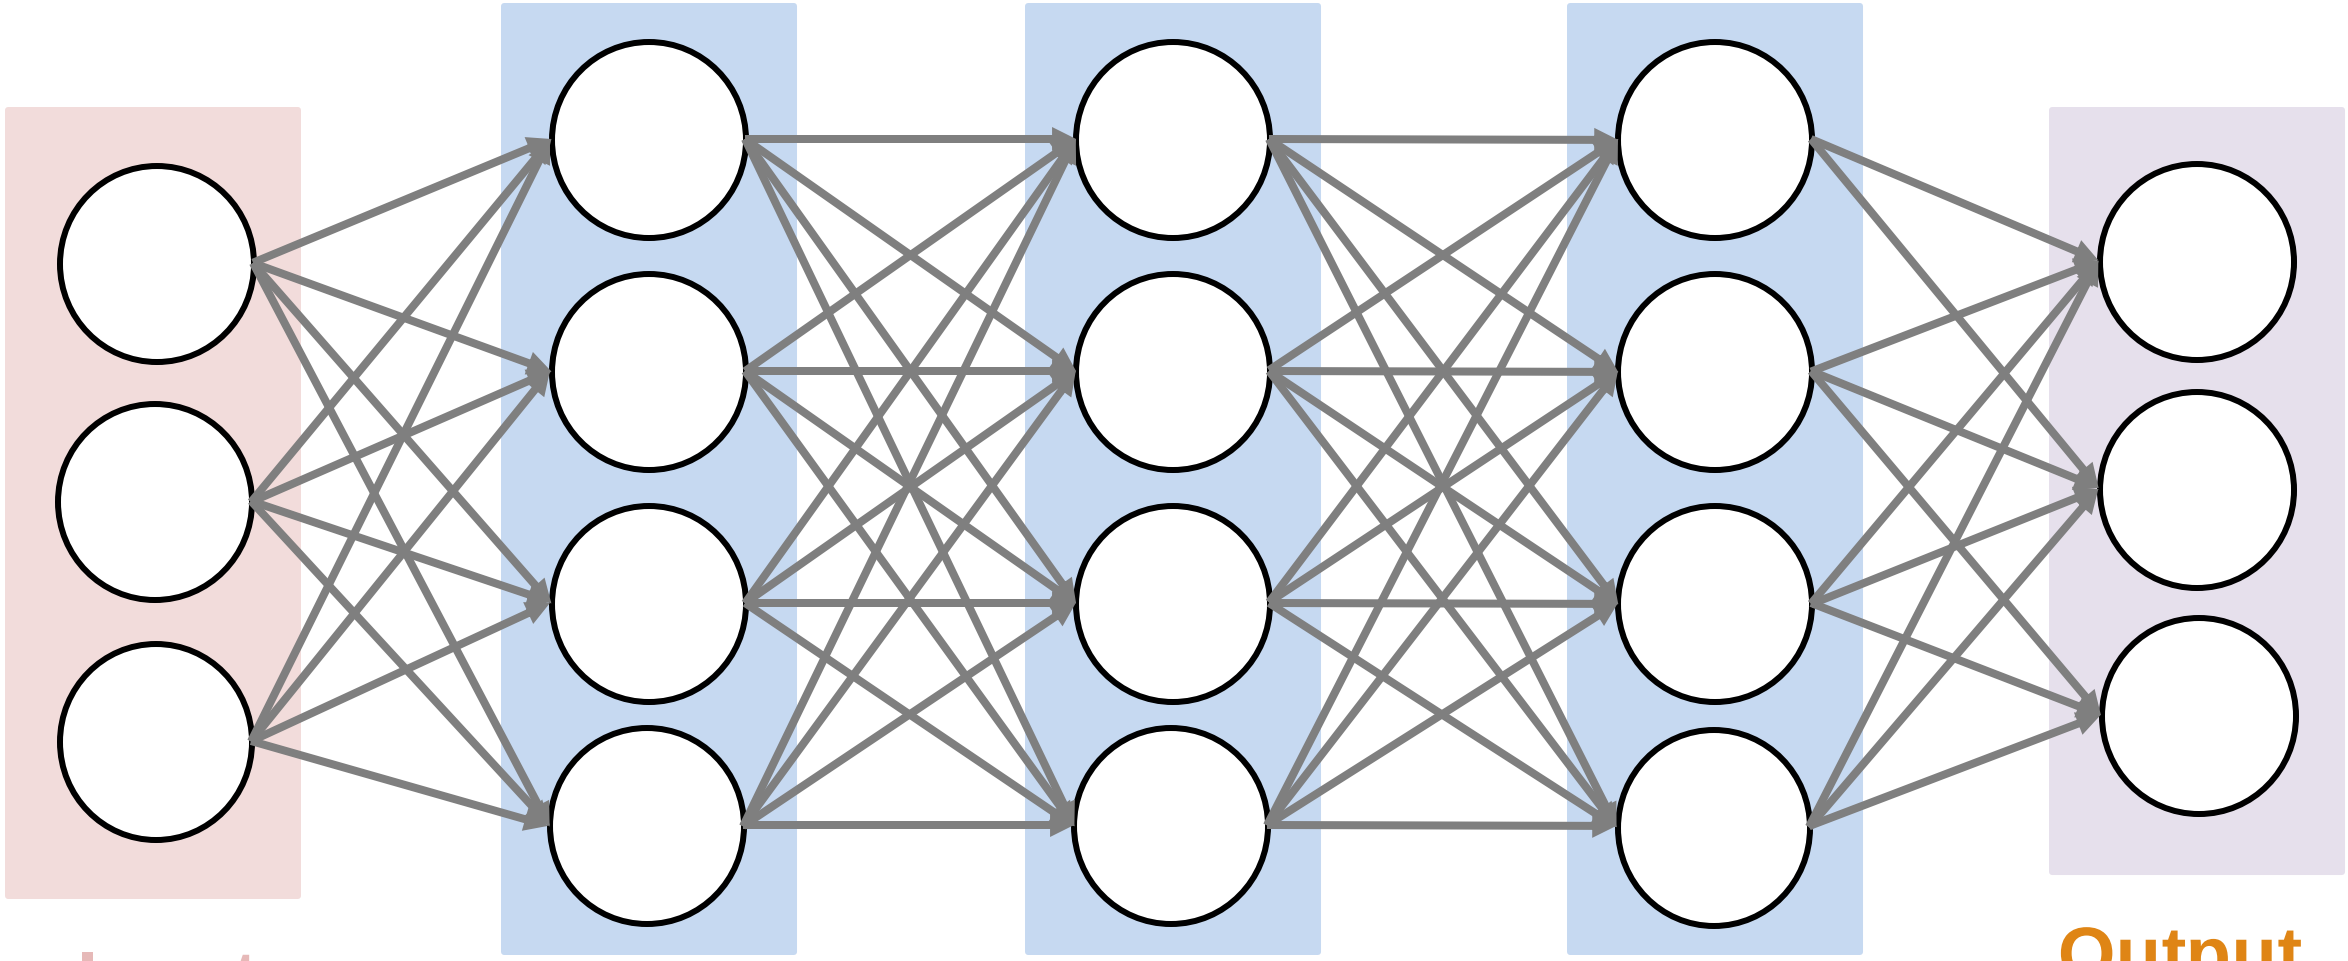

Thinking about Layers

- 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

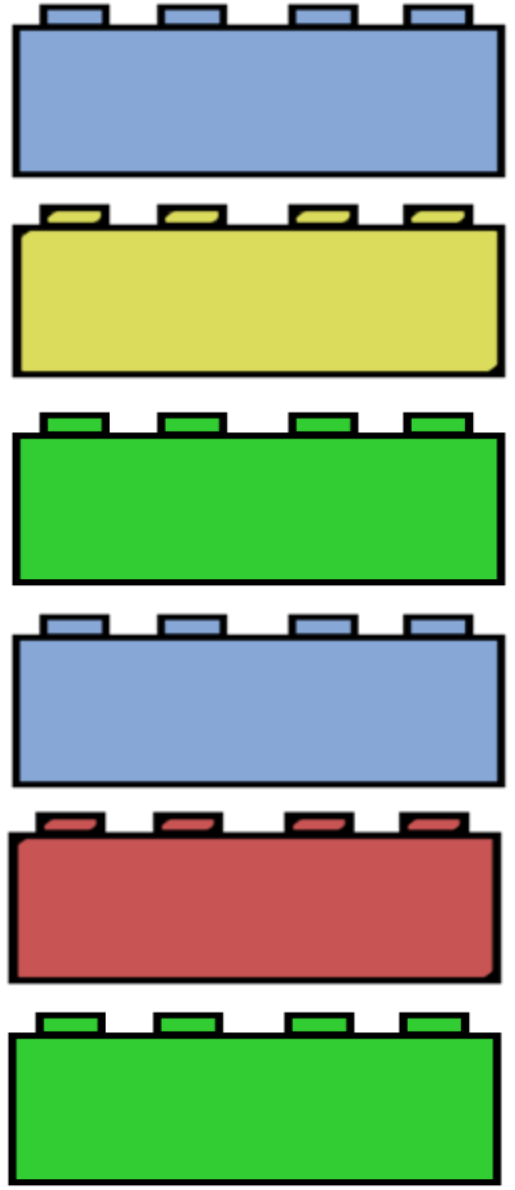

Thinking about Layers

- Layers are like blocks
  - Stack 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

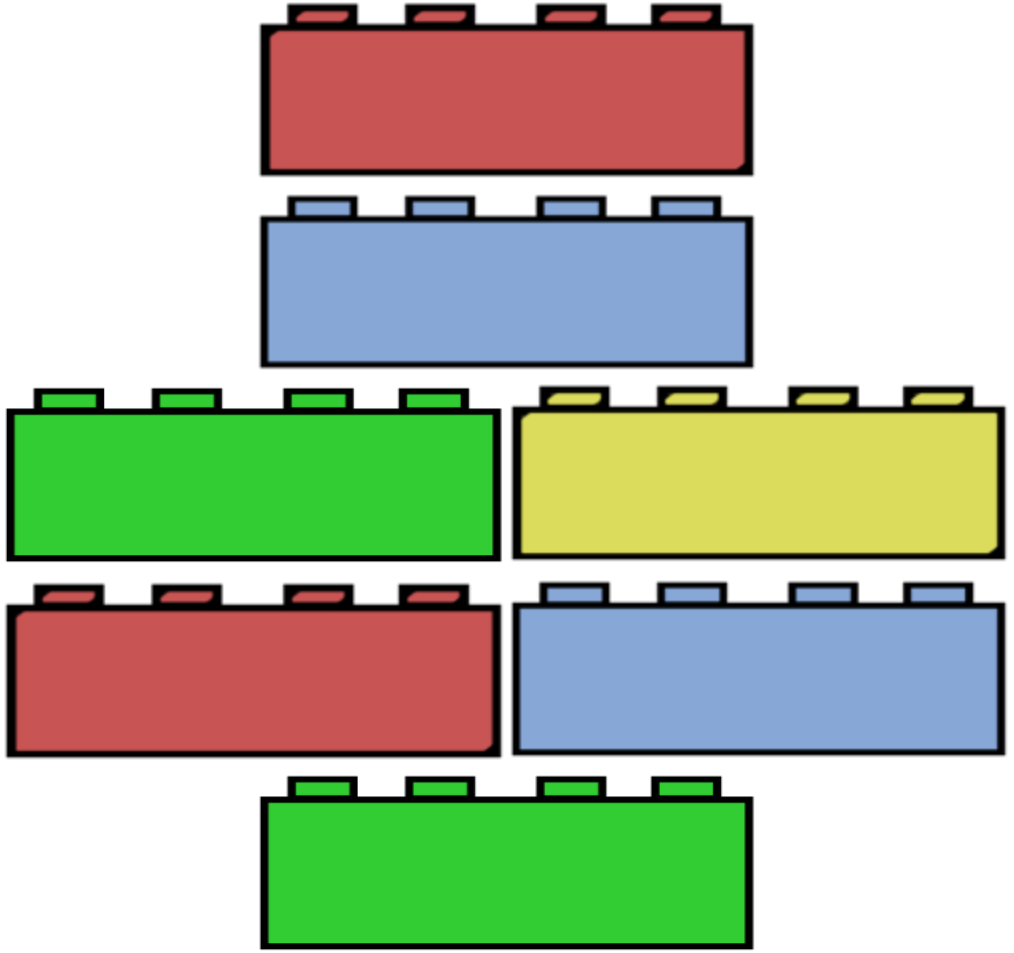

Convolutional Neural Networks (CNNs)

- Special layer combinations that make them great for image classification
  - Convolution Layer
  - Max Pooling Layer
  - ReLU Layer

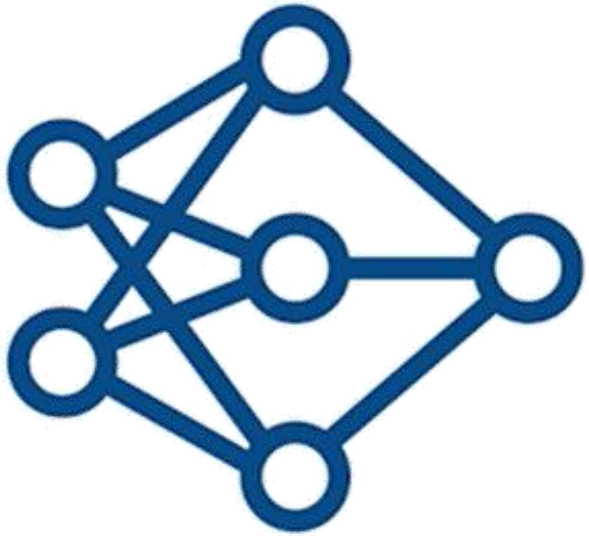

Convolution Layers Search for Patterns

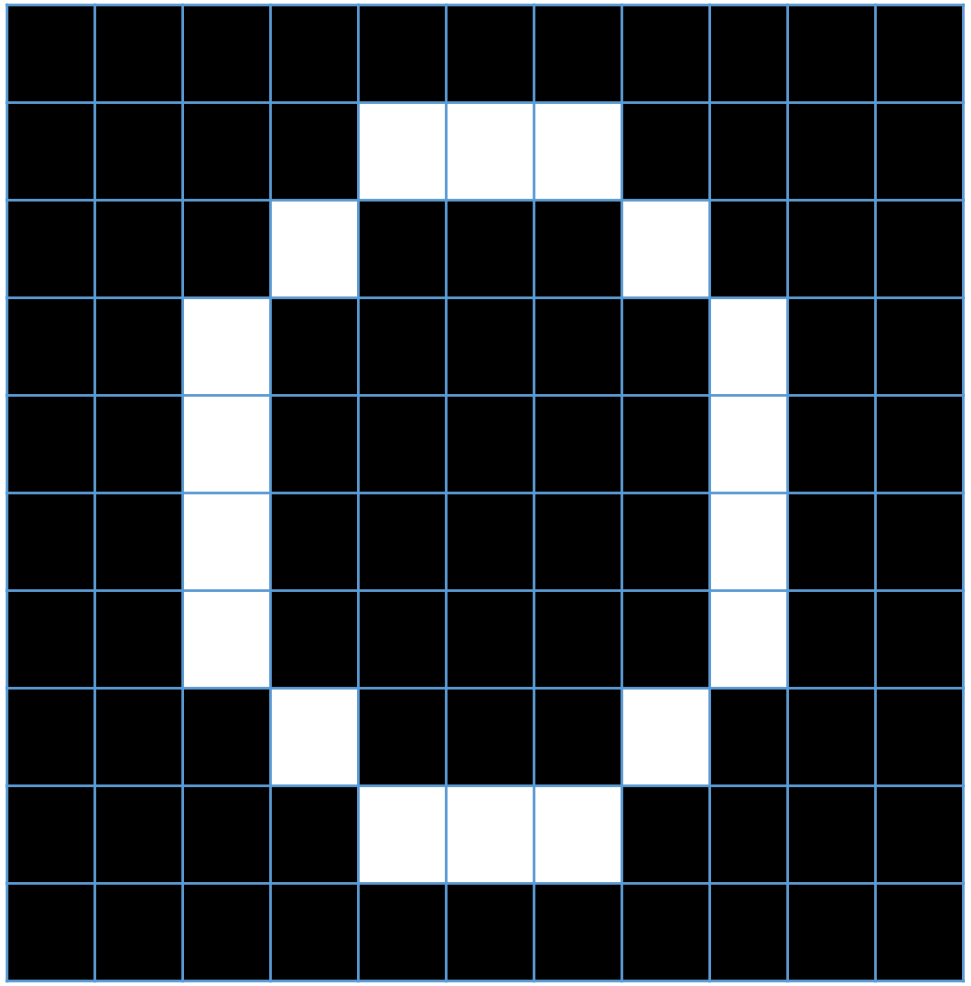

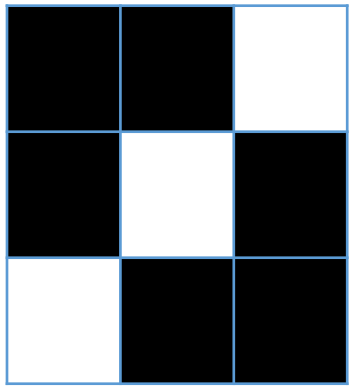

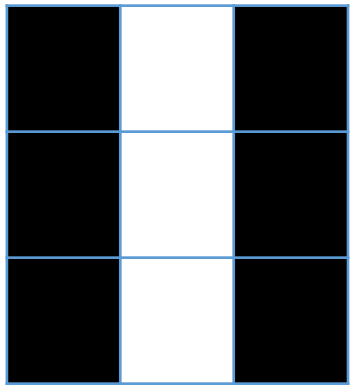

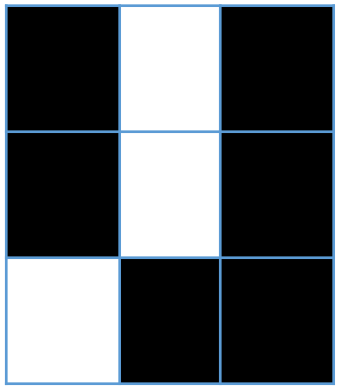

These patterns would be common in the number 0

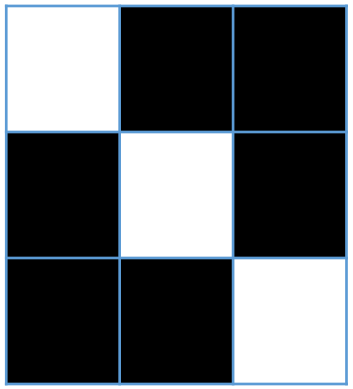

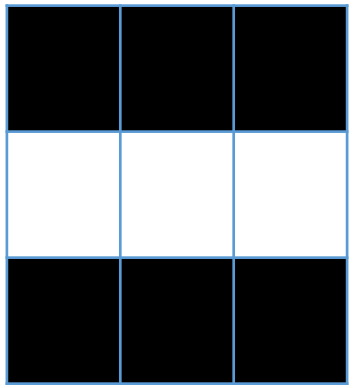

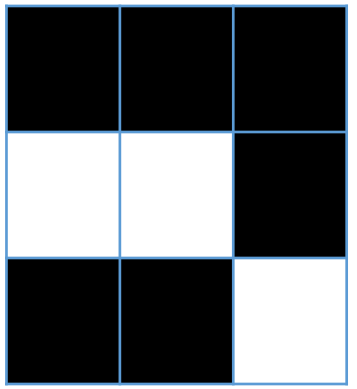

All patterns are compared to the patterns on a new image.

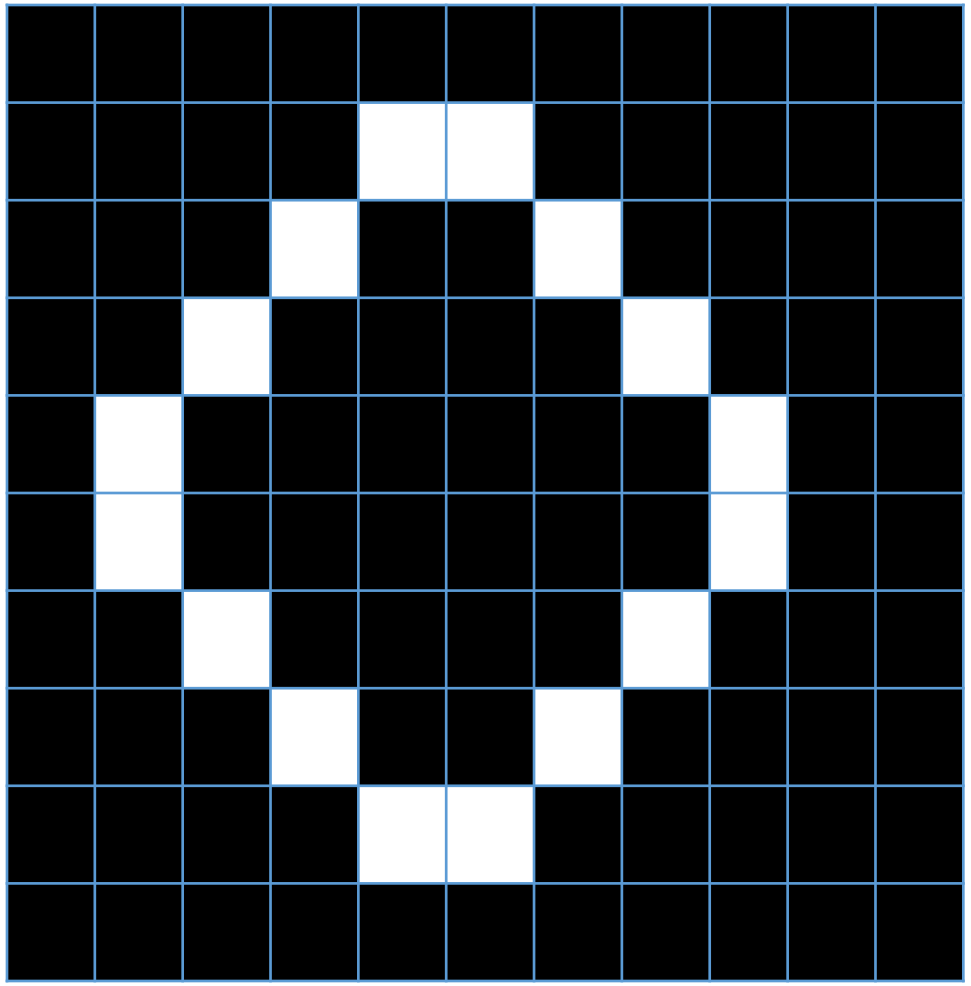

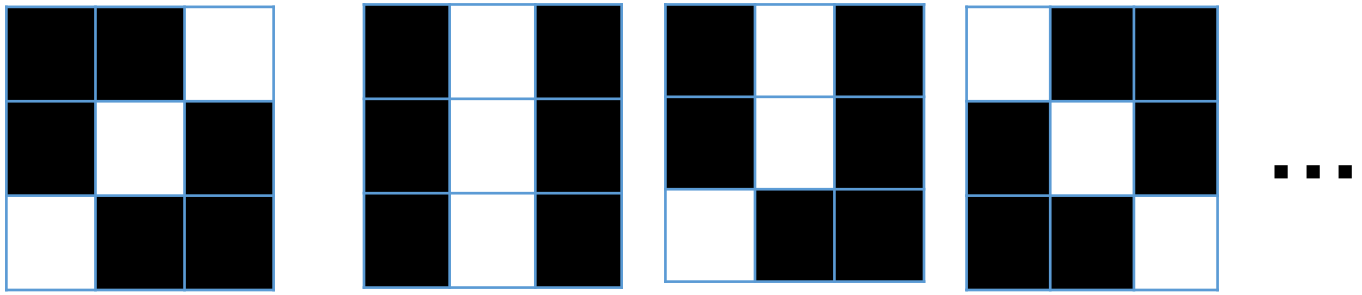

- Pattern starts at left corner
  - Perform comparison
  - Slide over one pixel
- Reach end of image
- Repeat for next pattern

Convolution Layers Search for Patterns

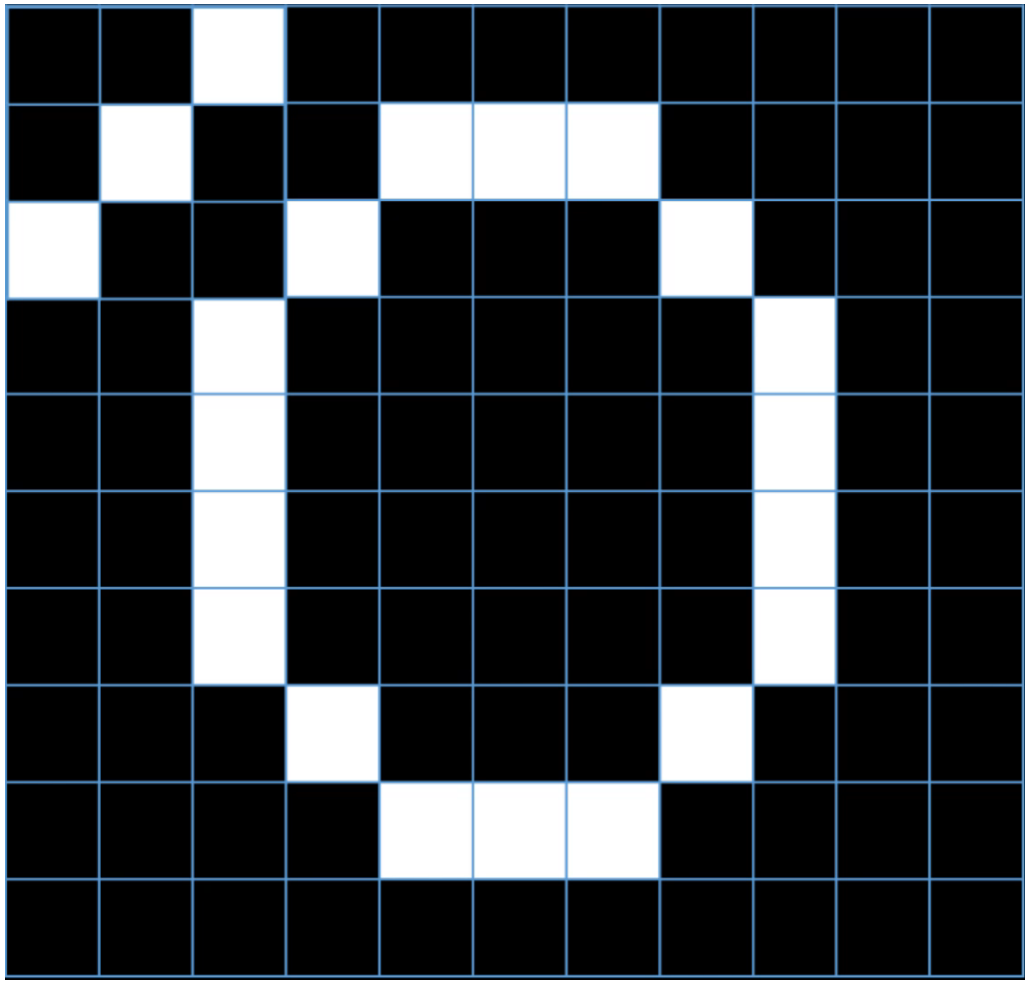

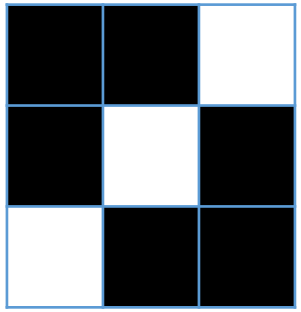

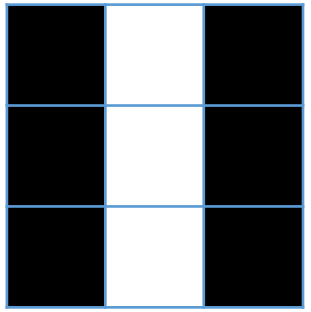

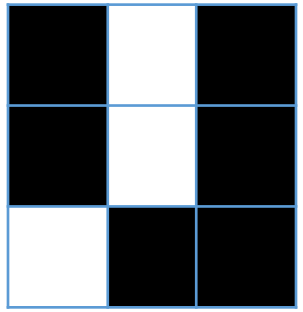

These patterns would be common in the number 0

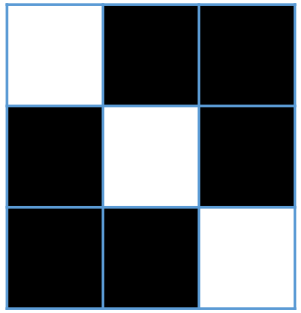

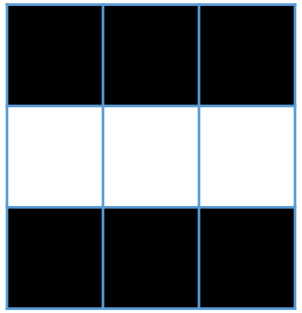

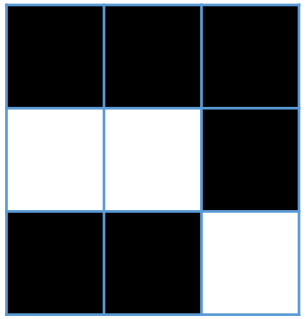

Good pattern matching in convolution improves chances that object will classify properly

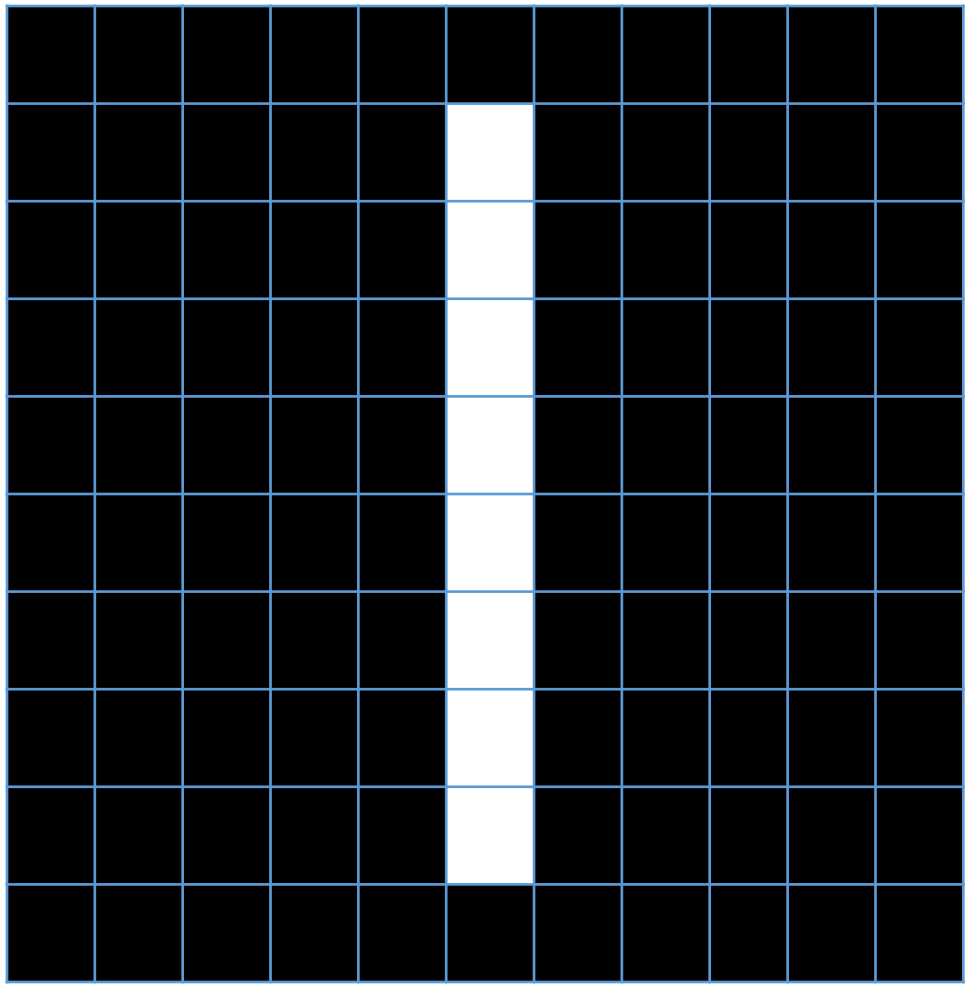

- This image would not match well against the patterns for the number zero
- It would only do very well against this pattern

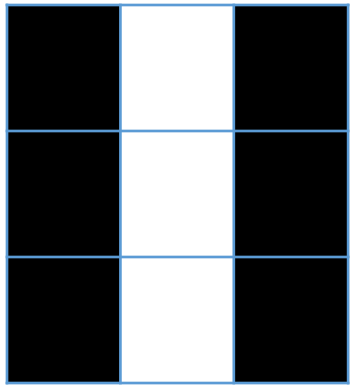

Max Pooling is a down-sampling operation

Shrink large images while preserving important information

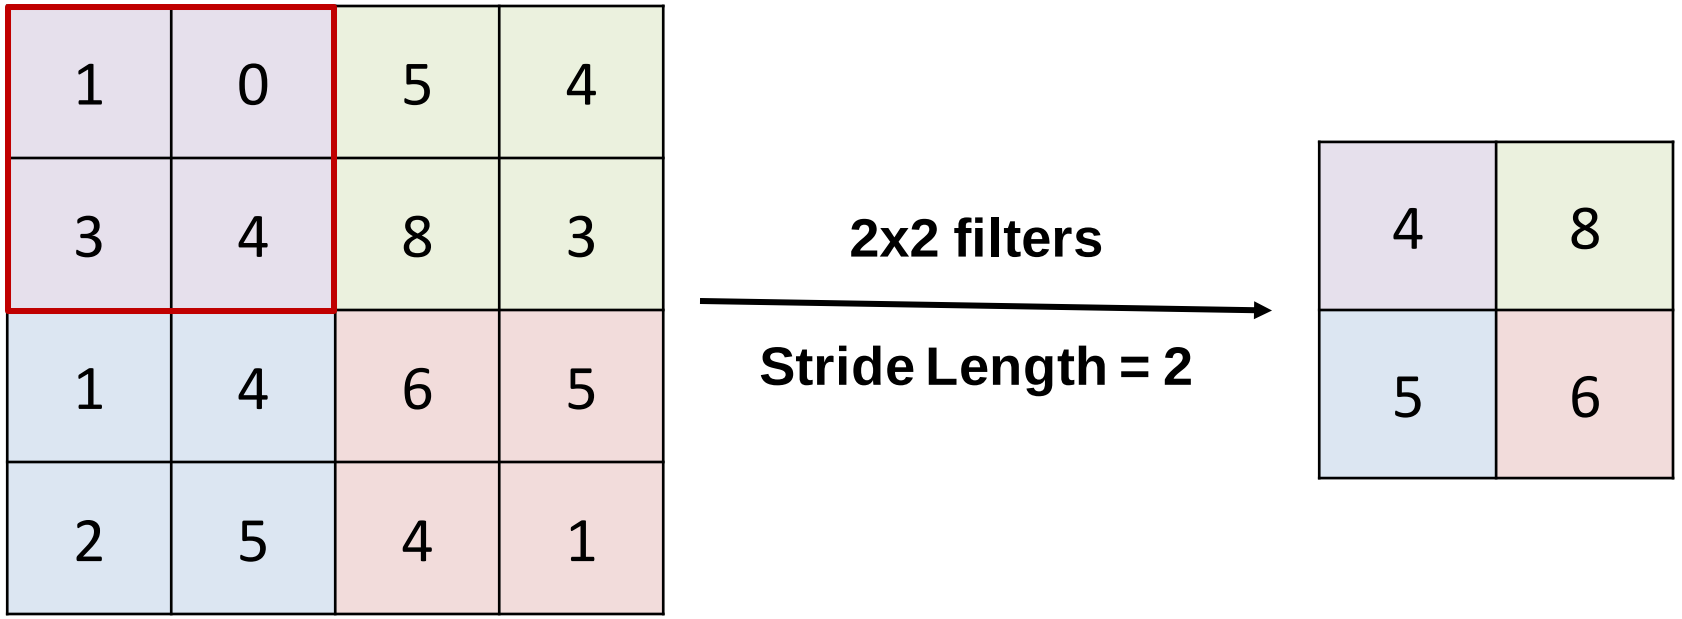

Rectified Linear Units Layer (ReLU)

Typically converts negative numbers to zero

-1	0	5	4
3	-4	-8	3
1	4	6	-5
-2	-5	4	1

0	0	5	4
3	0	0	3
1	4	6	0
0	0	4	1

CNNs 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)

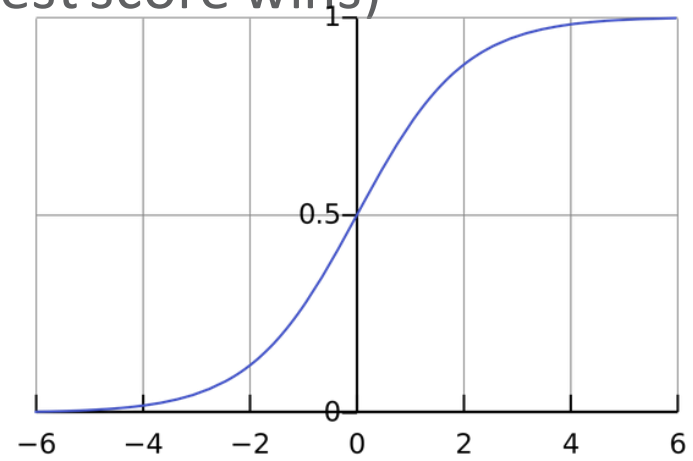

- Softmax Layer
  - Turns scores into probabilities.

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

Questions?

Pre-trained Models

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

Import & Export Models Between Frameworks

- Keras-Tensorflow Importer
- Caffe Model Importer
- ONNX Model Converter

Model Exchange with MATLAB

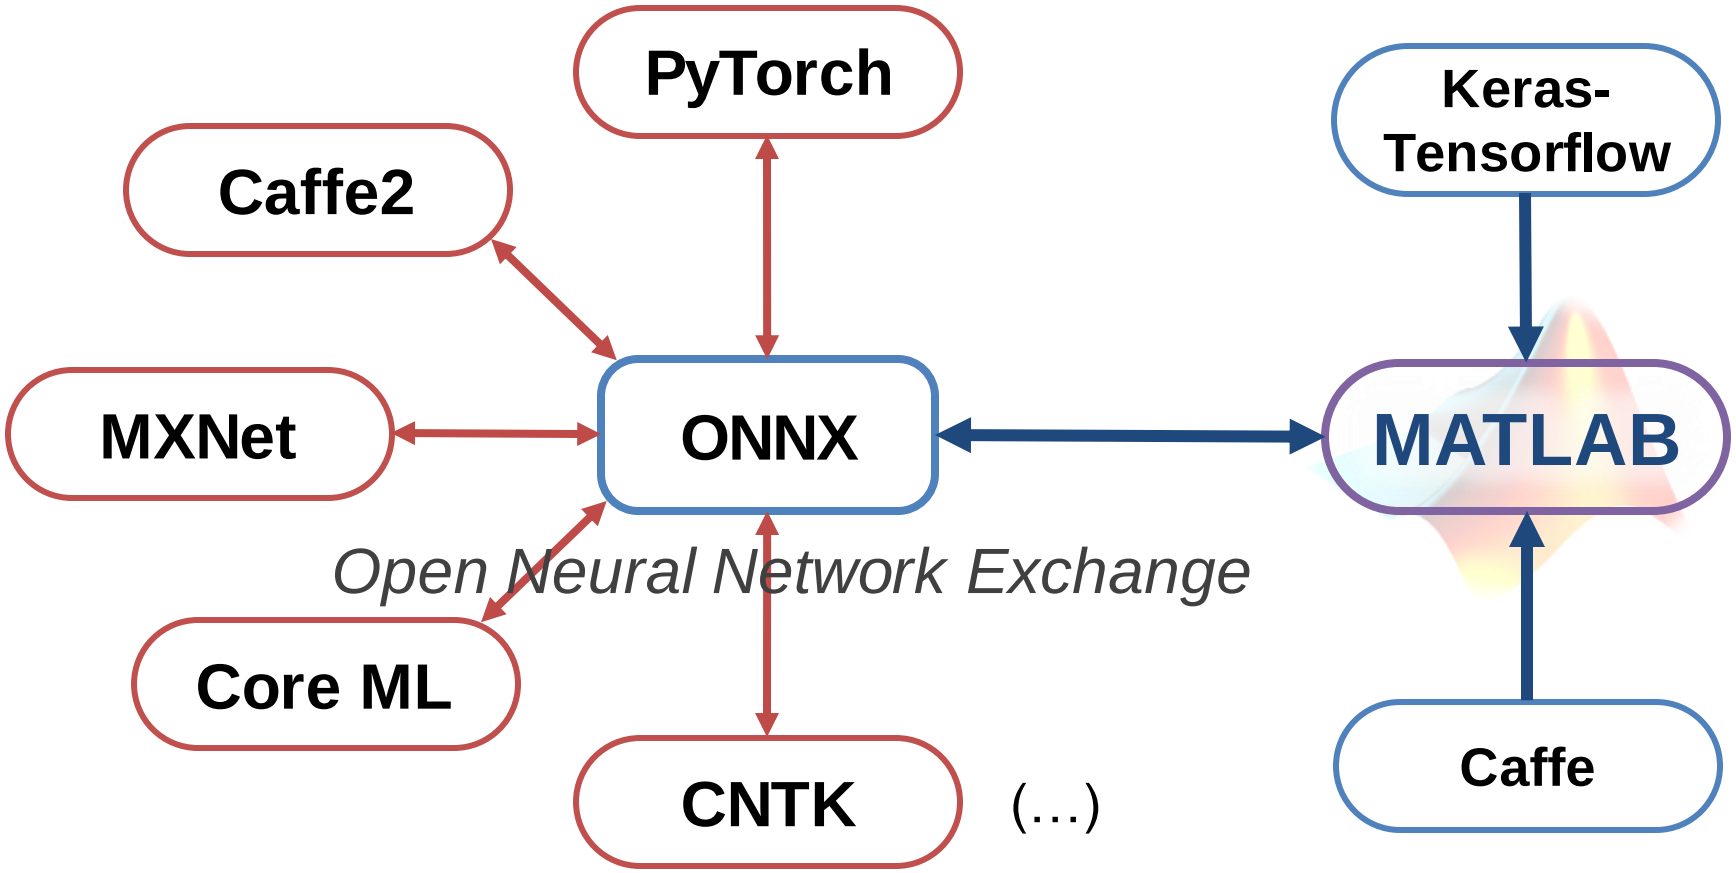

Let's try it out!

Exercise: ***Work_ExploringPretrainedNetworks.mlx***
in folder 02-PretrainedModelExercise

Takeaways (Pretrained Models)

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

Deep Learning Workflow

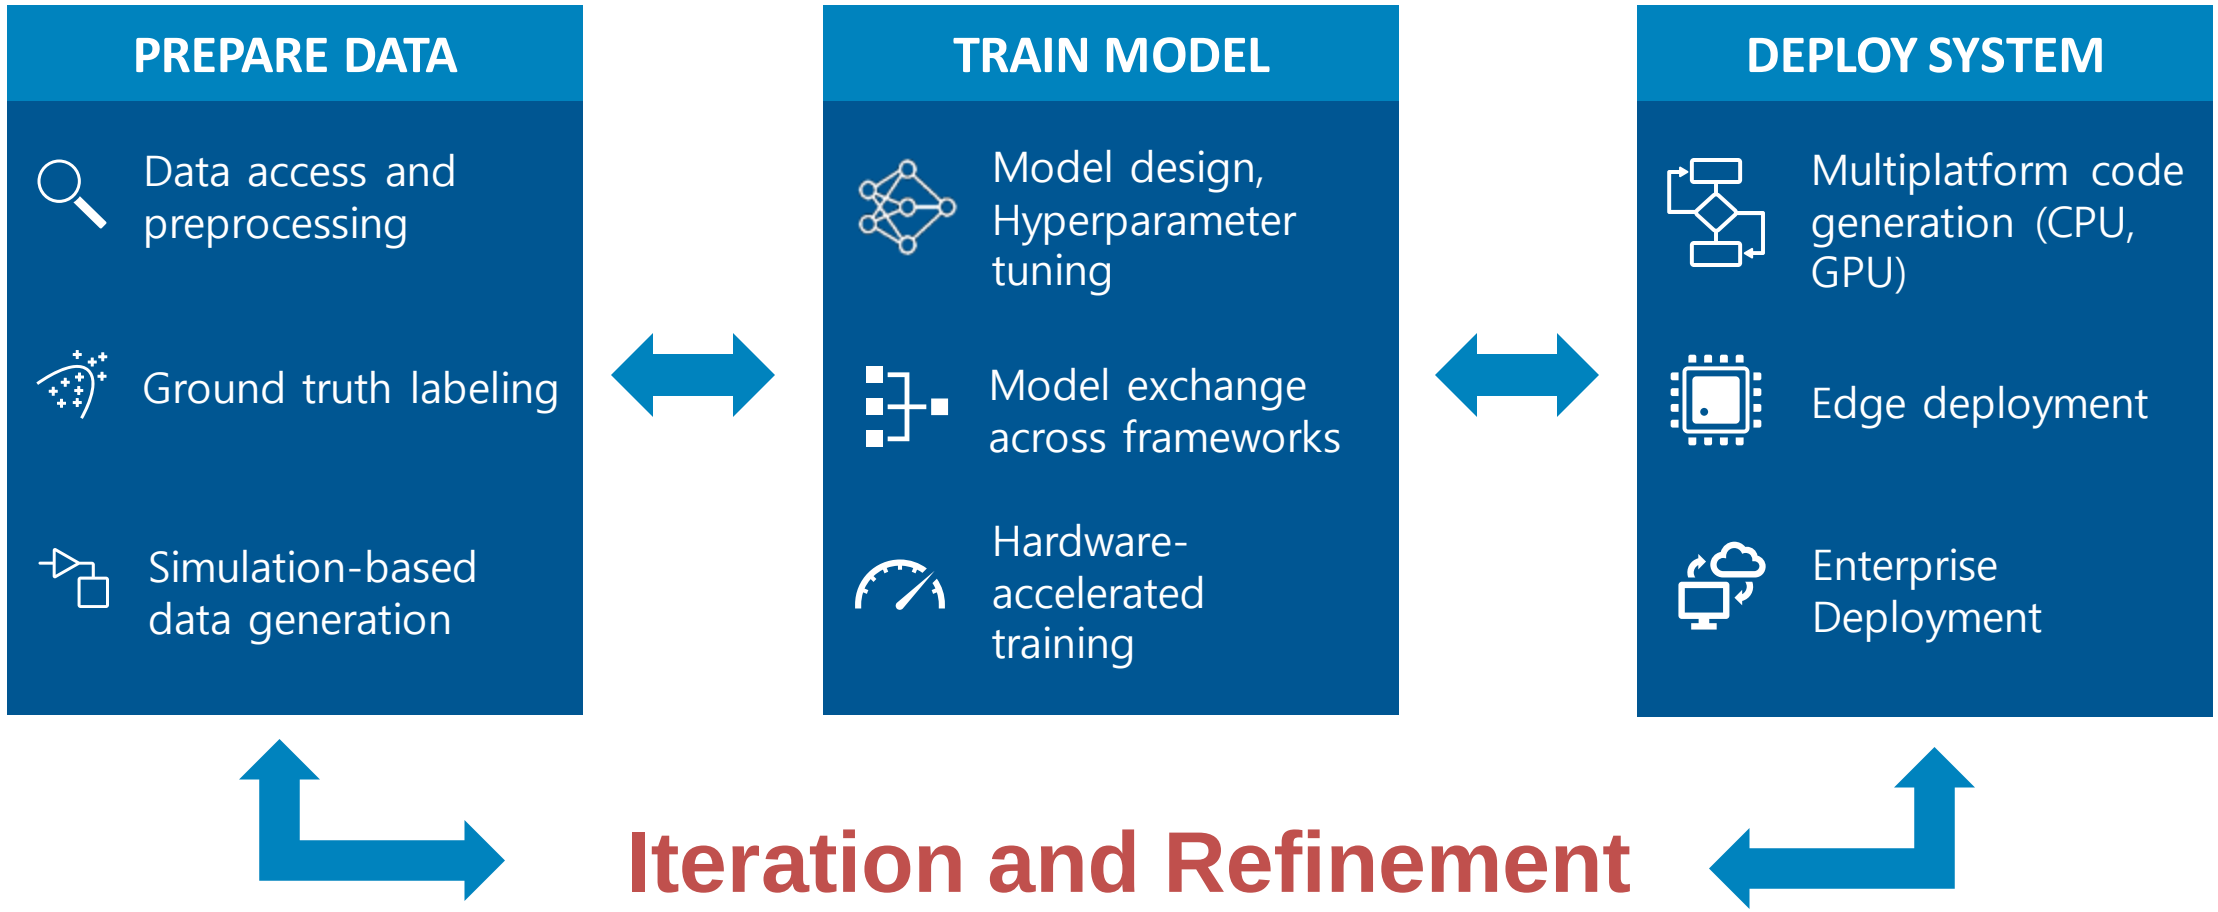

3 Components to Train a Network

Data

How much data?

It depends…but
A LOT

Network Architecture

Define Inputs and layers for deep learning

Training Options

Influence training time and accuracy

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

Let's try it out!

Exercise: ***MNIST_HandwritingRecognition.mlx***

in folder 03-MNISTExercise

Takeaways (MNIST)

- 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
  - Initial learning rate

Two Approaches for Deep Learning

1. Train a Deep Neural Network from Scratch

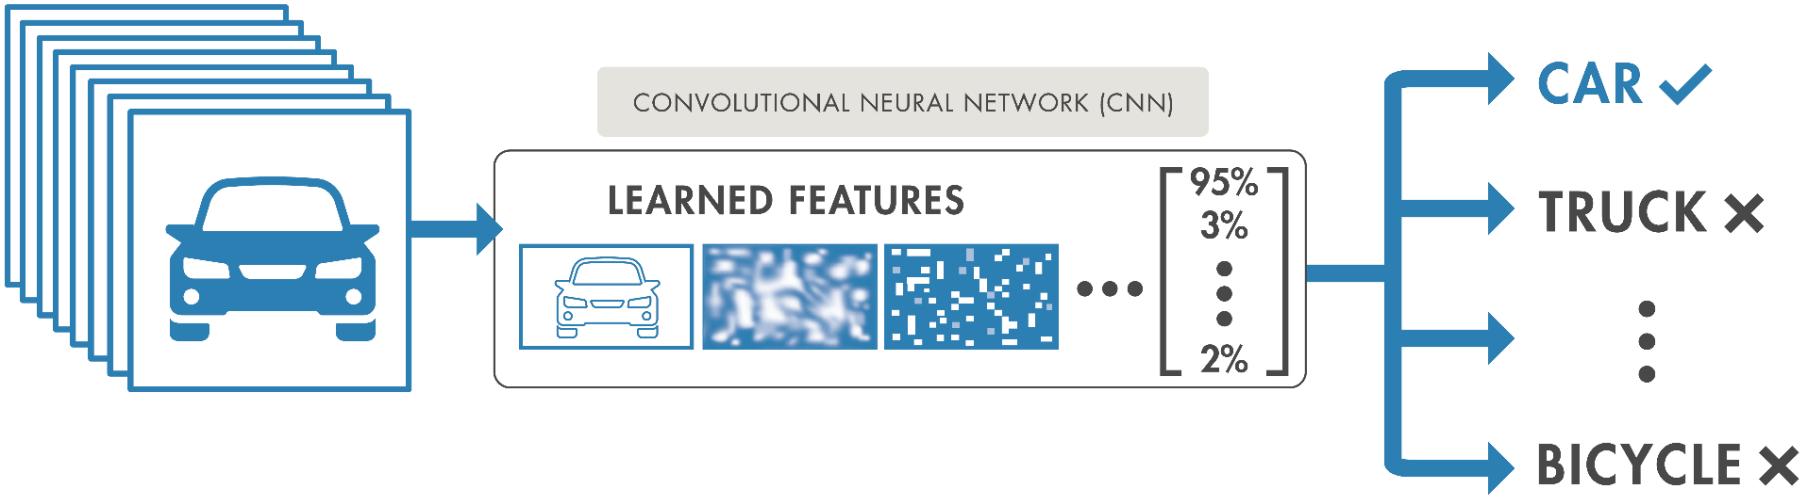

2. Fine-tune a pre-trained model (transfer learning)

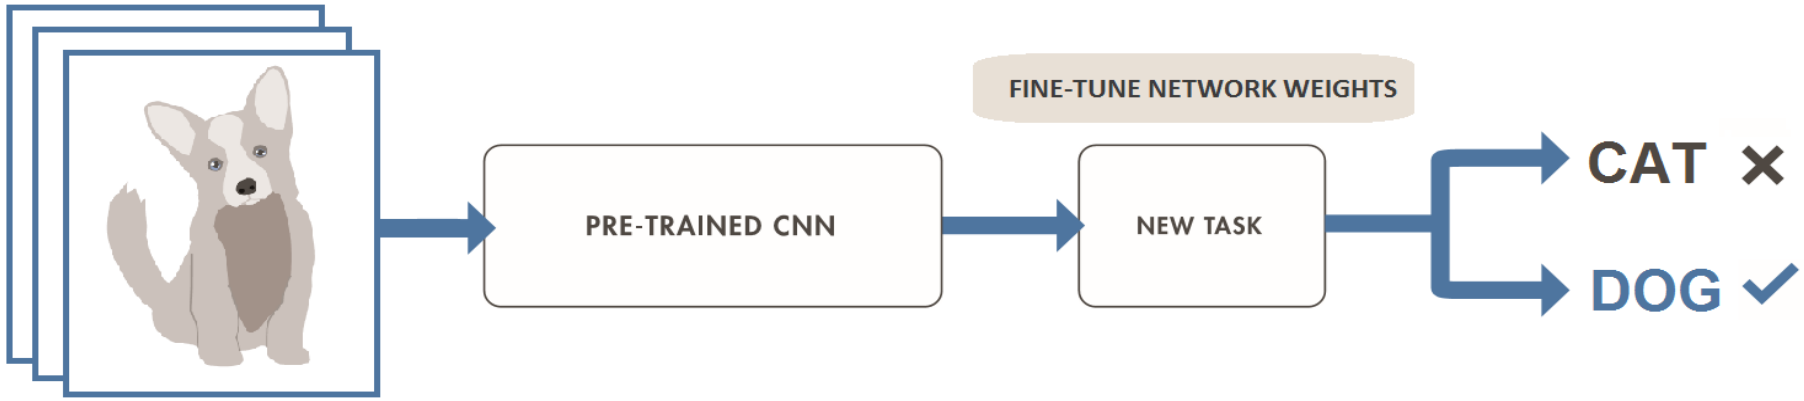

Transfer Learning Workflow

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

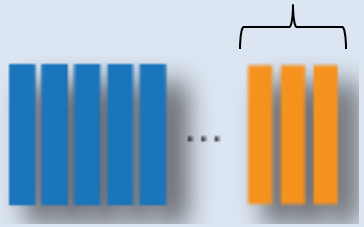

Fewer classes
Learn faster

Train network

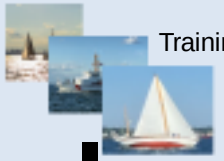

Training images

Training options

100s images
10s classes

Predict and assess network accuracy

Test images

Trained Network

Deploy results

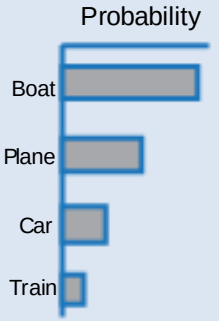

Transfer Learning Workflow – Step 1

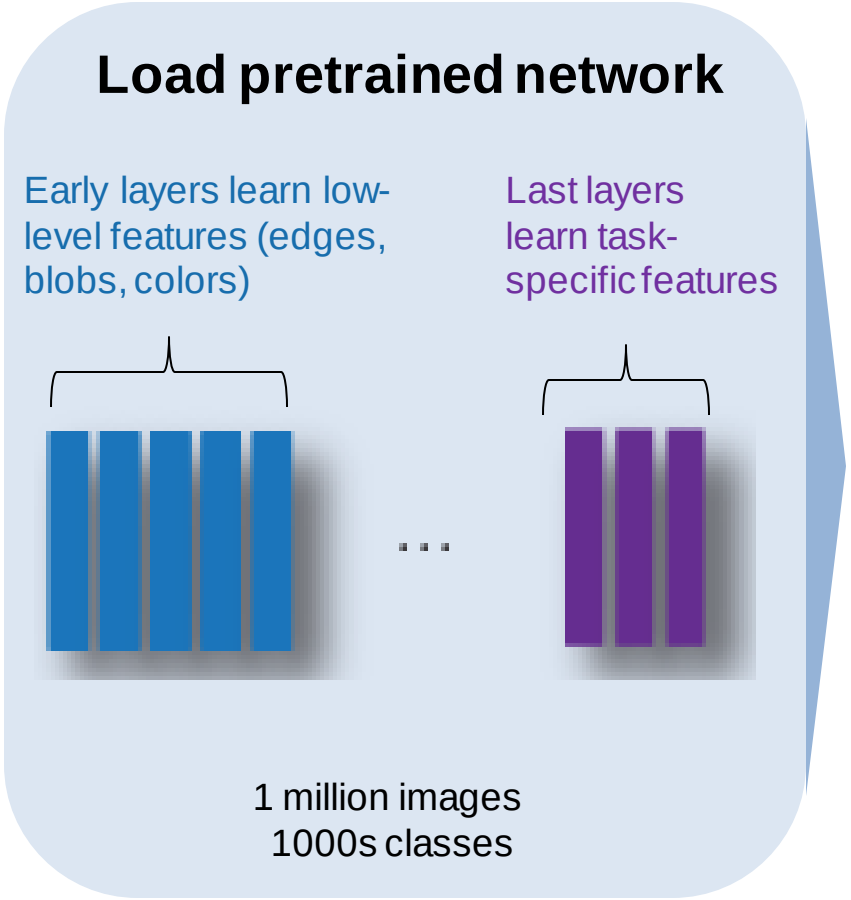

Transfer Learning Workflow – Step 2

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

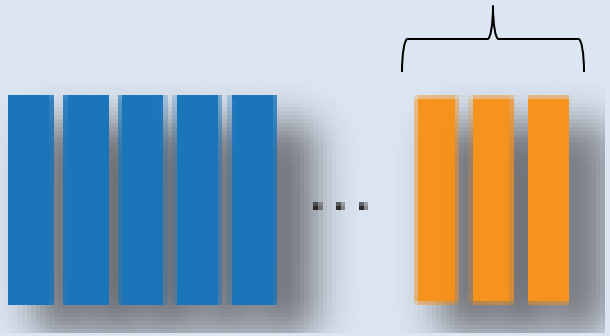

Fewer classes
Learn faster

Transfer Learning Workflow – Step 3

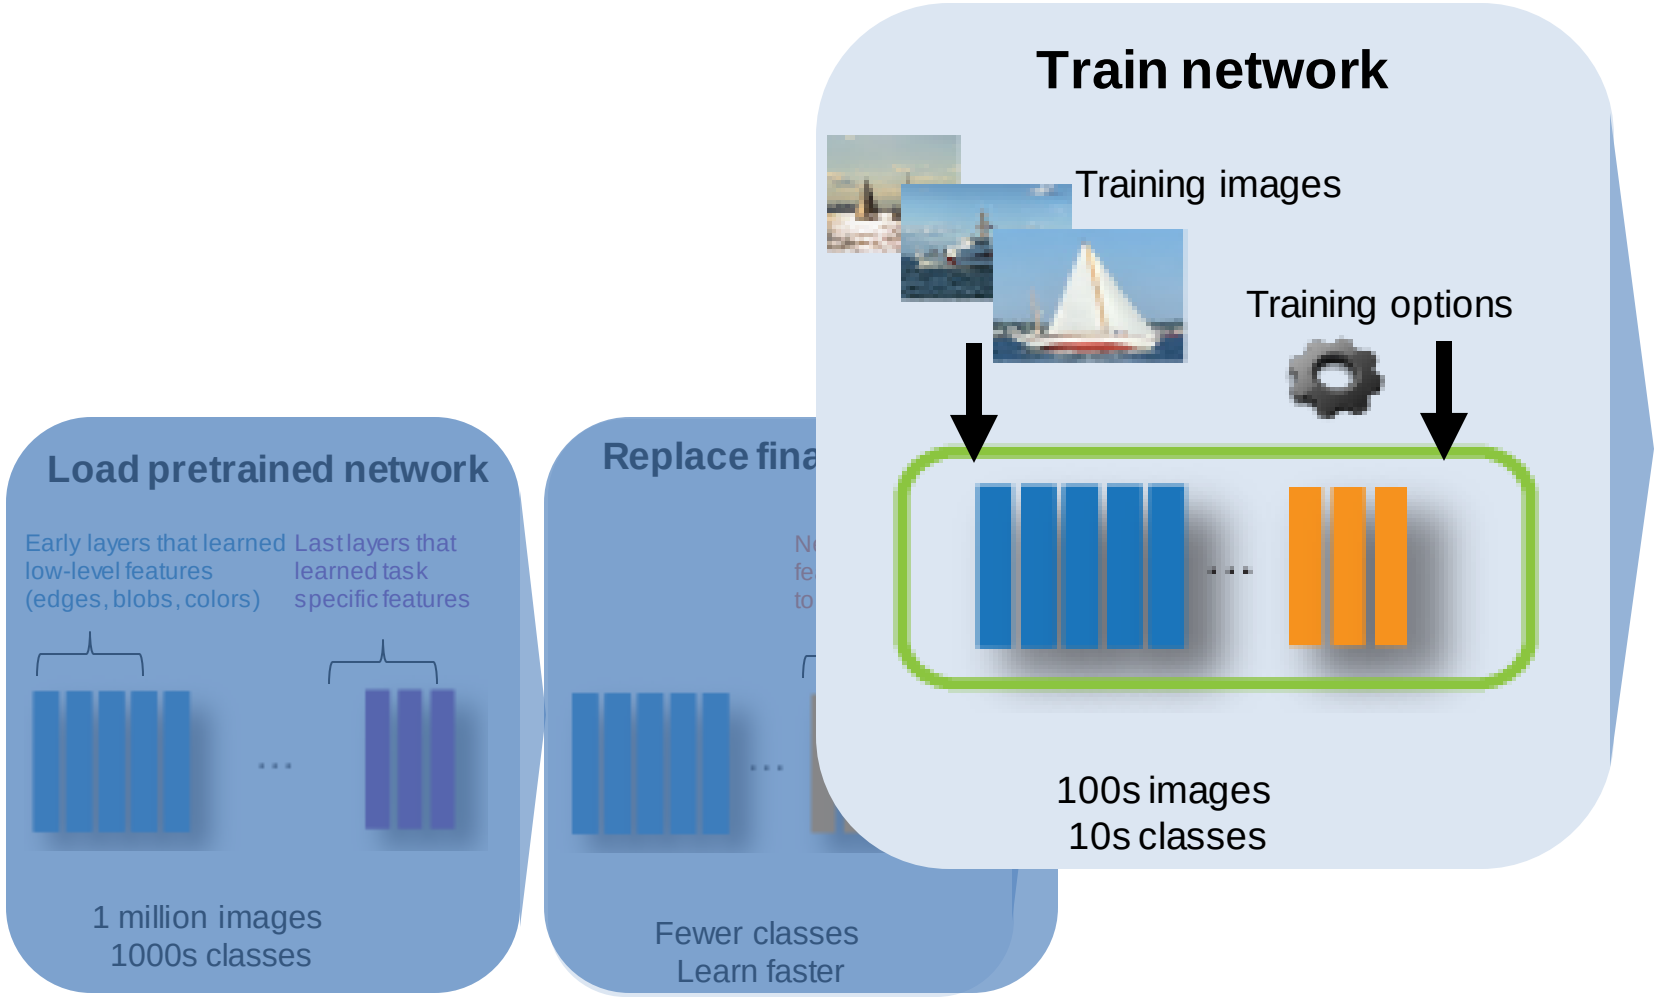

Transfer Learning Workflow – Step 4

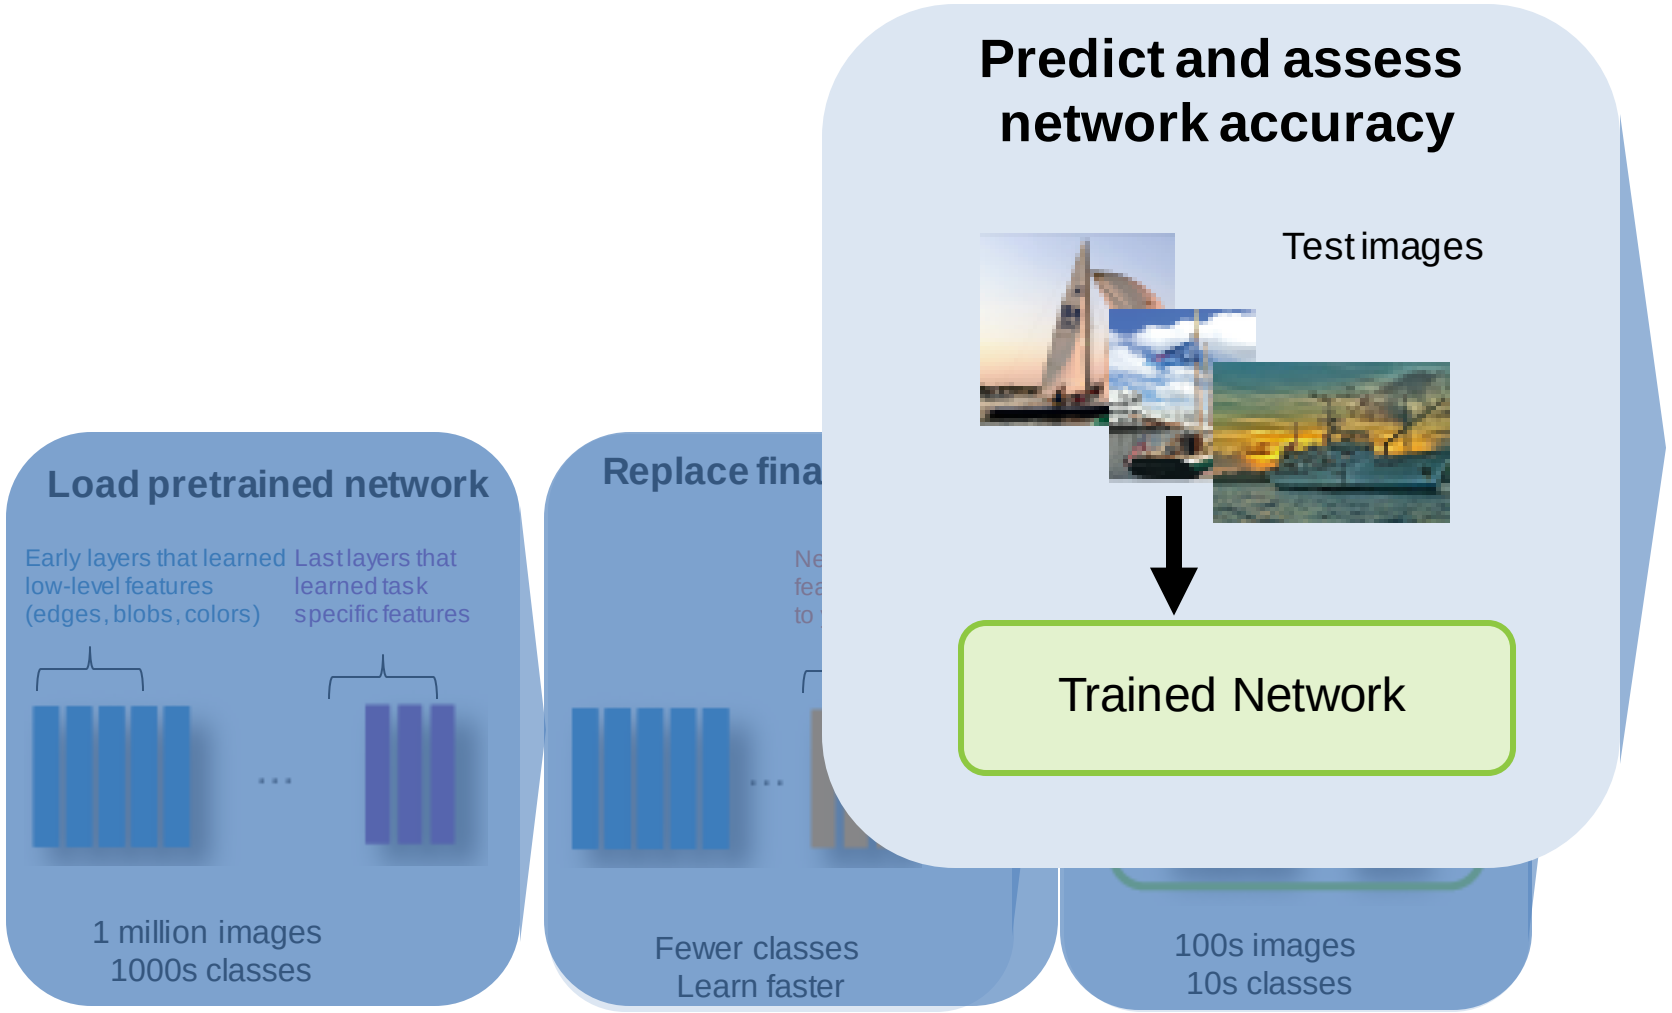

Transfer Learning Workflow – Step 5

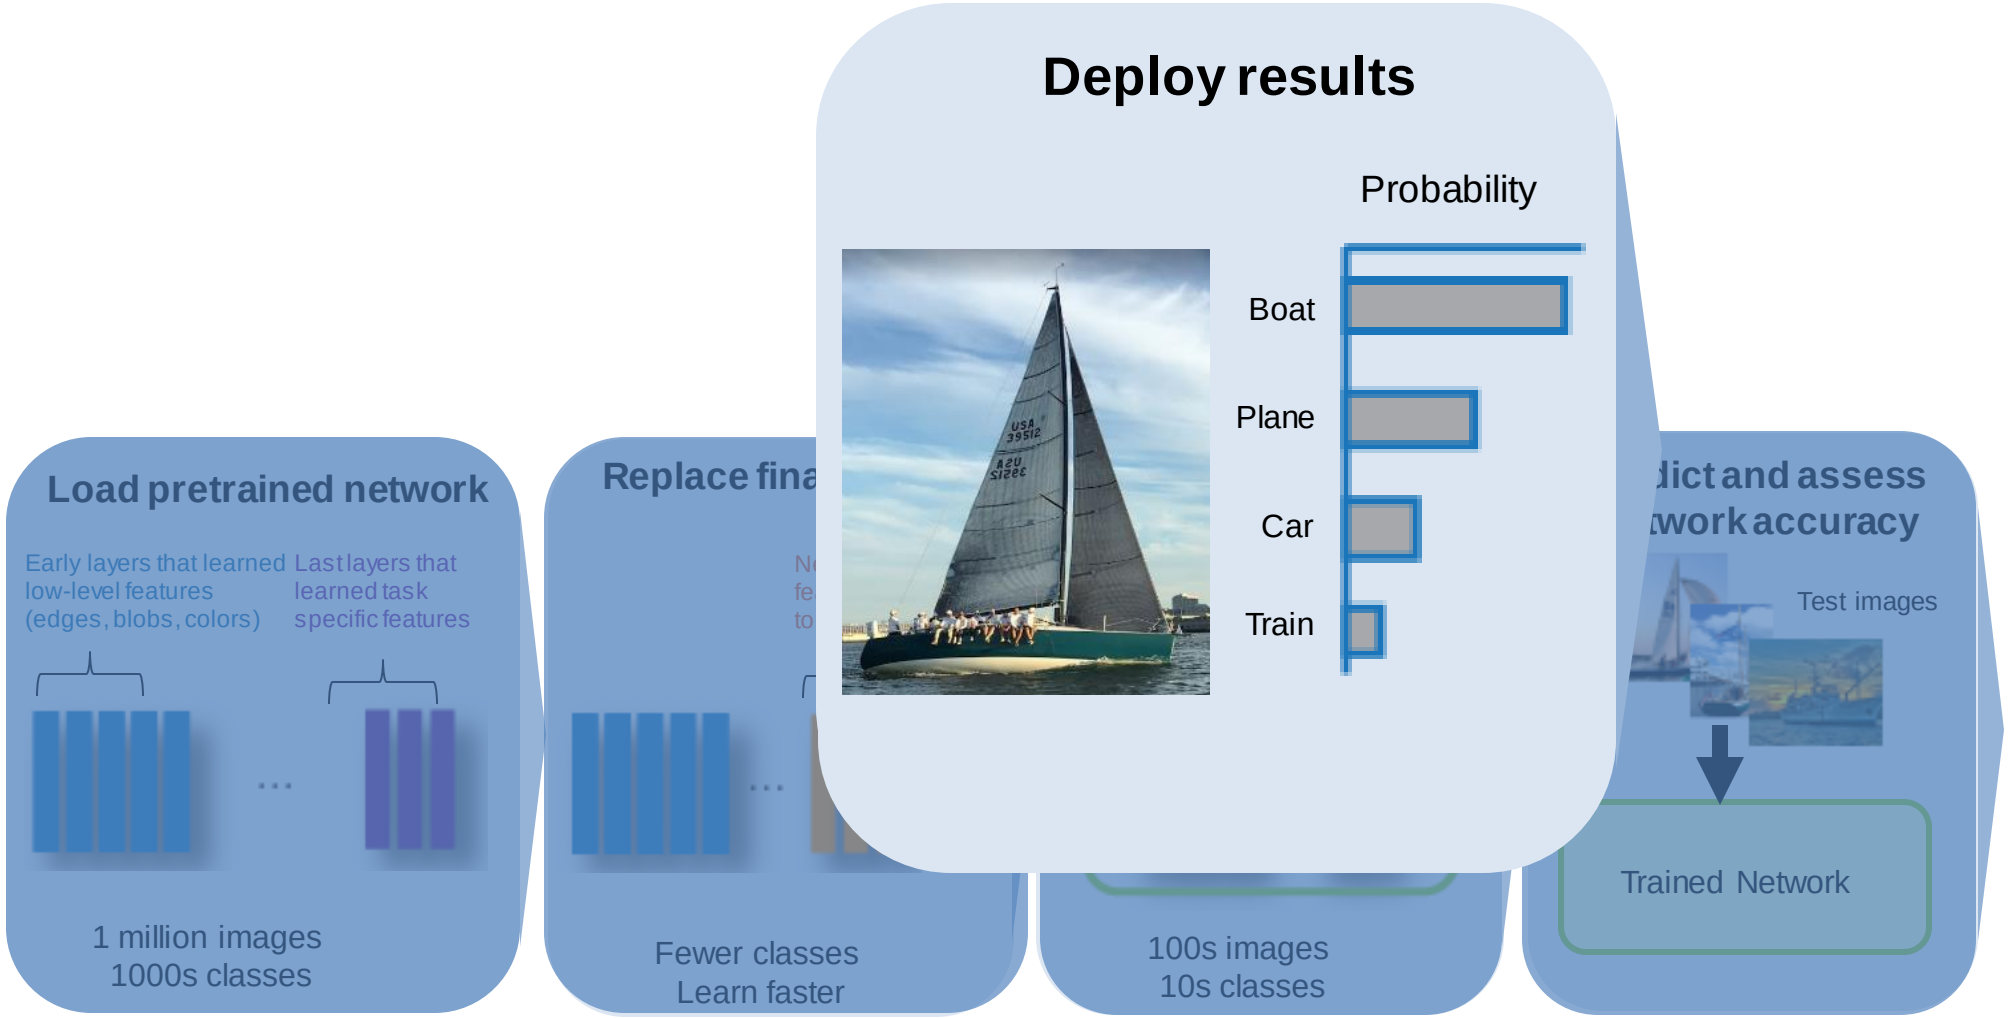

Transfer Learning Workflow

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

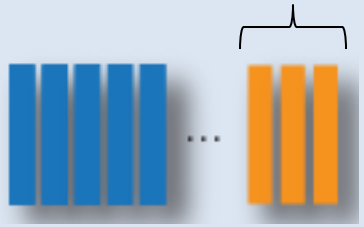

Fewer classes
Learn faster

Train network

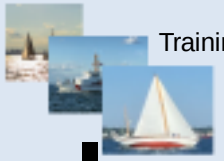

Training images

Training options

100s images
10s classes

Predict and assess network accuracy

Test images

Trained Network

Deploy results

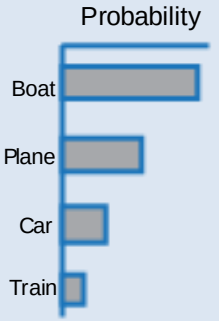

Let's try it out!

Exercise: ***Work_SeeFoodTransferLearning.mlx***
in folder 04-TransferLearningExercise

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
- MATLAB lets you visualize activations in a network

Which Type of Network for Time Series, Signals, and Text?

Recurrent Neural Networks

Classification

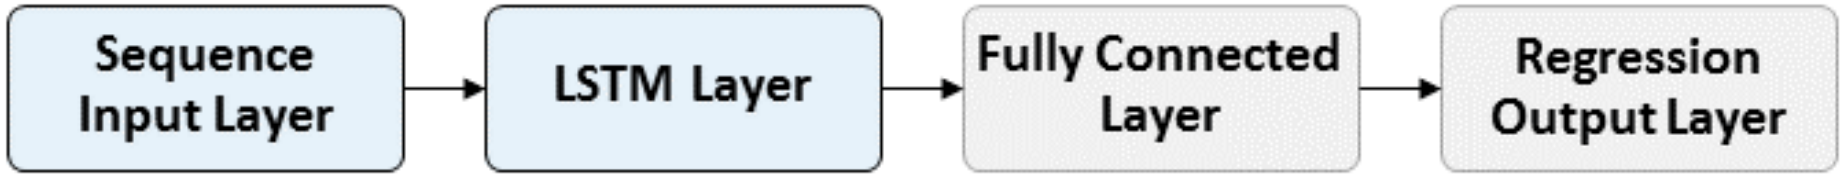

Regression

Types of Datasets

Numeric Data

ID	WC_TA	RE_TA	EBIT_TA	MVE_BVTD	S_TA	Industry	Rating
62394	0.013	0.104	0.036	0.447	0.142	3	BB
48608	0.232	0.335	0.062	1.969	0.281	8	A
42444	0.311	0.367	0.074	1.935	0.366	1	A
48631	0.194	0.263	0.062	1.017	0.228	4	BBB
43768	0.121	0.413	0.057	3.647	0.466	12	AAA
39255	-0.117	-0.799	0.01	0.179	0.082	4	CCC
62236	0.087	0.158	0.049	0.816	0.324	2	BBB
39354	0.005	0.181	0.034	2.597	0.388	7	AA
40326	0.47	0.752	0.07	11.596	1.12	8	AAA
51681	0.11	0.337	0.045	3.835	0.812	4	AAA

Machine Learning or **LSTM**

Time Series Signal Text

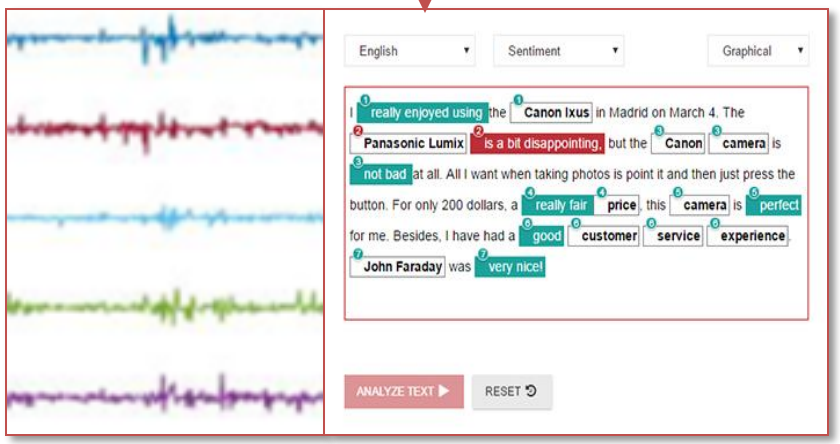

LSTM or CNN

Image Data

CNN

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

I was born in France…

… I speak _____ ?

Take into account previous data when making new predictions

Recurrent Neural Networks

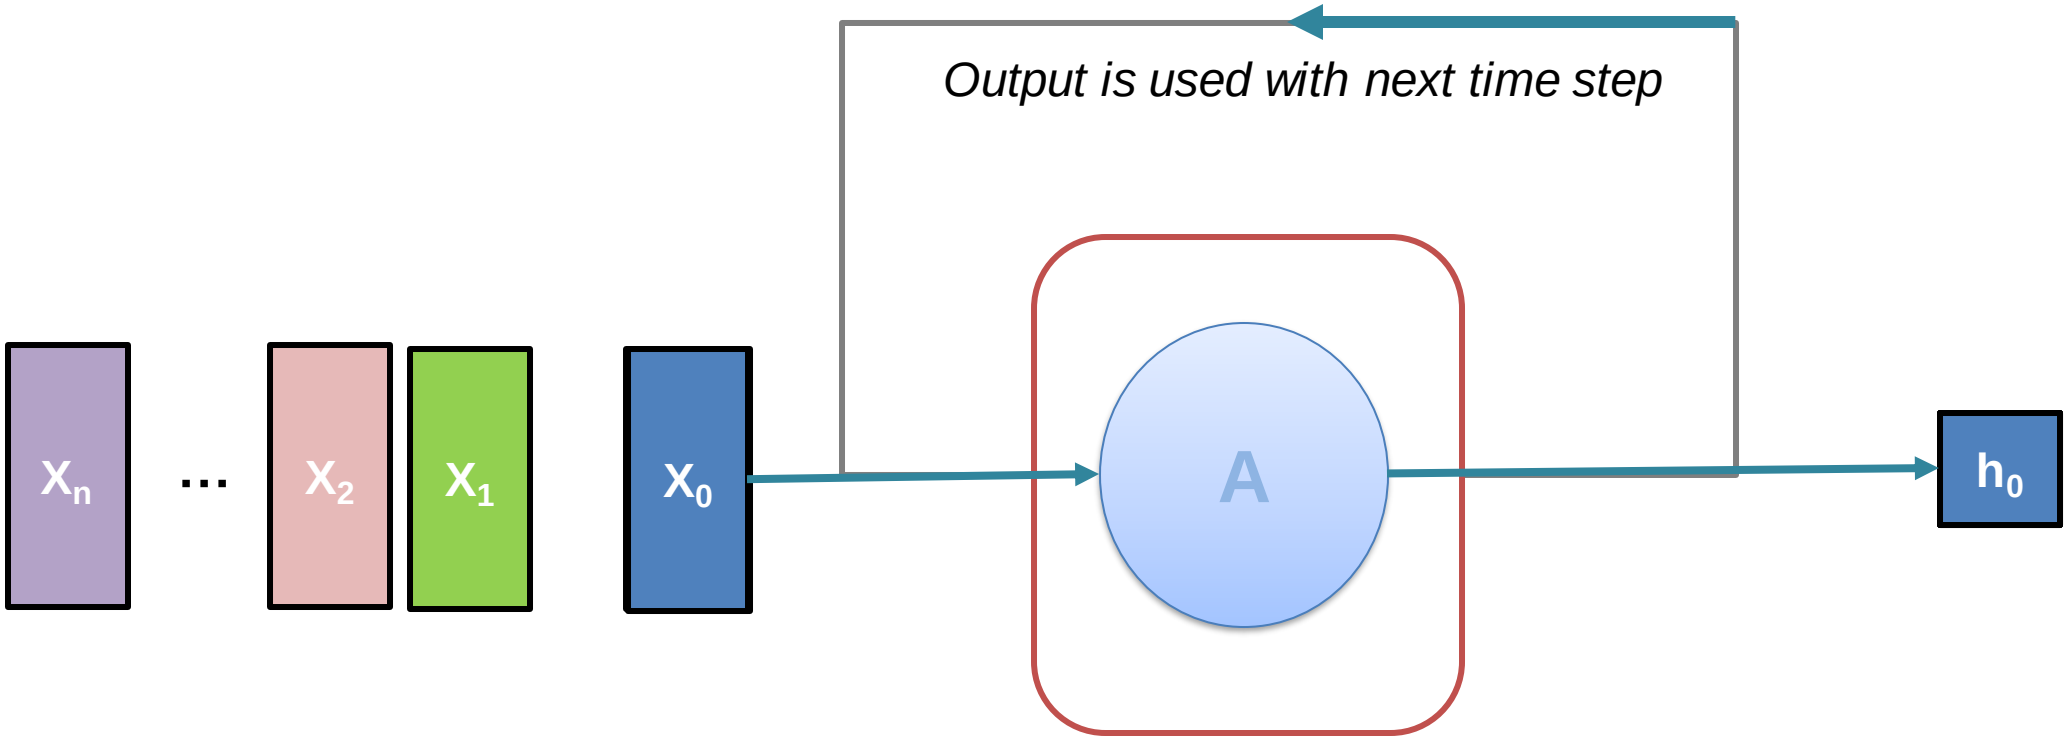

I was born in France…

[2000 words]

… I speak _____ ?

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

Long Short-Term Memory Network

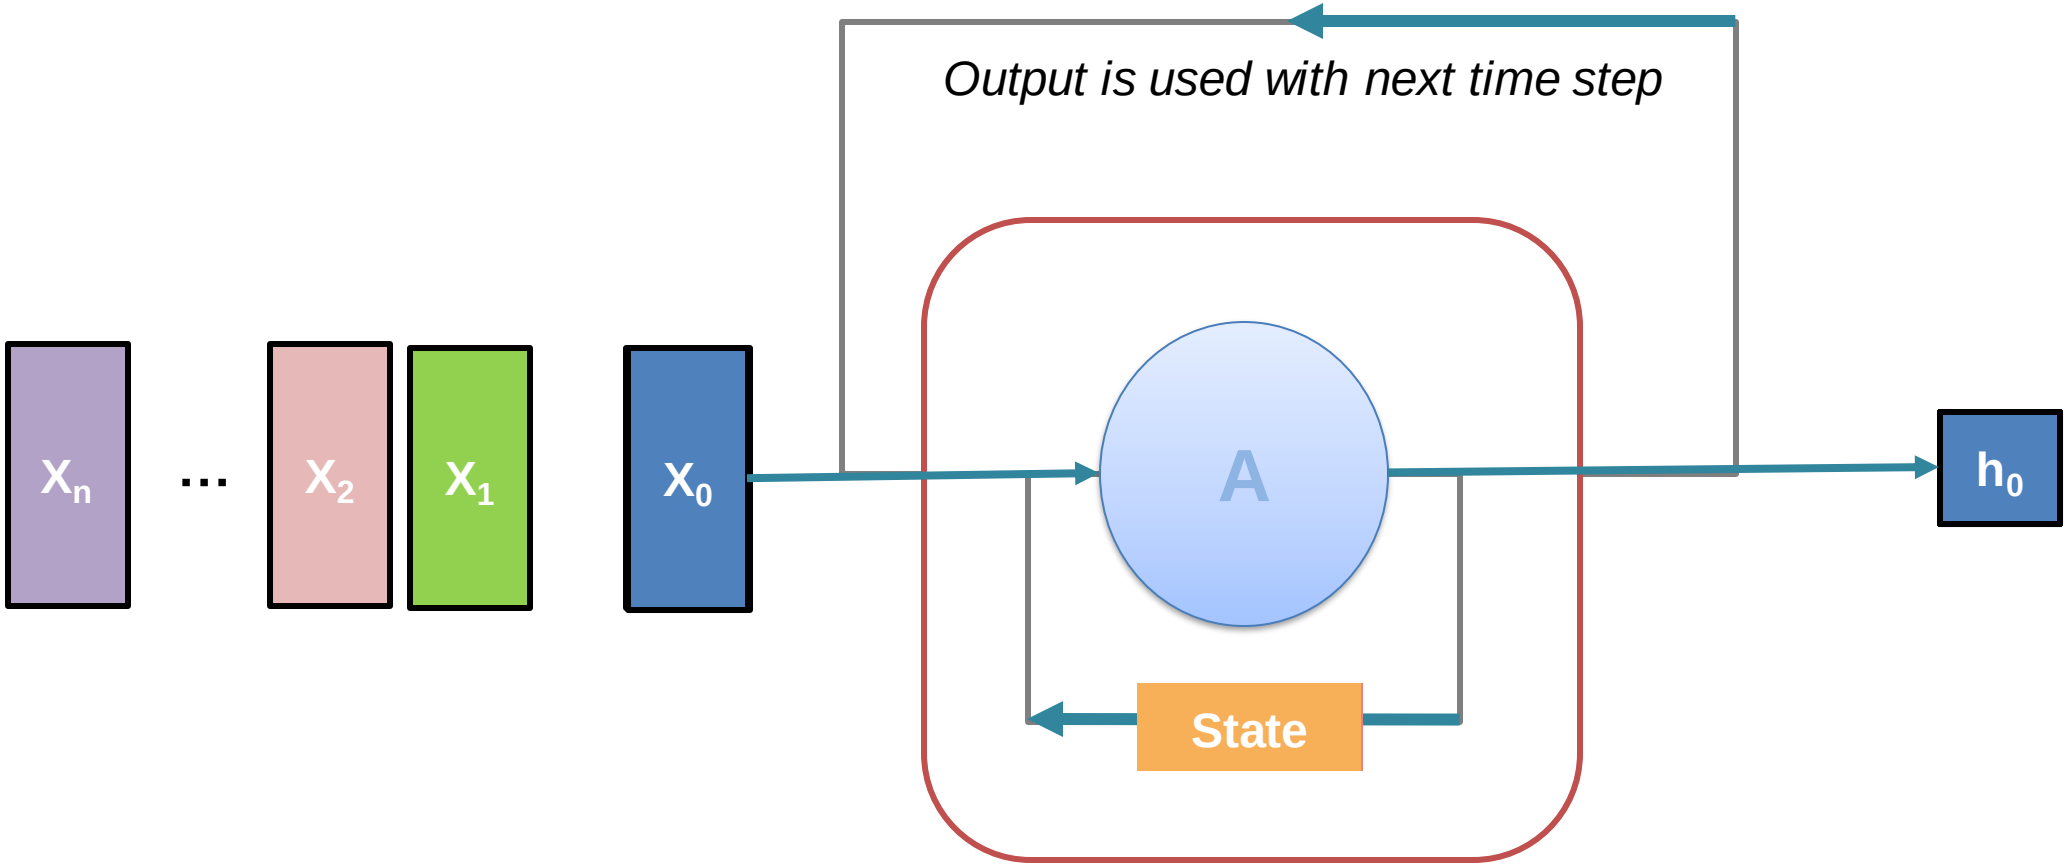

Examples in MATLAB Documentation

Sequence Classification Using Deep Learning

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

Open Live Script

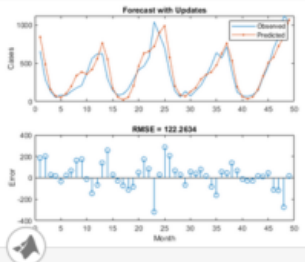

Time Series Forecasting Using Deep Learning

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

Open Live Script

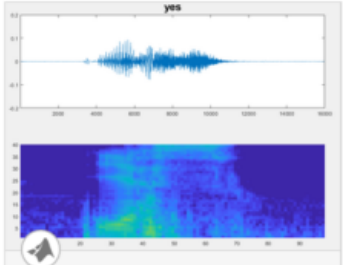

Speech Command Recognition Using Deep Learning

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

Open Script

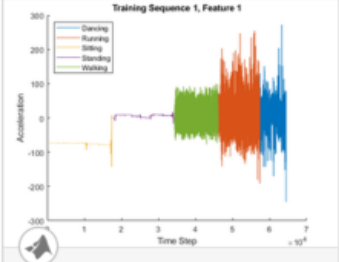

Sequence-to-Sequence Classification Using Deep Learning

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

Open Live Script

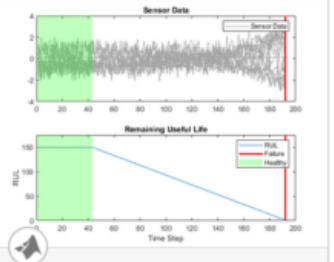

Sequence-to-Sequence Regression Using Deep Learning

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

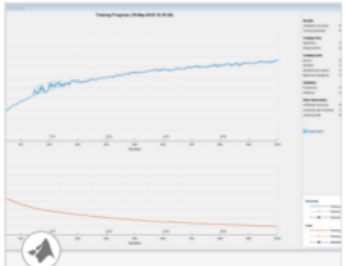

Generate Text Using Deep Learning

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

Highlights

Sequence Classification Using Deep Learning

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

Open Live Script

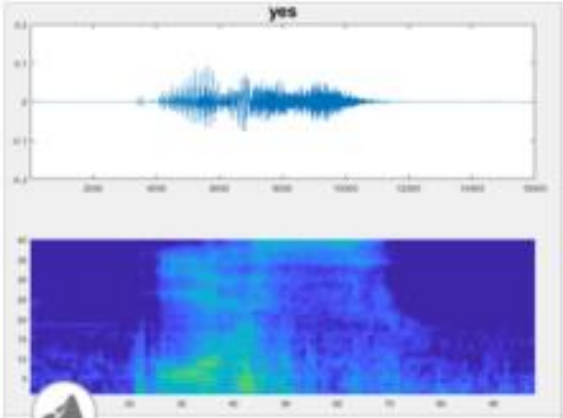

Speech Command Recognition Using Deep Learning

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

Open Script

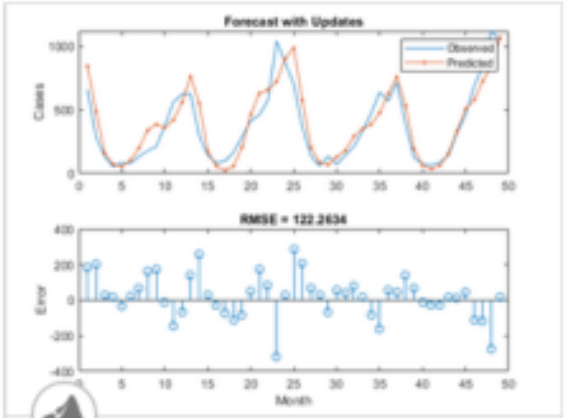

Time Series Forecasting Using Deep Learning

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

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

Let's try it out!

Exercise: ***Work_TimeSeriesForecastingLSTMs.mlx***
in folder 06-LSTMExercise

Deployment process

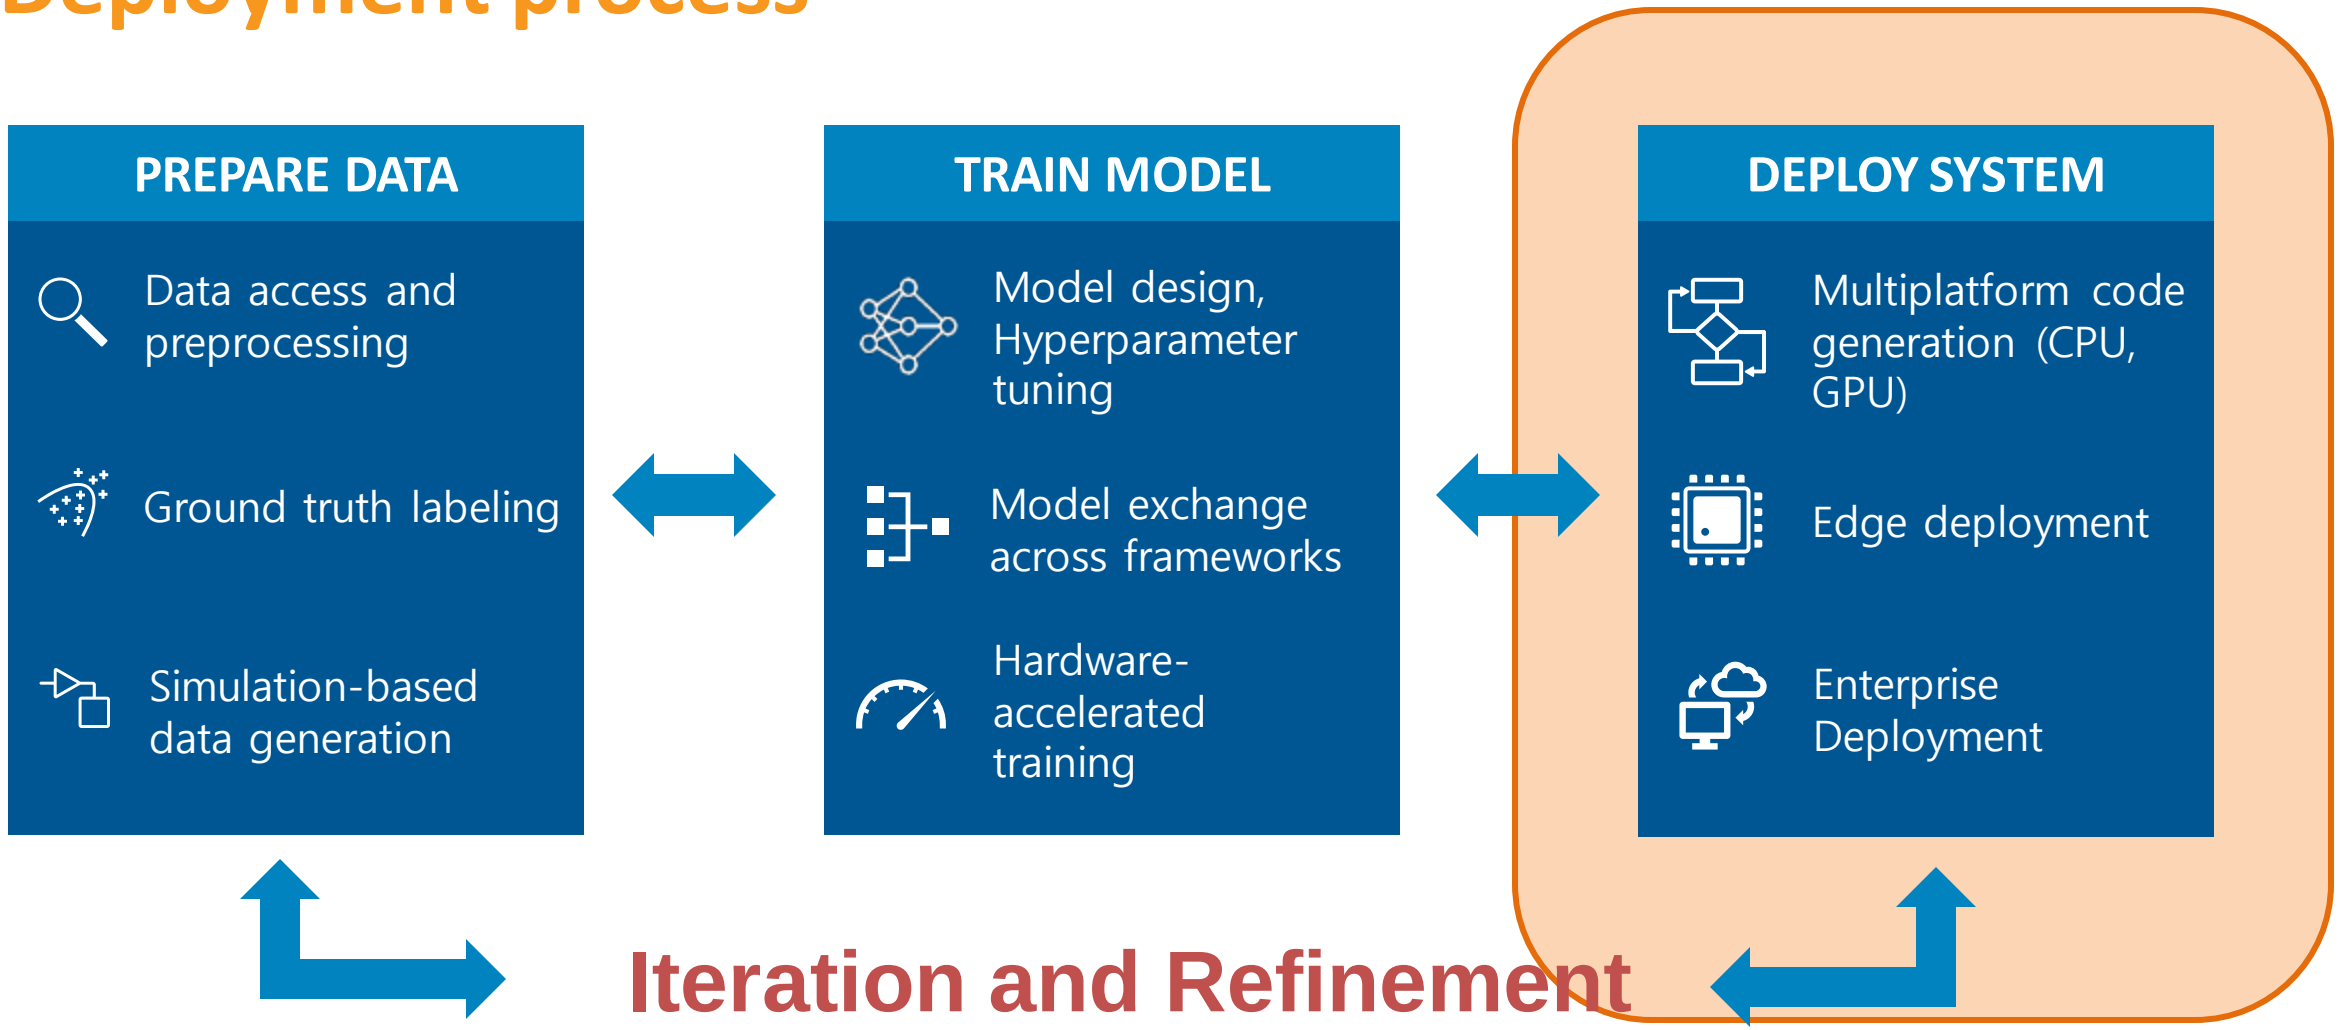

Deployment and Scaling for A.I.

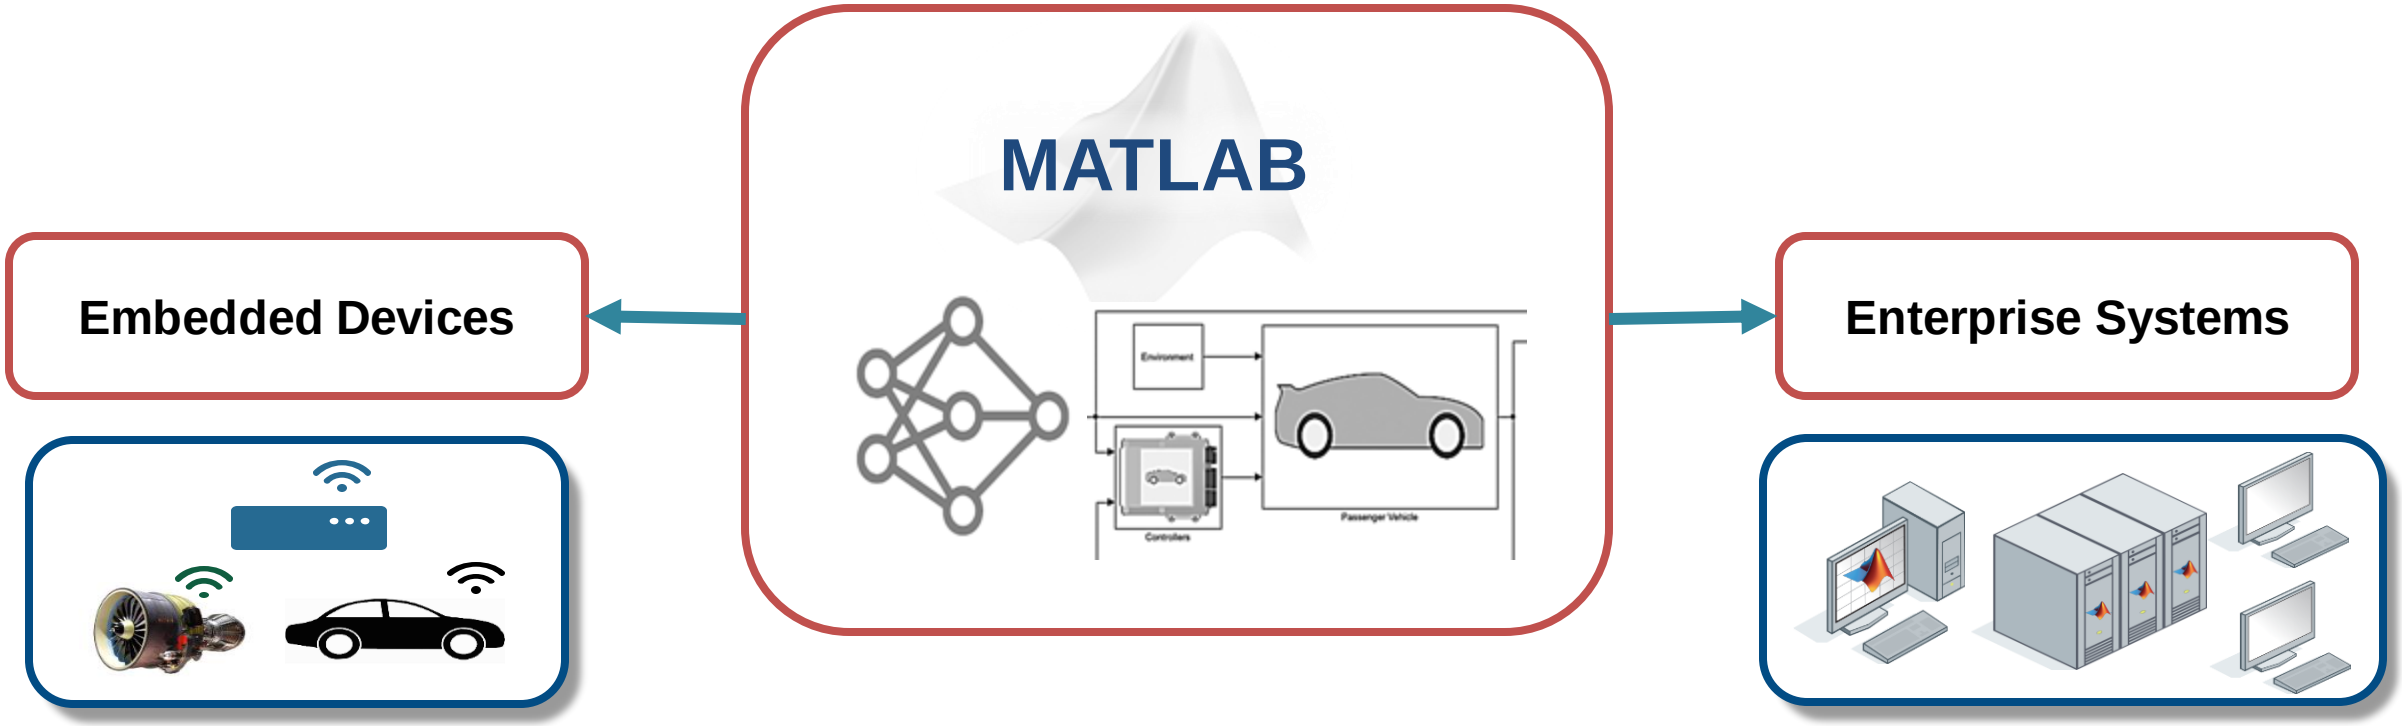

Deploying Deep Learning Models for Inference

Deep Learning Networks

Coder Products

MATLAB Coder

GPU Coder

MATLAB Coder

Intel MKL-DNN Library

NVIDIA TensorRT & cuDNN Libraries

ARM Compute Library

Deployment Target Examples

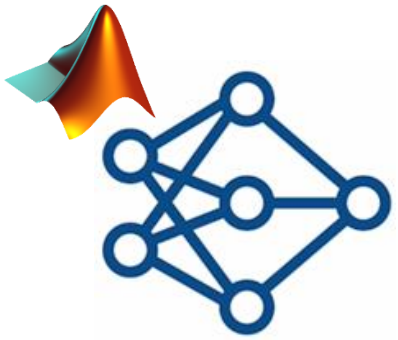

Deep Learning Networks

Coder Products

MATLAB Coder

Desktop CPU

GPU Coder

NVIDIA

NVIDIA TensorRT & cuDNN Libraries

MATLAB Coder

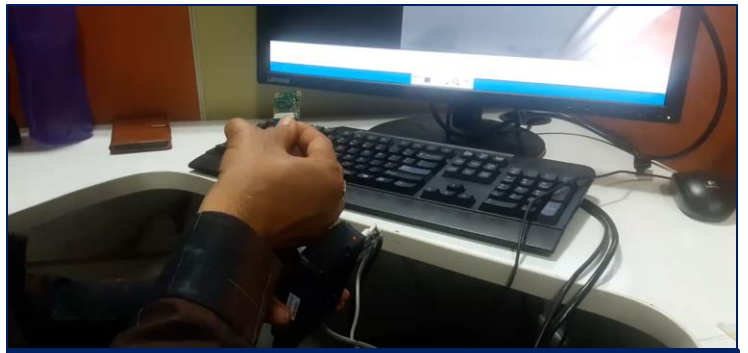

Raspberry Pi board

With GPU Coder, MATLAB is fast

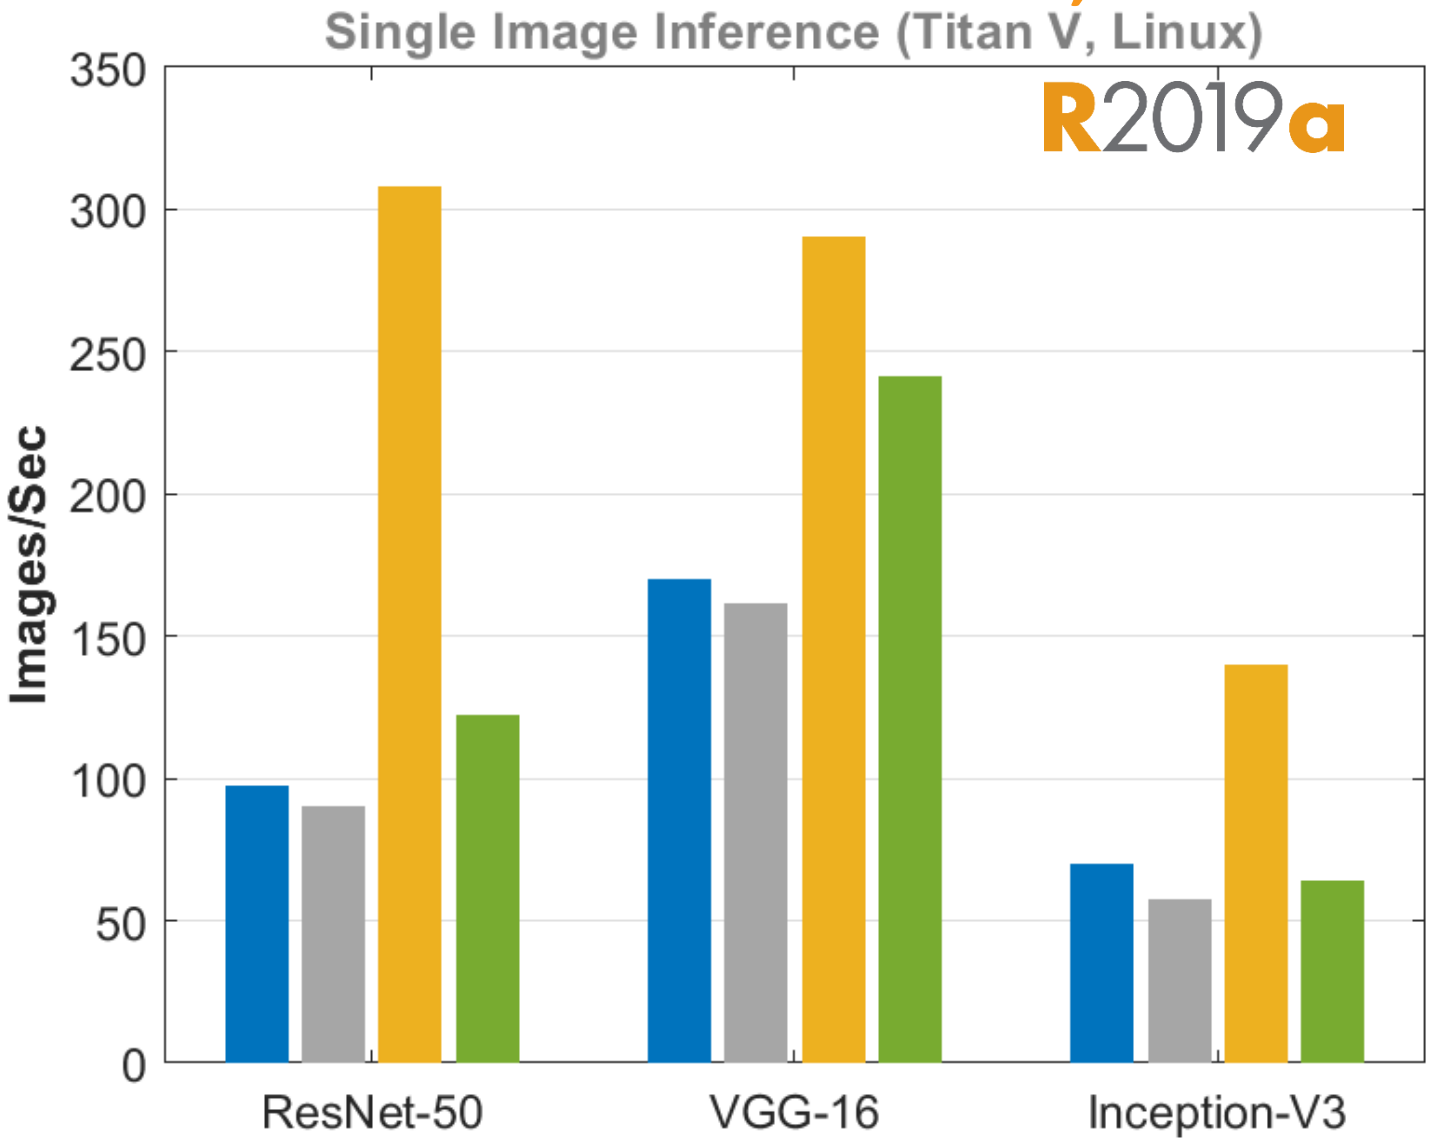

GPU Coder is faster than TensorFlow, MXNet and Pytorch

- TensorFlow
- MXNet
- GPU Coder
- PyTorch

Intel® Xeon® CPU 3.6 GHz - NVIDIA libraries: CUDA10 - cuDNN 7 - Frameworks: TensorFlow 1.13.0, MXNet 1.4.0 PyTorch 1.0.0

GPU Code Generation with YOLOv2

YOLOv2

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

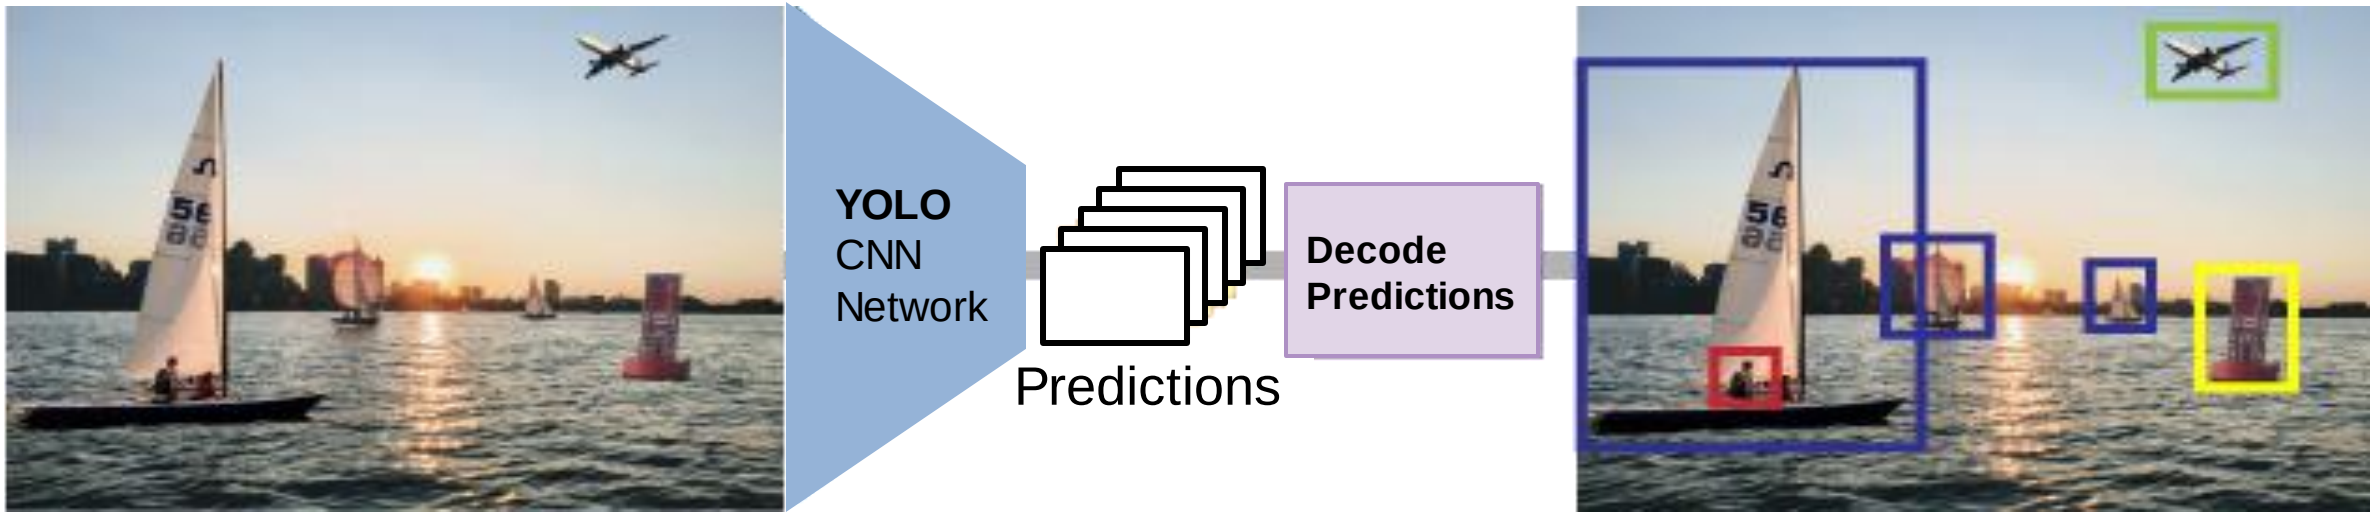

GPU Coder Demo

Deploying our deep network on a GPU

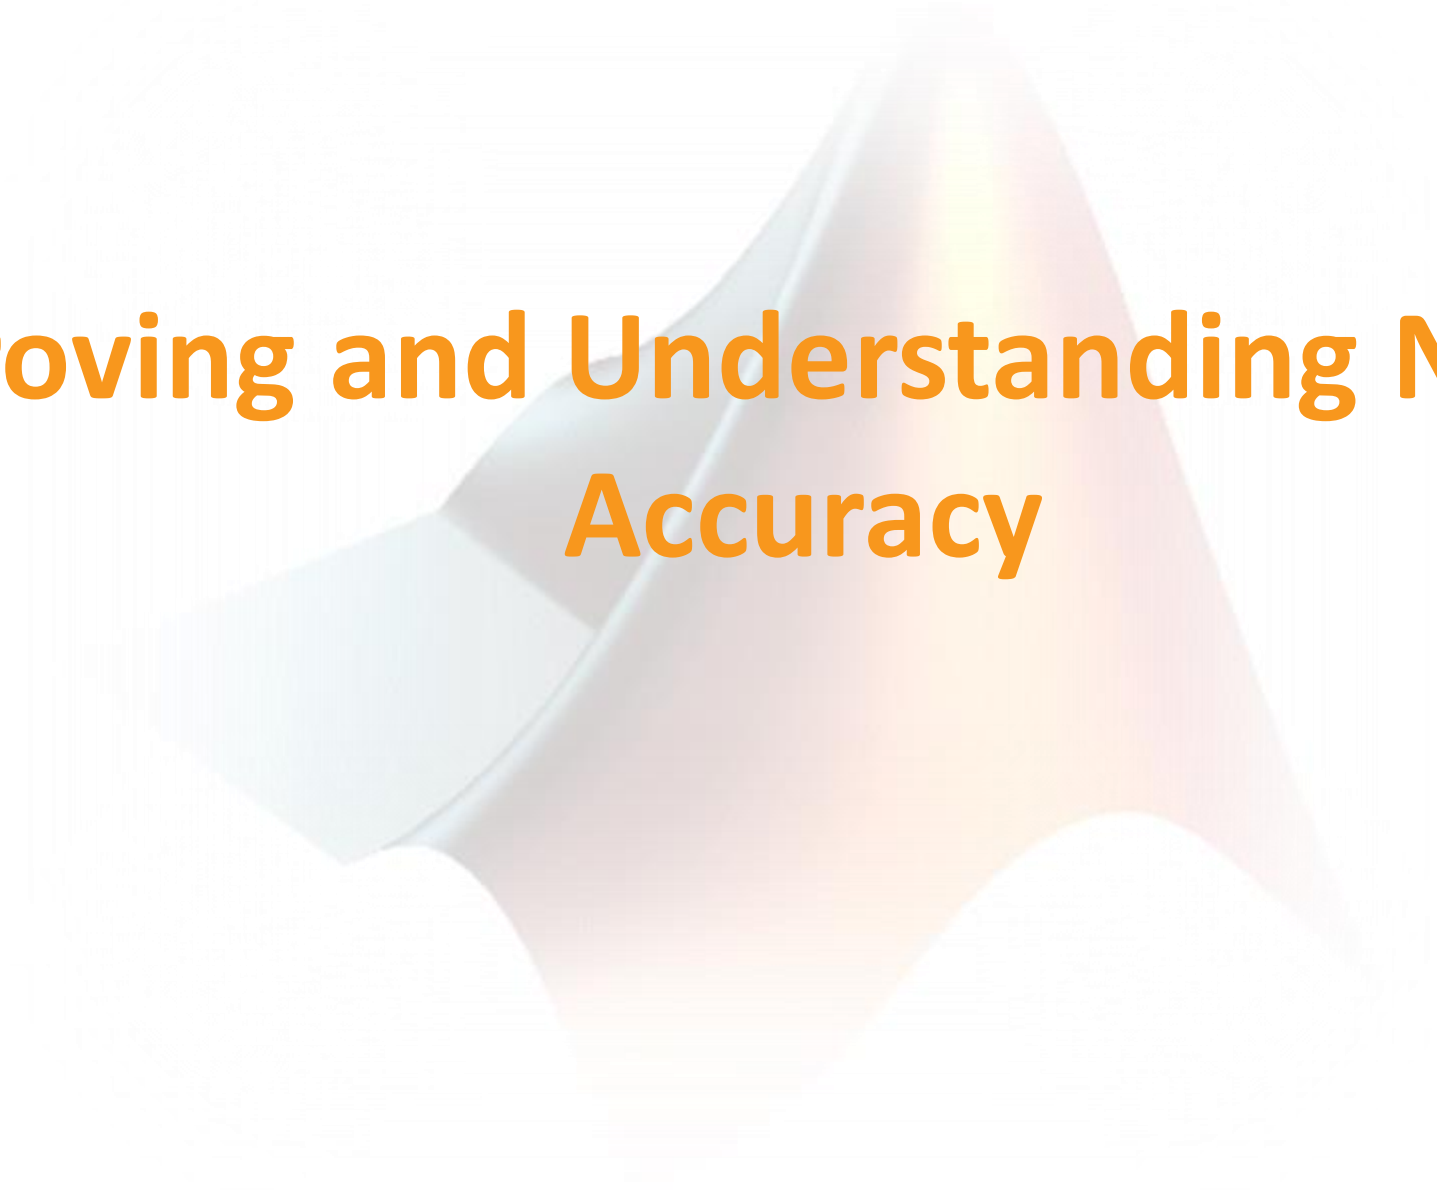

Improving and Understanding Network Accuracy

Common Network Architectures - Signal Processing

Convolutional Neural Networks (**CNN**)

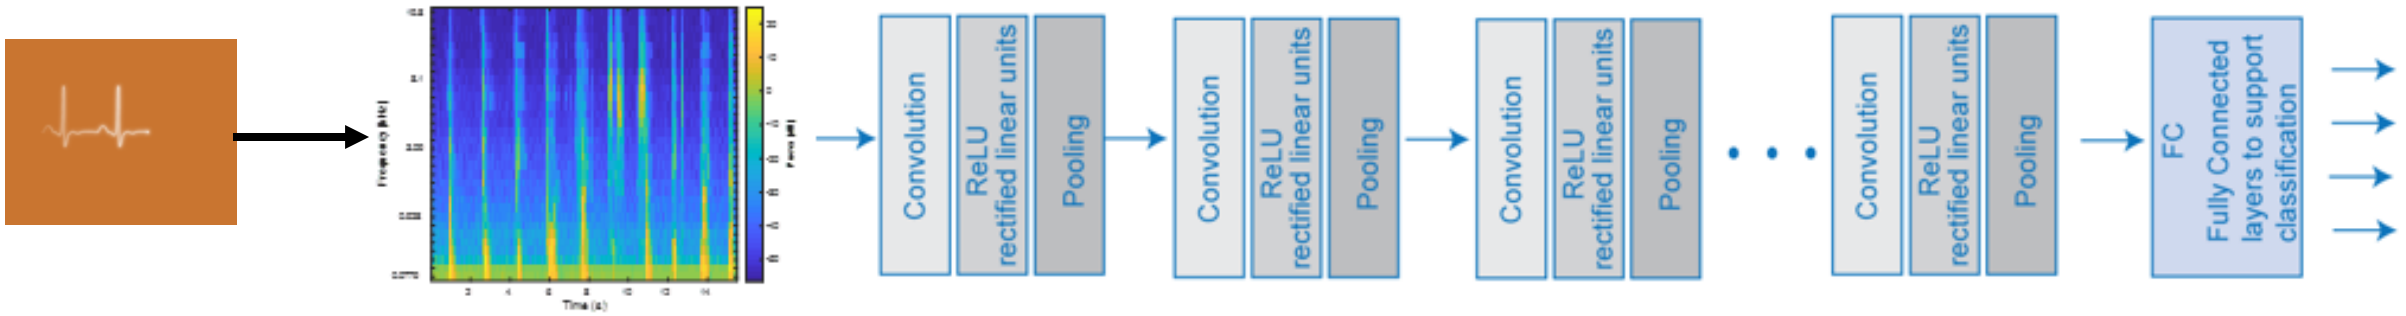

Time-Frequency Transformation

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

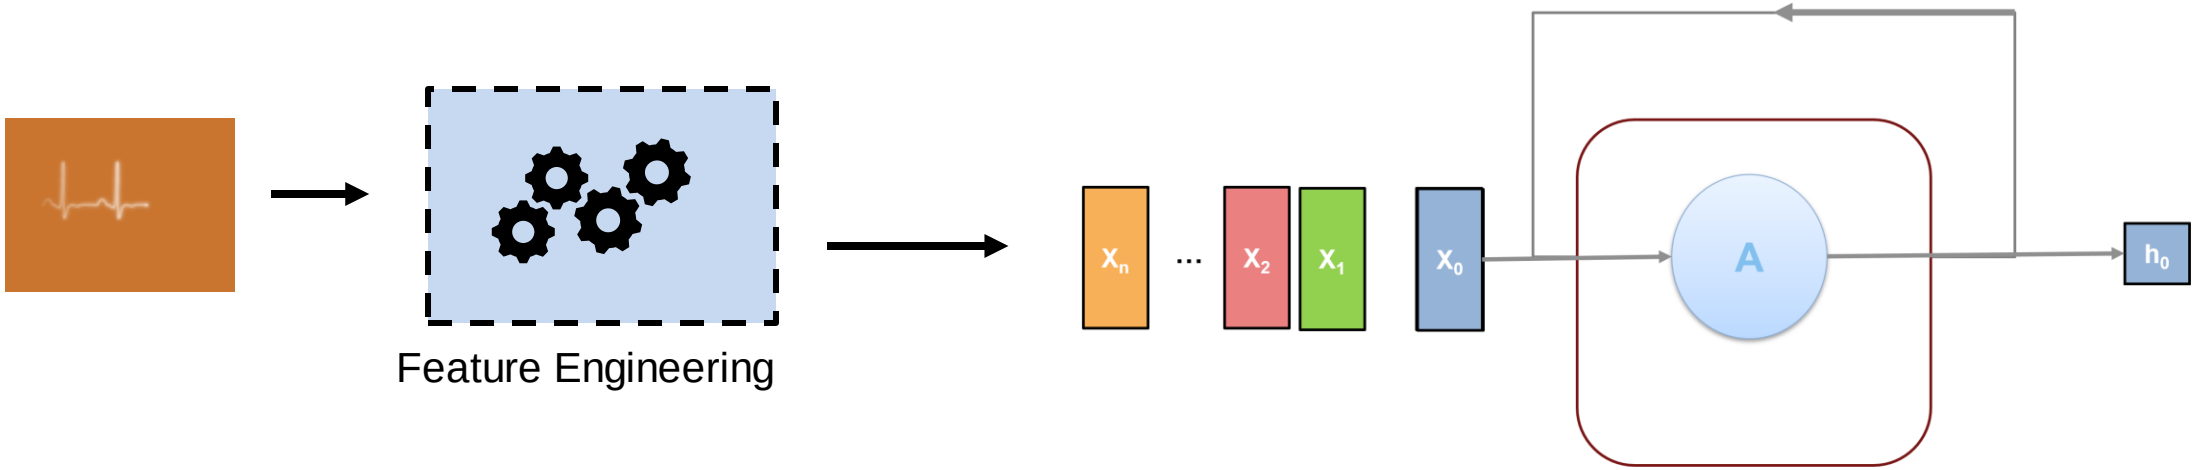

Feature Engineering

Speech Recognition Example

Audio signal → Spectrogram → Image Classification algorithm

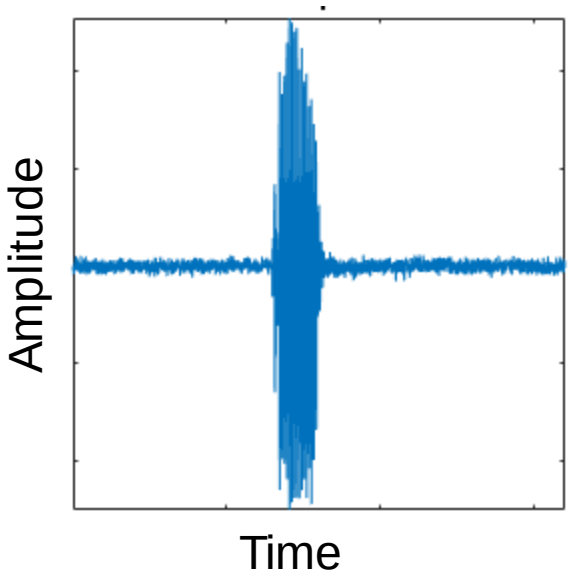

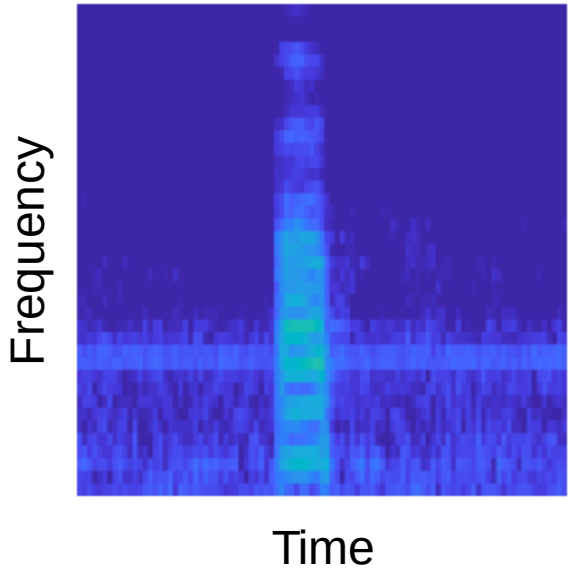

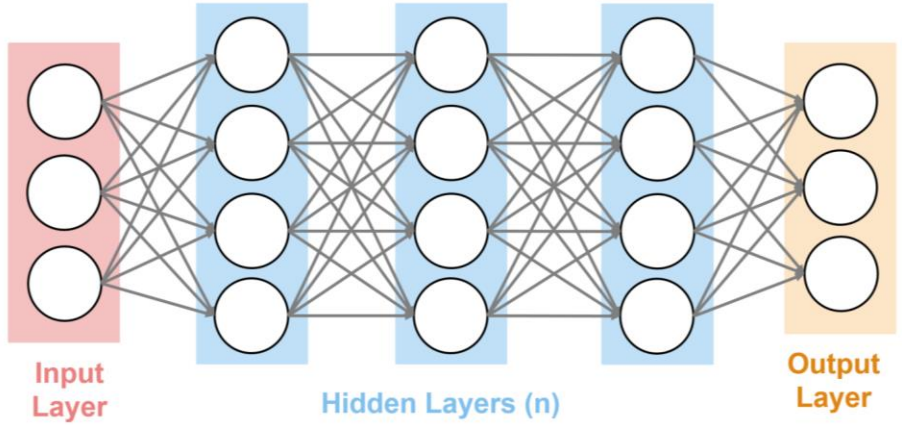

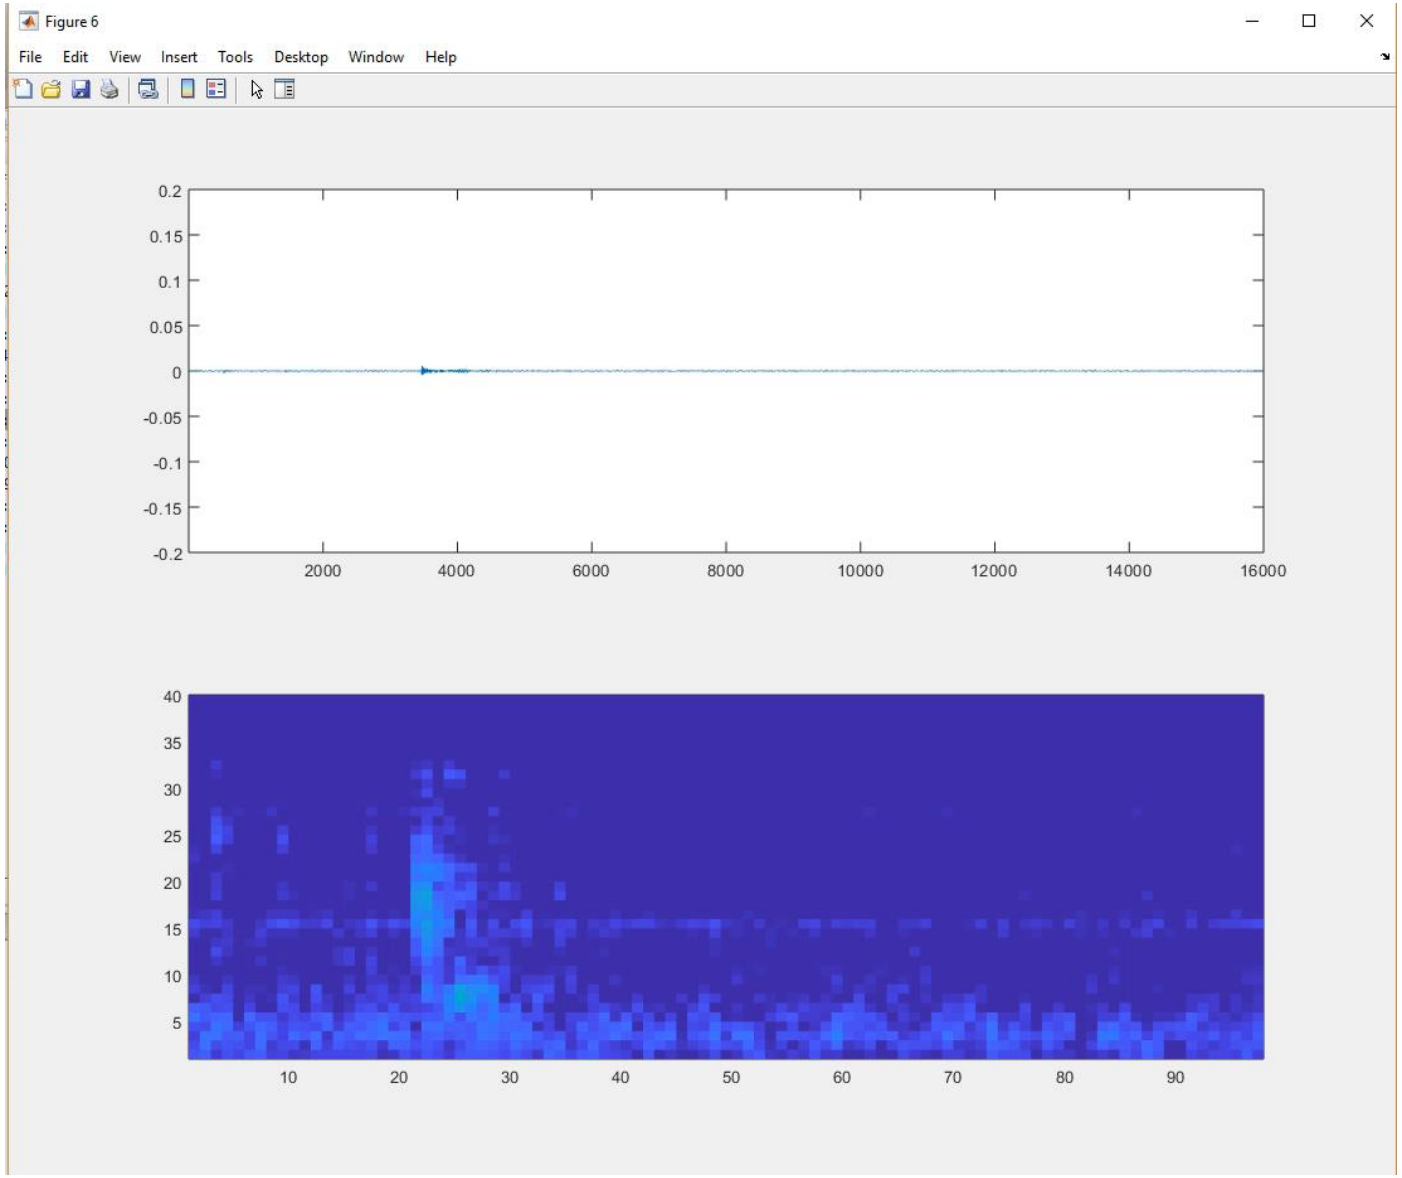

Speech Recognition using CNNs

Example: 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

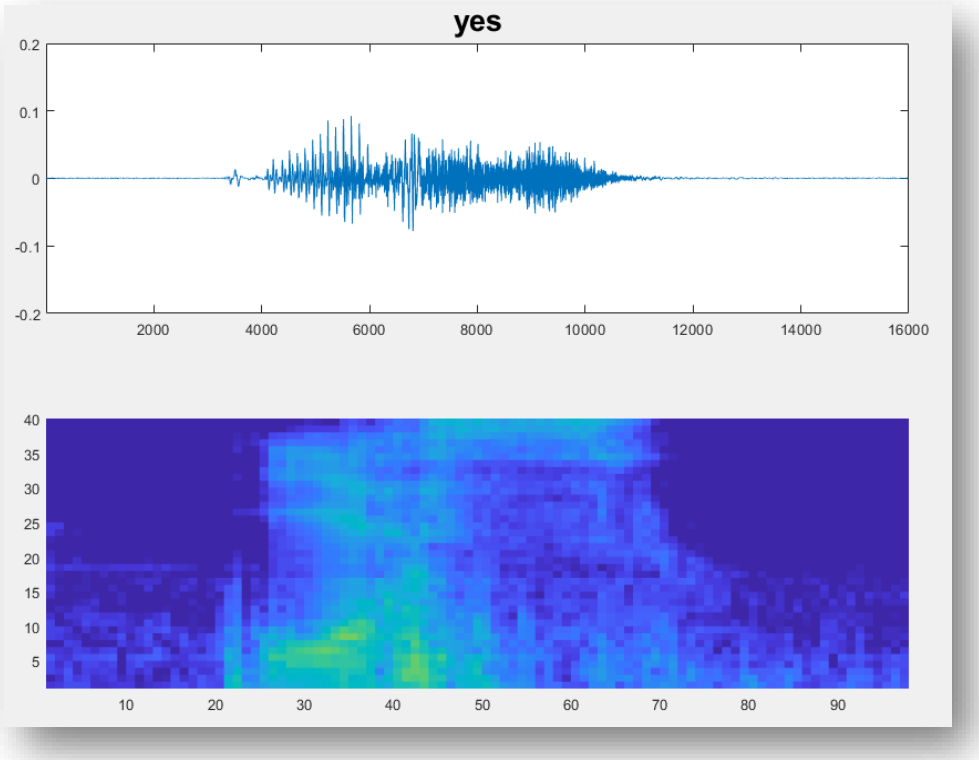

What is Reinforcement Learning?

- What is Reinforcement Learning?
  - Type of machine learning that trains an **‘agent’** through repeated interactions with an environment
- How does it work?
  - Through a trial & error process that uses a reward system to maximize success

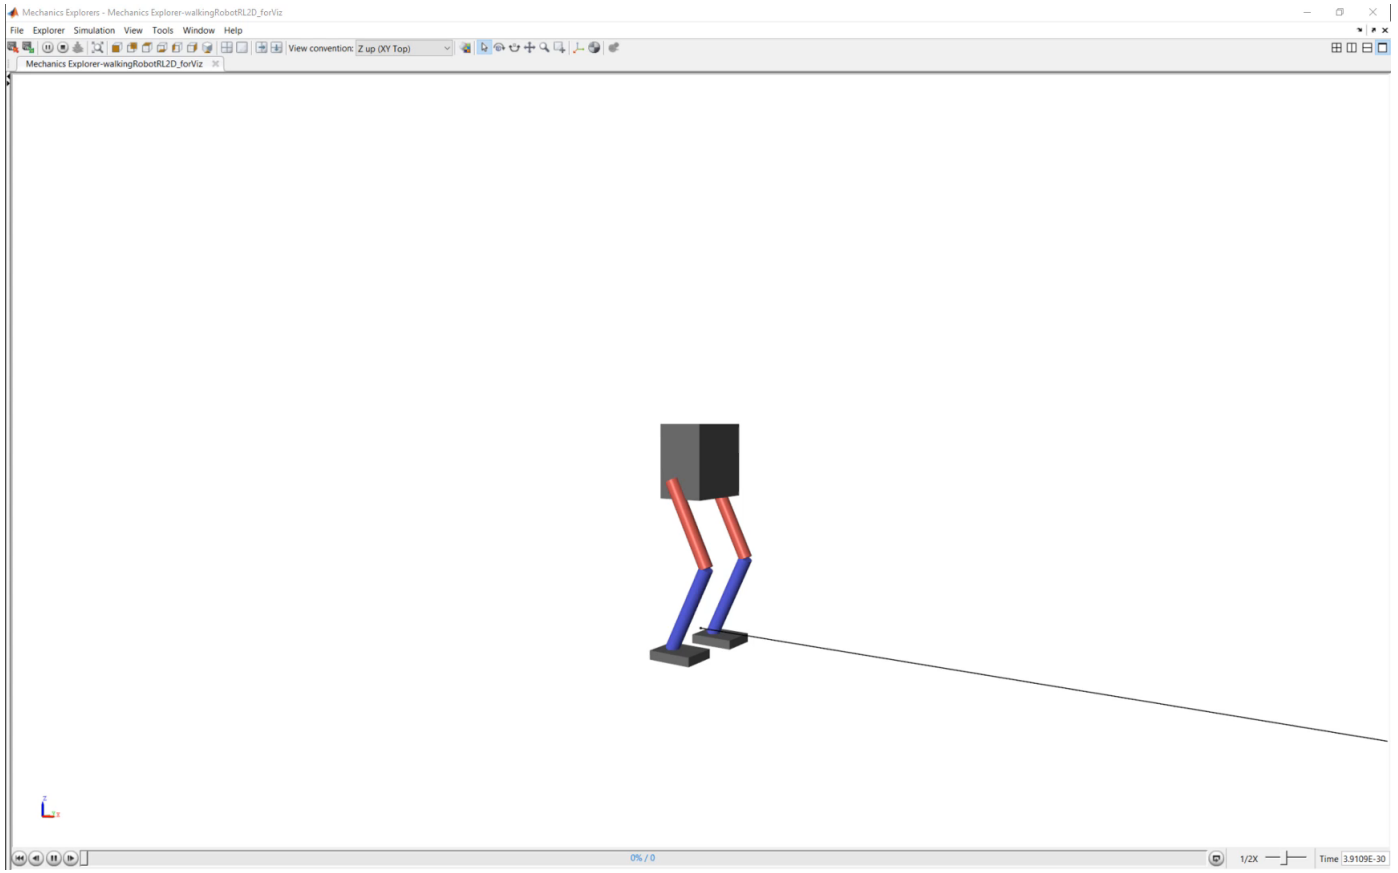

Reinforcement Learning enables the use of Deep Learning for Controls and Decision Making Applications

Controls

Robotics

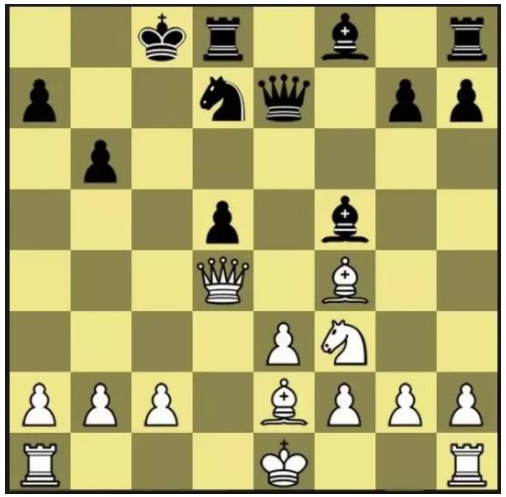

A.I. Gameplay

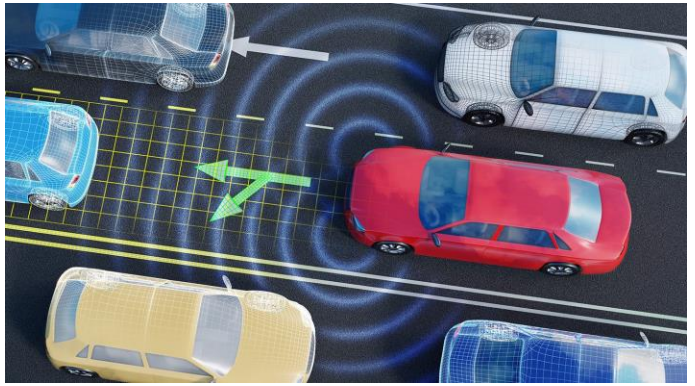

Autonomous driving

How Does Reinforcement Learning Work?

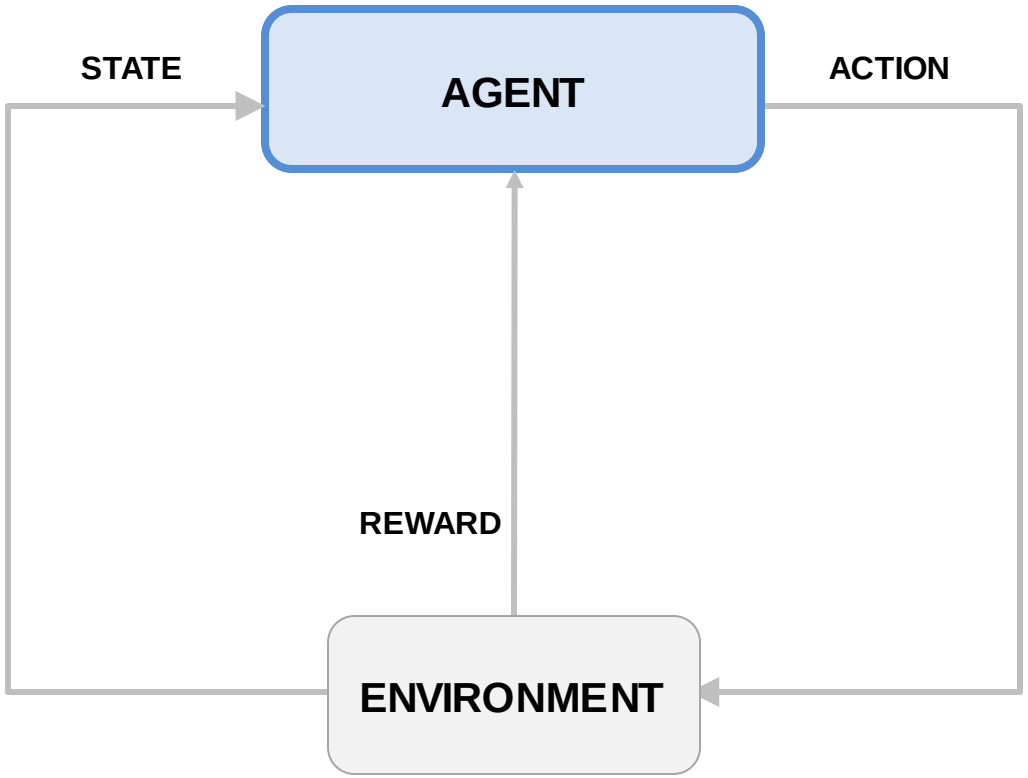

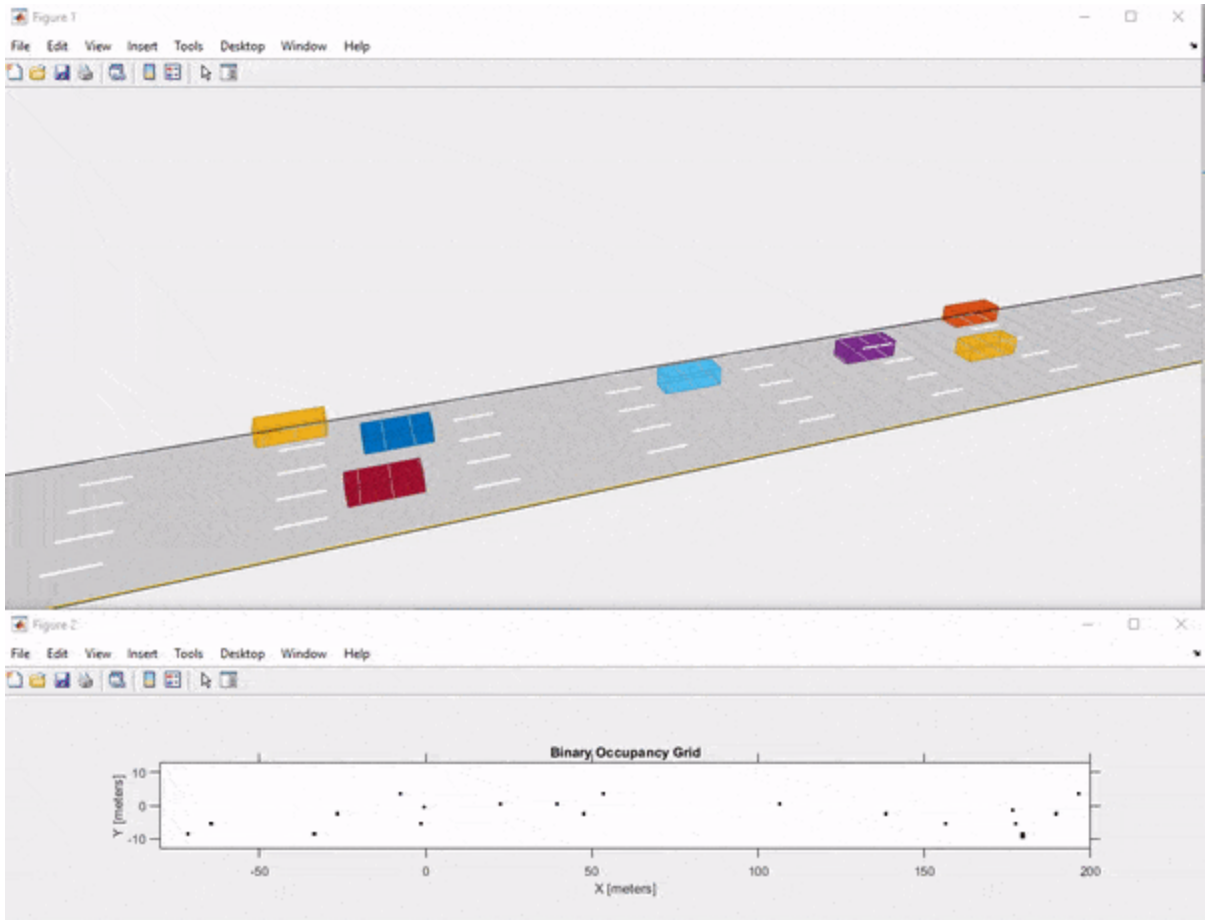

POP QUIZ

Results will be reported to your manager

What is the difference between Machine Learning and Deep Learning?

A. Deep learning is machine learning done really far underground.

B. I don't know, I didn't pay attention, I actually don't even work here, I just show up to these things.

C. Machine learning requires manual feature extraction while deep learning automatically extracts features making it end-to-end learning

What is the difference between Machine Learning and Deep Learning?

A. Deep learning is machine learning done really far underground.

B. I don't know, I didn't pay attention, I actually don't even work here, I just show up to these things.

C. Machine learning requires manual feature extraction while deep learning automatically extracts features making it end-to-end learning

Which of the following is not an application of deep learning?

A. Image classification
B. Speech recognition
C. Automated driving
D. Filtering applications like rain removal
E. Recognizing people's faces on your phone's photo app
F. Building a hotdog/not-hotdog classifier
G. None of the above

Which of the following is not an application of deep learning?

A. Image classification
B. Speech recognition
C. Automated driving
D. Filtering applications like rain removal
E. Recognizing people's faces on your phone's photo app
F. Building a hotdog/not-hotdog classifier
G. None of the above

Which of the following is NOT a layer in deep networks?

A. Fully Connected Layer
B. Softmax Layer
C. Classification Layer
D. Convolution Layer
E. ReLu Layer
F. MaxPooling Layer
G. Banana Layer (classifies all objects as Banana)

Which of the following is NOT a layer in deep networks?

A. Fully Connected Layer
B. Softmax Layer
C. Classification Layer
D. Convolution Layer
E. ReLu Layer
F. MaxPooling Layer
G. Banana Layer (classifies all objects as Banana)

What does the Fully Connected Layer do?

A. Calculates a score for each category
B. Ensures your layered sandwiches stay Fully Connected
C. Saves you 15% or more on car insurance

What does the Fully Connected Layer do?

A. Calculates a score for each category
B. Ensures your layered sandwiches stay Fully Connected
C. Saves you 15% or more on car insurance

How do we perform transfer learning?

A. Change every other layer of our network to a softmax layer
B. Transfer all data from the CPU to the GPU
C. Load in a pre-trained network, modify the last few layers, and train it on our data.

How do we perform transfer learning?

A. Change every other layer of our network to a softmax layer
B. Transfer all data from the CPU to the GPU
C. Load in a pre-trained network, modify the last few layers, and train it on our data.

What are three hyperparameters that have a major impact on training time and accuracy?

A. Network Architecture
B. Mini Batch Size
C. Learning Rate
D. Flux Capacitor

What are three hyperparameters that have a major impact on training time and accuracy?

A. Network Architecture
B. Mini Batch Size
C. Learning Rate
D. Flux Capacitor

What is loss?

A. The opposite of a win

B. The state or feeling of grief when deprived of someone or something of value

C. A measurement of error between predicted labels and actual labels. Loss has an inverse relationship with score, and our goal is to minimize loss.

D. All of the above

What is loss?

A. The opposite of a win

B. The state or feeling of grief when deprived of someone or something of value

C. A measurement of error between predicted labels and actual labels. Loss has an inverse relationship with score, and our goal is to minimize loss.

D. All of the above

Which of the following statements is false?

A. MATLAB makes it easy to import pre-trained models through add-ons and model importers
B. MATLAB supports the entire deep learning workflow including labeling, training, and deployment
C. MATLAB has visual training plots that allow you to see accuracy and loss during training
D. We do a great job of **subtly** marketing MATLAB's deep learning capabilities

Which of the following statements is false?

A. MATLAB makes it easy to import pre-trained models through add-ons and model importers

B. MATLAB supports the entire deep learning workflow including labeling, training, and deployment

C. MATLAB has visual training plots that allow you to see accuracy and loss during training

D. We do a great job of **subtly** marketing MATLAB's deep learning capabilities

Free Seminar: ADAS and Automated Driving Development Using MATLAB and Simulink

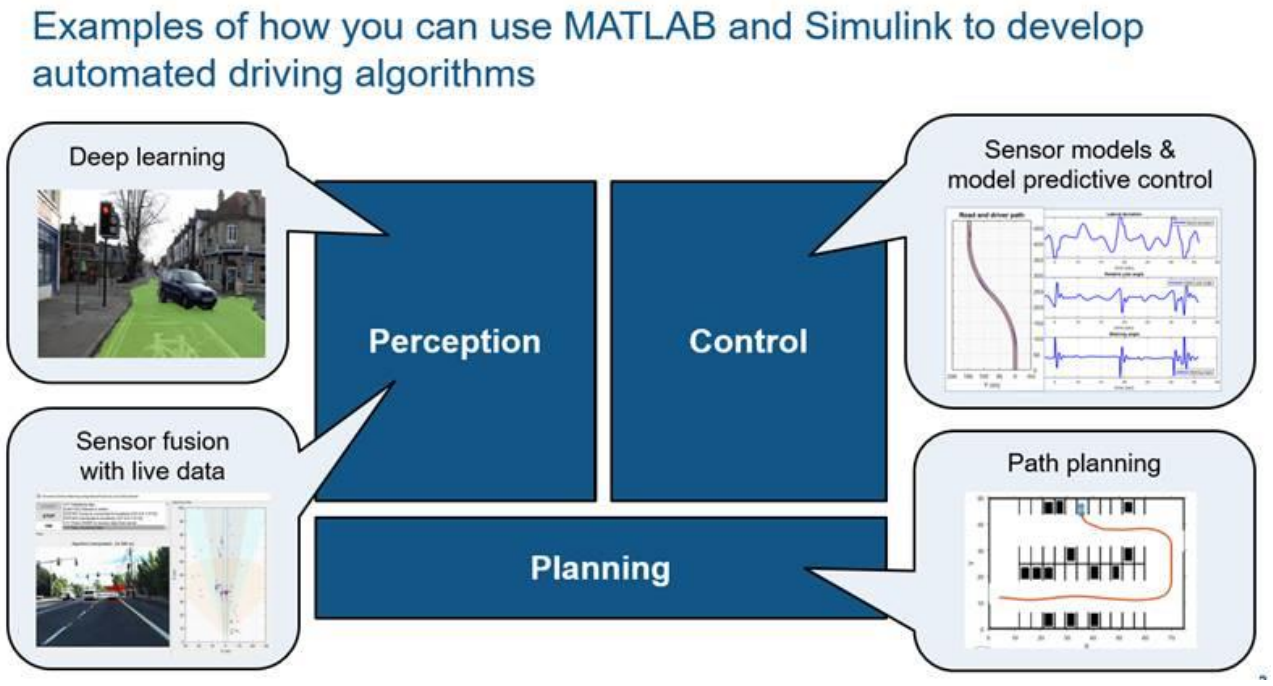

Location	Venue	Start Date	End Date
Santa Clara, CA	MathWorks Office (Mission Towers, Floor 1)	29 Mar 2018 - 9:00 AM	29 Mar 2018 - 12:00 PM

Overview

Do you have an **open and flexible visualization** tool to gain insight from vision, radar, and LiDAR data?
How quickly can you apply the latest **deep learning** research to vision perception development?
Are you able to design **sensor fusion and control** in simulation, before going to the test vehicle?

In this seminar, MathWorks engineers will demonstrate several new technologies to accelerate ADAS and automated driving development with MATLAB and Simulink.

Highlights

They will introduce the latest ADAS and automated driving development tools from MathWorks, including

- Visualize recorded and live sensor data
- Framework for sensor fusion algorithm design and test
- Deep learning for LiDAR and camera processing
- Control design in simulation

They will demonstrate new products and ADAS-extensions of existing MathWorks products, including

- Automated Driving System Toolbox
- Model Predictive Control Toolbox
- Vehicle Network Toolbox
- Robotics Systems Toolbox
- GPU Coder

Who Should Attend

Engineers and managers working on ADAS and automated driving system, algorithm, and software development.

Questions?

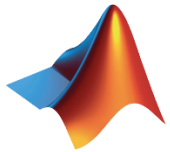

MathWorks® can help you do Deep Learning

Free resources

- **Guided evaluations with a MathWorks deep learning engineer**
- Proof-of-concept projects
- **Deep learning hands-on workshop**
- Seminars and technical deep dives
- Deep learning onramp course

More options

- Consulting services
- Training courses
- Technical support
- Advanced customer support
- Installation, enterprise, and cloud deployment
- Deep Learning Paid Training

Automated Object Detection

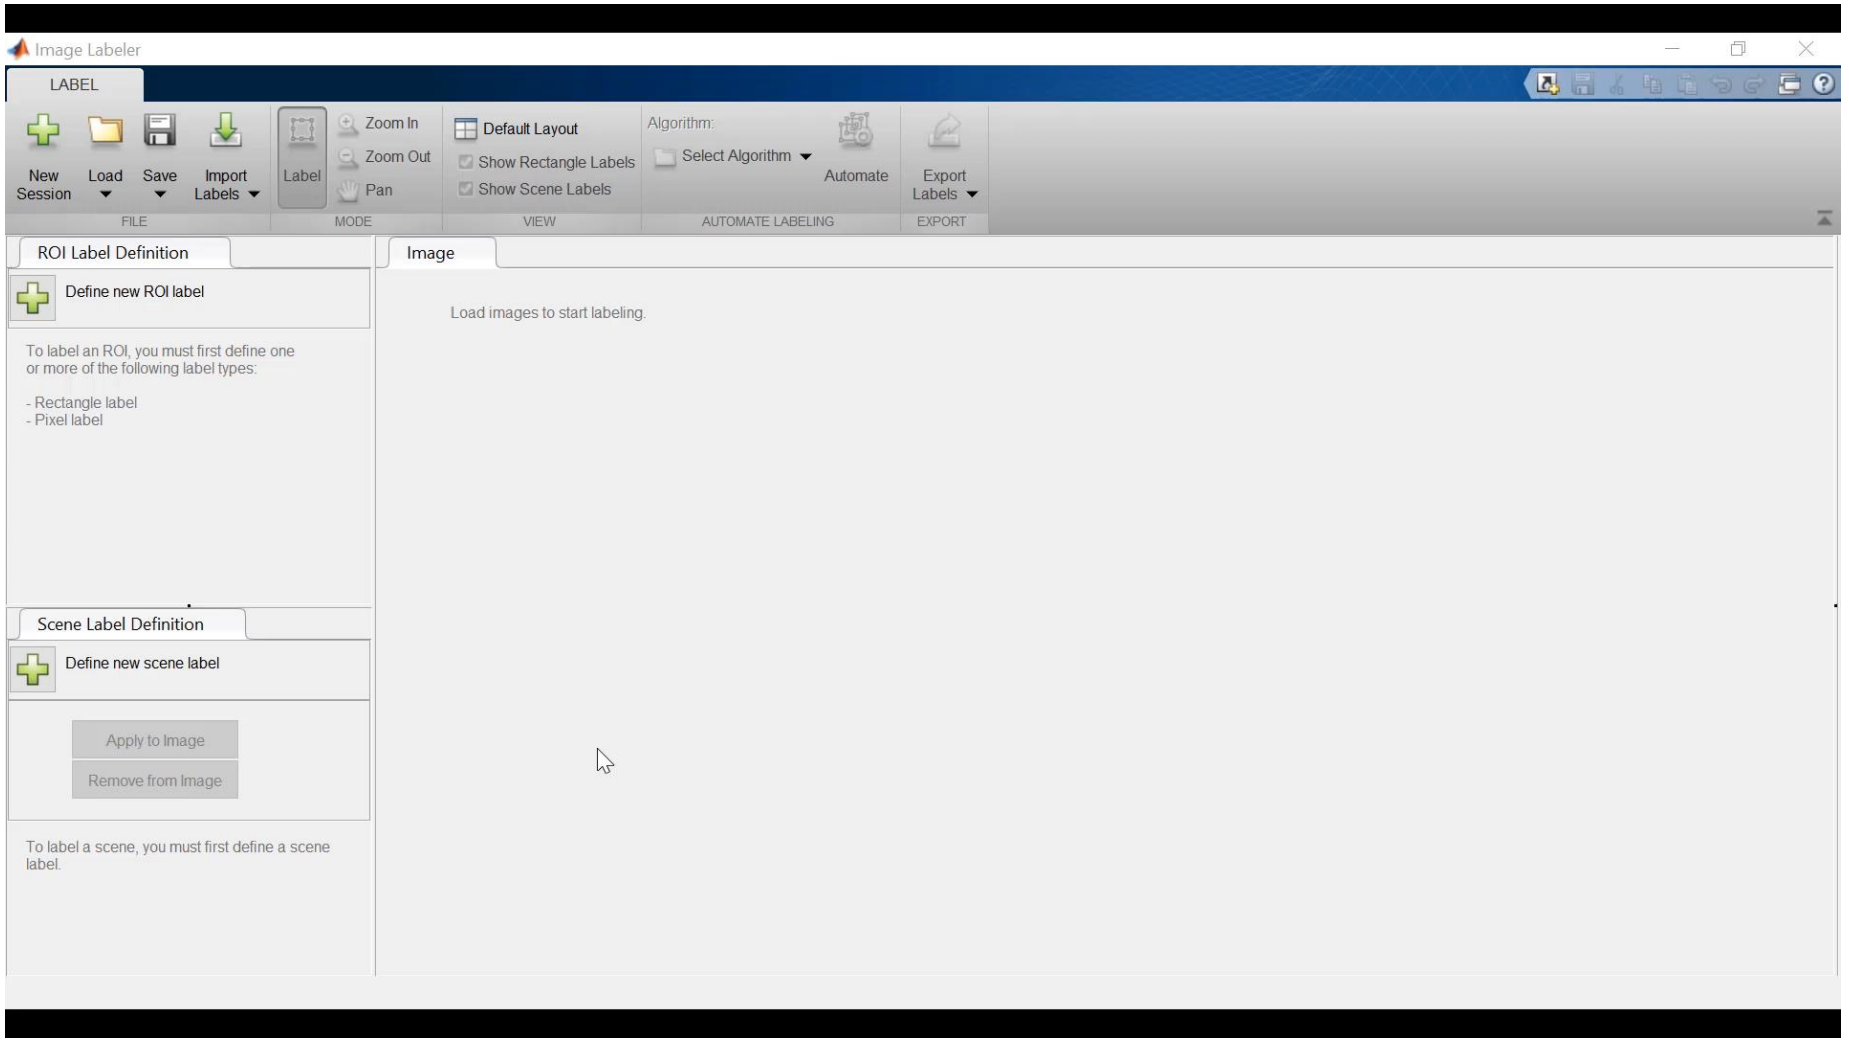

Visualizations for Understanding Network Behavior

Filters

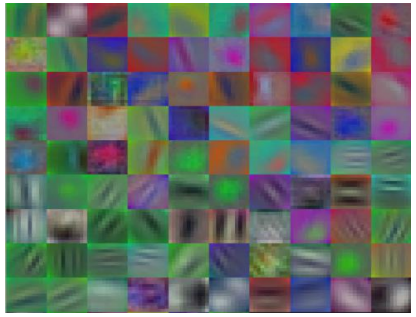

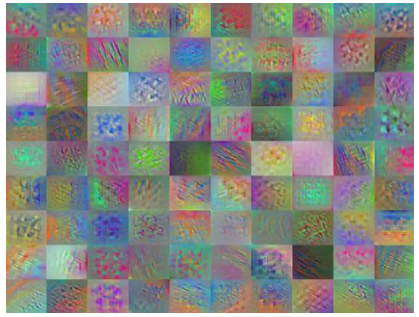

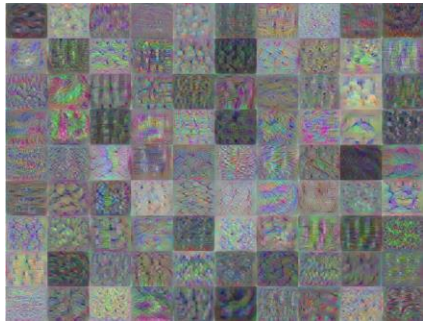

• • •

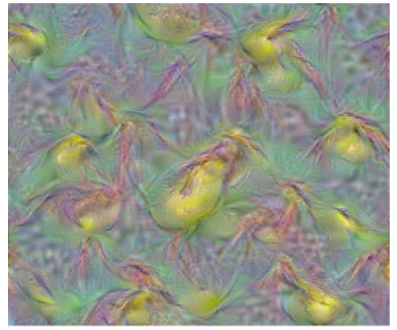

Deep Dream

Activations

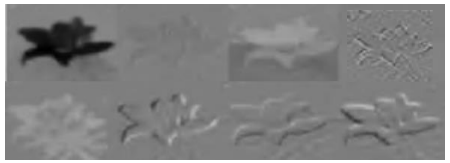

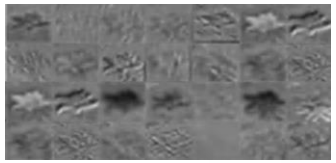

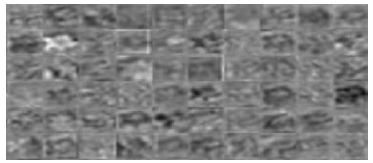

- Custom visualizations
  - Example: Class Activation Maps

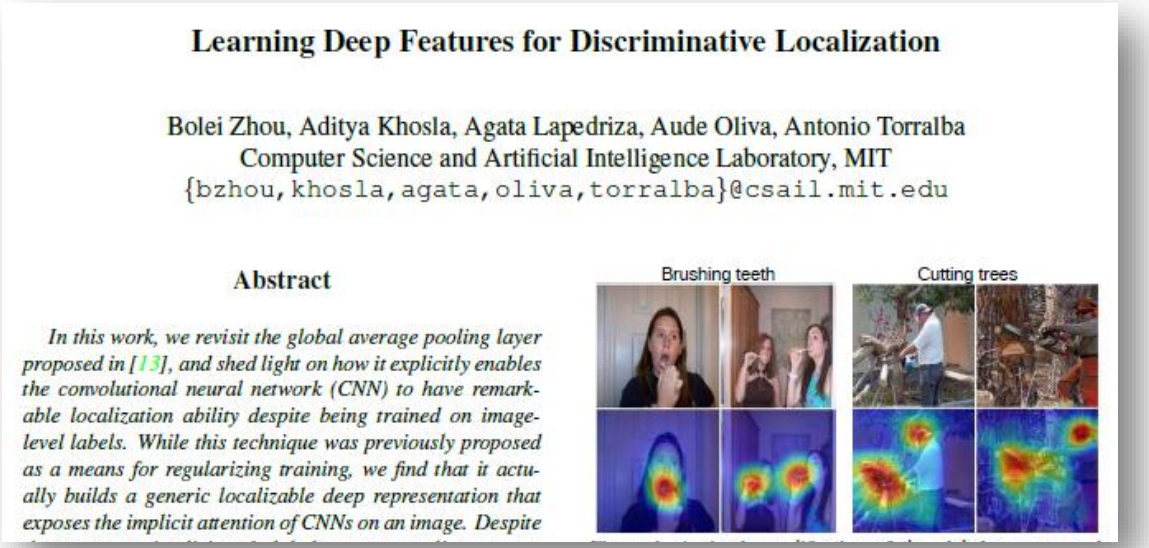

Learning Deep Features for Discriminative Localization

Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, Antonio Torralba
Computer Science and Artificial Intelligence Laboratory, MIT
{bzhou, khosla, agata, oliva, torralba}@csail.mit.edu

Abstract

In this work, we revisit the global average pooling layer proposed in [13], and shed light on how it explicitly enables the convolutional neural network (CNN) to have remarkable localization ability despite being trained on image-level labels. While this technique was previously proposed as a means for regularizing training, we find that it actually builds a generic localizable deep representation that exposes the implicit attention of CNNs on an image. Despite

http://cnnlocalization.csail.mit.edu/Zhou_Learning_Deep_Features_CVPR_2016_paper.pdf

A Practical Example of Reinforcement Learning

Training a Self-Driving Car

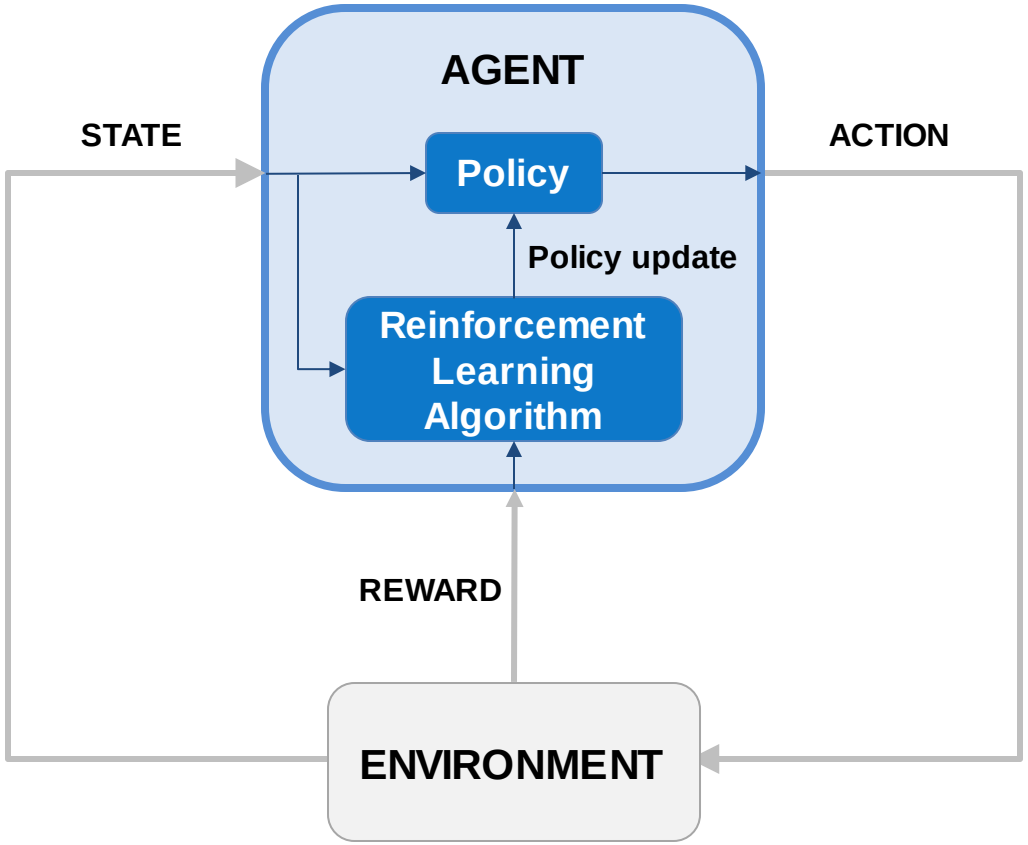

- Vehicle's computer learns how to drive… (**agent**)
- using sensor readings from LIDAR, cameras,… (**state**)
- that represent road conditions, vehicle position,… (**environment**)
- by generating steering, braking, throttle commands,… (**action**)
- based on an internal state-to-action mapping… (**policy**)
- that tries to optimize driver comfort & fuel efficiency… (**reward**).

- The policy is updated through repeated trial-and-error by a **reinforcement learning algorithm**

What is Reinforcement Learning?

- What is Reinforcement Learning?
  - Type of machine learning that trains an **'agent'** through repeated interactions with an environment
- How does it work?
  - Through a trial & error process that maximizes success

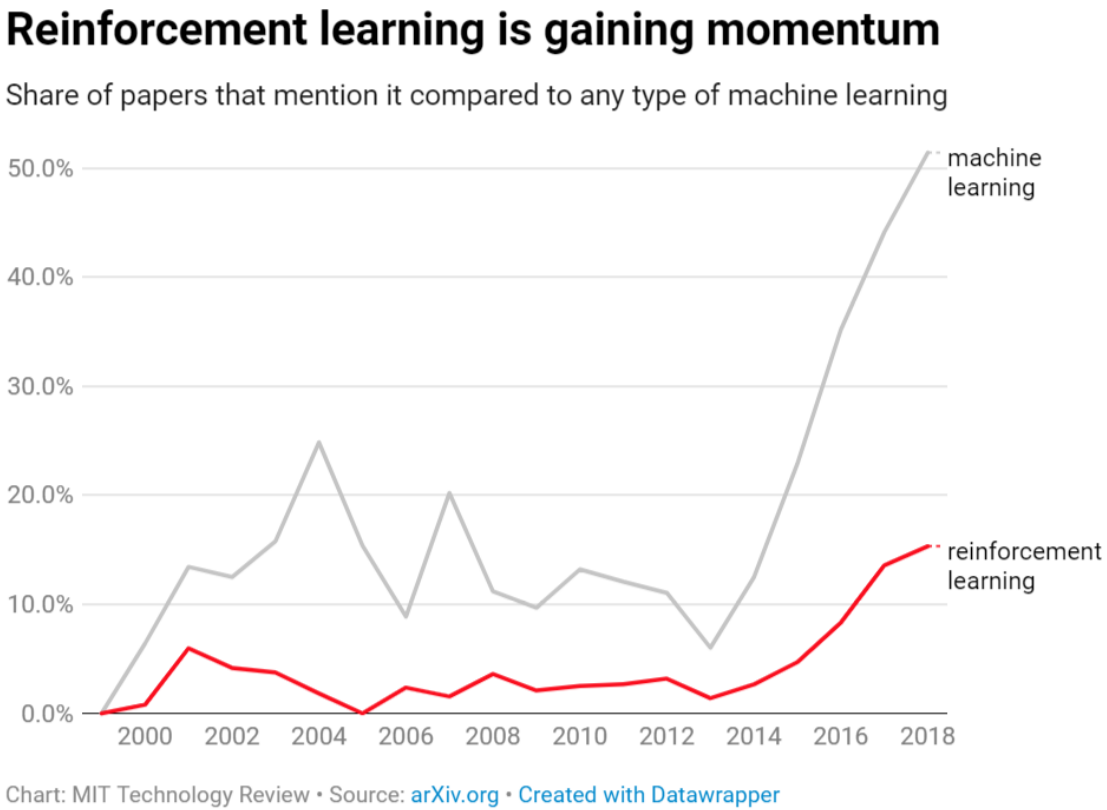

Reinforcement Learning Applications

- Why should you care about Reinforcement Learning?
  - It enables the use of deep learning for controls and decision-making applications

Controls

Game Play

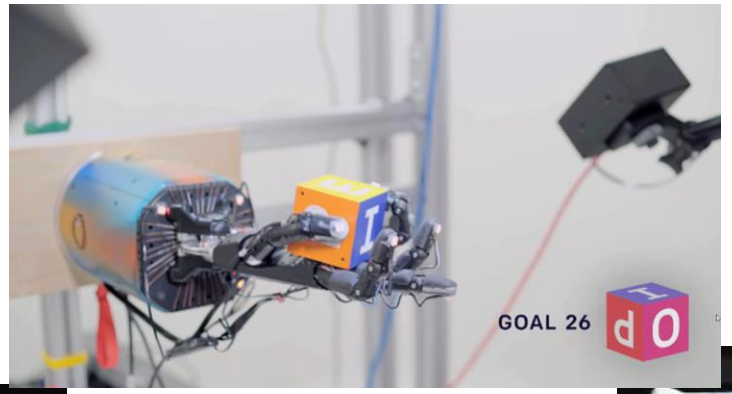

Robotics

Autonomous driving

Reinforcement Learning vs Machine Learning vs Deep Learning

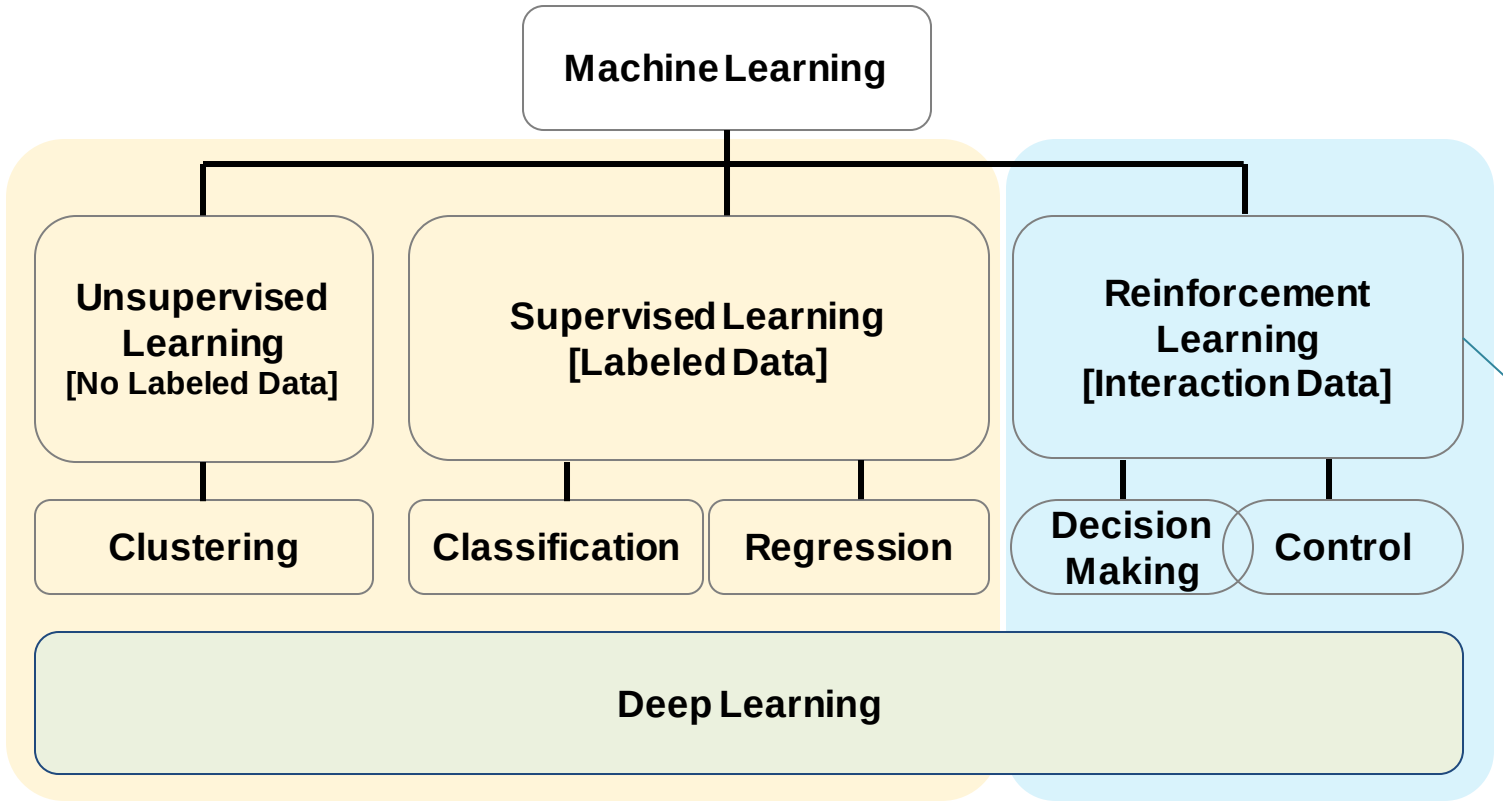

Reinforcement learning:

- Learning through trial & error [*interaction*]
- Complex problems typically need deep learning [*Deep Reinforcement Learning*]
- It's about learning a **behavior** or accomplishing a **task**

Reinforcement Learning vs Machine Learning vs Deep Learning

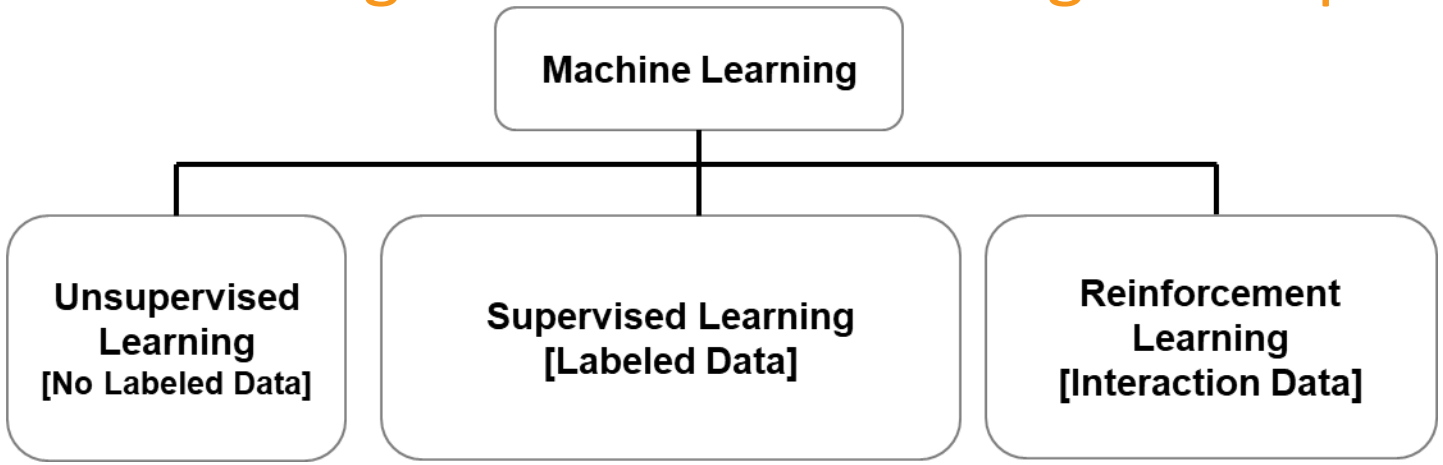

Reinforcement Learning vs Machine Learning vs Deep Learning

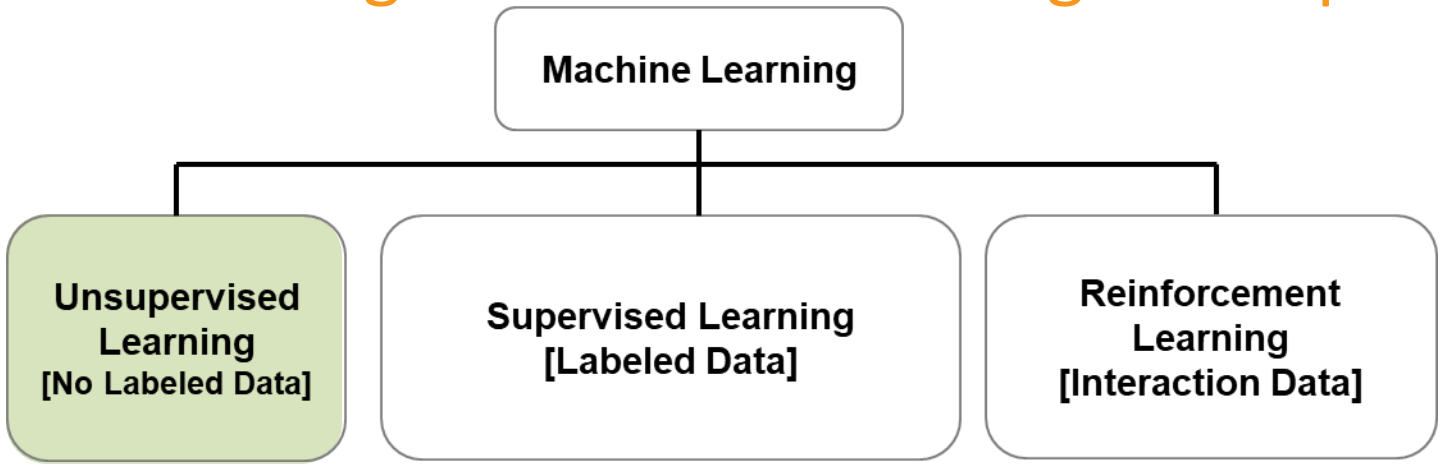

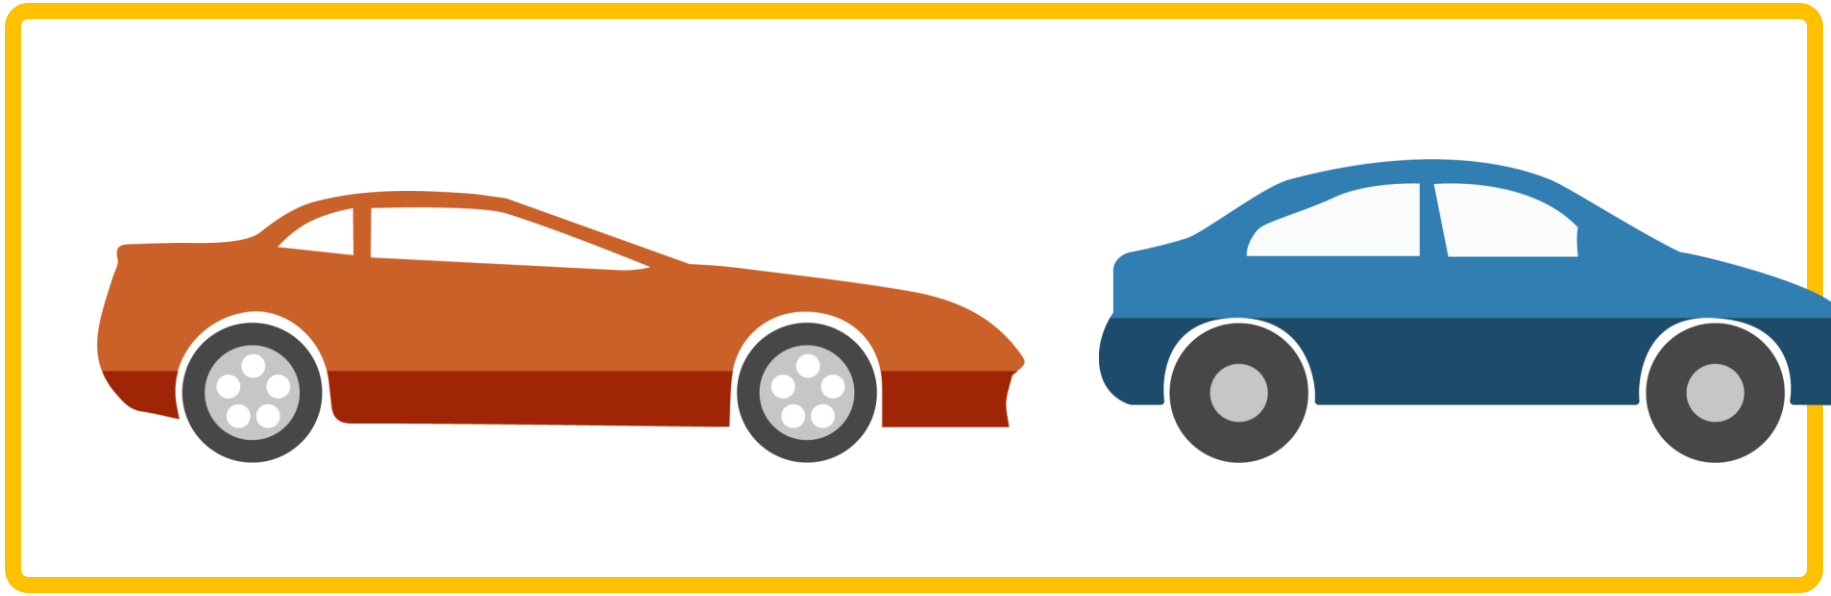

Reinforcement Learning vs Machine Learning vs Deep Learning

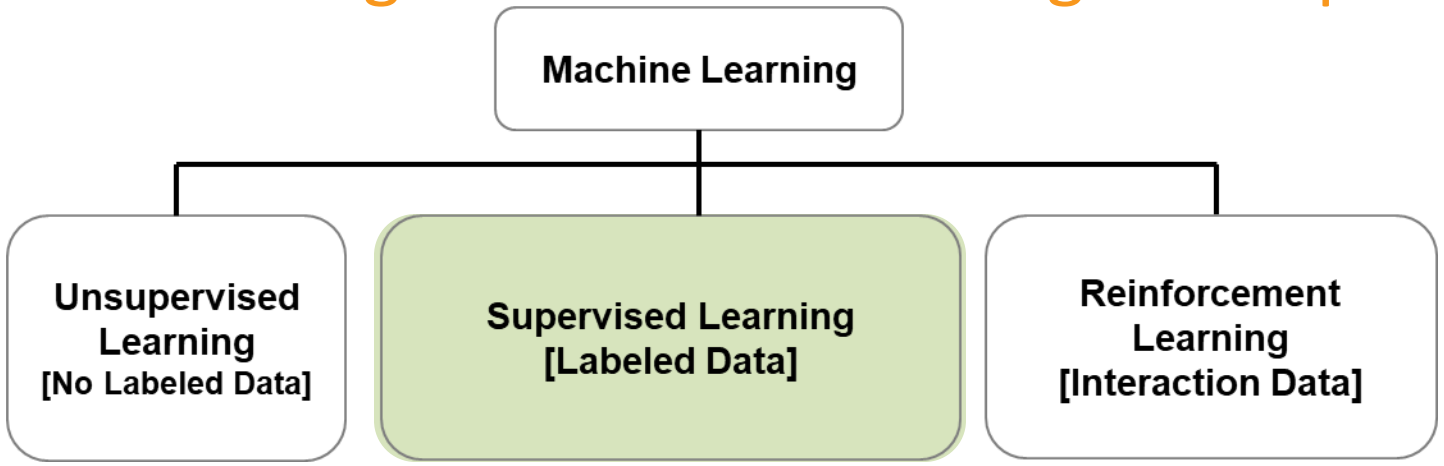

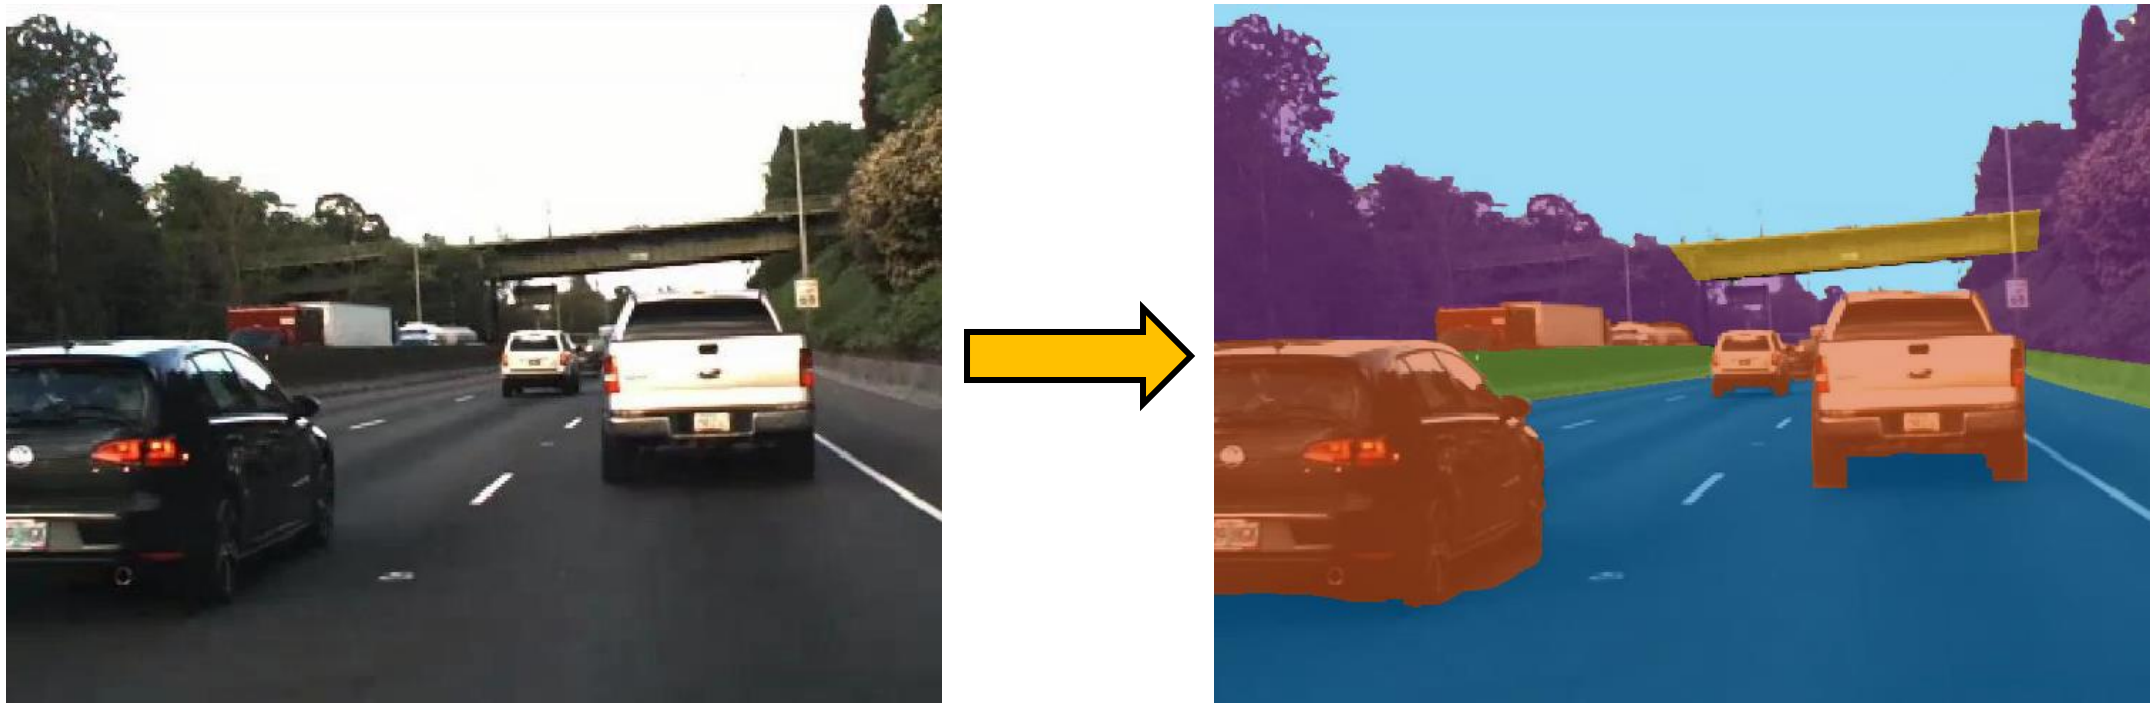

Reinforcement Learning vs Machine Learning vs Deep Learning

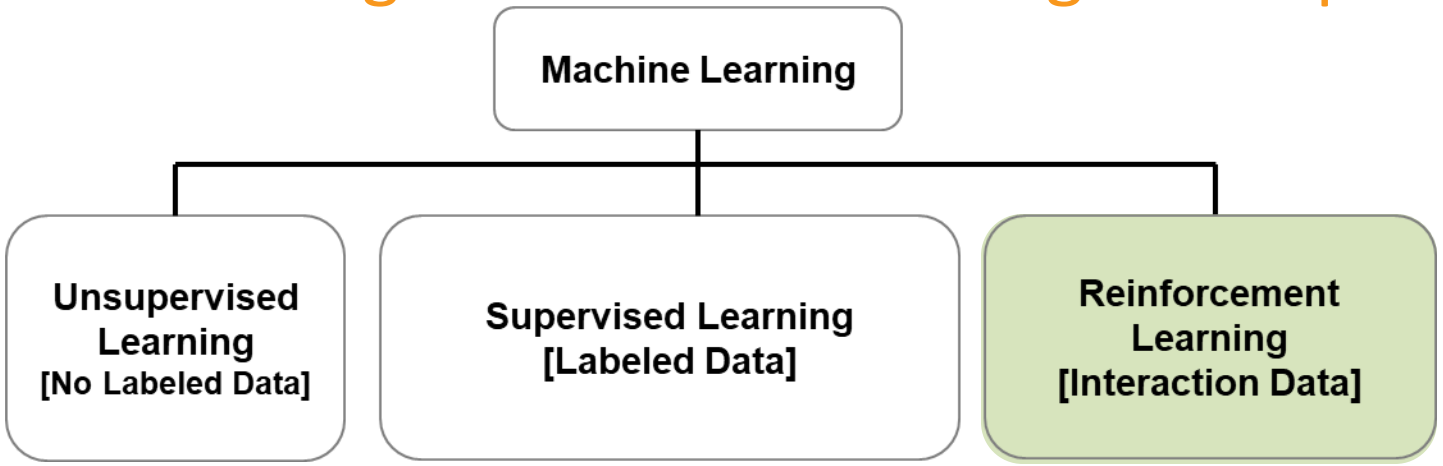

Resources for Learning More

- Documentation
  - https://www.mathworks.com/help/reinforcement-learning/
- Examples
  - https://www.mathworks.com/help/reinforcement-learning/examples.html
- Product Page
  - https://www.mathworks.com/products/reinforcement-learning.html
- Tech Talks
  - https://www-integ1.mathworks.com/videos/tech-talks/controls.html

Dynamic Solutions. Precise Results.