


Hands-on Virtual Lab: Deep Learning



Reece Teramoto
Application Engineer

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

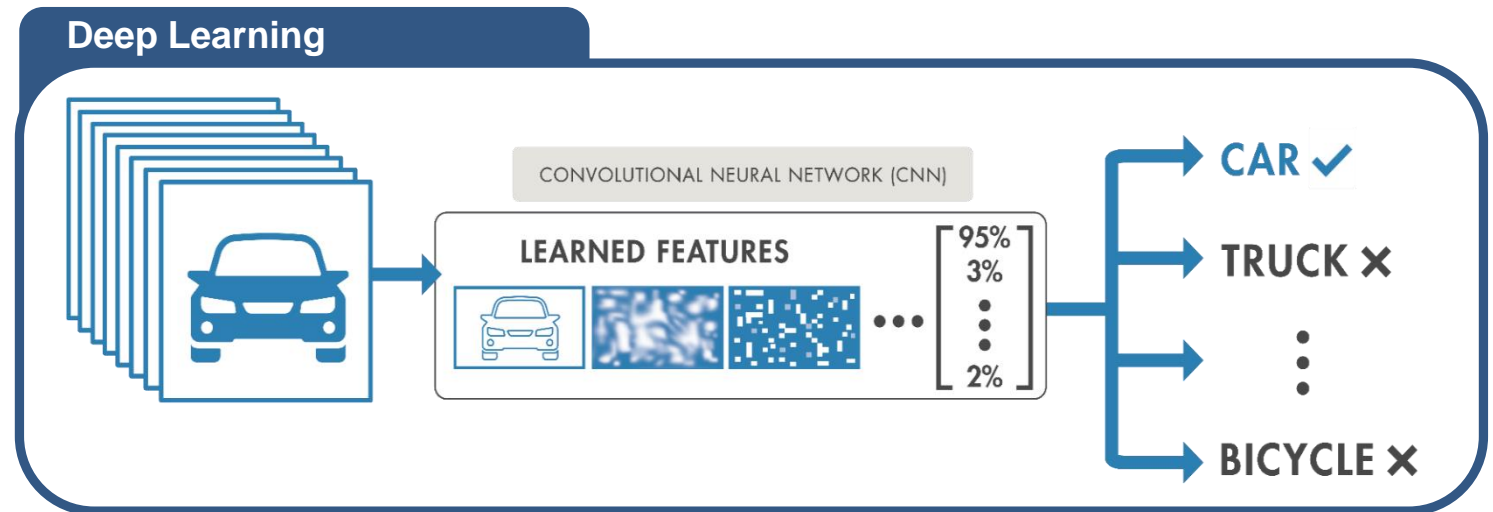



Optional: Deploying Deep Networks– OR – Improving Network Accuracy

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



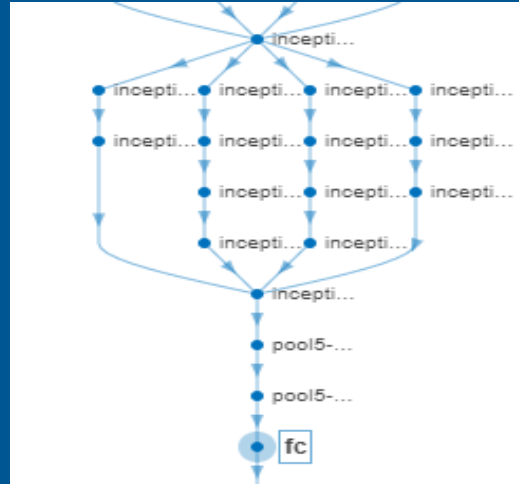
Deep Learning Workflow

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

- \$10m's spent on land seismic acquisition each year.
- Terrain type very important to daily shot target.
- Currently manually drawn polygons on satellite/drone images + direct site visits - weeks.
- We replace whole workflow with DL semantic segmentation approach (segnet).

Copyright of Shell Global Solutions (UK)

Radar image with rough polygons overlaid

Terrain Recognition with Hyperspectral Data

CNNs for Digital Pathology Analysis

Tumor ROI in yellow.

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

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




Deep Learning Networks Take in Numeric Data



199	206	208	201	188	178	165	164	180
202	205	202	188	176	169	178	186	183
203	206	189	178	181	183	182	154	87
203	192	184	186	177	167	153	181	192
191	182	176	166	153	141	136	180	227
166	165	154	154	138	137	169	170	211
158	150	145	183	144	156	158	154	179
143	51	98	144	129	130	143	178	123
107	50	33	95	152	173	192	159	87
104	100	84	120	132	172	131	64	94
119	101	97	81	90	109	87	106	111
127	122	110	97	108	120	133	131	134
111	117	108	119	131	143	146	141	156
126	122	113	119	139	142	155	161	151
129	126	130	111	103	130	149	149	156
138	128	136	144	136	129	134	122	145
154	133	134	141	168	150	126	127	151

Images are a numeric matrix



Signals are numeric vectors

The Bird Flies = [0 13 5 6]
 The Leaf Is Brown = [13 3 11 2]

Text is processed as numeric vectors

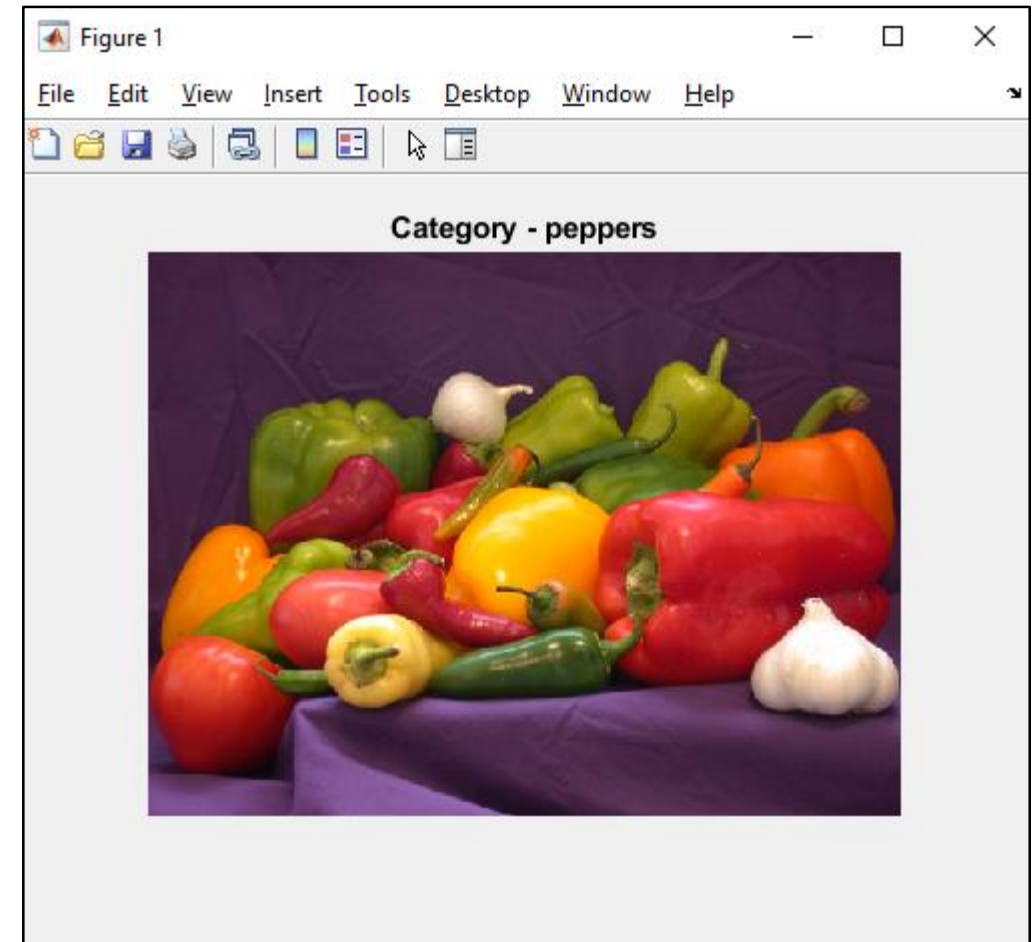
Exercise 1 – Deep Learning in 6 Lines of Code

Purpose:

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

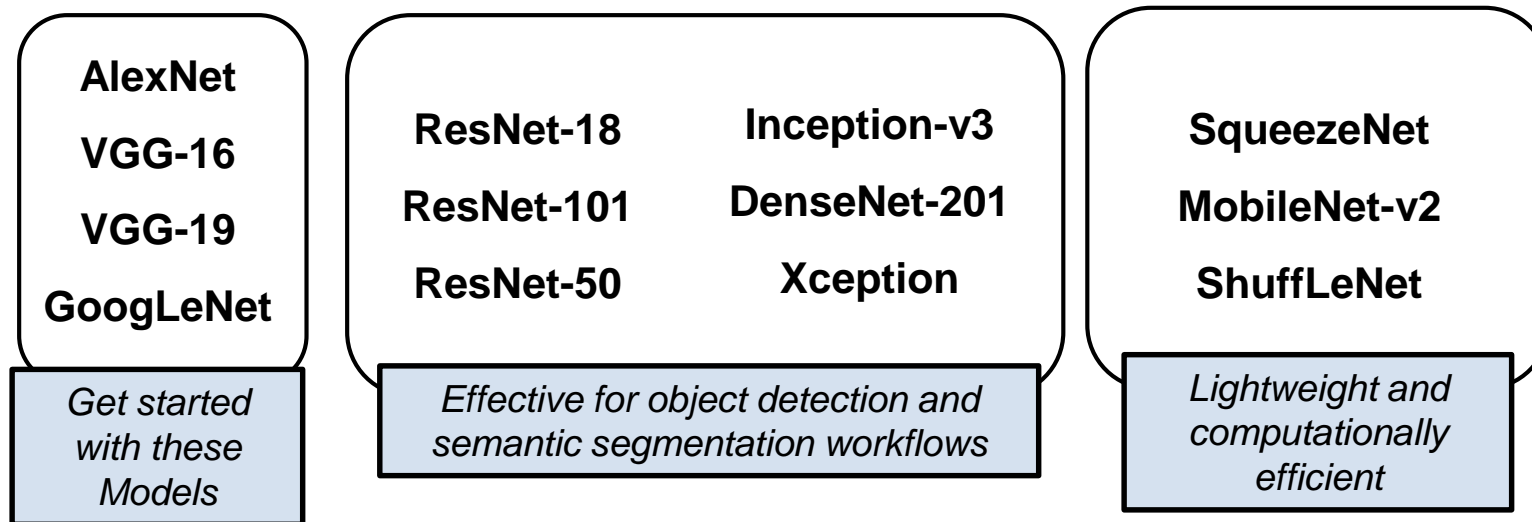
To Do:

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 [HERE](#)

Access Pretrained Models from Within MATLAB or Import from the Web



Exercise 2 – Pretrained Models

Purpose:

- Classify Images using pretrained models.
- See how different network architectures affect results.
- Use datastores to access data efficiently

To Do:

- Open `work_pretrainednetworks.mlx`.

Pick a network to load in.

net = googlenet



Inspect

```
% Insp
% analy
net.La
% Insp
% Look
class
```

- alexnet
- googlenet
- vgg16
- resnet50
- xception
- inceptionv3
- squeezenet
- nasnetlarge

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






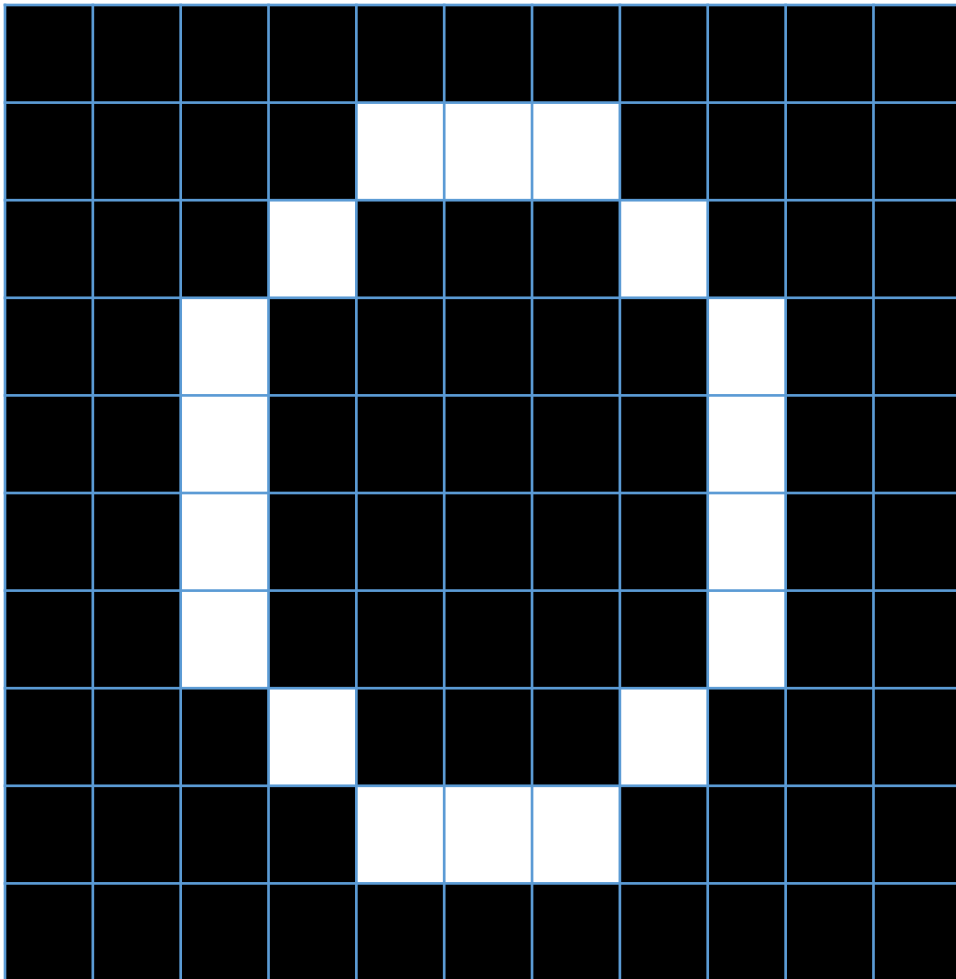
Creating Layer Architectures

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




- Convolution Layer
- ReLU Layer
- Max Pooling 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

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



Rectified Linear Units Layer (ReLU)

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


Max Pooling is a down-sampling operation

Shrink large images while preserving important information



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

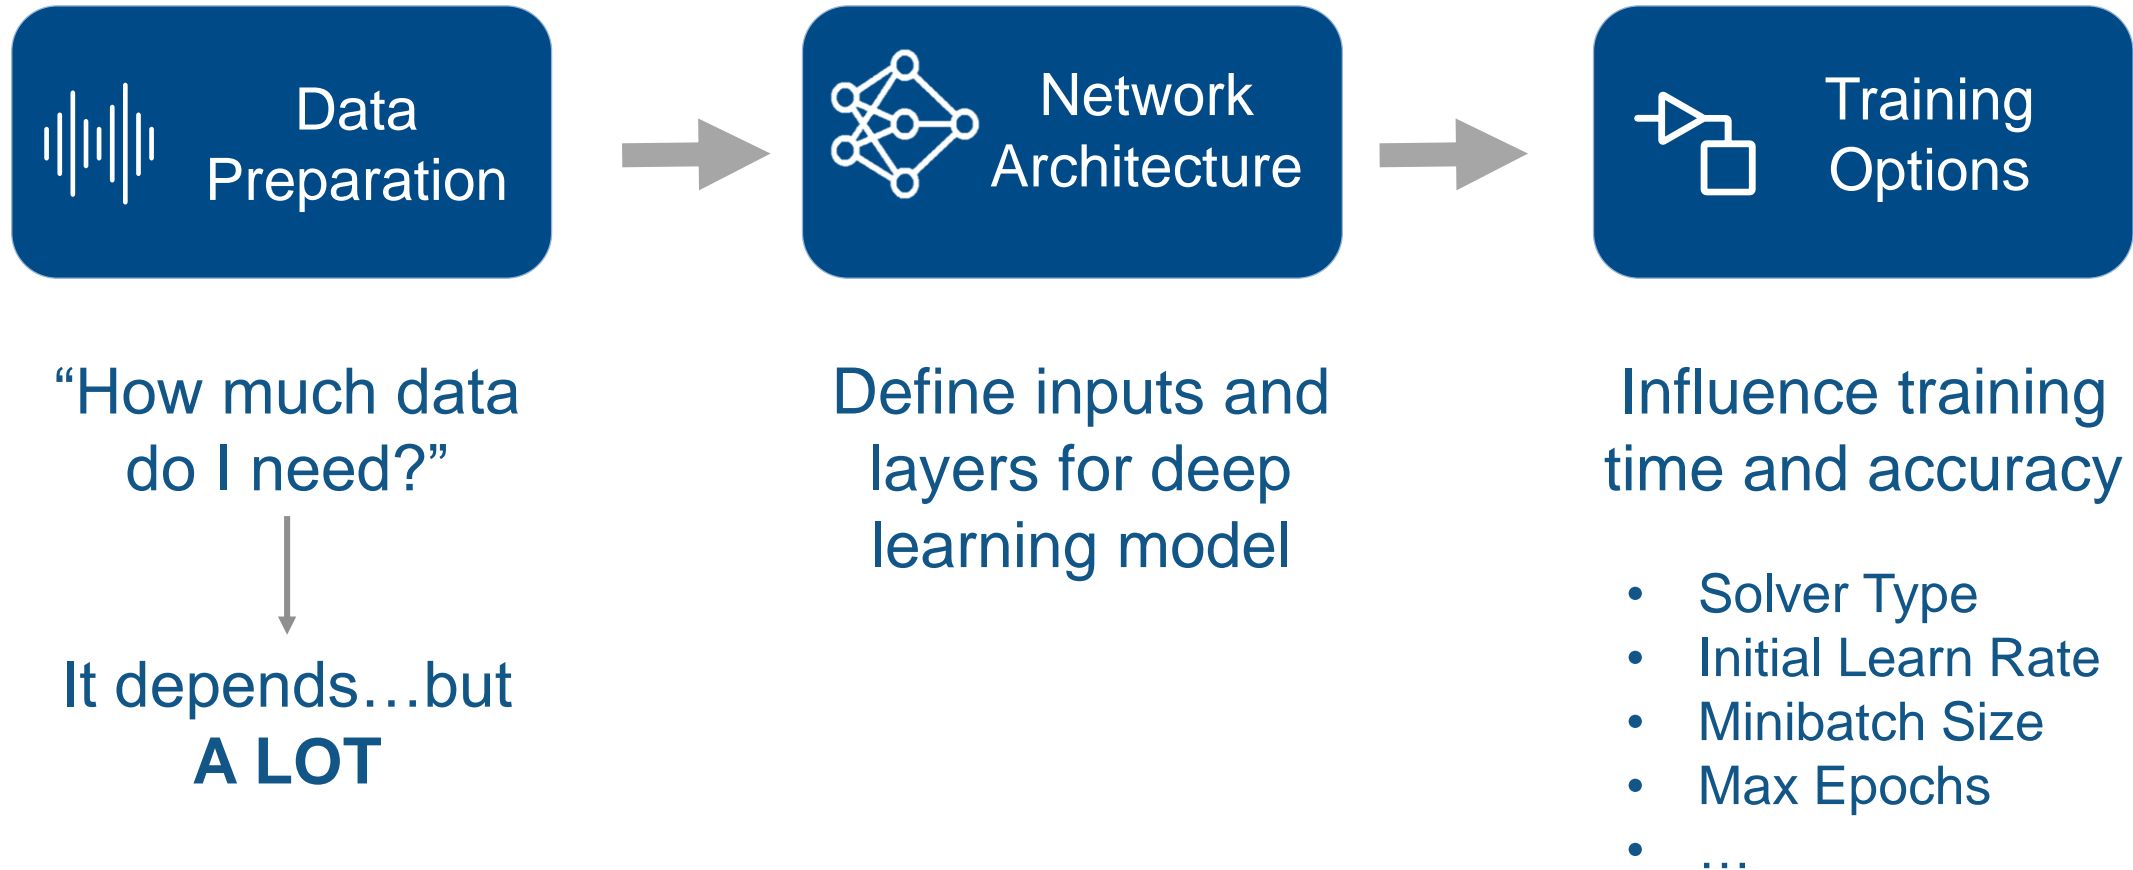
- LSTM
- BiLSTM
- Word Embedding (...)

Normalization

- Dropout
- Batch normalization
- (...)

Research papers and [doc examples](#) can provide guidelines for creating architecture.

3 Components to Train any Network




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



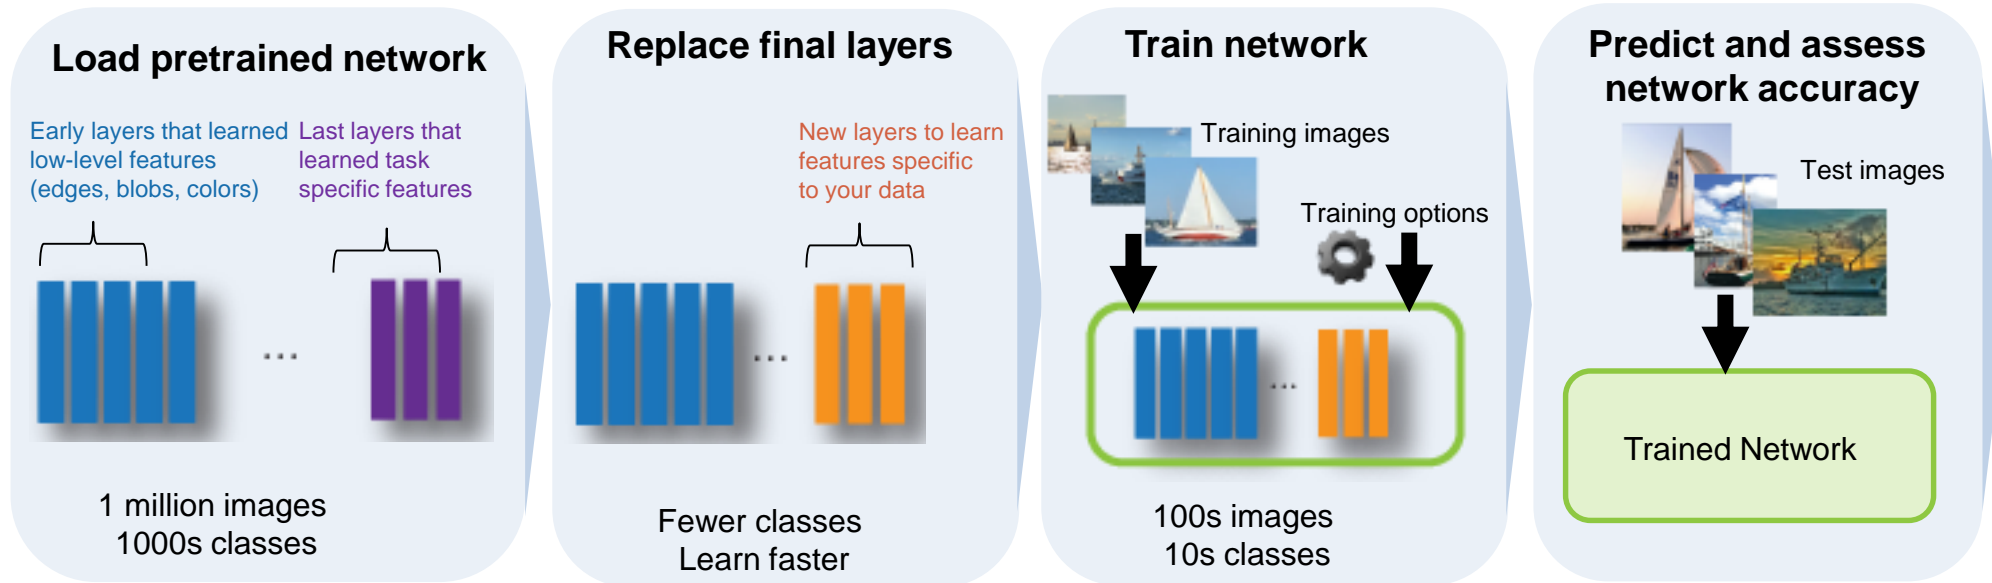
Sources: <http://yann.lecun.com/exdb/mnist/>
<https://rodrigob.github.io/are-we-there-yet/build/classification-datasets-results>

Experiment Manager – Run, Track, and Analyze Multiple Deep Learning Experiments

The screenshot displays the Experiment Manager interface. The top toolbar includes buttons for New, Save, Duplicate, Layout, Run, Stop, Training Plot, Confusion Matrix, Filter, and Export. The left sidebar shows the Experiment Browser with a tree view of experiments like 'Baseline Establishment' and 'Baseline Tuning'. The main area shows 'Result Details' for 'Baseline Tuning' with a progress bar and a summary table. Below this is a detailed table of 16 trials.

Trial	Status	Progress	Elapsed Time	myInitialLearn...	convFilterSize	Training Accu...	Training Loss	Validation Ac..
1	Complete	100.0%	0 hr 0 min 16 sec	1.0000e-6	3.0000	12.5000	2.6441	10.
2	Complete	100.0%	0 hr 0 min 15 sec	1.0000e-5	3.0000	25.7813	2.1228	20.
3	Complete	100.0%	0 hr 0 min 14 sec	0.0001	3.0000	64.8438	1.0878	42.
4	Complete	100.0%	0 hr 0 min 16 sec	0.0005	3.0000	90.6250	0.4648	49.
5	Complete	100.0%	0 hr 0 min 15 sec	1.0000e-6	4.0000	11.7188	2.4967	6.
6	Complete	100.0%	0 hr 0 min 15 sec	1.0000e-5	4.0000	23.4375	2.1213	14.
7	Complete	100.0%	0 hr 0 min 17 sec	0.0001	4.0000	72.6563	1.0283	39.
8	Running	30.7%	0 hr 0 min 4 sec	0.0005	4.0000			
9	Queued	0.0%		1.0000e-6	5.0000			
10	Queued	0.0%		1.0000e-5	5.0000			
11	Queued	0.0%		0.0001	5.0000			
12	Queued	0.0%		0.0005	5.0000			
13	Queued	0.0%		1.0000e-6	6.0000			
14	Queued	0.0%		1.0000e-5	6.0000			
15	Queued	0.0%		0.0001	6.0000			
16	Queued	0.0%		0.0005	6.0000			

Transfer Learning Workflow




Transfer Learning Workflow – Step 1

Load pretrained network

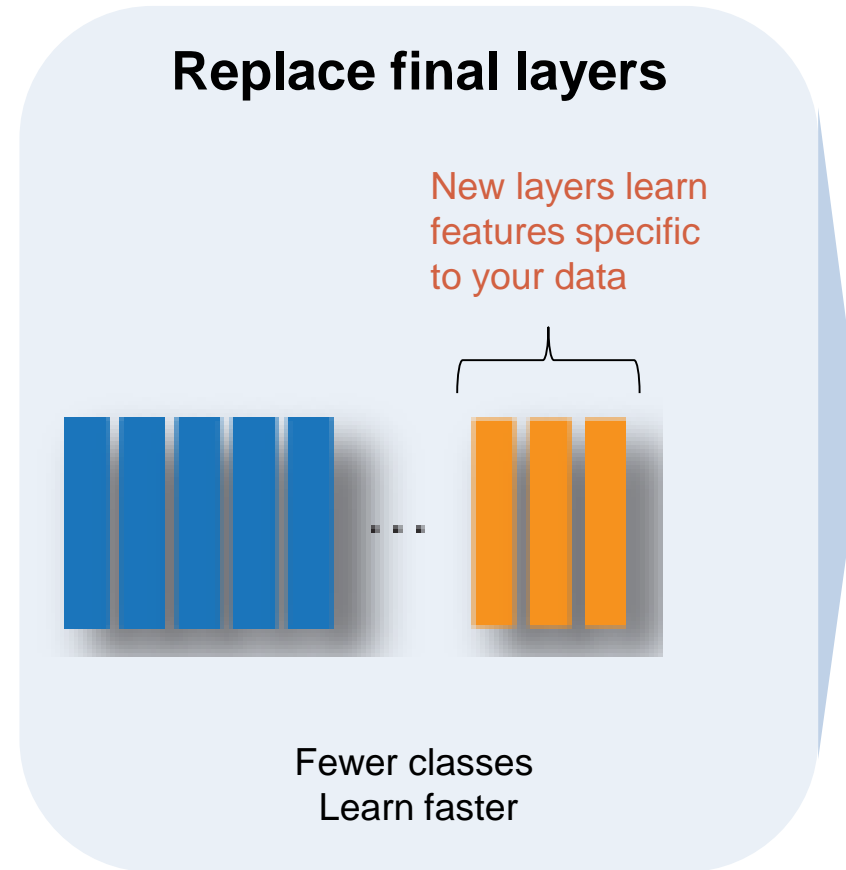
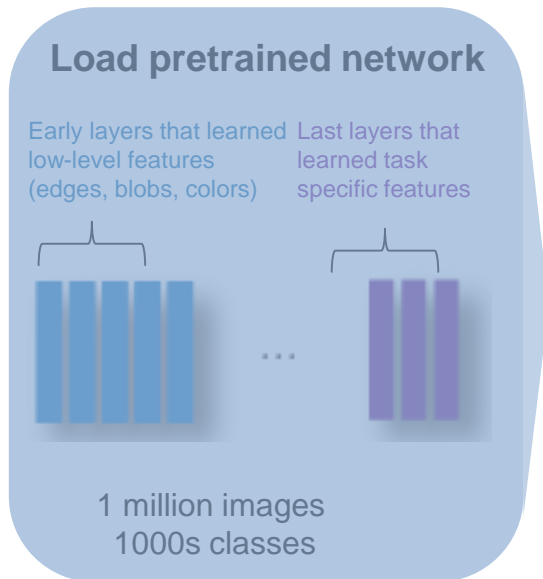
Early layers learn low-level features (edges, blobs, colors)

Last layers learn task-specific features

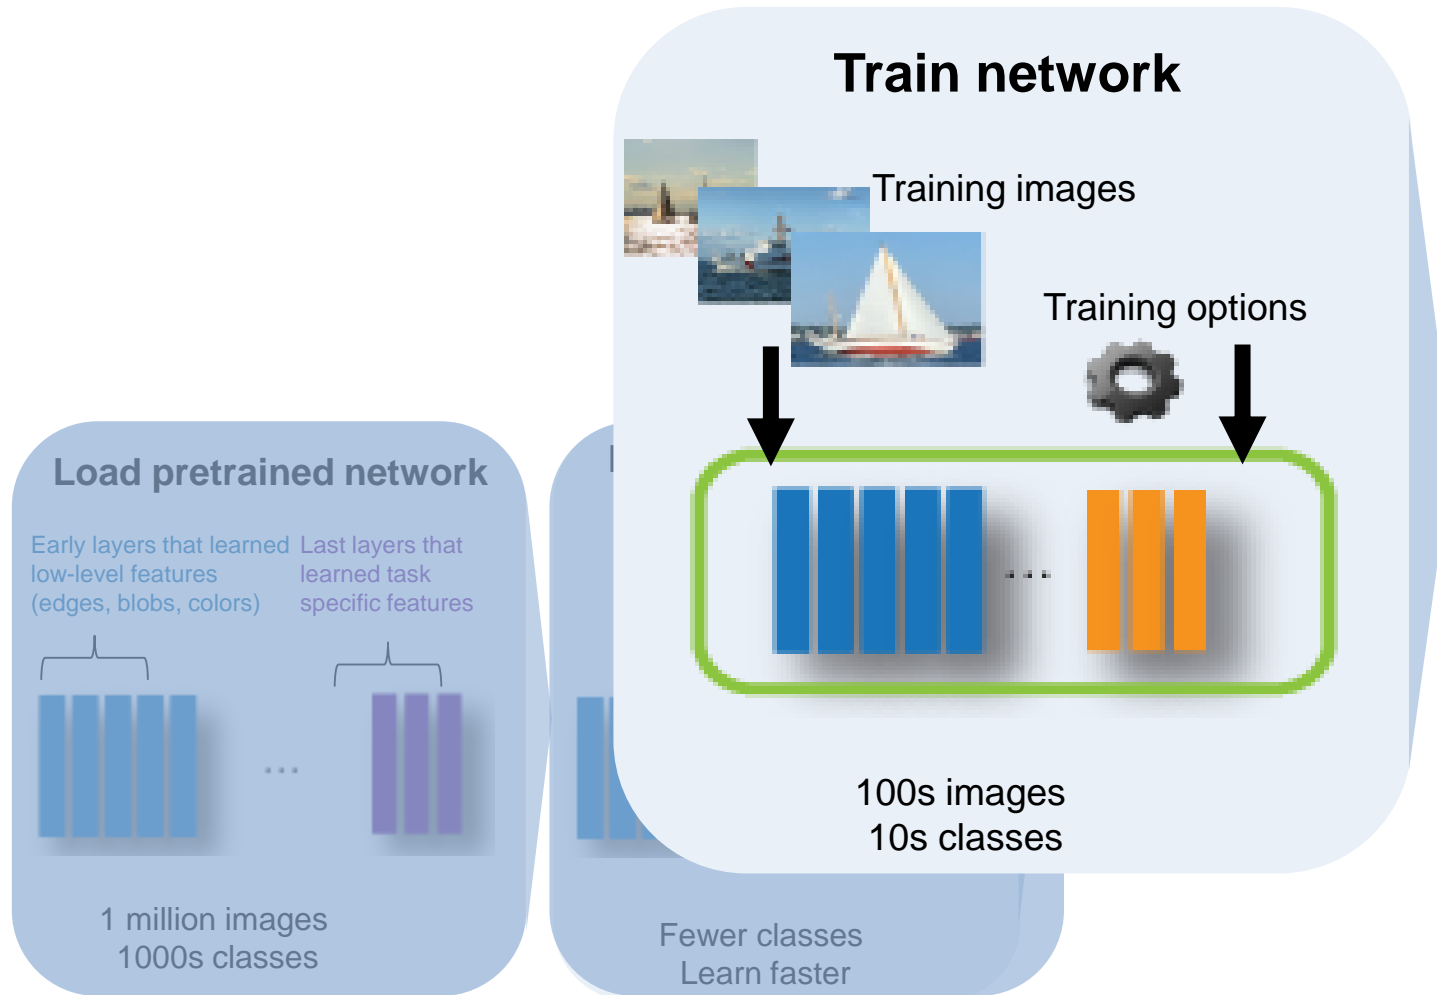


1 million images
1000s classes

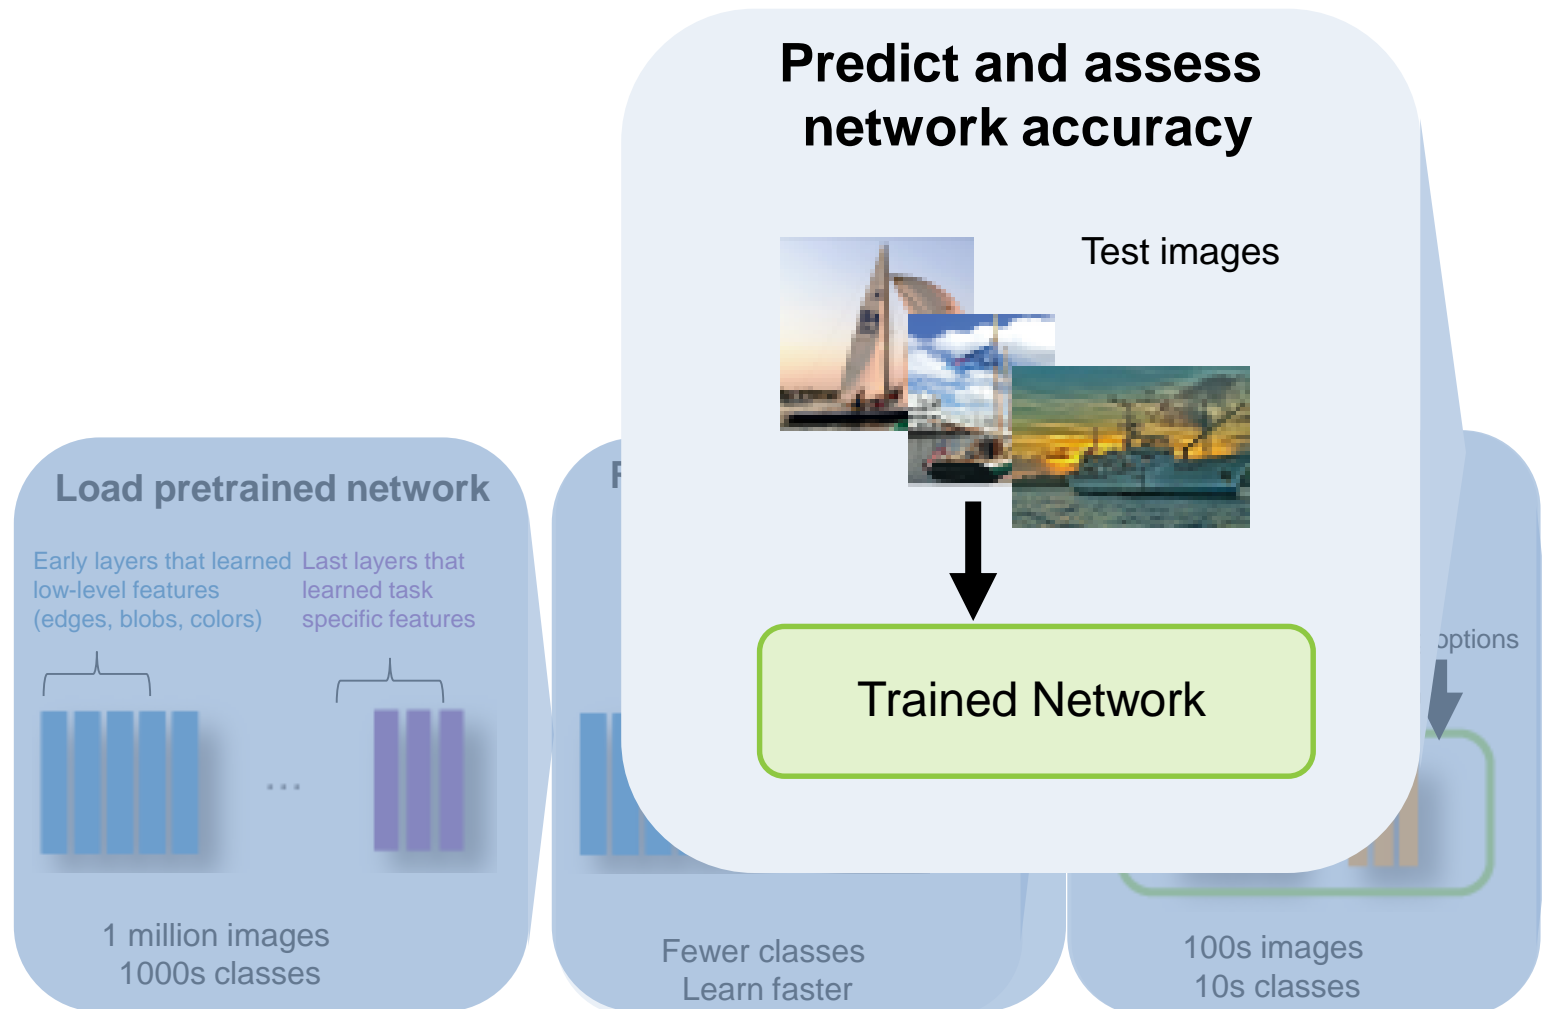
Transfer Learning Workflow – Step 2



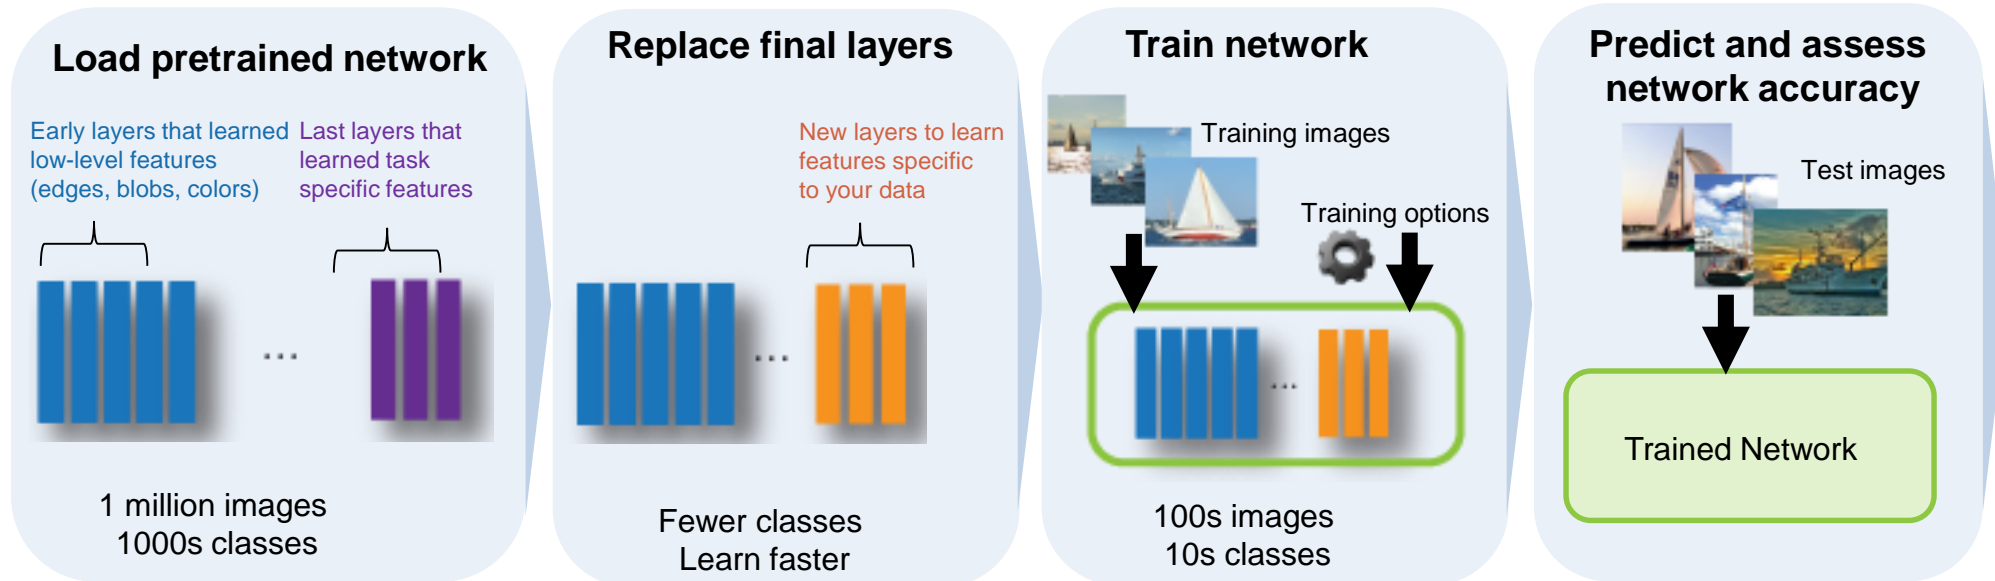
Transfer Learning Workflow – Step 3



Transfer Learning Workflow – Step 4



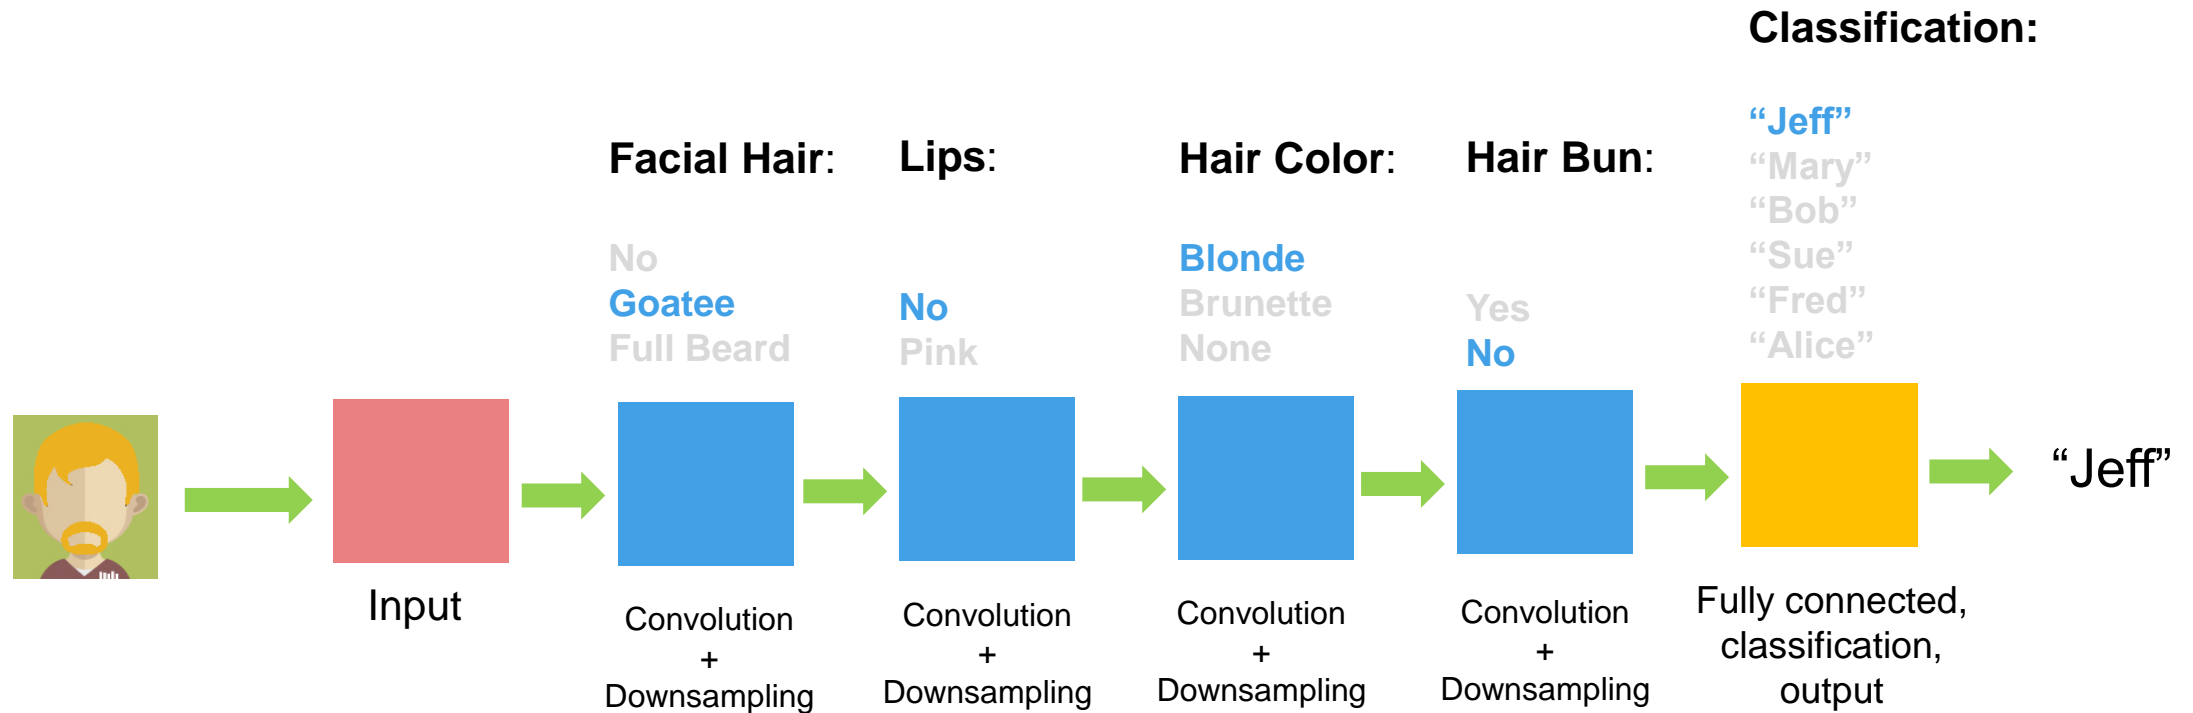
Transfer Learning Workflow



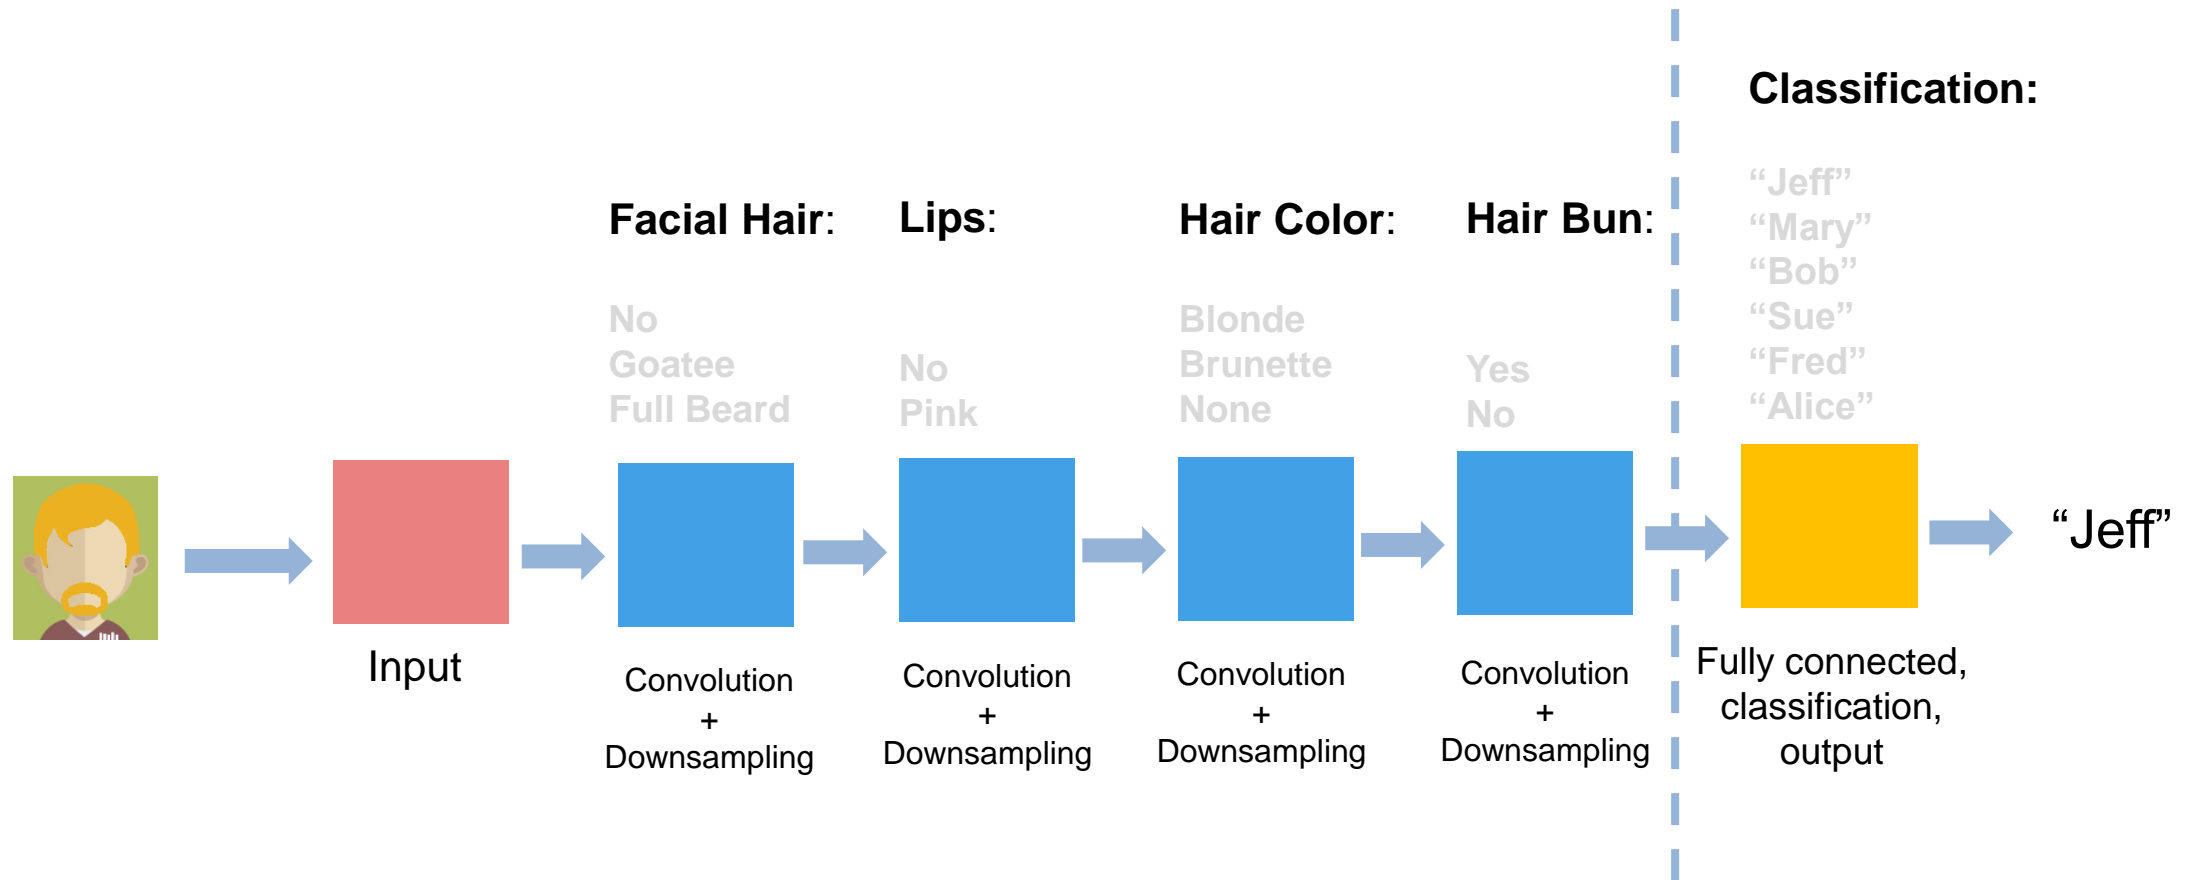
Transfer Learning



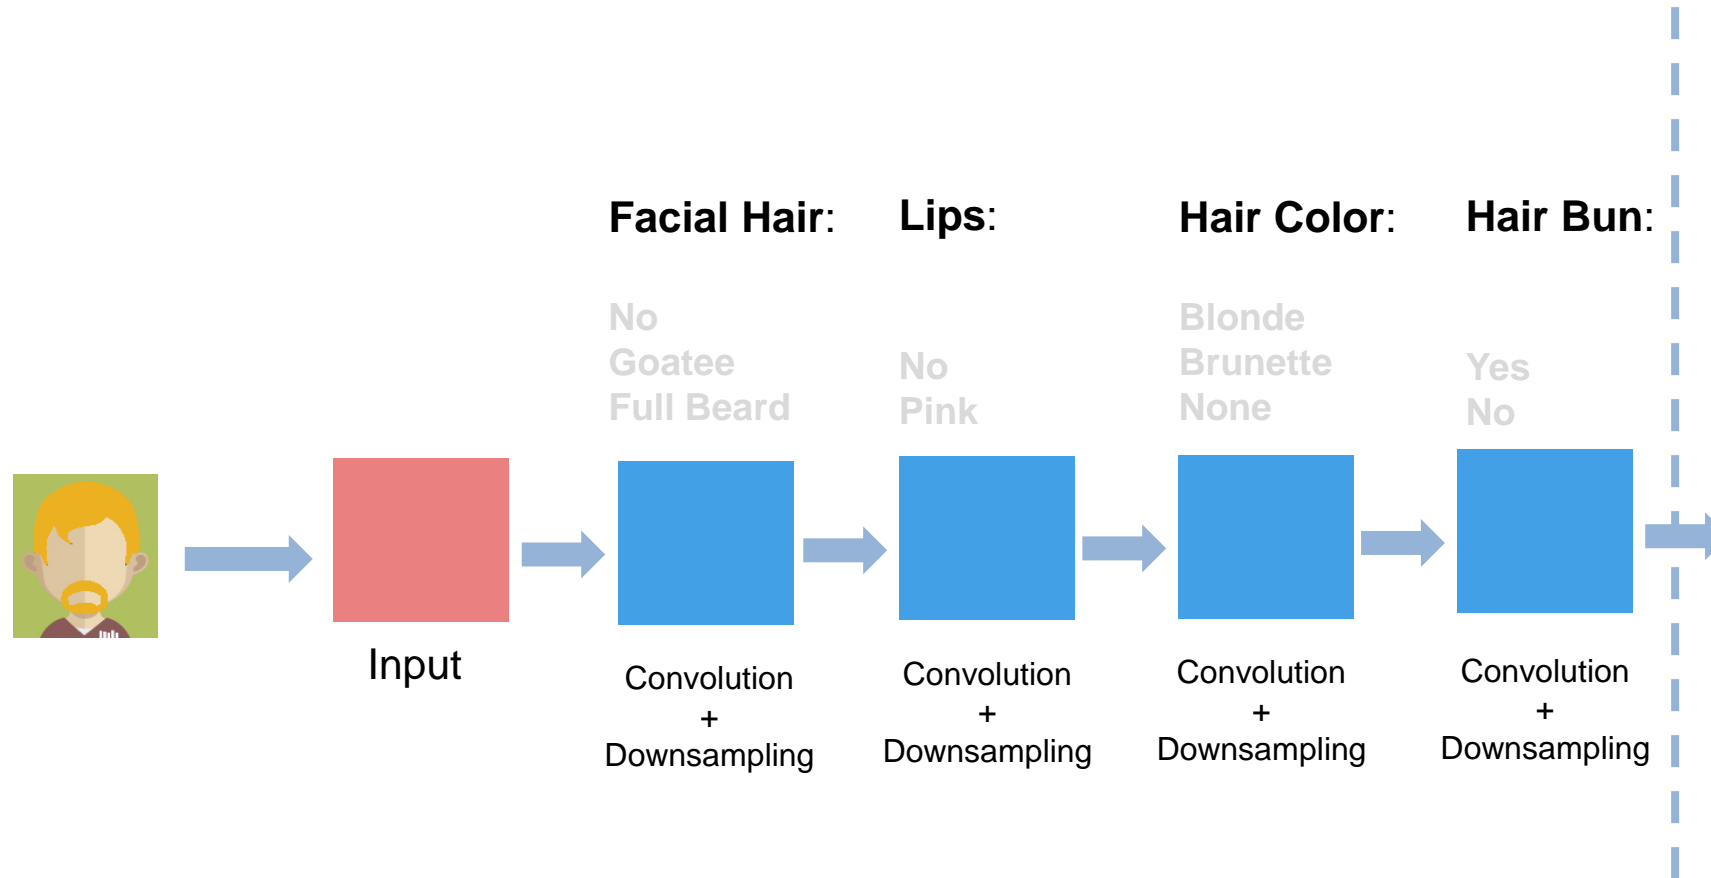
Transfer Learning



Transfer Learning



Transfer Learning



Transfer Learning



Transfer Learning





Exercise 4 – Transfer Learning

Purpose:

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

To Do:


1. Open `work_pretrainednetworks.mlx`.



Techniques Covered so Far

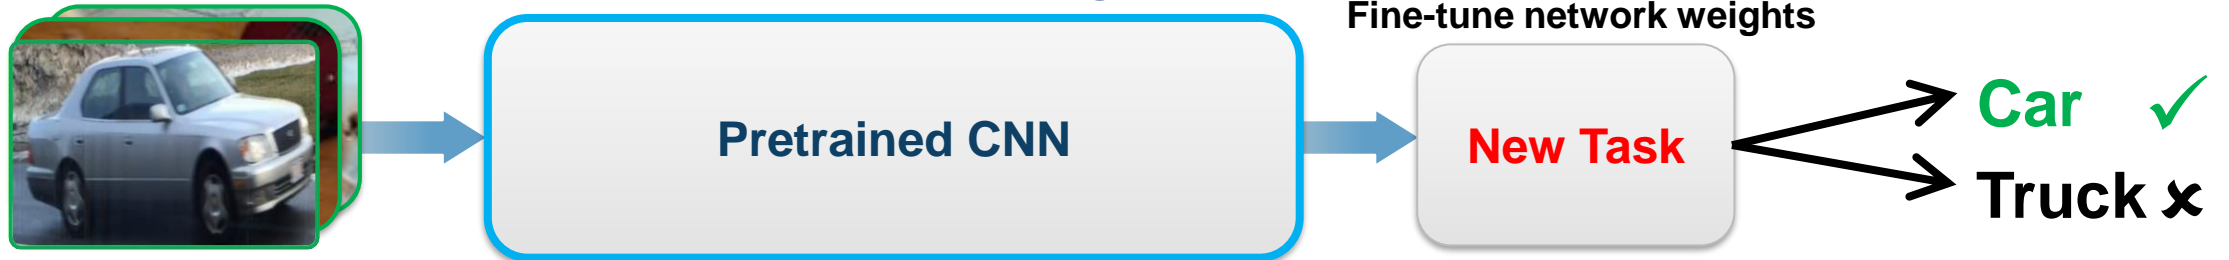
1. Train a Deep Neural Network from Scratch

Convolutional Neural Network (CNN)



2. Fine-tune a pretrained model (transfer learning)

Fine-tune network weights



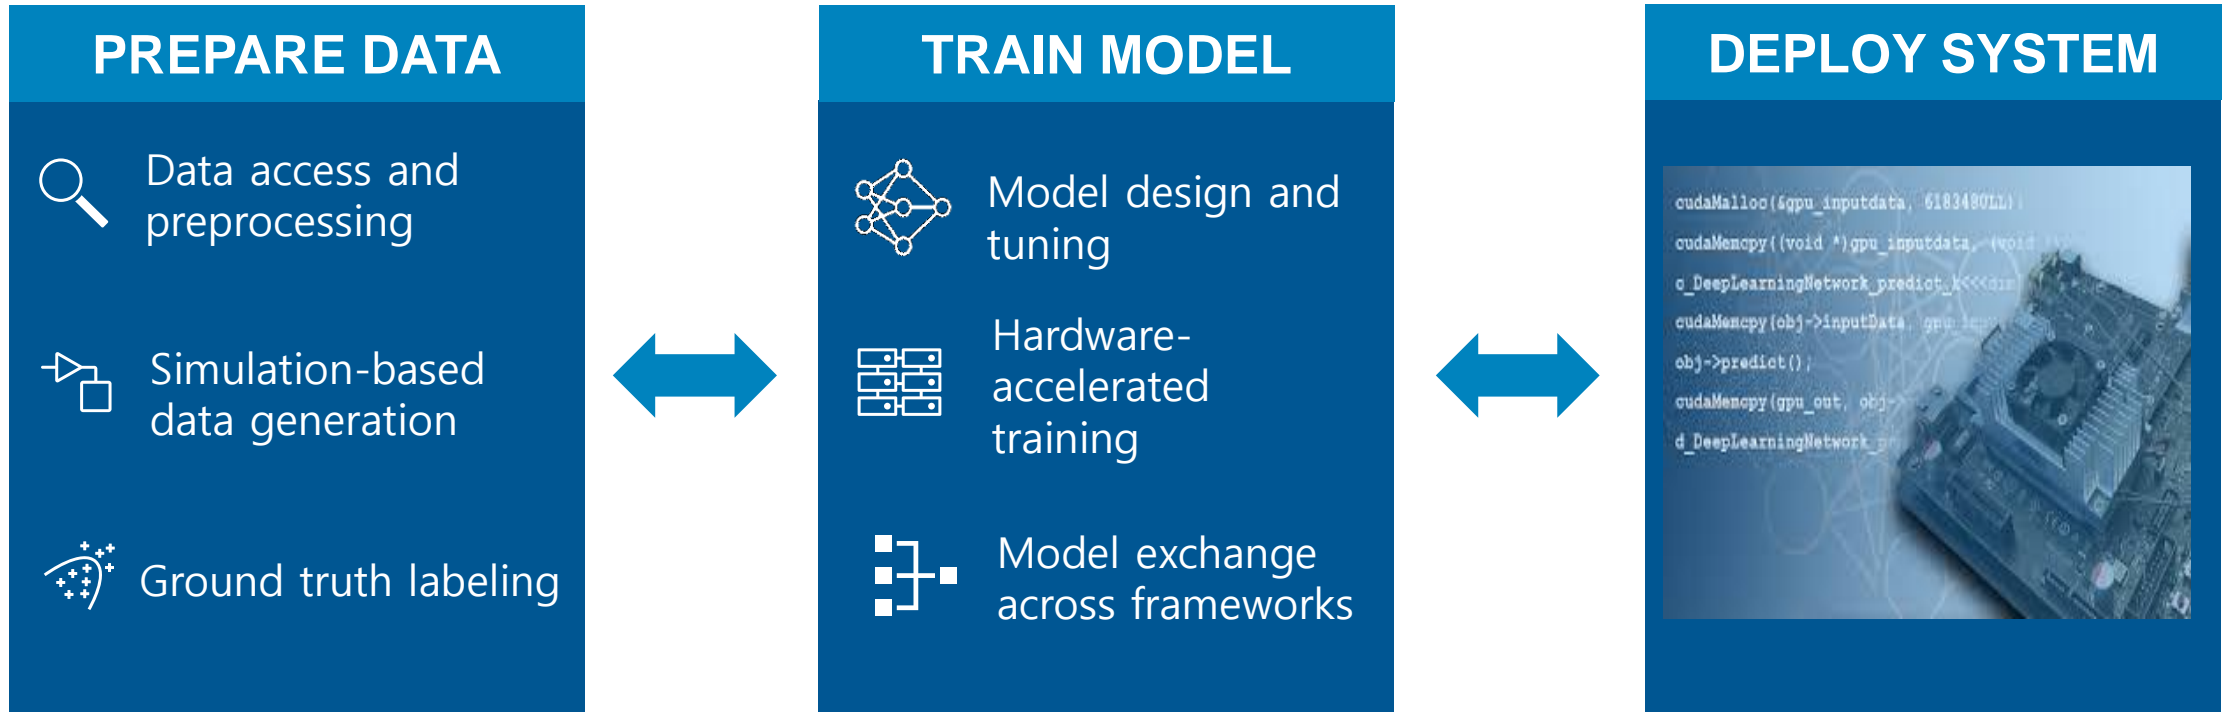
Deep Learning and Machine Learning Combined

3. Extract features with a pretrained CNN model



Click [HERE](#) to learn more about Machine Learning with MATLAB

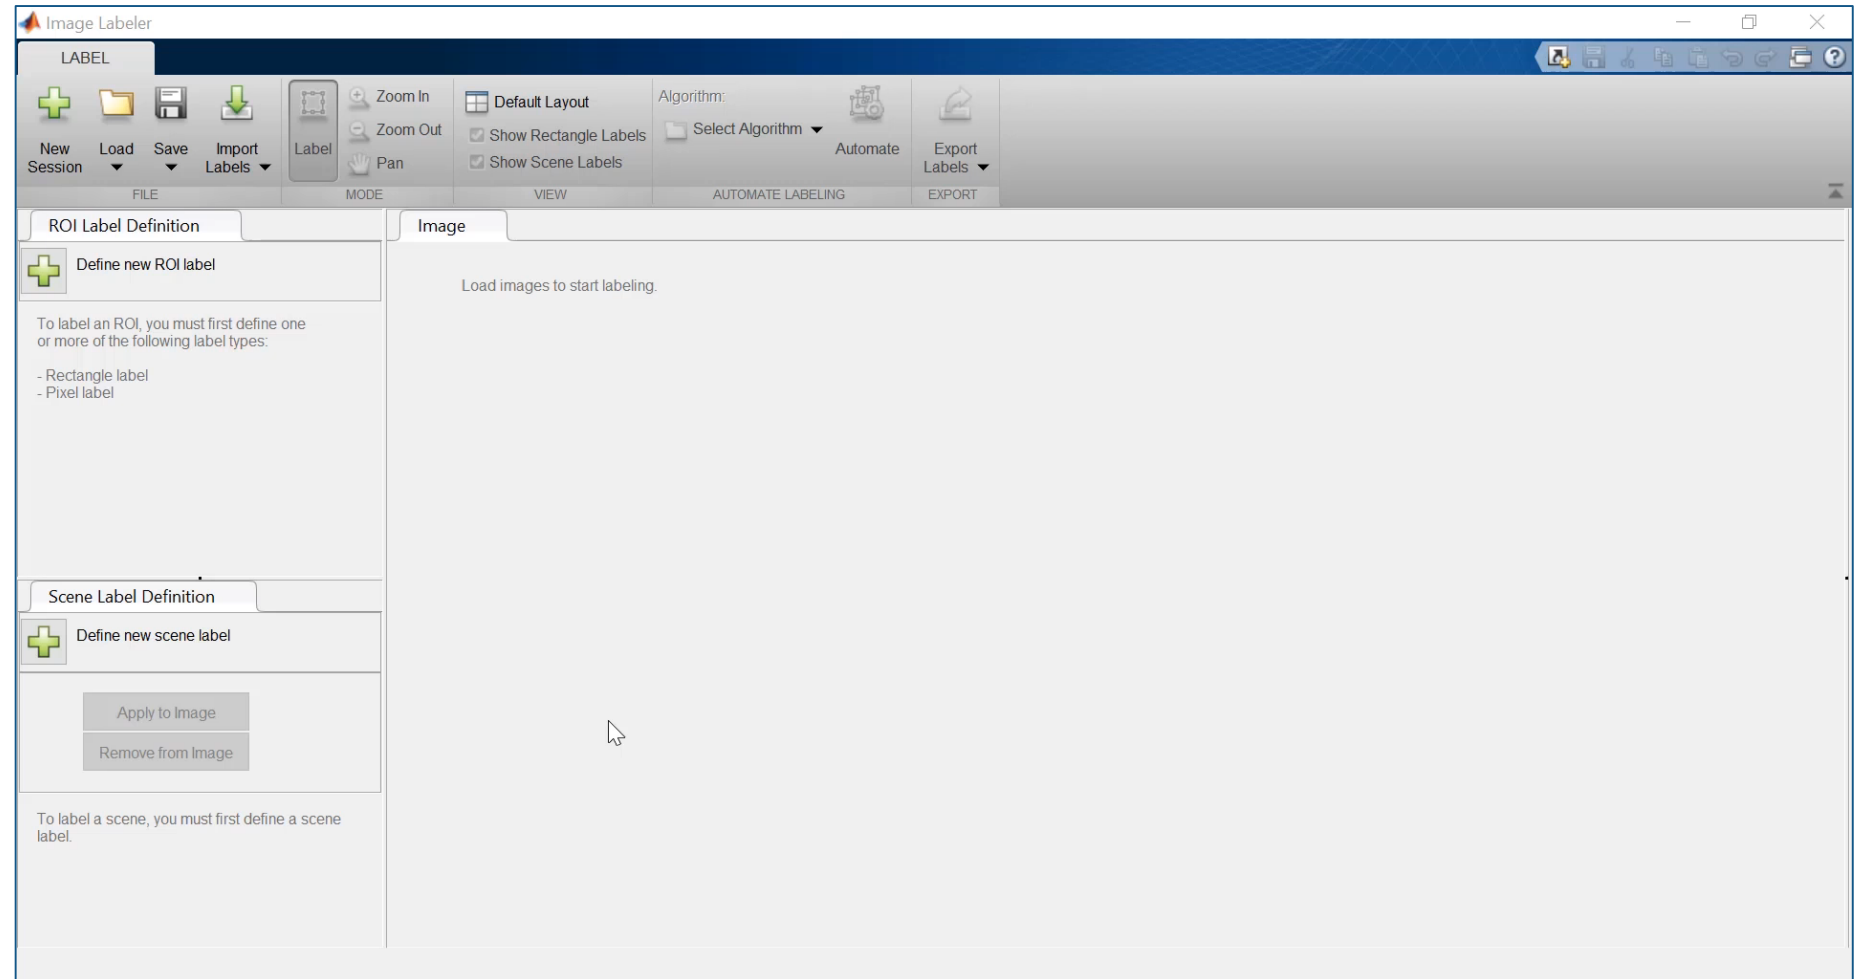
Deep Learning Workflow – Prepare Data



How do I label my data?

Image Labeler
+ Video labeler

Signal Labeler
+ Audio Labeler



How do I label my data?

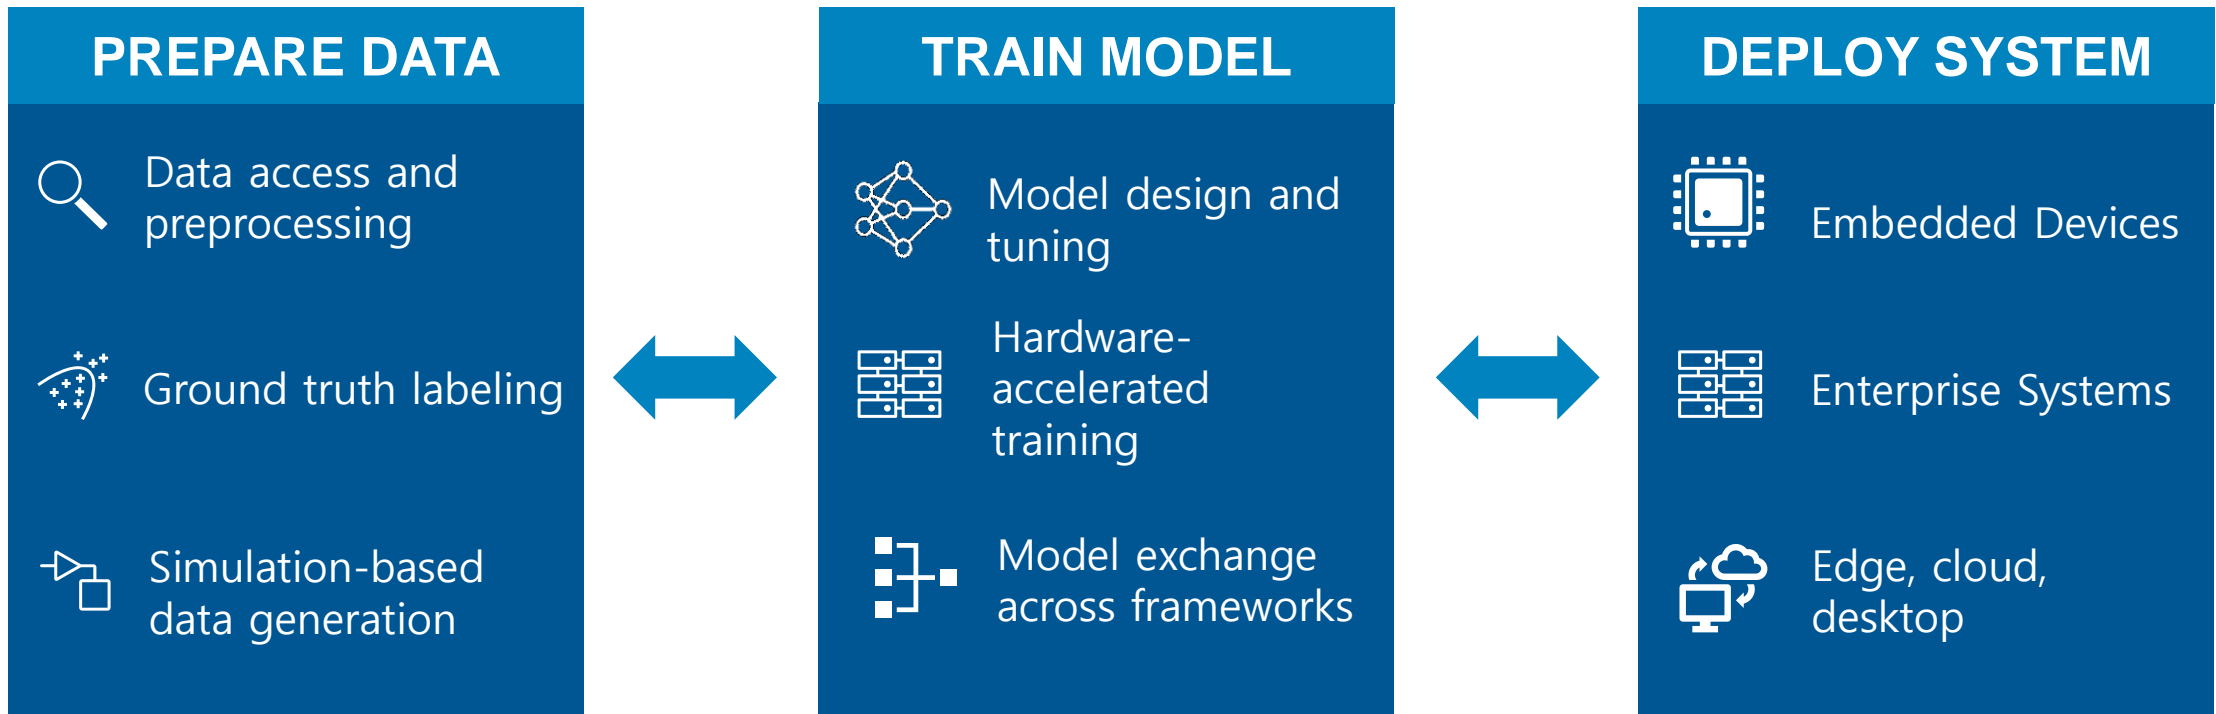
Image Labeler
+ Video labeler

Signal Labeler
+ Audio Labeler

The screenshot displays the 'Audio Labeler' application window. The title bar reads 'Audio Labeler - Counting-16-44p1-mono-15secs.wav'. The interface is divided into several sections:


- Top Bar:** Contains 'LABEL' and 'RECORD' tabs, a 'Cleanup' button, and standard window controls.
- Toolbars:** Includes 'FILE' (Load, Save, Import), 'DEVICE' (Audio Player, Settings), 'VIEW' (Default Layout, Legend), 'AUTOMATION' (Speech Detector, Speech to Text), and 'EXPORT' (Export).
- Data Browser:** A list of audio files, with 'Counting-16-44p1-mono-15secs.wav' selected.
- File Labels:** A panel for defining file-level labels, currently empty with a message: 'To label an audio file, you must first import or add a file label definition.'
- Audio File Info:** Displays metadata for the selected file: 'Counting-16-44p1-mono-15secs.wav', Channels: 1, Sample Rate: 44100 Hz, Duration: 15.534 s, Compression: Uncompressed, Bit Depth: 16 bits/sample, Location: C:\MATLAB\R2019b_Bash\toolb...
- ROI Labels:** A panel for defining regions of interest, currently empty with a message: 'To label a region of interest, you must first import or add an ROI label definition.'
- Main View:** A large black area showing a blue audio waveform. The x-axis is labeled with '5', '10', and '15' seconds. The y-axis represents amplitude. Playback controls are visible at the bottom of this area.
- Status Bar:** Shows 'Ready' on the left and 'Samples Underrun = 0' on the right.

Deep Learning Workflow – Deploy System

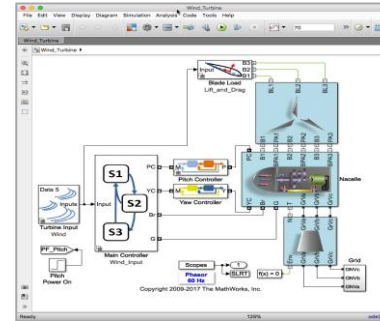



Deployment and Scaling for A.I.

Embedded Systems




MATLAB




Enterprise Systems




Embedded Deployment – Automatic Code Generation



MATLAB Code

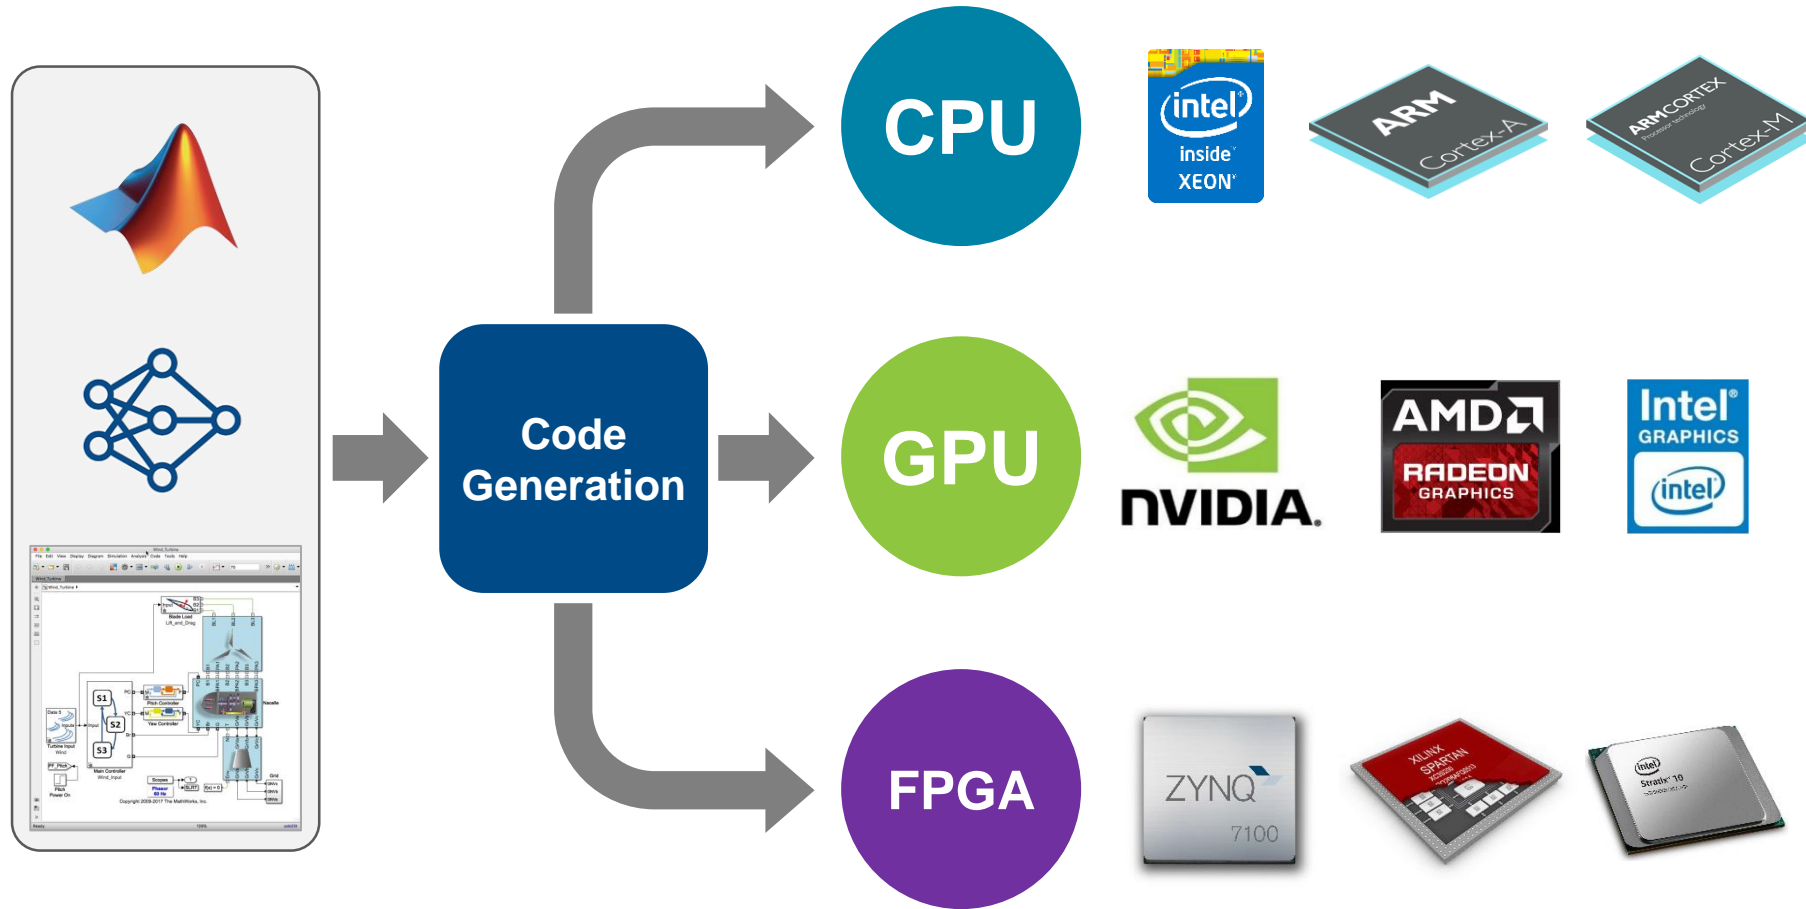


**Auto-generated Code
(C/C++/CUDA)**




**Deployment
Target**

Deploying Models for Inference

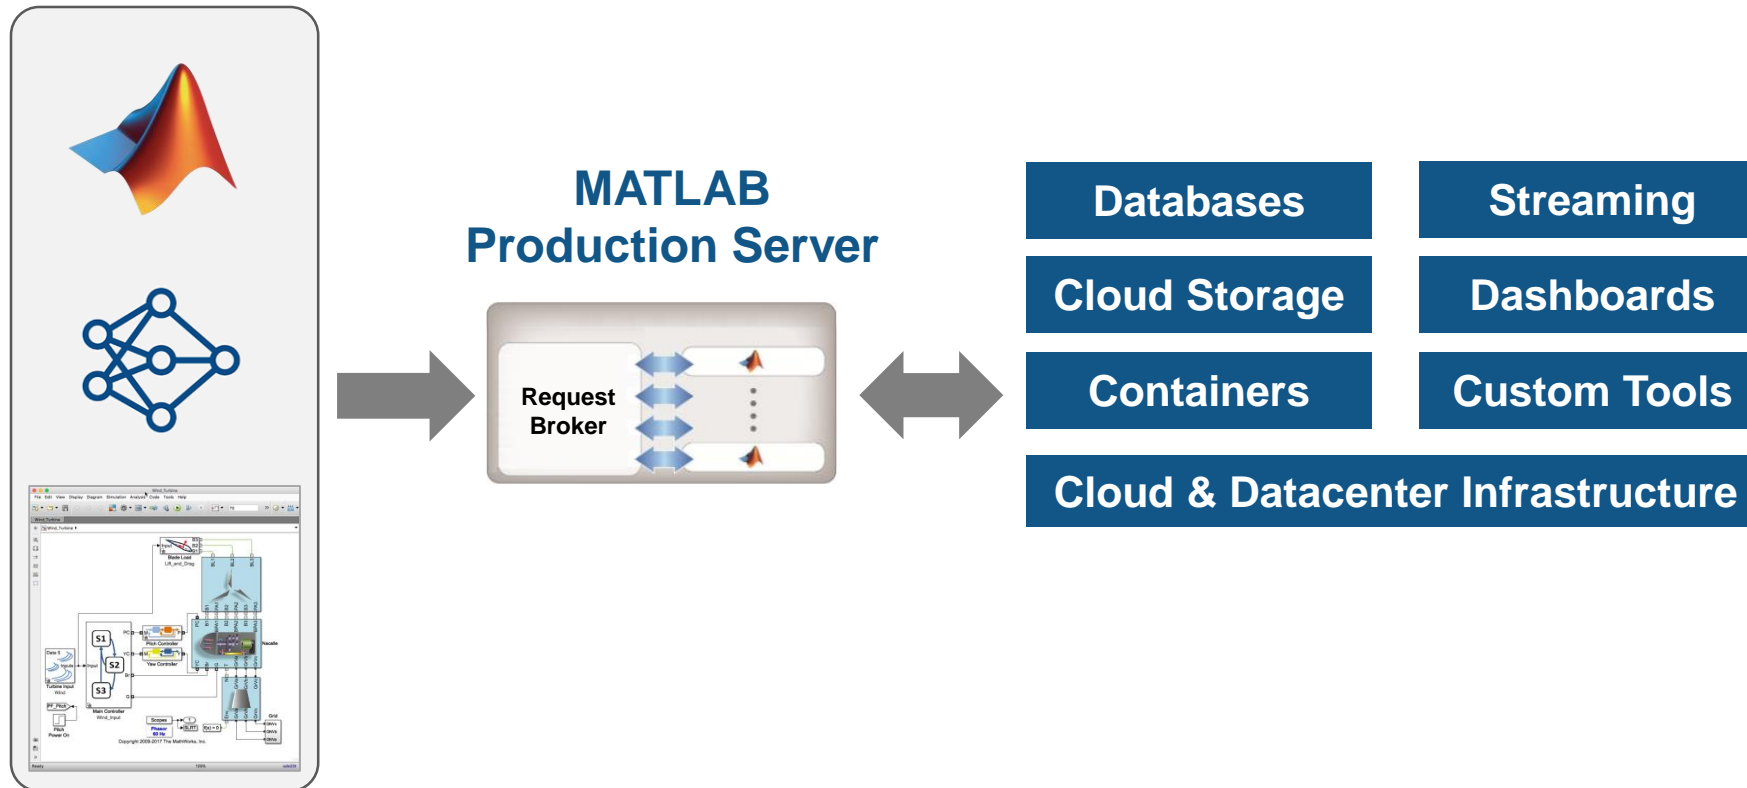


GPU Coder Inference Performance with ResNet-50 on Titan V

Batch 1



Deploy to Enterprise IT Infrastructure




Generate GPU Code for Deep Networks

GPU Coder

Generate Code for Deploying Deep Networks


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

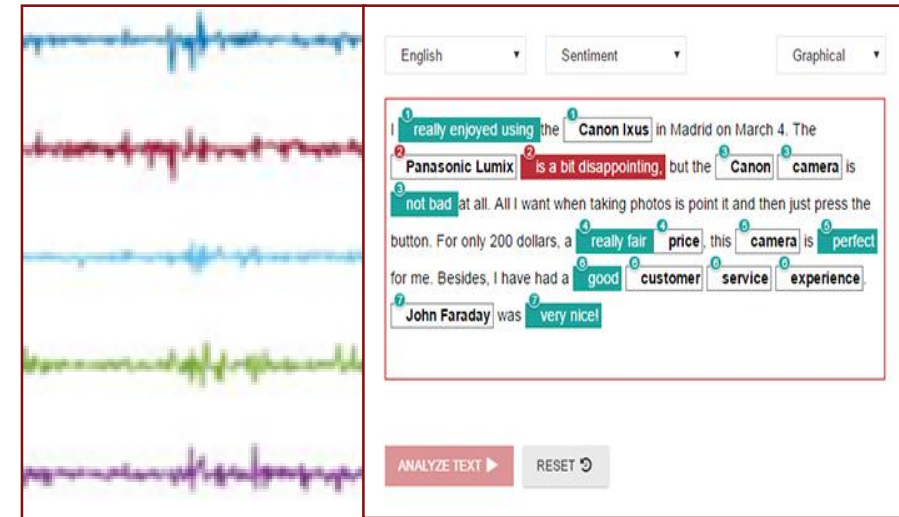
Selecting a Network Architecture

Image
Data



CNN


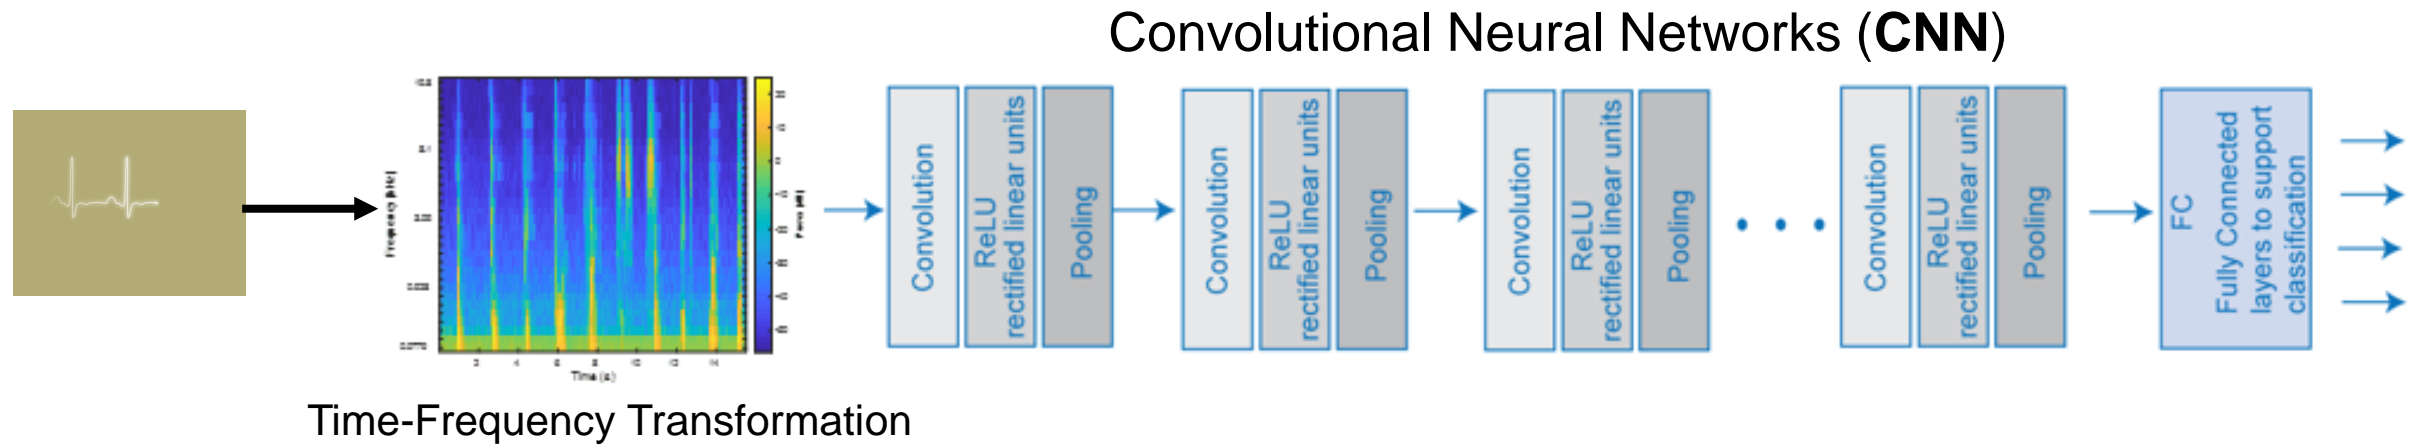
Signal or
Text Data



LSTM or CNN

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

Signal Processing Architectures

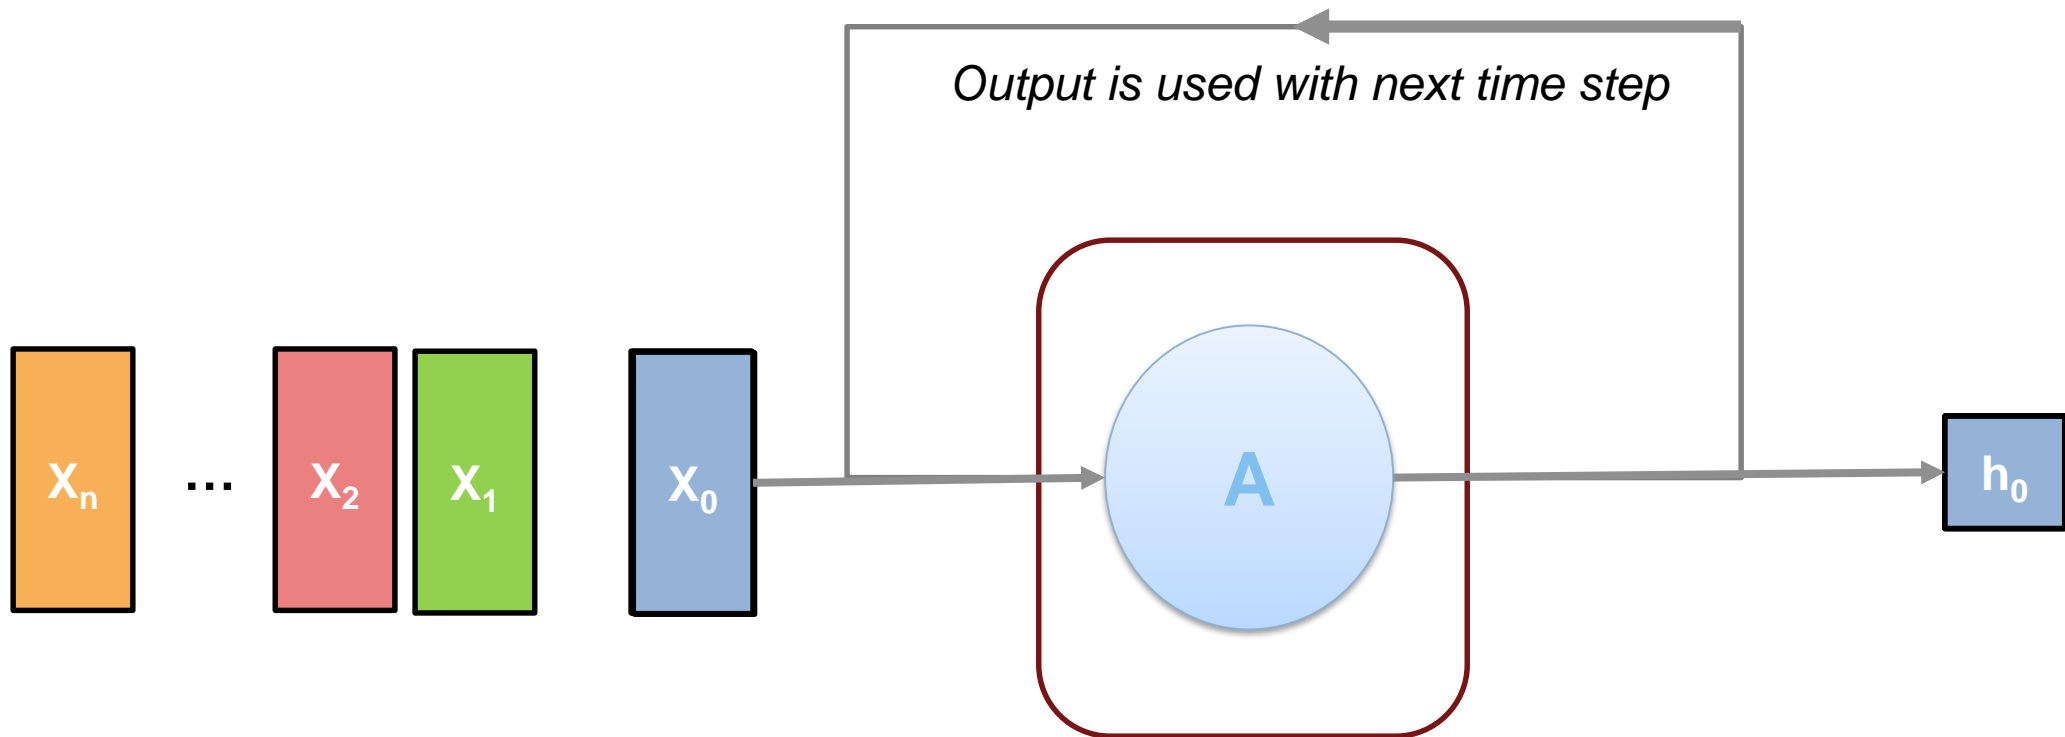


I was born in France...

... I speak _____ ?

Recurrent Neural Networks

Take into account previous data when making new predictions



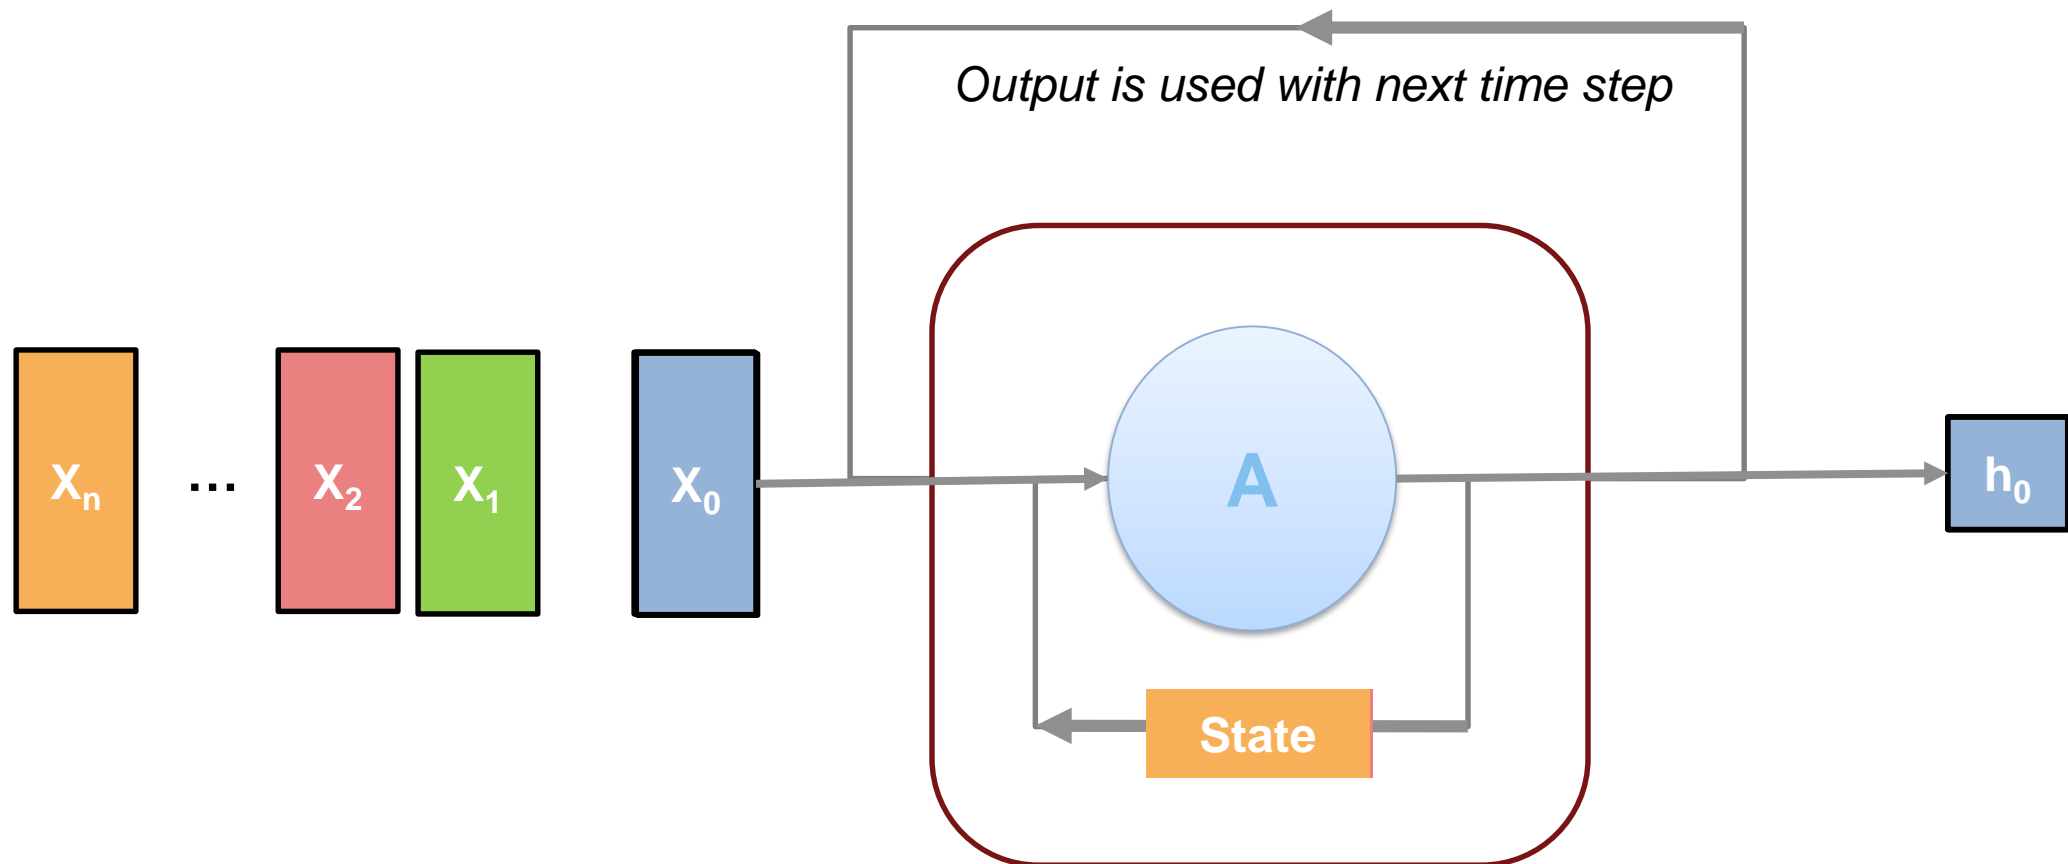
I was born in France...

[2000 words]

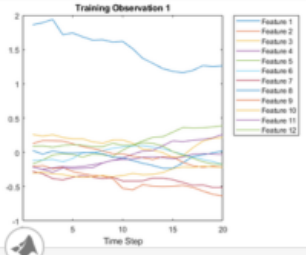
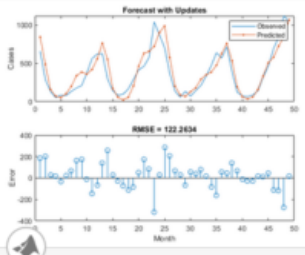
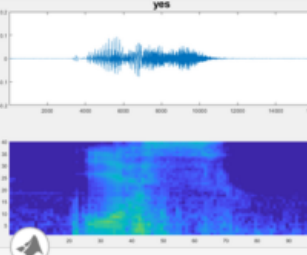
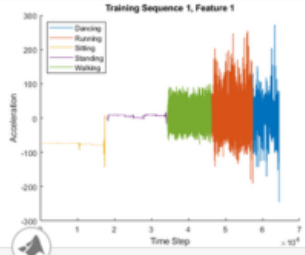
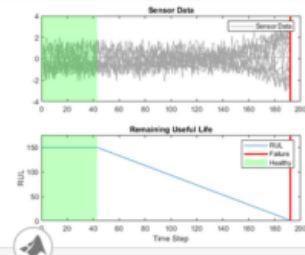

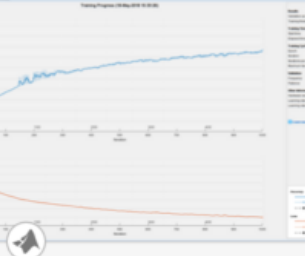
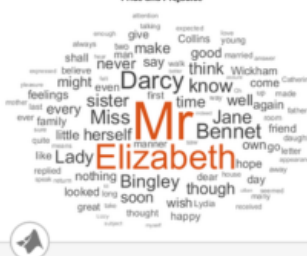

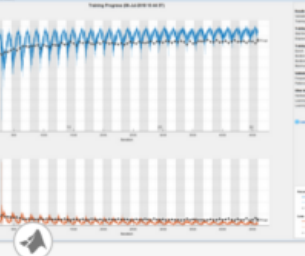
... I speak _____ ?

Long Short-Term Memory Network

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



Examples in MATLAB Documentation

 <p>Sequence Classification Using Deep Learning</p> <p>Classify sequence data using a long short-term memory (LSTM) network.</p> <p>Open Live Script</p>	 <p>Time Series Forecasting Using Deep Learning</p> <p>Forecast time series data using a long short-term memory (LSTM) network.</p> <p>Open Live Script</p>	 <p>Speech Command Recognition Using Deep Learning</p> <p>Train a simple deep learning model that detects the presence of speech commands in audio. The example uses the Speech Commands</p> <p>Open Script</p>	 <p>Sequence-to-Sequence Classification Using Deep Learning</p> <p>Classify each time step of sequence data using a long short-term memory (LSTM) network.</p> <p>Open Live Script</p>	 <p>Sequence-to-Sequence Regression Using Deep Learning</p> <p>Predict the remaining useful life (RUL) of engines by using deep learning.</p> <p>Open Live Script</p>
 <p>Classify Text Data Using Deep Learning</p> <p>Classify text descriptions of weather reports using a deep learning long short-term memory (LSTM) network.</p> <p>Open Live Script</p>	 <p>Generate Text Using Deep Learning</p> <p>Train a deep learning long short-term memory (LSTM) network to generate text.</p> <p>Open Live Script</p>	 <p>Pride and Prejudice and MATLAB</p> <p>Train a deep learning LSTM network to generate text using character embeddings.</p> <p>Open Live Script</p>	 <p>Word-By-Word Text Generation Using Deep Learning</p> <p>Train a deep learning LSTM network to generate text word-by-word.</p> <p>Open Live Script</p>	 <p>Classify Out-of-Memory Text Data Using Custom Mini-Batch Datasore</p> <p>Classify out-of-memory text data with a deep learning network using a custom mini-batch datasore.</p> <p>Open Live Script</p>

Exercise – ECG Signal Classification

Purpose:

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



To Do:

1. Open `work_ClassifyECGSignals.mlx`.

The spectral entropy measures how spiky flat the spectrum of a signal is. A signal with a spiky spectral entropy. The pentropy function estimates the spectral entropy based on a power spectrum which results in 255 time windows for a signal of 9000 samples. The 255-long time windows are used to calculate the spectral entropy.

Visualize the spectral entropy for each type of signal.

```
[pentropyA,tA2] = pentropy(aFib,fs);  
[pentropyN,tN2] = pentropy(normal,fs);  
  
plotPentropy(tN2,pentropyN,tA2,pentropyA);
```



MathWorks Engineering Support



Training



Consulting



Onsite Workshops and Seminars



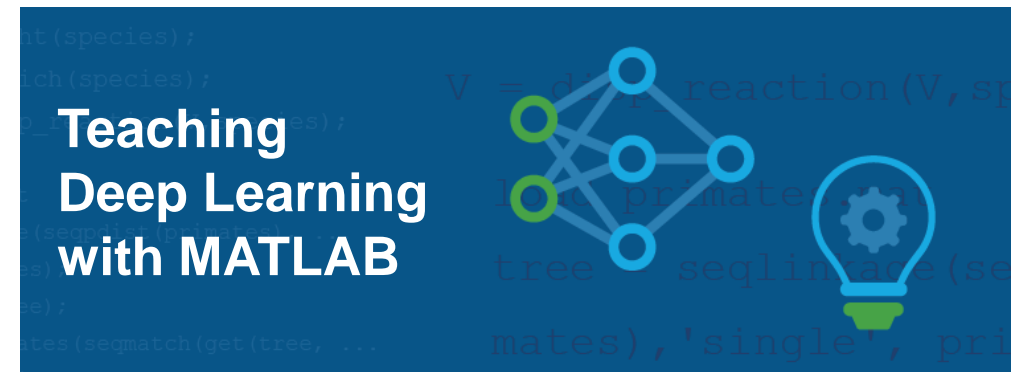

Guided Evaluations



Technical Support

Further Learning and Teaching

- [Deep Learning Onramp](#)
 - 2 hr online tutorial
- [Deep Learning Workshop](#)
 - 3 hr hands on session
 - Contact us to schedule
- [Deep Learning Training](#)
 - 16 hr in depth course
 - Online or Instructor Lead
- [Teaching Deep Learning with MATLAB](#)
 - Curriculum support



Thank you!