

The slide features a dark blue background with a large, light blue triangular shape on the right side. Inside this triangle, there is a 3D wireframe plot of a surface, colored with a gradient from yellow at the top to blue at the bottom. In the background, there are faint white lines representing a signal waveform and a circuit diagram with various components and arrows.

# MATLAB EXPO 2017

## Introduction to Machine Learning and Deep Learning

Conor Daly

# Machine learning in action






## CamVid Dataset

1. Segmentation and Recognition Using Structure from Motion Point Clouds, ECCV 2008
2. Semantic Object Classes in Video: A High-Definition Ground Truth Database, Pattern Recognition Letters

# Machine learning is everywhere

- Image recognition
- Speech recognition
- Stock prediction
- Medical diagnosis
- Predictive maintenance
- Language translation
- and more...

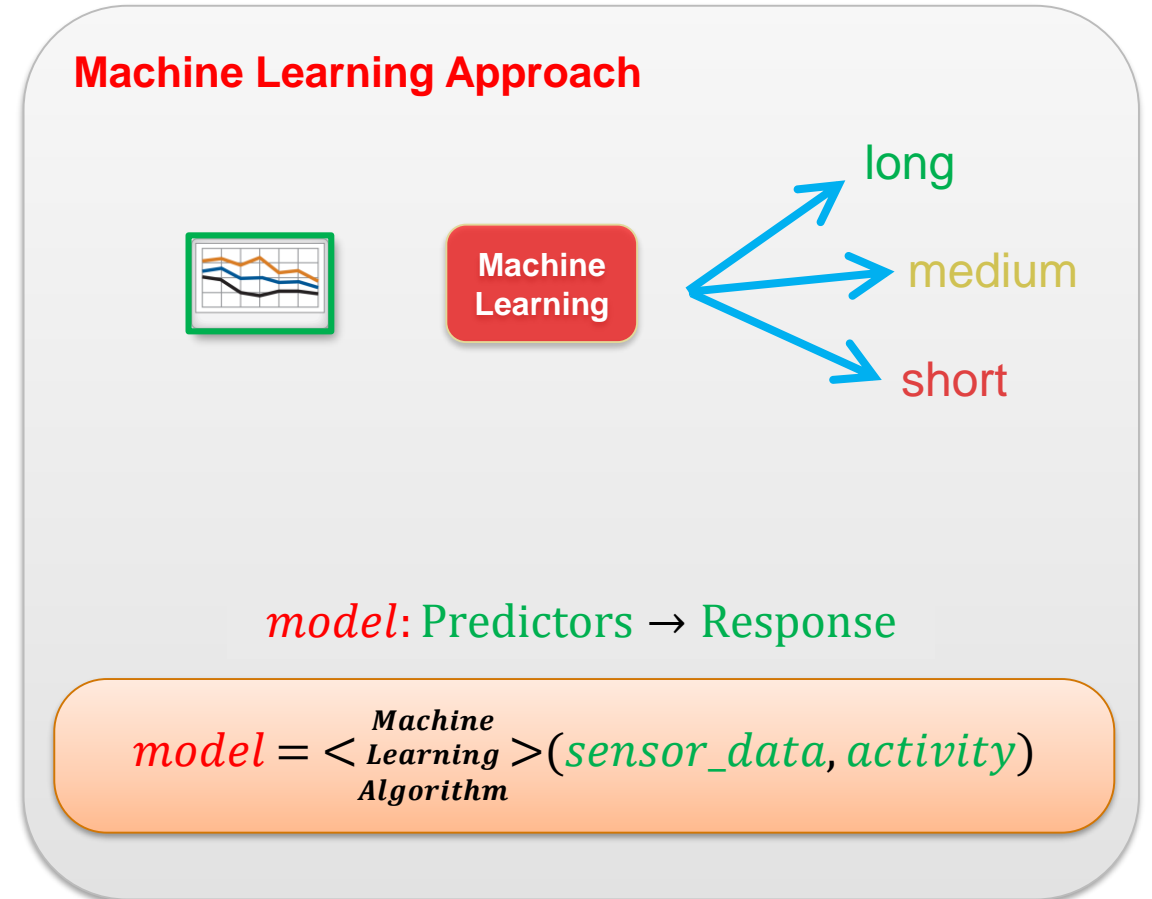
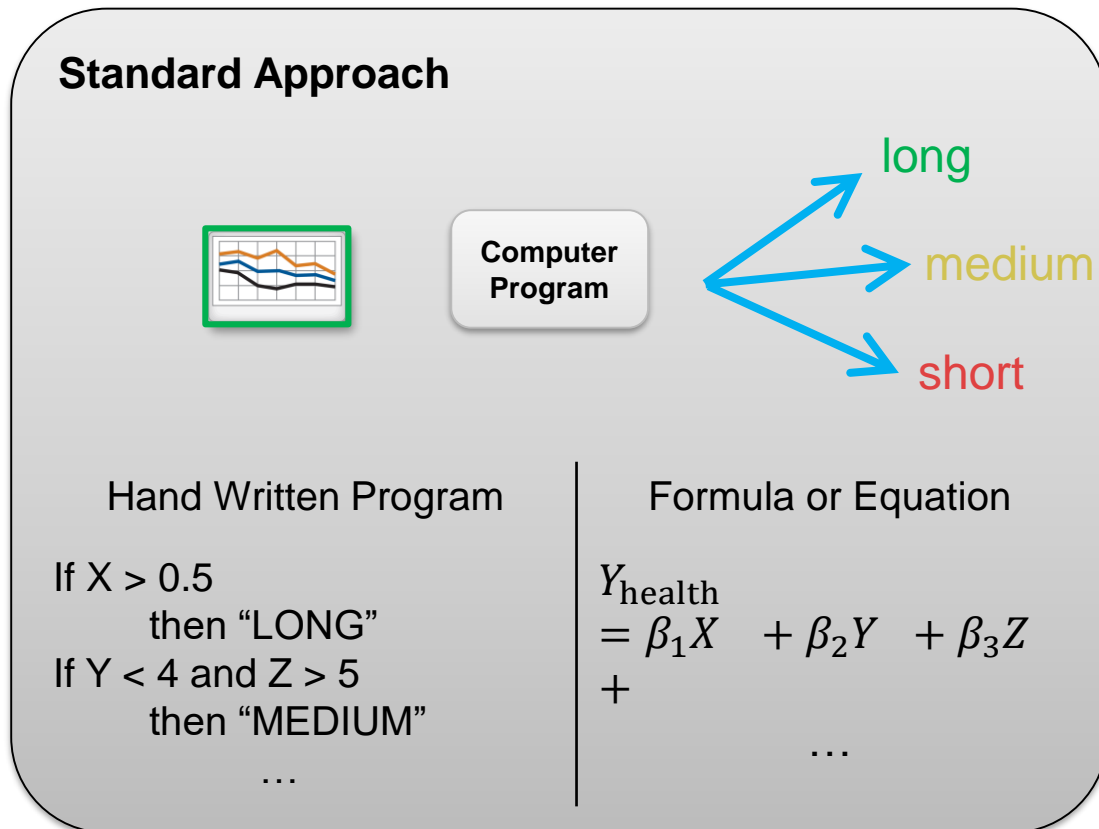


# Agenda

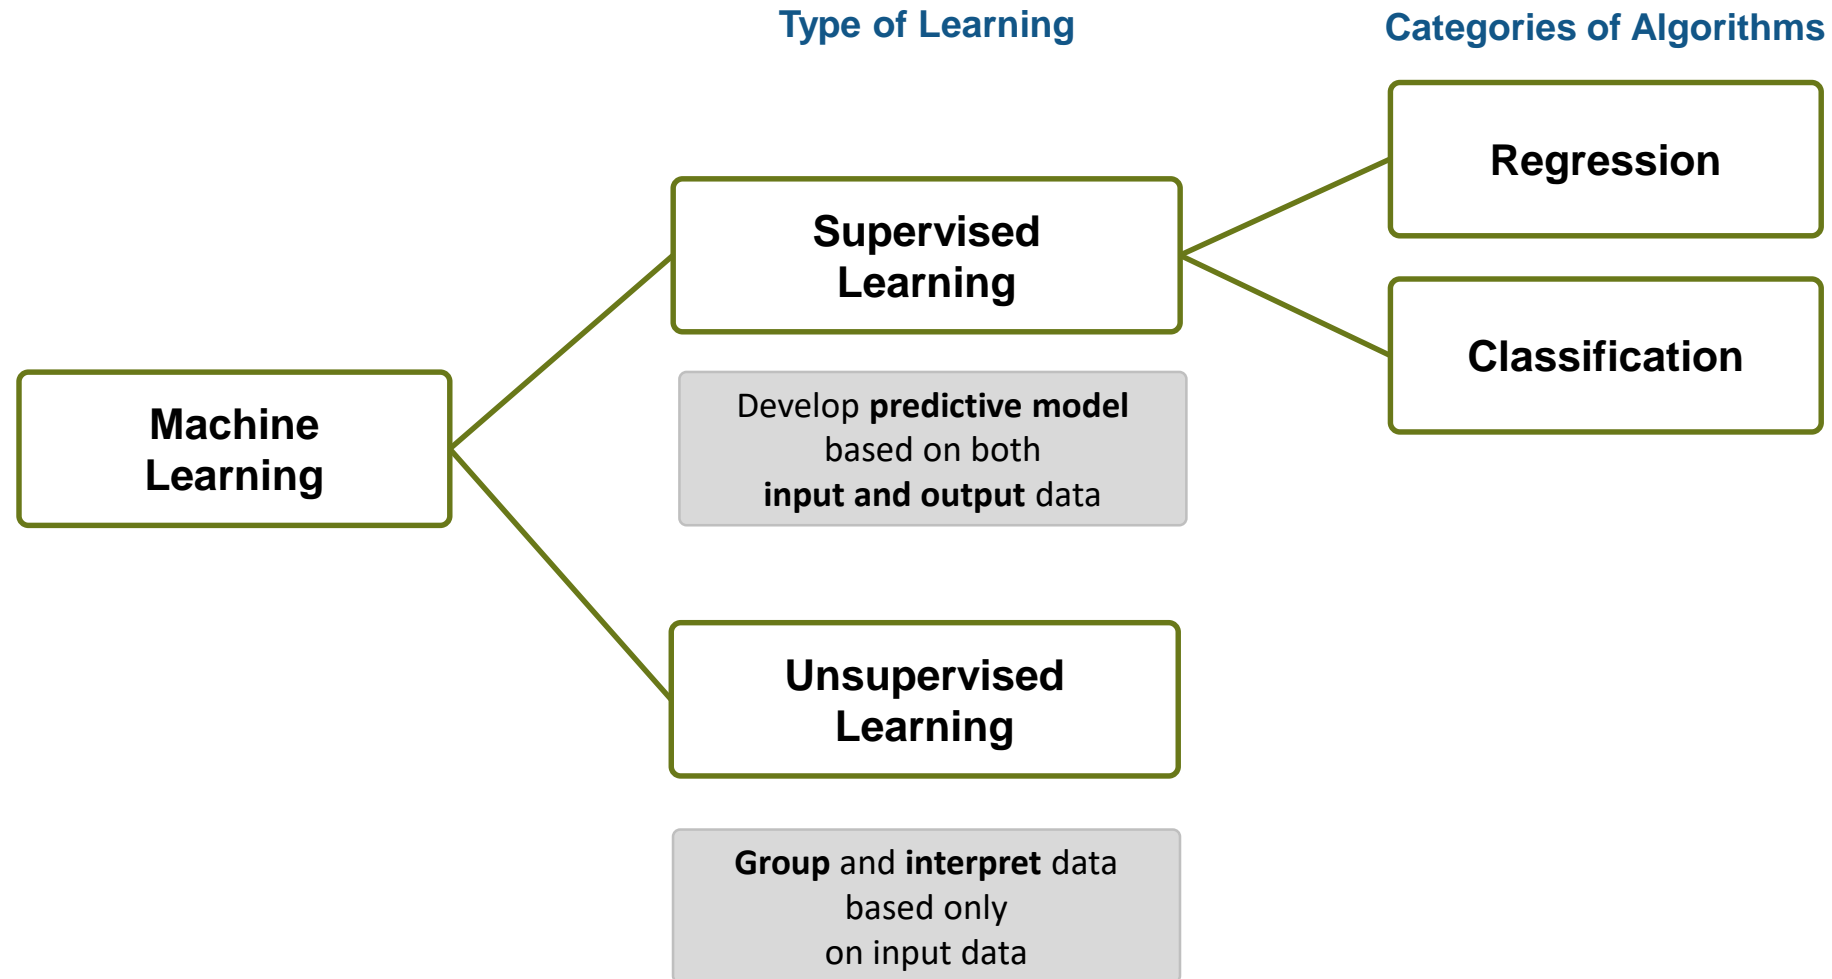
1. **Machine learning – predictive maintenance**
2. Deep learning – build a digits classifier
3. Predictive maintenance revisited – a deep learning approach

# What is machine learning?

Machine learning uses **data** and produces a **model** to perform a **task**



# Machine Learning: problem specific overview



# Predictive maintenance of turbofan engine

Sensor data from 100 engines of the same model


## Motivation

- Import and analyze historical sensor data
- Train model to predict when failures will occur
- Deploy model to run on live sensor data
- Predict failures in real time


Data provided by NASA PCoE

<http://ti.arc.nasa.gov/tech/dash/pcoe/prognostic-data-repository/>

MATLAB EXPO 2017




# Use historical data to predict when failures will occur





# Preprocessing and classifying our input data



# Agenda

1. Machine learning – predictive maintenance
2. **Deep learning – build a digits classifier**
3. Predictive maintenance revisited – a deep learning approach

# Can you tell the difference? Japanese or Blenheim Spaniel?




Blenheim Spaniel



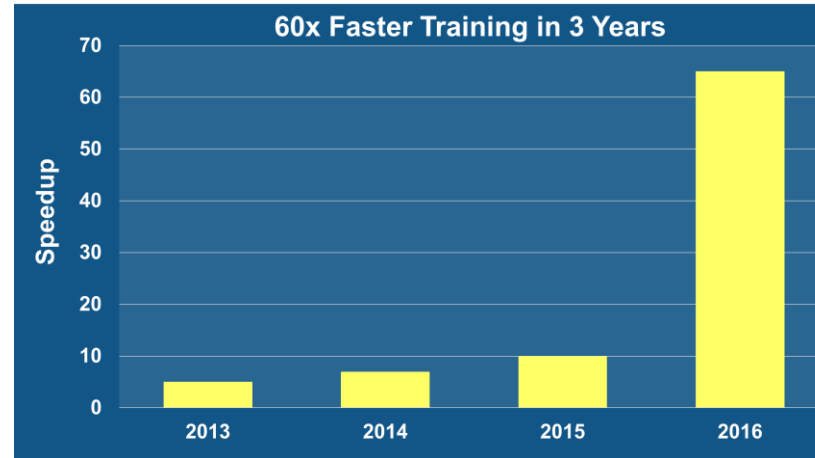
Japanese Spaniel

# Why is deep learning so popular now?




# Deep learning enablers

Acceleration with GPUs



Massive sets of labeled data




Availability of state of the art models from experts



# Machine learning vs deep learning


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

Deep learning algorithms also **scale with data** – traditional machine learning **saturates**



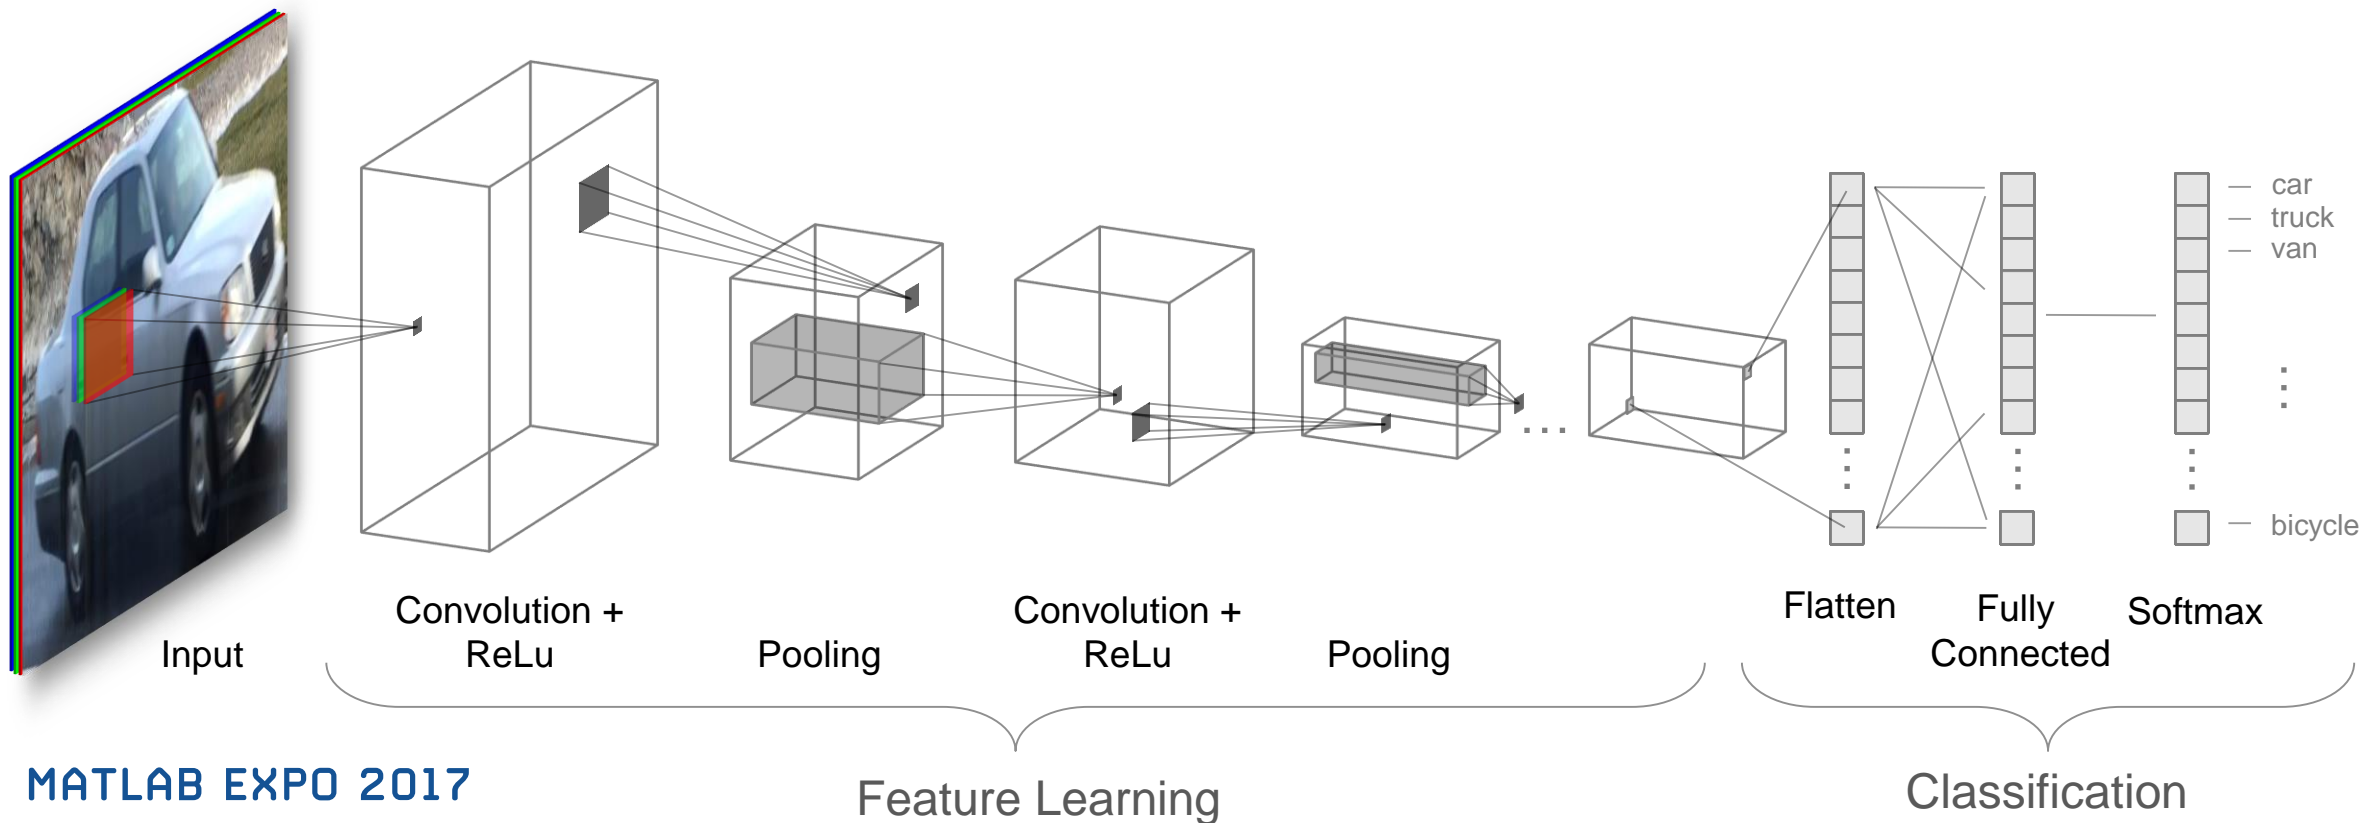
# Deep learning and neural networks

- Deep learning == neural networks
- Data flows through network in **layers**
- **Layers** provide transformation of data



# Convolutional neural networks


- Train “deep” neural networks on structured data (e.g. images, signals, text)
- Implements Feature Learning: Eliminates need for “hand crafted” features
- Trained using GPUs for performance



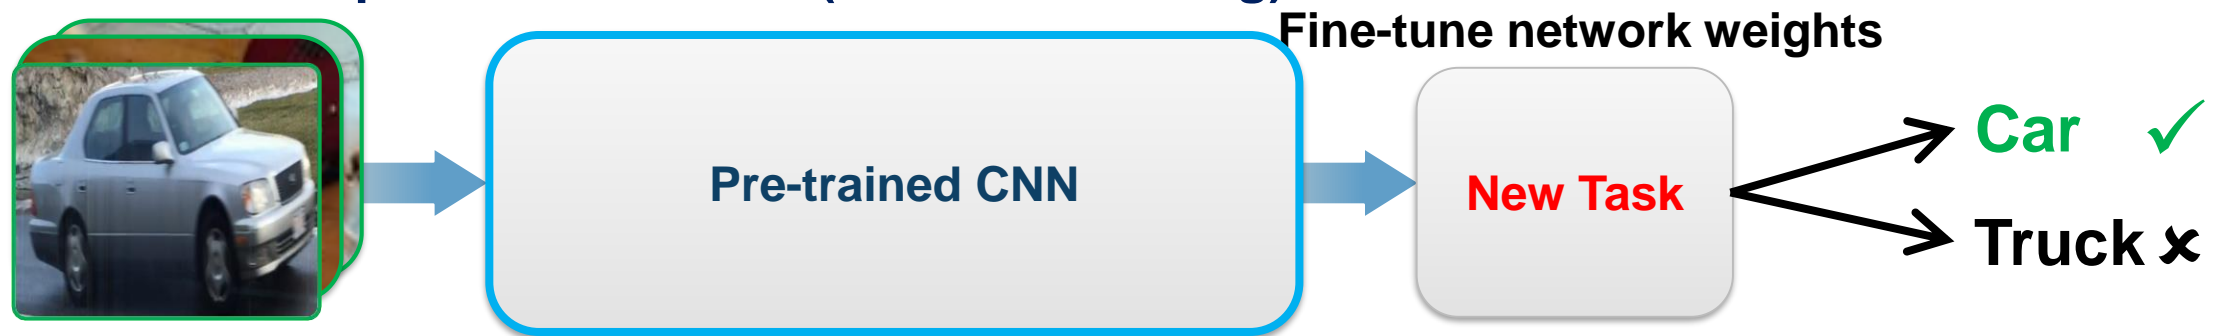


# Two approaches for deep learning

## 1. Train a Deep Neural Network from Scratch




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



# Two deep learning approaches

## Approach 1: Train a Deep Neural Network from Scratch



### Recommended when:


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

## Two deep learning approaches

### Approach 2: Fine-tune a pre-trained model (transfer learning)

#### CNN trained on massive sets of data

- Learned robust representations of images from larger data set
- Can be fine-tuned for use with *new data or task* with small – medium size datasets




#### Recommended when:

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

# Digits classification

<b>What?</b>	A set of 'handwritten' digits from 0-9 (c.f. MNIST)
<b>Why?</b>	An easy task for machine learning beginners
<b>How many?</b>	60,000 training images 10,000 test images
<b>Best results?</b>	99.79% accuracy



# Agenda

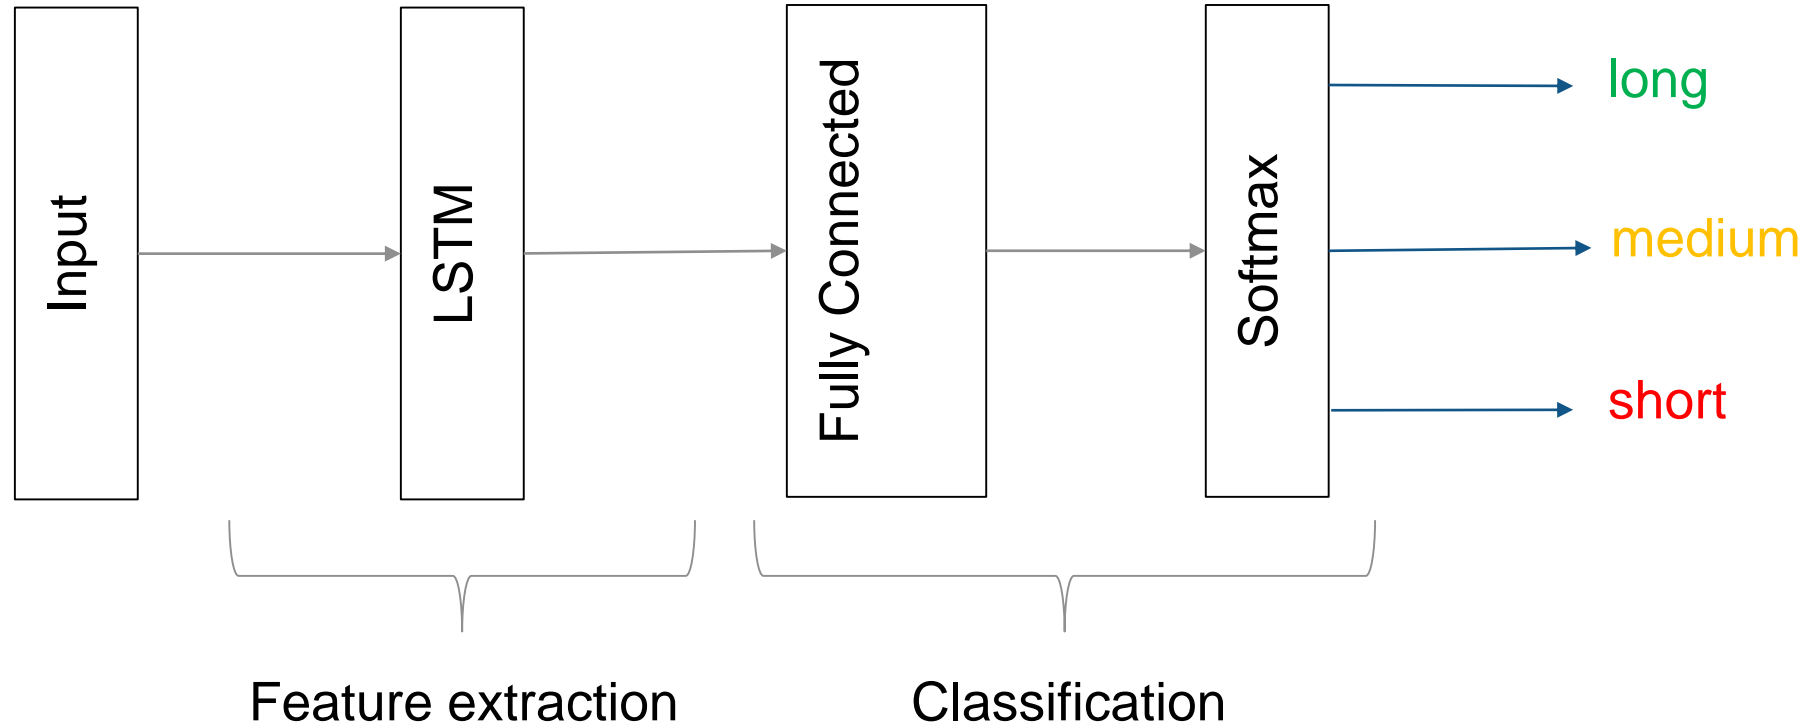
1. Machine learning – predictive maintenance
2. Deep learning – build a digits classifier
3. **Predictive maintenance revisited – a deep learning approach**

# Tackle time series data with LSTM

- How can we apply deep learning to time series data?
- One approach is to use Long Short-Term Memory (LSTM) neural networks
- These networks learn long-term temporal dependencies
- LSTMs work well with sequential input data, for example:
  - Time series
  - Text
  - Video

# LSTM classification networks

- LSTM is used to extract time series features



# Shakespearean LSTM

```
startPhrase = 'Accelerating the pace of ';
numberOfCharactersToGenerate = 500;

generatedText = iGenerateText( startPhrase, numberOfCharactersToGenerate, net )
```

```
generatedText =
    'Accelerating the pace of me!

    DUKE VINCENT:
    Boy to hear you!

    MENENIUS:
    I'll be him, and he shall see how in Rome rags,
    Bearing the part of my father's royal dead,
    And where be leave to gain.

    ISABELLA:
    Set now fair ground indeed the last of all.

    FLORIZELO:
    No mark bench, say. Grepert, sir, and boss
    of proclait is extraity, for the senators
    of more you over it bleeding.
```

- Text generated from deep LSTM network
- Network has learned long-term text style of Shakespeare
- E.g. punctuation, character-name capitalization



# Thanks for listening!

Today we've looked at:

1. Machine learning – predictive maintenance
2. Deep learning – build a digits classifier
3. Predictive maintenance revisited – a deep learning approach

What to learn more/try it for yourself?

Try **MATLAB Deep Learning Onramp**