AI in Embedded Systems

Class 2: Toolkit - Algorithms

February 25, 2020
Louis W. Giokas

Presented by:
CEC CONTINUING EDUCATION CENTER
Digi-Key ELECTRONICS

This Week's Agenda

Monday	Overview and Requirements
Tuesday	Toolkit - Algorithms
Wednesday	Toolkit - Hardware
Thursday	Systems
Friday	Application Examples

Course Description

This topic will cover a new approach to developing embedded systems which includes AI/Machine Learning(ML) as an element of the suite of tools available. As microcontroller devices become more powerful, these techniques are now within reach. We will look at the types of AI/ML algorithms available and appropriate in embedded systems. We will also look at how interaction with higher level systems, such as cloud analytics, can be integrated to create systems that evolve and improve over time and space.

Today's Agenda

- Overview
- Basic Statistical Techniques
- Machine Learning
- Deep Learning
- Conclusion/Next Class

Overview

- We will review in detail the algorithmic landscape making up the software and analytic side of the system. This class will provide an overview of the topic of AI/ML and some history, as well as examples of detailed algorithms.

Overview

- AI consists of statistical techniques that rely on data to guide decision making
- Models of the data are used, and then actual data is used to "train" these models
  - This consists of letting the data guide the selection of parameters of the models
- This is in contrast to selecting the parameters by hand or by some external mechanism

Overview

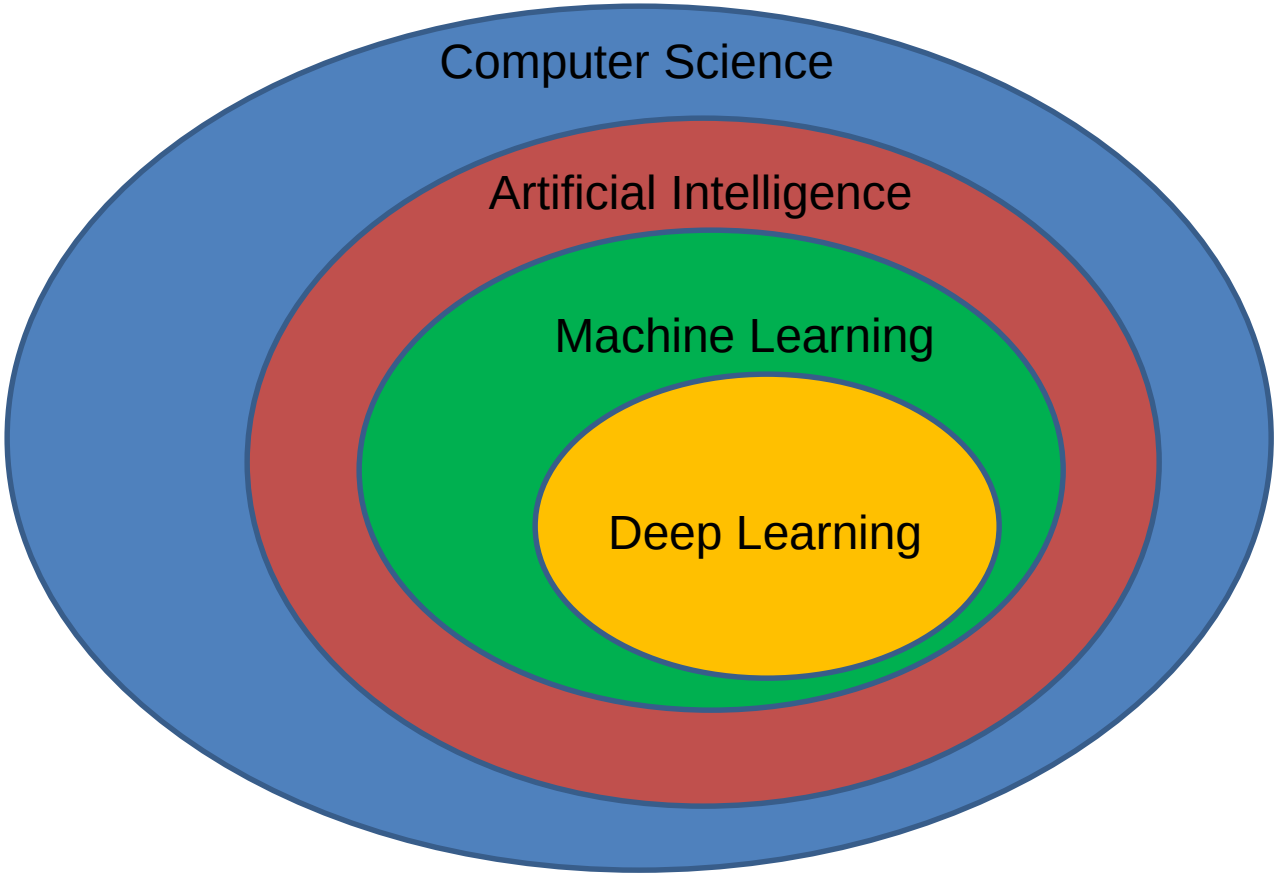
Computer Science
Artificial Intelligence
Machine Learning
Deep Learning

Overview

- We will not be looking at AI techniques outside of Machine Learning, although they may have some application
  - Not often used in the embedded context
  - An example of this is Expert Systems
    - Work with experts to extract their decision-making process in particular situations
    - Code this process to automate

Machine Learning

- Machine Learning (ML) encompasses a number of areas
  - Data Mining
  - Inference
  - Prediction
- The group of techniques is sometimes called Statistical Learning, or sometimes Data Science
  - Originally developed in the statistics community
  - Picked up and expanded by the computer science community

Machine Learning

- We often speak of ML as learning from the data
  - Extract patterns, trends and understand what the data is telling us
  - For example, this can be used to find patterns in the data we did not know beforehand
    - Data Mining
  - We can also cluster data to then predict if a new measurement is expected, or somehow an outlier

Machine Learning

- Typically use a training data set
  - Take, say 80% of the data to train the algorithm
  - Then take the remaining data to measure the accuracy of the parameters set
    - Gives a measure of predictive power
    - Will typically will run this process many times, with random subsets of the data
    - Retrain the algorithm until the desired accuracy is obtained

Machine Learning

- Prediction algorithm examples
  - Least Squares and Nearest Neighbors
  - Linear Regression
- Classification algorithm examples
  - Linear Discriminant Analysis (LDA)
  - Logistic Regression
  - Separating Hyperplanes
    - Can handle situations in higher dimensions

Machine Learning

- Tree Based Methods: Partition the feature space into a set of rectangles

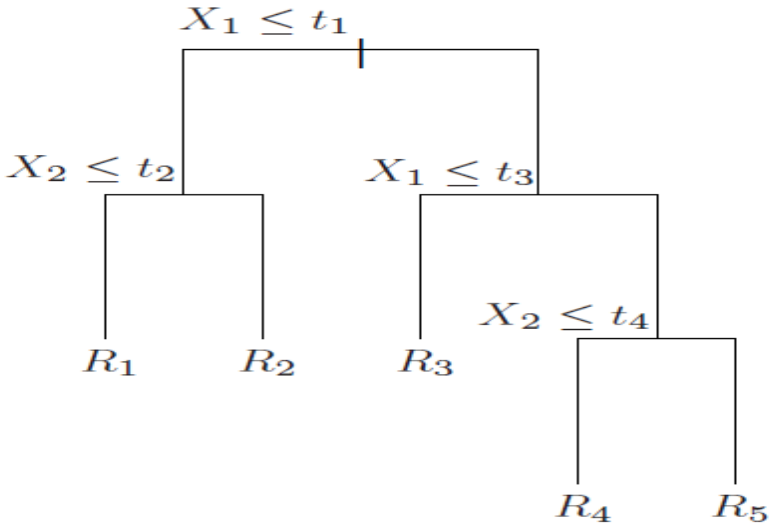

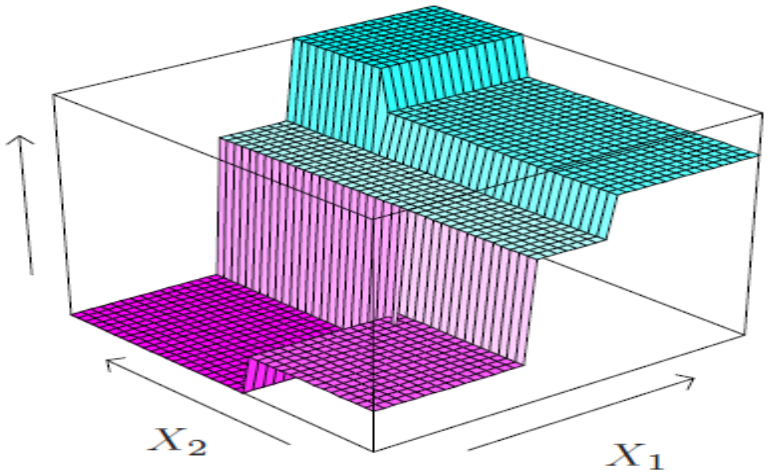

Machine Learning

- Additive Models: combine multiple predictions

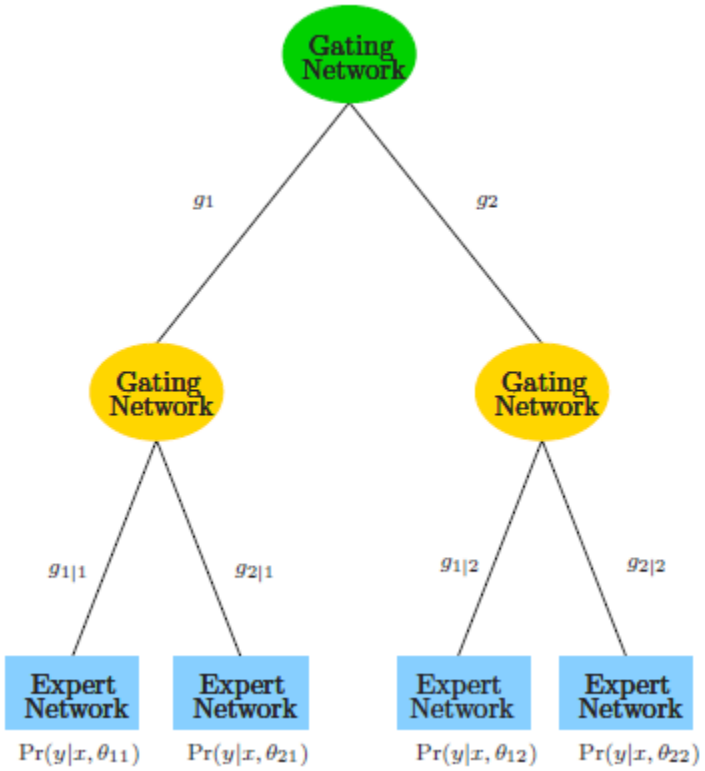

Machine Learning

- Support Vector Machines (SVM): classify new data in complex situations

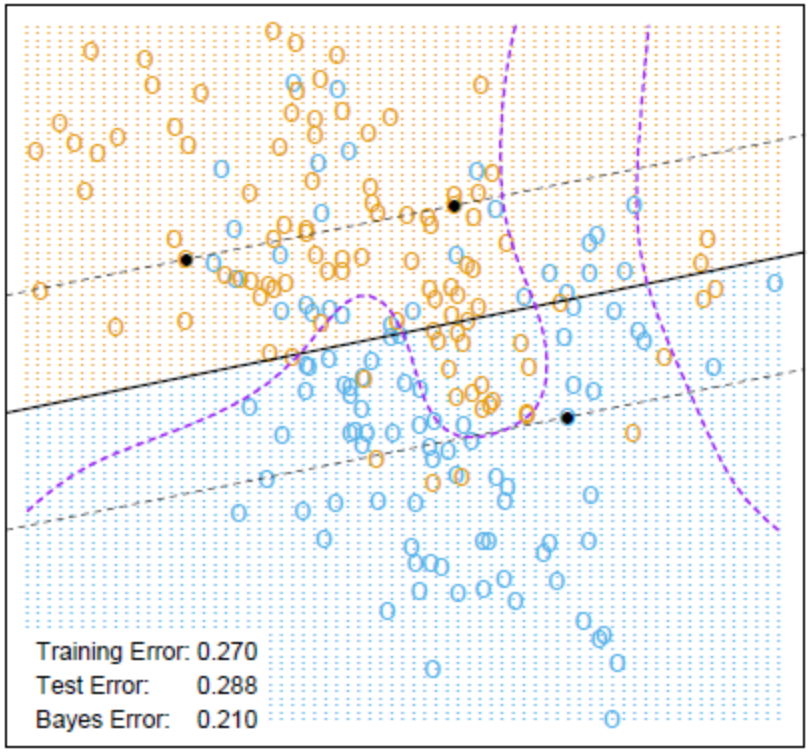

$C = 10000$

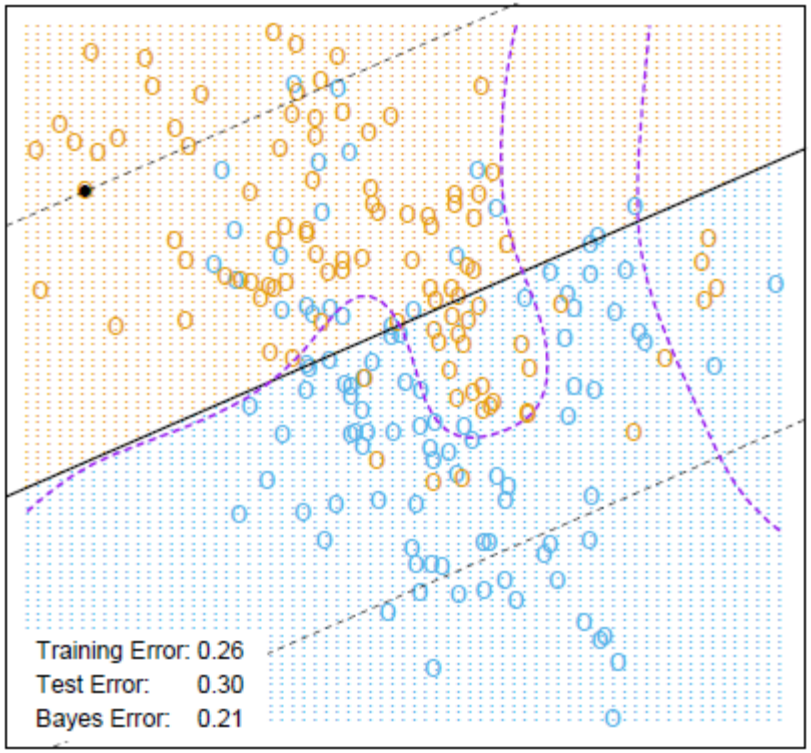

$C = 0.01$

Machine Learning

- Bayesian Classifiers
  - Based on Bayes Theorem
  - Conditional probability based on prior distributions
  - Basis of Kalman Filtering
    - Expected value

$$\hat{f}(\mathbf{x}_{\text{new}}) = \sum_{\ell=1}^{L} w_\ell \mathrm{E}(Y_{\text{new}} | \mathbf{x}_{\text{new}}, \hat{\theta}_\ell)$$

Deep Learning

- Deep Learning (DL) is another method of ML
  - Based on Neural Networks
    - Really the same thing
    - Deep Learning is just a term used to differentiate newer systems from early ones because of the “depth” of the networks commonly used
  - DL is very useful in image classification and understanding and other signal processing tasks
  - DL has much higher processing requirements than other ML techniques

Deep Learning

- DL algorithms and structures can be very complicated to develop
- Major considerations include:
  - Depth of the network
  - Shape of the network
  - Node response function
  - Training functions and approaches
    - Backpropagation
    - Gradient Descent

Deep Learning

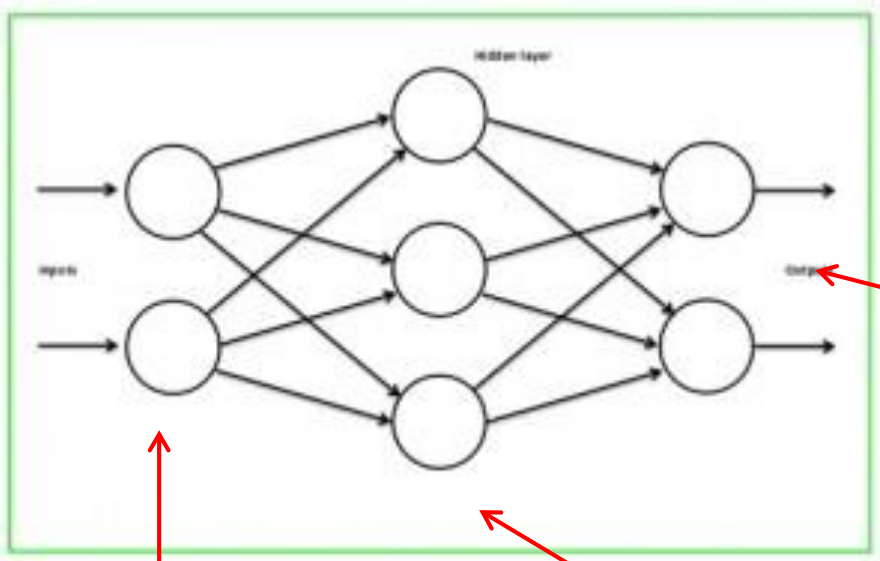

Output Layer:
As many as there are items you want to identify

Input Layer: as many as there are elements in your signal

Hidden Layer: Can have multiple

Deep Learning

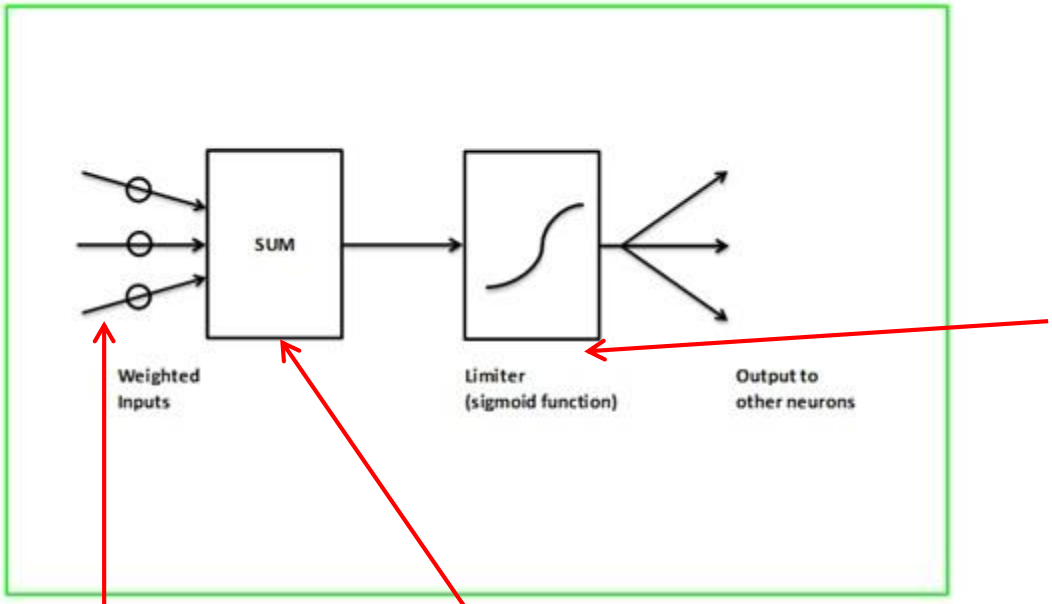

A sigmoid function is applied to the sum. This limits the output to a value between -1 and 1. Signal is then sent to all connected neurons in the next layer.

Weighted inputs are added together

All inputs are multiplied by their weight

Deep Learning

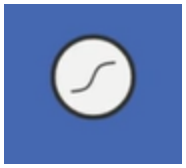

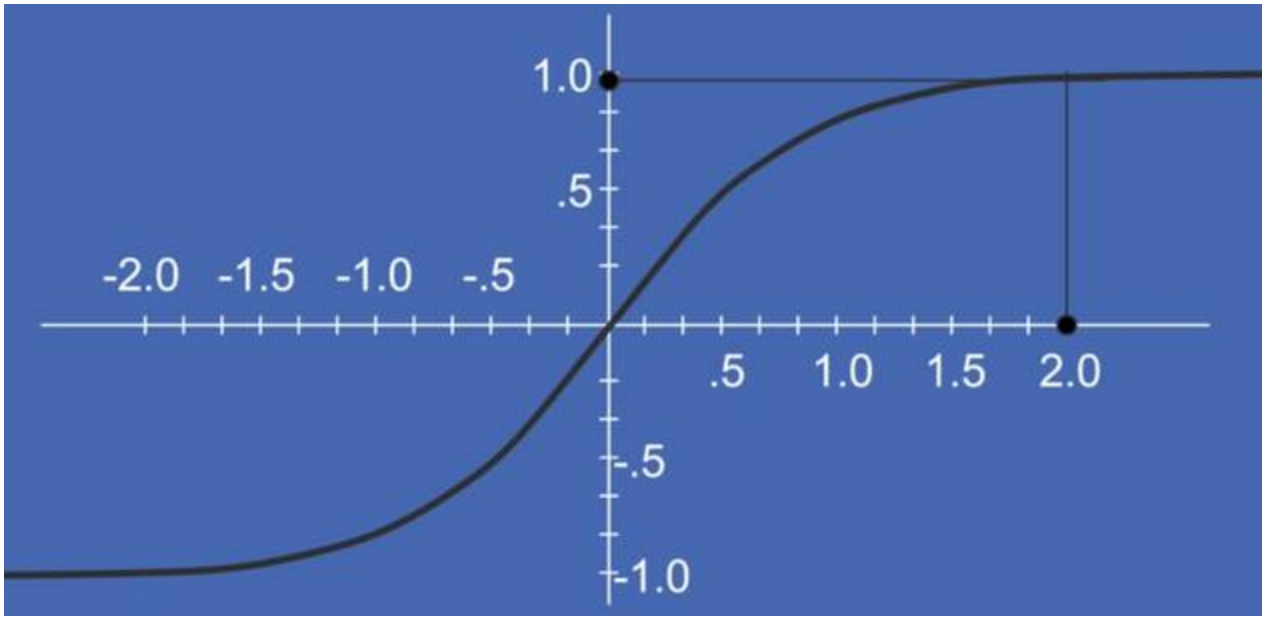

Sigmoid function: ensures that the output is between -1 and 1. x axis is weighted input. y axis is output.

Deep Learning

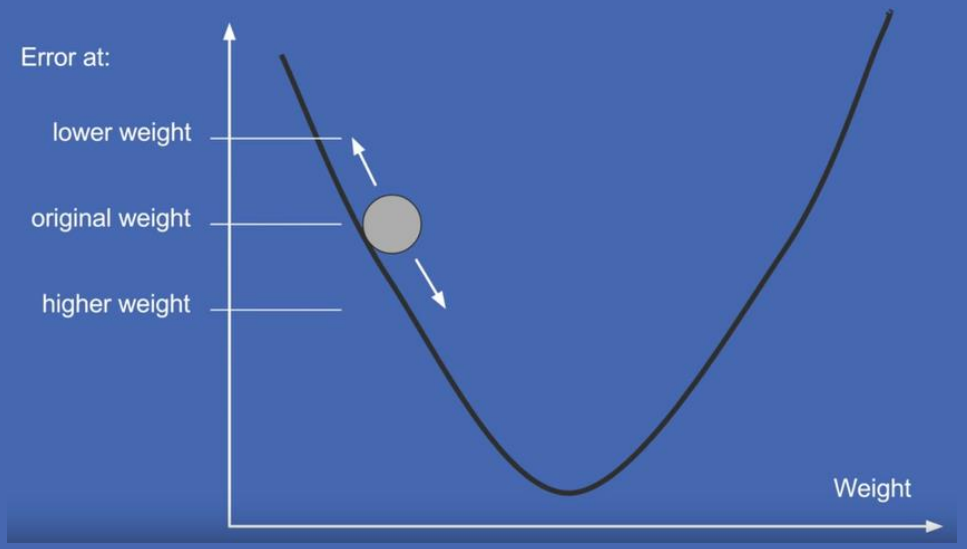

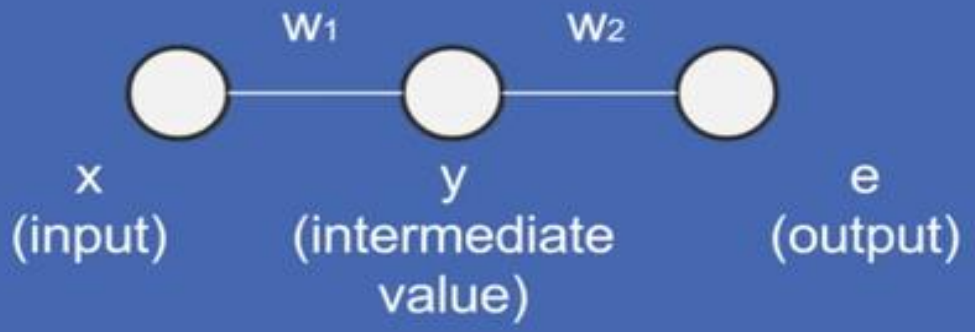

$$y = x * w_1$$
$$\frac{\partial y}{\partial w_1} = x$$

$$e = y * w_2$$
$$\frac{\partial e}{\partial y} = w_2$$

$$e = x * w_1 * w_2$$
$$\frac{\partial e}{\partial w_1} = x * w_2$$

Deep Learning

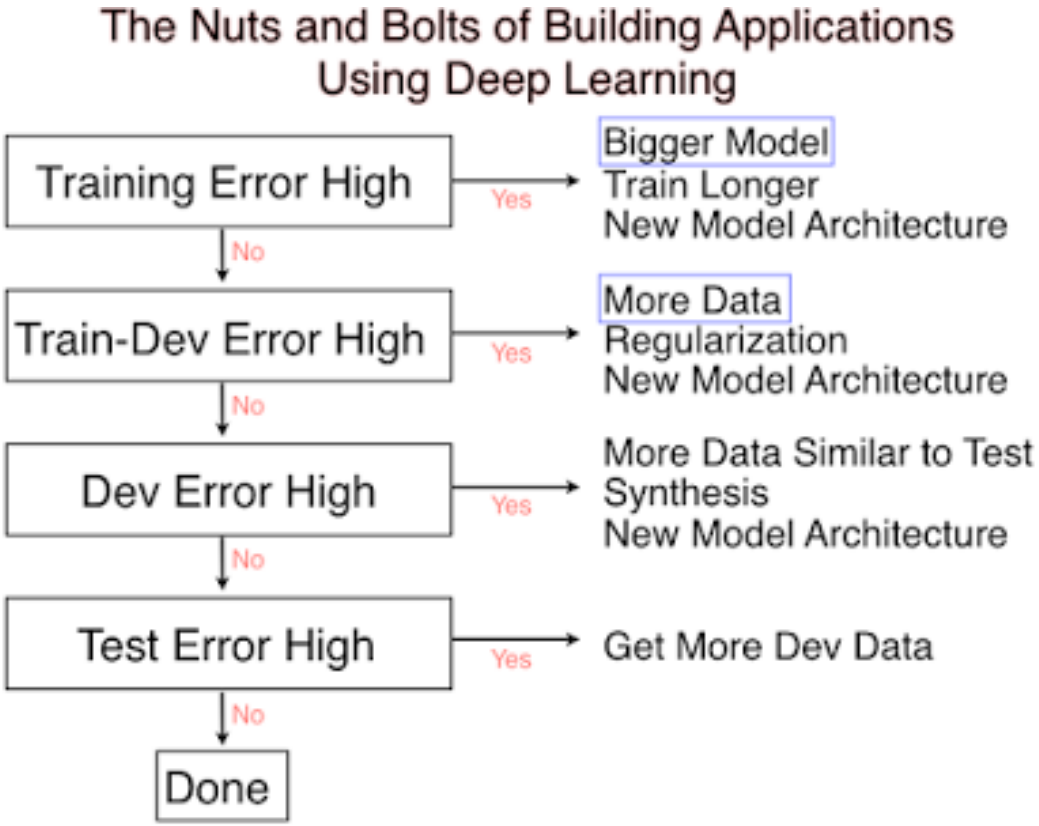

Deep Learning

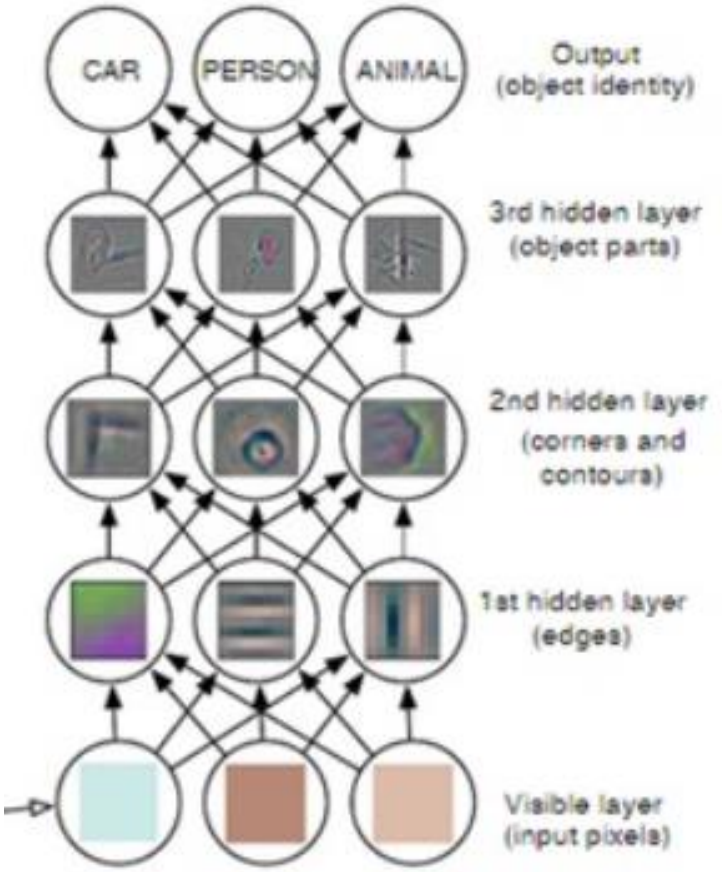

Conclusion/Next Class

- Today we have looked at a sampling of the types of algorithms that might be used in an embedded application
- ML and DL algorithms are based in statistical concepts
- These algorithms are a menu of choices that you can use/must select to match your application and processing environment
- Tomorrow we will look at that processing environment