Machine Learning with MATLAB

Debbi Cohen
RPI Account Manager

Adam Sifounakis
Application Engineer

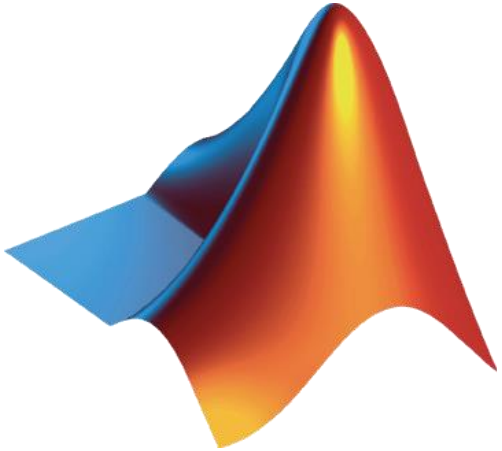

March 30, 2016

Goals

- Overview of machine learning
- Machine learning models & techniques available in MATLAB
- Using MATLAB to streamline your machine learning workflow

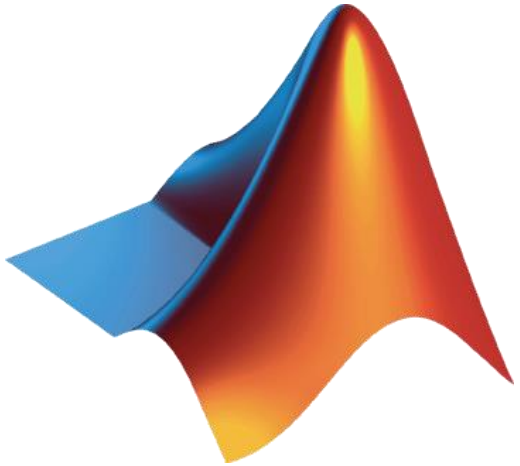

Machine Learning

Characteristics and Examples

- Characteristics
  - Lots of data (many variables)
  - System too complex to know the governing equation
    (e.g., black-box modeling)

- Examples
  - Pattern recognition *(speech, images)*
  - Financial algorithms *(credit scoring, algo trading)*
  - Energy forecasting *(load, price)*
  - Biology *(tumor detection, drug discovery)*

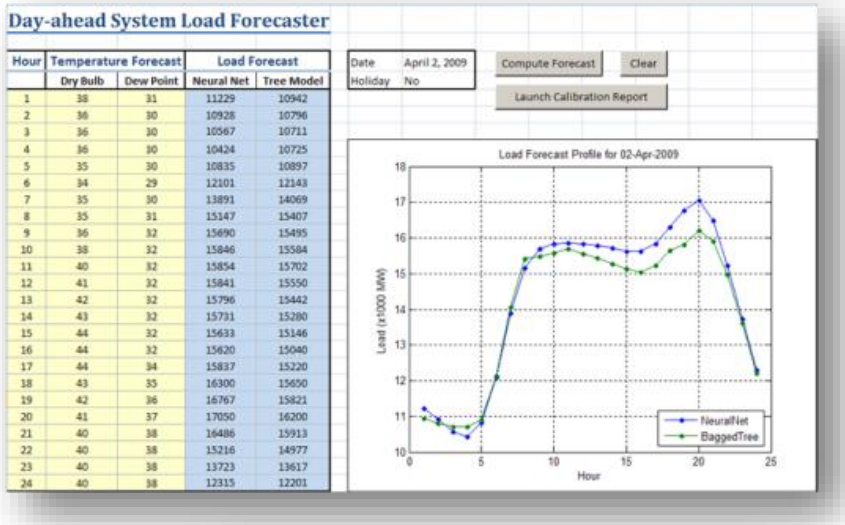

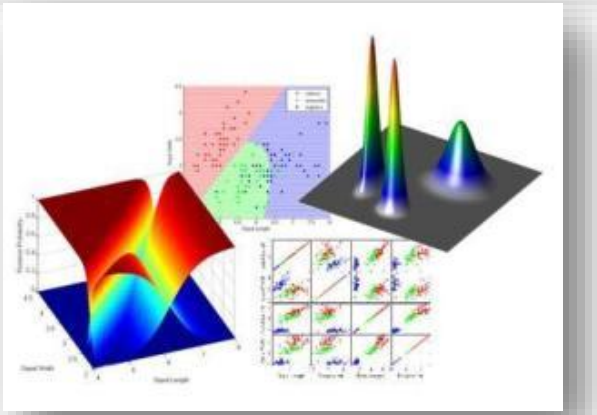

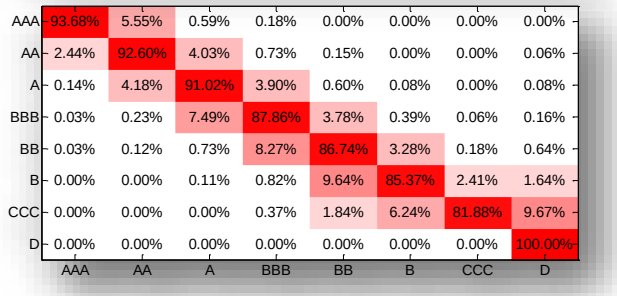

Challenges – Machine Learning

- Significant technical expertise required
- No “one size fits all” solution
- Time required to conduct the analysis
- Locked into Black Box solutions

Overview – Machine Learning

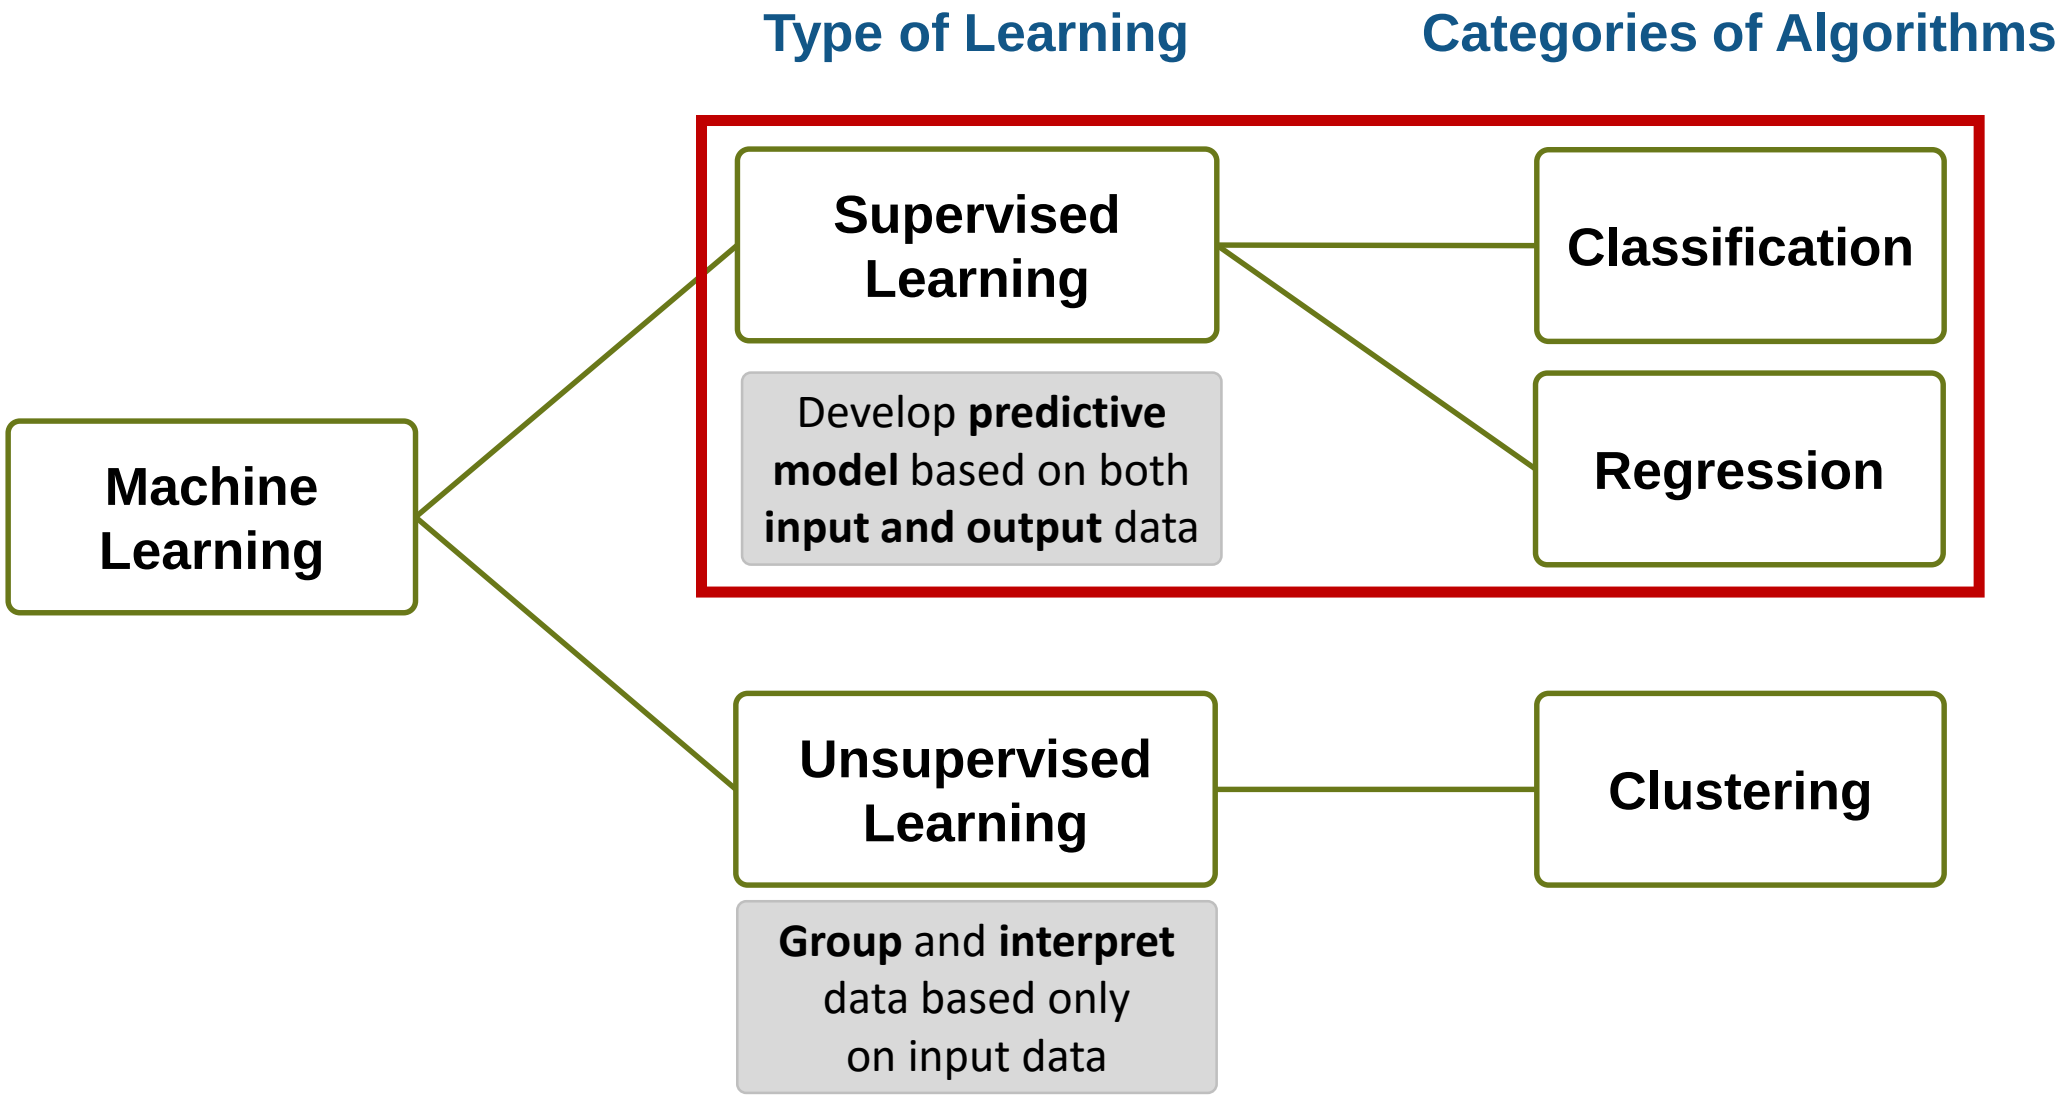

Supervised Learning

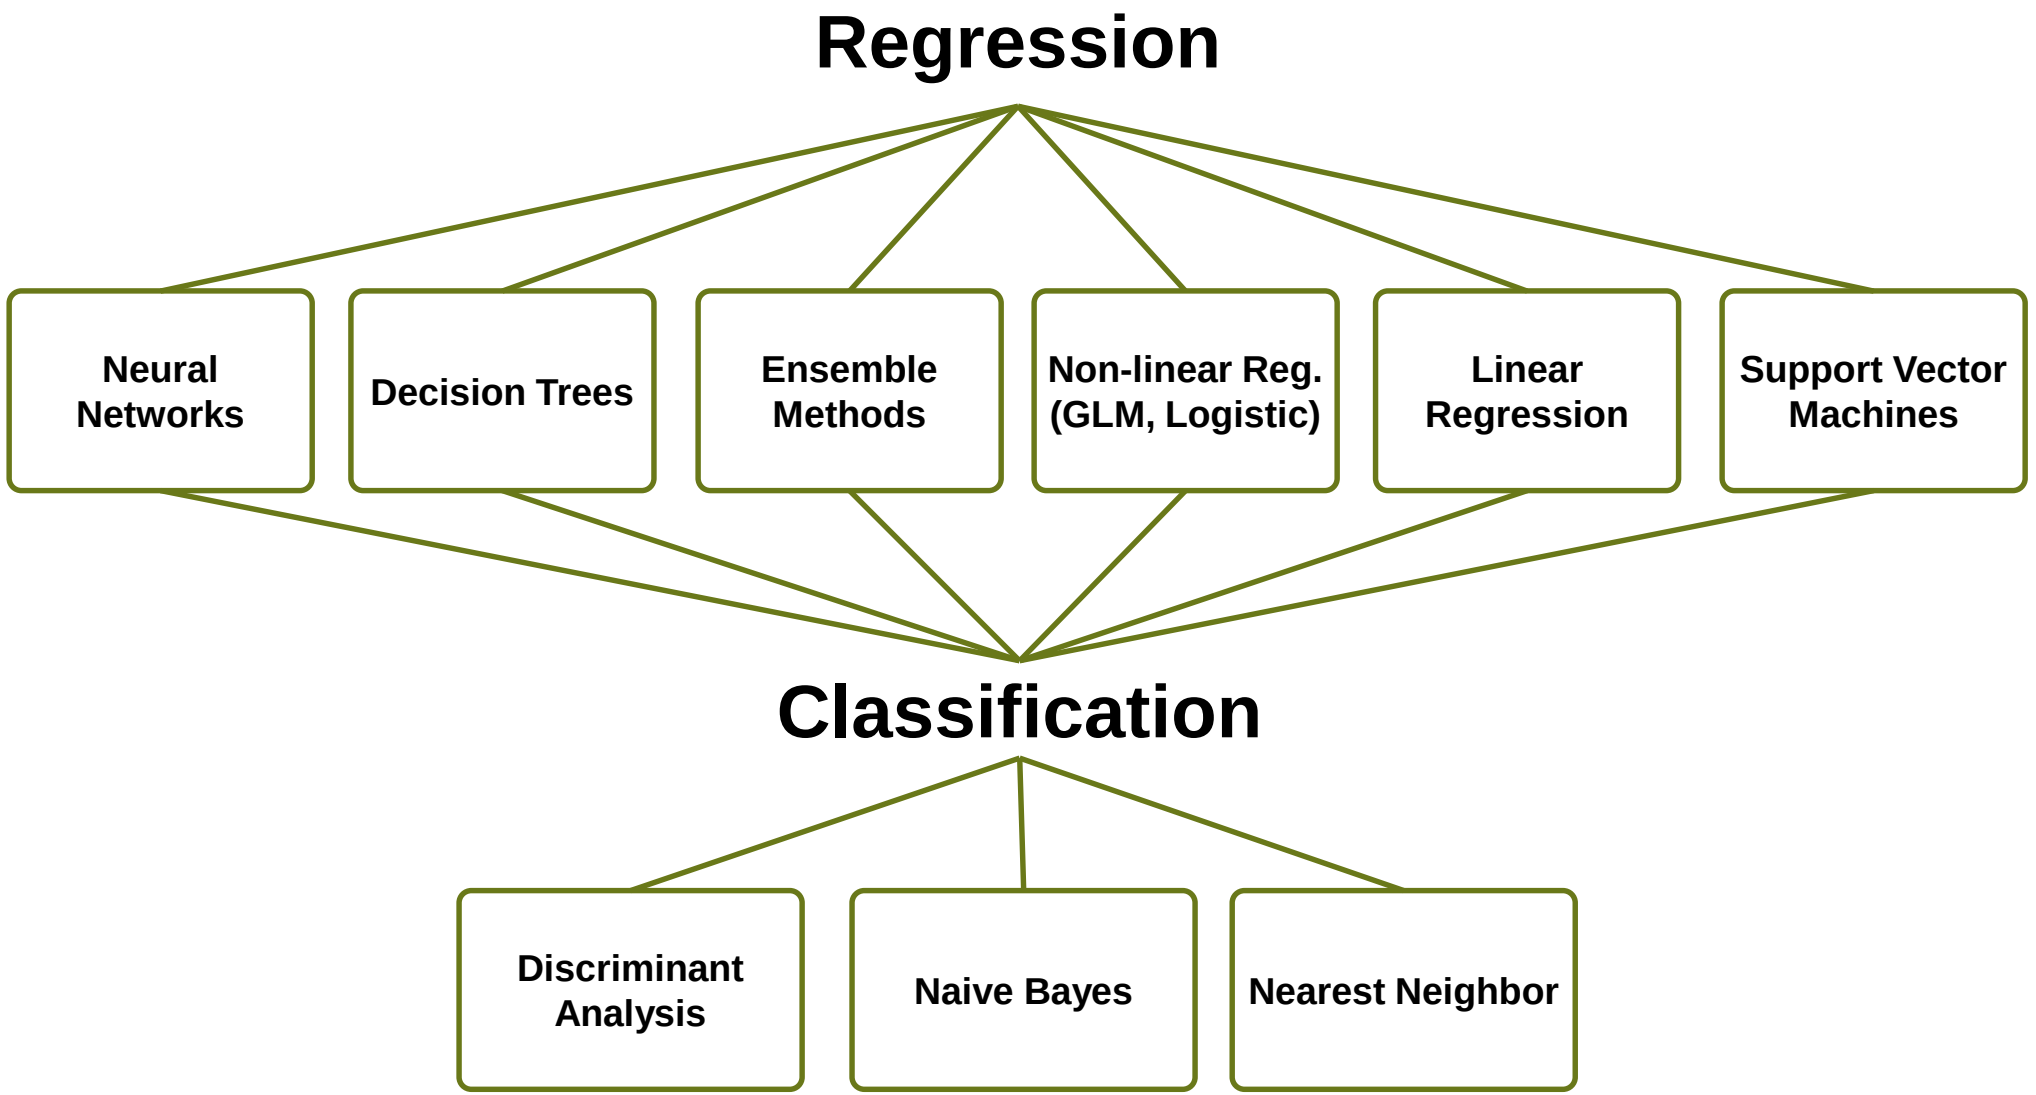

Demo: Robot Failures Identification

- Challenge:
  - Find natural groupings among large number of predictors
- Approach:
  - Train a classifier using different models
  - Measure accuracy and compare models
  - Reduce model complexity
  - Use classifier for prediction

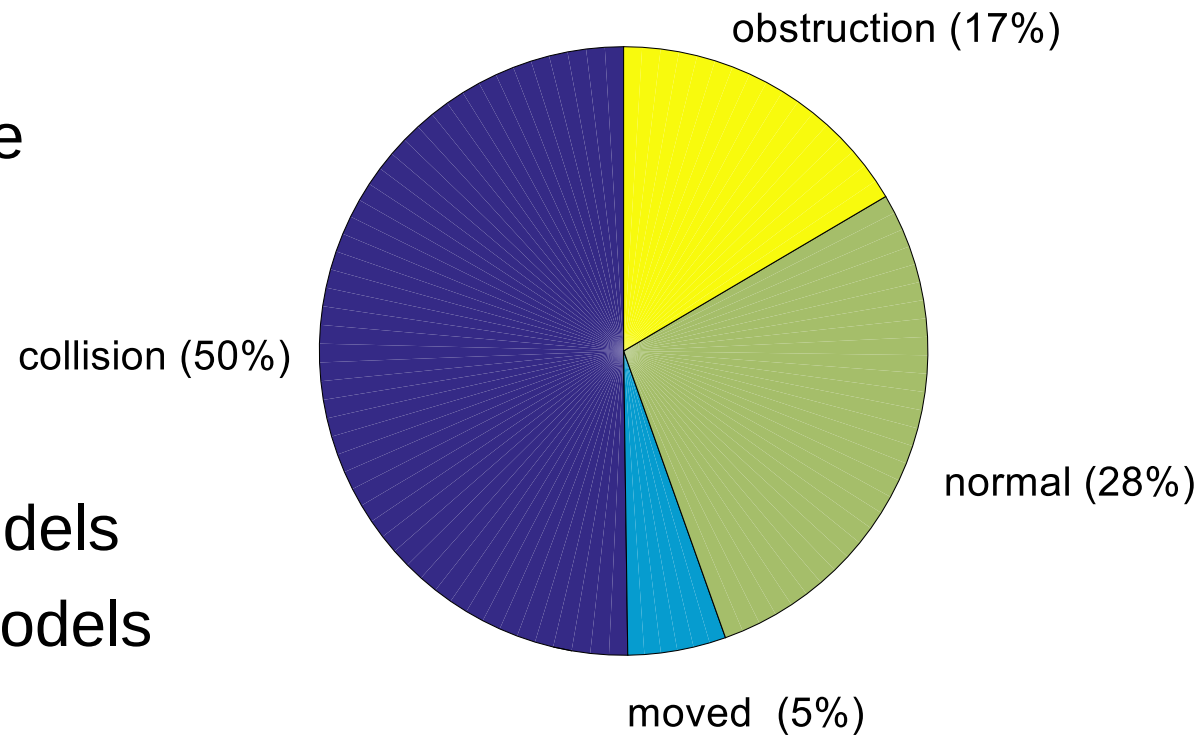

Decision Trees – What are they?

- Branching tree structure:

OR

- Nested if statements:

```
Decision tree for classification
1  if x1<1 then node 2 else node 3
2  class = 2
3  if x2<2 then node 4 else node 5
4  class = 1
5  class = 2
```

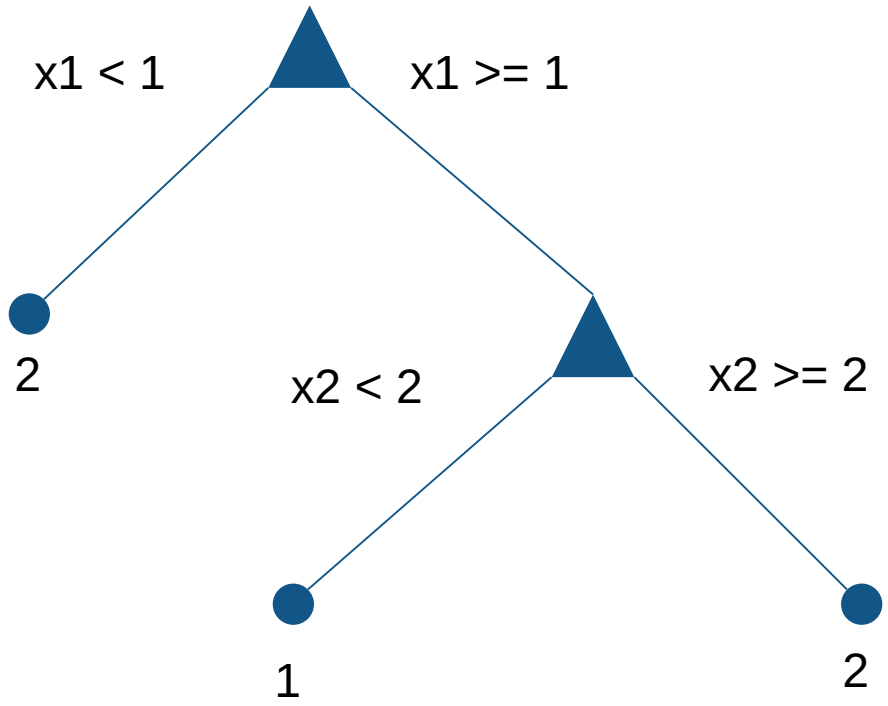

Decision Trees – How do they work?

Design a classifier for the following situation:

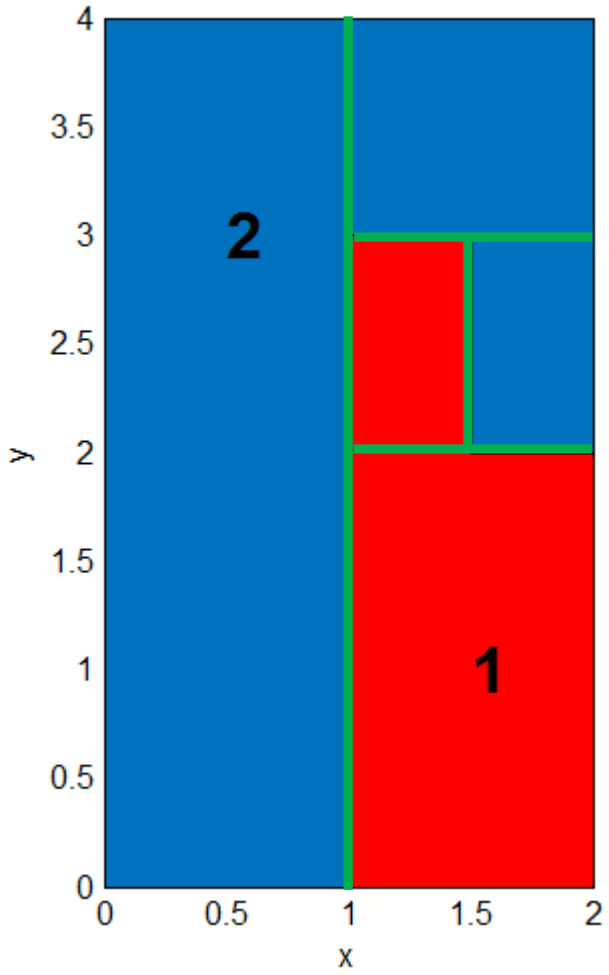

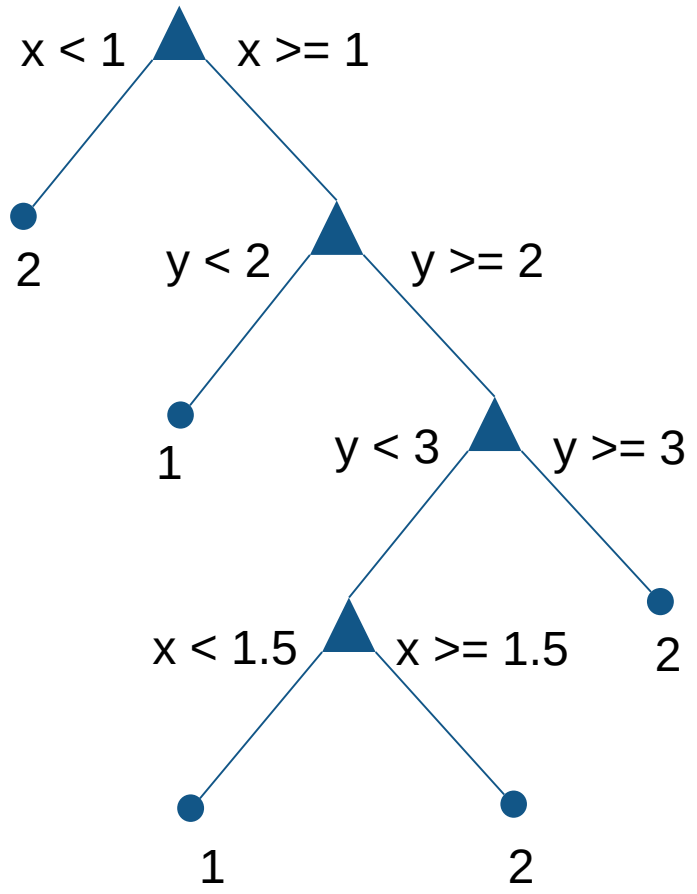

Machine Learning Workflow

Train: Iterate until you find the best model

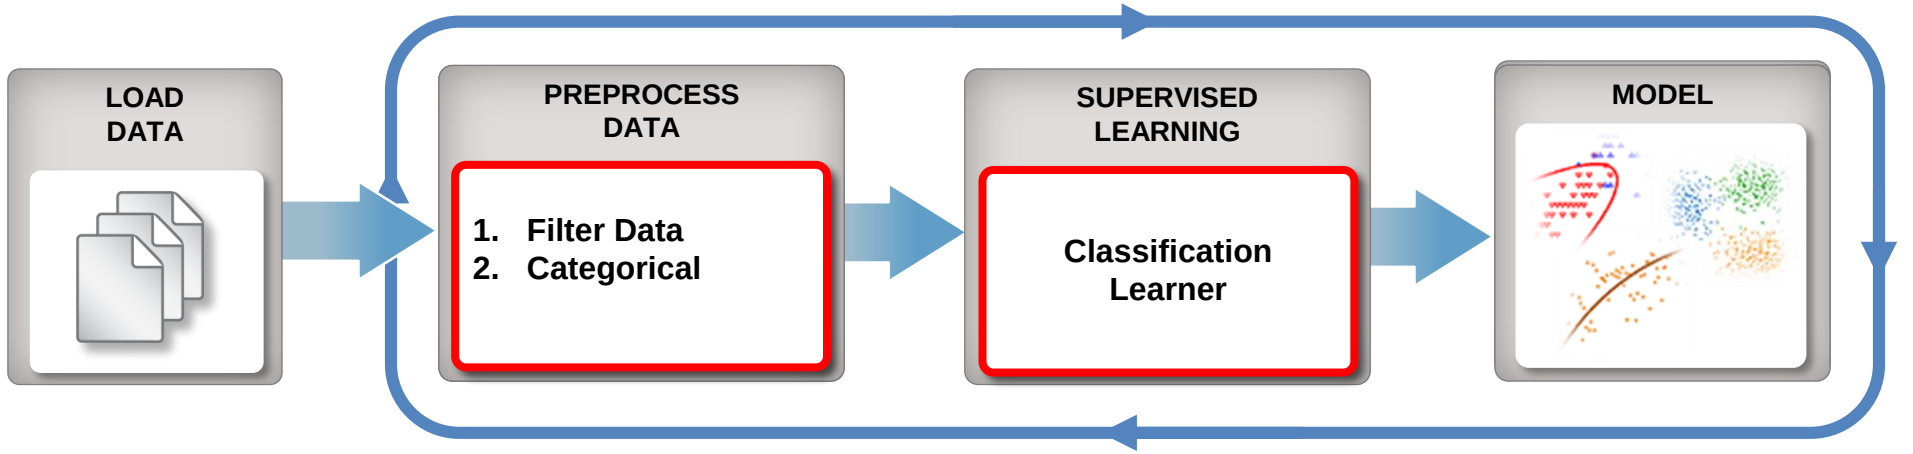

Predict: Integrate trained models into applications

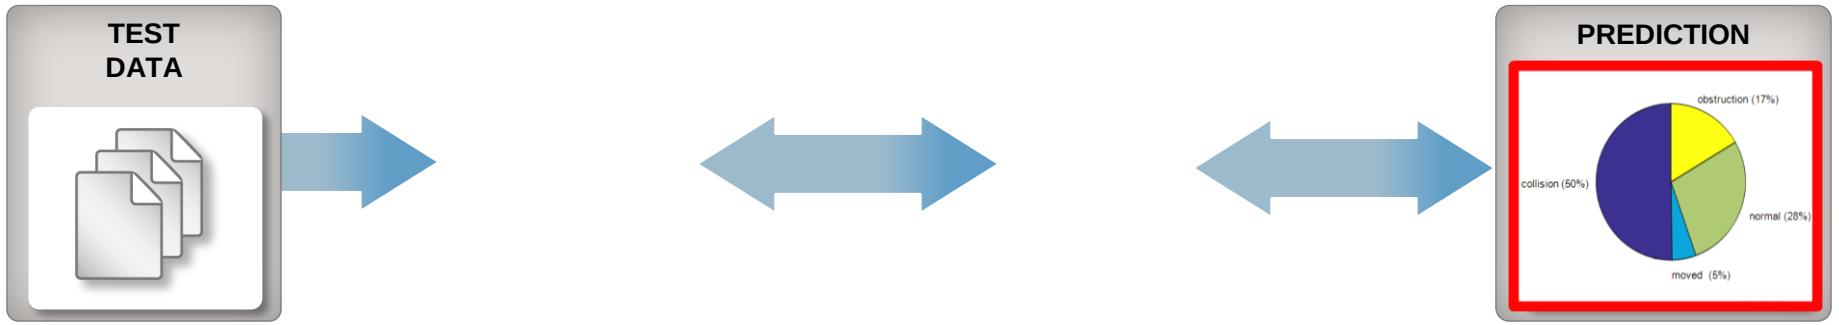

Demo: Predicting Fuel Economy

- Goal:
  - Study the relationships between fuel economy, horsepower, and type of vehicle
- Approach:
  - Access data from Excel
  - Create a predictive model
  - Improve the model accuracy

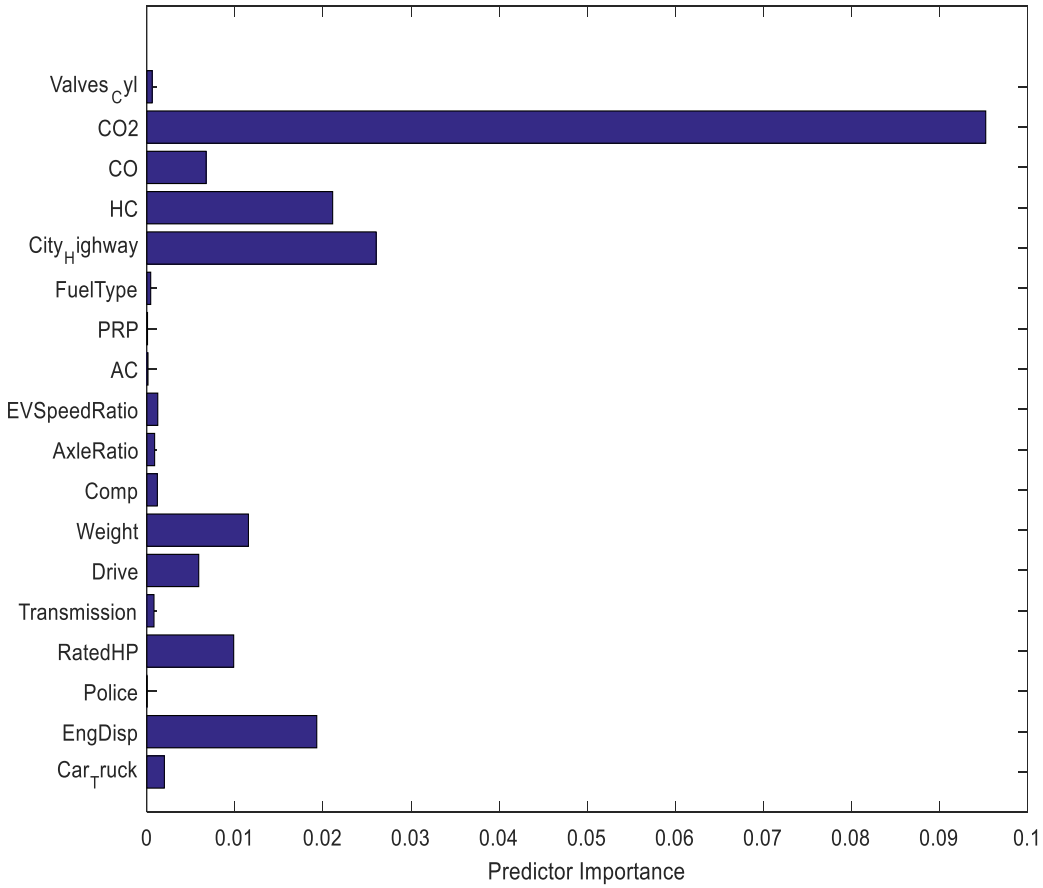

Machine Learning Workflow

Train: Iterate until you find the best model

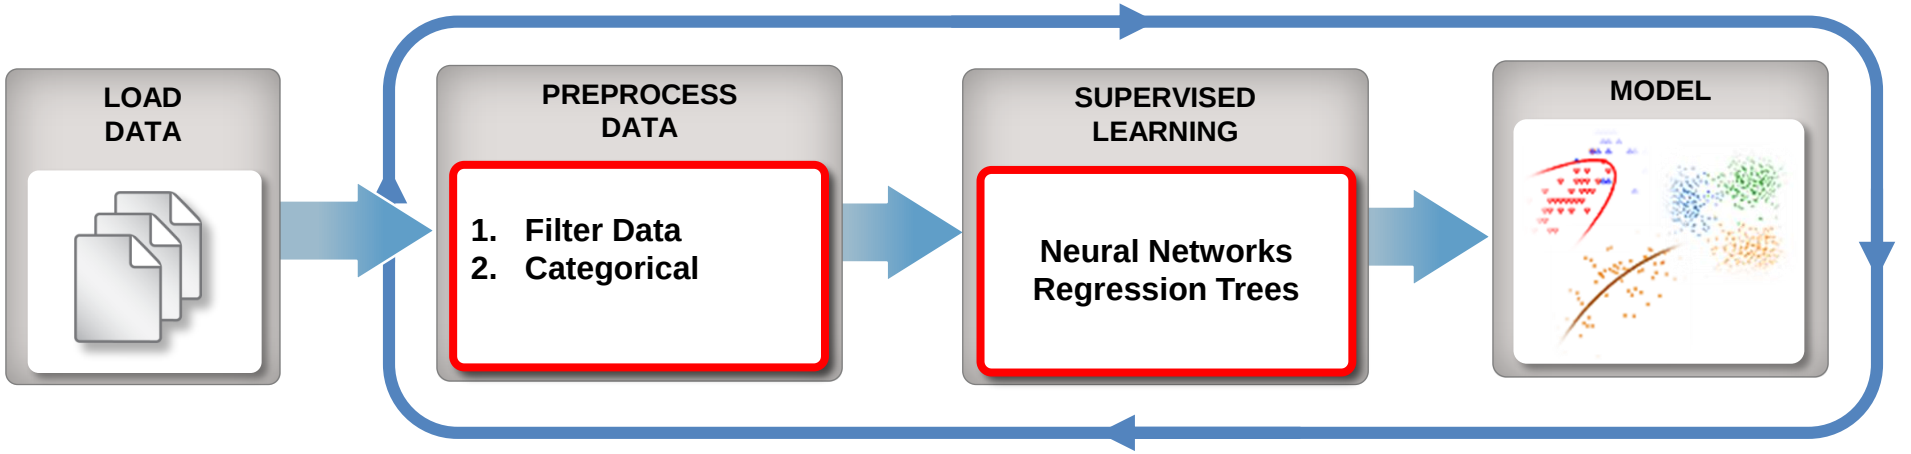

Predict: Integrate trained models into applications

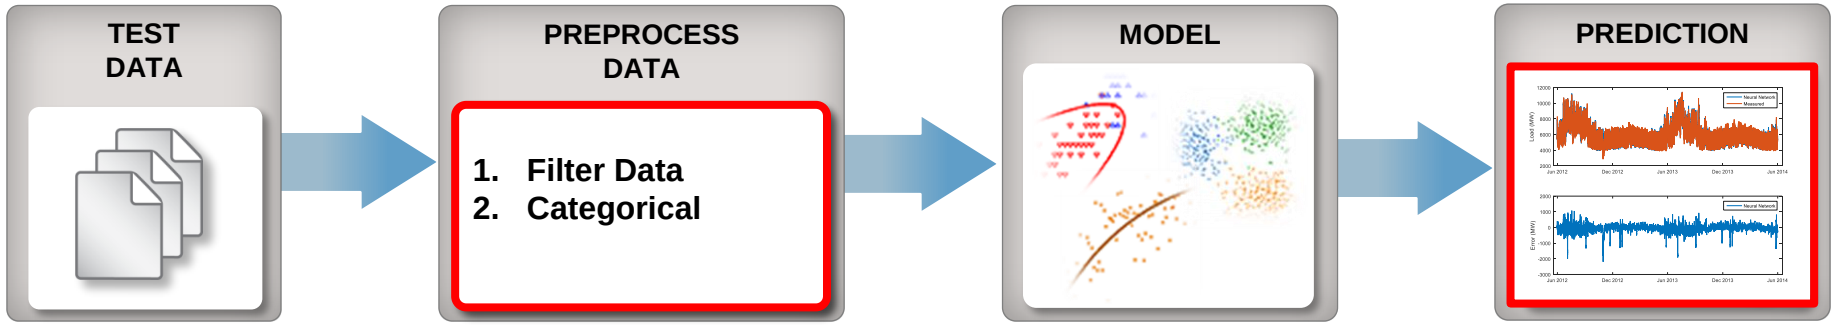

Overview – Machine Learning

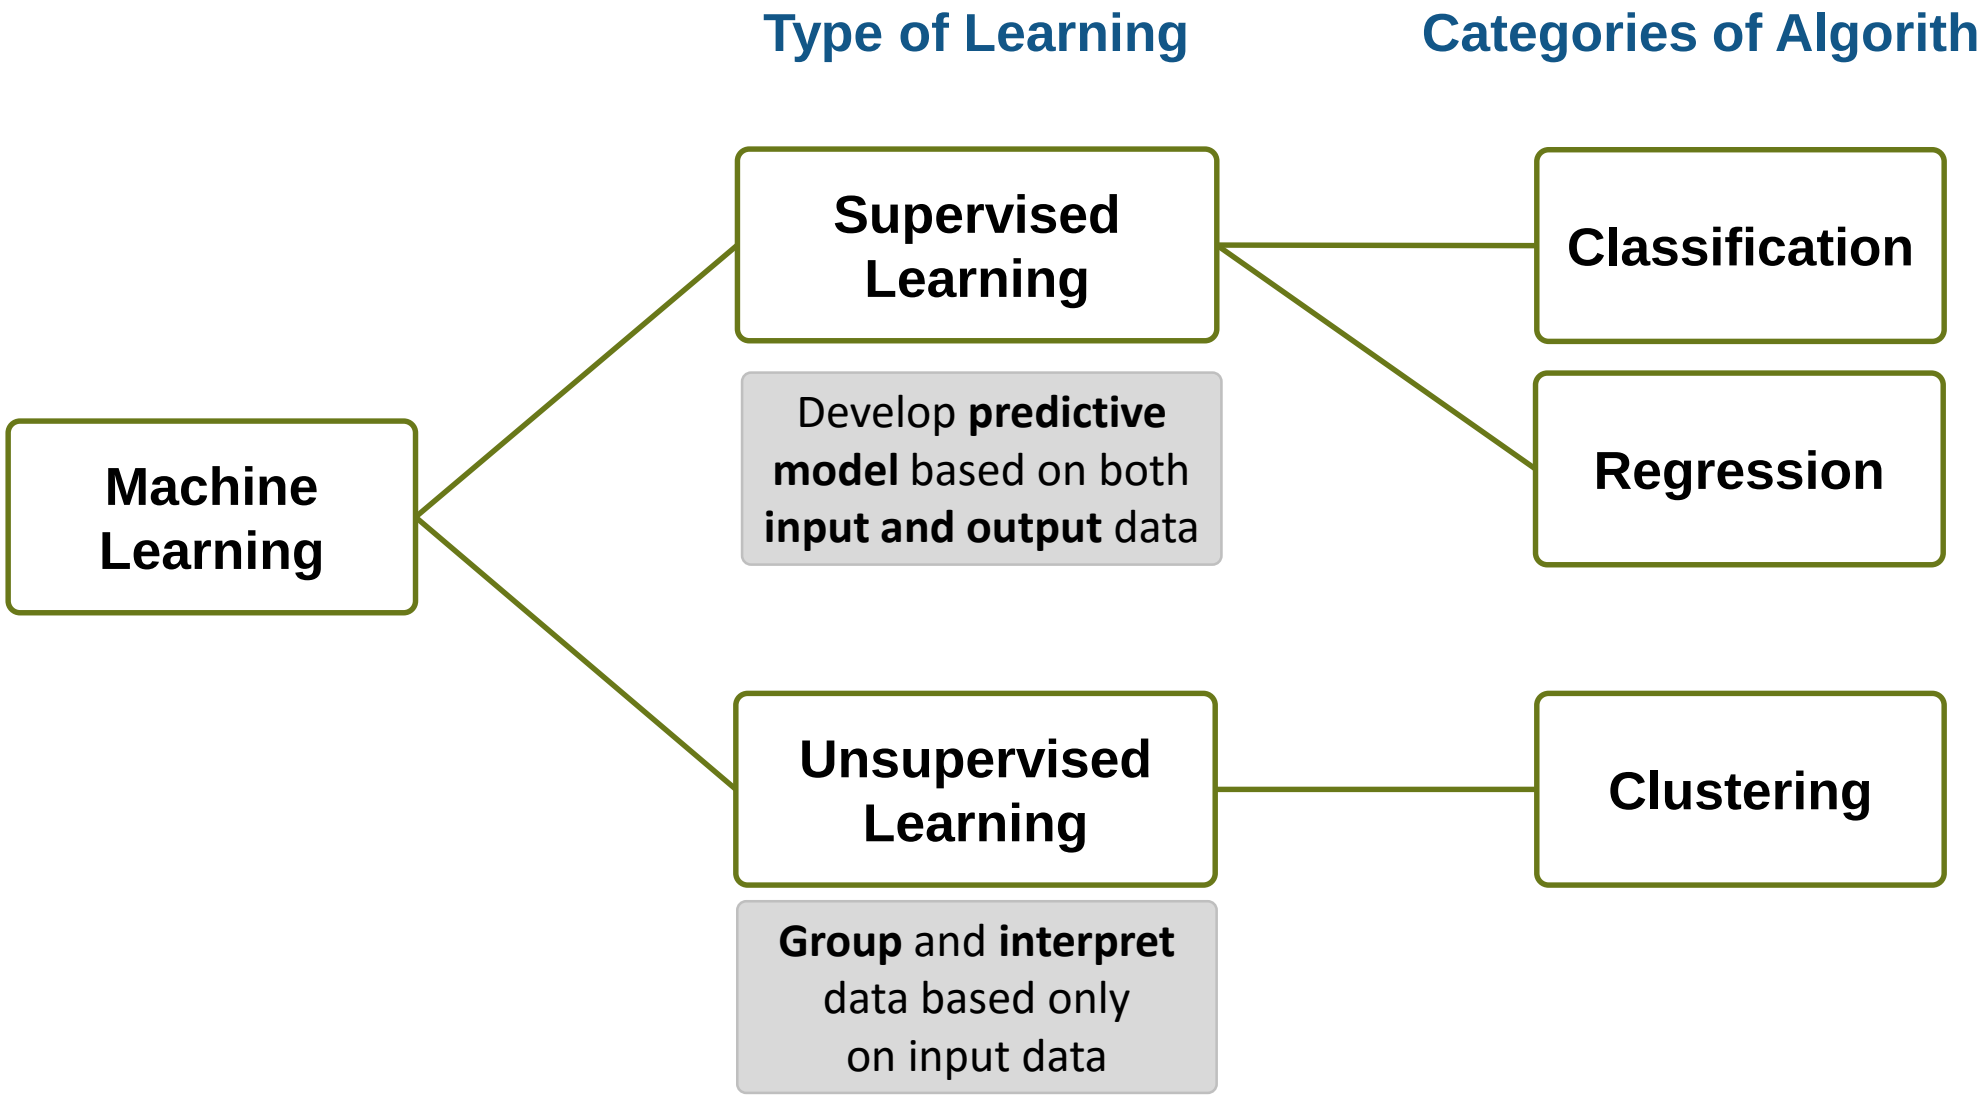

Unsupervised Learning

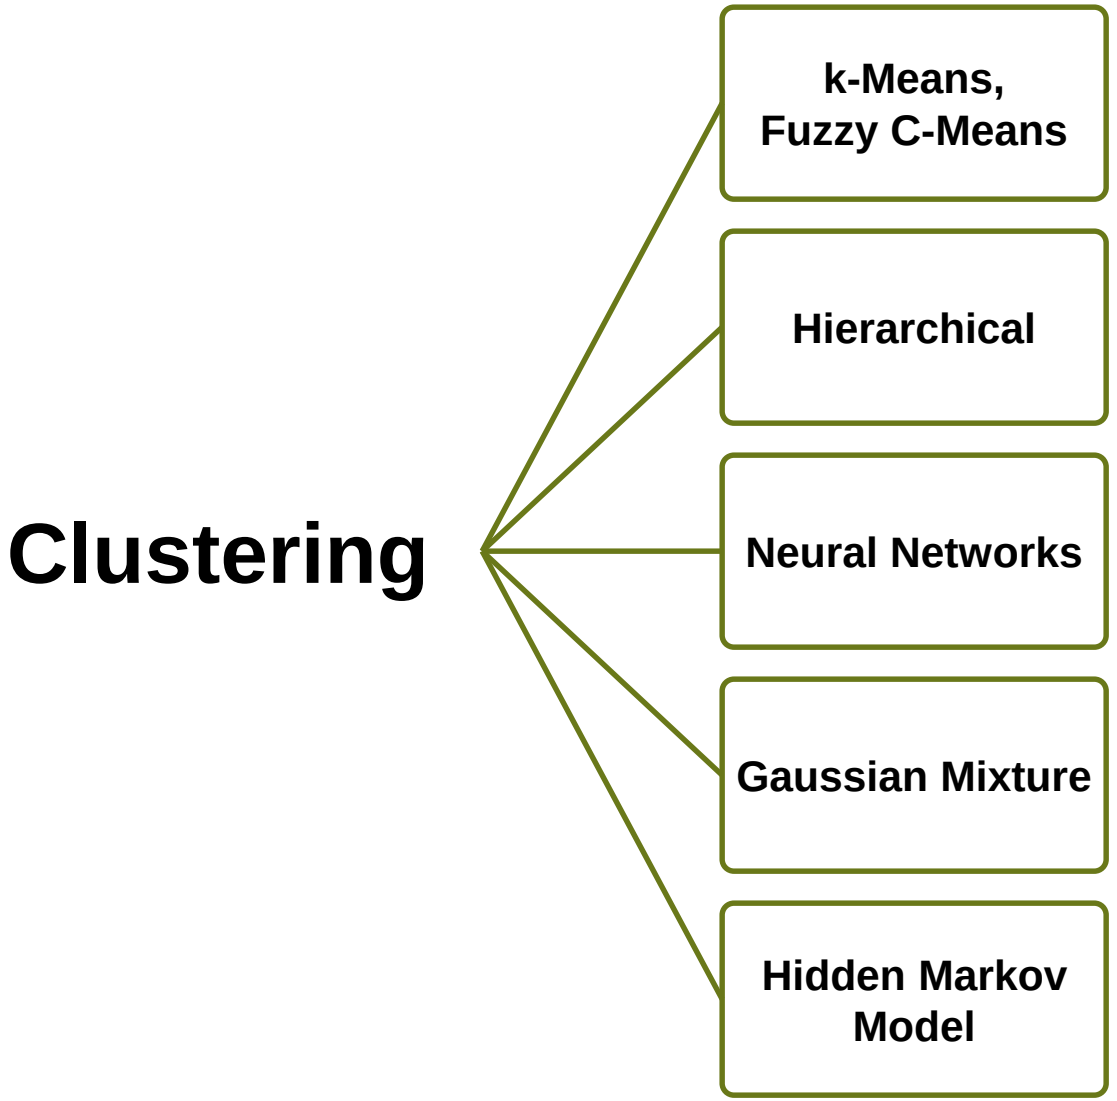

Example: Human Activity Recognition

- Challenge:
  - Find natural groupings among large number of predictors
  - Build predictive model for classifying new data points
- Approach:
  - Reduce dimensionality and visualize structure of data
  - Evaluate different clustering techniques to identify groups of behaviors
  - Determine types of activities for new data points using classification techniques

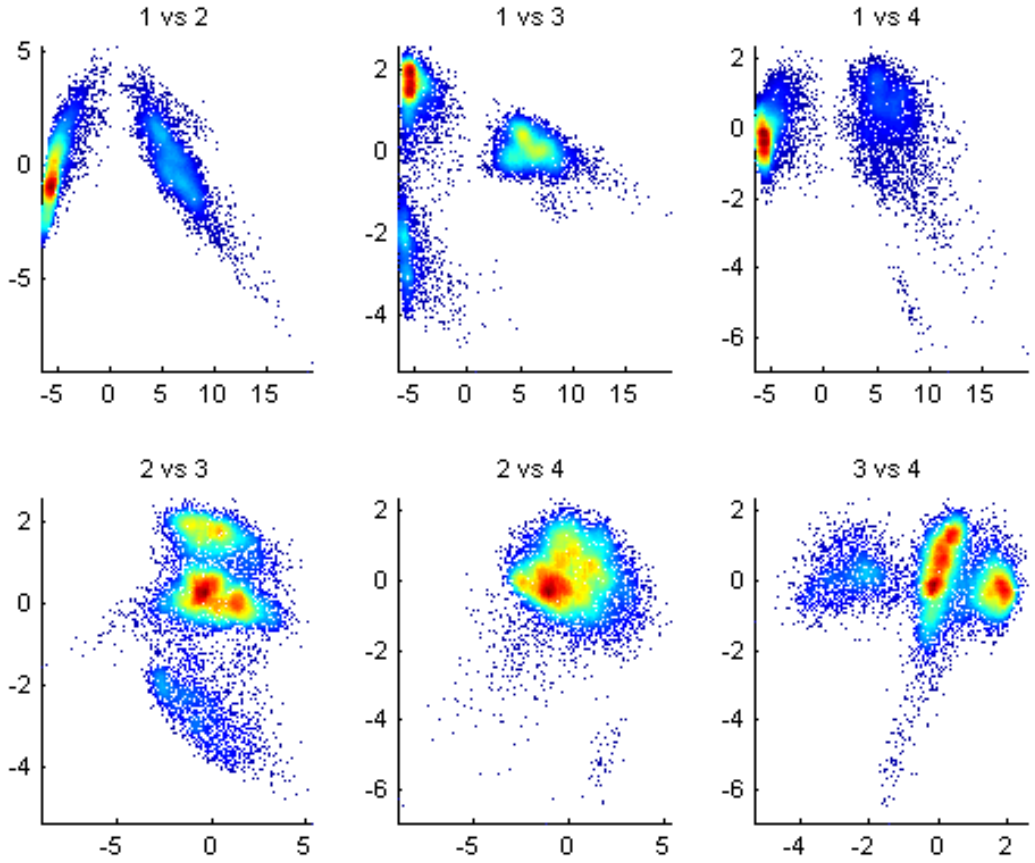

Principal Components Analysis – what is it doing?

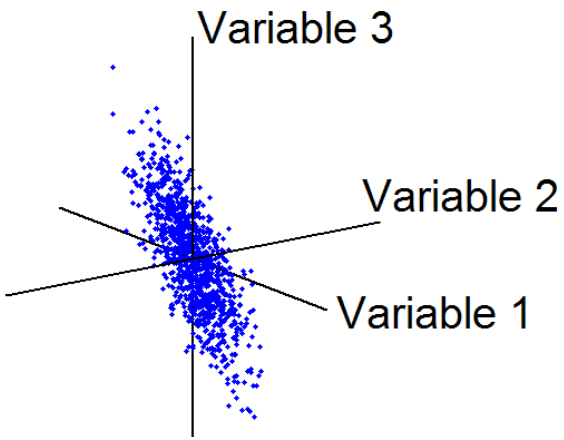

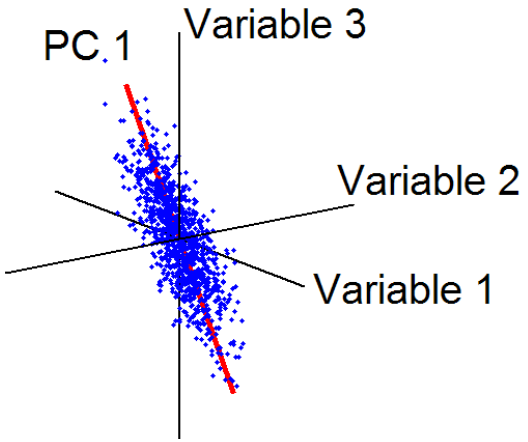

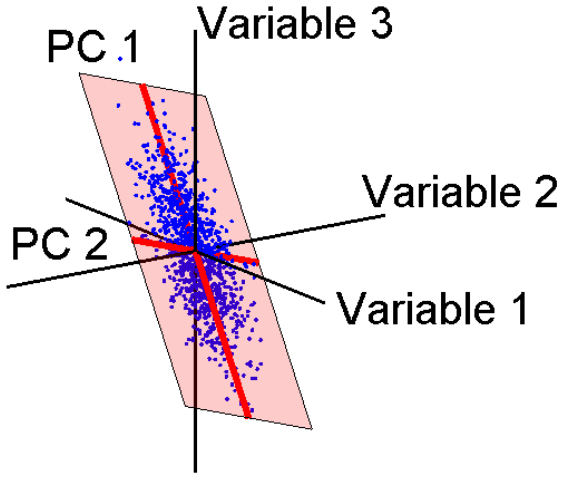

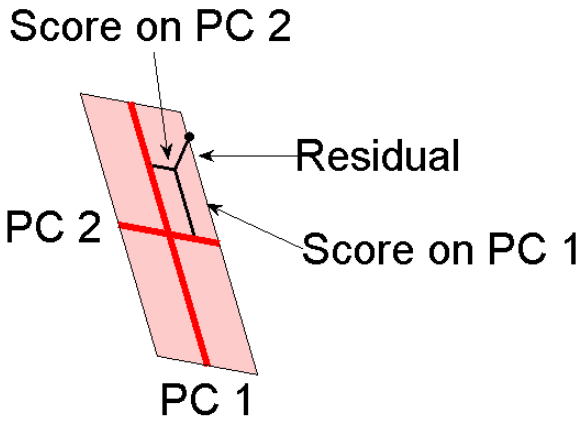

K-Means Cluster Analysis – what is it doing?

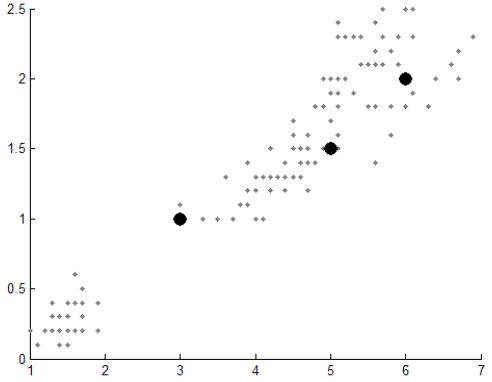

Randomly pick K cluster centroids

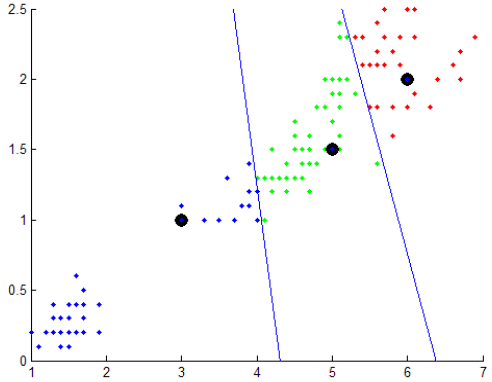

Assign points to the closest centroid

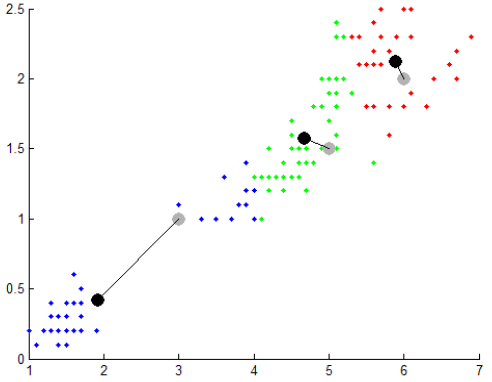

Recalculate positions of cluster centroids

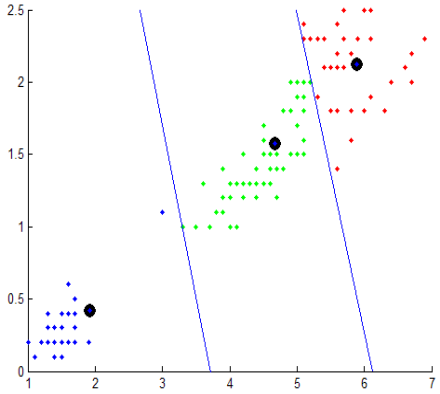

Reassign points to the closest centroid

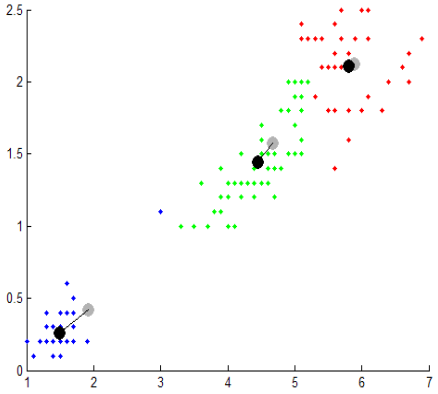

Recalculate positions of cluster centroids

. . . . . .

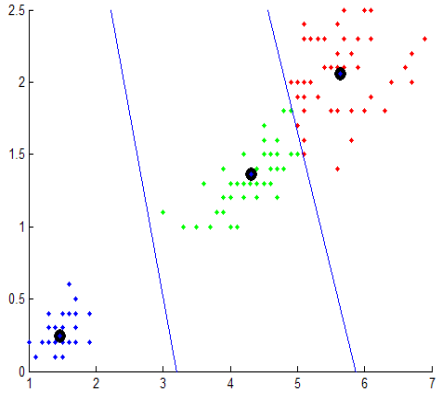

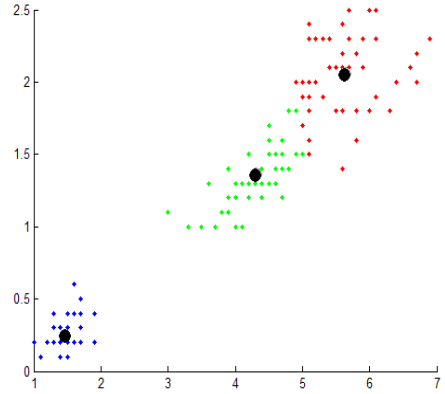

Repeat until centroid positions converge

Machine Learning with MATLAB

- Interactive environment
  - Visual tools for exploratory data analysis
  - Easy to evaluate and choose best algorithm
  - Apps available to help you get started
    (e.g,. neural network tool, curve fitting tool)

- Multiple algorithms to choose from
  - Classification
  - Regression
  - Clustering

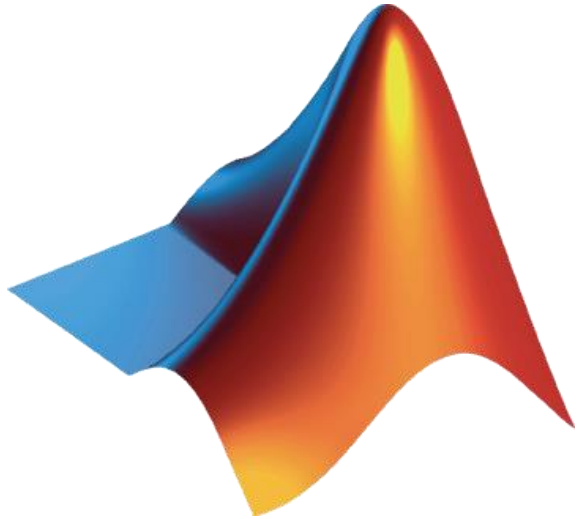

Deployment

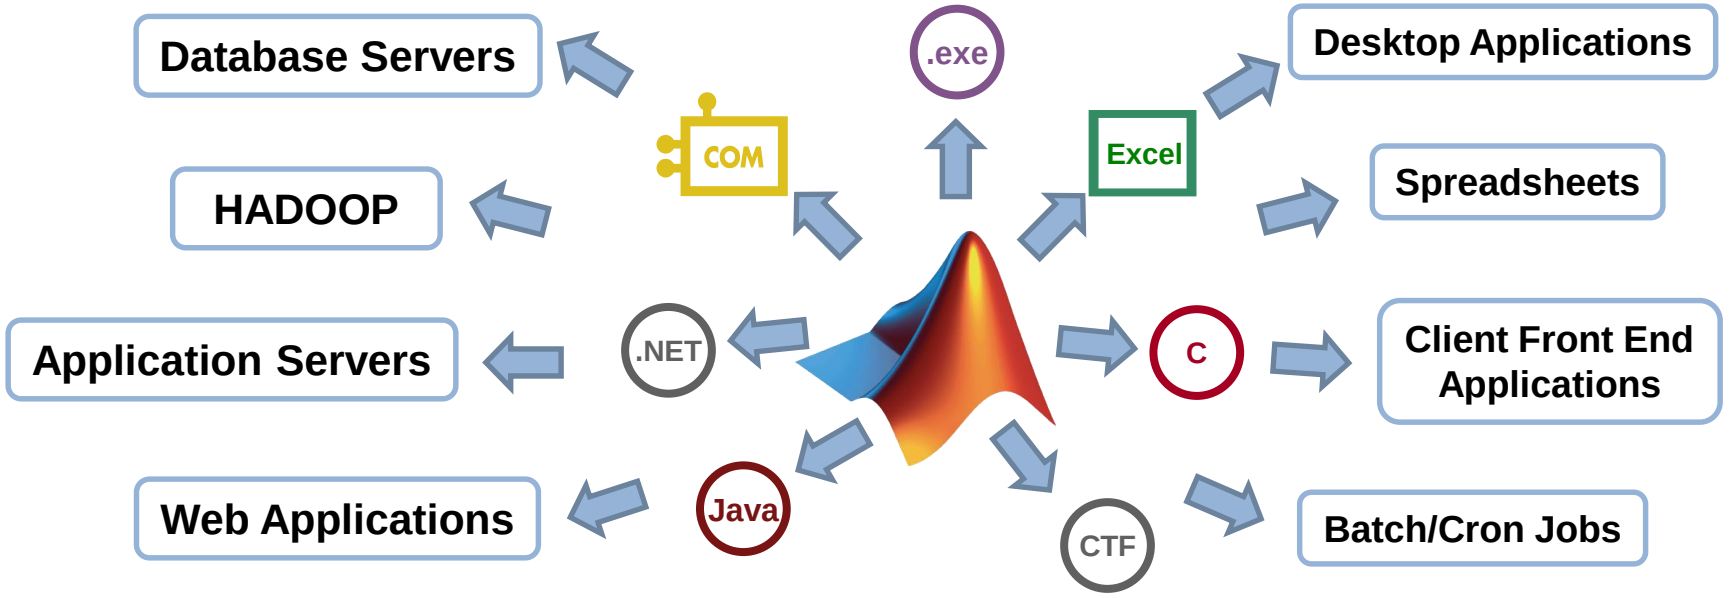

- Royalty-free deployment
- Point-and-click workflow
- Unified process for desktop and server apps

Key Takeaways

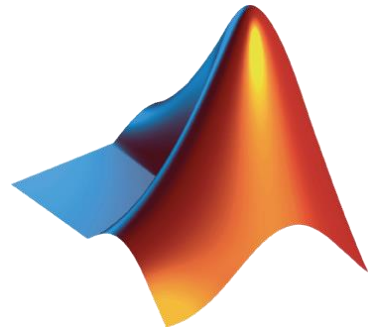

Consider Machine Learning when:

- Hand written rules and equations are too complex
  - *Face recognition, speech recognition, recognizing patterns*
- Rules of a task are constantly changing
  - *Fraud detection from transactions, anomaly in sensor data*
- Nature of the data changes and the program needs to adapt
  - *Automated trading, energy demand forecasting, predicting shopping trends*

“Essentially, all models are wrong, but some are useful.”

– George Box

Learn More: Machine Learning with MATLAB

mathworks.com/machine-learning

Classification Examples


Basket Selection Using Stepwise Regression

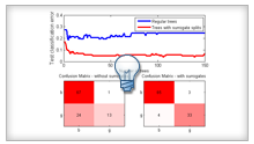
Classification in the Presence of Missing Data

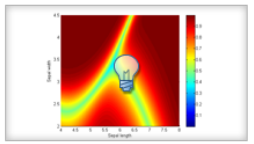
Classification Probability


Digit Classification Using HOG Features

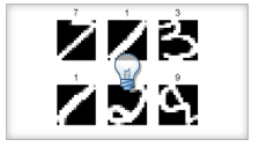
Handwriting Recognition Using Bagged Classification Trees

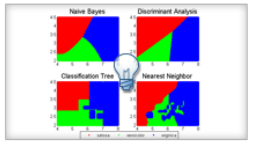
Visualize Decision Surfaces for Different Classifiers

Regression Examples

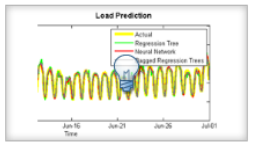
Electricity Load Forecasting

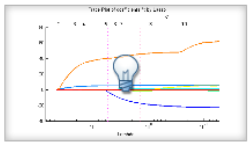
Lasso Regularization

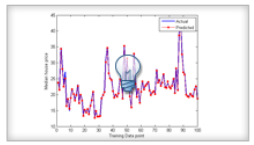
Regression with Boosted Decision Trees

Clustering Examples

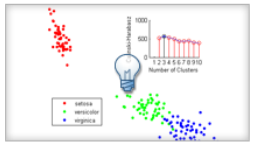
Cluster Evaluation


Cluster Genes Using K-Means and Self-Organizing Maps


Color-Based Segmentation Using K-Means Clustering

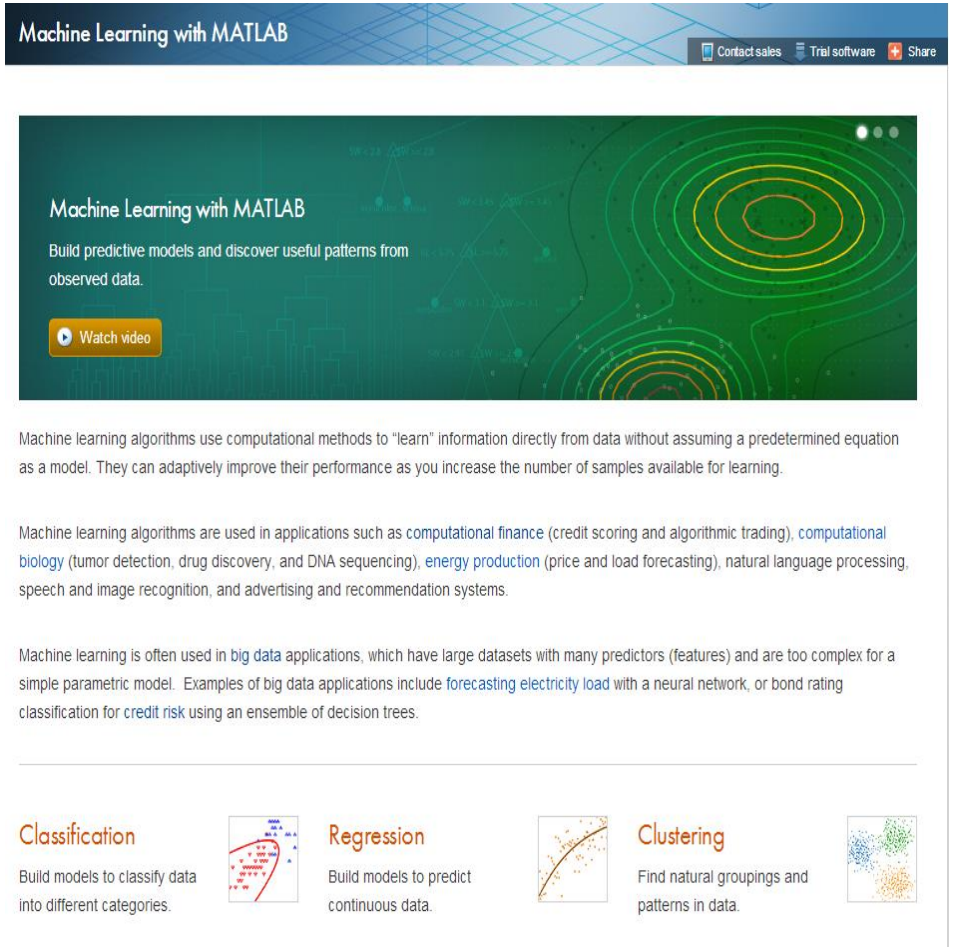

Questions?