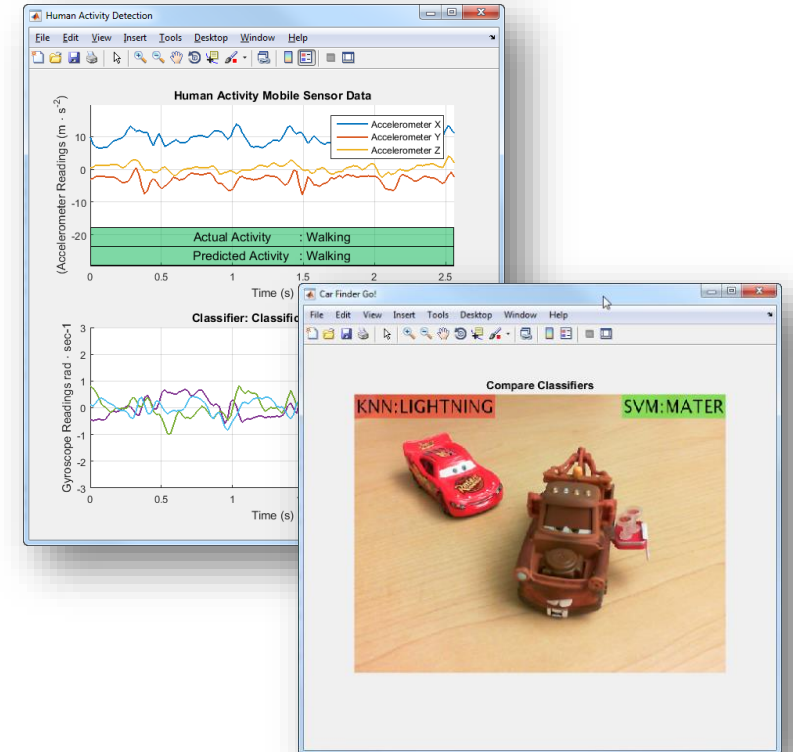


Machine Learning Made Easy

David Willingham
Senior Application Engineer






Agenda

- Machine Learning
 - What is Machine Learning and why do we need it?
 - Common challenges in Machine Learning
- Example 1: Human activity learning using mobile phone data
 - Learning from sensor data
- Example 2: Real-time car identification using images
 - Learning from images
- Summary & Key Takeaways

Machine Learning is Everywhere


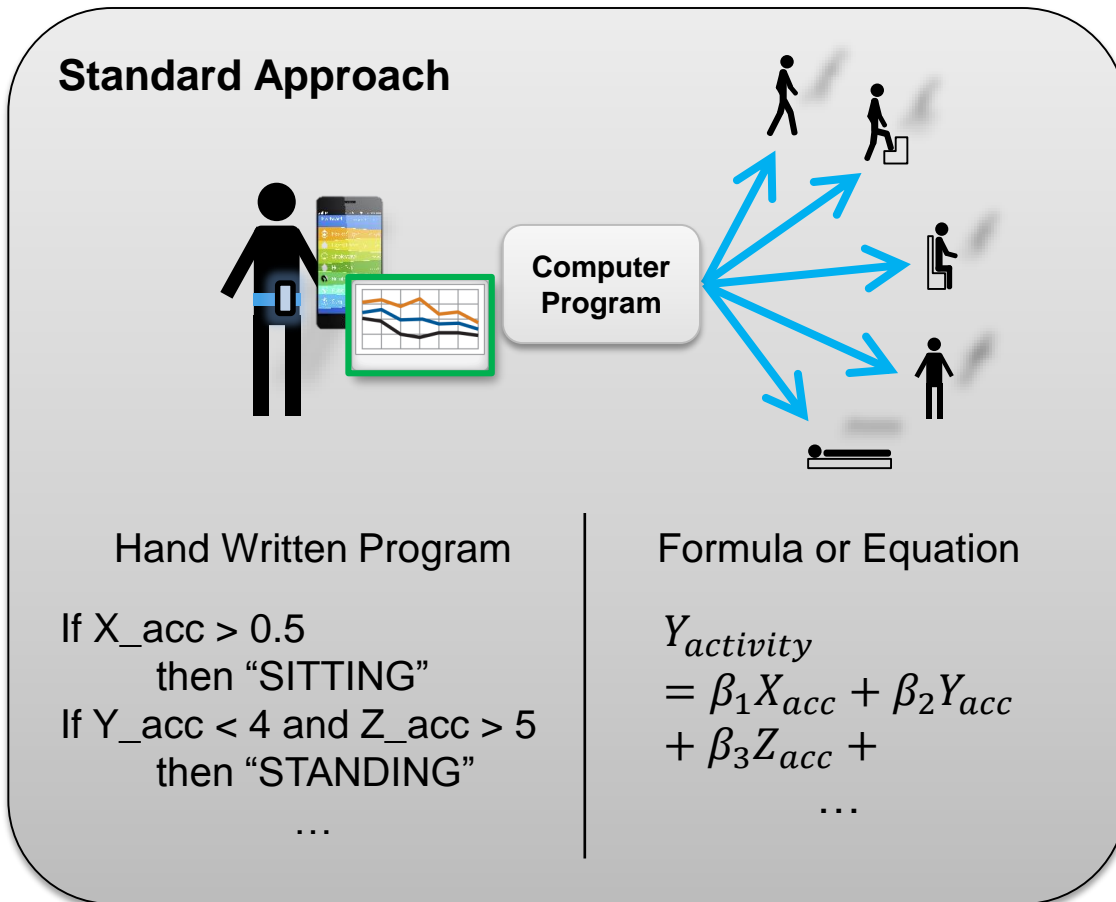
- Image Recognition
- Speech Recognition
- Stock Prediction
- Medical Diagnosis
- Data Analytics
- Robotics
- and more...




Machine Learning

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

Task: Human Activity Detection

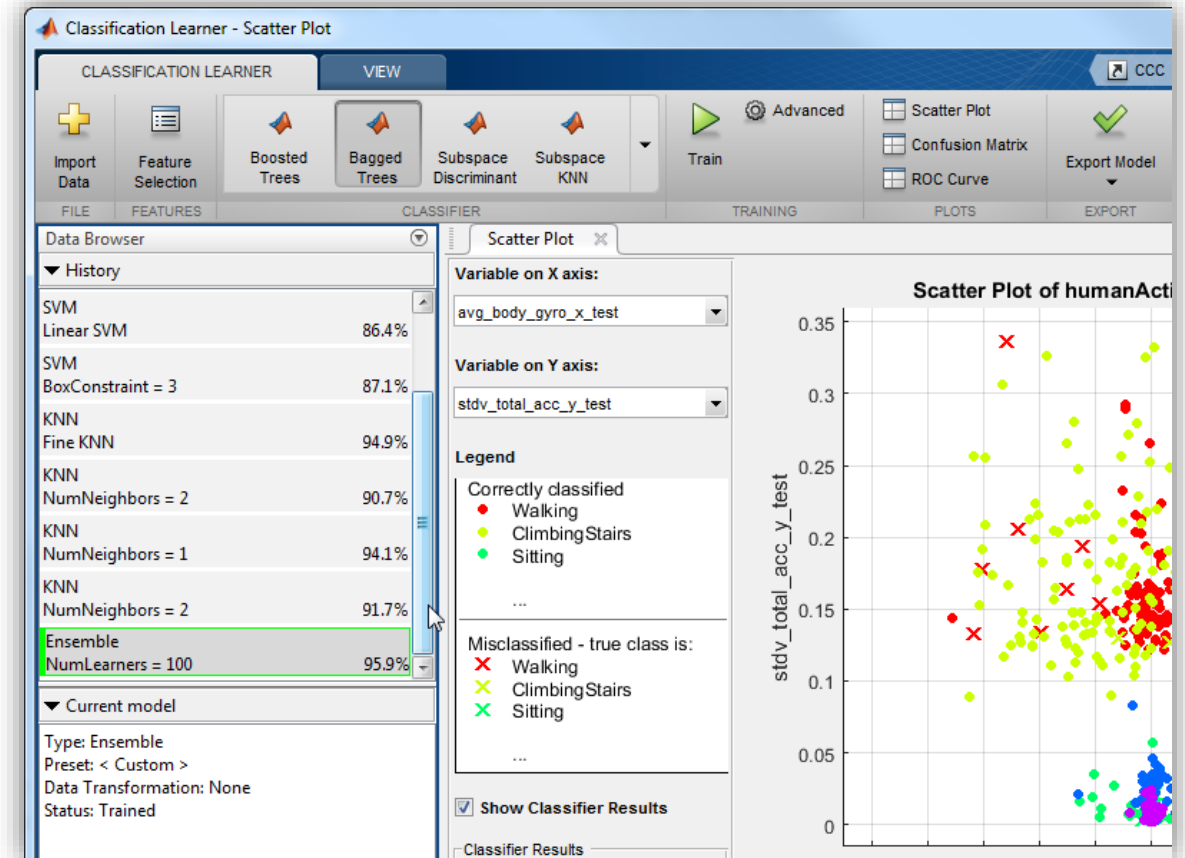


Example: Human Activity Learning Using Mobile Phone Data



Data:

- 3-axial Accelerometer data
- 3-axial Gyroscope data





**“essentially, all models are wrong,
but some are useful”
– George Box**


Challenges in Machine Learning

Hard to get started


Steps	Challenge
Access, explore and analyze data	Data diversity Numeric, Images, Signals, Text – not always tabular

Machine Learning Workflow

Train: Iterate till you find the best model




Predict: Integrate trained models into applications



Agenda






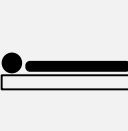
- Machine Learning
 - What is Machine Learning and why do we need it?
 - Common challenges in Machine Learning
- Example 1: Human activity learning using mobile phone data
 - Learning from sensor data
- Example 2: Real-time car identification using images
 - Learning from images
- Summary & Key Takeaways



Example 1: Human Activity Learning Using Mobile Phone Data

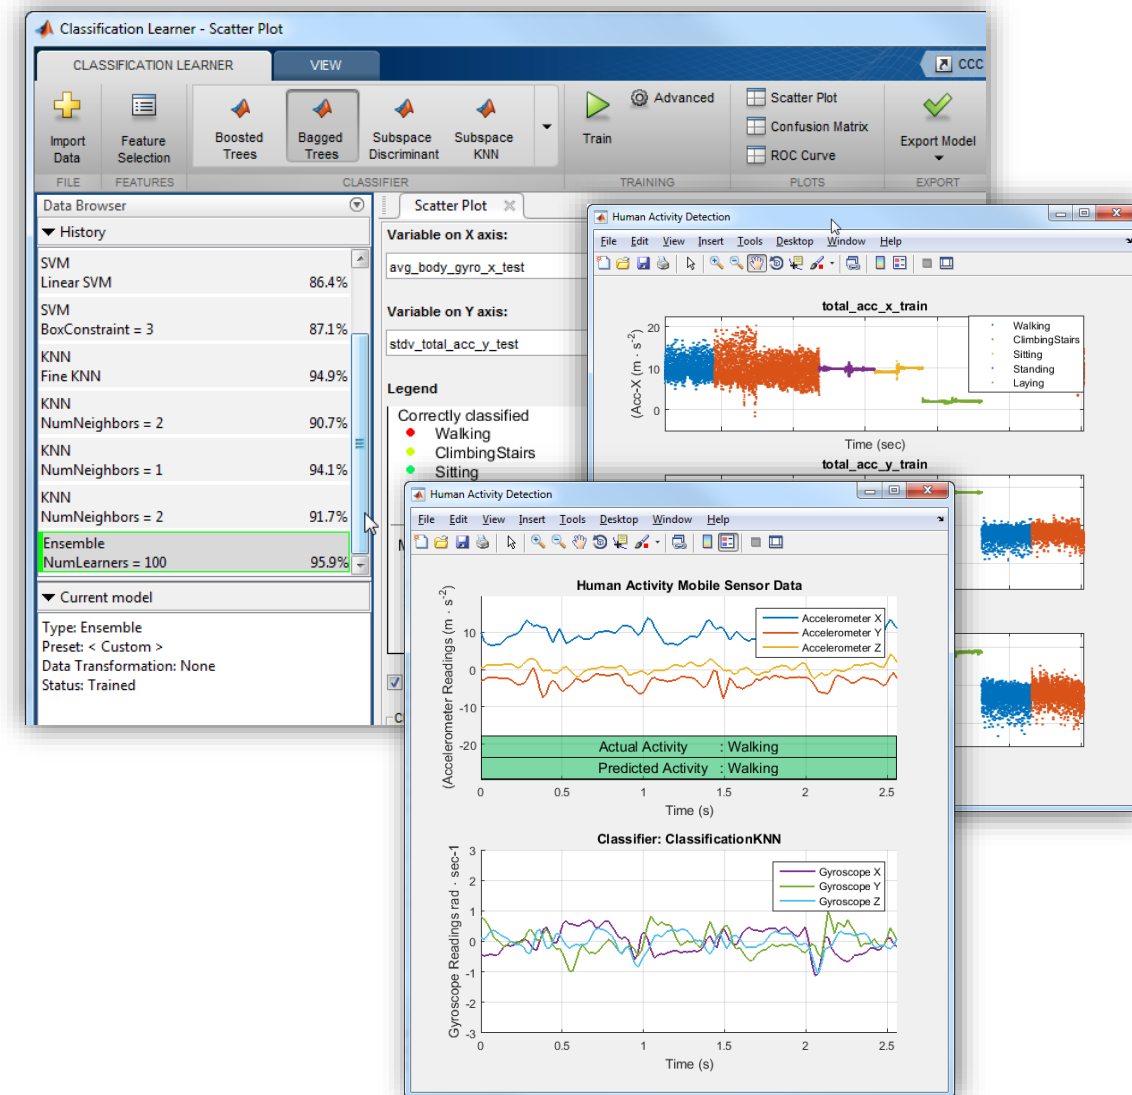
Objective: Train a classifier to classify human activity from sensor data

Data:

Predictors	3-axial Accelerometer and Gyroscope data	
Response	Activity:	    


Approach:

- Extract features from raw sensor signals
- Train and compare classifiers
- Test results on new sensor data

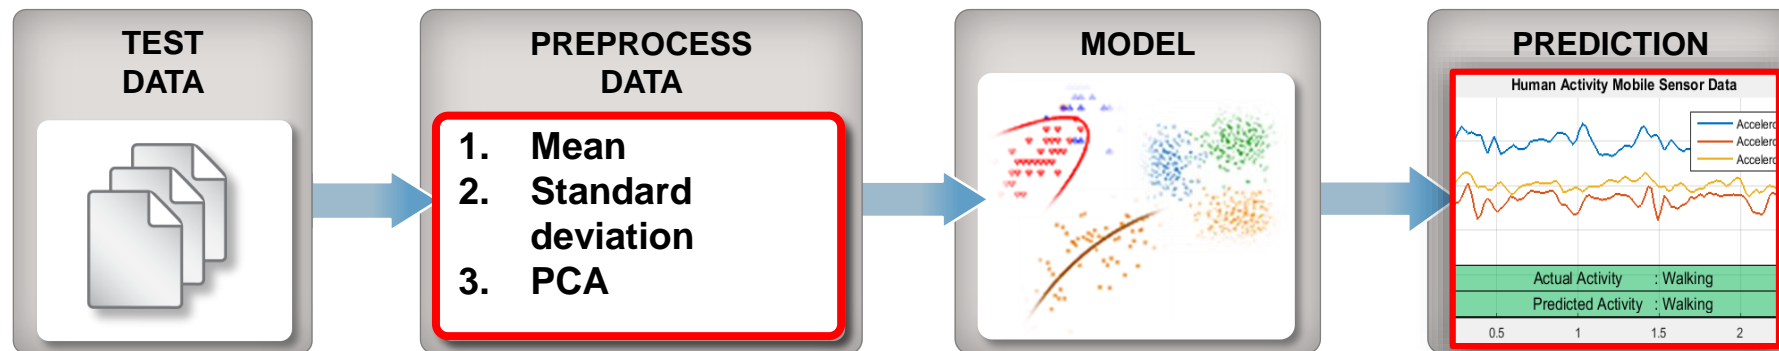


Machine Learning Workflow for Example 1

Train: Iterate till you find the best model




Predict: Integrate trained models into applications



Agenda

- Machine Learning
 - What is Machine Learning and why do we need it?
 - Common challenges in Machine Learning
- Example 1: Human activity learning using mobile phone data
 - Learning from sensor data
- Example 2: Real-time car identification using images
 - Learning from images
- Summary & Key Takeaways



Example 2: Real-time Car Identification Using Images


Objective: Train a classifier to identify car type from a webcam video

Data:

Predictors	Several images of cars: 
Response	NIGEL, LIGHTNING, SANDDUNE, MATER

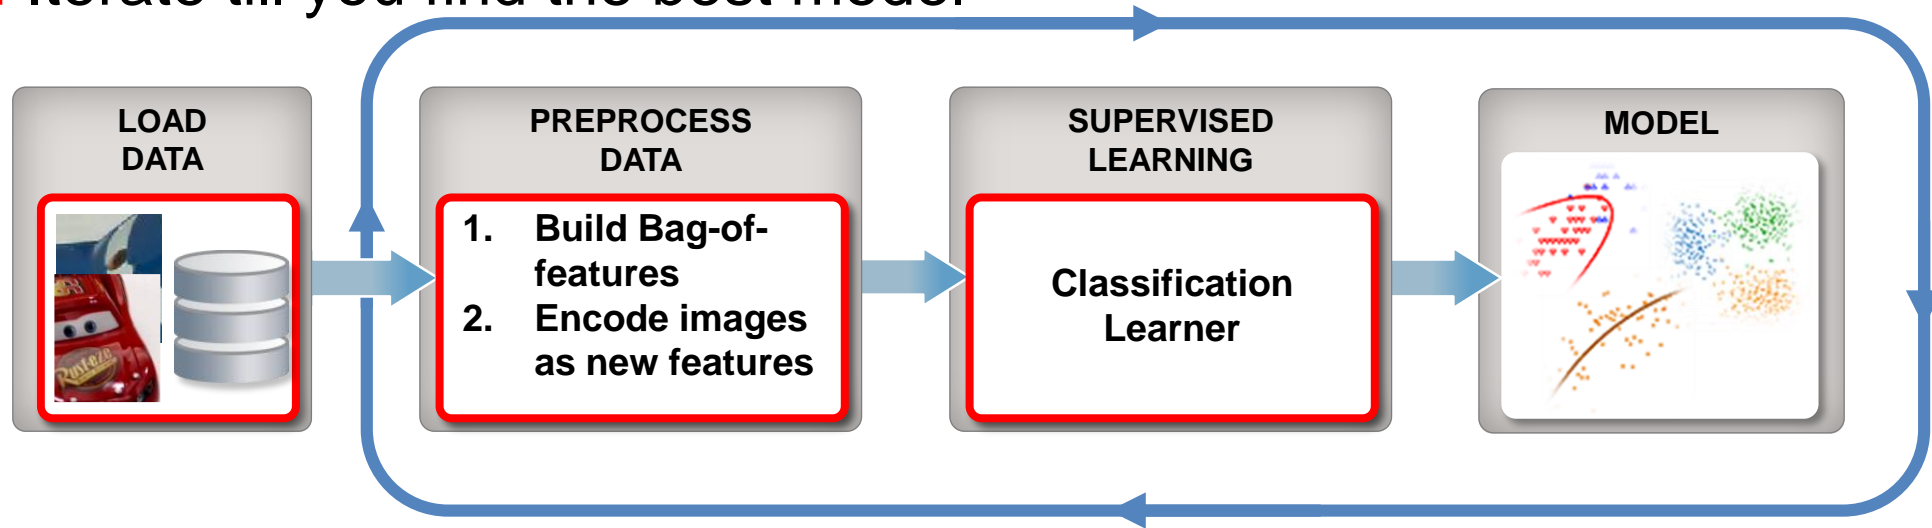
Approach:

- Extract features using Bag-of-words
- Train and compare classifiers
- Classify streaming video from a webcam

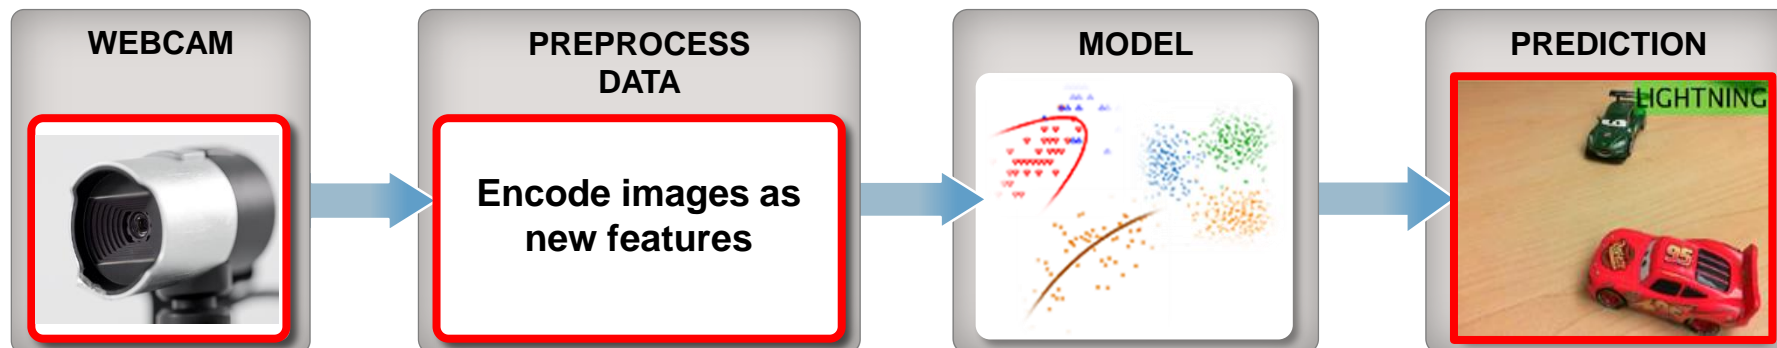


Machine Learning Workflow for Example 2

Train: Iterate till you find the best model




Predict: Integrate trained models into applications




Agenda

- Machine Learning
 - What is Machine Learning and why do we need it?
 - Common challenges in Machine Learning
- Example 1: Human activity learning using mobile phone data
 - Learning from sensor data
- Example 2: Real-time car identification using images
 - Learning from images
- Summary & Key Takeaways




Challenges in Machine Learning


Steps	Challenge
Accessing, exploring and analyzing data	Data diversity
Preprocess data	Lack of domain tools
Train models	Time consuming
Assess model performance	Avoid pitfalls Over Fitting, Speed-Accuracy-Complexity
Iterate	

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*



- MATLAB for Machine Learning



Email me if you have further questions

Additional Resources

Documentation:

The screenshot shows the MathWorks documentation interface for R2015a. The main heading is "Machine Learning" with a sub-heading "Supervised, unsupervised, and ensemble learning". The page provides a general overview of machine learning, explaining that it aims to build models that make decisions based on evidence in the presence of uncertainty. It details supervised learning (classification and regression) and unsupervised learning (clustering and dimensionality reduction). A "Machine Learning Basics" section lists several key topics for further reading, such as "Steps in Supervised Learning", "What Are Linear Regression Models?", and "Introduction to Cluster Analysis". The page also includes a "Was this topic helpful?" feedback form at the bottom.

mathworks.com/machine-learning

The screenshot displays the "Machine Learning with MATLAB Webinar" page. It features a large background image of a scatter plot with blue and red data points. The main text encourages users to learn how to get started using machine learning tools to detect patterns and build predictive models. A prominent "View webinar" button is visible. Below the main text, there are three detailed sections: "Classification" (building models to classify data), "Regression" (building models to predict continuous data), and "Clustering" (finding natural groupings in data). Each section includes a brief description and a list of relevant algorithms. For example, the Classification section lists support vector machine (SVM), boosted and bagged decision trees, and the Regression section lists linear model, nonlinear model, regularization, stepwise regression, and boosted.

Q & A