ECE595 / STAT598: Machine Learning I
Course Overview

Spring 2020

Stanley Chan

School of Electrical and Computer Engineering
Purdue University

PURDUE
UNIVERSITY

Machine Learning: Your Interpretation?

https://www.kaggle.com/benhamner/nips-2015-papers/discussion/22778

Elements of Learning?

- Data
- Computer
- Algorithm

What is Learning? What is NOT learning?

- You are given a bag of US coins.
- Your task: Build a classifier.
- Four classes: Penny, Nickel, Dime, Quarter.

Approach 1: Learning

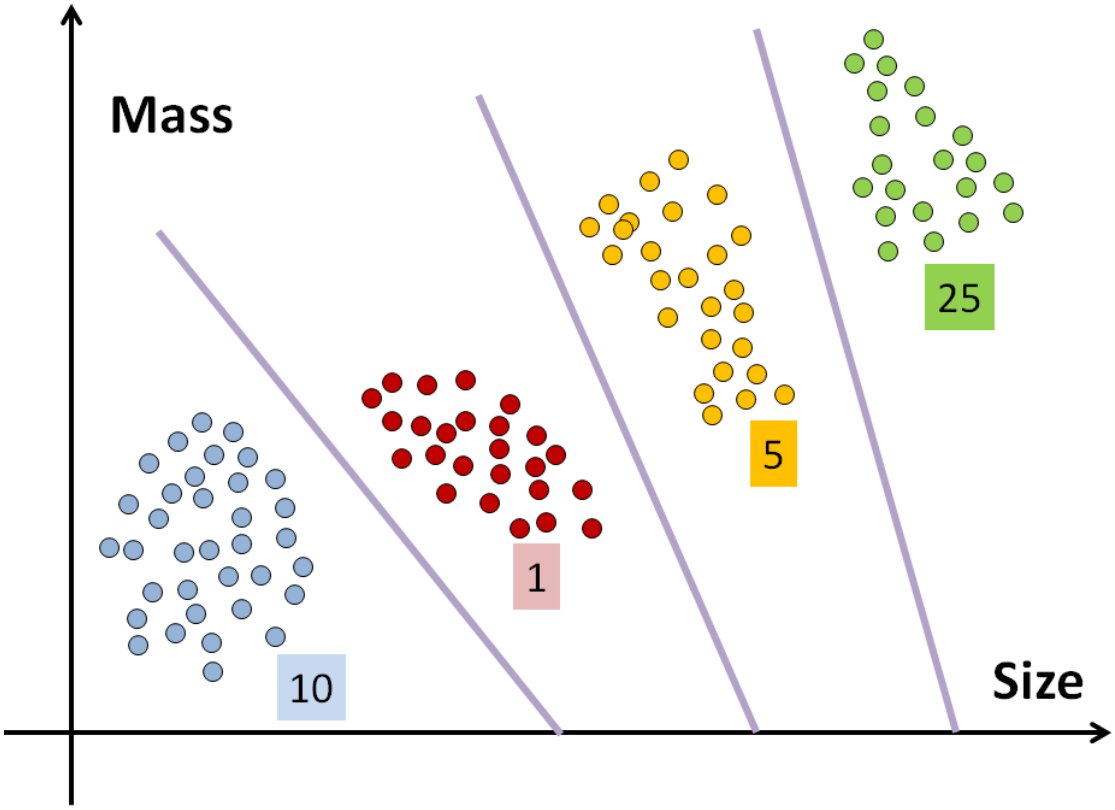

- You measure mass and size.
- Put each coin to its class. Plot a 2D histogram.
- Create the classifier.

Approach 2: Design

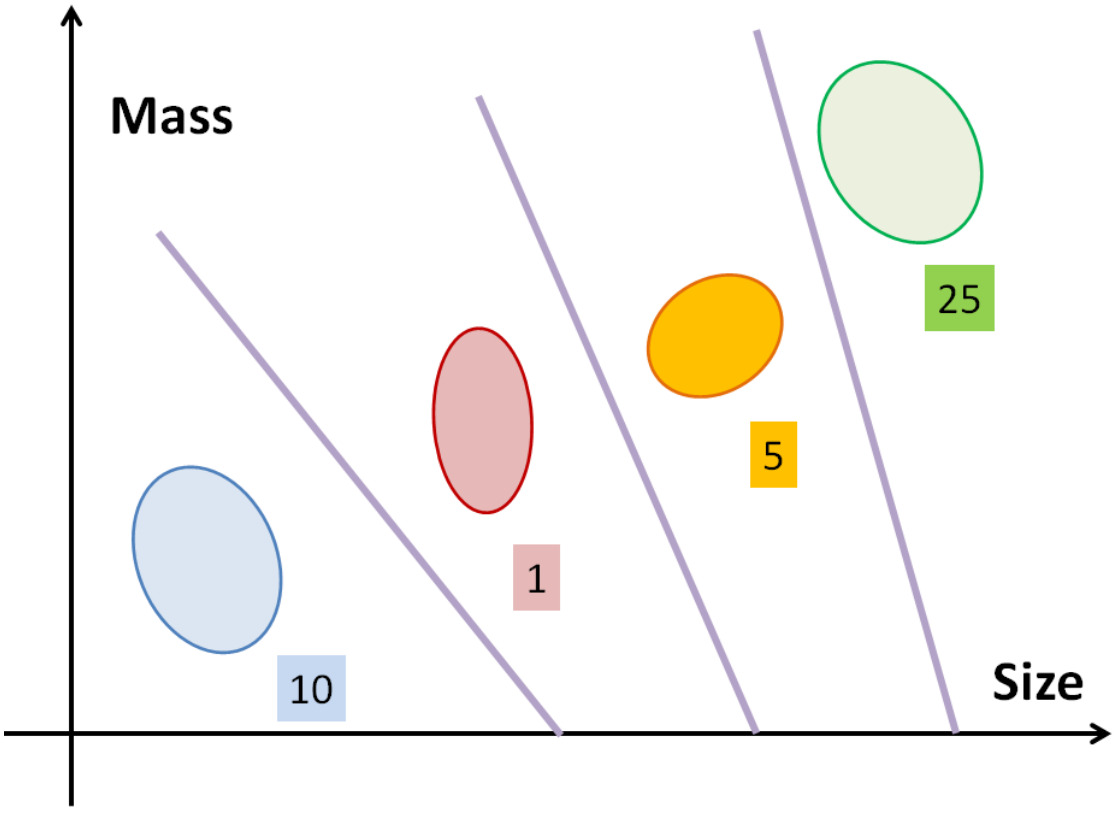

- You go to United States Mint to ask ideal mass and size of the coins.
- You ask them to give you the measurement error. Plot 2D distribution.
- Create the classifier.

Which one fits learning? Which one fits design?

- Determining the age at which a particular medical test should be performed.
- Classifying numbers into primes and non-primes.
- Detecting potential fraud in credit card charges.
- Determining the time it would take a falling object to hit the ground.
- Determining the optimal cycle for traffic lights in a busy intersection.

Machine Learning Model

- Data points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N$.
- Labels $y_1, \ldots, y_N$.

- Where does a data point $\boldsymbol{x}_n$ come from?
- How is a label y_n defined?
- What do we mean by a learning algorithm?
- What is a classifier?
- How to evaluate a classifier?

Machine Learning Model

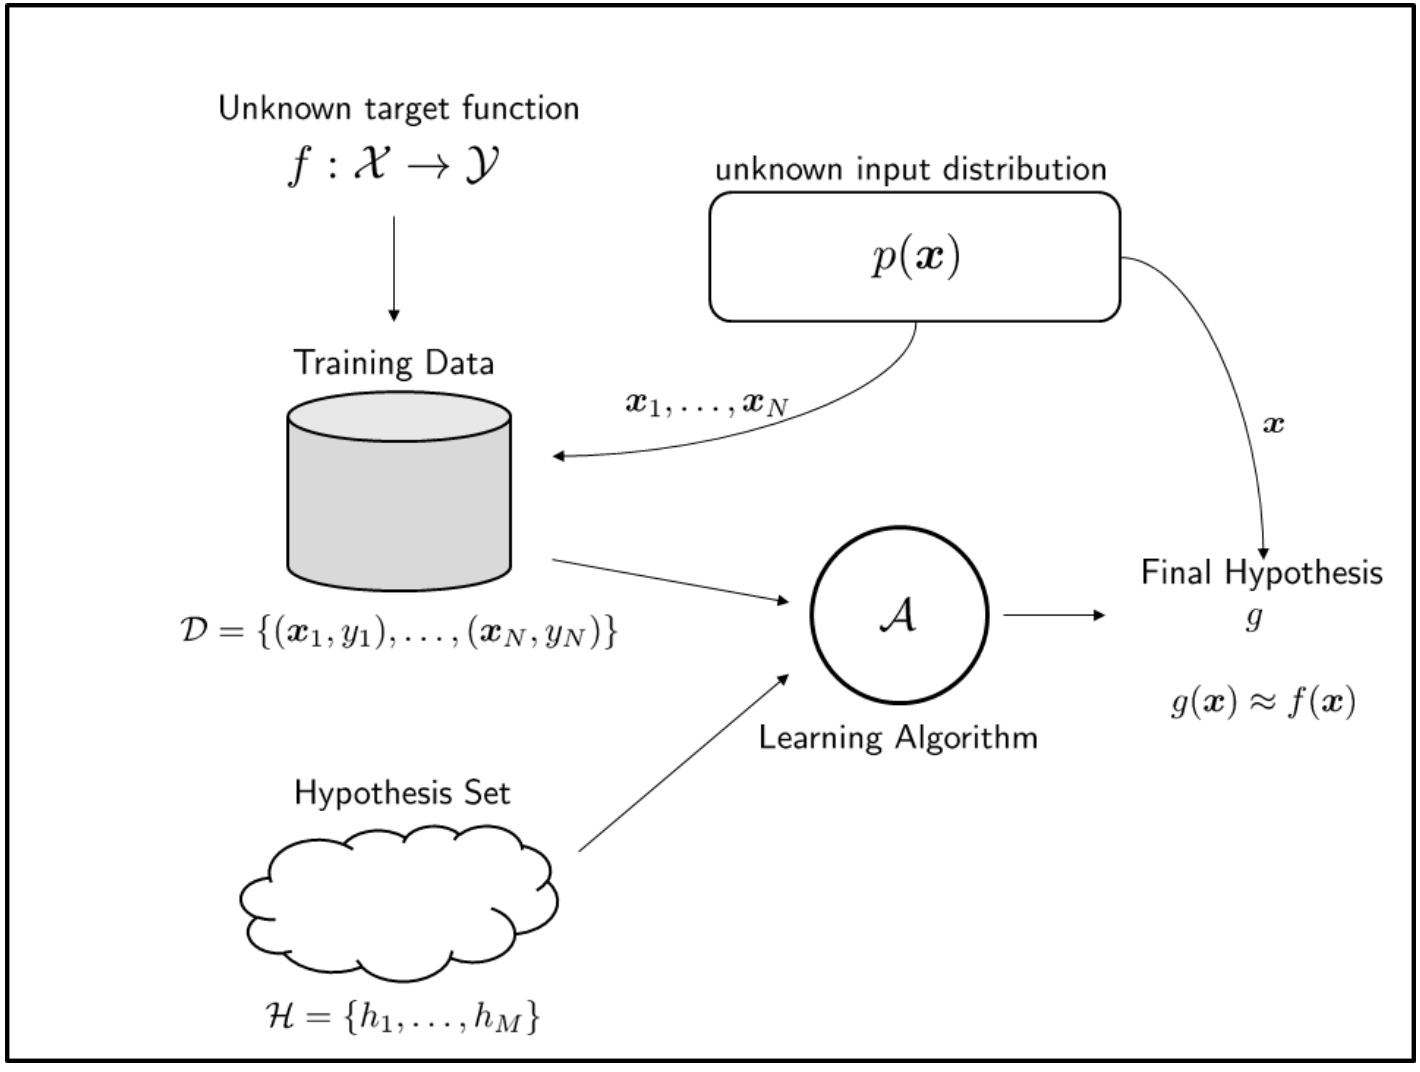

See Learning from Data (Chapter 1).

Learning Algorithm

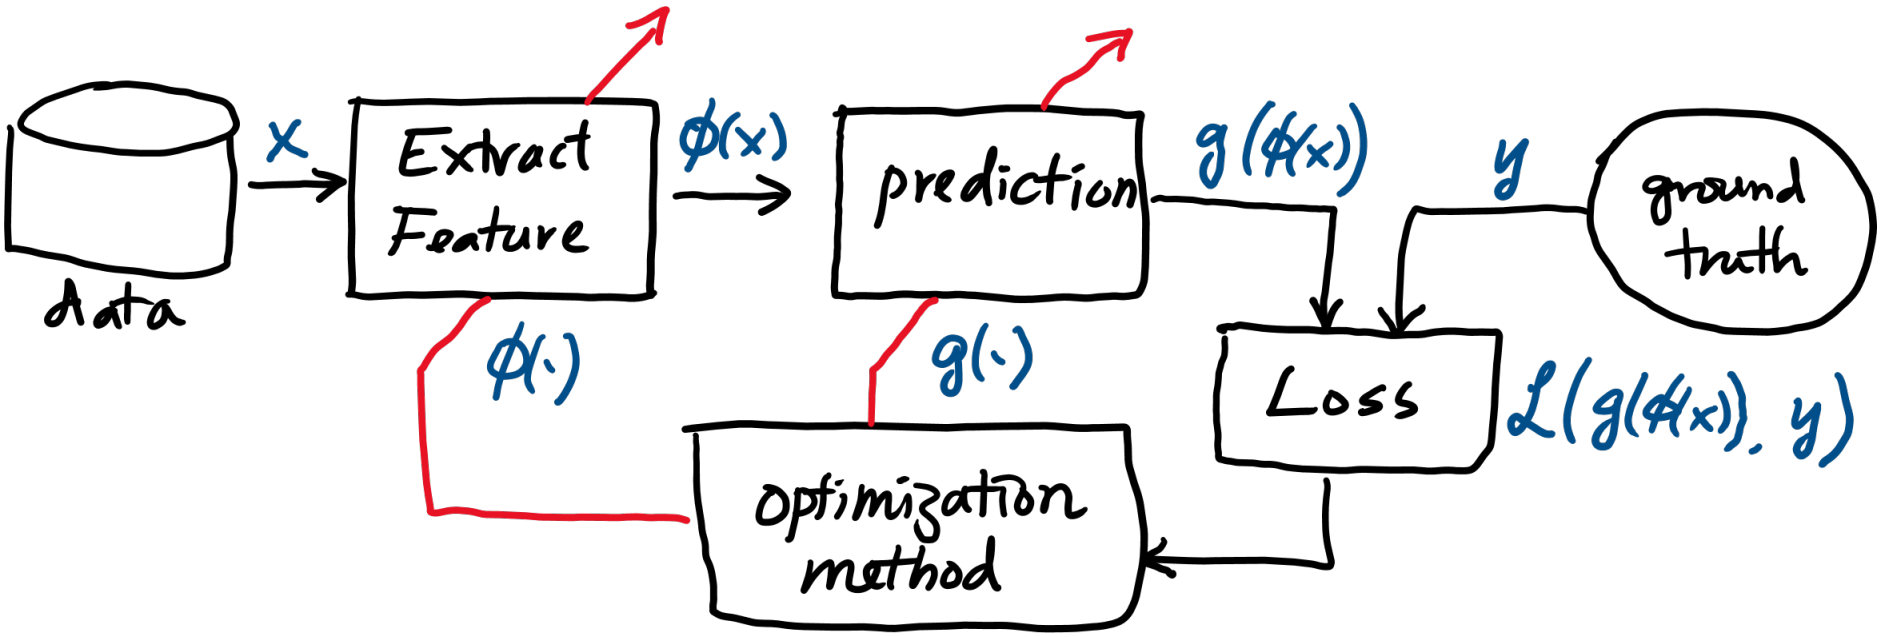

Types of Learning

- Supervised Learning: Labels available.
- Unsupervised Learning: No label.

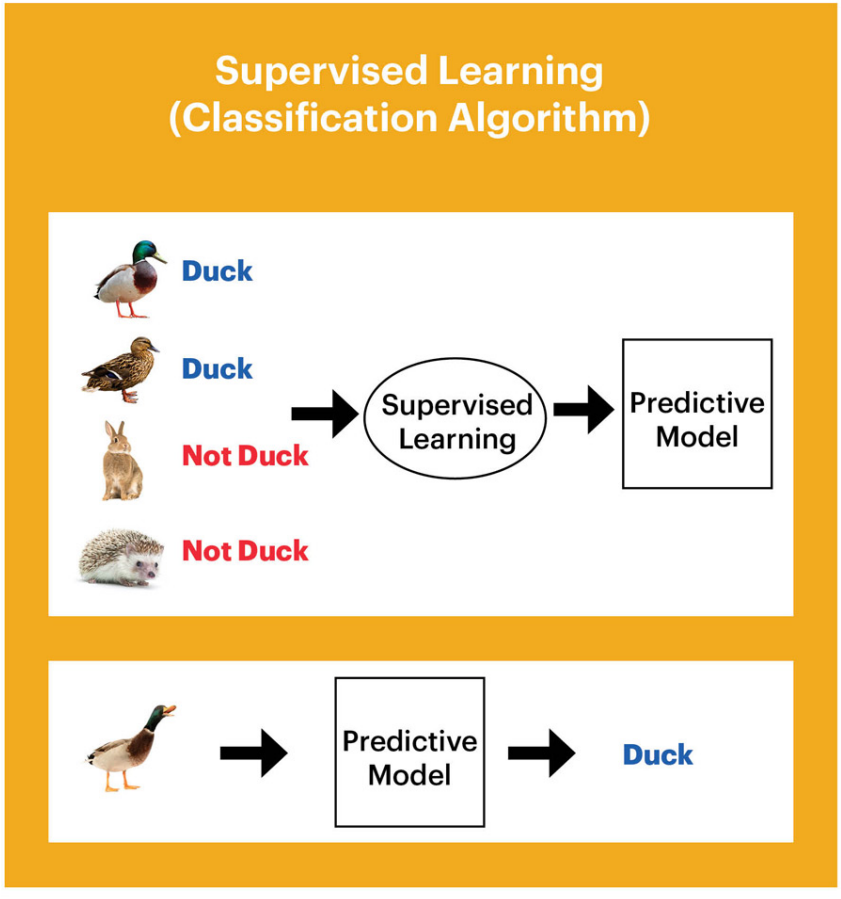

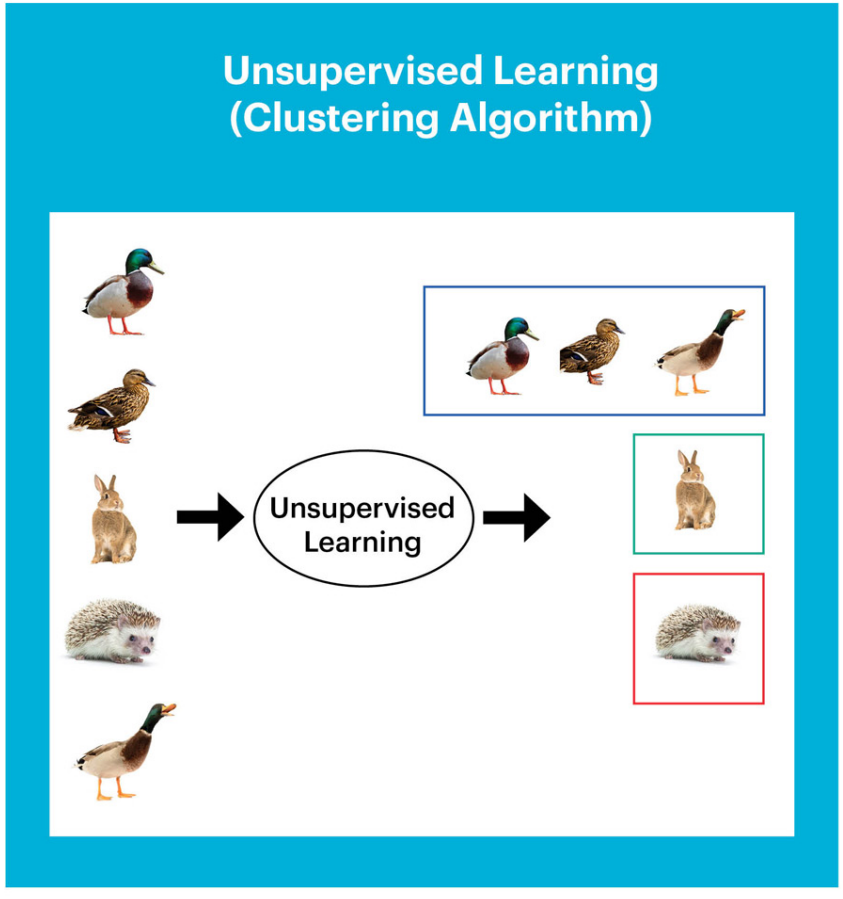

Outline of ECE 595

- **Part 1: Mathematical Background** (2 weeks)
  - Linear Regression and Optimization
  - Please review linear algebra, probability, optimization in the Tutorial Note.
- **Part 2: Classification** (5 weeks)
  - Methods to train linear classifiers
  - Feature analysis, Geometry, Bayesian decision rule, logistic regression, perceptron algorithm, support vector machine
- **Part 3: Handling Uncertainty** (3 weeks)
  - Imperfect data: noisy label, unbalanced data, missing data, knowledge transfer
  - Robustness study: adversarial attack and defense
- **Part 4: Learning Theory** (5 weeks)
  - Evaluation of a classifier.
  - Feasibility of learning, VC dimension, bias-variance, validation

Textbook and References

- **Pattern Classification**, by Duda, Hart and Stork, 2000.
- **Pattern Recognition and Machine Learning**, by Bishop, 2006.
- **Learning from Data**, by Abu-Mostafa, Magdon-Ismail and Lin, 2012.
- **Elements of Statistical Learning**, by Hastie, Tibshirani and Friedman, 2009.

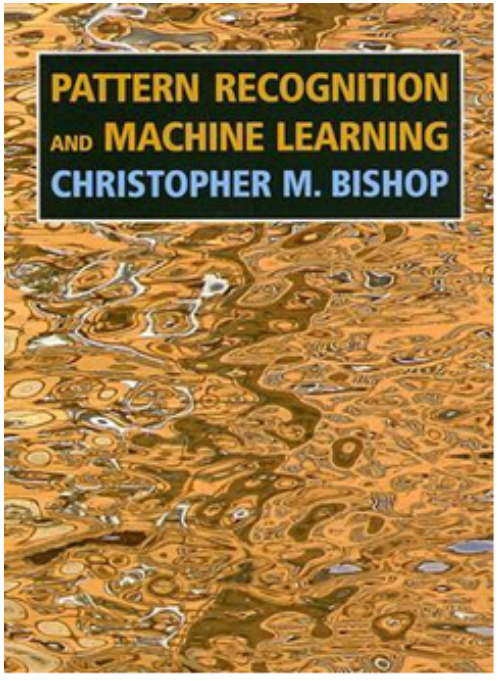

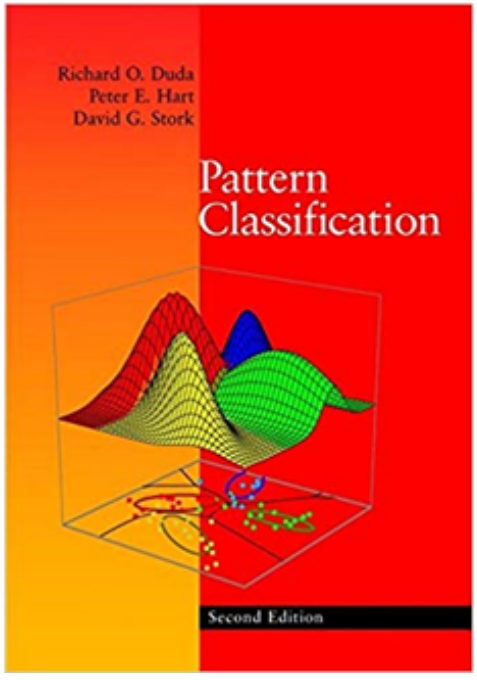

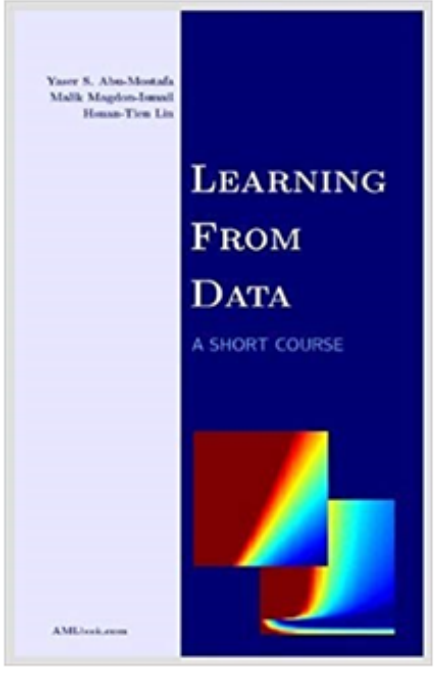

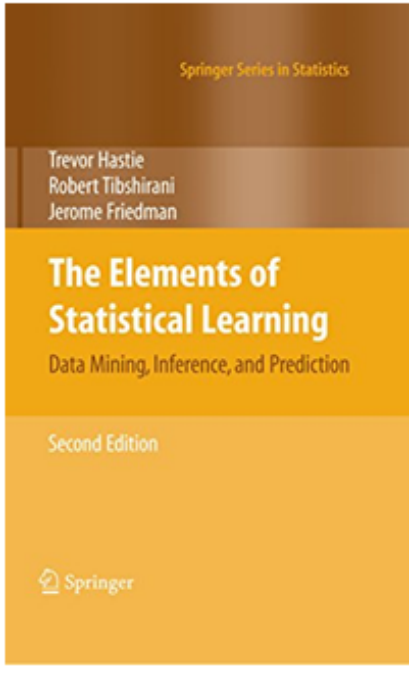

Pre-Requisites

- **Linear Algebra**:
  - Matrix-vector multiplication $\boldsymbol{A}\boldsymbol{x}$, transpose $\boldsymbol{A}^T$, symmetric matrices $\boldsymbol{A} = \boldsymbol{A}^T$, norm $\|\boldsymbol{x}\|$, trace $\text{Tr}(\boldsymbol{A})$, inverse $\boldsymbol{A}^{-1}$, determinant $|\boldsymbol{A}|$, eigenvalue and eigenvector $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^T$.
  - Gilbert Strang, Linear Algebra and Its Applications, 5th Edition
- **Probability**:
  - Random variable X, probability density function $p(x)$, cumulative distribution function $F(x)$, expectation $\mathbb{E}[X]$, variance $\text{Var}[X]$, function of random variable $\mathbb{E}[g(X)]$, joint Gaussian, Law of Large Number, Central Limit Theorem.
  - Dimitri Bertsekas, Introduction to Probability, Athena Scientific, 2008, 2nd Edition.
- **Optimization**:
  - Convex function, convex set, operations which preserve convexity, Lagrange multiplier, KKT conditions, primal optimal, dual optimal, complementary slackness, constrained optimization, duality theorem.
  - Stephen Boyd and Lieven Vandenberghe, Convex Optimization, Cambridge 2004.

Reading List

Tutorials

- Tutorial on Linear Algebra
  `https://engineering.purdue.edu/ChanGroup/ECE595/files/Tutorial01_handout.pdf`
- Tutorial on Probability
  `https://engineering.purdue.edu/ChanGroup/ECE595/files/Tutorial02_handout.pdf`