Foundations of Modern Machine Learning

Syllabus Outline

1 Representation and Learning
- 1.1 Feature Vectors, Feature Spaces
- 1.2 Feature Extraction
- 1.3 Learning Problem Formulation

2 Appreciating and Interpreting Data
- 2.1 Dimensionality Reduction
- 2.2 Data Visualization

3 Classification
- 3.1 Nearest Neighbour methods
- 3.2 Linear Classifiers, Perceptrons, Gradient Descent
- 3.3 Multi-class classifiers
- 3.4 Decision Trees

4 Experimentation Methods
- 4.1 Training, Testing and Validation
- 4.2 Overfitting and Generalization
- 4.3 Feature Engineering
- 4.4 Performance Metrics

5 Probabilistic Methods
- 5.1 Bayes and Naive Bayes Classifiers
- 5.2 Mixture Models
- 5.3 MLE and MAP Estimates

6 Unsupervised Learning and Clustering
- 6.1 K-Means, EM and Mixture Model Fitting
- 6.2 Similarity Metrics, Criterion Functions
- 6.3 Graph-based clustering; Hierarchical Clustering

7 Regression
- 7.1 Linear and Logistic Regression
- 7.2 Regularization

8 Neural Networks
- 8.1 Multilayer Perceptrons
- 8.2 Back-propagation, Training strategies

9 Deep Neural Networks
- 9.1 Convolutional Neural Networks
- 9.2 Recurrent Neural Networks
- 9.3 Autoencoders