

RAI 832 Machine Learning (3-0)

Textbook: Introduction to Machine Learning, by Ethem Alpaydin, MIT Press, 2004.

ISBN-10: 0-262-01211-1, ISBN-13: 978-0-262-01211-9

Reference Book: Machine Learning, by Tom M. Mitchell, McGraw-Hill Science / Engineering / Math; 1 edition (March 1, 1997). ISBN-10: 0070428077, ISBN-13: 978-0070428072 **Objective:**

When you have completed this course, you should be able to apply machine learning algorithms to solve both IID and sequential data problems of moderate complexity. You should also be able to read current research papers in machine learning and understand the issues raised by current research in supervised learning.

Pre-Requisite:

CSE 860 Artificial Intelligence (or equivalent) **Course**

Outcome:

The students graduating from this course are expected to have a thorough knowledge base of Machine Learning algorithms and methodologies as well as their applications in robotics and machine intelligence.

Course Outline:

This course will present an introduction to algorithms for machine learning and data mining. These algorithms lie at the heart of many leading edge computer applications including optical character recognition, speech recognition, text mining, document classification, pattern recognition, computer intrusion detection, and information extraction from web pages. Every machine learning algorithm has both a computational aspect (how to compute the answer) and a statistical aspect (how to ensure that future predictions are accurate). Algorithms covered include linear classifiers (Gaussian maximum likelihood, Naive Bayes, and logistic regression) and non-linear classifiers (neural networks, decision trees, support-vector machines, nearest neighbor methods). The class will also be introduced to techniques for learning from sequential data and advanced ensemble methods such as bagging and boosting.

Topics	Allocated Periods
<p>Overview/Introduction to machine learning</p> <p>Hypothesis spaces</p> <p>Space of Algorithms, Linear Threshold Classifiers</p> <p>Project details</p> <p>Perceptrons</p> <p>Logistic Regression</p> <p>Linear Discriminant Analysis</p> <p>Off-The-Shelf Learning Algorithms</p> <p>Decision Trees</p>	<p>45</p>
<p>Nearest Neighbor</p> <p>Neural networks</p> <p>Bayesian Learning</p> <p>Support Vector Machines</p> <p>Learning Theory</p> <p>Learning Theory finished</p> <p>Bias/Variance Theory & Ensemble Methods</p> <p>Preventing Over-fitting: Penalty and Hold-out methods</p> <p>Hold-Out and Cross-validation Methods</p> <p>Penalty methods: decision trees, neural nets, SVMs</p> <p>Evaluating and Comparing Classifiers</p> <p>Unsupervised Learning</p>	