Course 3	**Supervised Pattern Recognition**
Program	1. Introduction 1.1. Supervised classification 1.2. Semisupervised classification 1.3. Partially supervised classification 1.4. Unsupervised classification 2. Assessing the performance of supervised classification algorithms 2.1. Error generalization 2.2. Area under the ROC curve 2.3. Brier score 2.4. Holdout method 2.5. k-fold cross-validation 2.6. Bootstrapping 3. Preprocessing 3.1. Introduction 3.2. Imputation 3.3. Discretization 3.4. Dimensionality reduction 4. Classification techniques 4.1. Bayesian network classifiers 4.2. Nearest neighbour classifier 4.3. Rule induction 4.4. Classification trees 4.5. Logistic regression 5. Combining multiple classifiers 5.1. Introduction 5.2. Basic methods: - Fusion of label outputs - Fusion of continuous-valued outputs - Stacked generalization - Cascading 5.3. Advanced methods: - Bagging - Randomization - Boosting - Hybrid classifiers 6. Comparing supervised classification algorithms 6.1. Two classifiers in the same database 6.2. More than two classifiers in the same database 6.3. Two classifiers in multiple databases 6.4. More than two classifiers in multiple databases Practical demonstration: WEKA
Bibliography	• Webb, A. (2002) *Statistical Pattern Recognition*, Wiley, 2nd ed. • Kuncheva, L. (2004) *Combining Pattern Classifiers*, Wiley. • Murphy, K. (2012) *Machine Learning: A Probabilistic Perspective*, The MIT Press.
Readings before coming	Recommended listening: “An Introduction to Pattern Classification” by E. Yom Tov, IBM Haifa Research Lab at: http://videolectures.net/mlss03_tov_ipc/