



CS4375: Recap

Nicholas Ruozzi

University of Texas at Dallas

Supervised Learning



- Regression & classification
- Discriminative methods
 - k-NN
 - Decision trees
 - Perceptron
 - SVMs & kernel methods
 - Logistic regression
 - Neural networks
- Parameter learning
 - Maximum likelihood estimation
 - Expectation maximization

Bayesian Approaches



- MAP estimation
- Prior/posterior probabilities
- Bayesian networks
 - Naive Bayes

- Clustering
 - k -means
 - Spectral clustering
 - Hierarchical clustering
- Expectation maximization
 - Soft clustering
 - Mixtures of Gaussians

Learning Theory



- PAC learning
- VC dimension
- Bias/variance tradeoff
- Chernoff bounds
- Sample complexity

Optimization Methods



- Gradient descent
 - Stochastic gradient descent
 - Subgradient methods
- Coordinate descent
- Lagrange multipliers and duality

- Dimensionality Reduction
 - PCA
 - Matrix Factorizations

- Bootstrap sampling
- Bagging
- Boosting

Other Learning Topics



- Evaluation metrics
- Reinforcement learning



Questions about the course content?

For the final...



- You should understand the basic concepts and theory of all of the algorithms and techniques that we have discussed in the course
- There is no need to memorize complicated formulas, etc.
 - For example, if I ask for the sample complexity of a scheme, I will give you the generic formula
- However, you should be able to derive the algorithms and updates
 - E.g., Lagrange multipliers and SVMs, the EM algorithm, etc.

For the final...



- No calculators, books, notes, etc. will be permitted
 - As before, if you need a calculator, you have done something terribly wrong
- The exam will be in roughly the same format
 - Expect true/false questions, short answers, and two-three long answer questions
- Exam will emphasize the new material, but **ALL** material will be tested
- Take a look at the practice exams!

Monday, 12/9/2019

11:00AM - 1:45PM

ECSS 2.203

Related Courses at UTD



- Human Language Technologies (CS 4395)
- Introduction to Computer Vision (CS 4391)
- Artificial Intelligence (CS 4364)
- Computational Methods for Data Scientists (CS 4371)
- Intelligent Systems Design (CS 4315)

ML Related People



- Vincent Ng (NLP)
- Vibhav Gogate (MLNs, Sampling, Graphical Models)
- Sriraam Natarajan (Graphical Models & Reinforcement Learning)
- Sanda Harabagiu (NLP & Health)
- Dan Moldovan (NLP)
- Nicholas Ruozzi (Graphical Models & Approx. Inference)
- Rishabh Iyer (Submodular Functions)
- Jessica Ouyang (NLP)



Please evaluate the course!

eval.utdallas.edu