Machine Learning Principles

Course Number: 01:198:461 section 06
Instructor: Bernhard Firner
Email: bfirner@cs.rutgers.edu
Office hours: Mondays and Wednesdays from 2:40-3:40 PM in Hill 273.

Lectures are held on Mondays and Wednesdays from 3:50PM to 5:10PM in SEC 202.

Recitations are Mondays from 7:45 to 8:40PM in SEC 203.

Lecture Notes

- Slides will be posted on https://firner.com/presentations/cs461
  - They may be updated after the lecture to correct typos or add last-minute content
  - All code is copyable from the slides
- Recitation slides and materials will be posted on https://daizedong.github.io/teaching/2025-fall-cs461

Resources

- Recommended Reading
  - Machine Learning: A Probabilistic Perspective by Murphy
    - This is a big book with a lot of concepts and a fair amount of math
    - It has a wealth of information on everything *but* deep learning
  - Machine Learning: The Art and Science of Algorithms that Make Sense of Data by Flach
    - This is a more approachable book, with straightforward explanations and algorithms
    - It stops short of deep learning, but is a nice, practical book otherwise
- For Neural Networks (starting with results from the 90s)
  - I will select readable research papers and highlight sections
- Suggested reading will be included in the slides for each lecture

Assessment

- 20% Quizzes (in class)
  - 6 quizzes, lowest score dropped
- 50% Assignments
- 30% Final Exam

Quizzes and the final exam will be done on paper and are meant to assess your mastery of course topics. Assessments may cover:

- Questions about graphs and how to interpret them
- Advantages or disadvantages of various ML techniques
- Understanding of fundamental equations
- Implementation details of important techniques in ML

Assignments will be assessed a late penalty of 20% for up to two days late (at 10% per day) and a penalty of 50% at one week (6% per day after the second day). If you have trouble meeting an assignment deadline, please reach out before the actual due date.

Schedule

Class	Date	Topic	Note
1	09/03	missed	
2	09/08	Linear Regression & Least Squares, Ridge Regression, Lasso, Complexity, Bias, and Variance	
3	09/10	Logistic Regression, Bias Variance Tradeoff, & Regularization	
4	09/15	Normalization, logistic regression and decision boundaries, exemplar vs non-exemplar approaches	
5	09/17	Dealing with discrete and tabular data Decision Trees, Gini Impurity	HW 1
6	09/22	Parametric Vs Nonparametric models, Clustering, HAC, K-Nearest Neighbors, K-Means	
7	09/24	Generative probabilistic models, K-Means++, the EM Algorithm, Gaussian Mixture Models	Quiz 1
8	09/29	Gaussian Mixture models and the EM Algorithm	
9	10/01	Latent Linear, PCA, dimensionality reduction, latent spaces	
10	10/06	Bagging (power of ensembles)	HW 2
11	10/08	Random Forests and AdaBoost, Ensemble Model Indepedence	Quiz 2
12	10/13	Sequence modelling with Markov chains	
13	10/15	Hidden Markov models, Forwards-Backwards Algorithm	
14	10/20	Hidden Markov models and applications, language modelling	
15	10/22	Spatial modelling; the Perceptron, Kernel functions	Quiz 3
16	10/27	Kernel functions, the kernel trick and linear classifiers	HW 3.1 (NLP Part 1)
17	10/29	Kernels, polynomial and radial basis functions	
18	11/03	Large Margins and Support Vector Machines, the Dual Form	HW 3.2 (NLP Part 2)
19	11/05	Before NNs; RBMs	Quiz 4
20	11/10	Neural Networks, Nonlinear Activations, SGD, Universal Approximation Theorem, Vanishing Gradients	HW 4.1 (Convolutions 1)
21	11/12	Convolutional Neural Networks, LeNet	
22	11/17	DNN capacity, AlexNet, and the deep learning explosion	HW 4.2 (Convolutions 2)
23	11/19	ResNets and DNNs as black box functions	Quiz 5
24	11/24	Temporal neural networks An overview of LSTMs, RNNs, Attention & Transformers	Break
25	12/01	Lessons from ResNeXt	HW 5
26	12/03	End-to-end learning and Unsupervised techniques	
27	12/08	Reinforcement learning	
28	12/10	Wrapup and review	

Table 1: Machine Learning Course Schedule