

Introduction to the theory and applications of *Machine Learning From Data*. The (adaptive) outline is:

1. Foundations

- (i) *The Learning Problem*: What is learning?
- (ii) *Training Versus Testing*: Can we learn?
- (iii) *The Linear Model*: How to learn?
- (iv) *Overfitting*: How to learn well?
- (v) *Three Learning Principles*: Lessons learned: Occam; sampling bias; snooping.

2. Techniques (we will cover a few)

- (vi) *Similarity Based Methods*
- (vii) *Neural Networks*
- (viii) *Support Vector Machines*
- (ix) *Aggregation Methods*: methods for combining models.
- (x) *Learning Aides*: methods for improving the performance.

3. Other Paradigms (we won't have time)

- (xi) *Bayesian Learning*
- (xii) *Graphical Models*
- (xiii) *Reinforcement Learning*
- (xiv) *Unlabeled Data*

Text Book: *Learning From Data* by Abu-Mostafa, Magdon-Ismail, Lin.

We will be following the foundational approach in this textbook. However, there are other suggested textbooks at the end of this document which you can look to for similar information.

Book Forum (free): book.caltech.edu/bookforum. A forum where readers discuss exercises from the book.

Learning Outcomes. Students entering this course should have mathematical preparation in multivariate calculus, probability and linear algebra, plus experience with computer algorithms, data-structures and programming. This course will develop an understanding of the limits to which information can be learned from data, and how. The student will gain an understanding of basic tradeoffs in learning from data, and the general pitfalls. At the end of the course, the student should be able to: implement, as well as critique basic models of learning from data; formulate a learning problem precisely, in terms of inputs and outputs; select a learning model and algorithm, run it on the data, and interpret the results; provide some measures of how effective the learning was; read and critique recent published literature in the field.

In addition, students in the 6100-level course should be able to apply more advanced mathematical techniques involving linear algebra and mathematical programming to solve suitably formulated learning tasks.

Prerequisites

Mathematics: Familiarity at the level of Assignment 0 is expected (eg. MATH 1020 & 2010, MATP 4600.) Students without adequate preparation have had serious difficulty, with adverse effects on their grade.

Computing Skills: Ability to program and develop algorithms in some programming language, (eg. CSCI-2300). We will *not* be offering “debugging help”.

Policies

Grade: homework 100% (online version)

Homeworks will be due roughly every week, and will closely follow lectures. Some requirements on the homeworks for the 4100 level differ from the 6100 level (students at the 6100-level will be required to demonstrate more advanced use of linear algebra and linear programming). Thus, you should read the homeworks *carefully*.

There will be separate curves for the 4100 level and the 6100 level courses. Historically, the threshold for A has been approximately 95%.

Collaboration and Academic Dishonesty

Discussion is allowed. Copying (from **anywhere** other than the class text) is not. You should write and understand all solutions yourself. Do not destroy any of your code. In the event of strange results, we may ask you to submit your code.

You are expected to treat your work with pride, and to respect the work of others. If someone else’s work is used to produce any work you will hand in to this course, to the extent that you use their results or techniques, you should: (i) indicate how this third party work was used to solve your tasks; and, (ii) acknowledge the original authors of the work in your submission.

Plagiarizing someone else’s work is a **serious issue**. In cases of academic dishonesty, the minimum penalty will be an automatic grade of F, in addition to other institute mandated protocols.

Late Assignments

Assignments are generally due one week after being handed out. Late assignments are penalized 20% per day (except in institute established illness or emergency).

Other Suggested Texts and Related Reading

LEARNING

Bishop, *Neural Networks for Pattern Recognition*
Mitchell, *Machine Learning*
Sutton, Barto, *Reinforcement Learning*
Vapnik, *Statistical Learning Theory*
Duda, Hart, *Pattern Classification and Scene Analysis*
Ripley, *Pattern Recognition and Neural Networks*
Murphy, *Machine Learning*
Mohri, et al., *Foundations of Machine Learning*

GENERAL MATHEMATICS

Olofsson, *Probability, Statistics and Stochastic Processes*
Pitman, *Probability*
Degroot, *Probability and Statistics*
Khuri, *Advanced Calculus with Applications in Statistics*
Lang, *Undergraduate Algebra*
Horn and Johnson, *Matrix Analysis*
Golub & Van Loan, *Matrix Computations*