Course Information

Mohammad Emtiyaz Khan
エムティ
Team Leader, Approximate Bayesian Inference Team
RIKEN Center for Advanced Intelligence Project (AIP), Tokyo.
Visiting Professor at TUAT, Tokyo
Part-time Lecturer at Waseda University

Goals

- By the end of this part, you should be able to
  - Describe the goal this course
  - Summarize what you will learn
  - Understand the course evaluation
  - Know the pre-requisite for the course and prepare yourself for the difficult times ahead

Couse Goals

- Define Regression and Classification, and explain the main differences between them
- Describe a few models and algorithms for them.
- Implement and apply these methods.
- Derive the theory behind ML methods taught in the course and generalize them to new problems.
- Continue to work through difficulties or initial failure to find optimal solutions.
- Assess one's own level of skill acquisition, and plan their on-going learning goals.

(A rough) Outline of the Course

- [Apr 14] Supervised ML-I: Regression, linear regression, cost function, gradient descent, least squares.
- [Apr 21] Evaluation methods: Overfitting, regularization, bias-variance, relation to deep learning
- [Apr 28] Supervised ML-II: Classification, logistic regression, neural networks, (kernel ridge)
- [May 12] Unsupervised ML: K-means, PCA

Structure of Each Lecture Day

- Four lecture days (with 8 lectures)
  - April 14, 21, 28 and May 12.
- From 2pm to 5:30
- Two lectures each of length 1hr 30 minutes, with a 30 minutes break in between.
- We will have multiple breaks in between. Sometimes for discussion.

Course Evaluation

- A final exam (40%)
- Submit weekly reports (30%)
  - Read the lecture notes, try to do labs offline and submit a short report (at most 2-3 pages).
  - In the report, explain what you learnt and what you found difficult to understand
  - For labs, try to do as much as possible. It is not mandatory.
- Class participation (30%)
  - Participate in in-class quizes
  - Take notes and, during each break, write down a 1-minute summary of what you learnt.
  - Submit this at the end of each lecture day.

A bit more about the final exam

- It will contain questions on what you have learned during the lectures and exercise sessions.
- You are allowed to bring one cheat sheet (A4 size paper both sides can be used) and a calculator.
- No collaborations. No cell phones. No laptops etc.

The Teaching Team

- Toru Asahi
  (tasahi@waseda.jp)
- Yusuke Maruyama
  (y_maruyama@aoni.waseda.jp)

Teaching Assistant

- Wataru Onodera
  (wonodera0501@moegi.waseda.jp)
- Yuta Okamoto
  (yutaokamoto48@ruri.waseda.jp)

Resources

- Course Webpage
  - https://emtiyaz.github.io/teaching/waseda18_ml/ml.html
  - The course material will also be available on the internal website of Waseda University
- Lecture notes
  - During the lectures, I will use lecture notes (with blank space for you to take note if you want).
  - These will be available on the course webpage one day before the lecture. We will also provide a printed version before each lecture day.
  - An annotated copy (with my annotations from the class) will be available after the lecture.

Books

- Books (not required but recommended)
  - G. James, D. Witten, T. Hastie and R. Tibshirani: An introduction to statistical learning (free download from http://www-bcf.usc. edu/~gareth/ISL/).
  - T. Hastie, R. Tibshirani and J. Friedman: Elements of statistical learning (download from http://statweb.stanford.edu/~tibs/ElemStatLearn/).
  - C. Bishop: Pattern Recognition and Machine Learning.
  - K. Murphy: Machine Learning: A Probabilistic Perspective

What not to expect!

- You will not learn ALL advanced methods.
- You will not learn ALL the details.
- This course is not about big data or largescale methods.
- This is not a course about numerical optimization, neither is it about statistics. We will use both of these and learn basic techniques only.
- We will not teach the pre-requisite for ML. You have to learn that on your own.
- This course does not teach you all that you need to know to be able to apply machine learning, but this course will get you started for sure

Prerequisite (must know)

- Matrix calculus.
  - How to take derivative with respect to vectors and matrices.
    - https://atmos.washington.edu/~dennis/MatrixCalculus.pdf
    - https://en.wikipedia.org/wiki/Matrix_calculus
  - You can learn more about it from wikipedia or Matrix Cookbook
    - http://www.imm.dtu.dk/pubdb/views/edoc_download.php/3274/pdf/imm3274.pdf
- Basic Probability
  - Normal distribution
  - Read Chapter 2 in Bishop's book on Machine Learning

Prerequisite (must know)

- Matrix algebra.
  - Basics: Vector and matrix multiplication, (https://en.wikipedia.org/wiki/Matrix_multiplication)
  - More advance topics (see Wikipedia): Matrix inversion and determinants, rank, null and range space, eigenvalue decomposition.
  - There is also a handout posted on the course webpage.

Goals

- By the end of this part, you should be able to
  - Describe the goal this course
  - Summarize what you will learn
  - Understand the course evaluation
  - Know the pre-requisite for the course and prepare yourself for the difficult times ahead