



STAT3612 Statistical Machine Learning (Fall 2020)

Course Outline

| | |
|------------------------|--|
| Instructor: | Dr. Aijun Zhang (RR224) |
| Email: | ajzhang@umich.edu |
| Lecture Hours: | Tuesday 1:30pm – 4:20pm (Zoom, & T1) |
| Tutors: | Mr. Yuyang Xu (RR114) and Mr. Yifeng Guo (RR114) |
| Tutorial Hours: | TBD |
| Course Website: | http://moodle.hku.hk & Github: ajzhanghku/Stat3612 |

Course Objectives:

Machine learning is the study of computer algorithms that build models of observed data in order to make predictions or decisions. Statistical machine learning emphasizes the importance of statistical theory and methodology in the algorithmic development. This course provides a comprehensive and practical coverage of essential machine learning concepts and a variety of learning algorithms under supervised and unsupervised settings. The course materials are presented with lots of examples and reproducible codes.

Contents and Topics:

Data science, data exploration, generalized linear models, variable selection, basis expansion, regularization, cross-validation, tree-based methods, kernel methods, neural networks, dimension reduction, principal component analysis, cluster analysis, stochastic optimization, interpretable machine learning.

Intended Learning Outcomes:

1. Get familiar with the workflow of a data science or machine learning project;
2. Understand and apply a wide range of statistical machine learning methods, and recognize their characteristics, strengths and weaknesses;
3. Identify and use appropriate techniques for a particular data science project;
4. Evaluate the quality of the resulting model in terms of prediction accuracy and model explainability;
5. Apply Python/R programming for solving data-scientific problems.

Prerequisites:

STAT2602 (Probability and Statistics II) or STAT3902 (Statistical Models)

Assessment:

| | | |
|-----------------|--|-----|
| Assignments: | Hands-on data analysis and coding (3 sets) | 30% |
| In-class Tests: | Problem solving skills (2 times) | 40% |
| Group Project: | Proposal, oral presentation and written report | 30% |

Software/Programming:

Python (Jupyter Notebook), R (RStudio), TensorFlow/Keras.

References:

1. James, G., Witten, D, Hastie, T. and Tibshirani R. (2013). *An Introduction to Statistical Learning with Applications in R*, Springer, New York. <http://faculty.marshall.usc.edu/gareth-james/ISL/>
2. Hastie, T, Tibshirani, R. and Friedman, J. (2009). **The Elements of Statistical Learning: Data Mining, Inference, and Prediction**. Second Edition, Springer, New York. <https://web.stanford.edu/~hastie/ElemStatLearn/>
3. Géron, A. (2017). *Hands-On Machine Learning with Scikit-Learn and TensorFlow*, OReilly. <https://github.com/ageron/handson-ml>
4. Chollet, F. (2018). *Deep Learning with Python*. Manning. <https://github.com/fchollet/deep-learning-with-python-notebooks>

For Future Data Scientists:

