CS4033/5033: Machine Learning

Fall 2023

Time: MWF 10:30am-11:20am

Location: Dale Hall 0206

Instructor: Chao Lan (clan@ou.edu)

Teaching Assistant: Shayan Shafaei (shayan.shafaei@ou.edu)

Office Hours:

– Instructor: Thu 9:30am-12:30pm at DEH 210.

– TA: Thu 3:30pm-5pm on Zoom (ID: 4551860899, pwd: 63471781).

---

1. Course Description

This course introduces the fundamentals of machine learning, with a focus on the mathematical aspects of predictive models and learning techniques. It covers the design, interpretation, optimization and enhancement of models for different tasks including regression, classification and clustering. Topics include

Regression

– Linear regression model and learning techniques (least square, ridge, Lasso).

– Kernel regression model and learning techniques.

– Probabilistic interpretation of linear regression using MLE and MAP.

Classification

– Linear classification models (logistic regression, LDA, LSVM) and learning techniques.

– Nonlinear classification models (naive Bayes, k-NN, tree, neural network) and learning techniques.

– Theoretical properties such as bias-variance trade-off, generalization, model complexity.

Model Enhancement

– Hyper-parameter optimization using K-fold cross validation.

– Multi-class classification using 1v1 and 1vA strategies.

– Ensemble methods (bagging, boosting).

Other Topics

– Clustering (K-means)

– Dimensionality Reduction (PCA)

– Collaborative Filtering (matrix factorization)

– Sequential Prediction (Markov model)

– Algorithmic Ethics

2. Main Reference

[PML] K.P. Murphy. Probabilistic Machine Learning: An Introduction. MIT Press, March 2022.

[ELS] T. Hastie, R. Tibshirani and J. Friedman. The Elements of Statistical Learning. Springer, 2009.

[PRML] C. Bishop. Pattern Recognition and Machine Learning. Springer, 2006.

[COV] S. Boyd and L.Vandenberghe. Convex Optimization. Cambridge University Press, 2004.

[ALS] L. Wasserman. All of Statistics. Springer, 2004.

** All reference books have free copies online.*

3. Assignment

There will be approximately 20 assignments, each containing some written and/or programming tasks. Written tasks expect students to derive the mathematics of certain machine learning techniques and present the results in Latex-created PDF files, and programming tasks expect students to implement certain machine learning techniques from scratch in Python and evaluate their performance on real-world data.

<u>One-Week Late Submission Policy</u>

Students are encouraged to submit assignments by the deadlines, but we allow submission of an assignment one week after its deadline with 25% penalty on the earned score e.g., if a student earns 100 in a late submission, this score will become 75 when being integrated into the final score.

<u>Group Work and Use of AI tools</u>

Students can collaborate on the assignments but need to clarify the collaborator in submission and be ready to independently defend their submitted answers upon request.

Currently we do not have any regulation on the use of generative AI tools such as ChatGPT in completing assignments, but this may change during the semester to comply with any new university policy. In addition, the whole course is designed with an assumption that students will continue the practice of traditional coding and analysis. It is not allowed to use AI tools during exams.

4. Exam

There are two in-semester written exams which cover the topics of Regression, Classification and Model Enhancement. Both exams will be closed-book and taken in class. Students cannot collaborate on exams, and cannot use AI tools during exams.

There is no final exam. Instead, each student is expected to select and read a research paper published at ICML, NeurIPS, UAI or AISTATS within the recent five years (no earlier than 2018), and submit a two-page review for the paper.

<u>Bonus-Point Presentation</u>

Each student can propose to give a 15-minute in-class presentation of the reviewed paper to earn a bonus point. Students who plan to present must email a formal request to the instructor and receive an approval email by Nov 1, 2023. Under two circumstances, a student may receive a negative point (which will be added to the final score): (1) decides not to present after receiving a formal approval email and (2) delivers a low-quality presentation or demonstrate a poor understanding of the presented paper.

5. Point Distribution

Assignments contribute to 50% of the final score. Each student's lowest score will be dropped in calculation.

Each exam contributes to 20% of the final score. Exam scores may be lightly curved.

Paper review contributes to 10% of the final score.

Bonus-point presentation contributes up to 10% (or down to -10%) of the final score.

In general, a student gets A for earning 90+, B for 80+, C for 70+, D for 60+ and F for 60-.

6. Pre-requisites

Besides official pre-requisites, students are expected to have working knowledge on linear algebra and probability. We will have ad-hoc reviews and exercises on related math subjects in class but they cannot guarantee full understanding of the lectured subjects. Students are also expected to have sufficient programming skills to implement lectured algorithms from scratch in Python.

7. Others

We will handle other issues such as academic misconduct according to the university's policies.