

CS4033/5033: Machine Learning

Fall 2024

Title	Name	Email	Office Hour	Office Location
Instructor	Chao Lan	clan@ou.edu	MW 11am-12pm	DEH 210
TA	Shayan Shafaei	shayan.shafaei@ou.edu	Tu 3-4:30pm	Zoom: 4551860899

1. Course Description

This course introduces the fundamentals of machine learning, with a focus on the mathematical design, interpretation and optimization of predictive models for regression and classification tasks. Topics include linear and kernel regression, classification, dimensionality reduction, clustering, ensemble methods, learning theory and other topics. A tentative schedule of the course is listed below.

Week 1	Introduction & Math Review	
Week 2	Linear Regression	HW1 & HW2
Week 3		
Week 4	Kernel Regression	HW3
Week 5	Topics & Exam 1	
Week 6	Classification	HW4 & HW5
Week 7		
Week 8	PCA & Clustering	HW 6
Week 9	Topics & Exam 2	
Week 10	Ensemble Methods	HW7
Week 11	Learning Theory	HW8
Week 12	Other Topics	HW9 & HW10
Week 13		
Week 14		Presentation Request Due
Week 15	Thanksgiving, No Class	
Week 16	Paper Presentation	Paper Review Due

2. Main Reference

- [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.

3. Assignment

There are 10 assignments, each containing some written and/or programming tasks. Written tasks normally ask students to derive the mathematics of learning techniques, while programming tasks normally ask students to implement learning techniques from scratch. Written tasks must be done using Latex (Overleaf) and programming tasks must be done using Python.

Late Submission Policy

After the deadline of an assignment, students can still submit the work for grading within a week with 25% penalty on the earned score. Submissions one week after the deadline will not be graded.

Group Work and Use of AI tools

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 restriction 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.

4. Exams, Paper Review and Bonus Presentation

There are two closed-book exams during the semester. One focuses on regression while the other focuses on classification. During the exams, students cannot collaborate or use any AI tool.

There is no final exam. Instead, each student needs to select and review a research paper published between 2019-2024 at the main conference (not workshops) of ICML, NeurIPS, UAI or AISTATS. A two-page review report must be submitted by the end of the 16th week. See the schedule for more clarity.

Bonus-Point Presentation

Each student can propose to give a 15-minute in-class presentation of the reviewed paper in the 16th week.

Students who want to present must email a formal request to the instructor and attach a completed version of the presentation slides (which could be modified later) before the official Thanksgiving holiday begins.

All requests will be reviewed by the instructor and at most 9 students will be selected to present.

Below is the grading rubric:

- Any student who submits the request with a completed version of the presentation slides by the deadline will receive 1 point bonus even if he/she is not selected for presentation.
- A selected student can earn up to 5 bonus points depending on the quality of presentation. However, if a student is selected but does not present, he/she gets a -5 bonus points.
- Selection criterion depends on the topic of the paper, the quality of the attached slides and the student's overall performance in the course e.g., an application paper with little math is unlikely to be selected, and a student who misses half of the assignments is unlikely to be selected.

5. Final Grade

Assignments contribute to 70% of the final score. Each student's lowest score will be dropped in calculation.

Each exam contributes to 10% of the final score.

Paper review contributes to 10% of the final score.

Bonus points are directly added to 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-.