

COS 575 - Machine Learning

Salimeh Yasaei Sekeh

University of Maine, Spring semesters

Instructor:

Salimeh Yasaei Sekeh

Required textbook: None.

Recommended texts:

- Murphy, Machine Learning: A Probabilistic Perspective, MIT Press, 2012, available online.
- Hastie, Tibshirani, and Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Springer, Second Edition, available online.
- Bishop, Pattern Recognition and Machine Learning, Springer, 2006.
- Mohri, Rostamizadeh, and Talwalkar, Foundations of Machine Learning, MIT Press, 2012, available online.
- Duda, Hart, and Stork, Pattern Classification, Wiley, 2001.

The pdf format of textbooks will be uploaded to Blackboard.

Additional references:

- Scholkopf and Smola, Learning with Kernels, MIT Press, 2002.
- Mardia, Kent, and Bibby, Multivariate Analysis, Academic Press, 1979 (good for PCA, MDS, and factor analysis).
- Boyd and Vandenberghe, Convex Optimization, Cambridge University Press, 2004.
- Shalev-Shwartz and Ben-David, Understanding machine learning: from theory to algorithms, Cambridge University Press, 2014.
- Sutton and Barto, Reinforcement Learning: An Introduction, MIT Press, 1998.
- DasGupta, Probability for Statistics and Machine Learning, Springer 2011.

Course Prerequisites:

- MAT 126, MAT 127
- STS 232 (or STS 434 or STS 332 or STS 435)

Background Prerequisites:

- Probability: jointly distributed random variables, multivariate densities and mass functions, expectation, independence, conditional distributions, Bayes rule, the multivariate normal distribution.
- Linear algebra: rank, nullity, linear independence, inner products, orthogonality, positive (semi-) definite matrices, eigenvalue decompositions.
- Multivariable calculus: partial derivatives, gradients, chain rule.
- Programming skills in Python.
- It is expected that students will have a good working knowledge of these topics. Students with most but not all of this background should be able to catch up during the semester with some additional effort. Certain topics will be briefly reviewed.

Grading:

- Homework: 35% - Lowest 1 dropped.
- Midterm exam: 20%
- Final exam: 20%
- Project: 25%
- Extra credit: 5-10% to students who answer questions in Blackboard and significantly enhance the course experience through their contributions.
- The grading scale for the final mark is as follows:

Letter Grades	Numerical Range
A	94-100
A-	90-94
B+	87-90
B	84-87
B-	80-84
C+	77-80
C	74-77
C-	70-74
D+	67-70
D	64-67
D-	60-64
F	0-60

Course Schedule:

The table (below) provides the initial distribution of topics discussed over the weeks in the semester. This schedule is tentative and subject to change during the semester at the instruction discretion. All changes will be announced in class or on the course website. Students are responsible for making sure they are informed about announcements.

Week	Class Tue/Thu	Materials
1	26/01 28/01	Syllabus, Introduction on Machine Learning, Real world problems K-nearest neighbor classifier, Linear Algebra Review
2	02/02 04/02	KNN classification on Iris Data set (Python), Bayes Classifier Naive Bayes, Probability Review
3	09/02 11/02 12/02	Linear Discriminant Analysis (LDA) Example in LDA (Python), Logistic Regression (LR) Homework 1 Due Date
4	16/02 18/02	Quadratic Discriminant Analysis (QDA), Cross Validation (CV) CV (Python), Newton Method (Gradient Descent), Regularized LR
5	23/02 25/02 26/02	Linear Regression, MLE OLS Regression, Ridge Regression Homework 2 Due Date
6	02/03 04/03	GD & SGD, Lasso, Bias and Variance, Robust Reg. (Python) MM Algorithm, Separating Hyperplane, SVM
7	09/03 11/03 12/03	Feature Transformation, Kernels Kernel Ridge Regression (KRR), Kernel SVM Homework 3 Due Date
8	16/03 18/03	KRR (Python), Unconstrained Optimization Kernel SVM (Python), Clustering (K-Means)
9	23/03 25/03	Midterm Exam Dimension Reduction Method, Principle Component Analysis (PCA)
10	30/03 01/04 02/04	PCA (Python), Face Detection, Kernel PCA Gaussian Mixture Models, Expectation Maximization Algorithm Project Proposal Due
11	06/04 08/04 09/04	Latent Variable Models Neural Networks, Backpropagation Homework 4 Due Date
12	13/04 15/04	Deep Learning, AlexNet, ResNet, VGG Architectures Autoencoder, Adversarial Machine Learning
13	20/04 22/04 23/04	Generative Adversarial Networks, Recurrent Neural Network Decision Tree, Ensemble Method, Bagging, Random Forest Homework 5 Due Date
14	27/04 29/04 30/04	Image Segmentation, Online Learning Reinforcement Learning Final Project Report Due

Additions between COS 575 - ML and COS 475 - ML courses:

A major addition between these two level machine learning courses is in covering advance topics. For instance in COS 575 the instructor will definitely discuss Generative Adversarial Networks and On-Line learning as they are two top and hot topics in modern Machine Learning applications such as fake deep faces/videos and streaming online recognitions, respectively. In COS 475 machine learning course these very advanced topics will not be covered.

The second main difference between COS 475 and COS 575 is assignments. Due to the undergraduate and graduate levels of students in 400 level and 500 level classes, the assignments for the graduate level includes additional and advanced context while for undergraduate level the assignments includes problems based on covered topics in the class with extensive explanation so that students could establish and improve their understandings of the material discussed in the class through homework problems.

Details:

- **Homeworks:**

Homeworks will be assigned bi-weekly. Applications will be developed through programming exercises, including face recognition, spam filtering, handwritten digit recognition, image compression, and image segmentation. Python is the supported languages in the course. You must complete your programming assignments in this language.

- **Exams:**

You may use two cheat sheets (front and back), and no other materials are allowed. Please notify us the first week of class if you have a conflict.

- **Final Project:**

There will be a final project. You will be asked to form groups of maximum 2 people. The project must explore a methodology or application not covered in the lectures. You will be asked to select a paper on a methodology not covered in class, and implement the method.

- **Blackboard:**

Most questions you have about the course, both logistical and technical, should be posted to Blackboard. Questions about how to solve homework problems are encouraged, but responses should provide hints as opposed to detailed answers. You may indicate that your questions is for instructors only your question is of a sensitive nature or may disclose solutions to the class.

- **Standards of Conduct:**

Please,

Arrive to class on time. If you must enter a class after lecture has clearly begun, please do so quietly. Focus on class material during class time. Sleeping, doing work for another class, checking email, and exploring the internet are unacceptable and can be disruptive. Only use personal electronic devices during class for viewing or taking notes. If you elect to use a laptop during class, please type very quietly. Few things are more annoying than sitting next to someone who is pounding on their keyboard during a lecture. Refrain from eating during class. Avoid audible and visible signs of restlessness. These are both rude and disruptive to the rest of the class. Don't pack your bags to leave until the instructor has dismissed class. Thank you.

- **Collaboration on homeworks:**

Each student will prepare the final write-up/coding of his or her homework solutions without reference to any other person or source, aside from the student's own notes or scrap work. Students may consult classmates for the purpose of brainstorming, but not for obtaining the details of solutions. Under no circumstances may you copy solutions or code from a classmate or other source.

Campus Policies:

- **Academic Honesty Statement:**

Academic honesty is very important. It is dishonest to cheat on exams, to copy term papers, to submit papers written by another person, to fake experimental results, or to copy or reword parts of books or articles into your own papers without appropriately citing the source. Students committing or aiding in any of these violations may be given failing grades for an assignment or for an entire course, at the discretion of the instructor. In addition to any academic action taken by an instructor, these violations are also subject to action under the University of Maine Student Conduct Code. The maximum possible sanction under the student conduct code is dismissal from the University.

- **Students Accessibility Services Statement:**

If you have a disability for which you may be requesting an accommodation, please contact Student Accessibility Services, 121 East Annex, 581.2319, as early as possible in the term. Students who have already been approved for accommodations by SAS and have a current accommodation letter should meet with me privately during the first two weeks of class. All discussions will remain confidential.

- **Course Schedule Disclaimer (Disruption Clause):**

In the event of an extended disruption of normal classroom activities, the format for this course may be modified to enable its completion within its programmed time frame. In that event, you will be provided an addendum to the syllabus that will supersede this version.

<https://umaine.edu/citl/teaching-resources-2/required-syllabus-information/#Schedule>

- **UMaine Student Code of Conduct:**

All students are expected to conform to the UMaine Student Code of Conduct.

<https://www.maine.edu/board-of-trustees/policy-manual/section-501/>

- **Observance of Religious Holidays/Events:**

The University of Maine recognizes that when students are observing significant religious holidays, some may be unable to attend classes or labs, study, take tests, or work on other assignments. If they provide adequate notice (at least one week and longer if at all possible), these students are allowed to make up course requirements as long as this effort does not create an unreasonable burden upon the instructor, department or University. At the discretion of the instructor, such coursework could be due before or after the examination or assignment. No adverse or prejudicial effects shall result to a student's grade for the examination, study, or course requirement on the day of religious observance. The student shall not be marked absent from the class due to observing a significant religious holiday. In the case of an internship or clinical, students should refer to the applicable policy in place by the employer or site.

<https://umaine.edu/citl/teaching-resources-2/required-syllabus-information/#Observance>

- **Sexual Discrimination Reporting:**

<https://umaine.edu/citl/teaching-resources-2/required-syllabus-information/#Reporting-Long>

- **COVID-19 Syllabus Statement:**

Please read the University of Maine COVID-19 Syllabus Statement available on

CITL website (link below):

<https://umaine.edu/citl/2020/08/17/suggested-syllabus-language-for-covid19-is-available/>

1 Course Syllabus

Topics to be covered:

1.1 Introduction

- Overview
- Linear Algebra Review
- Probability Review
- Unconstrained optimization

1.2 Classification

- K-nearest neighbors (KNN)
- Bayes Classifiers
- Linear Discriminant Analysis (LDA)
- Quadratic Discriminant Analysis (QDA)
- The Naive Bayes
- Logistic Regression

1.3 Separating Hyperplanes

1.4 Regression

- Linear Regression
- Least Squares
- Probabilistic Interpretation (connection to MLE)
- Ridge Regression
- Robust Regression

1.5 Empirical Risk Minimization

1.6 Regularization

- L2 Regularization
- L1 Regularization, Sparsity and Feature Selection
- Bias-Variance Tradeoff

1.7 Kernel Methods

- Positive Definite Symmetric (PSD) Kernels
- Kernel Ridge Regression

1.8 Support Vector Machine

- Constrained Optimization Review
- Support Vector Machine (SVM)
- Model Selection

1.9 Unsupervised Learning

- Principle Components Analysis (PCA)
- Clustering, K-Means
- Gaussian Mixture Models
- The Expectation Maximization Algorithm
- Latent variable models

1.10 Neural Networks and Deep Learning

1.11 Decision Trees

1.12 Ensemble methods

1.13 Kernel Density Estimation (KDE)

1.14 Advanced Topics

- Generative Adversarial Networks
- Continual Learning
- Overparameterization neural network
- Learning Theory
- On-Line Learning
- Markov Decision Processes
- Advanced Deep Learning
- Graphical Models
- Reinforcement Learning
- Boosting

- Graphical Neural Networks
- Multi-armed bandits
- Bayesian linear regression
- Multi-armed bandits
- NLDR / Euclidean embedding
- Sparsity and low-rank models
- ...