CSE4/574: Introduction to Machine Learning
Fall 2021

Instructor: Mingchen Gao

Syllabus

Machine learning is concerned with the question of how to make computers learn from experience. The ability to learn is not only central to most aspects of intelligent behavior, but machine learning techniques have become key components of many software systems. For examples, machine learning techniques are used to create spam filters, to analyze customer purchase data, to understand natural language, or to detect fraudulent credit card transactions.

This course will introduce the fundamental set of techniques and algorithms that constitute machine learning as of today, ranging from classification methods like decision trees, support vector machines and neural networks, over structured models like hidden Markov models, to clustering and matrix factorization methods for recommendation. The course will not only discuss individual algorithms and methods, but also tie principles and approaches together from a theoretical perspective, as well as the gaining fundamentals of applying machine learning techniques to real-world problems. In particular, the course will cover the following topics:

Generative Models, Bayesian Learning, Linear Regression, Logistic Regression, Perceptron, Neural Networks, Convolutional Neural Network, Recurrent Neural Network, Reinforcement Learning, Graphical Models, Clustering, Latent Linear Models, Support Vector Machines, Decision Tree, Boosting, Random Forest, Hidden Markov Models.

The prerequisites of CSE4/574 are Linear Algebra, Probability, and Statistics. We are going to use Python for the programing assignments.

Tentative schedule for the topics covered in this course is listed as follows:

- Introduction, chapter 1
- Probability, chapter 2
- Generative models, chapters 3.1-3.3, 4.1
- Naive Bayesian, chapters 3.5, 4.2
- Linear regression, chapter 7, **first assignment**
- Logistic regression, chapter 8
- Optimization
- Perception
- Neural network, **second assignment**
- Convolutional neural network, Recurrent network
- Max margin/support vector Machine, chapter 14.5, **third assignment**
- Kernel method , chapter 14
- K-means clustering, chapter 11.4
- Spectral clustering, chapter 25.4
- Mixture models/expectation maximization, chapter 11
- Factor analysis, chapter 12.1
- PCA, SVD, chapter 12.2

Readings

- **MURPHY: Kevin Murphy, Machine Learning: A Probabilistic Perspective, MIT Press, 2012.**
- BISHOP: Chris Bishop, Pattern Recognition and Machine Learning, Springer, 2006.
- MITCHELL: Tom Mitchell, Machine Learning. McGraw-Hill, 1997.
- HASTIE: Trevor Hastie, Robert Tibshirani and Jerome Friedman, The Elements of Statistical Learning. Springer, 2009.

Grading Scheme

- Grading Scheme
  - Short weekly quizzes using Gradiance (12) – 20%
  - Programming Assignments (3) – 30%
  - Mid-term Exam (in-class, open book/notes) – 20%
  - Final Exam (in-class, open book/notes) – 30%
  - Final grade (*Tentative*):

A	$[92.5, 100]$	C+	$[67.5, 72.5)$
A-	$[87.5, 92.5)$	C	$[62.5, 67.5)$
B+	$[82.5, 87.5)$	C-	$[57.5, 62.5)$
B	$[77.5, 82.5)$	D+	$[52.5, 57.5)$
B-	$[72.5, 77.5)$	D	$[47.5, 52.5)$

- Gradiance quizzes
  - Will be released every Monday at 9.00 AM EST
  - Due next Sunday at 11.59 PM EST
  - Gradiance 0 will not be evaluated (warm up)
- All assignments are electronically due on Wednesdays by 08.59 AM EST through UBLearns.

Accessibility Services and Special Needs

If you have a disability and may require some type of instructional and/or examination accommodation, please inform me early in the semester so that we can coordinate the accommodations you may need. If you have not already done so, please contact the Office of Accessibility Services; 25 Capen Hall; email: stu-accessibility@buffalo.edu Phone: 716-645-2608 (voice); 716-645-2616 (TTY); and on the web at http://www.student-affairs.buffalo.edu/ods/. All information and documentation is confidential.

The University at Buffalo and the School of Engineering and Applied Sciences are committed to ensuring equal opportunity for persons with special needs to participate in and benefit from all of its programs, services and activities.

Academic Integrity

This course will operate with a zero-tolerance policy regarding cheating and other forms of academic dishonesty. Any act of academic dishonesty will subject the student to penalty, including the high probability of failure of the course (i.e., assignment of a grade of F). It is expected that you will behave in an honorable and respectful way as you learn and share ideas. Therefore, recycled papers, work submitted to other courses, and major assistance in preparation of assignments without identifying and acknowledging such assistance are not acceptable. All work for this course must be original for this course. Please familiar yourself with the CSE Departments' and University at Buffalo's Academic Integrity Policy and Procedure outlined at `http://www.cse.buffalo.edu/shared/policies/academic.php`.