

ECE 492-045 Special Topics: Introduction to Machine Learning Course Syllabus (Fall 2019)

Time & Location: MW4:30–5:45 pm, 1226 EB2

Instructor: Dr. Chau-Wai Wong

Office hour: Tu4:30–5:45 pm or by appointment, 2072 EB2

Email: chauwai.wong@ncsu.edu

<https://people.engr.ncsu.edu/cwong9/> [or Google his name]

Teaching Assistant: Runze Liu

Office hour: Th3:15–4:15 pm, 2091 EB2

Email: rliu10@ncsu.edu

Objective: To introduce fundamental concepts in machine learning, and to provide students with hands-on experiences in machine learning applications.

Course Description: With the availability of huge datasets and the recent advancement in computational power, machine learning as a predictive tool has been increasingly successful in virtually all aspects of our life. In order to achieve the best results out of applying such tool, both solid understandings of the underlying principles and hands-on experiences are needed. This course introduces fundamental concepts in class and exposes students to real-world applications via well-guided homework programming problems.

Learning Outcomes: By the end of this course, the students should be able to:

1. Explain basic concepts of and understand mathematical principles behind machine learning algorithms.
2. Derive and implement from scratch basic but fundamental machine learning algorithms such as PCA and Logistic regression.
3. Make informed choices about using off-the-shelf machine learning packages.
4. Effectively visualize and clean data.
5. Use and work across different programming languages when there is a need.

Prerequisites: ST 300-level or above, and ECE 301/CSC 316. Talk to the instructor if not in ECE/CSC.

Topics: Linear statistical models, Bayesian classifiers, neural networks (NN), support vector machine (SVM), classification/decision tree, clustering, principal component analysis (PCA), naive Bayes, topic model, hidden Markov model (HMM).

Textbooks:

G. James, D. Witten, T. Hastie and R. Tibshirani, Introduction to Statistical Learning with Applications in R, Springer, 2013. [Online]

T. Hastie, R. Tibshirani, J. Friedman, The Elements of Statistical Learning, (12th Printing, 2017), 2nd Ed., Springer. [Online]

Reference Books:

E. Alpaydin, Introduction to Machine Learning, 3rd Ed., MIT Press, 2009.

Workload & Grading: There will be weekly homework assignments that contains both written problems and programming problems (40%), two midterm exams (20%×2), and one final exam (20%). Programming will be in Python, R, and Matlab.

Auditing Policy: Students must register this course for auditing. Auditing students must turn in all homework assignments. However, they are not required to take the exams.

Late/Missed Assignments/Exams:

All assignments must be turned in at the beginning of class on the date they are due. A penalty of 20 out of 100 pts per day will be assessed for all homework assignments turned in late. You are expected to turn in your assignments on time for any anticipated absences that you will have so please plan accordingly. Emergency or unanticipated absences will be handled on a case by case basis. There will be no make-up exams. Missing grades for missed exams with valid excuses and documentation will be replaced by the final exam grade. Students who believe they have valid excuses to miss assignments or exams must comply with University Attendance Regulations REG 02.20.03 on NC State's website.

Students with Disabilities:

Reasonable accommodations will be made for students with verifiable disabilities. In order to take advantage of available accommodations, students must register with Disability Services Office. For more information on NC State's policy on working with students with disabilities, please see REG 02.20.01 on NC State's website.

Academic Integrity: All the provisions of the code of student conduct apply to this course as appropriate. See

<https://studentconduct.dasa.ncsu.edu/code/>

Course Evaluation: Students will be notified during the last week of the class to complete course evaluations. All evaluations are confidential; instructors will never know how any one student responded to any question, and students will never know the ratings for any particular instructors.