

CS722/822 Machine Learning
Fall 2018, MW 3:00 pm – 4:15 pm
Location: CONST 2065

Instructor:

Jiangwen Sun, Ph.D., Assistant Professor, CS Department
Office: E&CS 3204, Phone: 683-7712, Email: jsun@odu.edu
Office Hours: W 11:00 am – 12:00 pm (other times by appointment)

Teaching Assistant:

TBD

Prerequisite:

Knowledge of linear algebra, probability theory, calculus and optimization. Basic computer skills and programming experience with one or more following languages: Python (preferred), MATLAB, R, Java and C++.

Course Description:

The objective of the course is to introduce basic and advanced concepts in machine learning to students, and enable them to use machine learning methods in real-life applications, and review state of the art literature in machine learning. This course covers basic topics of machine learning, such as supervised learning (classification, regression, feature selection, etc.), unsupervised learning (clustering, dimension reduction, or component analysis, etc.) and some advanced/emerging topics, such as deep learning, multi-task learning and multi-view data analysis. Usually the emerging topics will be identified by reviewing the latest publications in top venues, such as ICML, NIPS, SIGKDD, etc. Because of the diversity in machine learning topics, the materials covered in this course may vary among semesters.

Approaches:

The course consists of lectures, paper reviews, quizzes, homeworks and projects. Lectures will serve as the vehicle for the instructor to introduce concepts and knowledge to students. Paper reviews are used to inform students of the latest research topics and techniques. Quizzes are used to test if certain basic concepts have been mastered. A course project will be used for students to get profound hands-on experience by programming certain machine learning algorithms identified from the recent literature. Participation in lectures is encouraged during the class. The course may also contain guest lectures from related fields and respective experts.

Students are encouraged to form study groups to facilitate discussion. Each group is expected to consist of two to three students. During paper review process, each group can select a paper from a list provided by the instructor or the recent machine learning venues provided by the instructor (ICML, NIPS, SIGKDD, etc.), study and present the paper in class. As part of the course, the students will work on a term project where each group can choose to implement the algorithms discussed in the paper they choose to present. Each team is required to present in the classroom and submit a project report, which describes whether the algorithm could be replicated in their implementation as stated in the original paper, if not, any potential reason why. This exercise will help the team gain much deeper insights into certain algorithms and promote collaborations among team members. Works beyond that described above will be considered for extra credit, such as applying the algorithm to datasets (problems) not included in the original paper, comparing with more relevant algorithms, etc.

Topics and Tentative Schedule:

| Week | Topics | Note |
|--------------------------------|--|------------------------------------|
| Week 1 (Aug. 27 ~ Aug. 31) | Introduction: supervised and unsupervised learning problems, different learning tasks, etc. Review of basic mathematics (matrix computation, norms, probability, mean, variance, etc.) | |
| Week 2 (Sep. 3 ~ Sep. 7) | Continue review of basic mathematics Supervised learning – regression: overview, least square, overfitting, regularization (ridge, Lasso), probabilistic interpretation, (generalized) gradient descent | No class on Sep. 3 (Labor Day) |
| Week 3 (Sep. 10 ~ Sep 14) | Continue regression | |
| Week 4 (Sep. 17 ~ Sep 21) | Continue regression Supervised Learning – classification: overview, KNN, model evaluation, linear discriminant analysis, logistic regression, support vector machine, neural network | |
| Week 5 (Sep. 24 ~ Sep. 28) | In-class quiz (45 minutes) Continue classification | |
| Week 6 (Oct. 1 ~ Oct. 5) | Continue classification | |
| Week 7 (Oct. 8 ~ Oct. 12) | Continue classification | No class on Oct. 8 (Fall holiday) |
| Week 8 (Oct. 15 ~ Oct. 19) | Continue classification | |
| Week 9 (Oct. 22 ~ Oct. 26) | In-class quiz (45 minutes) Paper review: student teams present papers from the latest ML venues | |
| Week 10 (Oct. 29 ~ Nov. 2) | Paper review: student teams present papers from the latest ML venues Unsupervised learning – clustering: overview, K-means, hierarchical clustering, DBSCAN, clustering evaluation, spectral clustering | |
| Week 11 (Nov. 5 ~ Nov 9) | Continue clustering | |
| Week 12 (Nov. 12 ~ Nov. 16) | Continue clustering | |
| Week 13 (Nov. 19 ~ Nov. 23) | Continue clustering | No class on Nov. 21 (Thanksgiving) |
| Week 14 (Nov. 26 ~ Nov. 30) | In-class quiz (45 minutes) Unsupervised learning – dimension reduction: PCA, CCA | |
| Week 15 (Dec. 3 ~ Dec. 7) | Advanced topics – deep learning | |
| Additional time | Project presentations: Student teams present their final term projects | |
| Final Week | A final comprehensive quiz | |

Textbooks:

- Introduction to Data Mining by Pang-Ning Tan, Michael Steinbach, and Vipin Kumar, ISBN-10: 0321321367
- Pattern Recognition and Machine Learning (Information Science and Statistics) by Christopher M. Bishop, ISBN-10: 0387310738
- Deep Learning (Adaptive Computation and Machine Learning Series) by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, ISBN-10: 0262035618
- Pattern Classification (2nd Edition) by Richard O. Duda, Peter E. Hart and David G. Stork, ISBN-10: 0471056693

None of the textbooks will be required. However, having one or two of them may complement and expand the materials discussed in lectures. Lectures will come with slide files and tutorial/review papers for students to study after lectures.

Attendance Policy:

Students are expected to attend classes regularly.

Grading Policy:

- Close-book quizzes (3 in-class quizzes and 1 final quiz): 40%
- Paper review and presentation (1): 10%
- Non-programming homework assignment (2-3): 5%
- Programming-involved homework assignment (3): 15%
- Term Project (1): a team can only consist of two to three persons, 30% (implementing selected algorithms from a recent publication)

Drop Policy:

As per University guidelines. See the University Calendar for drop dates.

Disabilities:

Old Dominion University is committed to ensuring equal access to all qualified students with disabilities in accordance with the Americans with Disabilities Act. The Office of Educational Accessibility (OEA) is the campus office that works with students who have disabilities to provide and/or arrange reasonable accommodations.

- If you experience a disability which will impact your ability to access any aspect of my class, please present me with an accommodation letter from OEA so that we can work together to ensure that appropriate accommodations are available to you.
- If you feel that you will experience barriers to your ability to learn and/or testing in my class but do not have an accommodation letter, please consider scheduling an appointment with OEA to determine if academic accommodations are necessary.

The Office of Educational Accessibility is located at 1021 Student Success Center and their phone number is (757)683-4655. Additional information is available at the OEA website: <http://www.odu.edu/educationalaccessibility/>

Honor Code:

Students are expected to follow the ODU Honor Code for all assignments and exams. Any violations will be dealt with strictly according to university policy. Despite that this course requires a lot of interaction, and thus discussions of ideas are encouraged, **the work that you turn in must be your own.**

Academic Dishonesty:

Old Dominion University is committed to students' personal and academic success. In order to achieve this vision, students, faculty, and staff work together to create an environment that provides the best opportunity for academic inquiry and learning. All students must be honest and forthright in their academic studies. Your work in this course and classroom behavior must align with the expectations outlined in the Code of Student Conduct, which can be found at www.odu.edu/oscai. The following behaviors along with classroom disruptions violate this policy, corrupt the educational process, and will not be tolerated.

Cheating: Using unauthorized assistance, materials, study aids, or other information in any academic exercise.

Plagiarism: Using someone else's language, ideas, or other original material without acknowledging its source in any academic exercise.

Fabrication: Inventing, altering or falsifying any data, citation or information in any academic exercise.

Facilitation: Helping another student commit, or attempt to commit, any Academic Integrity violation, or failure to report suspected Academic Integrity violations to a faculty member.

Academic dishonesty will be reported to the Office of Student Conduct & Academic Integrity and may result in sanctions up to and including expulsion from the University.