Advanced Machine Learning

DATA 442/642

Fall 2021

Instructor:	Zois Boukouvalas	**Time:**	T 5:30 PM - 8:00 PM
Email:	boukouva@american.edu	**Room:**	Don Myers Building - 119
Office:	Don Myers Building - 222		

Office hours: After class, or T 4:00pm - 5:30pm via Zoom, or by appointment. Always feel welcome to stop by my office hours. You are also encouraged to ask me questions online via email. If you are having **ANY** trouble with the class, please come see me about it as soon as possible. **Do not wait until it is too late!**

Course Pages: I will use Canvas (https://american.instructure.com/) to post any supplementary materials, suggested readings/practice exercises, assignments, and announcements. Sometimes I may also use my personal website (https://zoisboukouvalas.github.io/).

Materials:

- Machine Learning A Bayesian and Optimization Perspective, Sergios Theodoridis, Elsevier. (2nd Edition) **(Required)**
- Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow. **(Recommended)**
- Mathematics For Machine Learning, Deisenroth, Faisal, Ong, Cambridge University Press. **(Recommended)**

Software: The main software for this course is Python. In addition, I will introduce Matlab since many algorithms and demonstrations from the book are in Matlab. If you do not know Python yet, (https://learnpython.org) and (https://docs.python.org/3/tutorial) are good places to start. Matlab is free for AU students and faculty. For more information on how to use Matlab using your AU account visit (https://www.mathworks.com/academia/tah-portal/american-university-40555344.html). The Youtube videos (https://www.youtube.com/playlist?list=PLgEkkJcwLzjeL2Ywe0g66GaggDQxmV05p) provide a great introduction to programming with Matlab.

Course description: Advanced machine learning will be an upper level machine learning course focusing on recent advances in machine learning such as Tensor/Matrix Factorization Methods, Neural Networks and Deep Learning, Probabilistic Graphical Models, and Bayesian Learning. The course will cover theoretical foundations as well as methods used in the machine learning discipline that have achieved a considerable attention and led to a great deal of industrial and academic interest. The course will introduce the mathematical definitions of the relevant machine learning models and derive their associated optimization algorithms. It will cover a range of applications of advanced machine learning in natural language processing, image processing, data fusion, and machine translation. Through this course students will have the chance to design and build a research project that demonstrates how computational learning algorithms can solve difficult tasks in research areas they are interested in.

Pre-requisites: STAT 427/627 "Statistical Machine Learning" or equivalent.

Course plan: We will cover the following topics.

- A tour of the classics, basic concepts of linear algebra, optimization, and their connection to machine learning.
- Sparsity aware learning.
- Learning in reproducing Hilbert spaces.
- Bayesian learning and the Expectation Maximization algorithm.

- Neural networks and deep learning.
- Latent variable methods and introduction to multi-modal learning.

Class structure: This class will be a blend of lecture, class discussion, and labs. I want you all to be involved during class and please do not hesitate to ask questions whenever something is unclear to you. You are expected to attend all class meetings, as I believe that attending class regularly contributes greatly to your performance in the course. It is understandable that you may have to miss class on a rare occasion. You are responsible for any assignments or papers given out during any missed class.
Data scientists must learn to discover solutions for themselves. You should expect to have to research (use Google, stackoverflow, etc) to do your assignments. All you need to do the assignment will NOT have been provided to you in the lectures and course book. This is an essential part of becoming a data scientist!

Assignments & Grading:

Assignments (20%): During the semester I will assign, collect, and grade assignments. There will be five formal assignments throughout the semester. You may receive assistance from other students in the class and me, but your submissions must be composed of your own thoughts, coding and words. I expect you to get ideas from online resources such as stackoverflow or github when you get stuck. Please cite your source when you do so and be specific about what you have added to it. **I will not accept late assignments.**

Labs (20%): 45-minute labs at the end of each class. Each lab covers the material of the lecture.

Exams (20%): We will have two exams. No make-up exams will be given unless you have an extremely compelling excuse such as observance of a religious holiday (in which case you need to let me know in advance) or a documented medical emergency. You will have three days to complete each exam.

Reading quizzes (15%): We will have six take home reading quizzes.

(20% Poster/Presentation + 5% Project Proposal): You will have to prepare a project using the tools and methods learned in the class. You are expected to submit a mid-semester research proposal in order to get your topic approved. The class projects will be presented as a poster presentation. You should prepare a poster, and be prepared to give a very short explanation (10 minutes), in front of the poster, about your work. At the poster session (online), you'll also have an opportunity to see what everyone else did for their projects. You will also need to submit your poster as a PDF the day before the presentation. There are three type of projects:

- **Application project**. Pick an application that interests you, and explore how best to apply learning algorithms to solve it.
- **Algorithmic project**. Pick a problem or family of problems, and try to develop a novel variant of an existing algorithm, to solve it.
- **Theoretical project**. Prove some interesting/non-trivial properties of a new or an existing learning algorithm.

Please schedule a meeting with me if you would like to see or discuss your grade at any point during the semester.

Important dates:

Midterm	October 12, 2021
Project Proposal	October 26, 2021
Project Presentations	December 7, 2021
Final Exam	TBA

Emergency preparedness: In the event of an emergency, students should refer to the AU Web site `http://www.american.edu/emergency` and the AU information line at (202) 885-1100 for general university-wide information. In case of a prolonged closure of the University, I send updates to you by email and will post all announcements on Canvas.

Support services: A wide range of services is available to support you in your efforts to meet the course requirements.

1. Mathematics & Statistics Tutoring Lab (Don Myers Building) provides tutoring in Intermediate Mathematics and Statistics. http://www.american.edu/cas/mathstat/tutoring.cfm

2. Academic Support and Access Center offers study skills workshops, individual instruction, tutor referrals, Supplemental Instruction, writing support, and technical and practical support and assistance with accommodations for students with physical, medical, or psychological disabilities. Writing support is also available in the Writing Center, Battelle-Tompkins 228.

3. Center for Diversity & Inclusion (X3651, MGC 201) is dedicated to enhancing LGBTQ, Multicultural, First Generation, and Women's experiences on campus and to advance AU's commitment to respecting & valuing diversity by serving as a resource and liaison to students, staff, and faculty on issues of equity through education, outreach, and advocacy.

4. The Office of Advocacy Services for Interpersonal and Sexual Violence (X7070) provides free and confidential advocacy services for anyone in the campus community who is impacted by sexual violence (sexual assault, dating or domestic violence, and stalking).

Additional notes:

1. I expect you to be courteous to me and your fellow classmates during the online lectures/meetings.

2. Please let me know during the first week of classes if you have any special needs that require accommodations.

3. Please be sure that you are familiar with AU's Academic Integrity Code, as I am required to report any cases of academic dishonesty to the dean of CAS. For your review: http://www.american.edu/academics/integrity/.