Course Syllabus: COMPSCI 689
Machine Learning – Spring 2023

Instructor: Brendan O'Connor
Office Hours: Wed 3-4pm, CS 238 + zoom
Contact: Piazza private message

First Class: Feb 7, 2023
Time: Tu/Th 10:00-11:15am
Location: CS room 142

Webpage: https://people.cs.umass.edu/~brenocon/cs689_2023/

Course Description: Machine learning is the computational study of artificial systems that can adapt to novel situations, discover patterns from data, and improve performance with practice. This course will cover the mathematical foundation of supervised and unsupervised learning. The course will provide a state-of-the-art overview of the field, with an emphasis on implementing and deriving learning algorithms for a variety of models from first principles. 3 credits.

Course Goals and Required Background: Machine learning at the PhD level aims to prepare students to participate in machine learning research. It requires both strong mathematical foundations and the ability to implement algorithms with a high degree of precision and computational efficiency. Specifically, the course requires a solid undergraduate-level background in linear algebra, vector calculus, multi-variate probability, and Python programming. Students will also need to be comfortable learning new programming libraries and programming languages on their own.

Learning Outcomes: The primary learning outcomes in this course are as follows:

1. To be able to correctly define core supervised and unsupervised machine learning tasks.
2. To gain familiarity with optimization-based frameworks for developing machine learning algorithms including exact and iterative non-linear numerical optimization methods.
3. To gain familiarity with select classical machine learning models and algorithms including what task they solve, how they are derived, and what their computational scalability properties are.
4. To be able to develop customized machine learning models and prediction and/or inference methods for solving a specified task within multiple generalized modeling frameworks including neural networks and generalized probabilistic models.
5. To be able to explain concepts including generalization, capacity control, overfitting, training error, validation error, test error and to understand when to apply different approaches and metrics for evaluating the performance of machine learning methods.
6. To be able to design valid machine learning experiments for assessing model performance including while optimizing model hyper-parameters.
7. To be able to clearly communicate the results of machine learning experiments using appropriately selected and correctly formatted tables and plots and to be able to correctly interpret such output.
8. To gain practical experience developing efficient and numerically stable implementations of learning and inference methods for both classical and customized models.
9. To gain practical experience using current machine learning software and tools based on Python including Numpy, SciPy, and PyTorch and to understand key facilities of these tools such as automatic differentiation.

10. To understand the limitations of optimization-based machine learning.

Instructional Modalities: In spring 2023, COMPSCI 689 will be hybrid with both a traditional and online (UWW) section. The traditional in-person class meeting will be live-streamed on Zoom for online participants. Slides, demo code, and video recordings will be available to students and will generally be posted by the end of the day following each lecture. Assignment submission will be online. Exams will be in-person for the normal section, and completed remotely, with timing, for online.

Course Materials and Technologies: The course will use a freely available textbook and a variety of UMass supported educational technologies.

- **Primary Textbook:** *Machine Learning: A Probabilistic Perspective* by Kevin Murphy, 2012. (Free for UMass students via the UMass Library; see Piazza.)
- **Course Website:** Piazza will be the main course website. See link on public course webpage. It will include lecture notes, assignments, readings, demos, and links to online resources.
- **Course Announcements and Discussion Boards:** Official announcements for the course will go out through Piazza. Piazza will also host discussion forums for the class. Students will be added to Piazza by course staff at the start of the semester.
- **Assignment Submission:** Assignments will be submitted using Gradescope. Students will be added to Gradescope by course staff prior to the release of the first assignment.
- **Lecture Videos** will be accessible via Piazza.

Coursework and Grading Plan: Coursework will consist of four homework assignments (worth 12.5% each), a midterm exam (worth 25%), and a final exam (worth 25%). These items are described in more detail below. The likely letter grade mapping for the course is shown below to the right.

Highest	Lowest	Letter
100.00 %	93.00 %	A
92.99 %	90.00 %	A-
89.99 %	87.00 %	B+
86.99 %	83.00 %	B
82.99 %	80.00 %	B-
79.99 %	75.00 %	C+
74.99 %	70.00 %	C
69.99 %	0.00 %	F

- **Homework:** Homework assignments will generally be split into two parts. The first part of each assignment will consist of derivations and other written problems. The second part of each assignment will consist of implementation, experimentation, and/or analysis problems. There will be a total of two weeks to complete each assignment. The first part of each assignment will be due at the end of the first week of each assignment period, after which solutions to the first part of the assignment will be distributed. Solutions to the second part of each assignment will be due at the end of the second week of each assignment period. Scans of clear, handwritten solutions are acceptable. Typeset solutions are encouraged. Code will be submitted to Gradescope as Python files (not notebooks). If any autograding is used within Gradescope, student code must run within the Gradescope environment to receive credit assigned through autograding. Students should expect to spend around 10 hours for each part of each homework assignment. Time needed to revise background material is in addition to this time estimate.
- **Exams:** There is one midterm and one final exam, both two hours long. The midterm take place 7-9pm, likely in early April. The final exam will be scheduled by the university during the final exam period at the end of the semester (see SPIRE). Students should not make travel plans to leave campus before the end of the final exam period until the final exam for the course has been scheduled. The midterm exam will cover material up to mid-semester. The final exam will be non-cumulative. It will cover material

from mid semester until the end of the semester. Exams are nearly closed-book; for each exam, a student is allowed "cheat sheet" of notes, front and back of a normal (letter or A4) sized sheet of paper. All electronic devices are prohibited in exams. In-person section students must attend exams in person. Online section students must complete exams under the same conditions.

Course Policies: Students should make sure they are familiar with all course policies and the relevant University policies linked to below. By staying enrolled in this course, students agree to be bound by all applicable policies.

- **Course Community Code of Conduct:** The instructor and the course staff are committed to providing a friendly, safe and welcoming environment for all, regardless of gender identity and expression, sexual orientation, disability, personal appearance, body size, race, ethnicity, age, religion, nationality, or other similar characteristic. Please be courteous, respectful, and professional in all of your interactions with other students, TAs, and graders in all mediums of communication including but not limited to in-person, email, video meetings, chat, discussion forums, and re-grade submissions.

  Demeaning, insulting or harassing any member of the course community over any medium of communication is not acceptable behavior, including in person, through official course platforms and through personal/private platforms (social media, email, DM, text, etc.). Students who engage in such behavior will be warned at most once before the behavior is reported to the Dean of Students office. If you feel you have been or are being harassed or made uncomfortable by a member of this course community, please contact a member of the course staff immediately (or if you do not feel safe doing so, contact the Chair of the Faculty of CICS, Erik Learned-Miller (chair@cs.umass.edu), or the Dean of Students office). Whether you've been at UMass for years or are a newcomer, we care about making this course a safe and welcoming place for all.

- **Accommodation Statement:** The University of Massachusetts Amherst is committed to providing an equal educational opportunity for all students. If you have a documented physical, psychological, or learning disability on file with Disability Services (DS), you may be eligible for reasonable academic accommodations to help you succeed in this course. If you have a documented disability that requires an accommodation, please notify the instructor within the first two weeks of the semester so that we may make appropriate arrangements.

- **Class Attendance Policy:** Students are encouraged to attend class meetings unless they are feeling even mildly ill. Students who are not able to attend class meetings for any reason can stay up to date by watching lecture videos and reading posted course notes, slides, and course announcements. This course does not have a graded participation component.

- **Exam Absence Policy:** A makeup exam time will be provided to students who are unable to attend regular scheduled exams according to University policy (e.g., in the case of illness, religious observances, official University travel, and other extenuating circumstances). Note that a makeup exam will be provided in the case of official University travel (e.g., to present at a conference), but conflicting research deadlines (e.g., a conference paper submission deadline) are not grounds for requesting a makeup exam. When students are aware of an exam conflict ahead of time, they should contact the course staff in writing as soon as possible and no later than one week before the exam date to arrange a time for a makeup exam. In the case of illness or other unforeseen extenuating events, students should contact course staff in writing when they are able. Providing documentation for unforeseen exam absences is greatly appreciated.

- **Late Homework Policy:** Assignments must be submitted by the indicated due date in order to count for credit unless there are extenuating circumstances warranting an extension or exemption. Potential circumstances warranting an extensions or exemption include illness, religious observances, official university travel, and other extenuating circumstances. Job-related issues are not acceptable circumstances. Students with known conflicts with assignment submission deadlines should contact the instructor in writing as soon as possible to inquire about the possibility of an accommodation. In the case of illness or other unforeseen extenuating events, students should contact course staff in writing when they are able. Providing documentation for unforeseen extenuating events is greatly appreciated. An accommodation may take the form of allowing a late submission or waiving the requirement to complete all or part of an assignment. The type of accommodation given is at the instructor's discretion.

- **Re-grading Policy:** Errors in grading can occur despite the best efforts of the course staff. If you believe you've found a grading error, please submit a re-grade request. Re-grade requests must be submitted no later than one week after the graded material is returned. Note that re-grading may result in your original grade increasing, decreasing or remaining unchanged as appropriate.

- **Homework Collaboration Policy:** Homework reports and code are considered individual work. You may discuss the problems with other students, but your report and code must be your own work. The list below describes example scenarios that are not permitted and are very likely to result in violations of the course's academic honesty policy:
  - Sharing completed or in-progress reports or code with another student
  - Copying all or part of completed or in-progress reports or code requested from another student
  - Using another student's code or code output to help debug your code
  - Working with another student while you both write-up or program a solution to a problem
  - Posting completed solutions on public code repositories (e.g., GitHub, etc.) during or after the course

- **Academic Honesty Policy:** Examples of cheating include:
  - Copying any solution materials (derivations, code, method descriptions) in whole or in part from external sources or from other students is considered cheating. Examples of external sources include books, web pages, homework "help" services such as Chegg, automated generation such as ChatGPT, etc.
  - Sharing your code, code output, or solutions with other students is also considered cheating.
  - Using another student's code or code output to debug your code is considered cheating. This includes sharing material with students in future years, publicly posting solutions or code, or requesting materials from students who took the course in previous years.
  - Any collaboration practice that results in solution reports or code appearing to be copied from other students (e.g., collaboration indistinguishable from copying) will be investigated as potential cheating.
  - Detected cheating on homework assignments will result in a grade penalty on the assignment up to and including the full value of the assignment. Any misuse autograding software or other course platforms is grounds for an F in the course. Cheating on exams is grounds for an F in the course.

    All instances of suspected cheating will be dealt with through official UMass Amherst Academic Honesty Procedures. Students are expected to be familiar with the relevant policies and procedures: http://www.umass.edu/honesty/. Any person who has reason to believe that a student has committed academic dishonesty should bring such information to the attention of the appropriate course instructor as soon as possible. Instances of academic dishonesty not related to a specific course should be brought to the attention of the appropriate department Head or Chair. Since students are expected to be familiar with this policy and the commonly accepted standards of academic integrity, ignorance of such standards is not normally sufficient evidence of lack of intent.

http://www.umass.edu/dean_students/codeofconduct/acadhonesty/

- **Course Material Intellectual Property Policy:** The instructor and the University share intellectual property rights for all course materials including lecture slides, lecture audio/video recordings, demo code, assignment handouts, and exam materials. Students are allowed to keep copies of this material for personal use, but are prohibited from distributing it to other individuals and/or posting it in part or in whole on publicly accessible sites including on slide share sites and sites such as Chegg. Students are not permitted to make their own lecture recordings (audio or video). Official recordings will be made available to all students.
- **Course Communication Policy:** Clarification questions about course material or assignments and logistics questions should be submitted to Piazza as public posts. Questions about your solutions to homework assignments and other personal matters should be submitted to Piazza as private posts, which will be viewable by course staff (instructor and graders). For private questions about any personal matters that you do not want to be viewed by any other course staff, please email the instructor at breno-con@cs.umass.edu. Homework grade requests will be submitted on Gradescope. We aim to answer Piazza posts and emails within 24 hours outside of weekends.
- **SAT/Fail Request Policy:** Students must request SAT/Fail grading on Piazza by the last day of classes.

Approximate Schedule: (Subject to change over the semester)

Lecture	Topics
Lecture 1	Course Overview: Supervised and Unsupervised Learning
Unit 1: Optimization-Based Supervised Learning	
Lecture 2	Linear Regression
Lecture 3	Linear Regression and Generalization
Lecture 4	Linear Classification and Logistic Regression
Lecture 5	Numerical Optimization
Lecture 6	Basis Function Expansion and Regularization
Lecture 7	Regularization and Support Vector Machines
Lecture 8	Optimization for Support Vector Machines
Lecture 9	Supervised Learning Experiment Design
Lecture 10	Neural Networks (I)
Lecture 11	Neural Networks (II)
Lecture 12	Optimization for Neural Networks
Lecture 13	Neural Network Architectures
Lecture 14	Probabilistic Supervised Learning (I)
Lecture 15	Probabilistic Supervised Learning (II)
Lecture 16	PyTorch
Unit 2: Optimization-Based Unsupervised Learning	
Lecture 17	Joint Probability Models
Lecture 18	Probabilistic Mixture Models
Lecture 19	Latent Linear Models
Lecture 20	Autoencoders
Lecture 21	EM and Direct Gradients
Unit 3: Advanced Topics	
Lecture 22	Transformers (I)
Lecture 23	Transformers (II)
Lecture 24	Markov Chain Monte Carlo Methods

Lecture 25	Advanced Topics
Lecture 26	Advanced Topics
Lecture 27	Course Wrap-Up