

5630/6630/6636 Syllabus

1 Course Objective

This course aims at providing students basic concepts and popular algorithms in machine learning (ML) and modern AI. By the end of this course, the student will **(1)**. demonstrate proficiency in concepts, techniques, and applications of machine learning **(2)**. demonstrate proficiency in software development for machine learning **(3)**. demonstrate proficiency in oral and written communication in the context of machine learning projects **(4)**. apply machine learning algorithms By the end of the course, 6000-level students will **(5)**. demonstrate critical thinking and analytical skills through a summary and evaluation of state-of-the-art research reported in the literature

2 Basic Information

- **Modality: Online.**
- Instructor: Bo Liu (boliu@auburn.edu)
- Time: 1:00-1:50pm on MWF
- Lecturer Office Hour: TBD on Friday (starting from the 2nd week) **on Zoom.**
- **Zoom ID: 283-127-3778**
- **Virtual Classroom: <https://auburn.zoom.us/j/2831273778>**

- TA: Liangliang Xu (lzx0014@auburn.edu)
- TA Office Hour: 3:30 - 5:00 pm Tuesday/Thursday. (starting from the 2nd week) **on Zoom.**
- **Zoom ID: 928 861 9894**
- **Virtual TA Office: <https://auburn.zoom.us/j/9288619894>**

3 Textbook

- Machine Learning: a Probabilistic Perspective, Kevin Patrick Murphy. MIT Press (**required**)

- Machine Learning, Tom M. Mitchell (optional)
- Introduction to Machine Learning, Third Edition. Ethem Alpaydin. MIT Press (optional)
- The Elements of Statistical Learning, T. Hastie, R. Tibshirani, and J. Friedman (optional)
- A Course in Machine Learning, Hal Daume III (available online <http://ciml.info/>) (optional)

4 Grading

- Weekly homework assignments (20%). One every week for a total of about 10 for the semester (2% each). The first homework will not be graded but is required. The homework is due on the date the homework is listed on the schedule.
- Midterms (30%). Two midterms in total, and each accounts for 15%. Closed-book, in-class, and your only helper is your one-page, A-4 size, hand-written cheatsheet.
- Mini-projects (30%). Two in total (Each 12.5%) to be done individually containing both theory and programming aspects.
- Final project (20%). To be done in groups of two or more. Each team will submit an initial proposal, make a mid-term and final presentation, and submit a written report. Details will be posted in due course.
- Class participation (5%). Questions/answers in class, etc. to encourage active class involvement. The maximum bonus point per class is 1 point, and the total is capped with 5 points. Remote students automatically have this participation part and are graded on a different track.
- **There is no final exam (0%).**

Grading Scheme:

- A: Above top 20% or >90
- B: Above top 40% or >80
- C: Above top 70% or >70
- D: Above top 90% or >60
- F: <= 60

5 Topics

Major topics:

There are four basic tasks of machine learning, aka, *representation, prediction, inference, and decision-making*. This class will cover 2 out of 4 topics (aka, all but decision-making and inference).

- **Representation: Ground-true Data is Dirty Data**
 - * Dimensionality reduction for data denoising
- **Data! Data!**
 - “*We don’t have better algorithms than anyone else; we just have more data,*” *Google’s Chief Scientist Peter Norvig*
 - * Decision Tree
 - * Linear Model, Generalized Linear Model
 - * Naive Bayes Classification
 - * SVM and Kernel Learning
 - * ANN and Deep Learning
- **Advanced Topics: Meta-Methods**
 - “*Three ordinary minds outwit an outstanding one,*” *Chinese proverb*
 - * Beyond binary classification
 - * Ensemble methods
 - * Adaboost and boosting (optional)

Minor topics (Optional):

- Reasoning under uncertainty
- Graphical Models

6 Late Penalties

All assignments are due at the beginning of class on the due date. One (1) day late, defined as a 24-hour period from the deadline (weekday or weekend), will result in 30% of the total points for the assignment deducted. So, for example, if a 120-point assignment is due on a Wednesday at 8 am, but it is handed in between Wednesday 8:01 am and Thursday 8 am, 36 points will be deducted from the point you received from this homework or project. A total of two (2) free late days may be used throughout the semester without penalty. You can choose to use zero, one, or two free late days toward anyone late homework, depending on availability. You must specify on the first page of your late homework how many free late days to use.

7 Recommended Pre-requisite

- COMP 3270 equivalent (you should be very familiar with algorithm basics, understanding, design, synthesizing, and analyzing techniques.)
- COMP 5600/6600 or equivalent (you should be very familiar with kNN, k-Means, and decision-tree algorithms).
- Programming language: Python.
- Maths: undergrad-level calculus, linear algebra, statistics.
- Welcomed (not required): functional analysis, continuous optimization.

8 Academic Honesty Policy

Auburn University's academic honesty policy applies to all assignments and exams for this class. The full-text of the academic honesty policy can be found at: [honesty](#) In addition to the general guidelines mentioned in the above policy, for this course, I require that: In examinations, you must work individually with no communication with others and use only materials/tools that have been explicitly allowed. For homework, you may discuss problems with your no more than three colleagues, but final solutions need to be worked out, written, and submitted individually. Any external material used should be clearly cited. In your writings (example term papers, homework solutions, proposals, etc.), no more than one or two sentences may be used verbatim from any source. **READ THESE INSTRUCTIONS CAREFULLY!** If any aspect of the academic honesty policy and guidelines for this course are unclear, please ask me for clarification. Lack of awareness or understanding of this policy will not be an acceptable excuse or defense against disciplinary action.

9 Related courses

Here is a list of similar courses offered at various places:

- Machine learning at UMD by Hal Daume III
- Machine learning at CMU by Tom Mitchell
- Machine learning at UW by Carlos Guestrin
- Machine learning at UW by Luke Zettlemoyer

10 Academic Contingency Policies

If normal class are disrupted due to illness, emergency, or crisis situation (such as an H1N1 flu/Covid-19 outbreak), the syllabus and other course plans and assignments may be modified to allow completion of the course. If this occurs, an addendum to your syllabus and/or course assignments will replace the original materials.