Welcome

Instructor:
Charles Isbell, **isbell@cc.gatech.edu (mailto:isbell@cc.gatech.edu)**
Dean's Suite, College of Computing Building

Creators of Online Material:
Prof. Charles Isbell, **isbell@cc.gatech.edu (mailto:isbell@cc.gatech.edu?subject=CS%207641)**
Prof. Michael Littman, **mlittman@cs.brown.edu (mailto:littman@cs.brown.edu?subject=CS%207641)**

TAs:
Dan Boros, **boros@gatech.edu (mailto:boros@gatech.edu)**
John Mansfield, **jmansfield6@gatech.edu (mailto:jmansfield6@gatech.edu)**

Office Hours:
See Ed Discussion for details.

Required Text:
Machine Learning
by Tom Mitchell, McGraw Hill, 1997

Video Lectures:
"Modules" in Canvas
Beta: "Ed Lessons" in Canvas

General Information

Machine Learning is a three-credit course on, well, Machine Learning. Machine Learning is that area of Artificial Intelligence that is concerned with computational artifacts that modify and improve their performance through experience. The area is concerned with issues both theoretical and practical. This particular class is a part of a series of classes in Machine Learning and takes care to present algorithms and approaches in such a way that grounds them in larger systems. We will cover a variety of topics, including statistical supervised and unsupervised learning methods, randomized search algorithms, Bayesian learning methods, and reinforcement learning. The course also covers theoretical concepts such as inductive bias, the PAC and Mistake-bound learning frameworks, minimum description length principle, and Ockham's Razor. In order to ground these methods, the course includes some programming and involvement in a number of projects.

Objectives

There are four primary objectives for the course:

- To provide a broad survey of approaches and techniques in ML
- To develop a deeper understanding of several major topics in ML
- To develop the design and programming skills that will help you to build intelligent, adaptive artifacts
- To develop the basic skills necessary to pursue research in ML

The last objective is the core one: you should develop enough background that you can pursue any desire you have to learn more about specific techniques in ML, either to pursue ML as a research career or to apply ML techniques in other research areas in interesting (as opposed to uninteresting) ways.

Prerequisites

The official prerequisite for this course is an introductory course in artificial intelligence. In particular, those of you with experience in general representational issues in AI, some AI programming, and at least some background (or barring that, willingness to pick up some background) in statistics and information theory should be fine. Any student who did well in an introductory AI course should be fine. You will note that most semi-modern AI courses suggests at least some tentative background in some machine learning techniques as well. Having said all that, the most important prerequisite for enjoying and doing well in this class is your interest in the material. I say that every semester and I know it sounds trite, but it's true. In the end, it will be your own motivation to understand the material that gets you through it more than anything else. If you are not sure whether this class is for you, please talk to me.

Resources

- **Readings.** The textbook for the course is *Machine Learning* by Tom Mitchell. We will follow the textbook quite closely for most of the semester, so it is imperative that you have a copy of the book. We will also use supplemental readings as well, but those will be provided for you.
- **Computing.** Even though you absolutely will not need it, you will have access to CoC clusters for your programming assignments. You can test your code on the Shuttles cluster, using your GT username and password to log in - you will not need a CoC account for this course. More info can be found here: https://support.cc.gatech.edu/facilities/general-access-servers (https://support.cc.gatech.edu/facilities/general-access-servers) .
- **Web.** We will use the class canvas page and Ed Discussions to post last-minute announcements, so check them early and often.

Statement of Academic Honesty

At this point in your academic careers, I feel that it would be impolite to harp on cheating, so I won't. You are all adults, more or less, and are expected to follow the university's code of academic conduct (you know, **the honor code** (http://www.honor.gatech.edu/)). Furthermore, at least some of you are researchers-in-training, and I expect that you understand proper attribution and the importance of intellectual honesty.

Please note that unauthorized use of any previous semester course materials, such as tests, quizzes, homework, projects, videos, and any other coursework, is prohibited in this course. In particular, you are not allowed to use old exams. Using these materials will be considered a direct violation of academic policy and will be dealt with according to the GT Academic Honor Code. Furthermore, I do not allow copies of my exams out in the ether (so there should not be any out there for you to use anyway). Just as you are not to use the previous material you are not to share current material—including lecture material—with others either now or in the future. My policy on that is strict. If you violate the policy in any shape, form or fashion you will be dealt with according to the GT Academic Honor Code. I also have several... friends... from Texas who will help me personally deal with you.

Readings and Lectures

My research area is machine learning, and I'm deeply into the area. Given that and my enormous lung capacity, and my tendency to get distracted, it turns out that I can ramble on about the material for days on end, even with an editor to try to make me concise; however, that rather misses the point.

The online lectures are meant to summarize the readings and stress the important points. You are expected to critically read any assigned material. Your active participation in the material, the lectures, and various forums are crucial in making the course successful. This is less about my teaching than about your learning. My role is merely to assist you in the process of learning more about the area.

To help you to pace yourself, I have provided a nominal schedule (see the Syllabus link on the left) that tells you when we would be covering material if we were meeting twice a week during the term. I recommend you try to keep that pace.

Scoring and Grading

Your final grade is determined by how you do on three components: assignments, a midterm and a final exam.

- **Assignments.** There will be four scored assignments, one for the first section, two for the second, and one for the third. They will be about programming and analysis. Generally, they are designed to give you deeper insight into the material and to prepare you for the exams. The programming will be in service of allowing you to run and discuss experiments, do analysis, and so on. In fact, the programming is incidental, as you shall see.

When your assignments (projects and exams) are scored, you will receive feedback explaining your errors (and your successes!) in some level of detail. This feedback is for your benefit, both on this assignment and for future assignments. It is considered a part of your learning goals to internalize this feedback. This is one of many learning goals for this course, such as understanding how to analyze data or the differences between each algorithm or bias/variances in each of them.

If you are convinced that your score is in error in light of the feedback—as opposed to, say, that you scored lower than you expected—you may request a rescoring within a week of the score and feedback being returned to you. You will need to send a private Ed Discussions post to the head TA. You will need to provide sufficient explanation as to why you think the TA made a mistake (a rescoring request is only valid if it includes an explanation of where the TA made an actual error). Be concrete and specific. After the first assignment is returned, we will share painfully detailed directions on exactly how to do this sort of thing. We will not consider requests that do not follow those directions.

It is important to note that because we consider your ability to internalize feedback a learning goal, we also assess it. This ability is considered 10% of each assignment. We default to assigning you full credit. If you request a rescore and do not receive at least 5 points as a result of the request, you will lose those 10 points.

- **Midterm.** There will be a written, closed-book midterm roughly halfway through the term. The midterm will be administered via whatever our proctoring solution is this term.
- **Final Exam.** There will be a written, closed-book final exam at the end of the term. The final exam will also be administered via whatever our proctoring solution is this term.

Due Dates

All scored assignments are due by the time and date indicated. Here "time and date" means **Eastern Time**. If you are in another time zone, you should probably go to settings on Canvas and set your time zone appropriately. I will not accept late assignments or makeup exams. You will earn zero credit for any late assignment. The only exceptions will require: a **note** from an appropriate authority and **immediate notification** of the problem when it arises. Naturally, your excuse must be acceptable. If a meteor landed on your bed and destroyed your assignment, I need a signed note from the meteor.

Numbers

Component	
Assignments	50%

Midterm	25%
Final	25%

We intend to provide additional problem sets that you will be able to turn in. We will provide answers but will not score them or even look at them. So why would you ever bother? First, they will help prepare you for the exams, which will be scored. Second, in case you are right on the edge of two grades, I will remember that you did this or not.

Office Hours and Other Channels

I love my assignments. As you will discover they are wonderfully open-ended, much more so than many of you will be used to. It is therefore very important that in addition to watching the lectures and reading the, um, readings that you attend (or if you cannot or do not want to attend them as they happen that you later watch) office hours and regularly check Ed Discussions. You should consider your participation in both required.

Disclaimer

I reserve the right to modify any of these plans as need be during the course of the class; however, I won't do anything capriciously, anything I do change won't be too drastic, and you'll be informed as far in advance as possible.

Reading List

Required Text:

- **Tom Mitchell, Machine Learning. McGraw-Hill, 1997. (http://www.cs.cmu.edu/afs/cs.cmu.edu/user/mitchell/ftp/mlbook.html)**

Optional Text:

- **Larry Wasserman, All of Statistics. Springer, 2010. (http://www.stat.cmu.edu/~larry/all-of-statistics/)** (Read Part 1 for an intro to Probability Theory)
- Richard Sutton and Andrew Barto, Reinforcement Learning: An introduction. (for Reinforcement Learning) (**Nov 5, 2017 version (http://incompleteideas.net/book/bookdraft2017nov5.pdf)**)

A List:

- Linear Algebra
  - **Linear Algebra and Eigenproblems (https://github.com/pushkar/4641/raw/master/downloads/Eigenproblems.fm.pdf)**
- **ML is the ROX (https://classroom.udacity.com/courses/ud262/lessons/3625438937/concepts/last-viewed)**
  - Mitchell Ch 1

- [Decision Trees (https://classroom.udacity.com/courses/ud262/lessons/313488098/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/313488098/concepts/last-viewed)
  - Mitchell Ch 3
- [Regression and Classification (https://classroom.udacity.com/courses/ud262/lessons/312357973/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/312357973/concepts/last-viewed)
- [Neural Networks (https://classroom.udacity.com/courses/ud262/lessons/315142919/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/315142919/concepts/last-viewed)
  - Mitchell Ch 4
- [Instance-Based Learning (https://classroom.udacity.com/courses/ud262/lessons/666010252/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/666010252/concepts/last-viewed)
  - Mitchell Ch 8
- [Ensemble Learning (https://classroom.udacity.com/courses/ud262/lessons/367378584/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/367378584/concepts/last-viewed)
  - [Schapire's Introduction (https://github.com/pushkar/4641/raw/master/downloads/boosting.ps)](https://github.com/pushkar/4641/raw/master/downloads/boosting.ps)
  - [Jiri Matas and Jan Sochman's Slides (https://github.com/pushkar/4641/raw/master/downloads/adaboost_matas.pdf)](https://github.com/pushkar/4641/raw/master/downloads/adaboost_matas.pdf)
- [Kernel Methods and SVMs (https://classroom.udacity.com/courses/ud262/lessons/386608826/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/386608826/concepts/last-viewed)
  - [An introduction to SVMs for data mining (https://www.cc.gatech.edu/classes/AY2008/cs7641_spring/handouts/yor12-introsvm.pdf)](https://www.cc.gatech.edu/classes/AY2008/cs7641_spring/handouts/yor12-introsvm.pdf)
  - [Christopher Burges tutorial on SVMs for pattern recognition (http://research.microsoft.com/pubs/67119/svmtutorial.pdf)](http://research.microsoft.com/pubs/67119/svmtutorial.pdf)
  - [Scholkopf's NIPS tutorial slides on SVMs and kernel methods (https://github.com/pushkar/4641/raw/master/downloads/svm-scholkopf.ps)](https://github.com/pushkar/4641/raw/master/downloads/svm-scholkopf.ps)
- [Computational Learning Theory (https://classroom.udacity.com/courses/ud262/lessons/383498973/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/383498973/concepts/last-viewed)
  - Mitchell Ch 7
- [VC Dimensions (https://classroom.udacity.com/courses/ud262/lessons/417758568/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/417758568/concepts/last-viewed)
  - Mitchell Ch 7
- [Bayesian Learning (https://classroom.udacity.com/courses/ud262/lessons/454308909/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/454308909/concepts/last-viewed)
  - Mitchell Ch 6
- [Bayesian Inference (https://classroom.udacity.com/courses/ud262/lessons/478818537/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/478818537/concepts/last-viewed)
- [Randomized Optimization (https://classroom.udacity.com/courses/ud262/lessons/521298714/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/521298714/concepts/last-viewed)
  - Mitchell Ch 9

    - [No Free Lunch Theorem (https://ml-cs7641.s3.us-east-1.amazonaws.com/nfl-optimization-explanation.pdf)](https://ml-cs7641.s3.us-east-1.amazonaws.com/nfl-optimization-explanation.pdf)
- [Clustering (https://classroom.udacity.com/courses/ud262/lessons/644878538/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/644878538/concepts/last-viewed)
    - Mitchell Ch 6
    - [Intuitive Explanation of EM (http://www.cc.gatech.edu/~dellaert/em-paper.pdf)](http://www.cc.gatech.edu/~dellaert/em-paper.pdf)
    - [Statical View of EM (https://github.com/pushkar/4641/raw/master/downloads/em.pdf)](https://github.com/pushkar/4641/raw/master/downloads/em.pdf)
    - [Jon Kleinberg's Impossibility Theorem for Clustering (https://www.cs.cornell.edu/home/kleinber/nips15.pdf)](https://www.cs.cornell.edu/home/kleinber/nips15.pdf)
- [Feature Selection (https://classroom.udacity.com/courses/ud262/lessons/627968607/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/627968607/concepts/last-viewed)
    - [ICA: Algorithms and Applications (http://mlsp.cs.cmu.edu/courses/fall2012/lectures/ICA_Hyvarinen.pdf)](http://mlsp.cs.cmu.edu/courses/fall2012/lectures/ICA_Hyvarinen.pdf)
    - [Restructuring High Dimensional Data by Charles and Paul Viola (https://www.cc.gatech.edu/~isbell/papers/isbell-ica-nips-1999.pdf)](https://www.cc.gatech.edu/~isbell/papers/isbell-ica-nips-1999.pdf)
- [Feature Transformation (https://classroom.udacity.com/courses/ud262/lessons/649069103/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/649069103/concepts/last-viewed)
- [Information Theory (https://classroom.udacity.com/courses/ud262/lessons/672178843/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/672178843/concepts/last-viewed)
    - [Charles Isbell's Note on Information Theory (https://www.cc.gatech.edu/~isbell/tutorials/InfoTheory.fm.pdf)](https://www.cc.gatech.edu/~isbell/tutorials/InfoTheory.fm.pdf)
    - [An Introduction to Information Theory and Entropy (https://github.com/pushkar/4641/raw/master/downloads/gentle_intro_to_information_theory.pdf)](https://github.com/pushkar/4641/raw/master/downloads/gentle_intro_to_information_theory.pdf)
- [Markov Decision Processes (https://classroom.udacity.com/courses/ud262/lessons/684808907/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/684808907/concepts/last-viewed)
- [Reinforcement Learning (https://classroom.udacity.com/courses/ud262/lessons/643978935/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/643978935/concepts/last-viewed)
    - Mitchell Ch 13
    - [Richard Sutton and Andrew Barto, Reinforcement Learning: An introduction. MIT Press, 1998. (http://incompleteideas.net/book/bookdraft2017nov5.pdf)](http://incompleteideas.net/book/bookdraft2017nov5.pdf)
    - [Reinforcement Learning: A Survey (https://github.com/pushkar/4641/raw/master/downloads/kaelbling96reinforcement.pdf)](https://github.com/pushkar/4641/raw/master/downloads/kaelbling96reinforcement.pdf)
- [Game Theory (https://classroom.udacity.com/courses/ud262/lessons/668248596/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/668248596/concepts/last-viewed)
    - [Andrew Moore's slides (http://www.cs.cmu.edu/~awm/tutorials.html)](http://www.cs.cmu.edu/~awm/tutorials.html) [(http://www.cs.cmu.edu/~awm/tutorials.html)](http://www.cs.cmu.edu/~awm/tutorials.html)
- [Outro (https://classroom.udacity.com/courses/ud262/lessons/1571328670/concepts/last-viewed)](https://classroom.udacity.com/courses/ud262/lessons/1571328670/concepts/last-viewed)

Software

- **WEKA** **(http://www.cs.waikato.ac.nz/ml/weka/)** Machine learning software in JAVA that you can use for your projects
- **Data Mining with Weka** **(https://weka.waikato.ac.nz/)** A MOOC Course
- **ABAGAIL** **(https://github.com/pushkar/ABAGAIL)** Machine learning software in JAVA. This is hosted on my github, so you can contribute too
- **scikit-learn** **(http://scikit-learn.org/stable/)** A popular python library for supervised and unsupervised learning algorithms
- **pybrain** **(http://pybrain.org/)** A popular python library for artifical neural networks
- **Murphy's MDP Toolbox for Matlab** **(http://www.cs.ubc.ca/~murphyk/Software/MDP/mdp.html)**
- **MATLAB Clustering Package** **(http://www.cc.gatech.edu/~dellaert/FrankDellaert/Software.html)** By **Frank Dellaert** **(http://www.cc.gatech.edu/~dellaert/FrankDellaert/Frank_Dellaert/Frank_Dellaert.html)**

Datasets

- **UCI Machine Learning Repository** **(http://archive.ics.uci.edu/ml/)** An online repository of data sets that can be used for machine learning experiments.
- **Stanford Large Network Dataset** **(http://snap.stanford.edu/data/)** Dataset of large social and information networks.
- **Vision Benchmark Suite** **(http://www.cvlibs.net/datasets/kitti/index.php)** Autonomous car dataset