CSCI447/547
Machine Learning
MWF, 11:00AM-11:50AM, SS 362

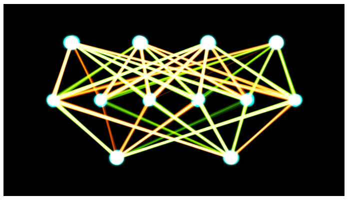

Instructor: Doug Brinkerhoff
E-mail: douglas1.brinkerhoff@umontana.edu
Office: ISB 406B
Office Hours: MWF, 1:00PM-2:00PM
(E-mail for an appointment, or my door is always open when I'm in)

Course Description: As a society we have reached a point where the amount of information available to us exceeds our capacity to analyze it without the assistance of the very computers that have made the collection of such vast sums of data possible. In this course we will explore the techniques required to turn these data into predictive models of varying complexity, from the modest linear regression to the vaunted deep neural network, with many other methods in between.

Course Objectives: At the completion of this course, students will be able to:

1. Understand what the phrase 'machine learning' means in the context of contemporary data analysis.
2. Understand fundamental principles such as Occam's Razor, inductive bias, and overfitting.
3. Select error metrics appropriate for the structure and subject of problems being considered.
4. Implement a bestiary of machine learning algorithms, both from scratch and with the assistance of high performance libraries like Google's TensorFlow.
5. Apply these algorithms to problems in their primary field of study.

Course Organization: Class time will be roughly proportioned between interactive lectures and collaborative work sessions. There will be 4 assignments on which you are encouraged to work with your classmates. You will have three weeks to work on each, but they will be extensive and likely difficult! Grad students will have the enviable experience of answering an extra advanced problem on each assignment.

Each student will be required to propose and execute a project. For undergraduates this may be either the implementation of a new machine learning algorithm, the application of machine learning to a non-trivial dataset, or an in depth report on a contemporary or classic paper on machine learning. This last option will not be available to grad students and the expectations will be somewhat higher. The latter several weeks will be dedicated to presenting the results of these projects. This project will take the place of a final exam.

Computers, Software, and Online Material: If possible, I would suggest bringing a laptop to class. Alternatively, the software necessary for the course will be available on the computers in SS362. A tentative list of the software that we'll be using is as follows:

1. Python (2.7 or 3, your choice)
2. Numpy/Scipy/Matplotlib: http://www.scipy.org/install.html
3. Jupyter: http://jupyter.org/install
4. scikit-learn: http://scikit-learn.org/stable/install.html

5. tensorflow: https://www.tensorflow.org/install/

All teaching material will be presented in Python, and I think that everyone will be happier if you do your assignments in Python as well. An important exception to this is the final project, where you are welcome to use whatever language/tools that you want (subject to approval of your proposal).

We will utilize many online resources throughout the course, from software libraries to external reading. These will generally be linked from the course's Moodle page. All internal course material (assignments, lecture notes and slides, Jupyter notebooks, etc.) will be available on Moodle.

Prerequisite(s): Officially, CSCI232: Data Structures and Algorithms. In reality, this course requires a commitment to making up any knowledge gaps that the student might have with respect to the course material. Because of the nature of the subject, ML borrows heavily from topics in calculus, statistics, discrete math, and programming. It is unlikely that anyone is going to be comfortable with the course material all the time. Don't get too bent out of shape about it.

Text(s):

1. *Pattern Recognition and Machine Learning*, Christopher M. Bishop, Springer, ISBN-13:978-0387310732
2. *Bayesian Reasoning and Machine Learning*, David Barber, http://web4.cs.ucl.ac.uk/staff/D.Barber/pmwiki/pmwiki.php?n=Brml.HomePage

Grade Distribution:

Course Participation and Attendance	20%
Assignments	40%
Project Proposal	10%
Project	20%
Project Presentation	10%

Letter Grade Distribution:

>= 93.00	A	73.00 - 76.99	C
90.00 - 92.99	A-	70.00 - 72.99	C-
87.00 - 89.99	B+	67.00 - 69.99	D+
83.00 - 86.99	B	63.00 - 66.99	D
80.00 - 82.99	B-	60.00 - 62.99	D-
77.00 - 79.99	C+	<= 59.99	F

(Note that I reserve the right to adjust these grades to your advantage to account ffor errors that I make in scaling assignment difficulty.)

Attendance Policy: Attendance is generally required, as participation is a significant portion of your grade. However, I also understand that there are lots of good reasons for not being able to make it to class; just talk to me about it and chances are we can work something out.

Late Assignments: I will not accept late assignments unless an extension was agreed upon well in advance of the due date.

Academic Integrity: All students must practice academic honesty. Academic misconduct is subject to an academic penalty by the course instructor and/or a disciplinary sanction by the University. All students need to be familiar with the Student Conduct Code. I will follow the guidelines given there. In cases of academic dishonesy, I will seek out the maximum allowable penalty. In general, there will not be many options for cheating in this course: Assignments will not come from the book, you are generally welcome to collaborate with your classmates, and any piracy of your final project will be painfully obvious for me and for everyone else. Look, this is a 400/500 level class, and if you're reading this you're probably looking to have a career in CS or a related field. When you're at a job interview, don't be sitting there regretting that you didn't learn anything in Machine Learning because you were cheating the whole time. Nobody wants that.

Disabilities: Students with disabilities may request reasonable modifications by contacting me. The University of Montana assures equal access to instruction through collaboration between students with disabilities, instructors, and Disability Services for Students. Reasonable means the University permits no fundamental alterations of academic standards or retroactive modifications.

Tentative Course Schedule: The following is subject to change according to the rate at which we proceed through the material, the moon and tides, and the results of my horoscope for the week.

Date	Topic	Assignments	Reading
Jan. 22 – Jan. 26	What does it mean to learn? • The vocabulary of optimization • Linear regression as a first example	HW1 Assigned	Bishop Ch. 1, Bishop Ch. 3
Jan. 29 – Feb. 02	Drawing a line in the sand. • Classification versus regression • Logistic regression		Bishop Ch. 4
Feb. 05 – Feb. 09	Known knowns, known unknowns, and unknown unknowns. • The meaning of probability • Maximum likelihood • Bayes' theorem and Naive Bayes	Dataset/paper idea due, Proposal assigned	Barber Ch. 1, Barber Ch. 10
Feb. 12 – Feb. 16	Bayesian belief networks • Graphs • Inference	HW 1 due, HW 2 assigned	Barber Ch. 2, Barber Ch. 3
Feb. 19 – Feb. 23	A mind like a silicone trap, Pt. 1 • The perceptron • Artificial Neural Networks		Bishop Ch.5.1–5.2

Feb. 26 – Mar. 02	A mind like a silicone trap, Pt. 2 • The backpropagation algorithm • Regularization		Bishop Ch.5.3–5.7
Mar. 05 – Mar. 09	Dimensionality Reduction • Feature selection • Principal component analysis	HW 2 due, HW 3 assigned	Barber Ch. 15
Mar. 12 – Mar. 16	It's a beautiful day in the k-nearest neighborhood. • K-nearest neighbor classification		Barber Ch. 14
Mar. 19 – Mar. 23	Tending the orchard. • Decision trees • Random forests	Proposal due	Bishop Ch. 14.4, Wikipedia
Mar. 26 – Mar. 30	Spring Break '018, No parents!		
Apr. 02 – Apr. 06	Support Vector Machines/Kernel Methods. • Gaussian process regression • Kernel development	HW 3 due, HW 4 assigned	Bishop Ch. 6
Apr. 09 – Apr. 13	Computers gone wild. • Unsupervised Algorithms • Clustering • Mixture of Gaussians/Expectation Maximization		Bishop Ch. 9
Apr. 16 – Apr. 20	Presentations 1		
Apr. 23 – Apr. 27	Presentations 2	HW 4 due	
Apr. 30 – May. 04	Presentations 3	Project due	