Probabilistic Models of Human and Machine Intelligence
CSCI 5822
Spring 2021

Instructor:
Rebecca Morrison
ECOT 820
rebeccam@colorado.edu

Class details:
Class meetings: MWF 1:50 – 2:40, ECCR 1B40/Zoom
TA: Tzu-Chi Yen, tzuchi.yen@colorado.edu
IS: Chou Yi, yi.chou@colorado.edu
Tech Co-Pilot: Michela Puni Nimako, michela.puninimako@colorado.edu
Office hours:
Mon 3:00 – 5:00p (CY)
Tu 2:00 – 4:00p (TY)
Wed 9:15 – 10:45a (RM)
and by appointment

Materials:
Bayesian Reasoning and Machine Learning by David Barber
Available online: http://web4.cs.ucl.ac.uk/staff/D.Barber/textbook/200620.pdf
Lecture notes (R. Morrison)
Canvas for announcements, video recordings of the class, and to post and collect homework
Zulip for questions about the course and student-led discussions
Various papers and book chapters (will be available on Canvas)

*******Email to instructor should be used as a last resort.*******

Course description

This course will introduce the basic concepts of probabilistic models: directed, undirected, and factor graphs, and the computations we can do with them. Probabilistic modeling—at the intersection of graph theory and probability—is an incredibly powerful area of mathematics and computer science with countless applications. In this class, we will learn how to leverage known probabilistic structures to make sense of large data, and how to learn unknown structure. Topics include the fundamentals of Bayesian inference; expressiveness of graphical models; conditioning, marginalization, and message passing; junction trees, triangulation, and moralization; the exponential family; maximum likelihood estimation; learning with hidden variables; fundamentals of machine learning; Gaussian processes; discrete and continuous Markov models; sampling methods; and variational inference.

Course outline, by week*

Week 0 (1/15)	Introductions
Week 1 (1/20, 1/22)	Ch 1: Probabilistic Reasoning
Week 2 (1/25, 1/27, 1/29)	Ch 2: Basic Graph Concepts, Ch 3: Belief Networks
Week 3 (2/1, 2/3, 2/5)	Ch 4: Graphical Models
Week 4 (2/8, 2/10, 2/12)	Ch 5: Efficient Inference in Trees
Week 5 (2/15, 2/19)	Ch 6: The Junction Tree Algorithm
Week 6 (2/22, 2/24, 2/26)	Ch 8: Statistics for Machine Learning
Week 7 (3/1, 3/3, 3/5)	Ch 9: Learning as Inference
Week 8 (3/8, 3/10, 3/12)	Ch 10: Naive Bayes, Ch 11: Learning with Hidden Variables
Week 9 (3/15, 3/17, 3/19)	Ch 12: Bayesian Model Selection, Ch 13: Machine Learning Concepts
Week 10 (3/22, 3/24, 3/26)	"Break Week:" No new book material (Paper discussion)
Week 11 (3/29, 3/31, 4/2)	Ch 14: Nearest Neighbor Classification, Ch 18: Bayesian Linear Models
Week 12 (4/5, 4/7, 4/9)	Ch 19: Gaussian Processes, Ch 20: Mixture Models
Week 13 (4/12, 4/14, 4/16)	Ch 23: Discrete-State Markov Models, Ch 24: Continuous-State Markov Models
Week 14 (4/19, 4/21, 4/23)	Ch 27: Sampling
Week 15 (4/26, 4/28)	Ch 28: Deterministic Approximate Inference
Finals slot (TBA)	Reserve for Project Presentations

*Note that this schedule is approximate and subject to modifications.

Course work and grading

Grades will be determined based on homeworks (70%) and a final project (30%).

- **Homework:** Homeworks will be assigned approximately every two weeks on Canvas. You will be asked to either upload a pdf and/or a Jupyter notebook. You may complete homework assignments alone or in groups of two. The homework grade with the lowest score will be dropped. If you are very sure that the homework has been misgraded, you may contact Tzu-Chi directly. Otherwise, please accept the grade and try to understand what you could do better or make clearer next time.
  - **Late policy:** Because of the large class size, late homework will not be accepted.
- **Final project:** For the final project, please work in groups of size at least two, and not bigger than five. You will have the option to either work on something related to your own research, or complete an assigned project. More information will be given during break week, but these projects will not be much longer than the assignments (3–5 pages). We will use the finals slot for very quick final presentations.

Note: You may choose whether or not to attend class, and I aim to make class worth your time. If you do attend, I expect you to arrive on time and stay until the end.