CSC 535: Probabilistic Graphical Models

Tue / Thurs 5:00pm – 6:15pm : Gould-Simpson 701

Course Description

Probabilistic graphical modeling and inference is a powerful modern approach to representing the combined statistics of data and models, reasoning about the world in the face of uncertainty, and learning about it from data. It cleanly separates the notions of representation, reasoning, and learning. It provides a principled framework for combining multiple sources of information, such as prior knowledge about the world, with evidence about a particular case in observed data. This course will provide a solid introduction to the methodology and associated techniques, and show how they are applied in diverse domains ranging from computer vision to computational biology to computational neuroscience.

Instructor and Contact Information

Jason Pacheco, GS 724, pachecoj@cs.arizona.edu

Office Hours: Friday 3-5pm

Course Homepage: http://www.pachecoj.com/courses/csc535_fall24

D2L: https://d2l.arizona.edu/d2l/home/1493375

Piazza: https://piazza.com/arizona/fall2024/csc535

Instructor Homepage: http://www.pachecoj.com

Course Format and Teaching Methods

The course will consist of regular in-person lectures. In-class discussion as well as Q&A is encouraged.

Course Objectives

The broad objectives of this course are to develop a solid fundamental understanding of probabilistic graphical models, learn how to apply them to diverse problems, and build a toolkit of useful statistical models and related algorithms. Assignments and exams will develop and evaluate both conceptual understanding and applying the methodology to practical problems.

Concepts that students are expected to learn include: Bayesian methodology, conditional independence, modeling and inference as distinct activities, model selection, Bayesian decision making, directed graphical models (Bayes nets), sampling probability distributions from Bayes nets (ancestral sampling), undirected graphical models (Markov random fields, factor graphs), relationships between model types and the space of probability distributions, causality, statistical clustering, statistical inference, exact inference on graphs using message passing, expressing model learning as inference, approximate inference for missing value problems using expectation maximization (EM), variational inference, sampling probability distributions using Markov chain Monte Carlo (MCMC), and how MCMC can used for inference.

Commonly used models that students will learn about include Naïve Bayes, Gaussian mixture models (GMM), hidden Markov models (HMM), and linear dynamic systems (LDS). Generally applicable algorithms that students will learn about include sum-product (includes forward/backward for HMM as a special case), max-sum (includes Viterbi as a special case), K-means clustering, expectation maximization (EM), variational inference, Kalman filter, Metropolis Hastings, Gibbs sampling, and particle filter.

Topics

Introductory foundations
- Probabilistic foundations
- Introduction to the Bayesian methodology and introductory examples
- Actions and decisions
- Model selection

Graphical representation of probabilistic models
- Representing models using directed graphs (Bayes nets)
- Representing models using undirected graphs (Markov Random fields)
- Factor graphs

Examples of graphical models
- Naïve Bayes
- Gaussian Mixture Models (GMM)
- Hidden Markov Models (HMM)
- Linear Dynamical Systems (LDS)

Inference for graphical models
- Sum product algorithm
- Max sum algorithm
- Expectation maximization (EM)
- Markov chain Monte Carlo (MCMC) methods

Expected Learning Outcomes

Specific skills that students will develop as a result of this course include the following:

- Basic knowledge of applied statistics and probability including, both, frequentist (classical) as well as Bayesian methods
- Creating both directed and undirected graphical models for data
- Identifying conditional independencies in graphical models
- Specifying distributions for parameters of model components that link the model to data
- Applying exact inference methods to compute marginal probabilities and maximally probable configurations given a model (sum-product and max-sum algorithms, respectively)
- Applying approximate inference to learn model parameters using expectation maximization (EM algorithm), variational inference, and various Markov chain Monte Carlo methods including Metropolis Hastings sampling, Gibbs sampling, and Hamiltonian Monte Carlo.

Makeup Policy for Students Who Register Late

Students must complete assignments on time. No makeup opportunity will be offered for students. Late registrants to the course will receive a zero grade for any missed assignments.

Course Communications

All course communications will be conducted via Piazza. Assignments will be submitted via D2L and feedback will be provided on that platform.

Required Texts or Readings

This class will select reading material from the following textbook:

Murphy, Kevin, "Machine Learning: A Probabilistic Perspective." 2012. (electronic copy available through the UA library webpage)

Assignments and Examinations: Schedule/Due Dates

- 08/29 : HW1 Assigned : Probability Primer
- 09/10 : HW1 Due : HW2 Assigned : Bayesian Probability + Independence
- 09/19 : HW2 Due : HW3 Assigned : Sum-Product Belief Propagation
- 10/08 : HW3 Due
- 10/10 : Midterm
- 10/22 : HW4 Assigned : Kalman Filter + Expectation Maximization
- 11/07 : HW4 Due : HW5 Assigned : Importance Sampling + Particle Filter
- 11/26 : HW5 Due : HW6 Assigned : MCMC / Gibbs Sampling
- 12/10 : HW6 Due
- 12/18 : Final Exam Due

Final Examination

The final exam will be take-home. It will be due 12/18, the day of the scheduled examination period.

Grading Scale and Policies

Assignment grading. Assignment deliverables will generally consist of two parts: 1) all code developed in response to the assignments; and 2) a report, in PDF format explaining what has been done, what the results were, commenting on the results, and answering any questions posed in the assignment. The instructor will provide a document that details the expectations of the report. Assignments will be graded with respect to four criteria: 1) reproducibility (the ease by which the grader can run the code to get the reported results); 2) completeness (the extent that the work done and sufficient effort was applied); 3) correctness; and 4) the exposition (clarity, insight, and conformance to the guidelines provided). The weight of these four criteria will vary among the assignments, but students are advised that the fourth criterion will generally have substantive weight.

Grading breakdown

Assignments: 65%
Midterm: 15%
Final Exam: 20%

90% guarantees an A, 80% guarantees a B, 70% a C, and 60% a D.

Incomplete (I) or Withdrawal (W):

Requests for incomplete (I) or withdrawal (W) must be made in accordance with University policies, which are available at https://catalog.arizona.edu/policy/courses-credit/grading/grading-system.

Dispute of Grade Policy:
Any grading disputes should be communicated to the professor within one week of having received the grade. The professor will announce in class and on Piazza when grading is complete for each assignment. If a student has not received a grade on a submitted assignment it must be communicated to the adviser within one week of this announcement. If no assignment was submitted it will receive a score of zero.

Tentative Scheduled Topic and Activities

All assignments, due dates, and readings are available on the course webpage. The following are scheduled topics by week:

Week 1: Introduction + Probability Primer
Week 2: Bayesian Probability
Week 3: Probabilistic Graphical Models (PGMs)
Week 4: PGMs Continued
Week 5: Message Passing Inference (Belief Propagation)
Week 6: Message Passing Inference (Variable Elimination, Junction Tree)
Week 7: Parameter Learning + Expectation Maximization
Week 8: Dynamical Systems
Week 9: Dynamical Systems Continued
Week 10: Monte Carlo Inference (Monte Carlo Methods)
Week 11: Monte Carlo Inference (Markov Chain Monte Carlo)
Week 12: Variational Inference
Week 13: Variational Inference Continued
Week 14: Variational Autoencoder
Week 15: Wrapup

Classroom Behavior Policy

To foster a positive learning environment, students and instructors have a shared responsibility. We want a safe, welcoming, and inclusive environment where all of us feel comfortable with each other and where we can challenge ourselves to succeed. To that end, our focus is on the tasks at hand and not on extraneous activities (e.g., texting, chatting, reading a newspaper, making phone calls, web surfing, etc.).

Students are asked to refrain from disruptive conversations with people sitting around them during lecture. Students observed engaging in disruptive activity will be asked to cease this behavior. Those who continue to disrupt the class will be asked to leave lecture or discussion and may be reported to the Dean of Students.

Safety on Campus and in the Classroom

For a list of emergency procedures for all types of incidents, please visit the website of the Critical Incident Response Team (CIRT): https://cirt.arizona.edu/case-emergency/overview

Also watch the video available at

https://arizona.sabacloud.com/Saba/Web_spf/NA7P1PRD161/app/me/ledetail;spfurl=common%2Flearningeventdetail%2Fcrtfy000000000003841

University-wide Policies link

Links to the following UA policies are provided here: https://catalog.arizona.edu/syllabus-policies

- Absence and Class Participation Policies
- Threatening Behavior Policy
- Accessibility and Accommodations Policy
- Code of Academic Integrity
- Nondiscrimination and Anti-Harassment Policy

Department-wide Syllabus Policies and Resources link

Links to the following departmental syllabus policies and resources are provided here, https://www.cs.arizona.edu/cs-course-syllabus-policies :

- Department Code of Conduct
- Class Recordings
- Illnesses and Emergencies
- Obtaining Help
- Preferred Names and Pronouns
- Confidentiality of Student Records
- Additional Resources
- Land Acknowledgement Statement

Subject to Change Statement

Information contained in the course syllabus, other than the grade and absence policy, may be subject to change with advance notice, as deemed appropriate by the instructor.