CSC 535

Probabilistic graphical models

MW 12:30-1:45, Gould Simpson, 701

Description of the Course

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 source 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 molecular biology to astronomy.

Course Prerequisites or Co-requisites

MATH 223 and MATH 313 or equivalent math background. MATH 464 or alternative course that covers basic discrete and continuous probability. CSC 345 or equivalent preparation in algorithms, data structures, and programming.

Instructor and Contact Information

Instructor:

Kobus Barnard, GS 708, kobus@cs.arizona.edu

Office Hours: TBA and by appointment

Web information:

Course home page: http://vision.cs.arizona.edu/teaching/cs535/spring18

Instructor home page: http://kobus.ca

We will use D2L and Piazza for this course

Course Format and Teaching Methods

Lecture only.

Course Objectives and Expected Learning Outcomes

The broad objectives are 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), 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), 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) (includes Kalman Filtering as a special case). Generally applicable algorithms that students will learn about include sum-product (includes alpha-beta for HMM as a special case), max-sum (includes Viterbi as a special case), K-means clustering, expectation-maximization (EM), Metropolis Hastings, Gibbs sampling, and Hamiltonian Monte Carlo.

Specific skills that students will develop through homework assignments include: **1)** creating both directed and undirected graphical models for data; **2)** identifying conditional independencies in graphical models; **3)** specifying distributions for parameters of model components that link the model to data; **4)** applying exact inference methods to compute marginal probabilities and maximally probable configurations given a model (sum-product and max-sum algorithms, respectively); and **5)** applying approximate inference to learn model parameters using expectation maximization (EM algorithm), and various Markov chain Monte Carlo methods including Metropolis Hastings sampling, Gibbs sampling, and Hamiltonian Monte Carlo.

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)
  - Causality
  - Factor graphs

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

- Inference for graphical models (interspersed with examples of graphical models)
  - Sum product algorithm
  - Max sum algorithm
  - Expectation maximization (EM)
  - Markov chain Monte Carlo (MCMC) methods including Metropolis Hastings, Gibbs sampling, and Hamiltonian Monte Carlo.

Absence and Class Participation Policy

The UA's policy concerning Class Attendance, Participation, and Administrative Drops is available at http://catalog.arizona.edu/policy/class-attendance-participation-and-administrative-drop

The UA policy regarding absences for any sincerely held religious belief, observance or practice will be accommodated where reasonable: http://policy.arizona.edu/human-resources/religious-accommodation-policy.

Absences preapproved by the UA Dean of Students (or dean's designee) will be honored. See https://deanofstudents.arizona.edu/absences

Good attendance is expected. Poor attendance in the first few weeks may lead to the student being dropped from the class. Exams must be attended at their appointed time unless you have permission in advance to do otherwise. If you are not able to make an exam time due to extenuating circumstances, the instructor must be contacted in advance to verify that alternative arrangements are justified.

Makeup Policy for Students Who Register Late

In consultation with the instructor, students who register late can makeup for assignments that are past their due date with optional parts of subsequent assignments or alternative assignments. However, such students are still responsible for the intellectual content of the past due assignments which can be relevant for subsequent assignments as well as exams.

Course Communications

Online communication will be conducted using D2L, Piazza, and, where applicable, official UA e-mail addresses.

Required Texts or Readings

This class will select material from the following two text books:

1) Probabilistic Graphical Models: Principles and Techniques, Daphne Koller and Nir Friedman

2) Pattern Recognition and Machine Learning, Chris Bishop (we will make significant use of chapter 8 which is available on-line from the author).

Neither book is required as all course material will be available within the lecture notes, assignments, and supplementary readings.

The IVILAB has a copy of both of these text books that are available for short term loan, and the UA library has a copies.

Required or Special Materials

Matlab: Most students will use Matlab for most assignments. While there will be some flexibility in the choice of programming languages, unless there is a reason to do otherwise (please consult with the instructor), students are advised to use Matlab. Matlab is available on a number of student accessible computers across campus including the CS machine general purpose instructional computer "lectura". Students wishing to use Matlab on a personal computer can download and install it through the U. Arizona web pages (http://softwarelicense.arizona.edu/mathworks-matlab).

Required Extracurricular Activities

None.

Assignments and Examinations: Schedule/Due Dates

There will be 10 assignments, one midterm, and a final, as detailed in the table below. For maximum flexibility, I will post assignments as soon as we have covered material that enables students to start on them. Due dates are nominally midnight, with grace until 8am the following morning.

	Description	Weight	Due	Graded
HW1	Introduction	4	Jan 23	Jan 28
HW2	Bayesian probability	7	Feb 02	Feb 08
HW3	Linear models, AIC, and BIC.	7	Feb 13	Feb 18
HW4	Bayesian networks (directed PGMs)	7	Feb 23	Mar 01
HW5	Markov Random Fields (undirected PGMs)	7	Mar 13	Mar 18
Midterm	All material up to and including week 7 and HW4	10	Mar 14	Mar 18
HW6	Sum-product max-sum algorithm (Part I)	5	Mar 21	Mar 25

	Withdraw deadline		Mar 27	
HW7	Sum-product max-sum algorithm (Part II)	7	Mar 31	April 05
HW8	Hidden Markov Models, Linear dynamical systems	7	April 11	April 15
HW9	Gaussian Mixture Models, Expectation Maximization	7	April 21	April 26
HW10	Markov chain Monte Carlo sampling	7	May 02	May 04
Final	All course material	25	May 09	May 11

Final Examination or Project

The final exam will occur on May 9 in Gould-Simpson 701, from 10:30am to 12:30pm.

For the U. Arizona Final Exam Regulations see https://www.registrar.arizona.edu/courses/final-examination-regulations-and-information.

For the U. Arizona and Final Exam Schedule, http://www.registrar.arizona.edu/schedules/finals.htm

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 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 (see above table for more detail).

Assignments: 65%

Midterm: 10%

Final Exam: 25%

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

Requests for incomplete (I) or withdrawal (W) must be made in accordance with University policies, which are available at http://catalog.arizona.edu/policy/grades-and-grading-system#incomplete and http://catalog.arizona.edu/policy/grades-and-grading-system#Withdrawal, respectively.

Dispute of Grade Policy. Students wishing to dispute a grade on an assignment or exam should contact the instructor within two weeks of the date that the assignment or exam was returned to the students.

Scheduled Topics

Week 01: Course mechanics, introduction to models connecting theory to data
Week 02: Technical introduction
Week 03: Probability review
Week 04: Sampling probability distributions, Bayesian update, conjugate priors
Week 05: Predictive distribution, model selection, cross validation, classification, decision making
Week 06: Directed graphical models
Week 07: Causality, undirected graphical models
Week 08: Exact inference for graphical models, factor graphs, message passing
Week 09: Sum-product algorithm and max-sum algorithms
Week 10: Clustering, Expectation Maximization for Gaussian Mixture Model
Week 11: Expectation Maximization more formally, Hidden Markov Model (HMM)
Week 12: Inference for HMMs, HMM example, Linear Dynamical Systems (LDS)
Week 13: Sampling based inference, rejection sampling, importance sampling, Markov chain Monte Carlo (MCMC), Metropolis Hastings method
Week 14: Matrix algebra interpretation of MCMC, Gibbs sampling, Hamiltonian MCMC sampling
Week 15: Advanced topics

Bibliography

All material for assignments and exams will be in the class notes that will be posted on D2L. However, for additional study by interested students, the following will be referred to. The three text books are available in the library or short term loan from the instructor. The papers are all available on-line.

[1] C. Andrieu, N. d. Freitas, A. Doucet, and M. I. Jordan., "An introduction to MCMC for machine learning.," Machine Learning, vol. 50, pp. 5-43, 2003.

[2] C. M. Bishop, Pattern recognition and machine learning: Springer, 2006.

[3] S. Borman, "The expectation maximization algorithm-a short tutorial," in Submitted for publication, 2004, pp. 1-9.

[4] B. Efron, "Why isn't everyone a Bayesian?," The American Statistician, vol. 40, pp. 1-5, 1986.

[5] R. E. Kass and A. E. Raftery, "Bayes factors," Journal of The American Statistical Association, vol. 90, pp. 377-395, 1995.

[6] D. Koller and N. Friedman, Probabilistic Graphical Models: The MIT Press, 2009.

[7] K. Murphy, Machine Learning: A Probabilistic Perspective: The MIT Press, 2012.

[8] R. M. Neal, "Probabilistic inference using Markov Chain Monte Carlo methods," 1993.

[9] R. M. Neal, "MCMC using Hamiltonian dynamics," in Handbook of Markov Chain Monte Carlo, MCMC using Hamiltonian dynamics, 2011, pp. 113-162.

[10] L. R. Rabiner, "A tutorial on Hidden Markov Models and selected applications in speech recognition," Proceedings of the IEEE, vol. 77, pp. 257-286, 1989.

Department of Computer Science Code of Conduct

The Department of Computer Science is committed to providing and maintaining a supportive educational environment for all. We strive to be welcoming and inclusive, respect privacy and confidentiality, behave respectfully and courteously, and practice intellectual honesty. Disruptive behaviors (such as physical or emotional harassment, dismissive attitudes, and abuse of department resources) will not be tolerated. The complete Code of Conduct is available on our department web site. We expect that you will adhere to this code, as well as the UA Student Code of Conduct, while you are a member of this class.

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.

Some learning styles are best served by using personal electronics, such as laptops and iPads. These devices can be distracting to other learners. Therefore, students who prefer to use electronic devices for note-taking during lecture should sit towards the back of the class, or an area of the classroom agreed upon between the instructor and students.

Threatening Behavior Policy

The UA Threatening Behavior by Students Policy prohibits threats of physical harm to any member of the University community, including to oneself. See http://policy.arizona.edu/education-and-student-affairs/threatening-behavior-students.

Accessibility and Accommodations

At the University of Arizona we strive to make learning experiences as accessible as possible. If you anticipate or experience physical or academic barriers based on disability or pregnancy, you are welcome to let me know so that we can discuss options. You are also encouraged to contact Disability Resources (520-621-3268) to explore reasonable accommodation.

If our class meets at a campus location: Please be aware that the accessible table and chairs in this room should remain available for students who find that standard classroom seating is not usable.

Code of Academic Integrity

Students are encouraged to share intellectual views and discuss freely the principles and applications of course materials. However, graded work/exercises must be the product of independent effort unless otherwise instructed. Students are expected to adhere to the UA Code of Academic Integrity as described in the UA General Catalog. See http://deanofstudents.arizona.edu/academic-integrity/students/academic-integrity.

The University Libraries have some excellent tips for avoiding plagiarism, available at http://www.library.arizona.edu/help/tutorials/plagiarism/index.html.

Sharing solution keys with others (e.g., students who might take the class in a future term, or who are taking the class in a future term) is considered by the instructor to be a serious violation of academic integrity.

Selling class notes and/or other course materials to other students or to a third party for resale is not permitted without the instructor's express written consent. Violations to this and other course rules are subject to the Code of Academic Integrity and may result in course sanctions. Additionally, students who use D2L or UA e-mail to sell or buy these copyrighted materials are subject to Code of Conduct Violations for misuse of student e-mail addresses. This conduct may also constitute copyright infringement.

UA Nondiscrimination and Anti-harassment Policy

The University is committed to creating and maintaining an environment free of discrimination; see http://policy.arizona.edu/human-resources/nondiscrimination-and-anti-harassment-policy

Our classroom is a place where everyone is encouraged to express well-formed opinions and their reasons for those opinions. We also want to create a tolerant and open environment where such opinions can be expressed without resorting to bullying or discrimination of others.

Additional Resources for Students

UA Academic policies and procedures are available at http://catalog.arizona.edu/policies

Student Assistance and Advocacy information is available at http://deanofstudents.arizona.edu/student-assistance/students/student-assistance

Confidentiality of Student Records

Please refer to:

http://www.registrar.arizona.edu/personal-information/family-educational-rights-and-privacy-act-1974-ferpa?topic=ferpa

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.