CSC 535: Probabilistic Graphical Models

MW 12:30pm – 1:45pm : Gould-Simpson 701

Description of 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 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.

Course Prerequisites

MATH 223 and MATH 313or 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:
Jason Pacheco, GS 724, Email: pachecoj@cs.arizona.edu
Office Hours: Tuesday, 10:00-11:30am

Web Information:
Course Homepage: http://www.pachecoj.com/courses/csc535_spring22
D2L: https://d2l.arizona.edu/d2l/home/1135240
Piazza: https://piazza.com/arizona/spring2022/csc535
Instructor Homepage: http://www.pachecoj.com

Obtaining Help

- **Advising:** If you have questions about your academic progress this semester, or your chosen degree program, consider contacting your graduate program coordinator and faculty advisor. Your program coordinator, faculty advisor, and the Graduate Center can guide you toward university resources to help you succeed.**Computer Science students** are encouraged to email gradadvising@cs.arizona.edu for advising related questions.
- **Life challenges:** If you are experiencing unexpected barriers to your success in your courses, please note the Dean of Students Office is a central support resource for all students and may be helpful. The Dean of Students Office can be reached at 520-621-2057 or DOS-deanofstudents@email.arizona.edu.
- **Physical and mental-health challenges**: If you are facing physical or mental health challenges this semester, please note that Campus Health provides quality medical and mental health care. For medical appointments, call (520-621-9202. For After Hours care, call (520) 570-7898. For the Counseling & Psych Services (CAPS) 24/7 hotline, call (520) 621-3334.

Class Recordings

Lectures will be recorded via Zoom and available through D2L. For lecture recordings, which are used at the discretion of the instructor, students must access content in D2L only. Students may not modify content or re-use content for any purpose other than personal educational reasons. All recordings are subject to government and university regulations. Therefore, students accessing unauthorized recordings or using them in a manner inconsistent with UArizona values and educational policies (Code of Academic Integrity and the Student Code of Conduct) are also subject to civil action.

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)
- 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
- Sum product algorithm
- Max sum algorithm
- *Expectation maximization (EM)*
- *Markov chain Monte Carlo (MCMC) methods including Metropolis Hastings and Gibbs sampling*

Expected Learning Outcomes

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

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)
5) 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.

Absence and Class Participation Policy

Participating in the course and attending lectures and other course events are vital to the learning process. As such, attendance is required at all lectures and discussion section meetings. To request a disability-related accommodation to this attendance policy, please contact the Disability Resource Center at (520) 621-3268 or drc-info@email.arizona.edu. If you are experiencing unexpected barriers to your success in your courses, the Dean of Students Office is a central support resource for all students and may be helpful. The Dean of Students Office is located in the Robert L. Nugent Building, room 100, or call 520-621-7057.

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

Illnesses and Emergencies

- If you feel sick, or may have been in contact with someone who is infectious, stay home. Except for seeking medical care, avoid contact with others and do not travel.
- Notify your instructor(s) if you will be missing up to one week of course meetings and/or assignment deadlines.
- If you must miss the equivalent of more than one week of class and have an emergency, the Dean of Students is the proper office to contact (DOS-deanofstudents@email.arizona.edu). The Dean of Students considers the following as qualified emergencies: the birth of a child, mental health hospitalization, domestic violence matter, house fire, hospitalization for physical health (concussion/emergency surgery/coma/COVID-19 complications/ICU), death of immediate family, Title IX matters, etc.
- Please understand that there is no guarantee of an extension when you are absent from class and/or miss a deadline.

Statement on compliance with COVID-19 mitigation guidelines: As we enter the Spring semester, your and my health and safety remain the university's highest priority. To protect the health of everyone in this class, students are required to follow the university guidelines on COVID-19 mitigation. Please visit www.covid19.arizona.edu.

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

Means by which online communication will be conducted (e.g., official UA e-mail address, D2L)

Required Texts and Materials

This class will select reading material from the following two textbooks:

1) Bishop, Chris, "Pattern Recognition and Machine Learning" (https://www.microsoft.com/en-us/research/uploads/prod/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf)
2) Murphy, Kevin, "Machine Learning: A Probabilistic Perspective." 2012. (electronic copy available through the UA library webpage)

Scheduled Topics/Activities

Week 1: Introduction
Week 2: Probability Primer
Week 3: Bayesian Probability and Statistics
Week 4: Probabilistic Graphical Models
Week 5: Message Passing Inference (Variable Elimination, Sum-Product Belief Propagation)
Week 6: Message Passing Inference (Max-Product Belief Propagation, Junction Tree)
Week 7: Parameter Learning (Expectation Maximization)
Week 8: Midterm
Week 9: Dynamical Systems
Week 10: Monte Carlo Methods (Importance Sampling, Sequential Monte Carlo)
Week 11: Monte Carlo Methods (Markov chain Monte Carlo)
Week 12: Exponential Families and Variational Inference

Week 13: Variational Inference (Mean Field, Stochastic Variational)
Week 14: Topic Models
Week 15: Bayesian Nonparametrics
Week 16: Course Wrapup

Final Examination or Project

This class will have a take-home final exam that will be due **no later than** Wednesday, May 11. Final Exam Regulations and Final Exam Schedule: https://registrar.arizona.edu/faculty-staff-resources/room-class-scheduling/schedule-classes/final-exams

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.

Department of Computer Science Grading Policy:

1. Instructors will explicitly promise when every assignment and exam will be graded and returned to students. These promised dates will appear in the syllabus, associated with the corresponding due dates and exam dates.
2. Graded homework will be returned before the next homework is due.
3. Exams will be returned "promptly", as defined by the instructor (and as promised in the syllabus).
4. Grading delays beyond promised return-by dates will be announced as soon as possible with an explanation for the delay.

Incomplete (I) or Withdrawal (W):

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

Any grading disputes should be communicated with the professor **within one week** of receiving 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.

Bibliography

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

Wainwright, Martin J., and Michael I. Jordan. "Graphical models, exponential families, and variational inference." *Foundations and Trends® in Machine Learning* 1.1–2 (2008): 1-305.

Andrieu, Christophe, N. d. Freitas, A. Doucet, M. I. Jordan. "An introduction to MCMC for machine learning." *Machine learning* 50.1-2 (2003): 5-43.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. "Latent dirichlet allocation." *Journal of machine Learning research* 3.Jan (2003): 993-1022.

Murphy, Kevin P. *Machine learning: a probabilistic perspective* . MIT press, 2012.
Neal, Radford M. "MCMC using Hamiltonian dynamics." *Handbook of markov chain monte carlo* 2.11 (2011): 2.

Hoffman, Matthew D., and Andrew Gelman. "The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research* 15.1 (2014): 1593-1623.

Snoek, Jasper, Hugo Larochelle, and Ryan P. Adams. "Practical bayesian optimization of machine learning algorithms." *Advances in neural information processing systems* . 2012.

Shahriari, Bobak, et al. "Taking the human out of the loop: A review of Bayesian optimization." *Proceedings of the IEEE* 104.1 (2015): 148-175.

Kingma, Diederik P., and Max Welling. "Auto-encoding variational bayes." *arXiv preprint arXiv:1312.6114* (2013).

Rezende, Danilo Jimenez, Shakir Mohamed, and Daan Wierstra. "Stochastic backpropagation and approximate inference in deep generative models." *arXiv preprint arXiv:1401.4082* (2014).

Rasmussen, Carl Edward. "Gaussian processes in machine learning." *Summer School on Machine Learning* . Springer, Berlin, Heidelberg, 2003.

Teh, Y., M. Jordan, M. Beal, et al. Hierarchical Dirichlet processes.Journal of the American Statistical As-sociation, 101(476):1566–1581, 2007.

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.

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 barriers based on disability or pregnancy, please contact the Disability Resource Center (520-621-3268, https://drc.arizona.edu/) to establish reasonable accommodations.

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.

Uploading material from this course to a website other than D2L (or the class piazza) is strictly prohibited and will be considered a violation of the course policy and a violation of the code of academic integrity. Obtaining material associated with this course (or previous offerings of this course) on a site other than D2L (or the class piazza), such as Chegg, Course Hero, etc. or accessing these sites during a quiz or exam is a violation of the code of academic integrity. Any student determined to have uploaded or accessed material in an unauthorized manner will be reported to the Dean of Students for a Code of Academic Integrity violation, with a recommended sanction of a failing grade in the course.

The University Libraries have some excellent tips for avoiding plagiarism, available at https://new.library.arizona.edu/research/citing/plagiarism.

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.

Nondiscrimination and Anti-harassment Policy

The University of Arizona is committed to creating and maintaining an environment free of discrimination. In support of this commitment, the University prohibits discrimination, including harassment and retaliation, based on a protected classification, including race, color, religion, sex, national origin, age, disability, veteran status, sexual orientation, gender identity, or genetic information. For more information, including how to report a concern, please 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

Visit the UArizona COVID-19 page for regular updates.

Campus Health

http://www.health.arizona.edu/

Campus Health provides quality medical and mental health care services through virtual and in-person care. Voluntary, free, and convenient COVID-19 testing is available for students on Main Campus. COVID-19 vaccine is available for all students at Campus Health.

Phone: 520-621-9202

Counseling and Psych Services (CAPS)

https://health.arizona.edu/counseling-psych-services

CAPS provides mental health care, including short-term counseling services.

Phone: 520-621-3334

The Dean of Students Office's Student Assistance Program

http://deanofstudents.arizona.edu/student-assistance/students/student-assistance

Student Assistance helps students manage crises, life traumas, and other barriers that impede success. The staff addresses the needs of students who experience issues related to social adjustment, academic challenges, psychological health, physical health, victimization, and relationship issues, through a variety of interventions, referrals, and follow up services.

Email: DOS-deanofstudents@email.arizona.edu

Phone: 520-621-7057

Survivor Advocacy Program

https://survivoradvocacy.arizona.edu/

The Survivor Advocacy Program provides confidential support and advocacy services to student survivors of sexual and gender-based violence. The Program can also advise students about relevant non-UA resources available within the local community for support.

Email: survivoradvocacy@email.arizona.edu

Phone: 520-621-5767

Campus Pantry

Any student who has difficulty affording groceries or accessing sufficient food to eat every day, or who lacks a safe and stable place to live and believes this may affect their performance in the course, is urged to contact the Dean of Students for support. In addition, the University of Arizona Campus Pantry is open for students to receive supplemental groceries at no cost. Please see their website at: *campuspantry.arizona.edu* for open times.

Furthermore, please notify me if you are comfortable in doing so. This will enable me to provide any resources that I may possess.

Preferred Gender Pronoun

This course affirms people of all gender expressions and gender identities. If you prefer to be called a different name than what is on the class roster, please let me know. Feel free to

correct instructors on your preferred gender pronoun. If you have any questions or concerns, please do not hesitate to contact me directly in class or via email (instructor email). If you wish to change your preferred name or pronoun in the UAccess system, please use the following guidelines:

Preferred name: University of Arizona students may choose to identify themselves within the University community using a preferred first name that differs from their official/legal name. A student's preferred name will appear instead of the person's official/legal first name in select University-related systems and documents, provided that the name is not being used for the purpose of misrepresentation. Students are able to update their preferred names in UAccess.

Pronouns: Students may designate pronouns they use to identify themselves. Instructors and staff are encouraged to use pronouns for people that they use for themselves as a sign of respect and inclusion. Students are able to update and edit their pronouns in UAccess.

More information on updating your preferred name and pronouns is available on the Office of the Registrar site at https://www.registrar.arizona.edu/.

Safety on Campus and in the Classroom

Familiarize yourself with the UA Critical Incident Response Team plans:
https://cirt.arizona.edu/

Department of Computer Science Evacuation Plan for Gould-Simpson:
https://drive.google.com/file/d/1iR1IcGcV_BgbGnEFBzZ2-do0FbLC3cvo/view?usp=sharing

Also watch the video available at https://ua-saem-aiss.narrasys.com/#/story/university-of-arizona-cert/active-shooter

Confidentiality of Student Records

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

Land Acknowledgement Statement

We respectfully acknowledge the University of Arizona is on the land and territories of Indigenous peoples. Today, Arizona is home to 22 federally recognized tribes, with Tucson being home to the O'odham and the Yaqui. Committed to diversity and inclusion, the University strives to build sustainable relationships with sovereign Native Nations and Indigenous communities through education offerings, partnerships, and community service.

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.