CSC 696H: Advanced Topics in Artificial Intelligence

(Topic: Probabilistic Methods in Machine Learning)

Mon / Wed : 2:00pm – 3:15pm : Gould-Simpson, Rm 701

Course Description

The aim of this course is to explore advanced techniques in probabilistic graphical models (PGMs) and statistical machine learning (ML) more broadly. Students will develop the ability to apply these techniques to their own research. Students will learn to perform statistical inference and reasoning in complex probabilistic statistical models. The course will survey state-of-the-art ML research including: variational inference, Bayesian Deep Learning, representation learning, and uncertainty quantification. Upon conclusion of this course students will be capable of developing new methods and advancing the state-of-the-art in ML and PGM research.

Instructor and Contact Information

Instructor:

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

Office Hours: TBD

Web Information:

Course Homepage: http://pachecoj.com/courses/csc696h_spring2024

D2L: https://d2l.arizona.edu/d2l/home/1415589

Piazza: https://piazza.com/arizona/spring2024/csc696h1

Instructor Homepage: http://pachecoj.com

Course Format and Teaching Methods

Each class will combine instructor lectures with student presentations of assigned readings. In addition, each student must present the results of a term project of their choosing during the final period.

Course Objectives

The primary objective of this course is to provide students a deeper understanding and familiarity with advanced topics in PGMs and statistical ML, as well as recent research in these areas. Our primary focus will be algorithms for Bayesian posterior inference in hierarchical statistical models. Planned topics will include:

- Introductory foundations
  - Probability and Statistic Primer
  - Bayesian Inference
  - Probabilistic Graphical Models
  - Exponential Families
- Inference
  - Monte Carlo Methods
  - Variational Inference
  - Implicit Inference

Bayesian Deep Learning
- Bayesian Neural Networks
- Monte Carlo Dropout
- Variational Autoencoder

Representation Learning
- Information Bottleneck
- Information Dropout

Uncertainty Quantification
- Variational Information Bounds
- Mutual Information Neural Estimation
- Contrastive Predictive Coding

Expected Learning Outcomes

Upon conclusion of this course, students will have learned to read, understand, and critique machine learning research articles. Students will demonstrate their understanding of articles through class presentations and critical summaries. Through assigned readings, students will be familiar with state-of-the-art research in key areas of statistical machine learning and PGMs. Students will apply these techniques to their own research, which will be assessed through term projects.

Makeup Policy for Students Who Register Late

Late registrants are required to complete all reading assignments that were missed. They will be required to submit reading summaries of the assignments and to view lecture recordings for the aforementioned topics.

Course Communications

Online communication will be conducted via Piazza. Grade feedback will be provided via D2L.

Required Texts or Readings

Required readings will be made available electronically, and on a weekly basis, as the course progresses.

Assignments and Examinations: Schedule/Due Dates

There will be a required reading associated with most lectures. The student will be required to summarize the reading in a brief writeup to be submitted with each lecture. There will also be a final project report due on the final exam date Fri, 5/3.

Final Examination

The course will not have a final exam but will instead have a semester project. The final project report will be due on the scheduled exam date Fri, 5/3.

Grading Scale and Policies

Students shall submit a critical summary for each assigned reading. To receive full credit, summaries must demonstrate that the student has adequately read and critiqued the material. Each student will additionally select among assigned papers and prepare two separate 1hr

presentations to the class in which they explain key technical details of the reading. To receive full credit for class participation, students must attend and participate in the discussion of all classes. Students should contact the instructor regarding absences for make-up. Finally, term project grading will be assessed based on how well the idea is conceived, planned, executed, and presented.

Grading Breakdown

Paper presentation:	20%
Critical reading summaries:	10%
Term project proposal:	30%
Term project (presentation and writeup):	40%

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: The student is required to notify the instructor of any grading dispute within 1 week of receiving the grade.

Scheduled Topic and Activities

Week 1: Introduction + Course Mechanics
Week 2: Probability Primer
Week 3: Inference
Week 4: Inference (Variational)
Week 5: Inference (Implicit Models)
Week 6: Inference (Implicit Models)
Week 7: Bayesian Deep Learning
Week 8: Bayesian Deep Learning (Monte Carlo Dropout)
Week 9: Bayesian Deep Learning (Bayesian Neural Networks)
Week 10: Bayesian Deep Learning (Variational Autoencoders)
Week 11: Bayesian Deep Learning (Variational Dropout)
Week 12: Representation Learning (Information Bottleneck)
Week 13: Representation Learning (Information Dropout)
Week 14: Uncertainty Quantification
Week 15: Uncertainty Quantification (CPC / MINE / DAD)
Week 16: Project Presentations
Week 17: Project Presentations

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/common/learningeventdetail/crtfy0000000000003560

University-wide Policies link

Links to the following UA policies are provided here, http://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.