ASEN 6519 Probabilistic Algorithms for Aerospace Autonomy
Spring 2019 Course Syllabus

General Information

Instructor: Prof. Nisar Ahmed (`nisar.ahmed@colorado.edu`)

Time and Location: Tues & Thurs 5:00 pm - 6:15 pm, Duane G2B47.

Course Canvas Website: `canvas.colorado.edu` (will be used for posting all course materials and announcements)

Office Hours: TBD (other times by appointment only)

Course Textbook:

Mykel Kochenderfer, *Decision Making Under Uncertainty: Theory and Application*, MIT Press, 1st edition, 2015. ISBN-13: 978-0262029254.

Optional resources: the following texts are recommended (but not required) for a deeper treatment of the core probabilistic AI, machine learning and pattern recognition concepts to be covered in this course, as well as many other important/useful topics not covered:

David Barber, *Bayesian Reasoning and Machine Learning*, Cambridge University Press, 2012: **available free online**: http://web4.cs.ucl.ac.uk/staff/D.Barber/pmwiki/pmwiki.php?n=Brml.HomePage

Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 1st edition, 2007. ISBN 978-0387310732.

Finn V. Jensen and Thomas D. Nielsen, *Bayesian Networks and Decision Graphs*, Springer, 2nd edition 2007. ISBN 978-0387682815.

Sebastian Thrun, Wolfram Burgard, and Dieter Fox, *Probabilistic Robotics*, MIT Press, 2005. ISBN 978-0262201629.

Sergios Theodoridis, *Machine Learning: A Bayesian and Optimization Perspective*, Academic Press, 1st edition, 2015. ISBN 978-0128015223.

Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Pearson, 3rd edition, 2009. ISBN 978-0136042594.

Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, Bradford, 1st edition, 1998. ISBN 978-0262193986.

Course Details

Description Autonomous systems learn, adapt, and make decisions under uncertainty. Probabilistic algorithms for machine learning and decision-making are crucial for modern autonomous robotic and unmanned aerospace systems, and are being adopted by many other engineering/science disciplines as well. **This advanced graduate course will provide an introduction to modern probabilistic machine learning and AI techniques that allow autonomous systems to make decisions under uncertainty.** As this course is designed with aerospace engineering graduate students in mind, it will present topics from a rigorous engineering viewpoint (as opposed to a more conventional rigorous computer science perspective) and build on fundamental knowledge of probability, estimation and control theory, with a strong emphasis on applications via programming exercises/projects and discussions of current research. Major topics to be covered include:

- probabilistic graphical models and applications (hidden Markov models, Bayes nets, and others such as Markov random fields, decision graphs, Bayesian nonparametric models, etc. as time permits);
- model learning and parameter estimation for decision making and pattern recognition applications;
- inference methods: exact methods for basic probabilistic models and commonly used approximate methods for more complex models (Monte Carlo);
- sequential decision making and dynamic programming: Markov decision processes (MDPs) and partially observable MDPs (POMDPs);
- reinforcement learning and related topics, e.g. multi-armed bandits, inverse reinforcement learning;
- survey of advanced/current research topics (as time and interest permits): deep learning and neural networks; fully Bayesian inference and learning techniques; probabilistic programming; human-autonomy interaction and explainable/introspective AI.

While these concepts will be covered in a mathematically rigorous way, students will apply these concepts to application problems developed from their own research. Students will develop their own software that could, for instance, serve as the basis for decision-making on board an unmanned ground robot, air vehicle, spacecraft, or other application platform/system that connects to their own research problems. Students will be expected to develop and refine their project application throughout the semester, by incorporating course material into their problem and culminating in a short final project presentation and report. Toward the latter part of the course, students will also lead discussions of work from leading robotics, machine learning, and AI conferences and journals.

Prerequisites: Previous coursework in probability/statistics at least at level of ASEN 5044: Statistical Estimation for Dynamical Systems is **required**; formal linear algebra and control theory at least at level of ASEN 5014: Linear Control Systems highly recommended. Students must be comfortable with a technical programming language for software exercises and projects (e.g. Matlab/Octave, Python, C/C++, C#, Java, Julia, R, or similar).

Grading and Project Assignments Course work will largely be project-oriented. There will be no exams. Several required topical exercises related to the lectures will be posted to ensure that students demonstrate understanding of the course material, as well as to provide periodic feedback and guidance as students try to integrate/explore concepts into their final projects. These exercises will consist of short theoretical and programming problems for toy applications, as well as questions to guide the development of final project applications. All exercises will be graded on a binary 'satisfactory' (S)/'unsatsifactory' (U) scale. To receive full credit for these, students must submit and receive a 'satisfactory' grade on at least 3 out of the 5 exercises to be posted. Students will be allowed to resubmit 'unsatisfactory' assignments for a regrade, as long as the initial assignments are submitted in a timely manner. Exercises are expected to be posted every other week, to coincide with major lecture topics being covered around that time. The final project will be developed over the course of the semester via the exercises. Students have the option of working together in groups of two (max) on the exercises and the final project if they so choose, though some level of individual contributions/work will be expected on group projects. In any case, students are highly encouraged to collaborate with one another and to constantly think about how best to connect course material to their own research.

Grading breakdown: exercises: 20%; final project: 60% (30% in class presentation + 30% report); class participation: 20%. Note that any group exercise submissions and final project submissions will result in the same grade for both group members.

Benefits and Learning Objectives This course will enable students to:

1. combine knowledge of probability and statistics with engineering/science domain knowledge to formulate mathematically sound models of uncertain dynamical systems and autonomous decision making problems;

2. identify and design algorithms that enable autonomous learning, inference and decision making with probabilistic models, and explain their similarities/differences.

3. define, explain and demonstrate the importance of fundamental tools, including (but not limited to): probabilistic graphical models; point estimation and maximum likelihood techniques; fully Bayesian and Monte Carlo inference/learning methods; dynamic programming; MDPs/POMDPs; reinforcement learning;

4. identify and explain strengths/weaknesses of state of the art autonomous reasoning algorithms;

5. develop and implement software to simulate and evaluate the performance of autonomous agents for real-world/research applications.

Tentative Course Schedule

Week(s)	Topic
1	Course intro & overview
1-2	Probabilistic graphical models: HMMs, Bayesian Nets, etc.
2-4	Basic inference techniques and maximum likelihood parameter learning
5-8	Online decision making under uncertainty: MDPs, POMDPs
8-10	Reinforcement Learning
–	SPRING BREAK
11-15	Fully Bayesian inference/learning and other advanced topics; final project short talks

General Policies If you qualify for accommodations because of a disability, please submit to your professor a letter from Disability Services in a timely manner (for exam accommodations provide your letter at least one week prior to the exam) so that your needs can be addressed. Disability Services determines accommodations based on documented disabilities. Contact Disability Services at 303-492-8671 or by e-mail at dsinfo@colorado.edu.

If you have a temporary medical condition or injury, see Temporary Medical Conditions: Injuries, Surgeries, and Illnesses guidelines under Quick Links at Disability Services website and discuss your needs with your professor.

Campus policy regarding religious observances requires that faculty make every effort to deal reasonably and fairly with all students who, because of religious obligations, have conflicts with scheduled exams, assignments or required attendance. In this class, accommodations can be made two weeks in advance. See full campus policy details at http://www.colorado.edu/policies/fac_relig.html

Students and faculty each have responsibility for maintaining an appropriate learning environment. Those who fail to adhere to such behavioral standards may be subject to discipline. Professional courtesy and sensitivity are especially important with respect to individuals and topics dealing with differences of race, color, culture, religion, creed, politics, veteran's status, sexual orientation, gender, gender identity and gender expression, age, disability, and nationalities. Class rosters are provided to the instructor with the student's legal name. I will gladly honor your request to address you by an alternate name or gender pronoun. Please advise me of this preference early in the semester so that I may make appropriate changes to my records. See policies at

- http://www.colorado.edu/policies/classbehavior.html
- http://www.colorado.edu/studentaffairs/judicialaffairs/code.html#student_code

The University of Colorado Boulder (CU-Boulder) is committed to maintaining a positive learning, working, and living environment. The University of Colorado does not discriminate on the basis of race, color, national origin, sex, age, disability, creed, religion, sexual orientation, or veteran status in admission and access to, and treatment and employment in, its educational programs and activities. (Regent Law, Article 10, amended 11/8/2001). CU-Boulder will not tolerate acts of discrimination or harassment based upon Protected Classes or related retaliation against or by any employee or student. For purposes of this CU-Boulder policy, "Protected Classes" refers to race, color, national origin, sex, pregnancy, age, disability, creed, religion, sexual orientation, gender identity, gender expression, or veteran status. Individuals who believe they have been discriminated against should contact the Office of Discrimination and Harassment (ODH) at 303-492-2127 or the Office of Student Conduct (OSC) at 303-492-5550. Information about the ODH, the above referenced policies, and the campus resources available to assist individuals regarding discrimination or harassment can be obtained at http://www.colorado.edu/odh

All students of the University of Colorado at Boulder are responsible for knowing and adhering to the academic integrity policy of this institution. Violations of this policy may include: cheating, plagiarism, aid of academic dishonesty, fabrication, lying, bribery, and threatening behavior. All incidents of academic misconduct shall be reported to the Honor Code Council (honor@colorado.edu; 303-735-2273). Students who are found to be in violation of the academic integrity policy will be subject to both academic sanctions from the faculty member and non-academic sanctions (including but not limited to university probation, suspension, or expulsion). Other information on the Honor Code can be found at http://www.colorado.edu/policies/honor.html and at http://www.colorado.edu/academics/honorcode/