Introductory Information

Course: 582636, Probabilistic Models (4 cr)
Semester: Spring 2014, Period 3-3

Class Meeting: Exactum D122, Tuesday and Thursday, 16:15 - 18
Lecturer: Brandon Malone
Email: brandon.malone@cs.helsinki.fi
Office: Exactum, A323
Office Hours: Wednesday, 10 - 12, or by appointment

Lab Meeting: Exactum C222, Wednesday, 16:15 - 18
Instructor: Quan Nguyen (Eric)
Email: quan.nguyen@cs.helsinki.fi
Office: Exactum, A339

Course website: https://www.cs.helsinki.fi/en/courses/582636/2014/k/k/1

Course Description

Course Overview: This course provides an introduction to probabilistic modeling from a computer scientist's perspective. Many of the research issues in Artificial Intelligence, Computational Intelligence, Informatics and Machine Learning/Data Mining can be viewed as topics in the "science of uncertainty," which addresses the problem of optimal processing of incomplete information, i.e., plausible inference, and this course shows how the probabilistic modeling framework forms a theoretically elegant and practically useful solution to this problem. The course focuses on the "degree-of-belief" interpretation of probability and illustrates the use of Bayes' Theorem as a general rule of belief-updating. As a concrete example of methodological tools based on this approach, we will study probabilistic graphical models focusing in particular on (discrete) Bayesian networks and on their applications in different probabilistic modeling tasks.

This is an introductory course, and only elementary knowledge on probability theory is required as a prerequisite. Furthermore, we will approach the class from a computer scientist's perspective. Consequently, we will often emphasize algorithms and data structures over probability theory. Nevertheless, the probability theory underpins everything, so we will cover that in detail, as well.

Course Technical Information: This course belongs to the Algorithms and Machine Learning sub-programme in the Master's programme of the department, and together with *582637 Project in Probabilistic Models (2 cr)*, it forms one of the three optional courses of the sub-programme.

For students in the old Intelligent Systems specialisation area: this course, together with the project work *582637 Project in Probabilistic Models (2 cr)*, replaces the course *Three Concepts: Probability (6 cr)*.

Course Prerequisites (suggested): *582630 Design and analysis of algorithms (4 cr)* and *582631 Introduction to machine learning (4 cr)*.

Course Objectives: The objectives for this course are available here:
http://www.cs.helsinki.fi/u/bmmalone/probabilistic-models-spring-2014/LearningObjectives.pdf.

Textbooks and Other Resources: *Modeling and Reasoning with Bayesian Networks*, by Adnan Darwiche, ISBN-13 9780521884389. *Probabilistic Graphical Models: Principles and Techniques*, by Daphne Koller and Nir Friedman, ISBN-13 9780262013192. It is not necessary to purchase the textbooks, but they are both good references for probabilistic (graphical) models.

Furthermore, we will use several of the videos from Daphne Koller's Coursera course on probabilistic graphical models (https://www.coursera.org/course/pgm). Other online materials, including recorded tutorials, survey papers, etc., will also be incorporated into the class. All of these will be made available by the week for which the resource is required on the course website. Slides from each lecture will be made available after class.

Course Delivery: For this course, we will use the 'flipped classroom' model. That is, we will have videos and other materials available before class meetings. Students will be expected to watch, read, etc., this material before class meetings. Then, during class, rather than spending time covering basic material, we can focus on the more challenging aspects of the content. In particular, class meetings are about 2 hours long. Most meetings will be split across two or three topics. For each topic, I will spend 15 to 20 minutes recapping the material given before class. We will then split into small groups to discuss particular aspects of the topic in more detail, again for 15 to 20 minutes. Then, we will come back together as a whole and review the findings of the groups. Of course, some topics are more amenable to this format than others, so individual meetings may vary.

Student Responsibilities

Learning is a two-way street. I commit to providing you with all of the tools necessary to be successful in this course, including fair assessments, appropriate deadlines and timely feedback. In return, I expect you to take responsibility for your own learning. This includes attending class, watching the videos, reading the papers and completing the exercises.

This is a 4 cr course, so that means you should spend about 110 hours working on this course. We have about 25 hours of in-class time, as well as another 10 hours of lab time, so that means you should expect to spend about 75 hours outside of class meetings on this course.

Evaluation

The grade for this course will be determined by two components: five weekly exercises and a final exam. The maximum number of points that can be earned from the exercises is 25 and from the course exam is 35, so the total maximum is 60 points. The following list gives the number of points required for each grade:

1. 30 - 35 points
2. 36 - 41 points
3. 42 - 47 points
4. 48 - 53 points
5. 54 - 60 points

Additionally, you must score at least 20 points on the final exam to pass the course.

Please alert me as soon as possible if you need the exercises or the final exam in a language other than English!

Weekly exercises: The weekly exercise sets will be posted online at least one week in advance. Typically, the exercise sets will cover the material from class, although they may sometimes include material we do not discuss. In these cases, sufficient references will be given to make the exercises clear. The solutions must be submitted to the lab instructor (Eric) before the lab meeting of the given week electronically. If you handwrite your solutions, please scan the solutions. Even though we will work in small groups, each student is expected to submit their own set of solutions.

There will be five exercise sets. In order to acclimate everyone to the grading, the initial sets will be worth fewer points than the later ones. **The following list gives the due dates and point values of all sets**:

1. January 22, 3 points
2. January 29, 5 points
3. February 5, 5 points
4. February 12, 6 points
5. February 19, 6 points

Every effort will be made to grade and return the exercise sets by the next lab meeting. Furthermore, the lab meetings will mostly consist of forming small groups and discussing the exercises from the previous week. Ideally, this feedback will allow all students to identify areas on which they should focus.

Final exam: The final exam period is **Monday, February 24, 16:00 - 18:30 in B123** (N.B. it starts at 16:00, NOT 16:15). Students will be allowed to use a simple calculator and one A4 handwritten (not photocopied) sheet of notes. The exam is worth 35 points, and students must score at least 20 points on the exam to receive a grade of 1 or more.

Class participation: Students are not required to attend class meetings; however, the meetings are the primary venue to address challenges encountered. Furthermore, during the small group work, everyone in the group is expected to participate. The small groups also offer an excellent opportunity to discuss questions with peers.

Late work: Exercise sets may not be submitted late. The final exam cannot be taken late. For extenuating circumstances, conferences, etc., please contact me in advance to make arrangements.

Tentative Schedule

The following list gives a tentative schedule of topics that will be covered in this course. The course website gives more information, including necessary resources and a summary of which objectives will be covered during each meeting. The course website also gives the exercise sets, relevant objectives, and solutions (after the sets have been returned).

- January 14: Introduction
- January 16: Refresher on probability theory
- January 21: Bayesian networks
- January 23: Special models
- January 28: Inference in simple models (belief propagation algorithm)
- January 30: Inference in general models (join tree algorithm)
- February 4: Parameter inference (multinomial distribution)
- February 6: Predictive inference
- February 11: Hidden variables and missing values (discrete Bayesian networks)
- February 13: Parameter inference (topic models)
- February 18: Model selection
- February 20: Structure learning

Course Policies

Attendance: Attendance is not required for this course; nevertheless, it is highly encouraged to take full advantage of the lecture, discussion and lab sessions.

Class Behavior: Please be respectful of the rest of your classmates. Arrive to class on-time. Make arrangements with me if you need to leave early. Similarly, please turn the volume of your cell phones off.

Withdrawal: Please check the registrar's web page at http://www.helsinki.fi/studying/enrolment_and_registration.shtml for the last day to withdraw from the course without a grade, with a "W" and from the university.

Incompletes: Except under extenuating circumstances, a grade of "Incomplete" will not be given in this course.

Academic Dishonesty: From the University of Helsinki blog: "Cheating means breaching good scientific and study practice through dishonesty. A person who cheats in their research or studies violates the accepted rules of good work conduct either intentionally, in order to mislead, or due to carelessness (ignorance or negligence). Ignorance is no defence; if you do not know what to do, seek guidance."

For additional information please visit: http://blogs.helsinki.fi/alakopsaa/what-constitutes-cheating/?lang=en

For purposes of this class, students are allowed to discuss general strategies for solving the problem using examples from the textbook, lecture or other approved material, but students may not share work. If a student uses text or ideas from any source, he or she must mention the source in the relevant references section for that assignment. Both sharing work with another student and failing to mention the text and ideas taken from another source will constitute plagiarism. Any student involved in plagiarism will be reported. The professional conduct policy of the CS department and UH will be followed are enforced. This course will involve work in small groups, though, and it is acceptable for students in the same group to have similar answers; simple comments in your own words, e.g., for what happens during each step in a derivation is sufficient to distinguish the responses of different group members.

Acknowledgements

Much of the text was adapted from previous versions by Petri Myllymäki and Pekka Parviainen.

This syllabus was designed with text and ideas from the following sources:
MSU, CSE 1284, designed by Dr. T.J. Jankun-Kelly and Mrs. Lisa Henderson
http://www.gccaz.edu/facdev/syllabus.htm
https://tle.wisc.edu/forum/syllabus-statements-about-special-needs-diversity-and-academic-honesty