Course Number and Title: ISTA 410/510– Bayesian Modeling and Inference

Course Description:
Bayesian modeling and inference is a powerful modern approach to representing the statistics of the world, 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 with examples ranging from computer vision to social sciences to astronomy.

Prerequisite(s): Undergraduates: ISTA 130 (or equivalent programming course), Math 215 (Linear Algebra---implies calculus), ISTA 116 or ISTA 4421 (or other probability, statistics, or machine learning courses), or permission of instructor.

Grad students: Exposure to the basics of probability, calculus, linear algebra, and some experience with programming.

Students missing aspects of technical preparation will be expected to study/review background material as needed.

Units: 3 (lecture)

(1) **Instructor information**:

Kobus Barnard
kobus@sista.arizona.edu
Gould-Simpson 927A

(2) **Office hours:**

Arrange via email 18 hours in advance. Good times to propose will be posted.

(3) **Course Description and Objectives**:

To develop a solid fundamental understanding of Bayesian models and associate inference methods and how to apply them to diverse problems. Skills developed will include: 1) creating graphical models for data; 2) specifying distributions for parameters of model components that link the model to data; 3) exact inference methods to estimate model parameters for discrete Bayes nets with tree topology (sum product and max sum algorithms); 5) Markov random fields; 6) commonly used models including Gaussian Mixture Model (GMM), Hidden Markov Model (HMM), and Linear dynamical systems (LDS) 7) learning model parameters from data using approximate inference methods including Expectation Maximization (EM), and MCMC sampling including Metropolis-Hastings (MH), Gibbs, and Hybrid Monte Carlo (HMC);

Topics:

- Probabilistic foundations
- Introduction to the Bayesian methodology and introductory examples
- Representing models using graphs
- Inference for discrete Bayes nets (sum product and max sum algorithms)
- Markov random fields
- Gaussian Mixture Model (GMM)
- Expectation Maximization (EM)
- Hidden Markov Models (HMM) and Linear dynamical systems
- MCMC sampling (subtopics include Metropolis-Hastings (MH), Gibbs sampling, and Hybrid Monte Carlo (HMC))

(4) **Grade policies:**

For both undergraduates and graduate students, the grading breakdown will be as follows:

Assignments:	60%	(there will be 6-10 assignments)
Midterms:	20%	(there will be two midterms).
Final Exam:	20%	

As a function of the enrolled students, and the discretion of the instructor, any subset of the exams can be take-home exams. In 2013, all exams were take-home, which worked well.

Assignments and take-home exams have typically have basic components and advanced components. Undergraduate students will be responsible for the basic components. Graduate students will be responsible for both components.

Graduate students will be expected to perform at a higher level on any in-class exams for the same exam grade. In addition, undergraduate students will not be responsible for exam questions specific to grad student assignment topics.

For both undergraduates and graduate students, 90% guarantees an A, 80% guarantees a B, 70% a C, and 60% a D.

(5) **Absence policy:**

Good attendance is expected. 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.

All holidays or special events observed by organized religions will be honored for those students who show affiliation with that particular religion. Absences pre-approved by the UA Dean of Students (or Dean's designee) will be honored. If the final exam coincides with convocation, and students in the class would like to attend, an alternative final exam time will be provided.

(6) **Required Text(s):**

There is no required text. All needed material will be posted on-line. However, this course will draw material from three excellent and highly recommended books:

Pattern Recognition and Machine Learning by Chris Bishop.

Probabilistic Graphical Models: Principles and Techniques by Daphne Koller and Nir Friedman. This text is very extensive. Its introduction (http://mitpress.mit.edu/books/chapters/0262013193chap1.pdf) gives a flavor of the course.

Machine Learning: A Probabilistic Perspective by Kevin Murphy

(7) **Number of required examinations and papers:**

See grade policies above.

(8) **Policies on expected classroom behavior:**

Mature adult behavior is expected. Students are requested to keep their cell phones, pagers, and similar devices on mute or manner mode.

(9) **Policies against plagiarism:**

Students are bound by the Code of Academic Integrity (see http://catalog.arizona.edu/policies/974/acacode.htm). This codes prohibits all forms of student academic dishonesty, including but not limited to cheating, fabrication, and plagiarism. Violations can result in serious penalties, including expulsion from the University.

Exams and written assignment must be the sole work of the student (or student team). Students may help each other with the problem analysis and general strategies relevant to the programming assignments, but detailed help or code sharing is not permitted. All code in programming assignments will be assumed to have been written by the student (or student team) unless attribution is given. An obvious exception to this rule is sample code which has been provided by the instructor for this course through the course web page tree. Such code does not require attribution (we know where it came from). It is also permissible to include with attribution code from external sources provided that the code is published, has not been solicited, and was not written for course requirement for this or a similar course given elsewhere.

(10) **Policies against threatening behavior:**

Students are reminded that the Student Code of Conduct (5-308.F.11) dictates that no person or organization may interfere with or threaten University-sponsored classroom activities (see http://web.arizona.edu/~policy/threaten.shtml).

(11) **Required extracurricular activities:**

There are no required extracurricular activities.

(12) **Required or Special Materials:**

None.

(13) **Notice of Potentially Offensive Material:**

None.

(14) **Notice to students with disabilities:**

Students with disabilities, who may require academic adjustments or reasonable accommodations in order to participate fully in course activities or to meet course requirements, must first register with the Disability Resource Center (DRC) (http://drc.arizona.edu).

DRC staff will qualify students for services, and provide a letter to the instructor listing accommodations to be made. This letter should be submitted by the student directly to the instructor as soon as possible during the first week of classes.

The student should meet as soon as possible with the instructor by appointment or during office hours to discuss accommodations and how course requirements and activities may impact your ability to fully participate.

(15) **"Subject to Change" Statement:**

The instructors reserve the right to change with advance notice where appropriate the content of the course. This right does not apply to posted grading and absence policies.