Welcome to ISTA 410/510

Bayesian Modeling and Statistics

Today

Course mechanics, syllabus, etc.
Brief course outline
Introduce the topic

Course mechanics

Course page is at: http://www.sista.arizona.edu/classes/ista410/spring12
(Linked from instructor's home page http://kobus.ca

Lectures and assignments will require either connecting from a UA machine, OR a login id ("me") and password ("bayes4fun").

Significant communication for the course will happen using the class mail list (ista410@listserv.arizona.edu).

Tentatively we will use the above mail list as a discussion board. But we could try Piazza.com if there is interest. (?)

Course mechanics (II)

Group office hours: Friday 10:30-11:30 in GS 919

For an individual appointment send email, with proposed availability (if possible) during likely open times as described at:
http://kobus.ca/calendar_info.html.

Current list of times

Monday: 9:30 - 10:30
Tuesday: 9:30 - 10:00
Thursday: 9:30 - 10:00

Course mechanics (III)

- Lecture note previews will be posted sporadically for those who want to look at them
  - The longer in the future the material is, the less accurate it will be!
- Official PDFs for lectures will be posted on the class web page after class.
- Videos for lectures will also be linked.

Course mechanics (IV)

- This course requires a SISTA/CS account. Apparently, if you have a UA email and are registered, an account is automatically created for you.
- The 9th floor lab will be available for this course.
- Assignments will be posted on the web page.
- Assignments will be handed in using the "turnin" program on the machine "lectura".
- The course will have both written and programming assignments
  - Key deliverable will typically be a PDF with answers, results, etc.
  - If programming was involved, code needs to be submitted.
  - Recommended programming language is Matlab
  - C/C++ is also an option (library support is available)
  - Others languages can also be used (but I won't look at the code)

Course mechanics (V)

Books and materials
- No required text (all material will be lecture notes and assignments)
- Important reference is Bishop (key chapters will be put online)
- Good reference (too extensive to be our text) is Koller and Friedman

Co-convened course
- Grad students will have longer assignments
- Grad students will be expected to do more/better on exams

Grade distribution
- Assignments: 60% (there will be 4-8 assignments)
- Midterms: 20% (there will be two midterms, likely take home)
- Final Exam: 20% (likely take home)

Participation in experiments
- Extra credit, TBA (an alternative will be available)

Additional policies and procedures available in syllabus linked from class page

Course outline

Blurb: To develop a solid fundamental understanding of Bayesian 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) applying inference methods to estimate model parameters; 4) setting up learning model structure from data; and 5) applying Bayesian methods to decision making processes.

Topics:
- Probabilistic foundations
- Introduction to the Bayesian methodology and introductory examples
- Representing models using graphs
- Inference for graphical models
- Learning model structure
- Actions and decisions

What is this course about?

Recommended reading

http://mitpress.mit.edu/books/chapters/0262013193chap1.pdf

Introductory Questions

What is a model?

What is a statistical model?

What is a parameterized model?

Why do we like probabilistic / statistical models?

What is a Bayesian model?

What do we mean by inference?

Simple example

Developed on white board

We observe a sequence of points ((x,y) coordinates) from an unknown physical process or sensor

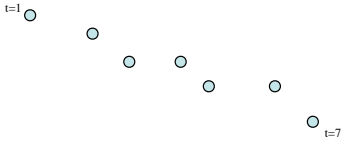

What are the statistical dependencies? The points do not seem independent!

We might declare a plausible model that the points are independent, conditioned on a line model.