Bayesian Inference (PSYCH 4KK3)

Winter Semester, 2012

Instructor
Dr. Daniel Goldreich
goldrd@mcmaster.ca
ext. 28666

Description

This course explores the far-reaching applications in science of a simple but remarkably insightful equation. First introduced in 1764, Bayes' formula elegantly embodies the scientific method, in which predictions made from hypotheses are put to the test, and inferences regarding the relative merits of the hypotheses are drawn from the resulting data. Bayes' equation is increasingly used as a powerful research tool by scientists in all disciplines. It provides a sophisticated method for drawing inferences from data, used both for statistical analysis and as a model of human brain function.

Students will apply Bayes' formula to gain insight into core concepts of scientific reasoning, such as induction and deduction, and heuristics such as Occam's Razor. They will learn modern statistical analysis techniques, derived from Bayes' formula, that provide a coherent and attractive alternative to the conventional p-value methods. This conceptually sophisticated course is presented in an interactive and mathematically accessible fashion.

The course meets twice per week in room BSB 106: Tuesdays 11:30 AM - 1:20 PM, and Thursdays 11:30 AM - 12:20 PM. An optional tutorial is held on Mondays, 9:30 - 10:20 AM, in ABB 162. Students are strongly encouraged to attend tutorial.

Lectures will be interspersed with critical thinking activities called Bayesian Brain Bogglers, which collectively account for 10% of the course grade. In addition, students will be graded on weekly homework exercises, two midterm tests, and one comprehensive final exam (see below for grading details).

The course web site plays an integral role in this course. The web site contains this syllabus, links to the online course textbook, homework assignments, and Bayesian inference information on the web, a form for anonymous feedback to the instructor, and a password-protected system for students to check their marks. The web site also contains "Bayes-Engage," an online discussion forum for the students.

Graduate students taking Psych 730 are required to attend all lectures and complete all work along with the undergrads in Psych 4KK3 (i.e, in-class bogglers, weekly homework assignments, midterm tests and final exam). The graduate students will write somewhat more advanced versions of the midterm tests and final exam than the undergraduates. In addition, the graduate students will complete three extra assignments involving extensive reading and/or data analysis.

Objectives

Upon completion of this course, the student will understand:

- the concept of conditional probabilities
- how to work with the sum and product rules of probability theory
- the meaning of each element of Bayes' formula: prior, likelihood, and posterior
- key features of the scientific method which naturally emerge from the application of Bayes' formula, such as Occam's razor, the importance of open-mindedness, and the importance of skepticism.
- how Bayesian inference differs from the traditional (frequentist) statistical method
- how to apply Bayesian inference to simple two-and three- hypothesis situations drawn from clinical diagnostic testing scenarios (sensitivity and specificity problems), genetics testing (inferring genotype from phenotype), and other such scenarios.
- how to use Bayesian inference to estimate the values of binomial, Poisson, and Gaussian parameters
- how to apply Bayesian inference to linear and non-linear curve-fitting
- how to perform Bayesian model comparison and Bayesian model averaging

Materials

Required Textbook: Goldreich, D. (2011) Introduction to Bayesian Inference (online)

Students are responsible for reading each book chapter in its entirety. Some homework and exam questions may be based on material in the book that was not covered during lectures.

Note: The online book is a work in progress and may be modified slightly throughout the semester. If you are planning on printing the book, please wait until we reach the corresponding course section before printing each chapter, to ensure that you print the most recent version.

Required Calculator: The McMaster Standard Calculator (Casio FX-991) is permitted in exams and for in-class Brain Boggler exercises unless otherwise specified. This calculator is available at the McMaster Bookstore. No other calculator is permitted.

Evaluation

The student is responsible for understanding all material covered in class (lectures and Brain Boggler exercises) and in the homework assignments. Any material covered in class or in homework assignments may appear on exams, including material that does not appear in the assigned readings. The evaluation scheme (below) applies equally to Psych 730 students and Psych 4KK3 students.

The student's course percentage score is a weighted average of the following five items:

Item	Weight

Bayes-Engage Forum Participation	3%
Bayesian Brain Bogglers	7%
Homework Assignments	15%
Midterm Test 1	15%
Midterm Test 2	15%
Comprehensive Final Exam	45%

Bayes-Engage Forum

Students are expected to participate throughout the semester in the online Bayes-Engage discussion forum. The forum serves to facilitate student understanding of Bayesian inference by allowing students to discuss homework assignments, course material, and related Bayesian topics while not in class.

Each student is expected to contribute at least 12 substantive forum posts during the semester, and to post at least once in each of the three forum categories. A substantive post is not merely a question or a request for assistance (such posts are certainly encouraged, but do not receive credit). Rather, a substantive post is one that contributes significantly to other students' understanding of Bayesian inference. Substantive posts include: an intelligent and articulate comment about a specific concept covered in lecture or in the textbook; a correct and well-reasoned answer to another student's posted question about a challenging Bayesian topic or problem; or an astute observation regarding Bayesian inference in everyday life.

At the end of the semester, the number of substantive posts made by each student will be tallied. If this number is greater than or equal to 12, the student will receive a Bayes-Engage percentage score of 100%. If the number of substantive posts, *n,* is less than 12, then the student's Bayes-Engage percentage score will be (*n* / 12) * 100%. For instance, if a student has made 10 substantive posts, the student's Bayes-Engage percentage score will be 83.33%.

Each student is expected to make at least one substantive post in each of the three forum categories. For each category in which a student has not made a substantive post, 10% will be deducted from the student's Bayes-Engage percentage score. For instance, if a student has made 10 substantive posts, but these have occurred in two categories only, with no substantive posts in the third category, then the student's final Bayes-Engage course percentage score will be 83.33% - 10% = 73.33%.

Bayesian Brain Bogglers

The brain bogglers are in-class exercises designed to encourage critical thinking about Bayesian inference.

Some of the brain boggler exercises will require a calculator; the student should bring the McMaster Standard calculator (Casio FX-991) to class.

On average, about one brain boggler will be given in each hour of class time.

Unless otherwise announced, each brain boggler exercise is worth 2 points. Each answer will receive either zero, half, or full-credit (0, 1, or 2 points).

The student's running brain boggler total is reported online in the check-marks system.

The student's brain boggler percentage score, used in the student's course percentage score calculation, is the number of brain boggler points earned divided by the total number of points possible.

Homework

Homework assignments will be linked from the schedule table (above) each Tuesday, and are due the following Tuesday at the beginning of class. Late homework returns will not be accepted.

Students are encouraged to type their answers whenever possible.

Students may collaborate on homework assignments, if they wish, and collaborating students may choose to submit a single, group answer (maximum group size: 3 students). If this is done, then the name and student number of each contributing student must appear on the homework submission. Each student in the group will receive the same mark on that homework assignment. Students may change groups throughout the semester, or choose to submit their assignments individually.

The point value of each question will be indicated on the homework assignment.

The score for the entire homework assignment will always be reported on a 0-to-100% scale. For example, suppose an assignment has three questions, worth 2, 4, and 4 points each. If a student earns half credit on the first question, full credit on the second, and half credit on the third, then the student's score will be 1 + 4 + 2 = 7, and the score will be reported as 70%.

The student's homework assignment scores are reported online in the check-marks system.

At the end of the term, the student's lowest homework assignment score will be dropped. The average of the student's remaining homework assignment scores will then be calculated. This is the student's course homework score.

Graduate students taking Psych 730 will complete three additional homework assignments (identified as GS HW A, B, and C in the schedule table above). These assignments, which will involve extra reading and/or data analysis, are intended to deepen the graduate students' understanding of the material covered in each of the three course sections. These three homework assignments will be weighted equally with the other assignments in the calculation

of the student's course homework score (i.e., whereas undergrads complete 10 homework assignments, grad students complete 13). Grad students must complete all three GS HW assignments (A, B, and C). The lowest GS HW grade will not be dropped (only the lowest of the 10 other assignments will be dropped).

Tests and Exam

The student should bring the McMaster standard calculator (Casio FX-991) to all exams. Only the McMaster standard calculator will be allowed.

Midterm test 1 covers material from the first course segment (basic Bayesian inference and parameter estimation). Midterm test 2 covers material from the second course segment (curve fitting, robust estimation & advanced parameter estimation). The final exam is comprehensive; it covers material from all three course segments.

Graduate students taking Psych 730 will write somewhat more advanced versions of the tests and exam than the undergraduates taking Psych 4KK3.

Each test / exam is given a mark on a scale from 0 - 100%.

The student's test marks are reported online in the check-marks system.

Create-a-Question: As an optional exercise, students are encouraged to try to generate an excellent exam question of their own. Please email the instructor with your proposed exam question and answer. Questions that are carefully and articulately worded, and that probe student understanding of important concepts, will be considered for inclusion. The instructor will not inform you in advance of the exam whether your question will be used, and will not provide feedback as to whether your answer is correct. If it is used on an exam, your name will not be attached to your question, but a note will be attached to indicate that the question was student-generated. Furthermore, if it is used, your question may be edited and/or otherwise modified by the instructor.

Course Grade Calculation

The following formula is used to calculate the student's course percentage score:

Course percentage score = (course Bayes-Engage percentage score)(0.03) + (course Brain Boggler percentage score)(0.07) + (course Homework score)(0.15) + (Midterm test 1score) (0.15) + (Midterm test 2 score) (0.15) + (Final Exam score) (0.45)

The student's course grade will be determined from the student's course percentage score, as follows:

Letter Grade	**Percentage Score**

A+	90-100
A	85-89
A-	80-84
B+	77-79
B	73-76
B-	70-72
C+	67-69
C	63-66
C-	60-62
D+	57-59
D	53-56
D-	50-52
F	0-49

Missed Work

If a student misses a lecture or a test due to illness, personal circumstances, or late registration, it is the student's responsibility to notify the instructor and to complete a McMaster student absence form, or to submit appropriate documentation (e.g. note from physician) to the appropriate Faculty/Program office. A mark of zero will be entered for all missed Brain Boggler exercises, homework assignments, and tests until the instructor receives notification from the appropriate university authority. (If a student is unable to attend a class, then the student may submit the homework assignment by email, prior to the class period, to the instructor. Late email submissions will not be accepted.) It is the student's responsibility to learn all material that the student missed for any reason; this can be done by consulting with classmates, the teaching assistants, and/or the instructor.

Missed class periods: If university approval is obtained for missed class periods, the total number of Brain Boggler points in the course used to calculate the student's Brain Boggler percentage score will be reduced by the number of excused Brain Bogglers that the student missed. Similarly, if university approval is obtained for missed homework, the total number of homework points in the course used to calculate the student's homework percentage score will be reduced by the number of excused homework assignments that the student missed.

Missed Tests: If university approval is obtained, tests missed due to illness or personal circumstances may be made up, or the course grade may be redistributed such that more weight is applied to the other test / exams, at the discretion of the instructor. Make-up tests, if given, may differ in format and specific content from the original test. For example, a test consisting of written questions (1-hour) and oral questions (1-hour) may be given in place of a 2-hour written original test.

Audio Recording

Students may make audio recordings of the lectures, after asking permission from the instructor. Audio recordings must be used for personal study purposes only, and not posted or distributed to others in any way. Audio recordings should be destroyed when the semester ends. Photography in the classroom (taking pictures of lecture slides, or any other pictures) is strictly prohibited, as is video recording.

Academic Integrity

As a student, you are expected to behave honestly and ethically at all times.

According to McMaster University's Academic Integrity Policy, you are engaging in academic dishonesty if you "knowingly act or fail to act in a way that results or could result in unearned academic credit or advantage" (Academic Integrity Policy, p. 6).

This behaviour can result in serious consequences, such as a grade of zero on an assignment, loss of credit with a notation on the transcript that reads “Grade of F assigned for academic dishonesty," and/or suspension or expulsion from the university.

It is your responsibility to understand what constitutes academic dishonesty. The following are just three forms of academic dishonesty:

1. Plagiarism.
2. Improper collaboration.
3. Copying or using unauthorized aids in tests and examinations.

For more information on academic dishonesty and academic integrity, please read the Academic Integrity Policy: http://www.mcmaster.ca/academicintegrity

Online Privacy

This course makes use of an online discussion forum. You should be aware that your chosen forum username will be apparent to all other students in the course. The technology used in the course web pages is designed to keep private all additional information you provide (your name, email address, student number, and passwords). However, you should be aware that submission of such information on this or any other web site always carries some risk that the information will become public (for example, if the web site is "hacked"). Your continuation in this course will be interpreted to indicate that you consent to the disclosure of your chosen username in the online forum, and that you accept the risk of submitting the additional requested information. If you have any questions or concerns about the privacy of your information, please discuss these with the instructor.

Weekly Schedule

The table below shows weekly topics, assigned readings, and homework assignments (HW = homework assignment required of all students; GS HW = homework assignment required of grad students only). Students are encouraged to read the assigned material prior to each lecture. At least 12-hours prior to each lecture, the instructor will link each topic description to a PDF file containing the majority of the slides that will be shown in lecture (two PDF versions will be provided: a version with one slide per page, and a small slides version with four slides per page).

Week	Topic	Homework Assignment
	Basic Bayesian Inference, Parameter Estimation (Chs. 1 - 2)	
Jan 3 / 5	Slides: Large Small	--
Jan 10 / 12	Slides: Large Small	HW 1 (Due Jan 10)
Jan 17 / 19	Slides: Large Small	HW 2 (Due Jan 17)
Jan 24 / 26	Slides: Large Small	HW 3 (Due Jan 24)
Jan 31 / Feb 2	Review (Jan 31); Midterm Test 1 (Feb 2)	HW 4 (Due Jan 31)
	Curve Fitting, Robust Estimation, Adv Param Est (Chs. 3 - 6)	
Feb 7 / 9	Slides: Large Small	GS HW A (Due Feb 7)
Feb 14 / 16	Slides: Large Small	HW 5 (Due Feb 14)
Feb 28 / Mar 1	Slides: Large Small	HW 6 (Due Feb 28)
Mar 6 / 8	Review (March 6); Midterm Test 2 (March 8)	HW 7 (Due Mar 6)
	Model Comp & Avg, Bayesian Brain, Decision Theory (Chs. 7 - 10)	
Mar 13 / 15	Slides: Large Small	GS HW B (Due Mar 13)
Mar 20 / 22	Slides: Large Small	HW 8 (Due Mar 20)
Mar 27 / 29	Slides: Large Small	HW 9 (Due Mar 27)
Apr 3	Review	HW 10 (Due April 3) GS HW C (Due April

		3)
TBA	Final Exam (comprehensive, 3 hours)	

Note: The course schedule may be modified during the semester, at the discretion of the instructor. Modifications will be made directly to the table above and announced in class.

Help Outside the Classroom: Office Hours and Tutorial

Resource	**Day/Time**	**Place**
Tutorial	Monday, 9:30 - 10:20 AM	ABB 162
Dr. Goldreich office hour	Monday, 4:00 - 5:00 PM	PC 413
Luxi Li office hour	By appt. (lil54@mcmaster.ca)	By appt.
Ryan Peters office hour	By appt. (petersrm@mcmaster.ca)	By appt.

Note: Please do not email the assistants or instructor with specific Bayesian inference questions; we will not answer such questions by email. Instead, specfic content questions should be asked in office hours, in the tutorial sessions, or in class. In addition, students are strongly encouraged to post any such questions on the Bayes-Engage forum, where other students can help to answer them.
