PSYC 757: Introduction to Bayesian Statistics

Professor: Patrick E. McKnight, Ph.D.	Office Hours: Thurs 10:30-11:30am
Email: pmcknigh at gmu dot edu	Office Location: David King 2064/65
Class Location: ONLINE	Lecture: per schedule
Course Website	Fall 2019

Course Overview: Psychology (PSYC) 757 is a course devoted to introducing graduate psychology students to Bayesian statistics. The course was redesigned from previous incarnations to increase student involvement and reduce the workload to a manageable level. Also, in previous versions, I attempted to implement a standard lecture (2010), problem-focused (2013), and flipped-learning (2016) formats - each offered some advantages to students but none satisfied my goal (see below). The past two years (2018-2019), I implemented a group, project-based approach with a blend of flipped-learning, lecture, and problem-focused pieces; this year, I plan to do the same given the enthusiastic feedback from students. **My goal with this class is to provide every student the opportunity to implement basic Bayesian methods to standard social/behavioral science data analysis problems.**

Course Objectives: The purpose of the class remains focused on Bayesian statistics, however, unlike many Bayesian devotees, I consider Bayesian statistics a tool. Just like other tools, Bayesian methods have suitable and unsuitable applications. Students enrolled in my course may find my approach more challenging than standard graduate lectures because I emphasize mastery (i.e., learning) over achievement (i.e., grades). I demand more from you (and me) so that we all learn together. If this format is more difficult, why would I implement it? Simple. I want students to learn, not just get grades for taking a class. The more my students learn, the more future professional opportunities await them.

Reading Materials: I require no textbooks for this course for two reasons. Reason 1: the material becomes outdated quickly given the quick pace of Bayesian tool development. Reason 2: an abundance of online and free resources exist so that you can continue to learn with the most updated materials. I link below and on the course website all the relevant material you ought to read. Feel free to send me any other resources you find useful.

Grading Policy: I intend to assign one of three grades to you based upon your group's performance over 10 projects. All students will participate in a single group for the duration of the semester. I assign students to the groups. Research indicates that group work in graduate classes translates well to professional activities. Moreover, your group will be graded as a whole - just like your professional work will be viewed as an end product by a team. Do your best on each assignment and grades will not be an issue. I have 10 projects assigned throughout the semester (roughly 1 assignment per week). Your group's performance will be judged by both me and the rest of the class. I reserve the right to overrule the class if they collectively judge groups too stringently. The grades for all 10 projects will be communicated to students by demand. Your individual grade will be based upon the average of those projects throwing out the two worst grades. Thus, you will be graded on 80% (8 out of 10 projects) of your work. Each project is worth 10 points for a total of 80 points for the course.

Project Presentations: Every Tuesday, teams present their project results to the class in 3 minute sessions. Each group will present their results using the following format:

Slide Number	Content
1	The Problem
2	The Solution
3	The Results

4	Lessons Learned

Please do not deviate from the format provided; no more and no fewer slides permitted. Think like scientists when presenting the group project. You may select one presenter or divide the slide presentation among your group members but the transitions must be swift because you only have 3 minutes to present your project.

Graduate Grading Scale

Grade	Quality Points	Grade Points	Graduate Course Evaluation
A+	4.00	78+	Satisfactory/Passing + Letter
A	4.00	75	Satisfactory/Passing
A-	3.67	70	Satisfactory/Passing
B+	3.33	68	Satisfactory/Passing
B	3.00	65	Satisfactory/Passing
B-	2.67	60	Satisfactory*/Passing
C	2.00	55	Unsatisfactory/Passing
F	0.00	> 55	Unsatisfactory/Failing

* Although a B- is a satisfactory grade for a course, students must maintain a 3.00 average in their degree program and present a 3.00 GPA on the courses listed on the graduation application.

Projects: The following projects and the descriptions for each will be found on the linked documents. Check back throughout the semester for the most up-to-date project details - they may change as I gather more information on student progress.

Project Due Date	Project Purpose
9/4 (1)	Start using R to solve basic probability problems
9/11 (2)	Compute conditional probabilities in R
9/18 (3)	Use bayess package to solve one normal distribution problem
9/25 (4)	Apply Bayes Rule
10/2 (5)	Show me the money!: Plotting your results
10/16 (6)	Compare frequentists and subjectivists results from a single source
10/23 (7)	Use Stan to solve a simple Bayesian problem
10/30 (8)	Calculate and interpret correlations using Bayesian methods
11/6 (9)	Stan and R for general linear model estimation and inference
11/20 (10)	R, Stan, and Python applications for Bayesian inference

Rating forms: Rate Others then Rate your Own Team

Technology in Use: I use Google Sites, Groups, Forms, Docs, and Sheets for the course - as I do for all my other courses. You will not find any of the material on Blackboard because I have found that platform too unreliable over the years. Google is a toaster - it simply works. I highly recommend you use Google Hangouts for your team work outside regular class hours but use whatever tool works best for you and the rest of your team.

Students with Disabilities: If you are a student with a disability and you need academic accommodations, please see me and contact the Disability Resource Services (DRS) at 703-993-2474.

Honor Code: George Mason University has an Honor Code, which requires all members of this community to maintain the highest standards of academic honesty and integrity. Cheating, plagiarism, lying, and stealing are all prohibited. It is every student's responsibility to familiarize himself or herself with the Honor Code. The Honor Code is available at: http://oai.gmu.edu/the-mason-honor-code-2/ All violations of the Honor Code will be reported to the Honor Committee.

Official Communications via GMU E-mail: Mason uses electronic mail to provide official information to students. Examples include communications from course instructors, notices from the library, notices about academic standing, financial aid information, class materials, assignments, questions, and instructor feedback. Students are responsible for the content of university communication sent to their Mason e-mail account and are required to activate that account and check it regularly. NOTE: We will use Google Groups extensively throughout the course to communicate ideas, questions, and answers to one another.

Important Dates: **8/31** (Last day to add classes), **9/8** (Final drop deadline with no tuition penalty), **9/16-9/28** (Web withdrawal period with 100% liability), **10/12** (Fall break - Mon meets on Tues), **11/25-11/29** (Thanksgiving recess), **12/5** (Last day of classes)

Tentative Schedule (*dates below represent the week not just that Tuesday*)

Date	Readings	Objective	Optional but strongly recommended
8/25	Syllabus, Jeffrey (chapters 1-5), and R setup with Rstudio	Introduce course, probability theory, and R	Rstudio install, R tutorial #1, R tutorial #2, R tutorial #3, How to become a Bayesian, and for the really ambitious R for Data Science
9/1	Count Bayesie - Bayes' theorem with Lego	Bayes theorem, subjective probability, and updating priors	Kahn Academy Probability and Combinatorics, Bayes Theorem video, and this video
9/8	Bayesian Essentials with R (pages 1-61) requires GMU login	Normal distribution theory	Solutions to problems, bayess package
9/13	The Likelihood and Updating your priors	Simple(r) distributions and update functions	Plot different distributions
9/22	Graphics in R using ggplot2	Better displays	Hadley Wickham's Book
9/29	Flam (2014), Cartoon Explanation, Little (2005), and Meehl (1997)	Frequentists vs. Subjectivists	Select from the host of Goodman articles listed here
10/6	Spiegelhalter & Rice (2009)	Bayesian methods - from start to finish	Stan tutorial, Simple English
10/13*	Baath (2013) & Baath (2013a)	Correlations, predictions, and	Spurious correlations

		spurious relationships	
10/20	Bayesian Essentials (pages 65-137)	Linear models with Bayesian methods	Bayesian Linear Regression without tears in R
10/27	Gelman (2011) and Morey (2016)	Evidence and Bayes' Factor	Replicability-Index, Goodman (1999), and Johnson (2013), Feynman (1974)
11/3*	Bayesian Computation with Stan and Farmer Jöns	Stan	Choose from a host of tutorial papers
11/10	Introduction to Bayesian Inference (with Python)	Python applications	Think Bayes - long but worth reading throughout the semester
11/17	Peruse the Bayesian Inference Task View on the CRAN	Using R packages for Bayesian methods	Try BayesianTools, Using R for Bayesian Statistics
12/1	Replication of Psychological Science, Estimating the reproducibility of psychological science	Replication and Reproducibility of Science	Bayesian meta-analysis in R

* No class on Tuesday so this is a Thursday only week.