BMB 510. Data Analysis and Scientific Inference

An introductory course in the analysis of data and scientific inference for graduate students in Biochemistry, Molecular Biophysics, and related quantitative biomedical research areas. The course will stress fundamental principles of data analysis, best practice in presenting data, and how to draw sound scientific inferences from the data. The approach throughout is Bayesian. The overall goal is to provide students the tools to carry out rigorous and reproducible scientific research.

Section I

Presenting and summarizing data: Using Python/Jupyter Notebooks to do mathematical manipulations of data, calculate averages, standard deviations, etc, make histograms, box plots, scatter plots and line plots, perform numerical computations on data arrays.

Section II

1) Role of randomness, common pitfalls and errors in data analysis.
2) Review of probability theory and the tools used for manipulating probabilities. Introduction to key concepts of probability density functions, cumulative probability distributions.
3) Principles of parameter estimation. Emphasis will be on robust approaches to obtaining credible intervals for parameter estimates that are valid even with small amounts of data and/or non-normal distributions, and ways to correctly incorporate results from previous experiments and other prior information.
3) Examples of Parameter Estimation will include: fraction/proportion parameters, population sizes, rate/time constant/decay length parameters, counting data with and without background, differences in parameters between two sets of measurements, linear regression.

Section III

Higher-level aspects of analysis of data , including experimental design, quantitative comparison of models, mixture models and clustering.

Intended Students

First year graduate students in BMB and other BGS graduate groups with suitable background in the mathematical and physical sciences. The course is intended as an alternative to BIOMED 611 for these more quantitative students in order to fulfill the BGS biostatistics requirement. The course will be offered spring semester, yearly.

Requirements

Required for BMB students. All others by permission of the instructor. Students are required to bring a laptop to each class. The miniconda programming environment https://conda.io/miniconda.html equipped with the matplotlib, numpy, jupyter notebook, pandas and xlrd modules will be used throughout the class. A Python code development environment module (such as Idle or Spyder) will also be useful for novice programmers). Help with software installation will be provided. Experience with the Python programming language is not required, but as students acquire it, it will help the students connect the lecture material to the programs they will run. Methods will be taught by example, and students will run the examples themselves either on data provided by the instructor, or on suitable data from their own work.

Textbook

Sharp, Kim A. (2018) *Being Less Wrong: A Bayesian approach to data analysis and Scientific Inference* (hard copies will be provided)

Additional Texts

Iversen, Gudmund R. 1984. *Bayesian Statistical Inference*. Sage Publications.
Sivia D, Skilling J (2006) *Data Analysis, a Bayesian Tutorial* (Oxford University Press, Oxford).

Evaluation

Grades will be based on homework (40%), final exam (40%), participation in class discussions (20%)
Homework: Students will be required to run data analysis examples either in class or as homework, and email their results to the TA to be graded. Final exam format: Each Q will involve analysis of data followed by a short text answer with interpretation/discussion of the results.

Expectations upon successful completion of the course

The students will understand the different kinds of probability: joint, conditional, marginal, and how they are used to analyze data; know which kind of analysis to apply depending on the type of data and what question is being asked; know how to obtain the usual statistical quantities – mean, variance, differences in mean and variance, etc., especially the importance of having credible intervals on every quantity they obtain and how to interpret the results of their analysis; recognize the confounding effects of random variation, noise, small sample size, and non-normal distributions; understand the principles of i) experimental design in the context of structural, physical, mechanistic experiments that form the core of modern biophysics and biochemistry research. ii) the quantitative comparison of models or hypotheses.

Instructors

Director: Kim Sharp, Ph.D, sharpk@pennmedicine.upenn.edu   TA: Neha Srikumar, Neha.Srikumar@pennmedicine.upenn.edu

Schedule

11.00-12.15 Tues, Friday. Classes will be taught via Zoom

Date	Topic	pages in BLW text[1]	
F 22 Jan	Python Notebook I: Doing math with Python, coding up statistical equations		
T 26	Python Notebook II: Data Structures		
F 29	Python Notebook III: Data Structures/Control Flow		
T 2 Feb	Python Notebook IV: Data Input and Plotting with MatPlotLib		
F 5	Python Notebook V: Building more complex programs: Defs and Modules		
T 9	Python Notebook VI: Tables and Data Manipulation using Pandas		
F 12	No Class- Upenn Engagement day		
T 16	Python Notebook VII: Numerical Computing using Numpy		
F 19	Data Presentation: central tendency, spread, Averaging, Simpson P'dox, Plotting	pp29-34	
T 23	Data Presentation: linear regression, correlation coefficients. Gary Smith, what the luck. Hotelling Review. Contingency tables, Ioannides paper	pp29-34 pp59-62 p25	
F 26	No Class B&B and BMB mini-retreat		
T 2 Mar	Probability Basics. Read Howard Wainer	pp15-21	
F 5	Probability Basics, Bayes Rule	pp15-21	
T 9	Parameter Estimation: Population ID	pp35-37	
F 12	Parameter Estimation: Proportion/Fraction	pp38-40	
T 16	Parameter Estimation: Rates	pp42-46	
F 19	Parameter Estimation: Rate+Bkgnd, Decay length, windowing	pp42-46	
T 23	Multi-Parameter Estimation: Difference in Proportion/fraction parameter. Thompson sampling	pp51-52	

F 26	Multi-Parameter Estimation: Mean and Variance	pp53-55	
T 30	<mark>No Class- Upenn Engagement day</mark>		
F 2 Apr	Multi-Parameter Estimation: Difference in means, variance, multi- comparisons	pp55-59	
T 6	Multi-Parameter Estimation: Difference in Fractional/Proportion parameter		
F 9	Using Prior information, Estimating differences in mean, proportion using statistical summmaries	pp47-50	
T 13	Decay rate revisited: Survival analysis	pp62-65	
F 16	Discrete/non-parametric data. Estimating Population Size, Rank Tests	pp40-42	
T 20	Clustering, Mixture models, hierarchical models.	pp67-75	
F 23	Review		
T 27	Final Exam		
F 30			

[1]"Being Less Wrong", Sharp, 2018