STAT371: Bayesian Data Analysis

General Course Information

This document provides a general description of STAT371, and should be treated as a broad overview and subject to minor change from one year to another. A more detailed Course Information document is provided each semester.

Lectures

There are 30 lectures per semester. These are usually delivered three lectures per week for 5 weeks and then 2 lectures per week thereafter, but subject to change as needed.

Tutorials/Practicals

There is a tutorial or practical each week. These alternate between practicals designed to develop skills and tutorials to go over questions realted to the assignments.

Website

Copies of the lecture slides and assignments will be available via a course website, along with other resources such as announcements and recorded lectures. Please note that we will work on the whiteboard and have in-class discussions, so recorded lectures are not a substitute for in-person attendance.

Prerequisites

STAT260 and STAT270 are required to provide the basics of computing, probability, and inference. Mathematical skills required for STAT270 (e.g. differentiation, integration, matrix algebra) are also assumed. Familiarity with aspects of theoretical statistics (e.g. likelihoods, probability density functions, cumulative density functions) that are taught in STAT270 are crucial. Prior programming experience in R allows us to focus on more advanced modelling and accessing new programming tools (JAGS and Stan).

Content Overview

We introduce the theory and application of Bayesian inference, a formally niche topic that has been transformed into a powerful modelling tool due to modern computational methods. As an introduction, we will initially look at using Bayesian methods to solve the types of problems encountered in STAT110 and STAT210. Then, we will then extend the methods to more complex data structures and models. Models will be fitted using modern Bayesian software including JAGS, Stan, and R. This course is aimed at motivated and engaged students with a good understanding of intermediate probability and inference (STAT 270) with an interest in applied statistical modelling and computing.

Learning Outcomes

At the end of STAT371 a student should be able to independently use R with JAGS or Stan to complete Bayesian statistical analyses. The student will be familiar with Bayesian programs and will be able to input, manipulate, plot and analyse a variety of different data types. Finally, they should be able to communicate their results to others and understand the ethical and scientific importance of reproducible research.

The only way to learn data analysis is by doing it. The lectures are designed to illustrate relevant concepts and explain common difficulties. The assignment exercises are designed to give the students hands-on experience and understanding of the material and computational methods. Successful students in this class are those who spend time on assignment questions and revising the lectures in R, JAGS, and Stan (not just reading the lectures but implementing the concepts discussed). It is important to practise key programming concepts from the start of lectures. To be successful, students should expect to work hard and remain engaged throughout the semester.

Modules

The topics covered in STAT371 include:

- Theory of Bayesian inference
- Markov chain Monte Carlo methods
- Modelling using JAGS and Stan
- Advanced scientific modelling

Outline (subject to minor changes)

Week	Topics
1	Overview, probability and inference review
	Introduction to Bayesian inference
2	Simple models
	R implementation with jags
3	Linear, GLM models using Bayes
	brms package using Stan
4	Markov chain Monte Carlo
	Priors
5	More MCMC
6	Scientific models I
	Model fit
7	Scientific models I
	Overdispersion, random effects
8	Scientific models II
	Time to failure data, censoring, truncation
9	Scientific models II
	Data splitting, data augmentation
10	Scientific models III
	Calibration models
11	Scientific models III
	Model selection, nonlinear models
12	Scientific models III
	Model selection, WAIC, loo package
13	Course Review

Grading

The course has three component: exam (E), assignments (A) and mid-term test (T). The exam is worth 60% of your final grade, the assignments are worth 25% and the test is worth 15%.

Terms requirements:

- You must achieve a minimum score of 40% on the assignments in order to pass the paper.
- You must achieve a minimum score of 40% on the exam in order to pass the paper.

Thus, for a student who meets terms, the final mark (F) is: $F = (60E + 25A + 15T)/100$.

Assignments

There are normally five assignments throughout the semester, approximately fortnightly. Assignments will be graded for technical correctness, presentation, and communication. To get full marks on your assignment, you may be required to explain your working/reasoning to your instructor.

Mid-term Test

The mid-term test will examine core knowledge from the first 5 weeks and occurs during a normal lecture period.

Final Exam

The final exam will be 3 hours in duration. A practice exam is made available close to the exam date.

Late Work

Unless you make arrangements with the lecturer beforehand, all late work turned in after the due date will receive a 10% penalty per day and will not be accepted after 2-3 days.