STAT 520C – TOPICS IN BAYESIAN ANALYSIS AND DECISION THEORY

"Principles and Operating Characteristics of Bayesian Procedures,"

2022/23 – Second half of Term 2 (1.5 credits)
_____

Instructor: Dr. Paul Gustafson (e-mail: gustaf@stat.ubc.ca, office: ESB 3192)
Lectures: Tues, Thurs, 12:30 - 14:00, ESB 4192
                (first lecture: March 8th, last lecture: April 13th)

Prerequisite: Open to any interested graduate students in the Department of Statistics. Graduate students from other departments are welcome, provided they have sufficient statistical and mathematical backgrounds (roughly, statistical theory to at least the level of UBC STAT 305, more ideally to the level of STAT 460). Students from other units should consult the instructor about suitability.

Text: There is no textbook. Readings may be assigned, ideally from texts available via the UBC library e-book collection.

Course description: This half-term course is about the *operating characteristics* of Bayesian statistical procedures. Nowadays, there are lots of software packages to implement Bayesian analyses: the user inputs the statistical model, the prior distribution, and the data, then the software outputs the posterior distribution. We are going to leave this as a black-box and not worry about the algorithms being used "under the hood." Rather, we are going to discuss how well Bayesian inference works in various settings.

So to head off confusion, this is not a course on applied Bayesian data analysis. Nor is it a course on computing algorithms for Bayesian inference. It's about the *conceptual stuff* in between, trying to quantify and characterize what Bayesian inference can (and can't) achieve.

As a further distinction, the course will not be particularly theoretical. We will use enough math (and simulation) as needed to make conceptual points, but it's *not* a course about theorems and their proofs.

Examples (not an exhaustive list) of questions we will explore include:

- In what senses are Bayesian estimation procedures optimal?

- In what senses are Bayesian interval estimates calibrated?

- How can Bayesian procedures synthesize multiple sources of evidence?

- What are the principles and operating characteristics of Bayesian model selection procedures?

- How are Bayesian model selection and Bayesian prediction intertwined?

- What are some of the fundamental ideas around "borrowing of strength" in a Bayesian hierarchical model?

Lecture format: Mostly traditional lecturing, perhaps some small-group, in-class activities. I will post "pre" versions of the lecture slides in advance of each lecture. But these will have gaps that we will fill in during the class.

Evaluation: based on class participation (10%), short reflective writing exercises (20%), and, in lieu of a final exam, a final project (70%, due at the end of the April exam period). I will post a separate document concerning the reflective writing exercises and the final project.