Bayesian Statistics

Augusto Fasano

Course Description

The course provides an introduction to the Bayesian approach to statistical inference. Unlike with the classical frequentist approach to statistics, a Bayesian statistical model can incorporate prior beliefs about the object of inference that are formally combined with the information provided by the collected data. These in turn produce an updated, or posterior, model, which formalises the process of learning from experience in a mathematically coherent way.

The course will start from foundational and methodological aspects of Bayesian statistics and will move to more sophisticated constructions and related computational considerations in the second part. By the end of the course, students will be able to formulate and implement the most used Bayesian models and develop their own constructions. Computations will be carried out by making use of the statistical software `R`.

Prerequisites

Basic Probability theory: discrete and continuous random variables. Frequentist parametric Statistics: point estimation, hypothesis testing, interval estimation and linear regression. Familiarity with the `R` software (www.r-project.org) is not required, but some basic programming skills with some programming language can be useful.

Outline of the Course

1. The Bayesian approach to statistics.
   (a) The likelihood principle, the prior and posterior distributions.
   (b) The Bayes Theorem for dominated models.
   (c) Interpretation of the scientific inference via the Bayesian approach.
2. The three fundamental problems of inference.
   (a) Point estimation, hypothesis testing and interval estimation.
   (b) Comparison between the Bayesian and frequentist methods.
3. Prior distributions.
   (a) Noninformative prior distributions. Conjugate and mixture distributions.
   (b) Examples of univariate conjugate models: Bernoulli beta model, Poisson gamma model, normal normal-inverse-gamma model.
4. Computational methods for non-conjugate Bayesian models.
   (a) General outline of Markov chain Monte Carlo methods.
   (b) Gibbs sampling and the Metropolis-Hastings algorithm.
   (c) Checking convergence of a Markov Chain.
   (d) (optional) Hamiltonian Monte Carlo with `stan`.
   (e) Approximate methods: Variational Bayes and Expectation Propagation.
5. Introduction to more advanced models.
   (a) Multivariate linear models and generalized linear models in the Bayesian framework.
   (b) A Bayesian view of some discrete-choice models from econometrics.
   (c) Dynamic linear models for time-series analysis.

Textbooks

Main references:

- Hoff P. (2009), A First Course in Bayesian Statistical Methods, Springer Texts in Statistics.
- Gelman, A., Carlin, J.B., Stern H.S., Dunson, D.B., Vehtari, A., Rubin, D.B. (2013), Bayesian Data Analysis, Third Edition.

Evaluation

60% assignment + 40% written exam.