

STATS 730: Statistical inference

Course outline 2014

Lecturer:

Russell Millar, Room 229

Assessment:

60% final exam + 20% midterm + 20% assignments, or 100% final exam, whichever is to your advantage.

Course overview:

STATS 730 gives you general-purpose skills to model real data, using likelihood-based statistical inference under the frequentist paradigm. It begins with a gentle introduction to maximum likelihood and the notation that will be used throughout, followed by simple and not-so-simple (e.g., finite mixture model) iid examples. The essential properties and tools of maximum-likelihood inference are then presented. Maximum likelihood is then applied in a wide variety of settings with examples in both R and SAS, and ADMB where needed. (Students may choose any of these languages for their homeworks.) The course concludes by looking at extensions of maximum likelihood for models for more challenging situations, including quasi-likelihood, conditional likelihood, and mixture models.

STATS 730 provides the tools and skills used by many other graduate courses on offer in this department, and of invaluable use to students undertaking MSc projects or beginning PhD study. It gives strong exposure to statistical programming in both R and SAS, and brief exposure to ADMB - the most powerful optimization software available today.

Course content:

The content of **STATS 730** will closely follow Chapters 1–4, 7 and 10 of Millar (2011), and pick bits from Chapters 5 and 6 and other material as required.

Introduction to likelihood

Essential concepts and iid examples

Hypothesis tests and confidence intervals/regions

Essential tools and tricks

Maximizing the likelihood

Applications

Generalized linear models and extensions

Mixed effects models

Reference

Millar, R. B. 2011. Maximum likelihood estimation and inference: With examples in R, SAS and ADMB. Wiley.