Bayesian Learning 2025

Professor: Hedibert Freitas Lopes
Teaching assistant: Guilherme Piantino

Syllabus: The ultimate goal of this course is to enable graduates to critically decide between the classical or Bayesian approach, or a combination of both, when faced with real-world decision-making problems under uncertainty. Areas where these real-world problems arise, as examples discussed throughout the course, include microeconomics, macroeconomics, finance, quantitative marketing, among many others. With this objective in mind, we will study the basic ingredients of the Bayesian paradigm: formulation of the binomial model-prior, model comparison and combination, computational aspects, and Bayesian decision-making. In the second part of the course, the Bayesian approach to traditional linear regression and logistic regression models will be introduced, as well as their modern versions where priors are treated as regularization mechanisms and sparsity inducers. Sparsity will be present throughout the 2nd and 3rd parts of the course when dealing with highly dimensional and/or highly complex models. In the third and final part of the course, we will present several statistical models currently used for this purpose, such as mixture models, hierarchical models, factor models, and regression tree models, as well as models based on neural networks and models that use texts and documents as data (text modeling). All, it is worth mentioning, under the unified and coherent Bayesian approach. All calculations during the course will be performed using the R statistical package.

Homework assignments

HW1 + HW2 + HW3: Choose one of the three STAN examples we worked in class (AR(1), GARCH(1,1) or SV-AR(1)) and run assuming Gaussian errors. Then, compare the results of both models, i.e. the model with Student's t errors and the model with Gaussian errors. The idea of this exercise is for you to familiarize yourself with STAN/RSTAN for Bayesian modeling. **HW4:** Replicate the PCA/FA exercise we worked out in the class of June 2nd 2025. **HW5:** In 2021, I worked on an example of predictive modeling where the response is binary (Bernoulli). See if you can replicate the exercise and understand the comparison between the various models: GLM (logistic regression), Bayesian GLM, Ridge/Lasso/Elastic net logistic regressions, CART, Random forest and BART. Here is some of the output for the Intensive Care Unit (ICU) data (from R package "aplore") and here is my R code

Final project – paper presentation:
Lista de artigos científicos
Papers assigned to each student

Additional examples developed/discussed in class

Mixture of Gaussian prior with unknown weights
Weight vs height: outlier, Gaussian vs Student's t errors
Learning the degree of freedom
Binomial vs Negative Binomial
Nonlinear regression: SIR, MCMC and MC methods
HMC: A toy example introduction
Stan/rstan for HMC-based posterior inference
HMC via ChatGPT (supervised by me!)
HMC & STAN – Nonlinear regression illustration (R script + STAN script)
HMC & STAN – AR(1) model with Student's t errors
HMC & STAN – GARCH(1,1) model with Student's t errors
HMC & STAN – SV-AR(1) model with Gaussian or Student's t errors
- HMC for (Physics) Dummies (Vats, 2023)
- MCMC using Hamiltonian dynamics (Neal, 2012)
- The Metropolis-Hastings-Green algorithm (Geyer, 2003)
- Introduction to statistical computing and HMC (Michael Betancourt, 2018)
- Geometric foundations of HMC (Michael Betancourt, 2014)
- Riemann Manifold Langevin and HMC Methods (Girolami & Calderhead, 2011)
- HMC for Hierarchical Models (Betancourt & Girolami, 2013)
- The No-U-Turn Sampler: Adaptively Setting Path Lengths in HMC (Hoffman and Gelman, 2014)

AR(p) models: Bayesian updating with conjugate prior vs Gibbs sampler
Return on education: Gaussian linear model with conjugate prior
Gaussian linear regression model: Gibbs sampler for conditionally conjugate prior
Subset selection via BIC and shrinkage/sparsity via regularizing priors – Stock and Watson's (2002) macro data
Logistic regression: GLM vs Bayesian GLM (R packages: UPG & stan_glm) + Graphical summaries of the results
Two more dataset for logistic regression
Default of credit card clients (dataset)
Example of principal component analysis and factor analysis
Inflation forecasting (Graphical output) + (dataset)
Example of latent Dirichlet allocation in R
Example of neural net in R (chatGPT assisted)

Course notes (+ R code & references)

Bayesian ingredients

- The Monty Hall problem
- Flipping one of three coins three times and observing three heads (R code)
- Chapter 1 of "Bayes' Rule: A tutorial Introduction to Bayesian Analysis (Slide)
- Tiago Mendonca's shiny for the physicists example
- Phisycists A, B, C and D: Normal model and 4 priors
- Histórias da Matemática: Da Contagem nos Dedos à Inteligência Artificial (by Marcelo Viana, IMPA)
    - Provavelmente não somos bons em probabilidade
    - Paradoxos de probabilidade
    - O segredo para ganhar o jogo
    - O problema dos testes falsos positivos
    - O paradoxo de Simpson
    - Esperança para a estrela solitária

Bayesian computation

- Monte Carlo Integration – Excerpts from chapter 3 of Gamerman and Lopes (2006)
- Banana-shape posterior: Posterior inference via SIR
- Learning the number of degree of freedom of a Student's t (Rmarkdown)
- Bayesian regression with the normal-gamma (NG) prior
- A bit of Monte Carlo simulation and integration
- Back to the physicists example
- Normal vs skew-normal distributions: an exercise in Bayesian learning
- The potential fallacy of using the prior as proposal in SIR algorithms
- Bayesian learning of correlation in bivariate normal data via SIR algorithm
- Smith and Gelfand (1992) Bayesian Statistics without Tears. TAS, 46(2), 84-88.
- Hamiltonian Monte Carlo: a brief introduction

Bayesian linear regression

- Boston housing data: ML and Bayesian inference
- Motorcycle example
- Regressing weight on height: SIR for Gaussian and Student's t models
- Bayesian regularization
- iid Bernoulli or logit regression or probit regression? (Rmarkdown)
- Bayesian Poisson regression model versus i.i.d. Poisson model
- Another example of Poisson regression
- A few R packages for Bayesian inference in linear models
- Stan/rstan: Bayes Sparse Regression

Bayesian classification via logistic regression

- Sparse logistic regression for the spam/ham dataset (data)
- Marketing campaigns of a Portuguese banking institution (data)
- Sparse logistic regression: comparison of regularization and Bayesian implementations
- Gelman et al. (2008) A weakly informative default prior distribution for logistic and other regression models, *AOAS*, 2(4), 1360-1383.
- Polson et al. (2013) Bayesian Inference for logistic models using Pólya-Gamma latent variables, *JASA*, 108, 1339-1349.

Other important modeling structures

Factor models (Additional material)

- Indice de Desenvolvimento Humano
- US market (data)
- Bayesian factor models
- Reducao de ruido com PCA e FA

Time-varying variance/covariance

- Stochastic volatility
- Factor Stochastic Volatility(FSV)
- FSV x Dynamic Conditional Correlation (DCC)

Finite mixture of distributions

Machine Learning 1: Tree models

- Classification and regression trees (CART)
    - CART step-by-step for the Boston housing data
    - CART for the motorcycle data
    - ROC and AUC
- Bayesian CART
- Bootstrap aggregating (bagging)
    - Ilustracao do bootstrap
- Boosting (weak/stronger learners)
- Random forests
    - Predicao da covid-19 com dados publicos do Hospital Israelita Albert Einstein/SP
- Bayesian additive regression trees (BART) + (Rob's notes) + (a list of slides/examples/people/papers)

Machine Learning 2: Modeling text

- Latent Dirichlet Allocation (LDA)
- Short list of slides and papers
- Example 1 & 2: Twitter & BBC
- Example 3: Dickens, Wells, Verne or Austen?

Machine Learning 3: Neural nets

- Neural Networks
- short list of slides and papers

Additional supporting material

- Stan/rstan for posterior inference: Hamiltonian MC (HMC) methods – by Hedibert Lopes (February 2021)
- MC and MCMC: Key References – by Hedibert Lopes (February 2021)
- R packages for Bayesian linear regression – by Hedibert Lopes (February 2020)
- R packages for Bayesian Econometrics – by Hedibert Lopes (March 2014)
- CRAN Task View on Bayesian Inference (July 2023)
- Mathematics for Machine Learning – by Deisenroth, Faisal and Ong (2020)
- Conceitos e analises estatisticas com R e JASP – by Luis Anunciação (September 2021)
- Data Science, Marketing and Business by Pedro Fernandes & Paulo Marques (October 2019)
- Aprendizado de Máquina: Uma Abordagem Estatística (by Rafael Izbicki & Tiago Mendonça)
- Estatística e Ciência de Dados (by Pedro Morettin & Julio Singer)