21D008

Probabilistic Inference in Machine Learning

Term 2

3 ECTS

ELECTIVE

Pre-requisites

Familiarity with Statistical Inference (as given by the Foundations in Data Science course or similar), Bayesian Inference (as given by the Statistical Inference course or similar), R programming.

Overview and Objectives

The combination of Bayesian Statistics and Machine Learning provides a rich set of data analysis tools to tackle challenging problems. It provides a Probabilistic Machine Learning framework that moves the focus from purely making predictions to also portraying the uncertainty associated with those predictions, and on the elements that build the actual predictor itself. In this manner, one gains insights into the studied phenomenon, how predictions are obtained, and how reliable they are. These two features, interpretability and the ability to quantify uncertainty, are essential to permit the use of machine learning in sensitive applications, such as healthcare and public policy. Enabled by computational advances, the framework allows considering highly flexible models, possibly with infinitely many parameters. Therefore, one can describe phenomena in situations that are too complex to capture with standard parametric models.

The course covers state-of-the-art methods related to non-parametric regression, latent variables, non-parametric hierarchical models, graphical models and their regression on further predictors, and generalized Bayes notions related to model-free inference. It also discusses the main computational advances that enabled their use, which provide a valuable toolset beyond the models considered in this course.

Course Outline

Advanced Bayesian computation: Markov Chain Monte Carlo, Sequential Monte Carlo, Variational Bayes, Expectation Propagation. Applications such as latent variable models for text data analysis

Advanced graphical models: Bayesian structural learning for graphical models, regression of graphical structure on covariates

Non-parametric Bayesian inference: advanced Gaussian processes, Dirichlet process mixtures, Bayesian Deep Learning

Further topics: Bayesian optimization, generalized Bayesian Inference and PAC-Bayes

21D008

Probabilistic Inference in Machine Learning

Term 2

3 ECTS

ELECTIVE

Learning Outcomes

- Build and fit probabilistic and advanced Bayesian Machine learning models
- Learn to reason about attaching uncertainty and model interpretation to machine learning forecasts
- Learn computational algorithms and strategies to fit complex probability models
- Learn and apply supervised and semi-supervised algorithms to predict random phenomena

Evaluation

Two homework assignments (30% each) and final project (40%)

Materials

Bishop is a great general reference on Bayesian Machine Learning. Barber is more oriented towards computer scientists and has lots of material on graphical models. Gelman et al is the main primer on Bayesian inference, and describes both parametric and some of the non-parametric models seen in this course. Rasmussen & Williams is an excellent reference for Gaussian processes.

C. M. Bishop. Pattern Recognition and Machine Learning.

D. Barber. Bayesian Reasoning and Machine Learning. Freely available online.

A Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, D.B. Rubin. (2013). Bayesian data analysis. CRC press.

C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. Freely available online.