

Bayesian Machine Learning

Giorgio Maria Di Nunzio¹

¹ Department of Information Engineering
Email: dinunzio@dei.unipd.it

Timetable: 20 hrs. Lectures on Wednesday and Friday, 10:30-12:30, starting from March 9th, 2022 (see on <https://phd.dei.unipd.it/course-catalogues/>)

Course requirements: Basics of Probability Theory. Basics of R Programming.

Examination and grading: Homework assignments and final project.

SSD: Information Engineering

Aim: The course will introduce fundamental topics in Bayesian reasoning and how they apply to machine learning problems. In this course, we will present pros and cons of Bayesian approaches and we will develop a graphical tool to analyse the assumptions of these approaches in classical machine learning problems such as classification and regression.

Course contents:

Introduction of classical machine learning problems.

1. Mathematical framework
2. Supervised and unsupervised learning

Bayesian decision theory

1. Two-category classification
2. Minimum-error-rate classification
3. Bayes decision theory
4. Decision surfaces

Estimation

1. Maximum Likelihood Estimation
2. Expectation Maximization
3. Maximum A Posteriori
4. Bayesian approach

Graphical models

1. Bayesian networks
2. Two-dimensional visualization

Evaluation

1. Measures of accuracy

References:

1. J. Kruschke, Doing Bayesian Data Analysis: A Tutorial Introduction With R and Bugs, Academic Press 2010
2. Christopher M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics), Springer 2007
3. Richard O. Duda, Peter E. Hart, David G. Stork, Pattern Classification (2nd Edition), Wiley-Interscience, 2000
4. Yaser S. Abu-Mostafa, Malik Magdon-Ismail, Hsuan-Tien Lin, Learning from Data, AML-Book, 2012 (supporting material available at <http://amlbook.com/support.html>)
5. David J. C. MacKay, Information Theory, Inference and Learning Algorithms, Cambridge University Press, 2003 (freely available and supporting material at <http://www.inference.phy.cam.ac.uk/mackay/>)
6. David Barber, Bayesian Reasoning and Machine Learning, Cambridge University Press, 2012 (freely available at <http://web4.cs.ucl.ac.uk/staff/D.Barber/pmwiki/pmwiki.php?n=>)
7. Kevin P. Murphy, Machine Learning: A Probabilistic Perspective, MIT Press, 2012 (supporting material <http://www.cs.ubc.ca/~murphyk/MLbook/>)
8. Richard McElreath, Statistical Rethinking, CRC Press, 2015 (supporting material <https://xcelab.net/rm/statistical-rethinking/>)