BOOK REVIEW

Simon Rogers and Mark Girolami: A first course in machine learning

CRC Press, Boca Raton, 2012, xx + 285 pp., US $69.95, GB £36.99, € 51.80, ISBN 978-143982414-6

Thoralf Mildenberger

This book offers an introduction to machine learning for students with rather limited background in mathematics and statistics. The seven chapters are titled "Linear Modelling: A Least Squares Approach", "Linear Modelling: A Maximum Likelihood Approach", "The Bayesian Approach to Machine Learning", "Bayesian Inference", "Classification", "Clustering", and "Principal Components Analysis and Latent Variable Models". This choice of topics is somewhat unbalanced, as nearly all techniques and concepts have their origin not in computer science but in statistics, the only exception being Support Vector Machines and kernel methods. Among many other methods, neural networks, trees and boosting are missing.

The book is well written and focusses on explaining the main concepts at a very basic level, keeping in mind the limited mathematical background of the intended audience. There are also useful references for further reading at the end of each chapter, and MATLAB code implementing the methods is available online along with the data sets. The code also seems to work well with free alternatives to MATLAB like Octave and FreeMat.

Sadly, the limited focus of the book makes it unsuitable for a machine learning course for students that already have some background in statistics. On the other hand, for an introductory course on (multivariate) statistics, there is too much "classical" material which is missing. Although the book contains some very nice material, it is therefore not the best choice as a textbook for a machine learning course for students in mathematics, statistics or computer science.

T. Mildenberger (✉)
IDP Institute of Data Analysis and Process Design, Zurich University of Applied Sciences, Winterthur, Switzerland
e-mail: thoralf.mildenberger@zhaw.ch