Machine Learning Methods in Environmental Sciences

Neural Networks and Kernels

William W. Hsieh
University of British Columbia
Vancouver, B.C., Canada

Cambridge University Press

August 31, 2008

Preface

Machine learning is a major subfield in computational intelligence (also called artificial intelligence). Its main objective is to use computational methods to extract information from data. Machine learning has a wide spectrum of applications including handwriting and speech recognition, object recognition in computer vision, robotics and computer games, natural language processing, brain-machine interfaces, medical diagnosis, DNA classification, search engines, spam and fraud detection, and stock market analysis. Neural network methods, generally regarded as forming the first wave of breakthrough in machine learning, became popular in the late 1980s, while kernel methods arrived in a second wave in the second half of the 1990s.

In the 1990s, machine learning methods began to infiltrate the environmental sciences. Today, they are no longer an exotic fringe species, since their presence is ubiquitous in the environmental sciences, as illustrated by the lengthy bibliography of this book. They are heavily used in satellite data processing, in general circulation models (GCM) for emulating physics, in the post-processing of GCM model outputs, in weather and climate prediction, air quality forecasting, analysis and modelling of environmental data, oceanographic and hydrological forecasting, ecological modelling, and in the monitoring of snow, ice and forests, etc.

This book presents machine learning methods (mainly neural network and kernel methods) and their applications in the environmental sciences, written at a level suitable for beginning graduate students and advanced undergraduates. It is also aimed at researchers and practitioners in environmental sciences, who having been intrigued by exotic terms like neural networks, support vector machines, self-organizing maps, evolutionary computation, etc., are motivated to learn more about these new methods and to use them in his/her own work. The reader is assumed to know multivariate calculus, linear algebra and basic probability.

Chapters 1–3, intended mainly as background material for students, cover the standard statistical methods used in environmental sciences. The machine learning methods of later chapters provide powerful nonlinear generalizations for many of these standard linear statistical methods. The reader already familiar with the background material of Chapters 1–3 can start directly with Chapter 4, which introduces neural network methods. While Chapter 5 is a relatively technical chapter on nonlinear optimization algorithms, Chapter 6 on learning and generalization is essential to the proper use of machine learning methods — in particular, Section 6.10 explains why a nonlinear machine learning method often outperforms a linear method in weather applications but fails to do so in climate applications. Kernel methods are introduced in Chapter 7. Chapter 8 covers nonlinear classification, Chapter 9, nonlinear regression, Chapter 10, nonlinear principal component analysis, and Chapter 11, nonlinear canonical correlation analysis. Chapter 12 broadly surveys the applications of machine learning methods in the environmental sciences (remote sensing, atmospheric

science, oceanography, hydrology, ecology, etc.). For exercises, the student could test the methods on data from their own area or from some of the web sites listed in Appendix A. Codes for many machine learning methods are also available from sites listed in Appendix A.

On a personal note, writing this book has been both exhilarating and gruelling. When I first became intrigued by neural networks through discussions with Dr. Benyang Tang in 1992, I recognized that the new machine learning methods would have major impact on the environmental sciences. However, I also realized that I had a steep learning curve ahead of me, as my background training was in physics, mathematics and environmental sciences, but not in statistics nor computer science. By the late 1990s I became convinced that the best way for me to learn more about machine learning was to write a book. What I thought would take a couple of years turned into a marathon of over eight years, as I desperately tried to keep pace with a rapidly expanding research field. I managed to limp pass the finish line in pain, as repetitive strain injury from overusage of keyboard and mouse struck in the final months of intensive writing!

I have been fortunate in having supervised numerous talented graduate students, post-doctoral fellows and research associates, many of whom taught me far more than I taught them. I received helpful editorial assistance from the staff at the Cambridge University Press and from Max Ng. I am grateful for the support from my two university departments (Earth and Ocean Sciences, and Physics and Astronomy), the Peter Wall Institute of Advanced Studies, the Natural Sciences and Engineering Research Council of Canada and the Canadian Foundation for Climate and Atmospheric Sciences.

Without the loving support from my family (my wife Jean and my daughters, Teresa and Serena), and the strong educational roots planted decades ago by my parents and my teachers, I could not have written this book.

Notation used in this book

In general, vectors are denoted by lower case bold letters (e.g. $\mathbf{v}$), matrices by upper case bold letters (e.g. $\mathbf{A}$) and scalar variables by italics (e.g. x or J). A column vector is denoted by $\mathbf{v}$, while its transpose $\mathbf{v}^{\mathrm{T}}$ is a row vector, i.e. $\mathbf{v}^{\mathrm{T}} = (v_1, v_2, \ldots, v_n)$ and $\mathbf{v} = (v_1, v_2, \ldots, v_n)^{\mathrm{T}}$, and the inner or dot product of two vectors $\mathbf{a} \cdot \mathbf{b} = \mathbf{a}^{\mathrm{T}}\mathbf{b} = \mathbf{b}^{\mathrm{T}}\mathbf{a}$. The elements of a matrix $\mathbf{A}$ is written as A_{ij} or $(\mathbf{A})_{ij}$. The probability for discrete variables is denoted by the upper case P, whereas the probability density for continuous variables is denoted by the lower case p. The expectation is denoted by $\mathrm{E}[\ldots]$ or $\langle \ldots \rangle$. The natural logarithm is denoted by ln or log.

Acronyms:

AO = Arctic Oscillation
BNN = Bayesian neural network

CART = classification and regression tree
CCA = canonical correlation analysis
CDN = conditional density network
EC = evolutionary compution
ENSO = El Niño-Southern Oscillation
EOF = empirical orthogonal function
EEOF = extended empirical orthogonal function
GCM = general circulation model (or global climate model)
GA = genetic algorithm
GP = Gaussian process model
IC = information criterion
LP = linear projection
MAE = mean absolute error
MJO = Madden-Julian Oscillation
MLP = multi-layer perceptron neural network
MLR = multiple linear regression
MOS = model output statistics
MSE = mean square error
MSSA = multichannel singular spectrum analysis
NAO = North Atlantic Oscillation
NN = neural network
NLCCA = nonlinear canonical correlation analysis
NLPC = nonlinear principal component
NLPCA = nonlinear principal component analysis
NLSSA = nonlinear singular spectrum analysis
PC = principal component
PCA = principal component analysis
PNA = Pacific-North American teleconnection
POP = principal oscillation pattern
QBO = Quasi-Biennial Oscillation
RBF = radial basis function
RMSE = root mean square error
SSA = singular spectrum analysis
SLP = sea level pressure
SOM = self-organizing map
SST = sea surface temperature (sum of squares in Chapter 1)
SVD = singular value decomposition
SVM = support vector machine
SVR = support vector regression

Contents