Harvey Mudd College
Computer Science 152

Neural Networks

Fall 2007

Trailer: Can a computer be taught to read words aloud, recognize faces, perform a medical diagnosis, drive a car, play a game, balance a pole, predict physical phenomena? The answer to all of these is yes. All these applications and others have been demonstrated using varieties of the computational model known as "neural networks", the subject of this course.

The course will develop the theory of a number of neural network models. Participants will exercise the theory through both pre-developed computer programs and ones of their own design.

Instructor: Professor Bob Keller, x 18483, 1253 Olin, keller at cs.hmc.edu

Catalog Description: Modeling, simulation, and analysis of artificial neural networks. Relationship to biological neural networks. Design and optimization of discrete and continuous neural networks. Backpropagation, and other gradient descent methods. Hopfield and Boltzmann networks. Unsupervised learning. Self-organizing feature maps. Applications chosen from function approximation, signal processing, control, computer graphics, pattern recognition, time-series analysis. Relationship to fuzzy logic, genetic algorithms, and artificial life.

Prerequisites: CS 60 (Principles of Computer Science) and Mathematics 12 (Multi-variate Calculus and Linear Algebra), or permission of the instructor.

Text: Raul Rojas, ***Neural Networks - A Systematic Introduction***, Springer-Verlag, Berlin, New-York, 1996 (502 p.,350 illustrations). This book is no longer in print. However, it may be obtained legally on the web in pdf form at the following site:

http://page.mi.fu-berlin.de/rojas/neural/index.html

Note that both chapter and whole-book files are available, and there is an errata sheet (very small), plus the table of contents is on the website.

Course Requirements and Grading:

There will be some homework, mostly programming assignments, but no exams. The assignments will constitute about 40% of your grade. 40% of your grade is from a substantial final project involving either creation of a working neural network application or a research paper. The grade on the project will be determined by the comprehensiveness and degree to which you explored competing approaches. The projects will be presented orally. 10% of your grade will be based on a preliminary presentation you make, ideally on material related to the your project. Finally, 10% will be based on general participation, which includes attendance.

CS 152 Topic Outline (Approximate) [Reading is in brackets; additional reading will be provided as needed. either on the web or as handouts]:

Week 1 (1 day):
o Biological motivation [chapter 1]
o Historical overview
o Perceptrons [chapter 3]

Week 2:
o Perceptron convergence theorem [chapter 4]
o Adalines
o Least Mean Square (LMS) algorithm
o Adaptive filtering

Week 3:
o Backpropagation [chapter 7]

Week 4: Backpropagation performance optimization [chapter 8]
o Conjugate gradient
o Newton's method
o Levenberg-Marquardt
o Quickprop
o Resilient backpropagation (RPROP)

Week 5: Sequential networks [section 7.4]
o Time series
o Backpropagation through time
o Real-time recurrent learning
o Control applications
o Finite Impulse Response (FIR) MLP
o Temporal difference Method

Week 6: Radial basis function networks [section 16.2.5]

Week 7: Support Vector Machines (SVM)

Week 8: Presentations by students

Week 9: Supervised and unsupervised Hebbian learning [chapter 12]

- o Supervised Hebbian learning
- o Pseudoinverse rule
- o Unsupervised Hebbian learning

Week 10: Associative learning (supplementary material)
- o Unsupervised Hebb rule
- o Hebb rule with decay
- o Instar rule
- o Kohonen rule
- o Outstar rule

Week 11: Competitive networks [chapters 5, 15]
- o Hamming network
- o Self-Organizing feature maps (SOM)
- o Counterpropagation networks (CPN)
- o Learning vector quanitization (LVQ)
- o Principal component networks (PCAN)
- o Independent component analysis (ICAN)
- o Bi-directional associative memory (BAM)
- o Adaptive resonance theory (ART)

Week 12: Physics-based networks [chapter 13]
- o Hopfield networks
- o Solving constraint-satisfaction problems
- o Spin-glass model and simulated annealing
- o Boltzman networks
- o Cauchy networks
- o Cascade correlation learning
- o Helmholtz machine

Week 13: Other soft computing concepts [chapters 11, 17]
- o Fuzzy logic
- o Evolutionary computation
- o Artificial life

Weeks 14-15: Final Presentations by Students