AI and its Applications

Eötvös Loránd University, Budapest, Hungary, 2005 – 2007 Fall, Spring
Instructor: Barnabás Póczos

This course provides a brief introduction to AI and its applications. We cover basic methods such as machine learning (supervised and unsupervised learning, reinforcement learning, graphical models, Markov chain Monte Carlo methods), knowledge representation, and feature selection. We discuss topics and applications in bioinformatics, computer vision, robotics, planning, and games. The course project assignments require programming in Matlab, C++, or Java.

Graphical Models

Bayes nets: directed and undirected graphical models, inference, loopy extension.

Hidden Markov models: brute-force likelihood evaluation, forward algorithm, backward algorithm, Viterbi algorithm, Baum-Welch algorithm, parameter estimation with EM.

Kalman-filtering and Extended Kalman filtering: Linear Dynamical System parameter estimation with EM and Extended Kalman-filtering.

Variational methods: free energy.

Support Vector Machines

Foundations: classification, margin, primal problem, Wolfe dual, KKT conditions, support vectors, hard SVM, soft SVM, nonlinear SVM, kernel map, Mercer conditions.

Computational Learning Theory

Expected risk. Empirical risk. VC dimension. Risk bound. Generalization error.
VC dimension of hyper planes. VC dimension of SVMs. Structural risk minimization.

Markov Chain Monte Carlo methods

Sampling from uniform distribution. Sampling from general distributions.
Importance sampling and rejection sampling. Curse of dimensionality.
Metropolis-Hastings algorithm. Gibbs sampling. Simulated annealing.
Sampling from graphs. Sampling from permutations.

Reinforcement Learning

Markov decision process. Value functions. Policy evaluation. Policy iteration. Value iteration.
Bellman equation. Dynamic programming. Monte Carlo methods. On/Off policy methods.
Temporal difference learning. Eligibility traces. Q-learning. SARSA. Function approximation.
Exploration vs exploitation. POMDP.

Robotics

Planning, Monte Carlo planning. Control. Path finding. A and A^* algorithms.
Simultaneous localization and mapping (SLAM).

Features

Feature selection, construction, extraction.
Optimization with the Cross Entropy method, genetic algorithms, and simulated annealing.

AI in games

Game theory. Minimax algorithms. Alpha-beta pruning.
Upper confidence bounds applied to trees (UCT).
AI algorithms for Tetris, Ms. Pac Man, Neverwinter Nights, Chess, Go, Backgammon...

AI in medical data processing

EEG, fMRI, microarray data processing. AI systems for diagnosis in cancer research.