

15-780 review session

Topics

- Probabilistic inference
- Bayesian networks
- Density estimation
- HMMs
- Decision trees
- Neural networks
- MDPs
- Reinforcement learning

Probabilistic inference

- Axioms
- Joint distribution, conditional probability
- Chain rule
- Bayes rule

Bayesian networks

- Meaning of nodes, edges (and lack of edges)
- CPTs
- Joint probability distributions
- From variables to Bayesian networks
- Inference complexity
- Inference methods
- Structure learning (briefly)

Density estimation

- Relation to Bayesian networks
- Maximum likelihood estimation (MLE)
- Binomial variables
- Discrete variables
- Normally distributed variables (continuous)

HMMs

- Why HMMs?
- Markov assumption
- Formal model and structure of HMMs
- Inference in HMMs
 - No observation
 - With observations (including Viterbi)
- Structure learning (Baum-Welch)
- Expectation maximization

Decision trees

- Classification
- Entropy
- Conditional entropy
- Information gain
- Overfitting, pruning using test data
- Continuous values

Neural networks

- Regression
- Linear regression
- Gradient descent
- Perceptron
- Sigmoid function for classification
- Multi-layer neural networks

MDPs

- Rewards, actions
- Discounted rewards
- Sum of expected rewards with no actions
- Determining the optimal policy:
 - Value iteration
 - Policy iteration

Reinforcement learning

- Model, assumptions
- Determining sum of expected rewards from a state
 - Supervised learning
 - Certainty-Equivalent (CE)
 - Temporal difference (TD, model free)
- Determining optimal policy
 - Update CE
 - Q learning (updating TD)
- How to explore the set of actions