

CS 540 Final Exam Topics

1. Constraint Satisfaction

Problem formulation in terms of variables, domains and constraints, constraint graph, depth-first search, backtracking with consistency checking, most constrained variable heuristic, most constraining variable heuristic, least constraining value heuristic, min-conflicts heuristic, min-conflicts algorithm, forward checking algorithm, arc consistency algorithm (AC-3), combining search with CSP inference.

2. Support Vector Machines

Maximum margin, definition of margin, kernel trick, support vectors, slack variables.

3. Neural Networks

Perceptron, LTU, activation functions, bias input, input units, output units, Perceptron learning rule, Perceptron learning algorithm, Perceptron convergence theorem, epoch, weight space, input space, linearly separable, credit assignment problem, multi-layer feed-forward networks, hidden units, sigmoid function, ReLU, softmax, back-propagation algorithm, gradient descent search in weight space, stochastic gradient descent, parameter setting using a tuning set, deep learning, convolutional neural networks, pooling.

4. Reasoning under Uncertainty

Random variable, mutually exclusive, prior probability, 3 axioms of probability, joint probability, conditional probability, posterior probability, full joint probability distribution, degrees of freedom, summing out, marginalization, normalization, product rule, chain rule, conditionalized version of chain rule, Bayes's rule, conditionalized version of Bayes's rule, addition/conditioning rule, independence, conditional independence, naïve Bayes classifier, add-1 smoothing, Laplace smoothing.

5. Bayesian Networks

Bayesian network DAG, conditional probability tables, space saving compared to full joint probability distribution table, conditional independence property defined by a Bayesian network, inference by enumeration from a Bayesian network, naïve Bayes classifier as a Bayesian network.

6. Speech Recognition

Phones, phonemes, speech recognition using Bayes's rule, language model, acoustic model, bigram model, trigram model, first-order Markov assumption, probabilistic finite state machine, first-order Markov model, state transition matrix, π vector, computing conditional probabilities from a Markov model, hidden Markov model, observation likelihood matrix, computing joint probabilities and conditional probabilities from an HMM by enumeration. (Nothing on Forward algorithm, Viterbi algorithm, Forward-Backward algorithm, Siri, particle filters, tracking in video.)

7. Computer Vision

Viola-Jones face detection algorithm, boosting ensemble learning, AdaBoost algorithm, weak classifier, weighted-majority classification, decision stump.