CMSC 671, Final Exam Content

Planning	Concepts
Classical Planning	
Graph-based planning	
Scheduling	NO
Hierarchical planning	
Multi-agent planning	NO
Semantic Web and OWL	**NO**
Uncertainty and Decision Making	
Bayes' Rule; Probabilistic Reasoning	Exercises, see HW5 (small)
Bayesian Networks	Exercises, see HW5 HW5 (small)
Probabilistic reasoning over time	Concepts
Decision theory	Concepts
Machine Learning	
Machine Learning: Decision Trees	Exercise: build a decision tree; class exercise (complete)
Naive Bayes	Exercise: apply Naïve Bayes; class exercise (complete)
Learning Probabilistic Models	Concepts
Reinforcement learning	Exercise: calculate Q's; class exercise (small)
Multiagent Systems (Introduction)	Concepts

Concepts means you should know the concepts, descriptions, and general idea of any algorithm discussed. Semantic Web and OWL will not be on the exam.

CMSC 671, Final Exam
Detailed Content

- **Planning**
- Classical Planning
  - The planning problem, what is planning?, representation, PDDL (is a first-order logic language?), preconditions, applicability of actions, effects, state-space planning (progression, regression, and heuristics).
- Graph-based planning
  - Planning graph components, valid plans, basic GraphPlan algorithm, searching for a plan.
- Hierarchical planning
  - Hierarchical decomposition or HTN, assumptions, refinements, HLA, finding a solution plan, searching for a plan.

- **Uncertainty and Decision Making**
- Bayes' Rule; Probabilistic Reasoning
  - Decision making with uncertainty, Bayesian reasoning, joint probabilities, full joint probability, conditional probabilities, inference from the joint, independence, conditional independence, Bayes' rule, Bayesian inference, limitations.
- Bayesian Networks
  - Bayesian networks, conditional independence and chaining
  - Inference tasks
  - Inference methods
    - exact inference: enumeration and variable elimination
    - approximate inference:
      - sampling methods: direct sampling, rejection sampling, and MCMC)
- Probabilistic reasoning over time
  - States and observations, transition model and sensor model, Markov processes, assumptions, inference tasks, Probabilistic temporal models
- Decision theory
  - Utilities, rational agents, maximum expected utility, bounded rationality, expected monetary value (EVM), money vs utility, decision networks, evaluating decision networks, value of information
- **Machine Learning**
- Decision Trees
- Naive Bayes
- Learning Probabilistic Models
  - Learning Bayesian Networks: learning parameters, learning structure, learning with missing values.
  - Clustering
- Reinforcement learning

- **Multiagent Systems (Introduction)**
  - Cooperative MAS, competitive MAS, application domains, examples, approaches used, learning in MAS.