

# Final Exam Review / Study Guide

CS 2710, Fall 2010

# Scope and Format

- Primarily Chapters 10 and on (sections from the syllabus), but still need to know basics from first part of course
  - Lectures (notes/in class)
  - Textbook
  - Homework assignments
- Closed book, in-class exam
  - no make-up possibilities
- Question types same as midterm
  - True/false
  - short-answer
  - representation/problem solving

# KR/Planning

- Representation and Inference
  - STRIPS
  - POP
  - GraphPlan
  - SatPlan
- Planning and KR Translation/Comparison
  - English -> Representation
  - Representation A -> Representation B
  - Relation to Search

# Uncertainty

- Motivation/Differences from Logic
- Representation
  - random variables, atomic events
  - prior and conditional probability
  - definitions and axioms
  - distributions
- Inference via full joint distribution
- Bayes Rule
- (Conditional) Independence

# Bayesian Networks

- Translation from English into a network
- Compute the probability of some outcome using the network
- Explanation and comparison of networks
  - Are two events in the network (conditionally) independent ?
  - If multiple networks can model a scenario, why is one better than the other?
  - Comparison with FOL solution

# Decision making

- A person wants to buy a car and she must decide between two used cars, car A and B. Car A costs \$1,500, but its value is estimated to be at \$2,000 if it is of good quality, so the buyer can make \$500 by buying the car. If the quality is bad, the costs of repairs are \$700 in which case the buyer would lose \$200. Car B costs \$1,150 which is \$250 below its market value. Even if it is in bad condition, the repairs will cost at most \$150.
- The buyer knows that the chances that car A is of good quality are 0.7 and the chances that car B is of good quality are 0.8.
- **1. [4 points]** Assume that the buyer's objective is to maximize the expected monetary value of her purchase. What is the best car to buy? Justify your answer.
- **2. [1 point]** Why does money not always behave as a utility function? Explain in one or two sentences using the context of the car buying example above.

# Machine Learning

- Decision tree learning
- Information gain computation
- Learning curves
- Training/development/testing/cross-validation
- Evaluation metrics

# Communication

- (P)CFGs – writing and using



# Summary

- You should be able to formalize/represent a problem intuitively described in English
- You should be able to solve such a problem, once represented
- You should know the correct terminology
  
- You should be able to translate one formal representation into another
- You should be able to compare, contrast, and evaluate all the different representation and reasoning methods (e.g., with respect to expressiveness/tractability tradeoffs)